Math. Struct. in Comp. Science (2018), *vol.* 28, *pp.* 1485–1505. © Cambridge University Press 2018
doi:10.1017/S0960129518000221 First published online 25 July 2018

Finding small counterexamples for abstract rewriting properties

HANS ZANTEMA

Department of Computer Science, TU Eindhoven, P.O. Box 513,
5600 MB Eindhoven, The Netherlands.
Email: H.Zantema@tue.nl, and
Institute for Computing and Information Sciences, Radboud University
Nijmegen, P.O. Box 9010, 6500 GL Nijmegen, The Netherlands

Received 1 February 2012; revised 1 November 2017

Rewriting notions like termination, normal forms and confluence can be described in an abstract way referring to rewriting only as a binary relation. Several theorems on rewriting, like Newman's lemma, can be proved in this abstract setting. For investigating possible generalizations of such theorems, it is fruitful to have counterexamples showing that particular generalizations do not hold. In this paper, we develop a technique to find such counterexamples fully automatically, and we describe our tool Carpa that follows this technique. The basic idea is to fix the number of objects of the abstract rewrite system, and to express the conditions and the negation of the conclusion in a satisfiability (SAT) formula, and then call a current SAT solver. In case the formula turns out to be satisfiable, the resulting satisfying assignment yields a counterexample to the encoded property. We give several examples of finite abstract rewrite systems having remarkable properties that are found in this way fully automatically.

1. Introduction

Rewriting occurs in several flavors: first-order term rewriting, graph rewriting, string rewriting, higher order term rewriting, conditional rewriting, rewriting with respect to strategies or priorities, and so on. But in all kinds of rewriting one considers a set A of objects that may be rewritten, and one considers a *rewrite relation* that is a binary relation on A that describes rewrite steps: *abstract rewriting*. Most times there is a notion of *computation*: apply rewrite steps as long as possible. If after a finite number of steps an object is obtained on which no rewrite step is possible any more, such an object is called a *normal form*, the result of the computation. One can state that computation coincides with rewriting to normal form. In this framework, several things can go wrong. For instance, computation may go on forever, without reaching a normal form. This can be avoided by requiring that the rewrite system is *terminating*, that is, it does not allow infinite computations. In doing rewriting with respect to a terminating rewrite system, one will always reach a normal form. But without extra restrictions, this normal form may be not unique, while in most applications, it is desirable that the result of a computation is unique, that is, does not depend on the choice of a rewrite step for objects that allow different rewrite steps. These non-unique normal forms can be avoided by requiring that

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

the rewrite system is *confluent*, that is, if an object x can be rewritten to both y and z, then there exists an object w such that both y and z can be rewritten to w.

A huge amount of research has been done on the basic notions termination and confluence and several variants. In particular, several results exist for concluding termination (Bachmair and Dershowitz 1986; Bellegarde and Lescanne 1990; Doornbos and von Karger 1998; Van Oostrom and Zantema 2012) and confluence (Van Oostrom 1994, 2008) in abstract rewriting, that is, for binary relations on arbitrary sets, independent on the structure of the objects. When searching for theorems for rewriting properties it would be very convenient to have machinery for automatically finding counterexamples for variants of the theorems, for instance, to check whether conditions are essential. In case this machinery finds such a counterexample after removing a condition, one may conclude that this condition is essential, and inspecting the counterexample may help to understand why this is the case.

Developing such a machinery is exactly the topic of this paper. We focus on the situation where the set A is finite. In applications of rewriting, most time the sets of objects are infinite. But in case some property does not hold, often this can be shown by a finite counterexample, and often the smallest possible counterexample gives the best insight why the property does not hold. So we will not only focus on counterexamples in which A is finite, we will also focus on counterexamples in which $\#A$ is as small as possible. As our abstract rewrite relations are just binary relations on a set A, in this paper, we will mainly speak about binary relations rather than rewrite systems.

As a binary relation on a set of n elements can be expressed by n^2 Boolean variables, our problem area can be seen as a class of constraint problems on Boolean variables. Our goal is to express rewriting properties like termination and confluence in propositional logic in such a way that we may express our problems as propositional SAT(isfiability) problems, by which we may exploit current powerful SAT solvers.

We succeeded in expressing all abstract rewriting properties of our interest. Often this can be done in several ways. For instance, a relation R on a finite set is terminating if and only if an irreflexive transitive relation S exists such that $R \subseteq S$, as we will see in Theorem 4.1. This characterization is the basis of how termination is characterized in our machinery. Other correct characterizations of termination of R are obtained by replacing the requirement of S being transitive by $R \cdot S \subseteq S$ or by $S \cdot R \subseteq S$. Alternatively, one also checks that a relation R on a finite set is terminating if and only if R^+ is irreflexive. Here, R^+ is the transitive closure of R, for which we develop a characterization in propositional logic. Having many such correct characterizations for a single notion is typical in our investigations. In this paper, the choice among these characterizations was guided by efficiency of the implementation. For the chosen characterizations we prove correctness. None of the proofs are very deep; the main effort was in finding the right characterizations.

Based on these characterizations, we developed our tool `Carpa` (Counter examples for Abstract Rewriting Produced Automatically). In this tool, a list of properties for a number of binary relations can be entered, and then for a given number $n = \#A$ it either gives an example in which these properties hold, or concludes that such an example does not exist. Typically, when wondering whether a property P follows from conditions $C_1, \ldots, C_k$,

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

then one enters properties $C_1, \ldots, C_k$ and $\neg P$ in `Carpa`. In case an example is found, this is a counterexample showing that P does not follow from the conditions. Internally, the tool first builds a formula expressing the desired properties based on the characterizations presented in this paper, then it calls an external SAT solver. In case the SAT solver concludes that the formula is unsatisfiable, the tool concludes that there is no solution, and in case the SAT solver concludes that the formula is satisfiable, the tool investigates the corresponding satisfying assignment and extracts an example out of it that satisfies the given properties. Although internally the SAT solving is crucial, the user of `Carpa` does not see this. He only sees either the automatic creation of an example of a set of binary relations that satisfies the given list of desired properties, or the observation that it does not exist.

This paper is organized as follows. In Section 2, some preliminaries are given. In Section 3, we describe the basic encoding of relations in propositional formulas. In Section 4, we describe how to characterize termination. In Section 5, we describe how to specify transitive closures. Based on this, in Section 6, we describe how to characterize confluence. Completeness is defined to be the conjunction of termination and confluence. In Section 7, we see how this can be characterized much more efficiently than by taking the conjunction of the characterizations of termination and confluence. In all of these sections, we give several examples, for many of which it would be a hard job to find them by hand. These examples were all found by our tool `Carpa` as it is presented in Section 8. Most times these examples are the smallest possible, again shown by our implementation by yielding an unsatisfiable formula when decreasing the value of $n = \#A$. We conclude in Section 9.

2. Preliminaries

We start by recalling some basic notions that we will use throughout the paper.

A *binary relation* R on a set A is defined to be a subset of $A \times A$. For $x, y \in A$, we will use xRy as an abbreviation of $(x, y) \in R$.

A binary relation R on a set A is called *reflexive* if $\forall x \in A : xRx$ holds.

A binary relation R on a set A is called *irreflexive* if $\forall x \in A : \neg(xRx)$ holds.

A binary relation R on a set A is called *symmetric* if $\forall x, y \in A : xRy \rightarrow yRx$ holds.

A binary relation R on a set A is called *transitive* if $\forall x, y, z \in A : (xRy \wedge yRz) \rightarrow yRz$ holds.

For two binary relations R, S on a set A their *composition* $R \cdot S$ is defined to be the relation

$$R \cdot S \;=\; \{(x, z) \in A \times A \mid \exists y \in A : (xRy \wedge yRz)\}.$$

The *identity relation* I on A is defined by $I = \{(x, x) \mid x \in A\}$.

For $i \geqslant 0$, the relation R^i is defined inductively by

$$R^0 = I, \quad R^{i+1} = R \cdot R^i.$$

It is easily proved by induction that $R^i \cdot R^j = R^{i+j}$ for all $i, j \geqslant 0$.

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

For a binary relation, R on a set A its *transitive closure* R^+ is defined by

$$R^+ = \bigcup_{i=1}^{\infty} R^i.$$

It is the smallest transitive relation that contains R, as stated in the following well-known lemma.

Lemma 2.1. If R and S are binary relations on a set A for which $R \subseteq S$ and S is transitive, then $R^+ \subseteq S$.

For a binary relation, R on a set A its *transitive reflexive closure* R^* is defined by

$$R^* = \bigcup_{i=0}^{\infty} R^i = I \cup R^+.$$

An element $x \in A$ is called a *normal form* with respect to R if $\neg(xRy)$ for all $y \in A$. For an element, $x \in A$ an element $y \in A$ is called a *normal form of* x with respect to R if xR^*y and y is a normal form with respect to R.

A binary relation R on a set A is called *terminating* or *well-founded* if no infinite sequence $a_1, a_2, \ldots$ of elements in A exists such that $a_i R a_{i+1}$ for all $i \geqslant 1$. For A being finite this is equivalent to irreflexivity of R^+. With respect to a terminating binary relation, every element has at least one normal form.

The union of two terminating relations does not need to be terminating. For instance, for $R = \{(1,2)\}$ and $S = \{(2,1)\}$ both R and S are terminating but $R \cup S$ is not. The following theorem due to Doornbos and von Karger (1998) gives an extra condition by which termination of the union can be concluded. In Example 5.1, we will show that this does not hold for a variant of this condition. Proofs of this theorem are given in Doornbos and von Karger (1998) and Van Oostrom and Zantema (2012).

Theorem 2.1. Let R and S be terminating relations on a set A for which

$$R \cdot S \subseteq R \cup (S \cdot (R \cup S)^*).$$

Then, $R \cup S$ is terminating.

For a binary relation, R its *inverse* R^{-1} is defined by

$$(x, y) \in R^{-1} \iff (y, x) \in R.$$

A binary relation R is called *confluent* if $(R^*)^{-1} \cdot R^* \subseteq R^* \cdot (R^*)^{-1}$. It is easy to see that with respect to a confluent binary relation, every element has at most one normal form.

A binary relation R is called *locally confluent* if $R^{-1} \cdot R \subseteq R^* \cdot (R^*)^{-1}$. It is well-known that there are non-terminating relations R that are locally confluent but not confluent; in Example 6.1, we will see how to find such an example fully automatically. The following well-known theorem is usually called Newman's Lemma. For a proof, we refer to standard texts like Baader and Nipkow (1998) and TERESE (2003).

Theorem 2.2. A terminating relation is confluent if and only if it is locally confluent.

A binary relation that is both terminating and confluent is called *complete*.

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

3. Basic encoding

We fix a number n. We will consider binary relations on the set $\{1, 2, \ldots, n\}$. These binary relations are numbered from 1 to m, and are denoted by $R_1, R_2, \ldots, R_m$. Typically, the first one, two or three are the relations referred to in the given property, the others are introduced for being able to express these properties. As we want to express the given property in a SAT problem, these m binary relations have to be expressed by Boolean variables. We do this by introducing n^2m Boolean variables $R(k, i, j)$ for $k = 1, \ldots, m$ and $i, j = 1, \ldots, n$. If these Boolean variables have Boolean values, they define the corresponding relation $R_1, R_2, \ldots, R_m$ as follows:

$$(i, j) \in R_k \iff R(k, i, j) \text{ is true.}$$

Using this encoding, the following standard properties of binary relations are expressed as propositional formulas in the variables $R(k, i, j)$.

A relation R_k is reflexive if and only if $\mathsf{refl}(k)$ defined by

$$\mathsf{refl}(k) \equiv \bigwedge_{i=1}^{n} R(k, i, i),$$

holds.

A relation R_k is irreflexive if and only if $\mathsf{irrefl}(k)$ defined by

$$\mathsf{irrefl}(k) \equiv \bigwedge_{i=1}^{n} \neg R(k, i, i),$$

holds.

A relation R_k is transitive if and only if $\mathsf{trans}(k)$ defined by

$$\mathsf{trans}(k) \equiv \bigwedge_{i=1}^{n} \bigwedge_{j=1}^{n} \bigwedge_{p=1}^{n} ((R(k, i, j) \wedge R(k, j, p)) \rightarrow R(k, i, p)),$$

holds.

Next, we express some basic set theoretic concepts. A relation R_k is a subset of a relation $R_{k'}$ if and only if $\mathsf{subset}(k, k')$ defined by

$$\mathsf{subset}(k, k') \equiv \bigwedge_{i=1}^{n} \bigwedge_{j=1}^{n} (R(k, i, j) \rightarrow R(k', i, j)),$$

holds.

A relation R_k is the union of a relation $R_{k'}$ and a relation $R_{k''}$ if and only if $\mathsf{union}(k, k', k'')$ defined by

$$\mathsf{union}(k, k', k'') \equiv \bigwedge_{i=1}^{n} \bigwedge_{j=1}^{n} (R(k, i, j) \leftrightarrow (R(k', i, j) \vee R(k'', i, j)),$$

holds.

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

A relation R_k is the intersection of a relation $R_{k'}$ and a relation $R_{k''}$ if and only if $\mathsf{intersect}(k,k',k'')$ defined by

$$\mathsf{intersect}(k,k',k'') \equiv \bigwedge_{i=1}^{n}\bigwedge_{j=1}^{n}(R(k,i,j) \leftrightarrow (R(k',i,j) \wedge R(k'',i,j)),$$

holds.

Next, we define relation composition. A relation R_k is the composition of a relation $R_{k'}$ and a relation $R_{k''}$ if and only if $\mathsf{compose}(k,k',k'')$ defined by

$$\mathsf{compose}(k,k',k'') \equiv \bigwedge_{i=1}^{n}\bigwedge_{j=1}^{n}(R(k,i,j) \leftrightarrow \bigvee_{p=1}^{n}(R(k',i,p) \wedge R(k'',p,j))),$$

holds. Note that the size of all of these formulas is $O(n^3)$.

Now, we are ready to give our first example.

Example 3.1. For any sub-relation R of a strict order, can we conclude that R^2 is transitive? We will build a formula expressing this question. Let $R_1 = R$, let R_2 be the strict order, that is, it is transitive and irreflexive, and let $R_3 = R^2$. Our formula will consist of the conjunction of all conditions and the negation of the conclusion:

$$\mathsf{subset}(1,2) \wedge \mathsf{trans}(2) \wedge \mathsf{irrefl}(2) \wedge \mathsf{compose}(3,1,1) \wedge \neg(\mathsf{trans}(3)).$$

Applying a SAT solver to this formula for various n shows that for $n \leqslant 4$ this formula is unsatisfiable, but for $n = 5$ it is satisfiable. After renumbering the elements $1,2,3,4,5$, the resulting satisfying assignment can be interpreted as $1 < 2 < 3 < 4 < 5$ for $<$ being the order R_2, and $R = R_1$ is a relation for which $1R2$, $2R3$, $3R4$ and $4R5$ all hold, showing $1R^2 3$ and $3R^2 5$, but for which $1R^2 5$ does not hold, indeed showing that R^2 is not transitive.

By our tool `Carpa` as described in Section 8, we can enter the above requirements directly, from which the above formula is generated.

4. Termination

In this section, we present a way how to express the property of termination (= well-foundedness) of a binary relation R on a finite set as a propositional formula. This can be done in several ways, for instance, by expressing that R^+ is irreflexive. However, this requires R^+ to be expressed, and this requires several auxiliary relations as we will see in the next section. We prefer the characterization specified in the next theorem, by which only one auxiliary relation is required.

Theorem 4.1. A binary relation R on a finite set is terminating if and only if a binary relation S on the same set exists that is transitive and irreflexive, and for which $R \subseteq S$.

Proof. (only if) Let R be terminating, and choose $S = R^+$. Then, S is transitive and satisfies $R \subseteq S$. So, it remains to prove that $S = R^+$ is irreflexive. Assume it is not, then there is an element a satisfying aR^+a. This yields an infinite reduction $aR^+aR^+aR^+\ldots$, contradicting termination of R.

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

(if) Assume S satisfies $R \subseteq S$ and is transitive and irreflexive. Assume R is not terminating. Then, R admits an infinite reduction, and since the set is finite there is an element a occurring more than once in this infinite reduction. Hence, aR^+a. By Lemma 2.1, we conclude aSa, contradicting that S is irreflexive. □

Note that in Theorem 4.1 finiteness is only used in the 'if'-direction. There it is essential: the relation $R = <$ on the natural numbers is both irreflexive and transitive, so choosing $S = R$ satisfies all requirements, while R is not terminating.

Theorem 4.1 can be used for expressing termination of R_k in a SAT formula: simply add a fresh relation $R_{k'}$, so choose k' to be one higher than the highest relation number in use, and generate

$$\mathsf{trans}(k') \wedge \mathsf{irrefl}(k') \wedge \mathsf{subset}(k, k').$$

For example, the part $\mathsf{subset}(1,2) \wedge \mathsf{trans}(2) \wedge \mathsf{irrefl}(2)$ of the formula in Example 3.1 exactly expresses that R_1 is terminating.

Theorem 4.1 is intended to be used for specifying termination. For instance, if combined with other properties the resulting formula is unsatisfiable then we conclude that no terminating relation exists on n elements moreover satisfying the specified properties. We stress that Theorem 4.1 cannot be used for the other way around: specifying non-termination. For A being finite it is the case that R is non-terminating if and only if R^+ is not irreflexive, but from $R \subseteq S$ and transitivity of S we cannot conclude that $S = R^+$, only $R^+ \subseteq S$. In Section 5, we will see how to fully specify R^+, by which non-termination can be specified.

5. Transitive closure

For a given binary relation R on a finite set A, we want to fully specify its transitive closure R^+. The relation R^+ can be seen as the least fixed point of the equation

$$X = R \cup (R \cdot X).$$

However, by only giving this equation R^+ is not fully specified: in general this equation has more solutions than only the least fixed point R^+. Stated in other words: there are more fixed points than only the least fixed point. For instance, for R being the identity relation I on a set A, not only $R^+ = I$ satisfies the above equation, but also $A \times A$. The extra requirement 'least' in 'least fixed point' is hard to express in SAT.

Instead, we will specify R^+ by non-recursive equations.

Theorem 5.1. Let R be a relation on a finite set A. Let $k \geqslant 1$ satisfy $2^k \geqslant \#A$. Let R_i be relations on A for $i = 1, 2, \ldots, k$, satisfying

$$R_1 = R \cup R^2, \quad \text{and} \quad R_{i+1} = R_i \cup R_i^2 \quad \text{for } i = 1, \ldots, k-1.$$

Then, $R_k = R^+$.

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

Proof. First, we prove the following claim.

Claim: For $i = 1, 2, \ldots, k$, we have

$$R_i = \bigcup_{j=1}^{2^i} R^j.$$

This claim is proved by induction on i. For $i = 1$, it holds by definition. For the induction step, we have to prove that

$$R_i \cup R_i^2 = \bigcup_{j=1}^{2^{i+1}} R^j,$$

using the induction hypothesis $R_i = \bigcup_{j=1}^{2^i} R^j$.

'$\subseteq$': If $(x, y) \in R_i$, then by the induction hypothesis, we have $(x, y) \in \bigcup_{j=1}^{2^i} R^j \subseteq \bigcup_{j=1}^{2^{i+1}} R^j$. If $(x, y) \in R_i^2$ then there exists a z such that $(x, z) \in R_i$ and $(z, y) \in R_i$. Using the induction hypothesis twice yields $(x, z) \in R^j$ and $(z, y) \in R^{j'}$ for $j, j' \leqslant 2^i$. Hence, $j + j' \leqslant 2^{i+1}$, so $(x, y) \in R^{j+j'} \in \bigcup_{j=1}^{2^{i+1}} R^j$.

'$\supseteq$': Let $(x, y) \in \bigcup_{j=1}^{2^{i+1}} R^j$, then $(x, y) \in R^j$ for some $j \leqslant 2^{i+1}$. If $j \leqslant 2^i$ then $(x, y) \in \bigcup_{j=1}^{2^i} R^j = R_i$. If $j > 2^i$ then one can write $j = j' + j''$ for $1 \leqslant j', j'' \leqslant 2^i$, so $R^{j'} \subseteq R_i$ and $R^{j''} \subseteq R_i$. Hence,

$$(x, y) \in R^j = R^{j'+j''} = R^{j'} \cdot R^{j''} \subseteq R_i \cdot R_i = R_i^2,$$

concluding the proof of the claim.

It remains to prove that for $R_k = \bigcup_{j=1}^{2^k} R^j$ and $R^+ = \bigcup_{j=1}^{\infty} R^j$, we have $R_k = R^+$. Here, '$\subseteq$' is obvious; for the converse we have to prove that if $(x, y) \in R^j$ for any $j > 0$, then $(x, y) \in R^{j'}$ for some $j' \leqslant 2^k$. Let $(x, y) \in R^j$ for some $j > 2^k$. Then there are $x_0, x_1, \ldots, x_j$ such that $x_0 = x$, $x_j = y$ and $x_i R x_{i+1}$ for all $i = 0, \ldots, j-1$. Since $j > 2^k \geqslant \#A$, among the j elements $x_0, \ldots, x_{j-1}$ from A at least one element occurs at least twice, by which $(x, y) \in R^{j'}$ can be concluded for some $j' < j$. Repeating this argument will yield $j' \leqslant 2^k$ for which $(x, y) \in R^{j'}$, concluding the proof. □

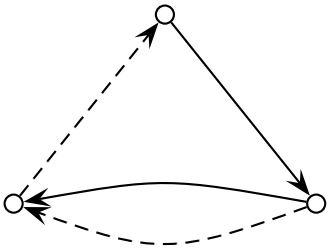
The bound $2^k \geqslant \#A$ in Theorem 5.1 is sharp, as is shown by the following example. Let $\#A = 2^k + 1$, say, $A = \{1, 2, \ldots, 2^k + 1\}$. Let R describe a cycle of length $2^k + 1$, say $R = \{(1, 2), (2, 3), \ldots, (2^k, 2^k + 1), (2^k + 1, 1)\}$. Then $(1, 1) \in R^{2^k+1} \subseteq R^+$, but $(1, 1) \notin R_k$ since 1 cannot be reached from 1 in less than $2^k + 1$ steps.

Note that for expressing termination of a relation R by means of Theorem 4.1 only one auxiliary relation is required, while for expressing the transitive closure of a relation R by means of Theorem 5.1 the number of required auxiliary relations is logarithmic in $\#A$.

Example 5.1. A direct consequence of Theorem 2.1 is that the union of two terminating relations R and S is terminating if $R \cdot S \subseteq R \cup S^+$. One may wonder whether this still holds if this requirement is relaxed to $R \cdot S \subseteq R^+ \cup S^+$. It turns out to be not. In our tool `Carpa/` one can enter these requirements, after which the termination requirements are encoded according Theorem 4.1, transitive closures are encoded according Theorem 5.1, and composition, union and subset are encoded as described in Section 3. The corresponding formula turns out to be satisfiable, and by inspecting the resulting satisfying assignment

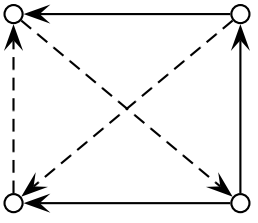
https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

`Carpa` generates the following counterexample for $\#A = 3$, in which R steps are denoted by solid arrows and S steps are denoted by dashed arrows:



Indeed, one easily checks that for both ways a solid arrow is followed by a dashed arrow, a path with the same start and end can be found either consisting of only solid arrows or only dashed arrows, showing $R \cdot S \subseteq R^+ \cup S^+$, while only the solid arrows are terminating, the same for dashed arrows, but the combination admits a cycle.

In this example, we see that R and S are not disjoint: at the bottom there is both an R step and an S step from right to left. If we moreover require disjointness of R and S then for $\#A = 3$ there is no solution any more, but for $\#A = 4$ the same approach yields the following example:



The following theorem shows that the bound $2^k \geqslant \#A$ in Theorem 5.1 may be omitted in case R satisfies some extra condition, namely, that $R_k = \bigcup_{j=1}^{2^k} R^j$ is transitive. For taking k large enough, namely, $k \geqslant \log_2(\#A)$, this always holds due to Theorem 5.1, but often this already holds for smaller values of k.

Theorem 5.2. Let R be a relation on a finite set A and let $k \geqslant 1$. Let R_i be relations on A for $i = 1, 2, \ldots, k$, satisfying

$$R_1 = R \cup R^2, \quad \text{and} \quad R_{i+1} = R_i \cup R_i^2 \quad \text{for } i = 1, \ldots, k-1.$$

Assume R_k is transitive. Then, $R_k = R^+$.

Proof. From the proof of Theorem 5.1, we use the claim, and conclude $R_k = \bigcup_{j=1}^{2^k} R^j$. Hence, $R_k \subseteq \bigcup_{j=1}^{\infty} R^j = R^+$. Conversely, we have $R \subseteq R_k$, and by Lemma 2.1 and transitivity of R_k we conclude $R^+ \subseteq R_k$. □

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

6. Confluence

For specifying confluence and local confluence of a binary relation R, we need $R^* = I \cup R^+$. In Section 5, we described how for a given relation R the relation R^+ and hence also R^* can be described using a logarithmic number of auxiliary relations by means of Theorem 5.1. In the remainder of this section, we assume that apart from R we also have access to the relation R^*.

Apart from composition it is convenient to specify peak and valley. For two binary relations R and S on a set A, we write $\mathsf{peak}(R,S)$ for $R^{-1} \cdot S$, so

$$(x,y) \in \mathsf{peak}(R,S) \iff \bigvee_{z \in A} (zRx \wedge zSy).$$

Similarly, we write $\mathsf{val}(R,S)$ for $R \cdot S^{-1}$, so

$$(x,y) \in \mathsf{val}(R,S) \iff \bigvee_{z \in A} (xRz \wedge ySz).$$

Now by definition a relation R is confluent if and only if

$$\mathsf{peak}(R^*,R^*) \subseteq \mathsf{val}(R^*,R^*),$$

and a relation R is locally confluent if and only if

$$\mathsf{peak}(R,R) \subseteq \mathsf{val}(R^*,R^*).$$

Example 6.1. When looking for a binary relation on four elements that is locally confluent but not confluent, our tool `Carpa` finds

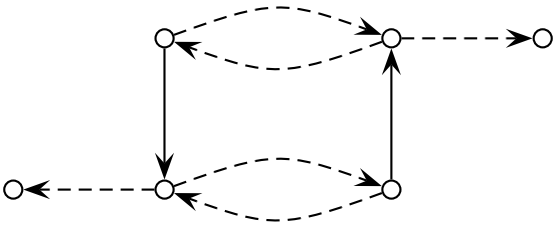


This is the well-known standard example of a locally confluent system that is not confluent, but now it has been found fully automatically. For details how this was done by our tool `Carpa` we refer to Section 8.

Example 6.2. By taking $R = \{(1,2)\}$ and $S = \{(1,3)\}$ one easily sees that the union of two confluent relations R and S does not need to be confluent. But in case moreover R,S-peaks converge, that is, $R^{-1} \cdot S \subseteq (R \cup S)^* \cdot ((R \cup S)^*)^{-1}$, one may conclude local confluence, and one may wonder whether then confluence of the union may be concluded. This is not the case: take $R = \{(1,2),(2,3)\}$ and $S = \{(2,1),(1,4)\}$: this is essentially the same example as in Example 6.1 in which the arrows from left to right are R steps and the arrows from right to left are S steps. One easily checks that in this example one has $R^{-1} \cdot S \subseteq (R^2)^{-1} \cup S^2$. Next, we wonder whether confluence of the union can be concluded if this requirement of peak convergence is strengthened to $R^{-1} \cdot S \subseteq S \cdot R^* \cdot (R^{-1})^*$. Roughly speaking, this means that the S steps always shift to the left and never disappear

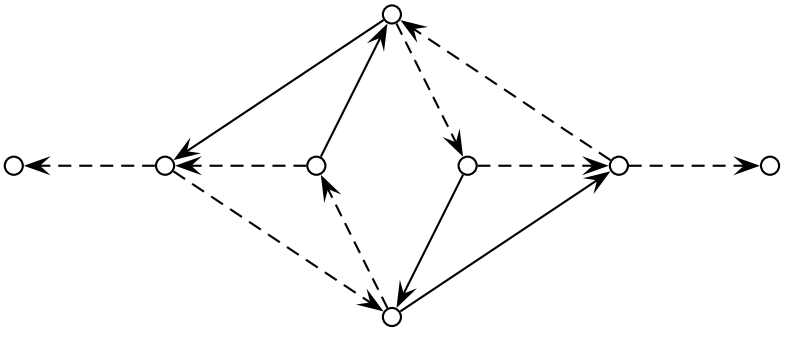
https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

or duplicate. Here, `Carpa` yields that this is not possible for $n = \#A \leqslant 5$, but for $n = 6$ the following example is generated:



Here, again R steps are denoted by solid arrows and S steps are denoted by dashed arrows. Indeed one easily checks that both the solid arrows and the dashed arrows are confluent, while the union is not since there are two distinct normal forms that are connected. Also, the R,S-peak requirement holds, even $R^{-1} \cdot S \subseteq S \cdot R$. Note that since local confluence easily follows from our requirements, due to Newman's Lemma (Theorem 2.2) in every such example $R \cup S$ will be non-terminating.

Example 6.3. In Stump et al. (2011), confluence was studied of a system combining simply typed lambda calculus and type computation, using the technique of decreasing diagrams (Van Oostrom 1994, 2008). This work was extended to Stump et al. (2013), where the abstract properties leading to confluence were further investigated. It turned out that for R being β-reduction in simply typed lambda calculus and S consisting of type computation steps, these relations are both confluent and satisfy $R^{-1} \cdot S \subseteq (S \cup R^*) \cdot (R^*)^{-1}$, while the goal was to prove that $R \cup S$ is confluent. So, in this setting it was a natural question whether this could already be concluded from these abstract properties. It turned out to be not, as is shown by the following example in which again R and S are denoted solid and dashed, respectively:



Originally, this example was found by expressing variants of the requirements in a propositional formula, applying a SAT solver on it and inspecting the resulting satisfying assignment. In fact this was the starting point of the research described in this paper and

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

the development of Carpa. Applying Carpa on a slightly stronger requirement, namely $R^{-1} \cdot S \subseteq (S \cup R) \cdot (R^*)^{-1}$, also yields this example.

7. Completeness

A binary relation is called *complete* if it is both terminating and confluent. Due to Theorem 2.2 (Newman's Lemma) this is equivalent to being both terminating and locally confluent. Since termination implies that every element has at least one normal form and confluence implies that every element has at most one normal form, completeness implies that every element has exactly one normal form, which is a desirable property in many situations. One way to specify completeness in propositional logic is simply by both specifying termination by the approach of Section 4 and specifying confluence or local confluence by the approach of Section 6. However, this needs one auxiliary relation for termination and a logarithmic number of auxiliary relations for (local) confluence. Alternatively, one could specify the transitive closure of the relation as described in Section 5, and require this to be irreflexive, and specify confluence based on this transitive closure. But also this approach needs a logarithmic number of auxiliary relations. The next theorem describes a way to specify completeness using only two auxiliary relations, which may be more efficient in the sense that SAT solvers find solutions much faster. In the discussion of Example 7.3, we will see an instance where this really pays off. The approach is based on discussions with Bas Joosten.

Theorem 7.1. A binary relation R on a set A is complete if and only if two binary relations S and T on A exist such that the following properties hold:

1. $R \subseteq S$,
2. S is transitive and irreflexive,
3. $\bigwedge_{x \in A}(xTx \vee \bigvee_{y \in A} xRy)$,
4. $\bigwedge_{x,y \in A}((xSy \wedge yTy) \rightarrow xTy)$,
5. $\bigwedge_{x,y,z \in A, y \neq z} \neg(xTy \wedge xTz)$.

Proof. (only if) Let R be complete, then every element $x \in A$ has a unique normal form with respect to R. Define $S = R^+$ and define T by xTy if and only if y is the normal form with respect to R. Now property 1 is obvious and property 2 follows from termination of R. For property 3, we have to prove that xTx holds for every normal form x; this holds since x is its own normal form. For proving property 4 assume that xSy and yTy. Then, xR^+y according to the definition of S, and y is a normal form with respect to R since yTy. Hence, y is a normal form of x, hence xTy, concluding the proof. Finally, for proving property 5 assume that xTy and xTz for $x \neq y$. Then due to the definition of T the element x has two distinct normal forms y and z, contradicting completeness.

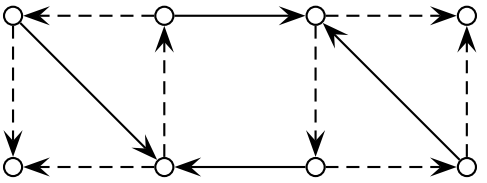
(if) Assume S and T satisfy properties 1–5. By properties 1 and 2 and Theorem 4.1, we conclude that R is terminating. It remains to prove that R is confluent. Assume xR^*y and xR^*z, we have to find w such that yR^*w and zR^*w. If $x = y$, we may choose $w = z$ and if $x = z$ we may choose $w = y$, so it remains to consider xR^+y and xR^+z. Since R is terminating, y has a normal form y' and z has a normal form z'. Since xR^+y and xR^+z, we conclude xR^+y' and xR^+z'. Due to Lemma 2.1, we obtain xSy' and xSz'. Since y'

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

and z' are normal forms, from property 3, we conclude that $y'Ty'$ and $z'Tz'$. Now from property 4, we conclude xTy' and xTz'. So using property 5 yields $y' = z'$. Now choosing $w = y' = z'$ concludes the proof. □

When specifying completeness of a relation in `Carpa`, internally Theorem 7.1 is exploited to express this requirement in a SAT formula. We give a few examples.

Example 7.1. In Example 6.2, we saw that for two confluent relations R and S satisfying $R^{-1} \cdot S \subseteq S \cdot R^* \cdot (R^{-1})^*$ the union may be non-confluent. But in the given example the relation S is not terminating. Now, we wonder whether the same holds for R and S both being terminating. We do so by replacing the confluence requirements for R and S by completeness and applying `Carpa`: for $n = \#A \leqslant 7$ there is no solution, but for $n = 8$ the following example is generated

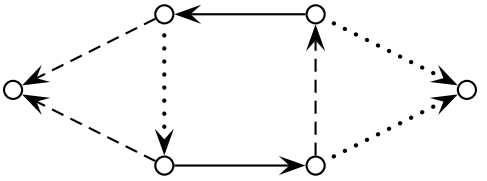


Here, again R steps are denoted by solid arrows and S steps are denoted by dashed arrows. Indeed all requirements hold, in fact even

$$R^{-1} \cdot S \subseteq S \cup (S \cdot R \cdot R^{-1}).$$

In Section 8, we describe in more detail how this example was obtained.

Example 7.2. If we have a set of binary relations of which the union of any two is complete, then it is easily seen that the union of all of them is locally confluent, as the union of any two is locally confluent. But can we conclude confluence of this union? This turns out to be not the case: `Carpa` finds the next example of three binary relations for which the union of any two of them is complete, but for which the union of all three is not confluent.

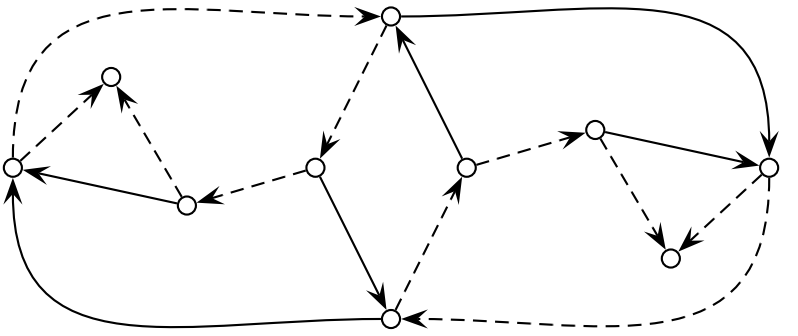


Here, the three relations are indicated by solid, dashed and dotted arrows, respectively.

Example 7.3. In Example 6.3, we saw that for two confluent relations R and S satisfying $R^{-1} \cdot S \subseteq (S \cup R^*) \cdot (R^{-1})^*$ the union may be non-confluent, being a stronger requirement than in Example 6.2. Both in Example 6.2 and Example 6.3, the relation S is not terminating. So also here it is a natural question whether confluence of the union may be concluded if moreover both R and S are terminating. Remember that in the origin of this question R corresponds to β-reduction in simply typed lambda calculus and S corresponds

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

to type computation, both being terminating. The following example, in which again R and S are denoted solid and dashed, respectively, shows that even for R and S both being terminating, the union does not need to be confluent.



This example was found several times for several variations of the set of requirements. A first approach before the representation of completeness based on Theorem 7.1 was invented used the approach of Sections 4 and 6 to express termination and confluence of R and S. In this approach, the SAT solver ran for hours without giving any result. This was the case for several variants of the problem, e.g., by slightly strengthening the peak requirement, or by requiring a term having two distinct normal forms rather than requiring non-confluence. By that time the result of Example 6.3 was already found, and the symmetry in this result was observed. So in a next attempt not only the given requirements were expressed in the formula, also a requirement of symmetry. More precisely, for $A = \{1, 2, 3, \ldots, 10\}$, for every $x, y \in A$ the requirements $(x, y) \in R \iff (11-x, 11-y) \in R$ and $(x, y) \in S \iff (11 - x, 11 - y) \in S$ were added. In this way, the SAT solver found a solution within seconds, from which the above example was extracted.

Using the representation of completeness based on Theorem 7.1 it turned out that without adding symmetry requirements also a solution could be found within seconds, for several variants of the specification of the problem. All solutions that we found on 10 elements turned out to coincide with the example given above, sometimes after removing redundant arrows.

One can wonder whether this number of 10 elements is minimal. For proving so, the formula for $n = 9$ should be unsatisfiable. After a few hours of computation indeed this was concluded for a formula expressing that R and S are complete, some element has two distinct normal forms with respect to $R \cup S$ and $R^{-1} \cdot S \subseteq (S \cup R) \cdot (R^{-1})^*$. Here, $k = 3$ was chosen: from the requirement that $R \cup S$ has distinct normal forms it can be concluded that $R^* = \bigcup_{i=0}^{8} R^i$ for $n = 9$.

8. The tool `Carpa`

We developed a tool called `Carpa` (Counter examples for Abstract Rewriting Produced Automatically) for entering a list of properties of binary relations, and then the tool either builds a set of binary relations on the specified number of elements that satisfies

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

these properties, or shows that this is impossible. The tool `Carpa` can be downloaded from

`http://www.win.tue.nl/~hzantema/carpa.html`

including the source code, a Linux executable, a file `Readme` with basic instructions, and encodings of all examples in this paper.

Internally `Carpa` does this via SAT solving and the techniques described in this paper. As the SAT solver, it uses `Yices` (Dutertre and de Moura), which is not only a SAT solver, but also an SMT solver (satisfiability modulo theories). We also developed a version of our tool generating the formulas in dimacs format and then calling the SAT solver `Minisat` instead. For simple examples both versions generate solutions within fraction of seconds; for harder examples the version generating formulas in SMT format and calling `Yices` turned out to be slightly more efficient. So we decided only to distribute this one.

We defined an input format in which all properties discussed in this paper can be specified directly, abstracting from the auxiliary relations that have to be introduced internally.

The input for `Carpa` always starts by two numbers n, m. Here, $n = \#A$ is the cardinality of the set A on which we search for binary relations, and m is the number of basic relations in the specification, identified by the numbers $1, \ldots, m$. So if we look for a single relation R with a given set of properties we choose $m = 1$, and if we look for two relations R and S with a given set of properties, we choose $m = 2$.

The rest of the input consists of a number of lines each being either a predicate or an assignment. In the following R, S refer to binary relations on A, identified by a number from 1 to m or a variable of the shape `x`i. Further x, y refer to elements of A, being numbers from 1 to n. The possible predicates are as follows:

— `subs`, where `subs`(R, S) means that $R \subseteq S$,
— `nsubs`, where `nsubs`(R, S) means that $\neg(R \subseteq S)$,
— `disj`, where `disj`(R, S) means that $R \cap S = \emptyset$,
— `trans`, where `trans`(R) means that R is transitive,
— `ntrans`, where `ntrans`(R) means that R is not transitive,
— `irr`, where `irr`(R) means that R is irreflexive,
— `nirr`, where `nirr`(R) means that R is not irreflexive,
— `symm`, where `symm`(R) means that R is symmetric,
— `sn`, where `sn`(R) means that R is terminating,
— `nsn`, where `nsn`(R) means that R is not terminating,
— `wn`, where `wn`(R) means that R is weakly normalizing (every element has at least one normal form),
— `nwn`, where `nwn`(R) means that R is not weakly normalizing,
— `cr`, where `cr`(R) means that R is confluent,
— `ncr`, where `ncr`(R) means that R is not confluent,
— `wcr`, where `wcr`(R) means that R is locally confluent,
— `nwcr`, where `nwcr`(R) means that R is not locally confluent,

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

— un, where un(R) means that R has the unique normal form property (every element has at least one normal form),
— nun, where nun(R) means that R does not have the unique normal form property,
— compl, where compl(R) means that R is complete,
— nf, where nf(x, R) means that x is a normal form with respect to R,
— red, where red(x, y, R) means that $(x, y) \in R$,
— nrrules, where nrrules(R, j) means that R has at most j elements, and
— nriter, where nriter(j) means that the number k used to define transitive closures based on Theorems 5.1 and 5.2 is replaced by j; its default value is $\lceil \log_2 n \rceil$.

It looks strange to have separate predicates for the negations of other predicates instead of having an operator for negation. We decided to do so since predicates like sn and compl internally introduce auxiliary relations as described in Sections 4 and 7, by which taking the negation of the generated formula is not equivalent to the negation of the intended property.

Assignments always consist of a variable name followed by the symbol '=', followed by an operation applied on a number of arguments. Here for variable names we always choose 'x' followed by a number, and the possible operations are as follows:

— union, where union(R, S) represents the relation $R \cup S$,
— inters, where inters(R, S) represents the relation $R \cap S$,
— comp, where comp(R, S) represents the relation $R \cdot S$,
— peak, where peak(R, S) represents the relation $R^{-1} \cdot S$,
— val, where val(R, S) represents the relation $R \cdot S^{-1}$,
— inv, where inv(R) represents the inverse R^{-1} of R,
— tc, where tc(R) represents the transitive closure R^+ of R,
— rc, where rc(R) represents the reflexive closure $R \cup I$ of R, and
— trc, where trc(R) represents the transitive reflexive closure $R^* = R^+ \cup I$ of R.

Here, the relations R, S should be either one of the basic relations, numbered $1, \ldots, m$, or a variable name that has been defined in an earlier assignment.

Our tool Carpa reads a list of requirements in this format, and builds a formula for it representing these requirements in the way as described in this paper. For every assignment a new binary relation is created. For every call of sn a new binary relation is created to represent S in Theorem 4.1, and to generate the corresponding requirements. For every call of compl two new binary relations are created to represent S and T in Theorem 7.1, and to generate the corresponding requirements. For every call of tc and trc k new binary relations are created to generate the requirements as described in Theorems 5.1 and 5.2. Call of cr and related properties are expanded internally according their standard definitions, calling trc.

In this format, we described the requirements for all examples as they occur in this paper, in fact all of the examples were found by applying our tool on the specifications written in this format.

For instance, for finding Example 6.1, a locally confluent relation on four elements that is not confluent, we choose $n = 4$ being the number of elements, $k = 2$ since that is the smallest value for which $2^k \geqslant n$, and $m = 1$ since we look for a single relation. Now local

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

confluence and non-confluence can be directly specified by `wcr` and `ncr`. Further, in order to avoid self-loops, we add the requirement that the relation is irreflexive, yielding the following input:

```
4
1
wcr(1)
ncr(1)
irr(1)
```

On this input `Carpa` first generates a propositional formula of 827 lines, describing exactly the given requirements. Then, it calls the SAT solver on this formula which yields satisfiability within a fraction of a second. Finally, `Carpa` inspects the satisfying assignment generated by the SAT solver and gives as output the desired relation:

```
Relation 1:
(1,2)
(1,3)
(2,1)
(2,4)
```

for which indeed coincides with Example 6.1.

Although internally the SAT solver plays a crucial role, the user of `Carpa` does not see this: he only calls `./carpa ex1` where `ex1` is a file containing the above input, and receives the above output immediately.

Internally requirements like `wcr` and `ncr` are transformed to a list of requirements first specifying transitive reflexive closure, then specifying new relations being peaks and valleys, and then specifying that such new relations are subsets or not. In this way, however, it may be the cases that the same relation is specified twice, making a formula being longer than necessary, and possibly making the SAT problem harder. In the above example, the transitive reflexive closure of the basic relation is specified twice. This can be avoided by specifying it only once explicitly, and then specify the peaks and valleys and subset requirements explicitly. So for the same problem, the following could be specified:

```
4
1
x1=trc(1)
x2=peak(1,1)
x3=val(x1,x1)
subs(x2,x3)
x4=peak(x1,x1)
nsubs(x4,x3)
irr(1)
```

Here, we introduce `x1=trc(1)`, being R^* for R being the basic relation indicated by number 1. Further, we need the local peak `x2=peak(1,1)`, the valley `x3=val(x1,x1)`, and the global peak `x4=peak(x1,x1)`. Local confluence states that `x2` should be a subset of `x3`, while the negation of confluence states that `x4` should not be a subset of `x3`. On

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

this input `Carpa` first generates a propositional formula of only 531 lines, and yields the same output, in an even smaller fraction of a second.

Next we consider Example 7.1; one way to achieve it is the following. Instead of non-confluence of the union we specify a slightly stronger requirement, namely that there exists an element with two distinct normal forms. As the input, we define

```
8
2
compl(1)
compl(2)
x1=union(1,2)
nf(2,x1)
nf(3,x1)
x2=tc(x1)
red(1,2,x2)
red(1,3,x2)
x1=trc(1)
x2=comp(2,x1)
x3=peak(1,2)
x4=val(x2,x1)
subs(x3,x4)
```

in which it is specified that the element 1 has two distinct normal forms 2 and 3 with respect to the union `x1` of the two basic relations. For the rest this input only consists of the requirements that 1 and 2 are complete, and that $R^{-1} \cdot S \subseteq S \cdot R^* \cdot (R^{-1})^*$ for R being 1 and S being 2. Note that in this example variable names are reused: at some point the union `x1` of 1 and 2 was not needed any more, and `x1` was redefined by `x1=trc(1)`. On this input `Carpa` generates a formula of 8119 lines within a fraction of a second, on which the SAT solver needs a few seconds to establish satisfiability. From the corresponding satisfying assignment `Carpa` generates the output

```
Relation 1:
(1,4)
(5,4)
(7,6)
(8,6)
Relation 2:
(1,3)
(4,3)
(4,8)
(5,7)
(6,2)
(6,5)
(7,2)
(8,1)
```

that indeed can be represented by the picture given in Example 7.1.

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

We conclude this section by some remarks on fine-tuning numbers of rules and numbers of iterations. Once a solution has been found it is natural to wonder whether this solutions is the smallest possible. For minimizing n, being the number of nodes in the graph, we typically decrease n until the resulting formula is unsatisfiable. But it is also natural to minimize the number of rules, begin the number of edges in the graph. For doing so, we add `nrrules(`R, j`)`, meaning that R has at most j elements, for decreasing numbers j until no solution is found any more. For instance, in Example 6.2 in the given solution R consists of two rules and S consists of six rules. Both numbers are shown to be minimal by showing that after adding either `nrrules(1,1)` or `nrrules(2,5)` to the input, in both cases `Carpa` reports `no solution`, in which 1 stands for R and 2 stands for S.

For the number k of iterations as occurring in Theorems 5.1 and 5.2, the default value is $\lceil \log_2 n \rceil$. This is always correct, but in some cases there are arguments by which a smaller value for k is correct too, and this may be more efficient. For instance, in Example 7.3 the number n is equal to 10, by which $\lceil \log_2 n \rceil = 4$. However, as the specification includes that $R \cup S$ has two distinct normal forms, in this case we can conclude that $P^* = \bigcup_{i=0}^{8} P^i$ for every $P \subseteq R \cup S$, by which choosing $k = 3$ is correct too. This is implemented by adding the statement `nriter(3)` before the transitive closures are computed, by which `Carpa` finds the solution much faster: in 16 s rather than 160 s. These figures involve instances not containing the requirements on numbers of rules.

9. Conclusions

This paper presents a method for automatically finding finite counterexamples for any list of abstract rewriting properties and describes a corresponding implementation. The basic idea is to fix the number n of elements of the set on which binary relations are searched for, and then build a propositional formula describing the properties. On this formula, a SAT solver is applied. If the formula is unsatisfiable then no example exists satisfying the given properties. If the formula is satisfiable then from the corresponding satisfying assignment an example is extracted satisfying the given properties. An implementation following this approach shows to be successful for various examples, typically up to around $n = 10$, including examples that are very hard to find by hand. The formulas are made in such a way that for every line in the list of properties at most $O(n^2 \log n)$ fresh Boolean variables are created, and the contribution to the size of the formula is at most $O(n^3 \log n)$. Although in SAT solving a restricted size of the formulas does not guarantee a quick solution at all, avoiding a combinatorial explosion in the size of the formula is important.

In our implementation `Carpa` we restricted to basic notions like termination, confluence, completeness, normal forms, transitive closures and properties that can be expressed as compositions, peaks and valleys, transitive closures and subset relations. The main reason for this is that the main properties of our interest can be expressed in these notions. As soon as other notions come up that can be expressed in our setting, our implementation may be easily extended accordingly.

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

A general observation in SAT solving is that among similar formulas some of which are satisfiable and other are not, proving unsatisfiability is harder than proving satisfiability. In our experiments, this was confirmed: for n being the smallest number for which there exists an example for a given list of properties, typically finding such an example by our tool is done much faster than proving that such an example does not exist for $n-1$.

In our encodings for termination, we needed one auxiliary relation, for completeness we needed two, and for transitive closure and confluence we needed $\log(n)$ auxiliary relations. We conjecture that it is not possible to fully specify transitive closure or confluence only using a constant number of auxiliary relations. It was observed by Bertram Felgenhauer that by mimicking the Floyd–Warshall algorithm, the transitive closure can be specified by a formula of size $O(n^3)$, but this needs $\Omega(n^3)$ auxiliary Boolean variables. Experiments on this alternative encoding for transitive closure do not show remarkable differences in efficiency.

The approach of this paper restricts to finite abstract reduction systems. For some combinations of properties, for instance termination of a relation and non-termination of its inverse, no finite abstract reduction system exists, while in term rewriting easily an example is found, in this example the single rule $f(a) \to a$. In a follow-up (Zantema 2013) of the research presented in this paper, it is investigated how to find such term rewriting systems automatically by the tool `Carpa+`.

We want to thank Bas Joosten for his contribution to Theorem 7.1, Bertram Felgenhauer and René Thiemann for fruitful discussions, and the anonymous reviewers for their detailed and fruitful comments.

References

Baader, F. and Nipkow, T. (1998). *Term Rewriting and All That*, Cambridge University Press.

Bachmair, L. and Dershowitz, N. (1986). Commutation, transformation and termination. In: Siekmann, J. (eds.) *Proceedings of the 8th International Conference on Automated Deduction (CADE-8)*, Lecture Notes in Computer Science, vol. 230, Springer, 5–20.

Bellegarde, F. and Lescanne, P. (1990). Termination by completion. *Applicable Algebra in Engineering, Communication and Computing* **1** (2) 79–96.

Doornbos, H. and von Karger, B. (1998). On the union of well-founded relations. *Logic Journal of the IGPL* **6** (2) 195–201.

Dutertre, B. and de Moura, L. Yices: An SMT solver. Available at `http://yices.csl.sri.com/`.

Stump, A., Kimmell, G. and El Haj Omar, R. (2011). Type preservation as a confluence problem. In: Schmidt-Schauß, M. (ed.) *Proceedings of the 22nd International Conference on Rewriting Techniques and Applications*, LIPIcs, vol. 10, 345–360.

Stump, A., Kimmell, G., Zantema, H. and El Haj Omar, R. (2013). A rewriting view of simple typing. *Logical Methods in Computer Science* **9** (1).

Terese. (2003). *Term Rewriting Systems*, Cambridge Tracts in Theoretical Computer Science, vol. 55, Cambridge University Press.

van Oostrom, V. (1994). Confluence by decreasing diagrams, *Theoretical Computer Science* **126** (2) 259–280.

van Oostrom, V. (2008). Confluence by decreasing diagrams, converted. In: Voronkov, A. (ed.) *Proceedings of the 19th Conference on Rewriting Techniques and Applications*, Lecture Notes in Computer Science, vol. 5117, Springer, 306–320.

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press

van Oostrom, V. and Zantema, H. (2012). Triangulation in rewriting. In: Tiwari, A. (eds.) *Proceedings of the 23rd International Conference on Rewriting Techniques and Applications*, volume 15 of *Leibniz International Proceedings in Informatics (LIPIcs)*, Leibniz International Proceedings in Informatics (LIPIcs), vol. 15, Dagstuhl, Germany, Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 240–255.

Zantema, H. (2013). Automatically finding particular term rewriting systems. Available via `http://www.win.tue.nl/~hzantema/carpa.html`.

https://doi.org/10.1017/S0960129518000221 Published online by Cambridge University Press