

Transforming SAT into Termination of Rewriting¹

Harald Zankl² Christian Sternagel³ Aart Middeldorp⁴

*Institute of Computer Science
University of Innsbruck
Innsbruck, Austria*

Abstract

In this paper we propose different translations from SAT to termination of term rewriting, i.e., we translate a propositional formula φ into a generic rewrite system \mathcal{R}^φ with the property that φ is satisfiable if and only if \mathcal{R}^φ is (non)terminating. Our experiments reveal that the generated rewrite systems are challenging for automated termination provers. Furthermore, a large class of them seems to be just unprovable by current methods implemented in termination analyzers.

Keywords: term rewriting, termination, semantic labeling, SAT solving

1 Introduction

Termination of term rewrite systems (TRSs) is an undecidable property [12]. Nevertheless, nowadays powerful (incomplete) algorithms exist that can prove termination of many rewrite systems as can be witnessed by the international termination competition.⁵ In 2004 Kurihara and Kondo were the first who encoded a termination method in propositional logic [19] and in 2006 the first tools (Jambox⁶ and Matchbox [21]) employed SAT-solving techniques in the competition. They surprised the community by the gains in power and speed. Their success was mainly due to the so-called matrix-method [8] which can effectively be implemented using SAT-solvers. But even for very simple and ancient methods like the lexicographic path order [13,6] (LPO) the recent development in the SAT community allows way faster implementations [4] than some years ago. A similar speedup [22] is achieved for the Knuth-Bendix order (KBO) [14]. This is remarkable because KBO orientability is

¹ This research is supported by FWF (Austrian Science Fund) project P18763.

² Email: harald.zankl@uibk.ac.at

³ Email: christian.sternagel@uibk.ac.at

⁴ Email: aart.middeldorp@uibk.ac.at

⁵ <http://www.lri.fr/~marche/termination-competition/>

⁶ Available from <http://joerg.endrullis.de>.

known to be decidable in polynomial time [18] whereas SAT is NP-complete [5]. In other words, the sophisticated algorithms for solving the computationally harder (unless $P = NP$) problem SAT outperform the dedicated methods for KBO [7,18]. In this paper we address the question whether a similar result also holds when translating the NP-complete SAT problem into the undecidable termination property of TRSs. However, the experiments reveal that at least for our translations the results are as expected. Concerning the transformation from SAT to termination, the dedicated SAT approaches perform much better. Even further, only the most simple propositional formulas produce TRSs which can be shown terminating by state-of-the-art termination provers. Therefore the translations can be used to generate a large set of difficult termination problems automatically.

The rest of the paper is organized as follows: In Section 2 propositional formulas are introduced and many-sorted rewriting is defined. In Section 3 we define TRSs \mathcal{U}^φ that are terminating if and only if the propositional formula φ is unsatisfiable. In Section 4 the dual problem is considered for many-sorted TRSs \mathcal{S}^φ and \mathcal{T}^φ that are terminating if and only if φ is satisfiable. That even simple propositional formulas produce TRSs where termination analysis is challenging is demonstrated in Section 5 where it also becomes apparent that narrowing [10] is one method which can handle small instances. We conclude in Section 6.

2 Preliminaries

In this section we fix basic notation concerning propositional logic, introduce many-sorted TRSs and define the model variant of semantic labeling [24] in a many-sorted setting. Aoto and Yamada [1] already generalized semantic labeling to many-sorted rewriting but just for the quasi-model case.

2.1 Propositional Logic

Let \mathcal{A} be a set of propositional variables (atoms). Sometimes we find it convenient to abbreviate the set of atoms p_1, \dots, p_n by \mathcal{A}_n . The set of propositional formulas $\mathcal{P}(\mathcal{A})$ is inductively defined by the following BNF

$$\varphi ::= p \in \mathcal{A} \mid (\varphi \wedge \psi) \mid (\neg\varphi)$$

Note that we do not allow disjunction (which does not pose a restriction but allows to keep the presentation concise). The following convention is used to reduce the number of parentheses: (i) Outermost parentheses are omitted, (ii) ‘ \wedge ’ is left-associative, and (iii) ‘ \neg ’ binds stronger than ‘ \wedge ’.

Let $\mathbb{B} := \{0, 1\}$. An *assignment* is a mapping $\alpha: \mathcal{A} \rightarrow \mathbb{B}$. It is lifted to an interpretation of formulas: $[\alpha]: \mathcal{P}(\mathcal{A}) \rightarrow \mathbb{B}$ with

$$[\alpha](\varphi) = \begin{cases} \alpha(p) & \text{if } \varphi = p \text{ for some } p \in \mathcal{A} \\ [\alpha](\psi) \cdot [\alpha](\chi) & \text{if } \varphi = \psi \wedge \chi \\ \overline{[\alpha](\psi)} & \text{if } \varphi = \neg\psi \end{cases}$$

Here $(\cdot): \mathbb{B} \times \mathbb{B} \rightarrow \mathbb{B}$ is defined as $x \cdot y = 1$ if and only if $x = y = 1$ and for $\overline{(\cdot)}: \mathbb{B} \rightarrow \mathbb{B}$

we have $\bar{x} = 1$ if and only if $x = 0$. A formula φ is *satisfiable* (*unsatisfiable*) if an (no) assignment α exists such that $[\alpha](\varphi) = 1$. This problem is known as the satisfiability problem (SAT). For a propositional formula φ its *depth* is defined as follows: $\text{depth}(p) = 0$ for $p \in \mathcal{A}$, $\text{depth}(\varphi \wedge \psi) = 1 + \max(\text{depth}(\varphi), \text{depth}(\psi))$, and $\text{depth}(\neg\varphi) = 1 + \text{depth}(\varphi)$. Similarly the set of variables $\mathcal{V}\text{ar}(\varphi)$ is defined recursively by: $\mathcal{V}\text{ar}(p) = \{p\}$ for $p \in \mathcal{A}$, $\mathcal{V}\text{ar}(\varphi \wedge \psi) = \mathcal{V}\text{ar}(\varphi) \cup \mathcal{V}\text{ar}(\psi)$, and $\mathcal{V}\text{ar}(\neg\varphi) = \mathcal{V}\text{ar}(\varphi)$. The well-known coincidence lemma states that when testing φ for satisfiability only the (finitely many) variables that actually occur in φ have to be considered which makes SAT decidable because the search space becomes finite. Furthermore, this allows us to relate assignments to substitutions (whose domain must be finite by definition) in the next section. Despite the fact that the search space for a satisfying assignment is finite, deciding SAT is difficult, more precisely, SAT is an NP-complete problem [5].

2.2 Many-Sorted Semantic Labeling

We assume basic familiarity with term rewriting [3]. Let \mathcal{S} be a non-empty set of *sorts*. An \mathcal{S} -sorted signature is a set of function symbols \mathcal{F} , where each $f \in \mathcal{F}$ of arity n is associated with the *function signature* $\text{sig}: \mathcal{F} \rightarrow \mathcal{S}^{n+1}$. Here the first n components of $\text{sig}(f)$ give the sort (type) of each argument and the last gives the sort of the function's result. In the following, we write $f : s_1 \times \cdots \times s_n \rightarrow s_{n+1}$, to express that f has (function) signature (s_1, \dots, s_{n+1}) .

An \mathcal{S} -sorted set A is a family of sets $\{A_s\}_{s \in \mathcal{S}}$. For an \mathcal{S} -sorted set \mathcal{V} of *variables* (where $\mathcal{V}_s \cap \mathcal{V}_t = \emptyset$ for $s \neq t$), let $\mathcal{T}(\mathcal{F}, \mathcal{V})_s$ denote the set of terms with sort s over \mathcal{F} and \mathcal{V} , which is defined inductively by the rules

$$\frac{x \in \mathcal{V}_s \quad f \in \mathcal{F} \quad f : s_1 \times \cdots \times s_n \rightarrow s \quad t_i \in \mathcal{T}(\mathcal{F}, \mathcal{V})_{s_i}}{x \quad f(t_1, \dots, t_n)}$$

This yields the \mathcal{S} -sorted set $\mathcal{T}(\mathcal{F}, \mathcal{V}) = \{\mathcal{T}(\mathcal{F}, \mathcal{V})_s\}_{s \in \mathcal{S}}$. Associated with every term $t \in \mathcal{T}(\mathcal{F}, \mathcal{V})$ is its *sort*, i.e., if $t \in \mathcal{T}(\mathcal{F}, \mathcal{V})_s$ then $\text{sort}(t) = s$. An \mathcal{S} -sorted TRS \mathcal{R} is an \mathcal{S} -sorted set of pairs $(l, r) \in \mathcal{R}_s$ —the so called *rewrite rules*—written as $l \rightarrow r$, such that there exists an $s \in \mathcal{S}$ with $l, r \in \mathcal{T}(\mathcal{F}, \mathcal{V})_s$ and the usual restrictions that l is not a variable and all variables in r also occur in l are satisfied. In the sequel we identify one-sorted TRSs with unsorted ones and feel free to omit sort information where it is not essential.

Let \mathcal{F} be an \mathcal{S} -sorted signature. An \mathcal{S} -sorted \mathcal{F} -algebra \mathcal{A} consists of an \mathcal{S} -sorted *carrier* A (where each $A_s \in A$ is non-empty) and a set of *interpretations* $\{f_{\mathcal{A}}\}_{f \in \mathcal{F}}$, such that for each function symbol $f : s_1 \times \cdots \times s_n \rightarrow s_{n+1}$ there is an interpretation $f_{\mathcal{A}} : A_{s_1} \times \cdots \times A_{s_n} \rightarrow A_{s_{n+1}}$. An \mathcal{S} -sorted *substitution* $\sigma : \mathcal{V} \rightarrow \mathcal{T}(\mathcal{F}, \mathcal{V})$ is a set of mappings $\sigma_s : \mathcal{V}_s \rightarrow \mathcal{T}(\mathcal{F}, \mathcal{V})_s$ for every $s \in \mathcal{S}$ such that $\sigma(x) \neq x$ only for finitely many $x \in \mathcal{V}$. An \mathcal{S} -sorted *assignment* $\alpha : \mathcal{V} \rightarrow A$ is a set of mappings $\alpha_s : \mathcal{V}_s \rightarrow A_s$ for every $s \in \mathcal{S}$. For every \mathcal{S} -sorted term t and assignment $\alpha : \mathcal{V} \rightarrow A$, a mapping $[\alpha]_{\mathcal{A}} : \mathcal{T}(\mathcal{F}, \mathcal{V}) \rightarrow A$ is defined inductively

$$[\alpha]_{\mathcal{A}}(t) = \begin{cases} \alpha_s(x) & \text{if } t = x \text{ and } \text{sort}(t) = s \\ f_{\mathcal{A}}([\alpha]_{\mathcal{A}}(t_1), \dots, [\alpha]_{\mathcal{A}}(t_n)) & \text{if } t = f(t_1, \dots, t_n) \end{cases}$$

An \mathcal{S} -sorted \mathcal{F} -algebra is a *model* of an \mathcal{S} -sorted TRS, if for all \mathcal{S} -sorted assignments α and rewrite rules $l \rightarrow r \in \mathcal{R}$ it holds that $[\alpha]_{\mathcal{A}}(l) = [\alpha]_{\mathcal{A}}(r)$. A *labeling* L chooses for every $f \in \mathcal{F}$ a set of *labels* L_f . The *labeled signature* is defined by

$$\mathcal{F}_{\text{lab}} = \{f \mid f \in \mathcal{F}, L_f = \emptyset\} \cup \{f_a \mid f \in \mathcal{F}, a \in L_f\}$$

where the arity and function signature of f_a and f coincide. A *labeling* ℓ for an \mathcal{S} -sorted algebra \mathcal{A} consists of a labeling L together with a *labeling function* $\ell_f: A_{s_1} \times \dots \times A_{s_n} \rightarrow L_f$ for every $f \in \mathcal{F}$ with $L_f \neq \emptyset$ and $\text{sig}(f) = s_1 \times \dots \times s_n \rightarrow s_{n+1}$. Let $A^{\mathcal{V}}$ denote the set of all \mathcal{S} -sorted assignments from \mathcal{V} to A . Let ℓ be a labeling for \mathcal{A} . For every assignment $\alpha \in A^{\mathcal{V}}$ a mapping $\text{lab}_{\alpha}: \mathcal{T}(\mathcal{F}, \mathcal{V}) \rightarrow \mathcal{T}(\mathcal{F}_{\text{lab}}, \mathcal{V})$ is defined inductively as follows

$$\text{lab}_{\alpha}(t) = \begin{cases} x & \text{if } t = x \\ f(\text{lab}_{\alpha}(t_1), \dots, \text{lab}_{\alpha}(t_n)) & \text{if } t = f(t_1, \dots, t_n) \text{ and } L_f = \emptyset \\ f_a(\text{lab}_{\alpha}(t_1), \dots, \text{lab}_{\alpha}(t_n)) & \text{if } t = f(t_1, \dots, t_n) \text{ and } L_f \neq \emptyset \end{cases}$$

with $a = \ell_f([\alpha]_{\mathcal{A}}(t_1), \dots, [\alpha]_{\mathcal{A}}(t_n))$. For any \mathcal{S} -sorted TRS \mathcal{R} over \mathcal{F} , together with an \mathcal{F} -algebra \mathcal{A} and a labeling ℓ , the \mathcal{S} -sorted TRS \mathcal{R}_{lab} over \mathcal{F}_{lab} is given by

$$\mathcal{R}_{\text{lab}} = \{\text{lab}_{\alpha}(l) \rightarrow \text{lab}_{\alpha}(r) \mid l \rightarrow r \in \mathcal{R}, \alpha \in A^{\mathcal{V}}\}$$

An \mathcal{S} -sorted TRS \mathcal{R} is *terminating* if it does not admit an infinite rewrite sequence $t_1 \rightarrow_{\mathcal{R}} t_2 \rightarrow_{\mathcal{R}} \dots$ starting at some $t_1 \in \mathcal{T}(\mathcal{F}, \mathcal{V})_s$ for some $s \in \mathcal{S}$.

Theorem 2.1 *Let \mathcal{R} be an \mathcal{S} -sorted TRS. Let the algebra \mathcal{A} be a model of \mathcal{R} and let ℓ be a labeling for \mathcal{A} . Then \mathcal{R} is terminating if and only if \mathcal{R}_{lab} is terminating.*

For the TRSs we are dealing with in the subsequent sections, many-sorted termination is equivalent to the one-sorted case [23] since the systems are non-collapsing. A TRS is collapsing if it contains a rule $l \rightarrow x$ for some variable x . Restricting to many-sortedness simplifies the proofs of Theorems 4.2 and 4.6 considerably.

3 Transforming Unsatisfiability to Termination

In the following we want to express SAT as a termination problem in rewriting, i.e., given a formula φ , we construct a TRS R^{φ} that is terminating if and only if φ is satisfiable. This transformation is addressed in the next section. First we focus on the simpler dual problem, namely the construction of a TRS \mathcal{R}^{φ} that is terminating if and only if φ is unsatisfiable.

For this purpose we consider a $\{\text{bool}\}$ -sorted signature $\mathcal{F} = \{\star, -\}$ containing a binary function symbol $(\star) : \text{bool} \times \text{bool} \rightarrow \text{bool}$ and a unary function symbol $- : \text{bool} \rightarrow \text{bool}$. Furthermore we assume that the propositional atoms in \mathcal{A} are contained in the set of term variables $\mathcal{V}_{\text{bool}}$. Although ‘ \star ’ will represent (on the term level) the same as ‘ \wedge ’ does on formulas, we use different function symbols because we want to clearly separate between the two different concepts. The same holds for the symbols ‘ $-$ ’ and ‘ \neg ’. The obvious encoding $\lceil \star \rceil: \mathcal{P}(\mathcal{A}) \rightarrow \mathcal{T}(\{\star, -\}, \mathcal{V})$ transforms formulas into terms as follows: $\lceil p \rceil = p$ for $p \in \mathcal{A}$, $\lceil \varphi \wedge \psi \rceil = \lceil \varphi \rceil \star \lceil \psi \rceil$,

and $\lceil \neg\varphi \rceil = -\lceil \varphi \rceil$. Now every well-formed formula in $\mathcal{P}(\mathcal{A})$ has a corresponding term representation in $\mathcal{T}(\mathcal{F}, \mathcal{V})$.

The next goal is to mimic the task of assignments for formulas on the term level. Thus the signature \mathcal{F} is extended by two constant symbols of sort `bool`, namely ‘ \perp ’ and ‘ \top ’. We say that an assignment $\alpha: \mathcal{A}_n \rightarrow \mathbb{B}$ and a substitution $\sigma: \mathcal{V}_n \rightarrow \{\perp, \top\}$ are *corresponding* if $\alpha(p_i) = 0$ if and only if $\sigma(p_i) = \perp$ for all $1 \leq i \leq n$ (here $\mathcal{V}_n = \mathcal{A}_n$).

In order to perform the work $[\alpha]$ does on formulas the six rewrite rules

$$\perp \star \perp \rightarrow \perp \quad \perp \star \top \rightarrow \perp \quad \top \star \perp \rightarrow \perp \quad \top \star \top \rightarrow \top \quad -\perp \rightarrow \top \quad -\top \rightarrow \perp$$

referred to as the TRS \mathcal{S}_{imp} are employed. The next lemma formalizes the interplay of assignments and substitutions.

Lemma 3.1 *Let $\varphi \in \mathcal{P}(\mathcal{A}_n)$ and $t \in \mathcal{T}(\mathcal{F}, \mathcal{V})$ such that $\lceil \varphi \rceil = t$. If the assignment α and the substitution σ are corresponding, then*

- $[\alpha](\varphi) = 0$ implies $t\sigma \rightarrow_{\mathcal{S}_{\text{imp}}}^* \perp$ and dually
- $[\alpha](\varphi) = 1$ implies $t\sigma \rightarrow_{\mathcal{S}_{\text{imp}}}^* \top$

Proof By induction on the structure of φ . □

The following example already contains the main idea for constructing nonterminating sequences.

Example 3.2 Consider the formula $\varphi = p_1 \wedge \neg p_2$ with the corresponding term $t = \lceil \varphi \rceil = p_1 \star (-p_2)$. Then the TRS \mathcal{S}_{imp} together with the rewriting rule

$$\text{unsat}(p_1, p_2, \top) \rightarrow \text{unsat}(p_1, p_2, p_1 \star (-p_2))$$

admits the cyclic reduction

$$\text{unsat}(\top, \perp, \top) \rightarrow \text{unsat}(\top, \perp, \top \star (-\perp)) \rightarrow \text{unsat}(\top, \perp, \top \star \top) \rightarrow \text{unsat}(\top, \perp, \top)$$

which proves nontermination of this TRS. The reason for nontermination is that for a satisfying assignment α (in this case $\alpha(p_1) = 1$ and $\alpha(p_2) = 0$) there is a *corresponding* substitution σ such that the term $\lceil \varphi \rceil \sigma = t\sigma$ rewrites to \top by Lemma 3.1.

The next theorem formally establishes the relation between satisfiable formulas and nontermination of corresponding TRSs.

Theorem 3.3 *Let $\varphi \in \mathcal{P}(\mathcal{A}_n)$. The parametrized TRS \mathcal{U}^φ that consists of all rules in \mathcal{S}_{imp} and*

$$\text{unsat}(p_1, \dots, p_n, \top) \rightarrow \text{unsat}(p_1, \dots, p_n, \lceil \varphi \rceil) \tag{1}$$

is terminating if and only if φ is unsatisfiable.

Proof For the direction from left to right assume \mathcal{U}^φ to be terminating and φ to be satisfiable to arrive at a contradiction. Since φ is satisfiable there must be a

$$\begin{array}{ll}
 \perp_{\mathcal{B}} = 0 & \top_{\mathcal{B}} = 1 \\
 \star_{\mathcal{B}}(x, y) = x \cdot y & -_{\mathcal{B}}(x) = \bar{x} \\
 \text{unsat}_{\mathcal{B}}(p_1, \dots, p_n, y) = 0 &
 \end{array}$$

Table 1
A model for the TRS \mathcal{U}^φ .

satisfying assignment α and a corresponding substitution σ . But then there is the cyclic reduction

$$t = \text{unsat}(\sigma(p_1), \dots, \sigma(p_n), \top) \rightarrow \text{unsat}(\sigma(p_1), \dots, \sigma(p_n), \lceil \varphi \rceil \sigma) \rightarrow^* t$$

where the first rewrite step is an application of rule (1) and the rest of the sequence holds by Lemma 3.1 since $\lceil \varphi \rceil \sigma \rightarrow_{\text{Simp}}^* \top$. Contradiction.

For the direction from right to left we assume unsatisfiability of φ and show termination of \mathcal{U}^φ . For this purpose we apply semantic labeling. Note that we consider \mathcal{U}^φ as one-sorted. The idea is to label the symbol unsat by the value which φ evaluates to—under all possible assignments. To obtain a model, the function symbols are interpreted in the Boolean algebra. The interpretation \mathcal{B} over the carrier \mathbb{B} depicted in Table 1 is a model for \mathcal{U}^φ . Next the labeling for \mathcal{U}^φ is defined. For this purpose only the function symbol unsat gets labeled, i.e., $L_\star = L_- = \emptyset$ ⁷ and $L_{\text{unsat}} = \mathbb{B}$. The labeling function $\ell_{\text{unsat}} : \mathbb{B}^{n+1} \rightarrow \mathbb{B}$ is defined as: $\ell_{\text{unsat}}(p_1, \dots, p_n, y) = y$. By assumption the formula φ evaluates to 0 under all assignments. Hence the labeled variant of rule (1) looks like

$$\text{unsat}_1(p_1, \dots, p_n, \top) \rightarrow \text{unsat}_0(p_1, \dots, p_n, \lceil \varphi \rceil) \quad (2)$$

Termination of the labeled system can then easily be shown by some basic method, e.g., LPO; choosing the precedence $\text{unsat}_1 > \text{unsat}_0, \star, -$ allows to orient rule (2) from left to right and $- > \perp, \top$ handles the rules in Simp . So $\mathcal{U}_{\text{lab}}^\varphi$ is terminating. Theorem 2.1 yields the termination of \mathcal{U}^φ . \square

4 Transforming Satisfiability to Termination

In the previous section the task was somehow simpler since there it sufficed to construct a nonterminating sequence if there exists a satisfying assignment. Hence by guessing a satisfying assignment for φ one could construct an infinite sequence in the TRS \mathcal{U}^φ . In this section the endeavor is more challenging, because one has to guarantee that one cycles if *no* satisfying assignment exists. Hence, the parametrized TRS will have to test all assignments before entering a loop if none of them satisfied the formula φ . Thus we have to provide the possibility to generate all assignments successively. The following three rules, referred to as \mathcal{N} ext do this job by representing assignments as bitlists (consequently the signature \mathcal{F} is extended by the binary function symbol $(::) : \text{bool} \times \text{list} \rightarrow \text{list}$, the constant $\text{nil} : \text{list}$, and the

⁷ Labeling constants is superfluous and hence we implicitly set $L_\perp = L_\top = \emptyset$.

unary function symbol $\text{next} : \text{list} \rightarrow \text{list}$):

$$\begin{aligned} \text{next}(\text{nil}) &\rightarrow \text{nil} \\ \text{next}(\perp :: xs) &\rightarrow \top :: xs \\ \text{next}(\top :: xs) &\rightarrow \perp :: \text{next}(xs) \end{aligned}$$

To ease notation we will encode lists over \perp and \top as natural numbers. Therefore, lists are interpreted as little endian representation of binary numbers where \perp corresponds to 0 and \top to 1. Let \mathcal{G} be the signature $\{\perp, \top, ::, \text{nil}\}$. The mapping $\text{enc} : \mathcal{T}(\mathcal{G}) \rightarrow \mathbb{N} \times \mathbb{N}$, $\text{enc}(\text{nil}) = (0, 0)$ and $\text{enc}(x :: xs) = (x + 2i, l + 1)$ where $\text{enc}(xs) = (i, l)$, uniquely associates lists with entries \perp or \top to pairs. The first component of the pair is the little endian representation of the bitlist whereas the second component is the length of the list. For convenience we denote (i, l) by \mathbf{i}_l . Furthermore if l is irrelevant or fixed we feel free to omit it. Taking the above conventions into account the bitlist $[\top; \perp; \top; \top]$ ⁸ can be written as $\mathbf{13}_4$ or more sloppily as $\mathbf{13}$. But we do not only identify these bitlists with natural numbers, they also encode substitutions. Hence, a bitlist $[t_1; \dots; t_n]$ gives rise to a substitution σ with $\sigma(p_i) = t_i$ for $1 \leq i \leq n$. Using this convention a term t indexed with a bold face integer denotes the result of applying the substitution to the term, i.e., $(p_1 \star ((-p_2) \star p_3))_{\mathbf{13}}$ denotes $\top \star ((-\perp) \star \top)$.

Lemma 4.1 *For a bitlist t , $\text{next}(t)$ rewrites to the successor of t :*

$$\text{If } \text{enc}(t) = \mathbf{i}_l \text{ then } \text{next}(t) \rightarrow_{\mathcal{N}_{\text{ext}}}^* t' \text{ with } \text{enc}(t') = (\mathbf{i} + \mathbf{1} \bmod \mathbf{2}^l)_l.$$

Proof By induction on the structure of t and unfolding the definition of \mathbf{i}_l . \square

To proceed we explicitly state the function signature sig , i.e., the sort for each function symbol, in the left column of Table 2. In the theorem below the variables p_1, \dots, p_n are of sort bool and xs is of sort list .

Theorem 4.2 *Let $\varphi \in \mathcal{P}(\mathcal{A}_n)$. Then the parametrized $\{\text{bool}, \text{list}\}$ -sorted TRS \mathcal{S}^φ that contains all rules of Simp , \mathcal{N}_{ext} , and additionally*

$$\text{sat}([p_1; \dots; p_n], \perp) \rightarrow \text{sat}(\text{next}([p_1; \dots; p_n]), \ulcorner \varphi \urcorner) \quad (3)$$

is terminating if and only if the formula φ is satisfiable.

Proof For the direction from left to right assume for the sake of a contradiction unsatisfiability of φ . The cyclic reduction

$$\begin{aligned} \text{sat}(\mathbf{0}, \perp) &\rightarrow \\ \text{sat}(\text{next}(\mathbf{0}), \ulcorner \varphi \urcorner_{\mathbf{0}}) &\rightarrow^* \text{sat}(\mathbf{1}, \ulcorner \varphi \urcorner_{\mathbf{0}}) \rightarrow^* \text{sat}(\mathbf{1}, \perp) \rightarrow^* \dots \rightarrow^* \\ \text{sat}(\text{next}(\mathbf{2}^n - \mathbf{1}), \ulcorner \varphi \urcorner_{\mathbf{2}^n - \mathbf{1}}) &\rightarrow^* \text{sat}(\mathbf{0}, \ulcorner \varphi \urcorner_{\mathbf{2}^n - \mathbf{1}}) \rightarrow^* \text{sat}(\mathbf{0}, \perp) \end{aligned}$$

proves nontermination of \mathcal{S}^φ where $\text{next}(\mathbf{i}_n) \rightarrow_{\mathcal{N}_{\text{ext}}}^* (\mathbf{i} + \mathbf{1} \bmod \mathbf{2}^n)_n$ follows from Lemma 4.1 and since we assumed that φ is unsatisfiable $\ulcorner \varphi \urcorner_{\mathbf{i}_n} \rightarrow_{\text{Simp}}^* \perp$ by Lemma 3.1, for all $0 \leq i < 2^n$.

⁸ To ease readability, lists $x :: (y :: (z :: \text{nil}))$ are abbreviated by $[x; y; z]$.

$\perp : \text{bool}$	$\perp_{\mathcal{A}} = 0$
$\top : \text{bool}$	$\top_{\mathcal{A}} = 1$
$\star : \text{bool} \times \text{bool} \rightarrow \text{bool}$	$\star_{\mathcal{A}}(x, y) = x \cdot y$
$- : \text{bool} \rightarrow \text{bool}$	$-_{\mathcal{A}}(x) = \bar{x}$
$\text{nil} : \text{list}$	$\text{nil}_{\mathcal{A}} = (0, 0)$
$:: : \text{bool} \times \text{list} \rightarrow \text{list}$	$::_{\mathcal{A}}(x, (i, l)) = (x + 2i, l + 1)$
$\text{next} : \text{list} \rightarrow \text{list}$	$\text{next}_{\mathcal{A}}((i, l)) = (i + 1 \bmod 2^l, l)$
$\text{sat} : \text{list} \times \text{bool} \rightarrow \text{bool}$	$\text{sat}_{\mathcal{A}}((i, l), b) = 0$

Table 2
A model for the $\{\text{bool}, \text{list}\}$ -sorted TRS \mathcal{S}^φ .

For the direction from right to left we will again give a proof using semantic labeling. The difference this time is that we exploit the many-sorted version of semantic labeling. Now for every sort $s \in \{\text{bool}, \text{list}\}$ we have to specify a carrier. The choices are $A_{\text{bool}} = \mathbb{B}$ and $A_{\text{list}} = P := \{(i, l) \in \mathbb{N} \times \mathbb{N} \mid i < 2^l\}$. Then the interpretation in the right column of Table 2 is a model for \mathcal{S}^φ . We show this for the $\mathcal{N}\text{ext}$ -rules. Let us fix an arbitrary value $(i, l) \in P$ for xs . The three $\mathcal{N}\text{ext}$ -rules generate the three equalities

$$(0 + 1 \bmod 2^0, 0) = (0, 0) \quad (4)$$

$$((0 + 2i) + 1 \bmod 2^{l+1}, l + 1) = (1 + 2i, l + 1) \quad (5)$$

$$((1 + 2i) + 1 \bmod 2^{l+1}, l + 1) = (0 + 2(i + 1 \bmod 2^l), l + 1) \quad (6)$$

Equation (4) is trivially valid. Since $i < 2^l$ by definition of P equation (5) holds since the modulo operation can be omitted. Validity of equation (6) is shown by case distinction. For $i = 2^l - 1$ it simplifies to $1 + 2(2^l - 1) + 1 \bmod 2^{l+1} = 0 + 2(2^l - 1 + 1 \bmod 2^l)$ where both sides equal 0. For the other case we know that $i < 2^l - 1$ and consequently $2i + 2 < 2^{l+1}$. Hence, the modulo operation does no harm and both sides evaluate to the same value.

The following sets of labels are employed: $L_\star = L_- = L_{::} = L_{\text{next}} = \emptyset$ and $L_{\text{sat}} = \mathbb{N} \times \mathbb{B}$. Then, the labeling function $\ell_{\text{sat}} : P \times \mathbb{N} \rightarrow \mathbb{N} \times \mathbb{B}$ with $\ell_{\text{sat}}((i, l), b) = (i, b)$ is used which produces the following labeled variants of rule (3)

$$\text{sat}_{i,0}([p_1; \dots; p_n], \perp) \rightarrow \text{sat}_{(i+1 \bmod 2^n), \varphi_{\mathbf{i}_n}}(\text{next}([p_1; \dots; p_n]), \lceil \varphi \rceil) \quad (7)$$

where $0 \leq i < 2^n$. In the right-hand side of the generic rule (7) the expression $\varphi_{\mathbf{i}_n}$ means that φ is evaluated by the assignment corresponding to the bitlist \mathbf{i}_n .

If for at least one assignment φ evaluates to 1 then the system can be proved terminating. Assume that the j -th assignment satisfies φ . Then the precedence

$$\begin{aligned} \text{sat}_{(j+1),0} &> \text{sat}_{(j+2),0} > \dots > \text{sat}_{(2^n-1),0} > \text{sat}_{0,0} > \text{sat}_{1,0} > \dots > \text{sat}_{j,0} \\ \text{sat}_{i,0} &> \text{sat}_{(i+1 \bmod 2^n),1} \quad (0 \leq i < 2^n) \\ \text{sat}_{j,0} &> \text{next}, \star, - > \perp, \top \end{aligned}$$

is well-founded and allows LPO to orient all rules of the labeled TRS $\mathcal{S}_{\text{lab}}^\varphi$ from left to right. Termination of \mathcal{S}^φ follows from Theorem 2.1. \square

As an example consider the transformation of the formula $p_1 \wedge \neg p_2$ below.

Example 4.3 The system $\mathcal{S}^{p_1 \wedge \neg p_2}$ gives rise to the labeled rules

$$\begin{aligned} \text{sat}_{0,0}([p_1; p_2], \perp) &\rightarrow \text{sat}_{1,0}(\text{next}([p_1; p_2]), p_1 \star (-p_2)) \\ \text{sat}_{1,0}([p_1; p_2], \perp) &\rightarrow \text{sat}_{2,1}(\text{next}([p_1; p_2]), p_1 \star (-p_2)) \\ \text{sat}_{2,0}([p_1; p_2], \perp) &\rightarrow \text{sat}_{3,0}(\text{next}([p_1; p_2]), p_1 \star (-p_2)) \\ \text{sat}_{3,0}([p_1; p_2], \perp) &\rightarrow \text{sat}_{0,0}(\text{next}([p_1; p_2]), p_1 \star (-p_2)) \end{aligned}$$

Note that because in the second line the term $(p_1 \star (-p_2))_1$ is interpreted as 1 and hence the system can easily be proved terminating by LPO with the precedence

$$\text{sat}_{2,0} > \text{sat}_{3,0} > \text{sat}_{0,0} > \text{sat}_{1,0} > \text{sat}_{2,1}, \text{next}, \star, - \quad - > \perp, \top$$

In this translation the TRS $\mathcal{S}_{\text{lab}}^\varphi$ gets exponentially larger (in the number of variables in φ) than the original unlabeled system. More precisely, rule (3) gives rise to 2^n different labeled variants due to the n Boolean variables in the list $[p_1; \dots; p_n]$. But the resulting TRS is still finite, in contrast to the one from the next subsection.

4.1 An Alternative Transformation

In the transformation \mathcal{S}^φ the formula φ gets assigned the values implicitly by pattern matching because the same variables p_1, \dots, p_n are used in the formula and in the assignment. One not so nice side-effect is that in rule (3) the list of variables occurring in φ must be specified as the first argument to sat . Here we present a different translation where the variables p_1, \dots, p_n are considered as constants v_1, \dots, v_n in the signature \mathcal{F} . For terms that represent formulas on the syntactic level, the sort `formula` is used, i.e., $v_i : \text{formula}$ for $1 \leq i \leq n$. Furthermore a close inspection of the systems \mathcal{U}^φ and \mathcal{S}^φ reveals that there is no clear separation between syntax and semantics when formulas are represented as terms. To differentiate these two concepts we employ different function symbols for the two layers. Once more the signature \mathcal{F} is augmented by a binary function symbol `and` : `formula` \times `formula` \rightarrow `formula` and a unary function symbol `not` : `formula` \rightarrow `formula`. Consequently also the encoding $\lceil * \rceil$ must now map formulas to their syntactic representation on the term level. Hence the function $\lceil * \rceil$ is redefined accordingly, i.e., $\lceil * \rceil : \mathcal{P}(\mathcal{A}_n) \rightarrow \mathcal{T}(\{v_1, \dots, v_n, \text{and}, \text{not}\})$ with $\lceil p_i \rceil = v_i$ for $1 \leq i \leq n$, $\lceil \varphi \wedge \psi \rceil = \text{and}(\lceil \varphi \rceil, \lceil \psi \rceil)$, and $\lceil \neg \varphi \rceil = \text{not}(\lceil \varphi \rceil)$. Thus, for the formula $p_1 \wedge \neg p_2$ the (syntactic) term representation $\lceil \varphi \rceil$ is `and(v1, not(v2))`.

In the TRS \mathcal{S}^φ the assignment was applied automatically by pattern matching of the variables. Now we employ separate rewrite rules that perform that step. Note that these rules at the same time execute the transformation from the syntactic to

the semantic level. The TRS $\mathcal{A}\text{ssign}$

$$\begin{aligned}
 \text{assign}(xs, \text{and}(x, y)) &\rightarrow \text{assign}(xs, x) \star \text{assign}(xs, y) \\
 \text{assign}(xs, \text{not}(x)) &\rightarrow - \text{assign}(xs, x) \\
 \text{assign}(xs, v_i) &\rightarrow \text{nth}(xs, s^i(0)) && 1 \leq i \leq n \\
 \text{nth}(\perp :: xs, 0) &\rightarrow \perp \\
 \text{nth}(\top :: xs, 0) &\rightarrow \top \\
 \text{nth}(b :: xs, s(j)) &\rightarrow \text{nth}(xs, j)
 \end{aligned}$$

performs the task of $[\alpha]$ on the term representation of propositional formulas. The way how assignments were generated in the previous subsection is no longer suitable. There all variables occurring in φ had to be specified in sat 's first argument. Since we want to get rid of that requirement the idea is to start with an empty assignment (empty list) and increase its length repeatedly. Hence in this section the assignments are no longer computed modulo some length but the *overflow* is simply taken into account by increasing the length of the list. The three rules below are referred to as the TRS $\mathcal{N}\text{ext2}$:

$$\begin{aligned}
 \text{next}(\text{nil}) &\rightarrow \top :: \text{nil} \\
 \text{next}(\perp :: xs) &\rightarrow \top :: xs \\
 \text{next}(\top :: xs) &\rightarrow \perp :: \text{next}(xs)
 \end{aligned}$$

Similar to before a more readable notation for bitlists is employed, i.e., they are identified with natural numbers as follows: $\text{enc}: \mathcal{T}(\mathcal{G}) \rightarrow \mathbb{N}$ with $\text{enc}(\text{nil}) = \text{enc}(\perp) = 0$, $\text{enc}(\top) = 1$, and $\text{enc}(x :: xs) = \text{enc}(x) + 2 \text{enc}(xs)$. This encoding is not injective because the lists $[\top; \perp; \perp]$ and $[\top]$ are both denoted by $\mathbf{1}$. In our setting these (more or less) leading zeros do not pose a problem.

Lemma 4.4 *For a bitlist t , $\text{next}(t)$ rewrites to the successor of t :*

$$\text{If } \text{enc}(t) = \mathbf{i} \text{ then } \text{next}(t) \rightarrow_{\mathcal{N}\text{ext2}}^* t' \text{ with } \text{enc}(t') = \mathbf{i} + \mathbf{1}.$$

Proof By induction on the structure of t and unfolding the definition of \mathbf{i} . \square

The desired property that the rules in $\mathcal{A}\text{ssign}$ evaluate the term representation $\lceil \varphi \rceil$ for a given bitlist \mathbf{i} is formalized in the lemma below.

Lemma 4.5 *Let $\varphi \in \mathcal{P}(\mathcal{A}_n)$ and let \mathbf{i} be the encoding of an assignment α with $[\alpha](\varphi) = 0$. Then $\text{assign}(\mathbf{i}, \lceil \varphi \rceil) \rightarrow_{\mathcal{A}\text{ssign} \cup \text{Simp}}^* \perp$.*

Proof By induction on the structure of $\lceil \varphi \rceil$ and unfolding the definition of \mathbf{i} . \square

Now, we will establish a theorem similar to Theorem 4.2. Again, we prove the theorem in a many-sorted setting. The full information is depicted in Table 3. The variables in the TRS are associated to sorts as follows: $b \in \mathcal{V}_{\text{bool}}$, $xs \in \mathcal{V}_{\text{list}}$, $j \in \mathcal{V}_{\text{nat}}$, and $x, y \in \mathcal{V}_{\text{formula}}$.

$\perp : \text{bool}$	$\top_{\mathcal{A}} = 1$
$\top : \text{bool}$	$\perp_{\mathcal{A}} = 0$
$\star : \text{bool} \times \text{bool} \rightarrow \text{bool}$	$\star_{\mathcal{A}}(x, y) = x \cdot y$
$- : \text{bool} \rightarrow \text{bool}$	$-_{\mathcal{A}}(x) = \bar{x}$
$\text{nil} : \text{list}$	$\text{nil}_{\mathcal{A}} = 0$
$:: : \text{bool} \times \text{list} \rightarrow \text{list}$	$::_{\mathcal{A}}(x, i) = x + 2i$
$\text{next} : \text{list} \rightarrow \text{list}$	$\text{next}_{\mathcal{A}}(i) = i + 1$
$v_i : \text{formula} \quad 1 \leq i \leq n$	$v_{i,\mathcal{A}} = p_i$
$\text{and} : \text{formula} \times \text{formula} \rightarrow \text{formula}$	$\text{and}_{\mathcal{A}}(x, y) = x \wedge y$
$\text{not} : \text{formula} \rightarrow \text{formula}$	$\text{not}_{\mathcal{A}}(x) = \neg x$
$0 : \text{nat}$	$0_{\mathcal{A}} = 0$
$s : \text{nat} \rightarrow \text{nat}$	$s_{\mathcal{A}}(x) = x + 1$
$\text{assign} : \text{list} \times \text{formula} \rightarrow \text{bool}$	$\text{assign}_{\mathcal{A}}(i, \varphi) = [\alpha_i](\varphi)$
$\text{nth} : \text{list} \times \text{nat} \rightarrow \text{bool}$	$\text{nth}_{\mathcal{A}}(i, j) = \alpha_i(p_j)$
$\text{sat} : \text{list} \times \text{bool} \rightarrow \text{bool}$	$\text{sat}_{\mathcal{A}}(i, b) = 0$

Table 3
A model for the $\{\text{bool}, \text{formula}, \text{list}, \text{nat}\}$ -sorted TRS \mathcal{T}^φ .

Theorem 4.6 *Let $\varphi \in \mathcal{P}(\mathcal{A}_n)$. Then the parametrized $\{\text{bool}, \text{formula}, \text{list}, \text{nat}\}$ -sorted TRS \mathcal{T}^φ consisting of the *Simp*-, *Next2*-, and *Assign*-rules plus additionally*

$$\text{sat}(xs, \perp) \rightarrow \text{sat}(\text{next}(xs), \text{assign}(xs, \ulcorner \varphi \urcorner)) \quad (8)$$

is terminating if and only if φ is satisfiable.

Proof Concerning the direction from left to right one can again construct a non-terminating reduction for any unsatisfiable formula φ . In order not to get stuck while evaluating $\text{assign}(\mathbf{i}, \ulcorner \varphi \urcorner)$ a sufficiently large \mathbf{i} is taken (e.g., $\mathbf{i} = 2^{n+1}$). Then there is the infinite sequence

$$\begin{aligned} \text{sat}(\mathbf{i}, \perp) &\rightarrow \text{sat}(\text{next}(\mathbf{i}), \text{assign}(\mathbf{i}, \ulcorner \varphi \urcorner)) \rightarrow^* \\ &\text{sat}(\mathbf{i} + \mathbf{1}, \perp) \rightarrow \text{sat}(\text{next}(\mathbf{i} + \mathbf{1}), \text{assign}(\mathbf{i} + \mathbf{1}, \ulcorner \varphi \urcorner)) \rightarrow^* \\ &\text{sat}(\mathbf{i} + \mathbf{2}, \perp) \rightarrow \text{sat}(\text{next}(\mathbf{i} + \mathbf{2}), \text{assign}(\mathbf{i} + \mathbf{2}, \ulcorner \varphi \urcorner)) \rightarrow^* \dots \end{aligned}$$

where the \rightarrow -steps are applications of rule 8 and the \rightarrow^* -steps can be performed because of Lemmata 4.4 and 4.5.

For the direction from right to left once more a semantic labeling approach is followed. The interpretation given in Table 3 models the $\{\text{bool}, \text{formula}, \text{list}, \text{nat}\}$ -sorted TRS \mathcal{T}^φ . What remains to be defined is an enumeration α_i of assignments as follows: $\alpha_i(p_j) = f^j(i) \bmod 2$ with $f^0(i) = i$ and $f^{j+1}(i) = f^j(\lceil i \div 2 \rceil)$. Checking that \mathcal{A} models \mathcal{T}^φ is straightforward.

Again, only the function symbol sat is labeled. Note that the labeled TRS T_{lab}^φ is infinite since all possible instances of bitlists are considered (compared to finitely

many bitlists of a specified length in the previous subsection). The labeling function $\ell_{\text{sat}}(i, b) = (i, b)$ gives rise to infinitely many rules of the following shape

$$\text{sat}_{i,0}(xs, \perp) \rightarrow \text{sat}_{(i+1),[\alpha_i](\varphi)}(\text{next}(xs), \text{assign}(xs, \ulcorner \varphi \urcorner))$$

Similar to before a precedence of the shape $\text{sat}_{i,0} > \text{sat}_{(i+1),0}$ if $[\alpha_i](\varphi) = 0$ and $\text{sat}_{i,0} > \text{sat}_{(i+1),1}$ if $[\alpha_i](\varphi) = 1$ for all $i \geq 0$ is needed which in general might not be well-founded since it can contain the infinite sequence

$$\text{sat}_{0,0} > \text{sat}_{1,0} > \text{sat}_{2,0} > \dots$$

but due to the assumption that φ is satisfiable, not all of these precedence comparisons are necessary. If $[\alpha_j](\varphi) = 1$ then there is no labeled rule which demands $\text{sat}_{j,0} > \text{sat}_{(j+1),0}$. Without loss of generality we can assume $0 \leq j < 2^n$. Due to the construction of α_j also $\alpha_{j+2^n}, \alpha_{j+2^{n+1}}, \dots$ satisfy φ and hence removing all superfluous comparisons $\text{sat}_{(j+2^{n+m}),0} > \text{sat}_{(j+2^{n+m+1}),0}$ for all $m \in \mathbb{N}$ produces a well-founded precedence (because for any $i \in \mathbb{N}$ one can find a $k \in \mathbb{N}$ such that $i \leq j + 2^{n+k}$). It follows that $\mathcal{T}_{\text{lab}}^\varphi$ is terminating. The termination of \mathcal{T}^φ follows from Theorem 2.1. \square

Although the transformations \mathcal{S}^φ and \mathcal{T}^φ look very similar at first, they are quite different. Concerning the number of rewrite rules, \mathcal{S}^φ does not depend on φ whereas \mathcal{T}^φ depends linearly on the number of variables in φ . On the other hand, the list of variables p_1, \dots, p_n must be given as an argument to sat in \mathcal{S}^φ . In Section 5 it becomes apparent that proving (non)termination automatically is much more challenging for \mathcal{T}^φ than for \mathcal{S}^φ . The main reason is that by separating syntax from semantics, there is less structure that can be exploited by termination tools. The nontermination proofs become more challenging because for \mathcal{S}^φ an infinite rewrite sequence can be captured by considering cyclic reductions of ground terms, i.e., $t \rightarrow^+ t$ for a ground term t (cf. the proof of Theorem 4.2). In contrast \mathcal{T}^φ really demands looping reductions, i.e., $t \rightarrow^+ C[t\sigma]$ where the context C is empty but t may no longer be ground since the lengths of the bitlists are increased.

5 Evaluation

For experimental results⁹ we considered all automated (non)termination analyzers that participated in the 2007 edition of the international termination competition for term rewrite systems augmented with TPA [15], a tool with strong support for termination proofs via semantic labeling. To our knowledge none of these tools supports analysis of sorted TRSs. Consequently we provide our examples unsorted. As already stated in the beginning, dropping sorts does not affect termination of the TRSs we propose. Furthermore we stress that the proofs of Theorems 4.2 and 4.6 can be modified to work on unsorted TRSs. For the TRS \mathcal{S}^φ this means that the interpretations range over the set of pairs P whereas the proof of Theorem 4.6 can be generalized to one sort by using the natural numbers as a carrier and representing formulas via a Gödel encoding [11].

⁹ Further details to be found at <http://colo6-c703.uibk.ac.at/ttt2/hz/sat2trs/>.

tool	2 variables, depth 3			3 variables, depth 4			4 variables, depth 5		
	\mathcal{S}^φ	\mathcal{T}^φ	\mathcal{U}^φ	\mathcal{S}^φ	\mathcal{T}^φ	\mathcal{U}^φ	\mathcal{S}^φ	\mathcal{T}^φ	\mathcal{U}^φ
	T/N	T/N	T/N	T/N	T/N	T/N	T/N	T/N	T/N
AProVE	81/19	0/0	19/81	34/0	0/0	10/88	14/0	0/0	5/79
Jambox	16/0	0/0	19/0	24/0	0/0	12/0	15/0	0/0	11/0
NTI	0/19	0/0	0/81	0/5	0/0	0/74	0/0	0/0	0/11
TPA	0/0	0/0	1/0	0/0	0/0	0/0	0/0	0/0	0/0
$\mathsf{T}\mathsf{T}\mathsf{T}_2$	10/0	0/0	0/0	6/0	0/0	0/0	5/0	0/0	0/0

Table 4
Experimental Results.

It turned out that even for rather small formulas (some of) our transformations produce rewrite systems whose termination analysis is challenging. We considered 100 randomly generated formulas of different shapes. Table 4 summarizes the results, e.g., formulas of depth three using two different propositional variables are considered in the leftmost block, etc. Every tool was run on all TRSs resulting from transforming the formula φ to \mathcal{S}^φ , \mathcal{T}^φ , and \mathcal{U}^φ for at most 60 seconds to analyze termination (T) or nontermination (N) of each system. Globally speaking, for TRSs originating from very small formulas AProVE [9] performs best. This is due to its support for narrowing which allows to exploit the structure of \mathcal{S}^φ and \mathcal{U}^φ . Jambox solves some instances by semantic labeling over Boolean models (which is very close to the way how we proved termination) and by the matrix method. The latter systems could also be handled by $\mathsf{T}\mathsf{T}\mathsf{T}_2$.¹⁰ NTI [20] supports only nontermination analysis, using an unfolding operator. Semantic labeling based on Boolean models and (quasi-)models over the naturals is implemented in TPA [17,16] which usually performs very well on standard examples. The experiments reveal that the latter is not powerful for the systems obtained from the transformations proposed in this paper.

But narrowing is expensive which can be seen by comparing the different blocks of Table 4. AProVE can handle all TRSs resulting from the \mathcal{S}^φ translation if formulas are of depth three but for depth four (five) the performance decreases to 34% (14%). For the other translations the effect is not so tremendous, well, for \mathcal{T}^φ the surprising outcome is that no tool could handle any system at all and the systems in \mathcal{U}^φ are generally a bit easier since they do not iterate over the assignments. Needless to say, the formulas φ which are considered for our experiments are a very trivial task for any SAT-solver.

We conclude this section by a sketch of how AProVE solves many instances by considering the TRS \mathcal{S}^φ for $\varphi = p_1 \wedge p_2$. After some preliminary analysis based on dependency pairs [2], AProVE concludes that any infinite sequence applies the rule

$$\text{sat}([p_1; p_2], \perp) \rightarrow \text{sat}(\text{next}([p_1; p_2]), p_1 \star p_2)$$

indefinitely. Narrowing the above rule at position 1 allows to replace it by the two

¹⁰<http://colo6-c703.uibk.ac.at/ttt2/>

rules

$$\begin{aligned} \text{sat}([\perp; p_2], \perp) &\rightarrow \text{sat}(\text{next}([\perp; p_2]), \perp \star p_2) \\ \text{sat}([\top; p_2], \perp) &\rightarrow \text{sat}(\text{next}([\top; p_2]), \top \star p_2) \end{aligned}$$

and narrowing these rules at position 1 gives

$$\begin{aligned} \text{sat}([\perp; \perp], \perp) &\rightarrow \text{sat}(\text{next}([\perp; \perp]), \perp \star \perp) \\ \text{sat}([\perp; \top], \perp) &\rightarrow \text{sat}(\text{next}([\perp; \top]), \perp \star \top) \\ \text{sat}([\top; \perp], \perp) &\rightarrow \text{sat}(\text{next}([\top; \perp]), \top \star \perp) \\ \text{sat}([\top; \top], \perp) &\rightarrow \text{sat}(\text{next}([\top; \top]), \top \star \top) \end{aligned}$$

After this state is reached the right-hand sides can be rewritten [10] using the Simp and Next rules which allows the dependency graph processor [2] to conclude termination.

6 Conclusion

In this paper we proposed three different transformations from propositional formulas φ to confluent—since orthogonal—term rewrite systems \mathcal{S}^φ , \mathcal{T}^φ , and \mathcal{U}^φ such that φ is satisfiable (unsatisfiable) if and only if \mathcal{S}^φ , \mathcal{T}^φ (\mathcal{U}^φ) is terminating. Although the systems can be proved (non)terminating by semantic labeling using intuitive models, state-of-the-art termination tools fail even on very small and simple TRSs. Especially the transformation \mathcal{T}^φ produces unsolvable rewrite systems which might be due to the fact that it preserves much less structure than \mathcal{S}^φ does. If tool authors investigate the reasons why the generated problems are that hard, new termination techniques could emerge.

References

- [1] Aoto, T. and T. Yamada, *Termination of simply typed term rewriting by translation and labelling*, in: *Proc. 14th International Conference on Rewriting Techniques and Applications*, LNCS **2706**, 2003, pp. 380–394.
- [2] Arts, T. and J. Giesl, *Termination of term rewriting using dependency pairs*, *Theoretical Computer Science* **236** (2000), pp. 133–178.
- [3] Baader, F. and T. Nipkow, “Term Rewriting and All That,” Cambridge University Press, 1998.
- [4] Codish, M., V. Lagoon and P. Stuckey, *Solving partial order constraints for LPO termination*, in: *Proc. 17th International Conference on Rewriting Techniques and Applications*, LNCS **4098**, 2006, pp. 4–18.
- [5] Cook, S., *The complexity of theorem-proving procedures*, in: *Proc. 3rd annual ACM symposium on theory of computing* (1971), pp. 151–158.
- [6] Dershowitz, N., *Orderings for term-rewriting systems*, *Theoretical Computer Science* **17** (1982), pp. 279–301.
- [7] Dick, J., J. Kalmus and U. Martin, *Automating the Knuth-Bendix ordering*, *Acta Informatica* **28** (1990), pp. 95–119.
- [8] Endrullis, J., J. Waldmann and H. Zantema, *Matrix interpretations for proving termination of term rewriting*, *Journal of Automated Reasoning* **40** (2008), pp. 195–220.
- [9] Giesl, J., P. Schneider-Kamp and R. Thiemann, *AProVE 1.2: Automatic termination proofs in the dependency pair framework*, in: *Proc. 3rd International Joint Conference on Automated Reasoning*, LNAI **4130**, 2006, pp. 281–286.

- [10] Giesl, J., R. Thiemann and P. Schneider-Kamp, *The dependency pair framework: Combining techniques for automated termination proofs*, in: *Proc. 11th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning*, LNAI **3452**, 2005, pp. 301–331.
- [11] Gödel, K., *Über formal unentscheidbare sätze der principia mathematica und verwandter systeme*, Monatshefte für Mathematik und Physik **38** (1931), pp. 173–198.
- [12] Huet, G. and D. Lankford, *On the uniform halting problem for term rewriting systems*, Technical Report 282, INRIA, Le Chesnay, France (1978).
- [13] Kamin, S. and J. Lévy, *Two generalizations of the recursive path ordering*, Unpublished manuscript, University of Illinois (1980).
- [14] Knuth, D. and P. Bendix, *Simple word problems in universal algebras*, in: J. Leech, editor, *Computational Problems in Abstract Algebra*, Pergamon Press, 1970 pp. 263–297.
- [15] Koprowski, A., *TPA: Termination proved automatically.*, in: *Proc. 17th International Conference on Rewriting Techniques and Applications*, LNCS **4098**, 2006, pp. 257–266.
- [16] Koprowski, A. and A. Middeldorp, *Predictive labeling with dependency pairs using SAT*, in: *Proc. 21st International Conference on Automated Deduction*, LNAI **4603**, 2007, pp. 410–425.
- [17] Koprowski, A. and H. Zantema, *Automation of recursive path ordering for infinite labelled rewrite systems*, in: *Proc. 3rd International Joint Conference on Automated Reasoning*, LNAI **4130**, 2006, pp. 332–346.
- [18] Korovin, K. and A. Voronkov., *Orienting rewrite rules with the Knuth-Bendix order*, Information and Computation **183** (2003), pp. 165–186.
- [19] Kurihara, M. and H. Kondo, *Efficient BDD encodings for partial order constraints with application to expert systems in software verification*, in: *Proc. 17th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, LNAI **3029**, 2004, pp. 827–837.
- [20] Payet, É., *Detecting non-termination of term rewriting systems using an unfolding operator*, in: *Proc. 11th International Symposium on Logic-Based Program Synthesis and Transformation*, LNCS **4407**, 2007, pp. 194–209.
- [21] Waldmann, J., *Matchbox: A tool for match-bounded string rewriting*, in: *Proc. 15th International Conference on Rewriting Techniques and Applications*, LNCS **3091**, 2004, pp. 85–94.
- [22] Zankl, H. and A. Middeldorp, *Satisfying KBO constraints*, in: *Proc. 18th International Conference on Rewriting Techniques and Applications*, LNCS **4533**, 2007, pp. 389–403.
- [23] Zantema, H., *Termination of term rewriting: Interpretation and type elimination*, Journal of Symbolic Computation **17** (1994), pp. 23–50.
- [24] Zantema, H., *Termination of term rewriting by semantic labelling*, Fundamenta Informaticae **24** (1995), pp. 89–105.