

Rocket-Fast Proof Checking for SMT Solvers

Michał Moskal

University of Wrocław, Poland

Abstract. Modern Satisfiability Modulo Theories (SMT) solvers are used in a wide variety of software and hardware verification applications. Proof producing SMT solvers are very desirable as they increase confidence in the solver and ease debugging/profiling, while allowing for scenarios like Proof-Carrying Code (PCC). However, the size of typical proofs generated by SMT solvers poses a problem for the existing systems, up to the point where proof checking consumes orders of magnitude more computer resources than proof generation. In this paper we show how this problem can be addressed using a simple term rewriting formalism, which is used to encode proofs in a natural deduction style. We formally prove soundness of our rules and evaluate an implementation of the term rewriting engine on a set of proofs generated from industrial benchmarks. The modest memory and CPU time requirements of the implementation allow for proof checking even on a small PDA device, paving a way for PCC on such devices.

1 Introduction

Satisfiability Modulo Theories (SMT) [14] solvers check satisfiability of a first order formula, where certain function and constant symbols are interpreted according to a set of background theories. These theories typically include integer or rational arithmetic, bit vectors and arrays. Some SMT solvers support only quantifier free fragments of their logics, other also support quantifiers, most often through instantiation techniques. SMT solvers are often based on search strategies of SAT solvers.

The usage of background theories, instantiation techniques and efficient handling of the Boolean structure of the formula differentiates SMT solvers from first-order theorem provers based on resolution. SMT solvers are efficient for larger mostly ground formulas. This makes them good tools for hardware and software verification.

SMT solvers typically either answer that the input formula is unsatisfiable, or give some description of a model, in which the formula might be satisfiable. In terms of software verification the first answer means that the program is correct, while the second answer means, that an assertion might be violated. The model description is used to identify a specific assertion and/or execution trace.

What is troubling is that we are trusting the SMT solver, when it says the program is correct. One problem is that there might be a bug in the SMT solver, whose implementation can be largely opaque to others than the developer.

The other problem is that we might want to provide the evidence of program being correct to someone else, like in Proof-Caring Code [13] scenarios.

It is therefore desirable for an SMT solver to produce the proof of the unsatisfiability of formulas. The problem is that in program verification, the queries are rather huge and so are the proofs. For example formulas in the AUFLIA division of the SMT problem library¹ contain up to 130 000 distinct subterms, with an average of 8 000. The proofs we have generated are on average five times bigger than the formulas. The most complicated proof we have encountered contains around 40 000 basic resolution steps and around 1 000 000 (sub)terms in size. What is worth noting however, is that state of the art SMT solvers are able to check a vast majority of such queries in under a second. As the general expectation is that proof checking should be faster than proof generation, it becomes clear that we need very efficient means of proof checking.

1.1 Contributions

The contributions of this paper are:

- we introduce a simple, yet expressive term rewrite formalism (Sect. 2), and show it is strong enough to encode and check proofs of theory and Boolean tautologies (Sect. 3), and also NNF/CNF conversions with skolemization (Sect. 4),
- we discuss two highly efficient implementations of the proposed rewrite system (Sect. 6). In particular we discuss performance issues (Sect. 6.2) and we describe techniques to help ensure soundness of the rewrite rules (Sect. 6.1).

There are two reasons to use term rewriting as a proof checking vehicle. One is that the term rewriting is a simple formalism, therefore it is relatively easy to reason about the correctness of an implementation of the proof checker. The bulk of soundness reasoning goes at term rewrite rules level, which is much better understood and simpler to reason about than a general purpose (often low level) programming language used to implement a proof checker.

The second reason is memory efficiency, which on modern CPUs is also a key to time efficiency. We encode proof rules as rewrite rules and handle non-local conditions (like uniqueness of Skolem functions) at the meta level, which allows for the rewrite rules to be local. The idea behind the encoding of the proof rules is to take proof terms and rewrite them into the formulas that they prove. This allows for memory held by the proof terms to be immediately reclaimed and reused for the next fragment of the proof tree read from a proof file.

1.2 Proof Search in SMT Solvers

This section gives description of a proof search of an SMT solver based on DPLL and E-matching. It applies to most of the current SMT solvers.

¹ Both the library and this particular division are described further in Sect. 6.2.

In order to check unsatisfiability of a formula, an SMT solver will usually first transform it into an equisatisfiable CNF formula, while simultaneously performing skolemization. Subsequent proof search alternates between Boolean reasoning, theory reasoning, and quantifier instantiation. For the Boolean part we use resolution. Also the final empty clause is derived using resolution. Theory reasoning produces conflict clauses, which are tautologies under respective theories, e.g., $\neg(a > 7) \vee a \geq 6$ or $\neg(c = d) \vee \neg(f(c) = 42) \vee \neg(f(d) < 0)$. Quantifier reasoning is based on instantiating universal quantifiers and producing tautologies like $\neg(\forall x. f(x) > 0 \rightarrow P(x)) \vee \neg(f(3) > 0) \vee P(3)$. It can be thought of as just another background theory.

To make this search procedure return a proof, we need proofs of: CNF translation, Boolean tautologies and theory tautologies. By taking these three together, we should obtain a proof that the formula is unsatisfiable.

2 Definitions

Let \mathcal{V} be an infinite, enumerable, set of variables. We use x and y (all symbols possibly with indices) as meta-variables ranging over \mathcal{V} . Let Σ be an infinite, enumerable set of function symbols, we use meta-variable f ranging over Σ . We define the set of *terms* \mathcal{T} , and the set of *patterns* $\mathcal{P} \subseteq \mathcal{T}$ as follows:

$$\begin{aligned} \mathcal{T} ::= & x \mid f(\mathcal{T}_1, \dots, \mathcal{T}_n) \mid \lambda x. \mathcal{T}_1 \mid \mathbf{cons} \cdot (\mathcal{T}_1, \mathcal{T}_2) \mid \mathbf{nil} \cdot () \mid \mathbf{build} \cdot (f, \mathcal{T}_1) \mid \\ & \mathbf{apply} \cdot (\mathcal{T}_1, \mathcal{T}_2) \mid \mathbf{fold} \cdot (\mathcal{T}_1) \\ \mathcal{P} ::= & x \mid f(\mathcal{P}_1, \dots, \mathcal{P}_n) \end{aligned}$$

where $n \geq 0$. The notion $s \cdot (\dots)$ stands for a *special form*, which have particular interpretations in the term rewrite system. We will use $t_1 :: t_2$ as a syntactic sugar for $\mathbf{cons} \cdot (t_1, t_2)$, and \mathbf{nil} for $\mathbf{nil} \cdot ()$.

The set of *free variables* of a term, $FV : \mathcal{T} \rightarrow \mathcal{P}(\mathcal{V})$, is defined as usual:

$$\begin{aligned} FV(x) &= \{x\} \\ FV(f(t_1, \dots, t_n)) &= \bigcup_{1 \leq i \leq n} FV(t_i) \\ FV(\lambda x. t) &= FV(t) \setminus \{x\} \\ FV(s \cdot (t_1, \dots, t_n)) &= \bigcup_{1 \leq i \leq n} FV(t_i) \end{aligned}$$

Note that it is also defined on \mathcal{P} , as $\mathcal{P} \subseteq \mathcal{T}$. Let $\mathcal{T}(A, B)$ be a set of terms built from function symbols from the set A and variables from the set $B \subseteq \mathcal{V}$ (i.e. if $t \in \mathcal{T}(A, B)$ then $FV(t) \subseteq B$). A *substitution* is a function $\sigma : \mathcal{V} \rightarrow \mathcal{T}$, which we identify with its homomorphic, capture free extension to $\sigma : \mathcal{T} \rightarrow \mathcal{T}$.

A *rewrite rule* is a pair (p, t) , where $p \in \mathcal{P}$, $t \in \mathcal{T}$ and $FV(t) \subseteq FV(p)$. Let \mathcal{R} be set of such rewrite rules, such that for distinct $(p, t), (p', t') \in \mathcal{R}$, p and p' do not unify. We define a *normal form* of a term t , with respect to \mathcal{R} as $\mathbf{nf}(t)$, with the rules below. Because the function defined below is recursive it is possible for it not to terminate. If the rules below do not result in a single unique normal form for term t , then we say that $\mathbf{nf}(t) = \otimes$. If term has \otimes as subterm, it is itself regarded as equal to \otimes . In practice this condition is enforced by limiting running time of the proof checker.

$$\begin{aligned}
 \mathbf{nf}(x) &= x \\
 \mathbf{nf}(f(t_1, \dots, t_n)) &= \begin{cases} \mathbf{nf}(t\sigma) & \text{for } (p, t) \in \mathcal{R} \text{ such that} \\ & \exists \sigma. p\sigma = f(\mathbf{nf}(t_1), \dots, \mathbf{nf}(t_n)) \\ f(\mathbf{nf}(t_1), \dots, \mathbf{nf}(t_n)) & \text{otherwise} \end{cases} \\
 \mathbf{nf}(\lambda x. t_1) &= \lambda x. \mathbf{nf}(t_1) \\
 \mathbf{nf}(\mathbf{apply} \cdot (\lambda x. t_1, t_2)) &= \mathbf{nf}(t_1[x := t_2]) \\
 \mathbf{nf}(\mathbf{build} \cdot (f, t_1 :: \dots :: t_n :: \mathbf{nil})) &= \mathbf{nf}(f(t_1, \dots, t_n)) \\
 \mathbf{nf}(\mathbf{build} \cdot (f, t_1)) &= \mathbf{build} \cdot (f, \mathbf{nf}(t_1)) \quad \text{if none of the above apply} \\
 \mathbf{nf}(s \cdot (t_1, \dots, t_n)) &= s \cdot (\mathbf{nf}(t_1), \dots, \mathbf{nf}(t_n)) \quad \text{if none of the above apply}
 \end{aligned}$$

where $t_1[x := t_2]$ denotes a capture free substitution of x with t_2 in t_1 ².

The semantics of the $\mathbf{fold} \cdot (t)$ is not defined above. Its role is to perform theory-specific constant folding on t . Folding is implemented either inside the proof checker or by an external tool called by the proof checker. In this paper we use integer constant folding (for example $\mathbf{nf}(\mathbf{fold} \cdot (\mathbf{add}(20, 22))) = 42$).

The signature used throughout this paper can be divided in four categories:

1. logical connectives: `false`, `implies`, `and`, `or`, `forall`, `neg`
2. theory specific symbols: `eq`, `add`, `leq`, `minus` and natural number literals (`0`, `1`, `2`, ...)
3. technical machinery: `lift_known`, `□`, `sk`
4. rule names

3 Boolean Deduction

Consider the logical system from Fig. 1. It is complete for Boolean logic with connectives \rightarrow and \perp . Three of the derivation rules there (`(mp)`, `(absurd)` and `(npp)`) fit a common scheme:

$$\frac{\Gamma \vdash \Xi_1(\psi_1, \dots, \psi_m) \dots \Gamma \vdash \Xi_n(\psi_1, \dots, \psi_m)}{\Gamma \vdash \Xi(\psi_1, \dots, \psi_m)} (r)$$

where Ξ_i and Ξ are formulas built from the Boolean connectives and formula meta-variables ψ_1, \dots, ψ_m , while (r) is the name of the rule. We call such rules *standard rules*. Additional Boolean connectives can be handled by adding more standard rules. To encode a standard derivation rule, we use the following rewrite:

$$r(\square(\Xi_1(x_1, \dots, x_m)), \dots, \square(\Xi_n(x_1, \dots, x_m)), x_{i_1}, \dots, x_{i_l}) \blacktriangleright \square(\Xi(x_1, \dots, x_m))$$

² The actual implementation uses de Bruijn indices, so the “capture free” part comes at no cost.

Proof rule	Rewrite rule
$\frac{\Gamma \vdash \psi_1 \rightarrow \psi_2 \quad \Gamma \vdash \psi_1}{\Gamma \vdash \psi_2} \text{ (mp)}$	$\text{mp}(\square(\text{implies}(x_1, x_2)), \square(x_1)) \blacktriangleright \square(x_2)$
$\frac{\Gamma \vdash \perp}{\Gamma \vdash \psi} \text{ (absurd)}$	$\text{absurd}(\square(\text{false}), x) \blacktriangleright \square(x)$
$\frac{\Gamma \vdash (\psi \rightarrow \perp) \rightarrow \perp}{\Gamma \vdash \psi} \text{ (nnpp)}$	$\text{nnpp}(\square(\text{implies}(\text{implies}(x, \text{false}), \text{false}))) \blacktriangleright \square(x)$
$\frac{\Gamma \cup \{\psi_1\} \vdash \psi_2}{\Gamma \vdash \psi_1 \rightarrow \psi_2} \text{ (assume)}$	$\begin{aligned} &\text{assume}(x_1, x_2) \blacktriangleright \\ &\quad \text{lift_known}(\text{implies}(x_1, \text{apply} \cdot (x_2, \square(x_1)))) \\ &\text{lift_known}(\text{implies}(x_1, \square(x_2))) \blacktriangleright \\ &\quad \square(\text{implies}(x_1, x_2)) \end{aligned}$
$\frac{\psi \in \Gamma}{\Gamma \vdash \psi} \text{ (assumption)}$	

Fig. 1. A complete system for \rightarrow and \perp

where x_{i_j} are additional technical arguments, to fulfill the condition that the left-hand side of a rule has to contain all the free variables in the right-hand side and r is a function symbol used to encode this particular rule. Therefore we can model (mp), (absurd) and (nnpp) using the rewrite rules listed in Fig. 1.

We are left with the (assume)/(assumption) pair, which is modeled using lambda expressions. There is no explicit rewrite for (assumption) rule. The term x_2 is expected to be of the form $\lambda y. t$, where t is a proof using y in places where (assumption) should be used. This is very similar to the encoding of the IMP-I rule in Edinburgh Logical Framework [10].

We call *restricted* the terms of the form $\square(\dots)$, $\text{lift_known}(\dots)$ or $\text{sk}(\dots)$ (the last one is used in the next section). We say that $P \in \mathcal{T}$ is a pre-proof (written $\text{preproof}(P)$), if it does not contain a restricted subterm, or a subterm which is a $s \cdot (\dots)$ special form.

Lemma 1. *For any pair (P, σ) , such that $\text{preproof}(P)$, $\forall x \in \mathcal{V}. x\sigma = x \vee \exists \phi. x\sigma = \square(\phi)$ and $\mathbf{nf}(P\sigma) = \square(\psi)$, there exists a derivation $\Gamma \vdash \psi$ where $\Gamma = \{\phi \mid x \in \mathcal{V}, x\sigma = \square(\phi)\}$.*

Proof. The proof is by induction on the size of P . Because $\square(\dots)$ is not a subterm of P , the head of P must be either:

1. a variable x , in which case $x\sigma$ is $\square(\psi)$ and the (assumption) rule can be used, since $\psi \in \Gamma$,
2. $P = r(P_1, \dots, P_n, t_1, \dots, t_m)$, where a rewrite, obtained from a derivation rule (r), is applicable to:

$$r(\mathbf{nf}(P_1\sigma), \dots, \mathbf{nf}(P_n\sigma), \mathbf{nf}(t_1), \dots, \mathbf{nf}(t_m))$$

We use the induction hypothesis on (P_i, σ) , where $\mathbf{nf}(P_i\sigma) = \square(\psi_i)$ and build the derivation using the (r) rule.

Proof rule	Rewrite rule
$\frac{}{\Gamma \vdash t = t}$ (eq_refl)	eq_refl(x) \blacktriangleright $\square(\text{eq}(x, x))$
$\frac{\Gamma \vdash t_1 = t_2 \quad \Gamma \vdash t_2 = t_3}{\Gamma \vdash t_1 = t_3}$ (eq_trans)	eq_trans($\square(\text{eq}(x_1, x_2)), \square(\text{eq}(x_2, x_3))$) \blacktriangleright $\square(\text{eq}(x_1, x_3))$
$\frac{\Gamma \vdash t_1 = t_2}{\Gamma \vdash t_2 = t_1}$ (eq_trans)	eq_symm($\square(\text{eq}(x_1, x_2))$) \blacktriangleright $\square(\text{eq}(x_2, x_1))$
$\frac{\Gamma \vdash t_1 = t_2}{\Gamma \vdash \psi(t_1) \rightarrow \psi(t_2)}$ (eq_sub)	eq_sub($\square(\text{eq}(x_1, x_2)), y$) \blacktriangleright $\square(\text{implies}(\text{apply} \cdot (y, x_1), \text{apply} \cdot (y, x_2)))$
$\frac{\Gamma \vdash x + x_1 \leq c_1 \quad \Gamma \vdash -x + x_2 \leq c_2}{x_1 + x_2 \leq c_1 + c_2}$ (utvpi_trans)	
utvpi_trans($\square(\text{leq}(\text{add}(x_1, x_2), x_3)), \square(\text{leq}(\text{add}(\text{minus}(x_1), y_2), y_3))$) \blacktriangleright $\square(\text{leq}(\text{add}(x_2, y_2), \text{fold} \cdot (\text{add}(x_3, y_3)))$)	

Fig. 2. The equality rules, and an example of an UTVPI rule

3. $P = \text{assume}(P_1, \psi)$, which rewrites to $\square(\dots)$ in two steps, through the `lift_known(...)` (which cannot be used explicitly because $\text{preproof}(P)$). $\text{apply} \cdot (P_1, \square(\psi))$ needs to be reduced to $\square(\dots)$ for the `lift_known(...)` to be applied, so $\text{nf}(P_1) = \lambda x. P_2$, for some P_2 . Because no rule can result in a rewrite to a lambda term (all of the rewrite rules have a term of the form $f(\dots)$ as their right hand side), then not only $\text{nf}(P_1)$, but also P_1 itself needs to start with a lambda binder. Therefore $P_1 = \lambda x. P_3$, for some P_3 . In this case we use the induction hypothesis on $(P_3, \sigma[x := \square(\psi)])$, and then use the (`assume`) rule to construct the implication.

There are no other cases, since no other rewrite rule has $\square(\dots)$ as the right-hand side. \square

Applying this lemma with (P, \emptyset) gives the theorem.

Theorem 1. *For any P , such that $\text{preproof}(P)$ and $\text{nf}(P\sigma) = \square(\psi)$ there exists a derivation $\vdash \psi$.*

Theory Conflicts. Proving theory conflicts clearly depends on the particular theory. Fig. 2 lists rules for the theory of equality. The encoding is the same as for the standard rules from Fig. 1. For arithmetic we currently support the UTVPI fragment [11] of integer linear arithmetic. It consists of inequalities of the form $ax + by \leq c$, where $a, b \in \{-1, 0, 1\}$ and c is an integer. The decision procedure closes set of such inequalities, with respect to a few rules, of the form similar to the one listed in Fig. 2. Again, the encoding is the same as for ordinary deduction rules.

4 Skolemization Calculus

Fig. 3 lists rules for a skolemization calculus. The \uplus is disjoint set union (i.e. $A \cup B$ if $A \cap B = \emptyset$ and undefined otherwise). The intuition behind $S; Q \vdash \psi \rightsquigarrow \phi$

Proof rule	Rewrite rule
$\frac{\emptyset; Q \vdash \neg\psi(f(Q)) \rightsquigarrow \phi}{\{f\}; Q \vdash \neg\forall x. \psi(x) \rightsquigarrow \phi}$ (skol)	$\mathbf{sk}(y, \mathbf{skol}(f, y_1), \mathbf{neg}(\mathbf{forall}(x))) \blacktriangleright \mathbf{sk}(y, y_1, \mathbf{neg}(\mathbf{apply} \cdot (x, \mathbf{build} \cdot (f, y))))$
$\frac{S; Q, x \vdash \psi(x) \rightsquigarrow \phi(x)}{S; Q \vdash \forall x. \psi(x) \rightsquigarrow \forall x. \phi(x)}$ (skip _v)	$\mathbf{sk}(y, \mathbf{skip}_v(y_1), \mathbf{forall}(x_1)) \blacktriangleright \mathbf{forall}(\lambda x. \mathbf{sk}(x :: y, y_1, \mathbf{apply} \cdot (x_1, x)))$
$\frac{}{\emptyset; Q \vdash \overline{\psi} \rightsquigarrow \overline{\psi}}$ (id)	$\mathbf{sk}(y, \mathbf{id}, x_1) \blacktriangleright x_1$
$\frac{S_1; Q \vdash \psi_1 \rightsquigarrow \phi_1 \quad S_2; Q \vdash \psi_2 \rightsquigarrow \phi_2}{S_1 \uplus S_2; Q \vdash \psi_1 \wedge \psi_2 \rightsquigarrow \phi_1 \wedge \phi_2}$ (rec _∧)	$\mathbf{sk}(y, \mathbf{rec}_\wedge(y_1, y_2), \mathbf{and}(x_1, x_2)) \blacktriangleright \mathbf{and}(\mathbf{sk}(y, y_1, x_1), \mathbf{sk}(y, y_2, x_2))$
$\frac{S_1; Q \vdash \psi_1 \rightsquigarrow \phi_1 \quad S_2; Q \vdash \psi_2 \rightsquigarrow \phi_2}{S_1 \uplus S_2; Q \vdash \psi_1 \vee \psi_2 \rightsquigarrow \phi_1 \vee \phi_2}$ (rec _∨)	$\mathbf{sk}(y, \mathbf{rec}_\vee(y_1, y_2), \mathbf{or}(x_1, x_2)) \blacktriangleright \mathbf{or}(\mathbf{sk}(y, y_1, x_1), \mathbf{sk}(y, y_2, x_2))$
$\frac{S_1; Q \vdash \neg\psi_1 \rightsquigarrow \phi_1 \quad S_2; Q \vdash \neg\psi_2 \rightsquigarrow \phi_2}{S_1 \uplus S_2; Q \vdash \neg(\psi_1 \vee \psi_2) \rightsquigarrow \phi_1 \wedge \phi_2}$ (rec _{¬∨})	$\mathbf{sk}(y, \mathbf{rec}_{\neg\vee}(y_1, y_2), \mathbf{neg}(\mathbf{or}(x_1, x_2))) \blacktriangleright \mathbf{and}(\mathbf{sk}(y, y_1, \mathbf{neg}(x_1)), \mathbf{sk}(y, y_2, \mathbf{neg}(x_2)))$
$\frac{S_1; Q \vdash \neg\psi_1 \rightsquigarrow \phi_1 \quad S_2; Q \vdash \neg\psi_2 \rightsquigarrow \phi_2}{S_1 \uplus S_2; Q \vdash \neg(\psi_1 \wedge \psi_2) \rightsquigarrow \phi_1 \vee \phi_2}$ (rec _{¬∧})	$\mathbf{sk}(y, \mathbf{rec}_{\neg\wedge}(y_1, y_2), \mathbf{neg}(\mathbf{and}(x_1, x_2))) \blacktriangleright \mathbf{or}(\mathbf{sk}(y, y_1, \mathbf{neg}(x_1)), \mathbf{sk}(y, y_2, \mathbf{neg}(x_2)))$
$\frac{S_1; Q \vdash \psi_1 \rightsquigarrow \phi_1}{S_1; Q \vdash \neg\neg\psi_1 \rightsquigarrow \phi_1}$ (rec _{¬¬})	$\mathbf{sk}(y, \mathbf{rec}_{\neg\neg}(y_1), \mathbf{neg}(\mathbf{neg}(x_1))) \blacktriangleright \mathbf{sk}(y, y_1, x_1)$

Fig. 3. The skolemization calculus

is that for each model M of $\forall Q. \psi$ there exists an extension of M on the symbols from S that satisfies $\forall Q. \phi$. We formalize it using second order logic with the following lemma:

Lemma 2. *Let $Q = \{x_1, \dots, x_n\}$ and $S = \{f_1 \dots f_n\}$. If $S; Q \vdash \psi \rightsquigarrow \phi$ where $\psi \in \mathcal{T}(\Sigma, Q)$, $\phi \in \mathcal{T}(\Sigma \uplus S, Q)$, then $\models \exists^2 f_1 \dots \exists^2 f_n. \forall x_1, \dots, x_n. \psi \rightarrow \phi$.*

The key point of the proof is that for the rules of the common form:

$$\frac{S_1; Q \vdash \Xi_1(\psi_1, \dots, \psi_k) \rightsquigarrow \phi_1 \quad \dots \quad S_m; Q \vdash \Xi_m(\psi_1, \dots, \psi_k) \rightsquigarrow \phi_m}{S_1 \uplus \dots \uplus S_m; Q \vdash \Xi(\psi_1, \dots, \psi_k) \rightsquigarrow \Xi'(\psi_1, \dots, \psi_k, \phi_1, \dots, \phi_m)} (r)$$

where ψ_j and ϕ_j range over first order formulas it is enough to show the following:

$$\begin{aligned} & \forall Q : \mathbf{Type}. \forall S_1 \dots S_n : \mathbf{Type}. \\ & \forall \psi_1 \dots \psi_k : Q \rightarrow \mathbf{Prop}. \\ & \forall \phi_1 : S_1 \times Q \rightarrow \mathbf{Prop}. \dots \forall \phi_m : S_m \times Q \rightarrow \mathbf{Prop}. \\ & \bigwedge_{i=1 \dots m} (\exists f_i : S_i. \forall x : Q. \Xi_i(\psi_1(x), \dots, \psi_k(x)) \rightarrow \phi_i(f_i, x)) \rightarrow \\ & (\exists f_1 : S_1. \dots \exists f_m : S_m. \forall x : Q. \Xi(\psi_1(x), \dots, \psi_k(x)) \rightarrow \\ & \quad \Xi'(\psi_1(x), \dots, \psi_k(x), \phi_1(f_1, x), \dots, \phi_m(f_m, x))) \end{aligned}$$

which is much like the formula from the lemma, except that there is only one Skolem constant f_i per premise and also there is only one free variable in all the formulas, namely x . However these symbols are of arbitrary type, so they can be thought of as representing sequences of symbols.

We prove such formulas for each rule, the reader can find the proof scripts for Coq proof assistant online [1].

Rewrite encoding. The common form of a rule is encoded as:

$$\mathbf{sk}(y, r(y_1, \dots, y_m), \Xi(x_1, \dots, x_k)) \blacktriangleright \Xi'(x_1, \dots, x_k, \mathbf{sk}(y, y_1, \Xi_1(x_1, \dots, x_k)), \dots, \mathbf{sk}(y, y_m, \Xi_m(x_1, \dots, x_k)))$$

The first argument of $\mathbf{sk}(\dots)$ is the list of universally quantified variables in scope. The second argument is a rule name, along with proofs of premises. The third argument is the formula to be transformed.

The encoding of non-common rules (as well as the common rules used here) is given in the Fig. 3.

Lemma 3. *If $\mathbf{preproof}(P)$, $\mathbf{nf}(\mathbf{sk}(x_1 :: \dots :: x_n :: \mathbf{nil}, P, \psi)) = \psi'$, and for each occurrence of $\mathbf{skol}(f)$ as a subterm of P , the function symbol f does not occur anywhere else in P nor in ψ , then there exists S , such that $S; x_1, \dots, x_n \vdash \rightsquigarrow \psi\psi'$.*

Proof. By structural induction over P . □

Theorem 2. *If $\mathbf{preproof}(P)$, $\mathbf{nf}(\mathbf{sk}(\mathbf{nil}, P, \psi)) = \psi'$, and for each occurrence of $\mathbf{skol}(f)$ as a subterm of P , the function symbol f does not occur anywhere else in P nor in ψ , and ψ' is unsatisfiable then ψ is unsatisfiable.*

Proof. By Lemmas 2 and 3. □

5 The Checker

The proof checker reads three files: (1) rewrite rules describing the underlying logic; (2) a query in SMT-LIB concrete syntax; and (3) the proof. The concrete syntax for both the rewrite rules and the proof term is similar to the one used in SMT-LIB. The proof term language includes the following commands:

- **let** $x := t_1$: bind the identifier x to the term $\mathbf{nf}(t_1)$
- **initial** $t_1 t_2$: check if $\mathbf{skol}(f)$ is used in t_1 only once with each f , that the f symbols do not occur in t_2 , compares t_2 against the query read from the SMT-LIB file and if everything succeeds, binds $\square(\mathbf{nf}(\mathbf{sk}(\mathbf{nil}, t_1, t_2)))$ to the special identifier **initial**; this command can be used only once in a given proof
- **final** t_1 : checks if $\mathbf{nf}(t_1) = \square(\mathbf{false})$, and if so reports success and exits
- **assert.eq** $t_1 t_2$: checks if $\mathbf{nf}(t_1) = \mathbf{nf}(t_2)$ (and aborts if this is not the case)
- **assert.ok** $t_1 t_2$: checks if $\mathbf{nf}(t_1) = \square(\mathbf{nf}(t_2))$ (and aborts if this is not the case)
- **print** t_1 : prints a string representation of t_1

The last three commands are used to debug the proofs.

The proofs, after initial skolemization, are structured as a sequence of clause derivations, using either resolution, theory conflicts, instantiation or CNF-conversion steps. All these clauses are **let**-bound, until we reach the empty clause. Basically we end up with a proof-tree in natural deduction, deriving the Boolean false constant from the initial formula. The tree is encoded as a DAG, because **let**-bound clauses can be used more than once.

Proof rule	Rewrite rule
$\frac{\Gamma \vdash \psi_1 \wedge \psi_2}{\Gamma \vdash \psi_1}$ ($\text{elim}_{\wedge 1}$)	$\text{elim}_{\wedge 1}(\Box(\text{and}(x_1, x_2))) \blacktriangleright$ $\Box(x_1)$
$\frac{\Gamma \vdash \psi_1 \wedge \psi_2}{\Gamma \vdash \psi_2}$ ($\text{elim}_{\wedge 2}$)	$\text{elim}_{\wedge 2}(\Box(\text{and}(x_1, x_2))) \blacktriangleright$ $\Box(x_2)$
$\frac{\Gamma \vdash \psi_1 \vee \psi_2}{\Gamma \vdash \neg \psi_1 \rightarrow \psi_2}$ (elim_{\vee})	$\text{elim}_{\vee}(\Box(\text{or}(x_1, x_2))) \blacktriangleright$ $\Box(\text{implies}(\text{neg}(x_1), x_2))$
$\frac{\Gamma \vdash \neg \psi \quad \Gamma \vdash \psi}{\Gamma \vdash \perp}$ (elim_{\neg})	$\text{elim}_{\neg}(\Box(\text{neg}(x_1)), \Box((x_1))) \blacktriangleright$ $\Box(\text{false})$
$\frac{\Gamma \vdash \psi}{\Gamma \vdash \neg \neg \psi}$ (add_{\neg})	$\text{add}_{\neg}(\Box(x_1)) \blacktriangleright$ $\Box(\text{neg}(\text{neg}(x_1)))$
$\frac{\Gamma \vdash \psi \rightarrow \perp}{\Gamma \vdash \neg \psi}$ (intro_{\neg})	$\text{intro}_{\neg}(\Box(\text{implies}(x_1, \text{false}))) \blacktriangleright$ $\Box(\text{neg}(x_1))$
$\frac{\Gamma \vdash \neg \psi \rightarrow \perp}{\Gamma \vdash \psi}$ ($\text{elim}_{\neg \rightarrow}$)	$\text{elim}_{\neg \rightarrow}(\Box(\text{implies}(\text{neg}(x_1), \text{false}))) \blacktriangleright$ $\Box(x_1)$
$\frac{\Gamma \vdash \forall x. \psi(x)}{\Gamma \vdash \psi(t)}$ (inst)	$\text{inst}(y, \Box(\text{forall}(x))) \blacktriangleright$ $\Box(\text{apply} \cdot (x, y))$

Fig. 4. Additional rules for the example

All those steps are best described through an example³. Fig. 4 lists rules not previously mentioned in this paper, that were used in the proof. The real proof system has more rules. As described in Sect. 6.1, we mechanically check all rules. Our example formula is:

$$P(c) \wedge (c = d) \wedge (\forall x. \neg P(x) \vee \neg(\forall y. \neg Q(x, y))) \wedge (\forall x. \neg Q(d, x))$$

The first step is the initial skolemization:

```

let  $q_1 := \text{forall}(\lambda x. \text{or}(\text{neg}(P(x)), \text{neg}(\text{forall}(\lambda y. \text{neg}(Q(x, y)))))$ 
let  $q_2 := \text{forall}(\lambda x. \text{neg}(Q(d, x)))$ 
let  $f_{in} := \text{and}(P(c), \text{and}(\text{eq}(c, d), \text{and}(q_1, q_2)))$ 
let  $sk := \text{rec}_{\wedge}(\text{id}, \text{rec}_{\wedge}(\text{id}, \text{rec}_{\wedge}(\text{skip}_{\vee}(\text{rec}_{\vee}(\text{id}, \text{skol}(f, \text{rec}_{\neg}(\text{id}))), \text{id})))$ ,  $\text{id}$ ))
initial  $sk f_{in}$ 

```

Here our expectation, as the proof generator, is that $\forall x. \neg P(x) \vee \neg(\forall y. \neg Q(x, y))$ will be replaced by $\forall x. \neg P(x) \vee Q(x, f(x))$, which we express as:

```

let  $q_3 := \text{forall}(\lambda x. \text{or}(\text{neg}(P(x)), Q(x, f(x))))$ 
let  $f_{sk} := \text{and}(P(c), \text{and}(\text{eq}(c, d), \text{and}(q_3, q_2)))$ 
assert_ok initial  $f_{sk}$ 

```

The first step of the actual proof is a partial CNF-conversion. Our CNF conversion uses Tseitin scheme, which introduces proxy literals for subformulas. This

³ The proof presented here is not the simplest possible of this very formula. However it follows the steps that our SMT solver does and we expect other SMT solvers to do.

produces equisatisfiable set of clauses, yet the proof maps the proxy literals back to the original subformulas. Then the defining clauses of proxy literals become just basic Boolean facts. We therefore derive clauses of the form $f_{sk} \rightarrow \neg\psi \rightarrow \perp$, where ψ is one of the conjuncts of f_{sk} , for example:

```
let c1 := assume(fsk, λf. assume(neg(eq(c, d)), λp. elim-(p, elim∧2(elim∧1(f))))))
assert_ok c1 implies(fsk, implies(neg(eq(c, d)), false))
```

and similarly we derive:

```
assert_ok c0 implies(fsk, implies(neg(P(c)), false))
assert_ok c2 implies(fsk, implies(neg(q3), false))
assert_ok c3 implies(fsk, implies(neg(q2), false))
```

Next we instantiate the quantifiers:

```
let c4 := assume(q2, λq. assume(Q(d, f(c)), λi. elim-(inst(f(c), q), i)))
assert_ok c4 implies(q2, implies(Q(d, f(c)), false))
let i1 := or(neg(P(c)), Q(c, f(c)))
let c5 := assume(q3, λq. assume(neg(i1), λi. elim-(i, inst(c, q))))
assert_ok c5 implies(q3, implies(neg(i1), false))
```

Then we need to classify i_1 :

```
let c6 := assume(i1, λi. assume(P(c), λo1. assume(neg(Q(c, f(c))), λo2.
  elim-(o2, mp(elim∨(i), add-(o1))))))
assert_ok c6 implies(i1, implies(P(c), implies(neg(Q(c, f(c))), false)))
```

Then we do some equality reasoning:

```
let c7 := assume(neg(Q(d, f(c))), λln. assume(eq(c, d), λe. assume(Q(c, f(c)), λlp.
  elim-(ln, mp(eq_sub(e, λx. Q(x, f(c))), lp))))))
assert_ok c7 implies(neg(Q(d, f(c))), implies(eq(c, d), implies(Q(c, f(c)), false)))
```

What remains is a pure Boolean resolution. The resolution is realized by assuming the negation of the final clause and then using unit resolution of the assumed literals and some previous clauses, to obtain new literals, and as a last step, the false constant. We first resolve c_4 with c_7 :

```
let c8 := assume(q2, λl1. assume(eq(c, d), λl2. assume(Q(c, f(c)), λl3.
  mp(mp(mp(c7, intro-(mp(c4, l1))), l2), l3))
assert_ok c8 implies q2, implies eq(c, d), implies Q(c, f(c)), false
```

and finally we derive (also through resolution) the false constant:

```
let kq2 := elim-(mp(c3, initial))
let kq3 := elim-(mp(c2, initial))
let kp := elim-(mp(c0, initial))
let ke := elim-(mp(c1, initial))
let kq := elim-(mp(mp(c6, elim-(mp(c5, kq3))), kp))
let c9 := mp(mp(mp(c8, kq2), ke), kq)
final c9
```

6 Implementation

We have implemented two versions of the proof checker: one full version in OCaml and a simplified one written in C. Proof generation was implemented inside the Fx7 [1] SMT solver, implemented in the Nemerle programming language. The solver came second in the AUFLIA division of 2007 SMT competition, being much slower, but having solved the same number of benchmarks as the winner, Z3 [6].

An important point about the implementation, is that at any given point, we need to store only terms, that can be referenced by **let**-bound name, and thus the memory used by other terms can be reclaimed. As in our encoding the proof terms actually rewrite to formulas that they prove, there is no need to keep the proof terms around. We suspect this to be the main key to memory efficiency of the proof checker. The C implementation exploits this fact, the OCaml one does not.

Both implementations use de Bruijn [5] indices in representation of lambda terms. We also use hash consing, to keep only a single copy of a given term. We cache normal forms of the terms, we remember what terms are closed (which speeds up beta reductions). Also a local memoization is used in function computing beta reduction to exploit the DAG structure of the term. The rewrite rules are only indexed by the head symbol, if two rules share the head symbol, linear search is used.

All the memoization techniques used are crucial (i.e., we have found proofs, where checking would not finish in hours without them).

The OCaml implementation is about 900 lines of code, where about 300 lines is pretty printing for Coq and Maude formats. The C implementation is 1500 lines. Both implementation include parsing of the proof and SMT formats and command line option handling. The implementations are available online along with the Fx7 prover.

6.1 Soundness Checking

The OCaml version of the checker has also a different mode of operation, where it reads the rewrite rules and generates corresponding formulas to be proven in the Coq proof assistant. There are three proof modes for rules:

- for simple facts about Boolean connectives, arithmetic and equality, the checker generates a lemma and a proof, which is just an invocation of appropriate tactic
- for other generic schemas of proof rules from Sect. 3 and 4, the checker produces proof obligations, and the proofs need to be embedded in the rule descriptions
- for non-generic proof rules, the user can embed both the lemma and the proof in the rule description file, just to keep them close

This semiautomatic process helps preventing simple, low-level mistakes in the proof rules. The checker provides commands to define all these kinds of rules and associated proofs.

Directory	Total	UNSAT	% UNSAT	Fake	% Fake	Fail	% Fail
front_end_suite	2320	2207	95.13%	101	4.35%	12	0.52%
boogie	908	866	95.37%	25	2.75%	17	1.87%
simplify	833	729	87.52%	44	5.28%	60	7.20%
piVC	41	17	41.46%	10	24.39%	14	34.15%
misc	20	16	80.00%	0	0.00%	4	20.00%
Burns	14	14	100.00%	0	0.00%	0	0.00%
RicartAgrawala	14	13	92.86%	0	0.00%	1	7.14%
small_suite	10	8	80.00%	0	0.00%	2	20.00%

Fig. 5. Results on the AUFLIA division of SMT-LIB

6.2 Performance Evaluation

When running Fx7 on a query there are five possible outcomes:

- it reports that the query is unsatisfiable, and outputs a proof
- it reports that the query is unsatisfiable, but because the proof generation is only implemented for the UTVPI fragment of linear arithmetic, the proof is correct only if we assume the theory conflicts to be valid (there is typically a few of them in each of such “fake” proofs)
- it reports the query is satisfiable, timeouts or runs out of memory

Tests were performed on AUFLIA benchmarks from the SMT-LIB [14]. This division includes first order formulas, possibly with quantifiers, interpreted under uninterpreted function symbols, integer linear arithmetic and array theories. They are mostly software verification queries. The machine used was a 2.66GHz Pentium 4 PC with 1GB of RAM, running Linux. The time limit was set to ten minutes.

The results are given in Fig. 5. The “Total” column refers to the number of benchmarks marked unsatisfiable in the SMT-LIB; “UNSAT” refers to the number of cases, where the benchmark was found unsatisfiable and a correct proof was generated; “Fake” is the number of benchmarks found unsatisfiable, but with “fake” proofs; finally “Fail” is the number of cases, where Fx7 was unable to prove it within the time limit. It should be the case that UNSAT + Fake + Fail = Total. The percentages are with respect to the Total.

With the **C implementation**, proof checking a single proof never took more than 7 seconds. It took more than 2 seconds in 4 cases and more than 1 second in 19 cases (therefore the average time is well under a second). The maximal amount of memory consumed for a single proof was never over 7MB, with average being 2MB.

We have also tested the C implementation on a Dell x50v **PDA** with a 624MHz XScale ARM CPU and 32MB of RAM, running Windows CE. It was about 6 times slower than the Pentium machine, but was otherwise perfectly capable of running the checker. This fact can be thought of as a first step on a way to PCC-like scenarios on small, mobile devices. Other devices of similar computing power and, what is more important, RAM amount include most smart phones and iPods.

The **OCaml implementation** was on average 3 times slower than the C version, it also tends to consume more memory, mostly because it keeps all the terms forever (which is because of our implementation, not because of OCaml).

We have also experimented with translating the proof objects into the **Maude syntax** [3]. We have implemented lambda terms and beta reduction using the built-in Maude integers to encode de Bruijn indices and used the standard equational specifications for the first order rules. The resulting Maude implementation is very compact (about 60 lines), but the performance is not as good as with the OCaml or C implementation — it is between 10 and 100 times slower than the OCaml one. It also tends to consume a lot more memory. The reason is mainly the non-native handling of lambda expressions. Beta reductions translate to large number of first order rewrites, which are then memoized, and we were unable to instrument Maude to skip memoization of those.

We have performed some experiments using **Coq metalogic** as the proof checker. We did not get as far as implementing our own object logic. The main obstacle we have found was the treatment of binders. Performing skolemization on a typical input results in hundreds of Skolem functions. When using a higher order logic prover, such functions are existentially quantified and the quantifiers need to be pushed through the entire formula to the beginning. Later, during the proof, we need to go through them to manipulate the formula. This puts too much pressure on the algorithms treating of binders in the higher order prover. In our approach Skolem functions are bound implicitly, so there is no need to move them around. This is especially important in SMT queries, where the vast majority of the input formula is ground and quantified subformulas occur only deep inside the input. We can therefore keep most of the formula binder-free. We were not able to perform any realistic tests, as Coq was running out of memory.

Both Maude and Coq are far more general purpose tools than just proof checkers. However relatively good results with Maude suggest that using a simple underlying formalism is beneficial in proof checking scenarios.

7 Related and Future Work

CVC3 [2] and Fx7 were the only solvers participating in the 2007 edition of the SMT competition to produce formal proofs. The proof generation in CVC3 is based on the LF framework. We are not aware of a published work evaluating proof checking techniques on large industrial benchmarks involving quantifiers.

Formalisms for checking SMT proofs have been proposed in the past, most notably using an optimized implementation [15] of Edinburgh Logical Framework [10]. However even with the proposed optimizations, the implementations has an order of magnitude higher memory requirements than our solution. Also the implementation of the checker is much more complicated.

Recently a Signature Compiler tool has been proposed [16]. It generates a custom proof checker in C++ or Java from a LF signature. We have run our proof checker on a 1:1 translation of the artificial EQ benchmarks from the paper. It is running slightly faster than the generated C++ checker. The memory

requirements of our implementation are way below the size of the input file on those benchmarks. The checkers remain to be compared on real benchmarks involving richer logics and quantifiers.

In context of the saturation theorem provers it is very natural to output the proof just as a sequence of resolution or superposition steps. What is missing here, is the proof of CNF translation, though proof systems has been proposed [8], [7] to deal with that.

Finally, work on integrating SMT solvers as decision procedures inside higher order logic provers include [12], [9], [4]. The main problem with these approaches is that proof generation is usually at least order of magnitude faster than proof checking inside higher order logic prover. The Ergo [4] paper mentions promising preliminary results with using proof traces instead of full proofs with Coq for theory conflicts. It is possible that using traces could also work for CNF conversion and skolemization. Yet another approach mentioned there is verifying the SMT solver itself.

An important remaining problem is the treatment of theory conflicts. One scenario here is to extend the linear arithmetic decision procedure to produce proofs. It should be possible to encode the proofs with just a minor extensions to the rewrite formalism. Another feasible scenario is to use a different SMT solver as a oracle for checking the harder (or all) theory conflicts. This can be applied also to other theories, like bit vectors or rational arithmetic.

8 Conclusions

We have shown how term rewriting can be used for proof checking. The highlights of our approach are (1) time and space efficiency of the proof checker; (2) simplicity of the formalism, and thus simplicity of the implementation; and (3) semiautomatic checking of proof rules. The main technical insight is that the proof rules can be executed locally. Therefore the memory taken by proofs trees can be reclaimed just after checking them and reused for the subsequent fragments of the proof tree.

The author wishes to thank Joe Kiniry, Mikoláš Janota, and Radu Grigore for their help during the work on the system, and Nikolaj Bjørner as well as anonymous TACAS reviewers for his help in getting the presentation of this paper better.

This work was partially supported by Polish Ministry of Science and Education grant 3 T11C 042 30.

References

1. Fx7 web page, <http://nemerle.org/fx7/>
2. Barrett, C., Berezin, S.: CVC Lite: A new implementation of the Cooperating Validity Checker. In: Alur, R., Peled, D.A. (eds.) CAV 2004. LNCS, vol. 3114, pp. 515–518. Springer, Heidelberg (2004)

3. Clavel, M., Durán, F., Eker, S., Lincoln, P., Martí-Oliet, N., Meseguer, J., Quesada, J.F.: Maude: Specification and programming in rewriting logic. *Theoretical Computer Science* (2001)
4. Conchon, S., Contejean, E., Kanig, J., Lescuyer, S.: Lightweight Integration of the Ergo Theorem Prover inside a Proof Assistant. In: *Second Automated Formal Methods workshop series (AFM 2007)*, Atlanta, Georgia, USA (November 2007)
5. de Bruijn, N.G.: Lambda-calculus notation with nameless dummies: a tool for automatic formula manipulation with application to the Church-Rosser theorem. *Indag. Math.* 34(5), 381–392 (1972)
6. de Moura, L., Bjorner, N.: Efficient E-matching for SMT solvers. In: *Proceedings of the 21st International Conference on Automated Deduction (CADE-21)*, Springer, Heidelberg (to appear, 2007)
7. de Nivelle, H.: Implementing the clausal normal form transformation with proof generation. In: *fourth workshop on the implementation of logics*, Almaty, Kazakhstan, University of Liverpool, University of Manchester, pp. 69–83 (2003)
8. de Nivelle, H.: Translation of resolution proofs into short first-order proofs without choice axioms. *Information and Computation* 199(1), 24–54 (2005)
9. Fontaine, P., Marion, J.-Y., Merz, S., Nieto, L.P., Tiu, A.: Expressiveness + automation + soundness: Towards combining SMT solvers and interactive proof assistants. In: *Hermanns, H., Palsberg, J. (eds.) TACAS 2006. LNCS, vol. 3920*, pp. 167–181. Springer, Heidelberg (2006)
10. Harper, R., Honsell, F., Plotkin, G.: A framework for defining logics. In: *Proceedings 2nd Annual IEEE Symp. on Logic in Computer Science, LICS 1987*, Ithaca, NY, USA, June, 22–25, 1987, pp. 194–204. IEEE Computer Society Press, New York (1987)
11. Harvey, W., Stuckey, P.: A unit two variable per inequality integer constraint solver for constraint logic programming (1997)
12. McLaughlin, S., Barrett, C., Ge, Y.: Cooperating theorem provers: A case study combining HOL-Light and CVC Lite. In: *Armando, A., Cimatti, A. (eds.) Proceedings of the 3rd Workshop on Pragmatics of Decision Procedures in Automated Reasoning (PDPAR 2005)*, Edinburgh, Scotland, January 2006. *Electronic Notes in Theoretical Computer Science*, vol. 144(2), pp. 43–51. Elsevier, Amsterdam (2006)
13. Necula, G.C.: Proof-carrying code. In: *Conference Record of POPL 1997: The 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, Paris, France, January 1997, pp. 106–119 (1997)
14. SMT-LIB: The Satisfiability Modulo Theories Library. <http://www.smt-lib.org/>
15. Stump, A., Dill, D.: Faster Proof Checking in the Edinburgh Logical Framework. In: *18th International Conference on Automated Deduction* (2002)
16. Zeller, M., Stump, A., Deters, M.: A signature compiler for the Edinburgh Logical Framework. In: *Proceedings of International Workshop on Logical Frameworks and Meta-Languages: Theory and Practice* (2007)