

A Model-Constructing Satisfiability Calculus

Leonardo de Moura¹ and Dejan Jovanović²

¹ Microsoft Research

² New York University

Abstract. We present a new calculus where recent model-based decision procedures and techniques can be justified and combined with the standard DPLL(T) approach to satisfiability modulo theories. The new calculus generalizes the ideas found in CDCL-style propositional SAT solvers to the first-order setting.

1 Introduction

Considering the theoretical hardness of SAT, the astonishing adeptness of SAT solvers when attacking practical problems has changed the way we perceive the limits of algorithmic reasoning. Modern SAT solvers are based on the idea of *conflict driven clause learning* (CDCL) [11,15,13]. The CDCL algorithm is a combination of an explicit backtracking search for a satisfying assignment complemented with a deduction system based on Boolean resolution. In this combination, the worst-case complexity of both components is circumvented by the components guiding and focusing each other. The generalization of the SAT problem into the first-order domain is called satisfiability modulo theories (SMT). The common way to solve an SMT problem is to employ a SAT solver to enumerate the assignment of the Boolean abstraction of the formula. The candidate Boolean assignment is then either confirmed or refuted by a *decision procedure* dedicated to reasoning about conjunctions of theory-specific constraints. This framework is commonly called DPLL(T) [10,14] and is employed by most of the SMT solvers today. Although DPLL(T) at its core relies on a CDCL SAT solver, this SAT solver is only used as a black-box. This can be seen as an advantage since the advances in SAT easily transfer to performance improvements in SMT. On the other hand, in the last few years the idea of direct model construction complemented with conflict resolution has been successfully generalized to fragments of SMT dealing with theories such as linear real arithmetic [4,12,9], linear integer arithmetic [7], nonlinear arithmetic [8], and floating-point [6]. All these procedures, although quite effective in their corresponding first-order domains, have not seen a more widespread acceptance due to their limitations in purely Boolean reasoning and incompatibility with DPLL(T).

In this paper we propose a *model-constructing satisfiability calculus* (mcSAT) that encompasses all the decision procedures above, including the decision procedures aimed at DPLL(T), while resolving the limitations mentioned above. The mcSAT framework extends DPLL(T) by allowing assignments of variables to concrete values, while relaxing the restriction that decisions, propagations, and explanations of conflicts must be in term of existing atoms.

2 A Model Based Abstract Procedure

We assume that the reader is familiar with the usual notions and terminology of first-order logic and model theory (for an introduction see e.g. [2]). We describe the new procedure as an abstract transition system in the spirit of Abstract DPLL [14]. The crucial difference of the system we present is that we are not restricted to Boolean decisions. Instead, we allow the model that the theory is trying to construct to be involved in the search and in explaining the conflicts, while allowing new literals to be introduced so as to support more complex conflict analyses.

The states in the transition system are pairs of the form $\langle M, \mathcal{C} \rangle$, where M is a sequence (usually called a *trail*) of trail elements, and \mathcal{C} is a set of clauses. Each trail element is either a *decided literal*, a *propagated literal*, or a *model assignment*. We refer to both decided literals and model assignments as *decisions*. A decided literal is a literal that we assume to be true. On the other hand, a propagated literal, denoted as $C \rightarrow L$, marks a literal L that is implied to be true in the current state by the clause C (the explanation). In both cases, we say that the literal L appears in M , and write this as $L \in M$. A model assignment, written as $x \mapsto \alpha$, is an assignment of a first-order uninterpreted symbol x to a value α .¹ Given a trail M that contains model assignments $x_{i_1} \mapsto \alpha_1, \dots, x_{i_k} \mapsto \alpha_k$, we can construct a first-order interpretation $v[M] = [x_{i_1} \mapsto \alpha_1, \dots, x_{i_k} \mapsto \alpha_k]$. Given a term t , the interpretation $v[M](t)$ is either a value of the term t under the assignment in M , or **undef** if the term cannot be fully evaluated.

The content of the trail implies an interpretation of literals and is the core of our procedure. In order to evaluate the value of some literal L with respect to a trail M , we define the functions value_B and value_T , the former interpreting the literal according to the Boolean assignment, and the latter interpreting the literal according to the model assignment of variables.

$$\text{value}_B(L, M) = \begin{cases} \text{true} & L \in M \\ \text{false} & \neg L \in M \\ \text{undef} & \text{otherwise} \end{cases} \quad \text{value}_T(L, M) = \begin{cases} \text{true} & v[M](L) = \text{true} \\ \text{false} & v[M](L) = \text{false} \\ \text{undef} & \text{otherwise} \end{cases}$$

We say that a trail M is *consistent* if the Boolean assignment and first-order model are not in conflict, i.e. when for all $L \in M$ we have that $\text{value}_T(L, M) \neq \text{false}$. Additionally we say that the trail M is *complete* when each asserted first-order literal $L \in M$ is justified by the first-order interpretation, i.e. $\text{value}_T(L, M) = \text{true}$. We use the predicate $\text{consistent}(M)$ to denote that M is consistent and $\text{complete}(M)$ to denote that M is complete. Note that if a trail M is consistent, this does not mean that the assertions on the trail are truly satisfiable (feasible), just that the current partial assignment does not refute any of the individual trail literals. When there is a set of literals on the

¹ The actual representation for values is theory specific and depends on the type of x . For example, for the theory of linear real arithmetic, the values are rational numbers. We never assign Boolean variables to values as they are considered literals.

trail M that, together with the model assignments from M , is not satisfiable, we call the trail *infeasible* and denote this with the predicate $\text{infeasible}(M)$. We have that $\neg \text{consistent}(M)$ implies $\text{infeasible}(M)$.

Since the values of $\text{value}_T(L, M)$ and $\text{value}_B(L, M)$ do not disagree for all $L \in M$, we define the value of a literal in a consistent state as

$$\text{value}(L, M) = \begin{cases} \text{value}_B(L, M) & \text{value}_B(L, M) \neq \text{undef}, \\ \text{value}_T(L, M) & \text{otherwise.} \end{cases}$$

Example 1. Consider the trail $M = \llbracket x > 0, x \mapsto 1, y \mapsto 0, z > 0 \rrbracket$. The model interpretation according to M is $v[M] = [x \mapsto 1, y \mapsto 0]$. Therefore we have that $\text{value}_T(x > 0, M) = \text{value}_B(x > 0, M) = \text{true}$, $\text{value}_T(x > 1, M) = \text{false}$, $\text{value}_T(z > 0, M) = \text{undef}$, $\text{value}_B(z > 0, M) = \text{true}$, $v[M](x + y + 1) = 2$, and $v[M](x + z) = \text{undef}$. The trail M is consistent, but $M' = \llbracket M, y < 0 \rrbracket$ is not because $\text{value}_T(y < 0, M') = \text{false}$ and $\text{value}_B(y < 0, M') = \text{true}$. The trail M is not complete as it does not interpret z and therefore $\text{value}_T(z > 0, M) = \text{undef}$. Finally, $M'' = \llbracket M, z < x \rrbracket$ is infeasible because $\{x \mapsto 1, z > 0, z < x\}$ is unsatisfiable, but M'' is consistent.

We extend the definition of value to clauses so that $\text{value}(C, M) = \text{true}$ if at least one literal of C evaluates to true , $\text{value}(C, M) = \text{false}$ if all literals evaluate to false , and $\text{value}(C, M) = \text{undef}$ otherwise. We say a clause C is satisfied by trail M if $\text{value}(C, M) = \text{true}$. A set of clauses \mathcal{C} is satisfied by M if M is complete (and therefore consistent), and all clauses $C \in \mathcal{C}$ are satisfied by M . We use the predicate $\text{satisfied}(\mathcal{C}, M)$ to denote that \mathcal{C} is satisfied by M ,

Given a set of clauses \mathcal{C}_0 , our procedure starts with the state $\langle \llbracket \rrbracket, \mathcal{C}_0 \rangle$ and performs transitions according to the rules we explain below. The goal is to either enter into a state sat denoting that the problem is satisfiable, or into a state unsat denoting that the problem is unsatisfiable. The states we traverse are either *search states* of the form $\langle M, \mathcal{C} \rangle$ or *conflict resolution states* of the form $\langle M, \mathcal{C} \rangle \vdash C$. In both types of states we keep the invariant that M is a consistent trail and $\mathcal{C}_0 \subseteq \mathcal{C}$. Additionally, in conflict resolution states the clause C is always a clause implied by \mathcal{C}_0 and refuted by the trail, i.e. $\mathcal{C}_0 \models C$ and $\text{value}(C) = \text{false}$. We call the clause C the *conflicting clause*.

To ensure termination, the transition system assumes existence of a finite set of literals \mathbb{B} that we call the *finite basis*. During a derivation of the system, any literal added to the trail will be from \mathbb{B} , and the clauses that the system uses will only contain literals from \mathbb{B} .² The minimal assumption is that \mathbb{B} must include all literals (and their negations) from the initial problem \mathcal{C}_0 , and the theory-specific decision procedure must ensure that for any \mathcal{C}_0 such a finite basis exists.

2.1 Clausal Rules

We start by presenting the set of search rules and conflict analysis rules that resemble those of abstract DPLL and are the backbone of CDCL-style SAT

² Our finite basis corresponds to the closure of the literal-generating function used in splitting-on-demand [1].

DECIDE		
$\langle M, \mathcal{C} \rangle \longrightarrow \langle \llbracket M, L \rrbracket, \mathcal{C} \rangle$	\mathbf{if}	$L \in \mathbb{B}, \text{value}(L, M) = \text{undef}$
PROPAGATE		
$\langle M, \mathcal{C} \rangle \longrightarrow \langle \llbracket M, C \rightarrow L \rrbracket, \mathcal{C} \rangle$	\mathbf{if}	$C = (L_1 \vee \dots \vee L_n \vee L) \in \mathcal{C}$ $\forall i : \text{value}(L_i, M) = \text{false}$ $\text{value}(L, M) = \text{undef}$
CONFLICT		
$\langle M, \mathcal{C} \rangle \longrightarrow \langle M, \mathcal{C} \rangle \vdash C$	\mathbf{if}	$C \in \mathcal{C}, \text{value}(C) = \text{false}$
SAT		
$\langle M, \mathcal{C} \rangle \longrightarrow \text{sat}$	\mathbf{if}	$\text{satisfied}(\mathcal{C}, M)$
FORGET		
$\langle M, \mathcal{C} \rangle \longrightarrow \langle M, \mathcal{C} \setminus \{C\} \rangle$	\mathbf{if}	$C \in \mathcal{C}$ is a learned clause.

Fig. 1. Clausal search rules

solvers. The clausal search rules are presented in Fig. 1 and the clausal conflict analysis rules are presented in Fig. 2.

Search Rules. The DECIDE rule can take any literal from the basis \mathbb{B} that does not yet have a value added to the trail. The PROPAGATE performs Boolean clause propagation by assigning to true the only unassigned literal from a clause where all other literals already evaluate to false.³ If we encounter a clause C such all literals in C evaluate to false, we use the CONFLICT to enter conflict resolution state. During conflict analysis we can learn new clauses, and these can be removed using the FORGET rule. If our trail is complete and satisfies all the clauses, we can finish the search using the SAT rule.

Conflict Analysis Rules. As in CDCL, the conflict analysis rules recover from a conflict encountered during the search, learn the reason of the conflict, and backtrack to an appropriate state to continue the search. The main analysis rule is the RESOLVE rule. This rule performs Boolean resolution of clauses C and D over the literal L , obtaining the clause $R = \text{resolve}(C, D, L)$. The clause R is a valid deduction and moreover evaluates to false in M . If the result of the resolution is an empty clause (denoted with false), we can deduce that the problem is unsatisfiable using the UNSAT rule. Since, in a conflict analysis state $\langle M, \mathcal{C} \rangle \vdash C$, the clause C is always a valid deduction, we can use the LEARN to add the clause C to set of clauses in order to aid in reducing the search space in the future. If this learned clause is at a later point not deemed useful, we can remove it using the FORGET rule. The CONSUME rules skips over those decided and propagated literals in the trail that are not relevant for the inconsistency of

³ Note that in both DECIDE and PROPAGATE we do not need to ensure that the new literal does not cause an infeasibility. Additionally, we can decide literals that do not yet exist in \mathcal{C} , enabling refinement decisions as found in [6].

RESOLVE	
$\langle \llbracket M, D \rightarrow L \rrbracket, C \rangle \vdash C \longrightarrow \langle M, C \rangle \vdash R$	if $\neg L \in C,$ $R = \text{resolve}(C, D, L)$
CONSUME	
$\langle \llbracket M, D \rightarrow L \rrbracket, C \rangle \vdash C \longrightarrow \langle M, C \rangle \vdash C$	if $\neg L \notin C$
$\langle \llbracket M, L \rrbracket, C \rangle \vdash C \longrightarrow \langle M, C \rangle \vdash C$	if $\neg L \notin C$
BACKJUMP	
$\langle \llbracket M, N \rrbracket, C \rangle \vdash C \longrightarrow \langle \llbracket M, C \rightarrow L \rrbracket, C \rangle$	if $C = L_1 \vee \dots \vee L_m \vee L$ $\forall i : \text{value}(L_i, M) = \text{false}$ $\text{value}(L, M) = \text{undef}$ N starts with a decision
UNSAT	
$\langle M, C \rangle \vdash \text{false} \longrightarrow \text{unsat}$	
LEARN	
$\langle M, C \rangle \vdash C \longrightarrow \langle M, C \cup \{C\} \rangle \vdash C$	if $C \notin C$

Fig. 2. Clausal conflict analysis rules

the conflicting clause. Finally, we exit conflict resolution using the BACKJUMP rule, if we have deduced a clause C such that implies a literal earlier in the trail, skipping over at least one decision.

2.2 Theory-Specific Rules

We now extend the rules to enable theory-specific reasoning, allowing deductions in the style of DPLL(T), but more flexible, and allowing for assignments of variables to particular concrete values (Figure 3). As in DPLL(T), the basic requirement for a theory decision procedure is to provide an explain function that can explain theory-specific propagations and infeasible states. In DPLL(T), the theory explanations are theory lemmas in form of clauses that only contain negations of literals asserted so far. The explanation function explain here has more flexibility. Given a literal L and a consistent trail M that implies L to be true, i.e. such that $\llbracket M, \neg L \rrbracket$ is infeasible, $\text{explain}(L, M)$ must return a valid theory lemma $E = L_1 \vee \dots \vee L_k \vee L$. The literals of the clause E must be from the finite basis \mathbb{B} , and all literals L_i must evaluate to false in M . Allowing explanations containing more than just the literals off the trail allows for more expressive lemmas and is crucial for model-based decision procedures. Limiting the literals to a finite basis, on the other hand, is important for ensuring the termination of the procedure.

Search Rules. We can propagate a literal L using the T-PROPAGATE rule, if it is implied in the current state, i.e. if adding the literal $\neg L$ to the trail M makes it infeasible. The side condition $\text{value}(L, M) = \text{undef}$ ensures that we cannot

T-PROPAGATE		
$\langle M, C \rangle$	$\longrightarrow \langle \llbracket M, E \rightarrow L \rrbracket, C \rangle$	if $L \in \mathbb{B}, \text{value}(L, M) = \text{undef}$ $\text{infeasible}(\llbracket M, \neg L \rrbracket)$ $E = \text{explain}(\llbracket M, \neg L \rrbracket)$
T-DECIDE		
$\langle M, C \rangle$	$\longrightarrow \langle \llbracket M, x \mapsto \alpha \rrbracket, C \rangle$	if $x \in \text{vars}_T(C)$ $v[M](x) = \text{undef}$ $\text{consistent}(\llbracket M, x \mapsto \alpha \rrbracket)$
T-CONFLICT		
$\langle M, C \rangle$	$\longrightarrow \langle M, C \rangle \vdash E$	if $\text{infeasible}(M)$ $E = \text{explain}(\text{false}, M)$
T-CONSUME		
$\langle \llbracket M, x \mapsto \alpha \rrbracket, C \rangle \vdash C$	$\longrightarrow \langle M, C \rangle \vdash C$	if $\text{value}(C, M) = \text{false}$
T-BACKJUMP-DECIDE		
$\langle \llbracket M, x \mapsto \alpha, N \rrbracket, C \rangle \vdash C$	$\longrightarrow \langle \llbracket M, L \rrbracket, C \rangle$	$C = L_1 \vee \dots \vee L_m \vee L$ if $\exists i : \text{value}(L_i, M) = \text{undef}$ $\text{value}(L, M) = \text{undef}$

Fig. 3. Theory search and conflict rules

produce an inconsistent trail ($\text{value}(L, M) = \text{false}$), nor include redundant information ($\text{value}(L, M) = \text{true}$). We use the T-DECIDE rule to assign an interpretation/value (α) to a variable x that occurs in our set of clauses ($x \in \text{vars}_T(C)$). The side condition $v[M](x) = \text{undef}$ ensures that we cannot assign a variable x that is already assigned in M , and $\text{consistent}(\llbracket M, x \mapsto \alpha \rrbracket)$ ensures the new trail is consistent. We require a consistent trail because the function value is not well defined for inconsistent trails. We use the T-CONFLICT rule to enter conflict resolution state, whenever we detect that the trail is infeasible.

Conflict Analysis Rules. The T-CONSUME rule is similar to the consume rules in Figure 2, as it skips over a decided model assignment $x \mapsto \alpha$ that is not relevant for the inconsistency of the conflicting clause. The assignment is not relevant because the conflicting clause C still evaluates to false after the assignment $x \mapsto \alpha$ is removed from the trail. We use the T-BACKJUMP-DECIDE rule when we reach an assignment $x \mapsto \alpha$ that is relevant for the inconsistency of the conflicting clause C , but the BACKJUMP rule is not applicable because C contains more than one literal that evaluates to undef after we remove the assignment $x \mapsto \alpha$ from the trail. The T-BACKJUMP-DECIDE rule may generate doubts about the termination argument for our abstract procedure, since it is just replacing a decision $x \mapsto \alpha$ with another decision L . In our termination proof, we justify that by assuming that Boolean decisions (L) have “bigger” weight than model assignment decisions ($x \mapsto \alpha$).

2.3 Producing Explanations

A crucial component of our framework is the explanation function `explain`. Given an infeasible trail M , it must be able to produce a valid theory lemma that is inconsistent with M using only literals from the finite basis.

If the infeasibility of M only depends on literals that already occur in M , then `explain` can simply return the inconsistent clause $\neg L_1 \vee \dots \vee \neg L_k$ where all literal L_i occur in M . For example, the trail $M = \llbracket x < 0, y > 1, x > 0 \rrbracket$ is infeasible, and `explain(false, M) = $\neg(x < 0) \vee \neg(x > 0)$` is a valid theory lemma that is inconsistent with M . In such cases, the T-CONFLICT and T-PROPAGATE rules correspond to the theory propagation and conflict rules from the DPLL(T) framework.

The more interesting case occurs when the infeasibility on a trail M also depends on decided model assignments $x \mapsto \alpha$ in M . Consider, for example, the infeasible (but consistent) trail

$$M = \llbracket y > 0, z > 0, x + y + z < 0, x \mapsto 0 \rrbracket .$$

It might be tempting to define `explain(false, M)` to produce the valid theory lemma

$$\neg(y > 0) \vee \neg(z > 0) \vee \neg(x + y + z < 0) \vee \neg(x = 0)$$

This naïve `explain` function that just replaces the assignments $x \mapsto \alpha$ with literals $x = \alpha$, is inadequate as it *does not* satisfy the finite basis requirement for theories that operate over infinite domains, such as the Integers or the Reals. Using such a function, in this example, we would be able to continue indefinitely by assigning x to 1, then 2, and so on.

In principle, for any theory that admits elimination of quantifiers, it is possible to construct an explanation function `explain` that satisfy the finite basis requirement. The basic idea is to eliminate all unassigned variables and produce an implied formula that is also inconsistent with the assigned variables in the infeasible trail. In the previous example, the variables y and z are unassigned, so we can separate the infeasible literals that the naïve `explain` function return into the literals from the trail and literals from assignments

$$A \equiv (y > 0) \wedge (z > 0) \wedge (x + y + z < 0) \quad B \equiv (x = 0) .$$

Given A , we use a quantifier elimination procedure to generate a CNF formula F of the form $C_1 \wedge \dots \wedge C_k$ that is equivalent to $(\exists y, z : A)$, and therefore also inconsistent with B while only using the assigned variables (in this case x). Note that B corresponds to the assignment $v[M]$, and each clause C_i evaluates to true or false under $v[M]$ because all variables occurring in C_i are assigned by $v[M]$. Moreover, at least one of these clauses must evaluate to false because F is inconsistent with B . Let I be the clause C_i that evaluates to false. We have that $A \implies I$ is a valid clause because

$$A \implies (\exists y, z : A) \iff F \implies I$$

Since I is inconsistent with B , the clause $A \implies I$ will also be inconsistent with the trail and can be used as an explanation. Moreover, I is an *interpolant*

for A and B . In this example, we obtain $0 < x$ by eliminating the variables y and z from A using Fourier-Motzkin elimination, resulting in the explanation $A \implies (0 < x)$. When solving a set of linear arithmetic \mathcal{C} , Fourier-Motzkin elimination is sufficient to define the explain function, as shown in [12,9]. Fourier-Motzkin elimination gives a finite-basis \mathbb{B} with respect to \mathcal{C} , and the basis can be obtained by closing \mathcal{C} under the application of Fourier-Motzkin elimination step. It is fairly easy to show that the closure is a finite set, since we always produce constraints with one variable less.

In the last example, $0 < x$ is an *interpolant* for A and B , but so is $x \neq 0$. The key point is that an arbitrary interpolation procedure does not guarantee a finite basis. Nonetheless, this observation is useful when designing explanation functions for more complex theories. For nonlinear arithmetic constraints, we describe how to produce an explain procedure that produces an interpolant based on cylindrical algebraic decomposition (CAD) in [8]. The theory of uninterpreted functions (UF) does not admit quantifier elimination, but *dynamic-ackermannization* [5] can be used to create an interpolant I between A and B without compromising the finite basis requirement. For example, the trail $M = \llbracket x \mapsto 0, y \mapsto 0, f(x) \neq f(y) \rrbracket$ is infeasible with respect to UF. Then, we have $A \equiv f(x) \neq f(y)$ and $B \equiv (x = 0) \wedge (y = 0)$. Using dynamic-ackermannization we have that $I \equiv x \neq y$ is an interpolant for A and B , and $f(x) \neq f(y) \implies x \neq y$ is a valid theory lemma that is also inconsistent with M . We satisfy the finite basis requirement because dynamic-ackermannization only “introduces” equalities between terms that already exist in the trail. Finally, an explain function for the theory of arrays can be built on top of the approach described in [3].

Example 2. For the sake of simplicity, we restrict ourselves to the case of linear arithmetic. We illustrate the search rules by applying them to the following clauses over Boolean and linear arithmetic atoms $\mathcal{C} = \{x < 1 \vee p, \neg p \vee x = 2\}$. We start the deduction from the initial state $\langle \llbracket \rrbracket, \mathcal{C} \rangle$ and apply the rules of the mcSAT system.

$$\begin{aligned} & \langle \llbracket \rrbracket, \mathcal{C} \rangle \\ & \downarrow \text{T-DECIDE } (x \mapsto 1) \\ & \langle \llbracket x \mapsto 1 \rrbracket, \mathcal{C} \rangle \\ & \downarrow \text{PROPAGATE } (p \text{ must be true, since } x < 1 \text{ evaluates to false in the current trail}) \\ & \langle \llbracket x \mapsto 1, (x < 1 \vee p) \rightarrow p \rrbracket, \mathcal{C} \rangle \\ & \downarrow \text{CONFLICT } (\neg p \text{ and } x = 2 \text{ evaluate to false in the current trail}) \\ & \langle \llbracket x \mapsto 1, (x < 1 \vee p) \rightarrow p \rrbracket, \mathcal{C} \rangle \vdash \neg p \vee x = 2 \\ & \downarrow \text{RESOLVE } (\text{resolving } x < 1 \vee p \text{ and } \neg p \vee x = 2) \\ & \langle \llbracket x \mapsto 1 \rrbracket, \mathcal{C} \rangle \vdash x < 1 \vee x = 2 \end{aligned}$$

In this state, the BACKJUMP rule is not applicable because in the conflicting clause, both $x < 1$ and $x = 2$ evaluate to `undef` after the model assignment $x \mapsto 1$ is removed from the trail. The intuition is that the model assignment was “premature”, and the clause $x < 1 \vee x = 2$ is indicating that we should decide $x < 1$ or $x = 2$ before we assign x .

$\langle \llbracket x \mapsto 1 \rrbracket, \mathcal{C} \rangle \vdash x < 1 \vee x = 2$
 \downarrow T-BACKJUMP-DECIDE
 $\langle \llbracket x < 1 \rrbracket, \mathcal{C} \rangle$
 \downarrow T-DECIDE (the model assignment must satisfy $x < 1$)
 $\langle \llbracket x < 1, x \mapsto 0 \rrbracket, \mathcal{C} \rangle$
 \downarrow PROPAGATE ($\neg p$ must be true, since $x = 2$ evaluates to **false** in the current trail)
 $\langle \llbracket x < 1, x \mapsto 0, (\neg p \vee x = 2) \rightarrow \neg p \rrbracket, \mathcal{C} \rangle$
 \downarrow SAT ($x \mapsto 0$ and $\neg p$ satisfy all clauses)
sat

Example 3. Now, we consider a set of linear arithmetic (unit) clauses

$$\mathcal{C} = \{x < 1, x < y, 1 < z, z < x\} .$$

To simplify the presentation of this example, with a small abuse of notation, we use $\rightarrow L$ instead of $L \rightarrow L$, whenever the literal L is implied by the unit clause L .

$\langle \llbracket \rrbracket, \mathcal{C} \rangle$
 \downarrow PROPAGATE $\times 4$ (propagate all unit clauses)
 $\langle \llbracket \rightarrow x < 1, \rightarrow x < y, \rightarrow 1 < z, \rightarrow z < x \rrbracket, \mathcal{C} \rangle$
 \downarrow T-DECIDE (the current trail is consistent with the model assignment $x \mapsto 0$)
 $\langle \llbracket \rightarrow x < 1, \rightarrow x < y, \rightarrow 1 < z, \rightarrow z < x, x \mapsto 0 \rrbracket, \mathcal{C} \rangle$
 \downarrow T-DECIDE (peek a value for y , keeping consistency, y s.t. $x < y$)
 $\langle \llbracket \rightarrow x < 1, \rightarrow x < y, \rightarrow 1 < z, \rightarrow z < x, x \mapsto 0, y \mapsto 1 \rrbracket, \mathcal{C} \rangle$
 \downarrow T-CONFLICT ($1 < z$ and $z < x$ implies that $1 < x$)
 $\langle \llbracket \rightarrow x < 1, \rightarrow x < y, \rightarrow 1 < z, \rightarrow z < x, x \mapsto 0, y \mapsto 1 \rrbracket, \mathcal{C} \rangle \vdash C$

In the state above, the conflict was detected by noticing that we can not pick a value for z , because the trail contains $1 < z$, $z < x$ and $x \mapsto 0$. The **explain** procedure “generates” the explanation clause $C \equiv \neg(1 < z) \vee \neg(z < x) \vee 1 < x$ by eliminating z using Fourier-Motzkin elimination, and this clause evaluates to **false** in the current trail. Note that, as in DPLL(T), we could have also explained the infeasibility trail by producing the clause $\neg(1 < z) \vee \neg(z < x) \vee \neg(x < 1)$ that uses only literals that already exist in the trail. However, this clause is weaker since $\neg(x < 1) \equiv (1 \leq x)$ is weaker than $1 < x$. We continue the deduction by analyzing the conflict.

$\langle \llbracket \rightarrow x < 1, \rightarrow x < y, \rightarrow 1 < z, \rightarrow z < x, x \mapsto 0, y \mapsto 1 \rrbracket, \mathcal{C} \rangle \vdash C$
 \downarrow T-CONSUME (the conflict does not depend on y)
 $\langle \llbracket \rightarrow x < 1, \rightarrow x < y, \rightarrow 1 < z, \rightarrow z < x, x \mapsto 0 \rrbracket, \mathcal{C} \rangle \vdash C$
 \downarrow BACKJUMP (after backtracking $x \mapsto 0$, the clause C implies $1 < x$)
 $\langle \llbracket \rightarrow x < 1, \rightarrow x < y, \rightarrow 1 < z, \rightarrow z < x, C \rightarrow 1 < x \rrbracket, \mathcal{C} \rangle$

After the application of the backjump rule, the newly asserted literal $1 < x$ is immediately in conflict with the literal $x < 1$ and we enter conflict resolution again, with the explanation obtained by Fourier-Motzkin elimination.

$$\langle \llbracket \neg x < 1, \neg x < y, \neg 1 < z, \neg z < x, C \rightarrow 1 < x \rrbracket, \mathcal{C} \rangle$$

$$\downarrow \text{T-CONFLICT } (1 < x \text{ and } x < 1 \text{ implies } \text{false})$$

$$\langle \llbracket \neg x < 1, \neg x < y, \neg 1 < z, \neg z < x, C \rightarrow 1 < x \rrbracket, \mathcal{C} \rangle \vdash \neg(1 < x) \vee \neg(x < 1)$$

$$\downarrow \text{RESOLVE } \times 3, \text{ CONSUME, RESOLVE, UNSAT}$$

unsat

Theorem 1. *Given a set of clauses \mathcal{C} , and assuming a finite basis explanation function `explain`, any derivation starting from the initial state $\langle \llbracket \cdot \rrbracket, \mathcal{C} \rangle$ will terminate either in a state **sat**, when \mathcal{C} is satisfiable, or in the **unsat** state. In the later case, the set of clauses \mathcal{C} is unsatisfiable.*

Proof. Assume we have a set of clauses \mathcal{C} , over the variables x_1, \dots, x_n , and a finite-basis explanation function `explain`. Starting from the initial state $\langle \llbracket \cdot \rrbracket, \mathcal{C} \rangle$, we claim that any derivation of the transition system (finite or infinite), satisfies the following properties

1. the only possible “sink states” are the **sat** and the **unsat** states;
2. all $\vdash C$ clauses are implied by the initial set of clauses \mathcal{C} ;
3. during conflict analysis $\text{value}(C, M) = \text{false}$ for the $\vdash C$ clauses;

Assuming termination, and the above properties, the statement can be proven easily. Since **sat** and **unsat** are the only sink states, the derivation will terminate in one of these states. Since the SAT rule is only applicable if the set of clauses \mathcal{C} is satisfied, we have that the original problem is indeed satisfiable. On the other hand, if we terminate in the **unsat** state, by the above properties, the conflicting clause **false** is implied by the initial \mathcal{C} . Given that **false** is implied by the original set of clauses, the initial clauses themselves must truly be unsatisfiable.

The first property in the list above is a fairly easy exercise in case analysis and induction, so we skip those and concentrate on the more interesting properties. Proving the properties of conflict analysis is also quite straightforward, via induction on the number of conflicts, and conflict analysis steps. Clearly, initially, we have that C evaluates to **false** (the precondition of the **CONFLICT** and **T-CONFLICT** rules), and is implied by \mathcal{C} by induction. Then, every new clause that we produce during conflict resolution is obtained by the Boolean resolve rule, which will produce a valid deduction. Additionally, since the clause we are resolving with is a proper explanation, it will have all literals except the one we are resolving evaluate to **false**. Therefore, the resolvent also evaluates to **false**. As we backtrack down the trail with the conflicting clause, by definition of `value` and the preconditions of the rules, the clause still remains **false**.

Now, let us prove that the system terminates. It is clear that the conflict analysis rules (always removing elements from the trail) always terminate in a finite number of steps, and return to the search rules (or the **unsat** state). For the sake of the argument, let us assume that there is a derivation that does not terminate, and therefore does not enter the **unsat** state. We can define a big-step transition relation $\longrightarrow_{\text{bs}}$ that covers a transition from a search state, applying one or more transitions in the conflict analysis rules, and returns to a search state.

By assumption, we have a finite-basis explanation function `explain`, so we can assume a set of literals \mathbb{B} from which all the clauses that we can see during the search are constructed. In order to keep progress of the search, we define a partial order $M_1 \prec M_2$ on trails. The trail contains three different kinds of element: model assignments decisions ($x \mapsto \alpha$), Boolean decisions (L) and propagations ($C \rightarrow L$). The basic idea is to consider that propagations are heavier than Boolean decisions that are heavier than model assignments. We capture that by defining a (weight) function w from trail elements into the set $\{1, 2, 3\}$.

$$w(C \rightarrow L) = 3, \quad w(L) = 2, \quad w(x \mapsto \alpha) = 1$$

We define the $M_1 \prec M_2$ using a lexicographical order based on the weights of the trail elements, i.e.

$$\begin{aligned} \llbracket \rrbracket &\prec M = \text{true if } M \neq \llbracket \rrbracket \\ M &\prec \llbracket \rrbracket = \text{false} \\ \llbracket a, M_1 \rrbracket &\prec \llbracket b, M_2 \rrbracket = w(a) < w(b) \vee (w(a) = w(b) \wedge M_1 \prec M_2) \end{aligned}$$

It is clear that $\llbracket \rrbracket$ is the minimal element, and any trail containing $|\mathbb{B}|$ propagations followed by n model assignments is maximal. It is easy to see that for all trails M and trail elements a we have that $M \prec \llbracket M, a \rrbracket$. Thus, any rule that adds a new element to the trail is essentially creating a “bigger” trail with respect to the partial order \prec . This simple observation covers most of our rules.

Now, we consider the big-step \rightarrow_{bs} transition from a state $\langle M_1, \mathcal{C}_1 \rangle$ into a state $\langle M_2, \mathcal{C}_2 \rangle$. If we return using `BACKJUMP`, then the trail M_1 is of the form $\llbracket M, L, N \rrbracket$ or $\llbracket M, x \mapsto \alpha, N \rrbracket$, and the trail M_2 is of the form $\llbracket M, C \rightarrow L' \rrbracket$. In both cases $M_1 \prec M_2$ because a $w(C \rightarrow L')$ is greater than $w(L)$ and $w(x \mapsto \alpha)$. Similarly, if we return using `T-BACKJUMP-DECIDE`, the trail M_1 is of the form $\llbracket M, x \mapsto \alpha, N \rrbracket$, and M_2 is of the form $\llbracket M, L \rrbracket$, once again $M_1 \prec M_2$ since $w(L) > w(x \mapsto \alpha)$.

Now, to justify the `FORGET` rule, we define a partial order $\langle M_1, \mathcal{C}_1 \rangle \leq \langle M_2, \mathcal{C}_2 \rangle$ as $M_1 \prec M_2 \vee (M_1 = M_2 \wedge |\mathcal{C}_1| > |\mathcal{C}_2|)$. Using this definition, we have that $\langle M, \mathcal{C} \rangle \leq \langle M, \mathcal{C} \setminus \{C\} \rangle$, and consequently the `FORGET` rule also produces a “bigger” state. The partial order \leq also have maximal elements $\langle M_{\text{max}}, \mathcal{C}_0 \rangle$, where M_{max} is a maximal element for \prec and \mathcal{C}_0 is the initial set of clauses. Since all rules are producing bigger states and we can not increase forever, the termination of the system follows. \square

3 Conclusion

We proposed a model-constructing satisfiability calculus (`mcSAT`) that encompasses all model-based decision procedures found in the SMT literature. The `mcSAT` framework extends `DPLL(T)` by allowing assignments of variables to concrete values and relaxing the restriction on creation of new literals. The model created during the search also takes part in explaining the conflicts, and the full model is readily available as a witness if the procedure reports a satisfiable answer. The new calculus also extends `nlsat`, proposed at [8], by removing unnecessary restrictions on the Boolean-level search rules – it allows efficient Boolean

constraint propagation found in state-of-the-art SAT solvers and incorporates theory propagation and conflict rules from the DPLL(T) framework. The `mcSAT` calculus allows SMT developers to combine existing and successful techniques from the DPLL(T) framework with the flexibility provided by the `nlsat` calculus. We also presented a correctness proof for our procedure that is much simpler than the one provided for `nlsat`.

In this article, we did not explore the theory combination problem. However, we believe `mcSAT` is the ideal framework for combining implementations for theories such as: arithmetic, bit-vectors, floating-point, uninterpreted functions, arrays and datatypes.

References

1. Barrett, C., Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Splitting on Demand in SAT Modulo Theories. In: Hermann, M., Voronkov, A. (eds.) LPAR 2006. LNCS (LNAI), vol. 4246, pp. 512–526. Springer, Heidelberg (2006)
2. Barrett, C., Sebastiani, R., Seshia, S.A., Tinelli, C.: Satisfiability modulo theories. In: Handbook of Satisfiability. IOS Press (2009)
3. Brummayer, R., Biere, A.: Lemmas on Demand for the Extensional Theory of Arrays. *Journal on Satisfiability, Boolean Modeling and Computation* 6 (2009)
4. Cotton, S.: Natural Domain SMT: A Preliminary Assessment. In: Chatterjee, K., Henzinger, T.A. (eds.) FORMATS 2010. LNCS, vol. 6246, pp. 77–91. Springer, Heidelberg (2010)
5. de Moura, L., Bjørner, N.: Model-based Theory Combination. In: Satisfiability Modulo Theories. ENTCS, vol. 198, pp. 37–49 (2008)
6. Haller, L., Griggio, A., Brain, M., Kroening, D.: Deciding floating-point logic with systematic abstraction. In: FMCAD (2012)
7. Jovanović, D., de Moura, L.: Cutting to the Chase Solving Linear Integer Arithmetic. In: Bjørner, N., Sofronie-Stokkermans, V. (eds.) CADE 2011. LNCS, vol. 6803, pp. 338–353. Springer, Heidelberg (2011)
8. Jovanović, D., de Moura, L.: Solving Non-linear Arithmetic. In: Gramlich, B., Miller, D., Sattler, U. (eds.) IJCAR 2012. LNCS, vol. 7364, pp. 339–354. Springer, Heidelberg (2012)
9. Korovin, K., Tsiskaridze, N., Voronkov, A.: Conflict Resolution. In: Gent, I.P. (ed.) CP 2009. LNCS, vol. 5732, pp. 509–523. Springer, Heidelberg (2009)
10. Krstić, S., Goel, A.: Architecting Solvers for SAT Modulo Theories: Nelson-Oppen with DPLL. In: Konev, B., Wolter, F. (eds.) FroCos 2007. LNCS (LNAI), vol. 4720, pp. 1–27. Springer, Heidelberg (2007)
11. Malik, S., Zhang, L.: Boolean satisfiability from theoretical hardness to practical success. *Communications of the ACM* 52(8), 76–82 (2009)
12. McMillan, K.L., Kuehlmann, A., Sagiv, M.: Generalizing DPLL to Richer Logics. In: Bouajjani, A., Maler, O. (eds.) CAV 2009. LNCS, vol. 5643, pp. 462–476. Springer, Heidelberg (2009)
13. Moskewicz, M.W., Madigan, C.F., Zhao, Y., Zhang, L., Malik, S.: Chaff: engineering an efficient SAT solver. In: DAC (2001)
14. Nieuwenhuis, R., Oliveras, A., Tinelli, C.: Solving SAT and SAT modulo theories: From an abstract Davis–Putnam–Logemann–Loveland procedure to DPLL(T). *Journal of the ACM* 53(6), 937–977 (2006)
15. Silva, J.P.M., Sakallah, K.A.: GRASP – a new search algorithm for satisfiability. In: ICCAD (1997)