15-414: Bug Catching: Automated Program Verification

Lecture Notes on SMT Solving: Nelson-Oppen

Ruben Martins

Carnegie Mellon University

Lecture 18
Tuesday, March 26, 2024

1 Introduction

In the previous lecture we studied decision procedures for two first-order theories: arrays and equality with uninterpreted functions (EUF). Both procedures assumed that the formula to be decided was in the conjunctive, quantifier-free fragment of either theory, which means that the procedures are unable to handle a formula with a disjunction, or a negation over other logical connectives.

In this lecture, we will see how to decide formulas with arbitrary logical structure (i.e., they need not be in a conjunctive fragment), in any first-order theory for which we have a decision procedure capable of handling conjunctive, quantifier-free formulas. In particular, we will show how we can solve formulas that combine multiple theories by using the Nelson-Oppen combination method and the DPLL(T) framework. [1]

Learning Goals

1. The Nelson-Oppen procedure extends the approach to combinations of theories, but they must be *stably infinite*, and in some cases *convex*.

2. DPLL(T) combines a conjunctive theory solver and DPLL to decide formulas in a given first-order theory. The first step is to over-approximate a formula using its Boolean abstraction.

[1] Lecture notes based on [BM07] and [KS16a].

2 Review: First-Order Theories

A first-order theory T is defined by the following components.

- It's signature Σ is a set of constant, function, and predicate symbols.
- It's set of axioms $\mathcal{A}$ is a set of closed first-order logic formulae in which only constant, function, and predicate symbols of Σ appear.

Definition 1 (T-valid)**.** A Σ-formula φ is valid in the theory T (T-valid), if *every interpretation I* that satisfies the axioms of T (i.e., $I \models A$ for every $A \in \mathcal{A}$) also satisfies φ (i.e., $I \models \varphi$).

Definition 2 (T-satisfiable)**.** Let T be a Σ-theory. A Σ-formula φ is T-satisfiable if there *exists an interpretation I* such that $I \models A$ and $I \models \varphi$.

Definition 3 (T-decidable)**.** A theory T is decidable if $T \models \varphi$ is decidable for every Σ-formula. That is, there exists an algorithm that always terminate with "yes" if φ is T-valid or with "no" if φ is T-invalid.

2.1 Example of Theories

Some theories that we will use throughout this lecture are:

- The theory of equality with uninterpreted functions (T_{E}).
- The theory of reals ($T_{\mathbb{R}}$).

The **theory of equality with uninterpreted functions** T_{E} is the simplest first-order theory. It's signature

$$\Sigma_{\mathsf{E}} : \{=, a, b, c, \ldots, f, g, h, \ldots, p, q, r, \ldots\}$$

consists of

- = (equality), a binary predicate;
- and all constant, function and predicate symbols.

The axioms of T_{E} are the following:

1. $\forall x. x = x$ (reflexivity)
2. $\forall x, y. x = y \rightarrow y = x$ (symmetry)
3. $\forall x, y, z. x = y \land y = z \rightarrow x = z$ (transitivity)
4. $\forall \bar{x}, \bar{y}. (\bigwedge_{i=1}^{n} x_i = y_i) \rightarrow f(\bar{x}) = f(\bar{y})$ (congruence)
5. $\forall \bar{x}, \bar{y}. (\bigwedge_{i=1}^{n} x_i = y_i) \rightarrow (p(\bar{x}) \leftrightarrow p(\bar{y}))$ (equivalence)

The **theory of reals** $T_{\mathbb{R}}$ has signature

$$\Sigma_{\mathbb{R}} : \{0, 1, +, -, \cdot, =, \geq\}$$

where

- 0 and 1 are constants;
- + (addition) and $\cdot$ (multiplication) are binary functions;
- - (negation) is a unary function;
- and = (equality) and $\geq$ (weak inequality) are binary predicates.

$T_{\mathbb{R}}$ has a complex axiomatization and we will not describe all its axioms here since they are not essential to the understanding of the Nelson-Oppen procedure and the DPLL(T) framework. We refer the inththerested student to [BM07] for a detailed reading on the axiomatization of the theory of reals.

3 Theory combination

Definition 4 (Theory combination)**.** Given two theories T_1 and T_2 with signatures Σ_1 and Σ_2, respectively, the theory combination $T_1 \oplus T_2$ is a $(\Sigma_1 \cup \Sigma_2)$-theory defined by the axiom set $T_1 \cup T_2$.

Definition 5 (The theory combination problem)**.** Let φ be a $\Sigma_1 \cup \Sigma_2$ formula. The theory combination problem is to decide whether φ is $T_1 \oplus T_2$-valid. Equivalently, the problem is to decide whether the following holds: $T_1 \oplus T_2 \models \varphi$.

Given a Σ-formula φ in T_{E} and a Σ-formula ψ in $T_{\mathbb{R}}$ can we check the satisfiability of $\varphi \cup \psi$ by checking the satisfiability of φ and ψ independently and combining the results? **No!** This is not a sound procedure for the theory combination problem. Consider the following counterexample:

$$\begin{aligned} \varphi &= f(x) \neq f(y) \\ \psi &= x + y = 0 \land x = 0 \end{aligned}$$

Both φ and ψ are satisfiable but φ implies that $x \neq y$ and ψ implies that $x = y$, therefore their combination is not satisfiable!

4 The Nelson-Oppen Combination Procedure

The Nelson-Oppen combination procedure solves the theory combination problem for theories T_1 and T_2 that comply with the following restrictions:

- Both theories T_1 and T_2 are quantifier-free (conjunctive) fragments.

- Equality (=) is the only symbol in the intersection of their signatures, i.e., $\Sigma_1 \cap \Sigma_2 = \{=\}$.
- Both theories are stably infinite.

Definition 6 (Stably infinite)**.** A theory T with signature Σ is stably infinite if, for every satisfiable Σ_T-formula φ, there is an interpretation that satisfies φ and that has a universe of infinite cardinality

Consider the theory $T_{a,b}$ with signature $\Sigma_T : \{a, b, =\}$ where both a and b are constants and with the following axiom:

- $\forall x.x = a \lor x = b$ (two)

Because of axiom (two), every interpretation I is such that the domain of I has at most two elements. Therefore, $T_{a,b}$ is not stably infinite. Note that most of the theories of interest for program verification are stably infinite, e.g. theory of equality of uninterpreted functions and theory of integers.

The Nelson-Oppen procedure for a formula φ that combines different theories consists of:

1. **Purification**: Purify φ into $F_1, \ldots, F_n$.
2. Apply the decision procedure for T_i to F_i. If there exists i such that F_i is unsatisfiable in T_i, then φ is unsatisfiable.
3. **Equality propagation**: If there exists i, j such that F_i T_i-implies an equality between variables of φ that is not T_j-implied by F_j, add this equality to F_j and go to step 2.
4. If all equalities have been propagated then the formula is satisfiable.

4.1 Purification and equality propagation

Purification is a satisfiability-preserving transformation of the formula, after which each atom is from a specific theory. In this case, we say that all the atoms are **pure**. More specifically, given a formula φ, purification generates an equisatisfiable formula φ' as follows:

1. Let $\varphi' := \varphi$.
2. For each "alien" subexpression ϕ in φ':
   - Replace ϕ with a new auxiliary variable a_ϕ
   - Constraint φ' with $a_\phi = \phi$.

Consider the following formula:

$$\varphi = f(x + g(y)) \leq g(a) + f(b)$$

This formula combines the theories T_{E} and $T_{\mathbb{R}}$. Below we show the purification of φ into φ' defined over $T_{\mathbb{R}}$ and φ'' defined over T_{E}

Purification	
φ' ($T_{\mathbb{R}}$)	φ'' (T_{E})
$u_4 = x + u_1 \land$ $u_5 \leq u_2 + u_3$	$u_1 = g(y) \land$ $u_2 = g(a) \land$ $u_3 = f(b) \land$ $u_5 = f(u_4)$

Observe that φ' only contains atoms from $T_{\mathbb{R}}$ and φ'' only contains atoms from T_{E}. A variable is shared if it occurs in both formulas and local otherwise. For example, $\{u_1, u_2, u_3, u_4, u_5\}$ are shared variables since they appear in both φ' and φ'' and variables $\{x, y, a, b\}$ are local to either φ' ($\{x\}$) or φ'' ($\{y, a, b\}$).

Consider another formula:

$$\phi = f(f(x) - f(y)) \neq f(z) \land x \leq y \land y + z \leq x \land 0 \leq z$$

We will show how to determine the satisfiability of ϕ with the Nelson-Oppen procedure. We start by doing purification and then perform equality propagation over the shared variables.

Purification	
ϕ' ($T_{\mathbb{R}}$)	ϕ'' (T_{E})
$x \leq y \land$ $y + z \leq x \land$ $0 \leq z \land$ $w = u - v$	$f(w) \neq f(z) \land$ $u = f(x) \land$ $v = f(y)$
Equality propagation	
$x = y \land$ $u = v \land$ $w = z$	$x = y \land$ $u = v \land$ $w = z \land$ unsat

Observe that $x \leq y$, $y + z \leq x$ and $0 \leq z$ implies that $x = y$ and $z = 0$. Therefore, we add $x = y$ to both formulas. Since $x = y$ this implies that $f(x) = f(y)$ and therefore $u = v$. Since $u = v$ and $w = u - v$ than this implies that $w = 0$ which means that $w = z$. However, if $w = z$ than $f(w) = f(z)$ but ϕ'' contains $f(w) \neq f(z)$. Hence, ϕ is unsatisfiable.

4.2 Convex theories

The Nelson-Oppen procedure described in the previous section is only valid for convex theories. Note that this procedure can be modified to handle non-convex theories but for simplification purposes, we omit that version.

Definition 7 (Convex theory). A Σ-theory T is convex if for every conjunctive Σ-formula φ:

$$\begin{aligned}&(\varphi \rightarrow \bigvee_{i=1}^{n} x_i = y_i) \text{ is } T\text{-valid for some finite } n > 1 \rightarrow \\ &(\varphi \rightarrow x_i = y_i) \text{ is } T\text{-valid for some i } \in \{1, \cdots, n\}\end{aligned}$$

where x_i, y_i, for $i \in \{1, \cdots, n\}$, are some variables.

In other words, in a convex theory T, if a formula T-implies a disjunction of equalities, it also T-implies at least one of these equalities separately.

An example of a nonconvex theory is the theory of integers ($T_{\mathbb{Z}}$). For instance, while

$$x_1 = 1 \wedge x_2 = 2 \wedge 1 \leq x_3 \wedge x_3 \leq 2 \rightarrow (x_3 = x_1 \vee x_3 = x_2)$$

holds, neither

$$x_1 = 1 \wedge x_2 = 2 \wedge 1 \leq x_3 \wedge x_3 \leq 2 \rightarrow x_3 = x_1$$

nor

$$x_1 = 1 \wedge x_2 = 2 \wedge 1 \leq x_3 \wedge x_3 \leq 2 \rightarrow x_3 = x_2$$

holds.

Consider the following formula defined over the theory of integers ($T_{\mathbb{Z}}$) and the theory of uninterpreted functions with equality (T_{E}):

$$\varphi = 1 \leq x \wedge x \leq 2 \wedge f(x) \neq f(1) \wedge f(x) \neq f(2)$$

We can see that this formula is unsatisfiable since x is either 1 or 2 but $f(x) \neq 1 \wedge f(x) \neq 2$ which means that x has to be different than 1 and 2. However, if we apply the Nelson-Oppen procedure described in the previous section we will **incorrectly** conclude that φ is satisfiable:

Purification	
φ' ($T_{\mathbb{Z}}$)	φ'' (T_{E})
$1 \leq x \wedge$ $x \leq 2 \wedge$ $z = 1$ $w = 2$	$f(x) \neq f(z)$ $f(x) \neq f(w)$
Equality propagation	
sat	sat

In practice, SMT solvers use an extended version of Nelson-Oppen that propagates implied disjunctions of equalities [KS16b, Chapter 10]. The details of this extension are beyond the scope of the lecture, but note that adding additional disjunctions to a formula will force algorithms like DPLL(T) to solve them by case-splitting, which can quickly become expensive. So, while it is possible to combine non-convex theories with others, one should be aware that doing so may make the solver's job intractible, and explore other options.

5 DPLL(T) framework

The Nelson-Oppen procedure allows us to solve conjunctive first-order theories. To handle disjunction, we could convert the formula to Disjunctive Normal Form (DNF). However, this conversion is usually too expensive and is not the most efficient way of solving disjunctive first-order theories. One of the strengths of the DPLL algorithm is its ability to handle disjunctions efficiently via unit propagation and clause-learning. We will now see how DPLL can be extended to account for first-order theories via the DPLL(T) framework. This approach is used in almost all modern SMT solvers.

The key idea behind this framework is to decompose the SMT problem into parts we can deal with efficiently:

- Use SAT solver to cope with the **Boolean structure** of the formula;
- Use dedicated conjunctive **theory solver** to decide satisfiability in the background theory.

5.1 Boolean abstraction

We define the Boolean abstraction of a Σ-formula φ recursively:

- <literal> ::= <atom>$_T$ | ¬ <atom>$_T$
- <formula> ::= <literal> $\mathcal{B}\,(l_T) \overset{\text{def}}{=} P_i$, where P_i is a fresh variable
- <formula> ::= ¬ <formula> $\mathcal{B}\,(\neg F) \overset{\text{def}}{=} \neg\mathcal{B}(F)$
- <formula> ::= <formula> ∧ <formula> $\mathcal{B}\,(F_1 \wedge F_2) \overset{\text{def}}{=} \mathcal{B}(F_1) \wedge \mathcal{B}(F_2)$
- <formula> ::= <formula> ∨ <formula> $\mathcal{B}\,(F_1 \vee F_2) \overset{\text{def}}{=} \mathcal{B}(F_1) \vee \mathcal{B}(F_2)$
- <formula> ::= <formula>→ <formula> $\mathcal{B}\,(F_1 \rightarrow F_2) \overset{\text{def}}{=} \mathcal{B}(F_1) \rightarrow \mathcal{B}(F_2)$
- <formula> ::= <formula>↔ <formula> $\mathcal{B}\,(F_1 \leftrightarrow F_2) \overset{\text{def}}{=} \mathcal{B}(F_1) \leftrightarrow \mathcal{B}(F_2)$

Given a Σ-formula φ:

$$\varphi : g(a) = c \wedge (f(g(a)) \neq f(c) \vee g(a) = d) \wedge c \neq d$$

The Boolean abstraction of φ is the following:

$$\begin{aligned}\mathcal{B}(F) &= \mathcal{B}(g(a) = c) \wedge \mathcal{B}(f(g(a)) \neq f(c) \vee g(a) = d) \wedge c \neq d) \\ &= \mathcal{B}(g(a) = c) \wedge \mathcal{B}(f(g(a)) \neq f(c) \vee g(a) = d)) \wedge \mathcal{B}(c \neq d) \\ &= \mathcal{B}(g(a) = c) \wedge \mathcal{B}(f(g(a)) \neq f(c)) \vee \mathcal{B}(g(a) = d) \wedge \mathcal{B}(c \neq d) \\ &= P_1 \wedge (\neg P_2 \vee P_3) \wedge \neg P_4\end{aligned}$$

Note that we can also define $\mathcal{B}^{-1}$ which maps from the Boolean variables back to the atoms in the original formula. For example $\mathcal{B}^{-}1(P_1 \wedge P_3 \wedge P_4)$ corresponds to the formula $g(a) = c \wedge g(a) = d \wedge c = d$.

We call $\mathcal{B}(\varphi)$ an abstraction of φ since it is an over-approximation of φ with respect to satisfiability. Observe the following properties of this over-approximation:

- If φ is satisfiable then $\mathcal{B}(\varphi)$ is also satisfiable;
- If $\mathcal{B}(\varphi)$ is satisfiable then φ is not necessarily satisfiable:

  $$\varphi : 1 \leq x \wedge x \leq 2 \wedge f(x) \neq f(1) \wedge f(x) \neq f(2)$$

  φ is unsatisfiable in the theory of integers ($T_{\mathbb{Z}}$) since x is either 1 or 2 but $f(x) \neq f(1) \wedge f(x) \neq f(2)$ implies that x must be different than 1 and 2. However, the Boolean abstraction $\mathcal{B}(\varphi) = P_1 \wedge P_2 \wedge P_3 \wedge P_4$ is satisfiable.
- If φ is unsatisfiable then $\mathcal{B}(\varphi)$ is not necessarily unsatisfiable:

  $$\varphi : 1 \leq x \wedge x \leq 2 \wedge f(x) \neq f(1) \wedge f(x) \neq f(2)$$

  The same example as for the previous case holds for this case as well. φ is unsatisfiable in the theory of integers ($T_{\mathbb{Z}}$) but $\mathcal{B}(\varphi)$ is satisfiable.
- If $\mathcal{B}(\varphi)$ is unsatisfiable then φ is also unsatisfiable.

Combining theory and SAT solvers. The Boolean abstraction provides us with a **lazy** way to solve SMT. In the next lecture, we will talk about the DPLL(T) algorithm in detail and see how it can be used to determine the satisfiability of a formula that combines multiple theories.

References

[BM07] Aaron R. Bradley and Zohar Manna. *The Calculus of Computation: Decision Procedures with Applications to Verification*. Springer-Verlag New York, Inc., Secaucus, NJ, USA, 2007.

[KS16a] Daniel Kroening and Ofer Strichman. *Decision Procedures - An Algorithmic Point of View*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2016.

[KS16b] Daniel Kroening and Ofer Strichman. *Decision Procedures: An Algorithmic Point of View*. Springer Publishing Company, Incorporated, 2 edition, 2016.