22c:295 Seminar in AI — Decision Procedures

Satisfiability Modulo Shostak Theories

Cesare Tinelli

tinelli@cs.uiowa.edu

The University of Iowa

Outline

- Decidability Modulo Theories
- The Shostak's Method

Sources:
Harrison, John. *Introduction to Logic and Automated Theorem Proving*. Unpublished manuscript. Used by permission.
Barrett, Clark. *Checking Validity of Quantifier-Free Formulas in Combinations of First-Order Theories*. PhD Dissertation. Stanford University, 2003.

The Decision Problem: Recap

- We are interested in proving the unsatisfi ability (or dually, validity) of fi rst-order formulas.
- The general decision problem is to provide a yes or no answer to any question of satisfi ability or validity.
- There is no decision procedure for arbitrary fi rst order formulas.
- However, we may be able to get a decision procedure in two special cases.
  - Restrict the syntax of the formula.
  - Restrict the models under consideration. For example, only check validity in models of some set T of axioms.

Satisfi ability Modulo Theories

We focus again on (un)satisfi ability in a specifi c theory.

We now consider a general method for a class of theories called *Shostak* theories.

Recall:
A formula φ is satisfi able if there exists a model M and a variable assignment s such that $\models_M \varphi[s]$.

$\Gamma \models \varphi$ means that for every model M and variable assignment s, if $\models_M \Gamma[s]$, then $\models_M \varphi[s]$.

Shostak's Method

Robert Shostak published a paper in 1984 which detailed a particular strategy for deciding validity of quantifier-free formulas in certain kinds of theories.

Unfortunately, the original procedure contained many errors and a number of papers have since been dedicated to correcting them.

We will look at a simplified version of Shostak's procedure which is easily proved correct, yet still contains most of the essential ideas introduced by the original paper.

Equations in Solved Form

A set $\mathcal{S}$ of equations is said to be in *solved form* iff the left-hand side of each equation in $\mathcal{S}$ is a variable which appears only once in $\mathcal{S}$.

We call the left-hand sides variables of a set in solved form *solitary* variables.

A set $\mathcal{S}$ of equations in solved form defines an idempotent substitution: the one which replaces each solitary variable with its corresponding right-hand side.

If X is an expression or set of expressions, we denote the result of applying this substitution to X by $\mathcal{S}(X)$.

Equations in Solved Form

An interesting property of equations in solved form is the following.

Solved Form Theorem If T is a theory with signature Σ and $\mathcal{S}$ is a set of Σ-equations in solved form, then $T \cup \mathcal{S} \models \varphi$ iff $T \models \mathcal{S}(\varphi)$.

Proof

Clearly, $T \cup \mathcal{S} \models \varphi$ iff $T \cup \mathcal{S} \models \mathcal{S}(\varphi)$.
Thus we only need to show that $T \cup \mathcal{S} \models \mathcal{S}(\varphi)$ iff $T \models \mathcal{S}(\varphi)$.
The "if" direction is trivial.
To show the other direction, assume that $T \cup \mathcal{S} \models \mathcal{S}(\varphi)$. Any model of T can be made to satisfy $T \cup \mathcal{S}$ by assigning any value to the non-solitary variables of $\mathcal{S}$, and then choosing the value of each solitary variable to match the value of its corresponding right-hand side.

(over)

Equations in Solved Form

Since none of the solitary variables occur anywhere else in $\mathcal{S}$ this assignment is well-defined and satisfies $\mathcal{S}$
By assumption then, this model and assignment also satisfy $\mathcal{S}(\varphi)$, but none of the solitary variables appear in $\mathcal{S}(\varphi)$, so the initial arbitrary assignment to non-solitary variables must be sufficient to satisfy $\mathcal{S}(\varphi)$.
Thus it must be the case that every model of T satisfies $\mathcal{S}(\varphi)$ with every variable assignment. $\square$

By setting φ to $\mathbf{F}$ (false), we obtain the following.

Corollary If T is a satisfiable theory with signature $\Sigma$and $\mathcal{S}$ is a set of Σ-equations in solved form, then $T \cup \mathcal{S}$ is satisfiable.

Shostak Theories

A consistent theory T with signature Σ is a *Shostak* theory if the following conditions hold.

1. Σ contains no predicate symbols.
2. T is *convex*, that is, for every conjunction φ of literals and set $x_1 \approx y_1, \ldots x_n \approx y_n$ of equations between variables, if $T \cup \varphi \models x_1 = y_1 \vee \cdots \vee x_n = y_n$, then $T \cup \varphi \models x_i \approx y_i$ for some $1 \leq i \leq n$.
3. T has a *canonizer* *canon*, a computable function from Σ-terms to Σ-terms, such that $T \models a \approx b$ iff $\mathit{canon}(a) = \mathit{canon}(b)$.

(over)

Shostak Theories

4. T has a *solver* *solve*, a computable function from Σ-equations to sets of formulas defined as follows:
   (a) If $T \models a \not\approx b$, then $\mathit{solve}(a \approx b) = \{\mathbf{F}\}$.
   (b) Otherwise, $\mathit{solve}(a \approx b)$ returns a set $\mathcal{S}$ of equations in solved form such that

   $$T \models (a \approx b) \leftrightarrow \exists \overline{w}. \mathcal{S}$$

   where $\overline{w}$ is the set of variables that appear in $\mathcal{S}$ but not in a or b.

Canonizer

The canonizer is used to determine whether a specific equality is entailed by a set of equations in solved form.

Theorem (canon) If $\mathcal{S}$ is a set of Σ-equations in solved form, then

$$T \cup \mathcal{S} \models a \approx b \text{ iff } \mathit{canon}(\mathcal{S}(a)) = \mathit{canon}(\mathcal{S}(b)).$$

Proof
By the **Solved Form Theorem**, $T \cup \mathcal{S} \models a \approx b$ iff $T \models \mathcal{S}(a) \approx \mathcal{S}(b)$. But $T \models \mathcal{S}(a) \approx \mathcal{S}(b)$ iff $\mathit{canon}(\mathcal{S}(a)) = \mathit{canon}(\mathcal{S}(b))$, by the definition of *canon*. □

Procedure Sh

The procedure below checks the satisfiability in T of a set Γ set of equalities and a set Δ of disequalities.

Sh$(\Gamma, \Delta, \mathit{canon}, \mathit{solve})$
1. $\mathcal{S}$:= $\emptyset$;
2. **while** $\Gamma \neq \emptyset$ **do begin**
3. Remove some equality $a \approx b$ from Γ;
4. a' := $\mathcal{S}(a)$; b' := $\mathcal{S}(b)$;
5. $\mathcal{S}'$:= $\mathit{solve}(a' \approx b')$;
6. **if** $\mathcal{S}' = \{\mathbf{F}\}$ **then return** *false*
7. **else** $\mathcal{S}$:= $\mathcal{S}'(\mathcal{S}) \cup \mathcal{S}'$;
8. **end**
9. **if** $\mathit{canon}(\mathcal{S}(a)) = \mathit{canon}(\mathcal{S}(b))$
   for some $a \not\approx b \in \Delta$ **then return** *false*
10. **else return** *true*

Correctness of Procedure Sh

Termination of the procedure is trivial since each step terminates and each time line 3 is executed the size of Γ is reduced.

The following five lemmas are needed before proving correctness.

Lemma 1 If T' is a theory, Γ and Θ are sets of formulas, and $\mathcal{S}$ is a set of equations in solved form, then for any formula φ,

$$T' \cup \Gamma \cup \Theta \cup \mathcal{S} \models \varphi \text{ iff } T' \cup \Gamma \cup \mathcal{S}(\Theta) \cup \mathcal{S} \models \varphi.$$

Proof Follows trivially from the fact that $\Theta \cup \mathcal{S}$ and $\mathcal{S}(\Theta) \cup \mathcal{S}$ are satisfied by exactly the same models and variable assignments. $\square$

Correctness of Procedure Sh

Lemma 2 If Γ is any set of formulas, then for any formula φ, and Σ-terms a and b,

$$T \cup \Gamma \cup \{a \approx b\} \models \varphi \text{ iff } T \cup \Gamma \cup \mathit{solve}(a \approx b) \models \varphi.$$

Proof

$\Rightarrow$: Given that $T \cup \Gamma \cup \{a \approx b\} \models \varphi$, suppose that $M \models_\rho T \cup \Gamma \cup \mathit{solve}(a \approx b)$.

It is easy to see from the definition of *solve* that $M \models_\rho a \approx b$ and hence by the hypothesis, $M \models_\rho \varphi$.

(over)

Correctness of Procedure Sh

Lemma 2 (cont.) If Γ is any set of formulas, then for any formula φ, and Σ-terms a and b,

$$T \cup \Gamma \cup \{a \approx b\} \models \varphi \text{ iff } T \cup \Gamma \cup \mathit{solve}(a \approx b) \models \varphi.$$

Proof

$\Leftarrow$: Given that $T \cup \Gamma \cup \mathit{solve}(a \approx b) \models \varphi$, suppose that $M \models_\rho T \cup \Gamma \cup \{a \approx b\}$.

Since $T \models (a \approx b) \leftrightarrow \exists \overline{w}.\ \mathit{solve}(a \approx b)$, there exists a modified assignment ρ^* which assigns values to all the variables in $\overline{w}$ and satisfies $\mathit{solve}(a \approx b)$ but is otherwise equivalent to ρ. Then, by the hypothesis, $M \models_{\rho^*} \varphi$.

But the variables in $\overline{w}$ are fresh variables, so they do not appear in φ, meaning that changing their values cannot affect whether φ is true. Thus, $M \models_\rho \varphi$.

$\square$

Correctness of Procedure Sh

Lemma 3 Let Γ, $\{a \approx b\}$, and $\mathcal{S}$ be sets of Σ-formulas, with $\mathcal{S}$ in solved form. If $\mathcal{S}' = \mathit{solve}(\mathcal{S}(a \approx b))$ and $\mathcal{S}' \neq \{\mathbf{F}\}$, then for every formula φ,

$$T \cup \Gamma \cup \{a \approx b\} \cup \mathcal{S} \models \varphi \text{ iff } T \cup \Gamma \cup \mathcal{S}' \cup \mathcal{S}'(\mathcal{S}) \models \varphi.$$

Proof

	$T \cup \Gamma \cup \{a \approx b\} \cup \mathcal{S} \models \varphi$	
iff	$T \cup \Gamma \cup \{\mathcal{S}(a \approx b)\} \cup \mathcal{S} \models \varphi$	by **Lemma 1**
iff	$T \cup \Gamma \cup \mathcal{S}' \cup \mathcal{S} \models \varphi$	by **Lemma 2**
iff	$T \cup \Gamma \cup \mathcal{S}' \cup \mathcal{S}'(\mathcal{S}) \models \varphi$	by **Lemma 1**

$\square$

Correctness of Procedure Sh

Lemma 4 During the execution of Procedure Sh, $\mathcal{S}$ is always in solved form.

Proof Clearly, $\mathcal{S}$ is in solved form initially. Consider one iteration. By construction, a' and b' do not contain any of the solitary variables of $\mathcal{S}$, and thus by the defi nition of *solve*, $\mathcal{S}'$ doesn't either. Furthermore, if $\mathcal{S}' = \{\mathbf{F}\}$ then the procedure terminates at line 6. Thus, at line 7, $\mathcal{S}'$ must be in solved form. Applying $\mathcal{S}'$ to $\mathcal{S}$ guarantees that none of the solitary variables of $\mathcal{S}'$ appear in $\mathcal{S}$, so the new value of $\mathcal{S}$ is also in solved form. □

Correctness of Procedure Sh

Lemma 5 Let Γ_n and $\mathcal{S}_n$ be the values of Γ and $\mathcal{S}$ after the while loop in Procedure Sh has been executed n times. Then for each n, and any formula φ, the following invariant holds:

$$T \cup \Gamma_0 \models \varphi \text{ iff } T \cup \Gamma_n \cup \mathcal{S}_n \models \varphi.$$

Proof The proof is by induction on n. For $n = 0$, the invariant holds trivially. Now suppose the invariant holds for some $k \geq 0$. Consider the next iteration.

	$T \cup \Gamma_0 \models \varphi$	
iff	$T \cup \Gamma_k \cup \mathcal{S}_k \models \varphi$	by Induction Hypothesis
iff	$T \cup \Gamma_{k+1} \cup \{a \approx b\} \cup \mathcal{S}_k \models \varphi$	by Line 3
iff	$T \cup \Gamma_{k+1} \cup \mathcal{S}' \cup \mathcal{S}'(\mathcal{S}_k) \models \varphi$	by **Lemmas 3** and **4**
iff	$T \cup \Gamma_{k+1} \cup \mathcal{S}_{k+1} \models \varphi$	by Line 7

□

Correctness of Procedure Sh

Theorem Let T be a Shostak theory with signature Σ, canonizer *canon*, and solver *solve*. For all sets Γ of Σ-equalities and sets Δ of Σ-disequalities, $T \cup \Gamma \cup \Delta$ is satisfi able iff $\mathtt{Sh}(\Gamma, \Delta, \mathit{canon}, \mathit{solve}) = \mathit{true}$.

Proof
$\Rightarrow$: Suppose $\mathtt{Sh}(\Gamma, \Delta, \mathit{canon}, \mathit{solve}) \neq \mathit{true}$.
Since the procedure terminates for all inputs, it must be that $\mathtt{Sh}(\Gamma, \Delta, \mathit{canon}, \mathit{solve}) = \mathit{false}$.
If the procedure terminates at line 9, then $\mathit{canon}(\mathcal{S}(a)) = \mathit{canon}(\mathcal{S}(b))$ for some $a \not\approx b \in \Delta$.
It follows from the **canon** theorem and **Lemma 5** that $T \cup \Gamma \models a \approx b$, so clearly $T \cup \Gamma \cup \Delta$ is not satisfi able.
The other possibility when $\mathtt{Sh}(\Gamma, \Delta, \mathit{canon}, \mathit{solve}) = \mathit{false}$ is that the procedure terminates at line 6.

(over)

Correctness of Procedure Sh

Theorem (cont) [...] For all sets Γ of Σ-equalities and sets Δ of Σ-disequalities, $T \cup \Gamma \cup \Delta$ is satisfi able iff $\mathtt{Sh}(\Gamma, \Delta, \mathit{canon}, \mathit{solve}) = \mathit{true}$.

Proof (cont.)
Suppose the loop has been executed n times and that Γ_n and $\mathcal{S}_n$ are the values of Γ and $\mathcal{S}$ at the end of the last loop.
It must be the case that $T \models a' \not\approx b'$, so $T \cup \{a' \approx b'\}$ is unsatisfi able.
Clearly then, $T \cup \{a' \approx b'\} \cup \mathcal{S}_n$ is unsatisfi able, so by **Lemma 1**, $T \cup \{a \approx b\} \cup \mathcal{S}_n$ is unsatisfi able. But $\{a \approx b\}$ is a subset of Γ_n, so $T \cup \Gamma_n \cup \mathcal{S}_n$ must be unsatisfi able. Thus by **Lemma 5**, $T \cup \Gamma$ is unsatisfi able.

(over)

Correctness of Procedure Sh

Theorem (cont) [...] $T \cup \Gamma \cup \Delta$ is satisfi able iff $\mathtt{Sh}(\Gamma, \Delta, \mathit{canon}, \mathit{solve}) = \mathit{true}$.

Proof

$\Leftarrow$: Suppose that $\mathtt{Sh}(\Gamma, \Delta, \mathit{canon}, \mathit{solve}) = \mathit{true}$. Then the procedure terminates at line 10.

By **Lemma 4**, $\mathcal{S}$ is in solved form. Let $\overline{\Delta}$ be the disjunction of equalities equivalent to $\neg(\Delta)$.

Since the procedure does not terminate at line 9, $T \cup \mathcal{S}$ does not entail any equality in $\overline{\Delta}$. By the convexity of T, it follows that $T \cup \mathcal{S} \not\models \overline{\Delta}$.

Now, since $T \cup \mathcal{S}$ is satisfi able by the corollary to the **Solved Form Theorem**, it follows that $T \cup \mathcal{S} \cup \Delta$ is satisfi able.

But by **Lemma 5**, $T \cup \Gamma \models \varphi$ iff $T \cup \mathcal{S} \models \varphi$, so in particular $T \cup \mathcal{S} \models \Gamma$. Thus $T \cup \mathcal{S} \cup \Delta \cup \Gamma$ is satisfi able, and hence $T \cup \Gamma \cup \Delta$ is satisfi able. $\square$