Decision Procedures
An Algorithmic Point of View

Deciding Combined Theories

Daniel Kroening and Ofer Strichman

So far we know how to…

- Decide Equality Logic with Uninterpreted Functions:
  $(x_1 = x_2) \vee \neg(f(x_2) = x_3) \wedge \ldots$

- Decide Disjunctive Linear arithmetic:

  $3x_1 + 5x_2 \geq 2x_3 \wedge x_2 \leq 4x_4 \ldots$

- What about a combined formula ?
  $(x_2 \geq x_1) \wedge (x_1 - x_3 \geq x_2) \wedge (x_3 \geq 0) \wedge f(f(x_1) - f(x_2)) \neq f(x_3)$

We also know how to…

- Decide bit-vector equations

  $a[32] \times b[32] = b[32] \times a[32]$

- But how shall we decide

  $f(a[32], b[1]) = f(b[32], a[1]) \wedge a[32] = b[32]$

More combination examples:

- Combining lists, arithmetic and Uninterpreted Functions:

$$(x_1 \leq x_2) \wedge (x_2 \leq x_1 + car(cons(0,x_1))) \wedge p(h(x_1) - h(x_2)) \wedge \neg p(0)$$

- Combining Arrays and Arithmetic:

$$x = \text{store}(v,i,e)[j] \wedge y = v[j] \wedge x > e \wedge x > y$$

Combining theories

- Approach #1: Reduce all theories to a common logic, if possible (e.g. Propositional Logic).
  - All un-quantified theories we saw so far are in NP.
  - We saw their direct translation to SAT (i.e. not through a Turing-machine).
- Approach #2: Combine decision procedures of the individual theories.
  - How? we will learn the Nelson-Oppen method*

* Greg Nelson and Derek Oppen, *simplification by cooperating decision procedures*, 1979

Reminders: theories and signatures

- First order logic –
  - Symbols (Boolean connectives and quantifiers over variables), Syntax (wff-s).
  - Axioms, inference rules.
- First order theories –
  - Additional axioms and symbols characterizing the theory.
  - The signature Σ of a theory $\mathcal{T}$ holds the set of functions and predicates of the theory.
- "First order quantifier-free theories with equality" – the equality predicate must be part of the signature.

The Theory-Combination problem

- Given theories $\mathcal{T}_1$ and $\mathcal{T}_2$ with signatures Σ_1 and Σ_2, the combined theory $\mathcal{T}_1 \oplus \mathcal{T}_2$
  - has signature $\Sigma_1 \cup \Sigma_2$ and
  - the union of their axioms.
- Let ϕ be a $\Sigma_1 \cup \Sigma_2$ formula.
- The problem: Does $\mathcal{T}_1 \oplus \mathcal{T}_2 \models \phi$?

The problem

- The Theory-Combination problem is undecidable (even when the individual theories are decidable).
- Under certain restrictions, it becomes decidable.
- We will assume the following restrictions:
  - $\mathcal{T}_1$ and $\mathcal{T}_2$ are decidable, quantifier-free first-order theories with equality.
  - Disjoint signatures (other than equality): $\Sigma_1 \cap \Sigma_2 = \emptyset$
  - More restrictions to follow…
- There are extensions to the basic algorithm that we will study, that partially overcomes each of these restrictions.

The Nelson-Oppen method (1)

- Purification: validity-preserving transformation of the formula after which predicates from different theories are not mixed.

1. Replace an `alien' sub-expression ϕ with a new auxiliary variable a
2. Constrain the formula with $a = \phi$

Transform $x_1 \leq f(x_1)$

… into $\underbrace{x_1 \leq a_1 \wedge a_1 = f(x_1)}$

Pure expressions, shared variables

Uninterpreted Functions

Arithmetic

The Nelson-Oppen method (2)

- After purification we are left with several sets of pure expressions $F_1 \ldots F_n$ such that:
  - F_i belongs to some 'pure' theory which we can decide.
  - Shared variables are allowed, i.e. it is possible that for some i,j, $vars(F_i) \cap vars(F_j) \neq \emptyset$.
  - ϕ is satisfiable $\leftrightarrow F_1 \wedge \ldots \wedge F_n$ is satisfiable

The Nelson-Oppen method* (3)

1. Purify ϕ into $F_1 \wedge \ldots \wedge F_n$.
2. If $\exists i.\ F_i$ is unsatisfiable, return `unsatisfiable' .
3. If $\exists i,j.\ F_i$ implies an equality not implied by F_j, add it to F_j and goto step 2.
4. Return `satisfiable'.

* *So far only for 'non-convex' theories – to be explained*

Example (1)

$$(x_1 \leq x_2) \land (x_2 \leq (x_1 + car(cons(0, x_1)))) \land p(h(x_1) - h(x_2)) \land \neg p(0)$$

- Purification:

$$(x_1 \leq x_2) \land (x_2 \leq x_1 + a_1) \land p(a_2) \land \neg p(a_5) \land$$

$$a_1 = car(cons(a_5, x_1)) \land$$

$$a_2 = a_3 - a_4 \quad \land$$

$$a_3 = h(x_1) \quad \land$$

$$a_4 = h(x_2) \quad \land$$

$$a_5 = 0$$

Example (1), cont'd

Arithmetic	EUF	Lists
$x_1 \leq x_2$	$a_3 = h(x_1)$	$a_1 = car(cons(a_5, x_1))$
$x_2 \leq x_1 + a_1$	$a_4 = h(x_2)$	
$a_2 = a_3 - a_4$	$p(a_2)$	
$a_5 = 0$	$\neg p(a_5)$	
$a_1 = a_5$	$a_1 = a_5$	$\boxed{a_1 = a_5}$
$\boxed{x_1 = x_2}$	$x_1 = x_2$	$x_1 = x_2$
$a_3 = a_4$	$\boxed{a_3 = a_4}$	$a_3 = a_4$
$\boxed{a_2 = a_5}$	$a_2 = a_5$	$a_2 = a_5$
	$\boxed{False}$	

Example(2)

$$(x_2 \geq x_1) \wedge (x_1 - x_3 \geq x_2) \wedge (x_3 \geq 0) \wedge f(f(x_1) - f(x_2)) \neq f(x_3)$$

- Purification:

$$(x_2 \geq x_1) \wedge (x_1 - x_3 \geq x_2) \wedge (x_3 \geq 0) \wedge f(a_1) \neq f(x_3) \wedge$$

$$a_1 = a_2 - a_3 \wedge$$

$$a_2 = f(x_1) \quad \wedge$$

$$a_3 = f(x_2)$$

Example (2) – cont'd

Arithmetic	EUF
$x_2 \geq x_1$	$f(a_1) \neq f(x_3)$
$x_1 - x_3 \geq x_2$	$a_2 = f(x_1)$
$x_3 \geq 0$	$a_3 = f(x_2)$
$a_1 = a_2 - a_3$	
$\boxed{x_3 = 0}$	$x_3 = 0$
$\boxed{x_1 = x_2}$	$x_1 = x_2$
$a_2 = a_3$	$\boxed{a_2 = a_3}$
$\boxed{a_1 = 0}$	$a_1 = 0$
	$\boxed{False}$

Wait, it's not so simple…

- Consider: ϕ: $1 \leq x \wedge x \leq 2 \wedge p(x) \wedge \neg p(1) \wedge \neg p(2\ )$

$x \in \mathbb{Z}$

Arithmetic over $\mathbb{Z}$	Uninterpreted predicates
$1 \ \leq \ x$ $x \leq 2$	$p(x)$ $\neg p(1)$ $\neg p(2)$

- Neither theories imply an equality, and both are satisfiable.
- But ϕ is unsatisfiable!

Some theories have it, some don't

- Definition: A theory $\mathcal{T}$ is *convex* if for all conjunctions ϕ it holds that

  $\phi \rightarrow \bigvee_{i=1..n} x_i = y_i$ for some $n > 1 \Leftrightarrow$

  $\phi \rightarrow x_i = y_i$ for some $i \in \{1..n\}$

  where x_i, y_i are some $\mathcal{T}$ variables.

- *Convex*: Linear Arithmetic over $\mathbb{R}$, EUF
- *Non-convex*: Almost anything else…

Convexity: examples

- Linear arithmetic over $\mathbb{R}$ is convex

  ϕ: $x_1 \leq 1 \wedge x_1 \geq 0$ implies an infinite disjunction of equalities,

  ϕ: $x_1 \leq 1 \wedge x_1 \geq 1 \rightarrow x_1 = 1$ implies a singleton

  ϕ: $x_1 \leq 1 \wedge x_1 \geq 2$ implies everything

- Linear arithmetic over $\mathbb{Z}$ is *not* convex

  ϕ: $1 \leq x_1 \wedge x_1 \leq 2$

  Although $\phi \rightarrow (x_1 = 1 \vee x_1 = 2)$

  It is *not* the case that $\phi \rightarrow x_1 = 1 \vee \phi \rightarrow x_1 = 2$

So why is convexity important ?

- Recall: $\phi: 1 \leq x \wedge x \leq 2 \wedge p(x) \wedge \neg p(1) \wedge \neg p(2)$

$x \in \mathcal{Z}$

Arithmetic over $\mathbb{Z}$	Uninterpreted predicates
$1 \leq x$ $x \leq 2$	$p(x)$ $\neg p(1)$ $\neg p(2)$

- Neither theories imply an equality, and both are satisfiable.

So why is convexity important ? (cont'd)

- But: $1 \leq x \wedge x \leq 2$ imply the disjunction $x = 1 \vee x = 2$
- Since the theory is non-convex we cannot propagate either x=1 or x=2.
- We can only propagate the disjunction itself.

So why is convexity important ? (cont'd)

- Propagate the disjunction and perform case-splitting.

<table>
<tr><th>Arithmetic over $\mathbb{Z}$</th><th colspan="2">Uninterpreted predicates</th></tr>
<tr><td>$1 \leq x$
$x \leq 2$
$x = 1 \vee x = 2$</td><td colspan="2">$p(x)$
$\neg p(1) \wedge \neg p(2)$
$x = 1 \vee x = 2$ Split!</td></tr>
<tr><td></td><td>$\langle \cdot \rangle \wedge x = 1$
$False$</td><td>$\langle \cdot \rangle \wedge x = 2$
$False$</td></tr>
</table>

So why is convexity important? (cont'd)

- Conclusion: when the theory is non-convex, we must case-split.
- This adds a splitting step in Nelson-Oppen.
- As a result:
  - Convex theories: Polynomial
  - Non-Convex theories: Exponential

The (full) Nelson-Oppen method

1. Purify ϕ into ϕ': $F_1 \wedge \ldots \wedge F_n$.
2. If $\exists i.\ F_i$ is unsatisfiable, return `unsatisfiable' .
3. If $\exists i, j.\ F_i$ implies an equality not implied by F_j, add it to F_j and goto step 2.
4. If $\exists i.\ F_i \rightarrow (x_1 = y_1 \vee \ldots \vee x_k = y_k)$ but $\forall j\ F_i \nrightarrow x_j = y_j$, apply recursively to $\phi' \wedge x_1 = y_1, \ldots, \phi' \wedge x_k = y_k$. If any of them is satisfiable, return 'satisfiable'. Otherwise return 'unsatisfiable'.
5. Return `satisfiable'.

Correctness is hard to prove…

- Theorem: N.O. returns unsatisfiable if and only if its input formula ϕ is unsatisfiable.

- We will prove this theorem for the case of combining two convex theories. The generalization is not hard. The proof is based on [NO79].

Correctness is hard to prove…

- (➔) N.O. returns 'unsatisfiable' $\rightarrow$ ϕ is unsatisfiable.
  (That's the simple side)
  - Assume ϕ is satisfiable and let α be a satisfying assignment of ϕ.
  - Let A = $\{a_1,\ldots,a_n\}$ be the purification (auxiliary) variables.
  - *Claim*: there exists an assignment to the A variables such that α extended with this assignment satisfies $F_1 \wedge F_2$.
  - Let α' be this extended assignment.
  - For each equality eq added in line 3, $\exists i.\ F_i \rightarrow eq$.
  - Since α' ⊨ F_i then also α'⊨ eq.
  - Hence for all $j \in \{1,2\}$, α' ⊨ $F_j \wedge eq$.
  - Thus, N.O. *does not* return unsat in this case.
  - In other words, if N.O. returns unsat, then ϕ is unsat.

Proof (←)

- (←) If N.O. returns 'satisfiable', ϕ is satisfiable. (This will require several definition and lemmas)

- Dfn: A residue of a formula ϕ, denoted Res(ϕ), is the strongest Equality Logic formula implied by ϕ.

Res($x = f(a) \wedge y = f(b)$) is $a = b \rightarrow x = y$

Res($x \leq y \wedge y \leq x$) is $x = y$

- Lemma 1: For any formula F, there exists a formula Res(F) (we will skip the proof of this Lemma)

Proof (←)

- Recall: the Logical symbols of a formula are those shared by all first-order theories. We consider `=‘ as a logical symbol. The Non-logical symbols are theory-specific.

- Dfn: The parameters of a formula ϕ, denoted $param(\phi)$, are the non-logical symbols in ϕ.
- Craig's Interpolation Lemma: if A and B are formulas such that $A \rightarrow B$, then there exists a formula H such that $A \rightarrow H$ and $H \rightarrow B$, and $param(H) \subseteq param(A) \cap param(B)$.

Proof (←)

- Lemma 2: if F_1 and F_2 are formulas with disjoint signatures, $\text{Res}(F_1 \wedge F_2) \leftrightarrow (\text{Res}(F_1) \wedge \text{Res}(F_2))$.
- Proof: (→)
  - $F_1 \rightarrow \text{Res}(F_1)$, $F_2 \rightarrow \text{Res}(F_2)$,
  - $F_1 \wedge F_2 \rightarrow \text{Res}(F_1) \wedge \text{Res}(F_2)$
  - $\text{Res}(F_1 \wedge F_2) \rightarrow \text{Res}(F_1) \wedge \text{Res}(F_2)$ // *

* The consequence (RHS) is Equality Logic, hence it is implied by the residue of the Antecedent (LHS).

Proof of Lemma 2 (🡐)

(1) ■ $F_1 \wedge F_2 \rightarrow \text{Res}(F_1 \wedge F_2)$

(2) □ $F_1 \rightarrow (F_2 \rightarrow \text{Res}(F_1 \wedge F_2))$

■ There exists an interpolant H such that

(3) $(F_1 \rightarrow H) \wedge (H \rightarrow (F_2 \rightarrow \text{Res}(F_1 \wedge F_2)))$

Can be rewritten as

(4) $(\text{Res}(F_1) \rightarrow H) \wedge (H \rightarrow (F_2 \rightarrow \text{Res}(F_1 \wedge F_2)))$

because H is an Equality Logic formula, and hence everything implied by F_1 is also implied by $\text{Res}(F_1)$.

Why is H an Equality Logic formula? because
param(RES($F_1 \wedge F_2$)) = {} //Equality Logic formula
and *param*(F_1) ∩ *param*(F_2) = {}

Proof of Lemma 2 (←)

(4)
- $(\text{Res}(F_1) \rightarrow H) \wedge (H \rightarrow (F_2 \rightarrow \text{Res}(F_1 \wedge F_2)))$
- Since $\text{Res}(F_1 \wedge F_2)$ is also an Equality Logic formula:

(5) $(\text{Res}(F_1) \rightarrow \text{H}) \wedge (\text{H} \rightarrow (\text{Res}(\text{F}_2) \rightarrow \text{Res}(F_1 \wedge F_2)))$

which implies

(6) $(\text{Res}(F_1) \rightarrow (\text{Res}(F_2) \rightarrow \text{Res}(F_1 \wedge F_2)))$

and hence

(7) $(\text{Res}(F_1) \wedge \text{Res}(F_2)) \rightarrow \text{Res}(F_1 \wedge F_2)$

- q.e.d (Lemma 2):

$$\text{Res}(F_1) \wedge \text{Res}(F_2) \leftrightarrow \text{Res}(F_1 \wedge F_2)$$

Lemma 3

- Lemma 3:
  - Let F_1 and F_2 be satisfiable Equality Logic formulas s.t.
    - $V = vars(F_1) \cup vars(F_2)$.
    - $\forall x,y \in V,\ (F_1 \rightarrow x=y \wedge F_2 \rightarrow x=y)$ or $(F_1 \nrightarrow x=y \wedge F_2 \nrightarrow x=y)$
  - Then, $F_1 \wedge F_2$ is satisfiable.
- Proof: Let
  - S = the set of all equalities implied by both F_1 and F_2
  - T = the rest of the possible equalities in V.
  - α = an assignment s.t. $\forall eq \in S .\ \alpha \models eq, \forall eq \in T .\ \alpha \not\models eq$
  - Claim: $\alpha \models F_1 \wedge F_2$

Proof of Lemma 3

- Falsely assume that $\alpha \not\models F_1$
- Then, $(F_1 \rightarrow \vee_{eq \in T} eq)$
  - (Can it be, alternatively, that F_1 implies a negation of one of the equalities in S ? no, because it implies $\wedge_{eq \in S}$ eq

- If T is empty, F_1 is false (*contradiction*)
- If $\exists eq \in T.\ F_1 \rightarrow eq$, then $eq \in S$ (*contradiction*)
- Otherwise, F_1 is non-convex (*contradiction*)
- q.e.d (Lemma 3)

Proof (←)

- Now suppose N.O. returns SAT although $F_1 \wedge F_2$ is unsatisfiable.
- $\text{Res}(F_1 \wedge F_2) = \text{false}$
- Hence, by Lemma 2, $\text{Res}(F_1) \wedge \text{Res}(F_2) = \text{false}$

Proof (←)

- On the other hand, in step 4, where we return 'Satisfiable', we know that
  - F_1 and F_2 are separately satisfiable
  - F_1 and F_2 imply exactly the same equalities.
  - Thus, $\text{Res}(F_1)$ and $\text{Res}(F_2)$ are satisfiable and imply the same equalities.
- Hence, according to Lemma 3, $\text{Res}(F_1) \wedge \text{Res}(F_2)$ is also satisfiable, i.e. $\text{Res}(F_1) \wedge \text{Res}(F_2) \neq \text{false}$ (contradiction).
- Q.E.D (N.O.)

More problems…

- *Definition*: A Σ-theory $\mathcal{T}$ is *Stably-infinite* if for every quantifier-free Σ-formula ϕ
  ϕ is satisfiable $\Leftrightarrow$
  ϕ can be satisfied by an interpretation with an infinite domain.

- Specifically, this means that no theory with a finite domain is stably infinite.

Problem: non-stably infinite theories

- Consider a theory $\mathcal{T}_1$:
  - Σ_1: A function f,
  - Axioms that only allow solutions with 2 distinct values.
- And a theory $\mathcal{T}_2$:
  - Σ_2: A function g,
  - Domain: $\mathbb{N}$

Recall that the combined theory $\mathcal{T}_1 \oplus \mathcal{T}_2$ has the union of the axioms.
Hence the solution to any formula $\phi \in \mathcal{T}_1 \oplus \mathcal{T}_2$ cannot have more than 2 distinct values.

So this formula is unsatisfiable:

$$\phi: \; f(x_1) \neq f(x_2) \;\wedge\; g(x_1) \neq g(x_3) \;\wedge\; g(x_2) \neq g(x_3)$$

Problem: non-stably infinite theories

ϕ: $f(x_1) \neq f(x_2) \;\wedge\; g(x_1) \neq g(x_3) \;\wedge\; g(x_2) \neq g(x_3)$

$\mathcal{T}_1$	$\mathcal{T}_2$
$f(x_1) \neq f(x_2)$	$g(x_1) \neq g(x_3)$ $g(x_2) \neq g(x_3)$

No equalities to propagate: Satisfiable !

Solution to non-stable infinite theories

- Nelson-Oppen method cannot be used.
- Recently a solution to this problem was suggested by Tinelli & Zarba [TZ05]
  - Assuming all combined theories are stably-finite (in particular, it has a small model property), it computes, if possible, the upper bound on the minimal satisfying assignment, and propagates this information between the theories.