The Nelson–Oppen Algorithm (Deterministic Version for Convex Theories)

Unsat:

$$\frac{F_1, F_2}{\bot}$$

if $\exists\vec{x}\, F_i$ is unsatisfiable w. r. t. $\mathcal{T}_i$ for some i.

Propagate:

$$\frac{F_1, F_2}{F_1 \land (x \approx y), F_2 \land (x \approx y)}$$

if x and y are two different variables appearing in
both F_1 and F_2 such that
$\mathcal{T}_1 \models \forall\vec{x}\,(F_1 \rightarrow x \approx y)$ and $\mathcal{T}_2 \not\models \forall\vec{x}\,(F_2 \rightarrow x \approx y)$
or $\mathcal{T}_2 \models \forall\vec{x}\,(F_2 \rightarrow x \approx y)$ and $\mathcal{T}_1 \not\models \forall\vec{x}\,(F_1 \rightarrow x \approx y)$.

Theorem 1.8 *If $\mathcal{T}_1$ and $\mathcal{T}_2$ are signature-disjoint theories that are convex w. r. t. equations and have no trivial models, then the deterministic Nelson–Oppen algorithm is terminating, sound and complete for deciding satisfiability of pure conjunctions of literals F_1 and F_2 over $\mathcal{T}_1 \cup \mathcal{T}_2$.*

Proof. Termination and soundness are obvious: there are only finitely many different equations that can be added, and each of them is entailed by given formulas.

For completeness, we have to show that every configuration that is irreducible by "Unsat" and "Propagate" is satisfiable w. r. t.. $\mathcal{T}_1 \cup \mathcal{T}_2$: Let F_1, F_2 be such a configuration. As it is irreducible by "Propagate", we have, for every equation $x \approx y$ between shared variables, $\mathcal{T}_1 \models \forall\vec{x}\,(F_1 \rightarrow x \approx y)$ if and only if $\mathcal{T}_2 \models \forall\vec{x}\,(F_2 \rightarrow x \approx y)$. Consequently, F_1 and F_2 are compatible with the same equivalence on the shared variables of F_1 and F_2. Moreover, each of the formulas F_i is $\mathcal{T}_i$-satisfiable, and since convexity implies stable infiniteness, F_i has a $\mathcal{T}_i$-model with a countably infinite universe. Hence, by the amalgamation lemma, $F_1 \land F_2$ is $(\mathcal{T}_1 \cup \mathcal{T}_2)$-satisfiable.

Corollary 1.9 *The deterministic Nelson–Oppen algorithm for convex theories requires at most $O(n^3)$ calls to the individual decision procedures for the component theories, where n is the number of shared variables.*

Iterating Nelson–Oppen

The Nelson–Oppen combination procedures can be iterated to work with more than two component theories by virtue of the following observations where signature disjointness is assumed:

Theorem 1.10 *If $\mathcal{T}_1$ and $\mathcal{T}_2$ are stably infinite, then so is $\mathcal{T}_1 \cup \mathcal{T}_2$.*

Proof. The non-deterministic Nelson–Oppen algorithm is sound and complete for $\mathcal{T}_1 \cup \mathcal{T}_2$, that is, an existentially quantified conjunction F over $\Sigma_1 \cup \Sigma_2$ is satisfiable if and only if in every derivation from the purified form of F there exists a branch leading to some irreducible constraint F_1, F_2 entailing F. The amalgamation lemma 1.4 constructs a model of cardinality ω for F from the models of F_1 and F_2.

Lemma 1.11 *A first-order theory $\mathcal{T}$ is convex w. r. t. equations if and ony if for every conjunction Γ of Σ-equations and non-equational Σ-literals and for all equations $x_i \approx x_i'$ $(1 \leq i \leq n)$, whenever $\mathcal{T} \models \forall\vec{x}\,(\Gamma \rightarrow x_1 \approx x_1' \vee \ldots \vee x_n \approx x_n')$, then there exists some index j such that $\mathcal{T} \models \forall\vec{x}\,(\Gamma \rightarrow x_j \approx x_j')$.*

Lemma 1.12 *Let $\mathcal{T}$ be a first-order theory that is convex w. r. t. equations and and has no trivial models. Let F is a conjunction of literals; let F^- be the conjunction of all negative equational literals in F and let F^+ be the conjunction of all remaining literals in F. If $\mathcal{T} \models \forall\vec{x}\,(F \rightarrow x \approx y)$, then $\exists\vec{x}\,F$ is $\mathcal{T}$-unsatisfiable or $\mathcal{T} \models \forall\vec{x}\,(F^+ \rightarrow x \approx y)$.*

Proof. $\mathcal{T} \models \forall\vec{x}\,(F \rightarrow x \approx y)$ is equivalent to $\mathcal{T} \models \forall\vec{x}\,(F^+ \rightarrow (\neg F^- \vee x \approx y))$. By convexity of $\mathcal{T}$ we know that $\mathcal{T} \models \forall\vec{x}\,(F^+ \rightarrow x \approx y)$ or $\mathcal{T} \models \forall\vec{x}\,(F^+ \rightarrow A)$ for some literal $\neg A$ in F^-. In the latter case, $\exists\vec{x}\,(F^+ \wedge \neg A)$ is $\mathcal{T}$-unsatisfiable; hence $\exists\vec{x}\,F$, that is, $\exists\vec{x}\,(F^+ \wedge F^-)$ is $\mathcal{T}$-unsatisfiable as well.

Theorem 1.13 *If $\mathcal{T}_1$ and $\mathcal{T}_2$ are convex w. r. t. equations and do not have trivial models, then so is $\mathcal{T}_1 \cup \mathcal{T}_2$.*

Proof. Suppose that $\mathcal{T}_1$ and $\mathcal{T}_2$ are convex w. r. t. equations and do not have trivial models. Assume furthermore that $\mathcal{T} \models \forall\vec{x}\,(\Gamma \rightarrow x_1 \approx x_1' \vee \ldots \vee x_n \approx x_n')$ for some conjunction Γ of $(\Sigma_1 \cup \Sigma_2)$-literals. Then $\exists\vec{x}\,(\Gamma \wedge x_1 \not\approx x_1' \wedge \ldots \wedge x_n \not\approx x_n')$ is $\mathcal{T}$-unsatisfiable, and we can detect this by some run of the deterministic Nelson–Oppen algorithm starting with $\exists\vec{x},\vec{y}\,(\Gamma_1 \wedge \Gamma_2 \wedge x_1 \not\approx x_1' \wedge \ldots \wedge x_n \not\approx x_n')$, where $\Gamma_1 \wedge \Gamma_2$ is the result of purifying Γ. This run consists of a sequence of "Propagate" steps followed by a final "Unsat" step, and without loss of generality, we use the "Propagate" rule only if "Unsat" cannot be applied. Consequently, whenever we add an equation $x \approx y$ that is entailed by F_1 w. r. t. $\mathcal{T}_1$ or by F_2 w. r. t. $\mathcal{T}_2$, then it is already entailed by the positive and the non-equational literals in F_1 or F_2. Furthermore, due to the convexity of $\mathcal{T}_1$ and $\mathcal{T}_2$, the final "Unsat" step depends on at most one negative equational literal in F_1 or F_2. We can therefore construct a similar Nelson–Oppen derivation that starts with only the positive and the non-equational literals in Γ_1 and Γ_2, plus the one negative equational literal that may be needed for the "Unsat" step. If this negative equational literal is one of the $x_j \not\approx x_j'$, then $\exists\vec{x}\,(\Gamma \wedge x_j \not\approx x_j')$ is $\mathcal{T}$-unsatisfiable and $\forall\vec{x}\,(\Gamma \rightarrow x_j \approx x_j')$ is $\mathcal{T}$-valid; if the negative equational literal is a literal from Γ or if no negative equational literal is needed at all, then $\exists\vec{x}\,\Gamma$ is $\mathcal{T}$-unsatisfiable, so $\forall\vec{x}\,(\Gamma \rightarrow x_j \approx x_j')$ is $\mathcal{T}$-valid for every j.

Extensions

Many-sorted logics:

$read/2$ becomes $read : array \times int \to data.$
$write/3$ becomes $write : array \times int \times data \to array.$
Variables: $x : data$

Only one declaration per function/predicate/variable symbol.
All terms, atoms, substitutions must be well-sorted.

Algebras:

Instead of universe $U_{\mathcal{A}}$, one set per sort: $array_{\mathcal{A}}$, $int_{\mathcal{A}}$.

Interpretations of function and predicate symbols correspond to their declarations: $read_{\mathcal{A}} : array_{\mathcal{A}} \times int_{\mathcal{A}} \to data_{\mathcal{A}}$

If we consider combinations of theories with shared sorts but disjoint function and predicate symbols, then we get essentially the same combination results as before.

However, stable infiniteness and/or convexity are only required for the shared sorts.

Non-stably infinite theories:

If we impose stronger conditions on one theory, we can relax the conditions on the other one.

For instance, EUF can be combined with any other theory; stable infiniteness is not required.

Other examples: “shiny theories” (Tinelli/Zarba 2003)

Non-disjoint combinations:

Have to ensure that both decision procedures interpret shared symbols in a compatible way.

Some results, e. g. by Ghilardi, using strong model theoretical conditions on the theories.

Another Combination Method

Shostak's method:

Applicable to combinations of EUF and *solvable* theories.

A Σ-theory $\mathcal{T}$ is called *solvable*, if there exists an effectively computable function *solve* such that, for any $\mathcal{T}$-equation $s \approx t$:

(A) $\mathit{solve}(s \approx t) = \bot$ if and only if $\mathcal{T} \models \forall \vec{x}\,(s \not\approx t)$;

(B) $\mathit{solve}(s \approx t) = \emptyset$ if and only if $\mathcal{T} \models \forall \vec{x}\,(s \approx t)$; and otherwise

(C) $\mathit{solve}(s \approx t) = \{x_1 \approx u_1, \dots, x_n \approx u_n\}$, where

- the x_i are pairwise different variables occurring in $s \approx t$;
- the x_i do not occur in the u_j; and
- $\mathcal{T} \models \forall \vec{x}\,((s \approx t) \leftrightarrow \exists \vec{y}\,(x_1 \approx u_1 \land \ldots \land x_n \approx u_n))$, where $\vec{y}$ are the variables occurring in one of the u_j but not in $s \approx t$, and $\vec{x} \cap \vec{y} = \emptyset$.

Additionally useful (but not required):

A canonizer, that is, a function that simplifies terms by computing some unique normal form

Main idea of the procedure:

If $s \approx t$ is a positive equation and $\mathit{solve}(s \approx t) = \{x_1 \approx u_1, \dots, x_n \approx u_n\}$, replace $s \approx t$ by $x_1 \approx u_1 \land \ldots \land x_n \approx u_n$ and use these equations to eliminate the x_i elsewhere.

Practical problem:

Solvability is a rather restrictive condition.