

2.11 Combining Decision Procedures

Problem:

Let \mathcal{T}_1 and \mathcal{T}_2 be first-order theories over the signatures Σ_1 and Σ_2 .

Assume that we have decision procedures for the satisfiability of existentially quantified formulas (or the validity of universally quantified formulas) w.r.t. \mathcal{T}_1 and \mathcal{T}_2 .

Can we combine them to get a decision procedure for the satisfiability of existentially quantified formulas w.r.t. $\mathcal{T}_1 \cup \mathcal{T}_2$?

General assumption:

Σ_1 and Σ_2 are disjoint.

The only symbol shared by \mathcal{T}_1 and \mathcal{T}_2 is built-in equality.

We consider only conjunctions of literals.

For general formulas, convert to DNF first and consider each conjunction individually.

Abstraction

To be able to use the individual decision procedures, we have to transform the original formula in such a way that each atom contains only symbols of one of the signatures (plus variables).

This process is known as *variable abstraction* or *purification*.

We apply the following rule as long as possible:

$$\frac{\exists \vec{x} (F[t])}{\exists \vec{x}, y (F[y] \wedge t \approx y)}$$

if the top symbol of t belongs to Σ_i and t occurs in F directly below a Σ_j -symbol or in a (positive or negative) equation $s \approx t$ where the top symbol of s belongs to Σ_j ($i \neq j$), and if y is a new variable.

It is easy to see that the original and the purified formula are equivalent.

Stable Infiniteness

Problem:

Even if the Σ_1 -formula F_1 and the Σ_2 -formula F_2 do not share any symbols (not even variables), and if F_1 is \mathcal{T}_1 -satisfiable and F_2 is \mathcal{T}_2 -satisfiable, we cannot conclude that $F_1 \wedge F_2$ is $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -satisfiable.

Example:

Consider

$$\mathcal{T}_1 = \{\forall x, y, z (x \approx y \vee x \approx z \vee y \approx z)\}$$

and

$$\mathcal{T}_2 = \{\exists x, y, z (x \not\approx y \wedge x \not\approx z \wedge y \not\approx z)\}.$$

All \mathcal{T}_1 -models have at most two elements, and all \mathcal{T}_2 -models have at least three elements.

Since $\mathcal{T}_1 \cup \mathcal{T}_2$ is contradictory, there are no $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -satisfiable formulas.

To ensure that \mathcal{T}_1 -models and \mathcal{T}_2 -models can be combined to $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -models, we require that both \mathcal{T}_1 and \mathcal{T}_2 are stably infinite.

A first-order theory \mathcal{T} is called *stably infinite*, if every existentially quantified formula that has a \mathcal{T} -model has also a \mathcal{T} -model with a (countably) infinite universe.

Note: By the Löwenheim–Skolem theorem, “countable” is redundant here.

Shared Variables

Even if $\exists \vec{x} F_1$ is \mathcal{T}_1 -satisfiable and $\exists \vec{x} F_2$ is \mathcal{T}_2 -satisfiable, it can happen that $\exists \vec{x} (F_1 \wedge F_2)$ is not $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -satisfiable, for instance because the shared variables x and y must be equal in all \mathcal{T}_1 -models of $\exists \vec{x} F_1$ and different in all \mathcal{T}_2 -models of $\exists \vec{x} F_2$.

Example:

Consider

$$F_1 = (x + (-y) \approx 0),$$

and

$$F_2 = (f(x) \not\approx f(y))$$

where \mathcal{T}_1 is linear rational arithmetic and \mathcal{T}_2 is EUF.

We must exchange information about shared variables to detect the contradiction.

The Nelson–Oppen Algorithm (Non-deterministic Version)

Suppose that $\exists \vec{x} F$ is a purified conjunction of Σ_1 and Σ_2 -literals.

Let F_1 be the conjunction of all literals of F that do not contain Σ_2 -symbols; let F_2 be the conjunction of all literals of F that do not contain Σ_1 -symbols. (Equations between variables are in both F_1 and F_2 .)

The Nelson–Oppen algorithm starts with the pair F_1, F_2 and applies the following inference rules.

Unsat:

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

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

Branch:

$$\frac{F_1, F_2}{F_1 \wedge (x \approx y), F_2 \wedge (x \approx y) \quad | \quad F_1 \wedge (x \not\approx y), F_2 \wedge (x \not\approx y)}$$

if x and y are two different variables appearing in both F_1 and F_2 such that neither $x \approx y$ nor $x \not\approx y$ occurs in both F_1 and F_2

“|” means non-deterministic (backtracking!) branching of the derivation into two sub-derivations. Derivations are therefore trees. All branches need to be reduced until termination.

Clearly, all derivation paths are finite since there are only finitely many *shared variables* in F_1 and F_2 , therefore the procedure represented by the rules is terminating.

We call a constraint configuration to which no rule applies *irreducible*.

Theorem 2.2 (Soundness) *If “Branch” can be applied to F_1, F_2 , then $\exists \vec{x}(F_1 \wedge F_2)$ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$ if and only if one of the successor configurations of F_1, F_2 is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.*

Corollary 2.3 *If all paths in a derivation tree from F_1, F_2 end in \perp , then $\exists \vec{x}(F_1 \wedge F_2)$ is unsatisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.*

For completeness we need to show that if one branch in a derivation terminates with an irreducible configuration F_1, F_2 (different from \perp), then $\exists \vec{x}(F_1 \wedge F_2)$ (and, thus, the initial formula of the derivation) is satisfiable in the combined theory.

As $\exists \vec{x}(F_1 \wedge F_2)$ is irreducible by “Unsat”, the two formulas are satisfiable in their respective component theories, that is, we have \mathcal{T}_i -models \mathcal{A}_i of $\exists \vec{x} F_i$ for $i \in \{1, 2\}$. We are left with combining the models into a single one that is both a model of the combined theory and of the combined formula. These constructions are called *amalgamations*.

Let F be a Σ_i -formula and let S be a set of variables of F . F is called *compatible* with an equivalence \sim on S if the formula

$$\exists \vec{z} \left(F \wedge \bigwedge_{x, y \in S, x \sim y} x \approx y \wedge \bigwedge_{x, y \in S, x \not\sim y} x \not\approx y \right) \quad (1)$$

is \mathcal{T}_i -satisfiable whenever F is \mathcal{T}_i -satisfiable. This expresses that F does not contradict equalities between the variables in S as given by \sim .

The formula $\bigwedge_{x, y \in S, x \sim y} x \approx y \wedge \bigwedge_{x, y \in S, x \not\sim y} x \not\approx y$ is called an *arrangement* of S .

Proposition 2.4 *If F_1, F_2 is a pair of conjunctions over \mathcal{T}_1 and \mathcal{T}_2 , respectively, that is irreducible by “Branch”, then both F_1 and F_2 are compatible with some equivalence \sim on the shared variables S of F_1 and F_2 .*

Proof. If F_1, F_2 is irreducible by the branching rule, then for each pair of shared variables x and y , both F_1 and F_2 contain either $x \approx y$ or $x \not\approx y$. Choose \sim to be the equivalence given by all (positive) variable equations between shared variables that are contained in F_1 .

Let $\Sigma = (\Omega, \Pi)$; let $\Sigma' = (\Omega', \Pi')$ with $\Omega' \subseteq \Omega$ and $\Pi' \subseteq \Pi$ be a subsignature of Σ .

Let \mathcal{A} be a Σ -algebra. Then the *reduct* $\mathcal{A}|_{\Sigma'}$ is the Σ' -algebra \mathcal{A}' with

$$\begin{aligned} U_{\mathcal{A}'} &= U_{\mathcal{A}}, \\ f_{\mathcal{A}'} &= f_{\mathcal{A}} \text{ for all } f \in \Omega', \text{ and} \\ P_{\mathcal{A}'} &= P_{\mathcal{A}} \text{ for all } P \in \Pi'. \end{aligned}$$

Lemma 2.5 (Amalgamation Lemma) *Let \mathcal{T}_1 and \mathcal{T}_2 be two stably infinite theories over disjoint signatures Σ_1 and Σ_2 . Furthermore let F_1, F_2 be a pair of conjunctions of literals over \mathcal{T}_1 and \mathcal{T}_2 , respectively, both compatible with some equivalence \sim on the shared variables of F_1 and F_2 . Then $F_1 \wedge F_2$ is $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -satisfiable if and only if each F_i is \mathcal{T}_i -satisfiable.*

Proof. The “only if” part is obvious.

For the “if” part, assume that each of the F_i is \mathcal{T}_i -satisfiable. That is, there exist models \mathcal{A}_i of \mathcal{T}_i and variable assignments β_i such that $\mathcal{A}_i, \beta_i \models F_i$. As the F_i are compatible with an equivalence \sim on their shared variables, we may assume that the β_i also satisfy the extended conjunctions in (1) with S the set of shared variables. In particular, whenever we have two shared variables x and y , $\beta_1(x) = \beta_1(y)$ if and only if $\beta_2(x) = \beta_2(y)$. Since the theories are stably infinite we may additionally assume that the \mathcal{A}_i have countably infinite universes, hence there are bijections ρ_i from the domain of \mathcal{A}_i to \mathbb{N} such that $\rho_1(\beta_1(x)) = \rho_2(\beta_2(x))$ for each shared variable x . Now define \mathcal{A} to be the algebra having \mathbb{N} as its domain; for f or P in Σ_i define $f_{\mathcal{A}}(n_1, \dots, n_k) = \rho_i(f_{\mathcal{A}_i}(\rho_i^{-1}(n_1), \dots, \rho_i^{-1}(n_k)))$ and $P_{\mathcal{A}}(n_1, \dots, n_k) \Leftrightarrow P_{\mathcal{A}_i}(\rho_i^{-1}(n_1), \dots, \rho_i^{-1}(n_k))$. Define $\beta(x) = \rho_i(\beta_i(x))$ if x is a variable occurring in F_i . By construction of the ρ_i this definition is independent of the choice of i . Clearly $\mathcal{A}|_{\Sigma_i}, \beta \models F_i$, for $i = 1, 2$, hence $\mathcal{A}, \beta \models F_1 \wedge F_2$. Moreover, the reducts $\mathcal{A}|_{\Sigma_i}$ are isomorphic (via ρ_i) to \mathcal{A}_i and thus are models of \mathcal{T}_i , so that \mathcal{A} is a model of $\mathcal{T}_1 \cup \mathcal{T}_2$ as required.

Theorem 2.6 *The non-deterministic Nelson–Oppen algorithm is terminating and complete for deciding satisfiability of pure conjunctions of literals F_1 and F_2 over $\mathcal{T}_1 \cup \mathcal{T}_2$ for signature-disjoint, stably infinite theories \mathcal{T}_1 and \mathcal{T}_2 .*

Proof. Suppose that F_1, F_2 is irreducible by the inference rules of the Nelson–Oppen algorithm. Applying the amalgamation lemma in combination with Prop. 2.4 we infer that F_1, F_2 is satisfiable w.r.t. $\mathcal{T}_1 \cup \mathcal{T}_2$.

Convexity

The number of possible equivalences of shared variables grows superexponentially with the number of shared variables, so enumerating all possible equivalences non-deterministically is going to be inefficient.

A much faster variant of the Nelson–Oppen algorithm exists for convex theories.

A first-order theory \mathcal{T} is called *convex w.r.t. equations*, if for every conjunction Γ of Σ -equations and non-equational Σ -literals and for all Σ -equations A_i ($1 \leq i \leq n$), whenever $\mathcal{T} \models \forall \vec{x} (\Gamma \rightarrow A_1 \vee \dots \vee A_n)$, then there exists some index j such that $\mathcal{T} \models \forall \vec{x} (\Gamma \rightarrow A_j)$.

Theorem 2.7 *If a first-order theory \mathcal{T} is convex w.r.t. equations and has no trivial models (i.e., models with only one element), then \mathcal{T} is stably infinite.*

Proof. We shall prove the contrapositive of the statement. Suppose \mathcal{T} is not stably infinite. Then there exists a \mathcal{T} -satisfiable conjunction of literals $\exists \vec{x} F$ that has only finite \mathcal{T} -models. As \mathcal{T} is a first-order theory and first-order logic is compact, all \mathcal{T} -models of $\exists \vec{x} F$ are bounded in cardinality by some number m .

Let y_1, \dots, y_{m+1} be fresh variables not occurring in F . Then the formula $F_0 = \exists \vec{x} F \wedge \exists y_1 \dots y_{m+1} \bigwedge_{1 \leq i < j \leq m+1} y_i \not\approx y_j$ is \mathcal{T} -unsatisfiable since it expresses the fact that $\exists \vec{x} F$ has a model with more than m elements. Therefore, $\mathcal{T} \models \neg F_0$.

We can write F in the form $F^+ \wedge F^-$, where F^- contains the negative equational literals in F and F^+ contains the rest. Then $\mathcal{T} \models \neg F_0$ can be written as $\mathcal{T} \models \forall \vec{x} \vec{y} (\neg F^+ \vee \neg F^- \vee \bigvee_{1 \leq i < j \leq m+1} y_i \approx y_j)$, or equivalently, $\mathcal{T} \models \forall \vec{x}, \vec{y} (F^+ \rightarrow (\neg F^- \vee \bigvee_{1 \leq i < j \leq m+1} y_i \approx y_j))$. Note that $\neg F^-$ is a disjunction of positive equational literals.

Assume that $\mathcal{T} \models \forall \vec{x}, \vec{y} (F^+ \rightarrow A)$ for some literal A of $\neg F^- \vee \bigvee_{1 \leq i < j \leq m+1} y_i \approx y_j$. If A is a literal of $\neg F^-$, then $\mathcal{T} \models \forall \vec{x}, \vec{y} (F^+ \rightarrow A) \models \forall \vec{x}, \vec{y} (F^+ \rightarrow \neg F^-) \models \forall \vec{x}, \vec{y} \neg F$, which cannot hold since F is \mathcal{T} -satisfiable. Otherwise A is a literal $y_i \approx y_j$, then $\mathcal{T} \models \forall \vec{x}, \vec{y} (F^+ \rightarrow y_i \approx y_j)$. This cannot hold either: Note that $\exists \vec{x} F$ and thus $\exists \vec{x} F^+$ is \mathcal{T} -satisfiable. So let \mathcal{A} be some \mathcal{T} -model of $\exists \vec{x} F^+$. By assumption, \mathcal{A} is not a trivial model, therefore there is an \mathcal{A} -assignment β to \vec{x}, \vec{y} that satisfies F^+ and maps y_i and y_j to two arbitrary different elements. Consequently, $\mathcal{A}, \beta \not\models (F^+ \rightarrow y_i \approx y_j)$.

Lemma 2.8 Suppose \mathcal{T} is convex, F a conjunction of literals, and S a subset of its variables. Let, for any pair of variables x_i and x_j in S , $x_i \sim x_j$ if and only if $\mathcal{T} \models \forall \vec{x}(F \rightarrow x_i \approx x_j)$. Then F is compatible with \sim .

Proof. We show that with this choice of \sim the constraint (1), that is,

$$\exists \vec{z} \left(F \wedge \bigwedge_{x,y \in S, x \sim y} x \approx y \wedge \bigwedge_{x,y \in S, x \not\sim y} x \not\approx y \right)$$

is \mathcal{T} -satisfiable whenever F is. Suppose, to the contrary, that F is \mathcal{T} -satisfiable but (1) is not, that is,

$$\mathcal{T} \models \forall \vec{z} \left(F \rightarrow \bigvee_{x,y \in S, x \sim y} x \not\approx y \vee \bigvee_{x,y \in S, x \not\sim y} x \approx y \right)$$

or, equivalently,

$$\mathcal{T} \models \forall \vec{z} \left(F^+ \wedge \bigwedge_{x,y \in S, x \sim y} x \approx y \rightarrow \neg F^- \vee \bigvee_{x,y \in S, x \not\sim y} x \approx y \right).$$

where F^- contains the negative equational literals in F and F^+ contains the rest. By convexity of \mathcal{T} , the antecedent implies one of the equations of the succedent.

Suppose that this equation A comes from $\neg F^-$. Then

$$\mathcal{T} \models \forall \vec{z} \left(F^+ \wedge \bigwedge_{x,y \in S, x \sim y} x \approx y \rightarrow A \right)$$

and therefore

$$\mathcal{T} \models \forall \vec{z} \left(F^+ \wedge \bigwedge_{x,y \in S, x \sim y} x \approx y \rightarrow \neg F^- \right)$$

which means

$$\mathcal{T} \models \forall \vec{z} \left((F^+ \wedge F^-) \wedge \bigwedge_{x,y \in S, x \sim y} x \approx y \rightarrow \perp \right)$$

which cannot hold since $F = (F^+ \wedge F^-)$ is \mathcal{T} -satisfiable and entails the equations $x \approx y$ with $x \sim y$.

So the equation A must come from the last part of the disjunction. In other words, there exists a pair of different variables x' and y' in S such that $x' \not\sim y'$ and

$$\mathcal{T} \models \forall \vec{z} \left(F^+ \wedge \bigwedge_{x,y \in S, x \sim y} x \approx y \rightarrow x' \approx y' \right).$$

Since

$$\mathcal{T} \models \forall \vec{z} \left(F \rightarrow \bigwedge_{x,y \in S, x \sim y} x \approx y \right),$$

we derive $\mathcal{T} \models \forall \vec{z} \left(F \rightarrow x' \approx y' \right)$, which is impossible.

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

Unsat:

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

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

Propagate:

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

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 2.9 *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, $F_1 \wedge F_2$ is $(\mathcal{T}_1 \cup \mathcal{T}_2)$ -satisfiable.

Corollary 2.10 *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 2.11 *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 2.5 constructs a model with a countably infinite universe for F from the models of F_1 and F_2 .

Lemma 2.12 *A first-order theory \mathcal{T} is convex w.r.t. equations if and only 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 $\mathcal{T} \models \forall \vec{x} (\Gamma \rightarrow x_1 \approx x'_1 \vee \dots \vee x_n \approx x'_n)$, then there exists some index j such that $\mathcal{T} \models \forall \vec{x} (\Gamma \rightarrow x_j \approx x'_j)$.*

Lemma 2.13 *Let \mathcal{T} be a first-order theory that is convex w.r.t. equations. 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 2.14 *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. Then clearly $\mathcal{T}_1 \cup \mathcal{T}_2$ cannot have trivial models either, since any such model would also be a trivial model of \mathcal{T}_1 and \mathcal{T}_2 .

Assume furthermore that $\mathcal{T} \models \forall \vec{x} (\Gamma \rightarrow x_1 \approx x'_1 \vee \dots \vee x_n \approx x'_n)$ for some conjunction Γ of $(\Sigma_1 \cup \Sigma_2)$ -equations and non-equational $(\Sigma_1 \cup \Sigma_2)$ -literals. Then $\exists \vec{x} (\Gamma \wedge x_1 \not\approx x'_1 \wedge \dots \wedge x_n \not\approx x'_n)$ is \mathcal{T} -unsatisfiable, and we can detect this by some run of the deterministic Nelson–Oppen algorithm starting with $\exists \vec{x}, \vec{y} (\Gamma_1 \wedge \Gamma_2 \wedge x_1 \not\approx x'_1 \wedge \dots \wedge x_n \not\approx x'_n)$, 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 by Lemma 2.13 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 at most one negative equational literal that may be needed for the “Unsat” step. If a negative equational literal is needed, it 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 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* : $array \times int \rightarrow data$.

write/3 becomes *write* : $array \times int \times data \rightarrow array$.

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:

$read_{\mathcal{A}} : array_{\mathcal{A}} \times int_{\mathcal{A}} \rightarrow data_{\mathcal{A}}$

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.

E.g.: Strongly polite theories, shiny theories, flexible theories.

Strong Politeness

A theory \mathcal{T} is called *smooth*, if every quantifier-free formula that has a \mathcal{T} -model with some (finite or infinite) cardinality κ_0 has also \mathcal{T} -models with cardinality κ for every $\kappa \geq \kappa_0$.

A theory \mathcal{T} is called *finitely witnessable*, if there is a computable function wit that maps every quantifier-free formula F to a quantifier-free formula G such that

1. F and $\exists \vec{w} G$ are \mathcal{T} -equivalent, where $\vec{w} = \text{var}(G) \setminus \text{var}(F)$,
2. if $G \wedge \Delta$ is \mathcal{T} -satisfiable for some arrangement Δ , then there is a \mathcal{T} -model \mathcal{A} and an assignment β such that $\mathcal{A}, \beta \models (G \wedge \Delta)$ and $U_{\mathcal{A}} = \{ \beta(x) \mid x \in \text{var}(G \wedge \Delta) \}$

A theory \mathcal{T} is called *strongly polite*, if it is smooth and finitely witnessable.

Theorem 2.15 (Barrett & Jovanović) *We can combine two theories \mathcal{T}_1 and \mathcal{T}_2 if one of them is strongly polite.*

Again, in the many-sorted case, smoothness and finite witnessability must hold for all the shared sorts.

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) $\text{solve}(s \approx t) = \perp$ if and only if $\mathcal{T} \models \forall \vec{x} (s \not\approx t)$;
- (B) $\text{solve}(s \approx t) = \emptyset$ if and only if $\mathcal{T} \models \forall \vec{x} (s \approx t)$; and otherwise
- (C) $\text{solve}(s \approx t) = \{x_1 \approx u_1, \dots, x_n \approx u_n\}$, where
 - the x_i are pairwise different variables occurring in $s \approx t$;

- the x_i do not occur in the u_j ; and
- $\mathcal{T} \models \forall \vec{x} ((s \approx t) \leftrightarrow \exists \vec{y} (x_1 \approx u_1 \wedge \dots \wedge x_n \approx u_n))$, 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 $\text{solve}(s \approx t) = \{x_1 \approx u_1, \dots, x_n \approx u_n\}$, replace $s \approx t$ by $x_1 \approx u_1 \wedge \dots \wedge x_n \approx u_n$ and use these equations to eliminate the x_i elsewhere.

Practical problem:

Solvability is a rather restrictive condition.

Literature

Harald Ganzinger: Shostak Light. Automated Deduction, CADE-18, LNCS 2392, pp 332–346, Springer, 2002.

Silvio Ghilardi: Model Theoretic Methods in Combined Constraint Satisfiability. Journal of Automated Reasoning, 33(3–4):221–249, 2005.

Dejan Jovanović and Clark Barrett: Polite theories revisited. Logic for Programming, Artificial Intelligence, and Reasoning, LNCS 6397, pp. 402–416, Springer, 2010.

Greg Nelson, Derek C. Oppen: Simplification by Cooperating Decision Procedures. ACM Transactions on Programming Languages and Systems, 1(2):245–257, 1979.

Robert E. Shostak: Deciding Combinations of Theories. Journal of the ACM, 31(1):1–12, 1984.