Chapter 7

Propositional Logic Modulo Theories

In Chapter 6 I have studied a number of decision procedures for conjunctions of literals of some specific first-order theory or fragment. In this chapter the decision procedures are extended in two different ways. Firstly, by considering conjunctions of literals over several first-order theories. The respective procedure is the Nelson-Oppen combination procedure for theories [71], Section 7.1. Secondly, I lift the procedure from conjunctions of literals to arbitrary boolean combinations of literals. The respective procedure is CDCL(T), Section 7.2.

7.1 Nelson-Oppen Combination

Here I discuss a basic variant of the Nelson-Oppen [71] (NO) combination procedure for two theories $\mathcal{T}_1$ and $\mathcal{T}_2$ (see Definition 3.17.1) over two respective signatures Σ_1 and Σ_2 that do not share any function, constant, or predicate symbols, but may share sorts. The idea of the procedure is to reduce satisfiability of a quantifier-free formula over $\Sigma_1 \cup \Sigma_2$ to satisfiability of two separate formulas over Σ_1 and Σ_2, respectively.

The underlying semantics is that a quantifier-free formula ϕ over $\Sigma_1 \cup \Sigma_2$ is satisfiable if there exists a $\Sigma_1 \cup \Sigma_2$ algebra $\mathcal{A}$ such that $\mathcal{A}(\beta) \models \phi$ for some assignment β, and $\mathcal{A}|_{\Sigma_1}$ is isomorphic to a model in $\mathcal{T}_1$ and $\mathcal{A}|_{\Sigma_2}$ is isomorphic to a model in $\mathcal{T}_2$. Here $\mathcal{A}|_\Sigma$ denotes the restriction of $\mathcal{A}$ to the symbols in Σ. With appropriate restrictions, see below, the problem of testing satisfiability of ϕ can actually be reduced to solving finitely many separate satisfiability problems in Σ_1 and Σ_2, respectively.

Note that both theories share the equality symbol, because it is part of the first-order operator language. It is needed to separate the theories by the introduction of extra variables, called *parameters* and to transfer results from reasoning in $\mathcal{T}_1$ to $\mathcal{T}_2$ and vice versa.

For example, consider a combination of $\mathcal{T}_1 = \{\mathcal{A}_{\mathrm{LRA}}\}$, Section 6.2, with EUF, $\mathcal{T}_2 = \{\top\}$, Section 6.1 with signatures $\Sigma_1 = \Sigma_{\mathrm{LA}}$ and $\Sigma_2 = (\{S, \mathrm{LA}\}, \{g, a, b, c\}, \emptyset)$ and ground formula

$$\phi = g(b) > 5 \wedge g(c) < 5 \wedge g(c) \approx a \wedge g(b) \approx a.$$

Note that for LRA I fixed the standard algebra, whereas for EUF I fixed

a set with one axiom, actually $\top$. So for EUF all first-order Σ_2-algebras are considered. For both theories Chapter 6 contains decision procedures, however, ϕ contains mixed atoms such as $g(b) > 5$ that cannot be processed by the respective decision procedures. So the first step is *purification* where all mixed atoms are translated into pure atoms of Σ_1, Σ_2, respectively.

$$\phi = x_{\mathrm{LA}} > 5 \wedge y_{\mathrm{LA}} < 5 \wedge g(c) \approx a \wedge g(b) \approx a \wedge g(b) \approx x_{\mathrm{LA}} \wedge g(c) \approx y_{\mathrm{LA}}$$

Note parameters, e.g., x_{LA}, y_{LA}, are always implicitly existentially quantified. Now the separated formulas considered for both theories are

$$\begin{array}{l}\phi_1 = x_{\mathrm{LA}} > 5 \wedge y_{\mathrm{LA}} < 5\\ \phi_2 = g(c) \approx a \wedge g(b) \approx a \wedge g(b) \approx x_{\mathrm{LA}} \wedge g(c) \approx y_{\mathrm{LA}}\end{array}$$

Any LRA procedure for ϕ_1 immediately returns true. Congruence closure applied to ϕ_2 generates $x_{\mathrm{LA}} \approx y_{\mathrm{LA}}$ for the two existentially quantified variables. Transferring this equation to the LRA procedure on $\phi_1 \wedge x_{\mathrm{LA}} = y_{\mathrm{LA}}$ results in false. Therefore, ϕ is not satisfiable.

The example exhibits another property required by the respective theories, they have to be *convex*: if a disjunction of equations is the consequence of the theory, actually one equation holds. This property holds for LRA but not for LIA. For example,

$$1 < x_{\mathrm{LIA}} \wedge x_{\mathrm{LIA}} < 4 \models_{\mathrm{LIA}} x_{\mathrm{LIA}} = 2 \vee x_{\mathrm{LIA}} = 3$$

but none of the two single disjuncts is a consequence. Therefore, the Nelson-Oppen combination procedure between LIA and EUF will not be able to detect unsatisfiability of the already purified formulas

$$\begin{array}{l}\phi_1 = 1 < x_{\mathrm{LIA}} \wedge x_{\mathrm{LIA}} < 4 \wedge 1 < y_{\mathrm{LIA}} \wedge y_{\mathrm{LIA}} < 4 \wedge 1 < z_{\mathrm{LIA}} \wedge z_{\mathrm{LIA}} < 4\\ \phi_2 = x_{\mathrm{LIA}} \not\approx y_{\mathrm{LIA}} \wedge y_{\mathrm{LIA}} \not\approx z_{\mathrm{LIA}} \wedge z_{\mathrm{LIA}} \not\approx x_{\mathrm{LIA}}.\end{array}$$

Definition 7.1.1 (Convex Theory)**.** A theory $\mathcal{T}$ is *convex* if for a conjunction ϕ of literals with $\phi \models_{\mathcal{T}} x_1 \approx y_1 \vee \ldots \vee x_n \approx y_n$ then $\phi \models_{\mathcal{T}} x_k \approx y_k$ for some k.

Another property needed for the Nelson-Oppen procedure to work is that the theory models always include models with an infinite domain. Consider the two theories

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

and

$$\mathcal{T}_2 = \{\exists x.(x \not\approx a \wedge x \not\approx b \wedge a \not\approx b)\}$$

that do not share any signature symbols. Models of $\mathcal{T}_1$ have at most two elements, models of $\mathcal{T}_2$ at least three. So the conjunction $(\mathcal{T}_1 \cup \mathcal{T}_2)$ is already unsatisfiable. In order to ensure that different models for the respective theory can be combined, the Nelson-Oppen procedure requires the existence of models with infinite cardinality.

Definition 7.1.2 (Stably-Infinite Theory)**.** A theory $\mathcal{T}$ is *stably-infinite* if for every quantifier-free formula ϕ, if $\mathcal{T} \models \phi$, then then there exists also a model $\mathcal{A}$ of infinite cardinality, such that $\mathcal{A} \models_{\mathcal{T}} \phi$

Definition 7.1.3 (Nelson-Oppen Basic Restrictions)**.** Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two theories. Then the *Nelson-Oppen Basic Restrictions* are:

1. There are decision procedures for $\mathcal{T}_1$ and $\mathcal{T}_2$.
2. Each decision procedure returns a complete set of variable identities as consequence of a formula.
3. $\Sigma_1 \cap \Sigma_2 = \emptyset$ except for common sorts.
4. Both theories are convex.
5. $\mathcal{T}_1$ and $\mathcal{T}_2$ are stably-infinite.

Actually, restriction 7.1.3-2 is not needed, because a given finite quantifier-free formula ϕ over $\Sigma_1 \cup \Sigma_2$ contains only finitely many different variables. Now instead of putting the burden to identify variables on the decision procedure, all potential variable identifications can be guessed and tested afterwards. The disadvantage of this approach is, of course, that there are exponentially many identifications with respect to a fixed number of variables. Therefore, assuming 7.1.3-2 results in a more efficient procedure and is also supported by many procedures from Section 6. Still I will also formulate the procedure with respect to guessing the identifications, Definition 7.1.6, Proposition 7.1.7, because it enables a more elegant proof of completeness.

Restriction 7.1.3-5 can be further relaxed to assume that the domains of all shared sorts of all models are either infinite or have the same number of elements.

The Nelson-Oppen restrictions and procedure can be extended from two so several theories in the obvious way.

Example 7.1.4. $\mathcal{T}_1$ may be LA with the standard LA model over $\mathbb{Q}$ as the only model in $\mathcal{C}_1$ and $\mathcal{T}_2$ is EUF over $\Sigma_2 = \{a, g, f\}$, where a is a constant, g has arity 1 and f arity 2, with all respective term-generated models in $\mathcal{C}_2$.

The goal of the Nelson-Oppen combination procedure is now to decide the satisfiability of a quantifier-free formula ϕ over $\Sigma_1 \cup \Sigma_2$. The variables are implicitly existentially quantified. It actually suffices to consider conjunctions of atoms, because for boolean combinations CDCL(NO), Section 7.2, does the job. The first step of the procedure is to apply purification, i.e., transform the formula ϕ into a satisfiability equivalent formula ϕ' such that no term of an atom in ϕ' contains symbols from Σ_1 and Σ_2. This can always be achieved by the introduction of fresh variables.

Example 7.1.5. Consider the atom $f(x_1, 0) \geq x_3$ with respect to the theories of Example 7.1.4. The satisfiability preserving purified formula for $f(x_1, 0) \geq x_3$ is $x_4 \geq x_3 \wedge x_4 \approx f(x_1, x_5) \wedge x_5 \approx 0$.

Let N be a set of $\Sigma_1 \cup \Sigma_2$ literals interpreted as the conjunction. Then purification amounts to the exhaustive application of the following rule.

ToDo: differentiate between preprocessing (purify) and the nelson-oppen algorithm, introduce sharing of common subterms, see also congruence closure

Purify $N \uplus \{L[t[s]_i]_p\} \Rightarrow_{\mathrm{NO}} N \uplus \{L[t[z]_i]_p, z \approx s\}$
if $t = f(t_1, \ldots, t_n)$, $s = h(s_1, \ldots, s_m)$, the function symbols f and h are from different signatures, $1 \leq i \leq n$, (i.e., $t_i = s$) and z is a fresh variable of appropriate sort

After exhaustive application of Purify to any set N of $\Sigma_1 \cup \Sigma_2$ literals the set N can actually be split into two sets $N = N_1 \cup N_2$ where N_1 is build over Σ_1, N_2 is build over Σ_2 and N_1 and N_2 only share variables. Variable equations are distributed in both N_1 and N_2. Now a Nelson-Oppen problem state is a five tuple (N_1, E_1, N_2, E_2, s) with $s \in \{\top, \bot, \mathrm{fail}\}$, the sets E_1 and E_2 contain variable equations, and N_1, N_2 literals over the respective signatures, where

$(N_1; \emptyset; N_2; \emptyset; \bot)$	is the start state for some purified set of atoms $N = N_1 \cup N_2$ where the N_i are built from the respective signatures only
$(N_1; E_1; N_2; E_2; \mathrm{fail})$	is a final state, where $N_1 \cup N_2 \cup E_1 \cup E_2$ is unsatisfiable
$(N_1; E_1; N_2; E_2; \bot)$	is an intermediate state, where $N_1 \cup E_2$ and $N_2 \cup E_1$ have to be checked for satisfiability
$(N_1; \emptyset; N_2; \emptyset; \top)$	is a final state, where $N_1 \cup N_2$ is satisfiable

Solve $(N_1; E_1; N_2; E_2; \bot) \Rightarrow_{\mathrm{NO}} (N_1'; E_1'; N_2'; E_2'; \bot)$
if $N_1' = N_1 \cup E_1 \cup E_2$ and $N_2' = N_2 \cup E_1 \cup E_2$ are both $\mathcal{T}_i$-satisfiable, respectively, E_1' are all new variable equations derivable from N_1', E_2' are all new variable equations derivable from N_2' and $E_1' \cup E_2' \neq \emptyset$

Success $(N_1; E_1; N_2; E_2; \bot) \Rightarrow_{\mathrm{NO}} (N_1'; \emptyset; N_2'; \emptyset; \top)$
if $N_1' = N_1 \cup E_1 \cup E_2$ and $N_2' = N_2 \cup E_1 \cup E_2$ are both $\mathcal{T}_i$-satisfiable, respectively, E_1' are all new variable equations derivable from N_1', E_2' are all new variable equations derivable from N_2' and $E_1' \cup E_2' = \emptyset$

Fail $(N_1; E_1; N_2; E_2; \bot) \Rightarrow_{\mathrm{NO}} (N_1; E_1; N_2; E_2; \mathrm{fail})$
if $N_1' = N_1 \cup E_1 \cup E_2$ or $N_2' = N_2 \cup E_1 \cup E_2$ is $\mathcal{T}_i$-unsatisfiable, respectively

In the definition of the rules all derived equalities between variables are added to N_1 and N_2 and the decision procedures are always called to test satisfiability and produce new variable equalities. In an implementation this is not needed, a decision procedure needs only to be called if a new equality was derived by the other decision procedure.

The EUF decision procedure can easily be extended to explicitly produce derived variable equalities. For the suggested LA procedures (Fourier-Motzkin, Simplex, Virtual Substitution) this requires some extra work.

As a first example, consider the formula over LA and EUF

$$f(x_1, 0) \geq x_3 \wedge f(x_1, 0) \leq x_3$$

which becomes after purification

$$x_4 \geq x_3 \wedge f(x_1, x_5) \approx x_4 \wedge x_5 \approx 0 \wedge x_6 \leq x_3 \wedge f(x_1, x_5) \approx x_6$$

and the respective NO derivation is

$$\begin{aligned}
&(\{x_4 \geq x_3, x_5 \approx 0, x_6 \leq x_3\}, \emptyset, \{f(x_1, x_5) \approx x_4, f(x_1, x_5) \approx x_6\}, \emptyset, \bot) \\
&\Rightarrow^{\text{Solve}}_{\text{NO}} (\{x_4 \geq x_3, x_5 \approx 0, x_6 \leq x_3\}, \emptyset, \\
&\qquad \{f(x_1, x_5) \approx x_4, f(x_1, x_5) \approx x_6\}, \{x_4 \approx x_6\}, \bot) \\
&\Rightarrow^{\text{Solve}}_{\text{NO}} (\{x_4 \approx x_6, x_4 \geq x_3, x_5 \approx 0, x_6 \leq x_3\}, \{x_4 \approx x_3, x_6 \approx x_3\}, \\
&\qquad \{f(x_1, x_5) \approx x_4, f(x_1, x_5) \approx x_6, x_4 \approx x_6\}, \emptyset, \bot) \\
&\Rightarrow^{\text{Success}}_{\text{NO}} (\{x_4 \approx x_6, x_4 \geq x_3, x_5 \approx 0, x_6 \leq x_3, x_4 \approx x_3, x_6 \approx x_3\}, \emptyset, \\
&\qquad \{f(x_1, x_5) \approx x_4, f(x_1, x_5) \approx x_6, x_4 \approx x_6, x_4 \approx x_3, x_6 \approx x_3\}, \emptyset, \top)
\end{aligned}$$

Note that the Purify rule was applied in the above example in a slightly different way where the variable x_5 is shared for both occurrences of the term $f(x_1, 0)$. For an actual implementation, it is desirable to share as many subterms as possible that way.

As a second example, consider the formula over LA and EUF

$$x - y \approx 0 \wedge g(x) \not\approx g(y)$$

which is already purified and the respective NO derivation is

$$\begin{aligned}
&(\{x - y \approx 0\}, \emptyset, \{g(x) \not\approx g(y)\}, \emptyset, \bot) \\
&\Rightarrow^{\text{Solve}}_{\text{NO}} (\{x - y \approx 0\}, \{x \approx y\}, \{g(x) \not\approx g(y)\}, \emptyset, \bot) \\
&\Rightarrow^{\text{Fail}}_{\text{NO}} (\{x - y \approx 0\}, \{x \approx y\}, \{g(x) \not\approx g(y)\}, \emptyset, \text{fail})
\end{aligned}$$

For EUF variable identities are anyway computed by the congruence closure algorithm when computing the equivalence classes by generating a terminating and confluent R (see Section 6.1). However, for LA and, e.g., the simplex algorithm (see Section 6.2.2), it only comes at additional cost to identify variable identities.

Definition 7.1.6 (Arrangement)**.** Given a (finite) set of parameters X, an *arrangement* A over X is a (finite) set of equalities and inequalities over X such that for all $x_1, x_2 \in X$ either $x_1 \approx x_2 \in A$ or $x_1 \not\approx x_2 \in A$.

Proposition 7.1.7 (Nelson-Oppen modulo Arrangement)**.** Let $\mathcal{T}_1$ and $\mathcal{T}_2$ be two theories satisfying the restrictions of Definition 7.1.3 except for restriction 2. Let ϕ be a conjunction of literals over $\Sigma_1 \cup \Sigma_2$. Let N_1 and N_2 be the purified literal sets out of ϕ. Then ϕ is satisfiable iff there is an arrangement A over vars(ϕ) such that $N_1 \cup A$ is $\mathcal{T}_1$-satisfiable and $N_2 \cup A$ is $\mathcal{T}_2$-satisfiable.

Note that it is not sufficient to consider just equalities for some arrangement, because in one theory these equalities might imply further equalities which are then not transferred into the other theory.

Theorem 7.1.8 (Nelson-Oppen is Sound, Complete and Terminating)**.** Let $\mathcal{T}_1, \mathcal{T}_2$ be two theories satisfying the Nelson-Oppen basic restrictions. Let ϕ be a conjunction of literals over $\Sigma_1 \cup \Sigma_2$ and N_1, N_2 be the result of purifying ϕ.
(i) All sequences $(N_1; \emptyset; N_2; \emptyset; \bot) \Rightarrow_{\text{NO}}^{*} \ldots$ are finite.
Let $(N_1; \emptyset; N_2; \emptyset; \bot) \Rightarrow_{\text{NO}}^{*} (N_1; E_1; N_2; E_2; s)$ be a derivation with finite state $(N_1; E_1; N_2; E_2; s)$,
(ii) If $s = \text{fail}$ then ϕ is unsatisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.
(iii) If $s = \top$ then ϕ is satisfiable in $\mathcal{T}_1 \cup \mathcal{T}_2$.

Proof. (i) The relation $\Rightarrow_{\text{NO}}$ terminates as soon as no new equations are derived or one combination of formulas and equations is unsatisfiable. There are only finitely many different equations over the common variables of N_1, N_2, so $\Rightarrow_{\text{NO}}$ terminates.

(ii) Clearly purification preservers satisfiability. The Solve rule only adds logical consequences of the respective theory. Hence, if rule Fail is applicable then clearly $N_1 \cup E_1$ ($N_2 \cup E_2$) is unsatisfiable, hence ϕ is not satisfiable. This proves soundness.

(iii) Completeness is more complicated. I show it for the Nelson-Oppen formulation modulo arrangements, Proposition 7.1.7, completeness of $\Rightarrow_{\text{NO}}$ is then implied by convexity of $\mathcal{T}_1, \mathcal{T}_2$. Assume that the theories $\mathcal{T}_1, \mathcal{T}_2$ are given by possibly countably infinite sets of first-order clauses, we also denote by $\mathcal{T}_1, \mathcal{T}_2$. Then $\mathcal{T}_1 \cup \mathcal{T}_2 \cup \{\phi\}$ is unsatisfiable iff $\mathcal{T}_1 \cup \mathcal{T}_2 \cup N_1 \cup N_2$ is unsatisfiable iff $(\mathcal{T}_1 \wedge N_1) \rightarrow (\neg\mathcal{T}_2 \vee \neg N_2)$ is valid. By Craig's interpolation Theorem 3.12.15, there exists a finite set of clauses H such that $\mathcal{T}_1 \wedge N_1 \rightarrow H$ and $H \rightarrow (\neg\mathcal{T}_2 \vee \neg N_2)$, or, reformulated, $(H \wedge \mathcal{T}_2) \rightarrow \neg N_2$. The symbols used in H are common non-variable symbols of N_1 and N_2. So H is a conjunction of clauses over equations with universally quantified variables y_j and shared parameters x_i. It has the form $\bigwedge\bigvee[\neg]t_i \approx t_j$ of equational literals where the t_i, t_j are universally quantified variables y_j or parameters x_i. An equation $y_i \approx y_j$ between universally quantified variables is true iff $i = j$ and therefore needs not to be considered. Now this CNF can be transformed into a DNF yielding $\bigvee\bigwedge t_i \approx t_j$, in summary, $(\mathcal{T}_1 \wedge N_1) \rightarrow (\bigvee\bigwedge t_i \approx t_j)$. Next, I prove by

contradiction that actually one conjunct $\bigwedge t_i \approx t_j$ is implied and no t_i, t_j is a universally quantified variable. Assume this is not the case, i.e., in each of the conjuncts there are equation(s) $[\neg]x_i \approx a_i$ needed to establish the overall truth of $(\mathcal{T}_1 \wedge N_1) \rightarrow (\bigvee \bigwedge t_i \approx t_j)$. Then $(\mathcal{T}_1 \wedge N_1) \rightarrow (\bigvee x_i \approx a_i)$, where I filter only the positive equations out of the conjuncts. But the formula $(\bigvee x_i \approx a_i)$ implies a finite model, contradicting that $\mathcal{T}_1$ is stably infinite. Therefore, if $\mathcal{T}_1 \cup \mathcal{T}_2 \cup N_1 \cup N_2$ is unsatisfiable, then there is an arrangement E of the parameters such that $(\mathcal{T}_1 \wedge N_1 \wedge E)$ or $(\mathcal{T}_2 \wedge N_2 \wedge E)$ is unsatisfiable. □

Exercises

(7.1) Apply Nelson-Oppen to the LRA, EUF combination and the formula

$$f(x_1, 0) \geq x_3 \wedge f(x_2, 0) \leq x_3 \wedge x_1 \approx x_2 \wedge x_3 - f(x_1, 0) \geq 1$$

where I assume appropriate sorts for the variables and function declarations.

(7.2) Apply Nelson-Oppen to the LRA, EUF combination and the formula

$$f(f(a)) \approx b, f(a) \approx g(b), d \geq 6, b < c + 3, g(b) \approx c, f(c) \approx d, c \geq 2$$

(7.3) Apply Nelson-Oppen to the LRA, EUF combination and the formula

$$g(g(x) - g(y)) \approx z \wedge f(0) > z + 2 \wedge x \approx y$$

where I assume appropriate sorts for the variables and function declarations.

(7.4) Apply Nelson-Oppen to the LRA, EUF combination and the formula

$$x_2 \geq x_1 \wedge x_1 - x_3 \geq x_2 \wedge g(f(x_1) - f(x_2)) \not\approx g(x_3) \wedge x_3 \geq 0$$

where I assume appropriate sorts for the variables and function declarations.

(7.5)∗ Let $\Sigma = (\Omega, \emptyset)$ be a signature without predicate symbols (except built-in equality). For two Σ-algebras $\mathcal{A}$ and $\mathcal{B}$, we define the product $\mathcal{A} \times \mathcal{B}$ as the Σ-algebra whose universe is the cartesian product of the universes of $\mathcal{A}$ and $\mathcal{B}$, and where $f_{\mathcal{A}\times\mathcal{B}}((a_1, b_1), \ldots, (a_n, b_n)) = (f_{\mathcal{A}}(a_1, \ldots, a_n), f_{\mathcal{B}}(b_1, \ldots, b_n))$.
A Σ-theory $\mathcal{T}$ is called closed under products, if the product of any two models of $\mathcal{T}$ is again a model of $\mathcal{T}$.
Prove: If $\mathcal{T}$ is closed under products, then it is convex.

(7.6)∗ Prove: If the axioms of the Σ-theory $\mathcal{T}$ are universally quantified equational Horn clauses (that is, clauses where all atoms are equations and at most one of the literals is positive), then $\mathcal{T}$ is convex. You may use the previous exercise.