3.6 Constraint Superposition

So far:

Refutational completeness proof for superposition is based on the analysis of inferences between ground instances of clauses.

Inferences between ground instances must be covered by inferences between original clauses.

Non-ground clauses represent the set of all their ground instances.

Do we really need *all* ground instances?

Constrained Clauses

A *constrained clause* is a pair (C, K), usually written as $C \llbracket K \rrbracket$, where C is a Σ-clause and K is a formula (called *constraint*).

Often: K is a boolean combination of *ordering literals* $s \succ t$ with Σ-terms s, t. (also possible: comparisons between literals or clauses).

Intuition: $C \llbracket K \rrbracket$ represents the set of all ground clauses $C\theta$ for which $K\theta$ evaluates to true for some fixed term ordering. Such a $C\theta$ is called a ground instance of $C \llbracket K \rrbracket$.

A clause C without constraint is identified with $C \llbracket \top \rrbracket$.

A constrained clause $C \llbracket \bot \rrbracket$ with an unsatisfiable constraint represents no ground instances; it can be discarded.

Constraint Superposition

Inference rules for constrained clauses:

Pos. Superposition:

$$\frac{D' \vee t \approx t' \llbracket K_2 \rrbracket \qquad C' \vee s[u] \approx s' \llbracket K_1 \rrbracket}{(D' \vee C' \vee s[t'] \approx s')\sigma \llbracket (K_2 \wedge K_1 \wedge K)\sigma \rrbracket}$$

where $\sigma = \mathrm{mgu}(t, u)$ and
u is not a variable and
$$\begin{aligned} K = (t \succ t' \ \wedge \ & s[u] \succ s' \\ \wedge \ & (t \approx t') \succ_C D' \\ \wedge \ & (s[u] \approx s') \succ_C C' \\ \wedge \ & (s[u] \approx s') \succ_L (t \approx t')) \end{aligned}$$

The other inference rules are modified analogously.

To work with effectively with constrained clauses in a calculus, we need methods to check the satisfiability of constraints:

Possible for LPO, KBO, but expensive.

If constraints become too large, we may delete some conjuncts of the constraint. (Note that the calculus remains sound, if constraints are replaced by implied constraints.)

Refutational Completeness

The refutational completeness proof for constraint superposition looks mostly like in Sect. 3.4.

Lifting works as before, so every ground infererence that is required in the proof is an instance of some inference from the corresponding constrained clauses. (Easy.)

There is one significant problem, though.

Case 2 in the proof of Thm. 3.9 does not work for constrained clauses:

If we have a ground instance $C\theta$ where $x\theta$ is reducible by $R_{C\theta}$, we can no longer conclude that $C\theta$ is true because it follows from some rule in $R_{C\theta}$ and some smaller ground instance $C\theta'$.

Example: Let $C \llbracket K \rrbracket$ be the clause $f(x) \approx a \llbracket x \succ a \rrbracket$, let $\theta = \{x \mapsto b\}$, and assume that $R_{C\theta}$ contains the rule $b \to a$.
Then θ satisfies K, but $\theta' = \{x \mapsto a\}$ does not, so $C\theta'$ is *not* a ground instance of $C \llbracket K \rrbracket$.

Solution (Nieuwenhuis and Rubio, 1992):

Assumption: We start the saturation with a set N_0 of *unconstrained* clauses; the limit N_* contains constrained clauses, though.

During the model construction, we ignore ground instances $C\theta$ of clauses in N_* for which $x\theta$ is reducible by $R_{C\theta}$.

We obtain a model R_∞ of all *variable irreducible* ground instances of clauses in N_*.

R_∞ is also a model of all *variable irreducible* ground instances of clauses in N_0.

Since all clauses in N_0 are unconstrained, every ground instance of a clause in N_0 follows from some rule in R_∞ and some smaller ground instance; so it is true in R_∞.

Consequently, R_∞ is a model of *all* ground instances of clauses in N_0.

Other Constraints

The approach also works for other kinds of constraints.

In particular, we can replace unification by equality constraints ($\leadsto$ "basic superposition"):

Pos. Superposition:
$$\frac{D' \vee t \approx t' \, [\![K_2]\!] \qquad C' \vee s[u] \approx s' \, [\![K_1]\!]}{D' \vee C' \vee s[t'] \approx s' \, [\![K_2 \wedge K_1 \wedge K]\!]}$$
where u is not a variable and
$K = (t = u)$

Note: In contrast to ordering constraints, these constraints are essential for soundness.

The Drawback

Constraints reduce the number of required inferences; however, they are detrimental to redundancy:

Since we consider only *variable irreducible* ground instances during the model construction, we may use only such instances for redundancy:

A clause is redundant, if all its variable irreducible ground instances follow from smaller variable irreducible ground instances.

Even worse, since we don't know R_∞ in advance, we must consider variable irreducibility w. r. t. arbitrary rewrite systems.

Consequence: Not every subsumed clause is redundant!

3.7 Hierarchic Superposition

The superposition calculus is a powerful tool to deal with formulas in *uninterpreted* first-order logic.

What can we do if some symbols have a *fixed interpretation*?

Can we combine superposition with decision procedures, e. g., for linear rational arithmetic? Can we integrate the decision procedure as a "black box"?

Sorted Logic

It is useful to treat this problem in sorted logic (cf. Sect. 1.8, page 29).

A many-sorted signature $\Sigma = (\Xi, \Omega, \Pi)$ fixes an alphabet of non-logical symbols, where

- Ξ is a set of sort symbols,
- Ω is a sets of function symbols,
- Π is a set of predicate symbols.

Each function symbol $f \in \Omega$ has a unique declaration $f : \xi_1 \times \cdots \times \xi_n \to \xi_0$; each predicate symbol $P \in \Pi$ has a unique declaration $P : \xi_1 \times \cdots \times \xi_n$ with $\xi_i \in \Xi$.

In addition, each variable x has a unique declaration $x : \xi$.

We assume that all terms, atoms, substitutions are well-sorted.

A many-sorted algebra $\mathcal{A}$ consists of

- a non-empty set $\xi_{\mathcal{A}}$ for each $\xi \in \Xi$,
- a function $f_{\mathcal{A}} : \xi_{1,\mathcal{A}} \times \cdots \times \xi_{n,\mathcal{A}} \to \xi_{0,\mathcal{A}}$ for each $f : \xi_1 \times \cdots \times \xi_n \to \xi_0 \in \Omega$,
- a subset $P_{\mathcal{A}} \subseteq \xi_{1,\mathcal{A}} \times \cdots \times \xi_{n,\mathcal{A}}$ for each $P : \xi_1 \times \cdots \times \xi_n \in \Pi$.

Hierarchic Specifications

A *specification* $SP = (\Sigma, \mathcal{C})$ consists of

- a signature $\Sigma = (\Xi, \Omega, \Pi)$,
- a class of term-generated Σ-algebras $\mathcal{C}$ closed under isomorphisms.

If $\mathcal{C}$ consists of *all* term-generated Σ-algebras satisfying the set of Σ-formulas N, we write $SP = (\Sigma, N)$.

A *hierarchic specification* $HSP = (SP, SP')$ consists of

- a base specification $SP = (\Sigma, \mathcal{C})$,
- an extension $SP' = (\Sigma', N')$,

where $\Sigma = (\Xi, \Omega, \Pi)$, $\Sigma' = (\Xi', \Omega', \Pi')$, $\Xi \subseteq \Xi'$, $\Omega \subseteq \Omega'$, and $\Pi \subseteq \Pi'$.

A Σ'-algebra $\mathcal{A}$ is called a model of $HSP = (SP, SP')$, if $\mathcal{A}$ is a model of N' and $\mathcal{A}|_\Sigma \in \mathcal{C}$, where the reduct $\mathcal{A}|_\Sigma$ is defined as $((\xi_\mathcal{A})_{\xi \in \Xi}, (f_\mathcal{A})_{f \in \Omega}, (P_\mathcal{A})_{P \in \Pi})$.

Note:

- no confusion: models of HSP may not identify elements that are different in the base models.
- no junk: models of HSP may not add new elements to the interpretations of base sorts.

Example:

$\Sigma = (\{rat\}, \Omega, \Pi)$,
where $\Omega = \{\, q \mid q \in \mathbb{Q} \,\} \cup \{\, q \cdot \_ \mid q \in \mathbb{Q} \,\} \cup \{+, -\}$ and $\Pi = \{<, >, \leq, \geq\}$.

$\mathcal{C} = $ isomorphy class of $\mathbb{Q}$.

$\Sigma' = (\{rat, data\}, \Omega', \Pi')$,
where $\Omega' = \Omega \cup \{f : rat \to data, \, g : data \to rat\}$ and $\Pi' = \Pi$.

$N' = \Sigma'$-clauses (axioms for f and g and negated theorem to be proved).

Goal:

Check whether N' has a model in which the sort rat is interpreted by $\mathbb{Q}$ and the symbols from Ω and Π accordingly.

Hierarchic Superposition

In order to use a prover for the base theory, we must separate the base and the non-base part of clauses:

A clause is called *abstracted*, if none of its atoms contains both a base (function or predicate) symbol and a non-base symbol.

Every clause can be transformed into an equivalent abstracted clause. From now on, we consider only abstracted clauses.

Note: There is no restriction for base *variables*.

A term (or atom or clause) that consists only of base symbols and variables of base sort is called a base term (atom, clause).

A substitution is called simple, if it maps every variable of a base sort to a base term.

The inference rules of the hierarchic superposition calculus correspond to the rules of of the standard superposition calculus with the following modifications:

- We consider only simple substitutions as unifiers.
- We perform only inferences on non-base literals.

While clauses that contain non-base literals are manipulated using superposition rules, base clauses have to be passed to the base prover.

This yields one more inference rule:

Constraint Refutation:
$$\frac{M}{\bot}$$

where M is a set of base clauses
that is inconsistent w.r.t. $\mathcal{C}$.

Problems

We can only apply the constraint refutation rule to finite sets M. if $\mathcal{C}$ is not compact, this is not sufficient.

Since we only consider simple substitutions, we will only obtain a model of all *simple ground instances*.

To show that we have a model of *all* instances, we need an additional condition called *sufficient completeness w.r.t. simple instances*.

A set N of clauses is called *sufficiently complete with respect to simple instances*, if for every model $\mathcal{A}'$ of the set of simple ground instances of N and every ground non-base term t of a base sort there exists a ground base term t such that $t' \approx t$ is true in $\mathcal{A}'$.

Note: Sufficient completeness w.r.t. simple instances ensures the absence of junk.

Completeness of Hierarchic Superposition

If the base theory is compact, the hierarchic superposition calculus is refutationally complete for sets of clauses that are sufficiently complete with respect to simple instances (Bachmair, Ganzinger, Waldmann, 1992).

Main proof idea: If the set of base clauses in N has some base model, represent this model by a set E of convergent ground equations and a set D of ground disequations. Then show: If N is saturated w.r.t. hierarchic superposition, then $E \cup D \cup \tilde{N}$ is saturated w.r.t. standard superposition, where $\tilde{N}$ is the set of simple ground instances of clauses in N that are reduced w.r.t. E.