A Superposition View on Nelson-Oppen

Thomas Hillenbrand

Max-Planck-Institut für Informatik, D-66123 Saarbrücken
hillen@mpi-sb.mpg.de

1 Introduction

Following a line of research to prove decidability and combination results in a resolution framework, we reinspect the Nelson-Oppen combination procedure [NO79,Opp80,TH96] and reestablish its correctness as an instance of the completeness of the superposition calculus [BG94].

In doing so, we employ superposition as a particular means of generating models. The starting-point is an observation by Bachmair that given two clause sets over disjoint signatures, if each of them is saturated and does not contain the empty clause, so is their union, which is therefore satisfiable. With a suitable fine-tuning of the calculus parameters, this result can be adapted to the Nelson-Oppen setting where the signatures share a finite number of free constants.

Our goal in this research is two-fold. Firstly, a number of attempts have been undertaken to generalize the Nelson-Oppen procedure, relaxing either the signature disjointness condition or the stably infiniteness requirement. Can these also be rephrased and maybe extended within the superposition framework? Secondly, extensions of superposition, like e. g. the chaining calculus for transitive relations [BG98], are promising candidates for stretching the limits of the combination scheme, say for reasoning about data types that are based on some ordering.

2 Preliminaries

2.1 Notions and notations

We work in a logic with built-in equality and assume without loss of generality that equality is the only predicate symbol. Therefore a *signature* Σ consists of a set of *function symbols* with associated *arities*. For most logical notions and notations related to superposition, we refer to [NR01]. To simplify the notation, the equality symbol $\simeq$ is supposed to be symmetric. Terms are built as usually over the signature and *variables*. A literal $s \bowtie t$ is either an equation $s \simeq t$ or a disequation $s \not\simeq t$. *Clauses*, *formulae*, *quantifier-free* formulae and *universal* formulae are defined as usually.

In this paper, we identify a *Σ-theory* with a set of closed first-order formulae over Σ. We employ the standard notions of *Σ-structure*, *interpretation*, *satisfiability*, *entailment*, *equivalence*, *validity*, *equisatisfiability*; also with respect to some theory. The component theories of the Nelson-Oppen procedure are *stably infinite:* For quantifier-free formulae, $\mathcal{T}$-satisfiability is equivalent to $\mathcal{T}$-satisfiability in infinite models. A related property of a theory is *convexity*, holding in case that if a clause is $\mathcal{T}$-valid, then also some Horn clause is with the same disequations, but only one of the equations.

3 Reconstructing Nelson-Oppen

3.1 Obtaining a clausal theory presentation

We have defined theories to consist of arbitrary first-order formulae, whereas the superposition calculus deals with formulae in clause normal form. Hence our first task is to find a clausal theory presentation. To ease the presentation of the Nelson-Oppen procedure, we will assume that theories are not only stably infinite, but also convex.

In the superposition framework, when combining clause sets over disjoint signatures, inferences different from the superposition rule are unary and do not produce conclusions in the combined signature. The superposition rule behaves the same as long as the terms in the clauses that are to be unified both start with a function symbol. On the other hand, clauses with positive isolated occurrences of variables are harmful. Consider e.g. the variable x in the clause $x \simeq a \vee C$ where x does not appear in C: Some instances of x are strictly maximal then, such that the inference is necessary. But if we apply some splitting rule, the malicious clause is broken into C and $x \simeq a$, the latter clause contradicting stable infiniteness.

More generally, let us call a clause C *unshielding* if $C \equiv x \simeq t \vee C'$ where x occurs neither in t nor in C'. Still $x \simeq t$ contradicts stable infiniteness. However, if we want to replace the clause C by C' within a superposition derivation, we have to apply the concept of redundancy, which demands that the replacement be not also smaller, but also logically implied. The implication holds if the theory contains an axiom $\forall x \exists y.\, x \not\simeq y$, or a Skolemized version thereof. It turns out that this can be required in our context without loss of generality.

The following lemma makes our considerations precise.

Lemma 1. Let $\mathcal{T}$ denote a stably infinite and convex Σ-theory and $\#$ a unary operator not in Σ. Then there exists a set T of Horn clauses over $\Sigma \cup \{\#\}$ such that, going from $\mathcal{T}$ to T,
(i) stable infiniteness is preserved,
(ii) convexity is preserved,
(iii) the sets of theory-satisfiable quantifier-free formulae over Σ are the same,
(iv) and every non-trivial $\mathcal{T}$-model can be extended to a T-model.
Moreover the operator $\#$ occurs in exactly one clause of T, which is $x \not\simeq \#(x)$.

Proof (outline). We will develop T in several steps, i.e. via a sequence of theories $\mathcal{T}_i$. One can show that, whenever going from $\mathcal{T}_i$ to $\mathcal{T}_{i+1}$, the properties given from (i) through (iv) are preserved.

The sequence now is the following: $\mathcal{T}_0$ is simply $\mathcal{T}$ itself. $\mathcal{T}_1$ is the set of $\mathcal{T}_0$-valid universal formulae. $\mathcal{T}_2$ is obtained adding the axiom $\forall x \exists y.\, x \not\simeq y$. The theory $\mathcal{T}_3$ is a standard clause normal form of $\mathcal{T}_2$ where the only non-universal formula $\forall x \exists y.\, x \not\simeq y$ is mapped onto the clause $x \not\simeq \#(x)$. By convexity, for every non-Horn clause $\bigvee_i (s_i \not\simeq t_i) \vee \bigvee_j (u_j \simeq v_j) \in \mathcal{T}_3$, there is a $\mathcal{T}_3$-equivalent Horn subclause $\bigvee_i (s_i \not\simeq t_i) \vee u_j \simeq v_j$. Replacing all clauses that way gives the set $\mathcal{T}_4$, which has all properties required for T. $\square$

3.2 Combining convex theories

To simplify the presentation, we restrict ourselves to the important special case that the component theories are not only stably infinite, but also convex. The Nelson-Oppen procedure then becomes somewhat more elegant in that no initial guessing of equivalences between the shared variables is needed, or in other terms: branching is not required.

The procedure deals with finitely many theories $\mathcal{T}_1, \ldots, \mathcal{T}_n$, each $\mathcal{T}_i$ being built over some signature Σ_i. The signatures are disjoint. Every theory $\mathcal{T}_i$ comes with a decision procedure for $\mathcal{T}_i$-satisfiability of quantifier-free formulae. Let $\mathcal{T}$ denote the union of the theories, and Σ the union of the signatures. Our aim is to decide $\mathcal{T}$-satisfiability of a finite set Γ of Σ-literals. This straightforwardly extends to a procedure for arbitrary quantifier-free formulae. To accommodate to the superposition setting, we replace the variables in Γ by free constant symbols from an additional, finite signature Σ_0, such that from now on Γ is ground.

We exploit as usually that Γ can be *purified*, in linear time, with the help of free constants from Σ_0 into an equivalent presentation such that each literal is built from symbols of Σ_0 and just one component signature Σ_i, $i > 0$. We partition Γ accordingly into $\bigcup_{i \geq 0} \Gamma_i$ where Γ_0 gathers the literals over Σ_0 (and variables), which are common to all component signatures.

Let us introduce two notions to be employed within the description of the procedure: We say that Γ *is component satisfiable* if every $\Gamma_i \cup \Gamma_0$ is $\mathcal{T}_i$-satisfiable, and that Γ *is component unsatisfiable* otherwise. Furthermore, a *Σ_0-conclusion of Γ* shall be any Σ_0-equation $a \simeq a'$ such that the literal set $\Gamma \cup \{a \not\simeq a'\}$ is component unsatisfiable.

Algorithm 2. *Nelson-Oppen procedure* for convex theories:
(0) Purify and partition Γ.
(1) Exhaustively, unless Γ is component unsatisfiable,
extend Γ with Σ_0-conclusions.
(2) Return satisfiable if Γ is component satisfiable,
and unsatisfiable otherwise.

Of the properties we will need to prove the combination procedure correct, let us briefly gather the evident ones:

Proposition 3. The notions of component satisfiability and Σ_0-conclusion are subject to the following properties:
(i) Component satisfiability is decidable.
(ii) For Γ, component unsatisfiability entails $\mathcal{T}$-unsatisfiability.
(iii) If $a \simeq a'$ is a Σ_0-conclusion of Γ, then Γ and $\Gamma \cup \{a \simeq a'\}$ are $\mathcal{T}$-equivalent.

(iv) The set of all Σ_0-conclusion of Γ is computable.

The less evident property is that, if the procedure claims satisfiability, one can indeed exhibit some $\mathcal{T}$-model. We will use a superposition-based model construction to this end.

Lemma 4. If Γ is closed under adding Σ_0-conclusions, then component satisfiability implies $\mathcal{T}$-satisfiability.

Proof. By Lem. 1, for every theory $\mathcal{T}_i$ over Σ_i there is a Horn clause set T_i over $\Sigma_i \cup \{\#_i\}$ with respect to which the same quantifier-free formulae over $\Sigma_i \cup \Sigma_0$ are satisfiable. Every non-trivial $\mathcal{T}_i$-model can be extended to a T_i-model, and T_i contains the operator $\#_i$ exactly in the clause $x \not\simeq \#_i(x)$.

Consider any theory $\mathcal{T}_i$. Since the literal set $\Gamma_i \cup \Gamma_0$ is $\mathcal{T}_i$-satisfiable, it is also T_i-satisfiable. Equivalently, the clause set $\Gamma_i \cup \Gamma_0 \cup T_i$ is satisfiable, and therefore has a logically equivalent presentation S_i which is saturated, in the sense of the superposition calculus (cf. [NR01, Sect. 3 and 4]). Our aim is to show that with a suitable fine-tuning of the calculus parameters the union $S := \bigcup_i S_i$ is immediately saturated and, since it does not contain the empty clause, thereby satisfiable according to the completeness of the superposition calculus.

The parameter setting is the following:

(i) We choose a lexicographic path ordering with some total precedence over $\Sigma \cup \Sigma_0 \cup \{\#_1, \ldots, \#_n\}$ such that symbols from Σ_0 are minimal.
(ii) Eager selection is applied to negative equations in Σ_0-clauses.
(iii) Every unshielding clause $y \simeq t \vee C$ produces, with a clause $x \not\simeq \#_i(x)$ under the unifier $\{y \mapsto \#_i(x)\}$, the superposition conclusion $x \not\simeq t \vee C$, which in turn produces C by equality resolution. Thereby the unshielding clause becomes redundant such that it can be dropped. This is done eagerly.

We now discuss which inference rule can produce conclusions in which symbols from different Σ_i are mixed. The rule *equality factoring* is not applicable since every clause is Horn. The rule *equality resolution* is unary. Every conclusion is already contained in that clause set S_i which holds the premise, or is redundant there, and therefore need not be added to S. Likewise, applications of the binary *superposition rule* are not critical if both premises belong to the same S_i.

So it remains to study an inference between a clause $C \vee l \simeq r \in S_i$ and another clause $D \vee s[l'] \bowtie t \in S_j$. By definition of the superposition rule, the subterm l' is not a variable. If l were a variable x, then any further occurrence of x within the first premise would violate the strict maximality condition on l. Hence the first premise were an unshielding clause; but such clauses are dropped.

Therefore both l and l' are headed by function symbols, which have to coincide since l and l' are unifiable. This common symbol has to belong to the signature part common to S_i and S_j, namely Σ_0. Since Σ_0 is minimal in the precedence, the maximality conditions imply that both premises contain function symbols only from Σ_0.

Because eager selection is applied to negative equations in Σ_0-clauses, the first premise $C \vee l \simeq r \in S_i$ of the superposition inference is simply an equation $l \simeq r$, or more precisely $a \simeq a'$ because unshielding clauses are dropped.

Since every non-trivial model of $\Gamma_i \cup \Gamma_0 \cup \mathcal{T}_i$ extends to one of S_i, and since $a \simeq a'$ holds in every model of S_i, we obtain that $a \simeq a'$ must be valid in $\Gamma_i \cup \Gamma_0 \cup \mathcal{T}_i$. Hence $a \simeq a'$ is a Σ_0-conclusion of Γ and contained in Γ_0, as well as in every S_k. In particular the first premise is also an element of S_j, so that the superposition inference is not needed. □

Theorem 5. For convex component theories, Alg. 2 decides $\mathcal{T}$-satisfiability of quantifier-free formulae.

Proof. As outlined in the beginning of the subsection, Step (0) is computable and transforms Γ into a $\mathcal{T}$-equisatisfiable set of literals. By Prop. 3(i),(iv) every iteration of Step (1) is computable; by Prop. 3(iii) it produces another $\mathcal{T}$-equisatisfiable presentation of Γ. Since Σ_0 is finite, Step (1) terminates, and so does the procedure. After Step (1), the resulting literal set Γ is closed under adding Σ_0-conclusions. If Γ is component satisfiable, then by Lem. 4 a $\mathcal{T}$-model exists. Otherwise it is $\mathcal{T}$-unsatisfiable by Prop. 3(ii). □

References

[BG94] L. Bachmair and H. Ganzinger. Rewrite-based equational theorem proving with selection and simplification. *Journal of Logic and Computation*, 4(3):217–247, 1994.

[BG98] L. Bachmair and H. Ganzinger. Ordered chaining calculi for first-order theories of transitive relations. *Journal of the ACM*, 45(6):1007–1049, 1998.

[NO79] G. Nelson and D. Oppen. Simplification by cooperating decision procedures. *ACM Transactions on Programming Languages and Systems*, 1(2):245–257, 1979.

[NR01] R. Nieuwenhuis and A. Rubio. Paramodulation-based theorem proving. In A. Robinson and A. Voronkov, editors, *Handbook of Automated Reasoning*, volume I, chapter 7, pages 371–443. Elsevier, 2001.

[Opp80] D. Oppen. Complexity, convexity and combination of theories. *Theoretical Computer Science*, 12:291–302, 1980.

[TH96] C. Tinelli and M. Harandi. A new correctness proof of the Nelson-Oppen combination procedure. In F. Baader and K. Schulz, editors, *Proceedings of the First International Workshop on Frontiers of Combining Systems*, pages 103–120. Kluwer Academic Publishers, 1996.