Chapter 5

First-Order Logic With Equality

In this Chapter I combine the ideas of Superposition for first-order logic without equality, Section 3.13, and Knuth-Bendix Completion, Section 4.4, to get a calculus for equational clauses. In Section 3.1 I already argued that any literal can be represented by an equation by "moving predicates to functions" and introducing a new sort Bool with specific constant true that is minimal in any considered ordering.

$$\begin{array}{lcl} P(t_1, \ldots, t_n) & \Rightarrow & f_P(t_1, \ldots, t_n) \approx \text{true} \\ \neg P(t_1, \ldots, t_n) & \Rightarrow & f_P(t_1, \ldots, t_n) \not\approx \text{true} \end{array}$$

The concentration on equational literals eases notation as I will show below. The constant true is minimal in the ordering, so the left hand side of a transformed literal is always strictly maximal. The freshly introduced functions f_P only occur at top level of a term, so a critical pair overlap between two such functions corresponds exactly to a Superposition Left (resolution) or Factoring inference of the superposition calculus for first-order logic without equality. Note that a literal true $\not\approx$ true can be simplified to $\bot$ and a literal true $\approx$ true to $\top$, respectively. So from now on I only equational clauses, i.e., there are no predicate symbols, $\Pi = \emptyset$.

Inference rules are to be read modulo symmetry of the equality symbol. First, I explain the ideas and motivations behind the superposition calculus with equality and its completeness proof for the ground case. At start I do not consider selection, it will be eventually added in the obvious way when considering clauses with variables.

The running example for this chapter is the theory of arrays $\mathcal{T}_{\text{Array}}$, see also Section 7.3, which consists of the following three axioms:

$$\begin{array}{l} \forall x_A, y_I, z_V . \, \text{read}(\text{store}(x, y, z), y) \approx z \\ \forall x_A, y_I, {y'}_I, z_V . (y \not\approx y' \rightarrow \text{read}(\text{store}(x, y, z), y') \approx \text{read}(x, y')) \\ \forall x_A, {x'}_A . \exists y_I . (\text{read}(x, y) \not\approx \text{read}(x', y) \vee x \approx x'). \end{array}$$

The goal is to decide for an additional set of ground clauses N over the above signature plus further constants of the three different sorts, whether $\mathcal{T}_{\text{Array}} \cup N$ is satisfiable. I will show that superposition can be turned into a decision procedure for this problem, following [**?**]. The superposition calculus including some array specific refinements, will always terminate on a clause set $\mathcal{T}_{\text{Array}} \cup N$. This results in an alternative decision procedure compared to the instantiation-based procedures used in the SMT (Satisfiability Modulo Theories) context, see Section 7.3.

5.1 Ground Superposition

The idea of the superposition calculus without equality was to restrict inferences to maximal literals, Section 3.13. Knuth-Bendix completion considers critical pairs between maximal sides of equations, Section 4.4. Superposition on equational clauses combines the two restrictions: inferences are between maximal left hand sides of maximal literals in the respective clauses. Since all considered orderings are total on ground terms, they maximality conditions can be stated positively.

The ground inference rules corresponding to Knuth-Bendix critical pair computation generalized to clauses. Superposition Left on first-order logic without equality is generalized to equational clauses an inferences below top atom positions. Then the ordering construction of Definition 3.12.1 is lifted to equational clauses. The multiset $\{s, t\}$ is assigned to a positive literal $s \approx t$, the multiset $\{s, s, t, t\}$ is assigned to a negative literal $s \not\approx t$. The *literal ordering* $\succ_L$ compares these multisets using the multiset extension of $\succ$. The *clause ordering* $\succ_C$ compares clauses by comparing their multisets of literals using the multiset extension of $\succ_L$. Eventually $\succ$ is used for all three orderings depending on the context.

Superposition Left $(N \uplus \{D \vee t \approx t', C \vee s[t] \not\approx s'\}) \Rightarrow_{\text{SUPE}} (N \cup \{D \vee t \approx t', C \vee s[t] \not\approx s'\} \cup \{D \vee C \vee s[t'] \not\approx s'\})$

where $t \approx t'$ is strictly maximal and $s \not\approx s'$ are maximal in their respective clauses, $t \succ t'$, $s \succ s'$

Superposition Right $(N \uplus \{D \vee t \approx t', C \vee s[t] \approx s'\}) \Rightarrow_{\text{SUPE}} (N \cup \{D \vee t \approx t', C \vee s[t] \approx s'\} \cup \{D \vee C \vee s[t'] \approx s'\})$

where $t \approx t'$ and $s \approx s'$ are strictly maximal in their respective clauses, $t \succ t'$, $s \succ s'$

The two rules are not yet sufficient to obtain completeness. There is no rule corresponding to Factoring and there is no way to apply reflexivity of equality, i.e., refute negative equations. The latter is solved by the below rule Equality Resolution.

Equality Resolution $(N \uplus \{C \vee s \not\approx s\}) \Rightarrow_{\text{SUPE}} (N \cup \{C \vee s \not\approx s\} \cup \{C\})$

where $s \not\approx s$ is maximal in the clause

Similar to Factoring on ground clauses, Equality Resolution is also a simplification on ground clauses, i.e., the parent clause becomes redundant with respect to the result of the derivation step. Once Equality Resolution is lifted to clauses with variables this is no longer the case, because the applied substitution may instantiate further literals in C.

It turns out that a direct adaption of the Factoring rule from superposition for first-order logic without equality is not sufficient for completeness. This becomes obvious in the context of the model construction. Basically, for the model construction the same ideas as in the completeness proof for superposition without equality apply, see Section 3.13. However, a Herbrand interpretation does not work for equality: the equality symbol $\approx$ must be interpreted by equality in the interpretation. The solution is to define a set E of ground equations and take $T(\Sigma,\emptyset)/E = T(\Sigma,\emptyset)/{\approx_E}$ as the universe. Then two ground terms s and t are equal in the interpretation if and only if $s \approx_E t$. If E is a terminating and confluent rewrite system R, then two ground terms s and t are equal in the interpretation, if and only if $s \downarrow_R t$.

Now the problem with the standard factoring rule is that in the completeness proof for the superposition calculus without equality, the following property holds: if $C = C' \vee A$ with a strictly maximal atom A is false in the current interpretation N_C with respect to some clause set, see Definition 3.12.5, then adding A to the current interpretation cannot make any literal in C' true. This does not hold anymore in the presence of equality. Let $b \succ c \succ d$. Assume that the current rewrite system (representing the current interpretation) contains the rule $c \to d$. Now consider the clause $b \approx c \vee b \approx d$ where $b \approx c$ is strictly maximal. A further needed inference rule to deal with clauses of this kind, is the below Equality Factoring rule, a generalization of the non-equational Factoring rule.

Equality Factoring $(N \uplus \{C \vee s \approx t' \vee s \approx t\}) \Rightarrow_{\text{SUPE}} (N \cup \{C \vee s \approx t' \vee s \approx t\} \cup \{C \vee t \not\approx t' \vee s \approx t'\})$

where $s \succ t'$, $s \succ t$ and $s \approx t$ is maximal in the clause

5.2 Superposition

The lifting from the ground case to the first-order case with variables is then identical to the case of superposition without equality: identity is replaced by unifiability, the mgu is applied to the resulting clause, and $\succ$ is replaced by $\not\preceq$. In addition, as in Knuth-Bendix completion, overlaps at or below a variable position are not considered. The consequence is that there are inferences between ground instances $D\sigma$ and $C\sigma$ of clauses D and C which are not ground instances of inferences between D and C. Such inferences have to be treated in a special way in the completeness proof and will be shown to be obsolete.

Until now I mostly described the ideas behind the superposition calculus and its completeness proof. Now, precise definitions and proofs will be given.

Inference rules are applied with respect to the commutativity of equality $\approx$. Selection of negative literals is considered as well.

Superposition Right $(N \uplus \{D \vee t \approx t', C \vee s[u] \approx s'\}) \Rightarrow_{\text{SUPE}} (N \cup \{D \vee t \approx t', C \vee s[u] \approx s'\} \cup \{(D \vee C \vee s[t'] \approx s')\sigma\})$
where σ is the mgu of t, u, u is not a variable $t\sigma \not\preceq t'\sigma$, $s\sigma \not\preceq s'\sigma$, $(t \approx t')\sigma$ strictly maximal in $(D \vee t \approx t')\sigma$, nothing selected and $(s \approx s')\sigma$ maximal in $(C \vee s \approx s')\sigma$ and nothing selected

Superposition Left $(N \uplus \{D \vee t \approx t', C \vee s[u] \not\approx s'\}) \Rightarrow_{\text{SUPE}} (N \cup \{D \vee t \approx t', C \vee s[u] \not\approx s'\} \cup \{(D \vee C \vee s[t'] \not\approx s')\sigma\})$
where σ is the mgu of t, u, u is not a variable $t\sigma \not\preceq t'\sigma$, $s\sigma \not\preceq s'\sigma$, $(t \approx t')\sigma$ strictly maximal in $(D \vee t \approx t')\sigma$, nothing selected and $(s \not\approx s')\sigma$ maximal in $(C \vee s \not\approx s')\sigma$ or selected

Equality Resolution $(N \uplus \{C \vee s \not\approx s'\}) \Rightarrow_{\text{SUPE}} (N \cup \{C \vee s \not\approx s'\} \cup \{C\sigma\})$
where σ is the mgu of s, s', $(s \not\approx s')\sigma$ maximal in $(C \vee s \not\approx s')\sigma$ or selected

Equality Factoring $(N \uplus \{C \vee s' \approx t' \vee s \approx t\}) \Rightarrow_{\text{SUPE}} (N \cup \{C \vee s' \approx t' \vee s \approx t\} \cup \{(C \vee t \not\approx t' \vee s \approx t')\sigma\})$
where σ is the mgu of s, s', $s'\sigma \not\preceq t'\sigma$, $s\sigma \not\preceq t\sigma$, $(s \approx t)\sigma$ maximal in $(C \vee s' \approx t' \vee s \approx t)\sigma$ and nothing selected

Proving soundness of the rules is not difficult, completeness, however, requires a non-trivial proof.

Theorem 5.2.1 (Superposition Soundness)**.** All inference rules of the superposition calculus are *sound*, i.e., for every rule $N \uplus \{C_1, \ldots, C_n\} \Rightarrow N \cup \{C_1, \ldots, C_n\} \cup \{D\}$ it holds that $\{C_1, \ldots, C_n\} \models D$.

The notion of redundancy does not change, i.e., a clause is redundant if it is implied by smaller clauses.

Definition 5.2.2 (Abstract Redundancy)**.** A clause C is *redundant* with respect to a clause set N if for all ground instances $C\sigma$ there are clauses $\{C_1, \ldots, C_n\} \subseteq N$ with ground instances $C_1\tau_1, \ldots, C_n\tau_n$ such that $C_i\tau_i \prec C\sigma$ for all i and $C_1\tau_1, \ldots, C_n\tau_n \models C\sigma$.

Given a set N of clauses $\text{red}(N)$ is the set of clauses redundant with respect to N.

Definition 5.2.3 (Saturation)**.** A clause set N is *saturated up to redundancy* if for every derivation $N \setminus \text{red}(N) \Rightarrow_{\text{SUPE}} N \cup \{C\}$ it holds $C \in (N \cup \text{red}(N))$.

For a set E of ground equations, $T(\Sigma, \emptyset)/E$ is an E-interpretation (or E-algebra) with universe $\{[t] \mid t \in T(\Sigma, \emptyset)\}$. Then for every *ground* equation $s \approx t$, $T(\Sigma, \emptyset)/E \models s \approx t$ holds if and only if $s \leftrightarrow_E^* t$, see Theorem 4.1.11. In particular,

if E is a convergent set of rewrite rules R and $s \approx t$ is a ground equation, then $T(\Sigma, \emptyset)/R \models s \approx t$ if and only if $s \downarrow_R t$. An equation or clause is valid (or true) in R if and only if it is true in $T(\Sigma, \emptyset)/R$.

Definition 5.2.4 (Partial Model Construction)**.** Given a clause set N and an ordering $\succ$ a (partial) model $N_\mathcal{I}$ can be constructed inductively over all ground clause instances of N as follows:

$$
\begin{array}{lcl}
N_C & := & \bigcup_{D \prec C}^{D \in \mathrm{grd}(\Sigma, N)} E_D \\
\\
E_D & := & \begin{cases}
\{s \approx t\} & \text{if } D = D' \vee s \approx t, \\
 & (i)\ s \approx t \text{ is strictly maximal in } D \\
 & (ii)\ s \succ t \\
 & (iii)\ D \text{ is false in } N_D \\
 & (iv)\ D' \text{ is false in } N_D \cup \{s \to t\} \\
 & (v)\ s \text{ is irreducible by } N_D \\
 & (vi) \text{ no negative literal is selected in } D' \\
\emptyset & \text{otherwise}
\end{cases} \\
N_\mathcal{I} & := & \bigcup_{C \in \mathrm{grd}(\Sigma, N)} N_C
\end{array}
$$

where N_D, $N_\mathcal{I}$, E_D are also considered as rewrite systems with respect to $\succ$. If $E_D \neq \emptyset$ then D is called *productive*.

Lemma 5.2.5 (Maximal Terms in Productive Clauses)**.** If $E_C = \{s \to t\}$ and $E_D = \{l \to r\}$, then $s \succ l$ if and only if $C \succ D$.

Corollary 5.2.6 (Partial Models are Convergent Rewrite Systems)**.** The rewrite systems N_C and $N_\mathcal{I}$ are convergent.

Proof. Obviously, $s \succ t$ for all rules $s \to t$ in N_C and $N_\mathcal{I}$. Furthermore, it is easy to check that there are no critical pairs between any two rules: Assume that there are rules $l \to r$ in E_D and $s \to t$ in E_C so that l is a subterm of s. As $\succ$ is a reduction ordering that is total on ground terms, $l \prec s$ holds and therefore $D \prec C$ and $E_D \subseteq N_C$. But then s would be reducible by N_C, contradicting condition Definition 5.2.4 (v). □

Lemma 5.2.7 (Ordering Consequences in Productive Clauses)**.** If $D \preceq C$ and $E_C = \{s \to t\}$, then $s \succ r$ for every term r occurring in a negative literal in D and $s \succeq l$ for every term l occurring in a positive literal in D.

Corollary 5.2.8 (Model Monotonicity True Clauses)**.** If D is true in N_D, then D is true in $N_\mathcal{I}$ and N_C for all $C \succ D$.

Proof. If a positive literal of D is true in N_D, then this is obvious. Otherwise, some negative literal $s \not\approx t$ of D must be true in N_D, hence $s \not\downarrow_{N_D} t$. As the rules in $N_\mathcal{I} \setminus N_D$ have left-hand sides that are larger than s and t, they cannot be used in a rewrite proof of $s \downarrow t$, hence $s \not\downarrow_{N_C} t$ and $s \not\downarrow_{N_\mathcal{I}} t$. □

Corollary 5.2.9 (Model Monotonicity False Clauses)**.** If $D = D' \vee s \approx t$ is productive, then D' is false and D is true in $N_\mathcal{I}$ and N_C for all $C \succ D$.

Proof. Obviously, D is true in $N_\mathcal{I}$ and N_C for all $C \succ D$. Since all negative literals of D' are false in N_D, it is clear that they are false in $N_\mathcal{I}$ and N_C. For the positive literals $s' \approx t'$ of D', condition Definition 5.2.4 (iv) ensures that they are false in $N_D \cup \{s \to t\}$. Since $s' \preceq s$ and $t' \preceq s$ and all rules in $N_\mathcal{I} \setminus N_D$ have left-hand sides that are larger than s, these rules cannot be used in a rewrite proof of $s' \downarrow t'$, hence $s' \not\downarrow_{N_C} t'$ and $s' \not\downarrow_{N_\mathcal{I}} t'$. □

Lemma 5.2.10 (Lifting Single Clause Inferences)**.** Let C be a clause and let σ be a substitution such that $C\sigma$ is ground. Then every equality resolution or equality factoring inference from $C\sigma$ is a ground instance of an inference from C.

Lemma 5.2.11 (Lifting Two Clause Inferences)**.** Let $D = D' \vee u \approx v$ and $C = C' \vee [\neg]s \approx t$ be two clauses (without common variables) and let σ be a substitution such that $D\sigma$ and $C\sigma$ are ground. If there is a superposition inference between $D\sigma$ and $C\sigma$ where $u\sigma$ and some subterm of $s\sigma$ are overlapped and $u\sigma$ does not occur in $s\sigma$ at or below a variable position of s then the inference is a ground instance of a superposition inference from D and C.

For the below theorem and the rest of the chapter I assume that clauses are variable disjoint and unifiers are idempotent.

Theorem 5.2.12 (Model Construction)**.** Let N be a set of clauses that is saturated up to redundancy and does not contain the empty clause. Then for every ground clause $C\sigma \in \mathrm{grd}(\Sigma, N)$ it holds that:

1. $E_{C\sigma} = \emptyset$ if and only if $C\sigma$ is true in $N_{C\sigma}$.
2. If $C\sigma$ is redundant with respect to $\mathrm{grd}(\Sigma, N)$ then it is true in $N_{C\sigma}$.
3. $C\sigma$ is true in $N_\mathcal{I}$ and in N_D for every $D \in \mathrm{grd}(\Sigma, N)$ with $D \succ C\sigma$.

Proof. The proof does not consider selection. The proof is by induction on the clause ordering $\succ$ and with the induction hypothesis that 1.–3. are already satisfied for all clauses in $\mathrm{grd}(\Sigma, N)$ that are smaller than $C\sigma$. Note that the "if" part of 1. is obvious from the construction and that condition 3. follows immediately from 1. and Corollaries 5.2.8 and 5.2.9. So it remains to show condition 2. and the "only if" part of 1.

(Condition 2) Case $C\sigma$ is redundant with respect to $\mathrm{grd}(\Sigma, N)$: If $C\sigma$ is redundant with respect to $\mathrm{grd}(\Sigma, N)$, then it follows from clauses in $\mathrm{grd}(\Sigma, N)$ that are smaller than $C\sigma$. By part 3. of the induction hypothesis, these clauses are true in $N_{C\sigma}$. Hence $C\sigma$ is true in $N_{C\sigma}$.

(Condition 1) If $E_{C\sigma} = \emptyset$ then $C\sigma$ is true in $N_{C\sigma}$.

(Condition 1.1) Case $x\sigma$ is reducible by $N_{C\sigma}$: Suppose there is a variable x occurring in C so that $x\sigma$ is reducible by $N_{C\sigma}$, say $x\sigma \to_{N_{C\sigma}} w$. Let the substitution σ' be defined by $x\sigma' = w$ and $y\sigma' = y\sigma$ for every variable $y \not\approx x$. The clause $C\sigma'$ is smaller than $C\sigma$. By part 3. of the induction hypothesis, it is true in $N_{C\sigma}$. By congruence, every literal of $C\sigma$ is true in $N_{C\sigma}$ if and only if the corresponding literal of $C\sigma'$ is true in $N_{C\sigma}$; hence $C\sigma$ is true in $N_{C\sigma}$.

(Condition 1.2) Case $C\sigma$ contains a maximal negative literal: Suppose that $C\sigma$ does not fall into Condition 2 and Condition 1.1 and that $C\sigma = C'\sigma \vee s\sigma \not\approx s'\sigma$, where $s\sigma \not\approx s'\sigma$ is maximal in $C\sigma$. If $s\sigma \approx s'\sigma$ is false in $N_{C\sigma}$, then $C\sigma$ is clearly true in $N_{C\sigma}$ and this part of the proof is done. So assume that $s\sigma \approx s'\sigma$ is true in $N_{C\sigma}$, that is, $s\sigma \downarrow_{N_{C\sigma}} s'\sigma$. without loss of generality, $s\sigma \succeq s'\sigma$.

(Condition 1.2.1) Case $s\sigma = s'\sigma$: If $s\sigma = s'\sigma$, then there is an *equality resolution* inference $N \uplus \{C'\sigma \vee s\sigma \not\approx s'\sigma\} \Rightarrow N \cup \{C'\sigma\}$. As shown in the Lifting Lemma, this is an instance of an *equality resolution* inference $N \uplus \{C' \vee s \not\approx s'\} \Rightarrow N \cup \{C'\theta\}$ where $C = C' \vee s \not\approx s'$ is contained in N and $\sigma = \theta \circ \rho$. without loss of generality, θ is idempotent, therefore $C'\sigma = C'\theta\rho = C'\theta\theta\rho = C'\theta\sigma$, so $C'\sigma$ is a ground instance of $C'\theta$. Since $C\sigma$ is not redundant with respect to $\mathrm{grd}(\Sigma, N)$, C is not redundant with respect to N. As N is saturated up to redundancy, the conclusion $C'\theta$ of the inference from C is contained in $N \cup \mathrm{red}(N)$. Therefore, $C'\sigma$ is either contained in $\mathrm{grd}(\Sigma, N)$ and smaller than $C\sigma$, or it follows from clauses in $\mathrm{grd}(\Sigma, N)$ that are smaller than itself (and therefore smaller than $C\sigma$). By the induction hypothesis, clauses in $\mathrm{grd}(\Sigma, N)$ that are smaller than $C\sigma$ are true in $N_{C\sigma}$, thus $C'\sigma$ and $C\sigma$ are true in $N_{C\sigma}$.

(Condition 1.2.2) Case $s\sigma \succ s'\sigma$: If $s\sigma \downarrow_{N_{C\sigma}} s'\sigma$ and $s\sigma \succ s'\sigma$, then $s\sigma$ must be reducible by some rule in some $E_{D\sigma} \subseteq N_{C\sigma}$. Let $D\sigma = D'\sigma \vee t\sigma \approx t'\sigma$ with $E_{D\sigma} = \{t\sigma \to t'\sigma\}$. Since $D\sigma$ is productive, $D'\sigma$ is false in $N_{C\sigma}$. Besides, by part 2. of the induction hypothesis, $D\sigma$ is not redundant with respect to $\mathrm{grd}(\Sigma, N)$, so D is not redundant with respect to N. Note that $t\sigma$ cannot occur in $s\sigma$ at or below a variable position of s, say $x\sigma = w[t\sigma]$, since otherwise $C\sigma$ would be subject to Case 1.1 above. Consequently, the *left superposition* inference $N \uplus \{D'\sigma \vee t\sigma \approx t'\sigma, C'\sigma \vee s\sigma[t\sigma] \not\approx s'\sigma\} \Rightarrow N \cup \{D'\sigma \vee C'\sigma \vee s\sigma[t'\sigma] \not\approx s'\sigma\}$ is a ground instance of a *left superposition* inference from D and C. By saturation up to redundancy, its conclusion is either contained in $\mathrm{grd}(\Sigma, N)$ and smaller than $C\sigma$, or it follows from clauses in $\mathrm{grd}(\Sigma, N)$ that are smaller than itself (and therefore smaller than $C\sigma$). By the induction hypothesis, these clauses are true in $N_{C\theta}$, thus $D'\sigma \vee C'\sigma \vee s\sigma[t'\sigma] \not\approx s'\sigma$ is true in $N_{C\sigma}$. Since $D'\sigma$ and $s\sigma[t'\sigma] \not\approx s'\sigma$ are false in $N_{C\sigma}$, both $C'\sigma$ and $C\sigma$ must be true.

(Condition 1.3) Case $C\sigma$ does not contain a maximal negative literal: Suppose that $C\sigma$ does not fall into Cases 1.1 and 1.2. Then $C\sigma$ can be written as $C'\sigma \vee s\sigma \approx s'\sigma$, where $s\sigma \approx s'\sigma$ is a maximal literal of $C\sigma$. If $E_{C\sigma} = \{s\sigma \to s'\sigma\}$ or $C'\sigma$ is true in $N_{C\sigma}$ or $s\sigma = s'\sigma$, then there is nothing to show, so assume that $E_{C\sigma} = \emptyset$ and that $C'\sigma$ is false in $N_{C\theta}$. without loss of generality, $s\sigma \succ s'\sigma$.

(Condition 1.3.1) Case $s\sigma \approx s'\sigma$ is maximal in $C\sigma$, but not strictly maximal: If $s\sigma \approx s'\sigma$ is maximal in $C\sigma$, but not strictly maximal, then $C\sigma$ can be written as $C''\sigma \vee t\sigma \approx t'\sigma \vee s\sigma \approx s'\sigma$, where $t\sigma = s\sigma$ and $t'\sigma = s'\sigma$. In this case, there is a *equality factoring* inference $N \uplus \{C''\sigma \vee t\sigma \approx t'\sigma \vee s\sigma \approx s'\sigma\} \Rightarrow N \cup \{C''\sigma \vee t'\sigma \not\approx s'\sigma \vee t\sigma \approx t'\sigma\}$. This inference is a ground instance of an inference from C. By induction hypothesis, its conclusion is true in $N_{C\sigma}$. Trivially, $t'\sigma = s'\sigma$ implies $t'\sigma \downarrow_{N_{C\sigma}} s'\sigma$, so $t'\sigma \not\approx s'\sigma$ must be false and $C\sigma$ must be true in $N_{C\sigma}$.

(Condition 1.3.2) Case $s\sigma \approx s'\sigma$ is strictly maximal in $C\sigma$ and $s\sigma$ is reducible: Suppose that $s\sigma \approx s'\sigma$ is strictly maximal in $C\sigma$ and $s\sigma$ is reducible by some rule in $E_{D\sigma} \subseteq N_{C\sigma}$. Let $D\sigma = D'\sigma \vee t\sigma \approx t'\sigma$ and $E_{D\sigma} = \{t\sigma \to t'\sigma\}$. Since $D\sigma$ is productive, $D\sigma$ is not redundant and $D'\sigma$ is false in $N_{C\sigma}$. Now proceed in essentially the same way as in Case 1.2.2: If $t\sigma$ occurred in $s\sigma$ at or below a variable position of s, say $x\sigma = w[t\sigma]$, then $C\sigma$ would be subject to Case 1.1 above. Otherwise, the *right superposition* inference $N \uplus \{D'\sigma \vee t\sigma \approx t'\sigma, C'\sigma \vee s\sigma[t\sigma] \approx s'\sigma\} \Rightarrow N \cup \{D'\sigma \vee C'\sigma \vee s\sigma[t'\sigma] \approx s'\sigma\}$ is a ground instance of a *right superposition* inference from D and C. By saturation up to redundancy, its conclusion is true in $N_{C\sigma}$. Since $D'\sigma$ and $C'\sigma$ are false in $N_{C\sigma}$, $s\sigma[t'\sigma] \approx s'\sigma$ must be true in $N_{C\sigma}$. On the other hand, $t\sigma \approx t'\sigma$ is true in $N_{C\sigma}$, so by congruence, $s\sigma[t\sigma] \approx s'\sigma$ and $C\sigma$ are true in $N_{C\sigma}$.

(Condition 1.3.3) Case $s\sigma \approx s'\sigma$ is strictly maximal in $C\sigma$ and $s\sigma$ is irreducible: Suppose that $s\sigma \approx s'\sigma$ is strictly maximal in $C\sigma$ and $s\sigma$ is irreducible by $N_{C\sigma}$. Then there are three possibilities: $C\sigma$ can be true in $N_{C\sigma}$, or $C'\sigma$ can be true in $N_{C\sigma} \cup \{s\sigma \to s'\sigma\}$, or $E_{C\sigma} = \{s\sigma \to s'\sigma\}$. In the first and the third case, there is nothing to show. Therefore assume that $C\sigma$ is false in $N_{C\sigma}$ and $C'\sigma$ is true in $N_{C\sigma} \cup \{s\sigma \to s'\sigma\}$. Then $C'\sigma = C''\sigma \vee t\sigma \approx t'\sigma$, where the literal $t\sigma \approx t'\sigma$ is true in $N_{C\sigma} \cup \{s\sigma \to s'\sigma\}$ and false in $N_{C\sigma}$. In other words, $t\sigma \downarrow_{N_{C\sigma} \cup \{s\sigma \to s'\sigma\}} t'\sigma$, but not $t\sigma \downarrow_{N_{C\sigma}} t'\sigma$. Consequently, there is a rewrite proof of $t\sigma \to^* u \,{}^*\!\leftarrow t'\sigma$ by $N_{C\sigma} \cup \{s\sigma \to s'\sigma\}$ in which the rule $s\sigma \to s'\sigma$ is used at least once. without loss of generality assume that $t\sigma \succeq t'\sigma$. Since $s\sigma \approx s'\sigma \succ t\sigma \approx t'\sigma$ and $s\sigma \succ s'\sigma$ it can be concluded that $s\sigma \succeq t\sigma \succ t'\sigma$. But then there is only one possibility how the rule $s\sigma \to s'\sigma$ can be used in the rewrite proof: $s\sigma = t\sigma$ must hold and the rewrite proof must have the form $t\sigma \to s'\sigma \to^* u \leftarrow^* t'\sigma$, where the first step uses $s\sigma \to s'\sigma$ and all other steps use rules from $N_{C\sigma}$. Consequently, $s'\sigma \approx t'\sigma$ is true in $N_{C\sigma}$. Now observe that there is an *equality factoring* inference $N \uplus \{C''\sigma \vee t\sigma \approx t'\sigma \vee s\sigma \approx s'\sigma\} \Rightarrow N \cup \{C''\sigma \vee t'\sigma \not\approx s'\sigma \vee t\sigma \approx t'\sigma\}$ whose conclusion is true in $N_{C\sigma}$ by saturation. Since the literal $t'\sigma \not\approx s'\sigma$ must be false in $N_{C\sigma}$, the rest of the clause must be true in $N_{C\sigma}$, and therefore $C\sigma$ must be true in $N_{C\sigma}$, contradicting the assumption. This concludes the proof of the theorem. □

Lemma 5.2.13 (Lifting Models)**.** Let N be a set of clauses with variables and let $\mathcal{A}$ be a term-generated Σ-algebra. Then $\mathcal{A}$ is a model of $\mathrm{grd}(\Sigma, N)$ if and only if it is a model of N.

Proof. ($\Rightarrow$) Let $\mathcal{A} \models \mathrm{grd}(\Sigma, N)$; let $(\forall \vec{x} C) \in N$. Then $\mathcal{A} \models \forall \vec{x} C$ iff $\mathcal{A}(\gamma[x_i \mapsto a_i])(C) = 1$ for all γ and a_i. Choose ground terms t_i such that $\mathcal{A}(\gamma)(t_i) = a_i$;

define σ such that $x_i\sigma = t_i$, then $\mathcal{A}(\gamma[x_i \mapsto a_i])(C) = \mathcal{A}(\gamma\circ\sigma)(C) = \mathcal{A}(\gamma)(C\sigma) = 1$ since $C\sigma \in G_\Sigma(N)$.

($\Leftarrow$) Let $\mathcal{A}$ be a model of N; let $C \in N$ and $C\sigma \in G_\Sigma(N)$. Then $\mathcal{A}(\gamma)(C\sigma) = \mathcal{A}(\gamma \circ \sigma)(C) = 1$ since $\mathcal{A} \models N$. □

Theorem 5.2.14 (Refutational Completeness: Static View)**.** Let N be a set of clauses that is saturated up to redundancy. Then N has a model if and only if N does not contain the empty clause.

Proof. If $\bot \in N$, then obviously N does not have a model. If $\bot \notin N$, then the interpretation $N_\mathcal{I}$ (that is, $T(\Sigma, \emptyset)/N_\mathcal{I}$) is a model of all ground instances in $\mathrm{grd}(\Sigma, N)$ according to Theorem 5.2.12.3. As $T(\Sigma, \emptyset)/N_\mathcal{I}$ is term generated, it is a model of N. □

So far, only inference rules that add new clauses to the current set of clauses have been considered, corresponding to the Deduce rule of Knuth-Bendix Completion. In other words, derivations of the form $N_0 \Rightarrow N_1 \Rightarrow N_2 \Rightarrow \ldots$, where each N_{i+1} is obtained from N_i by performing an inference from clauses in N_i. Under which circumstances can a clause during the derivation be deleted (or simplified)? Can additional clauses beyond the inferences be added?

Definition 5.2.15 (Superposition Run)**.** A *run* of the superposition calculus is a derivation $N_0 \Rightarrow_{\mathrm{SR}} N_1 \Rightarrow_{\mathrm{SR}} N_2 \Rightarrow_{\mathrm{SR}} \ldots$, so that

1. $N_i \models N_{i+1}$, and
2. all clauses in $N_i \setminus N_{i+1}$ are redundant with respect to N_{i+1}.

For a run, $N_\infty = \bigcup_{i\geq 0} N_i$ and $N_* = \bigcup_{i\geq 0} \bigcap_{j\geq i} N_j$. The set N_* of all *persistent* clauses is called the *limit* of the run.

In other words, during a run a new clause may be added if it follows from the old ones, and a clause may be deleted, if it is redundant with respect to the remaining ones.

Lemma 5.2.16 (Redundancy is Monotone)**.** If $N \subseteq N'$, then $\mathrm{red}(N) \subseteq \mathrm{red}(N')$.

Lemma 5.2.17 (Redundant Clauses Do not Contribute)**.** If $N' \subseteq \mathrm{red}(N)$, then $\mathrm{red}(N) \subseteq \mathrm{red}(N \setminus N')$.

Proof. Follows from the compactness of first-order logic and the well-foundedness of the multiset extension of the clause ordering. □

Lemma 5.2.18 (Redundancy is Monotone in Runs)**.** Let $N_0 \Rightarrow N_1 \Rightarrow_{\mathrm{SR}} N_2 \Rightarrow_{\mathrm{SR}} \ldots$ be a run. Then $\mathrm{red}(N_i) \subseteq \mathrm{red}(N_\infty)$ and $\mathrm{red}(N_i) \subseteq \mathrm{red}(N_*)$ for every i.

Corollary 5.2.19 (Redundancy is Monotone Modulo Persistent Clauses)**.** $N_i \subseteq N_* \cup \mathrm{red}(N_*)$ for every i.

Proof. If $C \in N_i \setminus N_*$, then there is a $k \geq i$ so that $C \in N_k \setminus N_{k+1}$, so C must be redundant with respect to N_{k+1}. Consequently, C is redundant with respect to N_*. □

Definition 5.2.20 (Fair Run)**.** A run is called *fair*, if $(N_* \setminus \mathrm{red}(N_*)) \Rightarrow_{\mathrm{SUPE}} (N_* \setminus \mathrm{red}(N_*)) \cup \{C\}$ then $C \in (N_i \cup \mathrm{red}(N_i))$ for some i.

Lemma 5.2.21 (Saturation of Fair Runs)**.** If a run is fair, then its limit is saturated up to redundancy.

Proof. If the run is fair, then the conclusion of every inference from non-redundant clauses in N_* is contained in some $N_i \cup \mathrm{red}(N_i)$, and therefore contained in $N_* \cup \mathrm{red}(N_*)$. Hence N_* is saturated up to redundancy. □

Theorem 5.2.22 (Refutational Completeness: Dynamic View)**.** Let $N_0 \Rightarrow_{\mathrm{SR}} N_1 \Rightarrow_{\mathrm{SR}} N_2 \Rightarrow_{\mathrm{SR}} \ldots$ be a fair run, let N_* be its limit. Then N_0 has a model if and only if $\bot \notin N_*$.

Proof. ($\Leftarrow$) By fairness, N_* is saturated up to redundancy. If $\bot \notin N_*$, then it has a term-generated model. Since every clause in N_0 is contained in N_* or redundant with respect to N_*, this model is also a model of $\mathrm{grd}(\Sigma, N_0)$ and therefore a model of N_0.

($\Rightarrow$) Obvious, since $N_0 \models N_*$. □

Historic and Bibliographic Remarks