3.3.1 Definition (Substitution (well-sorted))

A *well-sorted substitution* is a mapping $\sigma : \mathcal{X} \to T(\Sigma, \mathcal{X})$ so that

1. $\sigma(x) \neq x$ for only finitely many variables x and
2. $\mathrm{sort}(x) = \mathrm{sort}(\sigma(x))$ for every variable $x \in \mathcal{X}$.

The application $\sigma(x)$ of a substitution σ to a variable x is often written in postfix notation as $x\sigma$. The variable set $\mathrm{dom}(\sigma) := \{x \in \mathcal{X} \mid x\sigma \neq x\}$ is called the *domain* of σ.

The term set $\mathsf{codom}(\sigma) := \{x\sigma \mid x \in \mathsf{dom}(\sigma)\}$ is called the *codomain* of σ. From the above definition it follows that $\mathsf{dom}(\sigma)$ is finite for any substitution σ. The composition of two substitutions σ and τ is written as a juxtaposition $\sigma\tau$, i.e., $t\sigma\tau = (t\sigma)\tau$.

A substitution σ is called *idempotent* if $\sigma\sigma = \sigma$. A substitution σ is idempotent iff $\mathsf{dom}(\sigma) \cap \mathsf{vars}(\mathsf{codom}(\sigma)) = \emptyset$.

Substitutions are often written as sets of pairs $\{x_1 \mapsto t_1, \ldots, x_n \mapsto t_n\}$ if $dom(\sigma) = \{x_1, \ldots, x_n\}$ and $x_i\sigma = t_i$ for every $i \in \{1, \ldots, n\}$.

The *modification* of a substitution σ at a variable x is defined as follows:

$$\sigma[x \mapsto t](y) = \begin{cases} t & \text{if } y = x \\ \sigma(y) & \text{otherwise} \end{cases}$$

A substitution σ is identified with its extension to formulas and defined as follows:

1. $\bot\sigma = \bot$,
2. $\top\sigma = \top$,
3. $(f(t_1, \ldots, t_n))\sigma = f(t_1\sigma, \ldots, t_n\sigma)$,
4. $(P(t_1, \ldots, t_n))\sigma = P(t_1\sigma, \ldots, t_n\sigma)$,
5. $(s \approx t)\sigma = (s\sigma \approx t\sigma)$,
6. $(\neg\phi)\sigma = \neg(\phi\sigma)$,
7. $(\phi \circ \psi)\sigma = \phi\sigma \circ \psi\sigma$ where $\circ \in \{\vee, \wedge\}$,
8. $(Qx\phi)\sigma = Qz(\phi\sigma[x \mapsto z])$ where $Q \in \{\forall, \exists\}$, z and x are of the same sort and z is a fresh variable.

The result $t\sigma$ ($\phi\sigma$) of applying a substitution σ to a term t (formula ϕ) is called an *instance* of t (ϕ).

The substitution σ is called *ground* if it maps every domain variable to a ground term, i.e., the codomain of σ consists of ground terms only.

If the application of a substitution σ to a term t (formula ϕ) produces a ground term $t\sigma$ (a variable-free formula, $\mathrm{vars}(\phi\sigma) = \emptyset$), then $t\sigma$ ($\phi\sigma$) is called *ground instance* of t (ϕ) and σ is called *grounding* for t (ϕ). The set of ground instances of a clause set N is given by $\mathrm{grd}(\Sigma, N) = \{C\sigma \mid C \in N, \sigma \text{ is grounding for } C\}$ is the set of *ground instances* of N.

A substitution σ is called a *variable renaming* if $\mathrm{codom}(\sigma) \subseteq \mathcal{X}$ and for any $x, y \in \mathcal{X}$, if $x \neq y$ then $x\sigma \neq y\sigma$.

3.3.2 Lemma (Substitutions and Assignments)

Let β be an assignment of some interpretation $\mathcal{A}$ of a term t and σ a substitution. Then

$$\beta(t\sigma) = \beta[x_1 \mapsto \beta(x_1\sigma), \ldots, x_n \mapsto \beta(x_n\sigma)](t)$$

where $\operatorname{dom}(\sigma) = \{x_1, \ldots, x_n\}$.

Firstly, we define the classic Herbrand interpretations for formulas without equality.

3.5.1 Definition (Herbrand Interpretation)

A *Herbrand Interpretation* (over Σ) is a Σ-algebra $\mathcal{H}$ such that

1. $S^{\mathcal{H}} := T_S(\Sigma)$ for every sort $S \in \mathcal{S}$
2. $f^{\mathcal{H}} : (s_1, \ldots, s_n) \mapsto f(s_1, \ldots, s_n)$ where $f \in \Omega$, $\operatorname{arity}(f) = n$, $s_i \in S_i^{\mathcal{H}}$ and $f : S_1 \times \ldots \times S_n \to S$ is the sort declaration for f
3. $P^{\mathcal{H}} \subseteq (S_1^{\mathcal{H}} \times \ldots \times S_m^{\mathcal{H}})$ where $P \in \Pi$, $\operatorname{arity}(P) = m$ and $P \subseteq S_1 \times \ldots \times S_m$ is the sort declaration for P

3.5.2 Lemma (Herbrand Interpretations are Well-Defined)

Every Herbrand Interpretation is a Σ-algebra.

3.5.3 Proposition (Representing Herbrand Interpretations)

A Herbrand interpretation $\mathcal{A}$ can be uniquely determined by a set of ground atoms I

$$(s_1, \ldots, s_n) \in P^{\mathcal{A}} \text{ iff } P(s_1, \ldots, s_n) \in I$$

3.5.5 Theorem (Herbrand)

Let N be a finite set of Σ-clauses. Then N is satisfiable iff N has a Herbrand model over Σ iff $\mathrm{grd}(\Sigma, N)$ has a Herbrand model over Σ.

Here $\mathrm{grd}(\Sigma, N) = \{C\sigma \mid C \in N, \sigma$ grounding in $\Sigma\}$

First-Order Logic Theories

In Section 3.2 the semantics of a first-order formula is defined with respect to all algebras that assign meaning to the symbols of the signature. For many applications this is too crude. For example, let us assume we consider the signature of simple linear integer arithmetic without divisibility relations, $\Sigma_{\mathsf{LIA}} = (\{\mathsf{LIA}\}, \{0, 1, +, -\} \cup \mathbb{Z}, \{\leq, <, >, \geq\})$. Then a standard first-order algebra $\mathcal{A}$ is, e.g., $\mathsf{LIA}^{\mathcal{A}} = \{0, 1\}$, $0^{\mathcal{A}} = 0$, $1^{\mathcal{A}} = 1$, $k^{\mathcal{A}} = (|k| \mod 2)$ for all $k \in \mathbb{Z}$, $+^{\mathcal{A}}(0, 0) = 0$, $+^{\mathcal{A}}(1, 0) = +^{\mathcal{A}}(0, 1) = +^{\mathcal{A}}(1, 1) = 1$, $-^{\mathcal{A}}(0, 0) = -^{\mathcal{A}}(1, 1) = -^{\mathcal{A}}(0, 1) = 0$, $-^{\mathcal{A}}(1, 0) = 1$, and the relations $\leq$, $<$, $>$, $\geq$ are interpreted as usual over the domain $\{0, 1\}$. Obviously, $\mathcal{A}$ is not the standard interpretation of linear integer arithmetic, because the domain is not the integers, and , e.g., $\mathcal{A} \models 8 < 9$ but also $\mathcal{A} \models 10 < 9$.

Is there a way to fix the semantics to the intended interpretation? Actually, there are two: the syntactic way by requiring any algebra $\mathcal{A}$ of the signature Σ_{LIA} to satisfy a set of closed first-order formulas, called *axioms*, or the semantic way of fixing a set of algebras for Σ_{LIA}. In both cases, the set of algebras and axioms is a called a *theory* $\mathcal{T}$. For both cases I assume that the axioms are satisfiable and there is either at least on algebra in $\mathcal{T}$, respectively.

For the above example, the semantic way would be simply to fix the standard linear integer interpretation for $\mathcal{T} = \{\Sigma_{\mathsf{LIA}}\}$ as the only algebra to be considered. The syntactic way would mean to add enough formulas such that any algebra satisfying the formulas is the desired algebra. More concretely, the formulas

$$\begin{aligned}\mathcal{T} = \{&\{k \not\approx l \mid \text{for all } k, l \in \mathbb{Z},\ k \neq l\} \cup \\ &\{k < l \mid \text{for all } k, l \in \mathbb{Z},\ k < l\}\}\end{aligned}$$

Note, that the right hand side $\neq$ and $<$ are the standard relations on the integers. For any algebra $\mathcal{A}$ satisfying the infinitely many axioms of $\mathcal{T}$, $\mathcal{A} \models 8 < 9$ and $\mathcal{A} \models 9 < 10$ and $\mathsf{LIA}^{\mathcal{A}}$ will contain at least as many different elements as the integers. So $\mathsf{LIA}^{\mathcal{A}} = \mathbb{Z}$ is a possible domain of an algebra for $\mathcal{T}$, but also $\mathsf{LIA}^{\mathcal{A}} = \mathbb{Q}$ would satisfy the above axioms.

Fixing a set of algebras is actually the more general and powerful mechanism. However, it has also disadvantages. Given a finite set of axioms $\mathcal{T}$ proving with respect to $\mathcal{T}$ amounts to classical first-order theorem proving, e.g., validity is semi-decidable. Given a set $\mathcal{T}$ of algebras, proving with respect to the algebras is typically beyond first-order logic theorem proving, e.g., for $\mathcal{T} = \{\Sigma_{\mathsf{LIA}}\}$ theorem proving means inductive theorem proving, in general, hence, validity is no longer semi-decidable, but undecidable.

3.17.1 Definition (First-Order Logic Theory)

Given a first-order many-sorted signature Σ, a *theory* $\mathcal{T}$ is a non-empty set of Σ-algebras.

For some first-order formula ϕ over Σ we say that ϕ is $\mathcal{T}$*-satisfiable* if there is some $\mathcal{A} \in \mathcal{T}$ such that $\mathcal{A}(\beta) \models \phi$ for some β. We say that ϕ is $\mathcal{T}$*-valid* ($\mathcal{T}$*-unsatisfiable*) if for all $\mathcal{A} \in \mathcal{T}$ and all β it holds $\mathcal{A}(\beta) \models \phi$ ($\mathcal{A}(\beta) \not\models \phi$). In case of validity I also write $\models_{\mathcal{T}} \phi$.

Alternatively, $\mathcal{T}$ may contain a set of satisfiable axioms which then stand for all algebras satisfying the axioms.

The Σ-algebras can be restricted to term-generated models as long as there are "enough" constants (function) symbols in Σ, in general infinitely many are sufficient. Due to the Löwenheim-Skolem theorem different infinite cardinalities cannot be distinguished by first-order formulas.

Complexity

What is the complexity of deciding validity/unsatisfiability of a first-order formula?

1.3.6 Definition (PCP)

Given two finite lists of words $(u_1, \dots, u_n)$ and $(v_1, \dots, v_n)$ the *Post Correspondence Problem* (PCP) is to find a finite index list $(i_1, \dots, i_k)$, $1 \leq i_j \leq n$, so that $u_{i_1} u_{i_2} \dots u_{i_k} = v_{i_1} v_{i_2} \dots v_{i_k}$.

1.3.7 Theorem (Post 1946)

PCP is undecidable.

3.15.1 Theorem (First-Order Unsatisfiability is Undecidable)

Unsatisfiability of a set of first-order clauses is undecidable.

Proof.

(Construction) By a reduction of PCP, Definition 1.3.6, Theorem 1.3.7. So let $(u_1, \dots, u_n)$ and $(v_1, \dots, v_n)$ be finite words over alphabet $\{a, b\}$. The first-order signature contains two unary functions g_a and g_b, a constant ϵ, and a relation R of arity two, all over some sort S. Then a word pair u_i, v_i is encoded by first-order clauses

$$\neg R(x, y) \lor R(u_i(x), v_i(y))$$

where $u_i(x)$ and $v_i(x)$ stand for the encodings of the words through respective nested occurrences of g_a and g_b. The intended meaning of $R(x, y)$ is that the word pair x, y can be derived from the PCP. □

Proof.

(Ctd)
The empty pair is encoded by the ground clause

$$R(\epsilon, \epsilon)$$

but it is a trivial solution to the PCP that needs to be ruled out. This is done by the two clauses

$$\neg R(g_a(x), g_a(x)), \qquad \neg R(g_b(x), g_b(x)).$$

I call the clause set consisting of these clauses N. Now the PCP over the two word lists has a solution iff N is unsatisfiable. □