Functional Programming Languages

Prof. Dr. Peter Thiemann and Prof. Dr. Stefan Wehr

Universität Freiburg

November 2022

Signatures

Definition

- A signature Σ is a set of function symbols, where each $f \in \Sigma$ is associated with a natural number n called the arity of f.
- $\Sigma^{(n)}$ denotes the set of all n-ary elements of Σ.
- The elements of $\Sigma^{(0)}$ are also called constant symbols.

Example

Signature Σ_{pred} for predicate logic

$$\Sigma_{pred} = \{\mathbf{T}^{(0)}, \mathbf{F}^{(0)}, \neg^{(1)}, \wedge^{(2)}, \vee^{(2)}\}$$
$$\Sigma_{pred}^{(0)} = \{\mathbf{T}, \mathbf{F}\}$$
$$\Sigma_{pred}^{(1)} = \{\neg\}$$
$$\Sigma_{pred}^{(2)} = \{\wedge, \vee\}$$

Terms

Definition

Let Σ be a signature and X a set of variables such that $\Sigma \cap X = \emptyset$. The set $T(\Sigma, X)$ of all Σ-terms over X is inductively defined as

- $X \subseteq T(\Sigma, X)$,
- for all $n \in \mathbb{N}$, all $f \in \Sigma^{(n)}$, and all $t_1, \ldots, t_n \in T(\Sigma, X)$, we have $f(t_1, \ldots, t_n) \in T(\Sigma, X)$

Note:

- For a constant symbol $f \in \Sigma^{(0)}$, we often write the term $f()$ as f.
- From now on, we leave to variable set $X = \{x, x_1, x_2, \ldots, y, y_1, y_2, \ldots, z, z_1, z_2 \ldots\}$ implicit

Example

Suppose $\Sigma = \Sigma_{pred}$. Then

$$\vee(\neg(x_{42}), \wedge(\mathbf{T}, x_3)) \in T(\Sigma, X)$$

Example

Suppose $\Sigma = \Sigma_{pred}$. Then

$$\vee(\neg(x_{42}), \wedge(\mathbf{T}, x_3)) \in T(\Sigma, X)$$

Alternative notations

- Polish notation: $\vee \neg\, x_{42} \;\wedge\; \mathbf{T}\, x_3$
- Prefix notation as in Scheme: $(\vee\, (\neg\, x_{42})\, (\wedge\, \mathbf{T}\, x_3))$
- Infix notation (with implicit operator precedence): $\neg x_{42} \vee \mathbf{T} \wedge x_3$
- Tree notation

Unique Decomposition of Terms

Lemma

Let $t, s \in T(\Sigma, X)$ with $t = f(t_1, \ldots, t_n)$ and $s = g(s_1, \ldots, s_m)$. If $t = s$ then $f = g$, $n = m$, and $t_i = s_i$ for all $i \in \{1, \ldots, n\}$.

Proof. Clearly, $f = g$ and $t_1 \ldots t_n$ = $s_1 \ldots s_m$. We prove $n = m$ and $t_i = s_i$ by induction on the length k of $w = t_1 \ldots t_n$.

- Induction basis: $k = 0$, so nothing is to prove.
- Induction step: Suppose $k > 0$. Then $w = aw'$.
  - If $a \in X$ then $t_1 = a = s_1$ and $t_2 \ldots t_n = s_2 \ldots s_m$. The IH yields $n = m$ and $t_i = s_i$ for $2 \leq i \leq m$.
  - If $a \in \Sigma^{(p)}$ then $t_1 = a(t'_1, \ldots, t'_p)$ and $s_1 = a(s'_1, \ldots, t'_p)$. From the assumption $t_1 \ldots t_n$ = $s_1 \ldots s_m$ we get $t'_1 \ldots t'_p t_2 \ldots t_n = s'_1 \ldots s'_p s_2 \ldots s_m$. The IH now yields $n = m$ and $t_i = s_i$ for $2 \leq i \leq n$ and $t'_j = s'_j$ for $1 \leq j \leq p$. Hence, we also have $t_1 = s_1$. □

Positions and Size of Terms

Definition

Suppose $t \in T(\Sigma, X)$.

- The set of positions of term t is a set $\mathcal{P}os(t)$ of strings over the alphabet of natural numbers. It is inductively defined as follows:
  - If $t = x \in X$, then $\mathcal{P}os(t) := \{\epsilon\}$
  - If $t = f(t_1, ..., t_n)$, then

$$\mathcal{P}os(t) := \{\epsilon\} \cup \bigcup_{i=1}^{n} \{ip | p \in \mathcal{P}os(t_i)\}$$

- The position ϵ is called the root position of t, the function or variable at this position is called the root symbol of t.
- The size $|t|$ of t is the cardinality of $\mathcal{P}os(t)$.

Ordering Positions

Definition

We define an ordering $\preceq \subseteq \mathcal{P}os(t) \times \mathcal{P}os(t)$ inductively by

- $\epsilon \preceq p$, for all p;
- $ip \preceq jq$, if either $i < j$ or $i = j$ and $p \preceq q$.

(lexicographic ordering)

Subterms and Replacing

Definition (Subterm)

For $p \in \mathcal{P}os(t)$, the subterm of t at position p, denoted by $t|_p$, is defined by induction on the length of p:

$$\begin{aligned} t|_\epsilon &:= t \\ f(t_1, \ldots, t_n)|_{ip} &:= t_i|_p \end{aligned}$$

($ip \in \mathcal{P}os(t)$ implies that $t = f(t_1, \ldots, t_n)$ with $0 \leq i \leq n$.)

Subterms and Replacing

Definition (Subterm)

For $p \in \mathcal{P}os(t)$, the subterm of t at position p, denoted by $t|_p$, is defined by induction on the length of p:

$$
\begin{aligned}
t|_\epsilon &:= t \\
f(t_1, \ldots, t_n)|_{ip} &:= t_i|_p
\end{aligned}
$$

($ip \in \mathcal{P}os(t)$ implies that $t = f(t_1, \ldots, t_n)$ with $0 \leq i \leq n$.)

Definition (Replacing)

For $p \in \mathcal{P}os(t)$, we denote by $t[s]_p$ the term that is obtained from t by replacing the subterm at position p by s, i.e.

$$
\begin{aligned}
t[s]_\epsilon &:= s \\
f(t_1, \ldots, t_n)[s]_{ip} &:= f(t_1, ..., t_i[s]_p, \ldots, t_n)
\end{aligned}
$$

Examples

Suppose $t = \vee(\neg(x_{42}), \wedge(\mathbf{T}, x_3))$

- t in tree notation with position annotations:
- $\mathcal{P}os(t) = \{\epsilon, 1, 12, 2, 21, 22\}$
- $|t| = 6$ (number of nodes in the tree)
- $t|_2 = \wedge(\mathbf{T}, x_3)$
- $t[\neg(\mathbf{F})]|_2 = \vee(\neg(x_{42}), \neg(\mathbf{F}))$

An Induction Principle for Terms

Term Induction

To prove that a predicate P holds for all $t \in T(\Sigma, X)$, we have to show the following properties:

- Induction basis
  $P(x)$ holds for all $x \in X$ and $P(f)$ holds for all $f \in \Sigma^{(0)}$.
- Induction step
  Suppose $n > 0$, $f \in \Sigma^{(n)}$, and $t_1, \ldots, t_n \in T(\Sigma, X)$.
  Then $P(f(t_1, \ldots, t_n))$ holds assuming $P(t_1), \ldots, P(t_n)$.

Note: Term Induction can be seen as an instance of ordinary induction using the term size as the induction variable.

Example for Term Induction

Lemma

For all terms t, the set $\mathcal{P}os(t)$ is prefix closed, i.e. if $wv \in \mathcal{P}os(t)$ then $w \in \mathcal{P}os(t)$.

Proof. We prove the lemma by term induction.

- Basis: If $t = x$ or $t = f$ (where f is a constant), then $\mathcal{P}os(t) = \{\epsilon\}$, which is prefix closed.
- Step: see next slide

Example for Term Induction (cont.)

- Step: Suppose $t = f(t_1, \dots, t_n)$ with $f \in \Sigma^{(n)}$ for $n > 0$ and $\mathcal{P}os(t_i)$ is prefix closed for all $i = 1, \dots, n$. We have to show that

$$\mathcal{P}os(t) := \{\epsilon\} \cup \bigcup_{i=1}^{n} \{ip | p \in \mathcal{P}os(t_i)\}$$

  is prefix closed as well.
  Let $w \in \mathcal{P}os(t)$ and suppose that w' is a prefix of w, i.e. $w = w'w''$. We proceed by case distinction on the form of w.

  - Case $w = \epsilon$: Then $w' = \epsilon \in \mathcal{P}os(t)$.
  - Case $w = jv$: Then $v \in \mathcal{P}os(t_i)$ for some $i \in \{1, \dots, n\}$ and $v = v'v''$ with $w' = jv'$. By the IH, we get $v' \in \mathcal{P}os(t_i)$. Hence, $jv' \in \mathcal{P}os(t)$.

Substitutions

Definition

Let Σ be a signature.

- A $T(\Sigma, X)$-substitution is a function $\sigma : X \to T(\Sigma, X)$ such that $\sigma(x) \neq x$ for only finitely many xs.
- The domain of σ is $\mathcal{D}om(\sigma) := \{x \in X \mid \sigma(x) \neq x\}$.
- We write $\{x_1 \mapsto t_1, \ldots, x_n \mapsto t_n\}$ for a substitution that maps x_i to t_i and has domain $\mathcal{D}om(\sigma) = \{x_1, \ldots, x_n\}$.
- We write $\mathcal{S}ub(\Sigma, X)$ for the set of $T(\Sigma, X)$-substitutions.

Applying Substitutions to Terms

Definition

The extension of a $T(\Sigma, X)$-substitution σ to a mapping $\hat{\sigma} : T(\Sigma, X) \rightarrow T(\Sigma, X)$ on arbitrary terms is defined as follows:

- $\hat{\sigma}(x) := \sigma(x)$
- $\hat{\sigma}(f(t_1, \ldots, t_n)) := f(\hat{\sigma}(t_1), \ldots, \hat{\sigma}(t_n))$

Note

- We usually drop to distinction between σ and $\hat{\sigma}$.
- Applying the extension of a substitution σ to a term simultaneously replaces all occurrences of a variable by their respective σ-image

Example

A substitution on terms from $T(\Sigma_{pred}, X)$

$$
\begin{aligned}
\Sigma &= \Sigma_{pred} \\
\sigma &= \{x \mapsto \neg z, y \mapsto x \vee \mathbf{F}\} \\
t &= x \vee y \wedge z \\
\sigma(t) &= \neg z \vee (x \vee \mathbf{F}) \wedge z
\end{aligned}
$$

Composing Substitutions

Definition

The composition $\sigma\tau$ of two substitutions σ and τ is defined as $\sigma\tau(x) := \sigma(\tau(x))$.

Lemma

Composition of substitutions is an associative operation where the identity substitution is the unit.

Lemma

The extension of a composition $\sigma\tau$ is just the composition of the extensions; i.e. $\widehat{\sigma\tau} = \hat{\sigma}\hat{\tau}$.

Towards Reduction

Definition

Let Σ be a signature. A Σ-identity is a pair $(s, t) \in T(\Sigma, X) \times T(\Sigma, X)$. We write identities as $s \approx t$ and call s its left-hand side and t its right-hand side.

Reduction Relation

Definition

Let E be a set of Σ-identities. Define the reduction relation $\longrightarrow_E \; \subseteq T(\Sigma, X) \times T(\Sigma, X)$ by

$s \longrightarrow_E t$ if and only if

$$\exists (l, r) \in E, p \in \mathcal{P}os(s), \sigma \in \mathcal{S}ub(\Sigma, X),$$
$$s|_p = \sigma(l) \text{ and } t = s[\sigma(r)]_p.$$

Call $s|_p$ the redex and $\sigma(r)$ the reductum of the reduction (step).

Example

$$G := \{f(x, f(y, z)) \approx f(f(x, y), z), f(e, x) \approx x, f(i(x), x) \approx e\}$$

Reduce

$$f(i(e), f(e, e))$$

Leftmost Reduction

Definition

Let E be a set of Σ-identities. Define the leftmost reduction relation $\longrightarrow^{l}_{E} \subseteq T(\Sigma, X) \times T(\Sigma, X)$ by $s \longrightarrow^{l}_{E} t$ if and only if

- $s \longrightarrow_{E} t$ at position $p \in \mathcal{P}os(s)$ and
- for all positions $q \in \mathcal{P}os(s)$ such that $s \longrightarrow_{E} t'$ it must be $p \preceq q$.