

Some Basic Properties of Concurrent Computations

CS 524
Lecture 22

J. Meseguer

Recall, from Lecture 11, that, given a rewrite theory, $R = (\Sigma, E, R)$ [more generally, we could add to R a freeness map $\phi: \Sigma \rightarrow \mathcal{P}(\mathbb{N})$ to restrict rewriting under some arguments, as described in a later lecture], a computation specification is a labeled arrow:

$$[u] \xrightarrow{\alpha} [v]$$

provable in the logic $\mathcal{L}(R)$ of R by the Reflexivity, Congruence, Replacement and Transitivity inference rules in pg. 6 of Lecture 11.

Notes: 1. Recall that, if a freeness map $\phi: \Sigma \rightarrow \mathcal{P}(\mathbb{N})$ is given, the Congruence and Replacement inference rules are then restricted to forbid rewriting under frozen positions.

2. A somewhat subtle point worth keep in mind is that, since for any $t' \in [t]_E$ in an E-equality equivalence class we have $[t] = [t']$, the Reflexivity rule could have been stated in the more explicit

form: Reflexivity $\frac{}{[t] \xrightarrow{t'} [t]}$ where $t' \in [t] \in T_{\Sigma/E}$

that is, any $t' \in [t]$ can be used for an idle transition.

Finally, recall that a computation of R has the form:

$$[u] \xrightarrow{[\alpha]} [v]$$

where $[\alpha]$ is an equivalence class of proof terms by the equivalences defined in lecture 11, and computations then form a category \mathcal{T}_R .

1. Interleaving Descriptions

Suppose R defines a Petri net with places a, b, c, d, e and transitions:

$$[1]: a b \rightarrow b b a$$

$$[2]: b c \rightarrow b c c$$

$$[3]: d e \rightarrow d e a$$

Then we have the concurrent computation:

$$[a b b c d e] \xrightarrow{[1] [2] [3]} [a a b b b d e]$$

but we also have; for any permutation σ of the set $\{1, 2, 3\}$ computations

$$[a b c d e] \xrightarrow{[\sigma(1)]u_1; [\sigma(2)]u_2; [\sigma(3)]u_3} [a a b b \overset{\text{cc}}{d} e]$$

such that $[1] [2] [3] = [[\sigma(1)]u_1; \sigma[2]u_2; [\sigma(3)]u_3]$

For example, for $\sigma(1)=2$, $\sigma(2)=3$, $\sigma(3)=1$

we have:

$$\begin{array}{l}
 [a b c d e] \xrightarrow{[2] a b d e} [a b c c d e] \xrightarrow{[3] a b c c} [a a b c c d e] \rightarrow \\
 \xrightarrow{[1] a c c d e} [a a b b c c d e]
 \end{array}$$

What all such descriptions, called interleaving descriptions, have in common is that the computation is described in a sequentialized form, so that a single transition takes place at each step in the sequential composition.

The obvious question to ask is: Is this always possible. An the answer is, Yes!

Sequentialization lemma. Given a rewrite theory $\mathcal{R} = (\Sigma, E, R)$, any computation $[u] \xrightarrow{[\alpha]} [v]$

in $\mathcal{T}_{\mathcal{R}}$ is either:

- $[\alpha] = [u]$ and $[u] = [v]$ (idle computation), or
- $[\alpha] = [\beta_1; \dots; \beta_n]$, $n \geq 1$, where each β

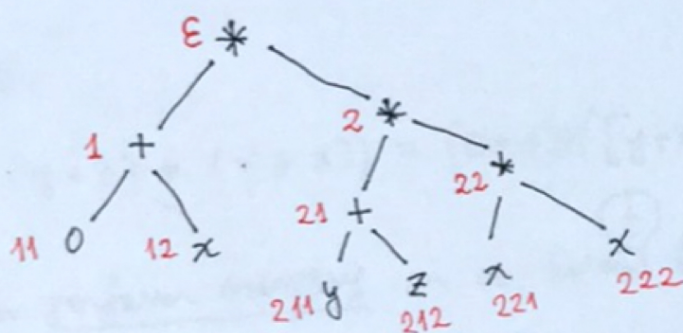
is a proof term of the form $t[l(u_1, \dots, u_n)]_p$ with $(l: u \rightarrow v) \in R$ involving variables x_1, \dots, x_n and $t, u_1, \dots, u_n \in \mathcal{T}_{\Sigma}$, and p a term position in t , as explained in more detail below. Such a proof term is called a one-step rewrite.

Term positions and related notation

We can easily describe the positions of a term by displaying it as a tree, and numbering its nodes as follows:

For example, consider the term: $(0+x) * ((y+z) * (x*x))$

Then its positions are:



That is, positions are strings in the monoid \mathbb{N}^*

Of course, we have a function

$$\text{pos} : T_{\Sigma} \longrightarrow \mathcal{P}(\mathbb{N}^*)$$

$$t \longmapsto \text{pos}(t)$$

assigning to each term t the set of its positions. For example:

$$\text{pos}((0+x) * ((y+z) * (x*x))) = \{ \epsilon, 1, 2, 11, 12, 21, 211, 212, 22, 221, 222 \}$$

Also, given a term t and a position $p \in \text{pos}(t)$ in it, then $t|_p$ denotes the subterm at that position. For

$$\text{example: } (0+x) * ((y+z) * (x*x))|_{21} = y+z$$

This notion of position allows us to decompose any term t such that $t|_p = u$ into a context and the subterm u "plugged" into such a context at position p . That is, t can be decomposed as:

$$t = t[u]_p \quad \text{when } u = t|_p$$

For example,

$$(0+x)*((y+z)*(x+x)) = (0+x)*([y+z]_{21}*(x+x))$$

We can also perform surgery on a term t by replacing a term $u = t|_p$ by another term v at the same position. For example, we may replace in the above term the subterm $x+y$ at position 21 by the term $z*(0+x)$ to get:

$$(0+x)*([z*(0+x)]_{21}*(x+x)) = 0+x*((z*(0+x))*(x+x))$$

In general terms, if $t = t[u]_p$ and v is another term, we can replace u by v at p to get $t[v]_p$.

Given two positions $p, q \in \text{Pos}(t)$ of a term t , any of the following exclusive possibilities can

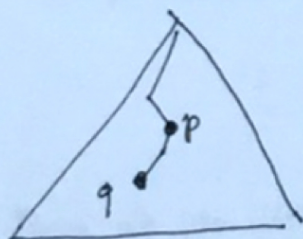
happen:

1. p and q [not necessarily distinct] by on the same path from the root of the term as a tree; we then call p and q nested positions, in the sense that either $p=q$, or the longer string, say q , is nested inside its prefix string p , i.e.

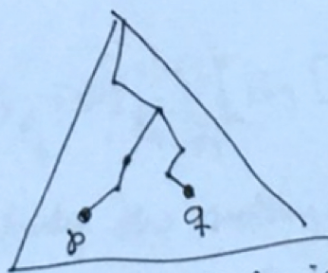
$$q = pq' \text{ for some } q',$$

maybe $q' = \varepsilon$, in case $p=q$

2. p and q are on different paths from the root. Then we call them disjoint positions.



nested positions



disjoint positions

Context notation can be generalized as follows:

1. If p_1, \dots, p_k are disjoint positions, we have a decomposition

$$t = \overset{\dagger}{t} [u_1]_{p_1} \dots [u_m]_{p_m} \quad \text{where } u_1 = t|_{p_1}, \dots, u_m = t|_{p_m}$$

If $q = p p'$ with $p' \neq \epsilon$ we also have
a decomposition

$$t = t [u [v]_{p'}]_p$$

where $u = t|_p$ and $v = t|_{pp'}$

Of course, this can be iterated to $p_1 p_2 \dots p_k$

with $t = t [u_1 [u_2 [\dots [u_k]_{p_k}]]_{p_2}]_{p_1}$

and $u_i = t|_{p_1 \dots p_i}$

and we can "mix and match" parallel decompositions
and nested ones, e.g.:

$$t [u_1 [u_2 [v]_{p_3} [v']_{p_4}]_{p_1} [u_3 [u_4 [v'']_{p_6} [v''']_{p_5}]_{p_2}]]_{p_2}$$

Now we can understand more clearly the notion of a
one-step rewrite step as a proof term of the form:

$$[t [u(v_1, \dots, v_n)]_p] \xrightarrow{t [l(v_1, \dots, v_n)]_p} [t [v(v_1, \dots, v_n)]_p]$$

where $l: u(x_1, \dots, x_n) \rightarrow v(x_1, \dots, x_n) \in R$

What this step does is to replace the term $u(v_1, \dots, v_n)$
at position p by the term $v(v_1, \dots, v_n)$ at same position

This notation has to be taken with a grain of salt, since, even if x_1, \dots, x_n appear in left-to-right order in $u(x_1, \dots, x_n)$, they may not appear that way in $v(x_1, \dots, x_n)$, since: (i) some variables may be dropped, and (ii) some variables may be repeated.

What $u(v_1, \dots, v_n)$ and $v(v_1, \dots, v_n)$ really mean is: $u\{x_1 \mapsto v_1, \dots, x_n \mapsto v_n\}$, $v\{x_1 \mapsto v_1, \dots, x_n \mapsto v_n\}$, but the shorthand notation $u(v_1, \dots, v_n)$ is of course simpler.

Sequential Systems

Sequential computation is a degenerate special case of concurrent computation. For example:

1. an automaton
 2. a Turing machine [deterministic or non-deterministic]
 3. a sequential programming language [with no parallelization optimizations such as speculative execution]
- are all sequential systems.

But if rewriting logic is a semantic framework for concurrent computation, it should be possible to characterize the sequential system subcase.

How? Quite simply:

Definition A rewrite theory $\mathcal{R} = (\Sigma, E, R)$ is called sequential if and only if:

1. It has no sideways parallelism, that is, no proof term of the form:

$$[u] \xrightarrow{\dagger [l(u_1, \dots, u_n)]_p [l'(v_1, \dots, v_n)]_q} [v]$$

with $l: u_i(x_1, \dots, x_n) \rightarrow u'(x_1, \dots, x_n)$, and p, q disjoint
 $l': v(x_1, \dots, x_m) \rightarrow v'(x_1, \dots, x_m)$ rules in R
 is possible at all.

2. It has no nested concurrency, i.e., no proof term of the form:

$$[u] \xrightarrow{\dagger [l(u_1, \dots, u_i [l'(v_1, \dots, v_n)]_q, \dots, u_n)]_p} [v]$$

with p and $p'q$ nested positions (for some p')
 is possible at all.

An obvious question is then the following:

If R is sequential, how do its computations in \mathcal{T}_R look like?

The answer can be given as follows:

Theorem. If $\mathcal{R} = (\Sigma, E, R)$ is sequential,
then, any computation $[u] \xrightarrow{[\alpha]} [v]$ has a
unique interleaving description as either

1. $[u] \xrightarrow{[u]} [v]$ idle computation, or

2. $[u] \xrightarrow{[\beta_1]} [u_1] \dots [u_{m-1}] \xrightarrow{[\beta_m]} [v]$
 $\underbrace{\hspace{15em}}_{[\alpha] = [\beta_1; \dots; \beta_m]} \quad m \geq 1$

where unique means up to E-equality, i.e.,

$$[\beta_1; \dots; \beta_m] =_E (\beta'_1; \dots; \beta'_m)$$

and of course up to associativity of ;: $(\beta_1; \beta_2); \beta_3 = \beta_1; (\beta_2; \beta_3)$

Proof Hints

To prove this theorem, a useful notion is that of an
idle proof term, which is a proof term built using
the inference rules of $L(R)$ using only the:

- Reflexivity
- Congruence, and
- Transitivity

rules, and showing that if $[u] \xrightarrow{\alpha} [v]$ with α an
idle proof term, then $[u] = [v]$ and $[\alpha] = [u]$.

Using this notion, the restrictions of a sequential \mathcal{R} , and
induction [structural or on depth] of proof terms, one can
prove this theorem.