Verifikation nebenläufiger Programme
Wintersemester 2004/05

Ulrich Hannemann   Jan Bredereke

15 Compositional Semantics of Synchronous Transition Diagrams

We want to apply a similar methodology for proving completeness in a compositional way as presented in Session 11. However, here the problem arises that we cannot even define the initial-final state behaviour of a location of a component of a parallel diagram in isolation, since input-output transitions are defined only in terms of the context of a parallel composition. Therefore we first define a *labelled* transition relation which also provides a semantics for input-output transitions, when considered in isolation, by defining their communication capabilities.

Definition 15.1 Let $P \equiv (L, T, s, t)$ be a sequential synchronous transition diagram. We define

$$\langle l; \sigma \rangle \xrightarrow{\langle\rangle} \langle l'; \sigma' \rangle$$

for $\sigma \models b$ and $\sigma' = f(\sigma)$, in case of an internal transition $l \xrightarrow{a} l' \in T$, with $a \equiv b \to f$, and where "$\langle\rangle$" expresses the empty sequence of communications.

In the case of an output transition $l \xrightarrow{a} l' \in T$, $a \equiv b; C!e \to f$, we have

$$\langle l; \sigma \rangle \xrightarrow{\langle (C,v) \rangle} \langle l'; \sigma' \rangle$$

if $\sigma \models b(\sigma)$, where $v = e(\sigma)$ and $\sigma' = f(\sigma)$.

In the case of an input transition $l \xrightarrow{a} l' \in T$, with $a \equiv b; C?x \to f$, we define

$$\langle l; \sigma \rangle \xrightarrow{\langle (C,v) \rangle} \langle l'; \sigma' \rangle$$

if $\models b(\sigma)$, where $\sigma' = f(\sigma : x \mapsto v)$, for v an *arbitrary* value in VAL.

Furthermore, we have the following rules for computing the reflexive and transitive closure:

$$\langle l; \sigma \rangle \xrightarrow{\langle\rangle} \langle l; \sigma \rangle$$

and

$$\frac{\langle l; \sigma \rangle \xrightarrow{\theta} \langle l'; \sigma' \rangle \quad \langle l'; \sigma' \rangle \xrightarrow{\theta'} \langle l''; \sigma'' \rangle}{\langle l; \sigma \rangle \xrightarrow{\theta \cdot \theta'} \langle l''; \sigma'' \rangle} .$$

Here θ and θ' denote sequences of communications, where a communication is represented by a pair of the form (C, v), called *communication record*, with C a channel and v a value. (The operation of concatenation is denoted by '$\cdot$'; we will frequently use, e.g., $\theta \cdot (C, v)$ to abbreviate $\theta \cdot \langle (C, v) \rangle$.)

■

One of the main points in the above definition is that an input transition is modelled by *guessing* locally the value received (in case the corresponding boolean guard is true) and assigning it to the specified local variable, after which the corresponding local state transformation is executed. The information about the communication, i.e., the guessed value and the channel involved, is attached to the transition itself in the form of a communication record. As we will see below, this information can be used in a parallel context to select the 'right' guesses, i.e., those guesses which correspond with the actual value sent.

Using the above transition relation we can now define the semantics of a *sequential* synchronous transition diagram, in which the value received in an input transition is selected by local guessing.

Definition 15.2 Let P be a sequential synchronous transition diagram, and l occur in P. We define

$$\mathcal{O}_l(P) \stackrel{\text{def}}{=} \{(\sigma, \sigma', \theta) \mid \langle s; \sigma\rangle \stackrel{\theta}{\longrightarrow} \langle l; \sigma'\rangle\}.$$

Note that we can now define the initial-final state semantics of P as $\mathcal{O}_t(P)$, also simply expressed as $\mathcal{O}(P)$.

■

For a sequence of communications θ and a set of channels $cset \subseteq CHAN$, we define the *projection* of θ onto $cset$, expressed by $\theta{\downarrow}cset$, as the sequence obtained from θ by deleting all records with channels not in $cset$.

Definition 15.3 (Projection of a sequence of communications to a set of channels) One can define $\theta{\downarrow}cset$, the projection of θ onto a set of channels $cset$, by induction on the length of θ:

$$\langle\rangle{\downarrow}cset \stackrel{\text{def}}{=} \langle\rangle,$$
$$(\langle(C,\mu)\rangle \cdot \theta'){\downarrow}cset \stackrel{\text{def}}{=} \begin{cases} \langle(C,\mu)\rangle \cdot \theta'{\downarrow}cset, & \text{if } C \in cset, \\ \theta'{\downarrow}cset, & \text{otherwise.} \end{cases}$$

Also we need to define the set of channels occurring in a trace θ.

Definition 15.4 (Channels occurring in a trace) The set of channels occurring in a sequence of communications, or *trace* θ, notation $Chan(\theta)$, is defined by

- $Chan(\langle\rangle) \stackrel{\text{def}}{=} \emptyset$
- $Chan(\theta \cdot (C,\mu)) \stackrel{\text{def}}{=} Chan(\theta) \cup \{C\}$.

■

In the following definition we extend the above semantics to a parallel composition of sequential synchronous diagrams.

Definition 15.5 (Compositional semantics of synchronous diagrams) Let $P \equiv P_1 \parallel \ldots \parallel P_n$ and $l = \langle l_1, \ldots, l_n\rangle$, l_i occurring in P_i. We define

$$\mathcal{O}_l(P) \stackrel{\text{def}}{=} \{(\sigma, \sigma', \theta) \mid (\sigma, \sigma'_i, \theta_i) \in \mathcal{O}_{l_i}(P_i),\ i = 1, \ldots, n\},$$

where

$$\sigma'(x) \stackrel{\text{def}}{=} \begin{cases} \sigma_i'(x), \text{ if } x \in var(P_i), \text{ for some } i = 1, \ldots, n, \\ \sigma(x), \text{ if } x \notin var(P_i), \text{ for all } i = 1, \ldots, n, \end{cases}$$

and θ_i denotes the projection $\theta{\downarrow}Chan(P_i)$ of θ along the channels of P_i. Similarly as above, the initial-final state semantics of a system of synchronous diagrams $P_1\| \ldots \|P_n$ is given by $\mathcal{O}_t(P_1\| \ldots \|P_n)$, which is also expressed by $\mathcal{O}(P_1\| \ldots \|P_n)$.

■

Alternatively, one could have defined the parallel operator by giving the definition of $P_1\|P_2$, proved its associativity and commutativity, and then observed that the meaning of $(\ldots((P_1\|P_2)\|P_3)\ldots P_n)$ amounts to the one given above.

Below we will refer to the compositional semantics defined above only as $\mathcal{O}$. We observe that in the above definition the requirement that the local histories θ_i can be obtained as the projection of one global history θ guarantees that an input on a channel indeed can be synchronised with a corresponding output. The following example illustrates this point.

Example 15.6 Consider the following parallel diagram:

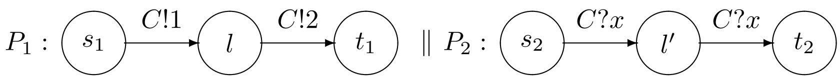

Now the process P_2 generates histories of the form $\langle(C, v), (C, w)\rangle$, where v and w are arbitrary values. On the other hand P_1 generates the history $\langle(C, 1), (C, 2)\rangle$. The requirement of the existence of a global history such that both the local histories of P_1 and P_2 can be obtained from it by projection along channel C thus restricts the choice of possible histories of P_2 to the 'right' one: $\langle(C, 1), (C, 2)\rangle$.

Consider next the diagram

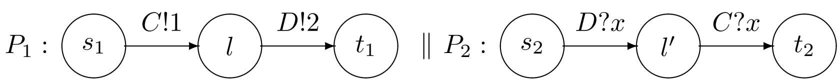

This diagram obviously deadlocks. This is also reflected by the fact that there exists no global history θ such that θ projected along the channels C and D equals each of the local histories of P_1 and P_2 (the local history of P_1 starts with a communication statement along C followed by one along D, whereas the local history of P_2 reverses the order of these communications).

■

The semantics $\mathcal{O}_l(P)$ also applies to the parallel composition $P_1\| \ldots \|P_n$ of sequential synchronous diagrams $P_1, \ldots, P_n$ in case the channels along which these processes communicate are *not* restricted to connecting processes from only the set $P_1, \ldots, P_n$. In this case $P_1\| \ldots \|P_n$ is called an *open* network. In the other case, in which every channel occurring in $P_1, \ldots, P_n$ connects exactly two different processes from $P_1, \ldots, P_n$, $P_1\| \ldots \|P_n$ is called *closed*.

It is important to realise that the above compositional semantics is consistent with our basic assumption that any communication involves only one sender and one receiver, under the condition that the channels are both uni-directional and one-to-one.

In fact, this semantics allows the generalisation of the one-to-one condition to channels with one sender and multiple receivers, while preserving unidirectionality. Consider, e.g., the network $C?x \parallel C?y \parallel C!0$, in which C connects two consumers with one producer. The above compositional semantics would generate for this network a global history which in fact models a multi-party communication interaction, i.e., the input produced is received by both consumers.

However, the condition of unidirectionality of the channels is necessary for this semantics to be consistent. Consider, e.g., the following network:

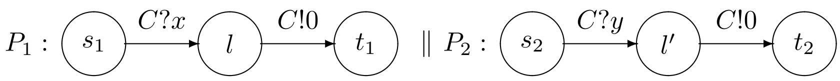

This network satisfies the one-to-one condition, since it connects exactly two processes, but violates unidirectionality of C. Both processes P_1 and P_2 act as producer and consumer with respect to C. It is easy to see that $\mathcal{O}_t(P_1 \parallel P_2)$ is nonempty, for $t = \langle t_1, t_2 \rangle$. This is due to the fact that the communication history does not indicate the direction of the communications, and consequently does not capture the different rôles of $C?x$ and $C!0$. However, according to our definition of closed product this diagram clearly deadlocks. Consequently the compositional semantics defined above is not correct with respect to the definition of closed product.

The following theorem states the correctness of the compositional semantics of a synchronous diagram $P \equiv P_1 \parallel \cdots \parallel P_n$, which does not contain any *external* channels, with respect to the initial-final state semantics $\mathcal{M}[\![P]\!]$. Here P denotes the closed product $P_1 \parallel \cdots \parallel P_n$. Its external channels are those channels which occur in some component P_i and which do not occur in the other components. The *internal* channels of such a system are those channels which connect two processes of that system.

Theorem 15.7 (Correctness of the compositional semantics)
Let $P_1 \parallel \cdots \parallel P_n$ be a synchronous diagram which does not contain external channels. We have that

$$\sigma' \in \mathcal{M}[\![P]\!]\sigma \quad \textit{iff} \quad \begin{array}{l}\textit{there exists a sequence of communications } \theta \\ \textit{such that } (\sigma, \sigma', \theta) \in \mathcal{O}_t(P),\end{array}$$

where P denotes the closed product $P_1 \parallel \cdots \parallel P_n$ and t denotes the exit location of P.

Proof
Let, for a location l of P, $\mathcal{M}_l[\![P]\!]\sigma$ denote all the resulting states of partial computations of P (viewed as a sequential transition diagram as formulated in Definition 13.1) which reach location l (note: $\mathcal{M}[\![P]\!]\sigma = \mathcal{M}_t[\![P]\!]\sigma$).

We first prove, for any location l of P, that the existence of a history θ such that $(\sigma, \sigma', \theta) \in \mathcal{O}_l(P)$ implies $\sigma' \in \mathcal{M}_l[\![P]\!]\sigma$. Let $l = \langle l_1, \ldots, l_n \rangle$. By the Definitions 15.1 and 15.5 we have that $(\sigma, \sigma', \theta) \in \mathcal{O}_l(P)$ iff $\langle s_i; \sigma \rangle \xrightarrow{\theta_i} \langle l_i; \sigma'_i \rangle$, for $i = 1, \ldots, n$, where s_i denotes the initial location of P_i, θ_i denotes the projection of θ along the channels of P_i, and σ'_i is obtained from σ' by assigning to the variables not belonging to process P_i their corresponding values in σ. If $\theta \neq \langle\rangle$ let the last communication of θ involve the channel C, with P_k and P_j the processes connected by C. Let $\theta = \theta' \cdot (C, v)$ and, for $i = 1, \ldots, n$, θ'_i denote the projection of θ' along the channels of P_i.

We proceed by induction on the sum of the lengths of the local computations

$$\langle s_i; \sigma\rangle \xrightarrow{\theta_i} \langle l_i; \sigma'_i\rangle.$$

For the starting location $s = \langle s_1, \ldots, s_n\rangle$ all local computations have a length of 0, $(\sigma, \sigma, \langle\rangle) \in \mathcal{O}_s(P)$ and $\sigma \in \mathcal{M}_s [\![P]\!] \sigma$.

We consider the different possibilities for the last transition that has *globally* taken place, separately – either it has been a local step of one process P_r or it has been the communication step along C.

Suppose first that there is an r, $1 \leq r \leq n$ such that the local computation of P_r ends with an internal transition, i.e., $\langle s_r; \sigma\rangle \xrightarrow{\theta_r} \langle l''_r; \tau\rangle \xrightarrow{\langle\rangle} \langle l_r; \sigma'_r\rangle$, for some location l''_r of P_r and state τ.

Let σ'' be obtained from σ' by assigning to all the variables of P_r their corresponding values in τ, i.e.,

$$\sigma''(x) \stackrel{\text{def}}{=} \begin{cases} \tau(x) & \text{if } x \in \mathit{var}(P_r), \\ \sigma'(x) & \text{if } x \notin \mathit{var}(P_r). \end{cases}$$

Moreover, let σ''_i, for $i = 1, \ldots, n$, be obtained from σ'' by assigning to the variables not belonging to process P_i their corresponding values in σ. We observe that $\sigma''_r = \tau$ and $\sigma''_i = \sigma'_i$, for $i \neq r$. Let $l''_i \stackrel{\text{def}}{=} l_i$, for $i \neq r$ and let $l'' \stackrel{\text{def}}{=} \langle l''_1, \ldots, l''_n\rangle$. For $i = 1, \ldots, n$, there exist computations $\langle s_i; \sigma\rangle \xrightarrow{\theta_i} \langle l''_i; \sigma''_i\rangle$, hence $(\sigma, \sigma'', \theta) \in \mathcal{O}_{l''}(P)$.

Since we have that the sum of the lengths of the following local computations $\langle s_i; \sigma\rangle \xrightarrow{\theta_i} \langle l''_i; \sigma''_i\rangle$, $i = 1, \ldots, n$, is one smaller than the sum of the lengths of the computations leading to l, we can apply the induction hypothesis: $\sigma'' \in \mathcal{M}_{l''} [\![P]\!] \sigma$, from which we derive by the internal transition $\langle l''_r; \sigma''_r\rangle \xrightarrow{\langle\rangle} \langle l_r; \sigma'_r\rangle$ and Definition 13.1 that $\sigma' \in \mathcal{M}_l [\![P]\!] \sigma$.

Suppose now that both the local computations of P_k and P_j end with a communication. Since both θ_k and θ_j are projections of θ along the channels of P_k and P_j, respectively, we have that

- for some state τ_k and location l''_k of P_k, $\langle s_k; \sigma\rangle \xrightarrow{\theta'_k} \langle l''_k; \tau_k\rangle$ and $\langle l''_k; \tau_k\rangle \xrightarrow{\langle(C,v)\rangle} \langle l_k; \sigma'_k\rangle$, with $\theta_k = \theta'_k \cdot (C, v)$, and, similarly,
- for some state τ_j and location l''_j of P_j, $\langle s_j; \sigma\rangle \xrightarrow{\theta'_j} \langle l''_j; \tau_j\rangle$ and $\langle l''_j; \tau_j\rangle \xrightarrow{\langle(C,v)\rangle} \langle l_j; \sigma'_j\rangle$, with $\theta_j = \theta'_j \cdot (C, v)$.

Let σ'' be obtained similarly as above from σ' by assigning to all the variables of P_k and P_j their corresponding values in τ_k and τ_j, respectively.

Moreover, let σ''_i, for $i = 1, \ldots, n$, be obtained from σ'' by assigning to the variables not belonging to process P_i their corresponding values in σ. We observe that $\sigma''_k = \tau_k$, $\sigma''_j = \tau_j$ and $\sigma''_i = \sigma'_i$, for $i \neq k$, $i \neq j$.

Let $l''_i \stackrel{\text{def}}{=} l_i$, for $i \neq k$, $i \neq j$ and let $l'' \stackrel{\text{def}}{=} \langle l''_1, \ldots, l''_n\rangle$. There exist computations $\langle s_i; \sigma\rangle \xrightarrow{\theta'_i} \langle l''_i; \sigma''_i\rangle$, for all $i = 1, \ldots, n$. Consequently, $(\sigma, \sigma'', \theta') \in \mathcal{O}_{l''}(P)$.

Since the sum of the lengths of these local computations is smaller than the sum of the lengths of the computations leading to location l, we can apply the induction hypothesis: $\sigma'' \in \mathcal{M}_{l''} [\![P]\!] \sigma$, from which we derive by Definition 13.1 and the transitions $\langle l''_k; \sigma''_k \rangle \stackrel{\langle (C,v) \rangle}{\longrightarrow} \langle l_k; \sigma'_k \rangle$ and $\langle l''_j; \sigma''_j \rangle \stackrel{\langle (C,v) \rangle}{\longrightarrow} \langle l_j; \sigma'_j \rangle$ that $\sigma' \in \mathcal{M}_l [\![P]\!] \sigma$.

Conversely, that $\sigma' \in \mathcal{M}_l [\![P]\!] \sigma$ implies $(\sigma, \sigma', \theta) \in \mathcal{O}_l(P)$ for some θ, can be proved by a straightforward induction on the length of the computation of P (viewed as a sequential transition diagram). This proof itself can be left as an exercise.

■