Verifikation nebenläufiger Programme
Wintersemester 2004/05
Ulrich Hannemann   Jan Bredereke

11 The Proof Method of Owicki & Gries

11.1 Completeness of the Method of Owicki & Gries

To establish the fact that $\models \{\varphi\}P_1\|\dots\|P_n\{\psi\}$ implies that $\{\varphi\}P_1\|\dots\|P_n\{\psi\}$ can be derived using the method of Owicki & Gries, one needs to define an appropriate assertion network.

The predicates of this network consist of strongest l-conditions of P_i w.r.t. precondition φ at label l, which are defined using a compositional semantics based on process-indexed sequences of states.

In order to define these concepts formally, we go through the following stages in a formal setting. We first formalise a simple, noncompositional, notion of initial-final-state behaviour $\mathcal{O}\llbracket P\rrbracket$ of a transition system P in a sequential setting, specialise this notion to the initial-final-state behaviour $\mathcal{O}_l\llbracket P\rrbracket$ at label l of P, and define strongest postconditions $SP(\varphi, P)$ and strongest l-conditions $SP_l(\varphi, P)$ w.r.t. this simple semantics. Due to the noncompositionality of this semantics w.r.t. concurrency these predicates are not interference free.

To obtain compositionality, we formally introduce *reactive sequences* [dBKPR91] in order to give a compositional characterisation $\mathcal{R}_l\llbracket P\rrbracket$ of $\mathcal{O}_l\llbracket P\rrbracket$. Using the compositional semantics $\mathcal{R}_l\llbracket P\rrbracket$ for shared-variable concurrency, a new notion of strongest postcondition, expressed by $SP'_l(\varphi, P)$ is defined. More precisely, first we consider the set of pairs $\langle\sigma, \theta\rangle$ with θ denoting a sequence of process-indexed states and σ the final state of the computation thus characterised, such that the projection $\theta[P](\sigma)$ of $\langle\sigma, \theta\rangle$ on P results in a sequence of state pairs describing transitions of P, i.e., such that $\theta[P](\sigma) \in \mathcal{R}_l\llbracket P\rrbracket$. Next the $SP'_l(\varphi, P)$ semantics is obtained by restricting this set to pairs $\langle\sigma, \theta\rangle$ of which the first state of θ satisfies φ in case θ is nonempty, and $\sigma \models \varphi$, if θ is empty.

Finally we prove that this choice of assertions satisfies the requirements imposed by the method of Owicki & Gries.

Strongest postcondition operators: sequential case

Definition 11.1 One defines the initial-final-state behaviour of transition system P by:

$$\mathcal{O}\llbracket P\rrbracket \overset{\text{def}}{=} \{(\sigma, \sigma') \mid \langle s; \sigma\rangle \to^* \langle t; \sigma'\rangle\},$$

where $\to^*$ denotes the reflexive transitive closure of the computation-step relation $\to$ between configurations defined in Session 4, where s and t are the entry and exit labels of a diagram representing P. □

Note that for all programs P one has that

$$\mathcal{O}\llbracket P\rrbracket = \{(\sigma, \sigma') \,|\, \sigma' \in \mathcal{M}\llbracket P\rrbracket\,\sigma \text{ and } \sigma, \sigma' \in \Sigma\}.$$

For an always terminating program P one even has

$$\mathcal{M}\,[\![P]\!]\,\sigma = \{\sigma' \mid (\sigma,\sigma') \in \mathcal{O}\,[\![P]\!]\,\}.$$

Next we define the initial-final-state behaviour at a location l of a transition diagram P. Observe that we have to switch from transition systems to transition diagrams because labels can only be identified by their names in transition diagrams.

Definition 11.2 Given a location l of P, define:

$$\mathcal{O}_l\,[\![P]\!] \stackrel{\text{def}}{=} \{(\sigma,\sigma') \mid \langle s;\sigma\rangle \to^* \langle l;\sigma'\rangle\}. \qquad \square$$

Note that $\mathcal{O}\,[\![P]\!] = \mathcal{O}_t\,[\![P]\!]$.

To relate these semantic notions to predicates, we introduce the notions of *strongest postcondition and strongest l-condition,* respectively, of a transition diagram P w.r.t. a precondition φ and a label $l \in P$.

Definition 11.3 (Strongest postcondition) Given a transition system P and a precondition φ, the strongest postcondition of P with respect to φ, expressed by $SP(\varphi, P)$, is defined as:

$$SP(\varphi,P) \stackrel{\text{def}}{=} \{\sigma' \mid \text{there exists } \sigma \text{ such that } (\sigma,\sigma') \in \mathcal{O}\,[\![P]\!] \text{ and } \sigma \models \varphi\}. \quad \square$$

We have the following basic property of strongest postcondition semantics.

Lemma 11.4 (Characterising property of SP) For any transition system P and predicate φ one has

$$\models \{\varphi\}\ P\ \{SP(\varphi,P)\}$$

and whenever $\models \{\varphi\}\ P\ \{\psi\}$ then $\models SP(\varphi,P) \to \psi$.

Definition 11.5 Given a transition diagram P, a location l of P, and a precondition φ, the strongest l-condition of P w.r.t. φ, expressed by $SP_l(\varphi, P)$, is defined as:

$$SP_l(\varphi,P) \stackrel{\text{def}}{=} \{\sigma' \mid \text{there exists } \sigma \text{ such that } (\sigma,\sigma') \in \mathcal{O}_l\,[\![P]\!] \text{ and } \sigma \models \varphi\}. \quad \square$$

Observe that $SP_t(\varphi,P) = SP(\varphi,P)$. That SP_l is not compositional w.r.t. $\|$ can be seen as follows. Let $P \equiv P_1\|\ldots\|P_n$. By Lemma 11.4, $\models \{\varphi\}\ P\ \{SP(\varphi,P)\}$ and $\models \{\varphi\}\ P\ \{\psi\}$ imply $\models SP(\varphi,P) \to \psi$. To derive $\{\varphi\}\ P\ \{SP(\varphi,P)\}$ we consider inductive assertion networks for $\{\varphi\}\ P_i\ \{SP(\varphi,P_i)\}$, P_i being a component of $P, i = 1,\ldots n$. However, in general $\bigwedge_i SP(\varphi,P_i)$ does not imply $SP(\varphi,P)$.

This is the lesson of Example 8.2, since for P_i as defined in Figure 3 of Session 8, one has that $SP(y = 0, P_i)$ is equivalent to $y = 1$, whereas for $P \equiv P_1\|P_2$, $SP(y = 0, P)$ is equivalent to $y = 2$. Hence $\not\models \bigwedge_i SP(\varphi,P_i) \to SP(\varphi,P_1\|\ldots\|P_n)$.

Reactive sequences

Next we investigate a definition of strongest l-condition in terms of a compositional semantics $\mathcal{R}_l$ based on reactive sequences. The additional information to make SP_l compositional will be encoded by auxiliary variables.

For the formal definition of reactive sequences we introduce the following alternative representation of a transition step.

Definition 11.6 (Reactive sequences) $l \overset{\langle\sigma,\sigma'\rangle}{\longrightarrow} l'$ iff $\langle l;\sigma\rangle \rightarrow \langle l';\sigma'\rangle$. The following axiom and rule allow us to compute the reflexive, transitive closure of $\overset{\langle\sigma,\sigma'\rangle}{\longrightarrow}$ and generate so-called reactive sequences, i.e., sequences of pairs of states:

$$l \overset{\epsilon}{\rightarrow} l$$

and

$$\frac{l \overset{w}{\rightarrow} l', l' \overset{w'}{\rightarrow} l''}{l \overset{w\cdot w'}{\rightarrow} l''}.$$

Here w and w' denote reactive sequences, ϵ the empty sequence, and the operation of concatenation is denoted by "$\cdot$". □

A reactive sequence models a computation of a transition diagram which takes into account possible interleavings by (other) parallel processes. These possible interleavings are made room for by "gaps", that is, subsequent pairs are such that the final state of the preceding pair differs from the initial state of the following pair, allowing for the insertion of interleaved pairs.

Now one can define the following compositional characterisation of $\mathcal{O}_l\,[\![P]\!]$.

Definition 11.7 $\mathcal{R}_l\,[\![P]\!] \overset{\text{def}}{=} \{w \mid s \overset{w}{\rightarrow} l\}$. □

Note that $\mathcal{O}_l\,[\![P]\!]$ is obtained from $\mathcal{R}_l\,[\![P]\!]$ by taking the initial state of the first pair and the final state of the last pair of *connected* sequences. Here connectedness means the absence of gaps in sequences, i.e., the final state of a preceding pair is the initial state of the next pair.

Next we need a definition of the interleaving operator $\tilde{\|}$ between sequences over some alphabet A.

Definition 11.8 (Interleaving) Let $a = a_k \ldots a_{k'}$, and $b = b_l \ldots b_{l'}$, be finite sequences over some alphabet A, with $k' \geq k-1$, and $l' \geq l-1$, and with $a = \epsilon$ if $k' = k-1$ and $b = \epsilon$ if $l' = l-1$. Then the operation of *interleaving* $\tilde{\|}$ the finite sequences a and b is defined as follows:

$$a\tilde{\|}b \overset{\text{def}}{=} \begin{cases} \{a\}, \text{ if } b = \epsilon, \\ \{b\}, \text{ if } a = \epsilon, \\ \{a_k \cdot (a_{k+1} \ldots a_{k'}\tilde{\|}b)\} \cup \{b_l \cdot (a\tilde{\|}b_{l+1} \ldots b_{l'})\}, \text{ otherwise,} \end{cases}$$

where for $a_i \in A$ and $s = s_1 \ldots s_m$, $a_i \cdot \{s \mid s \text{ satisfies } p\} \overset{\text{def}}{=} \{a_i \cdot s \mid s \text{ satisfies } p\}$, and $a_i \cdot s \overset{\text{def}}{=} a_i s_1 \ldots s_m$. □

The interleaving operator can be extended to sets of sequences by defining their pointwise extension $S\tilde{\|}T \stackrel{\text{def}}{=} \{s\tilde{\|}t \mid s \in S, t \in T\}$, and is commutative and associative, as, e.g., proved in [BK84].

Next we observe that $\mathcal{R}_l\, [\![P]\!]$ is compositional w.r.t. parallel composition.

Theorem 11.9 (Compositionality of interleaving)
Let $l = \langle l_1, \ldots, l_n\rangle$ be a location of $P \equiv P_1\| \ldots \|P_n$, then

$$\mathcal{R}_l\, [\![P]\!] = \mathcal{R}_{l_1}\, [\![P_1]\!] \,\tilde{\|}\, \ldots \tilde{\|}\mathcal{R}_{l_n}\, [\![P_n]\!]\,,$$

where $\tilde{\|}$ denotes the operation of interleaving.

■

This theorem follows immediately from the definition of $\mathcal{R}$. Observe that $\mathcal{R}_l\, [\![P]\!]$ extends and generalises the information given by $\mathcal{O}_l\, [\![P]\!]$. This is an example of a general principle that *compositional semantics are obtained from noncompositional ones by adding missing information in order to reconstruct the functional dependency between the semantics of a construct and the semantics of its components.*

Strongest postcondition operators: concurrent case

Next we define a strongest postcondition semantics based on the compositional semantics $\mathcal{R}$. To this end we introduce histories which record the sequence of state changes together with the active components responsible for these changes. Formally:

Definition 11.10 (Computation history) Given a set of process indices $(i \in)I$, a history θ is a sequence of indexed states (i, σ) indicating that process P_i makes a computation step in state σ. □

Remark 11.11 Given a sequence s and an element v, $s \cdot v$ denotes the sequence resulting from appending v to s. □

We also need the following projection operator which, given some set of sequential components and a final state, transforms a history into a reactive sequence consisting of all the computation steps involving one of the given components.

Definition 11.12 (Projection operator) We introduce the projection operator $\theta[I](\sigma)$ to denote the sequence of pairs of states which correspond to transitions of processes with indices from the set I:

$$\epsilon[I](\sigma) \stackrel{\text{def}}{=} \epsilon$$

$$(\theta \cdot (i, \sigma')\,)[I](\sigma) \stackrel{\text{def}}{=} \begin{cases} \theta[I](\sigma') \cdot \langle\sigma', \sigma\rangle, & \text{if } i \in I, \\ \theta[I](\sigma'), & \text{otherwise.} \end{cases} \qquad (1)$$

When the index set I contains the indices occurring in θ, then $\theta[I](\sigma)$ is connected. In that case the pair $\langle\sigma,\theta\rangle$ represents a connected reactive sequence with additional information about the identity of the active components. Let $\theta = (i_0,\sigma_0)\cdot\ldots\cdot(i_k,\sigma_k)\cdot\ldots\cdot(i_l,\sigma_l)$. Then process P_{i_m}, $0 \le m \le l$, transforms state σ_m into σ_{m+1}, where $\sigma_{l+1} = \sigma$.

Given a parallel program $P \equiv P_1 \parallel \cdots \parallel P_n$ we will identify below the set of indices $\{1,\ldots,n\}$ with the program P itself and the index i with its component P_i; correspondingly, one has notation $\theta[P](\sigma)$ for $\theta[\{1,\ldots,n\}](\sigma)$ and $\theta[P_i](\sigma)$ for $\theta[\{i\}](\sigma)$.

We are now able to introduce the strongest postcondition semantics based on the compositional semantics $\mathcal{R}$.

Definition 11.13 (Strongest l-condition for shared-variable concurrency) We associate with a location l of a transition system P the strongest l-condition of P based on the semantics $\mathcal{R}$:

$$SP'_l(\varphi,P) \overset{\text{def}}{=} \{\langle\sigma,\theta\rangle \mid \theta[P](\sigma) \in \mathcal{R}_l\,[\![P]\!] \textit{ and } \sigma' \models \varphi, \text{ with } \sigma' = \mathit{Init}(\sigma,\theta)\},$$

where

$$\mathit{Init}(\sigma,\theta) \overset{\text{def}}{=} \begin{cases} \text{the initial state of } \theta, & \text{if } \theta \text{ is non-empty,} \\ \sigma, & \text{otherwise.} \end{cases} \tag{2}$$

Note that if we restrict ourselves to histories referring only to the components of P, we have that

$$SP_l(\varphi,P) = \{\sigma \mid \text{there exists history } \theta \text{ referring only to the components of } P \text{ s.t. } \langle\sigma,\theta\rangle \in SP'_l(\varphi,P)\}.$$

Semantic completeness of the method of Owicki & Gries

In order to view the strongest l-condition of P as defined above as a state predicate (i.e., a predicate upon (program) states), we include histories of indexed states in the domain of values $\mathcal{D}$. A pair $\langle\sigma,\theta\rangle$ is then interpreted as a state which assigns the history θ to a distinguished history variable h, with all the other variables being assigned a value by σ.

Remark 11.14 For the sake of readability we express the application of a state-transformation f to a pair $\langle\sigma,\theta\rangle$ by $f(\sigma,\theta)$. □

Now we can prove completeness by associating with a location l of a component P_i of a parallel system $P \equiv P_1\|\ldots\|P_n$ the predicate $SP'_l(\varphi,P_i)$, *where we implicitly restrict ourselves to histories referring only to components of P.* In order to prove local correctness and interference freedom we augment the transition system as follows (we assume that the variable h does not already occur in P): Every transition $l \xrightarrow{a} l'$ in P_i, with $a = b \to f$, is transformed in a transition with action $b \to f \circ g$, where g is the (total) state transformation such that $g(\sigma,\theta) \overset{\text{def}}{=} \langle\sigma,\theta\cdot(i,\sigma)\rangle$, i.e., g only involves h. The resulting transition system is expressed by $P' \equiv P'_1 \parallel \ldots \parallel P'_n$. Note that we therefore associate with *each* location l of P'_i the predicate $SP'_l(\varphi,P_i)$, that is, the predicates of P'_i are defined in terms of P_i. In the following we show that the resulting assertion networks for P' are locally correct and free from interference.

Lemma 11.15 (Local correctness) Let P be a transition system in which the variable h does not already occur. For every transition $l \xrightarrow{a} l'$ of P' (the transition system P augmented with updates to the history variable h, as described above), we have

$$\models SP'_l(\varphi, P) \wedge b \to SP'_{l'}(\varphi, P) \circ f,$$

assuming $a = b \to f$.

Proof
Let $f(\sigma, \theta) = \langle \sigma', \theta' \rangle$. Note that $\theta' = \theta \cdot (i, \sigma)$, where i denotes the index of the active component of P', since f includes an update to the history variable as described above. Now, let $\langle \sigma, \theta \rangle \models SP'_l(\varphi, P) \wedge b$. Thus $\theta[P](\sigma) \in \mathcal{R}_l [\![P]\!]$. Since $\sigma \models b$ we derive that $\theta'[P](\sigma') = \theta[P](\sigma) \cdot \langle \sigma, \sigma' \rangle \in \mathcal{R}_{l'} [\![P]\!]$. Thus we conclude $\langle \sigma', \theta' \rangle \models SP'_{l'}(\varphi, P)$, or, equivalently, $\langle \sigma, \theta \rangle \models SP'_{l'}(\varphi, P) \circ f$.

■

Lemma 11.16 (Interference freedom) Let P and Q be two transition systems in which the variable h does not occur. Let $l \xrightarrow{a} l'$ be a transition of P' (the transition system P augmented with updates to the history variable h, as described above) and l'' be a location of Q, then:

$$\models SP'_{l''}(\varphi, Q) \wedge SP'_l(\varphi, P) \wedge b \to SP'_{l''}(\varphi, Q) \circ f,$$

assuming $a = b \to f$.

Proof
Actually we have already that

$$\models SP'_{l''}(\varphi, Q) \to SP'_{l''}(\varphi, Q) \circ f.$$

Let $\langle \sigma, \theta \rangle \models SP'_{l''}(\varphi, Q)$. By definition of SP' we have that $\theta[Q](\sigma) \in \mathcal{R}_{l''} [\![Q]\!]$. Moreover, we have that $f(\sigma, \theta) = \langle \sigma', \theta' \rangle$, where $\theta' = \theta \cdot (i, \sigma)$, with i the index of the active component of P'; and thus $\theta'[Q](\sigma') = \theta[Q](\sigma)$. From this we conclude $\langle \sigma', \theta' \rangle \models SP'_{l''}(\varphi, Q)$, i.e., $\langle \sigma, \theta \rangle \models SP'_{l''}(\varphi, Q) \circ f$.

■

Next we establish the remaining clauses of the method of Owicki & Gries.

Lemma 11.17 (Initialisation) Let $P \equiv P_1 \parallel \ldots \parallel P_n$ such that the variable h does not occur in P, and s_i denote the initial location of P_i. For any φ which does not refer to the variable h we have

$$\models \varphi \to \bigwedge_{i=1}^{n} SP'_{s_i}(\varphi, P_i) \circ f,$$

where $f(\sigma, \theta) \overset{\text{def}}{=} \langle \sigma, \epsilon \rangle$. (Here "$\wedge$" is assumed to bind stronger than "$\circ$", which binds in turn stronger than "$\to$".)

Lemma 11.18 (Finalisation) Let $P \equiv P_1 \parallel \ldots \parallel P_n$, let the variable h not occur in P, and t_i denote the final location of P_i. Furthermore suppose that

$\models \{\varphi\}P\{\psi\}$, where h does neither occur in φ nor in ψ. We have, restricting to histories which contain only references to components of P,

$$\models \bigwedge_{i=1}^{n} SP'_{t_i}(\varphi, P_i) \rightarrow \psi.$$

Proof
Let $\langle\sigma, \theta\rangle \models \bigwedge_i SP'_{t_i}(\varphi, P_i)$, and θ only refer to the components of P. It follows that $\theta[P_i](\sigma) \in \mathcal{R}_{t_i} [\![P_i]\!]$ for every $i \in \{1, \ldots, n\}$. Thus we obtain by the compositionality of $\mathcal{R}$ that $\theta[P](\sigma) \in \mathcal{R}_t [\![P]\!]$, where t denotes the final location of P, i.e., $t = \langle t_1, \ldots, t_n\rangle$. Next we observe that $\theta[P](\sigma)$ is a *connected* reactive sequence, since θ is assumed to contain only references to components of P. Thus it follows that $\sigma \in \mathcal{M} [\![P]\!] \sigma'$, where σ' is the initial state of $\theta[P](\sigma)$. Note furthermore that $\sigma' \models \varphi$, so we conclude by $\models \{\varphi\}P\{\psi\}$, that $\sigma \models \psi$ (and so, since h does not occur in ψ, $\langle\sigma, \theta\rangle \models \psi$).

■

Therefore we conclude:

Theorem 11.19 (Semantic completeness)
The proof method of Owicki & Gries is semantically complete.

■

It is interesting to observe that in the above completeness proof the compositionality of $\mathcal{R}$ is only used in the 'finalisation' lemma, namely in order to establish $\models \bigwedge_i SP'_{t_i}(\varphi, P_i) \rightarrow SP'_{\langle t_1,\ldots,t_n\rangle}(\varphi, P)$.

References

[BK84] J.A. Bergstra and J.W. Klop. Process algebra for synchronous communication. *Information and Control*, 60(1/3):109–137, 1984.

[dBKPR91] F.S. de Boer, J.N. Kok, C. Palamidessi, and J.J.M.M. Rutten. The failure of failures: towards a paradigm for asynchronous communication. In Baeten and Groote, editors, *CONCUR '91*, volume 527 of *LNCS*. Springer-Verlag, 1991.