

7 Verification by Rules

Next we look at an alternative approach to verification. We will use a rule-based reasoning to establish properties of computations.

We define *invariance property* as a property that can be specified by a formula $G\varphi$ where φ is a propositional formula. An invariance property p can be established by showing that

1. p holds initially and that
2. p is preserved during each step of all executions of the program.

Let the specifying formula for property p be φ_p . Thus we have to show that, for all computations σ ,

1. $(\mathcal{M}, \sigma) \models \varphi_p$ and
2. $(\mathcal{M}, \sigma) \models G(\varphi_p \rightarrow X\varphi_p)$.

Then, by an inductive argument, we have $(\mathcal{M}, \sigma) \models G\varphi_p$.

In general, we will consider *verification conditions* defined as follows. Let \mathcal{T} be a labelled transition system, φ and ψ be propositional formulas, and τ be a transition. The verification condition of φ and ψ relative to τ is the triple $\{\varphi\}\tau\{\psi\}$. We say that $\{\varphi\}\tau\{\psi\}$ holds if the following holds:

for all states s and t such that $s\tau t$, if φ is true at s , then ψ is true at t .

Note that if we can establish the verification condition $\{\varphi\}\tau\{\varphi\}$ for all transitions τ , then φ is preserved during all computations. Thus to prove that p is an invariance property, it suffices to show that

1. φ_p is true at all initial states and
2. the verification condition $\{\varphi_p\}\tau\{\varphi_p\}$ holds for all transitions τ .

7.1 Proving Invariance Properties

Let us fix a program P and an invariance property p with specifying formula φ_p . Let the labelled transition system associated with P be $\mathcal{T} = (S, L, R_a : a \in A)$ and the corresponding linear temporal model be $\mathcal{M} = (S, R, V)$. Assume that there are finitely many initial states in \mathcal{T} and that they can be defined by a propositional formula *init*, i.e., for every state s , s is an initial state iff *init* is true at s under the valuation V .

Consider the following rule.

$$\frac{\begin{array}{l} B1 \quad (\mathcal{M}, \sigma) \models \text{init} \rightarrow \varphi_p \\ B2 \quad \{\varphi_p\}R_a\{\varphi_p\} \text{ for all } a \in A \end{array}}{B3 \quad (\mathcal{M}, \sigma) \models G\varphi_p}$$

Rule INV-B

where σ is a computation. The basic invariance rule INV-B says, that the invariance formula φ_p is true in \mathcal{M} during computation σ if φ_p is true at every initial state and the verification condition $\{\varphi_p\}R_a\{\varphi_p\}$ holds for every transition R_a . Note that checking the verification condition involves checking all possible transitions between states, and it is not restricted to transitions between states *occurring in computations*.²

²In general, it is difficult to check which transitions between states occur in computations.

We can formulate Rule INV-B using temporal logic as follows.

$$\begin{array}{l} B1 \quad (\mathcal{M}, \sigma) \models \text{init} \rightarrow \varphi_p \\ B2 \quad (\mathcal{M}, \pi) \models \mathbf{G}(\varphi_p \rightarrow \mathbf{X}\varphi_p) \quad \text{for every run } \pi \\ \hline B3 \quad (\mathcal{M}, \sigma) \models \mathbf{G}\varphi_p \end{array}$$

Rule INV-B

An invariance formula φ is called *inductive* if both B1 and B2 in INV-B hold for φ .

Example 7.1 Consider the program MUX-SEM and the invariance formula $r_0 \vee r_1$. Define *init* as $\text{at}_{l_0} \wedge \text{at}_{m_0} \wedge r_1$. Thus *init* is true precisely at the initial state $(l_0, m_0, r = 1)$. Obviously, $(\mathcal{M}, \sigma) \models \text{init} \rightarrow r_0 \vee r_1$ for every run (computation) σ . Let $\text{move}(l_i, l_j)$ denote that process 1 makes a move from location l_i to l_j by executing some instruction (and similarly for process 2). Using the definition of a semaphore, it is routine to check that the invariance conditions

$$\{r_0 \vee r_1\} \text{move}(l_i, l_j) \{r_0 \vee r_1\} \quad \text{and} \quad \{r_0 \vee r_1\} \text{move}(m_i, m_j) \{r_0 \vee r_1\}$$

hold for any $i, j \in \{0, 1, 2, 3, 4\}$ (recall that a process cannot perform **request**(r) if $r = 0$). Thus we can apply INV-B to establish that the semaphore r remains non-negative during all computations.

Invariance properties are monotone, i.e., the following rule is valid: for (propositional) formulas φ and ψ ,

$$\begin{array}{l} (\mathcal{M}, \sigma) \models \mathbf{G}\varphi \\ (\mathcal{M}, \sigma) \models \mathbf{G}(\varphi \rightarrow \psi) \\ \hline (\mathcal{M}, \sigma) \models \mathbf{G}\psi \end{array}$$

Rule MON-I, monotonicity of invariance

If we have two invariant formulas φ and ψ , then their conjunction is invariant as well. That is, the following rule is valid:

$$\begin{array}{l} (\mathcal{M}, \sigma) \models \mathbf{G}\varphi \\ (\mathcal{M}, \sigma) \models \mathbf{G}\psi \\ \hline (\mathcal{M}, \sigma) \models \mathbf{G}(\varphi \wedge \psi) \end{array}$$

Rule CON-I, conjunction of invariances

A basic limitation to rule INV-B is when φ_p is invariant but not inductive. This happens when $(\mathcal{M}, \sigma) \models \mathbf{G}(\varphi \rightarrow \mathbf{X}\varphi)$ holds for all *computations* σ , but there is a *run* that falsifies the formula; see the following example.

Example 7.2 We want to use rule INV-B to establish that the program MUX-SEM has the *mutual exclusion* property, i.e., for every computation σ , we have $(\mathcal{M}, \sigma) \models \mathbf{G}\neg(\text{at}_{l_3} \wedge \text{at}_{m_3})$. But B2 requires that $\{\neg(\text{at}_{l_3} \wedge \text{at}_{m_3})\} \tau \{\neg(\text{at}_{l_3} \wedge \text{at}_{m_3})\}$ should hold for all transitions τ . This fails, however, for the transition $\text{move}(l_2, l_3)$ from state $(l_2, m_3, r = 1)$ to state $(l_3, m_3, r = 0)$.³

7.2 Assertion Strengthening

Consider the following rule. For all computations σ ,

$$\begin{array}{l} I1 \quad (\mathcal{M}, \pi) \models \mathbf{G}(\varphi \rightarrow \psi) \quad \text{for every run } \pi \\ I2 \quad (\mathcal{M}, \sigma) \models \text{init} \rightarrow \varphi \\ I3 \quad \{\varphi\} R_a \{\psi\} \quad \text{for all } a \in A \\ \hline I4 \quad (\mathcal{M}, \sigma) \models \mathbf{G}\psi \end{array}$$

Rule INV, general invariance

³Of course, the state $(l_2, m_3, r = 1)$, hence the transition, cannot occur in any computation, but this is what we would like to establish.

Rule INV states that if φ is inductive and φ implies ψ at each state in all runs, then ψ is true at each state in all computations. Again, we can replace I3 by:

$$\text{for every run } \pi, (\mathcal{M}, \pi) \models \mathbf{G}(\varphi \rightarrow \mathbf{X}\varphi).$$

Example 7.3 Assume we want to prove mutual exclusion $\psi : \neg(at_{l_3} \wedge at_{m_3})$ for MUX-SEM. Let ρ be the propositional formula expressing that precisely one of $at_{m_3} \vee at_{m_4}$, $at_{l_3} \vee at_{l_4}$, or r_1 holds. **DEFINE** ρ . We define φ as

$$(r_0 \vee r_1) \wedge \rho$$

We have to establish the following for every computation σ :

- I1: $(\mathcal{M}, \pi) \models \mathbf{G}(\varphi \rightarrow \psi)$ for every run π ,
- I2: $(\mathcal{M}, \sigma) \models \text{init} \rightarrow \varphi$,
- I3: $(\mathcal{M}, \pi) \models \mathbf{G}(\varphi \rightarrow \mathbf{X}\varphi)$ for every run π ,

so that we have

- $(\mathcal{M}, \sigma) \models \mathbf{G}\psi$

by rule INV.

Premise I1: The implication

$$[(r_0 \vee r_1) \wedge \rho] \rightarrow \neg(at_{l_3} \wedge at_{m_3})$$

is true at every state (since $(at_{l_3} \wedge at_{m_3})$ would contradict to ρ).

Premise I2: Obviously, the implication

$$(at_{m_0} \wedge at_{l_0} \wedge r_1) \rightarrow [(r_0 \vee r_1) \wedge \rho]$$

is true (since ρ is true in the initial state).

Premise I3: We have to show that

$$(\mathcal{M}, \pi_i) \models (r_0 \vee r_1) \wedge \rho \text{ implies } (\mathcal{M}, \pi_{i+1}) \models (r_0 \vee r_1) \wedge \rho$$

for every run π and $i \in \omega$. Consider the transitions of process 1. Transitions $move(l_0, l_1)$, $move(l_1, l_2)$ and $move(l_3, l_4)$ do not change the truth of ρ (since these transitions do not change the value of r , $(at_{l_3} \vee at_{l_4})$ and $(at_{m_3} \vee at_{m_4})$). Transition $move(l_2, l_3)$ can be made only if $r = 1$. But ρ implies that process 2 cannot be in its critical section in this case. Also, the value of r after the transition is 0, by the definition of **request**(r). It follows that ρ is true after the transition. Finally, transition $move(l_4, l_0)$ changes $at_{l_3} \vee at_{l_4}$ from true to false and r_1 from false to true (by the definition of **release**(r)), while the truth value of $at_{m_3} \vee at_{m_4}$ is unchanged. Hence ρ is “preserved” under all transitions of process 1. A similar argument shows that transitions of process 2 cannot change the truth value of ρ either.

Thus φ is true along all runs, and ψ is true during all computations.

Exercise 7.4 Show that MUX-SEM avoids deadlock, i.e., that at every moment at least one transition is enabled.

Exercise 7.5 The **exchange**(x, y) instruction is defined in Figure 8. Consider the program in Figure 9.

1. Show that MEX achieves mutual exclusion (i.e., only one of the n processes can be in the critical section at any time).

Hint: Let $\varphi \oplus \psi$ denote exclusive disjunction $(\neg\varphi \wedge \psi) \vee (\varphi \wedge \neg\psi)$ and $\bigwedge_i \gamma_i$ denote the conjunction $\gamma_0 \wedge \dots \wedge \gamma_i \wedge \dots \wedge \gamma_n$. Let $crit_i$ be an atomic proposition expressing that ‘process i is in its critical section’.

```

int temp

void exchange(int x, int y)
{
    temp = x;
    x = y;
    y = temp;
}

```

Figure 8: The **exchange** instruction

```

int n /* number of processes */
int bolt = 0
int key[n] = 1 /* a variable for each process */

void P(int i)
{
    while(true){
        while(key[i] != 0) exchange(key[i], bolt);
        critical;
        exchange(key[i], bolt);
        noncritical;
    }
}

```

Figure 9: MEX — mutual exclusion with **exchange**

- Show that bolt and all key[i] are either 0 or 1: that is,

$$(\text{bolt} = 0 \oplus \text{bolt} = 1) \wedge \bigwedge_i (\text{key}[i] = 0 \oplus \text{key}[i] = 1)$$

is inductive.

- Show that

$$\text{bolt} + \sum_i \text{key}[i] = n$$

is inductive.

- Show that ‘if process i is in its critical section, then key[i] = 0’: that is,

$$\text{crit}_i \rightarrow \text{key}[i] = 0$$

is inductive for every i .

- Using the above inductive formulas show that

$$\text{crit}_i \rightarrow (\text{bolt} = 1 \wedge \bigwedge_{j \neq i} \text{key}[j] = 1)$$

is invariant for every i .

- Using the above invariant formulas show that

$$\text{crit}_i \rightarrow \bigwedge_{j \neq i} \neg \text{crit}_j$$

is invariant for every i .

2. Show “progress”, i.e., ‘if there is no process in the critical section, then one process can enter the critical section’.

Hint: Define the formula $enabled(crit_i)$ as $key_i = 0 \vee bolt = 0$ and show that

$$\left(\bigwedge_i \neg crit_i\right) \rightarrow \left(\bigvee_i enabled(crit_i)\right)$$

is invariant.

Exercise 7.6 Recall Peterson’s solution, see Figure 10. There are two processes, identified by 0 and 1.

```

int t
int y0 = 0
int y1 = 0

void P(int i)
{
  L1 : while(true){
  L2 :  yi = 1;
  L3 :  t = i;
  L4 :  while{t = i & y1-i = 1}    /* loop/;
  L5 :  critical;
  L6 :  yi = 0;
  L7 :  noncritical;
  }
}

```

Figure 10: PET

1. Show that mutual exclusion holds for PET.

Hint: Let L_j^i be an atomic proposition for each $1 \leq j \leq 7$ and $0 \leq i \leq 1$ expressing that process i is at location L_j .

- Show that

$$(y_0 = 0 \oplus y_0 = 1) \wedge (y_1 = 0 \oplus y_1 = 1) \wedge (t = 0 \oplus t = 1)$$

is inductive.

- Show that

$$(L3_i \vee L4_i \vee L5_i \vee L6_i) \rightarrow y_i = 1$$

is inductive for both $i = 0$ and $i = 1$.

- Using the above inductive formulas show that

$$((L5_i \vee L6_i) \wedge L4_{1-i}) \rightarrow (t = 1 - i \wedge y_i = 1)$$

is invariant for $i = 0, 1$.

- Using the above invariant formulas show that

$$(L5_i \vee L6_i) \rightarrow \neg(L5_{1-i} \vee L6_{1-i})$$

is invariant for both $i = 0$ and $i = 1$.

2. Show that PET avoids “livelock”, i.e., at least one process can enter the critical section.

Hint: Define the formula $enabled(L5_i)$ as $L4_i \wedge (t \neq i \vee y_{1-i} \neq 1)$ and show that

$$(L4_0 \wedge L4_1) \rightarrow (enabled(L5_0) \vee enabled(L5_1))$$

is invariant.