

VALIDATION

Are we building the right program?
Correct

Does the design faithfully
capture the requirements?
Correct

VERIFICATION

Are we building the program right?
Correct




Does the design satisfy the

Labeled Transition Systems


STATE

snapshot of
the system "at a given time"



TRANSITION

evolution step of the system
"in time"



Transition System $TS = \langle S, Act, \rightarrow, I, AP, L \rangle$

Where: S is a set of states.

Act is a set of actions;

$\rightarrow \subseteq S \times Act \times S$ is a transition relation

$I \subseteq S$ a set of initial states

AP is a set of atomic propositions

$L: S \rightarrow \underbrace{2^{\text{AP}}}_{\text{powerset of AP}}$ a Labeling function


$$2^{\text{AP}} = \{A \mid A \subseteq \text{AP}\}$$

If Φ is a formula of propositional logic over $ad. \text{prop}$ in AP ,
then we write

$$S \models \Phi : \Leftrightarrow L(S) \models \Phi.$$

Example: Beverage Vending Machine

2



$$S = \{ \text{pay}, \text{select}, \text{soda}, \text{beer} \}$$

$$I = \{ \text{pay} \}$$

$$\text{Act} = \{ \text{accept-coin}, \tau, \text{get-soda}, \text{get-beer} \}$$

$$\rightarrow = \{ \langle \text{pay}, \text{accept-coin}, \text{select} \rangle, \langle \text{select}, \tau, \text{soda} \rangle, \langle \text{select}, \tau, \text{beer} \rangle, \langle \text{soda}, \text{give-soda}, \text{pay} \rangle, \langle \text{beer}, \text{give-beer}, \text{pay} \rangle \}$$

$\text{pay} \xrightarrow{\text{accept-coin}} \text{select}$, $\text{select} \xrightarrow{\tau} \text{soda}$,
 $\text{select} \xrightarrow{\tau} \text{beer}$, ...

$$AP = \{ \text{paid}, \text{drink} \}$$

$$L: S \rightarrow 2^{AP}$$

$\text{pay} \mapsto \emptyset$
$\text{soda}, \text{beer} \mapsto \{ \text{paid}, \text{drink} \}$
$\text{select} \mapsto \{ \text{paid} \}$

In general $L(s) := \{s\}$ is a simple reasonable choice for a Labelling

But e.g. if we only want to check properties that do not refer to the drink in question, like:

"The vending machine only delivers a drink after a coin is inserted"

Then the above labeling is fine.

Therefore in general: $L(s) := \{s\} \cap AP$.

Later in LTL:
 $\square (\text{drink} \rightarrow \text{paid})$

Direct predecessors and Successors

$$TS = \langle S, A_{\text{ct}}, \rightarrow, I, A_{\text{PL}} \rangle$$

For all $s \in S$ and $\alpha \in A$:

$$\text{Post}(s, \alpha) := \{s' \in S \mid s \xrightarrow{\alpha} s'\},$$

(α -successors of s)

$$\text{Pre}(s, \alpha) := \{s' \in S \mid s' \xrightarrow{\alpha} s\},$$

(α -predecessors of s)

$$\text{Post}(s) := \bigcup_{\alpha \in A_{\text{ct}}} \text{Post}(s, \alpha),$$

(successors of s)

$$\text{Pre}(s) := \bigcup_{\alpha \in A_{\text{ct}}} \text{Pre}(s, \alpha).$$

(predecessors of s)

Extensions to sets C of states:

For all $C \subseteq S$ and $\alpha \in A_{\text{ct}}$:

$$\text{Post}(C, \alpha) := \bigcup_{s \in C} \text{Post}(s, \alpha),$$


$$\text{Pre}(C, \alpha) := \bigcup_{s \in C} \text{Pre}(s, \alpha),$$

$$\text{Post}(C) := \bigcup_{\alpha \in A_{\text{ct}}} \text{Post}(C, \alpha),$$

$$\text{Pre}(C) := \bigcup_{\alpha \in A_{\text{ct}}} \text{Pre}(C, \alpha).$$

Terminal State:

$s \in S$ is called terminal : $\Leftrightarrow \text{Post}(s) = \emptyset$.



No terminal states

$$\begin{aligned} \text{Post}(\text{select}) &= \{ \text{soda}, \text{beer} \} \\ &= \text{Pre}(\text{pay}) \end{aligned}$$

$$\text{Pre}(\text{pay}, \text{give-beer}) = \{ \text{beer} \}$$

neither action- nor AP-deterministic

TS = $\langle S, \text{Act}, \rightarrow, I, \text{AP}, L \rangle$ where $\rightarrow \subseteq S \times \text{Act} \times S$, 14

$I \subseteq S$ (usually $I \neq \emptyset$)

$L: S \rightarrow 2^{\text{AP}}$ (usually $\text{AP} \subseteq S$
and often $L(s) = \{s\} \cap \text{AP}$)

Executions and Traces

Execution fragments are sequences $g \in S(\text{Act } S)^* \cup S(\text{Act } S)^\omega$

such that

if $g = s_0 a_1 s_1 a_2 s_2 a_3 s_3 \dots a_n s_n \dots \Rightarrow s_i \xrightarrow{a_{i+1}} s_i$ for all i

g is initial if $s_0 \in I$


g is maximal if $\begin{cases} g \text{ is infinite} \\ \text{or} \\ g = g_0 a_1 g_1 a_2 g_3 \dots a_n g_n \text{ and } \text{Post}(g_n) = \emptyset \end{cases}$

An execution is an initial and maximal execution fragment.

Reach(TS) = $\{s \in S / \exists s_0 \in I \exists g \text{ (initial) execution fragment } (g \text{ starts in } s_0 \text{ and ends in } s) \}$

reachable

states of TS



A trace of an execution fragment

$g = g_0 a_1 g_1 a_2 g_2 a_3 g_3 \dots$



is the sequence

$\tau = L(g_0) L(g_1) L(g_2) \dots$

initial execution fragment Π

$\langle \text{pay, a-c, select, } \tau, \text{beer, g-b, pay} \rangle$


$\langle \emptyset, \{ \text{paid} \}, \{ \text{paid, drinky} \}, \emptyset \rangle$ trace of Π



$M \parallel_H U$

Where $H = \{ \text{coin}, \text{soda-select}, \text{beer-select}, \text{get-soda}, \text{get-beer} \}$

via handshaking
communication



$$\begin{array}{c}
 \frac{S_1 \xrightarrow{a} M S_1' \quad S_2 \xrightarrow{a} U S_2'}{(if \ a \in H)} \quad \langle S_1, S_1' \rangle \rightarrow \langle \bar{S}_1, \bar{S}_1' \rangle \\
 \hline
 \langle S_1, S_1' \rangle \rightarrow \langle \bar{S}_1, \bar{S}_1' \rangle
 \end{array}$$

"Why do we need both actions and Labelling to express properties?"

Verification can be









$\left\{ \begin{array}{l} \text{action-based} \\ \text{state-based} \\ \text{action+state-based} \end{array} \right.$
(most-complex,
often too complex)

Execution-fragments are used for action-based verification, which is essential for modelling communication.

Notation

Given a sequence $\sigma = \sigma_0 \sigma_1 \sigma_2 \dots \in A^* \cup A^\omega$

- the length of σ is $|\sigma|$ (if σ is infinite, then $|\sigma| = \infty$)
- the i -th element of σ is $\sigma[i]$
- the last element of σ is $\text{last}(\sigma)$, provided that $\sigma \in A^*$
and hence is finite



Penie₄ || Stick₃ || Phile₃ || Stick₂ || Phile₂ || Stick₁ || Phile₁ || Stick₀
|| Penie₀ || Sticky


permits deadlock (e.g. if every philosopher picks up her right stick immediately)

Possibility to avoid deadlock:

8


make a stick only available for one philosopher at a time

change the stick-initial state for the 0, 2,



Example: A (simplified) 3-wheel slot machine


apple pear banana



3 wheels

set of fruits

Fruits



$$S = \{0, 1, 2, \dots, n+2\}$$

$$Act = \{\text{bet}, \text{pull}, \text{loose}, \text{win}, \text{payout}\}$$

For an interval $W = [i, j]$ with $2 \leq i \leq j \leq n+1$

$$\rightarrow = \{\langle 0, \text{bet}, 1 \rangle\} \cup \{1\} \times \{\text{pull}\} \times [2, n+1] \cup \{\langle n+2, \text{payout}, 0 \rangle\} \\ \cup W \times \{\text{win}\} \times \{n+2\} \cup (S \setminus W) \times \{\text{loose}\} \times \{0\}$$

$h \leq \max \{p(s) / s \in W\}$

$$AP = \bigcup_{i=1}^3 \{w_i = f \mid f \in \text{Fruits}\} \cup \bigcup_{s \in W} \{\text{price} = n / h \in \mathbb{N}\}$$

where $\text{Fruits} = \{\text{apple}, \text{pear}, \text{banana}, \dots\}$

$$L: S \rightarrow 2^{AP}$$

$$s \mapsto \begin{cases} \bigcup_{k=1}^3 \{w_k = c_k(s)\} & \dots s \in \{2, \dots, n+1\} \setminus W \\ \{\text{price} = p(s)\} \cup \bigcup_{k=1}^3 \{w_k = c_k(s)\} & \dots s \in W \\ \emptyset & \dots s \in \{0, n+2\} \end{cases}$$

where $c_1, c_2, c_3: \{2, \dots, n+1\} \rightarrow \text{Fruits}$ functions that assign fruits to states 2, ..., n+1
 $p: W \rightarrow \mathbb{N}$ prize function

the 3 fruits that identify the state

Such that

$$\{\langle c_1(s), c_2(s), c_3(s) \rangle \mid s \in \{2, \dots, n+1\}\} = \text{Fruits} \times \text{Fruits} \times \text{Fruits} \\ = \text{Fruits}^3$$

Nondeterminism

Crucial modeling concepts

10

for modeling interleaving of processes

(wanting to study all possibilities how parallel processes can execute in time)

for abstraction purposes, e.g. underspecification (leaving room for how implementations can resolve different possibilities for intended executions)

for interaction with an unknown environment

$TS = (S, Act, \rightarrow, I, AP, L)$ is

$$s \xrightarrow{\alpha} s_1 \quad s \xrightarrow{\alpha} s_2 \Rightarrow s_1 = s_2$$


action-deterministic : $\Leftrightarrow |I| \leq 1$, and $\forall s \in S, a \in Act : |\text{Post}(s, a)| \leq 1$

AP-deterministic : $\Leftrightarrow |I| \leq 1$, and

$$\forall s \in S \forall A \in 2^{AP} : |\{s' \in \text{Post}(s) / L(s') = A\}| \leq 1$$


$$s \xrightarrow{\alpha} L(s_1) \quad s \xrightarrow{\beta} L(s_2) \Rightarrow s_1 = s_2$$

Example:



action-deterministic ✓

but not AP-deterministic



not action-deterministic

but AP-deterministic ✓




Figure 2.11: Rules for handshaking.

whereas according to the latter type processes “communicate” via shared variables. In this subsection, we consider a mechanism by which concurrent processes interact via handshaking. The term “handshaking” means that concurrent processes that want to interact have to do this in a *synchronous* fashion. That is to say, processes can interact only if they are both participating in this interaction at the same time—they “shake hands”.

Information that is exchanged during handshaking can be of various nature, ranging from the value of a simple integer, to complex data structures such as arrays or records. In the sequel, we do not dwell upon the content of the exchanged messages. Instead, an abstract view is adopted and only communication (also called synchronization) actions are considered that represent the occurrence of a handshake and not the content.

To do so, a set H of *handshake actions* is distinguished with $\tau \notin H$. Only if both participating processes are ready to execute the same handshake action, can message passing take place. All actions outside H (i.e., actions in $Act \setminus H$) are independent and therefore can be executed autonomously in an interleaved fashion.

Definition 2.26. Handshaking (Synchronous Message Passing)

Let $TS_i = (S_i, Act_i, \rightarrow_i, I_i, AP_i, L_i)$, $i=1, 2$ be transition systems and $H \subseteq Act_1 \cap Act_2$ with $\tau \notin H$. The transition system $TS_1 \parallel_H TS_2$ is defined as follows:

$$TS_1 \parallel_H TS_2 = (S_1 \times S_2, Act_1 \cup Act_2, \rightarrow, I_1 \times I_2, AP_1 \cup AP_2, L)$$

where $L(\langle s_1, s_2 \rangle) = L_1(s_1) \cup L_2(s_2)$, and where the transition relation \rightarrow is defined by the rules shown in Figure 2.11. ■

Notation: $TS_1 \parallel TS_2$ abbreviates $TS_1 \parallel_H TS_2$ for $H = Act_1 \cap Act_2$.




Figure 3.1: An example of a deadlock situation.

Example 3.1. Deadlock for Fault Designed Traffic Lights


Consider the parallel composition of two transition systems

$$TrLight_1 \parallel TrLight_2$$

modeling the traffic lights of two intersecting roads. Both traffic lights synchronize by means of the actions α and β that indicate the change of light (see Figure 3.1). The apparently trivial error to let both traffic lights start with a red light results in a deadlock. While the first traffic light is waiting to be synchronized on action α , the second traffic light is blocked, since it is waiting to be synchronized with action β . ■

Example 3.2. Dining Philosophers

This example, originated by Dijkstra, is one of the most prominent examples in the field of concurrent systems.



Five philosophers are sitting at a round table with a bowl of rice in the middle. For the philosophers (being a little unworldly) life consists of thinking and eating (and waiting, as we will see). To take some rice out of the bowl, a philosopher needs two chopsticks. In between two neighboring philosophers, however, there is only a single chopstick. Thus, at any time only one of two neighboring philosophers can eat. Of course, the use of the chopsticks is exclusive and eating with hands is forbidden.

Note that a deadlock scenario occurs when all philosophers possess a single chopstick. The problem is to design a protocol for the philosophers, such that the complete system is deadlock-free, i.e., at least one philosopher can eat and think infinitely often. Additionally, a fair solution may be required with each philosopher being able to think and eat infinitely often. The latter characteristic is called freedom of *individual starvation*.

The following obvious design cannot ensure deadlock freedom. Assume the philosophers and the chopsticks are numbered from 0 to 4. Furthermore, assume all following calculations be “modulo 5”, e.g., chopstick $i-1$ for $i=0$ denotes chopstick 4, and so on.

Philosopher i has stick i on his left and stick $i-1$ on his right side. The action $request_{i,i}$ express that stick i is picked up by philosopher i . Accordingly, $request_{i-1,i}$ denotes the action by means of which philosopher i picks up the $(i-1)$ th stick. The actions $release_{i,i}$ and $release_{i-1,i}$ have a corresponding meaning.

The behavior of philosopher i (called process $Phil_i$) is specified by the transition system depicted in the left part of Figure 3.2. Solid arrows depict the synchronizations with the i -th stick, dashed arrows refer to communications with the $i-1$ th stick. The sticks are modeled as independent processes (called $Stick_i$) with which the philosophers synchronize via actions $request$ and $release$; see the right part of Figure 3.2 that represents the process of stick i . A stick process prevents philosopher i from picking up the i th stick when philosopher $i+1$ is using it.




Figure 3.2: Transition systems for the *i*th philosopher and the *i*th stick.

The complete system is of the form:

$$Phil_4 \parallel Stick_3 \parallel Phil_3 \parallel Stick_2 \parallel Phil_2 \parallel Stick_1 \parallel Phil_1 \parallel Stick_0 \parallel Phil_0 \parallel Stick_4$$

This (initially obvious) design leads to a deadlock situation, e.g., if all philosophers pick up their left stick at the same time. A corresponding execution leads from the initial state

$$\langle think_4, avail_3, think_3, avail_2, think_2, avail_1, think_1, avail_0, think_0, avail_4 \rangle$$

by means of the action sequence *request₄*, *request₃*, *request₂*, *request₁*, *request₀* (or any other permutation of these 5 request actions) to the terminal state

$$\langle wait_{4,0}, occ_{4,4}, wait_{3,4}, occ_{3,3}, wait_{2,3}, occ_{2,2}, wait_{1,2}, occ_{1,1}, wait_{0,1}, occ_{0,0} \rangle.$$

This terminal state represents a deadlock with each philosopher waiting for the needed stick to be released.

A possible solution to this problem is to make the sticks available for only one philosopher at a time. The corresponding chopstick process is depicted in the right part of Figure 3.3. In state *available_{i,j}* only philosopher *j* is allowed to pick up the *i*th stick. The above-mentioned deadlock situation can be avoided by the fact that some sticks (e.g., the first, the third, and the fifth stick) start in state *available_{i,i}*, while the remaining sticks start in state *available_{i,i+1}*. It can be verified that this solution is deadlock- and starvation-free.




Figure 3.3: Improved variant of the i th philosopher and the i th stick.

A further characteristic often required for concurrent systems is robustness against failure of their components. In the case of the dining philosophers, robustness can be formulated in a way that ensures deadlock and starvation freedom even if one of the philosophers is “defective” (i.e., does not leave the think phase anymore).¹ The above-sketched deadlock- and starvation-free solution can be modified to a fault-tolerant solution by changing the transition systems of philosophers and sticks such that philosopher $i+1$ can pick up the i th stick even if philosopher i is thinking (i.e., does not need stick i) independent of whether stick i is in state $available_{i,i}$ or $available_{i,i+1}$. The corresponding is also true when the roles of philosopher i and $i+1$ are reversed. This can be established by adding a single Boolean variable x_i to philosopher i (see Figure 3.4). The variable x_i informs the neighboring philosophers about the current location of philosopher i . In the indicated sketch, x_i is a Boolean variable which is true if and only if the i th philosopher is thinking. Stick i is made available to philosopher i if stick i is in location $available_i$ (as before), or if stick i is in location $available_{i+1}$ while philosopher $i+1$ is thinking.

Note that the above description is at the level of program graphs. The complete system is a channel system with *request* and *release* actions standing for handshaking over a channel of capacity 0. ■

¹Formally, we add a loop to the transition system of a defective philosopher at state $think_i$.




Figure 3.4: Fault-tolerant variant of the dining philosophers.

3.2 Linear-Time Behavior

To analyze a computer system represented by a transition system, either an action-based or a state-based approach can be followed. The state-based approach abstracts from actions; instead, only labels in the state sequences are taken into consideration. In contrast, the action-based view abstracts from states and refers only to the action labels of the transitions. (A combined action- and state-based view is possible, but leads to more involved definitions and concepts. For this reason it is common practice to abstract from either action or state labels.) Most of the existing specification formalisms and associated verification methods can be formulated in a corresponding way for both perspectives.

In this chapter, we mainly focus on the state-based approach. Action labels of transitions are only necessary for modeling communication; thus, they are of no relevance in the following chapters. Instead, we use the atomic propositions of the states to formulate system properties. Therefore, the verification algorithms operate on the *state graph* of a transition system, the digraph originating from a transition system by abstracting from action labels.