

An Introduction to Proofs about Concurrent Programs

K. V. S. Prasad
(for the course TDA384/DIT391)
Department of Computer Science
Chalmers University

Lecture 8, Monday 18 Sep 2017

These are only a rough sketch of notes, released now since it will be too late for this course if we wait till the notes are polished.

1 Examples from Chapter 3 of Ben-Ari's book

This chapter roughly parallels Chap 3 of Ben-Ari, which analyses a series of programs using state diagrams. Here, we study the same programs by analysing the text.

- **State diagram based proofs** - these are easy to understand in principle, though it is also easy to see that the diagrams quickly become too big for manual analysis. Large state diagrams are analysed mechanically by tools called *model checkers* such as SPIN. (Neither SPIN, nor its modelling language, PROMELA, are examinable material in our course).

So one goal for us is to learn, if only in principle, how to use a model checker. Typically, we say what properties we think a program has, and the model checker hunts for counter-examples. Assertions are one way to state program properties. These are often enough for safety properties, but will not do for liveness in general. For that, we need (linear) temporal logic (LTL), covered in later notes.

- **Syntactic proofs** (i.e., arguing from the program text). We do these here in parallel with state diagram proofs, but in stages. The first stage uses informal but hopefully rigorous arguments, with a little propositional calculus notation for compactness. Simple theorems of propositional calculus are assumed. Temporal aspects (arguing about coming or previous states) are first treated entirely informally. Later sections use LTL notation. It is possible to formalise the programming language semantics, but that is not included here. So our arguments will continue to be held together by informal steps.

The goal is first to get you to follow the informal reasoning. Can you make your own arguments? (By the way, no one does formal reasoning before doing the informal thing first).

1.1 Notation

Let the boolean $p2$ mean that process p is at label $p2$, etc. Abusing notation, we sometimes also write $p2$ to mean the label $p2$ itself.

Logical symbols: We use \vee for *inclusive or*, \wedge for *and*, \neg for *not*, \rightarrow for *implies*, and \leftrightarrow for *implies and is implied by*.

1.2 “Hardware processes”

We have so far worked with an abstract world implemented by run-time support (RTS). This world consists of entities that we can call “software processes”, events, messages, atomic actions, and so on. Here, a process can be marked *blocked* while waiting for some event, and be unblocked and marked *ready* by the RTS when the event occurs. So the command `await B` can be interpreted as `block until B`. Only ready processes are *scheduled* (given CPU time); blocked processes are not, since they cannot run.

For these software processes, we also introduced semaphores and other abstract synchronisation and communication structures. How these structures and software processes are implemented by the RTS is not a concern for us here¹. We only need to know what has been implemented. One striking feature of this abstract world is that even for a ready process, we do not know if it is actually running. A related matter is that we know nothing about the speed at which any process runs.

By contrast, the world in this chapter is simpler. The processes here can be called “hardware processes”. Once *spawned*, these simply run until they terminate: they do not block. We assume that each process runs on a separate dedicated CPU. We interpret `await B` to mean `loop skip ; that is, keep doing a skip (do nothing) until B becomes true`. This re-formulation is called *busy-waiting*.

1.3 Definitions: fairness, deadlock, livelock, starvation

Because only one CPU at a time can access a shared variable, we still face issues of scheduling—not a process onto a CPU, but a CPU to a shared variable by a bus arbitrator or similar. We assume *weak fairness*: a scenario is weakly fair if a continually enabled command will be executed at some point.

Since there is no blocked state, and no blocking command, processes are either running or terminated. This means in this set-up we cannot have *deadlock*, which we define as “everyone blocked”. We can have *livelock*, which we define as “everyone busy-waiting”. Note that these definitions differ from those of the textbook (I find those definitions confusing).

We agree with the textbook’s definition of (individual) *starvation*: a process can get stuck forever (busy)-waiting to enter its critical section. A special case is that of

¹Some detail can be found in Ben-Ari’s book. For more, see books on Operating Systems (OS), such as “[Operating Systems: Three Easy Pieces](#)”, by Remzi and Andrea Arpaci-Dusseau, 2015.

non-competitive starvation, or NC-starvation, where p starves if q loops in its NCS.

A working equivalence is that in deadlock and livelock, processes mutually starve each other. In individual starvation, a scenario exists where one particular process starves. The third attempt below shows a program that can livelock even though no process NC-starves, i.e., the only starvation possible is mutual.

In the proofs that follow, a basic idea we explore is that of INVARIANTS.

1.4 First attempt, Alg. 3.5, p. 53

integer <code>turn</code> := 1	
p	q
loop forever	loop forever
p1: await <code>turn=1</code>	q1: await <code>turn=2</code>
p2: <code>turn:=2</code>	q2: <code>turn:=1</code>

The program:

We write t for the variable `turn`, and let t_1 mean $t = 1$ and t_2 mean $t = 2$.

Then we have invariants: $T_1 = t_1 \vee t_2$ and $T_2 = \neg(t_1 \wedge t_2)$. The first is established by noting what values are assigned to t , and the second follows from the nature of variables—they cannot hold two values simultaneously.

Then it follows that $p_2 \rightarrow t_1$ because p has just got past p_1 , and any interference from q can only result in (re)-setting t to 1. Similarly, $q_2 \rightarrow t_2$.

1.4.1 Mutex

We have to show that $M = \neg(p_2 \wedge q_2)$ is invariant. We have $p_2 \rightarrow t_1 \rightarrow \neg t_2 \rightarrow \neg q_2$, and similarly $q_2 \rightarrow \neg p_2$, so M holds.

1.4.2 Livelock

Let $L = p_1 \wedge \neg t_1 \wedge q_1 \wedge \neg t_2$. Then L contradicts T_1 . Thus $\neg L$ is an invariant, and since L defines livelock, we have shown that livelock cannot happen.

1.4.3 Starvation

NC-starvation is possible. If q_1 loops in its NCS (before it executes the `await`, which it may, according to the conditions of the CS problem), the scenario p_1, p_2, q_1 starves p . In this scenario, t_2 holds forever, so p will get stuck in p_1 .

1.5 Second attempt, Alg. 3.7, p. 56

boolean <code>wantp</code> := false, <code>wantq</code> := false	
p	q
loop forever	loop forever
p1: await <code>wantq = false</code>	q1: await <code>wantp = false</code>
p2: <code>wantp := true</code>	q2: <code>wantq := true</code>
p3: <code>wantp := false</code>	q3: <code>wantq := false</code>

The program:

We write w_p for want_p and w_q for want_q .

Note that only p sets w_p and only q sets w_q . Let $T_1 = (p_1 \vee p_2) \leftrightarrow \neg w_p$, and $T_3 = p_3 \leftrightarrow w_p$. Then T_1 and T_2 are invariant. Similar invariants hold for q .

Note that we cannot claim $p_2 \rightarrow \neg w_q$ even though $\neg w_q$ is needed for p to get past p_2 , since we do not know where q is. It may just have executed q_2 .

1.5.1 Mutex

This would require that $(p_2 \vee p_3) \rightarrow \neg(q_2 \vee q_3)$. But to ensure anything about where q is, we have to ensure something about w_q . For example, $w_q \rightarrow \neg q_2$. The premise for the mutex statement tells us nothing about w_q . So we cannot prove mutex, and indeed it is easy to write a scenario where it is broken: p_1, q_1 .

1.5.2 Livelock

Let $L = p_1 \wedge w_q \wedge q_1 \wedge w_p$. Then L defines livelock, and contradicts T_1 , so $\neg L$ is invariant. That is, livelock cannot happen.

1.5.3 Starvation

As in the first attempt above, both the NCS and the pre-protocol are notated q_1 in the abbreviated program. Let $S = p_1 \wedge w_q \wedge q_1$. If q is looping in its NCS, q_1 will always be true. Can then S be always true? If so, it will show NC-starvation of p . But $q_1 \rightarrow \neg w_q$, so S is self-contradictory. That is, $\neg S$ is invariant, and p cannot starve this way.

But p can starve if it is only scheduled to look at w_q after q_2 . Is this weakly fair?

1.6 Third attempt, Alg. 3.8, p. 57

boolean wantp := false, wantq := false	
p	q
loop forever	loop forever
p1: non-critical section	q1: non-critical section
p2: wantp := true	q2: wantq := true
p3: await wantq = false	q3: await wantp = false
p4: critical section	q4: critical section
p5: wantp := false	q5: wantq := false

The program:

We write w_p for wantp and w_q for wantq . Again, only p sets w_p and only q sets w_q .

Let $T_1 = (p_1 \vee p_2) \leftrightarrow \neg w_p$, and $T_2 = (p_3 \vee p_4 \vee p_5) \leftrightarrow w_p$. Then T_1 and T_2 are invariant. Similar invariants hold for q .

Note that we cannot claim $p_4 \rightarrow \neg w_q$ even though $\neg w_q$ is needed for p to get past p_3 , since we do not know where q is. It may just have executed q_2 .

1.6.1 Mutex

We have to show that $M = \neg(p_4 \wedge q_4)$ is invariant. M holds at the start. Can we go from a state where M holds to one where it doesn't? Suppose p is at p_4 , and q is not already at q_4 . To get to q_4 , we need $\neg w_p$ so that q can get past q_3 . But this contradicts T_2 . So M is invariant: mutex is assured.

1.6.2 Livelock

Let $L = p_3 \wedge w_q \wedge q_3 \wedge w_p$; then L defines livelock. But L can be true; nothing in the invariants contradicts it, so livelock can happen. A scenario for this is: $p_1, q_1, p_2, q_2, p_3, q_3$.

1.6.3 Starvation

Let $S = p_3 \wedge w_q \wedge q_1$. If S can be true, p can be NC-starved. But T_1 says $q_1 \rightarrow \neg w_q$, which contradicts S . So $\neg S$ is invariant; NC-starvation cannot occur.

But can $p_3 \wedge w_q$ forever, thus starving p , in some other scenario? Since $w_q \leftrightarrow (q_3 \vee q_4 \vee q_5)$ is invariant, this means $(q_3 \vee q_4 \vee q_5)$. If either q_4 or q_5 can hold forever, individual starvation can result. But q has to pass q_4, q_5 in finite time. So there is no individual starvation, but mutual starvation is possible (livelock, with the case q_3).

1.7 Fourth attempt, Alg. 3.9, p. 59

boolean wantp := false, wantq := false	
p	q
loop forever	loop forever
p1: non-critical section	q1: non-critical section
p2: wantp := true	q2: wantq := true
p3: while wantq	q3: while wantp
p4: wantp := false	q4: wantq := false
p5: wantp := true	q5: wantq := true
p6: critical section	q6: critical section
p7: wantp := false	q7: wantq := false

The program:

Note that this program has dispensed with the `await` statement, writing out the *busy-waits* explicitly.

We write w_p for `wantp` and w_q for `wantq`. Again, only p sets w_p and only q sets w_q .

Let $T_1 = (p_1 \vee p_2 \vee p_5) \leftrightarrow \neg w_p$, and $T_2 = (p_3 \vee p_4 \vee p_6 \vee p_7) \leftrightarrow w_p$. Then T_1 and T_2 are invariant. Similar invariants hold for q .

Note that we cannot claim $p_4 \rightarrow \neg w_q$ even though $\neg w_q$ is needed for p to get past p_3 , since we do not know where q is. It may just have executed q_2 or q_5 .

1.7.1 Mutex

We have to show that $M = \neg(p_6 \wedge q_6)$ is invariant. M holds at the start. Can we go from a state where M holds to one where it doesn't? Suppose p is at p_6 , and q is not already at q_6 . To get to q_6 , we need $\neg w_p$ so that q can get past q_3 . But this contradicts T_2 , which says $p_6 \rightarrow w_p$. So M is invariant: mutex is assured.

1.7.2 Livelock

Let $L = p_3 \wedge w_q \wedge q_3 \wedge w_p$; then a path where states repeatedly satisfy L defines *extended livelock*. But L can be true; nothing in the invariants contradicts it, so livelock can happen. A scenario for this is: $p_1, q_1, p_2, q_2, p_3, q_3$, followed by the execution of the pre-protocol loops p_3, p_4, p_5 and q_3, q_4, q_5 in parallel.

1.7.3 Starvation

Let $S = p_3 \wedge w_q \wedge q_1$. If S can be true, p can be NC-starved. But T_1 says $q_1 \rightarrow \neg w_q$, which contradicts S . So $\neg S$ is invariant; NC-starvation cannot occur.

But can $p_3 \wedge w_q$ forever, thus starving p , in some other scenario? Since $w_q \leftrightarrow (q_3 \vee q_4 \vee q_6 \vee q_7)$ is invariant, this means $(q_3 \vee q_4 \vee q_6 \vee q_7)$. Suppose p is in its pre-protocol loop. Either q is also stuck in its pre-protocol loop, or it escapes. In the latter case, w_q is false in q_1 , so p is stuck forever only if the scheduler never lets p_3 execute when q_1 . Fair?