

# Formal Methods in Software Development

## Exercise 10 (January 29)

Wolfgang Schreiner  
Wolfgang.Schreiner@risc.jku.at

The result is to be submitted by the deadline stated above *via the Moodle interface* of the course as a `.zip` or `.tgz` file which contains

1. a PDF file with
  - a cover page with the course title, your name, Matrikelnummer, and email address,
  - a section for each part of the exercise with the requested deliverables and optionally any explanations or comments you would like to make;
2. the file with the Promela model used in the exercise.
3. the files with the LTL properties (Button “Save As” in the LTL Property Manager).

Email submissions are *not* accepted.

## Exercise 10: Model Checking Leader Election in Spin

We consider a system of  $n$  processes  $p_0, \dots, p_{n-1}$  and  $n$  channels  $c_0, \dots, c_{n-1}$  each of which may hold  $m$  messages. Each process  $p_i$  can send a message to channel  $c_{(i+1) \bmod n}$  and receive a message from channel  $c_{(i+n-1) \bmod n}$ , i.e., the processes are organized in a unidirectional “ring”.

From time to time a process may desire to be elected the “leader” of the ring. The core requirement is that at every moment the ring has at most one leader. We want to ensure this by the following protocol:

- Every process nondeterministically cycles through the states “idle” (no activity), “waiting” (having requested to be elected the leader), or “leader” (having been successfully elected the leader).
- If process  $p_i$  is in state “idle”, it may switch to state “waiting”, i.e., it may request its election to the leader. For this purpose, it sends its identifier  $i$  to its successor in the ring.
- Every process  $p_i$  in state “idle” or “waiting” is ready to receive a message from its predecessor; if such a message (a process identifier)  $j$  is received, it is handled as follows:
  - If  $p_i$  is in state “idle”, it forwards  $j$  to its successor.
  - If  $p_i$  is in state “waiting”, there are three cases:
    1. If  $j < i$ , then  $p_i$  does nothing (i.e., it discards  $j$ ).
    2. If  $j > i$ , then  $p_i$  forwards  $j$  to its successor and switches to state “idle” (i.e., it has lost the election).
    3. If  $j = i$ , then  $p_i$  switches to state “leader” (i.e., it has won the election).
- Every process  $p_i$  in state “leader” does not receive or send any message; however, it may switch to state “idle” and thus make room for another leader.

By this protocol, if there are multiple processes simultaneously competing for leadership, the one with the highest process identifier “wins” the election and becomes the leader.

Your tasks are as follows:

1. Implement above model for process number  $n = 4$  and  $m = 2$  in Promela<sup>1</sup>. For this, introduce by `#define N 4` and `#define M 2` constants for the ring size and the channel size, use a process type `proctype p(byte id)` (which is instantiated as `p(0), ..., p(4)`), a global channel array `chan c[N] = [M] of { byte }`, a type `mtype = {idle, ...}` of states, and a global state array `mtype s[N] = idle`.

**Make sure that your model does not allow repeated transitions where nothing changes (i.e., no message is received or sent); such “stuttering steps” unnecessarily violate the progress properties given below. In particular, the main loop must not contain true -> skip transitions; if no incoming message is available, the loop must be blocked.**

---

<sup>1</sup> If model checking with  $m = 2$  is not feasible, choose  $m = 1$  (you may then also try  $n = 5$ ); generally choose the largest model that you can reasonably check.

2. Run a simulation for several hundred steps. The simulation must not run into a deadlock.
3. Formulate in Spin LTL the property
 

Always, if process  $i$  is the “leader”, no other process is also leader.

and check it for  $i = 0$  and  $i = n - 1$ . Analyze the results in detail and explain whether they indicate an error in your model or not.
4. Formulate in Spin LTL the property
 

Some process becomes the “leader” infinitely often.

and check it. Analyze the results in detail and explain whether they indicate an error in your model or not.
5. Formulate in Spin LTL the property
 

Every process becomes the “leader” infinitely often.

and check it. Analyze the results in detail and explain whether they indicate an error in your model or not.
6. Formulate in Spin LTL the property
 

Some process eventually becomes the “leader” permanently.

and check it. Analyze the results in detail and explain whether they indicate an error in your model or not.
7. Extend your model by an additional state “busy” and allow every process in state “idle” to switch non-deterministically either to state “busy” (in which it does nothing except to switch back to state “idle”) or “waiting”.
 

Check in this model again the property

Some process becomes the “leader” infinitely often.

Analyze the results in detail and explain them.

Assume that in the extended model no process remains forever in an “idle” $\leftrightarrow$ “busy” cycle. Formulate this assumption in LTL and check again above property under this precondition. Analyze the results in detail and explain them.

Please use sufficiently many parentheses to make the parsing of formulas unique (do e.g. not write  $[ ]p | | q$  but write  $( [ ]p ) | | q$  or write  $[ ] ( p | q )$ ).

Please make sure that the model check truly elaborates the whole state space, i.e., that no message error: `max search depth too small` appears in the output. If such a message should appear, increase in “Advanced Parameters” the “Maximum Search Depth” such that the message goes away (if not, a model check result error:  $\emptyset$  is meaningless, because only execution paths up to a certain length have been investigated).

Also make sure that the checker does not run into a memory limit, i.e., that no message `pan: reached -DMEMLIM bound` appears in the output. If such a message should appear, consider

the hints following the message how to either expand the available memory or switch on the hash-compact or bitstate/supertrace options.

The deliverables of the exercise consist of

- The completed Promela model.
- Screenshots of (the final parts of) the simulation runs.
- The LTL properties (PLTL formulas plus definitions of the predicates).
- The output of Spin for each model check.
- Screenshots of counterexample simulations (if any).
- For each model check, an interpretation (did the requested property hold or not and why)?

Some hints/reminders on Promela are given below:

- The Promela version of `if (E) C1 else C2` is

```
if
  :: E -> C1
  :: else -> C2
fi
```

The Promela version of `if (E) C` is

```
if
  :: E -> C
  :: else -> skip
fi
```

The Promela version of `while (E) C` is

```
do
  :: E -> C
  :: else -> break
od
```

In all cases, if you forget the `else` branch, the system will *deadlock* in a state that is not allowed by all the conditions in the other branches.

- The expression `c ? [ M ]` is true if and only if a channel `c` holds a message of type `M`. The statement `c ? M` will then remove the message. A typical application is in

```
do/iff
  :: cond && c ? [ M ] ->
    c ? M;
    ...
  :: ...
od/iff;
```

where in a certain situation only a certain kind of message may be accepted.

- In the attached Promela model, the processes receive identifiers 1,2,3,... i.e.

`p[1]@label`

indicates that  $p(0)$  is at the position indicated by `label` (see also the simulations for the identifiers of the individual processes).