

## Homework Exercise #2

### Due: 11:30 a.m., September 24, 2014

1. We looked at Dijkstra's solution to the  $N$ -process mutual exclusion problem in class. We assume the processes are numbered  $1..N$ . Below, it is given using more modern computer language constructs. ( $\triangleright$  indicates a comment.)

**Shared variables:**

*Interested*[ $1..N$ ]     $\triangleright$  Boolean entries indicate which processes are interested in using the resource  
*Passed*[ $1..N$ ]     $\triangleright$  Boolean entries indicate which processes have passed Phase 1  
*K*     $\triangleright$  integer is the name of a process allowed to try for the resource next

**Local variable:**

*done*     $\triangleright$  Boolean variable that is used to determine when process can exit outer loop

**Code for process  $i$ :**

```

1   $\triangleright$  ENTRY SECTION
2  Interested[ $i$ ]  $\leftarrow$  true
3  loop
4      loop     $\triangleright$  Entry Phase 1
5          exit when  $K = i$ 
6          if not Interested[ $K$ ] then     $\triangleright$  try to go next
7               $K \leftarrow i$ 
8          end loop
9      Passed[ $i$ ]  $\leftarrow$  true     $\triangleright$  begin Entry Phase 2
10     done  $\leftarrow$  true
11     for  $j \leftarrow 1..N$  except  $i$      $\triangleright$  check if anyone else has passed Phase 1
12         if Passed[ $j$ ] then     $\triangleright$  go back to Phase 1
13             Passed[ $i$ ]  $\leftarrow$  false
14             done  $\leftarrow$  false
15         end for
16     exit when done
17 end loop

18  $\triangleright$  CRITICAL SECTION

19  $\triangleright$  EXIT SECTION
20 Passed[ $i$ ]  $\leftarrow$  false
21 Interested[ $i$ ]  $\leftarrow$  false

22  $\triangleright$  REMAINDER SECTION

```

For the critical section, we assume that any process that enters the critical section (and does not experience a stopping failure) eventually completes the critical section and begins executing its exit section.

- (a) The original paper lists many properties that the algorithm is supposed to satisfy, but does not really explain why all of them hold. For example, property (c) says “If any of the computers is stopped well outside its critical section, this is not allowed to lead to potential blocking of the others” and this point is not explicitly addressed in the proof of correctness. Let’s examine what “well outside” means for this algorithm. The system is said to be *blocked* in an infinite execution if there is some suffix of the execution where some process performs infinitely many steps of the entry section and no process ever enters the critical section. Give a precise description of the conditions when a stopping failure of a process could cause the system to block. (For example, a possible answer might be that blocking could occur if and only if a process stops when it is just about to execute line 5 or line 10, and there is a full moon.)

*Briefly* explain why your answer is correct. That is, explain how stopping under the conditions you give could cause the system to block, and explain why stopping in any circumstance that does not satisfy your conditions cannot cause the system to block.

- (b) The liveness property described in Dijkstra’s paper is a bit weak. Here is a stronger liveness property:

For every execution  $\alpha$  where no process stops taking steps and for every process  $p$ , if  $p$  is in its entry section at some time during  $\alpha$ , then there is some later time during  $\alpha$  where  $p$  enters its critical section.

Does Dijkstra’s algorithm satisfy this stronger property? Prove your answer is correct.