

Concurrent Programming TDA384/DIT391

Wednesday, 19 August 2020

Exam supervisor: N. Piterman (piterman@chalmers.se, 073 856 49 10)

(Exam set by N. Piterman, based on the course given January-March 2020)

Material permitted during the exam (hjälpmedel):

As the exam is run remotely we cannot really restrict your usage of material.

Grading: You can score a maximum of 70 points. Exam grades are:

<u>points in exam</u>	<u>Grade Chalmers</u>	<u>Grade GU</u>
28–41	3	G
42–55	4	G
56–70	5	VG

Passing the course requires passing the exam and passing the labs. The overall grade for the course is determined as follows:

<u>points in exam + labs</u>	<u>Grade Chalmers</u>	<u>Grade GU</u>
40–59	3	G
60–79	4	G
80–100	5	VG

The exam **results** will be available in Ladok within 15 *working* days after the exam's date.

Instructions and rules:

- You should be monitored on the dedicated zoom channel while taking the exam!
- Submit the exam solution as a **PDF** file on Canvas. The solution should be typeset using your favourite software. **No** scanned handwritten notes or diagrams are allowed.
- Please write your answers clearly and legibly: unnecessarily complicated solutions will lose points, and answers that cannot be read will not receive any points!
- Justify your answers, and clearly state any assumptions that your solutions may depend on for correctness.

- Answer each question on a new page. Glance through the whole paper first; five questions, numbered Q1 through Q5. Do not spend more time on any question or part than justified by the points it carries.
- Be precise. In your answers, try to use the programming notation and syntax used in the questions. You can also use pseudo-code, *provided* the meaning is precise and clear. If need be, explain your notation.

Q1 (9p). Below is the pseudo-code of a program with two threads, **p** and **q**. The variables **x** and **y** are shared between **p** and **q**. The function $f(\cdot, \cdot)$ accepts two integers as parameters and returns an integer. You do not need to know anything else about f .

int x = 1000; int y = 1000;	
p	q
p_1 : while (x>0) { p_2 : x = x - 1; }	q_1 : while (y>0) { q_2 : y = y - 1; q_3 : x = f(x,y) }

The labels p_1, p_2, q_1, q_2 and q_3 are given only for ease of reference.

(Part a) Construct a scenario for which the program terminates. (4p)

(Part b) Does the program terminate for all scenarios? (5p)

Q2 (11p). The pseudo-code below tries to solve the critical section (CS) problem with two threads, **p** and **q**. Remember that CS must exit after a finite time, but NCS may loop.

The label p_i can mean the command that follows p_i , or the proposition that thread **p** is at p_i , and *the next command p* will execute is p_i .

boolean t= false; boolean tp= false; boolean tq= false;	
p	q
<pre> while(true) { p1 //NCS (non-critical section) p2: tp= true; p3: t= false; p4: while(tq!= t) { }; p5 //CS (critical section) p6: tp= false; } </pre>	<pre> while(true) { q1 //NCS (non-critical section) q2: tq= true; q3: t= true; q4: while(tp== t) { }; q5 //CS (critical section) q5: tq= false; } </pre>

(Part a) Show that $(p_2 \iff \neg tp)$ is an invariant of the program. That is, it always holds. Show that it holds initially and that it is preserved under every transition of process p . (2p)

Use the invariant $(q_2 \iff \neg tq)$ without proof. Notice that these are equivalent to $((p_3 \vee p_4 \vee p_6) \iff tp)$ and $((q_3 \vee q_4 \vee q_6) \iff tq)$.

(Part b) Show that $(p_4 \implies \neg t \vee q_4)$ is an invariant of the program. Show that it holds initially and that it is preserved under every transition of **every process**. (3p)

The invariant $(q_4 \implies t \vee p_4)$ holds as well.

(Part c) Show that $(p_6 \implies \neg t \vee q_4)$ is an invariant of the program. (3p)

Use the invariant $(q_6 \implies t \vee p_4)$ without proof.

(Part d) Show that the program maintains mutual exclusion. (3p).

Q3 (17p). The program from Q2 is repeated below for convenience.

boolean t= false; boolean tp= false; boolean tq= false;	
p	q
<p>while(true) { <i>p</i>₁ //NCS (non-critical section) <i>p</i>₂: tp= true; <i>p</i>₃: t= false; <i>p</i>₄: while(tq!= t) { }; <i>p</i>₅ //CS (critical section) <i>p</i>₆: tp= false; }</p>	<p>while(true) { <i>q</i>₁ //NCS (non-critical section) <i>q</i>₂: tq= true; <i>q</i>₃: t= true; <i>q</i>₄: while(tp== t) { }; <i>q</i>₅ //CS (critical section) <i>q</i>₅: tq= false; }</p>

From the onset, we expect each state to be a quintuple, $(p_i, q_j, \mathbf{tp}, \mathbf{tq}, \mathbf{t})$, where i and j range over $\{2, 3, 4, 6\}$, and \mathbf{tp} , \mathbf{tq} , \mathbf{t} are Boolean. From **Q2** we know that \mathbf{tp} and \mathbf{tq} can be deduced from p_i and q_j . Hence, we use states of the form (p_i, q_j, \mathbf{t}) . As transitions into p_4 and q_4 set \mathbf{t} , we can ignore the value of \mathbf{t} when both p and q are in locations 2 or 3. Only 16 states are reachable.

Notation: We denote the value of \mathbf{t} by x when we do not care about it. For example, (p_2, q_2, x) corresponds to either $(p_2, q_2, \mathit{false})$ or $(p_2, q_2, \mathit{true})$.

Here is a partial state transition table for the program above. As mentioned, only 16 states are reachable from the initial state $(p_2, q_2, \mathit{false})$.

state	new state if p moves	new state if q moves
s1	$(2, 2, x)$	$(3, 2, x) = s3$
s2	$(2, 3, x)$	$(2, 4, \mathit{true}) = s5$
s3	$(3, 2, x)$	$(4, 2, \mathit{false}) = s7$ $(3, 3, x) = s4$
s4	$(3, 3, x)$	
s5	$(2, 4, \mathit{true})$	$(2, 6, \mathit{true}) = s6$
s6	$(2, 6, \mathit{true})$	$(2, 2, x) = s1$
s7	$(4, 2, \mathit{false})$	$(6, 2, \mathit{false}) = s8$
s8	$(6, 2, \mathit{false})$	$(2, 2, x) = s1$
s9	$(4, 3, \mathit{false})$	
s10	$(4, 4, \mathit{false})$	
s11	$(4, 4, \mathit{true})$	
s12	$(4, 6, \mathit{false})$	$(4, 2, \mathit{false}) = s7$
s13	$(6, 3, \mathit{false})$	$(2, 3, x) = s2$
s14	$(6, 4, \mathit{true})$	$(2, 4, \mathit{true}) = s5$
s15	$(3, 4, \mathit{true})$	
s16	$(3, 6, \mathit{true})$	$(3, 2, x) = s3$

(Part a) Fill in the blank entries in the table.

(8p)

(Part b) Explain why the protocol maintains mutual exclusion. (2p)

(Part c) Explain why under fair scheduling the protocol avoids starvation. (7p)

Q4 (14p). In this question we create a barrier in erlang. The setup should include a number (10) of clients each doing some work in rounds. They should all be doing the work on round i at the same time and none of them can proceed to round $i + 1$ before all are done with round i .

The role of the barrier is taken by a server that interacts with the clients. The server sends a message to the clients telling them to start the i -th round and collects messages that notify it of the end of the work of each client on round i . After all clients have completed round i , the server initiates round $i + 1$.

Once all clients have finished 100 rounds, the server should ensure that all clients have terminated before terminating itself.

Your task is to implement the server and the client.

(Part a). Implement the server. Explain the role of the elements of the server's state. (6p)

(Part b). Implement the client. Explain the role of the elements of the client's state (if exists). (4p)

(Part c). Implement the initialization. You can either (a) initialize the server with the IDs of all the clients and have it notify them to start round 1 or (b) initialize the clients directly at round 1. (4p)

Q5 (19p). You are designing two classes that are supposed to access the same resource as readers and writers. Writers repeatedly (forever) compute data and then write to the resource. Readers repeatedly (forever) read from the resource and then process the data they collected. Readers can work on the same data several times if it has not been changed since their last access. Readers can all access the resource simultaneously but writers need exclusive access.

Here is a code skeleton that includes both classes and their initialization.

```
class ReadersWriters {
    final static int NumReaders = 5;
    final static int NumWriters = 2;

    // Synchronization declarations to be defined...

    static class Reader extends Thread {
        public void run() {
            // add synchronization
            while (true) {
                read();
                process();
            }
        }

        private void read() {
            // This function reads from the resource. // No need to implement it.
        }

        private void process() {
            // This function processes the data acquired.
            // It may take a long time. // No need to implement it.
        }
    }

    static class Writer extends Thread {
        public void run() {
            // add synchronization
            while (true) {
                compute();
                write();
            }
        }
    }
}
```

```

        private void compute() {
            // This function prepares data to be written to the resource.
            // It may take a long time. // No need to implement it.
        }

        private void write() {
            // This function writes to the resource. // No need to implement it.
        }
    }

    // Starting the readers and writers
    public static void main(String[] args) {
        for (int i = 0; i<NumReaders; i++) {
            new Reader().start();
        }
        for (int i = 0; i<NumWriters; i++) {
            new Writer().start();
        }
    }
}

```

Your task is to implement the following parts:

(Part a). Write the declarations of the variables you will use for synchronization (you may use either locks or semaphores). Pay attention to types, initialization, and scope. (3p)

(Part b). Complete the implementation of the method `run()` of the class `Reader` according to the description above. (5p)

(Part c). Complete the implementation of the method `run()` of the class `Writer` according to the description above. (5p)

(Part d). You increase the number of readers to 1000 and notice that writers never get a chance to change the resource. Add an additional mechanism that will block new readers from accessing the resource once a writer is ready to replace the data. (6p)