

Concurrent Algorithms 2019

Midterm Exam

December 9th, 2019

Time: 1h45

Instructions:

- This midterm is “closed book”: no notes, electronics, or cheat sheets allowed.
- When solving a problem, do not assume any known result from the lectures, unless we explicitly state that you might use some known result.
- Keep in mind that only one operation on one shared object (e.g., a read or a write of a register) can be executed by a process in a single step. To avoid confusion (and common mistakes) write only a single atomic step in each line of an algorithm.
- Remember to write which variables represent shared objects (e.g., registers).
- Unless otherwise stated, we assume atomic multi-valued MRMW shared registers.
- Unless otherwise stated, we ask for *wait-free* algorithms.
- Unless otherwise stated, we assume a system of n asynchronous processes which might crash.
- For every algorithm you write, provide a short explanation of why the algorithm is correct.

- Make sure that your name and SCIPER number appear on **every** sheet of paper you hand in.
- You are **only** allowed to use additional pages handed to you upon request by the TAs.

Good luck!

Problem	Max Points	Score
1	2	
2	2	
3	3	
4	3	
Total	10	

Problem 1 (2 points)

Your tasks:

1. Write a wait-free algorithm that implements a safe MRSW binary register using (any number of) safe SRSW binary registers.
2. Write a wait-free algorithm that implements a regular MRSW binary register using (any number of) safe MRSW binary registers.

Problem 2 (2 points)

A *snapshot* object maintains an array of registers R of size n , has operations $scan()$ and $update_i()$, where i is the invoking process, and the following sequential specification:

```
1 upon  $update_i(v)$  do
2 |    $R_i \leftarrow v$ 
3 upon  $scan$  do
4 |   return  $R$ 
```

Figure 1: Sequential specification of the snapshot object.

The following algorithm (incorrectly) implements a wait-free atomic *snapshot* object using an array of n shared registers R . Each array element R_i contains a value ($R_i.val$), a timestamp ($R_i.ts$), and a copy of the entire array of values ($R_i.snapshot$).

```
1 upon  $update_i(v)$  do
2 |    $ts \leftarrow ts + 1$ 
3 |    $R_i \leftarrow (v, ts, scan())$ 
4 upon  $scan$  do
5 |    $t_1 \leftarrow collect(), t_2 \leftarrow t_1$ 
6 |   while true do
7 |     |    $t_3 \leftarrow collect()$ 
8 |     |   if  $t_3 = t_2$  then return  $\langle t_3[1].val, \dots, t_3[N].val \rangle$ 
9 |     |
10 |    |   for  $k \leftarrow 1$  to  $N$  do
11 |    |     |   if  $t_3[k].ts \geq t_1[k].ts + 1$  then return  $t_3[k].snapshot$ 
12 |    |     |
13 |    |     |    $t_2 \leftarrow t_3$ 
14 upon  $collect$  do
15 |   for  $j \leftarrow 1$  to  $N$  do
16 |     |    $x_j \leftarrow R_j$ 
17 |   return  $x$ 
```

Figure 2: Incorrect implementation of the snapshot object.

Your task: Give an execution of the algorithm which violates atomicity of the *snapshot* object.

Problem 3 (3 points)

Consider the linearizable and wait-free *log* object. The log object supports two operations: `append` and `getLog`. The sequential specification of the log object is shown below:

```
1 Given:
2 Sequential linked list  $L$  that is initially empty.
3
4 procedure append(obj)
5    $L.append(obj)$ 
6
7 procedure getLog()
8    $result[] \leftarrow \perp$ 
9    $k \leftarrow \text{length}(L)$ 
10   $i \leftarrow 1$ 
11  while  $i \leq k$  do
12     $result[i] \leftarrow \text{element}(L, i)$  // the  $\text{element}(L, i)$  function call returns the  $i$ -th element of list  $L$ 
13     $i \leftarrow i + 1$ 
14  return  $result$ 
```

Figure 3: Sequential specification of the log object.

Furthermore, consider a linearizable and wait-free *fetch-and-increment* object where its sequential specification is shown below:

```
1 Given:
2 Register  $R$  that is initially 0.
3
4 procedure fetchAndIncrement()
5    $old \leftarrow R$ 
6    $R \leftarrow old + 1$ 
7   return  $old$ 
```

Figure 4: Sequential specification of the fetch-and-increment object.

Is it possible to implement the linearizable and wait-free log object by using any number of read-write registers and fetch-and-increment objects? Explain your answer. You can use any known results from the lectures.

Problem 4 (3 points)

An atomic shared counter maintains an integer x , initially 0, and has two operations $\text{inc}()$ and $\text{read}()$. Its sequential specification is as follows:

```
1  $x$  integer, initially 0
2 upon  $\text{read}(x)$  do
3 |   return  $x$ 
4 upon  $\text{inc}(x)$  do
5 |    $x \leftarrow x + 1$ 
```

Consider the following, *incorrect*, implementation of an atomic binary obstruction-free consensus object from shared counters:

uses: C_0, C_1 – atomic shared counters initialized to 0

```
1 upon  $\text{propose}(v)$  do
2 |   while true do
3 |      $(x_0, x_1) \leftarrow \text{readCounters}()$ 
4 |     if  $x_0 > x_1$  then
5 |       |  $v \leftarrow 0$ 
6 |     else if  $x_1 > x_0$  then
7 |       |  $v \leftarrow 1$ 
8 |     if  $|x_0 - x_1| \geq 1$  then
9 |       | return  $v$ 
10 |    |  $C_v.\text{inc}()$ 
11 upon  $\text{readCounters}()$  do
12 |   while true do
13 |     |  $x_0 \leftarrow C_0.\text{read}()$ 
14 |     |  $x_1 \leftarrow C_1.\text{read}()$ 
15 |     |  $x'_0 \leftarrow C_0.\text{read}()$ 
16 |     | if  $x_0 = x'_0$  then
17 |       | return  $(x_0, x_1)$ 
```

Give an execution of the above algorithm that shows that the algorithm is not a correct implementation of an obstruction-free consensus object, i.e. an execution in which some property (obstruction-freedom, validity, or agreement) of obstruction-free consensus is violated.

