



Computer Engineering II

Exercise Sheet 4

Quiz

1 Quiz

- a) In the lecture we use an AtomicBoolean object to implement locking. The AtomicBoolean has a synchronized method (Slide 4/19). So we implement a locking algorithm using locking. Do our locking algorithms work in reality? If so, how can the AtomicBoolean be implemented?
- b) Is it possible to achieve mutual exclusion without any RMW primitives?
- c) What purpose does the extra test in a TTAS lock serve? Why is it still not scalable?

Basic

2 Spin Locks

A read-write lock is a lock that allows either multiple processes to read some resource, or one process to write some resource.

- a) Write a simple read-write lock using only spinning, one shared integer and the CAS operation. Do not use local variables (it is ok to have a variable within a method, but not outside).
- b) What is the problem with this lock?
Hint: What happens if a lot of processes access the lock repeatedly?

We now build a queue lock using only spinning, one shared integer, one local integer per process and the CAS operation.

- c) To prepare for this task, answer the following questions:
 - i) Head and tail of the queue have to be stored in the shared integer. What are the “head” and the “tail”, and how can they both be stored in one integer?
Hint: Could the head be a process id? Or is there a much easier solution?
 - ii) How would a process add itself to the queue?
Hint: You need the local integer of the process for this operation.
 - iii) When has a process acquired the lock?
 - iv) How does a process release the lock?
- d) Write down the lock using pseudo-code. Do not forget to initialize all variables.

Advanced

3 ALock2

Have a look at the source code below. It is a modified version of the ALock for n processes (slides 4/50 ff.).

```
class ALock2 implements Lock {
    boolean[] flags = {true, true, false, ..., false}; // flags[] has length n
    AtomicInteger next = new AtomicInteger(0);
    ThreadLocal<Integer> mySlot;

    public void lock() {
        mySlot = next.getAndIncrement();
        while (!flags[mySlot % n]) {}
        flags[mySlot % n] = false;
    }

    public void unlock() {
        flags[(mySlot + 2) % n] = true;
    }
}
```

- What was the intention of the author of “ALock2”?
- Will ALock2 work properly? Why not?
- Suggest a way to repair ALock2.

Hint: Don't worry about performance.

Mastery

4 MCS Queue Lock

See slides 4/55 ff.

- A developer suggests to add an `abort` flag to each node: if a process no longer wants to wait it sets this `abort` flag to `true`. If a process unlocks the lock, it may see the `abort` flag of the next node, jump over the aborted node, and check the successor's successor node. Modify the basic algorithm to support aborts.

Optional: sketch a proof for your answer.

Hint: Be aware of race-conditions!

- Assuming many processes may abort concurrently, does your answer from a) still work? Explain why. If it does not work: modify your algorithm to allow concurrent aborts.

Optional: sketch a proof for your answer.

- Instead of a `locked` and an `aborted` flag one could use an integer, and modify the integer with the CAS operation. What do you think about this idea? How is the algorithm affected? How is performance affected?

- The CLH lock is basically the same as an MCS lock. Conceptually the only difference is, that a process spins on the `locked` field of the predecessor node, not on its own node. What could be an advantage of CLH over MCS and what could be a disadvantage?