




Recoverable Algorithms for Non-Volatile Main Memory: Exercises and Solutions

Danny Hendler

Ben-Gurion University



Exercise #1

Shared C initially $\langle \text{null}, \text{init} \rangle$, $R[N][N]$ initially *null*

```
1.  proc CAS(old,new)
2.     $\langle j, \text{val} \rangle := C.\text{read}()$ 
3.    if  $\text{val} \neq \text{old}$ 
4.      return false
5.    if  $j \neq \text{null}$ 
6.       $R[j][p] := \text{val}$ 
7.       $\text{ret} := C.\text{cas}(\langle j, \text{val} \rangle, \langle p, \text{new} \rangle)$ 
8.      return ret
9.  proc READ()
10.    $\langle j, \text{val} \rangle := C$ 
11.   return val
```

```
12. proc CAS.RECOVER(old,new)
13.   if  $C = \langle p, \text{new} \rangle v$ 
14.      $\text{new} \in \{R[p][1], \dots, R[p][N]\}$ 
15.     return true
16.   else
17.     proceed from line 2
18. proc READ.RECOVER()
19.    $\langle j, \text{val} \rangle := C$ 
20.   return val
```

Does the order in which the two parts of the conjunction in the condition of line 13 matter for linearizability?

Exercise #1: solution

Yes, the order matters, the left side of the condition must be evaluated before the right side. Assume otherwise. Consider the following scenario, where we use distinct values v_1, v_2 (both $\neq \text{init}$).

1. Process p_1 performs a $\text{CAS}(\text{init}, v_1)$ operation, its cas operation in line 7 succeeds but then p_1 fails.
2. Process p_1 starts to recover and first evaluates the second part of the expression in line 12, which returns false.
3. Process p_2 performs a $\text{CAS}(v_1, v_2)$ operation to completion, so now $C = \langle p_2, v_2 \rangle$ holds.
4. Process p_1 next evaluates the first part of the expression in line 12 which returns false as well, so it attempts to re-execute the CAS and returns false in line 8.

From an initial state of init , we had a $\text{CAS}(\text{init}, v_1)$ that fails and a $\text{CAS}(v_1, v_2)$ which succeeds. Since init , v_1 , v_2 are distinct values, this is a non-linearizable execution.

Exercise #2

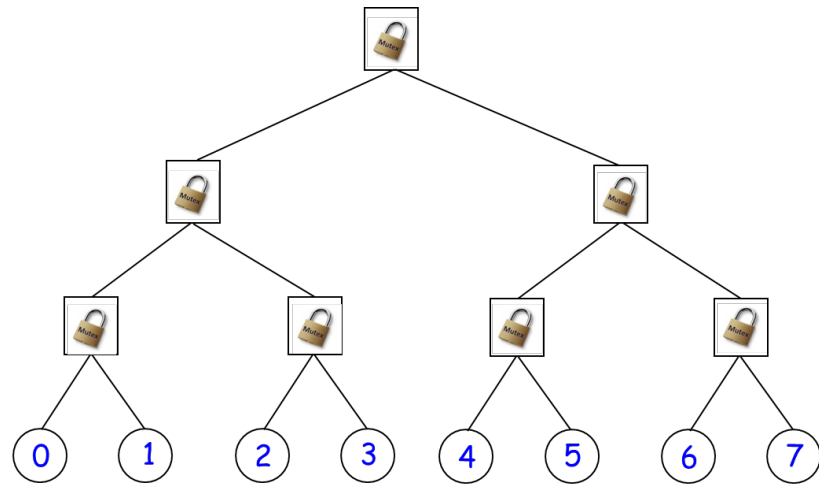
Exit for $p_{i \in \{0, \dots, n-1\}}$

14. $\text{node} := \text{root}$
15. repeat {
16. if i 'th leaf on node's left subtree
17. $\text{side} := \text{left}$
18. else
19. $\text{side} := \text{right}$
20. $\text{node}.\text{exit}(\text{side})$
21. $\text{node} := \text{node's child on path to the leaf}$ }
22. until node is a leaf

Recover for $p_{i \in \{0, \dots, n-1\}}$
(empty)

Shared variables:

A complete binary tree, where each node is an instance of the 2-process recoverable lock



Is Golab and Ramaraju's n -process algorithm still correct if locks are released bottom-up instead of top-down in the exit section?

Exercise #2: solution

No, it is not. Consider the following scenario in a tree with 8 leaves (of processes p0...p7) as shown in the previous slide.

1. Process p0 enters the critical section
2. Process p3 climbs up and waits for the lock of the root's left child in line 11 (slide 44 of my presentation)
3. Process p0 exits, starts by releasing the locks bottom up, so it performs line 14-15 (slide 44) on the root's left child and now p3 is able to climb to the root lock and enters the CS
4. Process p0 releases the root lock by performing lines 14-15 (slide 44) thus finishing its exit code
5. Process p7 now executes its entry code and is able to enter the CS together with p3, a violation of mutual exclusion.

Exercise #3

Extend the recoverable compare-and-swap algorithm to support write operations, in addition to read and CAS.

You may use an *mswap* operation that atomically swaps the values of two shared-memory variables, otherwise this exercise may require a lot of work...

We still assume that all values written by a process (either using write or CAS) are unique. We also assume that evaluation of disjunctive expressions is done in order (see exercise #1).

Exercise #3: solution, code for process p

Shared C initially $\langle \text{null}, \text{init} \rangle$, $R[N][N]$ initially *null*, $W[N]$ initially $\langle \text{null}, \text{null} \rangle$

```
1.  proc CAS(old,new)
2.     $\langle j, \text{val} \rangle := C.\text{read}()$ 
3.    if  $\text{val} \neq \text{old}$ 
4.      return false
5.    if  $j \neq \text{null}$ 
6.       $R[j][p] := \text{val}$ 
7.     $\text{ret} := C.\text{cas}(\langle j, \text{val} \rangle, \langle p, \text{new} \rangle)$ 
8.    return ret

9.  proc READ()
10.    $\langle j, \text{val} \rangle := C$ 
11.   return val

12. proc WRITE(new)
13.    $W[p] := \langle p, \text{new} \rangle$ 
14.    $\text{mswap}(C, W[p])$ 
15.    $\langle j, \text{val} \rangle := W[p]$ 
16.   if  $(j \neq \text{null})$ 
17.      $R[j, p] := \text{val}$ 
18.   return
```

```
19. proc CAS.RECOVER(old,new)
20.   if  $C = \langle p, \text{new} \rangle \vee$ 
       $\langle p, \text{new} \rangle \in \{W[1], \dots, W[N]\} \vee$ 
       $\text{new} \in \{R[p][1], \dots, R[p][N]\}$ 
21.     return true
22.   else
      proceed from line 2

23. proc READ.RECOVER()
24.    $\langle j, \text{val} \rangle := C$ 
25.   return val

26. proc WRITE.RECOVER(new)
27.   if  $C = \langle p, \text{new} \rangle \vee W[p] = \langle p, \text{new} \rangle \vee$ 
       $\text{new} \in \{R[p][1], \dots, R[p][N]\}$ 
28.     return
29.   else
30.     proceed from line 13
```