York University

EECS 6117

Homework Assignment #7 Due: March 27, 2023 at 10:00 p.m.

1. Recall the Michael-Scott lock-free queue discussed in class. Elements are stored in a singly-linked list with one dummy node at the head of the list. It uses two shared variables *Head* and *Tail* that point to nodes in shared memory. The other variables in the code below are local to each process.

```
1: function ENQUEUE(x)
 2:
        node \leftarrow new node containing x
        loop
 3:
             t \leftarrow Tail
 4:
                                                         \triangleright read Tail and store value in local variable t
             n \leftarrow t.next
                                                         \triangleright read t.next to see if Tail needs updating
 5:
             if t = Tail then
                                                         \triangleright reread Tail to see if it is still equal to t
 6:
 7:
                 if n = \text{null then}
                                                         \triangleright if, at line 5, Tail pointed to the last node t in list
 8:
                     exit when CAS(t.next, n, node) \triangleright try to append my node
 9:
                 else
                     CAS(Tail, t, n)
                                                         \triangleright help complete another enqueue before trying again
10:
                 end if
11:
             end if
12:
         end loop
13:
         CAS(Tail, t, n)
                                                        \triangleright finish my enqueue by advancing Tail pointer
14:
    end function
15:
16:
17: function DEQUEUE
        loop
18:
             h \leftarrow Head
19:
                                                         \triangleright read Head and store value in local variable h
             t \leftarrow Tail
20:
                                                         \triangleright read Tail and store value in local variable t
             n \leftarrow h.next
                                                         \triangleright read second element in list
21:
             if h = Head then
                                                         \triangleright reread Head to see if it is still equal to h
22:
                 if h = t then
23:
                                                         \triangleright check whether, at line 20, Head was equal to Tail
                     if n = \text{null then}
                                                         \triangleright check whether, at line 20, list contained only dummy node
24:
                         return null
                                                         \triangleright returning null indicates queue is empty
25:
                     end if
26:
                     CAS(Tail, t, n)
                                                         \triangleright help complete an enqueue before trying again
27:
28:
                 else
                     if CAS(Head, h, n) then
                                                         \triangleright try to advance Head to accomplish dequeue
29:
                         return n.value
30:
                                                         \triangleright return value in node that I advanced Head to
                     end if
31:
                 end if
32:
             end if
33:
34:
         end loop
35: end function
```

We argued in class that we could linearize

- an ENQUEUE operation when it performs a successful CAS on line 8,
- a DEQUEUE operation that returns null at its last execution of line 20, and
- a DEQUEUE operation that performs a successful CAS on line 29 at that CAS.

Linearizing in this way guaranteed the invariant that, at any time, the sequence of elements that should be in the queue (if the operations linearized so far were done sequentially in their linearization order) is exactly the same as the sequence of elements reachable from the node *Head* by following *next* pointers (excluding *Head* itself). We used this invariant to argue that each operation's response is consistent with the linearization ordering.

- (a) If line 14 were removed from the code, would the invariant above still be true? Would the linearization still be valid? Briefly justify your answer.
- (b) Now consider the code above (including line 14). Suppose we modify the linearization so that an ENQUEUE operation is linearized when *Tail* is changed to point to the node created by the ENQUEUE. DEQUEUE operations are linearized as before. Is this linearization still correct?

If your answer is yes, explain why the linearization point of each operation is between the invocation and response of that operation, give an invariant, and use the invariant to explain why each operation's response is consistent with the linearization ordering.

If your answer is no, give an execution for which this modified linearization is incorrect.

(c) Finally, consider the implementation with line 14 removed. Repeat part (b) for this implementation. (Is the modified linearization given in part (b) correct for the implementation with line 14 removed?)