

Homework 10: Data Synchronization

Introduction

The focus of this problem set is the theoretical side of the material taught in class. This problem set focuses on data synchronization and comparing lock-based and lock-free FIFO queue implementations.

1 Data synchronization

Figures 1 and 2 present two implementations of a FIFO queue. Figure 1 is a lock-based implementation, and Figure 2 is a lock-free implementation. In both queue implementations, a pool of nodes is allocated in advance. A call to `new_node()` grabs a free node from the pool of nodes, and `free_node(node)` returns the node `node` to the pool. In the implementation of the lock-free queue, the compare-and-swap instruction `CAS(addr, old_val, new_val)` is an atomic instruction that has the following effect:

```
if (*addr == old_val) {  
    *addr = new_val;  
    return true;  
}  
return false;
```

For the questions below, assume that `CAS` can operate on the entire `pointer_t`, that the compiler cannot change the order of instructions, and that there are always enough free nodes in the pool to perform all enqueue operations. Assume also that the nodes in the queue do not cross cache lines, and thus all writes are atomic.

Read both implementations carefully. Before you start answering the questions, you may find it helpful to draw diagrams of an empty queue and a queue with a few nodes. Using these diagrams, try to understand how nodes are inserted and deleted from the queue in both implementations.

Note that the first node added in initialization of both the lock-based and lock-free version of the deque is a dummy value and never dequeued. It is only used to denote an empty deque.

2 Check-off Questions

1. What are constraints on `enqueue` and `dequeue` in the FIFO queue? You do not need to look at the code yet.

2. What is the advantage of using two locks over one lock?
3. In the style of comments of the lock-based FIFO queue code, add comments to the lock-free code (on paper), explaining what each line does. The comments should be short and precise (not more than 10 words each). We have provided you a copy of the code in Figure 2.
4. Explain how a new node is inserted into the lock-free queue. How many successful `CAS`s are needed per node? What happens if the `CAS` in line 96 fails? How far can the tail lag behind? Is the program correct without line 96?
5. Carefully look at the code for the lock-free dequeue operation and answer the following questions:
 - (a) Line 104 checks what was already assigned in line 101. Why do we need line line 104?
 - (b) In line 111 the value of the node is read before the head is updated in line 112. Why is this important? What can happen if we change the order of the lines?
 - (c) What happens if the `CAS` in line 112 is unsuccessful?
6. Which implementation do you expect to run faster — the lock-based or the lock-free? Explain your answer in terms of cost of the synchronization primitives, contention, synchronization overhead, etc.
7. Show how to simplify the lock-based code if only one thread may enqueue nodes to the queue. Write the pseudocode and comment it. Explain in your own words why your solution is correct (i.e. any execution sequence keeps the FIFO ordering).
8. Show how to simplify the lock-free code if only one thread may dequeue nodes from the queue. Write the pseudocode and comment it. Explain in your own words why your solution is correct (i.e. any execution sequence keeps the FIFO ordering) and why it is non-blocking.
9. Explain how count is used to handle the ABA problem discussed in recitation.

```

11 struct node_t {
12     data_t value;
13     node_t* next;
14 };
15 struct queue_t {
16     node_t* head;
17     node_t* tail;
18     mutex_t h_lock;
19     mutex_t t_lock;
20 };
21
22 void initialize(queue_t* q, data_t value) {
23     node_t* node = new_node(); // Allocate a new node
24     node->value = value;
25     node->next = NULL;         // Make it the only node in the queue
26     q->head = node;           // Both head and tail point to it
27     q->tail = node;
28     q->h_lock = FREE;        // Locks are initially free
29     q->t_lock = FREE;
30 }
31
32 void enqueue(queue_t* q, data_t value) {
33     node_t* node = new_node(); // Allocate a new node
34     node->value = value;        // Copy enqueued value into node
35     node->next = NULL;         // Set next pointer of node to NULL
36     lock(&q->t_lock);          // Acquire t_lock to access tail
37     q->tail->next = node;      // Append node at the end of queue
38     q->tail = node;           // Swing tail to node
39     unlock(&q->t_lock);        // Release t_lock
40 }
41
42 bool dequeue(queue_t* q, data_t* pvalue) {
43     lock(&q->h_lock);          // Acquire h_lock to access head
44     node_t* node = q->head;    // Read head
45     new_head = node->next;     // Read next pointer
46     if (new_head == NULL) {    // Is queue empty?
47         unlock(&q->h_lock);    // Release h_lock before return
48         return false;          // Queue was empty
49     }
50
51     *pvalue = new_head->value; // Queue not empty. Read value
52     q->head = new_head;        // Swing head to next node
53     unlock(&q->h_lock);        // Release h_lock
54     free_node(node);          // Free node
55     return true;              // Dequeue succeeded
56 }

```

Figure 1: C-like pseudocode declaring, initializing, and adding for lock-based FIFO queue.

```
57 struct pointer_t {
58     node_t* ptr;
59     unsigned int count;
60 };
61 struct node_t {
62     data_t value;
63     pointer_t next;
64 };
65 struct queue_t {
66     pointer_t head;
67     pointer_t tail;
68 };
69
70 void initialize(queue_t* q, data_t value) {
71     node_t* node = new_node();
72     node->value = value;
73     node->next.ptr = NULL;
74     q->head.ptr = node;
75     q->tail.ptr = node;
76 }
77
78 void enqueue(queue_t* q, data_t value) {
79     node_t* node = new_node();
80     node->value = value;
81     node->next.ptr = NULL;
82     pointer_t tail;
83     while (true) {
84         tail = q->tail;
85         pointer_t next = tail.ptr->next;
86         if (tail == q->tail) {
87             if (next.ptr == NULL) {
88                 if (CAS(&tail.ptr->next, next, (struct pointer_t) {node, next.count + 1})) {
89                     break;
90                 }
91             } else {
92                 CAS(&q->tail, tail, (struct pointer_t) {next.ptr, tail.count + 1});
93             }
94         }
95     }
96     CAS(&q->tail, tail, (struct pointer_t) { node, tail.count + 1 });
97 }
```

Figure 2: C-like pseudocode for declaring, initializing, and adding to a lock-free FIFO queue.

```
98 bool dequeue(queue_t* q, data_t* pvalue) {
99     pointer_t head;
100     while (true) {
101         head = q->head;
102         pointer_t tail = q->tail;
103         pointer_t next = head.ptr->next;
104         if (head == q->head) {
105             if (head.ptr == tail.ptr) {
106                 if (next.ptr == NULL) {
107                     return false;
108                 }
109                 CAS(&q->tail, tail, (struct pointer_t) { next.ptr, tail.count + 1});
110             } else {
111                 *pvalue = next.ptr->value;
112                 if (CAS(&q->head, head, (struct pointer_t) { next.ptr, head.count + 1})) {
113                     break;
114                 }
115             }
116         }
117     }
118     free_node(head.ptr);
119     return true;
120 }
```

Figure 3: C-like pseudocode for dequeuing in a lock-free FIFO queue.

MIT OpenCourseWare
<https://ocw.mit.edu>

6.172 Performance Engineering of Software Systems
Fall 2018

For information about citing these materials or our Terms of Use, visit: <https://ocw.mit.edu/terms>