

# Assignment 8: Correctness conditions

ETH Zurich

## 1 Stack

### 1.1 Background

Figure 1 shows a history for three threads. Each time line corresponds to one thread. All the threads work on a single stack  $s$ . The query  $s.top(i)$  expects an element  $i$  to be on top of stack  $s$ . Note that  $s.top(i)$  does not remove the top item. The command  $s.push(i)$  pushes an element  $i$  on top of the stack  $s$ .




Figure 1: History

### 1.2 Task

1. Is the history shown in figure 1 linearizable? Justify your answer.
2. Is the history shown in figure 1 sequentially consistent? Justify your answer.

### 1.3 Solution

1. The history is not linearizable. The call to  $s.push(1)$  must happen before  $s.top(1)$ , because  $s.top(1)$  expects an element  $i$  on top of the stack. The call to  $s.push(0)$  must not happen before the call to  $s.top(1)$  took effect. The earliest point when  $s.top(1)$  can take effect is at the start of its time span. This is however already too late for  $s.push(0)$  to take effect.
2. The history is sequentially consistent. An equivalent legal sequential history is given by:
  - (a) Thread 2  $s.push(1)$
  - (b) Thread 2  $s:void$
  - (c) Thread 3  $s.top(1)$
  - (d) Thread 3  $s:1$
  - (e) Thread 1  $s.push(0)$
  - (f) Thread 1  $s:void$
  - (g) Thread 3  $s:top(0)$
  - (h) Thread 3  $s:0$

## 2 Non-linearizable queue

### 2.1 Background

This task has been adapted from [1]. The *AtomicInteger* class is a container for an integer value. One of its methods is `boolean compareAndSet(int expect, int update)`. This method compares the object's current value to *expect*. If the values are equal, then it atomically replaces the object's value with *update* and returns `true`. Otherwise, it leaves the object's value unchanged, and returns `false`. This class also provides `int get()` which returns the object's actual value.

Consider the following FIFO queue implementation. It stores its items in an array *items*, which, for simplicity, we will assume has unbounded size. It has two *AtomicInteger* fields. *head* is the index of the next slot from which to remove an item. *tail* is the index of the next slot in which to place an item.

```
class IQueue<T> {
    AtomicInteger head = new AtomicInteger(0);
    AtomicInteger tail = new AtomicInteger(0);
    T[] items = (T[]) new Object[Integer.MAX_VALUE];

    public void enq(T x) {
        int slot;
        do {
            slot = tail.get();
        } while (!tail.compareAndSet(slot, slot + 1));
        items[slot] = x;
    }

    public T deq() throws EmptyException {
        T value;
        int slot;

        do {
            slot = head.get();
            value = items[slot];
            if (value == null) {
                throw new EmptyException();
            }
        } while (!head.compareAndSet(slot, slot + 1));
        return value;
    }
}
```

### 2.2 Task

Give an example showing that this implementation is not linearizable.

### 2.3 Solution

The problem is that the two consecutive operations `tail.compareAndSet(slot, slot + 1)` and `items[slot] = x` do not happen atomically. The problem is illustrated in the following execution.

1. Thread 1 calls `q.enq(e1)` with a matching element *e1*.

2. Thread 1 executes `tail.compareAndSet(slot, slot + 1)` and exits the loop.
3. Thread 2 calls `q.enq(e2)` with a matching element `e2`.
4. Thread 2 executes `tail.compareAndSet(slot, slot + 1)`, exits the loop and executes `items[slot] = x`.
5. Thread 2 finishes the call to `q.enq(e2)`.
6. Thread 3 calls `q.deq()`.
7. Thread 3 executes `slot = head.get()` and `value = items[slot]`. The slot is the one set by thread 1. Note that thread 1 did not set the value for this slot yet.
8. Thread 3 throws an exception because `value` is `null`.
9. Thread 1 executes `items[slot] = x`.
10. Thread 1 finished the call to `q.enq(e1)`.

The execution is depicted in the following history.

1. Thread 1 `q.enq(e1)`
2. Thread 2 `q.enq(e2)`
3. Thread 2 `q:void`
4. Thread 3 `q.deq()`
5. Thread 3 `q:EmptyException()`
6. Thread 1 `q:void`

The history is not linearizable. In the history above the call by thread 2 precedes the call by thread 3. This precedence relation must be preserved in any equivalent sequential history. Furthermore such a history can not have any other dequeue operations other than the one by thread 3. Therefore any such history would be invalid because a dequeue operation after an enqueue operation must not throw an empty exception in the absence of other dequeue operations.

## 3 Binary search tree

### 3.1 Background

Listing 1 shows the class of a binary search tree. The class defines a feature `insert` to add a value to a tree and a feature `has` to check whether the tree contains a value.

Listing 1: Non-linearizable binary search tree

```
class BINARY_SEARCH_TREE
2
  create
4  make

6 feature -- Initialization
  make (a_value: INTEGER)
8    -- Initialize this node with 'a_value'.
```

```
10   do
11     left := Void
12     right := Void
13     value := a_value
14   end
15
16   feature -- Access
17     left: BINARY_SEARCH_TREE
18     -- The left sub tree.
19     right: BINARY_SEARCH_TREE
20     -- The right sub tree.
21     value: INTEGER
22     -- The value.
23
24   feature -- Basic operations
25     insert (a_new_value: INTEGER)
26     -- Insert 'a_new_value' into the tree.
27   require
28     tree_does_not_have_new_value: not has (a_new_value)
29   do
30     if a_new_value < Current.value then
31       if not left = Void then
32         left.insert (a_new_value)
33       else
34         left := create {BINARY_SEARCH_TREE}.make (a_new_value)
35       end
36     else
37       if not right = Void then
38         right.insert (a_new_value)
39       else
40         right := create {BINARY_SEARCH_TREE}.make (a_new_value)
41       end
42     end
43   end
44
45   has (a_value: INTEGER): BOOLEAN
46   -- Does the tree have 'a_value'?
47   do
48     if a_value = Current.value then
49       Result := True
50     else
51       if a_value < Current.value then
52         if not left = Void then
53           Result := left.has (a_value)
54         else
55           Result := False
56         end
57       else
58         if not right = Void then
59           Result := right.has (a_value)
60         else
61           Result := False
```

```
62     end
64     end
end
```

## 3.2 Task

1. Devise an execution sequence that demonstrates that the binary search tree from Listing 1 is not linearizable; provide a corresponding history and explain why this history is non-linearizable.
2. Using the feature *compare\_and\_swap*, develop a linearizable version of the binary search tree class. Provide only the changed features.

The feature *compare\_and\_swap* ( $\$entity$ ,  $test\_value$ ,  $new\_value$ ) sets the value of an entity to  $new\_value$  if and only if the entity currently has the value  $test\_value$ ; the feature call returns whether or not the test was successful. Here, the  $\$$  operator returns the address of the entity.

## 3.3 Solution

### 3.3.1 Task 1

Consider the following execution based on an entity  $n$  of type *BINARY\_SEARCH\_TREE*:

1. Thread 1: Calls **create**  $n.make$  (10) and finishes.
2. Thread 1: Calls  $n.insert$  (5). Continues until  $left := \mathbf{create} \{BINARY\_SEARCH\_TREE\}.make$  ( $a\_new\_value$ ).
3. Thread 2: Calls  $n.insert$  (4) and finishes.
4. Thread 1: Executes  $left := \mathbf{create} \{BINARY\_SEARCH\_TREE\}.make$  ( $a\_new\_value$ ). This overrides the value set by thread 2.
5. Thread 3: Executes  $n.has$  (4). Returns that the element could not be found.

The following history represents this execution:

1. Thread 1: create n.make (10)
2. Thread 1: n:void
3. Thread 1: n.insert (5)
4. Thread 2: n.insert (4)
5. Thread 2: n:void
6. Thread 1: n:void
7. Thread 3: n.has (4)
8. Thread 3: n:False

There are only two possible equivalent sequential histories that preserve the order relation. One of the histories lets the first thread insert its value first:

1. Thread 1: create n.make (10)
2. Thread 1: n:void
3. Thread 1: n.insert (5)
4. Thread 1: n:void
5. Thread 2: n.insert (4)
6. Thread 2: n:void
7. Thread 3: n.has (4)
8. Thread 3: n:False

The other one lets the second thread insert its value first:

1. Thread 1: create n.make (10)
2. Thread 1: n:void
3. Thread 2: n.insert (4)
4. Thread 2: n:void
5. Thread 1: n.insert (5)
6. Thread 1: n:void
7. Thread 3: n.has (4)
8. Thread 3: n:False

Both histories are illegal because the value 4 should not be reported as missing. Hence the implementation is not linearizable.

### 3.3.2 Task 2

Listing 2: Linearizable binary search tree

```

insert (a_new_value: INTEGER)
2   -- Insert 'a_new_value' into the tree.
   require
4   tree_does_not_have_new_value: not has (a_new_value)
   local
6   l_cached_sub_tree : BINARY_SEARCH_TREE -- The cached sub tree.
   do
8   if a_new_value < Current.value then
        l_cached_sub_tree := left
10    if not l_cached_sub_tree = Void then
            left.insert (a_new_value)
12    else
            if not compare_and_swap ($left, l_cached_sub_tree, create {
                BINARY_SEARCH_TREE}.make (a_new_value)) then
14                left.insert (a_new_value)
            end
16    end
end

```

```
18   else
19     l_cached_sub_tree := right
20     if not l_cached_sub_tree = Void then
21       right.insert (a_new_value)
22     else
23       if not compare_and_swap ($right, l_cached_sub_tree, create {
24         BINARY_SEARCH_TREE}.make (a_new_value)) then
25         right.insert (a_new_value)
26       end
27     end
28   end
29 end
```

The binary search tree is linearizable. When an insert operation runs in parallel to a number of insert operations, then the tree will contain the values of all these insert operations. When a has operation runs in parallel to a number of insert operations then the has operation might return true or false for the values of these insert operations. When the has operation runs after a number of insert operations then the has operation returns true for the values of these insert operations. For all other values the has operation returns false.

We now take a look at a history, by going from left to right; we look at the operations in the order in which they start. If multiple operations start at the same time, we look at these operations in a random order. As we go through the history, we construct a new sequential history, starting from an empty one.

- If the current operation is an insert operation then we add the insert operation to the end of the sequential history.
- If the current operation is a has operation that returns true for a value from an already added insert operation then we add the has operation to the end of the sequential history.
- If the current operation is a has operation that returns true for a value from a not yet added insert operation then the has operation must have run in parallel to a matching insert operation. In this case, we stall the has operation and add it after the insert operation as soon as we add the insert operation. This preserves the order relation because the insert operation ran in parallel to the has operation and hence we can choose the ordering between the two in the sequential history.
- If the current operation is a has operation that returns false and there is no insert operation that inserts this value then we add the has operation to the end of the sequential history.
- If the current operation is a has operation that returns false and there is an insert operation that inserts this value then either this insert operation must have run in parallel to the has operation or it must have run after the has operation.
  - In case the insert operation ran in parallel to the has operation, we have to distinguish two situations. On one hand, it could be that the insert operation has already been added, in which case we add the has operation in front of the insert operation. This preserves the order relation because the insert operation ran in parallel to the has operation and hence we can choose the ordering between the two in the sequential history. On the other hand, it could be that the insert operation has not been added, in which case we add the has operation to the end of the sequential history.
  - In case the insert operation ran after the has operation, we add the has operation to the end of the sequential history.

Because we add operations in the order in which they started with the only exception of operations that ran in parallel, the history produced in this way is an equivalent sequential history that preserves the order relation. The history is legal because no has operation returns true for a value that has not been added or false for a value that has been added. Hence, the binary search tree implementation is linearizable.

## 4 Practical sequential consistency

### 4.1 Background

One of the implicit simplifying assumptions behind many of the example programs presented in the course has been that sequential consistency is being followed. Recall that sequential consistency essentially means that the relative ordering of operations between threads does not have to be maintained, but the per-thread ordering of operations should be kept. However, this assumption is invalidated quite easily by both compilers and hardware without careful attention.

Compilers are free to reorder the instructions given in the program text, given that it does not change the output of the sequential program.

For example:

```
1  a := 1
   b := 2
```

can be rewritten to

```
2  b := 2
   a := 1
```

if the compiler thinks it would be faster, as the output of the sequential program is the same in either case.

### 4.2 Task

Consider this one-shot Peterson locking algorithm:

```
   enter1 := true
2  turn := 2
   if not enter2 or turn = 1 then
4     critical section
     enter1 := false
6  end
```

How does this locking algorithm break if the compiler (or CPU) can reorder reads and writes to independent variables? To see how, it may help to rewrite the algorithm so that intermediate expressions are computed and stored into temporary variables, for example, turning  $a + 1 = b$  into

```
   tmp1 := a + 1
2  tmp2 := tmp1 = b
```

### 4.3 Solution

Because the accesses to *enter1* and *enter2* are to independent locations, these can be reordered. For the first processor (similarly for the second) the program could be rewritten as



```
tmp := not enter2
2  enter1 := true
   turn := 2
4  if tmp or turn = 1 then
   critical section
6  enter1 := false
   end
```

Now, both processes will calculate *tmp* as *True*, since initially both entry variables are *False*. Therefore, both will enter the critical section at the same time, violating mutual exclusion.

## 5 Linearizability

### 5.1 Background

The following listing shows the class for a lock that is based on a faulty implementation of the bakery algorithm. The class defines a feature *lock* to acquire the lock and a feature *unlock* to release the lock. The feature *owns.lock* returns whether a thread currently owns the lock. The feature *make* creates the lock for threads  $1 \dots n$ .

```
class BAKERY_LOCK create make
2
feature {NONE} -- Implementation and initialization
4  numbers: ARRAY [INTEGER] -- The numbers.
   is_lock_owner: ARRAY [BOOLEAN] -- Which thread is the lock owner?
6
   make (n: INTEGER)
8     -- Initialize this lock for threads 1 to 'n'.
   do
10    create numbers.make (1, n); numbers.fill_with (0)
     create is_lock_owner.make (1, n); is_lock_owner.fill_with (False)
12   end

14 feature -- Basic operations
   lock (i: INTEGER)
16     -- Acquire this lock for thread 'i'.
   do
18     numbers [i] := 1 + max (numbers)
     for all j /= i do
20       await numbers [j] = 0 or (numbers [i], i) < (numbers [j], j)
     end
22     is_lock_owner [i] := True
   end

24   unlock (i: INTEGER)
26     -- Release this lock for thread 'i'.
   do
28     number [i] := 0
     is_lock_owner [i] := False
30   end

32   owns_lock (i: INTEGER): BOOLEAN
     -- Is thread 'i' the lock owner?
```

```
34  do Result := is_lock_owner [i] end
    end
```

## 5.2 Task

1. Devise an execution sequence that demonstrates that the bakery lock from the listing is non-linearizable; provide a corresponding history and explain why this history is non-linearizable.
2. Suggest a fix to make the bakery lock linearizable. Explain why the fix solves the issue.

## 5.3 Solution

Consider the following execution based on an entity  $l$  of type *BAKERY\_LOCK*:

1. Thread 1: Calls `create l.make (2)` and finishes.
2. Thread 1: Calls `l.lock (1)` and continues until it evaluated  $1 + \max(\text{numbers})$ . The evaluation results in 0.
3. Thread 2: Calls `l.lock (2)` and continues until it evaluated  $1 + \max(\text{numbers})$ . The evaluation results in 0.
4. Thread 2: Assigns 1 as its number.
5. Thread 2: Passes the wait statement because thread 1 has 0 as its number.
6. Thread 2: Finishes `l.lock (2)`.
7. Thread 1: Assigns 1 as its number.
8. Thread 1: Passes the wait statement because both threads have the same number, and thread 1 has a lower identifier than thread 2.
9. Thread 1: Finishes `l.lock (1)`.
10. Thread 1: Calls `l.owns_lock (1)` and finishes it with a positive result.
11. Thread 2: Calls `l.owns_lock (2)` and finishes it with a positive result.
12. Thread 1: Calls `l.unlock (1)` and finishes.
13. Thread 2: Calls `l.unlock (2)` and finishes.

The following history represents this execution:

1. Thread 1: create l.make (2)
2. Thread 1: l:void
3. Thread 1: l.lock (1)
4. Thread 2: l.lock (2)
5. Thread 2: l:void
6. Thread 1: l:void
7. Thread 1: l.owns\_lock (1)

8. Thread 1: l:True
9. Thread 2: l.owns\_lock (2)
10. Thread 2: l:True
11. Thread 1: l.unlock (1)
12. Thread 1: l:void
13. Thread 2: l.unlock (2)
14. Thread 2: l:void

There are only a few possible equivalent sequential histories that preserve the order relation. In all of them, both lock checks must occur after both lock operations; furthermore, both unlock operations must occur after both lock check operations. None of these histories is valid because in all of them two different threads own the lock at the same time, and this is against the validity constraints of the lock.

The problem can be fixed by noting when a thread is choosing its number. A thread can then only check the other numbers if no other thread is currently choosing a number. This is shown in the following listing:

```

class BAKERY_LOCK create make
2
feature {NONE} -- Implementation and initialization
4  numbers: ARRAY [INTEGER] -- The numbers.
   choosing: ARRAY [BOOLEAN] -- Which thread is choosing his number?
6  is_lock_owner: ARRAY [BOOLEAN] -- Which thread is the lock owner?

8  make (n: INTEGER)
   -- Initialize this lock for threads 1 to 'n'.
10 do
   create numbers.make (1, n); numbers.fill_with (0)
12   create choosing.make (1, n); numbers.fill_with (False)
   create is_lock_owner.make (1, n); is_lock_owner.fill_with (False)
14 end

16 feature -- Basic operations
   lock (i: INTEGER)
18   -- Acquire this lock for thread 'i'.
   do
20     choosing [i] := True
       numbers [i] := 1 + max (numbers)
22     choosing [i] := False
       for all j /= i do
24       await choosing [j] = False
         await numbers [j] = 0 or (numbers [i], i) < (numbers [j], j)
26     end
       is_lock_owner [i] := True
28   end

30   unlock (i: INTEGER)
   -- Release this lock for thread 'i'.
32   do

```

```
    number [i] := 0
34    is_lock_owner [i] := False
    end
36
    owns_lock (i: INTEGER): BOOLEAN
38    -- Is thread 'i' the lock owner?
    do Result := is_lock_owner [i] end
40 end
```

## References

- [1] Maurice Herlihy und Nir Shavit. The Art of Multiprocessor Programming. Morgan Kaufmann, 2008.