

Concurrent Programming TDA384/DIT391

Thursday, 26 October 2022

Exam supervisor: N. Piterman (piterman@chalmers.se, 073 856 49 10)

(Exam set by N. Piterman, based on the course given September-October 2022)

Material permitted during the exam (hjälpmedel): Two textbooks; four sheets of A4 paper with notes (potentially on both sides of each paper); English dictionary.

Visits to the exams rooms: The examiner will not be able to visit the exam rooms. Teaching assistants will visit the exam rooms instead. Invigilators are encouraged to allow students to call the examiner for questions.

Grading: You can score a maximum of 70 points. Exam grades are: between 28–41 (3), between 42–55 (4), 56 or more (5).

Passing the course requires passing the exam and passing the labs. The overall grade for the course is determined as follows: between 40–59 (3), between 60–79 (4), 80 or more (5).

The exam **results** will be available in Ladok within 15 *working* days after the exam's date.

Instructions and rules:

- Please write your answers clearly and legibly: unnecessarily complicated solutions will lose points, and answers that cannot be read will not receive any points!
- Justify your answers, and clearly state any assumptions that your solutions may depend on for correctness.
- Answer each question on a new page. Glance through the whole paper first; five questions, numbered Q1 through Q5. Do not spend more time on any question or part than justified by the points it carries.
- Be precise. In your answers, try to use the programming notation and syntax used in the questions. You can also use pseudo-code, *provided* the meaning is precise and clear. If need be, explain your notation.

Q1 (10p). On the next page there is an excerpt from the implementation of the simple *software transactional memory* machine presented in lecture 11. The full code is available in appendix.

Currently, the implementation does not distinguish between variables that are used for reading and for writing. That is, even if a variable is unchanged, when the result of a transaction is committed it will get a new version number.

(Part a). Change the program above so that a transaction will keep track of which variables actually changed and increase the version number of only those that have changed.

(5p)

(Part b). Give an example of transactions and their execution order that will expose the difference between the new and old implementations.

(5p)

```

1 % Variable: {name, version, value}
2 -record(var, {name, version = 0, value = undefined}).
3
4 % check out variable 'Name' from 'Tm'
5 pull(Tm, Name) ->
6 ...
7
8 % commit all variables in list 'Vars' to 'Tm'
9 push(Tm, Vars) when is_list(Vars) ->
10 ...
11
12 % read value of variable
13 read(#var{value = Value}) ->
14 Value.
15
16 % write 'Value' to 'Var'
17 write(Var = #var{}, Value) ->
18 Var#var{value = Value}.
19
20 stm(State = #state{storage = Storage},
21     {push, Vars}) ->
22     case try_push(Vars, Storage) of
23         {success, NewStorage} ->
24             {reply, State#state{storage =
25                 NewStorage}, success};
26         fail ->
27             {reply, State, fail}
28     end.
29
30 try_push([], Storage) ->
31     {success, Storage};
32 try_push([Var = #var{name = Name, version = Version} | Vars],
33         Storage) ->
34     case dict:find(Name, Storage) of
35         {ok, #var{version = Version}} ->
36             try_push(Vars,
37                 dict:store(Name, Var#var{version = Version + 1}, Storage));
38         _ -> fail
39     end.

```

Q2 (18p). In the bounded buffer problem, multiple producers and multiple consumers have access to a shared resource and collaborate in adding items to it (producers) and removing items from it (consumers).

If only one producer and one consumer are available, it is possible to implement a simpler data structure, where two bounded buffers are allocated, inbuffer and outbuffer. The producer continuously produces new data and add it to inbuffer until inbuffer is full. The consumer continuously consumes data from outbuffer until outbuffer is empty. When inbuffer is full and outbuffer is empty the two buffers are swapped.

Your task is to implement a class supporting this access pattern. A skeleton for the code and the way it would be used by two threads is available on the next page.

Notice that the code handles all exceptions by throwing them further, so you do not have to handle exceptions.

Hint: Think about *all* synchronization mechanisms we have learned.

Notice: Solutions should allow the producer/consumer to produce/consume *all* the buffer without synchronization. Synchronization should occur only at points of exchange of buffers.

(Part a). Write the declarations of the variables you will use for synchronization. Pay attention to types, initialization, and scope. (4p)

(Part b). Complete the implementation of the constructor. (2p)

(Part c). Complete the implementation of the method `read` according to the description above. (6p)

(Part d). Complete the implementation of the method `write` according to the description above. (6p)

Credits: this question is modified from an exam given at the university of Twente.

The rest of this page is left blank on purpose.

```

1 public class ProducerConsumer {
2     private int[] inbuffer, outbuffer;
3     private int readloc, writeloc;
4     final int size = 10;
5
6     // Introduce synchronization
7
8     ProducerConsumer() {
9         inbuffer = new int[size];
10        outbuffer = new int[size];
11        // Finalize initialization
12    }
13
14    int read() throws ... {
15        // Implement
16        return value;
17    }
18
19    void write(int val) throws ... {
20        // Implement
21    }
22
23    public static void main(String[] args) {
24
25        ProducerConsumer pc = new ProducerConsumer();
26
27        final Thread producer = new Thread(() -> {
28            while (true) {
29                try {
30                    int val = 0; // produce
31                    pc.write(val);
32                } catch (... e) {
33                    throw new RuntimeException(e);
34                }
35            }
36
37            final Thread consumer = new Thread(() -> {
38                while (true) {
39                    try {
40                        int val = pc.read(); // consume
41                    } catch (... e) {
42                        throw new RuntimeException(e);
43                    }
44                }
45                producer.start();
46                consumer.start();
47    }

```

Q3 (14p). We include below the implementation of a Queue that uses locking to ensure integrity of the data structure.

```
1 class SequentialNode<T> implements Node<T>
2 {
3     public T item; // value stored in node
4     public Node<T> next; // next node in chain
5
6     SequentialNode<T>(T item,Node<T> next) {
7         this.item = item;
8         this.next = next;
9     } }
10
11 class LockQueue<T> implements Queue<T>
12 { // access to front and back of queue
13     private Node<T> head = new SequentialNode<T>();
14     private Node<T> tail = new SequentialNode<T>();
15
16     // empty queue
17     public LockQueue() {
18         // value of sentinel does not matter
19         SequentialNode<T> sentinel = new SequentialNode<T>();
20         head.next = sentinel; tail.next = sentinel;
21     }
22
23     public void enqueue(T value) {
24         SequentialNode<T> newNode = new SequentialNode<T>(value,null);
25         synchronized (tail) {
26             SequentialNode<T> last = tail.next();
27             last.next = newNode;
28             tail.next = newNode;
29         } }
30
31     public T dequeue() throws EmptyException {
32         synchronized (head) {
33             SequentialNode<T> sentinel = head.next;
34             SequentialNode<T> first = sentinel.next;
35             if (first == null) throw new EmptyException();
36             T value = first.item;
37             head.next = first;
38         }
39         return value;
40     } }
```

This queue uses the same structure of the LockFreeQueue. It stores ref-

erences to the head and the tail of the queue. Furthermore, the linked list representing the queue is never empty; its first element (the sentinel) is always present but is not part of the queue. Thus, a linked list that contains only the sentinel represents the empty queue. However, **it uses normal reference and NOT atomic references.**

(Part a) Is it possible for two threads performing enqueue to be in a data race? What would be the consequences of the data race? (3p)

(Part b) Is it possible for two threads performing dequeue to be in a data race? What would be the consequences of the data race? (3p)

(Part c) Is it possible for a thread performing enqueue and a thread performing dequeue to be in a data race? What would be the consequences of the data race? (6p)

(Part d) Suggest how to fix the program so that the problems highlighted in (a), (b), and (c) above are resolved. (4p)

Q4 (16p). In this question we create an Erlang “gate”. This is a device that is either open or closed. When it is open it allows other processes to pass. When it is closed it blocks processes until it is opened. This is similar to the Erlang barrier shown in class whose code is given in the next page.

Implement an Erlang gate that has the following interface:

```
1 -module(gate).
2 -export([init/1,wait/1,allow/1,block/1]).
3
4 % initialize an open gate
5 init() ->
6     ...
7
8 % if the gate is open continue
9 % if the gate is closed block until the gate is opened
10 wait(Gate) ->
11     ...
12
13 % open the gate
14 open(Gate) ->
15     ...
16
17 % close the gate
18 close(Gate) ->
19     ...
```

That is, you should implement a server that is initialized as an open gate.

An open gate discards all open messages it receives. Once it receives a close message it closes. It allows arrivals at the gate to immediately pass through the gate.

A closed gate discards all close messages it receives. It blocks all arrivals at the gate. Once it receives an open message it releases all those who are waiting for the gate to open and becomes open.

(Part a). Implement the init function. (2p)

(Part b). Implement the wait function. (2p)

(Part c). Implement the open and close functions. (4p)

(Part d). Implement the server loop. (8p)

Credits: this question is modified from an exam given at the university of Twente.

For reference, here is the implementation of the Barrier.

```
1 -module(barrier).
2 -export([init/1,wait/1]).
3
4 % initialize barrier for 'Expected' processes
5 init(Expected) ->
6   spawn(fun () -> barrier(0, Expected, []) end).
7
8 % event loop of barrier for 'Expected' processes
9 %   Arrived: number of processes arrived so far
10 %   PidRefs: list of {Pid, Ref} of processes arrived so far
11 barrier(Arrived, Expected, PidRefs)
12   when Arrived == Expected ->    % all processes arrived
13     % notify all waiting processes
14     [To ! {continue, Ref} || {To, Ref} <- PidRefs],
15     % reset barrier
16     barrier(0, Expected, []);
17 barrier(Arrived, Expected, PidRefs) ->
18   receive % still waiting for some processes
19     {arrived, From, Ref} ->
20
21     % one more arrived: add {From, Ref} to PidRefs list
22     barrier(Arrived+1, Expected, [{From, Ref}|PidRefs])
23   end.
24
25 % block at 'Barrier' until all processes have reached it
26 wait(Barrier) ->
27   Ref = make_ref(),
28   % notify barrier of arrival
29   Barrier ! {arrived, self(), Ref},
30   % wait for signal to continue
31   receive {continue, Ref} -> through end.
```

Q5 (12p). The following program is a failed attempt at producing mutual exclusion between two threads using only shared variables. You are requested to analyze the transition table of this concurrent program.

boolean flagp= false; boolean flagq= false;	
p	q
p_1 while (true) {	q_1 while (true) {
p_2 await (!flagq);	q_2 await (!flagp)
p_3 : flagp= true;	q_3 : flagq= true;
p_4 : // critical section	q_4 : critical section
p_5 : flagp= false;	q_5 : flagq= false;
p_6 : }	q_6 : }

The state of the program is a quadruple $(p_i, q_j, \text{flagp}, \text{flagq})$, where i and j range over $\{2, 3, 5\}$, and flagp and flagq are Booleans.

Here is a partial state transition table for the program above. Only 9 states are reachable from the initial state $(p_2, q_2, \text{false}, \text{false})$.

	state	new state if p moves	new state if q moves
s1	$(2, 2, \text{false}, \text{false})$	FILL THIS ENTRY	FILL THIS ENTRY
s2	$(2, 3, \text{false}, \text{false})$	FILL THIS ENTRY	FILL THIS ENTRY
s3	$(2, 5, \text{false}, \text{true})$	FILL THIS ENTRY	FILL THIS ENTRY
s4	$(3, 2, \text{false}, \text{false})$	FILL THIS ENTRY	FILL THIS ENTRY
s5	$(3, 3, \text{false}, \text{false})$	FILL THIS ENTRY	FILL THIS ENTRY
s6	$(3, 5, \text{false}, \text{true})$	FILL THIS ENTRY	FILL THIS ENTRY
s7	$(5, 2, \text{true}, \text{false})$	FILL THIS ENTRY	FILL THIS ENTRY
s8	$(5, 3, \text{true}, \text{false})$	FILL THIS ENTRY	FILL THIS ENTRY
s9	$(5, 5, \text{true}, \text{true})$	FILL THIS ENTRY	FILL THIS ENTRY

(Part a) Fill in the blank entries in the table. (8p)

(Part b) Explain how to see from your table that the protocol does not fulfil mutual exclusion. (2p)

(Part c) Explain what goes wrong (a very short explanation is enough). (2p)

A Full code listings

A.1 Code for Q1

```
1 -module(stm).
2 -compile(exporet_all).
3
4 % Variable: {name, version, value}
5 -record(var, {name, version = 0, value = undefined}).
6
7 % initialize empty transactional memory and register as 'Tm'
8 start(Tm) ->
9   Storage = dict:new(), % empty key/value map
10  register(Tm,
11    gserver:start(#state{storage = Storage},
12      fun stm/2)).
13
14 % shutdown 'Tm'
15 stop(Tm) ->
16   gserver:stop(Tm).
17
18 % create variable 'Name' in 'Tm'
19 % using initial 'Value'
20 % return:
21 % ok -> success
22 % skip -> variable 'Name' already exists
23 create(Tm, Name, Value) ->
24   gserver:request(Tm, {create, Name, Value}).
25
26 % create variable 'Name' in 'Tm'
27 % with 'undefined' initial value
28 % return:
29 % ok -> success
30 % skip -> variable 'Name' already exists
31 create(Tm, Name) ->
32   create(Tm, Name, undefined).
33
34 % remove variable 'Name' from 'Tm'
35 % return:
36 % ok -> success
37 % skip -> variable 'Name' already exists
38 drop(Tm, Name) ->
39   gserver:request(Tm, {drop, Name}).
40
```

```

41 % check out variable 'Name' from 'Tm'
42 % return:
43 % variable -> variable 'Name' found
44 % not_found -> variable 'Name' not found
45 pull(Tm, Name) ->
46   gserver:request(Tm, {pull, Name}).
47
48 % commit all variables in list 'Vars' to 'Tm'
49 % return:
50 % success -> changes committed
51 % fail -> changes not committed, abort
52 push(Tm, Vars) when is_list(Vars) ->
53   gserver:request(Tm, {push, Vars});
54 % commit variable Var to 'Tm'
55 push(Tm, Var) ->
56   push(Tm, [Var]).
57
58 % read value of variable
59 read(#var{value = Value}) ->
60   Value.
61
62 % write 'Value' to 'Var'
63 write(Var = #var{}, Value) ->
64   Var#var{value = Value}.
65
66 stm(State = #state{storage = Storage},
67     {create, Name, Value}) ->
68   case dict:is_key(Name, Storage) of
69     true ->
70       % variable 'Name' already exists
71       {reply, State, skip};
72     false ->
73       % add variable
74       Var = #var{name = Name, value = Value},
75       {reply, State#state{storage =
76         dict:store(Name, Var, Storage)}, ok}
77   end;
78
79 stm(State = #state{storage = Storage},
80     {drop, Name}) ->
81   case dict:is_key(Name, Storage) of
82     true ->
83       {reply, State#state{storage =
84         dict:erase(Name, Storage)}, ok};

```

```

85     false ->
86         % variable 'Name' does not exist
87         {reply, State, skip}
88     end;
89
90 stm(State = #state{storage = Storage},
91     {pull, Name}) ->
92     case dict:is_key(Name, Storage) of
93     true ->
94         {reply, State, dict:fetch(Name, Storage)};
95     false ->
96         {reply, State, not_found}
97     end;
98
99 stm(State = #state{storage = Storage},
100    {push, Vars}) ->
101    case try_push(Vars, Storage) of
102    {success, NewStorage} ->
103        {reply, State#state{storage =
104            NewStorage}, success};
105    fail ->
106        {reply, State, fail}
107    end.
108
109 try_push([], Storage) ->
110    {success, Storage};
111 try_push([Var = #var{name = Name, version = Version} | Vars],
112    Storage) ->
113    case dict:find(Name, Storage) of
114    {ok, #var{version = Version}} ->
115        try_push(Vars,
116            dict:store(Name, Var#var{version = Version + 1}, Storage));
117    _ -> fail
118    end.

```