

Principles of Concurrent Programming TDA384/DIT391

Wednesday, 21 August 2019, 14:00–18:00

Exam set and supervised by: Sandro Stucki (sandros@chalmers.se, 076 420 8639)

Examiner: K. V. S. Prasad (prasad@chalmers.se)

Material permitted during the exam (hjälpmedel):

Two textbooks; four sheets of A4 paper with notes; English dictionary.

Grading: You can score a maximum of 70 points. Exam grades are:

Points in exam	Grade Chalmers	Grade GU
28–41	3	G
42–55	4	G
56–70	5	VG

Passing the course requires passing the exam and passing the labs. The overall grade for the course is determined as follows:

Points in exam + labs	Grade Chalmers	Grade GU
40–59	3	G
60–79	4	G
80–100	5	VG

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; there are five questions, numbered Q1 through Q6. 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. Filing cabinet

(10p)

The following pseudo code models a filing cabinet with the capacity to hold 100 files. (See Appendix A.1 for a full Java version of the code.)

```
1 class LockedCabinet {
2   int[] files = { 1, 1, ..., 1 }; // one hundred '1's
3   Lock lock = new ReentrantLock();
4
5   void checkout(int index) {
6     lock.lock();
7     while (files[index] <= 0) { lock.unlock(); lock.lock(); } // busy wait
8     files[index] = files[index] - 1;
9     lock.unlock();
10  }
11
12  void store(int index) {
13    lock.lock();
14    files[index] = files[index] + 1;
15    lock.unlock();
16  }
17 }
```

The field `files[i]` indicates the status of the i -th file: `files[i] = 1` means that the file is currently stored in the cabinet, `files[i] = 0` means the file has been checked out. All 100 files are initially stored in the cabinet. The methods `checkout()` and `store()` are used to checkout and return the file at position `index`. If a thread tries to checkout a file that is not in the cabinet, it waits for the file to be returned. To avoid multiple threads accessing the cabinet simultaneously, all read and write instructions to files are guarded by a lock.

(Part a) Why does the `while` loop on line 7 repeatedly acquire and release the lock? What could go wrong if the calls to `lock()` and `unlock()` in the loop were removed? Describe a scenario where this would happen. (2p)

(Part b) Assume that two threads t and u concurrently access the filing cabinet. Thread t calls `store(0)`, while thread u simultaneously calls `checkout(0)`. Can any *data races* occur in this situation? If so, give an example. Otherwise, explain why the code is free from data races. (2p)

Turn page for the rest of this question...

A semaphore-based implementation

Consider the following alternative, semaphore-based implementation of a filing cabinet.

```
1 class SemaphoreCabinet {
2     Semaphore[] files = {
3         new Semaphore(1), ..., new Semaphore(1) // one hundred semaphores
4     };
5     void checkout(int index) { files[index].acquire(); }
6     void store(int index)    { files[index].release(); }
7 }
```

Instead of an array of integers and a lock, SemaphoreCabinet uses an array of *semaphores* to keep track of files. The value of the semaphore `files[i]` indicates the status of file i . A value of 1 means the file is in the cabinet, 0 means it has been checked out. To checkout a file, a thread must acquire the corresponding semaphore; to return a file, it releases the semaphore. (See Appendix A.2 for a full Java version.)

(Part c) Assume we are given one instance `lc = new LockedCabinet()` of the lock-based implementation, and one instance `sc = new SemaphoreCabinet()` of the semaphore-based implementation. What is the maximum number of threads that can call `store()` on the lock-based instance `lc` concurrently without blocking? What about the semaphore-based instance `sc`? Justify your answers. (2p)

Assume that there are a hundred threads, each executing the following code:

```
for (int j = 0; j < 10000; ++j) { c.checkout(id); c.store(id); }
```

where `c` is either an instance of `LockedCabinet` or of `SemaphoreCabinet`, and `id` is the thread identifier ($id = 0, 1, \dots, 99$).

(Part d) On a system with $N = 100$ CPU cores, which choice of implementation for `c` will make the program complete faster, `LockedCabinet` or `SemaphoreCabinet`? How does your answer change if there is only $N = 1$ CPU core? For each case, give a rough estimate of the speedup/slowdown of the semaphore-based implementation compared to the lock-based implementation ($1\times, 10\times, 100\times, \dots$). Justify your answers. (4p)

Code and transition table for Q2 and Q3

The code below models a pair of threads t and u collaborating to solve two tasks. See Q2 for a detailed explanation of the code, and Appendix A.3 for the full Java program.

```

    boolean[] done = { false, false }; Lock[] lock = { ... };

```

thread t	thread u
1 while (true) {	1 while (true) {
2 lock[t].lock();	2 lock[u].lock();
3 if (!done[t]) {	3 if (!done[u]) {
4 done[t] = true;	4 done[u] = true;
5 } else {	5 } else {
6 lock[u].lock();	6 lock[t].lock();
7 if (!done[u]) done[u] = true;	7 if (!done[t]) done[t] = true;
8 else break; // exit loop	8 else break; // exit loop
9 lock[u].unlock();	9 lock[t].unlock();
10 }	10 }
11 lock[t].unlock();	11 lock[u].unlock();
12 } // t terminates	12 } // u terminates

Below is an incomplete state transition table for this program. Each state consists of the current values of the program counters (pc_t, pc_u), and the arrays done (d_t, d_u) and lock (l_t, l_u) for each thread. The value of l_i is the thread holding $\text{lock}[i]$. In the table, pc_t and pc_u only take the values 2, 6, 9, 11 or 12; $pc_t = 2$ means the *next* line executed by t is 2.

	current state $s = (pc_t, pc_u, d_t, d_u, l_t, l_u)$	s'(t) = next state if t moves	s'(u) = next state if u moves
s_1	(2, 2, F, F, -, -)	(11, 2, T, F, t, -) = s_2	(2, 11, F, T, -, u) = s_3
s_2	(11, 2, T, F, t, -)	fill this in = s_5	(11, 11, T, T, t, u) = s_4
s_3	(2, 11, F, T, -, u)	(11, 11, T, T, t, u) = s_4	(2, 2, F, T, -, -) = s_6
s_4	(11, 11, T, T, t, u)	(2, 11, T, T, -, u) = s_7	(11, 2, T, T, t, -) = s_8
s_5	fill this in	fill this in = s_9	(2, 11, T, T, -, u) = s_7
s_6	(2, 2, F, T, -, -)	(11, 2, T, T, t, -) = s_8	(2, 6, F, T, -, u) = s_{10}
s_7	(2, 11, T, T, -, u)	fill this in = s_{11}	(2, 2, T, T, -, -) = s_{13}
s_8	(11, 2, T, T, t, -)	(2, 2, T, T, -, -) = s_{13}	(11, 6, T, T, t, u) = s_{12}
s_9	fill this in	fill this in = s_{14}	fill this in = s_{11}
s_{10}	(2, 6, F, T, -, u)	(11, 6, T, T, t, u) = s_{12}	(2, 9, T, T, u, u) = s_{15}
s_{11}	fill this in	fill this in	fill this in = s_{16}
s_{12}	(11, 6, T, T, t, u)	(2, 6, T, T, -, u) = s_{17}	no move
s_{13}	(2, 2, T, T, -, -)	fill this in = s_{16}	(2, 6, T, T, -, u) = s_{17}
s_{14}	fill this in	(11, 2, T, T, t, -) = s_8	fill this in
s_{15}	(2, 9, T, T, u, u)	no move	(2, 11, T, T, -, u) = s_7
s_{16}	fill this in	fill this in = s_{19}	fill this in = s_{18}
s_{17}	(2, 6, T, T, -, u)	fill this in = s_{18}	(2, 12, T, T, u, u) = s_{20}
s_{18}	fill this in	fill this in	fill this in
s_{19}	fill this in	fill this in	fill this in
s_{20}	(2, 12, T, T, u, u)	fill this in	no move

This page left blank so that you can detach this sheet when doing Q2 and Q3 which refer to the code and transition table on page 5.

Q2. Collaborating threads – state transitions (20p)

The pseudo-code given on page 5 models a pair of threads t and u collaborating to solve two tasks. Each task is represented by a Boolean $done[i]$ indicating whether task i has been completed. Each thread is assigned a task and attempts to complete that task first. Once it has completed its task, it attempts to steal and complete the task assigned to the other thread. If it discovers that both tasks have been completed already, it terminates.

To ensure that only one thread works on a given task, the array $done$ is guarded by a pair of locks. Before a thread can read or modify $done[i]$, it must acquire $lock[i]$.

Page 5 also contains an incomplete state transition table modeling the behavior of the two threads. To keep the table small, only program positions relevant to the concurrent behavior of the program are tracked: lines where a lock is acquired (2 and 6) or released (9 and 11), and the end of the threads (line 12). Remember that $pc_t = 2$ means the *next* line executed by t is 2.

Your task is to complete the state transition table and prove or disprove that the code is free from *deadlocks* and *starvation*.

(Part a) Fill in the missing fields in rows $s_5, s_9, s_{11}, s_{14}, s_{16}, s_{18}, s_{19}$ and s_{20} of the transition table. Make sure your solution is consistent with the names given for the missing fields in rows s_2, s_7, s_{13} and s_{17} . Don't write directly into the table on page 5. Instead, write down the complete rows as $s_i : (X, Y, \dots), (U, V, \dots) \dots$ on a separate sheet of paper with the rest of your solutions. (8p)

(Part b) Does the table contain any final states? If so, list all of them, otherwise explain what the absence of final states means for the program. (1p)

(Part c) Are any of the fields in the state tuples redundant? Is it possible to remove any of the variables $pc_t, pc_u, d_t, d_u, l_t, l_u$ without collapsing two previously distinct states? (2p)

Deadlock

A thread is *blocked* in state s if it cannot move from s even though it has not yet terminated. For example u is blocked in s_{12} . A *deadlock* occurs if the program reaches a state where both threads are blocked.

(Part d) Can the program deadlock? Prove your answer from the state transition table. If your answer is yes, also give a trace (i.e. a sequence of states starting from the initial state s_1) that exhibits the deadlock. (2p)

Turn page for the rest of this question. . .

Starvation

Remember that a thread is *starving* if it is blocked indefinitely while other threads can make progress or terminate. We say a thread has become *immortal* if it is starving while all other threads have terminated.

(Part e) Can either of the threads become immortal? Prove your answer from the state transition table. If your answer is yes, also give a trace (i.e. a sequence of states starting from the initial state s_1) that exhibits starvation. (2p)

Releasing locks on termination

Assume that line 8 in the code for threads t and u is changed from

```
8  else break; // exit loop
```

to

```
8  else { lock[t].unlock(); lock[u].unlock(); break; }
```

As a consequence, the values of states s_{19} and s_{20} must be updated, and 3 new states have to be added to the transition table.

...
s_{19}	(12, 2, T, T, -, -)	no move	<i>fill this in = s_{21}</i>
s_{20}	(2, 12, T, T, -, -)	<i>fill this in = s_{22}</i>	no move
s_{21}	<i>fill this in</i>	<i>fill this in</i>	<i>fill this in = s_{23}</i>
s_{22}	<i>fill this in</i>	<i>fill this in = s_{23}</i>	<i>fill this in</i>
s_{23}	<i>fill this in</i>	<i>fill this in</i>	<i>fill this in</i>

(Part f) Fill in the missing fields in rows s_{21} – s_{23} in the updated state transition table. Make sure your solution is consistent with the updated rows s_{19} and s_{20} . (3p)

(Part g) Do your answers to Parts d and e change after the code update to line 8? If so, prove your new answers from the updated state transition table. (2p)

Q3. Collaborating threads – logical reasoning (8p)

Consider again the collaboration problem described in its original form on page 5 (*not* the updated version from Q2.f). See Q2 for an explanation of the code.

Here in Q3, you must argue from the program, not from the state transition table (though you may seek inspiration from it). You get full credit for correct reasoning, whether you use formal logic, everyday language, or a mixture. Formulas and logical laws make your argument concise and precise, and help you keep track. With everyday language, be careful not to be fuzzy, or to mistake wishful thinking for proof. Appendix B reviews briefly the notation of propositional logic and linear temporal logic.

In your reasoning you may assume the following basic invariants about t and u :

1. the threads are only ever observed at the following locations:

$$pc_t \in \{2, 6, 11, 9, 12\} \quad \text{and} \quad pc_u \in \{2, 6, 11, 9, 12\};$$

2. the done flags are Boolean-valued: either d_t or $\neg d_t$ and similarly for d_u .

3. the possible values of the locks l_t and l_u are $l_t \in \{-, t, u\}$ and $l_u \in \{-, t, u\}$;

4. at any given time, the program counters, and locks have exactly one value:

$$pc_t = i \wedge pc_t = j \text{ if and only if } i = j, \text{ and similarly for } pc_u, l_t, \text{ and } l_u.$$

Consider the following propositions:

- $B_t \equiv (pc_t = 2 \wedge l_t \neq -) \vee (pc_t = 6 \wedge l_u \neq -)$,
- $B_u \equiv (pc_u = 2 \wedge l_u \neq -) \vee (pc_u = 6 \wedge l_t \neq -)$,
- $W \equiv (pc_t = 6 \wedge pc_u = 6)$,
- $D \equiv B_t \wedge B_u$,
- $F_t \equiv B_t \rightarrow \diamond(pc_t = 12)$.

(Part a) Describe the propositions B_t , D and F_t in words. Which concurrency properties do D and F_t correspond to? (3p)

(Part b) Explain briefly why the following invariants hold at any point during the execution of the program:

(I1) whenever $pc_t = 6$, we must have $l_t = t$;

(I2) whenever $pc_u = 6$, we must have $l_u = u$. (1p)

(Part c) Show that, at any point in the execution of the program, if W holds, so does D . (*Hint*: you may use the invariants from Part b.) (2p)

(Part d) Give a counterexample to F_t , i.e. describe a program run where B_t becomes true, but $pc_t = 12$ can *never* become true afterwards. (2p)

Q4. Fork/join parallelism

(10p)

The following recursive Java method computes the sum of the values stored in an array `arr` between index `start` and index `end` (excluding `arr[end]`).

```
1  int sum(int[] arr, int start, int end) {
2      if (end <= start) return 0;
3      else return arr[start] + sum(arr, start + 1, end);
4  }
```

Based on this method, Billy has implemented the following naive fork-join task for summing up numbers in an array. (See Appendix A.4 for a full code listing.)

```
1  class SumTask extends RecursiveTask<Integer> {
2      protected int[] arr;
3      protected int start, end;
4
5      SumTask(int[] arr, int start, int end) {
6          this.arr = arr; this.start = start; this.end = end;
7      }
8
9      @Override public Integer compute() {
10         if (end <= start) return 0;
11         else {
12             SumTask t = new SumTask(arr, start + 1, end);
13             t.fork();
14             int s = t.join();
15             return arr[start] + s;
16         }
17     }
18 }
```

(Part a) Draw the dependency graph for a run of `SumTask(arr, 0, 8)` assuming `arr` has exactly 8 elements. Each node should represent a task instance, with leaf nodes representing the base cases (where `end <= start`). Edges should represent parent-child relationships between tasks. How many nodes does the graph contain? (2p)

(Part b) What is the approximate runtime of `SumTask(arr, 0, 8)` assuming each task takes about one unit of time to perform its computation? What is the maximum number of tasks that can be executed in parallel in this implementation (excluding parent tasks waiting for a child task to finish)? How can these numbers be inferred from the dependency graph? (4p)

Turn page for the rest of this question...

After taking a class on concurrent programming, Billy adjusts the `compute()` method of `SumTask` as follows:

```
1  @Override public Integer compute() {
2      if      (end <= start)    return 0;
3      else if (end == start + 1) return arr[start];
4      else {
5          int mid = (start + end) / 2;
6          SumTask t1 = new SumTask(arr, start, mid);
7          SumTask t2 = new SumTask(arr, mid, end);
8          t1.fork();
9          t2.fork();
10         int s1 = t1.join();
11         int s2 = t2.join();
12         return s1 + s2;
13     }
14 }
```

(Part c) What is the total number of tasks created by the updated implementation when running an instance of `SumTask(arr, 0, 8)`? Assume this instance is running on a machine with 8 or more CPU cores. Does the new implementation improve the total runtime compared to the naive one? If so, why? Justify your answers by drawing the dependency graph for `SumTask(arr, 0, 8)` using the new implementation of `compute()`. (4p)

Q5. Event counter

(12p)

The following Erlang function implements the event handler of a so-called *event counter*.

```
1 handler(Goal, Done, Clients) when Done == Goal -> % Goal reached
2   [C ! {done, N} || {C , N} <- dict:to_list(Clients)],
3   handler(Goal, 0, dict:new()); % reset server state
4 handler(Goal, Done, Clients) -> % Goal not yet reached
5 receive
6   {event, From} ->
7     From ! ok,
8     NewClients = dict:update_counter(From, 1, Clients),
9     handler(Goal, Done + 1, NewClients)
10 end.
```

An event counter is a server that receives and counts event messages from clients. The server keeps track of the total number of received events using the argument `Done`, and of the number of events received from particular clients using the dictionary `Clients`. When `Done` reaches a predefined number `Goal`, the server sends a `done` message to all its clients and resets its state (except for the `Goal` argument, which never changes).

Clients interact with the server by calling the functions `event/0` and `wait/0`.

```
11 event() ->
12 counter ! {event, self()},
13 receive ok -> ok end.
14
15 wait() ->
16 receive {done, N} -> N end.
```

Clients notify the server of events by calling `event()`. When a client calls `wait()`, it blocks until it receives a `done` message from the server.

See Appendix A.5 for a full code listing.

(Part a) List all the message types processed by the server. For each type, describe

- the *name* of the message (e.g. “**ok**”);
- the *sender* and *receiver* of the message (e.g. “from server to client”);
- the *payload* of the message (e.g. “the sender PID” or “no payload”);
- what *causes* a message to be sent (e.g. “a client calls the event function”). (3p)

(Part b) What is the meaning of the number returned by a call to `wait()`? (1p)

Turn page for the rest of this question...

Assume an event counter with Goal=2 has just been initialized and two processes P and Q are spawned from the main process using the following code:

```
Main = self(), % remember PID of main process
P = spawn(fun() -> event(), event(), Main ! {self(), wait()} end),
Q = spawn(fun() -> event(), Main ! {self(), wait()} end),
event(). % event generated by the main process
```

Here are two possible sequences of messages received by the Main process after P and Q terminate:

1. {P, 2}, {Q, 1}.
2. {Q, 1}, {P, 1}.

(Part c) For each sequence, give a corresponding sequence of messages received by the server that explains the messages seen by Main. Your answer should be of the form $(M_1, D_1), (M_2, D_2), \dots$, where each M_i is a message received by the server and D_i is the value of Done before M_i is processed by the server. (2p)

Resetting the server

Sometimes it is necessary to reset the event counter. For this, we introduce a new type of message – reset – that can be sent to the server from any other process. reset messages are handled as follows.

By the server: when the server receives a message of the form {reset, P} from process P, it first responds to P with an **ok** message. It then sends a reset message to all its clients (those listed in the Clients dictionary). Finally, it resets its internal state (Done and Clients).

By clients: clients that receive a reset message while waiting for a done message from the server (in a call to wait()) exit the wait function and return **error** to indicate that the server was reset without reaching its goal.

You are now asked to write a short piece of code. You may use Erlang, pseudo code, or a mixture of both. Syntax details, such as punctuation, are not important, but an experienced Erlang programmer should be able to understand your code. Your code *must use functional style*: use recursion, list comprehensions or higher-order functions instead of mutable state or loops. You are not expected to write more than 20 lines of code in total.

(Part d) Implement the reset mechanism according to the above specification. Describe the necessary changes to the functions wait/0 and handler/3. You don't need to repeat parts of the code that do not change. Describe also the new function reset/0 used to reset the server. Calls to reset() should block until an acknowledgment from the server has been received. (6p)

Q6. Lock-free filing cabinet

(10p)

Recall the semaphore-based filing cabinet implementation SemaphoreCabinet from Q1 (page 4). See Appendix A.2 for the full Java version.

```
1 class SemaphoreCabinet {
2     Semaphore[] files = {
3         new Semaphore(1), ..., new Semaphore(1) // one hundred times
4     };
5     void checkout(int index) { files[index].acquire(); }
6     void store(int index)    { files[index].release(); }
7 }
```

Here in Q6, you will be asked to implement a lock-free filing cabinet and to compare its performance to the original SemaphoreCabinet.

Remember that the AtomicInteger class from the Java standard library provides a compareAndSet method with the following signature.

```
boolean compareAndSet(int expectedValue, int newValue)
```

(Part a) Describe a lock-free implementation of SemaphoreCabinet. Your solution should use a CAS (compare-and-set) synchronization primitive, such as Java's AtomicInteger, but no semaphores or locks. Describe all fields and methods of the class that need to be added, changed or removed. You may use either Java code or pseudo code to do so. (6p)

Let c be a filing cabinet that is accessed by N different threads executing the following code:

```
1 for (int j = 0; j < 100; ++j) {
2     // switch between index = 0 and index = 1 in even/odd iterations
3     int index = j % 2;
4     c.checkout(index);
5     doWork();
6     c.store(index);
7 }
```

Assume that the method doWork() performs a non-trivial computation ($>10^6$ CPU cycles). Consider the following scenarios:

(Scenario 1) There are 32 CPU cores and $N = 32$ threads.

(Scenario 2) There are 2 CPU cores and $N = 100$ threads.

(Part b) Which of the implementations do you expect to perform better in Scenario 1, the semaphore-based one from Q1 or your lock-free version? In Scenario 2? Justify your answers for both scenarios. (4p)

Appendix A. Full code listings

Appendix A.1. Code for Q1

```
1 import java.util.concurrent.*;
2 import java.util.concurrent.locks.*;
3
4 class LockedCabinet {
5     final private int[] files = new int[100];
6     final private Lock lock = new ReentrantLock();
7     public LockedCabinet() {
8         for (int i = 0; i < 100; ++i) { files[i] = 1; }
9     }
10    public void checkout(int index) {
11        lock.lock();
12        try {
13            while (files[index] <= 0) { lock.unlock(); lock.lock(); }
14            files[index] = files[index] - 1;
15        } finally { lock.unlock(); }
16    }
17    public void store(int index) {
18        lock.lock();
19        try { files[index] = files[index] + 1; } finally { lock.unlock(); }
20    }
21    public static void main(String[] args) {
22        LockedCabinet lc = new LockedCabinet();
23        Thread[] ts = new Thread[100];
24        for (int i = 0; i < ts.length; ++i) {
25            final int id = i;
26            ts[i] = new Thread() {
27                @Override public void run() { // thread code for Q1
28                    for (int j = 0; j < 10000; ++j) {
29                        lc.checkout(id); lc.store(id);
30                    }
31                }
32            };
33            ts[i].start();
34        }
35        for (int i = 0; i < ts.length; ++i) {
36            try { ts[i].join(); } catch (InterruptedException ie) {}
37        }
38    }
39 }
```

Appendix A.2. Code for Q1 and Q6

```
1 import java.util.concurrent.*;
2
3 class SemaphoreCabinet {
4     final private Semaphore[] files = new Semaphore[100];
5     public SemaphoreCabinet() {
6         for (int i = 0; i < 100; ++i) { files[i] = new Semaphore(1); }
7     }
8     public void checkout(int index) {
9         try { files[index].acquire(); } catch (InterruptedException ie) {}
10    }
11    public void store(int index) {
12        files[index].release();
13    }
14    public static void main(String[] args) {
15        SemaphoreCabinet sc = new SemaphoreCabinet();
16        Thread[] ts = new Thread[100];
17        for (int i = 0; i < ts.length; ++i) {
18            final int id = i;
19            ts[i] = new Thread() {
20                @Override public void run() { // thread code for Q1
21                    for (int j = 0; j < 10000; ++j) {
22                        sc.checkout(id); sc.store(id);
23                    }
24                }
25            };
26            ts[i].start();
27        }
28        for (int i = 0; i < ts.length; ++i) {
29            try { ts[i].join(); } catch (InterruptedException ie) {}
30        }
31    }
32 }
```

Appendix A.3. Code for Q2 and Q3

```
1 import java.util.concurrent.*;
2 import java.util.concurrent.locks.*;
3
4 class TaskThread extends Thread {
5     private int id;
6     public static boolean[] done = { false, false };
7     public static Lock[] lock = { new ReentrantLock(), new ReentrantLock() };
8
9     TaskThread(int id) { this.id = id; }
10
11     @Override public void run() {
12         int me = id;
13         int other = 1 - id;
14         while (true) {
15             lock[me].lock();
16             if (!done[me]) {
17                 done[me] = true;
18                 System.out.println("Thread " + me + " completed task " + me);
19             } else {
20                 lock[other].lock();
21                 if (!done[other]) {
22                     done[other] = true;
23                     System.out.println("Thread " + me + " completed task " + other);
24                 } else {
25                     break; // exit loop
26                 }
27                 lock[other].unlock();
28             }
29             lock[me].unlock();
30         }
31         System.out.println("Thread " + me + " done.");
32     }
33
34     public static void main(String[] args) {
35         TaskThread t0 = new TaskThread(0);
36         TaskThread t1 = new TaskThread(1);
37         t0.start(); t1.start();
38         try {
39             t0.join(); t1.join();
40         } catch (InterruptedException i) {}
41     }
42 }
```

Appendix A.4. Code for Q4

```
1 class SumTaskNaive extends RecursiveTask<Integer> {
2     protected int[] arr;
3     protected int start, end;
4
5     static int sum(int[] arr, int start, int end) {
6         if (end <= start) return 0;
7         else return arr[start] + sum(arr, start + 1, end);
8     }
9
10    SumTaskNaive(int[] arr, int start, int end) {
11        this.arr = arr; this.start = start; this.end = end;
12    }
13    @Override public Integer compute() {
14        if (end <= start) return 0;
15        else {
16            SumTaskNaive t = new SumTaskNaive(arr, start + 1, end);
17            t.fork(); int s = t.join();
18            return arr[start] + s;
19        }
20    }
21    public static void main(String[] args) {
22        final int N = 200; int[] numbers = new int[N];
23        for (int i = 0; i < N; ++i) { numbers[i] = i; }
24        SumTaskNaive t = new SumTaskNaive(numbers, 0, N);
25        int res = ForkJoinPool.commonPool().invoke(t);
26        System.out.println("Sum: " + res);
27    }
28 }
29
30 class SumTask extends SumTaskNaive {
31    SumTask(int[] arr, int start, int end) { super(arr, start, end); }
32    @Override public Integer compute() {
33        if (end <= start) return 0;
34        else if (end == start + 1) return arr[start];
35        else {
36            int mid = (start + end) / 2;
37            SumTask t1 = new SumTask(arr, start, mid);
38            SumTask t2 = new SumTask(arr, mid, end);
39            t1.fork(); t2.fork();
40            int s1 = t1.join(); int s2 = t2.join();
41            return s1 + s2;
42        }
43    }
44 }
```

Appendix A.5. Code for Q5

```
1 -module(counter).
2 -export([init/1,event/0,wait/0,reset/0]).
3
4 % start the event counter process with a given Goal
5 init(Goal) ->
6   register(counter, spawn(fun () -> handler(Goal, 0, dict:new()) end)).
7
8 % notify the counter about an event
9 event() ->
10  counter ! {event, self()},
11  receive ok -> ok end.
12
13 % wait until the counter reaches its goal
14 wait() ->
15  receive {done, N} -> N end.
16
17 % synchronously reset the counter
18 reset() ->
19  todo.
20
21 % the server's event handler
22 handler(Goal, Done, Clients) when Done == Goal -> % Goal reached
23  [C ! {done, N} || {C , N} <- dict:to_list(Clients)],
24  handler(Goal, 0, dict:new()); % reset server state
25 handler(Goal, Done, Clients) -> % Goal not yet reached
26  receive
27    {event, From} ->
28    From ! ok,
29    NewClients = dict:update_counter(From, 1, Clients),
30    handler(Goal, Done + 1, NewClients);
31    {reset, From} ->
32    todo
33  end.
```

Appendix B. Linear Temporal Logic (LTL) notation

1. An atomic proposition such as $q2$ (process q is at label $q2$) *holds for* a state s if and only if process q is at $q2$ in s .
2. Let ϕ and ψ be formulas of LTL. Formulas are either atomic propositions, or are built up from other formulas using the following operators: \neg for “not”, \vee for “or”, \wedge for “and”, \rightarrow for “implies”, \square for “always”, and \diamond for “eventually”. A convenient abbreviation is ϕ iff ψ (i.e., ϕ if and only if ψ) for $(\phi \rightarrow \psi) \wedge (\psi \rightarrow \phi)$.

These operators have the obvious meanings, but two differ from what might be your interpretation of the names. First, $\phi \vee \psi$ (“ ϕ or ψ ”) is false iff both ϕ and ψ are false. This is an “inclusive or”, so $\phi \vee \psi$ is also true if both ϕ and ψ are true. Second, $\phi \rightarrow \psi$ (“ ϕ implies ψ ”) is false iff ϕ is true and ψ is false. So, in particular, $\phi \rightarrow \psi$ is true if ϕ is false. The meanings of the operators \square and \diamond are defined below.

3. A *path* is a possible future of the system, a possibly infinite sequence of states, each reachable from the previous state in the path. A state s *satisfies* formula ϕ if every path from s satisfies ϕ .

A path π satisfies $\square\phi$ if ϕ holds for the first state of π , and for all subsequent states in π . The path π satisfies $\diamond\phi$ if ϕ holds for some state in π .

Note that \square and \diamond are duals:

$$\square\phi \equiv \neg\diamond\neg\phi \quad \text{and} \quad \diamond\phi \equiv \neg\square\neg\phi.$$