Chalmers | GÖTEBORGS UNIVERSITET
Sandro Stucki, Computer Science and Engineering

Principles of Concurrent Programming TDA384/DIT391

Tuesday, 19 March 2019, 8:30–12:30

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 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. Concurrent counter array   *(14p)*

The following Java class implements a thread-safe array of counters:

```java
class LockedCounterArray {
  private int[] counters;
  private Lock lock = new ReentrantLock();

  LockedCounterArray(int size) { counters = new int[size]; /* all 0s */ }

  void increment(int index) {
    lock.lock();
    counters[index] = counters[index] + 1;
    lock.unlock();
  }

  int value(int index) {
    lock.lock();
    int cnt = counters[index];
    lock.unlock();
    return cnt;
  }
}
```

The number of counters in the array is given by the constructor argument `size`; the method `increment()` safely increments the counter at a given `index`; the method `value()` returns the value of counter `index`. See Appendix A.1 for the full Java code listing.

(Part a) Assume we are given a counter array `lca = new LockedCounterArray(100)`. What is the maximum number of threads that can call the method `increment()` on `lca` concurrently without blocking? What is the maximum number of counters that could safely be incremented concurrently in principle? Justify your answers.   *(2p)*

(Part b) How could the implementation of `LockedCounterArray` be adjusted to allow a maximal number of threads concurrent access to the counter array? Propose a lock-based solution. Describe which fields and methods of the class need to be changed and how. You may describe the changes either in words or by giving concrete Java code for new/updated fields and methods.   *(4p)*

Turn page for the rest of this question...

Eliminating locking overhead

Locking can introduce significant overhead in some situations. For the remainder of this question, assume that the size of `counters`, the number of threads trying to access the counters, and the number of processor cores all have the same value: 16.

(Part c) Explain what type of overhead your lock-based solution from Part b may suffer from under this assumption. How costly is locking compared to just incrementing the counter? Does it take orders of magnitude fewer CPU cycles, orders of magnitude more, or about the same? *(2p)*

(Part d) How could the implementation of `LockedCounterArray` be adjusted to remove locking overhead while still allowing a maximal concurrent access to the counter array? Describe which fields and methods of the class need to be changed and how. You may assume that strong synchronization primitives/instructions of the type you have seen in the course are supported by the underlying machine architecture. *(4p)*

(Part e) How does the performance of your solutions from Part b and Part d compare? Which of the two solutions would you expect to perform better in the given scenario? Briefly motivate your answer. *(2p)*

Code and transition table for Q2 and Q3

The code below implements barrier synchronization for two threads t and u using a pair of semaphores. (See Appendix A.2 for the full Java program.)

```
Semaphore[] barrier = { new Semaphore(0), new Semaphore(0) };
```

	thread t		thread u	
1	`boolean done = false;`		`boolean done = false;`	10
2	`do {`		`do {`	11
3	`    // synchronize`		`    // synchronize`	12
4	`    barrier[t].up();`		`    barrier[u].up();`	13
5	`    barrier[u].down();`		`    barrier[t].down();`	14
6	`    // do work`		`    // do work`	15
7	`    done = doWork();`		`    done = doWork();`	16
8	`} while (!done);`		`} while (!done);`	17
9	`// t terminates`		`// u terminates`	18

Below is an incomplete state transition table for this program. Each state consists of the current values of the program counters for both threads (pc_t, pc_u), the local variable `done` in each thread (d_t, d_u), and the barrier semaphores (b_i is the value of `barrier[i]`). The states in the table cover only a subset of program positions: statements that access the shared variable `barrier` (lines 4, 5 for t, and 13, 14 for u), and the ends of the threads (lines 9 and 18). Remember that $pc_t = 4$ means that the *next* line executed by t is 4.

	current state $s = (pc_t, pc_u, d_t, d_u, b_t, b_u)$	$s'(t)$ = next state if t moves	$s'(u)$ = next state if u moves
s_1	(4, 13, F, F, 0, 0)	*fill this in* $= s_2$	(4, 14, F, F, 0, 1) $= s_3$
s_2	*fill this in*	*fill this in*	*fill this in* $= s_4$
s_3	(4, 14, F, F, 0, 1)	*fill this in* $= s_4$	no move
s_4	*fill this in*	(4, 14, F, F, 1, 0) $= s_5$ (9, 14, T, F, 1, 0) $= s_{X1}$	*fill this in* $= s_6$ (5, 18, F, T, 0, 1) $= s_{Y1}$
s_5	(4, 14, F, F, 1, 0)	*fill this in* $= s_7$	(4, 13, F, F, 0, 0) $= s_1$ (4, 18, F, T, 0, 0) $= s_{Y2}$
s_6	*fill this in*	(4, 13, F, F, 0, 0) $= s_1$ (9, 13, T, F, 0, 0) $= s_{X2}$	*fill this in* $= s_8$
s_7	*fill this in*	*fill this in*	*fill this in* $= s_2$ (5, 18, F, T, 1, 0) $= s_{Y3}$
s_8	*fill this in*	(4, 14, F, F, 0, 1) $= s_3$ (9, 14, T, F, 0, 1) $= s_{X3}$	*fill this in*
s_{X1}	(9, 14, T, F, 1, 0)	no move	(9, 13, T, F, 0, 0) $= s_{X2}$ (9, 18, T, T, 0, 0) $= s_Z$
s_{X2}	(9, 13, T, F, 0, 0)	no move	(9, 14, T, F, 0, 1) $= s_{X3}$
s_{X3}	(9, 14, T, F, 0, 1)	no move	*fill this in*
s_{Y1}	(5, 18, F, T, 0, 1)	(4, 18, F, T, 0, 0) $= s_{Y2}$ (9, 18, T, T, 0, 0) $= s_Z$	no move
s_{Y2}	(4, 18, F, T, 0, 0)	(5, 18, F, T, 1, 0) $= s_{Y3}$	no move
s_{Y3}	(5, 18, F, T, 1, 0)	no move	no move
s_Z	(9, 18, T, T, 0, 0)	no move	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. Barriers – state transitions   *(20p)*

The pseudo-code given on page 5 implements barrier synchronization for two threads t and u using semaphores (see Appendix A.2 for the complete Java program).

Before doing any work, the threads synchronize at the barrier. After synchronization, each thread performs some computation by calling the doWork() method. The method returns **true** to indicate that the thread is done and should terminate, or **false** to indicate that more work needs to be done. Every call to doWork() terminates but not every invocation takes the same amount of time to do so.

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 (see page 5 for details). Remember that $pc_t = 4$ means the *next* line executed by t is 4.

Note that a thread sometimes has *more than one move*: depending on the value of the variable done, a thread may execute another iteration of the **while** loop or terminate. For example, thread u has two possible moves from state s_5: to s_1 if doWork() returns **false**, or to s_{Y2} if it returns **true**.

Your job 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 the transition table (rows s_1–s_8 and s_{X3}). Make sure your solution is consistent with the names given for the missing fields in rows s_1, s_3 and s_5. Don't write directly into the table on page 5. Instead, note the names and fields of the incomplete rows as $s_i : (X, Y, \dots), (U, V, \dots) \dots$ on a separate sheet of paper with the rest of your solutions.   *(6p)*

(Part b) List all the final states in the table.   *(1p)*

(Part c) Are any of the fields in the state tuples redundant? Is it possible to remove any of the variables $pc_t, pc_u, d_t, d_u, b_t, b_u$ 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_3.

(Part d) Can the program deadlock? In other words, is there a state where both threads are blocked? 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.   *(1p)*

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. Formally, a t-starvation path is a finite, non-empty sequence of states

$$s^1 \longrightarrow s^2 \longrightarrow s^3 \longrightarrow \cdots \longrightarrow s^n$$

where t is blocked in every state s^i, while u moves from s^i to s^{i+1} and either terminates in s^n or *loops forever* on the path, i.e. it moves from s^n back to s^1 and starts over. A u-starvation path is an analogous sequence of states where only t makes progress. (Note that a path can be of length $n = 1$.)

(Part e) Can either of the threads starve? That is, is there a t- or u-starvation path? 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 starvation (i.e. that ends in a starvation path). If your answer is no, briefly explain why there cannot be any t- or u-starvation paths. *(3p)*

Scenario: equal work for t and u

Now, assume that the doWork() methods of t and u are set up in such a way that the **while** loops of both threads execute *the same number of iterations*. For example, the two doWork() methods could execute tasks from a pair of work list l_t and l_u of the same length. Call this the *equal work* scenario.

(Part f) List all the states from the transition table that become *unreachable* in the *equal work* scenario. (A state is unreachable if no execution of the program will visit it.) Why are they unreachable? *(4p)*

(Part g) Does your answer to Part e change in the *equal work* scenario? If yes, prove your new answer from the updated state transition table (with unreachable states removed). If your new answer is yes, give a trace that ends in a starvation path. If your new answer is no, briefly explain why there cannot be any t- or u-starvation paths in this scenario. *(3p)*

Q3. Barriers – temporal logic   *(8p)*

Consider again the barrier synchronization problem described on page 5, and 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. a thread is in exactly one program location at once:

   $pc_t = i \wedge pc_t = j$ if and only if $i = j$, and similarly for pc_u;

2. the threads are only ever observed at three locations:

   if $pc_t = i$ then $i \in \{4, 5, 9\}$, and if $pc_u = i$ then $i \in \{13, 14, 18\}$.

3. the values of the semaphores b_t and b_u are non-negative: $b_t \geq 0$ and $b_u \geq 0$;

4. a semaphore has exactly one value:

   $b_t = i \wedge b_t = j$ if and only if $i = j$, and similarly for b_u.

Consider the following propositions:

- $B_t \equiv (pc_t = 5 \wedge b_u = 0)$ and $B_u \equiv (pc_u = 14 \wedge b_t = 0)$,

- $R_t \equiv (pc_t = 5 \rightarrow b_t > 0)$,

- $D \equiv \neg(B_t \wedge B_u)$,

- $S_t \equiv B_t \rightarrow \Diamond(pc_t = 4 \vee pc_t = 9)$.

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

(Part b) Show that D is true if R_t is.   *(2p)*

(Part c) Give a counterexample to S_t, i.e. describe a program run where B_t is true, but $pc_t = 4 \vee pc_t = 9$ *never* becomes true.   *(3p)*

Q4. Fork/join parallelism *(10p)*

Consider the following recursive Java method `max()`, which finds the maximum value stored in the array `arr` from index `start` to index $end - 1$ (excluding `arr[end]`).

```java
int max(int[] arr, int start, int end) {
  if (end - start < 1)       return Integer.MIN_VALUE;
  else if (end - start == 1) return arr[start];
  else {
    int mid = (start + end) / 2;
    int m1 = max(arr, start, mid);
    int m2 = max(arr, mid, end);
    return max2(m1, m2);
  }
}
```

The method `max()` uses of the helper function `max2(int x, int y)`, which computes the maximum value of two integers `x` and `y` (see Appendix A.3 for its implementation).

(Part a) Draw the dependency graph for a call to `max(arr, 0, 8)` assuming `arr` has exactly 8 elements. Each node should represent a recursive call to `max()` with the leaf nodes representing the base cases (where `end - start <= 1`). Edges should represent data dependencies between calls. How many nodes does the graph contain? *(2p)*

(Part b) What is the maximum number of calls to `max()` that can be executed in parallel in this scenario (excluding parent calls waiting for a recursive call to finish)? How can this be inferred from the dependency graph? *(2p)*

(Part c) What is the approximate runtime of a fully parallel version of `max(arr, 0, 8)`, assuming each call to `max()` takes about one unit of time? How can this be inferred from the dependency graph? *(2p)*

Turn page for the rest of this question. . .

11

The class `MaxTask` below implements a naive fork/join parallelization of `max()`. The constructor of `MaxTask` has the same arguments as the method `max()`; the `run()` method computes the maximum value and stores the result in the public `result` field. See Appendix A.3 for a full code listing.

```java
class MaxTask extends Thread {
  private int[] arr;
  private int start, end;
  public int result;

  MaxTask(int[] arr, int start, int end) {
    this.arr = arr; this.start = start; this.end = end;
  }

  public void run() {
    if (end - start < 1)        result = Integer.MIN_VALUE;
    else if (end - start == 1) result = arr[start];
    else {
      int mid = (start + end) / 2;
      MaxTask t1 = new MaxTask(arr, start, mid);
      MaxTask t2 = new MaxTask(arr, mid, end);
      t1.start(); t2.start();
      try { t1.join(); t2.join(); } catch (InterruptedException i) {}
      result = max2(t1.result, t2.result);
    }
  }
}
```

(Part d) What is the total number of threads started by `MaxTask` when running an instance of `MaxTask(arr, 0, 8)`? How does this compare to the maximum number of parallel calls to `max()` you computed in Part b? Are there any blocked threads at any point during the execution? If yes, what is the maximum number of blocked threads? *(2p)*

(Part e) Assume that an instance of `MaxTask(arr, 0, 1000000)` is being run on a machine with a 4-core processor. Will this result in an efficient parallel computation? How would the situation be improved if `MaxTask` was instead implemented as a `RecursiveTask` and scheduled using a `ForkJoinPool` from the `java.util.concurrent` package? *(2p)*

Q5. Erlang bowling alley *(18p)*

Consider the following event handler function of an Erlang server. The server models a clerk working at a bowling alley, handing out shoes to customers.

```erlang
1 handler(Stock) ->
2   receive
3     {request, From, N} when N =< Stock ->
4       From ! ok,
5       handler(Stock - N);
6     {return, From, N} ->
7       From ! ok,
8       handler(Stock + N);
9     {stock, From} ->
10      From ! {ok, Stock},
11      handler(Stock)
12  end.
```

The clerk manages `Stock` pairs of shoes and needs to handle three types of events: customers may request or return shoes for `N` players, and the manager (or anyone else) may ask how many pairs of shoes the clerk currently has in stock. For the remainder of this question, assume the following *shoe dogma*:

> customers **never** return **more** shoes to the clerk than they previously obtained!

In the following, let C_1, C_2, ... C_n be the PIDs of customers processes, `S` the PID of the server, and `B` the PID of the boss. Assume the clerk starts off with 30 pairs of shoes, i.e. the server is spawned using `S = spawn(fun() -> handler(30) end)`.

(Part a) Assume the server receives the following sequence of messages *in this order*:

$$\{\texttt{request, } C_1\texttt{, 8}\}, \{\texttt{request, } C_2\texttt{, 25}\}, \{\texttt{return, } C_1\texttt{, 5}\}, \{\texttt{stock, B}\}.$$

Give a sequence of *responses* sent by the server. Your answer should have the format $P_1!r_1, P_2!r_2, \ldots$ where P_i are PIDs and r_i are server responses. *(1p)*

(Part b) Is your answer to Part a the only possible sequence of responses? If so, explain why. Does the order in which the server receives the messages impact its responses? *(2p)*

(Part c) Now, consider instead the following sequences of messages, given in the order they are received by the clerk (starting with an initial stock of 30 pairs of shoes):

1. $\{\texttt{request, } C_1\texttt{, 40}\}, \{\texttt{request, } C_2\texttt{, 5}\}$

2. $\{\texttt{request, } C_2\texttt{, 8}\}, \{\texttt{request, } C_1\texttt{, 25}\}, \{\texttt{request, } C_3\texttt{, 10}\}, \{\texttt{return, } C_2\texttt{, 8}\}, \{\texttt{request, } C_4\texttt{, 9}\}, \{\texttt{return, } C_3\texttt{, 10}\}$

Are all messages in these sequences eventually processed and answered by the server? If not, is it possible for customer processes to send *additional* messages to the server in such a way that all messages are answered by the server in the end? Give separate answers for each sequence and justify them. Your answers must respect the *shoe dogma*. *(3p)*

More robust servers

You will be asked to write two short pieces 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. Each of your answers should fit into at most 10 lines of code.

(Part d) Let `Total` be the size of the clerk's initial stock, or equivalently, the total number of pairs of shoes in the bowling alley, counting both the clerk's stock and the shoes worn by customers. What happens when a customer asks for `Total + 1` pairs of shoes? *(1p)*

(Part e) How could the server be changed so that customers requesting shoes for more than `Total` players receive a message `{error, Total}`? How should the state of the server (i.e. the arguments of `handler`) be adjusted? Propose *one new clause* to be inserted at the beginning of the **receive** expression in `handler`. You may write this clause in Erlang or pseudo code, but it should be of the form `P -> E` or `P when G -> E`, where `P` is a message pattern, `G` is a guard condition, and `E` is an expression. Your solution should work for arbitrary values of `Total`. *(4p)*

(Part f) Assume now that the server has been running for a while, handling out shoes to customers, so that `Stock < Total`. Suddenly, a customer `C` arrives requesting `N` pairs of shoes – more than the clerk has in stock, but such that `Stock < N <= Total`. Describe a scenario where `C`'s request is delayed indefinitely, while other customers continuously obtain and return shoes. What concurrency property is violated here? *(2p)*

(Part g) To avoid this type of scenario, we change the `handler` function of the original server as follows, inserting a new clause in the **receive** expression:

```erlang
handler(Stock) ->
  receive
    {request, From, N} when N =< Stock ->
      % ...
    {request, From, N} ->                    % request where 'Stock < N'
      NewStock = recover(Stock, N),
      From ! ok, handler(NewStock - N);
    % ...
```

The new clause calls the helper function `recover`, which behaves like `handler` but takes an additional argument `Goal` and ignores any `request` messages until `Stock >= Goal`, at which point it simply returns the updated stock. Complete the following unfinished implementation of `recover` (only replace the bit marked "todo"): *(5p)*

```erlang
recover(Stock, Goal) ->
  receive
    {return, From, M} -> %%%%% Implement this clause!  %%%%%
      todo;
    {stock, From} -> From ! {ok, Stock}, recover(Stock, Goal)
  end.
```

14

Appendix A. Full code listings

Appendix A.1. Code for Q1

```java
1  import java.util.concurrent.*;
2  import java.util.concurrent.locks.*;
3
4  class LockedCounterArray {
5    final private int[] counters;
6    final private Lock lock = new ReentrantLock();
7
8    public LockedCounterArray(int size) {
9      counters = new int[size];  /* all 0s by default */
10   }
11   public void increment(int index) {
12     lock.lock();
13     try { counters[index] = counters[index] + 1; }
14     finally { lock.unlock(); }
15   }
16   public int value(int index) {
17     lock.lock();
18     try { return counters[index]; }
19     finally { lock.unlock(); }
20   }
21   public static void main(String[] args) {
22     final int N = 10;
23     final int size = Integer.parseInt(args[0]);
24     LockedCounterArray lca = new LockedCounterArray(size);
25
26     Thread[] ts = new Thread[N];
27     for (int i = 0; i < N; ++i) {
28       ts[i] = new Thread() {
29           @Override public void run() {
30             for (int j = 0; j < size; ++j) { lca.increment((j + N) % size); }
31           }
32         };
33       ts[i].start();
34     }
35     for (int i = 0; i < N; ++i) {
36       try { ts[i].join(); } catch (InterruptedException ie) {}
37     }
38     for (int j = 0; j < size; ++j) {
39       System.out.println("Counter " + j + ": " + lca.value(j));
40     }
41   }
42 }
```

Appendix A.2. Code for Q2 and Q3.

```java
1  import java.util.concurrent.*;
2
3  class BarrierThread extends Thread {
4    private int id, work;
5    public static Semaphore[] barrier = { new Semaphore(0), new Semaphore(0) };
6
7    BarrierThread(int id, int work) { this.id = id; this.work = work; }
8
9    boolean doWork() throws InterruptedException {
10     if (work > 0) {              // more work?
11       sleep(500 + id * 100);  // do "work"
12       work = work - 1;         // update work counter
13       return false;
14     } else return true;         // done!
15   }
16
17   @Override public void run() {
18     int me = id;
19     int other = 1 - id;
20     try {
21       boolean done = false;
22       do {
23         // synchronize: wait for both threads to reach the barrier.
24         System.out.println("Thread " + me + " synchronizing...");
25         barrier[me].release();     // I'm done.
26         barrier[other].acquire();  // Wait for the other thread.
27         System.out.println("Thread " + me + " ready to work!");
28         // do work.
29         done = doWork();
30       } while (!done);
31     } catch (InterruptedException i) {}
32   }
33
34   public static void main(String[] args) {
35     final int N1 = Integer.parseInt(args[0]);
36     final int N2 = Integer.parseInt(args[1]);
37     BarrierThread t0 = new BarrierThread(0, N1);
38     BarrierThread t1 = new BarrierThread(1, N2);
39     t0.start(); t1.start();
40     try {
41       t0.join(); t1.join();
42     } catch (InterruptedException i) {}
43   }
44 }
```

Appendix A.3. Code for Q4

```java
1  import java.util.concurrent.*;
2
3  class MaxTask extends Thread {
4    private int[] arr;
5    private int start, end;
6    public int result;
7
8    MaxTask(int[] arr, int start, int end) {
9      this.arr = arr; this.start = start; this.end = end;
10   }
11
12   static int max2(int x, int y) {
13     if (x > y) return x;
14     else       return y;
15   }
16
17   @Override public void run() {
18     if (end - start < 1)       result = Integer.MIN_VALUE;
19     else if (end - start == 1) result = arr[start];
20     else {
21       int mid = (start + end) / 2;
22       MaxTask t1 = new MaxTask(arr, start, mid);
23       MaxTask t2 = new MaxTask(arr, mid, end);
24       t1.start(); t2.start();
25       try { t1.join(); t2.join(); } catch (InterruptedException i) {}
26       result = max2(t1.result, t2.result);
27     }
28   }
29
30   public static void main(String[] args) {
31     final int N = 20; int[] numbers = new int[N];
32     for (int i = 0; i < N; ++i) { numbers[i] = i; }
33     MaxTask m = new MaxTask(numbers, 0, N);
34     m.run();
35     System.out.println("Max: " + m.result);
36   }
37 }
```

Appendix A.4. Code for Q5

```erlang
1  -module(shoes).
2  -export([init/1,request/1,return/1,stock/0]).
3
4  % start the clerk process managing a stock of N pairs of shoes
5  init(N) ->
6      register(clerk, spawn(fun () -> handler(N) end)).
7
8  % request N pairs of shoes
9  request(N) ->
10     clerk ! {request, self(), N},
11     receive ok -> ok end.
12
13 % return N pairs of shoes
14 return(N) ->
15     clerk ! {return, self(), N},
16     receive ok -> ok end.
17
18 % ask the clerk what the current stock is
19 stock() ->
20     clerk ! {stock, self()},
21     receive {ok, N} -> N end.
22
23 % the clerk's event handler
24 handler(Stock) ->
25   receive
26     {request, From, N} when N =< Stock ->
27         From ! ok,
28         handler(Stock - N);
29     {return, From, N} ->
30         From ! ok,
31         handler(Stock + N);
32     {stock, From} ->
33         From ! {ok, Stock},
34         handler(Stock)
35   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:

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