

G52CON: Concepts of Concurrency

Lecture 19 Revision

Brian Logan
School of Computer Science & IT
bsl@cs.nott.ac.uk

Outline of this lecture

- exam format and scope
- how to revise
- overview of the module
- sample exam question

Copies of past papers

- past papers for the 02-03, 03-04, 04-05, 05-06 and 06-07 sessions are available from the exams office web page:

– <http://www.nottingham.ac.uk/is/gateway/exampapers>

- the module changed in 05-06—exam format for that session differs
- other useful sources include the unassessed exercises, and exercises from textbooks

Exam format

There will be six questions:

- one compulsory multiple choice question
- choose *three* questions from the five remaining
- please do not attempt more than four questions
- note that the exam covers the *whole* syllabus, and is not limited to the material on the lecture slides

What you need to know

do we need to learn code?

- you will need to *understand* how the example programs given in the lectures and in the recommended reading work—simply memorising the code examples from the lectures won't help you very much.

is the material on the slides everything we need to know or should we support the information using the recommended reading?

- the material on the slides is **not** everything you need to know, and there is a warning on the module web page to this effect.

Revision

There are several ways to revise for the exam:

- first you should look at your own notes
- all the slides for each lecture are available on-line, as are the lists of suggested reading, the unassessed exercises and the model answers
- textbooks are useful, particularly if there is something in your notes you don't understand
- unassessed exercises and exercises from textbooks can also be very useful in giving you more practice at solving particular types of problems.

Outline Syllabus

The module covered of four main themes:

- introduction to concurrency;
- design of simple concurrent algorithms in Java;
- correctness of concurrent algorithms; and
- design patterns for common concurrency problems.

Topics covered

- mutual exclusion and condition synchronisation
- atomic actions
- mutual exclusion algorithms: Test-and-Set, Peterson's algorithm
- semaphores
- monitors
- distributed processing: message passing, RPC and rendezvous
- correctness of concurrent programs
- concurrent programming in Java (shared memory & distributed processing)

- common concurrency problems: e.g., Producer-Consumer, Readers and Writers, Client-Server

Lectures and topics

Lectures often cover more than one topic, e.g., a lecture about a concurrency primitive like monitors might contain:

- something about how monitors are implemented
- how they compare to other primitives
- the safety and liveness properties of solutions built using them
- examples of how they are used to solve a one of the common concurrency problems, e.g., bounded buffer problems

Mutual Exclusion & Condition Synchronisation

- mutual exclusion
- condition synchronisation
- interference
- critical sections
- classes of critical sections

Atomic actions

- process switching
- atomic actions
- memory accesses
- machine instructions
- multiprocessors
- mutual exclusion protocols

Mutual Exclusion algorithms

- archetypical mutual exclusion problem
- general form of a solution
- Test-and-Set algorithm
- Peterson's algorithm

Semaphores

- P and V operations
- general and binary semaphores
- blocking
- mutual exclusion & condition synchronisation with semaphores
- properties of semaphore solutions & problems of semaphores
- how to use semaphores

Monitors

- components of a monitor
- mutual exclusion & condition synchronisation in monitors
- operations on condition variables
- signalling disciplines
- how to use monitors

Distributed processing

- processes and channels
- message passing: asynchronous & synchronous message passing
- remote invocation: RPC & (extended) rendezvous
- modules & synchronisation

Correctness of concurrent programs

- safety properties: e.g., mutual exclusion, absence of deadlock, absence of unnecessary delay
- liveness properties: e.g., eventual entry
- properties of algorithms: e.g., Test-and-Set, Peterson's algorithm
- proof-based approaches to verification: e.g., assertional reasoning
- model-based approaches to verification: e.g., model checking, CTL specifications

Concurrent programming in Java

- Java threads
- monitors and Java
- Java memory model
- synchronisation: mutual exclusion & condition synchronisation
- distributed processing in Java: `java.rmi`

Producer-Consumer problems

- the role of buffering
- different sizes of buffer: e.g., single buffer, double buffer, bounded buffer
- implementations: e.g., semaphores & monitors
- applications

Readers and Writers problems

- Reader & Writer processes
- synchronisation requirements and scheduling policies: e.g., readers' preference, writers' preference, fair solutions
- implementations: e.g., semaphores & monitors
- applications

Client-Server problems

- data flow in concurrency problems
- patterns of communication
- examples:
 - chat server
 - time server using RPC

A sample question

(a) Briefly describe the major stages in the lifecycle of a thread and explain how the transitions between stages occur. (5)

(b) A savings account is accessed by several processes. A process making a deposit never has to delay (except for mutual exclusion), but a withdrawal has to wait until there are sufficient funds. Develop a Java `SavingsAccount` class which allows safe concurrent update of `SavingsAccount` objects. The class should have three public methods:

- `void deposit(int amount)` which adds amount to the current balance;
- `void withdraw(int amount)` which subtracts amount from the current balance;
- `int balance()` which returns the current balance.

Assume the arguments to `deposit` and `withdraw` are positive. Explain your answer. (10)

(c) Extend your solution to part (b) to incorporate an additional method

- `void transfer(SavingsAccount account, int amount)` which transfers amount from the account `account` to the balance of this account.

Explain clearly the problem(s) that must be addressed in designing the method, how your solution solves these problems and why it is correct. (10)

How to answer

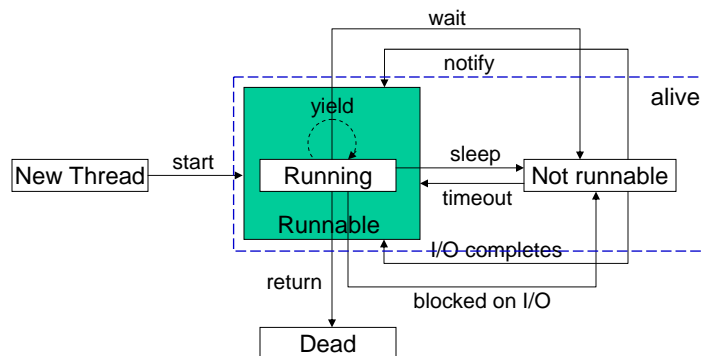
- the number of marks gives some indication of the level of detail required, or the likely difficulty in the case of a problem solving question
- read the question *carefully* and answer what is asked—if it asks for a solution using, e.g., semaphores, then a monitor solution will get zero marks
- *explain* your answer if you are asked to do so
- attempt *all* parts of the question

A sample answer

(a) There are three states in the lifecycle of a thread:

- creation: we make a new Thread by calling the constructor of a subclass of the Thread class or by passing an instance of a class which implements Runnable to the Thread constructor (1)
- alive: calling the start() method on the Thread object invokes its run() method as an independent activity. After a Thread has been started, it is said to be alive. A thread which is alive is either runnable or not runnable. Calling wait() or sleep() or blocking on I/O will make a running thread not runnable. Calls to notify(), timeouts and the completion of I/O will make a not runnable thread runnable. (3)
- termination: the Thread terminates when its run method completes, either by returning normally or by throwing an unchecked exception (RuntimeException, Error or one of their subclasses). (1)

Thread lifecycle



A sample answer

```
class SavingsAccount {
    private long balance = 0;

    public void synchronized deposit(int amount) {
        balance = balance + amount;
        notifyAll();
    }

    public void synchronized withdraw(int amount)
        throws InterruptedException {
        while (amount > balance)
            wait();
        balance = balance - amount;
    }

    public long synchronized balance() { return balance; }
}
```

© Brian Logan 2007

G52CON Lecture 20: Revision

25

A sample answer

(b) cont.

The current balance is held in the private long variable, balance.

The deposit and withdraw operations are implemented as synchronized methods. deposit simply increases the current balance by amount and signals all threads waiting to make a withdrawal that additional funds are now available. deposit never blocks (except for mutual exclusion), since making a deposit is always legal, whatever the balance.

Withdraw checks to see if there are sufficient funds to cover the withdrawal. If there are, it simply decreases the balance by amount. If there are insufficient funds, it waits until a deposit operation signals that there are more funds available and then tries again. Note that the call to wait() is enclosed in a while loop to ensure that the condition on which we are waiting—there being sufficient funds to cover the withdrawal—is actually true when we update the balance. Using notifyAll() wakes all processes waiting to make a withdrawal, allowing withdrawals to proceed up to but not exceeding the new balance. This could either be one large withdrawal, or several smaller withdrawals.

© Brian Logan 2007

G52CON Lecture 20: Revision

26

A sample answer

```
public void transfer(SavingsAccount other, int amount) {
    if (System.identityHashCode(this) <
        System.identityHashCode(other))
        this.doTransferFrom(other, amount);
    else
        other.doTransferTo(this, amount);
}

protected synchronized void doTransferFrom(SavingsAccount other, int amt){
    other.withdraw(amt);
    deposit(amt);
}

protected synchronized void doTransferTo(SavingsAccount other, int amt) {
    withdraw(amt);
    other.deposit(amt);
}
```

A sample answer

(c) cont.

The key problem is avoiding deadlock due to circular waiting, i.e., if two threads try to transfer money between the same two accounts, we don't end up with one thread holding the lock on each account and unable to obtain the lock on the other. (The problem of ensuring that there are sufficient funds in the account from which the money is being transferred is already solved by the withdraw method)

The problem can be solved by ensuring that threads always acquire the locks in the same order. One way to do this is to associate a unique tag with each object and require all threads to lock objects in tag order. For example, we can use the SavingsAccount's hash codes (e.g., System.identityHashCode()) to order them (while it is possible for two objects to have the same hash code, it is unlikely, and this avoids allocating a tag to each new SavingsAccount object created).