PAPER CODE NO.	EXAMINER : Boris Konev
COMP212	DEPARTMENT : Computer Science Tel. No. 0151 795 4260

UNIVERSITY OF
LIVERPOOL

May 2007 EXAMINATIONS

Bachelor of Arts: Year 2
Bachelor of Science : Year 2
No qualification aimed for: Year 1

Distributed Systems

TIME ALLOWED : Two Hours

INSTRUCTIONS TO CANDIDATES

Answer FOUR questions.

If you attempt to answer more questions than the required number of questions (in any section), the marks awarded for the excess questions answered will be discarded (starting with your lowest mark).

UNIVERSITY OF LIVERPOOL

Question 1

A. Give four examples of resources that may be shared in a distributed system. **2 marks**

B. Explain what is meant by (distribution) transparency. **4 marks**

C. What is the role of middleware in a distributed system? **4 marks**

D. Name two middleware communication services. **2 marks**

E. In Java RMI, all methods of a remote object must be declared able to throw instances of a particular subclass of the Exception class. What is this subclass? **3 marks**

F. Explain the role of the data level in a three-tired client/server architecture. **3 marks**

G. Consider a network consisting of 5 computers, A (coordinator), B, C, D, and E. At $14:40$ the coordinator decides to synchronise the clock of all computers in the network. At that moment, the clock of every computer in the network shows the following.

Computer	Clock
A	14:40
B	14:37
C	14:35
D	14:28
E	14:25

Apply the Berkley clock synchronisation algorithm to this situation, show the stages of computation, and write what will be the outcome of the synchronisation. Assume that the time needed for computation and for network communication is negligible. **5 marks**

H. Communication in Jini is based on which of the following:

a) message queueing,

b) Java RMI, or

c) remote procedure calls? **2 marks**

UNIVERSITY OF LIVERPOOL

Question 2

A. Classify the following protocols: IP, TCP, HTTP according to the ISO OSI classification. **3 marks**

B. Describe precisely what is meant by a scalable system. **3 marks**

C. Give at least one advantage of threads over processes. **2 marks**

D. Why is it sometimes so hard to hide the occurrence and recovery from failures in a distributed system? **4 marks**

E. Why are transport-level communication services often inappropriate for building distributed applications? **3 marks**

F. Processes in distributed systems are often divided into two groups: clients and servers.

(a) Describe the client-server model. **3 marks**

(b) Give a graphical representation of the request-reply interaction between a client and a server. **2 marks**

G. Consider a group of eight processors numbered 0 to 7. Previously process 7 was the coordinator, but it has crashed.

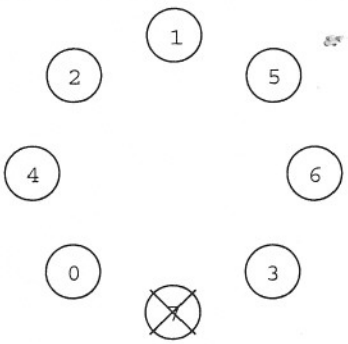

Process 6 is the first one to notice this, so it initiates a new election by the ***bully*** algorithm. Apply the algorithm to the described situation, show messages being sent between the processes, and find who will be the new coordinator. **5 marks**

UNIVERSITY OF LIVERPOOL

Question 3

A. **(a)** What is the purpose of a registry for object-oriented remote method invocation (RMI)? **2 marks**

(b) If a client invokes the same remote method for a second time, does it necessarily use the registry? Explain your reasoning. **2 marks**

B. Describe the difference between *persistent* and *transient* communication. **3 marks**

C. Why are threads (or processes) so important for distributed systems? **4 marks**

D. Give an example of a stateless protocol. **2 marks**

E. Imagine a Web server that maintains a table in which client IP addresses are mapped to the most recently accessed Web pages. When a client connects to the server, the server looks up the client in its table, and if found, returns the registered page. In your opinion, is this server stateful or stateless? Present an argument to support your opinion. **5 marks**

F. Is it possible to connect across a network two components of a distributed system written in different languages (e.g., COBOL and Java) using a *distributed object-based system*? Is it possible to connect such components by means of Java RMI? Give brief justifications for your answers. **4 marks**

G. Are processes using Jini for coordination required to coexist at the same time? Explain why/why not. **3 marks**

UNIVERSITY OF LIVERPOOL

Question 4

A. Give two examples of identifiers (here, by the term "identifiers" we mean those identifiers that refer to entities in distributed systems). **2 marks**

B. Which name server addresses do DNS name servers hold by default, and why? **4 marks**

C. Describe the relative advantages and disadvantages of iterative and recursive name resolution in a distributed naming service. **4 marks**

D. Recall that a *distributed commit protocol* aims at having an operation executed by all members of a process group or by none of them.

(a) Consider the following **one-phase protocol**: the coordinator simply tells all other processes to perform (locally) the operation. Give your reasoning why such a one-phase protocol will not achieve the goal of a distributed commit protocol. **4 marks**

(b) Describe the **two-phase distributed commit protocol** (2PC). Your description need not be complete, but should convey the idea on which the protocol is based. **4 marks**

E. Can one name belong to more than one namespace? If yes, give an example. **3 marks**

F. The distributed file system Coda allows a client to continue working with a shared file even if there is no network connection between the client and the server. Explain how this is made possible. **4 marks**

UNIVERSITY OF LIVERPOOL

Question 5

A. Name two election algorithms that can be used to agree on the new coordinator in a distributed system. **2 marks**

B. A clock is reading 10:15:35.0 (hr:min:sec) when it is discovered to be 10 sec fast. Are there any drawbacks in setting it back to the correct time at that point? **2 marks**

C. Give a description of a distributed token ring mutual exclusion algorithm. Illustrate your description with a drawing. **4 marks**

D. Do Lamport's Timestamps achieve total or partial ordering of events in a system where not all nodes participate in the communication? Explain your reasoning. **5 marks**

E. The following is Ricart and Agrawala's distributed mutual exclusion algorithm considered in the lectures.

- A process wanting to enter a critical section (CS) sends a message with (`cs name, process id, current time`) to all processes (including itself).
- When a process receives a CS request from another process, it reacts based on its current state with respect to the CS requested as follows:
  - If the receiver is not in the CS and it does not want to enter the CS, it sends an OK message to the sender.
  - If the receiver is in the CS, it does not reply and queues the request.
  - If the receiver wants to enter the CS but has not yet, it compares the timestamp of the incoming message with the timestamp of its message sent to everyone. The *lowest* timestamp wins.
    - If the incoming timestamp is *lower*, the receiver sends an OK message to the sender.
    - If its own timestamp is *lower*, the receiver queues the request and sends nothing.
  - After sending out requests asking permission to enter a critical section, a process waits until everyone else has given permission. As soon as all the permissions are in, it may enter the critical section. When it exits the section, it sends OK to all the processes on its queue and deletes them all from the queue.

Give reasons why the algorithm satisfies the

- Safety
- Liveness

properties. **6 marks**

F. What is the purpose of an Internet firewall? Describe the packet-filtering technique used in firewalls. **4 marks**

G. Is it possible to describe *semantics* of objects and services in CORBA IDL? **2 marks**

UNIVERSITY OF LIVERPOOL

Question 6

A. Is agreement possible in a system with unreliable communication? **2 marks**

B. What purpose is achieved by means of physical triple modular redundancy? **3 marks**

C. Can the model of triple modular redundancy handle Byzantine failures? **3 marks**

D. Consider two communication services for use in asynchronous distributed systems. In service A, messages may be lost, corrupted, duplicated or delayed. In service B, messages may be lost, delayed or delivered too fast but those that are delivered arrive ordered and with the correct contents.

(a) Describe the classes of failure exhibited by each service **2 marks each**

(b) Can service B be described as a reliable communication service? **2 marks**

E. If Alice wants to send secret information to Bob, should she know Bob's public or Bob's private key? **2 marks**

F. Describe what an *access control matrix* is. **3 marks**

G. Compare secret key encryption with public key encryption. Give two advantages of secret key over public key encryption and two advantages of public key over secret key encryption. Can a protocol use both secret and public key encryptions to have the advantages of both? **6 marks**