

Two hours

Question ONE is COMPULSORY

**UNIVERSITY OF MANCHESTER
SCHOOL OF COMPUTER SCIENCE**

Distributed Computing

Date: Friday 26th May 2017

Time: 09:45 - 11:45

**Please answer Question ONE and also TWO other Questions
from the remaining THREE Questions provided**

This is a CLOSED book examination

The use of electronic calculators is NOT permitted

[PTO]

1. **Compulsory**

- a) For each item in column one, choose the best match from column two. Each item in column two should be used only once.

<i>Column 1</i>	<i>Column 2</i>
1. Checkpoint	a. Backward Recovery
2. Edge Chasing	b. Cloud Computing
3. Eventual Consistency	c. Distributed Deadlock
4. Iterative/Recursive Resolution	d. Domain Name System
5. Publish-Subscribe	e. Messaging
6. Virtualisation	f. Replicated Storage

(6 marks)

- b) Give one example for each of: *physical redundancy*, *time redundancy*, *information redundancy*. (2 marks)
- c) When using Java RMI, what is the purpose of the *rmiregistry*? (2 marks)
- d) What is the worst-case scenario for the number of messages in the ring-based election algorithm? How many messages do you expect in this case? (2 marks)
- e) Explain briefly what *Little's Law* is. (2 marks)
- f) In the context of lab exercise 2, what would you do to launch a *denial of service attack* against the server? (2 marks)
- g) Suppose the C function below is to be made available to remote processes using RPC. What particular implementation problem does this highlight? Why is the proposed solution only partly satisfactory?

```
void doIt (int *p, int *q) { p++ ; q-- ; return ; }
```

(4 marks)

Answer any two from Questions TWO, THREE and FOUR

2. a) Explain briefly why some applications are not parallelisable. Describe Amdahl's law and explain what it can be used for. (4 marks)
- b) Explain briefly what the four properties commonly denoted by the acronym ACID are when referring to transactions. (4 marks)
- c) In the context of lab exercise 2, give an example of how a client can hold more than the maximum permitted number of reservations if the server may concurrently process message requests and the server code for processing requests to reserve a slot is not enclosed in an ACID transaction. (4 marks)
- d) A service is replicated onto 3 computers.
- The first computer, A, has a mean time between failures of 1 day.
 - The second computer, B, has a mean time between failures of 1.75 days.
 - The third computer, C, has a mean time between failures of 6 days.

When a failure occurs, it takes on average 6 hours to fix.

- i) What is the availability of the replicated service? (2 marks)
- ii) What would the availability of the replicated service be if only computers A and B were used? (2 marks)
- iii) Describe how in the general case of n computers, each with a mean time between failures f_i and a time to fix the failures t_i , you would choose the two computers that provide the highest availability. (4 marks)

[PTO]

3. a) Describe all the operations that take place during a Remote Procedure Call (RPC). (4 marks)

b) A client attempts to synchronize with a time server using Cristian’s algorithm. It records round-trip times (in msec) and timestamps returned by the server as follows:

Round-trip time (milliseconds)	Timestamp hh:mm:ss.msec
22	13:25:23.674
25	13:25:25.450
20	13:25:28.342

Which of these times should the client use to set its clock? To what time should it set it? If the minimum communication time (one way) is 8 msec, estimate the accuracy of the new time.

(6 marks)

c) In a system containing 7 computers, identified by the integers 1-7, the coordinator is chosen by the Bully algorithm to be the live one with the highest identifier. Assume for this part that all messages are delivered promptly, and that the computers and the network are entirely reliable.

At a certain point in time the coordinator (computer 7) and the computer with the second-highest identifier (computer 6) crash. How many messages in total are sent if the computer with identifier 1 is the computer discovering a crash and triggering an election? You need to count all types of messages that the algorithm sends. You can assume that computers with identifiers 6 and 7 remain crashed during the election.

(6 marks)

d) Suppose that a server, similar to the server used in lab exercise 2, implements a queue to store incoming client requests. Four threads, running on different CPUs (cores) on the server, concurrently fetch and then process requests from the queue, one request per thread at a time. Every request requires processing time of 50 milliseconds and, on average, there are about 5000 requests per hour sent to the server. How long do you expect to have requests waiting in the queue before their processing starts? (4 marks)

4. a) Give the reasoning that leads to the conclusion that the commanding generals of two separate armies cannot agree a plan of attack using solely a finite sequence of unreliable messages. (3 marks)
- b) Outline the Byzantine Generals problem and illustrate how one of the three being a traitor makes a solution impossible, whereas with one of four it is achievable. (6 marks)
- c) Explain the relevance of parts (a) and (b) above to distributed computing systems. (3 marks)
- d) In a client-server application, assume that each client request is added to a server queue and three servers can serve requests from the queue. Server A can serve 29 requests per second. Server B can serve 34 requests per second. Server C can serve 37 requests per second.
- (i) Assume that at a certain point in time there are 1000 requests in the queue. How would you allocate these requests to the servers to achieve load balancing? (2 marks)
- (ii) If at most 50 clients can operate in parallel, how many requests per second would you advise each client to make so that the server queue is not flooded with requests? (2 marks)
- (iii) Describe how in the general case of n servers, each capable of serving s_i requests per second, you would allocate x requests from the queue to achieve load balancing. State any assumptions you make. (4 marks)

END OF EXAMINATION