

Two hours

**UNIVERSITY OF MANCHESTER
SCHOOL OF COMPUTER SCIENCE**

Distributed Computing

Date: Wednesday 23rd May 2018

Time: 14:00 - 16:00

Please answer all Questions.

© The University of Manchester, 2018

This is a CLOSED book examination

The use of electronic calculators is NOT permitted

[PTO]

1. a) Indicate if the following statements are True or False:
- (i) In distributed systems, it is useful to assume that latency is zero.
 - (ii) The C in ACID stands for Concurrency.
 - (iii) Most client-server systems implement exactly-once semantics.
 - (iv) In Byzantine failures, faulty components may exhibit arbitrary behaviour. (4 marks)

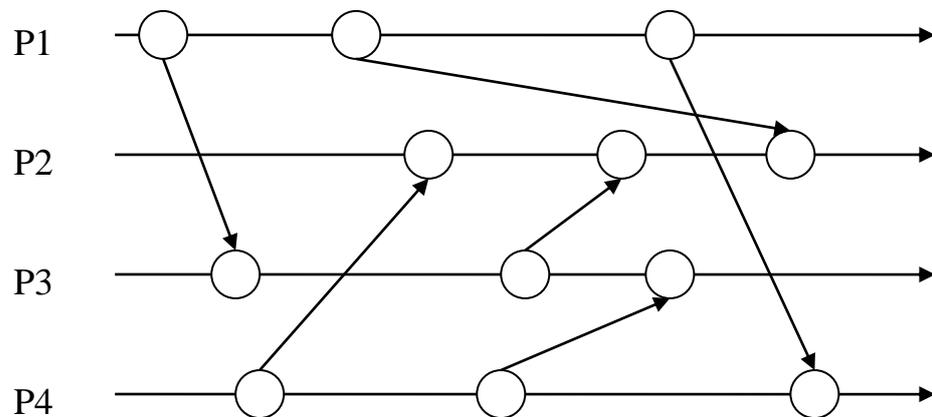
- b) For each item in column one, choose the best match from column two. Each item in column two should be used only once.

<i>Column 1</i>	<i>Column 2</i>
1. Checkpoint	a. Backward Recovery
2. Edge Chasing	b. Cloud Computing
3. Infrastructure as a Service (IaaS)	c. Distributed Deadlock
4. Marshalling/Unmarshalling	d. Distributed Transactions
5. Strict Consistency	e. Remote Procedure Call
6. Two-Phase Commit Protocol	f. Replicated Storage

(6 marks)

- c) In a distributed system, what is the purpose of an IDL? (2 marks)
- d) When using Java RMI, what is the purpose of the *rmiregistry*? (2 marks)
- e) Explain briefly what is meant by the term *middleware*. (2 marks)
- f) Why is it practically impossible to achieve exact synchronisation of clocks in a distributed system? (2 marks)
- g) In the context of lab exercise 2, what would you do to launch a *denial of service attack* against the server? (2 marks)

2. a) Describe clearly all the operations that take place during a Remote Procedure Call (RPC). (4 marks)
- b) Explain briefly what the four properties commonly denoted by the acronym ACID are when referring to transactions. (4 marks)
- c) Consider the figure below, which shows 4 processes and a number of communication events taking place over a period of time.



Calculate the value of Lamport clocks and vector clocks for each of the 12 events shown above. You can assume that all logical clocks start initially with zeros. (6 marks)

- d) The following two processes access the shared variables x , y , z . Each process accesses a different replica of the store used to hold these variables. Before any process starts executing, the value of all three variables, x , y , z , is 0 in all the replicas.

Process A
 $x=1;$
 if ($y==0$) $z++;$

Process B
 $y=1;$
 if ($x==0$) $z++;$

- (i) When both processes have completed executing the statements given, what are the possible values of z , if the replication uses the sequential consistency model? Justify your answer. (3 marks)
- (ii) When both processes have completed executing the statements given, what are the possible values of z , if the replication uses the causal consistency model? Justify your answer. (3 marks)

[PTO]

3. a) Explain briefly the relevance of the Byzantine Generals problem to distributed computing systems. (2 marks)
- b) Two computers are used to provide a replicated service. Each computer has a mean time between failures of 12 days; a failure takes on average 12 hours to fix. What is the availability of the replicated service? (3 marks)
- c) Consider a simple server that continuously receives and executes client requests without accessing other servers. Explain why it is generally not possible to set an upper bound on the time taken by such a server to complete a client request. (3 marks)
- d) Suppose a network of machines where one of the machines is used to provide (distributed) mutual exclusion service. When the machine supporting this service gets overloaded with other tasks it needs to find the least loaded machine in the network, and pass over the provision of the mutual exclusive access service to a process on that machine. Two algorithms are being considered for this. The first is to have the server ask each machine about its workload and then notify all the clients with the identity of the new server. The second is to use a ring-based election, initiated by the current server. Fully describe the latter, clearly stating any assumptions you make, (4 marks) and compare it with the former with respect to the number of messages passed. (4 marks)
- e) In the context of lab exercise 2, suppose that the following pseudo-code describes a client strategy to reserve the earliest common slot for two different slot tables: band and hotel. Explain why this strategy may not always achieve the desired result. State any assumptions you make.

```
// assume 1 sec delays and all requests are correctly implemented
cancel all current bookings;
repeat
    get availability for band;
    get availability for hotel;
    find earliest_common_slot;
    reserve earliest_common_slot for hotel;
    reserve earliest_common_slot for band;
until both hotel and band reservation successful;
print "reserved:" earliest_common_slot;
```

(4 marks)

END OF EXAMINATION