



THE UNIVERSITY  
OF QUEENSLAND

VENUE:

SEAT NUMBER:

STUDENT NUMBER:

**FINAL EXAMINATION**

**St Lucia Campus**

**First Semester, 2007**

**CSSE4004/7014 – Distributed Computing**

**PERUSAL TIME**                      10 mins. During perusal, write only on this exam paper.

**WRITING TIME**                      2:00 Hours

**EXAMINER**                          Jadwiga Indulska

**NO. OF PAGES** (*include title page and attachments*) 5 Double Sided

**Exam Type:**                          Open Book - Unrestricted Materials

**Permitted Materials:**            Calculator - Yes - Non-programmable calculators only  
Dictionary - Yes - Any unmarked paper dictionary is permitted  
Other – No electronic aids are permitted (e.g. laptops, phone)

**Answer:**                                In writing booklet

**Number of Questions:**            9

**Weighting/Marks:**                65%/100marks

**Special Instructions:**            Students must comply with the General Award Rules 1A.7 and 1A.8  
which outline the responsibilities of students during an examination.

**THIS EXAMINATION PAPER MUST NOT BE REMOVED  
FROM THE EXAMINATION ROOM**

**Question 1 (10 marks)**

Address the following issues about communication in distributed systems:

- (a) Explain how dynamic method invocation works. (3 marks)
- (b) Do message queuing systems provide transient and asynchronous communication? Explain your answer. (3 marks)
- (c) Routing tables in current application level messaging systems are configured manually. Describe a protocol which would provide automatic configuration of the routing tables. (4 marks)

**Question 2 (5 marks)**

Answer the following questions related to naming in distributed systems:

- (a) Is the name “www.qut.edu.au/index.html” location independent? Explain. (2 marks)
- (b) In a hierarchical location service with a depth of  $n$ , how many location records need to be updated at most when a mobile entity changes its location? (3marks)

**Question 3 (8 marks)**

Address the following issues about distributed systems:

- (a) What is the difference between a distributed computing middleware and a network operating system? Does the middleware depend on services of the network operating system? (4 marks)
- (b) A process in a distributed system runs on one node and accesses data from another node. After some time, for load balancing purposes, this process relocates to a different node. What kind of transparencies should be provided for this process in a distributed system? (4 marks)

**Question 4 (10 marks)**

Address the following issues related to synchronization in distributed systems:

- (a) Diagram 1 illustrates local times in three different nodes and also communication between three processes running on these nodes. Introduce total order to the communication events by applying Lamport’s “happened before” relationship. Redraw the diagram to show the changes. (6 marks)
- (b) Discuss pros and cons for centralized and distributed algorithms for mutual exclusion. (4 marks)




Diagram 1

**Question 5 (16 marks)**

Address the following issues related to the distributed computing platforms:

- Explain the role of Interface Definition Languages (IDL) in distributed computing. List advantages of using IDL. (4 marks)
- Explain the sequence of actions taken in CORBA if a client uses a persistent CORBA object reference to invoke a method when the server responsible for this object is not running.  
Explain the role of object references in such flexible invocations (7 marks)
- What kind of distributed applications are more suitable for Web Services than for .NET (or the RMI service in WCF)? Explain why. (5 marks)

**Question 6 (13 marks)**

The following question relates to distributed file systems and pervasive systems:

- Show the differences in the Coda design, compared to NFS, which stem from the fact that Coda was built for availability. (4 marks)
- Which feature of the naming used in Coda makes the naming more suitable to distributed systems than the naming used in NFS?  
Explain why. (4 marks)
- Two commonly used communication paradigms in context-aware applications are RMI and asynchronous notifications. Explain why. Support your explanation with application examples. (5 marks)

**Question 7 (13 marks)**

The following questions relate to fault tolerance:

- (a) Explain problems with designing fault tolerant multicasting. (6 marks)
- (b) There is a group of seven replicated processes producing some results. Assume that these processes can only crash silently; what is the maximum number of these processes that may crash while it is still possible to get the correct results? What is the maximum number of Byzantine processes in this group to achieve the correct results? Explain. (4 marks)
- (c) Assume that a group of processes has to reach an agreement on electing a coordinator. There is one process in the group which exhibit Byzantine failures. How many processes have to be in this group to reach the agreement? (3 marks)

**Question 8 (17 marks)**

The following questions relate to consistency and replication:

- (a) Explain the differences between sequential consistency and entry consistency. Which of these models would be better, from the performance point of view, for typical business applications operating on sets of data? (4 marks)
- (b) What kind of consistency is provided for replicated databases if all operation requests are delivered by totally ordered multicast? (3 marks)
- (c) Which of the following are valid combinations of read and write quorums for Gifford's scheme (quorum-based protocols)? Assume  $N = 8$ . Explain each answer.

$N_R = 2, N_w = 6$  (1 mark)

$N_R = 3, N_w = 6$  (1 mark)

$N_R = 1, N_w = 8$  (1 mark)

$N_R = 5, N_w = 4$  (1 mark)

- (d) Assume that you are to design an email application that needs to provide the *monotonic-read* consistency to mobile users. Would the Gifford's scheme based protocol or the primary-backup-local-write protocol be the right protocol to implement it? Explain why. (6 marks)

**Question 9 (8 marks)**

The following questions relate to processes:

- (a) Discuss the advantages/disadvantages of a multi-threaded server compared to a server which spawns multiple processes. (4 marks)
- (b) We want to compare the performance of a single-threaded and a multi-threaded file server. The following assumptions are made. It takes 10ms to get a request, dispatch it and do the rest of the necessary processing involved in serving the file, assuming the file is cached in main memory. If the file is not cached, a disk operation is needed in which case an additional 50 ms is required, during which the thread sleeps. We assume that for one third of all requests, the file can be served from the cache.
- a. How many requests per second can the single-threaded server handle? (2 marks)
- b. How many requests per second can the multi-threaded server handle? Assume that there is always enough threads. (2 marks)