



THE UNIVERSITY  
OF QUEENSLAND

VENUE:

SEAT NUMBER:

STUDENT NUMBER:

St. Lucia Campus

**FINAL EXAMINATION**

**First Semester, 2008**

**CSSE4004/CSSE7014- Distributed Computing**

**PERUSAL TIME** 10mins.

**WRITING TIME** 120mins

**EXAMINER** Prof. Jadwiga Indulska

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

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

**Permitted Materials:** Calculator - Yes - Any type of calculator is permitted  
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 (8 marks)**

Address the following issues about distributed systems:

- (a) Explain what a middleware for distributed systems adds to the functionality of network operating systems. Show the architecture of distributed systems and position middleware and network operating systems in it. (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 2 (12 marks)**

Address the following issues about communication in distributed systems:

- (a) Explain how dynamic method invocation works. (3 marks)
- (b) Do messaging systems such as IBM's WebSphere MQ provide transient and synchronous 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. (3 marks)
- (d) What are the pros and cons of using epidemic protocols for disseminating information? (3 marks)

**Question 3 (9 marks)**

Answer the following questions related to naming in distributed systems:

- (a) Is the name "www.uq.edu.au/index.html" location independent? Explain. (2 marks)
- (b) Suppose that it is known that a specific mobile entity will almost never move outside domain  $D$ , and if it does, it can be expected to return soon. How can this information be used to speed up the lookup operation in a hierarchical location service? (4 marks)
- (c) 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? (3 marks)

**Question 4 (12 marks)**

Address the following issues related to synchronization in distributed systems:

- (a) Give two examples of distributed applications which require total (logical) ordering of events. Explain why. (4 marks)
- (b) 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. (5 marks)
- (c) Discuss pros and cons for centralized and distributed mutual exclusion algorithms. (3 marks)




Diagram 1

**Question 5 (17 marks)**

Address the following issues related to distributed computing platforms:

- (a) Explain the role of Interface Definition Languages (IDL) in distributed computing. (3 marks)
- (b) 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 (5 marks)
- (c) What kind of distributed applications are more suitable for Web Services than for .NET Remoting (or the RMI service in WCF)? Explain why. (4 marks)
- (d) Explain why it is difficult to make various distributed computing platforms interoperable (i.e. why it would be difficult to develop a distributed applications where application components (objects) would be placed on various distributed computing middleware). (5 marks)

**Question 6 (8 marks)**

The following questions relate to fault tolerance:

- (a) Give two examples of *forward recovery* and two examples of *backward recovery*. (4 marks)
- (b) There is a group of five 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)

**Question 7 (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 15ms 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 60 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 are always enough threads. (2 marks)

**Question 8 (13 marks)**

The following questions relate to consistency and replication:

- (a) A consistency model uses a small number of locks. Each lock protects updates to a different variable.
  - a. What kind of consistency model is this? Explain. (3 marks)
  - b. Describe pros and cons for protecting each variable separately. (3 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  $N = 10$ . Explain each answer.

$N_R = 2, N_w = 8$  (1 mark)

$N_R = 1, N_w = 10$  (1 mark)

$N_R = 4, N_w = 7$  (1 mark)

$N_R = 5, N_w = 7$  (1 mark)

**Question 9 (13 marks)**

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

- (a) Show the differences in the Coda design, compared to NFS, which stem from the fact that Coda was built for availability. (4 marks)

- (b) Which feature of the naming approach used in Coda makes the naming more suitable to distributed systems than the naming used in NFS? Explain why. (4 marks)
- (c) Two commonly used communication paradigms in context-aware applications are RMI and asynchronous notifications. Explain why. Support your explanation with application examples. (5 marks)