

EXAMINATIONS — 2010
MID-YEAR

NWEN 401
Distributed Systems Design

Time Allowed: 3 Hours

Instructions: You must attempt ALL questions.

Answers for questions 1-2 must be written in the first booklet.

Answers for questions 3-4 must be written in the second booklet.

Answers for questions 5-6 must be written in the third booklet.

Each question is worth 25 marks.

There are a maximum of 150 marks in total.

Within each question, the marks for subparts are shown.

Non-programmable calculators without full alphabetic keys are permitted.

Paper foreign language to English dictionaries are permitted.

Question	Topic	Marks
1	Design Issues	25 marks
2	Communication and Time	25 marks
3	Distributed Algorithms	25 marks
4	Fault Tolerance	25 marks
5	P2P Middleware	25 marks
6	Content Distribution Networks	25 marks

Question 1. Design Issues

[25 marks]

- (a) [6 marks] Identify at least THREE (3) criteria for defining “scalability”.
- (b) [9 marks] Identify THREE (3) techniques for building scalable systems and explain each technique in a few sentences.
- (c) [10 marks] List at least TWO (2) scalability problems that might arise in pervasive systems, identify the underlying causes of the problems and briefly explain how scaling techniques might be applied to solve these problems.

Question 2. Communication and Time

[25 marks]

- (a) Communication.
 - (i) [5 marks] With the aid of a simple diagram, explain how RPC is implemented. Be sure to include information in your answer how the protocol operates and how stubs and skeletons are generated.
 - (ii) [5 marks] Explain how communication transparency promotes load distribution.
- (b) Logical time.
 - (i) [5 marks] With the aid of a simple diagram, illustrate how Lamports clocks can be used to impose a “happened-before” event ordering in a distributed system.
 - (ii) [5 marks] Again, using a simple diagram, illustrate how Vector clocks can be used to impose a “causal” event ordering in a distributed system.
 - (iii) [5 marks] Compare and contrast the “happened-before” ordering with the “causal” ordering.

Question 3. Distributed Algorithms

[25 marks]

(a) Elections.

(i) [5 marks] Outline the Bully algorithm.

(ii) [5 marks] Outline the Ring algorithm.

(iii) [5 marks] Compare and contrast the relative communication overheads and the fault tolerance provided by each algorithm.

(b) Consider the delivery of multicast messages within a closed group of processes that may fail by crashing.

(i) [5 marks] Outline an algorithm for implementing FIFO ordering of multicast messages within the group and explain why it works.

(ii) [5 marks] All-or-nothing delivery semantics are desirable for implementing reliable distributed systems based upon group communications. Outline an algorithm for reliable multicast with these semantics and explain why it works.

Question 4. Fault Tolerance

[25 marks]

(a) Consider the core dependability/fault tolerance concepts.

(i) [5 marks] A navigation satellite stops communicating with its base station due to a cosmic ray causing a bit flip in its electronics. Identify the “dependable service”, “fault”, “failure”, “vulnerability” and “error” in this scenario. Justify your answer in terms of the definitions of these concepts.

(ii) [5 marks] Identify at least TWO (2) different approaches to “error processing” in fault tolerance and discuss their respective limitations.

(b) Triple modular redundancy is a common technique for implementing hardware fault tolerance.

(i) [5 marks] Define triple modular redundancy and use a simple diagram to provide an example of its application.

(ii) [5 marks] Identify and explain at least TWO (2) arguments against applying triple modular redundancy to software systems.

(c) [5 marks] Consider distributed transactions. Under what circumstances could a worker in a transaction be forced to block in a ready state indefinitely? Identify how the design of the three-phase protocol resolves this problem.

Question 5. P2P Middleware

[25 marks]

(a) [6 marks] List THREE (3) ways by which a distributed middleware hides the underlying networked environment's complexity.

(b) P2P Messaging System

(i) [6 marks] With the aid of a simple diagram, explain how groups are organized in the P2P Messaging System (P2PMS).

(ii) [4 marks] What are the advantages of the proposed approach?

(c) The peer-to-peer communication in P2PMS is based on a multi-ring topology.

(i) [4 marks] What are the advantages of the ring topology?

(ii) [5 marks] What does the system parameter q define? If $q = 10$, and the total number of nodes (i.e. outer most ring) is 600 nodes, how many inner rings does the P2P network have and how many nodes are in the innermost ring? Explain your reasoning.

Question 6. Content Distribution Networks

[25 marks]

(a) [4 marks] List each of the key components of a CDN and write a sentence to describe each component.

(b) Surrogate server placement algorithms.

(i) [2 marks] What are the different classes of surrogate server placement algorithms?

(ii) [10 marks] Choose TWO (2) classes and briefly discuss the key features of each class.

(c) [9 marks] Clearly outline the steps in Request Routing using DNS redirection.
