
UNIVERSITI SAINS MALAYSIA

First Semester Examination
Academic Session 2004/2005

October 2004

CCS521 – Advanced Distributed Systems Concepts and Design

Duration : 2 hours

INSTRUCTION TO CANDIDATE:

- Please ensure that this examination paper contains **FOUR** questions in **FOUR** printed pages before you start the examination.
 - Answer **ALL** questions.
 - You can choose to answer either in Bahasa Malaysia or English.
-

ENGLISH VERSION OF THE QUESTION PAPER

1. (a) Distributed algorithms can be classified into two; one requiring the use of a coordinating process while the other has a non-fixed coordinator.
 - (i) State an example algorithm for each class. (2/25)
 - (ii) For each, discuss the main advantage and disadvantage of the algorithm. (4/25)
 - (b) Why is it sometimes difficult to hide the occurrence and recovery from failures in a distributed system? (4/25)
 - (c) (i) State **two (2)** necessary requirements of a name service. (2/25)
 - (ii) What is the main problem with multicast navigation in name services? (2/25)
 - (iii) Suggest **two (2)** ways to overcome the problem in 1(c)(ii). (6/25)
 - (iv) Propose a way to overcome the problem of stale information provided by local name resolvers that use caching. (5/25)
2. (a) Consider a critical region which is located inside a loop. Assume that within the critical region, there are shared data which are replicated at various nodes.
 - (i) Differentiate between eager release consistency model and lazy release consistency model using the above scenario as illustration. (6/25)
 - (ii) Evaluate both models in 2(a)(i) in terms of bandwidth utilization and number of messages generated. (9/25)

- (b) Consider the following execution of two processes. Assume that initially, all variables are set to zero.

P1: R(x)1; R(x)2; W(y)1

P2: W(x)1; R(y)1; W(x)2

- (i) Is the memory underlying the above execution sequentially consistent? Explain your answer.

(5/25)

- (ii) Is the memory underlying the above execution coherent? Explain your answer.

(5/25)

3. The *Election* interface provides two remote methods:

vote: with two parameters through which the client supplies the name of a candidate (a string) and the voter's number (an integer used to ensure each user votes once only). The voter's numbers are allocated sparsely from the range of integers to make them hard to guess.

result: with two parameters through which the server supplies the client with the name of a candidate and the number of votes for that candidate.

- (a) Which of the parameters of these two procedures are input parameters and which are output parameters?

(4/25)

- (b) Define the interface to the *Election* service in CORBA IDL and JAVA RMI. Compare the methods in the two languages for specifying *input* and *output* arguments.

(10/25)

- (c) The *Election* service must ensure that a vote is recorded whenever any users think they have cast a vote. Discuss the effect of 'maybe' call semantics on the *Election* service. Would 'at-least-once' call semantics be acceptable for the *Election* service or would you recommend 'at-most-once' call semantics?

(5/25)

- (d) Outline an implementation for the *Election* service that ensures that its records remain consistent when service is accessed concurrently by multiple clients.

(6/25)

4. (a) A client makes remote procedure calls to server. The client takes 10 milliseconds to compute the arguments for each request, and the server takes 20 milliseconds to process each request. The local operating system processing time for each send or receive operation is 1 milliseconds, and the network time to transmit each request or reply message is 5 milliseconds. Marshalling or unmarshalling takes 1 milliseconds per message.

Calculate the time taken by the client to generate and return from two requests:

- (i) If it is single-threaded, and (5/25)
- (ii) If it has two threads that can make requests concurrently on a single processor. (5/25)
- (b) Compare the static scheduling with the dynamic scheduling. Give one example for each of the scheduling. (5/25)
- (c) Highlight some existing research problems on knowledge grid. (5/25)
- (d) Compare the searching process in Napster and GNUTELLA. (5/25)