

---

UNIVERSITI SAINS MALAYSIA

First Semester Examination  
Academic Session 2002/2003

September 2002

**CCS521 – Advanced Distributed Systems Concepts and Design**

Duration : 3 hours

---

**INSTRUCTION TO CANDIDATE:**

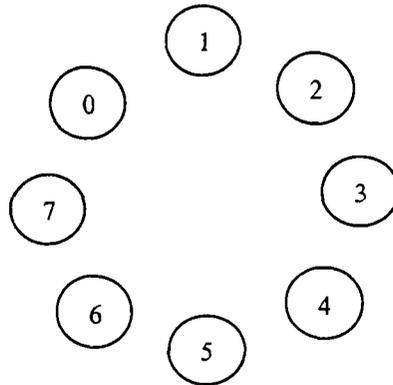
- Please ensure that this examination paper contains **FOUR** questions in **FOUR** printed pages before you start the examination.
  - Answer **ALL** questions.
  - You can choose to answer either in Bahasa Malaysia or English.
- 

ENGLISH VERSION OF THE QUESTION PAPER

1. (a) Differentiate between a distributed operating system and a network operating system. What is common between them, if any?

(5/25)

(b)



Consider the ring-based election algorithm. Two processes, 2 and 5, discover simultaneously that the previous coordinator, process 7, has crashed. Hence both processes 2 and 5 start an election simultaneously.

- (i) What is the number of messages circulating when both processes started the election? (Show your working.)
- (ii) Suggest an improvement to the algorithm to reduce the number of messages circulating.

(14/25)

- (c) Consider the central server algorithm for mutual exclusion. Suppose that there are no processes in the critical section and no processes queued to enter critical section. Assume a synchronous system and that every request is answered (with either a permission or denial).

- (i) If the coordinator crashes under the above circumstances, will anything fatal happen to the system?
- (ii) A process requests to enter critical section (under the above circumstances) but discovers that the coordinator has crashed. How does the process detect the coordinator failure? What can the process do to recover?
- (iii) The above algorithm maintains a first-come-first-serve queue of requests. Suggest another possible algorithm for handling the requests.

(6/25)

2. (a) Why is it useful to define the interfaces of an object in an Interface Definition Language?  
(5/25)
- (b) Write an IDL example for a calculator services that support addition and multiplication.  
(5/25)
- (c) A client makes remote procedure calls to server. The client takes 5 milliseconds to compute the arguments for each request, and the server takes 10 milliseconds to process each request. The local operating system processing time for each send or receive operation is 0.5 milliseconds, and the network time to transmit each request or reply message is 3 milliseconds. Marshalling or unmarshalling takes 0.5 milliseconds per message.
- Calculate the time taken by the client to generate and return from two requests:
- (i) if it is single-threaded, and
- (ii) if it has two threads that can make requests concurrently on a single processor.  
(7/25)
- (d) Draw a Directed Acyclic Graph in which there are four tasks (T1, T2, T3, T4), each with unit execution time 1,2,3,4 respectively. Task T2 and T3 must wait for the completion of T1. Task T4 must wait for the completion of T2 and T3. What is the optimal schedule time if the unit communication cost between processor P1 and P2 is 1, within the same processor is 0, and the unit message is 2 for each interprocess communication?  
(8/25)
3. (a) What are the differences between client-server computing and peer-to-peer computing?  
(6/25)
- (b) Compare the searching process in Napster and GNUTELLA.  
(6/25)
- (c) Explain **three (3)** of the protocols in JXTA.  
(6/25)
- (d) What kinds of applications require grid computing? How can grid computing support them?

4. (a) One important characteristic of a distributed system is transparency in terms of access, location, migration, concurrency and parallelism. As much as transparency is highly favoured, do you think it is a good idea to implement the highest degree of transparency possible? Give an explanation to support your answer.

(5/25)

- (b) Name resolution can be implemented by either using the iterative name resolution or the recursive name resolution.

(i) Briefly explain both the iterative and the recursive name resolution methods.

(ii) State **one (1)** advantage and **one (1)** disadvantage of both methods.

(iii) Assume that the local name resolver on the client side caches results of the name resolution. However, the result of this caching sometimes provide stale information to the client. Suggest a solution to this problem.

(10/25)

- (c) (i) One approach to implement consistency model is to introduce an explicit synchronization variable. Methods that employ this approach includes weak consistency, release consistency and entry consistency. Differentiate the three methods in terms of the way synchronization works.

(ii) Choose one of the methods in 4(c)(i) and explain in general the way it works. Use a diagram to aid your explanation.

(10/25)