

GUJARAT TECHNOLOGICAL UNIVERSITY
ME - SEMESTER– II(Old course) • EXAMINATION (Remedial) – WINTER- 2015

Subject Code: 1720201**Date: 09/12/2015****Subject Name: Distributed Operating Systems****Time: 2:30 pm to 5:00 pm****Total Marks: 70****Instructions:**

1. Attempt all questions.
2. Make suitable assumptions wherever necessary.
3. Figures to the right indicate full marks.

- Q.1**
- (a) Why are distributed operating systems more difficult to design than operating systems for centralized time-sharing systems? **04**
- (b) Concurrency transparency is a desirable goal for distributed systems. Do centralized systems have this property automatically? **03**
- (c) Explain the concept of parallelism transparency. **03**
- (d) Suppose that two variables, A and B are both located, by accident, on the same page of a page-based DSM system. However, both of them are unshared variables. Is false sharing possible? Justify. **02**
- (f) Reading the first block of a file from a remote file server is an idempotent operation. What about writing the first block? **02**

- Q.2**
- (a) What are the main advantages of an RPC system that allows the binding between a client and a server to change dynamically? What are the main issues involved in providing this flexibility? **07**
- (b) What is the difference between a connection-oriented and connectionless communication protocol? Indicate which of the two protocols is preferable for the transmission of the following types of information: (i) voice (ii) video **07**

OR

- (b) Give three different methods that may be used in ATM networks to allocate bandwidth to applications to make best use of the available bandwidth while satisfying the requests of the applications. Also give the relative advantages and limitations of the three methods. **07**
- Q.3**
- (a) Write the steps to explain a producer-consumer pair of processes for the following two cases: **07**
- (i) they use a single message buffer:
- (ii) they use a buffer that can accommodate up to n messages.
- The producer produces messages and puts them in the message buffer, and the consumer consumes messages from the message buffer. Assume that all messages are of fixed size.
- (b) The operations performed by a server are non idem potent. Describe a mechanism for implementing exactly-once IPC semantics in this case. **07**

OR

- Q.3**
- (a) Discuss for each of the following applications, at least once semantics or at most once semantics is best? **07**
- i) Reading and writing files from a file server
- ii) Compiling a program
- Describe a mechanism for implementing consistent ordering of messages for one-to-many, many-to-one and many-to-many communication.

- (b) When buffered communication is used, a primitive is normally available for user processes to create mailboxes. But whether this primitive must specify the size of the mailbox? **07**
- Q.4** (a) In a distributed system, parallelism improves performance and blocking system calls make programming easier. Explain how the concept of threads can be used to combine both advantages. **07**
- (b) What are the main causes of thrashing in a DSM system? What are the commonly used methods to solve the thrashing problem in a DSM system? **04**
- (c) Consider the behavior of two machines in a distributed system. Both have clocks that are supposed to tick 1000 times per millisecond. One of them actually does, but the other ticks only 990 times per millisecond. If UTC updates come in once a minute. What is the maximum clock skew that will occur? **03**

OR

- Q.4** (a) What are the main differences between the load-balancing and load-sharing approaches for process scheduling in distributed systems? Which of the various policies to be used in the implementation of the two approaches are different and which of them are same? **07**
- (b) Discuss the relative advantages and disadvantages of the various data-locating mechanisms that may be used in a DSM system that uses the NRMB and NRRMB strategy. **07**
- Q.5** (a) A system has three types of resources, R1, R2, R3 and their numbers of units are 3, 2, 2 respectively. Four processes P1, P2, P3, P4 are concurrently competing for these resources in the following manner: **04**
- i) P1 is holding one unit of R1 and is requesting for one unit of R2.
 - ii) P2 is holding two units of R2 and is requesting for one unit each of R1 and R3.
 - iii) P3 is holding one unit of R1 and is requesting for one unit of R2.
 - iv) P4 is holding two units of R3 and is requesting for one unit of R1.
- Determine which of the processes are deadlocked in this system state?
- (b) Why do some distributed systems use two-level naming? **03**
- (c) Discuss process management for Amoeba, V-system, Mach and chorus? **07**

OR

- Q.5** (a) Why do stateless servers have to include a file offset in each request? Is this also needed for stateful servers? **04**
- (b) When file systems replicate files, they do not normally replicate all files. Give an example of a kind of file that is not worth replicating. **03**
- (c) How inter process communication is done in Amoeba, V-system, Mach and chorus? **07**
