Hall Ticket Number:

Code No. : 41111

VASAVI COLLEGE OF ENGINEERING *(Autonomous)*, HYDERABAD

B.E. (CSE) IV Year I-Semester Main Examinations, December-2017

Distributed Systems

Time: 3 hours Max. Marks: 70

Note: Answer ALL questions in Part-A and any FIVE from Part-B

Part-A (10 × 2 = 20 Marks)

1. List the various transparency issues to be addressed during designing a distributed application?
2. What is the main disadvantage of distributed systems which exploit the infrastructure offered by Internet? How can this be overcome?
3. Consider the below space time graph of 4 processes (p1, p2, p3, p4) involved in a distributed transaction

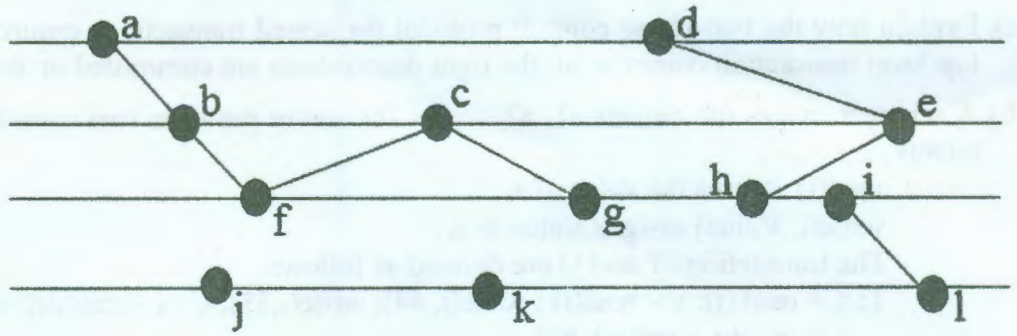

Compute the vector clocks of 'c', 'h', 'e' and 'l'. List the events which are concurrent.

4. Define the types of communication paradigm used between objects.
5. Even without a deadlock, a poor algorithm might lead to starvation. Give an example of a system which leads to starvation.
6. Explain the concepts of International Atomic Time and Coordinated Universal Time.
7. What is a lock manager? Explain how the granting of locks is implemented.
8. If up to f servers can exhibit Byzantine failures, how many servers can provide a correct service?
9. What is a file group? What is the purpose of a file group identifier?
10. Which services must be included to enable clients to bind to services in a given host by name?

Part-B (5 × 10 = 50 Marks)

11. a) What kind of scheduling avoids race conditions? Which factors decide whether a thread should be scheduled as non-preemptive? [4]

    b) A 'null' RMI that takes no parameters calls an empty procedure and returns no values and delays the caller for 2.0 milliseconds. Explain what contributes to this time. In the same RMI system, each 1K of user data adds an extra 1.5 milliseconds. A client wishes to fetch 32K of data form a file server. Should it use one 32K RMI or 32 1K RMIs? [6]

Contd... 2

12. a) Explain the various ways of navigation adopted in DNS name resolution. [6]

b) A client makes remote method invocations to a server. The client takes 5ms to compute the arguments for each request, and the server takes 10ms to process each request. The local operating system processing time for each send or receive operation is 0.5ms and the network time to transmit each request or reply message is 3ms. Marshalling or unmarshalling takes 0.5ms per message. Calculate the time taken by the client to generate and return from two requests ignoring context-switching times for: [4]

(i) if it is single threaded and

(ii) if it has two threads that can make requests concurrently on a single processor.

Is there a need for asynchronous invocation if the client and server processers are threaded? Why or why not?

13. a) Devise an algorithm to implement unreliable failure detector. [5]

b) How is the fault tolerance of different mutual exclusion algorithms evaluated? Name the type of failure that cannot be tolerated by the ring-based algorithm. [5]

14. a) Explain how the two-phase commit protocol for nested transactions ensure that if the top-level transaction commits, all the right descendants are committed or aborted. [4]

b) A server manages the objects a1, a2, ...an. The server provides two operations for its clients: [6]

read(i) returns the value of a_i;
write(i, Value) assigns Value to a_i;
The transactions T and U are defined as follows:
T: x = read (j); y = read(i) ; write(j, 44); write(i,33); U: x = read(k); write(I, 55);
y = -read(j_; write(k,66)
Give three serially equivalent interleavings of the transactions T and U.

15. a) Give the uses of the following NFS server operations: [4]
(i) read(fh,offset,count) *(ii)* write(fh,offset,count,data) *(iii)* lookup(dirfh,name)

b) Describe the overall architecture of Google Infrastructure. [6]

16. a) Explain thread-per-request and thread-per-connection architecture with respect to multi-threaded servers. Does switching present a problem for threads implementation? [5]

b) Explain the various types of Client Server communications. [5]

17. Answer any ***two*** of the following:

a) Explain whether the algorithm for reliable multicast over IP multicast works for open as well as closed groups. Given any algorithm for closed groups, how, simply, can we derive an algorithm for open groups? [5]

b) A clock is reading 11:34:26.0 (hr:min:sec) when it is discovered to be 6 seconds fast. Explain why it is undesirable to set it back to the right time at that point, and show (numerically) how it should be adjusted so as to be correct after 12 seconds have elapsed. [5]

c) Both GFS and Bigtable make the same core design choice – to have a single master. What are the repercussions of a failure of this single master in each case? [5]