

98-Comp-B10, Distributed Systems

3 hours duration

NOTES:

1. If doubt exists as to the interpretation of any question, the candidate is urged to submit with the answer paper a clear statement of any assumptions made.
2. This is a CLOSED BOOK examination.
Only non-programmable calculator is permitted
3. Answer any **five of the six** questions.
Only the first five questions as they appear in the answer book will be marked.
4. All questions carry equal weight.
5. Most questions require an answer in essay format. Clarity and organization of the answer are important.

Question # 1. Characteristics of distributed systems [20 Marks]

- a. Discuss the difference and similarities between distributed systems and centralized systems. [5 Marks]
- b. Name, Explain and illustrate (in graphical form) one distributed system architecture. You may use a major Internet applications as an example. [5 Marks]
- c. Discuss the resource sharing concept in Distributed Systems. Name 2 advantages and 2 disadvantages/challenges. [5 Marks]
- d. List the factors that contribute to the time delay taken to send a message between two processes over a network channel. Discuss the measures that can be taken to set a bound on the each factor's contribution to the total delay time. [5 Marks]

Question # 2. Fundamental concepts and mechanisms. [20 Marks]

- a. A client Web browser sends two 750-byte request messages to a web server (assume that the messages are sent back to back at the client side; i.e., the client sends directly the second request message after sending the first message), which produces a single response message with a size of 6000 bytes. Estimate the total time for the client to receive the response message from the server in each of the following cases, with the performance assumptions listed below:
 - i) Using connectionless (datagram) communication (for example, UDP); [4 Marks]
 - ii) Using connection-oriented communication (for example, TCP); [4 Marks]
 - iii) The server process is in the same machine as the client. [2 Marks]

*[Latency per packet (local or remote, incurred on both send and receive): 5 milliseconds
Connection setup time (TCP only): 5 milliseconds
Data transfer rate: 20 megabits per second; MTU: 2000 bytes
Server request processing time: 2 milliseconds; Assume that the network is lightly loaded.]*

- b. A sequence of packets are transmitted back to back for a source to a destination through a wide area network.
 - i) Is it possible for these packets to be received at the destination in an order that differs from that in which they were transmitted? Explain. [6 Marks]
 - ii) If the source and destination are in the same local area network, can these packets arrive to the destination out of order? [4 Marks]

Question # 3. Client-server systems & inter-process communications

- a. The concept of a remote procedure call (RPC) represents a major intellectual breakthrough in distributed computing.
 - i) Explain briefly the RPC concept and why it represents a major intellectual breakthrough in distributed computing. [3 Marks]
 - ii) Discuss the RPC design issues? Namely, explain the following: 1) programming with interfaces; 2) the call semantics associated with RPC; and 3) the key issue of transparency and how it relates to remote procedure calls. [3 Marks]

- b. A request-reply protocol is a form of communication that is designed to support message exchanges in typical client-server interactions. The request-reply protocol is based on a trio of communication primitives, *doOperation*, *getRequest* and *sendReply*.
- Explain these three communication primitives. [4 Marks]
 - The request-reply protocol masks the heterogeneity of operating systems and of computer networks. How is this accomplished? [3 Marks]
- c. A client makes remote method invocations to a server. The client requires 4 milliseconds to generate each request, and the server takes 10 milliseconds to serve each request. The local operating system processing time for each send or receive operation is 0.6 milliseconds, and the network time to transmit each request or response message is 3 milliseconds.
- Calculate the time taken by the client to generate and return from two requests:
- if it is single-threaded; [4 Marks]
 - if it has two threads that can make requests concurrently on a single processor. [3 Marks]

Question # 4. Security

- Encryption is the process of encoding a message in such a way as to hide its contents. What are then two main classes of encryption algorithms that are in general use? Provide a brief description of each class of algorithms. [6 Marks]
- Initial exchanges of public keys between two end communicating processes may be vulnerable to man-in-the-middle attacks. Discuss the potential defences that can be used to overcome this vulnerability. [6 Marks]
- Estimate the time required to crack a 56-bit DES (Data Encryption Standard) key by a brute-force attack using a 2000 MIPS (million instructions per second) computer, assuming that the inner loop for a brute-force attack program involves around 10 instructions per key value, plus the time to encrypt an 8-byte plaintext. Perform the same calculation for a 128-bit IDEA key. [8 Marks]

Question # 5.

- Name and discuss three requirements in the design of distributed file systems? [5 Marks]
- Which file system (i.e., NFS or AFS) can support more clients, given a server that runs on identical hardware? Explain. [5 Marks]
- Traditional NFS has a stateless server, allowing it to reboot without impacting existing client connections (much). What does NFS trade-off to gain this stateless server advantage? i.e., where does it lose in return? [5 Marks]
- Compare AFS and NFS from scalability point view? Are there any limits on AFS scalability, assuming that servers can be added as required? [5 Marks]

Question # 6.

- a. Discuss why replication is a key to the effectiveness of distributed systems. *[5 Marks]*
- b. What is the difference between active and passive replications? *[5 Marks]*
- c. What is the gossip architecture? Why does a replica manager need to keep both a 'replica' timestamp and a 'value' timestamp? *[5 Marks]*
- d. Does making some replica managers read-only improve the performance of a gossip system? Explain. *[5 Marks]*