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 seven** 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*

a. Why distributed system designs are being used? List three advantages of distributed systems.
b. Explain what is meant by a client program and a server program. Explain and illustrate (in graphical form) the client-server architecture of one major Internet applications (for example the Web, email or ftp).
a. Give a definition of middleware. Show in a small diagram its position in the architecture of a distributed system.

Question # 2. *Fundamental concepts and mechanisms*

a. Discuss the use of connectionless (UDP) and connection-oriented (TCP) communication for of each of the following application/application protocol:
    1. File transfer protocol;
    2. HTTP protocol;
    3. DNS protocol;
    4. Youtube.

b. A user application sends a 400-byte request message to a service, which serves the request and produces a response containing 6000 bytes. Estimate the total time to complete the request in each of the following cases, with the performance assumptions listed below:
    i) The user application uses the User Datagram Protocol (as a transport layer protocol);
    ii) The user application is built on top of the HTTP protocol;
    iii) The server process is in the same machine as the client.

    [Latency per packet (local or remote, incurred on both send and receive): 10 milliseconds
    Connection setup time (TCP only):6 milliseconds
    Data transfer rate: 10 megabits per second
    MTU: 1500 bytes
    Server request processing time: 2 milliseconds
    Assume that the network is lightly loaded.]

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

a. What are the differences between a local call and a remote call?
b. Explain "call by reference" vs. "call by value".
c. List three ways to invoke a method on a remote object.
d. What is the purpose of an Interface Definition Language?
e. Give explanation to the design choices that are relevant to decreasing the amount of reply data held at a server. Suppose the RRA protocol is in use. What is the length of time needed for servers to retain unacknowledged reply data? Do the servers have to continually send the reply in an attempt to obtain an acknowledgement?

Question # 4. *Operating systems for distributed architectures*

a. A client makes RMIs to a server. The client takes 5 ms to compute the arguments for each request, and the server takes 10 ms to process each request. The local OS processing time for each *send* or *receive* operation is 0.5 ms, and the network time to transmit each request or reply message is 3 ms. Marshalling or unmarshalling takes 0.5 ms per message.
Estimate 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. Is there a need for asynchronous RMI if processes are multi-threaded?

b. Network transmission time accounts for 20% of a null RPC and 80% of an RPC that transmits 800 user bytes (less than the size of a network packet). By what percentage will the times for these two operations improve if the network is upgraded from 5 megabits/second to 100 megabits/second?

Question # 5. *Security*

a. Explain the steps that a pair of users using PGP must take prior to exchanging email messages with privacy and authenticity guarantees. What scope is there to make the preliminary key negotiation invisible to users? (The PGP negotiation is an instance of the hybrid scheme.)
b. Estimate the time required to crack a 56-bit DES 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. Extrapolate your calculations to obtain the cracking time for a 200,000 MIPS parallel processor (or an Internet consortium with similar processing power).

Question # 6. *Distributed file systems*

a. How does AFS gain control when an open or close system call referring to a file in the shared file space is issued by a client?
b. What data must the NFS client module hold on behalf of each user-level process?
c. How does AFS deal with the risk that callback messages may be lost?
d. Compare AFS and NFS from stability point view? Are there any limits on AFS scalability, assuming that servers can be added as required? Are there any recent technological developments that would help to offer greater scalability?

Question # 7. *Principles of fault tolerance continual*

a. Discuss the functions of the implementation repository from the point of view of scalability and fault tolerance.
b. Explain why making some replica managers read-only may improve the performance of a gossip system.
a. Write pseudocode for a scheme for integrating two replicas of a file system directory that underwent separate updates during disconnected operation.