C12-R3: DISTRIBUTED SYSTEM

NOTE:

1. **Answer question 1 and any FOUR questions from 2 to 7.**
2. **Parts of the same question should be answered together and in the same sequence.**

Time: 3 Hours **Total Marks: 100**

1.

a) Give an example of a URL. List the three main components of a URL, stating how their boundaries are denoted and illustrating each one from your example. To what extent is a URL location transparent?

b) How does a newly installed personal computer connected to an Ethernet discover the IP addresses of local servers? How does it translate them to Ethernet addresses?

c) Java and other languages support exceptions, which are raised when an error occurs. How would you implement exceptions in RPCs and RMIs?

d) Explain what the kernel must provide for a user-level implementation of threads, such as in Java on Unix. Do page faults present a problem for user-level threads implementations? Justify.

e) What is a digital signature? What are its uses in the security of a distributed system? Give a method to create a digital signature.

f) Describe how a non-recoverable situation could arise if write locks are released after the last operation of a transaction but before its commitment.

g) Discuss whether message passing or DSM (Distributed share memory) is preferable for fault-tolerant applications.

(7x4)

2.

a) How do the client-server architecture of one or more major Internet applications involve the partitioning and/or replication of data amongst servers? Discuss with the help of an example.

b) Assume the RRA (request-reply-acknowledge) protocol is in use. How long should servers retain unacknowledged reply data? Should servers repeatedly send the reply in an attempt to receive an acknowledgement?

c) Outline the design of a scheme that uses message retransmissions with IP multicast to overcome the problem of dropped messages. Your scheme should take the following points into account:
   i) there may be multiple senders;
   ii) generally only a small proportion of messages are dropped;
   iii) unlike the request-reply protocol, recipients may not necessarily send a message within any particular time limit.

   Assume that messages that are not dropped arrive in sender ordering.

(6+6+6)

3.

a) Implement a simple client-server system using RPC. The server offers one procedure, *next*, which takes an integer as input and returns its successor as output. Write a stub procedure called *next* for use on the client side. Its job is to send the parameter to the server using UDP and wait for the response, timing out if the response takes too long. The server

procedure should listen on a known port, accept requests, carry them out, and send back the results.

b) Why do DNS root servers hold entries for two-level names such as *yahoo.com* and *engineering.edu*, rather than one-level names such as *edu* and *com*?

c) GNS does not guarantee that all copies of entries in the naming database are up-to-date. How are clients of GNS likely to become aware that they have been given an out-of-date entry? Under what circumstances might it be harmful?

(6+6+6)

4.

a) Which features of the AFS design make it more scalable than NFS? What are the limits on its scalability, assuming that servers can be added as required? Which recent developments offer greater scalability?

b) Draw and explain the typical infrastructure components for a simple multimedia conferencing application running on two personal computers.

c) In the following table some typical types of multimedia applications are given. Fill up the approximate data rates and frame or sample size and frequencies.

	Data Rate (approximate)	**Sample or Frame**	
		Size	**Frequency**
i) Telephone speech ii) CD-quality sound iii) Standard TV video (uncompressed) iv) Standard TV video (MPEG-1 compressed) v) HDTV video (uncompressed) vi) HDTV video (MPEG-2 compressed)			

(8+6+4)

5.

a) Estimate the time required to crack a 56-bit DES key by a brute-force attack using a 500 MIPS (million instruction per second) workstation, 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 128bit IDEA key. Extrapolate your calculations to obtain the cracking time for a 50,000 MIPS parallel processor (or an Internet consortium with similar processing power).

b) How would email be sent to a large list of recipients using PGP (Pretty Good Privacy) or a similar scheme? Suggest a scheme that is simpler and faster when the list is used frequently.

(10+8)

6.

a) Explain why serial equivalence requires that once a transaction has released a lock on an object, it is not allowed to obtain any more locks. A server manages the objects a_1, a_2, … a_n. The server provides two operations for its clients:

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 (i); write (j, 44);*
U: *write (i, 55); write (j, 66);*

Describe an interleaving of the transactions *T* and *U* in which locks are released early with

the effect that the interleaving is not serially equivalent.

b) The role of server and its status at the time of failure are given below. Write and discuss the action of the recovery manager with respect to the two-phase commit protocol for any transaction.

	Server Role	Status	Action of recovery manager
i)	Coordinator	prepared	
ii)	Coordinator	committed	
iii)	Participant	committed	
iv)	Participant	uncertain	
v)	Participant	prepared	
vi)	Coordinator	done	

(9+9)

7.

a) A file server uses caching, and achieves a hit rate of 80%. File operations in the server cost 5ms of CPU time when the server finds the requested block in the cache, and take an additional 15ms of disk I/O time otherwise. Explaining any assumptions you make, estimate the server's throughput capacity (average requests/sec) if it is:
   i) Single-threaded;
   ii) Two-threaded, running on a single processor;
   iii) Two-threaded, running on a two-processor computer.

b) Prove that Ivy's write-invalidate protocol guarantees sequential consistency.

c) In Ivy's dynamic distributed manager algorithm, what steps are taken to minimize the number of lookups necessary to find a page?

(6+6+6)