1a What is meant by an open distributed system? *5pt*

This means that the system offers well-defined interfaces, supports application portability, and can operate with other open distributed systems.

1b What is meant by a three-tiered client-server architecture? *5pt*

This is an architecture that consists of three types of machines: a user machine, an application server, and a database server.

1c Consider the following client-server organizations. Organization (e) used to be popular, but there is a trend towards going back to organization (a). Why? *5pt*

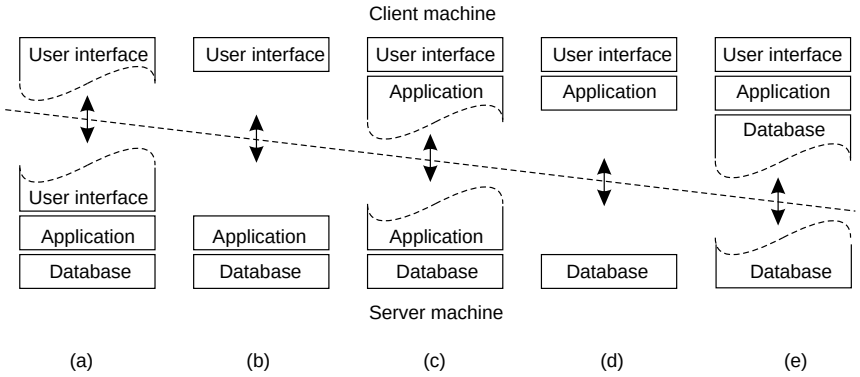

Organization (a) has thin clients, meaning that less things can go wrong at the client side. From a systems management perspective, this makes things a lot easier.

2a Remote procedure calls rely on a call-by-copy/restore parameter passing mechanism. Explain this mechanism. *5pt*

It means that parameter values are copied to the stack associated with the procedure, from where they can be manipulated. When the procedure execution is finished, the value is copied back to its original place, overwriting any value that had been produced concurrently while the procedure was being executed.

2b Consider a client performing an RPC, but the server crashes before it could send the response. What should the client do? *5pt*

It really depends. When the RPC involved an idempotent operation, the client can simply redo the RPC. If not, it depends on whether the crash happened before or after the operation was completed at the server.

3a What is the difference between an iterative and a concurrent server? *5pt*

An iterative server can handle only one request at a time, whereas a concurrent server can have multiple outstanding requests at the same time.

3b Explain how a superserver works. *5pt*

A superserver represents multiple services, normally identified by well-known port numbers. When a service request arrives at one of these ports, the superserver spawns a separate process to handle that request and just continues listening to the ports again.

4a Give a convincing example where using mobile agents is the obvious solution. *5pt*

There may be only one: when lots of remote data needs to be inspected. In that case, moving the code to the data rather than moving the data to the place where the inspection is initiated, will show to perforom best.

4b Give an example of receiver-initiated code migration. *5pt*

The case where a server returns an applet to a client in response to a HTTP request.

5a Explain how Lamport timestamps work. *5pt*

Every process P_i keeps a local counter C_i, which is initially set to 0. When a message m is sent, it receives a timestamp $T(m)$ that is set to C_i, after which C_i is incremented by 1. When process P_j receives a message, it sets its own value C_j to $\max\{C_j+1, T(m)+1\}$.

5b Explain how totally-ordered multicasting can be implemented with Lamport timestamps. *5pt*

See also book, page 255. The essence is as follows. When process P_i wants to update the replicas, it broadcasts the update m_i to itself and all other processes. Message m_i is timestamped P_i's current value of C_i. When a process P_j receives m_i it puts it in a local queue, ordered by $T(m_i)$, and broadcasts an acknowledgement. Only when m_i is at the head of the queue, and for every process there is a message queued with a higher timestamp, will m_i be passed to the application. Lamport assumes that messages are sent FIFO-wise, and that no messages are lost.

5c Explain what is meant by a distributed snapshot. *5pt*

A distributed snapshot is a representation of a consistent global state. Consistency refers to the fact that for each process a local state is recorded such that the recorded arrival of every message is backed up by the recorded sending of that message.

6a What is the essential difference between caching and replication? *5pt*

Caching takes place under the regime of a client's decision; replication is generally initiated by a server.

6b What is meant by active replication? *5pt*

With active replication, a number of processes execute update operations in the same order. They generally operate in lock-step mode. The main advantage is that there is no need to propagate state updates, but only ensure that the replicas execute the same operations in the same order.

7a What is the difference between independent and coordinated checkpointing? *5pt*

With independent checkpointing, each process in a distributed system checkpoints its state independently from other processes. In contrast, with coordinated checkpointing, a master process instructs processes to checkpoint their state; wait until everyone has finished; and instructs them to continue.

7b Explain what a piecewise deterministic execution model is. *5pt*

In such a model, the execution of a process is assumed to take place in intervals that are separated by nondeterministic events. Typically, events are certain calls to the operating system, the exchange of messages, or the occurrence of an interrupt. Between the events, it is assumed that a series of instructions are executed that can always be played back without leading to different results.

8a Does NFS version 4 follow the remote access model, or the upload/download model? Motivate your answer. *5pt*

One could argue that it follows the upload/download model, for the simple reason that clients are provided with a copy of the file they requested for, and which they can modify locally until the server recalls this so-called delegation.

8b Explain what is meant with session semantics in the context of distributed file systems. *5pt*

Session semantics is an issue when dealing with shared files. In practice, if no locking is applied, a client will open a file after which it is allowed to locally modify its copy. The effect of its modifications are made visible to other processes only after the file is closed again. The problem that we need to handle is when two or more processes have concurrently opened a file. In NFS, the last process to close a file will overwrite all the other process's changes.

***Grading:** The final grade is calculated by accumulating the scores per question (maximum: 90 points), and adding 10 bonus points. The maximum total is therefore 100 points.*