

1a Achieving full distribution transparency is impossible. Why? 5pt

A simple answer is that you cannot fully hide failures and their recovery: it is impossible to distinguish a slow computer from a failing one. More subtle is that striving for one type of transparency will generally lead to severe performance problems, thereby revealing that you actually may be working in a networked environment.

1b Explain the principle of a page-based distributed shared memory system. 5pt

Crucial is that you mention the combination of local virtual memory and that page faults result in fetching a copy of a page from another machine. For the rest of the answer: see book.

2a RPC support systems provide various tools, including a runtime library. What does this library contain? 5pt

In principle only two things: (1) a collection of messaging primitives that are to be invoked for sending and receiving messages that are wrapped RPC calls/responses, and (2) a collection of routines for transforming data structures into host-independent/network-independent representations.

2b What is generally the major difference between a traditional RPC, and a remote method invocation (RMI)? 5pt

RMI systems generally offer a global, systemwide object reference, allowing a process to refer to objects on remote machines. In other words, it supports a notion of a global pointer.

3a Consider a stateless server that needs to support operations requiring to keep information about clients (such as which files a client has opened). How can such operations be supported without violating the stateless design of the server? 5pt

One obvious solution is to have the client-side code keep the necessary state, for example in the form of cookies.

3b Stateful servers can often perform better than stateless servers. Explain why this is true by means of an example. 5pt

The best example is that of wide-area client-server interaction: a client asks to access a file; the server returns the entire file to allow local operations by the client. Another example is being able to do read-ahead operations on file blocks, although this is less convincing because the server can do this anyway.

- 4a A directory service is very similar to a Yellow Pages service: entities are looked up by specifying properties instead of exact names (such as in DNS). Why is it so difficult to have an efficient distributed directory service? 5pt

The main problem is that by specifying properties, we, in principle, need to do an exhaustive search through an entire database. This is doable on a single machine, but not when the service is distributed, as every query, in principle, has to be forwarded to every machine, after which the answers have to be collected and combined. With a naming service, the name resolution path is explicit. No, this answer is not in the book.

- 4b Explain what is meant by a closure mechanism in the context of name resolution. 5pt

A closure mechanism describes how and where a name resolution should start. It is generally hard coded into a system and cannot be described as a general mechanism (which would require a closure mechanism by itself).

- 5 What is the relation between a physical clock and a logical clock? 5pt

None

<p>Grading: The final grade is calculated by accumulating the scores per question (maximum: 45 points), and adding 5 bonus points. The maximum total is therefore 50 points.</p>
