1a Give five examples of policies that have been separated from mechanisms in distributed systems. *5pt*

(1) telling what your browser should do when it comes to caching Web pages. (2) deciding on which security keys you use for SSH, (3) setting protection bits in distributed file systems, (4) deciding how often your machine should synchronize with a time server, (5) decding on how and when you want your packages of system software to be automatically checked for updates. Similar examples easily come to mind.

1b Explain what is meant by geographic scalability problems and give an example how such a problem can be solved. *5pt*

The essence is that the behavior or performance of a system is influenced by the fact that components are placed far apart, incurring high communication latencies. There is not really that much you can do, except perhaps by replicating data close to clients, or concealing waiting times by letting the client do something else.

1c Web sites are often organized as a multi-tiered client-server system consisting of a Web server, and application server, and a database server. What are the advantages of such a scheme? *5pt*

The main advantage is that it can be easily managed. In addition, its modularity often makes it easier to replace parts, or to, for example, replicate and distribute the database.

2a What is the major disadvantage of using RPCs in comparison to messaging as in message-queuing systems? *5pt*

The main disadvantage is that due to the synchronous and nonpersistent nature of RPCs, whenever a fault occurs it has to be dealt with immediately. In MQ systems, message delivery can simply be delayed until the problem is solved.

2b Explain what subject-based routing means and how message brokers can be used to implement it. *5pt*

Subject-based routing refers to the fact that messages can be tagged with a subject as address and that only processes who have subscribed to the subject will be able to receive it. What is needed is that senders and subscribers are matched, which is typically something message brokers are good at. All messages are simply sent to the broker who maintains a list of (subject,receiver) pairs and simply forwards messages to the subscribed processes.

2c Consider a worldwide overlay network of many message brokers connected through IBM MQ series. Sketch a straightforward solution for supporting multicasting in such an environment. *5pt*

One possible solution is to organize the brokers into a single multicast backbone tree. Subscriptions for subject S are broadcast to all brokers, where each of them simple registers along which links messages on S should be forwarded when published. This induces a multicast subtree.

3a When using DNS in the following way, the user apparently is getting inconsistent answers. Explain this difference. *5pt*

```
seuss > host flits.cs.vu.nl; host 192.31.231.65
flits.cs.vu.nl has address 192.31.231.65
65.231.31.192.in-addr.arpa domain name pointer flits.few.vu.nl.

seuss > host entry.van-steen.net; host 213.10.169.34
entry.van-steen.net has address 213.10.169.34
34.169.10.213.in-addr.arpa domain name pointer ipd50aa922.speed.planet.nl.
```

It is actually very simple. Although one would expect that entry.van-steen.net is registered as uniquely belonging to 213.10.169.34, this is not the case. The fact is that the address belongs to the planet.nl domain, for which it has been registered under ipd50aa922.speed.planet.nl. However, the domain name entry.van-steen.net is registered with a separate DNS server, where it is resolved to 213.10.169.34.

3b User `maarten` in domain `van-steen.net` has a forwarding mail address to `steen@cs.vu.nl`. Show which messages are exchanged between a user agent sending a message to `maarten@van-steen.net`, DNS, mail servers, and the recipient user agent. *10pt*

1. SND requests MX record for van-steen.net. 2. DNS returns domain name. 3. SND requests IP address of mail server of van-steen.net; DNS returns IP address (4 is often returned along with 2). 5. SND sends mail to mail server MS1 at van-steen.net. 6. MS1 looks up mail server for cs.vu.nl (as before). 7. MS1 sends mail to mail server MS2 at cs.vu.nl. 8. MS2 forwards mail to user's mail box/user agent.

4a Explain how Lamport timestamps work and show by example that they do not necessarily capture (potential) causality. *5pt*

Lamport timestamps are explained in the book. For the example, simply take two concurrent events with different timestamps. It should be obvious that there is no potential causality.

4b Vector timestamps can be used for capturing potential causality. Fill in the vector clocks in the following figure and explain why the delivery of message m* at process P_2 is delayed. *5pt*

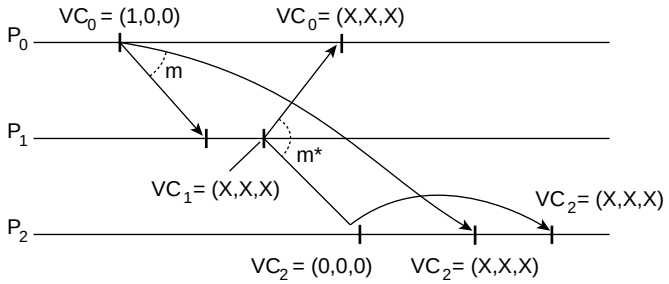

Because $VC_1 = (1,1,0)$ at the time m arrives at P_2, the latter notices it is missing a message that causally preceded m*. It will have to wait until that message (m) arrives.*

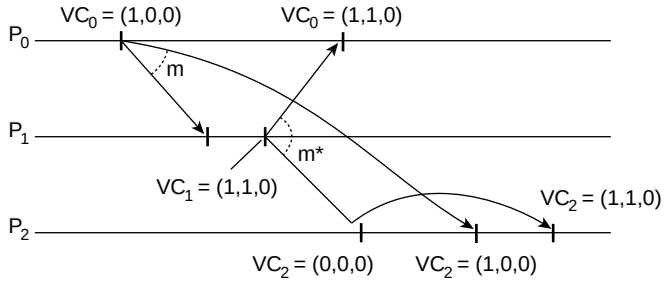

5a Explain what is meant by sequential consistency. Be precise! *5pt*

See book.

5b What is the essence of demanding that the set of synchronization variables adheres to sequential consistency when considering weak consistency models? *5pt*

This demand ensures that all processes see the same lock/release behavior for synchronization variables. As a consquence, everyone will have the same view on whether it is allowed to operate on the data protected by a synchronization variable.

5c Is it possible to have write-write conflicts in client-centric consistency models? Explain your answer. *5pt*

Yes. These models say nothing about concurrent access at all. Therefore, it could very well be possible that while Alice is modifying x at location A, that Bob is modifying x as well, but at another location B.

6a Explain what is meant by a virtual synchronous multicast. *5pt*

In such a multicast, the sender will have a view of who the recipients are. For VSM, it is guaranteed that a message that sent to such a view is delivered to all nonfaulty members of that view, or to none at all.

6b In virtual synchrony we can separate the ordering of message delivery at a single recipient from the fact that messages are subject to total-ordered delivery. What is the difference between the two? *5pt*

As already partly mentioned, the difference is between what happens at a single recipient, and what happens with respect to the whole group. Total-ordered delivery means that all recipients see the same messages in the same order. However, this delivery may not necessarily mean that, for example, potential causality is preserved.

7a What is the difference between two-phase locking and two-phase commit? *5pt*

The two have nothing to do with each other. 2PL refers to the fact that once a transaction releases a lock, it can no longer acquire one. 2PC refers to the way a distributed transaction instructs participants to abort or commit.

7b Two-phase commit (2PC) is said to be a blocking protocol for which reason three-phase commit was devised (3PC). How can this blocking happen? *5pt*

When the coordinator crashes in 2PC, the remaining participants may not be able to come to a final decision on the next step and will have to wait until the coordinator recovers.

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.*