

- 1a What is meant by geographical scalability and how can this problem be tackled? 5pt  
*In a geographically scalable system users and resources may lie so far apart that this fact may become difficult to hide. Solutions are generally to replicate data or attempt to hide communication latencies.*
- 1b One could argue that peer-to-peer systems solve the administrative scalability problem. Explain how. 5pt  
*The key to the answer lies in the fact that in many p2p systems end users are in control. As a result, there simply is not much formal administration to deal with, let alone administrative problems that are tightly coupled to organizational issues.*
- 1c Replication is widely applied as a scaling technique, but introduces a new scalability problem. Explain what this problem is and how it can be solved. 5pt  
*Replication requires that copies stay consistent, in turn introducing a global synchronization problem, which is inherently not scalable. The only solution is to weaken data consistency.*
- 2a Give a typical example of client-side software that is needed in a distributed system. 5pt  
*Any answer related to client-side stubs for implementing RPC, handling replication, or fault tolerance would be correct.*
- 2b Give at least one compelling reason why using mobile agents would be a good idea. 5pt  
*The single most compelling reason is performance: it may be much cheaper to send code to a remote site and do local data analysis, instead of having to ship the data to the requester. Another good reason is related to fault tolerance: by letting an agent represent a user and moving to another site when the home is down, at least some degree of failure transparency can be supported.*
- 3a Is the URL `http://www.cs.vu.nl/` location independent? 5pt  
What about `http://www.van-steen.net/`? Explain your answer!  
*It really depends whether these names are tightly coupled to the addresses they refer to. If the coupling is weak (e.g., because it is easy to change the mapping), they cannot be considered location dependent. Most important is that you cannot see by the name itself whether it is location dependent or not.*
- 3b Name spaces generally provide hard links and symbolic links. Explain the difference. 5pt  
*A hard link is a path name to a node in the name space. A symbolic link is a path name to a leaf node containing a path name; the latter is to be used to continue name resolution.*
- 4a Describe a simple, centralized scheme to accomplish totally-ordered reliable multicasting. 5pt  
*The simplest solution is perhaps to make use of a serializing node: all multicast requests are sent to that node, which attaches a sequence number to each packet. To assure reliability, a receiver is required to acknowledge incoming messages. An alternative is to let senders multicast messages to a group, but receivers are requested to delay delivery until a central sequencer has assigned a number to each multicasted packet.*
- 4b Ricart and Agrawala's algorithm for distributed mutual exclusion among  $N$  processes introduces  $N$  points of failures. How? 5pt  
*The key to this answer is that each process will be capable of accessing the shared resource only if it has heard from every other process. If a single process fails to respond to the multicasted request for access, all other processes will come to a grinding halt.*

- 5a Consider a Web server handing out leases to invalidate client-side caches. When should renewal-frequency based leases be applied in order to obtain maximal effectiveness? Include the lease period in your answer. 5pt
- These type of leases are related to the frequency at which clients request the server to verify the validity of their cache's content. If the client accesses the cache more often than that its content is to be refreshed, the server should handout a lease promising to invalidate the cache. The higher the read/write ratio of a document, the longer the lease period should be.*
- 5b What do leases such as those just described actually accomplish? 5pt
- The ultimate goal is to get client's off the server's back, while at the same time the server establishes strong consistency: clients will always get to see up-to-date information during the lease period.*
- 6a Explain how (scalable) NACK-based reliable multicasting works. 5pt
- In this type of multicasting, a sender attaches a sequence number to each multicasted packet. Instead of letting receivers ACK every packet, receivers remain still until they detect a missing packet. For scalability, receivers listen to a common feedback channel. When a receiver returns a NACK to the sender, any NACK that would have been sent by another receiver is suppressed, effectively establishing that only a single NACK reaches the receiver. The latter is then required to multicast the missing packet again.*
- 6b What is the best way to establish scalable ACK-based reliable multicasting? 5pt
- The only way to avoid feedback implosion is to aggregate ACKs. This can be done by building a feedback tree with the receivers as source, having the original sender as the only sink. Whenever an intermediate node receives an ACK from each of its children, it will forward a single ACK to its own parent.*
- 7 The remote-access model is becoming less popular in distributed file systems. Why? 5pt
- The answer is simple: in order to facilitate clients that are far away (and thus incur high latencies in communication), it makes more sense to allow for advanced client-side caching, thereby letting the client manipulate the requested file locally. The server takes care of consistency management. In the remote-access model, all file operations requested by a client are carried out at the server.*
- 8a Systems like Globule determine per Web document the strategy according to which that document will be replicated. What is the advantage of this approach in comparison to adopting a single, site-wide replication policy? 5pt
- The fact that the unit on which you decide to replicate is smaller, more or less automatically leads to better resource usage. Many content delivery networks reach efficiency in terms of client-perceived latency and consumed bandwidth by selective replication of documents. Globule does this as well, in addition to deciding per document how, where, and when it should be replicated.*
- 8b Considering that a content delivery network such as Akamai uses existing Web caching techniques for replication, what type of consistency does it provide to clients? 5pt
- This is a bit tricky. In fact, they provide strong consistency as modified (embedded) documents are simply renamed. As a result, when a client is directed to a replica server, that server will not find the document in its cache and fetch it from the main server.*
- 9a What is an important advantage of using a message broker for the implementation of a publish/subscribe system? 5pt
- Simplicity: you can simply keep a list of subscribers at a single place and manage/route messages based on such a list.*
- 9b What makes subject-based publish/subscribe systems so difficult to scale across a wide-area network? 5pt
- The real problem comes from trying to efficiently set up routes between publishers and subscribers. In principle, we are dealing with multicast trees where intermediate nodes act as routers. These routers need to know about subscribers that can be reached through their outgoing links. Setting up tables such that only those messages are forwarded across links known to lead to subscribers makes matters complicated as systems grow.*

**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.*