

Part I

This part covers the same material as the midterm exam.

- 1a Explain what is meant by administrative scalability and why it is often such a difficult problem to solve. 5pt
- 1b What are the main differences between a network operating system and a (middleware based) distributed system? 5pt
- 2a Give at least two examples of a middleware protocol. 5pt
- 2b What is the difference between passing an object by reference or an passing object by value in the case of a remote object invocation? 5pt
- 3a Motivate the extensive use of threads in distributed systems. 5pt
- 3b Consider a mobile agent having its own thread of control. Are we dealing with weak or strong mobility when the agent moves to another machine? Motivate your answer. 5pt
- 4a Explain what a closure mechanism is in naming systems, and give at least one example. 5pt
- 4b Iterative and recursive name resolution in a distributed naming system have different scalability properties. Explain what the most important difference is. 5pt
- 5 Explain what is meant by a consistent global state when dealing with distributed snapshots. 5pt

Part II

- 6a A file is replicated on 10 servers. List all combinations of read quorum and write quorum that are permitted by the quorum-based voting protocol when dealing with replicated writes. 5pt
- 6b Explain how an epidemic-based distribution protocol works for propagating updates. 5pt
- 6c Give an example of a causal-consistent execution of read and write operations in a distributed system. 5pt
- 6d It has been argued that middleware protocols for achieving causal consistency should be replaced by application-level protocols. What is a strong argument in favor of such a replacement? 5pt
- 7a Explain what a virtually synchronous reliable multicast is and what the advantage of this scheme is. 5pt
- 7b What is a pessimistic message logging scheme and why is it preferred over optimistic approaches? Be precise. 5pt
- 8a Sketch two different approaches for implementing a subject-based publish/subscribe system. 5pt
- 8b How does Jini realize referential decoupling of communicating processes? Can this method scale? 5pt
- 8c What is the (semantic) difference between writing a tuple instance to a transparently N times replicated JavaSpace, and writing that tuple instance N times to a nonreplicated JavaSpace? 5pt

Final grade: (1) Add, per part, the total points. (2) Let T denote the total points for the midterm exam ($0 \leq T \leq 45$); $D1$ the total points for part I; $D2$ the total points for part II. The final number of points E is equal to $\max\{T, D1\} + D2 + 10$.