

Review Questions

Instructions: You are not required to turn in answers to these review questions, but you should do them on your own after you read this section's papers to test your knowledge. Use Piazza or the TAs if you have questions or get stuck on any of these.

Lamport's paper

1. Lamport's happened before relationship by itself is sufficient to derive a total order of events in a distributed system. True or False? Justify your answer with examples if necessary.

2. The system of logical clocks proposed by Lamport is sufficient to develop any "real world" distributed algorithm. True or False? Justify your answer with examples if necessary.

3. Assume that messages arrive out-of-order in a distributed system. Does this violate the "happened before" relation? Justify your answer with examples if necessary.

4. Assume that messages arrive out-of-order in a distributed system. Can this result in erroneous behavior of the mutual exclusion algorithm proposed by Lamport? Justify your answer with examples if necessary.

If it could lead to erroneous results how can you fix it?

[Hint: consider using sequence numbers for messages between pairs of processes.]

5. Come up with a distributed mutual exclusion algorithm which reduces the message complexity to $2(N-1)$, where N is the number of processes competing for a lock.

Limits to low latency communication on high speed networks

1. Explain the methodology used in the Thekkath and Levy paper for teasing out the component times (reported in Table I) for a lower bound on cross machine communication.

2. As a designer of an RPC subsystem, what are the avenues available to you for shaving the marshaling and unmarshaling software overhead? Explain your answer with respect to the expected benefit, potential protection concerns for each such avenue.

3. As a designer of an RPC subsystem, what are the avenues available to you for shaving the control transfer software overhead? Explain your answer with respect to the expected benefit, potential protection concerns for each such avenue.

4. Discuss the usual functionalities found in popular protocol stacks such as TCP/IP and their appropriateness/inappropriateness for a transport meant for an RPC subsystem. Propose a minimal functionality transport stack for an RPC system which will still ensure that the higher level semantics of RPC are preserved.

5. Explain each component cost that appears in the critical path of RPC latency.
6. Consider all the potential context switches that can happen during an RPC call/return. Explain which of these appear in the critical path of RPC latency.
7. If the network controller does not have the ability to do scatter/gather DMA then there will always be at least 1 copy incurred for each of the call and return path of RPC. True or False? Justify your answer.
8. Consider an RPC subsystem.
 - (a) Give a breakdown of the costs involved in implementing an RPC subsystem.
 - (b) As a system designer, what are the avenues available to you for shaving each of the component cost you identified in part (a)? In discussing such avenues, you have to clearly state what assumptions you are making about the execution environment to shave the costs, and the pros and cons of your design choices.

Active Networks

1. Argue for and against the idea of executable capsules flowing through the routers of the Internet from the point of view of performance.
2. Argue for and against the idea of executable capsules flowing through the routers of the Internet from the point of view of protection and integrity of all the other traffic.
3. Argue for and against the idea of executable capsules flowing through the routers of the Internet from the point of view of protection and integrity of resource management.
4. Consider an application wanting to send the same message to N of its peers. Let T_m be the time for the application to construct and send a message. Let H_c be the number of common hops that the message has to traverse. Let H_u be the average number of unique hops to get to each of the intended destinations. Let T_a be the time required to send a message over one hop. Construct a simple mathematical model and show when an active network that implements a multicast service in the network may benefit over a vanilla Internet. Assume that a message fits in one IP packet.

Ensemble and Nuprl

1. The authors cite features that justify the choice of OCaml as the systems programming language. Yet when it comes to optimization opportunities of the layered systems code, they point to many of these same features as the ones to circumvent to achieve good performance! Is there a contradiction? Explain why or why not.

2. Distinguish between properties, abstraction behavioral specification, concrete behavioral specification, and implementation.
3. Using Lamport's distributed mutual exclusion algorithm as an example, develop an abstract specification and a concrete specification. You do not have to adhere to any specific syntax (such as IOA), and can invent your own so long as it is straightforward and understandable.
4. Ultimately this paper is about specializing of a subsystem (communication stack to be specific in this case). Other papers we have seen so far in this course (SPIN, Exokernel, and Thekkath and Levy) also deal with a similar theme. What are the similarities and differences between Ensemble/Nuprl and these other works?