CS 6210 Spring 2024 Test 2

(130 min Canvas Quiz)

Max Points:104

Distributed Systems [30 points]

Lamport's Logical Clock [8 points]

1. (8 points) A student has implemented a distributed algorithm using Lamport's happened-before relationship to timestamp events. The student is in the middle of debugging the program, and observes the following activities in the program:

P1's activities	P2's activities	P3's activities
E1: msg-send (to P2)	E5: local event	E9: msg-receipt (from P2)
E2: local event	E6: msg-receipt (from P1)	E10: local event
E3: msg-receipt (from P2)	E7: msg-send (to P3)	E11: msg-send (to P1)
E4: msg-receipt (from P3)	E8: msg-send (to P1)	

Please give the causal relationship between the following pairs of events with reasoning. (No credit without reasoning)

(a) (2 points) E1 and E10?

(b) (2 points) E2 and E8?

(c) (2 points) E8 and E11?

(d) (2 points) E2 and E11?

Lamport's ME Lock [4 points]

1. (4 points) Give 4 conditions on which the correctness of the ME lock relies on.

RPC Latency limits [10 points]

1. (5 points)

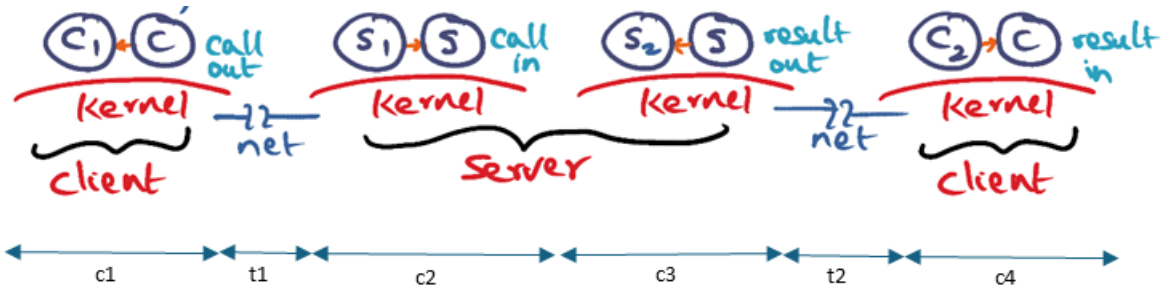

In the context of this question and its subparts, consider process context switches and transmission on physical links as the only sources of latency in RPCs (any other source of latency is negligible).

The figure above shows the timeline of a simplified RPC. Any time period demarcated by 'c' followed by a numeral represents the latency introduced by the corresponding context switch for the associated arrow. Similarly, any time period demarcated by 't' followed by a numeral represents the latency introduced by the corresponding transmission over the physical link for the associated arrow.

For this question and its subparts, only use variables provided in the above model of the RPC.

(a) (2 points) If the context switches not in the critical path of the RPC latency are overlapped with other activities in the critical path of the RPC latency, give a mathematical expression of the expected RPC latency.

(b) (3 points) Can we further optimize the RPC latency? If yes, give the mathematical expression of the optimized RPC latency and the mathematical condition under which such an optimization would be justified. If no, provide appropriate justification.

2. (3 points)

(a) (1 point) How does the server know that an incoming RPC call is not a duplicate?

(b) (1 point) What is the purpose of the server-side buffering of the RPC response in a reliable LAN setting?

(c) (1 point) How long will this buffered RPC response be kept on the server?

3. (2 points) Thekkath and Levy suggest having a shared descriptor between the kernel and the RPC client stub for reducing the number of copies in marshalling the arguments of the RPC. The alternative is for the client stub to be downloaded into the kernel. What is the downside of this alternative?

Active Networks [8 points]

1. (8 points) You have been given the design of the capsule that will enable the intermediate routers to execute the code for incoming packets for Active Networks. As discussed, if a node receives a

capsule whose TYPE has not been seen before, the node will request to load the code from the PREV active node.

(a) (1 point) Why is it possible for the PREV node to not have the code corresponding to the TYPE field?

(b) (1 point) How often will this occur (justify your answer)?

2. (2 points) Why does it make sense to "drop" the packet if the PREV node does not have the code?

3. (4 points) Software Defined Networking (SDN) is often considered as the modern-day evolution of Active Networks. Give two features of SDN that overcome the drawbacks of Active Networks.

Distributed Objects and Middleware [15 points]

SPRING OS [7 points]

1. [7 points] Assume you are managing a cluster of nodes in a datacenter that is using Spring OS as the network OS. The current deployment looks like this:

   A. Web servers replicated on 3 nodes
   B. SQL Database servers replicated on 2 nodes
   C. File server on a node accessed only by the web servers.

Each of the above servers are hosted on distinct nodes. The clients are expected to make requests to both the web servers and database servers.

(a) [3 points] You wish to add a web-server-load-balancer that balances the load of the client requests across the web servers. List the subcontracts needed to effect this architectural change.

(b) [2 points] While all clients can read from the database servers, only certain clients are allowed to write to the database servers. How can you accomplish this?

(c) [2 points] Your colleague has developed a new custom communication protocol to access the file server and wants to test it. How can he leverage the extensibility of Spring OS to perform his tests?

Enterprise JavaBeans [8 points]

1. (4 points) Consider the following design alternative for organizing an N-tier application using EJB.

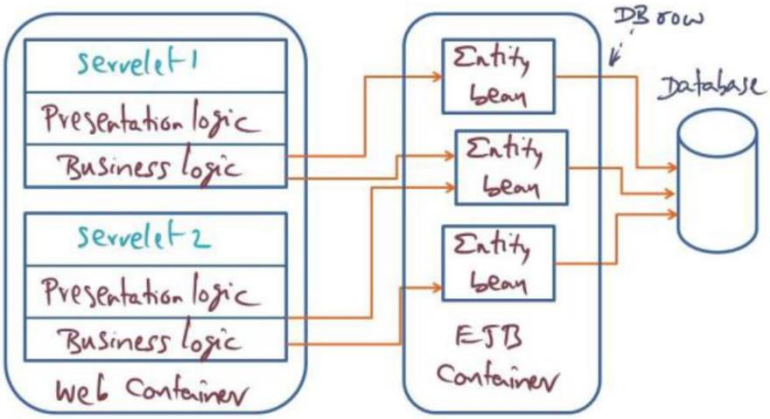

(a) (2 points) Describe one benefit of structuring an application this way.

(b) (2 points) Describe one disadvantage of structuring an application this way.

2. (4 points) Imagine you are the designer of an airline booking portal like expedia. You have been tasked with building a multi-tier software architecture for this portal. Enumerate and describe the tiers you will have in your service (concise bulleted list please). [Hint: You have used such services; now you have to turn around and ask yourself how you will design such a portal given all the knowledge you have gained through this course.]

Distributed Subsystems [39 points]

DSM [10 points]1. [10 points] Consider the following sequence of actions in the following time order happening in a Treadmarks DSM program. Assume a clean copy of X is with the owner at the start of the program, and the program starts execution at time T1. T2 and T3 represent increasing time order.

T1: Process P1:

acq(L1)

modify X

rel(L1)

T1: Process P2:

acq(L2)

modify X

rel(L2)

T1: Process P3:

acq(L3)

modify X

rel(L3)

T2: Process P4:

acq(L1)

modify Z

rel(L1)

T2: Process P5:

acq(L2)

modify X

rel(L2)

T3: Process P6:

acq(L2)

modify X

rel(L2)

a) [5 points] Note that the lock requests from P1, P2, and P3 occur at the same time T1 and all of them modify the same page X in their respective critical sections.

(i) (2 points) How would Treadmarks handle this situation?

(ii) (3 points) What actions will be taken by Treadmarks in executing each of these critical sections (at lock acquisition, during the critical section, and upon lock release)

b) [5 points] What actions are taken by Treadmarks when P6 tries to acquire lock L2 at time T3 (at lock acquisition, during the critical section, and at lock release)?

GMS [15 points]

1. (8 points) Node P faults on page X, which is not present in any of the peer nodes the cluster. Node R houses the globally oldest page Y. Explain how GMS would strive to handle this page fault. Your answer should cover all possible scenarios. Note that this question is NEITHER about the actual implementation of GMS, NOR the Geriatric algorithm.

2. (4 points)
For this question, assume that GMS implementation has a "master node" that is always up. A new node N joins the GMS cluster. What actions will ensue as a result of this new node joining the GMS cluster?

3. (3 points) Explain how the Geriatrics algorithm illustrates the principle "Think Globally but Act Locally".

DFS [14 points]

1. (6 points) You and your friend are tasked with implementing a disk-based distributed file system for the College of Computing. You are given a LAN cluster interconnected by 100 Gbps links. Your system should meet the following design objectives:

- Reduce disk accesses as much as possible
- Enable parallel processing of requests as much as possible
- Ensure no file server becomes a hotspot for requests
- Maximize the use of available compute and physical storage in the cluster to enhance performance

(a) [2 points] Your friend suggests storing the metadata for a file at the server that hosts the file on its local disk for simplicity of implementation. What are the downsides to this suggestion?

(b) [2 points] You decide to decouple the server location of the metadata of a file from the server that hosts the file on its local disk. How does this reinforce the design objectives?

(c) [2 points] You and your friend discuss striping files across distinct and disjoint subsets of node storages. How does this reinforce the design objectives?

2. [8 points] For the purposes of this question assume that the granularity of access is an ENTIRE file. Consider the following scenario in xFS. The contents of file F are in log segments that are striped across storage servers S1, S2, and S3. Node M1 is the statically assigned metadata manager for file F. Currently, file F is in the local cache of Node N1. Node N2 wants to make modifications to F. Assume that S1, S2, S3, M1, N1, and N2 are different nodes of the cluster.

List all the actions and the modifications to the data structures of xFS at the different nodes to satisfy the above modification request of N2.

[HIDDEN] Potpourri [20 points]