

# Systems Qualifier Exam Spring 2012

---

Score \_\_\_\_\_ /100

Name (please print): \_\_\_\_\_

SID: \_\_\_\_\_

**Instructions:**

- This is individual exam, to be completed closed book and closed notes.
- Each sub-question (e.g. (a), (b), (c), (d)) is worth an equal amount of points (e.g. 5 points)
- Please use the back sides of each page for additional room if needed.
- You have 120 minutes to complete this exam.

**Q1 (20 points): Communication**

(a) What is **persistent** and **transient communication** in distributed systems?

(b) What is **synchronous** and **asynchronous communication** in distributed systems?

(c) Communication between processes can be characterized along different dimensions. Give an example for each of the following combinations.

	transient	persistent
synchronous		
asynchronous		

(d) Suppose that you could make use of only asynchronous communication primitives. How would you implement *synchronous* communication?

Score \_\_\_\_\_ /15

**Q2 (15 points): Fault Tolerance**

(a) What is meant by the statement that distributed systems are characterized by the occurrence of **partial failures**?

(b) What is a **k fault-tolerant group**, and how does k depend on failure semantics?


(c) Does **TMR (triple modular redundancy)** generalize to five elements per group instead of three? If so, what properties does it have?

**Q3 (15 points): Synchronization**

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

(b) Explain how **Vector logic time** works.

(c) Please give the Lamport and Vector logical time of the following execution.



Score \_\_\_\_\_ /15

**Q4 (15 points): Processes & Threads**

(a) Describe what a process and a thread is, and the difference between them? Why are synchronization locks needed with threads?

(b) Describe what shared address space and message passing is, and the difference between them? In what environments would one be used over the other?

(c) Constructing a concurrent server by spawning a process has some advantages and disadvantages compared to multithreaded servers. Name two pros and cons.

**Q5 (20 points): Client-Server Architecture & Remote Procedure Calls**

- (a) In this problem you are to compare reading a file using a single-threaded file server with a multi-threaded file server. It takes 5 msec to get a request for work, dispatch it, and do the rest of the necessary processing, assuming the data are in the block cache. If a disk operation is needed, as is the case one-fourth of the time, an additional 8 msec is required. What is the throughput (requests/sec) if a single threaded server is required? What is the throughput if the server is multi-threaded?
- (b) Suppose that the time to perform a null RPC (i.e. 0 data bytes) is 0.5 msec, with an additional 4 msec required for every 1K of data. How long does it take to read 16K from the file server using a single 32K RPC? How about 32 1K RPCs?
- (c) What mechanism is used to search unstructured peer-to-peer overlays? What is the worst case scenario for locating the object in question in terms of number of nodes visited? What about structured overlays?
- (d) Outline an efficient implementation of globally unique identifiers.

**Q6 (15 points): Shared/Parallel/Distributed File Systems**

(a) Define the following, and give an example of each:

- local filesystem (FS)
  
- shared filesystem (SFS)
  
- parallel filesystem (PFS)
  
- distributed filesystem (DFS)

(b) Despite that the Google File System (GFS) scales well, it could be argued that the master is still a potential bottleneck. Define the purpose of the master. What would be a reasonable alternative to replace it? What properties would be improved and what might be worse? Please address scalability, performance, latency, and reliability.

(c) What is the default replication level in GFS? Explain the motivation behind this number.