

COSC 6360—Operating Systems
MATERIALS ON THE THIRD FALL 2020 QUIZ
October 26, 2020

You are only responsible for the materials discussed in class as they are *summarized* in the handouts and discussed in the PowerPoint presentations. I expect you to understand these summaries and to be able to comment around them. Always ask yourself *why* a specific technique was used and *which* problem it addressed.

Distributed Systems Issues

L. Lamport, Time, clocks and the ordering of events in a distributed system, *CACM*, 21(7) 558–565, July 1978.

- *We covered most aspects of the paper: skip the proof at the end.*

D. Ongaro and J. Ousterhout, In search of an understandable consensus algorithm, *Proc. 2014 USENIX Annual Technical Conf. (ATC 2014)*. pp. 305-319, June 2014. *Best Paper Award*

- *We covered most aspects of the paper: focus on the basic operation of Raft, the concept of a leader, how updates are propagated, how new leaders are elected, how new leaders force their updates to their followers.*

Security

T. Ylonen. SSH -- Secure login connections over the internet, *Proc. 6th USENIX Security Symp.*, pp. 37-42, July 1996.

- *Focus on the login protocol and the general overview of cryptography in the slides.*

Peer-to-Peer Systems

P. Shah and J.-F. Paris. Peer-to-Peer Multimedia Streaming Using BitTorrent, *Proc. 26th International Performance of Computers and Communication Conf.*, pp. 340-347, Apr. 2007.

- *Focus on the aspects discussed in the PowerPoint presentation. You should understand the sliding window protocol, how chunks are selected within the sliding window and the modified tit-for-tat policy.*

File Systems

D. A. Patterson, G. A. Gibson, and R. H. Katz. A case for redundant arrays of inexpensive disks (RAID), *Proc. SIGMOD 1988 International Conf. on Data Management*, pp. 109-116, June 1988.

- *You should understand the general motivation for RAID organizations and the small write problem in RAID 5. Skip the availability analysis: it is wrong.*

M. Rosenblum and J. K. Ousterhout, The design and implementation of a log-structured file system, *ACM Trans. on Computer Systems*, 10:1(26-52), Feb. 1992.

- *You should understand the motivations for a log-structured file system, its general organization (including the new i-node map) and the trade-offs to be considered in the segment cleaning process.*

Review questions

1. A system of physical clocks consists of two clocks, namely, one that is slow and loses 5 minutes every hour and another that is fast and advances by 5 minutes every hour. Assuming that the clocks are managed by Lamport's physical clock protocol, what will be the time marked by each clock at 2:00 PM given that
 - a) both clocks indicated the correct time at noon;
 - b) the processors on which the clock reside continuously exchange messages between themselves; and
 - c) the message transmission delays are negligible.
2. Redo problem 3 assuming that the processors on which the two clocks reside stopped communicating with each other at one pm.
3. What does RAFT do to minimize the chances of *split votes* during *leader elections*?
4. How does SSH authenticate a *new server* (in opposition to a *new user*)?
5. According to Shah et al., what is the main motivation for their *randomized tit-for-tat* policy?
6. A RAID level 5 array consists of n drives. What is the most efficient way to update a *single block* whose previous contents were *never read* into memory when $n = 8$?
7. What is the purpose of the *i-node map* of a log-structured file system? Where is it stored? How does the system retrieve it at boot time?
8. What is the cost of a write in a log-structured file system, when its segment cleaner has to clean three segments to produce two clean segments?

Answers: I. Both clocks will mark two ten. II. The fast clock will mark two ten and the slow clock will mark two. III. It uses randomized timers so that different servers will detect the failure of their leader at different times. IV It asks users whether to accept or to reject the public key of the new server. V. To speed up delivery of video data to newly arrived peers. VI. When the RAID array has eight disks, the best solution is to compute the new parity from the old value of the parity block and the old value of the block that is being updated. VII. The i-node map of a log-structured file system contains the addresses of all i-node blocks. It is stored on the log. The checkpoint area of the file system contains the addresses of all i-node map blocks at checkpoint time. VIII. Each write will cost $2/(1 - u) = 2/(2/3) = 3$ disk accesses.