

## 91.515. Fall 2011 Exam 2 Help File

This help file will enumerate the topics that you should focus on for the second exam. It does not contain sample questions, but the question format is similar to the first exam.

The text coverage includes the second half of chapter six, beginning with the token based algorithms, and continues through chapters 7,8,9,10 and 11. Although I have included chapters 10 and 11 in the reading assignments, I will not test over material in either chapter on this exam.

### Ch 6, Token based algorithms:

- You should have operational knowledge of each of the three algorithms discussed here.
- Potential problems may include:
  - A tree depiction of nodes that use Raymond's algorithm to share a token. A node may have a request queue and the tree will have a current token position at any given time. You should be able to determine which node the token would move to after its current user finished with it.
  - Using the Singhal algorithm with the SSV and SSN arrays for a site, and the TSV and TSN arrays for the token as it arrives at the site, you should be able to determine how the arrays are updated after the site finishes with the token, and where the token could go next. You should also be able to handle the same situation with the Suzuki-Kasami algorithm, given a site vector and a token vector.
- You should be able to contrast the behavior of the token based algorithms with those of the non-token based algorithms discussed in the first half of the chapter

### Ch7. Distributed Deadlock Detection:

- You should be able to describe why detection is the only feasible treatment of deadlock management in the distributed world.
- The chapter addresses centralized, distributed and hierarchical detection, and you should be able to contrast the three methods, pointing to the pros and cons of each
- Detection in distributed systems must meet progress and safety requirements.
- The work of Ho and Ramamoorthy is discussed in the first part of the chapter and then again in the last part of the chapter. You should be familiar with their algorithms.
- The distributed algorithms are divided into those that can solve the AND request model and those targeted for the OR request model. Specifically you should be able to work with the edge chasing algorithms described and with the diffusion computation work of Chandy et al

### Ch8. Agreement Protocols:

- You should understand the agreement and validity requirements of these algorithms.

- The number of faulty nodes in an N node network that can be tolerated and still allow non-faulty nodes to reach agreement is a fundamental ratio that you must know.
- The synchronous and asynchronous communication models are also critical in treating these agreement problems
- You should be able to briefly describe the agreement problems discussed in class, and have a detailed understanding of the Byzantine agreement problem
- The message complexity of the Lamport algorithm and the Dolev algorithm are important, as are the message round requirements of each.
- You must have a detailed operational knowledge of Lamport's algorithm, and the ability to describe the basic operation of Dolev's algorithm.

#### Ch9. Distributed File systems:

- Although the text considers various systems, the exam will only test over the NFS and CIFS (SMB) protocols.
- You should be able to contrast these two systems on the basic stateless vs stateful implementations that they deploy
- Details regarding coherence and coherence widows are important, as are locking details in each system
- In general, you should be able to describe the most favorable and least favorable attributes of the two protocols
- You should understand the problems associated with non-idempotent transactions with the NFS protocol, and some of the implementation details used by developers to address these issues (e.g. the create/unlink problem)