Due: 2 p.m., December 6

Problem 1

How many lookup calls are necessary to resolve a five-part pathname (for example /usr/users/sam/../abc.txt) for a file that is stored on an NFS server? What is the reason for performing the translation step-by-step?

Problem 2

Why should UFIDs be unique across all possible file systems? How is uniqueness for UFIDs ensured?

Problem 3

Explain in which respects DSM is suitable or unsuitable for client-server systems.

Problem 4

Discuss whether message passing or DSM is preferable for fault-tolerant applications.

Problem 5

Discuss whether the following operations are *idempotent*:

1. Pressing an elevator call button
2. Sorting a list
3. Appending to a file

Is it a necessary condition for idempotence that the operation should not be associated with any state?

Problem 6

Describe a distributed computation that can be implemented using MapReduce. (Please come up with something on your own, as opposed to an example you have seen elsewhere). Sketch the pseudocode for the map() and reduce() functions.

Problem 7

Provide an example of a distributed computation that would be difficult to implement in MapReduce, giving full reasons for your answer.