

CSc 552 Final Exam

December 17, 2002

Name:

SID:

You may not use books, notes, or calculators. You have 2 hours to complete the exam. The number in parenthesis at the beginning of each question indicates the number of points given to the question. Do all work on these sheets, using the backs if necessary.

Problem	Points Scored	Points Possible
1		20
2		10
3		15
4		25
5		15
6		15
Total		100

(f) Munin provides *release consistency* of shared variables.

(g) In the Andrew file system a user authenticates himself/herself to a Vice server by binding to it using his/her password as the key.

(h) A Plan-9 process has its own private namespace, making it difficult for processes to share files.

(i) The Cooperative File System (CFS) uses a virtual server abstraction to balance loads across physical servers with different capacities.

(j) With Scheduler Activations, a process is given one activation for every thread it creates.

-
2. (10 points) Describe a scenerio in which using Mach's external pagers can cause a deadlock.

3. (15 points) Suppose Bob wants to read the file */foo/bar/baz.txt* that Alice has stored on CFS. He wants to be sure that the file was created by Alice, and Charlie has not modified it. Describe how CFS enables Bob to do this.

4. (25 points) For each of the following scenerios, describe how it is handled in the NFS, Sprite, and AFS file systems, and whether or not it is possible for applications to read stale data.

(a) A source file is compiled on one client immediately after being edited and saved on another.

(b) One client opens and continually writes to a log file (keeping it open), while another client opens and reads from it.

(c) Two clients simultaneously open, write, and close the same file.

(d) One client deletes a file while another client is reading from it.

5. (15 points) What is the difference between *idempotent* and *non-idempotent* operations? Give an example of each. Why is this distinction important to *at-least-once* vs. *at-most-once* RPC semantics?

6. (15 points) LFS originally used a greedy cleaning policy, choosing to clean the segment containing the fewest live bytes. Surprisingly, this was found to lead to poor cleaning performance (high write cost). Explain why.