Stanford University
Computer Science Department
CS 240 Quiz 3
Spring 2002

This is an open-book exam. You have 50 minutes to answer 6 short answer questions and one long answer question. Write all of your answers directly on the paper. Make your answers as concise as possible. Sentence fragments ok.

NOTE: We will take off points if a correct answer also includes incorrect or irrelevant information. (I.e., don't put in everything you know in hopes of saying the correct buzzword.)

Question	Score
1 - 2	
3 - 5	
6 - 8	
9	
total	

Stanford University Honor Code

In accordance with both the letter and the spirit of the Honor Code, I did not cheat on this exam nor will I assist someone else cheating.

Name and Stanford ID:

Signature:

Short-attention-span questions (30 points)
Answer **6 of the first 8 questions** (i.e., skip two) and, in a sentence or two, say *why* your answer holds. Please state any assumptions you make in your answer. (5 points each).

1. You are cleaning segment 7, which contains a file block 10 immediately followed by its inode, 12. If we can reclaim the space for block 10 are we guaranteed that we can reclaim the space for 12? Similarly, if we can reclaim the space holding inode 12 do we know that the space for block 10 is free as well?

2. Figure 2 in the soft updates paper illustrates a specific write order. Redraw the last two steps ((c) and (d)) to show what happens if the last two writes were reversed and give an intuitive description of what has to happen.

3. If the block size is 8K, how many direct block pointers can you infer that an inode has in the soft updates version of FFS based on the performance curve in Figure 3?

4. Suppose you have a RAID 5 array with 6 disks. What is the best case number of simultaneous writes we could make? What is the worst case? (Please clearly state any assumptions you make.)

5. Assume you are still using RAID 5: what RAID-specific recovery will you have to do after a crash to get to consistent state? Is there a problem with crashing during recovery?

6. Disk drives typically only guarantee that individual 512-byte sectors are written out atomically. What implications does this have for logs?

7. Log checkpoints are supposed to allow for faster recovery after a failure. Describe a scenario where checkpoints can make recovery slower.

8. Butler makes the point that quantizing a resource into fixed sized chunks can significantly help resource management. (E.g., clothing is quantized into S, M, L, or XL rather than a continuous set of sizes.) Give three examples from the papers we read where such a trick has been used.

Problem 9: Yes, we have no chaos (20 points)

Assume you have a spiffy new solid-state device (SSD) that can hold 60 GB of memory, costs about as much as a normal disk, and can be accessed in the same amount of time as a RAM reference. You buy it of course and immediately start porting your traditional Unix file system to it (i.e., your file system has directories, inodes, indirect blocks, etc.) The new SSD device can only be written one byte at a time with a "store byte" instruction. On crash, the last byte is either written in its entirety, or not at all.

Explain, concretely, the steps you have to take to create a new file in a directory and what problems you have to guard against. Be sure to handle the case where you crash while doing so and how you would recover afterwords. You should only use the metadata structures that already exist, with possibly (very) minor modifications — *i.e., do not add a log.*