Stanford University
Computer Science Department
CS 240 Quiz 1
Winter 2003

April 21, 2003

This is an open-book exam. You have 50 minutes to answer eight out of ten questions. 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 - 10	
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:

Answer **8 of the following 10 questions** (i.e., skip two) and, in a sentence or two, say *why* your answer holds. (5 points each).

1. The Therac code in Figure 4 on page 34 shows an overflow error involving the "Class3" variable. If we assume the code in Figure 4 accurately represents the code in the Therac, describe the race that exists between Set-up Test and Lmtchk.

2. The set of locks that are acquired at any given time is a global property (e.g., switching threads has no effect on it). Given this, why does Eraser use per-thread locksets rather than a single global lockset? Give an intuition and a simple example of the problem a global lockset would cause.

3. The Eraser paper claims "A write access from a new thread changes the state from Exclusive or Shared to the Shared-Modified state..." But Figure 4 says that a write by any thread in the Shared state takes it to the Shared-Modified state. Which is right?

4. Your ex-140 partner loudly declares that if the semantics of the "wakeup-waiting switch" is good for naked notifies, it must be good for normal notifies, which should be replaced with them. Can you do this substitution and preserve correctness? (List any assumptions you must make.)

5. When does the coalesce daemon run in the superpage system? What would the effect of running it more often be?

6. Your ex-140 partner (fresh from their triumphant foray into monitor semantics) decides to get the same effect of a balloon driver using an application running on top of the guest OS. On ESX's request the application will aggressively touch many pages of memory, causing them to be "recently used" and the guest OS to page out other, less recently used pages. Is this a bright idea?

7. VMS: You increase the size of the free list to 80% of memory. Describe a workload that will run better, and one that will run worse.

8. Assume you want to add ESX memory taxation to the VMS system. In a few sentences, describe why this might make sense and how you would integrate it.

9. Scheduler activations: list the events that cause the kernel to do an upcall.

10. Assume you put all processes in the NUMA system on a single global run queue. Which of the NUMA benchmarks do you expect to be affected by this change and in what way? Do you expect this difference to be at all sensitive to how long the application takes to complete?