Stanford University
Computer Science Department
CS 240 Quiz 1
Spring 2003

April 24, 2003

This is an open-book exam. You have 50 minutes to answer eight 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	
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 all 8 of the following 8 questions and, in a sentence or two, say *why* your answer holds. (5 points each).

1. The Therac paper is a study in severed feedback loops, ranging from AECL not informing customers of problems in the Therac-25 to the video camera being off when a patient was zapped. From the paper, give three (other) examples of such missing feedback loops.

2. Assume we have two threads, T1 and T2, that execute the following code in the following order:

```
int *q = NULL; /* shared global variable */

        T1                               T2
  int *p = malloc(sizeof *p);
  *p = 0;
  q = p;
                                         while(q == NULL)
                                              ;
                                         *q = 1;
                                         q = 0;
 while(q != NULL)
     ;
 printf("p = %d\n", *p);
 exit();
                                          exit();
```

Assuming naive memory semantics (if you don't know what this means, you are just fine): What will Eraser do for this code? Is its behavior correct?

3. For the following snippet of Mesa monitor code:

```
ENTRY void foo() {
    signal(c);
}
```

If the code runs on a system with different priority processes, what is a potential way that it will generate useless overhead? Is there an automatic way to fix this problem?

4. What would Hoare-wakeup semantics correspond to in a message passing system?

5. Identify two potential weaknesses in the VMS memory system and describe experiments to measure them.

6. When the reference bit of a base page within a superpage is reset, the superpage management system demote the superpage speculatively, recursively performing the demotions with a probability $p = 1$. Suppose we set $p = 0.5$ instead. Give one reason why the system might perform better because of that, and one reason why the system might perform worse because of that.

7. When measuring active memory (for use in idle memory tax calculations), ESX Server keeps three statistical averages: a "slow" average, a "fast" average, and a "very fast" average. Explain the problem that could result if ESX stopped using the "slow" average.

8. Assume you can mark pages as non-executable. What SFI restrictions does this lift and how could SFI exploit this flexibility?