CS 4760	**Operating Systems**	Test 3
Name:	Fall 2013	Max Pts: 55

Important: This is an open book test. You can use any books, notes, or paper but no electronic device. *Do not log into the computer during the test, or use any electronic or communications device. Change your cell phones to slent mode*. Any calculations and rough work can be done on the back side of the test pages. If there is a syntax error in any program segment, just write it down and you will get full credit for the problem. You will lose five points for not writing your name.

1. [6 pt] The code that has to be shared across processes must be written as reentrant code or pure code. Such code cannot call any function that is not reentrant. Linux kernel is reentrant but includes calls to some nonreentrant functions. How does it ensure the sharing of code in those cases?

2. [6 pt] Compare the working set model and page fault frequency model to tune an OS to prevent thrashing.

3. [10 pt] Consider a machine with disk blocks of 4KiB. You have a new disk of size 4TiB. What is the theoretical maximum file size possible using the UFS scheme of file allocation?

4. [9 pt] Ignoring overhead for directories and file descriptors, consider a file system in which files are stored in blocks of 4KiB. For each of the following file sizes, calculate the percentage of wasted file space due to incomplete filling of the last block: 22,919 bytes, 495,260 bytes, 2,702,100 bytes.

5. [6 pt] What is the advantage to be gained by using a memory-mapped file? How will you increase the size of such a file in Unix?

6. [6 pt] Give at least two advantages of the bit-vector approach to manage free space in a disk. Why is it not recommended to use this approach in a large disk? Explain with an example.

7. [6 pt] How does the use of `defrag` utility improve the performance of a Windows machine?

8. [6 pt] What is security through obscurity? How can I implement that in Unix for my files?