NAME: _____ (**FIRST NAME FIRST**) **SCORE:** _____

COSC 4330/6310 FINAL JULY 25, 2012

This exam is **closed book**. You can have **one** page of notes. Your test has three pages.

1. Which of the seven following page replacement policies (3×5 points)

Global LRU ***Global FIFO*** ***Original Clock*** ***Berkeley Clock*** ***Mach*** ***Windows***

a) Simulate a ***page referenced-bit*** ? Berkeley Clock

b) Incorporate a ***global queue*** ? Mach and Windows

(I would also accept Gloval LRU and Gloal FIFO)

c) Were designed to handle ***real-time*** applications? Windows

2. For each of the statements below, indicate in one or two sentence whether the statement is ***true*** or ***false***, and ***why***: (3×5 points)

a) UNIX lets each user select the block size of his or her files.

FALSE, the block size is set by the system administrator for all files in a given disk partition.

b) Internal fragmentation is much more of an issue in file systems than in virtual memory systems.

TRUE, because the same disk partition might contain both very small and very small files.

c) In the Fast File System, each ***cylinder group*** contains ***either*** inodes ***or*** data blocks.

FALSE, each cylinder group contains both inodes and data blocks.

3. ***Questions with short answers*** (6×5 points)

a) How do you prevent ***thrashing*** in a virtual memory system?

By ensuring that all programs have their working set of pages in physical memory. This requires a good page replacement policy and enough physicial memory.

b) Why are TLB entries ***bigger*** than regular page table entries?

Because they must contain the virtual page number in addition to the page frame number.

c) How does the Fast File System ensure that all ***metadata updates*** are executed in the ***right order***?

By forcing all metadata updates to be BLOCKING: this ensures that these operations are executed one by one in strict sequence.

d) Describe the full contents of a UNIX ***directory entry***?

Each directory entry contains a name and an inode number (and NOTHING ELSE!)

e) In the Windows page replacement policy, what is the cost of ***reclaiming*** a page from the ***global queue***?

Two context switches.

f) What is the main advantage of using ***asynchronous log updates*** in a ***journaling file system***?

Faster updates.

4. An old virtual memory system has a virtual address space of four Gigabytes and a page size of two Kilobytes. (2×5 points)

a) How many bits are used for ***the page number***? 32 – 11 = 21 **bits**

b) How many bits are used for the ***byte offset***? $\log_2 2{,}048 = 11$ **bits**

5. A virtual memory system has 8 Kilobyte pages and 64-bit addresses and uses ***inverted page tables***.

a) Given that each page table entry occupies 24 bytes what is the fraction of main memory occupied by page tables? (5 points)

24/8K = 24/8,192 = a bit less than 3 percent (2.923 percent to be exact)

b) Why are these page tables entries occupying 24 bytes? (5 points)

Each page table entry must contain a page number, a page frame number and the address of the next entry. Since the system uses 64-bit addresses each of these entries occupies 8 bytes.

6. A UNIX Fast File System has 32-bit addresses, 8 Kilobyte blocks and 15 block addresses in each inode. How many file blocks can be accessed: (4×5 points)

a) Directly from the i-node? 12 **blocks**

You may, if you want, detail your computation here:

b) With one level of indirection? 8K/4 =2K **blocks**

You may, if you want, detail your computation here:

c) With two levels of indirection? 4G/8K –2K – 12 = 512K -2K – 12 = 510K -12 **blocks**

You may, if you want, detail your computation here:

(The maximum file size is 4GB!) (I would have accepted 512K as a correct answer.)

d) With three levels of indirection? ZERO **blocks**

You may, if you want, detail your computation here: