NAME: _____ (FIRST NAME FIRST) SCORE: _____

COSC 4330 FINAL EXAMINATION DECEMBER 16, 2009

This exam is **closed book**. You can have **one** page of notes. UH expels cheaters.

1. The block size of a 32-bit UNIX file system with 15 block addresses is 8 kilobytes. How many file blocks can be accessed:

a) directly from the i-node: 12 blocks
(5 points)

b) with one level of indirection: 8K/4 = 2,048 blocks
(5 points)

c) with two levels of indirection: 4G/8K – 2,048 – 12 =512K – 2,060 blocks
(5 points)

You may use this space to write down your computations for potential partial credit. .

Do not forget that the maximum file size of a 32-bit file system is 2^{32} bytes = 4GB.

2. A virtual memory system has 32-bit addresses. Given that 11 of these 32 bits are used by the bit offset,

a) What is the ***page size*** of the system**?** (5 points) 2^{11} =2.048 bytes

b) What is the maximum size of a process virtual address space? (5 points) 2^{32} bytes = 4 GB

c) What is the maximum size of a process page table? (5 points) $4x2^{21}$ bytes = 8 MB

You may use this space to write down your computations for potential partial credit.

If the bit offset occupies 11 bits, the page number uses the 32 – 11 = 21 remaining bits of the address.

3. What are the main advantage and the main disadvantage of ***access control lists*** over ***tickets***? (2×5 points)

Access control lists are more flexible than tickets and let the owner of the file revoke the access rights of individual users. Tickets only let the owner of the file revoke the access rights of all the users having a specific ticket.

On the other hand consulting an access control list verifying a ticket takes much more time than verifying a ticket.

4. What does an FFS ***cylinder group*** contain**?** (2×5 easy points) How do these cylinder groups improve the performance of the file system? (5 points)

Each FFS cylinder group has space allocated to (a) i-nodes and (b) data blocks.

Cylinder groups minimize seek times by grouping together the i-node associated with a given file and the data blocks of that file.

5. Describe in some detail a possible hashed page table organization for a virtual memory system with 64-bit addresses. (10 points for a short explanation with a good drawing)

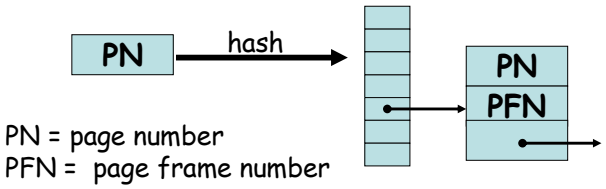

Unlike other page table organizations, hashed page tables only contain entries for the pages that a re currently in main memory. The page number is transformed by a hashing function into a bucket index. Each bucket contains a linked list containing the page table entries of all pages that hash into the same bucket. These page table entries contain the page number, its associated page frame number, various bits and the address of the next entry in the list.

Recalling that each address occupies 8 bytes, what would be the size of a page table entry? (5 points)

Each page table entry would occupy 3x8 = 24 bytes. _____

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

a) What is the main purpose of the UNIX `mount` primitive?

To make the directory tree of the mounted disk partition a sub-tree of the disk partition on which it was mounted.

b) What is the major disadvantage of using ***large block***s in a file system?

Excessive internal fragmentation.

c) What is the major disadvantage of letting the kernel handle ***TLB misses***?

TLB misses will occasion two context switches.

d) What is the main advantage of ***journaling file systems*** using ***synchronous log writes***?

No metadata update will be lost when the system crashes.

e) Why is the ***Windows page replacement policy*** better suited to the needs of ***real-time processes*** than the Clock policy or the Mach policy?

The windows page replacement policy allocates a minimum number of page frames to each process.

f) When a user creates a ***new file***, in which order should the file system write to disk (a) the ***directory block*** containing the new entry for the file and (b) the ***i-node block*** containing the new i-node?

We should create the i-node block to disk before the directory entry pointing to it.