

CS4410 Homework 4

Due 9am Thursday March 12

100 points as shown below. Very short answers, please.

- 1. (10 points).** What do we mean when we say that a virtual memory system is thrashing ?
- 2. (10 points).** Do threads in the same process share page table entries? What about TLB entries?
- 3. (10 points).** Assuming that a process is launched with no pages in memory. On a virtual memory system, what is the minimum number of page faults that it could occur before it terminates?
- 4. (10 points).** Suppose that an operating system supports very large virtual address spaces composed of paged variable-length segments that may be widely separated by gaps. Why would a traditional page table get large in this case?
- 5. (10 points).** Explain the idea of an inverted page table. Would it suffer from the same issues that you mentioned in responding to question 4?
- 6. (10 points).** Describe a situation in which “file-block prefetching” can harm performance.
- 7. (10 points).** Suppose that you create a file named `/tmp/xyz.jog` and then create a true link (on Linux) to it named “`my_file.jpg`”. This file now has two “types”. What will happen if you try and open `my_file`? Assume that the file is really a JPEG (a photo) and that “.jog” is a typo, not a file system extension for any existing application.
- 8. (10 points).** Same scenario as 7. You realize that the original file name was just a mistake, so you rename `/tmp/xyz.jog` as `/tmp/xyz.jpg`. But then you edit the photo, and the photo editing program deletes `/tmp/xyz.jpg` and then writes a new file, which it renames `/tmp/xyz.jpg`. What would you see if you print the photo under its other name: `my_file.jpg`?
- 9. (10 points).** Same scenario as 8, but now assume that the link is a symbolic link.
- 10. (10 points).** You measure the file access performance of your new computer and discover that reading a random byte from a random file seems to have more than one “cost”. In fact, the values cluster: there is one set of files for which this cost is usually 11ms. A second set of files cost more to access: 18ms. A third set of files cost even more to access: 25ms. Explain why this sort of “three costs model” isn’t surprising. What accounts for the extra costs in the second and third group?