

CS 4410: Operating Systems

Homework 8

- Homework may be done in pairs, or individually. If doing in pairs, **one** of you should upload to gradescope and add your partner to the group assignment in the upper right corner of the screen. (Do **not** just upload the assignment twice or it will be graded twice, which means grading will take longer.)
- The deadline is Tue, Nov 29 at [11:59AM](#).
- No late submissions will be accepted.
- [You must attribute every source used to complete this homework.](#)
- For some of the problems, you will need two integers. Here is the algorithm for computing these integers:
 - If you are working with a partner, let `var` be the lexicographically smaller of the two NetIDs.
 - Let `varInt` be the integral part of `var`. That is, if `var = rst12`, then `varInt = 12`.
 - If `varInt` is a single digit integer, let `varInt = 13 × varInt`
 - Let `Int1` be the first digit of `varInt`
 - Let `Int2` be the second digit of `varInt`
- Assume the storage unit convention: $1G = 2^{10} \times M = 2^{20} \times K = 2^{30}$ (bytes).
- **For all problems that use `Int1` or `Int2`, please write down the parameters (related variables and settings calculated from your NetID) before answering each question.**

1 Harold's Financial Files

Harold hides his confidential financial files by burying them inside a series of nested directories. It's tax season and Harold plans to access these files very frequently, so he asks you to calculate the performance of accessing a file with a very long path name.

For security reasons, Harold can't tell you the path name. He only tells you that there are $(\text{Int } 1 \bmod 3) + 3$ slashes ("/") in the path name of a UNIX FFS file system. Assume there is no file cache. Furthermore, assume that when the OS looks up a particular file in a directory, it knows exactly which data block of the directory it needs to access in order to find that specific entry. For example, if the file were listed in the second data block of the directory, the OS would not need to access the first data block.

1. What is the minimum number of disk blocks that must be read to fetch the first block of the file specified by the path name?
2. What is the maximum number of disk blocks that must be read to fetch the first block of the file specified by the path name?

(Hint: think about what happens when a directory has many entries.)

2 Fat Fast File System

VWware uses a variant of Fast File System (FFS) of UNIX, the Fat Fast File System (FFFS), in order to store large files. For each inode, there are $(12 + (\text{Int } 1 \bmod 3))$ direct, 2 indirect, 2 double indirect, 2 triple indirect, and 1 quadruple indirect pointers. Suppose the block size is $(8 + (\text{Int } 2 \bmod 4))K$, each pointer is 8 bytes, and each directory entry is 16 bytes.

1. What is the largest file that can be indexed using only the direct pointers?
2. To within 1% ¹, what is the maximum file size supported by FFFS?

¹If the the maximum file size were 10,097,248 bytes it would be acceptable to answer 10,000,000 bytes.

3 Wasted Space

Harold is not satisfied by FFFS described above because of internal fragmentation. He asks you to analyze FFFS fragmentation using the following case study: Consider an FFFS with only the following files and all the induced intermediate directories:

Path Name	File Size
/animal/felidae/tom.jpg	12K
/animal/rodent/jerry.jpg	10K
/animal/rodent/fievel.jpg	20K
/animal/primate/hominidae/human.jpg	30K
/animal/primate/code-monkey/hacker/ted.jpg	40K
/animal/primate/code-monkey/unknown/harold.jpg	100K

1. What is the internal fragmentation (in bytes/K/M/etc.) of the data blocks (just directories and files) on the disk?
2. Double the block size. Now what is the internal fragmentation?