

CSE 60641: Operating Systems

Home work assignment #4

Open book, open notes, individual effort. Google searches are okay - always cite your source.
All questions carry equal points. Brevity is encouraged,

1. "The trace driven analysis of the 4.2 BSD file system" paper made the following observations: a) 70% of all files are whole-file transfers, b) 75% of all files are open less than 0.5 seconds. c) about 20-30% of new written information is deleted within 30 seconds. For this question, assume that these observations were still true. Now, consider a modern disk that reorders read/write requests internally. Suppose the OS knows that the disk supports 1GB of cache but does not know how the disk manages this disk.
 - a) How would you change the design of the file system in order to use this disk cache to your advantage (and get good performance?) Clearly describe any assumptions about the disk behavior.
 - b) Traditional disks support operations such as a) read(block), b) write(block). Suppose you are allowed to introduce one new command into the disk controller, what would that command be such that you improve the performance of your file system?
2. (Duke Finals '99): In an effort to "improve" on the UNIX inode design, programmers of the Cluelix OS came up with the following structure:
 - Directory entries contain a pointer to the first block of the file. The block size is 4K bytes.
 - The first part of that block contains the information normally in the inode, with the following exceptions:
 - There is no need for the first direct block pointer, since the first (partial) block of data is here (suppose that leaves 2K bytes of data).
 - Instead of 12 direct block pointers, there are only 4.
 - Following the direct block pointers, there are 4 single indirect block pointers. Assume an index block can hold 1000 disk addresses.
 - Then come 4 double indirect block pointers.
 - Finally, there are 2 triple indirect block pointers.
 - a) For what ranges of file sizes is this structure better and why? (Note there may be more than one meaning for "better" that you might consider).
 - b) For what ranges of file sizes is the original UNIX structure better and why?
 - c) For what sizes does it make no difference?