

CS323 – Exercises
Week 9
16 May 2019

Exercise 1: (File system on-disk data structures)

What is the size in bytes of the biggest file that can be represented with an inode that contains PD direct pointers (disk addresses), PI pointers to an indirect block, and PDI pointers to double-indirect blocks? Assume that a disk block is S bytes and P pointers (disk addresses) fit into a disk block.

Exercise 2: (File System Data Structures)

Consider the following file system very much like the one discussed in class. It maintains on disk a device directory consisting of num_inodes inodes at a fixed disk address inode_start. Each inode contains a reference count, a length in bytes, and a number PD of disk addresses for the sectors with data belonging to that inode. For simplicity, assume that there are no indirect or double-indirect entries in the inode. The bitmap of free disk sectors is kept only in memory (i.e., not on disk). After a machine crash that information is obviously lost. Write pseudocode for a program that rebuilds the free list on crash recovery.

Exercise 3: (RAID)

Consider a 4-disks, 200GB-per-disks RAID array. What is the available data storage capacity for each of the RAID levels 0, 1, 3, 4, and 5?

Exercise 4: (RAID)

RAID level 3 stripes data at byte-level and is able to correct single-bit errors using only one parity drive. In contrast, RAID level 2 stripes data at bit-level, and uses Hamming error-correction codes (ECCs) stored across at least 2 additional drives.

What is the point of RAID level 2? After all, it also can only correct one error and takes more drives to do so.