CS 537: Intro to Operating Systems (Fall 2017)

Worksheet 13 - RAID and FFS

Due: Dec 13th 2017 (Wed) in-class OR email to Simmi before 11:59 pm

NOTE: No late day(s) allowed for this worksheet

1. In this question, we'll examine how long it takes to perform a small workload consisting of 12 writes to random locations within a RAID. Assume that these random writes are spread "evenly" across the disks of the RAID. To begin with, assume a simple disk model where each read or write takes D time units

   a. Assume we have a 4-disk RAID-0 (striping). How long does it take to complete the 12 writes?

   b. How long on a 4-disk RAID-1 (mirroring)?

   c. How long on a 4-disk RAID-4 (parity)?

   d. How long on a 4-disk RAID-5 (rotated parity)?

   e. Now assume we have a better disk model, in which it takes S time units to perform a random seek and R units of time to perform a full rotation; assume transfer is free. How long do the 12 random writes take to complete on a 4-disk RAID-0?

   f. How long on a 4-disk RAID-1 (mirroring)?

   g. How long on a 4-disk RAID-4 (parity)?

   h. How long on a 4-disk RAID-5 (rotated parity)?

2. This question is about the Berkeley Fast File System, FFS. Remember, the first file system that realized it was built on a mechanical disk? The question focuses on the flow of data to disk, i.e., where does FFS place files?

   Assume we have **10 cylinder groups** (or block groups) on disk. As you might recall, each group has four basic components: an inode bitmap (Bi), a data bitmap (Bd), some inodes (I), and some data blocks (D).

   This question focuses on which data structures are touched when certain operations occur. For example, if an application reads a 1-block file in the second block group, the flow would be: "2I, 2D", because to read that file, the application first has to read its inode and then its data block, both of which are in the second block group.

   Assume for these questions that all files and directories are 1 block in size unless otherwise specified. Also assume the root directory is found in block group 1.

   a. There is an existing file "/foo" in block group 3. An application wishes to read it from disk. What is the flow of reads from disk needed to satisfy this request? (Remember to start at the root)

   b. Now an empty file "/bar" is created in the same directory as foo (i.e., the root directory). What flow describes the writes to disk needed to create "/bar"? (Note: the exact ordering does not matter in this case.)

   c. Now assume that someone runs "ls" on the root directory and wants to obtain the size of each file or directory within. What is the flow of reads from disk?

   d. Finally, assume another file is getting created in the root directory, "/big". Assume it is a really big file. As it keeps growing and growing and eventually nearly filling the disk, what flow might describe the its sequence of writes to disk?