CSE451 File System Introduction and Disk Drivers
Spring 2001

Gary Kimura
Lecture #19
May 7, 2001

Today's Topics

- Where the does the file system and disk driver fit in the scheme of things
- General features and semantics presented in a disk driver (next we'll cover file systems)
- Implementation details

Your Job this week

- No homework this week
- Study for the midterm on Friday
  - Everything up through Memory Management
- Readings in Silberschatz
  - Chapter 13 and 11
- Work on your project

Secondary Storage

- Secondary Storage is usually:
  - anything outside of "primary memory"
  - storage that does not permit direct instruction execution or data fetch by load/store instructions
  - it's large
  - it's cheap
  - it's non-volatile
  - it's slow

The Big Picture

- One goal of the file system is to add useful structure and semantics to an otherwise blah disk.
- The major OS components in this area are:
  - The Disk Driver
  - The File System
  - The I/O System
  - The Software Cache Management System

The Disk Driver

- Abstractly the disk driver's task is to present to the rest of the operating systems a consistent physical media for storing and retrieving data.
  - It presents consecutively numbered sectors (Logical Blocks) to the rest of the OS
  - The numbering is referred to as the Logical Block Number (LBN)
- The Disk driver must deal with the mess of physical devices:
  - Various sizes and speeds
  - Errors
  - Bad blocks
  - Missed seeks
- The disk driver hides this mess

A Rather Old Physical Disk Structure

- Disk are made of metal platters with a read/write head flying over it.
- To read from a disk, we must specify:
  - cylinder #
  - surface #
  - sector #
  - transfer size
  - memory address
- Transfer time includes: seek, latency, and transfer time

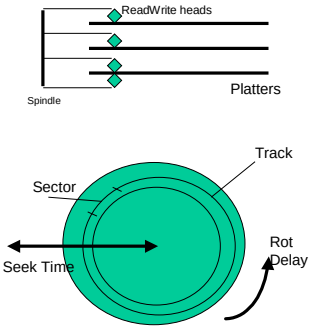

Some Typical Numbers

- Sector size: 512 bytes
- Cylinders per disk (tracks per platter): 6962†
- Platters: 3 - 12†
- Rotational speed: 10000 RPM
- Storage size: 4 - 18 GB
- Seek time: 5 - 12 ms
- Latency: 3 ms
- Transfer rate: 14 - 20 MB/sec

†Modern hardware often lies about its dimensional size

Disk Scheduling

- Because disks are slow and seeks are long and depend on distance, we can schedule disk accesses, e.g.:
  - FCFS (do nothing)
    - ok when load is low
    - long waiting times for long request queue
  - SSTF (shortest seek time first)
    - always minimize arm movement. maximize throughput.
    - favors middle blocks
  - SCAN (elevator) -- continue in same direction until done, then reverse direction and service in that order
  - C-SCAN -- like scan, but go back to 0 at end
- In general, unless there are request queues, it doesn't matter
- The File System may locate files strategically for performance reasons.

Bad Blocks

- With increasing densities, all disks have some bad blocks, and some go bad as time goes on.
- The File System can remove that block from its allocation map.
- On some disks, each cylinder contains a set of replacement blocks that the device can remap to replace other "logical" blocks on the same cylinder that are bad.

Other Disk Issues

- As disk technology has grown modern hardware system have changed some of the changes are
  - Automatically Bad Block revectoring
  - Constant density means more sectors on the outside cylinders which implies a performance win if we can mostly utilize the outside cylinders.
  - Hardware caching means that the seek algorithm is now out of the software systems direct control.
- Multi-volume sets is often handled by the disk driver
  - Stripping is then possible
  - RAID volumes are included in this scenario (another lecture)

Still to come

- File system details
- How to impose a structure to the disk
- How to design the file system driver
- Software caching
- etc.