

CSCE 351

Operating System Kernels

Hard and Floppy Disks


Steve Goddard
goddard@cse.unl.edu

<http://www.cse.unl.edu/~goddard/Courses/CSCE351>

1

Anatomy of a Disk

Basic components




2

Anatomy of a Disk

Example: Seagate 9GB Fast/Wide/Differential SCSI disk

◆ Specs:

- » 12 platters
 - » 22 heads
 - » variable # of sectors/track
 - » 7,200 RPM
 - ❖ average latency: 4.2 ms.
 - » Seek times
 - ❖ track-to-track: 1 ms
 - ❖ average: 7.9 ms
 - » 40MB/s peak transfer rate
- » 11 arms
 - » 4,735 tracks
 - » 512 bytes/sector




3

Disk Operations

Data transfer in units of sectors


Random access devices with non-uniform access times

- ◆ Present disk with a sector address
 - » $DA = (\text{drive, surface, track, sector})$
 - » May use Logical Block Addressing (LBA). $DA = \text{absolute sector}$
- ◆ Head moved to appropriate track
 - » “seek time”
- ◆ The appropriate head is enabled
- ◆ Wait for the sector to appear under the head
 - » “rotational latency”
- ◆ Read/write the sector
 - » “transfer time”

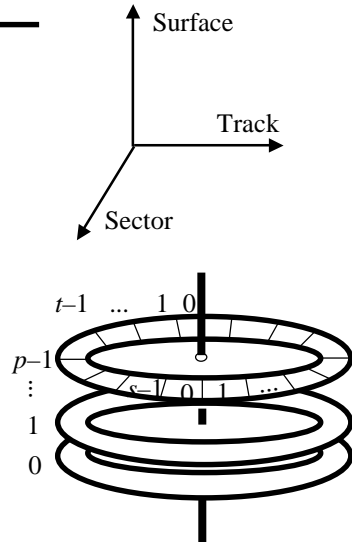


4

Device Driver Block to Sector Mappings




- ◆ Device driver translates block requests into cylinder, track, and sector requests.



5

Disk Head Scheduling Maximizing disk throughput

- ◆ In a multiprogramming/timesharing environment, disk I/O requests are queued up




- ◆ The OS maximizes disk I/O throughput by minimizing head movement through *disk head scheduling*

6

Disk Head Scheduling

Examples

- ◆ Assume a queue of requests exists to read/write tracks:
 - » 150 16 147 14 72 83 and the head is on track 65




- ◆ FCFS scheduling results in the head moving 550 tracks
 - » Can we do better?

7

Disk Head Scheduling

Minimizing head movement

- ◆ Greedy scheduling: *shortest seek time first*
 - » Rearrange queue from: 150 16 147 14 72 83
 - To: 72 83 147 150 16 14




- ◆ SSTF results in the head moving 221 tracks
 - » Can we do better?

8

Disk Head Scheduling

Optimal scheduling

- ◆ Rearrange queue from: **150 16 147 14 72 83**
To: **16 14 72 83 147 150**




- ◆ SCAN scheduling
 - » Move the head in one direction until all requests have been serviced and then reverse
 - » Results in the head moving 187 tracks

9

Disk Head Scheduling

Other variations

- ◆ C-SCAN scheduling (“Circular”-SCAN)
 - » Move the head in one direction until an edge of the disk is reached and then reset to the opposite edge



- ◆ LOOK scheduling
 - » C-SCAN except the head is reset when no more requests exist between the current head position and the approaching edge of the disk


10

Speeding Up Disk I/O

Disk architectures

- ◆ Disk striping

- » Blocks broken into sub-blocks that are stored on separate disks
 - ❖ similar to memory inter-leaving
- » Provides for higher disk bandwidth through a larger effective block size




11

Speeding Up Disk I/O

Disk architectures

- ◆ RAID (*redundant array of inexpensive disks*) disks
 - » Bit-wise striping of the disks (RAID-3) or
 - » Block-wise striping of the disks (RAID-5)
 - » Provides better performance & reliability
- ◆ Example: storing the bit-string 101

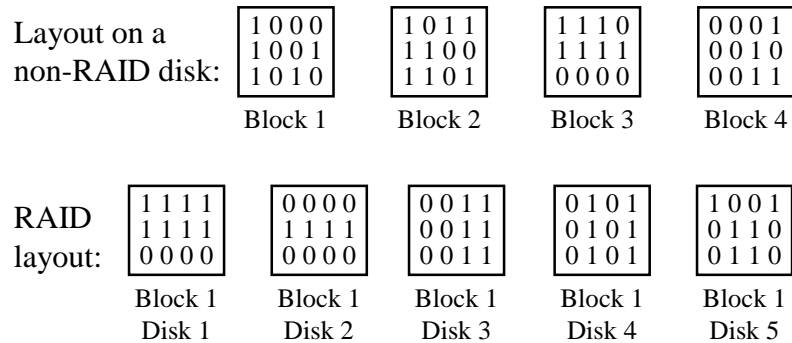


12

RAID Disks

Improving reliability & availability

- ◆ Block interleaved parity striping
 - » Allows one to recover from the crash of any one disk
 - » Example: storing 8, 9, 10, 11, 12, 13, 14, 15, 0, 1, 2, 3



13

Sources of Disk Access Error

- ◆ Programming error
 - » Return error to caller
- ◆ Transient checksum error
 - » Read again
- ◆ Permanent checksum error
 - » Block marked as **Bad Block**, return error
- ◆ Seek error
 - » Mechanical error
 - » Corrected by controller on HD
 - » Send RECALIBRATE command for FD
- ◆ Controller error
 - » Reset
 - » If all else fails, print a message and give up (panic!)

14

Track-at-a-Time Caching

- ◆ When a sector is read, the disk driver may read all sectors in the track and store the data in a buffer
- ◆ Advantages?
- ◆ Disadvantages?
- ◆ Some HD controllers do track-at-a-time caching in their own internal memory.
 - » Advantages?
 - » Disadvantages?
- ◆ Many HD now have caches in the drive as well

15

Floppy Disk Driver

- ◆ More complicated than HD driver
 - » Why?
- ◆ Strictly FCFS scheduling in MINIX
- ◆ The SEEK operation must be done by the driver
 - » When it fails (SEEK error), the driver sends a RECALIBRATE command
- ◆ Other problems:
 - » Removable media
 - » Multiple disk formats
 - » Motor control

16