Lecture 16: Survey of I/O systems

16.0 Main Points

How terminals and disks work.

Overhead, latency, bandwidth definitions

Concept of batching

16.1 Definitions

Overhead: CPU time to initiate operation (can't be overlapped)
Latency: time to initiate 1 byte operation
Bandwidth: rate of I/O transfer, once initiated

Rule of thumb:

programming abstraction: byte transfers

batch into blocks for efficiency -- pro-rate overhead, latency over larger unit.

16.2 Terminal: keyboard and display

Keyboard connects to computer via serial line; type characters, get characters back to display.

For example, RS-232 is bit serial: start bit, character code, stop bit. Typical bandwidth: 9600 baud (transmission rate in bits/sec) = 900 bytes/sec.

Same principle for modems, networks, at faster rate. For example, fast Ethernet = 100 Mb/s Ethernet.

What if interrupt per byte?

Suppose:

10 users (vi/emacs) (or 10 people using a modem)
900 interrupts/sec per user
overhead of handling interrupt is 100 usec

implies:

devote whole computer to simply handling interrupts!

Alternative: **Direct memory access** (DMA). Instead of interrupt on every byte, block transfer, interrupt CPU when block is done. Allows much higher transfer rates.

Another alternative: raster bitmap display, formed out of NxN bit values.

Raster displays are how TV sets, laser printers, and monitors work.

16.3 Disk Organization

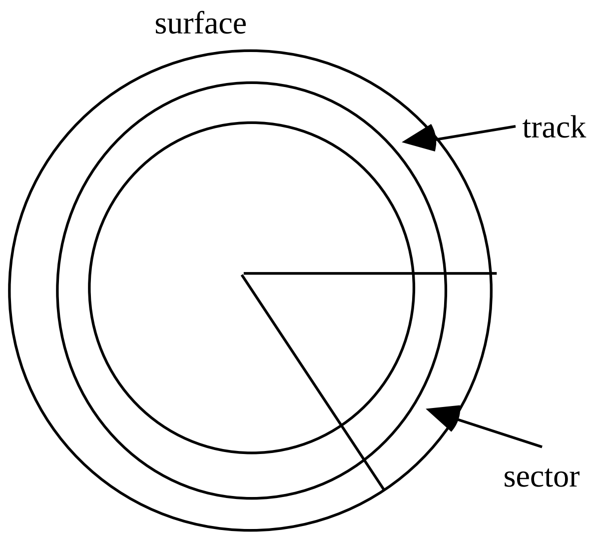

Disk **surface**: circular disk, coated with magnetic material

Tracks:concentric rings around disk surface, bits laid out serially along each track

Each track is split up into **sectors**: arc of track, also minimum unit of transfer

CD's come individually, but magnetic disks come organized in a **disk pack**, by ganging a set of **platters** together in a stack (each platter is 2 surfaces)

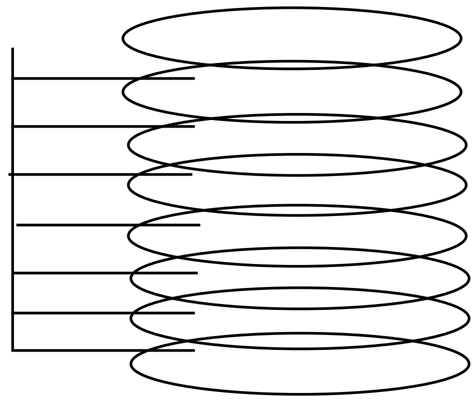

Disk is read via a comb -- 2 read/write "head assembly" at end of each arm (typically, only put disk heads to read inside surfaces, external surfaces are too exposed).

Cylinder: corresponding track on each surface

Disk operation is in terms of radial coordinates, not x,y,z: move arm to correct track, wait for disk to rotate under head, select head, then transfer as it's going by.

16.4 Disk Performance

To read or write disk block:

1. **seek**: position heads over cylinder

Time for seek depends on how fast you can move the arm. Typically, 10 msec to move all the way across disk

2. **rotational delay**: wait for sector to rotate underneath head: typically,120 cycles per second => 8 msec per rotation

3. select head (~ 0.5ms)

4. **transfer** bytes (typically, 10MB/s, 1KB/sector => 0.1ms)

Thus, overall time to do disk I/O: seek + rotational delay + select head + transfer

Latency: seek + rotational delay + select head
Bandwidth: transfer rate

What is rate of 1 KB random transfers?

What is rate of 1MB random transfers?

What is rate of 1 KB random transfers, within same cylinder?

What is rate of 1 KB transfers, if each is to next block?

Key to using disk effectively (and therefore to everything in file systems!):

    transfer in large chunks and/or
    place data to minimize seek and rotational delay.

16.5 Disk Tradeoffs

How do we choose disk sector size?

Need to synchronize head with rotation speed, track position:

Put 100-1000 bits in between each sector, to allow system to measure how fast disk is spinning (also allows tolerance to small changes in track length)

What if sector was 1 byte?
What if sector was 1KBytes?

What if sector was 1M bytes?

If 1MB sectors allow more efficient transfers, why not use them?

16.6 Disk Technology Trends

Typical disk today:

Disk capacity	9 GB
# of surfaces per pack	18
# of tracks per surface	8000
# of sectors per track	64
# of bytes per sector	1KB
# of revolutions per min	10000

Transfer rate: # of bytes rotating under head per second

1KB bytes/sector * 64 sectors/cycle * 160 cycles/sec

= 10 MB/sec

1. Disk data getting denser (bits per sq. inch) -- more capacity on same size disk

2. Disk capacity getting cheaper (factor of 2 per year since 1991)

3. Disks getting smaller, for similar capacity
   smaller => spin disk faster (less rotational delay, higher bandwidth
   smaller => less distance for head to travel (faster seeks)
   smaller => lighter weight (for portables)

Nevertheless, disk density ($ per byte) improving much faster than mechanical limitations (seek, rotational delay).

Key to improving density: get head close to surface; precise positioning of disk head

Heads are spring loaded, aerodynamically designed to fly as close to surface as possible (also, lightweight to allow for faster seeks)

What happens if head contacts surface?

Head crash -- scrapes off magnetic information

Problem -- space between head and surface gets so small, no room for air molecules. Instead, fill with special inert gas!