Lecture 2: Secondary Storage Devices

Last day: Introduction

- course outline
- main memory
- secondary storage
- Files

Today: Magnetic Disks

- disk organization
- access time
- example

Zoellick & Folk § 3.1

Magnetic Disks

- Bits of data (0's and 1's) are stored on circular magnetic platters called disks.

- A disk rotates rapidly (& never stops).

- A disk head reads & writes bits of data as they pass under the head.

- Often, several platters are organized into a disk pack.

A Disk Pack

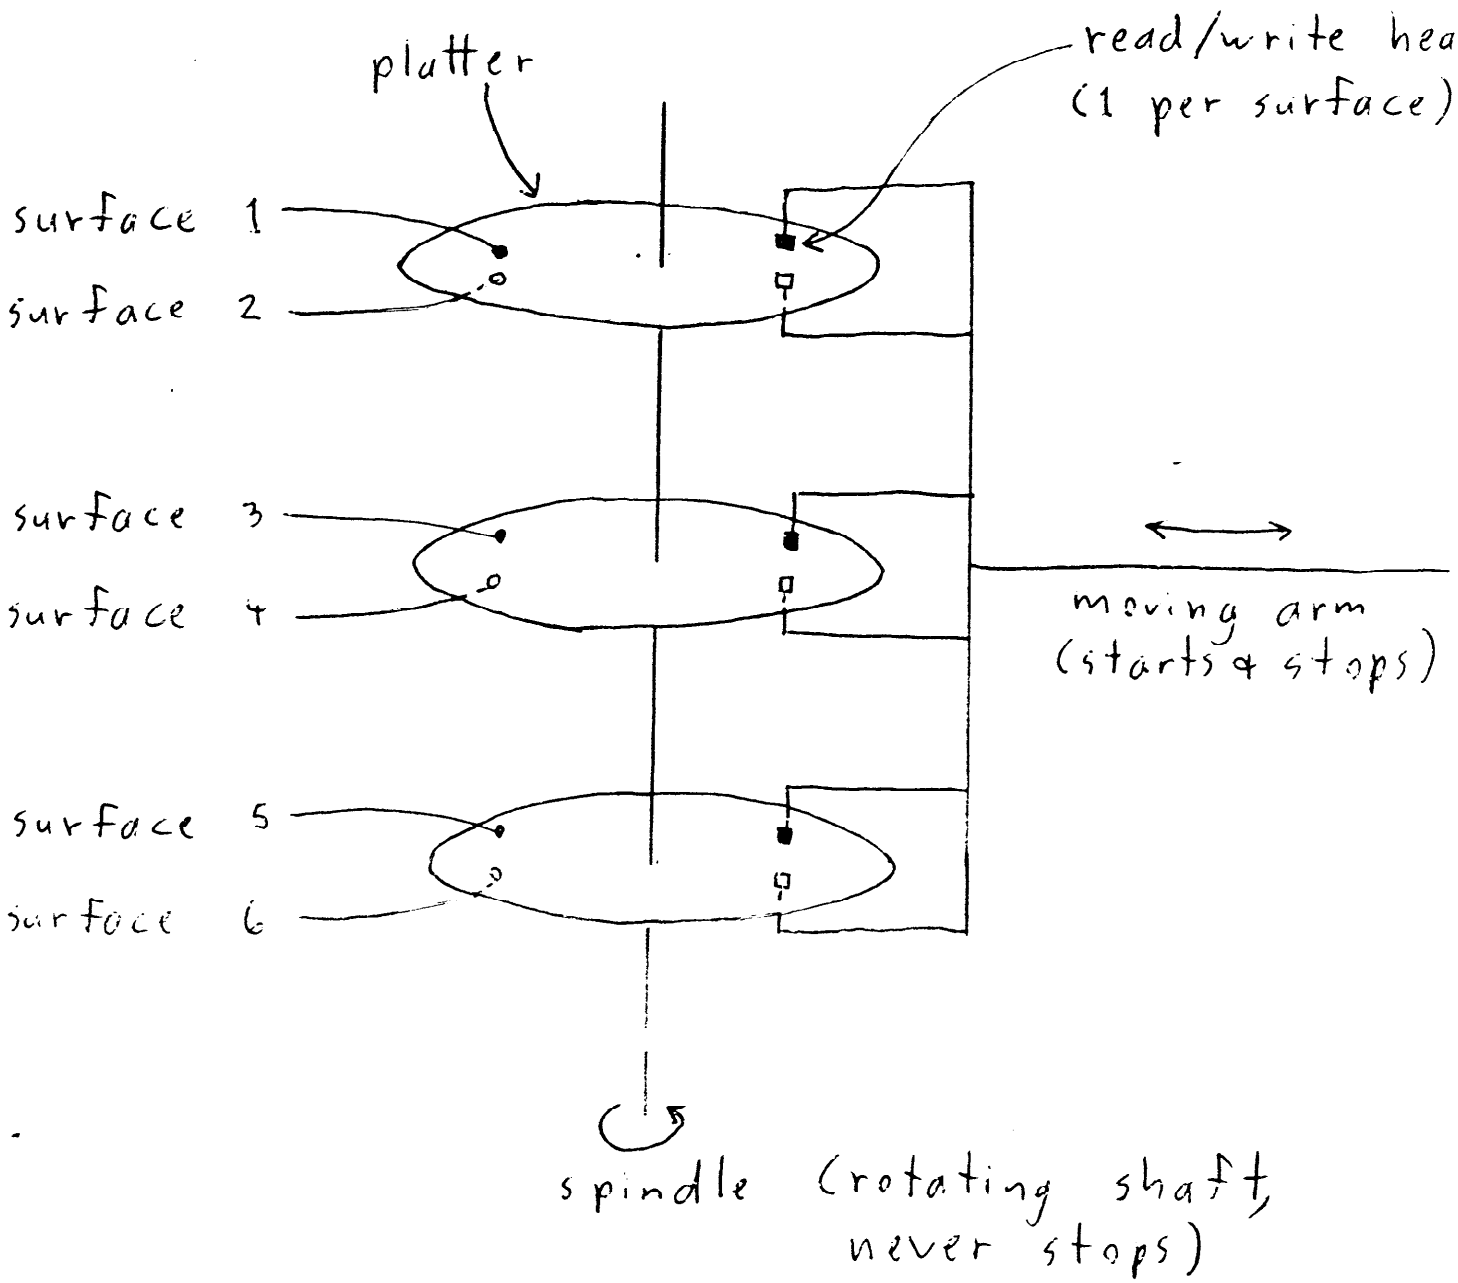

3 platters, 6 surfaces, 6 read/write heads.

Organization of Data on the Surface of a Platter

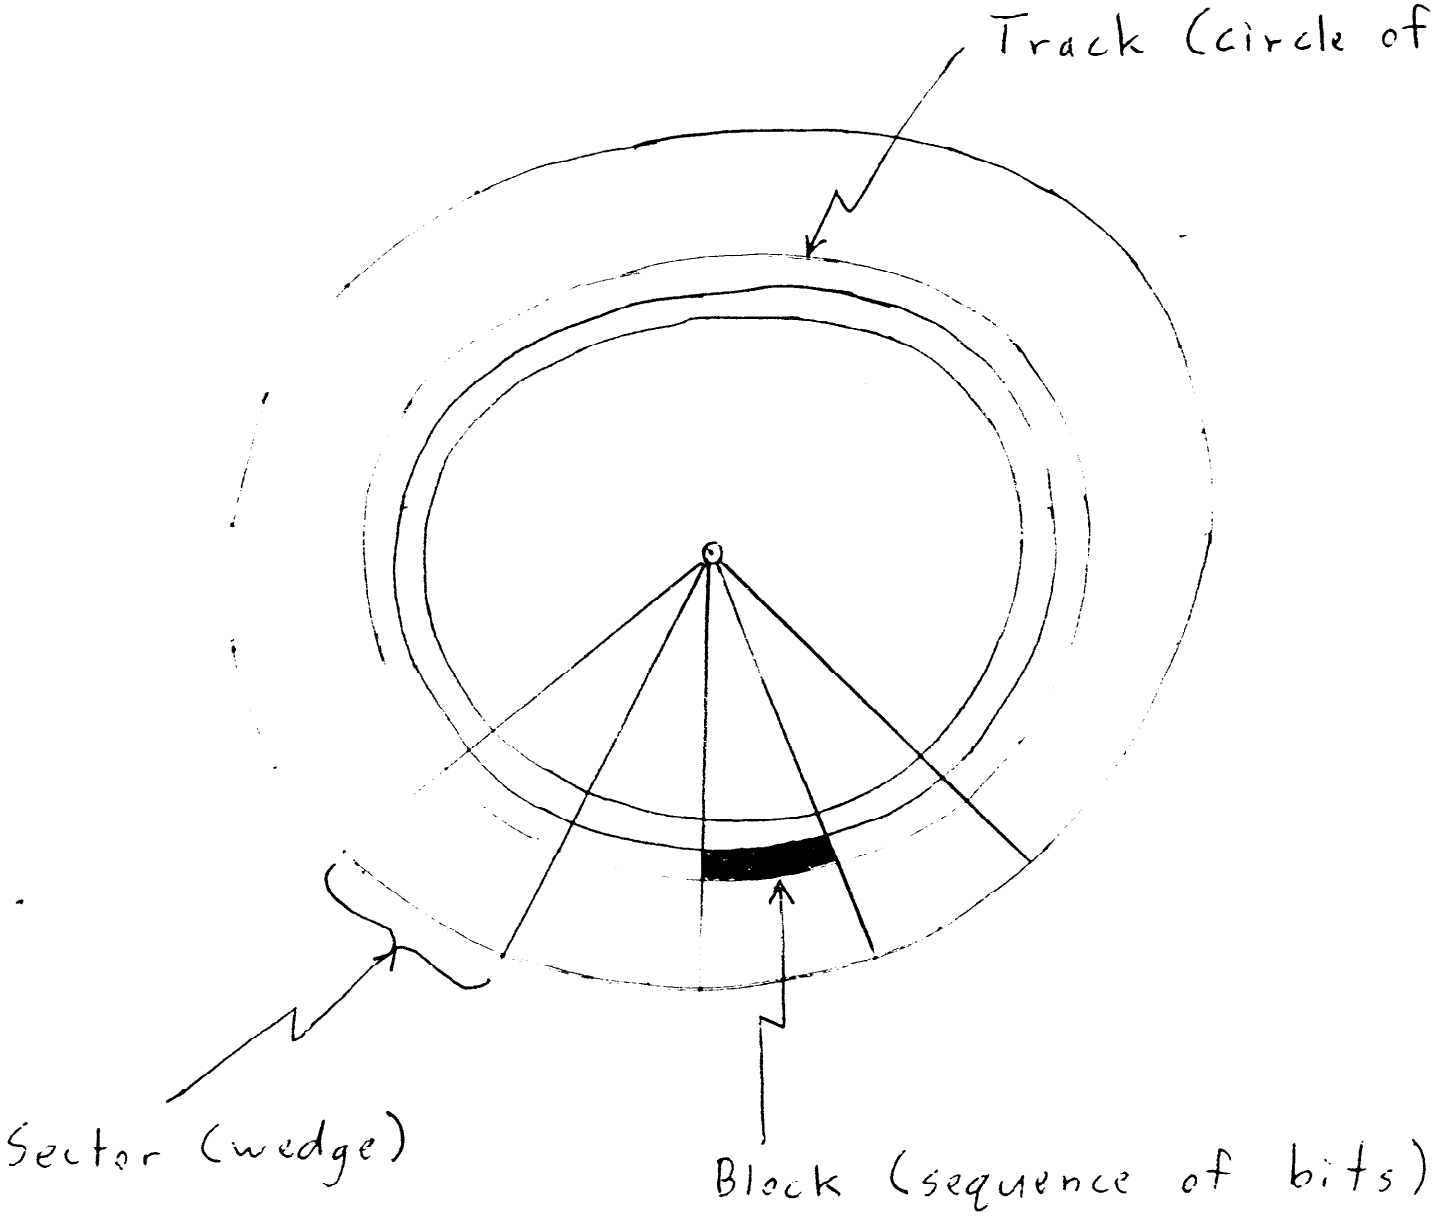

Note: # blocks per track = # sectors.

Accessing Data

- Each block of data is addressed by its track, sector & surface.

- eg. Track 126, Sector 4, Surface 3 identifies a single block of data.

- Given the address of a block, the disk head moves to the given track, & waits for the block to rotate into position under the head.

- It takes a long time to do this, but once in position, the block rotates past the head very quickly.

Consequently...

- It takes a long time to read the first bit in a block, but very little time to read the remaining bits in the block.

- For this reason, a disk reads/writes an entire block of data at a time, not just a single bit or byte.

Typically

- Only one head can read or write at a time
- All heads move together, i.e. they are all at the same radius

Cylinders

- A cylinder is the set of tracks at a given radius of a disk pack.
- ie, a cylinder is the set of tracks that can be accessed without moving the disk arm.
- Moving the disk arm is relatively slow.
- ∴ Once the heads are positioned at a cylinder, data in that cylinder can be accessed more quickly than data in other cylinders.
- ∴ We try to store related data (eg, a file) on the same cylinder.

Access Time

- Given the address of a block, it takes time to access (read or write) it.

- Access time has two main components:

(1) Seek Time: the time needed to move the read/write heads to the correct track.

Typically:

- 5 ms (milliseconds) to move from one track to the next (track-to-track),
- 50 ms maximum (from inside track to outside track).
- 30 ms average (from one random track to another random track).

(2) Latency (rotational delay):
the time needed for the desired bit to rotate into position under the disk head.

Typically: 8ms average.

Note:
- Min latency = 0
- Max latency = Time for one disk revolution.
- Average latency

$$= \frac{\text{min} + \text{max}}{2}$$

$$= \text{max}/2$$

= Time for $\frac{1}{2}$ disk rev.

Example

Given the Following disk:

- 20 surfaces
  25 sectors / surface
  800 tracks / surface
  512 bytes / block

- 3600 rpm (revolutions per minute)

- 7ms track-to-track seek time.
  28 ms avg. seek time.
  50 ms max. seek time.

Latency

3600 rpm $\Rightarrow$ 3600 rev/min.

$\Rightarrow$ 3600/60 rev/sec

$\Rightarrow$ 60 rev/sec

$\Rightarrow$ 1/60 sec/rev

$= .0167$ sec/rev

$= 16.7$ ms/rev

$\Rightarrow$ 16.7 ms max latency

$\Rightarrow \frac{16.7}{2} = 8.35$ ms avg. latency

Disk Capacity

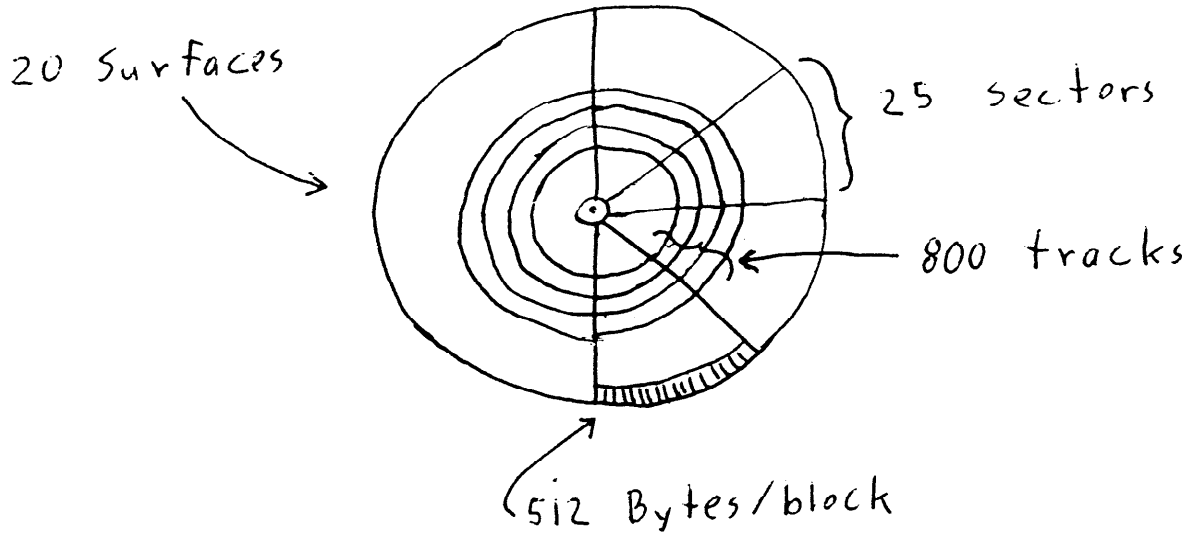

$$\#\text{Bytes/track} = \#\text{bytes/block} \times \#\text{blocks/track}$$

($\#$sectors)

$$= 512 \times 25$$

$$= 12{,}800 \text{ Bytes} = 12.8 \text{ KB (Kilo bytes)}$$

$$\#\text{Bytes/surface} = \#\text{bytes/track} \times \#\text{tracks/surface}$$

$$= 12{,}800 \times 800 = 10{,}240{,}000 \text{ Bytes}$$

$$= 10.24 \text{ MB (Mega Bytes)}$$

$$\#\text{Bytes/pack} = \#\text{bytes/surface} \times \#\text{surfaces/pack}$$

$$= 10.24 \text{ MB} \times 20 = 204.8 \text{ MB}$$

Read Time (cylinder by cylinder)

How long does it take to read the entire disk, one cylinder at a time?

Track read-time = 16.7 ms (1 revolution)

Cylinder read-time
= track read-time × (# surfaces / # tracks/cylinder)
= 16.7 ms × 20 = 334 ms

Total read-time
= (800 cylinder reads / total latency) + (799 cylinder switches / total seek-time)

= (800 × 334 ms) + (799 × 7 ms)

= 267 sec + 5.59 sec = 272.59 sec

Observation

read-time per byte

= 273 sec / 204.8 MB

= 1.33 sec / MB

= 1.33 μs / byte (micro sec / byte)

This is comparable to main-memory speeds!
Reading a disk sequentially is very fast.

In contrast, seek time for a random read of 1 byte is 28 ms, ie, 28,000 μs.

Data processing requires random access.
Usually don't read more than one block (512 bytes) sequentially at a time.
ie, Read a block, then while processing it,