Chap. 6 : External Memory

- Types of External Memory
  - —Magnetic Disk*
  - —Optical
    - CD-ROM
    - CD-Recordable (CD-R)
    - CD-R/W
    - DVD
  - —Magnetic Tape

Magnetic Disk

- Disk: circular platter of nonmagnetic material (substrate) coated with magnetizable material (iron)
- Substrate used to be aluminium
- Now glass or ceramic
  - —Improved surface uniformity
    - Increases reliability
  - —Reduction in surface defects
    - Reduced read/write errors
  - —Lower flight heights
  - —Better stiffness
  - —Better shock/damage resistance

Multiple Platters

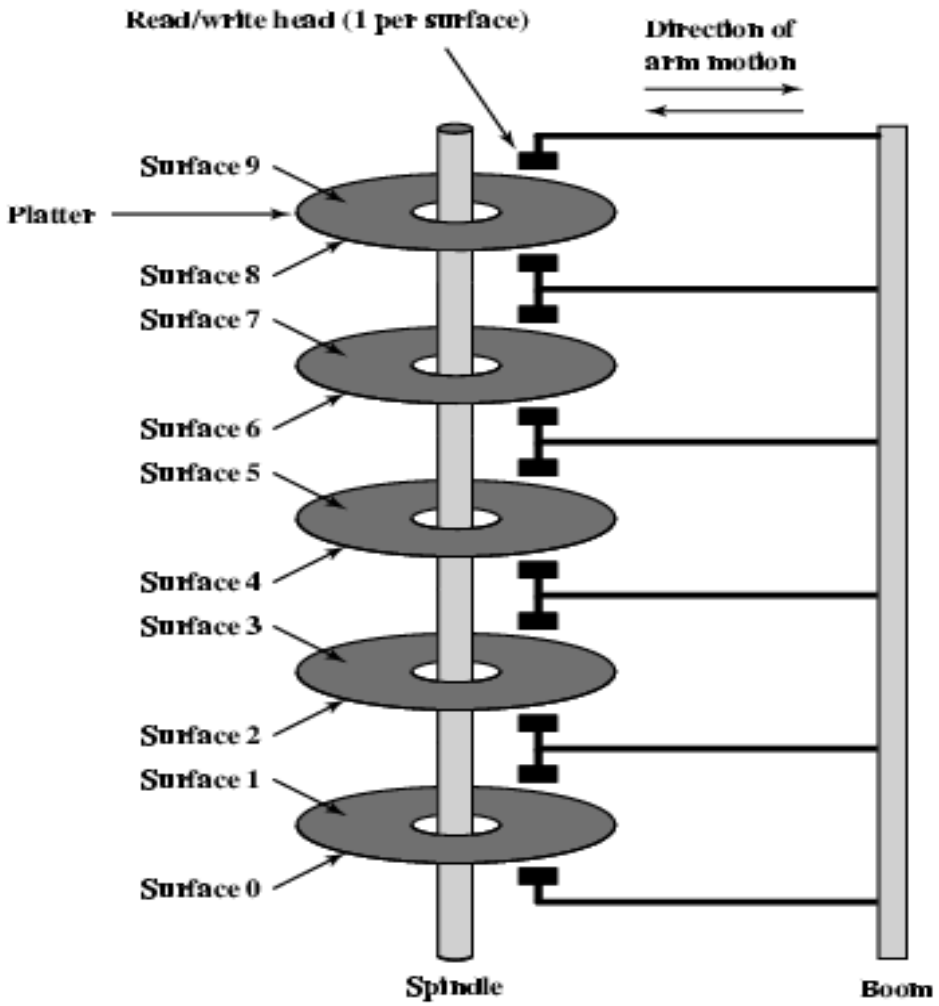

Inductive Write Read Head

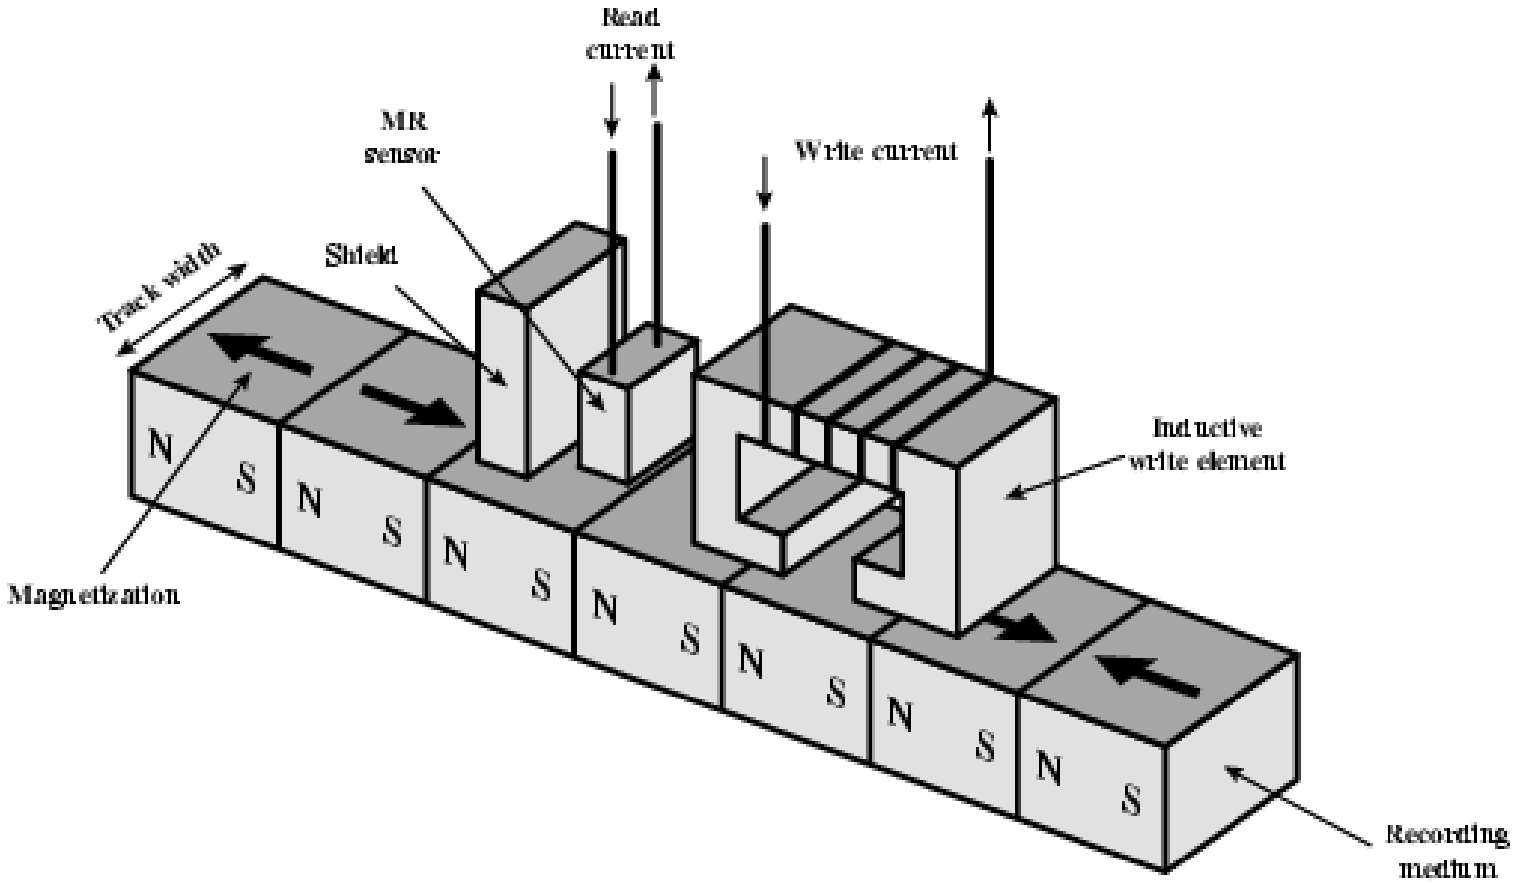

Read and Write Mechanisms

- Recording and retrieval via conductive coil called a head
- May be single read/write head or separate ones
- During read/write, head is stationary, platter rotates
- Write
  - Current through coil produces magnetic field
  - Pulses sent to head
  - Magnetic pattern recorded on surface below
- Read (traditional)
  - Magnetic field moving relative to coil produces current
  - Coil is the same for read and write
- Read (contemporary)
  - Separate read head, close to write head
  - Partially shielded magneto resistive (MR) sensor
  - Electrical resistance depends on direction of magnetic field
  - High frequency operation
    - Higher storage density and speed

Cylinders

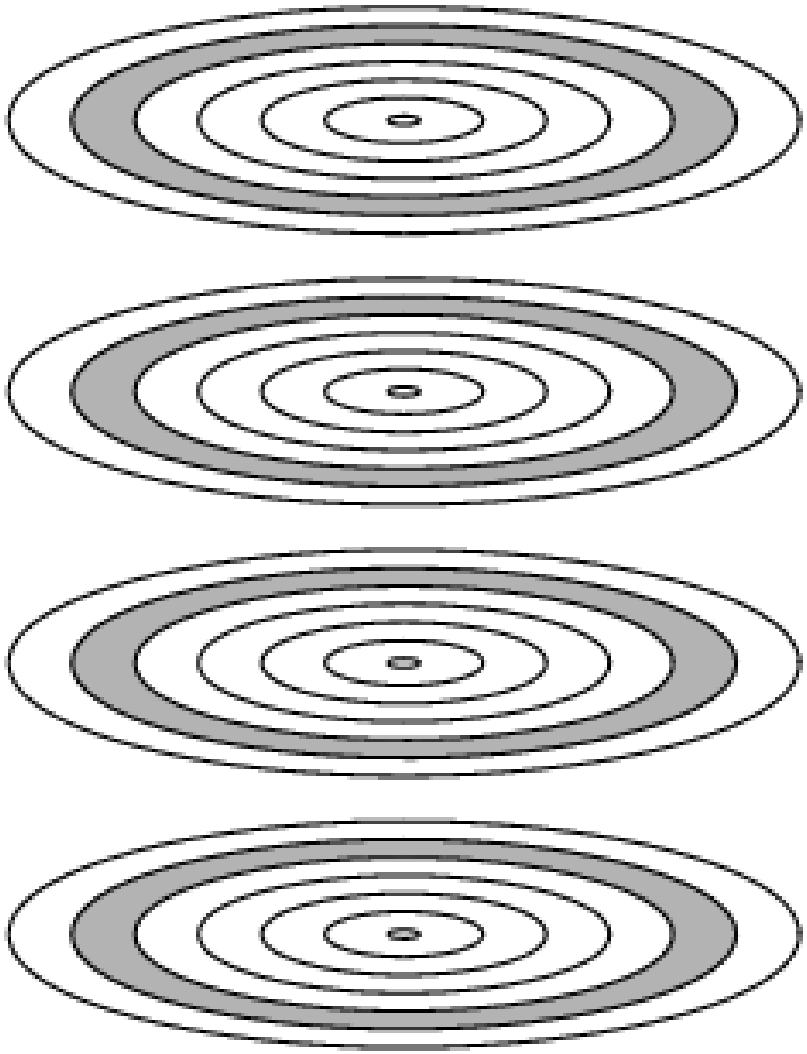

Disk Storage

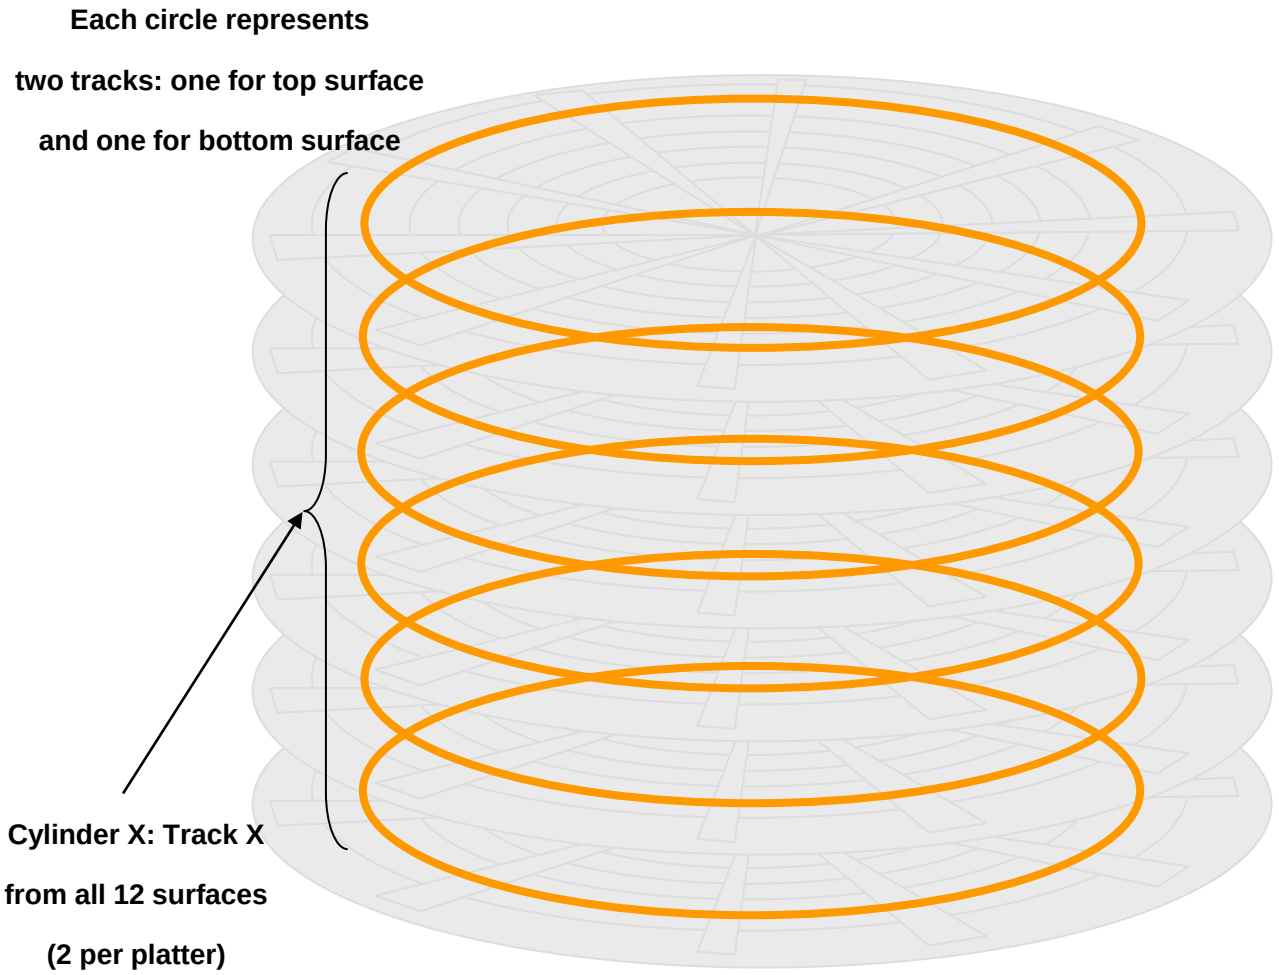

Disk Storage

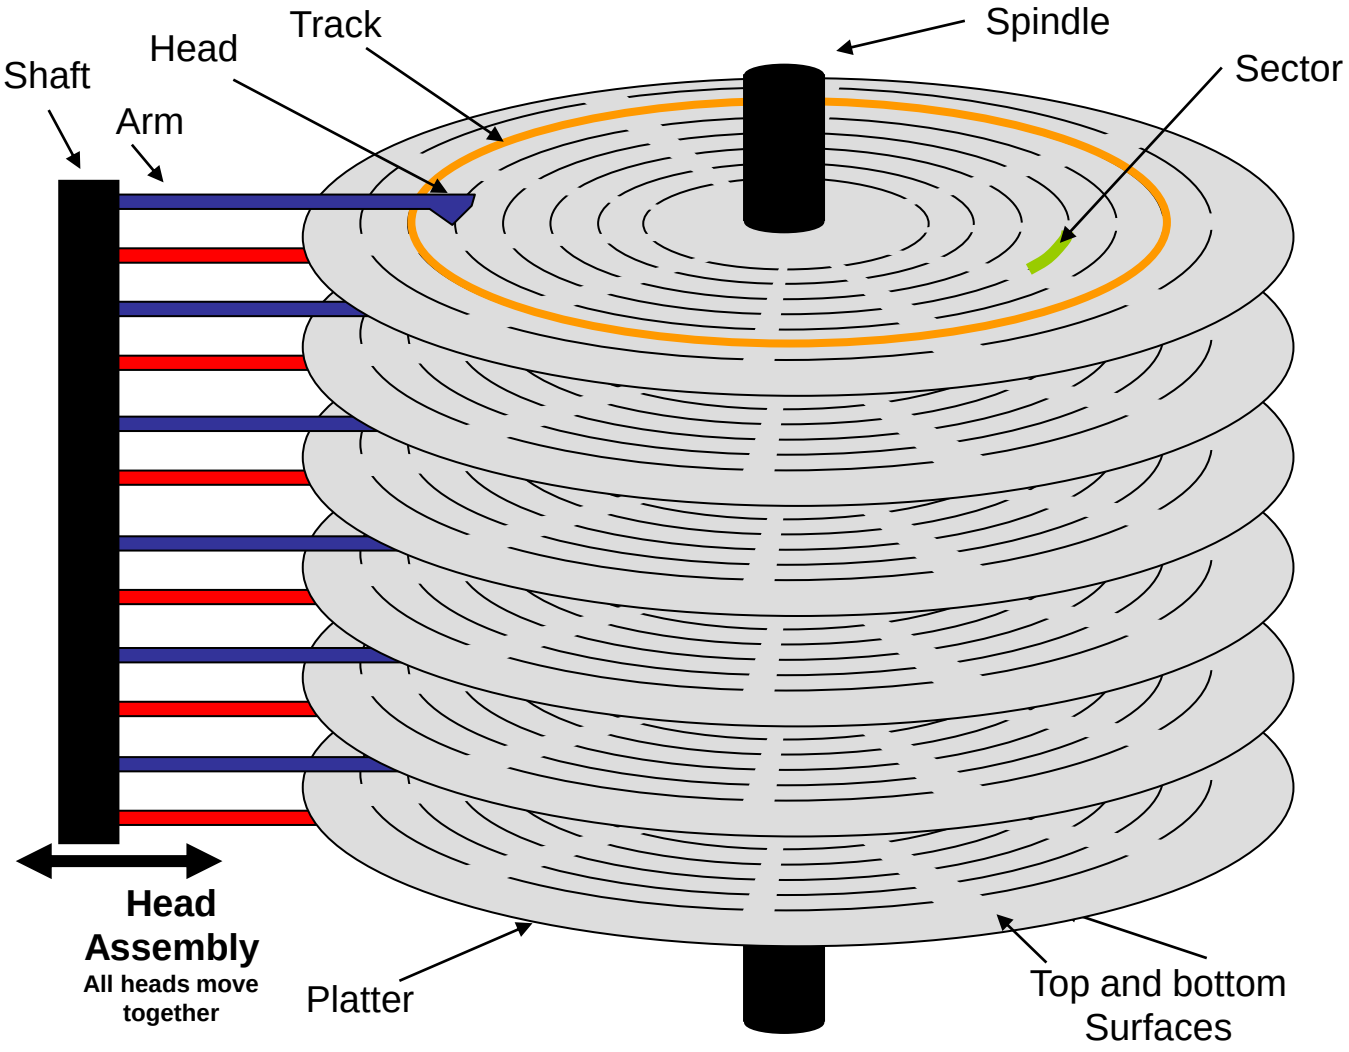

Data Organization and Formatting

- Concentric rings or tracks
  - —Gaps between tracks (inter-track gaps)
  - —Reduce gap to increase capacity
  - —Same number of bits per track (variable packing density)
  - —Constant angular velocity
- Tracks divided into sectors
  - —Gaps between sectors (inter-sector gaps or intra-track gaps)
- Min block size is one sector(data transfer unit)
- Sector size 512 bytes (almost always)
- May have more than one sector per block

Disk Data Layout

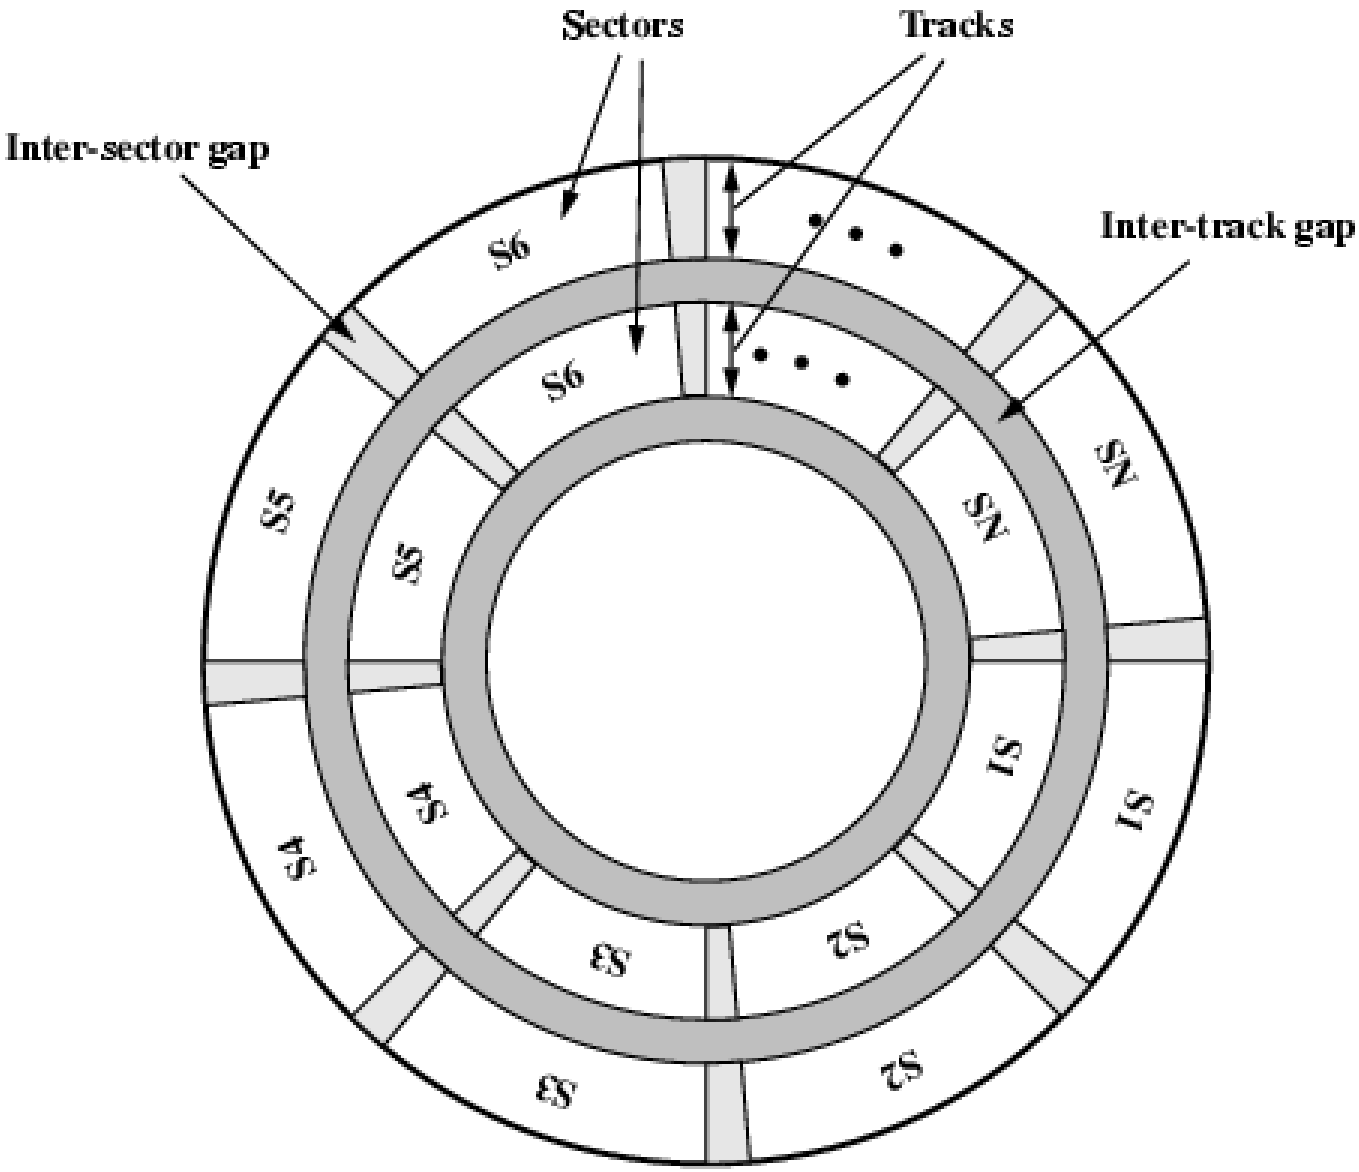

Example

Given a disk drive with the characteristics:

- Number of surfaces = 200
- Number of tracks per surface = 100
- Number of sectors per track = 50
- Bytes per sector = 256

- What is the total disk capacity?
- Total disk capacity=number of surfaces * tracks per surface * sectors per track *bytes per sector

  = 200*100*50*256 bytes

  = 256000000 bytes

  =256 MillionBytes (≠256 MB, MegaBytes=2^{20} bytes)

Disk Velocity

- Bit near centre of rotating disk passes fixed point slower than bit on outside of disk
- Increase spacing between bits in different tracks
- Rotate disk at constant angular velocity (CAV)
  - Gives pie shaped sectors and concentric tracks
  - Individual tracks and sectors addressable
  - Move head to given track and wait for given sector
  - Waste of space on outer tracks (Lower data density)
- Can use zones to increase capacity
  - Surface is divided into zones (grouping adjacent tracks)
  - Each zone has fixed bits per track
  - Zones farther from the center has more bits(sectors) than zones closer to the center
  - More complex circuitry

Disk Layout Methods Diagram

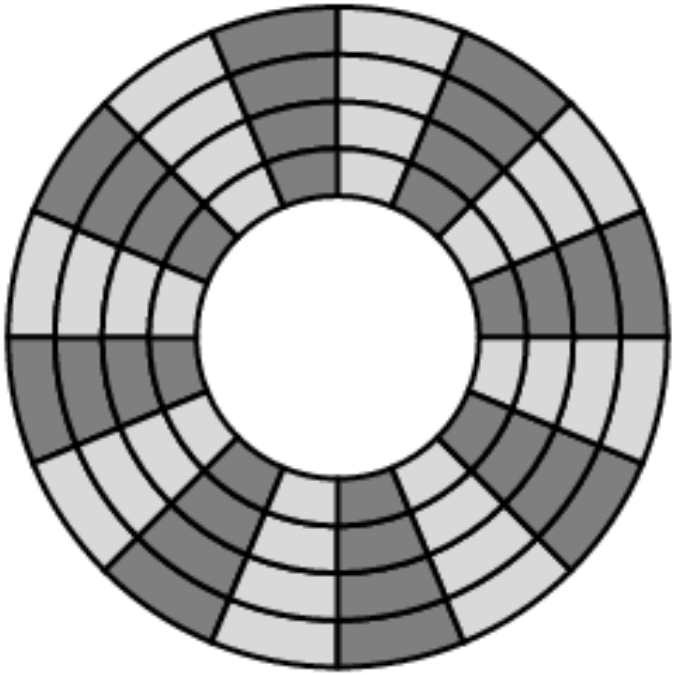

(a) Constant angular velocity

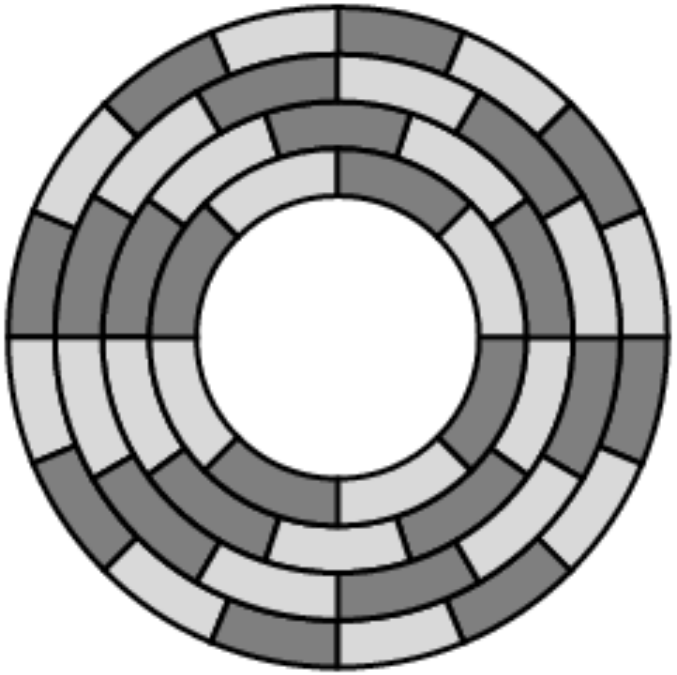

(b) Multiple zoned recording

Each zone: just a single track wide

Finding Sectors

- Must be able to identify
  - —start of track, and
  - —start and end of each sector
  - —By way of control data recorded on the disk(via format)
- Format disk
  - —Additional information: used by disk drive and not available to user
    - – marks tracks and sectors

ST506 format (track)

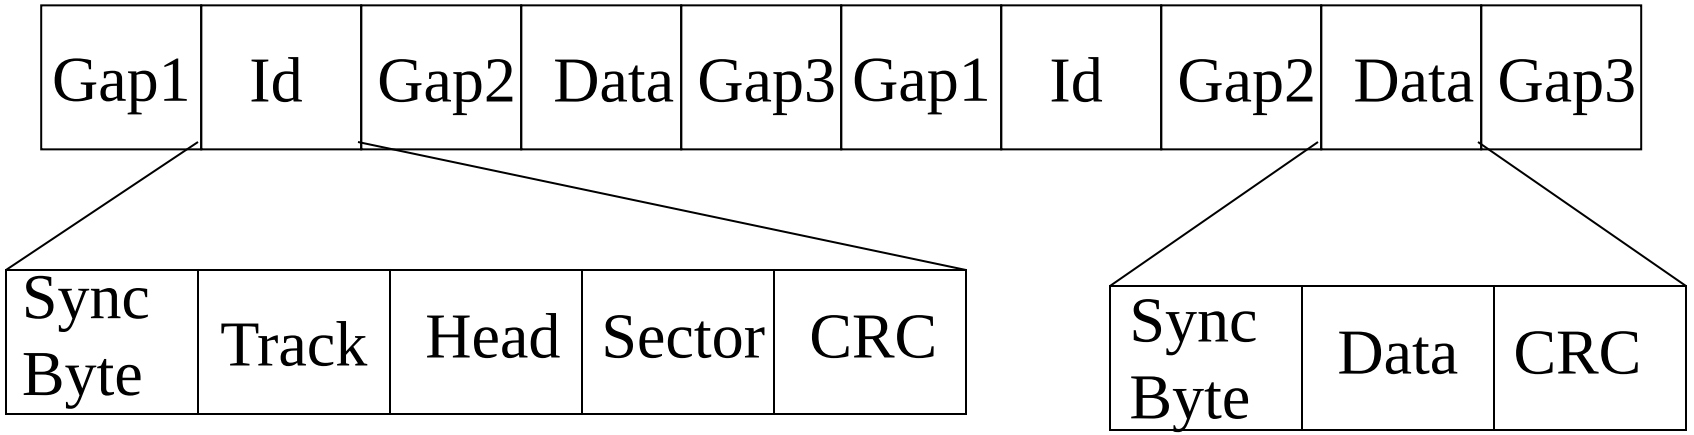

- Each track has 30 fixed length sectors of 600 bytes each
- Each sector has 512 bytes of data + control info(88 bytes)
- Id: unique address to locate a particular sector
- Synch:beginning of a field
- Track: identifies a track on a surface, head: fixed/movable
- CRC: error detecting code in ID and data fields

Disc Characteristics

- Fixed (rare) or movable head
  - Fixed:One read write head per track, heads mounted on fixed arm
  - Movable:One read write head per side, mounted on a movable arm
- Removable (R) or Fixed (F) Disk
  - R: can be removed from drive and replaced with another disk. Provides unlimited storage capacity. Easy data transfer b/w systems
  - F: Permanently mounted in the drive
- Single or double (usually) sided
- Single platter or multiple platter(reduces head movement, increases speed - transfer rate)
- Head mechanism
  - Contact (Floppy)
  - Fixed gap
  - Flying (Winchester)

Winchester Hard Disk

- Developed by IBM in Winchester (Virginia, USA)
- Sealed unit, one or more platters (disks)
- Heads fly on boundary layer of air as disk spins, rest on disk when power off
- Very small head to disk gap
- Getting more robust
  - —Automatic error corrections, remapping of bad sectors
- Universal
- Cheap
- Fastest external storage
- Getting larger all the time
  - —Multiple Gigabyte now usual

Disk Performance

Mechanizm for a disk drive to read data:3 steps

1. The head assembly has to move to the specific cylinder.
2. The disk must spin to bring the required sector under the head.
3. The data from the selected surface is read and transferred to the controller as the sector moves under the head

- disk is continuously spinning and the head is reading the bits off the surface as they pass under the head
- Mechanizm for a disk drive to write data?

Speed

- **Seek time**
  - —Time to move head to correct track
  - —Typically < 10 ms today
- **(Rotational) latency**
  - —Waiting time until start of data sector to rotate under head
- **(Disk) access time = Seek + Latency**
  - —The time it takes to get into position to read or write
  - —Once head is in position, read/write op is performed (data sector moves under the head)
- **Transfer time**
  - —The time required for the read/write of data (transfer)

Transfer Time and Rates

- Transfer time T = b / (Rn), b = bytes to transfer, r = revolutions/sec, N = number of bytes/track
- Average (data) access time (in disk) T_a

  $T_a = T_s + 1/(2r) + b/(rN)$,

  where T_s is average seek time, 1/(2r) average latency, and b/(rN) transfer time
- Given the specifications for a disk drive:
  - 256 bytes per sector
  - 2000 sectors per track
  - 1000 tracks per surface
  - 2 platters
  - Rotational speed 7500 RPM

What is the transfer rate of the disk?

- Each track has 2000 x 256 bytes = 512,000 bytes/track
- At 7500 rpm 1 rotation takes 7,500 rev/60 sec = 125 rev/sec or 0.008sec/rev
- Thus, we can transfer 512,000 bytes/0.008 sec=64,000,000 bytes/sec

Timing of Disk I/O Transfer

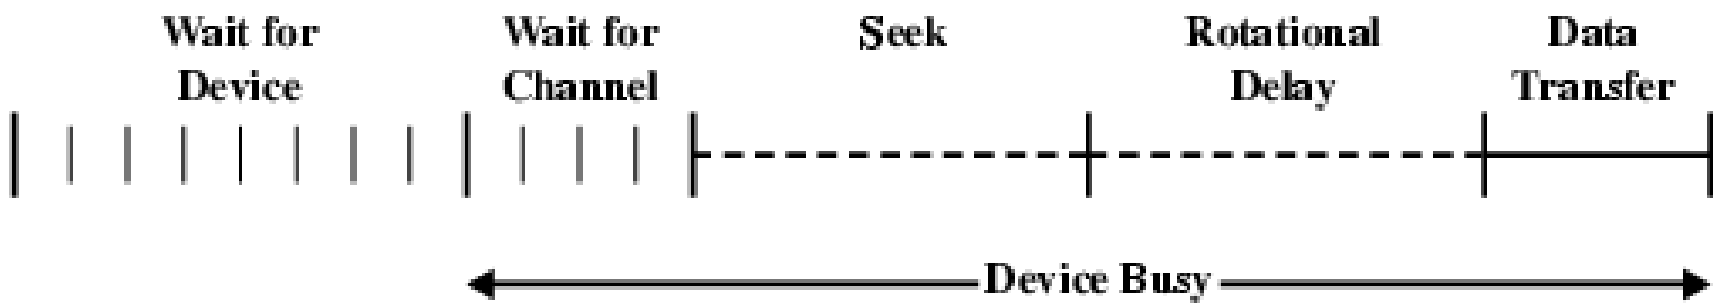

Typical Specs for a disk

Characteristics	Seagate Barracuda ES.2	Seagate Barracuda 7200.10	Seagate Barracuda 7200.9	Seagate	Hitachi Microdrive
Application	High-capacity server	High-performance desktop	Entry-level desktop	Laptop	Handheld devices
Capacity	1 TB	750 GB	160 GB	120 GB	8 GB
Minimum track-to-track seek time	0.8 ms	0.3 ms	1.0 ms	—	1.0 ms
Average seek time	8.5 ms	3.6 ms	9.5 ms	12.5 ms	12 ms
Spindle speed	7200 rpm	7200 rpm	7200	5400 rpm	3600 rpm
Average rotational delay	4.16 ms	4.16 ms	4.17 ms	5.6 ms	8.33 ms
Maximum transfer rate	3 GB/s	300 MB/s	300 MB/s	150 MB/s	10 MB/s
Bytes per sector	512	512	512	512	512
Tracks per cylinder (number of platter surfaces)	8	8	2	8	2

Design Choices in implementing a File System on a Disk Subsystem

- Some design constraints
  - —Files are of arbitrary size
  - —Files may be accesses sequentially or randomly
  - —Files need to be allocated initially
  - —Files need to be able to grow
  - —Space should be used efficiently

Contiguous Allocation

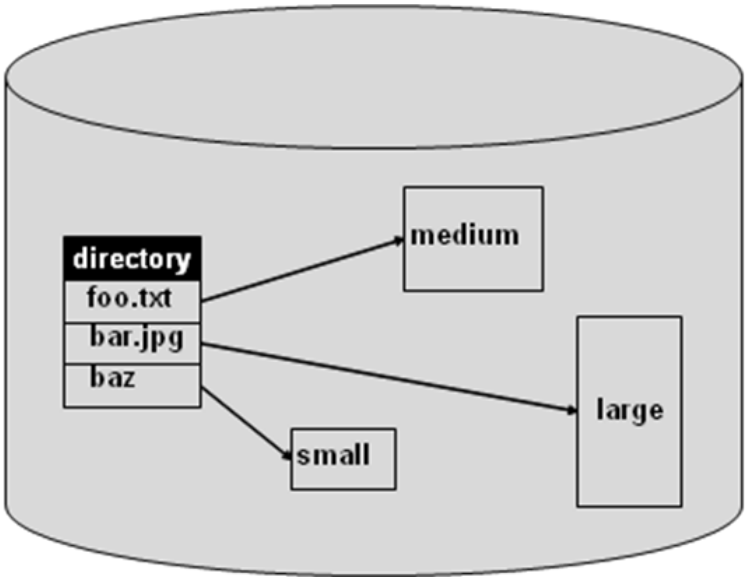

- At file creation time a set amount of space is allocated (may depend on file type)
- File cannot grow beyond that size
- Fragmentation a problem

Contiguous Allocation

- Free list
  - —Allocation may be by first or best fit
  - —Requires periodic compaction

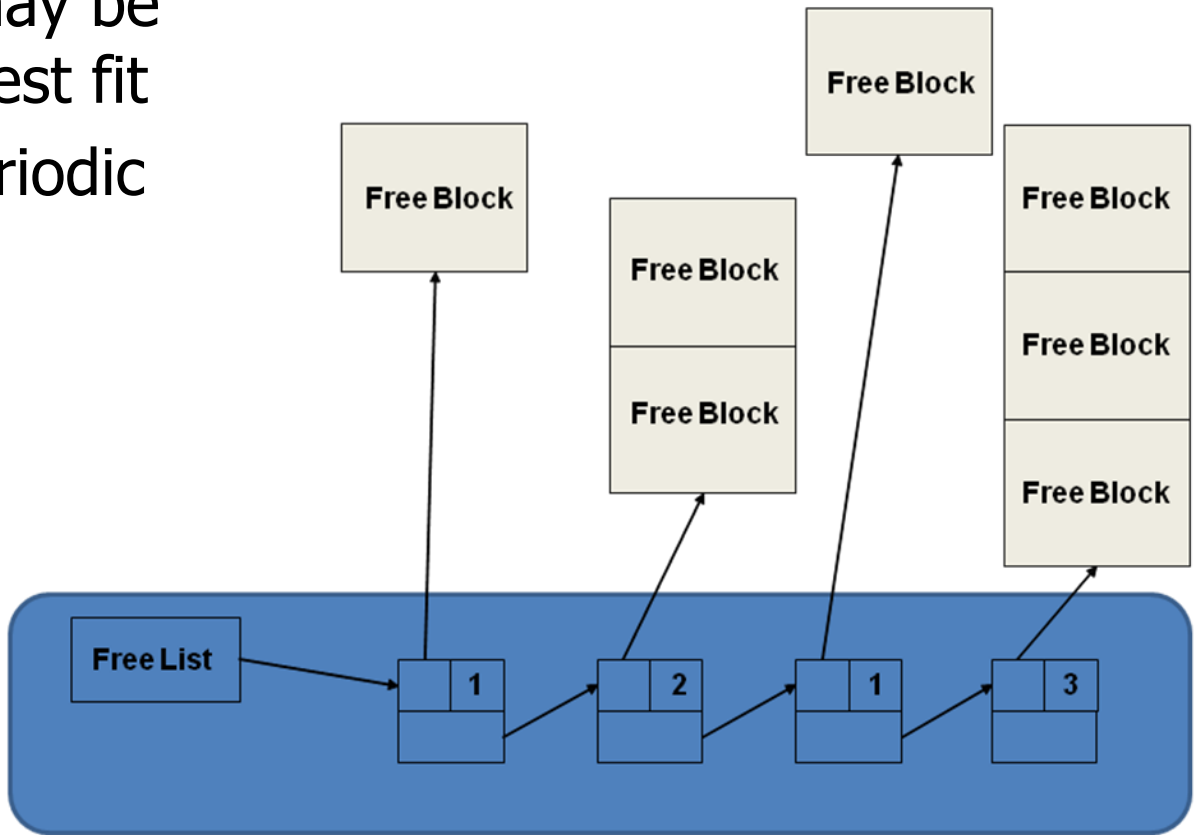

Contiguous Allocation with Overflow Area

- Modification of previous scheme to allow files to expand into a designated overflow area
- Random access suffers due to overflow area
- Despite limitations has been used extensively due to fast file access times

Linked Allocation

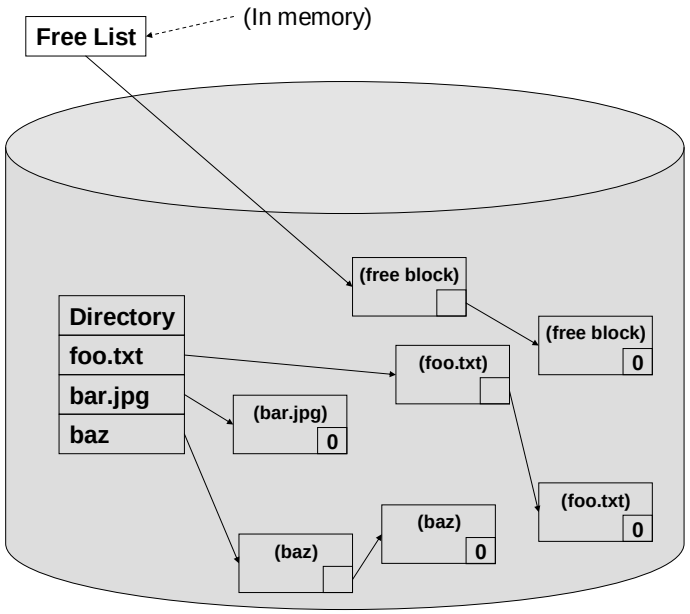

- Files not stored contiguously
- No compaction required
- Sequential and random access poor
- Susceptible to errors

File Allocation Table (FAT)

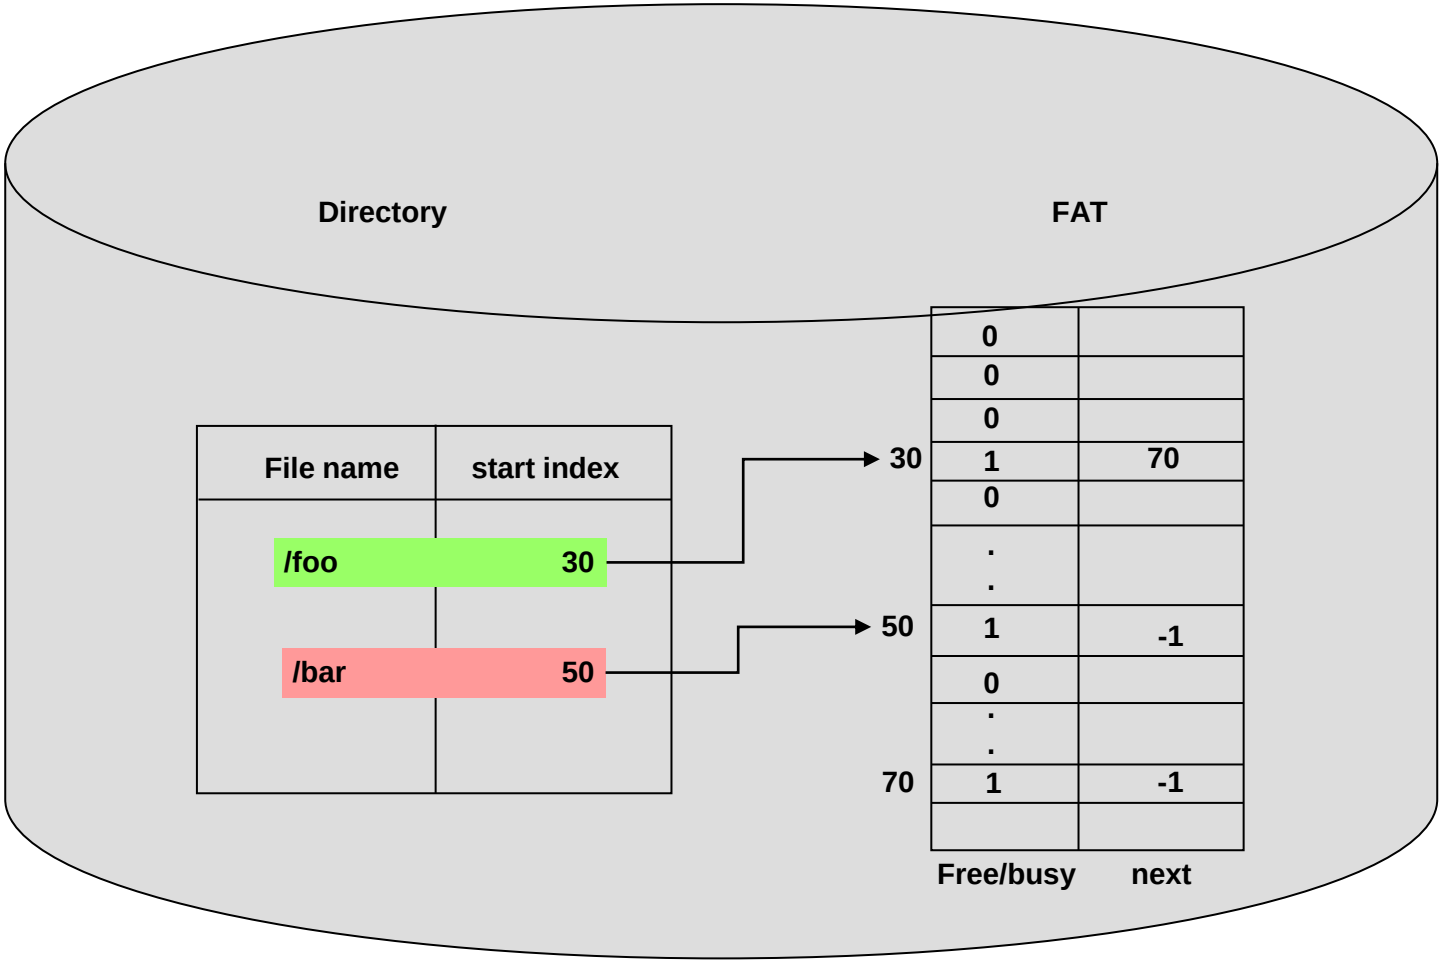

File Allocation Table (FAT)

- Divide disk into partitions
- Each partition has a FAT
- The directory just has a pointer into the starting sector entry in the FAT for each file.
- Less chance for errors than linked allocation
- FAT becomes big so clustering and partitioning may be necessary leading to other problems

Indexed Allocation

- Essentially breaks up FAT into one data structure per file
- Allocate an index disk block for each file called an i-node
- Directory entries now point to the i-node for that file
- Maintain free list as bit vector