Chapter 6A

Storage and Other I/O Topics

Introduction

- Users become apoplectic if their storage system crash or lose information
- The standard for dependability is much higher for storage than for computation
- Emphasis
  - I/O devices: dependability and cost
  - Processor and memory: performance and cost
- I/O systems must plan
  - for expandability of devices
  - for diversity of devices

Characteristics of I/O devices

- I/O devices are diverse and can be characterized by
  - Behavior: input (read only), output (write only, cannot be read), storage (read and write)
  - Partner: human or machine
  - Data rate: bytes/sec, transfers/sec
- I/O bus connections

Connection of I/O devices

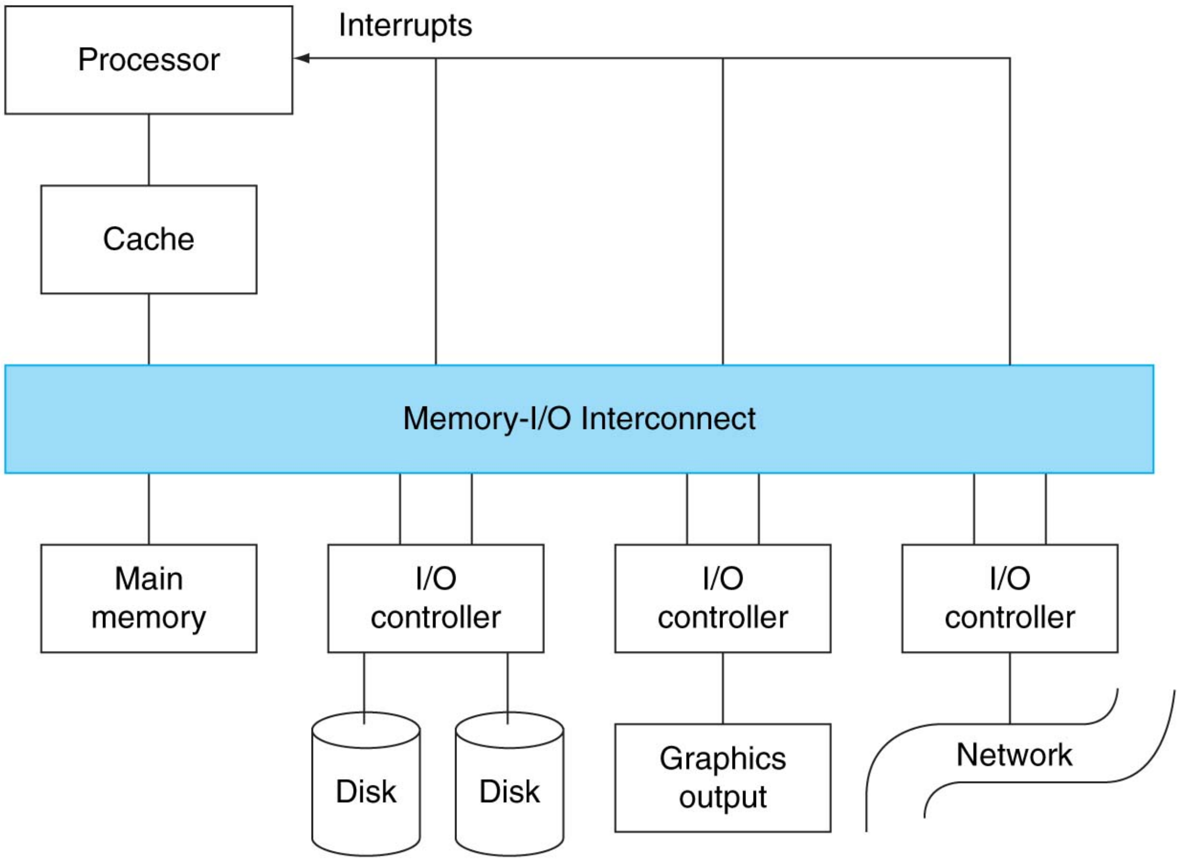

The connections between the I/O devices, processor, and memory are historically called *buses*, although the term means shared parallel wires and most I/O connections today are closer to dedicated serial lines. Communication among the devices and the processor uses both interrupts and protocols on the interconnect.

Diversity of I/O Devices

Device	Behavior	Partner	Data rate (Mbit/sec)
Keyboard	Input	Human	0.0001
Mouse	Input	Human	0.0038
Voice input	Input	Human	0.2640
Sound input	Input	Machine	3.0000
Scanner	Input	Human	3.2000
Voice output	Output	Human	0.2640
Sound output	Output	Human	8.0000
Laser printer	Output	Human	3.2000
Graphics display	Output	Human	800.0000–8000.0000
Cable modem	Input or output	Machine	0.1280–6.0000
Network/LAN	Input or output	Machine	100.0000–10000.0000
Network/wireless LAN	Input or output	Machine	11.0000–54.0000
Optical disk	Storage	Machine	80.0000–220.0000
Magnetic tape	Storage	Machine	5.0000–120.0000
Flash memory	Storage	Machine	32.0000–200.0000
Magnetic disk	Storage	Machine	800.0000–3000.0000

I/O System Characteristics

- Dependability is important
  - Particularly for storage devices
- I/O Performance measures
  - How much data can we move through the system in a certain time?
  - How many I/O operations can we do per unit of time?
- Desktops & embedded systems are mainly interested in
  - response time (response time)
  - diversity of devices
- Servers are interested in
  - throughput (bandwidth) &
  - expandability of devices

Dependability, Reliability, Availability

- Definition of Dependability
  - Computer system dependability is the quality of delivered service such that reliance can justifiably be placed on this service.
  - The service delivered by a system is its observed actual behavior as perceived by other system(s) interacting with this system's users.
  - Each module also has an ideal specified behavior, where a service specification is an agreed description of the expected behavior.
  - A system failure occurs when the actual behavior deviates form the specified behavior.

Dependability

- Users can see a system alternating between tow states of delivered services with respect to the service specification:
  1. Service accomplishment
  2. Service interruption

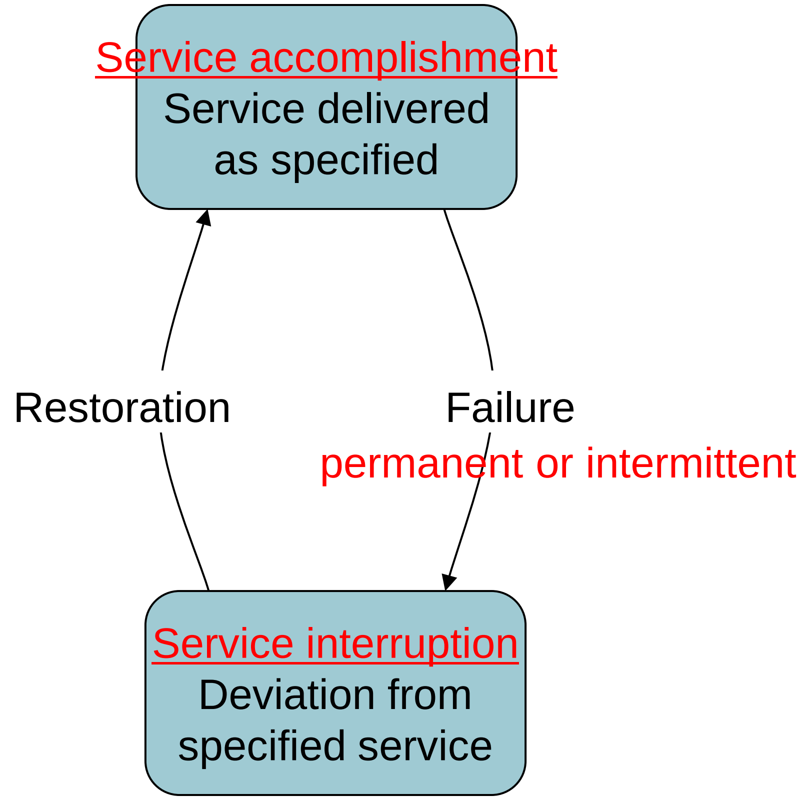

Dependability

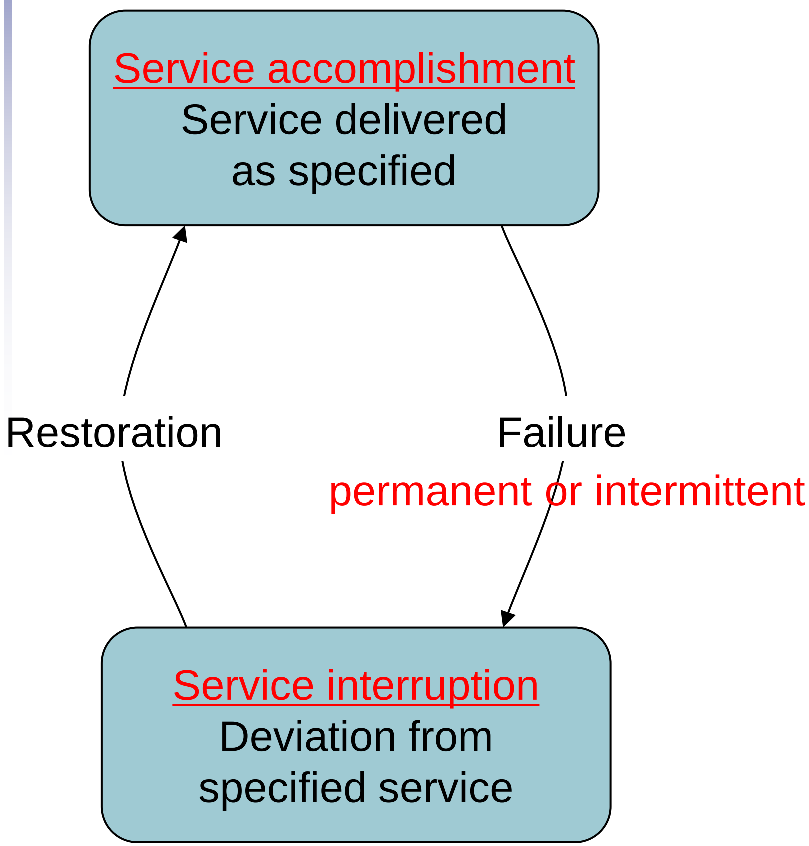

- Fault: failure of a component
  - May or may not lead to system failure
- Reliability: a measure of the continuous service accomplishment or, equivalently of the time to failure
- Availability: a measure of service accomplishment with respect to the alteration between the two states of accomplishment and interruption.

Dependability Measures

- Reliability: mean time to failure (MTTF)
- Service interruption: mean time to repair (MTTR)
- Mean time between failures
  - MTBF = MTTF + MTTR
- Availability = MTTF / MTBF = MTTF / (MTTF + MTTR)
- Improving Availability
  - Increase MTTF: fault avoidance, fault tolerance, fault forecasting
  - Reduce MTTR: improved tools and processes for diagnosis and repair

To increase MTTF

- Fault avoidance: preventing fault occurrence by construction
- Fault tolerance: using redundancy to allow the service to comply with the service specification despite fault occurring, which applies primarily to hardware faults.
- Fault forecasting: predicting the presence and creation of faults, which applies to hardware and software faults. Allowing the components to be replaced before it faults.

Reasons for failures

Human operators are a significant source of failures.

Operator	Software	Hardware	System	Year data collected
42%	25%	18%	Datacenter (Tandem)	1985
15%	55%	14%	Datacenter (Tandem)	1989
18%	44%	39%	Datacenter (DEC VAX)	1985
50%	20%	30%	Datacenter (DEC VAX)	1993
50%	14%	19%	U.S. public telephone network	1996
54%	7%	30%	U.S. public telephone network	2000
60%	25%	15%	Internet services	2002

Although it is difficult to collect data to determine whether operators are the cause of errors, since operators often record the reasons for failures, these studies did capture that data. There were often other categories, such as environmental reasons for outages, but they were generally small. The top two rows come from a classic paper by Jim Gray [1990], which is still widely quoted almost 20 years after the data was collected. The next two rows are from a paper by Murphy and Gent, who studied causes of outages in VAX systems over time ["Measuring system and software reliability using an automated data collection process," *Quality and Reliability Engineering International* 11:5, September–October 1995, 341–53]. The fifth and sixth rows are studies of FCC failure data about the U.S. public switched telephone network by Kuhn ["Sources of failure in the public switched telephone network," *IEEE Computer* 30:4, April 1997, 31–36] and by Patty Enriquez. The study of three Internet services is from Oppenheimer, Ganapath, and Patterson [2003].

Disk Storage

- Nonvolatile: the data remains even when power is removed.
- A collection of platters
  - Each platter has two recordable disk surfaces
- Tracks and sectors
- Rotating magnetic storage
  - 5400 ~ 15000 RPM

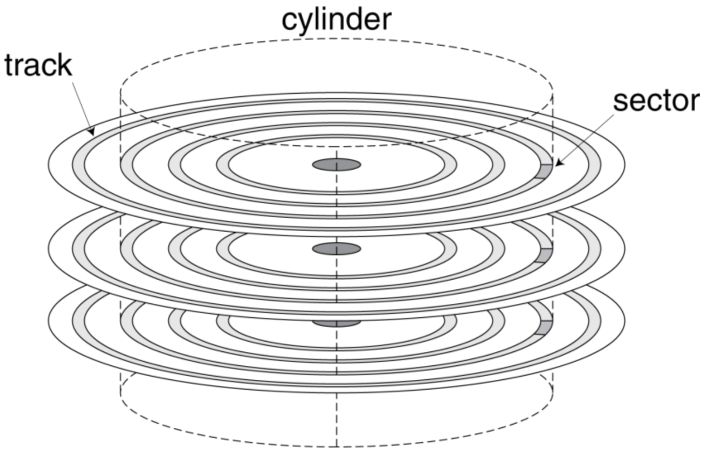

Magnetic Disk

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

Moving-head Disk Mechanism

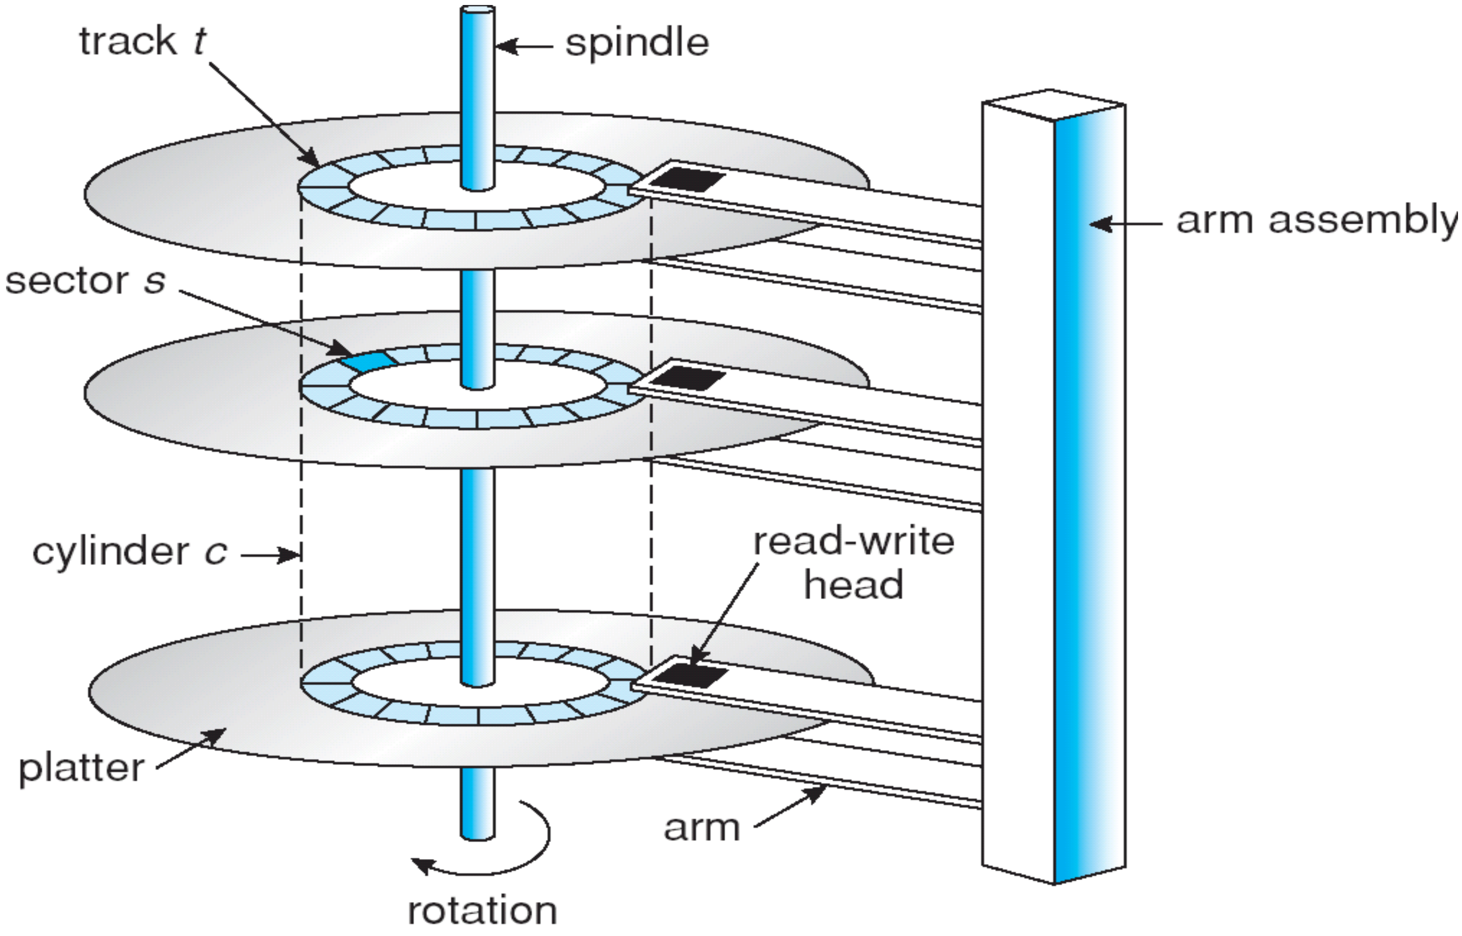

Cylinder

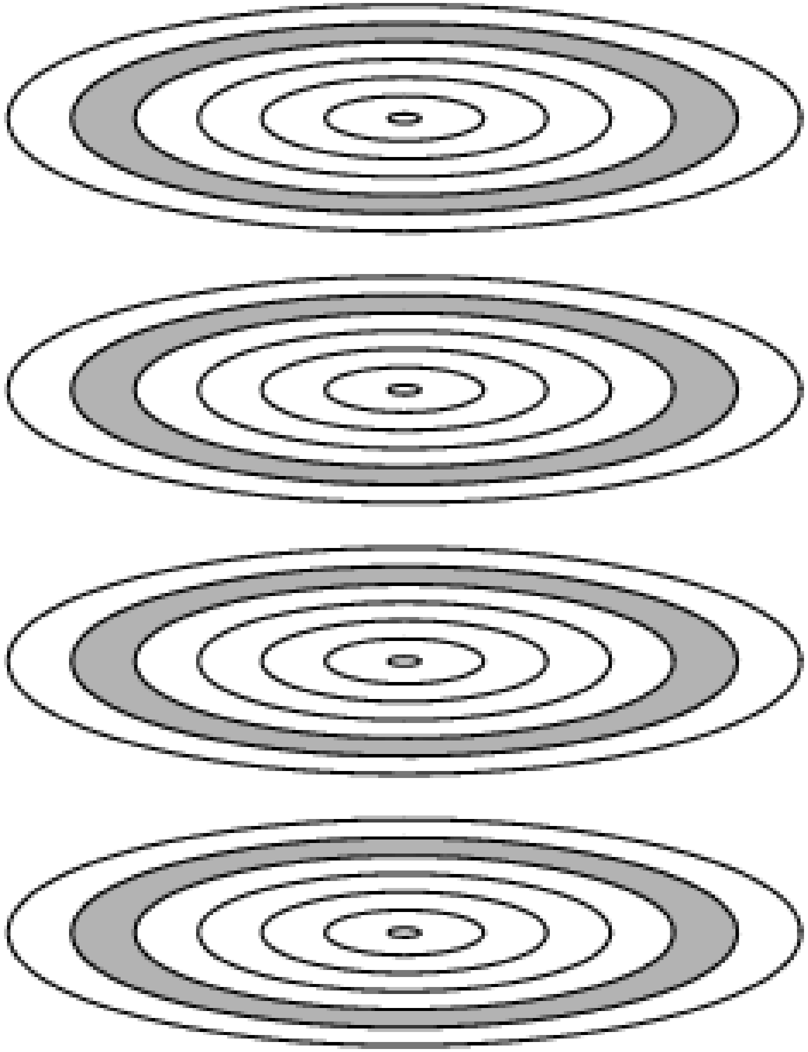

Read and Write Mechanisms

- Recording & 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
  - (MR) magneto resistive sensor
  - Electrical resistance depends on direction of magnetic field
  - High frequency operation
    - Higher storage density and speed

Inductive Write MR Read

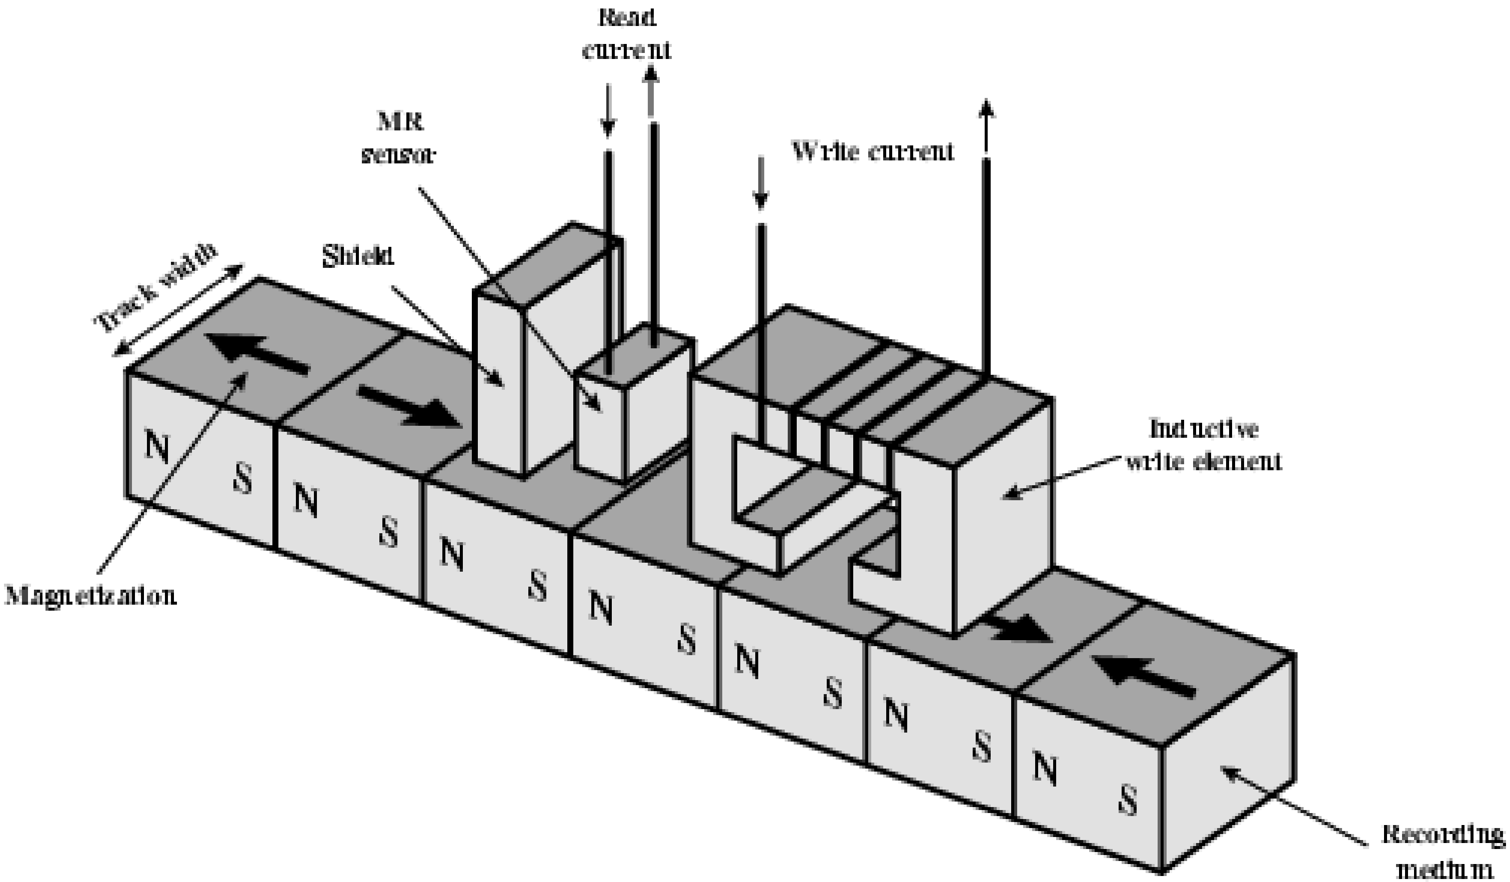

Disk Data Layout – not used at present

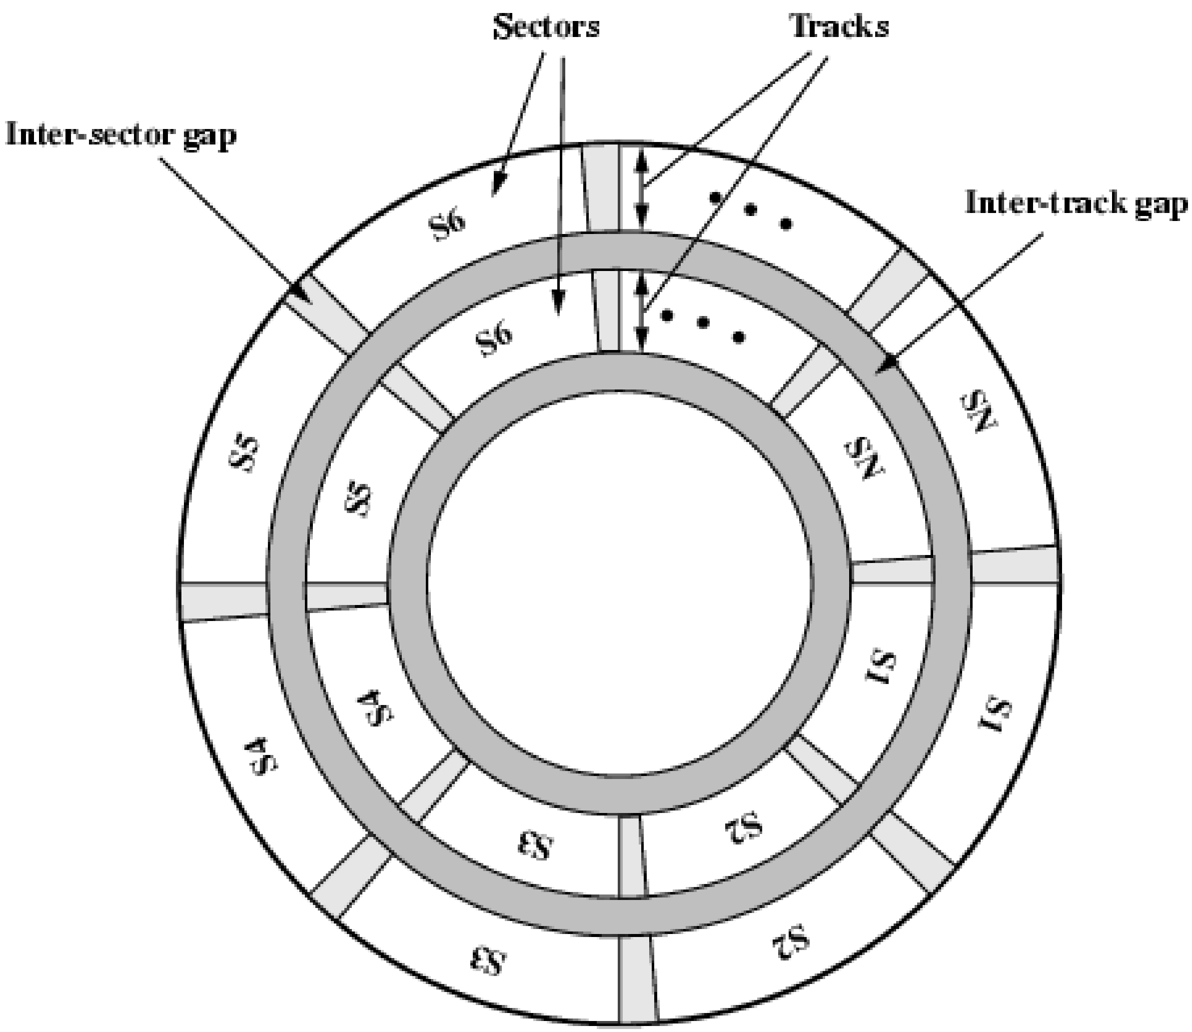

Disk Velocity

- Bit near centre of rotating disk passes fixed point slower than bit on outside of disk
- Different 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
  - Each zone has fixed bits per track
  - More complex circuitry

Disk Layout Methods Diagram

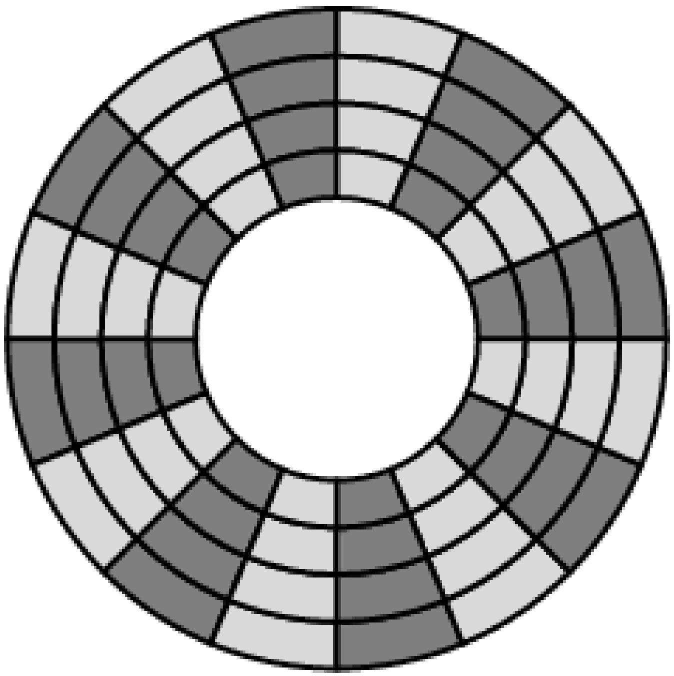

(a) Constant angular velocity

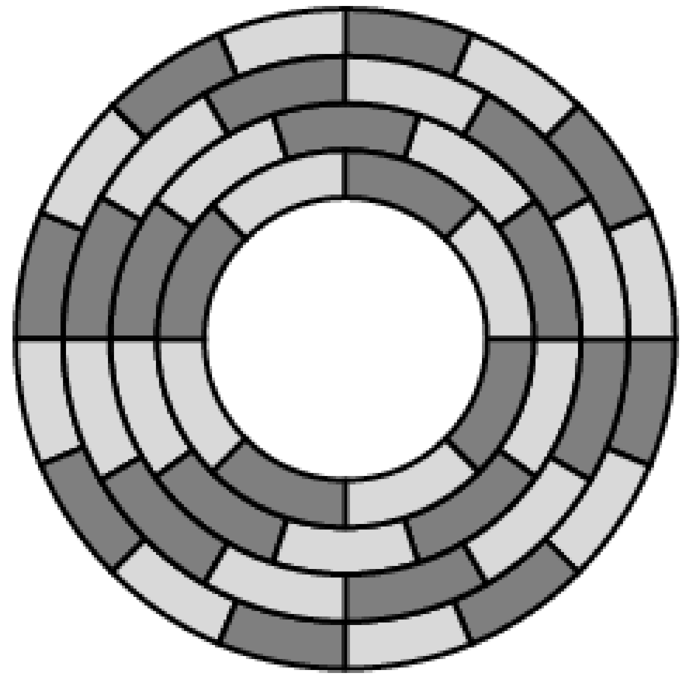

(b) Multiple zoned recording

Finding Sectors

- Must be able to identify start of track and sector
- Format disk
  - Additional information not available to user
  - Marks tracks and sectors

Disk Format Seagate ST506

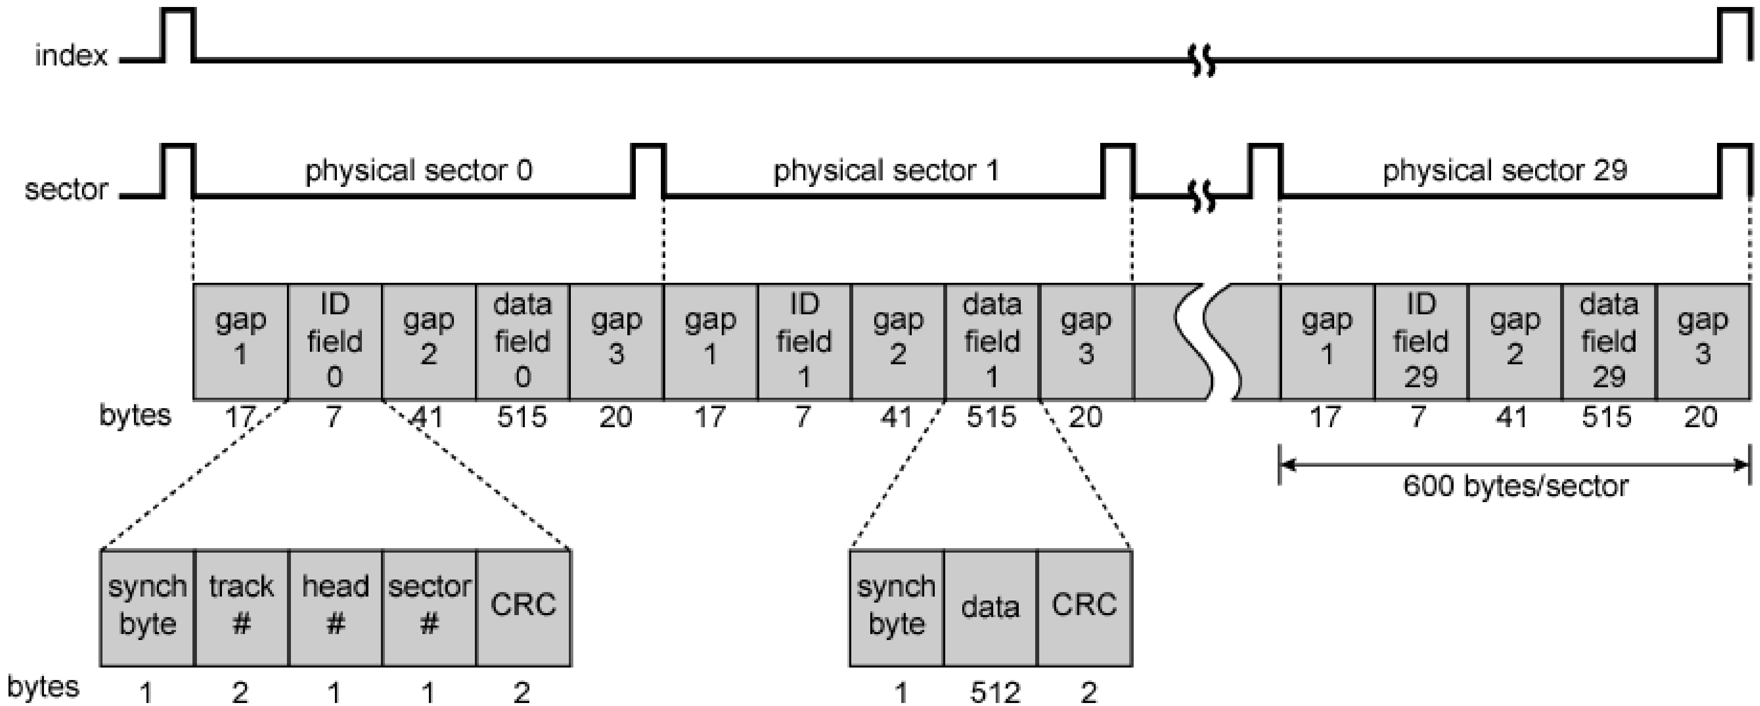

Cylinder skew

What is the problem? Assume the head is positioned at the inner most track, and system will read 33 consecutive sectors.
If no cylinder skew, when the arm seeks from inner most track to the 2^{nd} inner most track, sector 0 in 2^{nd} inner most track will pass the head.

Example: (Cylinder skew calculation)

For example, a 10,000-RPM drive rotates in 6 msec. If a track contains 300 sectors, a new sector pass under the head every 20 μsec. If the track-to-track seek time is 800 μsec, 40 sectors will pass by during the seek, so the cylinder skew should be 40 sectors.

cyliner skew = 40 / 300 x 360^o = 56^o

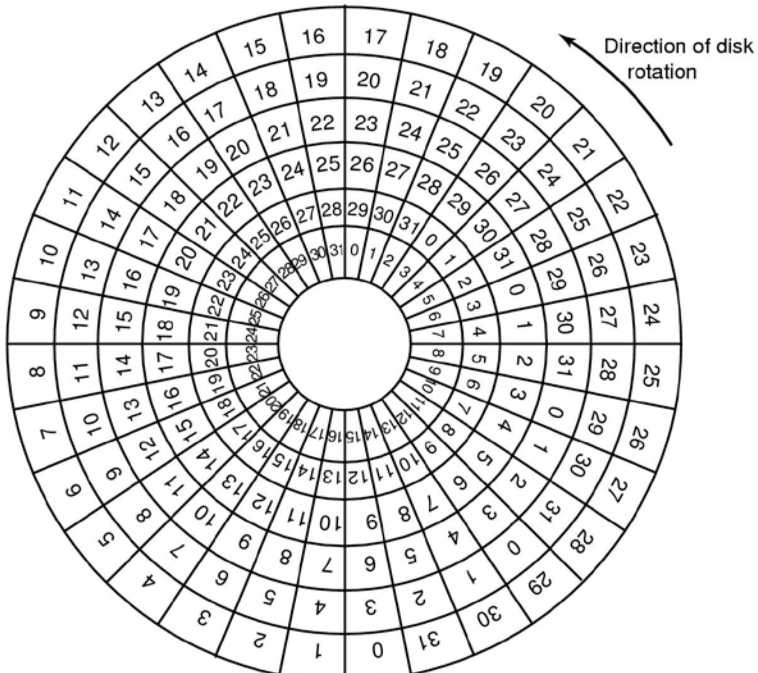

How to number the sectors

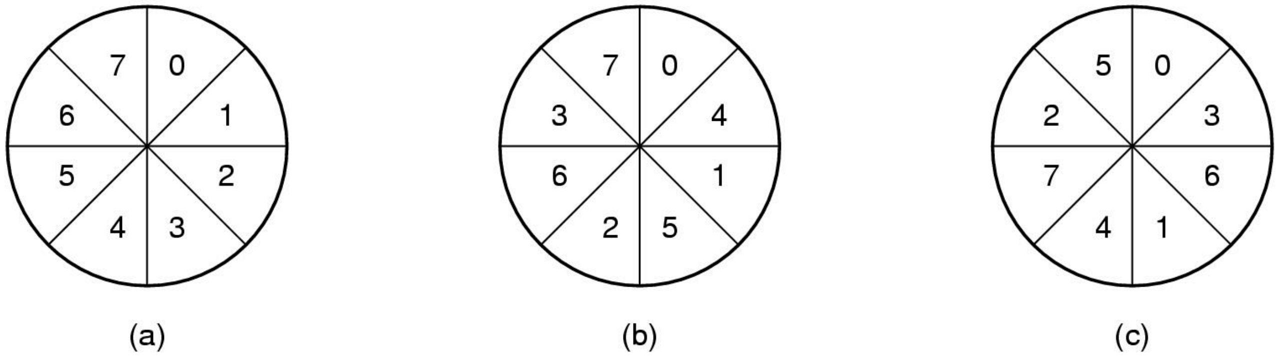

No interleaving | Single interleaving | Double interleaving

Problem of No interleaving:

Assume the buffer of the disk controller can only store one sector. When the buffer is full, it will send the data to the main memory. While sending data to the main memory, the next sector pass by the head. It will wait for an entire rotation to read next sector.

Disk Performance Issues

- Smart disk controller allocate physical sectors on disk
  - Present logical sector interface to host
- Connections
  - SCSI (small computer systems interface)
  - ATA (advanced technology attachment),
  - SATA: Serial ATA
  - SAS: Serial Attached SCSI
- Disk drives include caches
  - Prefetch sectors in anticipation of access
  - Avoid seek and rotational delay

Disk Performance Issues

- Manufacturers quote average seek time
  - Based on all possible seeks
  - Locality and OS scheduling lead to smaller actual average seek times
- Actual seek time depends on the application and scheduling of disk requests
  - May be only 25% ~ 33% of the advertised number because of locality of disk references
- Cylinder skew
- How to number sectors in a track

Components of Disk I/O Transfer

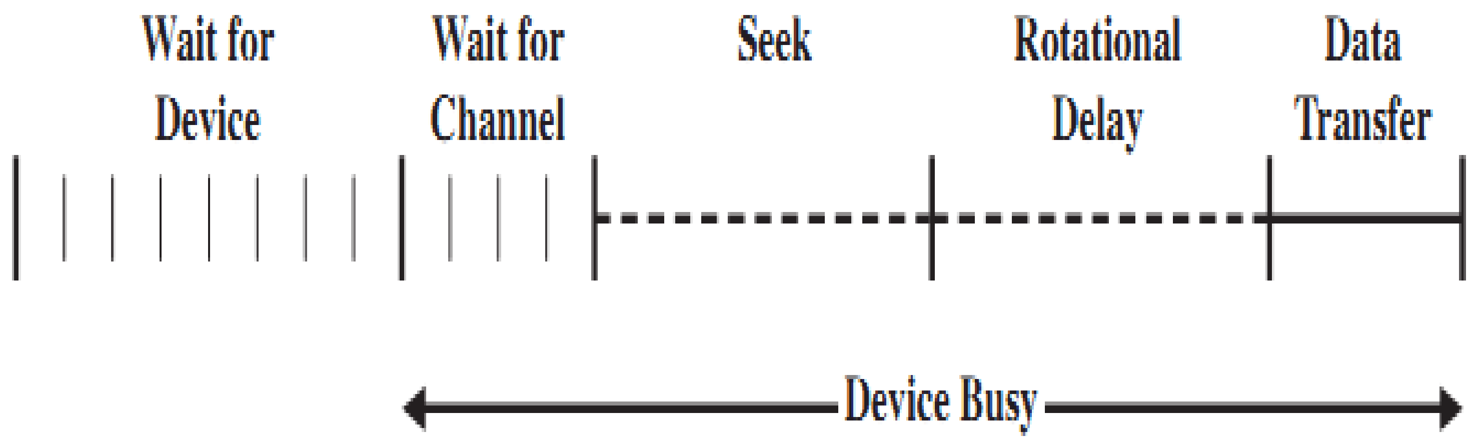

Disk Sectors and Access

- Each sector records
  - Sector ID
  - Data (512 bytes, 4096 bytes proposed)
  - Error correcting code (ECC)
    - Used to hide defects and recording errors
  - Synchronization fields and gaps
- Access to a sector involves
  - Queuing: delay if other accesses are pending
  - Seek: move the head over the proper track on a disk
  - Rotational latency: reach the right sector
  - Data transfer: read the sector
  - Controller overhead

Disk Access Example

- Given
  - 512B sector, 15,000rpm, 4ms average seek time, 100MB/s transfer rate, 0.2ms controller overhead, idle disk
- Average read time
  - 4ms average seek time
    + ½ / (15,000/60) = 2ms rotational latency
    + 512B / 100MB/s = 0.005ms transfer time
    + 0.2ms controller delay
    = 6.205ms
- If actual average seek time is 1ms
  - Rotational latency is the largest
  - Average read time = 3.2ms

Disk Arm Scheduling Algorithms

- The time to read or write a disk block is determined by 3 factors
  - the seek time: the time to move the arm to the proper cylinder (most important)
  - the rotational delay: the time for the proper sector to rotate under the head
  - the actual transfer time
- Three algorithms
  - FCFS (first come, first served): the disk driver accepts the request one at a time and carries then out in that order
  - SSF (shortest seek first)
  - Elevator algorithm

Disk Scheduling

- The operating system is responsible for using hardware efficiently
- For the disk drives, this means having a fast access time and disk bandwidth.
- Access time has two major components:
  - Seek time + Rotational
  - Minimize seek time ≈ seek distance.
- Disk bandwidth is the total number of bytes transferred, divided by the total time between the first request for service and the completion of last transfer.

Disk Structure

- Disk drives are addressed as large 1-dimensional arrays of logical blocks, where the logical block is the smallest unit of transfer.
- The 1-dimensional array of logical blocks is mapped into the sectors of the disk sequentially:
  - Sector 0 is the first sector of the first track on the outermost cylinder.
  - Mapping proceeds in order through that track, then the rest of the tracks in that cylinder, and then through the rest of the cylinders from outermost to innermost.

First Come First Serve (FCFS)

- Handle I/O requests sequentially.
- Fair to all processes.
- Approaches random scheduling in performance if there are many processes/requests.
- Suffers from global zigzag effect.

Shortest Seek Time First (SSTF)

- Selects the request with the minimum seek time from the current head position.
- Also called Shortest Seek Distance First (SSDF) – It's easier to compute distances.
- It's biased in favor of the middle cylinders requests.
- SSTF scheduling is a form of SJF scheduling; may cause starvation of some requests.

Elevator Algorithms

- Algorithms based on the common elevator principle.
- Four combinations of Elevator algorithms:
  – Service in both directions or in only one direction.
  – Go until last cylinder or until last I/O request.

Direction \ Go until	Go until the last cylinder	Go until the last request
Service both directions	**Scan**	**Look**
Service in only one direction	**C-Scan**	**C-Look**

Scan

- The disk arm starts at one end of the disk, and moves toward the other end, servicing requests until it gets to the other end of the disk, where the head movement is reversed and servicing continues.
- It moves in both directions until both ends.
- Tends to stay more at the ends so more fair to the extreme cylinder requests.

Look

- The disk arm starts at the first I/O request on the disk, and moves toward the last I/O request on the other end, servicing requests until it gets to the other extreme I/O request on the disk, where the head movement is reversed and servicing continues.
- It moves in both directions until both last I/O requests; more inclined to serve the middle cylinder requests.

C-Scan

- The head moves from one end of the disk to the other, servicing requests as it goes. When it reaches the other end, however, it immediately returns to the beginning of the disk, without servicing any requests on the return trip.
- Treats the cylinders as a circular list that wraps around from the last cylinder to the first one.
- Provides a more uniform wait time than SCAN; it treats all cylinders in the same manner.

C-Look

- Look version of C-Scan.
- Arm only goes as far as the last request in each direction, then reverses direction immediately, without first going all the way to the end of the disk.
- In general, Circular versions are more fair but pay with a larger total seek time.
- Scan versions have a larger total seek time than the corresponding Look versions.

Example

(a) FIFO (starting at track 100)		(b) SSTF (starting at track 100)		(c) SCAN (starting at track 100, in the direction of increasing track number)		(d) C-SCAN (starting at track 100, in the direction of increasing track number)	
Next track accessed	**Number of tracks traversed**	**Next track accessed**	**Number of tracks traversed**	**Next track accessed**	**Number of tracks traversed**	**Next track accessed**	**Number of tracks traversed**
55	45	90	10	150	50	150	50
58	3	58	32	160	10	160	10
39	19	55	3	184	24	184	24
18	21	39	16	90	94	18	166
90	72	38	1	58	32	38	20
160	70	18	20	55	3	39	1
150	10	150	132	39	16	55	16
38	112	160	10	38	1	58	3
184	146	184	24	18	20	90	32
Average seek length	55.3	**Average seek length**	27.5	**Average seek length**	27.8	**Average seek length**	35.8

Summary of Disk Scheduling Algorithms

- The performance depends heavily on the workload (number of requests). Under light load all algorithms perform the same.
  - If the queue seldom has more than one outstanding request, then all algorithms are effectively the same.
- Their performance also depends upon the file organization and the type of generated requests.
  - In a sequential processing and sequential file, the head movement will be minimum and therefore the seek time and latency time will be minimum so FCFS may perform better.
  - An indexed sequential file, on the other hand, may include blocks that may be scattered all over the disk and a sequential processing with FCFS will be very slow.
- SSTF is quite common
- SCAN and C-SCAN are good for heavy load

Magnetic Disk Trends

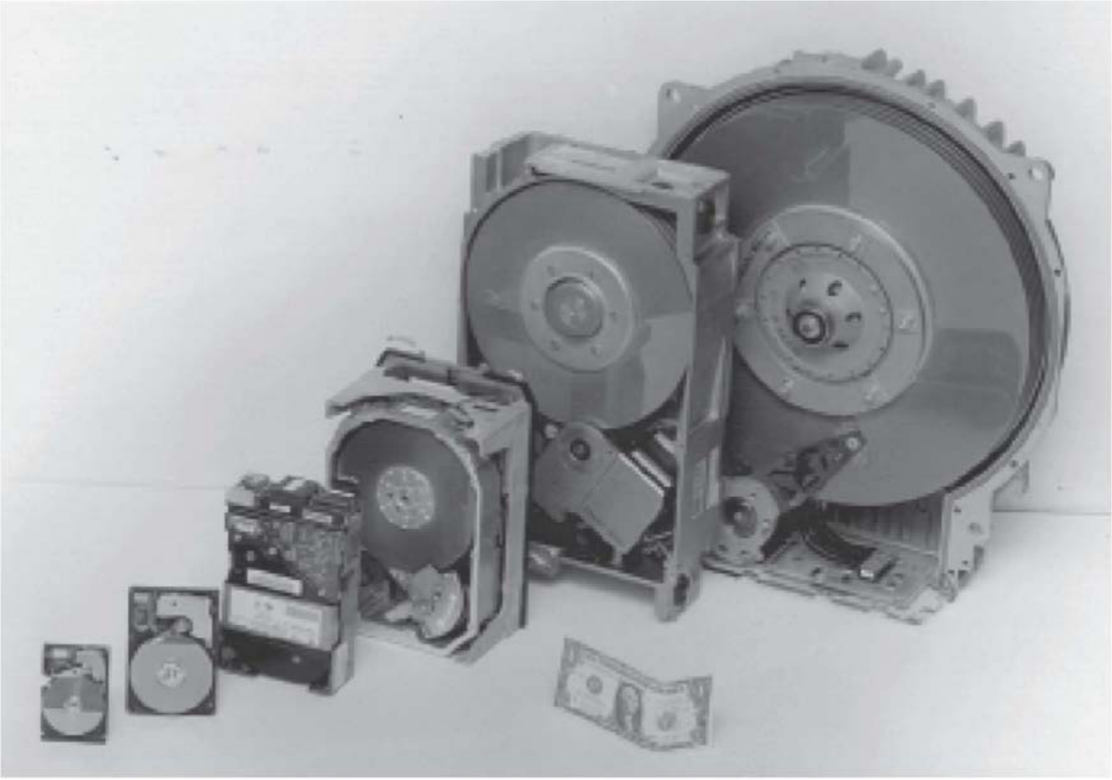

- The widest disk is the DEC R81 with four 14-inch diameter platters and storing 456 MB (1985)
- The 8-inch diameter Fujitsu disk stores 130 MB on six platters (1984).
- The Micropolis RD53 has five 5.25-inch platters and stores 85 MB.
- The IBM 0361with five platters with just 3.5 inches holds 320 MB (1988). In 2008, the most dense 3.5-inch disk with 2 platters held 1 TB in the same space. (density is about 3000 times!)
- The Conner CP 2045: two 2.5-inch platters containing 40 MB (1990)
- The smallest disk is the Integral 1820 with a single 1.8-inch platter contains 20 MB (1992)

Current Magnetic Disks

Characteristics	Seagate ST33000655SS	Seagate ST31000340NS	Seagate ST973451SS	Seagate ST9160821AS
Disk diameter (inches)	3.50	3.50	2.50	2.50
Formatted data capacity (GB)	147	1000	73	160
Number of disk surfaces (heads)	2	4	2	2
Rotation speed (RPM)	15,000	7200	15,000	5400
Internal disk cache size (MB)	16	32	16	8
External interface, bandwidth (MB/sec)	SAS, 375	SATA, 375	SAS, 375	SATA, 150
Sustained transfer rate (MB/sec)	73–125	105	79–112	44
Minimum seek (read/write) (ms)	0.2/0.4	0.8/1.0	0.2/0.4	1.5/2.0
Average seek read/write (ms)	3.5/4.0	8.5/9.5	2.9/3.3	12.5/13.0
Mean time to failure (MTTF) (hours)	1,400,000 @ 25°C	1,200,000 @ 25°C	1,600,000 @ 25°C	—
Annual failure rate (AFR) (percent)	0.62%	0.73%	0.55%	—
Contact start-stop cycles	—	50,000	—	>600,000
Warranty (years)	5	5	5	5
Nonrecoverable read errors per bits read	<1 sector per 10^{16}	<1 sector per 10^{15}	<1 sector per 10^{16}	<1 sector per 10^{14}
Temperature, shock (operating)	5°–55°C, 60 G	5°–55°C, 63 G	5°–55°C, 60 G	0°–60°C, 350 G
Size: dimensions (in.), weight (pounds)	1.0" × 4.0" × 5.8", 1.5 lbs	1.0" × 4.0" × 5.8", 1.4 lbs	0.6" × 2.8" × 3.9", 0.5 lbs	0.4" × 2.8" × 3.9", 0.2 lbs
Power: operating/idle/standby (watts)	15/11/—	11/8/1	8/5.8/—	1.9/0.6/0.2
GB/cu. in., GB/watt	6 GB/cu.in., 10 GB/W	43 GB/cu.in., 91 GB/W	11 GB/cu.in., 9 GB/W	37 GB/cu.in., 84 GB/W
Price in 2008, $/GB	~ $250, ~ $1.70/GB	~ $275, ~ $0.30/GB	~ $350, ~ $5.00/GB	~ $100, ~ $0.60/GB

SAS: serial attached SCSI
SATA: serial ATA

Seatgate 10K HDD

SED Model Number SED FIPS 140-2 Model Number	ST900MM0026 ST900MM0036[3]	ST600MM0026	ST450MM0026	ST300MM0026
Interface	6Gb/s SAS	6Gb/s SAS	6Gb/s SAS	6Gb/s SAS
Capacity				
Formatted 512 Bytes/Sector (GB)	900	600	450	300
External Transfer Rate (MB/s)	600	600	600	600
Performance				
Spindle Speed (RPM)	10K	10K	10K	10K
Average Latency (ms)	2.9	2.9	2.9	2.9
Sustained Transfer Rate Outer to Inner Diameter (MB/s) Cache, Multisegmented (MB)	204 to 125 64	204 to 125 64	204 to 125 64	204 to 125 64
Configuration/Reliability				
Disks	3	2	2	1
Heads	6	4	3	2
Nonrecoverable Read Errors per Bits Read	1 per 10E16	1 per 10E16	1 per 10E16	1 per 10E16
Annualized Failure Rate (AFR)	0.44%	0.44%	0.44%	0.44%
Power Management				
Typical Op (amps) +5V/+12V	0.41/0.49	0.41/0.44	0.40/0.43	0.41/0.40
Power Idle (watts)	3.9	3.4	3.4	3.0
Environmental				
Temperature, Operating (°C)	5 to 55	5 to 55	5 to 55	5 to 55
Temperature, Nonoperating (°C)	–40 to 70	–40 to 70	–40 to 70	–40 to 70
Shock, Operating: 11ms (Gs)	40	40	40	40
Shock, Nonoperating: 2ms (Gs)	400	400	400	400
Acoustics Idle (bels—sound power)	3.0	3.0	3.0	3.0
Vibration, Operating: <500Hz (Gs)	0.5	0.5	0.5	0.5
Vibration, Nonoperating: <500Hz (Gs)	3.0	3.0	3.0	3.0
Physical				
Height (in/mm, max)[4]	0.591/15.00	0.591/15.00	0.591/15.00	0.591/15.00
Width (in/mm, max)[4]	2.760/70.10	2.760/70.10	2.760/70.10	2.760/70.10
Depth (in/mm, max)[4]	3.955/100.45	3.955/100.45	3.955/100.45	3.955/100.45
Weight (lb/kg)	0.467/0.212	0.487/0.221	0.489/0.222	0.479/0.217
Carton Unit Quantity	30	30	30	30
Cartons per Pallet	50	50	50	50
Cartons per Layer	10	10	10	10
Warranty				
Limited Warranty (years)	5	5	5	5

Flash Storage

- Nonvolatile semiconductor storage
  - 100× – 1000× faster than disk
  - Smaller, lower power, more robust
  - But more $/GB (between disk and DRAM)

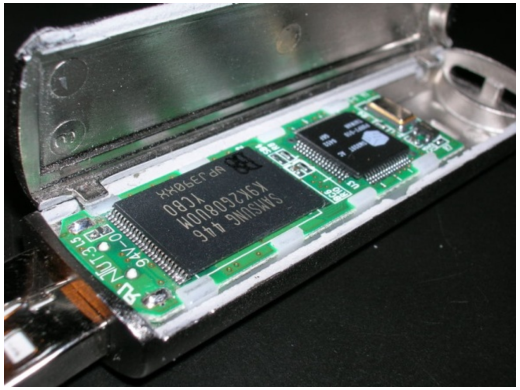

Flash Types

- NOR flash: bit cell like a NOR gate
  - Random read/write access
  - Used for instruction memory in embedded systems
- NAND flash: bit cell like a NAND gate
  - Denser (bits/area), but block-at-a-time access
  - Cheaper per GB
  - Used for USB keys, media storage, …
- Flash bits wears out after 1000's of accesses
  - Not suitable for direct RAM or disk replacement
  - Wear leveling: remap data to less used blocks

Flash key parameters

Table 1: Key parameters of the target large block NAND flash memory (K9WBG08U1M)

NAND flash memory organization	Block size Page size Number of pages per block	128 KB 2 KB 64
Access time for each operation	Read operation (1 page) Write operation (1 page) Erase operation (1 block)	25 usec 200 usec 2000 usec

- NAND flash memory consists of multiple blocks, and each block is composed of multiple pages.
- Each page is a unit of read and write operation, and each block is a unit of erase operation.
- Unlike a traditional hard disk drive, flash memory does not support overwrite operations because of its write-once nature.
- When the data at a specific page is modified, the new data value is written to another empty page and the page with the old data should be invalidated.

Flash Translation Layer

- This special feature of flash memory requires two storage management schemes.
- First, we need to provide an address mapping scheme, which maps the logical address from the file system to the physical address in flash memory by maintaining an address mapping table.
- Second, we need a garbage collection scheme to reclaim the invalidated pages.
- The garbage collection scheme should select a block which has many invalid pages and erase the block to be reused after migrating the valid pages in the block to a clean block.
- In order to support these two management tasks, a flash translation layer (FTL) is commonly used between the file system and flash memory devices.

Flash Types

- SLC/MLC/TLC
- Data retention time
- Flash cell endurance: maximum # of program/erase (p/e) cycles per physical sector
  - SLC/MLC/TLC = 50k/ 3~5k/ 1k
- Write amplification factor (WAF): Flash devices write in full blocks, which means that in order to write to a block that may already contain some data the Flash controller must move the existing data in the block (usually to main memory) and combine it with the new data and write all the data back to the Flash memory.
  - WAF=1 for hard disks
  - WAF=10 for low-end SSD
  - WAF=1.1 for X25-M SSD (Intel claims)
- Automatic bad sector remapping

Flash Translation Layer

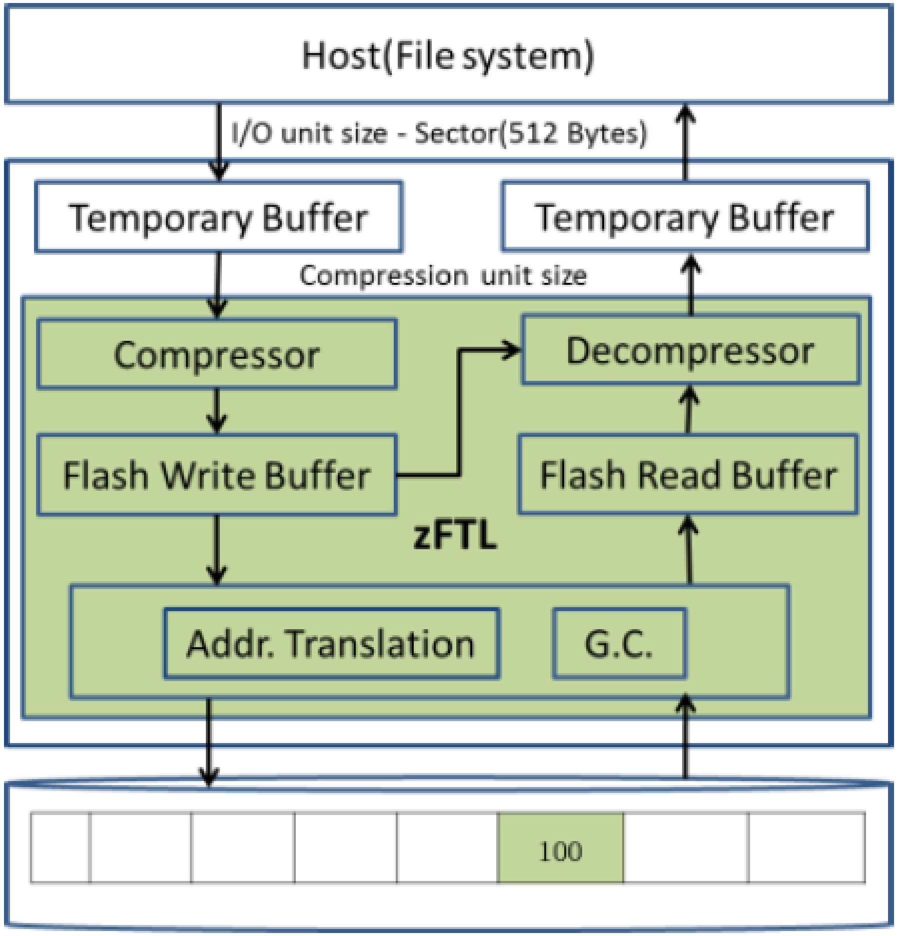

Figure 1. System architecture

Flash Storage Products

Characteristics	Kingston SecureDigital (SD) SD4/8 GB	Transend Type I CompactFlash TS16GCF133	RiDATA Solid State Disk 2.5 inch SATA
Formatted data capacity (GB)	8	16	32
Bytes per sector	512	512	512
Data transfer rate (read/write MB/sec)	4	20/18	68/50
Power operating/standby (W)	0.66/0.15	0.66/0.15	2.1/—
Size: height × width × depth (inches)	0.94 × 1.26 × 0.08	1.43 × 1.68 × 0.13	0.35 × 2.75 × 4.00
Weight in grams (454 grams/pound)	2.5	11.4	52
Mean time between failures (hours)	> 1,000,000	> 1,000,000	> 4,000,000
GB/cu. in., GB/watt	84 GB/cu.in., 12 GB/W	51 GB/cu.in., 24 GB/W	8 GB/cu.in., 16 GB/W
Best price (2008)	~ $30	~ $70	~ $300

FIGURE 6.6 Characteristics of three flash storage products. The CompactFlash standard package was proposed by Sandisk Corporation in 1994 for the PCMCIA-ATA cards of portable PCs. Because it follows the ATA interface, it simulates a disk interface, including seek commands, logical tracks, and so on.
The RiDATA product imitates an SATA 2.5-inch disk interface.

Flash Types

Characteristics	NOR Flash Memory	NAND Flash Memory
Typical use	BIOS memory	USB key
Minimum access size (bytes)	512 bytes	2048 bytes
Read time (microseconds)	0.08	25
Write time (microseconds)	10.00	1500 to erase + 250
Read bandwidth (MBytes/second)	10	40
Write bandwidth (MBytes/second)	0.4	8
Wearout (writes per cell)	100,000	10,000 to 100,000
Best price/GB (2008)	$65	$4

FIGURE 6.7 Characteristics of NOR versus NAND flash memory in 2008. These devices can read bytes and 16-bit words despite their large access sizes.

Samsung SSD 840 (250 GB)

- 21 nm TLC NAND flash
- 6 Gbps serial ATA drive
- Max. p/e cycles = 1k
- Read/Program/Erase= 75us/900~1350us/4.5ms
- Sequential reads: 540 MB/sec
- Sequential writes: 450 MB/sec
- 100k IOPS for reads
- 78k IPOS for writes
- ~$600

Intel 335 SSD

- 20 nm NAND flash
- 6 Gbps serial ATA drive
- 240 GB capacity
- Sequential reads: 500 MB/sec
- Sequential writes: 450 MB/sec
- Jointly developed with IM Flash technology
- 42k IOPS for reads
- 52k IPOS for writes
- 2.5-in form factor 99.5 mm case
- Hi-k/metal gate, planar cell