


Secondary Memory Management

- Motivation
 - Main memory is small (expensive) and volatile
 - Secondary memory is large (cheap) and persistent
 - Typically disks
- Operational basis for important OS components
 - Swapping
 - Virtual memory
 - File system

Disk Drives


- Physical structure
 - Media
 - hard or flexible, fixed or removable
 - Driver (mechanical)
 - disk rotation and head positioning
 - Controller (electronic)
 - operation and host interfacing
- Technologies
 - Magnetic
 - Optic
 - Optomagnetic



Disks

- Physical structure
 - Concentric tracks divided in sectors
 - Inter-sector gaps
 - Sectors are typically formatted to be 512 bytes-long
- Logical structure
 - Unidimensional, linear array of blocks
 - Block = 1 or n sectors
- Translation

$$\text{blk} = \text{sec} + \frac{\text{\#sec}}{\text{\#tra}} \times (\text{sur} + \text{cyl} \times \frac{\text{\#tra}}{\text{\#cyl}})$$



- Partition
 - Set of contiguous disk cylinders considered by the OS as an autonomous **logical disk**


Disk Scheduling

- Disk access time parameters
 - Seek: time to move the arm to a given cylinder
 - Latency: delay until a sector passes under the head
 - Transfer: time to transfer data from the disk controller to main memory
- Disk access requests
 - Disk address + memory address + size
 - Request queue
 - Order requests gathering those for the same cylinder
 - Order requests for different cylinders to reduce seek time
- Other performance factors
 - File organization (contiguous/disperse)
 - Control info location
 - Cache

Disk Scheduling Algorithms


■ First-Come First-Served (FCFS)

Queue: 98, 183, 37, 122, 14, 124, 65, 67 Seek: 640 tracks



■ Shortest Seek Time First (SSTF)


Queue: 98, 183, 37, 122, 14, 124, 65, 67 Seek: 236 tracks



Disk Scheduling Algorithms


■ Scan (Elevator)

Queue: 98, 183, 37, 122, 14, 124, 65, 67 Seek: 208 tracks



■ Circular Scan (C-SCAN)

Queue: 98, 183, 37, 122, 14, 124, 65, 67 Seek: 326 tracks



Redundant Array of Independent Disks

- RAID 0 (stripping)
 - Each block is broken down in sub-blocks
 - Each sub-block is stored on a different disk
 - High performance
- RAID 1 (shadowing/mirroring)
 - Each block is stored twice
 - High reliability
- RAID 5 (stripping + rotating parity)
 - High performance with good reliability

