Disk Hardware

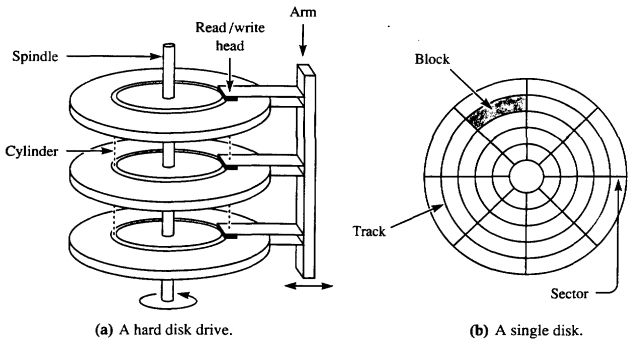

Diagram from *Computer Science*, Volume 2, J. Stanley Warford, Heath, 1991.

- Arm can move in and out
  - Read / write head can access a ring of data as the disk rotates
- Disk consists of one or more *platters*
  - Each platter is divided into rings of data, called *tracks*, and each track is divided into *sectors*
  - One particular platter, track, and sector is called a *block*

Disk Hardware (cont.)

- Typical disk today (Quantum XP39100):
  - Total capacity = 9.1GB
  - 10 platters (20 surfaces)
  - 5964 tracks per surface
  - 108-180 sectors per track
  - 512 bytes per sector
- Trends in disk technology
  - Disks are getting smaller, for similar capacity
    - Faster data transfer, lighter weight
  - Disk are storing data more densely
    - Faster data transfer
    - Density improving faster than mechanical limitations (seek time, rotational delay)
  - Disks are getting cheaper (factor of 2 per year since 1991)

Common File Access Patterns

- Sequential access
  - Data is processed in order, one byte at a time
  - Most accesses are of this form
  - Example: compiler reading a source file
- Direct / random access
  - Can access any byte in the file directly, without accessing any of its predecessors
  - Example: database
- Keyed access
  - Can access a byte based on a *key* value
  - Example: database search, dictionary
  - OS does not support keyed access
    - User program must determine the address from the key, then use random access (provided by the OS) into the file

UNIX File System

- A file descriptor (*inode*) represents a file, and contains information about:
  - File type
  - Access permissions
  - Owner
  - Size
  - Blocks where file is stored on disk
- A UNIX disk may be divided into *partitions*, each of which contains:
  - Some special blocks
    - Boot block — code for booting the system
    - Super block — size of disk, number of free blocks, list of free blocks, size of ilist, number of free inodes in ilist, etc.
  - Blocks storing the ilist (array of inodes)
  - Blocks storing directories and files

UNIX File System (cont.)

- High-level view:

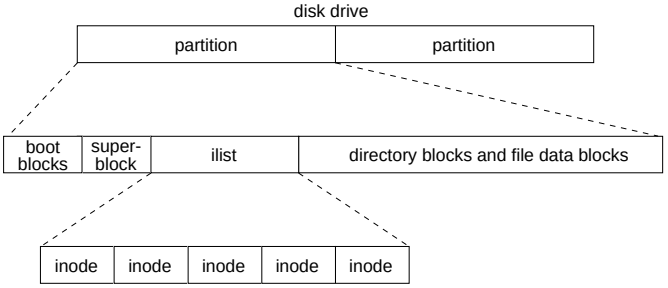

- Low-level view:

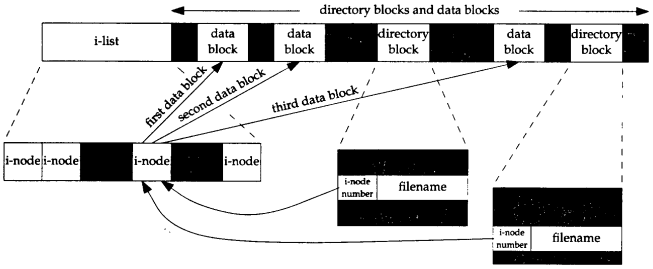

Diagram from *Advanced Programming in the UNIX Environment*, W. Richard Stevens, Addison Wesley, 1992.

Organization of Files (Contiguous and Linked/Chained Allocation)

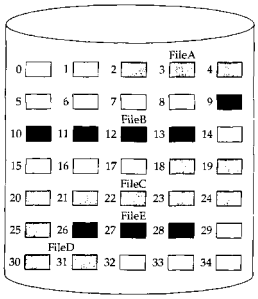

File Allocation Table

File Name	Start Block	Length
FileA	2	3
FileB	9	5
FileC	18	8
FileD	30	2
FileE	26	3

FIGURE 11.7 Contiguous file allocation

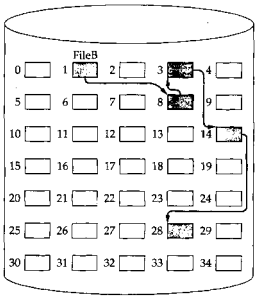

File Allocation Table

File Name	Start Block	Length
...	...	...
FileB	1	5
...	...	...

FIGURE 11.9 Chained allocation

Diagrams from *Operating Systems*, William Stallings, Prentice Hall, 1995.

Disk Hardware

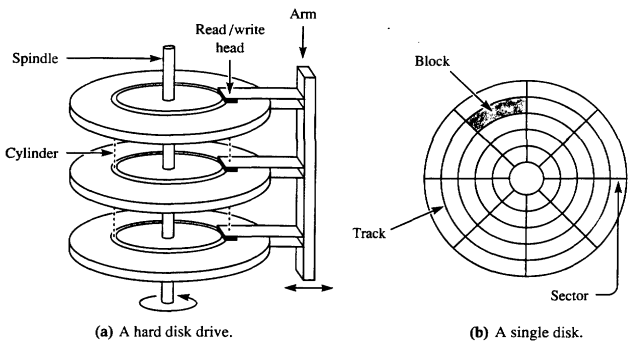

Diagram from *Computer Science*, Volume 2, J. Stanley Warford, Heath, 1991.

- *Seek time* — time required to position heads over the track / cylinder
  - Typically 10 ms to cross entire disk
- *Rotational delay* — time required for sector to rotate underneath the head
  - 120 rotations / second = 8 ms / rotation

Disk Access Times

- Some typical numbers:
  - 32-64 sectors per track
  - 1K bytes per sector
- *Data transfer rate* is number of bytes rotating under the head per second
  - 1 KB / sector * 32 sectors / rotation * 120 rotations / second = 4 MB / s
- Disk I/O time = seek + rotational delay + transfer
  - If head is at a random place on the disk
    - Avg. seek time is 5 ms
    - Avg. rotational delay is 4 ms
    - Data transfer rate for a 1KB is 0.25 ms
    - I/O time = 9.25 ms for 1KB

    ↪Real transfer rate is roughly 100 KB / s
  - In contrast, memory access may be 20 MB / s (200 times faster)

Selecting the Sector Size

- The read / write head needs to synchronize with the head as it rotates
  - Need 100-1000 bits between each sector to measure how fast disk is spinning
- If sector size is 1 byte
  - Only 1% of disk holds useful data
  - 1/1000 transfer rate as before = 100 B / s
- If sector size is 1 KB
  - 90% of disk holds useful data
  - Transfer rate is 100 KB / s
- If sector size is 1 MB
  - Almost all of disk holds useful data
  - Transfer rate is 4 MB / s (full disk transfer rate — seek and rotational latency usually won't matter anymore)