Operating Systems

(Storage Management: Disk management, swap space management, disk reliability, stable storage implementation, file concepts, directory structures, file system in Linux.)
Slides Set #19

By Prof K R Chowdhary

JNV University

January 24, 2024

Overview of Mass-Storage Structure

- The **file system** can be viewed logically consisting of three parts:
  - user and programmer interface to the file system,
  - internal data structures and algorithms used by the operating system to implement this interface,
  - file systems at the lowest level
- Disk-scheduling algorithms: Schedule the order of disk I/Os to maximize performance.
- **Magnetic Disks** (see fig.):
- A read-write head “flies” just above each surface of every platter. The heads are attached to a disk arm that moves all the heads as a unit.

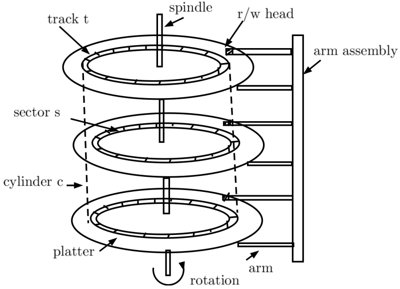

Overview of Mass-Storage Structure...

- When in use, a drive motor spins it at high speed.Common drives spin at 5,400, 7,200, 10,000, and 15,000 RPM.The positioning time, or random-access time, consists of two parts: called the **seek** time, and rotational **latency** time.
- Because the disk head flies on an extremely thin cushion of air (measured in microns), there is a danger that the head will make contact with the disk surface.
- A disk drive is attached to a computer by a set of wires called an I/O bus.
- **Solid-State Disks:** Sometimes old technologies are used in new ways as economics change or the technologies evolve.
- SSDs have the same characteristics as traditional hard disks but can be more reliable
- Because SSDs can be much faster than magnetic disk drives, standard bus interfaces can cause a major limit on throughput

Disk Management

Disk initialization, booting from disk, and bad-block recovery.

- **Disk Formatting**. Before a disk can store data, it must be divided into sectors that the disk controller can read and write. This low-level formatting – fills the disk with a special data structure.
- When the controller writes a sector of data during normal I/O, the ECC is updated with a value calculated from all the bytes in the data area.
- Most hard disks are low-level-formatted at the factory.
- Before it can write and read files, the OS needs to record its own data structures on the disk.

Disk Management...

- To increase efficiency, most file systems group blocks together into larger chunks, frequently called **clusters**.
- Some OS treats the disk partition as a large sequential array of logical blocks,
- **Boot Block.** For a computer to start running for instance, when it is powered up or rebooted – it must have an initial program to run. This program is **bootstrap program**.
- For most computers, the bootstrap is stored in read-only memory (ROM).
- The code in the boot ROM instructs the disk controller to read the boot blocks into memory and then starts executing that code.
- Windows allows a hard disk to be divided into partitions, and one partition – identified as the boot partition – contains the operating system and device drivers.

Swap-Space Management

- Swapping occurs when the amount of physical memory reaches a critically low point and processes are moved from memory to swap space to free available memory.
- Swap-space management is low-level task. Virtual memory uses disk space as an extension of main memory.
- **Swap-Space Use**. Swap space is used by OS depending on the memory-management algorithms in use.
- It may be safer to **overestimate** than to underestimate the amount of swap space required,

Swap-Space Management...

- **Bad Blocks**. Because disks have moving parts and small tolerances,they are prone to failure.
- Swapping occurs when the amount of physical memory reaches a critically low point
- Very few modern operating systems implement swapping like this. Systems now combine swapping with virtual memory techniques and swap pages,
- Swap-space management is another low-level task of the operating system.
- **Swap-Space Use**: Swap space is used in various ways by different operating systems. E.g., to hold an entire process image,

Swap-Space Management...

- It is better to overestimate the amount of swap space required,
- Some systems recommend the amount to be set aside for swap space. Solaris, e.g., sets swap space equal to the amount by which virtual memory **exceeds** physical memory.
- In the past, Linux suggested setting swap space to double the amount of physical memory.
- **Swap-Space Location**. A swap space can reside in one of two places: it can be in the normal file system, or it can be in a separate disk partition.
- Alternatively, swap space can be created in a separate raw partition.

Swap-space management

- An Example. The traditional UNIX kernel copied entire processes between contiguous disk regions and memory.
- **Stable-Storage Implementation**. Information residing in stable storage is never lost.

  A disk write results in one of three outcomes:

  1. **Successful completion**.
  2. **Partial failure**.
  3. **Total failure**.
- Whenever a failure occurs during writing of a block, the system needs to detect it and invoke a recovery procedure to restore the block