

Learning Objectives

- Understanding disk organization.
- Sectors, clusters and extents.
- Fragmentation.
- Disk access time.
- Improving disk access performance.

Secondary Storage Management


- Secondary storage devices:
 - have much longer access time than main memory
 - have access times that vary from one access to another (some accesses are relatively fast and other accesses are slower on the same device)
 - have a lot of more storage than main memory
 - have storage that is non-volatile

Disks

- Types of commonly used disks
 - hard disks
 - floppy disks
 - Iomega ZIP disks
 - Jaz disks

The Organization of Disks

- Data is stored on the surface of one or more platters.
- Disk storage units are:
 - tracks and cylinders
 - sectors



Disk Storage Capacity

- The amount of data that can be held on a disk depends on
 - how densely bits can be stored on the disk surface
- The capacity of the disk is a function of:
 - the number of cylinders
 - the number of tracks per cylinder
 - the capacity of a track

Track capacity = number of sectors per track × bytes per sector

Cylinder capacity = number of tracks per cylinder × track capacity

Drive capacity = number of cylinders × cylinder capacity

How is the Data Read from or Written to a Disk?

- The operating system sends control signals to the disk via a disk driver to read or to write data from a given sector of a given cylinder.
- The disk is rotating to position the needed sector under the *read/write* head (*rotational delay*)
- The read/write head is moving to the needed cylinder (*seek time*).

Specification of Disk Drives

- Capacity: e.g. 2 GB
- Minimum (track to track) seek time: e.g. 1 msec
- Average seek time : e.g. 12 msec (milliseconds)
- Maximum seek time: e.g. 22 msec
- Spindle speed: e.g. 5200 rpm (rotations per minute)
- Average rotational delay: e.g. 6 msec
- Mximum transfer rate: e.g. $2796 \text{ bytes/msec} = 2730 \text{ K/sec}$
- Bytes per sector: e.g. 512
- Sectors per track: e.g. 63
- Tracks per cylinder: e.g. 16
- Cylinders: 4092

Organizing Data by Sectors

- Consecutive physical sectors sometimes are not consecutive logically
 - this is called sector interleaving
- In the early 1990s, controller speeds improved so that disks can now offer non-interleaving (also known as 1:1 interleaving)

Clusters

- A cluster is a fixed number of consecutive (logical) disk sectors.
- Some operating systems view each file as a series of clusters.
- Clusters are designed to improve performance since all sectors in one cluster can be accessed without an additional seek.

Extents

- Extents of a file are those parts of the file which are stored in contiguous clusters.
- It is very beneficial to store the whole file in one extent (seek time is minimized).

Fragmentation

- Fragmentation is the wasted disk space due to the fact that the smallest organizational unit of a disk is one sector.
- If a sector size is 512 bytes then even if we need to store only one byte, we have to allocate to it one whole sector. Thus 511 bytes are wasted.

Blocks

- Some disk allow for storing data in *user defined blocks* instead of sectors.
- When the data on a disk is organized in blocks, this usually means that the amount of data transferred in a single I/O operation can vary.
- Blocks can be either variable or fixed length.
- Block organization can be more efficient than sector organization but it is much more complex.

Non-data Overhead

- Non-data overhead includes at the beginning of each sector:
 - sector address
 - track address
 - sector usability

The Cost of Disk Access

- Seek time
 - the time required to move the r/w head to the correct cylinder
- Rotational delay
 - the time required to rotate the disk so that the correct sector is positioned under the r/w head
- Transfer time
 - the time required to transfer the data:

$$\text{Transfer time} = \frac{\text{number of bytes transferred}}{\text{number of bytes on a track}} \times \text{rotation time}$$

Disks as Bottlenecks

- Disk speeds lag far behind
 - CPU
 - main memory
 - local network
- Computer programs spend most of time awaiting data from the disk

Improving Disk Performance

- Disk striping
 - splitting the parts of a single file on several drives
- RAID
 - Redundant Array of Inexpensive Disks
- RAM disk
- Disk caching
- Buffering