

SECONDARY STORAGE AND SYSTEM SOFTWARE - II

3.1.5 Organizing Tracks by Block

- Mainframe computers typically use variable size physical blocks for disk drives.
- Track capacity is dependent on block size, due to fixed overhead (gap and address block) per block.

3.1.6 The Cost of a Disk Access

direct access device

A data storage device which supports direct access.

direct access

Accessing data from a file by record position within the file, without accessing intervening records.

access time

The total time required to store or retrieve data.

transfer time

The time required to transfer the data from a sector, once the transfer has begun.

seek time

The time required for the head of a disk drive to be positioned to a designated cylinder.

rotational delay

The time required for a designated sector to rotate to the head of a disk drive.

- Access time of a disk is related to physical movement of the disk parts.
- Disk access time has three components: seek time, rotational delay, and transfer time.
- Seek time is affected by the size of the drive, the number of cylinders in the drive, and the mechanical responsiveness of the access arm.
- Average seek time is approximately the time to move across 1/3 of the cylinders.
- Rotational delay is also referred to as *latency*.
- Rotational delay is inversely proportional to the rotational speed of the drive.
- Average rotational delay is the time for the disk to rotate 180°.
- Transfer is inversely proportional to the rotational speed of the drive.
- Transfer time is inversely proportional to the physical length of a sector.
- Transfer time is roughly inversely proportional to the number of sectors per track.
- Actual transfer time may be limited by the disk interface.

3.1.8 Effect of Block Size

- Fragmentation waste increases as cluster size increases.
- Average access time decreases as cluster size increases.

3.1.9 Disk as a bottleneck

striping

The distribution of single files to two or more physical disk drives.

Redundant Array of Inexpensive Disks

An array of multiple disk drives which appears as a single drive to the system.

RAM disk

A virtual disk drive which actually exists in main memory.

solid state disk

A solid state memory array with an interface which responds as a disk drive.

cache

Solid state memory used to buffer and store data temporarily.

- Several techniques have been developed to improve disk access time.
- Striping allows disk transfers to be made in parallel.
- There are 6 versions, or levels, of RAID technology.
- RAID-0 uses striping.
- RAID-0 improves access time, but does not provide redundancy.
- RAID-1 uses mirroring, in which two drives are written with the same data.
- RAID-1 provides complete redundancy. If one drive fails, the other provides data backup.
- RAID-1 improves read access time, but slows write access time.
- RAM disks appear to programs as fast disk drives.
- RAM disks are volatile.
- Solid state disks appear to computer systems as fast disk drives.
- Solid state disks are used on high performance data base systems.
- Caching improves average access time.
- Disk caching can occur at three levels: in the computer main memory, in the disk controller, and in the disk drive.
- Windows operating systems use main memory caching.
- Disk controller caching requires special hardware.
- Most disk drives now contain caching memory.
- With caching, writes are typically reported as complete when the data is in the cache.
- The physical write is delayed until later.
- With caching, reads typically read more data than is requested, storing the unrequested data in the cache.
- If a read can be satisfied from data already in the cache, no additional physical read is needed.
- Read caching works on average because of program locality.