









Disk – Logical Structure

Older/smaller disks:

- Same number of sectors for all cylinders
- Constant angular velocity (CAV)
- Data more dense packed on inner cylinders

Newer/larger disks:

01011. C. Kessler, IDA. Linköpings universitet

- Different numbers of sectors for different groups of cylinders
- Still CAV
- Smaller difference between dense and sparse cylinders







Disk – Access Time



Both in the order of several ms

Disk access time has 2 major components:

- Seek time time to move R/W-heads to right cylinder
 Rotational latency time for the right block to appear
- Rotational speeds: ~400-12000 rpm and above

Example: Read a file of 128kB (= 256 blocks of 512 bytes)

- 12000 rpm = 200 rps → 5 ms/round
- Disk has 32 sectors (1 block/sector) → 5 ms / 32 = 156.25 µs/block to read
- Average seek time: 10ms
- Average rotational latency (12000rpm): 2.5ms
- File A all blocks in sequence: 10+2.5+256*0.156 = 52.5 ms

• File B – blocks scattered over disk: 256*(10+2.5+0.156) = 3240 ms











SSTF

Selects the request with the shortest seek time from the current head position.

May cause starvation of some requests.











Disk Management

- Low-level formatting, or physical formatting Dividing a disk into sectors that the disk controller can read and write.
- To use a disk to hold files, the OS still needs to record its own data structures on the disk.
 - *Partition* the disk into one or more groups of cylinders.
 - Logical formatting or "making a file system".
- Boot block initializes system.
 - The bootstrap is stored in ROM.
 - Bootstrap loader program.
- Methods such as sector sparing used to handle bad blocks.

Swap-Space Management

- Swap-space Virtual memory uses disk space as an extension of main memory
 - Less common now due to memory capacity increases
- Swap-space can be carved out of the normal file system, or, more commonly, it can be in a separate disk partition (raw)
- Swap-space management

DIU11, C. Kessler, IDA, Linköpings universitet.

- BSD allocates swap space when process starts; holds text segment (the program) and data segment
- Kernel uses swap maps to track swap-space use
- Solaris 2 allocates swap space only when a dirty page is forced out of physical memory
- What if a system runs out of swap space?
- Some systems allow multiple swap spaces
 Kessler, IDA, Linköpings universitet.





Operating System Issues

- Major OS jobs are
 - to manage physical devices and
 - to present a virtual machine abstraction to applications
- For hard disks, the OS provides two abstraction:
 - Raw disk an array of data blocks, no file system
 - Used exclusively and managed by an application program e.g., some database systems

- RAID administrative information
- Swap space (faster than if realized as a single large file)
- Virtual memory backing store
- "Cooked disk" / File system OS queues and schedules the interleaved requests from several applications.

1, C. Kessler, IDA, Linköpings universitet. 8.







Solid-State Disks	
Nonvolatile memory used like a hard drive	
 Many technology variations 	
Can be more reliable than HDDs	
More expensive (8 to 9 times) per MB	
Maybe have shorter life span	
Less capacity	
But much faster	
No moving parts, so no seek time or rotational latency	











RAID – Other Features

- Regardless of where RAID implemented, other useful features can be added
- Snapshot is a view of file system before a set of changes take place (i.e. at a point in time)
- Replication is automatic duplication of writes between separate sites
 - For redundancy and disaster recovery
 - Can be synchronous or asynchronous
- Hot spare disk is unused, automatically used by RAID production if a disk fails to replace the failed disk and rebuild the RAID set if possible
 - Decreases mean time to repair

DIU11, C. Kessler, IDA, Linköpings universitet.





