


Magnetic/Hard Disk (Movable head)

(Track #, Surface #, Sector #) "Physical" View of Disk



Steps to perform a disk read (or write):

- 1) Move Read/Write (R/W) heads over correct track/cylinder ← seek time
- 2) Wait for start of specified sector to spin to R/W head ← rotational delay
- 3) Read (or write) specified sector as it spins under the R/W head ← data transfer time


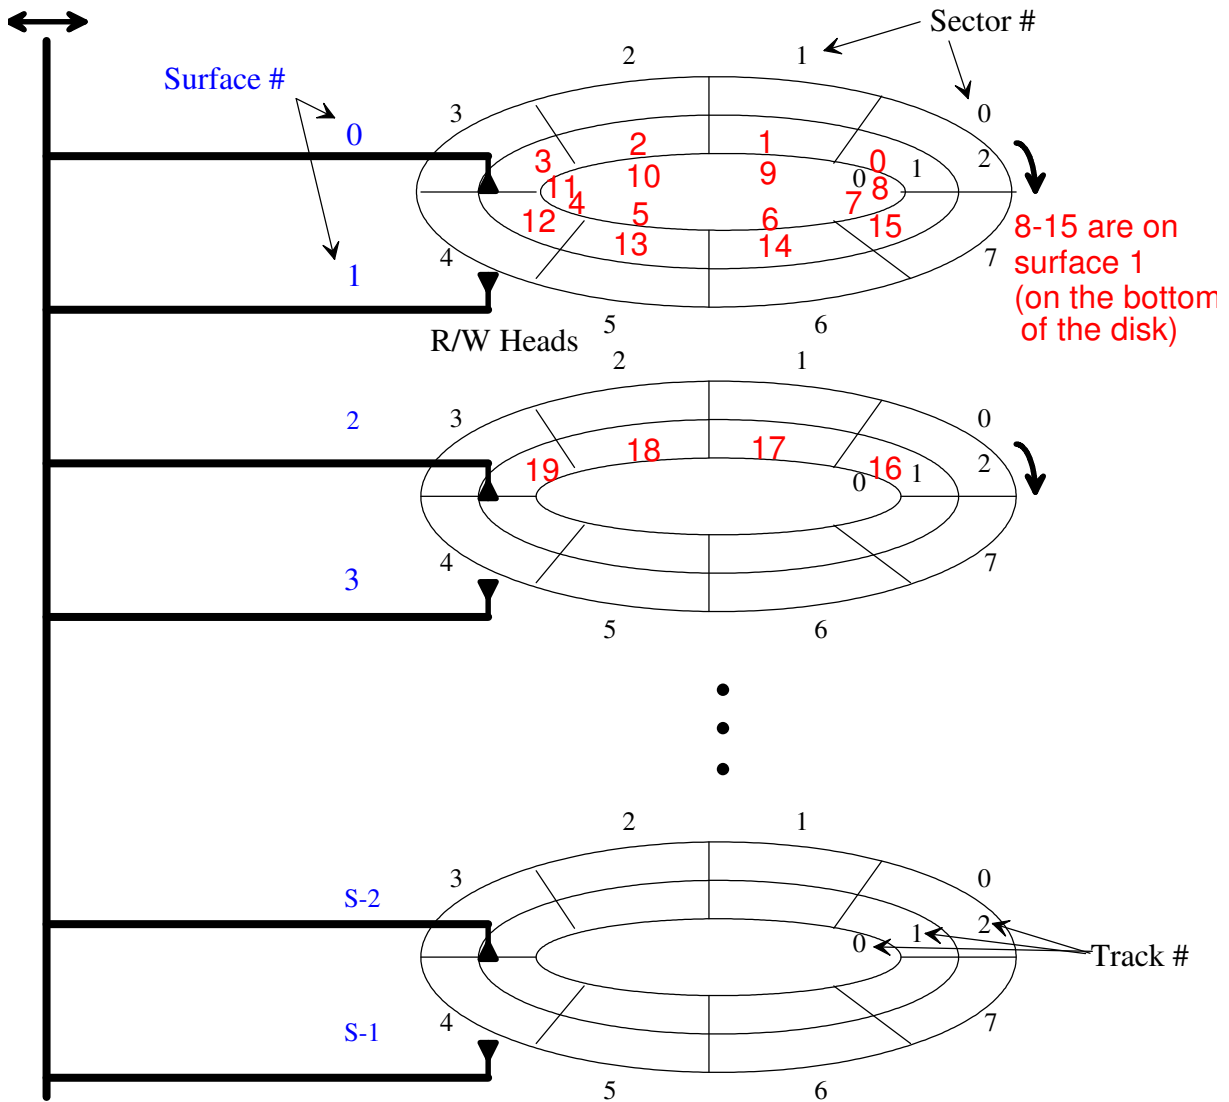



Figure 6.7 Timing of Disk I/O Transfer

Logical View of Disk as Linear Collection of Blocks




(track #, surface #, sector #) to Linear block # mapping




← All of cylinder 0 → × ← All of cylinder 1 →

Bits of linear block # : track # surface # sector #

Figure 6.3 Disk Layout Methods



(a) Constant angular velocity




(b) Multiple zoned recording


RAID - Redundant Array of Independent/Inexpensive Disks

- 6 levels in common use
- Not a hierarchy
- Set of physical disks viewed as single logical drive by O/S
- Data distributed across physical drives
- Can use redundant capacity to store parity information
- Selection of RAID level can optimize for **one** of the following:
 - (1) fast response time for many small I/O requests, e.g., database access, or
 - (2) very high data-transfer rates for single large I/O request, e.g.,


Figure 6.8 RAID Levels




(a) RAID 0 (non-redundant)




(b) RAID 1 (mirrored)




(c) RAID 2 (redundancy through Hamming code)




(d) RAID 3 (bit-interleaved parity)



(e) RAID 4 (block-level parity)



(f) RAID 5 (block-level distributed parity)



(g) RAID 6 (dual redundancy)