Hard Disk Drives Systems

Last updated 2/15/24

Hard Disk Drive - Systems

- Electronics

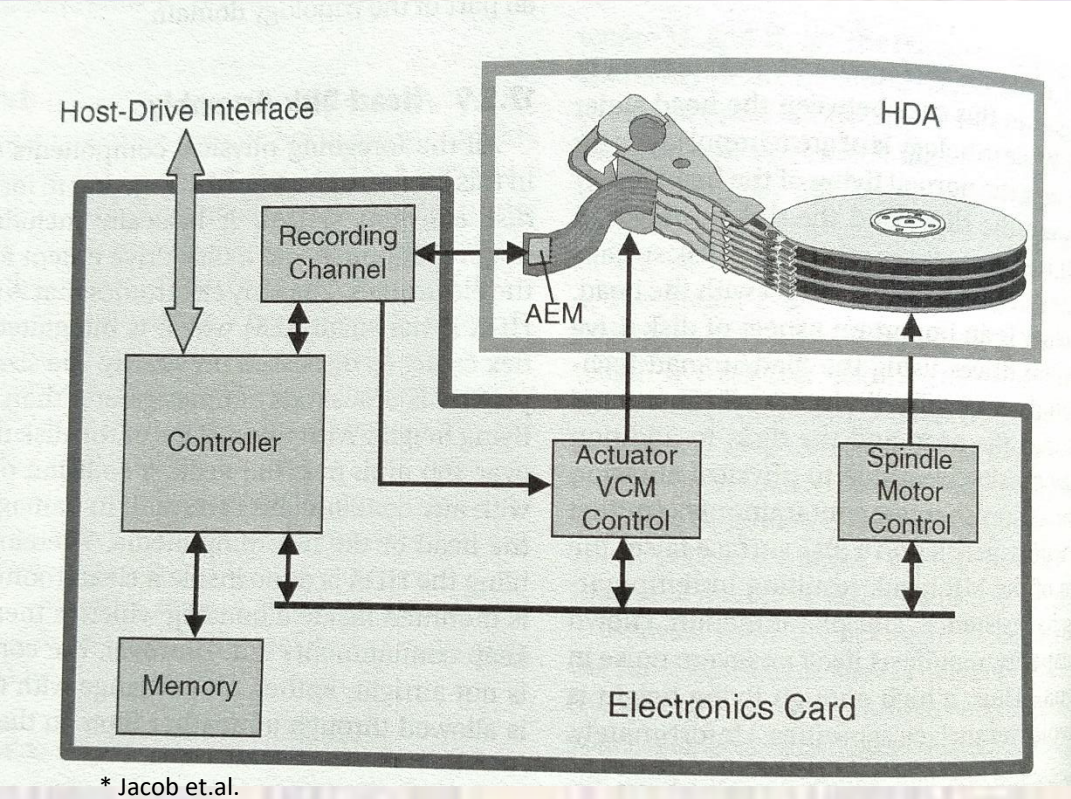

* Jacob et.al.

HDA – Head Disk Assembly
AEM – Arm Electronics Module
VCM – Voice Coil Module

© tj

Hard Disk Drive - Systems

- Controller

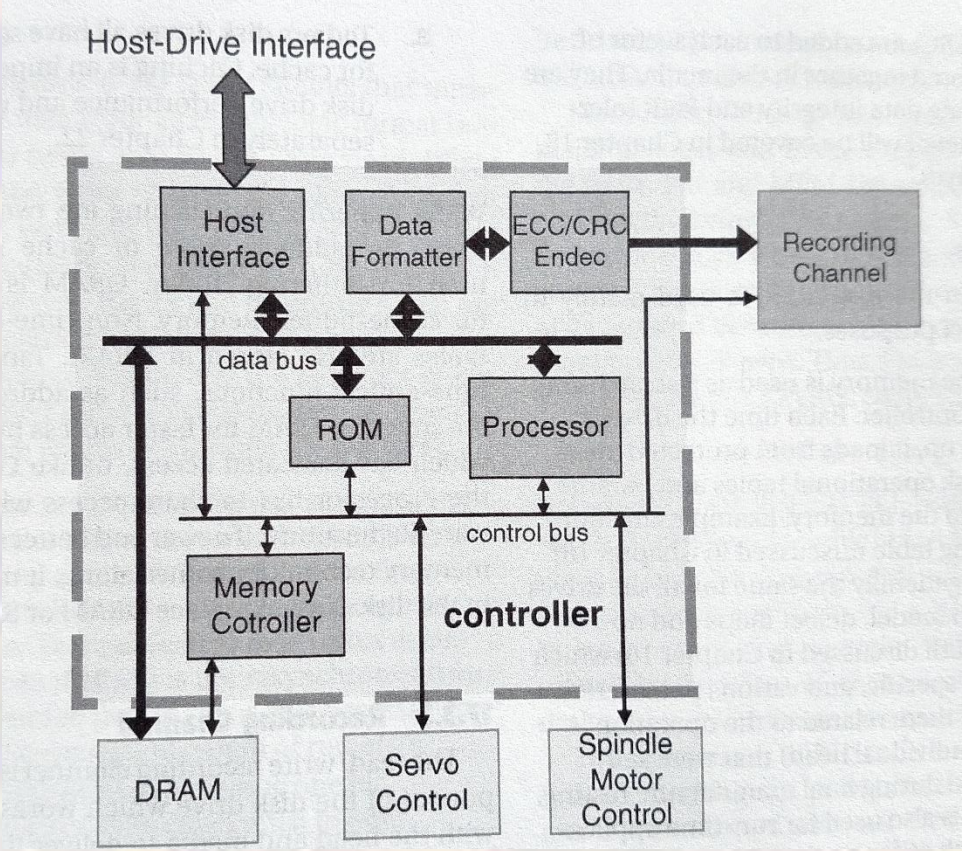

* Jacob et.al.

© tj

Hard Disk Drive - Systems

- Controller

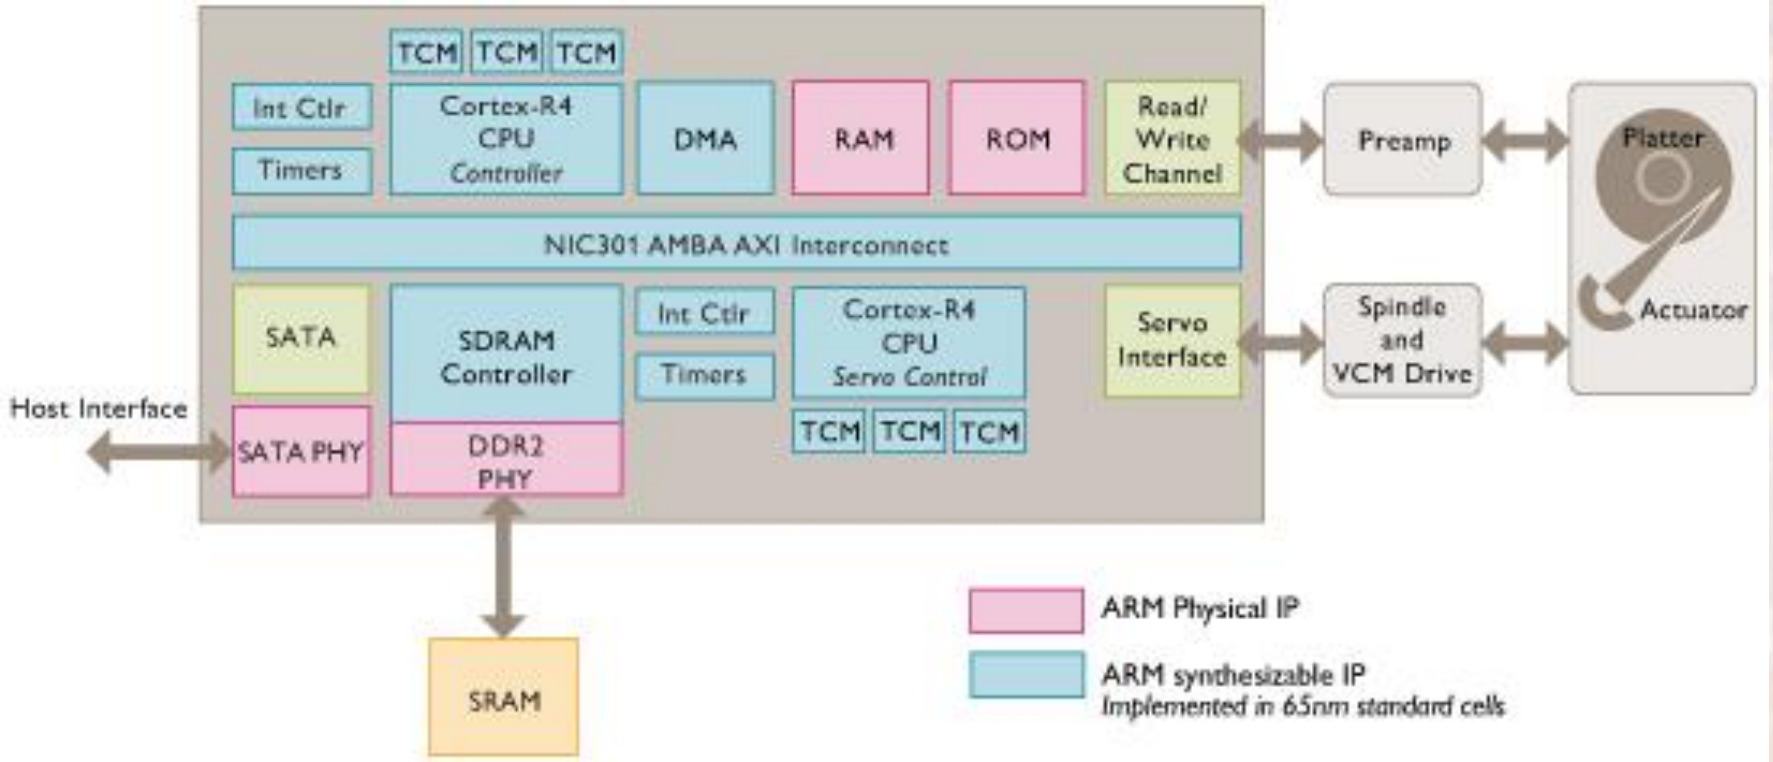

© tj

Hard Disk Drive - Systems

- Controller
  - Processor
    - Microcontroller (Arm Mx)
    - Manages the actions of the HDD
  - ROM
    - Stores firmware
  - Memory Controller
    - Manages the DRAM interface
    - DMA
    - Cache controller

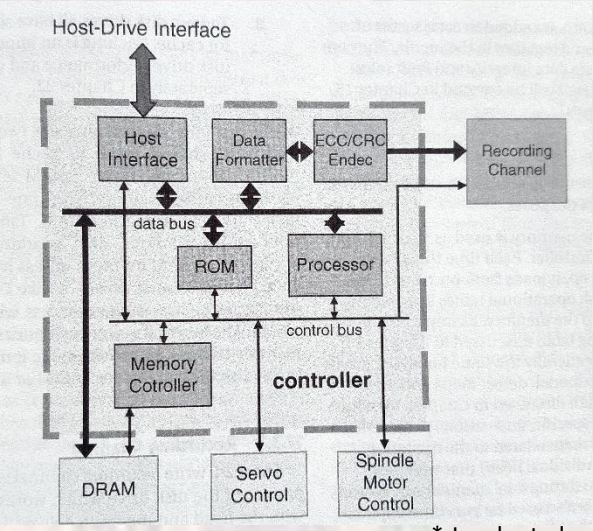

* Jacob et.al.

© tj

Hard Disk Drive - Systems

- Controller
  - Host Interface
    - Manages the external interface
    - Control registers
    - IDE, PATA, SCSI, SATA, SAS, USB
  - Data Formatter
    - Moves data to/from memory
    - Manages sector size
  - ECC/CRC
    - Adds error checking and correction bits
    - Checks for errors and performs correction

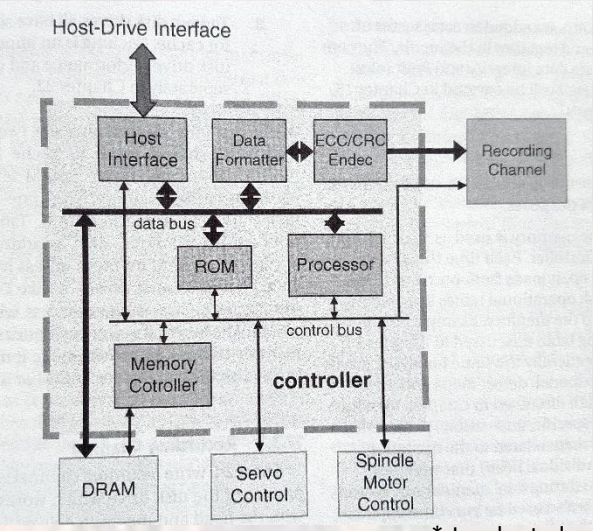

* Jacob et.al.

© tj

Hard Disk Drive - Systems

- Controller
  - DRAM
    - Processor operational memory
    - Buffer memory for R/W process
    - Disk Cache

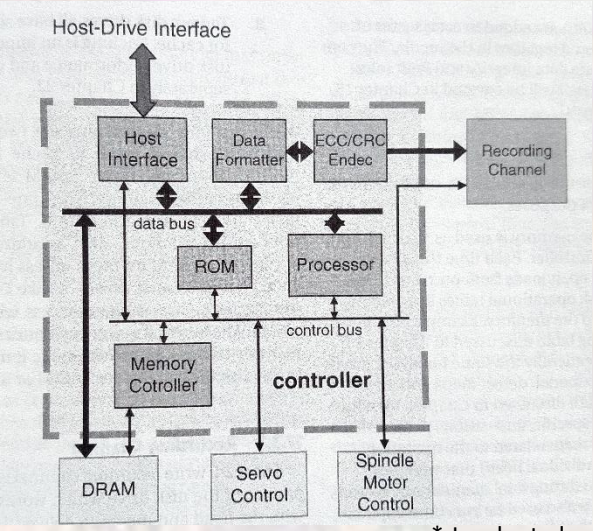

* Jacob et.al.

 © tj

Hard Disk Drive - Systems

- Recording Channel – write mode

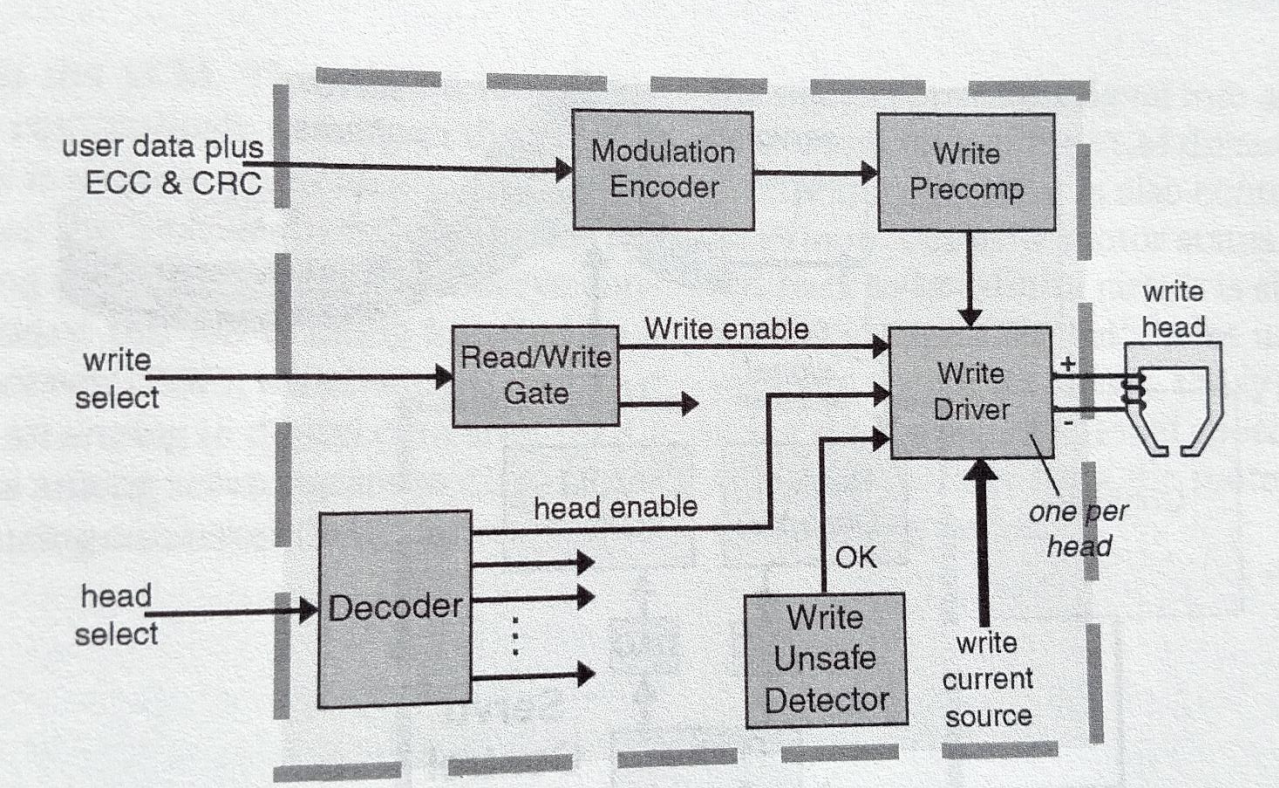

* Jacob et.al.

© tj

Hard Disk Drive - Systems

- Recording Channel – write mode

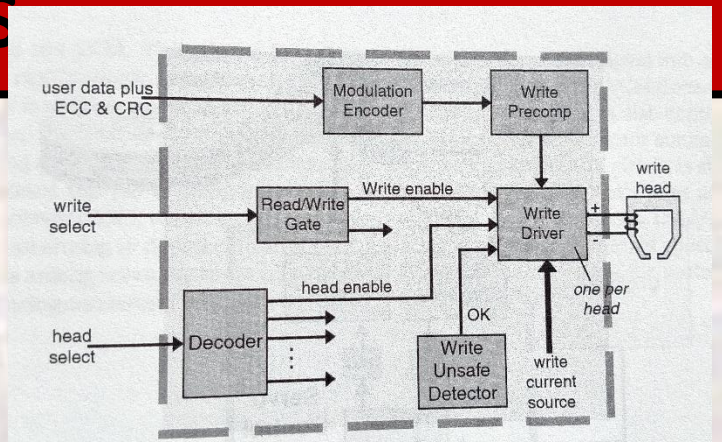

  - Modulation Encoding

    - Encodes the data to meet certain requirements
      - Sufficient transitions to allow clock recovery on reads
      - Limit errors on 1 bit from propagating indefinitely
      - Provide high data to coding ratio

    - NRZI – non-return to zero inverted
      - 0 represented by no transition
      - 1 represented by transition
      - Lacks any limit on 0's in a row → loss of clock

© tj

Hard Disk Drive - Systems

- Recording Channel – write mode

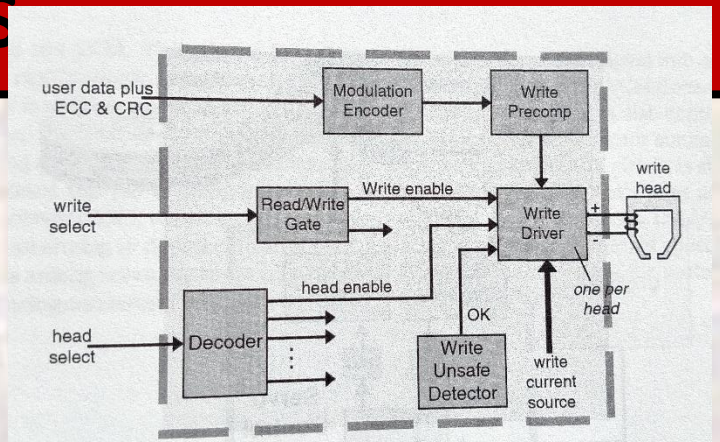

  - Modulation Encoding

    - RLL Codes – Run Length Limited
      - Limits the number of consecutive 0's or 1's
      - m/n(d,k)
        - m = # of data bits
        - n = # of encoded bits
        - d = minimum # of 0's (N's) required between two 1's (T's)
        - k = maximum # of 0's (N's) allowed in a row
      - Data Rate (DR) = (d+1)*m/n

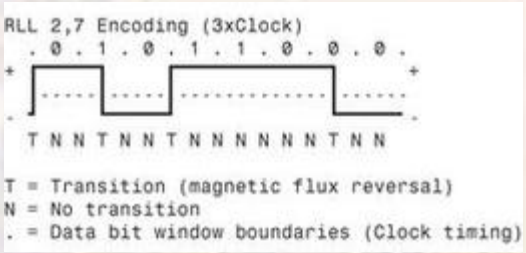

RLL 2,7 Data-to-Flux Transition Encoding	
Data Bit Values	Flux Encoding
10	NTNN
11	TNNN
000	NNNTNN
010	TNNTNN
011	NNTNNN
0010	NNTNNTNN
0011	NNNNTNNN
T = Flux transition, N = No flux transition	

© tj

Hard Disk Drive - Systems

- Recording Channel – write mode

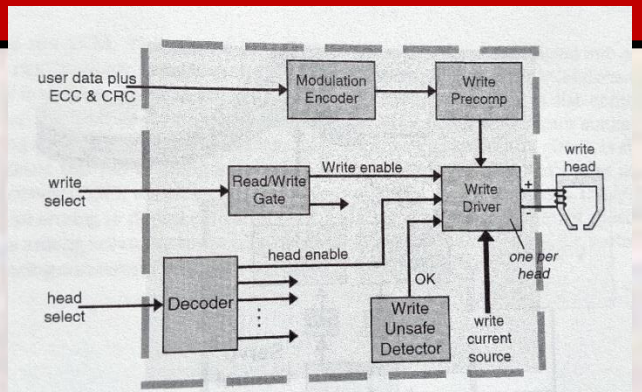

  - Write Pre-compensation (equalization)

    - Reduce Inter-Symbol-Interference

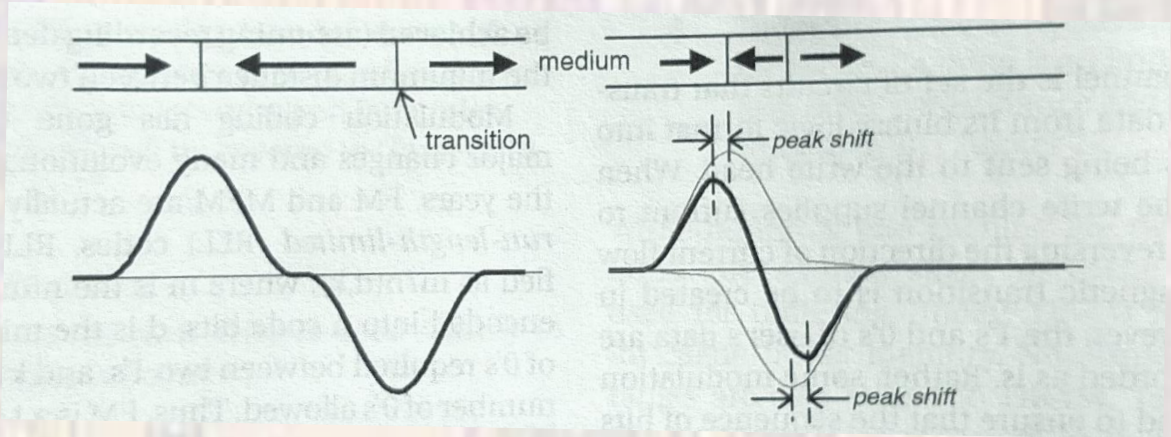

    - Delay 1st transition and speed up 2nd transition

© tj

Hard Disk Drive - Systems

- Recording Channel – write mode
  - Logic
    - Read/write
    - Correct track
    - Which head

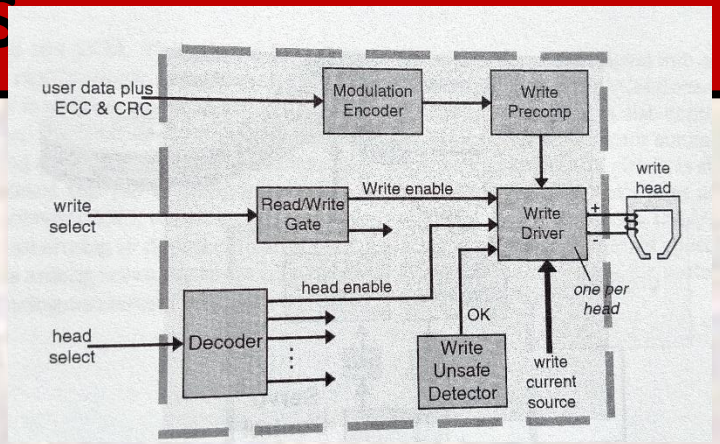

* Jacob et.al.

© tj

Hard Disk Drive - Systems

- Recording Channel – read mode

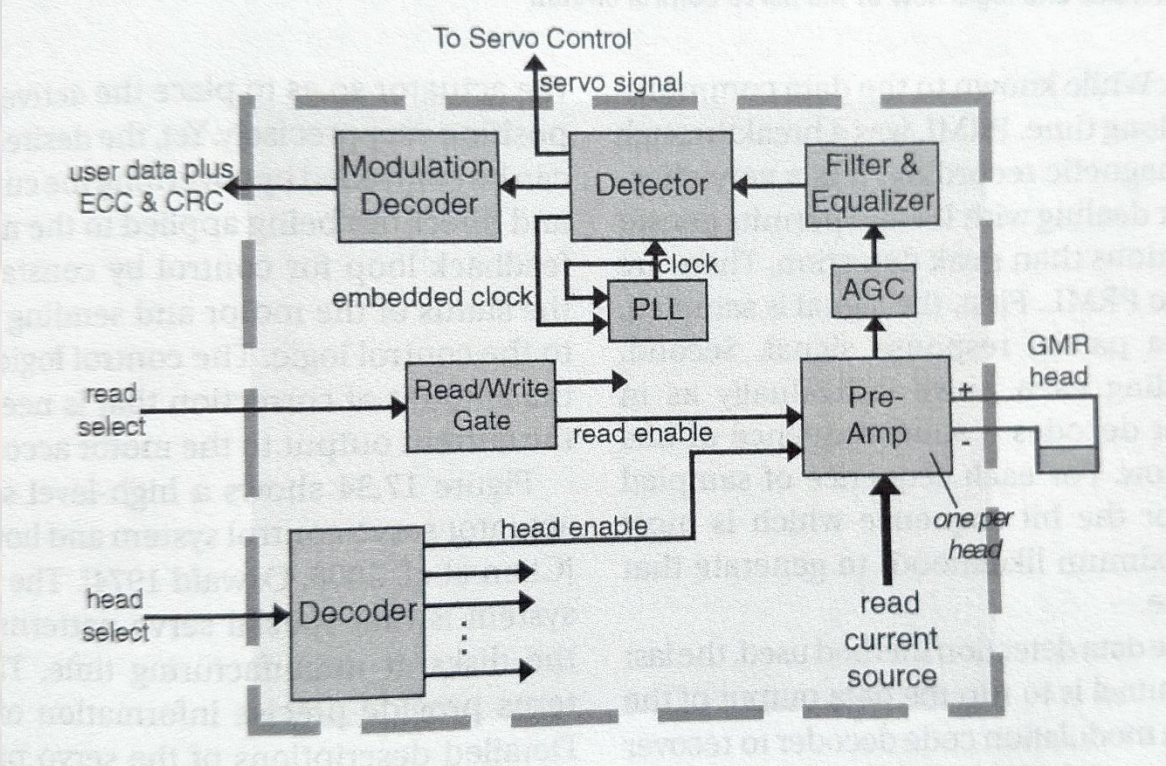

* Jacob et.al.

© tj

Hard Disk Drive - Systems

- Recording Channel – read mode
  - Pre-amp
    - Typical GMR signals < 1mV
    - Add gain to get a manageable signal level
  - AGC
    - Automatic gain control
    - Set peaks to a given desired value
  - Filter & Equalizer
    - Reduce high frequency noise
    - Sharpen pulses

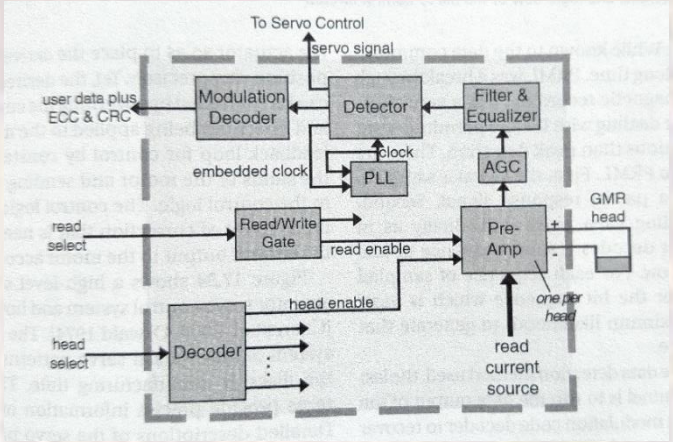

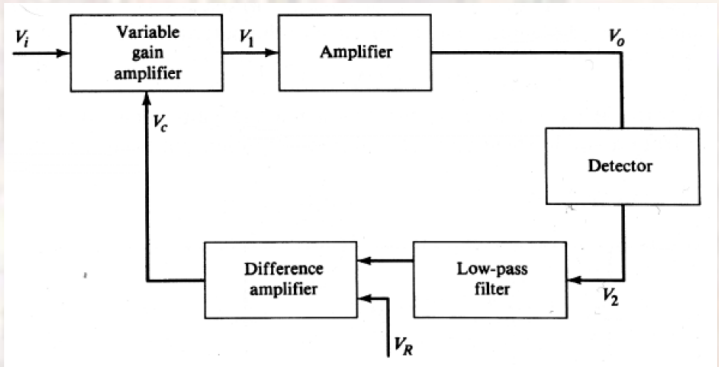

© tj

Hard Disk Drive - Systems

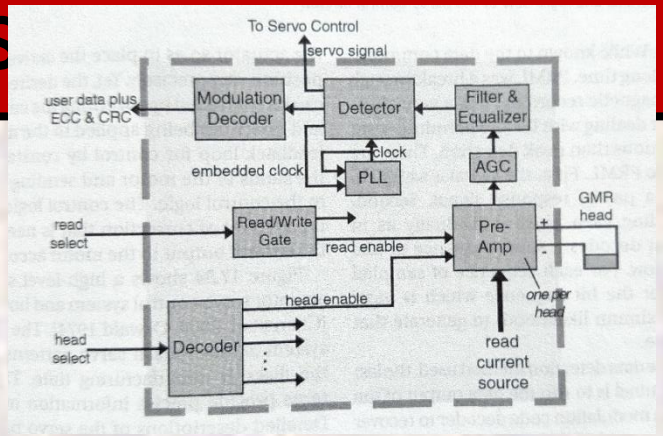

- Recording Channel – read mode
  - Detector
    - PRML – partial response maximum likelihood
      - Sample signal – partial response
      - Look at several bits worth of samples at a time
      - Choose the most likely bit pattern – maximum likelihood
    - Pick off the servo bits → Servo Controller
    - Use all bits → PLL → Clock
    - Pick off data bits → Decoder
  - Decoder
    - Reverse the RLL encoding

© tj

Hard Disk Drive - Systems

- Servo Controller

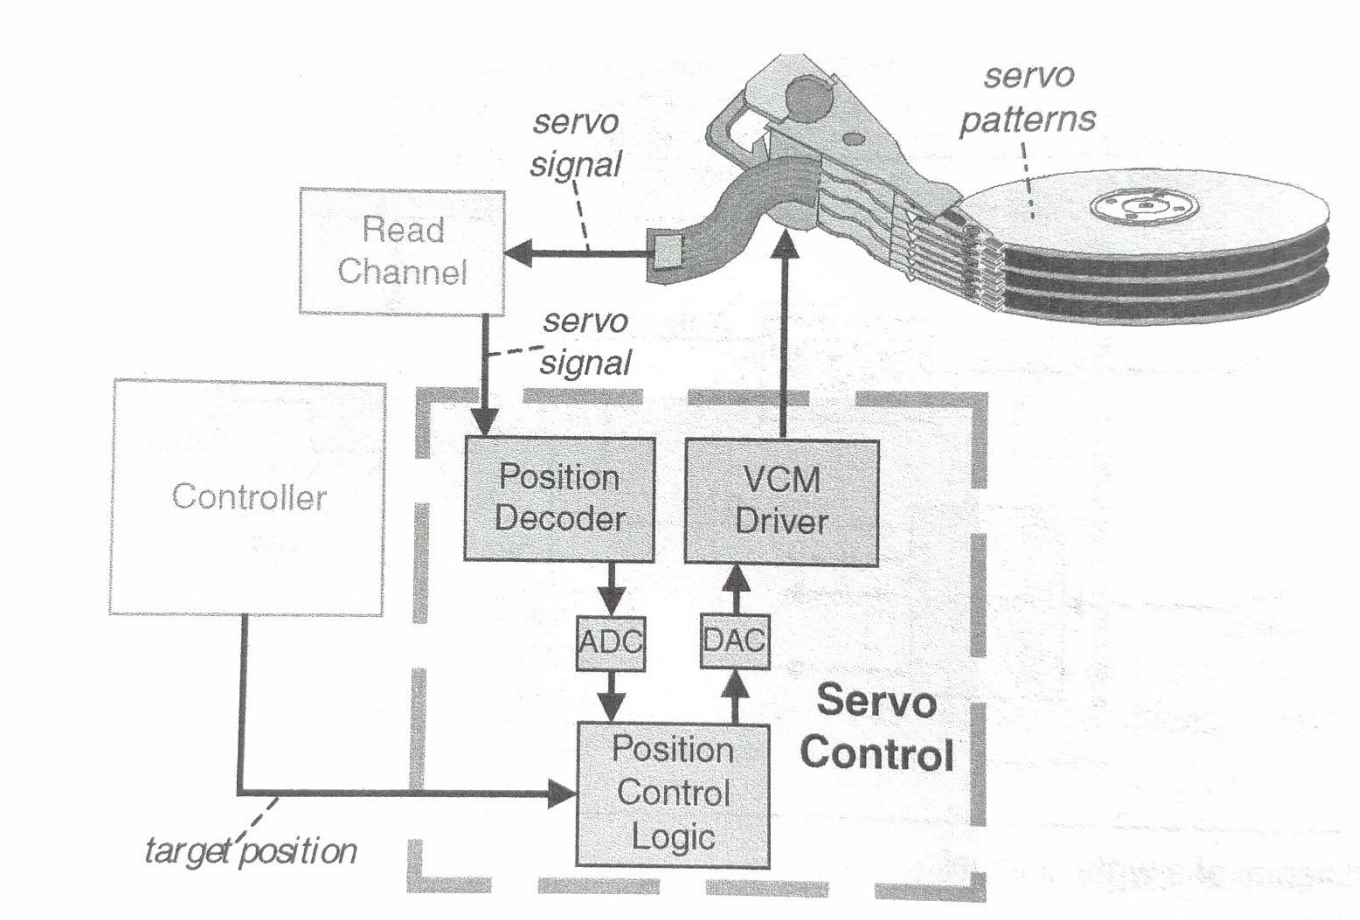

* Jacob et.al.

© tj

Hard Disk Drive - Systems

- Data
  - Data is stored in Fixed size blocks
    - 512 Bytes (user data)
    - ~ 544 Bytes after encoding
    - ~ 40 Bytes of ECC
    - ~ 2 Bytes of CRC
    - ~ 590 Bytes total for data
    - Some systems support 4KB user data blocks

© tj

Hard Disk Drive - Systems

- Overhead
  - Preamble (sync)
    - 10 bytes
    - Establish a baseline for the clock recovery PLL
    - Used to get AGC in range
  - Data Sync (address mark)
    - Special pattern – 3-4 bytes
    - Indicate beginning of data
  - Flush Pad
    - Extra bytes at the end to gracefully terminate the read channel at the end of the read

© tj

Hard Disk Drive - Systems

- Sector
  - Data + overhead
  - Fixed size

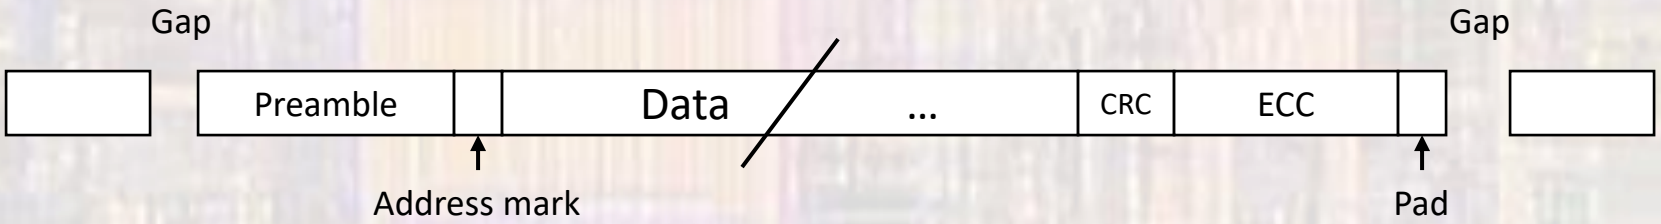

© tj

Hard Disk Drive - Systems

- Sector
  - Logical Sector
    - The size the host expects for data
    - 512B or 4KB
  - Physical Sector
    - The actual size the hard drives uses for sectors
    - The hard drive can collect multiple groups of data into a single entity
      - 4 – 512B host data blocks → 1 - 2K data block on the disk
      - Only 1 set of overhead for 4 host data blocks → higher density on disk
      - Must always look like 512B or 4KB at the external interface

© tj

Hard Disk Drive - Systems

- Sector
  - Physical sector size tradeoffs
    - Sequential sector configuration
      - Large files expect to be stored in sequential sectors
      - R/W over time leaves file size holes in the sector mapping
      - Large files cannot find big enough holes
      - Over time – lots of small holes get created – External Fragmentation

      → Logically sequential blocks and physically non-contiguous sectors
    - Large sector sizes
      - Small files or the ends of large files may not fill the sector – Internal Fragmentation

© tj

Hard Disk Drive - Systems

- Sector
  - Host has a file of X size (sequential)
  - The controller breaks the file into 512B blocks (sequential)
  - The controller maps the N 512B blocks into N physical sectors (non-contiguous)

© tj

Hard Disk Drive - Systems

- Tracks and Cylinders
  - Tracks
    - Concentric circles
    - Spiral
    - Track pitch < 20u inches
    - Sectors are numbered 1- N on any given track
    - Tracks are numbered 0 – M, with 0 at the outside edge
  - Cylinder
    - All the tracks with the same ID number up and down the stack
    - Some cylinders at the very outside edge are reserved for system use and are not available for data

© tj

Hard Disk Drive - Systems

- Sector Addressing
  - Internal Addressing
    - Each sector on the disk has a unique identifier (number) from 0 to N-1 where N is the total number of sectors on the disk drive
    - Also called Physical Block Address or Absolute Block Address
    - Each sector also has a CHS address
      - Cylinder
      - Head
      - Sector
      - Represents the sector in 3-D space
  - Both of these have been replaced with a method called GPT

© tj

Hard Disk Drive - Systems

- Sector Addressing
  - External Addressing
    - Logical Block Address
      - Host uses the logical address for the block
      - Controller maps the logical block address to a physical block address (PBA) in CHS format

© tj

Hard Disk Drive - Systems

- Sector Addressing
  - Logical to Physical Mapping
    - Sequential logical blocks naturally map to sequential physical sectors

      until the end of a track is reached

    - Cylinder mode
      - Go to the next track in the same cylinder

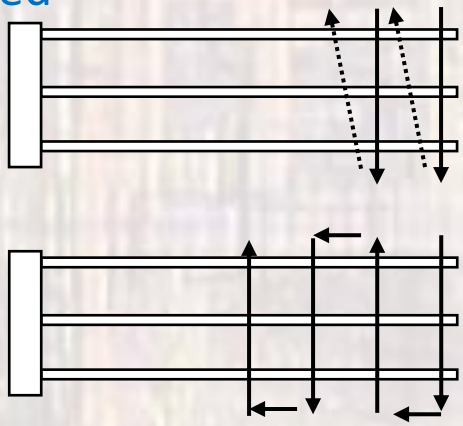

© tj

Hard Disk Drive - Systems

- Sector Addressing
  - Logical to Physical Mapping
    - Serpentine Format
      - Advance through tracks on a single disk
    - Banded Serpentine

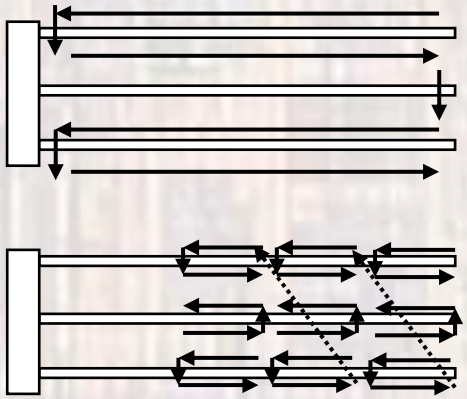

© tj

Hard Disk Drive - Systems

- Sector Addressing
  - Skewing
    - Stagger the first sector of each track relative to it's predecessor
    - Track Skew and Cylinder Skew

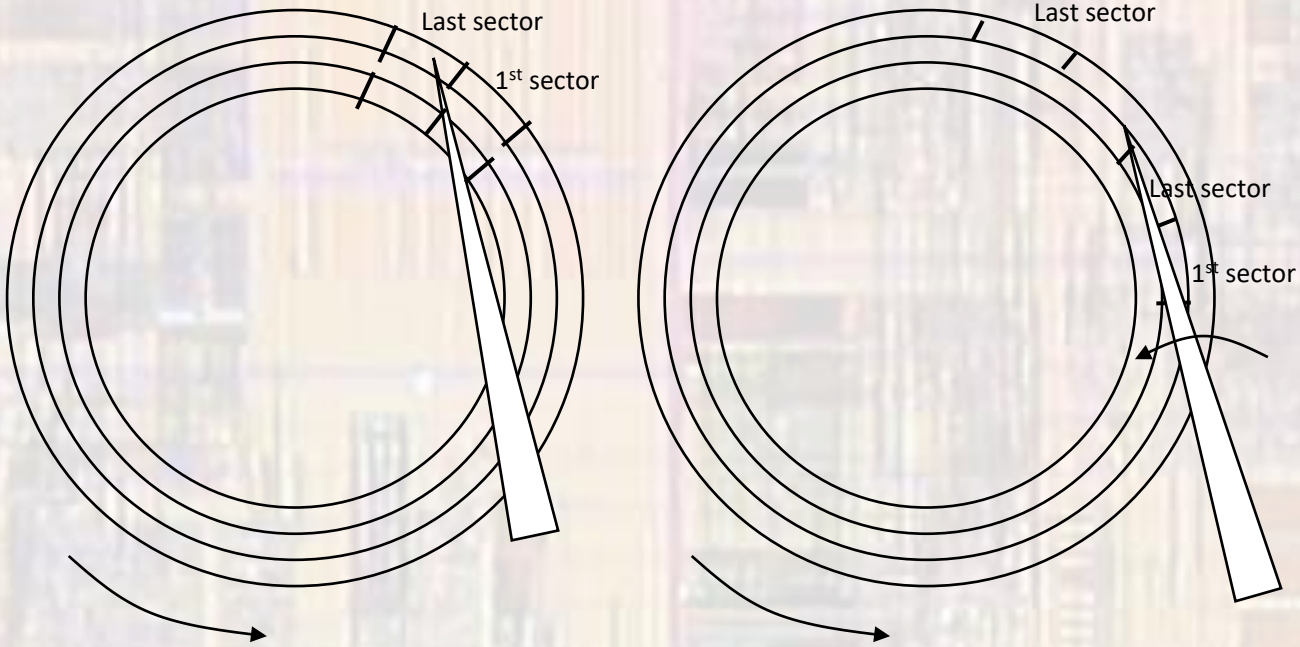

 © tj

Hard Disk Drive - Systems

- Cylinder Speed
  - Constant angular velocity
    - Difficult to modulate the rotational speed of the disk – fixed RPM
  - Fixed RPM → differences in linear speed for different tracks
    - Put the same number of sectors in each track
      - Constant bit rate
      - Poor bit density as you go further out
    - Use a fixed linear bit density
      - More sectors as you go out
      - Different bit rates – higher at the outside

© tj

Hard Disk Drive - Systems

- Zoned-Bit Recording (ZBR)
  - ZBR
    - Compromise between fixed number of sectors and fixed linear bit density
    - Fixed linear density
    - Limited number of different bit rates

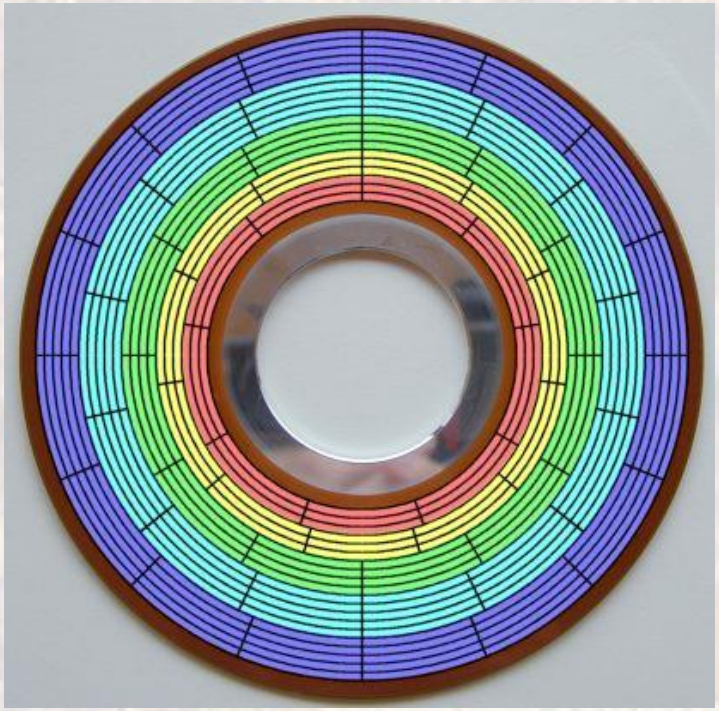

© tj

Hard Disk Drive - Systems

- Servo
  - How does the drive align the head with the tracks?

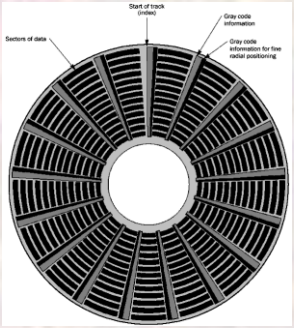

  - Dedicated Servo
    - One surface of one platter is dedicated to servo control
    - Special patterns allow the servo head to align and identify it's location
  - Embedded Servo
    - Stripes on the surface of each disk
    - Special patterns allow each head to align and identify it's location

Servo Burst

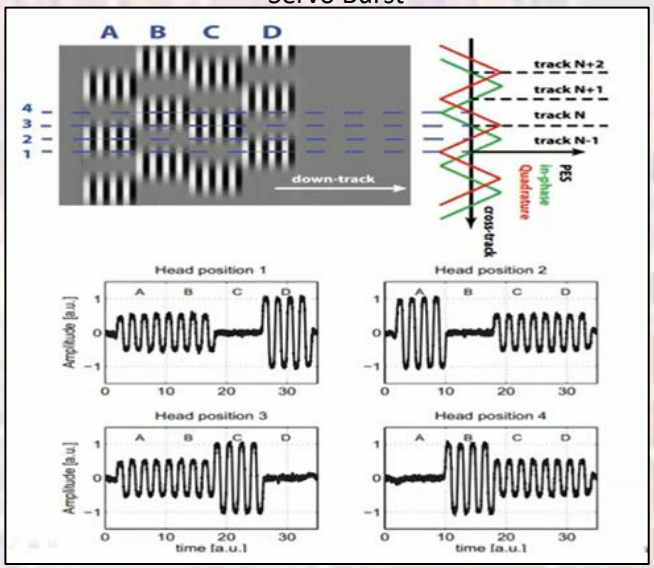

timing

AGC		Track ID	Servo Burst

© tj

Hard Disk Drive - Systems

- Performance

	ST4000NM0023 ST4000NM0043 ST4000NM0063	ST3000NM0023 ST3000NM0043 ST3000NM0063	ST2000NM0023 ST2000NM0043 ST2000NM0063	ST1000NM0023 ST1000NM0043 ST1000NM0063	
Drive capacity	4TB	3TB	2TB	1TB	(fomatted, rounded off value)
Read/write data heads	10	8	5	3	
Bytes/track	1,668,096				Bytes (average, rounded off values)
Bytes/surface	400,000				MB (unfomatted, rounded off values)
Tracks/surface (total)	320,800				Tracks (user accessible)
Tracks/in	305,000				TPI (average)
Peak bits/in	1,904,000				BPI
Areal density	578				Gb/in2
Internal data rate	2210				Mb/s (max)
Disk rotation speed	7200				RPM
Avg rotational latency	4.16				ms

© tj

Hard Disk Drive - Systems

- Performance

Maximum Internal data rate*	2.21 Gb/s
Sustained transfer rate	83 to 175 MB/s **
SAS Interface maximum instantaneous transfer rate	600MB/s* per port (dual port = 1200MB/s*)
Logical block sizes 512 (default), 520 or 528.	
Read/write consecutive sectors on a track	Yes
Flaw reallocation performance impact (for flaws reallocated at format time using the spare sectors per sparing zone reallocation scheme.)	Negligible
Average rotational latency	4.16ms

© tj

Hard Disk Drive - Systems

- Performance

Models	ST6000DM001, ST5000DM002	ST4000DM000
Interface	SATA	
Recording method	TGMR	
Recording density (kFCI)	1981	1807
Track density (ktracks/inch avg)	320	
Areal density (Gb/in^2)	633	625
Internal data transfer rate (Mb/s max)	1981	1813
Average data rate, read/write (MB/s)	180	146
Maximum sustained data transfer rate, OD read (MB/s)	220	180
I/O data-transfer rate (MB/s max)	600	

© tj

Hard Disk Drive - Systems

- Interface
  - Historical
    - IDE, PATA, SCSI
    - Parallel Interfaces
  - Current
    - SATA, SAS (Serial SCSI)
    - Serial Interfaces
    - Point to point
    - Protocol Based

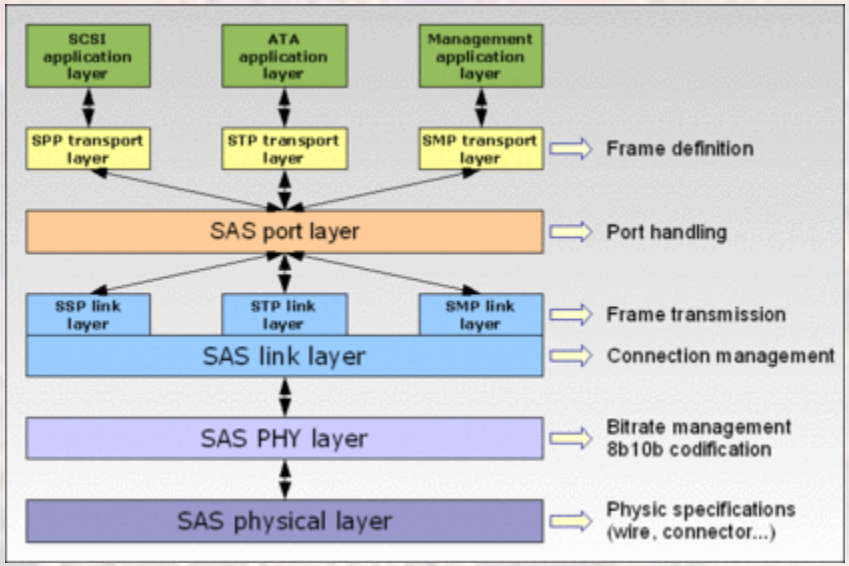

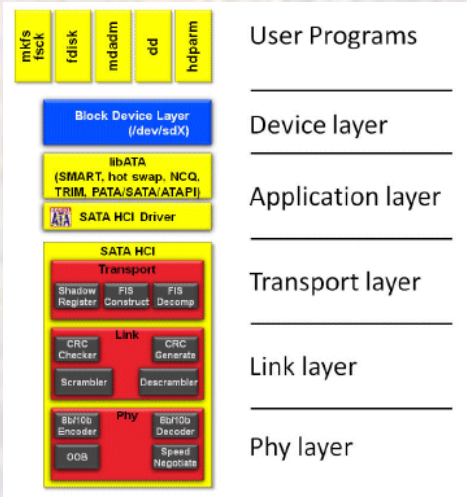

© tj