溫故之新: 하드디스크와 플래시메모리

민 상 렬 (symin@snu.ac.kr)
(+ Mobile Embedded Systems Lab 연구원들)
서울대학교 공과대학
컴퓨터공학부

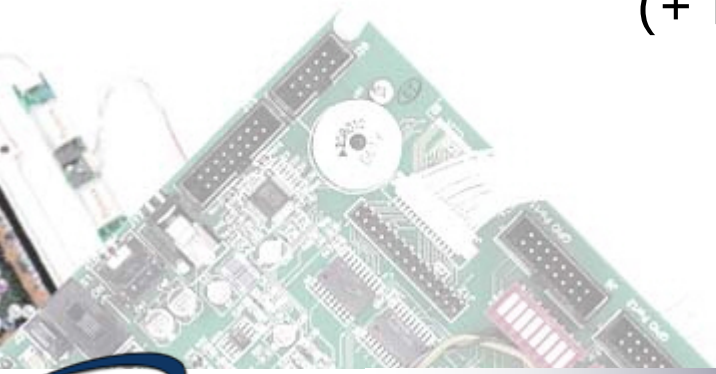

Outline

- HDD Basics and Demo
- Flash Memory Basics and Demo
- Storage Trends
- Conclusions

Outline

- HDD Basics and Demo
- Flash Memory Basics and Demo
- Storage Trends
- Conclusions

HDD internals

- Electronic components
- Mechanical components

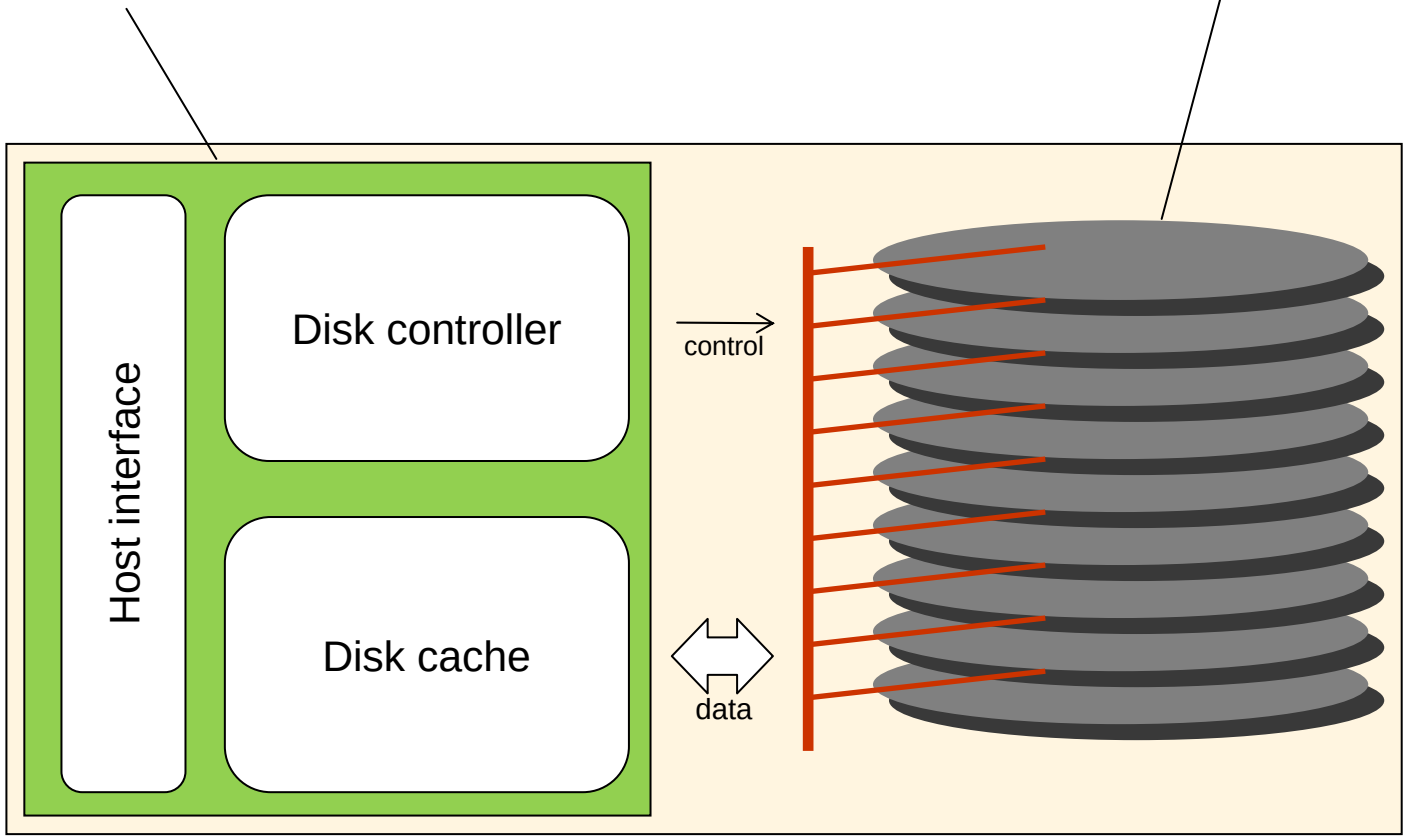

Mechanical components

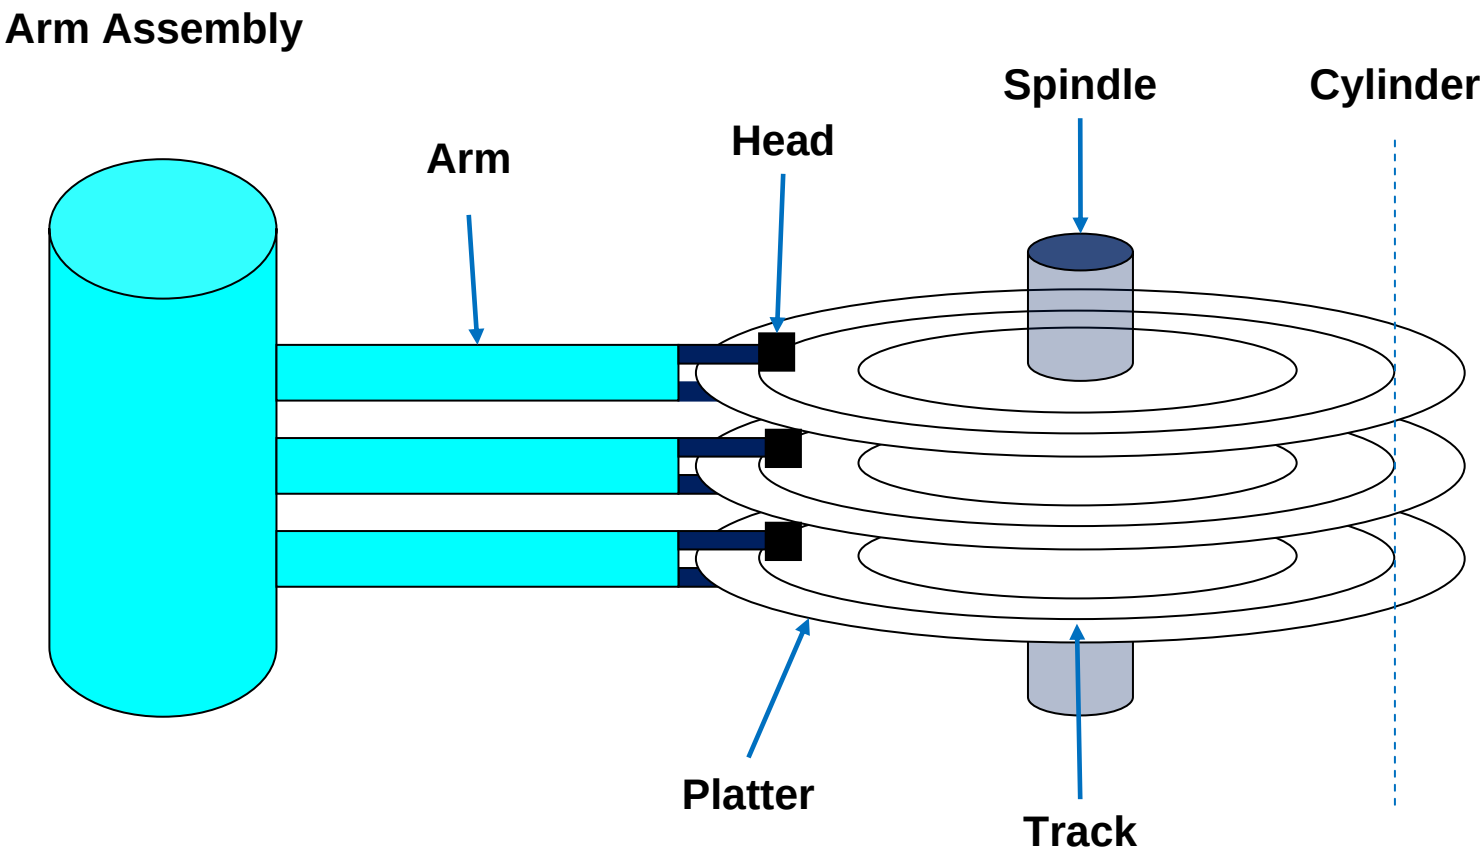

Source: *"ABCs of Disk Drives,"* Sudhanva Gurumurthi

Data layout

- Rotating disks consist of platters, each with two surfaces
- Each surface consists of concentric rings called tracks
- Each track consists of sectors separated by gaps

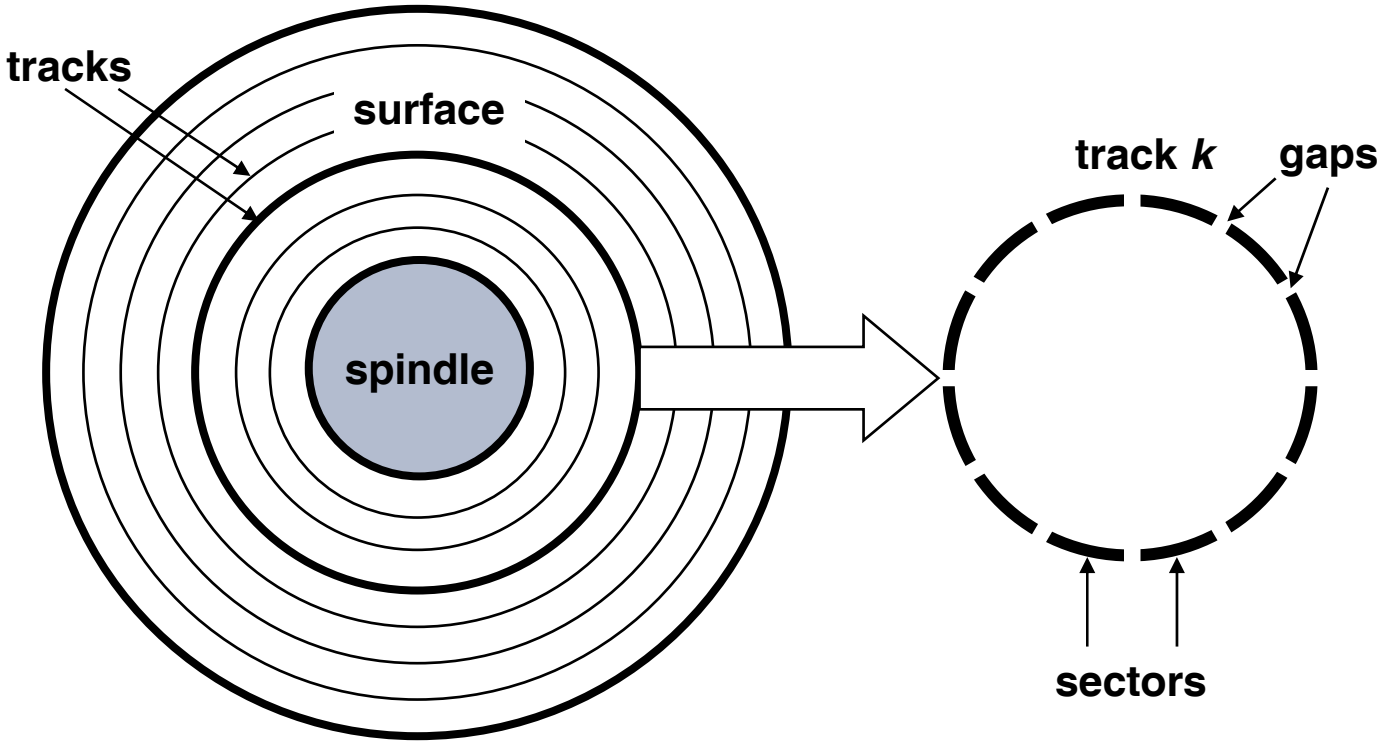

Source:
"http://camars.kaist.ac.kr/~joon/course/sep562_2006_1/notes/10_11%20Memory_Hierarchy.ppt"

Disk operation

The disk surface spins at a fixed rotational rate

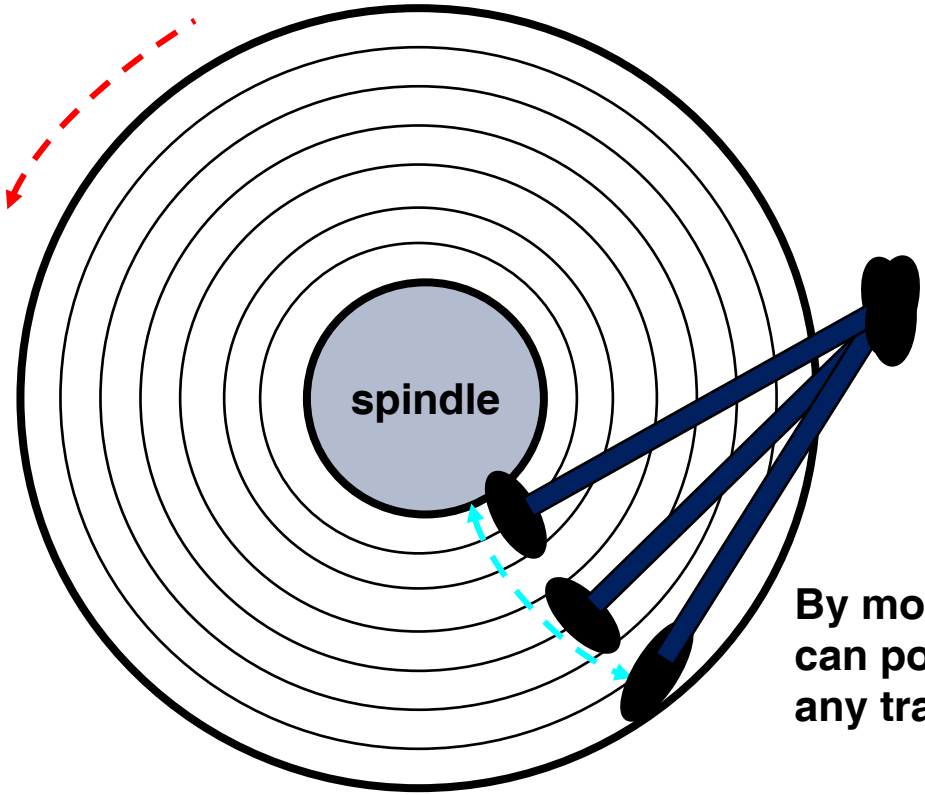

The head is attached to the end of the arm and flies over the disk surface on a thin cushion of air

By moving radially, the arm can position the head over any track

Source:
"http://camars.kaist.ac.kr/~joon/course/sep562_2006_1/notes/10_11%20Memory_Hierarchy.ppt"

Disk operation details

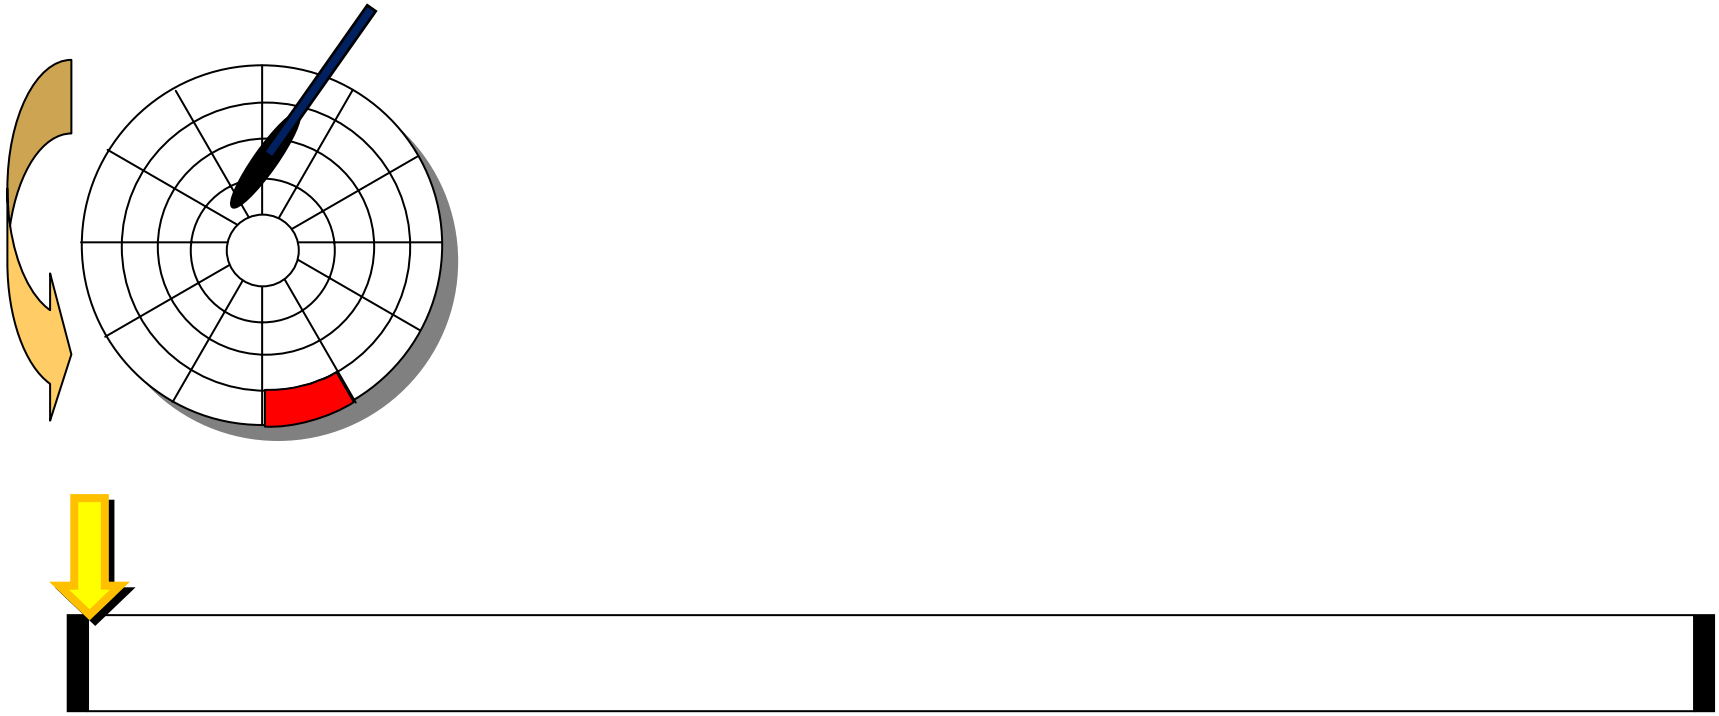

Source: *"http://www.cs.duke.edu/~chase/cps110/slides/files1.ppt"*

Disk operation details

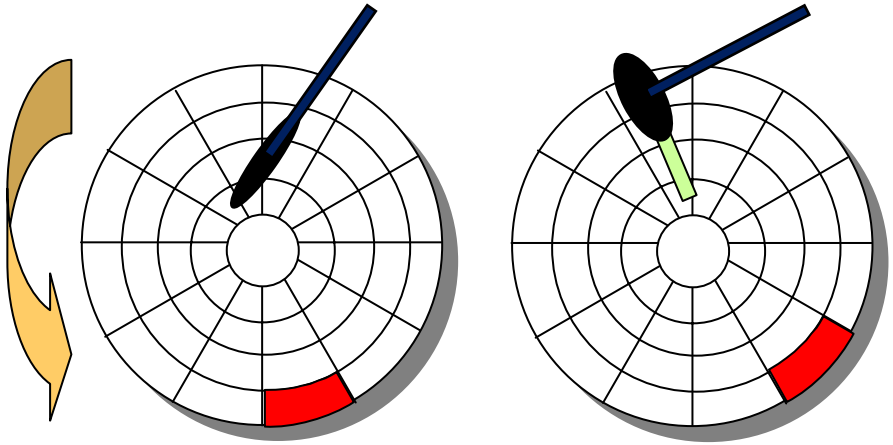

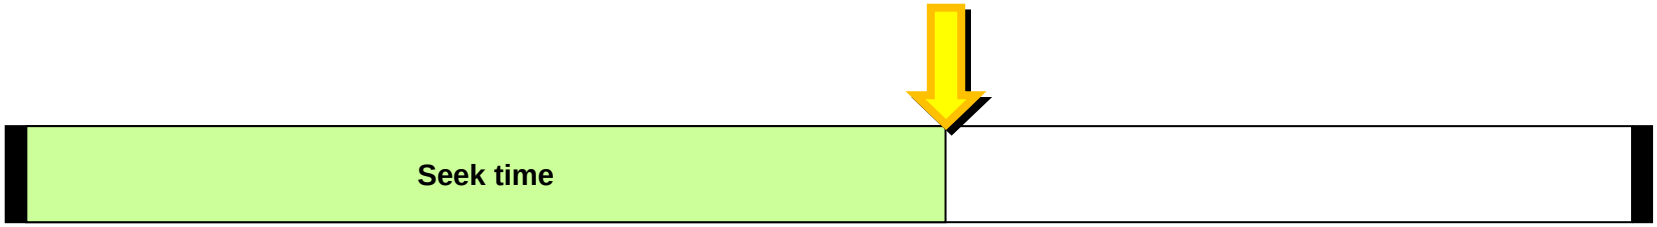

Source: *"http://www.cs.duke.edu/~chase/cps110/slides/files1.ppt"*

Disk operation details

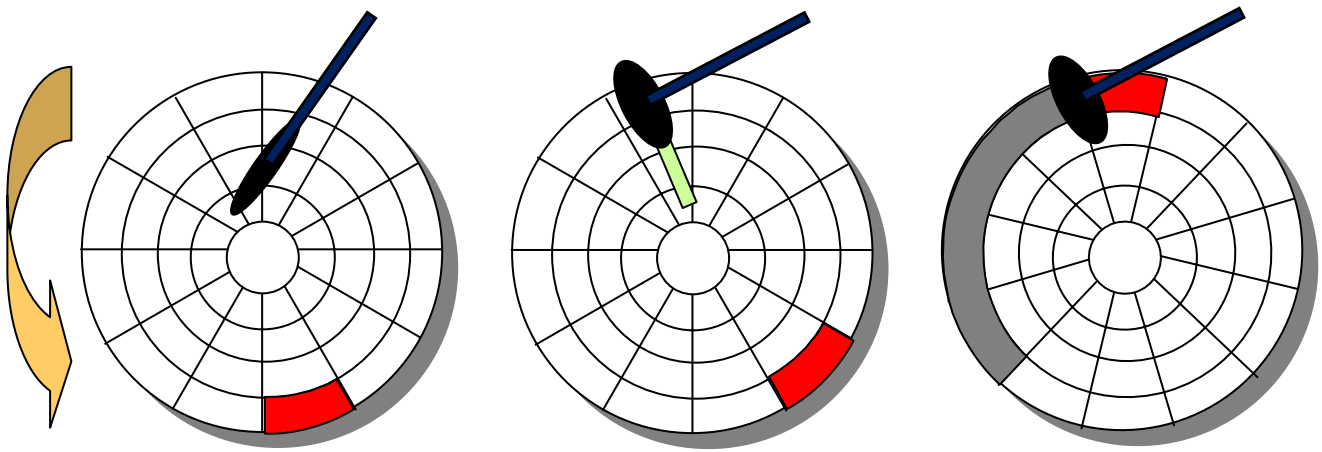

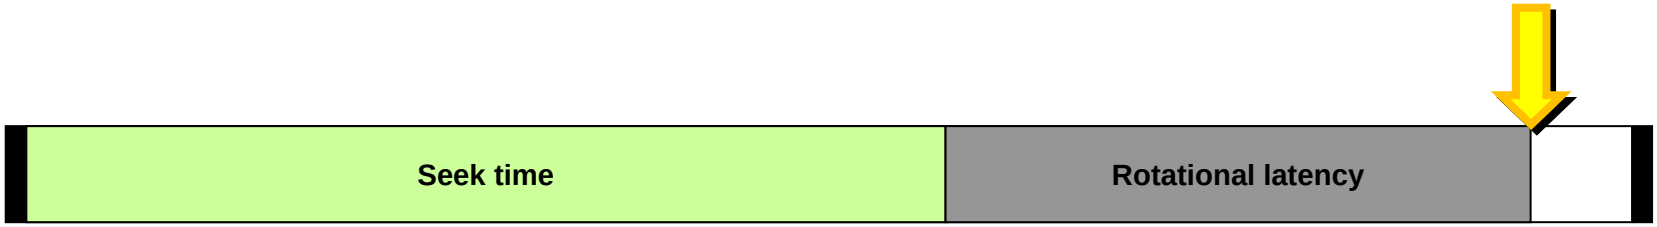

Source: *"http://www.cs.duke.edu/~chase/cps110/slides/files1.ppt"*

Disk operation details

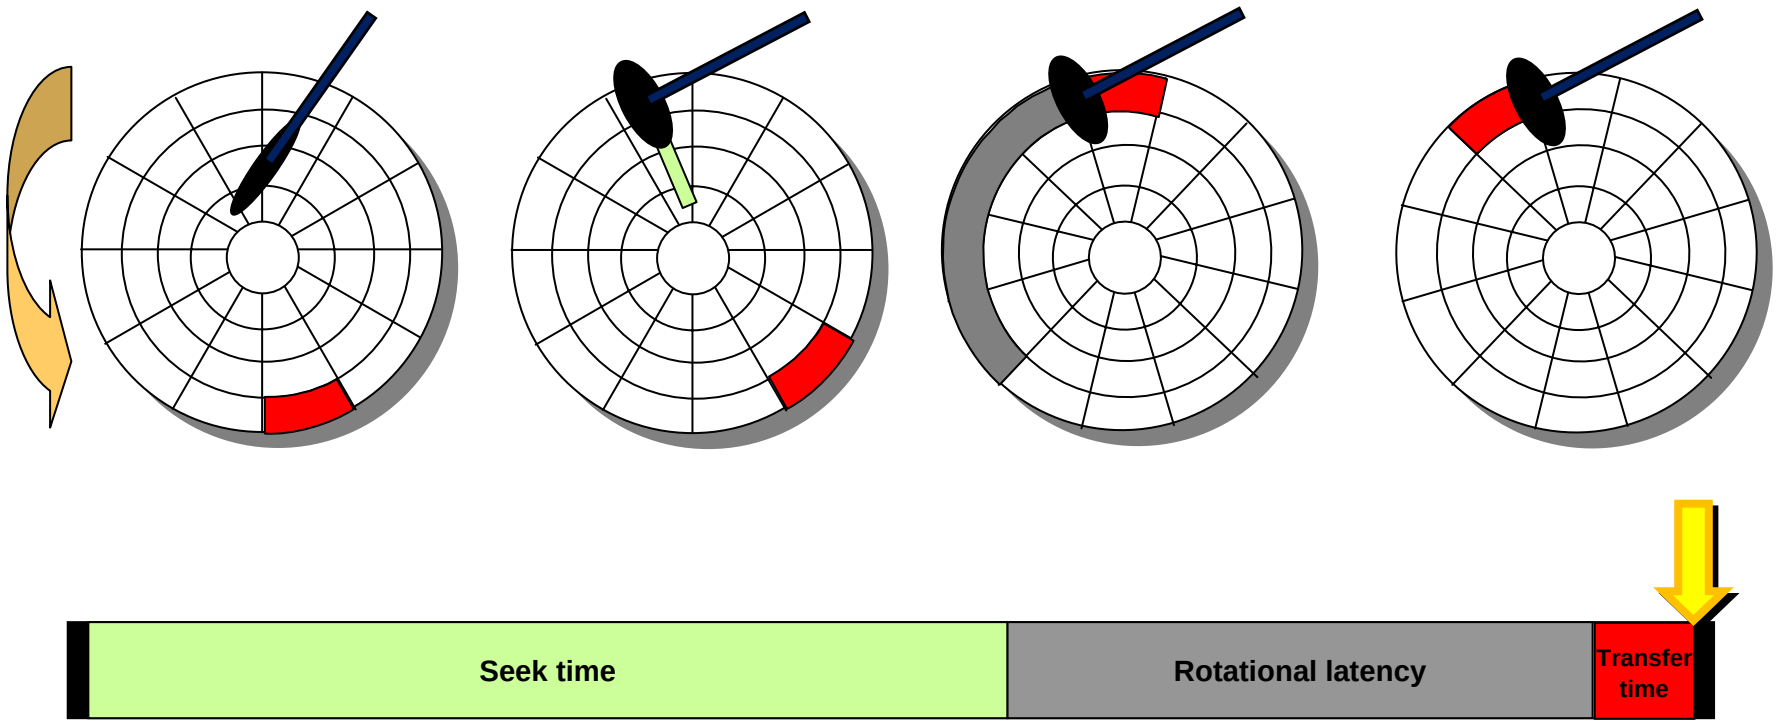

Source: *"http://www.cs.duke.edu/~chase/cps110/slides/files1.ppt"*

Disk access time

- Disk access time
  - Seek time + Rotational latency + Transfer time
- Seek time
  - Time to position heads over cylinder containing target sector
  - 0 ~ 25 ms
- Rotational latency
  - Time waiting for first bit of target sector to pass under r/w head
  - Full rotation: 4 ~ 12 ms (15000 ~ 5400 RPM)
- Transfer time
  - Time to read the bits in the target sector
  - 1 sector transfer: 1.3 ~ 12.8 us (380 ~ 40 MB/s transfer rate)

Electronic components

- Presenting a simple abstract view of the complex sector geometry

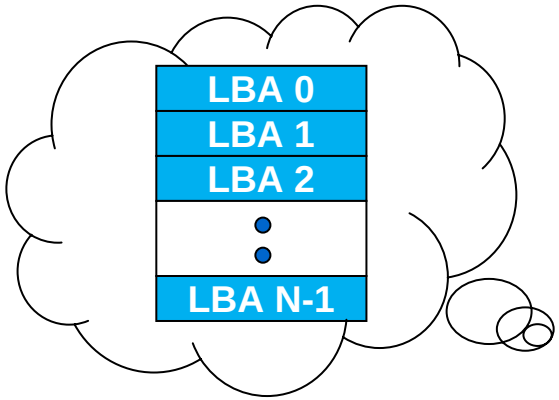

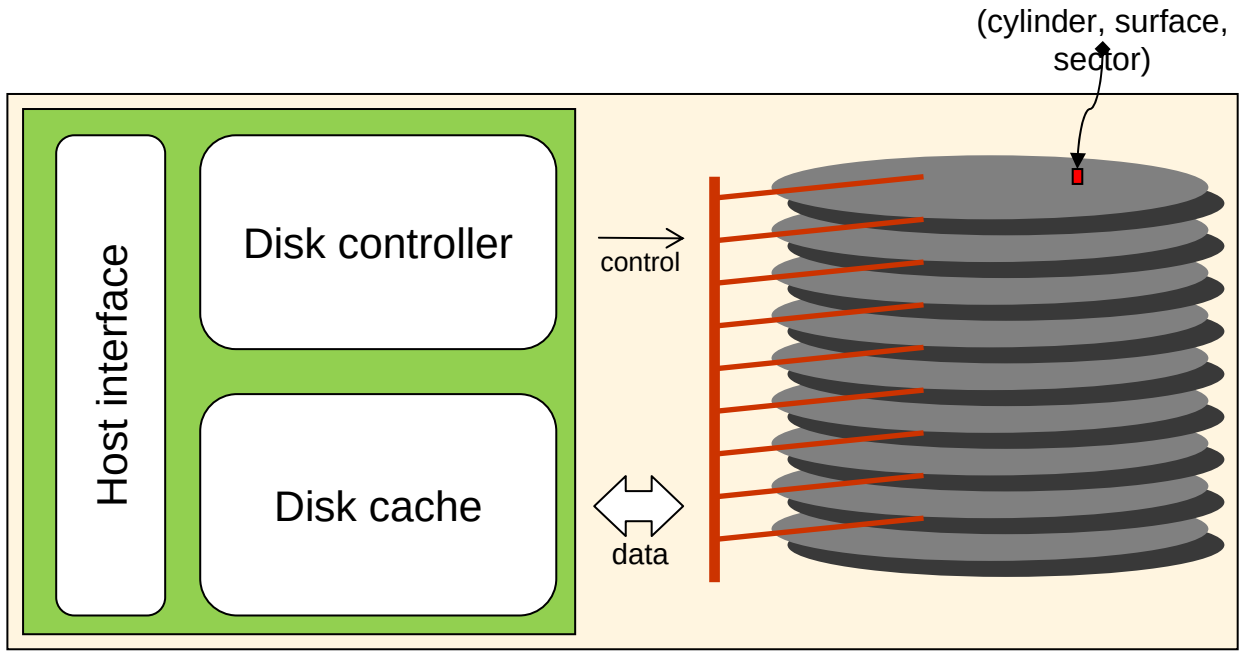

Electronic components

- Disk controller
  - Controlling the overall system
  - Major functions
    - Host interface
    - Request translation (LBA $\leftrightarrow$ [cylinder, surface, sector])
    - Reliability mechanism (e.g. ECC, bad sector handling)
    - Performance improvement (e.g. request scheduling and disk caching)
    - Power management (e.g. spin down of spindle motor)
  - Typically, embedded processor (such as ARM) + logic circuits

Outline

- HDD Basics and Demo
  - Demo
- Flash Memory Basics and Demo
- Storage Trends
- Conclusions

Demo HDD Specification

Model Name: SAMSUNG MP0402H (2.5 in)

- Size:
  - total 78,236,550 sectors
  - 40,057,113,600 bytes ≈ 37.30 GB
- Interface: ATA-6 (supports UDMA100)
- Buffer: 8MB DRAM
- Performance brief:
  - Avg. Seek time: 12 ms
  - Avg. Rotational Latency: 5.6 ms (5400 RPM)

- reference url:
  http://www.samsung.com/Products/HardDiskDrive/SpinPointMSeries/HardDiskDrive_SpinpointMseries_MP0402H_sp.htm

Demo I – Power-on sequence

Demo II – Sequential read/write

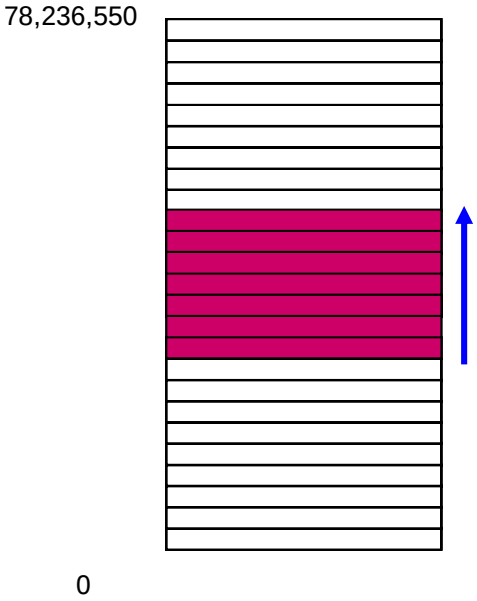

- Access pattern
  - read/write data whose address increases continuously

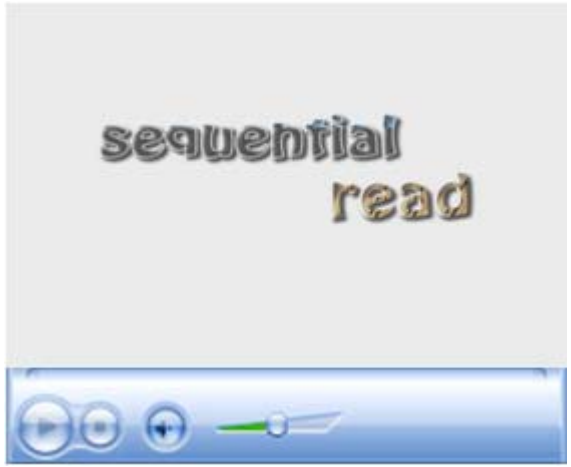

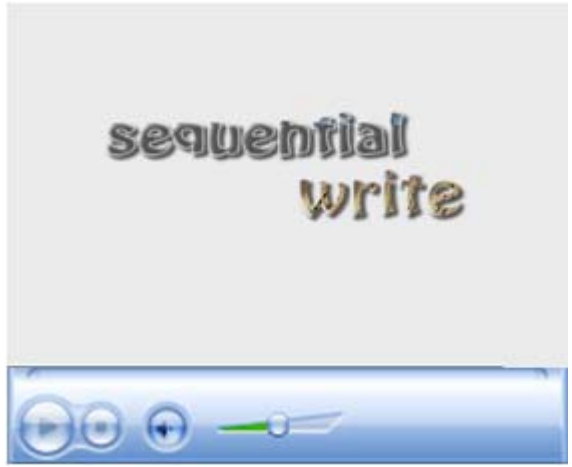

Demo III – Read/Write with a stride

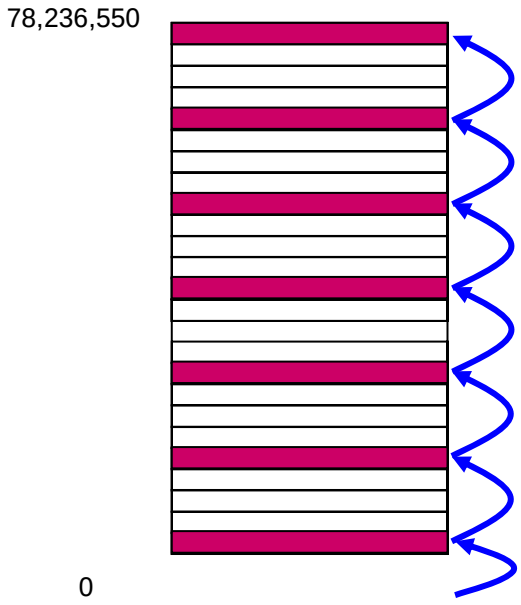

- Access pattern
  - read/write data whose address increases with a regular interval

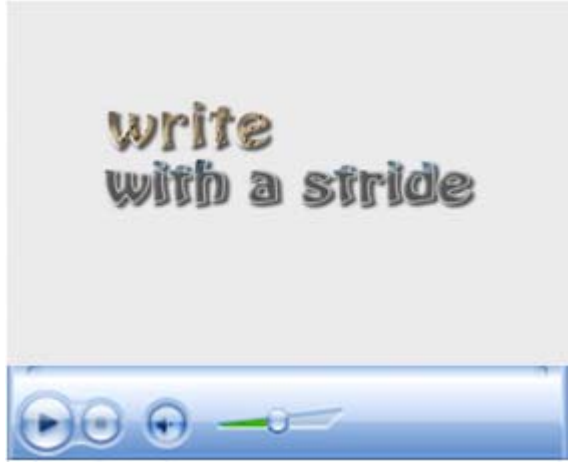

Demo IV – Read/Write in a convergent manner

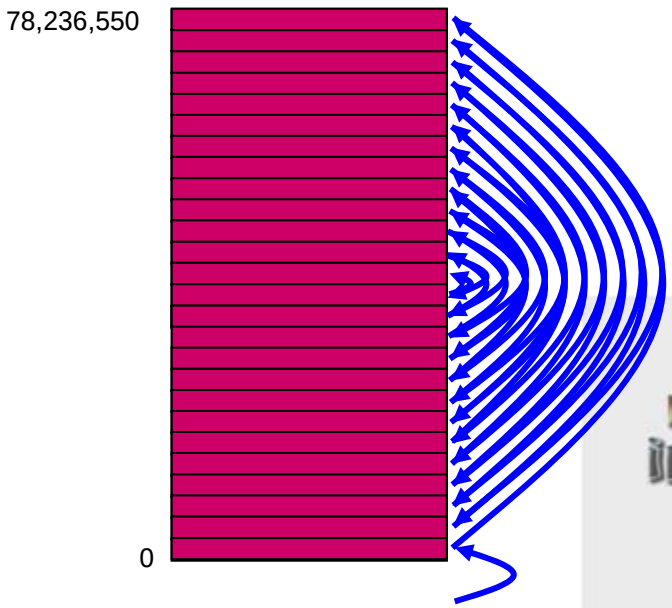

- Access pattern
  - read/write data whose address is not overlapped and is in a convergent manner

Demo V – Random read/write

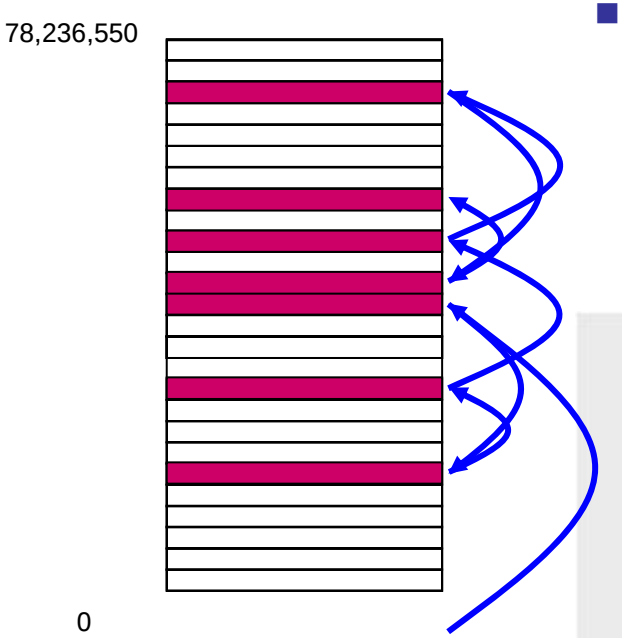

- Access pattern
  - read/write random addresses

Demo VI – Effect of read caching/write buffering

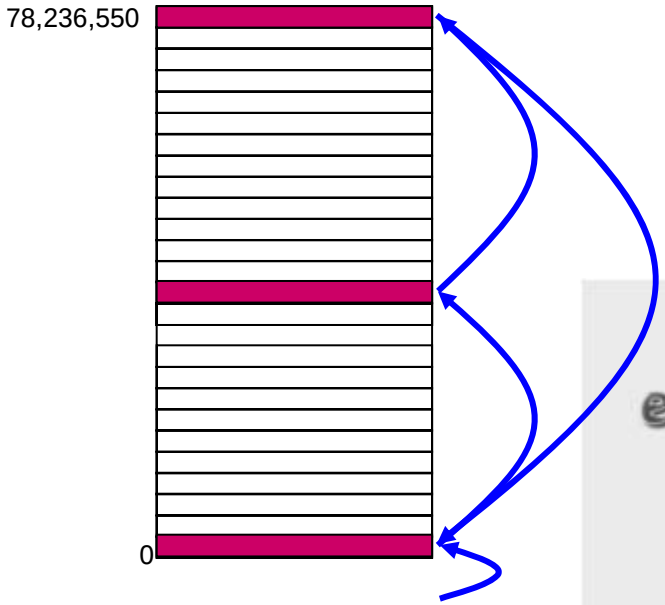

- Access pattern
  - access on the fixed addresses repeatedly

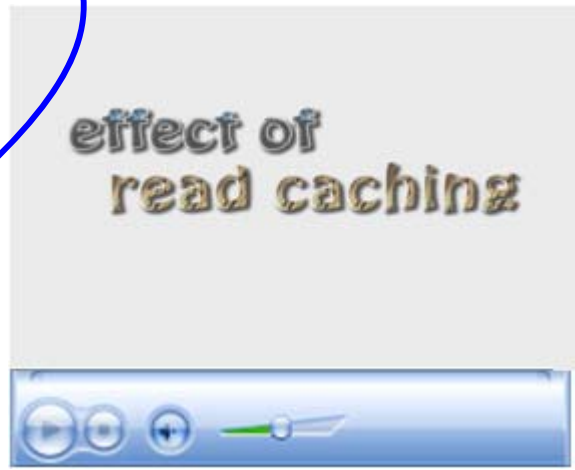

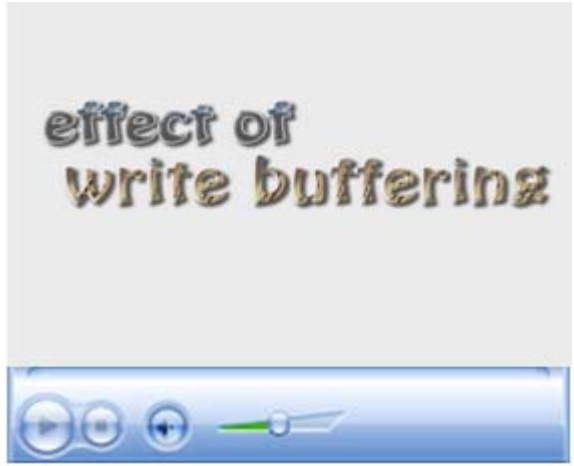

Demo VII – Windows XP start-up

HDD performance trends (1)

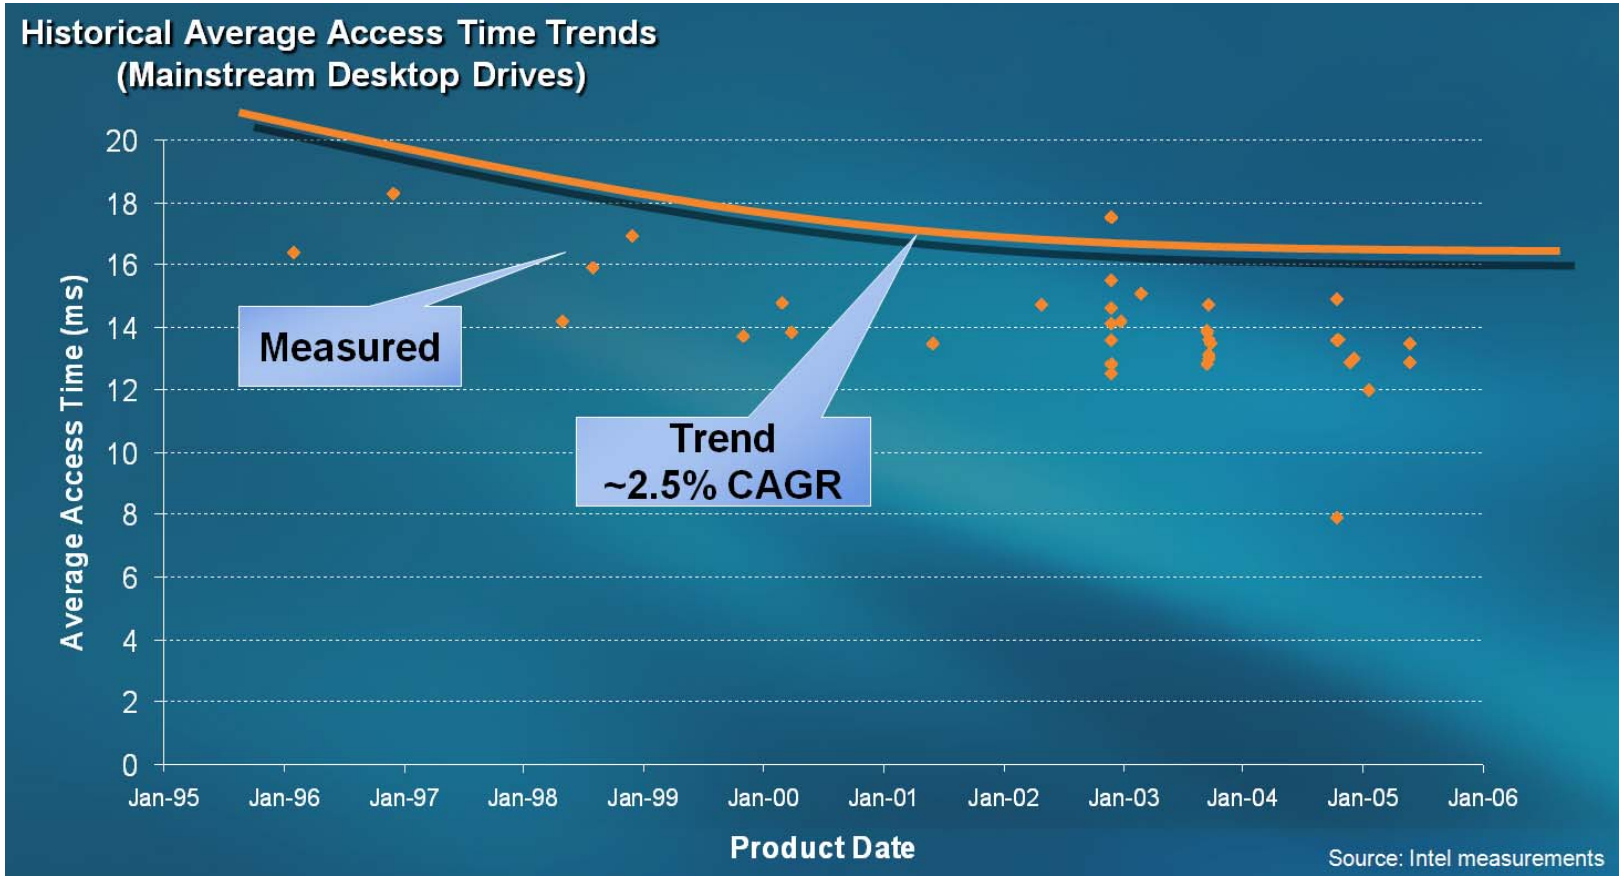

- HDD access time trends are fairly flat due to mechanical nature of device

HDD Performance trends (2)

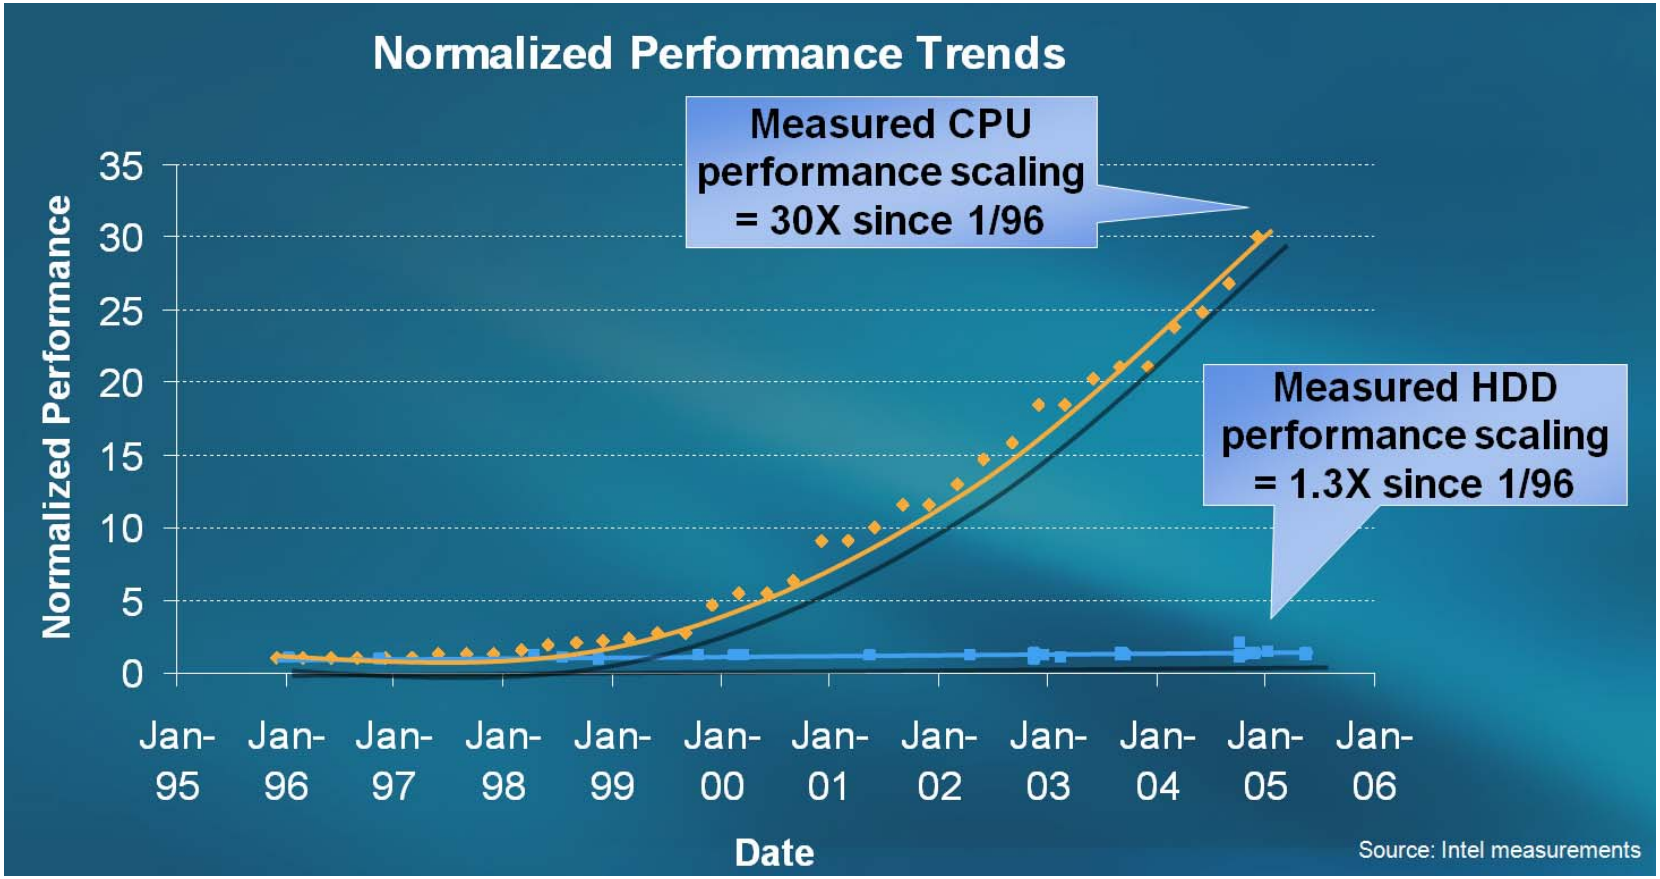

- A workload that was 5% disk bound in ‘96 would be 55% disk bound in ‘05

HDD density trends

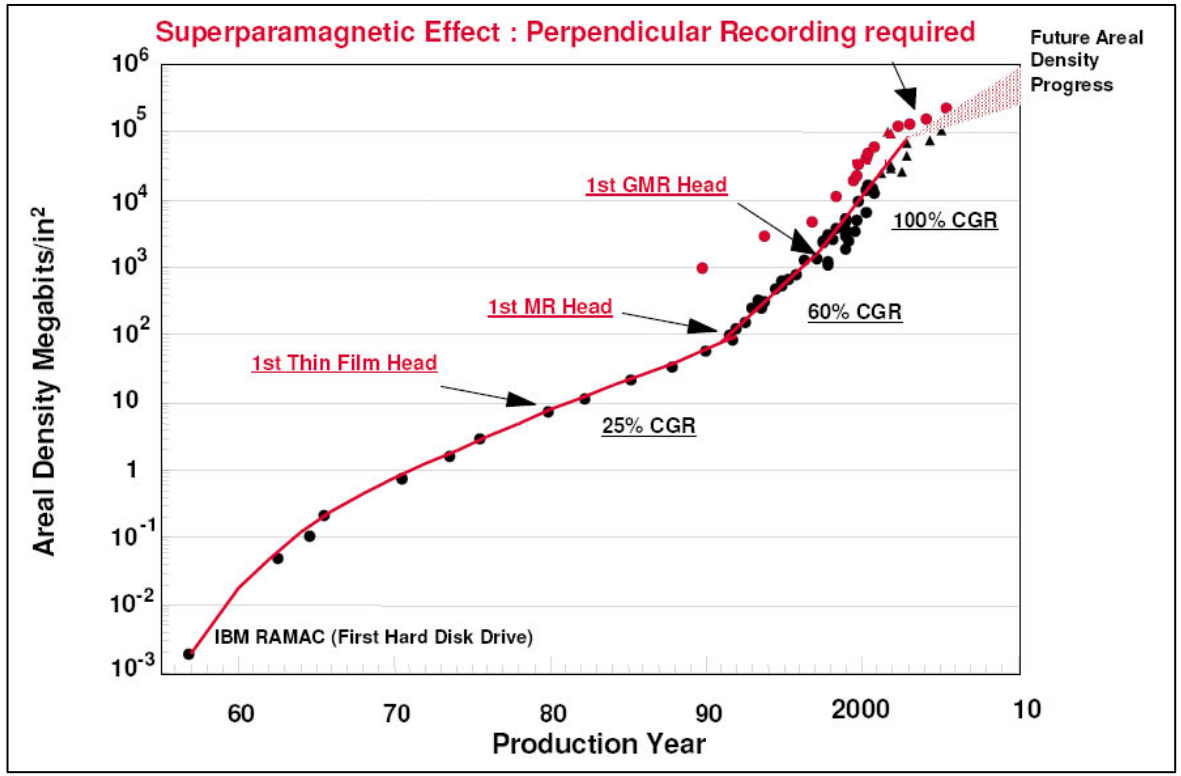

Source: Hitachi Global Storage Technologies

HDD Summary

The Ugly
- Latent sector errors

The Bad
- High latency
- High power consumption
- Low reliability
- Large form factor
- Limited parallelism

The Good
- High capacity
- Low cost

Outline

- HDD Basics and Demo
- Flash Memory Basics and Demo
- Storage Trends
- Conclusions

Conventional MOS Transistor

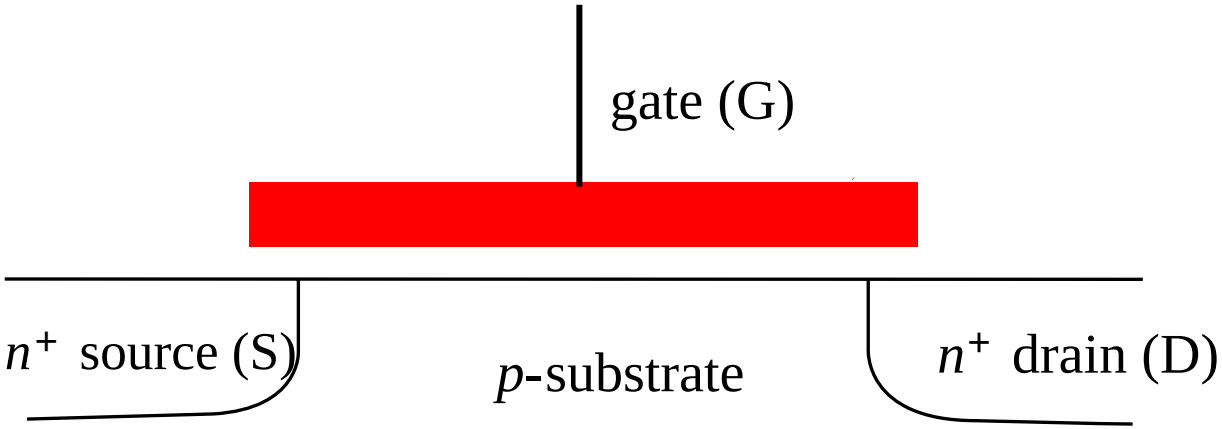

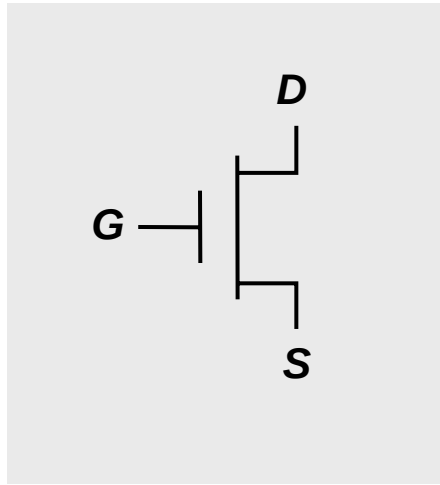

Schematic symbol

Conventional MOS Transistor: A Constant-Threshold Transistor

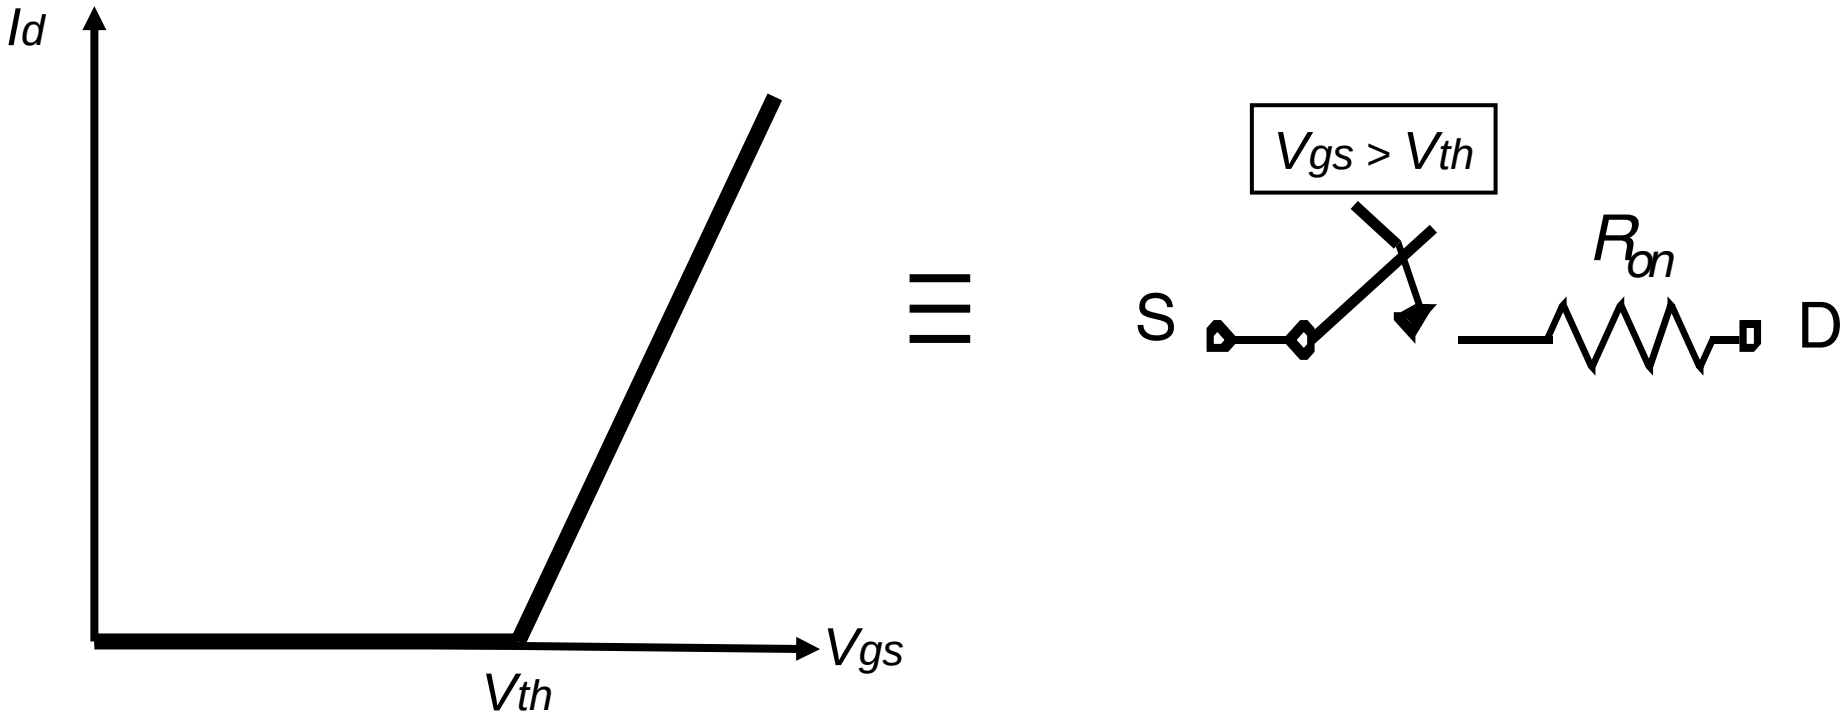

Flash Memory

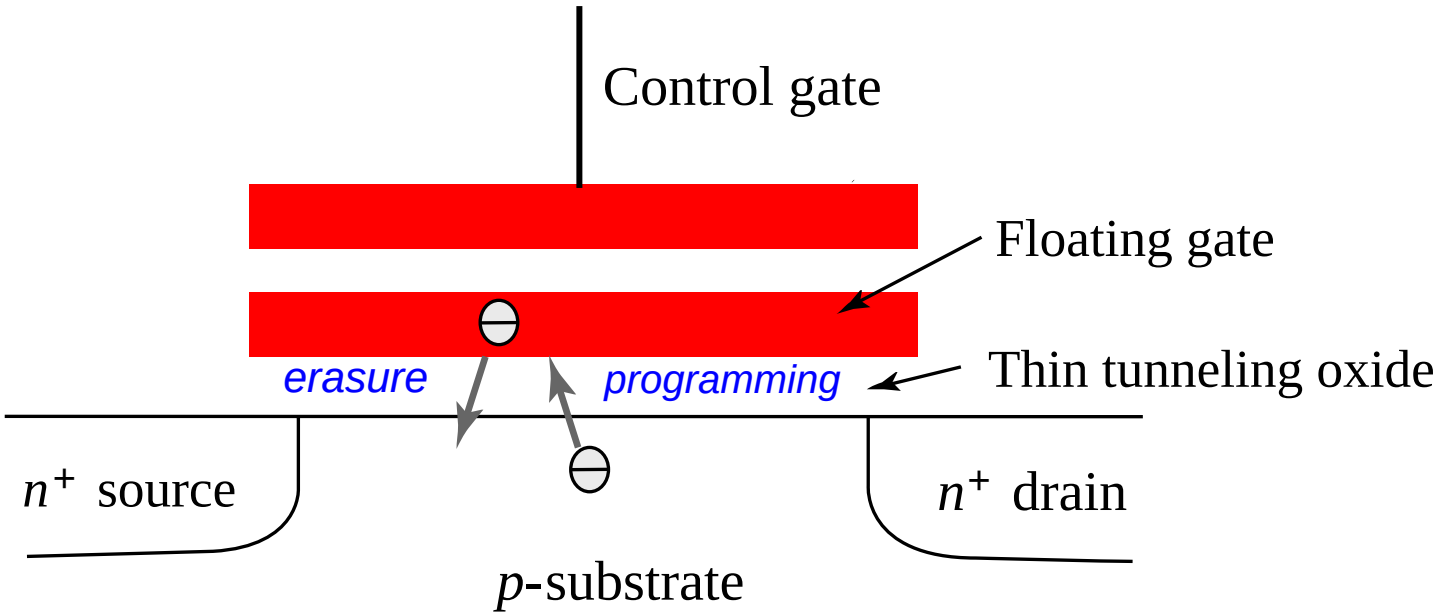

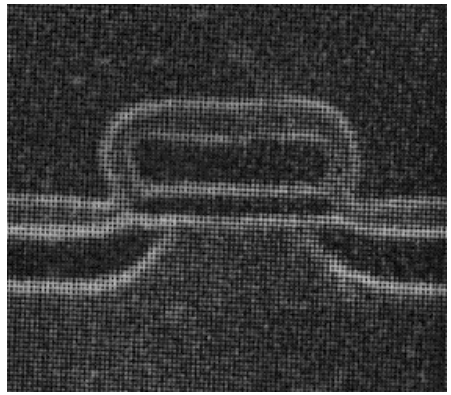

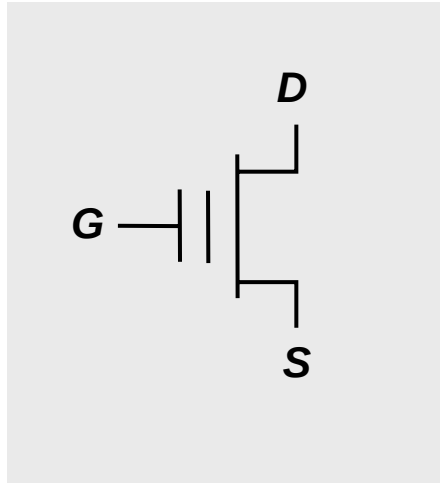

Schematic symbol

Flash Memory

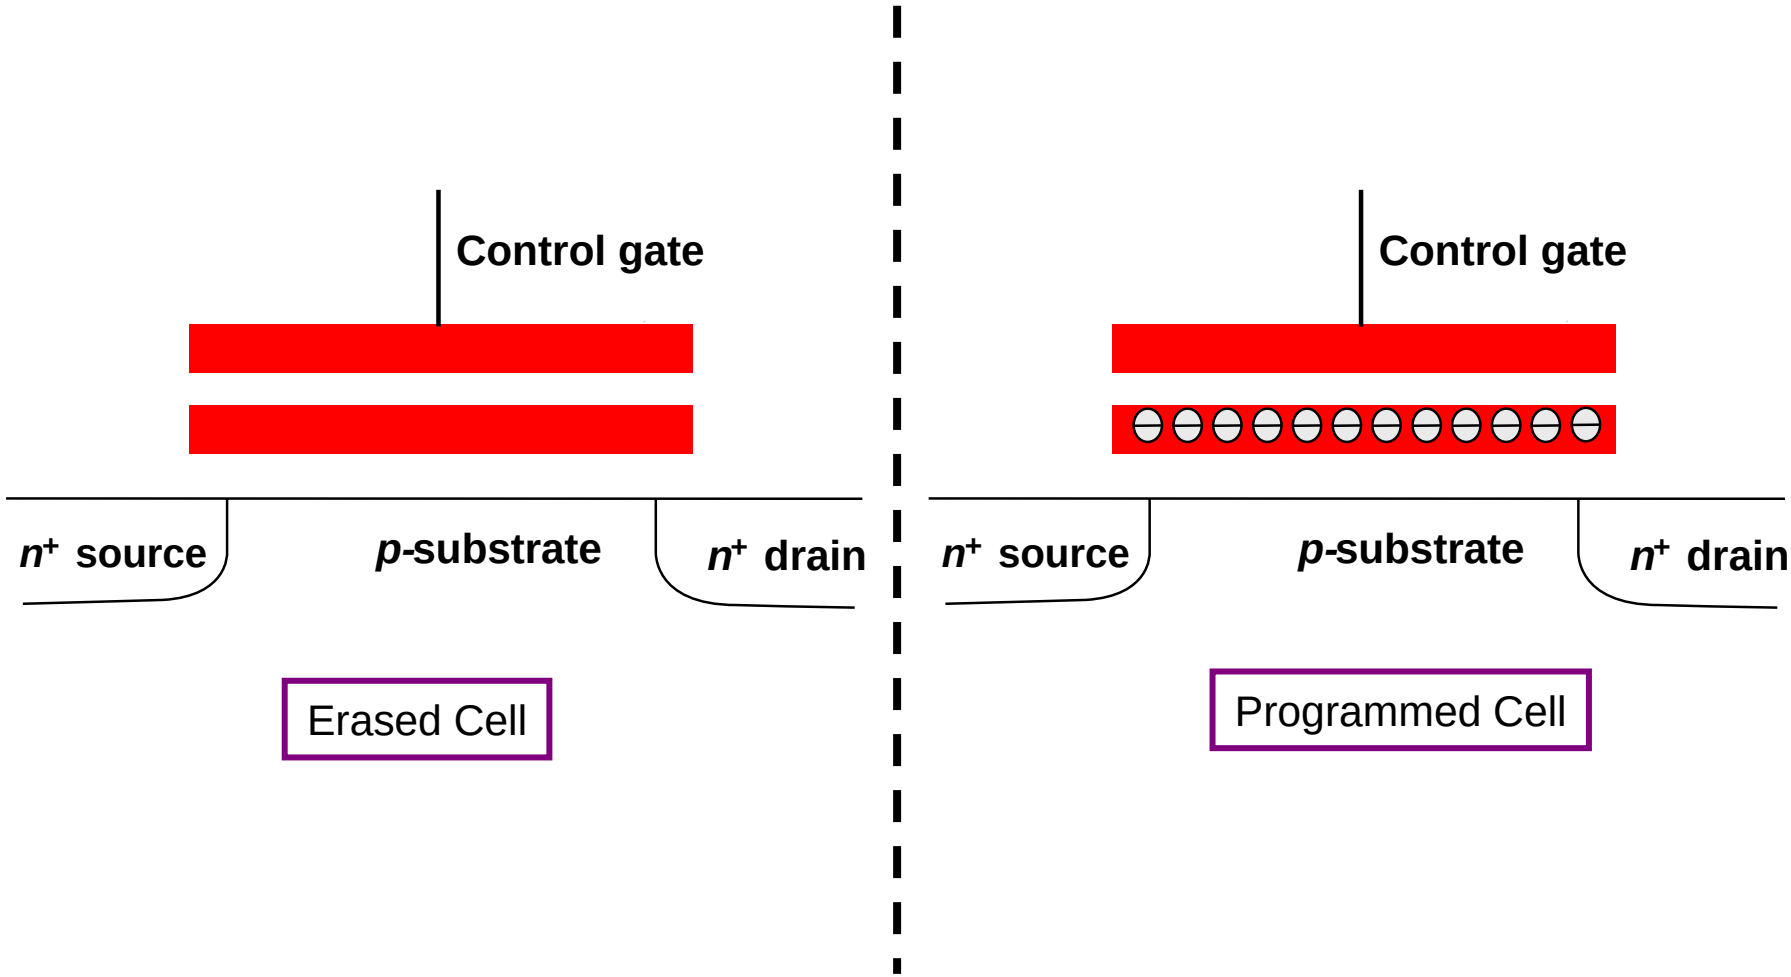

Flash Memory: A "Programmable-Threshold" Transistor

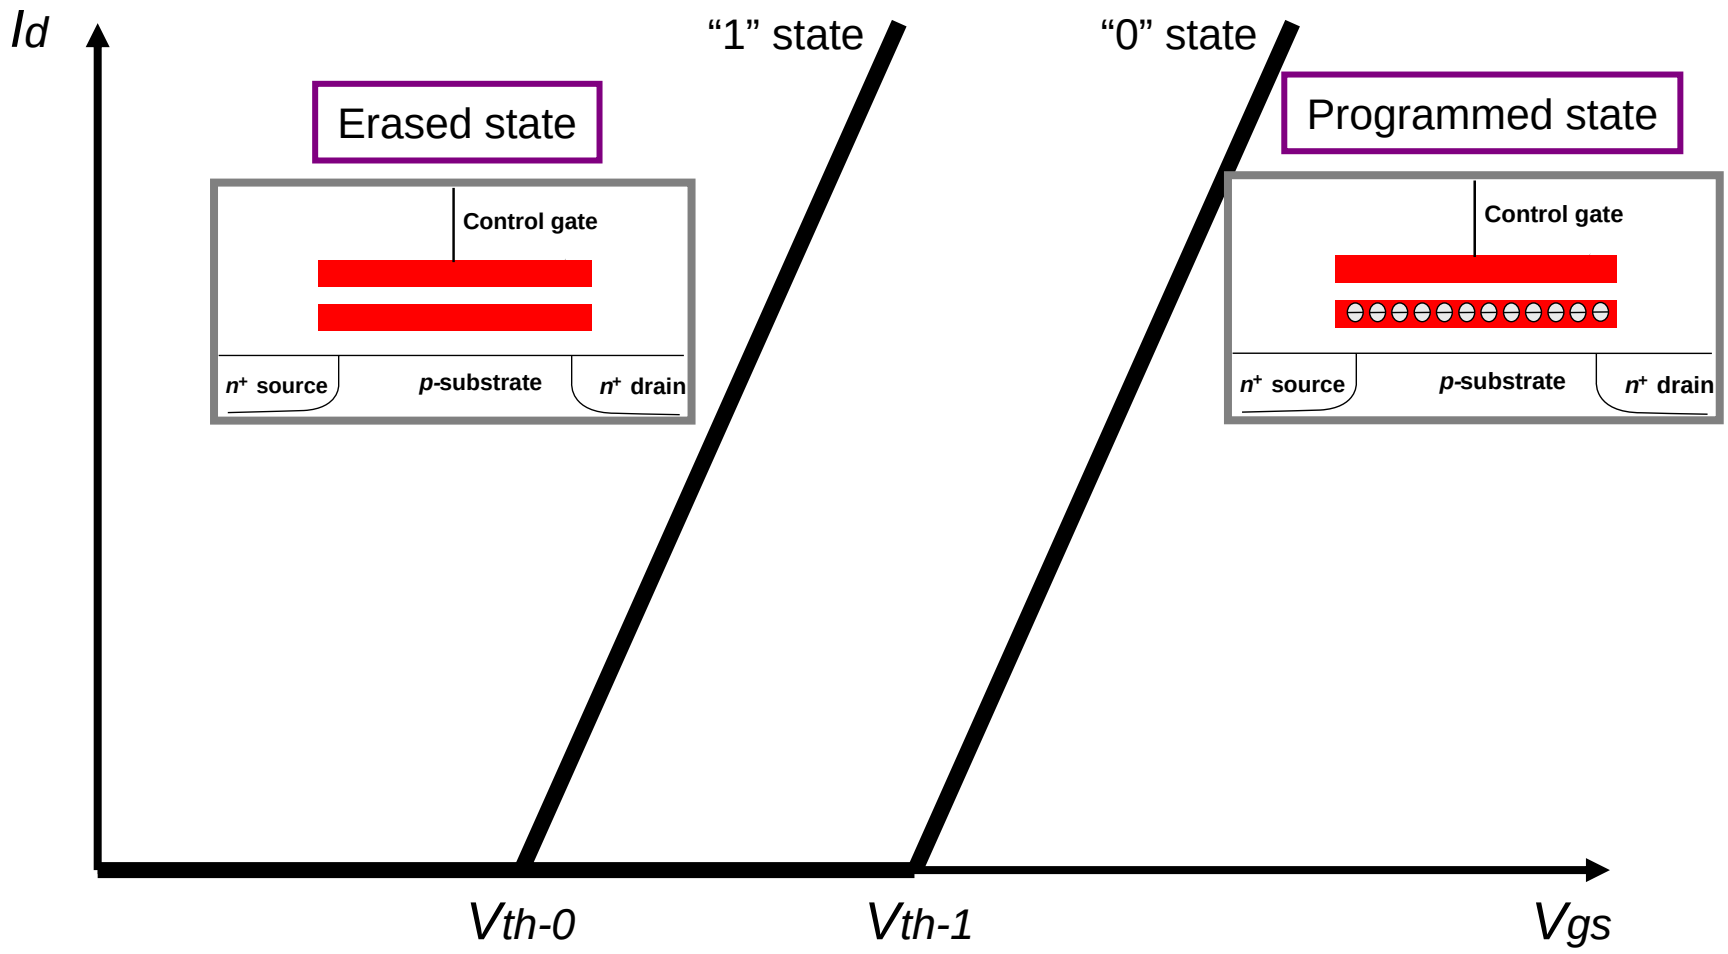

More Bits Per Transistor

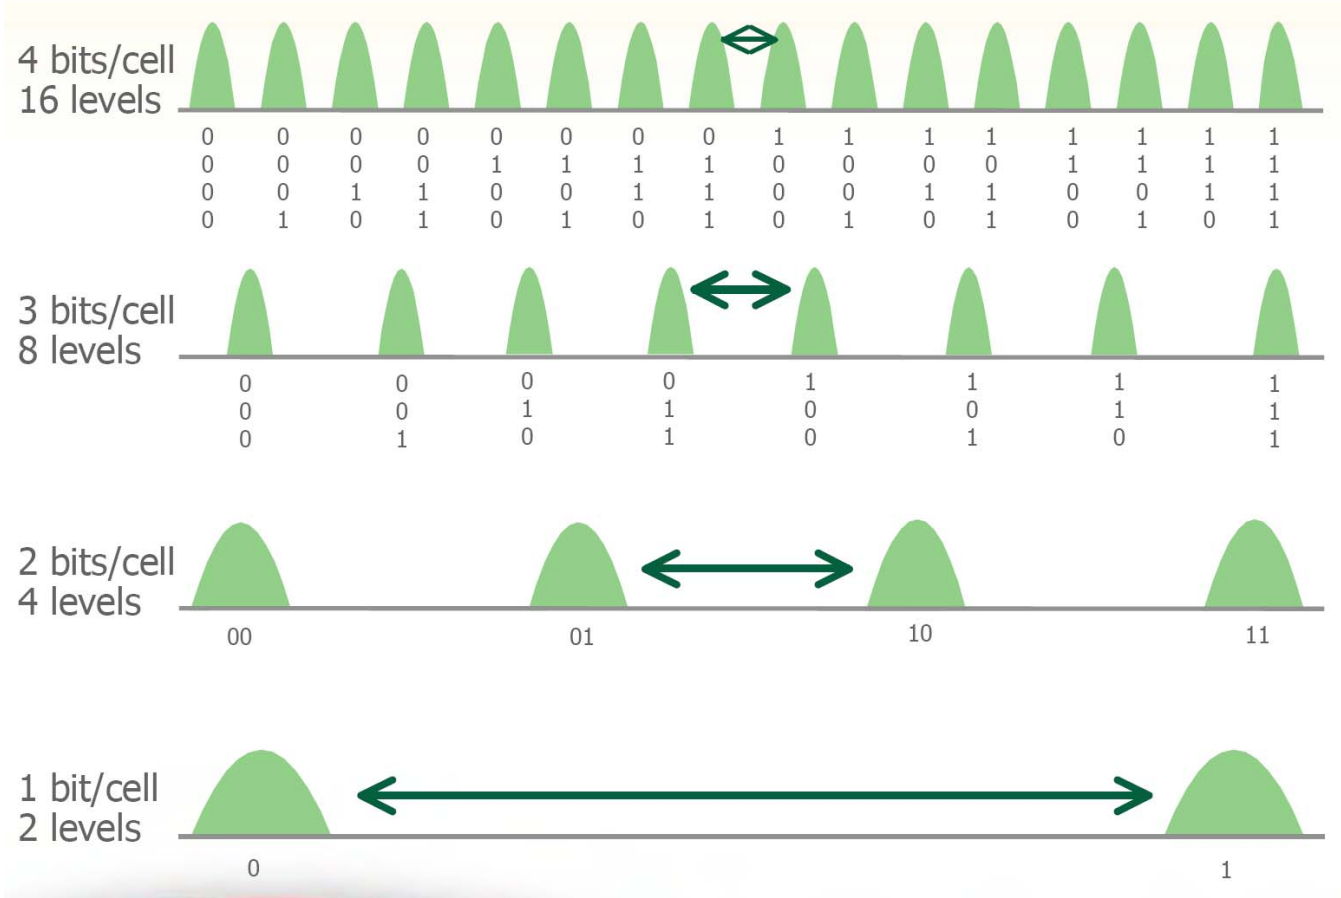

Source: Eli Harari (SanDisk), "NAND at Center Stage," Flash Memory Summit 2007.

NAND Flash Memory Interface

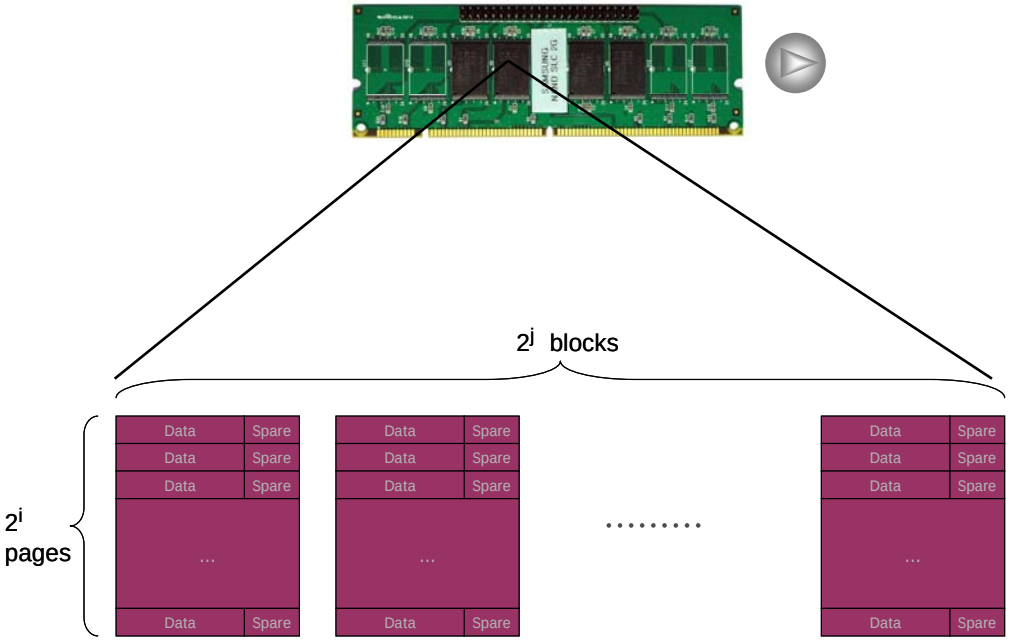

- Read physical page
  - (chip #, block #, page #)
  - ~ 20 us
- Write physical page
  - (chip #, block #, page #)
  - ~ 200 us
- Erase block
  - (chip#, block #)
  - ~ 2 ms

Why (NAND) Flash Memory?

- Advantages of Flash Memory over HDD
  - Low latency
  - Low power consumption
  - Tolerant to shock & vibration
  - Silent operation
  - Small size
  - Abundant parallelism
  - ...

- Single NAND Flash Memory Chip Density Trends

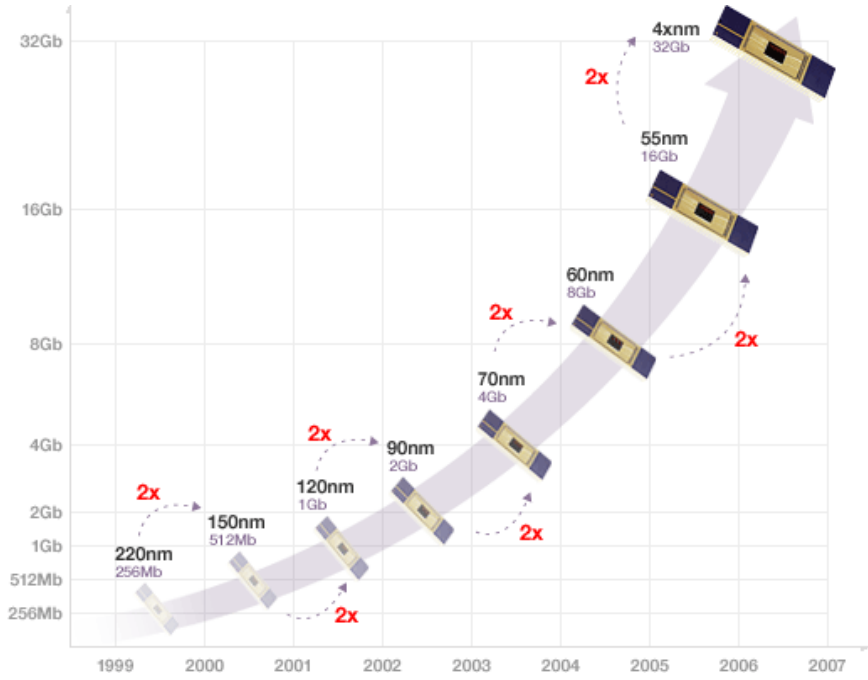

Source: Samsung Electronics

(More) NAND Flash Memory Trends

$ / MB	DRAM	NAND Flash
2000	$0.97	$1.35
2001	0.22	0.43
2002	0.22	0.25
2003	0.17	0.21
2004	0.17	0.10
2005	0.11	0.05
2006	0.096	0.021
2007	0.057	0.012
2008	~0.025	<0.005
CAGR	**-32.1 % / yr**	**-50.0 % / yr**

Source: Lane Mason (Denali Software), "NAND FlashPoint Platform"

(More) NAND Flash Memory Trends

Millions GB	DRAM	NAND Flash
2000	30	1.1
2001	50	1.6
2002	71	4.6
2003	98	14.6
2004	158	68
2005	240	200
2006	340	600
2007	645	1600
2008	1000	4000
CAGR	**+60.0%/yr**	**+150%/yr**

Source: Lane Mason (Denali Software), "NAND FlashPoint Platform"

Solid State Disk

- Provides an interface identical to a hard disk, but uses flash memory as a storage medium

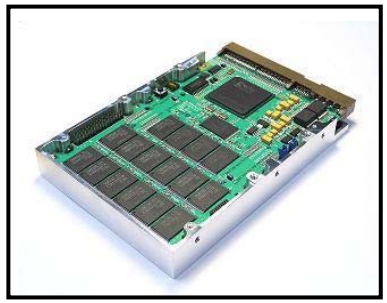

Identical Interface

≡

Solid State Disk: Form Factor Agnostic

	Standard FF			Special FF	
	1.8''	2.5''	1.0''	SLIM	So DIMM
					
Density	4~64GB	4~64GB	4~16GB	4~64GB	8~16GB
Dimension (H x W x T)	78.5x54x8.0	100.2x70x9.5	30x40x4.0	70.6x53.6x: 3.0: 16/32GB 2.5: 4~8GB	53.6x70.6x3.0
Connector	ZIF/IDE 50pin	IDE 44pin	ZIF 35pin	ZIF 40pin	200pin
Weight	44g	46g	TBD	20g	TBD
Market	Notebook	Sub-Note / Tablet	DVC/GPS/ UMPC	UMPC	Custom

Source: Jim Elliot (Samsung Electronics), "SSD: The Next Killer App in NAND Flash," Flash Memory Summit 2007.

Flash memory summary

The Good
- Low latency
- Low power consumption
- High Reliability
- Small form factor
- Massive parallelism

....

FROM *THE DARK NIGHT*

The Bad
- No in-place updating
- Limited endurance
- Bad blocks
- Write disturbance
- Read disturbance

The Ugly
- Retention errors
- Paired page problem

Outline

- HDD Basics and Demo
- Flash Memory Basics and Demo
- Storage Trends
- Conclusions

Storage Trends

Tape Is Dead

Disk Is Tape

- 1 TB disks are available
- 10+ TB disks are predicted in 5 years
- But: ~5..15 **hours to read (sequential)**

  ~15..150 **days to read (random)**
- Need to treat most of disk as Cold-storage archive

Source: Jim Gray (Microsoft), "Tape is Dead, Disk is Tape, Flash is Disk, RAM Locality is King"

Storage Trends

Disk Is Tape

Flash Is Disk

- 1995 16 Mb NAND flash chips

  2005 16 Gb NAND flash chips
- 2012 1 Tb NAND flash chips

  == 128 GB chip

  == 1 TB or 2 TB solid state disk for ~$400

  or 128 GB solid state disk for ~$40

  or 32 GB solid state disk for ~$5

Source: Jim Gray (Microsoft), "Tape is Dead, Disk is Tape, Flash is Disk, RAM Locality is King"

Disk is Tape / Flash is Disk

Poor Reliability
Carnegie Mellon & Google study show up to 8.6% annual failure rate for HDD in controlled environment

Heat
Rotating platters & moving heads need power → produces heat

Low Performance
Low IOPS performance → High redundancy to compensate for low performance per drive

High TCO
Initial purchase cost low, but maintenance, space, cooling & replacement will increase TCO substantially

Source: Esther Spanjer (Adtron), "Enterprise SSD: The next killer app," Flash Memory Summit 2007.

Disk is Tape / Flash is Disk

- **Performance**

1 SSD

35~50 HDDs

Source: Jim Gray (Microsoft), “Tape is Dead, Disk is Tape, Flash is Disk, RAM Locality is King”

Disk is Tape / Flash is Disk

- **Power Consumption**

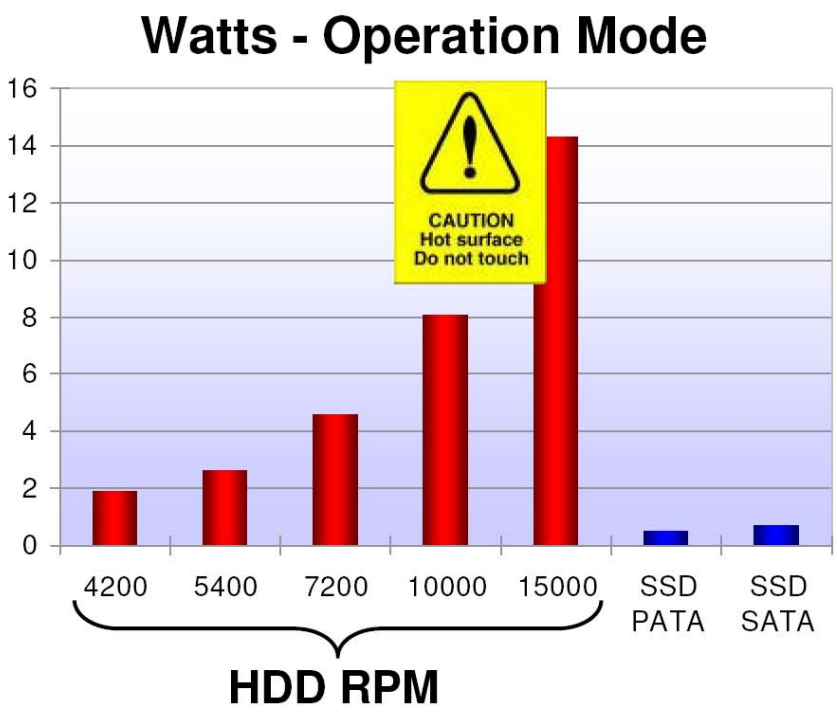

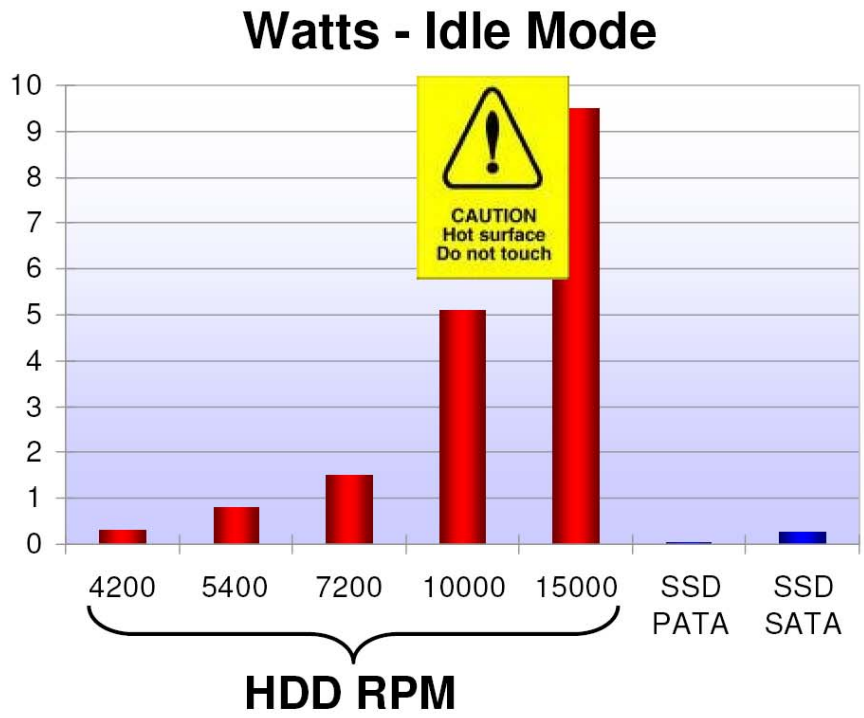

Source: Jim Elliot (Samsung Electronics), "SSD: The Next Killer App in NAND Flash," Flash Memory Summit 2007.

Future Outlook

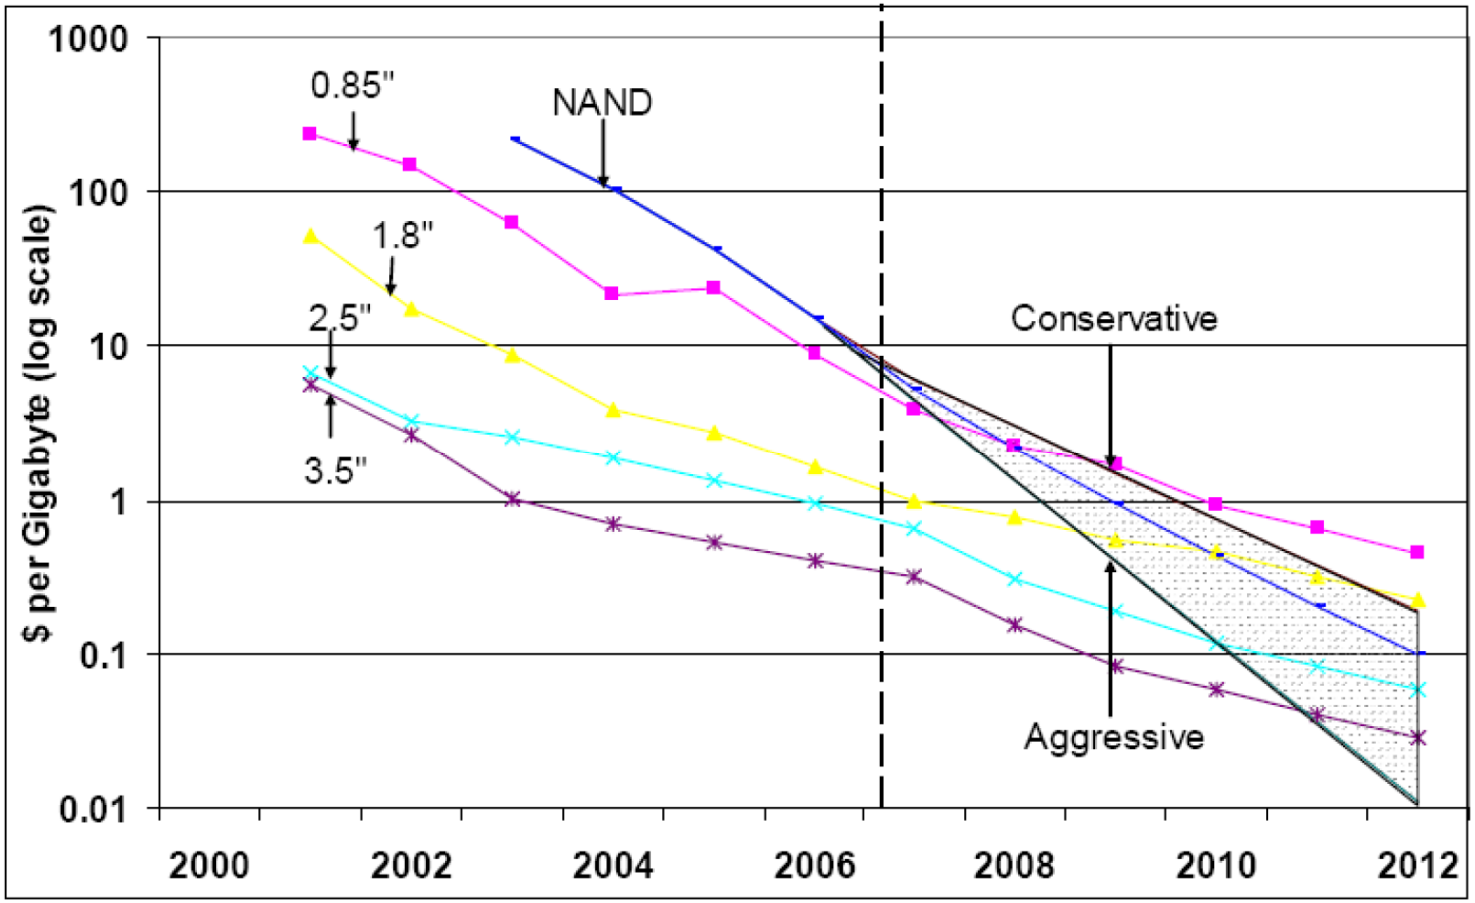

Source: Scott Deutsch (SanDisk), "Bringing Solid State Drives to Mainstream Notebooks," Flash Memory Summit 2007.

Outline

- HDD Basics and Demo
- Flash Memory Basics and Demo
- Storage Trends
- Conclusions

Conclusions

- In the animal world
  - Survival of the fittest

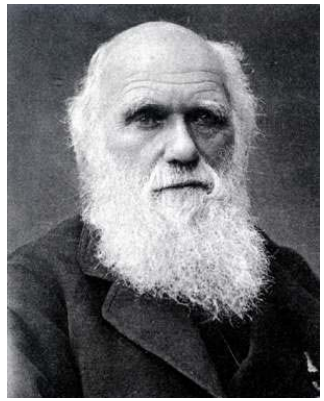

- In the memory world
  - Survival of the fastest or cheapest

	Volatile	Non-volatile
Fastest	SRAM	FRAM?
Cheapest	DRAM	NAND Flash HDD

Conclusions

- From the history

	IBM 360/85	IBM 360/91
Clock Rate	80 ns	60 ns
Memory Speed	1040 ns	750 ns
Memory Interleaving	4 way	8 way
Additional Features	Cache Memory	Register Renaming, Out-of-order Execution, *etc*

But, IBM 360/85 faster on 8 of 11 programs!

Source: David Patterson, *et al.,* "A Case for Intelligent DRAM: IRAM", Hot Chips VIII, August, 1996

The Ultimate Limit – HDD

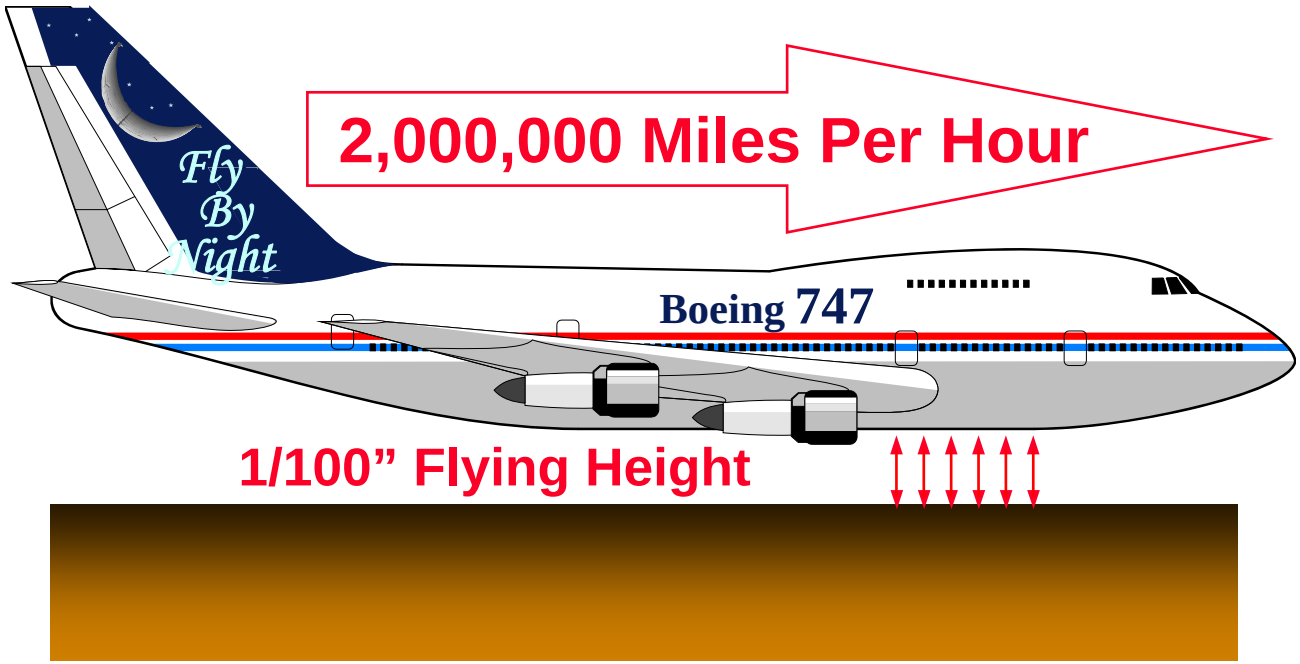

Source: Richard Lary, The New Storage Landscape: Forces shaping the storage economy, 2003.

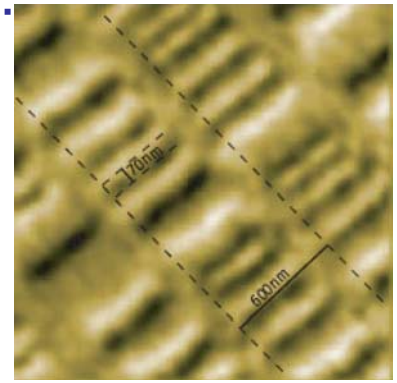

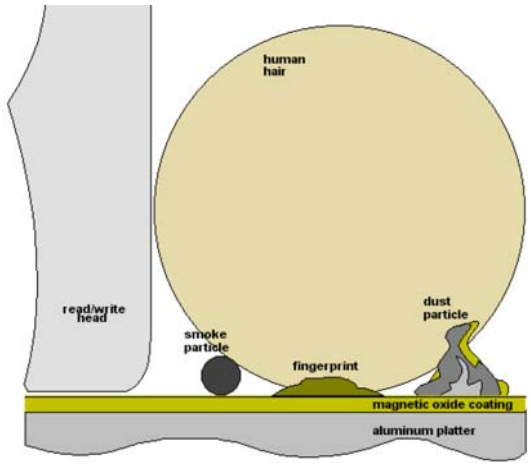

Source: B. Parhami, Dependable Computing: A Multilevel Approach

The Ultimate Limit – Flash Memory

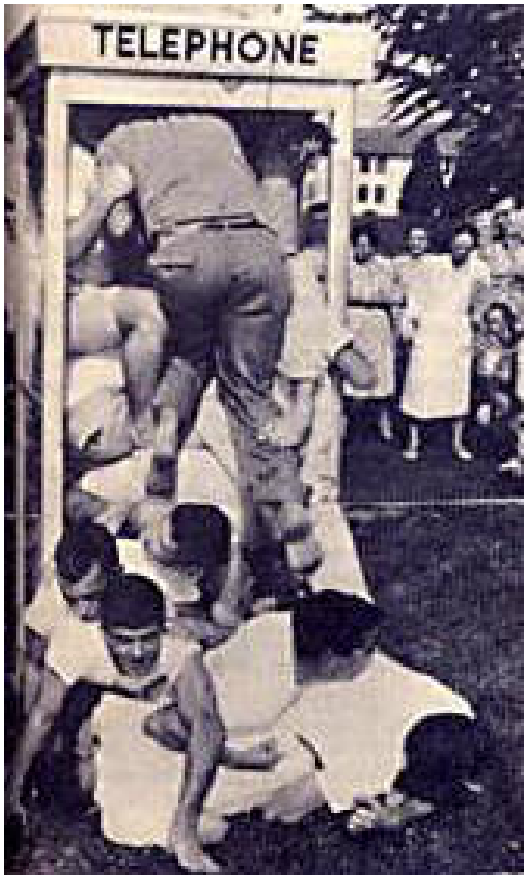

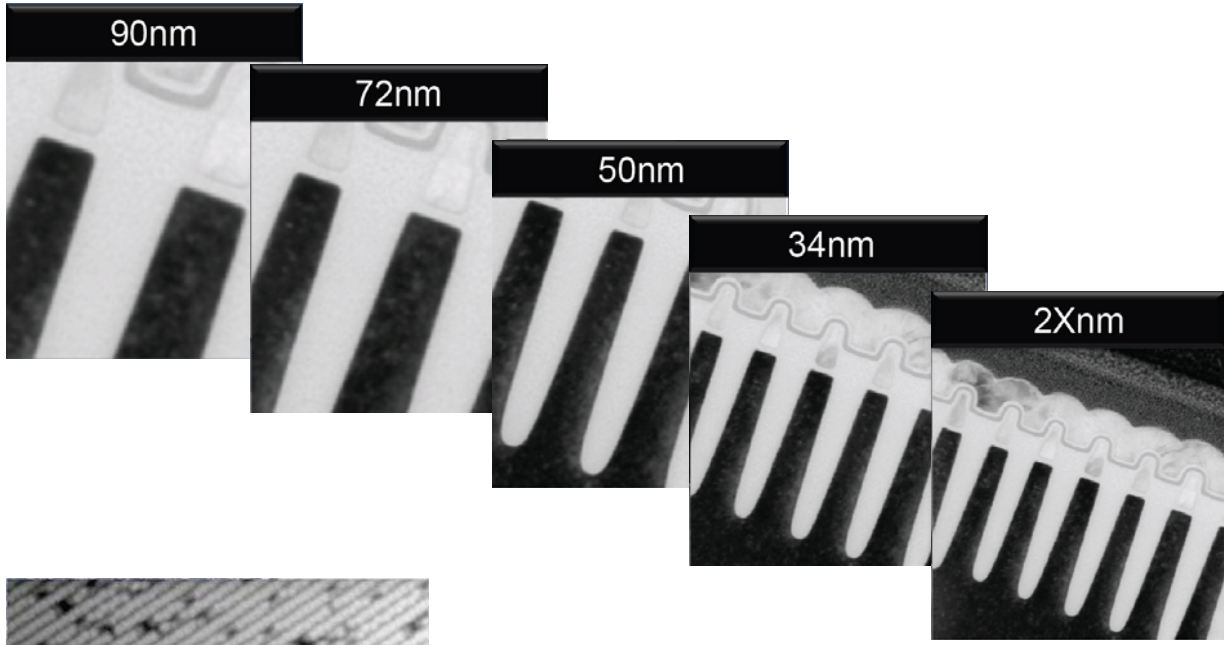

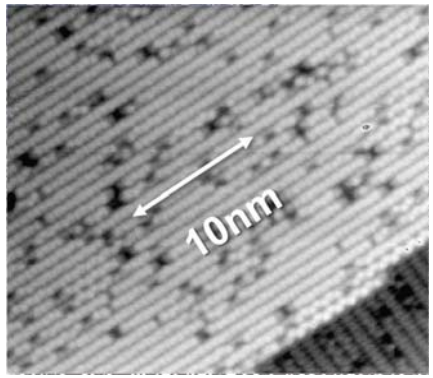

Scanning tunneling microscope image of a silicon surface showing 10 nm is ~20 atoms across

Source: B. Shirley, “The Many Flavors of NAND … and More to Come,” Flash Memory Summit 2009

Outline

- HDD Basics and Demo
- Flash Memory Basics and Demo
- Storage Trends
- Conclusions
- (More Demos)

Flash Memory Software Development Platforms

Embedded Platform

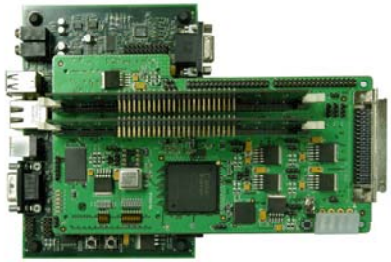

Embedded Flash Memory 소프트웨어 솔루션 개발용

SSD Platform

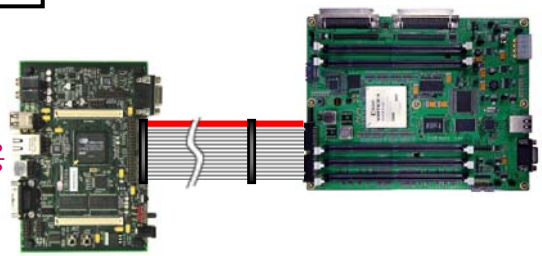

Solid State Disk 소프트웨어 개발용

Flash / NV-RAM Modules

Samsung SLC NAND

RAMTRON FRAM (serial)

Samsung MLC NAND

RAMTRON FRAM (parallel)

Samsung OneNAND

FREESCALE MRAM (parallel)

Hynix MLC NAND

Samsung Phase-change RAM

Embedded Platform

Features

- **For embedded Flash memory software development**
- **FPGA-based**
- **NAND slot x 2**
- **DAQ (Data Acquisition) interface**

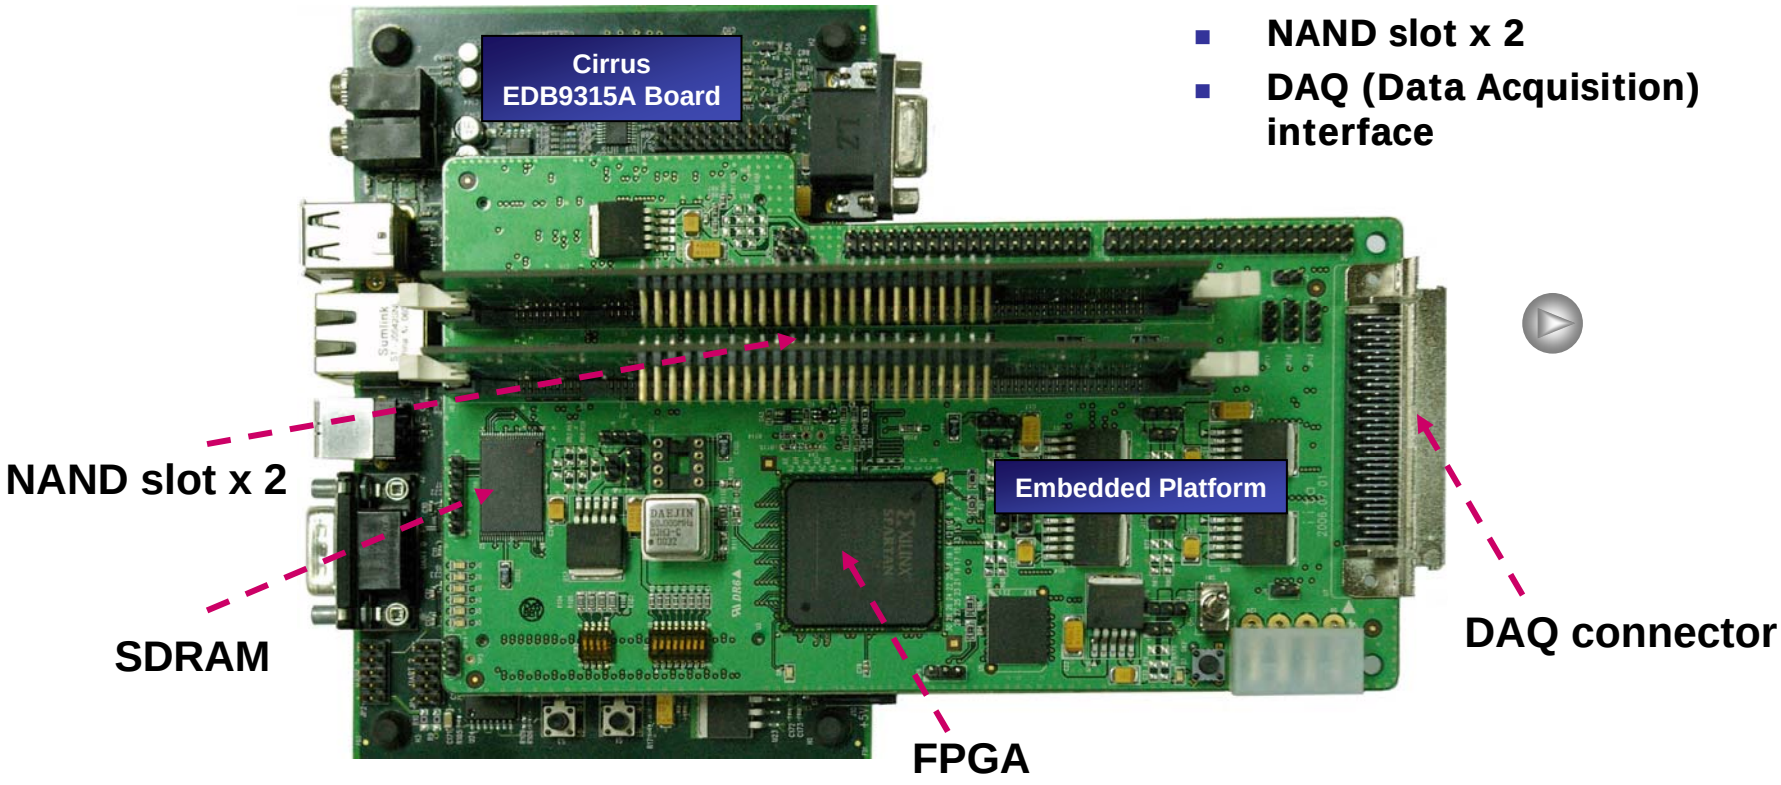

SSD (Solid State Disk) Platform

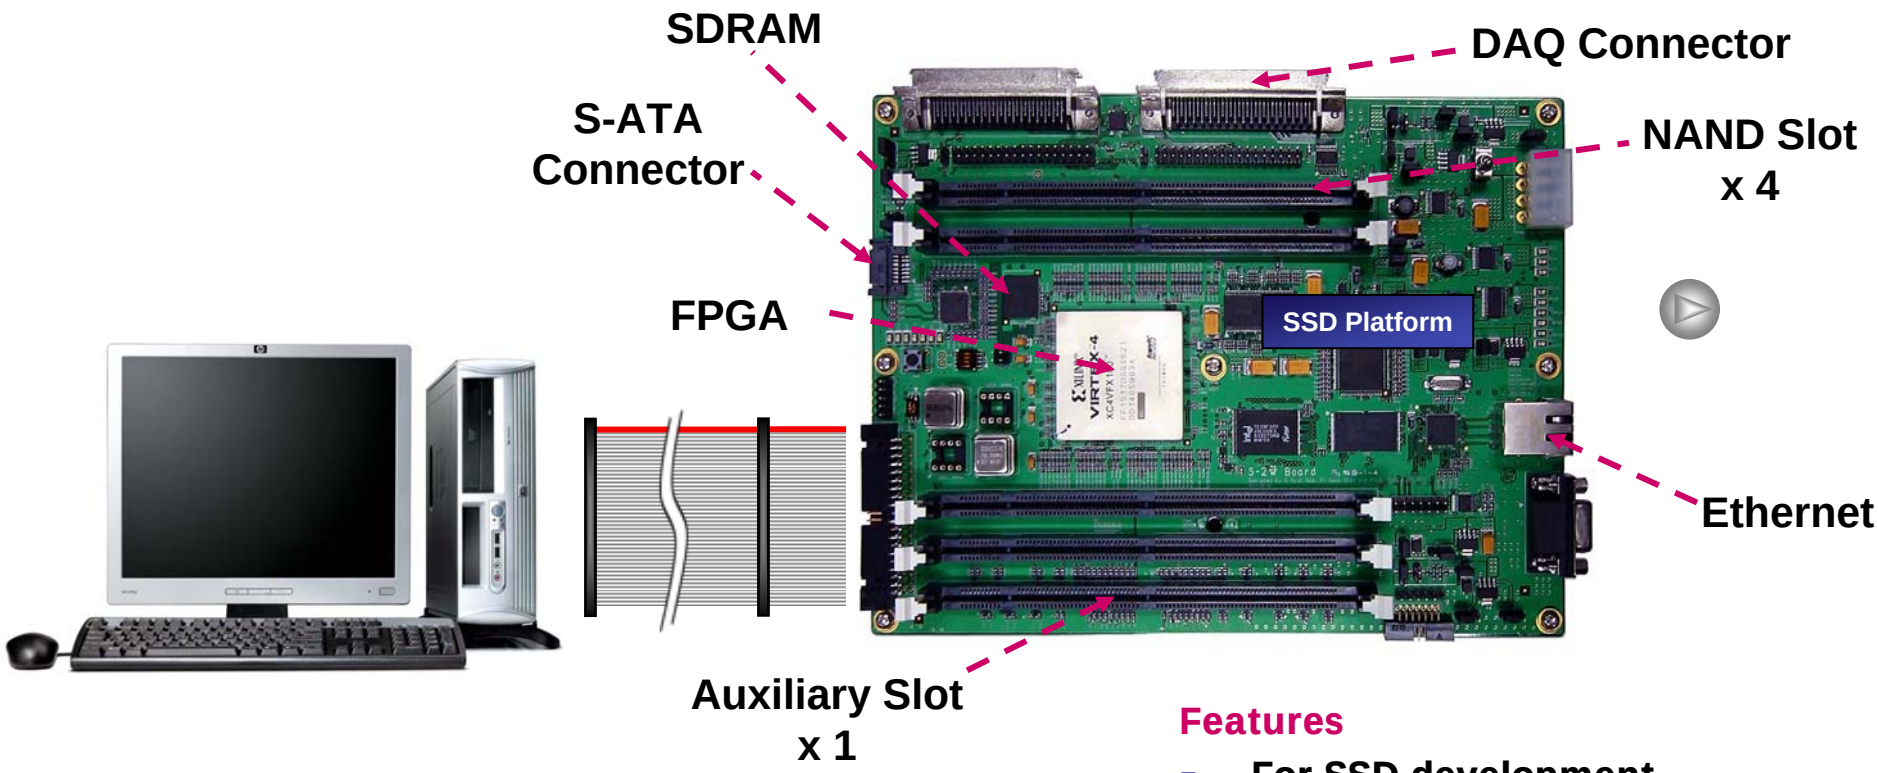

Features

- For SSD development
- FPGA-based
- SSD interface (P-ATA, S-ATA)
- NAND slot x 4