Current Trends in Flash Memory Technology

2006. 01. 25

Sang Lyul Min
Eyee Hyun Nam
Seoul National University

Agenda

- Overview of Portable Storage Technologies
- High Performance Flash Memory Controller
- Hybrid Hard Disk Drive
- Conclusions

Portable Storage Applications

The Contenders for Portable Storage Market

Flash Drive

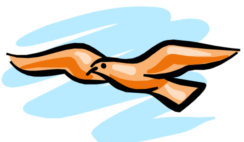

Toshiba

Samsung

Hitachi

SanDisk

LexarMedia

Portable Storage

Micro Drive

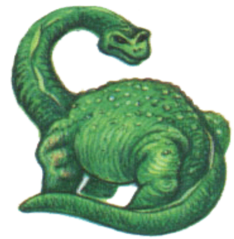

IBM

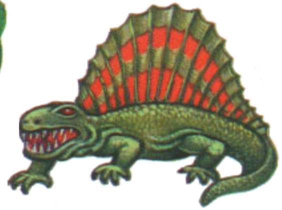

Hitachi

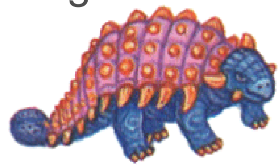

Seagate

Cost Comparison (2004. 8)

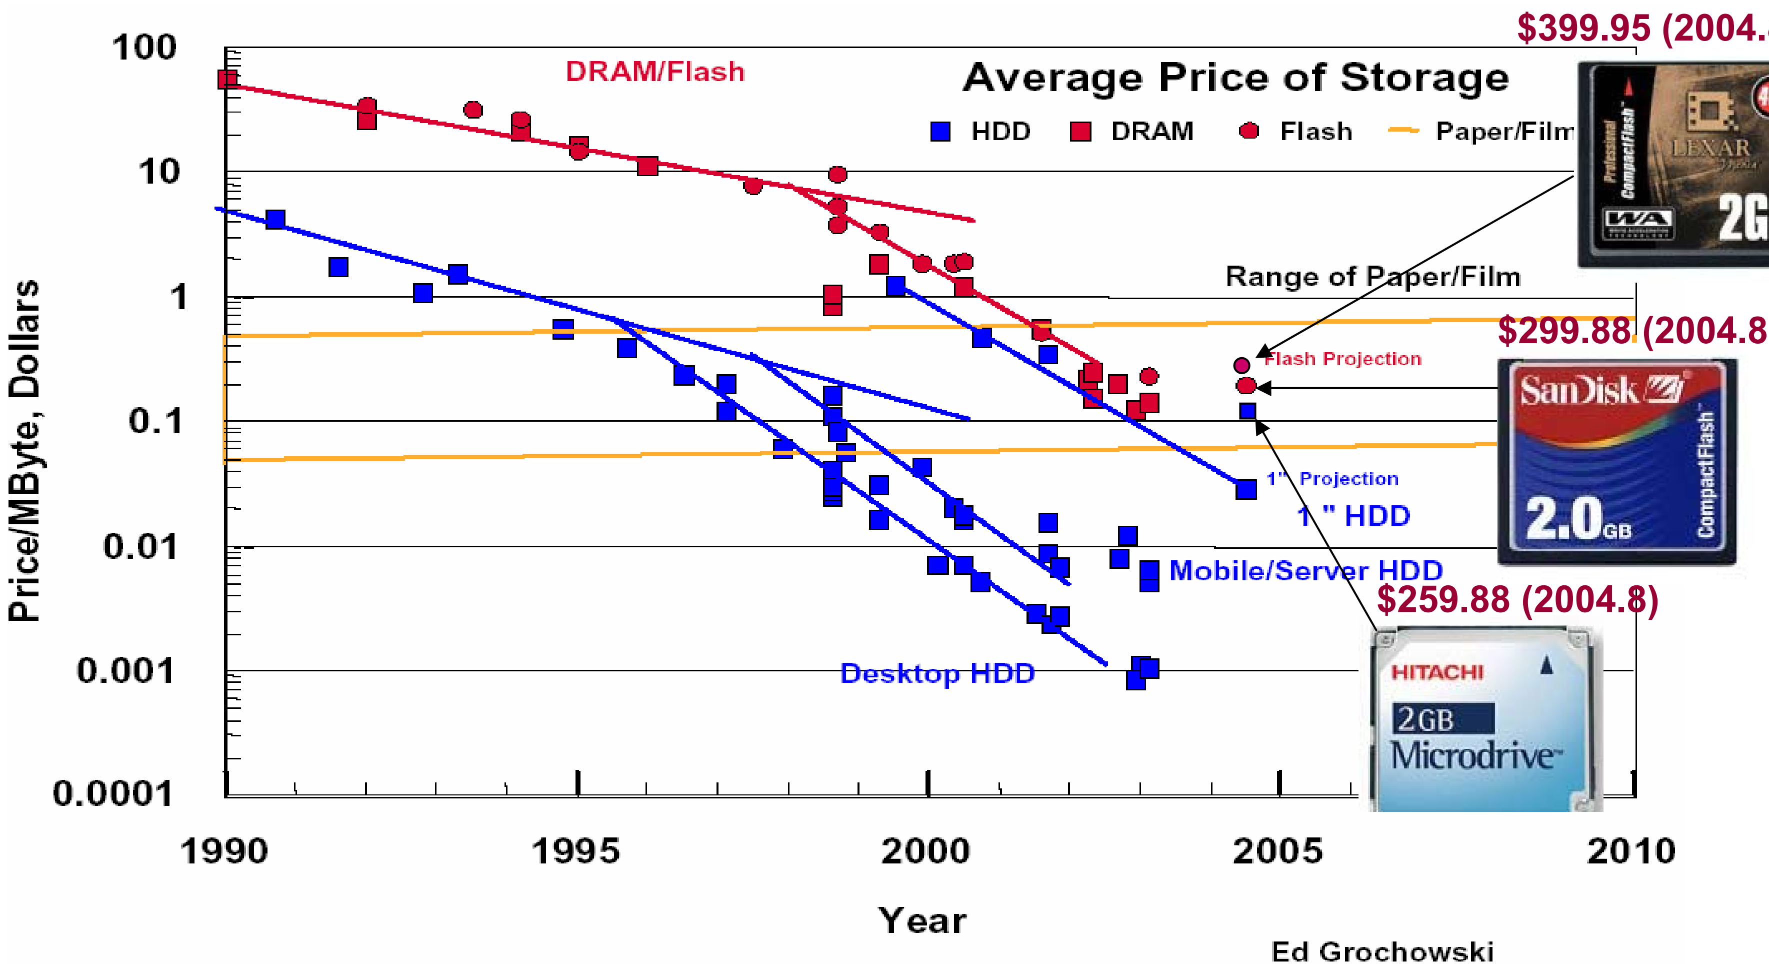

Source: http://www.hitachigst.com/hdd/technolo/overview/chart03.html

Cost Comparison (2005. 4)

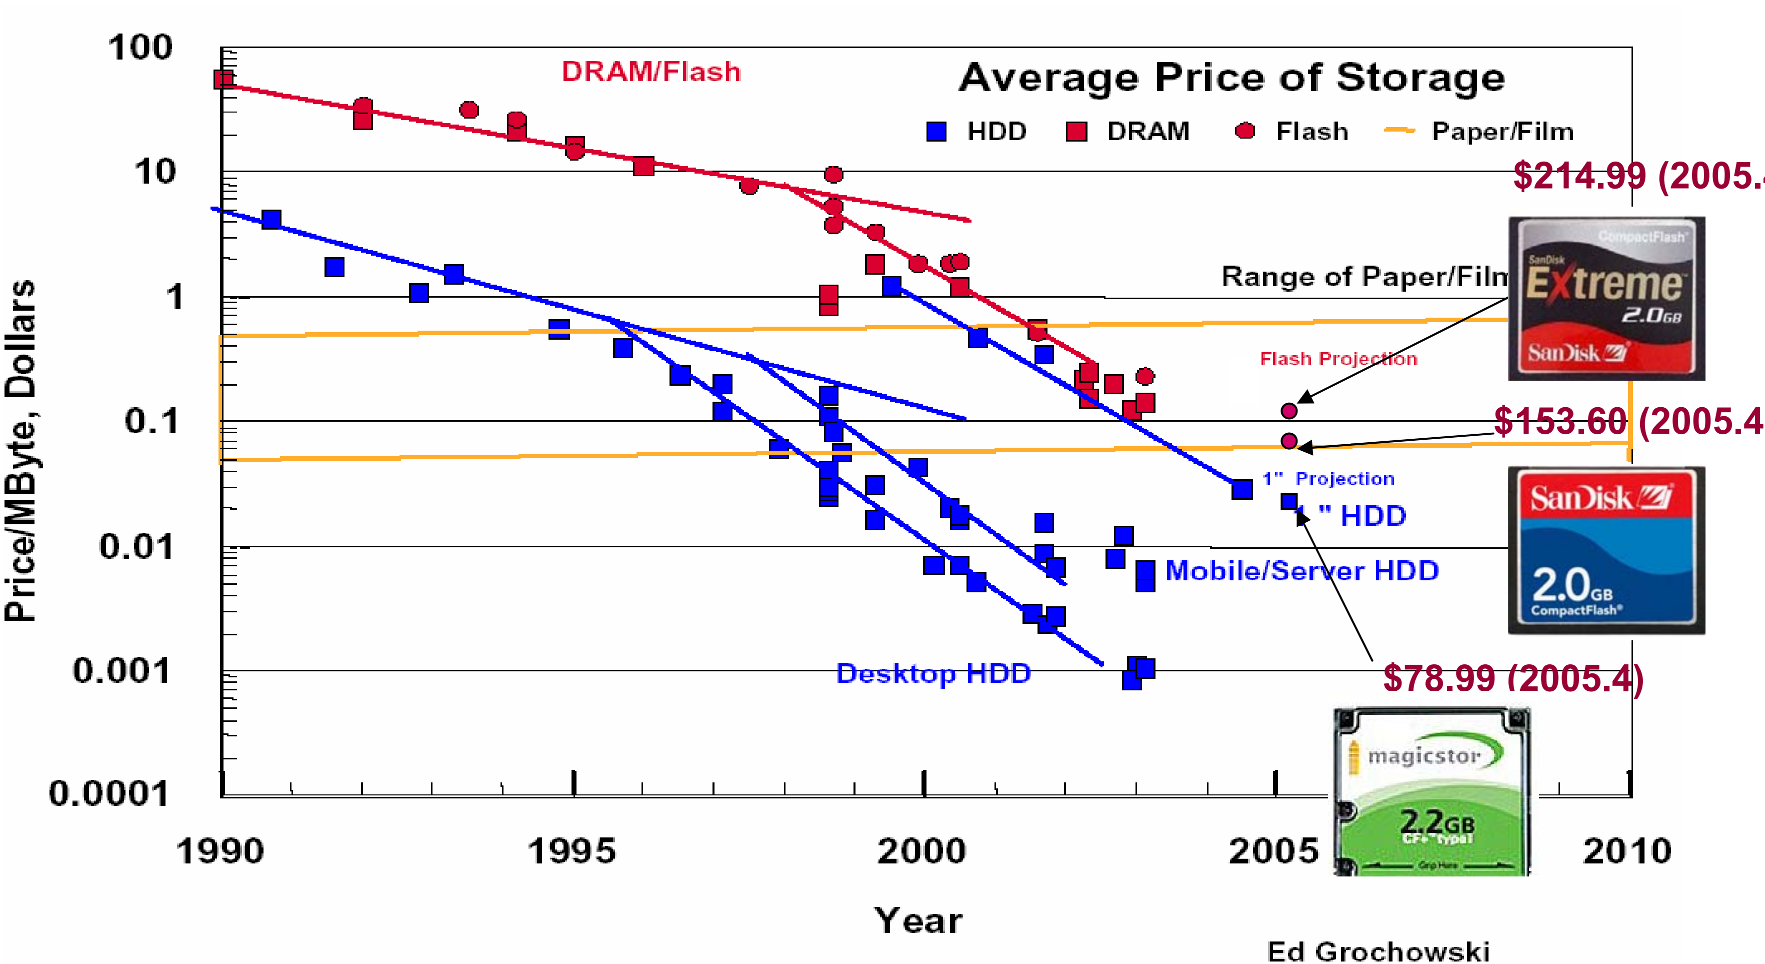

Source: http://www.hitachigst.com/hdd/technolo/overview/chart03.html

NAND Flash Memory Basics

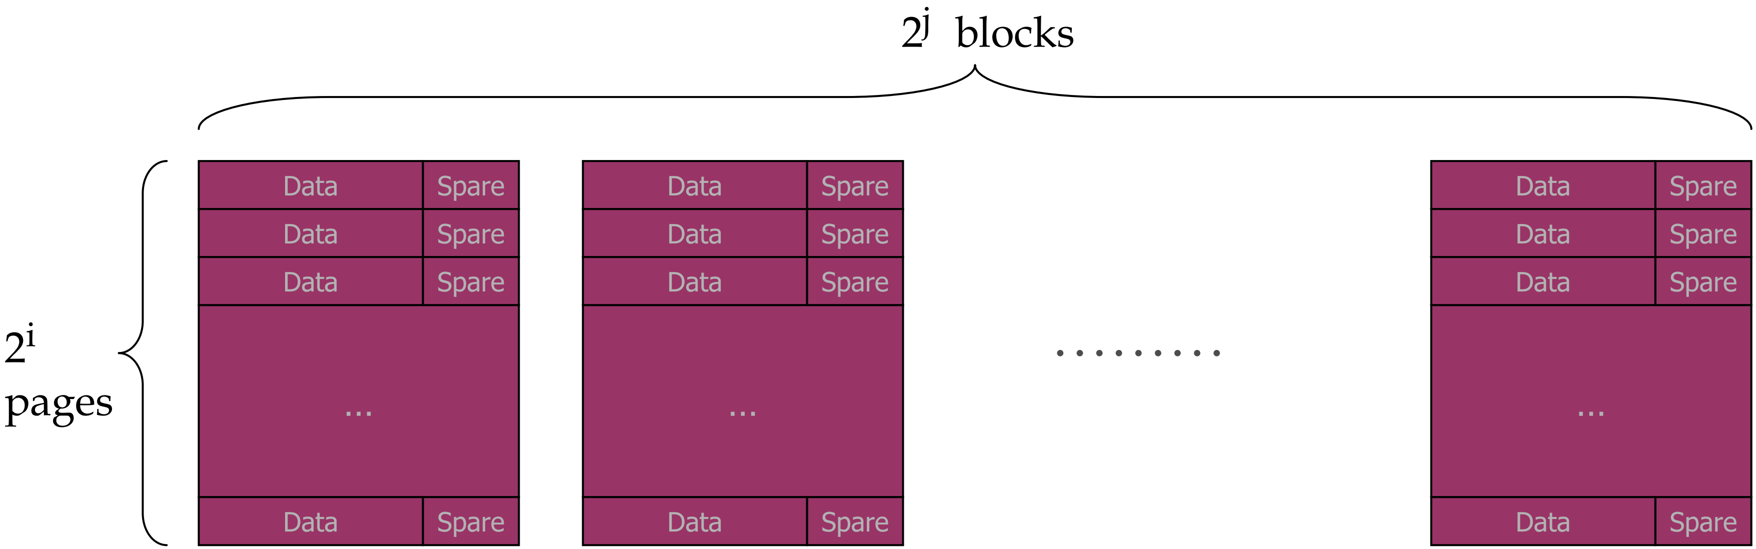

- Read physical page
  - (chip #, block #, page #)
  - ~ 25 us
- Write physical page
  - (chip #, block #, page #)
  - ~ 300 us
- Erase block
  - (chip#, block #)
  - ~ 2 ms

FTL (Flash Translation Layer)

- Definition
  - Software layer that makes flash memory appear to the system like a disk drive

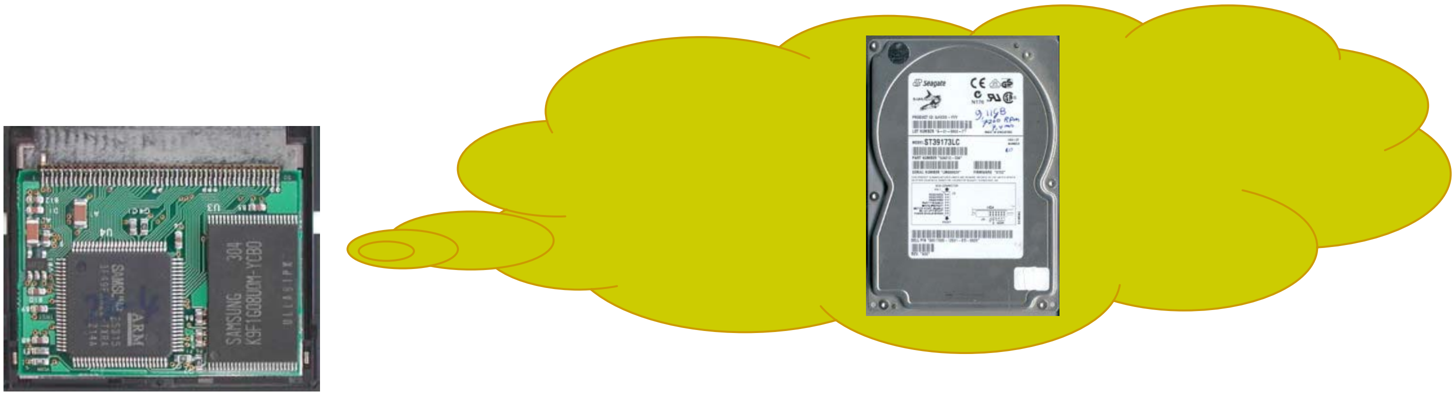

- Challenges in FTL
  - No overwrite is allowed without erasing
  - Asymmetry in read and write speeds

Logical interface for a disk drive

- Operations
  1. Identify drive(): returns N
  2. Read sectors(start sector #, # of sectors)
  3. Write sectors(start sector #, # of sectors)

Block level mapping

- Logical blocks

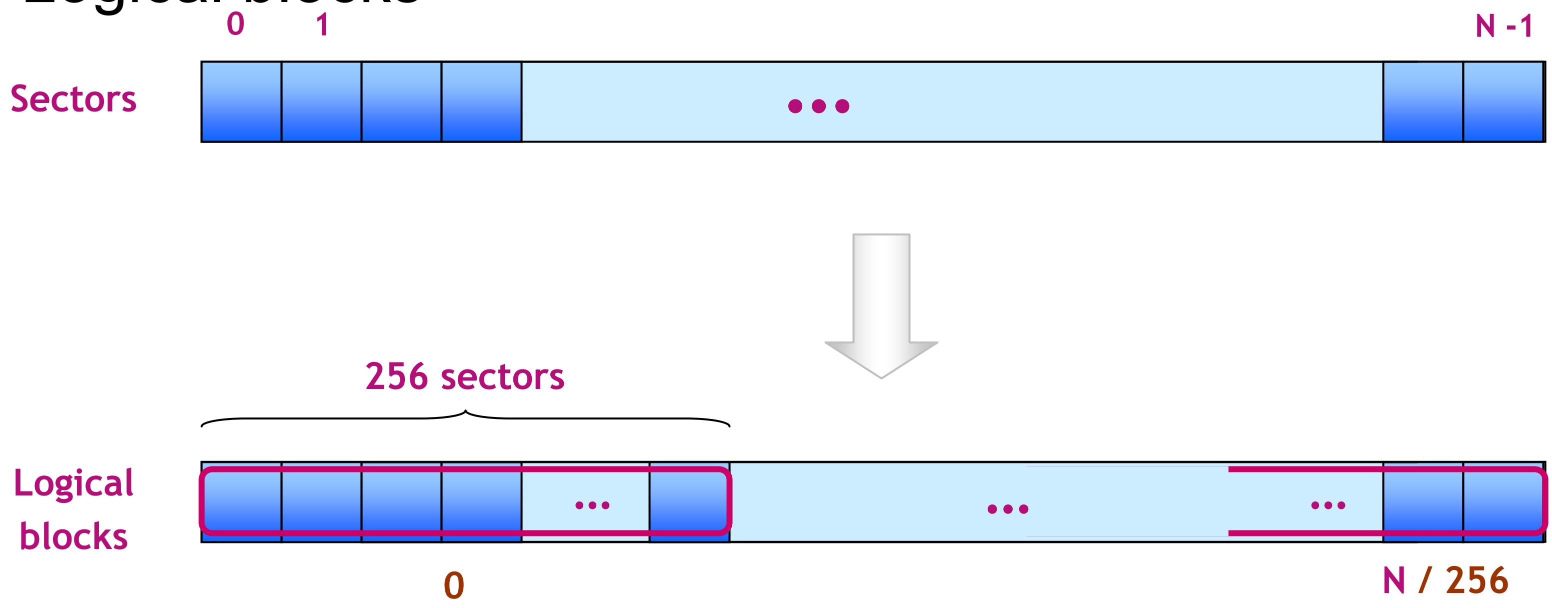

Block level mapping

- Logical to physical block mapping

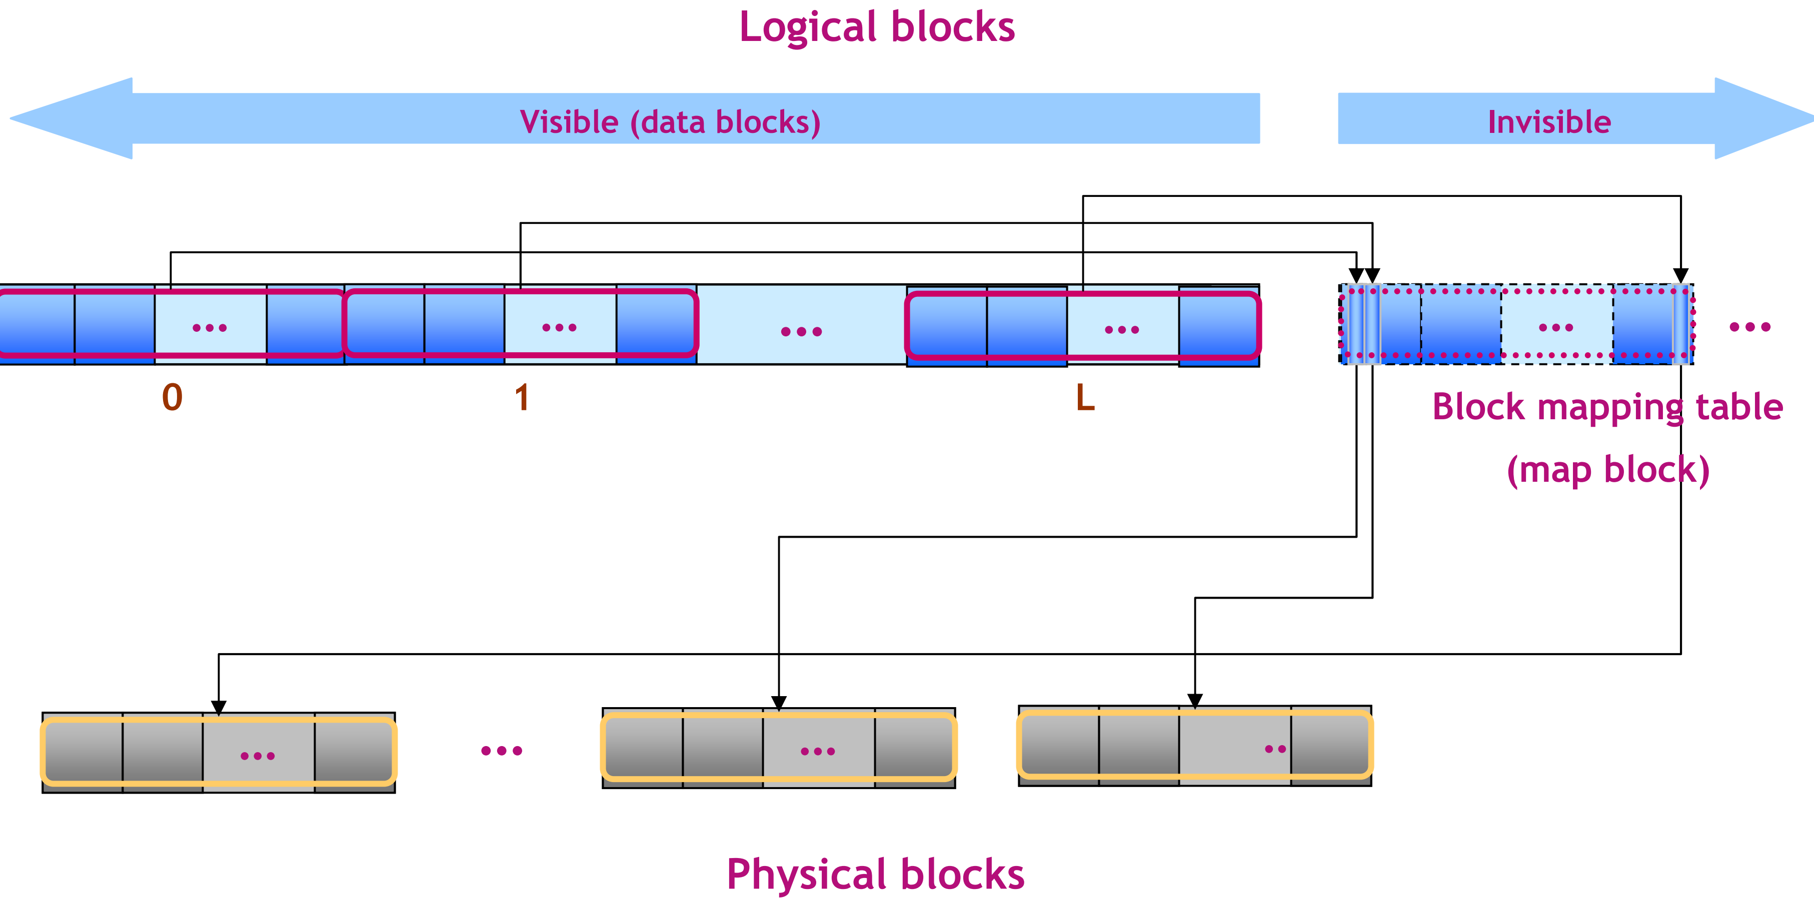

Read procedure

- Ex. read 3 sectors from 255

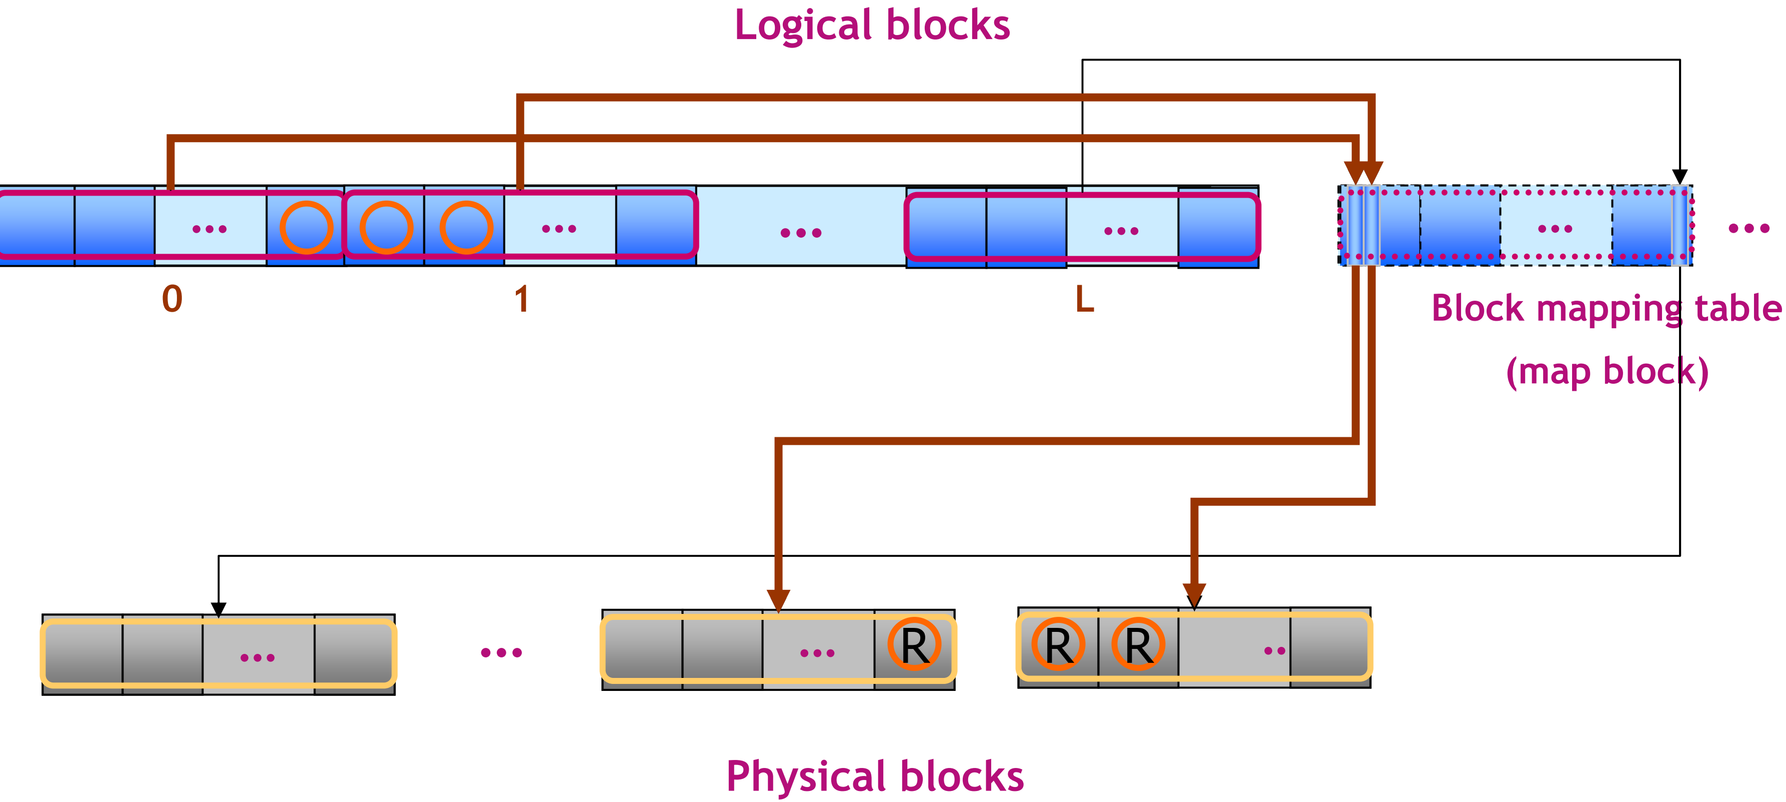

Write procedure (Data block update)

- Ex. write 3 sectors from 255

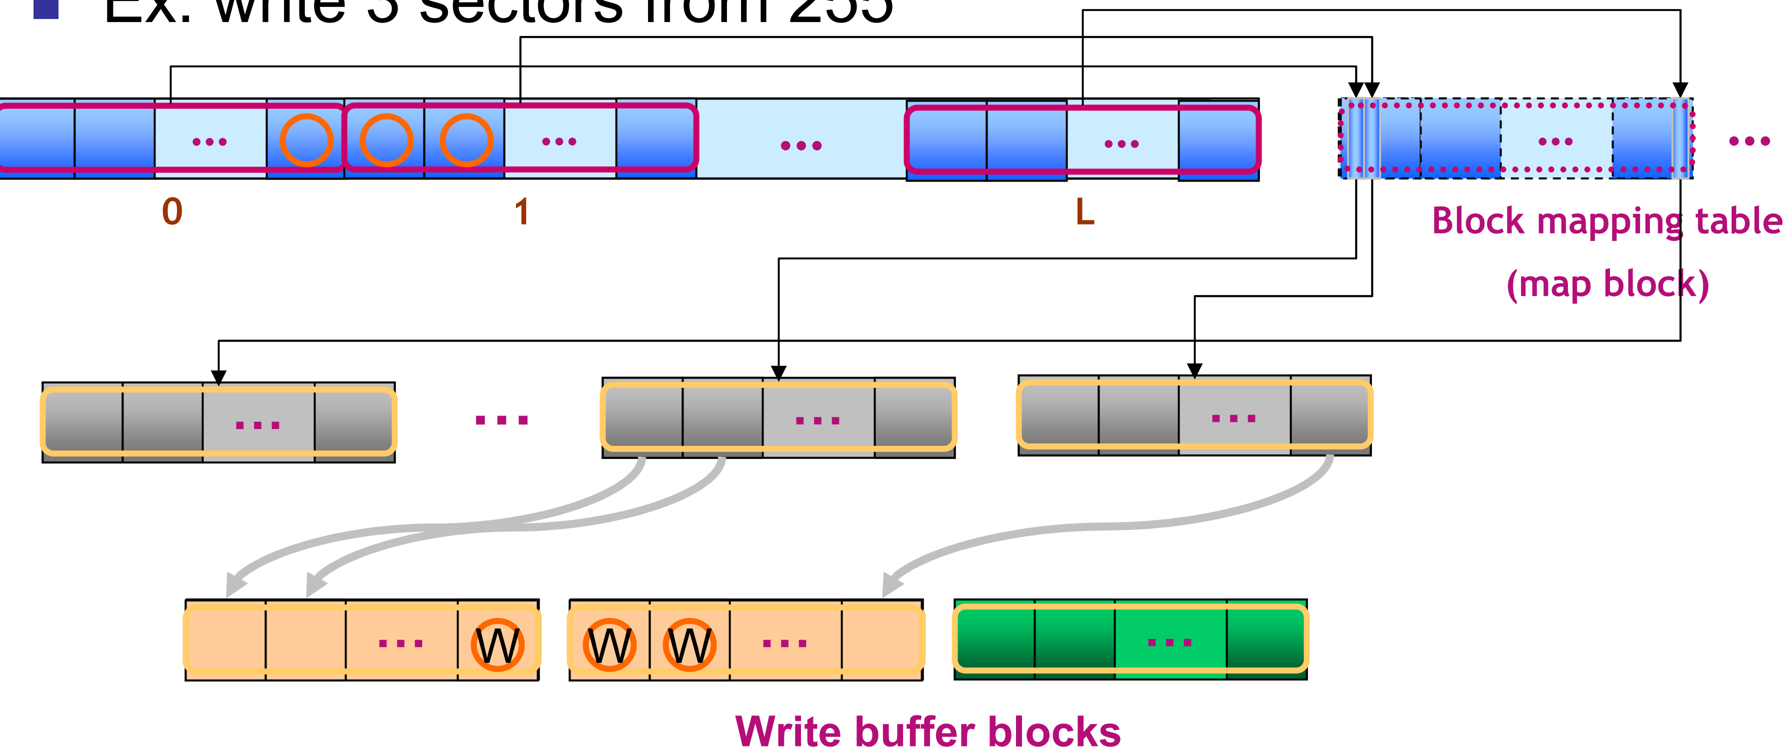

Still, update of mapping information is needed

Write procedure (Map block update)

- Ex. write 3 sectors from 255

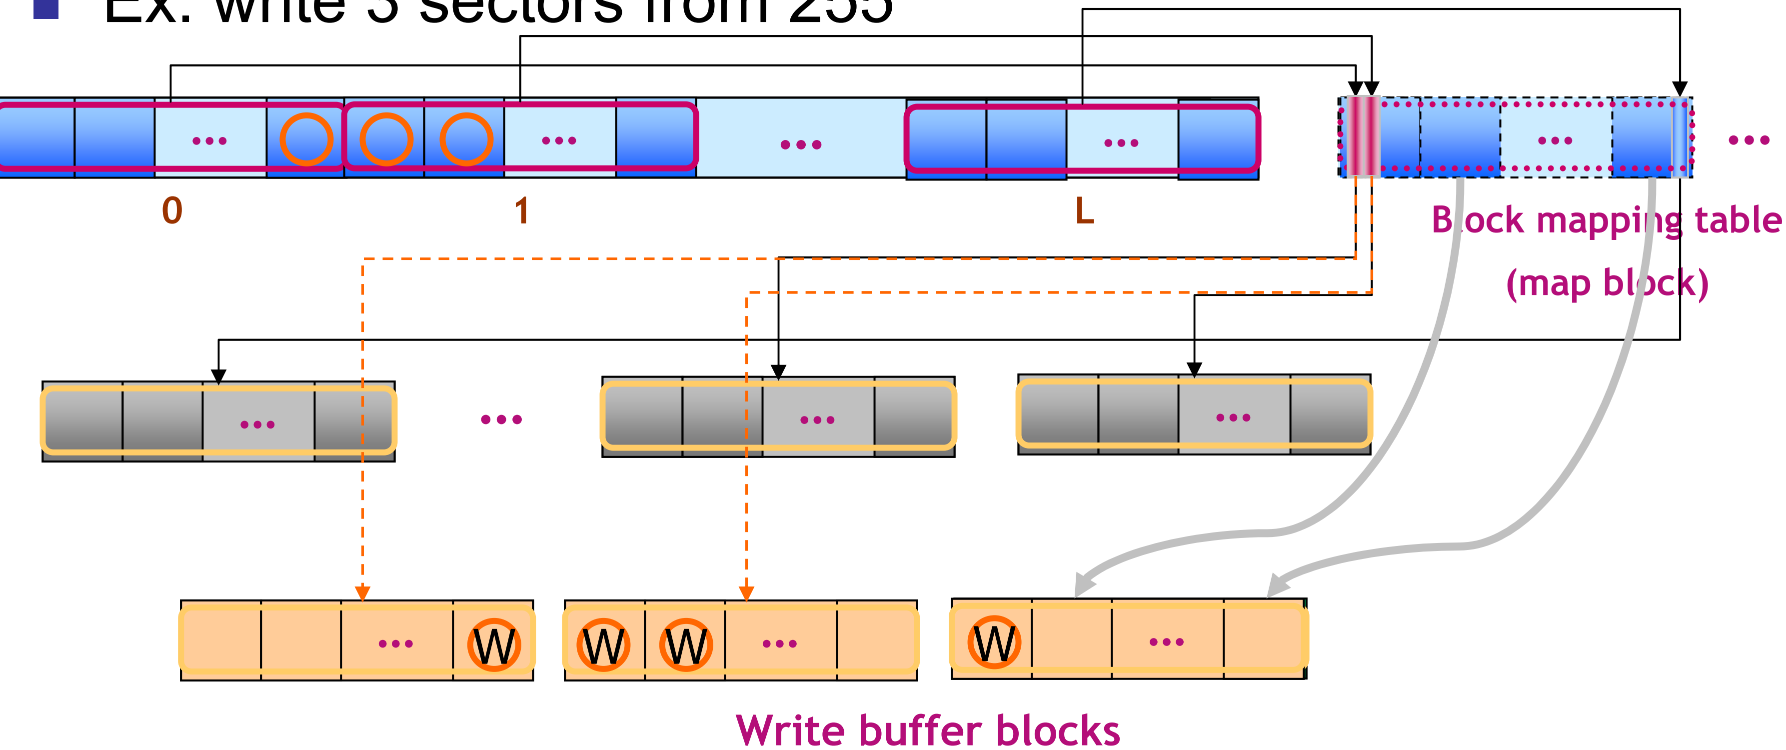

Still, somewhere we need to keep the addresses of new map and write buffer blocks (i.e., logging)

Inside Flash Drive

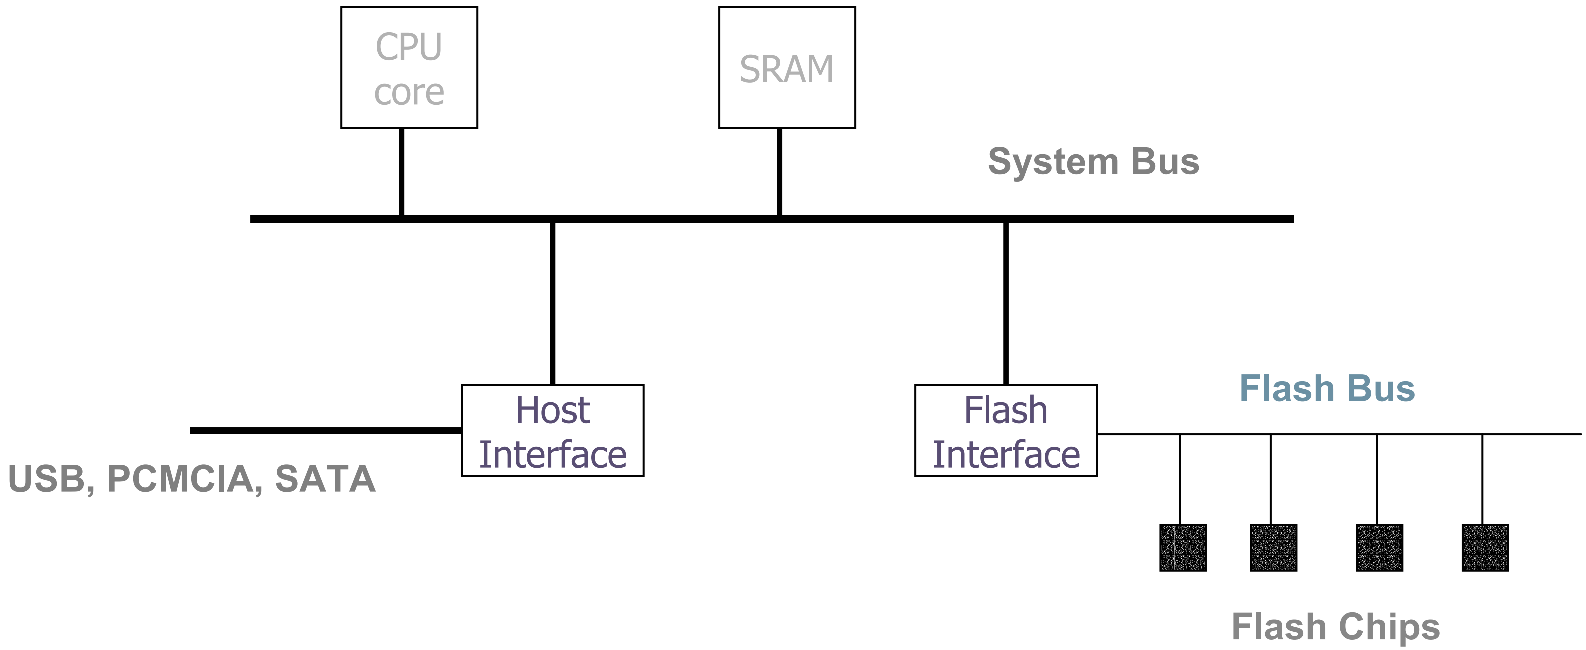

Host Interface Performance

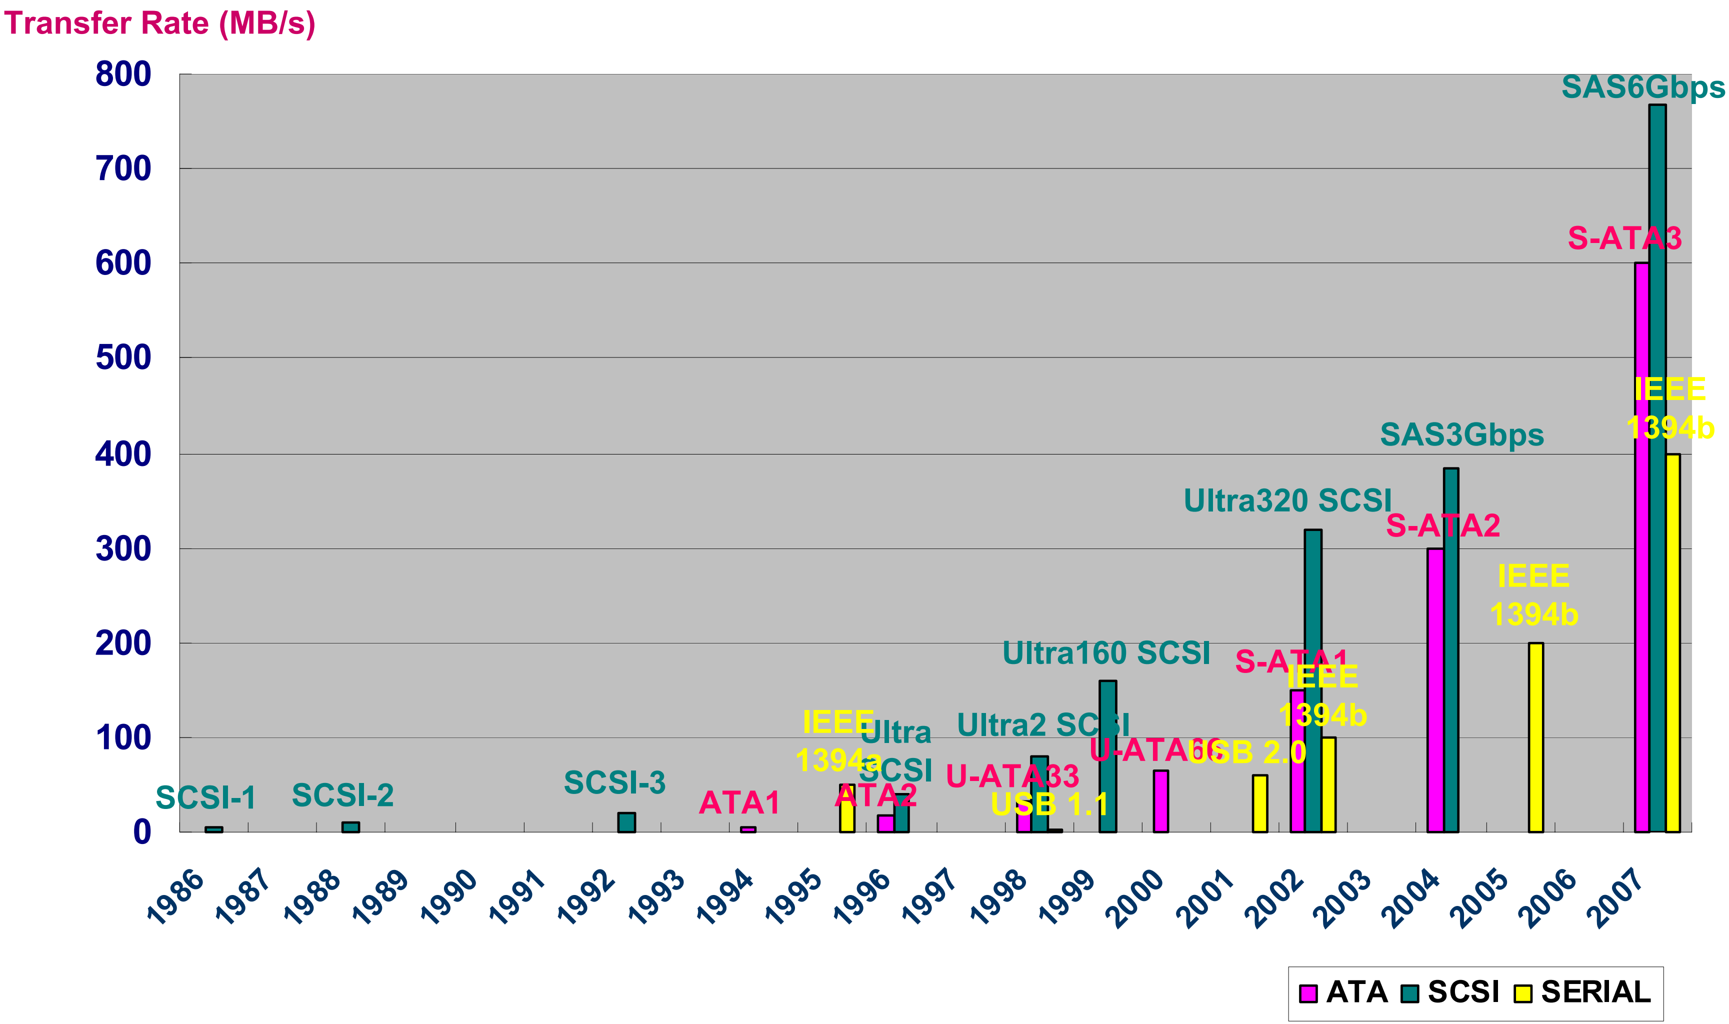

Flash Chip Bandwidth

- Write bandwidth = 2KB/300us = 6.7MB/s per chip
- Read bandwidth = 2KB/25us = 80MB/s per chip
- Erase bandwidth = 128KB/2ms = 64MB/s per chip

Flash bus bandwidth picture

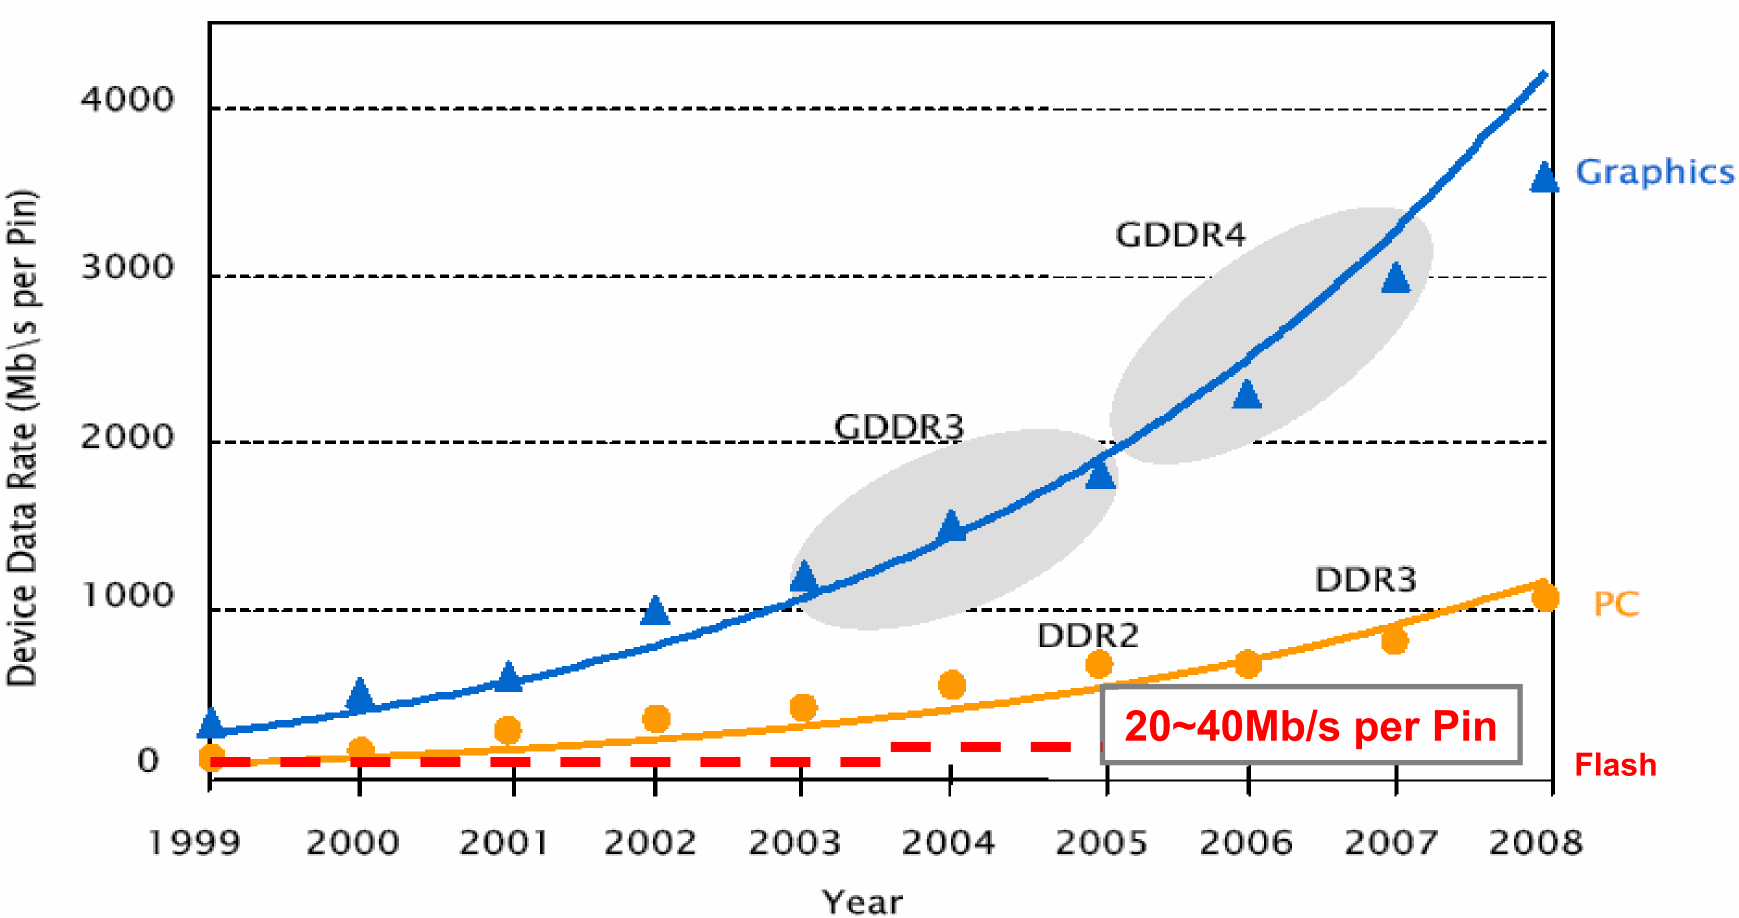

Source: Terry Lee, Micron Technology, Inc, VTF (VIA Technology Forum) 2003


Flash Bus & Flash Chips
Host Interface & System Bus
http://laughingriot.incide.net

Agenda

- Overview of Portable Storage Technologies
- High Performance Flash Memory Controller
- Hybrid Hard Disk Drive
- Conclusions

Techniques for High Performance Flash Drive

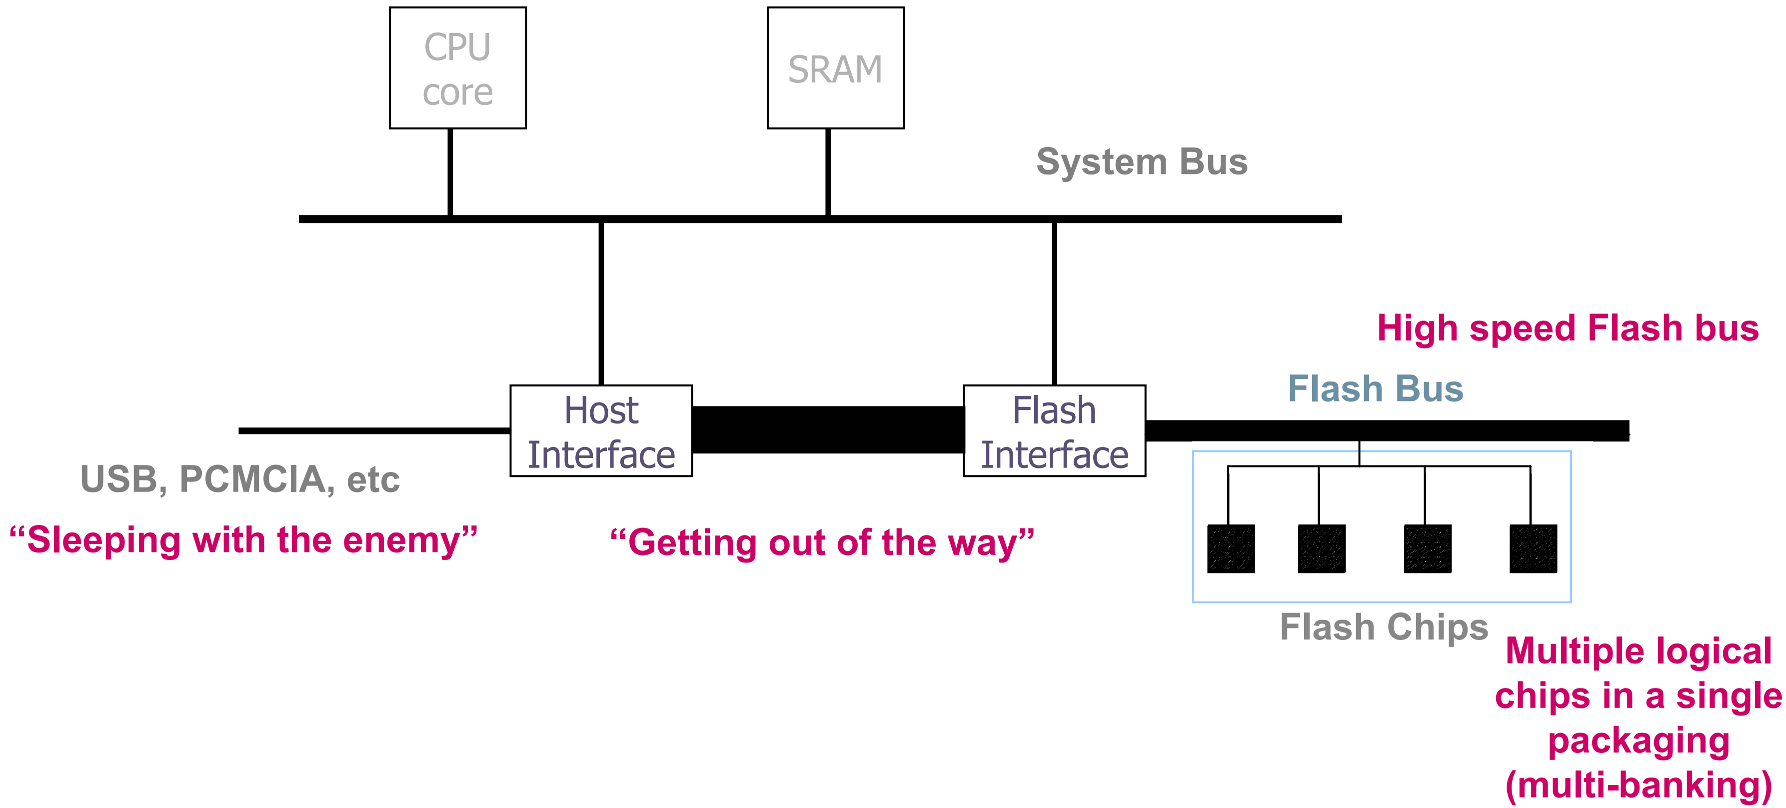

Evaluation board from FAST group (Version 1 – Home-made)

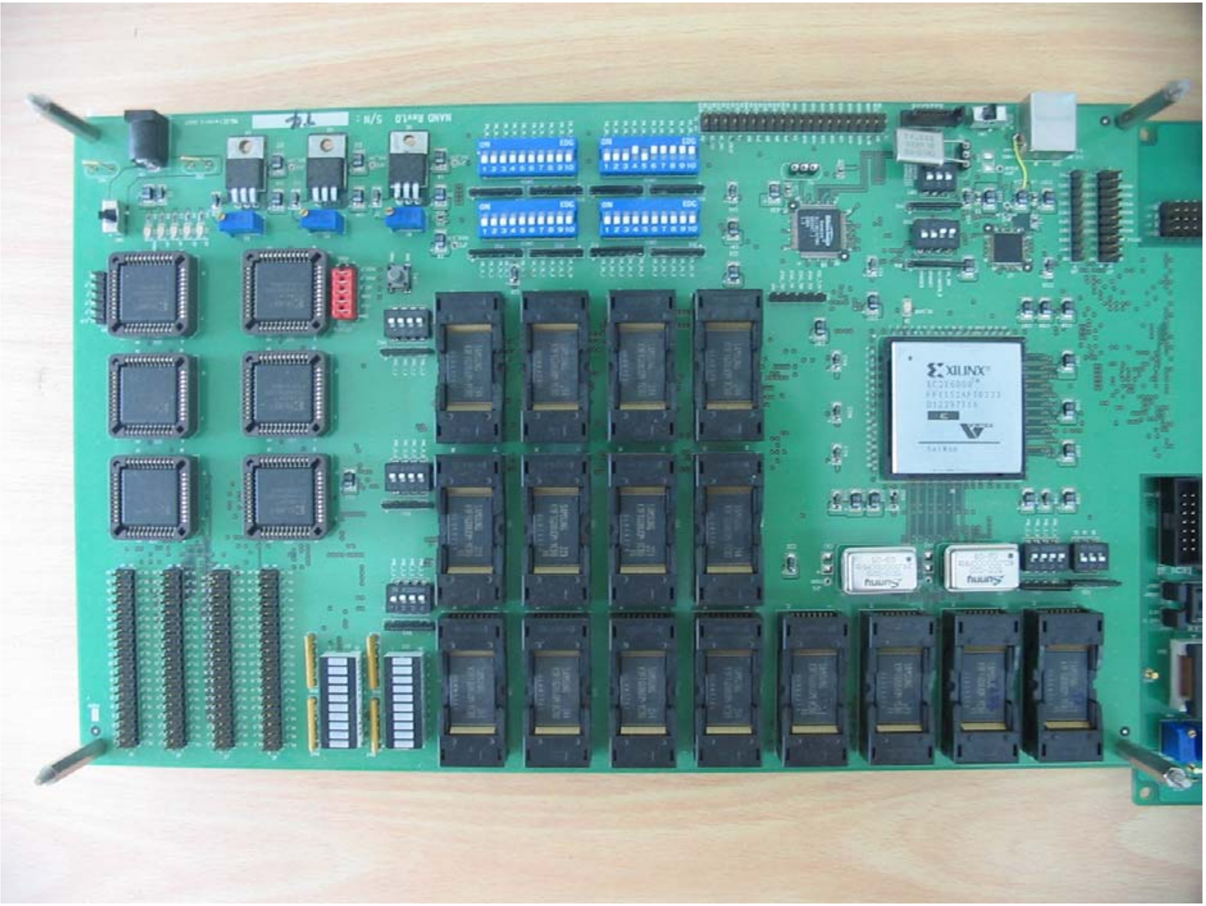

Evaluation board from FAST group (Version 2 with Samsung Electronics)

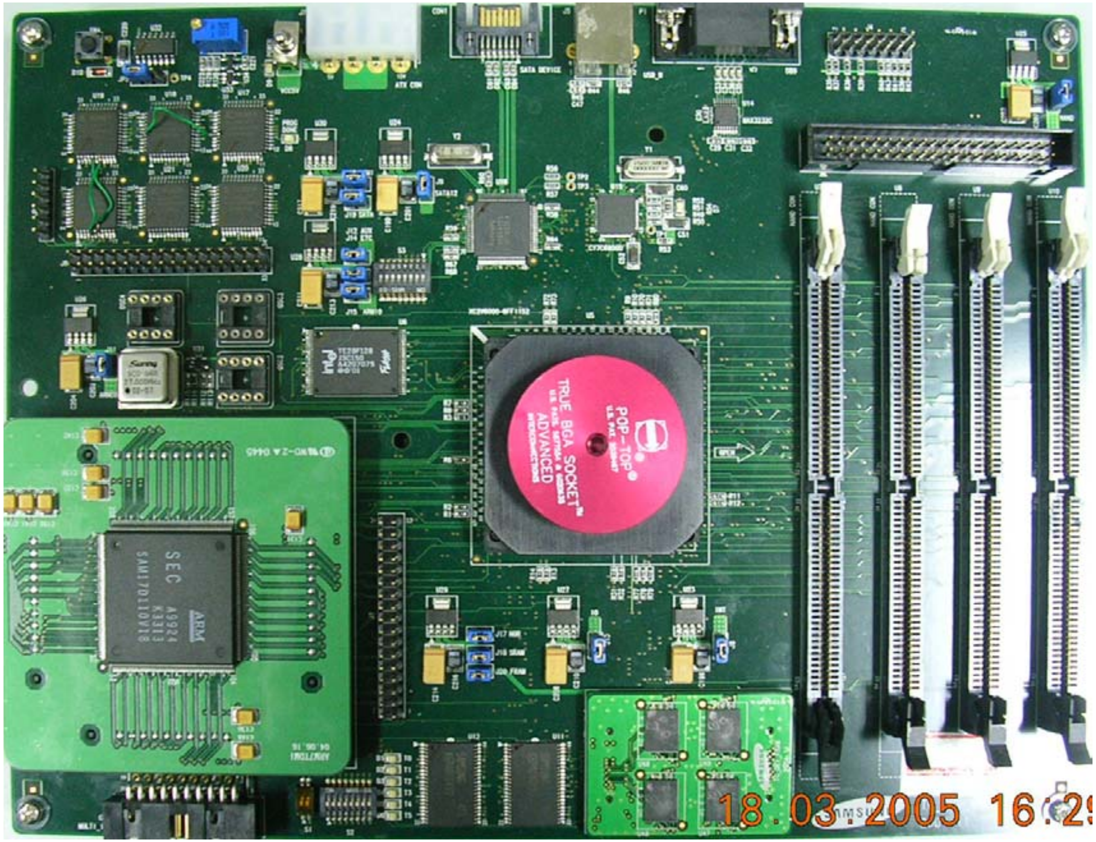

PCMARK'04 HDD Benchmarks

- PCMARK'04 HDD Benchmark: a Storage System Benchmark
  - Based on real usage
  - Consists of
    - Windows XP Startup
      - Contains disk activities occurring at operating system start-up
    - Application Loading
      - Contains disk activities from loading (opening and closing) MS Word, Acrobat Reader, Windows Media Player, etc
    - General Hard Disk Drive Usage
      - Contains disk activities while executing MS Word, Winzip, Winamp, Internet Explorer, Picture Viewers, etc
    - File Copying
      - Contains disk activities from copying 400MB of files

Performance Comparison (PCMark'04)

	HDD (2.5 in)	Chameleon (FRAM 2Mbytes)
Overall score	2499	8306
XP Startup	4.646	17.0
Application loading	3.670	15.4
File Copying	16.921	17.3
General HDD Usage	3.045	11.4

Why High Performance Flash Drive?

A Critical Point in VLSI Technology

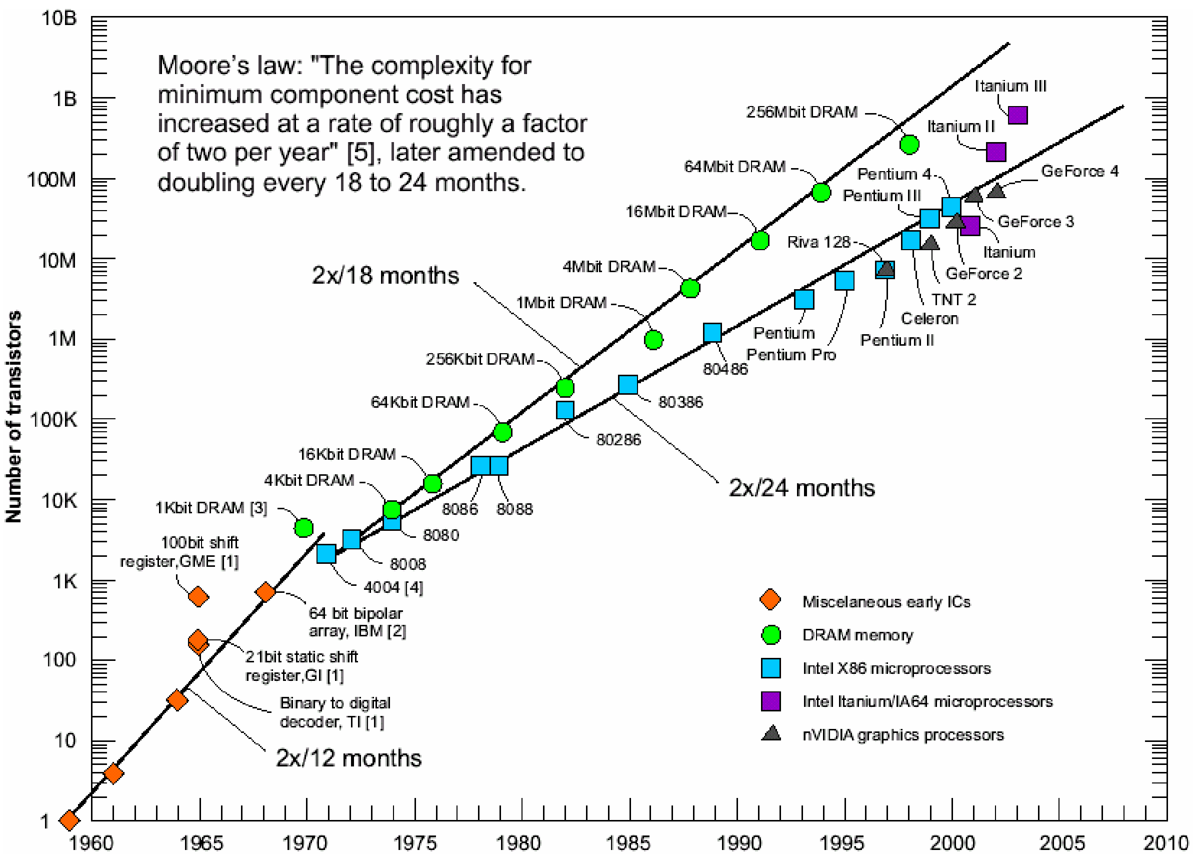

Source: www.icknowledge.com

A Critical Point in VLSI Technology and RISC Architecture

- Integration of processors on a single chip
  - The critical point ("*epoch*")
  - Argued for different architectures (RISC)
    - small repertoire of instructions in a uniform format
    - Pipelined execution
    - Cache memory
    - Load/store architecture
- The rest is history
  - Large/multi-level caches
  - Co-processors
  - Superscalar
  - Speculation
  - Simultaneous Multi-threading
  - etc

From "Single-Chip Multiprocessors: the Rebirth of Parallel Architecture" by Prof. Guri Sohi

A Critical Point in Portable Storage

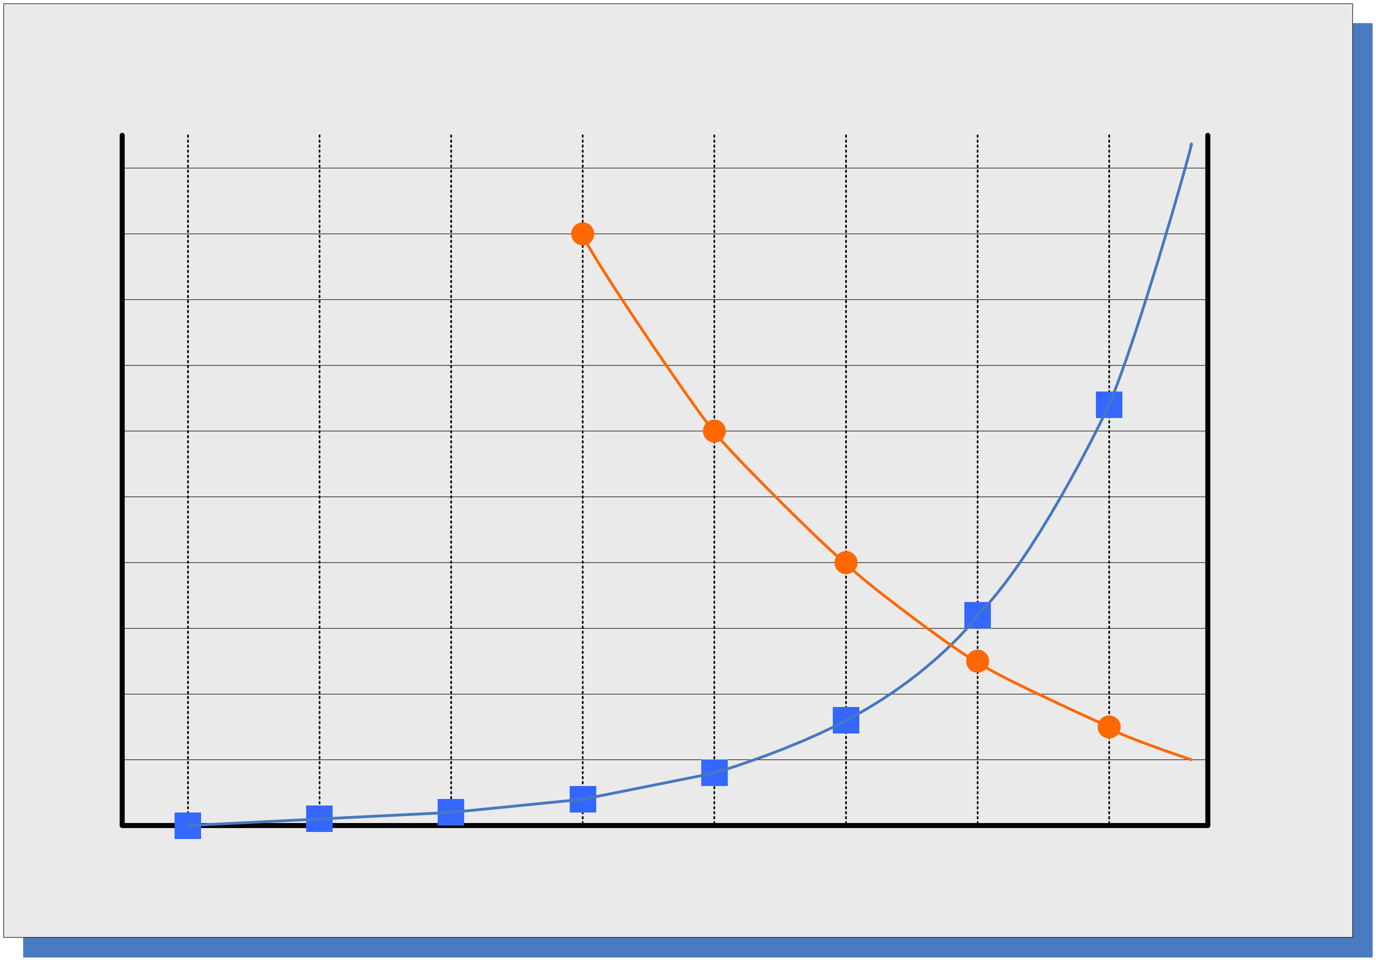

Density

Technical Impact of Large/High Performance Portable Storage: Stateless PC

What’s inside mobile storage?
1 Virtual Machine Monitors
2. Operating systems
3. File systems
4. Preference profile
5. Digital DNA

state

Stateless PC

Stateless Mobile Phone Analogy

state

What's inside sim card?
1 PIN number
2. Serial number
3. Encryption key
4. Network provider info
5. Phone book
6. Address book

Stateless GSM Mobile Phone

Agenda

- Overview of Portable Storage Technologies
- High Performance Flash Memory Controller
- Hybrid Hard Disk Drive
- Conclusions

Hybrid Hard Disk Drive

HDD

+

NAND Flash

=

HDD with
reduced
power consumption
and start-up time

Why Hybrid Hard Disk Drive?

1. Power consumption aspects:
   - In a laptop PC, HDD consumes
     - ~10% (~2W) total power when disk platters are spinning
     - ~1% (~0.2W) total power when disk platters are idle
2. Cost aspects:
   - 128MB Flash write buffer
     - < $8 in 2006
     - < $4 in 2008
3. Reliability aspects:
4. Performance aspects:

Source: Clark Nicholson, "Improved Disk Drive Power Consumption Using Solid State Non-Volatile Memory", WinHEC2004.

Hybrid Hard Disk Drive Block Diagram

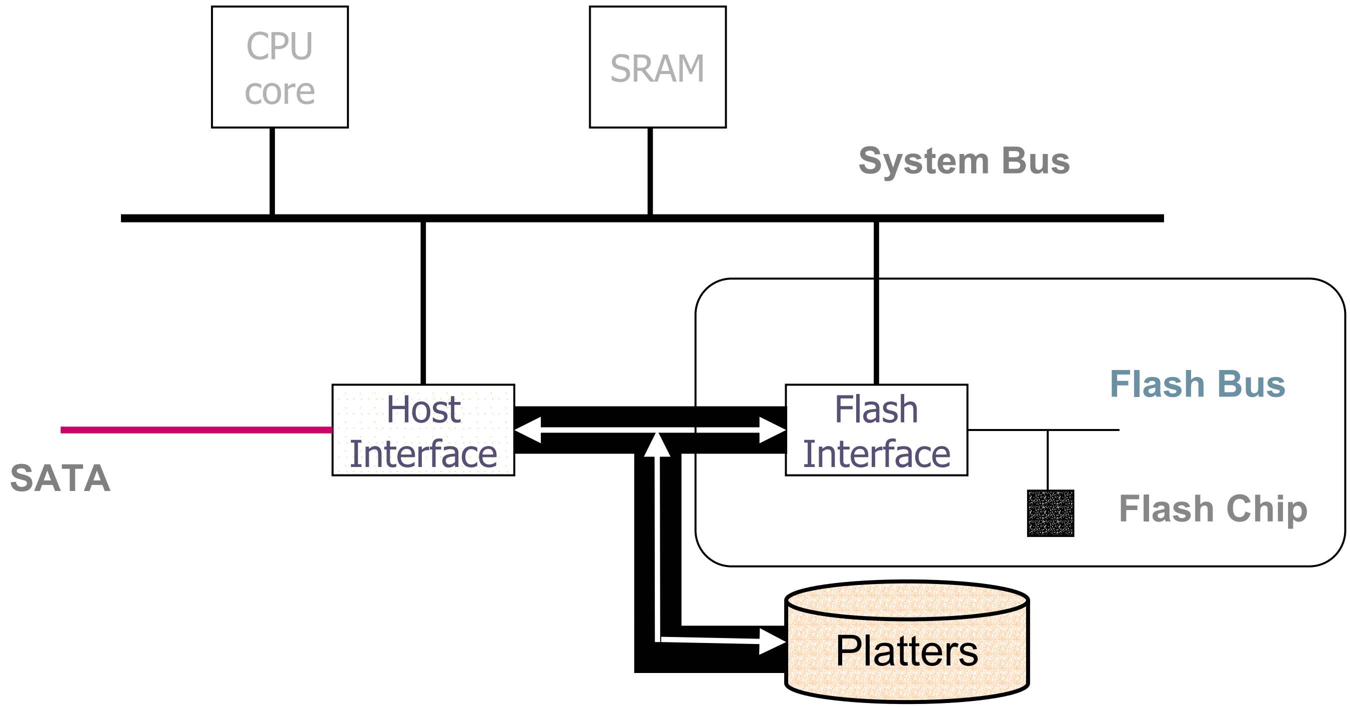

Key Benefits of Hybrid Hard Disk Drive

- 87% reduction in power can be achieved (1.75W)
- Assumptions
  - P_{avg} active = ~2W (measured)
  - P_{avg} with Flash write buffer and "Longhorn" kernel = 0.25W (calculated)
    - T_{off} = 600s @ .18W
    - T_{on} = 18s @ 2.5W
      - T_{on} = spin up time (5s) + Flash buffer flush time (13s)
        - Flash buffer size = 128MB

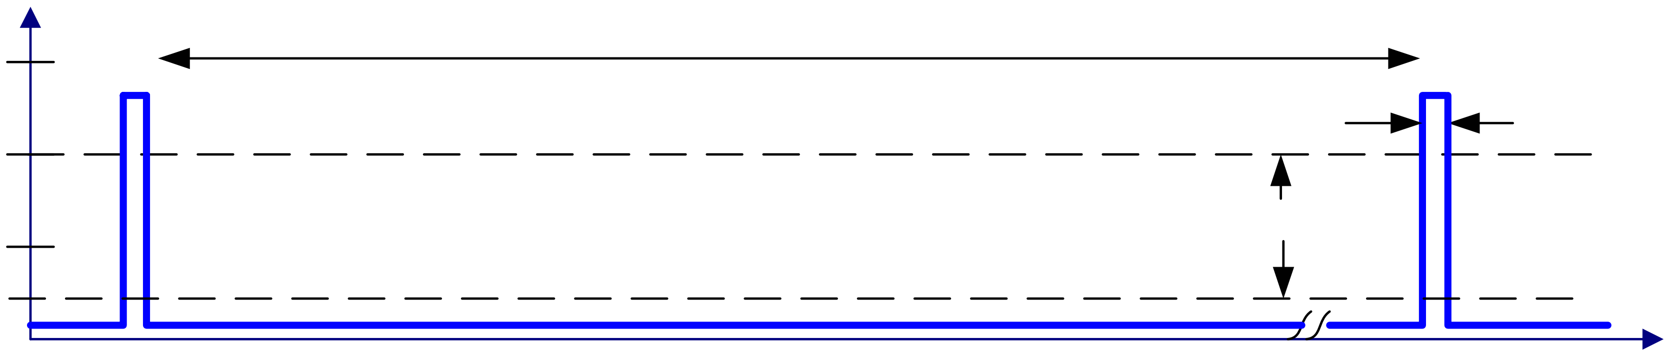

Source: Clark Nicholson, "Improved Disk Drive Power Consumption Using Solid State Non-Volatile Memory", WinHEC2004.

Key Considerations

1. Correctness: should preserve the semantics of HDD
2. Fault tolerance and graceful degradation: should operate correctly despite partial/total failure in flash memory
3. Power efficiency: should reduce the power consumption as much as possible
4. Reliability: should improve the reliability as much as possible
5. Performance: should improve the user-perceived performance as much as possible

Agenda

- Overview of Portable Storage Technologies
- High Performance Flash Memory Controller
- Hybrid Hard Disk Drive
- Conclusions

Conclusions

- In the animal world
  - Survival of the fittest

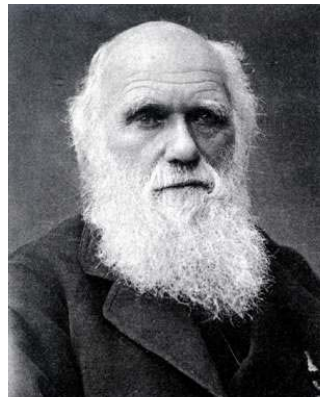

- In the memory world
  - Survival of the fastest or cheapest

	Volatile	Non-volatile
Fastest	SRAM	**FRAM, PRAM, or MRAM?**
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 – Micro Drive

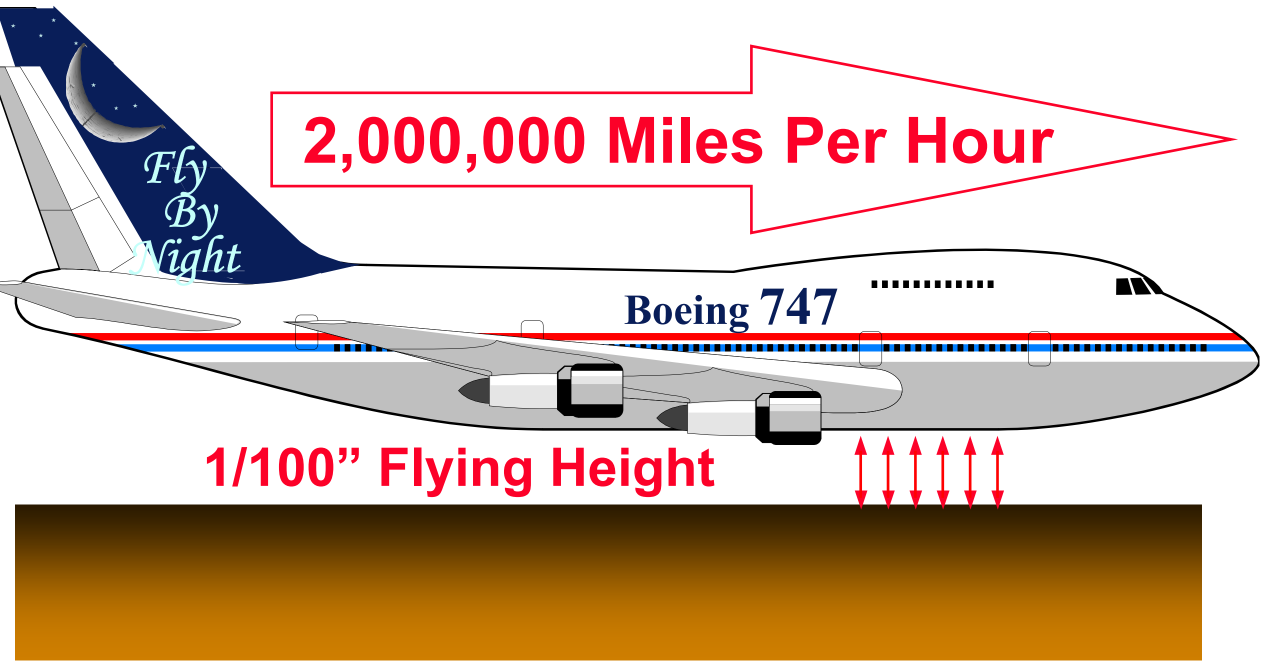

Source: Richard Lary, The New Storage Landscape: Forces shaping the storage economy, 2003.

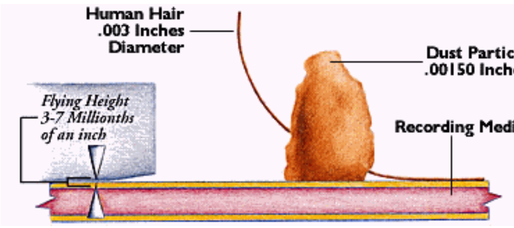

Source: http://www.hitachigst.cor

The Ultimate Limit – Flash Drive

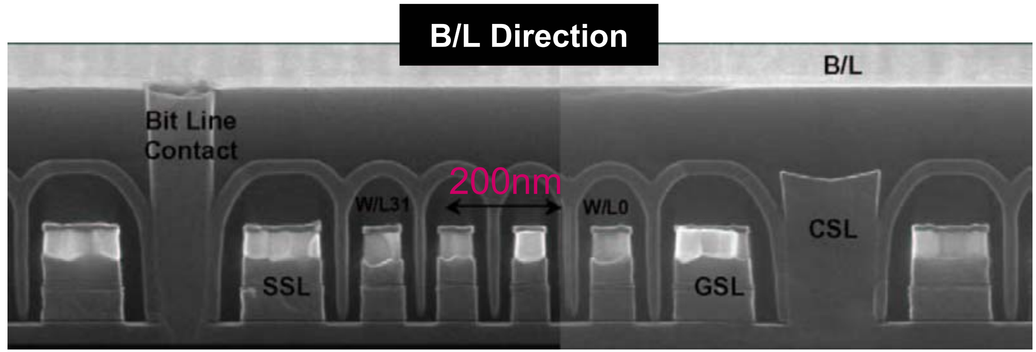

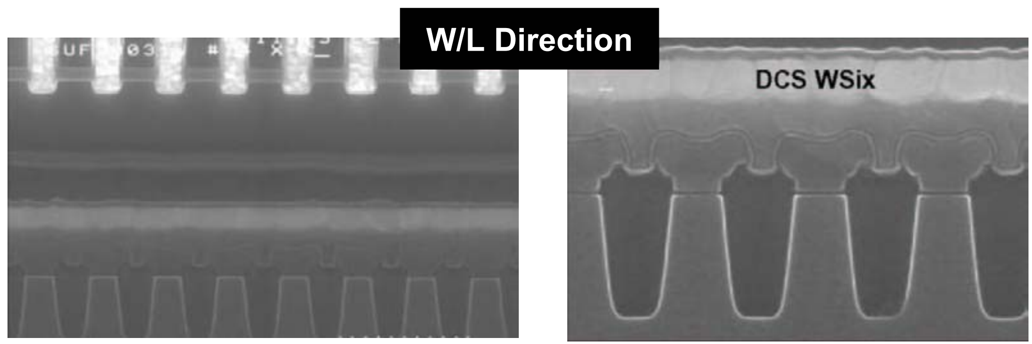

Source: K. Kim *et al.* IEDM Tech. Dig., 2002, pp. 919-922