

# SSD Application and Tech Trends

鍾慶龍 (Jeff Chung)

技術行銷副理

**AboUnion Technology Corp.**

March 19, 2009

# Within the Presentation...

- This presentation may make forward-looking statements during this presentation. Actual results may differ from those statements.
- This presentation may contain product information in development phase. The information here is subject to change without notice.
- The names, marks, photos, and brands are the property of their respective owners.

# Outline

- About AboUnion
- SSD Application
- SSD Tech Trends
- SSD Life time
- OS optimize for SSD

# About AboUnion

# Company Profile

- Founded March 2002
- Headquarter in Hsinchu, Taiwan
- Capital NT\$ 311M
- 70+ employees
- Chairman Mu-Lin Chen
- Certification ISO 9001-2000
- Design and manufacture flash memory storage product: USB flash drive, SD card, MMC card, Solid-State Drive, SSD module

# Milestone

2002 COB type MMC card solution

File patents



2003 Join MMCA

Window type MMC, file patents

SD card turnkey solution

2004 Join SDA



License CPRM

Association

2005 RS MMC, Mini SD, and Micro SD COB

ISO 9001 certification

2006 File patents


2008 SSD controller and module

File patents



友鑫科技股份有限公司  
AboUnion Technology Corp.


# SSD Product



2.5 inch PATA SSD



1.8 inch PATA SSD

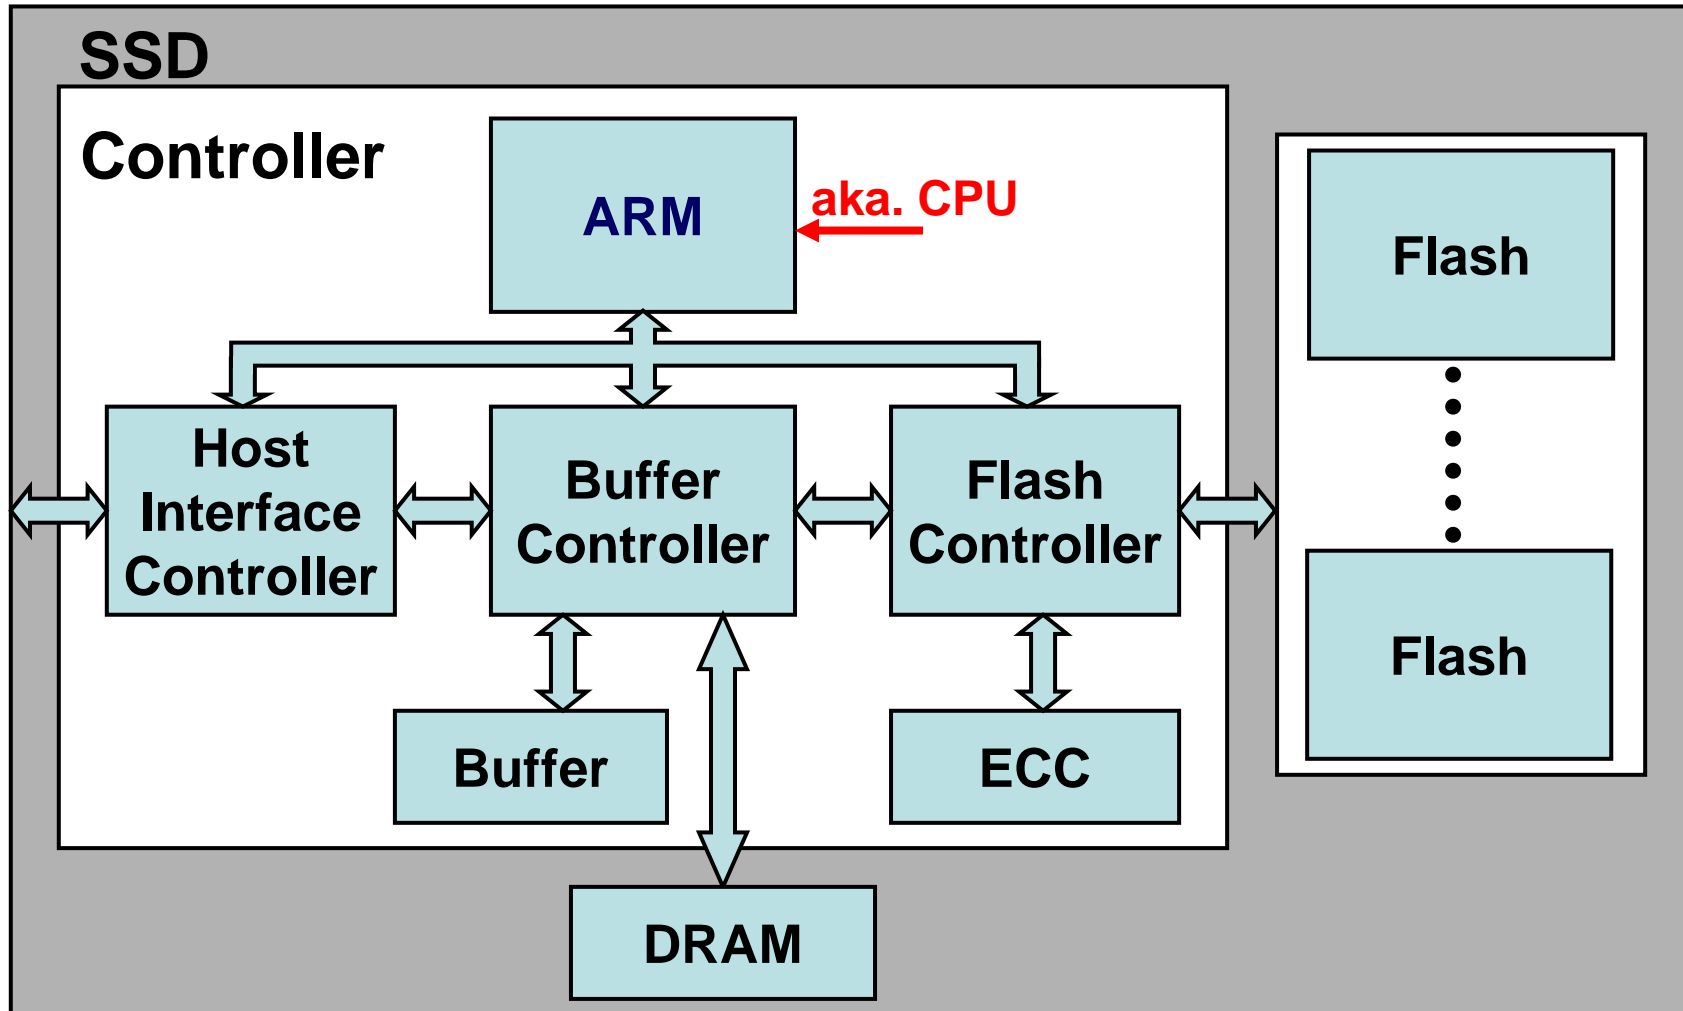


2.5 inch SATA SSD



友鑫科技股份有限公司  
AboUnion Technology Corp.

# AboUnion SSD Architecture



# SSD Application

# SSD V.S HDD



- Drop-in SATA or PATA HDD I/F
- No moving parts: NAND Silicon
- Multiple channels
- No form factor limitations
- No variation in performance
- Limitation in write cycles



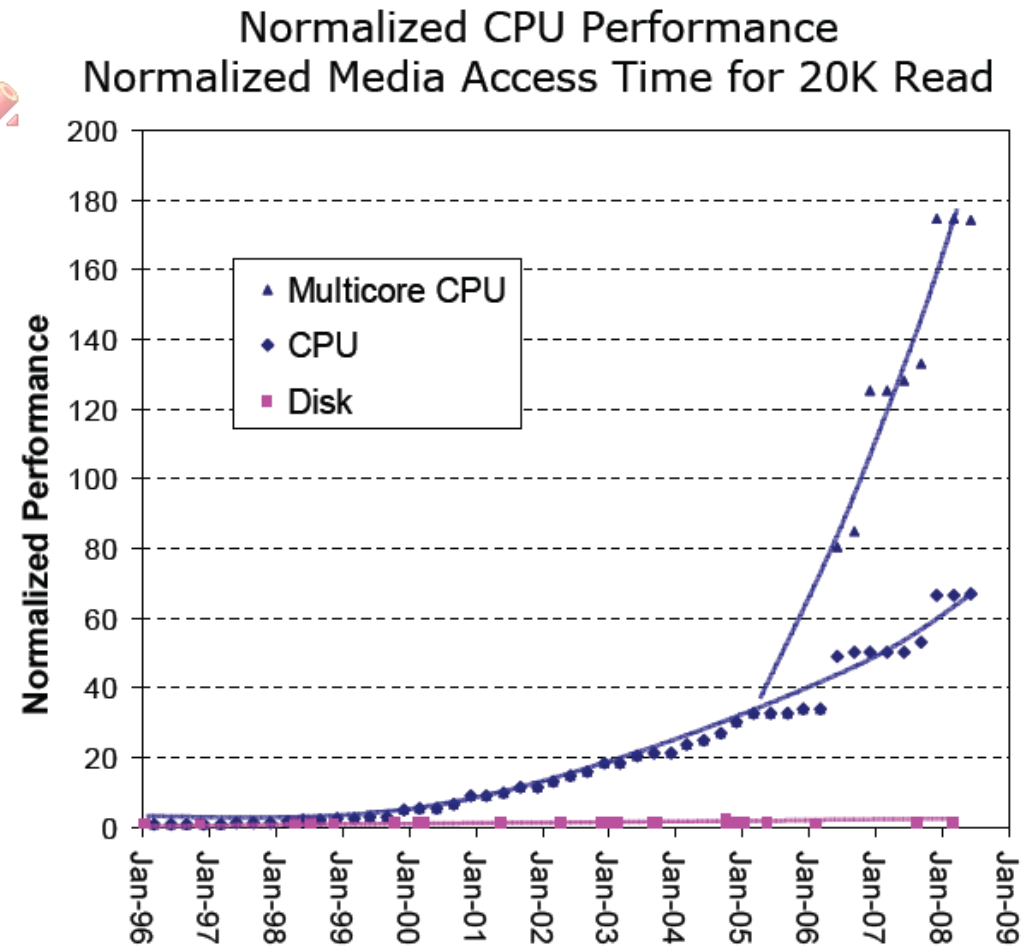
- Standard Storage
- Rotating magnetic media
- Single channel
- Mechanical FF limitations
- Performance spindle variation
- Random failures



# HOST I/O Bottleneck

- Mechanical latency dominates hard disk drive (HDD) media access for typical OS transfer sizes
- **1.3X** mechanical access improvement in 13 years
- **175X** CPU performance improvement in 13 years
- Solid-State Drives solve the storage performance gap

Source: Intel measurement



# SSD Market Info







Amazon.com: ssd - Amazon.com: Electronics - Windows Internet Explorer

http://www.amazon.com/s/qid=1236751156/ref=sr\_st?keywords=ssd&rs=172282&page=1&rh=i%3Aaps%2Ck%3Assd%2C1%3Aelectronics%2Cn%3A17

Amazon.com: ssd - Amazon.com: Electronics

Related Searches: [sd](#), [solid state drive](#), [ssd sata](#).

Showing 1 - 24 of 276 Results Sort by **Bestselling**

<p>Any Department</p> <p><b>Electronics</b></p> <ul style="list-style-type: none"> <li>Computers &amp; Add-Ons (214)</li> <li>Accessories &amp; Supplies (82)</li> <li>Camera &amp; Photo (28)</li> <li>Security &amp; Surveillance (1)</li> <li>Office Electronics (1)</li> </ul> <p><b>Shipping Option</b> <a href="#">(What's this?)</a></p> <p><b>Any Shipping Option</b></p> <ul style="list-style-type: none"> <li>Prime Eligible (69)</li> <li>Free Super Saver Shipping (69)</li> </ul> <p><b>Brand</b></p> <p><b>Any Brand</b></p> <ul style="list-style-type: none"> <li>OCZ (18)</li> <li>Transcend (18)</li> <li>Patriot Memory (3)</li> <li>Intel (8)</li> <li>Cavalry Storage (3)</li> <li>OCZ Technology (4)</li> <li>Lexar (2)</li> <li><a href="#">See more...</a></li> </ul> <p><b>Avg. Customer Review</b></p> <p><b>Any Avg. Customer Review</b></p> <ul style="list-style-type: none"> <li>★★★★☆ &amp; Up (20)</li> <li>★★★★☆ &amp; Up (30)</li> <li>★★★★☆ &amp; Up (37)</li> <li>★★★★☆ &amp; Up (38)</li> </ul> <p><b>Price</b></p> <p><b>Any Price</b></p> <ul style="list-style-type: none"> <li>\$0-\$25 (25)</li> <li>\$25-\$50 (22)</li> </ul>	<p>1. </p> <p><b>Thermaltake BlacX eSATA Hard Drive USB Docking Station</b></p> <p>Buy new: <del>\$68.26</del> <b>\$42.99</b></p> <p>16 Used &amp; new from \$37.99</p> <p>Get it by <b>Thursday, Mar 12</b> if you order in the next <b>15 hours</b> and choose one-day shipping.</p> <p>Eligible for <b>FREE</b> Super Saver Shipping.</p> <p>★★★★☆ (39)</p>	<p>2. </p> <p><b>Canon 104 Black Toner Cartridge</b></p> <p>Buy new: <del>\$73.99</del> <b>Click to see price</b></p> <p>103 Used &amp; new from \$27.71</p> <p>Get it by <b>Thursday, Mar 12</b> if you order in the next <b>16 hours</b> and choose one-day shipping.</p> <p>Eligible for <b>FREE</b> Super Saver Shipping.</p> <p>★★★★☆ (30)</p>	<p>3. </p> <p><b>Intel SSDSA2MH080G1C5 X25-M 80GB MLC 2.5-Inch 9.5mm Solid State Drive</b></p> <p>Buy new: <del>\$399.99</del> <b>\$379.99</b></p> <p>21 Used &amp; new from \$379.99</p> <p>Get it by <b>Thursday, Mar 12</b> if you order in the next <b>16 hours</b> and choose one-day shipping.</p> <p>Eligible for <b>FREE</b> Super Saver Shipping.</p> <p>★★★★☆ (14)</p>
	<p>4. </p> <p><b>OCZ OCZSSD2-1SLD60G Solid Series SATA2 2.5-Inch Solid State Drive</b></p> <p>Buy new: <del>\$100.99</del> <b>\$154.00</b></p> <p>5 Used &amp; new from \$125.00</p>	<p>5. </p> <p><b>OCZ OCZSSD2-1VTX120G 120GB SATA 2 Vertex Series Solid State Drive (Black)</b></p> <p>Buy new: <del>\$441.99</del> <b>\$431.99</b></p> <p>2 Used &amp; new from \$431.99</p>	<p>6. </p> <p><b>Western Digital VelociRaptor 300 GB Bulk/OEM Hard Drive 2.5 Inch, 16 MB Cache, 10,000 RPM SATA II WD3000GLFS</b></p> <p>Buy new: <del>\$349.99</del> <b>\$319.99</b></p>

Source: [www.amazon.com](http://www.amazon.com) 2009/3/12




友鑫科技股份有限公司  
AboUnion Technology Corp.

# Foreign SSD Spec

Vendor	M-brand		I-brand		S-brand	
Debut	Jul'08	Jul'08	Sep'08	Dec'08	Oct'08	Oct'08
Model	3500 2.5"	7500 2.5"	X-18/25M	X-25E	Half Slim 2.5"/1.8"	2.5"
Application	PC/NB	Enterprise	NB/PC	Enterprise	Netbook/ NB/PC	NB/PC/ Enterprise
Interface	SATA-II	SATA-II	SATA-II	SATA-II	SATA-II	SATA-II
NAND	SLC	SLC	MLC(50nm)	SLC	MLC	SLC
ECC Support	BCH-7	BCH-7	>BCH-8	na	na	na
R/W Performace(Seq R/W MB/s, IOPS/4KB)	SR:100 SW:100 RR:10,000 RW:60	SR:130 SW:120 RR:12,000 RW:130	SR:250 SW:70	SR:250 SW:170 RR:35,000 RW:3,300	SR:90 SE:70	SR:100 SW:80 RW:1,743
DRAM Cache	Yes	Yes	Yes(16MB)	Yes	Yes	Yes
Power(W)	Seq R:1.41 Seq W: 3.03 Idle:1.13	Seq R:2.06 Seq W:3.5 Idle:1.44	Active:0.15 Idle:0.06	Active:2.4 Idle:0.06	Active:0.43	Seq R:0.48 Seq W: 0.46 Idle:0.2
Pricing(US\$)	399(32GB)	779(32GB)	595(80GB)	695(32GB)	399(64GB)	729(64GB)
Other	ARM 7 SDRAM 4-ch separate Biit:25sec	ARM 7 SDRAM 4-ch separate	Xeon 10-ch	Xeon 10-ch	Boot:36sec	na

# Portable PC first overtake industrial application



Source : Web-Foot Research , 2007/10 , DIGITIMES , 2008/8



友鑫科技股份有限公司  
AboUnion Technology Corp.

# Industrial Application

**For OEM**

**(With Setup Utility)**

Full R/W  
authority



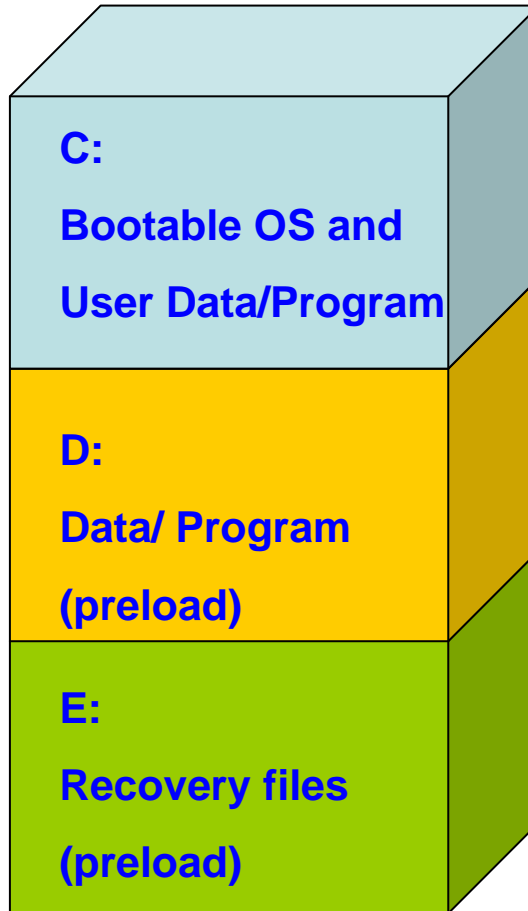
Full R/W  
authority



Full R/W  
authority



**SSD 16G (3 partition)**




**For End User**


Full R/W  
authority



Read only

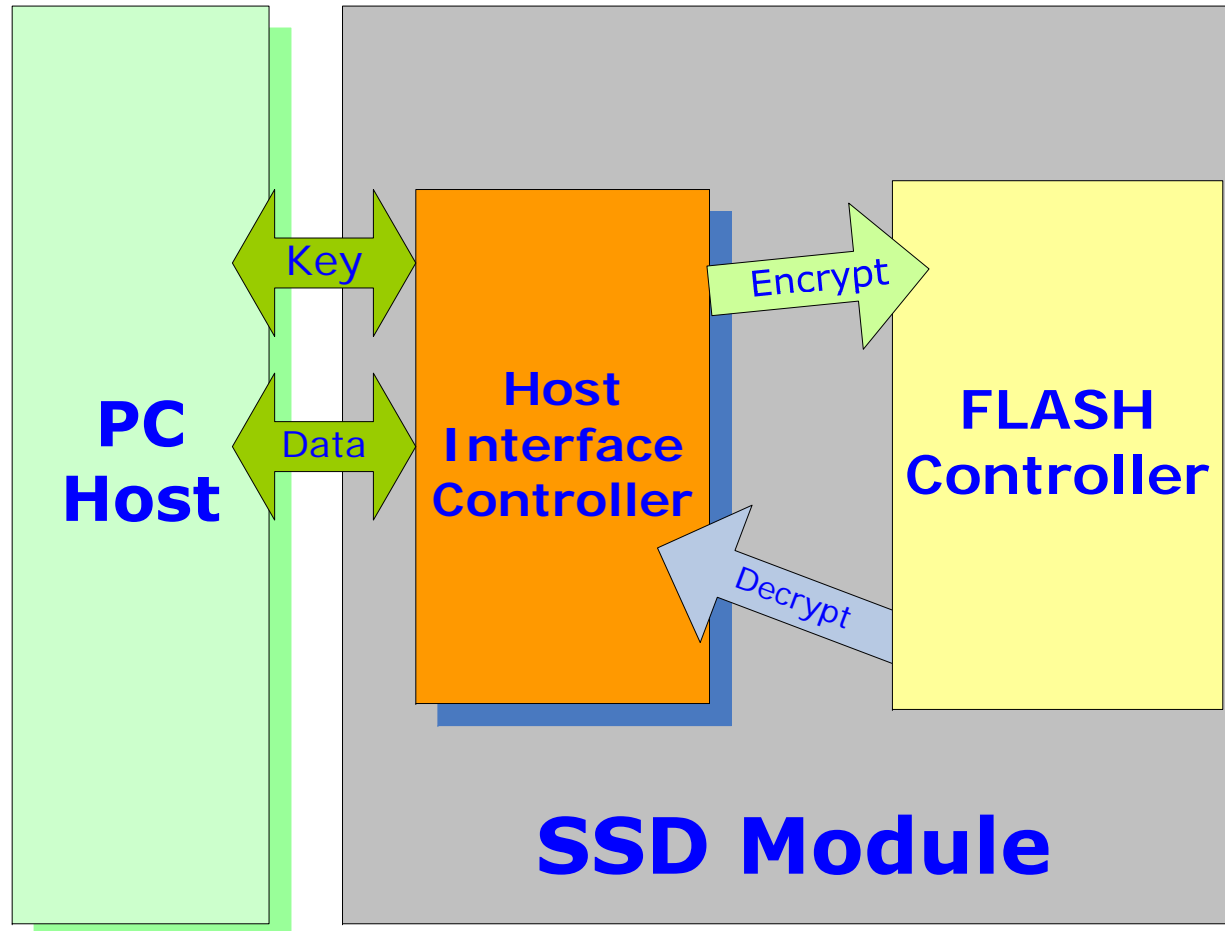


Invisible partition



友鑫科技股份有限公司  
AboUnion Technology Corp.

# Military Application



# SSD Tech Trends

# SSD Composition

## **Controller**

- CPU:8bit=>32bit
- DRAM Cache: 0/16/32MB
- ECC:6/8/12/24...
- Flash channel#: 2/4/8/10...

## **Host Interface**

- PATA
- SATA
- USB2.0
- USB3.0
- PCIe



## **Firmware**

- Wear leveling
- Bad block management
- Garbage collection
- Spare block

## **Flash**

- SLC/MLC/X3MLC
- Endurance:100k,10K,5K, 1K
- Process: 5x,4X,3Xnm
- DDR-Flash

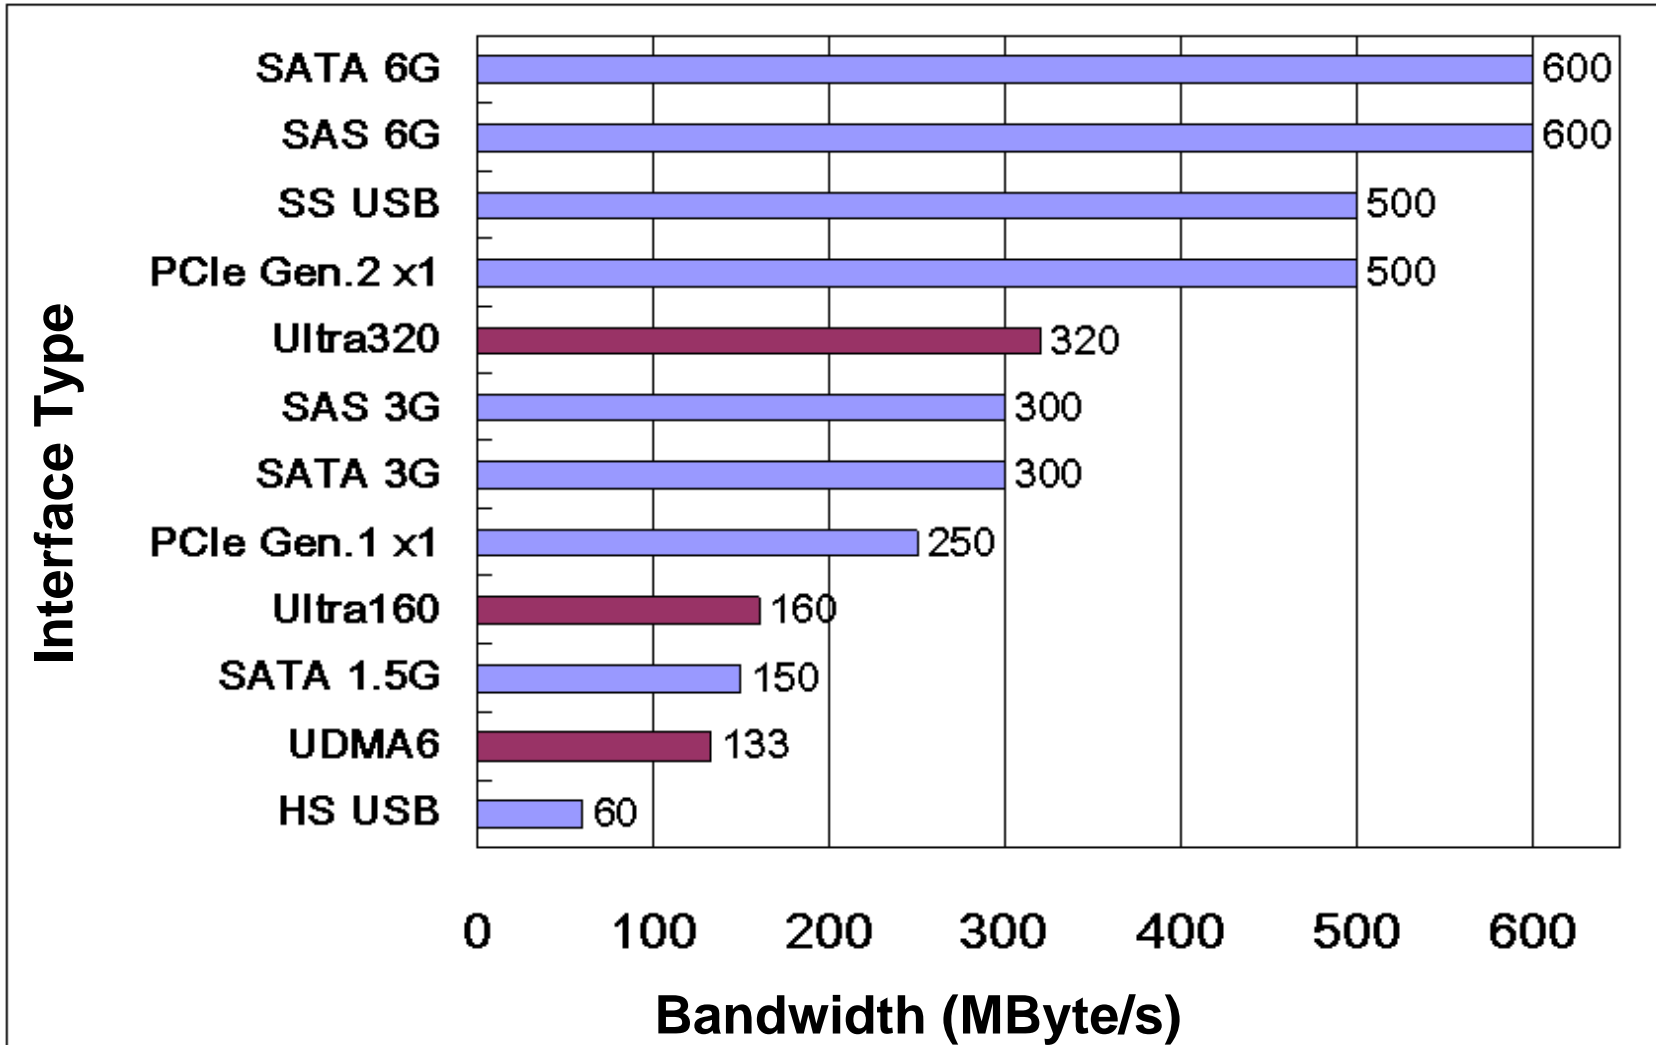


友鑫科技股份有限公司  
AboUnion Technology Corp.

# Performance Target

- **Individual Item**
  - Sequential access
  - Random access,
  - IOPS
  - Access time
- **Synthetic Benchmark**
  - Look forward to characterizing SSD effectively

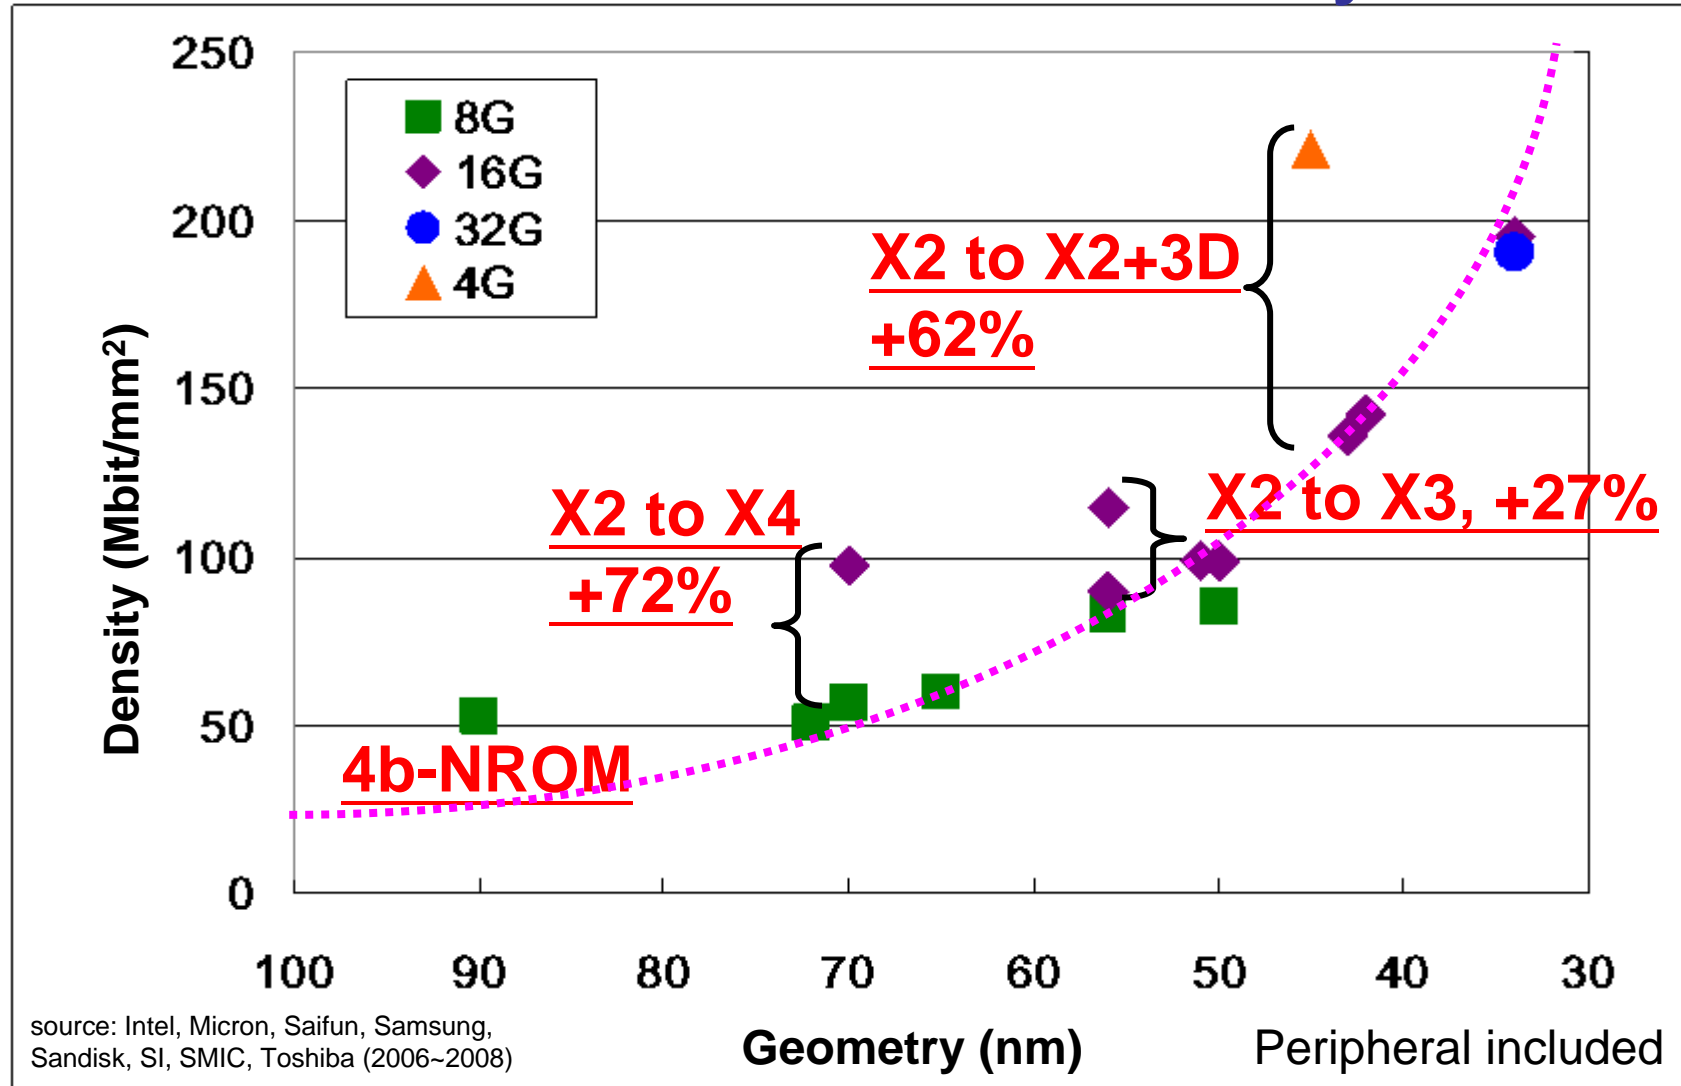
# Host Interface Trend



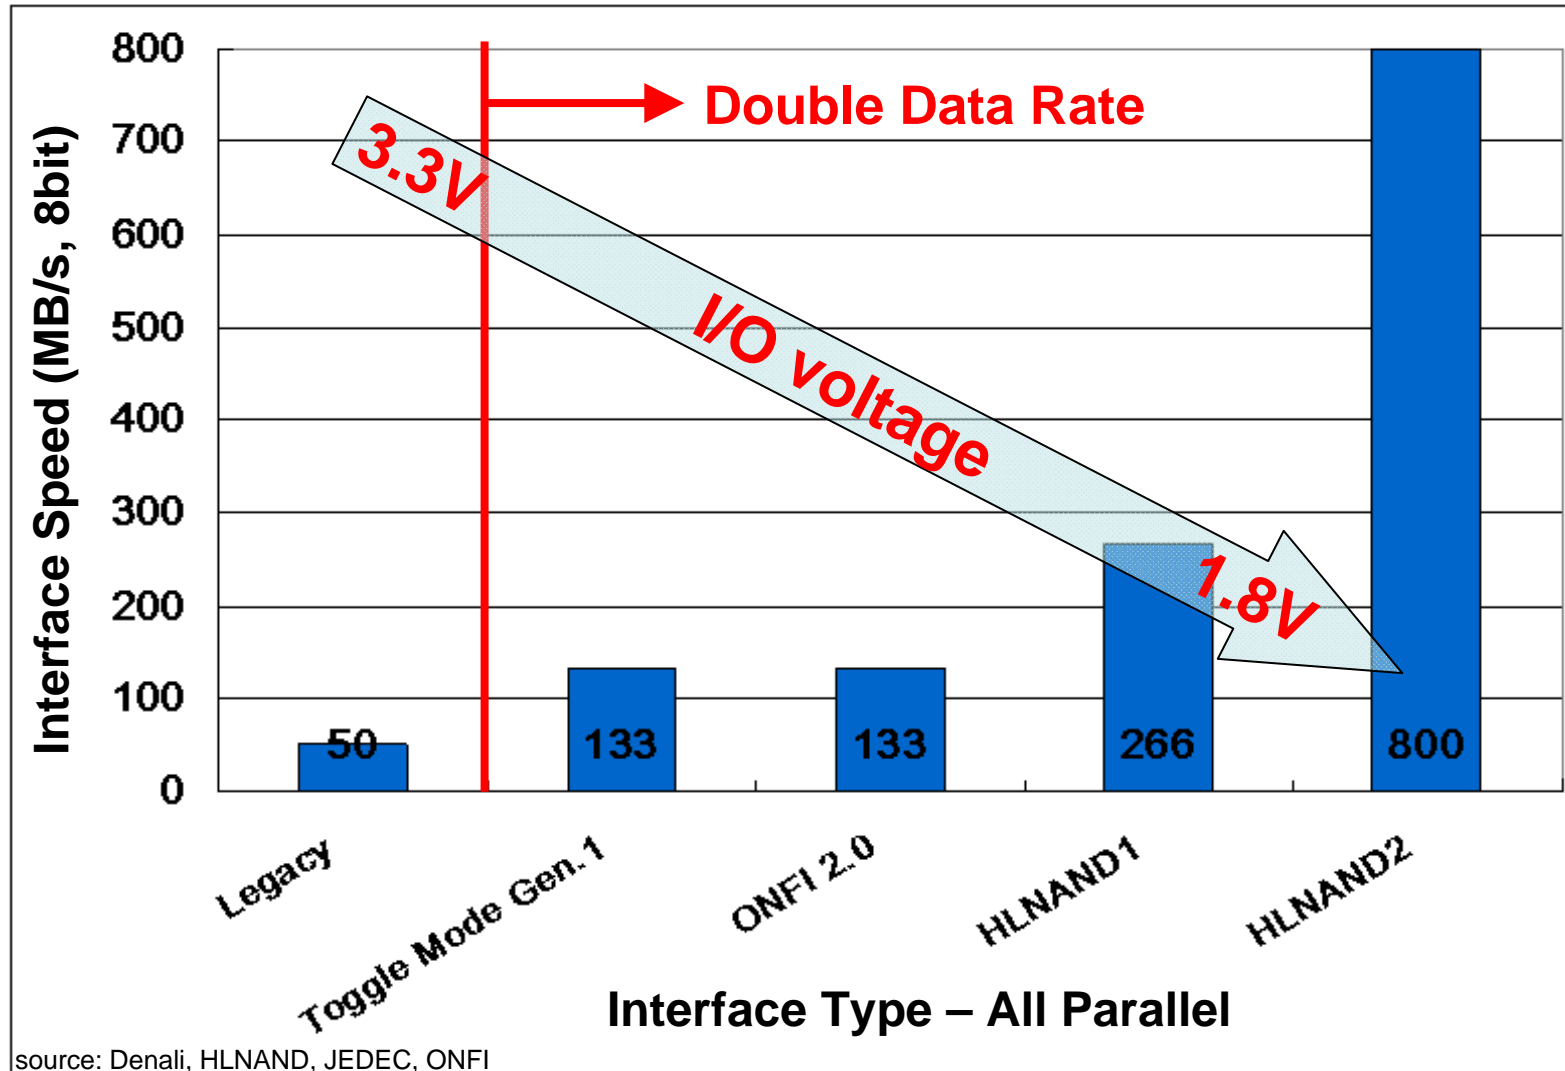
# NAND Flash Trends

- **Price Down**
  - Keep shrinking
  - More bits per cell
- **Capacity Up**
  - 32G bit mono die is coming
- **Throughput Up**
  - High speed interface (ONFI 2.0)
  - More bits read/programmed at once
- **Reliability .....**
  - More prone to disturbance
  - Higher bit error rate
  - Reduced endurance specification
  - Shorter data retention

# NAND Flash Density



# NAND Flash Interface



# SSD Lifetime

# How Long Will My SSD Last?


- SSD program endurance defines lifetime
- Flash P/E cycle
  - MLC: 10,000 / 5,000 / 3,000 / next?
  - SLC: 1,000,000 / 500,000 / 100,000 / next?
- It's the controller's job to translate these numbers into SSD lifetime
  - A good controller values every single P/E cycle

# What is S.M.A.R.T?

- **S**elf **M**onitoring **A**nalysis & **R**eporting **T**echnology
- Part of T13 (ATA Specification)

## What does it mean for SSDs?

- Read attribute threshold – pass/fail
  - Erase count
  - Spare blocks remaining
- Enables a prompt for user/service action
  - “Back-up your data”
  - “Contact your service representative”



# Wear Leveling (WL)

- Even distribution of P/E cycle
  - A good controller can make the wear-out happens gracefully
- Many algorithms have been proposed
  - A powerful controller performs best, such as a fast and globally optimized garbage collection
- But for long lifetime, “evenly distributed” alone is not enough

# Wear-leveling Type:

- ***Dynamic wear-leveling***
- Spreads writes across all free blocks
- Write to least –used cycle count block first
  
- ***Static wear-leveling***
  - Spreads low cycle count “static” data to high cycle count blocks
  - Automatically swap low cycle count ststic data to high cycle count block after reaching threshold
  - **A better wear-leveling method**

# SSD Lifetime Estimation

- $P/E \text{ Cycles} = (\text{Host writes}) / (\text{Drive capacity})$
- $P/E \text{ Cycles} = (\text{Host writes}) * (\text{Write amplification factor}) * (\text{Wear leveling factor}) / (\text{Drive capacity})$

# Write Amplification (WA)

- Proposed by Intel, defined as (lower is better):

$$WA = \frac{\text{data write to flash}}{\text{data write from host}}$$

- It's easy to maintain low WA during sequential access
- Random access is very prone to high WA
  - A powerful controller keeps WA low, such as having sufficient resources to pack random access efficiently
- With DRAM cache, WA below 1.0 is possible

# SSD Lifetime Example

- LOW P/E cycle is not that evil, HIGH WA is
- For example, assume evenly distributed WL engaged:
  - An 128G byte SSD using MLC flash rated 1,000 P/E cycle and having constant WA of 1.5 can be written 85T byte data before wear-out.

$$128 * 1,000 / 1.5 = 85,333 \text{ (G byte)}$$

- This translates to 4.6 years lifetime if 50G byte written per day.

$$85,333 / 50 / 365 = 4.67 \text{ (year)}$$

# Wear Leveling Experiment 1/2

**測試軟體:**  
BurnInTest V5.3 Pro


**測試實物:**  
AUC\_SSD 64G

**實際測試時數:**  
380h40m59s


磁碟區	配置	類型	檔案系統	狀態	容量	可用
AUC_SSD_64G	磁碟分割	基本	NTPS	狀況良好	62.63 GB	62.5
DOS98SE (E:)	磁碟分割	基本	FAT32	狀況良好	8.03 GB	8.03
WINXPP_SP2 (C:)	磁碟分割	基本	NTPS	狀況良好 (系統)	81.97 GB	71.8
WIN2KP_SP4 (F:)	磁碟分割	基本	FAT32	狀況良好	31.09 GB	24.3

Test Name	Cycle	Operations	Critical	Serious	Warning	Info
Disk (D: AUC_SSD_	8105	11,598 Trillion	0	0	0	0

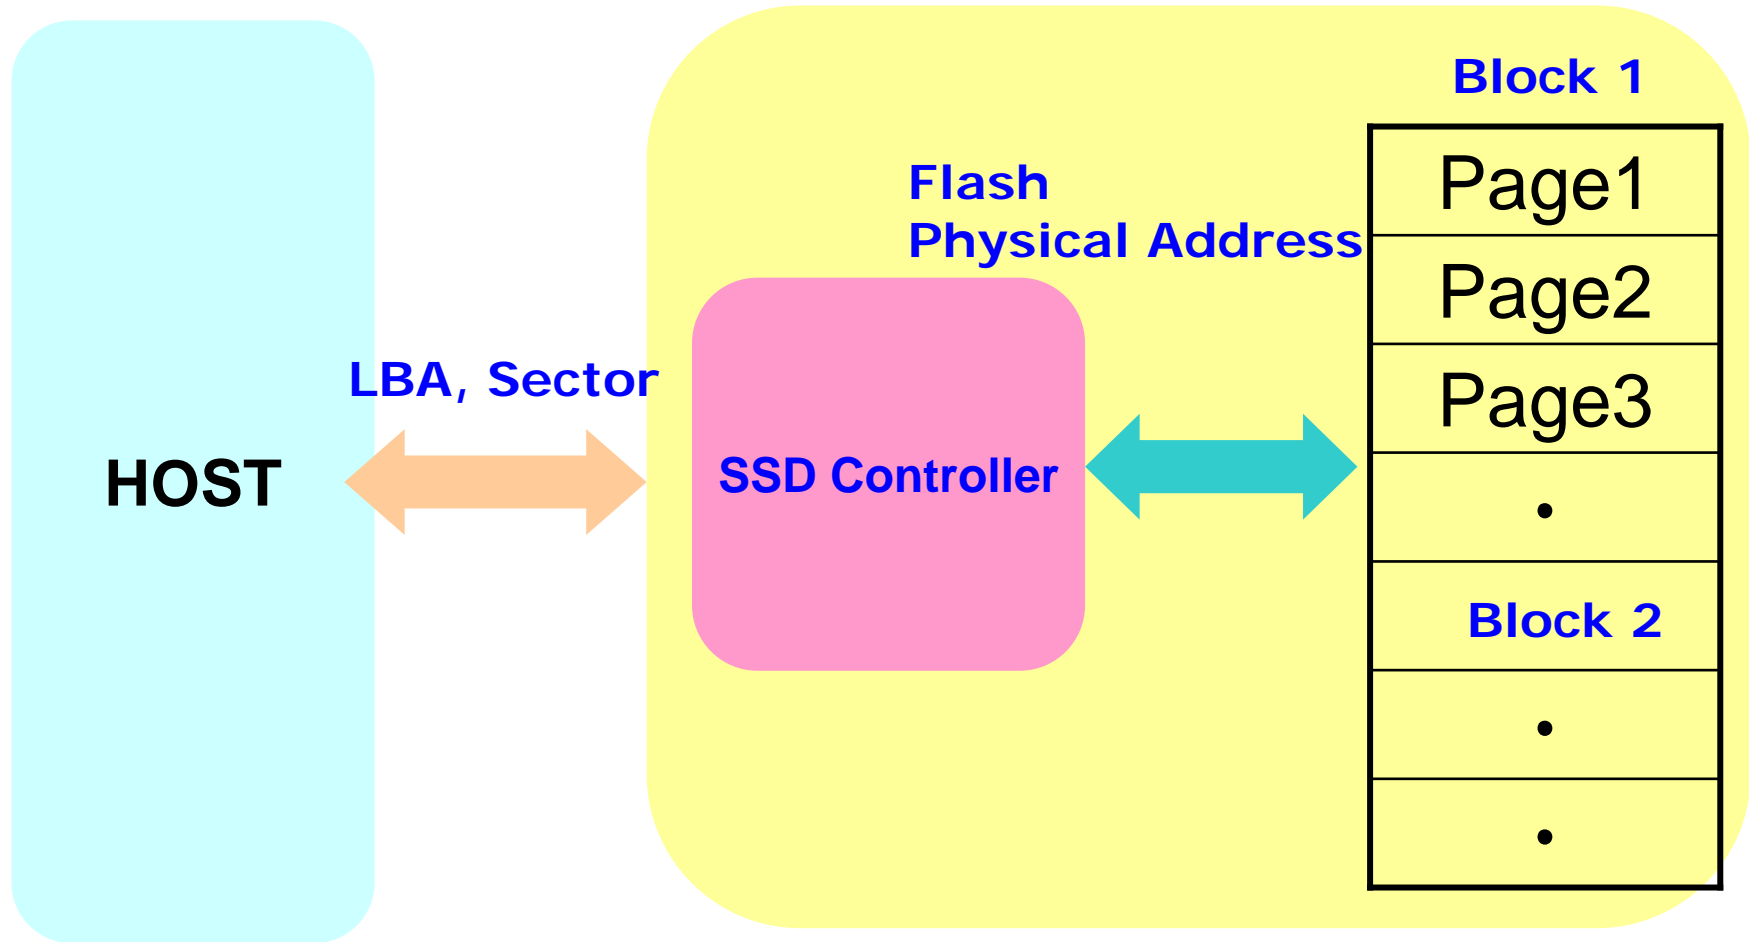
# Wear Leveling Experiment 2/2



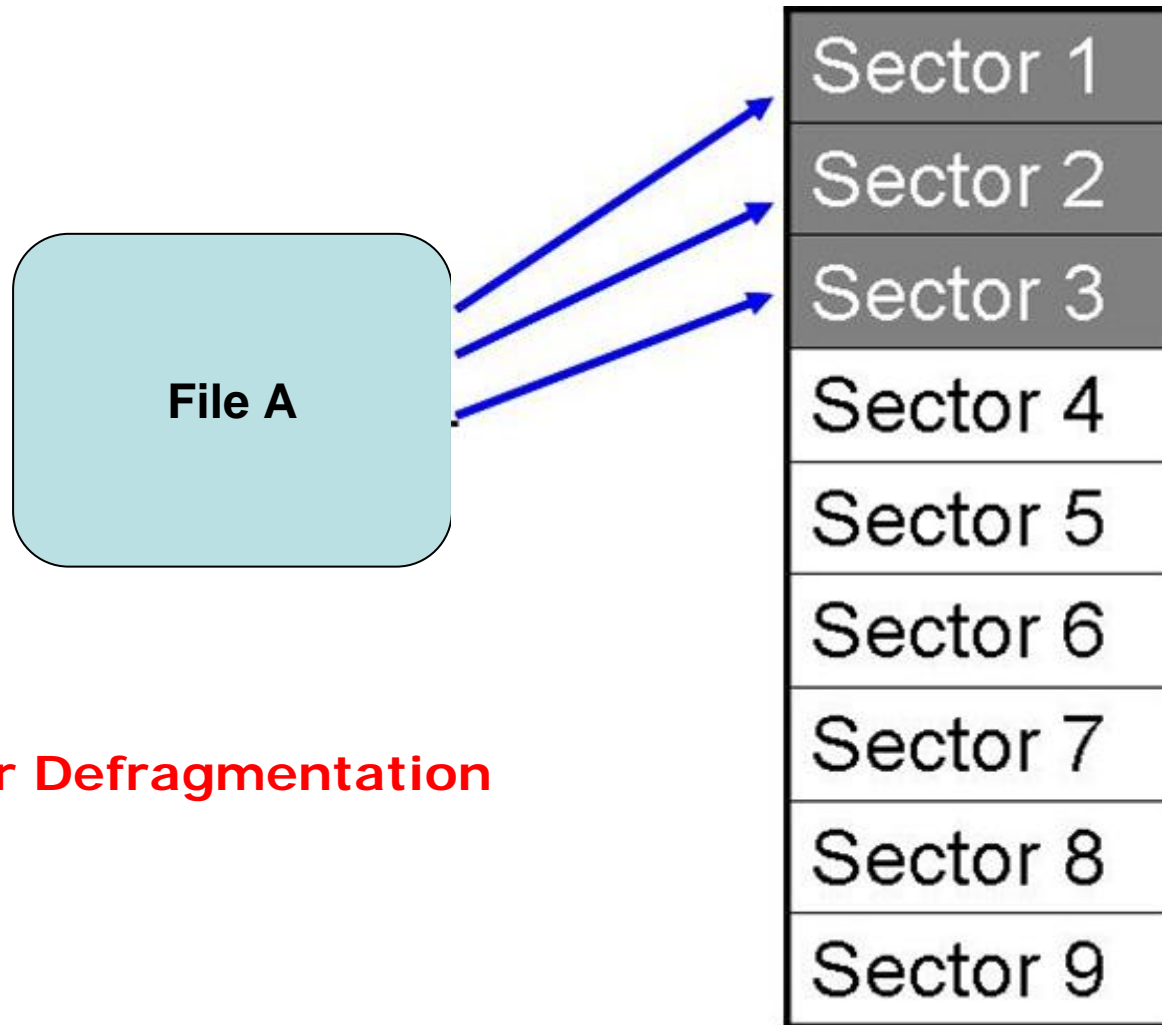
	CE0	CE1
MAX	3719	3308
MIN	5	3223
AVG	3667.37	3253.611
DIFFERENCE	3714	85



# SSD L2P Mapping



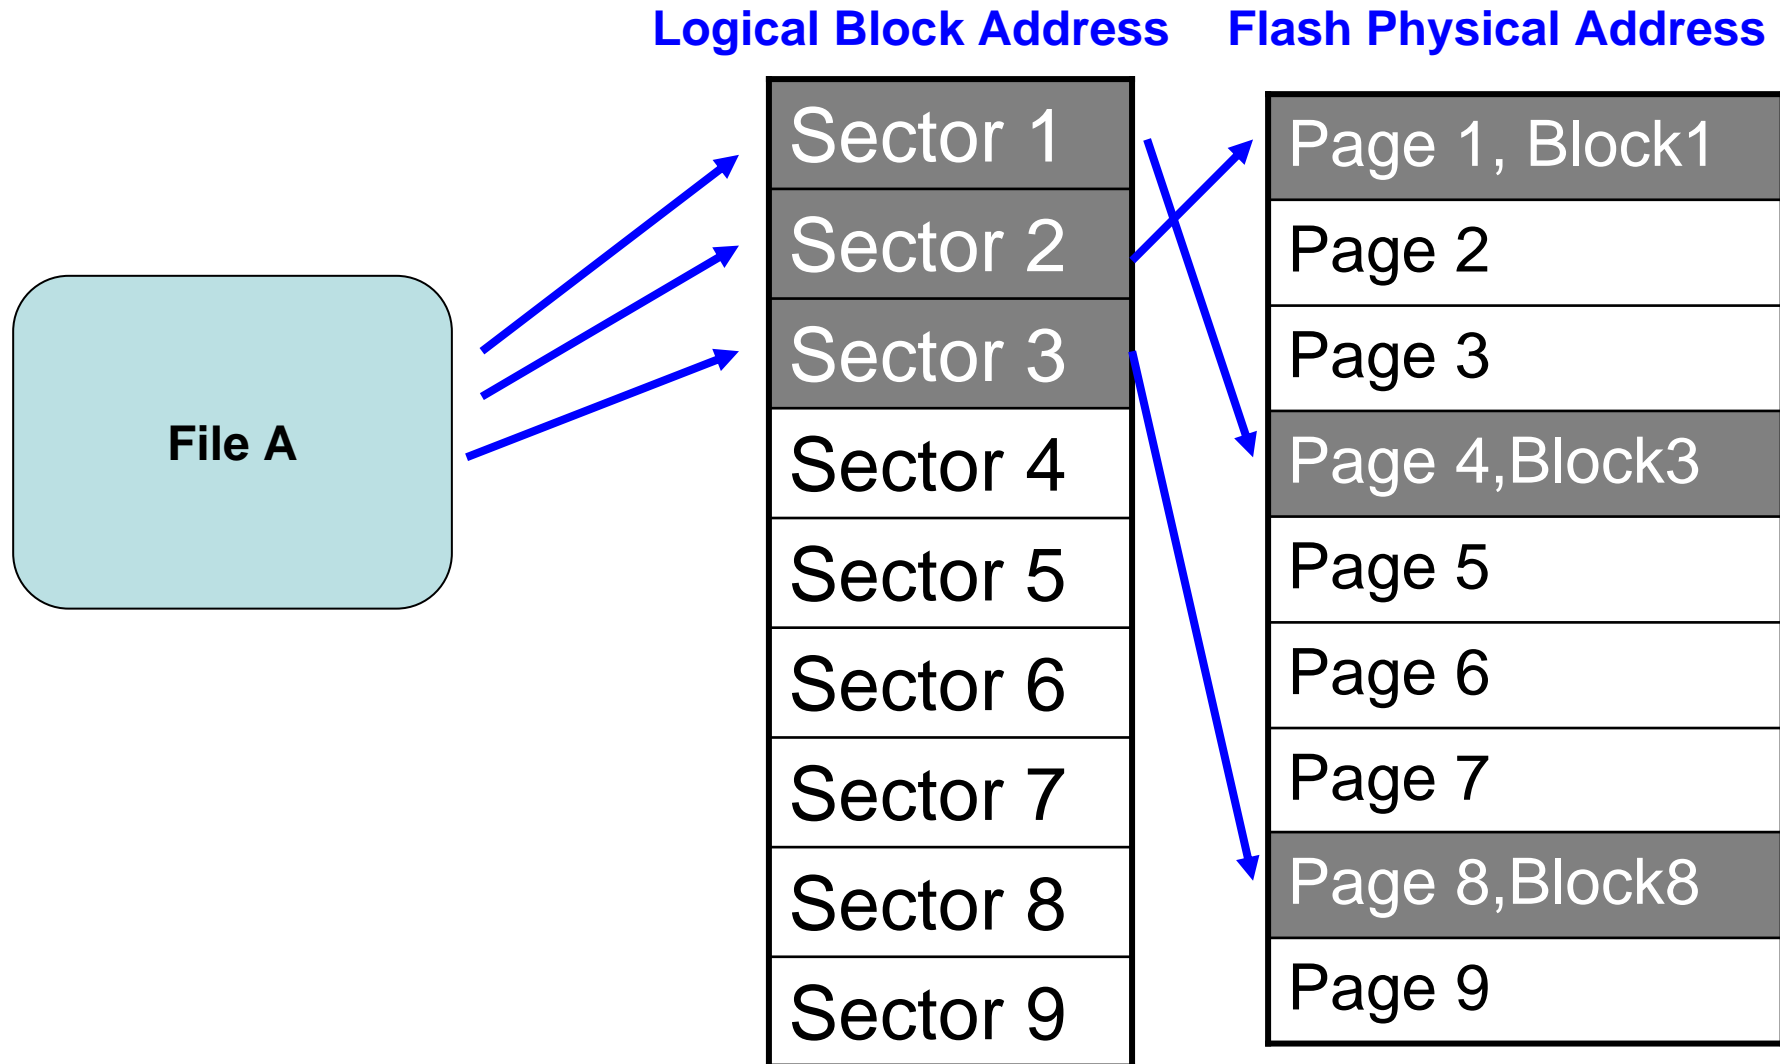
# HDD Defragmentation



**After Defragmentation**



# SSD Defragmentation



# OS optimize for SSD

- [Identifying SSD in Windows7](#)

Identification in system must be standard

New proposal for SSD identification is in T13

- [Turning off “Defragger” on SSD](#)

can save write cycles and add performance to SSD

- [New “trim” feature :](#)

Reducing overhead of SSD write operation in Windows with “Remove-On-Delete” that will cut down on the amount of data to be deleted (to further extend an SSD’s lifespan) and delete garbage data in advance (to speed up the writing of data, which is the activity where SSDs are slowest). Windows 7 will also use particular ATA commands which further improve SSD write speeds.

- [File system optimizations](#)

Minimizing random write can be a key for boosting SSD performance and reliability. SSD Page and SSD Block come to play, not sector or LBA

[Windows 7 logo evaluation kit \(available soon\) to ensure your device compliance](#)

# The Solid State Drive Identify Data

## 7.16.7.77 Word 217: Nominal media rotation rate

Word 217 indicates the nominal media rotation rate of the device and is defined in table 38.

Table 38 — Nominal Media Rotation Rate

Value	Description
0000h	Rate not reported
0001h	Non-rotating media (e.g., solid state device)
0002h-0400h	Reserved
0401h-FFFEh	Nominal media rotation rate in rotations per minute (rpm) (e.g., 7200 rpm = 1C20h)
FFFFh	Reserved

**Reporting non-rotating media will allow windows7 to set defrag off as default ; improving the device endurance by reducing write.**

**Source: T13,ATA8-ACS**



友鑫科技股份有限公司  
AboUnion Technology Corp.

# Summary of Device Statistic Information

1. Device Statistic Information Header
2. Defect Block Size
3. Number of Defect Blocks in the Solid State Media
4. Percentage of Defect Block in the Solid State Media
5. Number of Solid State Media Erase/Program Cycles
6. Percentage of the Solid State Media Remaining Endurance
7. Percentage of Spare Blocks Remaining in Solid State Media

**Source: T13/e06184r3, Samsung, IBM, Microsoft**

# To Realize New-Generation SSD

- To maximize performance
  - We need a powerful controller
- To maximize lifetime
  - We still need a powerful controller
- As long as WL/WA kept even/low respectively no matter whether sequential or not, MLC SSD finds its way, let alone SLC SSD.
  - A powerful controller makes this happen

# Very Heavy Loading on CPU!

- To Maximize Performance
- To Maximize Lifetime
- High Computation Power

**ARM based architecture is the one for new-generation SSD.**

- Address large DRAMs and many peripherals
- Address DRAM

- Low Power



# Not All Controllers Are Created Equal

Controller	AUC625	AUC635
Feature		

**A Powerful hardware is, actually, a platform.**

**AboUnion ARM based SSD  
provides all these for you!**

**The software does matter.**

Status	Sampling NOW	ES 2Q'09
--------	--------------	----------

# AboUnion Service

- **OEM/ODM**
- **Turnkey Solution**
- **No Brand Strategy**

**AboUnion provide you high performance and high reliability SSD solution.**

# Thank You!

[Jeffchung@abounion.com.tw](mailto:Jeffchung@abounion.com.tw)

