

NVRAMOS '11

November 8, 2011, Jeju

SSD Characterization: From Energy Consumption's Perspective

Youjip Won

Hanyang University



Outline

- Motivation
- Related Works
- SSD Organization and Energy Consumption
- Channels, Ways and Clusters
- Trivia in Measurement Methodology
- Case Study
- Power Budget
- Conclusion



Motivation


Understanding of Internal Mechanism of Storage Device is very important!

□ Hard disk


- ◆ Sector layout: cylinder serpentine vs. surface serpentine vs. hybrid serpentine
- ◆ Number of zones
- ◆ Degree of track skew
- ◆ Disk scheduling algorithm

Characterizing HDD: Sector Layout

- Jongmin Gim et al, ACM ToS 6, 2 (July 2010)



Hitachi HDT7725032VLA360 for 2000 tracks



Characterizing Storage Devices

HDD Characterization is via measuring **Seek time** and **Rotational Latency**.

Characterizing SSD... what do we use?...



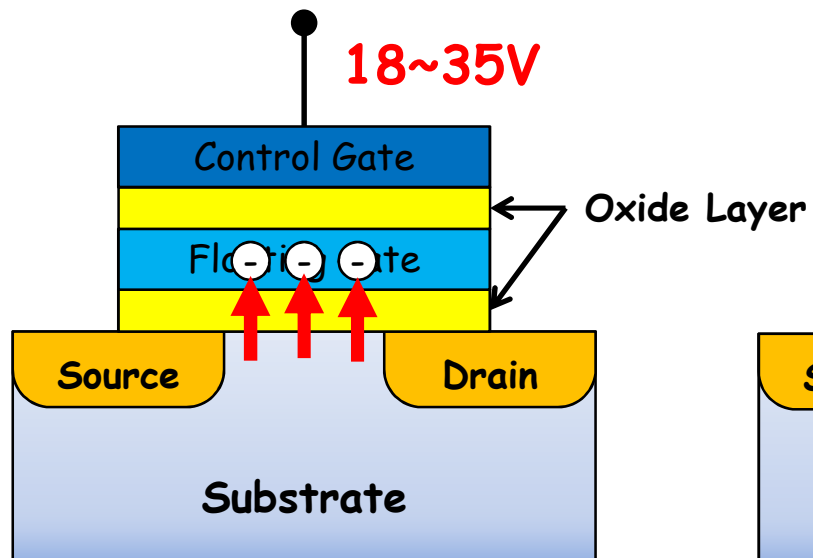
How do we figure out the internals of SSD?

- What is available
 - ◆ the number of channels
 - ◆ the number of chips/packages per channels

- What is not available?
 - ◆ Sector placement , Garbage collection algorithm


Operations on NAND Flash Cell

program



Programmed: 1 → 0

erase



Erased : 0 → 1

Operations on NAND Flash Cell



SSD Characterization

We will use "Energy Consumption"





Related Works


Related Works

- Dongkun Shin et al, NVRAMOS 2010 Spring

Identifying the relationship between workload characteristics and aggregate power consumption for each workload

Applied mixed workloads(random, sequential, etc) to SSDs with different request sizes and varied the file systems
Measured the Power Consumption(measured voltage change)

- Laura M. Grupp et al, MICRO 2009



Custom Board + Flash Memory

Read

Program

Erase

Basic Operation



Power Consumption of Flash Memory Basic Operations

Examine the average power consumption(W) and energy(J)


Related Works

- Euseong Seo et al, HotPower'08

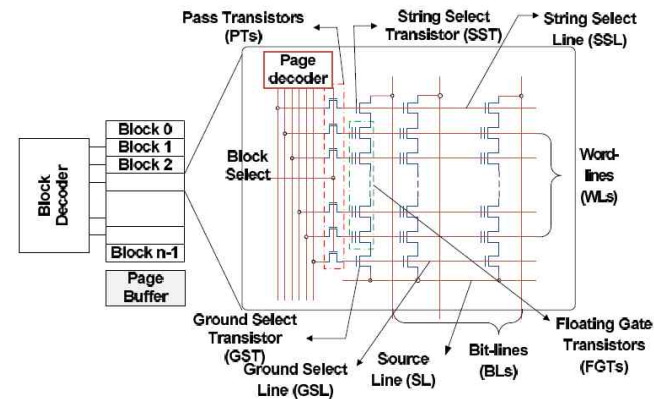
Analyzed the power consumption patterns of the SSDs

different hardware configurations to the various **atomic** operations and the combination of file systems and workloads.

- Vidyabhushan Mohan et al, Date '10



CACTI 5.3

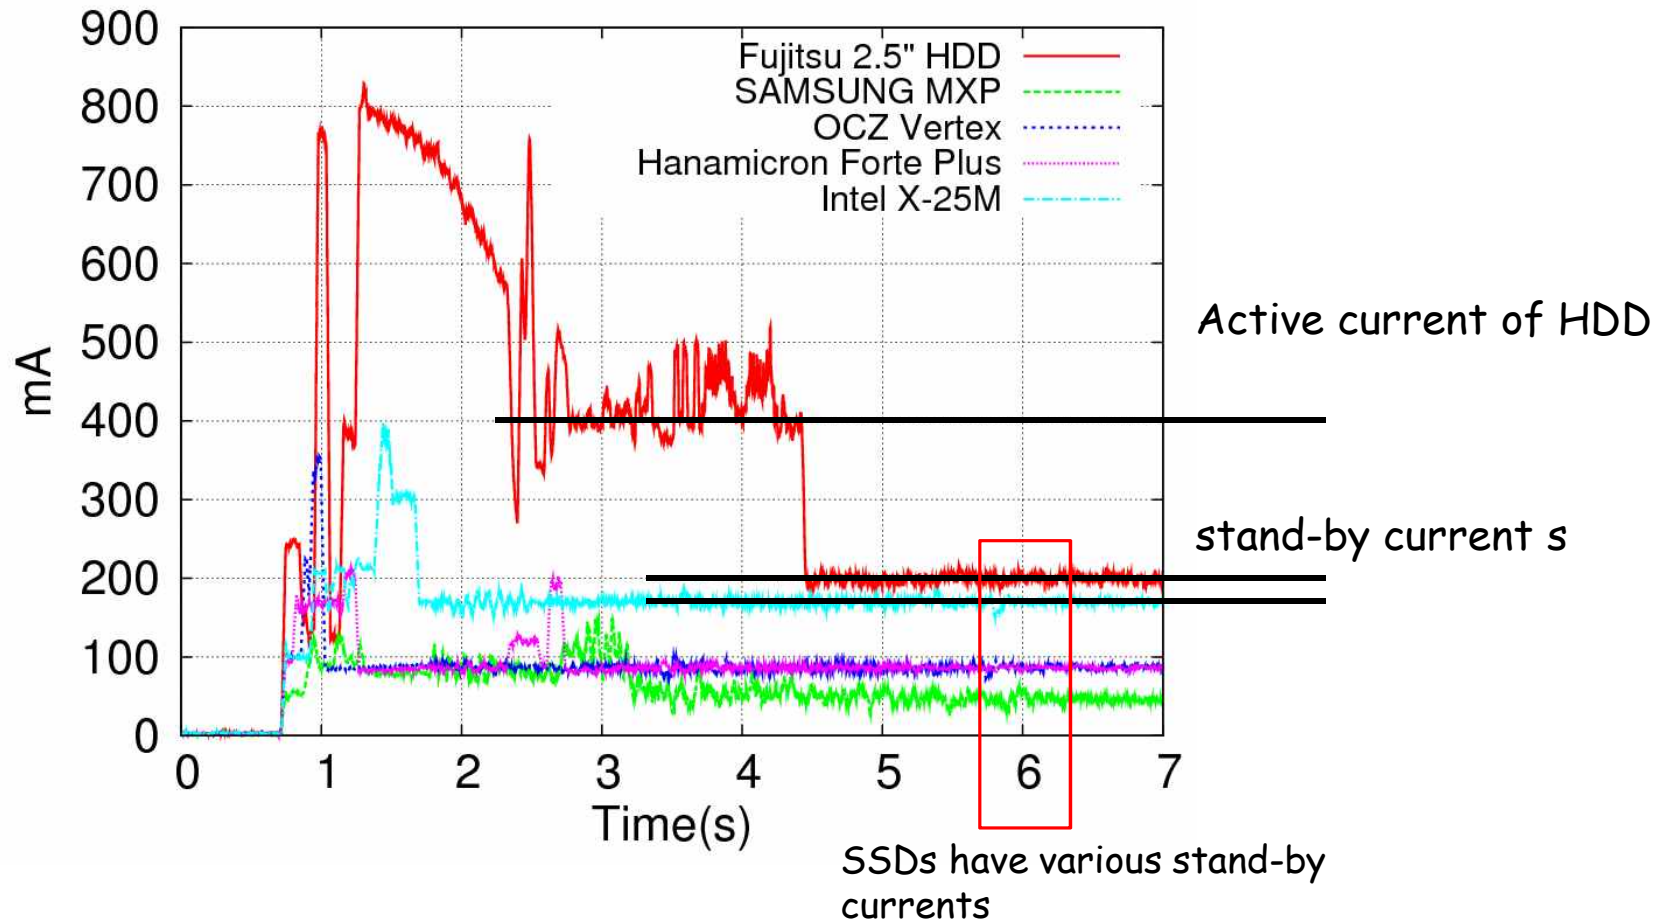


developed a detailed power model for the NAND flash chip itself with CACTI 5.3

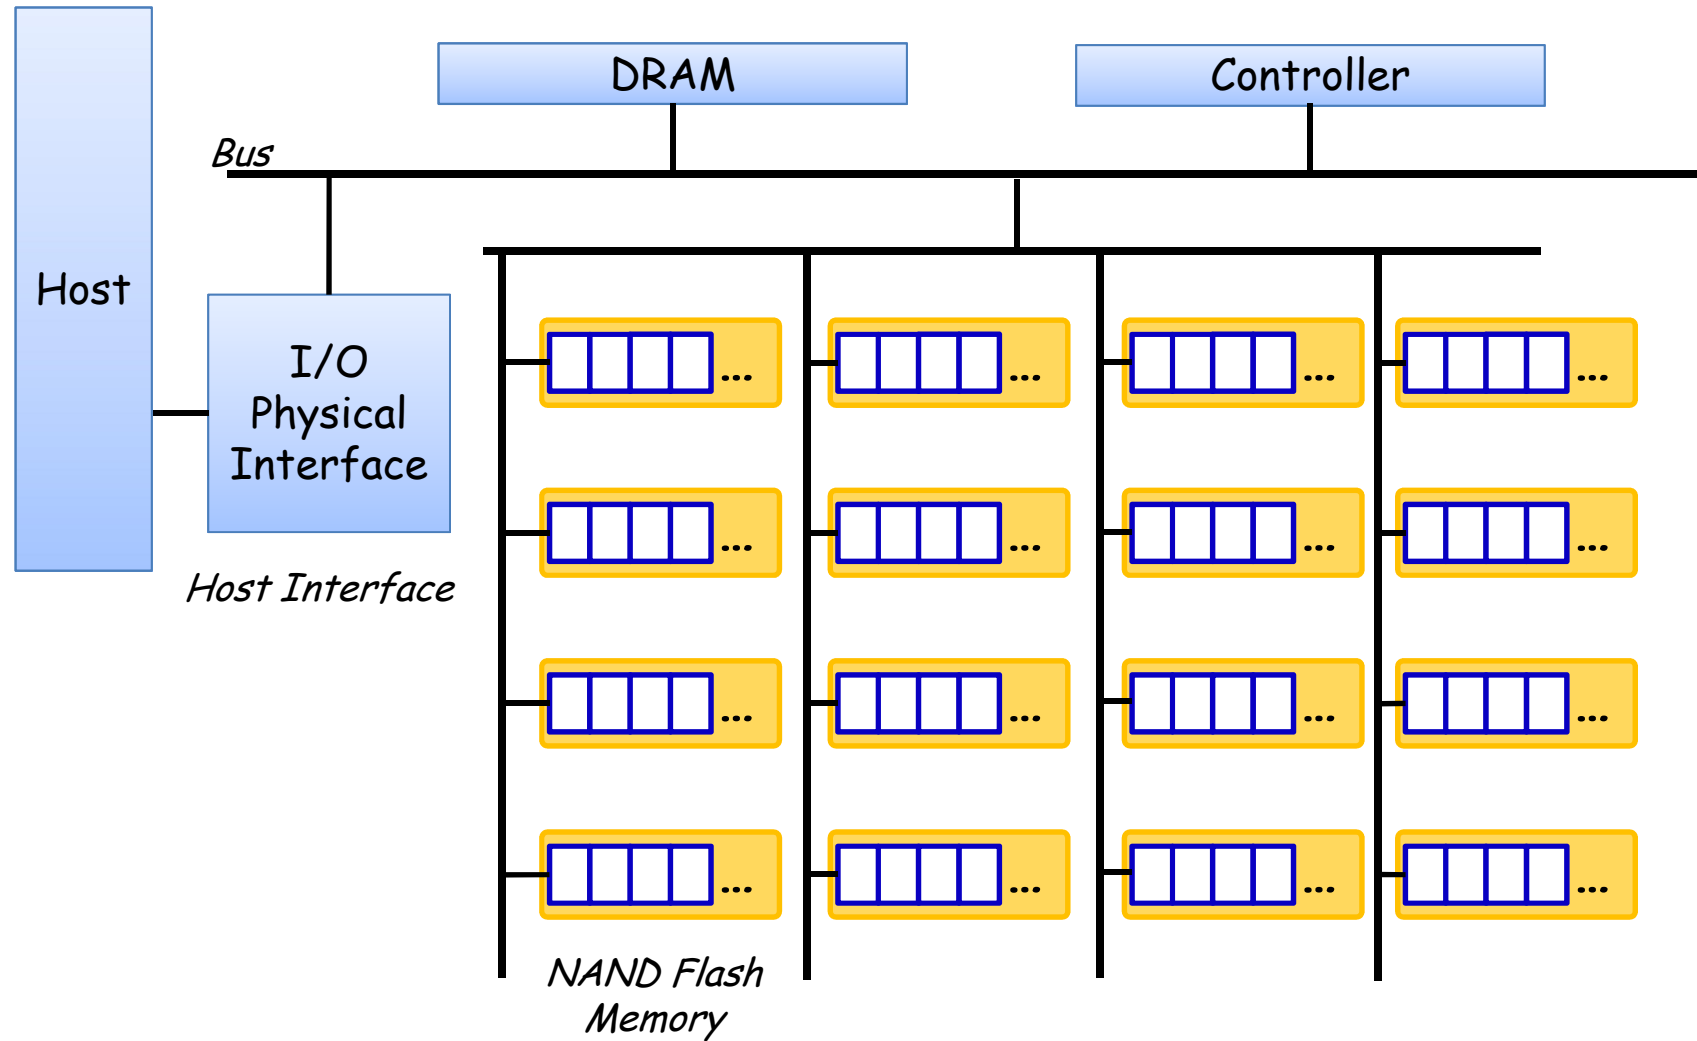


SSD Organization and Energy Consumption

Power Consumption of Storage Devices: System Boot



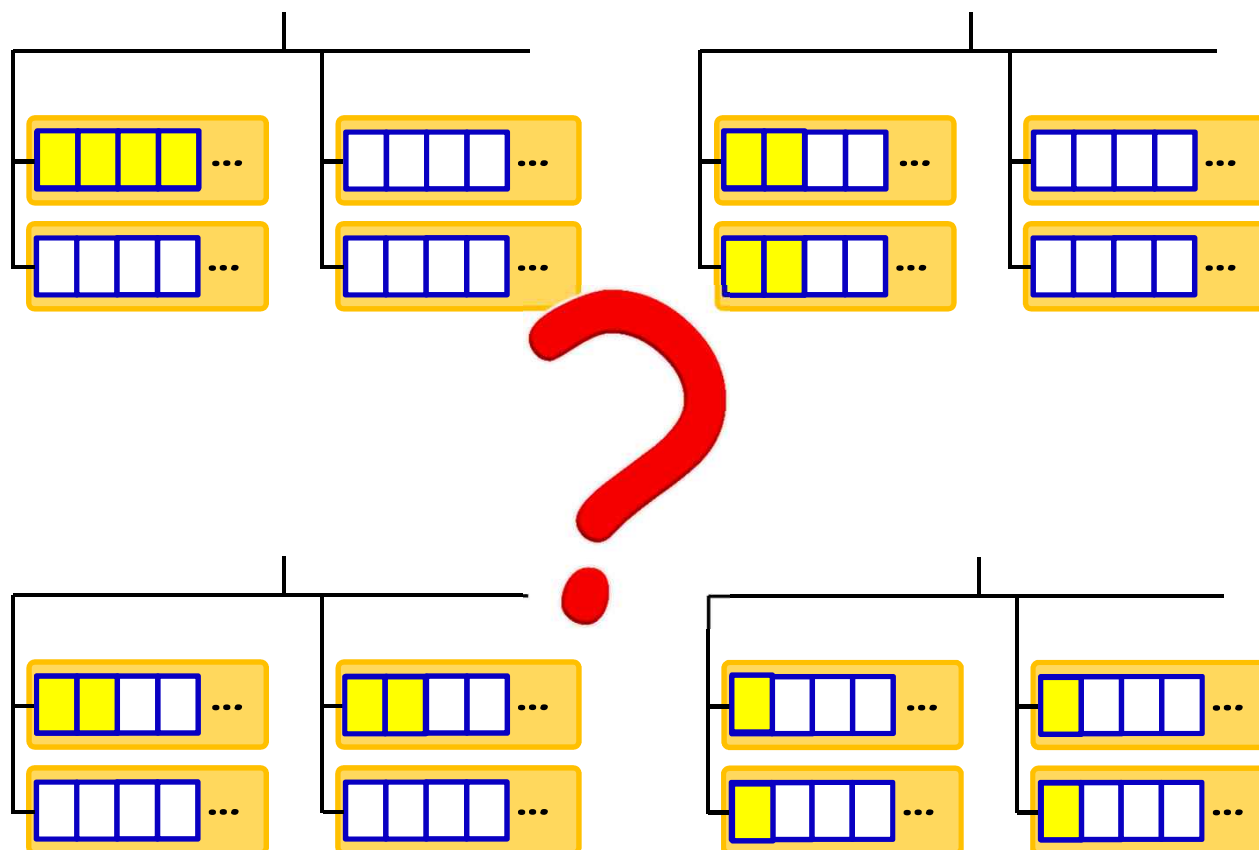
SSD Organization




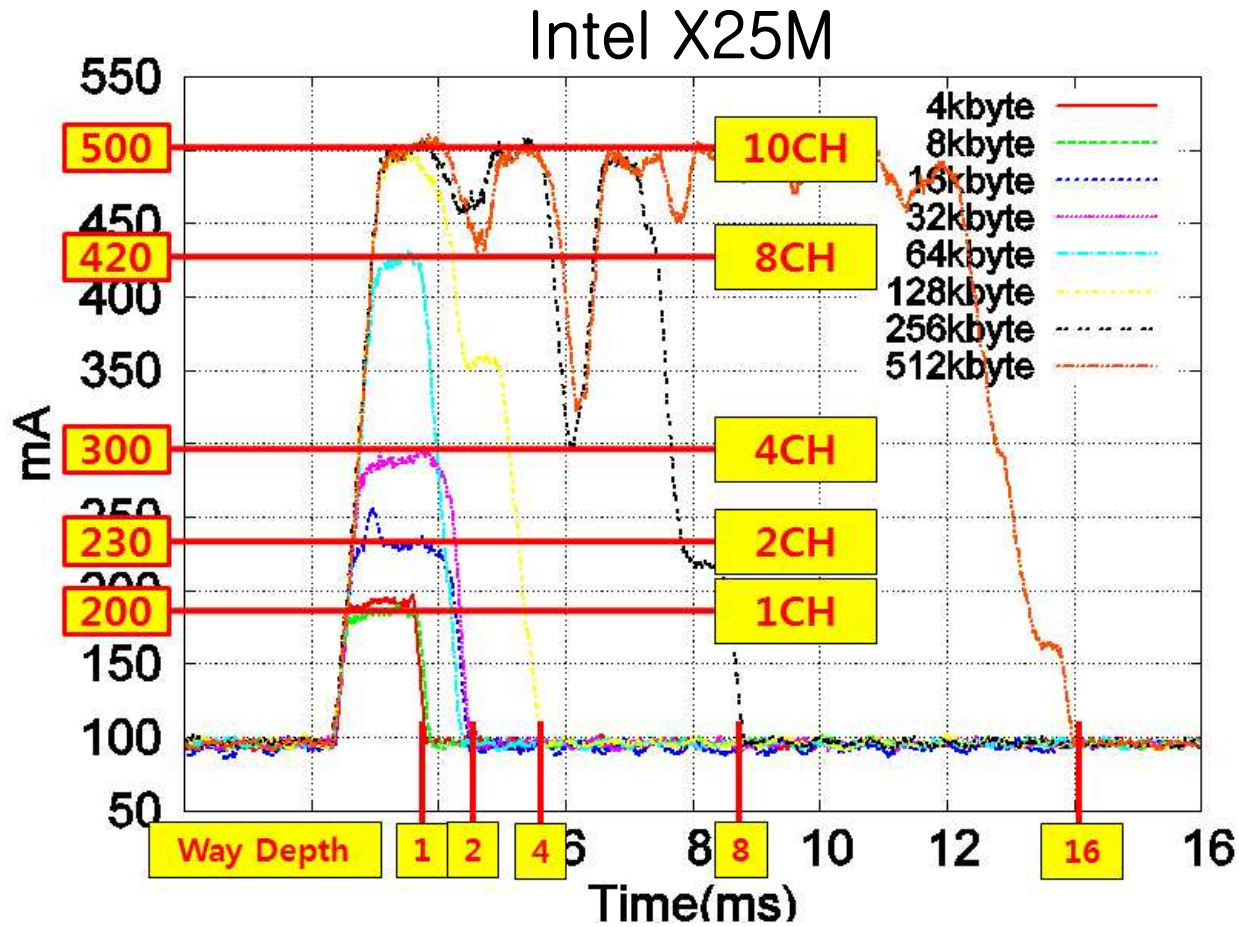
Characteristics of SSD Behavior

 =Page Write

 =Write Complete



IO Size vs. Energy Consumption




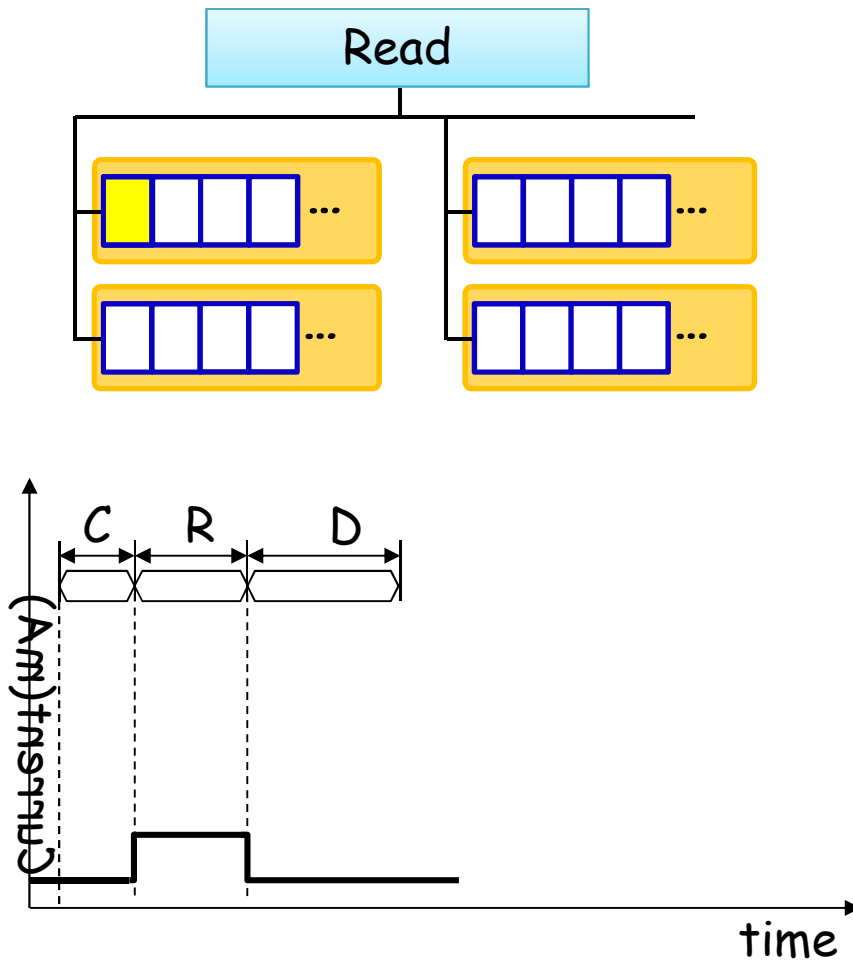


Power consumption vs. Channels, Ways, and Clusters

Read vs. Write


■ = Page Operation

■ = Operation Complete



Writing 4 pages: 1 Channel X 1 Way

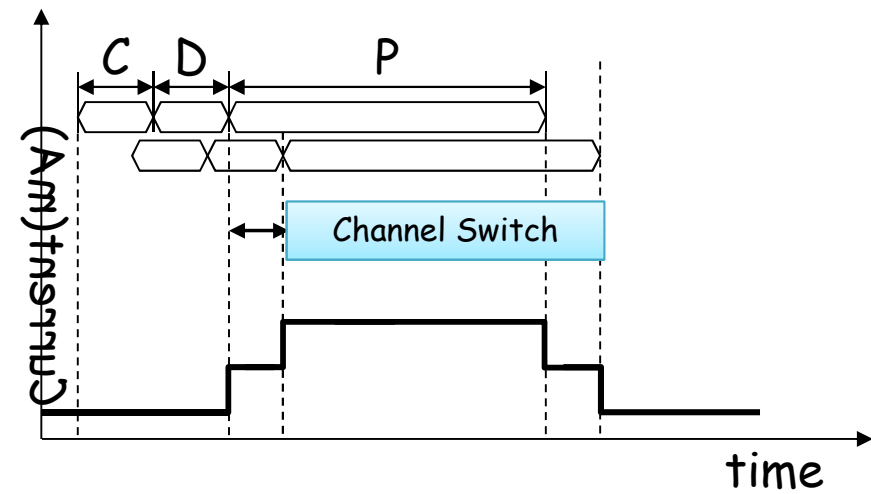

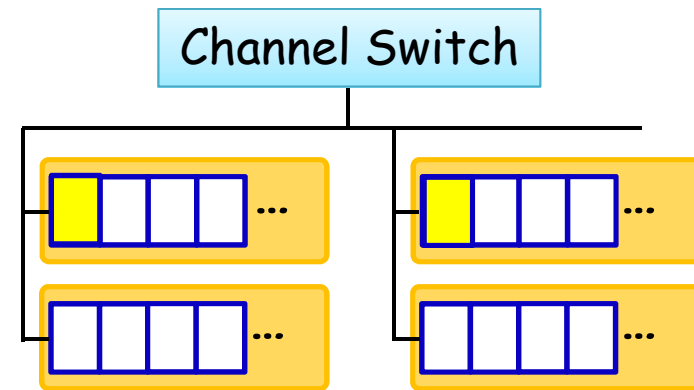
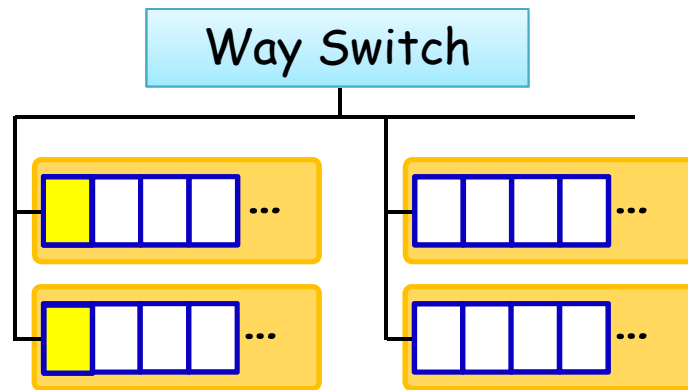
-  =Page Write
-  =Write Complete



Way switch vs. Channel switch Delay


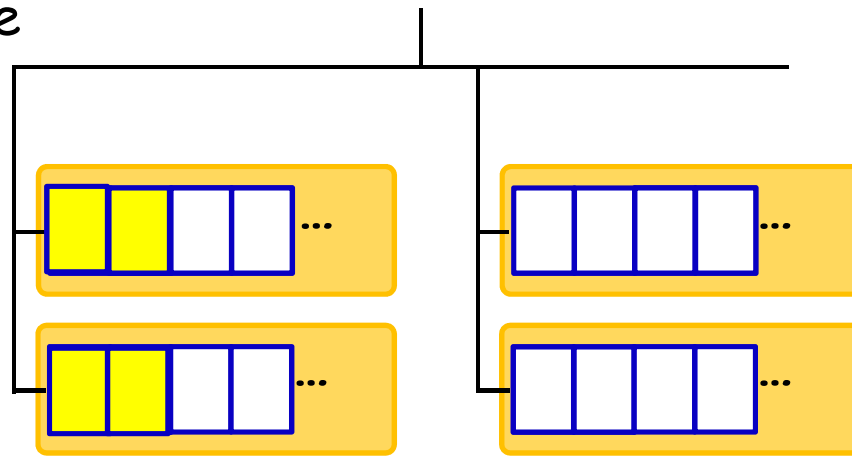
■ = Page Write

■ = Write Complete




Writing 4 pages: 1 Channel X 2 Way

-  =Page Write
-  =Write Complete



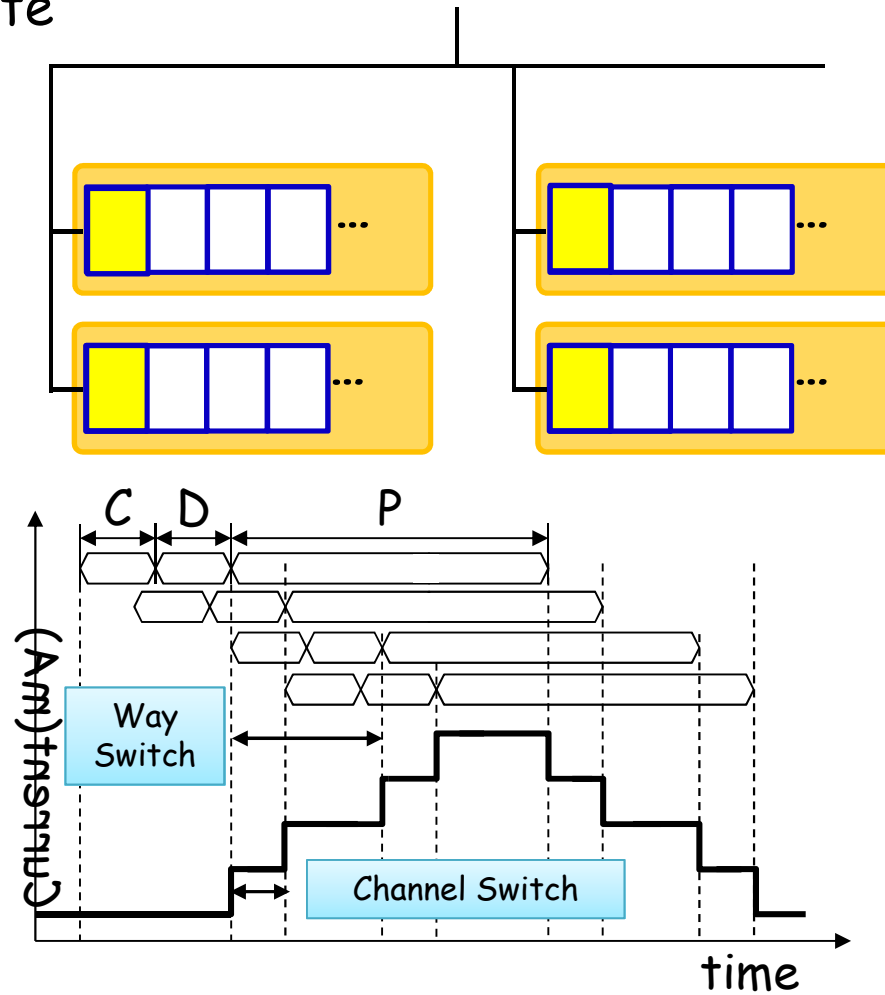
Writing 4 pages: 2 Channel X 1 Way

-  =Page Write
-  =Write Complete



Writing 4 pages: 2 Channel X 2 Way


-  =Page Write
-  =Write Complete



Cluster

- Cluster : Write Unit of SSD

 =Page Write



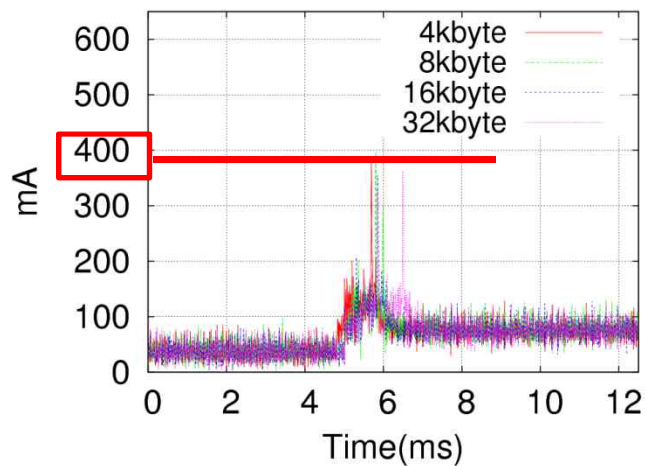


Trivia in Measurement Methodology


Trivia of Measurement Methodology

Sampling interval should be smaller than Read/Program Latency

- Smoothing the data to filter out measurement noise.



Smoothing



Trivia of Measurement Methodology

- Deactivate DRAM cache
 - ◆ SATA command 82h
 - Samsung MXP and Intel X-25M SSDs
 - ◆ Trigger DRAM flush
 - OCZ Vertex and Hanamicron Forte+



Case Studies

Measurement

Model	Vendor	Size	Channels	DRAM Size	Package	Type
X25M	Intel	80GB	10	16MB	20	MLC
MXP	SAMSUNG	128GB	8	128MB	16	MLC
Vertex	OCZ	60GB	8	64MB	16	MLC
Forte+	Hanamicron	32GB	8	32MB	8	MLC


- Oscilloscope (TDS3032)
- High resolution current probe(TCP202)
- Current probe to power line(Vdd) of the SSD
- Sampling interval(10samples MA): 10usec
- Each request measured 10 times



Case Study : Intel X25M

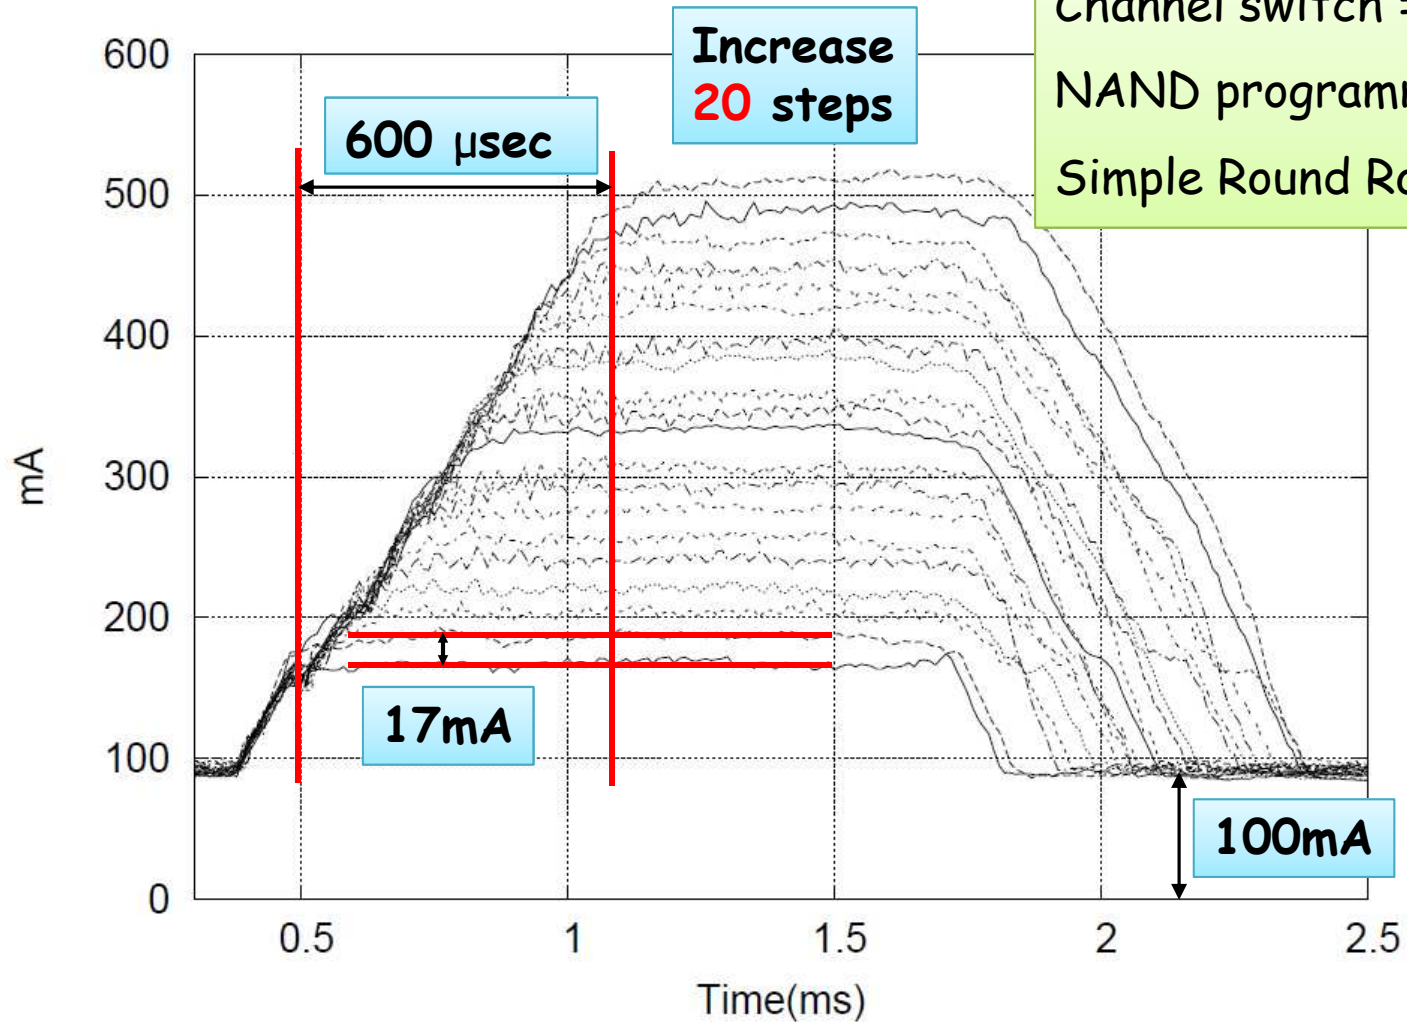
Intel X25M

Capacity	80GB
No. of Channels	10
Packages/Channel	2
Package	4 GB



Case Study : Intel X25M Write

- IO Size: 4KB to 80KB

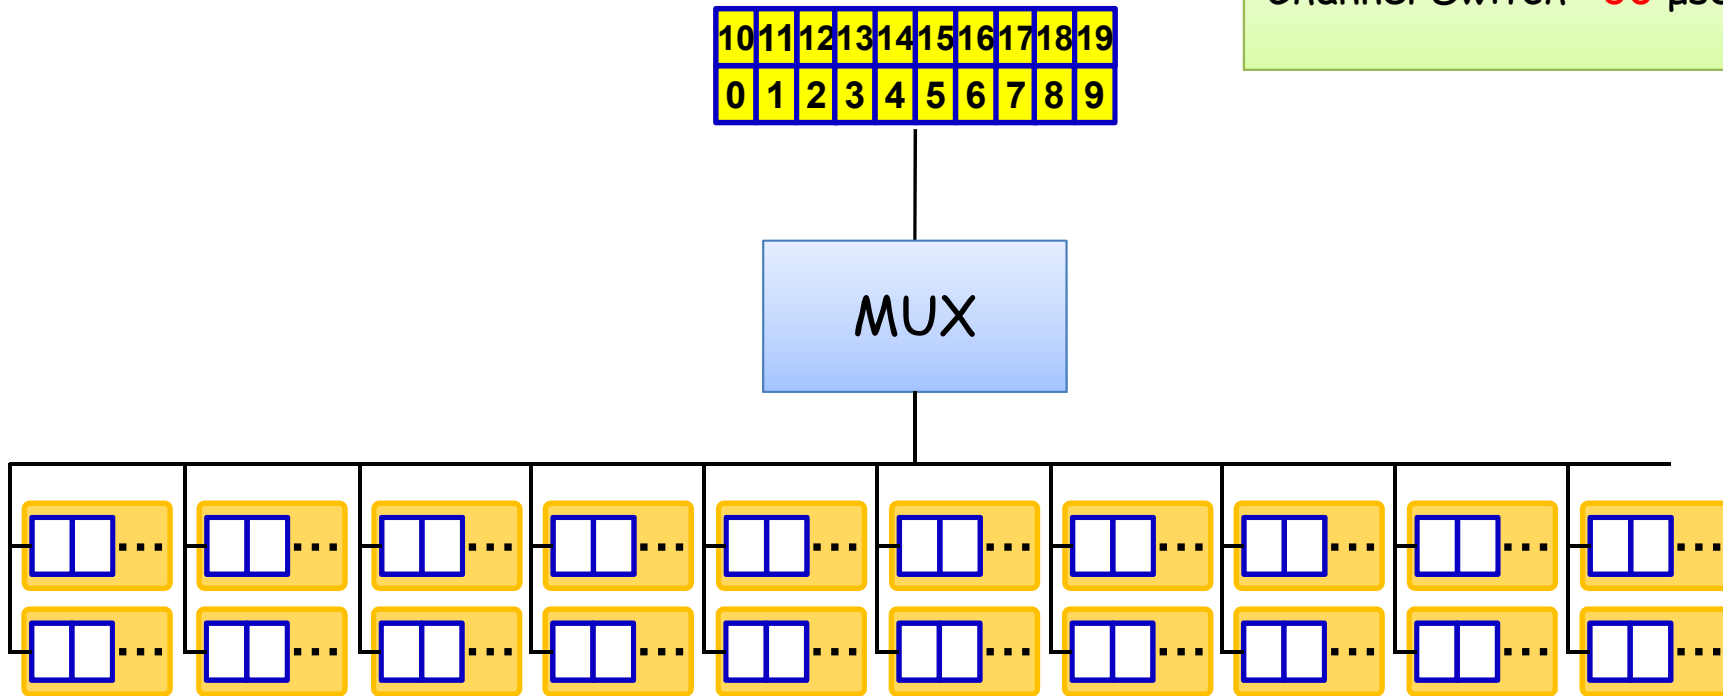


Case Study : Intel X25M Write

- 4KB to 80KB Sequential Write

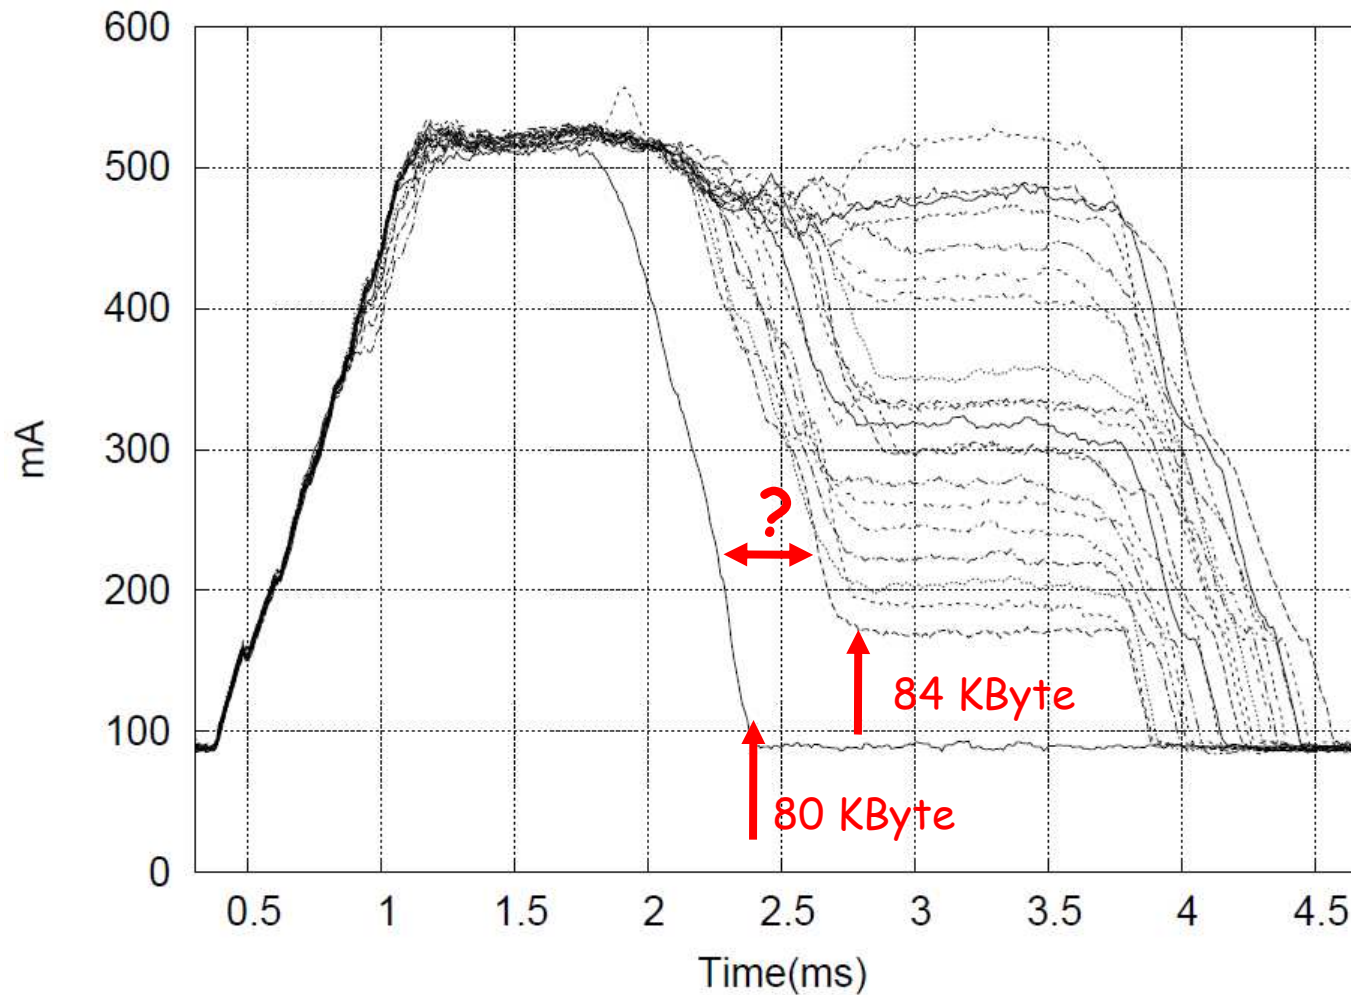
■ = Page Write

Channel Switch = 30 μ sec



Case Study : Intel X25M Write

- IO Size: 80KB to 160KB Write






Case Study : Intel X25M Write

- = 1 Page Write
- = Write Complete

10	11	12	13	14	15	16	17	18	19
0	1	2	3	4	5	6	7	8	9


MUX



Case Study : Intel X25M Read


- IO Size: 4KB to 16KB Read

Same Peak Current




Case Study : Intel X25M Read

- IO Size: 16KB to 28KB Read




Case Study : Intel X25M Read

- IO Size: 28KB to 48KB Read




Case Study : Intel X25M Read

- IO Size: 48KB to 52KB Read



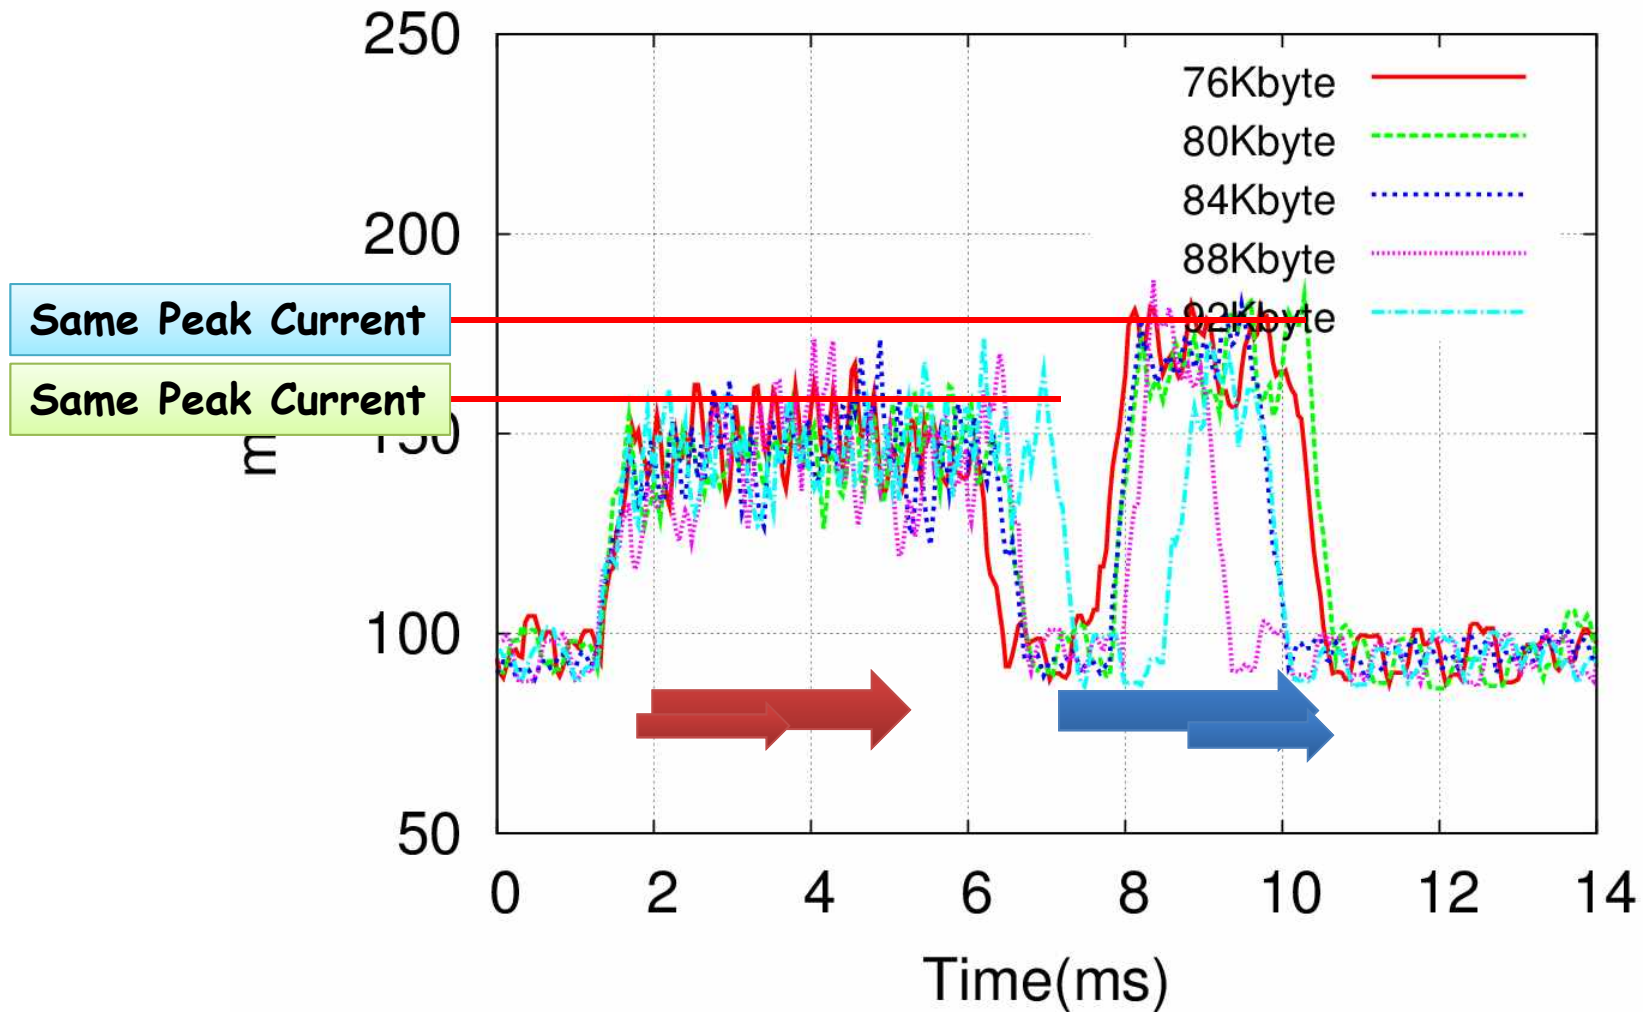
Case Study : Intel X25M Read

- IO Size: 56KB to 72KB Read




Case Study : Intel X25M Read

- IO Size: 76KB to 92KB Read



Case Study : Intel X25M Read

- IO Size: 96KB to 100KB Read



Case Study : Intel X25M Summary

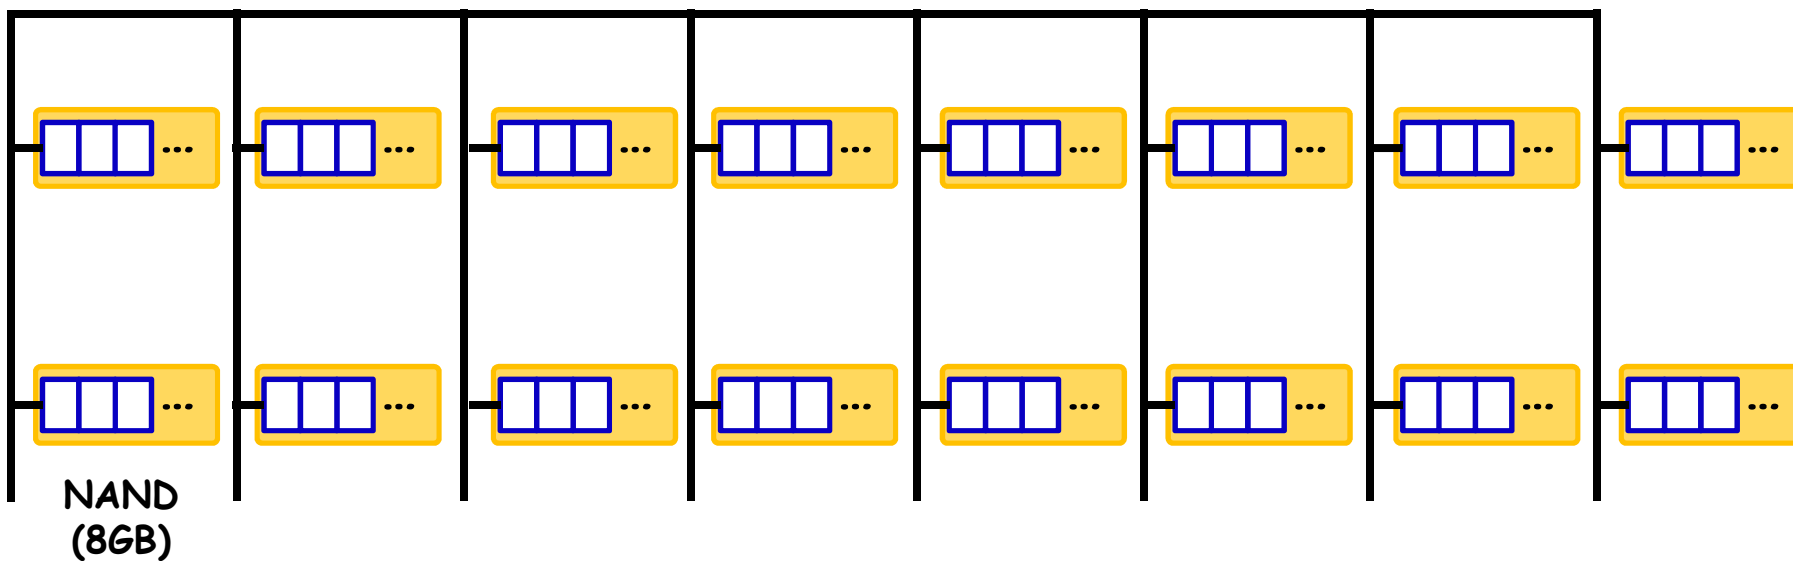
	Write	Read
✓ Cluster Size :	4KB	48KB
✓ Operation:	17mA	
✓ Peak Current :	500mA	180mA
✓ Channel switch:	30 μ sec	
✓ Policy	Round Robin	

- Different from Write, peak current remains the same.
- Only duration changes.

Case Study : SAMSUNG MXP


Samsung MXP

Capacity	128GB
No. of Channels	8
Packages/Channel	2
Package	8 GB




Case Study : SAMSUNG MXP

- Write Operation : From 4Kb to 32KB




Case Study : SAMSUNG MXP

- Write Operation : From 36Kb to 64KB




Case Study : SAMSUNG MXP

- Write Operation : From 68Kb to 96KB



Case Study : SAMSUNG MXP

- Write Operation : 100Kb to 112KB




Case Study : SAMSUNG MXP



What we still do not know of

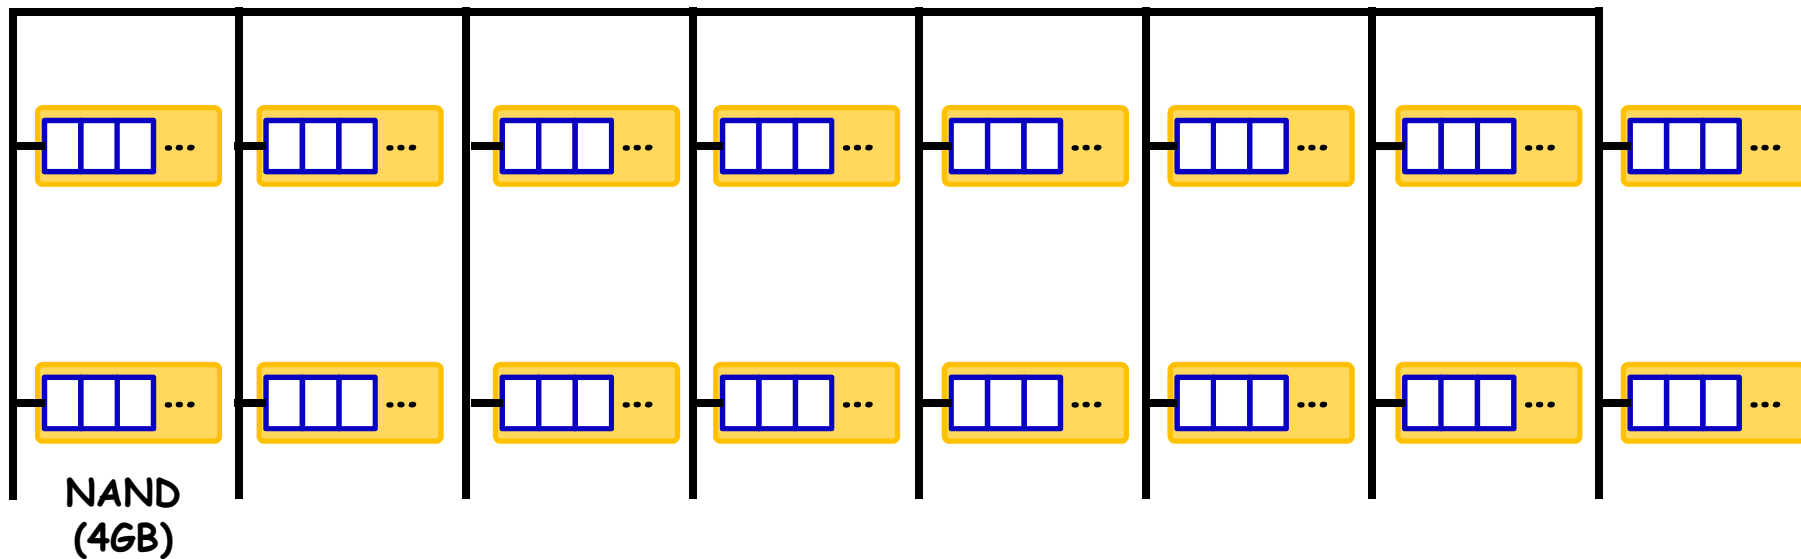
Peak of 128 Kbyte < Peak Current of 124 KByte



Case Study : OCZ Vertex





OCZ Vertex

Capacity	60GB
No. of Channels	8
Packages/Channel	2
Package	4 GB

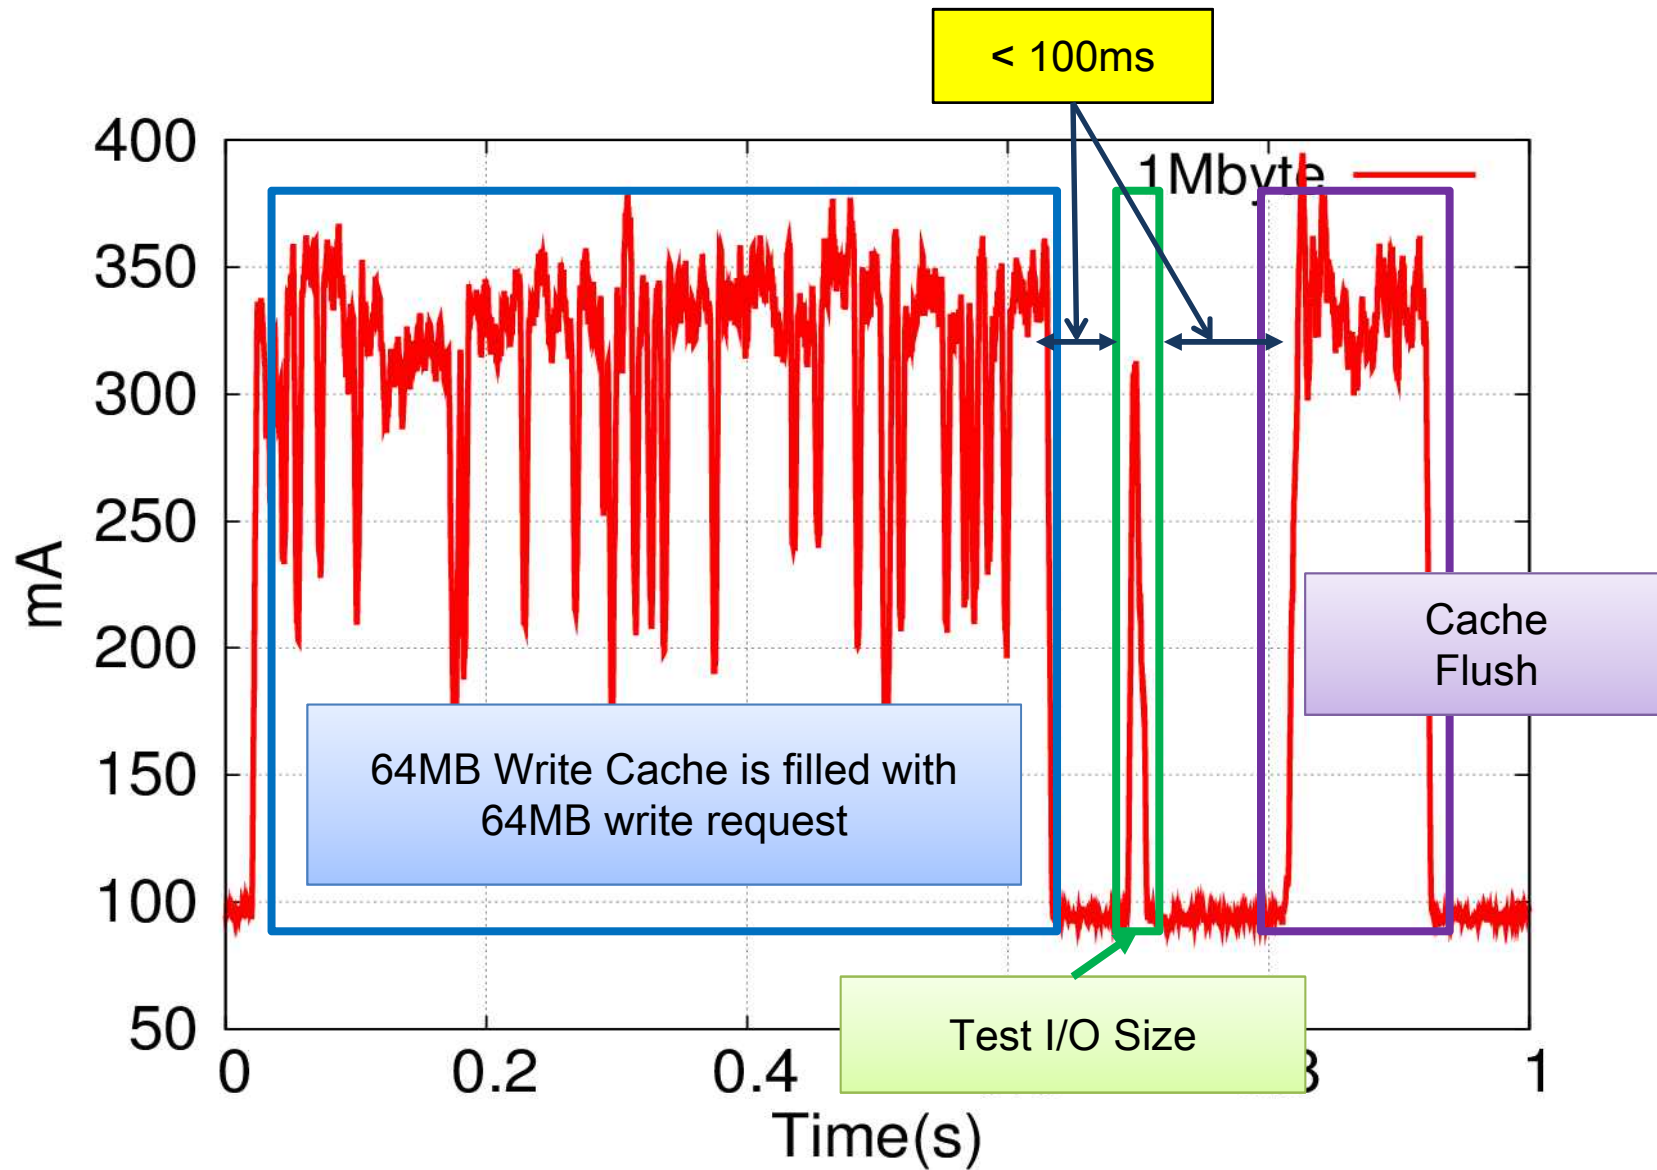


Case Study : OCZ Vertex

- write-caching off, 1KB write, 10 counts

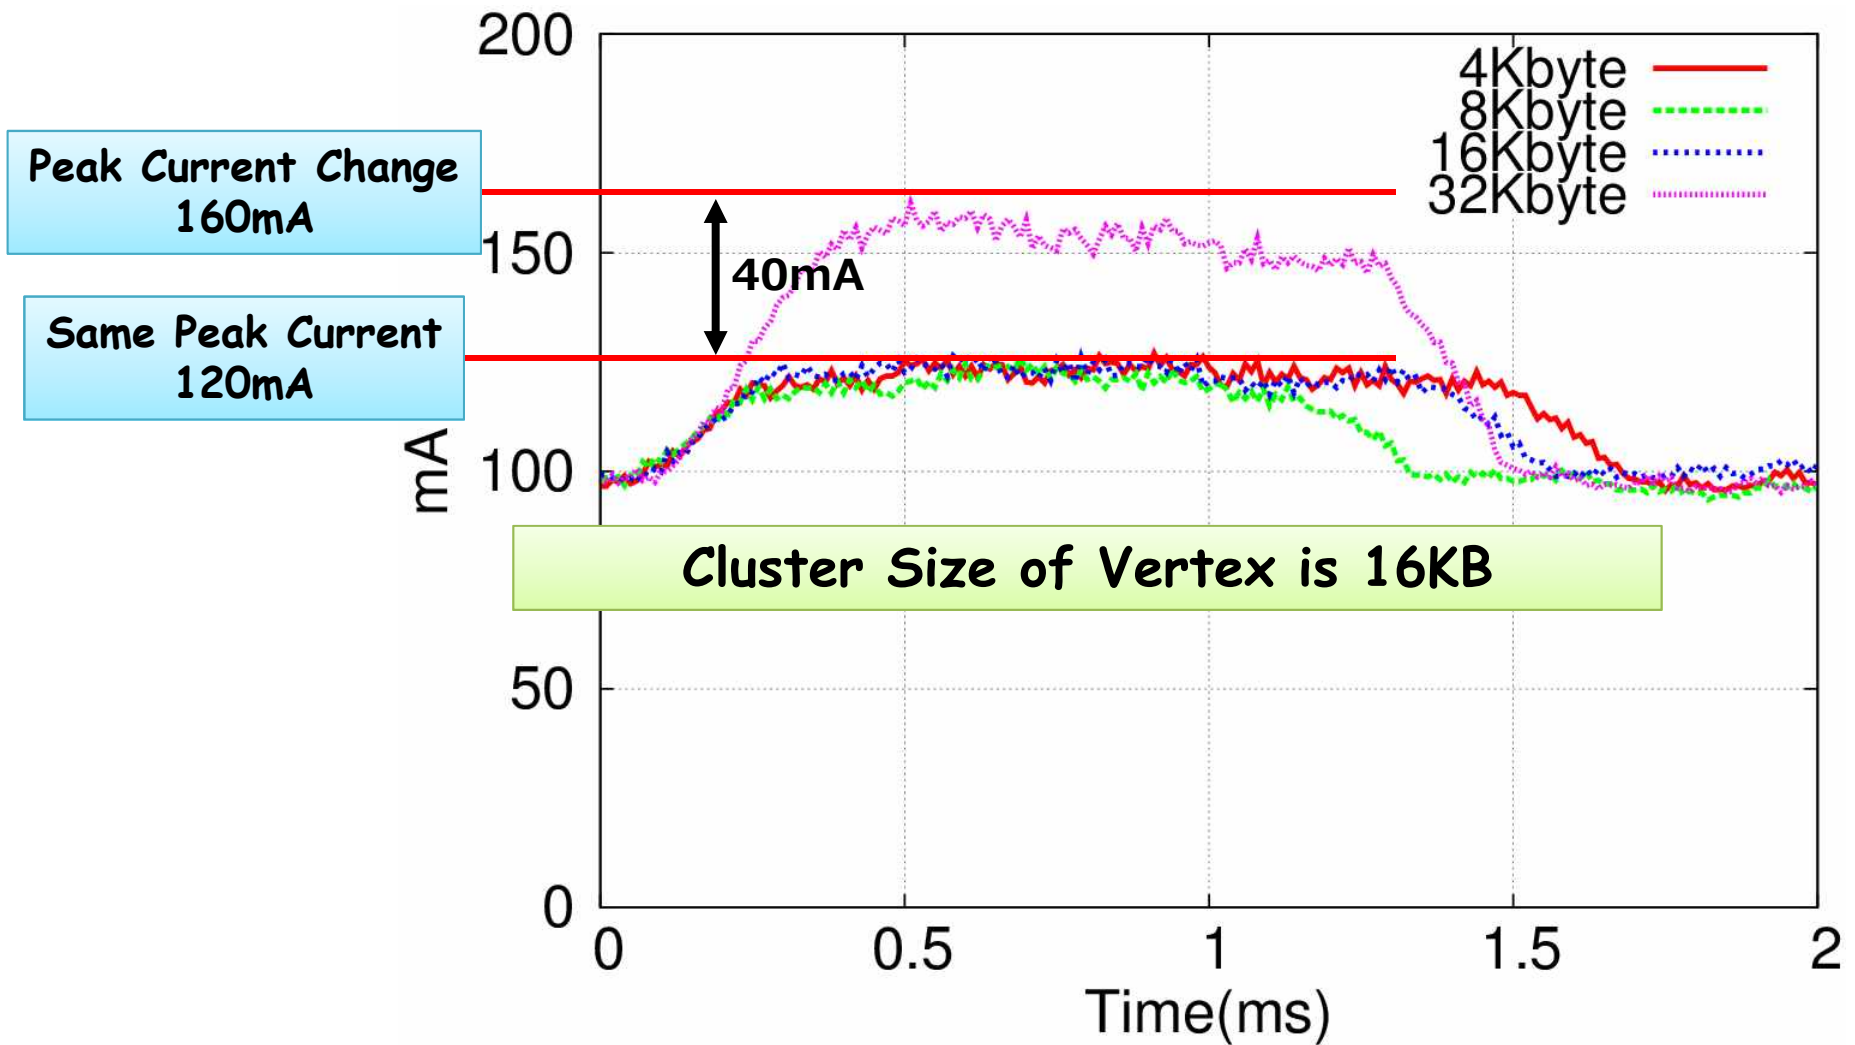


Case Study : OCZ Vertex




Case Study : OCZ Vertex

- IO Size: 4KB to 32KB




Case Study : OCZ Vertex

- write-caching off, 1counts, IO Size: 64k to 512k



Case Study : OCZ Vertex


- IO Size: 4KB to 4096KB Write



Case Study : Hanamicron Forte+


Hanamicron Forte+

Capacity	32GB
No. of Channels	8
Packages/Channel	1
Package	4 GB




Case Study : Hanamicron Forte+

- IO Size: 4KB to 32KB



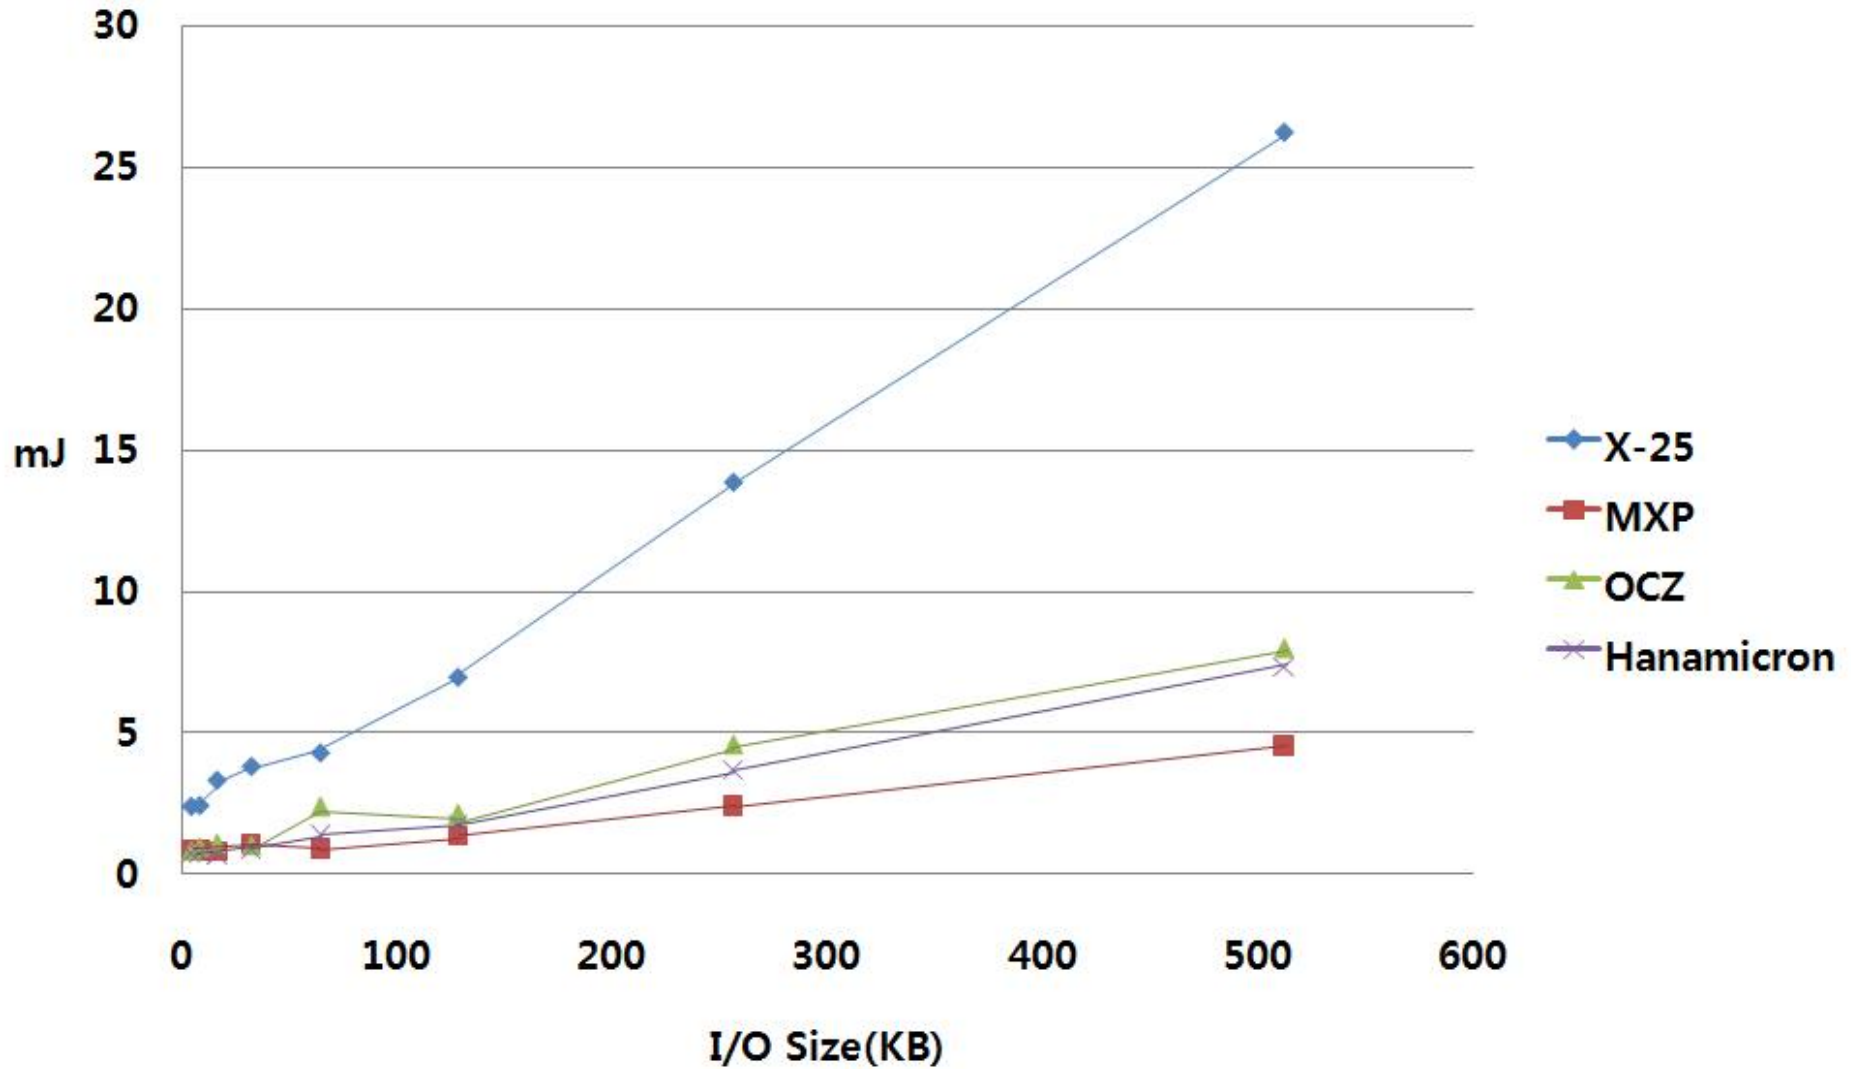
Case Study : Hanamicon Forte+

- IO Size: 4KB to 512KB Write




Comparison: Write Energy


□ Energy Consumption



Comparison: Peak and Duration of 512K Write



Forth Coming Problem in Multi-channel SSD

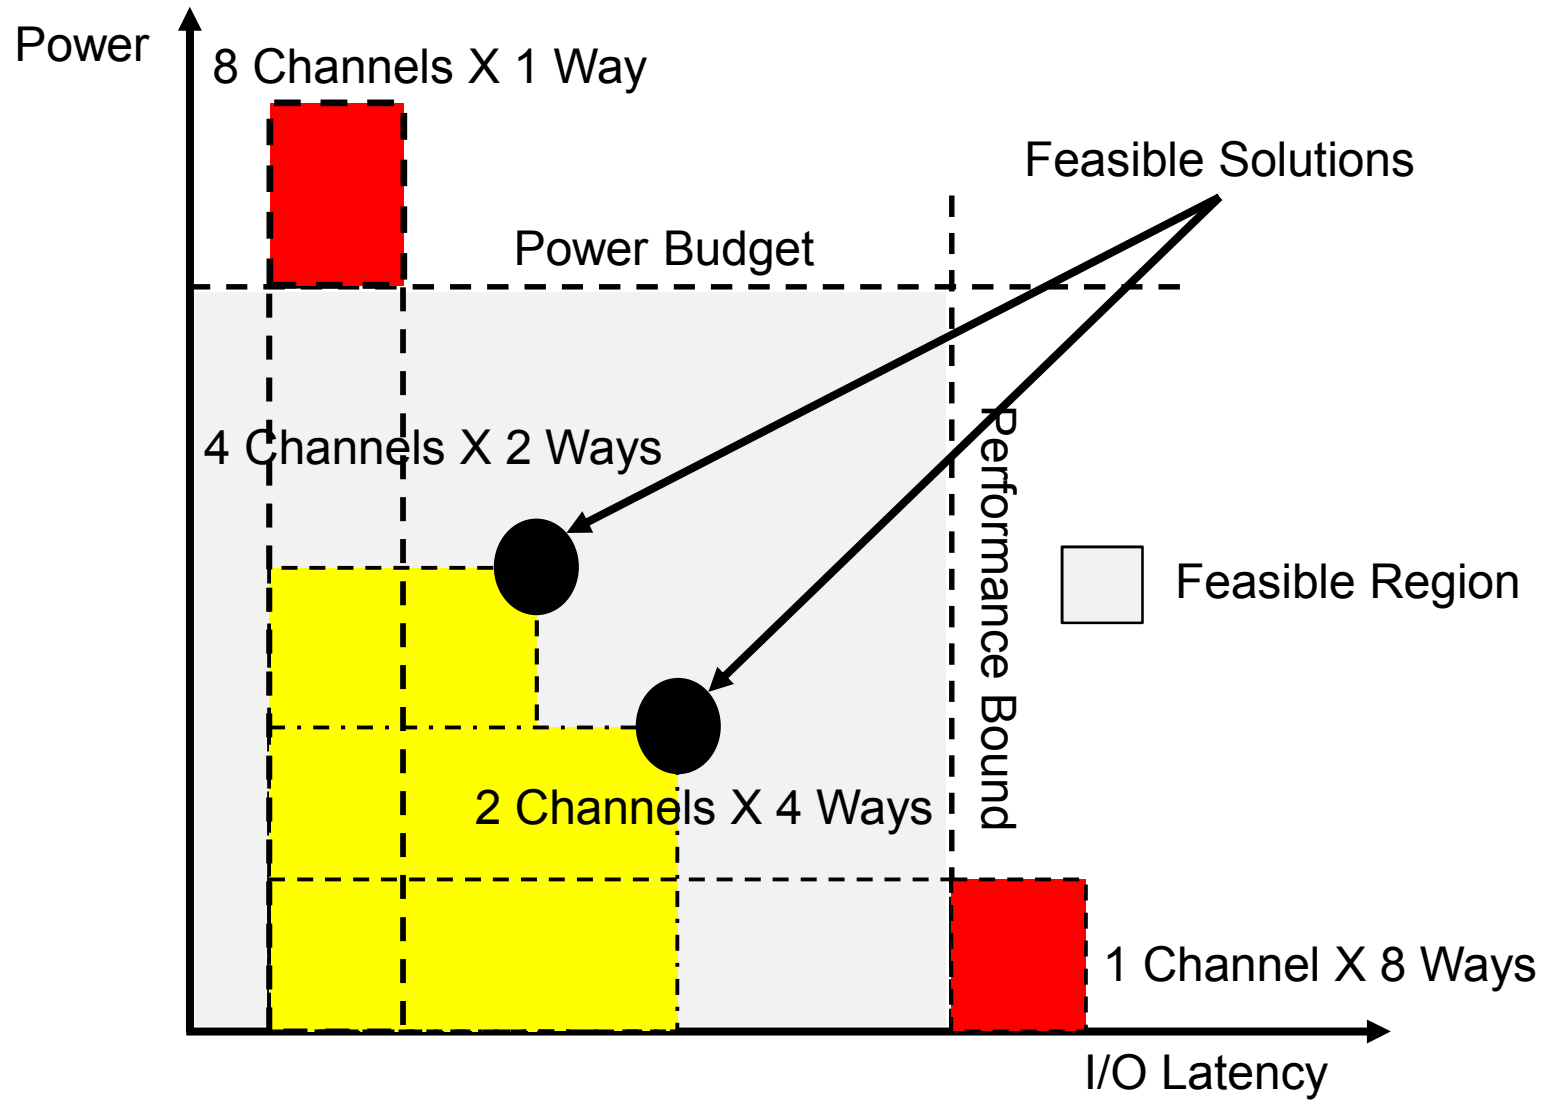


Model Name	Release
SATA3035 (Mtron)	2008.01
Vertex (OCZ)	2009.03
Vertex2 (OCZ)	2010.07
REVO Drive X2 (OCZ)	2011.01





Forth Coming Problem in Multi-channel SSD

- 10 channel: peak 500 mA
- 16 Channel: peak 800 - 900 mA → SSD is no long Green.
- Further, excessive peak current can cause...
 - ✓ supply voltage drop
 - ✓ ground bounce
 - ✓ signal noise
 - ✓ black-out
 - ✓ Etc...

Power Budget



Summary

	X25M	MXP	OCZ	Hanamicron
✓ Cluster Size :	4KB	32KB	16KB	16KB
✓ programming:	17 mA	35mA	15~20mA	15~20mA
✓ Peak Current :	500mA	350mA	350mA	220mA
✓ Channel switch:	30 μ sec			
✓ Energy				
✓ Performance				
✓ Small Write				
✓ Large write				
✓ Standby current				

Summary

- Energy Consumption is very good tool to characterize SSD.
- For larger number of channels, peak current will soon be a significant issue.
- We introduce the notion of Power Budget to resolve this issue.

