Flash's Role in Big Data, Past Present, and Future

Tutorial: Fast Storage for Big Data Hot Chips Conference August 25, 2013 Memorial Auditorium Stanford University



OBJECTIVE ANALYSIS

Semiconductor Market Research

- Market consulting/research firm
 - Market analysis, strategies, white papers
- Highly-respected lead analysts
 - Jim Handy: Memories & SSDs
 - Lane Mason: Memory chips
 - Tom Starnes: Processors
- Industry experience & 25+ years in field
- Reports, Competitive Analysis, Consulting

Haven't We Seen You Before?

 HotChips 2010: The Inevitable Rise of NVM in Computing



Agenda

- The problem
- SSDs as a solution
- The role of future memories
- Tomorrow's computing architecture

Big Data

- Data is exploding. By 2020:
 - Annual data production will be 35 ZB (CSC)
 - 50 billion "things" on the Internet (Cisco)
- Systems aren't keeping pace
 - CPUs moving along nicely
 - HDD/DRAM/Flash capacities still growing
 - Interfaces are not moving fast enough
- Result: Break the system into smaller chunks

The DRAM/HDD Speed Gap



From: Solid State Drives in the Enterprise OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com

"Sharding"

• Big Data? Subdivide the problem!

Sharded System



Agenda

- The problem
- SSDs as a solution
- The role of future memories
- Tomorrow's computing architecture

Add Flash, Cut Server Count

Sharded System

Single Server with SSD



Other Flash Benefits

- Less power & cooling
- Improved reliability

 The fewer things there are, the fewer will fail
- Floor space reduction
- Lower licensing fees
- Faster error recovery
 - RAID rebuilds
 - Backup restore
 - Snapshots

Agenda

- The problem
- SSDs as a solution
- The role of future memories
- Tomorrow's computing architecture

How to Maintain this Inertia?



From: Hybrid Drives: How, Why, & When?

NAND's Scaling Limit

- NAND will reach a limit
 - Too few electrons per gate
 - Needs constant shrinks for cost reductions
 - 4-bit/cell hard to make
 - This may be the maximum possible
- Other technologies will scale past NAND – PCM, MRAM, RRAM, FRAM....
 - Not yet clear which will win

An NVM Timeline

- 2013: 1Xnm planar cell
 Requires Hi-k gate oxide
- 2015: 1Ynm planar cell
 May be the last planar cell
- 2017: 3D NAND in volume
- 2019: 3D "The Next Generation"
- 2021: Final (?) 3D generation
- 2023: NAND yields to new technology

How Alternatives Will Emerge



New Memories are Better

<u>NAND</u>

- Serial read
- Erase before write
- Block erase/page write
- Slow write
- Inherent bit errors
- Wear

New Memory

- Random read
- Overwrite
- Byte write
- Fast write
- Lower error rates
- Low/no wear

Opens pathway to "Storage Class Memory"

Impact of New Memories



Agenda

- The problem
- SSDs as a solution
- The role of future memories
- Tomorrow's computing architecture

NVM Won't Cross HDD \$/GB



HDD Pricing Courtesy of PriceG2

An Evolving Computing Framework



NAND Fits in Computers Today



From: Solid State Drives in the Enterprise OBJECTIVE ANALYSIS – www.OBJECTIVE-ANALYSIS.com

Flash as Memory

Hard Way

- It's nonvolatile
 - Write new code
 - Create new topologies
- Worry about coherency
- Everything is new



- It's cheaper than DRAM
 - It's faster than HDD
- Who cares about volatility?
- Handle coherency like
 you do with DRAM

Use existing code

Flash More Economical than DRAM





Summary

- Flash belongs in <u>all</u> computers today
 So does HDD
- Flash vs. DRAM, <u>not</u> flash vs. HDD
- New NVMs will require new computing architectures

– "Flash as memory" can be used today



Thank You!

