Flash Memory Summit Review 08212008

By Ron Dennison

Introduction

The Flash Memory Summit was held last week in Santa Clara.
Well over a thousand people attended sessions which covered all aspects of flash memory technology and markets. I focused primarily on the enterprise application and technology side.
Links to some of the presentations can be found at
http://www.flashmemorysummit.com/English/Conference/Presentations_Chrono.html

Overview

The flash memory market is growing rapidly and currently stands at about $24B. However the difficulty is that no company except Samsung has made a profit in this market in its history. This has been largely due to rapidly falling prices (>40%/year) and low factory utilization which has resulted in oversupply. We currently have an oversupply situation. The analysts' consensus is that oversupply and pricing pressure will continue into (through?) 2009. OEM pricing stands at $2/GB for MLC NAND and $3.50/GB for SLC NAND.
A major new market seen by all competitors is the solid state disk (SSD) market which they hope will increase demand and stabilize prices. However this market is nascent and should it develop fully Eli Harari (chairman of SanDisk) believes 10 new megafabs (~$8-15B each?) will be required by 2013. The problem is that there is currently no source or rationale for this investment.
There is great interest from the enterprise storage market in this technology because of potential performance.

Enterprise Flash Sessions:

Various concepts and statements which had varying degrees of buyin by the panelists are reviewed below.

Reliability

1. Standards needs 1 umbrella group: JEDEC--- Steffen Helmold, Seagate
2. Data retention can be lowered in some applications to <1 year (< 1 month??) to trade-off performance for data retention in low latency applications--- Steve Garceau, STEC.
3. Reliability and durabiliity are going in the wrong direction; as geometry shrinks, number of electrons per gate decreases, increasing error rate (data volatility) and wear of the insulator tunneling barrier (durability). Wear leveling and more ECC help but there is no real answer except new technology. Scaling doesn't work well anymore!

4. Wearout must be handled by system firmware: automatic reallocation and S.M.A.R.T. reporting to host. Longer write times and more read errors indicate approaching failure.
5. Read Disturbance causing errors is a major issue for NAND.
6. Everything must be hot pluggable, planned downtime is an oxymoron--- Rob Peglar, Xiotech.

Cost and Performance

1. COST IS THE ELEPHANT IN THE ROOM!
2. Tier 0 storage is major potential use for NAND in enterprise
   Fits between conventional hard disk drives (HDD) and DRAM
3. Best Interface:
   a. SATA, SAS, 10G Ethernet for HDD replacment (storage space) versus,
   b. Memory (PCIe, DRAM, [NVMHCI (see below) + ONFI], or custom for applicance space): flash can be crippled by the interface (e.g. array speed is ~330 MB/s but slows to 30-34 MB/s due to bus limit therefore use ONFI I/F to increase bus speed to 200MB/s resulting in transfer rates of 120 MB/s for SLC NAND and 88 MB/s for MLC NAND--- Michael Abraham, MICRON). Backward compatible for existing controllers www.onfi.org www.nand.com www.micron.com/highsspeednand www.micron.com/roducts/nand/nand_webinars
4. Hybrid Storage Pool: Write biased SSD(SLC). Read biased SSD (MLC): **SUN** and BitMICRO advocate.
5. Virtualization lends itself to use of flash.
6. NAND technology is going in wrong direction for maximum performance: page size, block size (array write size), etc. are increasing
7. Increasing performance demands parallel access: multi-page writes and multi-block erases.
8. The correct cost metric is (No Consensus):
   a. $/IOP?
   b. $/GB?
   c. $/Watt?
   d. Total Cost of Ownership (TCO)?
9. The correct performance metric is (No Consensus):
   a. IOPS?
   b. Write IOPS?
   c. Random Write IOPS?
   d. IOPS/GB?
   e. Watts/IOP (high powerformance and low power is an oxymoron--- Marius Tudor, BitMICRO?)
10. DRAM SSD's still alive but DRAM cost is 3x NAND going to 5x; they will suffer.
11. There are currently 7 "real" enterprise SSD manufacturers

12. Spansion EcoRAM: replaces DRAM with specialized NOR in Virident internet servers with 8x improvement in power and price for similar performance. Jay Silverman, Spansion. Backed off on non-volatility of data to improve access time.
13. Unstructured data is a bad fit for SSD's, Structured data is a good fit--- Rob Peglar, Xiotech. True random workloads are rare--- Lawrence Sdchwartz, EMC.
14. SSD's will not replace HDD's because of difficulty in using all the IOPS available--- Adam Levanthal, **SUN**. Advocates Hybrid Storage Pool as combo of DRAM, SLC NAND, MLC NAND and HDD. Capitalizes on natural flash properties of quick reads and slow writes. ZFS operating system from SUN written to optimize use of Hybrid Storage Pool.
15. Video Servers are the natural market for flash for rapid ingest and high demand programming combined with HDD for lower demand material--- Tim Dodge, **Concurrent**.
16. **EMC** (Lawrence Schwartz) major proponent of Tier 0 enterprise storage. Believes flash is reliable enough now for job. Big advantage is time to data. Random Read Miss is much better for flash than HDD. Requires supporting infrastructure and software which is simply implemented and well integrated with current and future storage trends (automatic tiering, virtualization, thin provisioning, etc.). Showed TCO for Fortune 500 companies as 50% of traditional HDD + DRAM cache solutions.
17. Global shared memory is a better architecture than clustering.
18. **Network Appliance** is building its own controller to implement Reid Solomon ECC with unmasked consumer chips--- CSUSC Assoc. Prof. Ethan Miller
19. SATA is great due to plug compatibility but can do better over time with Non-Volatile Memory Host Controller Interface (NVMHCI). NVMHCI 1.0 completed and downloadable at www.intel.com/standards/nvmhci . Logical interface only, optimized for cache, has spare area for Metadata (mainly used for ECC but can be used for Metadata). Only 8 commands; one read, one write, flush; ID, get features, set features, get status, data set management. Is a register level interface that defines how to communicate with registers. Amber Huffman, Intel Storage Technologies Group

The Future

1. Best application of flash is replacement for DRAM in servers--- Michael Cornwell, **SUN**; best application for NAND in next 5 years is storage arrays--- Robin Harris, Storage Mojo
2. Tier 1 SAS will disappear, IF flash pricing comes down to meet HDD pricing [current differential is 8 to 10x in $/GB for the enterprise--- RD]--- Steve Garceau, STEC
3. Various presentations were given on future replacements for Flash. These included IBM's spin transfer torque domain wall motion based Racetrack, Grandis' Spin Transfer Torque RAM, SanDisk's 3D crosspoint diode array using RRAM material (only ROM so far, due to lack of "right" material but there may

be an announcement soon), and Numonyx's Chalcogenide glass Phase Change Memory (PCM). See the presentations link above.

The Market

Views of some of the panelists were:

Jeff Janukowicz, IDC

- SSD Market 2007 $396M revenue
- Mostly military, industrial, rugged, enterprise
- Major growth to be enterprise due desire for performance, better utilization, faster access
- 2nd Major growth PC's using MLC
- SSD mkt growth 85% CAGR

Alan Niebel, Web-Feet Research

Flash Memory Forecast

- 2008 decline in revenue vs 2007, return to growth in 2009 due $2/GB MLC.
- NAND surpasses DRAM in 2011, EcoRAM good potential growth
- Mobile is handset vs portable (laptop)
- Heirarchy HDD, SSD, EcoRAM, DRAM, CPU
- SSD definition includes the following interfaces: FC, SAS, ATA, SATA, SCSI, CE-ATA, [PCCard (PCMCIA)], USB, PCIe.

Jim Cantore, JLC Associates

- 12 years for replacement NVM's 3-4yr cycles 4-3yr cycles

Jim Handy

- NAND price/GB is below DRAM from mid 2004
- This makes putting NAND between DRAM and HDD reasonable
- Oversupply due Capex likely to extend into 2009

Jim McGregor

Market Changes Drive Flash Demand

- Impact of business models (who and how to make money) and mobile internet (eg. $10/mo for access will limit)
- Plethora of mobile devices
  - Move to mobile society
  - Internet bi-directional rich data content
  - Changes in intenet usage eg. Social networking
- New technologies
  - High performance processors and graphics
  - Advanced uses, browsers
  - Access to wireless broadband
- New Business Models = major market inflection points within 5 years
  - EVDO Sprint

Downloadable manuals

Flash demand: SSDs in mainstream mobile PCs in 2009/2010 (~25%), Increasing internet access in mobile devices (~48%)