Enabling MLC Flash SSD In Enterprise Storage

Gary Tressler
IBM Corporation
gtressle@us.ibm.com

August 17, 2010

Agenda

- Enterprise MLC SSD Enablement: Current Trend
- MLC NAND Scaling Trend & Outlook
- Enterprise Support Challenges: MLC SSD Controllers
- Enterprise Support Challenges: MLC Flash Media
- Dwell Time: Key Lever for SSD Reliability Improvement
- SSD Small Factor Innovation Required for Enterprise
- Summary

MLC SSD Enablement: Current Trend

- MLC Flash component endurance support has been extended to between consumer MLC and SLC for a given technology
  - Achieved via early signal processing implementation or program time degradation
- Power off data retention now optimized for Enterprise SSD
  - Duration requirement reduced from traditional 1-year to 3-months
    - Aligned with Enterprise field maintenance strategies
    - Further optimization may be required as technology advances
  - Temperature support reduced from 55'C to 40'C
    - Consistent with most data center environments
- ECC support per sector enhanced by industry SSD controllers
  - Required to satisfy 3x/2xnm MLC targets
  - Continuing improvements required in architecture / support level
- **NET: *MLC Flash provides leadership Enterprise SSD performance and lifespan today based on existing storage infrastructure***
  - **Higher capacity and lower cost structure enabled**

MLC NAND Scaling Trend & Outlook

	2009	2010	2011	2012	2013	2014	2015
Technology (nm)	3x	2x	2y	2y/1x	1x	1x	1y
Cell Architecture	Floating Gate	Floating Gate	Floating Gate	Floating Gate	Floating Gate	Charge Trap Device / 3D	Charge Trap Device / 3D
Bit Line Architecture	Shared Bit Line / All Bit Line	Shared Bit Line / All Bit Line	All Bit Line	All Bit Line	All Bit Line	<New Sensing Scheme>	<New Sensing Scheme>
MLC Density (Gb)	32	64	64	128	128	256	256

- Floating gate NAND cell scaling likely to continue to 1xnm technology node range in ~ 2013, based on conventional immersion lithography
- NAND scaling expected to continue beyond 2013 @ lower 1x nm NAND nodes, based on charge trap device and 3-D NAND cell & process architecture, reaching MLC NAND densities of 256Gb & higher
- NAND scaling rate and pace is based on consumer and not enterprise reliability requirements
- Ongoing enhancements (e.g., signal processing) required to achieve Enterprise-level MLC reliability requirements vs. time

NET: ***NAND scaling expected to continue via floating gate extension & charge trap / 3D cell architectures to 1xnm technology node range and 256Gb MLC density***

- ***Ongoing reliability enhancements required to satisfy Enterprise-level MLC requirements vs. time***

Enterprise Support Challenges: MLC SSD Controllers

Controller Feature	Description
Advanced wear leveling	▪ Continuing innovation to spread Flash usage as evenly as possible to help maximize SSD life
Flash usage management	▪ Implementation of various potential techniques such as compression, de-duplication, etc. to help limit Flash wear
Forward-looking ECC architecture / coverage	▪ Advanced architecture capability ▪ Multiple Flash technology node support
Write amplification optimization	▪ Provide maximum media lifespan capability
Advanced signal processing support	▪ Flash component partnership likely required ▪ Key next-generation Flash reliability improvement mechanism
System-level throttling capability	▪ Innovative methods to help manage Flash bandwidth to achieve improved SSD lifespan
Response time improvement	▪ Drive average and maximum response time improvements with increased SSD complexity
Trim command support / utilization	▪ Vehicle for improving mature SSD performance / lifespan
Power consumption reduction	▪ Manage power with increased performance ▪ Subsystem-level RAID level power will increase with Flash optimization
End-to-end data integrity support	▪ Fundamental requirement for Enterprise SSD

Enterprise Support Challenges: MLC NAND Flash Media

- Providing Enterprise support while driving aggressive technology advancement / cost reduction for Consumer applications with reduced reliability requirements is critical Flash industry challenge
- MLC NAND technology lifespan must be extended to support Enterprise product requirements
- New MLC component designs must offer additional spare area per sector to accommodate advanced ECC algorithms required for Enterprise
- Next key technology transition is advanced signal processing access capability provided by MLC suppliers to SSD controller designers / manufacturers
- Component read disturb specifications must be better understood by controller companies / users
  - JEDEC read disturb specification needed
- Drive-to goal for Enterprise Flash storage must be *to* ***apply consumer MLC to enterprise SSD applications with sufficient controller Flash management and reliability***

Dwell Time: Key Lever For SSD Reliability Improvement

- Definition
  - Dwell time – elapsed time between program operations for individual Flash component cells
    - Typically 2-10 min. for current industry component qualifications based on density and endurance targets – result based on qualification goals not Enterprise system requirements
- Concept
  - Increased dwell time allows for improved Flash cell recovery (de-trapping)
  - Initial test data shows increased dwell time improves retention UBER and RBER at both 25'C and 85'C
    - Data on different suppliers and various technology nodes
  - NET: Flash component endurance can be increased at constant data retention with extended dwell time
- SSD / system applications
  - Flash component dwell time is on average significantly longer in system environment behind host and SSD controllers with Flash management architectures
  - System- and storage subsystem-level characterization in process
- **Recommendation**
  - **Drive Flash component dwell time specification for Enterprise SSD based on available component-, SSD- and system-level data**
    - **Significant reliability improvement opportunity exists**

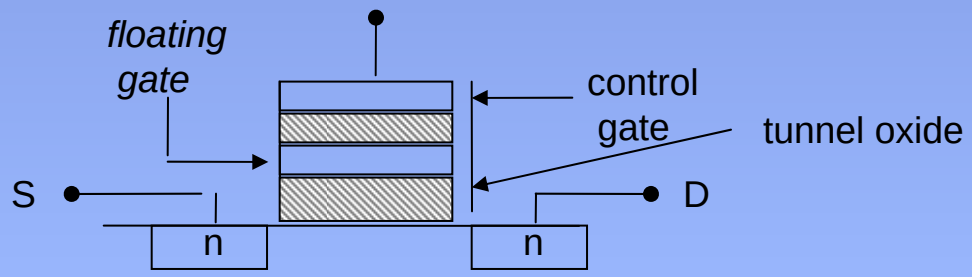

Small Form Factor Innovation Required for Enterprise

- **SSD**
  - ***1.8" Enterprise SFF SSD*** enables high capacity and throughput per system space allocated
    - MLC Flash media is key to value proposition
    - Enterprise data integrity requirements must be satisfied
    - SATA interface now, migrate to SAS by YE11
    - Industry defined Micro SAS connector SFF (8486)
    - Connector power requirements must be defined for 6 and 12Gbps interface speeds, enabling maximum SSD performance
    - Application opportunities
      - Boot drives for systems requiring small footprint and minimum space (e.g., blades)
      - Low-end rack servers
      - PCI-e adapter card building block
  - High performance ***one-chip SSD controller*** solutions without dependence on external caches required
  - New approaches needed for ***super capacitor*** designs (if required) to help improve reliability and reduce space occupancy
    - Discrete capacitors
    - Other backup mechanisms (e.g., phase change memory)

Small Form Factor Innovation Required for Enterprise

- **NAND Flash Media**
  - **12mm x 18mm BGA** body size support must be drive-to goal with advanced technology and increased density
    - Enables increased number of controller channels / Flash sites
  - Packaging innovation required to support **greater than 8 dice per Flash package** for increased capacity with similar reliability

Summary

- MLC Flash currently provides leadership Enterprise SSD performance and lifespan today based on existing storage infrastructure
  - Higher capacity and lower cost structure enabled
- NAND scaling expected to continue via floating gate extensions & charge trap / 3D cell architectures to 1xnm technology node range and 256Gb MLC density
  - Ongoing reliability enhancements required to satisfy Enterprise-level MLC requirements vs. time
- Enterprise SSD controller and Flash media challenges must be overcome to ensure MLC support through time
- Flash component dwell time is likely new lever for Enterprise SSD reliability improvement
- Small form factor innovation is required for Enterprise storage to enable improved capacity, performance and reliability