NAND Flash Architecture and Specification Trends

Michael Abraham (mabraham@micron.com)
NAND Solutions Group Architect
Micron Technology, Inc.

Abstract

- NAND Flash is quickly moving to sub-20nm lithographies, making it the fastest scaling semiconductor technology ever!
- What impact do these shrinks have to NAND's architecture, performance, and reliability in system solutions?
- Learn how to prepare for these changes and counteract some of them through improved system design.
- Also, take a look at innovative NAND technologies that improve performance and reliability.

Topics

- NAND Flash Architecture Trends
- The Cloud and Clients
- Enterprise Application Requirements
- ECC and SSD Topologies

NAND Process Migration: Shrinking Faster than Moore's Law

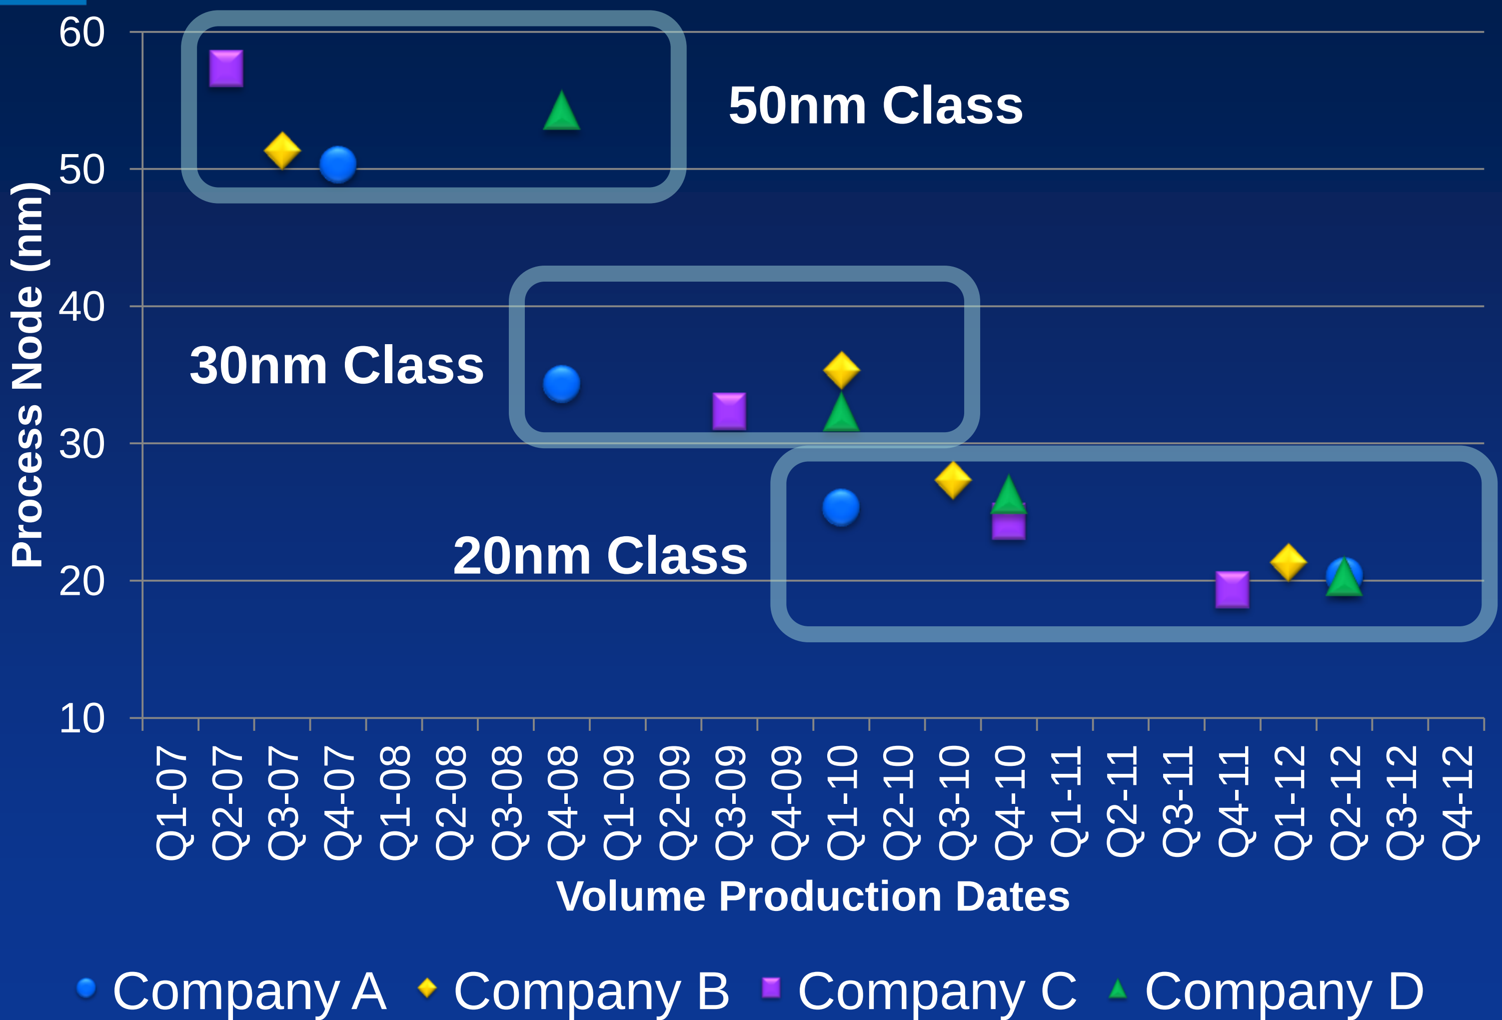

Data based on publicly available information

Memory Organization Trends

NAND block size is increasing

- Larger page sizes and more planes increase sequential throughput
- More pages per block reduce die size
- As ECC requirements increase, the spare area per NAND page is increasing

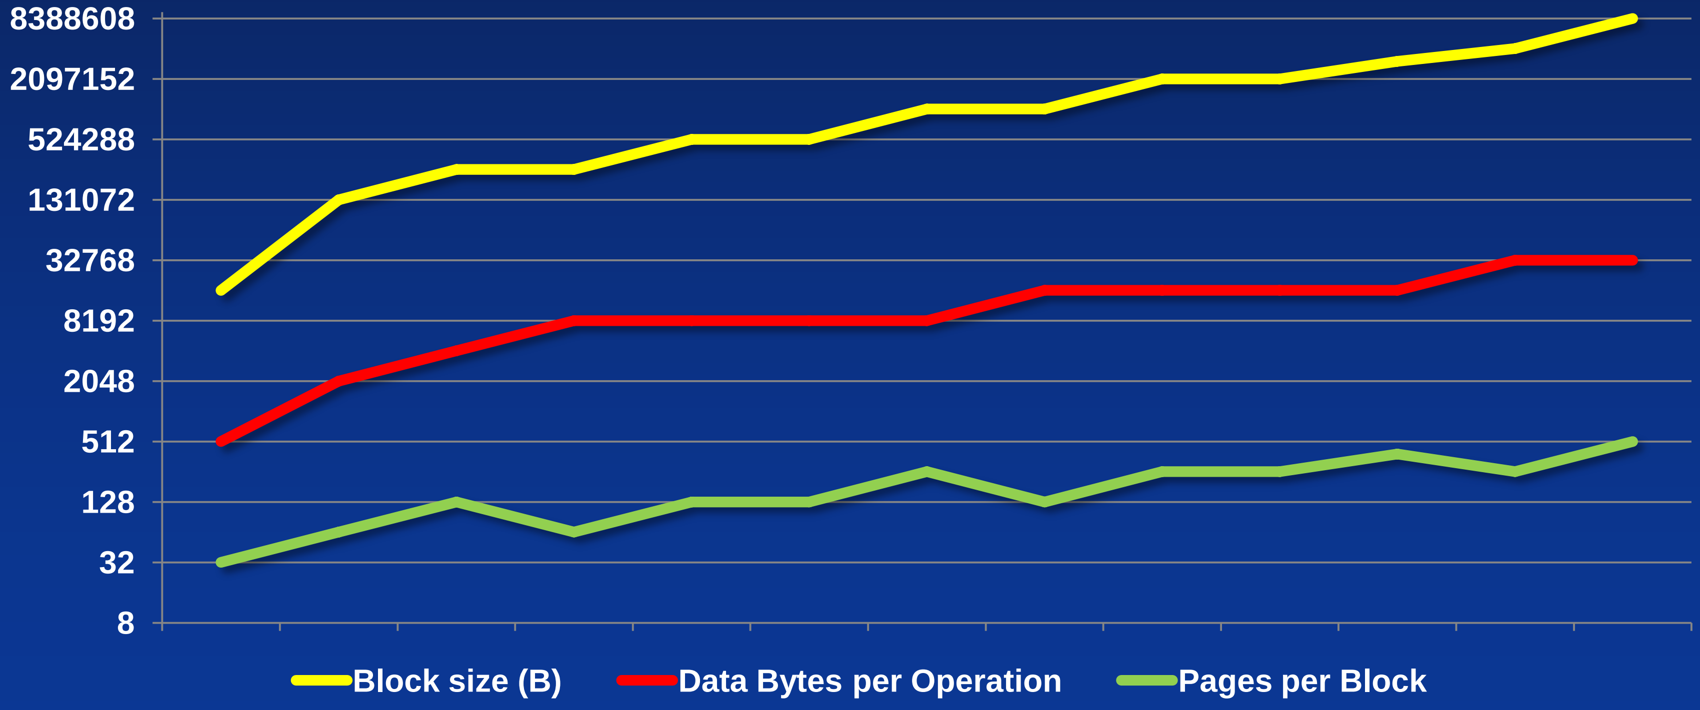

Consumer-grade NAND Flash: Endurance and ECC Trends

- ECC improves data retention and endurance
- Process shrinks lead to less electrons per floating gate
- To adjust for increasing RBERs, ECC is increasing exponentially to achieve equivalent UBERs
- ECC algorithms are transitioning from BCH to LDPC and codeword sizes are increasing

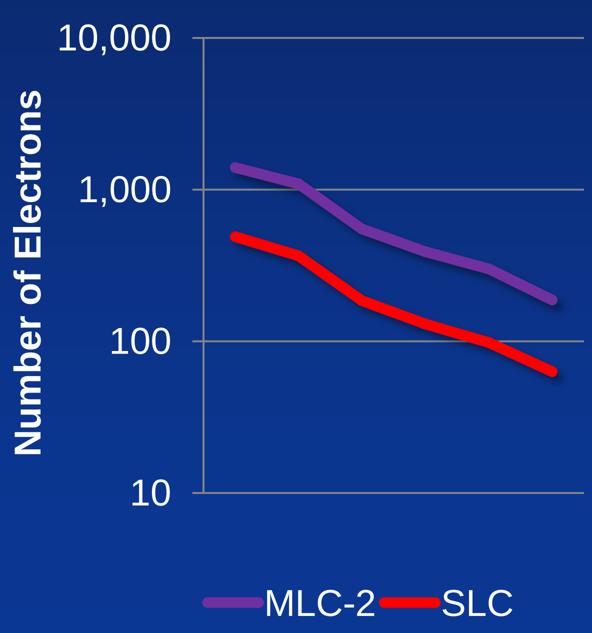

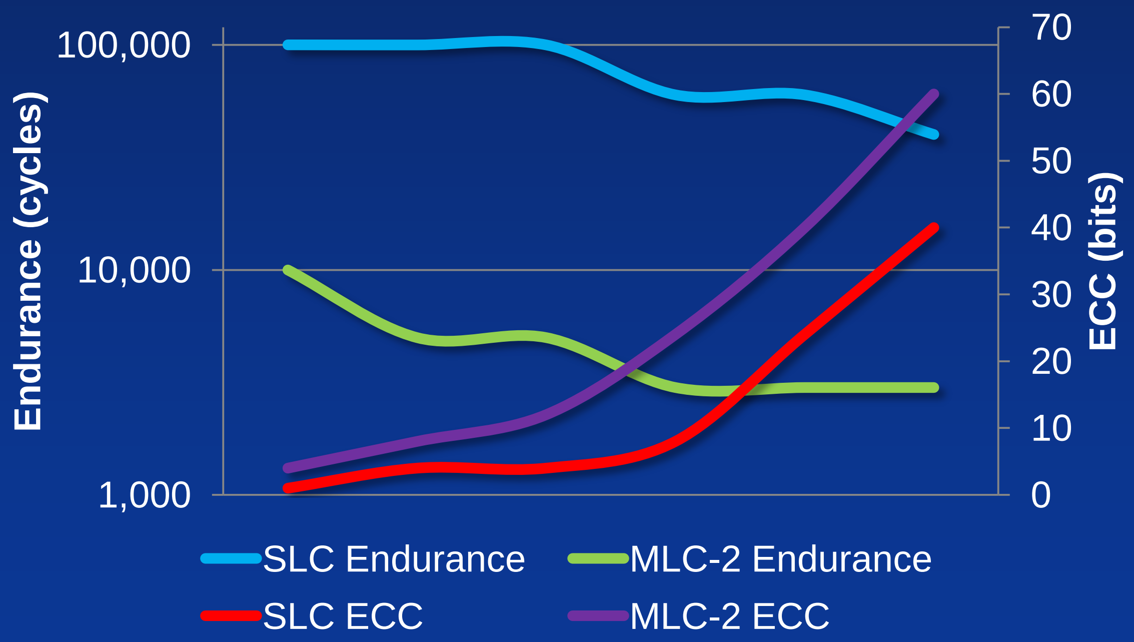

Larger Page Sizes Improve Sequential Write Performance

- For a fixed page size across process nodes, write throughput decreases as the NAND process shrinks
- NAND vendors increase the page size to compensate for slowing array performance
- Write throughput decreases with more bits per cell

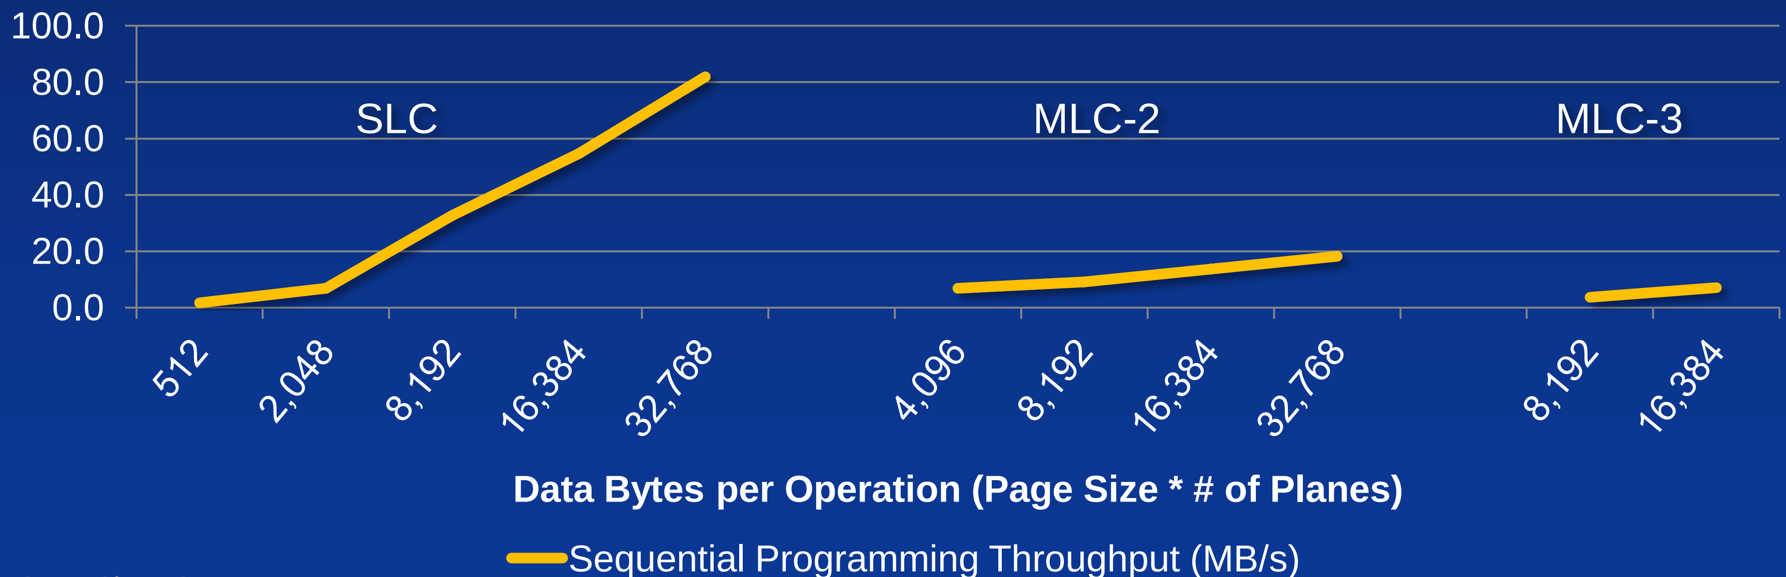

More Pages Per Block Affect Random Write Performance

- As block copy time increases, random performance decreases.
- Key factors that impact NAND Flash random write performance
  1. Number of pages per block
  2. Increase of tPROG
  3. Increase in I/O transfer time due to larger page sizes (effect not shown below)
- Impact to system product random performance
  - Some card interfaces have write timeout specs at 250ms.
  - To improve random performance, block management algorithms manage pages or partial blocks.

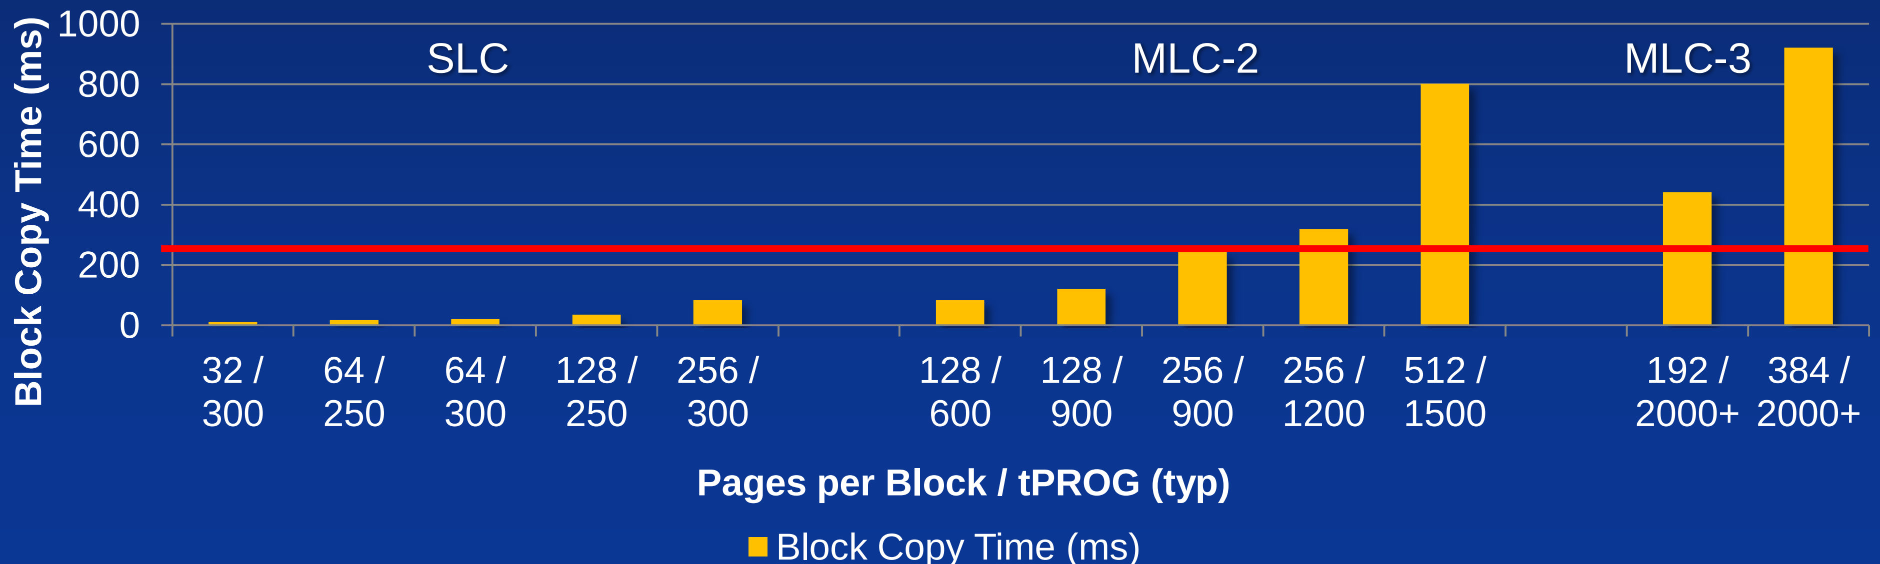

Larger Monolithic NAND Densities Increase Random Read Latencies

- Most applications favor read operations over write operations
- Most read operations are 4KB data sectors
- As monolithic NAND density increases, less NAND die are being used for a fixed system density
- As tPROG increases, the latency of random 4KB sector reads becomes more variable in mixed-operation environments as the probability of needing to read from a die that is busy increases

NAND Flash TAM by Density (Units)

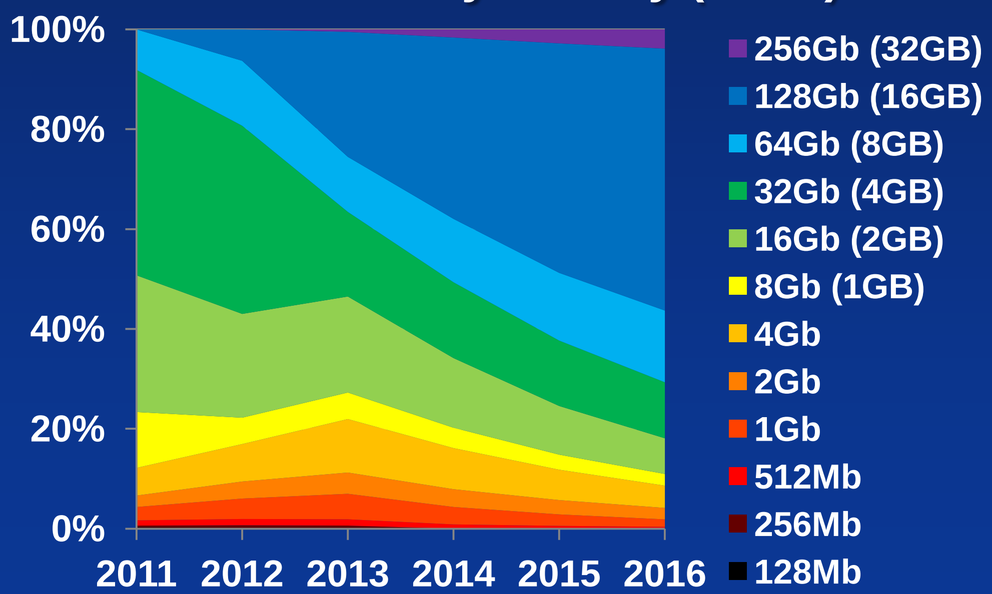

4KB Random Read Latency

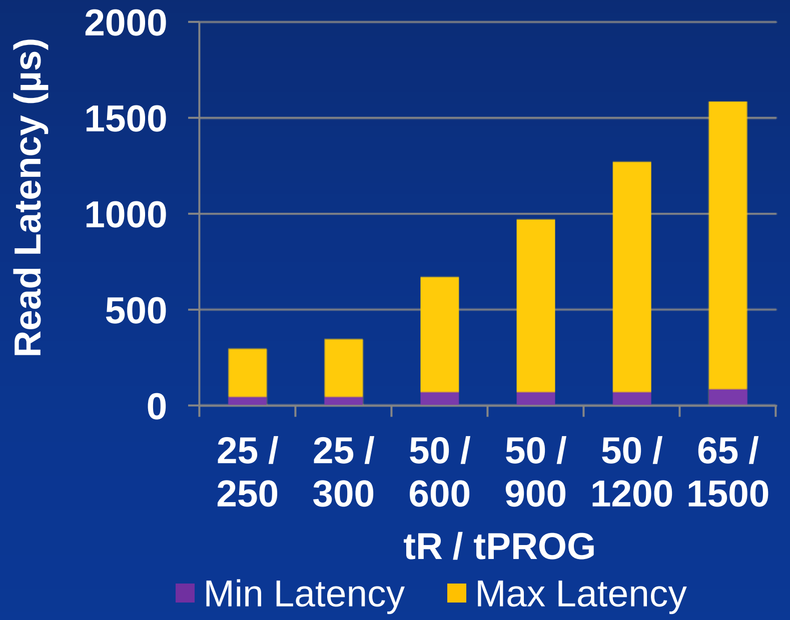

Source: iSuppli 2Q12

NAND Interface Trends for High-Performance Applications

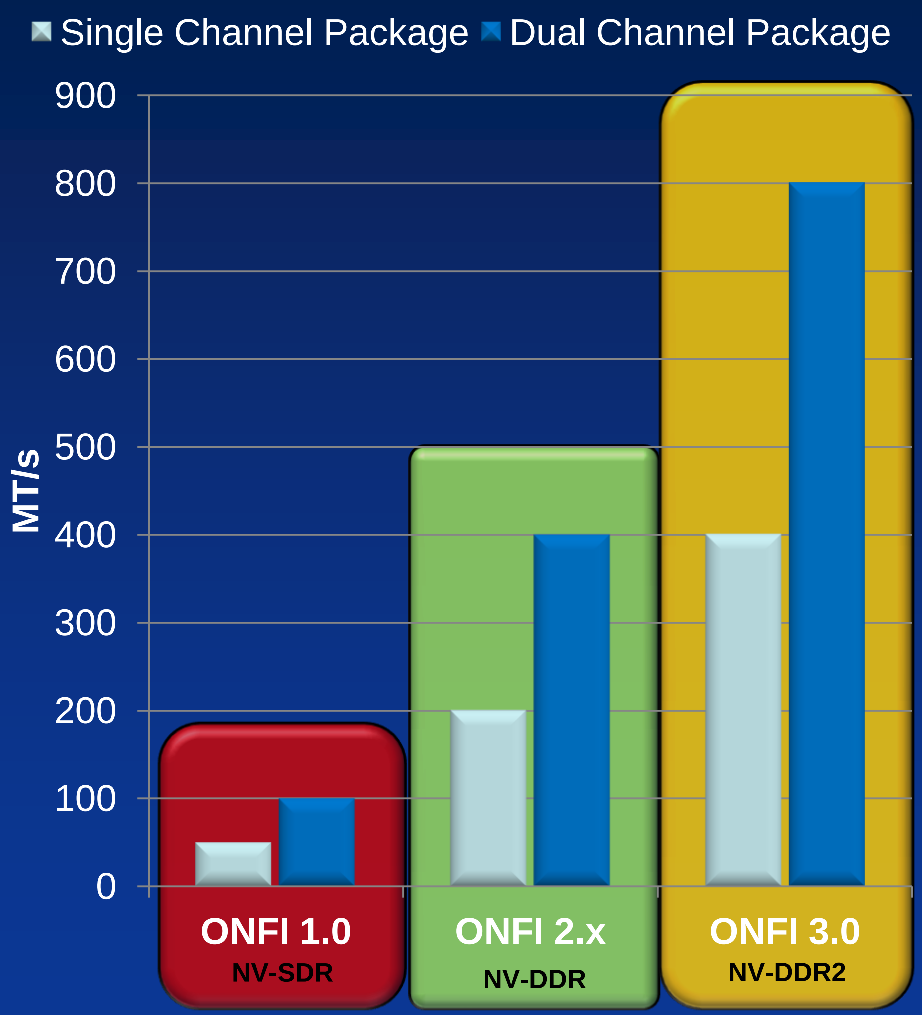

- Applications
  - Transitioned to 200MT/s interface
  - Shifting to 400MT/s interface
- Packaging
  - Typically BGA
  - 2 channel widely available
  - 4 channel being standardized
- ONFI 3.0 compatible components are available

The Cloud's Impact on NAND System Architectures

Storage Comparison

Client Storage

- Information stored locally, on the device
- Consumer or SSD grade NAND Flash

Cloud Storage

- Information stored in hosted server farms or data centers
- SLC or Enterprise grade NAND Flash

Comparison of NAND Flash by Application Requirement

Application Requirement	Client Storage / Consumer	Client Storage / SSD Grade	Cloud Storage / Enterprise Grade
NAND Cell	MLC-2 → MLC-3	MLC-2 → MLC-3	SLC → MLC-2
Endurance / Cycling	Up to 3K	Up to 3K	Up to 100K (SLC) Up to 30K (MLC-2)
DPM	Consumer grade	Better	Best
I/O Channel Throughput	40 → 200 MT/s	40 → 400 MT/s	133 MT/s → 400 MT/s
UBER	1E-14	Less	Less
Data retention at max cycling	1 year	1 year	Less
NAND Package Placements	Typically 1 to 4	4 to 16	Up to 32

How Do Enterprise Applications Meet Enterprise Requirements?

Application Requirement	Controller	SSD/Enterprise-grade NAND Flash
Higher system density	Ability to handle many NAND Flash – some controllers up to 256 die	Lower DPM
More throughput	Page-based block management, DRAM cache, Overprovisioning, More I/O channels, Faster I/O channels, Multiple ECC engines, Simultaneous, mixed operations	Faster I/O channel
Low latency reads	DRAM cache, Use smaller monolithic NAND densities	

How Do Enterprise Applications Meet Enterprise Requirements? (Part 2)

Application Requirement	Controller	SSD/Enterprise-grade NAND Flash
Higher endurance / reliability	Higher ECC	More ECC required, Lower UBER, Higher endurance
More consistent use over time	Balanced block management to reduce write amplification and provide even wear leveling so NAND die and blocks wear evenly	
Power within budget for parallel operations	Block management throttles parallelism as needed	Peak power reduction

ECC and Algorithms

- Enterprise-grade NAND requires more ECC than consumer-grade NAND Flash to achieve higher endurance and lower UBER
- Providing more ECC to a consumer-grade NAND Flash does not necessarily improve endurance, though it can improve data retention
- ECC requirements are going to increase to the point that it will be a significant amount of real estate on a multi-channel controller

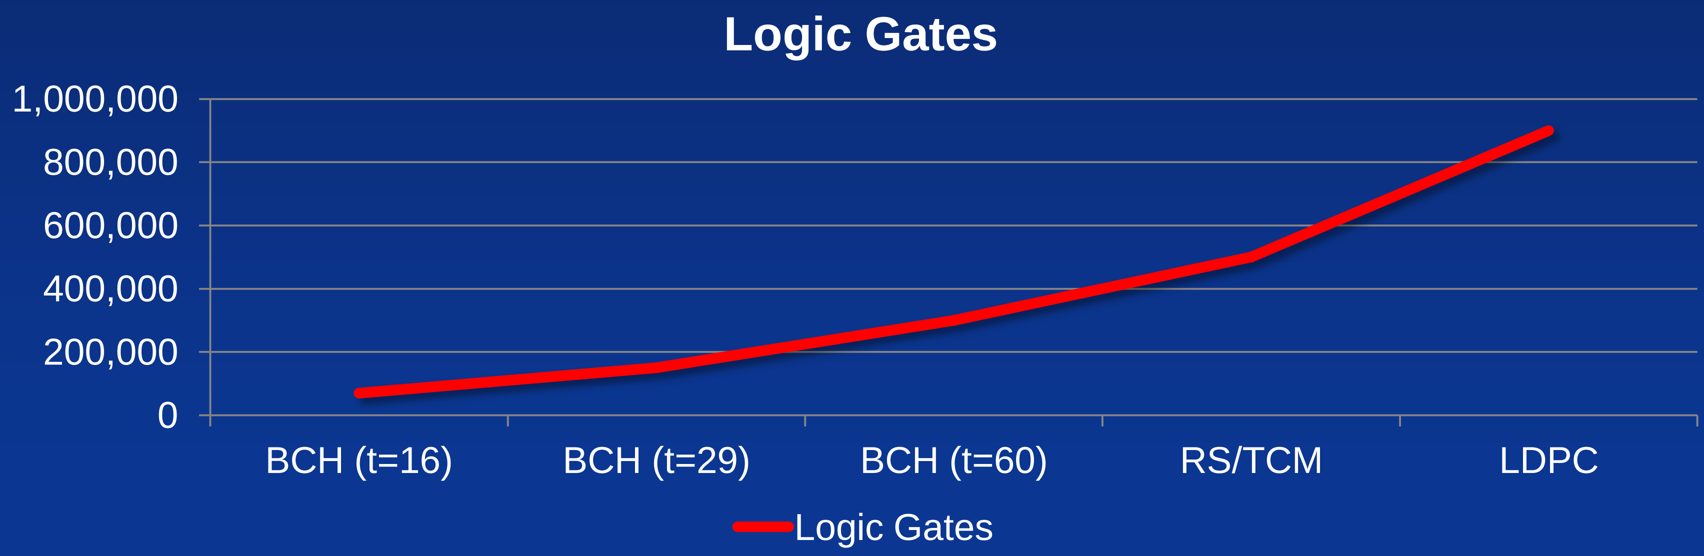

How to Handle Increasing ECC?

- ECC is NAND Flash technology dependent and is implemented in hardware
- Block management and drivers are not technology dependent and can be updated in software/firmware
- ECC Free Solution
  - Tightly couples ECC to the NAND technology
  - Also covers NAND aggregation, reducing channel loading
  - Block management performed in processor can use DRAM buffer and results in a higher performance than a fully managed solution

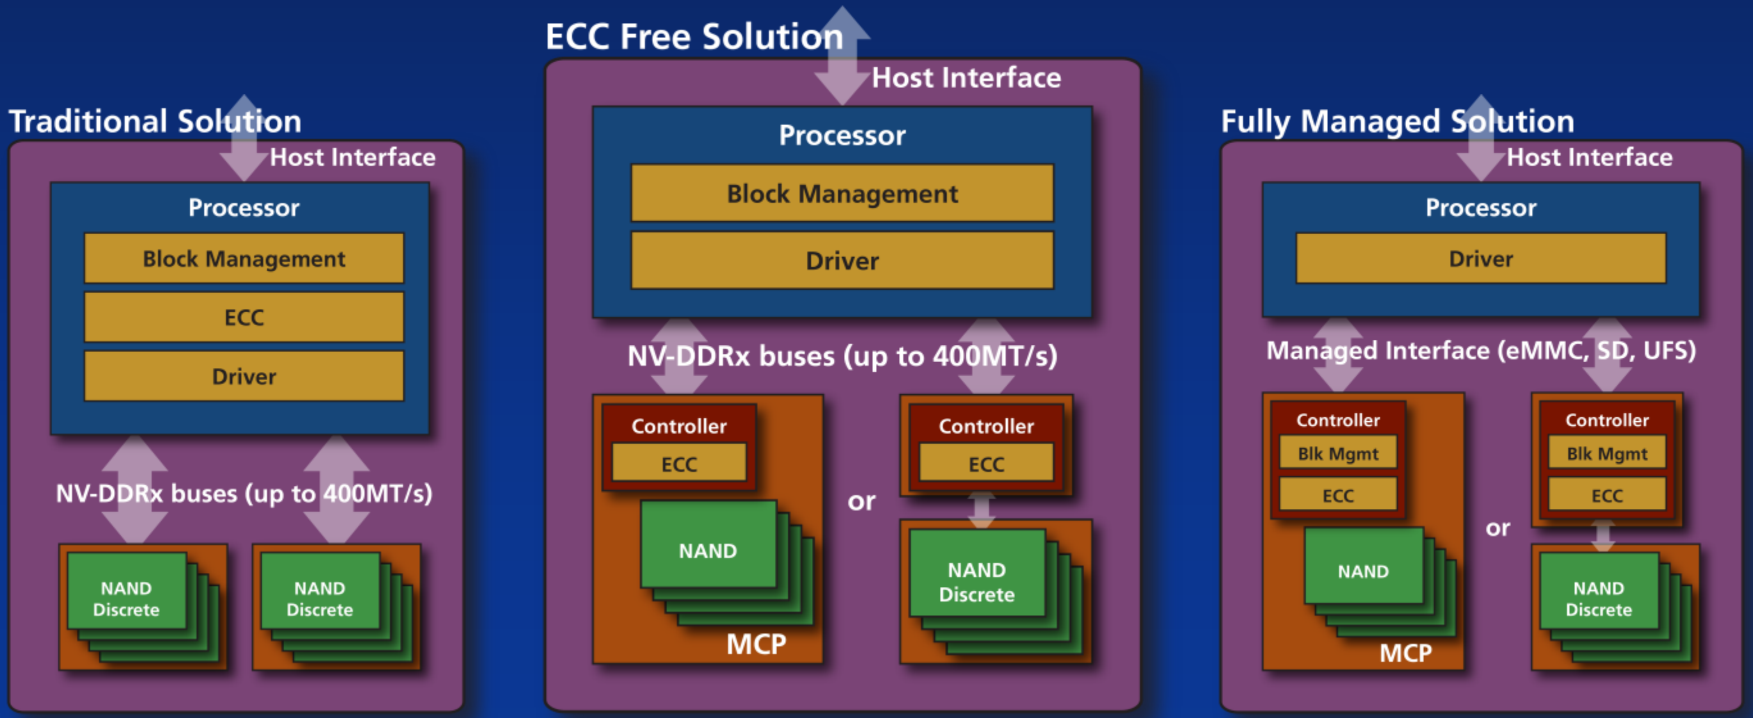

Questions?

Other Micron Presentations

- NAND uses in high performance platforms
  - Tutorial A-11 – Tuesday, August 21st @ 8:30 am
- NAND flash architecture and specification trends
  - Tutorial B-11 – Tuesday, August 21st @ 8:30 am
- MLC media discussion
  - Tutorial C-11 – Tuesday, August 21st @ 8:30 am
- Next-generation storage and the mobile computing ecosystem
  - Session 101-B – Tuesday, August 21st @ 8:30 am
- Why ECC-free NAND is the best solution for high-performance applications
  - Session 102-A – Tuesday, August 21st @ 10:10 am
- How ONFI standards are fueling high-performance SSDs
  - Session 102-C – Tuesday, August 21st @ 10:10 am
- The need for differentiated MLC solutions
  - Tutorial F-21 – Wednesday, August 22nd @ 8:30 am
- Virtualized SSD storage for enterprise systems
  - Tutorial H-22 – Wednesday, August 22nd @ 4:30 pm
- Performance trade-offs of flash-based client storage solutions
  - Tutorial A-31 – Thursday, August 23rd @ 8:30 am
- Phase Change Memory – Panel Discussion
  - Session 302-D – Thursday, August 23rd @ 9:50 am
- 2.5-inch PCIe interface for enterprise flash cache – Panel Discussion
  - Session 303-B – Thursday, August 23rd @ 3:10 pm

About Michael Abraham

- Architect in the NAND Solutions Group at Micron
- Covers advanced NAND and PCM interfaces and system solutions
- IEEE Senior Member
- BS degree in Computer Engineering from Brigham Young University