An ONFi Update

OPEN NAND
FLASH INTERFACE

Jim Cooke
NAND Marketing
Micron Technology, Inc.

Outline

- *ONFi Overview*
- ONFi Provides Speed Enhancements
- The Path to Higher Performance

Why ONFi ?

- NAND was the only commodity memory with no standard interface
- The Open NAND Flash Interface (ONFi) Workgroup was formed in May 2006 to drive standardization for the raw NAND Flash interface
- ONFi has been busy…

Why ONFi ?

- NAND was the only commodity memory with no standard interface
- The Open NAND Flash Interface (ONFi) Workgroup was formed in May 2006 to drive standardization for the raw NAND Flash interface
- ONFi has been busy…

ONFi 1.0 *(Dec '06)*

Defined a standard electrical and protocol interface for NAND, including the base command set.

Why ONFi ?

- NAND was the only commodity memory with no standard interface
- The Open NAND Flash Interface (ONFi) Workgroup was formed in May 2006 to drive standardization for the raw NAND Flash interface
- ONFi has been busy…

Block Abstracted NAND 1.0 *(Jul '07)*

Defined a managed NAND solution that utilizes the raw NAND interface.

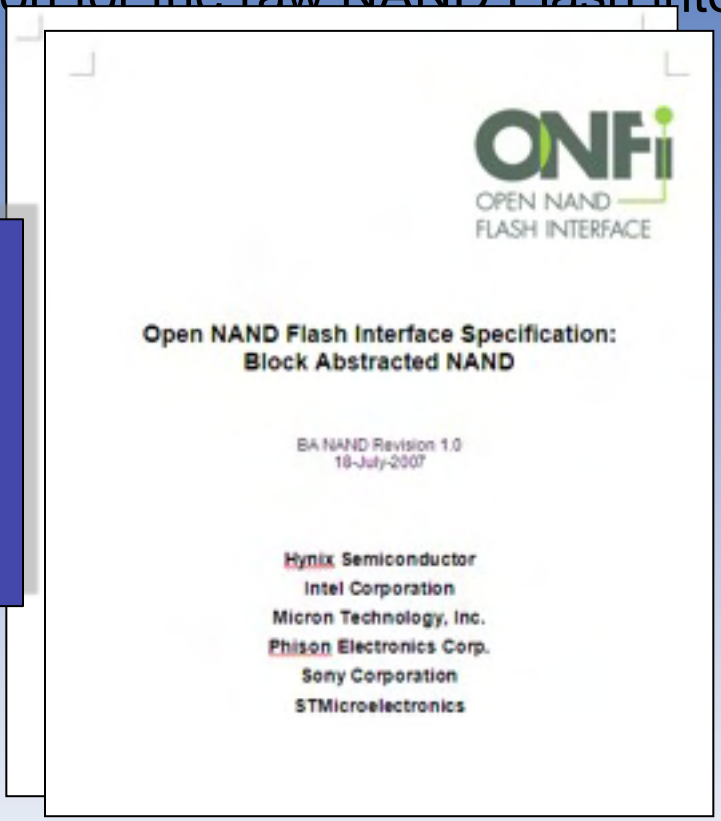

Why ONFi ?

- NAND was the only commodity memory with no standard interface
- The Open NAND Flash Interface (ONFi) Workgroup was formed in May 2006 to drive standardization for the raw NAND Flash interface
- ONFi has been busy…

ONFi 2.0 *(Feb '08)*

Defined a high speed DDR interface, tripling the traditional NAND bus speed from 40 MB/s to 133 MB/s.

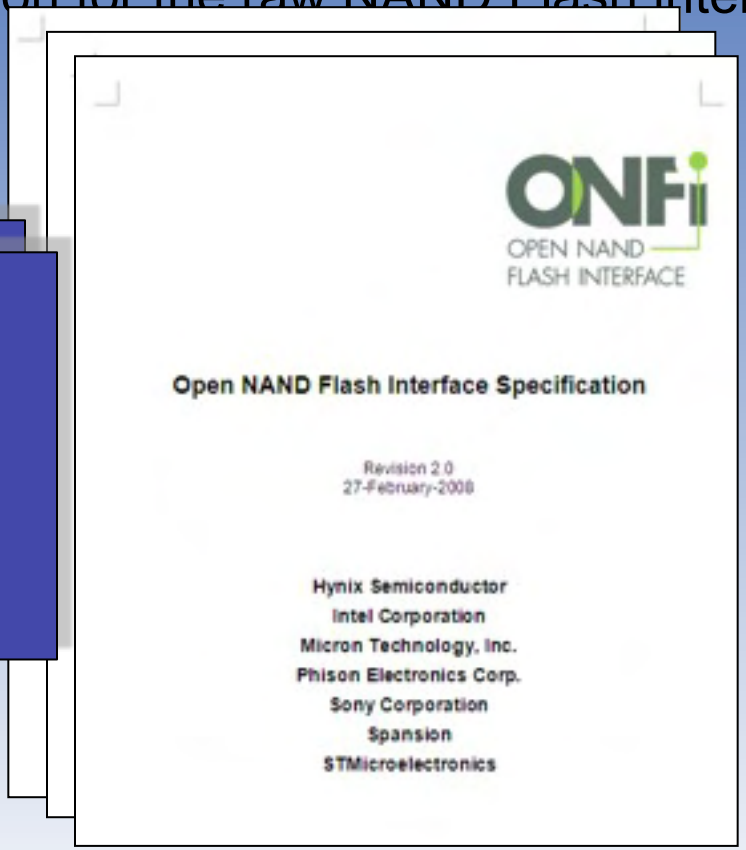

Why ONFi ?

- NAND was the only commodity memory with no standard interface
- The Open NAND Flash Interface (ONFi) Workgroup was formed in May 2006 to drive standardization for the raw NAND Flash interface
- ONFi has been busy…

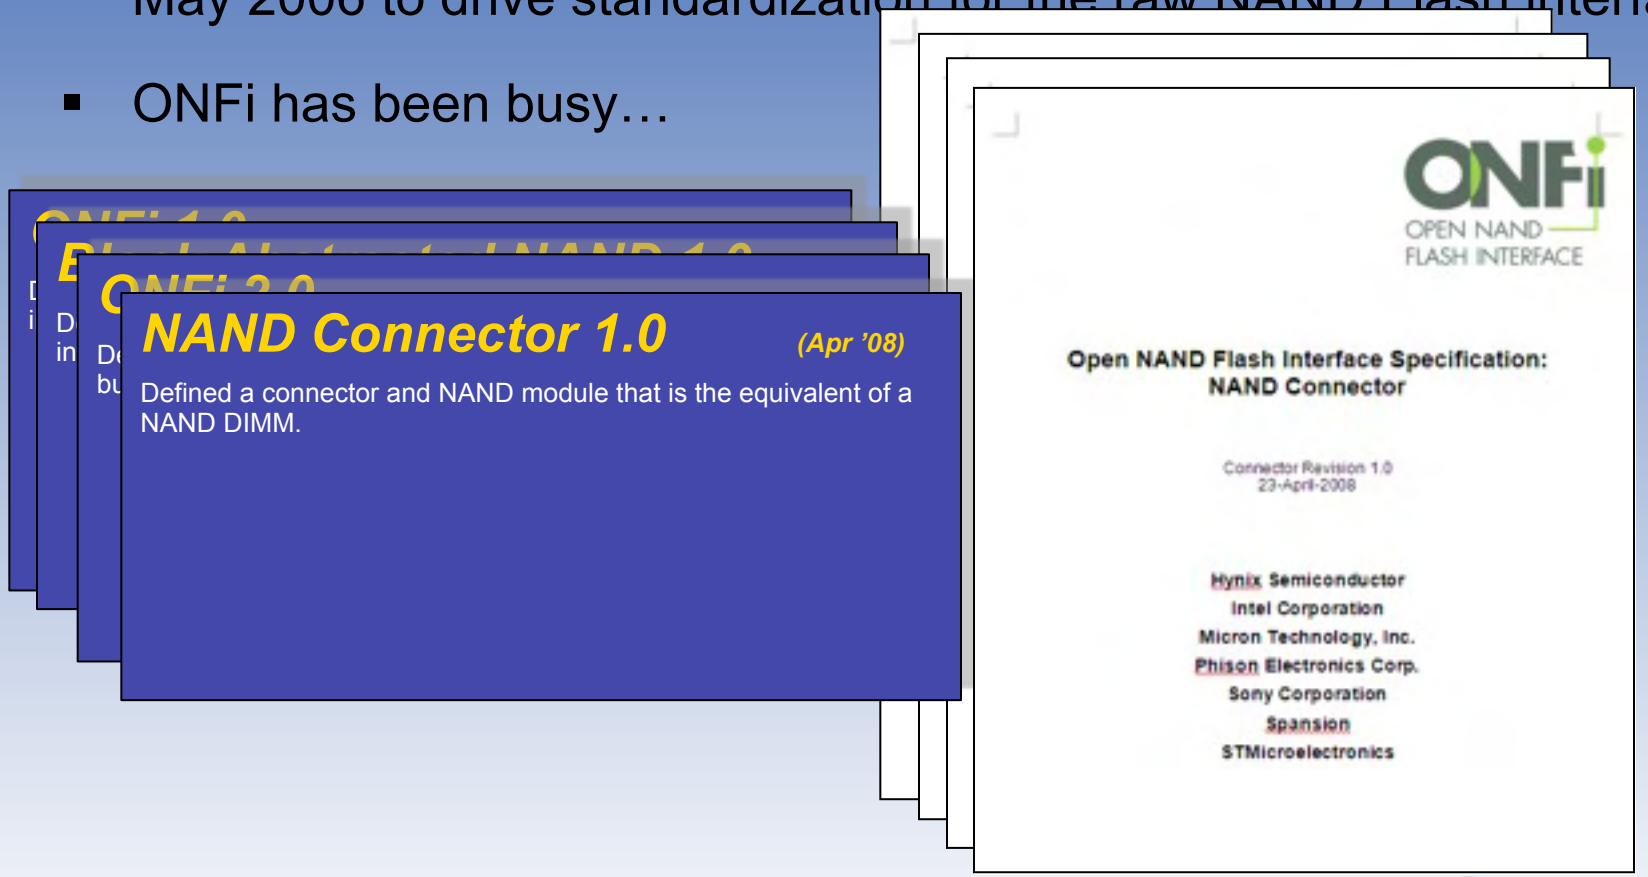

Why ONFi ?

- NAND was the only commodity memory with no standard interface
- The Open NAND Flash Interface (ONFi) Workgroup was formed in May 2006 to drive standardization for the raw NAND Flash interface
- ONFi has been busy…

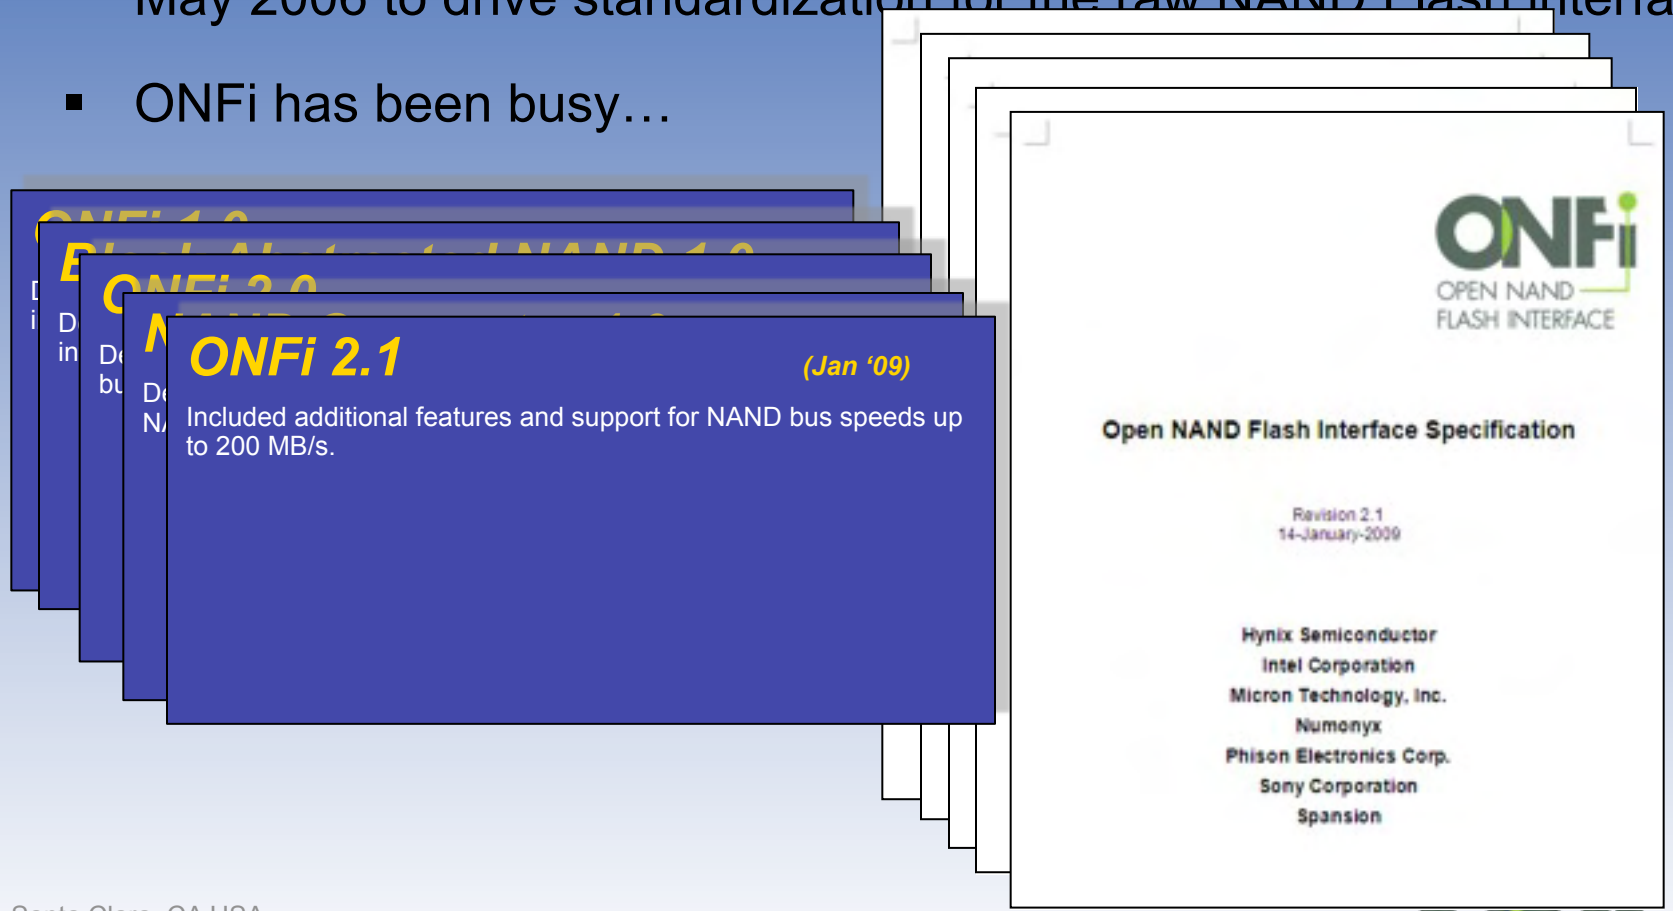

Why ONFi ?

- NAND was the only commodity memory with no standard interface
- The Open NAND Flash Interface (ONFi) Workgroup was formed in May 2006 to drive standardization for the raw NAND Flash interface
- ONFi has been busy…

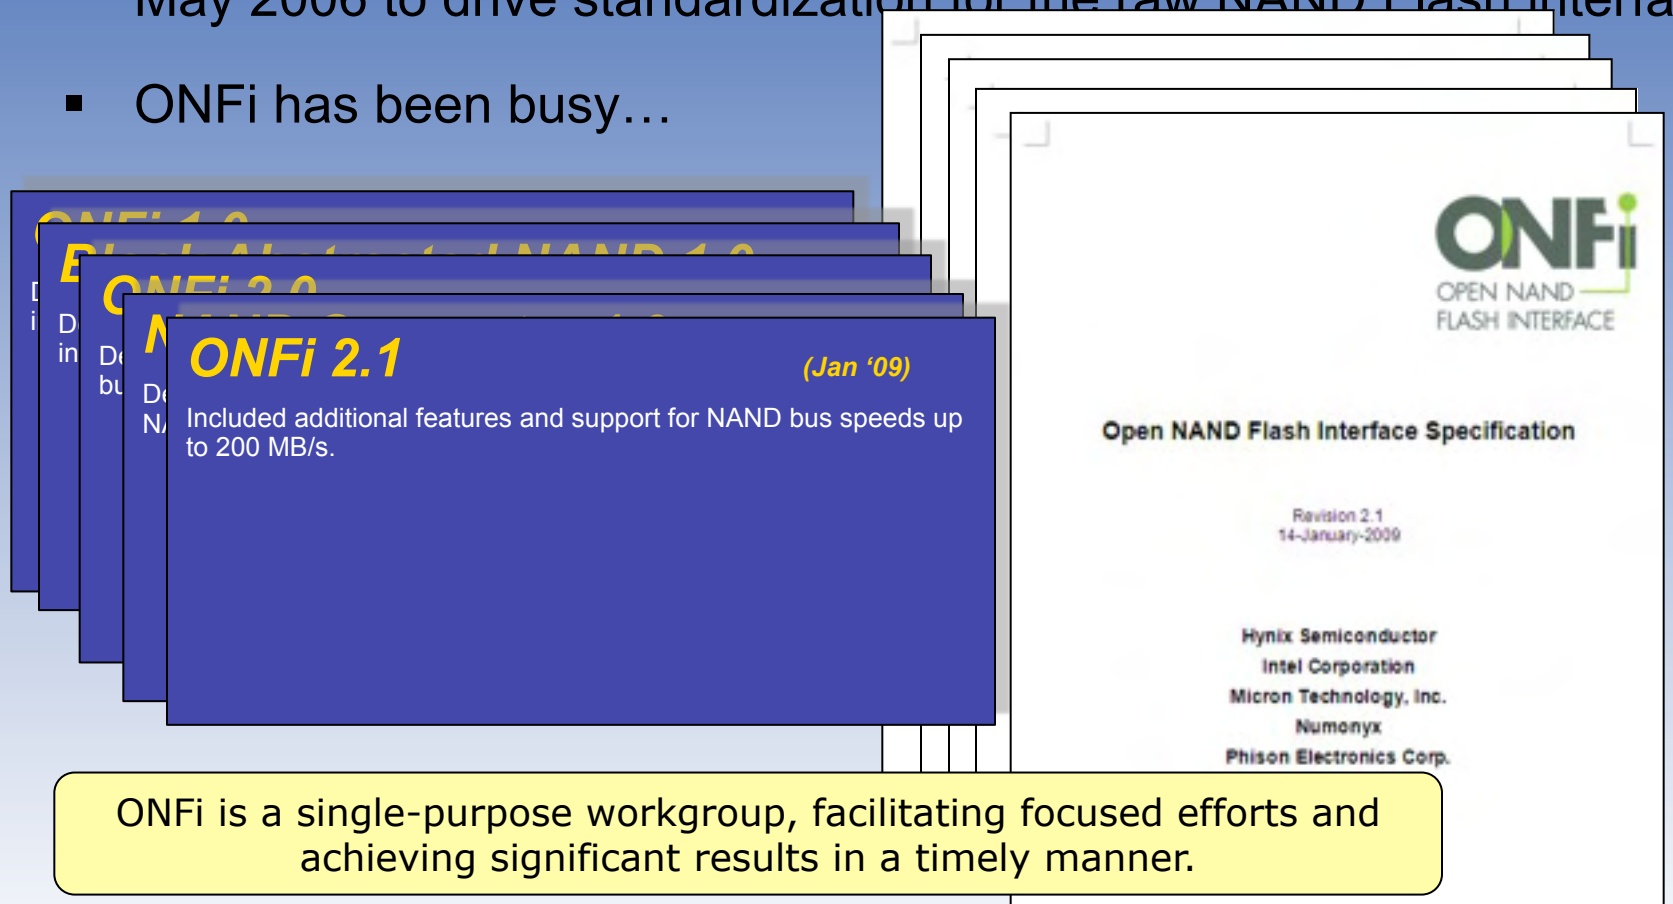

ONFi is a single-purpose workgroup, facilitating focused efforts and achieving significant results in a timely manner.

ONFi
OPEN NAND FLASH INTERFACE

hynix SONY numonyx Micron intel SanDisk PHISON SPANSION

ONFi
Members

A-Data
Aleph One
Arasan Chip Systems
Avid Electronics
Chipsbank
Data I/O
Software
Entorian
Foxconn
Genesys Logic
Hitachi GST
Indilinx
Jinvani Systech
Lotes
Macronix
Metaram
NVidia
Powerchip Semi.
Qimonda
Shenzhen Netcom
Silicon Motion
STEC
Solid State System
Tandon
Teradyne, Inc.
Tyco
Virident Systems

Afa Technologies
Anobit Tech.
ASMedia Technology
BitMicro
Cypress
Datalight
FCI
Fresco Logic
Hagiwara Sys-Com
Hyperstone
Inphi
Kingston Technology
Marvell
Moai Electronics
Orient Semiconductor
Power Quotient International
Sandforce
Sigmatel
Silicon Storage Tech
Skymedi
Super Talent Electronics
Tanisys
Testmetrix
UCA Technology
WinBond

Alcor Micro
Apacer
ATI
Biwin Technology
DataFab Systems
Denali
ENE Technology
FormFactor
Fusion Media Tech
HiperSem
InCOMM
Intelliprop
ITE Tech
Lauron Technologies
LSI
Mentor Graphics
Molex
P.A. Semi
Prolific Technology
Seagate
Silicon Integrated Systems
Silicon Systems
Smart Modular Tech.
Synopsys
Telechips
Transcend Information
University of York

Continued Innovation: ONFi 2.1

- ONFi 2.1 was ratified in January 2009 and contains a plethora of new features
- New capabilities include:
  - 166 MB/s and 200 MB/s speeds
  - Power management features
  - Enhanced ECC information
  - New commands for increased performance and functionality
    - Small data move
    - Change row address

ONFI 2.0 tripled the legacy NAND interface speed.
ONFI 2.1 delivers a ***5X*** speed boost over legacy NAND!

Adding Speed to Block Abstracted NAND

- Block Abstracted NAND is a managed solution, using the existing NAND bus
- BA NAND revision 1.1 adds the high-speed interface capabilities of ONFi 2.1 to this solution
  - Enables 200 MB/s performance, per 8-bit NAND channel
- BA NAND revision 1.1 ratified in July

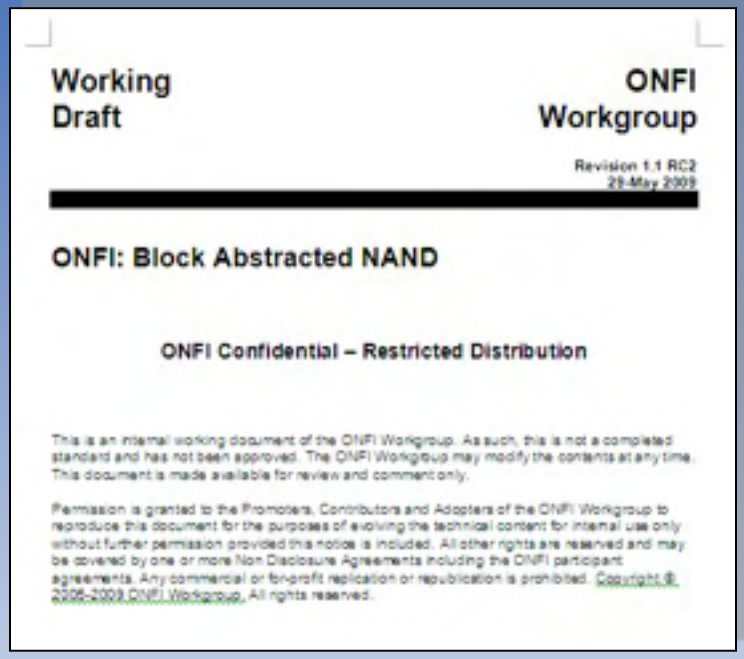
Working Draft

ONFI Workgroup

Revision 1.1 RC2
29-May-2009

ONFI: Block Abstracted NAND

ONFI Confidential – Restricted Distribution

This is an internal working document of the ONFI Workgroup. As such, this is not a completed standard and has not been approved. The ONFI Workgroup may modify the contents at any time. This document is made available for review and comment only.

Permission is granted to the Promoters, Contributors and Adopters of the ONFI Workgroup to reproduce this document for the purposes of evolving the technical content for internal use only without further permission provided this notice is included. All other rights are reserved and may be covered by one or more Non Disclosure Agreements including the ONFI participant agreements. Any commercial or for-profit replication or republication is prohibited. Copyright © 2006-2009 ONFI Workgroup. All rights reserved.

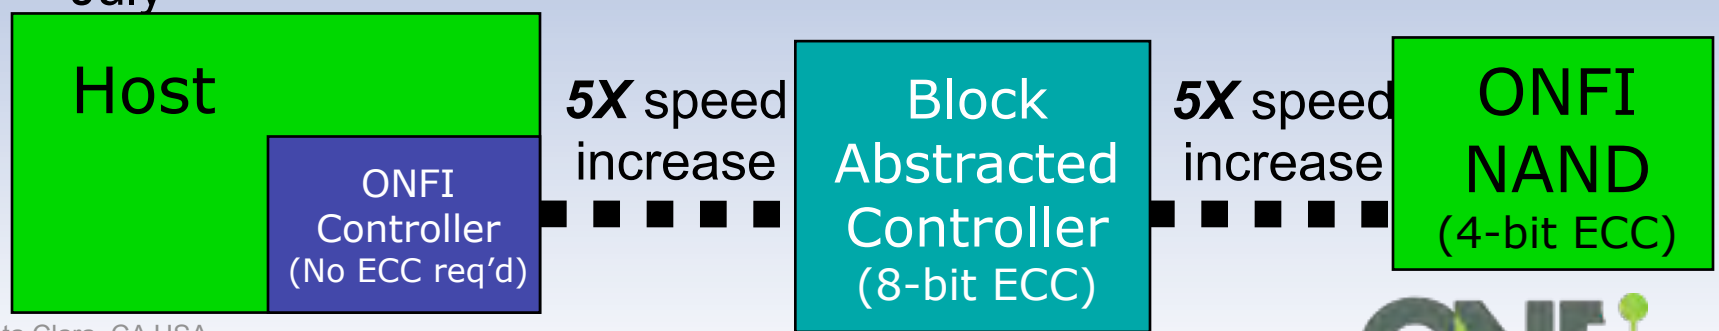

JEDEC and ONFi Are Collaborating

- JEDEC and ONFi are collaborating to increase the global reach of NAND standardization
- Any JEDEC member may participate in the ONFi-JEDEC Joint Taskgroup
  - The ONFi 2.1 specification has been contributed to the joint work

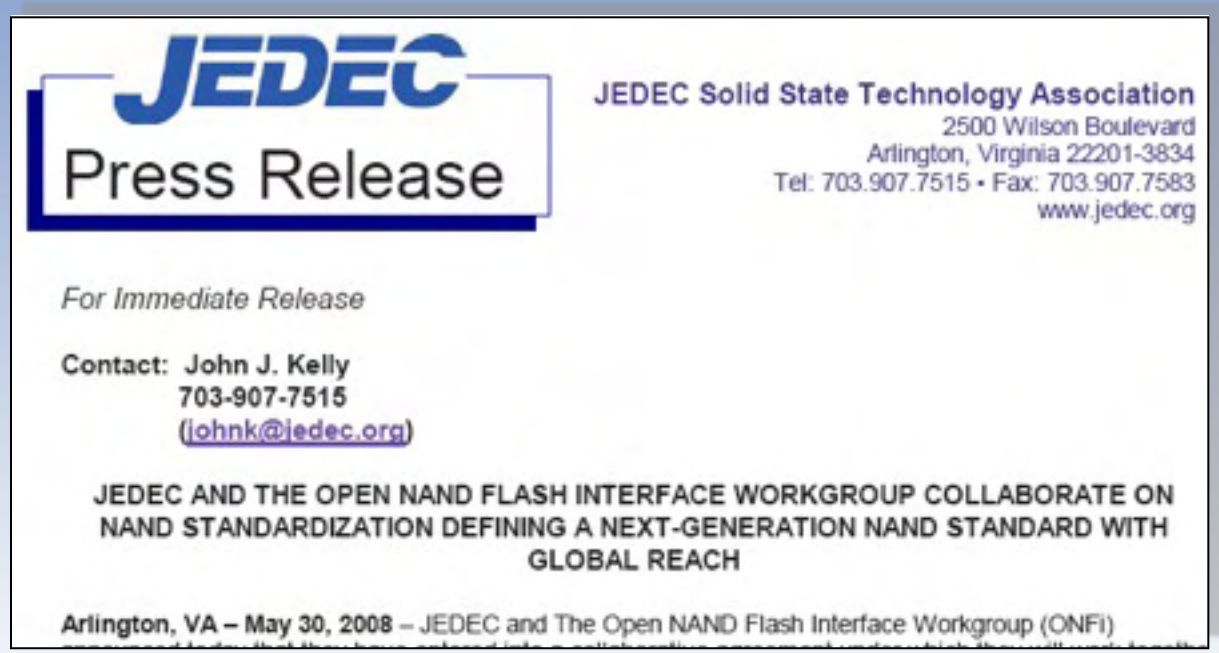

JEDEC

Press Release

JEDEC Solid State Technology Association
2500 Wilson Boulevard
Arlington, Virginia 22201-3834
Tel: 703.907.7515 • Fax: 703.907.7583
www.jedec.org

For Immediate Release

Contact: John J. Kelly
703-907-7515
(johnk@jedec.org)

JEDEC AND THE OPEN NAND FLASH INTERFACE WORKGROUP COLLABORATE ON NAND STANDARDIZATION DEFINING A NEXT-GENERATION NAND STANDARD WITH GLOBAL REACH

Arlington, VA – May 30, 2008 – JEDEC and The Open NAND Flash Interface Workgroup (ONFi)

Outline

- ONFi Overview
- *ONFi Provides Speed Enhancements*
- The Path to Higher Performance

Problem: Latency Increases as Lithography Shrinks

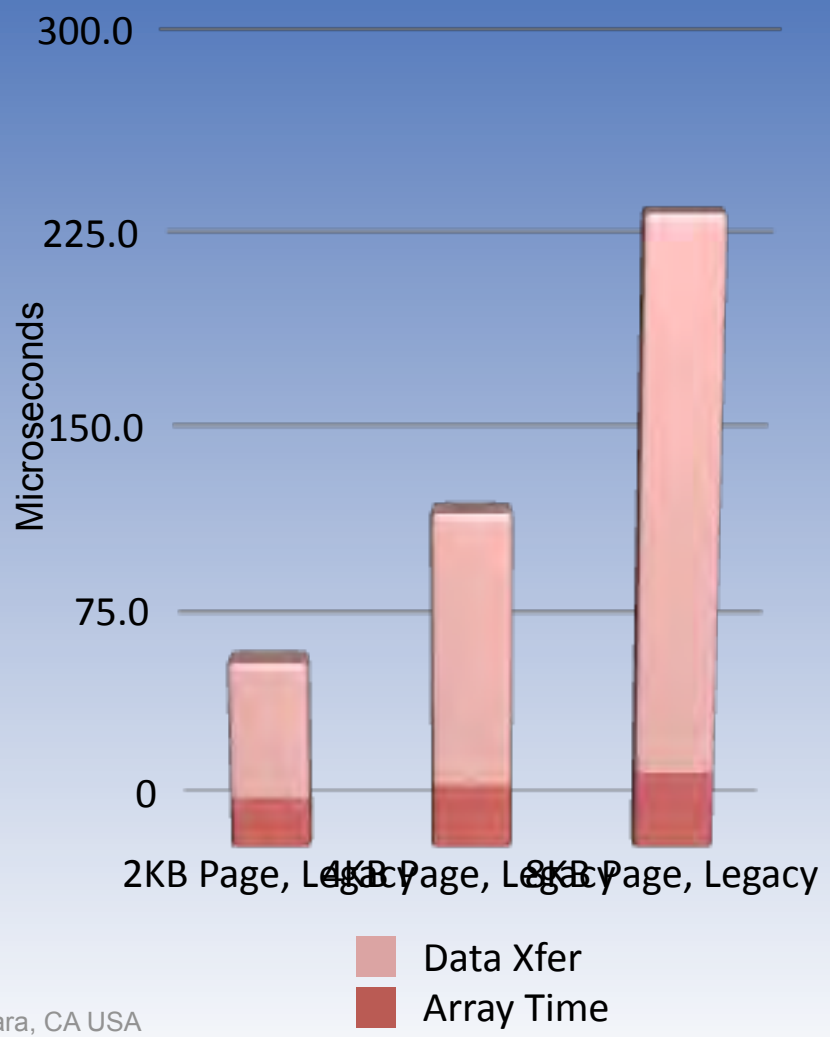

- NAND component performance is determined by two elements:
  - NAND array access time
  - Data transfer across the bus
- For legacy reads, performance is artificially limited to 40 MB/s
- As NAND page sizes increase, latency becomes large
  - Especially for small reads

Solution: Continually Scale the NAND Bus

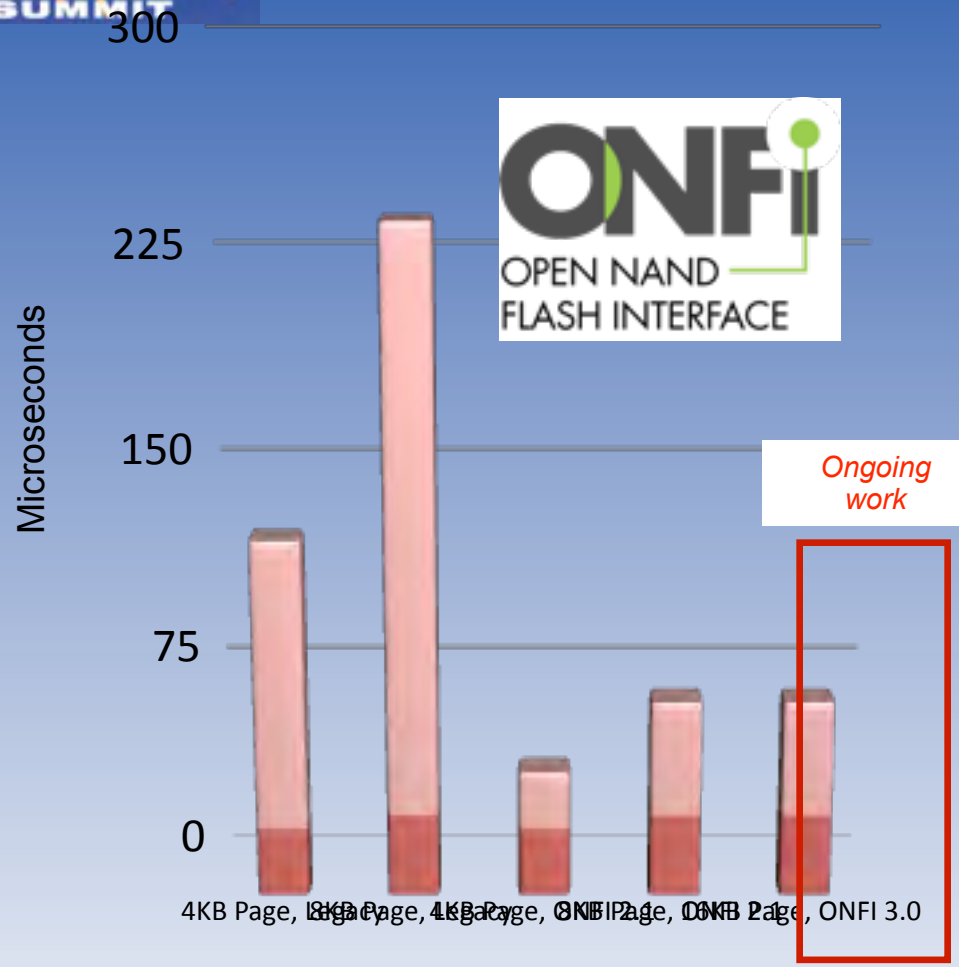

- The ONFI source synchronous interface delivers 200 MB/s in performance
- Reads balanced between array and data transfer
  - Pipelined reads enable NAND array times to be “hidden”
  - 200 MB/s is excellent fit for 8KB NAND pages
- What about the future?
  - For 16KB NAND pages, we need to continue scaling…
- ONFI 3.0: Target is 400 MB/s

ONFI provides the high-speed interface needed for SSD and cache designs. Work to reach 400 MB/s is underway.

High-Speed ONFi in Action

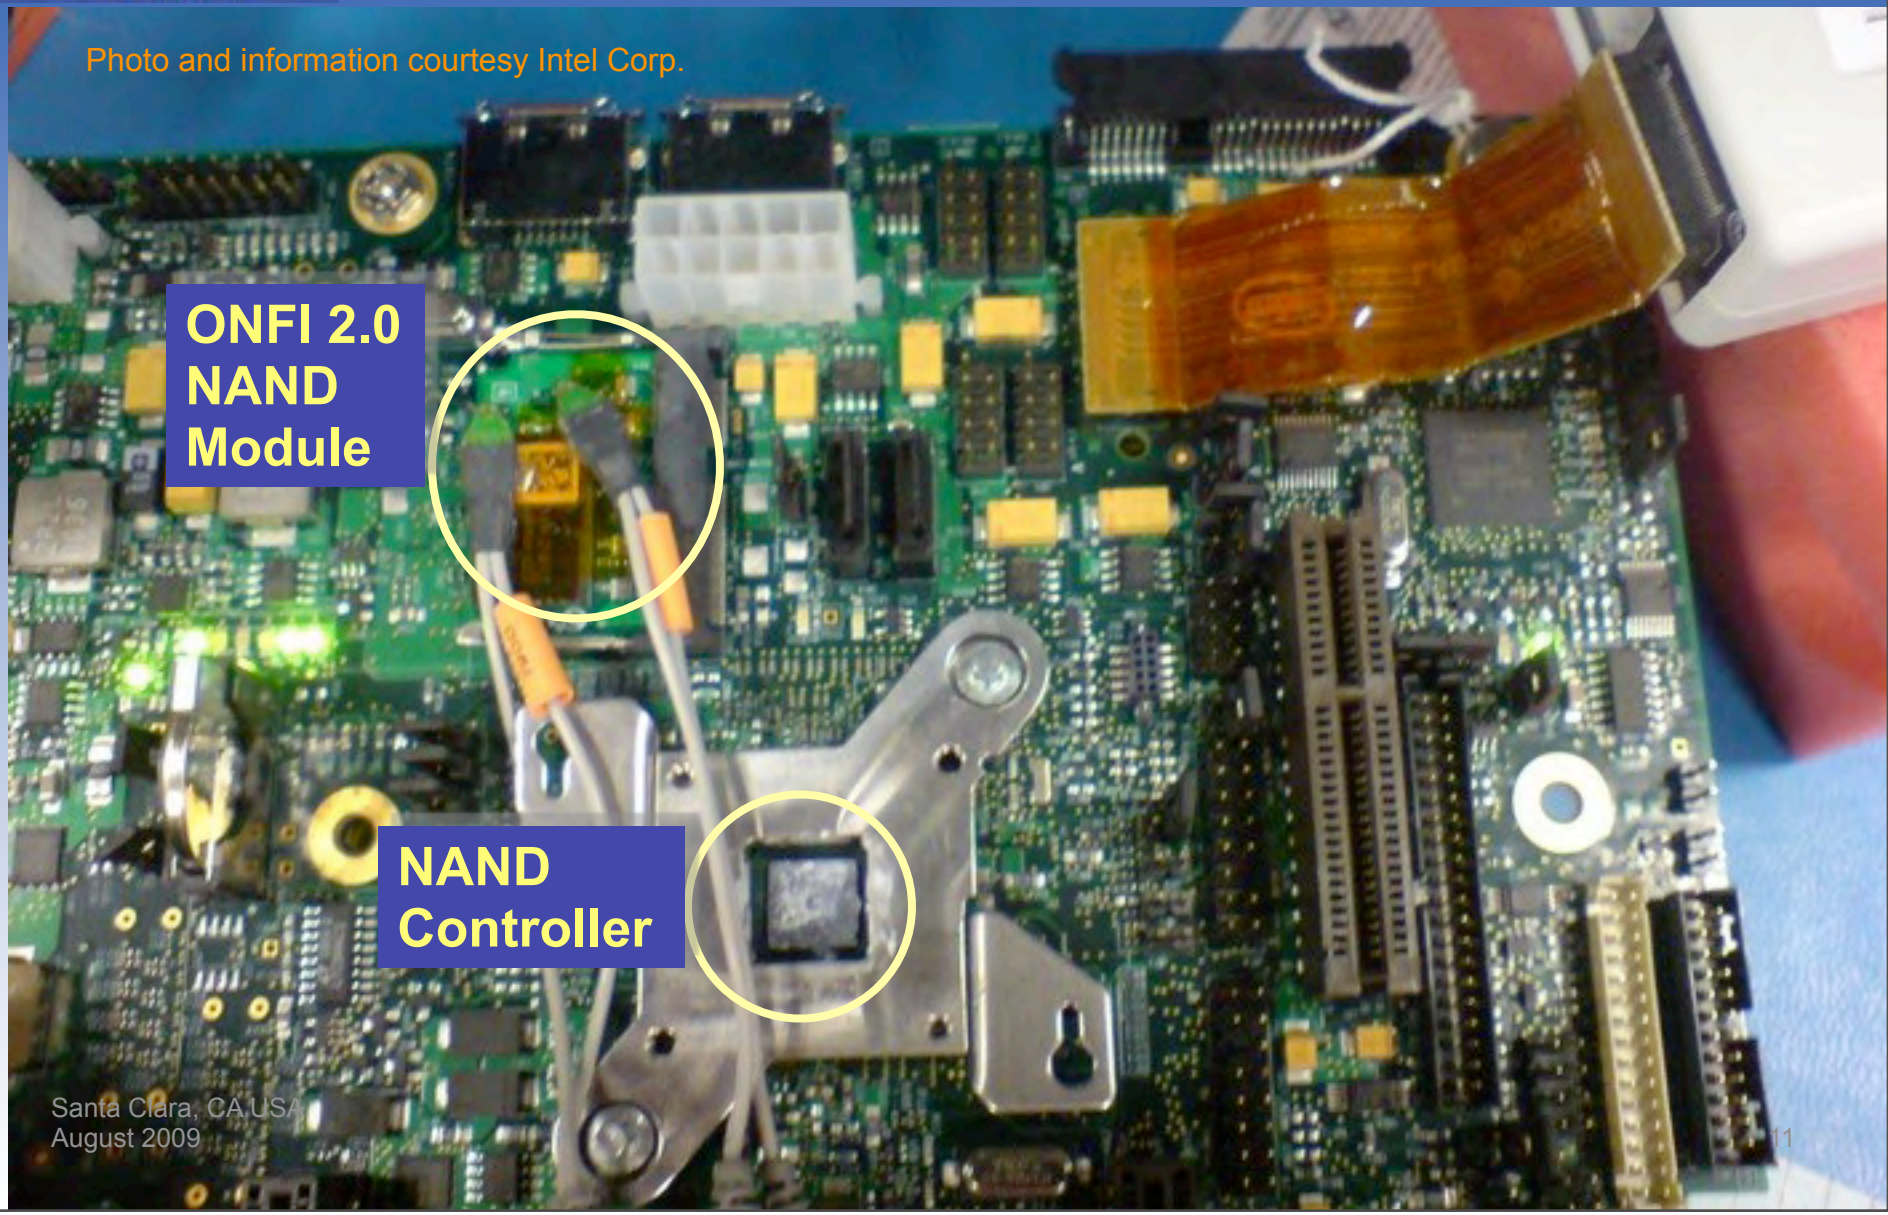

High-Speed ONFi in Action

Photo and information courtesy Intel Corp.

ONFI 2.0 NAND Module

NAND Controller

Excellent Data Eyes

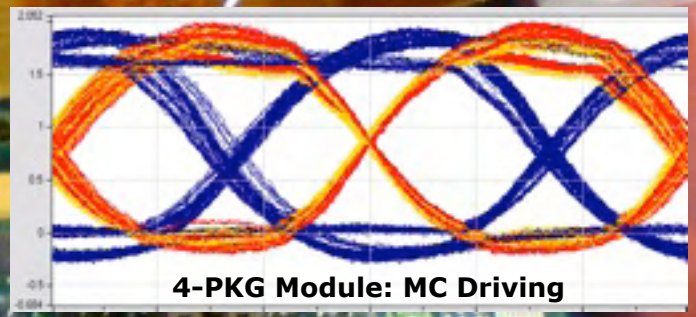

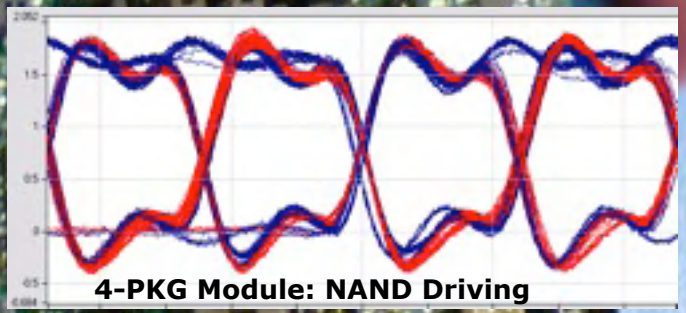

High-Speed ONFi in Action

Photo and information courtesy Intel Corp.

ONFI 2.0 NAND Module

NAND Controller

Excellent Data Eyes

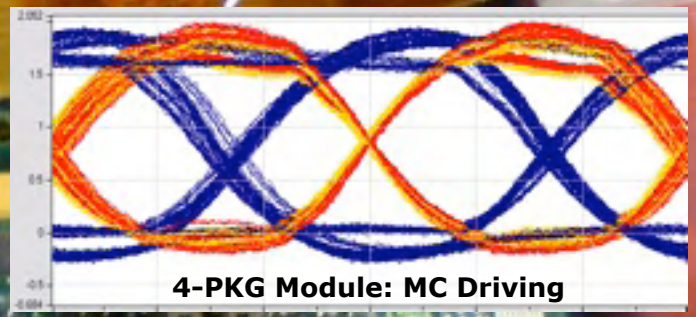

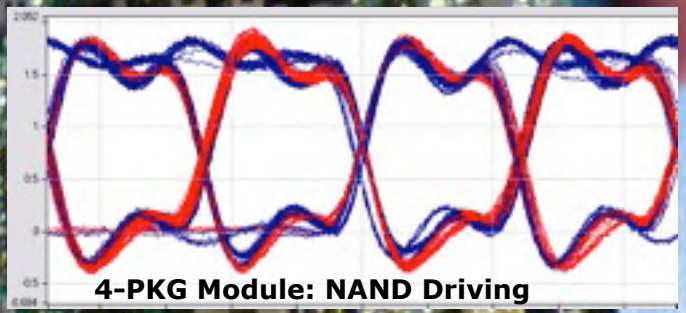

Successful operation utilizing the ONFi NAND module is achieved at 166 MT/s.

ONFi 2.1: The Speed You Need

- ONFi defines speeds ranging from 20 MB/s to 200 MB/s, to allow applications to balance performance and power
- The new speeds defined in ONFi 2.1 are 166 MB/s (mode 4) and 200 MB/s (mode 5)

Parameter	Mode 0		Mode 1		Mode 2		Mode 3		Mode 4		Mode 5		Unit
	50		30		20		15		12		10		ns
	~20		~33		~50		~66		~83		~100		MHz
	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	Min	Max	
tAC	—	20	—	20	—	20	—	20	—	20	—	20	ns
tADL	100	—	100	—	70	—	70	—	70	—	70	—	ns
tCADf	25	—	25	—	25	—	25	—	25	—	25	—	ns
tCADs	45	—	45	—	45	—	45	—	45	—	45	—	ns
tCAH	10	—	5	—	4	—	3	—	2.5	—	2	—	ns
tCALH	10	—	5	—	4	—	3	—	2.5	—	2	—	ns
tCALS	10	—	5	—	4	—	3	—	2.5	—	2	—	ns
tCAS	10	—	5	—	4	—	3	—	2.5	—	2	—	ns
tCH	10	—	5	—	4	—	3	—	2.5	—	2	—	ns
tCK(avg) or tCK	50	—	30	—	20	—	15	—	12	—	10	—	ns
...													

High-Density Scalability

- By providing multiple output drive strength settings, many NAND devices can share the I/O bus while maintaining I/O throughput
- DDR200 shown below

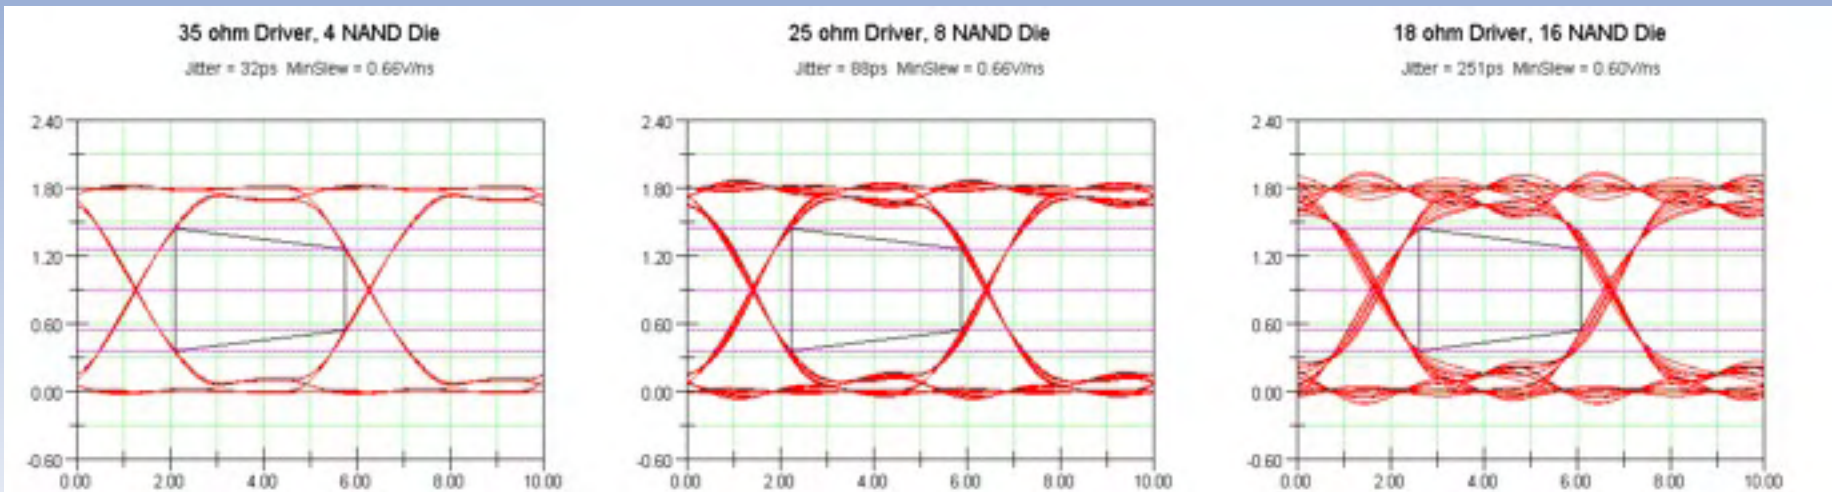

Simulation results courtesy of Micron Technology

Micron 34nm SLC NAND
Sustained Bandwidth

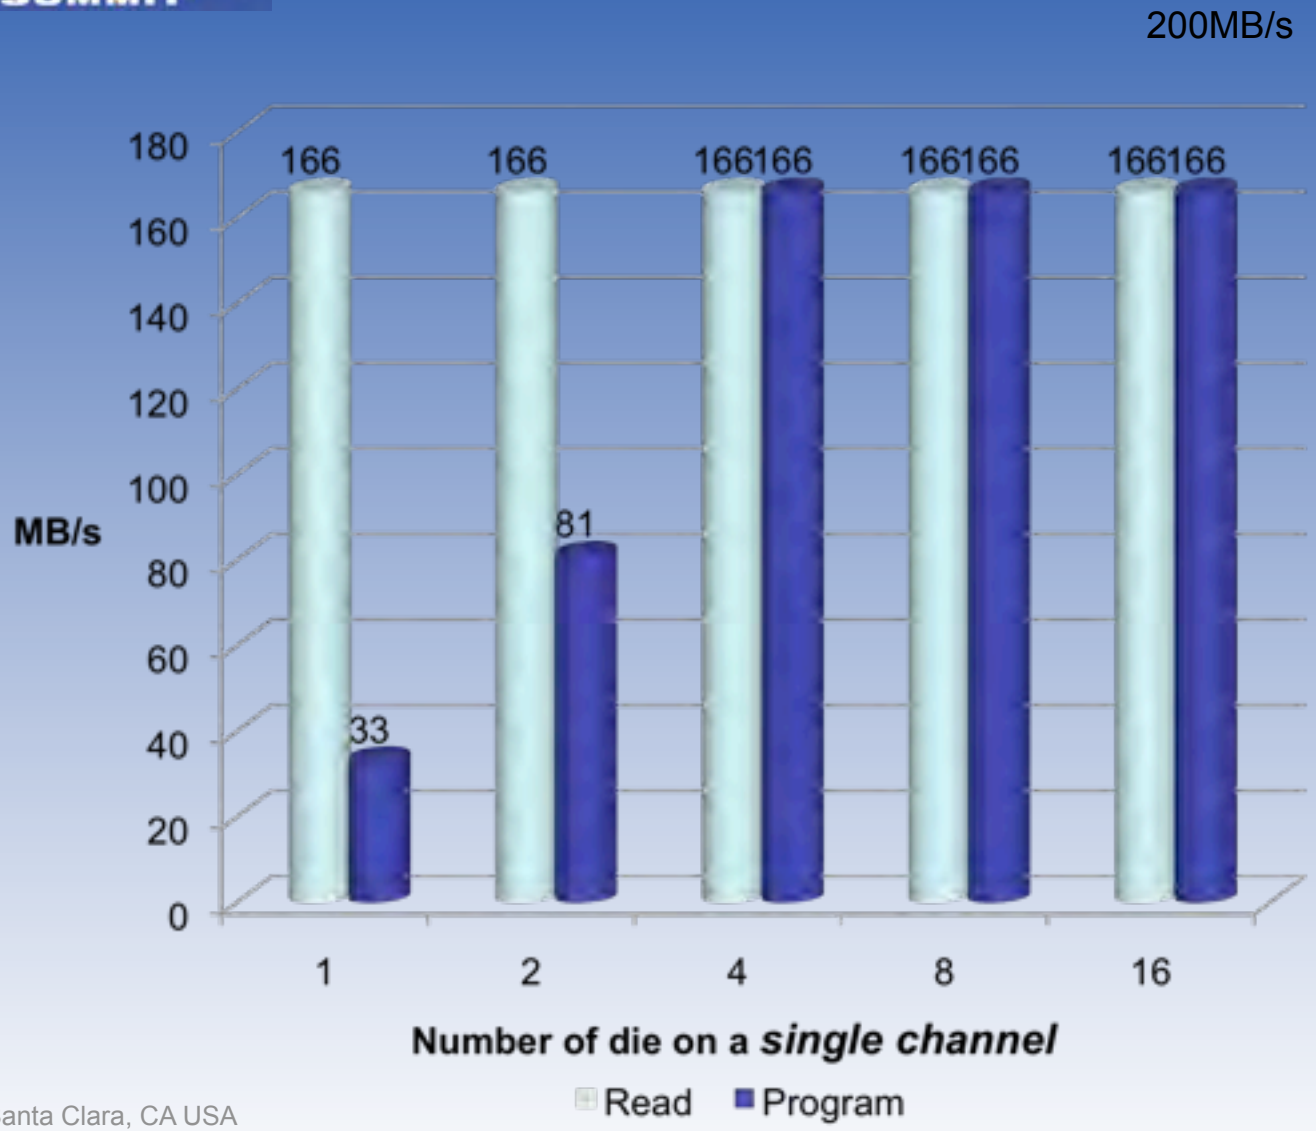

ONFi 2.1 sync supports transfers up to 200 MB/s

Traditional asynchronous interface is limited to less than 40 MB/s

Micron 34nm MLC NAND
Sustained Bandwidth

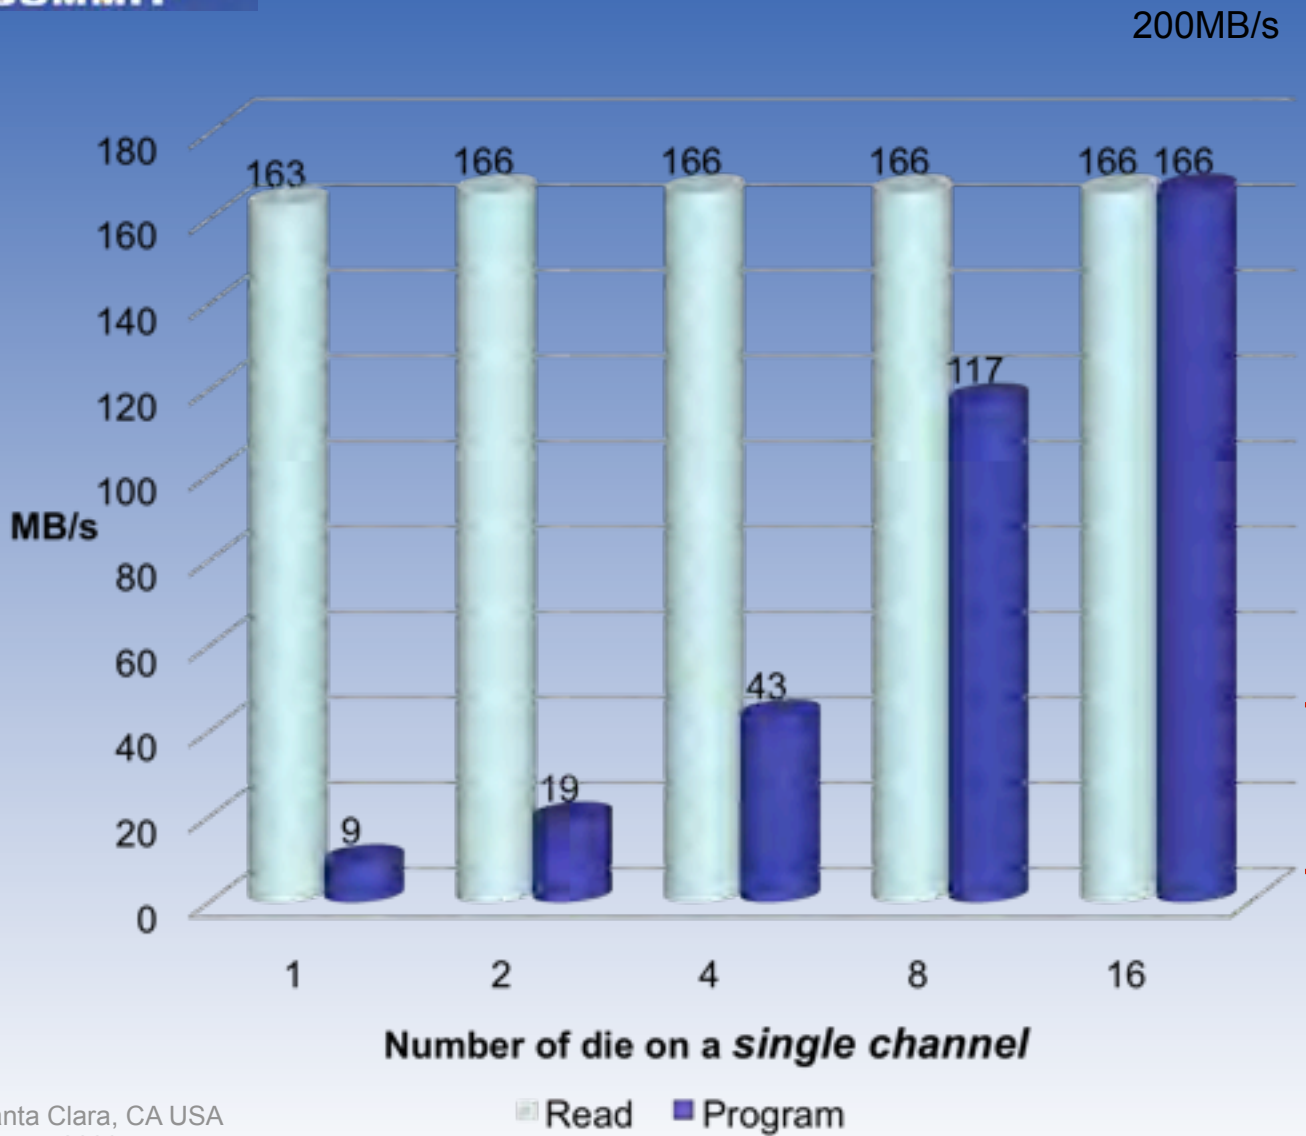

ONFi 2.1 sync supports transfers up to 200 MB/s

Traditional asynchronous interface is limited to less than 40 MB/s

Saving Power Wherever Possible

- To decrease power during writes to the NAND, ONFi 2.1 allows the host to stop the clock during writes
  - Savings of 10s of milliwatts are possible with this technique

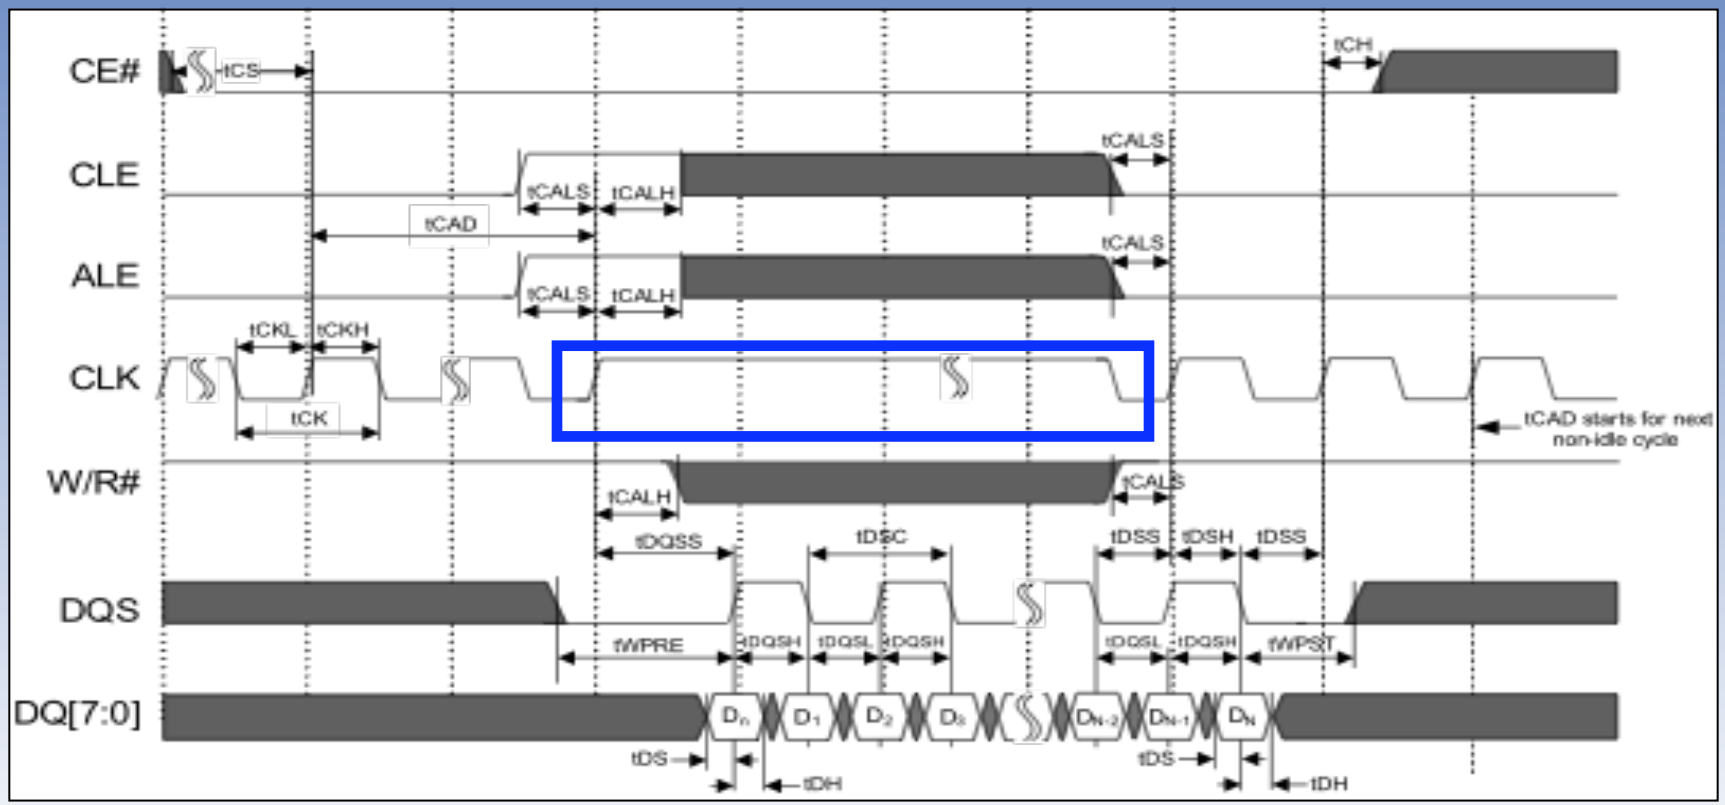

ECC Requirements Growing

- ECC is required to correct for bit errors that naturally occur with NAND
- As the raw bit error (RBER) of NAND increases, the amount of ECC applied goes up
- Communicating error correction needs effectively is critical

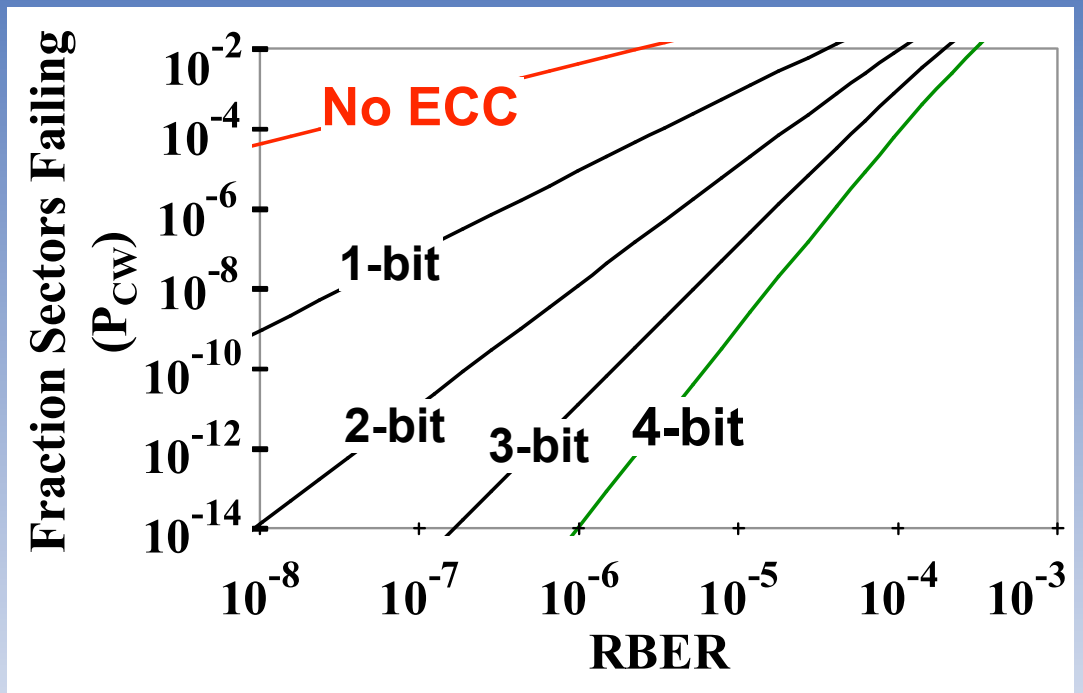

Communicating ECC Effectively

- For ECC, a set of parameters have to be looked at holistically as they are intimately related; the parameters are:
  - ECC codeword size
  - Number of bits to correct anywhere within the codeword
  - Number of program/erase cycles (distributed)
  - Number of valid blocks
- ONFi 2.1 allows the NAND device to communicate these parameters together in the *Extended ECC Information*
- The device may communicate multiple sets of these parameters if the NAND may be used in different environments (e.g., server, client, netbook)
  - The NAND device could indicate parameters for three different block endurance values (e.g., 1,000, 10,000 and 100,000 cycles)
  - The device can communicate the needed codeword size, which likely will grow from 512B to 1KB to avoid large increases in spare area

Byte	Definition
0	Number of bits ECC correctability
1	Codeword size
2-3	Bad block maximum per LUN
4-5	Block endurance
6-7	Reserved

Interleaved Reads for More Performance

- A logical unit (LUN) may support interleaved addressing
- This allows two or more of the same type of operations to execute concurrently
- ONFi 2.1 has added support for interleaved reads
- Enables deeper pipelining of reads, especially in MLC situations where the array time may be longer

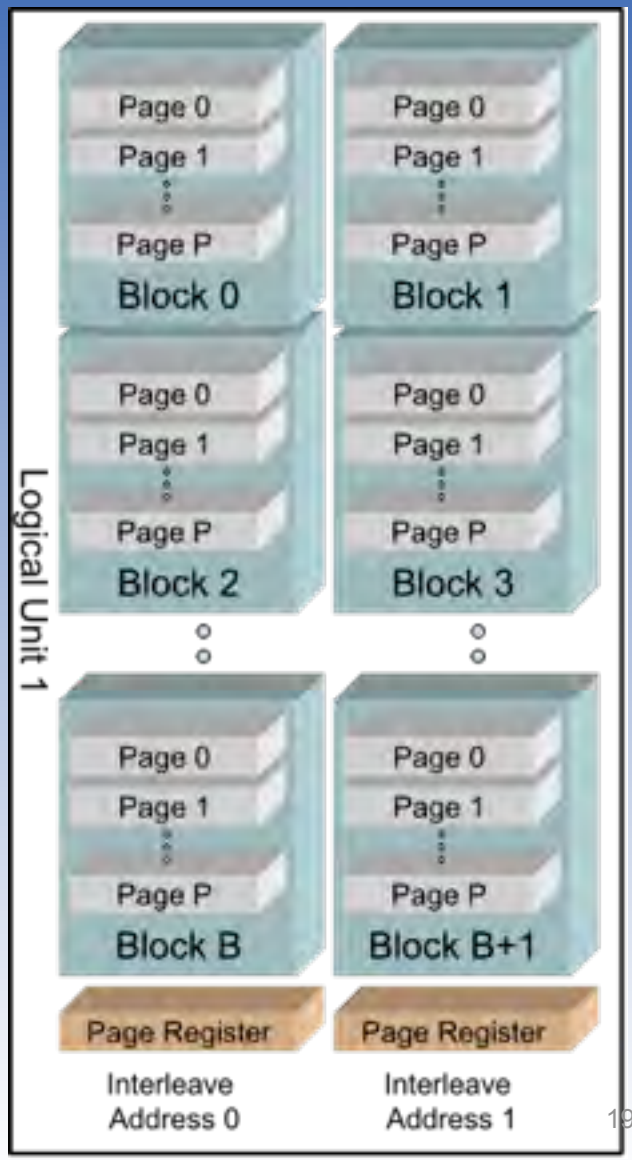

Enabling Low-Cost Controllers

- For copyback operations, low-cost controllers are forced to have internal SRAM that is a minimum of the NAND page size
- This requirement is necessary to allow the controller to correct ECC errors from the pages read as part of the copyback
- ONFi 2.1 has eliminated this storage requirement by introducing the SMALL DATA MOVE command

Small Data Move Detail

- SMALL DATA MOVE allows a portion of data to be written, followed by more data output

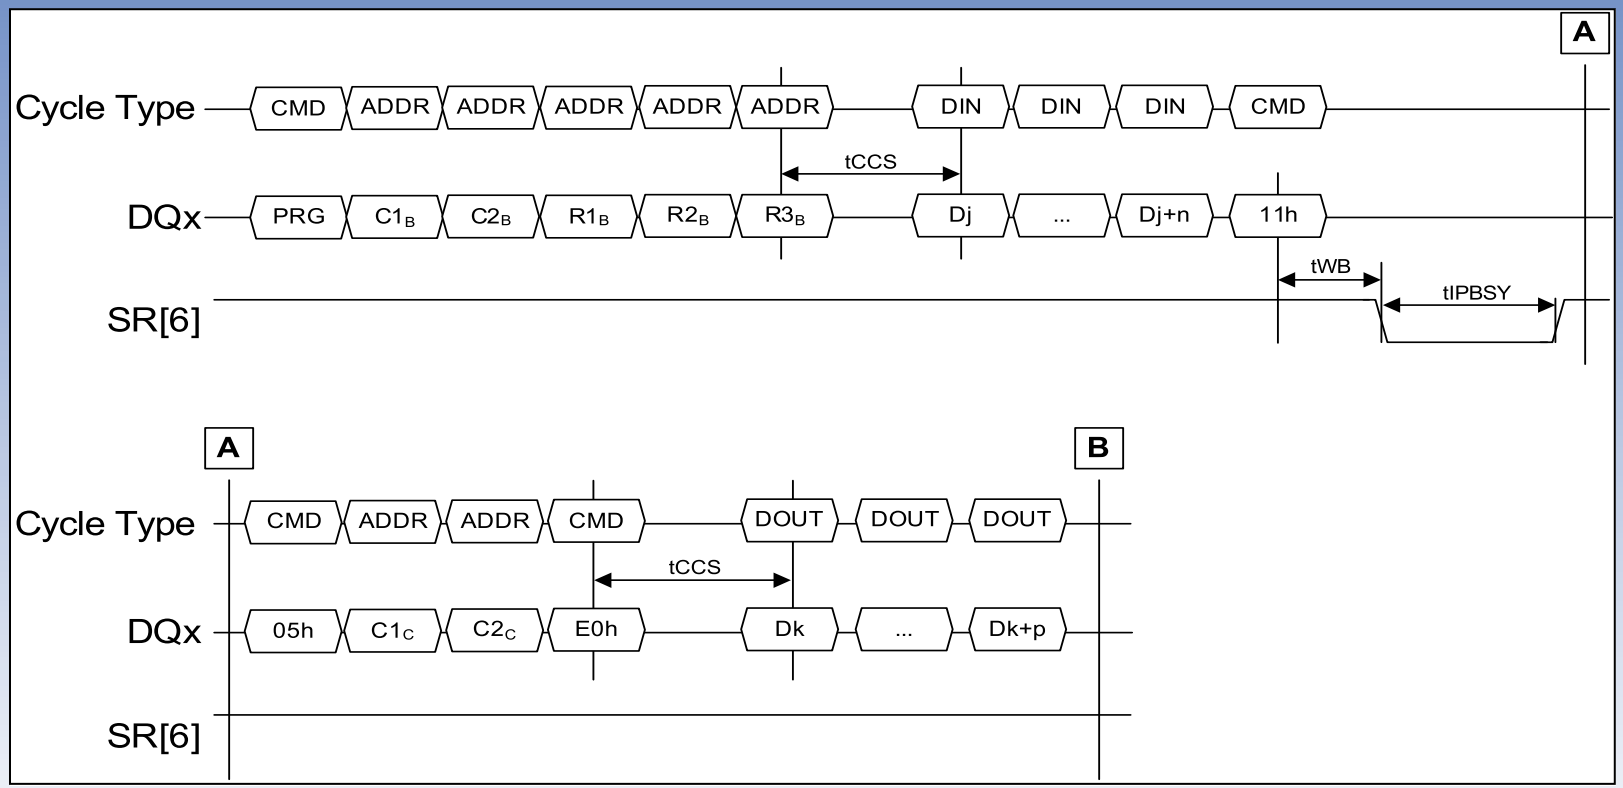

Making NAND Better for Card Applications

- In MMC and SD applications, there are commands that allow the host to write indefinitely
  - When the host stops writing, it may be in the middle of a page, which is awkward for the NAND

- ONFi 2.1 introduces the CHANGE ROW ADDRESS command
  - Change row address changes the location the page is being written to, avoiding writing an incomplete page to the final location
  - Firmware can then perform read/modify/write actions and write the complete page to the originally targeted location

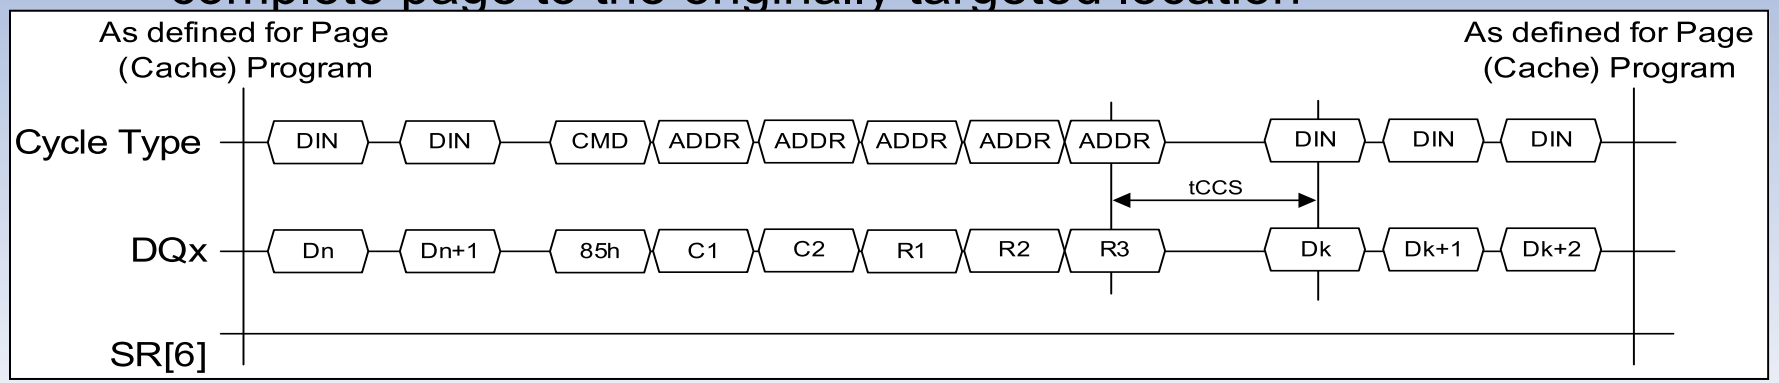

Outline

- ONFi Overview
- ONFi Provides Speed Enhancements
- *The Path to Higher Performance*

The Path to 400 MT/s

- Key enablers for ONFi 3.0 include:
  - A shorter channel (controller distance to the NAND)
  - Wider spacing between signals
  - On-die termination
  - Complementary clock and DQS signals

- Reaching particular speed grades requires a combination of enablers

The Path to 400 MT/s

- Key enablers for ONFi 3.0 include:
  - A shorter channel (controller distance to the NAND)
  - Wider spacing between signals
  - On-die termination
  - Complementary clock and DQS signals

- Reaching particular speed grades requires a combination of enablers

The Path to 400 MT/s

- Key enablers for ONFi 3.0 include:
  - A shorter channel (controller distance to the NAND)
  - Wider spacing between signals
  - On-die termination
  - Complementary clock and DQS signals

- Reaching particular speed grades requires a combination of enablers

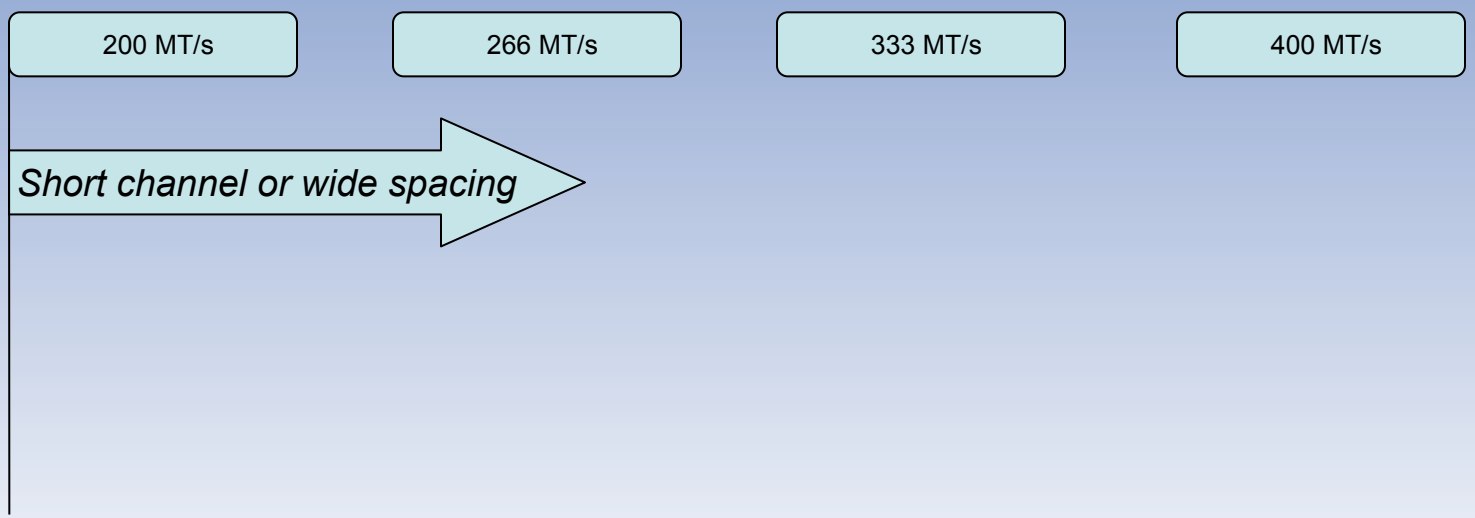

The Path to 400 MT/s

- Key enablers for ONFi 3.0 include:
  - A shorter channel (controller distance to the NAND)
  - Wider spacing between signals
  - On-die termination
  - Complementary clock and DQS signals

- Reaching particular speed grades requires a combination of enablers

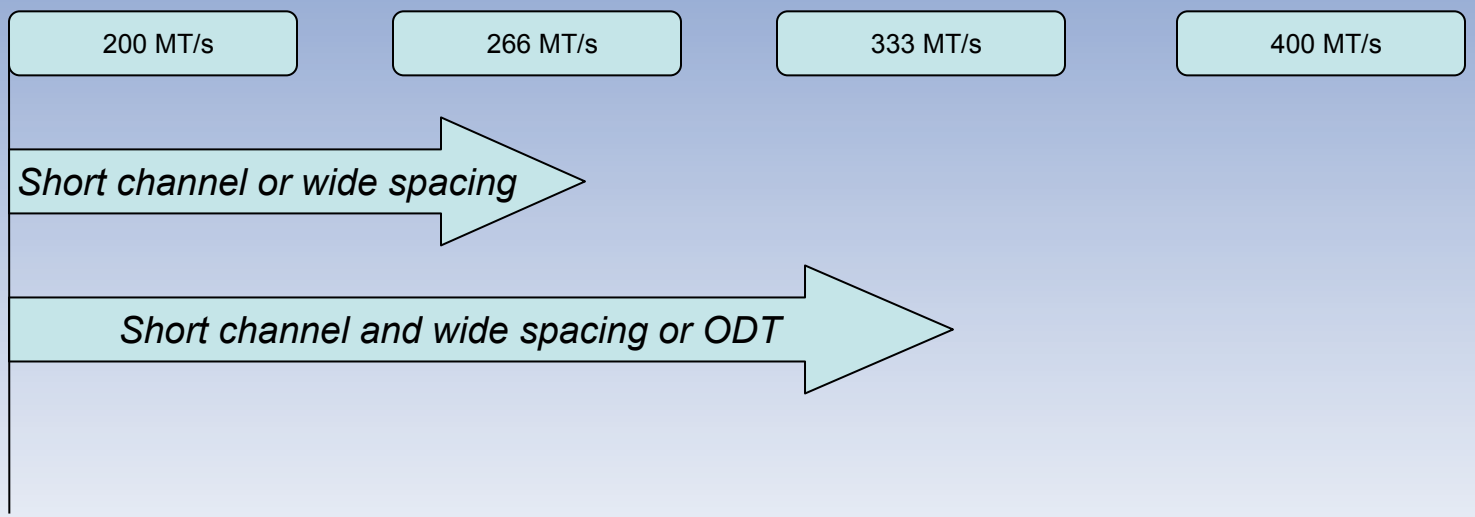

The Path to 400 MT/s

- Key enablers for ONFi 3.0 include:
  - A shorter channel (controller distance to the NAND)
  - Wider spacing between signals
  - On-die termination
  - Complementary clock and DQS signals

- Reaching particular speed grades requires a combination of enablers

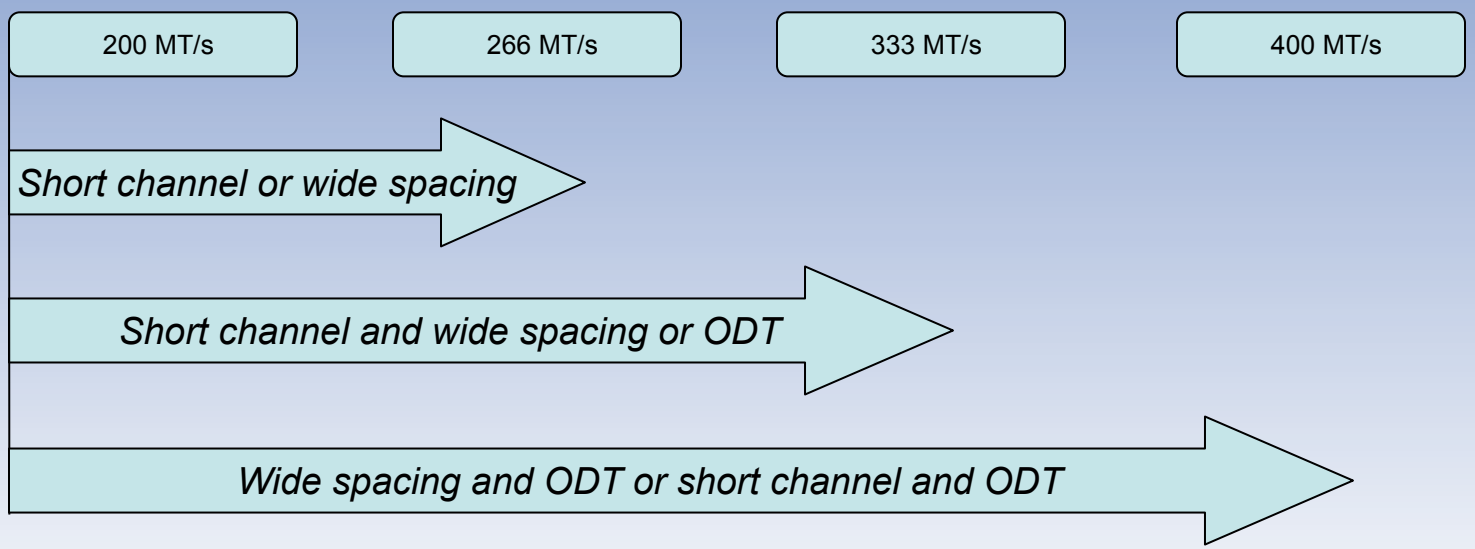

The Magic of On-Die Termination

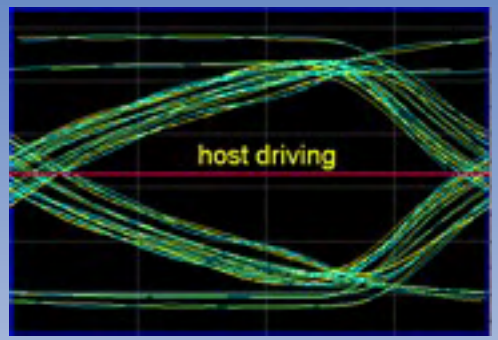

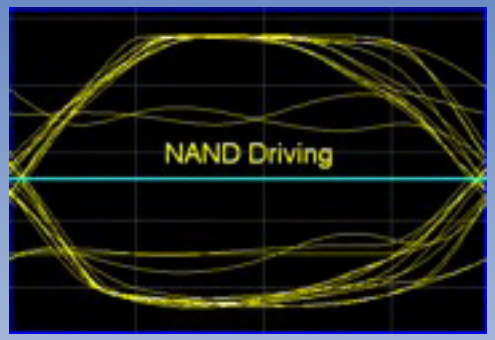

The Magic of On-Die Termination

- Data eye at 266 MT/s without ODT
  - Passable, but getting close to the edge…

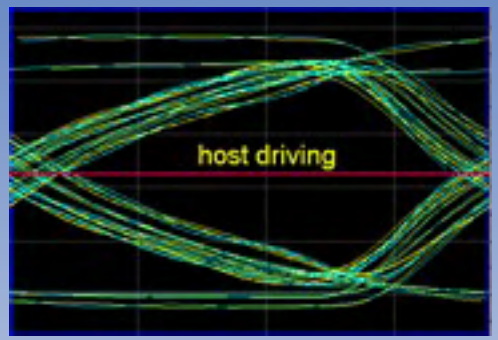

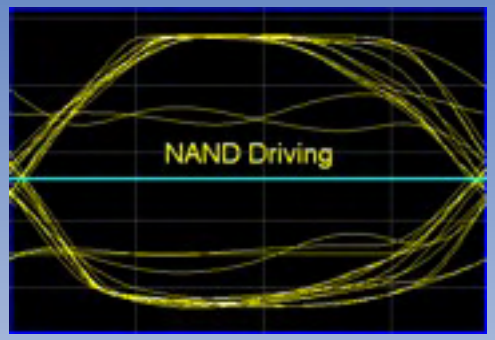

The Magic of On-Die Termination

- Data eye at 266 MT/s without ODT
  - Passable, but getting close to the edge…

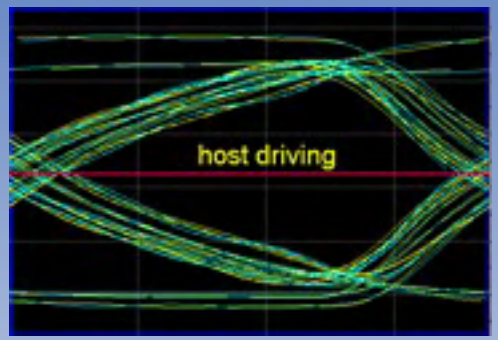

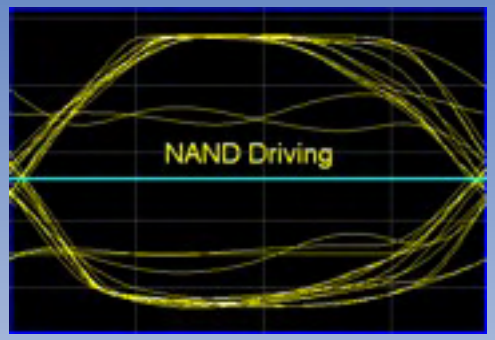

The Magic of On-Die Termination

- Data eye at 266 MT/s without ODT
  - Passable, but getting close to the edge…

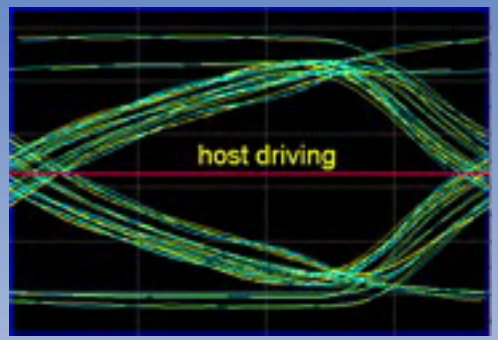

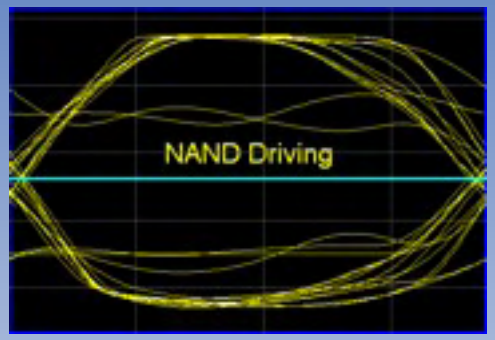

The Magic of On-Die Termination

- Data eye at 266 MT/s without ODT
  - Passable, but getting close to the edge…

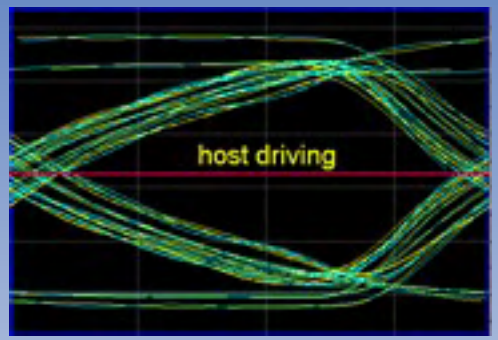

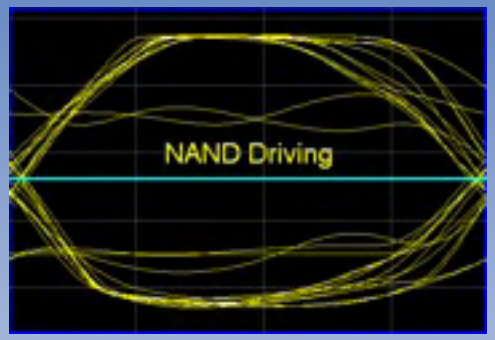

The Magic of On-Die Termination

- Data eye at 266 MT/s without ODT
  - Passable, but getting close to the edge…

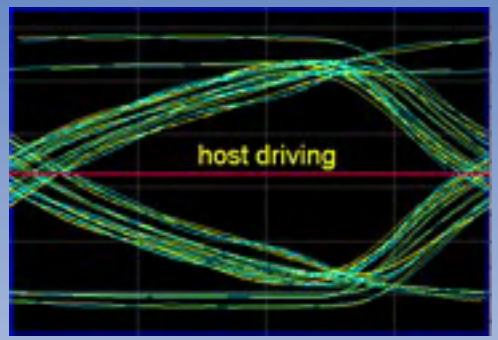

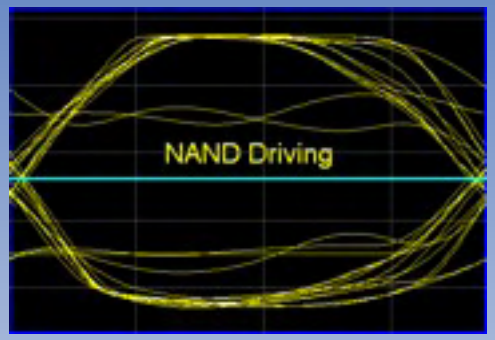

The Magic of On-Die Termination

- Data eye at 266 MT/s without ODT
  - Passable, but getting close to the edge…

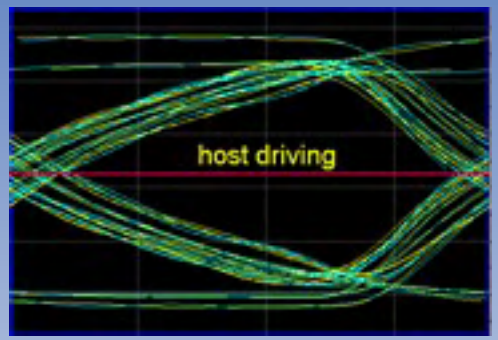

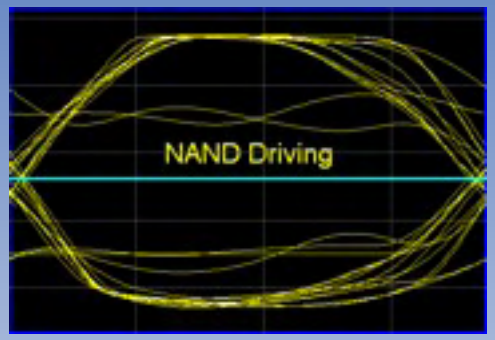

- Data eye at 400 MT/s with ODT “magic”
  - On-die termination is the key enabler for 400 MT/s operation

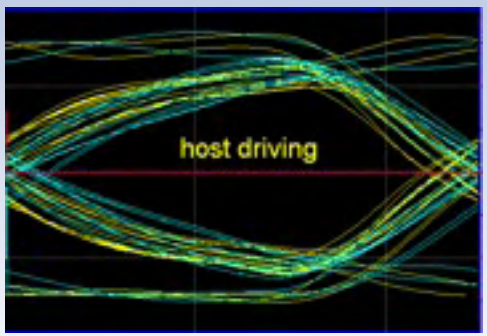

Trend Toward Longer Array Times

- The trend is toward NAND array times getting longer
  - Longer with lithography transitions
  - Longer with more bits per cell
- However, quality of service is a key metric, especially in enterprise
- What happens if a high-priority read comes in while doing an erase on the associated die/plane?

Operation	Estimated Array Time Ranges		
	SLC	MLC, 2-bit	MLC, 3-bit
Read	20–30µs	40–60µs	80–200µs
Program	200–300µs	0.3–2.5ms	0.3–10ms
Erase	0.5–3ms	2–10ms	2–30ms
Note: Represents a potential range of array times across lithography and implementation, being used as a guide by the ONFi technical team.			

Interrupting for High-Priority Reads

- It's desirable to have a method to interrupt erases, and potentially programs, when a high-priority read comes in
- ONFi 2.2 is pursuing two potential solutions
  - Erase Suspend
    - When the suspend is issued, the erase is stopped at the next erase pulse boundary
    - Advantage: Always makes forward progress on the erase
    - Implemented in NOR
  - Reset LUN
    - Resets only a particular LUN (die), so that another operation can quickly be issued
    - No forward progress is guaranteed
    - Advantage: Can be used to stop an erase or a program, no state to maintain to continue the erase

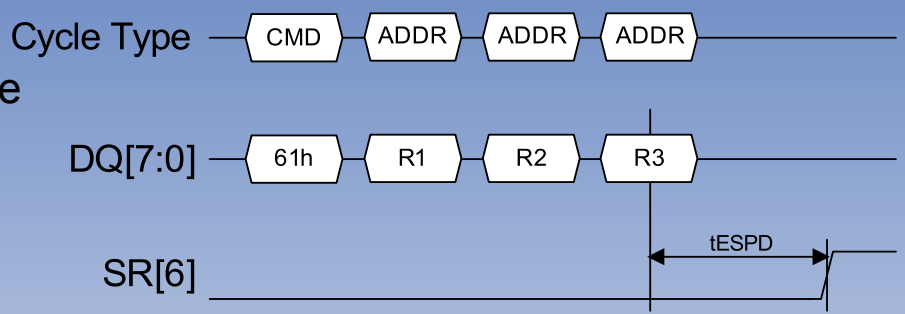

ONFi 2.2 is standardizing a method to ensure great quality of service for high-priority reads.

Summary

- ONFi is a single-purpose organization, enabling significant results in a timely manner
  - Five specifications have been delivered since the group was formed in 2006
  - ONFi 2.1 was released in January with speed grades of up to 200 MB/s
- ONFi is pursuing future performance
  - ONFi 2.2 will include a mechanism to stop a long array operation in order to satisfy a high-priority demand read quickly
  - ONFi 3.0 will enable speeds of up to 400 MB/s (another doubling)

Get involved with ONFi. Visit onfi.org!

Thank You!