



Flash Performance Enhancements through ONFI

Amber Huffman
Intel Corporation
ONFI Technical Chair



SAN JOSE

SEP 10 - 14

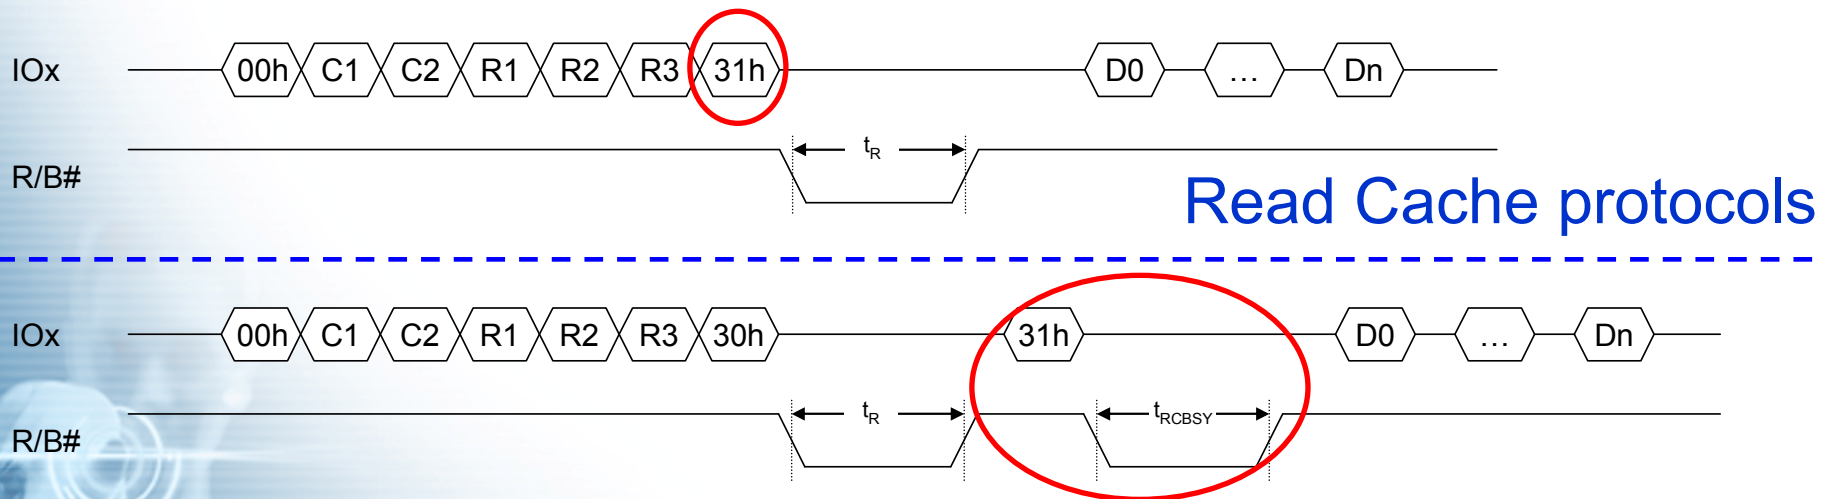
DRAM FLASH CE-ATA
PCI Express USB SATA

Agenda

- Current state of NAND interoperability
- ONFI changes the paradigm
- History of Read performance enhancements
- High performance applications and utility of traditional Read performance enhancements
- How ONFI does it better

The State of NAND Interoperability

- NAND is the only commodity memory with no standard i/f
- Basic NAND commands are similar amongst vendors
 - Read, Program, Erase, Reset, Read Status
- Timings vary from vendor to vendor
 - Often supporting a few vendors easy, all is more difficult
- Enhanced NAND commands vary widely



ONFI Paves the Way

- ONFI establishes a standard interface for NAND
 - A new paradigm is enabled: Host vendors no longer wait for the datasheet for the latest NAND part and then start making changes
- ONFI paves the way to:
 - Shorter NAND and product qualification times
 - More innovation since basic NAND command sequences always work
 - The host utilizing high performance command sequences

ONFI
OPEN NAND
FLASH INTERFACE

SONY

Micron[®]

ONFI Founders

hyunix

intel

PHISON
Knows What You Need

ST



SAN JOSE SEP 12 - 14

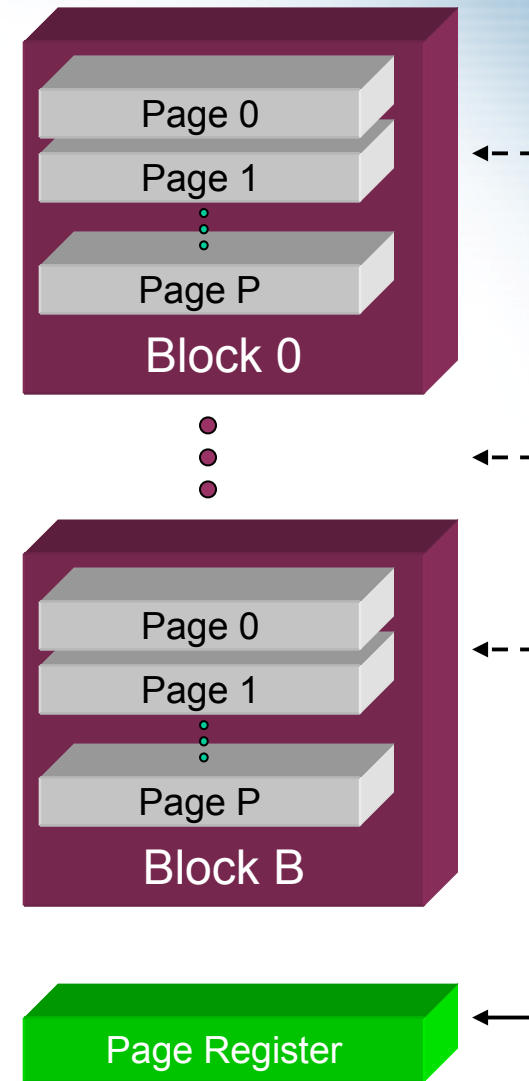


History of Read Performance Enhancements

- Cache Operations
 - Read Cache, variant 1
 - Read Cache, variant 2
- Multiple plane operations
 - Reads


Page Register Resource

- The page register is a critical data staging area in NAND operations
- Reads:
 1. Data flows from array to page register
 2. Data flows from page register to host
- Programs:
 1. Data flows from host to page register
 2. Data flows from page register to array
- The host is “stuck” until the page register is available again



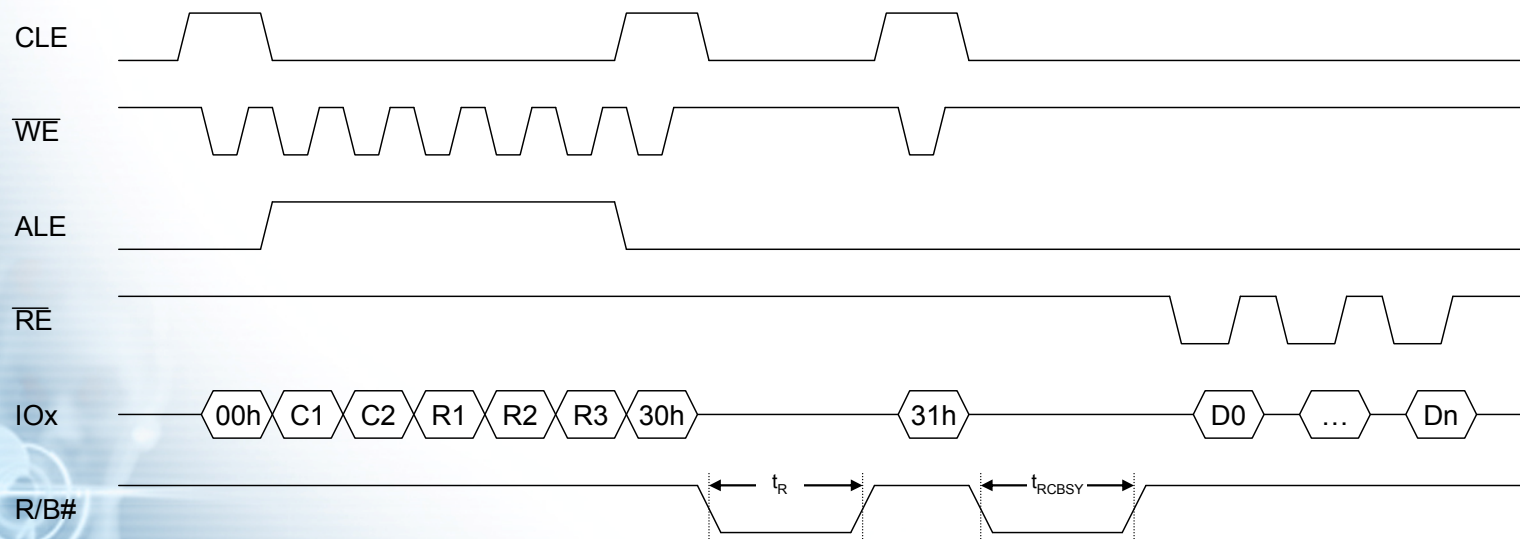
Removing the Page Register Bottleneck

- While the host is reading data from the page register, the next page could be read from the array
But... there is only one page register...
- The NAND vendors added a cache register to solve this issue



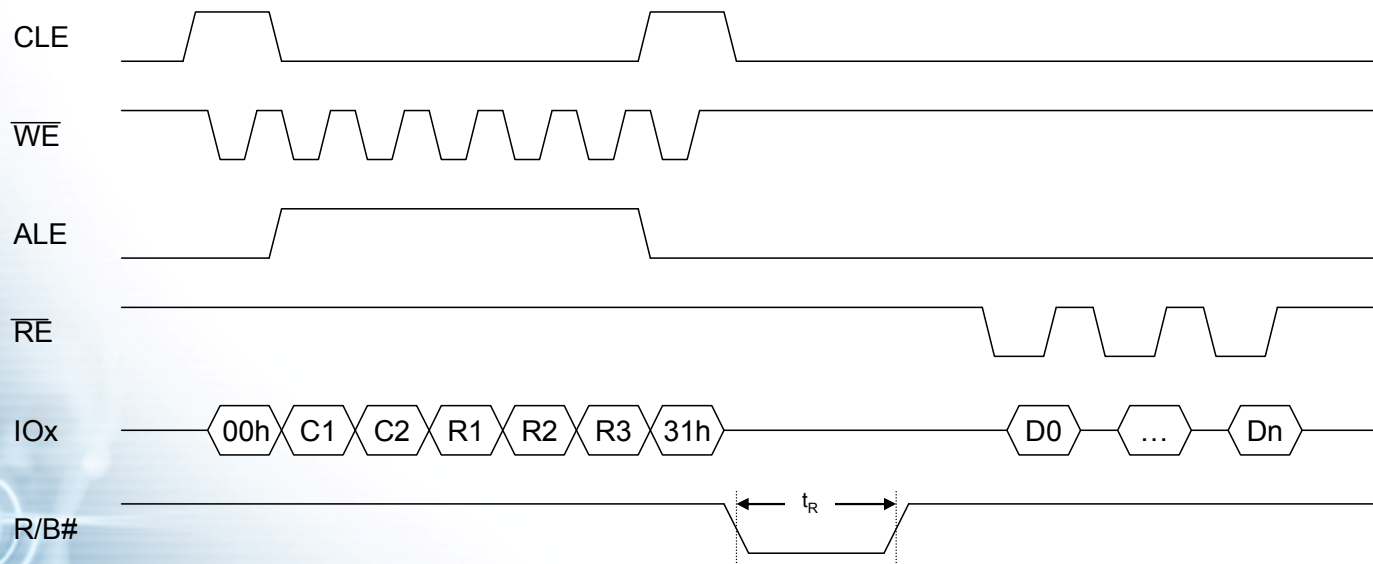
Read Cache, variant 1

- The sequence is:
 - Issue normal read command
 - Wait t_R time
 - Then issue 31h command to tell the NAND to start a background read operation to the cache register
- Limitation:
 - Enhancement only applies for sequential page reads



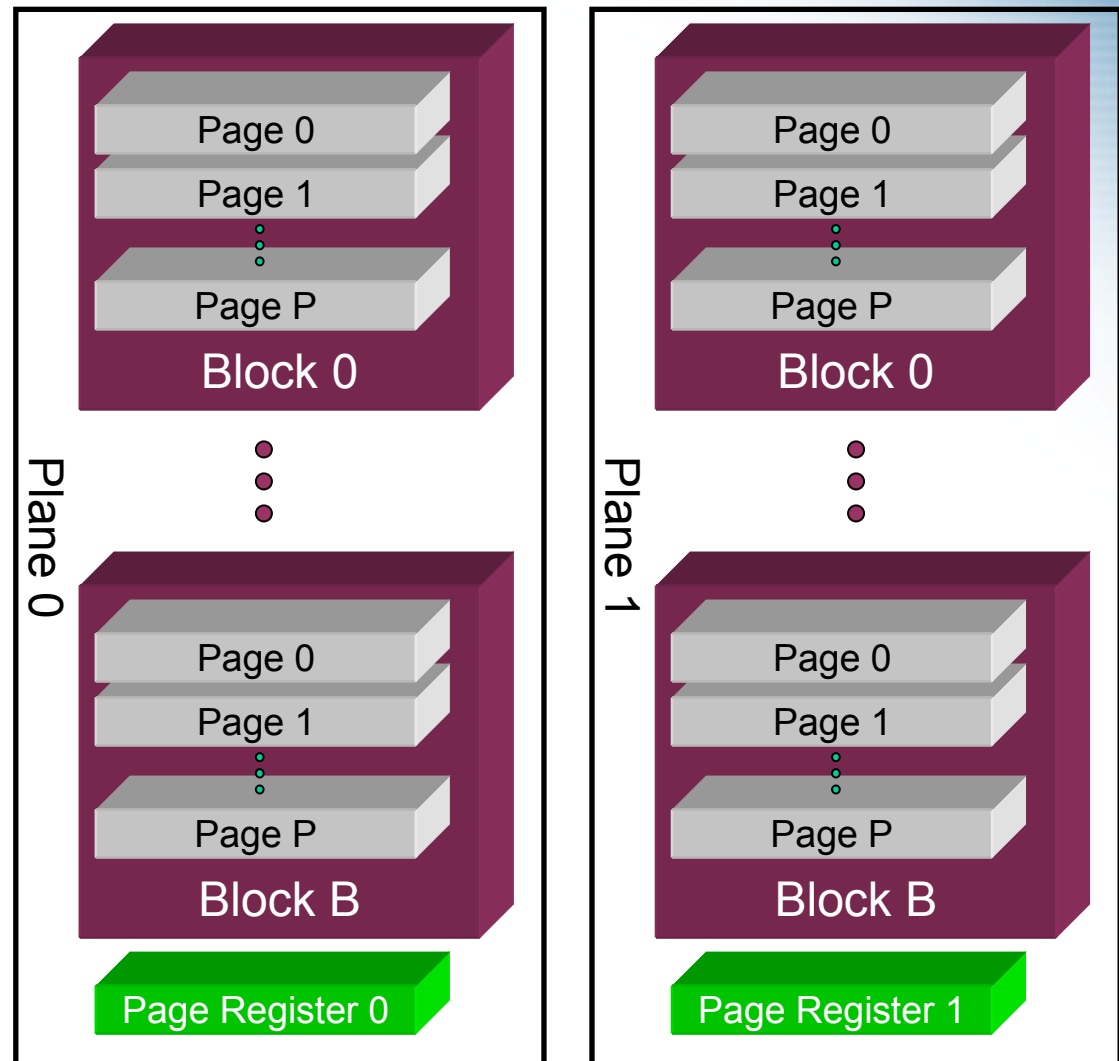
Read Cache, variant 2

- The sequence is:
 - Issue a modified Read command; finish with 31h instead of 30h
 - Each time the end of a page is reached, the NAND automatically starts reading the next page
- Limitation:
 - Enhancement only applies for sequential page reads



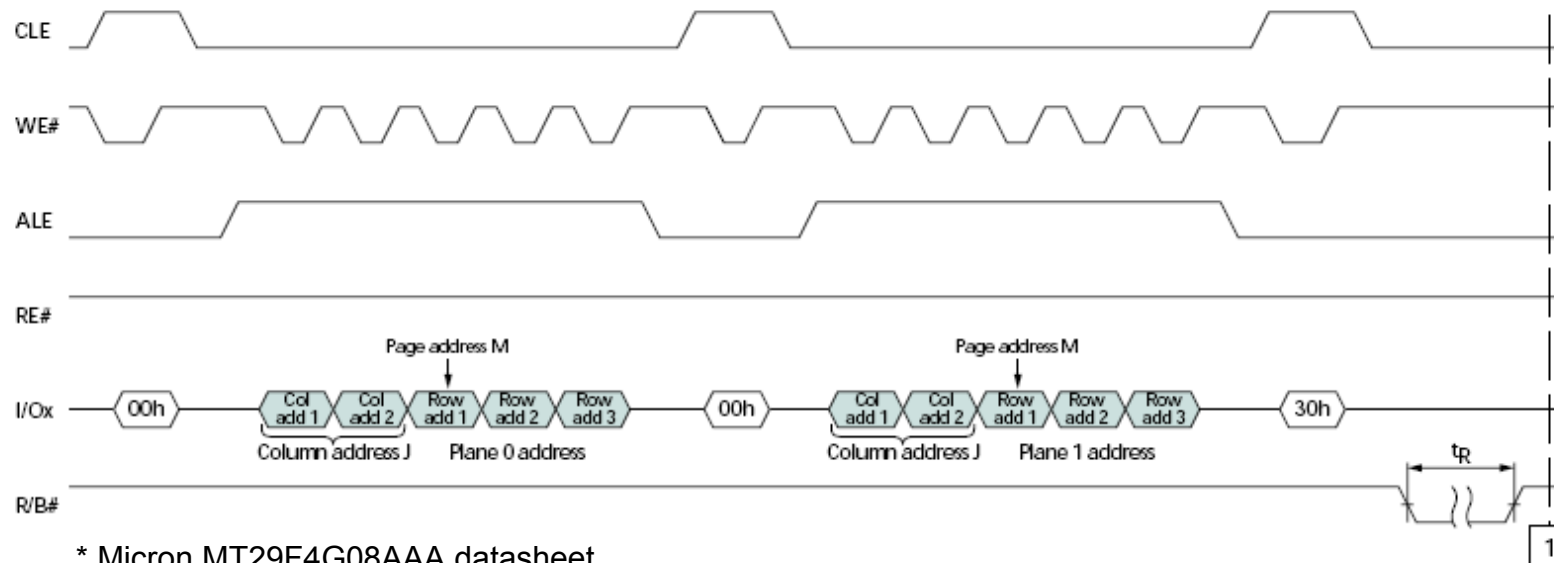
Multiple Plane Operations

- NAND vendors have started splitting the array into “planes” within a die
- Allows simultaneous operations of the same type to different block addresses
 - Reads
 - Programs
 - Erases



Multiple Plane Read Operations

- Multiple plane read operations are when multiple reads are issued at roughly the same time
- Limitation:
 - Page address for reads have to be the same



Behavior of High Performance Applications

- Apps, like Robson, receive frequent random reads
- To achieve high read performance for high performance applications, handling random and non-deterministic read patterns well is required

Analyze Trace - [PC05_XPStart_cut.raw]

Process: Reads Writes Select Drive: DISK #0

Index	Start	Finish	Service Time	OP Code	Length	LBA
238	01.329450	01.329597	00.000147	Read	8	3a4a97
239	01.329648	01.339210	00.009562	Read	8	6459f7
240	01.339826	01.339950	00.000124	Read	8	3a4a9f
241	01.340233	01.351330	00.011097	Read	8	3f7b5f
242	01.351539	01.357840	00.006301	Read	8	452c47
243	01.357881	01.358007	00.000126	Read	8	3a4aa7
244	01.358053	01.373354	00.015301	Read	8	58deb7
245	01.373399	01.386810	00.013411	Read	8	3a4ab7
246	01.386851	01.386973	00.000122	Read	8	452c3f

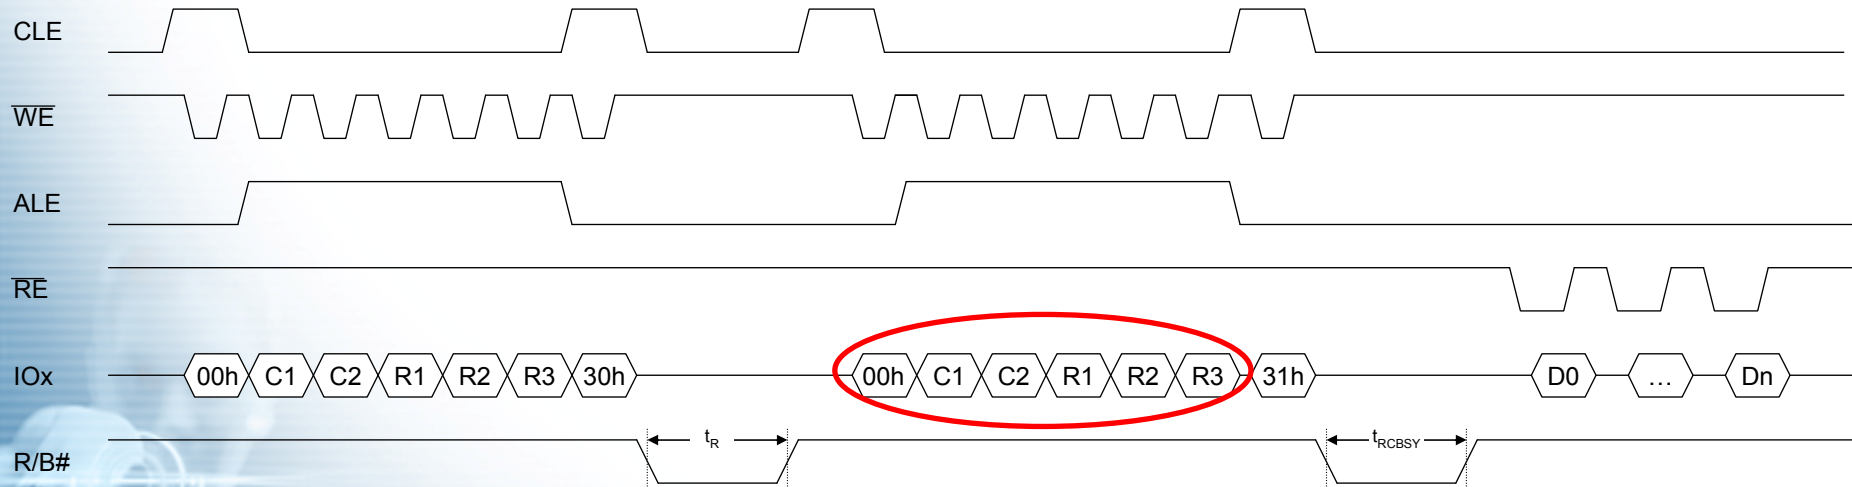
2: **None of the reads are issued in sequential order**

High Performance Apps and Read Enhancements

- Recall limitations with Read enhancements:
 - Read Cache: Must be sequential page within block
 - Multi-plane Reads: Needs to be same page in each block
- Recall high performance app read patterns:
 - Non-deterministic
 - Often highly random
- Result: Traditional read enhancements have little value in high performance applications!

ONFI Simplifies while Expanding Utility

- ONFI defines one read enhancement
- Read Cache Enhanced
 - Allows 00h-Cmd 5-Address cycles to precede 31h of Read Cache (variant 1) to deliver a random read cache
 - Host can choose to issue 31h alone for sequential behavior, thereby preserving any previous host investment in Read Cache



ONFI Delivers Read Performance Benefits

- Read Cache Enhanced delivers dramatic performance benefits in any read sequence (sequential or random)
 - Typical SLC performance boosted by 30%+
 - Typical MLC performance boosted by 50%+

Normal Page Read

		READ time / Block (ns)	data /Block (byte)	Read Performance (MB/s)
Page Read	SLC	5,668,480	135,168	23.85
	MLC	12,936,960	270,336	20.90

Cache Read – Sequential Pages

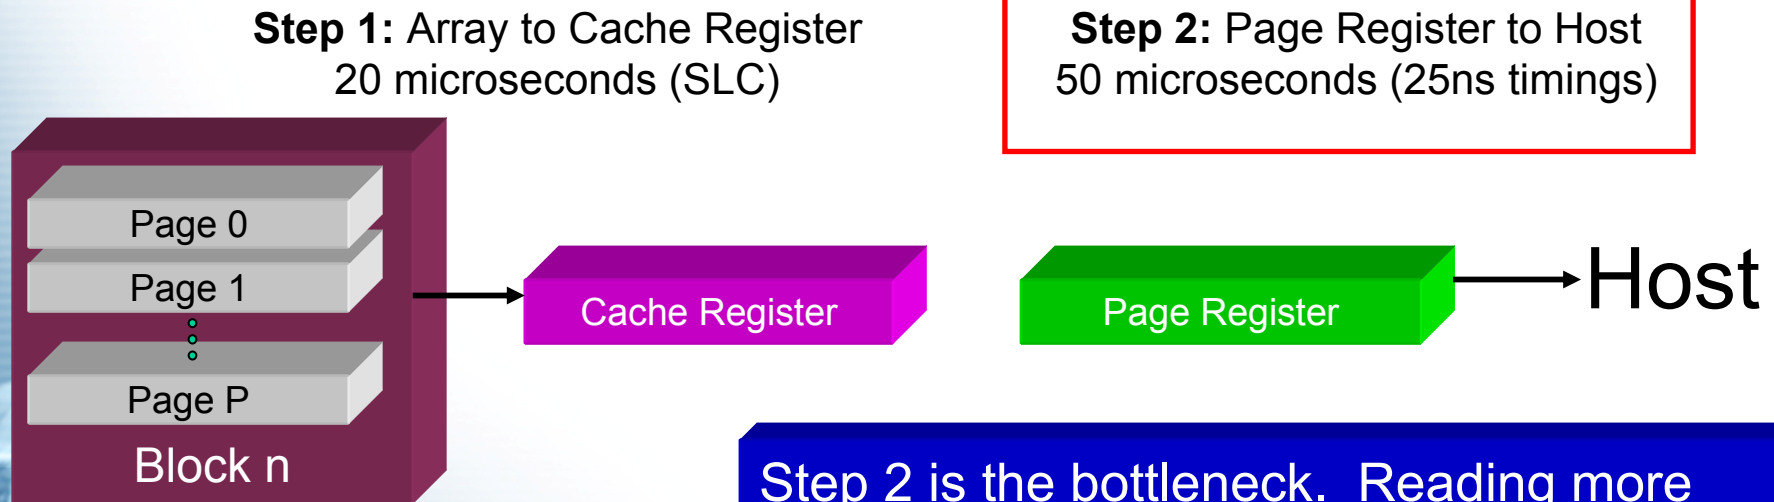
		READ time / Block (ns)	data /Block (byte)	Read Performance (MB/s)	vs. Page Read (%)
Read Cache Enhanced	SLC	4,274,170	135,168	31.62	32.6%
	MLC	8,533,710	270,336	31.68	51.6%

Cache Read – Non-Sequential Pages

		READ time / Block (ns)	data /Block (byte)	Read Performance (MB/s)	vs. Page Read (%)
Read Cache Enhanced	SLC	4,285,570	135,168	31.54	32.3%
	MLC	8,554,710	270,336	31.60	51.2%

ONFI Eliminates Interleaved Reads

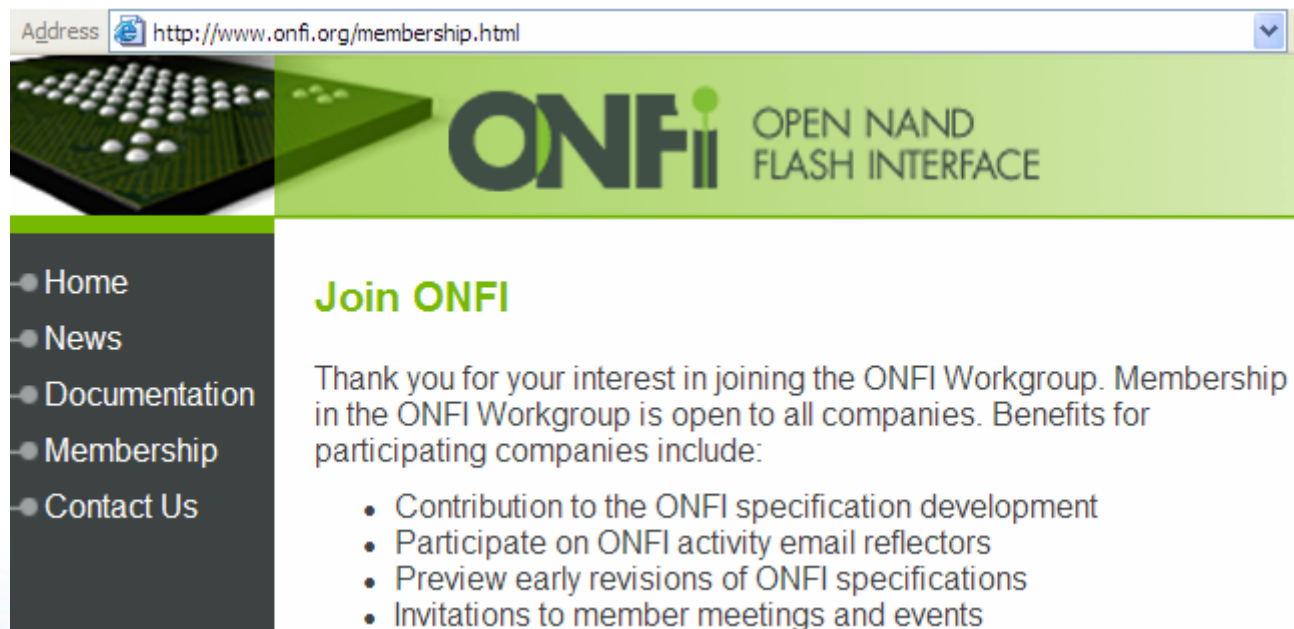
- ONFI has chosen to not implement interleaved reads
 - Multiple plane operations map well to interleaved operations in ONFI
- Why?
 - Limitation of same page address makes utility low
 - Rife with bus contention issues
 - Read Cache Enhanced provides all of the benefits



Step 2 is the bottleneck. Reading more than one page in the background from the array is pointless. Still serialized to host...

Summary

- ONFI delivers standardized high performance command sequences that deliver real value
 - Simplifies numerous command sequences in the industry
 - Delivers value to high performance applications
- Join ONFI and help deliver more tangible benefits to the NAND industry



The screenshot shows a web browser window with the address bar containing <http://www.onfi.org/membership.html>. The page features a green header with the ONFI logo and the text "OPEN NAND FLASH INTERFACE". A navigation menu on the left lists: Home, News, Documentation, Membership, and Contact Us. The main content area is titled "Join ONFI" and contains the following text: "Thank you for your interest in joining the ONFI Workgroup. Membership in the ONFI Workgroup is open to all companies. Benefits for participating companies include:" followed by a bulleted list of benefits.

- Contribution to the ONFI specification development
- Participate on ONFI activity email reflectors
- Preview early revisions of ONFI specifications
- Invitations to member meetings and events