



The Inconvenient Truths of NAND Flash Memory

Jim Cooke (jcooke@micron.com)

Applications Engineering Director

Micron Technology, Inc.

- NAND Flash Differences
 - MLC vs. SLC
 - Architecture, Features, and Performance Comparisons
- NAND Error Modes
 - Program Disturb
 - Read Disturb
 - Data Retention
 - Endurance
 - Wear Leveling
 - ECC Fixes Almost Everything

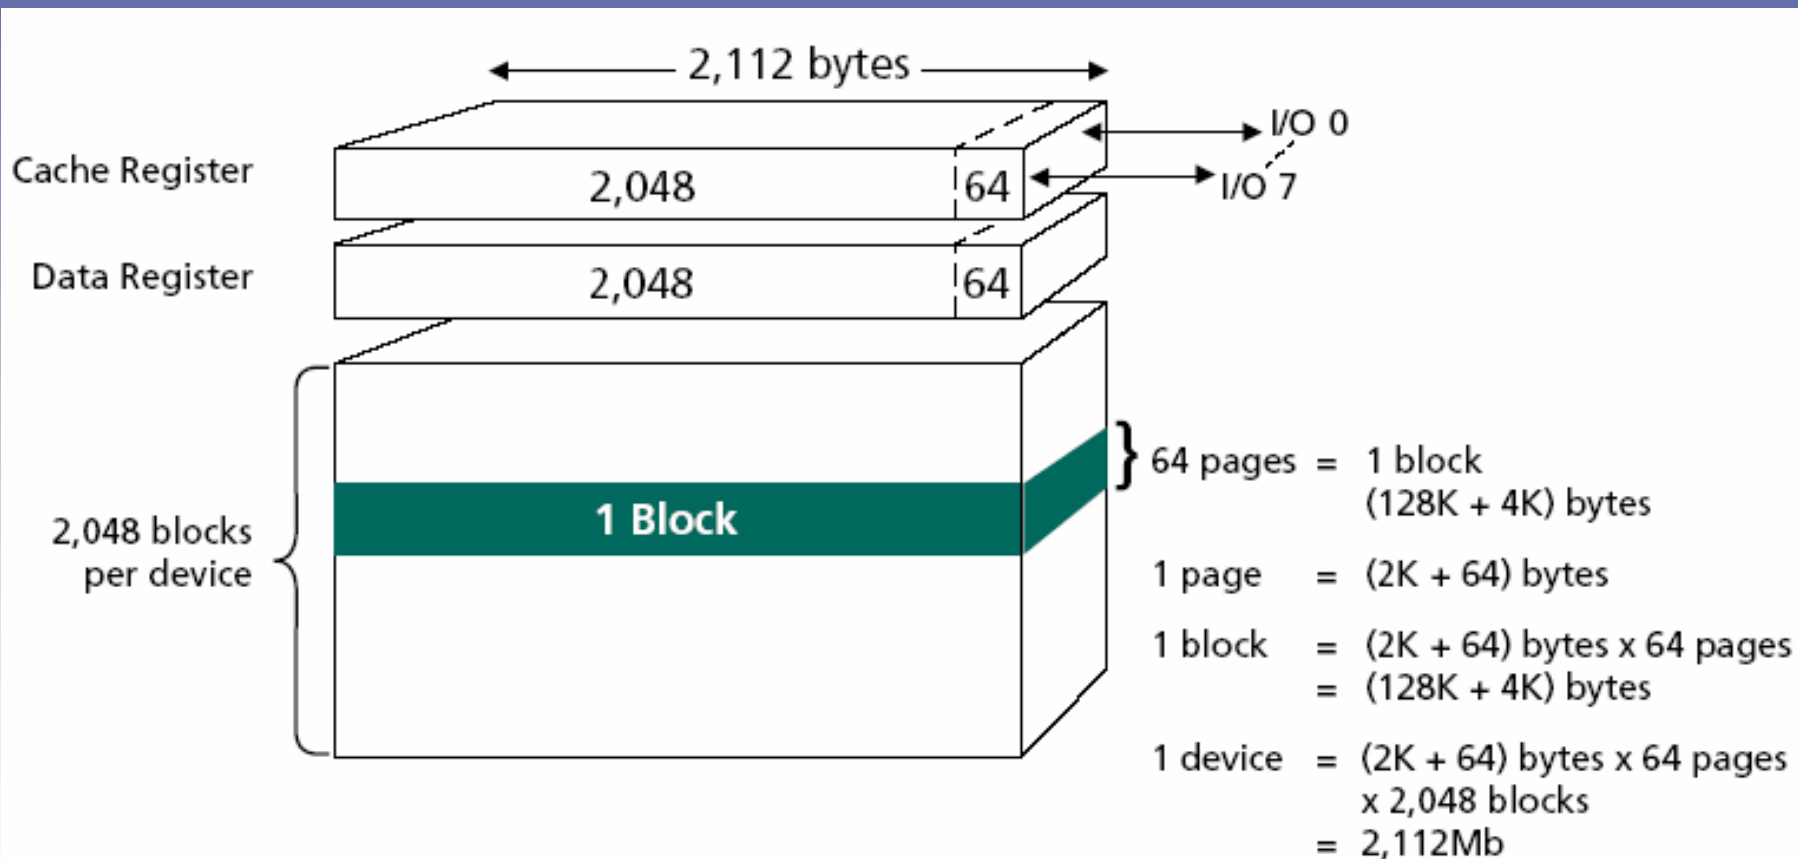
All NAND Flash Devices Are Not Created Equal

- Differences include:
 - Cell types
 - Architectural
 - Performance
 - Timing parameters
 - Command set
- Open NAND Flash Interface (ONFI) helps to address many of these

Cell Types – MLC vs. SLC

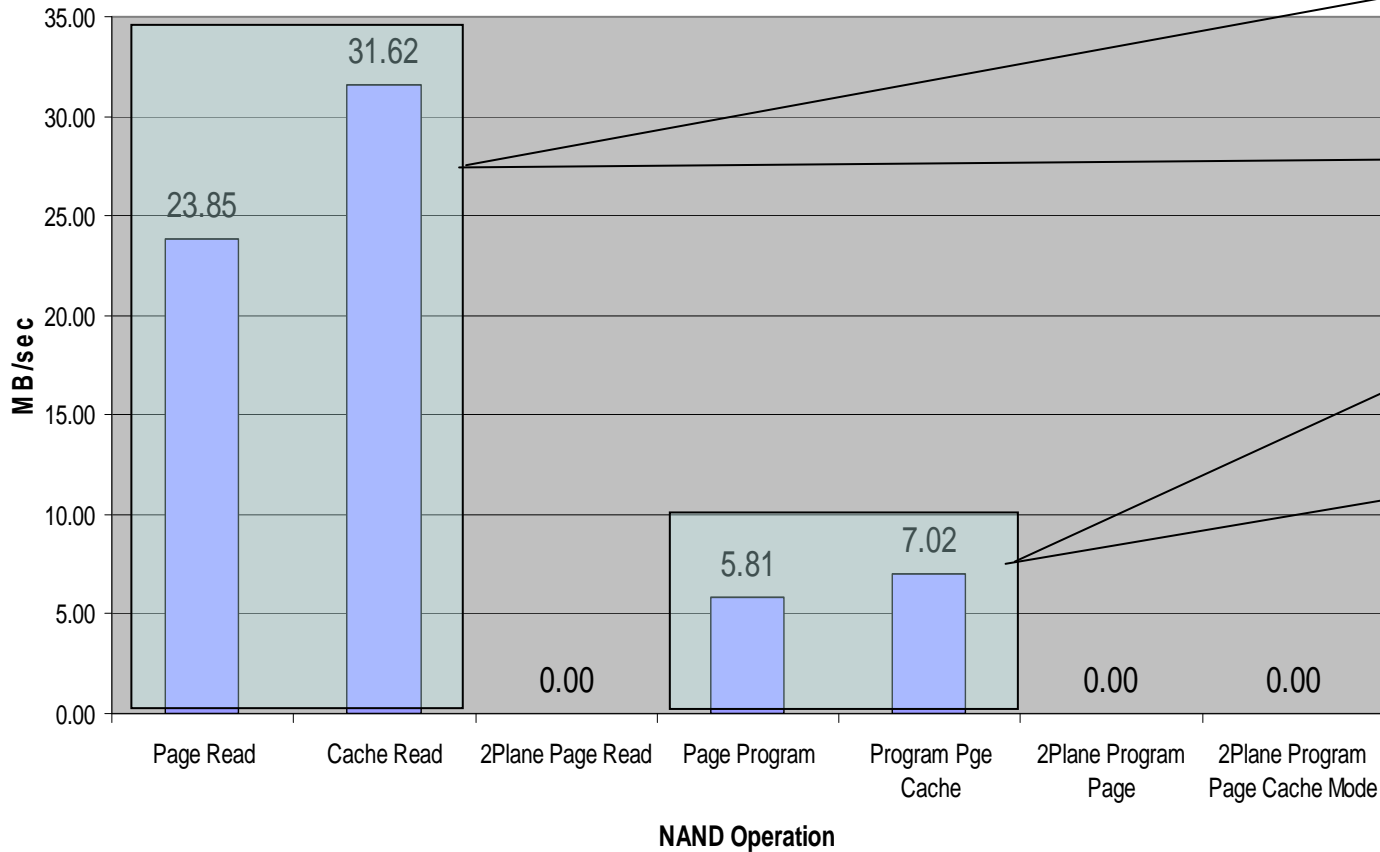
Features	MLC	SLC
Bits per cell	2	1
Voltage	3.3V	3.3V, 1.8V
Data width (bits)	x8	x8, x16
Architecture		
Number of planes	2	1 or 2
Page size	2,112–4,314 bytes	2,112 bytes
Pages per block	128	64
Reliability		
NOP (partial page programming)	1	4
ECC (per 512 bytes)	4+	1
Endurance (ERASE / PROGRAM cycles)	<10K	<100K
Array Operations		
t _R (Max)	50μs	25μs
t _{PROG} (Typ)	600–900μs	200–300μs
t _{BERS} (Typ)	2ms	1.5–2ms

2Gb, 2K-Page SLC NAND Architecture



2Gb, 2K-Page, 72nm, SLC Performance

Micron (72nm SLC) 2Gb die 2K Page Performance




Symbol	Time	Units
tR	25	us
tDCBSYR1	3	us
tDCBSYR2	3	us
tRC	30	ns
tRC (C)	30	ns
tRC	50	ns
tPROG	240	us
tCBSY	300	us
tDBSY	0.5	us
tWC	30	ns
tWC (C)	30	ns
tWC	45	ns
PS	2112	Byte
NP	64	Pages

Two-Plane Features

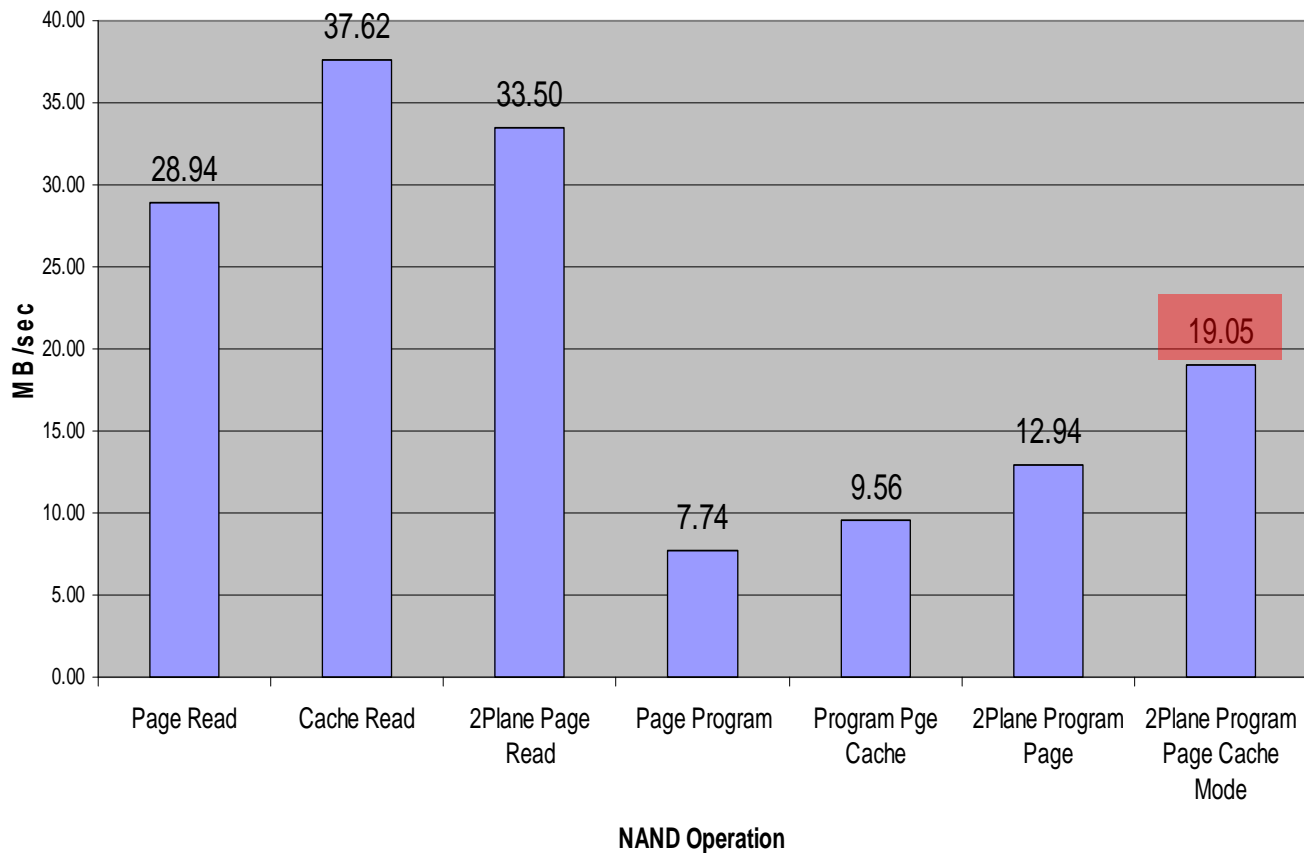
- Device is divided into two physical planes, odd/even blocks
- Users have the ability to:
 - Concurrently access two pages for read
 - Erase two blocks concurrently
 - Program two pages concurrently
- The page addresses of blocks from both planes must be the same during two-plane READ/PROGRAM/ERASE operations

4Gb, Two-Plane, 2K-Page SLC NAND Architecture



4Gb, 2K-Page SLC NAND Performance

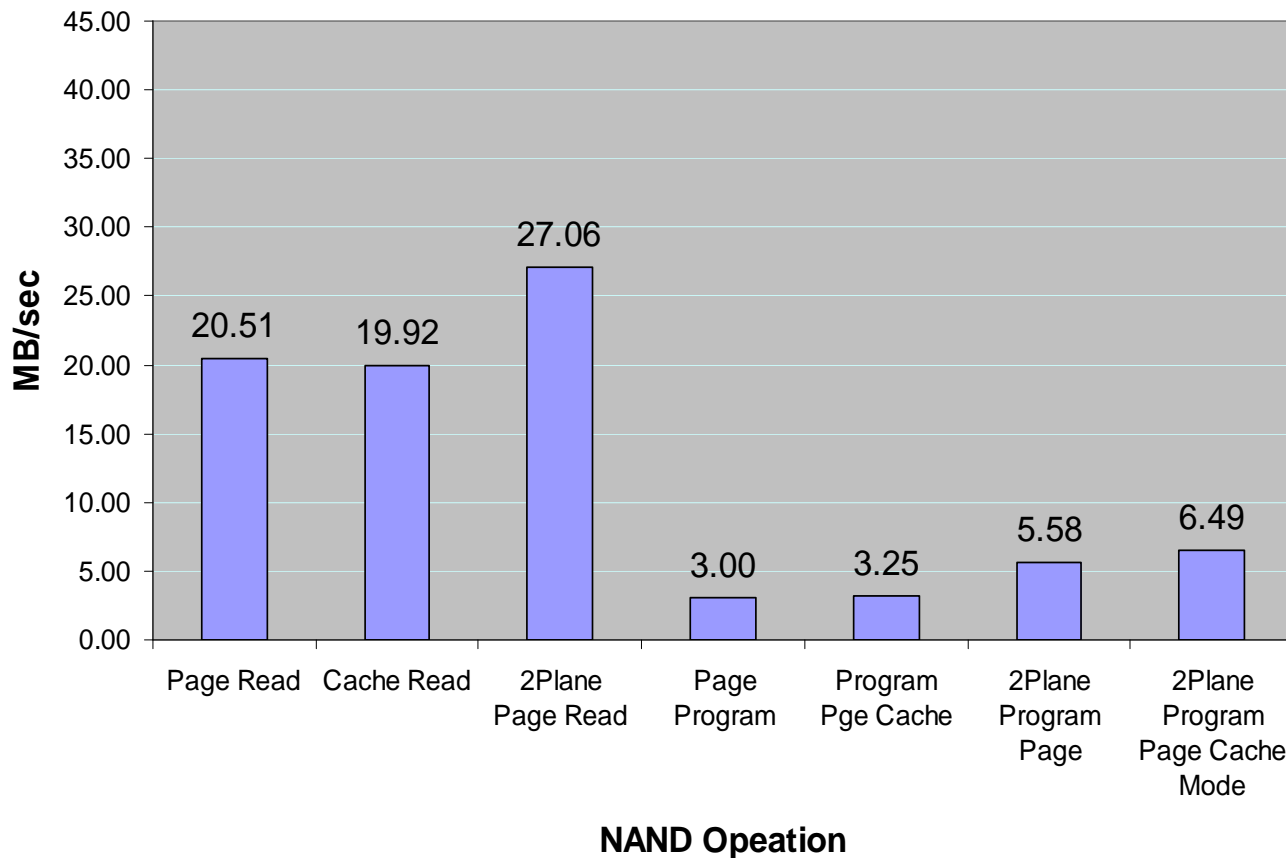
Micron (72nm SLC) 4Gb die 2K Page Performance



Symbol	Time	Units
tR	20	us
tDCBSYR1	3	us
tDCBSYR2	3	us
tRC	25	ns
tRC (C)	25	ns
tRC		ns
tPROG	220	us
tCBSY	3	us
tDBSY	0.5	us
tWC	25	ns
tWC (C)	25	ns
tWC		ns
PS	2112	Byte
NP	64	Pages

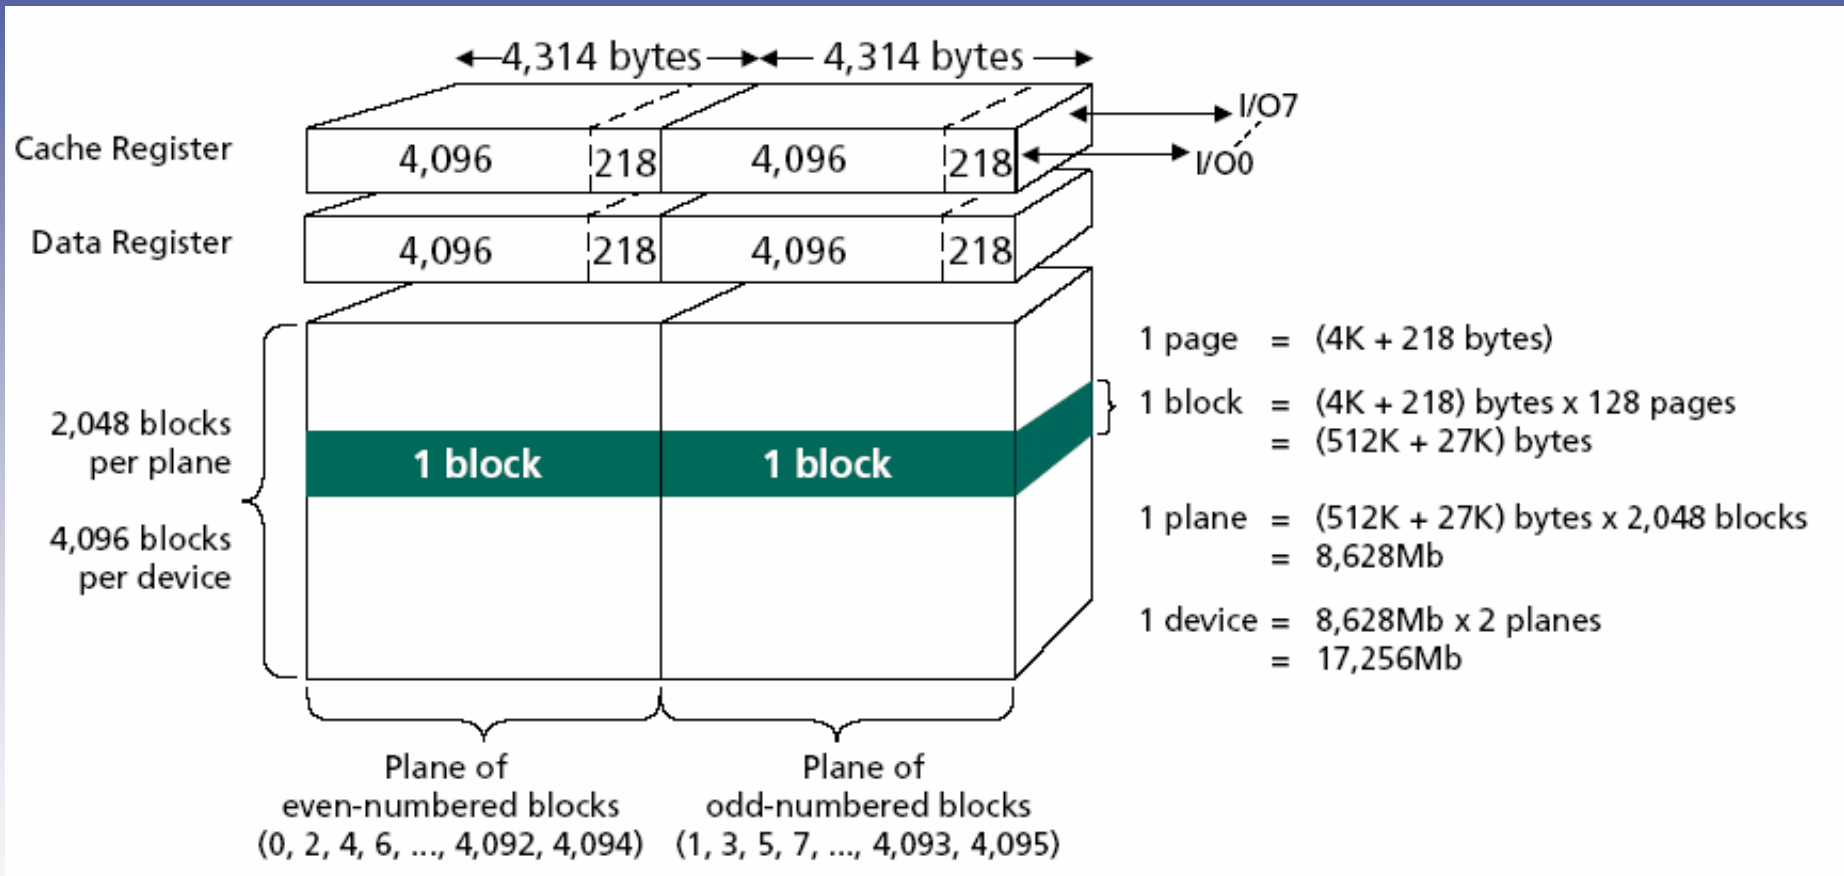
8Gb, 2K-Page MLC Performance

Micron (72nm MLC) 8Gb die 2K Page Performance



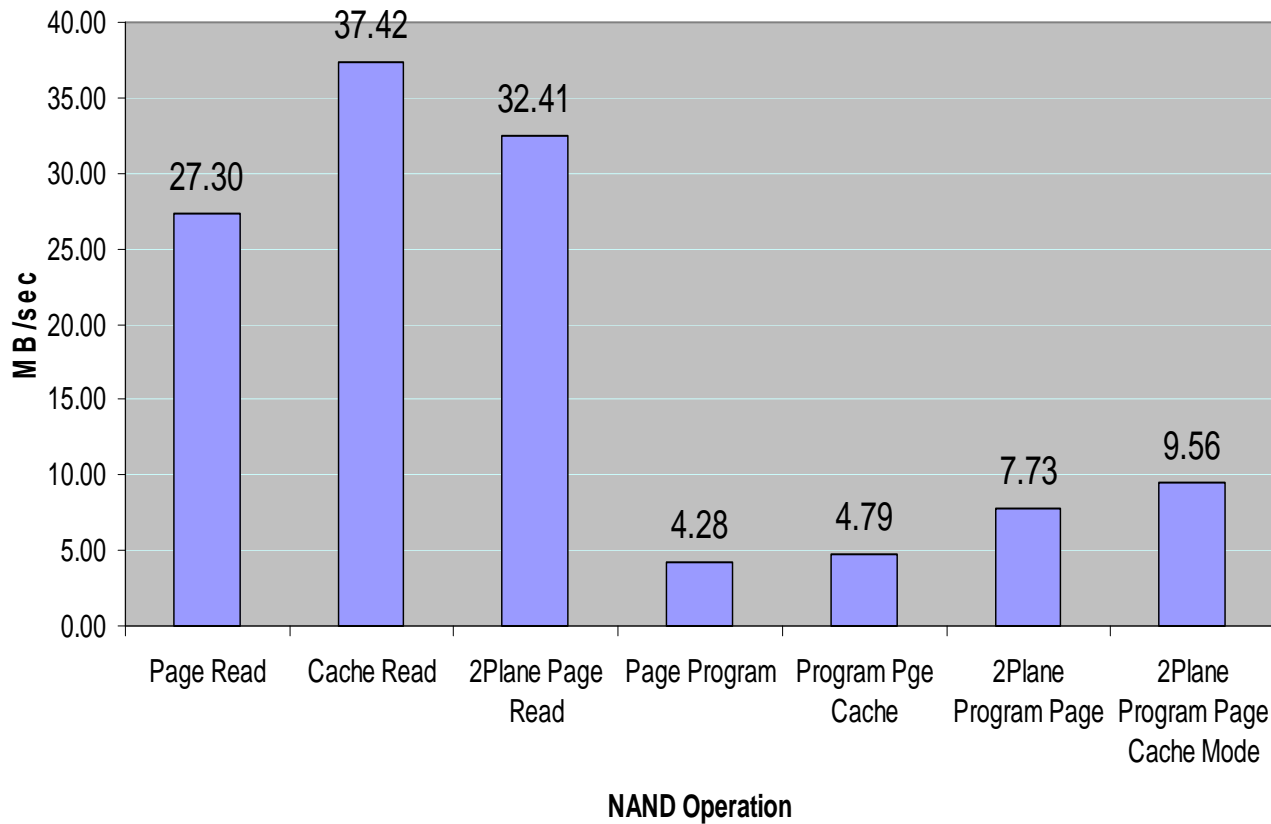
Symbol	Time	Units
tR	50	us
tDCBSYR1	7	us
tDCBSYR2	7	us
tRC	25	ns
tRC (C)	25	ns
tRC	25	ns
tPROG	650	us
tCBSY	30	us
tDBSY	0.5	us
tWC	25	ns
tWC (C)	25	ns
tWC	25	ns
PS	2112	Byte
NP	128	Pages

16Gb, Two-Plane, 4K-Page MLC NAND Architecture



Two-Plane, 4K-Page MLC NAND Architecture

Micron (55nm MLC) 16Gb die 4K Page Performance



Symbol	Time	Units
tR	50	us
tDCBSYR1	7	us
tDCBSYR2	7	us
tRC	25	ns
tRC (C)	25	ns
tRC	50	ns
tPROG	900	us
tCBSY	3	us
tDBSY	0.5	us
tWC	25	ns
tWC (C)	35	ns
tWC	45	ns
PS	4314	Byte
NP	128	Pages

Open NAND Flash Interface



- Future Micron NAND Flash devices support the ONFI specification
- Micron is a founding member of ONFI
- The ONFI 1.0 specification is available at <http://www.onfi.org/>
- ONFI founders:




NAND Error Modes

- Program Disturb
- Read Disturb
- Data Retention
- Endurance

Let's Get Orientated: NAND Architecture

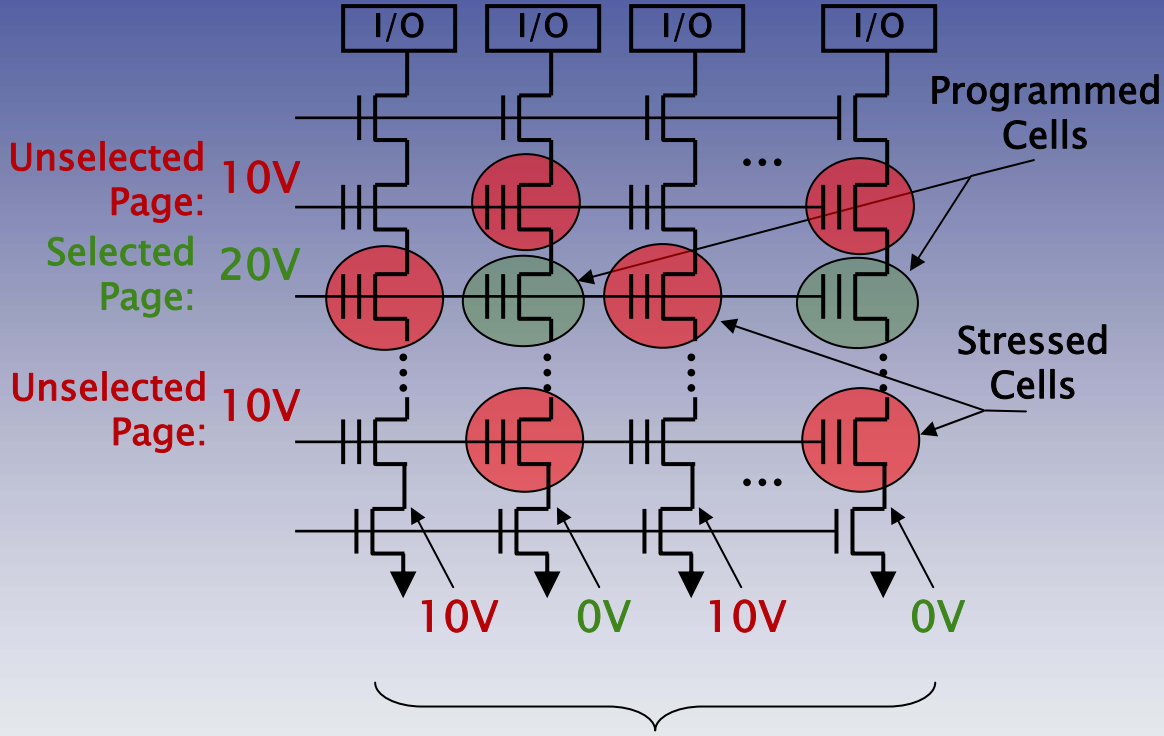
- NAND architecture is based on independent *blocks*
- Blocks are the smallest erasable units
- Pages are the smallest programmable units
 - Partial pages can be programmed in some devices



* Typical for 4Gb SLC

Program Disturb

- Cells *not* being programmed receive elevated voltage stress
- Stressed cells
 - Are always in the block being programmed
 - Can either be on pages *not* selected, or in a selected page but not supposed to be programmed
- Charge collects on the floating gate causing the cell to appear to be weakly programmed
- Does not damage cells; ERASE returns cells to undisturbed levels
- Disturbed bits are effectively managed with error correction codes (ECC)
- Partial-page programming accelerates disturbance



Strings being programmed are grounded; others are at 10V.


Note: Circuit structures and voltages are representative only. Details vary by manufacturer and technology node.

Reducing Program Disturb

- Program pages in a block sequentially, from page 0 to page 63 (SLC) or 127 (MLC)
- Minimize partial-page programming operations (SLC)
- It is mandatory to restrict page programming to a single operation (MLC)
- Use ECC to recover from program disturb errors

Read Disturb

- Cells *not* being read receive elevated voltage stress
- Stressed cells are
 - Always in the block being read
 - Always on pages not being read
- Charge collects on the floating gate causing the cell to appear to be weakly programmed
- Does not damage cells; ERASE returns cells to undisturbed levels
- Disturbed bits are effectively managed with ECC




Note: Circuit structures and voltages are representative only. Details vary by manufacturer and technology node.

Reducing Read Disturb

- Rule of thumb for excessive reads per block between ERASE operations
 - SLC – 1,000,000 READ cycles
 - MLC – 100,000 READ cycles
- If possible, read equally from pages within the block
- If exceeding the rule-of-thumb cycle count, then move the block to another location and erase the original block
- Establish ECC threshold to move data
- Erase resets the READ DISTURB cycle count
- Use ECC to recover from read disturb errors

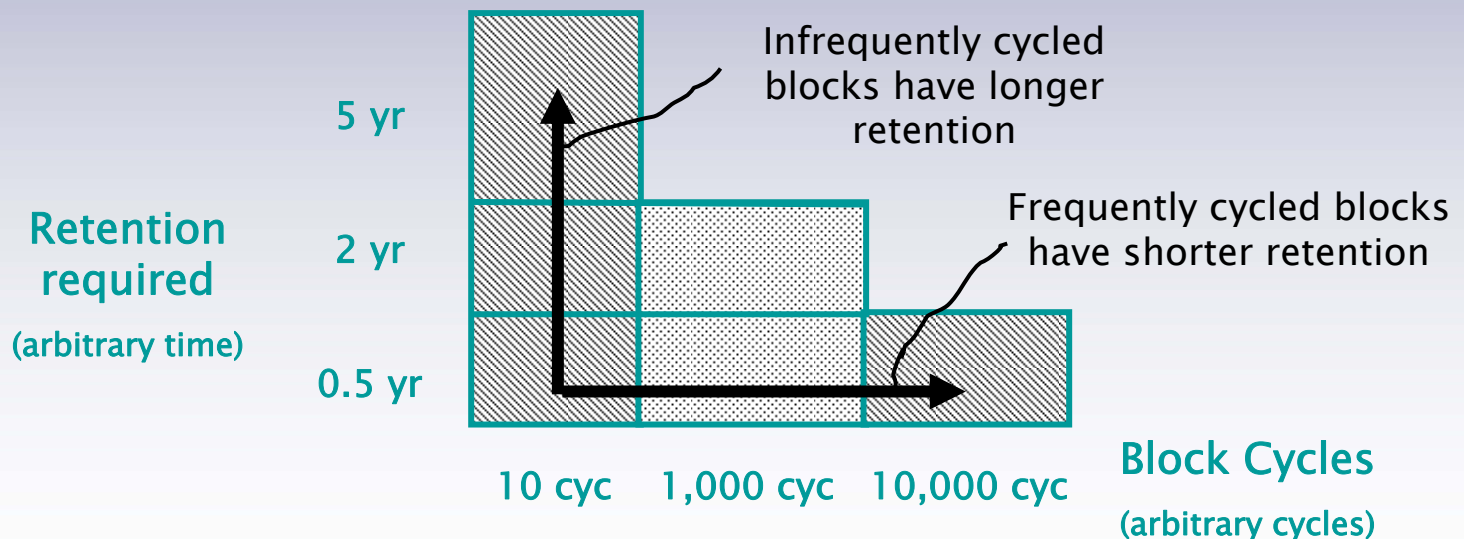
- Charge loss/gain occurs on the floating gate over time; device threshold voltage trends to a quiescent level
- Cell is undamaged; block can be reliably erased and reprogrammed



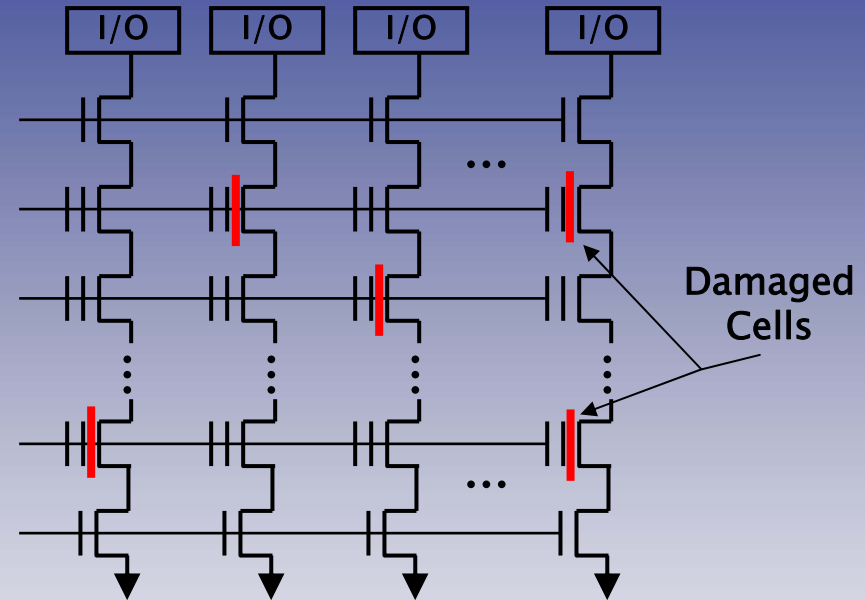
Note: Circuit structures and voltages are representative only. Details vary by manufacturer and technology node.

Improving Data Retention

- Limit PROGRAM/ ERASE cycles in blocks that require long retention
- Limit READs to reduce read disturb



- PROGRAM/ERASE cycles cause charge to be trapped in the dielectric
- Causes a permanent shift in cell characteristics— not recovered by erase
- Observed as failed program or erase status
- Blocks that fail should be retired (marked as bad and no longer used)



Note: Circuit structures and voltages are representative only. Details vary by manufacturer and technology node.

Endurance Recommendations

- Always check pass/fail status (SR0) for PROGRAM and ERASE operations
 - Note: READ operations do not set SR0 to fail status
- If fail status after PROGRAM, move all block data to an available block and mark the failed block bad
- Use ECC to recover from errors
- Write data equally to all good blocks (wear leveling)
- Protect block management/meta data in spare area with ECC

Wear Leveling

- Wear leveling is a plus on SLC devices where blocks can support up to 100,000 PROGRAM/ERASE cycles
- Wear leveling is imperative on MLC devices where blocks typically support less than 10,000 cycles
- If a block was erased and reprogrammed every minute, the 10,000 cycling limit would be exceeded in just 7 days!

$$60 \times 24 \times 7 = 10,080$$

- Rather than cycling the same block, wear leveling involves distributing the number of blocks that are cycled

Wear Leveling (continued)

- An 8Gb MLC device contains 4,096 independent blocks
- Using the previous example, if the cycles were distributed over 4,096 blocks, each block would be programmed less than 3 times (vs. 10,800 cycles if the same block is cycled)
- If perfect wear leveling was performed on a 4,096-block device, a block could be erased and programmed every minute, every day for 77 years!

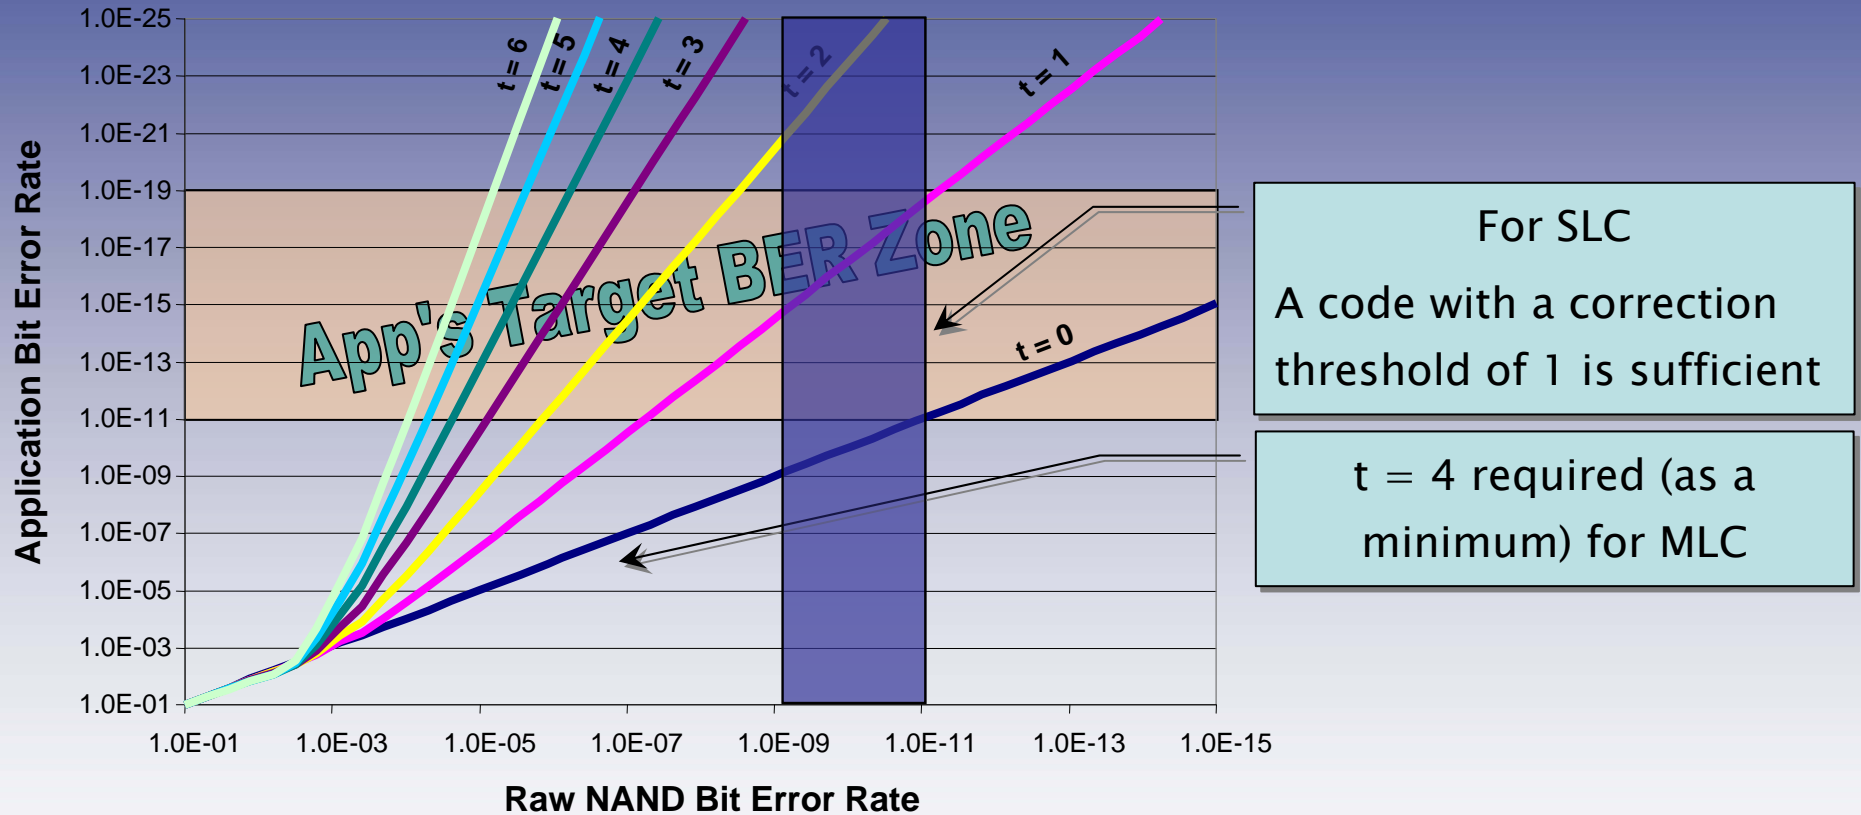
$$\frac{10,000 \times 4,096}{60 \times 24} = \frac{40,960,000}{1,440} = 28,444 \text{ days} = 77.9 \text{ years}$$

- Consider static vs. dynamic wear leveling

ECC Can Fix Everything (well, almost)

- Understand the target data-error rate for your particular system
- Understand the use model that you intend for your system
- Design the ECC circuit to improve the raw-bit error rate (BER) of the NAND Flash, under *your* use conditions, to meet the system's target BER

ECC Code Selection Is Becoming Even More Important



- As the raw NAND Flash BER increases, it becomes more important to match the ECC to the application's target BER

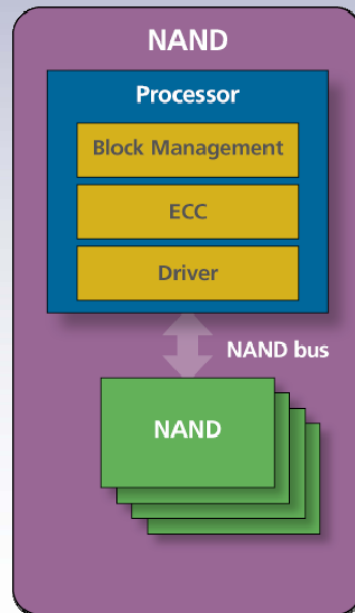
Another Option: Embedded MMC (eMMC)

- The complexities of future MLC require increased attention; the ECC algorithm, for example, is becoming more and more complex, moving from 4+ bits to 8+ bits in the future
- A managed interface addresses the complexities of current and future NAND Flash devices; this means the host does not need to know the details of NAND Flash block sizes, page sizes, planes, new features, process generation, MLC vs. SLC, wear leveling, ECC requirements, etc.
- Embedded MMC (eMMC) is the next logical step in the NAND Flash evolution for embedded applications because it turns a program/erase/read device with bad blocks and bad bits (NAND Flash) into a simple write/read memory


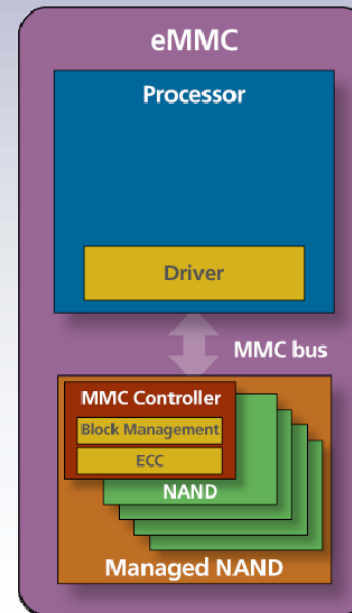
Micron Solutions: Managed NAND (eMMC)

- MLC NAND + MMC 4.2 version controller device
- High-speed solution:
 - Host-selectable x1, x4, and x8 I/Os
 - 52 MHz clock speed (MAX) – 416 Mb/s data rate (MAX)
- Fully backward compatible with previous MMC systems
- ECC, wear leveling, and block management

Direct NAND Interface



Managed NAND



Conclusions

- NAND Flash is the lowest cost, nonvolatile memory available today
- Complexities of MLC NAND require increased logic and software design
- All of these complexities are addressed through the use of the controller included with eMMC



Thank You

©2007 Micron Technology, Inc. All rights reserved. Products are warranted only to meet Micron's production data sheet specifications. Information, products and/or specifications are subject to change without notice. All information is provided on an "AS IS" basis without warranties of any kind. Dates are estimates only. Drawings not to scale. Micron and the Micron logo are trademarks of Micron Technology, Inc. All other trademarks are the property of their respective owners.

Santa Clara, CA USA
August 2007