Jin-Soo Kim
(*jinsoo.kim@snu.ac.kr*)

Systems Software &
Architecture Lab.

Seoul National University

Spring 2024

Flash Memory

Flash Memory Basics

- Two states based on the presence of electrons

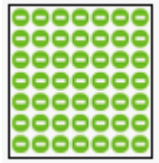

0 = Electrons present

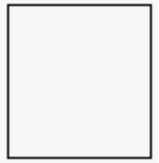

1 = No electrons

- Challenges
  - How to attract or expel electrons?
  - How to find out whether there are electrons or not?
  - How to keep electrons without any power?

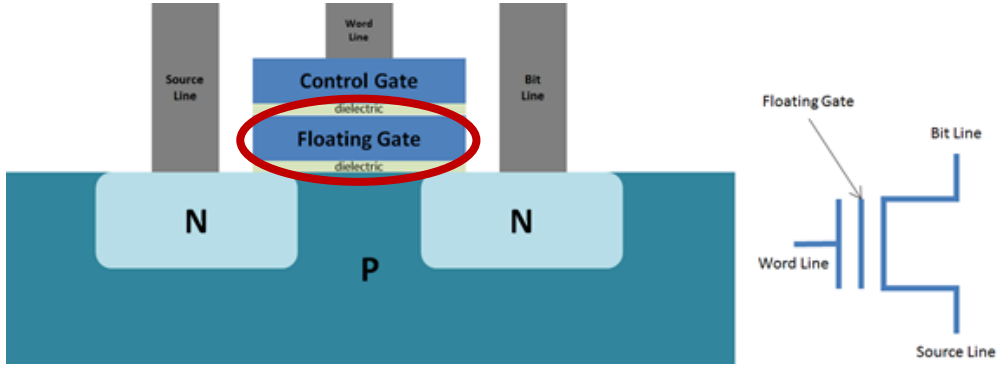

Flash Memory Characteristics

- Erase-before-write
  - Read
  - Write or Program: 1 → 0
  - Erase: 0 → 1
- Bulk erase
  - Read/program unit
    - NOR: byte or word
    - NAND: page
  - Erase unit: block

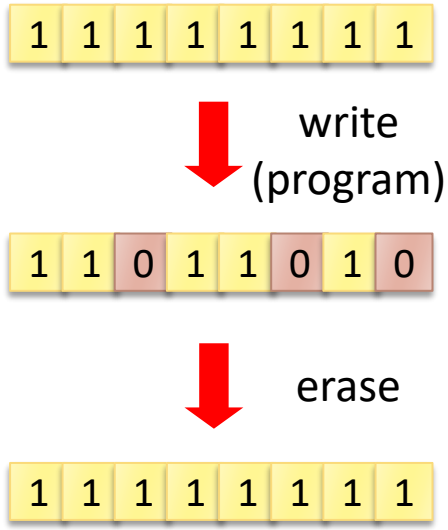

Logical View of NAND Flash

- A collection of blocks
- Each block has a number of pages
- The size of a block or a page depends on the technology (but, it's getting larger)

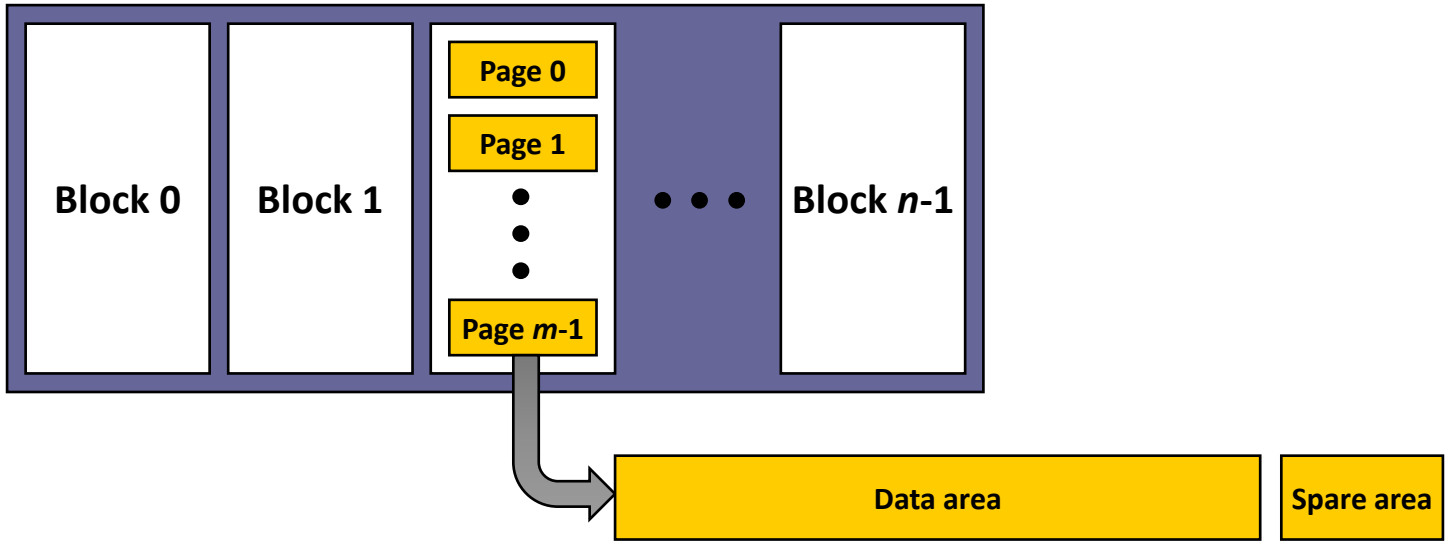

Plane

- Each plane has its own page register and cache register
- Pages can be programmed or read at once
- Optional feature: 1, 2, 4, 8, … planes

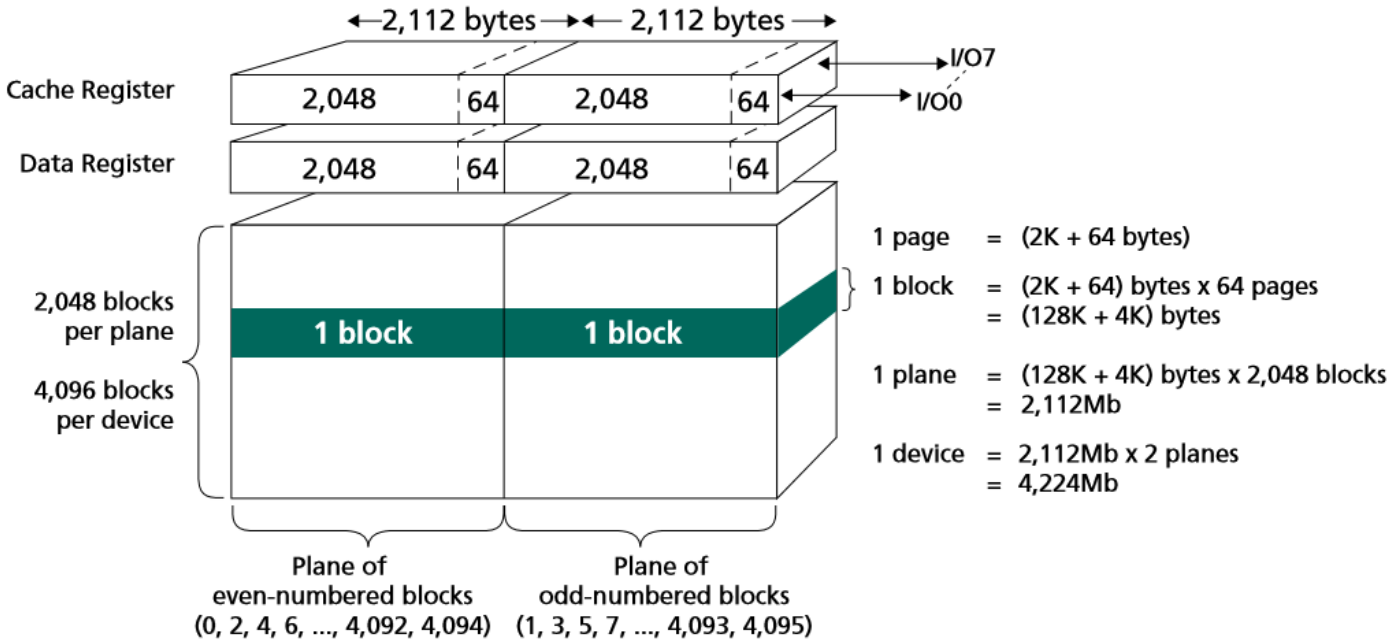

Die / Chip

- Each chip has multiple dies (can be stacked)
- + extra circuits, chip enable signal, ready/busy signal

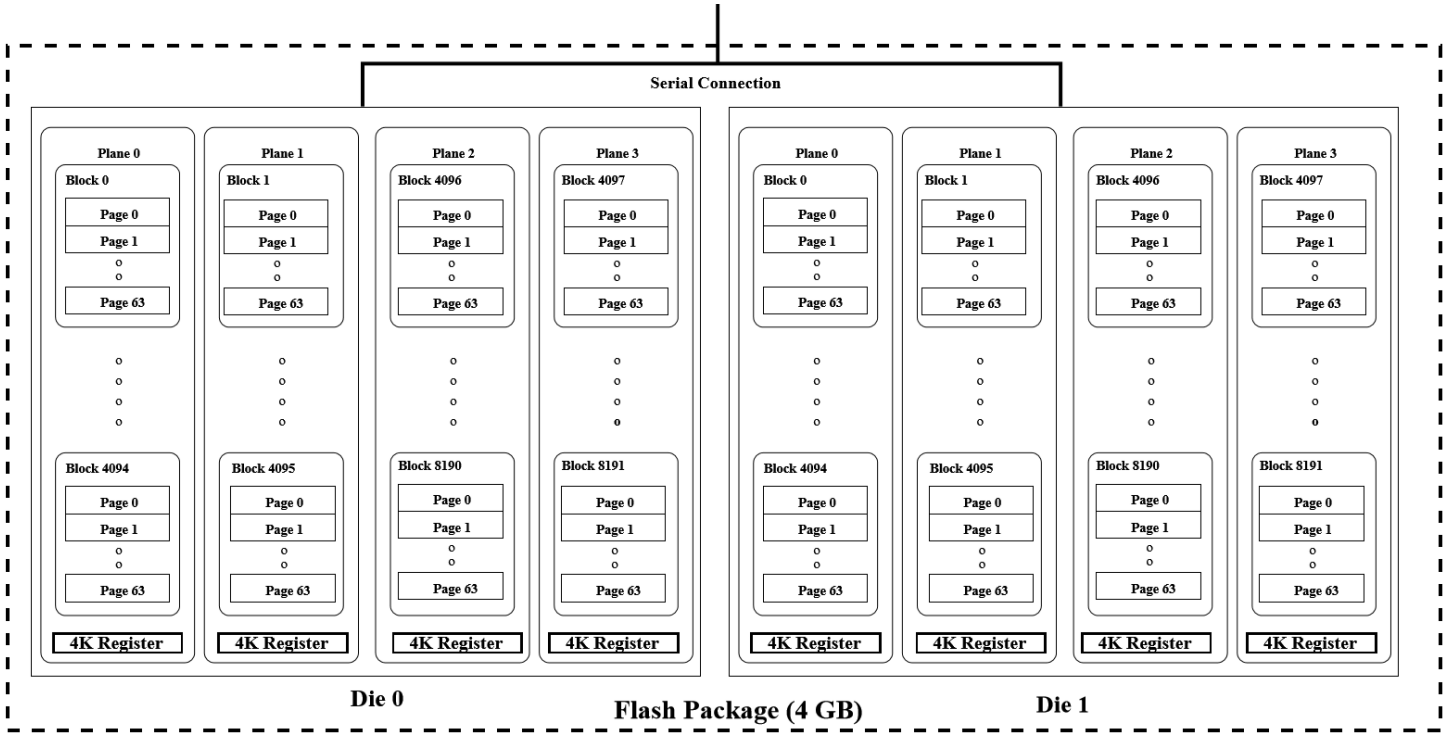

N. Agrawal, Design Tradeoffs for SSD Performance, USENIX ATC, 2008.

NAND Flash Types

- SLC NAND
  - Single Level Cell (1 bit/cell)
- MLC NAND
  - Multi Level Cell (2 bits/cell)
- TLC NAND
  - Triple Level Cell (3 bits/cell)
- QLC NAND
  - Quad Level Cell (4 bits/cell)

- 3D NAND (or V-NAND)

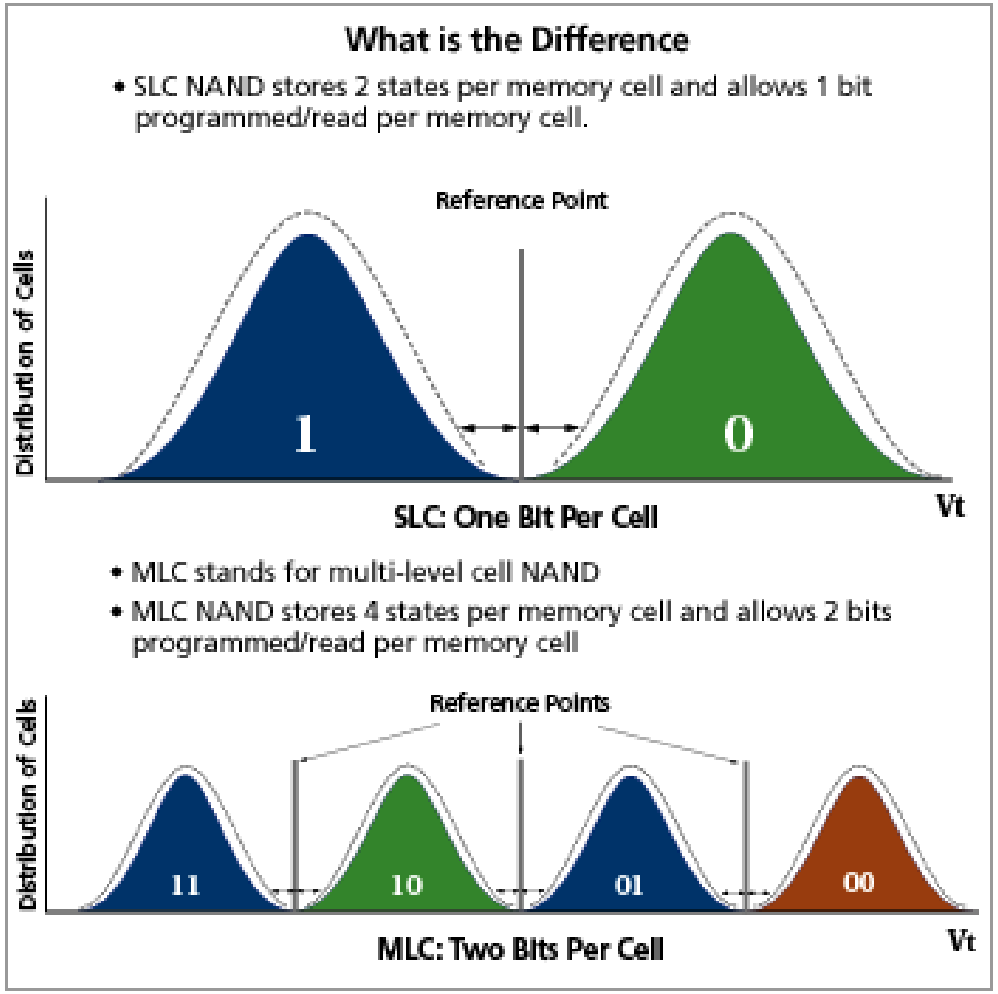

Source: Micron Technology, Inc.

Characteristics of NAND Flash

Erase-Before-Write

- In-place update (overwrite) is not allowed
- Pages must be erased before new data is programmed
- The erase unit is much larger than the read/write unit
  - Read/write unit: page (4KB, 8KB, 16KB, …)
  - Erase unit: block (64-512 pages)

- What if there are live pages in the block we wish to erase?

Limited Lifetime

- The number of times NAND flash blocks can reliably be programmed and erased (P/E cycle) is limited
  - SLCs: 50,000 ~ 100,000
  - MLCs: 1,500 ~ 5,000
  - eMLCs (Enterprise MLCs): 10,000 ~ 30,000
  - TLCs: < 1,000
  - QLCs: ???
- High voltage applied to cell degrades oxide
  - Electrons are trapped in oxide
  - Break down of the oxide structure
- Requires _____

Flash Endurance

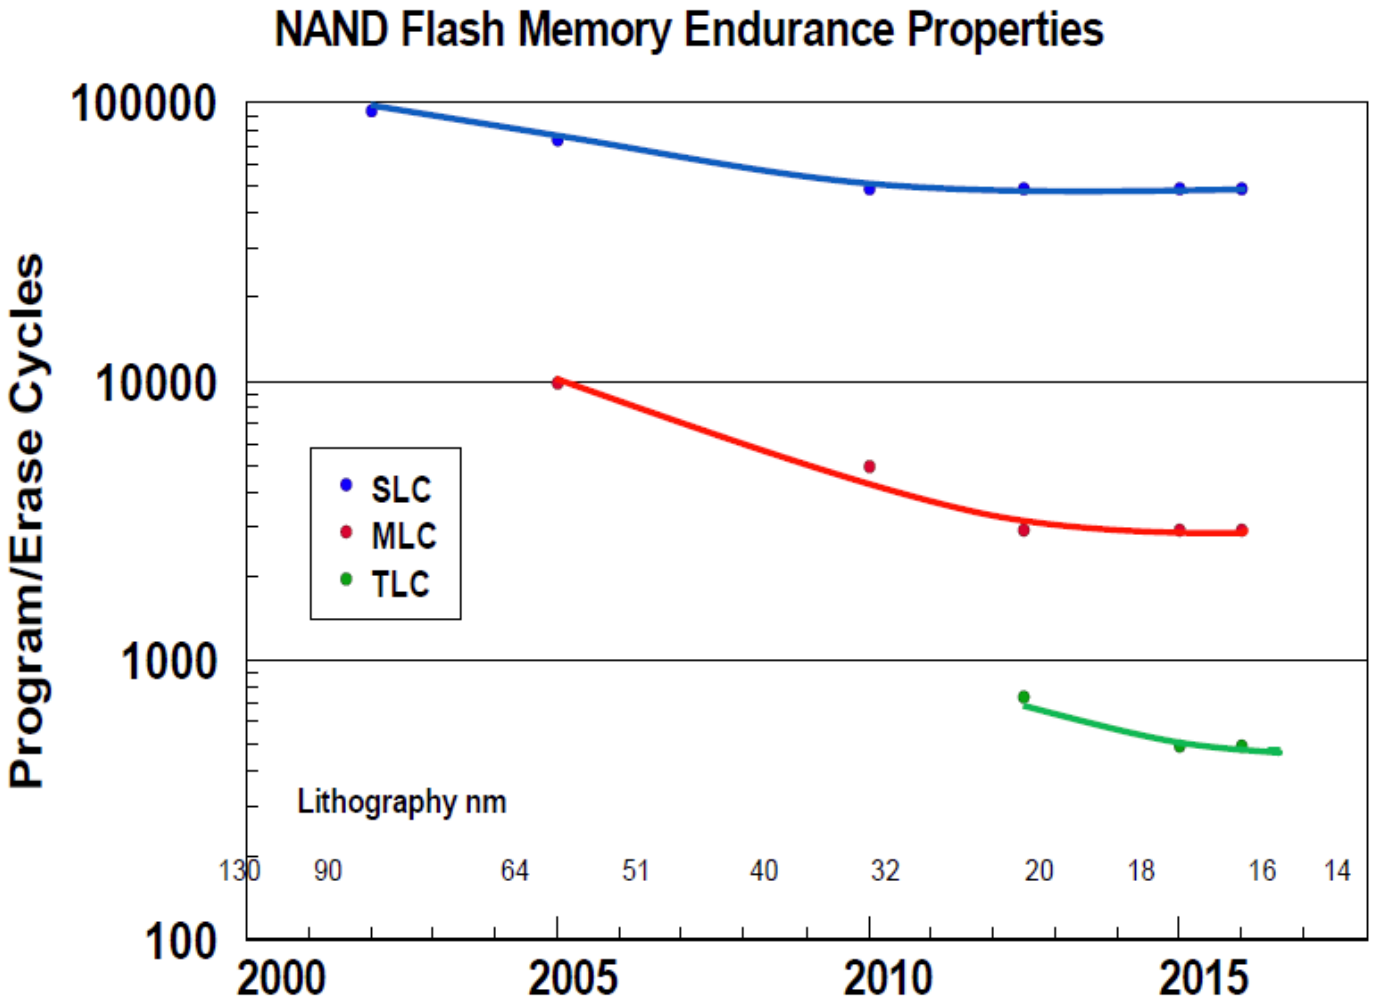

Old Data

E. Grochowski et al., Future Technology Challenges for NAND Flash and HDD Products, FMS, 2012.

Asymmetric Read/Write Latency

- Reading a page is faster than programming it
- Usually more than 10x
  - e.g., 1ynm MLC[1]: Read 45μs, Program 1350μs, Erase 4ms
- Programming a page should go through multiple steps of Program & Verify phases

- As the technology shrinks, read/write latency tends to increase
- MLC and TLC make it even worse

[1] *D. Sharma, System Design for Mainstream TLC SSD, FMS, 2014.*

MLC Programming

- LSB programmed first
  - Cell cannot move to the lower voltage before erase

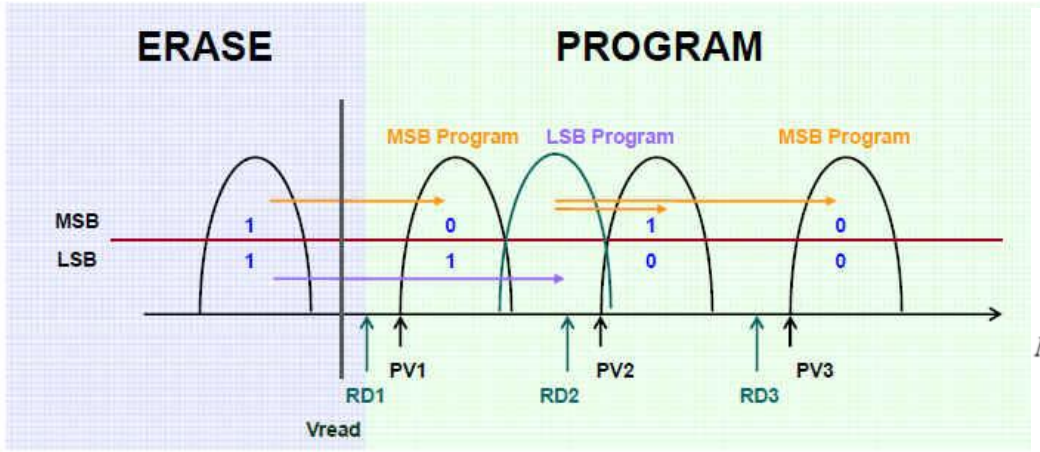

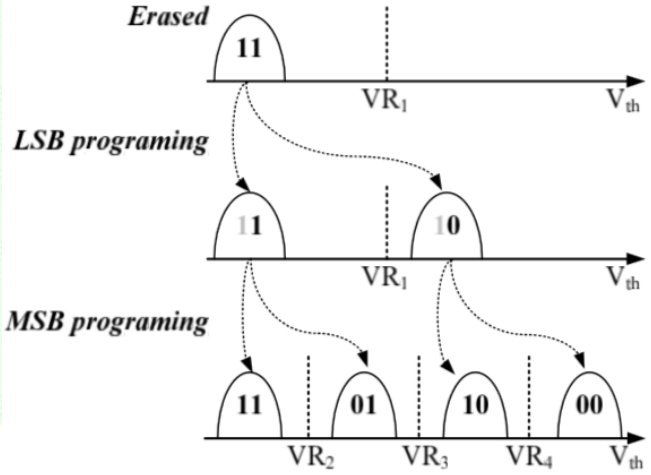

Program : "1"(Erase) → "0"(Program)

LSB Program : 1) Erase → Erase, 2) Erase → LSB

MSB Program: 1) Erase → Erase, 2) Erase → PV1, 3) LSB → PV2, 4) LSB → PV3

Paired Pages in MLC

- One cell represents two or three bits in paired pages
  - LSB: low voltage, fast program, less error
  - MSB: high voltage, slow program, more error
- Performance difference

- LSB page can be corrupted when MSB page programming is interrupted

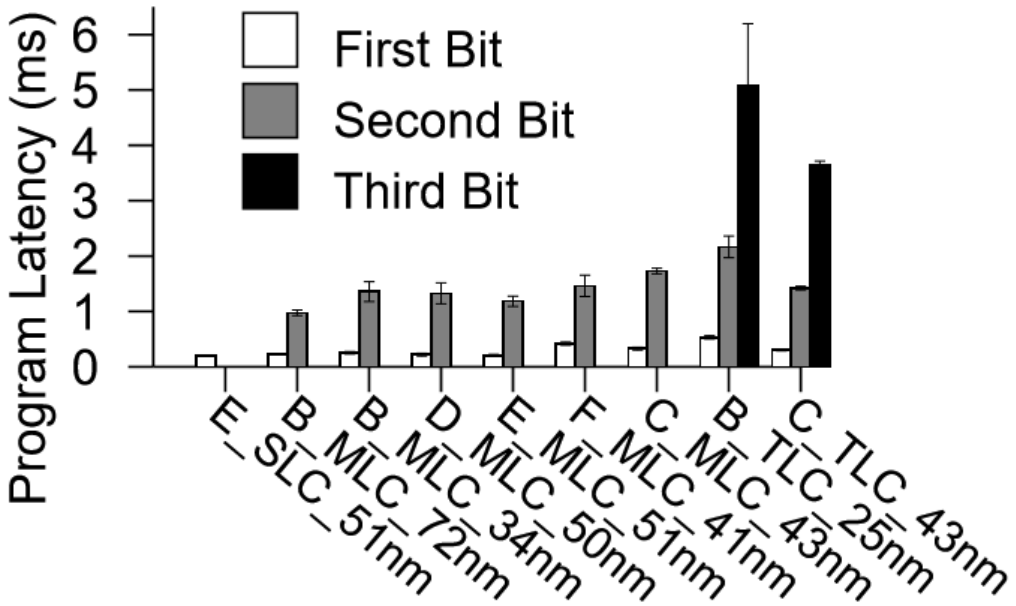

L. M. Grupp et al., The Harey Tortoise: Managing Heterogeneous Write Performance in SSDs, USENIX ATC, 2013.

TLC One-Shot Programming

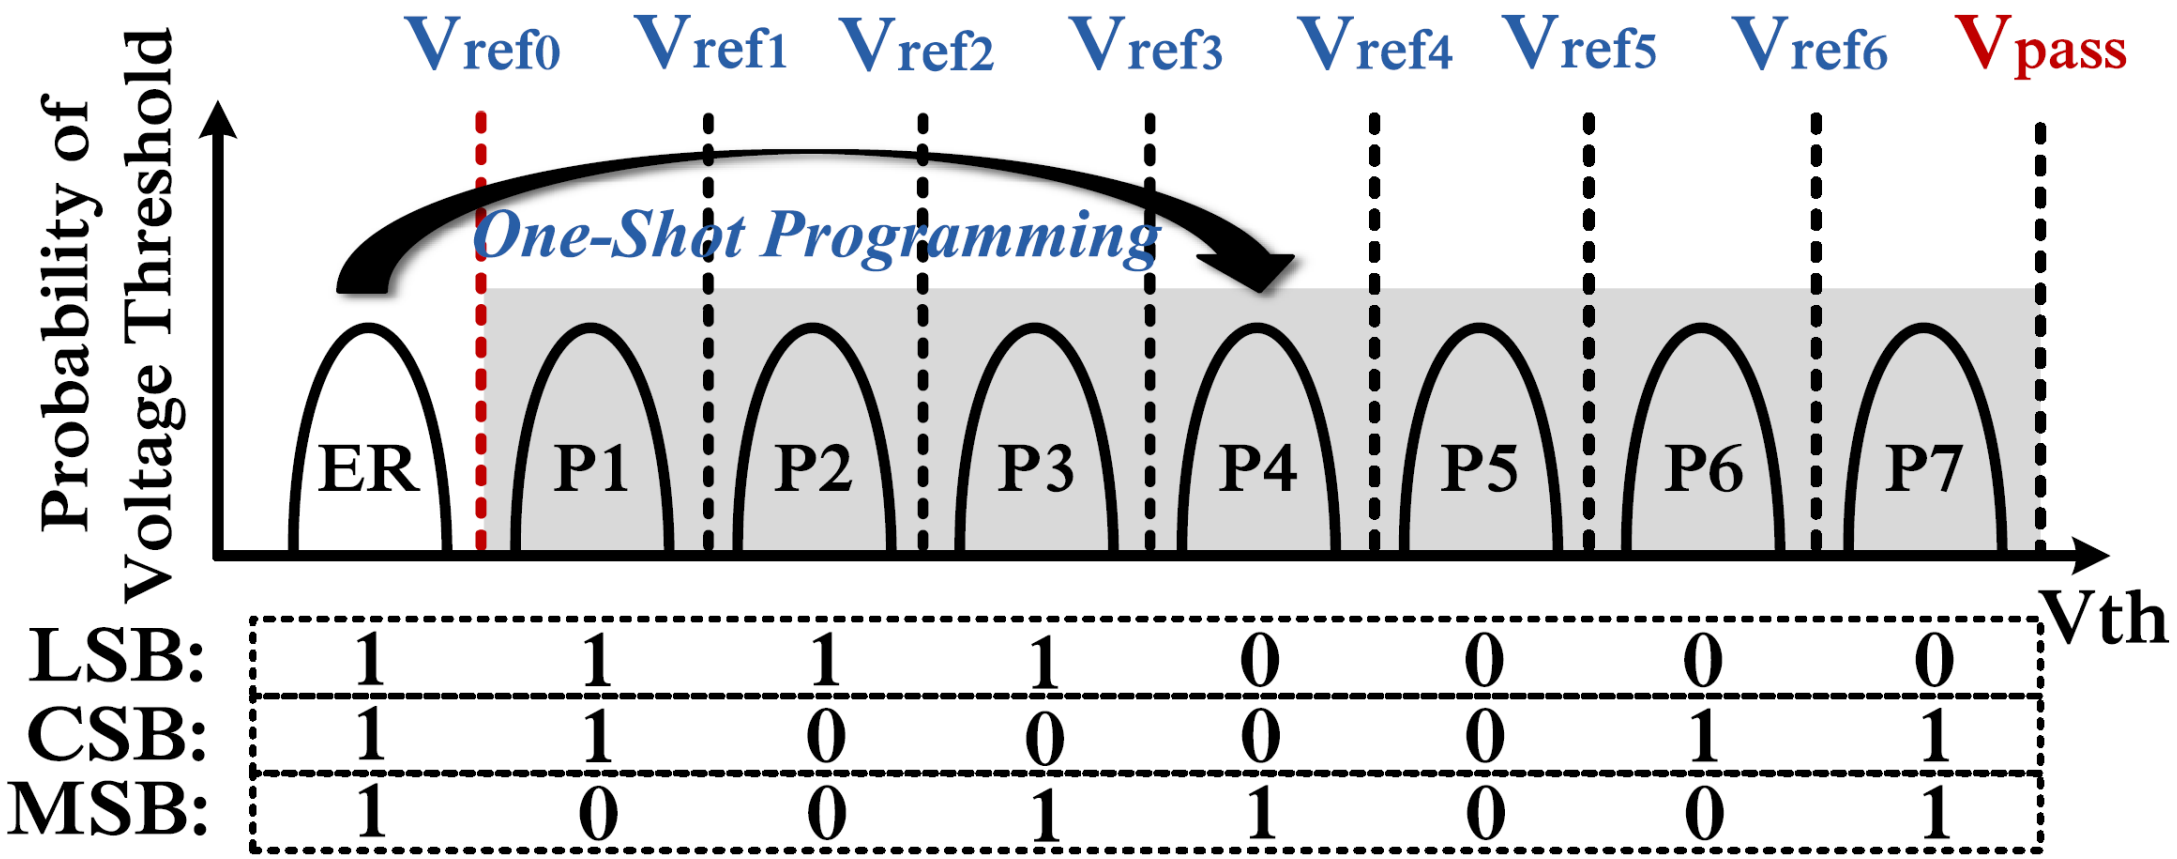

C. Gao et al., "Reprogramming 3D TLC Flash Memory based Solid State Drives", ACM TOS 2022.

Bit Errors

- Bits are flipping frequently
- Error Correction Code (ECC) in spare area

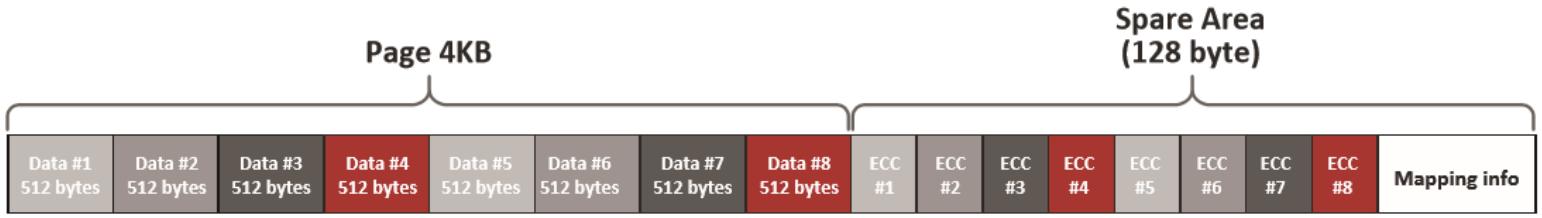

Error Correction Level	Bits Required in the NAND Flash Spare Area		
	Hamming	Reed-Solomon	BCH
1	13	18	13
2	N/A	36	26
3	N/A	54	39
4	N/A	72	52
5	N/A	90	65
6	N/A	108	78
7	N/A	126	91
8	N/A	144	104
9	N/A	162	117
10	N/A	180	130

Source: Micron Technology, Inc.

ECC Requirements

- Endurance continues to deteriorate
- Stronger ECCs are required: RS, BCH, LDPC

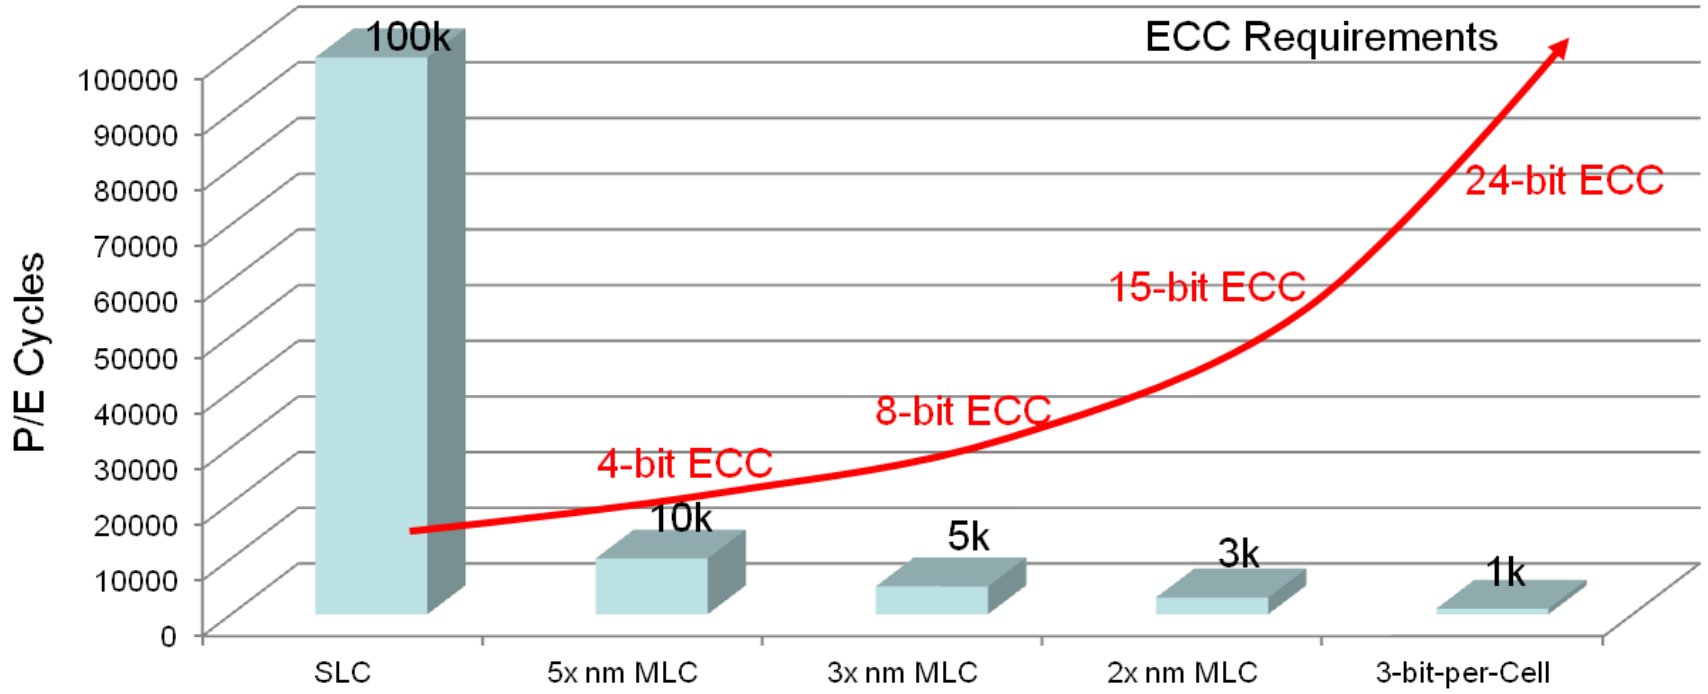

Y. Cai et al., Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis, DATE, 2012.

Reliability

- Write disturbance
  - When a page is programmed, adjacent calls receive elevated voltage stress
- Read disturbance
  - Repeated reading from one page can alter the values stored in other unread pages
- _____ error
  - Threshold voltage shifts down due to charge leakage from the floating gate

Write disturbance

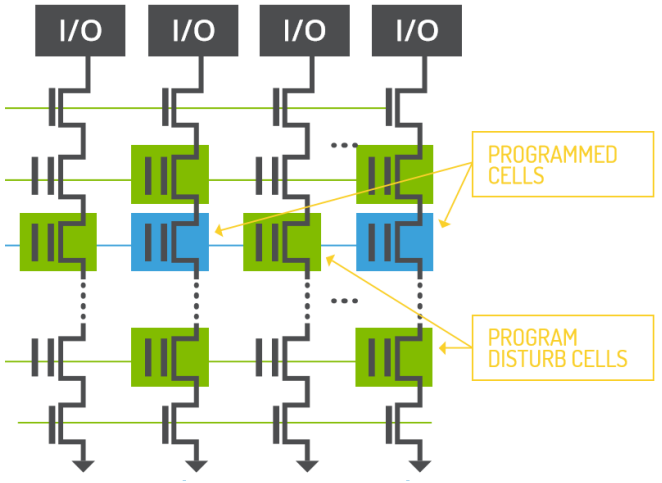

Read disturbance

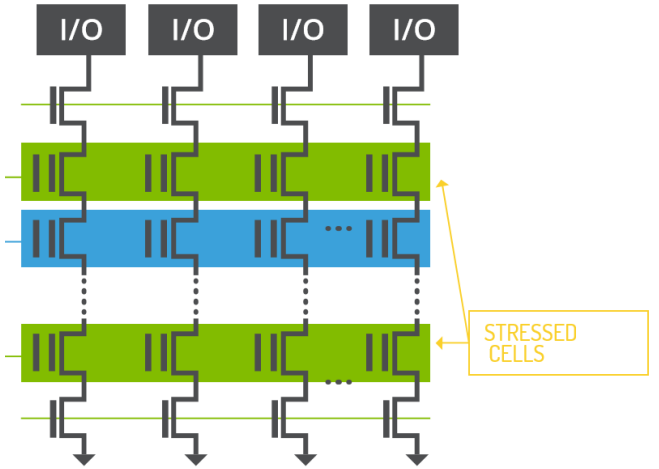

Bad Blocks

- Initial bad blocks
  - Due to production yield constraints and the pressure to keep costs low
  - SLCs: up to 2%
  - MLCs: up to 5%
- Run-time bad blocks
  - Read, write, or erase failure
  - Permanent shift in the voltage levels of the cells due to trapped electrons
- Requires run-time bad block management

Page Programming Constraints

- NOP
  - The number of partial-page programming is limited
  - 1 / sector for most SLCs (4 for 2KB page)
  - 1 / page for most MLCs and TLCs
- Sequential page programming
  - Pages should be programmed sequentially inside a block
  - For large block SLCs, MLCs, and TLCs
- SLC mode
  - Possible to use only LSB pages in MLCs and TLCs
  - Faster and more reliable, higher P/E cycles

Beauty and the Beast

- NAND Flash memory is a beauty
  - Small, light-weight, robust, low-cost, low-power non-volatile device
- NAND Flash memory is a beast
  - Much slower program/erase operations
  - No in-place-update
  - Erase unit > write unit
  - Limited lifetime
  - Bit errors, bad blocks, …
- Software support is essential for performance and reliability!