NAND Flash Media Management Algorithms

Erich F. Haratsch

Seagate

Outline

- NAND Flash Scaling Trends
- ECC
- Hard and Soft Decision Decoding
- Read Voltage Calibration
- Redundant Silicon Elements
- Summary

NAND Scaling Trends

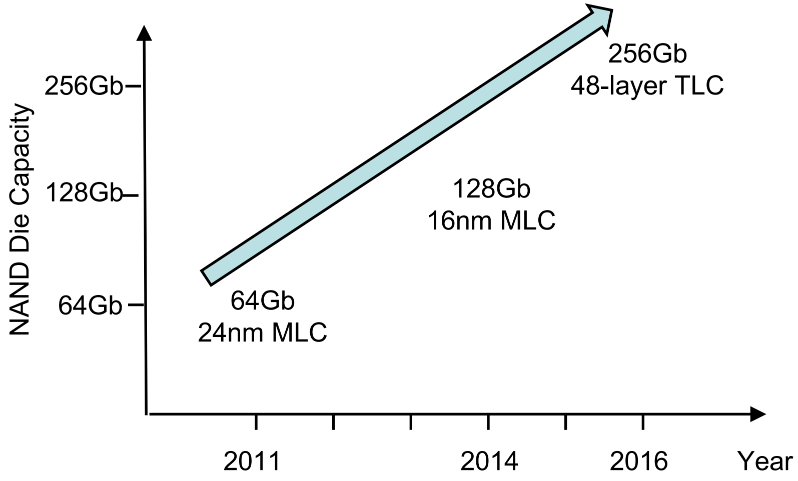

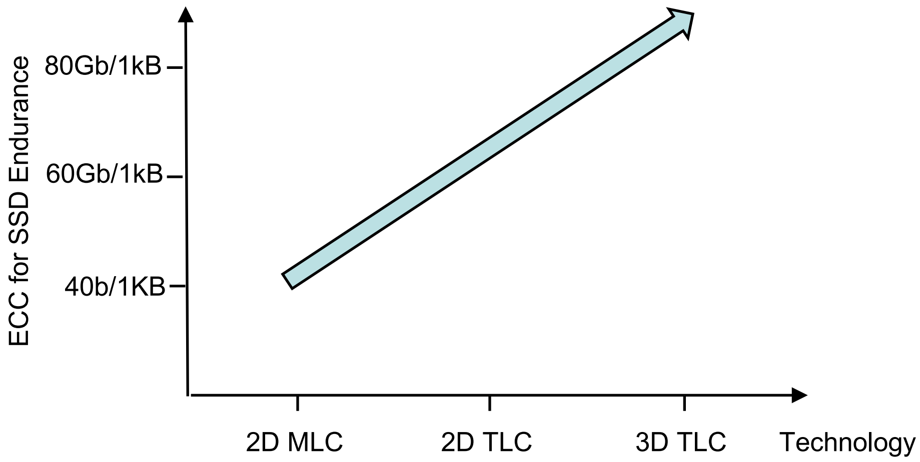

- 3D NAND may extend beyond 100 layers
- 3D NAND extends scaling towards 1Tb die capacity

- Required ECC for SSD-grade endurance exceeds 60b/1kB for 2D TLC
- 3D NAND relies on strong ECC to make TLC mainstream for SSDs

NAND Impairments

Impairment	Effect	Mitigation
Program/Erase Cycling	Voltage shift/widening	ECC Read Voltage Calibration
Retention	Voltage shift/widening	ECC Read Voltage Calibration
Media Defects	Page, block, plane, die failure	Redundant Silicon Elements

- Presented Flash media management algorithms can help to mitigate Read Disturb and Intercell Interference as well

ECC: BCH Codes

- Conventional SSD Controllers use BCH Codes
- BCH codes are algebraic codes, defined by:
  - Code word length
  - Error correction capability per code word
  - For example: 40bit error correction over 1kB code words
- Many SSD controllers implement BCH codes with 1kB code words

ECC: BCH Codes

- BCH codes typically support hard-decision decoding only
- Error recovery by read retry
- Individual hard decision decoding attempts for different read voltages

Reading from Flash: Hard Decision Decoding

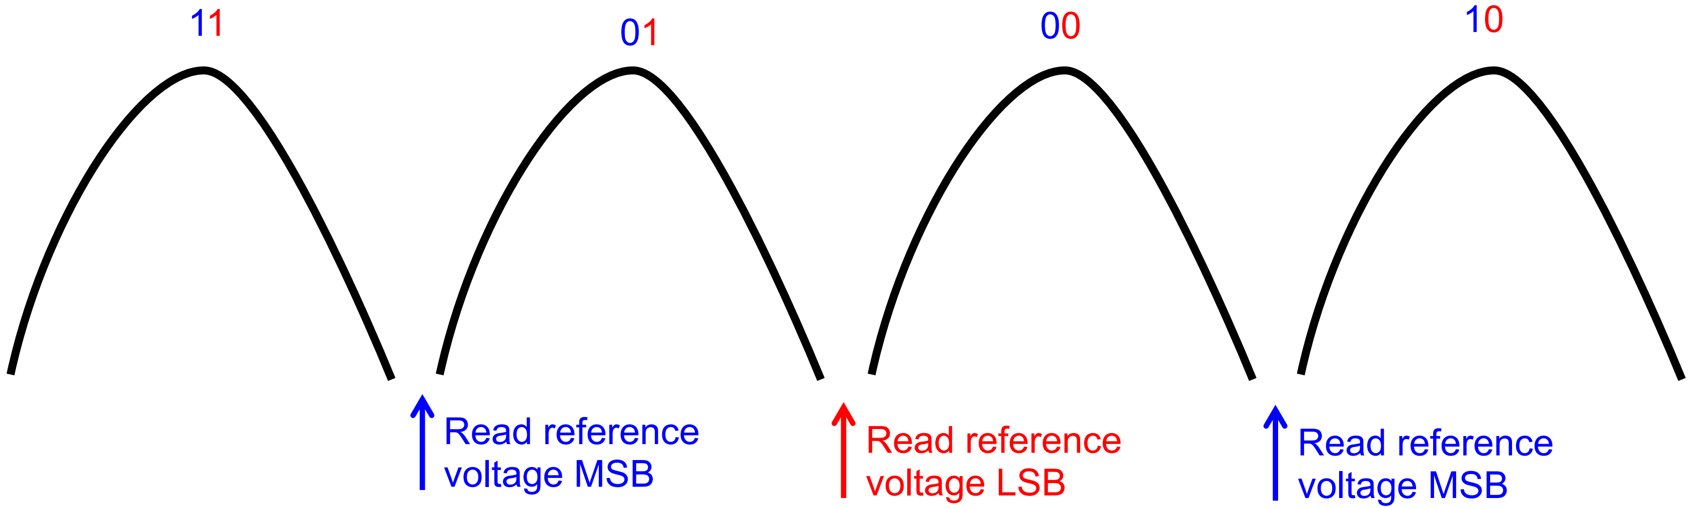

- NAND Flash Memory compares read voltage with read reference voltage to generate hard decision
- One reference voltage for MSB page, 2 reference voltages for LSB page
- Hard decision is used for decoding

Voltage Distribution Shift and Widening

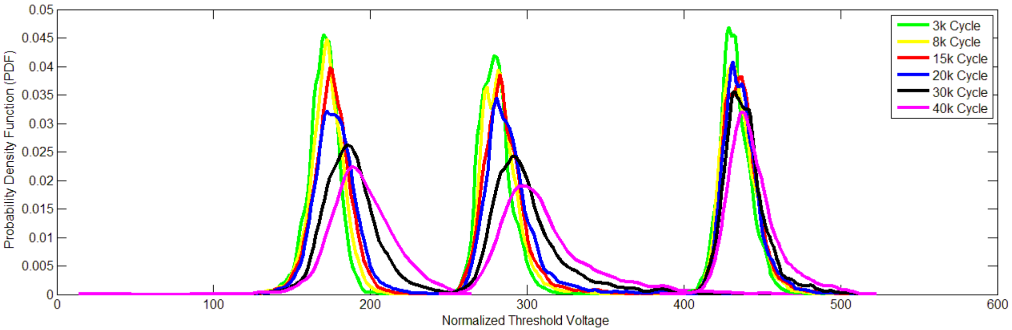

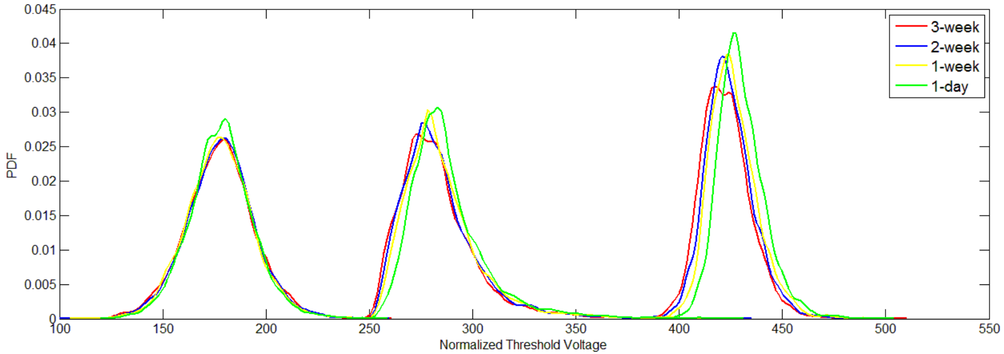

- P/E cycling increases right tails of distributions
- Retention increases left tails of distributions
- Default read reference voltages are misplaced as a result

Read Retry Algorithm

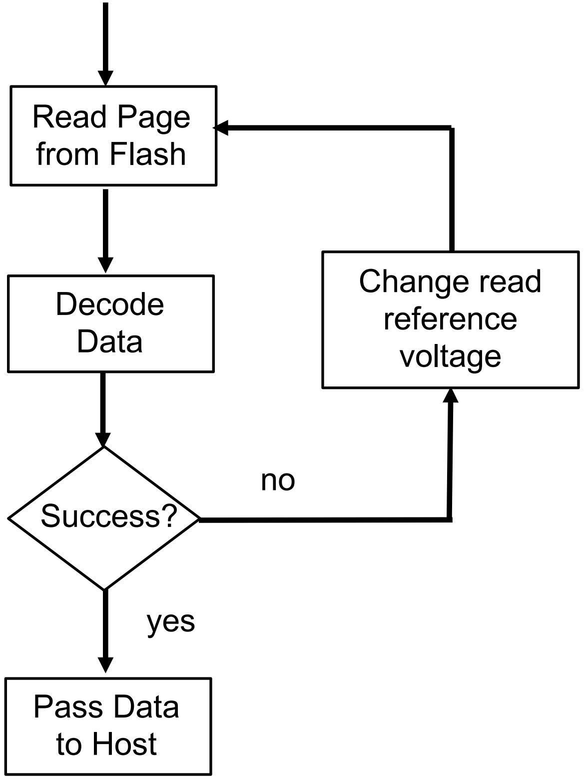

- Default read reference voltage optimized for typical condition
- Read retry algorithm cycles through several individual read decoding steps
- Retry steps use read reference voltages optimized for program/erase cycling, retention, read disturb, etc.

Low-Density Parity Check (LDPC) Codes

- Defined by a sparse (low density) parity check matrix H
- Are represented with a bi-partite graph
- Support hard and soft decision decoding

$$H = \begin{array}{c} \\ \\ \\ \\ \end{array}\begin{array}{cc} & \begin{array}{cccccc} b_1 & b_2 & b_3 & b_4 & b_5 & b_6 \end{array} \\ & \begin{bmatrix} 1 & 1 & 0 & 1 & 0 & 1 \\ 0 & 1 & 1 & 1 & 1 & 0 \\ 1 & 0 & 1 & 0 & 1 & 1 \end{bmatrix} \begin{array}{c} c_1 \\ c_2 \\ c_3 \end{array} \end{array}$$

Bi-Partite Graph:

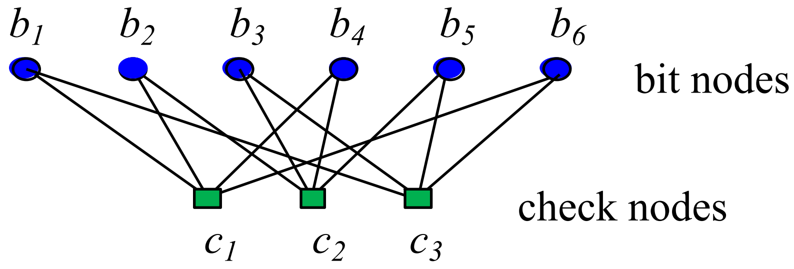

Soft Decision Decoding

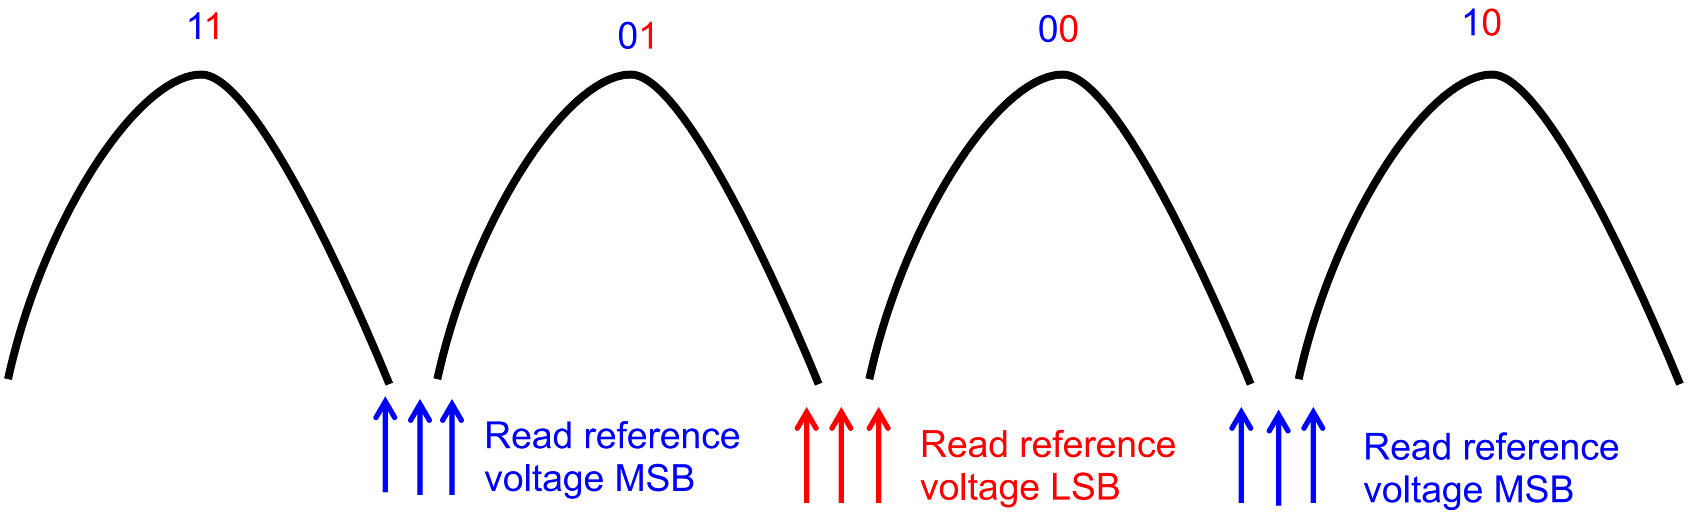

- Multiple read operations with different reference voltages to generate soft decision
- LDPC decoder uses soft decision during error recovery

Hard/Soft LDPC vs. BCH

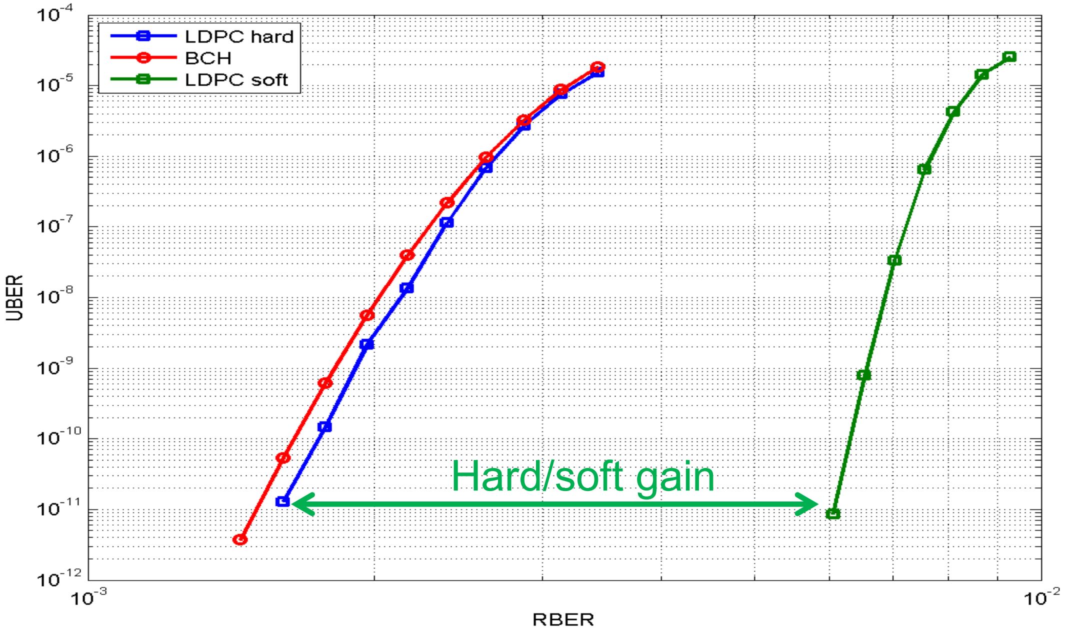

- Soft-decision LDPC decoding has significantly better error correction than BCH decoding

Soft LDPC Levels

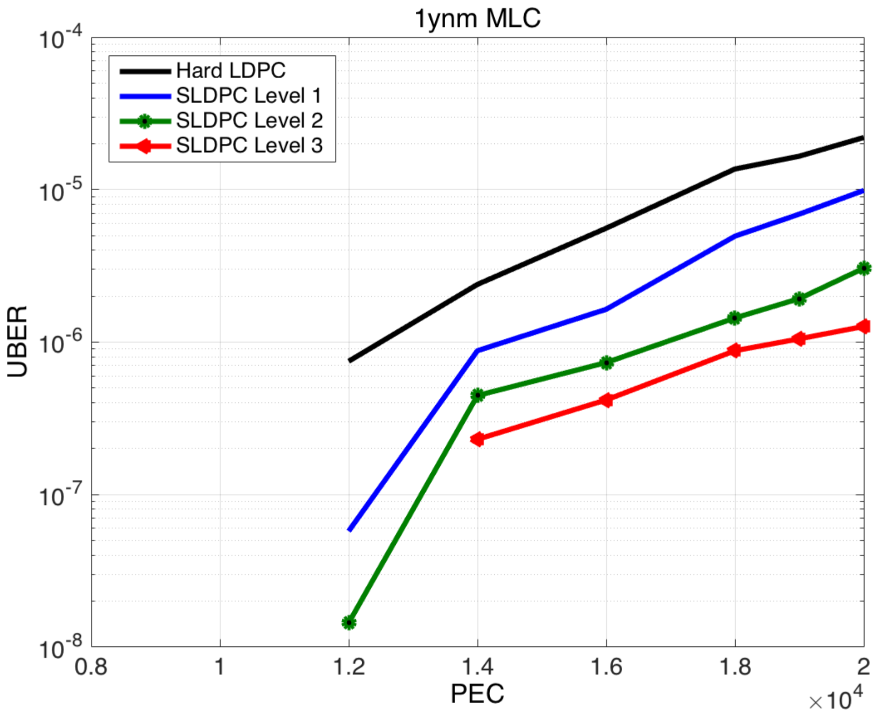

Read voltage placements for soft LDPC:

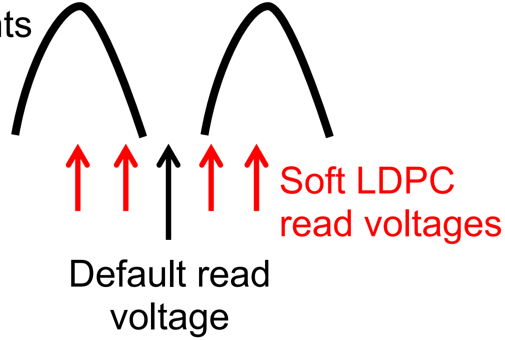

- Sequence of retries with varying read voltage settings
- Computation of soft information (LLRs) based on multiple read decisions

Optimizing LDPC Error Correction

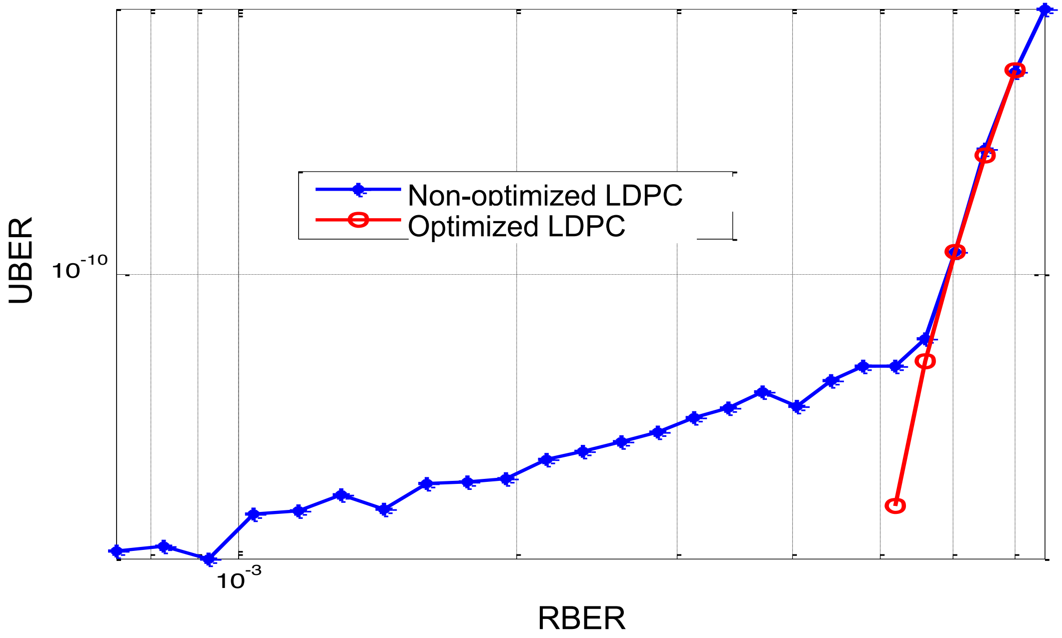

- LDPC code parameters and decoding algorithm need to be optimized for good performance at low error rates

Adaptive Code Rates

- Beginning of Life: use less ECC to increase overprovisioning
- End of life: increase ECC to maintain reliability

Conventional Error Correction:
Stores fixed ECC in spare field

Adaptive ECC (BOL):
Stores ECC in a portion of spare field and increase OP

Adaptive ECC (EOL):
Stores ECC in spare field and uses some of the NAND page

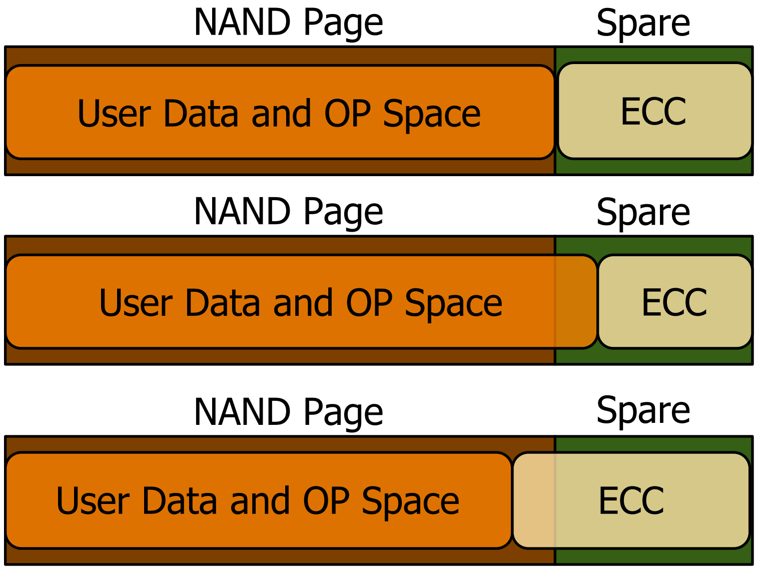

Adaptive ECC allows for more free space @ BOL = More OP and less write amplification

Switching Code Rates

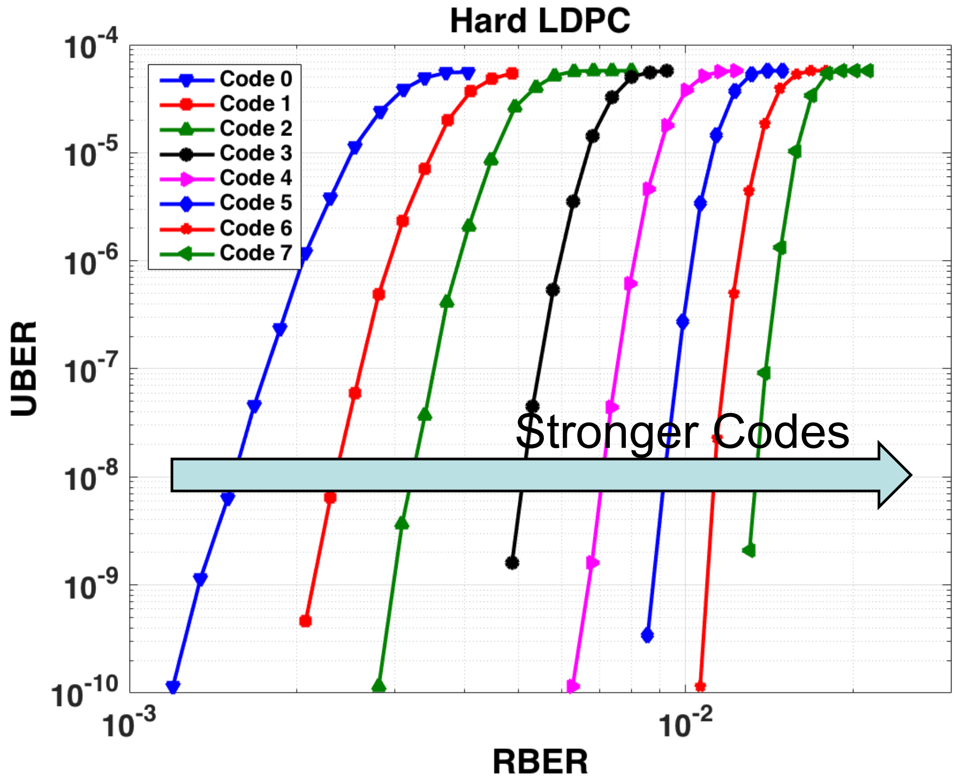

- Multiple LDPC codes cover wide RBER range
- As NAND flash ages, controller switches to the next stronger code
- Read performance improves, since stronger LDPC codes decode data faster

Read Voltage Calibration

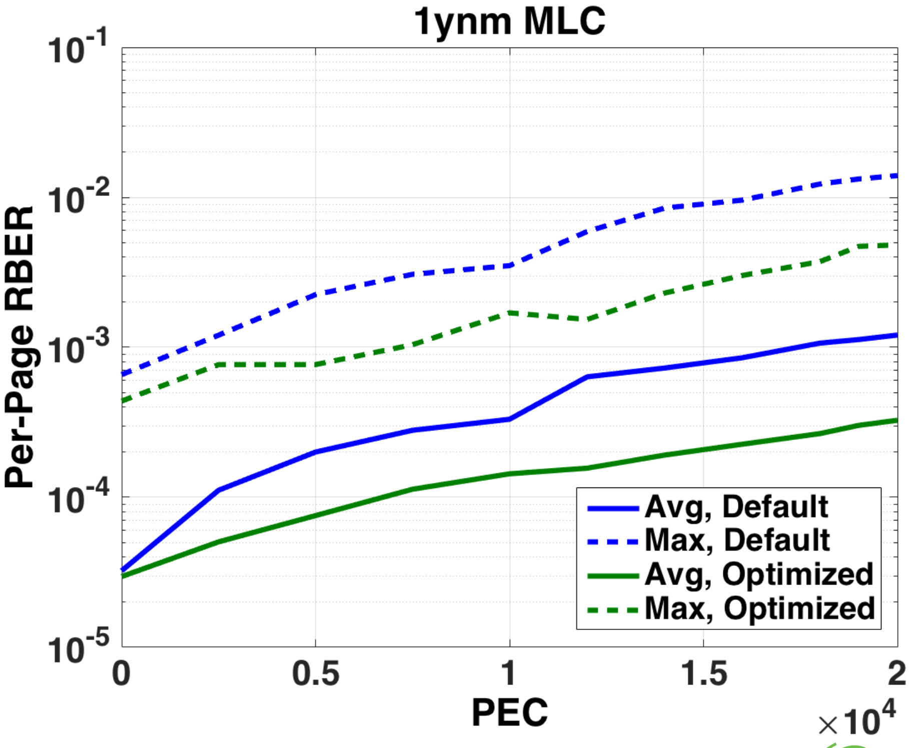

Voltage distributions before/after cycling:

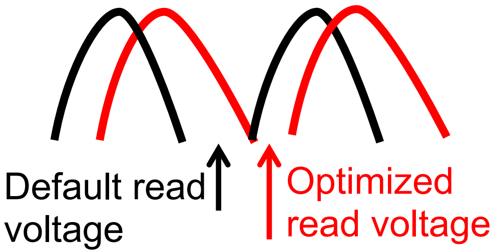

- Optimized read voltages reduce retry rate and extend endurance
- Optimum read voltages shift as a function of endurance, retention and read disturb

Media Failures

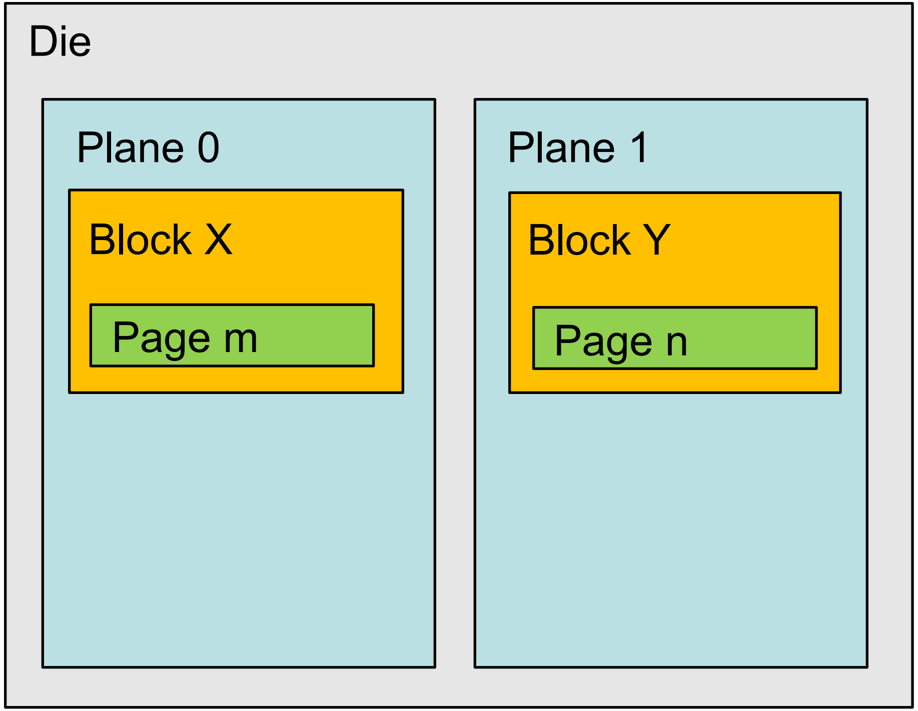

- Pages, blocks, planes or the whole die can fail
- ECC cannot recover data from such catastrophic failures
- Need RAID-like protection inside SSD

RAISE™: Redundant Array of Independent Silicon Elements

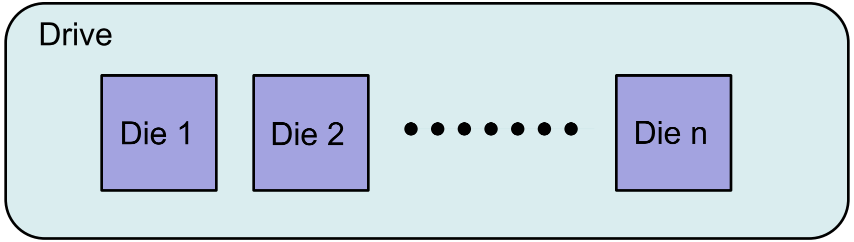

- RAID-like data protection within the drive
- Write data across multiple dies with additional protection
- Corrects full page, block or die failures when all soft LDPC steps fail

SSD Controller: Block Diagram

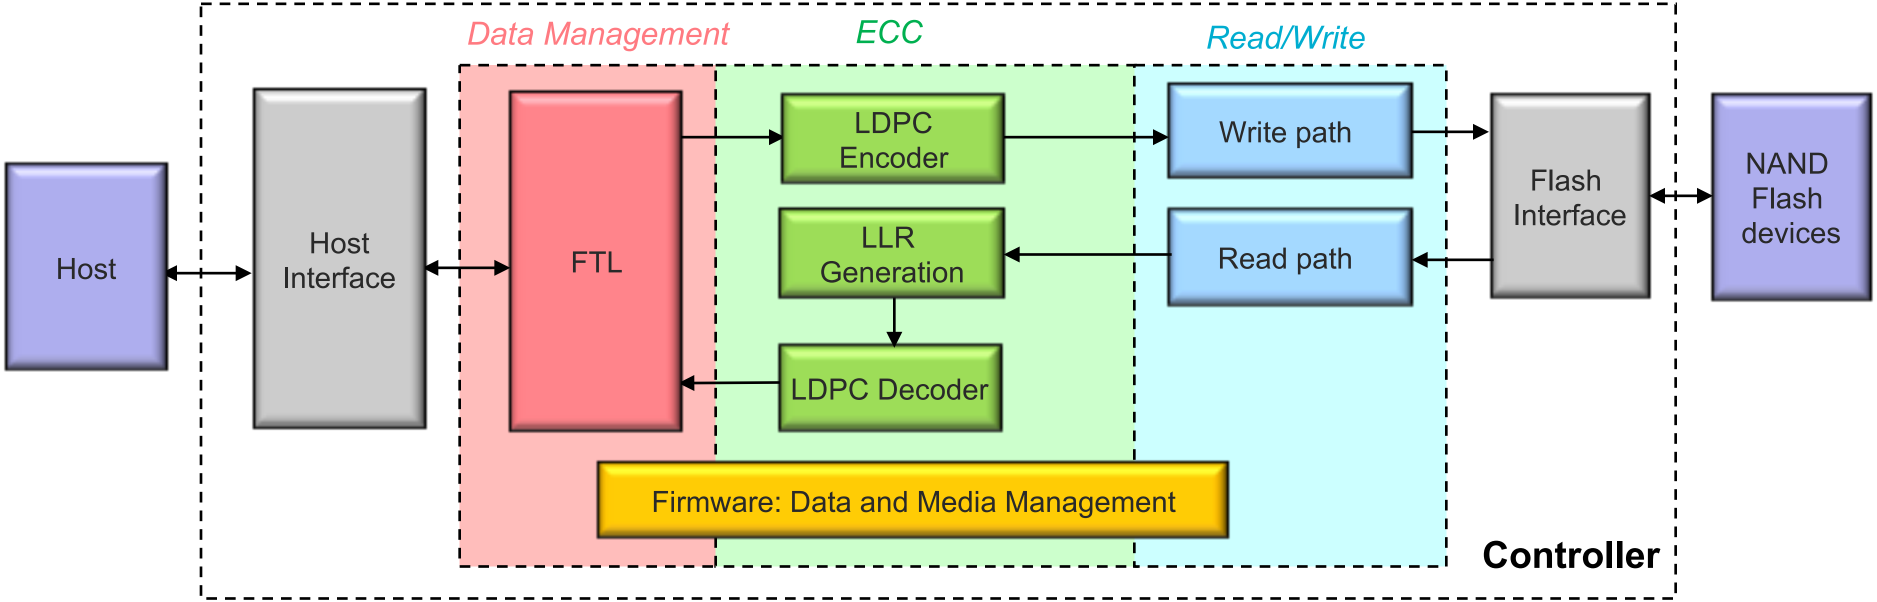

Multi-Level Error Correction

- Hard-decision LDPC decoding is on-the-fly error correction method
- Progressively apply stronger decoding methods such as soft-decision LDPC decoding and signal processing
- Specialized noise handling techniques for P/E cycling, retention, read disturb, etc.
- Optimize time-to-data

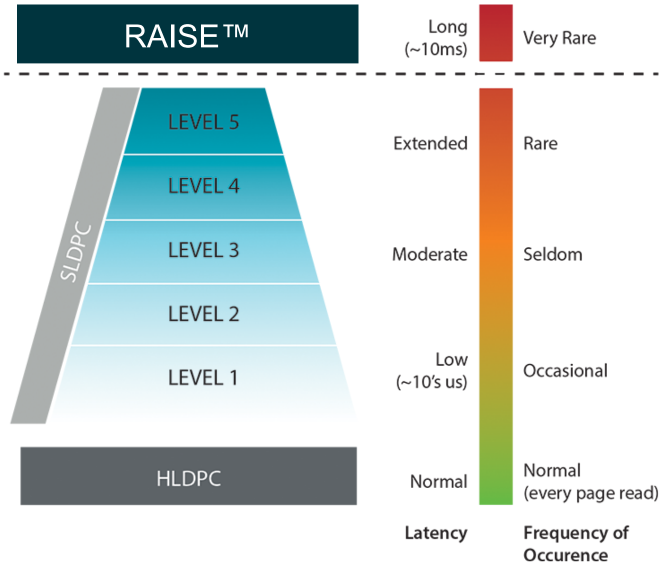

Conclusion

- Latest memory geometries demand intelligent NAND management features
- 3D NAND will still rely on strong ECC and advanced NAND management features to make TLC mainstream for SSD applications

Thank You! Questions?

Visit Seagate Booth #505

Learn about Seagate's ever-expanding portfolio of SSDs, Flash solutions and system level products for every segment

www.seagate.com/flash