Architectural Techniques for Improving NAND Flash Memory Reliability

Thesis Oral

Yixin Luo

Committee:

Onur Mutlu (Chair)

Phillip B. Gibbons

James C. Hoe

Erich F. Haratsch, Seagate

Yu Cai, SK Hynix

Carnegie Mellon

Presented in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Storage Technology Drivers - 2018

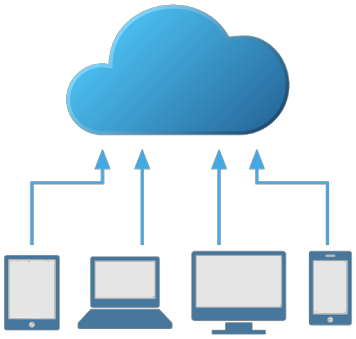

Faster access to **larger amounts** of persistent data

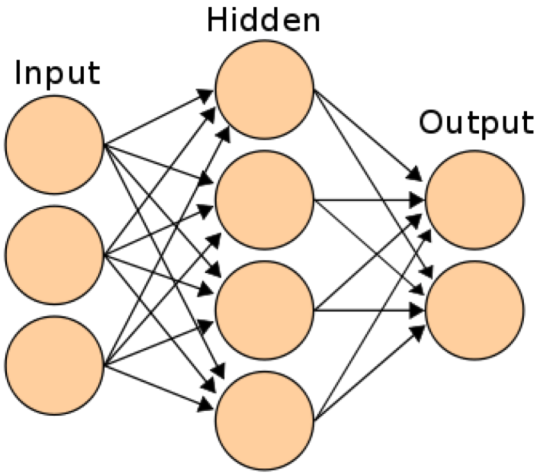

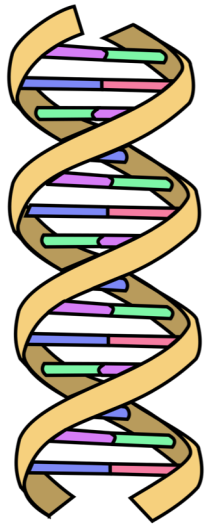

Flash-Memory-Based Solid-State Drive (SSD)

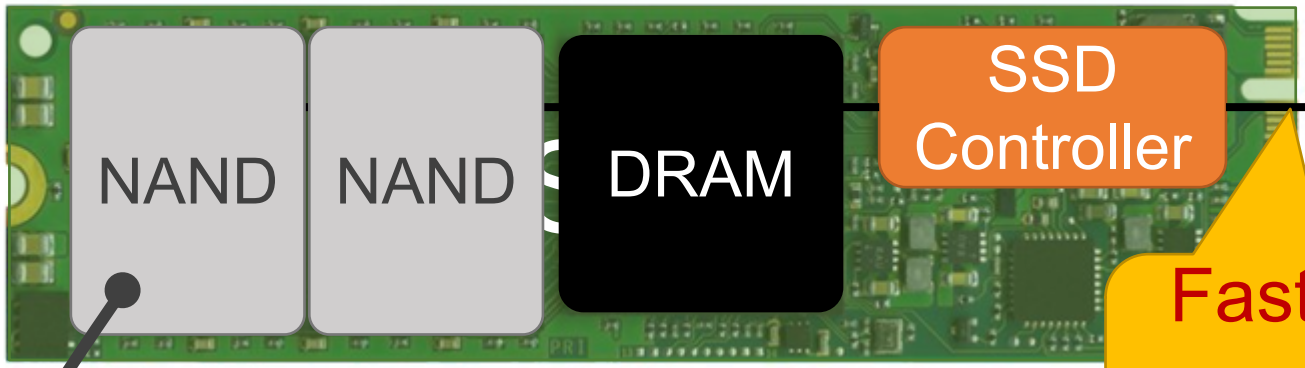

SSD Controller

HOST

Fast access
~50 µs
>100K IOPS

NAND Flash Memory
(Flash Chip)

Large data
Scaling

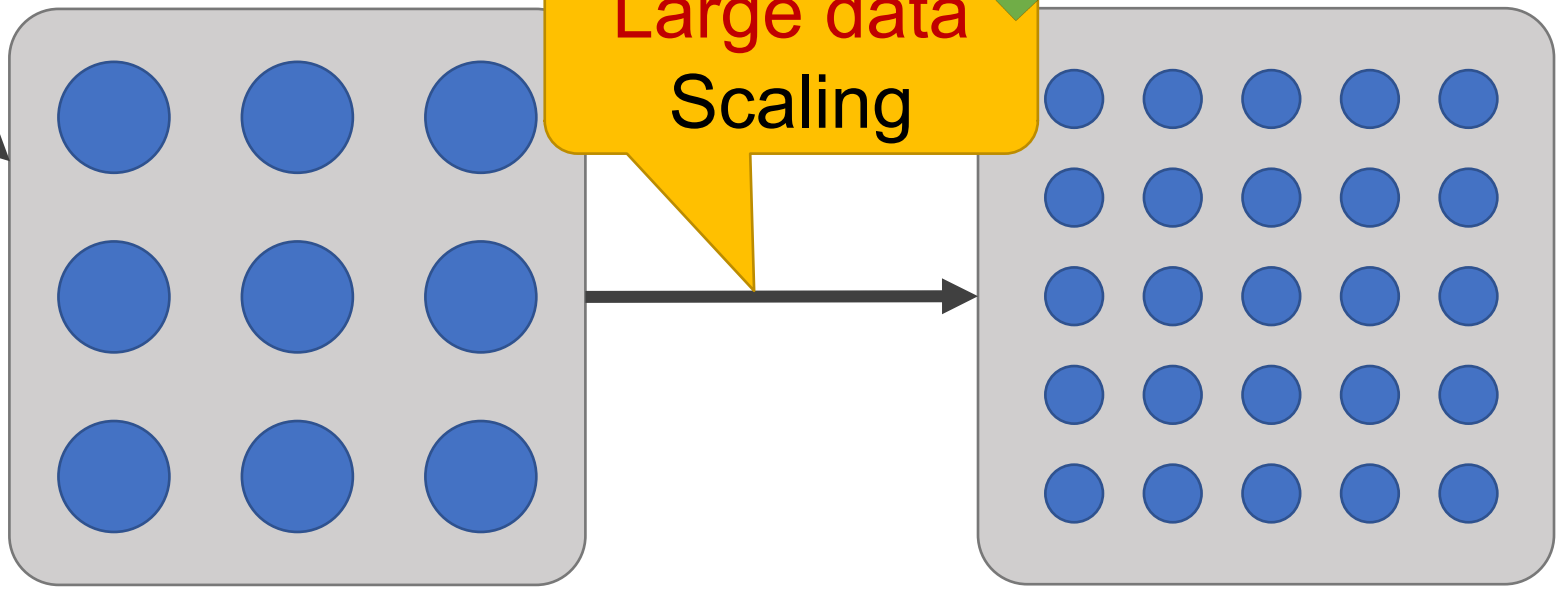

Scaling Degrades Reliability

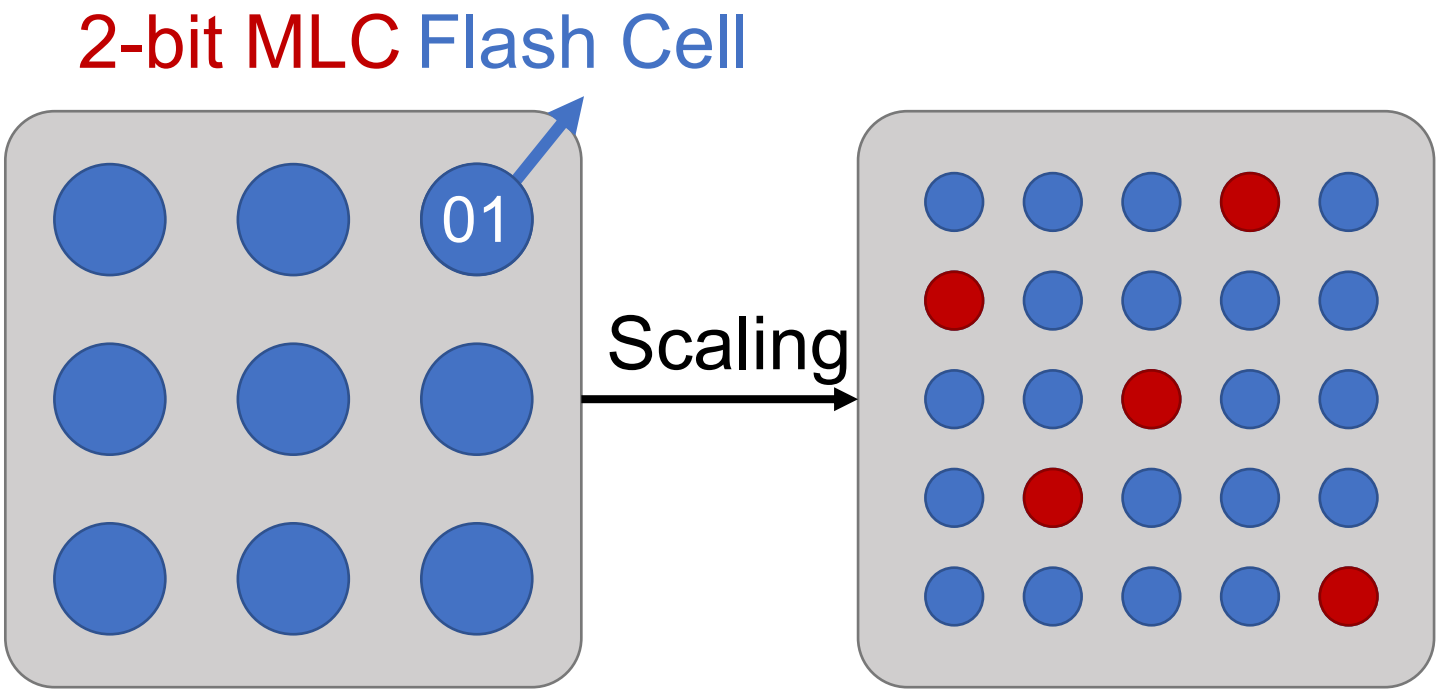

Scaling:
Smaller cell size
Smaller distance b/w cells

→ Bit flips or Raw bit errors

Degraded Flash Reliability

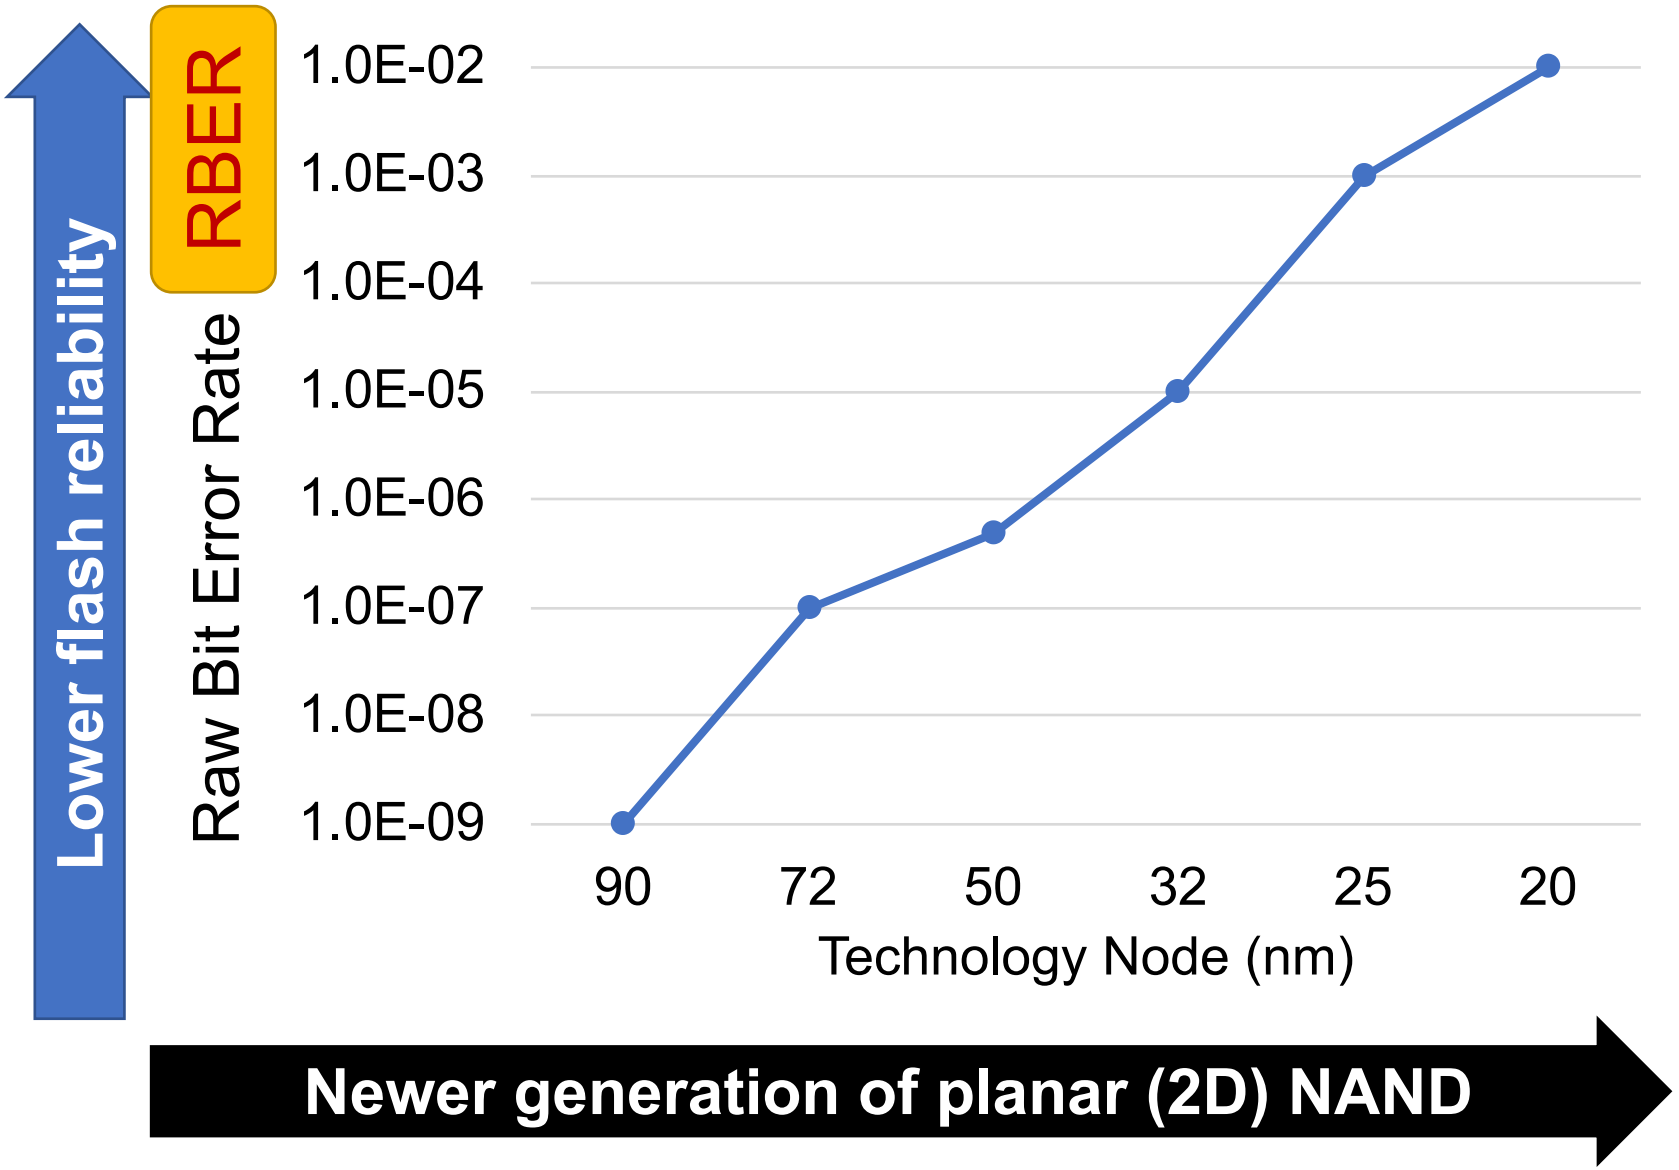

Source: Hagop Nazarian and Sylvain Dubois. The drive for SSDs: Whats holding back NAND flash? https://www.edn.com/Home/PrintView?contentItemId=4424905

Problem:
The Cost of Flash Reliability

Error Correction Code (ECC)

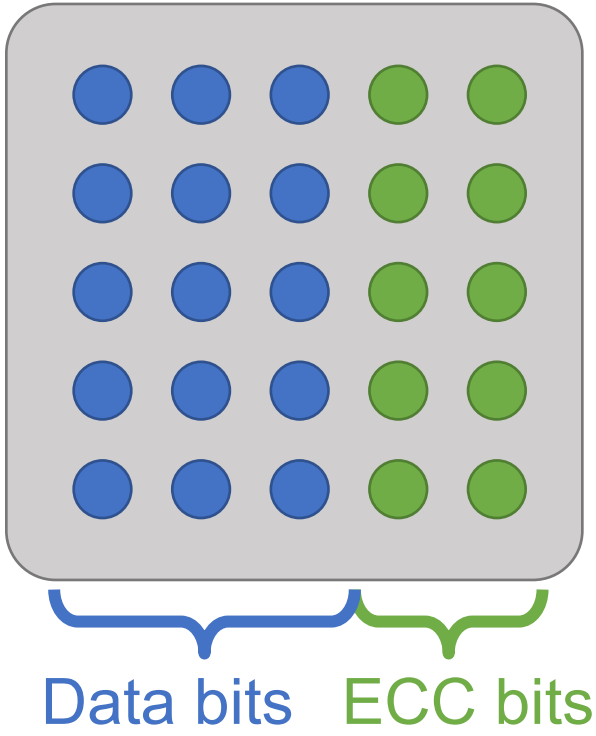

More ECC bits are required to correct more raw bit errors

Increased Cost to Improve Flash Reliability

High ECC cost, BUT NOT enough!

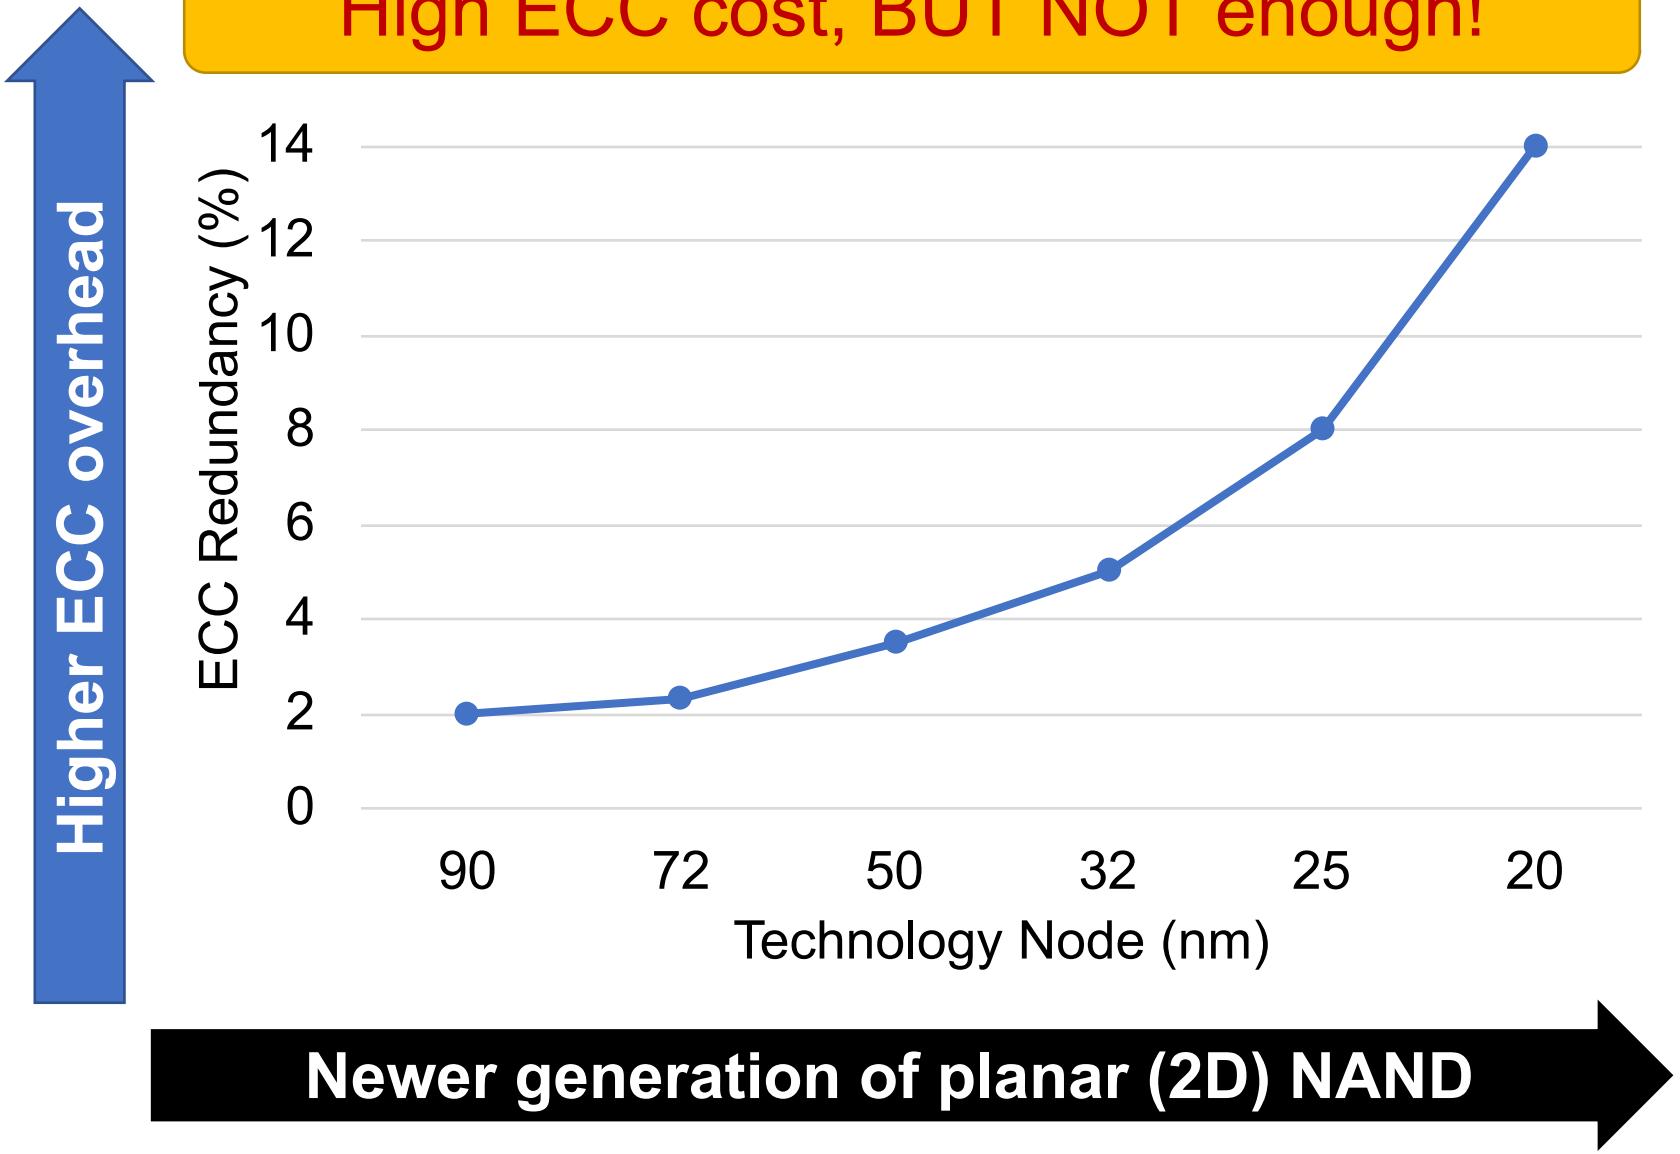

Source: Hagop Nazarian and Sylvain Dubois. The drive for SSDs: Whats holding back NAND flash? https://www.edn.com/Home/PrintView?contentItemId=4424905

P/E Cycle Lifetime

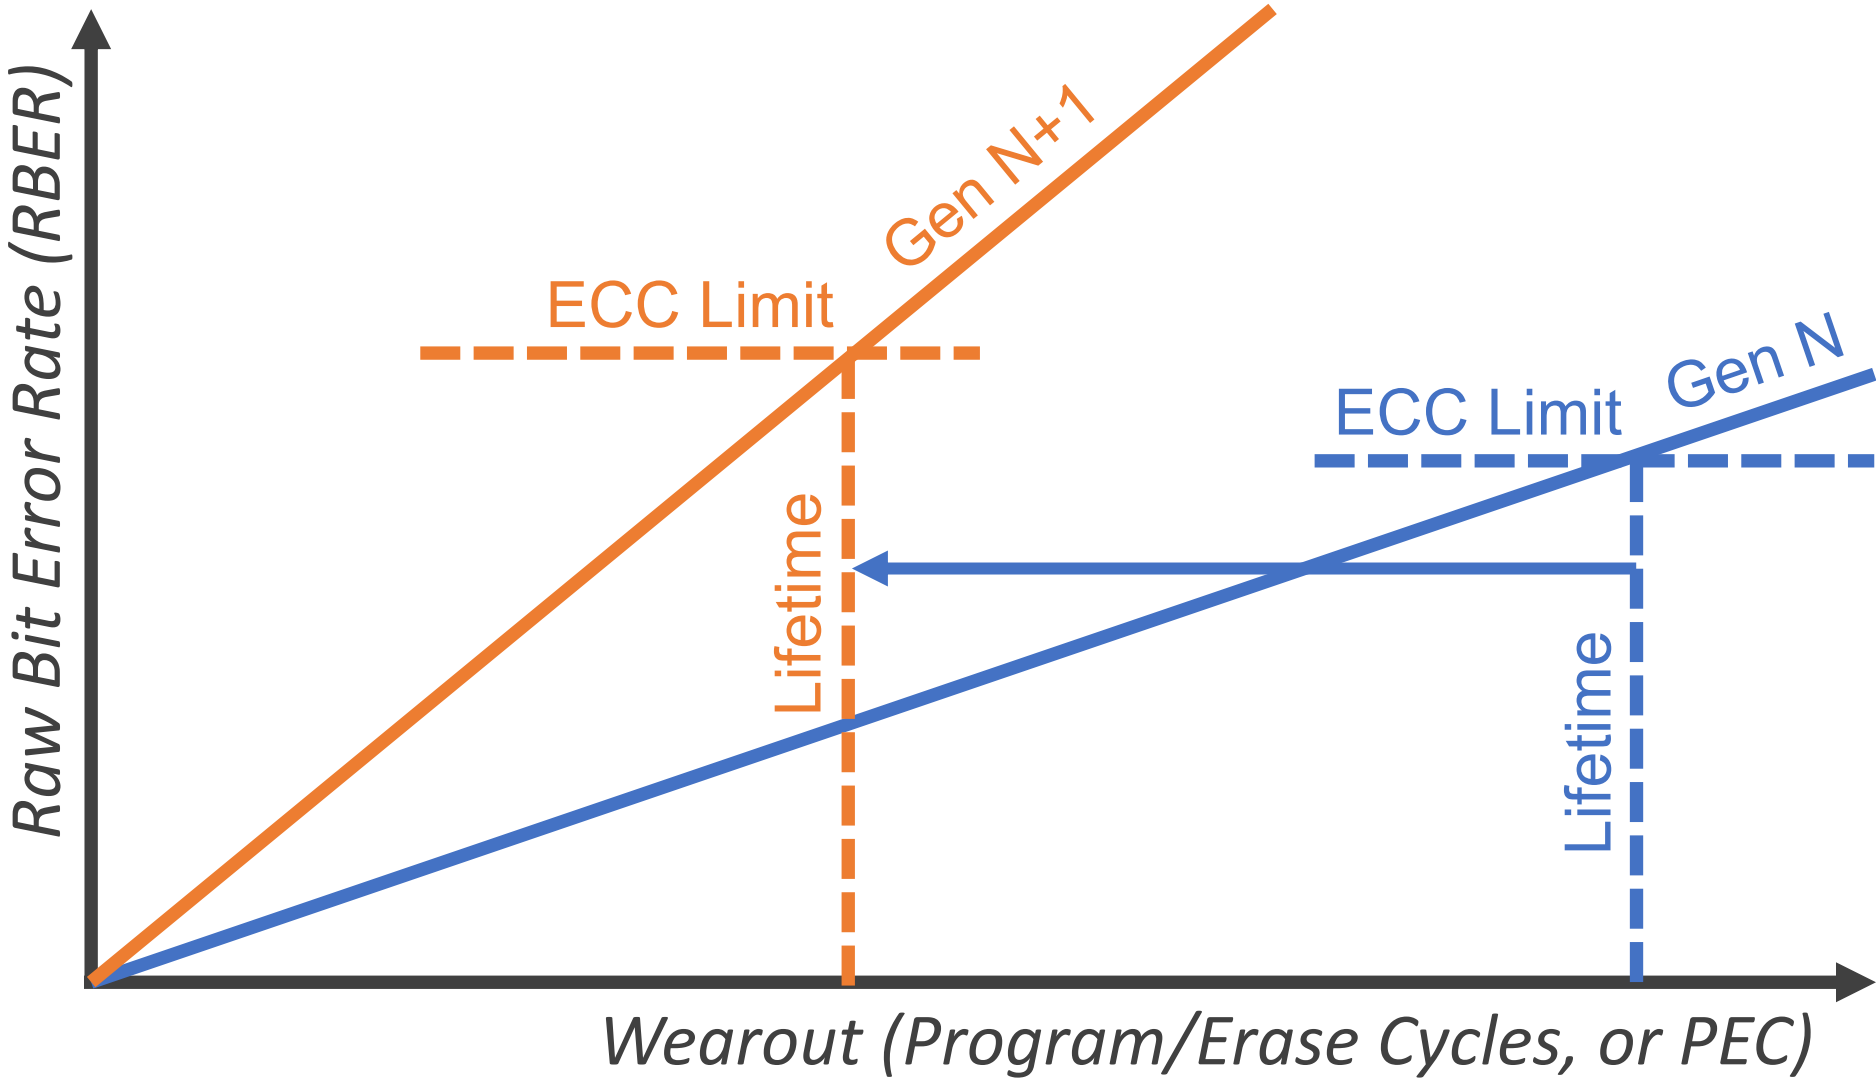

Degrading P/E Cycle Lifetime

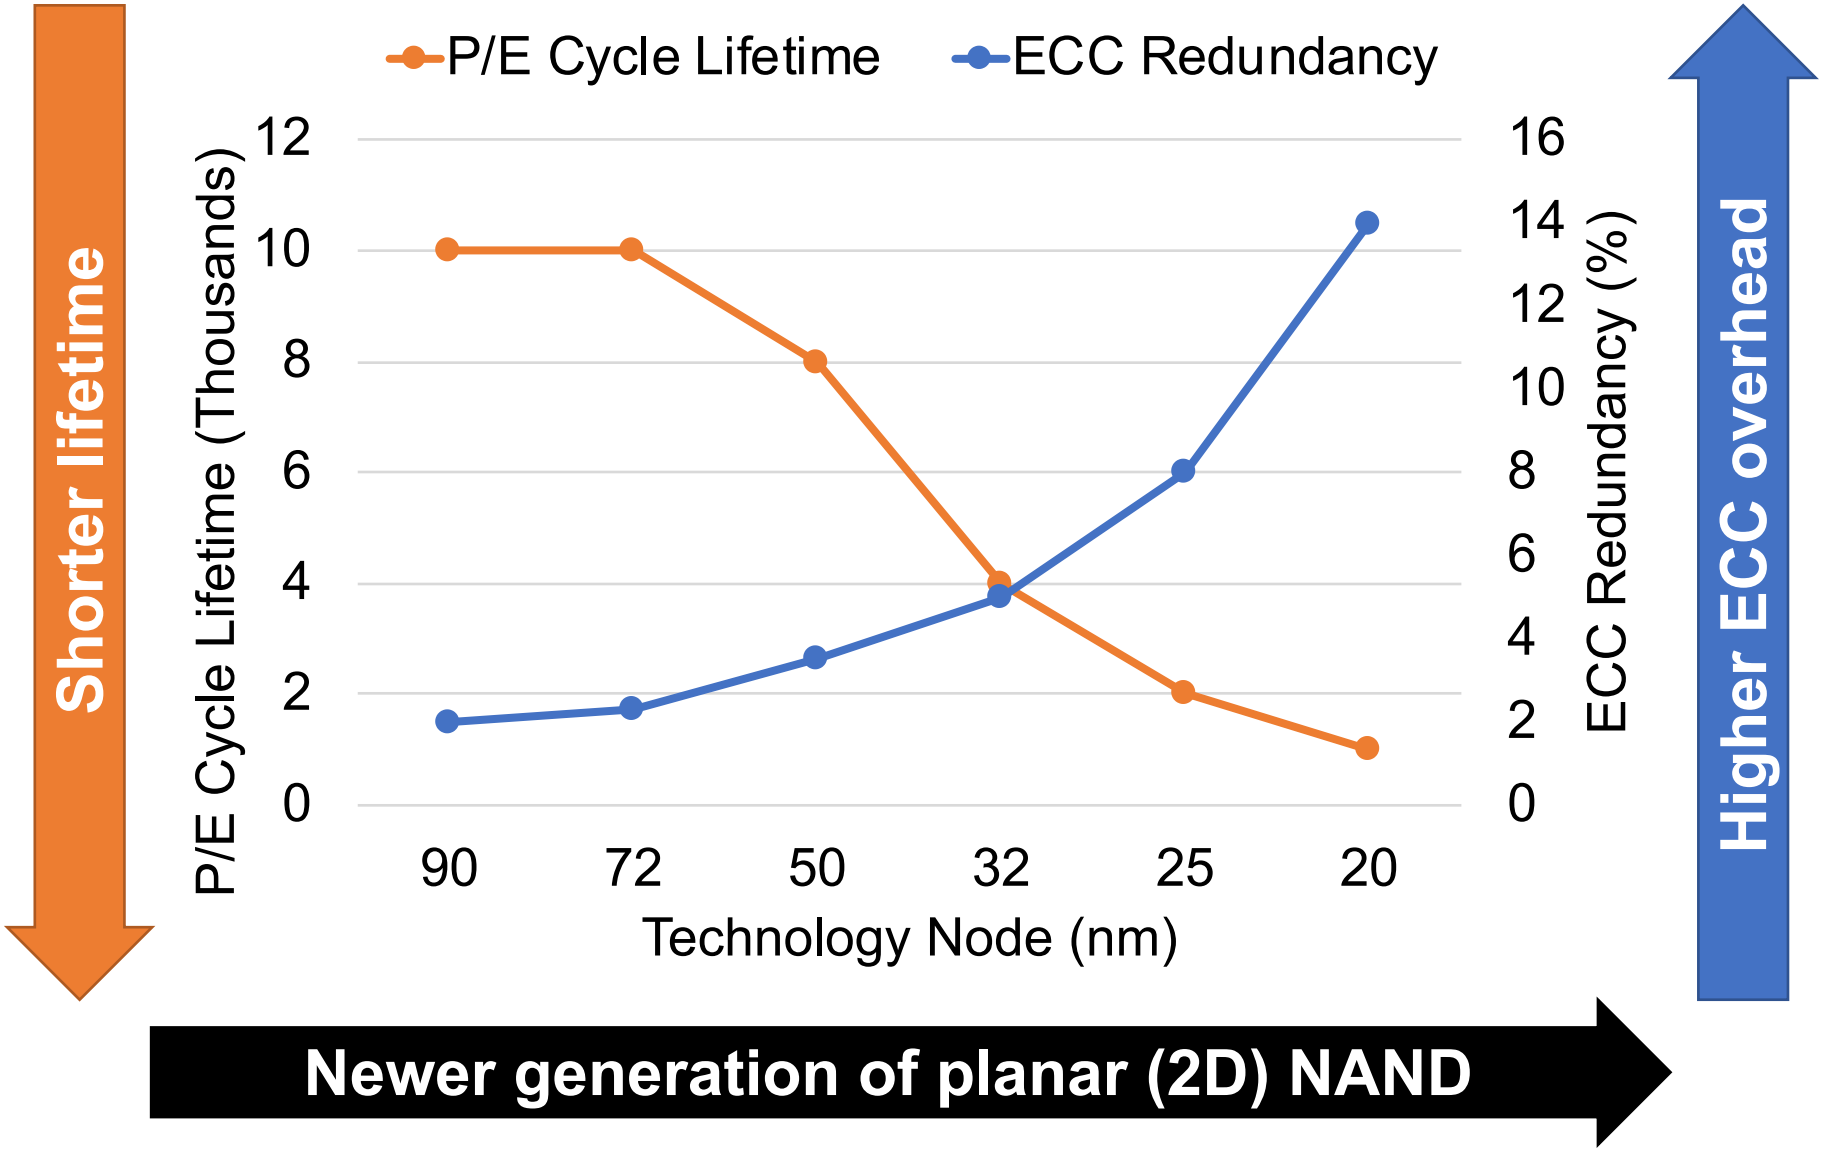

Goal:
Improve Flash Reliability at A Low Cost

Opportunities to Improve Flash Reliability

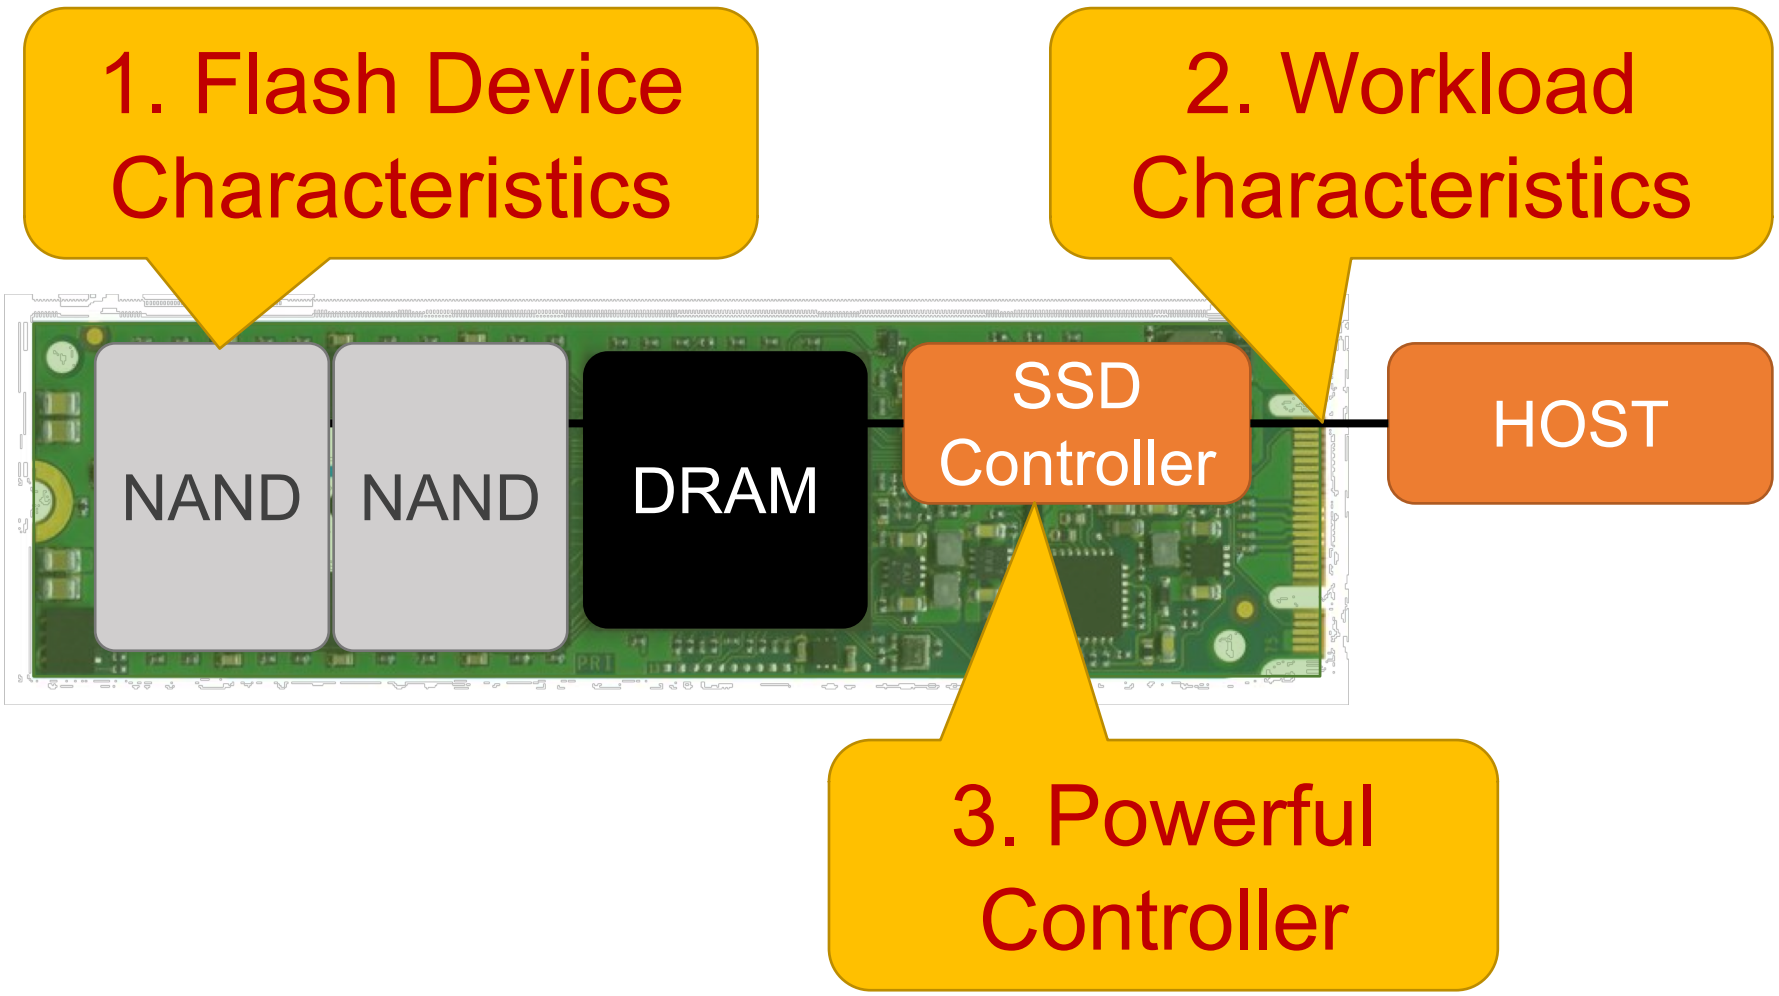

Thesis Statement

- NAND flash memory reliability can be improved
  - at low cost and with low performance overhead
- by deploying various architectural techniques that are aware of
  - higher-level application behavior and
  - underlying flash device characteristics

Contributions

Improve NAND flash memory reliability at low cost, using

1. Access pattern awareness
   - WARM [MSST'15]

2. Flash error awareness
   - Online Flash Channel Modeling [JSAC'16]

3. 3D NAND error and variation awareness
   - Understanding 3D NAND Errors, LI-RAID [under submission]

4. Self-recovery and temperature awareness
   - HeatWatch [HPCA'18]

Contributions

Improve NAND flash memory reliability at low cost, using

1. Access pattern awareness
   - WARM: Write-hotness Aware Retention Management [MSST'15]
     - *Retention: flash cell charge leakage over time*
     - *Write-hot data requires short retention time guarantee*

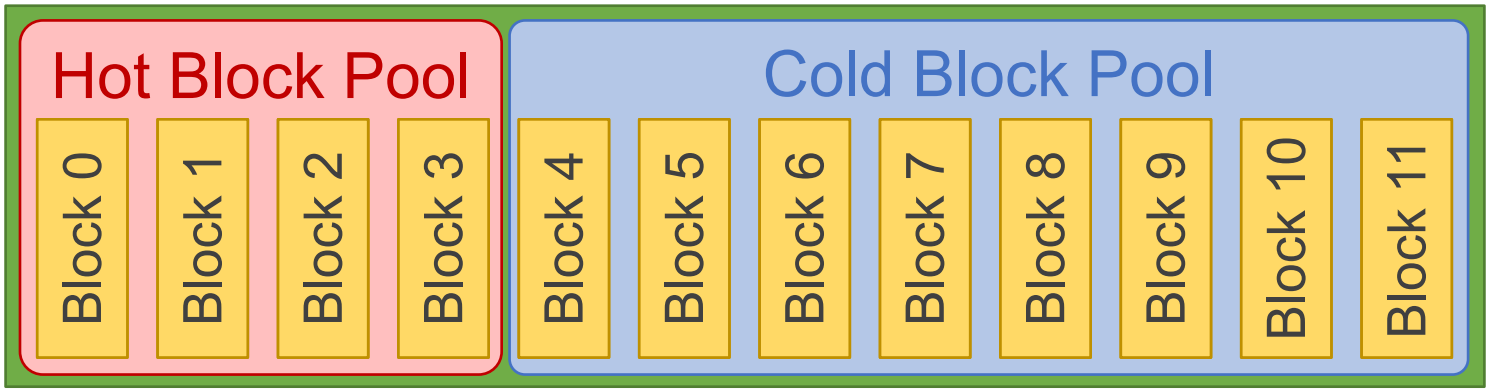

- *Improves flash lifetime by 12.9x*

Contributions

Improve NAND flash memory reliability at low cost, using

1. Access pattern awareness
   - WARM [MSST'15]

2. Flash error awareness
   - Online Flash Channel Modeling [JSAC 2016]
     - *Existing models designed for offline analysis*
     - *Accurate and easy-to-compute model*
       - Static threshold voltage distribution
       - Dynamically adjust to wearout
     - *Multiple applications*
       - Improves flash lifetime by up to 69.9%

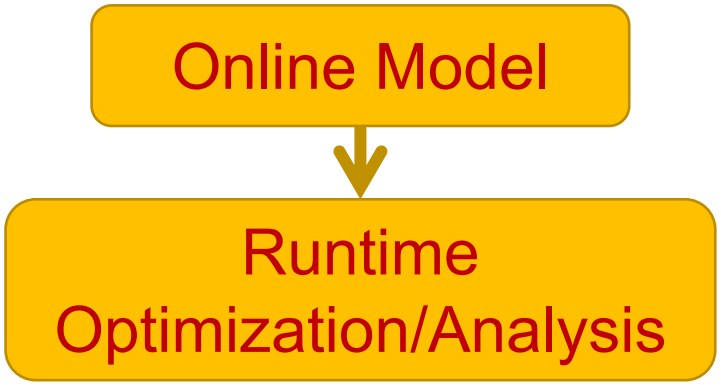

Flash Error Related Works

Planar (2D) NAND Errors

Data Retention	MSST'15, HPCA'15, ICCD'12
P/E Cycling	JSAC'16, GLOBECOM'14
Read Disturb	DSN'15, GLSVLSI'14, APSys'13
Two-Step Programming	HPCA'17, GLOBECOM'14
Program Interference	SIGMETRICS'14, ICCD'13

3D NAND

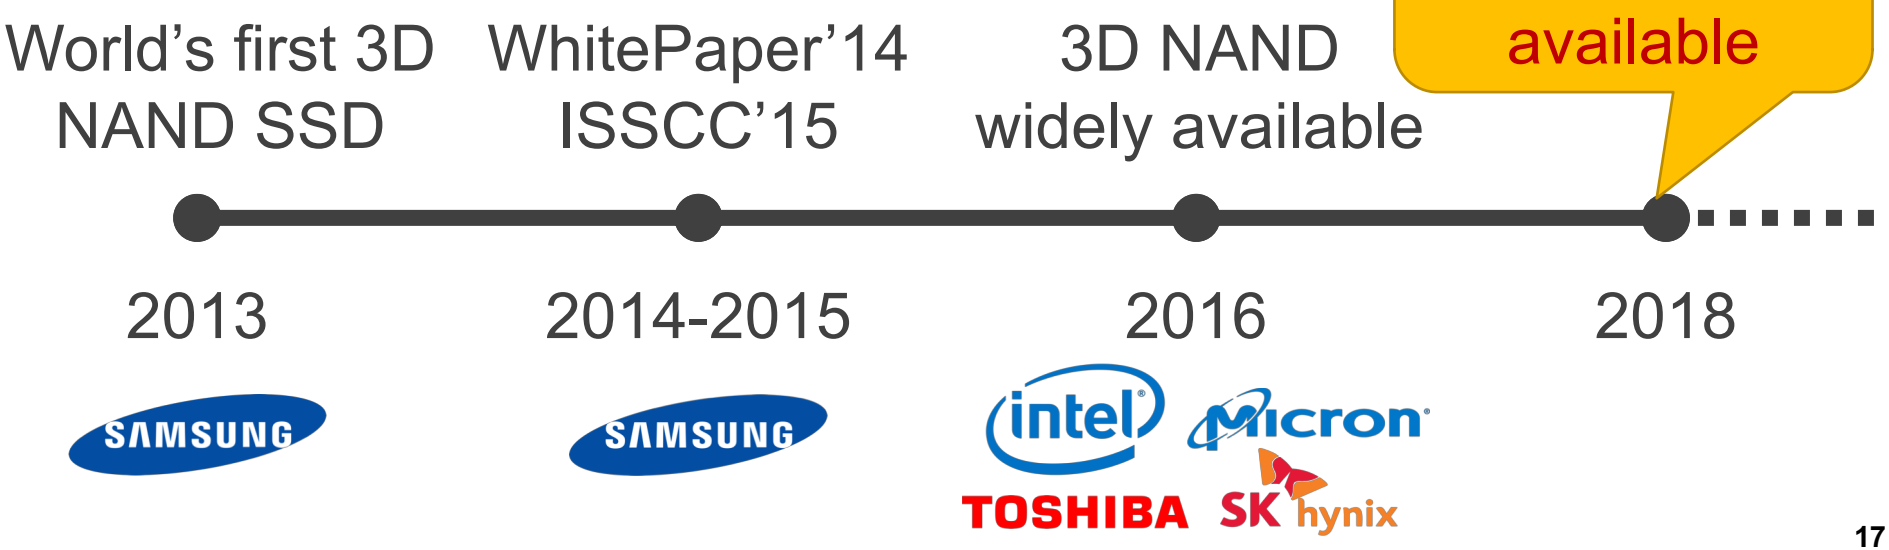

Contributions

Improve NAND flash memory reliability at low cost, using

1. Access pattern awareness
   - WARM [MSST'15]

2. Flash error awareness
   - Online Flash Channel Modeling [JSAC 2016]

3. 3D NAND error and variation awareness
   - Understanding 3D NAND Errors, LI-RAID [under submission]

4. Self-recovery and temperature awareness
   - HeatWatch [HPCA 2018]

Focus of this talk

1. Flash Device Characteristics

Understanding 3D NAND Errors: Through Characterization

Characterization Methodology

- Real flash chips
  - 3D NAND: 30-39 layer MLC 3D NAND flash chips
  - 2D NAND: 15-19 nm MLC NAND flash chips
- Using a modified firmware version in the SSD controller
  - Control the read reference voltage of the flash chip
  - Bypass ECC to get raw NAND data (with raw bit errors)
- Using a heat chamber to control SSD temperature

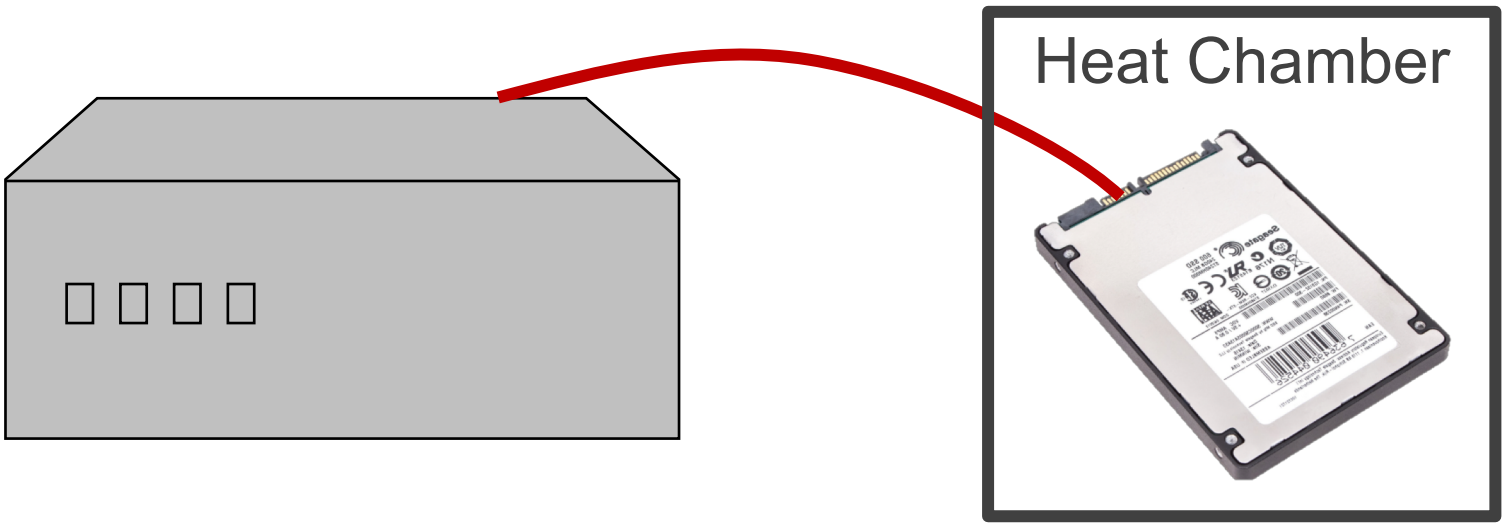

Characterization Methodology Cont'd

- 5 months to collect the data, even more for analysis
- Collected >180GB compressed data
- Characterize threshold voltage rather than raw bit error rate
  - Cannot be done without our methodology
  - Enables deeper understanding and new techniques
- Rigorous experiments to study 7 types of errors
  - P/E cycling, program interference, read disturb, read variation, retention, retention interference, process variation
- Develop insights into data through statistical modeling and analysis using python scripts

3D NAND Error Characteristics

Attribute	Observation in 3D NAND	Cause of Difference	Future Trend
Retention	HeatWatch dominate all errors		
Process Variation	New layer-to-layer	LI-RAID	
P/E Cycling	Other errors become less significant because of larger process technology		
Programming			
Program interference			
Read disturb			

2. Workload Characteristics

HeatWatch:

Mitigate 3D NAND Retention Using Self-Recovery and Temperature Awareness

3. Powerful Controller

Retention Errors

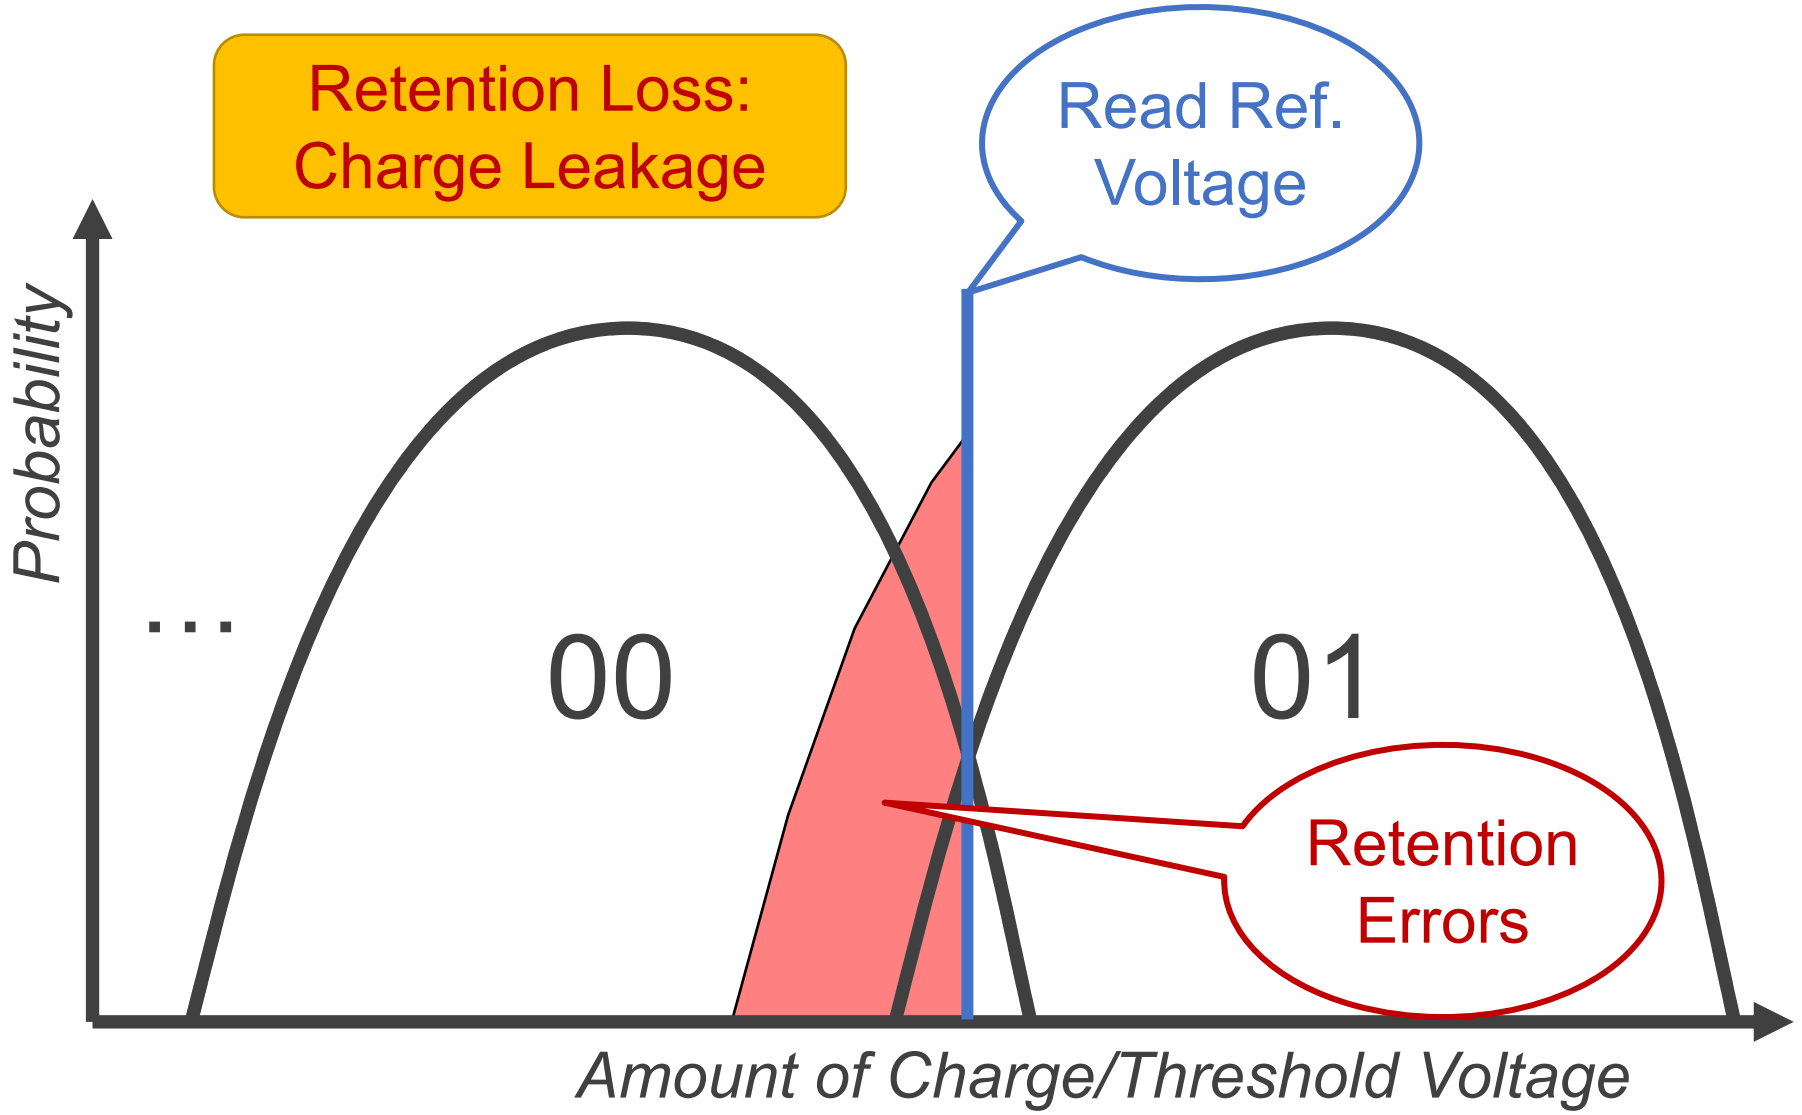

Charge → Voltage → Bit Values

Retention Errors Dominates

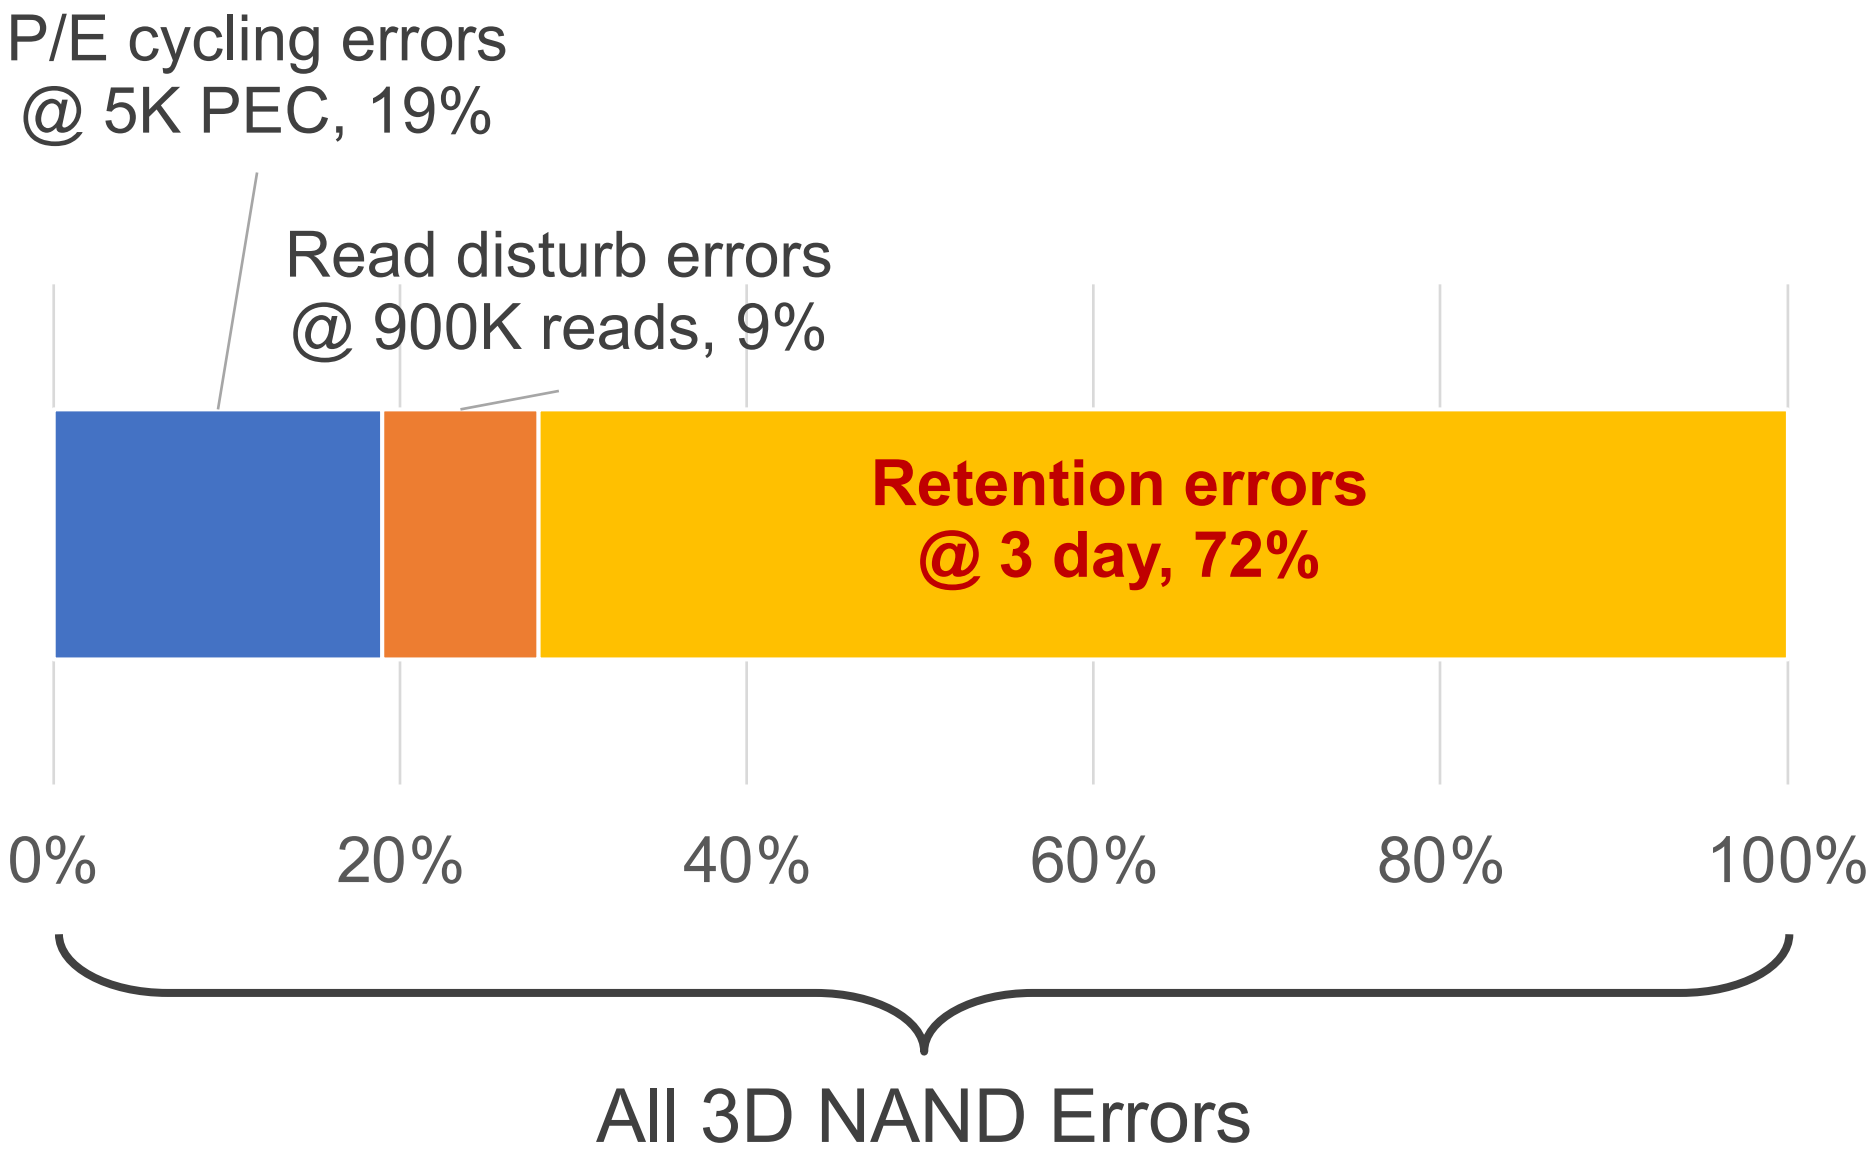

Mitigating Retention Errors

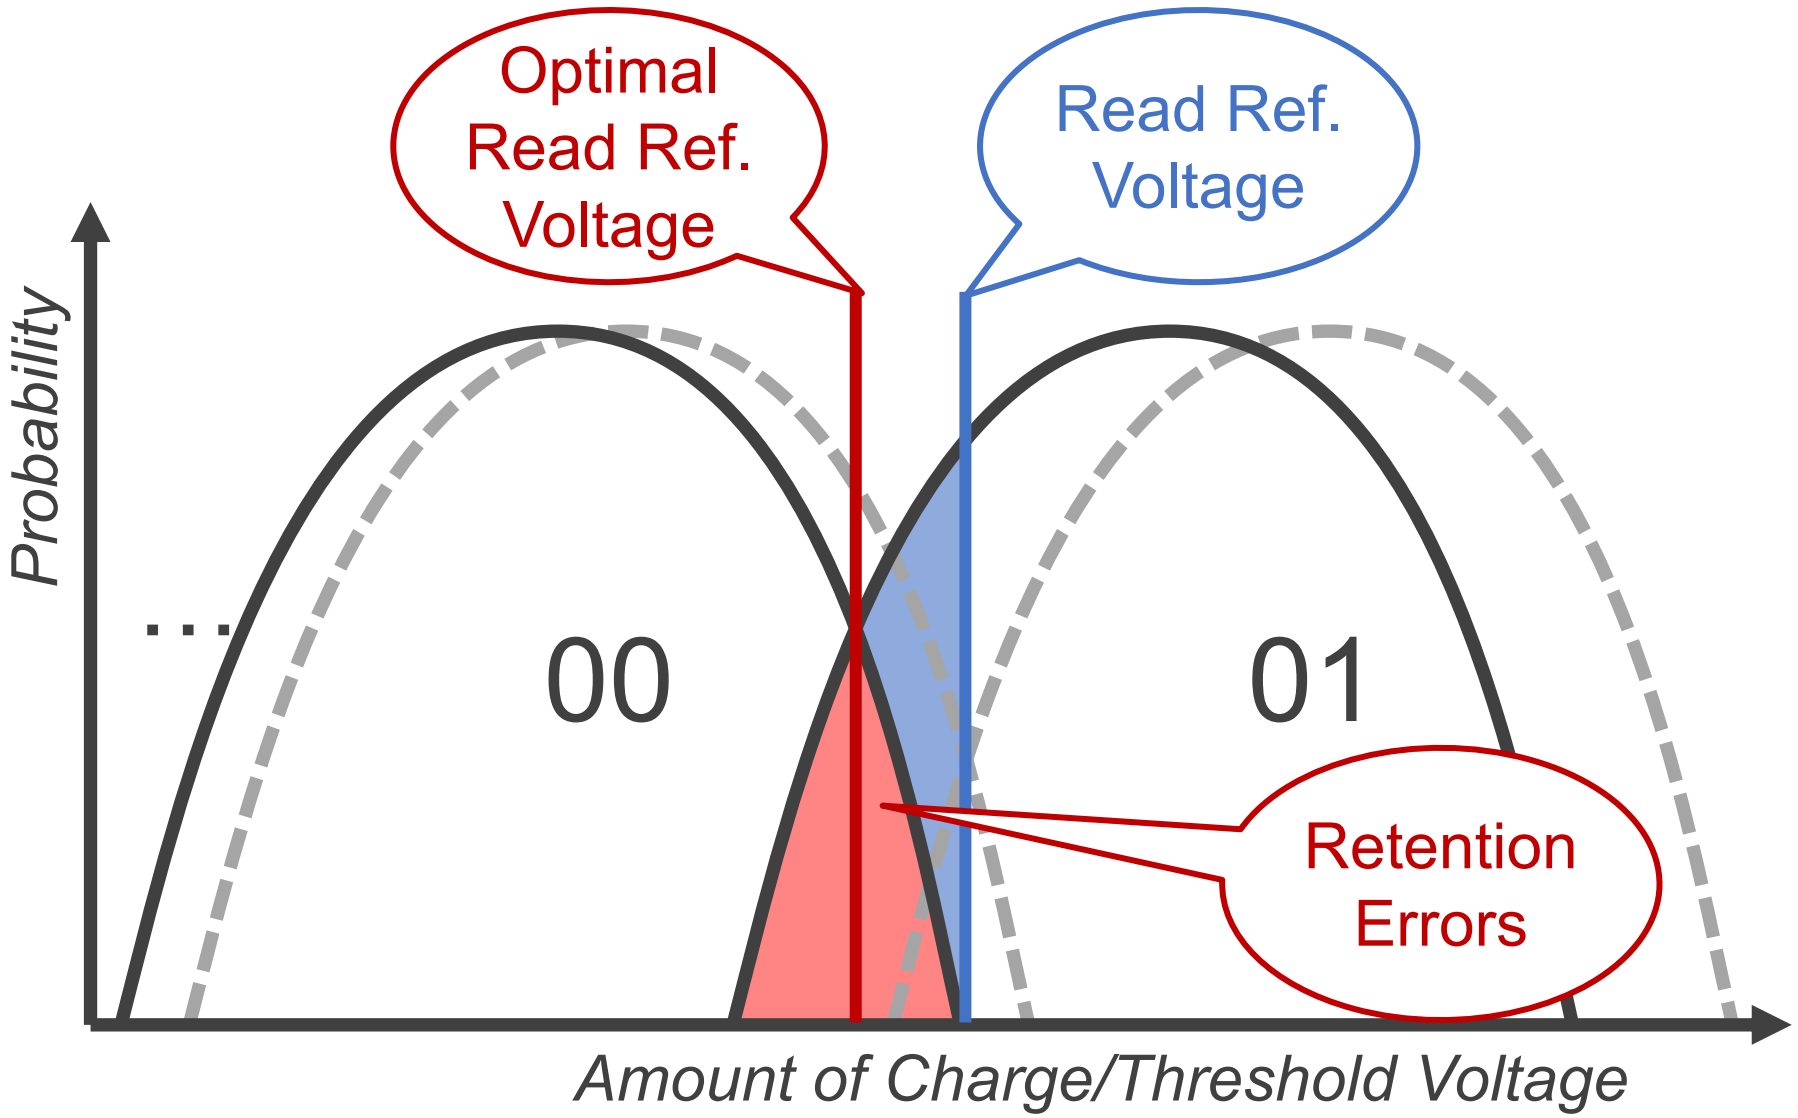

Predicting The Optimal Read Ref. Voltage

$$V_{opt} = V_0 + \Delta V$$

1. Initial Voltage Before Retention

2. Voltage Shift due to Retention Loss

1. Predicting V_0

Conventional Model

- Wearout (PEC)
  - Power-law model [JSAC'16]

HeatWatch Model

- 3D NAND Wearout (PEC)
  - Linear model

3D NAND Wearout Effect

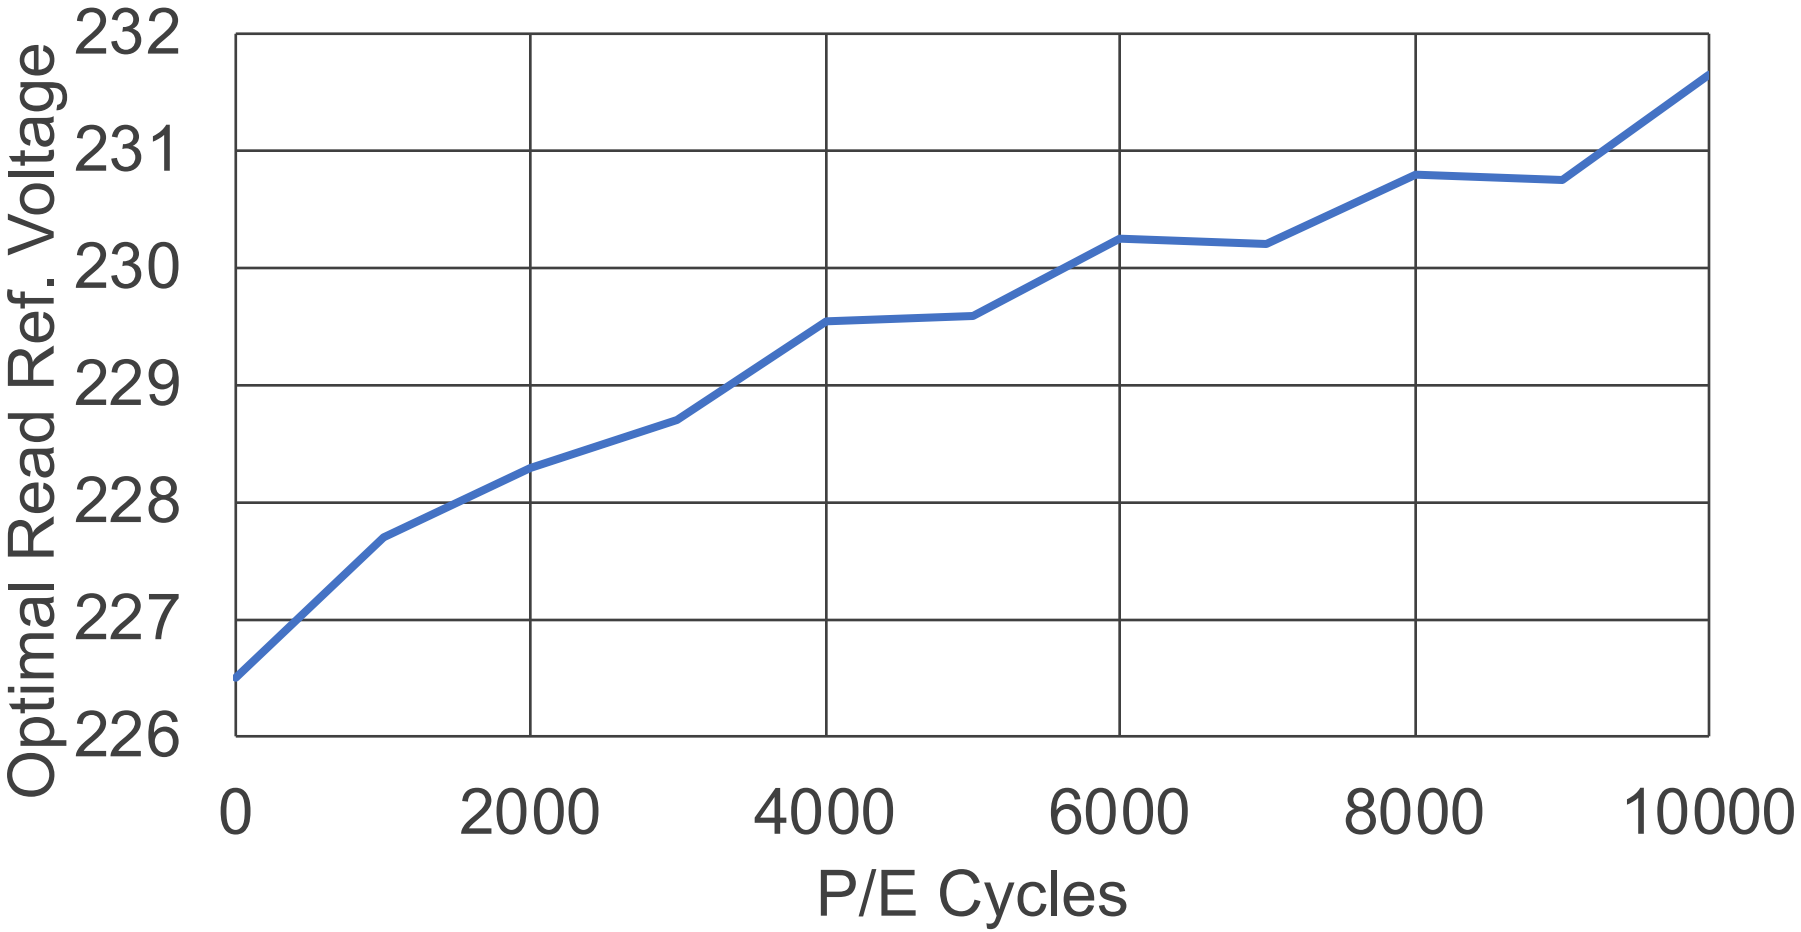

3D NAND wearout follows a linear trend

Predicting V_0

Conventional Model

- Wearout (PEC)
  - Power-law model [JSAC'16]

HeatWatch Model

- 3D NAND Wearout (PEC)
  - Linear model

- Prog. Temperature (T_p)

Programming Temperature Effect

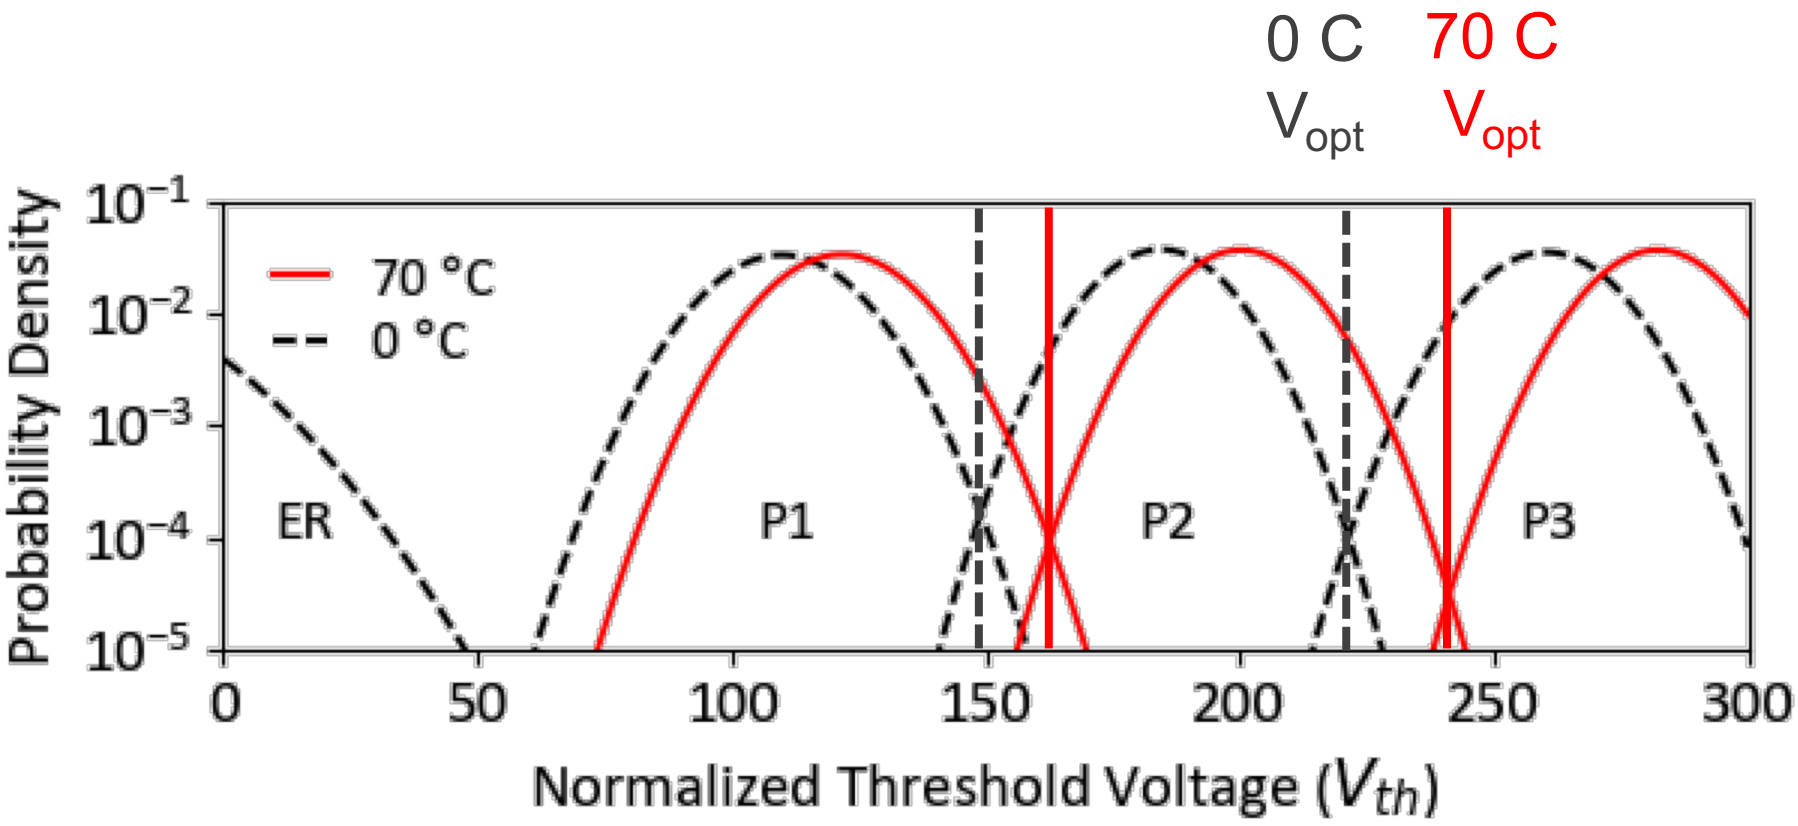

A higher temperature increases the optimal read reference voltage

Predicting The Optimal Read Ref. Voltage

Program Variation Component

$$Y_0 = A \cdot T_p \cdot PEC + B \cdot T_p + C \cdot PEC + D$$

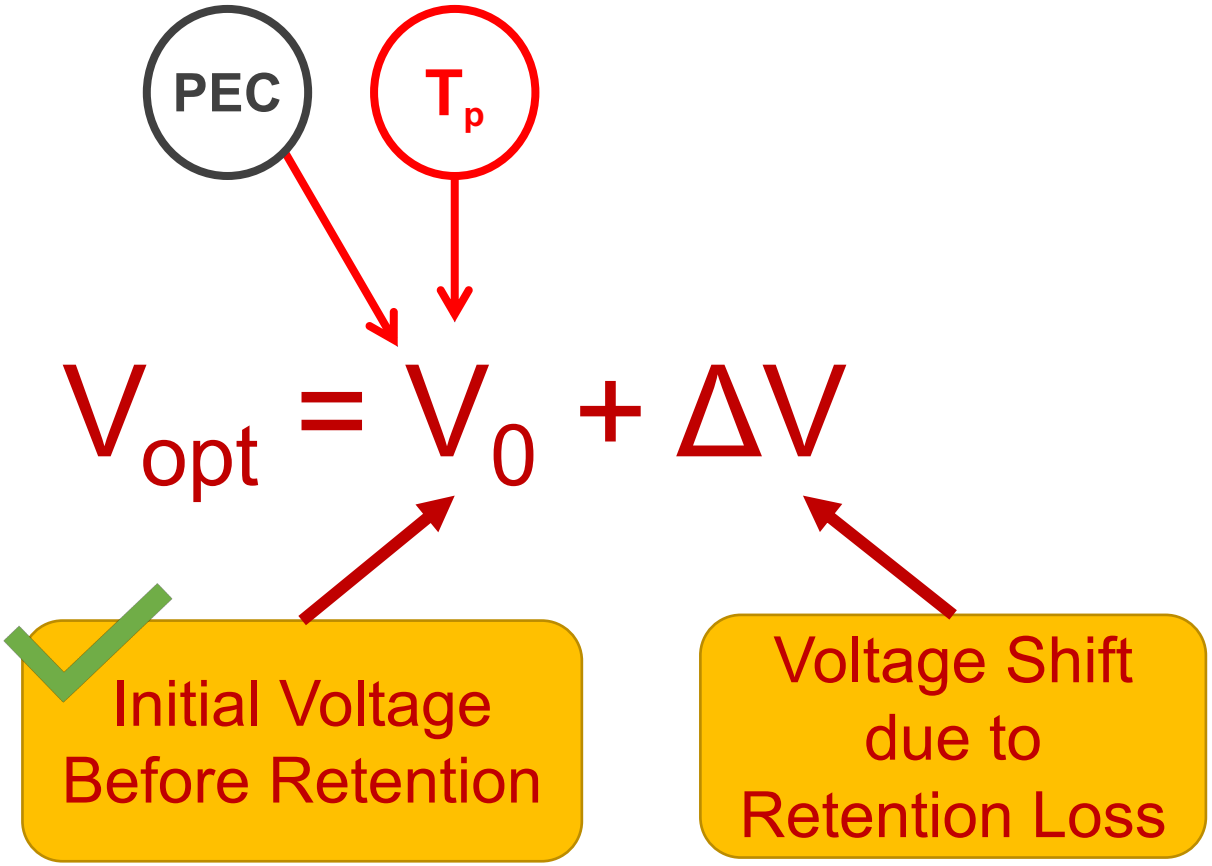

Predicting ΔV

Conventional Model

- Wearout (PEC)
- Retention Time (t_r)

HeatWatch Model

- 3D NAND Wearout (PEC)
- Retention Time (t_r)

- Dwell Time (t_d)
  - Idle time between program cycles

Self-Recovery Effect

Long dwell time slows down retention

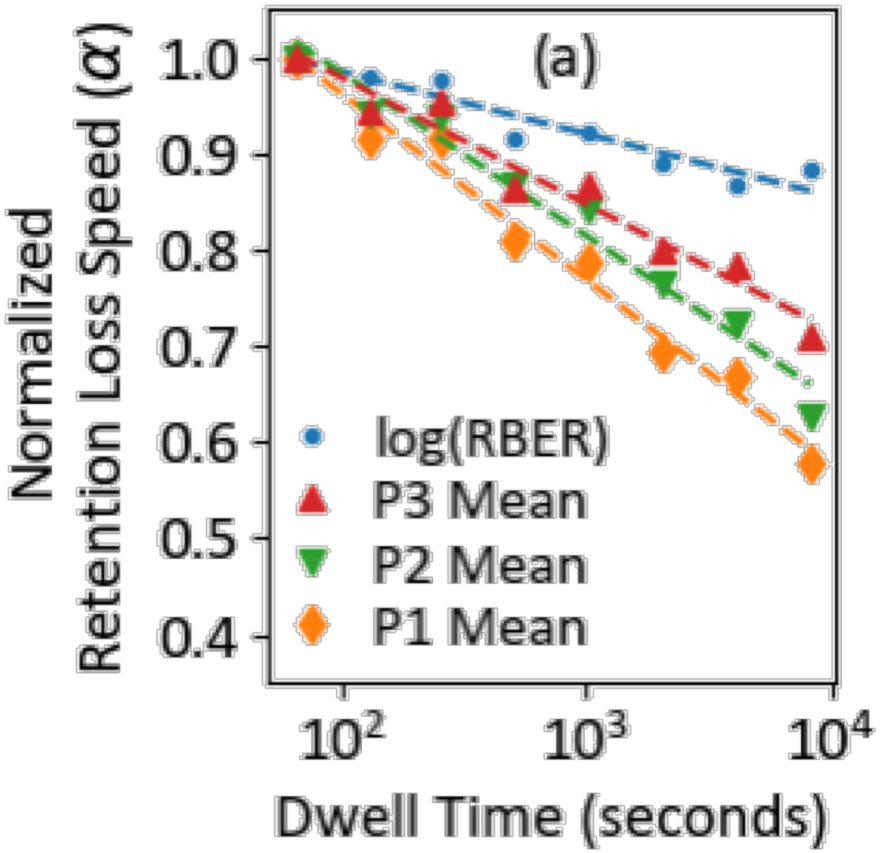

Self-Recovery Component

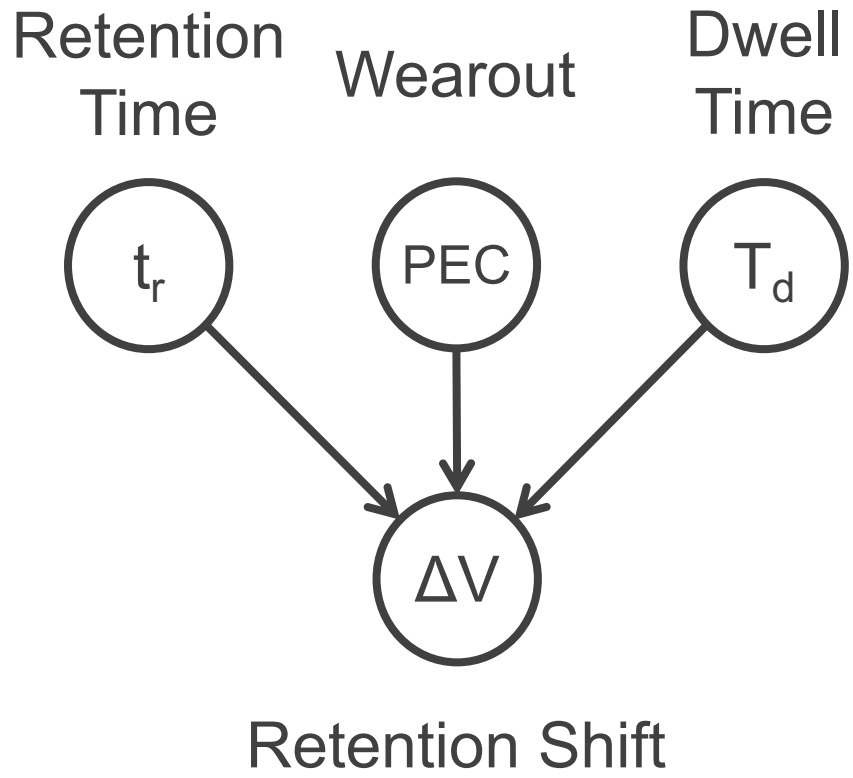

$$\Delta Y(t_{er}, t_{ed}, PEC) = b \cdot (PEC + c) \cdot \ln\left(1 + \frac{t_{er}}{t_0 + a \cdot t_{ed}}\right)$$

Predicting ΔV

Conventional Model

- Wearout (PEC)
- Retention Time (t_r)

HeatWatch Model

- 3D NAND Wearout (PEC)
- Retention Time (t_r)

- Dwell Time (t_d)
  - Idle time between program cycles

- Retention & Dwell Temperature (T_r & T_d)

Retention Temperature Effect

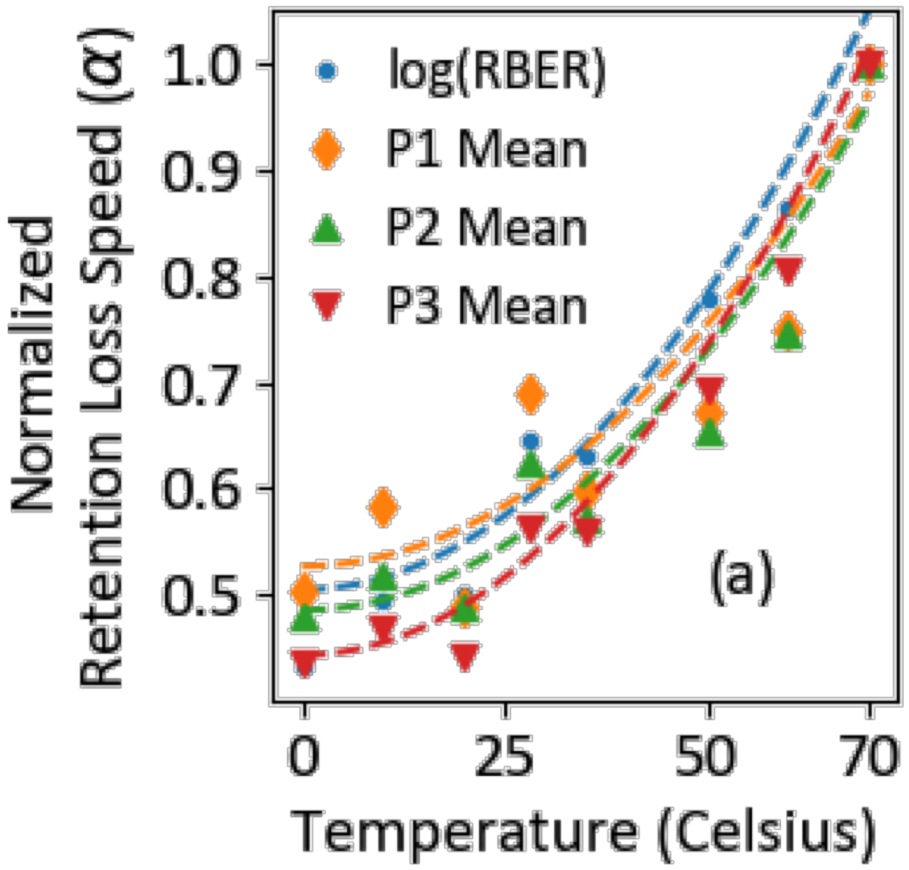

High temperature accelerates retention

Predicting ΔV

Conventional Model

- Wearout (PEC)
- Retention Time (t_r)

- Arrhenius Law with known activation energy (E_a) [JEDEC'10][ZPC1889]

HeatWatch Model

- 3D NAND Wearout (PEC)
- Retention Time (t_r)

- Dwell Time (t_d)
  - Idle time between program cycles

- Retention & Dwelling Temperature (T_r & T_d)
  - E_a for 3D NAND?

Effective Retention/Dwell Time Component

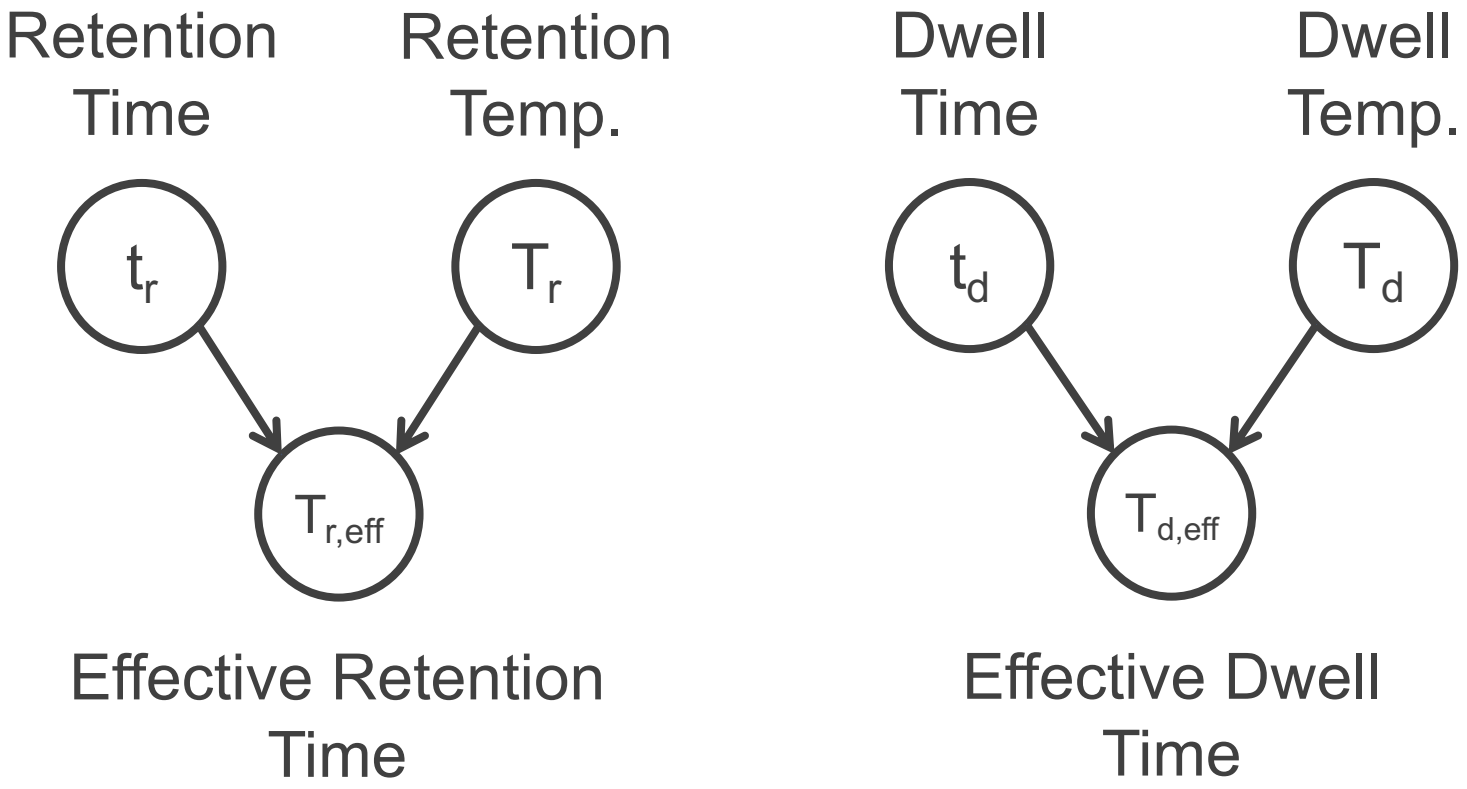

E_a = 1.04 eV
95% CI: 1.01 – 1.08 eV

Predicting The Optimal Read Ref. Voltage

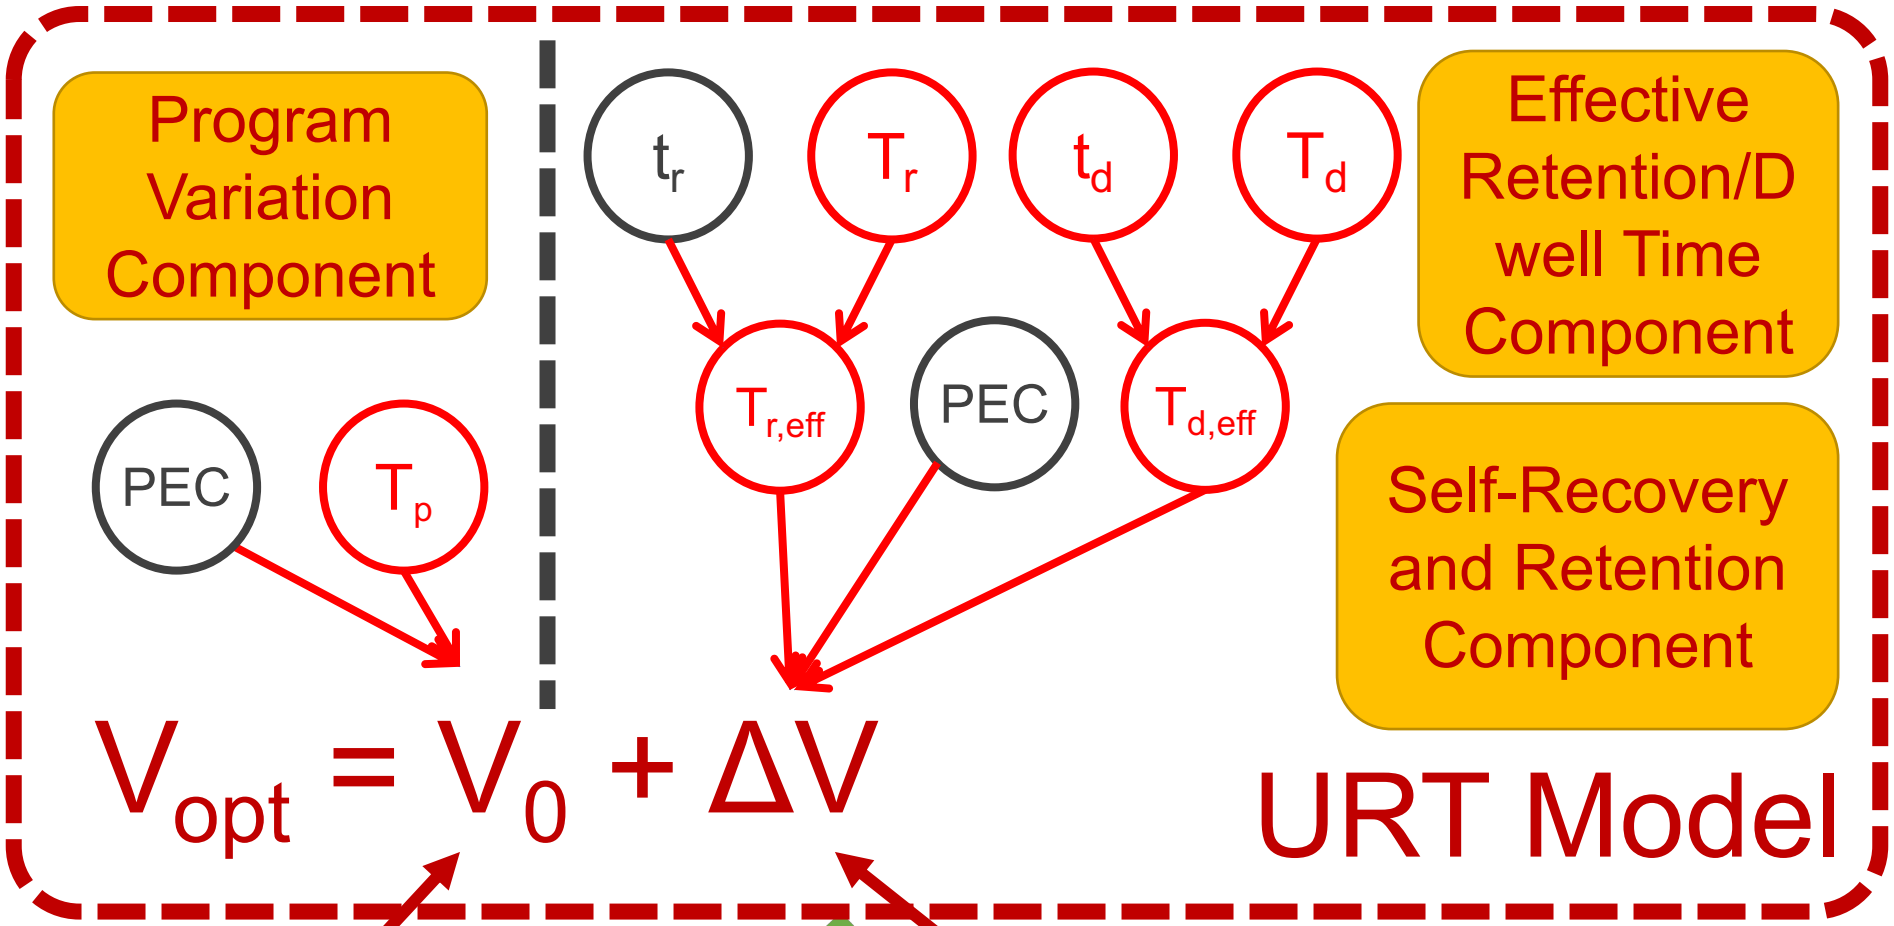

HeatWatch Mechanism

- Key Idea: Adapt to workload characteristics using URT model
- Tracking Components (Efficiently track URT parameters)
  - Tracking SSD temperature
  - Tracking dwell time
  - Tracking PEC and retention time

- Prediction Components (Accurately predict V_{opt} using URT)
  - Predicting the optimal read reference voltage
  - Fine-tuning URT model parameters online

Tracking SSD Temperature

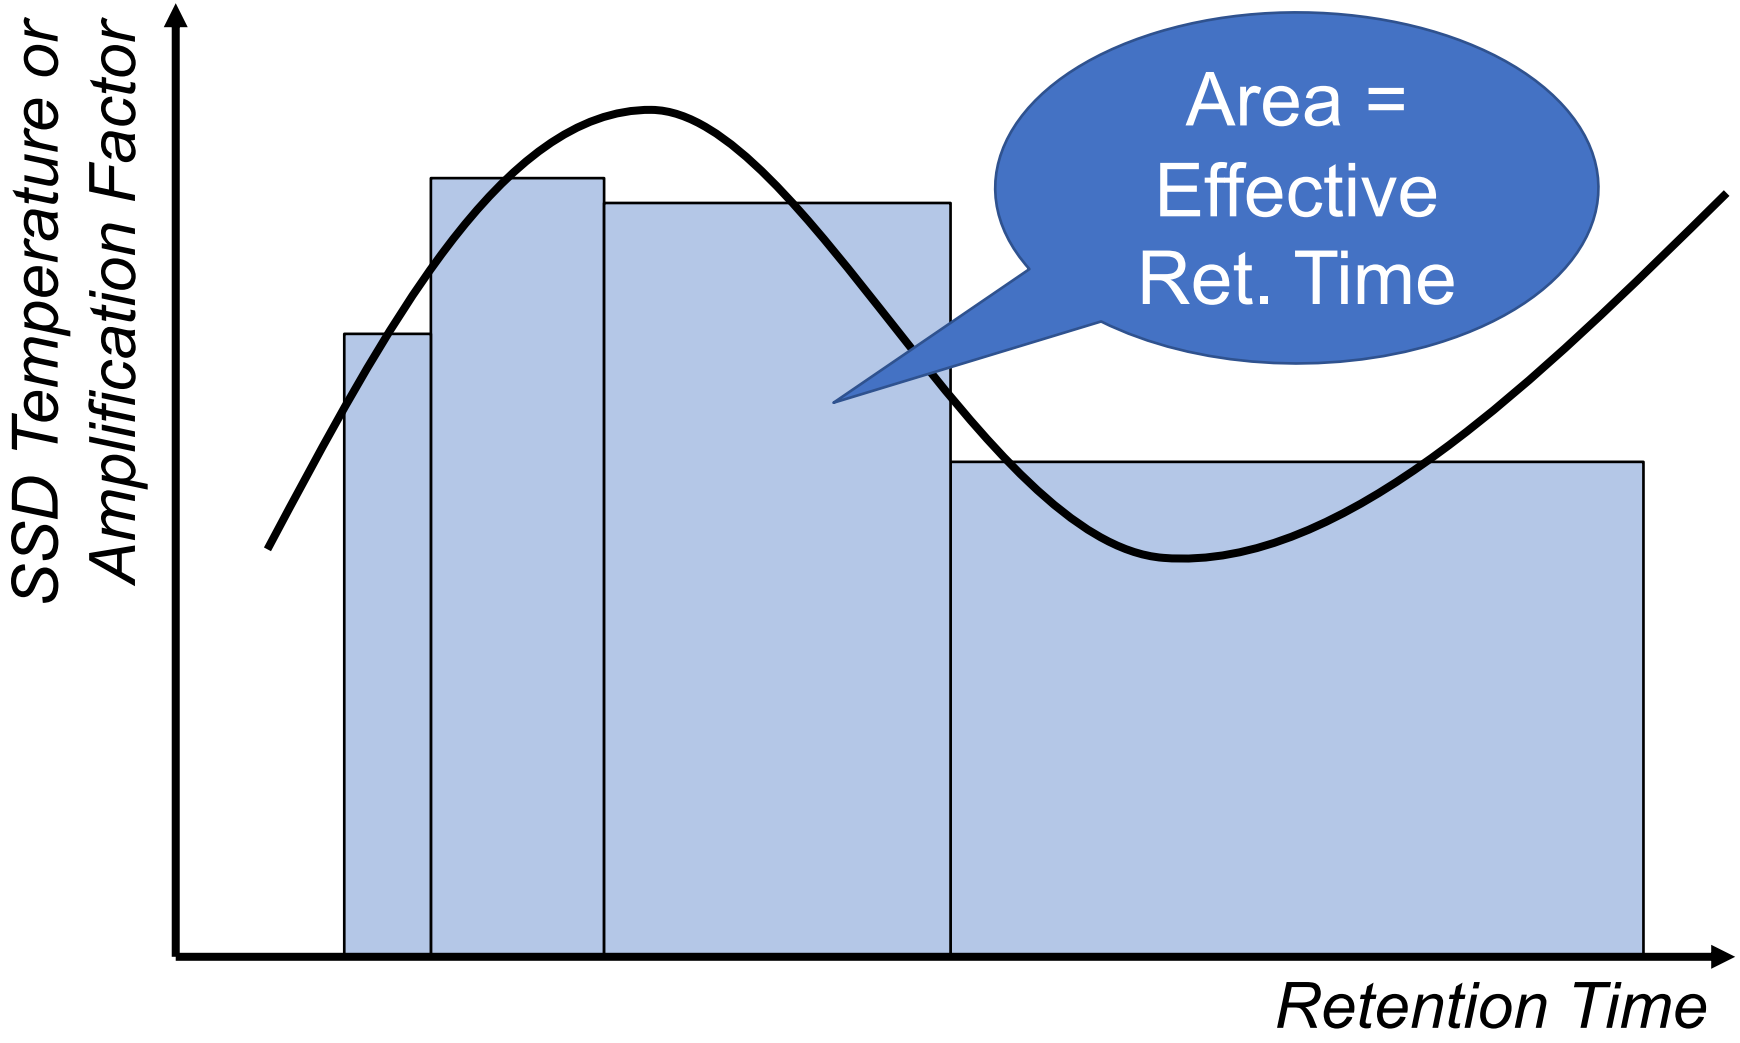

HeatWatch Mechanism Cont'd

- Key Idea: Adapt to workload characteristics using URT model
- Tracking Components (Efficiently track URT parameters)
  - Tracking SSD temperature
    - *Precompute and store, use existing sensors*
  - Tracking dwell time
    - *Only for the last 20 PEC*
  - Tracking PEC and retention time
    - *Log write timestamp per flash block*

- Prediction Components (Accurately predict V_{opt} using URT)
  - Predicting the optimal read reference voltage
  - Fine-tuning URT model parameters online

HeatWatch Mechanism Cont'd

- Key Idea: Adapt to workload characteristics using URT model
- Tracking Components (Efficiently track URT parameters)
  - Tracking SSD temperature
  - Tracking dwell time
  - Tracking PEC and retention time
  - Storage Overhead: <1.6MB for 1TB SSD

- Prediction Components (Accurately predict V_{opt} using URT)
  - Predicting the optimal read reference voltage
    - *Modeling error: 4.9%*
  - Fine-tuning URT model parameters online
    - *Use periodic sampling*
  - Latency Overhead: <1%

Evaluation Methodology

- 28 real-workload traces
  - Real dwell time, retention time
  - MSR-Cambridge
- Temperature Model:

  Trigonometric function + Gaussian noise
  - Periodic temperature variation within each day
  - Small transient temperature variation

Flash Lifetime Improvements

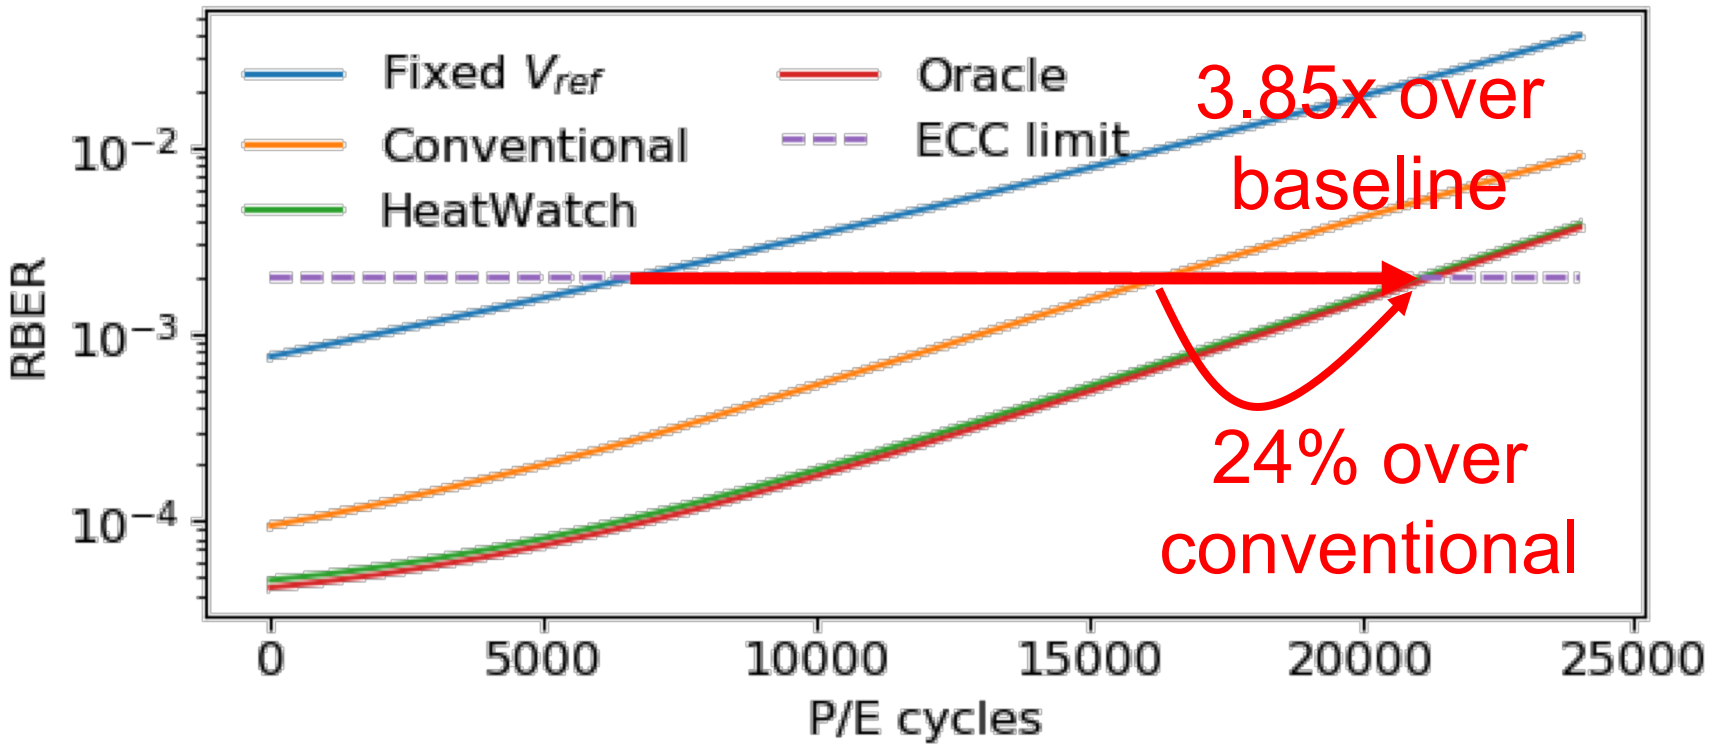

1. Flash Device Characteristics

LI-RAID:
Mitigate Process Variation

Layer-to-Layer Process Variation

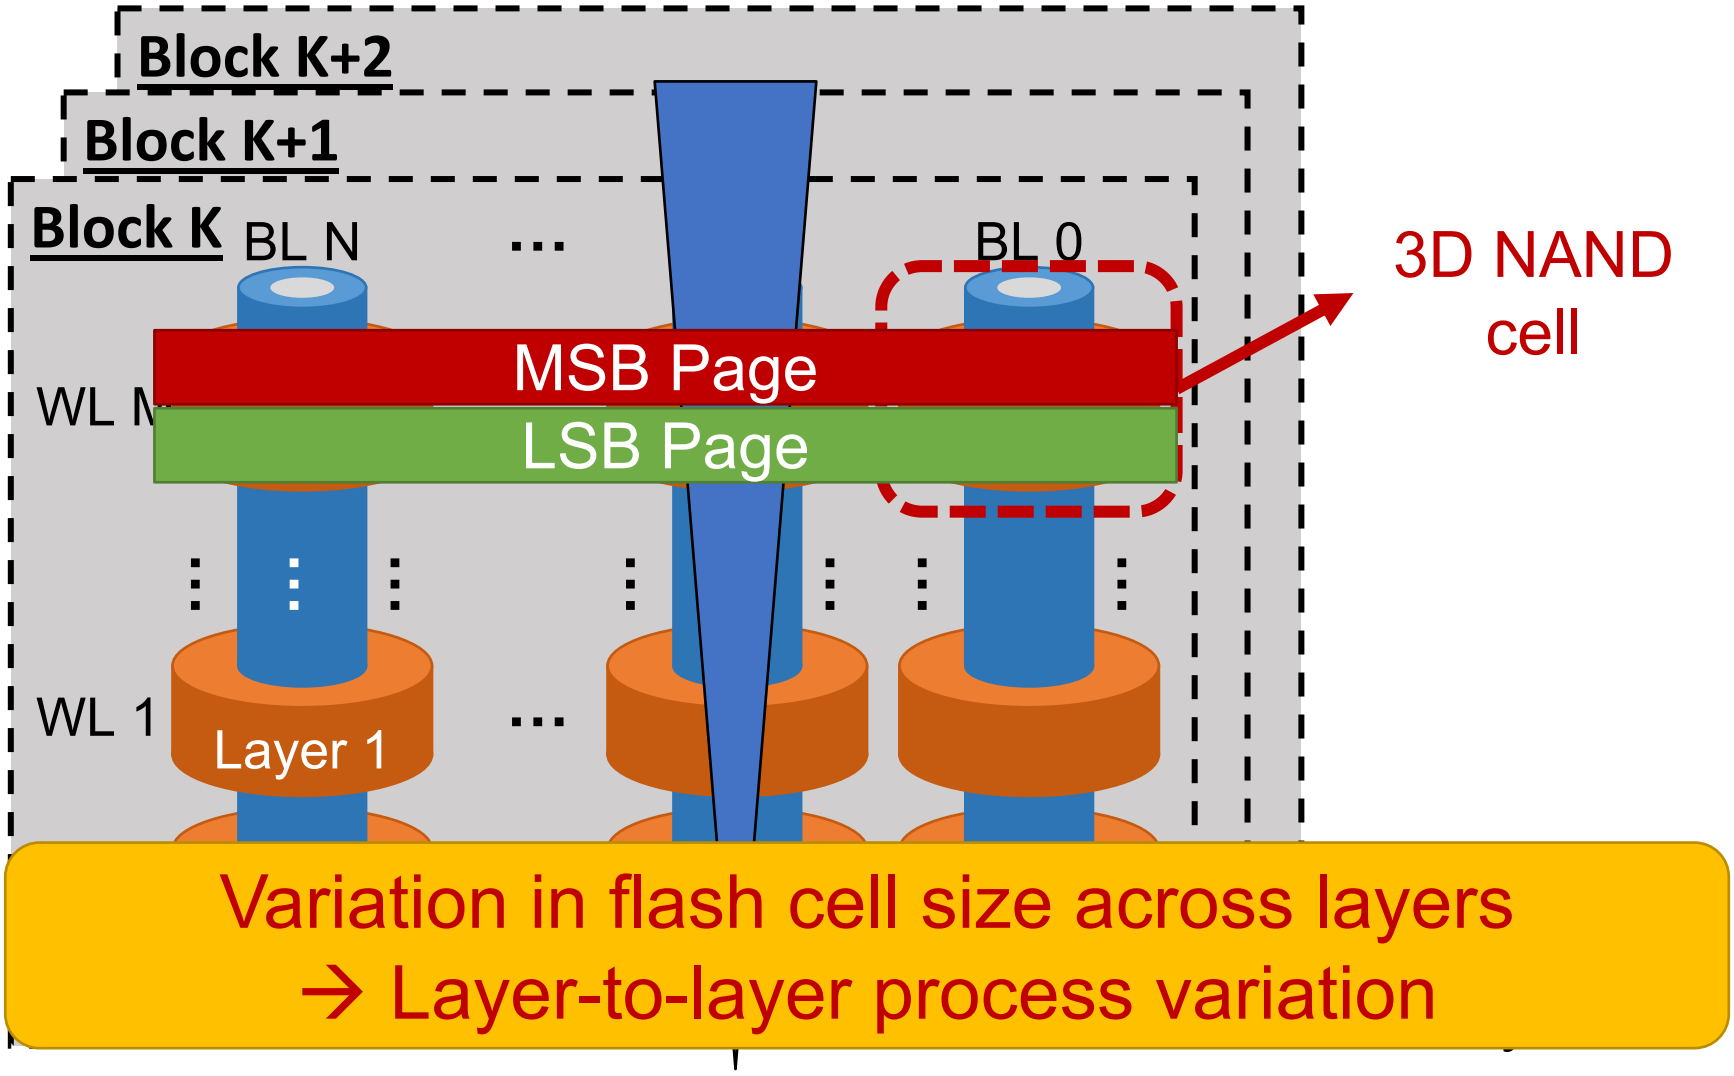

Tail RBER Problem

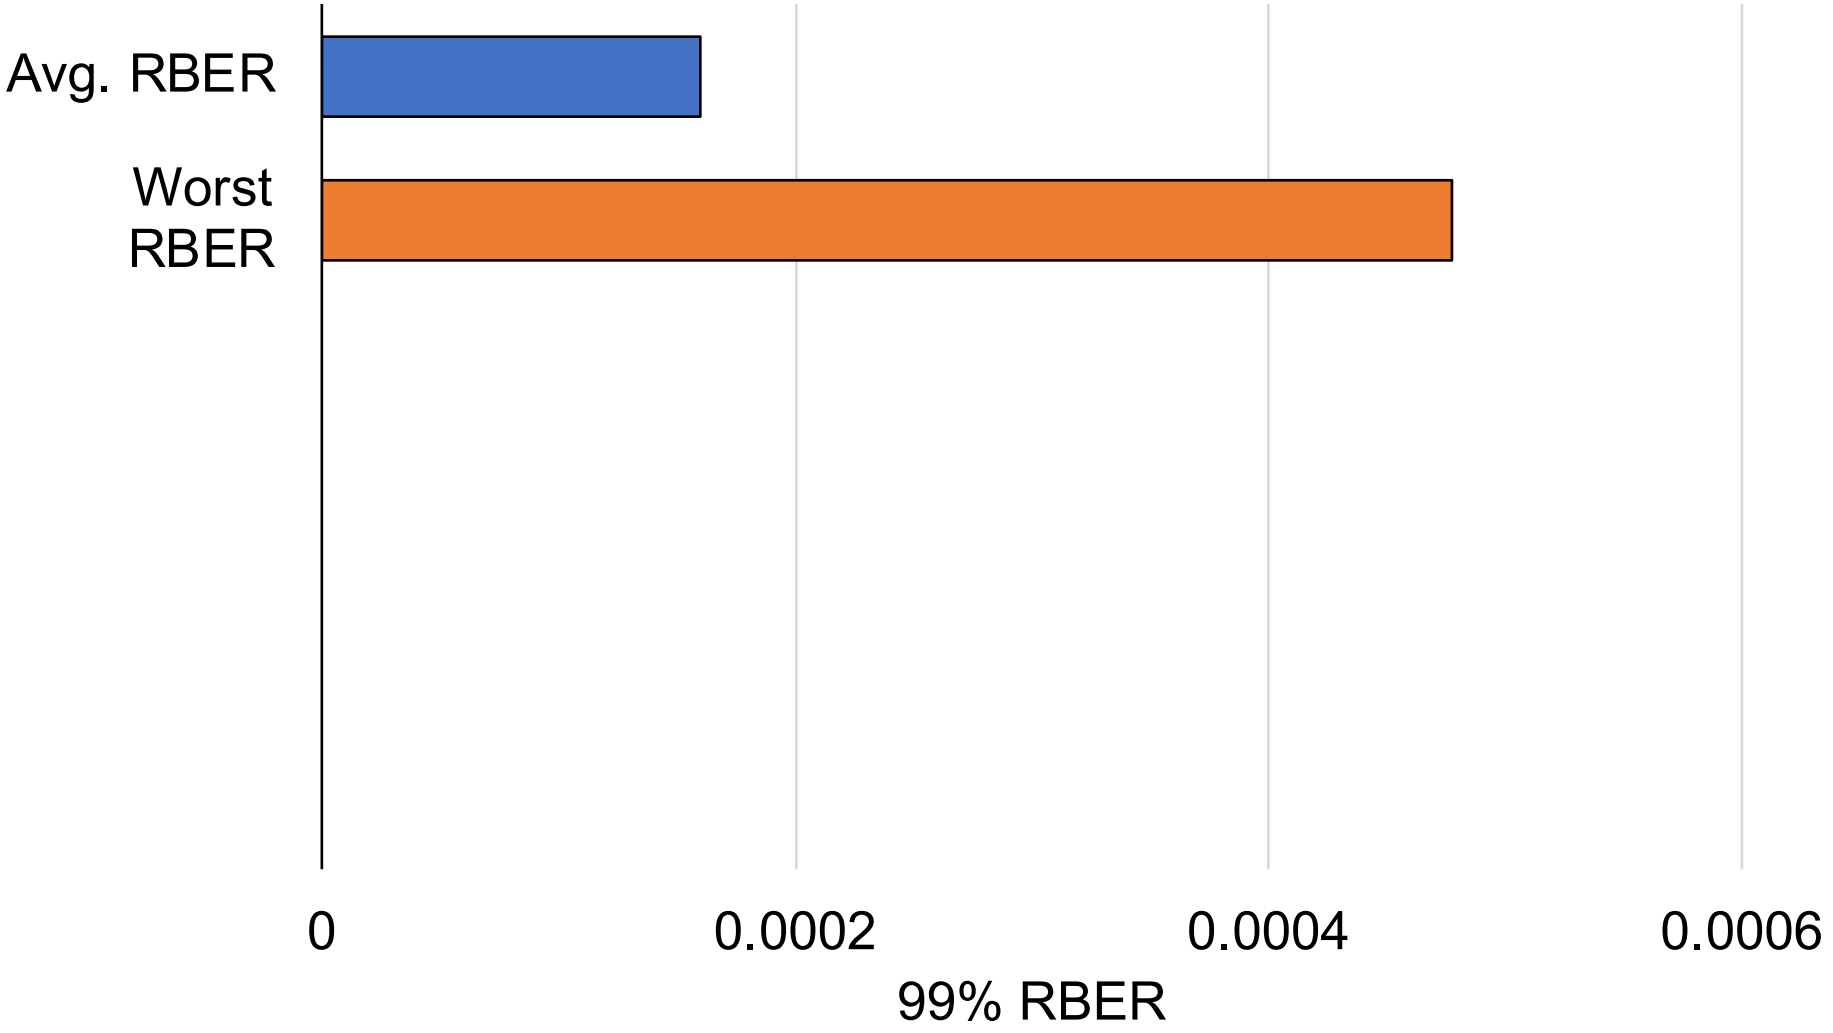

Adapting Optimal Read Ref. Voltage to Layer

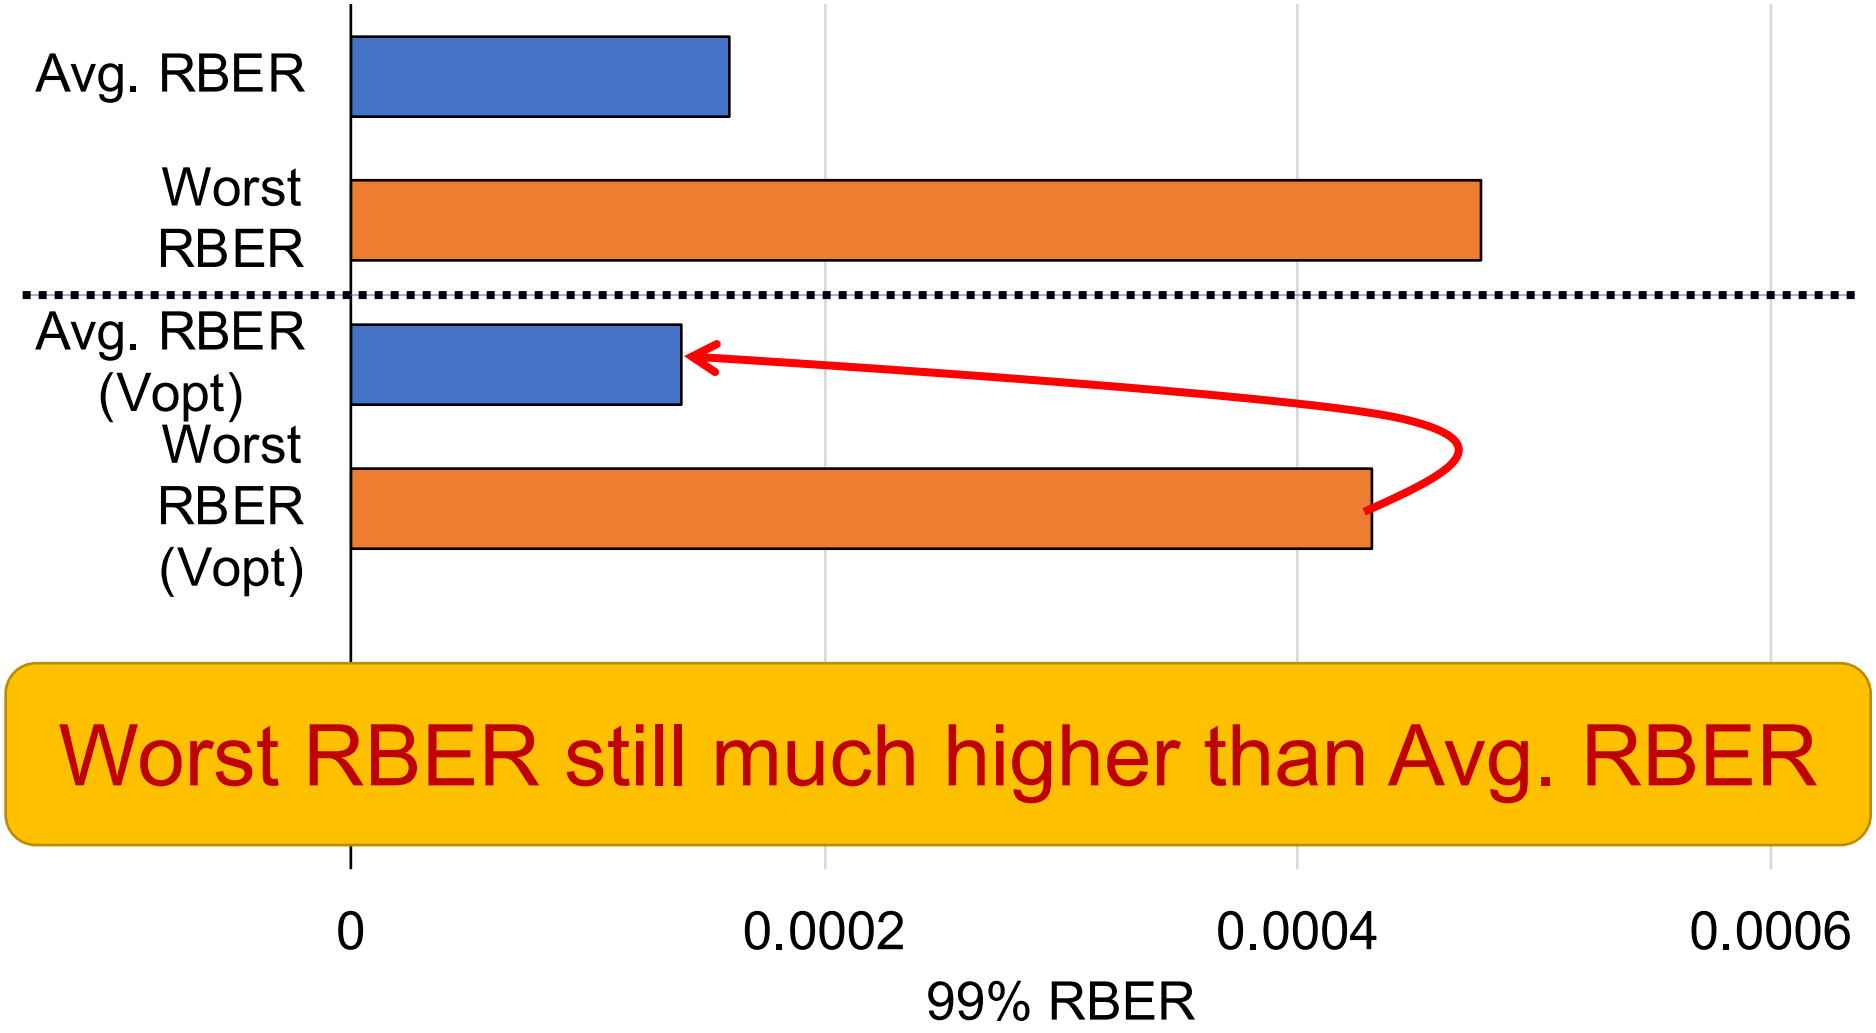

Conventional RAID for SSD

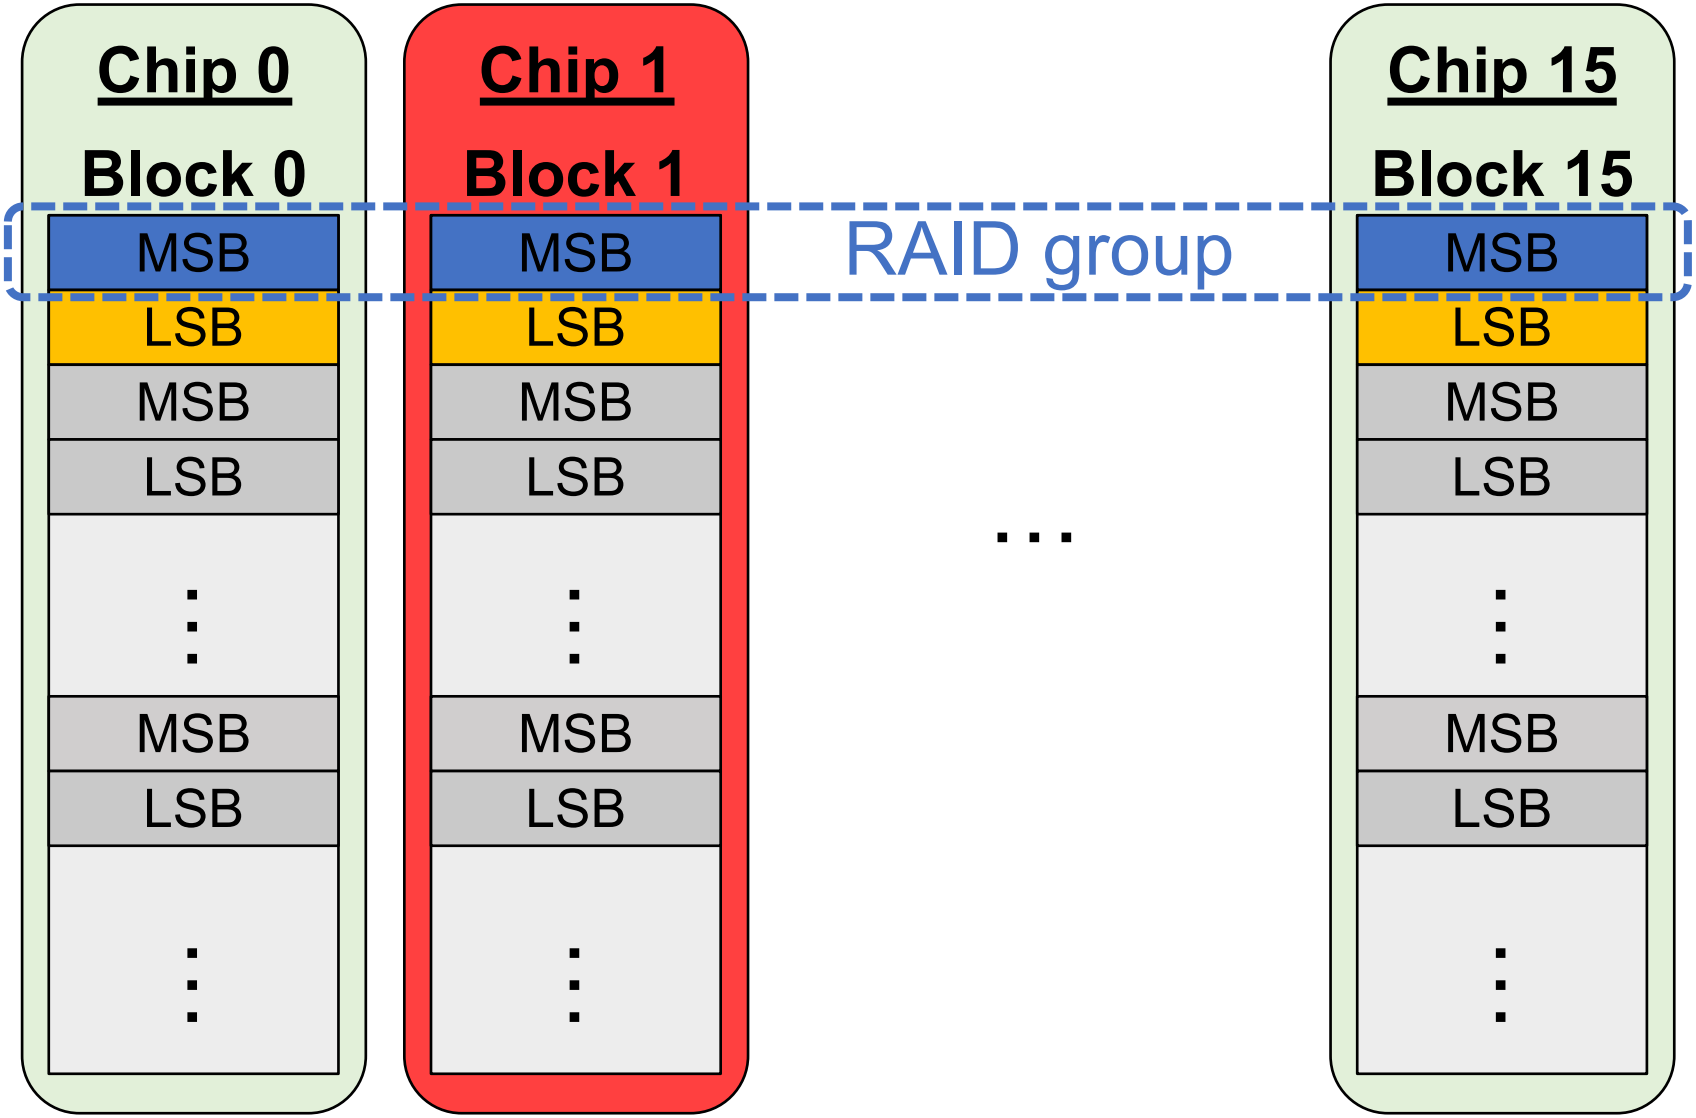

Layer-to-Layer Process Variation

- Limitations with conventional RAID
  - 1. Layer-to-layer process variation agnostic
    - *Middle layers have higher error rate*

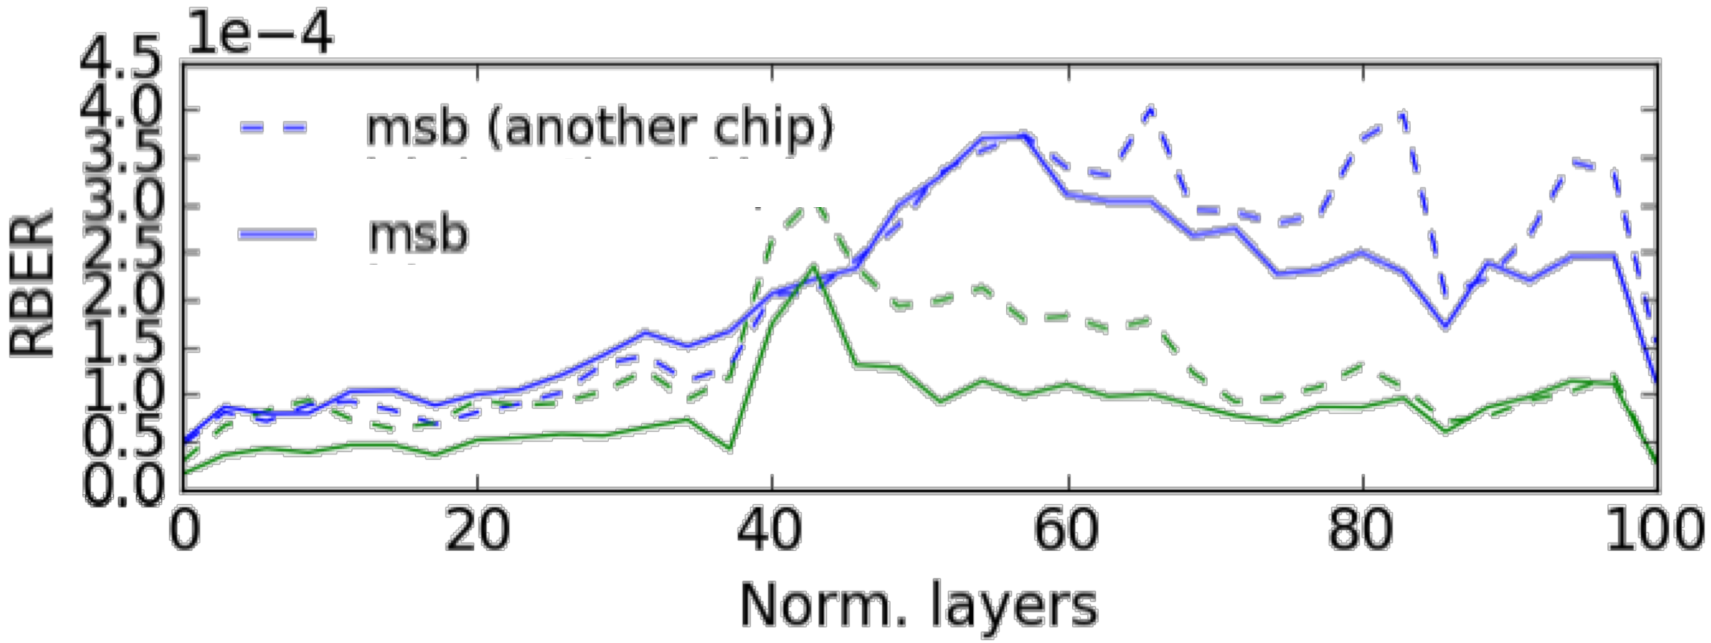

Conventional RAID for SSD

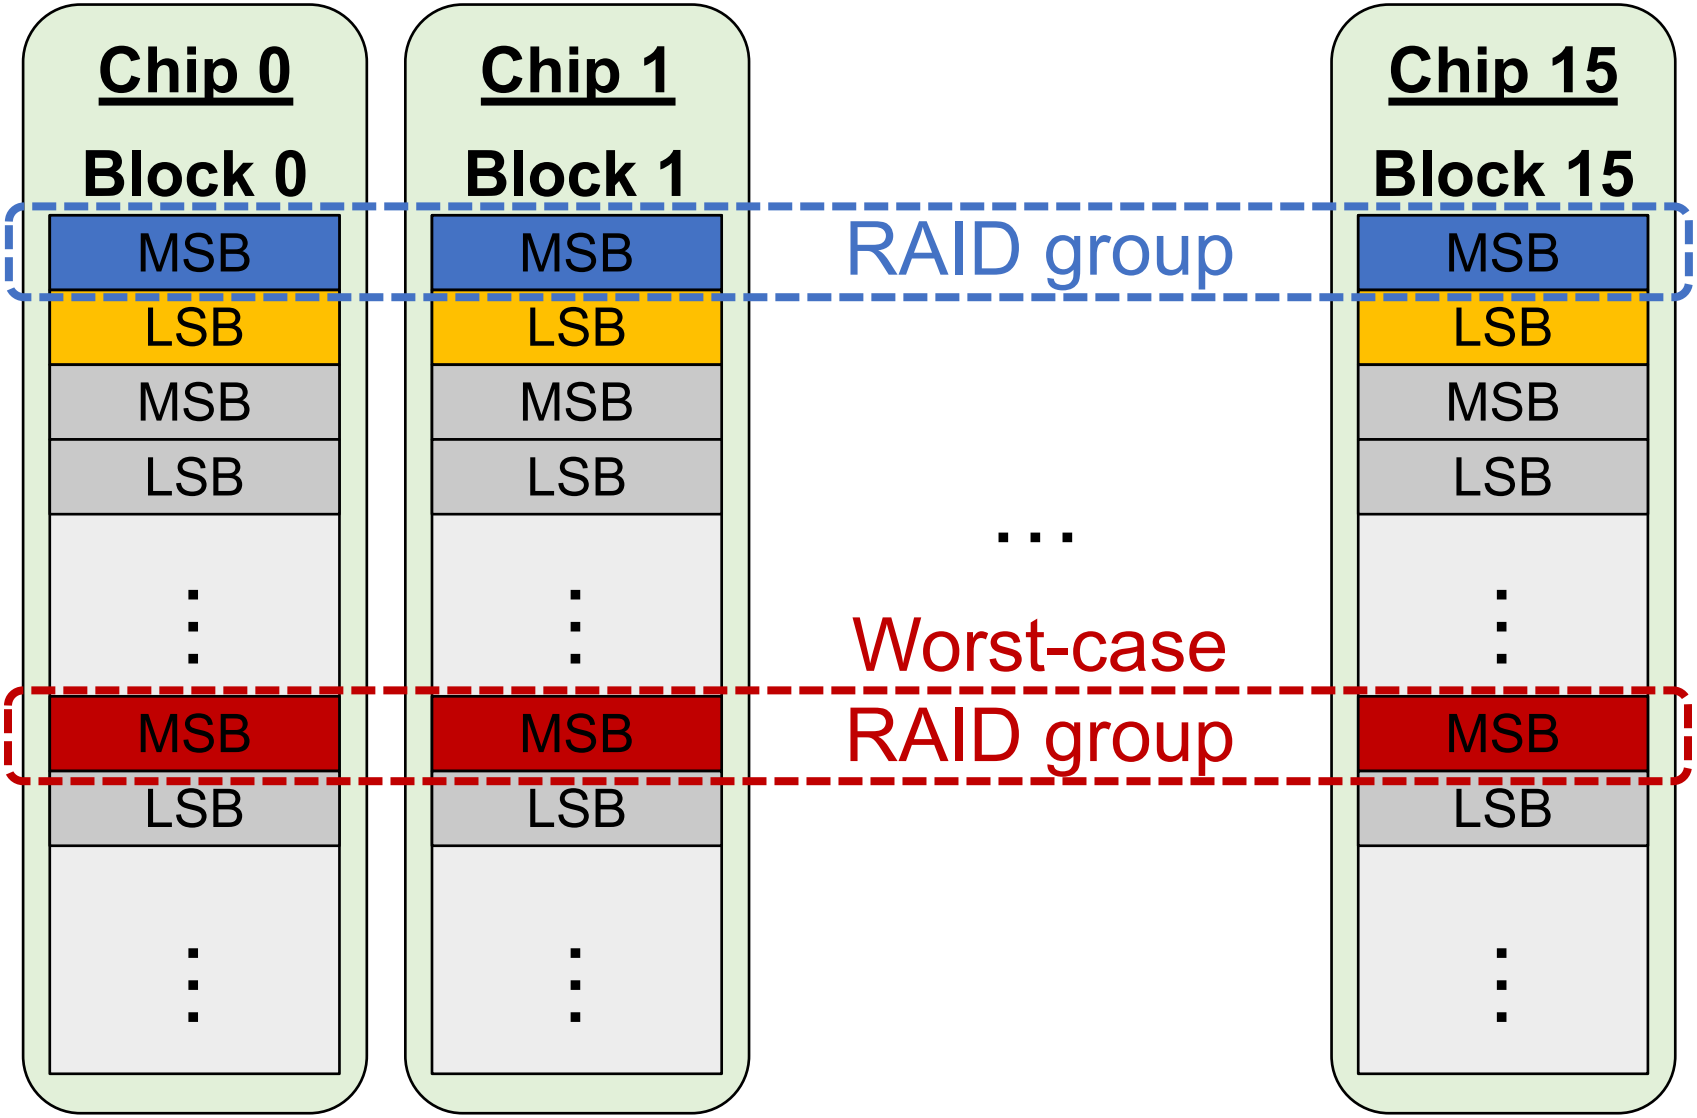

LI-RAID: Tolerating Layer-to-Layer Variation

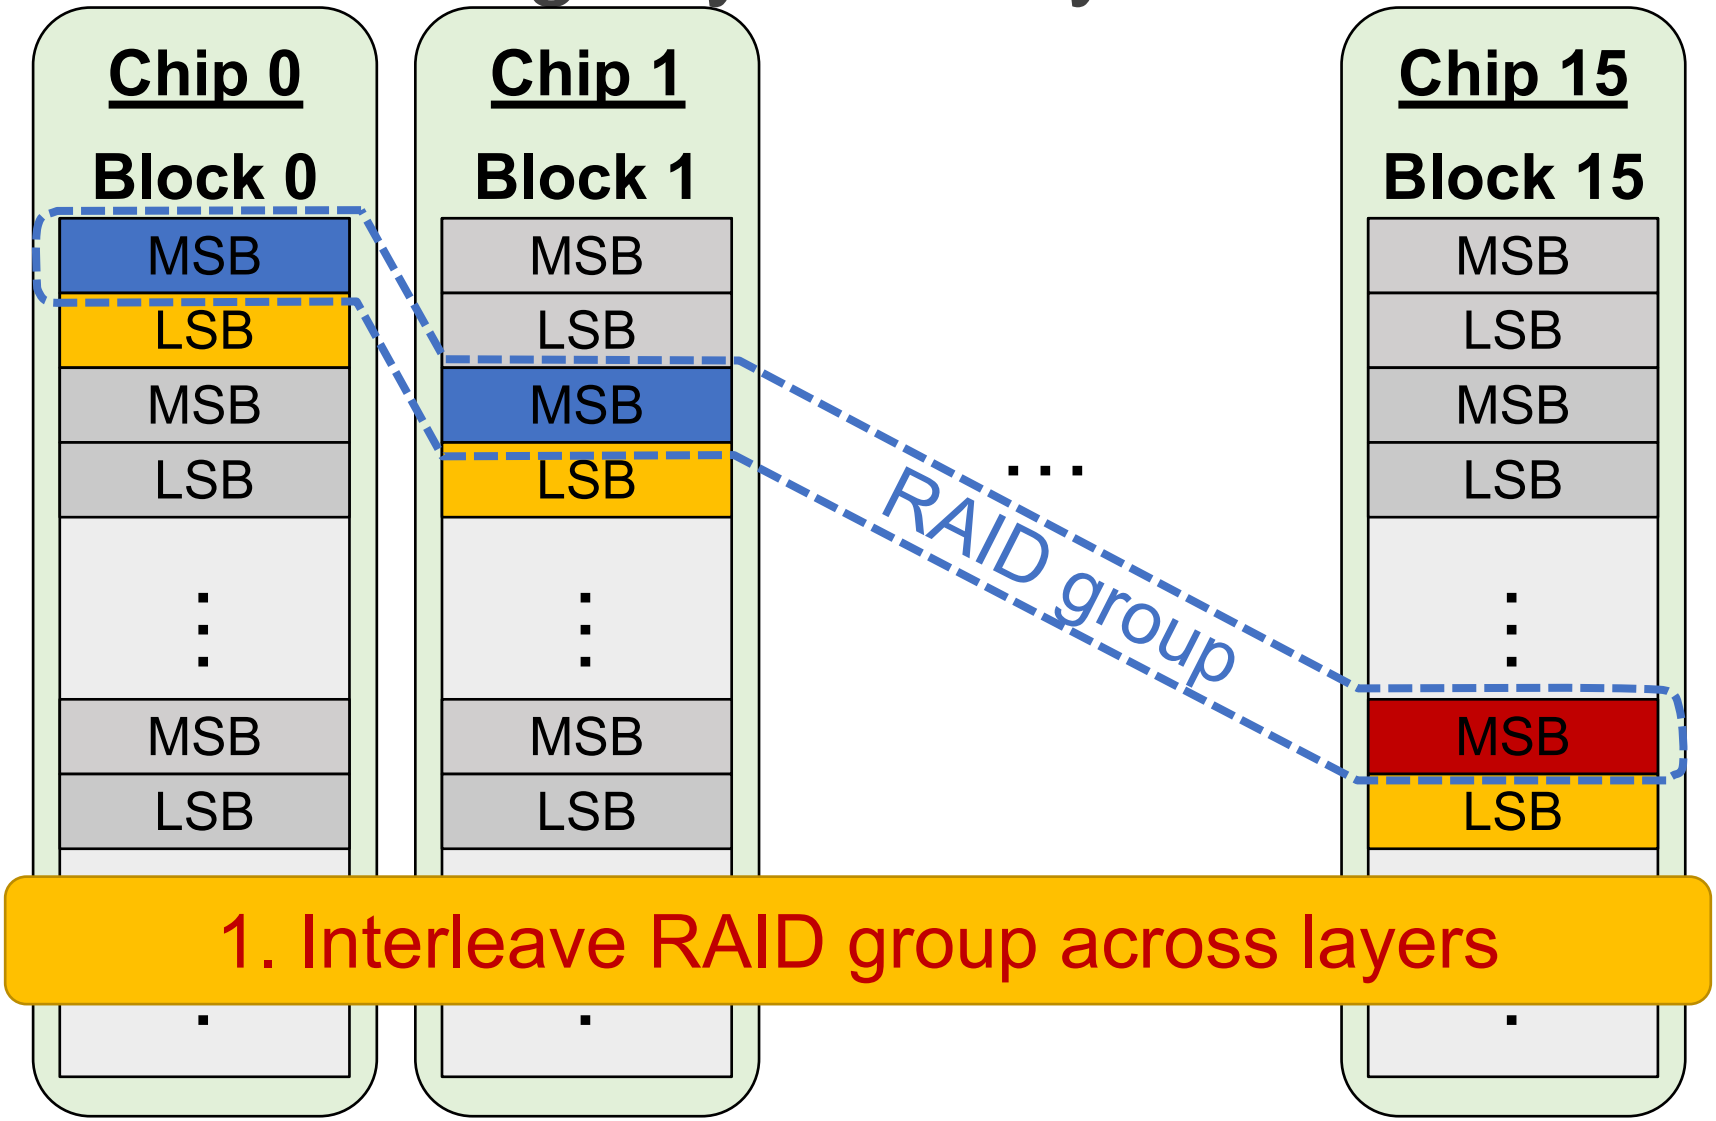

MSB-LSB Page Error Rate Variation

- Limitations with conventional RAID
  - 1. Layer-to-layer process variation agnostic
    - *Middle layers have higher RBER*
  - 2. MSB or LSB page agnostic
    - *MSB pages have higher RBER*

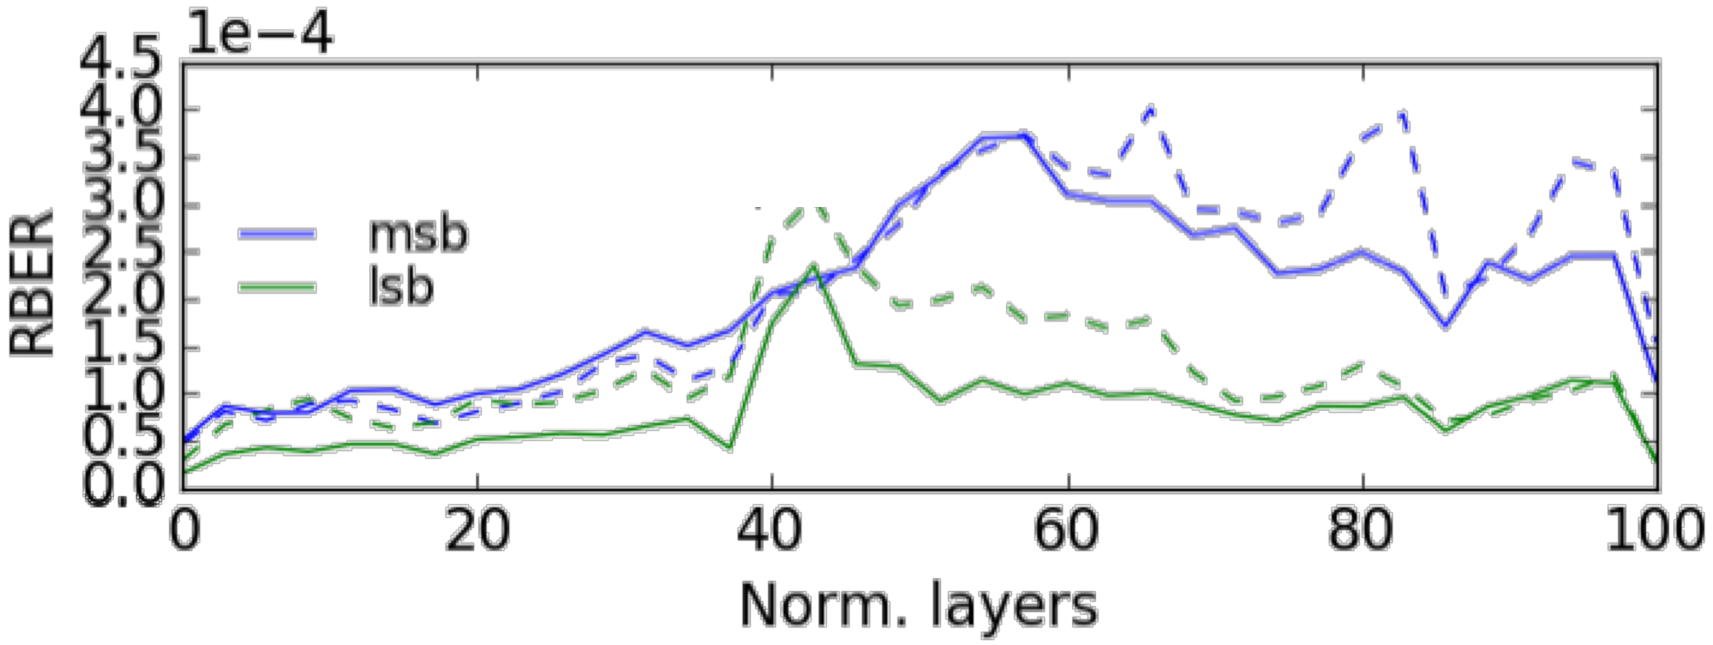

LI-RAID:
Tolerating MSB-LSB Error Rate Variation

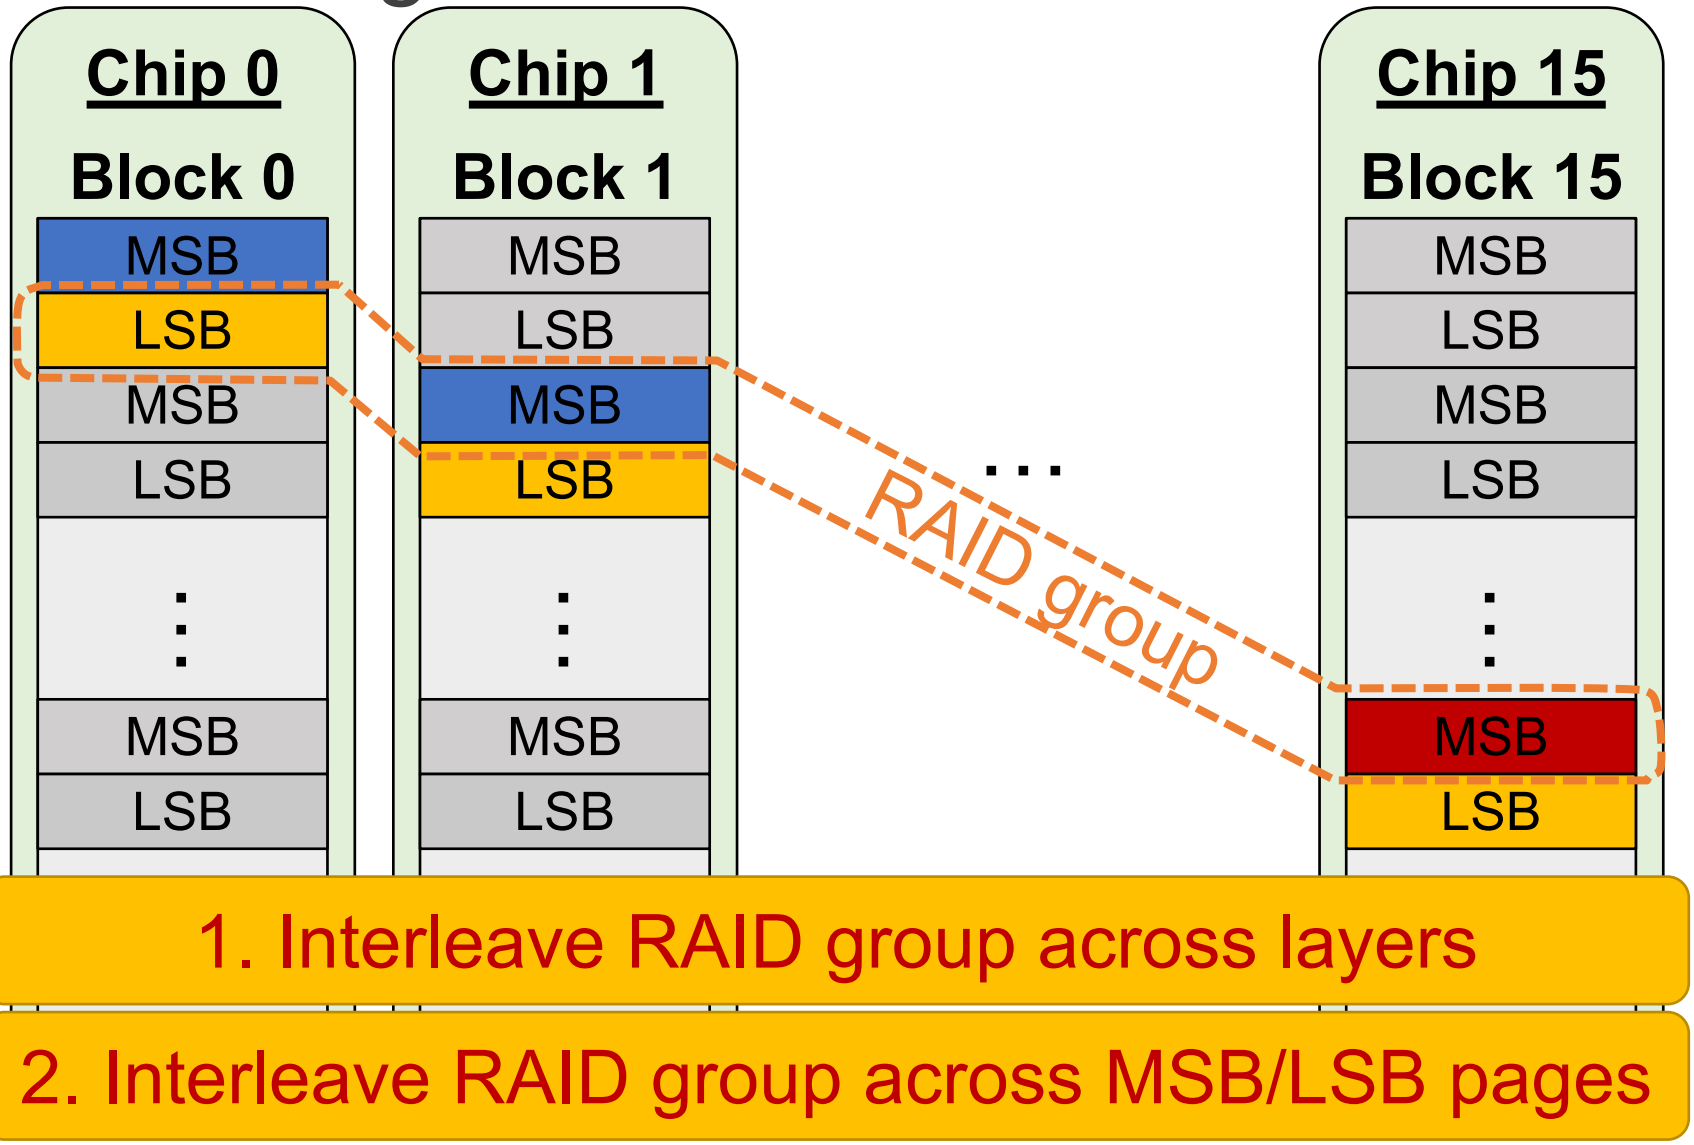

LI-RAID Evaluation

- Methodology
  - Based on characterization data at 10,000 P/E cycles
- Reliability
  - Improves MTTF by 9.1x over conventional RAID
- Overhead
  - No additional overhead on top of conventional RAID

Conclusions and Future Work

Goal: Improve Flash Reliability At Low Cost

3D NAND Errors, LI-RAID [under submission]

WARM [MSST ‘16]
HeatWatch [HPCA ‘18]

1. Flash Device Characteristics

2. Workload Characteristics

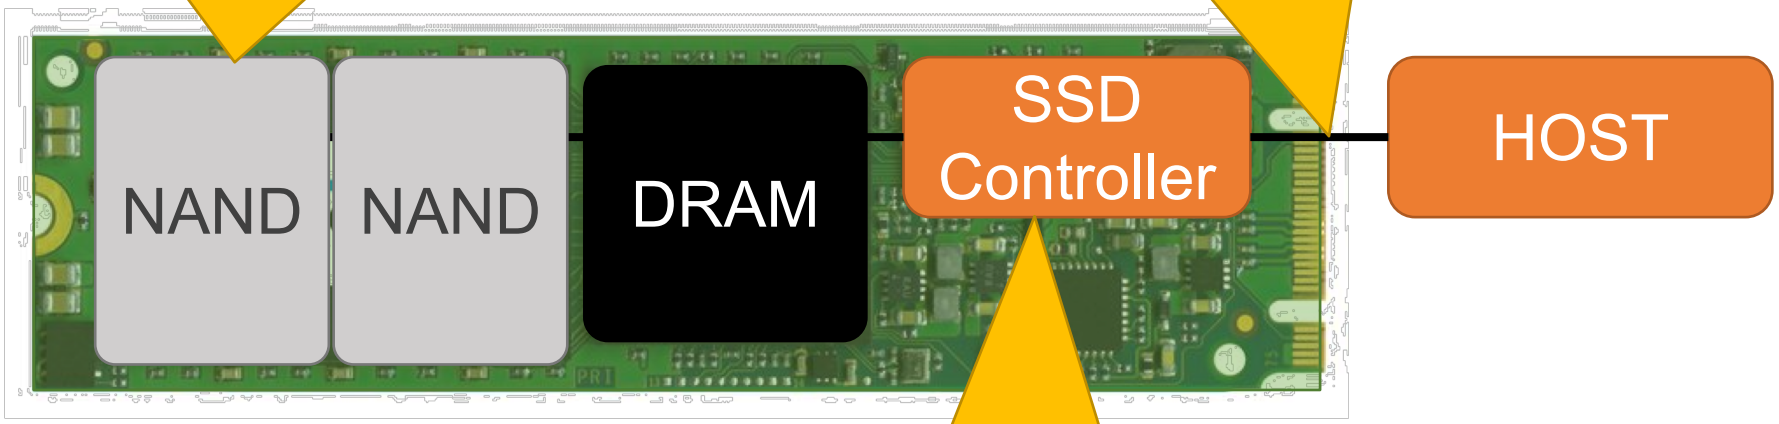

Online Flash Channel Modeling [JSAC ‘16]

3. Powerful Controller

Lessons Learned

- Specialization helps
  - Device characteristics
  - Workload characteristics
- Data-driven approach
  - Model-based techniques
  - Online model vs. fixed model
- Observation-driven research
  - Derive new insights through real characterization
  - New observation inspires new techniques

Future Research Directions

Manage unreliable cells in an SSD

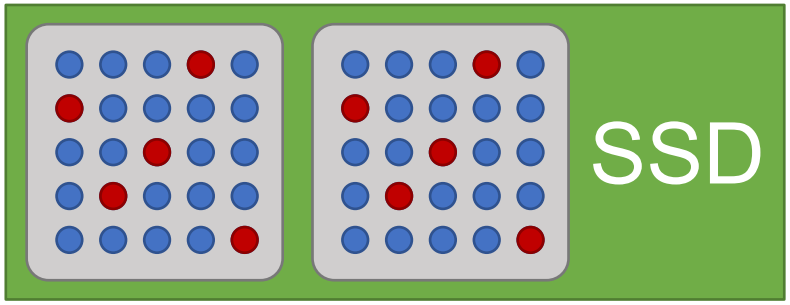

Manage unreliable SSDs in a data center

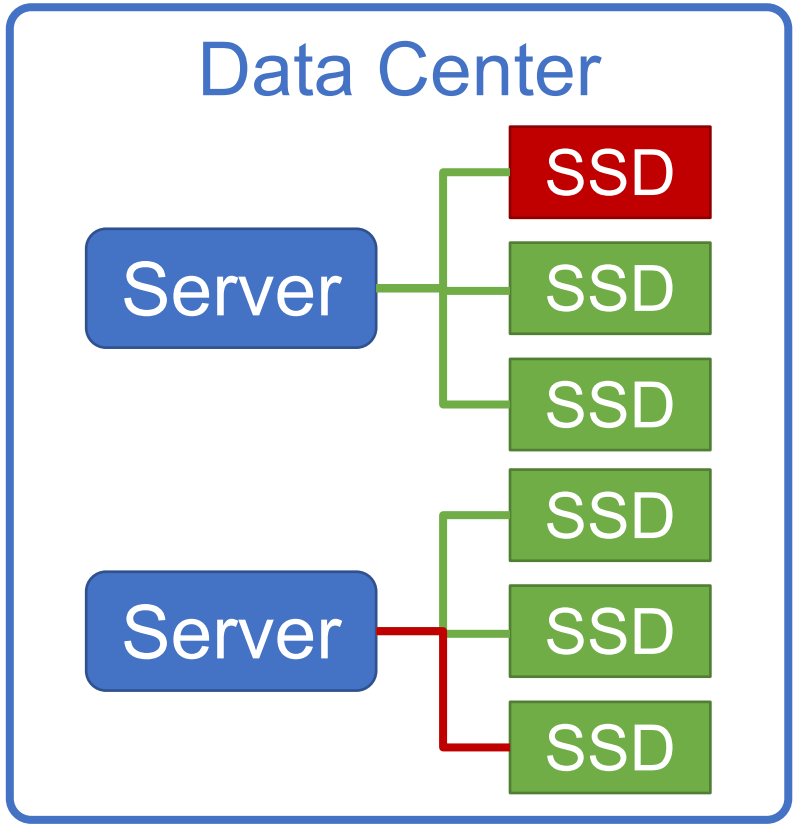

Future Research Directions

Data helps storage

- New models using machine learning/deep learning
- New techniques using reinforcement learning

Storage helps data

- Accommodate new applications: e.g., AI, DNA sequencing
- Accommodate new technologies: NVM, NVDIMM-F, zNAND
- Accommodate new storage architectures: distributed storage, single-level storage

Other Works During PhD

- Other NAND Flash Memory Reliability Works
  - [ProcIEEE '17], [HPCA '17], [DFRWS EU '17], [DSN '15], [HPCA '15]
- Heterogeneous-Reliability Memory
  - [DSN '14] [arXiv '17]
- Single-Level Storage
  - [WEED '13]
- Processing In Memory
  - [MICRO '13]

Acknowledgements

- Onur Mutlu
- Erich Haratsch and Yu Cai
- Phil Gibbons and James Hoe
- SAFARI
- Saugata Ghose
- Intern mentors and colleagues @ Seagate & MSR
- Deb!
- PDL and CALCM
- Friends
- Family

References (In Thesis)

- NAND flash-based SSD reliability
  - Enabling Accurate and Practical Online Flash Channel Modeling for Modern MLC NAND Flash Memory
    Yixin Luo, Saugata Ghose, Yu Cai, Erich F. Haratsch, and Onur Mutlu
    IEEE JSAC, 2016
  - WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management
    Yixin Luo, Yu Cai, Saugata Ghose, Jongmoo Choi, and Onur Mutlu
    MSST-31, 2015
  - Error Patterns in 3D NAND Flash Memory Devices: Characterization, Modeling, and Mitigation
    Yixin Luo, Saugata Ghose, Yu Cai, Erich F. Haratsch, and Onur Mutlu
    under submission, 2017
  - HeatWatch: Improving 3D NAND Flash Memory Device Reliability by Exploiting Self-Recovery and Temperature-Awareness
    Yixin Luo, Saugata Ghose, Yu Cai, Erich F. Haratsch, and Onur Mutlu
    HPCA-24, 2018

References (Thesis Related)

- NAND flash-based SSD reliability
  - Error Characterization, Mitigation, and Recovery in Flash Memory Based Solid-State Drives
    Yu Cai, Saugata Ghose, Erich F. Haratsch, Yixin Luo, and Onur Mutlu
    Proceedings of the IEEE, 2017 (Invited Paper)
  - Vulnerabilities in MLC NAND Flash Memory Programming: Experimental Analysis, Exploits, and Mitigation Techniques
    Yu Cai, Saugata Ghose, Yixin Luo, Ken Mai, Onur Mutlu, and Erich F. Haratsch
    HPCA-23, 2017
  - Improving the Reliability of Chip-Off Forensic Analysis of NAND Flash Memory Devices
    Aya Fukami, Saugata Ghose, Yixin Luo, Yu Cai, and Onur Mutlu
    DFRWS EU, 2017 — Best Paper Award
  - Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery
    Yu Cai, Yixin Luo, Erich F. Haratsch, Ken Mai, and Onur Mutlu
    HPCA-21, 2015 — Best Paper Runner Up
  - Read Disturb Errors in MLC NAND Flash Memory: Characterization, Mitigation, and Recovery
    Yu Cai, Yixin Luo, Saugata Ghose, Erich F. Haratsch, Ken Mai, and Onur Mutlu
    DSN-45, 2015

References (Other Works During PhD)

- Heterogeneous-Reliability Memory
  - Characterizing Application Memory Error Vulnerability to Optimize Data Center Cost via Heterogeneous-Reliability Memory
    Yixin Luo, Sriram Govindan, Bikash Sharma, Mark Santaniello, Justin Meza, Aman Kansal, Jie Liu, Badriddine Khessib, Kushagra Vaid, and Onur Mutlu
    DSN-44, 2014
  - Using ECC DRAM to Adaptively Increase Memory Capacity
    Yixin Luo, Saugata Ghose, Tianshi Li, Sriram Govindan, Bikash Sharma, Bryan Kelly, Amirali Boroumand, Onur Mutlu
    under submission, 2017
- Processing in memory
  - RowClone: Fast and Energy-Efficient In-DRAM Bulk Data Copy and Initialization
    Vivek Seshadri, Yoongu Kim, Chris Fallin, Donghyuk Lee, Rachata Ausavarungnirun, Gennady Pekhimenko, Yixin Luo, Onur Mutlu, Michael A. Kozuch, Phillip B. Gibbons, and Todd C. Mowry
    MICRO-46, 2013
- Single-Level Storage
  - A Case for Efficient Hardware-Software Cooperative Management of Storage and Memory
    Justin Meza, Yixin Luo, Samira Khan, Jishen Zhao, Yuan Xie, and Onur Mutlu
    WEED-5, 2013

Architectural Techniques for Improving NAND Flash Memory Reliability

Thesis Oral

Yixin Luo

Committee:

Onur Mutlu (Chair)

Phillip B. Gibbons

James C. Hoe

Erich F. Haratsch, Seagate

Yu Cai, SK Hynix

Carnegie Mellon

Presented in partial fulfilment of the requirements for the degree of Doctor of Philosophy

Backup Slides

Error Correction Code (ECC)

- Key Idea:
  - Use redundant bits to encode data bits

- Pros:
  - Avoids silent data corruption (Error Detection)
  - Increases data reliability (Error Correction)

- Cons:
  - Requires redundant ECC bits (High Cost)
  - Treats all errors as random, does not take advantage of error characteristics (Not Specialized)

NAND Flash Errors

- P/E cycling
  - Wear out
- Retention
  - Charge leakage
- Program interference
  - Coupling
- Read disturb
  - Weak programming
- Process variation

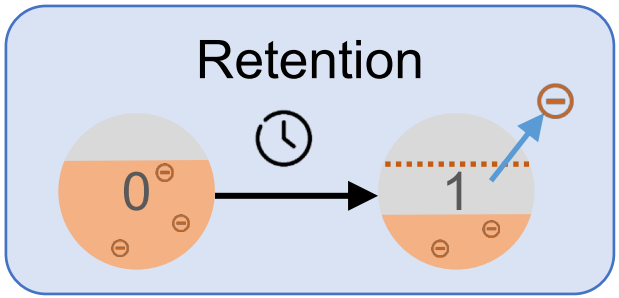

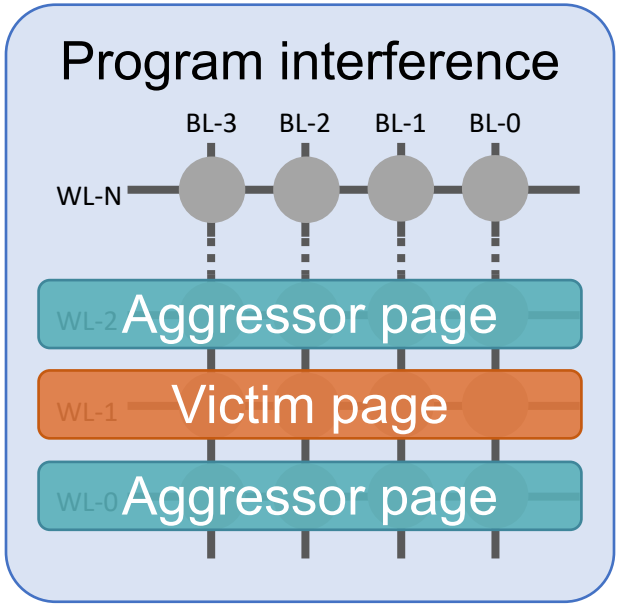

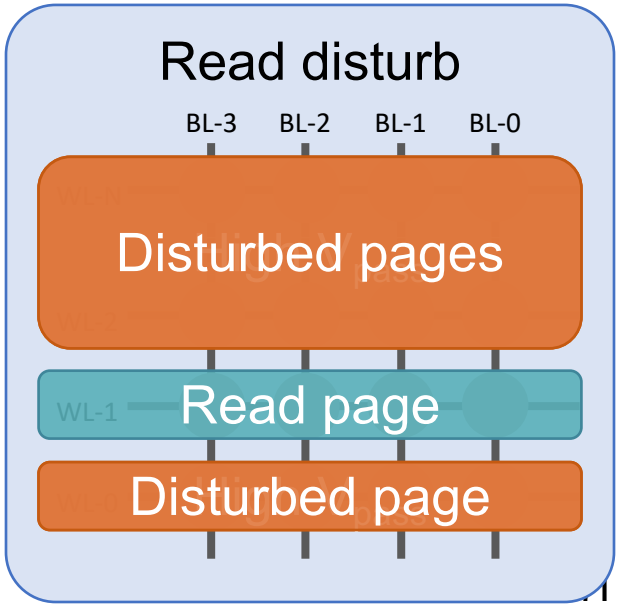

Future of Solid-State Drives (SSDs)

Capacity/Density

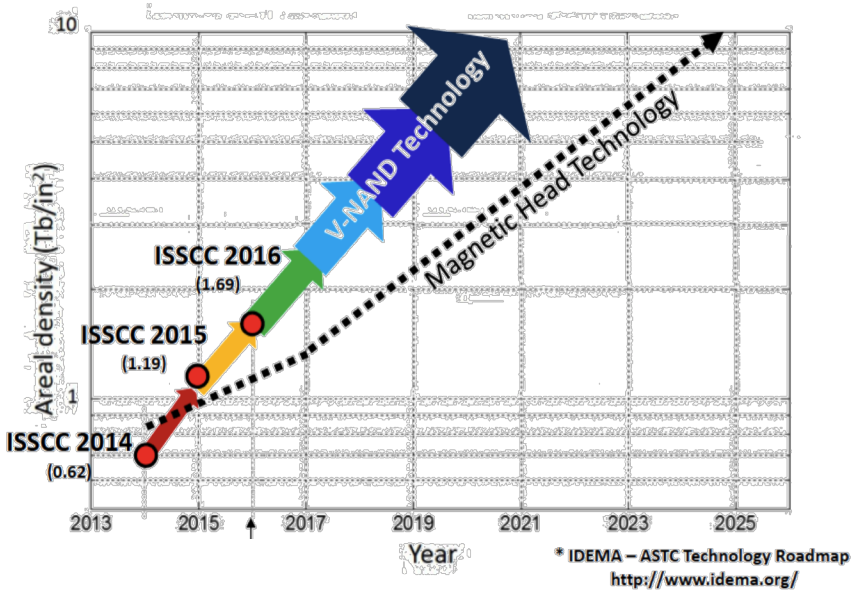

Cost

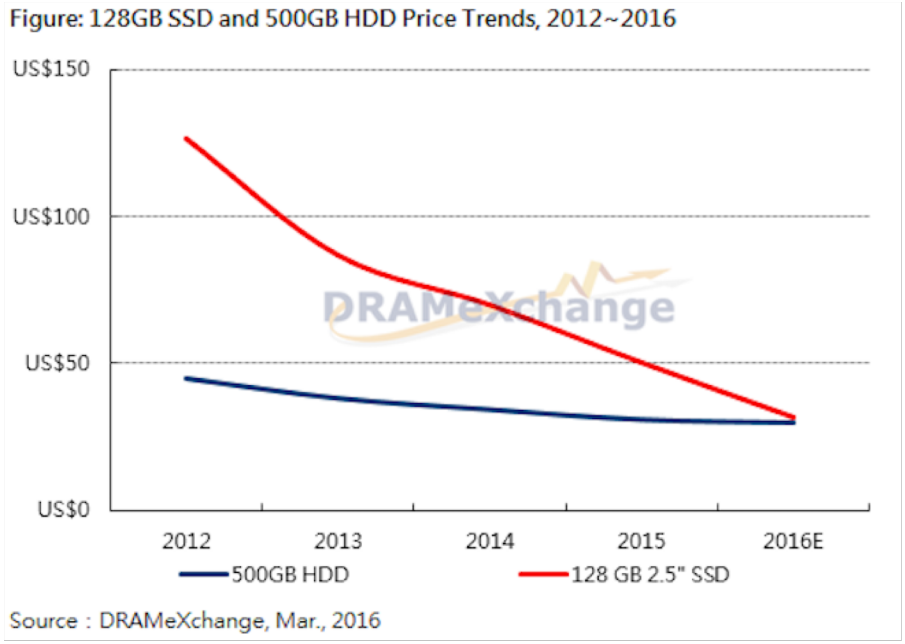

Image Source:
1. https://www.computerworld.com/article/3030642/data-storage/flash-memorys-density-surpasses-hard-drives-for-first-time.html
2. https://www.pcworld.com/article/3040591/storage/ssd-prices-plummet-again-close-in-on-hdds.html

Why do we care about SSD lifetime?

- Because SSD lifetime is an indicator of ***SSD reliability***
  - **Lifetime** – Errors increase with write cycles
    - *We are actually reducing SSD errors!*

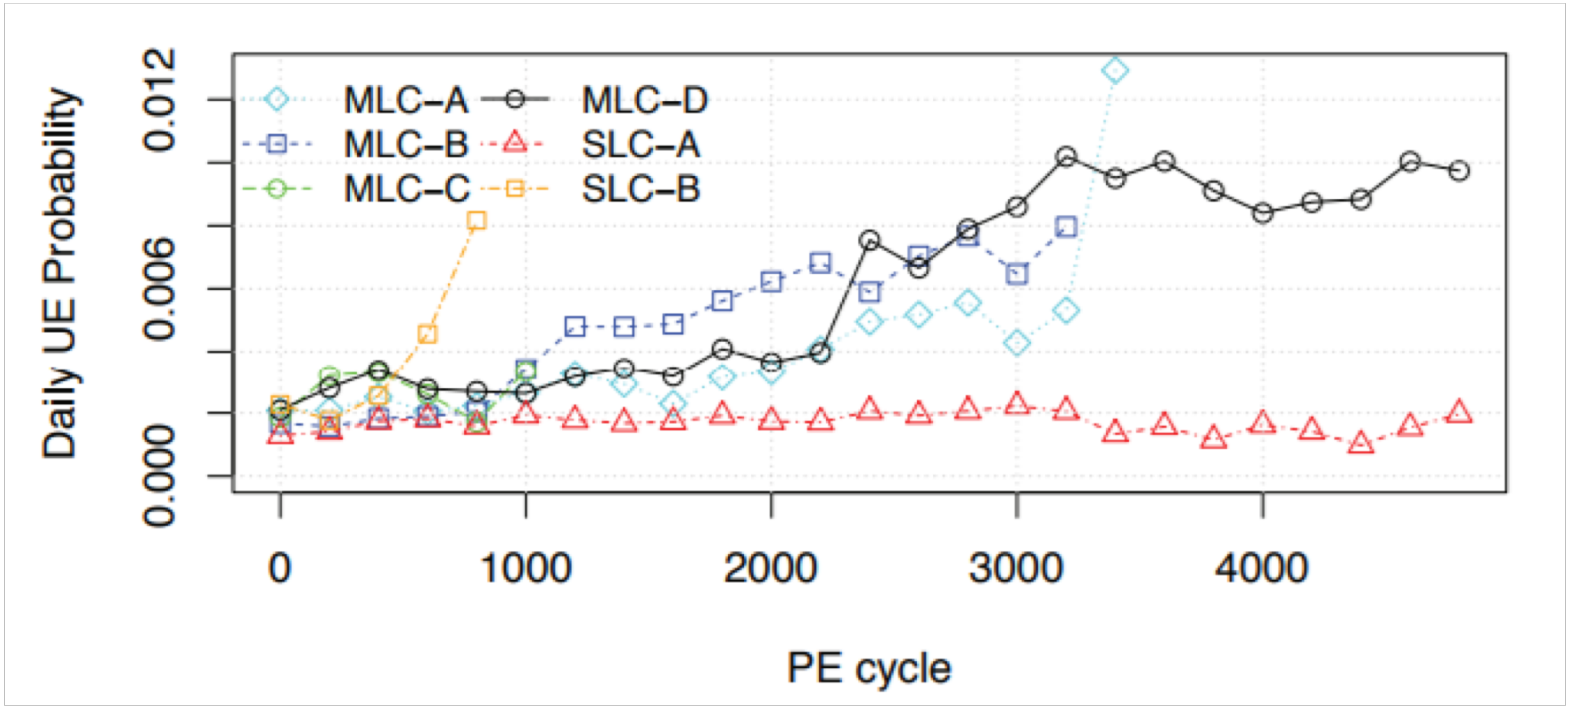

Image Source:
1. Bianca Schroeder, et al., "Flash Reliability in Production: The Expected and the Unexpected", FAST 2016.

Why do we care about SSD lifetime?

- Because SSD lifetime is an indicator of ***SSD reliability***
  - **Lifetime** – Errors increase with write cycles
    - *We are actually reducing SSD errors!*
  - **UE**: Uncorrectable errors or data corruption – Errors can lead to error correction failure
  - **Data retention** – Errors increase with retention time
  - **Performance**
    - *When SSD lifespan is fixed, limits drive writes per day*
    - *We can trade-off reliability for performance (Samsung zNAND)*
  - **Cost** – Errors increase as areal density increases

Mitigating Retention Improves Lifetime

Retention errors affect SSD reliability and lifetime.

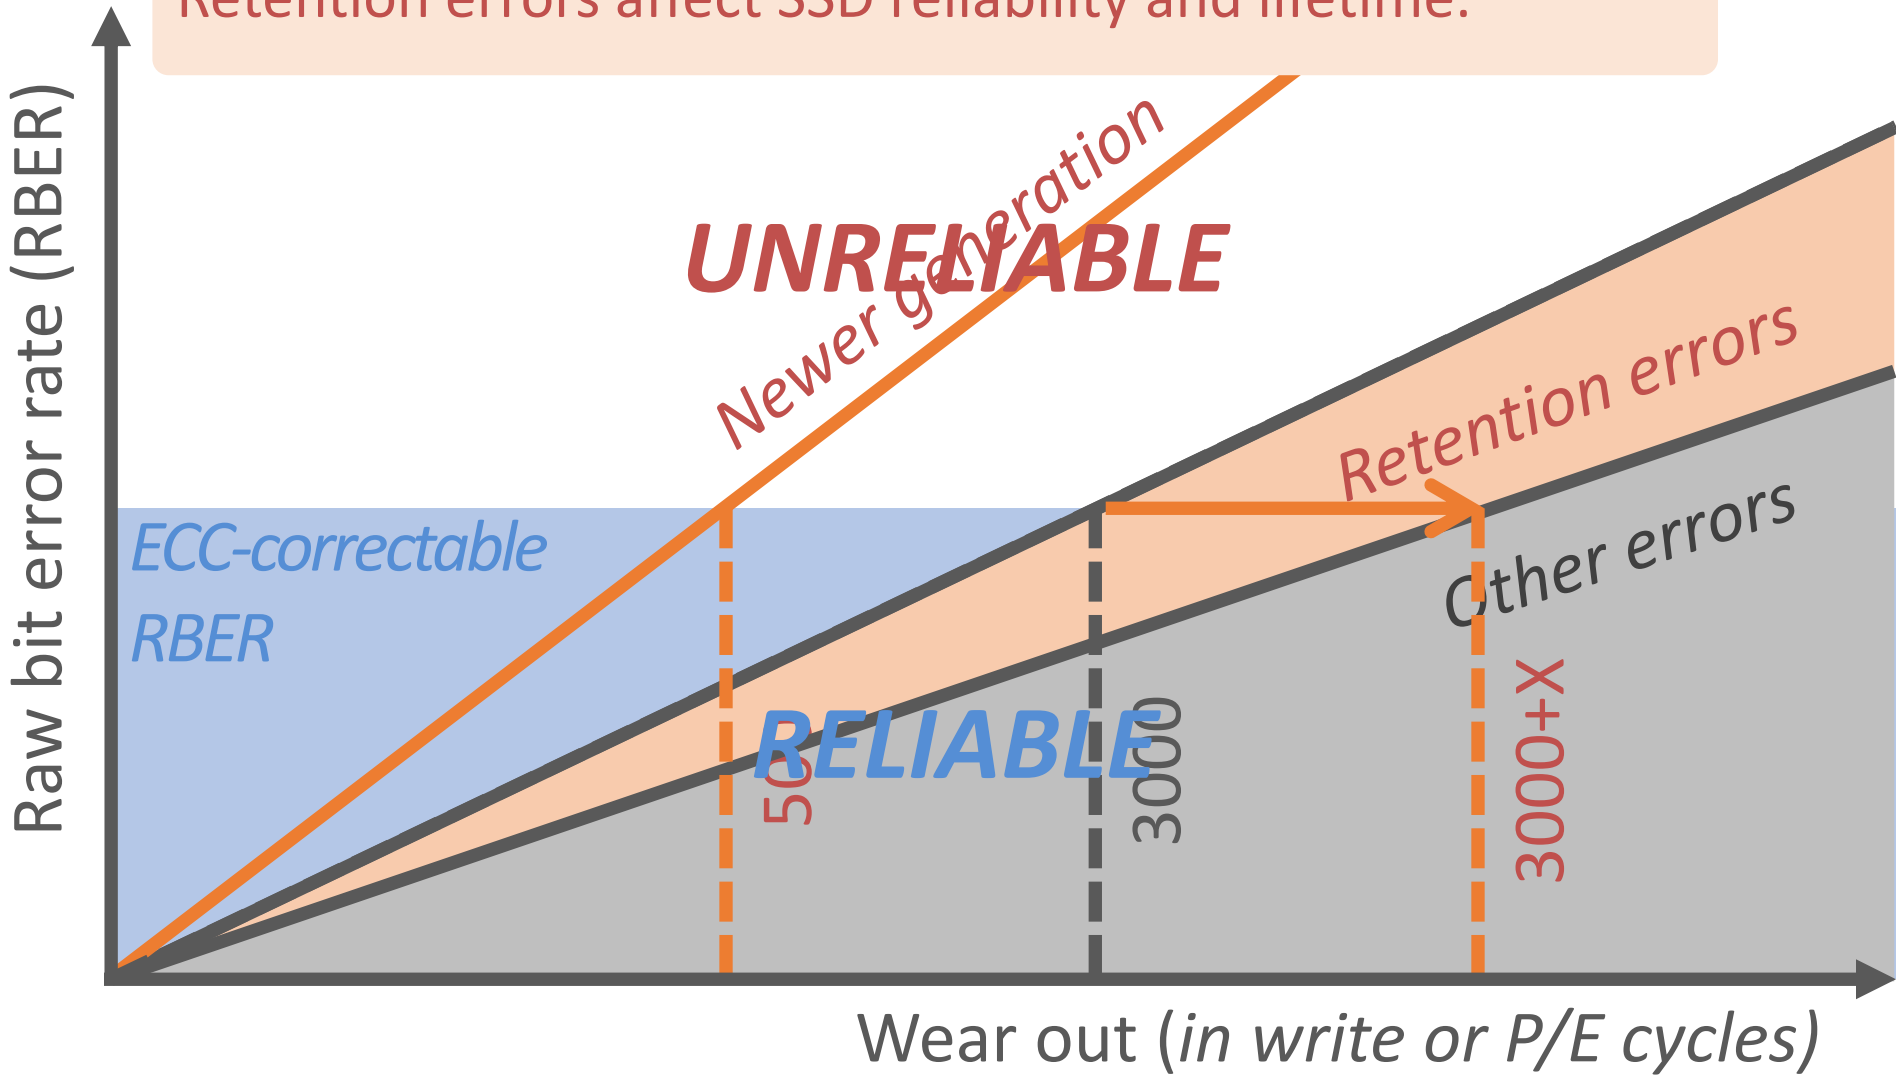

Conventional Retention Model Drawbacks

- Not designed for 3D NAND

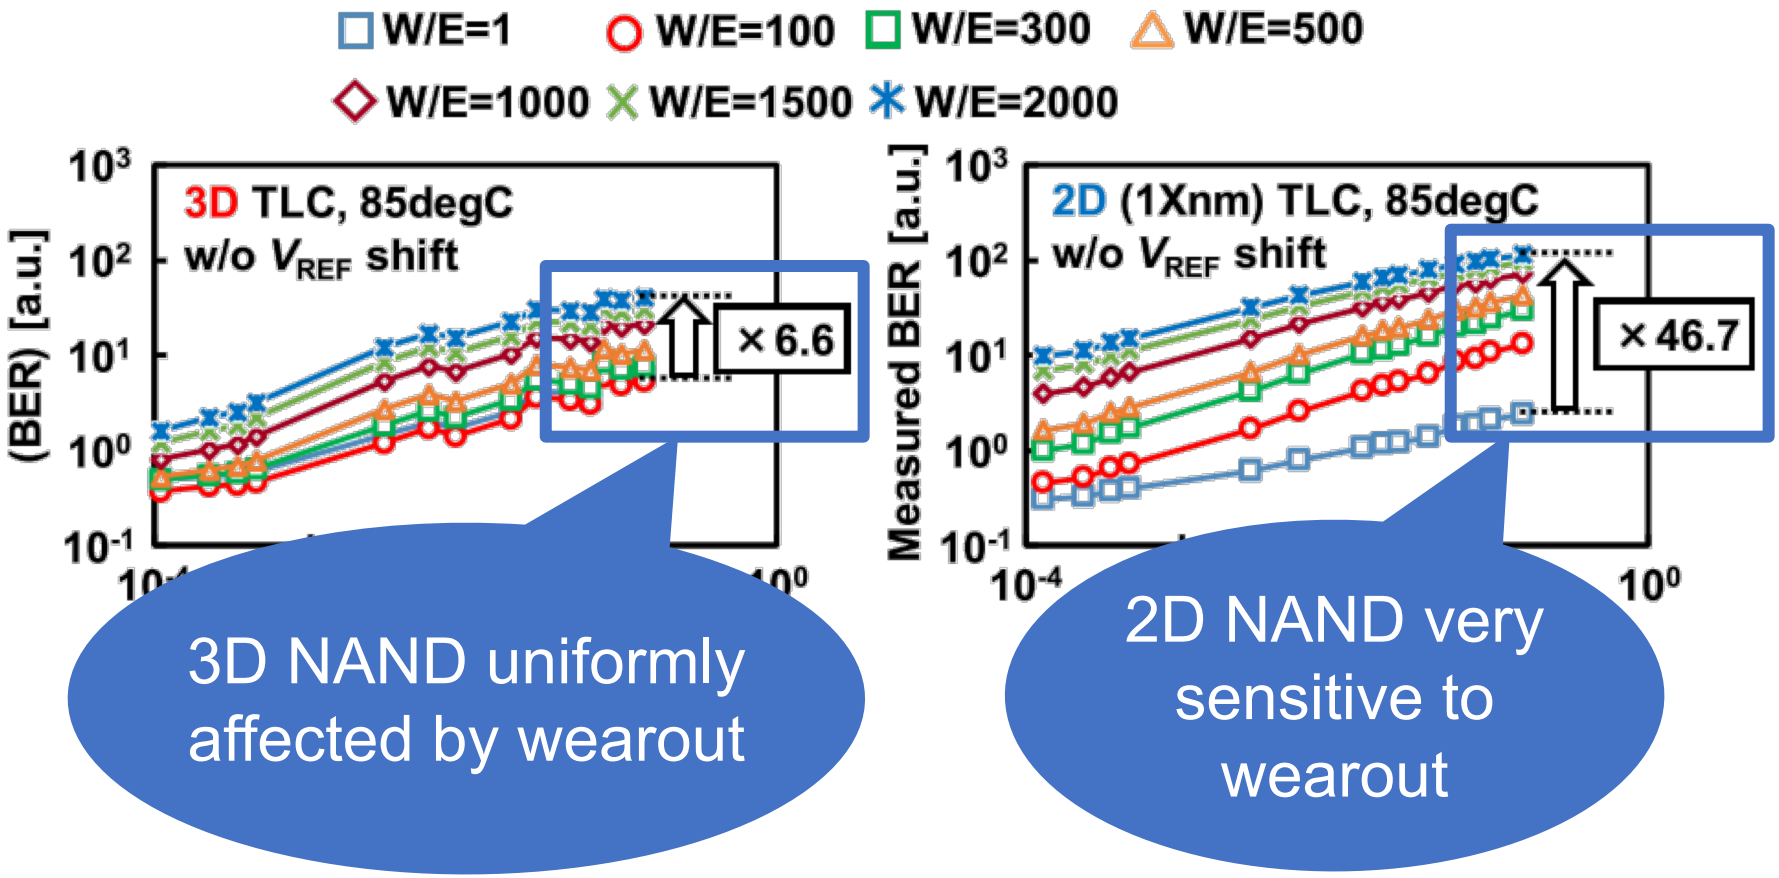

Source: K. Mizoguchi, et al., "Data-Retention Characteristics Comparison of 2D and 3D TLC NAND Flash Memories," IMW, 2017.

Conventional Retention Model Drawbacks

- Not designed for 3D NAND
- Retention temperature agnostic

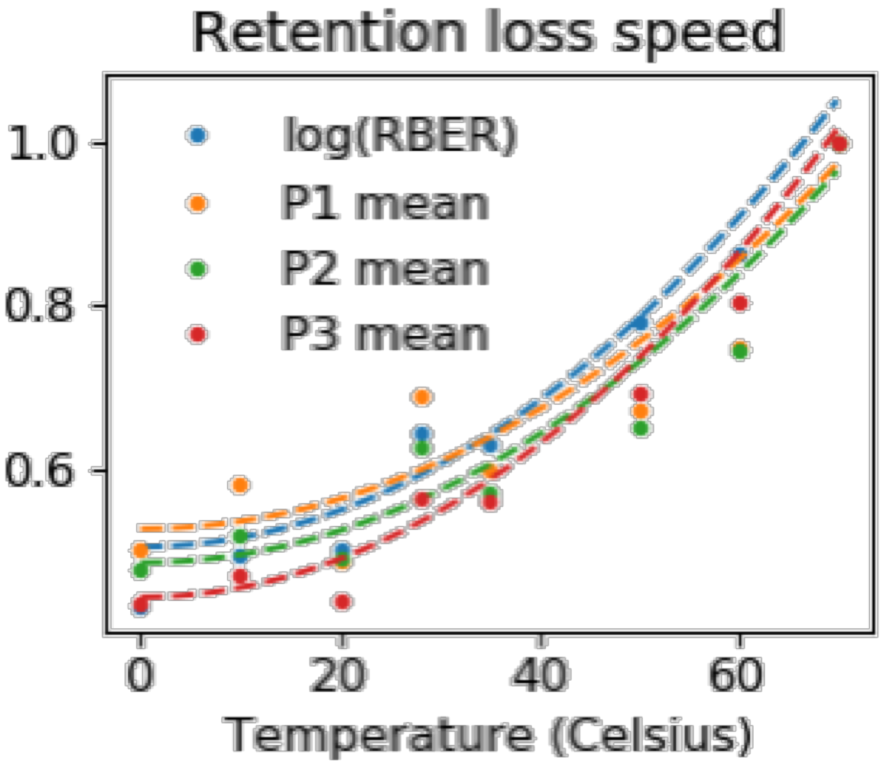

Conventional Retention Model Drawbacks

- Not designed for 3D NAND
- Retention temperature agnostic
- Dwell time agnostic
  - Dwell time between program cycles & temperature

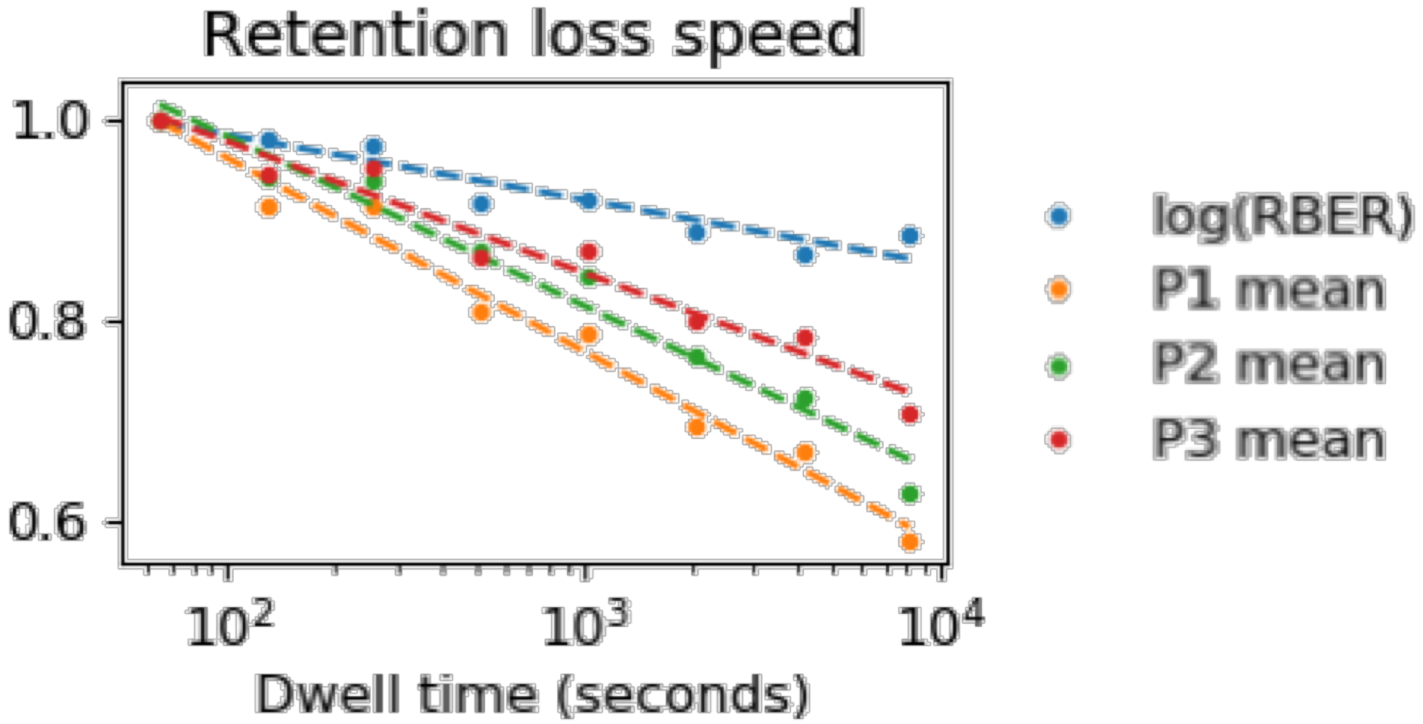

Conventional Retention Model Drawbacks

- Not designed for 3D NAND
- Retention temperature agnostic
- Dwell time agnostic
- Programming temperature agnostic

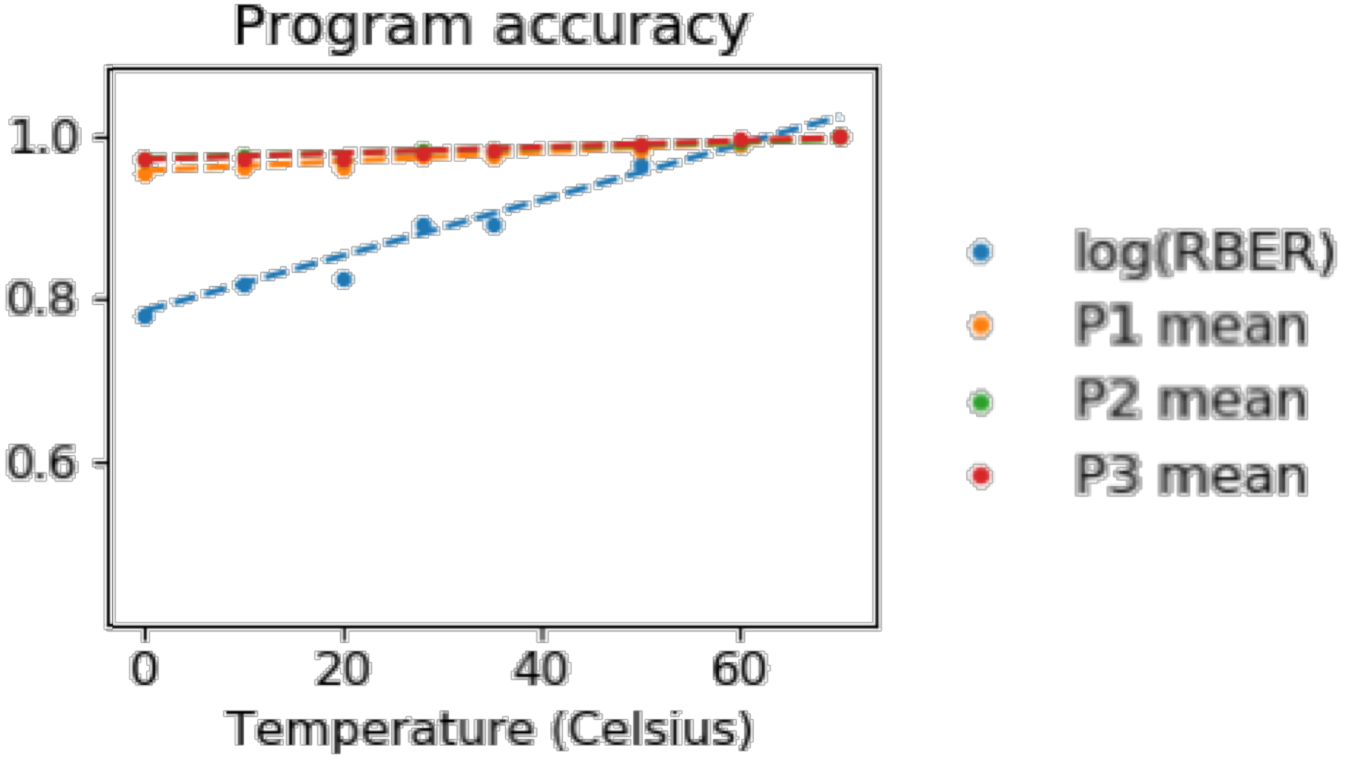

Optimal Read Ref. Voltage for Process Variation

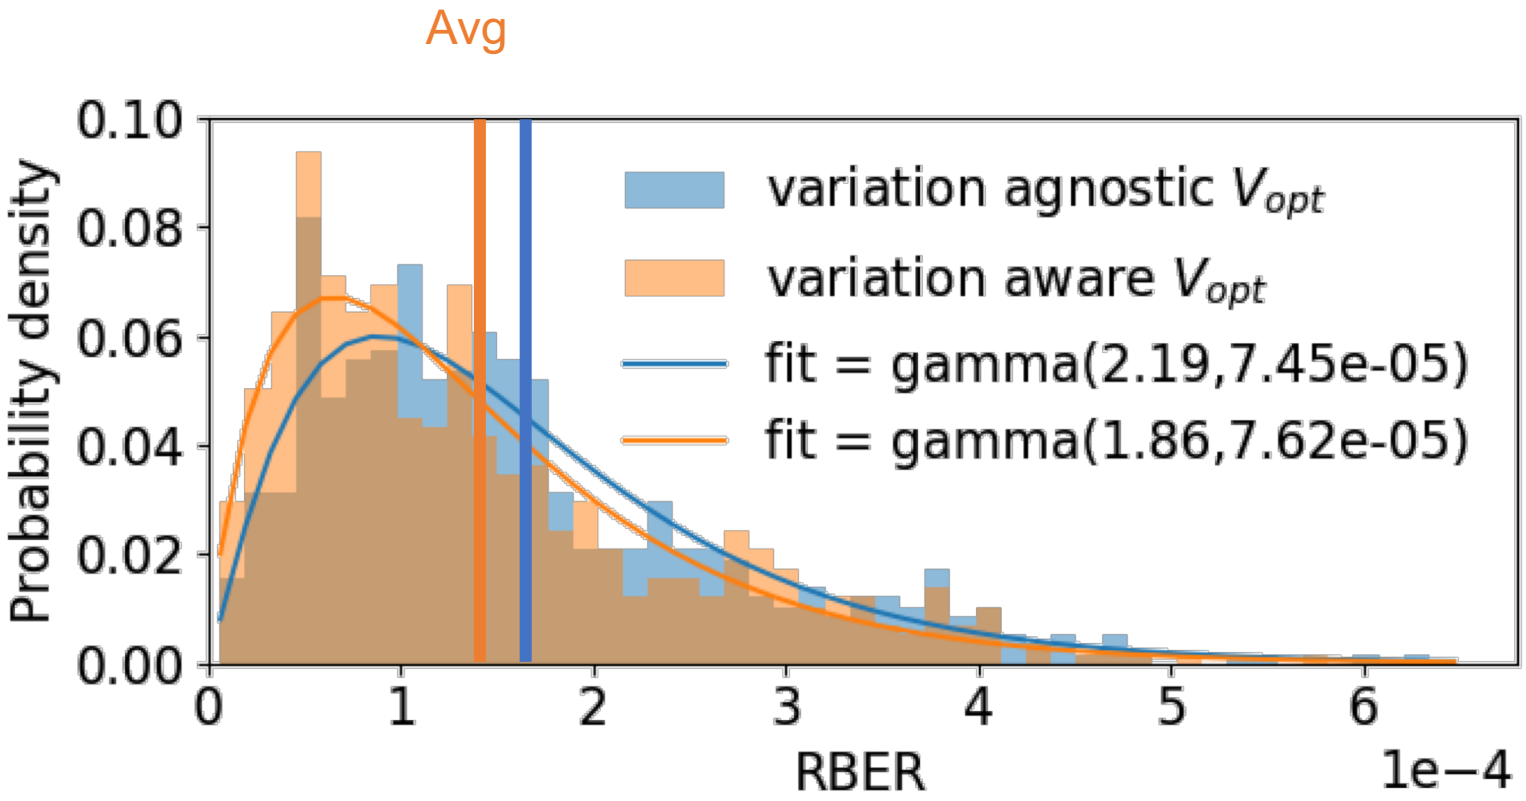

Threshold Voltage Distribution

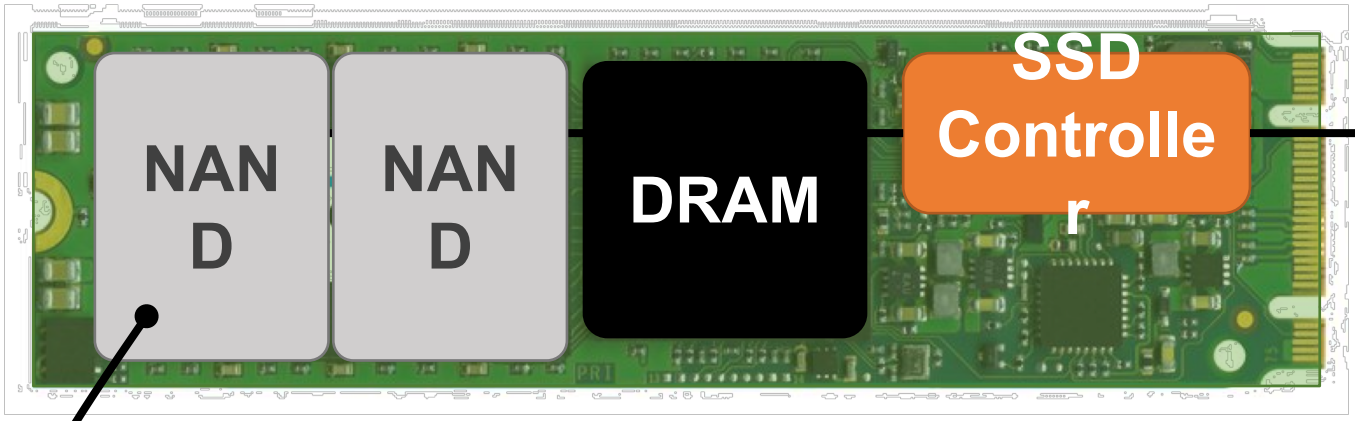

NAND Flash Memory
(Flash Chip)

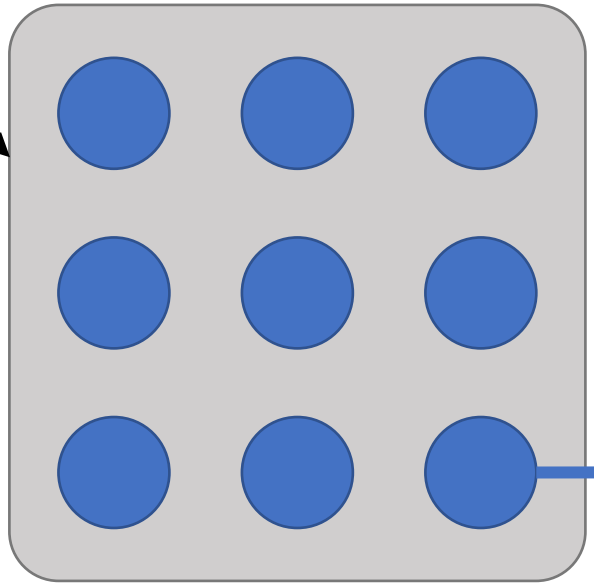

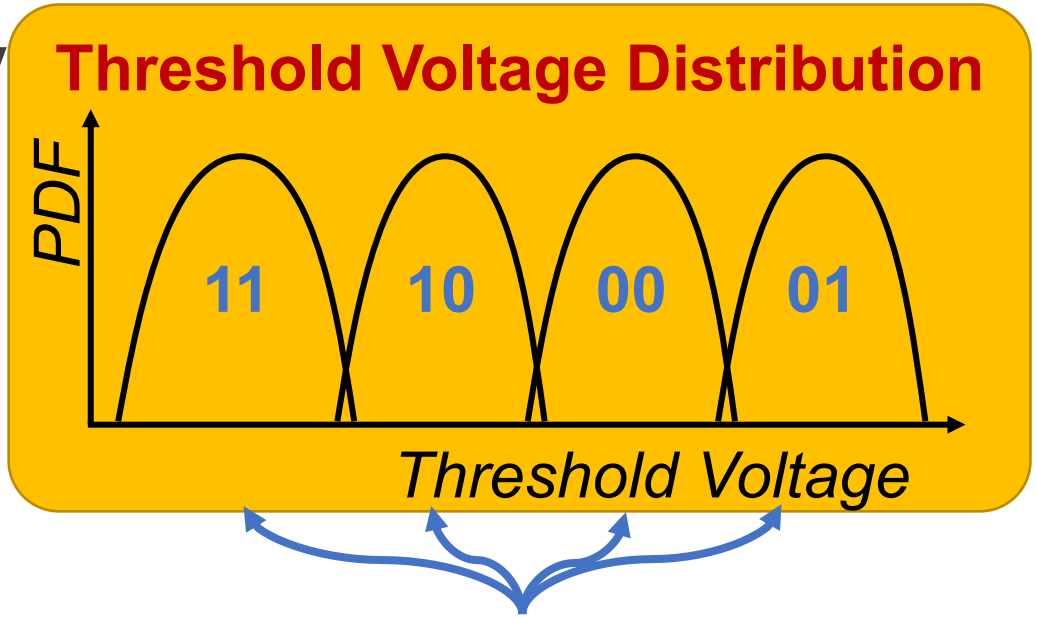

Flash Cell

Reading From A Flash Cell

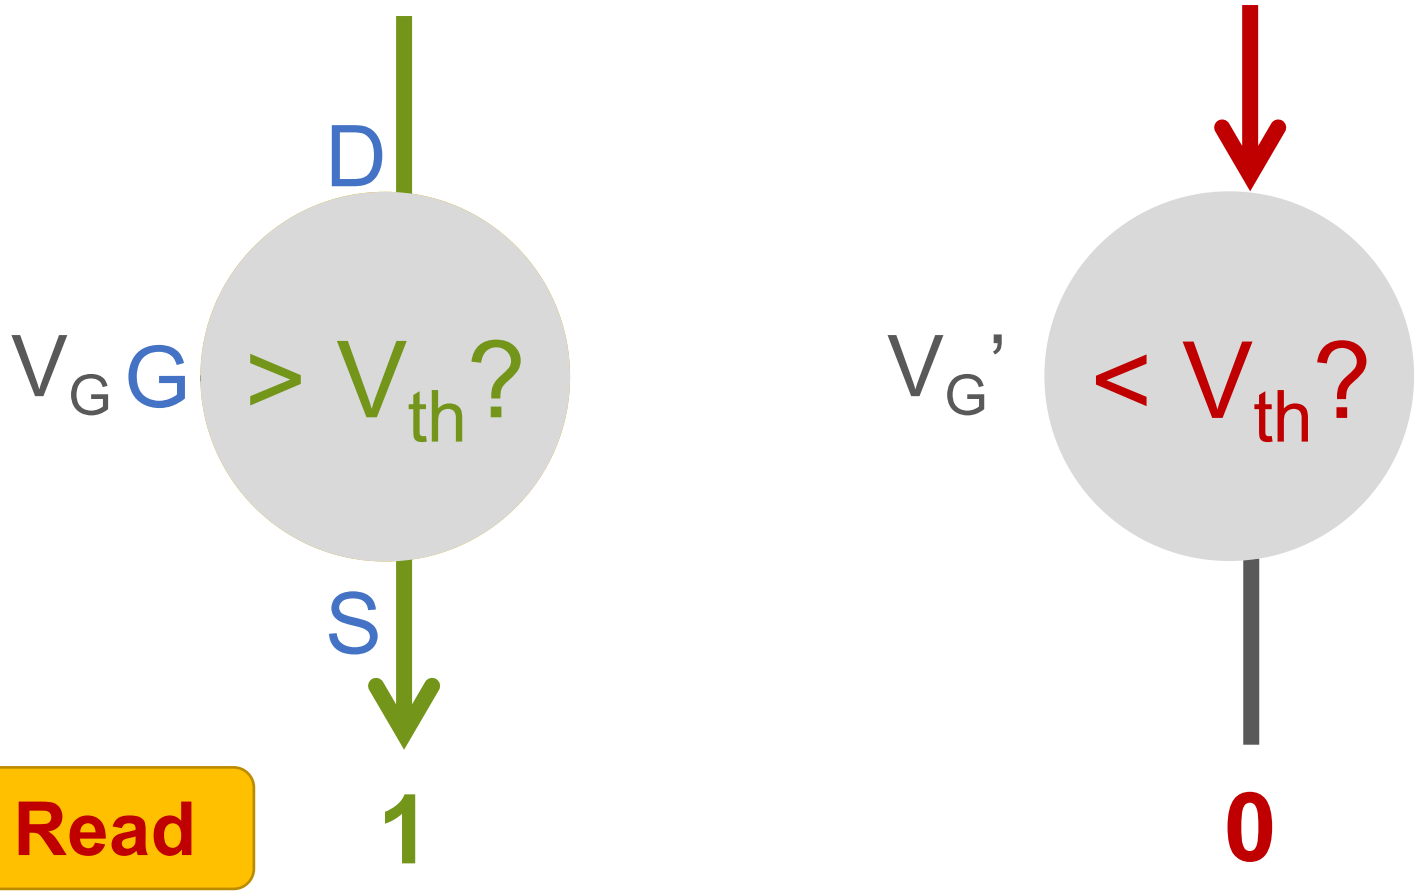

Reading From A Flash Cell

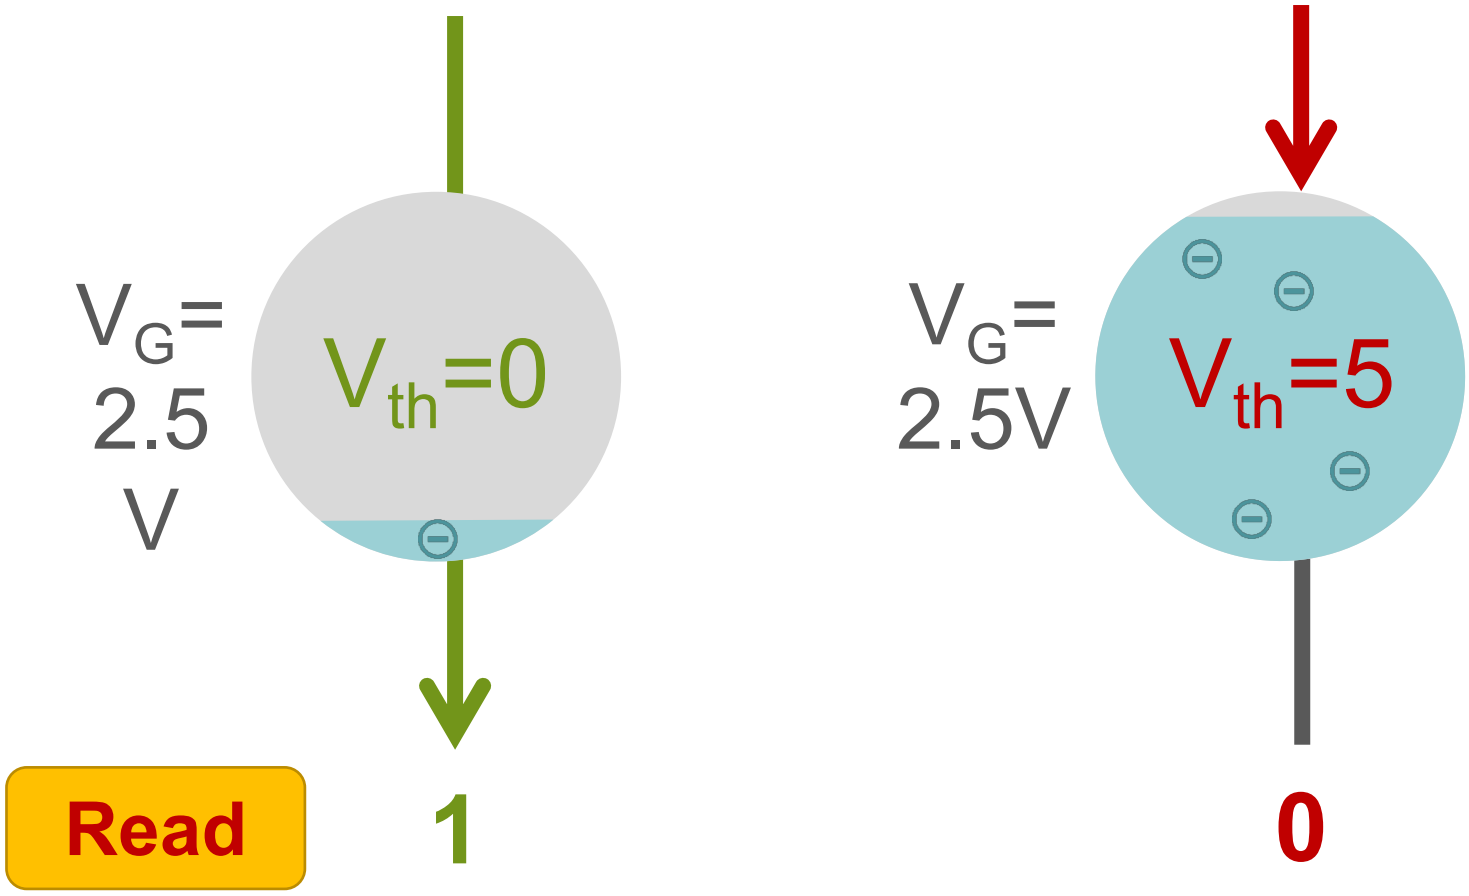

Writing To A Flash Cell

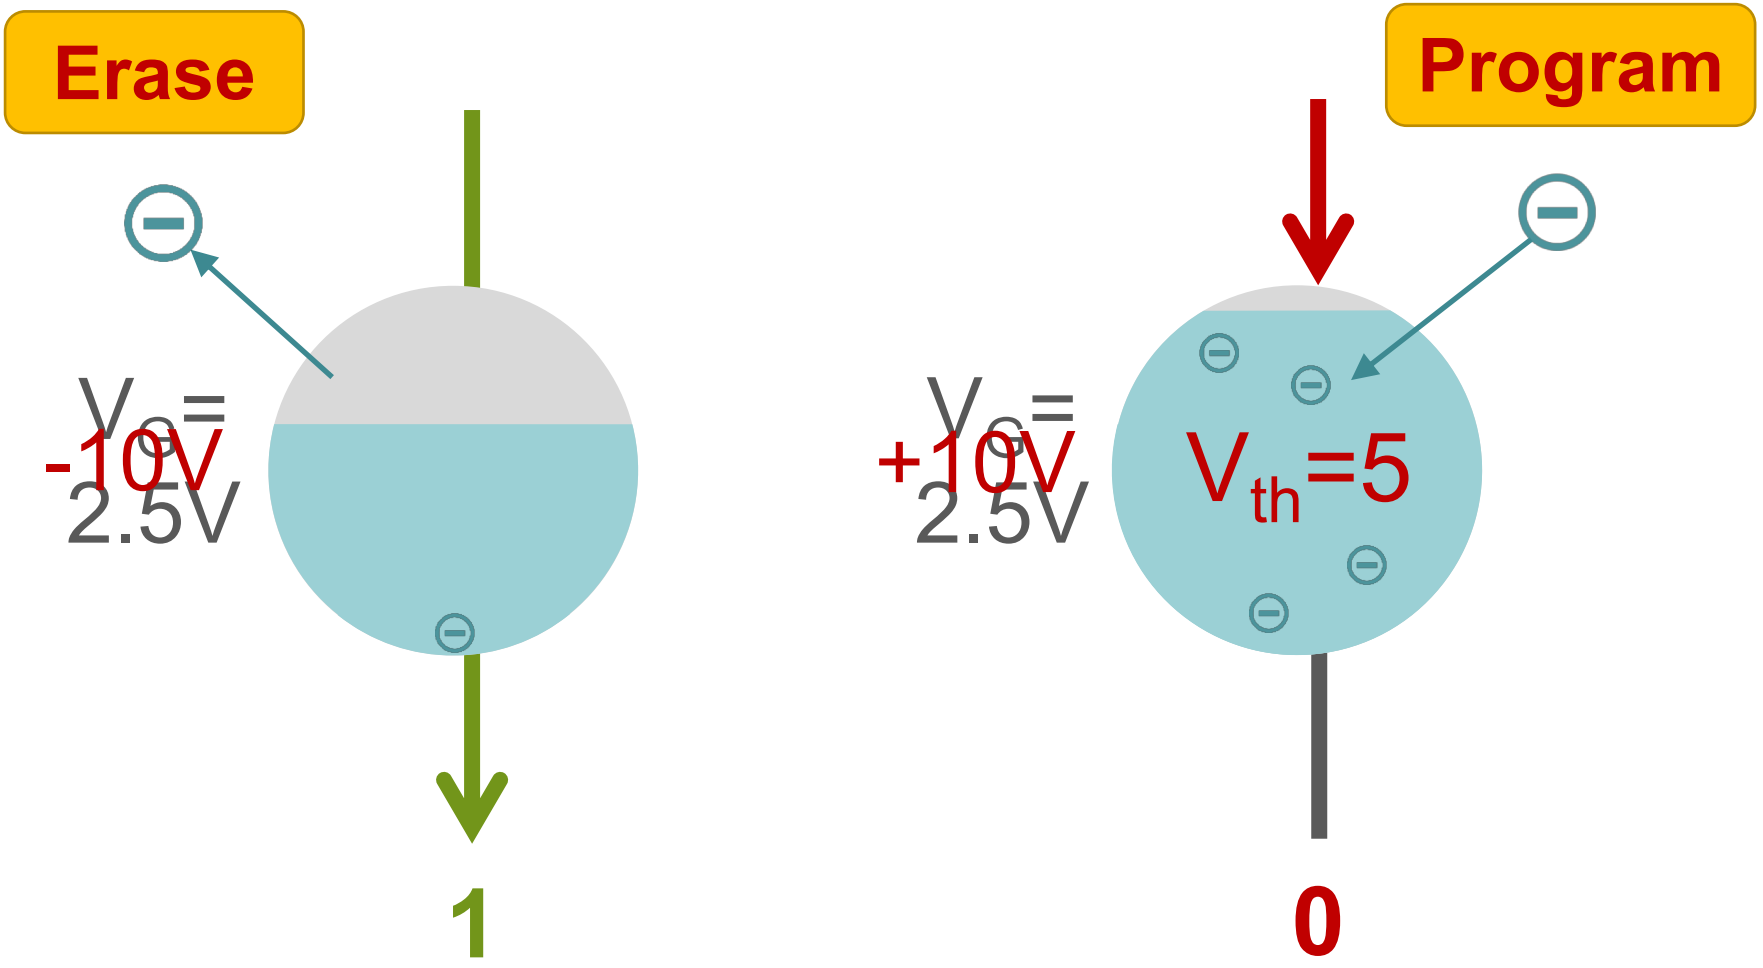

Threshold Voltage Distribution

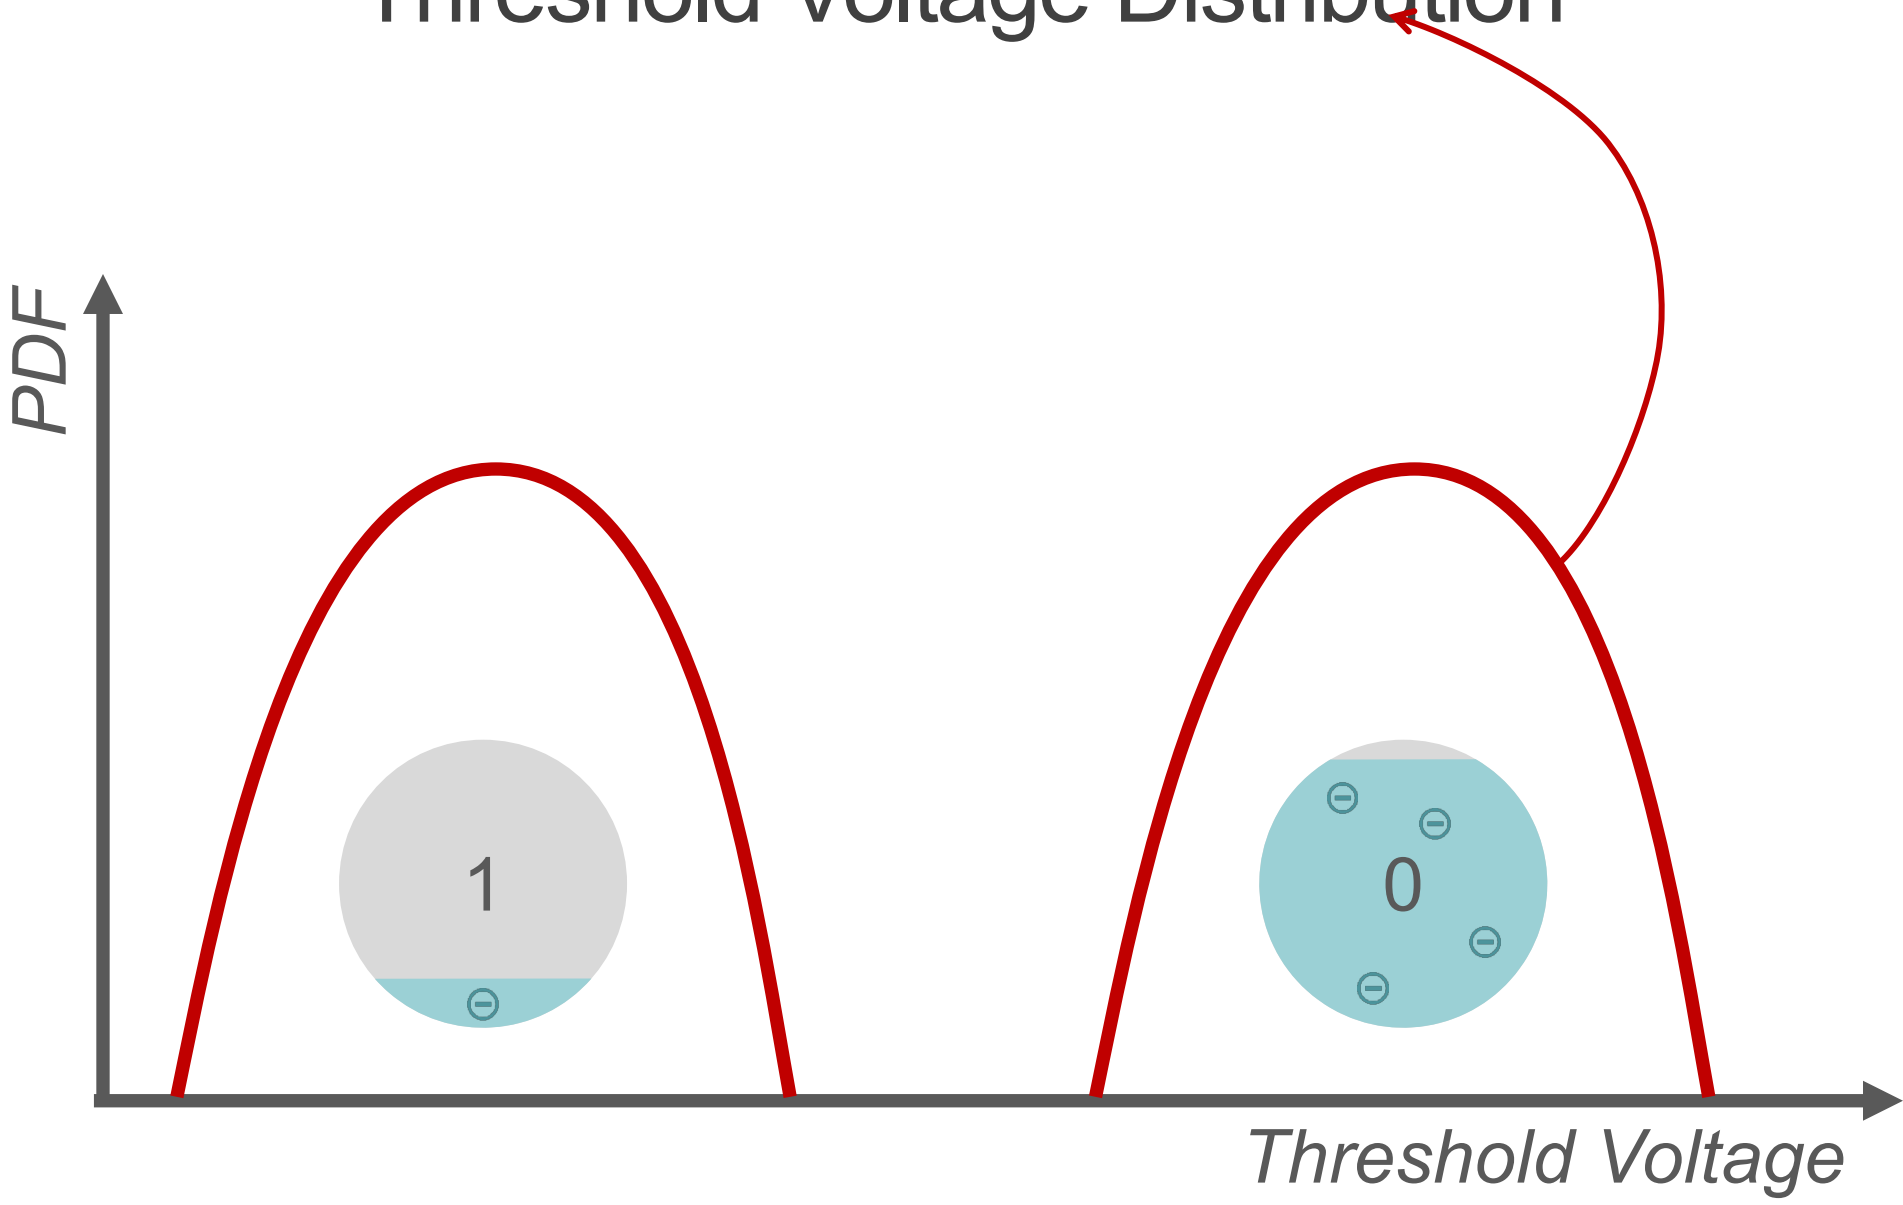

Read Reference Voltage (V_{ref})

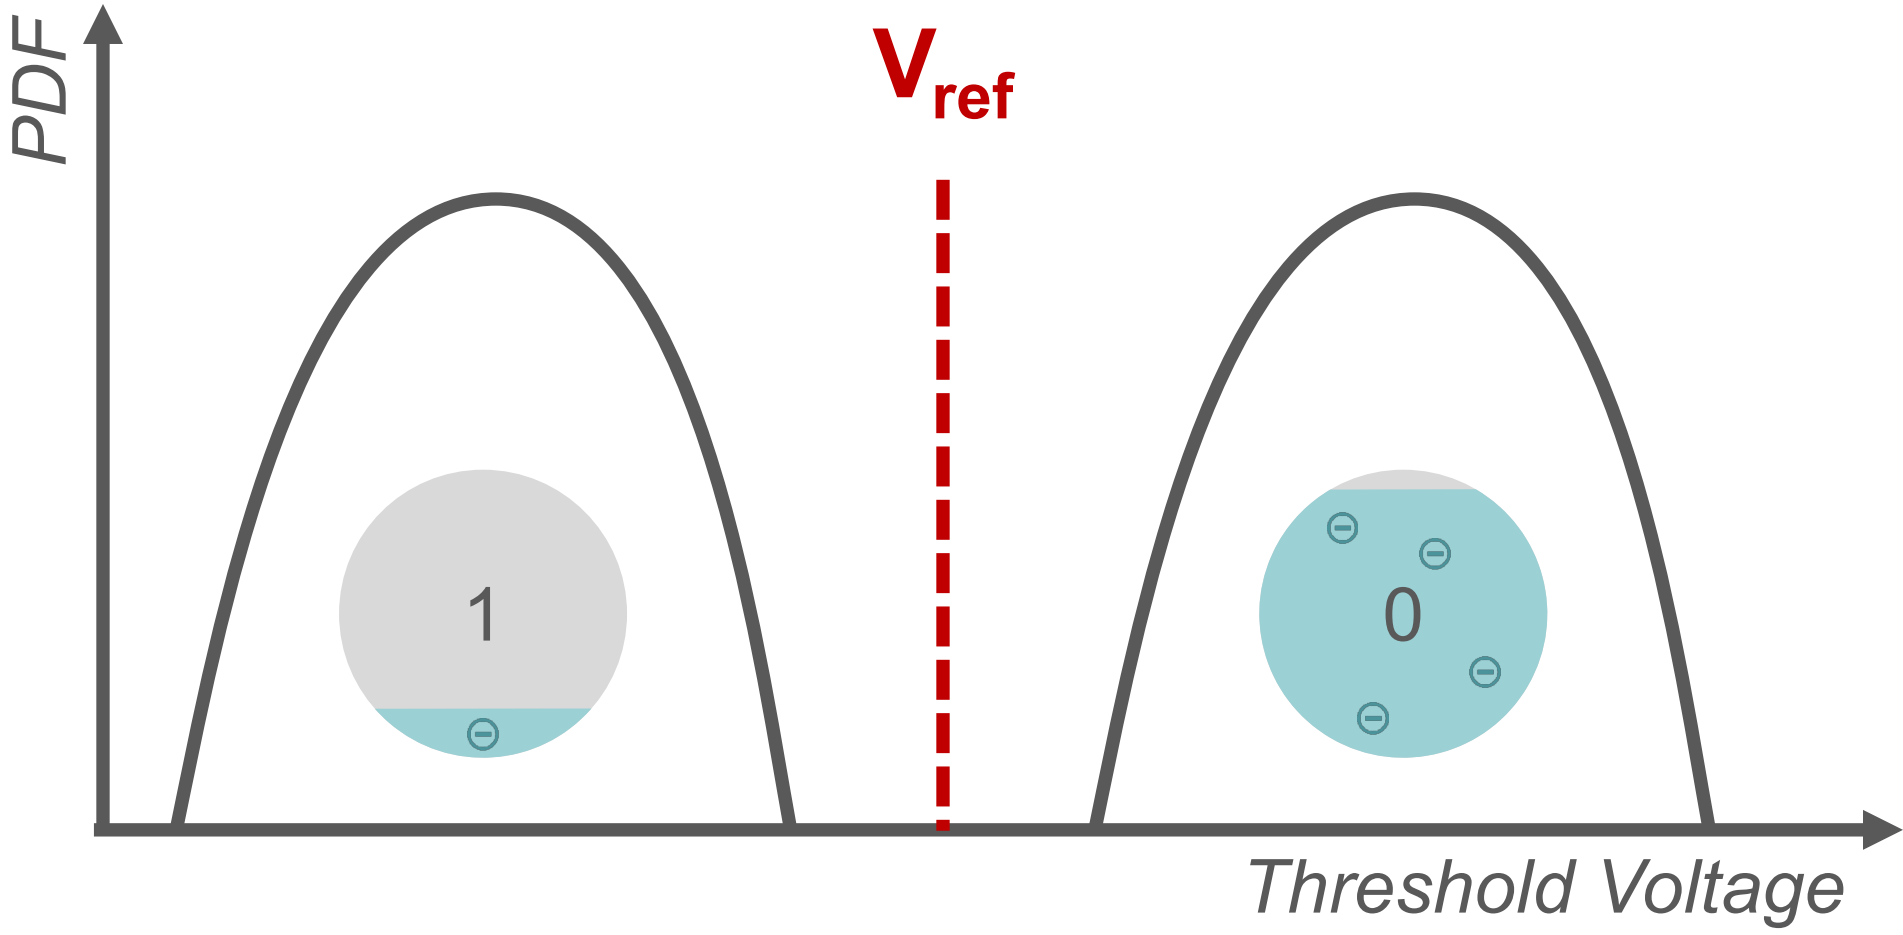

Multi-Level Cell (MLC)

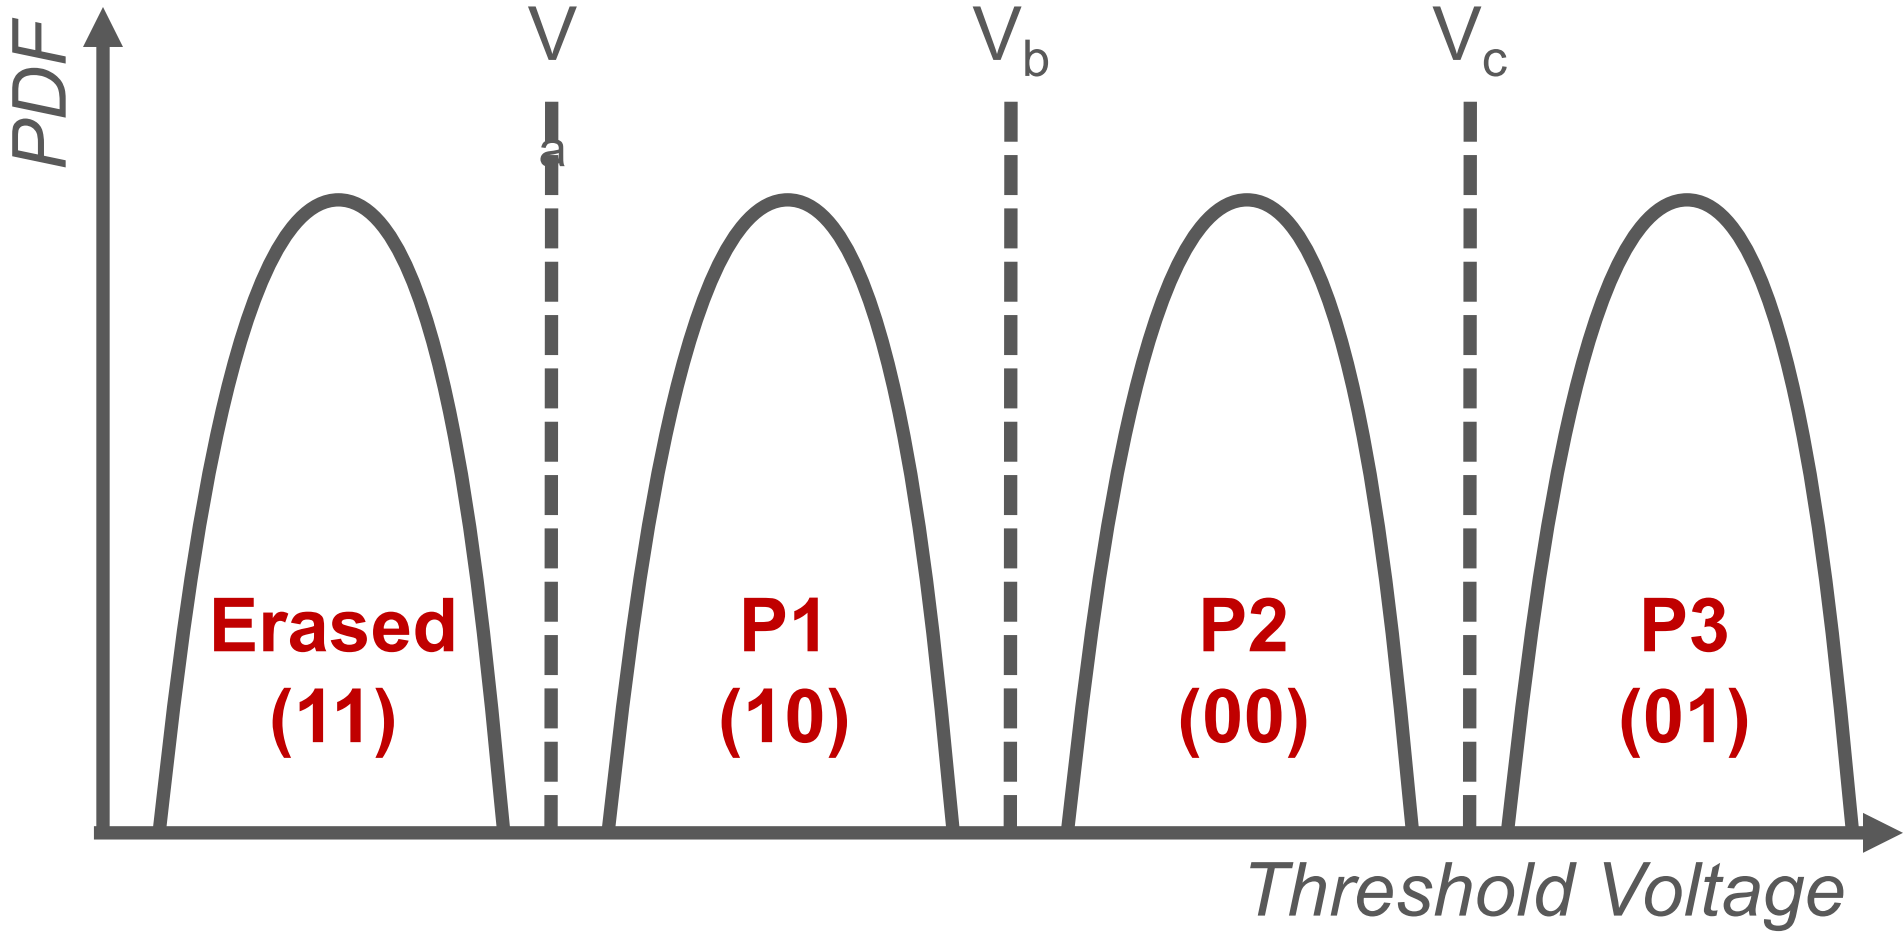

Background:

- Flash Reliability Background
- 3D NAND vs. Planar NAND

Threshold Voltage Distribution

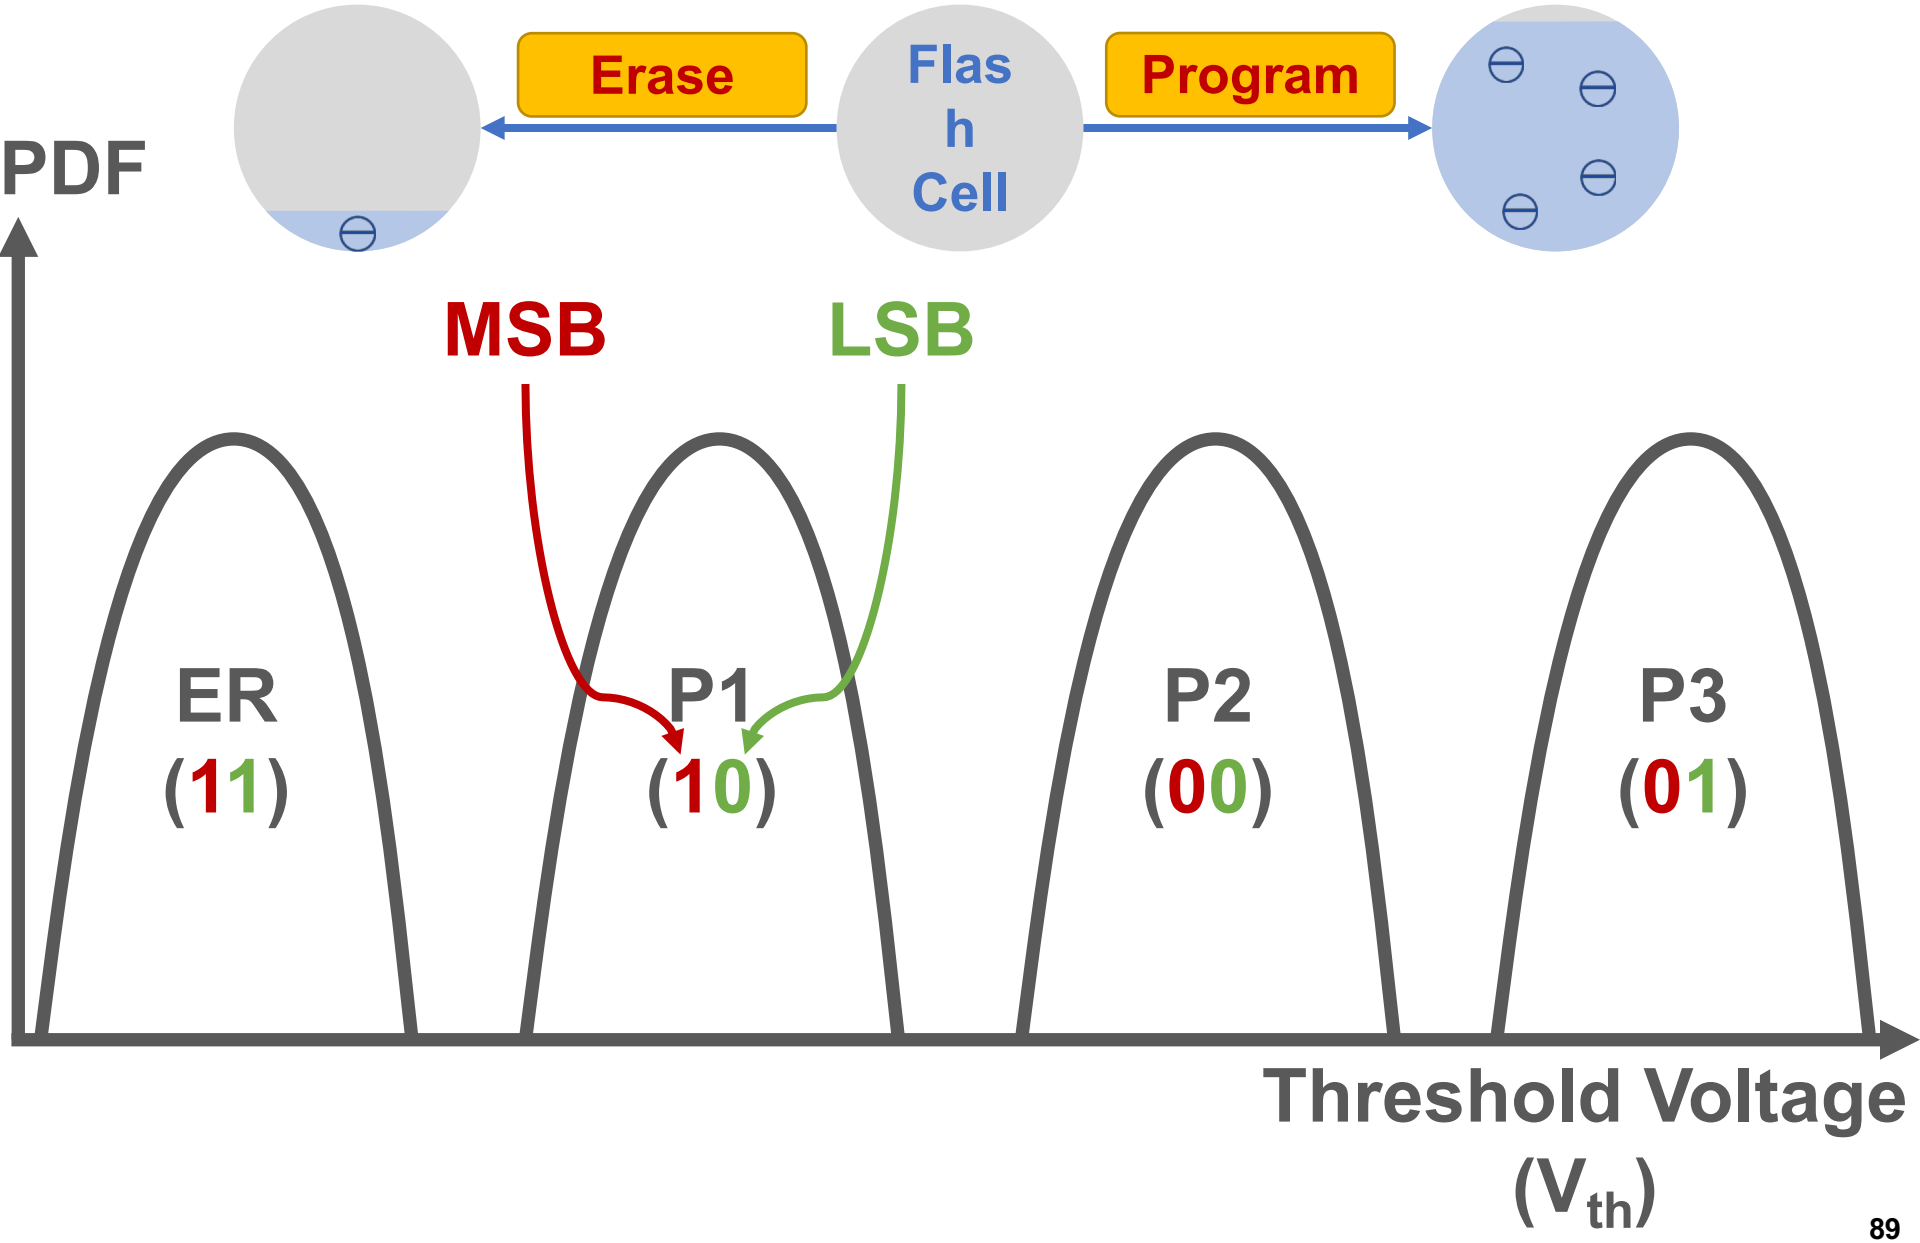

Flash Block Organization

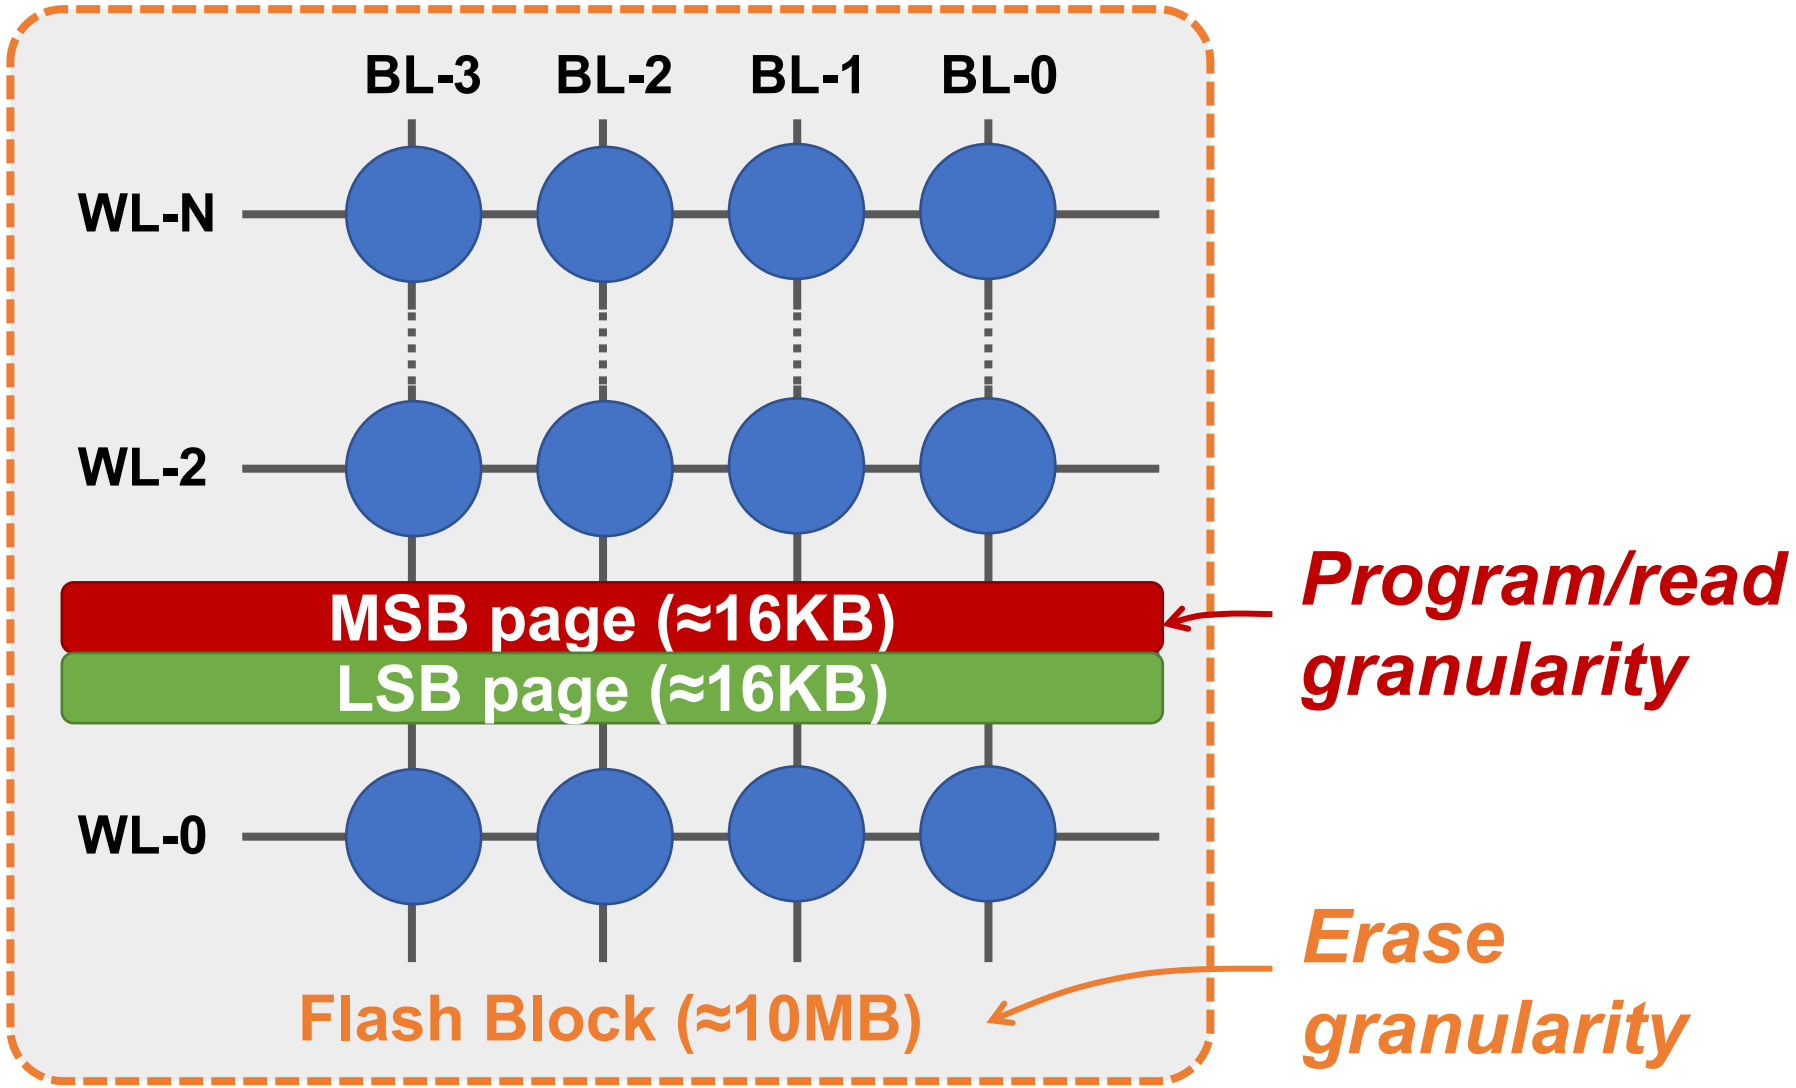

Read Reference Voltage

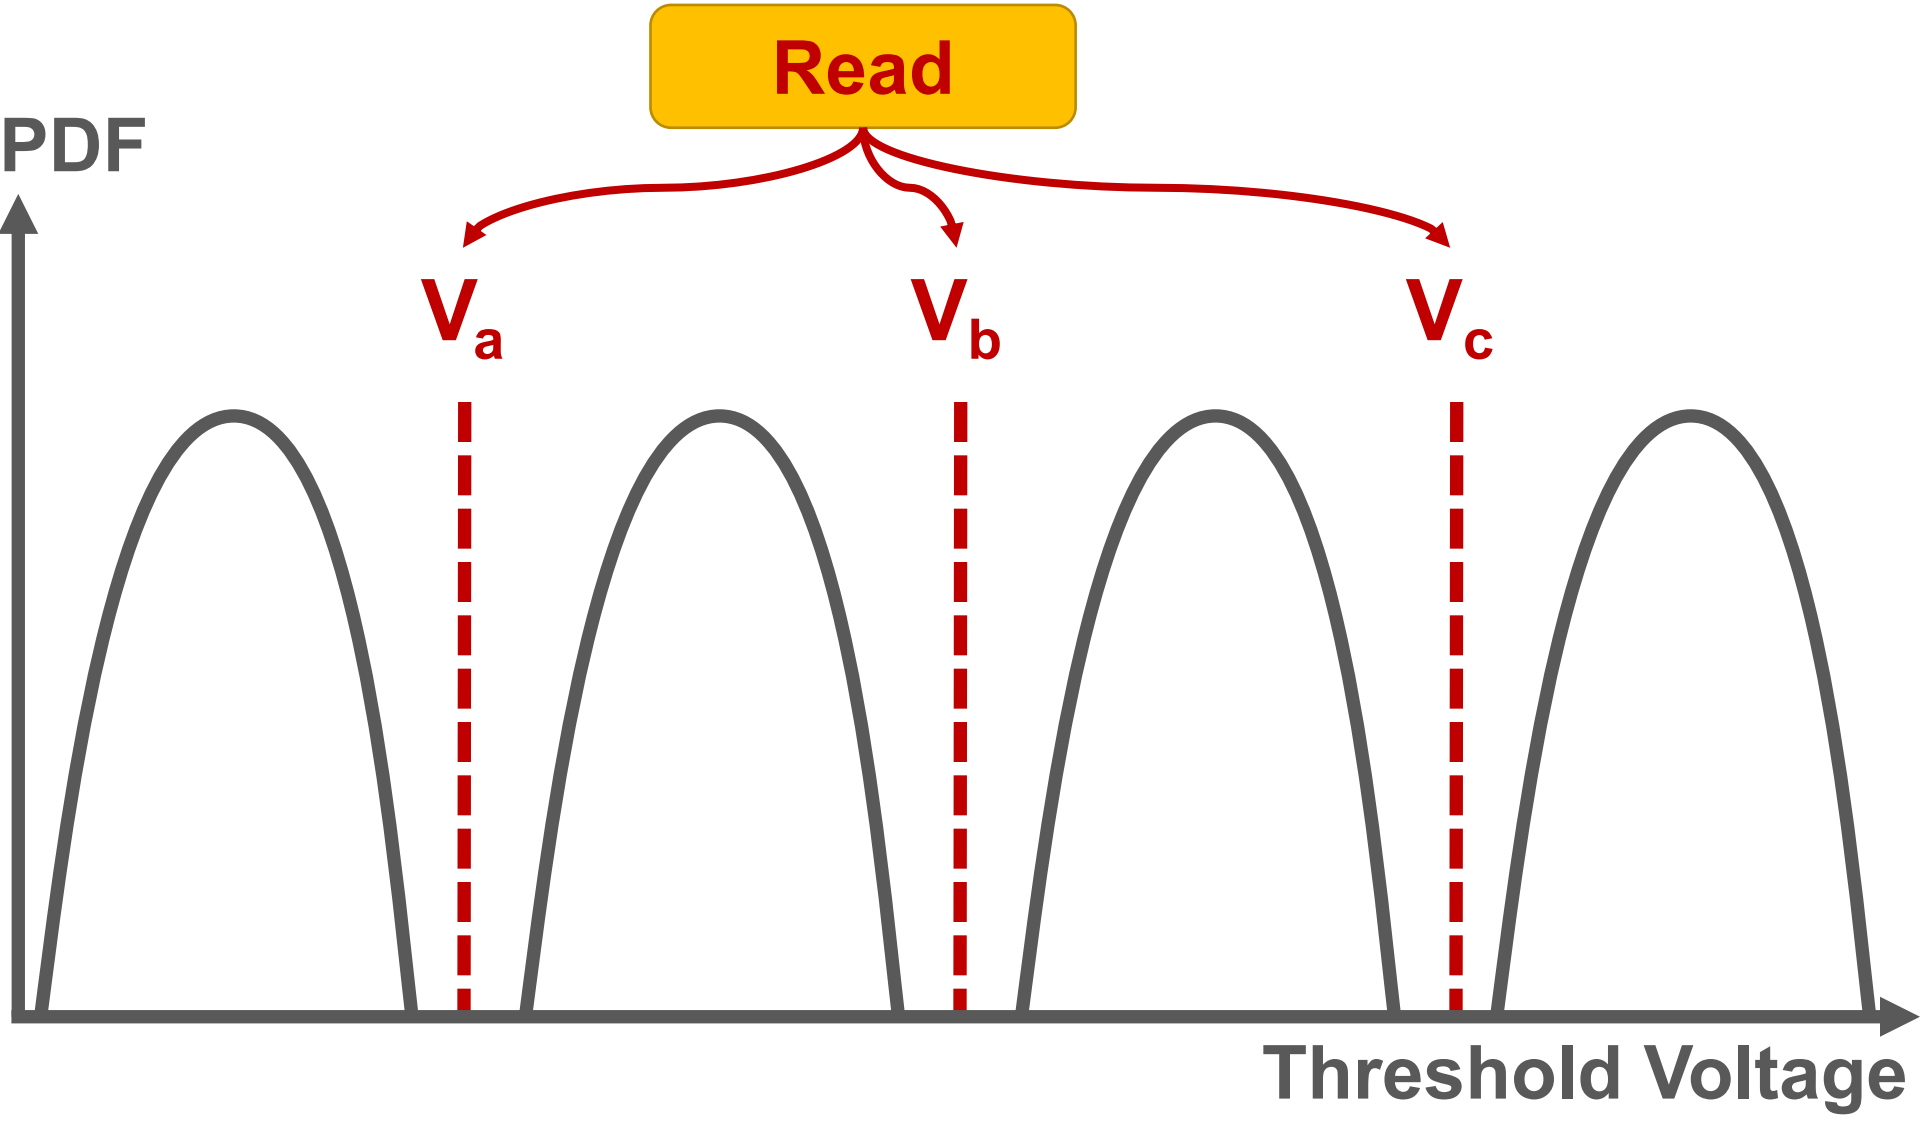

Common Types of Flash Errors

- P/E cycling [Yu+ DATE'13]
  - Wear out
- Program interference [Yu+ ICCD'13]
  - Coupling
- Program [Yu+ HPCA'17]
  - Two-step programming
- Read disturb [Yu+ DSN'15]
  - Weak programming
- Retention [Yu+ HPCA'15]
  - Charge leakage

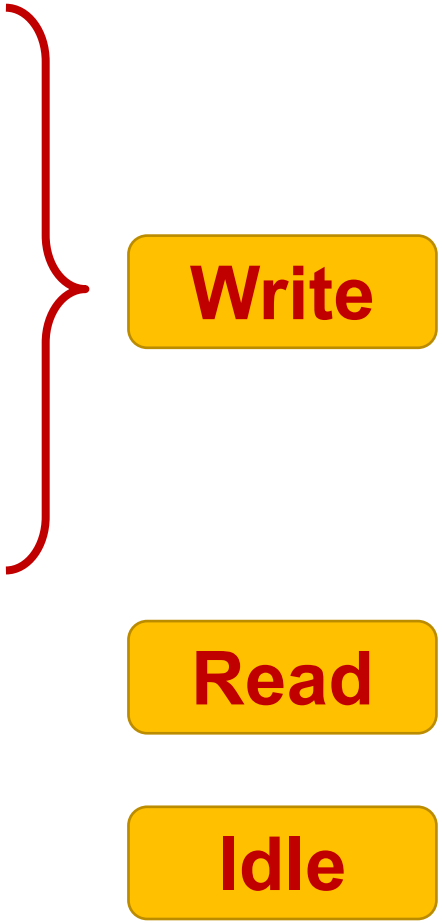

Raw Bit Errors

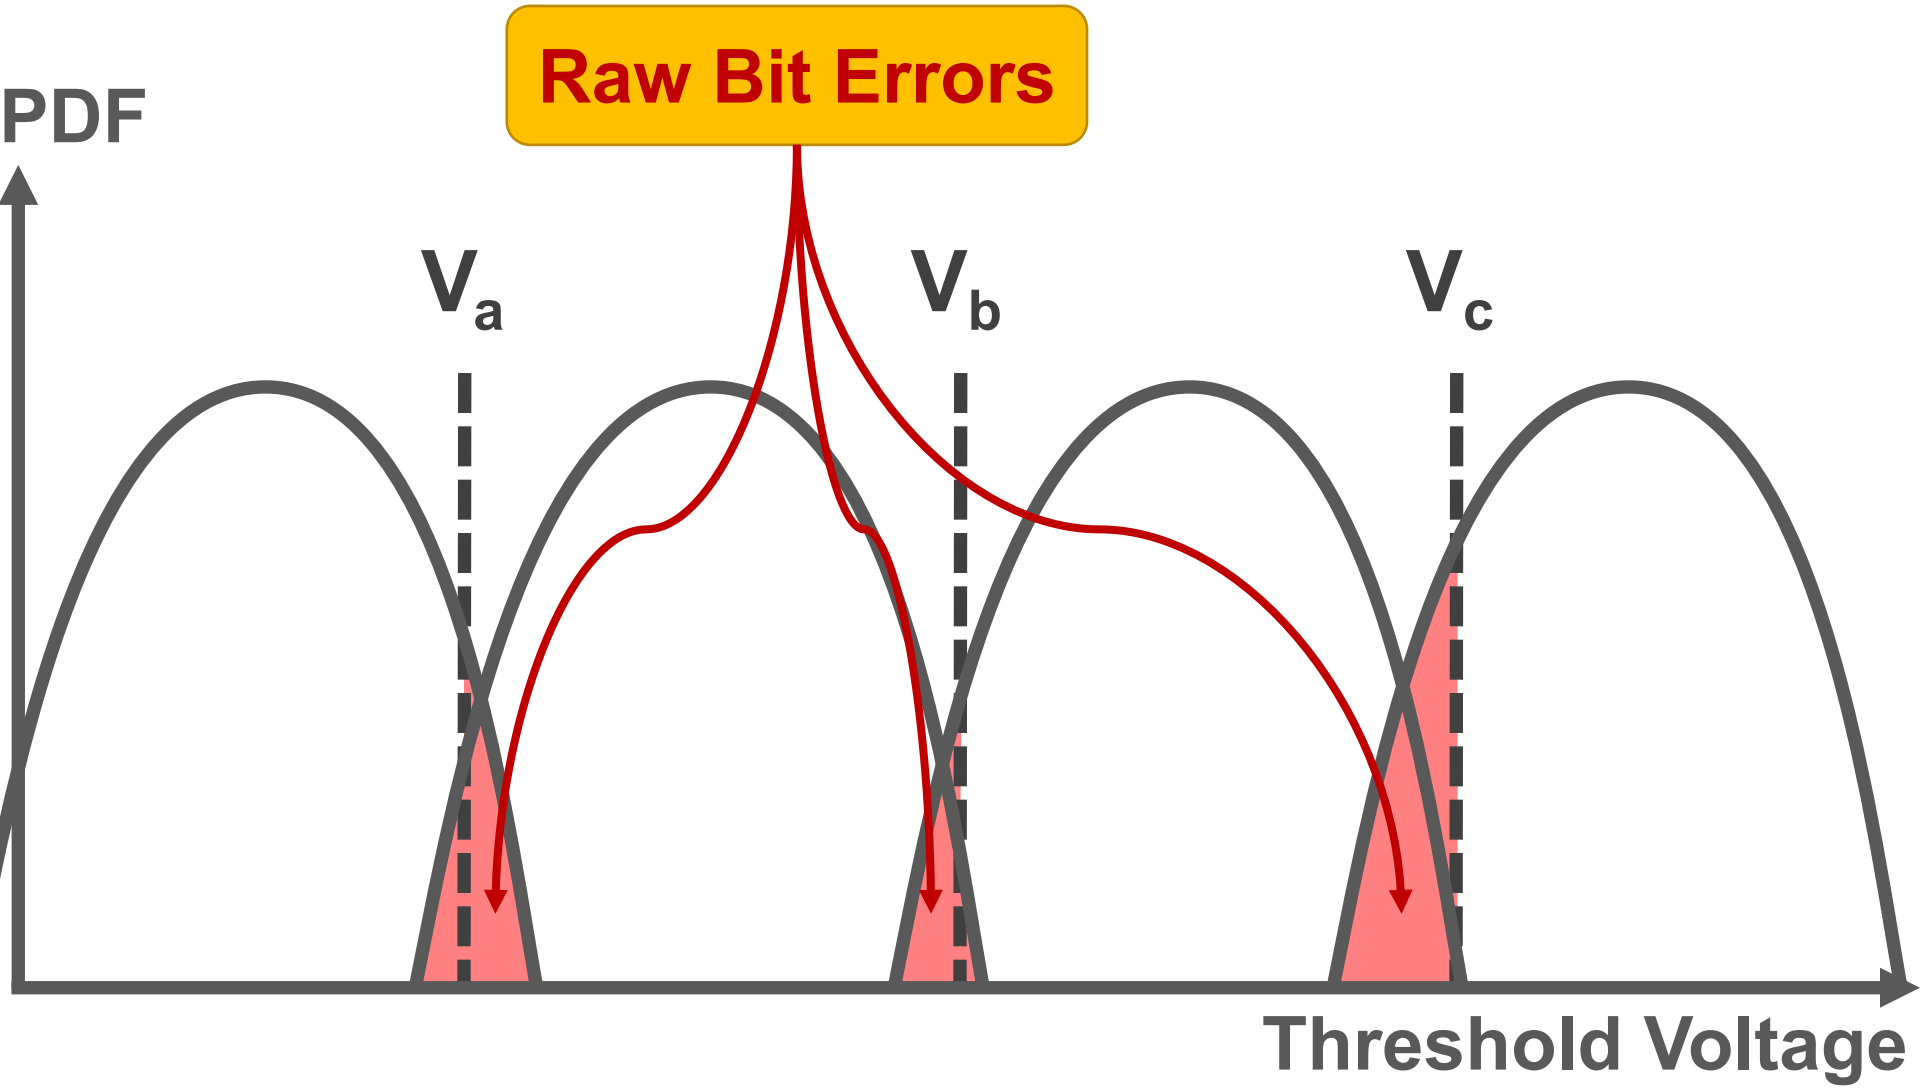

Flash Reliability Summary

- Flash operations
- → Various types of noise
  - → Threshold voltage distribution shift
    - → Raw bit errors

- Scaling
  - Smaller cells → bigger shifts
  - Smaller distance between cells → bigger noise

- Solution?

3D NAND Flash Memory Scaling

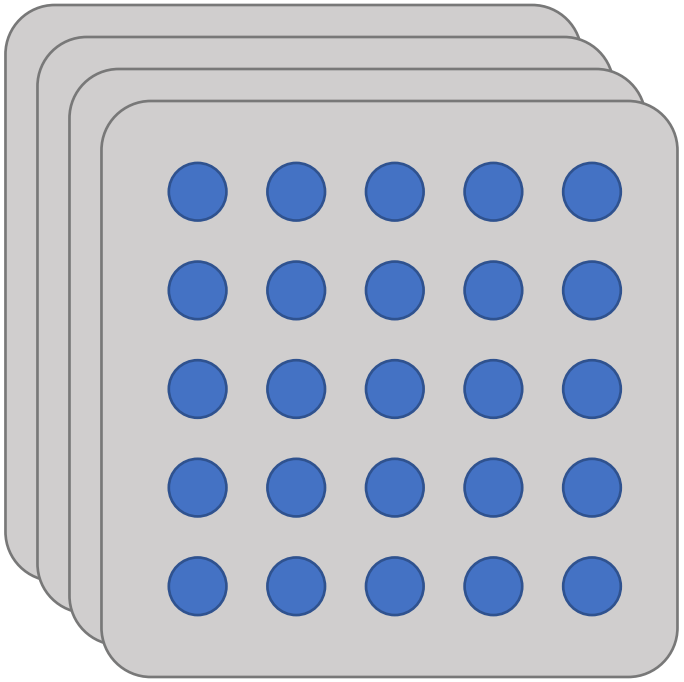

3D NAND vs. Planar NAND Differences

- Flash cell design
- Flash chip organization
- Larger manufacturing process

These differences fundamentally affect various types of flash errors!

Flash Cell Design

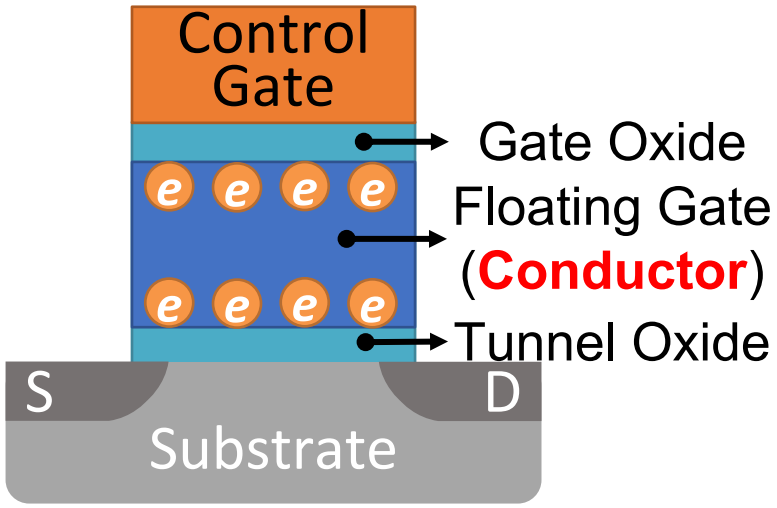

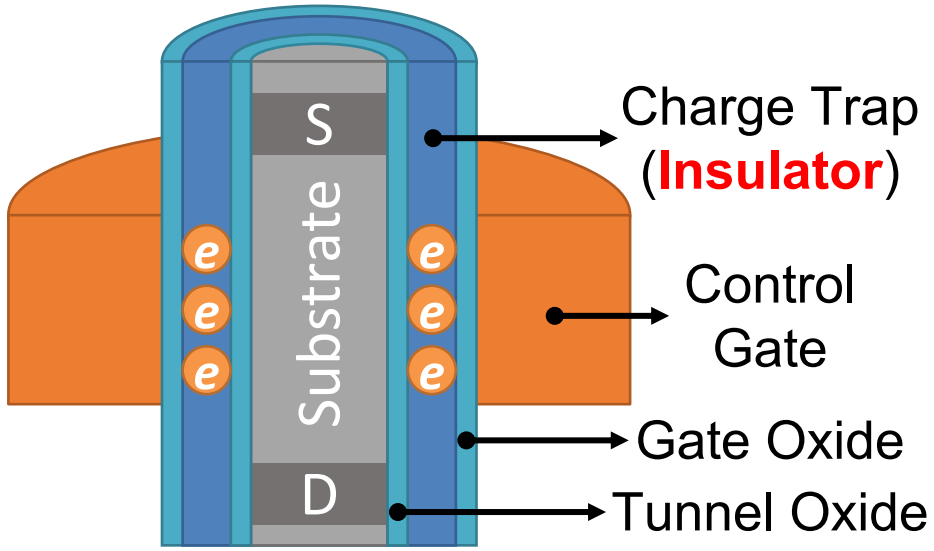

Floating-Gate Cell

3D Charge-Trap Cell

Charges stored in insulator, thinner tunnel oxide
→ Faster data retention

Flash Chip Organization

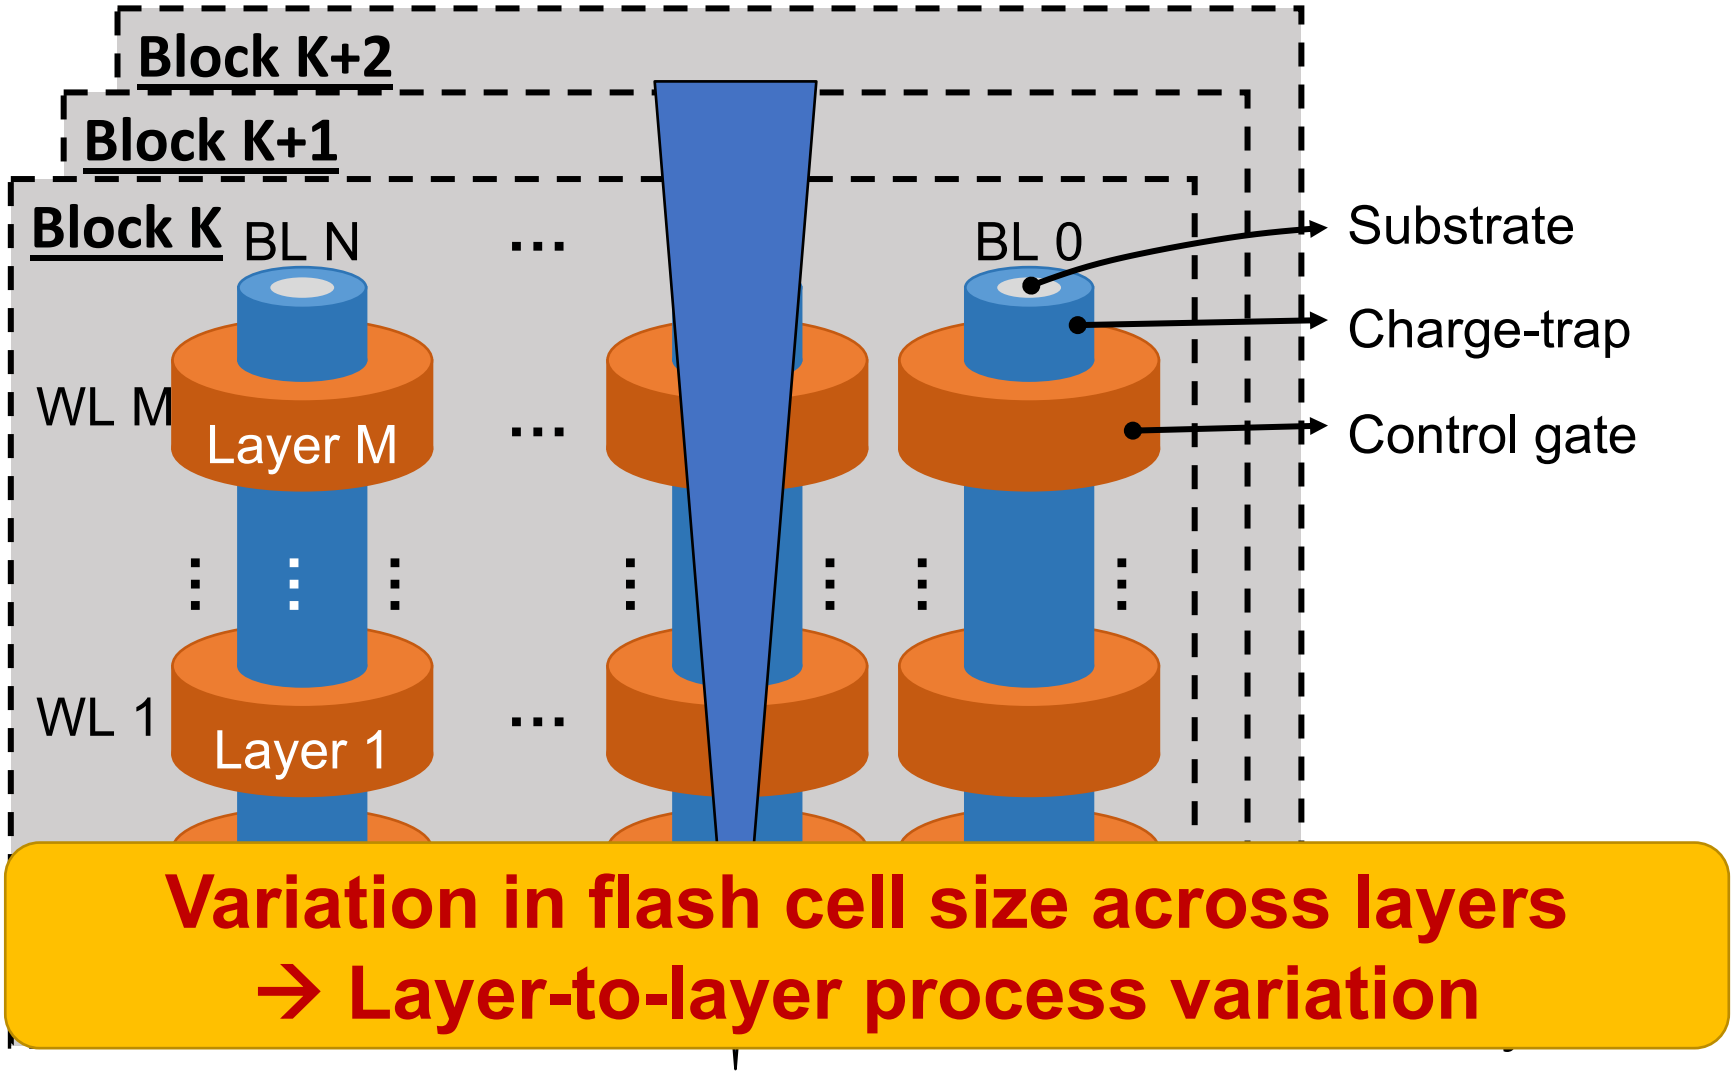

Variation in flash cell size across layers
→ Layer-to-layer process variation

Summary of Differences

- **Flash cell design**
  - Faster data retention
- **Flash chip organization**
  - Layer-to-layer process variation
- **Larger manufacturing process**
  - More resistant to other types of errors

Threshold Voltage Distribution Characterization Methodology

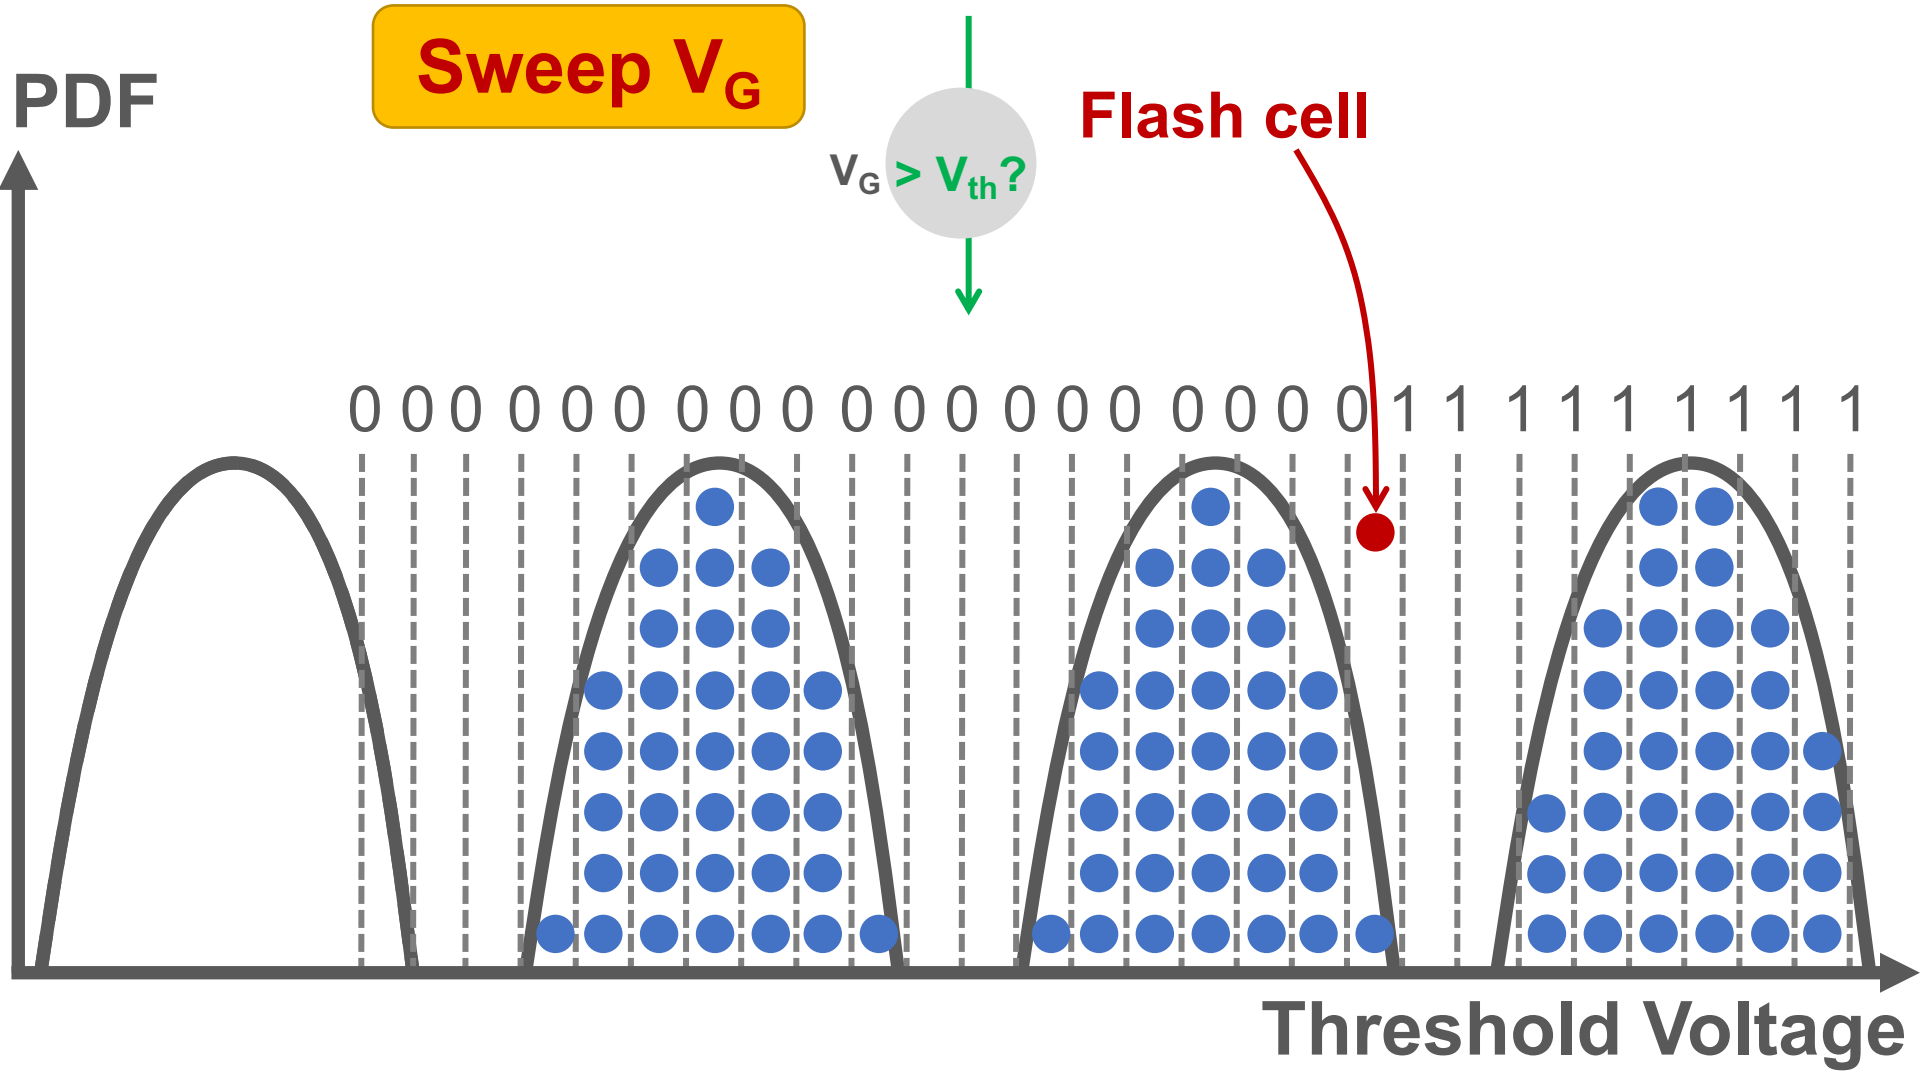

Retention Errors

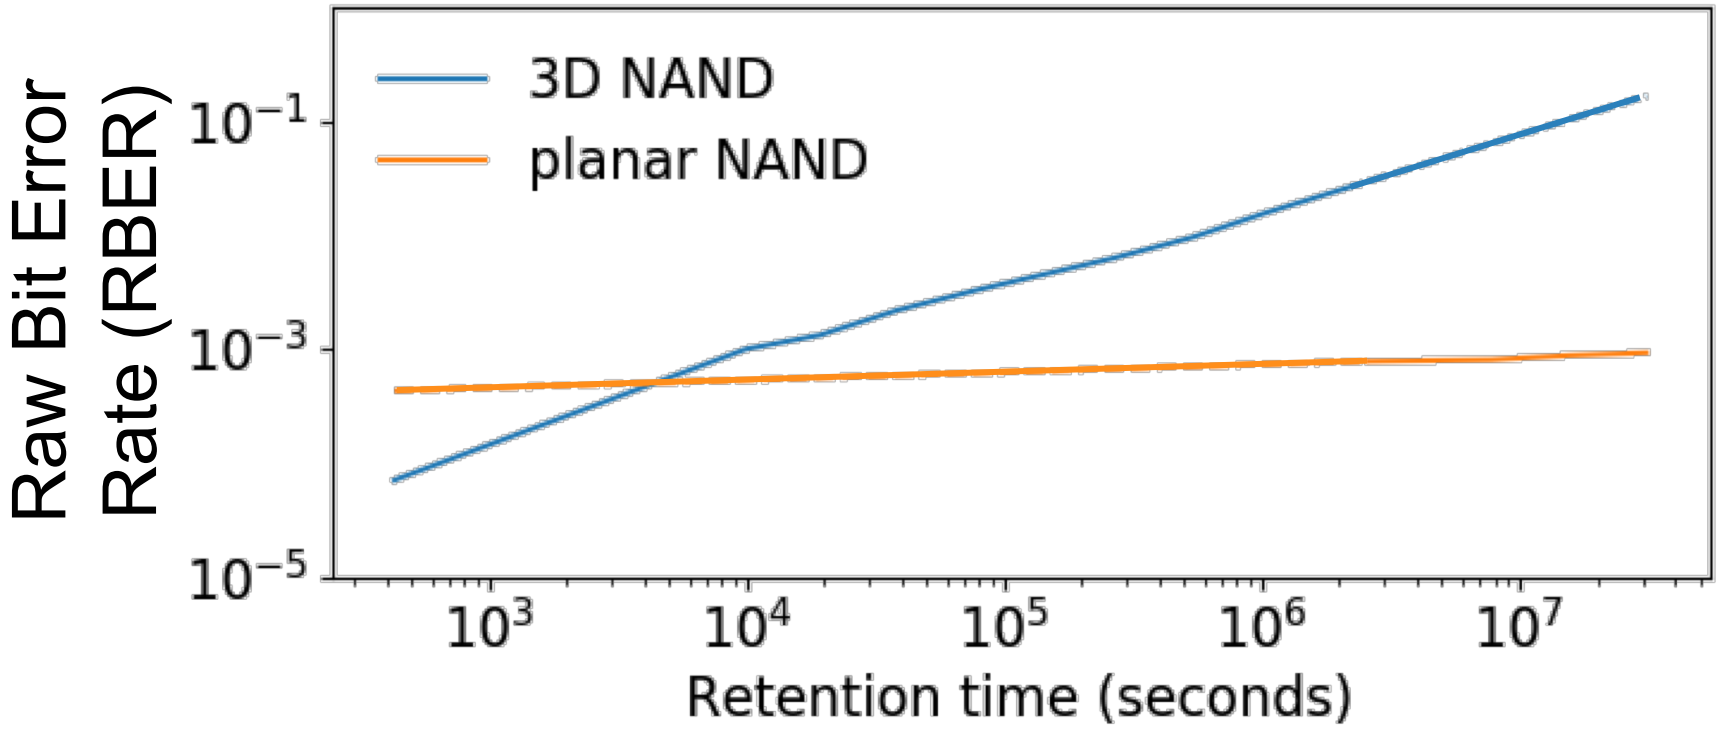

Retention errors increase faster in 3D NAND

Importance of Prediction Accuracy

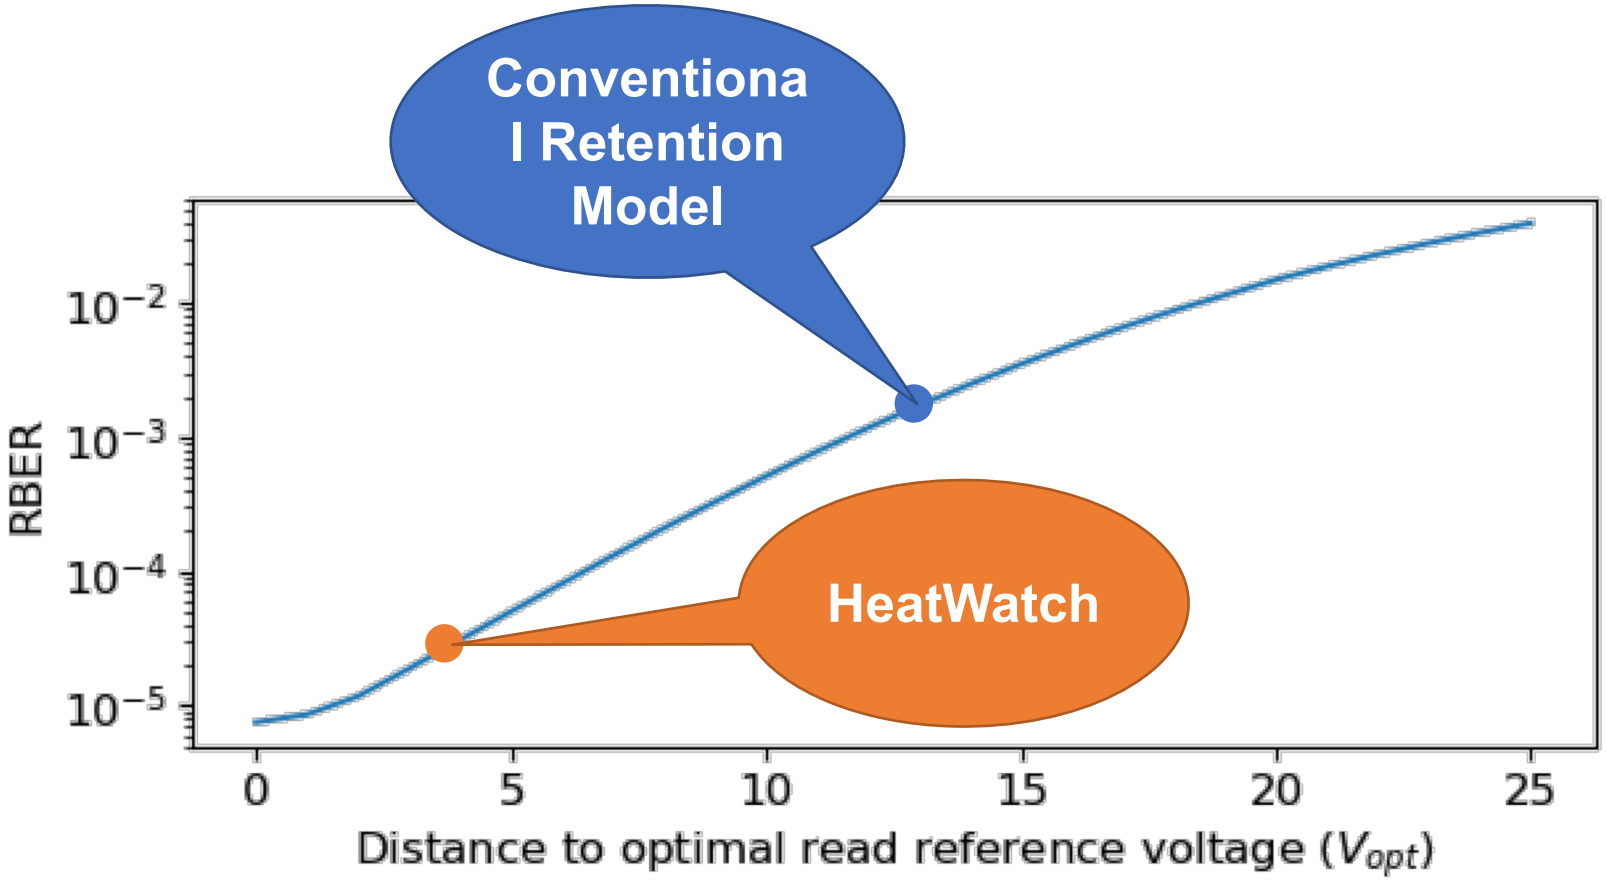

Raw Bit Error Rate Variation Within A Block

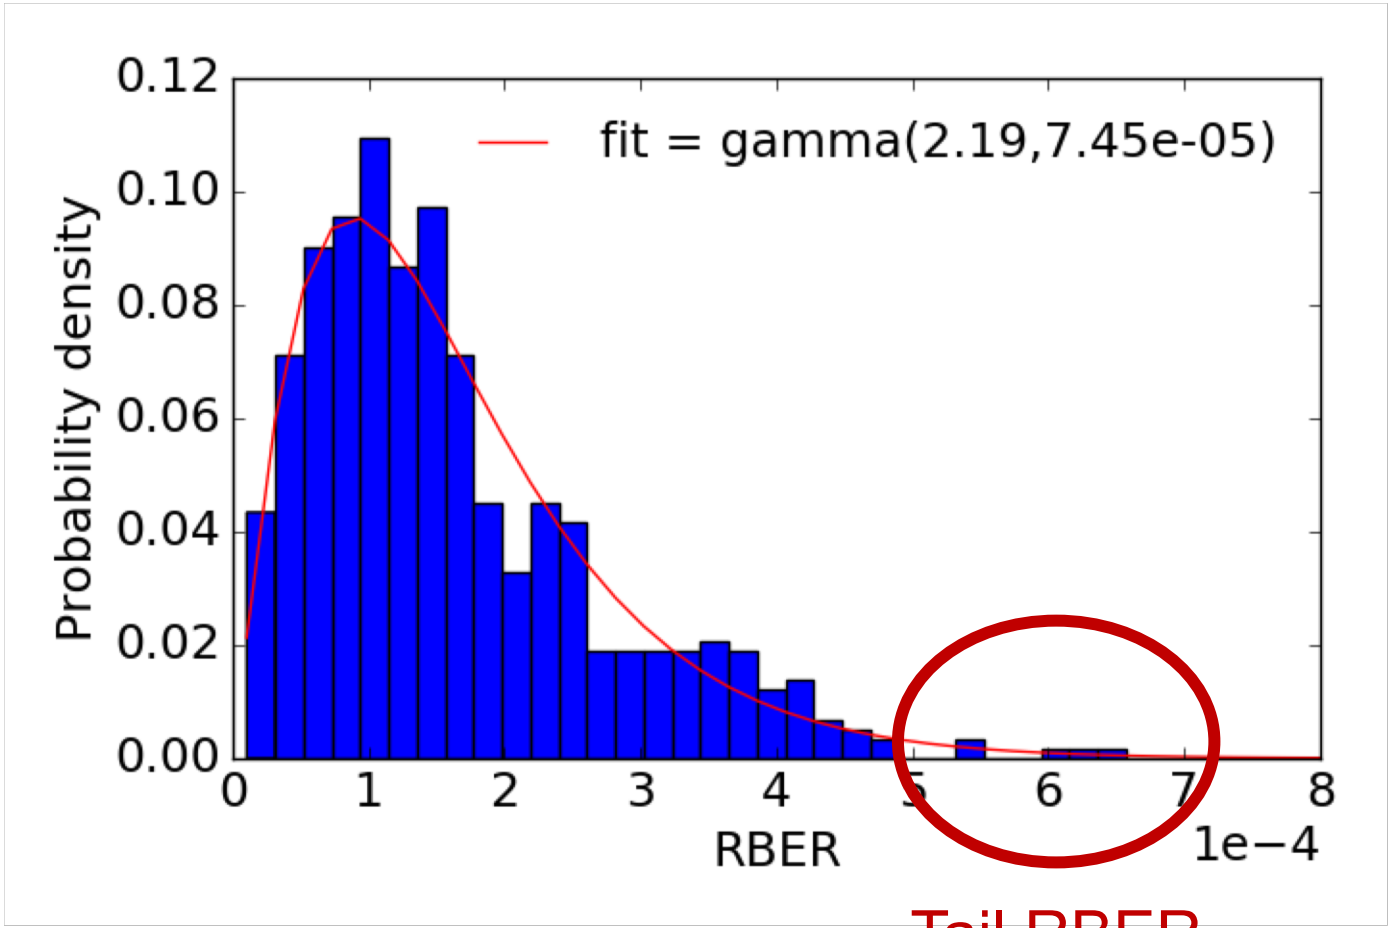

Tail RBER

MSB pages @ middle layers

Characterization Summary

- **Retention errors**
  - Increase much faster
  - Dominate SSD errors
- **Layer-to-layer process variation**
  - Error rate much higher than average in the MSB pages on the middle layers

HeatWatch Summary

- Dwell time and temperature affect retention
- Conventional retention model is insufficient
- HeatWatch
  - Uses a new unified retention model
    - *Unifies: PEC, t_{ret}, T_{ret}, t_{dwell}, T_{dwell}, T_{prog}*
  - Efficiently computes effective retention/dwell time
    - *Combines: t_{ret} & T_{ret}, t_{dwell} & T_{dwell}*
- Results
  - Improves flash lifetime by 3.85 times
  - < 1.6 MB memory for 1TB SSD

Rising Popularity of NAND Flash Memory

Data Centers and Servers

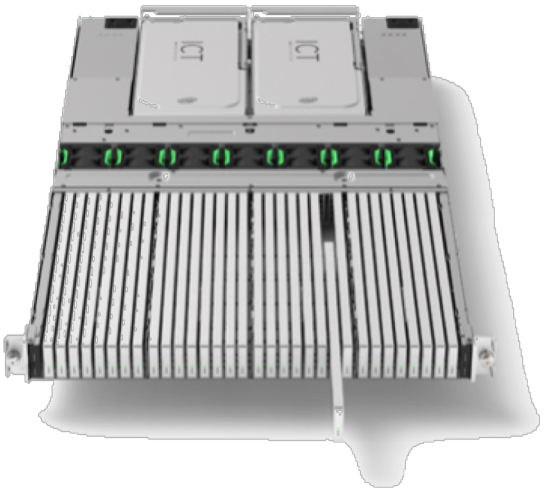

Personal Computers

Mobile Devices

Storage Technology Drivers - 2018

- Internet (User data, Cloud storage)
- Camera (4K, VR, Drones, Light field)
- AI (Machine learning, Self-driving)
- IoT (Sensor data)
- Bioinformatics (DNA sequencing, Health monitoring)

Primary Storage Demands

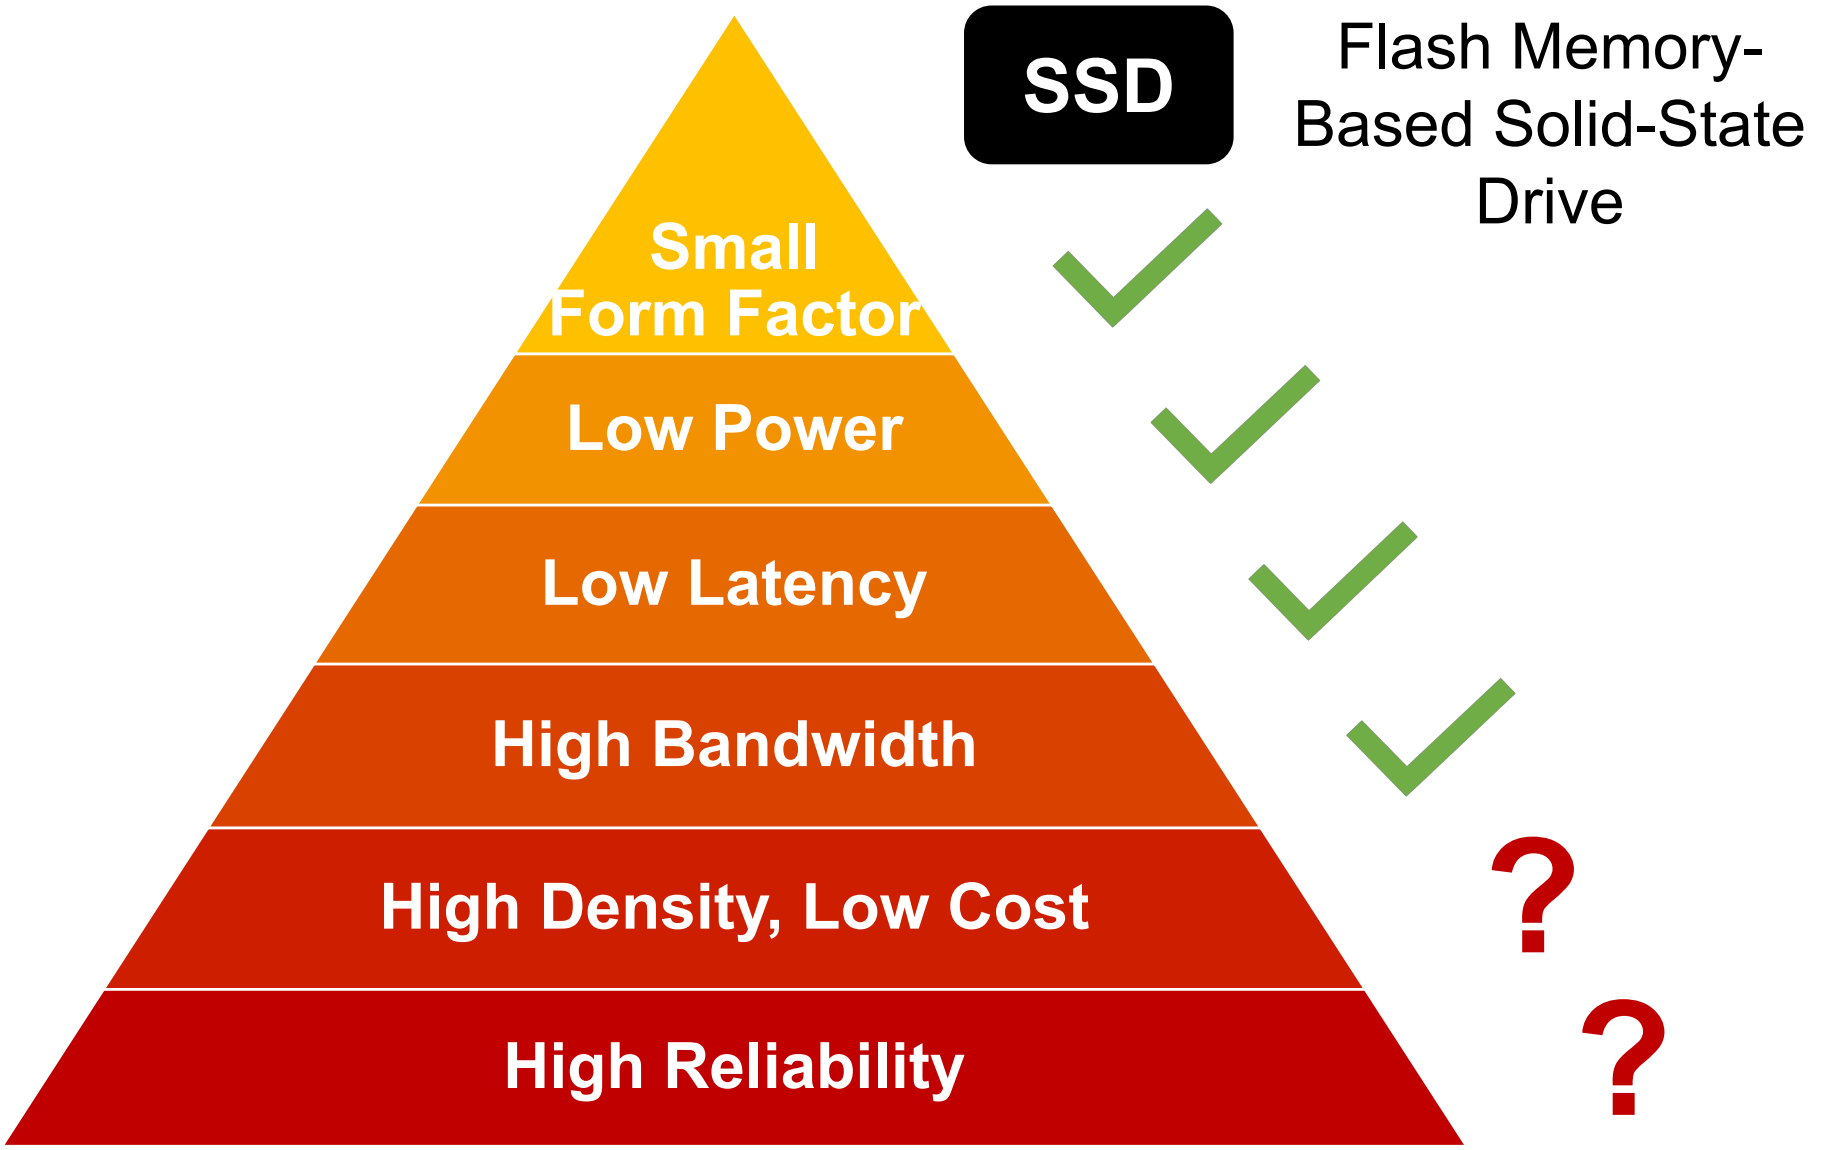

Degraded Flash Reliability Increases Cost

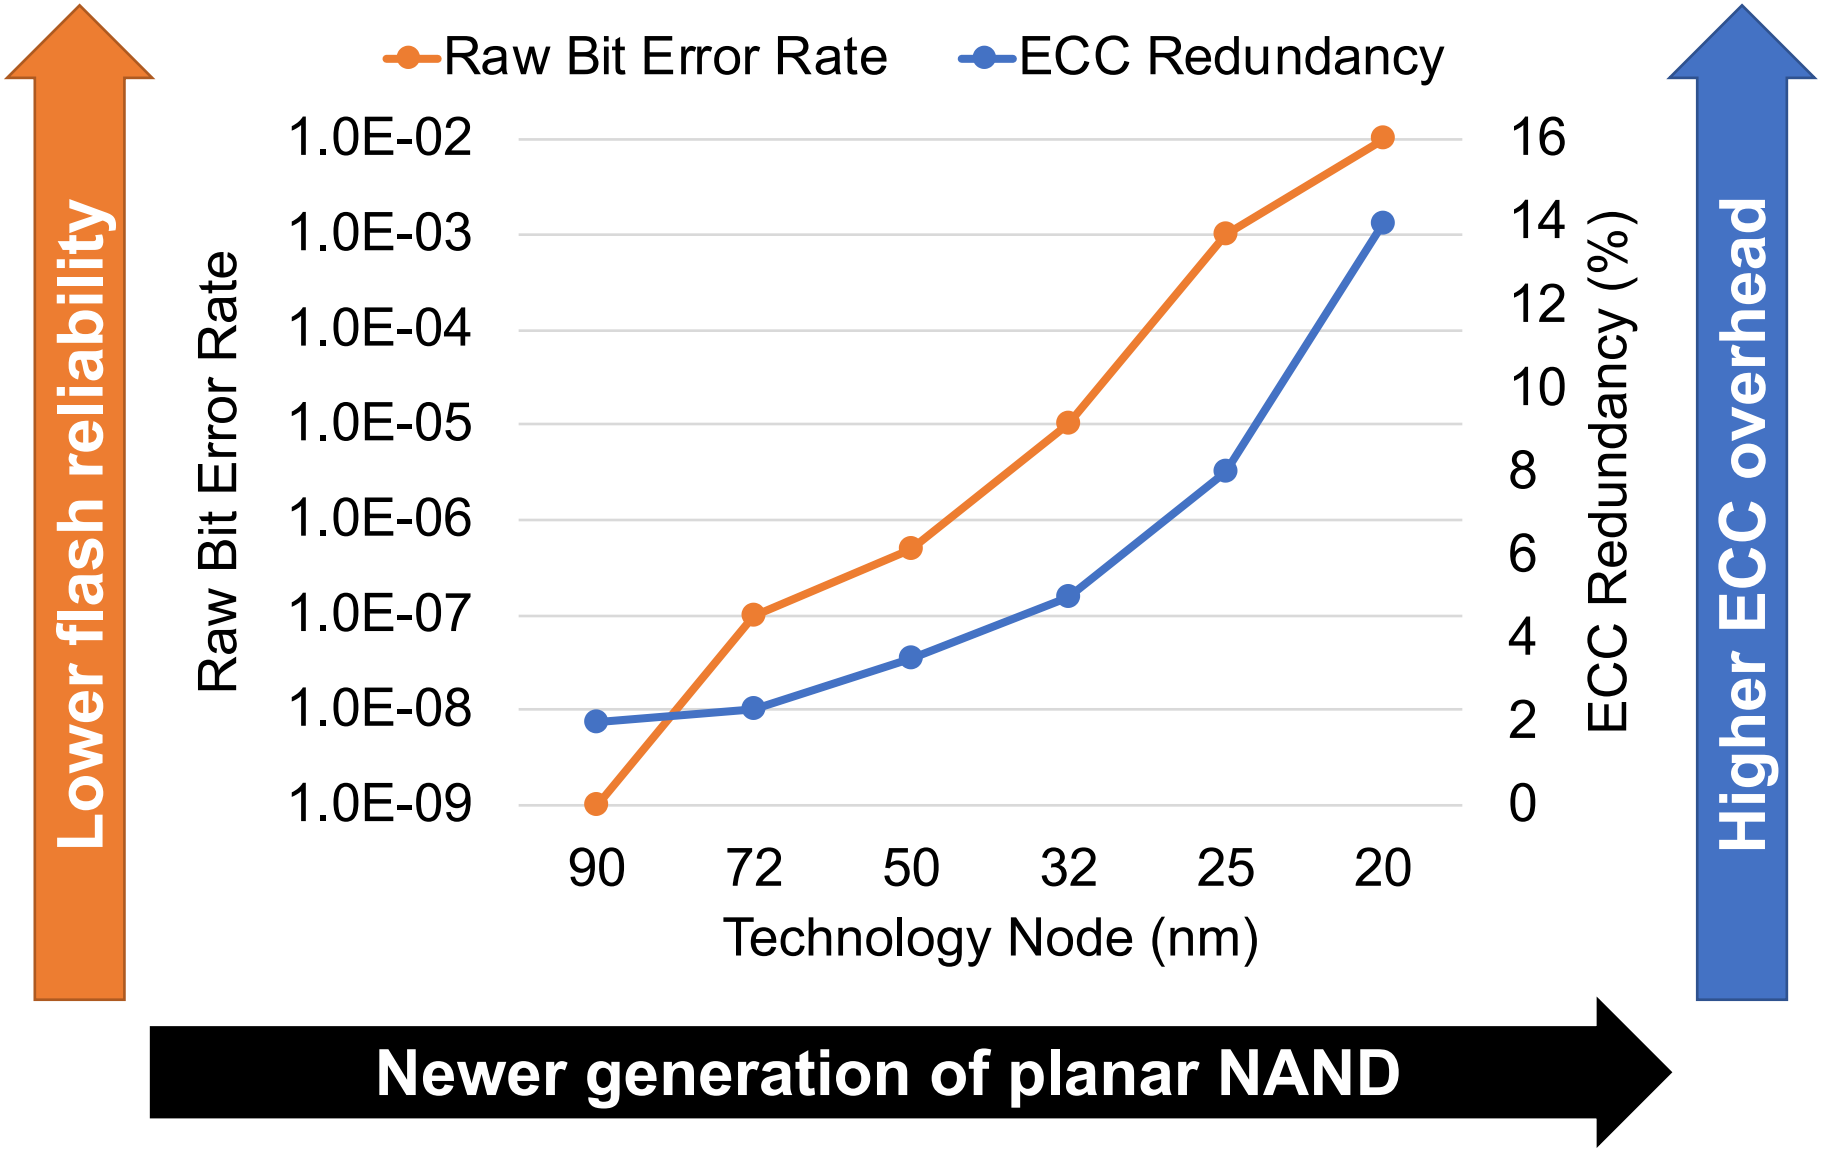

Causes of Raw Bit Errors

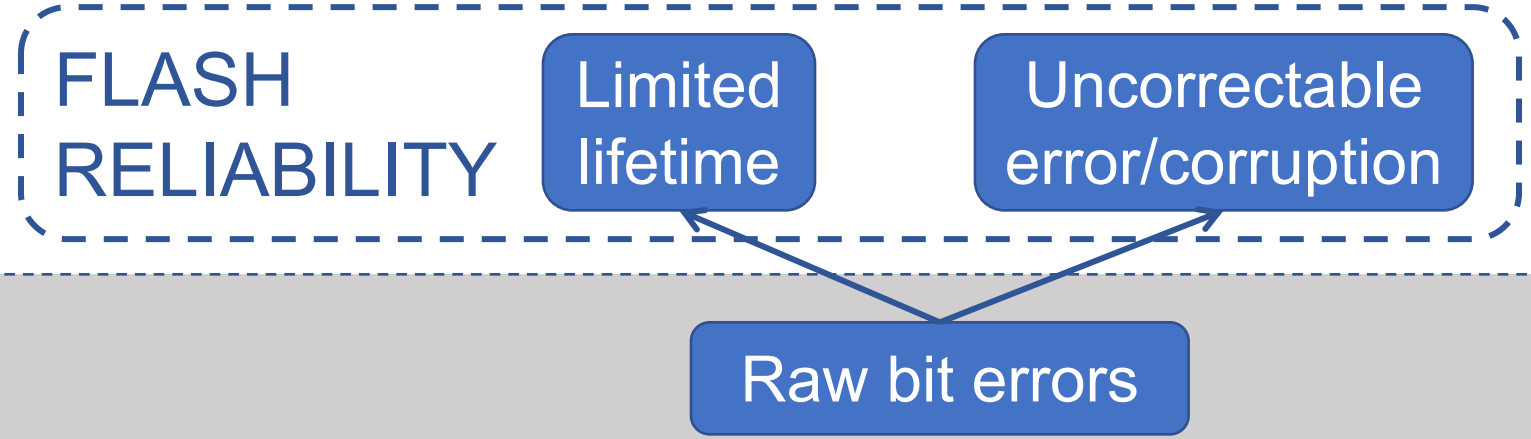

Causes of Raw Bit Errors

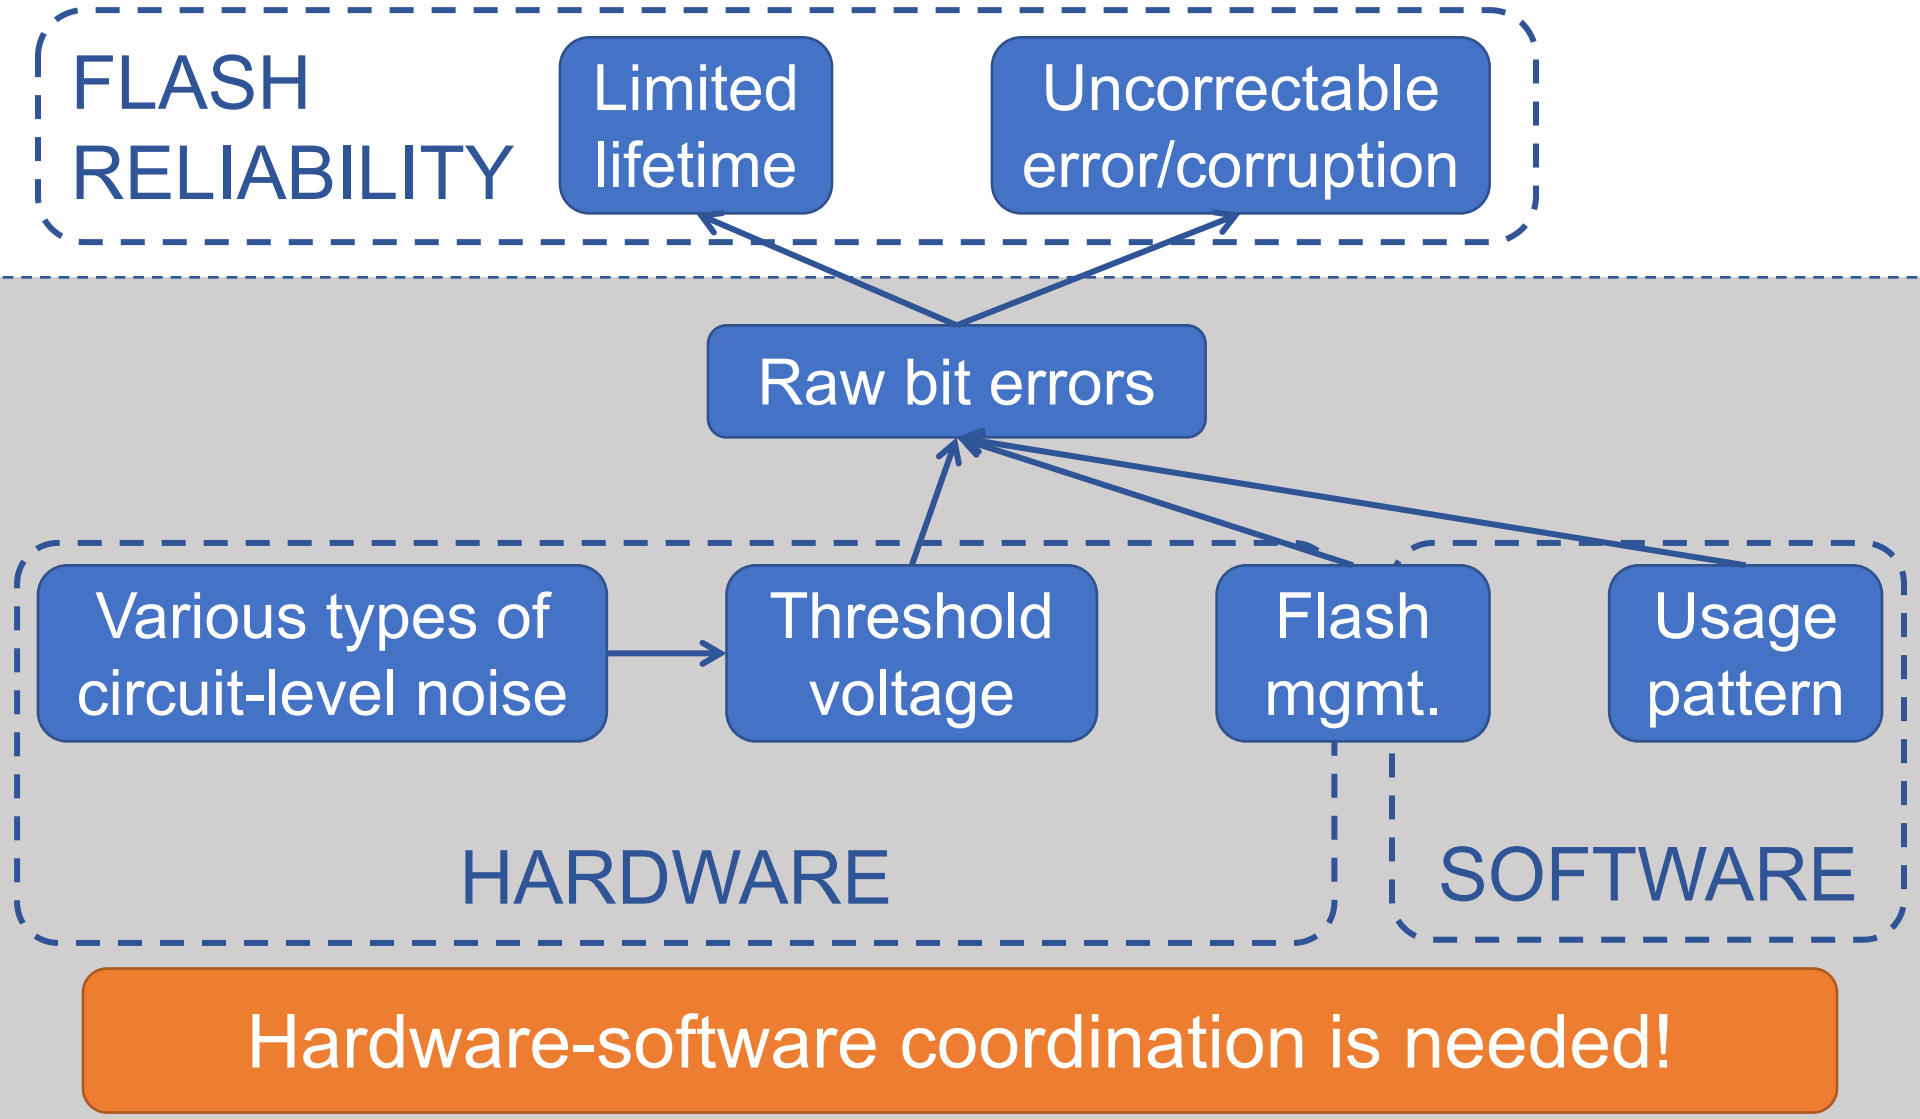

Hardware-software coordination is needed!

Future Research Directions

- SSD Errors At Scale
  - Problem
    - *Characterizing process variation requires lots of flash devices*
  - Directions
    - *Understanding other component failures*
    - *Deploy our proposed techniques at scale*
    - *Predicting and Preventing SSD Failures*
    - *Understanding and tolerating reliability variation across SSDs*
- Enabling Cold Storage in SSD
  - Problem
    - *Cost/GB is higher for SSD than for HDD*
  - Directions
    - *Identifying suitable data for cold storage*
    - *Increase SSD retention time and capacity*

SSD Reliability-Cost Trade-off

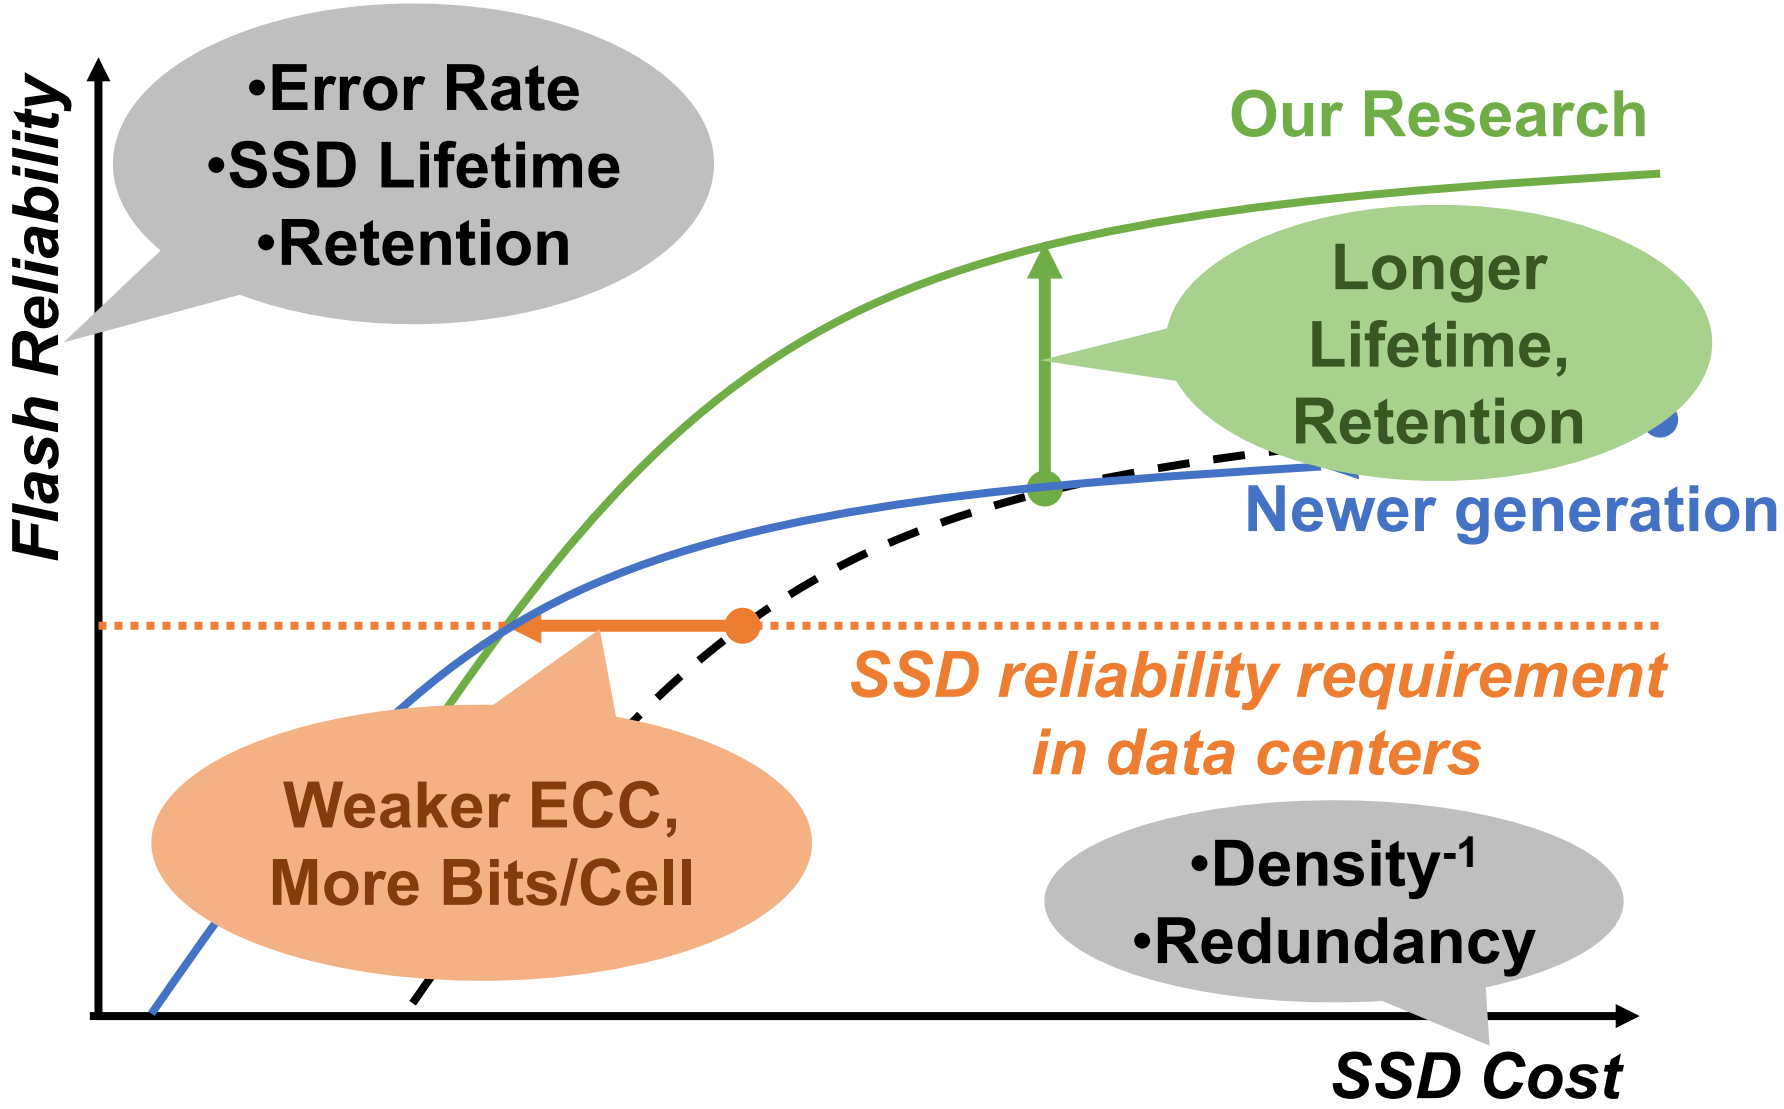

Summary

- Goal: Improve SSD reliability at low cost
- 3D NAND changes flash error characteristics
- **Real 3D NAND chips characterization**
  - Identify retention and process variation problems
- **HeatWatch**
  - Predict V_{opt} using dwell time and temperature
  - Improve lifetime by 3.85x, < 1.6 MB memory
- **Layer-Interleaved RAID**
  - Interleave layers and bits within each RAID group
  - Reduce 99% RBER by 66.9%

References

- *Y. Luo, et al., "**HeatWatch: Optimizing 3D NAND Read Operations With Self-Recovery and Temperature Awareness**," to appear HPCA, 2018*
- *Y. Luo, et al., "**Error Patterns in 3D NAND Flash Memory Devices: Characterization, Modeling, and Mitigation**," under submission, 2018*

Our other related work in this area:

- *Y. Luo, et al., "**Enabling Accurate and Practical Online Flash Channel Modeling for Modern MLC NAND Flash Memory**," IEEE JSAC, 2016*
- *Y. Luo, et al., "**WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management**," MSST, 2015*
- *Y. Cai, et al., "**Error Characterization, Mitigation, and Recovery in Flash Memory Based Solid-State Drives**," Proceedings of the IEEE, 2017 **(Invited Paper)***
- *Y. Cai, et al., "**Vulnerabilities in MLC NAND Flash Memory Programming: Experimental Analysis, Exploits, and Mitigation Techniques**," HPCA, 2017*
- *A. Fukami, et al., "**Improving the Reliability of Chip-Off Forensic Analysis of NAND Flash Memory Devices**," DFRWS EU, 2017 — **Best Paper Award***
- *Y. Cai, et al., "**Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery**," HPCA, 2015 — **Best Paper Runner Up***
- *Y. Cai, et al., "**Read Disturb Errors in MLC NAND Flash Memory: Characterization, Mitigation, and Recovery**," DSN, 2015*