

Modeling and Mitigating Early Retention Loss and Process Variation in 3D Flash

Saugata Ghose

Carnegie Mellon University



August 7, 2019
Santa Clara, CA

Carnegie Mellon

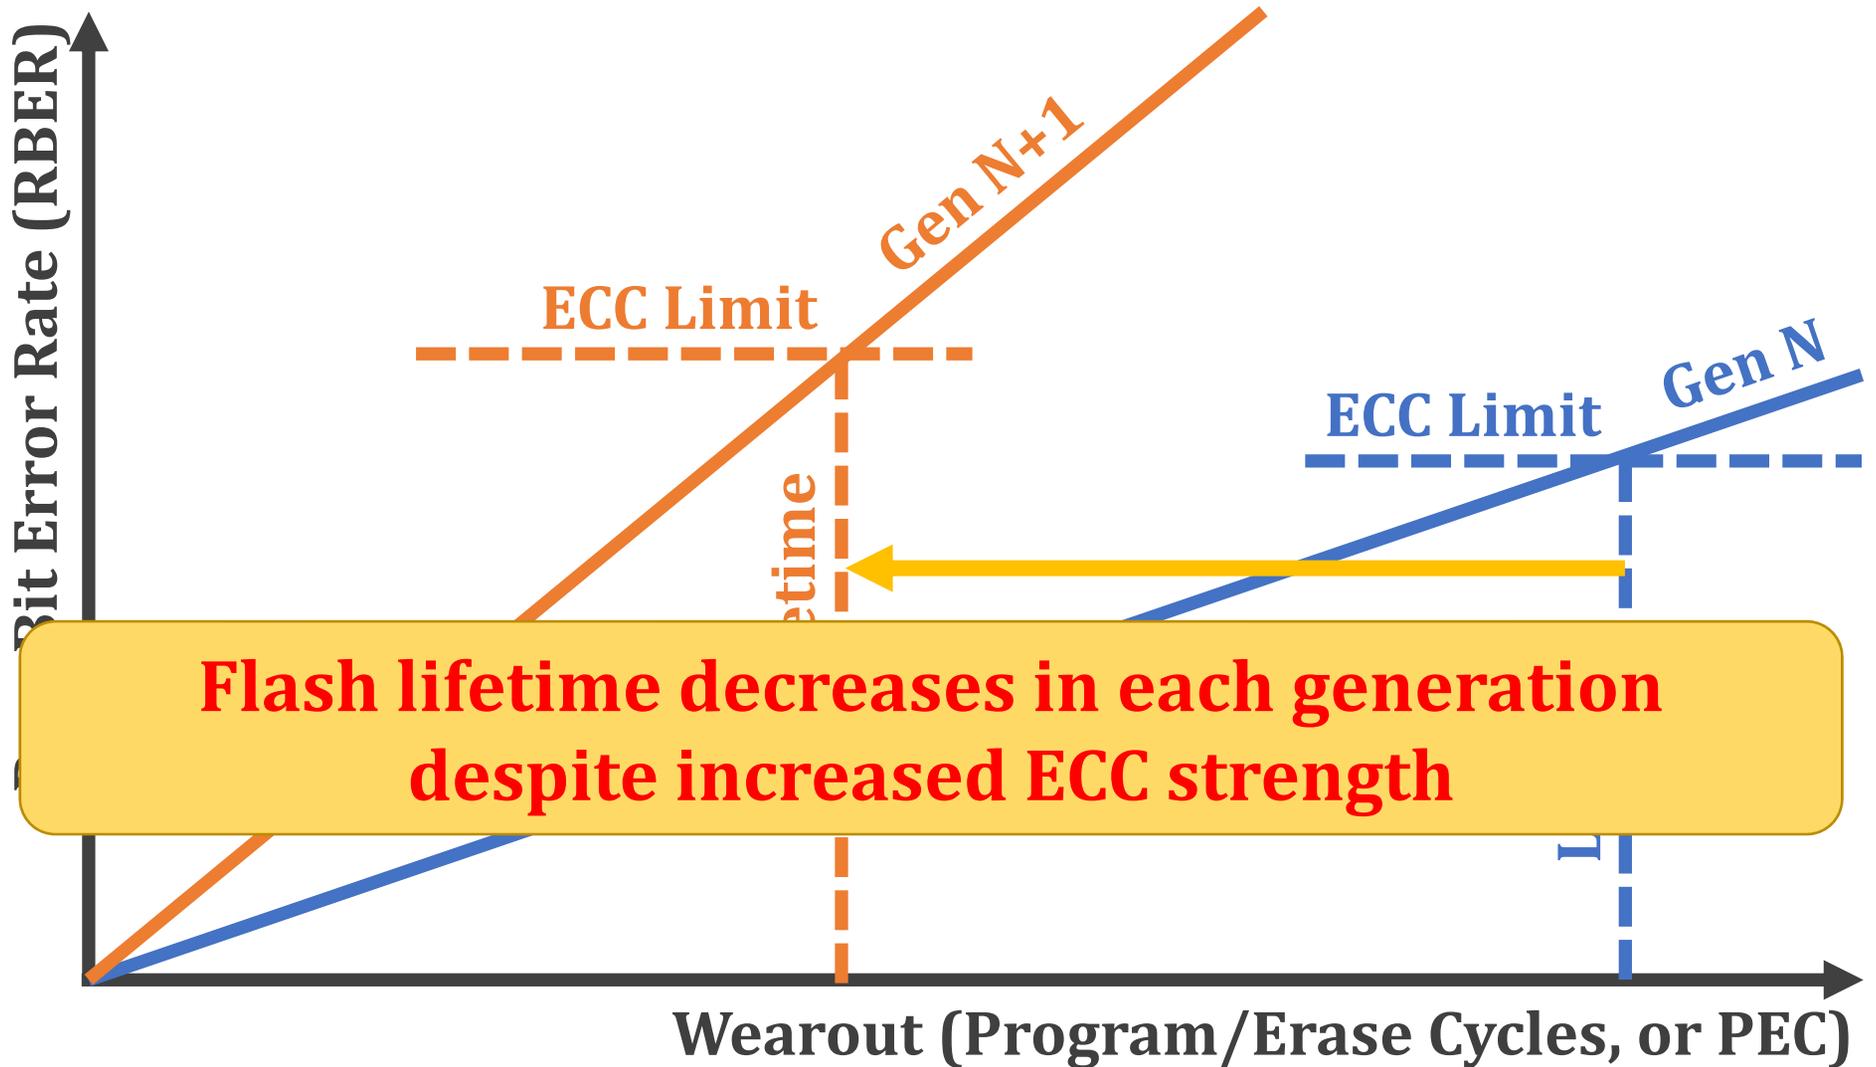


ETH zürich

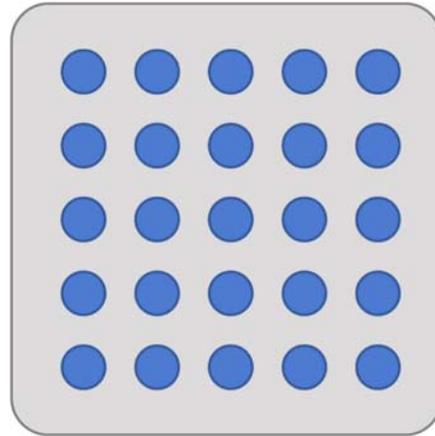
SAFARI



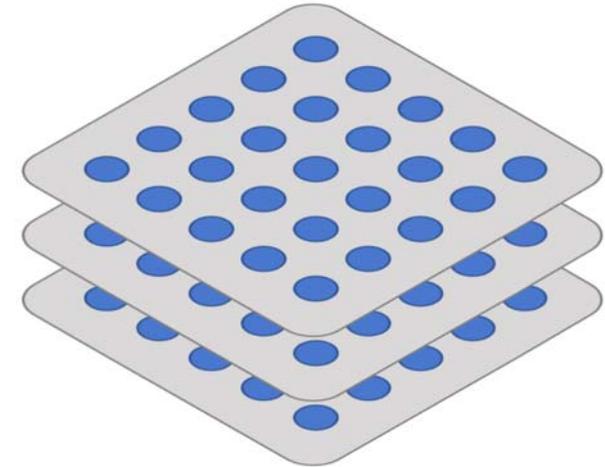
NAND Flash Memory Lifetime Problem



Planar vs. 3D NAND Flash Memory



**Planar NAND
Flash Memory**



**3D NAND
Flash Memory**

Scaling

Reduce flash cell size,
Reduce distance b/w cells

Increase # of layers

Reliability

Scaling hurts reliability

Not well studied!

Executive Summary

- Problem: 3D NAND error characteristics are **not well studied**
- Goal: *Understand & mitigate* 3D NAND errors to improve lifetime
- **Contribution 1: Characterize** real 3D NAND flash chips
 - **Process variation:** *21× error rate difference across layers*
 - **Early retention loss:** *Error rate increases by 10× after 3 hours*
 - **Retention interference:** *Not observed before in planar NAND*
- **Contribution 2: Model** RBER and threshold voltage
 - *RBER (raw bit error rate) variation model*
 - *Retention loss model*
- **Contribution 3: Mitigate** 3D NAND flash errors
 - **LaVAR:** *Layer Variation Aware Reading*
 - **LI-RAID:** *Layer-Interleaved RAID*
 - **ReMAR:** *Retention Model Aware Reading*
 - *Improve flash lifetime by 1.85× or reduce ECC overhead by 78.9%*

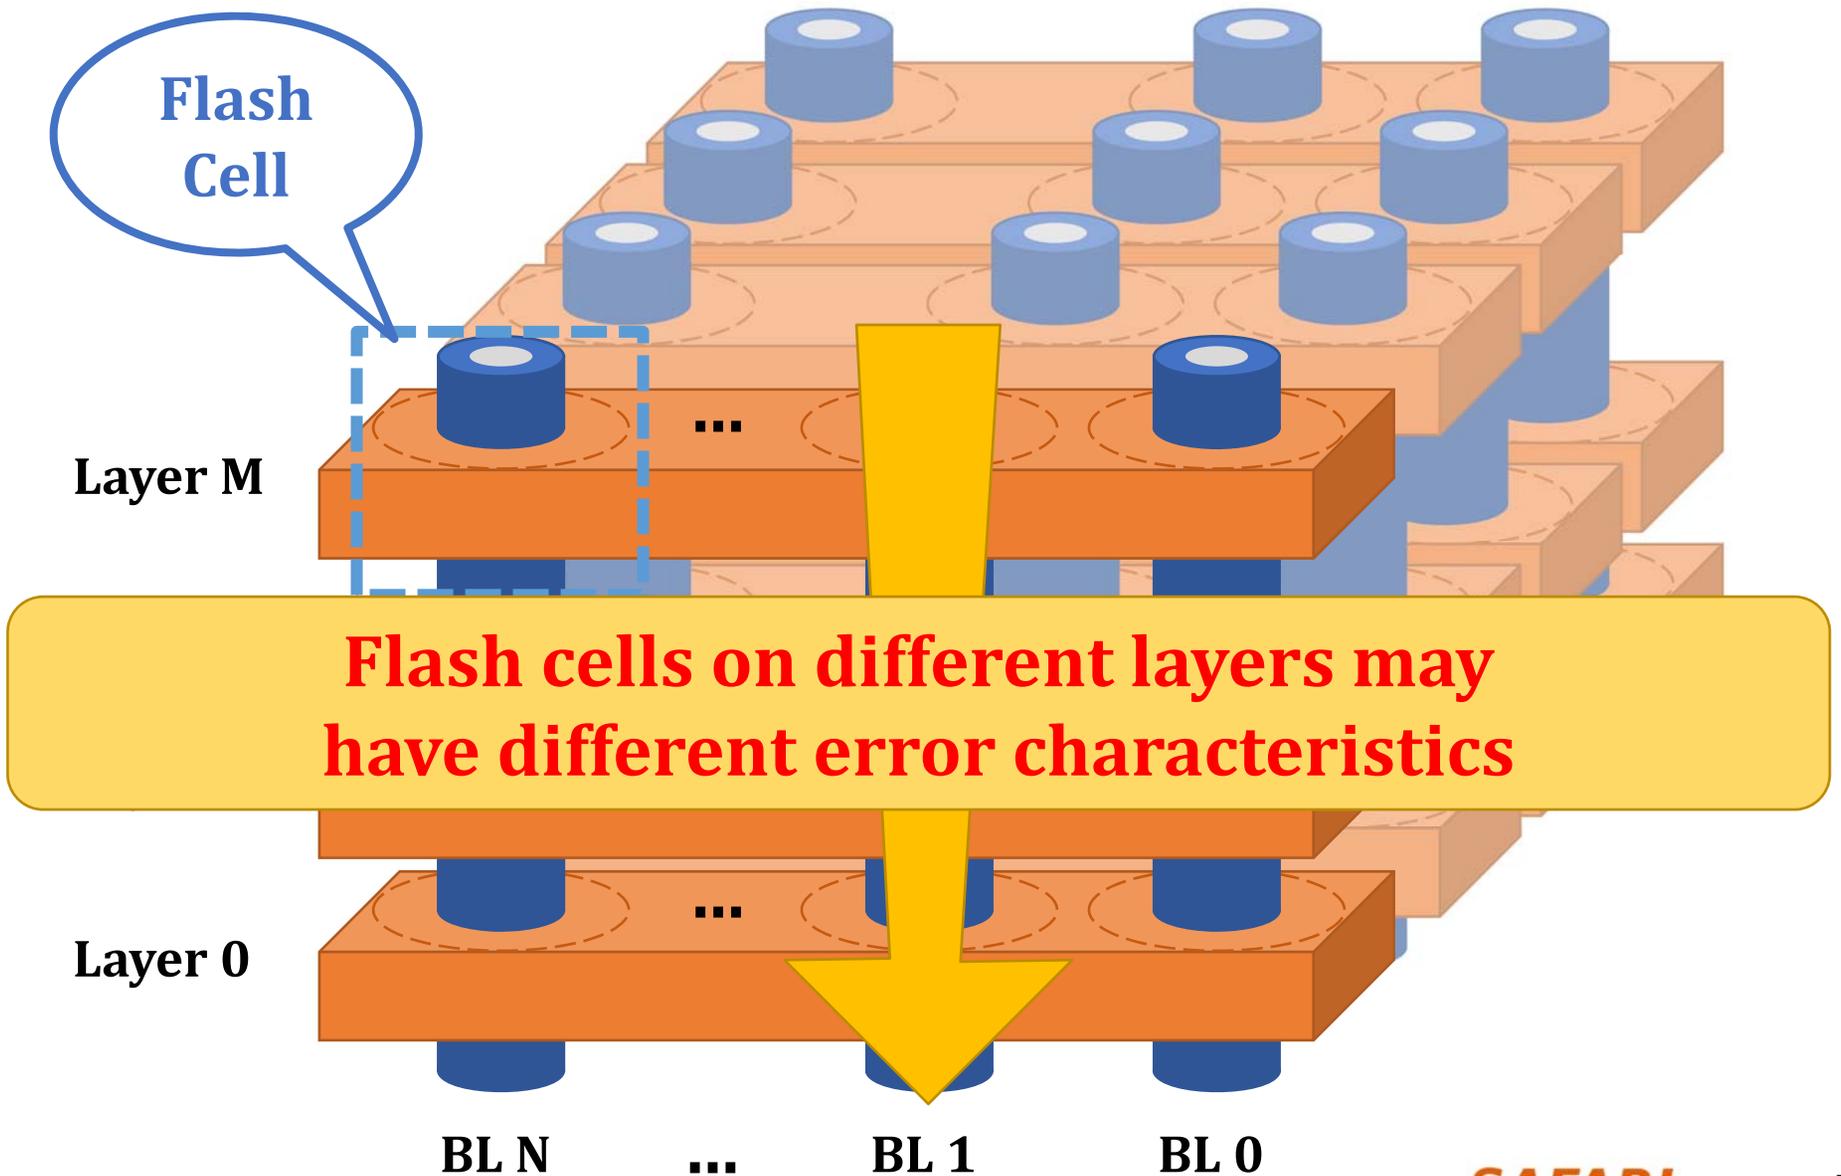
Agenda

- Background & Introduction
- Contribution 1: Characterize real 3D NAND flash chips
- Contribution 2: Model RBER and threshold voltage
- Contribution 3: Mitigate 3D NAND flash errors
- Conclusion

Agenda

- Background & Introduction
- **Contribution 1: Characterize real 3D NAND flash chips**
 - Process variation
 - Early retention loss
 - Retention interference
- Contribution 2: Model RBER and threshold voltage
- Contribution 3: Mitigate 3D NAND flash errors
- Conclusion

Process Variation Across Layers

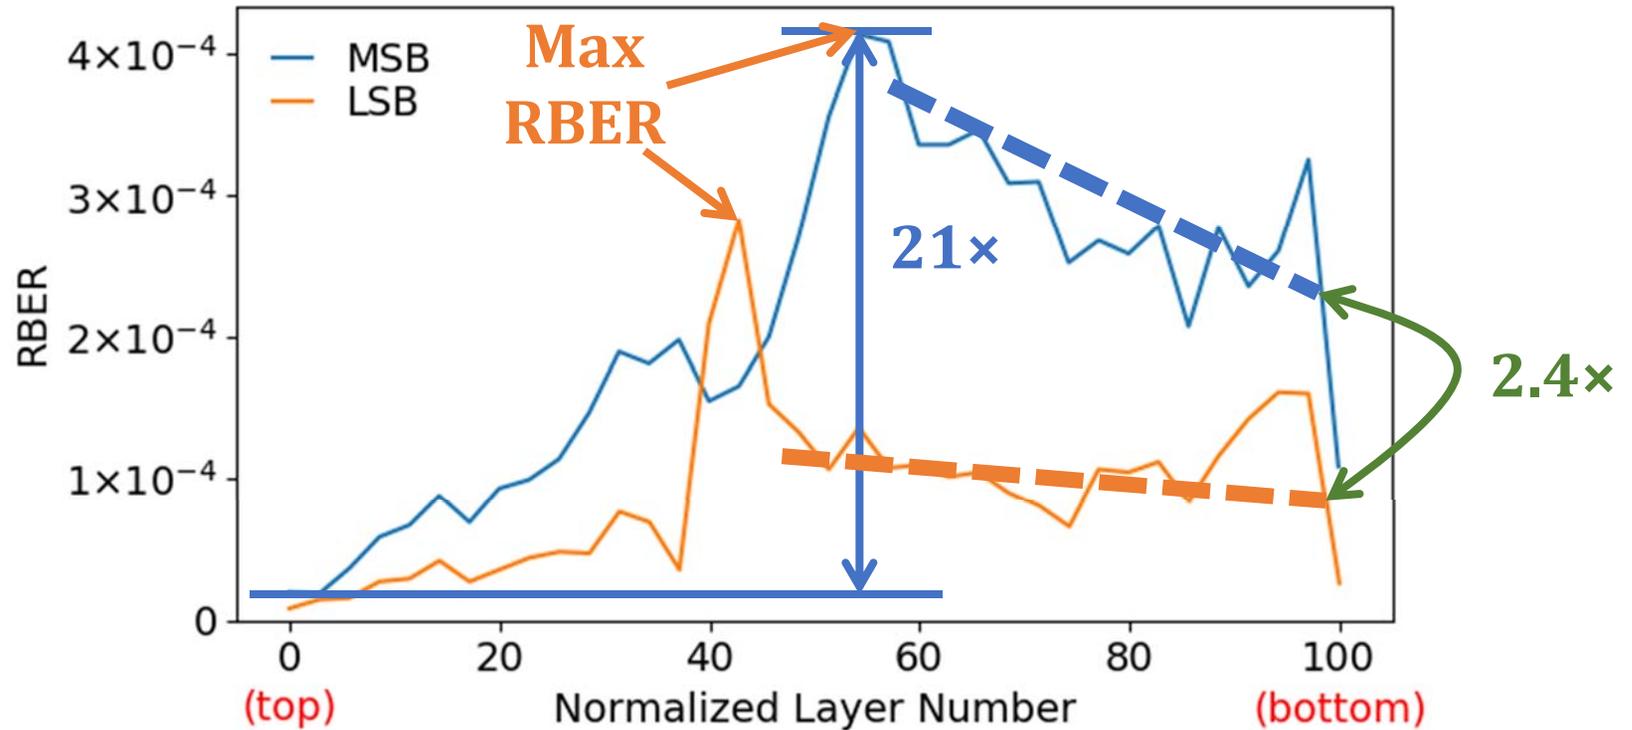


Characterization Methodology

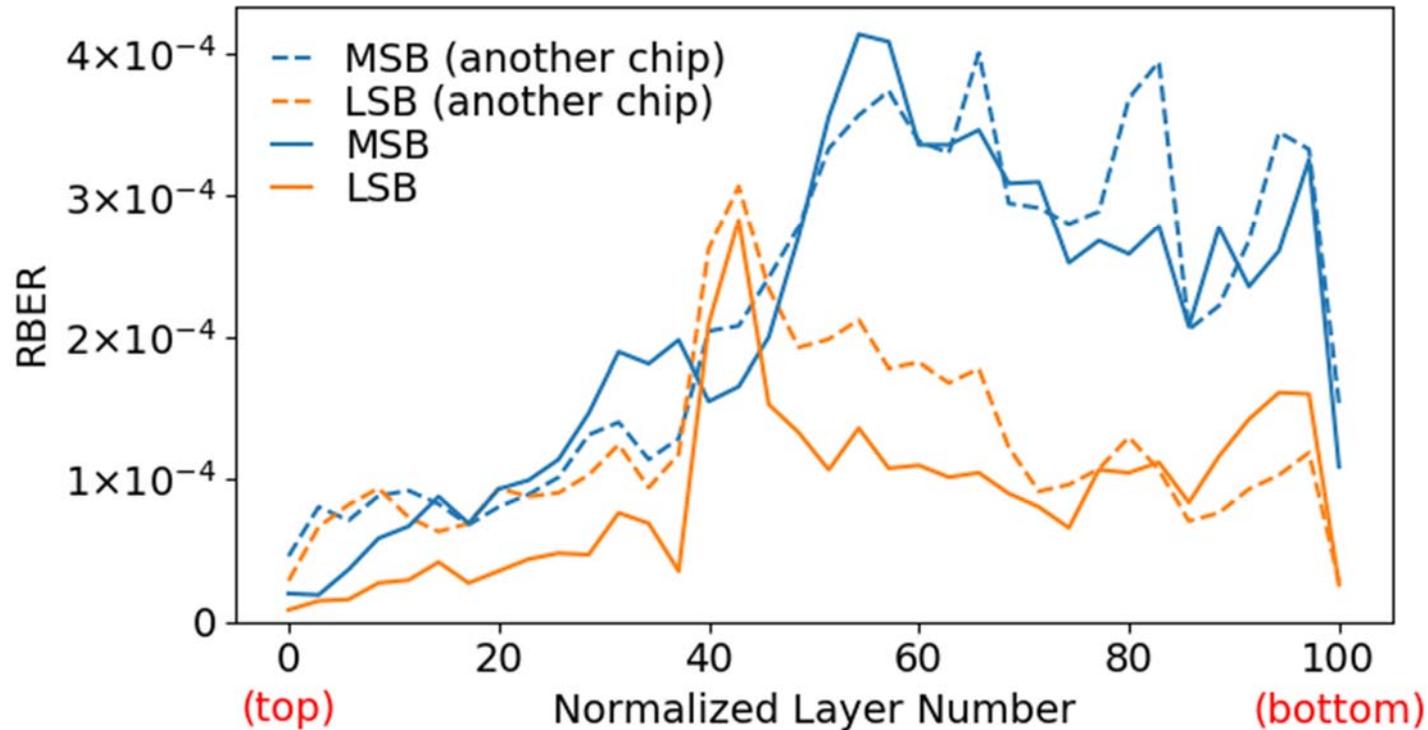
- **Modified firmware version in the flash controller**
 - Controls the read reference voltage of the flash chip
 - Bypasses ECC to get raw data (with raw bit errors)
- **Analysis and post-processing of the data on the server**



Layer-to-Layer Process Variation



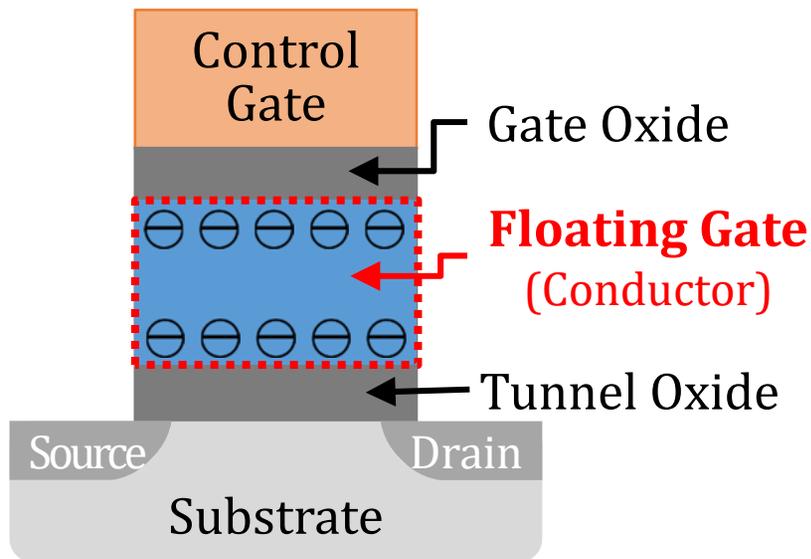
Layer-to-Layer Process Variation



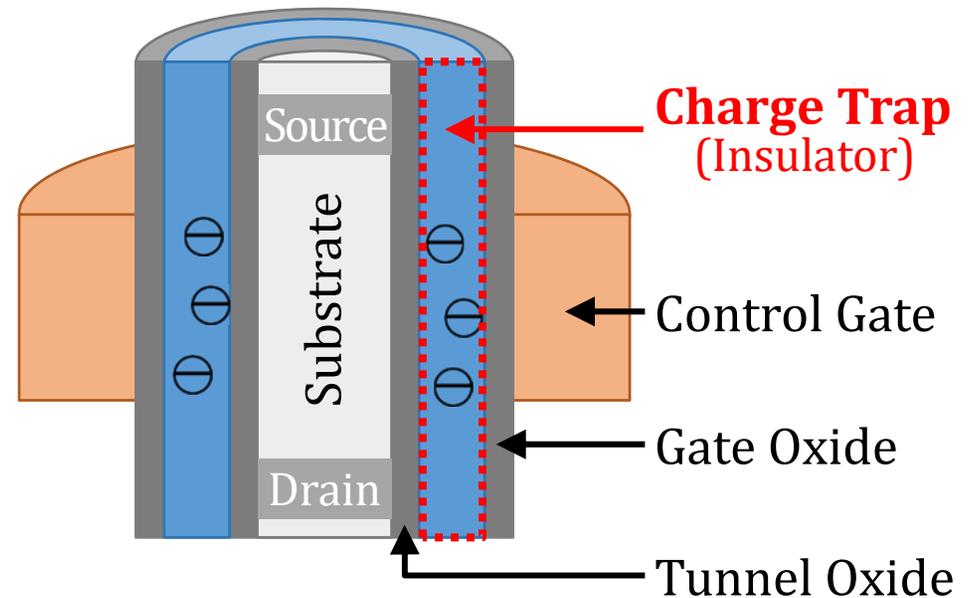
**Large RBER variation
across layers and LSB-MSB pages**

Retention Loss Phenomenon

Planar NAND Cell

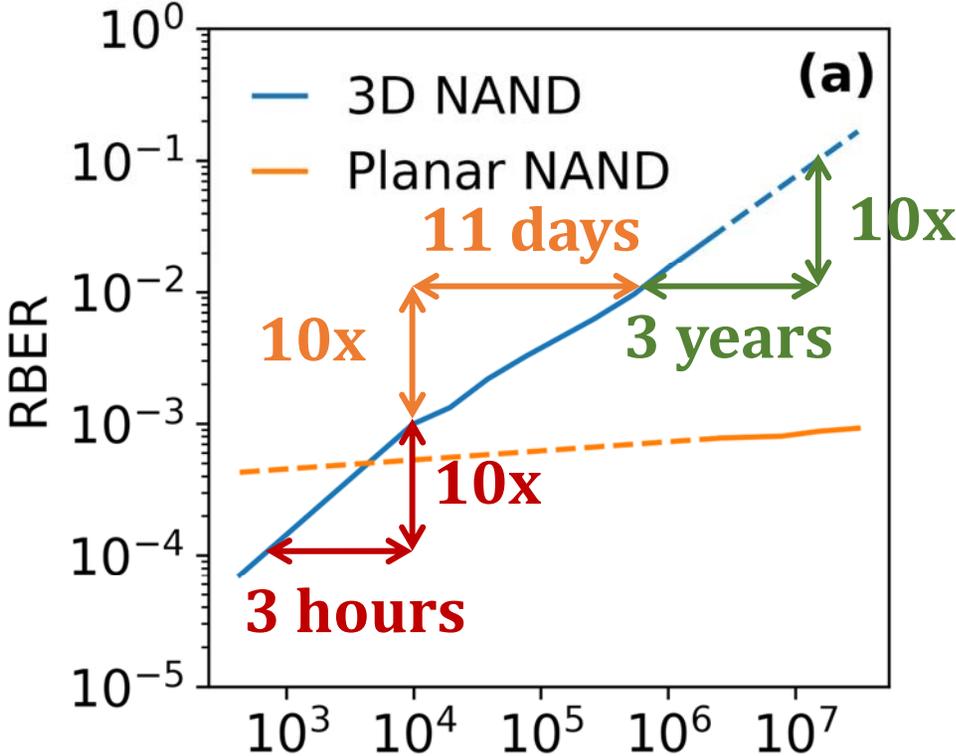


3D NAND Cell



**Most dominant type of error in planar NAND.
Is this true for 3D NAND as well?**

Early Retention Loss



Retention errors increase quickly immediately after programming

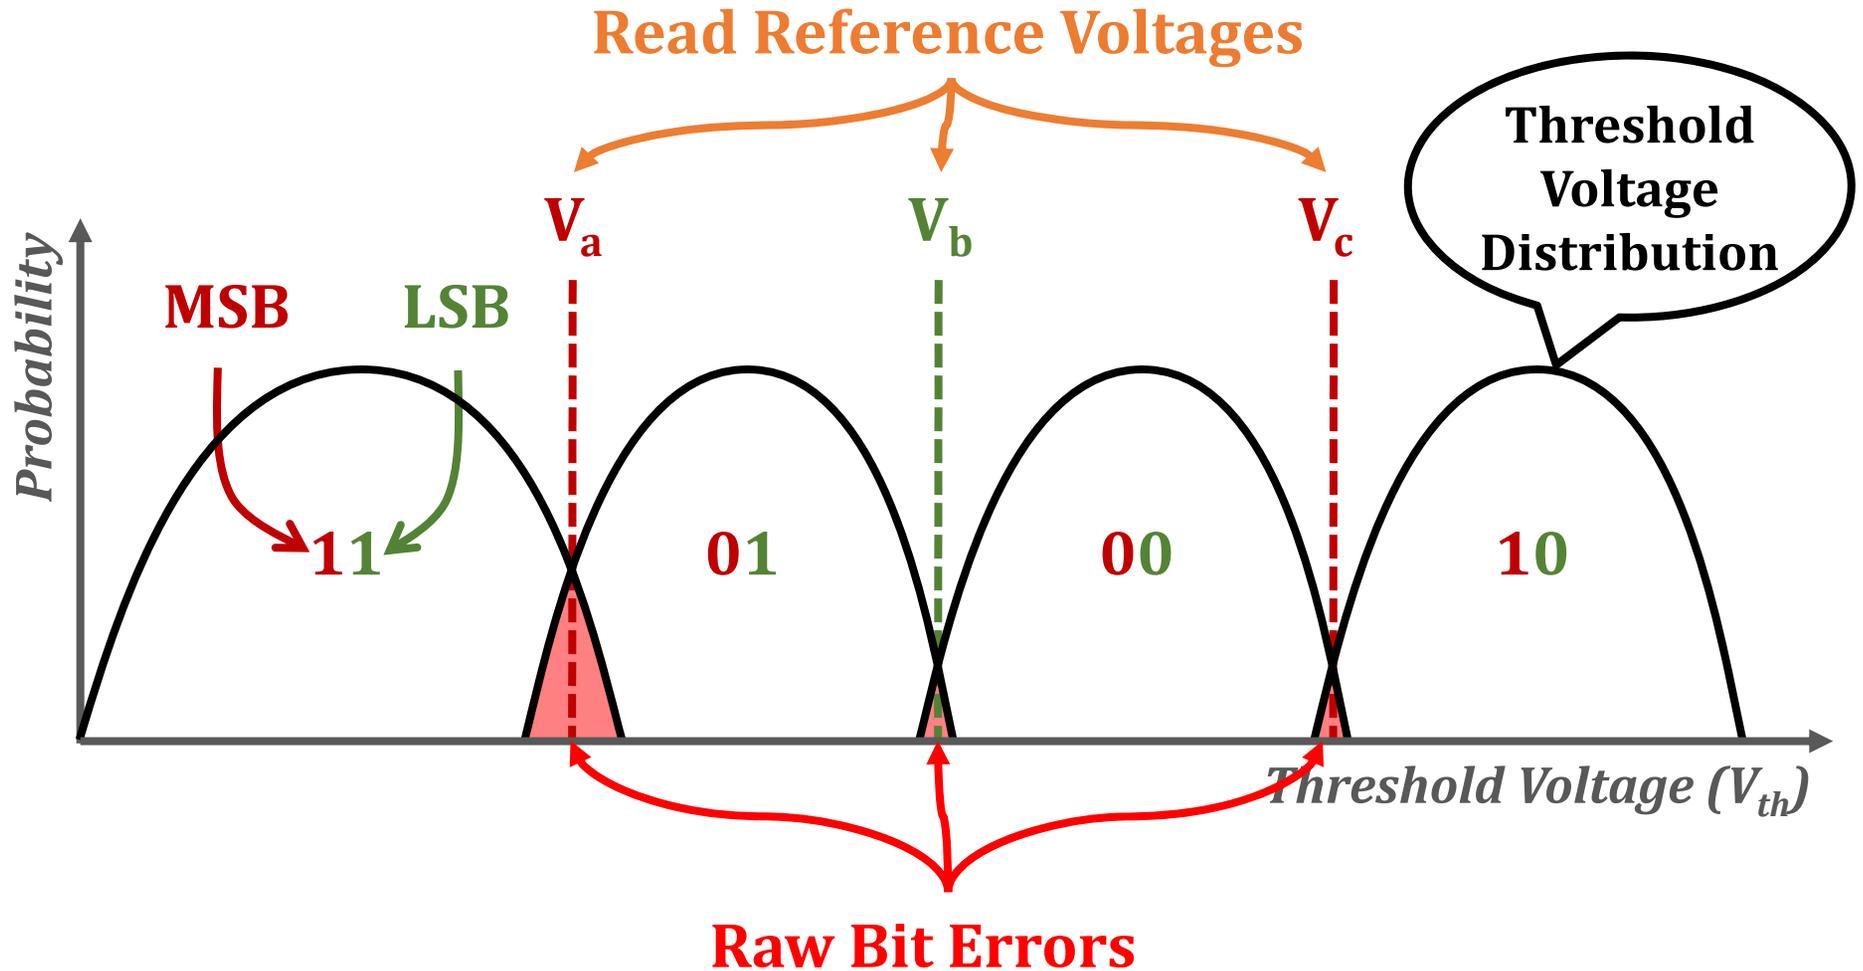
Characterization Summary

- **Layer-to-layer process variation**
 - Large RBER variation across layers and LSB-MSB pages
 - → Need new mechanisms to tolerate RBER variation!
- **Early retention loss**
 - RBER increases quickly after programming
 - → Need new mechanisms to tolerate retention errors!
- **Retention interference**
 - Amount of retention loss correlated with neighbor cells' states
 - → Need new mechanisms to tolerate retention interference!
- **More *threshold voltage* and *RBER* results in the paper:**
3D NAND P/E cycling, program interference, read disturb, read variation, bitline-to-bitline process variation
- **Our approach** based on insights developed via our experimental characterization: Develop **error models**, and build online **error mitigation mechanisms** using the models

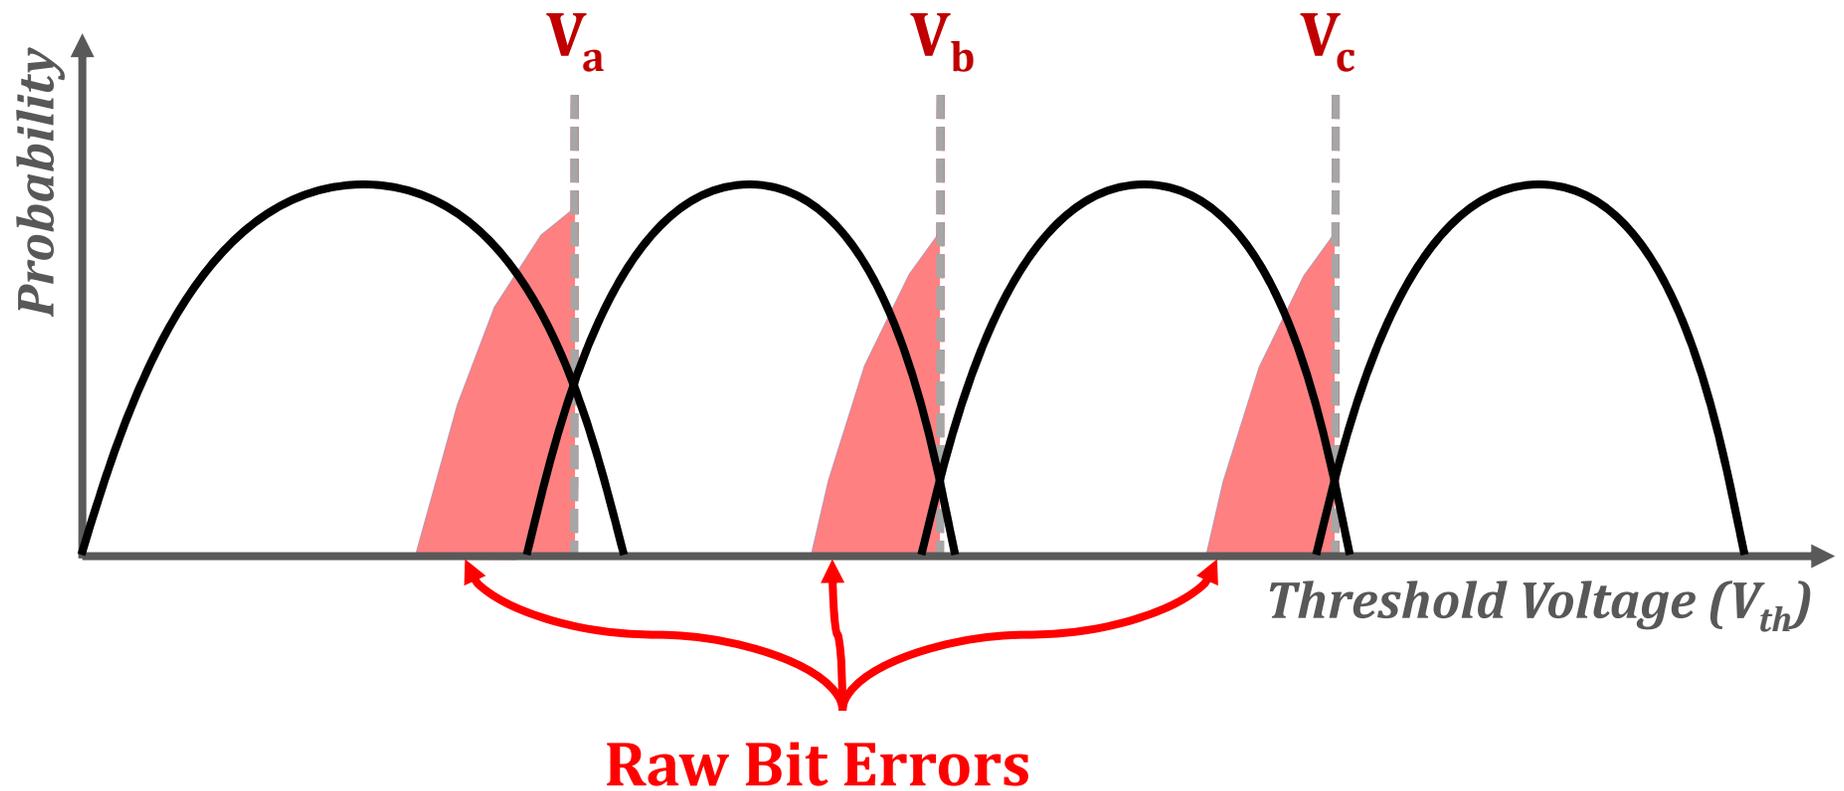
Agenda

- Background & Introduction
- Contribution 1: Characterize real 3D NAND flash chips
- Contribution 2: Model RBER and threshold voltage
 - Retention loss model
 - RBER variation model
- Contribution 3: Mitigate 3D NAND flash errors
- Conclusion

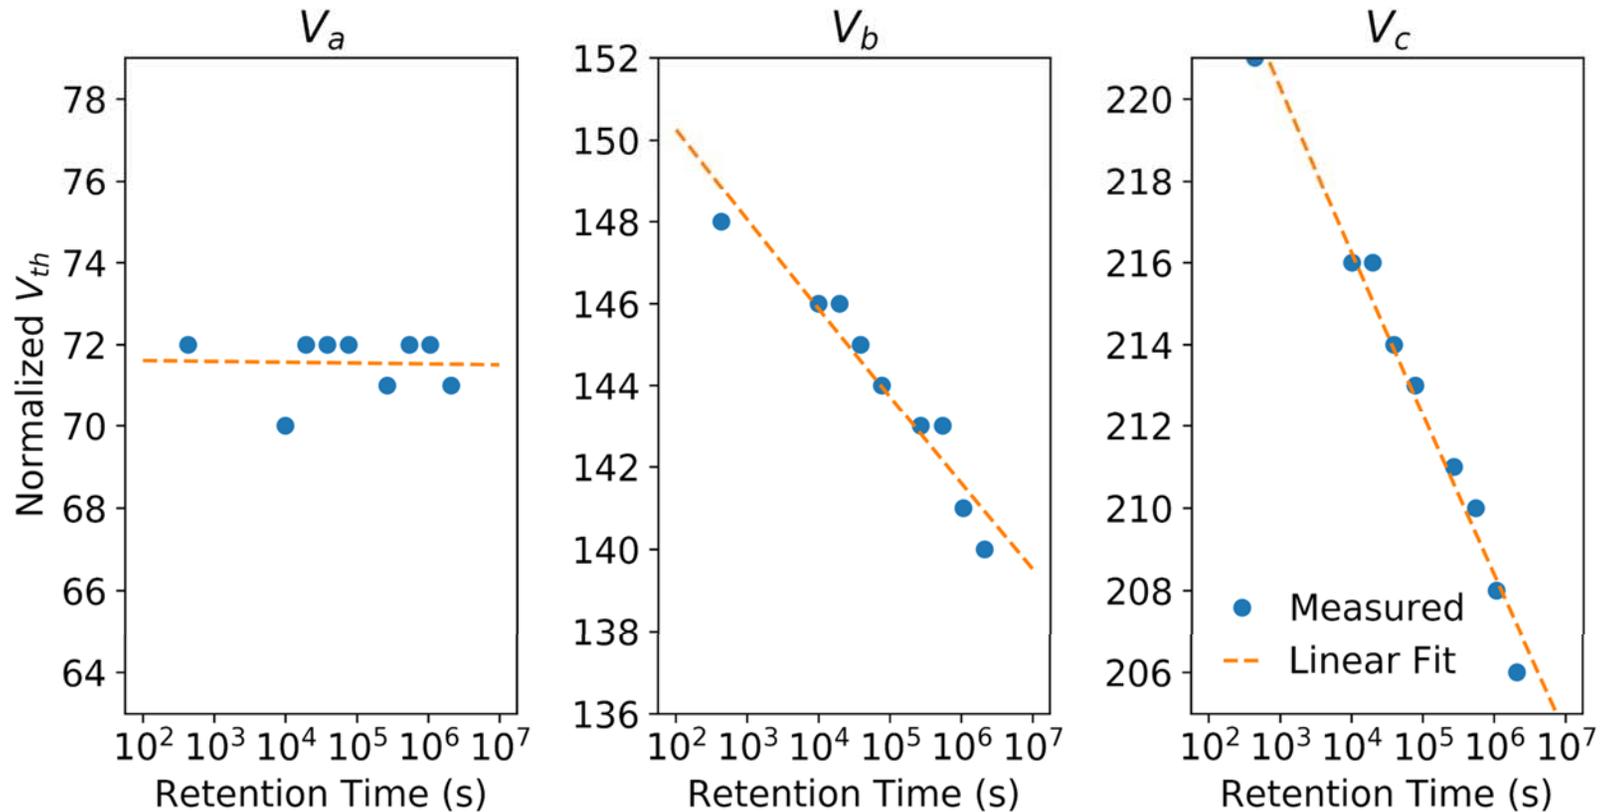
What Do We Model?



Optimal Read Reference Voltage



Retention Loss Model

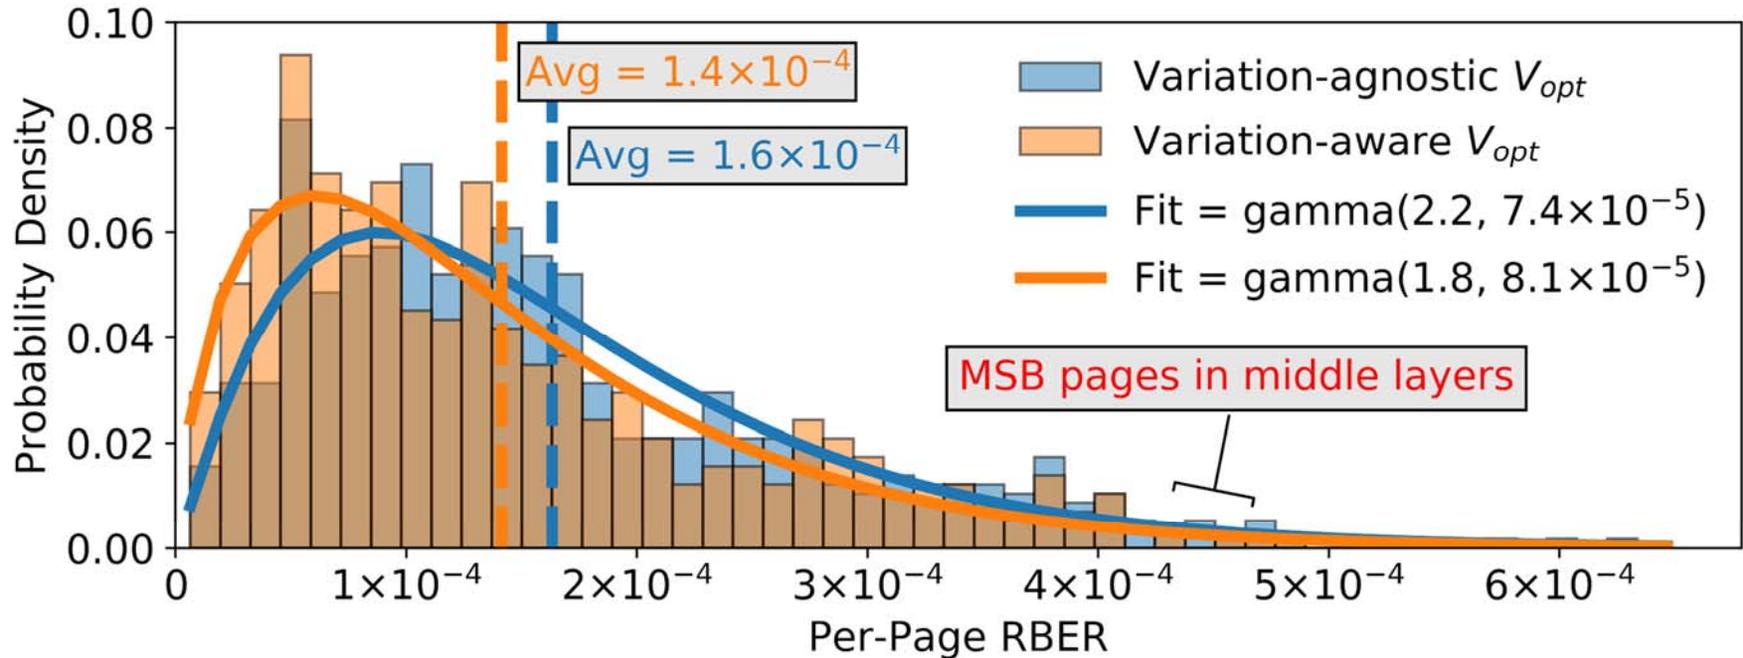


Early retention loss can be modeled as a simple linear function of $\log(\text{retention time})$

Retention Loss Model

- Goal: Develop a simple linear model that can be used online
- Models
 - Optimal read reference voltage (V_b and V_c)
 - Raw bit error rate ($\log(RBER)$)
 - Mean and standard deviation of threshold voltage distribution (μ and σ)
- As a function of
 - Retention time ($\log(t)$)
 - P/E cycle count (PEC)
- e.g., $V_{opt} = (\alpha \times PEC + \beta) \times \log(t) + \gamma \times PEC + \delta$
- Model error <1 step for V_b and V_c
- Adjusted $R^2 > 89\%$

RBER Variation Model



Variation-agnostic V_{opt}

- Same V_{ref} for all layers optimized for the entire block

RBER distribution follows gamma distribution

KL-divergence error = 0.09

Agenda

- Background & Introduction
- Contribution 1: Characterize real 3D NAND flash chips
- Contribution 2: Model RBER and threshold voltage
- Contribution 3: Mitigate 3D NAND flash errors
 - LaVAR: Layer Variation Aware Reading
 - LI-RAID: Layer-Interleaved RAID
 - ReMAR: Retention Model Aware Reading
- Conclusion

LaVAR: Layer Variation Aware Reading

- **Layer-to-layer process variation**
 - Error characteristics are different in each layer
- **Goal:** Adjust read reference voltage **for each layer**
- **Key Idea:** Learn a **voltage offset (Offset)** for each layer
 - $V_{opt}^{Layer\ aware} = V_{opt}^{Layer\ agnostic} + Offset$
- **Mechanism**
 - **Offset:** Learned once for each chip & stored in a table
 - *Uses (2 × Layers) Bytes memory per chip*
 - $V_{opt}^{Layer\ agnostic}$: Predicted by any existing V_{opt} model
 - *E.g., ReMAR [Luo+Sigmetrics'18], HeatWatch [Luo+HPCA'18], OFCM [Luo+JSAC'16], ARVT [Papandreou+GLSVLSI'14]*
- Reduces RBER on average by **43%**
(based on our characterization data)

LI-RAID: Layer-Interleaved RAID

- **Layer-to-layer process variation**
 - Worst-case RBER much higher than average RBER
- **Goal:** Significantly reduce worst-case RBER
- **Key Idea**
 - Group flash pages on *less reliable layers* with pages on *more reliable layers*
 - Group *MSB pages* with *LSB pages*
- **Mechanism**
 - Reorganize RAID layout to eliminate worst-case RBER
 - **<0.4% storage overhead**

Conventional RAID

<i>Wordline #</i>	<i>Layer #</i>	<i>Page</i>	Chip 0	Chip 1	Chip 2	Chip 3
<i>0</i>	<i>0</i>	<i>MSB</i>	Group 0	Group 0	Group 0	Group 0
<i>0</i>	<i>0</i>	<i>LSB</i>	Group 1	Group 1	Group 1	Group 1
<i>1</i>	<i>1</i>	<i>MSB</i>	Group 2	Group 2	Group 2	Group 2
<i>1</i>	<i>1</i>	<i>LSB</i>	Group 3	Group 3	Group 3	Group 3
<i>2</i>	<i>2</i>	<i>MSB</i>	Group 4	Group 4	Group 4	Group 4
<i>2</i>	<i>2</i>	<i>LSB</i>	Group 5	Group 5	Group 5	Group 5
<i>3</i>	<i>3</i>	<i>MSB</i>	Group 6	Group 6	Group 6	Group 6
<i>3</i>	<i>3</i>	<i>LSB</i>	Group 7	Group 7	Group 7	Group 7

**Worst-case RBER in any layer
limits the lifetime of conventional RAID**

LI-RAID: Layer-Interleaved RAID

<i>Wordline #</i>	<i>Layer #</i>	<i>Page</i>	Chip 0	Chip 1	Chip 2	Chip 3
0	0	MSB	Group 0	Blank	Group 4	Group 3
0	0	LSB	Group 1	Blank	Group 5	Group 2
1	1	MSB	Group 2	Group 1	Blank	Group 5
1	1	LSB	Group 3	Group 0	Blank	Group 4
2	2	MSB	Group 4	Group 3	Group 0	Blank
2	2	LSB	Group 5	Group 2	Group 1	Blank
3	3	MSB	Blank	Group 5	Group 2	Group 1
3	3	LSB	Blank	Group 4	Group 3	Group 0

Any page with worst-case RBER can be corrected by other reliable pages in the RAID group

LI-RAID: Layer-Interleaved RAID

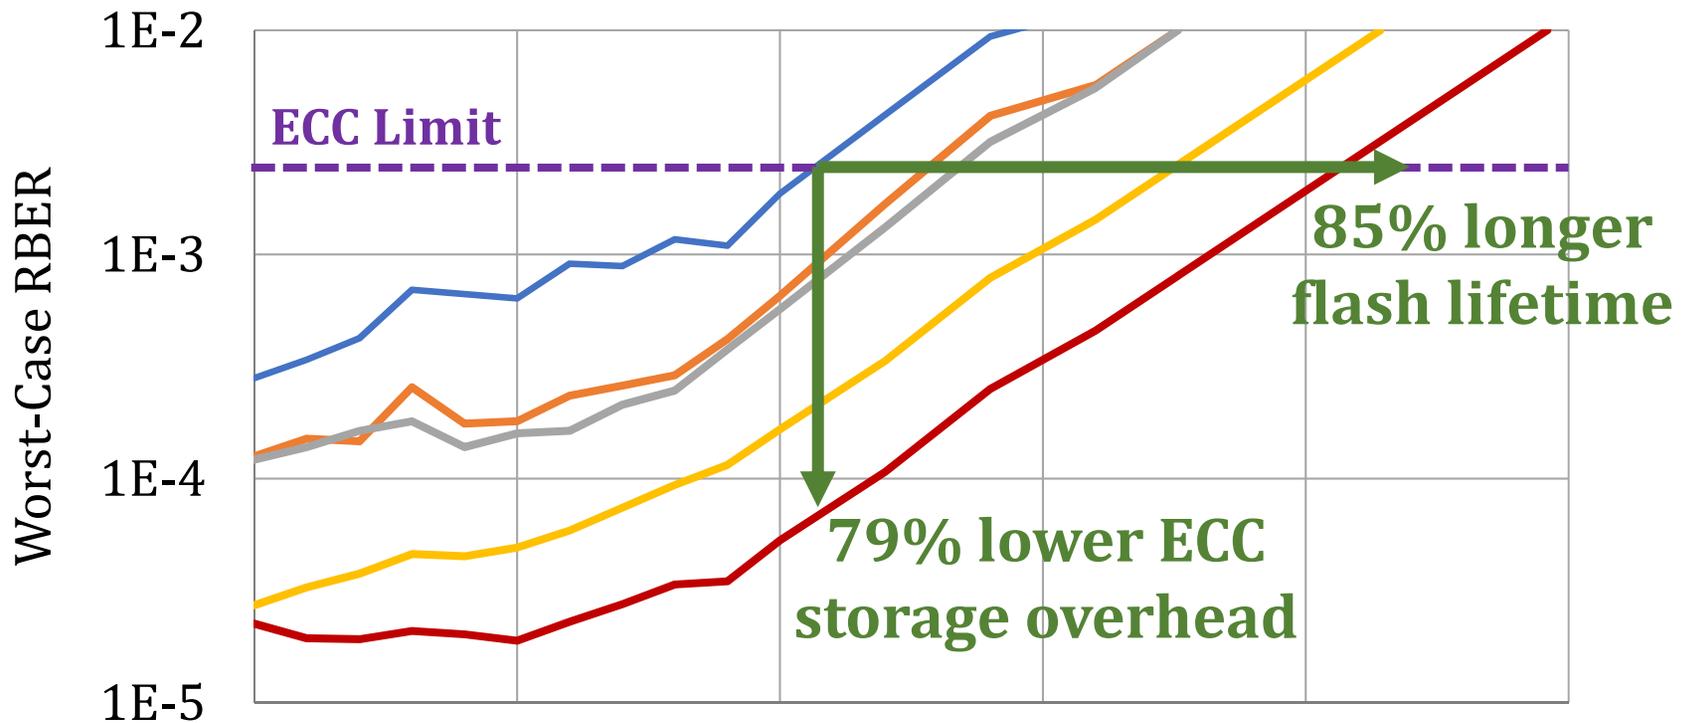
- **Layer-to-layer process variation**
 - Worst-case RBER much higher than average RBER
- **Goal:** Significantly reduce worst-case RBER
- **Key Idea**
 - Group flash pages on *less reliable layers* with pages on *more reliable layers*
 - Group *MSB pages* with *LSB pages*
- **Mechanism**
 - Reorganize RAID layout to eliminate worst-case RBER
 - **<0.8% storage overhead**
- Reduces worst-case RBER by **66.9%**
(based on our characterization data)

ReMAR: Retention Model Aware Reading

- **Early retention loss**
 - Threshold voltage shifts quickly after programming
- **Goal: Adjust read reference voltages based on retention loss**
- **Key Idea:** Learn and use a retention loss model online
- **Mechanism**
 - Periodically characterize and learn retention loss model online
 - Retention time = Read timestamp - Write timestamp
 - *Uses 800 KB memory to store program time of each block*
 - Predict retention-aware V_{opt} using the model
- Reduces RBER on average by **51.9%**
(based on our characterization data)

Impact on System Reliability

— Baseline — State-of-the-art — LaVAR
— LaVAR + LI-RAID — This Work



LaVAR, LI-RAID, and ReMAR improve flash lifetime or reduce ECC overhead significantly

Error Mitigation Techniques Summary

- **LaVAR: Layer Variation Aware Reading**
 - Learn a V_{opt} offset for each layer and apply *layer-aware V_{opt}*
- **LI-RAID: Layer-Interleaved RAID**
 - Group flash pages on *less reliable layers* with pages on *more reliable layers*
 - Group *MSB pages* with *LSB pages*
- **ReMAR: Retention Model Aware Reading**
 - Learn retention loss model and apply *retention-aware V_{opt}*
- **Benefits:**
 - Improve flash lifetime by **1.85×** or reduce ECC overhead by **78.9%**
- **ReNAC (in paper):** Reread a failed page using V_{opt} based on the *retention interference* induced by neighbor cell

Agenda

- Background & Introduction
- Contribution 1: Characterize real 3D NAND flash chips
- Contribution 2: Model RBER and threshold voltage
- Contribution 3: Mitigate 3D NAND flash errors
- **Conclusion**

Conclusion

- Problem: 3D NAND error characteristics are **not well studied**
- Goal: *Understand & mitigate* 3D NAND errors to improve lifetime
- **Contribution 1: Characterize** real 3D NAND flash chips
 - **Process variation:** *21× error rate difference across layers*
 - **Early retention loss:** *Error rate increases by 10× after 3 hours*
 - **Retention interference:** *Not observed before in planar NAND*
- **Contribution 2: Model** RBER and threshold voltage
 - *RBER (raw bit error rate) variation model*
 - *Retention loss model*
- **Contribution 3: Mitigate** 3D NAND flash errors
 - **LaVAR:** *Layer Variation Aware Reading*
 - **LI-RAID:** *Layer-Interleaved RAID*
 - **ReMAR:** *Retention Model Aware Reading*
 - *Improve flash lifetime by 1.85× or reduce ECC overhead by 78.9%*

Modeling and Mitigating Early Retention Loss and Process Variation in 3D Flash

Saugata Ghose

Carnegie Mellon University

Download our SIGMETRICS 2018 Paper at

**[http://ece.cmu.edu/~saugatag/papers/
18sigmetrics_3dflash.pdf](http://ece.cmu.edu/~saugatag/papers/18sigmetrics_3dflash.pdf)**



References to Papers and Talks

Our FMS Talks and Posters

- FMS 2019
 - Saugata Ghose, **Modeling and Mitigating Early Retention Loss and Process Variation in 3D Flash**
 - Saugata Ghose, **Enabling Fairness and Enhancing Performance in Modern NVMe Solid State Drives**
- FMS 2018
 - Yixin Luo, **HeatWatch: Exploiting 3D NAND Self-Recovery and Temperature Effects**
 - Saugata Ghose, **Enabling Realistic Studies of Modern Multi-Queue SSD Devices**
- FMS 2017
 - Aya Fukami, **Improving Chip-Off Forensic Analysis for NAND Flash**
 - Saugata Ghose, **Vulnerabilities in MLC NAND Flash Memory Programming**
- FMS 2016
 - Onur Mutlu, [ThyNVM: Software-Transparent Crash Consistency for Persistent Memory](#)
 - Onur Mutlu, [Large-Scale Study of In-the-Field Flash Failures](#)
 - Yixin Luo, [Practical Threshold Voltage Distribution Modeling](#)
 - Saugata Ghose, [Write-hotness Aware Retention Management](#)
- FMS 2015
 - Onur Mutlu, [Read Disturb Errors in MLC NAND Flash Memory](#)
 - Yixin Luo, [Data Retention in MLC NAND Flash Memory](#)
- FMS 2014
 - Onur Mutlu, [Error Analysis and Management for MLC NAND Flash Memory](#)

Our Flash Memory Works (I)

- **Summary of our work in NAND flash memory**
 - Yu Cai, Saugata Ghose, Erich F. Haratsch, Yixin Luo, and Onur Mutlu, [Error Characterization, Mitigation, and Recovery in Flash Memory Based Solid-State Drives](#), *Proceedings of the IEEE*, Sept. 2017.
- **Overall flash error analysis**
 - Yu Cai, Erich F. Haratsch, Onur Mutlu, and Ken Mai, [Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis](#), DATE 2012.
 - Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Adrian Cristal, Osman Unsal, and Ken Mai, [Error Analysis and Retention-Aware Error Management for NAND Flash Memory](#), ITJ 2013.
 - Yixin Luo, Saugata Ghose, Yu Cai, Erich F. Haratsch, and Onur Mutlu, [Enabling Accurate and Practical Online Flash Channel Modeling for Modern MLC NAND Flash Memory](#), *IEEE JSAC*, Sept. 2016.

Our Flash Memory Works (II)

- **3D NAND flash memory error analysis**
 - Yixin Luo, Saugata Ghose, Yu Cai, Erich F. Haratsch, and Onur Mutlu, [Improving 3D NAND Flash Memory Lifetime by Tolerating Early Retention Loss and Process Variation](#), SIGMETRICS 2018.
 - Yixin Luo, Saugata Ghose, Yu Cai, Erich F. Haratsch, and Onur Mutlu, [HeatWatch: Improving 3D NAND Flash Memory Device Reliability by Exploiting Self-Recovery and Temperature-Awareness](#), HPCA 2018.
- **Multi-queue SSDs**
 - Arash Tavakkol, Juan Gomez-Luna, Mohammad Sadrosadati, Saugata Ghose, and Onur Mutlu, [MQSim: A Framework for Enabling Realistic Studies of Modern Multi-Queue SSD Devices](#), FAST 2018.
 - Arash Tavakkol, Mohammad Sadrosadati, Saugata Ghose, Jeremie Kim, Yixin Luo, Yaohua Wang, Nika Mansouri Ghiasi, Lois Orosa, Juan G. Luna and Onur Mutlu, [FLIN: Enabling Fairness and Enhancing Performance in Modern NVMe Solid State Drives](#), ISCA 2018.

Our Flash Memory Works (III)

- **Flash-based SSD prototyping and testing platform**
 - Yu Cai, Erich F. Haratsh, Mark McCartney, Ken Mai, [FPGA-based solid-state drive prototyping platform](#), FCCM 2011.
- **Retention noise study and management**
 - Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Adrian Cristal, Osman Unsal, and Ken Mai, [Flash Correct-and-Refresh: Retention-Aware Error Management for Increased Flash Memory Lifetime](#), ICCD 2012.
 - Yu Cai, Yixin Luo, Erich F. Haratsch, Ken Mai, and Onur Mutlu, [Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery](#), HPCA 2015.
 - Yixin Luo, Yu Cai, Saugata Ghose, Jongmoo Choi, and Onur Mutlu, [WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management](#), MSST 2015.
 - Aya Fukami, Saugata Ghose, Yixin Luo, Yu Cai, and Onur Mutlu, [Improving the Reliability of Chip-Off Forensic Analysis of NAND Flash Memory Devices](#), *Digital Investigation*, Mar. 2017.

Our Flash Memory Works (IV)

- **Program and erase noise study**
 - Yu Cai, Erich F. Haratsch, Onur Mutlu, and Ken Mai, [Threshold Voltage Distribution in MLC NAND Flash Memory: Characterization, Analysis and Modeling](#), DATE 2013.
 - Y. Cai, S. Ghose, Y. Luo, K. Mai, O. Mutlu, and E. F. Haratsch, [Vulnerabilities in MLC NAND Flash Memory Programming: Experimental Analysis, Exploits, and Mitigation Techniques](#), HPCA 2017.
- **Cell-to-cell interference characterization and tolerance**
 - Yu Cai, Onur Mutlu, Erich F. Haratsch, and Ken Mai, [Program Interference in MLC NAND Flash Memory: Characterization, Modeling, and Mitigation](#), ICCD 2013.
 - Yu Cai, Gulay Yalcin, Onur Mutlu, Erich F. Haratsch, Osman Unsal, Adrian Cristal, and Ken Mai, [Neighbor-Cell Assisted Error Correction for MLC NAND Flash Memories](#), SIGMETRICS 2014.

Our Flash Memory Works (V)

- **Read disturb noise study**

- Yu Cai, Yixin Luo, Saugata Ghose, Erich F. Haratsch, Ken Mai, and Onur Mutlu, [Read Disturb Errors in MLC NAND Flash Memory: Characterization and Mitigation](#), DSN 2015.

- **Flash errors in the field**

- Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu, [A Large-Scale Study of Flash Memory Errors in the Field](#), SIGMETRICS 2015.

- **Persistent memory**

- Jinglei Ren, Jishen Zhao, Samira Khan, Jongmoo Choi, Yongwei Wu, and Onur Mutlu, [ThyNVM: Enabling Software-Transparent Crash Consistency in Persistent Memory Systems](#), MICRO 2015.

Referenced Papers and Talks

- All are available at
 - <https://safari.ethz.ch/publications/>
 - <https://www.ece.cmu.edu/~safari/talks.html>
- And, many other previous works on
 - Challenges and opportunities in memory
 - NAND flash memory errors and management
 - Phase change memory as DRAM replacement
 - STT-MRAM as DRAM replacement
 - Taking advantage of persistence in memory
 - Hybrid DRAM + NVM systems
 - NVM design and architecture