Data Retention in MLC NAND Flash Memory: Characterization, Optimization, and Recovery

Yu Cai, Yixin Luo, Erich F. Haratsch, Ken Mai, Onur Mutlu
Carnegie Mellon University, LSI Corporation

2015 IEEE 21st International Symposium on High Performance Computing Architecture (HPCA 2015)

Why flash memory?

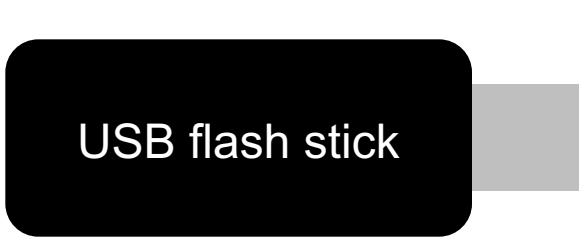

Solid State Drive

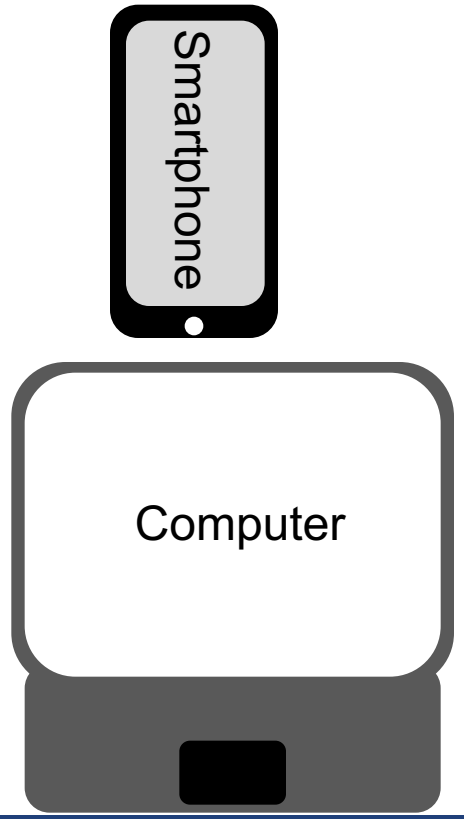

Devices using Flash Memory

NAND Flash Memory

- We **erase** in blocks consisting of multiple pages.
- We **program** and **read** in pages.

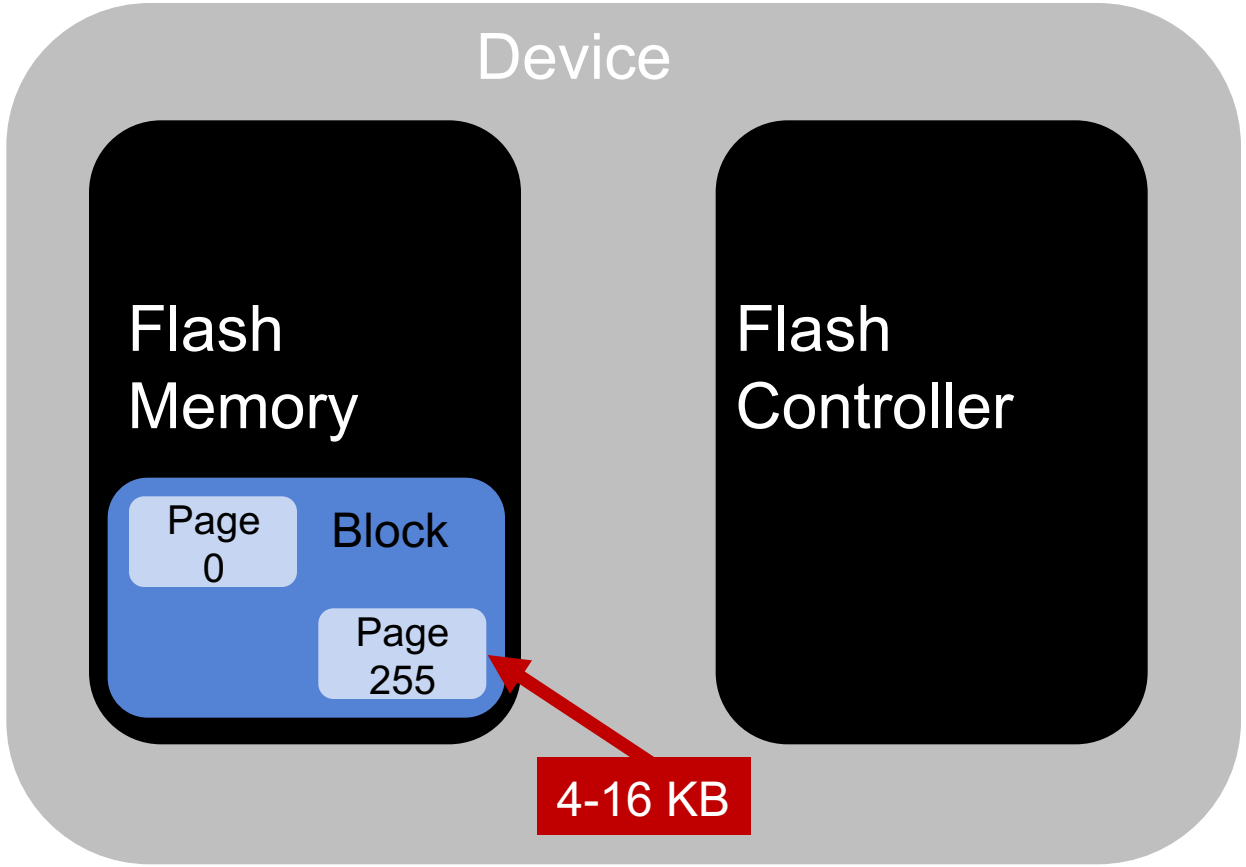

NAND Flash Memory

- We **erase** in blocks consisting of multiple pages.
- We **program** and **read** in pages.

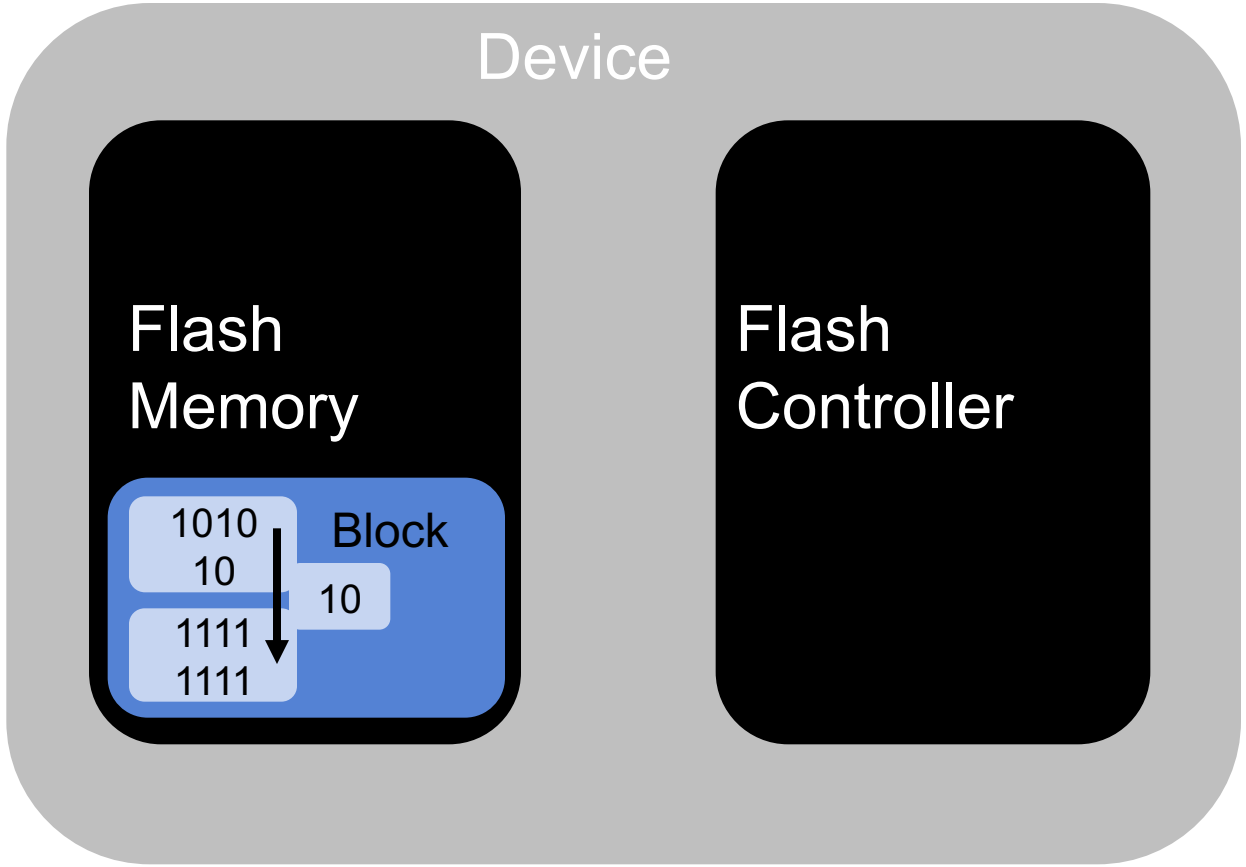

NAND Flash Memory

- We **erase** in blocks consisting of multiple pages.
- We **program** and **read** in pages.

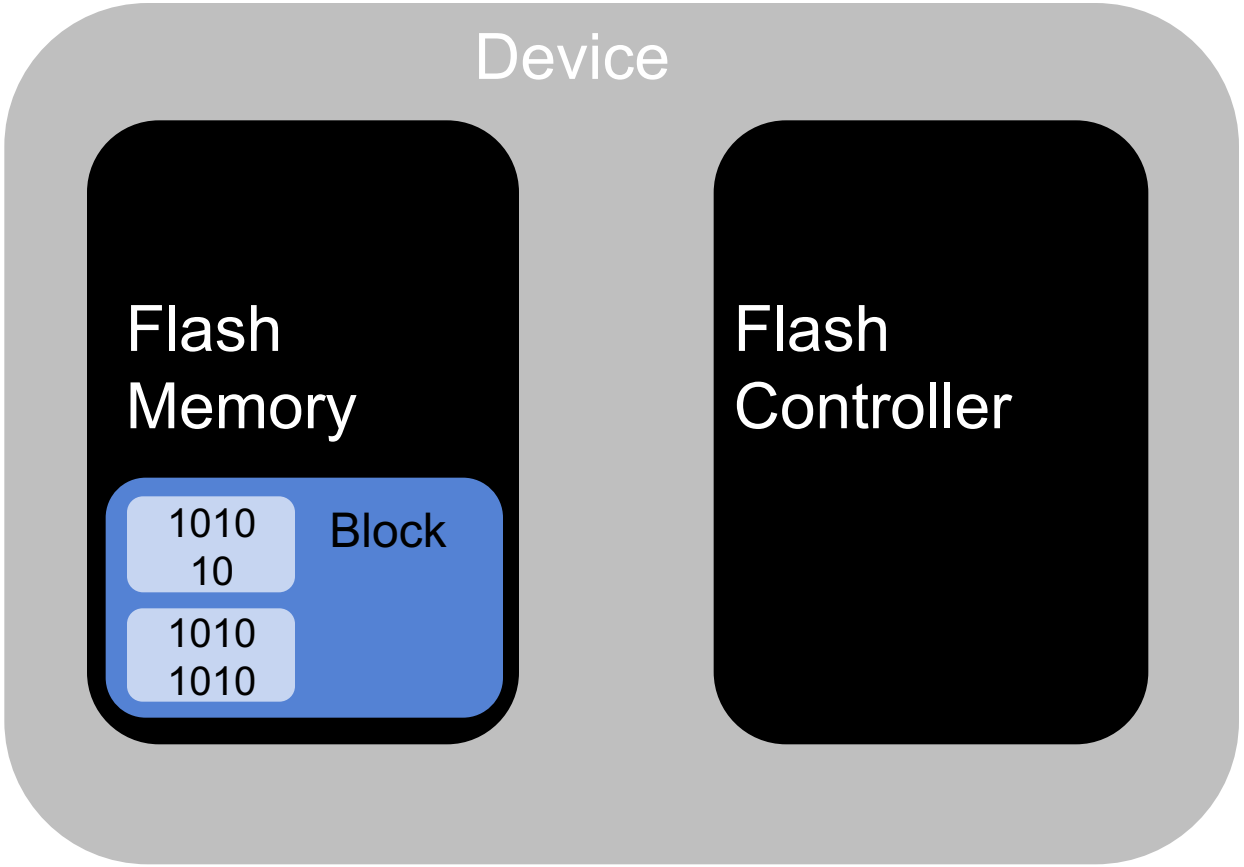

NAND Flash Memory

- We **erase** in blocks consisting of multiple pages.
- We **program** and **read** in pages.
- After programming we **erase old** copied pages.
- Flash controller **manages** operations.
- ECC controller **corrects** data.

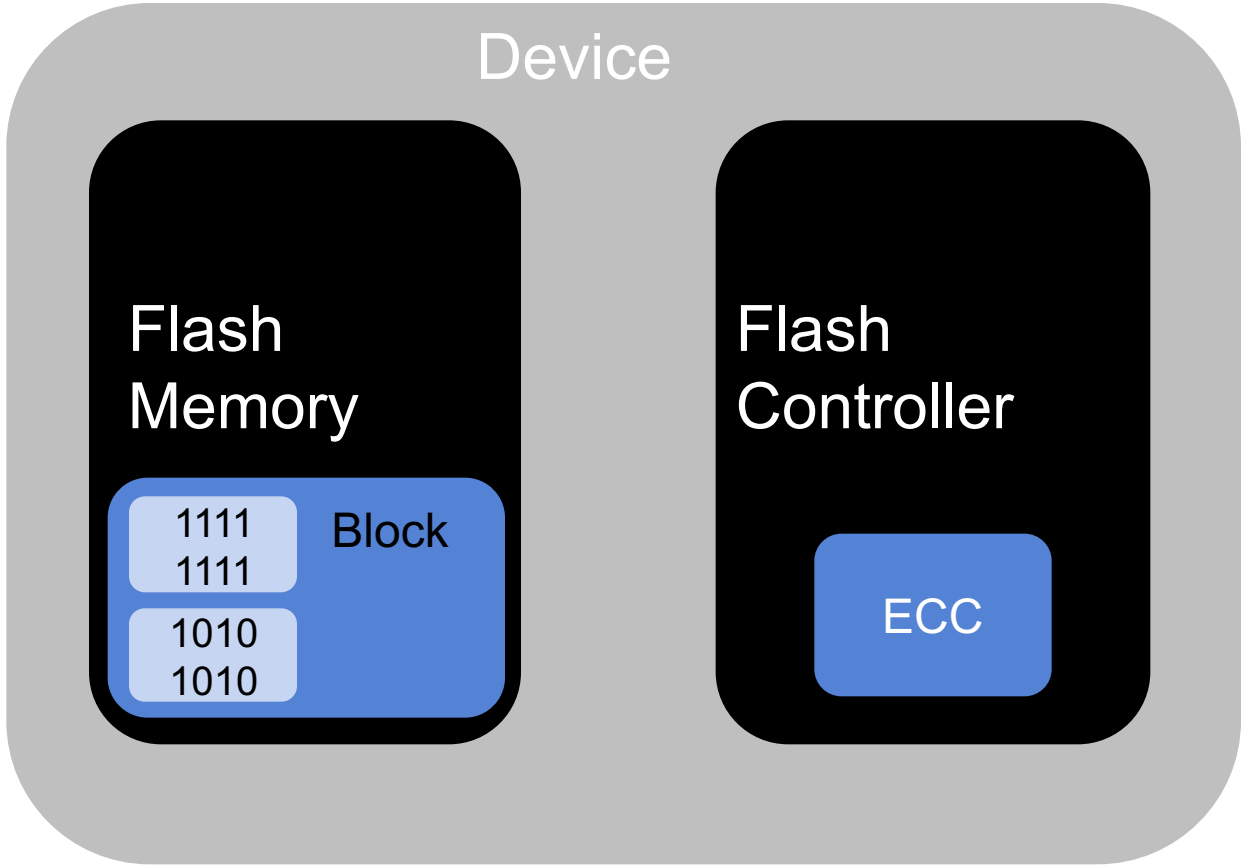

Flash Memory Cells

- A cell stores charge as electrons in the floating gate.
- We **program** cells by applying a high **positive** voltage to our control gate.
- The positive charge **attracts** electrons through the tunnel oxide from the substrate.

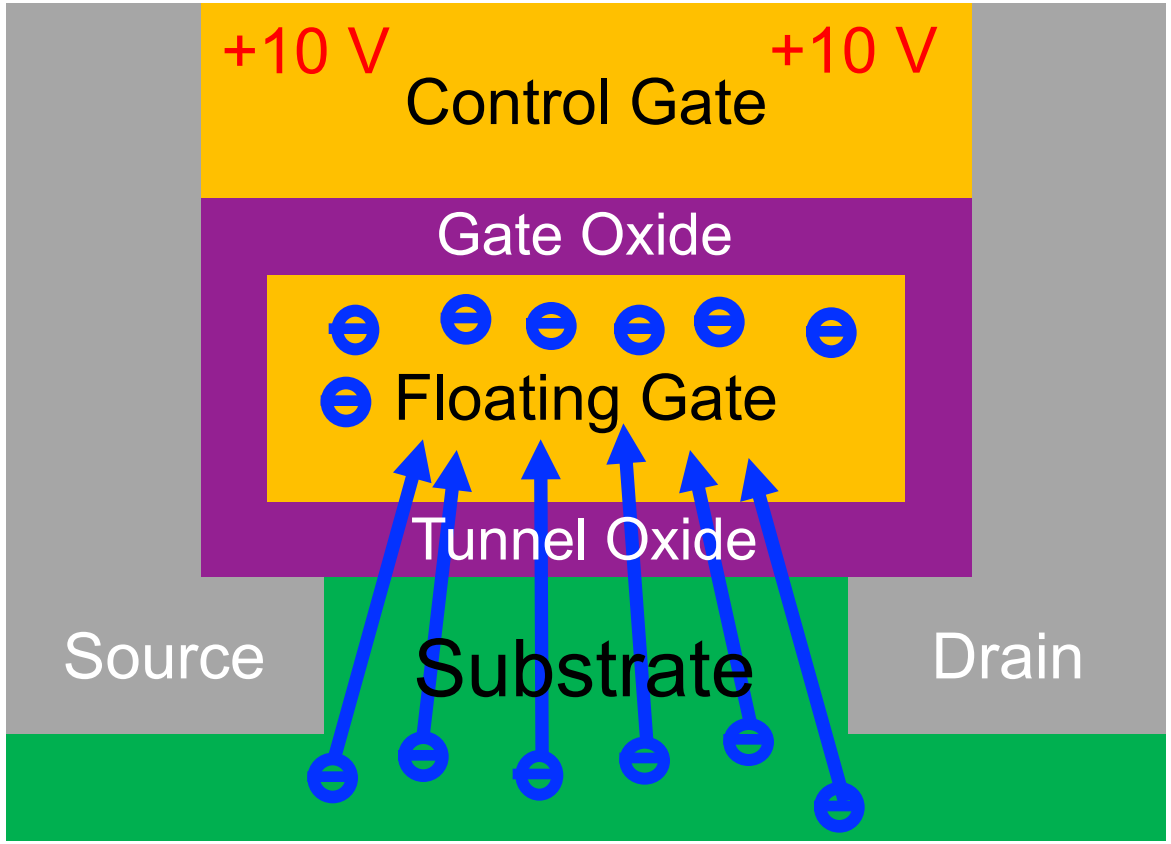

Erasing Cells

- We **erase** cells by applying a high **negative** voltage.

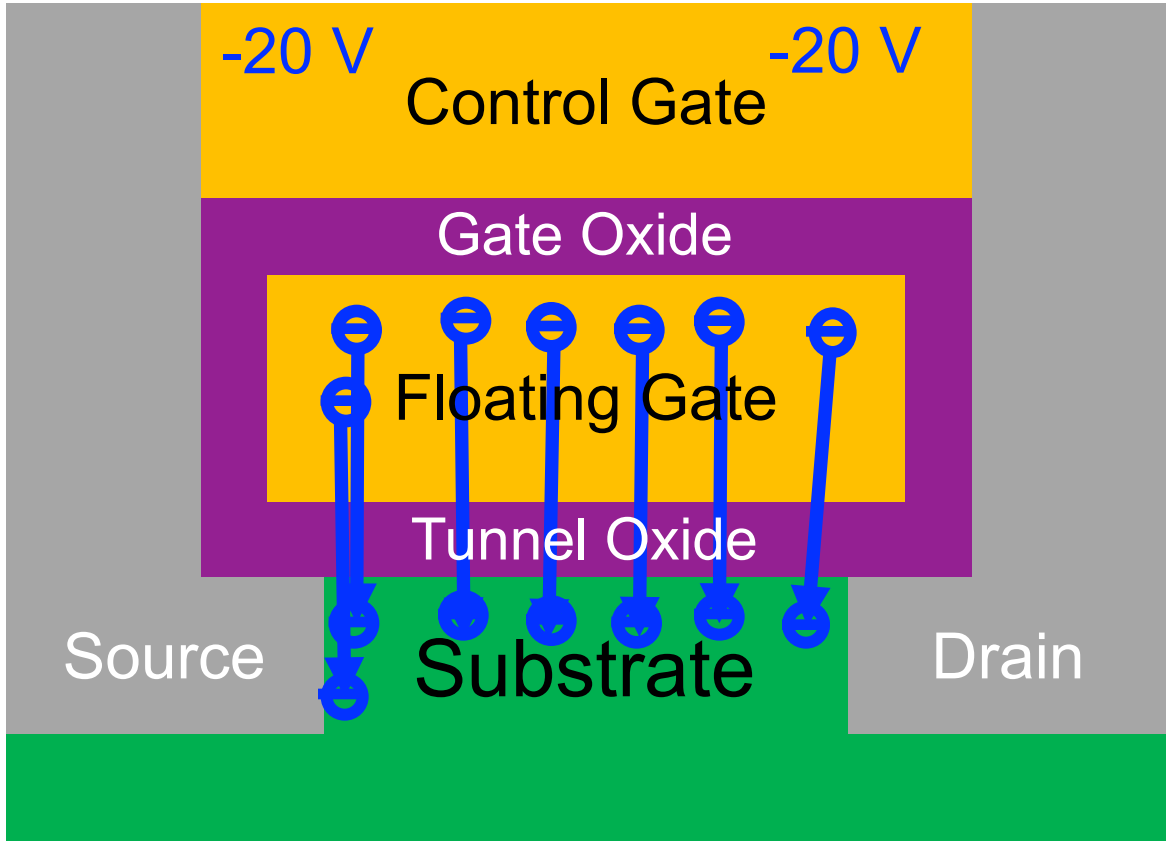

Trap Assisted Tunnelling

- Repeated **program**/**erase** cycles **trap** electrons in the tunnel oxide.
- An **electric field** is created by the trapped charge.
- Charge from the floating gate **leaks** to the substrate

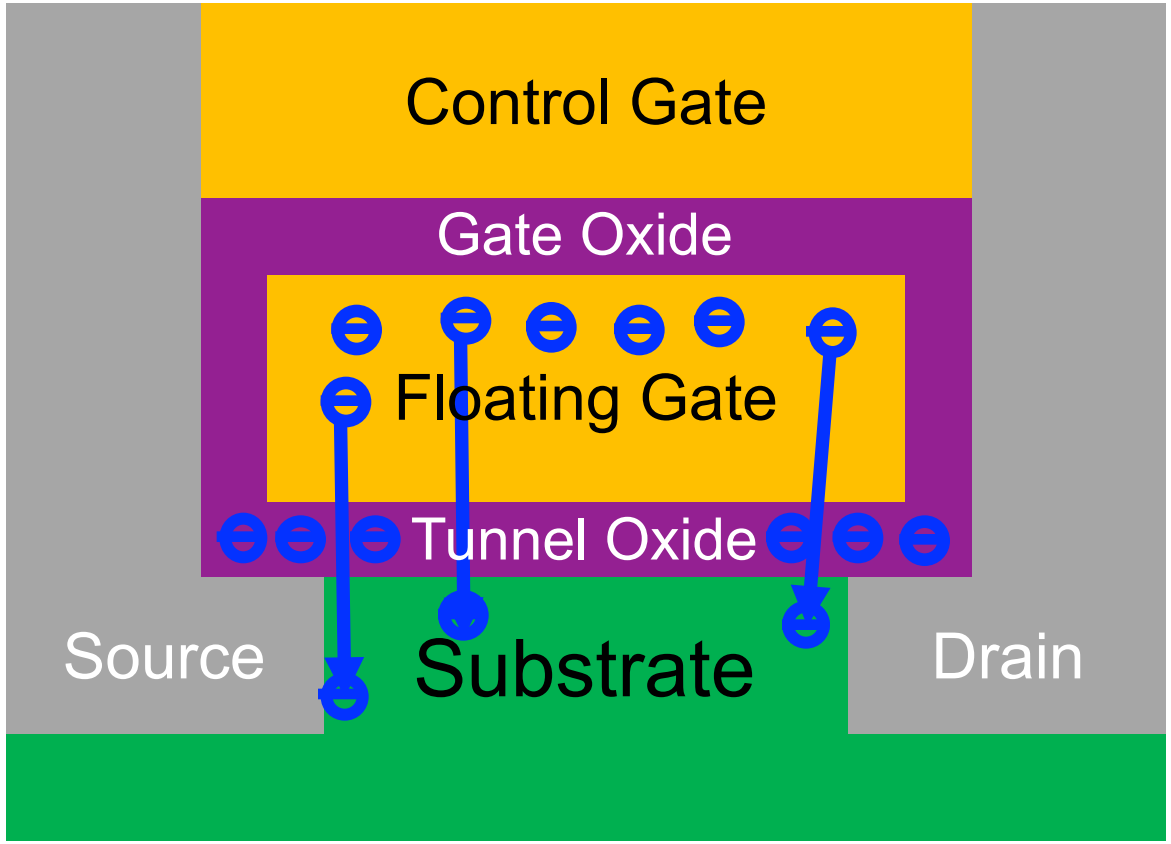

Charge De-Trapping

- Trapped charge **leaves** tunnel oxide.
- **Increases** floating gate **charge**.

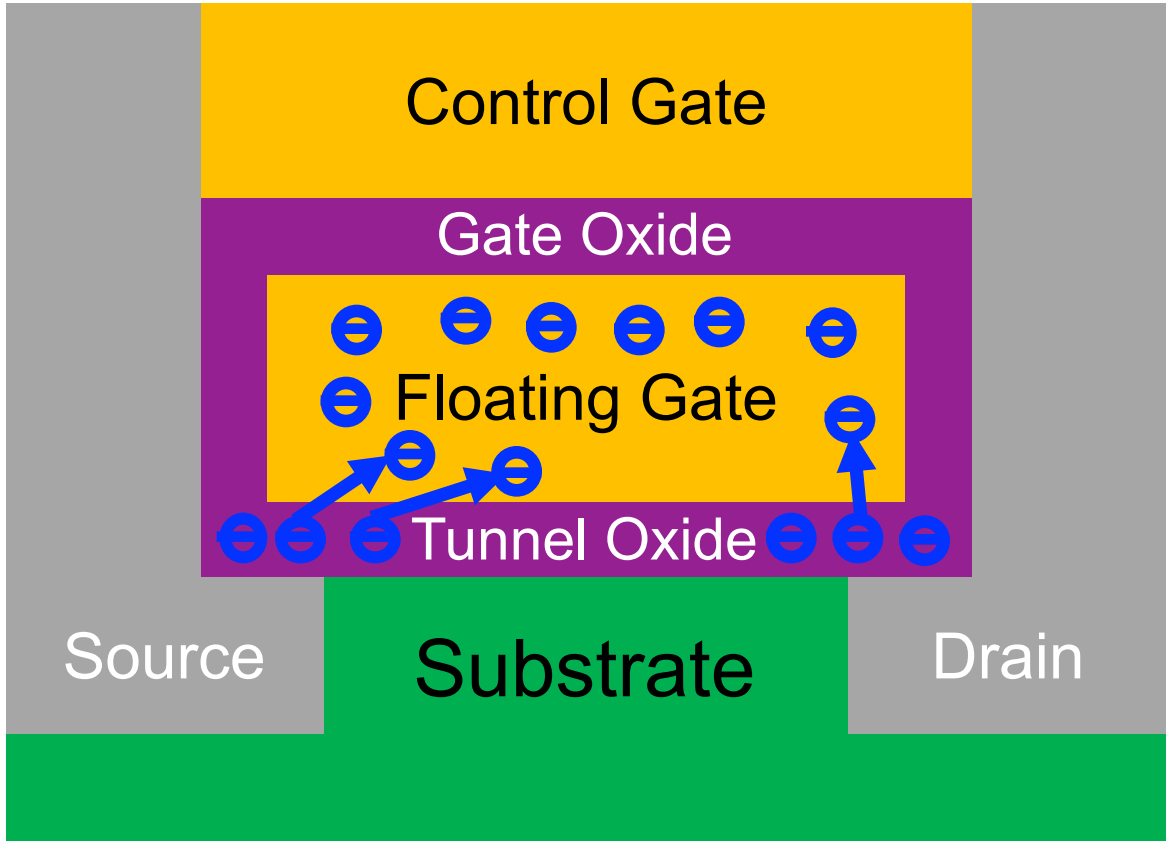

Charge De-Trapping

- Trapped charge **leaves** tunnel oxide.
- Increases floating gate charge.
- Or **decreases** if de-trapping towards the **substrate**.

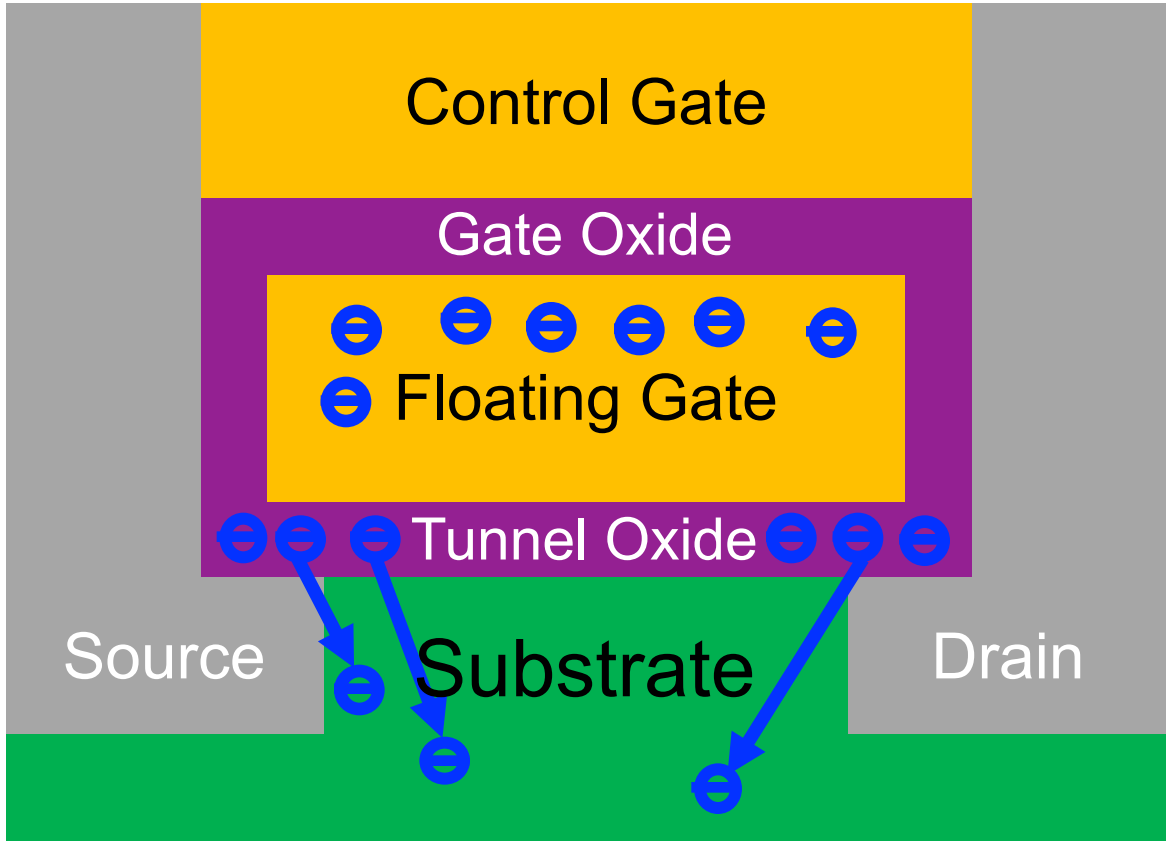

Charge De-Trapping

- Trapped charge **leaves** tunnel oxide.
- Increases floating gate charge.
- Or **decreases** if de-trapping towards the substrate.
- Leaves a **positive** charge that attracts charge from the floating gate.

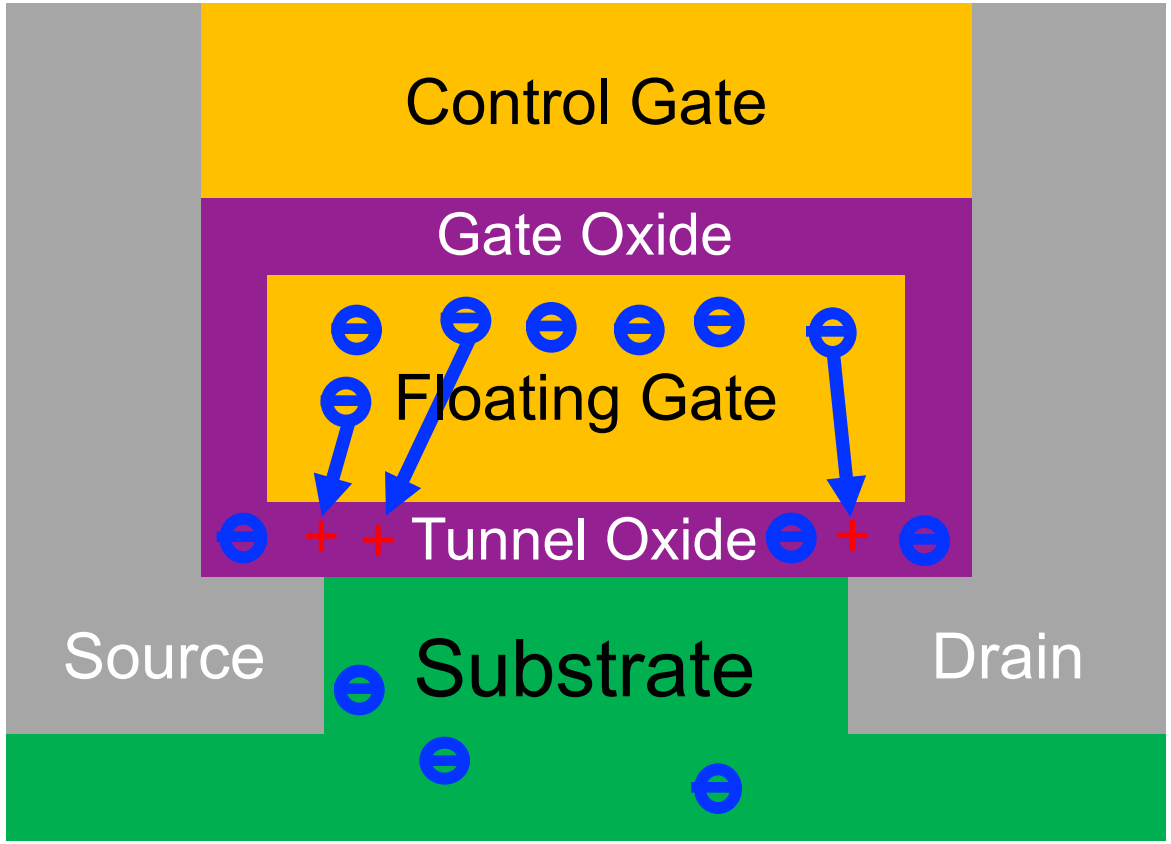

Charge De-Trapping

- Trapped charge **leaves** tunnel oxide.
- Increases floating gate charge.
- Or **decreases** if de-trapping towards the substrate.
- Leaves a **positive** charge that attracts charge from the floating gate.

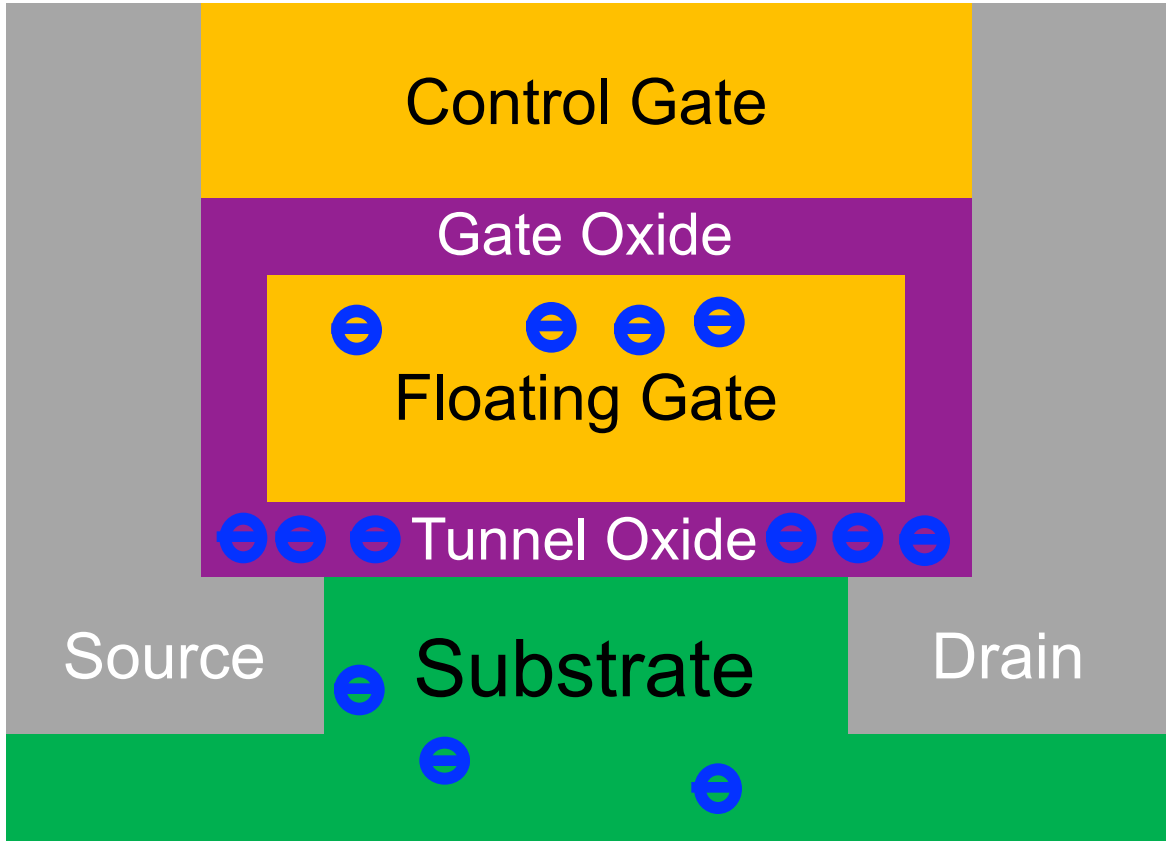

Retention Loss

- Use charge as **indicator** for **bit** values.
- Assign **0** to a **high** charge and **1** to a **low** charge.
- **Read reference voltage** separates differently charge cells.
- Charge leaks over time caused by **trap assisted tunnelling** or **charge de-trapping**.
- Changed values introduce **retention errors**.

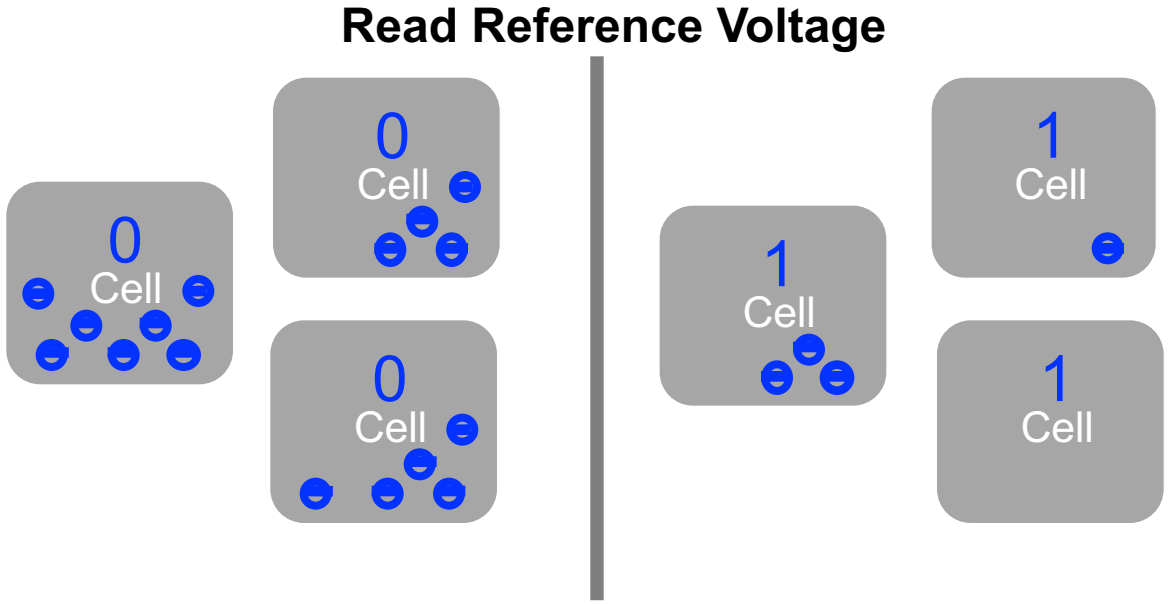

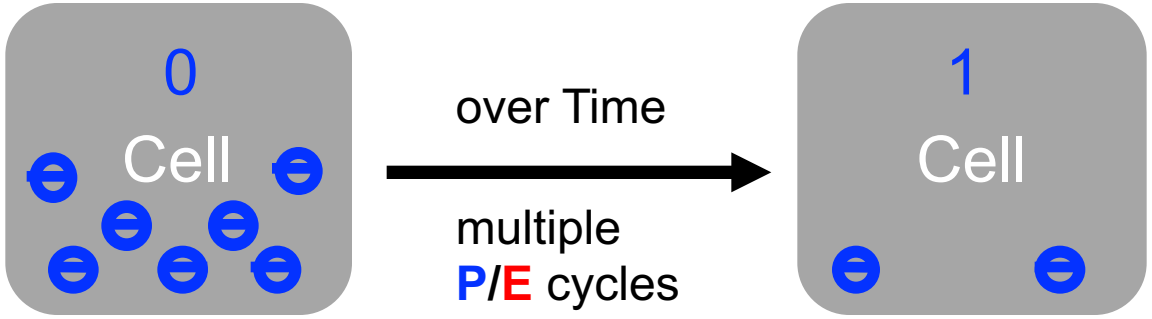

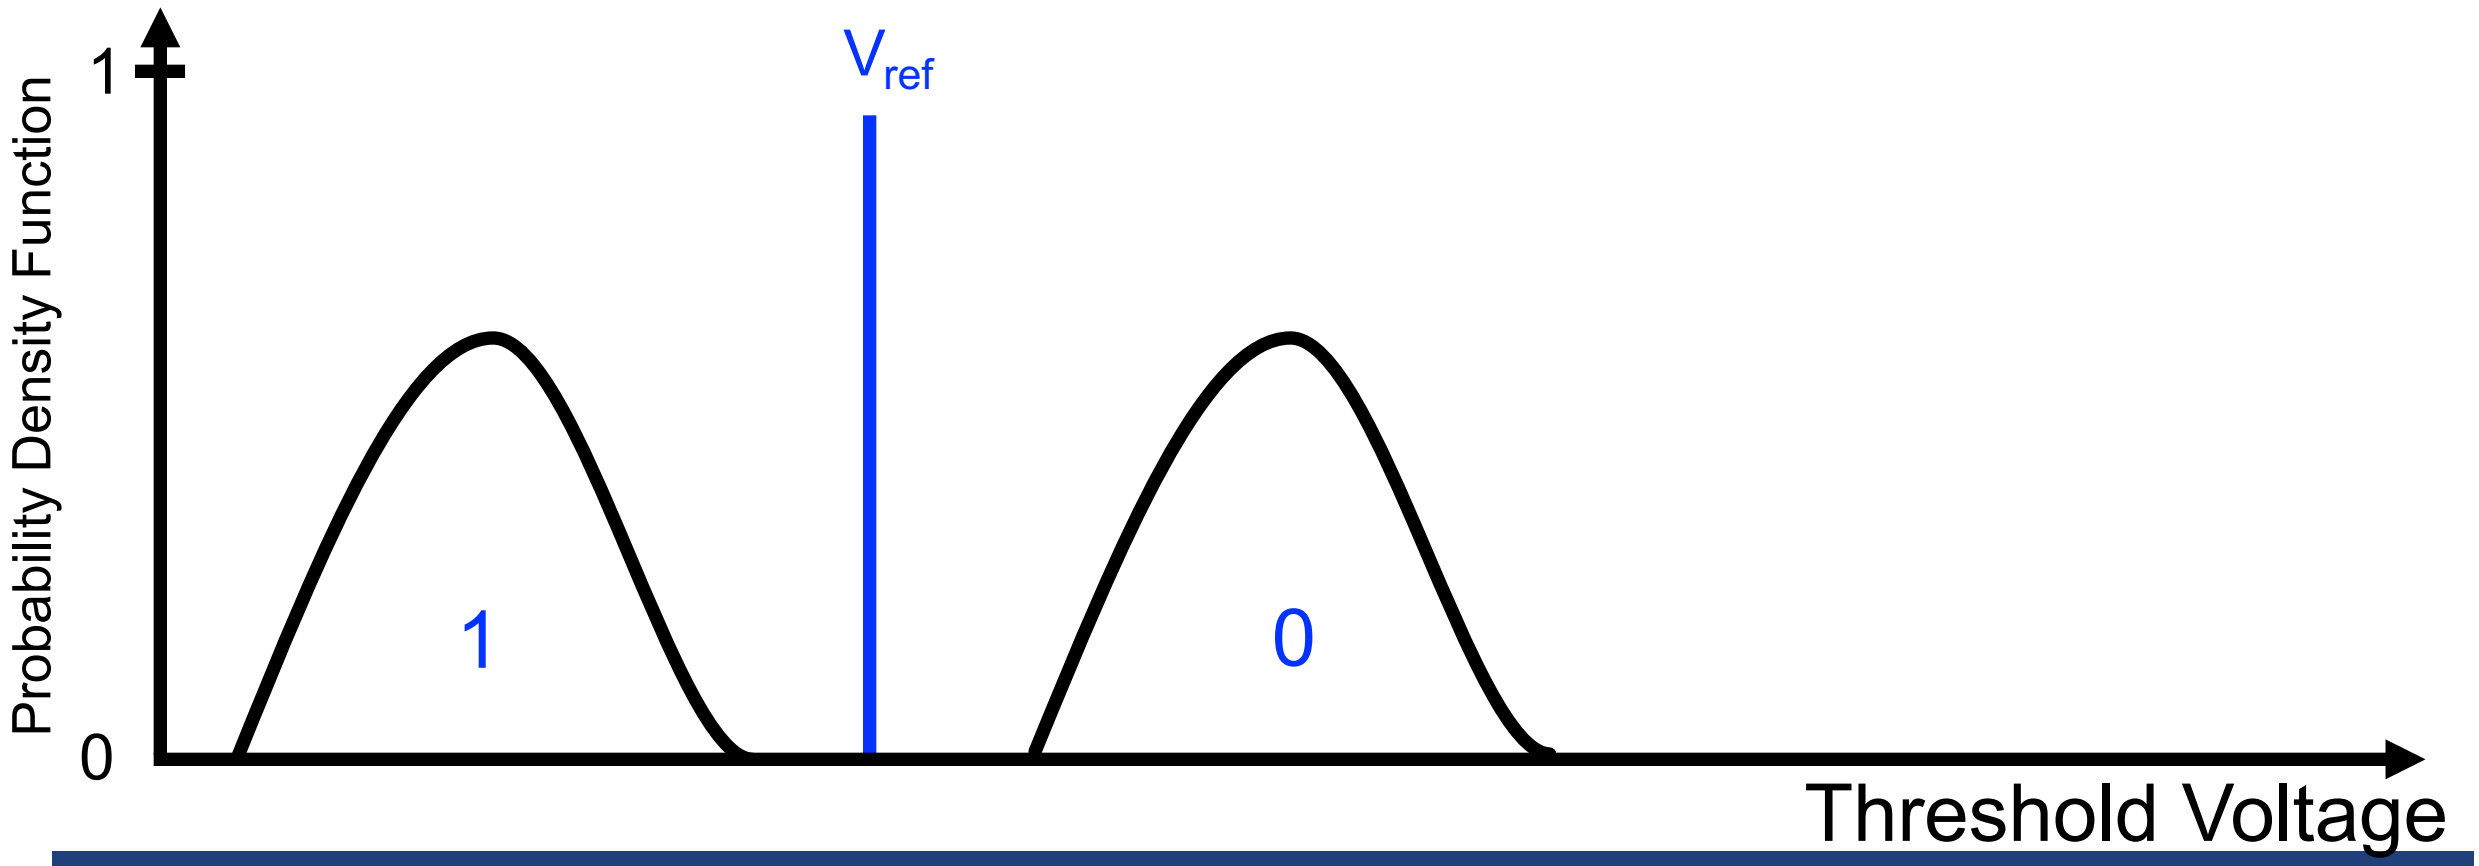

Threshold Voltage Distribution

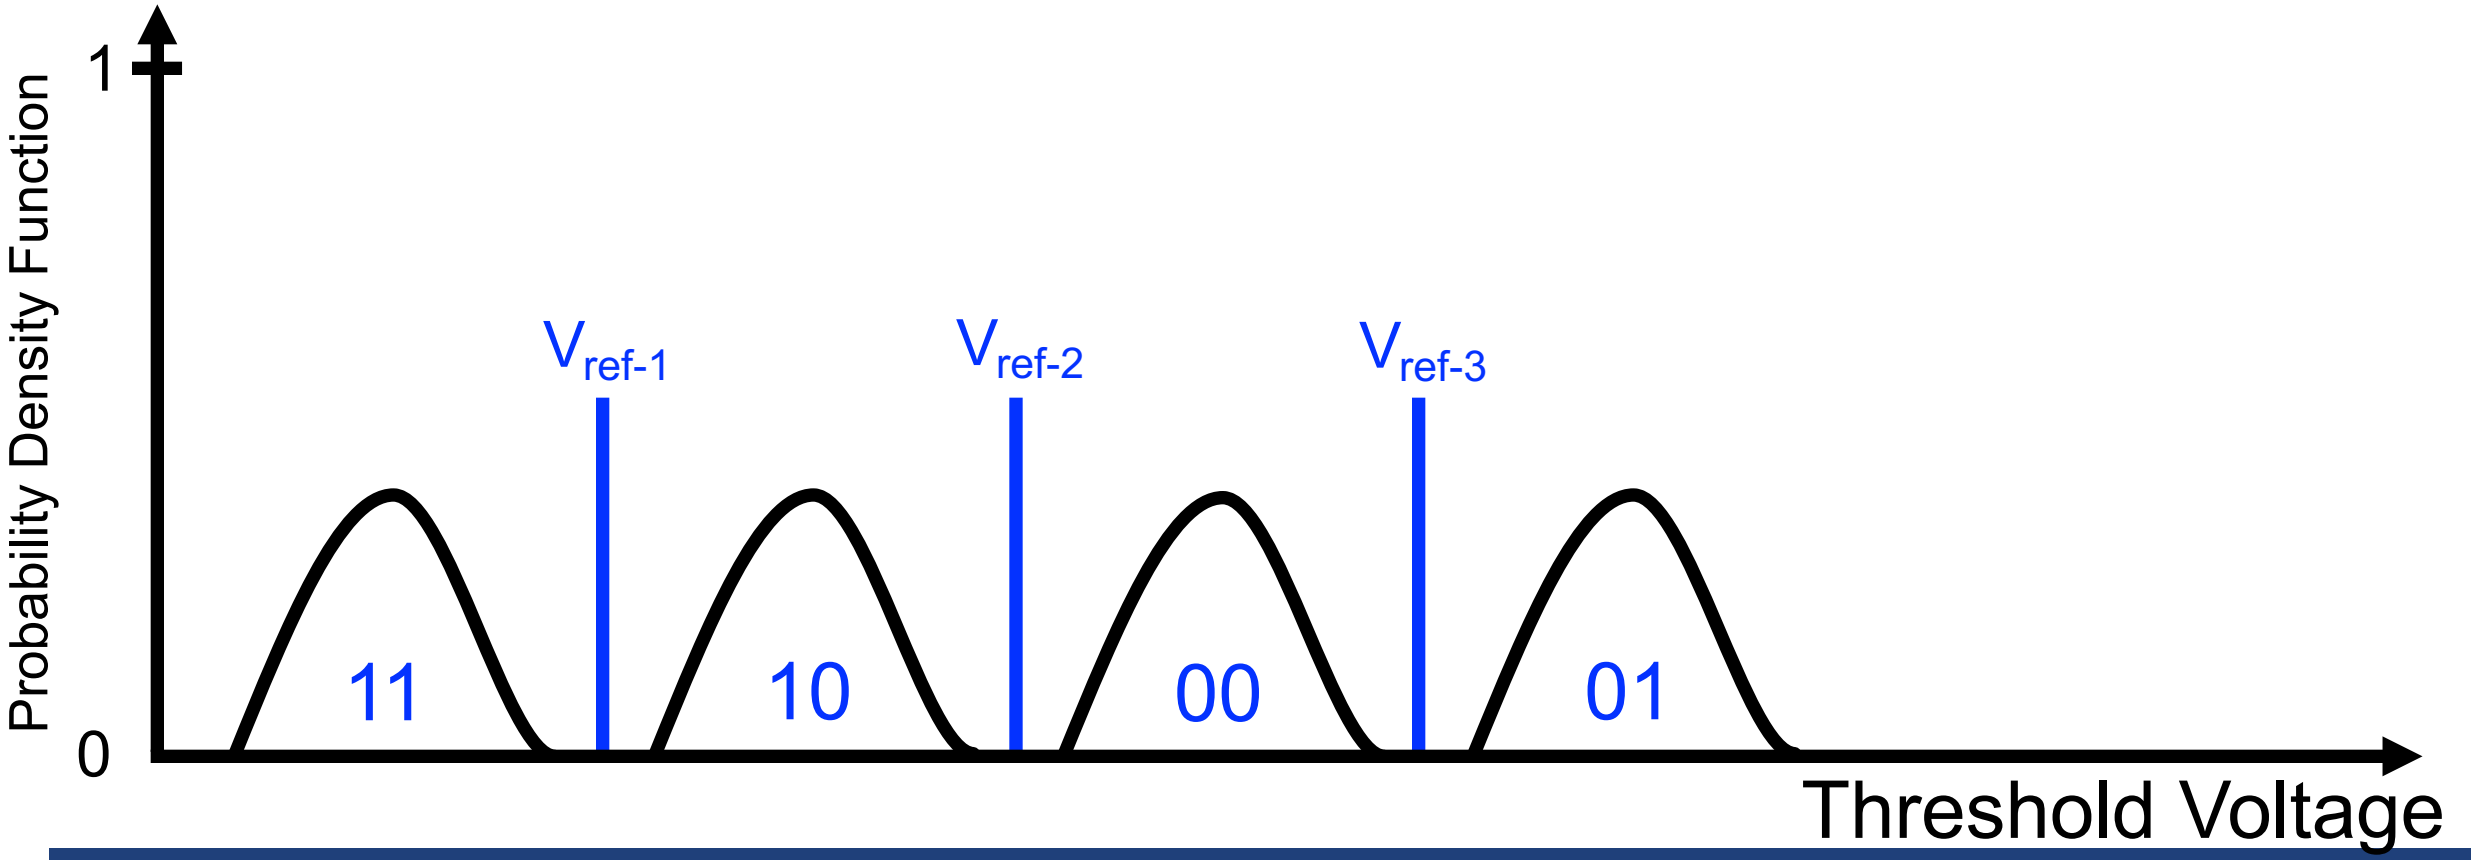

Threshold Voltage Distribution in Multi Level Cell Flash Memory

Read-Retry

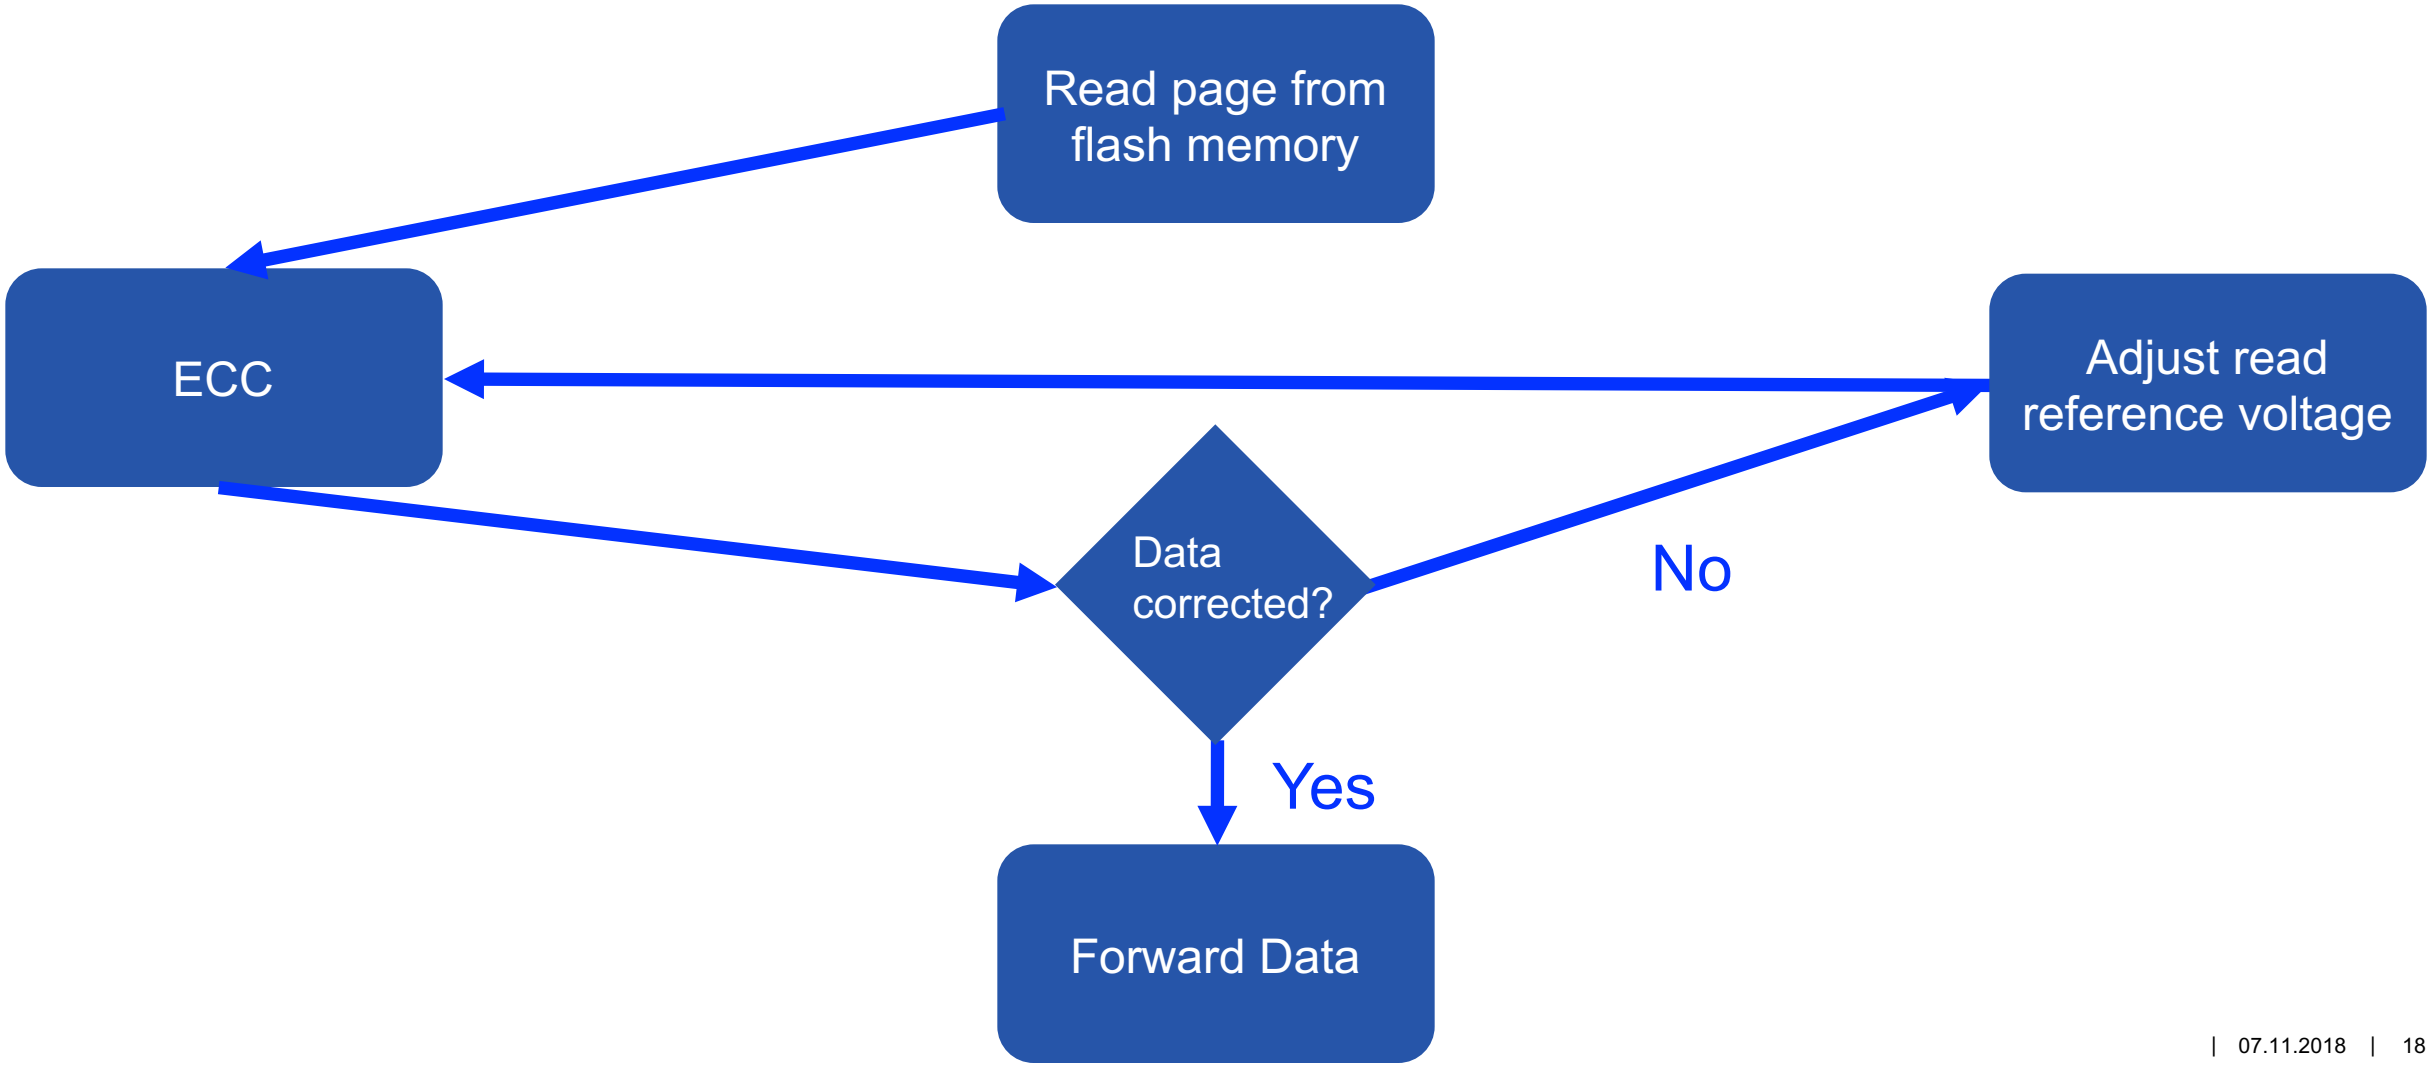
Read page from flash memory
ECC
Adjust read reference voltage
Data corrected?
No
Yes
Forward Data

Executive Summary

- **Problem**
  - **Density** of flash memory rises and diminishes **lifetime**.
  - Correcting errors increases **read latency**.
- **Goal**
  - Deepen understanding of **voltage threshold distributions** of flash memory.
  - **Improve** both lifetime and system performance.
  - **Recover** non-correctable data.
- **Method**
  - **Retention Optimized Reading**
    - **Improved Read-Retry**
  - **Retention Failure Recovery**
- **Result**
  - **Lifetime** improvement by 64%.
  - **Read latency** reduction by 70.4%.
  - **Raw bit error rate** drop by 50%.

Problem

- Multi level cell
  - **Higher error rate** due to smaller threshold windows.
- Lifetime
  - Retention errors:
    - **Limit** the time **flash memory** can be read from.
    - May lead to **loosing data**.
- Read Latency
  - Retention errors:
    - Introduce **overhead** by error correction codes.
    - Increase number of **read-retries**.

Goal

- Building a strong understanding, characterization, and analysis of **threshold voltage distribution** over retention age.
- Introduce a **dynamic technique** improving lifetime and read latency.
- Devise a new mechanism to **recover** non-correctable data.

FPGA-Based Flash Memory Testing Platform

- Different amounts of **program**/**erase** cycles for multiple groups of flash memory.
- Data of retention ages ranging from **0 to 40** days.
- All experiments were conducted under room temperature (**20°C**).

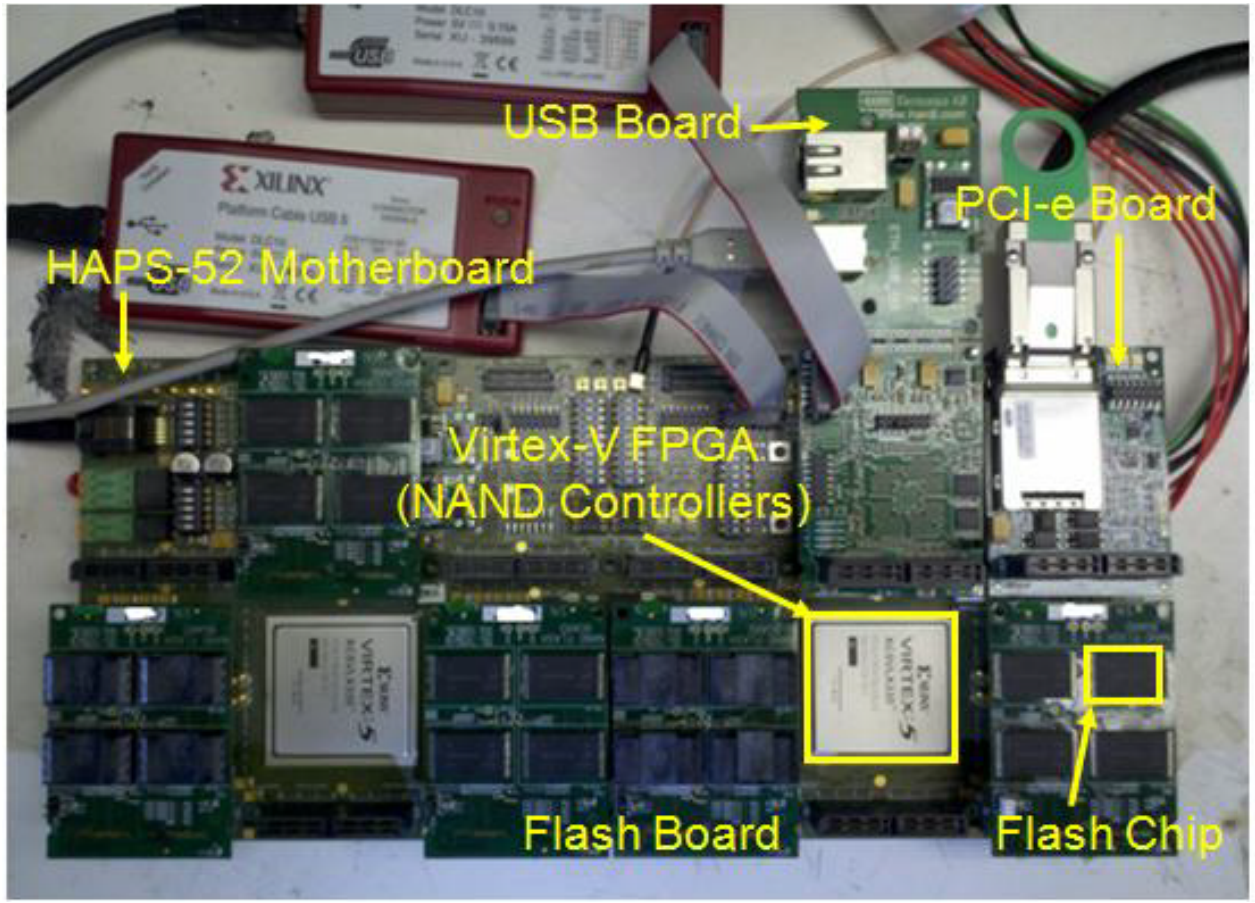

Source: Y. Cai et al., "FPGA-Based Solid-State Drive Prototyping Platform", FCCM 2011

Retention Optimized Reading

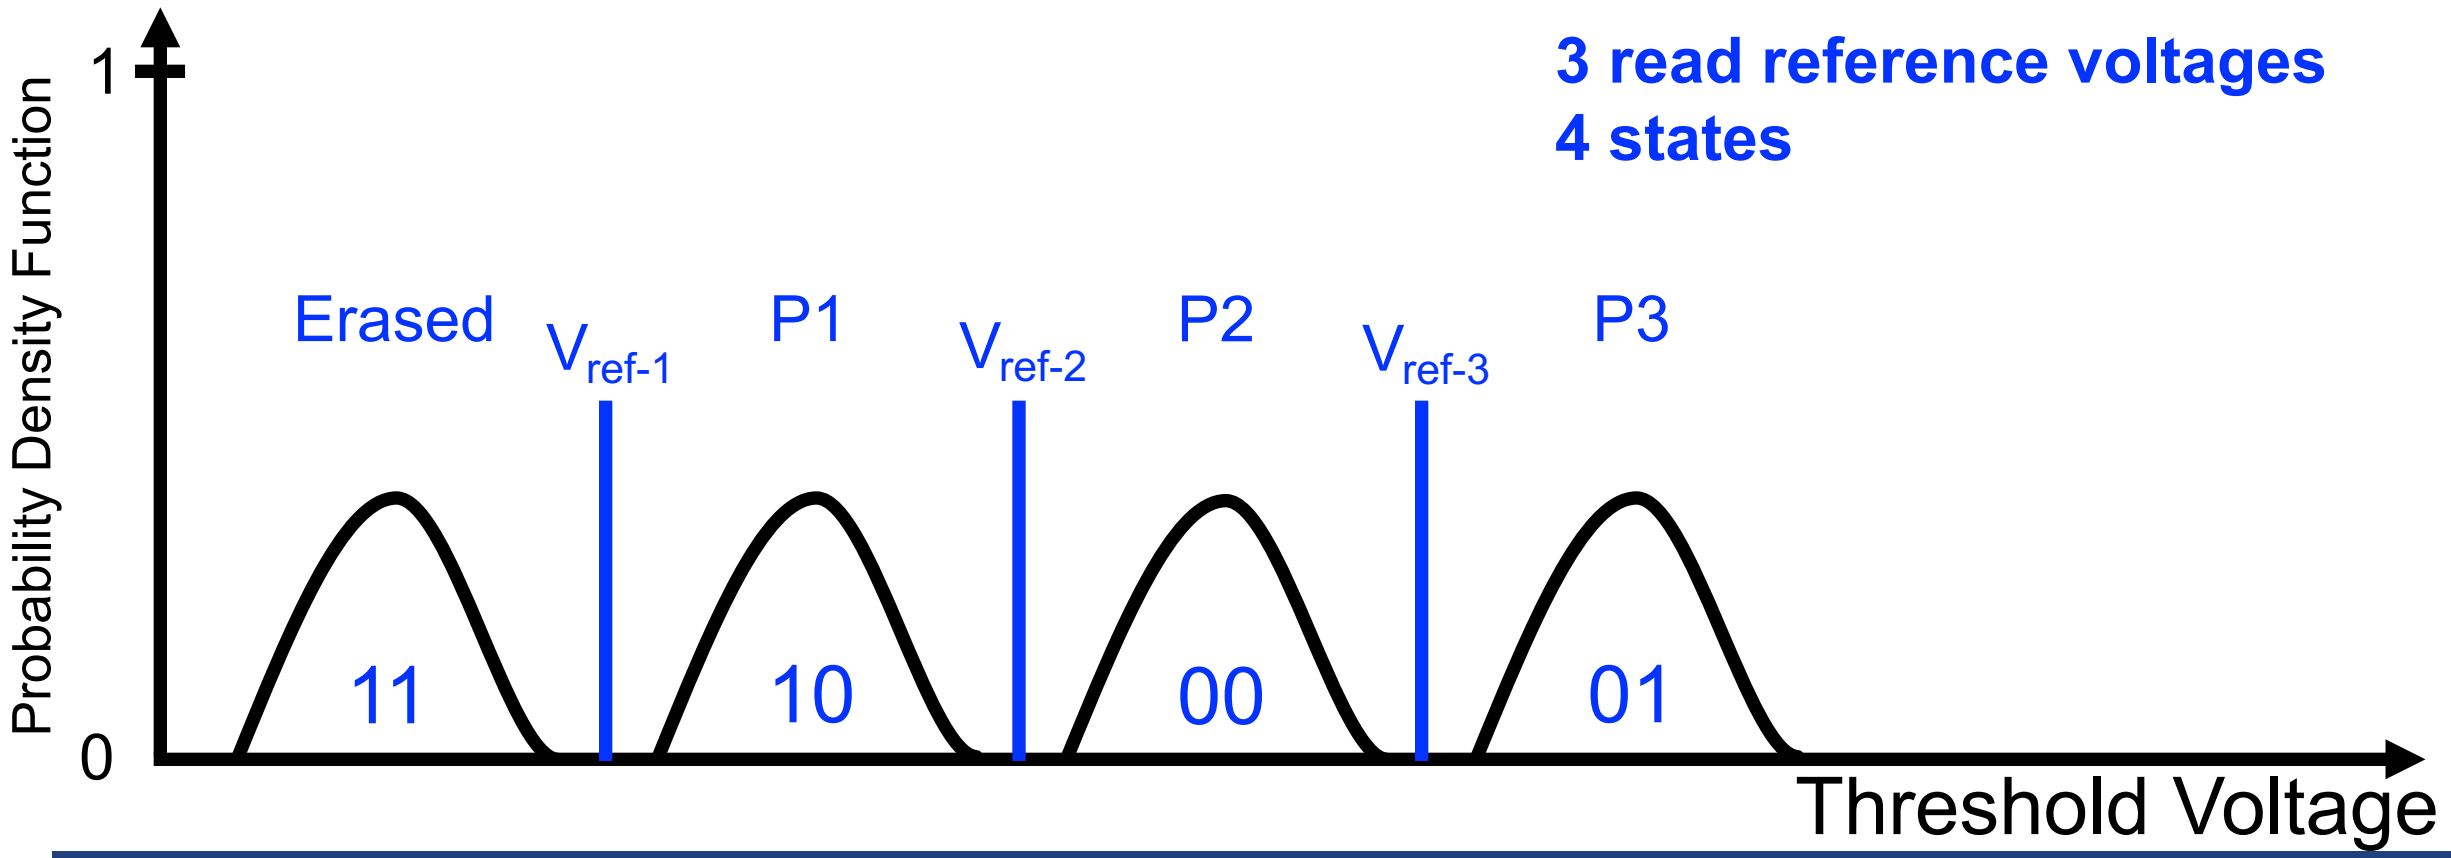

Threshold Voltage Distribution over Time

Source: Slides adapted from Data Retention in MLC NAND Flash Memory… Yixin Luo

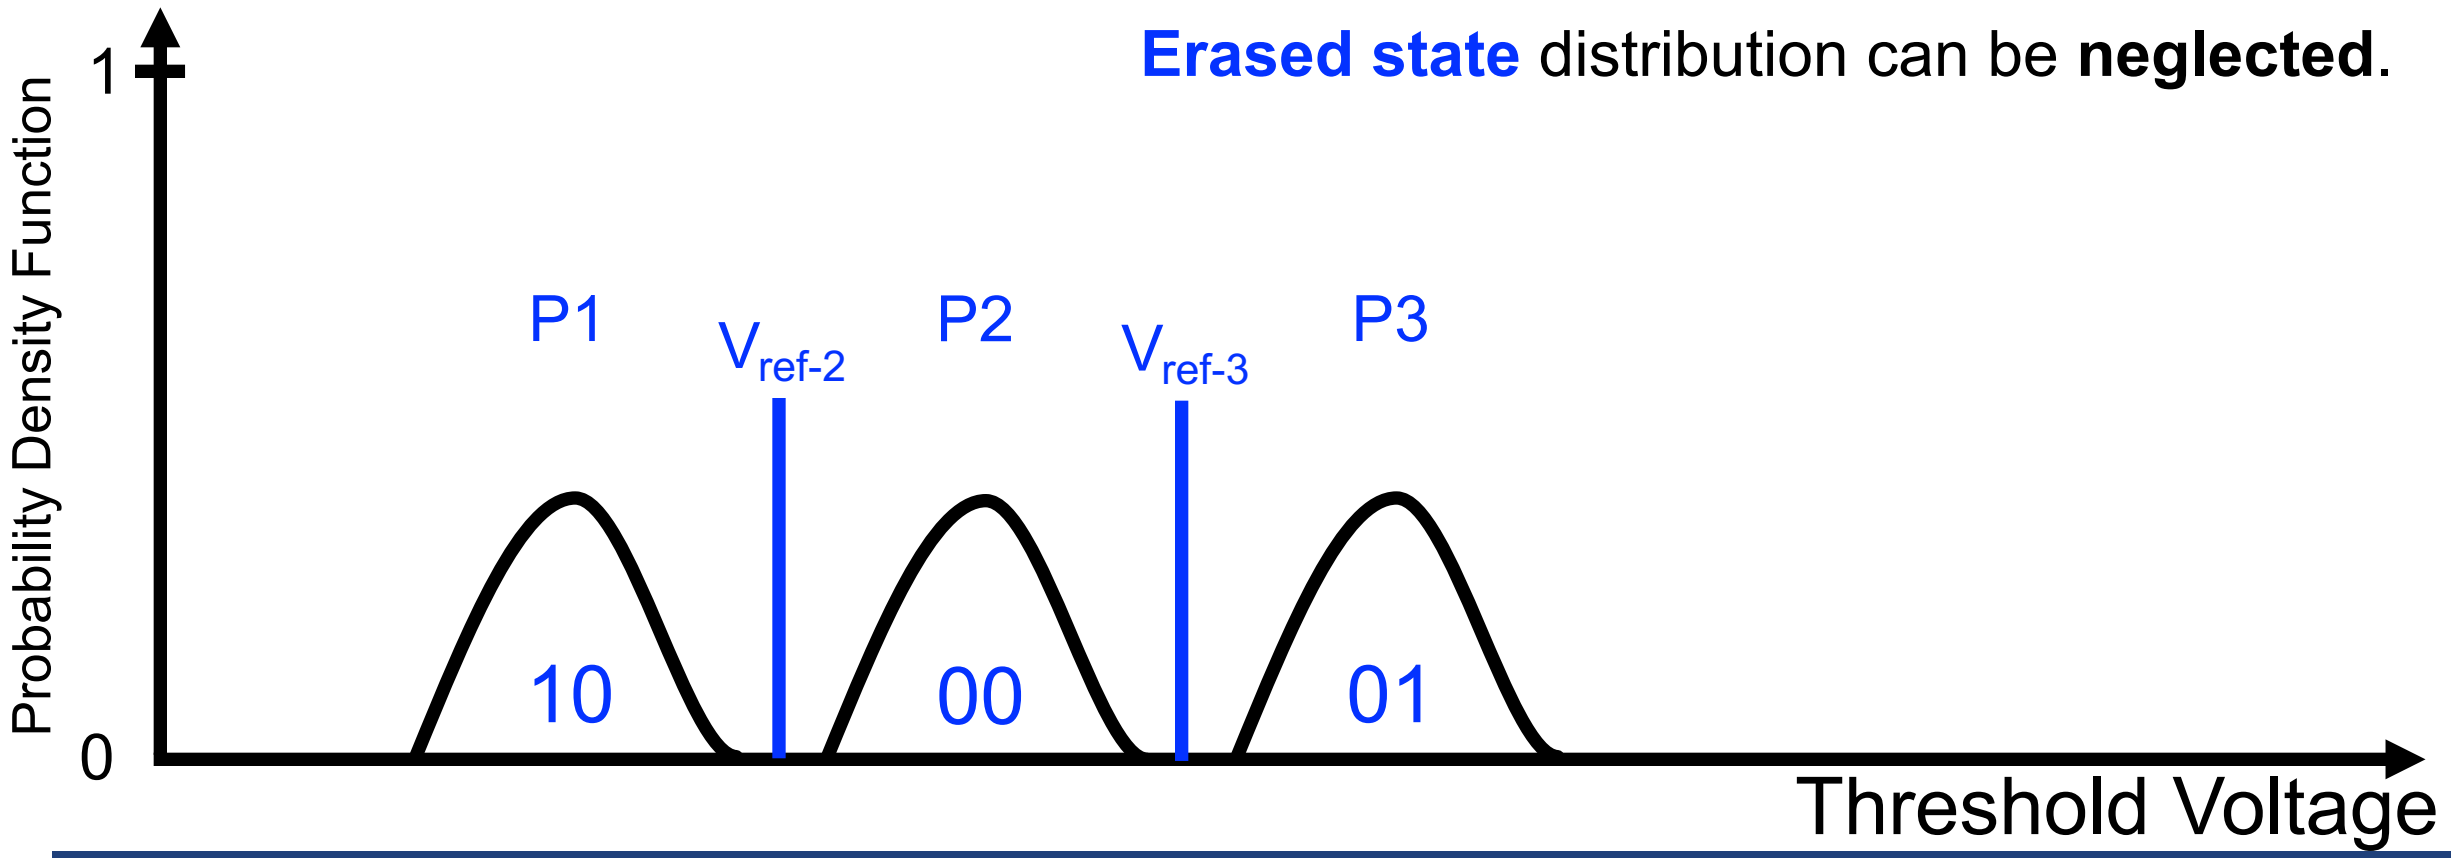

Threshold Voltage Distribution over Time

Source: Slides adapted from Data Retention in MLC NAND Flash Memory… Yixin Luo

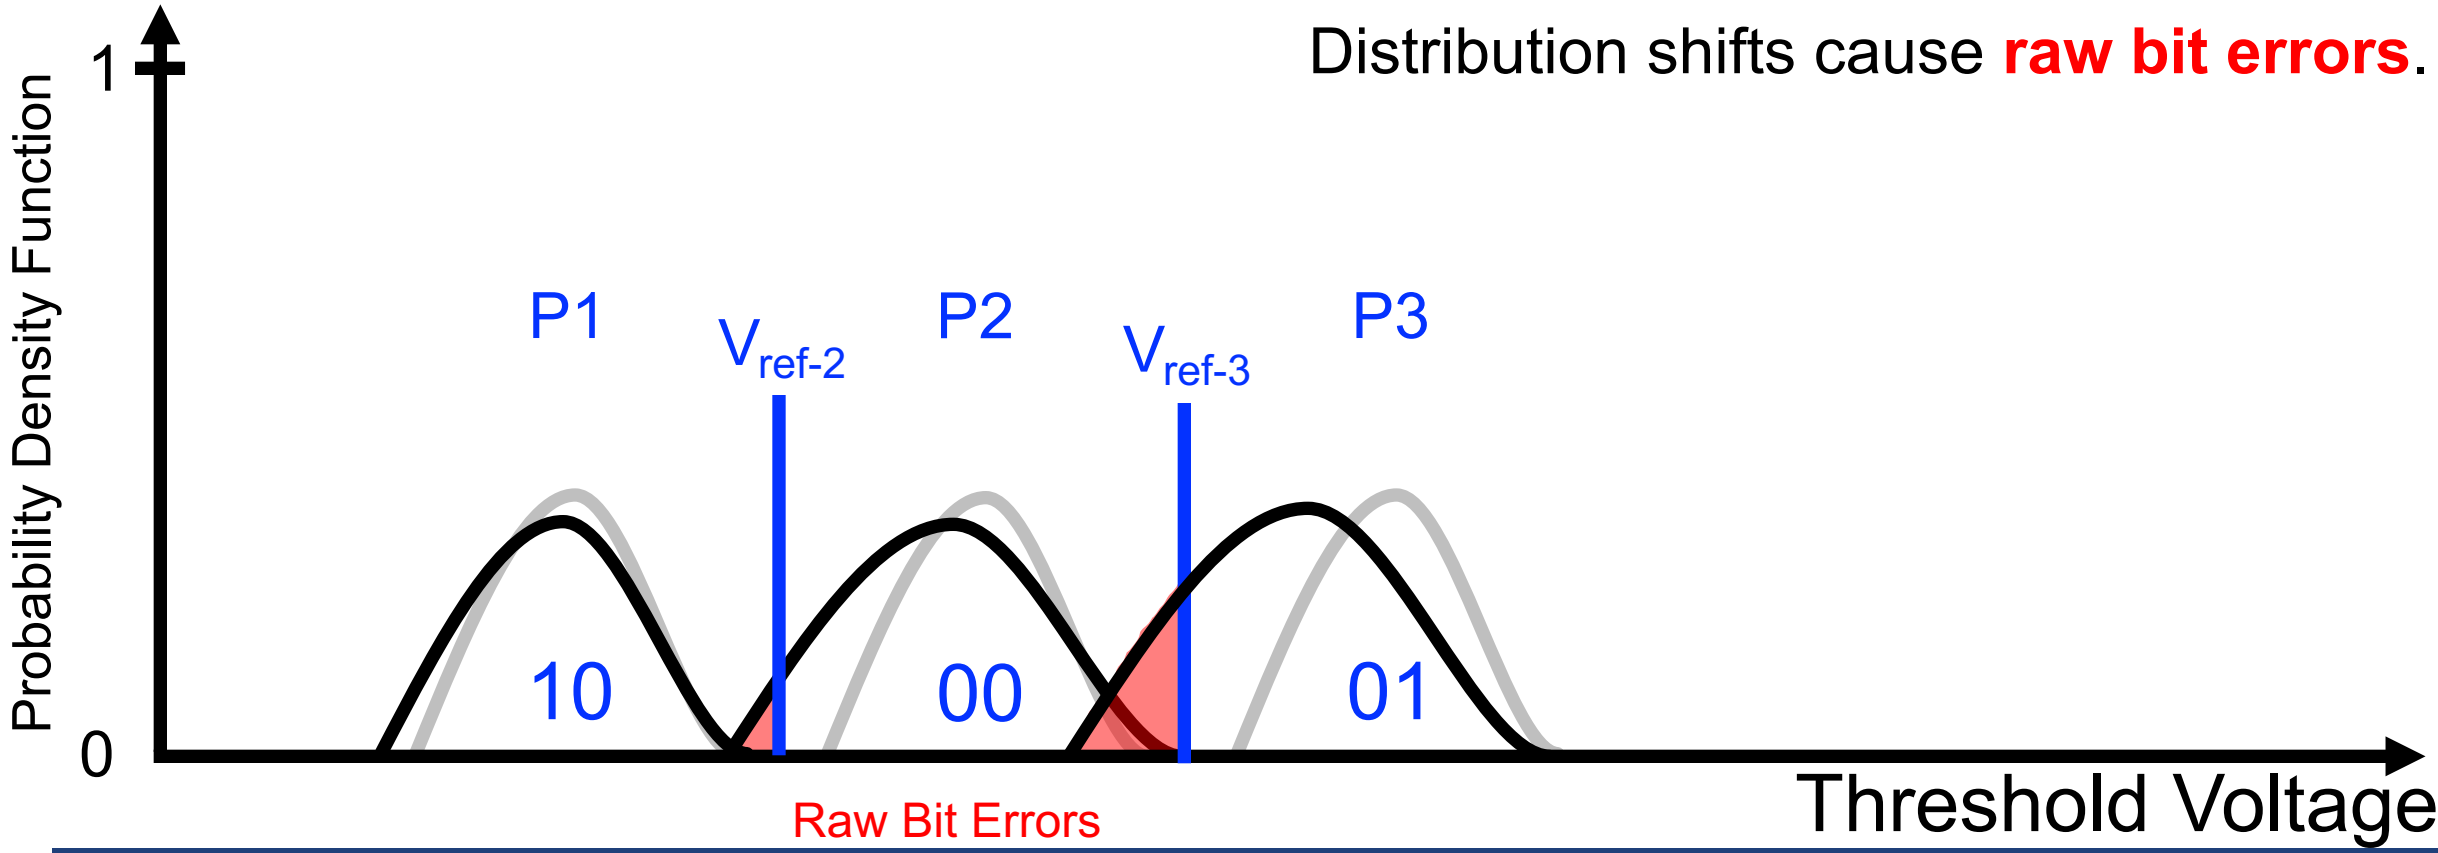

Threshold Voltage Distribution over Time

Source: Slides adapted from Data Retention in MLC NAND Flash Memory… Yixin Luo

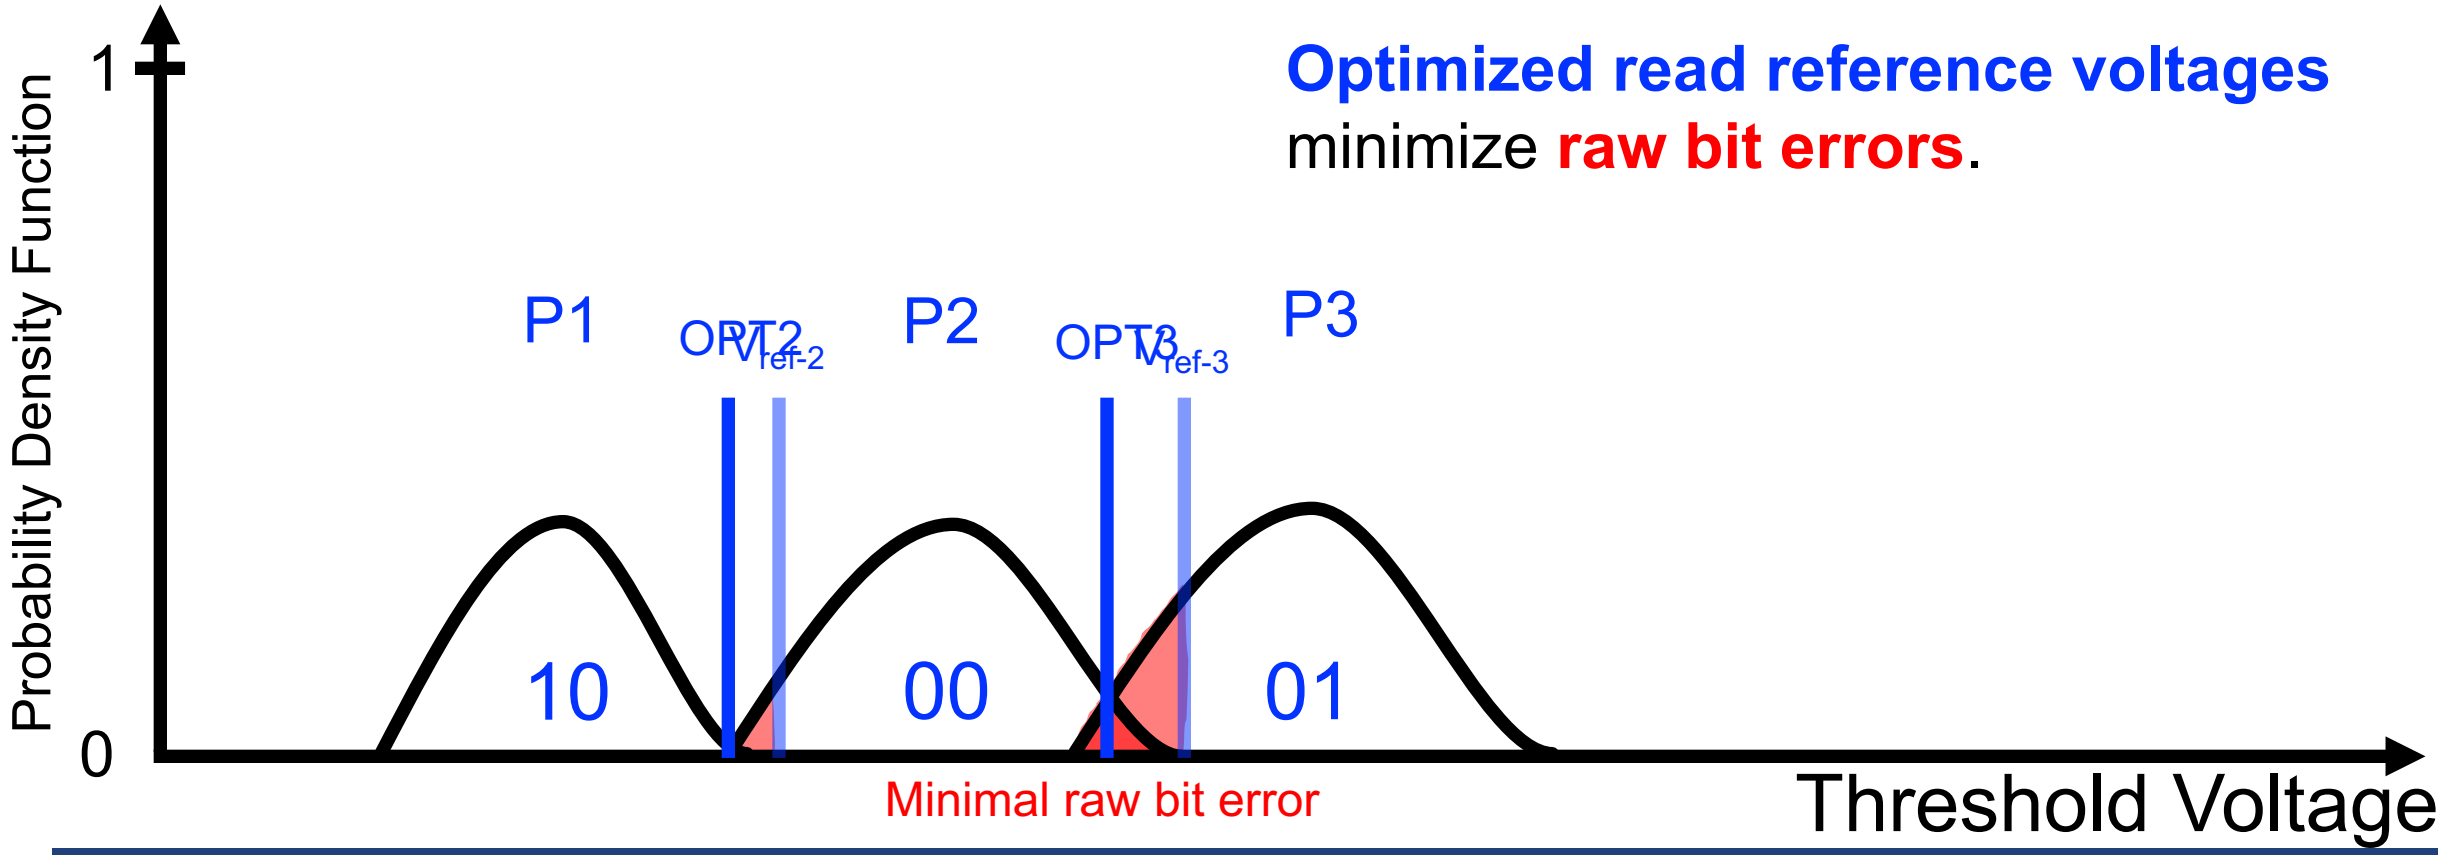

Optimized read reference voltages minimize **raw bit errors**.

Threshold Voltage Distribution over Time

Source: Slides adapted from Data Retention in MLC NAND Flash Memory… Yixin Luo

Retention Optimized Reading

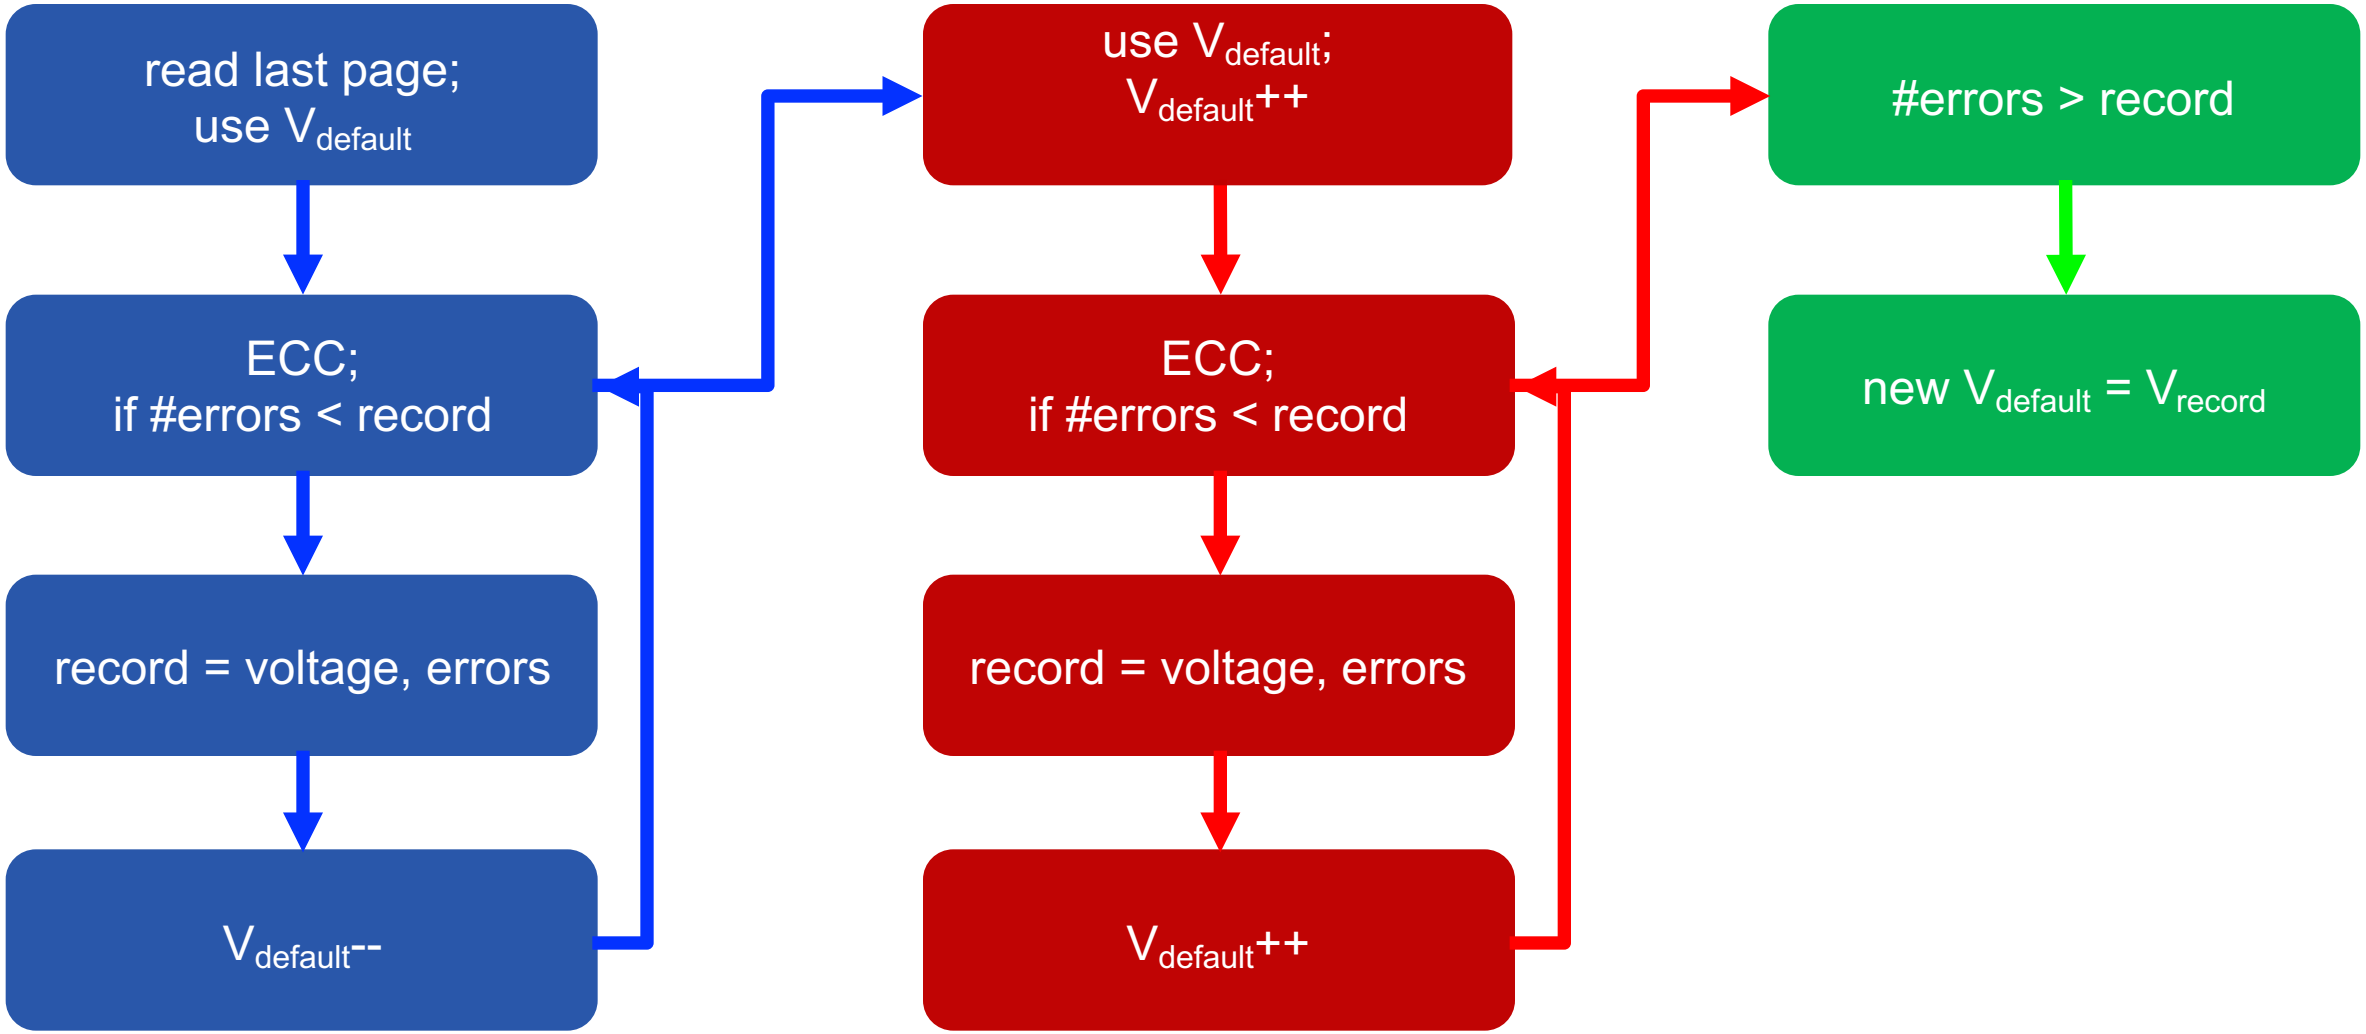

Improved Read-Retry

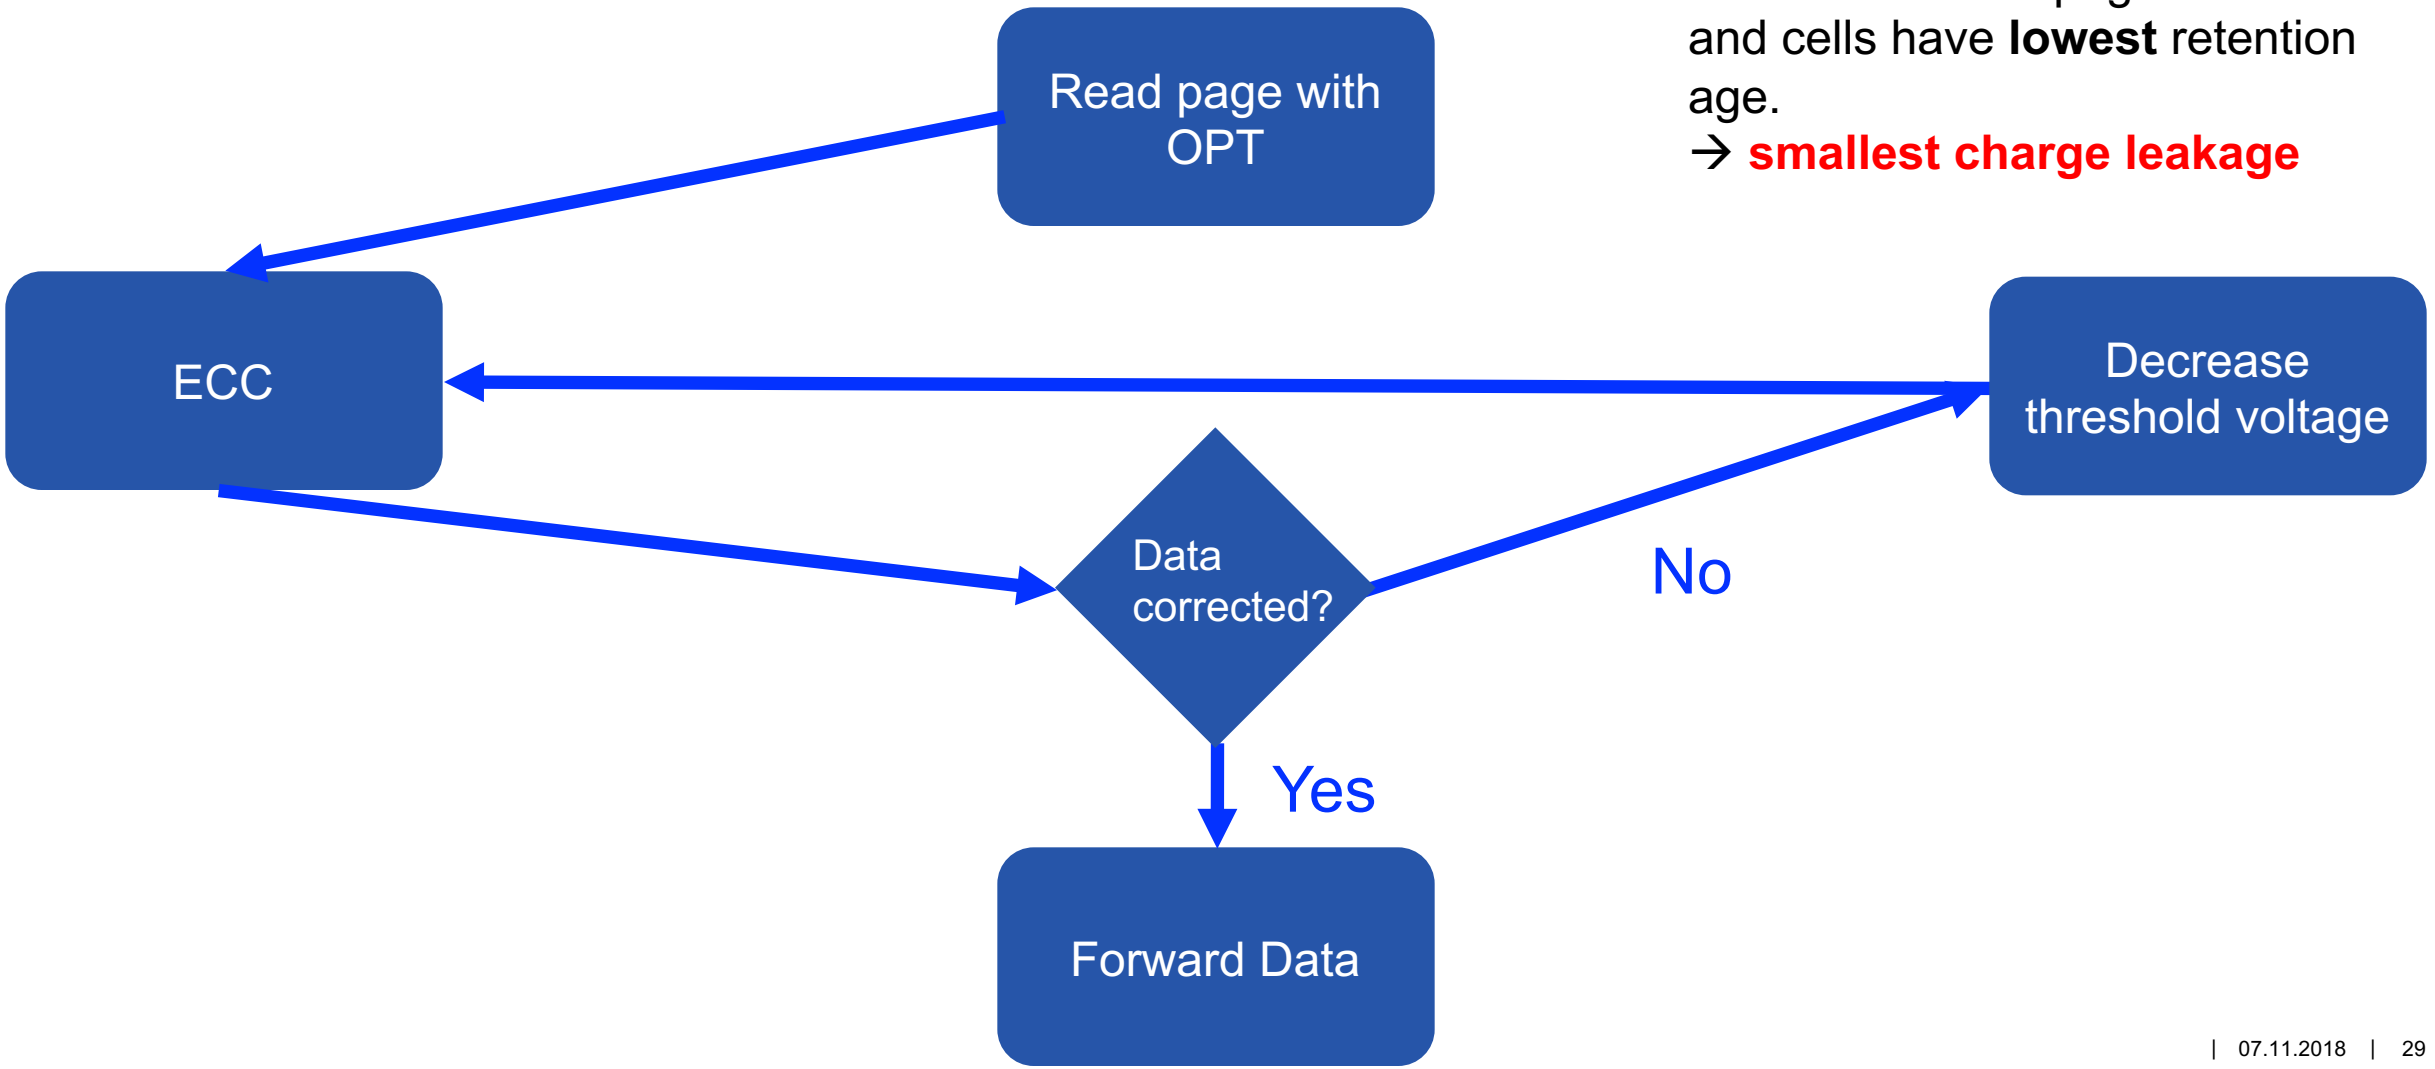

OPT is from **last** page of a block and cells have **lowest** retention age.
→ **smallest charge leakage**

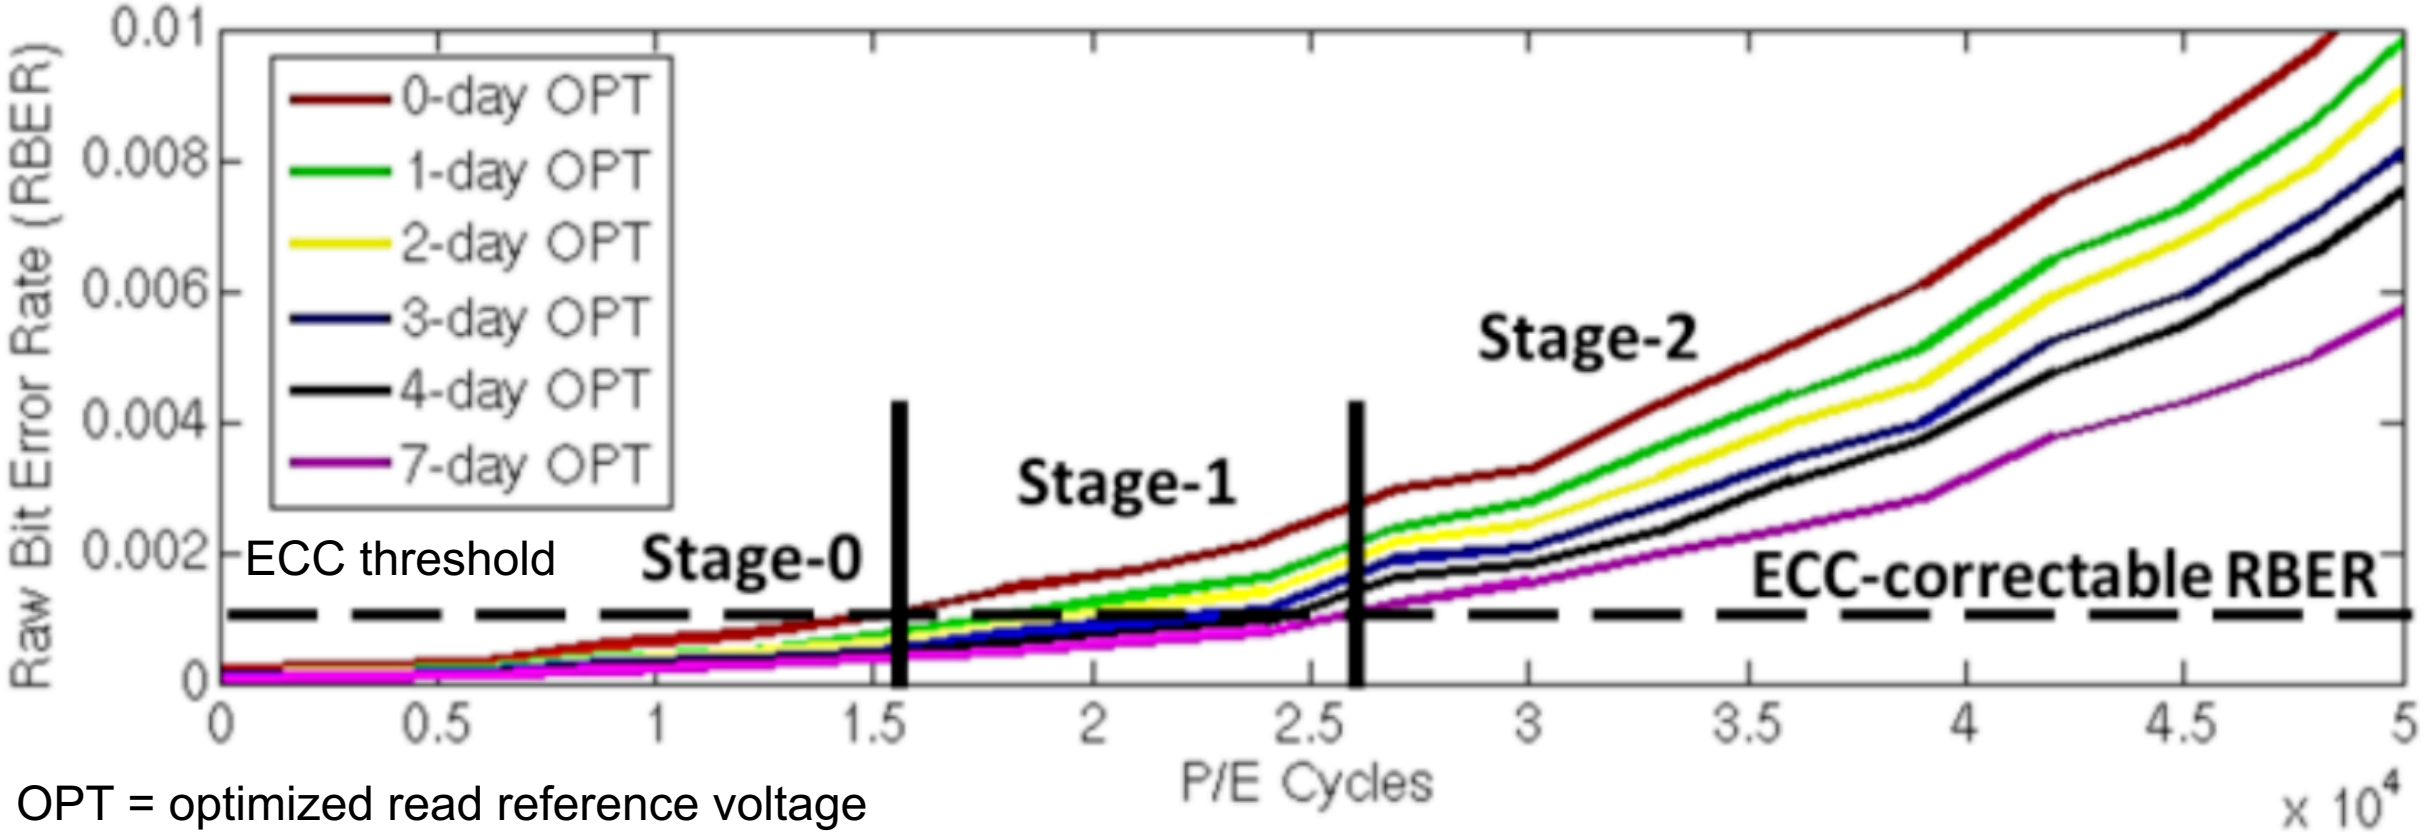

OPT = optimized read reference voltage

Raw Bit Error Rate to Program/Erase Cycles

Source: Y. Cai et al., “Data retention in MLC NAND flash memory:… “ in IEEE 21st Int. Symp. HPCA, 2015

Evaluation

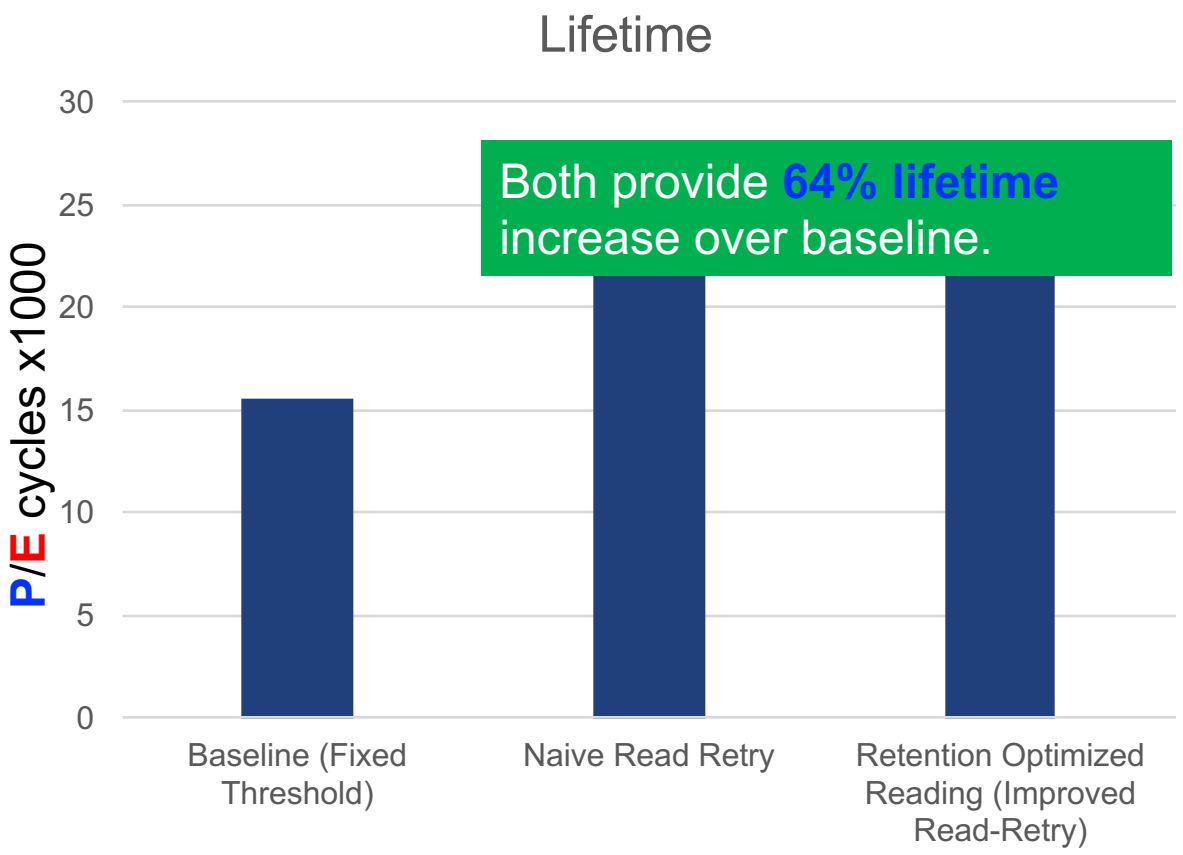
Lifetime
P/E cycles x1000
30
25
20
15
10
5
0
Both provide 64% lifetime increase over baseline.
Baseline (Fixed Threshold)
Naive Read Retry
Retention Optimized Reading (Improved Read-Retry)

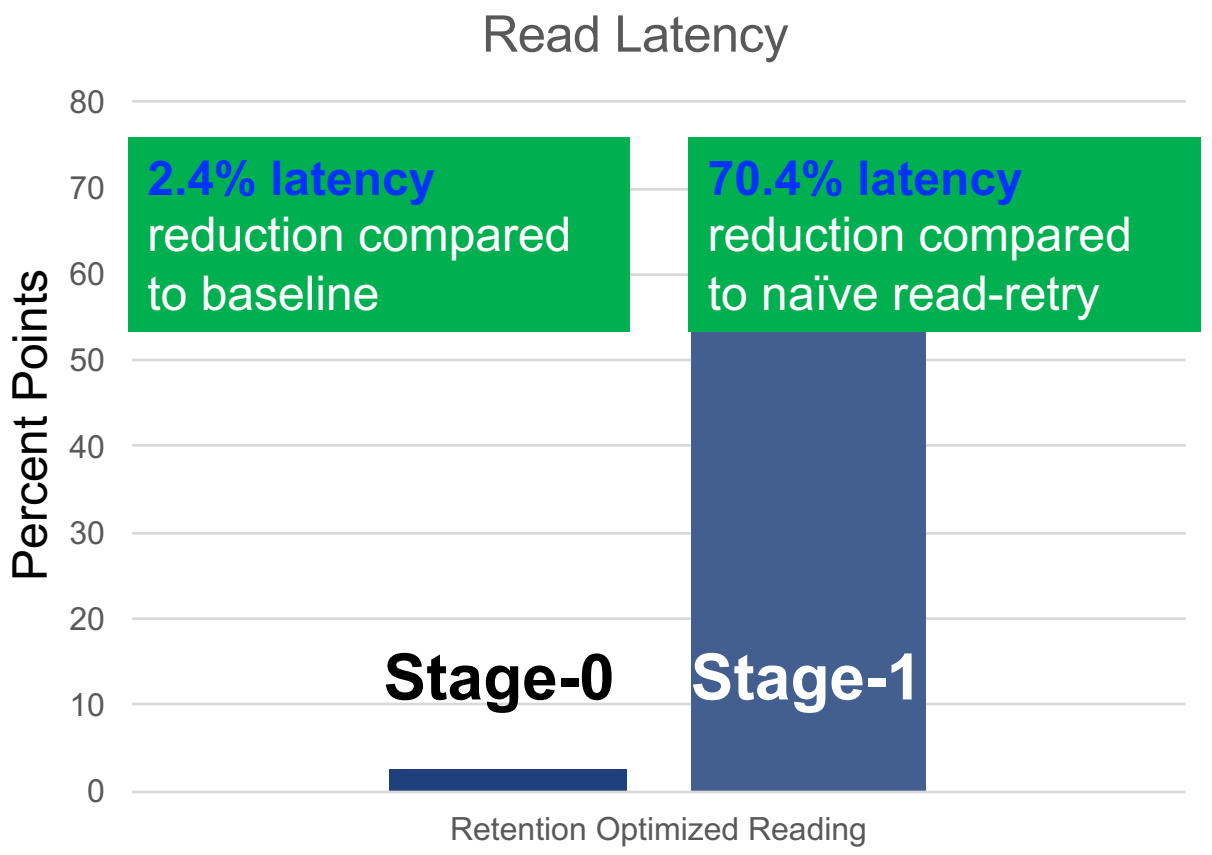
Read Latency
Percent Points
80
70
60
50
40
30
20
10
0
2.4% latency reduction compared to baseline
70.4% latency reduction compared to naïve read-retry
Stage-0
Stage-1
Retention Optimized Reading

Evaluation

- We have a storage overhead of 768 KB out of 512 GB. → **0.00015% overhead**
- **Execution overhead** depends on **program/erase cycles**, **retention age** and **amount of data** written.

Retention Age	P/E Cycles	Latency
1 day	8000	3 seconds
7 days	8000	15 seconds
30 days	8000	23 seconds

Assuming flash capacity is full (512 GB).

ETH zürich

Retention Failure Recovery

Fast and Slow Leaking Cells

- Separate cells into fast and slow leaking cells.
- Over the same time t **fast leaking cells leak more charge** than **slow leaking cells**.
- Threshold separating cells is the **average threshold voltage shift.**

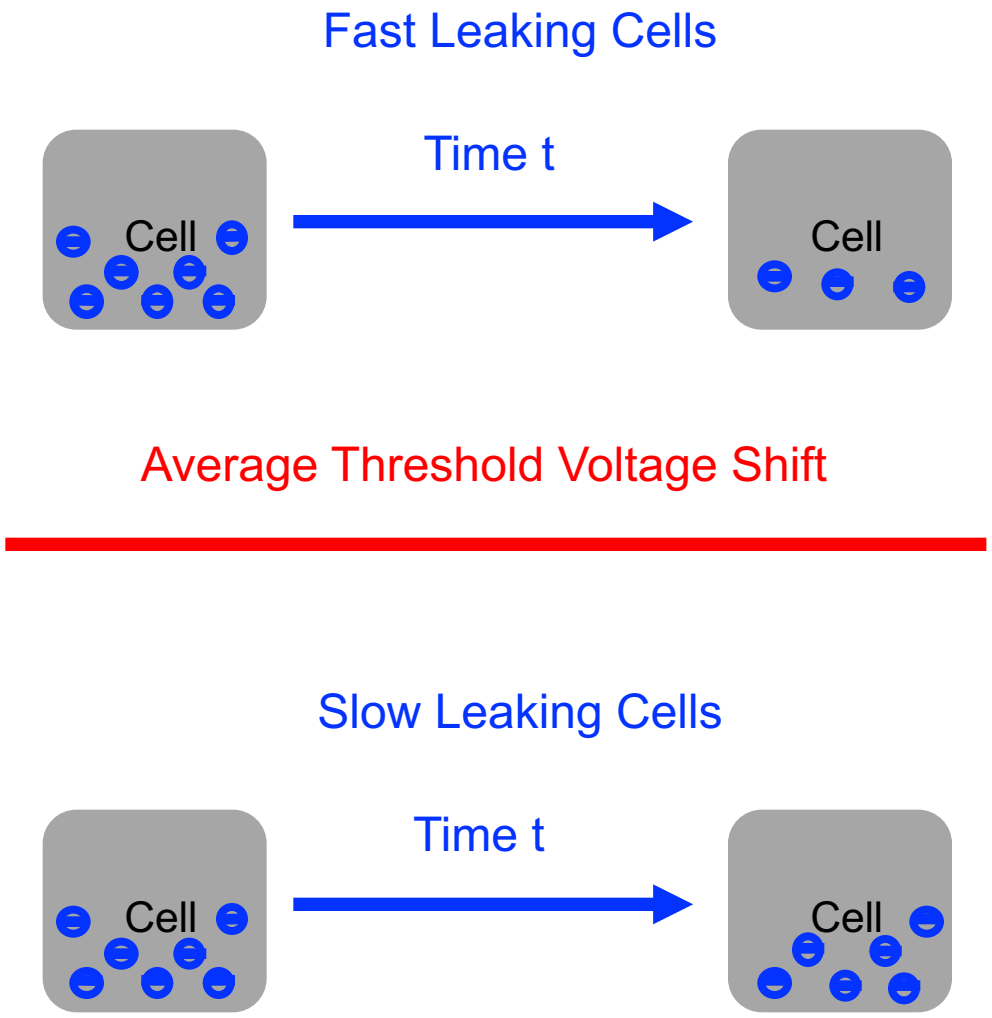

Retention Failure Recovery

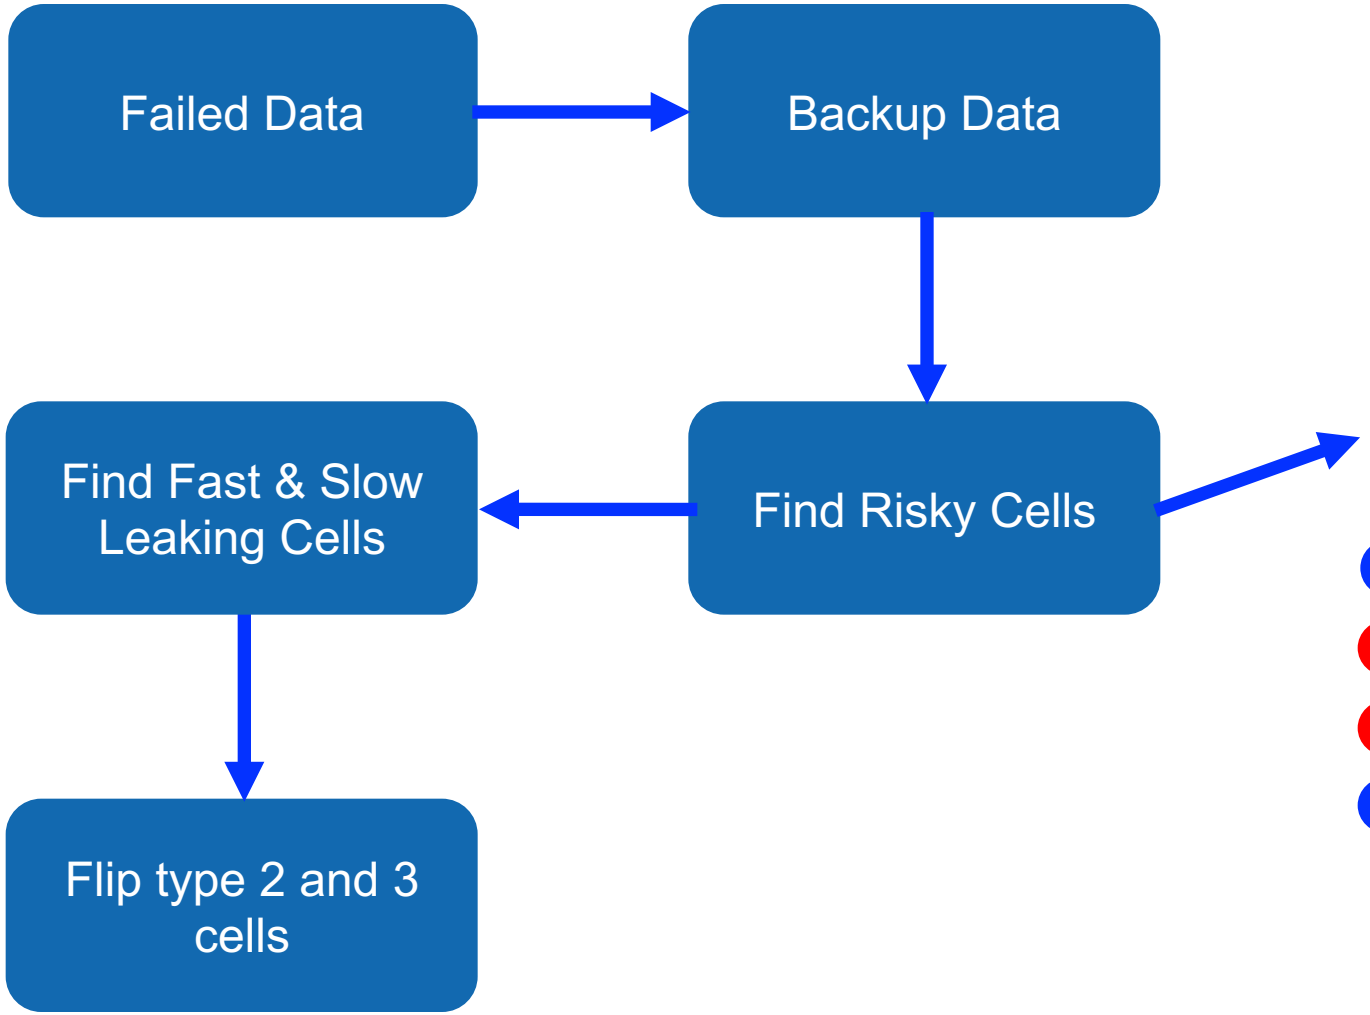

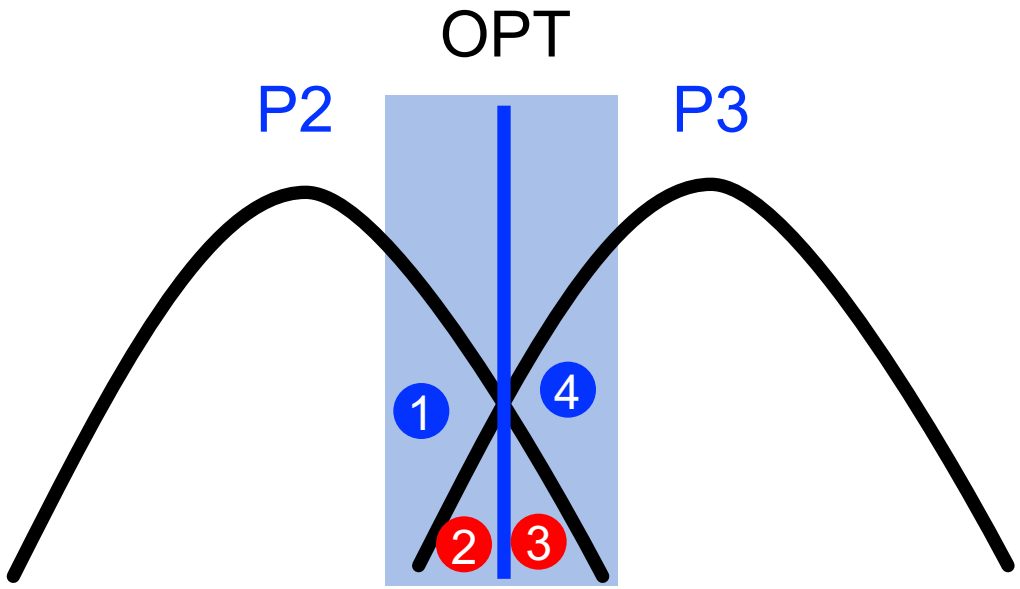

1. Cells correctly in P2
2. Fast leaking cells from P3 wrongly in P2
3. Slow leaking cells from P2 wrongly in P3
4. Cells correctly in P3

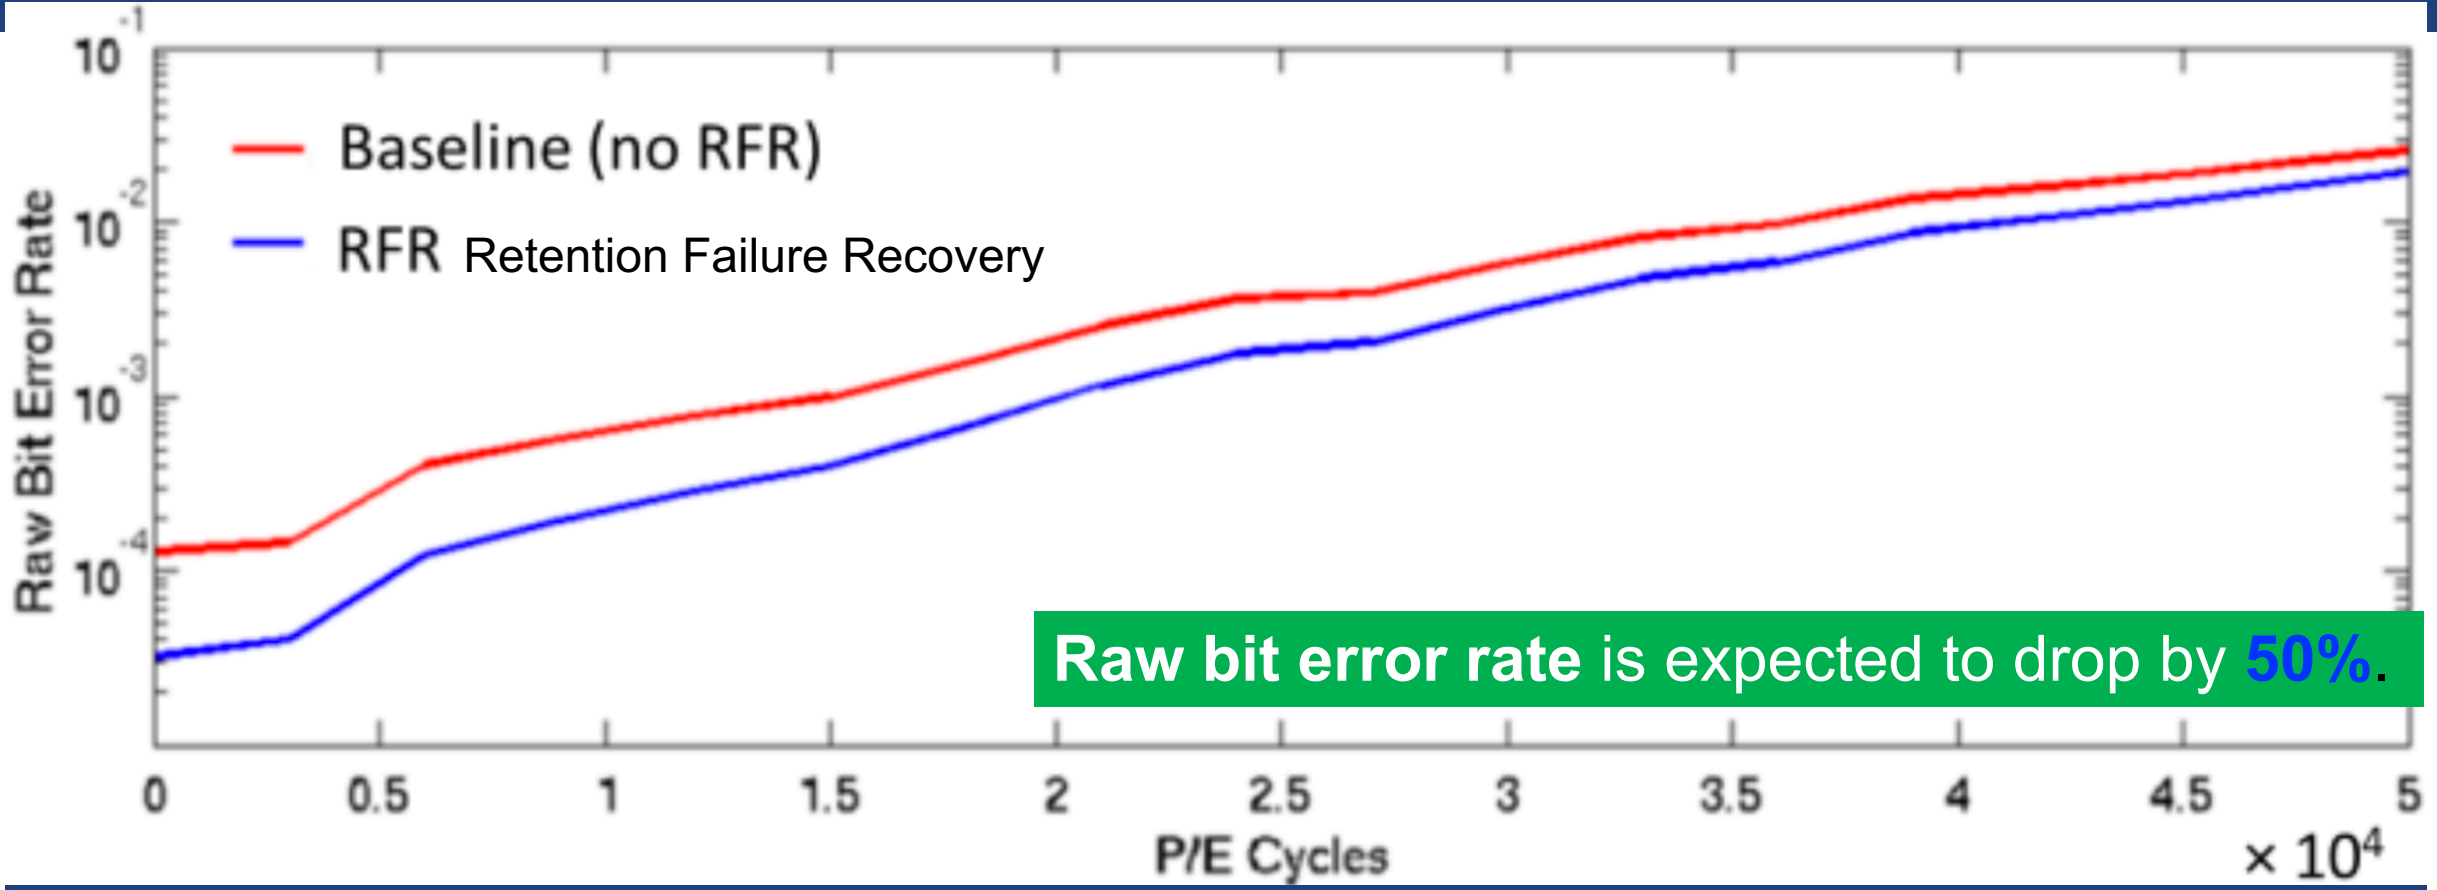

Raw bit error rate is expected to drop by **50%**.

Evaluation

Source: Y. Cai et al., “Data retention in MLC NAND flash memory:… “ in IEEE 21st Int. Symp. HPCA, 2015

Executive Summary

- **Problem**
  - **Density** of flash memory rises and diminishes **lifetime**.
  - Correcting errors increases **read latency**.
- **Goal**
  - Deepen understanding of **voltage threshold distributions** of flash memory.
  - **Improve** both lifetime and system performance.
  - **Recover** non-correctable data.
- **Method**
  - **Retention Optimized Reading**
    - **Improved Read-Retry**
  - **Retention Failure Recovery**
- **Result**
  - **Lifetime** improvement by 64%.
  - **Read latency** reduction by 70.4%.
  - **Raw bit error rate** drop by 50%.

Strengths

- **Retention optimized reading** enhances memory **lifetime** under **low overhead**.
- **Retention failure recovery** decreases **raw bit error rate**.
- Mechanisms **complement** each other, but can be implemented **individually**.
- We may adjust **ECC** capabilities to increase power efficiency.
- **Paper**
  - Presents a **simple** and intuitive **algorithm**.
  - Conducts research with **high potential impact**.

Weaknesses

- How does **temperature** affect threshold voltage shifts?
- **How many** flash memory **devices** were used?
- How does retention failure recovery affect **storage overhead**?
- The paper is **hard to understand** in detail and covers a lot of topics.
- Why was retention optimized reading **not compared** to adaptive voltage threshold?[1]
- The paper has many **similarities** with previously published papers.
- **Figure explanations** are quite sparsely provided.

[1]Papandreou et al., "Using Adaptive Read Voltage Thresholds to Enhance the Reliability of MLC NAND…", Proceedings of the 24th edition of the great lakes symposium on VLSI, 2014

Key Takeaways

- Retention errors **limit** flash memory **lifetime**.
- Read-retry **increases** read **latency**.
- We gained a clear understanding of **threshold voltage distributions**.
- **Retention optimized reading** improves **lifetime** and **read latency**.
- **Retention failure recovery** reduces **errors**.

Open Discussion

- In what **order** should we assign our **2 bit values** to our **4 states**?
  - They are often assigned this way: **Erased** - **11**, **P1** - **10**, **P2** - **00**, **P3** - **01**.
  - Because if the **threshold voltage** were to shift to the left we only get one bit error.

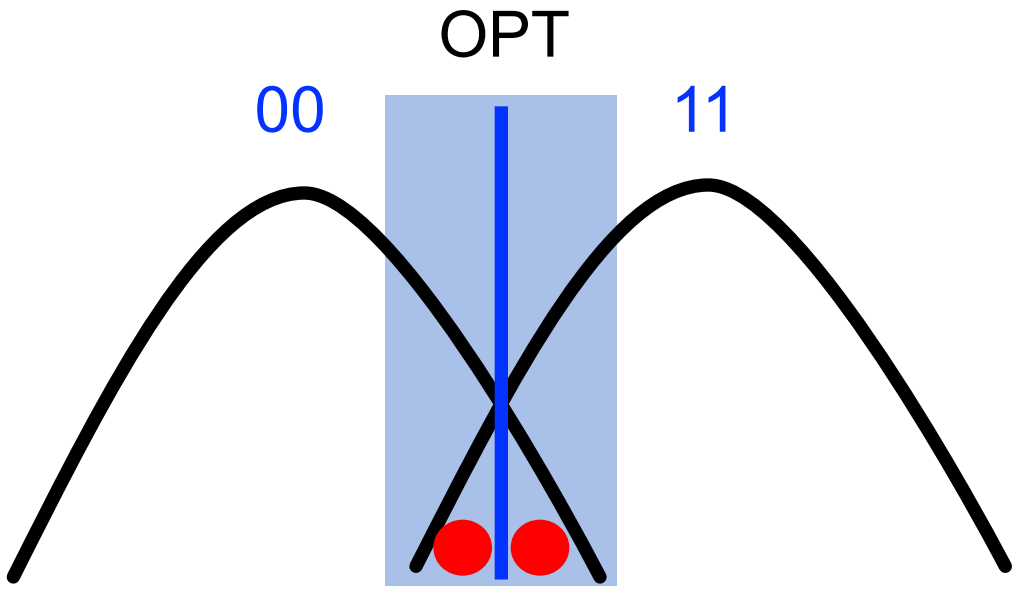

Open Discussion

- How should we assign our **2 bit values** to **pages**?

Row Index	LSB of the 2^{17} cells	MSB of the 2^{17} cells
0	Page 0	Page 2
1	Page 1	Page 4
2	Page 3	Page 6
…	…	…
127	Page 253	Page 255

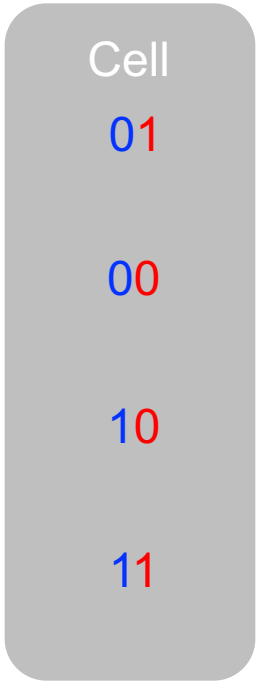

Source: Table adapted from Wang, Wei, et al. "Reducing MLC flash memory retention errors through programming initial step only.", MSST 31st Symposium on. IEEE, 2015

Open Discussion

- We have seen that reducing the number of **read-retries** has a great impact on **read latency**.
- Can you think of yet another method to **reduce** the number of **read-retries**?
  - My idea would be to use **binary search** implemented into our current read-retry mechanism.

Improved Read-Retry

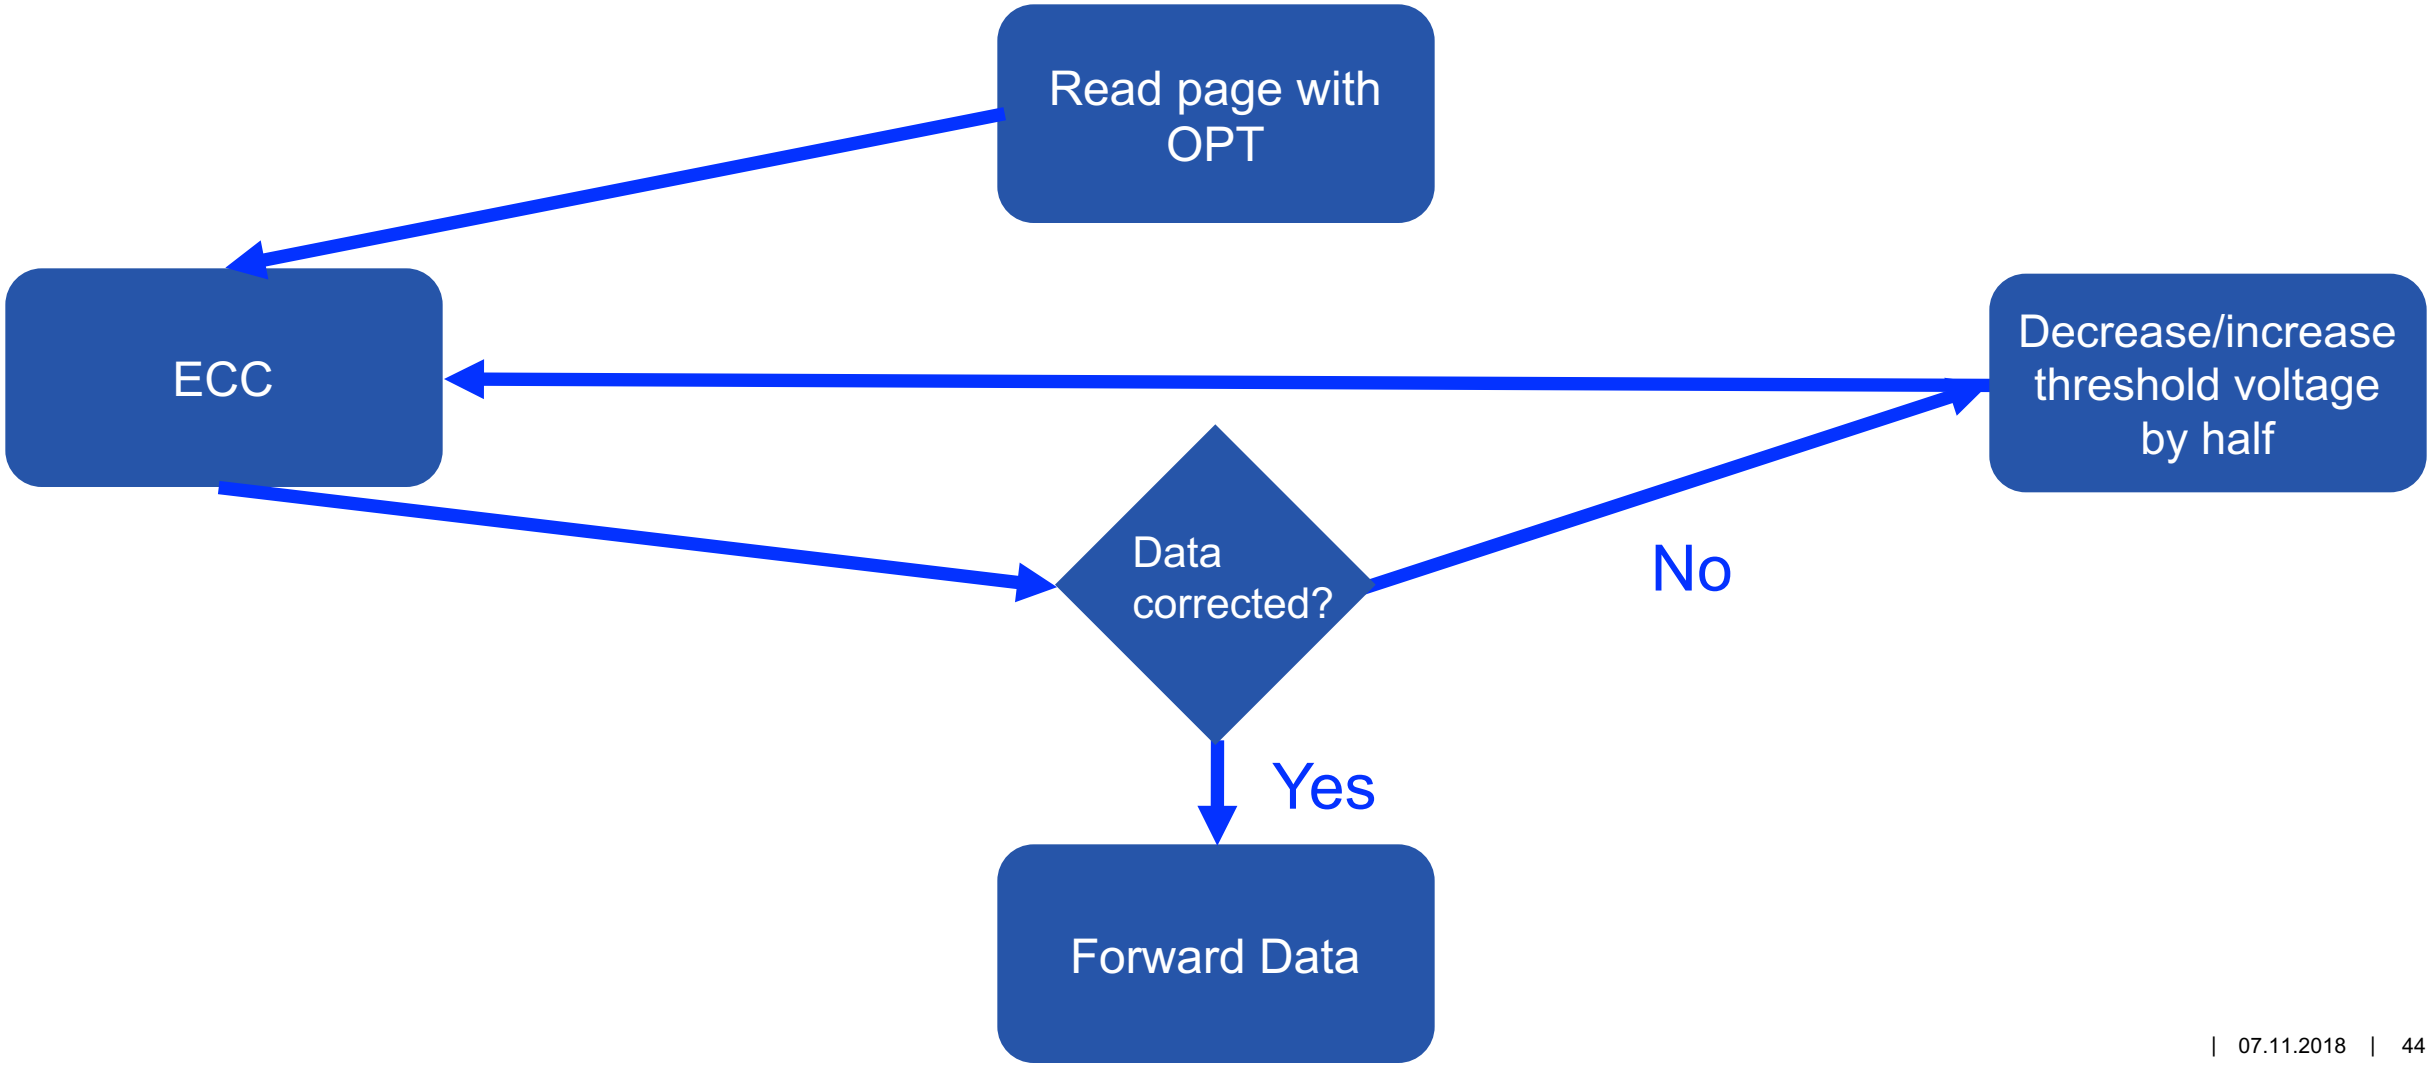

Additional Papers

- *Bez et al., "Introduction to Flash Memory", 2003*
- *Cai et al., "Flash Correct-and-Refresh: Retention-Aware Error Management for Increased Flash Memory Lifetime", 2012*
- *Cai et al., "Error Analysis and Retention-Aware Error Management For NAND Flash Memory", 2013*
- *Cai et al., "Threshold Voltage Distribution in MLC NAND Flash Memory: Characterization, Analysis, and Modeling", 2013*
- *Papandreou et al., "Using Adaptive Read Voltage Thresholds to Enhance the Reliability of MLC NAND Flash Memory Systems", 2014*
- *Aslam et al., "Read and Write Voltage Signal Optimization for Multi-Level-Cell (MLC) NAND Flash Memory", 2016*
- *Coutet et al., "Influence of temperature of storage, write and read operations on multiple level cells NAND flash memories", 2018*

Big Thanks to Giray & Mohammed for their support.

No, really, thanks.

ETH zürich

Backup Slides

Flash Correct-and-Refresh

- Read page with **fixed read reference voltage**.
- **Error correction** informs about range of actual **voltage threshold**.
- Identify cells in a **wrong state**.
- Identify **right shift errors** and **left shift errors**.
- **Left shift errors** are caused by **retention loss**.
- **Right shift errors** are cause by **cell-to-cell interference** when **programming** other cells.

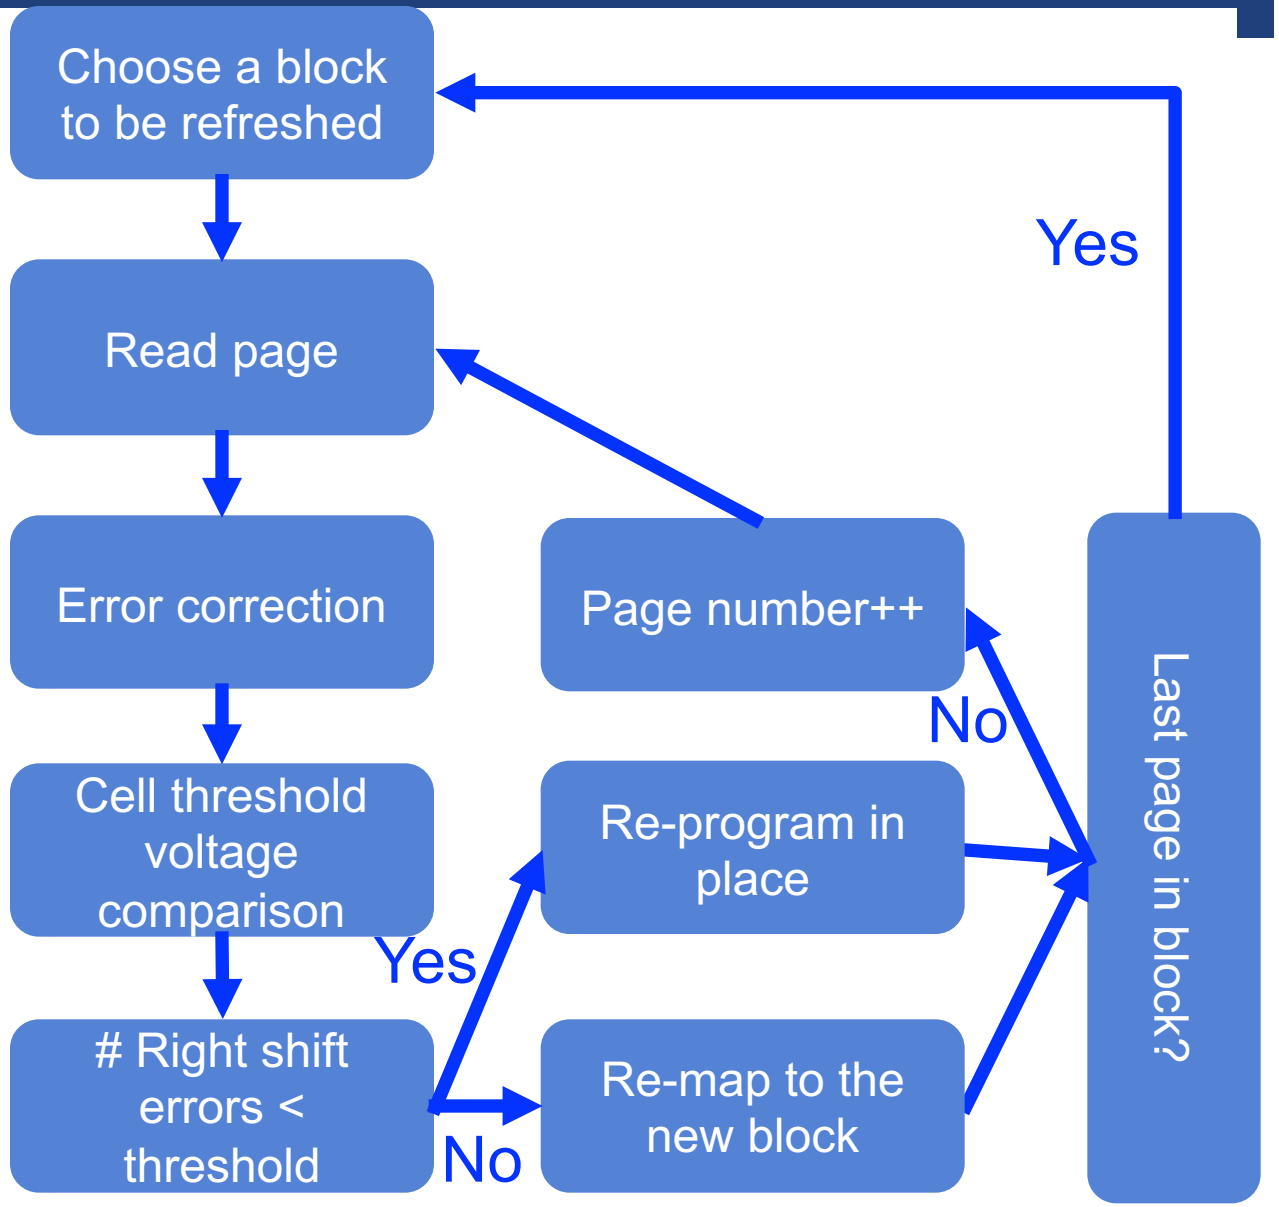

Source: Figure adapted from Y. Cai et al., "Flash Correct-and-Refresh:…", 2012.

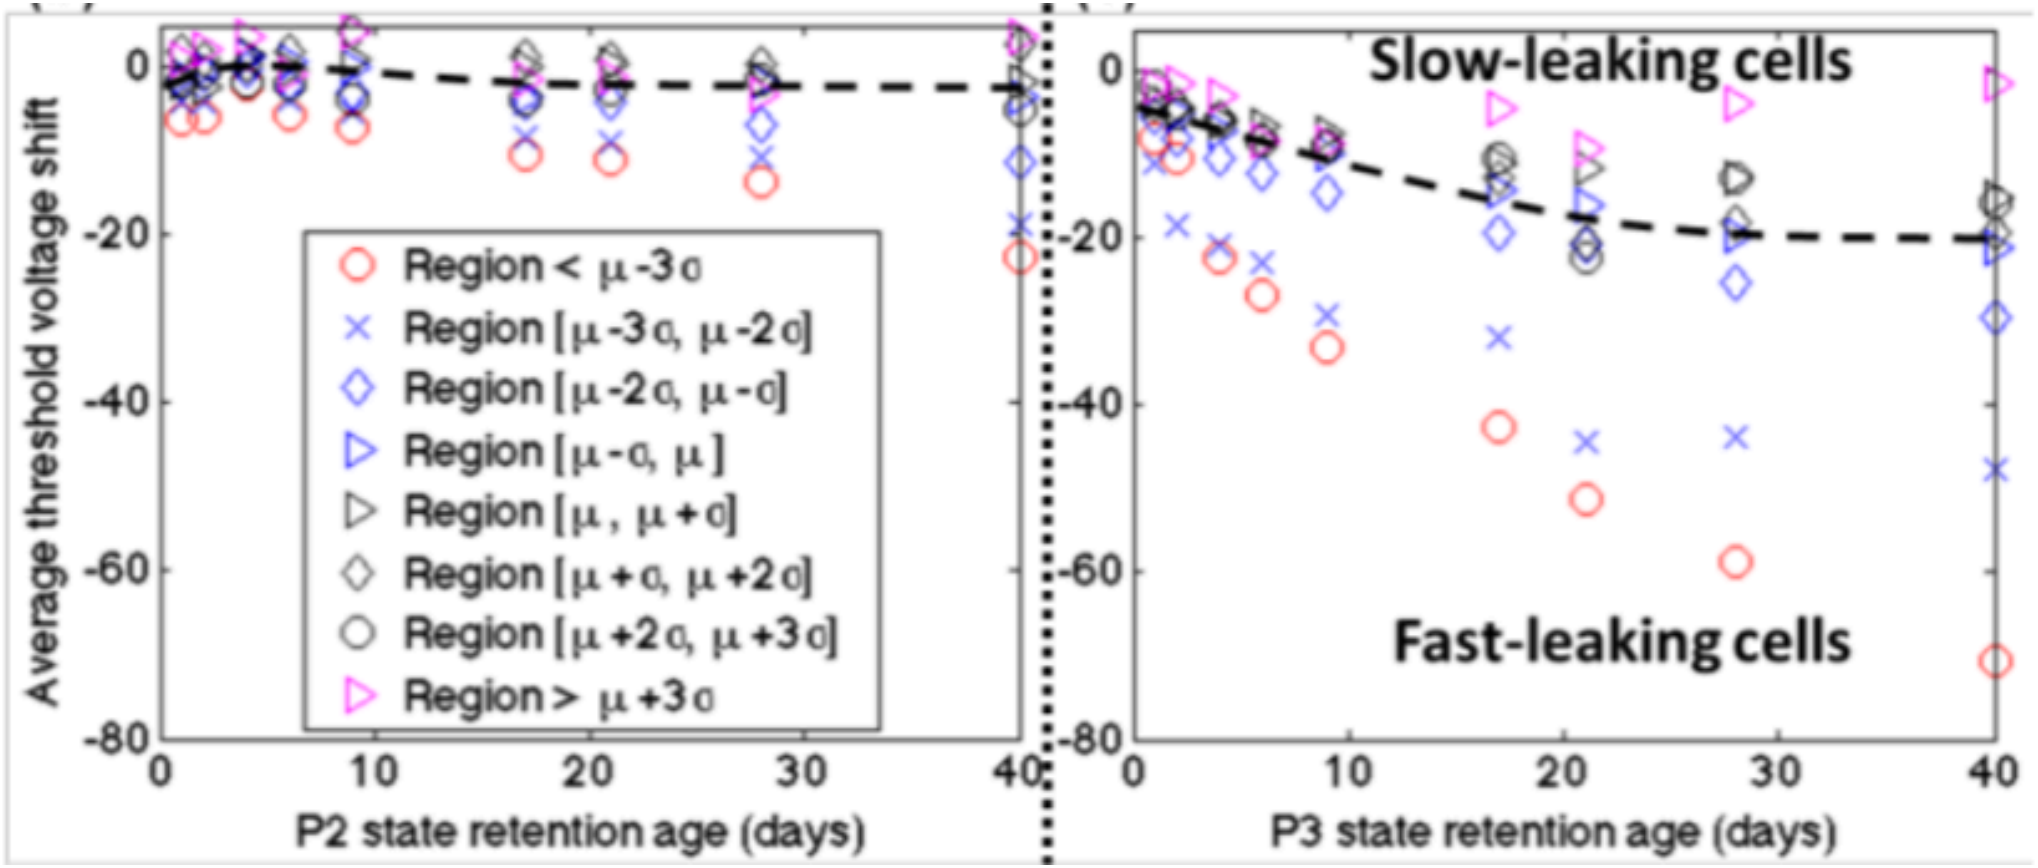

Fast and Slow Leaking Cells