Read Disturb Errors in MLC NAND Flash Memory

Characterization, Mitigation & Recovery

Onur Mutlu
onur@cmu.edu

(joint work with Yu Cai, Yixin Luo, Saugata Ghose, Erich Haratsch, Ken Mai)

August 12, 2015

Flash Memory Summit 2015, Santa Clara, CA

SAFARI

Carnegie Mellon

SEAGATE

DSSC

Paper

- Presented at IEEE/IFIP DSN 2015 Conference in June 2015.
- Full paper for details:

  - 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"
    Proceedings of the
    45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (***DSN***), Rio de Janeiro, Brazil, June 2015.

  - http://users.ece.cmu.edu/~omutlu/pub/flash-read-disturb-errors_dsn15.pdf

Executive Summary

- ***Read disturb errors*** limit flash memory lifetime today
  - Apply a *high pass-through voltage (V_{pass})* to multiple pages on a read
  - Repeated application of V_{pass} can alter stored values in unread pages
- We **characterize read disturb** on real NAND flash chips
  - Slightly lowering V_{pass} greatly reduces read disturb errors
  - Some flash cells are more prone to read disturb
- **Technique 1:** Mitigate read disturb errors online
  - ***V_{pass} Tuning*** dynamically finds and applies a lowered V_{pass} per block
  - Flash memory lifetime improves by 21%
- **Technique 2:** Recover after failure to prevent data loss
  - ***Read Disturb Oriented Error Recovery*** (RDR) selectively corrects cells more susceptible to read disturb errors
  - Reduces raw bit error rate (RBER) by up to 36%

Outline

- Background (Problem and Goal)
- Key Experimental Observations
- Mitigation: V_{pass} Tuning
- Recovery: Read Disturb Oriented Error Recovery
- Conclusion

Outline

- Background (Problem and Goal)
- Key Experimental Observations
- Mitigation: V_{pass} Tuning
- Recovery: Read Disturb Oriented Error Recovery
- Conclusion

NAND Flash Memory Background

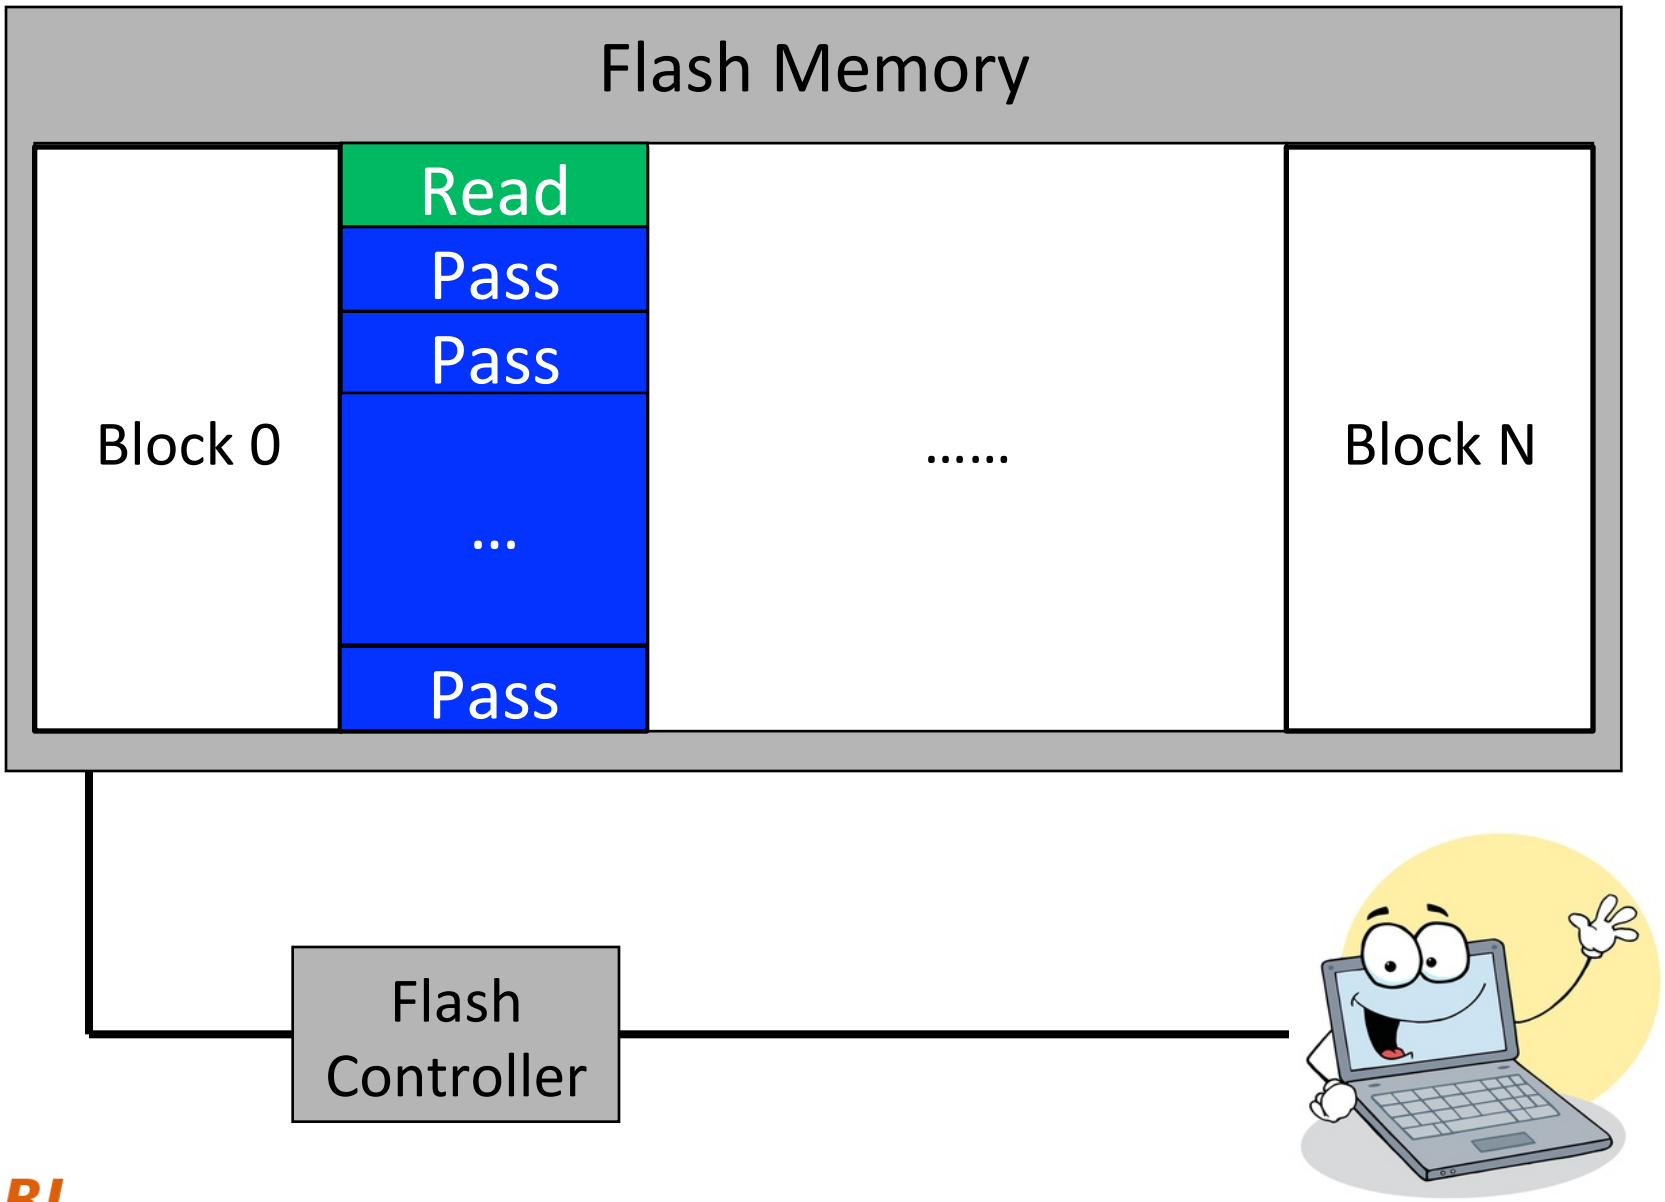

Flash Cell Array

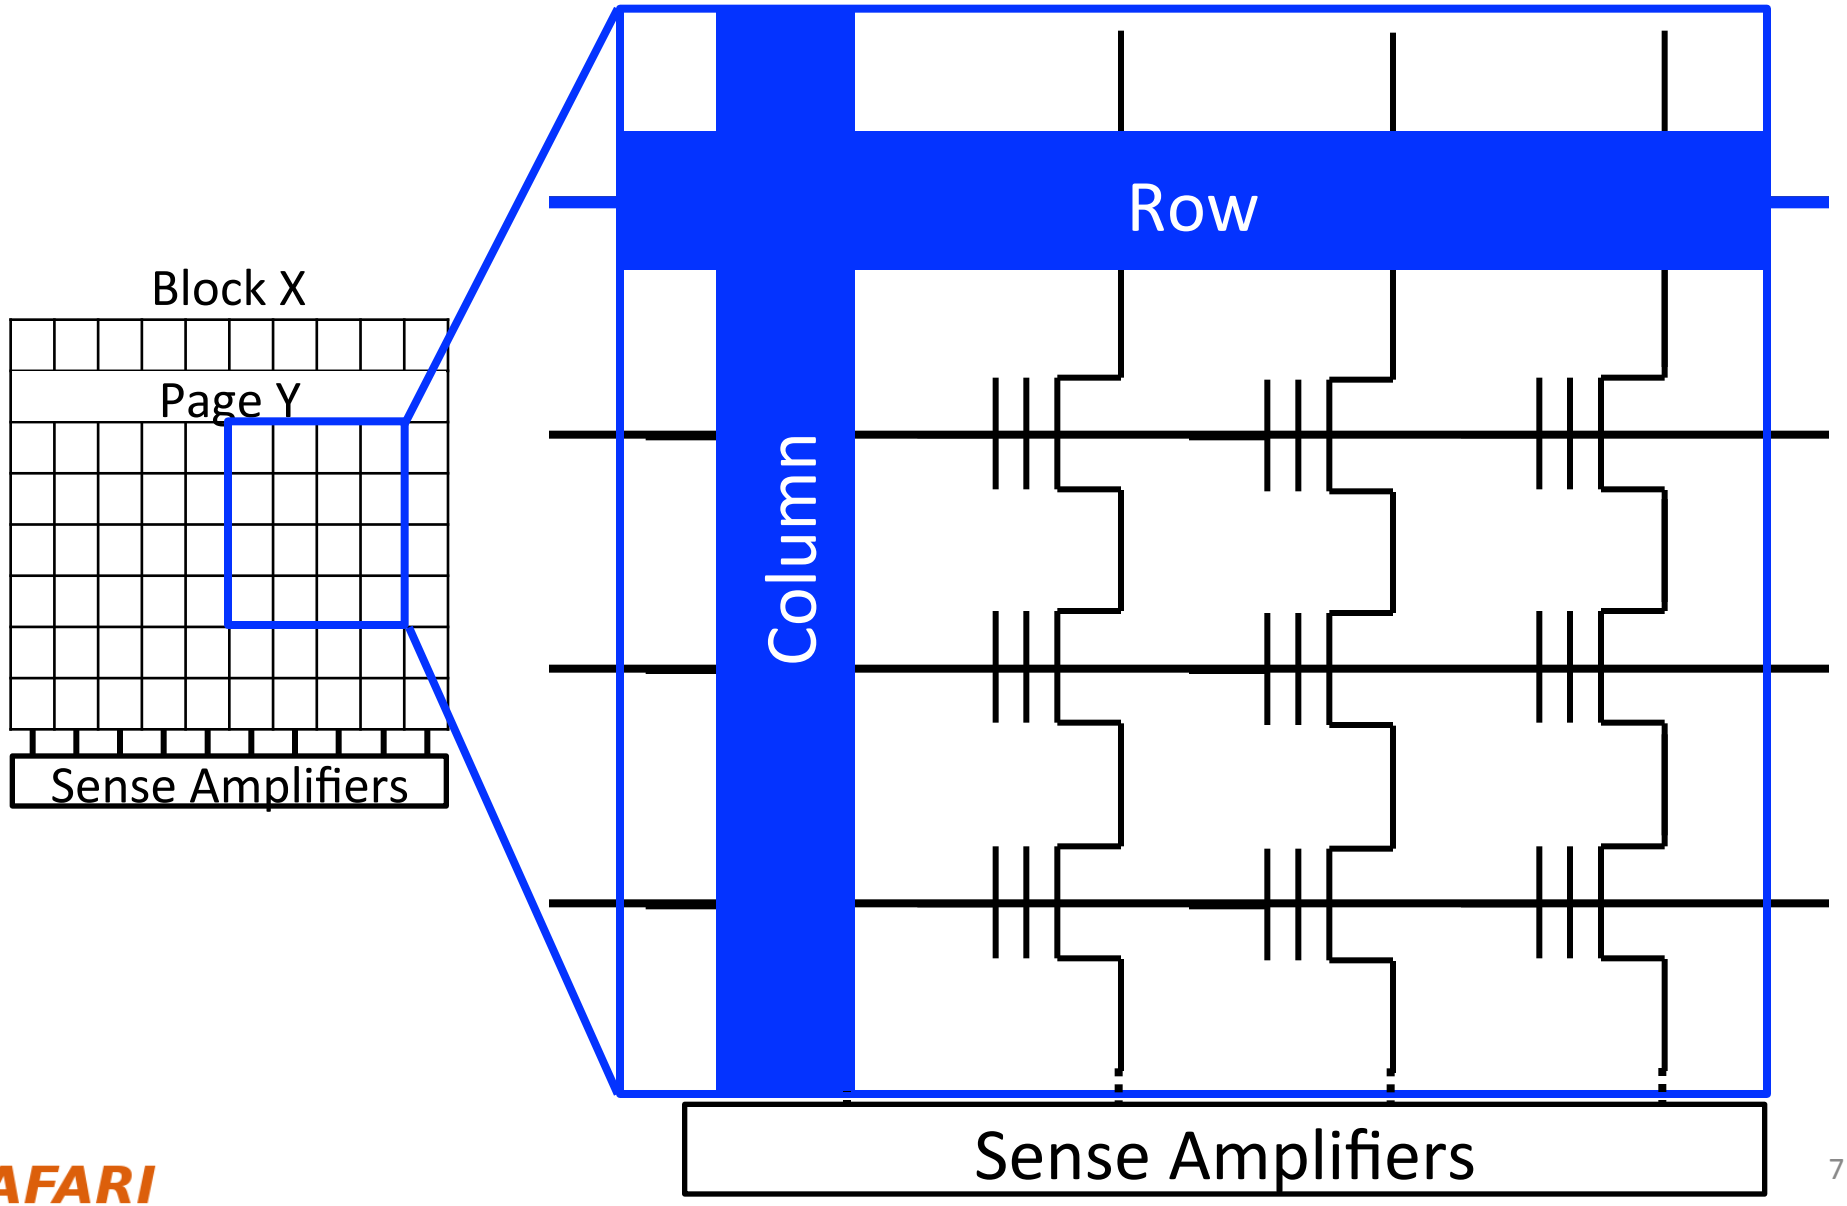

Flash Cell

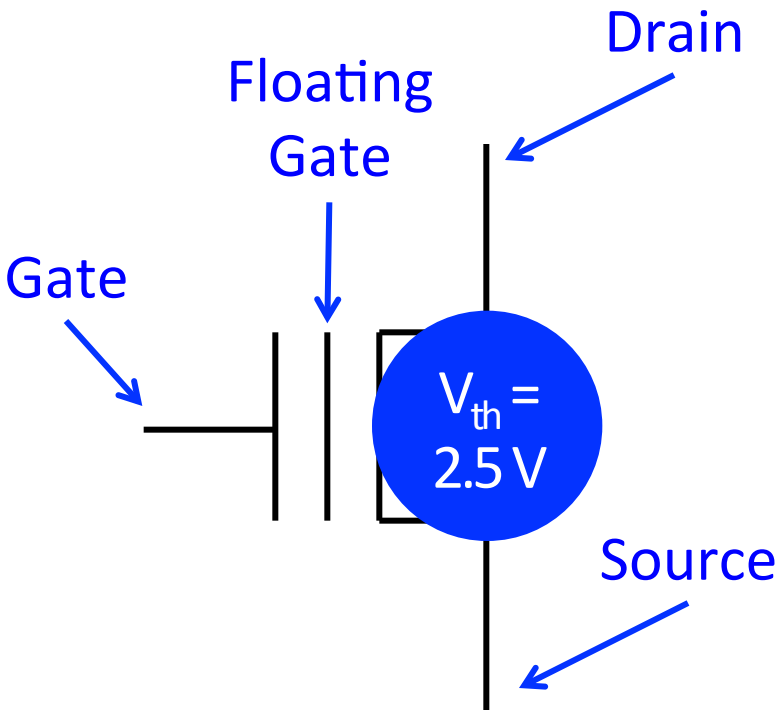

Floating Gate Transistor
(Flash Cell)

Flash Read

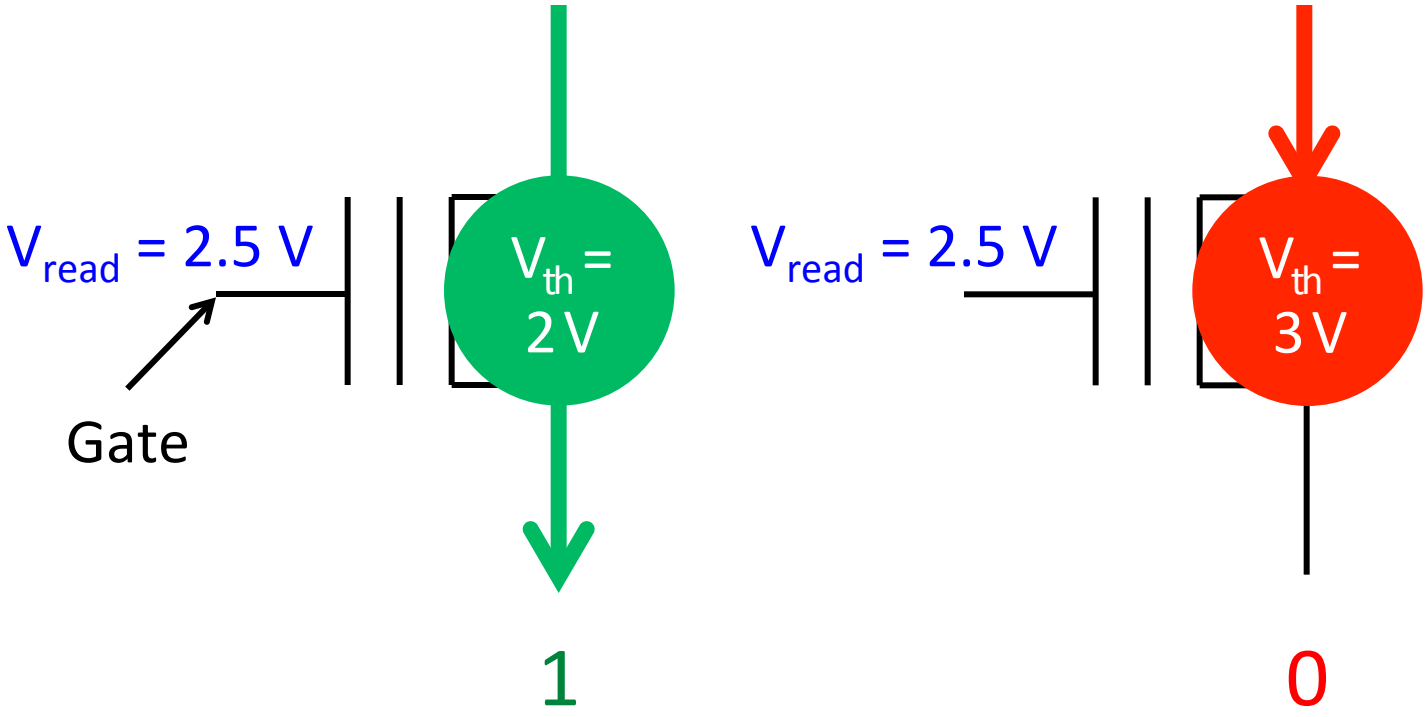

Flash Pass-Through

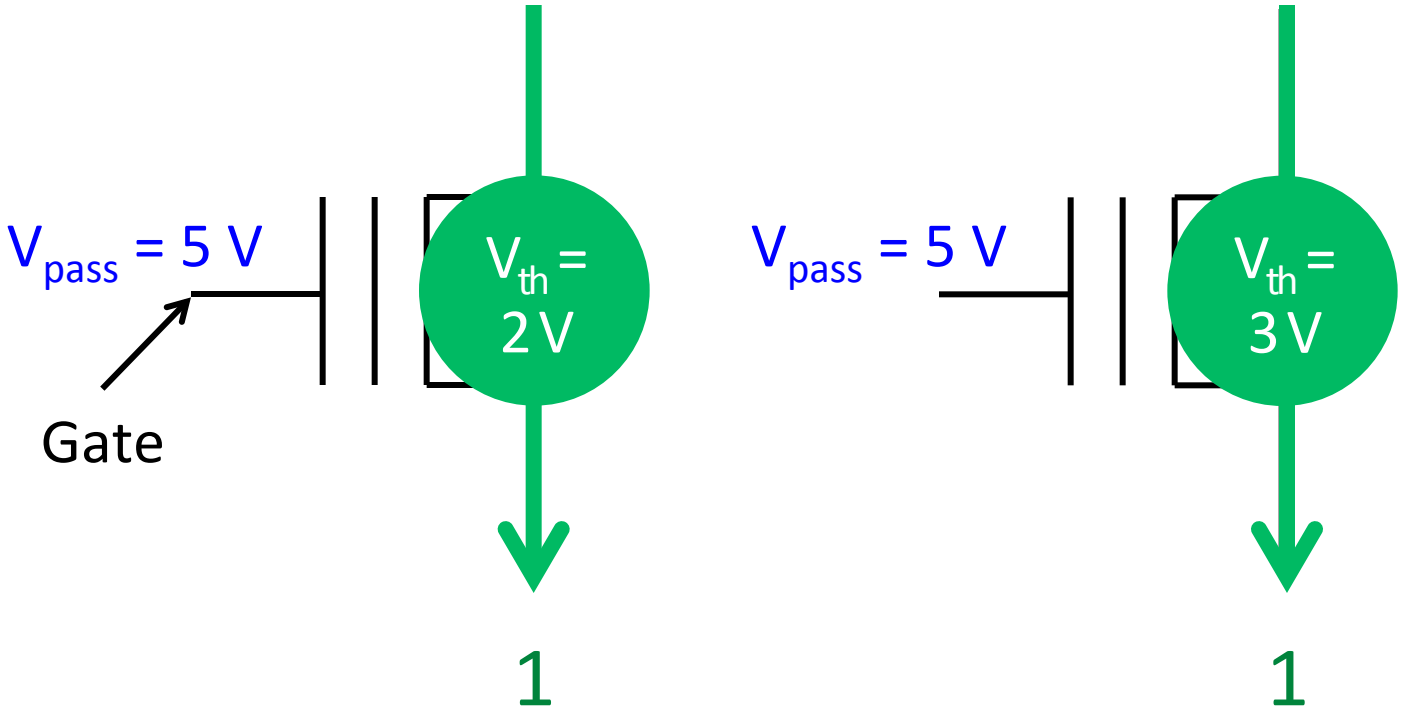

Read from Flash Cell Array

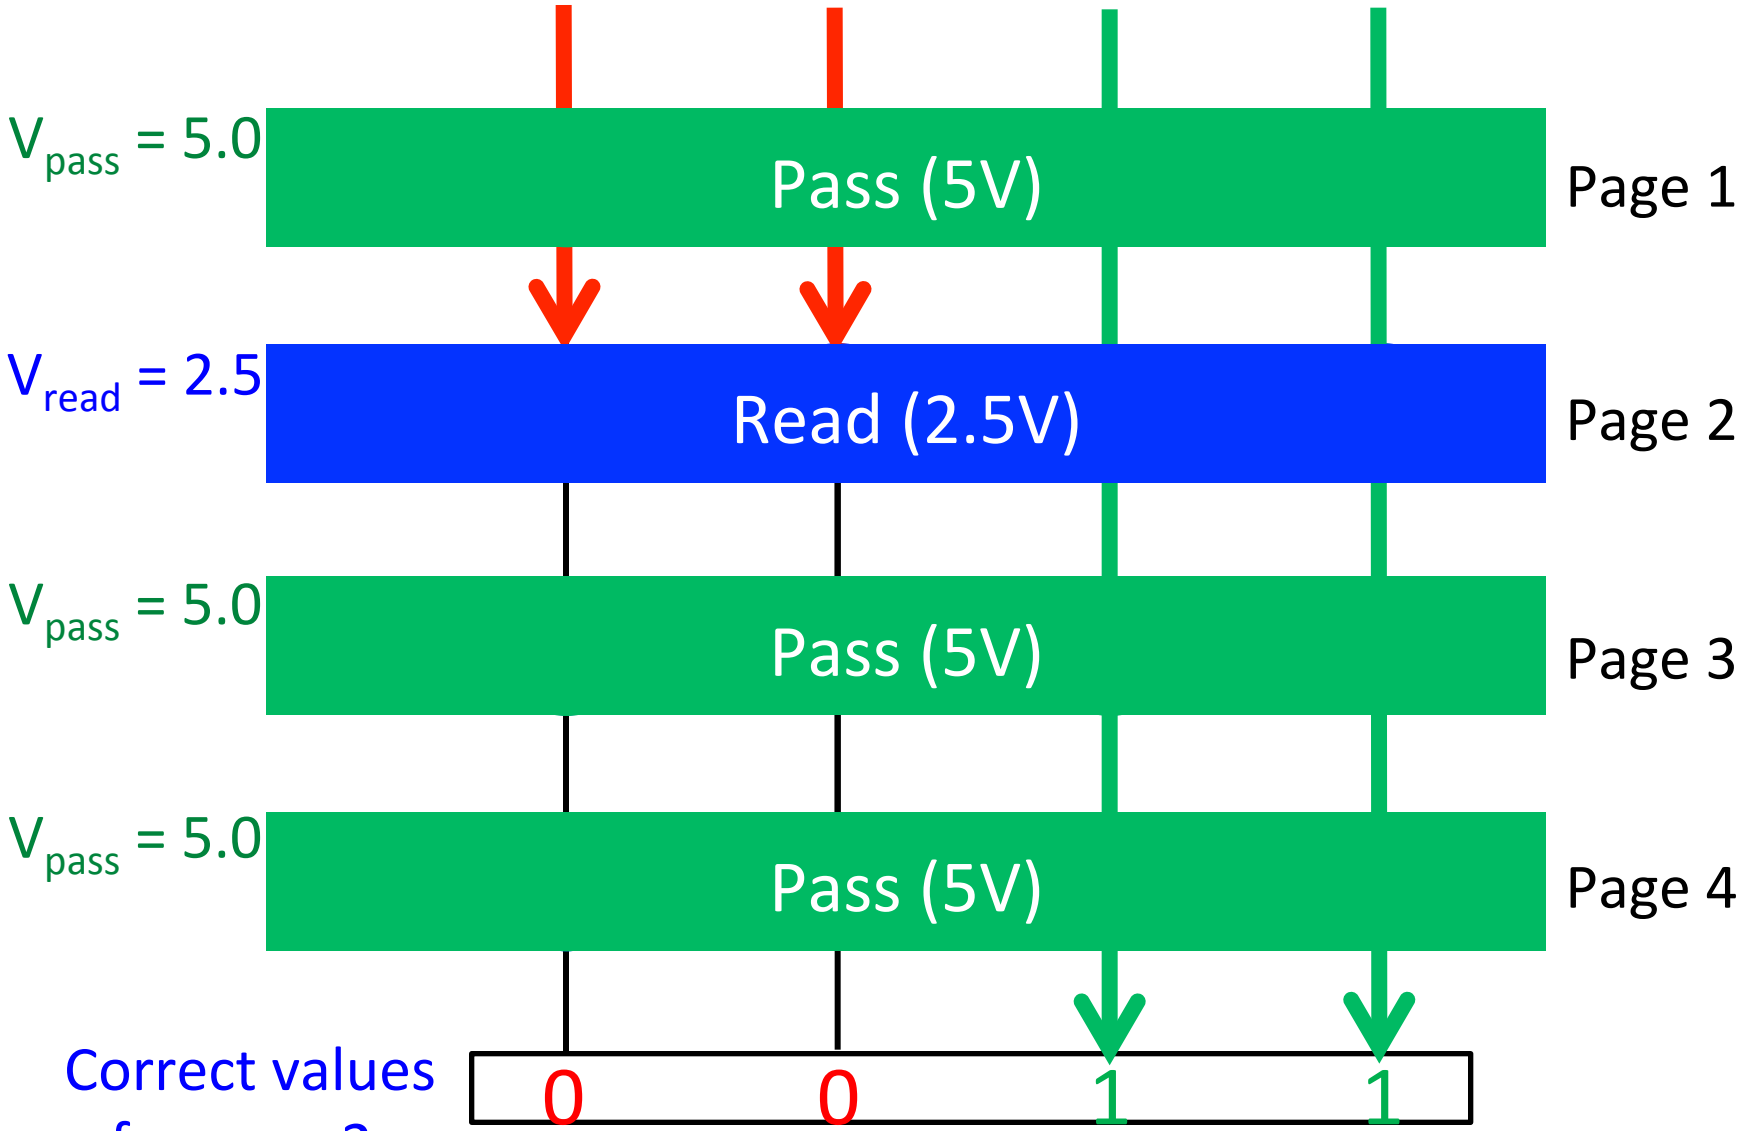

Read Disturb Problem: "Weak Programming" Effect

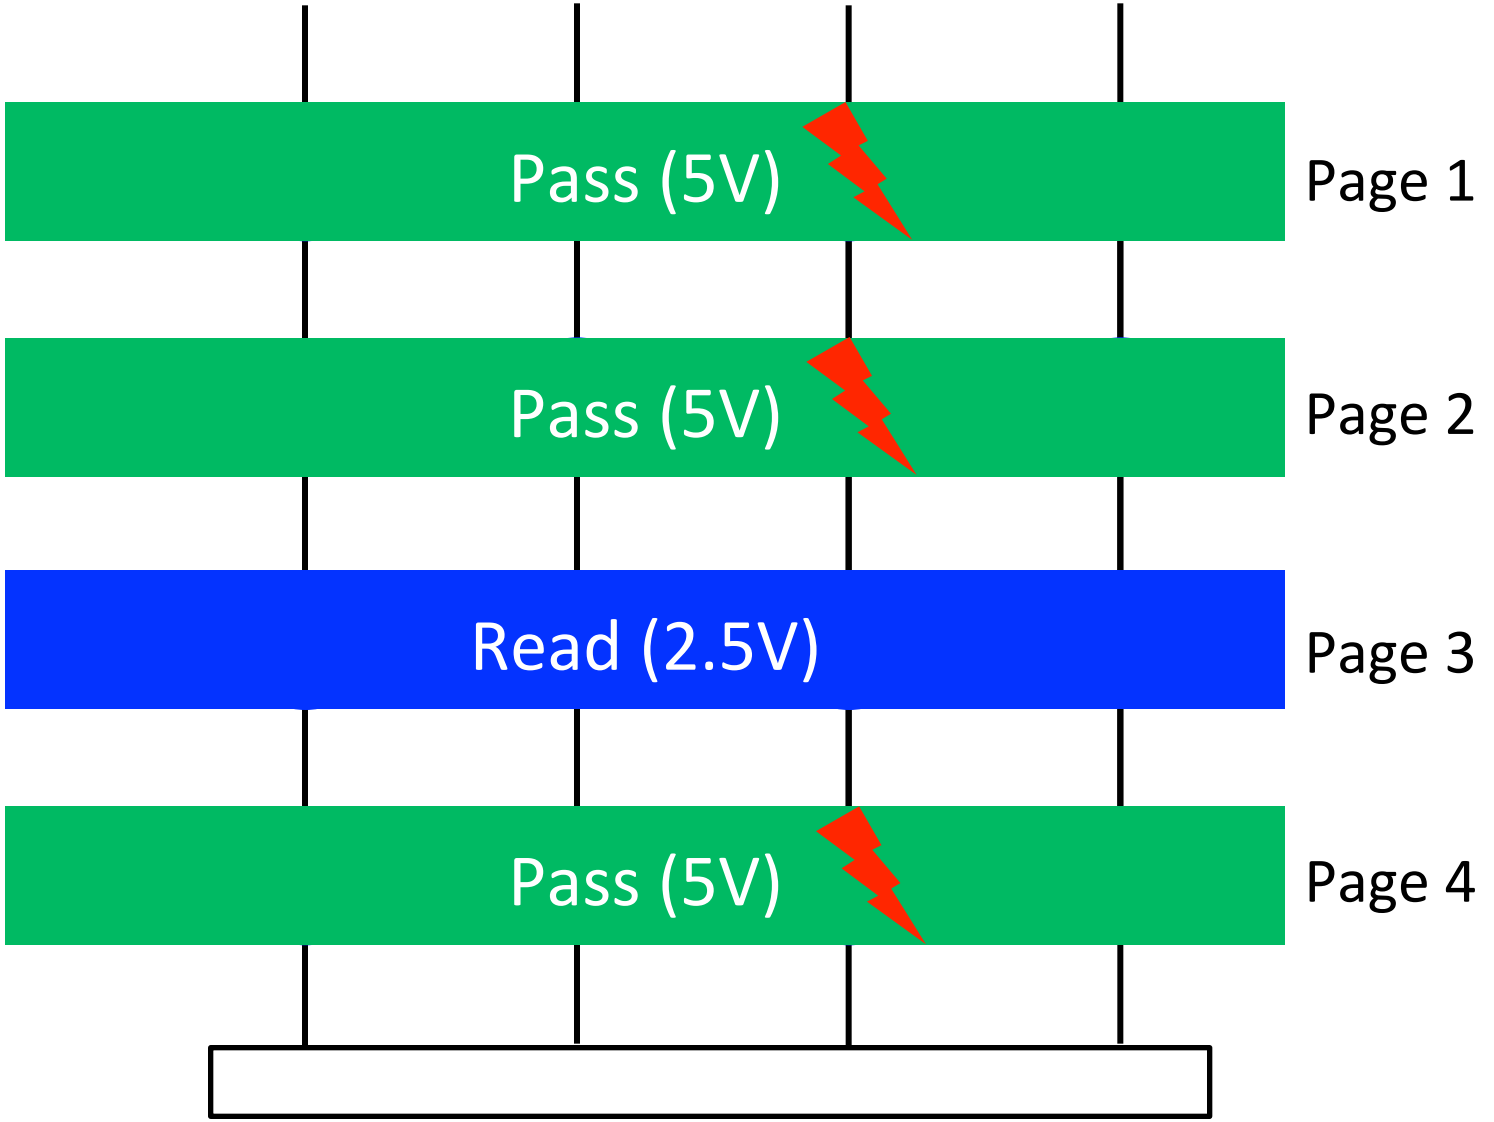

Repeatedly read page 3 (or any page other than page 2)

Read Disturb Problem: "Weak Programming" Effect

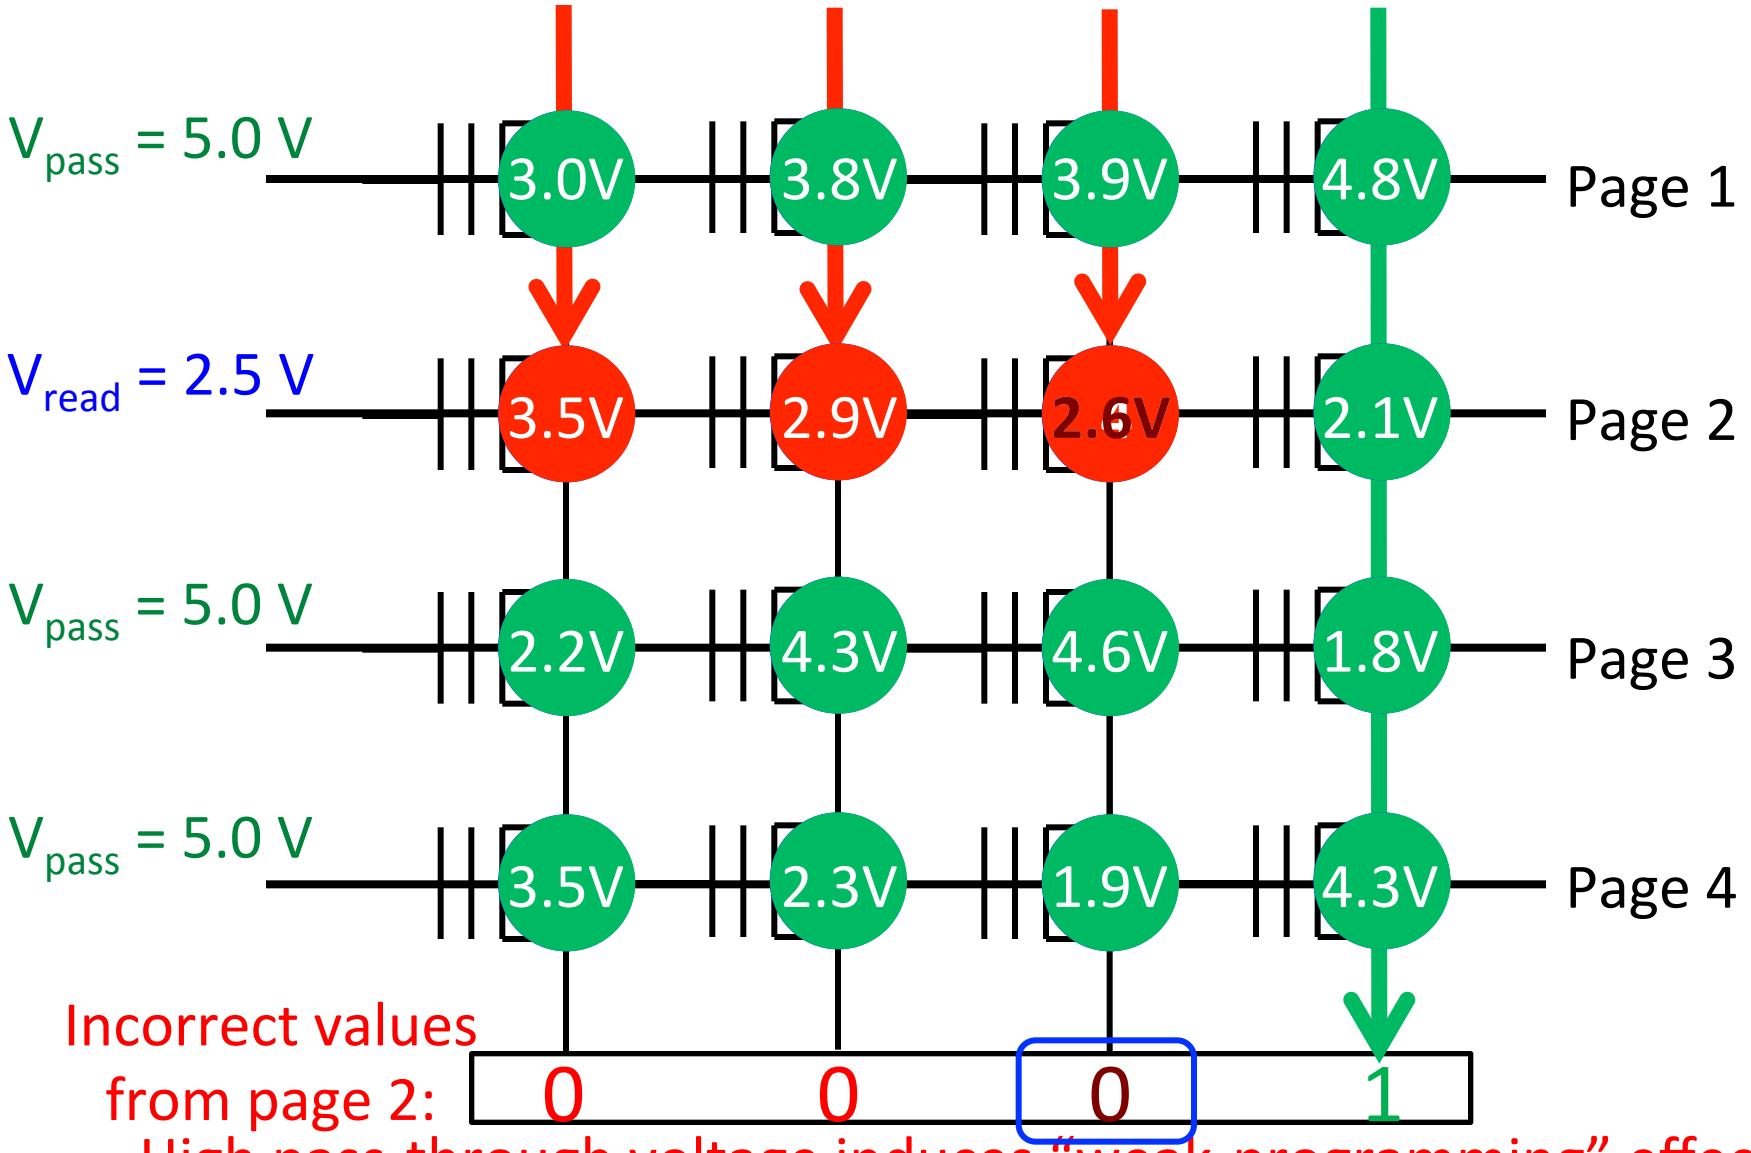

High pass-through voltage induces "weak-programming" effect

Read disturb errors: Reading from one page can alter the values stored in other unread pages

Goal: Mitigate and Recover Read Disturb Errors

Outline

- Background (Problem and Goal)
- Key Experimental Observations
- Mitigation: V_{pass} Tuning
- Recovery: Read Disturb Oriented Error Recovery
- Conclusion

Methodology

- FPGA-based flash memory testing platform [Cai+, FCCM '11]

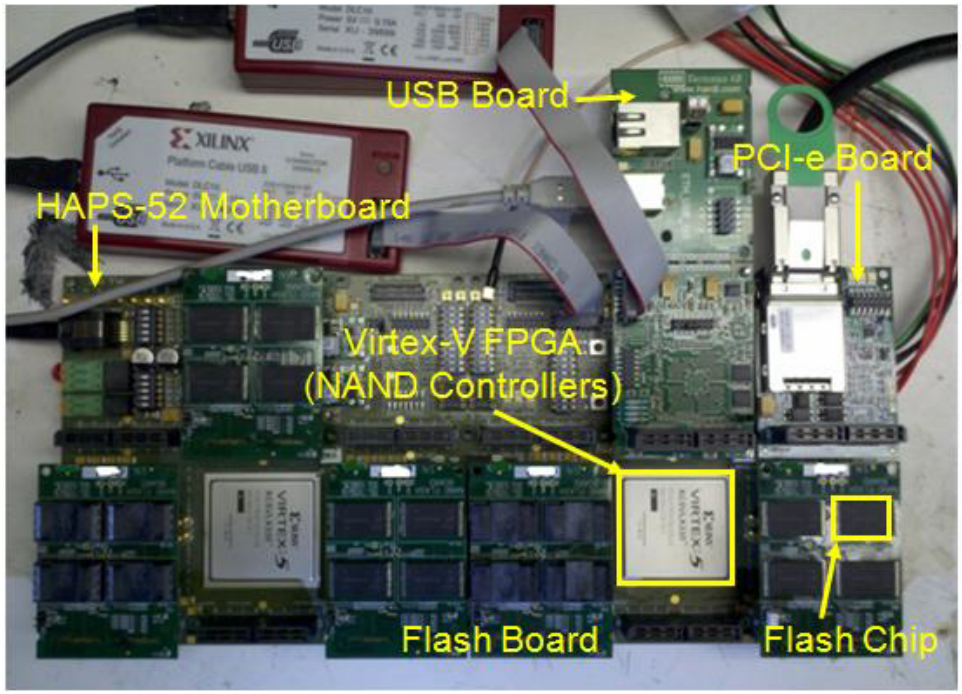

- Real 20- to 24-nm MLC NAND flash chips
- 0 to 1M read disturbs
- 0 to 15K Program/Erase Cycles (PEC)

Experimental Infrastructure

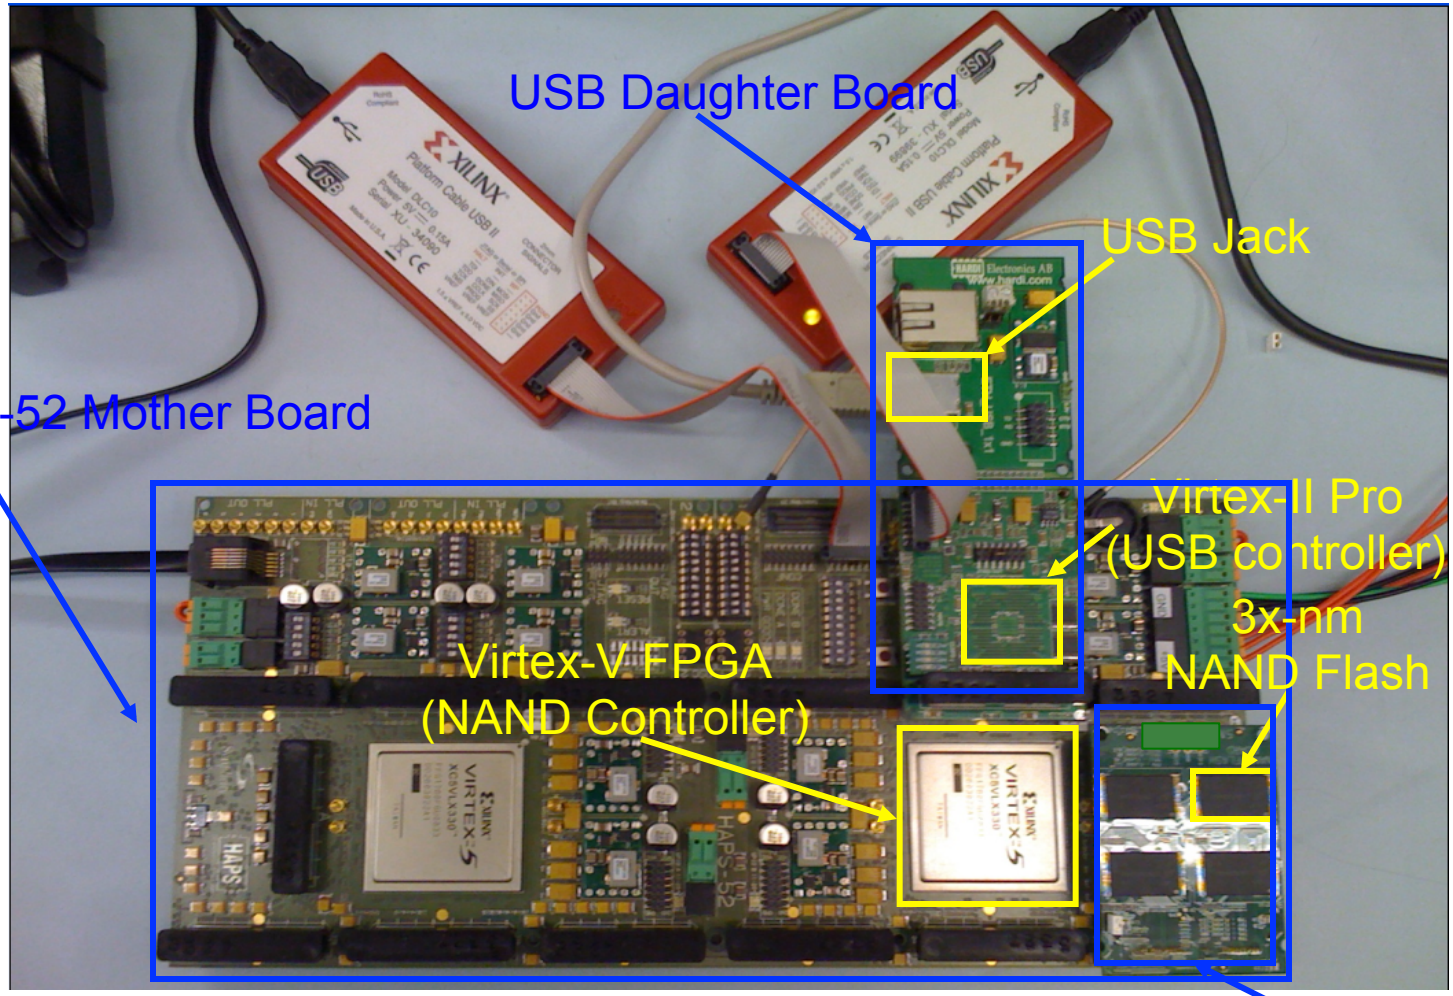

[Cai+, DATE 2012, ICCD 2012, DATE 2013, ITJ 2013, ICCD 2013, SIGMETRICS 2014, HPCA 2015, DSN 2015, MSST 2015]

Read Disturb Effect on V_{th} Distribution

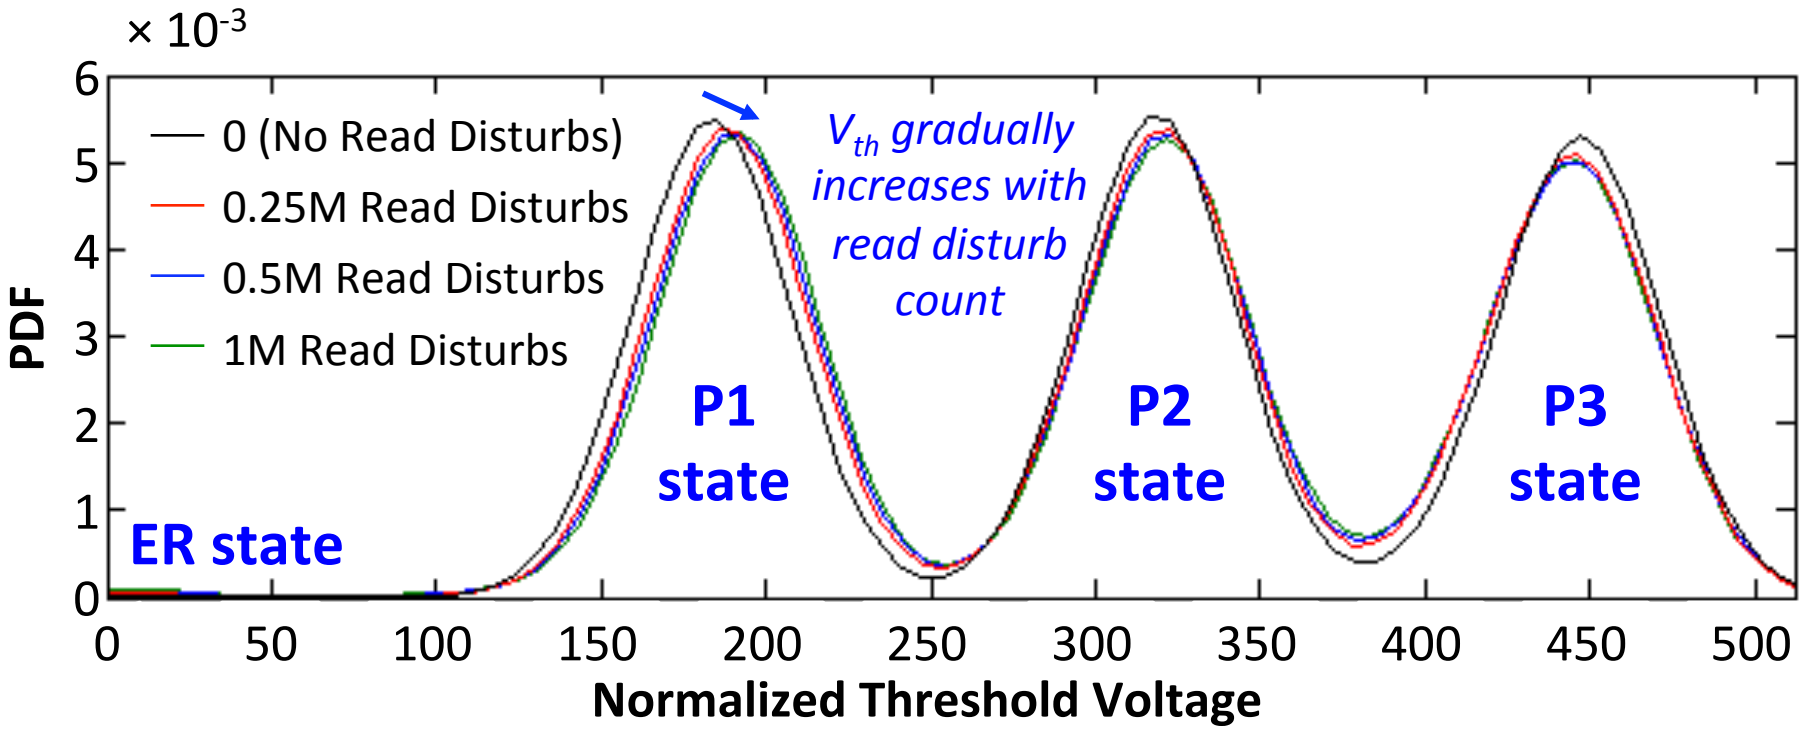

Other Experimental Observations

- Lower threshold voltage states are affected more by read disturb
- Wear-out increases read disturb effect

Key Observation 1: Slightly lowering V_{pass} greatly reduces read disturb errors

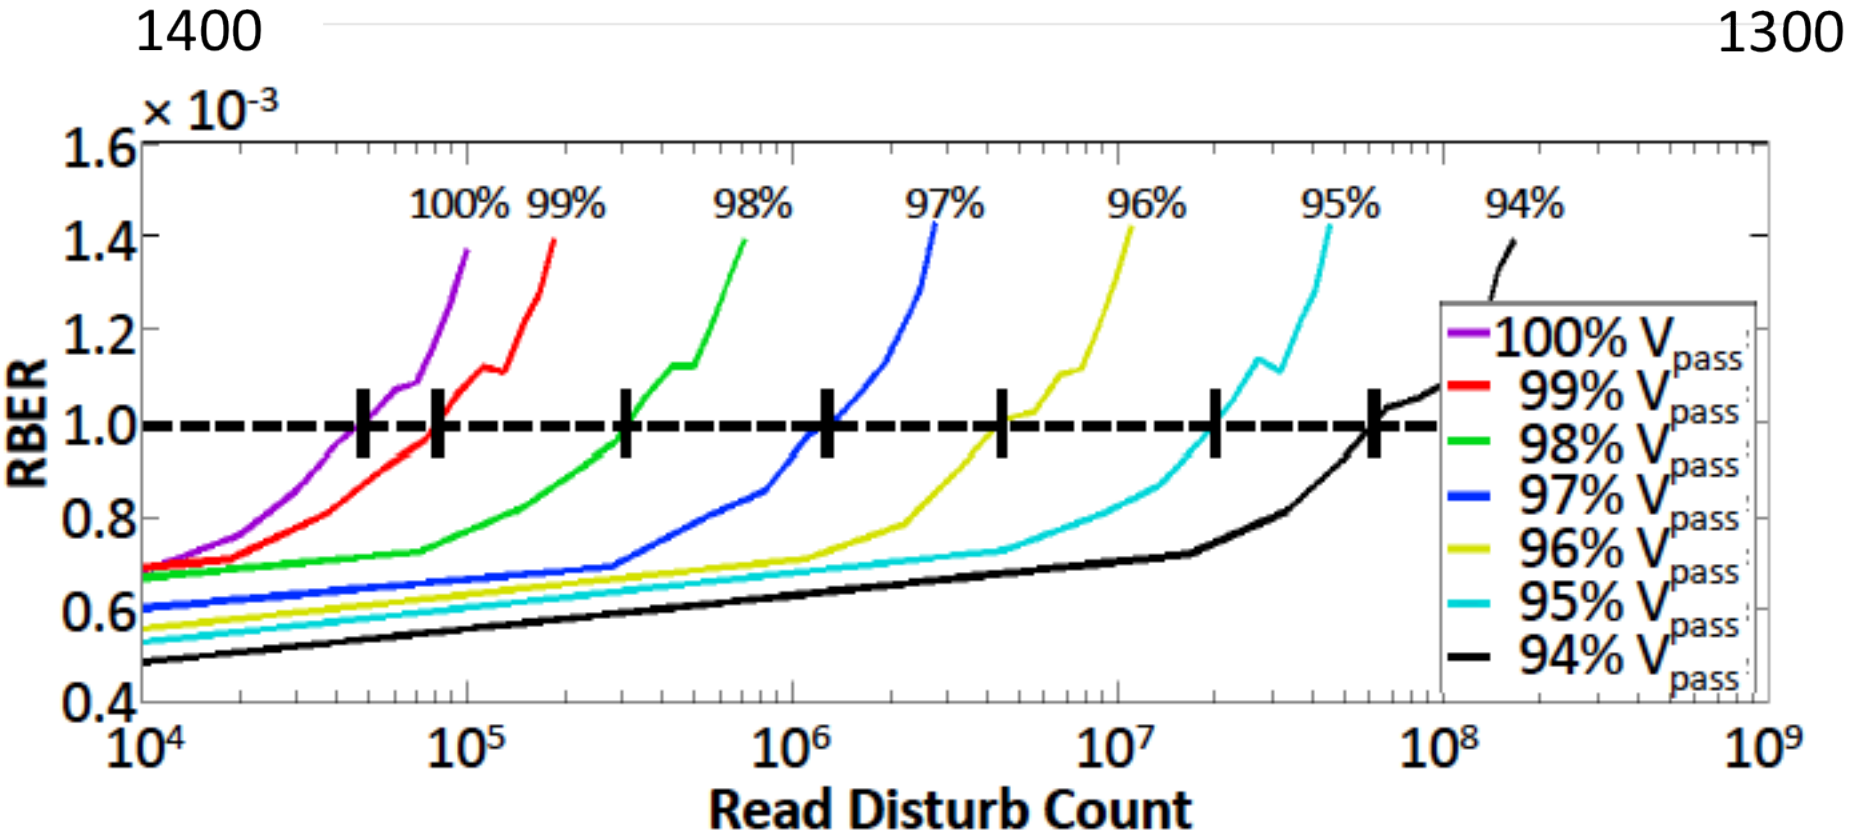

Fig. 11. Raw bit error rate vs. read disturb count for different V_{pass} values, for flash memory under 8K P/E cycles of wear.

Percentage of Vpass Reduction

Outline

- Background (Problem and Goal)
- Key Experimental Observations
- Mitigation: V_{pass} Tuning
- Recovery: Read Disturb Oriented Error Recovery
- Conclusion

Read Disturb Mitigation: V_{pass} Tuning

- Key Idea: Dynamically find and apply a lowered V_{pass}

- Trade-off for lowering V_{pass}
  - + Allows more read disturbs
  - – Induces more read errors

Read Errors Induced by V_{pass} Reduction

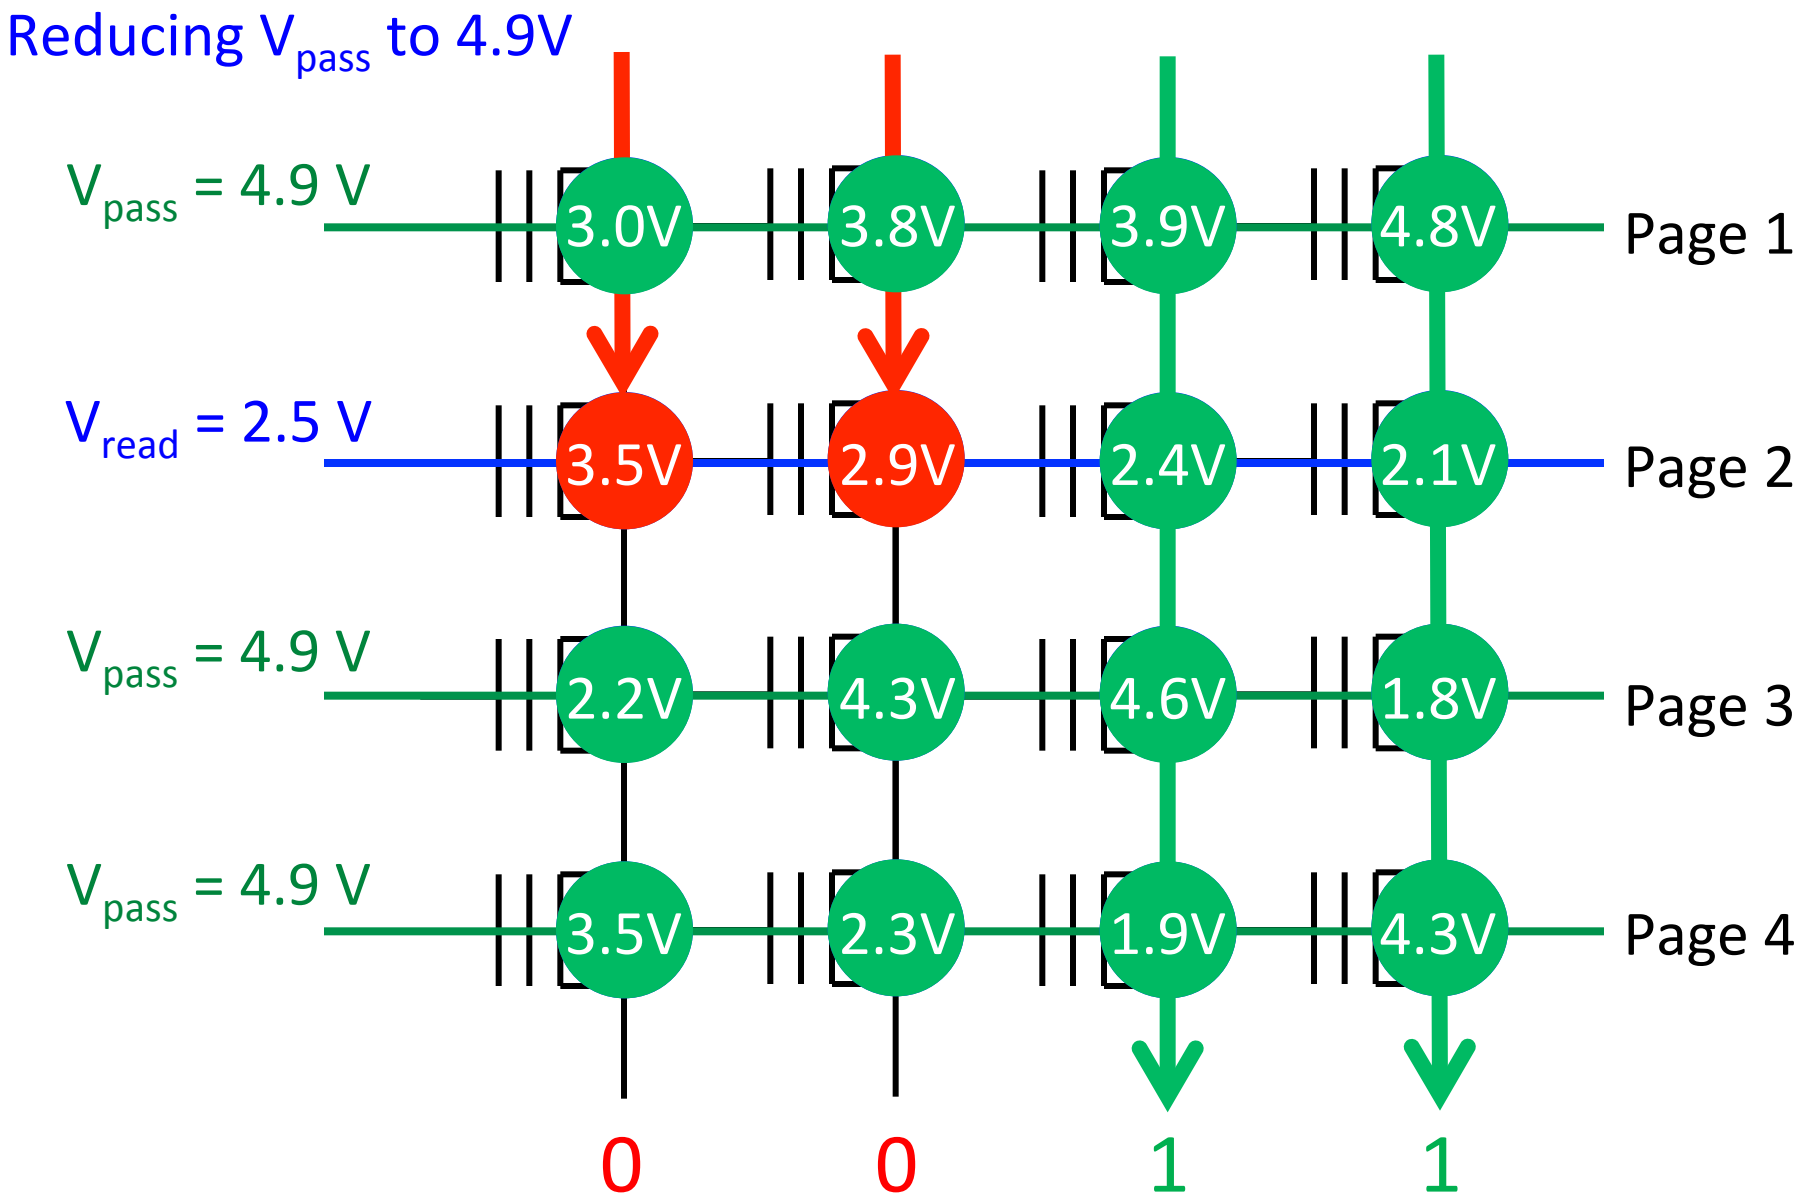

Read Errors Induced by V_{pass} Reduction

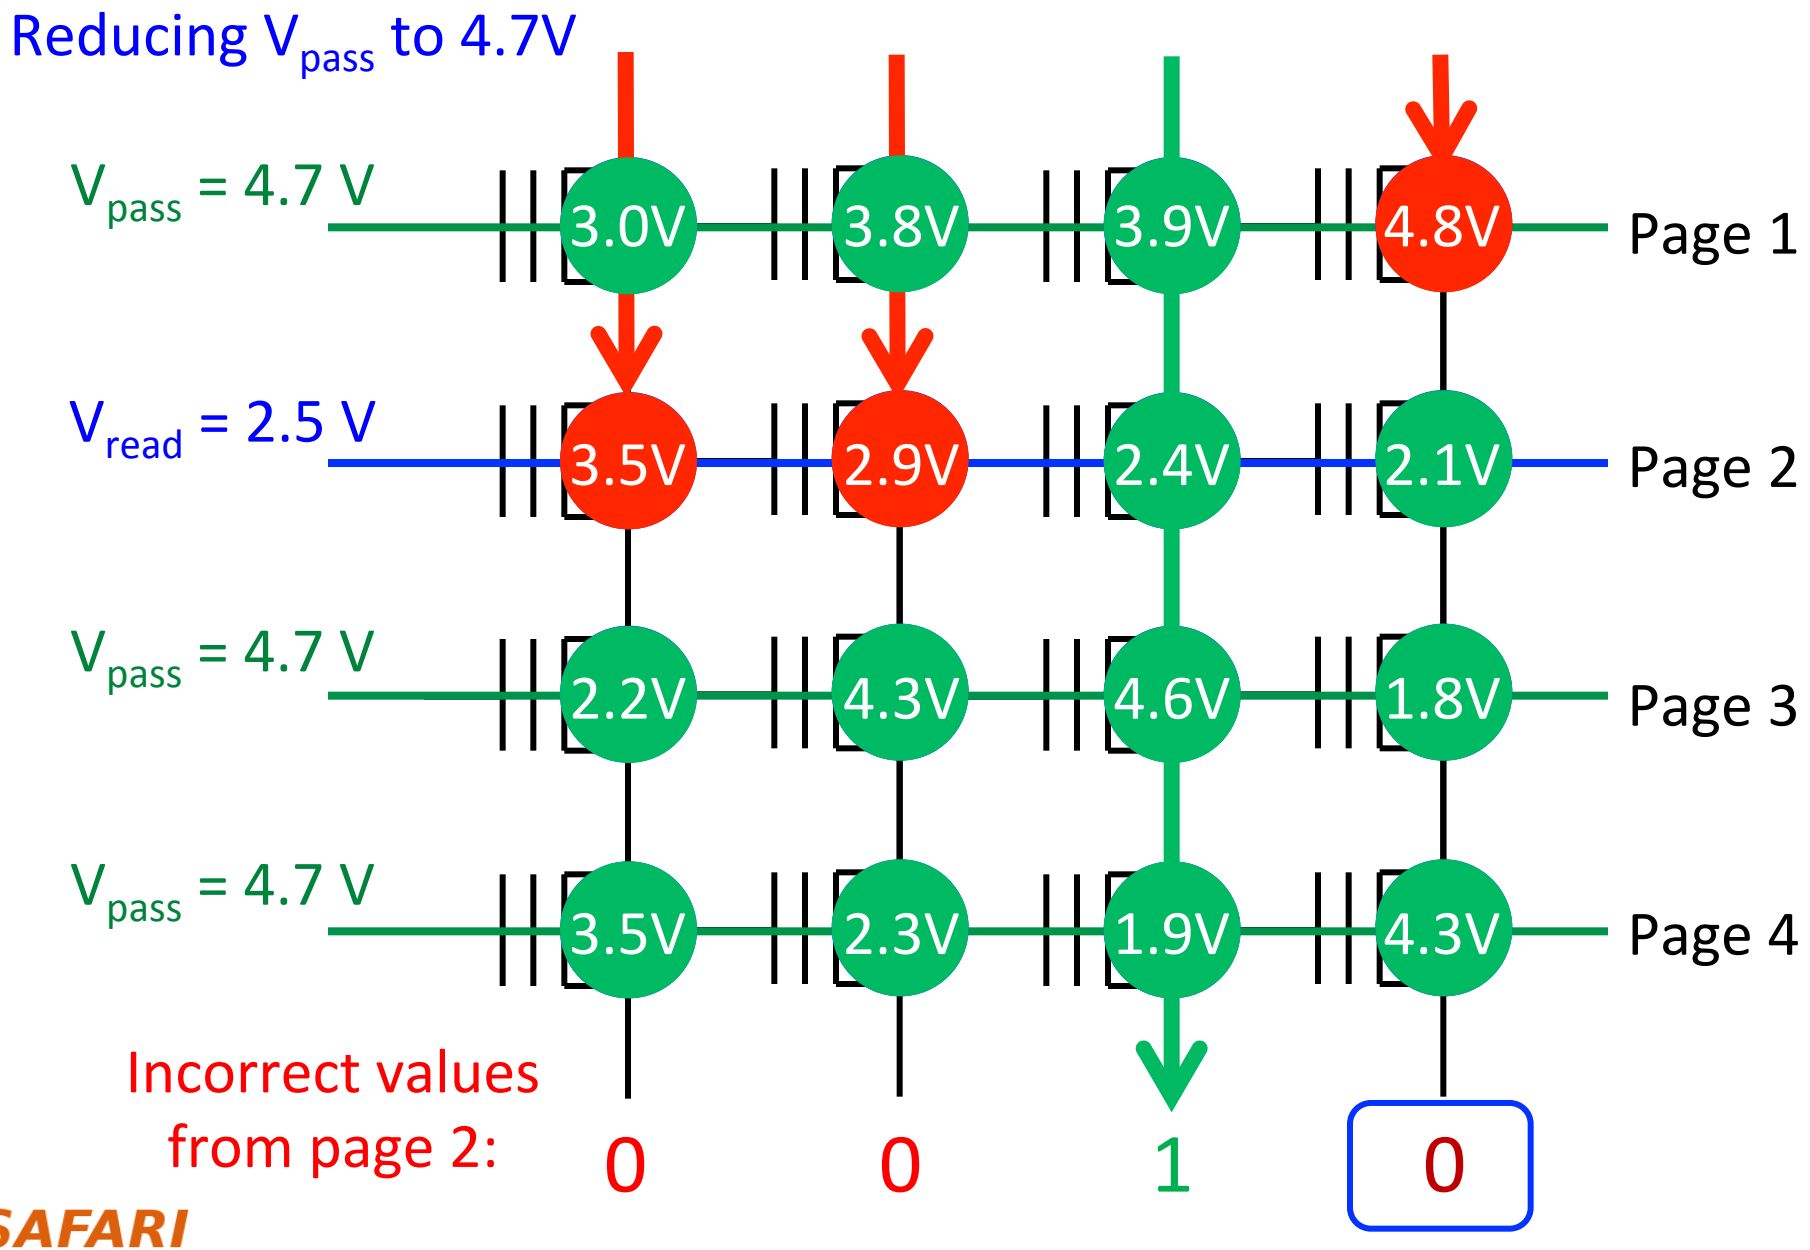

Utilizing the Unused ECC Capability

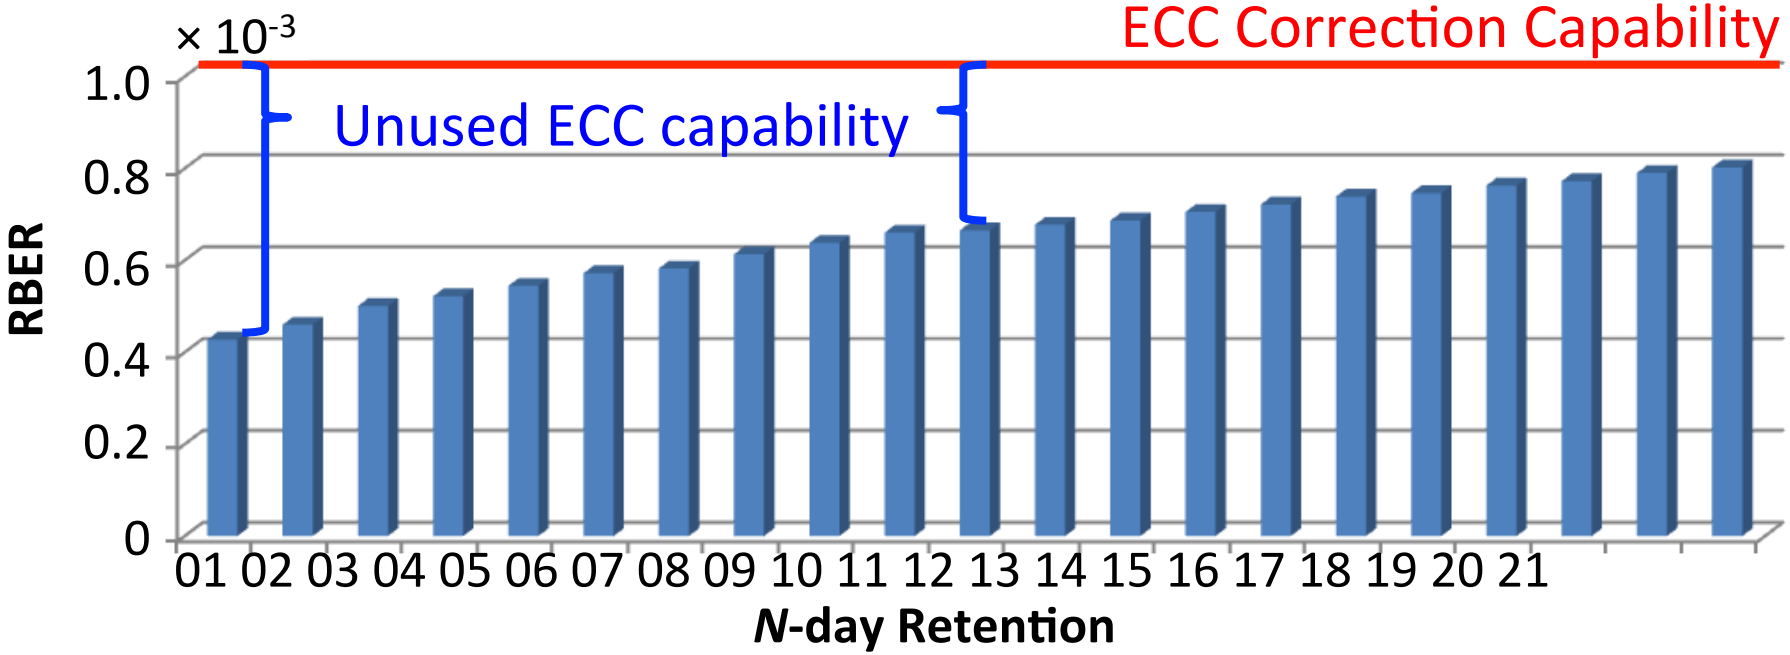

1. ECC provisioned for high retention "age"

2. Unused ECC capability can be used to fix read errors

3. Unused ECC capability decreases over retention age

Dynamically adjust V_{pass} so that read errors fully utilize the unused ECC capability

V_{pass} Reduction Trade-Off Summary

- Today: Conservatively set V_{pass} to a high voltage
  - – Accumulates more read disturb errors at the end of each refresh interval
  - + No read errors

- Idea: Dynamically adjust V_{pass} to unused ECC capability
  - + Minimize read disturb errors
  - ○ Control read errors to be tolerable by ECC
  - ○ If read errors exceed ECC capability, read again with a higher V_{pass} to correct read errors

V_{pass} Tuning Steps

- Perform once for each block every day:
  1. Estimate *unused ECC capability (using retention age)*
  2. Aggressively reduce V_{pass} until *read errors exceeds ECC capability*
  3. Gradually increase V_{pass} until read error becomes just less than ECC capability

Evaluation of V_{pass} Tuning

- 19 real workload I/O traces
- Assume 7-day refresh period
- Similar methodology as before to determine acceptable V_{pass} reduction

- Overhead for a 512 GB flash drive:
  - 128 KB storage overhead for per-block V_{pass} setting and worst-case page
  - 24.34 sec/day average V_{pass} Tuning overhead

V_{pass} Tuning Lifetime Improvements

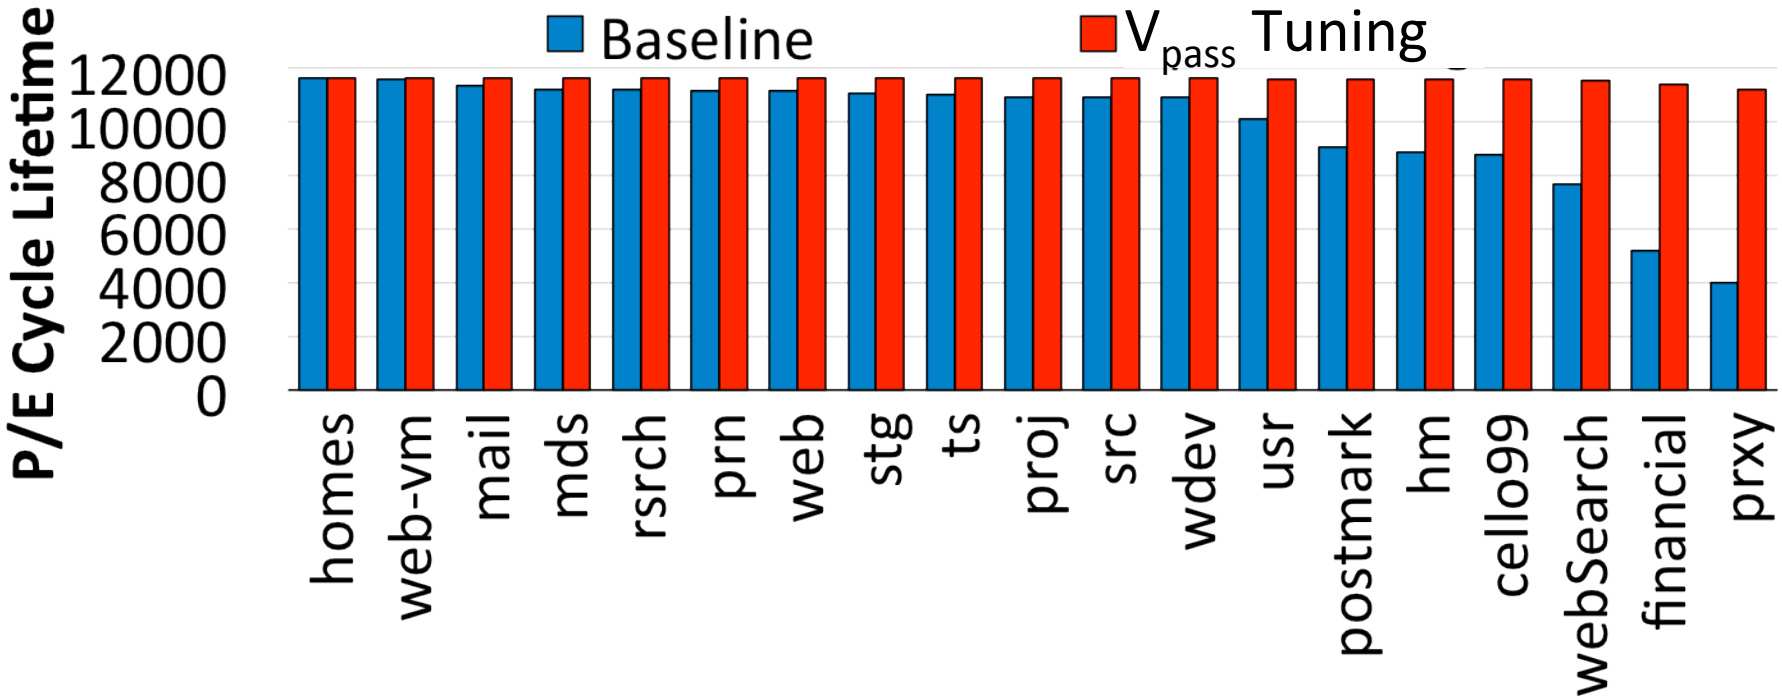

Average lifetime improvement: 21.0%

Outline

- Background (Problem and Goal)
- Key Experimental Observations
- Mitigation: V_{pass} Tuning
- Recovery: Read Disturb Oriented Error Recovery
- Conclusion

Read Disturb Resistance

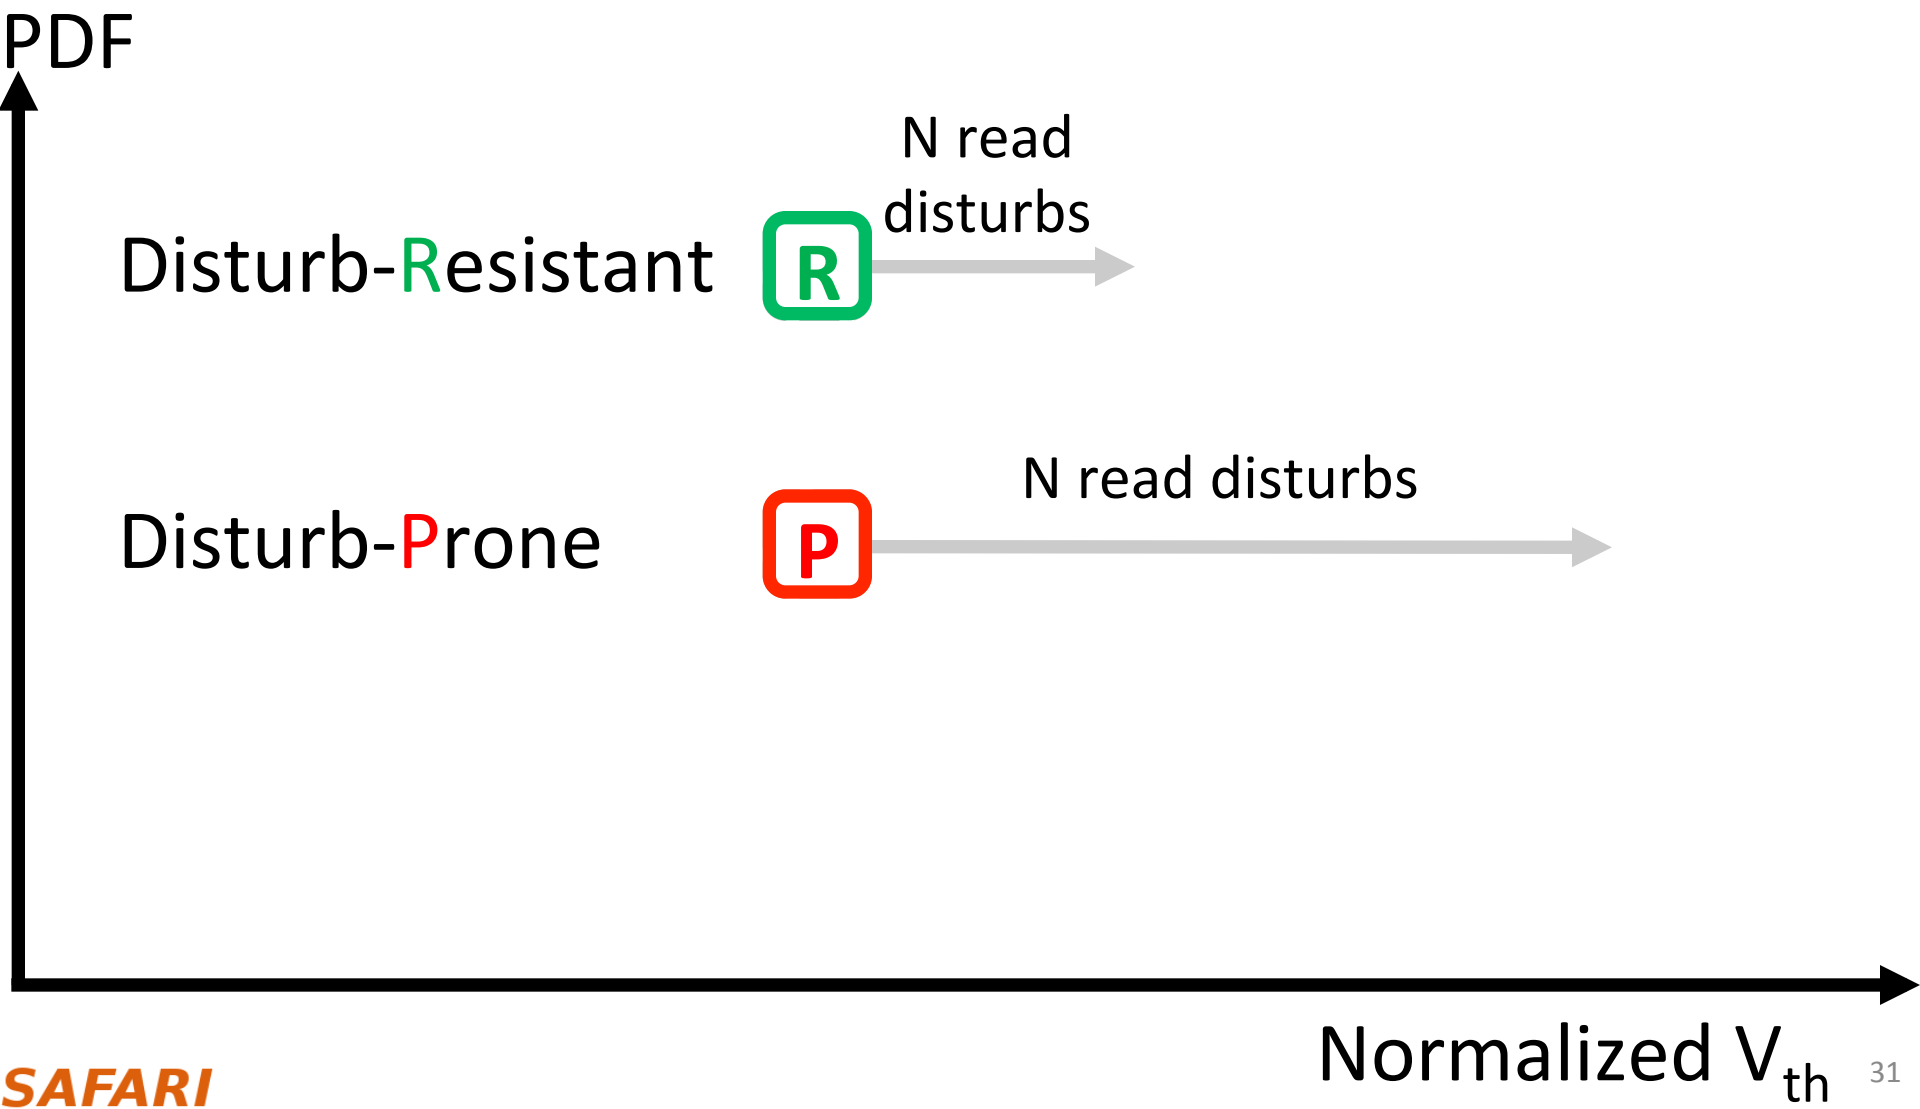

Observation 2: Some Flash Cells Are More Prone to Read Disturb

After 250K read disturb:

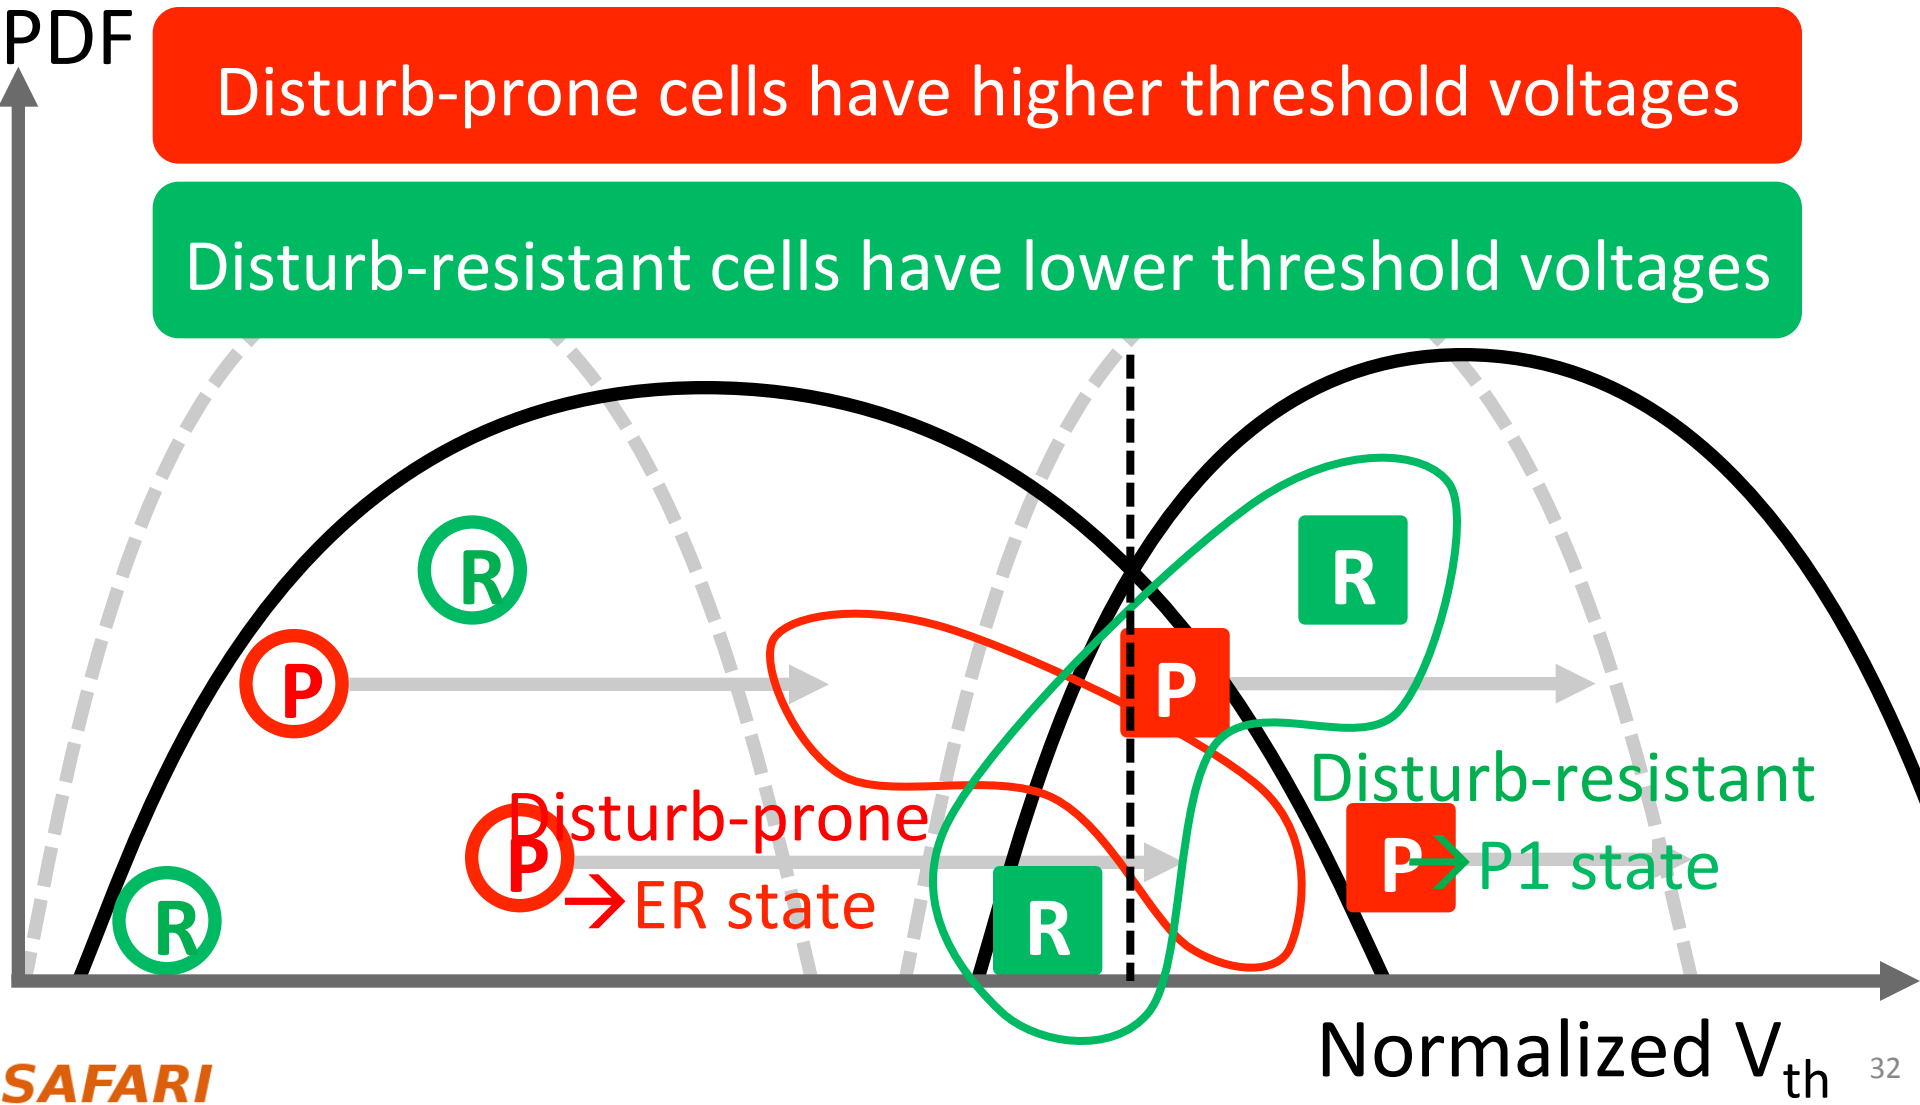

Read Disturb Oriented Error Recovery (RDR)

- Triggered by an uncorrectable flash error
  - Back up all valid data in the faulty block
  - Disturb the faulty page 100K times (more)
  - Compare V_{th}'s before and after read disturb
  - Select cells susceptible to flash errors ($V_{ref}-\sigma<V_{th}<V_{ref}-\sigma$)
  - Predict among these susceptible cells
    - Cells with more V_{th} shifts are disturb-prone → Higher V_{th} state
    - Cells with less V_{th} shifts are disturb-resistant → Lower V_{th} state

RDR Evaluation

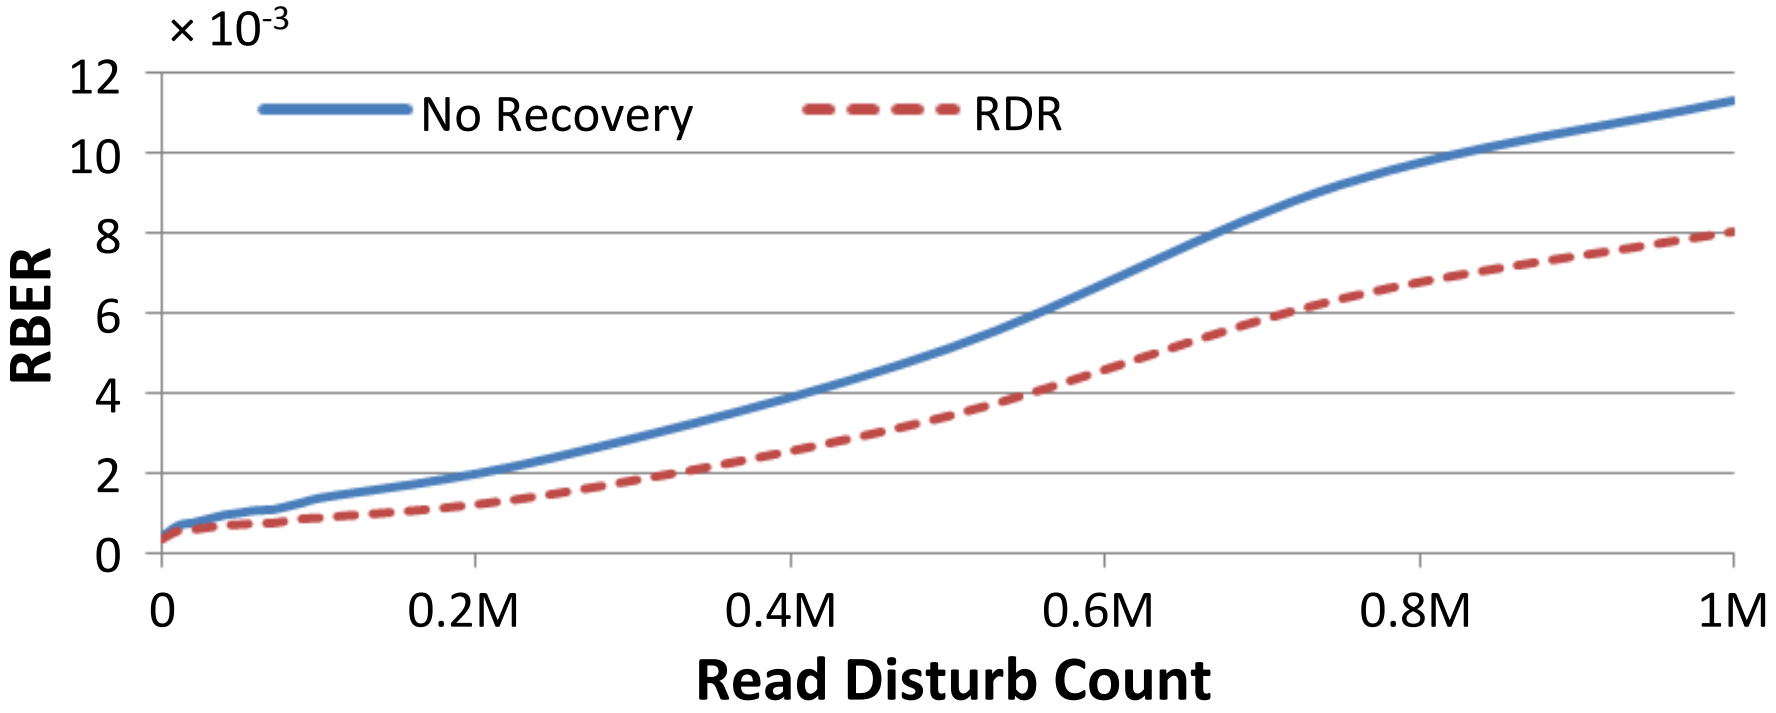

Reduces total error counts by up to 36% @ 1M read disturbs
ECC can be used to correct the remaining errors

Outline

- Background (Problem and Goal)
- Key Experimental Observations
- Mitigation: V_{pass} Tuning
- Recovery: Read Disturb Oriented Error Recovery
- Conclusion

Executive Summary

- ***Read disturb errors*** limit flash memory lifetime today
  - Apply a *high pass-through voltage (V_{pass})* to multiple pages on a read
  - Repeated application of V_{pass} can alter stored values in unread pages
- We **characterize read disturb** on real NAND flash chips
  - Slightly lowering V_{pass} greatly reduces read disturb errors
  - Some flash cells are more prone to read disturb
- **Technique 1:** Mitigate read disturb errors online
  - ***V_{pass} Tuning*** dynamically finds and applies a lowered V_{pass} per block
  - Flash memory lifetime improves by 21%
- **Technique 2:** Recover after failure to prevent data loss
  - ***Read Disturb Oriented Error Recovery*** (RDR) selectively corrects cells more susceptible to read disturb errors
  - Reduces raw bit error rate (RBER) by up to 36%

Two Other Recent Works on Flash Errors

- **"Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery"**
  - **A study of Data Retention Mechanisms, Errors, and Correction Mechanisms**
  - **Will be presented by Yixin Luo (CMU PhD Student) tomorrow morning @ 8:30am**
  - **Forum E-31: Flash Controller Design Options**

- **"A Large-Scale Study of Flash Memory Errors in the Field"**
  - **Study of flash-based SSD errors in Facebook data centers over the course of 4 years**
  - **First large-scale field study of flash memory reliability**

And, A Work on Read Disturb in DRAM

- **RowHammer**

- **"Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors"**
  - **Most modern DRAM chips are vulnerable to read disturb errors**
  - **These errors have slipped into the field, creating new security vulnerabilities**
    - Exploiting the DRAM rowhammer bug to gain kernel privileges (Seaborn, March 2015)

- **"The DRAM RowHammer Problem (and Its Reliability and Security Implications)"**

Referenced Papers and Talks

- All are available at
  http://users.ece.cmu.edu/~omutlu/projects.htm
  http://users.ece.cmu.edu/~omutlu/talks.htm

- And, many other previous works on NAND flash memory errors and management

Flash Memory SUMMIT

Thank you.

Feel free to email me with any questions & feedback

onur@cmu.edu
http://users.ece.cmu.edu/~omutlu/

Read Disturb Errors in MLC NAND Flash Memory

Characterization, Mitigation & Recovery

Onur Mutlu
onur@cmu.edu

(joint work with Yu Cai, Yixin Luo, Saugata Ghose, Erich Haratsch, Ken Mai)

August 12, 2015

Flash Memory Summit 2015, Santa Clara, CA

SAFARI

Carnegie Mellon

SEAGATE

DSSC

References to Papers and Talks

Our FMS Talks and Posters

- *Onur Mutlu, Error Analysis and Management for MLC NAND Flash Memory, FMS 2014.*
- *Onur Mutlu, Read Disturb Errors in MLC NAND Flash Memory, FMS 2015.*
- *Yixin Luo, Data Retention in MLC NAND Flash Memory, FMS 2015.*

- *FMS 2015 posters:*
  - WARM: Improving NAND Flash Memory Lifetime with Write-hotness Aware Retention Management
  - Read Disturb Errors in MLC NAND Flash Memory
  - Data Retention in MLC NAND Flash Memory

Our Flash Memory Works (I)

1. Retention noise study and management

1) 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.

2) 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.

3) 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.

2. Flash-based SSD prototyping and testing platform

4) Yu Cai, Erich F. Haratsh, Mark McCartney, Ken Mai,
FPGA-based solid-state drive prototyping platform, FCCM 2011.

Our Flash Memory Works (II)

3. Overall flash error analysis

5) Yu Cai, Erich F. Haratsch, Onur Mutlu, and Ken Mai,
Error Patterns in MLC NAND Flash Memory: Measurement, Characterization, and Analysis, DATE 2012.

6) 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.

4. Program and erase noise study

7) Yu Cai, Erich F. Haratsch, Onur Mutlu, and Ken Mai,
Threshold Voltage Distribution in MLC NAND Flash Memory: Characterization, Analysis and Modeling, DATE 2013.

Our Flash Memory Works (III)

5. Cell-to-cell interference characterization and tolerance

8) Yu Cai, Onur Mutlu, Erich F. Haratsch, and Ken Mai, **Program Interference in MLC NAND Flash Memory: Characterization, Modeling, and Mitigation**, ICCD 2013.

9) 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.

6. Read disturb noise study

10) 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.

7. Flash errors in the field

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

Referenced Papers and Talks

- All are available at
  http://users.ece.cmu.edu/~omutlu/projects.htm
  http://users.ece.cmu.edu/~omutlu/talks.htm

- And, many other previous works on NAND flash memory errors and management

Related Videos and Course Materials

- Undergraduate Computer Architecture Course Lecture Videos (2013, 2014, 2015)
- Undergraduate Computer Architecture Course Materials (2013, 2014, 2015)
- Graduate Computer Architecture Course Materials (Lecture Videos)
- Parallel Computer Architecture Course Materials (Lecture Videos)
- Memory Systems Short Course Materials (Lecture Video on Main Memory and DRAM Basics)

Flash Memory SUMMIT

Additional Slides on Read Disturb in NAND Flash

Read Disturb Induced RBER Increases Faster with Higher PEC

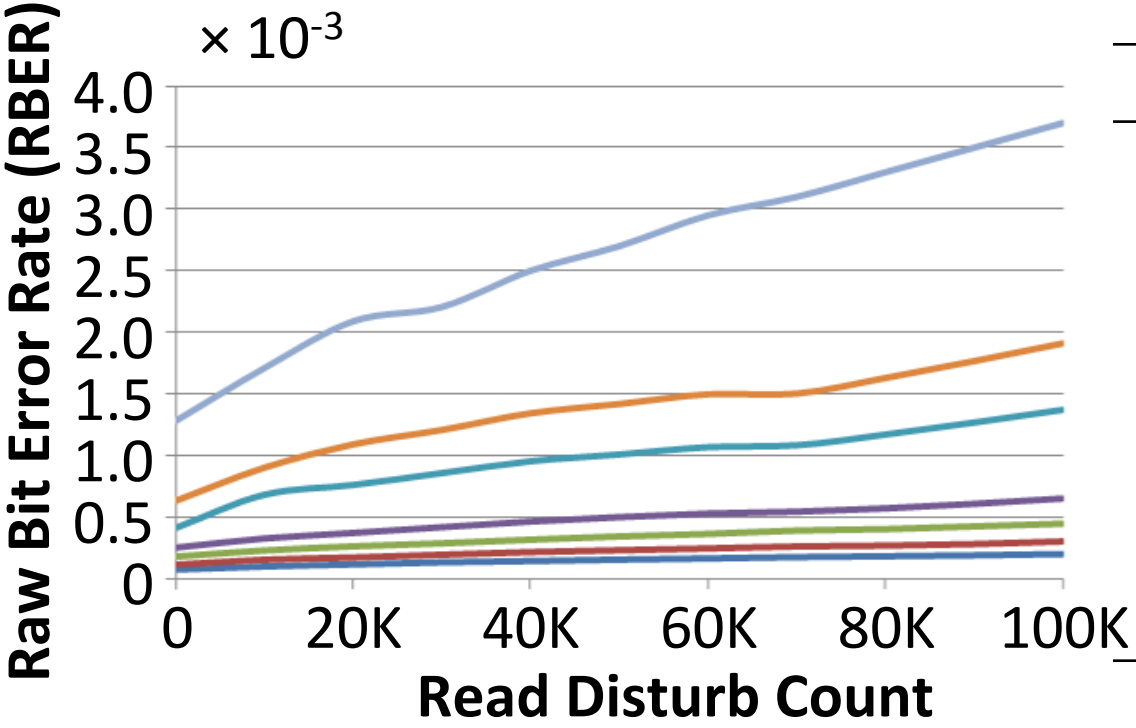

PEC	Slope
15K	1.90×10^{-8}
10K	9.10×10^{-9}
8K	7.50×10^{-9}
5K	3.74×10^{-9}
4K	2.37×10^{-9}
3K	1.63×10^{-9}
2K	1.00×10^{-9}

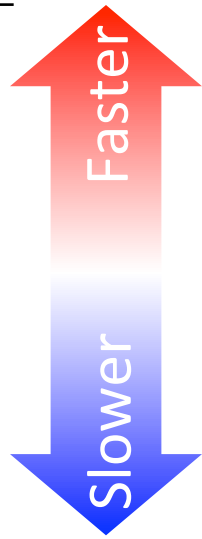

Threshold Voltage Increases with Read Disturb Count

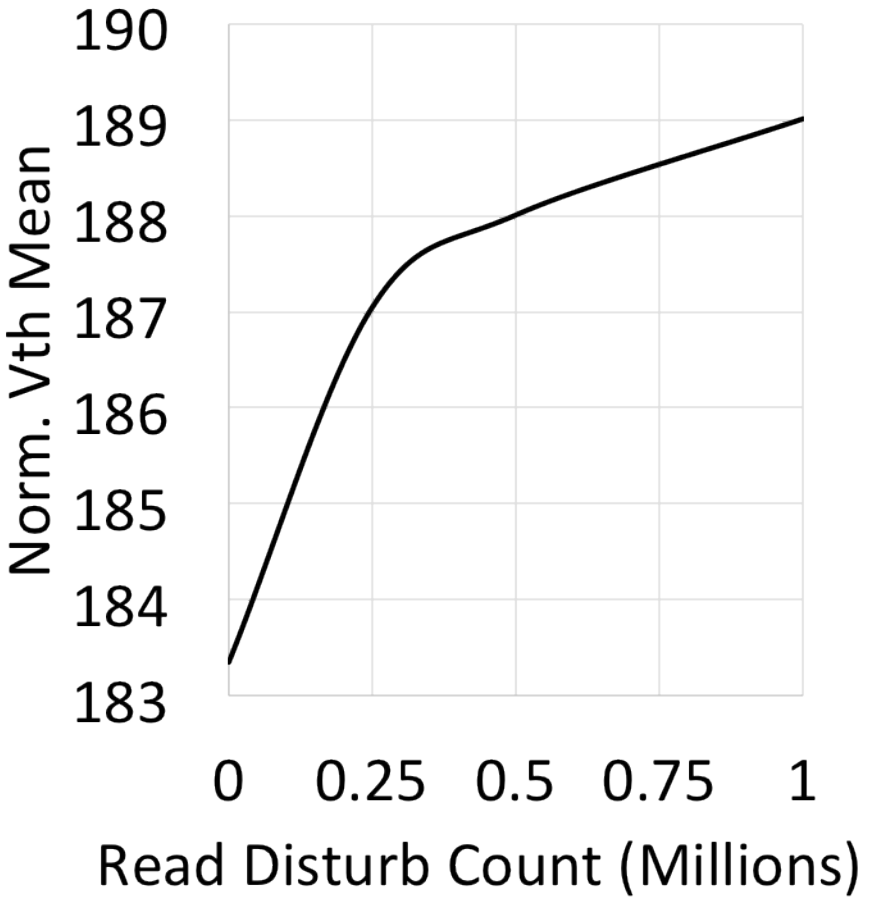

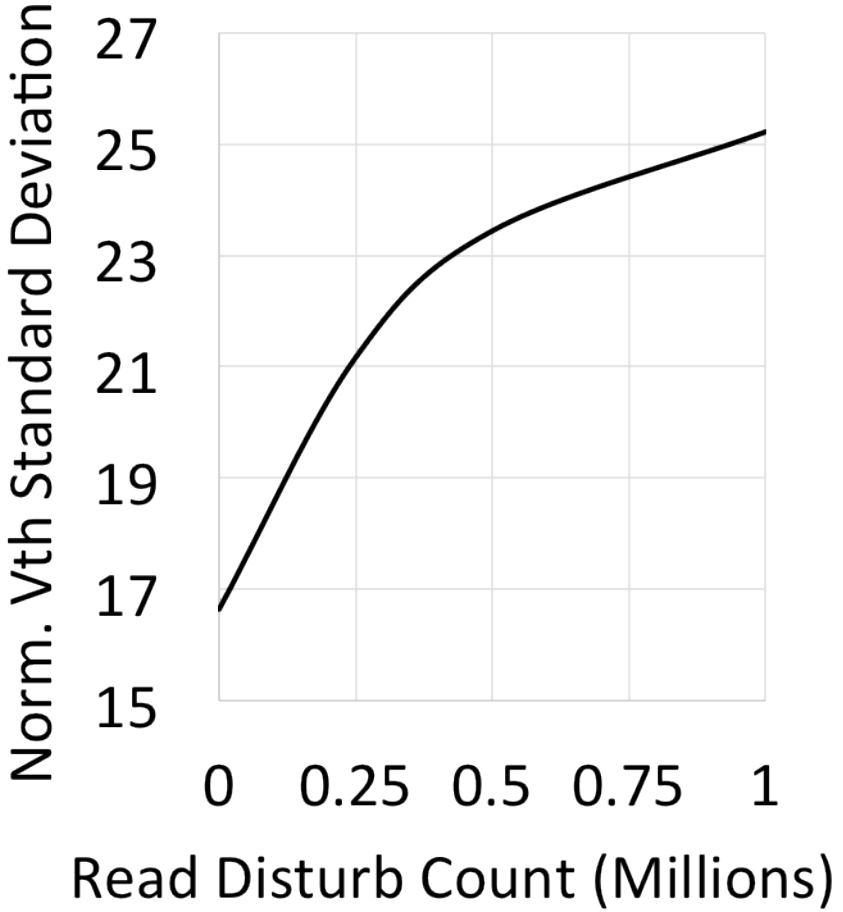

Showing results for P1 state @ 8K PEC, other states have similar trends

Lower Voltage States Are More Prone to Read Disturb

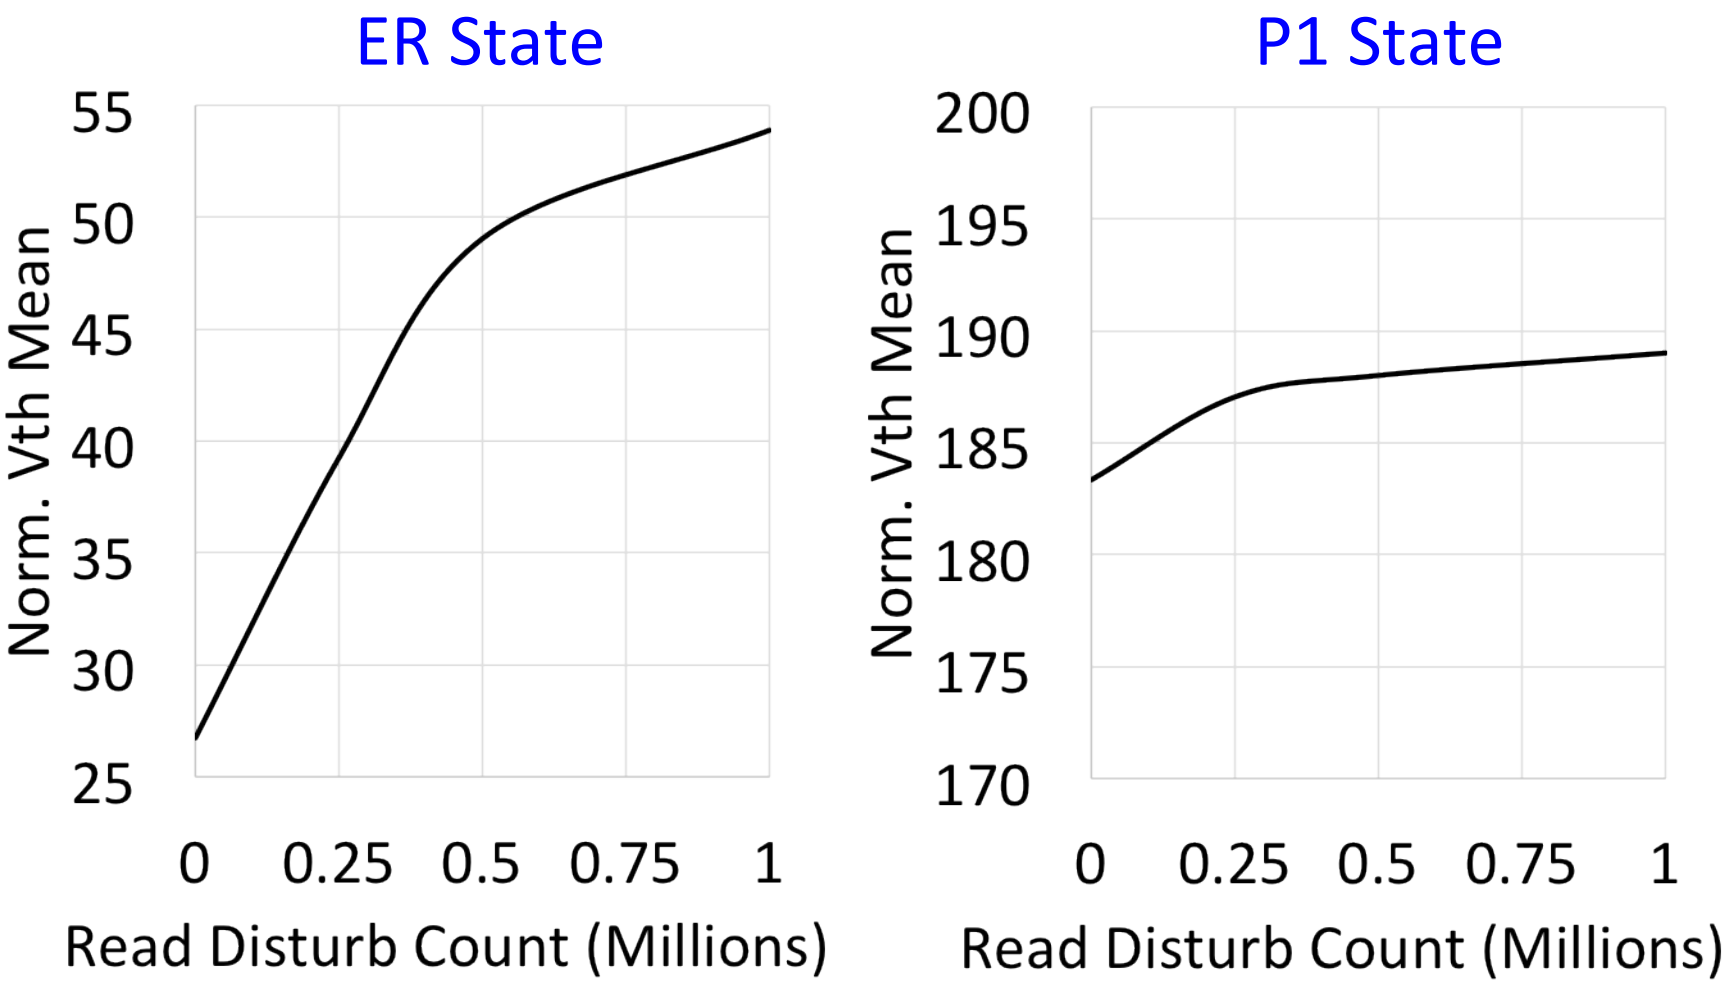

Reducing V_{pass} Increases Tolerable Read Disturb Count

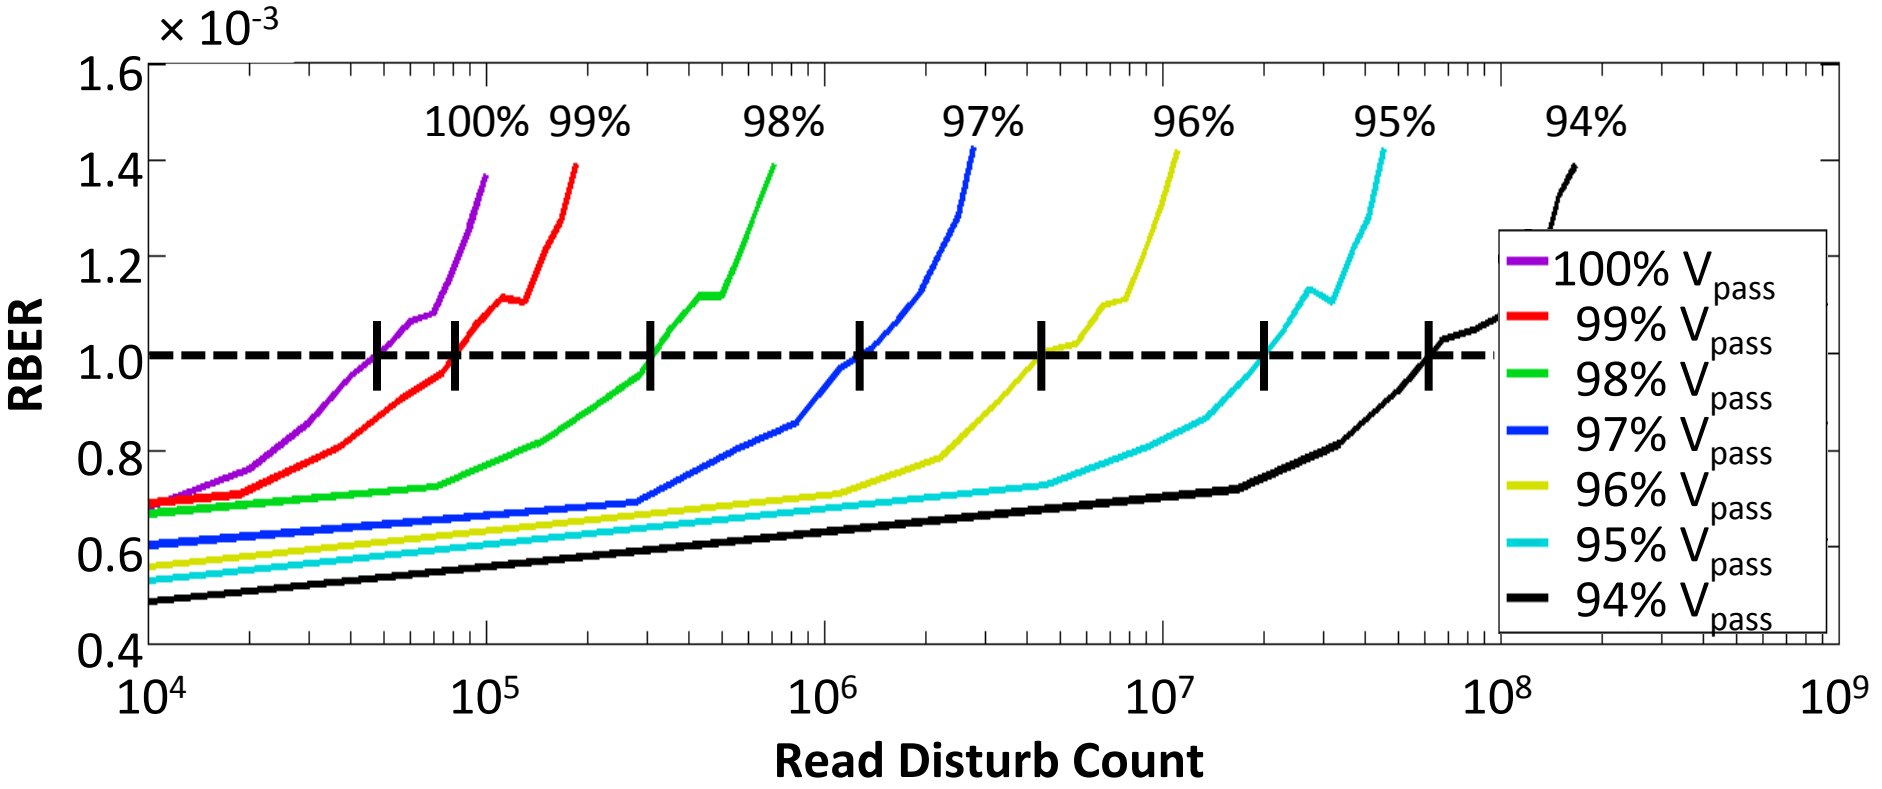

Pct. V_{pass} Value	100%	99%	98%	97%	96%	95%	94%
Rd. Disturb. Cnt.	1x	1.7x	6.8x	22x	100x	470x	1300x

Pass-Through Voltage Reduction Induced Read Error

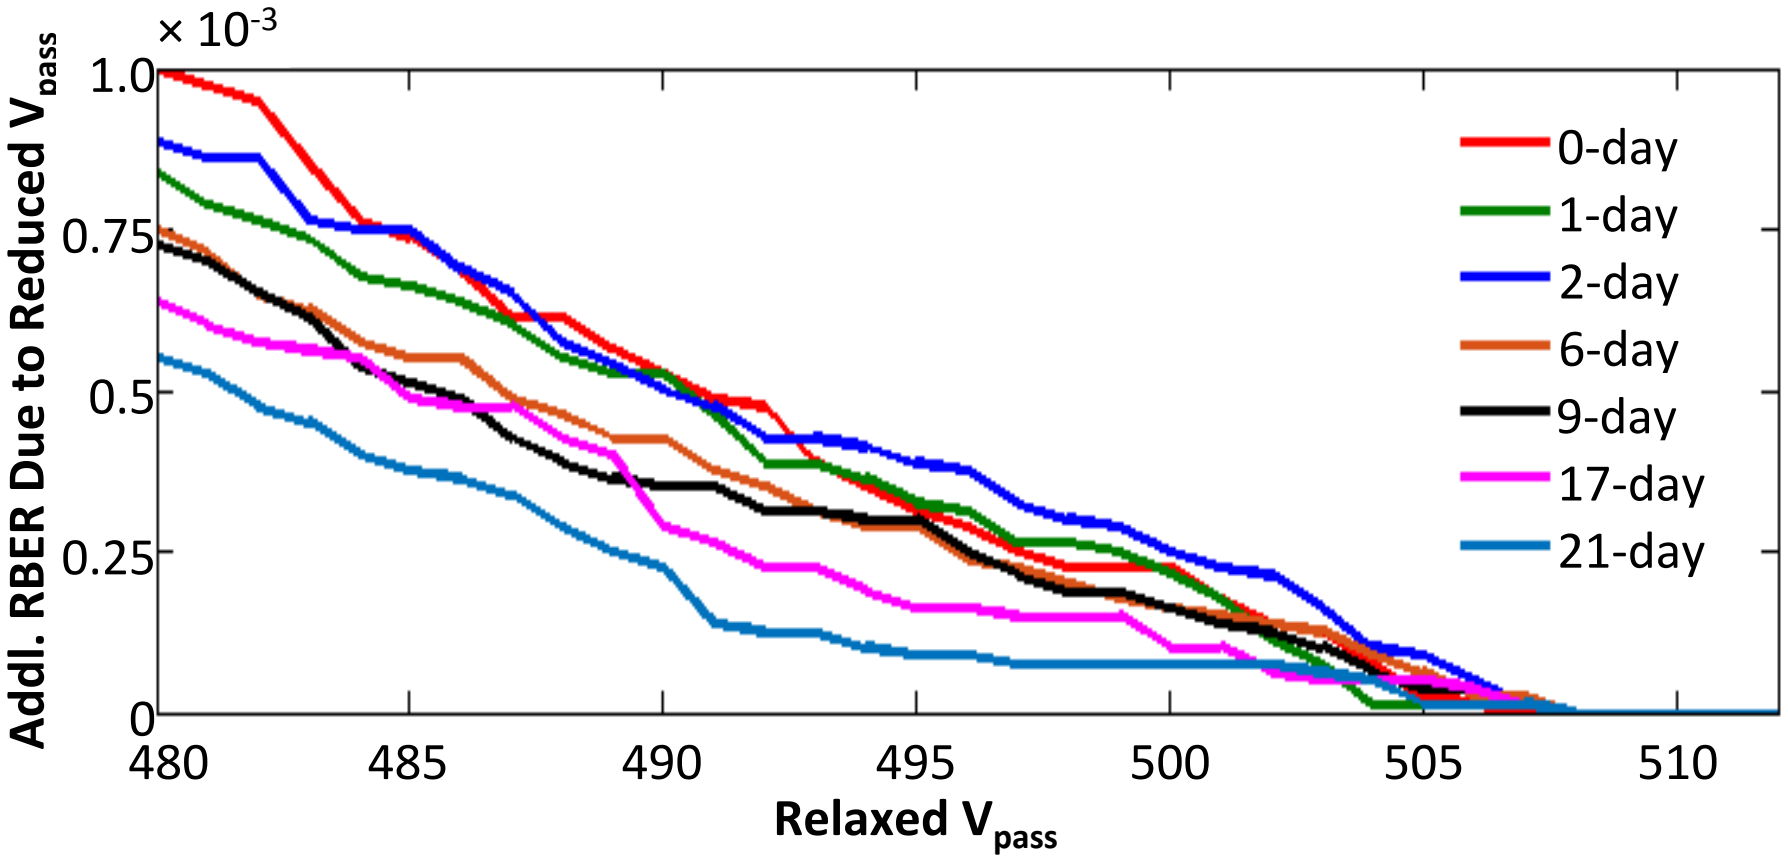

Read Errors Induced by V_{pass} Reduction

- Will generate a read error only if:
  - $Max(V_{th}) > V_{pass}$
  - Correct read value is 1

- These errors do not affect lifetime
  - can usually be tolerated by the unused ECC capability
- These errors are temporary
  - can be corrected (if necessary) by reading with the default V_{pass}

Illustration of V_{pass} Tuning Results

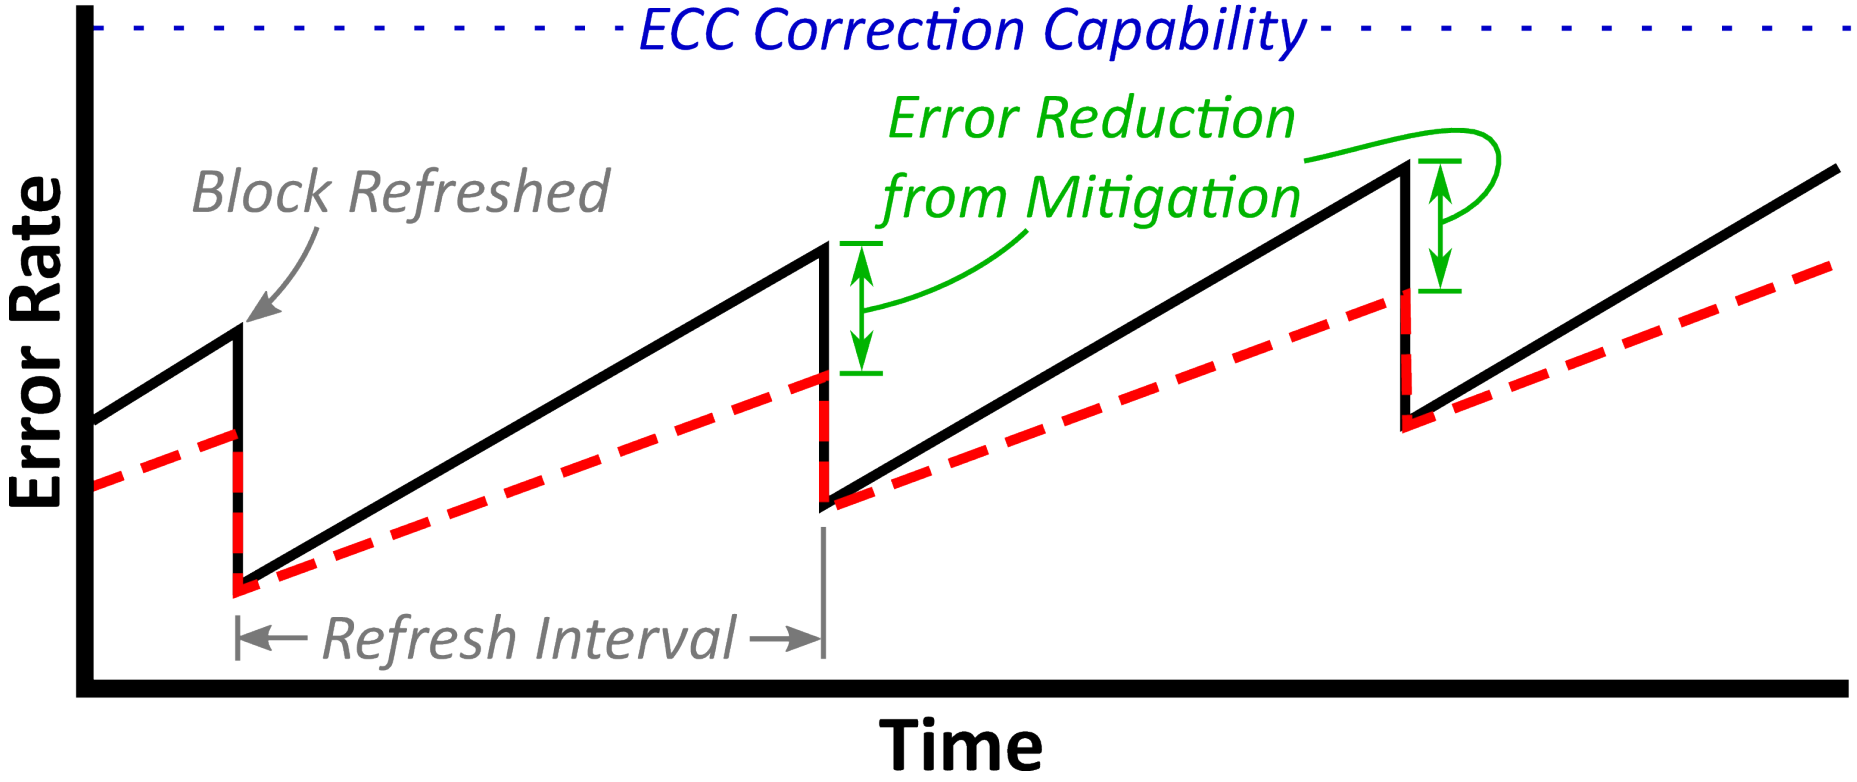

Some Flash Cells Are More Prone to Read Disturb

Flash Memory SUMMIT

Predict to be P1 state
- Area I is correct
- Area II is 50/50

Predict to be ER state
- Area III is correct
- Area IV is 50/50

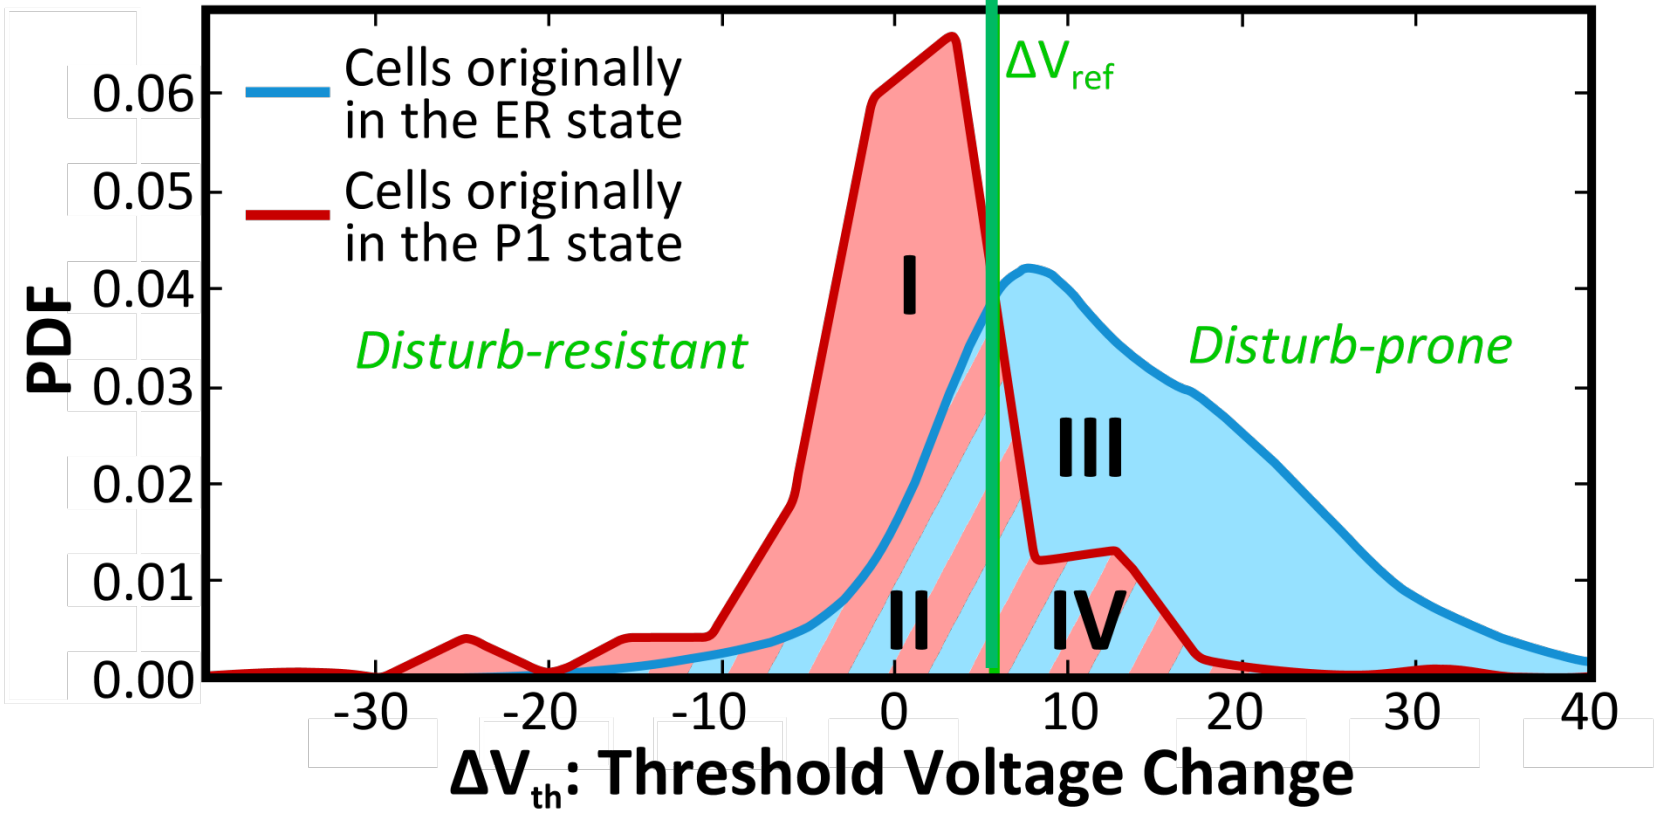

Showing ΔV_{th} with 8K PEC from 250K to 350K read disturbs

Read Disturb in DRAM

"Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors" (Kim et al., ISCA 2014)

"The DRAM RowHammer Problem (and Its Reliability and Security Implications)" (Mutlu, 2015)

Modern DRAM is Prone to Disturbance Errors

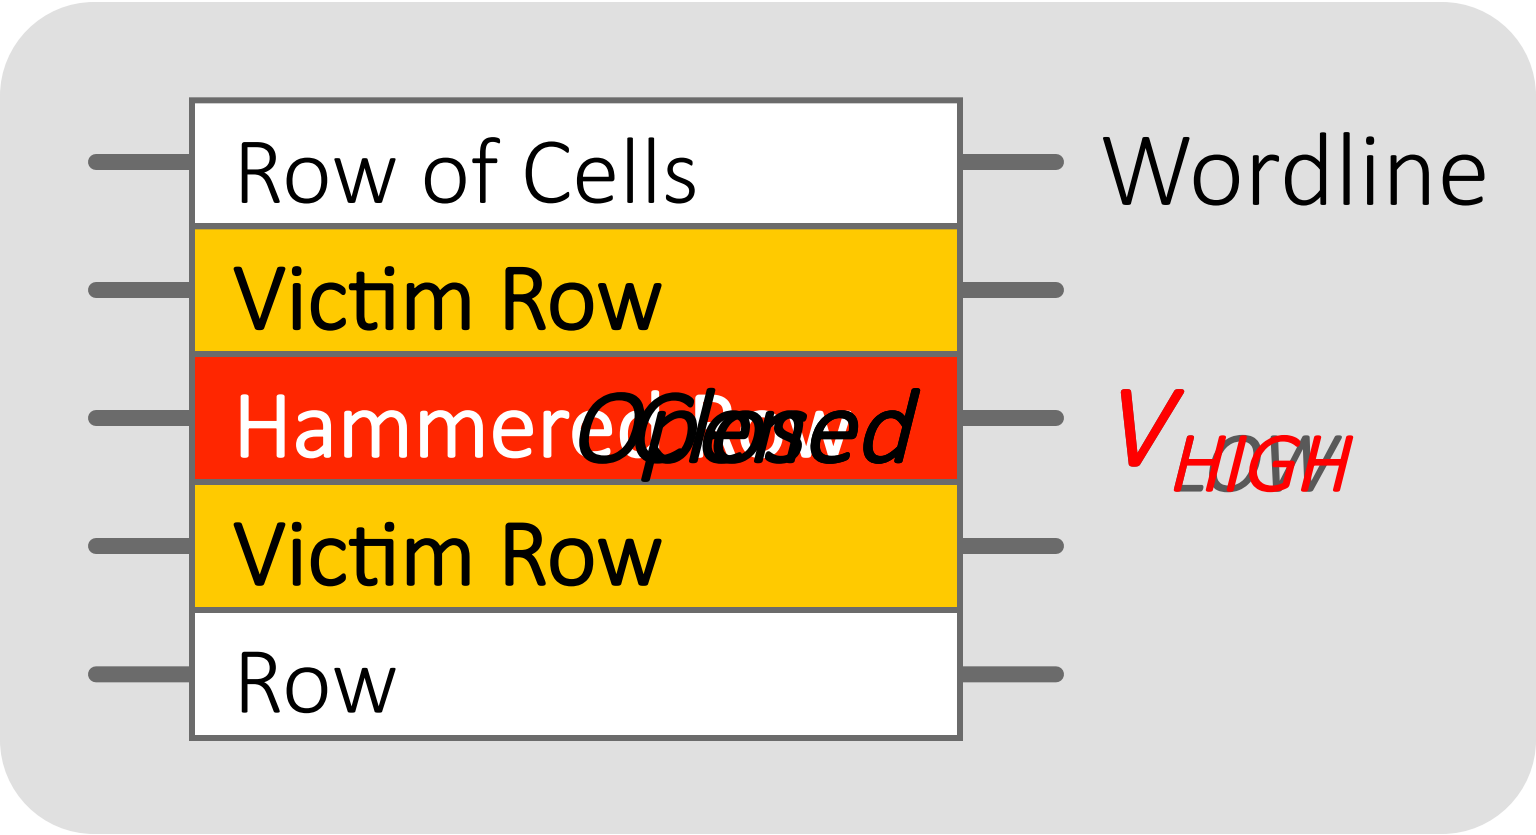

Repeatedly opening and closing a row enough times within a refresh interval induces disturbance errors in adjacent rows in most real DRAM chips you can buy today

Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors, (Kim et al., ISCA 2014)

Most DRAM Modules Are at Risk

A company

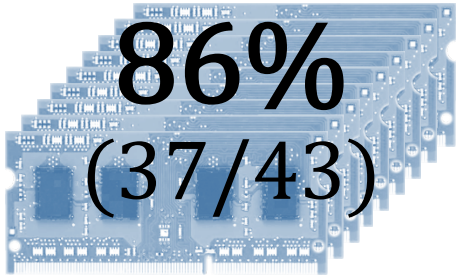

Up to
1.0×10^7
errors

B company

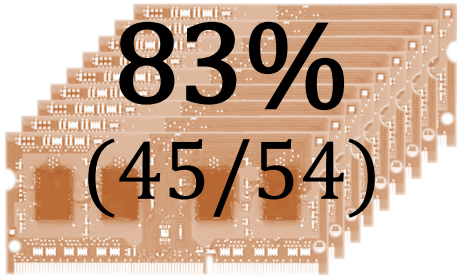

Up to
2.7×10^6
errors

C company

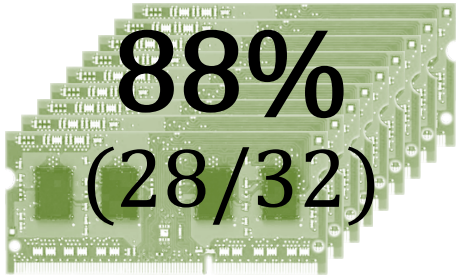

Up to
3.3×10^5
errors

Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors, (Kim et al., ISCA 2014)

A Simple Program Can Induce Many Errors

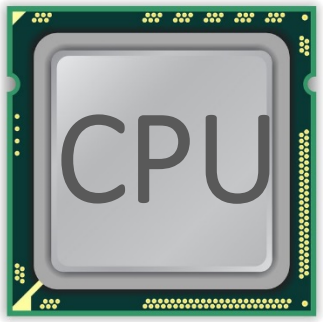

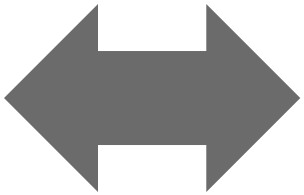

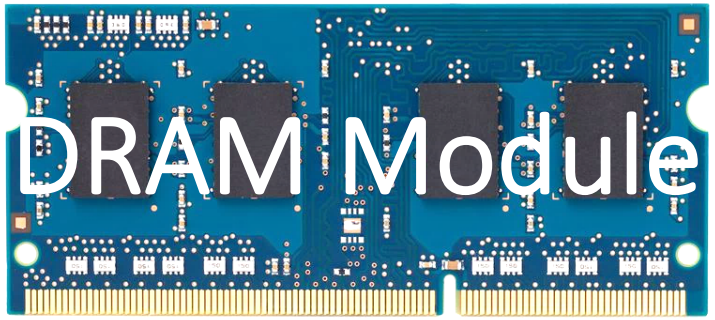


```
loop:
   mov (X), %eax
   mov (Y), %ebx
   clflush (X)
   clflush (Y)
   mfence
   jmp loop
```

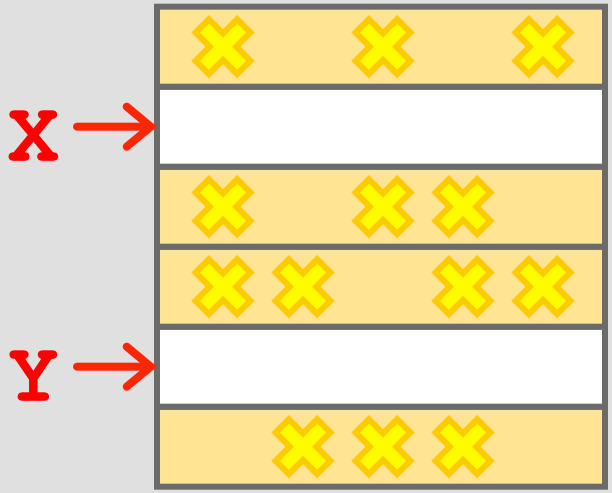

Download from: https://github.com/CMU-SAFARI/rowhammer

Observed Errors in Real Systems

CPU Architecture	Errors	Access-Rate
Intel Haswell (2013)	22.9K	12.3M/sec
Intel Ivy Bridge (2012)	20.7K	11.7M/sec
Intel Sandy Bridge (2011)	16.1K	11.6M/sec
AMD Piledriver (2012)	59	6.1M/sec

- *A real reliability & security issue*
- *In a more controlled environment, we can induce as many as ten million disturbance errors*

Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors, (Kim et al., ISCA 2014)

Experimental DRAM Testing Infrastructure

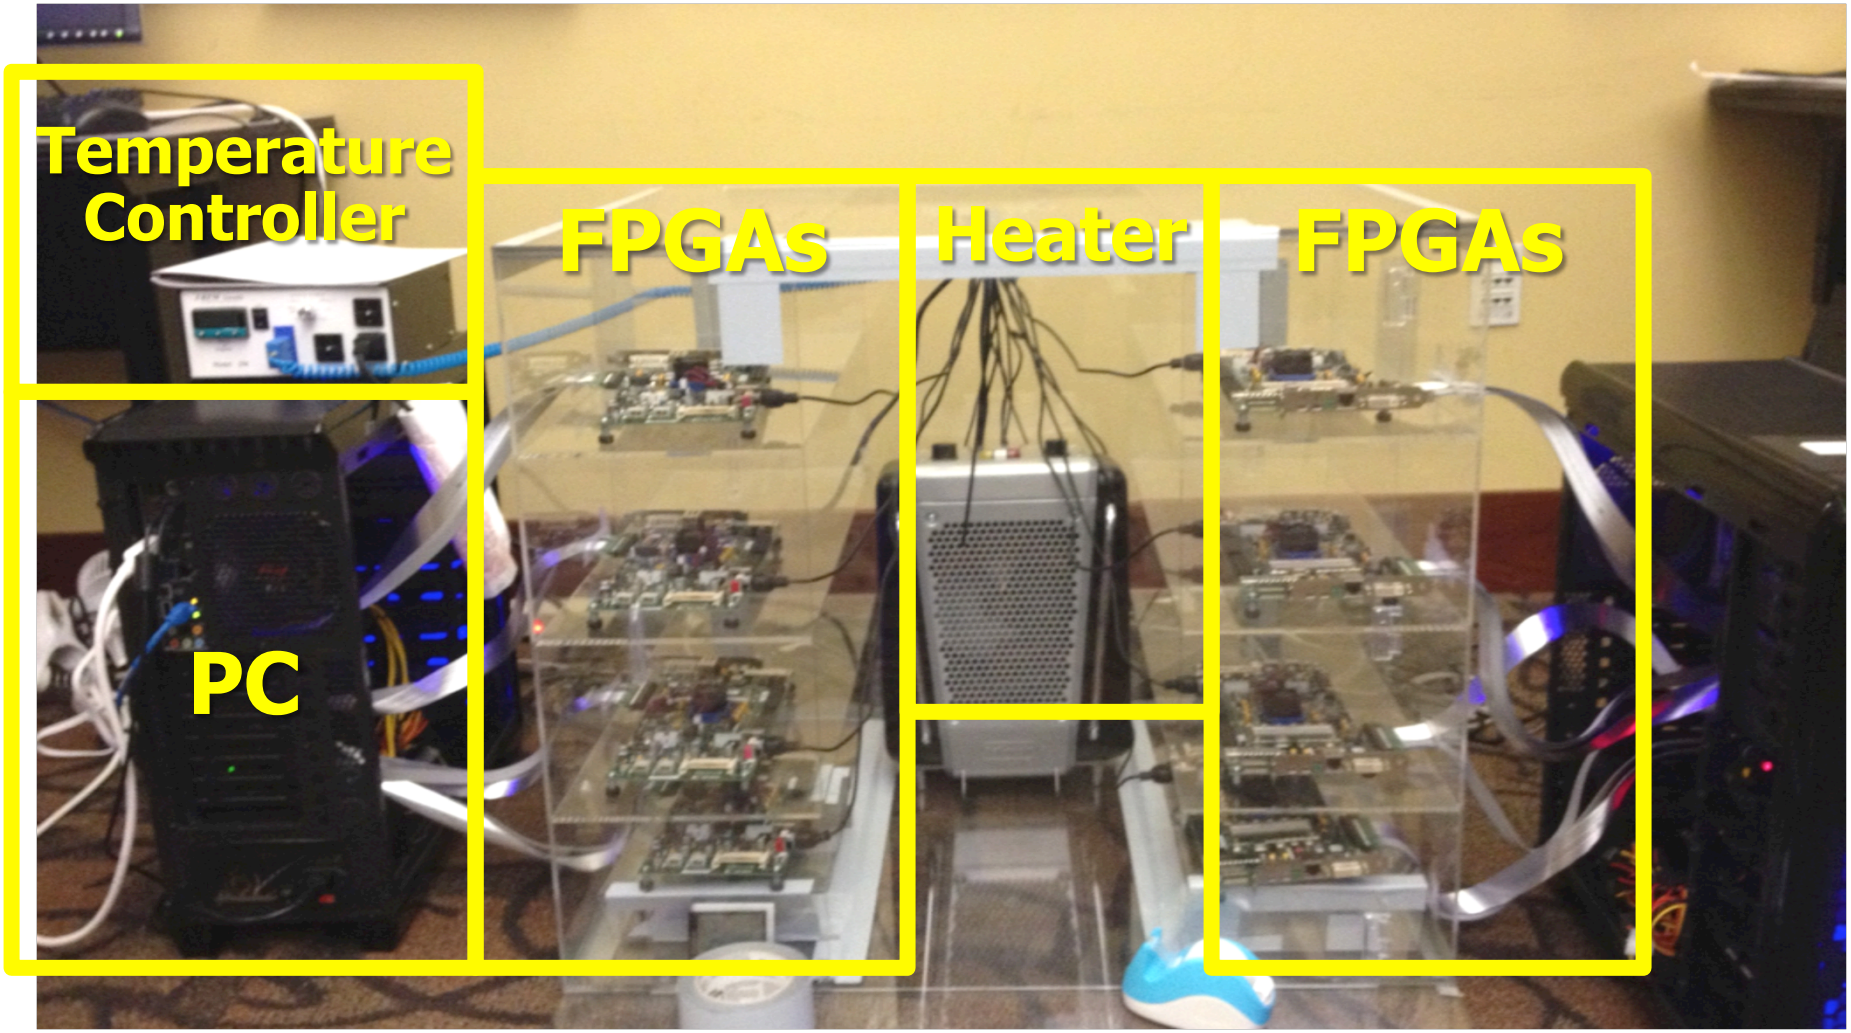

Kim+, "Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors," ISCA 2014.

One Can Take Over an Otherwise-Secure System

Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors

Abstract. Memory isolation is a key property of a reliable and secure computing system — an access to one memory address should not have unintended side effects on data stored in other addresses. However, as DRAM process technology

Project Zero

[Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors] (Kim et al., ISCA 2014)

News and updates from the Project Zero team at Google

[Exploiting the DRAM rowhammer bug to gain kernel privileges] (Seaborn, 2015)

Monday, March 9, 2015

Exploiting the DRAM rowhammer bug to gain kernel privileges

RowHammer Security Attack Example

- "Rowhammer" is a problem with some recent DRAM devices in which repeatedly accessing a row of memory can cause bit flips in adjacent rows (Kim et al., ISCA 2014).
  - Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors (Kim et al., ISCA 2014)
- We tested a selection of laptops and found that a subset of them exhibited the problem.
- We built two working privilege escalation exploits that use this effect.
  - Exploiting the DRAM rowhammer bug to gain kernel privileges (Seaborn, 2015)
- One exploit uses rowhammer-induced bit flips to gain kernel privileges on x86-64 Linux when run as an unprivileged userland process.
- When run on a machine vulnerable to the rowhammer problem, the process was able to induce bit flips in page table entries (PTEs).
- It was able to use this to gain write access to its own page table, and hence gain read-write access to all of physical memory.

Security Implications

It's like breaking into an apartment by repeatedly slamming a neighbor's door until the vibrations open the door you were after

Our Other Works on DRAM Errors

[An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms] (Liu et al., ISCA 2013)

[The Efficacy of Error Mitigation Techniques for DRAM Retention Failures: A Comparative Experimental Study] (Khan et al., SIGMETRICS 2014)

[Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors] (Kim et al., ISCA 2014)

[Adaptive-Latency DRAM: Optimizing DRAM Timing for the Common-Case] (Lee et al., HPCA 2015)

[AVATAR: A Variable-Retention-Time (VRT) Aware Refresh for DRAM Systems] (Qureshi et al., DSN 2015)