

# Computer Architecture

## Lecture 5c: Secure and Reliable Memory

Prof. Onur Mutlu

ETH Zürich

Fall 2020

1 October 2020

# Revisiting RowHammer in 2020 (I)

---

- Jeremie S. Kim, Minesh Patel, A. Giray Yaglikci, Hasan Hassan, Roknoddin Azizi, Lois Orosa, and Onur Mutlu,  
**"Revisiting RowHammer: An Experimental Analysis of Modern Devices and Mitigation Techniques"**  
*Proceedings of the 47th International Symposium on Computer Architecture (ISCA)*, Valencia, Spain, June 2020.  
[[Slides \(pptx\)](#)] [[pdf](#)]  
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]  
[[Talk Video](#) (20 minutes)]  
[[Lightning Talk Video](#) (3 minutes)]

## Revisiting RowHammer: An Experimental Analysis of Modern DRAM Devices and Mitigation Techniques

Jeremie S. Kim<sup>§†</sup>      Minesh Patel<sup>§</sup>      A. Giray Yağlıkçı<sup>§</sup>  
Hasan Hassan<sup>§</sup>      Roknoddin Azizi<sup>§</sup>      Lois Orosa<sup>§</sup>      Onur Mutlu<sup>§†</sup>  
<sup>§</sup>*ETH Zürich*      <sup>†</sup>*Carnegie Mellon University*



# Future Memory Reliability/Security Challenges

# Future of Main Memory

---

- DRAM is becoming less reliable → more vulnerable

# Large-Scale Failure Analysis of DRAM Chips


---

- Analysis and modeling of memory errors found in all of Facebook's server fleet
- Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu,  
**"Revisiting Memory Errors in Large-Scale Production Data Centers: Analysis and Modeling of New Trends from the Field"**  
*Proceedings of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, Rio de Janeiro, Brazil, June 2015.  
[[Slides \(pptx\)](#)] [[pdf](#)] [[DRAM Error Model](#)]

## Revisiting Memory Errors in Large-Scale Production Data Centers: Analysis and Modeling of New Trends from the Field

Justin Meza   Qiang Wu\*   Sanjeev Kumar\*   Onur Mutlu  
Carnegie Mellon University   \* Facebook, Inc.

# DRAM Reliability Reducing



*Intuition:  
quadratic  
increase in  
capacity*

# Aside: SSD Error Analysis in the Field

---

- First large-scale field study of flash memory errors
- Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu,  
**"A Large-Scale Study of Flash Memory Errors in the Field"**  
*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems*  
**(SIGMETRICS)**, Portland, OR, June 2015.  
[[Slides \(pptx\)](#)] [[pdf](#)] [[Coverage at ZDNet](#)]

## A Large-Scale Study of Flash Memory Failures in the Field

Justin Meza  
Carnegie Mellon University  
meza@cmu.edu

Qiang Wu  
Facebook, Inc.  
qw@fb.com

Sanjeev Kumar  
Facebook, Inc.  
skumar@fb.com

Onur Mutlu  
Carnegie Mellon University  
onur@cmu.edu

# Future of Main Memory

---

- DRAM is becoming less reliable → more vulnerable
- Due to difficulties in DRAM scaling, other problems may also appear (or they may be going unnoticed)
- Some errors may already be slipping into the field
  - Read disturb errors (RowHammer)
  - Retention errors
  - Read errors, write errors
  - ...
- These errors can also pose security vulnerabilities

# DRAM Data Retention Time Failures

---

- Determining the data retention time of a cell/row is getting more difficult
- Retention failures may already be slipping into the field

# Analysis of Data Retention Failures [ISCA'13]

---

- Jamie Liu, Ben Jaiyen, Yoongu Kim, Chris Wilkerson, and Onur Mutlu,  
**"An Experimental Study of Data Retention Behavior in Modern DRAM  
Devices: Implications for Retention Time Profiling Mechanisms"**  
*Proceedings of the 40th International Symposium on Computer Architecture  
(ISCA), Tel-Aviv, Israel, June 2013. [Slides \(ppt\)](#) [Slides \(pdf\)](#)*

## An Experimental Study of Data Retention Behavior in Modern DRAM Devices: Implications for Retention Time Profiling Mechanisms

Jamie Liu<sup>\*</sup>  
Carnegie Mellon University  
5000 Forbes Ave.  
Pittsburgh, PA 15213  
[jamiel@alumni.cmu.edu](mailto:jamiel@alumni.cmu.edu)

Ben Jaiyen<sup>\*</sup>  
Carnegie Mellon University  
5000 Forbes Ave.  
Pittsburgh, PA 15213  
[bjaiyen@alumni.cmu.edu](mailto:bjaiyen@alumni.cmu.edu)

Yoongu Kim  
Carnegie Mellon University  
5000 Forbes Ave.  
Pittsburgh, PA 15213  
[yoonguk@ece.cmu.edu](mailto:yoonguk@ece.cmu.edu)

Chris Wilkerson  
Intel Corporation  
2200 Mission College Blvd.  
Santa Clara, CA 95054  
[chris.wilkerson@intel.com](mailto:chris.wilkerson@intel.com)

Onur Mutlu  
Carnegie Mellon University  
5000 Forbes Ave.  
Pittsburgh, PA 15213  
[onur@cmu.edu](mailto:onur@cmu.edu)



# Two Challenges to Retention Time Profiling

---

- Data Pattern Dependence (DPD) of retention time
- Variable Retention Time (VRT) phenomenon

<https://www.youtube.com/watch?v=v702wUnaWGE>

# Industry Is Writing Papers About It, Too

## DRAM Process Scaling Challenges

### ❖ Refresh


- Difficult to build high-aspect ratio cell capacitors decreasing cell capacitance
- Leakage current of cell access transistors increasing

### ❖ tWR

- Contact resistance between the cell capacitor and access transistor increasing
- On-current of the cell access transistor decreasing
- Bit-line resistance increasing

### ❖ VRT

- Occurring more frequently with cell capacitance decreasing



# Industry Is Writing Papers About It, Too

## DRAM Process Scaling Challenges

### ❖ Refresh


- Difficult to build high-aspect ratio cell capacitors decreasing cell capacitance

THE MEMORY FORUM 2014

## Co-Architecting Controllers and DRAM to Enhance DRAM Process Scaling

Uksong Kang, Hak-soo Yu, Churoo Park, \*Hongzhong Zheng,  
\*\*John Halbert, \*\*Kuljit Bains, SeongJin Jang, and Joo Sun Choi

*Samsung Electronics, Hwasung, Korea / \*Samsung Electronics, San Jose / \*\*Intel*



# Keeping Future Memory Secure

# How Do We Keep Memory Secure?

---

- DRAM
- Flash memory
- Emerging Technologies
  - Phase Change Memory
  - STT-MRAM
  - RRAM, memristors
  - ...

# Many Errors and Their Mitigation [PIEEE'17]

**Table 3** List of Different Types of Errors Mitigated by NAND Flash Error Mitigation Mechanisms

Mitigation Mechanism	Error Type				
	<i>P/E Cycling</i> [32,33,42] (§IV-A)	<i>Program</i> [40,42,53] (§IV-B)	<i>Cell-to-Cell Interference</i> [32,35,36,55] (§IV-C)	<i>Data Retention</i> [20,32,34,37,39] (§IV-D)	<i>Read Disturb</i> [20,32,38,62] (§IV-E)
<b>Shadow Program Sequencing</b> [35,40] (Section V-A)			X		
<b>Neighbor-Cell Assisted Error Correction</b> [36] (Section V-B)			X		
<b>Refresh</b> [34,39,67,68] (Section V-C)				X	X
<b>Read-Retry</b> [33,72,107] (Section V-D)	X			X	X
<b>Voltage Optimization</b> [37,38,74] (Section V-E)	X			X	X
<b>Hot Data Management</b> [41,63,70] (Section V-F)	X	X	X	X	X
<b>Adaptive Error Mitigation</b> [43,65,77,78,82] (Section V-G)	X	X	X	X	X

Cai+, "Error Characterization, Mitigation, and Recovery in Flash Memory Based Solid State Drives," Proc. IEEE 2017.

# Solution Direction: Principled Designs

---

Design fundamentally secure  
computing architectures

Predict and prevent  
such safety issues


# Architecting Future Memory for Security

---


- **Understand:** Methods for vulnerability modeling & discovery
  - ❑ Modeling and prediction based on real (device) data and analysis
  - ❑ Understanding vulnerabilities
  - ❑ Developing reliable metrics
- **Architect:** Principled architectures with security as key concern
  - ❑ Good partitioning of duties across the stack
  - ❑ Cannot give up performance and efficiency
  - ❑ Patch-ability in the field
- **Design & Test:** Principled design, automation, (online) testing
  - ❑ Design for security
  - ❑ High coverage and good interaction with system reliability methods



# Understand and Model with Experiments (DRAM)



# Understand and Model with Experiments (Flash)



[DATE 2012, ICCD 2012, DATE 2013, ITJ 2013, ICCD 2013, SIGMETRICS 2014, HPCA 2015, DSN 2015, MSST 2015, JSAC 2016, HPCA 2017, DFRWS 2017, PIEEE 2017, HPCA 2018, SIGMETRICS 2018]

NAND Daughter Board

Cai+, "Error Characterization, Mitigation, and Recovery in Flash Memory Based Solid State Drives," Proc. IEEE 2017.





*Proceedings of the IEEE, Sept. 2017*

## Error Characterization, Mitigation, and Recovery in Flash-Memory-Based Solid-State Drives

*This paper reviews the most recent advances in solid-state drive (SSD) error characterization, mitigation, and data recovery techniques to improve both SSD's reliability and lifetime.*

By YU CAI, SAUGATA GHOSE, ERICH F. HARATSCH, YIXIN LUO, AND ONUR MUTLU

# Understanding Flash Memory Reliability

---

- Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu,  
**"A Large-Scale Study of Flash Memory Errors in the Field"**  
*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Portland, OR, June 2015.*  
[[Slides \(pptx\)](#)] [[pdf](#)] [[Coverage at ZDNet](#)] [[Coverage on The Register](#)]  
[[Coverage on TechSpot](#)] [[Coverage on The Tech Report](#)]

## A Large-Scale Study of Flash Memory Failures in the Field

Justin Meza  
Carnegie Mellon University  
[meza@cmu.edu](mailto:meza@cmu.edu)

Qiang Wu  
Facebook, Inc.  
[qw@fb.com](mailto:qw@fb.com)

Sanjeev Kumar  
Facebook, Inc.  
[skumar@fb.com](mailto:skumar@fb.com)

Onur Mutlu  
Carnegie Mellon University  
[onur@cmu.edu](mailto:onur@cmu.edu)

# NAND Flash Vulnerabilities [HPCA'17]

*HPCA, Feb. 2017*

## Vulnerabilities in MLC NAND Flash Memory Programming: Experimental Analysis, Exploits, and Mitigation Techniques

Yu Cai<sup>†</sup>   Saugata Ghose<sup>†</sup>   Yixin Luo<sup>††</sup>   Ken Mai<sup>†</sup>   Onur Mutlu<sup>§†</sup>   Erich F. Haratsch<sup>‡</sup>  
<sup>†</sup>Carnegie Mellon University   <sup>‡</sup>Seagate Technology   <sup>§</sup>ETH Zürich

*Modern NAND flash memory chips provide high density by storing two bits of data in each flash cell, called a multi-level cell (MLC). An MLC partitions the threshold voltage range of a flash cell into four voltage states. When a flash cell is programmed, a high voltage is applied to the cell. Due to parasitic capacitance coupling between flash cells that are physically close to each other, flash cell programming can lead to cell-to-cell program interference, which introduces errors into neighboring flash cells. In order to reduce the impact of cell-to-cell interference on the reliability of MLC NAND flash memory, flash manufacturers adopt a two-step programming method, which programs the MLC in two separate steps. First, the flash memory partially programs the least significant bit of the MLC to some intermediate threshold voltage. Second, it programs the most significant bit to bring the MLC up to its full voltage state.*

*In this paper, we demonstrate that two-step programming exposes new reliability and security vulnerabilities. We expe-*

*belongs to a different flash memory page (the unit of data programmed and read at the same time), which we refer to, respectively, as the least significant bit (LSB) page and the most significant bit (MSB) page [5].*

*A flash cell is programmed by applying a large voltage on the control gate of the transistor, which triggers charge transfer into the floating gate, thereby increasing the threshold voltage. To precisely control the threshold voltage of the cell, the flash memory uses incremental step pulse programming (ISPP) [12, 21, 25, 41]. ISPP applies multiple short pulses of the programming voltage to the control gate, in order to increase the cell threshold voltage by some small voltage amount ( $V_{step}$ ) after each step. Initial MLC designs programmed the threshold voltage in one shot, issuing all of the pulses back-to-back to program both bits of data at the same time. However, as flash memory scales down, the distance between neighboring flash cells decreases, which*

[https://people.inf.ethz.ch/omutlu/pub/flash-memory-programming-vulnerabilities\\_hpca17.pdf](https://people.inf.ethz.ch/omutlu/pub/flash-memory-programming-vulnerabilities_hpca17.pdf)

# 3D NAND Flash Reliability I [HPCA'18]

---

- 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"**

*Proceedings of the 24th International Symposium on High-Performance Computer Architecture (HPCA), Vienna, Austria, February 2018.*

[\[Lightning Talk Video\]](#)

[\[Slides \(pptx\) \(pdf\)\]](#) [\[Lightning Session Slides \(pptx\) \(pdf\)\]](#)

## HeatWatch: Improving 3D NAND Flash Memory Device Reliability by Exploiting Self-Recovery and Temperature Awareness

Yixin Luo<sup>†</sup>      Saugata Ghose<sup>†</sup>      Yu Cai<sup>‡</sup>      Erich F. Haratsch<sup>‡</sup>      Onur Mutlu<sup>§†</sup>  
<sup>†</sup>*Carnegie Mellon University*      <sup>‡</sup>*Seagate Technology*      <sup>§</sup>*ETH Zürich*

# 3D NAND Flash Reliability II [SIGMETRICS'18]

---

- 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"**  
*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Irvine, CA, USA, June 2018.*  
[[Abstract](#)]  
[[POMACS Journal Version \(same content, different format\)](#)]  
[[Slides \(pptx\)](#) ([pdf](#))]

## Improving 3D NAND Flash Memory Lifetime by Tolerating Early Retention Loss and Process Variation

Yixin Luo<sup>†</sup>

Saugata Ghose<sup>†</sup>

Yu Cai<sup>†</sup>

Erich F. Haratsch<sup>‡</sup>

Onur Mutlu<sup>§†</sup>

<sup>†</sup>Carnegie Mellon University

<sup>‡</sup>Seagate Technology

<sup>§</sup>ETH Zürich



# Recall: Collapse of the “Galloping Gertie”

---





# Another Example (1994)





# Yet Another Example (2007)

---



Source: Morry Gash/AP,  
<https://www.npr.org/2017/08/01/540669701/10-years-after-bridge-collapse-america-is-still-crumbling?t=1535427165809>



# A More Recent Example (2018)

---



In-Field Patch-ability  
(Intelligent Memory)  
Can Avoid Such Failures

# Final Thoughts on RowHammer

# Aside: Byzantine Failures

---

- This class of failures is known as **Byzantine failures**
- Characterized by
  - **Undetected erroneous computation**
  - Opposite of “fail fast (with an error or no result)”
- “erroneous” can be “malicious” (intent is the only distinction)
- Very difficult to detect and confine Byzantine failures
- **Do all you can to avoid them**
- Lamport et al., “The Byzantine Generals Problem,” ACM TOPLAS 1982.

# Aside: Byzantine Generals Problem

---

## The Byzantine Generals Problem

LESLIE LAMPORT, ROBERT SHOSTAK, and MARSHALL PEASE  
SRI International

---

Reliable computer systems must handle malfunctioning components that give conflicting information to different parts of the system. This situation can be expressed abstractly in terms of a group of generals of the Byzantine army camped with their troops around an enemy city. Communicating only by messenger, the generals must agree upon a common battle plan. However, one or more of them may be traitors who will try to confuse the others. The problem is to find an algorithm to ensure that the loyal generals will reach agreement. It is shown that, using only oral messages, this problem is solvable if and only if more than two-thirds of the generals are loyal; so a single traitor can confound two loyal generals. With unforgeable written messages, the problem is solvable for any number of generals and possible traitors. Applications of the solutions to reliable computer systems are then discussed.

Categories and Subject Descriptors: C.2.4. [**Computer-Communication Networks**]: Distributed Systems—*network operating systems*; D.4.4 [**Operating Systems**]: Communications Management—*network communication*; D.4.5 [**Operating Systems**]: Reliability—*fault tolerance*

General Terms: Algorithms, Reliability

Additional Key Words and Phrases: Interactive consistency



# RowHammer, Revisited

- One can **predictably induce bit flips** in commodity DRAM chips
  - >80% of the tested DRAM chips are vulnerable
- First example of how a **simple hardware failure mechanism** can create a **widespread system security vulnerability**

**WIRED**

Forget Software—Now Hackers Are Exploiting Physics

BUSINESS	CULTURE	DESIGN	GEAR	SCIENCE
----------	---------	--------	------	---------

ANDY GREENBERG SECURITY 08.31.16 7:00 AM

SHARE



SHARE  
18276



TWEET

# FORGET SOFTWARE—NOW HACKERS ARE EXPLOITING PHYSICS

# RowHammer: Retrospective

---

- New mindset that has enabled a renewed interest in HW security attack research:
  - ❑ Real (memory) chips are vulnerable, in a simple and widespread manner  
→ this causes real security problems
  - ❑ Hardware reliability → security connection is now mainstream discourse
- Many new RowHammer attacks...
  - ❑ Tens of papers in top security venues
  - ❑ **More to come** as RowHammer is getting worse (DDR4 & beyond)
- Many new RowHammer solutions...
  - ❑ Apple security release; Memtest86 updated
  - ❑ Many solution proposals in top venues (latest in ISCA 2019)
  - ❑ Principled system-DRAM co-design (in original RowHammer paper)
  - ❑ **More to come...**

# Perhaps Most Importantly...

---

- RowHammer enabled a shift of mindset in mainstream security researchers
  - General-purpose hardware is fallible, in a widespread manner
  - Its problems are exploitable
- This mindset has enabled many systems security researchers to examine hardware in more depth
  - And understand HW's inner workings and vulnerabilities
- It is no coincidence that two of the groups that discovered Meltdown and Spectre heavily worked on RowHammer attacks before
  - **More to come...**

# Summary: RowHammer

---

- DRAM reliability is reducing
- Reliability issues open up security vulnerabilities
  - Very hard to defend against
- **Rowhammer is a prime example**
  - First example of how a simple hardware failure mechanism can create a widespread system security vulnerability
  - Its implications on system security research are tremendous & exciting
- Bad news: RowHammer is getting worse.
- **Good news: We have a lot more to do.**
  - We are now fully aware hardware is easily fallible.
  - We are developing both attacks and solutions.
  - We are developing principled models, methodologies, solutions.

# For More on RowHammer...

---

- Onur Mutlu and Jeremie Kim,  
**"RowHammer: A Retrospective"**  
*IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems (TCAD) Special Issue on Top Picks in Hardware and Embedded Security*, 2019.  
[[Preliminary arXiv version](#)]

## RowHammer: A Retrospective

Onur Mutlu<sup>§‡</sup>      Jeremie S. Kim<sup>‡§</sup>  
§ETH Zürich      ‡Carnegie Mellon University

# RowHammer in 2020 (I)

---

- Jeremie S. Kim, Minesh Patel, A. Giray Yaglikci, Hasan Hassan, Roknoddin Azizi, Lois Orosa, and Onur Mutlu,  
**"Revisiting RowHammer: An Experimental Analysis of Modern Devices and Mitigation Techniques"**  
*Proceedings of the 47th International Symposium on Computer Architecture (ISCA)*, Valencia, Spain, June 2020.  
[[Slides \(pptx\)](#)] [[pdf](#)]  
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]  
[[Talk Video](#) (20 minutes)]  
[[Lightning Talk Video](#) (3 minutes)]

## Revisiting RowHammer: An Experimental Analysis of Modern DRAM Devices and Mitigation Techniques

Jeremie S. Kim<sup>§†</sup>      Minesh Patel<sup>§</sup>      A. Giray Yağlıkçı<sup>§</sup>  
Hasan Hassan<sup>§</sup>      Roknoddin Azizi<sup>§</sup>      Lois Orosa<sup>§</sup>      Onur Mutlu<sup>§†</sup>  
<sup>§</sup>*ETH Zürich*      <sup>†</sup>*Carnegie Mellon University*

# RowHammer in 2020 (II)

---

- Pietro Frigo, Emanuele Vannacci, Hasan Hassan, Victor van der Veen, Onur Mutlu, Cristiano Giuffrida, Herbert Bos, and Kaveh Razavi, **"TRRespass: Exploiting the Many Sides of Target Row Refresh"** *Proceedings of the 41st IEEE Symposium on Security and Privacy (S&P)*, San Francisco, CA, USA, May 2020.  
[[Slides \(pptx\)](#)] [[pdf](#)]  
[[Talk Video](#) (17 minutes)]  
[[Source Code](#)]  
[[Web Article](#)]  
***Best paper award.***

## TRRespass: Exploiting the Many Sides of Target Row Refresh

Pietro Frigo<sup>\*†</sup>   Emanuele Vannacci<sup>\*†</sup>   Hasan Hassan<sup>§</sup>   Victor van der Veen<sup>¶</sup>  
Onur Mutlu<sup>§</sup>   Cristiano Giuffrida<sup>\*</sup>   Herbert Bos<sup>\*</sup>   Kaveh Razavi<sup>\*</sup>



# RowHammer in 2020 (III)

---

- Lucian Cojocar, Jeremie Kim, Minesh Patel, Lillian Tsai, Stefan Saroiu, Alec Wolman, and Onur Mutlu,  
**"Are We Susceptible to Rowhammer? An End-to-End Methodology for Cloud Providers"**  
*Proceedings of the 41st IEEE Symposium on Security and Privacy (S&P)*, San Francisco, CA, USA, May 2020.  
[[Slides \(pptx\)](#)] [[pdf](#)]  
[[Talk Video](#) (17 minutes)]

## Are We Susceptible to Rowhammer? An End-to-End Methodology for Cloud Providers

Lucian Cojocar, Jeremie Kim<sup>§†</sup>, Minesh Patel<sup>§</sup>, Lillian Tsai<sup>‡</sup>,  
Stefan Saroiu, Alec Wolman, and Onur Mutlu<sup>§†</sup>  
Microsoft Research, <sup>§</sup>ETH Zürich, <sup>†</sup>CMU, <sup>‡</sup>MIT



Rowhammer

# Some History

# Some More Historical Perspective

---

- RowHammer is the first example of a circuit-level failure mechanism causing a widespread system security vulnerability
- It led to a large body of work in security attacks, mitigations, architectural solutions, ...
- Work building on RowHammer still continues
  - See MICRO 2020, ISCA 2020,
- Initially, it was dismissed by some reviewers
  - Rejected from MICRO 2013 conference

# Initial RowHammer Reviews (MICRO 2013)

## #66 Disturbance Errors in DRAM: Demonstration, Characterization, and Prevention

ON  
'e  
or

**Rejected (R2)**



863kB

Friday 31 May 2013 2:00:53pm PDT

b9bf06021da54cddf4cd0b3565558a181868b972

You are an **author** of this paper.

### + ABSTRACT

We demonstrate the vulnerability of commodity DRAM chips to disturbance errors. By repeatedly reading from one DRAM address, we show that it is possible to corrupt the data stored [\[more\]](#)

### + AUTHORS

Y. Kim, R. Daly, J. Lee, J. Kim, C. Fallin, C. Wilkerson, O. Mutlu  
[\[details\]](#)

**KEYWORDS:** DRAM; errors


### + TOPICS

[Review #66A](#)  
[Review #66B](#)  
[Review #66C](#)  
[Review #66D](#)  
[Review #66E](#)  
[Review #66F](#)

OveMer	Nov	WriQua	RevExp
--------	-----	--------	--------

1	4	4	4
5	4	5	3
2	3	5	4
1	2	3	4
4	4	4	3
2	4	4	3

# Reviewer A

**Review #66A** Modified Friday 5 Jul 2013 3:59:18am PDT  [Plain text](#)

## OVERALL MERIT (?)

**1.** Reject

## PAPER SUMMARY

This work tests and studies the disturbance problem in DRAM arrays in isolation.

## PAPER STRENGTHS

- + Many results and observations.
- + Insights on how the may happen

## PAPER WEAKNESSES

- Whereas they show disturbance may happen in DRAM array, authors don't show it can be an issue in realistic DRAM usage scenario
- Lacks architectural/microarchitectural impact on the DRAM disturbance analysis

## NOVELTY (?)

**4.** New contribution.

## WRITING QUALITY (?)

**4.** Well-written



# Reviewer A -- Security is Not “Realistic”

---

## COMMENTS FOR AUTHORS

I found the paper very well written and organized, easy to understand. The topic is interesting and relevant.

However, I'm not fully convinced that the disturbance problem is going to be an issue in a realistic DRAM usage scenario (main memory with caches). In that scenario the 64ms refresh interval might be enough. Overall, the work presented, the experimentation and the results are not enough to justify/claim that disturbance may be an issue for future systems, and that microarchitectural solutions are required.

I really encourage the authors to address this issue, to run the new set of experiments; if the results are positive, the work is great and will be easily accepted in a top notch conference. Test scenario in the paper (open-read-close a row many times consecutively) that is used to create disturbances is not likely to show up in a realistic usage scenario (check also rebuttal question).



# Rebuttal to Reviewer A

---

\_\_\_\_\_ WILL IT AFFECT REAL WORKLOADS ON REAL SYSTEMS?  
(A, E) \_\_\_\_\_

Malicious workloads and pathological access-patterns can bypass/thrash the cache and access the same DRAM row a very large number of times. While these workloads may not be common, they are just as real. Using non-temporal

# Reviewer A -- Demands

---

To make sure that correct information and messages are given to the research community, it would be good if the conclusions drawn in the paper were verified with the actual DRAM manufacturers, although I see that it can be difficult to do. In addition, knowing the technology node of each tested DRAM would make the paper stronger and would avoid speculative guesses.

## REVIEWER EXPERTISE (?)

4. Expert in area, with highest confidence in review.

# Reviewer C

## **Review #66C**

Modified Friday 12 Jul 2013 7:38:57am

 [Plain text](#)

PDT

### **OVERALL MERIT (?)**

**2.** Weak reject

### **PAPER SUMMARY**

This paper presents a rigorous study of DRAM module errors which are observed to be caused through repeated access to the same address in the DRAMs.

### **PAPER STRENGTHS**

The paper's measurement methodology is outstanding, and the authors very thoroughly dive into different test scenarios, to isolate the circumstances under which the observed errors take place.

### **PAPER WEAKNESSES**

This is an excellent test methodology paper, but there is no micro-architectural or architectural content.

### **NOVELTY (?)**

**3.** Incremental improvement.

### **WRITING QUALITY (?)**

**5.** Outstanding

### **QUESTIONS TO ADDRESS IN THE REBUTTAL**

My primary concern with this paper is that it doesn't have (micro-)architectural content, and may not spur on future work.

# Reviewer C -- Leave It to DRAM Vendors

---

## COMMENTS FOR AUTHORS

This is an extremely well-written analysis of DRAM behavior, and the authors are to be commended on establishing a robust and flexible characterization platform and methodology.

That being said, disturb errors have occurred repeatedly over the course of DRAM's history (which the authors do acknowledge). History has shown that particular disturbances, and in particular hammer errors, are short-lived, and are quickly solved by DRAM manufacturers. Historically, once these types of errors occur at a particular lithography node/DRAM density, they must be solved by the DRAM manufacturers, because even if a solution for a systemic problem could be asserted for particular markets (e.g., server, where use of advanced coding techniques, extra chips, etc. is acceptable), there will always be significant DRAM chip volume in single-piece applications (e.g., consumer devices, etc.) where complex architectural solutions aren't an option. The authors have identified a contemporary disturb sensitivity in DRAMs, but as non-technologists, our community can generally only observe, not correct, such problems.

## REVIEWER EXPERTISE (?)

4. Expert in area, with highest confidence in review.

# Reviewer D -- Nothing New in RowHammer

## **Review #66D**

Modified Thursday 18 Jul 2013 12:51pm

 [Plain text](#)

PDT

### OVERALL MERIT (?)

**1.** Reject

### REVIEWER EXPERTISE (?)

**4.** Expert in area, with highest confidence in review.

### PAPER SUMMARY

The authors demonstrate that repeated activate-precharge operations on one wordline of a DRAM can disturb a few cells on adjacent wordlines. They showed that such a behavior can be caused for most DRAMs and all DRAMs of recent manufacture they tested.

### PAPER STRENGTHS

DRAM errors are getting more likely with newer generations and it is necessary to investigate their cause and mitigation in computer systems, as such the paper addresses a subtopic of a relevant problem.

### PAPER WEAKNESSES

The mechanism investigated by the authors is one of many well known disturb mechanisms. The paper does not discuss the root causes to sufficient depth and the importance of this mechanism compared to others. Overall the length of the sections restating known information is much too long in relation to new work.

### NOVELTY (?)

**2.** Insignificant novelty.  
Virtually all of the ideas are published or known.

### WRITING QUALITY (?)

**3.** Adequate

# ISCA 2014 Submission

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

N

Accepted



639kB

21 Nov 2013 10:53:11pm CST |

f039be2735313b39304ae1c6296523867a485610

You are an **author** of this paper.

### + 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 [\[more\]](#)

### + AUTHORS

Y. Kim, R. Daly, J. Kim, J. Lee, C. Fallin, C. Wilkerson, O. Mutlu  
[\[details\]](#)

### + TOPICS

[Review #41A](#)  
[Review #41B](#)  
[Review #41C](#)  
[Review #41D](#)  
[Review #41E](#)  
[Review #41F](#)

OveMer	Nov	WriQua	RevConAnd
8	4	5	3
7	4	4	3
6	4	4	3
2	2	5	4
3	2	3	3
7	4	4	3

# Reviewer D

---

## **Review #41D**

Modified 19 Feb 2014 8:47:24pm



[Plain text](#)

CST

### **OVERALL MERIT (?)**

**2.** Reject

---

### **PAPER SUMMARY**

The authors

- 1) characterize disturbance error in commodity DRAM
- 2) identify the root cause such errors (but it's already a well know problem in DRAM community).
- 3) propose a simple architectural technique to mitigate such errors.

---

### **PAPER STRENGTHS**

The authors demonstrated the problem using the real systems

---

### **PAPER WEAKNESSES**

- 1) The disturbance error (a.k.a coupling or cross-talk noise induced error) is a known problem to the DRAM circuit community.

- 2) What you demonstrated in this paper is so called DRAM row hammering issue - you can even find a Youtube video showing this! - <http://www.youtube.com/watch?v=i3-qOSnBcdo>

- 2) The architectural contribution of this study is too insignificant.



## NOVELTY (?)

**2.** Insignificant novelty.  
Virtually all of the ideas  
are published or known.

---

## WRITING QUALITY (?)

**5.** Outstanding

---

## REVIEWER CONFIDENCE AND EXPERTISE (?)

**4.** Expert in area, with highest confidence in review.

---

## QUESTIONS FOR AUTHORS

1. There are other sources of disturbance errors How can you guarantee the errors observed by you are not from such errors?

2. You did you best on explaining why we have much fewer 1->0 error but not quite satisfied. Any other explanation?

3. Can you elaborate why we have more disturbed cells over rounds while you claim that disturbed cells are not weak cells? I'm sure this is related to device again issues

---

## DETAILED COMMENTS

This is a well written and executed paper (in particular using real systems), but I have many concerns:

1) this is a well-known problem to the DRAM community (so no novelty there); in DRAM community people use

# Reviewer D Continued...

---

2) what you did to incur disturbance is is so called "row hammering" issues - please see <http://www.youtube.com/watch?v=i3-qQSnBcdo> - a demonstration video for capturing this problem...

3) the relevance of this paper to ISCA. I feel that this paper (most part) is more appropriate to conferences like International Test Conference (ITC) or VLSI Test Symposium or Dependable Systems and Networks (DSN) at most. This is because the authors mainly dedicated the effort to the DRAM circuit characterization and test method in my view while the architectural contribution is very weak - I'm not even sure this can be published to these venues since it's a well known problem! I also assume techniques proposed to minimize disturbance error in STT-RAM and other technology can be employed here as well.

# Rebuttal to Reviewer D


\_\_\_\_Reviewer D (Comments)\_\_\_\_

---

- 1. As we acknowledge in the paper, it is true that different types of DRAM coupling phenomena have been known to the DRAM circuits/testing community. However, there is a clear distinction between circuits/testing techniques confined to the \*foundry\* versus characterization/solution of a problem out in the \*field\*. The three citations (from 10+ years ago) do \*not\* demonstrate that disturbance errors exist in DIMMs sold then or now. They do \*not\* provide any real data (only simulated ones), let alone a large-scale characterization across many DIMMs from multiple manufacturers. They do \*not\* construct an attack on real systems, and they do \*not\* provide any solutions. Finally, our paper \*already\* references all three citations, or their more relevant equivalents. (The second/third citations provided by the reviewer are on bitline-coupling, whereas we cite works from the same authors on wordline-coupling [2, 3, 37].)

- 2. We were aware of the video from Teledyne (a test equipment company) and have \*already\* referenced slides from the same company [36]. In terms of their content regarding "row hammer", the video and the slides are identical: all they mention is that "aggressive row activations can corrupt adjacent rows". (They then advertise how their test equipment is able to capture a timestamped DRAM access trace, which can then be post-processed to identify when the number of activations exceeds a user-set threshold.) Both the video and slides do \*not\* say that this is a real problem affecting DIMMs on the market now. They do \*not\* provide any quantitative data, \*nor\* real-system demonstration, \*nor\* solution.

# Reviewer E

**Review #41E** Modified 7 Feb 2014 11:08:04pm CST  [Plain text](#)

OVERALL MERIT (?)

**3.** Weak Reject

## PAPER SUMMARY

This paper studies the row disturbance problem in DRAMs. The paper includes a thorough quantitative characterization of the problem and a qualitative discussion of the source of the problem and potential solutions.

## PAPER STRENGTHS

+ The paper provides a detailed quantitative characterization of the “row hammering” problem in memories.

## PAPER WEAKNESSES

- Row Hammering appears to be well-known, and solutions have already been proposed by industry to address the issue.
- The paper only provides a qualitative analysis of solutions to the problem. A more robust evaluation is really needed to know whether the proposed solution is necessary.

NOVELTY (?)

**2.** Insignificant novelty.  
Virtually all of the ideas are published or known.

WRITING QUALITY (?)

**3.** Adequate

REVIEWER CONFIDENCE AND EXPERTISE (?)

**3.** Knowledgeable in area, and significant confidence in

but there are numerous mentions of hammering in the literature, and clearly industry has studied this problem for many years. In particular, Intel has a patent application on a memory controller technique that addresses this exact problem, with priority date June 2012:

<http://www.google.com/patents/WO2014004748A1?cl=en>

The patent application details sound very similar to solution 6 in this paper, so a more thorough comparison with solution 7 seems mandatory.

My overall feeling is that while the reliability characterization is important and interesting, a better target audience for the characterization work would be in a testing/reliability venue. The most interesting part of this paper from the ISCA point of view are the proposed solutions, but all of these are discussed in a very qualitative manner. My preference would be to see a much shorter characterization section with a much stronger and quantitative evaluation and comparison of the proposed solutions.

# Rebuttal to Reviewer

---

\*Nevertheless\*, we were able to induce a large number of DRAM disturbance errors on all the latest Intel/AMD platforms that we tested: Haswell, Ivy Bridge, Sandy Bridge, and Piledriver. (At the time of submission, we had tested only Sandy Bridge.) Importantly, the patents do \*not\* provide quantitative characterization  
\*nor\* real-system demonstration.

[R1] "Row Hammer Refresh Command." US20140006703 A1  
[R2] "Row Hammer Condition Monitoring." US20140006704 A1

\_\_\_\_Reviewer E (Comments)\_\_\_\_

After our paper was submitted, two patents that had been filed by

Intel were made public (one is mentioned by the reviewer [R1]).

Together, the two patents describe what we posed as the \*sixth\*

potential solution in our paper (Section 8). Essentially, the memory controller maintains a table of counters to track the number of activations to recently activated rows [R2].

And if one of the counters exceeds a certain threshold, the memory controller notifies the DRAM chips using a special command [R1].

The DRAM chips would then refresh an entire "region" of rows that

includes both the aggressor and its victim(s) [R1]. For the

patent [R1] to work, DRAM manufacturers must cooperate and

implement this special command. (It is a convenient way of

circumventing the opacity in the logical-physical mapping. If

implemented, the same command can also be used for our \*seventh\*

solution.) The limitation of this \*sixth\* solution is the storage

overhead of the counters and the extra power required to associatively search through them on every activation

(Section

8). That is why we believe our \*seventh\* solution to be more

attractive. We will cite the patents and include a more concrete

comparison between the two solutions.



# Suggestions to Reviewers

---

- Be fair; you do not know it all
- Be open-minded; you do not know it all
- Be accepting of diverse research methods: there is no single way of doing research
- Be constructive, not destructive
- Do not have double standards...

**Do not block or delay scientific progress for non-reasons**

# An Interview on Research and Education

---

- Computing Research and Education (@ ISCA 2019)
  - [https://www.youtube.com/watch?v=8ffSEKZhmvo&list=PL5Q2soXY2Zi\\_4oP9LdL3cc8G6NIjD2Ydz](https://www.youtube.com/watch?v=8ffSEKZhmvo&list=PL5Q2soXY2Zi_4oP9LdL3cc8G6NIjD2Ydz)
- Maurice Wilkes Award Speech (10 minutes)
  - [https://www.youtube.com/watch?v=tcQ3zZ3JpuA&list=PL5Q2soXY2Zi8D\\_5MGV6EnXEJHnV2YFBJI&index=15](https://www.youtube.com/watch?v=tcQ3zZ3JpuA&list=PL5Q2soXY2Zi8D_5MGV6EnXEJHnV2YFBJI&index=15)


# More Thoughts and Suggestions

---

- Onur Mutlu,  
**"Some Reflections (on DRAM)"**  
*Award Speech for ACM SIGARCH Maurice Wilkes Award, at the **ISCA** Awards Ceremony, Phoenix, AZ, USA, 25 June 2019.*  
[[Slides \(pptx\)](#)] [[pdf](#)]  
[[Video of Award Acceptance Speech \(Youtube; 10 minutes\)](#)] [[Youku; 13 minutes](#)]  
[[Video of Interview after Award Acceptance \(Youtube; 1 hour 6 minutes\)](#)] [[Youku; 1 hour 6 minutes](#)]  
[[News Article on "ACM SIGARCH Maurice Wilkes Award goes to Prof. Onur Mutlu"](#)]
  
- Onur Mutlu,  
**"How to Build an Impactful Research Group"**  
*57th Design Automation Conference Early Career Workshop (**DAC**), Virtual, 19 July 2020.*  
[[Slides \(pptx\)](#)] [[pdf](#)]

# Aside: A Recommended Book

---



Raj Jain, "[The Art of Computer Systems Performance Analysis](#)," Wiley, 1991.

## 10.8 DECISION MAKER'S GAMES

Even if the performance analysis is correctly done and presented, it may not be enough to persuade your audience—the decision makers—to follow your recommendations. The list shown in Box 10.2 is a compilation of reasons for rejection heard at various performance analysis presentations. You can use the list by presenting it immediately and pointing out that the reason for rejection is not new and that the analysis deserves more consideration. Also, the list is helpful in getting the competing proposals rejected!

There is no clear end of an analysis. Any analysis can be rejected simply on the grounds that the problem needs more analysis. This is the first reason listed in Box 10.2. The second most common reason for rejection of an analysis and for endless debate is the workload. Since workloads are always based on the past measurements, their applicability to the current or future environment can always be questioned. Actually workload is one of the four areas of discussion that lead a performance presentation into an endless debate. These “rat holes” and their relative sizes in terms of time consumed are shown in Figure 10.26. Presenting this cartoon at the beginning of a presentation helps to avoid these areas.




FIGURE 10.26 Four issues in performance presentations that commonly lead to endless discussion.

Raj Jain, "The Art of Computer Systems Performance Analysis," Wiley, 1991.



**Box 10.2 Reasons for Not Accepting the Results of an Analysis**

1. This needs more analysis.
2. You need a better understanding of the workload.
3. It improves performance only for long I/O's, packets, jobs, and files, and most of the I/O's, packets, jobs, and files are short.
4. It improves performance only for short I/O's, packets, jobs, and files, but who cares for the performance of short I/O's, packets, jobs, and files; its the long ones that impact the system.
5. It needs too much memory/CPU/bandwidth and memory/CPU/bandwidth isn't free.
6. It only saves us memory/CPU/bandwidth and memory/CPU/bandwidth is cheap.
7. There is no point in making the networks (similarly, CPUs/disks/...) faster; our CPUs/disks (any component other than the one being discussed) aren't fast enough to use them.
8. It improves the performance by a factor of  $x$ , but it doesn't really matter at the user level because everything else is so slow.
9. It is going to increase the complexity and cost.
10. Let us keep it simple stupid (and your idea is not stupid).
11. It is not simple. (Simplicity is in the eyes of the beholder.)
12. It requires too much state.
13. Nobody has ever done that before. (You have a new idea.)
14. It is not going to raise the price of our stock by even an eighth. (Nothing ever does, except rumors.)
15. This will violate the IEEE, ANSI, CCITT, or ISO standard.
16. It may violate some future standard.
17. The standard says nothing about this and so it must not be important.
18. Our competitors don't do it. If it was a good idea, they would have done it.
19. Our competition does it this way and you don't make money by copying others.
20. It will introduce randomness into the system and make debugging difficult.
21. It is too deterministic; it may lead the system into a cycle.
22. It's not interoperable.
23. This impacts hardware.
24. That's beyond today's technology.
25. It is not self stabilizing.
26. Why change—it's working OK.


Raj Jain, "The Art of Computer Systems Performance Analysis," Wiley, 1991.



# A Fun Reading: Food for Thought

---

- <https://www.livemint.com/science/news/could-einstein-get-published-today-11601014633853.html>



A similar process of professionalization has transformed other parts of the scientific landscape. (Central Press/Getty Images)

THE WALL STREET JOURNAL.

## Could Einstein get published today?

3 min read . Updated: 25 Sep 2020, 11:51 AM IST

The Wall Street Journal

Scientific journals and institutions have become more professionalized over the last century, leaving less room for individual style

# Computer Architecture

## Lecture 5c: Secure and Reliable Memory

Prof. Onur Mutlu

ETH Zürich

Fall 2020

1 October 2020

# Backup Slides

# Read Disturb in Flash Memory

# Many Errors and Their Mitigation [PIEEE'17]

**Table 3** List of Different Types of Errors Mitigated by NAND Flash Error Mitigation Mechanisms

Mitigation Mechanism	Error Type				
	<i>P/E Cycling</i> [32,33,42] (§IV-A)	<i>Program</i> [40,42,53] (§IV-B)	<i>Cell-to-Cell Interference</i> [32,35,36,55] (§IV-C)	<i>Data Retention</i> [20,32,34,37,39] (§IV-D)	<i>Read Disturb</i> [20,32,38,62] (§IV-E)
<b>Shadow Program Sequencing</b> [35,40] (Section V-A)			X		
<b>Neighbor-Cell Assisted Error Correction</b> [36] (Section V-B)			X		
<b>Refresh</b> [34,39,67,68] (Section V-C)				X	X
<b>Read-Retry</b> [33,72,107] (Section V-D)	X			X	X
<b>Voltage Optimization</b> [37,38,74] (Section V-E)	X			X	X
<b>Hot Data Management</b> [41,63,70] (Section V-F)	X	X	X	X	X
<b>Adaptive Error Mitigation</b> [43,65,77,78,82] (Section V-G)	X	X	X	X	X

Cai+, "Error Characterization, Mitigation, and Recovery in Flash Memory Based Solid State Drives," Proc. IEEE 2017.



*Proceedings of the IEEE, Sept. 2017*



## Error Characterization, Mitigation, and Recovery in Flash-Memory-Based Solid-State Drives

*This paper reviews the most recent advances in solid-state drive (SSD) error characterization, mitigation, and data recovery techniques to improve both SSD's reliability and lifetime.*

By YU CAI, SAUGATA GHOSE, ERICH F. HARATSCH, YIXIN LUO, AND ONUR MUTLU

<https://arxiv.org/pdf/1706.08642>

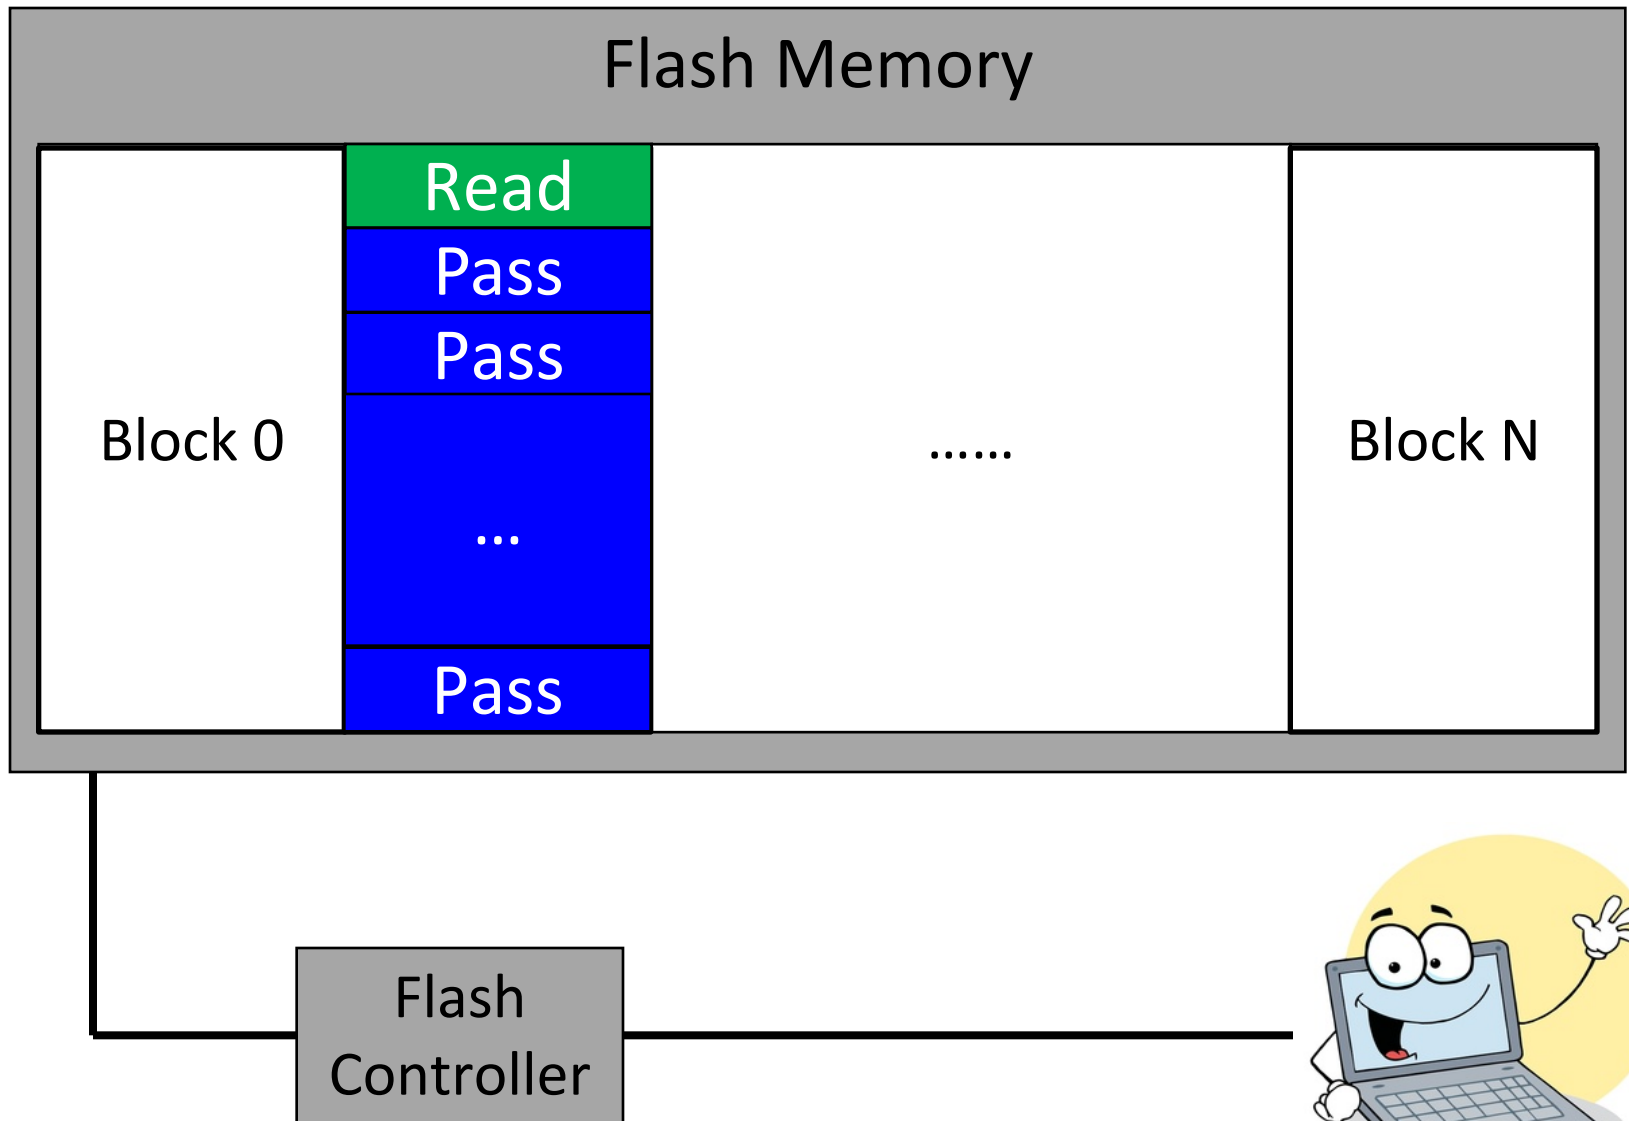


# One Issue: Read Disturb in Flash Memory


---

- All scaled memories are prone to read disturb errors
- DRAM
- SRAM
- Hard Disks: Adjacent Track Interference
- NAND Flash


# NAND Flash Memory Background



# Flash Cell Array




# Flash Cell




Floating Gate Transistor  
(Flash Cell)

# Flash Read

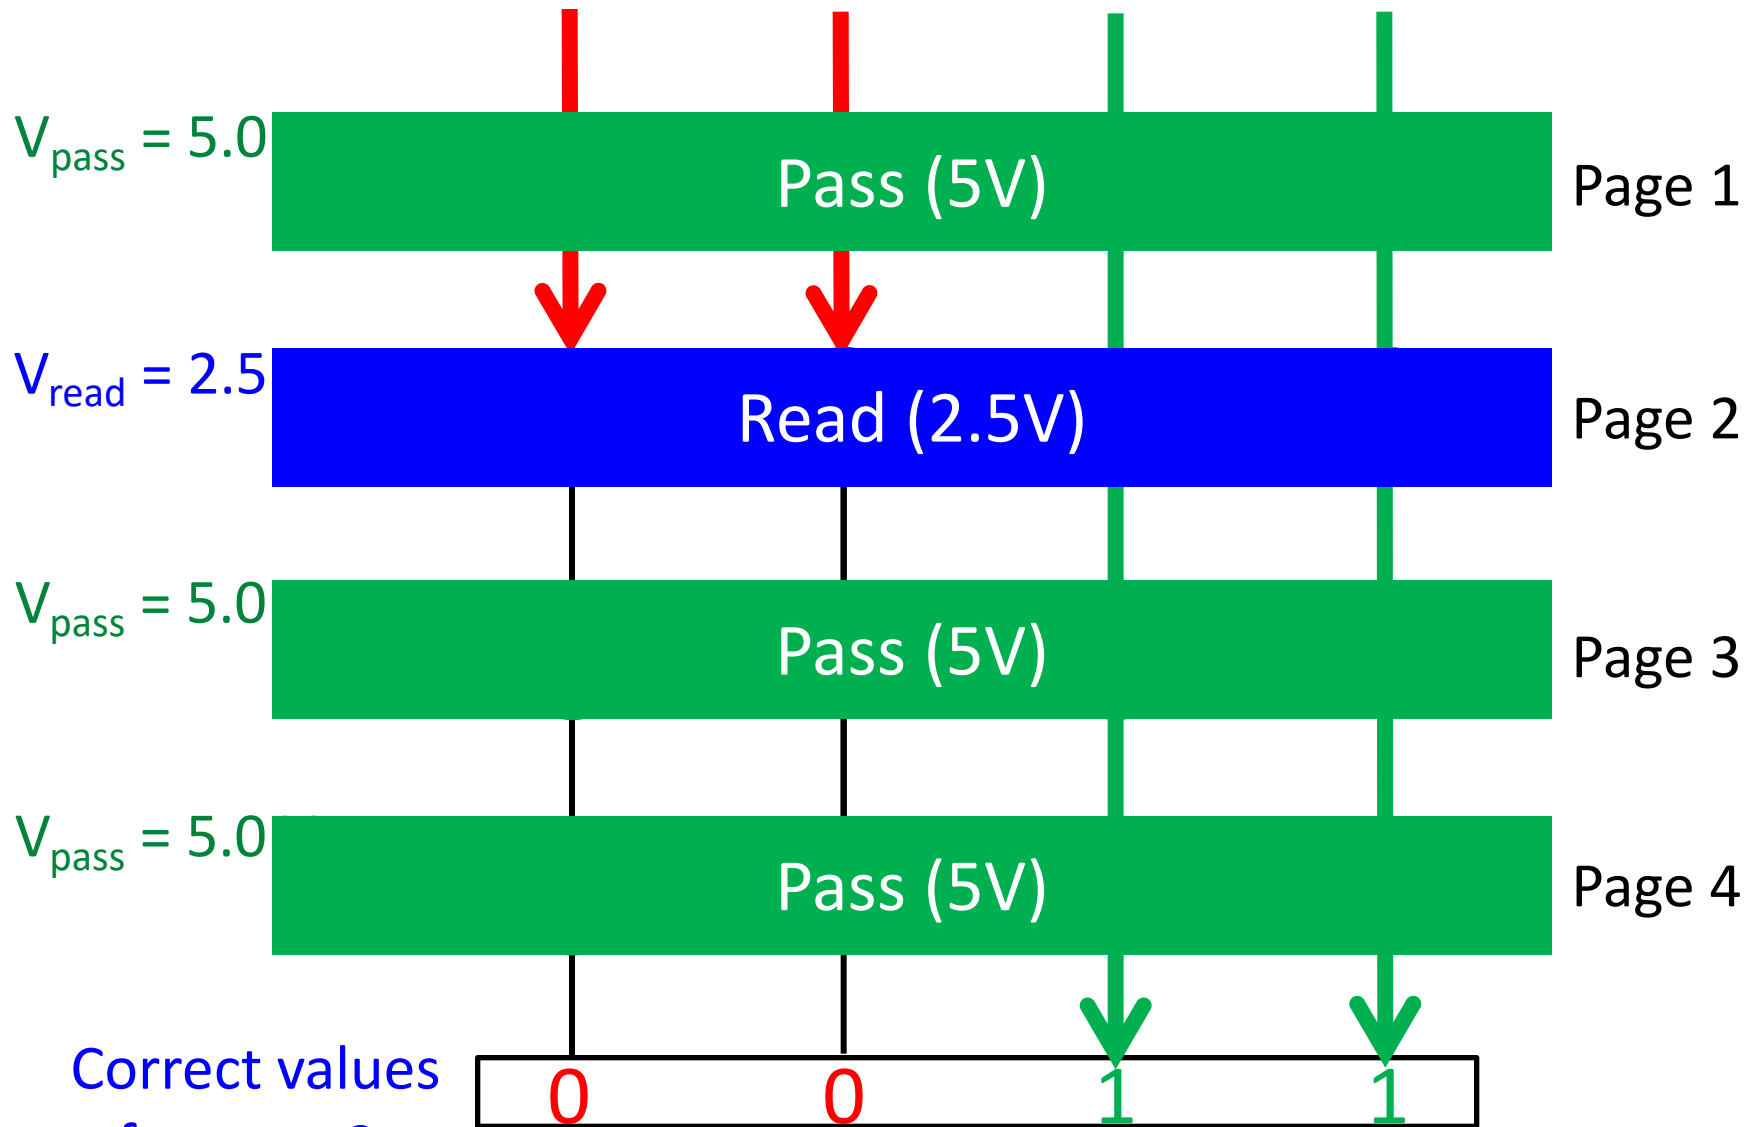


# Flash Pass-Through

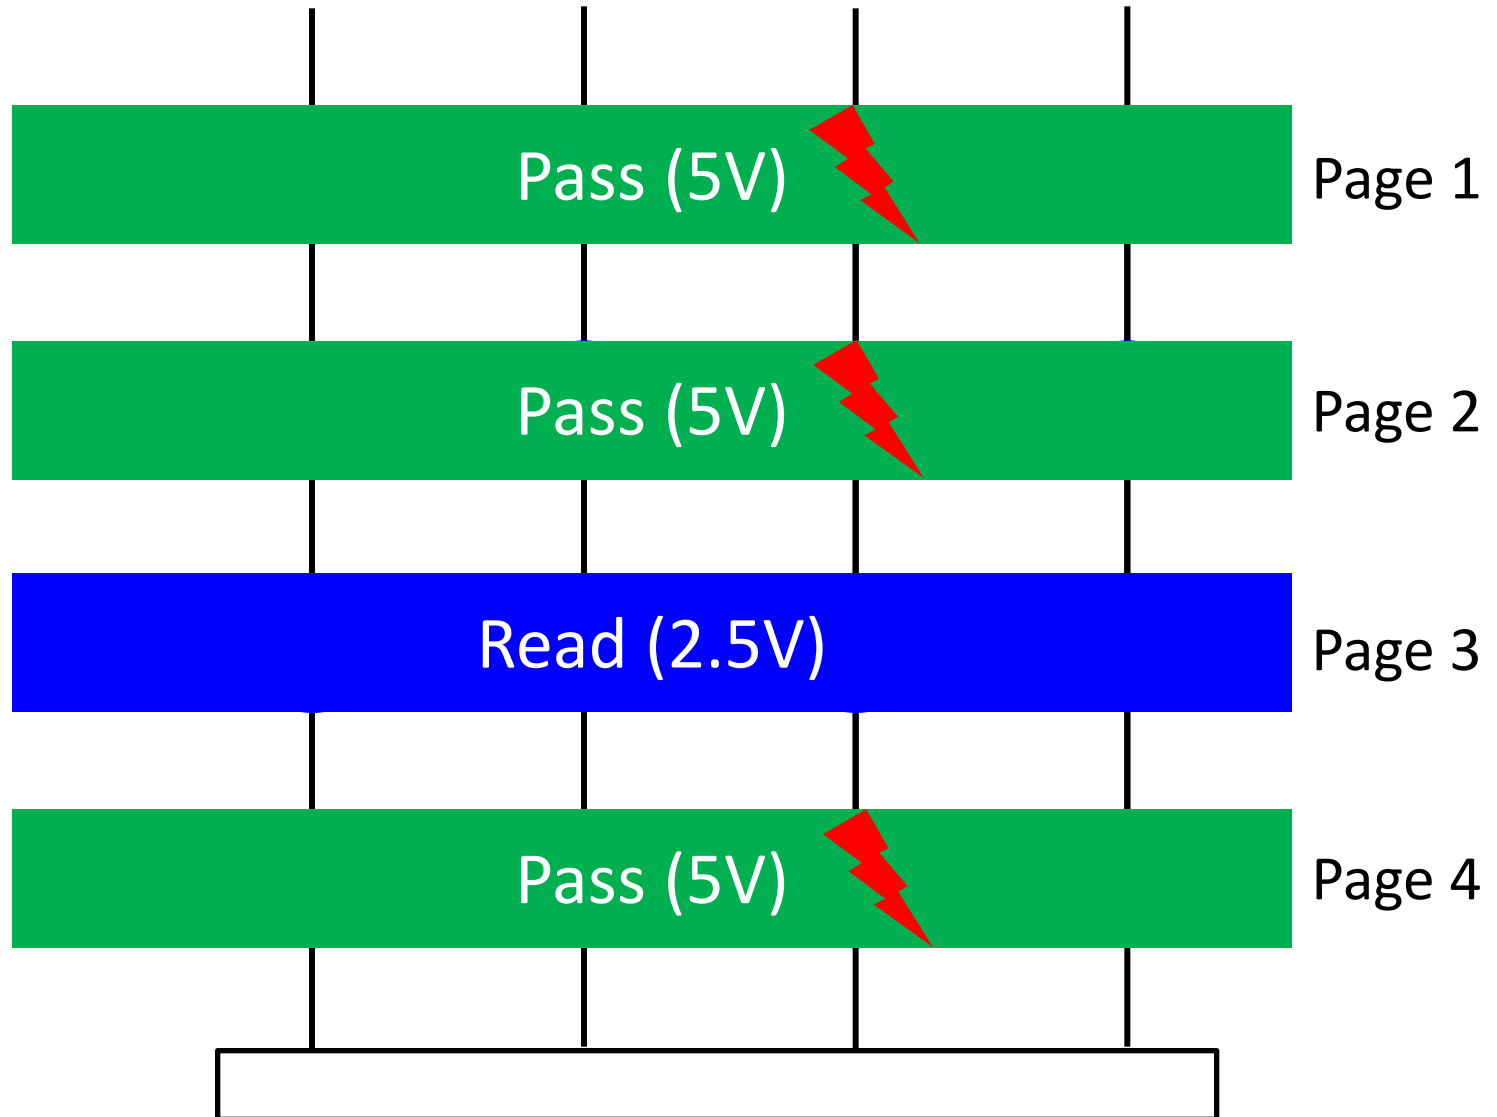




# Read from Flash Cell Array




# Read Disturb Problem: “Weak Programming” Effect



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

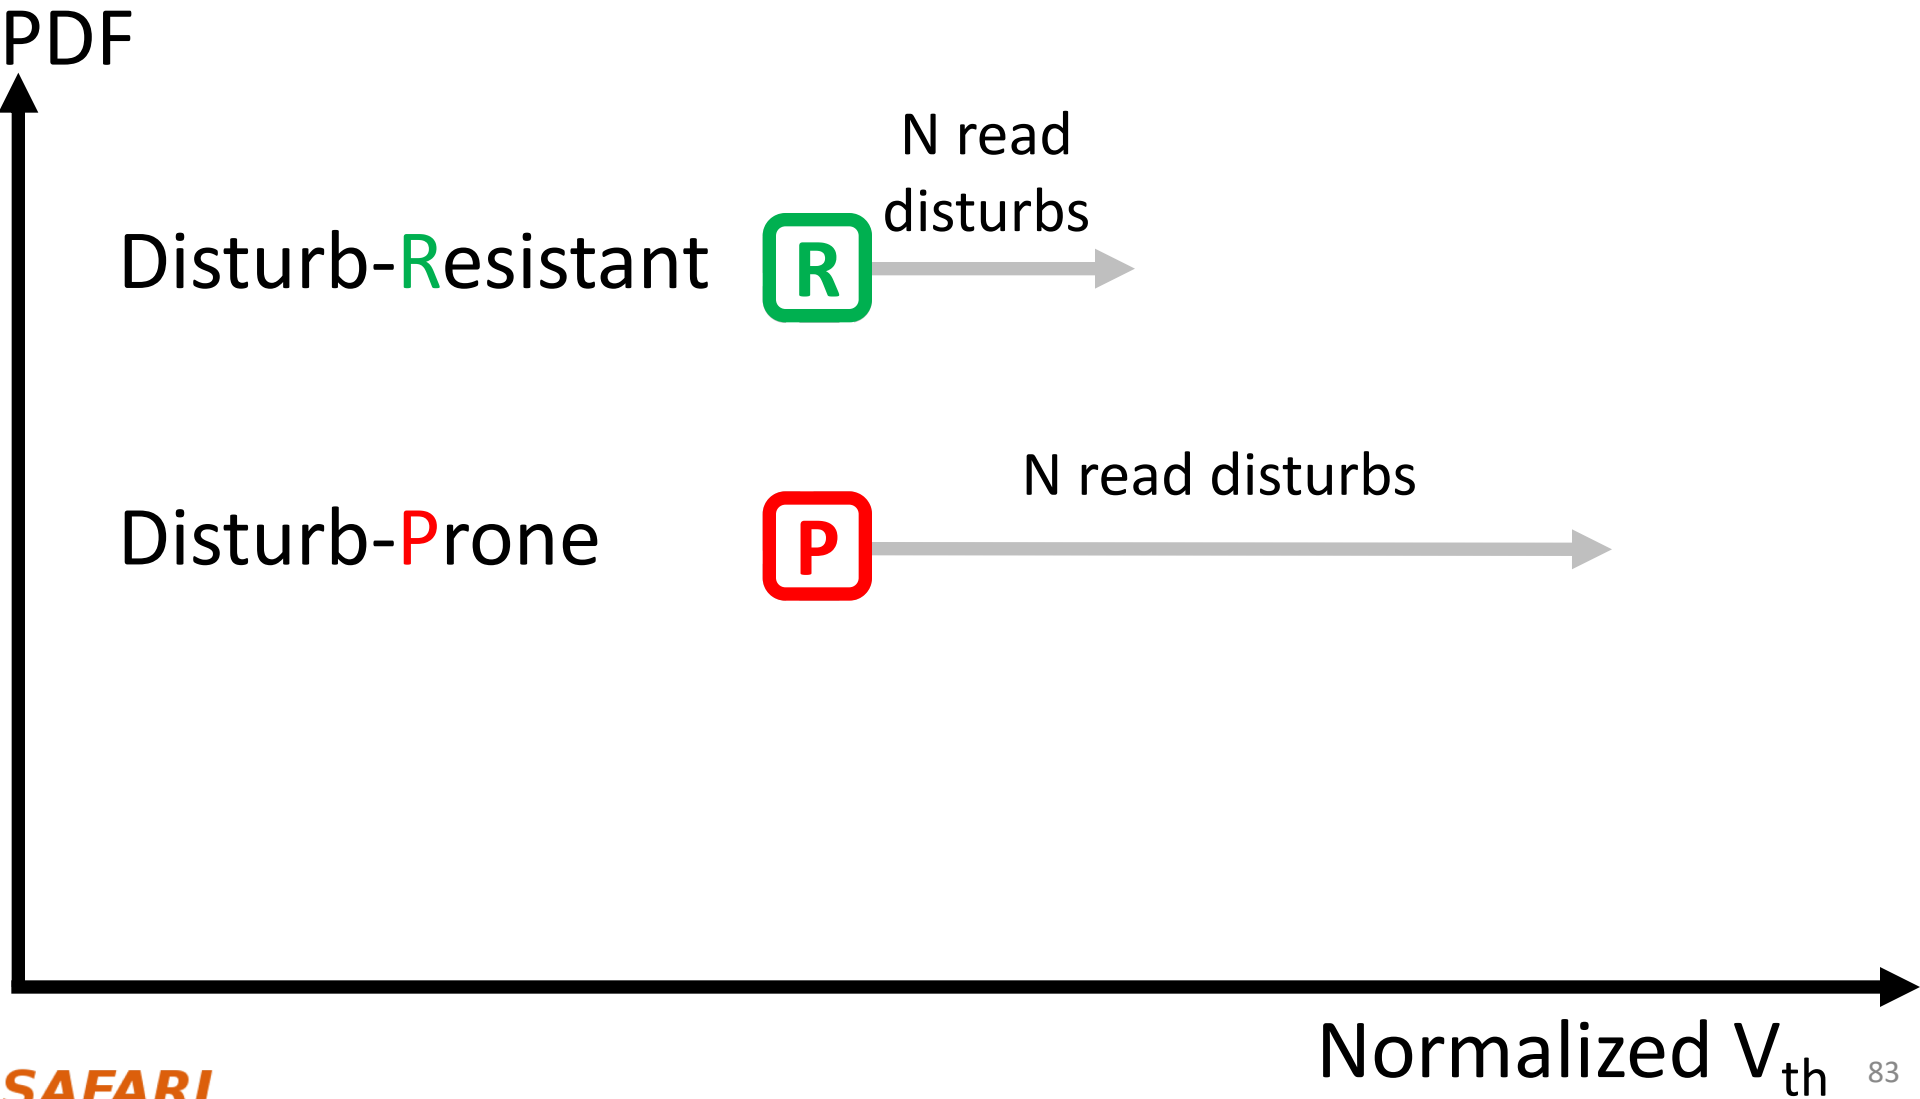
# Read Disturb Problem: “Weak Programming” Effect



# Executive Summary [DSN'15]

- **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%**

# Read Disturb Prone vs. Resistant Cells




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

After 250K read

disturbs:

Disturb-prone cells have higher threshold voltages

Disturb-resistant cells have lower threshold voltages





# 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

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

# More on Flash Read Disturb Errors [DSN'15]

---

- 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.

## Read Disturb Errors in MLC NAND Flash Memory: Characterization, Mitigation, and Recovery

Yu Cai, Yixin Luo, Saugata Ghose, Erich F. Haratsch\*, Ken Mai, Onur Mutlu  
Carnegie Mellon University, \*Seagate Technology  
[yucaicai@gmail.com](mailto:yucaicai@gmail.com), {[yixinluo](mailto:yixinluo@cmu.edu), [ghose](mailto:ghose@cmu.edu), [kenmai](mailto:kenmai@cmu.edu), [onur](mailto:onur@cmu.edu)}@cmu.edu

# Data Retention in Flash Memory

---

- Yu Cai, Yixin Luo, Erich F. Haratsch, Ken Mai, and Onur Mutlu,  
**"Data Retention in MLC NAND Flash Memory: Characterization, Optimization and Recovery"**  
*Proceedings of the 21st International Symposium on High-Performance Computer Architecture (HPCA)*, Bay Area, CA, February 2015.  
[[Slides \(pptx\)](#)] [[pdf](#)]

## 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

yucaicai@gmail.com, yixinluo@cs.cmu.edu, erich.haratsch@lsi.com, {kenmai, omutlu}@ece.cmu.edu

# Large-Scale SSD Error Analysis [SIGMETRICS'15]

---

- First large-scale field study of flash memory errors
- Justin Meza, Qiang Wu, Sanjeev Kumar, and Onur Mutlu,  
**"A Large-Scale Study of Flash Memory Errors in the Field"**  
*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, Portland, OR, June 2015.  
[[Slides \(pptx\)](#)] [[pdf](#)] [[Coverage at ZDNet](#)] [[Coverage on The Register](#)]  
[[Coverage on TechSpot](#)] [[Coverage on The Tech Report](#)]

## A Large-Scale Study of Flash Memory Failures in the Field

Justin Meza  
Carnegie Mellon University  
meza@cmu.edu

Qiang Wu  
Facebook, Inc.  
qw@fb.com

Sanjeev Kumar  
Facebook, Inc.  
skumar@fb.com

Onur Mutlu  
Carnegie Mellon University  
onur@cmu.edu