

The Story of RowHammer

Onur Mutlu

omutlu@gmail.com

<https://people.inf.ethz.ch/omutlu>

20 September 2022

University of British Columbia

SAFARI

ETH zürich

Carnegie Mellon

How Reliable/Secure/Safe is This Bridge?



Collapse of the “Galloping Gertie”



How Secure Are These People?



Security is about preventing unforeseen consequences

What Is RowHammer?

- 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

An “Early” Position Paper [IMW’13]


- Onur Mutlu,
"Memory Scaling: A Systems Architecture Perspective"
*Proceedings of the 5th International Memory Workshop (**IMW**), Monterey, CA, May 2013. Slides
(pptx) (pdf)
EETimes Reprint*

Memory Scaling: A Systems Architecture Perspective

Onur Mutlu
Carnegie Mellon University
onur@cmu.edu
<http://users.ece.cmu.edu/~omutlu/>

The DRAM Scaling Problem


- DRAM stores charge in a capacitor (charge-based memory)
 - Capacitor must be large enough for reliable sensing
 - Access transistor should be large enough for low leakage and high retention time
 - Scaling beyond 40-35nm (2013) is challenging [ITRS, 2009]



- DRAM capacity, cost, and energy/power hard to scale

As Memory Scales, It Becomes Unreliable

- Data from all of Facebook's servers worldwide
- Meza+, "Revisiting Memory Errors in Large-Scale Production Data Centers," DSN'15.



*Intuition:
quadratic
increase
in
capacity*

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.

Infrastructures to Understand Such Issues




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)



SoftMC: Open Source DRAM Infrastructure

- Hasan Hassan et al., “**SoftMC: A Flexible and Practical Open-Source Infrastructure for Enabling Experimental DRAM Studies**,” HPCA 2017.
- Flexible
- Easy to Use (C++ API)
- Open-source
github.com/CMU-SAFARI/SoftMC



SoftMC: Open Source DRAM Infrastructure

- <https://github.com/CMU-SAFARI/SoftMC>


SoftMC: A Flexible and Practical Open-Source Infrastructure for Enabling Experimental DRAM Studies

Hasan Hassan^{1,2,3} Nandita Vijaykumar³ Samira Khan^{4,3} Saugata Ghose³ Kevin Chang³
Gennady Pekhimenko^{5,3} Donghyuk Lee^{6,3} Oguz Ergin² Onur Mutlu^{1,3}

¹*ETH Zürich* ²*TOBB University of Economics & Technology* ³*Carnegie Mellon University*
⁴*University of Virginia* ⁵*Microsoft Research* ⁶*NVIDIA Research*

Data Retention in Memory [Liu et al., ISCA 2013]

- Retention Time Profile of DRAM looks like this:



RAIDR: Heterogeneous Refresh [ISCA'12]

- Jamie Liu, Ben Jaiyen, Richard Veras, and Onur Mutlu,
"RAIDR: Retention-Aware Intelligent DRAM Refresh"
*Proceedings of the 39th International Symposium on
Computer Architecture (ISCA)*, Portland, OR, June 2012.
Slides (pdf)

RAIDR: Retention-Aware Intelligent DRAM Refresh

Jamie Liu Ben Jaiyen Richard Veras Onur Mutlu
Carnegie Mellon University

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^{*}
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
jamiel@alumni.cmu.edu

Ben Jaiyen^{*}
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
bjaiyen@alumni.cmu.edu

Yoongu Kim
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
yoonguk@ece.cmu.edu

Chris Wilkerson
Intel Corporation
2200 Mission College Blvd.
Santa Clara, CA 95054
chris.wilkerson@intel.com

Onur Mutlu
Carnegie Mellon University
5000 Forbes Ave.
Pittsburgh, PA 15213
onur@cmu.edu

Mitigation of Retention Issues [SIGMETRICS'14]

- Samira Khan, Donghyuk Lee, Yoongu Kim, Alaa Alameldeen, Chris Wilkerson, and Onur Mutlu,
"The Efficacy of Error Mitigation Techniques for DRAM Retention Failures: A Comparative Experimental Study"
*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Austin, TX, June 2014. [[Slides \(pptx\)](#)] [[pdf](#)] [[Poster \(pptx\)](#)] [[pdf](#)] [[Full data sets](#)]*

The Efficacy of Error Mitigation Techniques for DRAM Retention Failures: A Comparative Experimental Study

Samira Khan^{†*}
samirakhan@cmu.edu

Donghyuk Lee[†]
donghyuk1@cmu.edu

Yoongu Kim[†]
yoongukim@cmu.edu

Alaa R. Alameldeen^{*}
alaa.r.alameldeen@intel.com

Chris Wilkerson^{*}
chris.wilkerson@intel.com

Onur Mutlu[†]
onur@cmu.edu

[†]Carnegie Mellon University

^{*}Intel Labs

Mitigation of Retention Issues [DSN'15]

- Moinuddin Qureshi, Dae Hyun Kim, Samira Khan, Prashant Nair, and Onur Mutlu,
"AVATAR: A Variable-Retention-Time (VRT) Aware Refresh for DRAM Systems"
Proceedings of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Rio de Janeiro, Brazil, June 2015.
[[Slides \(pptx\)](#)] [[pdf](#)]

AVATAR: A Variable-Retention-Time (VRT) Aware Refresh for DRAM Systems

Moinuddin K. Qureshi [†]	Dae-Hyun Kim [†]	Samira Khan [‡]	Prashant J. Nair [†]	Onur Mutlu [‡]
[†] Georgia Institute of Technology {moin, dhkim, pnair6}@ece.gatech.edu			[‡] Carnegie Mellon University {samirakhan, onur}@cmu.edu	

Mitigation of Retention Issues [DSN'16]

- Samira Khan, Donghyuk Lee, and Onur Mutlu,
"PARBOR: An Efficient System-Level Technique to Detect Data-Dependent Failures in DRAM"
Proceedings of the 45th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Toulouse, France, June 2016.
[[Slides \(pptx\)](#)] [[pdf](#)]

PARBOR: An Efficient System-Level Technique to Detect Data-Dependent Failures in DRAM

Samira Khan^{*}

^{*}University of Virginia

Donghyuk Lee^{†‡}

[†]Carnegie Mellon University

Onur Mutlu^{*†}

[‡]Nvidia

^{*}ETH Zürich

Mitigation of Retention Issues [MICRO'17]

- Samira Khan, Chris Wilkerson, Zhe Wang, Alaa R. Alameldeen, Donghyuk Lee, and Onur Mutlu,
"Detecting and Mitigating Data-Dependent DRAM Failures by Exploiting Current Memory Content"
Proceedings of the 50th International Symposium on Microarchitecture (MICRO), Boston, MA, USA, October 2017.
[\[Slides \(pptx\) \(pdf\)\]](#) [\[Lightning Session Slides \(pptx\) \(pdf\)\]](#) [\[Poster \(pptx\) \(pdf\)\]](#)

Detecting and Mitigating Data-Dependent DRAM Failures by Exploiting Current Memory Content

Samira Khan^{*} Chris Wilkerson[†] Zhe Wang[†] Alaa R. Alameldeen[†] Donghyuk Lee[‡] Onur Mutlu^{*}
^{*}University of Virginia [†]Intel Labs [‡]Nvidia Research ^{*}ETH Zürich

Mitigation of Retention Issues [ISCA'17]

- Minesh Patel, Jeremie S. Kim, and Onur Mutlu,
"The Reach Profiler (REAPER): Enabling the Mitigation of DRAM Retention Failures via Profiling at Aggressive Conditions"
Proceedings of the 44th International Symposium on Computer Architecture (ISCA), Toronto, Canada, June 2017.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Session Slides \(pptx\)](#)] [[pdf](#)]
- First experimental analysis of (mobile) LPDDR4 chips
- Analyzes the complex tradeoff space of retention time profiling
- Idea: enable fast and robust profiling at higher refresh intervals & temperatures

The Reach Profiler (REAPER): Enabling the Mitigation of DRAM Retention Failures via Profiling at Aggressive Conditions

Minesh Patel^{§‡} Jeremie S. Kim^{‡§} Onur Mutlu^{§‡}
[§]ETH Zürich [‡]Carnegie Mellon University

Mitigation of Retention Issues [DSN'19]

- Minesh Patel, Jeremie S. Kim, Hasan Hassan, and Onur Mutlu,
"Understanding and Modeling On-Die Error Correction in Modern DRAM: An Experimental Study Using Real Devices"
Proceedings of the 49th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Portland, OR, USA, June 2019.
[Source Code for EINSim, the Error Inference Simulator]
Best paper award.

Understanding and Modeling On-Die Error Correction in Modern DRAM: An Experimental Study Using Real Devices

Minesh Patel[†] Jeremie S. Kim^{‡†} Hasan Hassan[†] Onur Mutlu^{†‡}

[†]*ETH Zürich* [‡]*Carnegie Mellon University*

Mitigation of Retention Issues [MICRO'20]

- Minesh Patel, Jeremie S. Kim, Taha Shahroodi, Hasan Hassan, and Onur Mutlu, **"Bit-Exact ECC Recovery (BEER): Determining DRAM On-Die ECC Functions by Exploiting DRAM Data Retention Characteristics"**

*Proceedings of the 53rd International Symposium on Microarchitecture (**MICRO**), Virtual, October 2020.*

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

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

[[Talk Video](#) (15 minutes)]

[[Lightning Talk Video](#) (1.5 minutes)]

Best paper award.

Bit-Exact ECC Recovery (BEER): Determining DRAM On-Die ECC Functions by Exploiting DRAM Data Retention Characteristics

Minesh Patel[†] Jeremie S. Kim^{††} Taha Shahroodi[†] Hasan Hassan[†] Onur Mutlu^{†‡}

[†]*ETH Zürich* [‡]*Carnegie Mellon University*

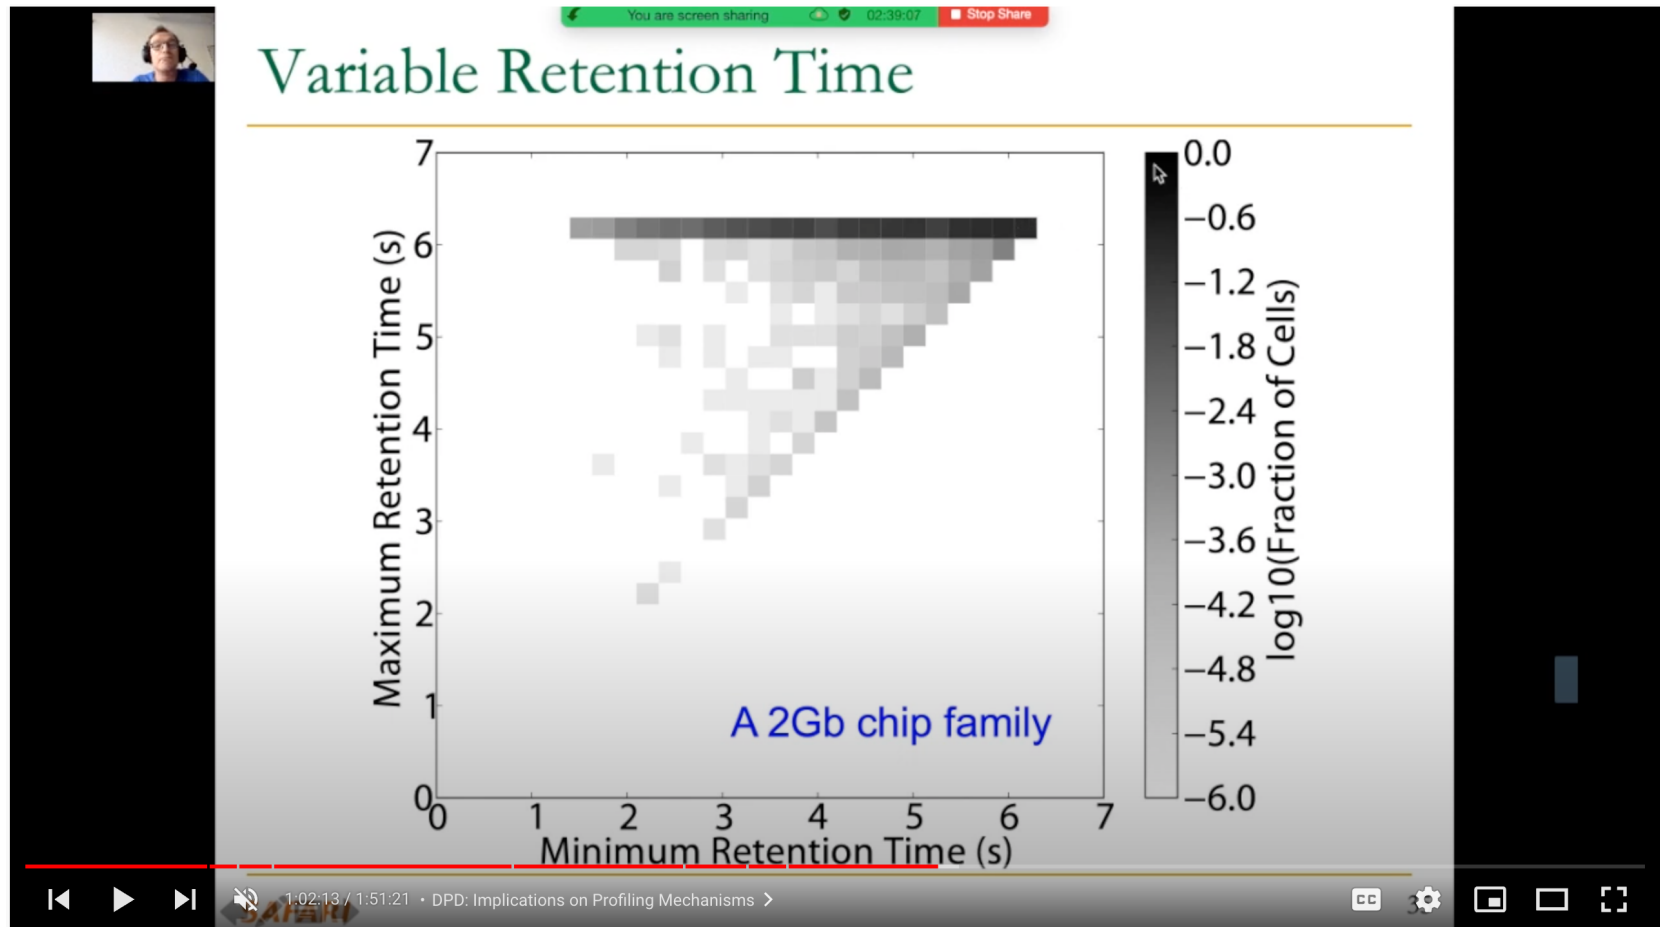
Mitigation of Retention Issues [MICRO'21]

- Minesh Patel, Geraldo F. de Oliveira Jr., and Onur Mutlu,
"HARP: Practically and Effectively Identifying Uncorrectable Errors in Memory Chips That Use On-Die Error-Correcting Codes"
Proceedings of the 54th International Symposium on Microarchitecture (MICRO), Virtual, October 2021.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Short Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (20 minutes)]
[[Lightning Talk Video](#) (1.5 minutes)]
[[HARP Source Code \(Officially Artifact Evaluated with All Badges\)](#)]



HARP: Practically and Effectively Identifying Uncorrectable Errors in Memory Chips That Use On-Die Error-Correcting Codes

More on DRAM Refresh & Data Retention



ETH ZÜRICH

Computer Architecture - Lecture 2b: Data Retention and Memory Refresh (ETH Zürich, Fall 2020)

3,204 views • Sep 19, 2020

43 0 SHARE SAVE ...



Onur Mutlu Lectures
19.1K subscribers


ANALYTICS

EDIT VIDEO

SoftMC: Enabling DRAM Infrastructure

- Hasan Hassan et al., “**SoftMC: A Flexible and Practical Open-Source Infrastructure for Enabling Experimental DRAM Studies**,” HPCA 2017.

- Flexible
- Easy to Use (C++ API)
- Open-source
github.com/CMU-SAFARI/SoftMC



A Curious Phenomenon

A Curious Discovery [Kim et al., ISCA 2014]

One can
predictably induce errors
in most DRAM memory chips

DRAM RowHammer

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


Modern DRAM is Prone to Disturbance Errors




Repeatedly reading a row enough times (before memory gets refreshed) induces **disturbance errors** in adjacent rows in **most real DRAM chips you can buy today**

Most DRAM Modules Are Vulnerable


A company



B company



C company




Up to
 1.0×10^7
errors

Up to
 2.7×10^6
errors

Up to
 3.3×10^5
errors

Recent DRAM Is More Vulnerable



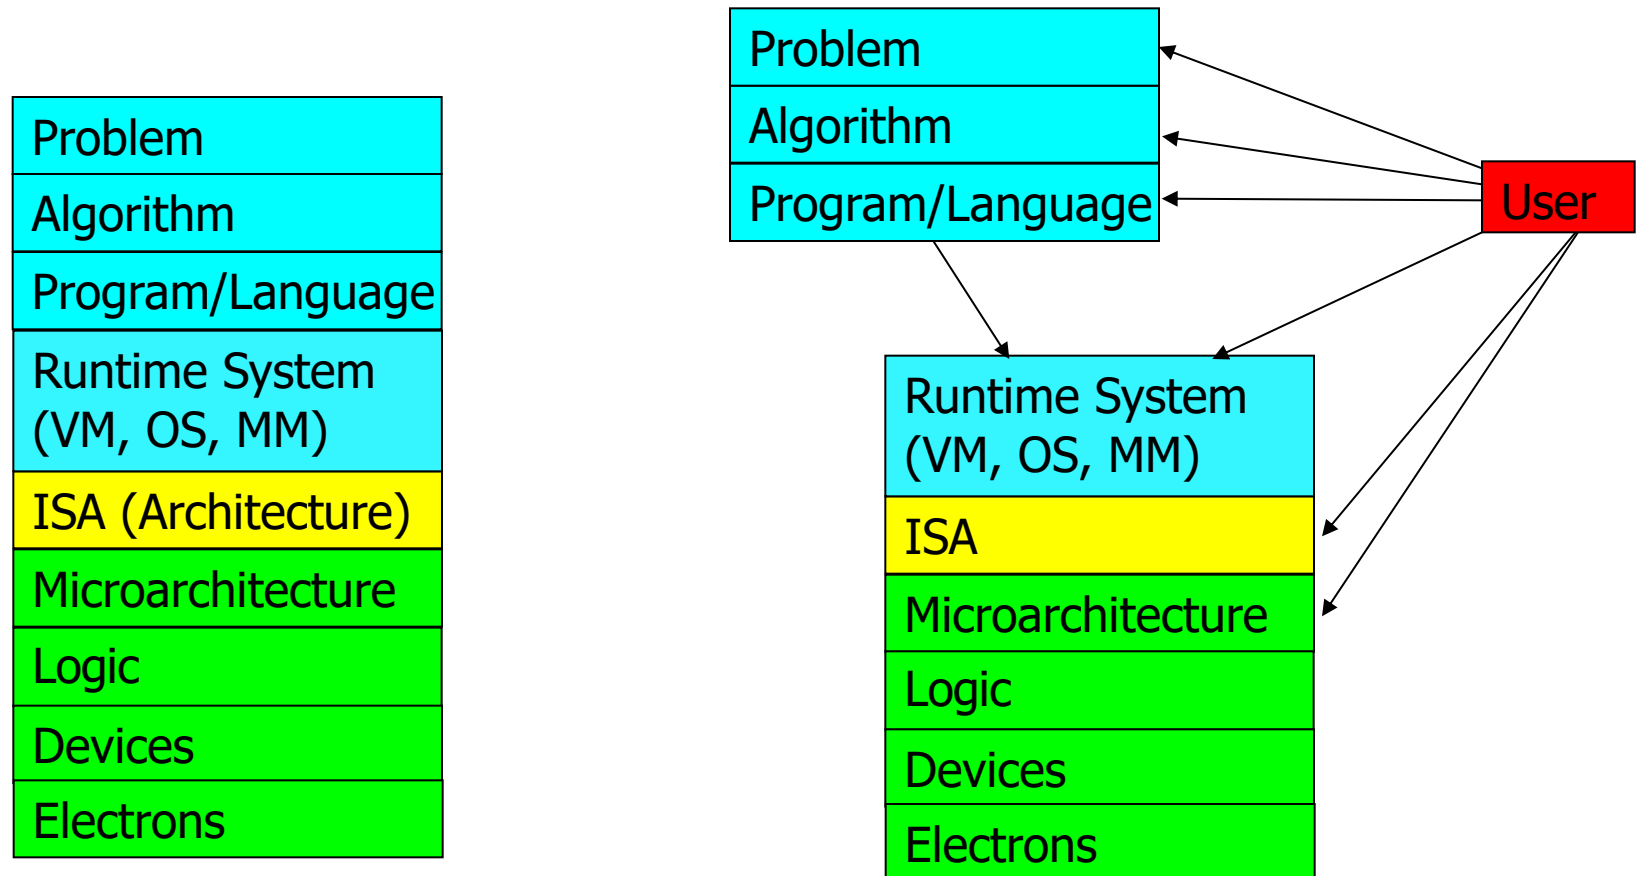
All modules from 2012-2013 are vulnerable

Why Is This Happening?




- DRAM cells are too close to each other!
 - They are not electrically isolated from each other
- Access to one cell affects the value in nearby cells
 - due to **electrical interference** between
 - the cells
 - wires used for accessing the cells
 - Also called cell-to-cell coupling/interference
- Example: When we activate (apply high voltage) to a row, an adjacent row gets slightly activated as well
 - Vulnerable cells in that slightly-activated row lose a little bit of charge
 - If RowHammer happens enough times, charge in such cells gets drained

Higher-Level Implications


- This simple circuit level failure mechanism has enormous implications on upper layers of the transformation hierarchy




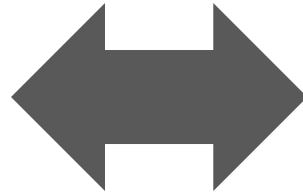
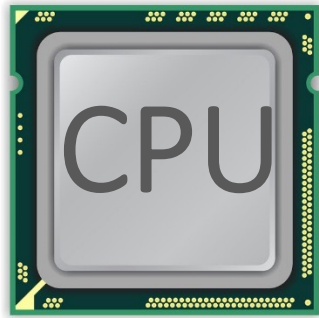
A Simple Program Can Induce Many Errors



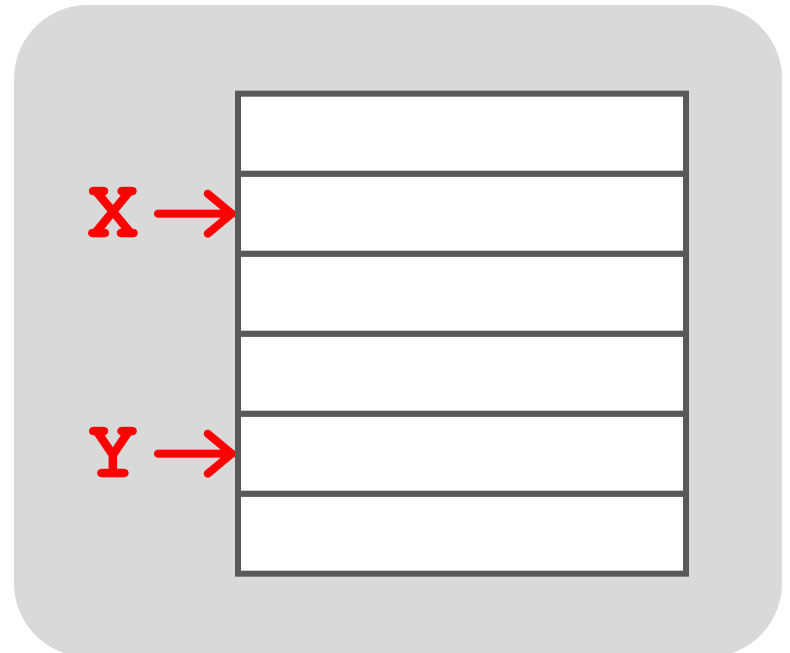
```
loop:  
  mov  (X), %eax  
  mov  (Y), %ebx  
  clflush (X)  
  clflush (Y)  
  mfence  
  jmp  loop
```






A Simple Program Can Induce Many Errors




1. Avoid *cache hits*
 - Flush **X** from cache
2. Avoid *row hits* to **X**
 - Read **Y** in another row



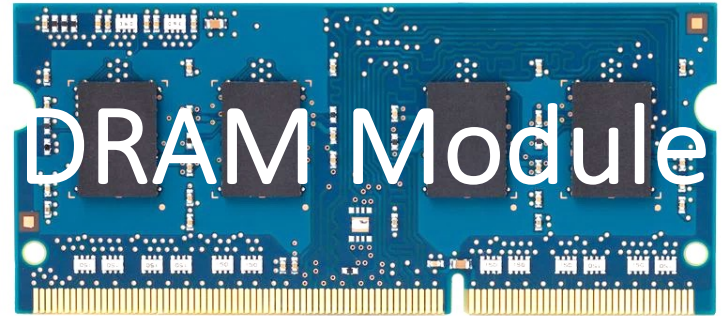
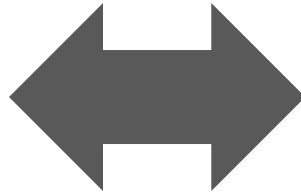
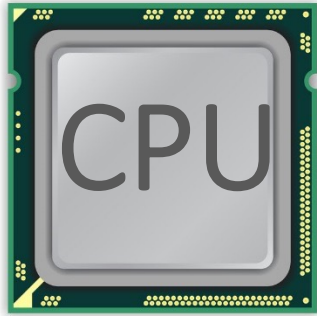
A Simple Program Can Induce Many Errors




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



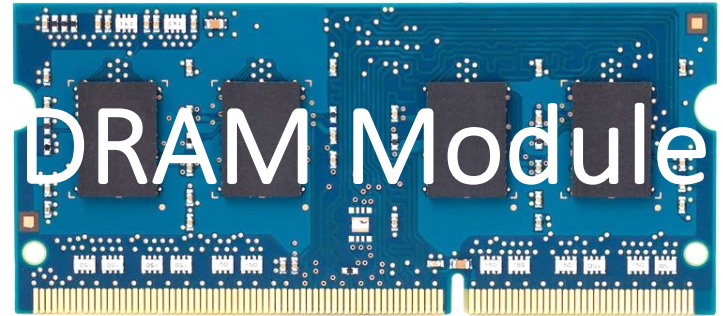
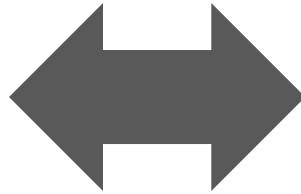
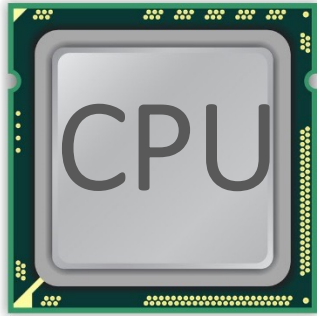
A Simple Program Can Induce Many Errors




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



A Simple Program Can Induce Many Errors



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



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

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



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

More Security Implications (I)

“We can gain unrestricted access to systems of website visitors.”

www.iaik.tugraz.at ■

Not there yet, but ...



ROOT privileges for web apps!

29

Daniel Gruss (@lavados), Clémentine Maurice (@BloodyTangerine),
December 28, 2015 — 32c3, Hamburg, Germany



GATED
COMMUNITIES

Rowhammer.js: A Remote Software-Induced Fault Attack in JavaScript (DIMVA'16)

More Security Implications (II)

"Can gain control of a smart phone deterministically"



Drammer: Deterministic Rowhammer
Attacks on Mobile Platforms, CCS'16 46

More Security Implications (III)

- Using an integrated GPU in a mobile system to remotely escalate privilege via the WebGL interface. [IEEE S&P 2018](#)



TECHNICA

[BIZ & IT](#) [TECH](#) [SCIENCE](#) [POLICY](#) [CARS](#) [GAMING & CULTURE](#)

"GRAND PWINING UNIT" —

Drive-by Rowhammer attack uses GPU to compromise an Android phone

JavaScript based GLitch pwns browsers by flipping bits inside memory chips.

DAN GOODIN - 5/3/2018, 12:00 PM

Grand Pwning Unit: Accelerating Microarchitectural Attacks with the GPU

Pietro Frigo
Vrije Universiteit
Amsterdam
p.frigo@vu.nl

Cristiano Giuffrida
Vrije Universiteit
Amsterdam
giuffrida@cs.vu.nl

Herbert Bos
Vrije Universiteit
Amsterdam
herbertb@cs.vu.nl

Kaveh Razavi
Vrije Universiteit
Amsterdam
kaveh@cs.vu.nl

More Security Implications (IV)

- Rowhammer over RDMA (I) [USENIX ATC 2018](#)

 ars TECHNICA

[BIZ & IT](#) [TECH](#) [SCIENCE](#) [POLICY](#) [CARS](#) [GAMING & CULTURE](#)

THROWHAMMER —

Packets over a LAN are all it takes to trigger serious Rowhammer bit flips

The bar for exploiting potentially serious DDR weakness keeps getting lower.

DAN GOODIN - 5/10/2018, 5:26 PM

Throwhammer: Rowhammer Attacks over the Network and Defenses

Andrei Tatar
VU Amsterdam

Radhesh Krishnan
VU Amsterdam

Elias Athanasopoulos
University of Cyprus

Cristiano Giuffrida
VU Amsterdam

Herbert Bos
VU Amsterdam

Kaveh Razavi
VU Amsterdam

More Security Implications (V)

■ Rowhammer over RDMA (II)



Nethammer—Exploiting DRAM Rowhammer Bug Through Network Requests



Nethammer: Inducing Rowhammer Faults through Network Requests

Moritz Lipp
Graz University of Technology

Misiker Tadesse Aga
University of Michigan

Michael Schwarz
Graz University of Technology

Daniel Gruss
Graz University of Technology


Clémentine Maurice
Univ Rennes, CNRS, IRISA

Lukas Raab
Graz University of Technology

Lukas Lamster
Graz University of Technology

More Security Implications (VI)

- IEEE S&P 2020



RAMBleed

RAMBleed: Reading Bits in Memory Without Accessing Them

Andrew Kwong
University of Michigan
ankwong@umich.edu

Daniel Genkin
University of Michigan
genkin@umich.edu

Daniel Gruss
Graz University of Technology
daniel.gruss@iaik.tugraz.at

Yuval Yarom
University of Adelaide and Data61
yval@cs.adelaide.edu.au

More Security Implications (VII)

■ USENIX Security 2019

Terminal Brain Damage: Exposing the Graceless Degradation in Deep Neural Networks Under Hardware Fault Attacks

Sanghyun Hong, Pietro Frigo[†], Yiğitcan Kaya, Cristiano Giuffrida[†], Tudor Dumitraş

University of Maryland, College Park

[†]Vrije Universiteit Amsterdam



A Single Bit-flip Can Cause Terminal Brain Damage to DNNs

One specific bit-flip in a DNN's representation leads to accuracy drop over 90%

Our research found that a specific bit-flip in a DNN's bitwise representation can cause the accuracy loss up to 90%, and the DNN has 40-50% parameters, on average, that can lead to the accuracy drop over 10% when individually subjected to such single bitwise corruptions...

[Read More](#)

More Security Implications (VIII)

■ USENIX Security 2020

DeepHammer: Depleting the Intelligence of Deep Neural Networks through Targeted Chain of Bit Flips

Fan Yao

University of Central Florida

fan.yao@ucf.edu

Adnan Siraj Rakin

Arizona State University

asrakin@asu.edu

Deliang Fan


Arizona State University

dfan@asu.edu

Degrade the inference accuracy to the level of Random Guess

Example: ResNet-20 for CIFAR-10, 10 output classes

Before attack, **Accuracy: 90.2%** After attack, **Accuracy: ~10% (1/10)**



More Security Implications (IX)

- Rowhammer on MLC NAND Flash (based on [Cai+, HPCA 2017])



Security

Rowhammer RAM attack adapted to hit flash storage

Project Zero's two-year-old dog learns a new trick

By [Richard Chirgwin](#) 17 Aug 2017 at 04:27

17 SHARE ▼

**From random block corruption to privilege escalation:
A filesystem attack vector for rowhammer-like attacks**

Anil Kurmus

Nikolas Ioannou

Matthias Neugschwandtner

Nikolaos Papandreou

Thomas Parnell

IBM Research – Zurich

More Security Implications?



A RowHammer Survey Across the Stack

- 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](#)]
[[Slides from COSADE 2019 \(pptx\)](#)]
[[Slides from VLSI-SOC 2020 \(pptx\) \(pdf\)](#)]
[[Talk Video](#) (1 hr 15 minutes, with Q&A)]

RowHammer: A Retrospective

Onur Mutlu^{§‡} Jeremie S. Kim^{‡§}
[§]ETH Zürich [‡]Carnegie Mellon University

Understanding RowHammer

First RowHammer Analysis


- Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu,
"Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors"
Proceedings of the 41st International Symposium on Computer Architecture (ISCA), Minneapolis, MN, June 2014.
[[Slides \(pptx\) \(pdf\)](#)] [[Lightning Session Slides \(pptx\) \(pdf\)](#)] [[Source Code and Data](#)] [[Lecture Video](#) (1 hr 49 mins), 25 September 2020]
One of the 7 papers of 2012-2017 selected as Top Picks in Hardware and Embedded Security for IEEE TCAD ([link](#)).

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

Yoongu Kim¹ Ross Daly* Jeremie Kim¹ Chris Fallin* Ji Hye Lee¹
Donghyuk Lee¹ Chris Wilkerson² Konrad Lai Onur Mutlu¹

¹Carnegie Mellon University ²Intel Labs

RowHammer Infrastructure (2012-2014)



Tested DRAM Modules from 2008-2014 (129 total)

Manufacturer	Module	Date*	Timing†		Organization		Chip			Victims-per-Module			RI _{th} (ms)
		(yy-ww)	Freq (MT/s)	t _{RC} (ns)	Size (GB)	Chips	Size (Gb)‡	Pins	DieVersion§	Average	Minimum	Maximum	Min
Total of 43 Modules	A ₁	10-08	1066	50.625	0.5	4	1 × 16	B	0	0	0	–	
	A ₂	10-20	1066	50.625	1	8	1 × 8	F	0	0	0	–	
	A ₃₋₅	10-20	1066	50.625	0.5	4	1 × 16	B	0	0	0	–	
	A ₆₋₇	11-24	1066	49.125	1	4	2 × 16	D	7.8 × 10 ¹	5.2 × 10 ¹	1.0 × 10 ²	21.3	
	A ₈₋₁₂	11-26	1066	49.125	1	4	2 × 16	D	2.4 × 10 ²	5.4 × 10 ¹	4.4 × 10 ²	16.4	
	A ₁₃₋₁₄	11-50	1066	49.125	1	4	2 × 16	D	8.8 × 10 ¹	1.7 × 10 ¹	1.6 × 10 ²	26.2	
	A ₁₅₋₁₆	12-22	1600	50.625	1	4	2 × 16	D	9.5	9	1.0 × 10 ¹	34.4	
	A ₁₇₋₁₈	12-26	1600	49.125	2	8	2 × 8	M	1.2 × 10 ²	3.7 × 10 ¹	2.0 × 10 ²	21.3	
	A ₁₉₋₃₀	12-40	1600	48.125	2	8	2 × 8	K	8.6 × 10 ⁶	7.0 × 10 ⁶	1.0 × 10 ⁷	8.2	
	A ₃₁₋₃₄	13-02	1600	48.125	2	8	2 × 8	–	1.8 × 10 ⁶	1.0 × 10 ⁶	3.5 × 10 ⁶	11.5	
	A ₃₅₋₃₆	13-14	1600	48.125	2	8	2 × 8	–	4.0 × 10 ¹	1.9 × 10 ¹	6.1 × 10 ¹	21.3	
	A ₃₇₋₃₈	13-20	1600	48.125	2	8	2 × 8	K	1.7 × 10 ⁶	1.4 × 10 ⁶	2.0 × 10 ⁶	9.8	
	A ₃₉₋₄₀	13-28	1600	48.125	2	8	2 × 8	K	5.7 × 10 ⁴	5.4 × 10 ⁴	6.0 × 10 ⁴	16.4	
	A ₄₁	14-04	1600	49.125	2	8	2 × 8	–	2.7 × 10 ⁵	2.7 × 10 ⁵	2.7 × 10 ⁵	18.0	
	A ₄₂₋₄₃	14-04	1600	48.125	2	8	2 × 8	K	0.5	0	1	62.3	
Total of 54 Modules	B ₁	08-49	1066	50.625	1	8	1 × 8	D	0	0	0	–	
	B ₂	09-49	1066	50.625	1	8	1 × 8	E	0	0	0	–	
	B ₃	10-19	1066	50.625	1	8	1 × 8	F	0	0	0	–	
	B ₄	10-31	1333	49.125	2	8	2 × 8	C	0	0	0	–	
	B ₅	11-13	1333	49.125	2	8	2 × 8	C	0	0	0	–	
	B ₆	11-16	1066	50.625	1	8	1 × 8	F	0	0	0	–	
	B ₇	11-19	1066	50.625	1	8	1 × 8	F	0	0	0	–	
	B ₈	11-25	1333	49.125	2	8	2 × 8	C	0	0	0	–	
	B ₉	11-37	1333	49.125	2	8	2 × 8	D	1.9 × 10 ⁶	1.9 × 10 ⁶	1.9 × 10 ⁶	11.5	
	B ₁₀₋₁₂	11-46	1333	49.125	2	8	2 × 8	D	2.2 × 10 ⁶	1.5 × 10 ⁶	2.7 × 10 ⁶	11.5	
	B ₁₃	11-49	1333	49.125	2	8	2 × 8	C	0	0	0	–	
	B ₁₄	12-01	1866	47.125	2	8	2 × 8	D	9.1 × 10 ⁵	9.1 × 10 ⁵	9.1 × 10 ⁵	9.8	
	B ₁₅₋₃₁	12-10	1866	47.125	2	8	2 × 8	D	9.8 × 10 ⁵	7.8 × 10 ⁵	1.2 × 10 ⁶	11.5	
	B ₃₂	12-25	1600	48.125	2	8	2 × 8	E	7.4 × 10 ⁵	7.4 × 10 ⁵	7.4 × 10 ⁵	11.5	
	B ₃₃₋₄₂	12-28	1600	48.125	2	8	2 × 8	E	5.2 × 10 ⁵	1.9 × 10 ⁵	7.3 × 10 ⁵	11.5	
Total of 32 Modules	B ₄₃₋₄₇	12-31	1600	48.125	2	8	2 × 8	E	4.0 × 10 ⁵	2.9 × 10 ⁵	5.5 × 10 ⁵	13.1	
	B ₄₈₋₅₁	13-19	1600	48.125	2	8	2 × 8	E	1.1 × 10 ⁵	7.4 × 10 ⁴	1.4 × 10 ⁵	14.7	
	B ₅₂₋₅₃	13-40	1333	49.125	2	8	2 × 8	D	2.6 × 10 ⁴	2.3 × 10 ⁴	2.9 × 10 ⁴	21.3	
	B ₅₄	14-07	1333	49.125	2	8	2 × 8	D	7.5 × 10 ³	7.5 × 10 ³	7.5 × 10 ³	26.2	
	C ₁	10-18	1333	49.125	2	8	2 × 8	A	0	0	0	–	
	C ₂	10-20	1066	50.625	2	8	2 × 8	A	0	0	0	–	
	C ₃	10-22	1066	50.625	2	8	2 × 8	A	0	0	0	–	
	C ₄₋₅	10-26	1333	49.125	2	8	2 × 8	B	8.9 × 10 ²	6.0 × 10 ²	1.2 × 10 ³	29.5	
	C ₆	10-43	1333	49.125	1	8	1 × 8	T	0	0	0	–	
	C ₇	10-51	1333	49.125	2	8	2 × 8	B	4.0 × 10 ²	4.0 × 10 ²	4.0 × 10 ²	29.5	
	C ₈	11-12	1333	46.25	2	8	2 × 8	B	6.9 × 10 ²	6.9 × 10 ²	6.9 × 10 ²	21.3	
	C ₉	11-19	1333	46.25	2	8	2 × 8	B	9.2 × 10 ²	9.2 × 10 ²	9.2 × 10 ²	27.9	
	C ₁₀	11-31	1333	49.125	2	8	2 × 8	B	3	3	3	39.3	
	C ₁₁	11-42	1333	49.125	2	8	2 × 8	B	1.6 × 10 ²	1.6 × 10 ²	1.6 × 10 ²	39.3	
	C ₁₂	11-48	1600	48.125	2	8	2 × 8	C	7.1 × 10 ⁴	7.1 × 10 ⁴	7.1 × 10 ⁴	19.7	
	C ₁₃	12-08	1333	49.125	2	8	2 × 8	C	3.9 × 10 ⁴	3.9 × 10 ⁴	3.9 × 10 ⁴	21.3	
	C ₁₄₋₁₅	12-12	1333	49.125	2	8	2 × 8	C	3.7 × 10 ⁴	2.1 × 10 ⁴	5.4 × 10 ⁴	21.3	
	C ₁₆₋₁₈	12-20	1600	48.125	2	8	2 × 8	C	3.5 × 10 ³	1.2 × 10 ³	7.0 × 10 ³	27.9	
	C ₁₉	12-23	1600	48.125	2	8	2 × 8	E	1.4 × 10 ⁵	1.4 × 10 ⁵	1.4 × 10 ⁵	18.0	
	C ₂₀	12-24	1600	48.125	2	8	2 × 8	C	6.5 × 10 ⁴	6.5 × 10 ⁴	6.5 × 10 ⁴	21.3	
	C ₂₁	12-26	1600	48.125	2	8	2 × 8	C	2.3 × 10 ⁴	2.3 × 10 ⁴	2.3 × 10 ⁴	24.6	
	C ₂₂	12-32	1600	48.125	2	8	2 × 8	C	1.7 × 10 ⁴	1.7 × 10 ⁴	1.7 × 10 ⁴	22.9	
	C ₂₃₋₂₄	12-37	1600	48.125	2	8	2 × 8	C	2.3 × 10 ⁴	1.1 × 10 ⁴	3.4 × 10 ⁴	18.0	
	C ₂₅₋₃₀	12-41	1600	48.125	2	8	2 × 8	C	2.0 × 10 ⁴	1.1 × 10 ⁴	3.2 × 10 ⁴	19.7	
	C ₃₁	13-11	1600	48.125	2	8	2 × 8	C	3.3 × 10 ⁵	3.3 × 10 ⁵	3.3 × 10 ⁵	14.7	
	C ₃₂	13-35	1600	48.125	2	8	2 × 8	C	3.7 × 10 ⁴	3.7 × 10 ⁴	3.7 × 10 ⁴	21.3	

* We report the manufacture date marked on the chip packages, which is more accurate than other dates that can be gleaned from a module.

† We report timing constraints stored in the module's on-board ROM [33], which is read by the system BIOS to calibrate the memory controller.

‡ The maximum DRAM chip size supported by our testing platform is 2Gb.


§ We report DRAM die versions marked on the chip packages, which typically progress in the following manner: $\mathcal{M} \rightarrow \mathcal{A} \rightarrow \mathcal{B} \rightarrow \mathcal{C} \rightarrow \dots$.

Table 3. Sample population of 129 DDR3 DRAM modules, categorized by manufacturer and sorted by manufacture date

RowHammer Characterization Results

1. Most Modules Are at Risk
2. Errors vs. Vintage
3. Error = Charge Loss
4. Adjacency: Aggressor & Victim
5. Sensitivity Studies
6. Other Results in Paper
7. Solution Space


4. Adjacency: Aggressor & Victim



Note: For three modules with the most errors (only first bank)

Most aggressors & victims are adjacent


① Access Interval (Aggressor)



Note: For three modules with the most errors (only first bank)

Less frequent accesses \rightarrow Fewer errors



2 Refresh Interval



Note: Using three modules with the most errors (only first bank)

More frequent refreshes → Fewer errors

3 Data Pattern



10x Errors

Errors affected by data stored in other cells

6. Other Key Observations [ISCA'14]

- *Victim Cells \neq Retention-Weak Cells*
 - Almost no overlap between them
- *Errors are repeatable*
 - Across ten iterations of testing, >70% of victim cells had errors in every iteration
- *As many as 4 errors per cache-line*
 - Simple ECC (e.g., SECDED) cannot prevent all errors
- *Cells affected by two aggressors on either side*
 - Double sided hammering

Major RowHammer Characteristics (2014)

- Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu,
"Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors"
Proceedings of the 41st International Symposium on Computer Architecture (ISCA), Minneapolis, MN, June 2014.
[[Slides \(pptx\) \(pdf\)](#)] [[Lightning Session Slides \(pptx\) \(pdf\)](#)] [[Source Code and Data](#)] [[Lecture Video](#) (1 hr 49 mins), 25 September 2020]
One of the 7 papers of 2012-2017 selected as Top Picks in Hardware and Embedded Security for IEEE TCAD ([link](#)).

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

Yoongu Kim¹ Ross Daly* Jeremie Kim¹ Chris Fallin* Ji Hye Lee¹
Donghyuk Lee¹ Chris Wilkerson² Konrad Lai Onur Mutlu¹

¹Carnegie Mellon University ²Intel Labs

RowHammer is Getting Much Worse (2020)

- 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^{§†} Minesh Patel[§] A. Giray Yağlıkçı[§]
Hasan Hassan[§] Roknoddin Azizi[§] Lois Orosa[§] Onur Mutlu^{§†}
[§]*ETH Zürich* [†]*Carnegie Mellon University*

New RowHammer Dimensions (2021)

- Lois Orosa, Abdullah Giray Yaglikci, Haocong Luo, Ataberk Olgun, Jisung Park, Hasan Hassan, Minesh Patel, Jeremie S. Kim, and Onur Mutlu,
"A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses"
*Proceedings of the 54th International Symposium on Microarchitecture (**MICRO**), Virtual, October 2021.*
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Short Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (21 minutes)]
[[Lightning Talk Video](#) (1.5 minutes)]
[[arXiv version](#)]

A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses

Lois Orosa*
ETH Zürich

A. Giray Yağlıkçı*
ETH Zürich

Haocong Luo
ETH Zürich

Ataberk Olgun
ETH Zürich, TOBB ETÜ

Jisung Park
ETH Zürich

Hasan Hassan
ETH Zürich

Minesh Patel
ETH Zürich

Jeremie S. Kim
ETH Zürich

Onur Mutlu
ETH Zürich

RowHammer vs. Wordline Voltage (2022)

- A. Giray Yağlıkçı, Haocong Luo, Geraldo F. de Oliveira, Ataberk Olgun, Minesh Patel, Jisung Park, Hasan Hassan, Jeremie S. Kim, Lois Orosa, and Onur Mutlu, **"Understanding RowHammer Under Reduced Wordline Voltage: An Experimental Study Using Real DRAM Devices"**
Proceedings of the 52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Baltimore, MD, USA, June 2022.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]
[[arXiv version](#)]
[[Talk Video](#) (34 minutes, including Q&A)]
[[Lightning Talk Video](#) (2 minutes)]

Understanding RowHammer Under Reduced Wordline Voltage: An Experimental Study Using Real DRAM Devices

A. Giray Yağlıkçı¹ Haocong Luo¹ Geraldo F. de Oliveira¹ Ataberk Olgun¹ Minesh Patel¹
Jisung Park¹ Hasan Hassan¹ Jeremie S. Kim¹ Lois Orosa^{1,2} Onur Mutlu¹
¹*ETH Zürich* ²*Galicia Supercomputing Center (CESGA)*

RowHammer Solutions

Two Types of RowHammer Solutions

■ Immediate

- ❑ To protect the vulnerable DRAM chips in the field
- ❑ Limited possibilities

■ Longer-term

- ❑ To protect future DRAM chips
- ❑ Wider range of protection mechanisms

■ Our ISCA 2014 paper proposes both types of solutions

- ❑ Seven solutions in total
- ❑ PARA proposed as best solution → already employed in the field

Some Potential Solutions (ISCA 2014)

- Make better DRAM chips

Cost

- Refresh frequently

Power, Performance

- Sophisticated ECC

Cost, Power

- Access counters

Cost, Power, Complexity

Apple's Security Patch for RowHammer

- <https://support.apple.com/en-gb/HT204934>

Available for: OS X Mountain Lion v10.8.5, OS X Mavericks v10.9.5

Impact: A malicious application may induce memory corruption to escalate privileges

Description: A disturbance error, also known as Rowhammer, exists with some DDR3 RAM that could have led to memory corruption. This issue was mitigated by increasing memory refresh rates.

CVE-ID

CVE-2015-3693 : Mark Seaborn and Thomas Dullien of Google, working from original research by Yoongu Kim et al (2014)

HP, Lenovo, and many other vendors released similar patches

Our Solution to RowHammer

- **PARA: Probabilistic Adjacent Row Activation**
- **Key Idea**
 - After closing a row, we activate (i.e., refresh) one of its neighbors with a low probability: $p = 0.005$
- **Reliability Guarantee**
 - When $p=0.005$, errors in one year: 9.4×10^{-14}
 - By adjusting the value of p , we can vary the strength of protection against errors


Advantages of PARA

- *PARA refreshes rows infrequently*
 - Low power
 - Low performance-overhead
 - Average slowdown: **0.20%** (for 29 benchmarks)
 - Maximum slowdown: **0.75%**
- *PARA is stateless*
 - Low cost
 - Low complexity
- *PARA is an effective and low-overhead solution to prevent disturbance errors*


Requirements for PARA

- If implemented in **DRAM chip** (done today)
 - Enough slack in timing and refresh parameters
 - Plenty of slack today:
 - Lee et al., “**Adaptive-Latency DRAM: Optimizing DRAM Timing for the Common Case**,” HPCA 2015.
 - Chang et al., “**Understanding Latency Variation in Modern DRAM Chips**,” SIGMETRICS 2016.
 - Lee et al., “**Design-Induced Latency Variation in Modern DRAM Chips**,” SIGMETRICS 2017.
 - Chang et al., “**Understanding Reduced-Voltage Operation in Modern DRAM Devices**,” SIGMETRICS 2017.
 - Ghose et al., “**What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study**,” SIGMETRICS 2018.
 - Kim et al., “**Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines**,” ICCD 2018.
- If implemented in **memory controller**
 - Need coordination between controller and DRAM
 - Memory controller should know which rows are physically adjacent

Probabilistic Activation in Real Life (I)



Probabilistic Activation in Real Life (II)



Seven RowHammer Solutions Proposed

- Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu,
"Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors"
Proceedings of the 41st International Symposium on Computer Architecture (ISCA), Minneapolis, MN, June 2014.
[\[Slides \(pptx\) \(pdf\)\]](#) [\[Lightning Session Slides \(pptx\) \(pdf\)\]](#) [\[Source Code and Data\]](#) [\[Lecture Video\]](#) (1 hr 49 mins), 25 September 2020]
One of the 7 papers of 2012-2017 selected as Top Picks in Hardware and Embedded Security for IEEE TCAD ([link](#)).


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

Yoongu Kim¹ Ross Daly* Jeremie Kim¹ Chris Fallin* Ji Hye Lee¹
Donghyuk Lee¹ Chris Wilkerson² Konrad Lai Onur Mutlu¹

¹Carnegie Mellon University ²Intel Labs

Main Memory Needs
Intelligent Controllers
for Security, Safety,
Reliability, Scaling

Aside: Intelligent Controller for NAND 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]

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>

Detailed Lectures on RowHammer

■ Computer Architecture, Fall 2021, Lecture 5

- RowHammer (ETH Zürich, Fall 2021)
- <https://www.youtube.com/watch?v=7wVKnPj3NVw&list=PL5Q2soXY2Zi-Mnk1PxjEIG32HAGILkTOF&index=5>

■ Computer Architecture, Fall 2021, Lecture 6

- RowHammer and Secure & Reliable Memory (ETH Zürich, Fall 2021)
- <https://www.youtube.com/watch?v=HNd4skQrt6I&list=PL5Q2soXY2Zi-Mnk1PxjEIG32HAGILkTOF&index=6>

<https://www.youtube.com/onurmutlulectures>

First RowHammer Analysis

- Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu,
"Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors"
Proceedings of the 41st International Symposium on Computer Architecture (ISCA), Minneapolis, MN, June 2014.
[[Slides \(pptx\) \(pdf\)](#)] [[Lightning Session Slides \(pptx\) \(pdf\)](#)] [[Source Code and Data](#)] [[Lecture Video](#) (1 hr 49 mins), 25 September 2020]
One of the 7 papers of 2012-2017 selected as Top Picks in Hardware and Embedded Security for IEEE TCAD ([link](#)).

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

Yoongu Kim¹ Ross Daly* Jeremie Kim¹ Chris Fallin* Ji Hye Lee¹
Donghyuk Lee¹ Chris Wilkerson² Konrad Lai Onur Mutlu¹

¹Carnegie Mellon University ²Intel Labs

Retrospective on RowHammer & Future

- Onur Mutlu,
"The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser"

*Invited Paper in Proceedings of the Design, Automation, and Test in Europe Conference (**DATE**), Lausanne, Switzerland, March 2017.*

[Slides (pptx) (pdf)]

The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser

Onur Mutlu
ETH Zürich
onur.mutlu@inf.ethz.ch
<https://people.inf.ethz.ch/omutlu>

A More Recent RowHammer Retrospective

- 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](#)]
[[Slides from COSADE 2019 \(pptx\)](#)]
[[Slides from VLSI-SOC 2020 \(pptx\) \(pdf\)](#)]
[[Talk Video](#) (1 hr 15 minutes, with Q&A)]

RowHammer: A Retrospective

Onur Mutlu^{§‡} Jeremie S. Kim^{‡§}
[§]ETH Zürich [‡]Carnegie Mellon University

RowHammer in 2020-2022

Revisiting RowHammer

RowHammer is Getting Much Worse

- 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^{§†} Minesh Patel[§] A. Giray Yağlıkçı[§]
Hasan Hassan[§] Roknoddin Azizi[§] Lois Orosa[§] Onur Mutlu^{§†}
[§]*ETH Zürich* [†]*Carnegie Mellon University*

Key Takeaways from 1580 Chips


- **Newer DRAM chips are much more vulnerable to RowHammer (more bit flips, happening earlier)**
- There are new chips whose weakest cells fail after **only 4800 hammers**
- Chips of newer DRAM technology nodes can exhibit RowHammer bit flips 1) in **more rows** and 2) **farther away** from the victim row.
- **Existing mitigation mechanisms are NOT effective at future technology nodes**

DRAM Testing Infrastructures

Three separate testing infrastructures

1. **DDR3:** FPGA-based SoftMC [Hassan+, HPCA'17]
(Xilinx ML605)
2. **DDR4:** FPGA-based SoftMC [Hassan+, HPCA'17]
(Xilinx Virtex UltraScale 95)
3. **LPDDR4:** In-house testing hardware for LPDDR4 chips

All provide fine-grained control over DRAM commands, timing parameters and temperature



1580 DRAM Chips Tested


DRAM type-node	Number of Chips (Modules) Tested			
	Mfr. A	Mfr. B	Mfr. C	Total
DDR3-old	56 (10)	88 (11)	28 (7)	172 (28)
DDR3-new	80 (10)	52 (9)	104 (13)	236 (32)
DDR4-old	112 (16)	24 (3)	128 (18)	264 (37)
DDR4-new	264 (43)	16 (2)	108 (28)	388 (73)
LPDDR4-1x	12 (3)	180 (45)	N/A	192 (48)
LPDDR4-1y	184 (46)	N/A	144 (36)	328 (82)

1580 total DRAM chips tested from **300** DRAM modules

- **Three** major DRAM manufacturers {A, B, C}
- **Three** DRAM *types or standards* {DDR3, DDR4, LPDDR4}
 - LPDDR4 chips we test implement on-die ECC
- **Two** technology nodes per DRAM type {old/new, 1x/1y}
 - Categorized based on manufacturing date, datasheet publication date, purchase date, and characterization results

Type-node: configuration describing a chip's type and technology node generation: **DDR3-old/new, DDR4-old/new, LPDDR4-1x/1y**


3. Hammer Count (HC) Effects



RowHammer bit flip rates **increase**
when going **from old to new** DDR4 technology node generations


**RowHammer bit flip rates (i.e., RowHammer vulnerability)
increase with technology node generation**

5. First RowHammer Bit Flips per Chip



Newer chips from each DRAM manufacturer
are more vulnerable to RowHammer

5. First RowHammer Bit Flips per Chip



In a DRAM type, **HC_{first}** reduces significantly from old to new chips, i.e., **DDR3: 69.2k to 22.4k**, **DDR4: 17.5k to 10k**, **LPDDR4: 16.8k to 4.8k**

There are chips whose weakest cells fail after only **4800 hammers**

Newer chips from a given DRAM manufacturer
more vulnerable to RowHammer

RowHammer is Getting Much Worse

- 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^{§†} Minesh Patel[§] A. Giray Yağlıkçı[§]
Hasan Hassan[§] Roknoddin Azizi[§] Lois Orosa[§] Onur Mutlu^{§†}
[§]*ETH Zürich* [†]*Carnegie Mellon University*

Detailed Lecture on Revisiting RowHammer

- Computer Architecture, Fall 2020, Lecture 5b
 - RowHammer in 2020: Revisiting RowHammer (ETH Zürich, Fall 2020)
 - <https://www.youtube.com/watch?v=gR7XR-Eepcg&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=10>

<https://www.youtube.com/onurmutlulectures>

TRRespass

Industry-Adopted Solutions Do Not Work

- 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](#)]
[[Lecture Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#)] (17 minutes)
[[Lecture Video](#)] (59 minutes)
[[Source Code](#)]
[[Web Article](#)]
Best paper award.
Pwnie Award 2020 for Most Innovative Research. [Pwnie Awards 2020](#)


TRRespass: Exploiting the Many Sides of Target Row Refresh

Pietro Frigo^{*†} Emanuele Vannacci^{*†} Hasan Hassan[§] Victor van der Veen[¶]
Onur Mutlu[§] Cristiano Giuffrida^{*} Herbert Bos^{*} Kaveh Razavi^{*}


TRRespass

- First work to show that TRR-protected DRAM chips are vulnerable to RowHammer in the field
 - Mitigations advertised as secure are not secure
- Introduces the Many-sided RowHammer attack
 - Idea: Hammer many rows to bypass TRR mitigations (e.g., by overflowing proprietary TRR tables that detect aggressor rows)
- (Partially) reverse-engineers the TRR and pTRR mitigation mechanisms implemented in DRAM chips and memory controllers
- Provides an automatic tool that can effectively create many-sided RowHammer attacks in DDR4 and LPDDR4(X) chips

Example Many-Sided Hammering Patterns



(a) Assisted double-sided



(b) 4-sided

Fig. 12: Hammering patterns discovered by *TRRespass*. Aggressor rows are in red (■) and victim rows are in blue (■).

BitFlips vs. Number of Aggressor Rows




Fig. 10: Bit flips vs. number of aggressor rows. Module C_{12} : Number of bit flips in bank 0 as we vary the number of aggressor rows. Using SoftMC, we refresh DRAM with standard t_{REFI} and run the tests until each aggressor rows is hammered 500K times.




Fig. 11: Bit flips vs. number of aggressor rows. Module A_{15} : Number of bit flips in bank 0 as we vary the number of aggressor rows. Using SoftMC, we refresh DRAM with standard t_{REFI} and run the tests until each aggressor rows is hammered 500K times.




Fig. 13: Bit flips vs. number of aggressor rows. Module A_{10} : Number of bit flips triggered with N -sided RowHammer for varying number of N on Intel Core i7-7700K. Each aggressor row is one row away from the closest aggressor row (i.e., VAVAVA... configuration) and aggressor rows are hammered in a round-robin fashion.

TRRespass Vulnerable DRAM Modules

TABLE II: *TRRespass* results. We report the number of patterns found and bit flips detected for the 42 DRAM modules in our set.

Module	Date (yy-ww)	Freq. (MHz)	Size (GB)	Organization			MAC	Found Patterns	Best Pattern	Corruptions			Double Refresh
				Ranks	Banks	Pins				Total	1 → 0	0 → 1	
$\mathcal{A}_{0,1,2,3}$	16-37	2132	4	1	16	×8	UL	—	—	—	—	—	—
\mathcal{A}_4	16-51	2132	4	1	16	×8	UL	4	9-sided	7956	4008	3948	—
\mathcal{A}_5	18-51	2400	4	1	8	×16	UL	—	—	—	—	—	—
$\mathcal{A}_{6,7}$	18-15	2666	4	1	8	×16	UL	—	—	—	—	—	—
\mathcal{A}_8	17-09	2400	8	1	16	×8	UL	33	19-sided	20808	10289	10519	—
\mathcal{A}_9	17-31	2400	8	1	16	×8	UL	33	19-sided	24854	12580	12274	—
\mathcal{A}_{10}	19-02	2400	16	2	16	×8	UL	488	10-sided	11342	1809	11533	✓
\mathcal{A}_{11}	19-02	2400	16	2	16	×8	UL	523	10-sided	12830	1682	11148	✓
$\mathcal{A}_{12,13}$	18-50	2666	8	1	16	×8	UL	—	—	—	—	—	—
\mathcal{A}_{14}	19-08 [†]	3200	16	2	16	×8	UL	120	14-sided	32723	16490	16233	—
$\mathcal{A}_{15}^{\ddagger}$	17-08	2132	4	1	16	×8	UL	2	9-sided	22397	12351	10046	—
\mathcal{B}_0	18-11	2666	16	2	16	×8	UL	2	3-sided	17	10	7	—
\mathcal{B}_1	18-11	2666	16	2	16	×8	UL	2	3-sided	22	16	6	—
\mathcal{B}_2	18-49	3000	16	2	16	×8	UL	2	3-sided	5	2	3	—
\mathcal{B}_3	19-08 [†]	3000	8	1	16	×8	UL	—	—	—	—	—	—
$\mathcal{B}_{4,5}$	19-08 [†]	2666	8	2	16	×8	UL	—	—	—	—	—	—
$\mathcal{B}_{6,7}$	19-08 [†]	2400	4	1	16	×8	UL	—	—	—	—	—	—
\mathcal{B}_8^{\diamond}	19-08 [†]	2400	8	1	16	×8	UL	—	—	—	—	—	—
\mathcal{B}_9^{\diamond}	19-08 [†]	2400	8	1	16	×8	UL	2	3-sided	12	—	12	✓
$\mathcal{B}_{10,11}$	16-13 [†]	2132	8	2	16	×8	UL	—	—	—	—	—	—
$\mathcal{C}_{0,1}$	18-46	2666	16	2	16	×8	UL	—	—	—	—	—	—
$\mathcal{C}_{2,3}$	19-08 [†]	2800	4	1	16	×8	UL	—	—	—	—	—	—
$\mathcal{C}_{4,5}$	19-08 [†]	3000	8	1	16	×8	UL	—	—	—	—	—	—
$\mathcal{C}_{6,7}$	19-08 [†]	3000	16	2	16	×8	UL	—	—	—	—	—	—
\mathcal{C}_8	19-08 [†]	3200	16	2	16	×8	UL	—	—	—	—	—	—
\mathcal{C}_9	18-47	2666	16	2	16	×8	UL	—	—	—	—	—	—
$\mathcal{C}_{10,11}$	19-04	2933	8	1	16	×8	UL	—	—	—	—	—	—
$\mathcal{C}_{12}^{\ddagger}$	15-01 [†]	2132	4	1	16	×8	UT	25	10-sided	190037	63904	126133	✓
$\mathcal{C}_{13}^{\ddagger}$	18-49	2132	4	1	16	×8	UT	3	9-sided	694	239	455	—

[†] The module does not report manufacturing date. Therefore, we report purchase date as an approximation.

[‡] Analyzed using the FPGA-based SoftMC.

[◊] The system runs with double refresh frequency in standard conditions. We configured the refresh interval to be 64 *ms* in the BIOS settings.

UL = Unlimited

UT = Untested

TRRespass Vulnerable Mobile Phones

TABLE III: LPDDR4(X) results. Mobile phones tested against *TRRespass* on ARMv8 sorted by production date. We found bit flip inducing RowHammer patterns on 5 out of 13mobile phones.

<i>Mobile Phone</i>	<i>Year</i>	<i>SoC</i>	<i>Memory (GB)</i>	<i>Found Patterns</i>
Google Pixel	2016	MSM8996	4 [†]	✓
Google Pixel 2	2017	MSM8998	4	—
Samsung G960F/DS	2018	Exynos 9810	4	—
Huawei P20 DS	2018	Kirin 970	4	—
Sony XZ3	2018	SDM845	4	—
HTC U12+	2018	SDM845	6	—
LG G7 ThinQ	2018	SDM845	4 [†]	✓
Google Pixel 3	2018	SDM845	4	✓
Google Pixel 4	2019	SM8150	6	—
OnePlus 7	2019	SM8150	8	✓
Samsung G970F/DS	2019	Exynos 9820	6	✓
Huawei P30 DS	2019	Kirin 980	6	—
Xiaomi Redmi Note 8 Pro	2019	Helio G90T	6	—

[†] LPDDR4 (not LPDDR4X)

TRRespass Based RowHammer Attack

TABLE IV: Time to exploit. Time to find the first exploitable template on two sample modules from each DRAM vendor.

<i>Module</i>	τ (ms)	<i>PTE</i> [81]	<i>RSA-2048</i> [79]	<i>sudo</i> [27]
\mathcal{A}_{14}	188.7	4.9s	6m 27s	—
\mathcal{A}_4	180.8	38.8s	39m 28s	—
\mathcal{B}_1	360.7	—	—	—
\mathcal{B}_2	331.2	—	—	—
\mathcal{C}_{12}	300.0	2.3s	74.6s	54m16s
\mathcal{C}_{13}	180.9	3h 15m	—	—

τ : Time to template a single row: time to fill the victim and aggressor rows + hammer time + time to scan the row.

TRRespass Key Results

- 13 out of 42 tested DDR4 DRAM modules are vulnerable
 - From all 3 major manufacturers
 - 3-, 9-, 10-, 14-, 19-sided hammer attacks needed
- 5 out of 13 mobile phones tested vulnerable
 - From 4 major manufacturers
 - With LPDDR4(X) DRAM chips
- These results are scratching the surface
 - TRRespass tool is not exhaustive
 - There is a lot of room for uncovering more vulnerable chips and phones

RowHammer is still
an open problem

Security by obscurity
is likely not a good solution

Detailed Lecture on TRRespass

- Computer Architecture, Fall 2020, Lecture 5a
 - RowHammer in 2020: TRRespass (ETH Zürich, Fall 2020)
 - https://www.youtube.com/watch?v=pwRw7QqK_qA&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=9

<https://www.youtube.com/onurmutlulectures>

Industry-Adopted Solutions Do Not Work

- 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](#)]
[[Lecture Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#)] (17 minutes)
[[Lecture Video](#)] (59 minutes)
[[Source Code](#)]
[[Web Article](#)]
Best paper award.
Pwnie Award 2020 for Most Innovative Research. [Pwnie Awards 2020](#)

TRRespass: Exploiting the Many Sides of Target Row Refresh

Pietro Frigo^{*†} Emanuele Vannacci^{*†} Hasan Hassan[§] Victor van der Veen[¶]
Onur Mutlu[§] Cristiano Giuffrida^{*} Herbert Bos^{*} Kaveh Razavi^{*}

How to Guarantee That a Chip is RowHammer-Free?

Hard to Guarantee RowHammer-Free Chips

- 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^{§†}, Minesh Patel[§], Lillian Tsai[‡],
Stefan Saroiu, Alec Wolman, and Onur Mutlu^{§†}
Microsoft Research, [§]ETH Zürich, [†]CMU, [‡]MIT

Uncovering TRR Almost Completely

Industry-Adopted Solutions Are Very Poor

- Hasan Hassan, Yahya Can Tugrul, Jeremie S. Kim, Victor van der Veen, Kaveh Razavi, and Onur Mutlu,
"Uncovering In-DRAM RowHammer Protection Mechanisms: A New Methodology, Custom RowHammer Patterns, and Implications"
Proceedings of the 54th International Symposium on Microarchitecture (MICRO), Virtual, October 2021.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Short Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (25 minutes)]
[[Lightning Talk Video](#) (100 seconds)]
[[arXiv version](#)]

Uncovering In-DRAM RowHammer Protection Mechanisms: A New Methodology, Custom RowHammer Patterns, and Implications

Hasan Hassan[†]

[†]ETH Zürich

Yahya Can Tuğrul^{†‡}

Kaveh Razavi[†]
[‡]TOBB University of Economics & Technology

Jeremie S. Kim[†]

Onur Mutlu[†]

Victor van der Veen^σ

^σQualcomm Technologies Inc.

U-TRR Summary & Key Results

Target Row Refresh (TRR):


a set of **obscure**, **undocumented**, and **proprietary** RowHammer mitigation techniques

We **cannot** easily study the *security properties* of TRR

Is TRR fully secure? How can we validate its security guarantees?

U-TRR

A new methodology that leverages *data retention failures* to uncover the inner workings of TRR and study its security




TRR **does not provide security** against RowHammer

U-TRR can **facilitate** the development of **new RowHammer attacks** and **more secure RowHammer protection** mechanisms

Overview of U-TRR

U-TRR: A new methodology to *uncover* the inner workings of TRR

Key idea: Use **data retention failures** as a side channel to **detect when a row is refreshed** by TRR




Analyzing TRR-Protected DDR4 Chips

FPGA Board


**DRAM
Module**

**Temperature
Controller**



* *SoftMC [Hassan+, HPCA'17] enhanced for DDR4*

U-TRR Analysis Summary



new RowHammer access patterns
that **circumvent TRR**



Key Takeaways


All 45 modules we test are vulnerable

99.9% of rows in a DRAM bank
experience **at least one RowHammer bit flip**

ECC is ineffective: up to **7 RowHammer bit flips**
in an 8-byte dataword

Module	Date (yy-ww)	Chip Density (Gbit)	Organization			HC_{first}^{\dagger}	Our Key TRR Observations and Results							
			Ranks	Banks	Pins		Version	Aggressor Detection	Aggressor Capacity	Per-Bank TRR	TRR-to-REF Ratio	Neighbors Refreshed	% Vulnerable DRAM Rows [†]	Max. Bit Flips per Row per Hammer [†]
A0	19-50	8	1	16	8	16K	A_{TRR1}	Counter-based	16	✓	1/9	4	73.3%	1.16
A1-5	19-36	8	1	8	16	13K-15K	A_{TRR1}	Counter-based	16	✓	1/9	4	99.2% - 99.4%	2.32 - 4.73
A6-7	19-45	8	1	8	16	13K-15K	A_{TRR1}	Counter-based	16	✓	1/9	4	99.3% - 99.4%	2.12 - 3.86
A8-9	20-07	8	1	16	8	12K-14K	A_{TRR1}	Counter-based	16	✓	1/9	4	74.6% - 75.0%	1.96 - 2.96
A10-12	19-51	8	1	16	8	12K-13K	A_{TRR1}	Counter-based	16	✓	1/9	4	74.6% - 75.0%	1.48 - 2.86
A13-14	20-31	8	1	8	16	11K-14K	A_{TRR2}	Counter-based	16	✓	1/9	2	94.3% - 98.6%	1.53 - 2.78
B0	18-22	4	1	16	8	44K	B_{TRR1}	Sampling-based	1	✗	1/4	2	99.9%	2.13
B1-4	20-17	4	1	16	8	159K-192K	B_{TRR1}	Sampling-based	1	✗	1/4	2	23.3% - 51.2%	0.06 - 0.11
B5-6	16-48	4	1	16	8	44K-50K	B_{TRR1}	Sampling-based	1	✗	1/4	2	99.9%	1.85 - 2.03
B7	19-06	8	2	16	8	20K	B_{TRR1}	Sampling-based	1	✗	1/4	2	99.9%	31.14
B8	18-03	4	1	16	8	43K	B_{TRR1}	Sampling-based	1	✗	1/4	2	99.9%	2.57
B9-12	19-48	8	1	16	8	42K-65K	B_{TRR2}	Sampling-based	1	✗	1/9	2	36.3% - 38.9%	16.83 - 24.26
B13-14	20-08	4	1	16	8	11K-14K	B_{TRR3}	Sampling-based	1	✓	1/2	4	99.9%	16.20 - 18.12
C0-3	16-48	4	1	16	x8	137K-194K	C_{TRR1}	Mix	Unknown	✓	1/17	2	1.0% - 23.2%	0.05 - 0.15
C4-6	17-12	8	1	16	x8	130K-150K	C_{TRR1}	Mix	Unknown	✓	1/17	2	7.8% - 12.0%	0.06 - 0.08
C7-8	20-31	8	1	8	x16	40K-44K	C_{TRR1}	Mix	Unknown	✓	1/17	2	39.8% - 41.8%	9.66 - 14.56
C9-11	20-31	8	1	8	x16	42K-53K	C_{TRR2}	Mix	Unknown	✓	1/9	2	99.7%	9.30 - 32.04
C12-14	20-46	16	1	8	x16	6K-7K	C_{TRR3}	Mix	Unknown	✓	1/8	2	99.9%	4.91 - 12.64


Effect on Individual Rows



All 45 modules we tested are vulnerable to our new RowHammer access patterns

Our RowHammer access patterns cause bit flips in more than 99.9% of the rows

Bypassing ECC with New RowHammer Patterns



Modules from all three vendors have many **8-byte data chunks** with **3 and more (up to 7) RowHammer bit flips**

Conventional DRAM ECC **cannot protect** against our **new RowHammer access patterns**

Many Observations & Results in the Paper

- More observations on the TRRs of the three vendors
- Detailed description of the crafted access patterns
- Hammers per aggressor row sensitivity analysis
- Observations and results for individual modules
- ...

Module	Date (yy-ww)	Chip Density (Gbit)	Organization			HC_{first}^{\dagger}	Our Key TRR Observations and Results							
			Ranks	Banks	Pins		Version	Aggressor Detection	Aggressor Capacity	Per-Bank TRR	TRR-to-REF Ratio	Neighbors Refreshed	% Vulnerable DRAM Rows [†]	Max. Bit Flips per Row per Hammer [†]
A0	19-50	8	1	16	8	16K	$ATRR_1$	Counter-based	16	✓	1/9	4	73.3%	1.16
A1-5	19-36	8	1	8	16	13K-15K	$ATRR_1$	Counter-based	16	✓	1/9	4	99.2% - 99.4%	2.32 - 4.73
A6-7	19-45	8	1	8	16	13K-15K	$ATRR_1$	Counter-based	16	✓	1/9	4	99.3% - 99.4%	2.12 - 3.86
A8-9	20-07	8	1	16	8	12K-14K	$ATRR_1$	Counter-based	16	✓	1/9	4	74.6% - 75.0%	1.96 - 2.96
A10-12	19-51	8	1	16	8	12K-13K	$ATRR_1$	Counter-based	16	✓	1/9	4	74.6% - 75.0%	1.48 - 2.86
A13-14	20-31	8	1	8	16	11K-14K	$ATRR_2$	Counter-based	16	✓	1/9	2	94.3% - 98.6%	1.53 - 2.78
B0	18-22	4	1	16	8	44K	$BTRR_1$	Sampling-based	1	✗	1/4	2	99.9%	2.13
B1-4	20-17	4	1	16	8	159K-192K	$BTRR_1$	Sampling-based	1	✗	1/4	2	23.3% - 51.2%	0.06 - 0.11
B5-6	16-48	4	1	16	8	44K-50K	$BTRR_1$	Sampling-based	1	✗	1/4	2	99.9%	1.85 - 2.03
B7	19-06	8	2	16	8	20K	$BTRR_1$	Sampling-based	1	✗	1/4	2	99.9%	31.14
B8	18-03	4	1	16	8	43K	$BTRR_1$	Sampling-based	1	✗	1/4	2	99.9%	2.57
B9-12	19-48	8	1	16	8	42K-65K	$BTRR_2$	Sampling-based	1	✗	1/9	2	36.3% - 38.9%	16.83 - 24.26
B13-14	20-08	4	1	16	8	11K-14K	$BTRR_3$	Sampling-based	1	✓	1/2	4	99.9%	16.20 - 18.12
C0-3	16-48	4	1	16	x8	137K-194K	$CTRR_1$	Mix	Unknown	✓	1/17	2	1.0% - 23.2%	0.05 - 0.15
C4-6	17-12	8	1	16	x8	130K-150K	$CTRR_1$	Mix	Unknown	✓	1/17	2	7.8% - 12.0%	0.06 - 0.08
C7-8	20-31	8	1	8	x16	40K-44K	$CTRR_1$	Mix	Unknown	✓	1/17	2	39.8% - 41.8%	9.66 - 14.56
C9-11	20-31	8	1	8	x16	42K-53K	$CTRR_2$	Mix	Unknown	✓	1/9	2	99.7%	9.30 - 32.04
C12-14	20-46	16	1	8	x16	6K-7K	$CTRR_3$	Mix	Unknown	✓	1/8	2	99.9%	4.91 - 12.64

Uncovering TRR Can Help Future Solutions

- Hasan Hassan, Yahya Can Tugrul, Jeremie S. Kim, Victor van der Veen, Kaveh Razavi, and Onur Mutlu,
"Uncovering In-DRAM RowHammer Protection Mechanisms: A New Methodology, Custom RowHammer Patterns, and Implications"
Proceedings of the 54th International Symposium on Microarchitecture (MICRO), Virtual, October 2021.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Short Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (25 minutes)]
[[Lightning Talk Video](#) (100 seconds)]
[[arXiv version](#)]

Uncovering In-DRAM RowHammer Protection Mechanisms: A New Methodology, Custom RowHammer Patterns, and Implications

Hasan Hassan[†]

Yahya Can Tuğrul^{†‡}

Jeremie S. Kim[†]

Victor van der Veen^σ

Kaveh Razavi[†]

Onur Mutlu[†]

[†]ETH Zürich

[‡]TOBB University of Economics & Technology

^σQualcomm Technologies Inc.

New RowHammer Characteristics

RowHammer Has Many Dimensions

- Lois Orosa, Abdullah Giray Yaglikci, Haocong Luo, Ataberk Olgun, Jisung Park, Hasan Hassan, Minesh Patel, Jeremie S. Kim, and Onur Mutlu,
"A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses"
*Proceedings of the 54th International Symposium on Microarchitecture (**MICRO**), Virtual, October 2021.*
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Short Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (21 minutes)]
[[Lightning Talk Video](#) (1.5 minutes)]
[[arXiv version](#)]

A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses

Lois Orosa*
ETH Zürich

A. Giray Yağlıkçı*
ETH Zürich

Haocong Luo
ETH Zürich

Ataberk Olgun
ETH Zürich, TOBB ETÜ

Jisung Park
ETH Zürich

Hasan Hassan
ETH Zürich


Minesh Patel
ETH Zürich

Jeremie S. Kim
ETH Zürich

Onur Mutlu
ETH Zürich

Our Goal


Provide insights into **three fundamental properties**



Temperature



Aggressor Row
Active Time



Victim DRAM Cell's
Physical Location


To find **effective and efficient** attacks and defenses

DRAM Testing Infrastructures

Two separate testing infrastructures

1. **DDR3:** FPGA-based SoftMC (Xilinx ML605)

2. **DDR4:** FPGA-based SoftMC (Xilinx Virtex UltraScale+ XCU200)



Fine-grained control over **DRAM commands**,
timing parameters and **temperature ($\pm 0.1^{\circ}\text{C}$)**

DRAM Chips Tested

Two DRAM standards

Mfr.	DDR4 DIMMs	DDR3 SODIMMs	# Chips	Density	Die	Org.
A (Micron)	9	1	144 (8)	8Gb (4Gb)	B (P)	x4 (x8)
B (Samsung)	4	1	32 (8)	4Gb (4Gb)	F (Q)	x8 (x8)
C (SK Hynix)	5	1	40 (8)	4Gb (4Gb)	B (B)	x8 (x8)
D (Nanya)	4	-	32 (-)	8Gb (-)	C (-)	x8 (-)

4 Major Manufacturers


272 DRAM Chips in total

Summary of The Study & Key Results

- **272 DRAM chips** from **four major manufacturers**
- **6 major takeaways** from **16 novel observations**
- A RowHammer bit flip is **more likely to occur**
 - 1) in a **bounded range of temperature**
 - 2) if the aggressor row is **active for longer time**
 - 3) in **certain physical regions** of the DRAM module under attack
- Our novel observations can inspire and aid future work
 - Craft **more effective attacks**
 - Design **more effective and efficient defenses**

Example Attack Improvement 3: Bypassing Defenses with Aggressor Row Active Time

Activating aggressor rows as frequently as possible:



Keeping aggressor rows active for a longer time:



Reduces the minimum activation count to induce a bit flip **by 36%**

Bypasses defenses that do not account for this reduction

Key Takeaways from Spatial Variation Analysis

Key Takeaway 5


RowHammer vulnerability **significantly varies** across DRAM rows and columns due to **design-induced** and **manufacturing-process-induced** variation

Key Takeaway 6

The distribution of **the minimum activation count to observe bit flips (HC_{first})** exhibits **a diverse set of values in a subarray** but **similar values across subarrays** in the same DRAM module

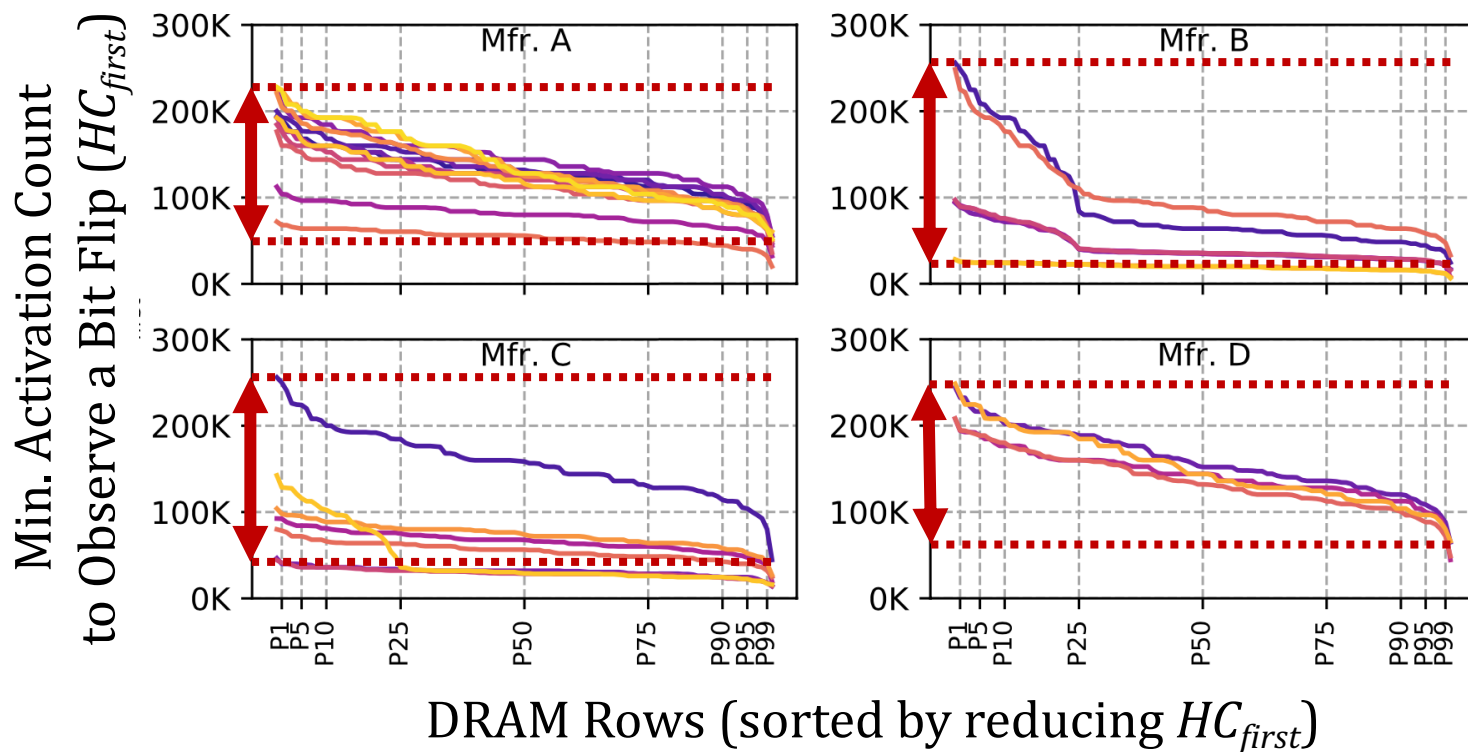
Spatial Variation across Rows

The **minimum activation count** to observe bit flips (HC_{first}) across **DRAM rows**:




The RowHammer vulnerability
significantly varies across DRAM rows

Spatial Variation across Rows



The RowHammer vulnerability
significantly varies across DRAM rows


Spatial Variation across Rows



OBSERVATION 12

A small fraction of DRAM rows are **significantly more vulnerable** to RowHammer than **the vast majority** of the rows

Spatial Variation across Columns




OBSERVATION 13

Certain columns are **significantly more vulnerable** to RowHammer than other columns

Example Defense Improvements

- **Example 1: Leveraging variation across DRAM rows**



- **Example 2: Leveraging variation with temperature**

- A DRAM cell experiences **bit flips** within a **bounded temperature range**



- A row can be **disabled** within the row's **vulnerable temperature range**



Many More Analyses In The Paper

- Lois Orosa, Abdullah Giray Yaglikci, Haocong Luo, Ataberk Olgun, Jisung Park, Hasan Hassan, Minesh Patel, Jeremie S. Kim, and Onur Mutlu,
"A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses"
Proceedings of the 54th International Symposium on Microarchitecture (MICRO), Virtual, October 2021.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Short Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (21 minutes)]
[[Lightning Talk Video](#) (1.5 minutes)]
[[arXiv version](#)]

A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses

Lois Orosa*
ETH Zürich

A. Giray Yağlıkçı*
ETH Zürich

Haocong Luo
ETH Zürich

Ataberk Olgun
ETH Zürich, TOBB ETÜ

Jisung Park
ETH Zürich

Hasan Hassan
ETH Zürich

Minesh Patel
ETH Zürich

Jeremie S. Kim
ETH Zürich

Onur Mutlu
ETH Zürich

More RowHammer Analysis

RowHammer vs. Wordline Voltage (2022)


- A. Giray Yağlıkçı, Haocong Luo, Geraldo F. de Oliveira, Ataberk Olgun, Minesh Patel, Jisung Park, Hasan Hassan, Jeremie S. Kim, Lois Orosa, and Onur Mutlu, **"Understanding RowHammer Under Reduced Wordline Voltage: An Experimental Study Using Real DRAM Devices"**
Proceedings of the 52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Baltimore, MD, USA, June 2022.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]
[[arXiv version](#)]
[[Talk Video](#) (34 minutes, including Q&A)]
[[Lightning Talk Video](#) (2 minutes)]

Understanding RowHammer Under Reduced Wordline Voltage: An Experimental Study Using Real DRAM Devices

A. Giray Yağlıkçı¹ Haocong Luo¹ Geraldo F. de Oliveira¹ Ataberk Olgun¹ Minesh Patel¹
Jisung Park¹ Hasan Hassan¹ Jeremie S. Kim¹ Lois Orosa^{1,2} Onur Mutlu¹
¹*ETH Zürich* ²*Galicia Supercomputing Center (CESGA)*

Updated DRAM Testing Infrastructure

FPGA-based SoftMC (Xilinx Virtex UltraScale+ XCU200)



Fine-grained control over DRAM commands,
timing parameters ($\pm 1.5\text{ns}$), temperature ($\pm 0.1^\circ\text{C}$),
and wordline voltage ($\pm 1\text{mV}$)

Summary

We provide *the first* RowHammer characterization **under reduced wordline voltage**

Experimental results with 272 *real DRAM chips* show that **reducing wordline voltage**:

1. Reduces RowHammer vulnerability

- **Bit error rate** caused by a RowHammer attack reduces by **15.2% (66.9% max)**
- A row needs to be activated **7.4% more times (85.8% max)** to induce *the first* bit flip

2. Increases row activation latency

- More than **76%** of the tested DRAM chips **reliably operate** using **nominal** timing parameters
- Remaining **24%** **reliably operate** with **increased** (up to 24ns) row activation latency

3. Reduces data retention time

- **80%** of the tested DRAM chips **reliably operate using nominal refresh rate**
- Remaining **20%** **reliably operate** by
 - Using **single error correcting codes**
 - **Doubling the refresh rate** for **a small fraction (16.4%) of DRAM rows**

Reducing wordline voltage can **reduce RowHammer vulnerability**
without significantly affecting **reliable DRAM operation**

New RowHammer Solutions

BlockHammer Solution in 2021

- A. Giray Yaglikci, Minesh Patel, Jeremie S. Kim, Roknoddin Azizi, Ataberk Olgun, Lois Orosa, Hasan Hassan, Jisung Park, Konstantinos Kanellopoulos, Taha Shahroodi, Saugata Ghose, and Onur Mutlu,

"BlockHammer: Preventing RowHammer at Low Cost by Blacklisting Rapidly-Accessed DRAM Rows"

Proceedings of the 27th International Symposium on High-Performance Computer Architecture (HPCA), Virtual, February-March 2021.

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

[[Short Talk Slides \(pptx\)](#) ([pdf](#))]

[[Intel Hardware Security Academic Awards Short Talk Slides \(pptx\)](#) ([pdf](#))]

[[Talk Video](#) (22 minutes)]

[[Short Talk Video](#) (7 minutes)]

[[Intel Hardware Security Academic Awards Short Talk Video](#) (2 minutes)]

[[BlockHammer Source Code](#)]

Intel Hardware Security Academic Award Finalist (one of 4 finalists out of 34 nominations)

BlockHammer: Preventing RowHammer at Low Cost by Blacklisting Rapidly-Accessed DRAM Rows

A. Giray Yağlıkçı¹ Minesh Patel¹ Jeremie S. Kim¹ Roknoddin Azizi¹ Ataberk Olgun¹ Lois Orosa¹
Hasan Hassan¹ Jisung Park¹ Konstantinos Kanellopoulos¹ Taha Shahroodi¹ Saugata Ghose² Onur Mutlu¹

¹ETH Zürich


²University of Illinois at Urbana-Champaign

RowHammer Solution Approaches

- More robust DRAM chips **and/or** error-correcting codes
- Increased refresh rate




- Physical isolation




Cost, Power, Performance, Complexity

- Reactive refresh



- Proactive throttling

SAFARI



Fewer activations allowed for aggressive applications

Two Key Challenges

1

Scalability

with worsening RowHammer vulnerability

2

Compatibility

with commodity DRAM chips

Our Goal

To prevent RowHammer **efficiently and scalably**
without knowledge of or modifications to DRAM internals

BlockHammer



Key Idea

Selectively throttle memory accesses
that may cause **RowHammer** bit-flips

BlockHammer: Practical Throttling-based Mechanism



- A RowHammer attack hammers Row A
- **BlockHammer** detects a RowHammer attack using **area-efficient Bloom filters**
- **BlockHammer** **selectively throttles** accesses from within **the memory controller**
- Bit flips **do not** occur
- BlockHammer can *optionally* **inform the system software** about the attack



Physical
Row Layout

BlockHammer is **compatible with commodity DRAM chips**
No need for **proprietary info** of or **modifications** to DRAM chips

BlockHammer

Overview of Approach

RowBlocker

Tracks row activation rates using area-efficient Bloom filters

Blacklists rows that are activated at a high rate

Throttles activations targeting a blacklisted row

No row can be activated at a high enough rate to induce bit-flips

AttackThrottler

Identifies threads that perform a RowHammer attack


Reduces memory bandwidth usage of identified threads

Greatly reduces the performance degradation and energy wastage a RowHammer attack inflicts on a system


Evaluation: BlockHammer

Scaling with RowHammer Vulnerability

- System throughput (weighted speedup)
- Job turnaround time (harmonic speedup)
- Unfairness (maximum slowdown)
- DRAM energy consumption



BlockHammer's performance and energy overheads remain **negligible** (<0.6%)



BlockHammer scalably provides **much higher performance** (71% on average) and **lower energy consumption** (32% on average) than state-of-the-art mechanisms

Key Results: BlockHammer

- **Competitive** with state-of-the-art mechanisms **when there is no attack**
- **Superior** performance and DRAM energy **when RowHammer attack present**
- **Better hardware area scaling with RowHammer vulnerability**
- **Security Proof**
- Addresses **Many-Sided Attacks**
- Evaluation of **14 mechanisms** across four desirable properties
 - Comprehensive Protection
 - Compatibility with Commodity DRAM Chips
 - Scalability with RowHammer Vulnerability
 - Deterministic Protection

BlockHammer is the only solution that satisfies all four desirable properties

Approach	Mechanism	Comprehensive Protection	Compatible w/ Commodity DRAM Chips	Scaling with RowHammer Vulnerability	Deterministic Protection
Increased Refresh Rate [2, 73]		✓	✓	✗	✓
Physical Isolation	CATT [14]	✗	✗	✗	✓
	GuardION [148]	✗	✗	✗	✓
	ZebRAM [78]	✗	✗	✗	✓
Reactive Refresh	ANVIL [5]	✗	✗	✗	✓
	PARA [73]	✓	✗	✗	✗
	PRoHIT [137]	✓	✗	✗	✗
	MRLoc [161]	✓	✗	✗	✗
	CBT [132]	✓	✗	✗	✓
	TWiCe [84]	✓	✗	✗	✓
	Graphene [113]	✓	✗	✓	✓
Proactive Throttling	Naive Thrott. [102]	✓	✓	✗	✓
	Thrott. Supp. [40]	✓	✗	✗	✓
	BlockHammer	✓	✓	✓	✓

More in the Paper: BlockHammer



[Full Paper](#)

- Using **area-efficient Bloom filters** for RowHammer detection
- Security Proof
 - Mathematically represent **all possible** access patterns
 - **No row can be activated high-enough times** to induce bit-flips
- BlockHammer prevents **many-sided attacks**
 - TRRespass [Frigo+, S&P'20]
 - U-TRR [Hassan+, MICRO'21]
 - BlackSmith [Jattke+, S&P'22]
 - Half-Double [Kogler+, USENIX Security'22]
- System Integration
 - **BlockHammer** can detect **RowHammer attacks** with **high accuracy** and **inform system software**
 - Measures **RowHammer likelihood of each thread**
- **Hardware complexity** analysis

Summary: BlockHammer

- BlockHammer is **the first work to practically enable throttling-based RowHammer mitigation**
- BlockHammer is implemented in **the memory controller** (*no proprietary information of / no modifications to DRAM chips*)
- BlockHammer is *both* **scalable with worsening RowHammer** and **compatible with commodity DRAM chips**
- BlockHammer is **open-source** along with **six state-of-the-art mechanisms**: <https://github.com/CMU-SAFARI/BlockHammer>



Main Memory Needs
Intelligent Controllers
for Security, Safety,
Reliability, Scaling

More RowHammer in 2020-2022

RowHammer in 2020 (I)

MICRO 2020

Submit Work ▾

Program ▾

Attend

Session 1A: Security & Privacy I

5:00 PM CEST – 5:15 PM CEST

Graphene: Strong yet Lightweight Row Hammer Protection

Yeonhong Park, Woosuk Kwon, Eojin Lee, Tae Jun Ham, Jung Ho Ahn, Jae W. Lee (Seoul National University)

5:15 PM CEST – 5:30 PM CEST

Persist Level Parallelism: Streamlining Integrity Tree Updates for Secure Persistent Memory

Alexander Freij, Shougang Yuan, Huiyang Zhou (NC State University); Yan Solihin (University of Central Florida)

5:30 PM CEST – 5:45 PM CEST

PThammer: Cross-User-Kernel-Boundary Rowhammer through Implicit Accesses

Zhi Zhang (University of New South Wales and Data61, CSIRO, Australia); Yueqiang Cheng (Baidu Security); Dongxi Liu, Surya Nepal (Data61, CSIRO, Australia); Zhi Wang (Florida State University); Yuval Yarom (University of Adelaide and Data61, CSIRO, Australia)

RowHammer in 2020 (II)

S & P



Home

Program ▼

Call For... ▼

Attend ▼

Workshops ▼

Session #5: Rowhammer

Room 2

Session chair: Michael Franz (UC Irvine)

RAMBleed: Reading Bits in Memory Without Accessing Them

Andrew Kwong (University of Michigan), Daniel Genkin (University of Michigan), Daniel Gruss (Data61)

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

Lucian Cojocar (Microsoft Research), Jeremie Kim (ETH Zurich, CMU), Minesh Patel (ETH Zurich, Microsoft Research), Onur Mutlu (ETH Zurich, CMU)

Leveraging EM Side-Channel Information to Detect Rowhammer Attacks

Zhenkai Zhang (Texas Tech University), Zihao Zhan (Vanderbilt University), Daniel Balasubramanian (Vanderbilt University), Peter Volgyesi (Vanderbilt University), Xenofon Koutsoukos (Vanderbilt University)

TRRespass: Exploiting the Many Sides of Target Row Refresh

Pietro Frigo (Vrije Universiteit Amsterdam, The Netherlands), Emanuele Vannacci (Vrije Universiteit Amsterdam, The Netherlands), Onur Mutlu (ETH Zürich), Cristiano Giuffrida (Vrije Universiteit Amsterdam, The Netherlands), Kaveh Razavi (Vrije Universiteit Amsterdam, The Netherlands)

RowHammer in 2020 (III)

29TH USENIX
SECURITY SYMPOSIUM

ATTEND

PROGRAM

PARTICIPATE

SPONSORS

ABOUT

DeepHammer: Depleting the Intelligence of Deep Neural Networks through Targeted Chain of Bit Flips

Fan Yao, *University of Central Florida*; Adnan Siraj Rakin and Deliang Fan, *Arizona State University*

AVAILABLE MEDIA   

Show details ▶

RowHammer in 2021 (I)

HotOS XVIII

The 18th Workshop on Hot Topics in Operating Systems

31-May 1 June–3 June 2021, Cyberspace, People's Couches, and Zoom

Stop! Hammer Time: Rethinking Our Approach to Rowhammer Mitigations

RowHammer in 2021 (II)

30TH USENIX
SECURITY SYMPOSIUM

[ATTEND](#)

[PROGRAM](#)

[PARTICIPATE](#)

[SPONSORS](#)

[ABOUT](#)

SMASH: Synchronized Many-sided Rowhammer Attacks from JavaScript

RowHammer in 2021 (III)



Session 10A: Security & Privacy III

Session Chair: Hoda Naghibijouybari (Binghamton)

9:00 PM CEST – 9:15 PM CEST

A Deeper Look into RowHammer's Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses

Lois Orosa, Abdullah Giray Yaglikci, Haocong Luo (ETH Zurich); Ataberk Olgun (TOBB University of Economics and Technology); Jisung Park, Hasan Hassan, Minesh Patel, Jeremie S. Kim, Onur Mutlu (ETH Zurich)

 [Paper](#)

9:15 PM CEST – 9:30 PM CEST

Uncovering In-DRAM RowHammer Protection Mechanisms: A New Methodology, Custom RowHammer Patterns, and Implications

Hasan Hassan (ETH Zurich); Yahya Can Tugrul (TOBB University of Economics and Technology); Jeremie S. Kim (ETH Zurich); Victor van der Veen (Qualcomm); Kaveh Razavi, Onur Mutlu (ETH Zurich)

 [Paper](#)

RowHammer in 2022 (I)

MAY 22-26, 2022 AT THE HYATT REGENCY, SAN FRANCISCO, CA

43rd IEEE Symposium on Security and Privacy

BLACKSMITH: Scalable Rowhammering in the Frequency Domain

**SpecHammer: Combining Spectre and Rowhammer
for New Speculative Attacks**

**PROTRR: Principled yet Optimal In-DRAM
Target Row Refresh**

RowHammer in 2022 (II)



Randomized Row-Swap: Mitigating Row Hammer by Breaking Spatial Correlation between Aggressor and Victim Rows

RowHammer in 2022 (III)


HPCA 2022

The 28th IEEE International Symposium on High-Performance Computer Architecture (HPCA-28), Seoul, South Korea

SafeGuard: Reducing the Security Risk from Row-Hammer via Low-Cost Integrity Protection

Mithril: Cooperative Row Hammer Protection on Commodity DRAM Leveraging Managed Refresh

RowHammer in 2022 (IV)



IRPS 2022

The Price of Secrecy: How Hiding Internal DRAM Topologies Hurts Rowhammer Defenses

Stefan Saroiu, Alec Wolman, Lucian Cojocar
Microsoft

RowHammer in 2022 (V)



Half-Double: Hammering From the Next Row Over

Andreas Kogler¹ Jonas Juffinger^{1,2} Salman Qazi³ Yoongu Kim³ Moritz Lipp^{4*}
Nicolas Boichat³ Eric Shiu⁵ Mattias Nissler³ Daniel Gruss¹

¹*Graz University of Technology* ²*Lamarr Security Research* ³*Google*
⁴*Amazon Web Services* ⁵*Rivos*

RowHammer in 2022 (VI)



HAMMERSCOPE: Observing DRAM Power Consumption Using Rowhammer

**When Frodo Flips:
End-to-End Key Recovery on FrodoKEM via Rowhammer**

RowHammer in 2022 (VII)



AQUA: Scalable Rowhammer Mitigation by Quarantining Aggressor Rows at Runtime

Anish Saxena, Gururaj Saileshwar (Georgia Institute of Technology); Prashant J. Nair (University of British Columbia); Moinuddin Qureshi (Georgia Institute of Technology)

HiRA: Hidden Row Activation for Reducing Refresh Latency of Off-the-Shelf DRAM Chips

Abdullah Giray Yaglikci (ETH Zürich); Ataberk Olgun (TOBB University of Economics and Technology); Lois Orosa, Minesh Patel, Haocong Luo, Hasan Hassan (ETH Zürich); Oguz Ergin (TOBB University of Economics and Technology); Onur Mutlu (ETH Zürich)

RowHammer in 2022 (VII)

- **To appear at MICRO 2022**

HiRA: Hidden Row Activation for Reducing Refresh Latency of Off-the-Shelf DRAM Chips

A. Giray Yağlıkçı¹ Ataberk Olgun^{1,2} Minesh Patel¹ Haocong Luo¹ Hasan Hassan¹
Lois Orosa^{1,3} Oğuz Ergin² Onur Mutlu¹

¹*ETH Zürich*

²*TOBB University of Economics and Technology*

³*Galicia Supercomputing Center (CESGA)*

RowHammer in 2022 (VIII)

A Case for Transparent Reliability in DRAM Systems

Minesh Patel[†] Taha Shahroodi^{‡‡} Aditya Manglik[†] A. Giray Yağlıkçı[†]
Ataberk Olgun[†] Haocong Luo[†] Onur Mutlu[†]

[†]*ETH Zürich* [‡]*TU Delft*

<https://arxiv.org/pdf/2204.10378.pdf>

RowHammer in 2022 (IX)

A Case for Self-Managing DRAM Chips: Improving Performance, Efficiency, Reliability, and Security via Autonomous in-DRAM Maintenance Operations

Hasan Hassan

Ataberk Olgun

A. Giray Yağlıkçı

Haocong Luo

Onur Mutlu

ETH Zürich

<https://arxiv.org/pdf/2207.13358.pdf>


Self-Managing DRAM: Overview

Self-Managing DRAM (SMD)

enables autonomous in-DRAM maintenance operations

Key Idea:

Prevent the memory controller from accessing DRAM regions that are *under maintenance* by **rejecting** row activation (**ACT**) commands



Leveraging the ability to *reject an ACT*, a **maintenance operation** can be implemented **completely within a DRAM chip**

SMD-Based Maintenance Mechanisms

DRAM Refresh

Fixed Rate (SMD-FR)

uniformly refreshes **all** DRAM rows with a **fixed** refresh period

Variable Rate (SMD-VR)

skips refreshing rows that can **retain their data for longer** than the default refresh period

RowHammer Protection

Probabilistic (SMD-PRP)

Performs **neighbor** row refresh with a **small probability** on every row activation

Deterministic (SMD-DRP)

keeps track of most **frequently activated** rows and performs **neighbor** row refresh when activation count threshold is exceeded

Memory Scrubbing

Periodic Scrubbing (SMD-MS)

periodically **scans** the **entire** DRAM for errors and corrects them

Self-Managing DRAM: Summary

The three major DRAM maintenance operations:

- ❖ Refresh
- ❖ RowHammer Protection
- ❖ Memory Scrubbing

Implementing new **maintenance mechanisms** often requires **difficult-to-realize changes**

Our Goal

- ① Ease the process of enabling new DRAM maintenance operations
- ② Enable more efficient in-DRAM maintenance operations

Self-Managing DRAM (SMD)


Enables implementing new **in-DRAM** maintenance mechanisms with **no further changes** in the *DRAM interface* and *memory controller*

SMD-based *refresh*, *RowHammer protection*, and *scrubbing* achieve **9.2% speedup** and **6.2% lower DRAM energy** vs. conventional DRAM

Talk on Self-Managing DRAM

Problem: The Rigid DRAM Interface

The **Memory Controller** manages DRAM maintenance operations



Changes to maintenance operations are often reflected to the memory controller design, DRAM interface, and other system components



Implementing new maintenance operations
(or modifying the existing ones) is difficult-to-realize



SAFARI Live Seminars 2022

SAFARI Live Seminar - Improving DRAM Performance, Reliability, and Security by Understanding DRAM

1,039 views • Streamed live on Sep 15, 2022

37 DISLIKE SHARE DOWNLOAD CLIP SAVE ...



Onur Mutlu Lectures
27.6K subscribers

ANALYTICS

EDIT VIDEO

Much More in Our Preprint...

A Case for Self-Managing DRAM Chips: Improving Performance, Efficiency, Reliability, and Security via Autonomous in-DRAM Maintenance Operations

Hasan Hassan

Ataberk Olgun

A. Giray Yağlıkçı

Haocong Luo

Onur Mutlu

ETH Zürich

<https://arxiv.org/pdf/2207.13358.pdf>

RowHammer in 2023

MAY 22-26, 2023 AT THE HYATT REGENCY, SAN FRANCISCO, CA

44th IEEE Symposium on Security and Privacy

CSI:Rowhammer – Cryptographic Security and Integrity against Rowhammer

Jonas Juffinger^{*†}, Lukas Lamster[†], Andreas Kogler[†], Maria Eichlseder[†], Moritz Lipp[‡], Daniel Gruss^{*†}

^{*}Lamarr Security Research, [†]Graz University of Technology, [‡]Amazon Web Services

More to Come...

Future Memory Reliability/Security Challenges

Future of Main Memory Security

- 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

Future of Main Memory Security

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

Main Memory Needs
Intelligent Controllers
for Security, Safety,
Reliability, Scaling

Intelligent
Memory Controllers
Can Avoid Many Failures
& Enable Better Scaling

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

A Case for Transparent Reliability in DRAM Systems

Minesh Patel[†] Taha Shahroodi^{‡†} Aditya Manglik[†] A. Giray Yağlıkçı[†]
Ataberk Olgun[†] Haocong Luo[†] Onur Mutlu[†]

[†]*ETH Zürich* [‡]*TU Delft*

<https://arxiv.org/pdf/2204.10378.pdf>

Better Coordination of DRAM & Controller

A Case for Self-Managing DRAM Chips: Improving Performance, Efficiency, Reliability, and Security via Autonomous in-DRAM Maintenance Operations

Hasan Hassan

Ataberk Olgun

A. Giray Yağlıkçı


Haocong Luo

Onur Mutlu


ETH Zürich

<https://arxiv.org/pdf/2207.13358.pdf>

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



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

Collapse of the “Galloping Gertie” (1940)



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)



A Most Recent Example (2022)



A Most Recent Example (2022)



A Most Recent Example (2022)



A Most Recent Example (2022)



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

An Early Proposal for Intelligent Controllers [IMW'13]

- Onur Mutlu,
"Memory Scaling: A Systems Architecture Perspective"
Proceedings of the 5th International Memory Workshop (IMW), Monterey, CA, May 2013. Slides
(pptx) (pdf)
EETimes Reprint

Memory Scaling: A Systems Architecture Perspective

Onur Mutlu
Carnegie Mellon University
onur@cmu.edu
<http://users.ece.cmu.edu/~omutlu/>

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*



Final Thoughts on RowHammer

Before RowHammer (I)

Using Memory Errors to Attack a Virtual Machine

Sudhakar Govindavajhala * Andrew W. Appel
Princeton University
{sudhakar,appel}@cs.princeton.edu

We present an experimental study showing that soft memory errors can lead to serious security vulnerabilities in Java and .NET virtual machines, or in any system that relies on type-checking of untrusted programs as a protection mechanism. Our attack works by sending to the JVM a Java program that is designed so that almost any memory error in its address space will allow it to take control of the JVM. All conventional Java and .NET virtual machines are vulnerable to this attack. The technique of the attack is broadly applicable against other language-based security schemes such as proof-carrying code.

We measured the attack on two commercial Java Virtual Machines: Sun's and IBM's. We show that a single-bit error in the Java program's data space can be exploited to execute arbitrary code with a probability of about 70%, and multiple-bit errors with a lower probability.

Our attack is particularly relevant against smart cards or tamper-resistant computers, where the user has physical access (to the outside of the computer) and can use various means to induce faults; we have successfully used heat. Fortunately, there are some straightforward defenses against this attack.

7 Physical fault injection

If the attacker has physical access to the outside of the machine, as in the case of a smart card or other tamper-resistant computer, the attacker can induce memory errors. We considered attacks on boxes in form factors ranging from a credit card to a palmtop to a desktop PC.

We considered several ways in which the attacker could induce errors.⁴

IEEE S&P 2003

Before RowHammer (II)

Using Memory Errors to Attack a Virtual Machine

Sudhakar Govindavajhala * Andrew W. Appel
Princeton University
{sudhakar,appel}@cs.princeton.edu




Figure 3. Experimental setup to induce memory errors, showing a PC built from surplus components, clip-on gooseneck lamp, 50-watt spotlight bulb, and digital thermometer. Not shown is the variable AC power supply for the lamp.

IEEE S&P 2003

After RowHammer

A simple memory error
can be induced by software

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 & architecture 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 ASPLOS 2022)
 - ❑ 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...**

Conclusion

Summary: RowHammer

- Memory 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

A RowHammer Survey Across the Stack

- 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](#)]
[[Slides from COSADE 2019 \(pptx\)](#)]
[[Slides from VLSI-SOC 2020 \(pptx\) \(pdf\)](#)]
[[Talk Video](#) (1 hr 15 minutes, with Q&A)]

RowHammer: A Retrospective

Onur Mutlu^{§‡} Jeremie S. Kim^{‡§}
§ETH Zürich ‡Carnegie Mellon University

Detailed Lectures on RowHammer

- Computer Architecture, Fall 2021, Lecture 5

- RowHammer (ETH Zürich, Fall 2021)
- <https://www.youtube.com/watch?v=7wVKnPj3NVw&list=PL5Q2soXY2Zi-Mnk1PxjEIG32HAGILkTOF&index=5>

- Computer Architecture, Fall 2021, Lecture 6

- RowHammer and Secure & Reliable Memory (ETH Zürich, Fall 2021)
- <https://www.youtube.com/watch?v=HNd4skQrt6I&list=PL5Q2soXY2Zi-Mnk1PxjEIG32HAGILkTOF&index=6>

<https://www.youtube.com/onurmutlulectures>



Funding Acknowledgments

- Alibaba, AMD, [ASML](#), [Google](#), [Facebook](#), [Hi-Silicon](#), HP Labs, [Huawei](#), IBM, [Intel](#), [Microsoft](#), Nvidia, Oracle, Qualcomm, Rambus, Samsung, Seagate, [VMware](#), [Xilinx](#)
- NSF
- NIH
- GSRC
- [SRC](#)
- CyLab
- [EFCL](#)

Thank you!

Acknowledgments

SAFARI

SAFARI Research Group

safari.ethz.ch

Think BIG, Aim HIGH!

<https://safari.ethz.ch>

SAFARI Research Group

- <https://safari.ethz.ch/safari-newsletter-december-2021/>

SAFARI
SAFARI Research Group

Think Big, Aim High

ETH zürich



View in your browser
December 2021



Comp Arch (Fall'21)

Trace: readings - start - schedule

Home

Announcements

Materials

- Lectures/Schedule
- Lecture Buzzwords
- Readings
- HWs
- Labs
- Exams
- Related Courses
- Tutorials

Resources

- Computer Architecture FS20: Course Webpage
- Computer Architecture FS20: Lecture Videos
- Digitaltechnik SS21: Course Webpage
- Digitaltechnik SS21: Lecture Videos
- Moodle
- HotCRP
- Verilog Practice Website (HDLBits)

Fall 2021 Edition:

- <https://safari.ethz.ch/architecture/fall2021/doku.php?id=schedule>

Fall 2020 Edition:

- <https://safari.ethz.ch/architecture/fall2020/doku.php?id=schedule>

Youtube Livestream (2021):

- https://www.youtube.com/watch?v=4yfkM_5EFg0&list=PL5Q2soXY2Zi-Mnk1PxjEIG32HAGILkTOF

Youtube Livestream (2020):

- <https://www.youtube.com/watch?v=c3mPdZA-Fmc&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN>


Master's level course

- Taken by Bachelor's/Masters/PhD students
- Cutting-edge research topics + fundamentals in Computer Architecture
- 5 Simulator-based Lab Assignments
- Potential research exploration
- Many research readings

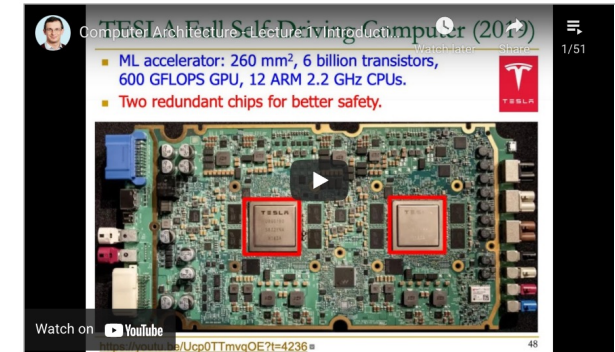
<https://www.youtube.com/onurmutlulectures>

Lecture Video Playlist on YouTube

Livestream Lecture Playlist



Recorded Lecture Playlist



Fall 2021 Lectures & Schedule

Week	Date	Livestream	Lecture	Readings	Lab	HW
W1	30.09 Thu.	YouTube Live	L1: Introduction and Basics PDF PPT	Required Mentioned	Lab 1 Out	HW 0 Out
	01.10 Fri.	YouTube Live	L2: Trends, Tradeoffs and Design Fundamentals PDF PPT	Required Mentioned		
W2	07.10 Thu.	YouTube Live	L3a: Memory Systems: Challenges and Opportunities PDF PPT	Described Suggested		HW 1 Out
			L3b: Course Info & Logistics PDF PPT			
			L3c: Memory Performance Attacks PDF PPT	Described Suggested		
	08.10 Fri.	YouTube Live	L4a: Memory Performance Attacks PDF PPT	Described Suggested	Lab 2 Out	
			L4b: Data Retention and Memory Refresh PDF PPT	Described Suggested		
			L4c: RowHammer PDF PPT	Described Suggested		

DDCA (Spring 2022)

Spring 2022 Edition:

- <https://safari.ethz.ch/digitaltechnik/spring2022/duku.php?id=schedule>

Spring 2021 Edition:

- <https://safari.ethz.ch/digitaltechnik/spring2021/duku.php?id=schedule>

Youtube Livestream (Spring 2022):

- <https://www.youtube.com/watch?v=cpXdE3HwvK0&list=PL5Q2soXY2Zi97Ya5DEUpMpO2bbAoaG7c6>

Youtube Livestream (Spring 2021):

- https://www.youtube.com/watch?v=LbC0EZY8yw4&list=PL5Q2soXY2Zi_uej3aY39YB5pfW4SJ7LIN

Bachelor's course

- 2nd semester at ETH Zurich
- Rigorous introduction into "How Computers Work"
- Digital Design/Logic
- Computer Architecture
- 10 FPGA Lab Assignments

SAFARI

<https://www.youtube.com/onurmutlulectures>



Digital Design and Computer Architecture -
Spring 2021

Search
Recent Changes Media Manager Site

Trace: - schedule

Home

Announcements

Materials

- Lectures/Schedule
- Lecture Buzzwords
- Readings
- Optional HWs
- Labs
- Extra Assignments
- Exams
- Technical Docs

Resources

- Computer Architecture (CMU) SS15: Lecture Videos
- Computer Architecture (CMU) SS15: Course Website
- Digitaltechnik SS18: Lecture Videos
- Digitaltechnik SS18: Course Website
- Digitaltechnik SS19: Lecture Videos
- Digitaltechnik SS19: Course Website
- Digitaltechnik SS20: Lecture Videos
- Digitaltechnik SS20: Course Website
- Moodle

Lecture Video Playlist on YouTube

Livestream Lecture Playlist

Recorded Lecture Playlist

Spring 2021 Lectures/Schedule

Week	Date	Livestream	Lecture	Readings	Lab	HW
W1	25.02 Thu.	YouTube Live	L1: Introduction and Basics Q2A (PDF) Q2B (PPT)	Required Suggested Mentioned		
	26.02 Fri.	YouTube Live	L2a: Tradeoffs, Metrics, Mindset Q2A (PDF) Q2B (PPT)	Required		
			L2b: Mysteries in Computer Architecture Q2A (PDF) Q2B (PPT)	Required Mentioned		
W2	04.03 Thu.	YouTube Live	L3a: Mysteries in Computer Architecture II Q2A (PDF) Q2B (PPT)	Required Suggested Mentioned		

The Story of RowHammer

Onur Mutlu

omutlu@gmail.com

<https://people.inf.ethz.ch/omutlu>

20 September 2022

University of British Columbia

SAFARI

ETH zürich

Carnegie Mellon

Backup Slides for Further Info

HiRA: Hidden Row Activation

for Reducing Refresh Latency of Off-the-Shelf DRAM Chips

Abdullah Giray Yağlıkçı

Ataberk Olgun Minesh Patel Haocong Luo Hasan Hassan

Lois Orosa Oğuz Ergin Onur Mutlu

SAFARI

ETH zürich

 **CESGA**
Centro de Supercomputación de Galicia

 **TOBB ETÜ**
University of Economics & Technology

Executive Summary

- **Problem:** *Periodic and preventive* refreshes cause **increasingly significant performance degradation** as DRAM chip density increases
- **Goal:** Reduce the **performance overhead** of *periodic and preventive* refreshes in off-the-shelf DRAM chips
- **Key Idea:** *Refresh* a DRAM row **concurrently with** *refreshing or activating* another row **in the same bank**, leveraging subarray-level parallelism
- **HiRA:** Hidden Row Activation
 - **Concurrently** opens two rows in *electrically isolated subarrays* in **quick succession**
 - *Refreshes* a DRAM row **concurrently with** *refreshing or activating* any of the **32% of the rows** in the same bank
 - **51.4% reduction** on the *overall latency* of refreshing two rows
- **HiRA-MC:** Buffers refresh requests for a time slack to leverage the parallelism HiRA provides
 - **12.6% speedup** by reducing *periodic* refresh's performance overheads
 - **3.73x speedup** by reducing *preventive* refresh's performance overheads

HiRA: Hidden Row Activation for Reducing Refresh Latency of Off-the-Shelf DRAM Chips

A. Giray Yağlıkçı¹ Ataberk Olgun¹ Minesh Patel¹ Haocong Luo¹ Hasan Hassan¹
Lois Orosa^{1,3} Oğuz Ergin² Onur Mutlu¹

¹ETH Zürich ²TOBB University of Economics and Technology ³Galicia Supercomputing Center (CESGA)

DRAM is the building block of modern main memory systems. DRAM cells must be periodically refreshed to prevent data loss. Refresh operations degrade system performance by interfering with memory accesses. As DRAM chip density increases with technology node scaling, refresh operations also increase because: 1) the number of DRAM rows in a chip increases; and 2) DRAM cells need additional refresh operations to mitigate bit failures caused by RowHammer, a failure mechanism that becomes worse with technology node scaling. Thus, it is critical to enable refresh operations at low performance overhead. To this end, we propose a new operation, Hidden Row Activation (HiRA), and the HiRA Memory Controller (HiRA-MC) to perform HiRA operations.

As DRAM density increases with technology node scaling, the performance overhead of refresh also increases due to three major reasons. First, as the DRAM chip density increases, more DRAM rows need to be periodically refreshed in a DRAM chip [55, 57–61]. Second, as DRAM technology node scales down, DRAM cells become smaller and thus can store less amount of charge, requiring them to be refreshed more frequently [10, 20, 67, 102, 103, 118, 122–124]. Third, with increasing DRAM density, DRAM cells are placed closer to each other, exacerbating charge leakage via a disturbance error mechanism called RowHammer [79, 84, 119, 120, 133, 134, 167, 180, 183], and thus requiring additional refresh operations (called *preventive* refreshes) to avoid data corruption due to RowHammer [2, 3, 5–7, 29, 33, 42, 63, 66, 76, 82, 84, 97, 98, 107, 135, 141].

HiRA: Hidden Row Activation

for Reducing Refresh Latency of Off-the-Shelf DRAM Chips

Abdullah Giray Yağlıkçı

Ataberk Olgun Minesh Patel Haocong Luo Hasan Hassan

Lois Orosa Oğuz Ergin Onur Mutlu

SAFARI

ETH zürich

 **CESGA**
Centro de Supercomputación de Galicia

 **TOBB ETÜ**
University of Economics & Technology

Understanding RowHammer

Root Causes of Disturbance Errors

- *Cause 1: Electromagnetic coupling*
 - Toggling the wordline voltage briefly increases the voltage of adjacent wordlines
 - Slightly opens adjacent rows → Charge leakage
- *Cause 2: Conductive bridges*
- *Cause 3: Hot-carrier injection*

Confirmed by at least one manufacturer

RowHammer Solutions

Naive Solutions

1 *Throttle accesses to same row*

- Limit access-interval: $\geq 500\text{ns}$
- Limit number of accesses: $\leq 128\text{K}$ ($=64\text{ms}/500\text{ns}$)

2 *Refresh more frequently*

- Shorten refresh-interval by $\sim 7\times$

Both naive solutions introduce significant overhead in performance and power

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*



Revisiting RowHammer

Revisiting RowHammer

An Experimental Analysis of Modern Devices and Mitigation Techniques

Jeremie S. Kim

Minesh Patel

A. Giray Yağlıkçı

Hasan Hassan

Roknoddin Azizi

Lois Orosa

Onur Mutlu

SAFARI

Key Conclusions


- We characterized **1580 DRAM** chips of different DRAM types, technology nodes, and manufacturers.
- We studied **five** state-of-the-art RowHammer mitigation mechanisms and an ideal refresh-based mechanism
- We made **two key observations**
 1. **RowHammer is getting much worse.** It takes much fewer hammers to induce RowHammer bit flips in newer chips
 - e.g., **DDR3**: 69.2k to 22.4k, **DDR4**: 17.5k to 10k, **LPDDR4**: 16.8k to 4.8k
 2. **Existing mitigation mechanisms do not scale** to DRAM chips that are more vulnerable to RowHammer
 - e.g., 80% performance loss when the hammer count to induce the first bit flip is 128
- We **conclude** that it is **critical** to do more research on RowHammer and develop scalable mitigation mechanisms to prevent RowHammer in future systems

DRAM Testing Infrastructures

Three separate testing infrastructures

1. **DDR3:** FPGA-based SoftMC [Hassan+, HPCA'17]
(Xilinx ML605)
2. **DDR4:** FPGA-based SoftMC [Hassan+, HPCA'17]
(Xilinx Virtex UltraScale 95)
3. **LPDDR4:** In-house testing hardware for LPDDR4 chips

All provide fine-grained control over DRAM commands, timing parameters and temperature



DRAM Chips Tested


DRAM type-node	Number of Chips (Modules) Tested			
	Mfr. A	Mfr. B	Mfr. C	Total
DDR3-old	56 (10)	88 (11)	28 (7)	172 (28)
DDR3-new	80 (10)	52 (9)	104 (13)	236 (32)
DDR4-old	112 (16)	24 (3)	128 (18)	264 (37)
DDR4-new	264 (43)	16 (2)	108 (28)	388 (73)
LPDDR4-1x	12 (3)	180 (45)	N/A	192 (48)
LPDDR4-1y	184 (46)	N/A	144 (36)	328 (82)

1580 total DRAM chips tested from **300** DRAM modules

- **Three** major DRAM manufacturers {A, B, C}
- **Three** DRAM *types or standards* {DDR3, DDR4, LPDDR4}
 - LPDDR4 chips we test implement on-die ECC
- **Two** technology nodes per DRAM type {old/new, 1x/1y}
 - Categorized based on manufacturing date, datasheet publication date, purchase date, and characterization results

Type-node: configuration describing a chip's type and technology
node generation: **DDR3-old/new, DDR4-old/new, LPDDR4-1x/1y**

3. Hammer Count (HC) Effects




RowHammer bit flip rates **increase**
when going **from old to new** DDR4 technology node generations

**RowHammer bit flip rates (i.e., RowHammer vulnerability)
increase with technology node generation**

5. First RowHammer Bit Flips per Chip


What is the minimum Hammer Count required to cause bit flips (HC_{first})?



We note the different DRAM types on the x-axis: **DDR3**, **DDR4**, **LPDDR4**.

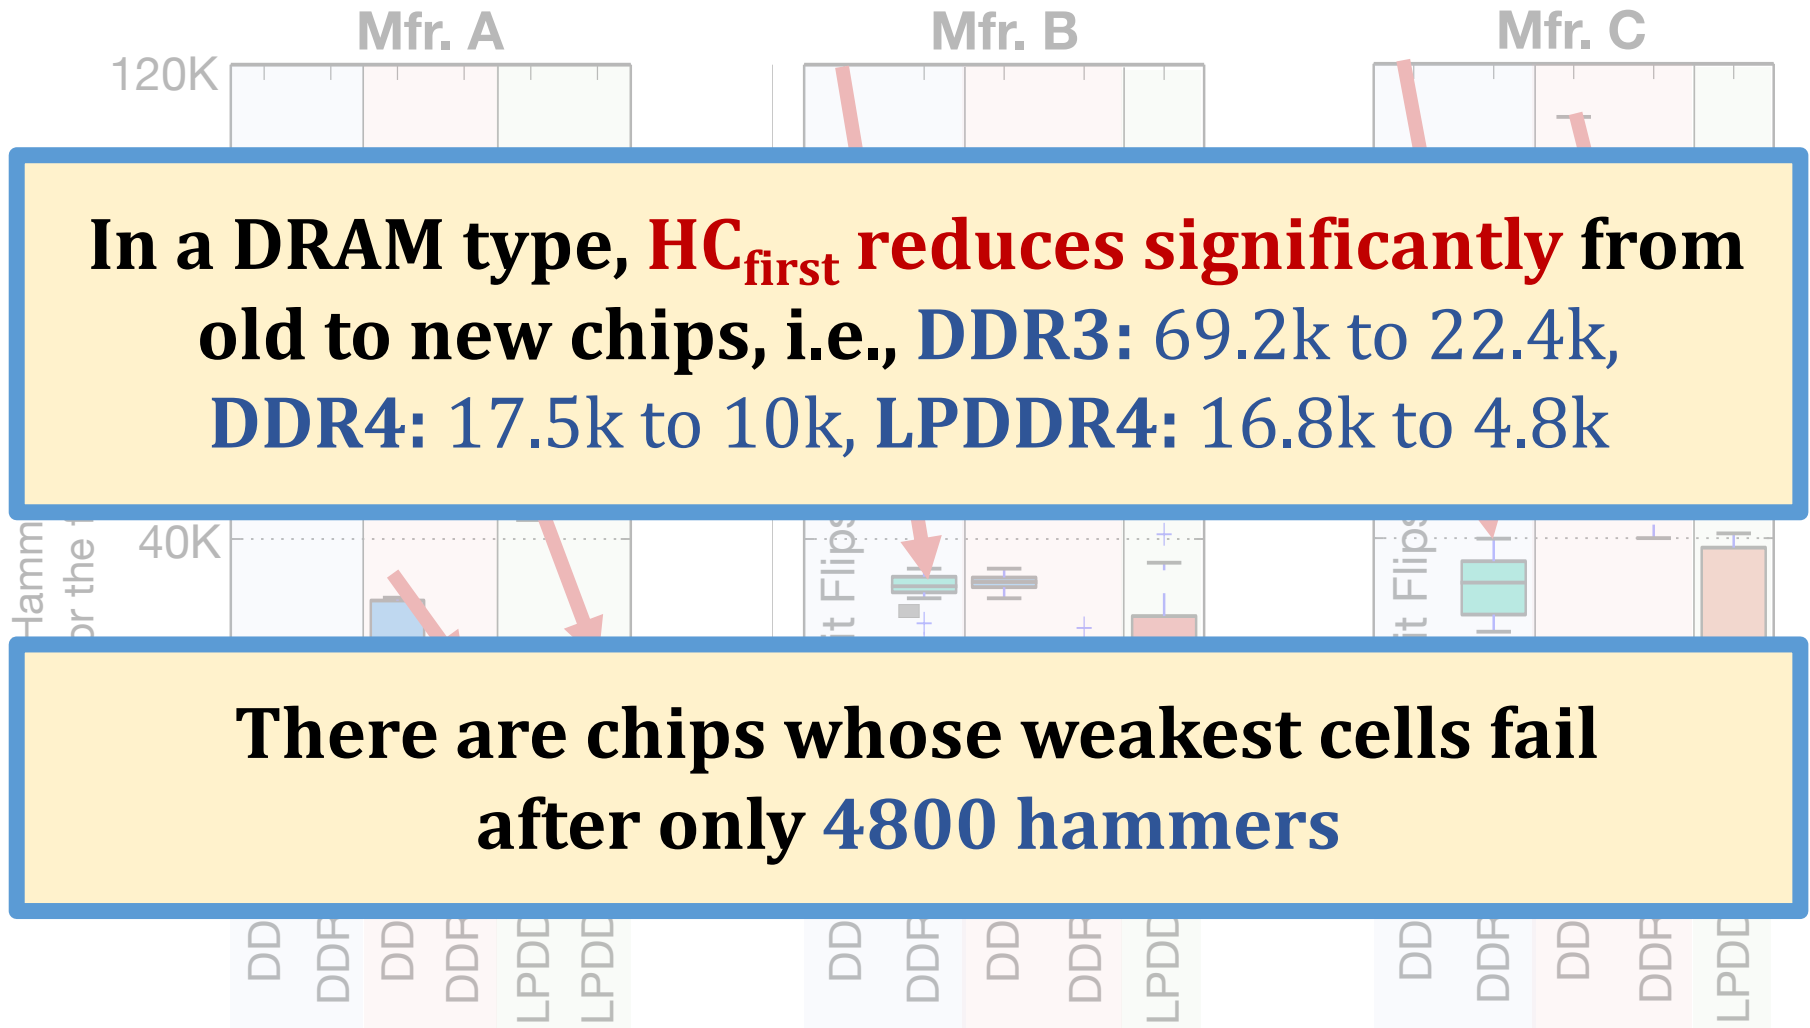
We focus on trends across chips of the same DRAM type to draw conclusions

5. First RowHammer Bit Flips per Chip



Newer chips from a given DRAM manufacturer
more vulnerable to RowHammer

5. First RowHammer Bit Flips per Chip




Newer chips from a given DRAM manufacturer
more vulnerable to RowHammer

Key Takeaways from 1580 Chips


- Chips of newer DRAM technology nodes are **more vulnerable** to RowHammer
- There are chips today whose weakest cells fail after **only 4800 hammers**
- Chips of newer DRAM technology nodes can exhibit RowHammer bit flips 1) in **more rows** and 2) **farther away** from the victim row.

Mitigation Mechanism Evaluation




PARA, ProHIT, and MRLoc mitigate RowHammer bit flips in **worst chips** today with reasonable system performance (92%, 100%, 100%)

Mitigation Mechanism Evaluation



Only PARA's design scales to low HC_{first} values
but has **very low normalized system performance**

Mitigation Mechanism Evaluation



Ideal mechanism is **significantly better** than any existing mechanism for $HC_{first} < 1024$

Significant opportunity for developing a RowHammer solution with **low performance overhead** that supports low HC_{first}

Key Takeaways from Mitigation Mechanisms

- Existing RowHammer mitigation mechanisms can prevent RowHammer attacks with **reasonable system performance overhead** in DRAM chips today
- Existing RowHammer mitigation mechanisms **do not scale well** to DRAM chips more vulnerable to RowHammer
- There is still **significant opportunity** for developing a mechanism that is **scalable with low overhead**

RowHammer Solutions Going Forward

Two promising directions for new RowHammer solutions:

1. DRAM-system cooperation

- We believe the DRAM and system should cooperate more to provide a **holistic** solution can prevent RowHammer at **low cost**

2. Profile-guided

- Accurate **profile of RowHammer-susceptible cells** in DRAM provides a powerful substrate for building **targeted** RowHammer solutions, e.g.:
 - Only increase the refresh rate for rows containing RowHammer-susceptible cells
- A **fast and accurate** profiling mechanism is a key research challenge for developing low-overhead and scalable RowHammer solutions

Revisiting RowHammer

An Experimental Analysis of Modern Devices and Mitigation Techniques

Jeremie S. Kim

Minesh Patel

A. Giray Yağlıkçı

Hasan Hassan

Roknoddin Azizi

Lois Orosa

Onur Mutlu

SAFARI

Detailed Lecture on Revisiting RowHammer

- Computer Architecture, Fall 2020, Lecture 5b

- RowHammer in 2020: Revisiting RowHammer (ETH Zürich, Fall 2020)
- <https://www.youtube.com/watch?v=gR7XR-Eepcg&list=PL5Q2soXY2Zi9xidyIgBxUz7xRPS-wisBN&index=10>

<https://www.youtube.com/onurmutlulectures>

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^{§†} Minesh Patel[§] A. Giray Yağlıkçı[§]
Hasan Hassan[§] Roknoddin Azizi[§] Lois Orosa[§] Onur Mutlu^{§†}
[§]*ETH Zürich* [†]*Carnegie Mellon University*

Executive Summary

- **Motivation**: Denser DRAM chips are more vulnerable to RowHammer but no characterization-based study demonstrates how vulnerability scales
- **Problem**: Unclear if existing mitigation mechanisms will remain viable for future DRAM chips that are likely to be more vulnerable to RowHammer
- **Goal**:
 1. Experimentally demonstrate how vulnerable modern DRAM chips are to RowHammer and study how this vulnerability will scale going forward
 2. Study viability of existing mitigation mechanisms on more vulnerable chips
- **Experimental Study**: First rigorous RowHammer characterization study across a broad range of DRAM chips
 - 1580 chips of different DRAM {types, technology node generations, manufacturers}
 - We find that RowHammer vulnerability worsens in newer chips
- **RowHammer Mitigation Mechanism Study**: How five state-of-the-art mechanisms are affected by worsening RowHammer vulnerability
 - Reasonable performance loss (8% on average) on modern DRAM chips
 - Scale poorly to more vulnerable DRAM chips (e.g., 80% performance loss)
- **Conclusion**: it is critical to research more effective solutions to RowHammer for future DRAM chips that will likely be even more vulnerable to RowHammer

Motivation

- Denser DRAM chips are **more vulnerable** to RowHammer
- Three prior works [Kim+, ISCA'14], [Park+, MR'16], [Park+, MR'16], **over the last six years** provide RowHammer characterization data on real DRAM
- However, there is **no comprehensive experimental study** that demonstrates **how vulnerability scales** across DRAM types and technology node generations
- It is **unclear whether current mitigation mechanisms will remain viable** for future DRAM chips that are likely to be more vulnerable to RowHammer

Goal

1. **Experimentally demonstrate** how vulnerable modern DRAM chips are to RowHammer and **predict how this vulnerability will scale** going forward
2. Examine the viability of current mitigation mechanisms on **more vulnerable chips**

Effective RowHammer Characterization

To characterize our DRAM chips at **worst-case** conditions, we:


1. Prevent sources of interference during core test loop

- We disable:
 - **DRAM refresh**: to avoid refreshing victim row
 - **DRAM calibration events**: to minimize variation in test timing
 - **RowHammer mitigation mechanisms**: to observe circuit-level effects
- Test for **less than refresh window (32ms)** to avoid retention failures

2. Worst-case access sequence

- We use **worst-case** access sequence based on prior works' observations
- For each row, **repeatedly access the two directly physically-adjacent rows as fast as possible**

Testing Methodology



DRAM_RowHammer_Characterization():

foreach row in DRAM:

set *victim_row* to row

set *aggressor_row1* to *victim_row* - 1

set *aggressor_row2* to *victim_row* + 1

Disable DRAM refresh

Refresh *victim_row*

for $n = 1 \rightarrow HC$: // core test loop

activate *aggressor_row1*

activate *aggressor_row2*

Enable DRAM refresh

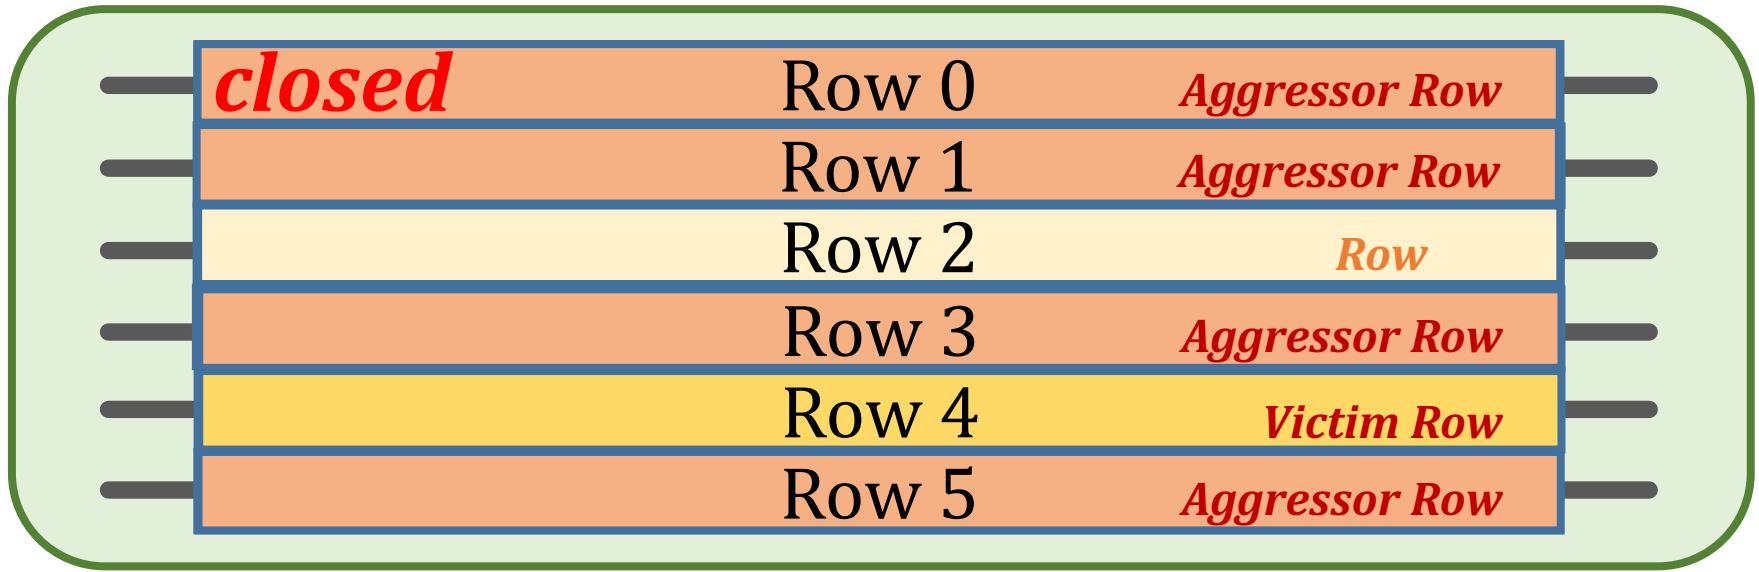
Record RowHammer bit flips to storage

Restore bit flips to original values

Disable refresh to **prevent interruptions** in the core loop of our test **from refresh operations**

Induce RowHammer bit flips on a **fully charged row**

Testing Methodology



DRAM_RowHammer_Characterization():

foreach row in DRAM:

set *victim_row* to row

set *aggressor_row1* to *victim_row* - 1

set *aggressor_row2* to *victim_row* + 1

Disable DRAM refresh

Refresh *victim_row*

for $n = 1 \rightarrow HC$: // core test loop

activate *aggressor_row1*

activate *aggressor_row2*

Enable DRAM refresh

Record RowHammer bit flips to storage

Restore bit flips to original values

Disable refresh to **prevent interruptions** in the core loop of our test **from refresh operations**

Induce RowHammer bit flips on a **fully charged row**

Core test loop where we alternate accesses to adjacent rows

1 Hammer (HC) = two accesses

Prevent further retention failures

Record bit flips for analysis

1. RowHammer Vulnerability

Q. Can we induce RowHammer bit flips in all of our DRAM chips?

All chips are vulnerable, except many DDR3 chips

- A total of 1320 out of all 1580 chips **(84%)** are vulnerable
- Within **DDR3-old** chips, **only 12%** of chips (24/204) are vulnerable
- Within **DDR3-new** chips, **65%** of chips (148/228) are vulnerable

Newer DRAM chips are more vulnerable to RowHammer

2. Data Pattern Dependence


Q. Are some data patterns more effective in inducing RowHammer bit flips?

- We test **several data patterns** typically examined in prior work to identify the worst-case data pattern
- The worst-case data pattern is **consistent across chips** of the same manufacturer and DRAM type-node configuration
- We use the **worst-case data pattern** per DRAM chip to characterize each chip at **worst-case conditions** and **minimize the extensive testing time**

[More detail and figures in paper]

3. Hammer Count (HC) Effects


Q. How does the Hammer Count affect the number of bit flips induced?



Hammer Count = 2 Accesses,
one to each adjacent row of victim

4. Spatial Effects: Row Distance


Q. Where do RowHammer bit flips occur relative to aggressor rows?



The number of RowHammer bit flips that occur in a given row decreases as the distance from the **victim row (row 0)** increases.

4. Spatial Effects: Row Distance

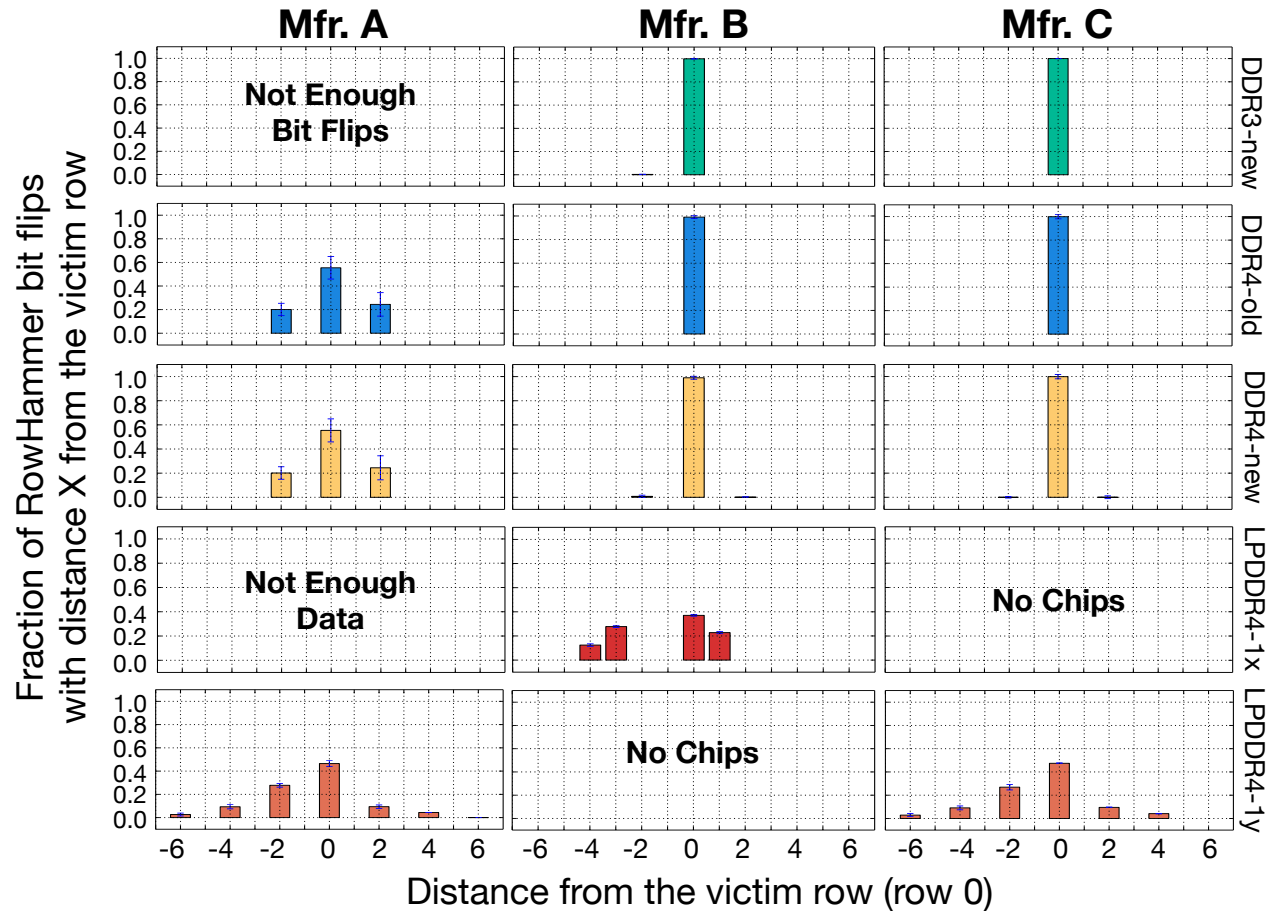
We normalize data by inducing a bit flip rate of 10^{-6} in each chip



Chips of newer DRAM technology nodes can exhibit RowHammer bit flips 1) in **more rows** and 2) **farther away** from the victim row.

4. Spatial Effects: Row Distance

We plot this data for each DRAM type-node configuration per manufacturer




[More analysis in the paper]

4. Spatial Distribution of Bit Flips

Q. How are RowHammer bit flips spatially distributed across a chip?

We normalize data by inducing a bit flip rate of 10^{-6} in each chip




The distribution of RowHammer bit flip density per word **changes significantly in LPDDR4 chips** from other DRAM types

At a bit flip rate of 10^{-6} , a 64-bit word can contain up to **4 bit flips**.
Even at this very low bit flip rate, a **very strong ECC** is required

4. Spatial Distribution of Bit Flips


We plot this data for each DRAM type-node configuration per manufacturer



[More analysis in the paper]

5. First RowHammer Bit Flips per Chip

What is the minimum Hammer Count required to cause bit flips (HC_{first})?



Evaluation Methodology

- **Cycle-level simulator:** Ramulator [Kim+, CAL'15]

<https://github.com/CMU-SAFARI/ramulator>

- 4GHz, 4-wide, 128 entry instruction window
- 48 8-core workload mixes randomly drawn from SPEC CPU2006 ($10 < \text{MPKI} < 740$)


- **Metrics to evaluate mitigation mechanisms**

1. **DRAM Bandwidth Overhead:** fraction of total system DRAM bandwidth consumption from mitigation mechanism
2. **Normalized System Performance:** normalized weighted speedup to a 100% baseline

Evaluation Methodology

- We evaluate **five** state-of-the-art mitigation mechanisms:
 - **Increased Refresh Rate** [Kim+, ISCA'14]
 - **PARA** [Kim+, ISCA'14]
 - **ProHIT** [Son+, DAC'17]
 - **MRLoc** [You+, DAC'19]
 - **TWiCe** [Lee+, ISCA'19]
- and **one** ideal refresh-based mitigation mechanism:
 - **Ideal**
- **More detailed descriptions in the paper on:**
 - Descriptions of mechanisms in our paper and the original publications
 - How we scale each mechanism to more vulnerable DRAM chips (lower **HC_{first}**)


Mitigation Mech. Eval. (Increased Refresh)




Substantial overhead for high HC_{first} values.

This mechanism does not support $HC_{first} < 32k$ due to the **prohibitively high refresh rates** required


Mitigation Mechanism Evaluation (PARA)



Mitigation Mechanism Evaluation (ProHIT)




Mitigation Mechanism Evaluation (MRLoc)



Models for **scaling** ProHIT and MRLoc for $HC_{first} < 2k$ are **not provided** and how to do so is **not intuitive**


Mitigation Mechanism Evaluation (TWiCe)



TWiCe does not support $HC_{first} < 32k$.

We evaluate an **ideal scalable version (TWiCe-ideal)** assuming it solves **two critical design issues**

Mitigation Mechanism Evaluation (Ideal)



Ideal mechanism issues a refresh command to a row **only right before** the row can potentially experience a RowHammer bit flip

Additional Details in the Paper


- **Single-cell RowHammer bit flip probability**
- More details on our **data pattern dependence** study
- Analysis of **Error Correcting Codes (ECC)** in mitigating RowHammer bit flips
- Additional **observations** on our data
- **Methodology details** for characterizing DRAM
- Further discussion on comparing data across different infrastructures
- **Discussion on scaling** each mitigation mechanism

RowHammer Reviews

Initial RowHammer Reviews

Disturbance Errors in DRAM: Demonstration, Characterization, and Prevention

Rejected (R2)

 863kB

Friday 31 May 2013 2:00:53pm PDT

b9bf06021da54cddf4cd0b3565558a181868b972

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

+ ABSTRACT

+ AUTHORS

	OveMer	Nov	WriQua	RevExp
Review #66A	1	4	4	4
Review #66B	5	4	5	3
Review #66C	2	3	5	4
Review #66D	1	2	3	4
Review #66E	4	4	4	3
Review #66F	2	4	4	3

Missing the Point **Reviews from Micro 2013**

PAPER WEAKNESSES

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

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

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.

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-gQSnBcdo>

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

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.

Final RowHammer Reviews

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

Accepted



639kB

21 Nov 2013 10:53:11pm CST |

f039be2735313b39304ae1c6296523867a485610

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

	OveMer	Nov	WriQua	RevConAnd
Review #41A	8	4	5	3
Review #41B	7	4	4	3
Review #41C	6	4	4	3
Review #41D	2	2	5	4
Review #41E	3	2	3	3
Review #41F	7	4	4	3

Some More 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 HPCA 2021, MICRO 2020, ISCA 2020, S&P 2020, SEC 2020
- Initially, RowHammer 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.

Top Pick Reviews

Review #54D Modified 1 Jan 2015 4:13:18pm PST  [Plain text](#)

SHORT PAPER SUMMARY

This paper observes through experimental measurements that

DRAM cells in a row can flip if a neighboring row is repeatedly

open and closed. One of the solutions proposed is: every time

1/29/15

<https://sampa.cs.washington.edu/microcrp/paper>

that a row is open and closed, refresh a neighboring row with a certain probability.

CHANCE OF IMPACT (?)

3. Minor impact

OVERALL MERIT (?)

2. Weak reject (Happy to discuss but unlikely to be chosen.)

COMMENTS FOR AUTHOR

Interesting paper for those interested in DRAM issues. I wonder if it is possible to gain an insight into why this happens.

I seem to remember that, during the presentation at ISCA, it was pointed out that DRAM manufacturers have already fixed the problem. So where is the novelty and long term impact?

An

SHORT PAPER SUMMARY

The paper explores how activating a row in a DRAM can corrupt nearby rows.

CHANCE OF IMPACT (?)

3. Minor impact

OVERALL MERIT (?)

3. Weak accept (Would consider for an honorable mention.)

COMMENTS FOR AUTHOR

This is a cute paper that explores DRAM errors in rows caused by accessing nearby rows. The results are certainly a bit surprising.

I can see this being a problem for PCs, where there is no ECC. Even if PCs, there are certain pieces that are

1/29/15

<https://sampa.cs.washington.edu/microcrp/paper/>

protected somehow (e.g., text is duplicated by Windows). What would be interesting to understand though is that whether the error rate from these "disturbance errors" is greater than the DRAM connector errors. If not, then do we really care?

I poked around a bit and DRAM vendors have already solved this problem. DRAM row hammering appears to be a known problem.

I don't see this as a real problem in running machines, but this could be made worse by malware. For example, a program running on a machine could be crashing other programs in an Amazon data center. Presumably, Amazon would find the offending program fairly easily though.

SAFAI

ew

Review #54E

Modified 4 Jan 2015 4:40:44am PST

[Plain text](#)

SHORT PAPER SUMMARY

This paper identifies a new class of DRAM errors called "disturbance" errors. The authors provide a characterization of such errors using DRAM chips dating back to 2008 and show that the disturbance error incidence is a relatively recent phenomenon (after 2010). Finally, the authors explore a set of possible mitigation solutions, while advocating one of them, called PARA (probabilistic adjacent row activation).

CHANCE OF IMPACT (?)

3. Minor impact

OVERALL MERIT (?)

2. Weak reject (Happy to discuss but unlikely to be chosen.)

COMMENTS FOR AUTHOR

The authors should be given due credit for identifying and characterizing an emerging new class of DRAM errors. However, it is not clear if this class of errors is significant enough in the future, given the many other modes of failure that DRAM vendors and users are primarily concerned with. As a reader of this paper, I could not but get the feeling that this is an interesting new DRAM error class, but could not find convincing arguments from the paper as to why this would constitute one of the key, first-order error behaviors affecting DRAMS of the future. The mitigation solution offered is simple and effective (I like it); but I was not

1/2

<https://sampa.cs.washington.edu/microcrp/paper/>

convinced that this paper will be cited in the future as one that opened up a brand new area of research and consequent use in practice.

SHORT PAPER SUMMARY

This paper makes the observation that when a DRAM row is opened (activated) and closed (precharged) repeatedly, it introduces disturbance errors in adjacent DRAM rows. The paper tests 129 DRAM modules from three manufacturers providing a wealth of information: 110 of the tested DRAM modules exhibit disturbance errors, and the trend seems to

1/29/15

<https://sampa.cs.washington.edu/microcrp/paper>

be increasing over the years. The paper then introduces a mechanism to prevent DRAM disturbance errors using a probabilistic approach. The paper also includes an FPGA-based testbed to analyze DRAM chips.

CHANCE OF IMPACT (?)

3. Minor impact

OVERALL MERIT (?)

4. Accept (Would argue for at least honorable mention.)

COMMENTS FOR AUTHOR

This is a great piece of work. It makes the point that disturbance errors occur in real DRAM chips, and that the problem is consistent across DRAM manufacturers and is getting worse over time. The paper characterizes a large number of real DRAM chips, clearly demonstrating the problem. The paper provides an FPGA-based testbed, and proposes a probabilistic mechanism to prevent DRAM disturbance errors. This is a very well executed piece of work overall.

While this is the first piece of work in the scientific literature to describe and characterize the problem, the problem of DRAM disturbance errors seems to be well-known to industry (as acknowledged in the paper). This somewhat reduces the significance of the work for consideration as a Top Pick.

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
- Maurice Wilkes Award Speech (10 minutes)
 - 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

Byzantine Failures

After RowHammer: 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

Some Selected Readings

Selected Readings on RowHammer (I)

- Our first detailed study: Rowhammer analysis and solutions (June 2014)
 - Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, Donghyuk Lee, Chris Wilkerson, Konrad Lai, and Onur Mutlu,
"Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors"
Proceedings of the 41st International Symposium on Computer Architecture (ISCA), Minneapolis, MN, June 2014. [[Slides \(pptx\)](#)] [[pdf](#)] [[Lightning Session Slides \(pptx\)](#)] [[pdf](#)] [[Source Code and Data](#)]
- Our Source Code to Induce Errors in Modern DRAM Chips (June 2014)
 - <https://github.com/CMU-SAFARI/rowhammer>
- Google Project Zero's Attack to Take Over a System (March 2015)
 - [Exploiting the DRAM rowhammer bug to gain kernel privileges](#) (Seaborn+, 2015)
 - <https://github.com/google/rowhammer-test>
 - **Double-sided Rowhammer**

Selected Readings on RowHammer (II)

- Remote RowHammer Attacks via JavaScript (July 2015)
 - <http://arxiv.org/abs/1507.06955>
 - <https://github.com/IAIK/rowhammerjs>
 - Gruss et al., DIMVA 2016.
 - **CLFLUSH-free Rowhammer**
 - “A fully automated attack that requires nothing but a website with JavaScript to **trigger faults on remote hardware**.”
 - “We can gain unrestricted access to systems of website visitors.”

- ANVIL: Software-Based Protection Against Next-Generation Rowhammer Attacks (March 2016)
 - <http://dl.acm.org/citation.cfm?doid=2872362.2872390>
 - Aweke et al., ASPLOS 2016
 - **CLFLUSH-free Rowhammer**
 - Software based monitoring for rowhammer detection

Selected Readings on RowHammer (III)

- **Dedup Est Machina: Memory Deduplication as an Advanced Exploitation Vector** (May 2016)
 - <https://www.ieee-security.org/TC/SP2016/papers/0824a987.pdf>
 - Bosman et al., IEEE S&P 2016.
 - Exploits Rowhammer and Memory Deduplication to overtake a browser
 - “We report on the **first reliable remote exploit for the Rowhammer vulnerability** running entirely in Microsoft Edge.”
 - “[an attacker] ... can reliably “own” a system with all defenses up, even if the software is entirely free of bugs.”

- **CAN't Touch This: Software-only Mitigation against Rowhammer Attacks targeting Kernel Memory** (August 2017)
 - <https://www.usenix.org/system/files/conference/usenixsecurity17/sec17-brasser.pdf>
 - Brasser et al., USENIX Security 2017.
 - Partitions physical memory into security domains, user vs. kernel; limits rowhammer-induced bit flips to the user domain.

Selected Readings on RowHammer (IV)

- **A New Approach for Rowhammer Attacks** (May 2016)
 - <https://ieeexplore.ieee.org/document/7495576>
 - Qiao et al., HOST 2016
 - **CLFLUSH-free RowHammer**
 - “Libc functions memset and memcpy are found capable of rowhammer.”
 - Triggers RowHammer with malicious inputs but benign code

- **One Bit Flips, One Cloud Flops: Cross-VM Row Hammer Attacks and Privilege Escalation** (August 2016)
 - https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_xiao.pdf
 - Xiao et al., USENIX Security 2016.
 - **“Technique that allows a malicious guest VM to have read and write accesses to arbitrary physical pages on a shared machine.”**
 - Graph-based algorithm to reverse engineer mapping of physical addresses in DRAM

Selected Readings on RowHammer (V)

- Curious Case of RowHammer: Flipping Secret Exponent Bits using Timing Analysis (August 2016)
 - https://link.springer.com/content/pdf/10.1007%2F978-3-662-53140-2_29.pdf
 - Bhattacharya et al., CHES 2016
 - Combines timing analysis to perform **rowhammer on cryptographic keys** stored in memory

- DRAMA: Exploiting DRAM Addressing for Cross-CPU Attacks (August 2016)
 - https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_pessl.pdf
 - Pessl et al., USENIX Security 2016
 - **Shows RowHammer failures on DDR4 devices despite TRR solution**
 - Reverse engineers address mapping functions to improve existing RowHammer attacks

Selected Readings on RowHammer (VI)

- Flip Feng Shui: Hammering a Needle in the Software Stack (August 2016)
 - https://www.usenix.org/system/files/conference/usenixsecurity16/sec16_paper_razavi.pdf
 - Razavi et al., USENIX Security 2016.
 - Combines memory deduplication and RowHammer
 - **"A malicious VM can gain unauthorized access to a co-hosted VM running OpenSSH."**
 - Breaks OpenSSH public key authentication

- Drammer: Deterministic Rowhammer Attacks on Mobile Platforms (October 2016)
 - <http://dl.acm.org/citation.cfm?id=2976749.2978406>
 - Van Der Veen et al., ACM CCS 2016
 - **Can take over an ARM-based Android system deterministically**
 - Exploits predictable physical memory allocator behavior
 - Can deterministically place security-sensitive data (e.g., page table) in an attacker-chosen, vulnerable location in memory

Selected Readings on RowHammer (VII)

- When Good Protections go Bad: Exploiting anti-DoS Measures to Accelerate Rowhammer Attacks (May 2017)
 - <https://web.eecs.umich.edu/~misiker/resources/HOST-2017-Misiker.pdf>
 - Aga et al., HOST 2017
 - “A virtual-memory based cache-flush free attack that is sufficiently fast to **rowhammer with double rate refresh.**”
 - Enabled by Cache Allocation Technology

- SGX-Bomb: Locking Down the Processor via Rowhammer Attack (October 2017)
 - <https://dl.acm.org/citation.cfm?id=3152709>
 - Jang et al., SysTEX 2017
 - “Launches the Rowhammer attack against enclave memory to trigger the processor lockdown.”
 - **Running unknown enclave programs on the cloud can shut down servers shared with other clients.**

Selected Readings on RowHammer (VIII)

- **Another Flip in the Wall of Rowhammer Defenses** (May 2018)
 - <https://arxiv.org/pdf/1710.00551.pdf>
 - Gruss et al., IEEE S&P 2018
 - **A new type of Rowhammer attack which only hammers one single address**, which can be done without knowledge of physical addresses and DRAM mappings
 - Defeats static analysis and performance counter analysis defenses by running inside an SGX enclave

- **GuardION: Practical Mitigation of DMA-Based Rowhammer Attacks on ARM** (June 2018)
 - https://link.springer.com/chapter/10.1007/978-3-319-93411-2_5
 - Van Der Veen et al., DIMVA 2018
 - Presents RAMPAGE, a DMA-based RowHammer attack against the latest Android OS

Selected Readings on RowHammer (IX)

- Grand Pwning Unit: Accelerating Microarchitectural Attacks with the GPU (May 2018)
 - <https://www.vusec.net/wp-content/uploads/2018/05/glitch.pdf>
 - Frigo et al., IEEE S&P 2018.
 - The first end-to-end remote Rowhammer exploit on mobile platforms that use our GPU-based primitives in orchestration to **compromise browsers on mobile devices in under two minutes.**
- Throwhammer: Rowhammer Attacks over the Network and Defenses (July 2018)
 - https://www.cs.vu.nl/~herbertb/download/papers/throwhammer_atc18.pdf
 - Tatar et al., USENIX ATC 2018.
 - “[We] show that **an attacker can trigger and exploit Rowhammer bit flips directly from a remote machine by only sending network packets.**”

Selected Readings on RowHammer (X)

- **Nethammer: Inducing Rowhammer Faults through Network Requests** (July 2018)
 - <https://arxiv.org/pdf/1805.04956.pdf>
 - Lipp et al., arxiv.org 2018.
 - “Nethammer is the first truly **remote Rowhammer attack**, without a single attacker-controlled line of code on the targeted system.”
- **ZebRAM: Comprehensive and Compatible Software Protection Against Rowhammer Attacks** (October 2018)
 - <https://www.usenix.org/system/files/osdi18-konoth.pdf>
 - Konoth et al., OSDI 2018
 - A new pure-software protection mechanism against RowHammer.

Selected Readings on RowHammer (XI.A)

- PassMark Software, memtest86, since 2014
 - <https://www.memtest86.com/troubleshooting.htm#hammer>

Why am I only getting errors during Test 13 Hammer Test?

The Hammer Test is designed to detect RAM modules that are susceptible to disturbance errors caused by charge leakage. This phenomenon is characterized in the research paper **Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors** by Yoongu Kim et al. According to the research, a significant number of RAM modules manufactured 2010 or newer are affected by this defect. In simple terms, susceptible RAM modules can be subjected to disturbance errors when repeatedly accessing addresses in the same memory bank but different rows in a short period of time. Errors occur when the repeated access causes charge loss in a memory cell, before the cell contents can be refreshed at the next DRAM refresh interval.

Starting from MemTest86 v6.2, the user may see a warning indicating that the RAM may be vulnerable to high frequency row hammer bit flips. This warning appears when errors are detected during the first pass (maximum hammer rate) but no errors are detected during the second pass (lower hammer rate). See **MemTest86 Test Algorithms** for a description of the two passes that are performed during the Hammer Test (Test 13). When performing the second pass, address pairs are hammered only at the rate deemed as the maximum allowable by memory vendors (200K accesses per 64ms). Once this rate is exceeded, the integrity of memory contents may no longer be guaranteed. If errors are detected in both passes, errors are reported as normal.

The errors detected during Test 13, albeit exposed only in extreme memory access cases, are most certainly real errors. During typical home PC usage (eg. web browsing, word processing, etc.), it is less likely that the memory usage pattern will fall into the extreme case that make it vulnerable to disturbance errors. It may be of greater concern if you were running highly sensitive equipment such as medical equipment, aircraft control systems, or bank database servers. It is impossible to predict with any accuracy if these errors will occur in real life applications. One would need to do a major scientific study of 1000 of computers and their usage patterns, then do a forensic analysis of each application to study how it makes use of the RAM while it executes. To date, we have only seen 1-bit errors as a result of running the Hammer Test.

Selected Readings on RowHammer (XI.B)

- PassMark Software, memtest86, since 2014
 - <https://www.memtest86.com/troubleshooting.htm#hammer>

Detection and mitigation of row hammer errors

The ability of MemTest86 to detect and report on row hammer errors depends on several factors and what mitigations are in place. To generate errors adjacent memory rows must be repeatedly accessed. But hardware features such as multiple channels, interleaving, **scrambling**, Channel Hashing, NUMA & XOR schemes make it nearly impossible (for an arbitrary CPU & RAM stick) to know which memory addresses correspond to which rows in the RAM. Various mitigations might also be in place. Different BIOS firmware might set the refresh interval to different values (tREFI). The shorter the interval the more resistant the RAM will be to errors. But shorter intervals result in higher power consumption and increased processing overhead. Some CPUs also support pseudo target row refresh (pTRR) that can be used in combination with pTRR-compliant RAM. This field allows the RAM stick to indicate the MAC (Maximum Active Count) level which is the RAM can support. A typical value might be 200,000 row activations. Some CPUs also support the Joint Electron Design Engineering Council (JEDEC) Targeted Row Refresh (TRR) algorithm. The TRR is an improved version of the previously implemented pTRR algorithm and does not inflict any performance drop or additional power usage. As a result the row hammer test implemented in MemTest86 maybe not be the worst case possible and vulnerabilities in the underlying RAM might be undetectable due to the mitigations in place in the BIOS and CPU.



Security Implications (ISCA 2014)

- *Breach of memory protection*
 - OS page (4KB) fits inside DRAM row (8KB)
 - Adjacent DRAM row → Different OS page
- *Vulnerability: disturbance attack*
 - By accessing its own page, a program could corrupt pages belonging to another program
- *We constructed a proof-of-concept*
 - Using only user-level instructions

Keeping Future Memory Secure