

Computer Architecture

Lecture 12a: Low-Latency Memory II

Prof. Onur Mutlu

ETH Zürich

Fall 2021

5 November 2021

SAFARI Live Seminar Nov 7

SAFARI Live Seminars in Computer Architecture

Damla Senol Cali, Bionano Genomics

Accelerating Genome Sequence Analysis via Efficient Hardware/Algorithm Co-Design

SAFARI
SAFARI Research Group

Our Goal & Approach

- **Our Goal:**
Accelerating genome sequence analysis by **efficient hardware/algorithm co-design**
- **Our Approach:**
 - (1) Analyze the **multiple steps** and the **associated tools** in the genome sequence analysis pipeline,
 - (2) Expose the **tradeoffs** between accuracy, performance, memory usage and scalability, and
 - (3) Co-design **fast and efficient algorithms** along with **scalable and energy-efficient customized hardware accelerators** for the key bottleneck steps of the pipeline

Damla Senol Cali

SAFARI

10

7 Sun
Nov
2021



Sunday, November 07 at 6:00 pm Zurich time (CET)

[Damla Senol Cali, Bionano Genomics](#)

[Accelerating Genome Sequence Analysis via Efficient Hardware/Algorithm Co-Design](#)

Livestream on YouTube [Link](#)

[Abstract & Speaker Bio](#)


SAFARI Live Seminar Nov 8

SAFARI Live Seminars in Computer Architecture

Gennady Pekhimenko, University of Toronto
Machine Learning Tools in Action


ETH zürich

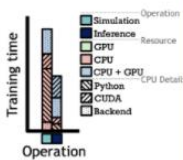
SAFARI
SAFARI Research Group



RL-Scope: Cross-Stack Profiling for Deep Reinforcement Learning Workloads

8 Mo Nov 2021





Category	On-policy $\pi(a s)$	Off-policy $V(S), Q(s,a)$
Framework	PyTorch	TensorFlow
Language	Python	C++
Performance	3.5x more simulation-bound	2.3x more CPU inflation!
Transitions	More transitions	More CPU inflation!
Training time	$\geq 30\%$ of training time!	100%
GPU usage	GPU usage is low (< 14%)	True GPU-bound time is low!

RL \neq Supervised Learning

Monday, November 08 at 4:00 pm Zurich time (CET)

[Gennady Pekhimenko](#), [University of Toronto](#), [EcoSystem group](#)

[Machine Learning Tools in Action](#)

Livestream on YouTube [Link](#)

[Abstract & Speaker Bio](#)

SAFARI Live Seminar Nov 11

SAFARI Live Seminars in Computer Architecture

Serghei Mangul, Mangul Lab, USC

Opportunities and challenges of computational data-driven immunology

SAFARI
SAFARI Research Group

ETH zürich



Opportunities and challenges of computational data-driven immunology

Serghei Mangul, Ph.D
Assistant Professor,
University of Southern California

<https://mangul-lab-usc.github.io/>

11 Thu Nov 2021

USC University of Southern California

Thursday, November 11 at 11:00 am Zurich time (CET)

[Serghei Mangul](#), University of Southern California, [Mangul Lab](#)

[Opportunities and challenges of computational data-driven immunology](#)

Livestream on YouTube [Link](#)

Note: this talk will take place in person at ETH Zentrum. Room info will be [posted soon!](#)

[Abstract & Speaker Bio](#)

SAFARI Live Seminars


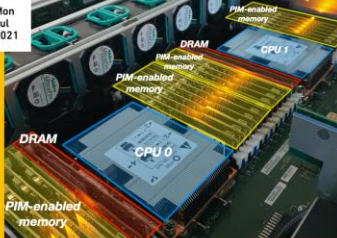
SAFARI Live Seminars in Computer Architecture

Dr. Juan Gómez Luna, ETH Zurich

Understanding a Modern Processing-in-Memory Architecture: Benchmarking and Experimental Characterization

SAFARI
SAFARI Research Group

12 Mon Jul 2021

SAFARI Live Seminars in Computer Architecture

Dr. Andrew Walker, Schiltron Corporation & Nexgen Power Systems

An Addition to Low Cost Per Memory Bit – How to Recognize it and What to Do About it

SAFARI
SAFARI Research Group

19 Mo Jul 2021




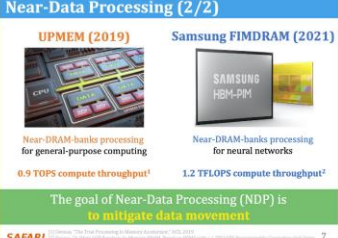

SAFARI Live Seminars in Computer Architecture

Geraldo F. Oliveira, ETH Zurich

DAMOV: A New Methodology and Benchmark Suite for Evaluating Data Movement Bottlenecks

SAFARI
SAFARI Research Group

22 Do Jul 2021


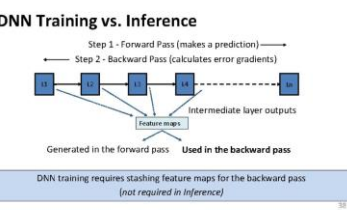
SAFARI Live Seminars in Computer Architecture

Gennady Pekhimenko, University of Toronto

Efficient DNN Training at Scale: from Algorithms to Hardware

SAFARI
SAFARI Research Group

5 Do Aug 2021


SAFARI Live Seminars in Computer Architecture

Jawad Haj-Yahya, Huawei Research Center Zurich

Power Management Mechanisms in Modern Microprocessors and Their Security Implications

SAFARI
SAFARI Research Group

16 Mo Aug 2021




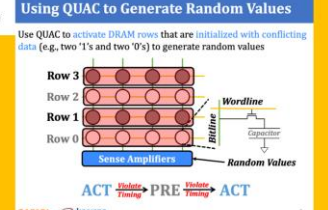

SAFARI Live Seminars in Computer Architecture

Ataberk Olgun, TOBB & ETH Zurich

QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAM Chips

SAFARI
SAFARI Research Group

15 Mi Sep 2021


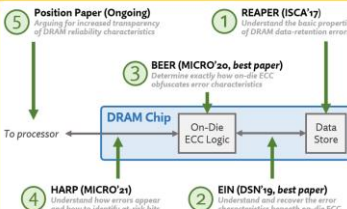
SAFARI Live Seminars in Computer Architecture

Minesh Patel, ETH Zurich

Enabling Effective Error Mitigation in Memory Chips That Use On-Die ECCs

SAFARI
SAFARI Research Group

21 Tue Sep 2021


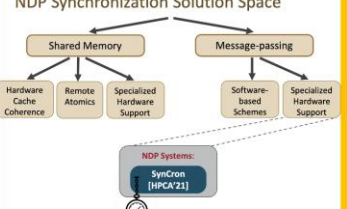



SAFARI Live Seminars in Computer Architecture

Christina Giannoula, National Technical University of Athens
Efficient Synchronization Support for Near-Data-Processing Architectures

SAFARI
SAFARI Research Group

27 Mo Sep 2021


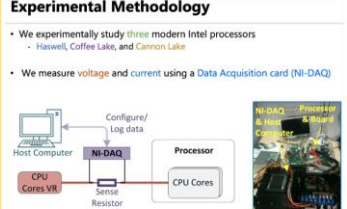
SAFARI Live Seminars in Computer Architecture

Jawad Haj-Yahya, Huawei Research Center Zurich

Security Implications of Power Management Mechanisms in Modern Processors, Current Studies and Future Trends

SAFARI
SAFARI Research Group

4 Mo Okt 2021

<https://safari.ethz.ch/safari-seminar-series/>

Four Key Current Directions

- Fundamentally **Secure/Reliable/Safe** Architectures
- Fundamentally **Energy-Efficient** Architectures
 - **Memory-centric** (Data-centric) Architectures
- Fundamentally **Low-Latency and Predictable** Architectures
- Architectures for **AI/ML, Genomics, Medicine, Health, ...**

Why the Long Memory Latency?

- Reason 1: Design of DRAM Micro-architecture
 - Goal: Maximize capacity/area, not minimize latency
- Reason 2: “One size fits all” approach to latency specification
 - Same latency parameters for all temperatures
 - Same latency parameters for all DRAM chips
 - Same latency parameters for all parts of a DRAM chip
 - Same latency parameters for all supply voltage levels
 - Same latency parameters for all application data
 - ...

More on AL-DRAM

- Donghyuk Lee, Yoongu Kim, Gennady Pekhimenko, Samira Khan, Vivek Seshadri, Kevin Chang, and Onur Mutlu,
"Adaptive-Latency DRAM: Optimizing DRAM Timing for the Common-Case"
Proceedings of the 21st International Symposium on High-Performance Computer Architecture (HPCA), Bay Area, CA, February 2015.
[[Slides \(pptx\) \(pdf\)](#)] [[Full data sets](#)]

Adaptive-Latency DRAM: Optimizing DRAM Timing for the Common-Case

Donghyuk Lee Yoongu Kim Gennady Pekhimenko
Samira Khan Vivek Seshadri Kevin Chang Onur Mutlu

Carnegie Mellon University

Different Types of Latency Variation

- AL-DRAM exploits latency variation
 - Across time (different temperatures)
 - Across chips

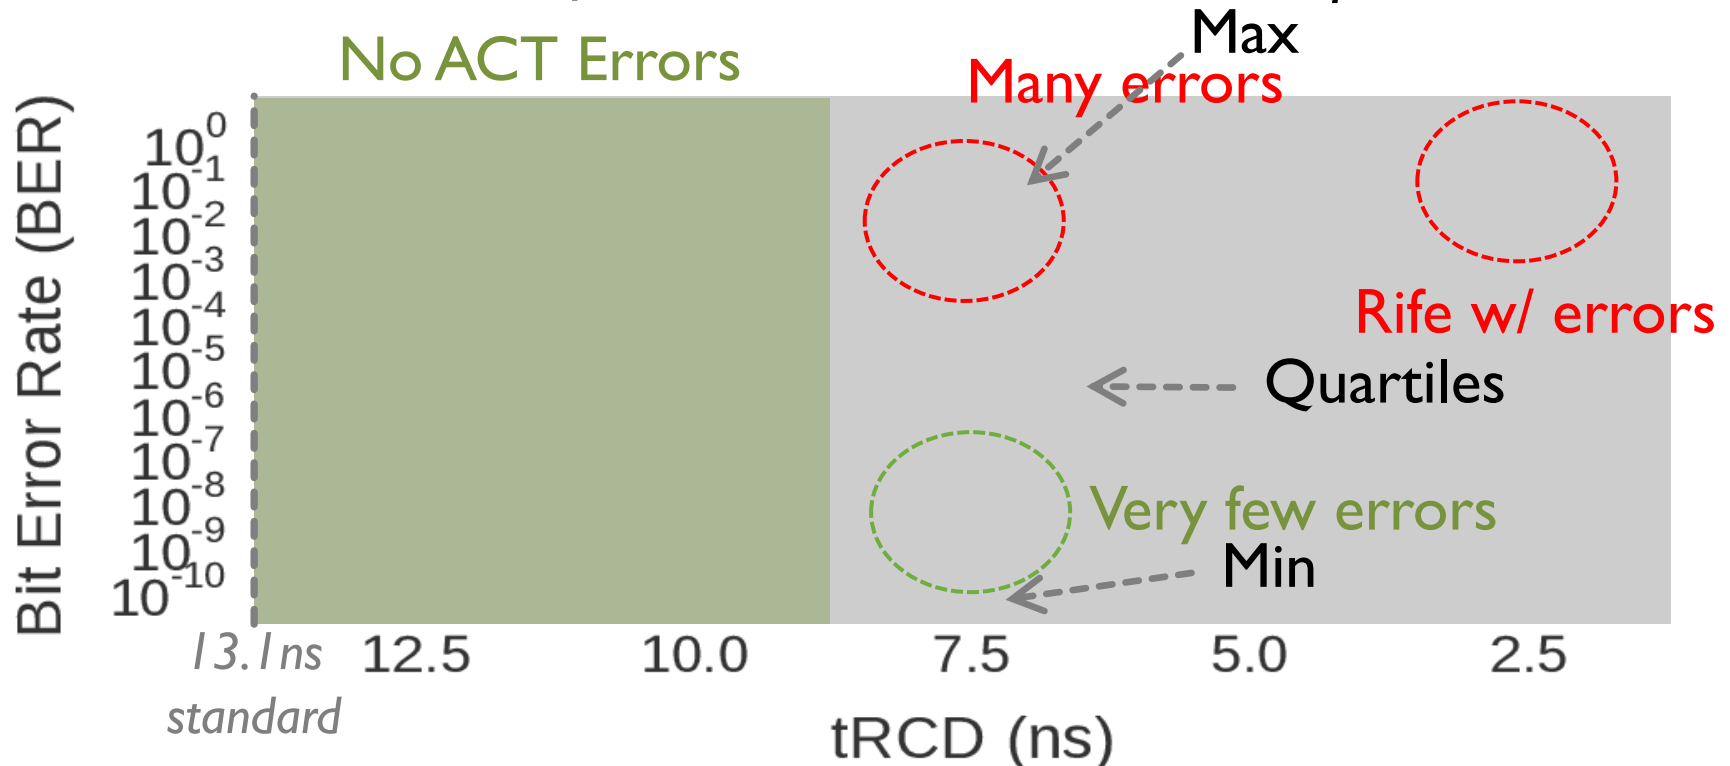
- Is there also latency variation within a chip?
 - Across different parts of a chip

Why the Long Memory Latency?

- Reason 1: Design of DRAM Micro-architecture
 - Goal: Maximize capacity/area, not minimize latency
- Reason 2: “One size fits all” approach to latency specification
 - Same latency parameters for all temperatures
 - Same latency parameters for all DRAM chips
 - Same latency parameters for all parts of a DRAM chip
 - Same latency parameters for all supply voltage levels
 - Same latency parameters for all application data
 - ...

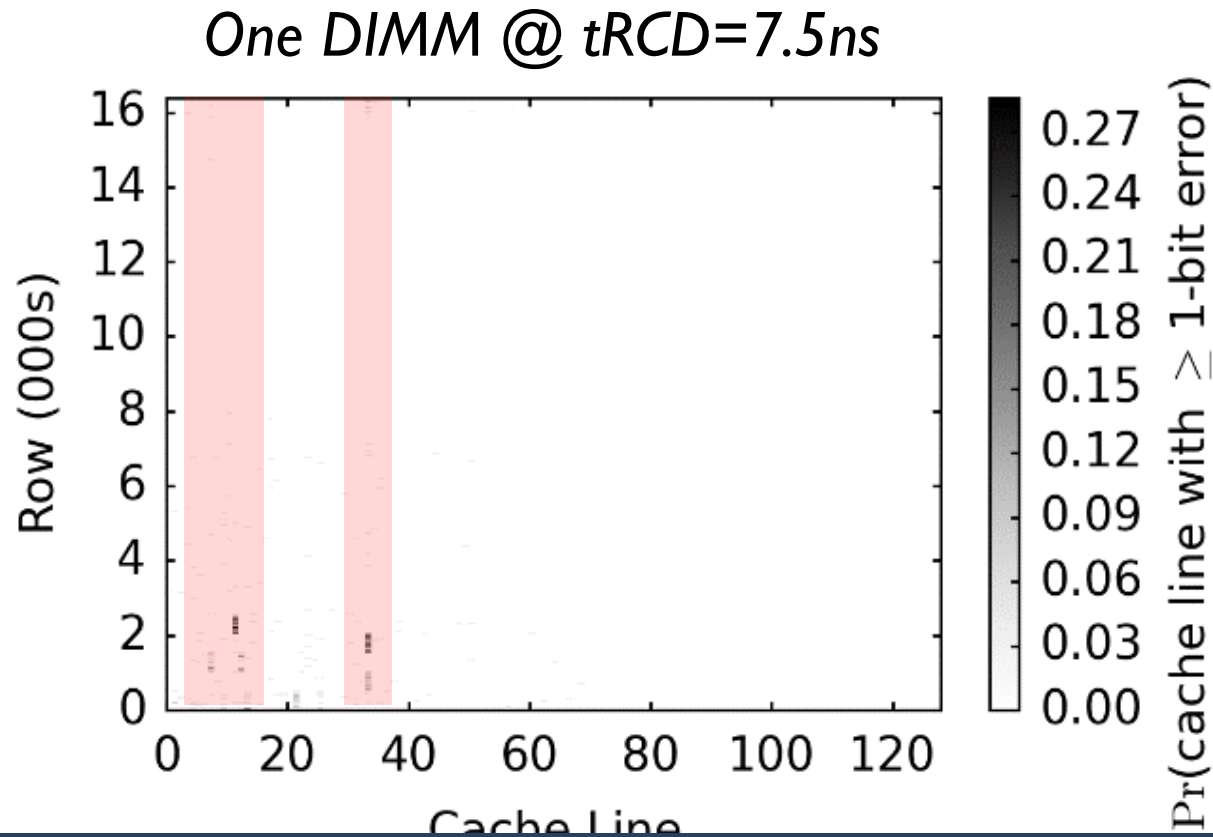
Variation in Activation Errors

Results from 7500 rounds over 240 chips



Modern DRAM chips exhibit significant variation in activation latency

Spatial Locality of Activation Errors

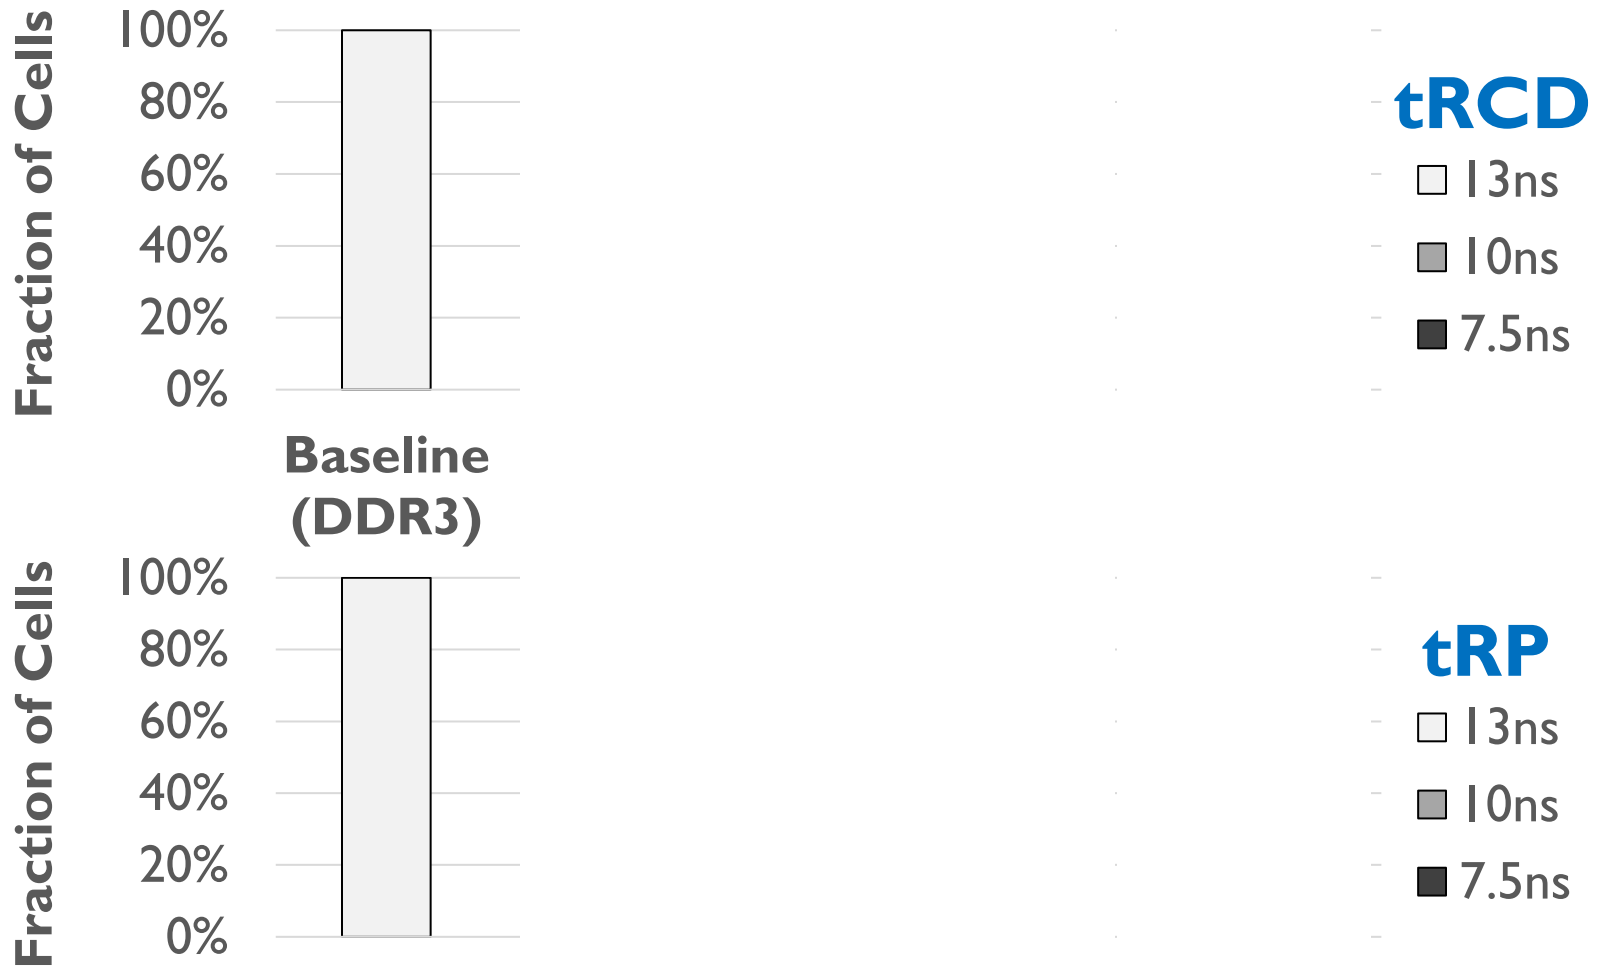


Activation errors are concentrated at certain columns of cells

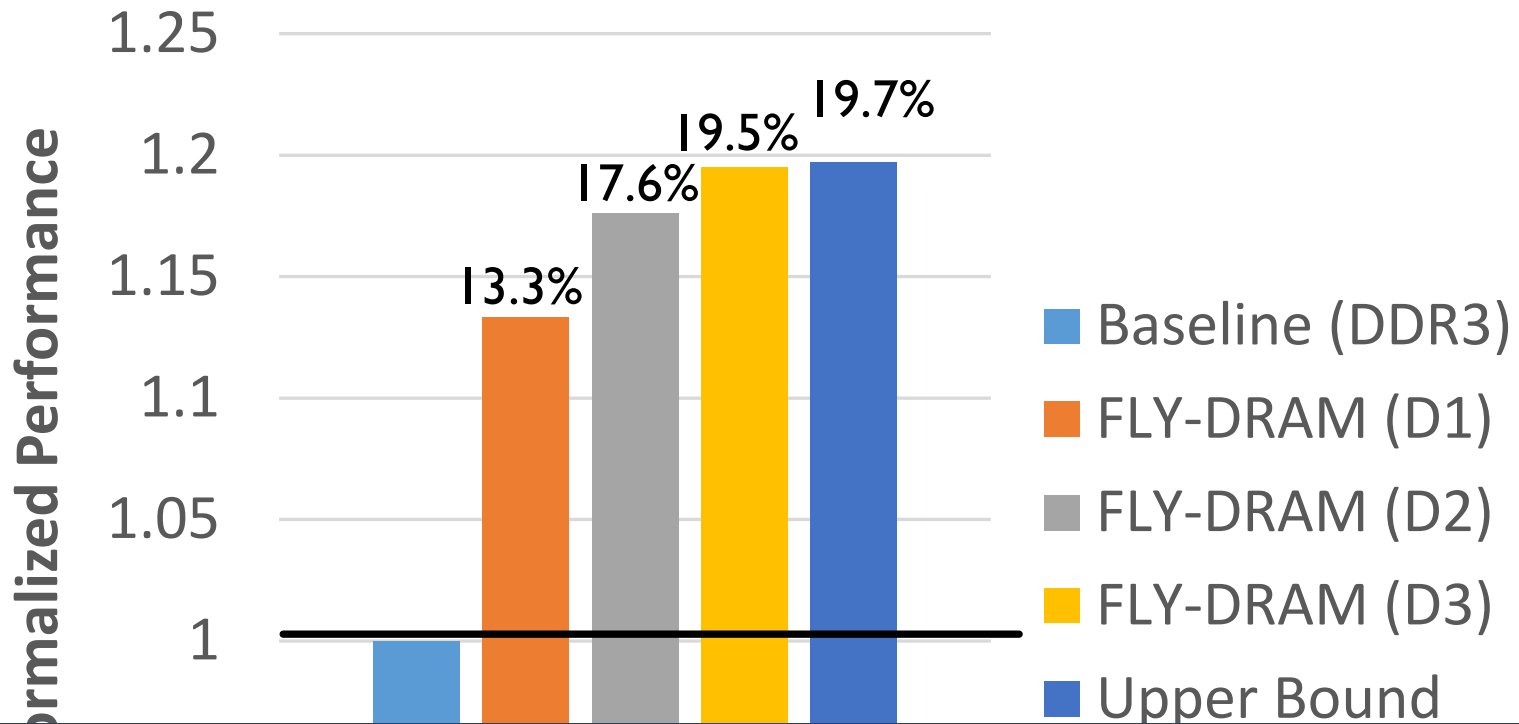
Mechanism to Reduce DRAM Latency

- **Observation:** DRAM timing errors (slow DRAM cells) are concentrated in certain DRAM regions
- **Flexible-Latency (FLY) DRAM**
 - A software-transparent design that reduces latency
- **Key idea:**
 - 1) Divide memory into regions of different latencies
 - 2) *Memory controller:* Use lower latency for regions without slow cells; higher latency for other regions

FLY-DRAM Configurations



Results



**FLY-DRAM improves performance
by exploiting spatial latency variation in DRAM**

FLY-DRAM: Advantages & Disadvantages

■ Advantages

- + Reduces latency significantly
- + Exploits significant within-chip latency variation

■ Disadvantages

- Need to determine reliable operating latencies for different parts of a chip → higher testing cost
- More complicated controller

Analysis of Latency Variation in DRAM Chips

- Kevin Chang, Abhijith Kashyap, Hasan Hassan, Samira Khan, Kevin Hsieh, Donghyuk Lee, Saugata Ghose, Gennady Pekhimenko, Tianshi Li, and Onur Mutlu,

"Understanding Latency Variation in Modern DRAM Chips: Experimental Characterization, Analysis, and Optimization"

*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Antibes Juan-Les-Pins, France, June 2016.*

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

[[Source Code](#)]

Understanding Latency Variation in Modern DRAM Chips: Experimental Characterization, Analysis, and Optimization

Kevin K. Chang¹

Abhijith Kashyap¹

Hasan Hassan^{1,2}

Saugata Ghose¹

Kevin Hsieh¹

Donghyuk Lee¹

Tianshi Li^{1,3}

Gennady Pekhimenko¹

Samira Khan⁴

Onur Mutlu^{5,1}

¹Carnegie Mellon University ²TOBB ETÜ ³Peking University ⁴University of Virginia ⁵ETH Zürich

SAFARI

Putting It All Together: Solar-DRAM

Solar-DRAM: Putting It Together

- Jeremie S. Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu,
"Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines"
Proceedings of the 36th IEEE International Conference on Computer Design (ICCD), Orlando, FL, USA, October 2018.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (16 minutes)]

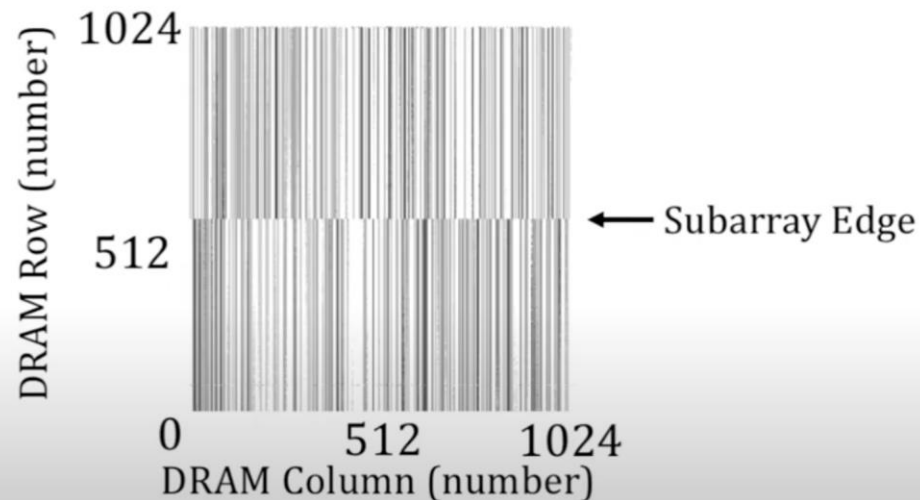
Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines

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

More on Solar DRAM

Spatial Distribution of Failures

How are activation failures spatially distributed in DRAM?



Activation failures are **highly constrained**
to local bitlines (i.e., subarrays)

Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines - ICCD 2018

101 views • Oct 23, 2018

4 0 SHARE SAVE ...

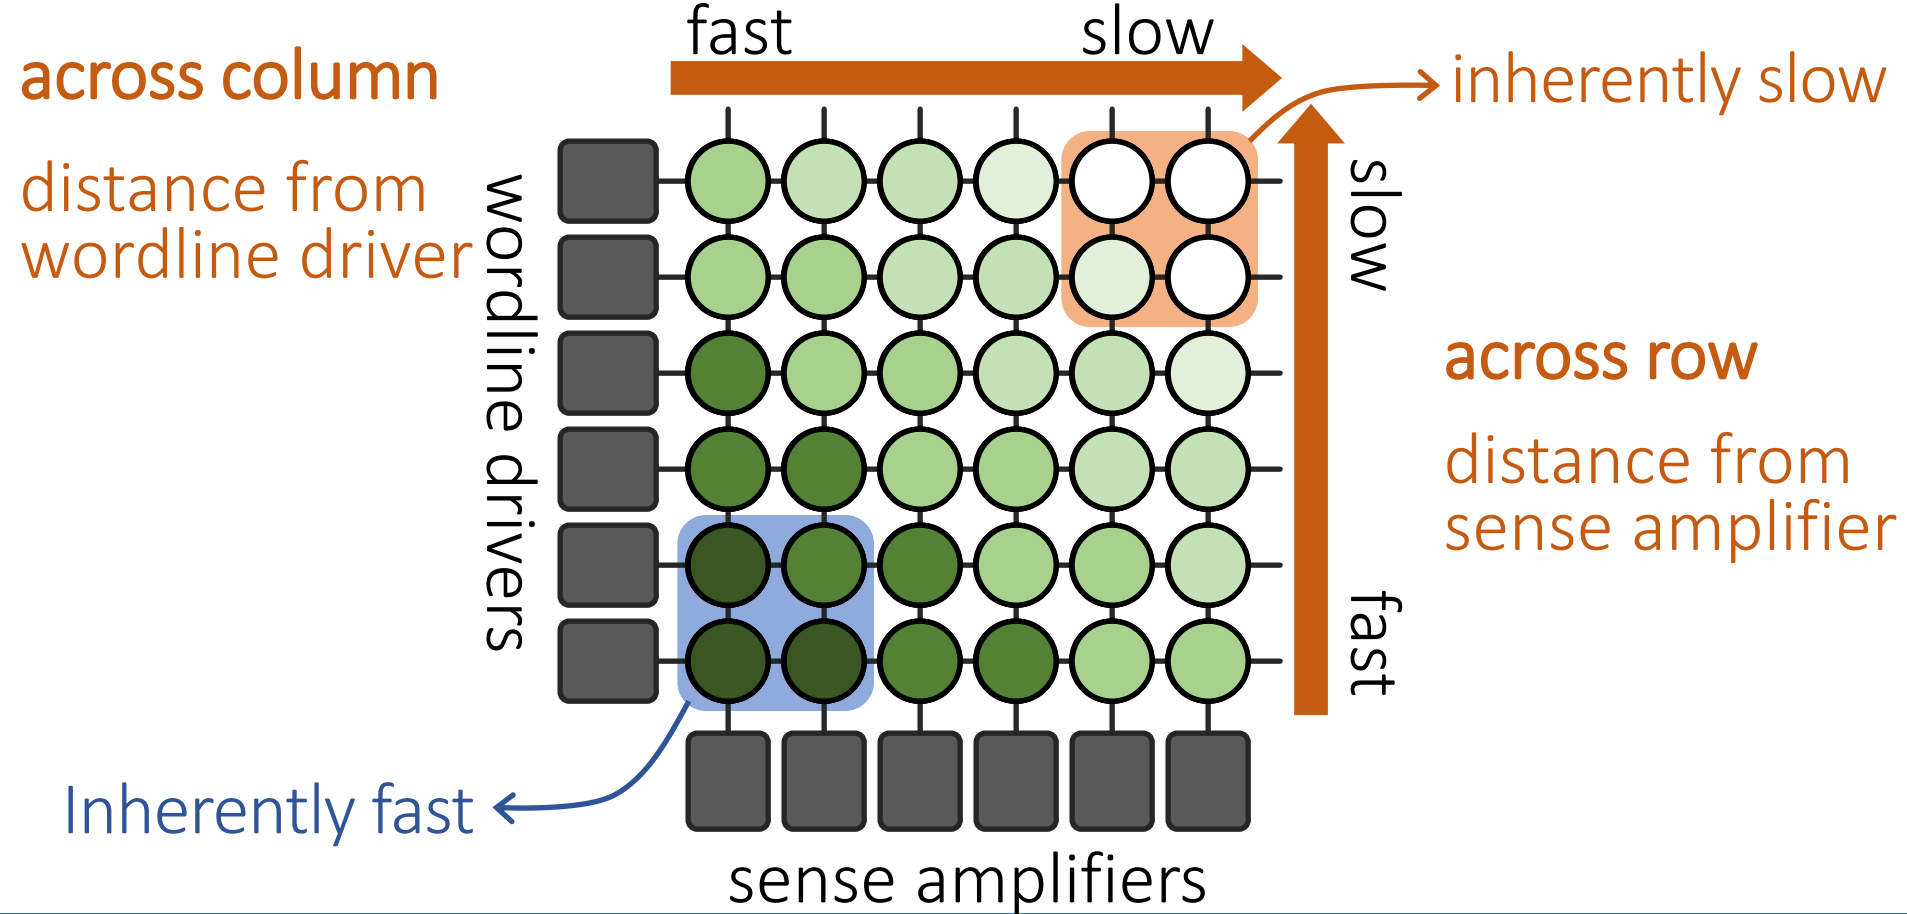


Jeremie Kim
18 subscribers

SUBSCRIBE

Why Is There Spatial Latency Variation Within a Chip?

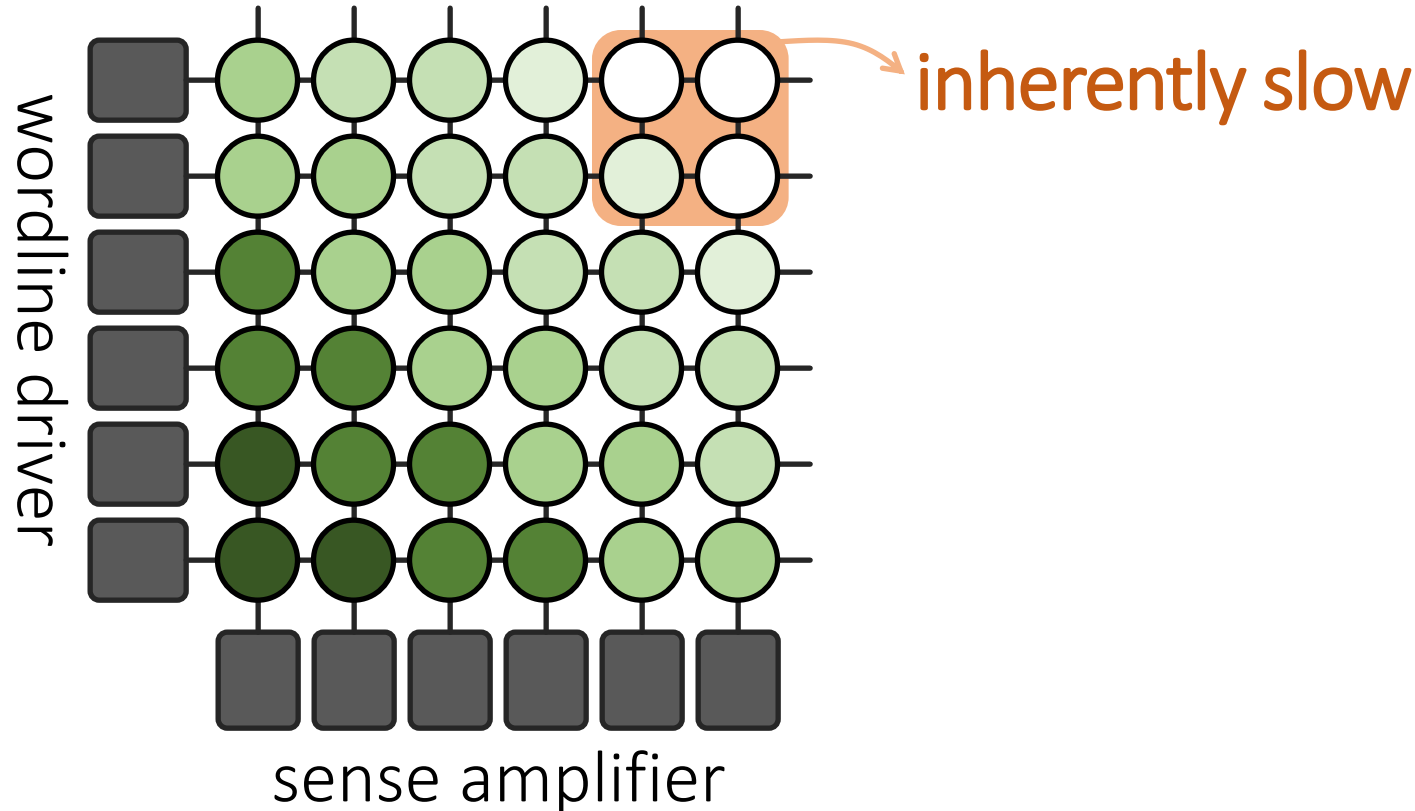
What Is Design-Induced Variation?



Systematic variation in cell access times
caused by the ***physical organization*** of DRAM

DIVA Online Profiling

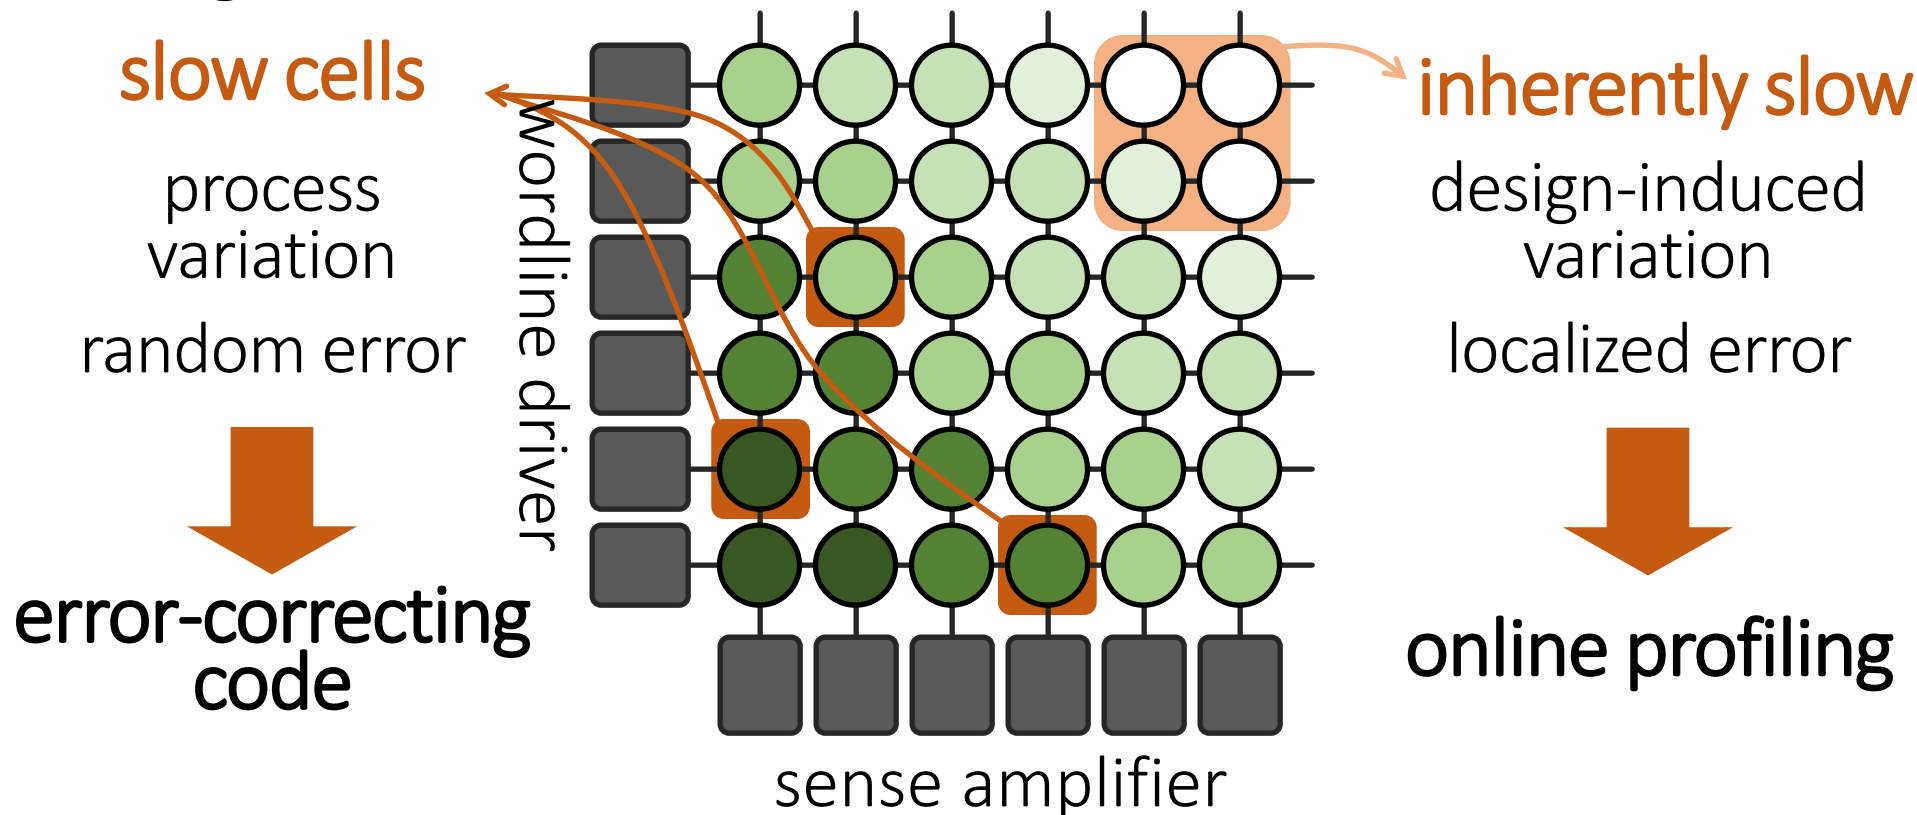
Design-Induced-Variation-Aware



Profile *only slow regions* to determine min. latency
→ *Dynamic* & *low cost* latency optimization

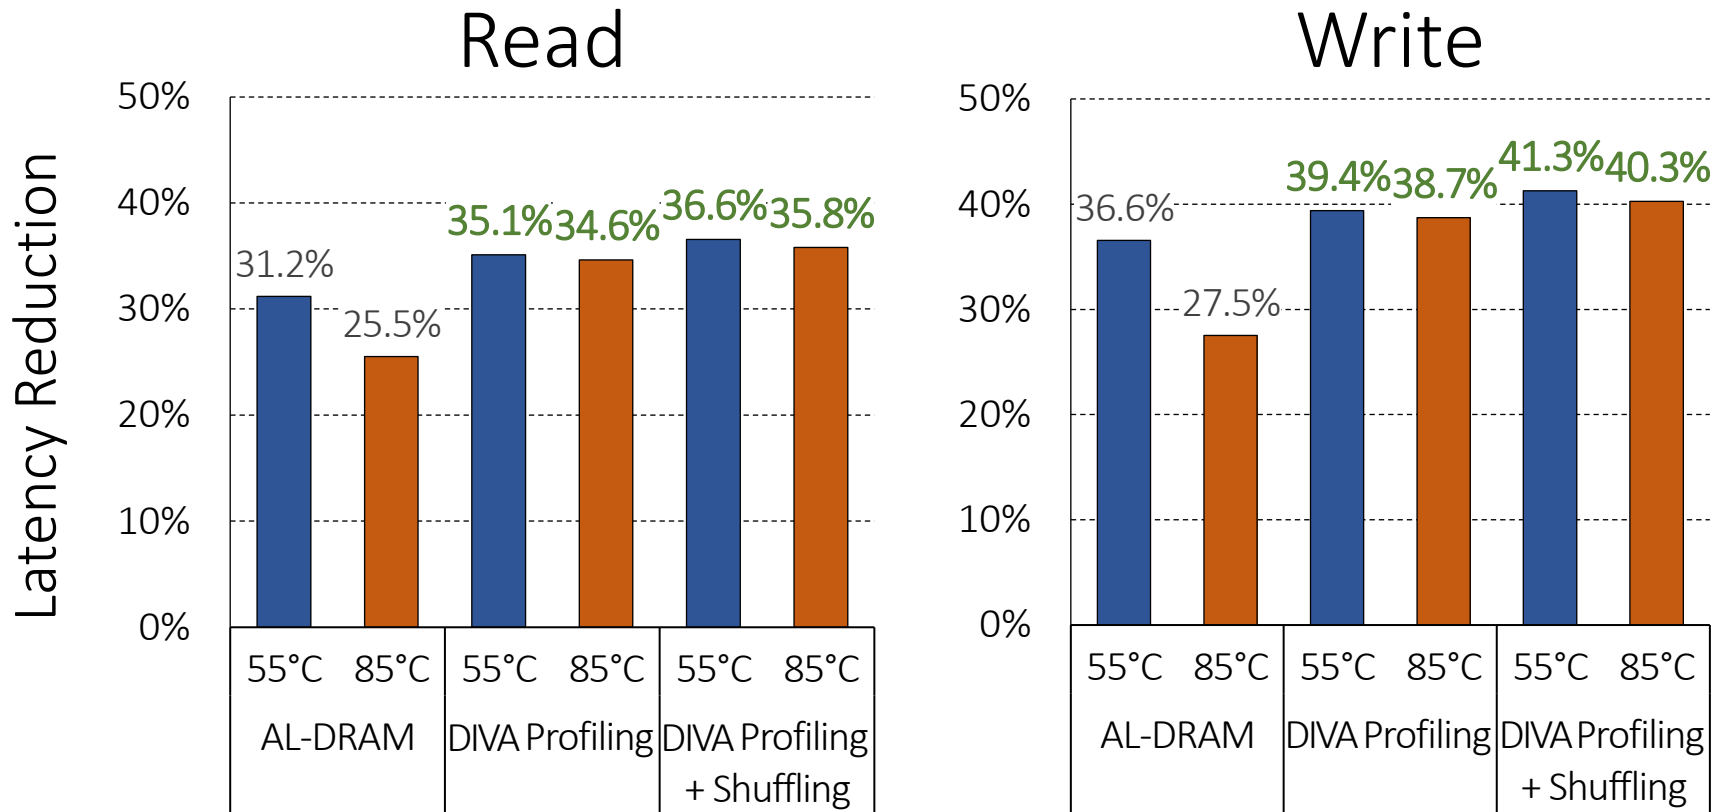
DIVA Online Profiling

Design-Induced-Variation-Aware



Combine **error-correcting codes** & **online profiling**
→ **Reliably** reduce DRAM latency

DIVA-DRAM Reduces Latency



DIVA-DRAM *reduces latency more aggressively*
and uses ECC to correct random slow cells

DIVA-DRAM: Advantages & Disadvantages

■ Advantages

- ++ Automatically finds the lowest reliable operating latency at system runtime (lower production-time testing cost)
- + Reduces latency more than prior methods (w/ ECC)
- + Reduces latency at high temperatures as well

■ Disadvantages

- Requires knowledge of inherently-slow regions
- Requires ECC (Error Correcting Codes)
- Imposes overhead during runtime profiling
- More complicated memory controller (capable of profiling)

Design-Induced Latency Variation in DRAM

- Donghyuk Lee, Samira Khan, Lavanya Subramanian, Saugata Ghose, Rachata Ausavarungnirun, Gennady Pekhimenko, Vivek Seshadri, and Onur Mutlu,
"Design-Induced Latency Variation in Modern DRAM Chips: Characterization, Analysis, and Latency Reduction Mechanisms"
*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Urbana-Champaign, IL, USA, June 2017.*

Design-Induced Latency Variation in Modern DRAM Chips: Characterization, Analysis, and Latency Reduction Mechanisms

Donghyuk Lee, NVIDIA and Carnegie Mellon University

Samira Khan, University of Virginia

Lavanya Subramanian, Saugata Ghose, Rachata Ausavarungnirun, Carnegie Mellon University

Gennady Pekhimenko, Vivek Seshadri, Microsoft Research

Onur Mutlu, ETH Zürich and Carnegie Mellon University

Why the Long Memory Latency?

- Reason 1: Design of DRAM Micro-architecture
 - Goal: Maximize capacity/area, not minimize latency
- Reason 2: “One size fits all” approach to latency specification
 - Same latency parameters for all temperatures
 - Same latency parameters for all DRAM chips
 - Same latency parameters for all parts of a DRAM chip
 - Same latency parameters for all supply voltage levels
 - Same latency parameters for all application data
 - ...

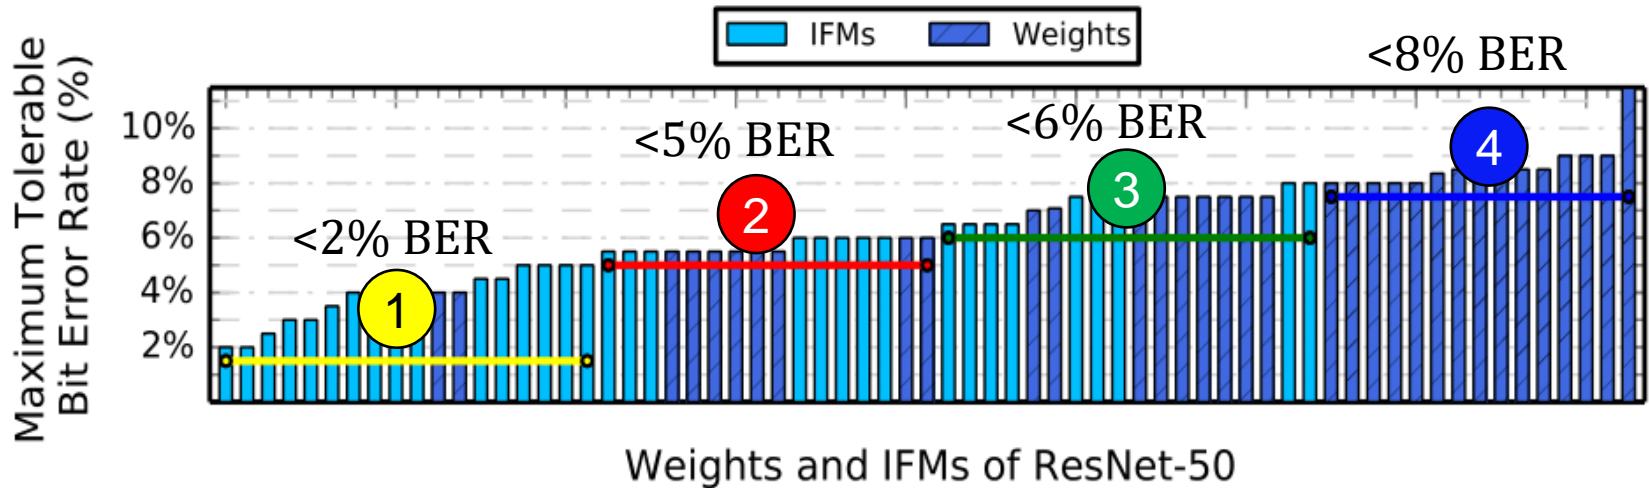
Data-Aware DRAM Latency for DNN Inference

- Deep Neural Network evaluation is very DRAM-intensive (especially for large networks)
1. Some data and layers in DNNs are very tolerant to errors
 2. Reduce DRAM latency and voltage on such data and layers
 3. While still achieving a user-specified DNN accuracy target by making training DRAM-error-aware

**Data-aware management of DRAM latency and voltage
for Deep Neural Network Inference**

Example DNN Data Type to DRAM Mapping

Mapping example of ResNet-50:



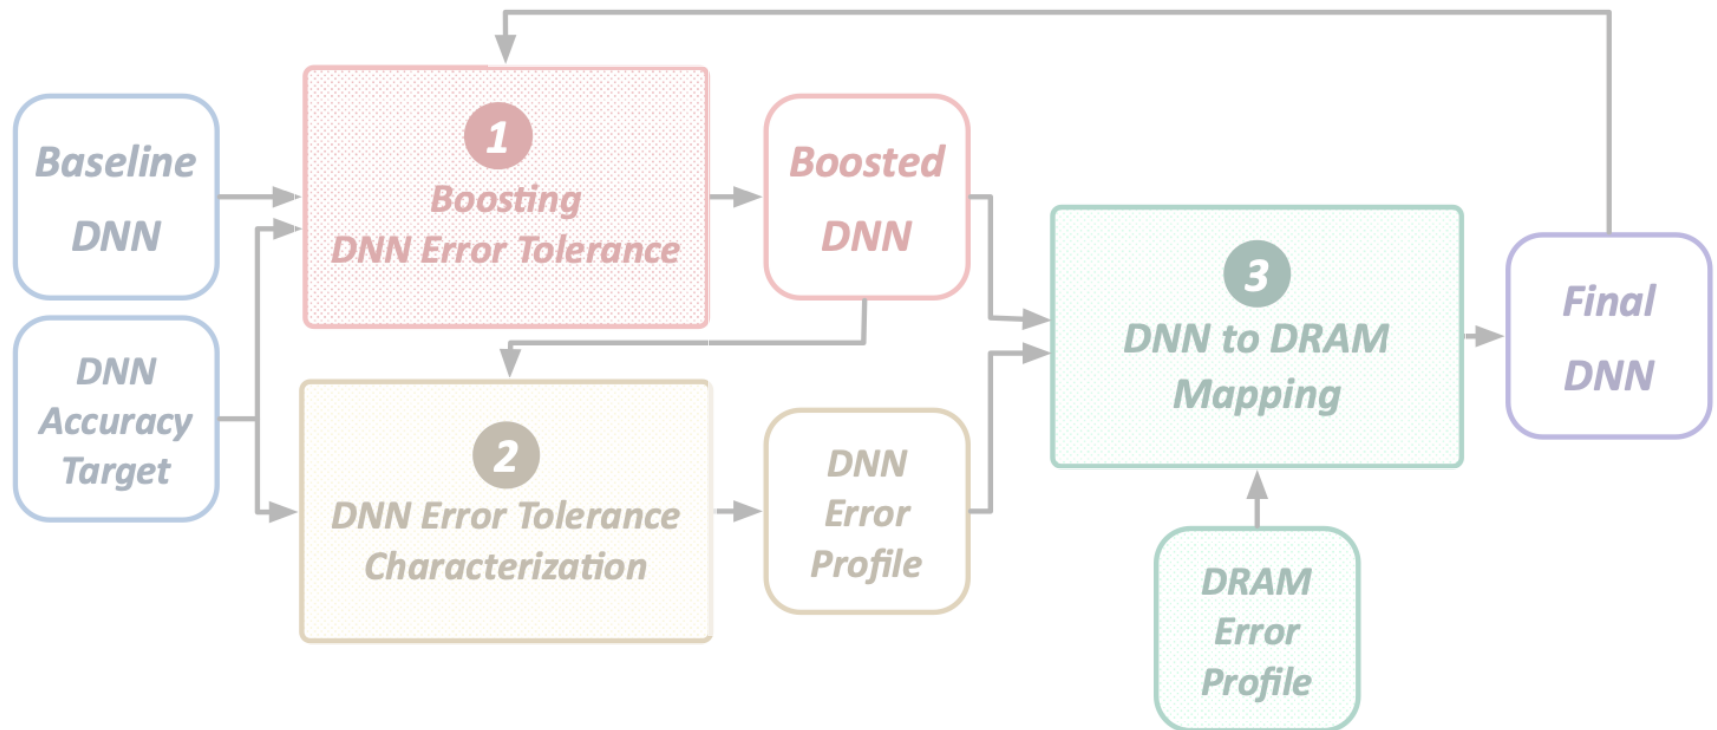
Map more error-tolerant DNN layers
to DRAM partitions with lower voltage/latency

4 DRAM partitions with different error rates

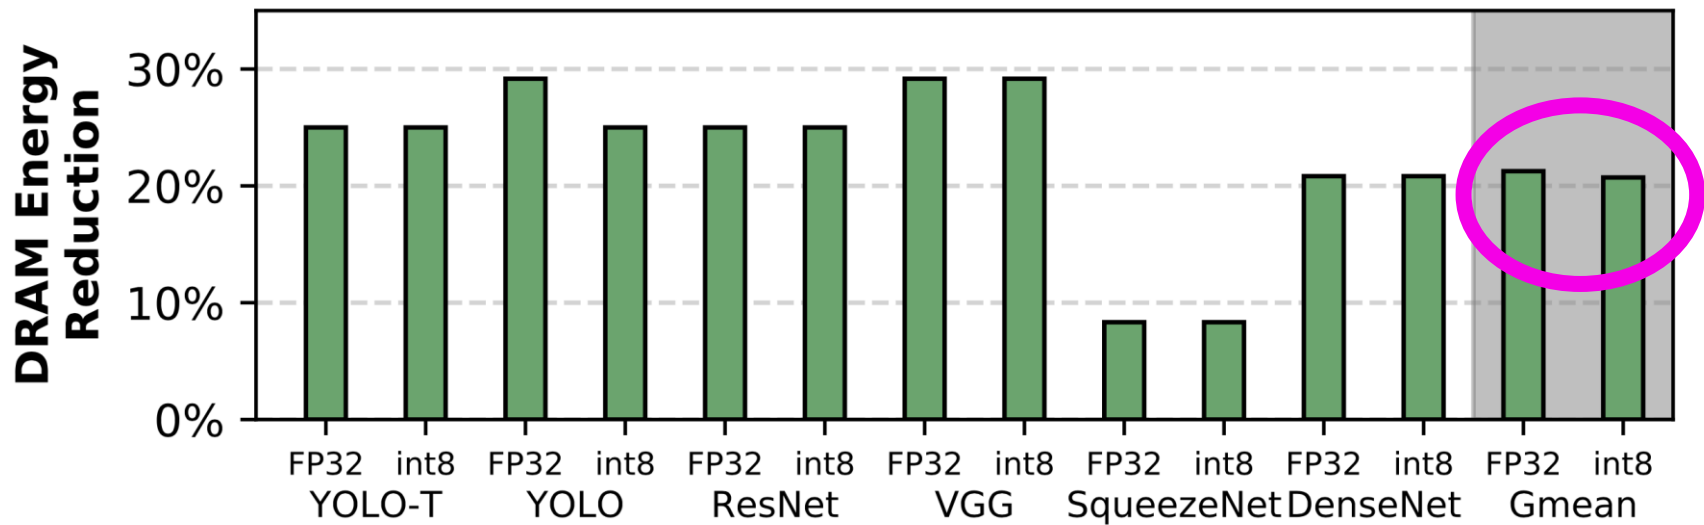
EDEN: Overview

Key idea: Enable **accurate, efficient** DNN inference using **approximate DRAM**

EDEN is an **iterative** process that has **3 key steps**

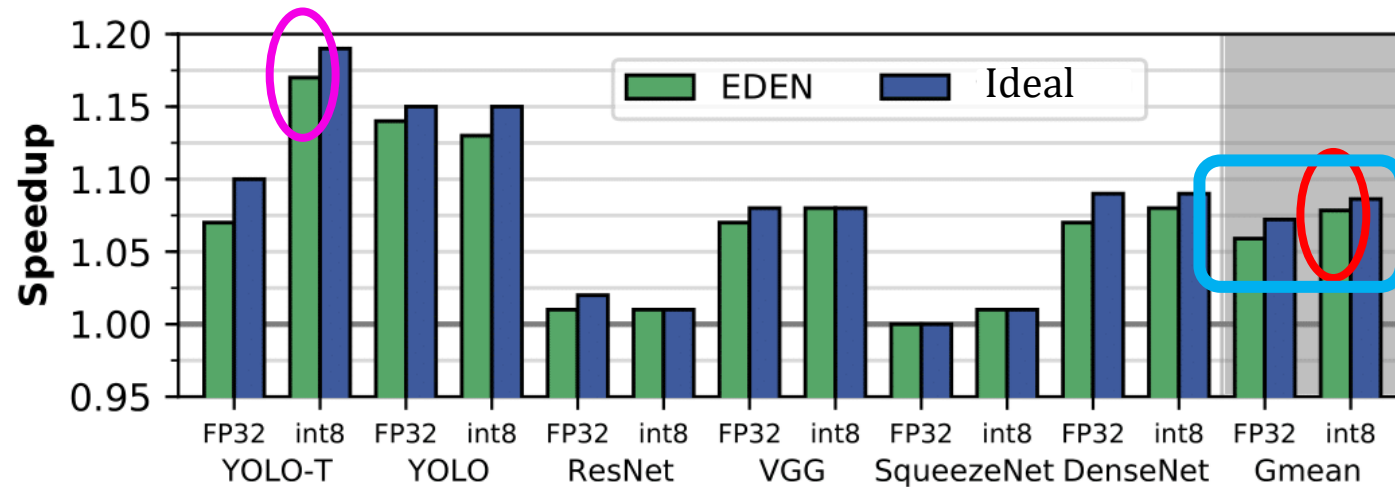


CPU: DRAM Energy Evaluation



Average 21% DRAM energy reduction
maintaining accuracy within 1% of original

CPU: Performance Evaluation



Average 8% system speedup
Some workloads achieve 17% speedup

EDEN achieves **close to the ideal** speedup
possible via tRCD scaling

GPU, Eyeriss, and TPU: Energy Evaluation

- GPU: average **37% energy reduction**
- Eyeriss: average **31% energy reduction**
- TPU: average **32% energy reduction**

EDEN: Data-Aware Efficient DNN Inference

- Skanda Koppula, Lois Orosa, A. Giray Yaglikci, Roknoddin Azizi, Taha Shahroodi, Konstantinos Kanellopoulos, and Onur Mutlu,
"EDEN: Enabling Energy-Efficient, High-Performance Deep Neural Network Inference Using Approximate DRAM"
Proceedings of the 52nd International Symposium on Microarchitecture (MICRO), Columbus, OH, USA, October 2019.
[[Lightning Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Lightning Talk Video](#) (90 seconds)]

EDEN: Enabling Energy-Efficient, High-Performance Deep Neural Network Inference Using Approximate DRAM

Skanda Koppula Lois Orosa A. Giray Yağlıkçı
Roknoddin Azizi Taha Shahroodi Konstantinos Kanellopoulos Onur Mutlu
ETH Zürich

More on EDEN

EDEN: Overview

Key idea: Enabling **accurate, efficient** DNN inference using **approximate DRAM**

EDEN is an **iterative** process that has **3 key steps**

```
graph LR; BaselineDNN[Baseline DNN] --> Step1[1 Boosting DNN Error Tolerance]; DNNAccuracyTarget[DNN Accuracy Target] --> Step1; Step1 --> BoostedDNN[Boosted DNN]; BoostedDNN --> Step2[2 DNN Error Tolerance Characterization]; Step2 --> DNNErrorProfile[DNN Error Profile]; DNNErrorProfile --> Step3[3 DNN to DRAM Mapping]; DRAMErrorProfile[DRAM Error Profile] --> Step3; Step3 --> FinalDNN[Final DNN]; FinalDNN --> Step1;
```

SAFARI

11:33 / 38:02 • EDEN Overview >

ETH ZÜRICH

Computer Architecture - Lecture 11d: EDEN: Reducing Memory Energy in DNNs (ETH Zürich, Fall 2019)

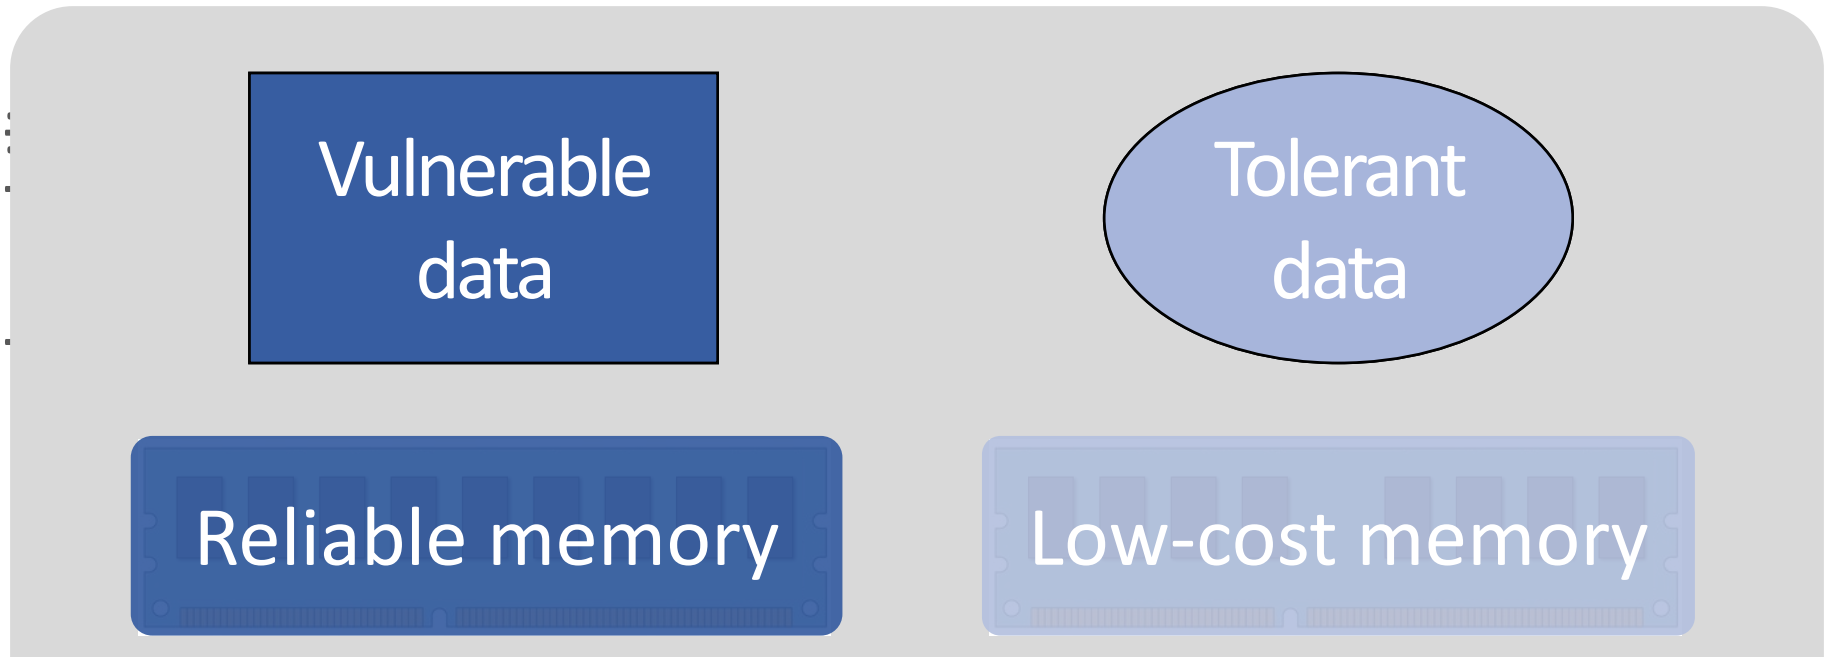
438 views • Oct 31, 2019

Onur Mutlu Lectures
19.7K subscribers

5 0 SHARE + SAVE ...

SUBSCRIBED

Recall: Exploiting Memory Error Tolerance with Hybrid Memory Systems



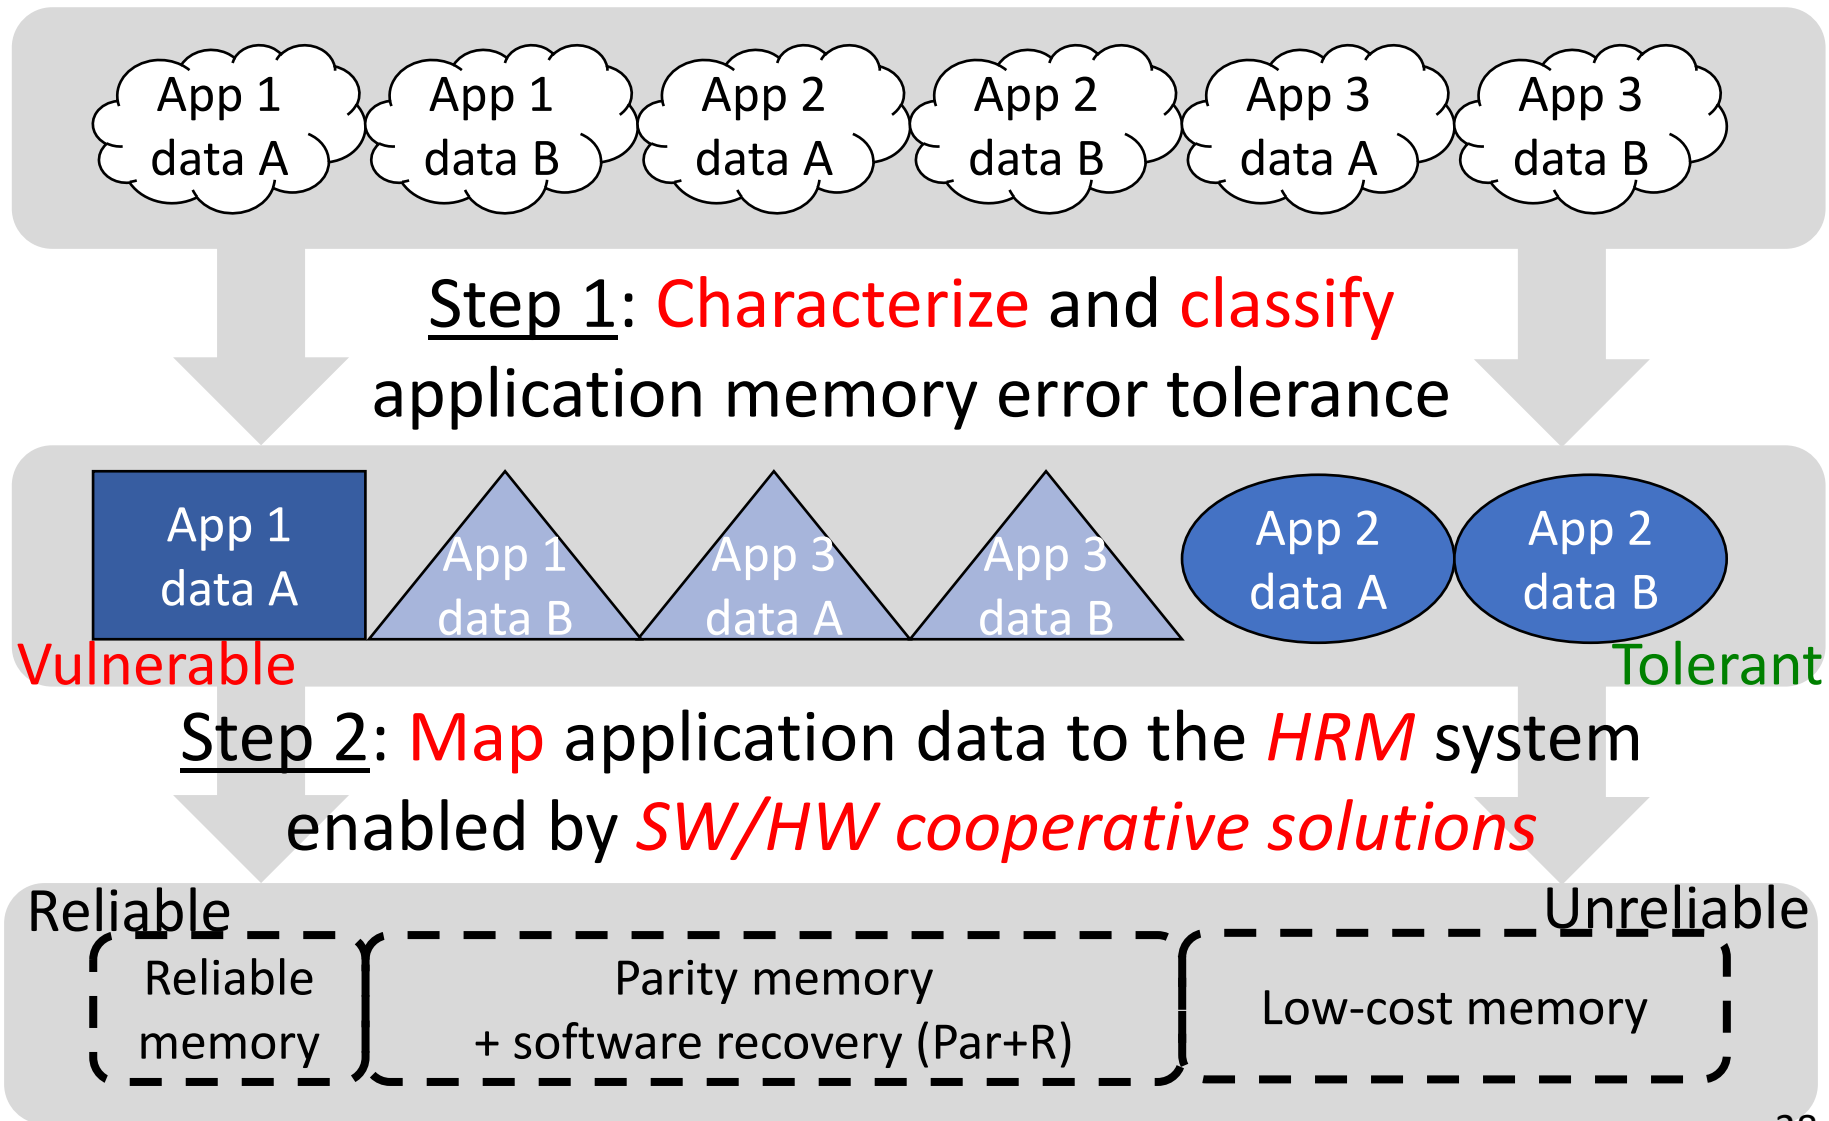
On Microsoft's Web Search workload

Reduces server hardware **cost** by **4.7 %**

Achieves single server **availability** target of **99.90 %**

Heterogeneous-Reliability Memory [DSN 2014]

Recall: Heterogeneous-Reliability Memory



More on Heterogeneous-Reliability Memory

- Yixin Luo, Sriram Govindan, Bikash Sharma, Mark Santaniello, Justin Meza, Aman Kansal, Jie Liu, Badriddine Khessib, Kushagra Vaid, and Onur Mutlu,
"Characterizing Application Memory Error Vulnerability to Optimize Data Center Cost via Heterogeneous-Reliability Memory"
Proceedings of the 44th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN), Atlanta, GA, June 2014. [[Summary](#)]
[[Slides \(pptx\)](#)] [[pdf](#)] [[Coverage on ZDNet](#)]

Characterizing Application Memory Error Vulnerability to Optimize Datacenter Cost via Heterogeneous-Reliability Memory

Yixin Luo Sriram Govindan* Bikash Sharma* Mark Santaniello* Justin Meza
Aman Kansal* Jie Liu* Badriddine Khessib* Kushagra Vaid* Onur Mutlu

Carnegie Mellon University, yixinluo@cs.cmu.edu, {meza, onur}@cmu.edu

*Microsoft Corporation, {srgovin, bsharma, marksan, kansal, jie.liu, bknessib, kvaid}@microsoft.com

Why the Long Memory Latency?

- Reason 1: Design of DRAM Micro-architecture
 - Goal: Maximize capacity/area, not minimize latency
- Reason 2: “One size fits all” approach to latency specification
 - Same latency parameters for all temperatures
 - Same latency parameters for all DRAM chips
 - Same latency parameters for all parts of a DRAM chip
 - Same latency parameters for all supply voltage levels
 - Same latency parameters for all application data
 - ...

Understanding & Exploiting the Voltage-Latency-Reliability Relationship

Analysis of Latency-Voltage in DRAM Chips

- Kevin Chang, A. Giray Yaglikci, Saugata Ghose, Aditya Agrawal, Niladrish Chatterjee, Abhijith Kashyap, Donghyuk Lee, Mike O'Connor, Hasan Hassan, and Onur Mutlu,

"Understanding Reduced-Voltage Operation in Modern DRAM Devices: Experimental Characterization, Analysis, and Mechanisms"

*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Urbana-Champaign, IL, USA, June 2017.*

Understanding Reduced-Voltage Operation in Modern DRAM Chips: Characterization, Analysis, and Mechanisms

Kevin K. Chang[†] Abdullah Giray Yağlıkçı[†] Saugata Ghose[†] Aditya Agrawal[¶] Niladrish Chatterjee[¶]
Abhijith Kashyap[†] Donghyuk Lee[¶] Mike O'Connor^{¶,‡} Hasan Hassan[§] Onur Mutlu^{§,†}

[†]Carnegie Mellon University

[¶]NVIDIA

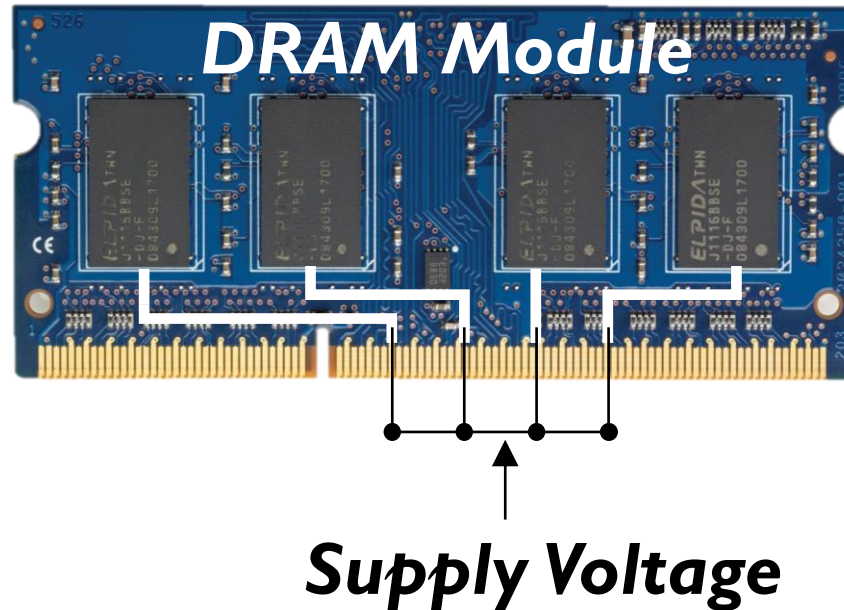
[‡]The University of Texas at Austin

[§]ETH Zürich

Key Questions

- How does reducing voltage affect ***reliability*** (errors)?
- How does reducing voltage affect ***DRAM latency***?
- How do we design a new DRAM energy reduction mechanism?

Supply Voltage Control on DRAM



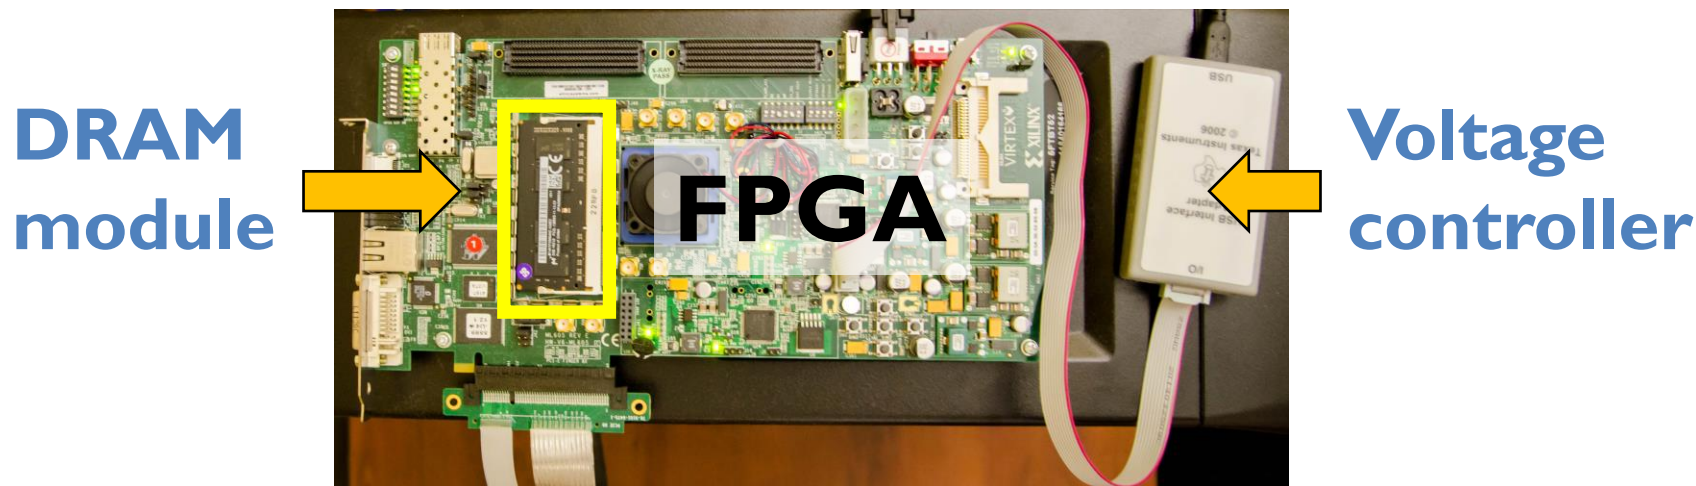
Adjust the *supply voltage* to every chip on the same module

Custom Testing Platform

SoftMC [Hassan+, HPCA'17]: FPGA testing platform to

- 1) Adjust supply voltage to DRAM modules
- 2) Schedule DRAM commands to DRAM modules

Existing systems: DRAM commands not exposed to users

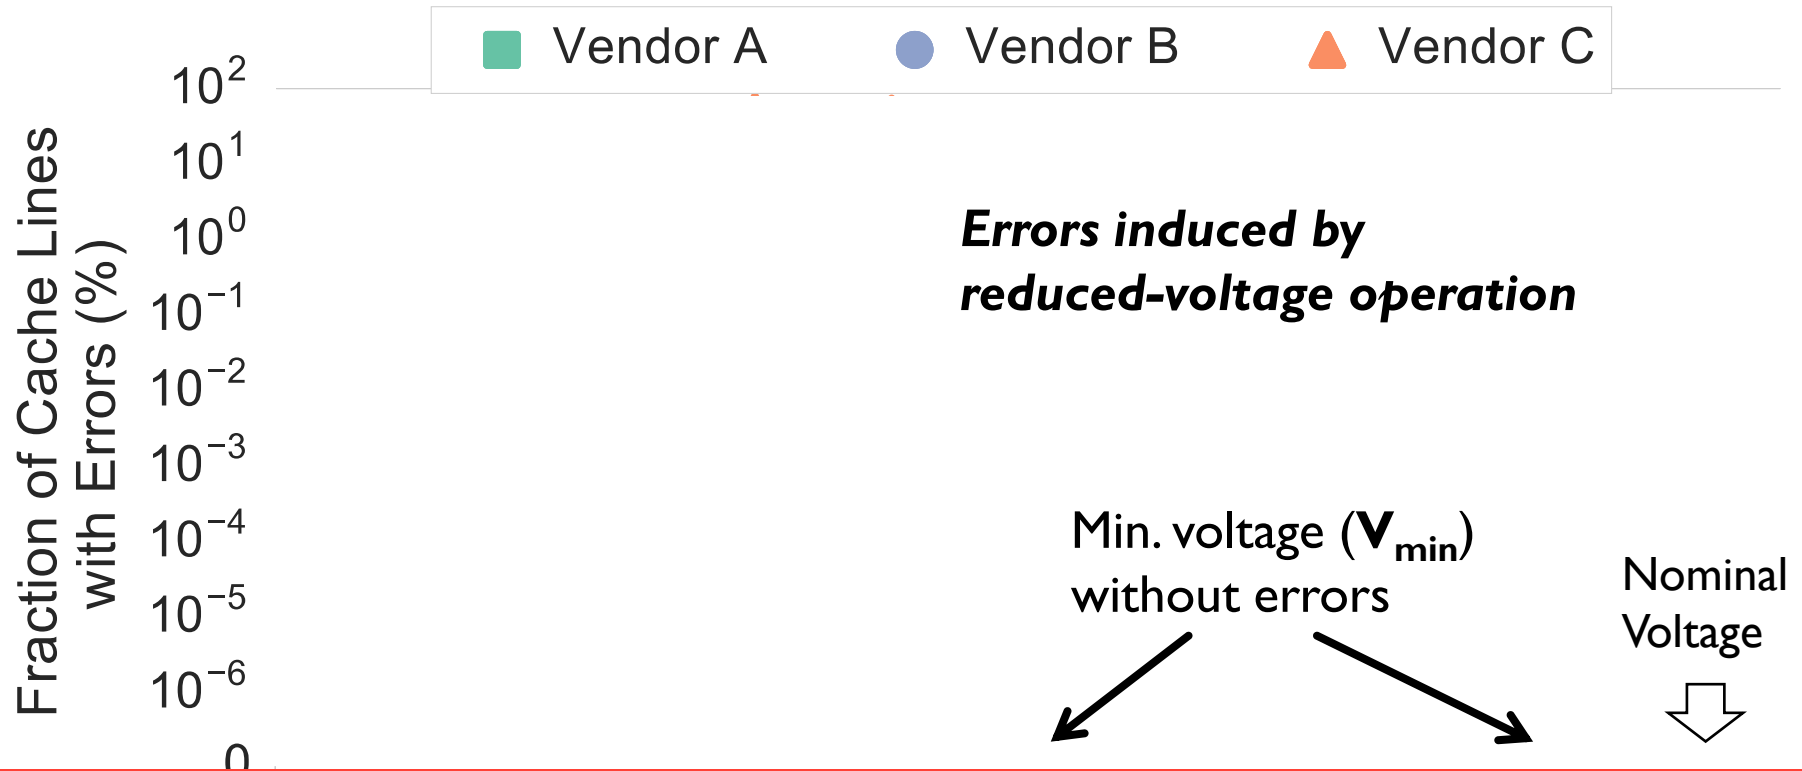


<https://github.com/CMU-SAFARI/DRAM-Voltage-Study>

Tested DRAM Modules

- **124 DDR3L** (low-voltage) DRAM chips
 - **31 SO-DIMMs**
 - **1.35V** (DDR3 uses 1.5V)
 - Density: 4Gb per chip
 - Three major vendors/manufacturers
 - Manufacturing dates: 2014-2016
- Iteratively read every bit in each 4Gb chip under a wide range of supply voltage levels: 1.35V to 1.0V (**-26%**)

Reliability Worsens with Lower Voltage

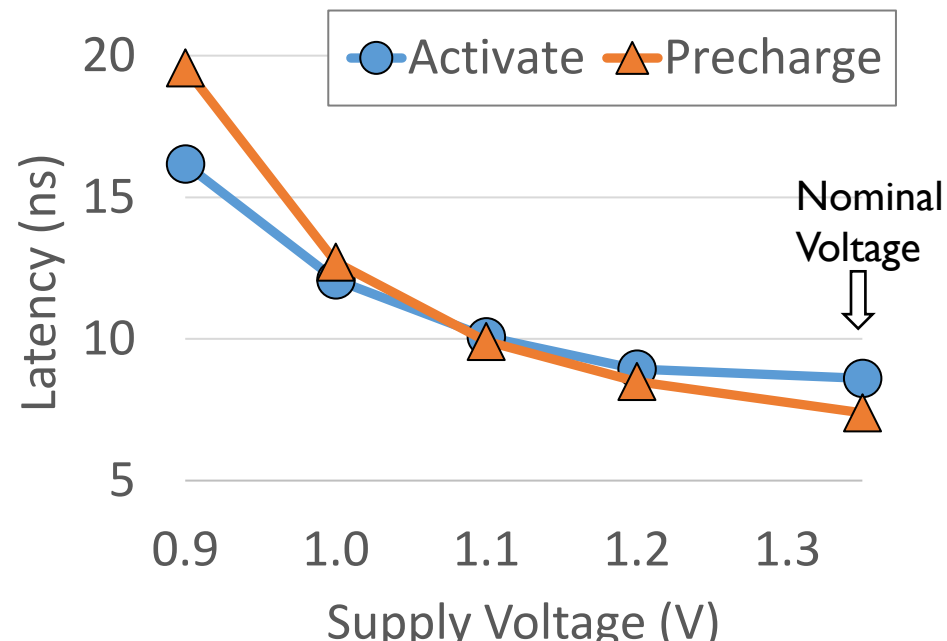
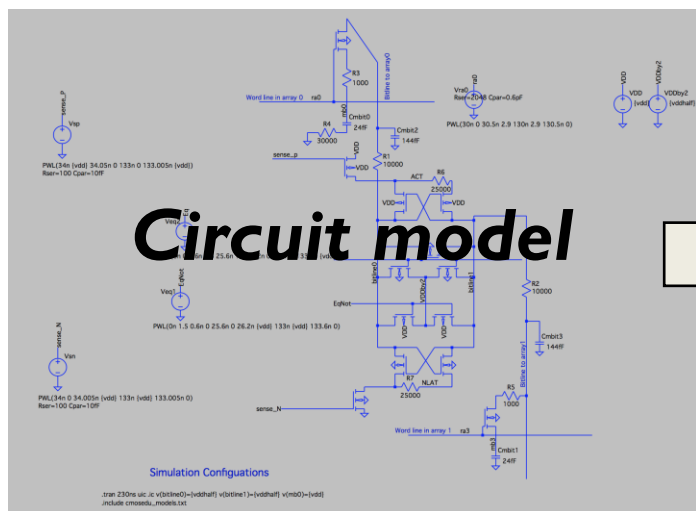


Reducing voltage below V_{\min} causes an increasing number of errors

Source of Errors

Detailed circuit simulations (SPICE) of a DRAM cell array to model the behavior of DRAM operations

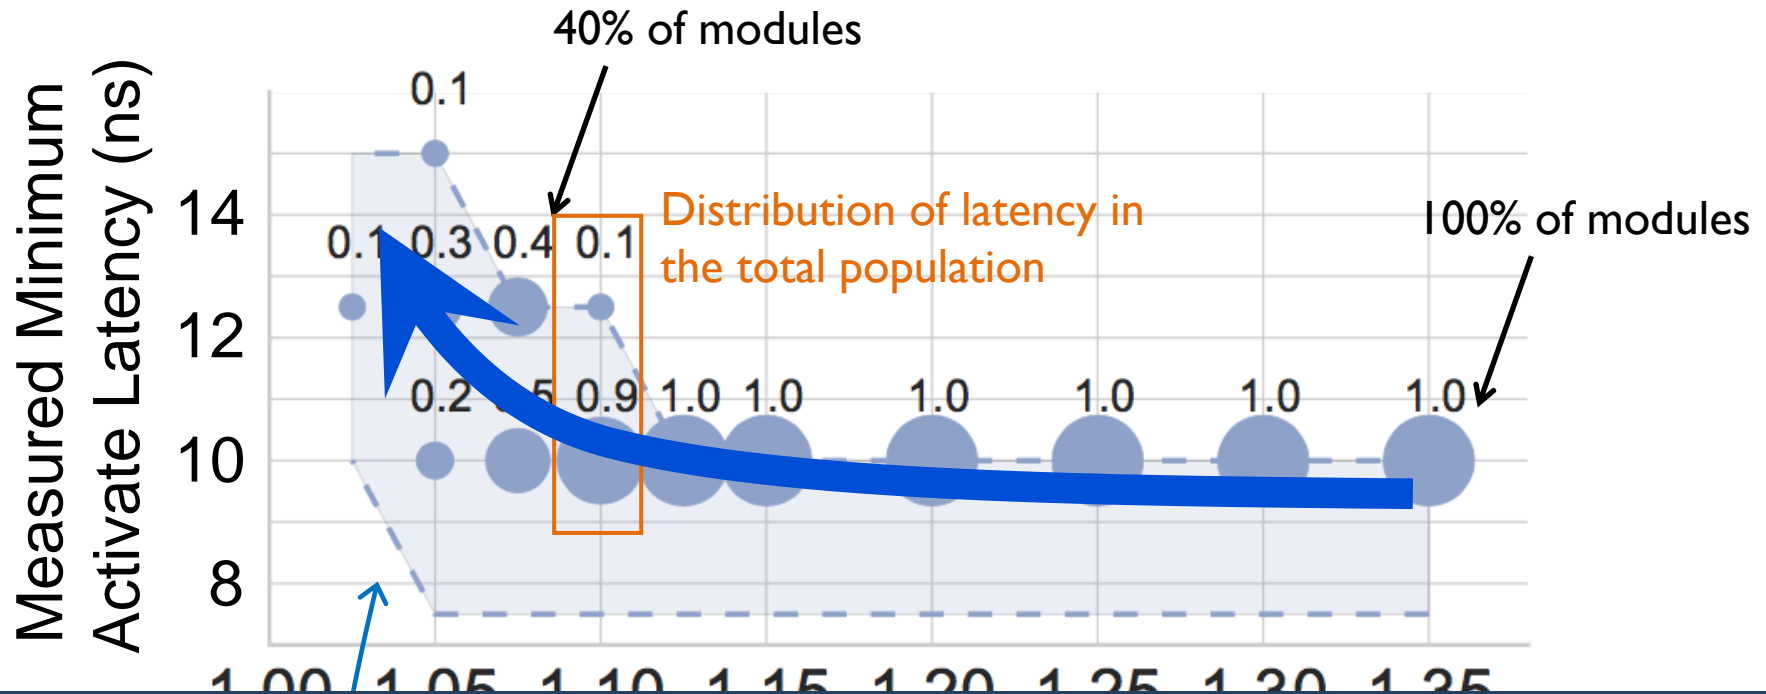
<https://github.com/CMU-SAFARI/DRAM-Voltage-Study>



Reliable low-voltage operation requires higher latency

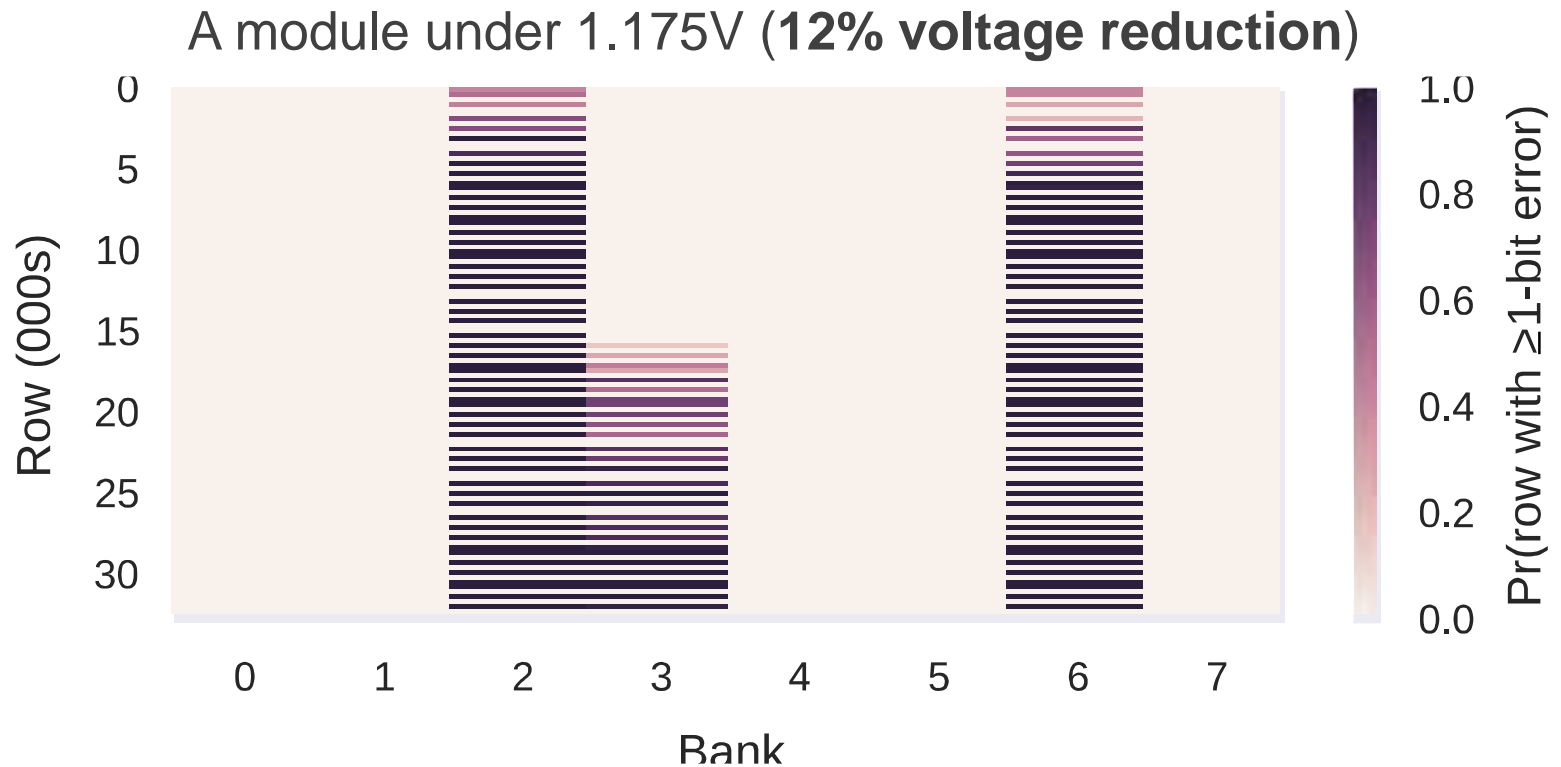
DIMMs Operating at Higher Latency

Measured minimum latency that *does not* cause errors in DRAM modules



DRAM requires longer latency to access data **without errors** at lower voltage

Spatial Locality of Errors



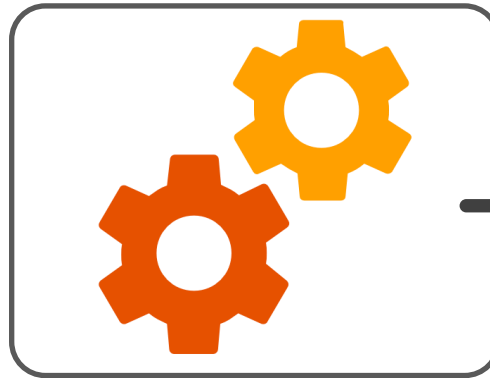
Errors concentrate in certain regions

Voltron Overview

Voltron



User specifies the
performance loss target

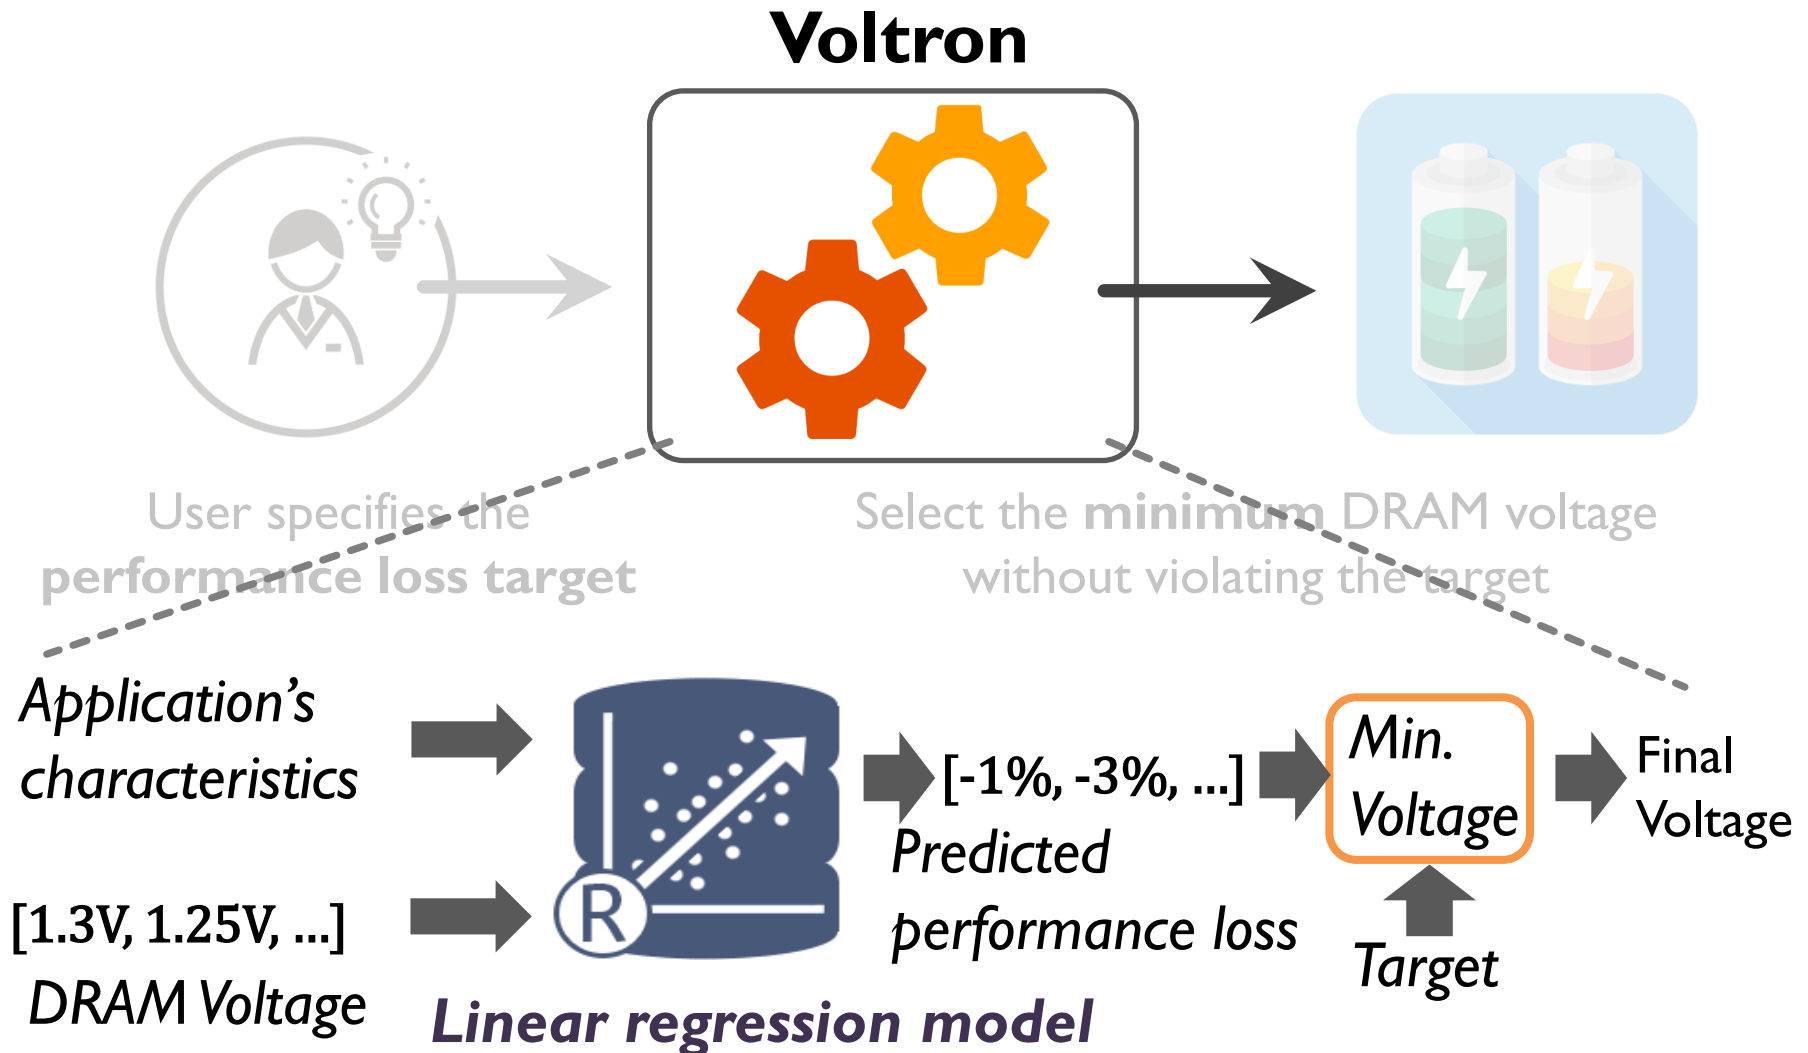


Select the **minimum** DRAM voltage
without violating the target

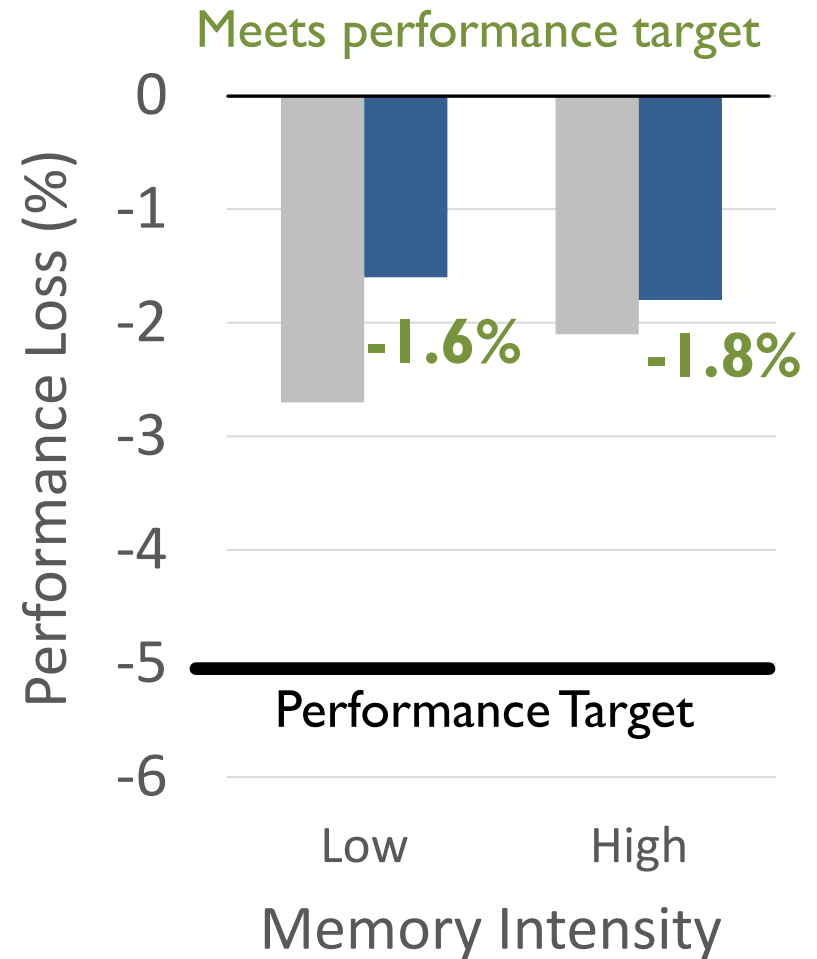
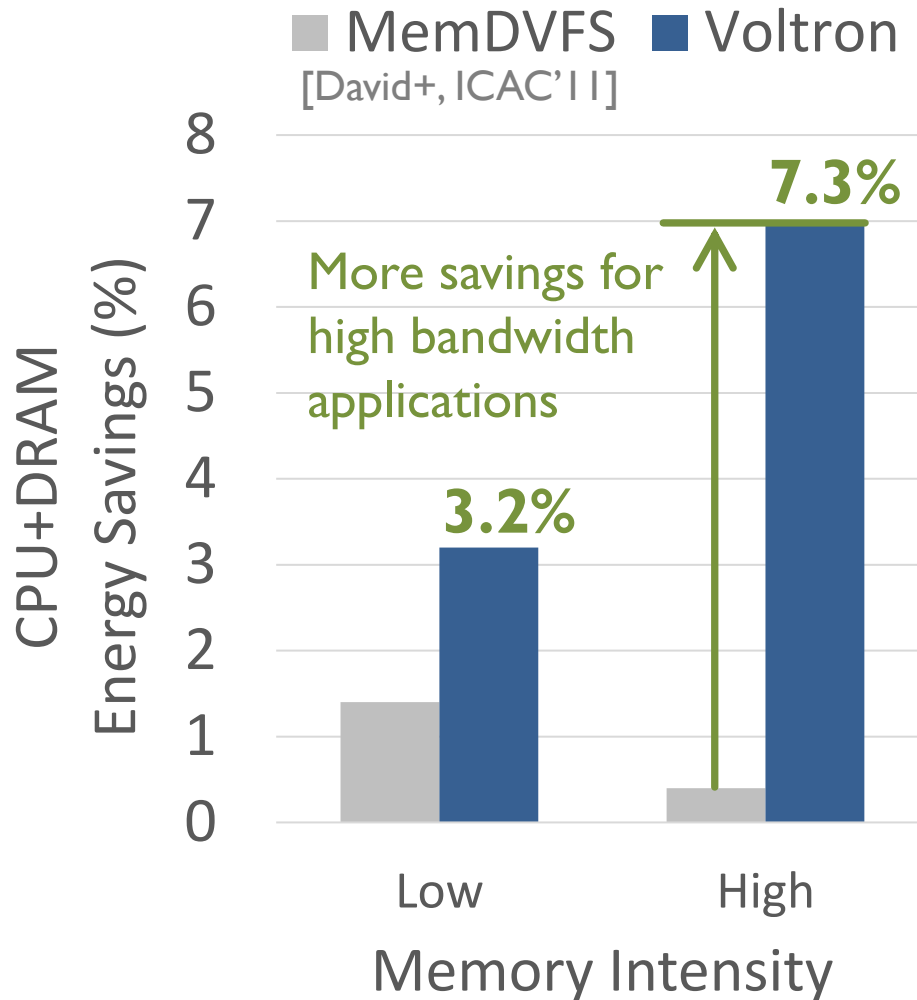


How do we predict performance loss due to increased latency under low DRAM voltage?

Linear Model to Predict Performance



Energy Savings with Bounded Performance



Voltron: Advantages & Disadvantages

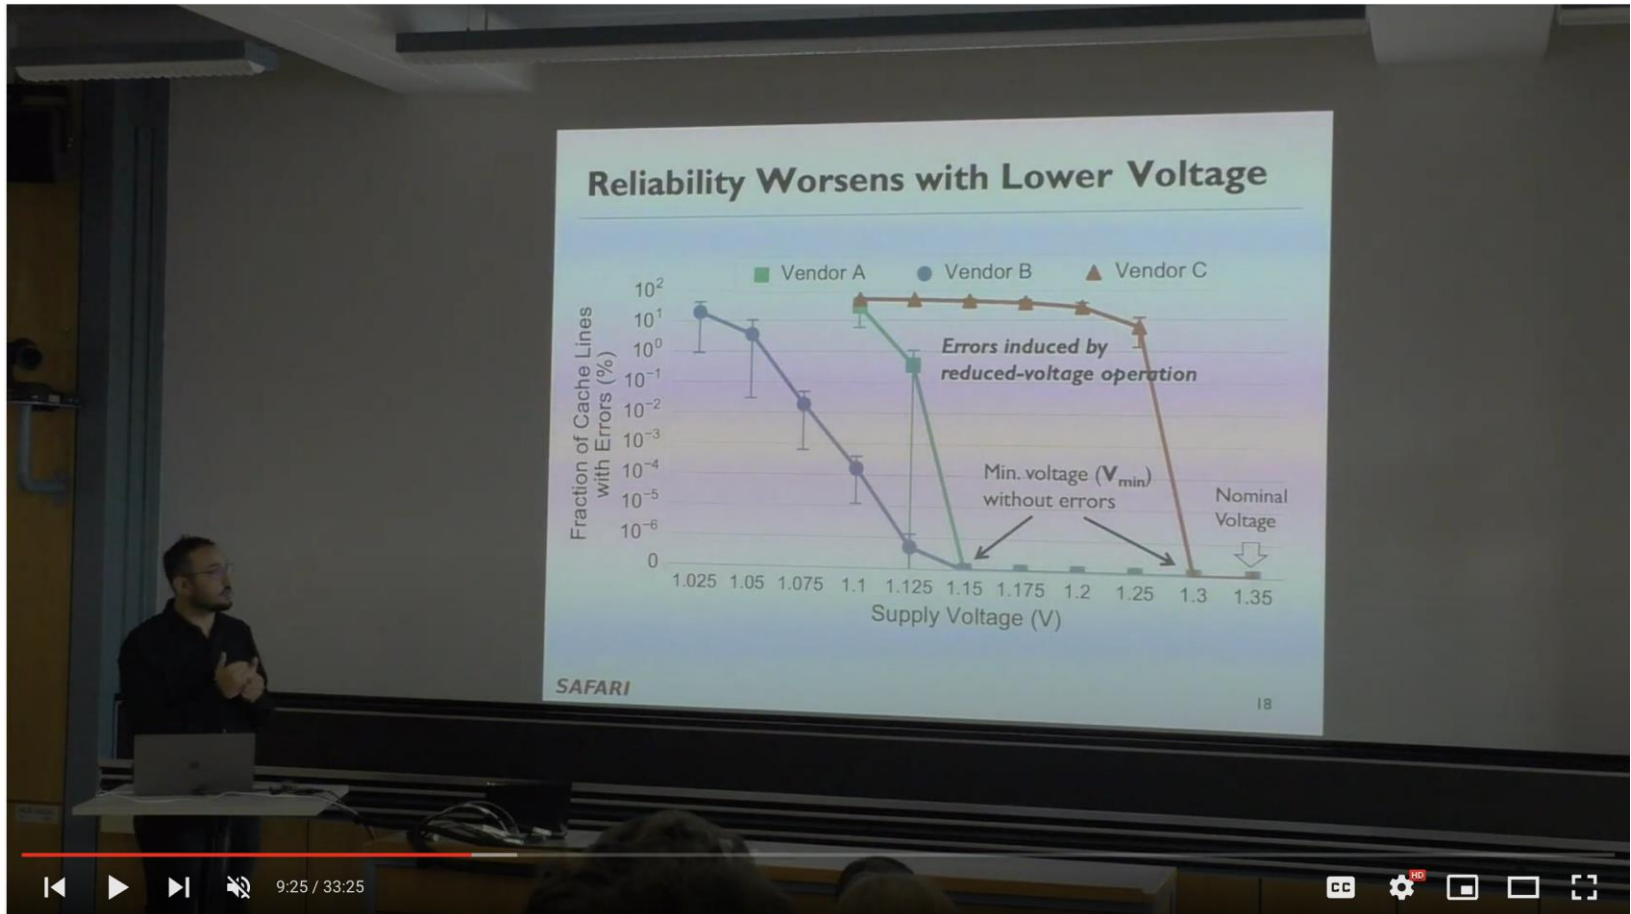
■ Advantages

- + Can trade-off between voltage and latency to improve energy or performance
- + Can exploit the high voltage margin present in DRAM

■ Disadvantages

- Requires finding the reliable operating voltage for each chip → higher testing cost
- More complicated memory controller

More on Voltron



ETH ZÜRICH

Computer Architecture - Lecture 11c: Voltron: Reducing DRAM Energy (ETH Zürich, Fall 2019)

409 views • Oct 31, 2019

7 0 SHARE SAVE ...



Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED



Reducing Memory Latency to Support Security Primitives

Using Memory for Security

- Generating True Random Numbers (using DRAM)
 - Kim et al., HPCA 2019
 - Olgun et al., ISCA 2021
- Evaluating Physically Unclonable Functions (using DRAM)
 - Kim et al., HPCA 2018
- Quickly Destroying In-Memory Data (using DRAM)
 - Orosa et al., arxiv 2019 + ISCA 2021

DRAM Latency PUFs

- Jeremie S. Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu,
"The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern DRAM Devices"
Proceedings of the 24th International Symposium on High-Performance Computer Architecture (HPCA), Vienna, Austria, February 2018.
[[Lightning Talk Video](#)]
[[Slides \(pptx\)](#)] [[pdf](#)] [[Lightning Session Slides \(pptx\)](#)] [[pdf](#)]
[[Full Talk Lecture Video](#) (28 minutes)]

The DRAM Latency PUF:

Quickly Evaluating Physical Unclonable Functions

by Exploiting the Latency-Reliability Tradeoff in Modern Commodity DRAM Devices

Jeremie S. Kim^{†§}

Minesh Patel[§]

Hasan Hassan[§]

Onur Mutlu^{§†}

[†]Carnegie Mellon University

[§]ETH Zürich

More on DRAM Latency PUFs

DRAM Latency PUF Key Idea

- A cell's latency failure probability is inherently related to **random process variation** from manufacturing
- We can provide **repeatable and unique device signatures** using latency error patterns

High % chance to fail with reduced t_{RCD}

Row Decoder

SA SA SA SA SA SA SA

SAFARI

19/45

ETH ZÜRICH

Computer Architecture - Lecture 11a: DRAM Latency PUF (ETH Zürich, Fall 2019)

449 views • Oct 31, 2019

6 0 SHARE SAVE ...



Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED



DRAM Latency True Random Number Generator

- Jeremie S. Kim, Minesh Patel, Hasan Hassan, Lois Orosa, and Onur Mutlu,
"D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput"
Proceedings of the 25th International Symposium on High-Performance Computer Architecture (HPCA), Washington, DC, USA, February 2019.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Full Talk Video](#) (21 minutes)]
[[Full Talk Lecture Video](#) (27 minutes)]
Top Picks Honorable Mention by IEEE Micro.

D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput

Jeremie S. Kim^{‡§}

Minesh Patel[§]

Hasan Hassan[§]

Lois Orosa[§]

Onur Mutlu^{§‡}

[‡]Carnegie Mellon University

[§]ETH Zürich

D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput

Jeremie S. Kim Minesh Patel

Hasan Hassan Lois Orosa Onur Mutlu

HPCA 2019

SAFARI

ETH zürich

Carnegie Mellon

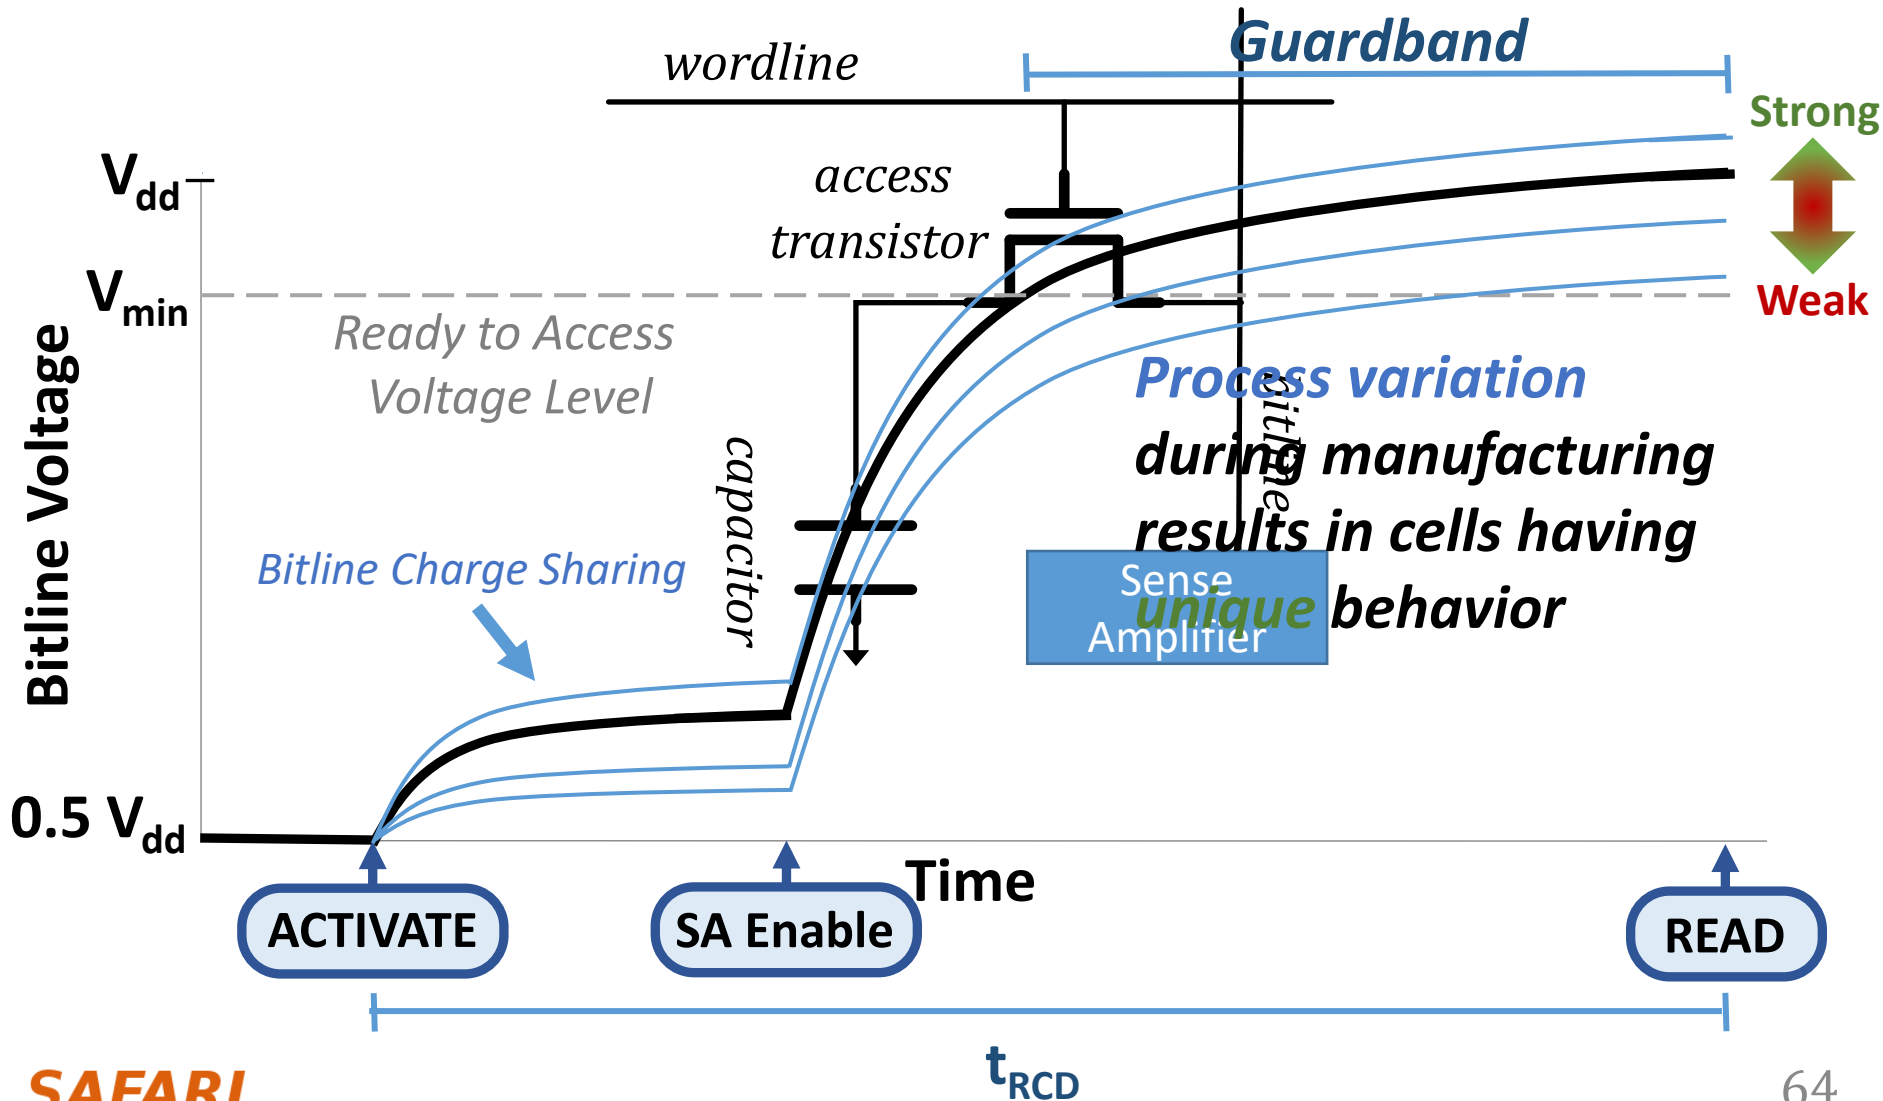
D-RaNGe Executive Summary

- **Motivation**: High-throughput true random numbers enable system security and various randomized algorithms.
 - Many systems (e.g., IoT, mobile, embedded) do not have dedicated **True Random Number Generator (TRNG)** hardware but have DRAM devices
- **Problem**: Current DRAM-based TRNGs either
 1. do **not** sample a fundamentally non-deterministic entropy source
 2. are **too slow** for continuous high-throughput operation
- **Goal**: A novel and effective TRNG that uses **existing** commodity DRAM to provide random values with 1) **high-throughput**, 2) **low latency** and 3) no adverse effect on concurrently running applications
- **D-RaNGe**: Reduce DRAM access latency **below reliable values** and exploit DRAM cells' failure probabilities to generate random values
- **Evaluation**:
 1. Experimentally characterize **282 real LPDDR4 DRAM devices**
 2. **D-RaNGe (717.4 Mb/s)** has significantly higher throughput (**211x**)
 3. **D-RaNGe (100ns)** has significantly lower latency (**180x**)

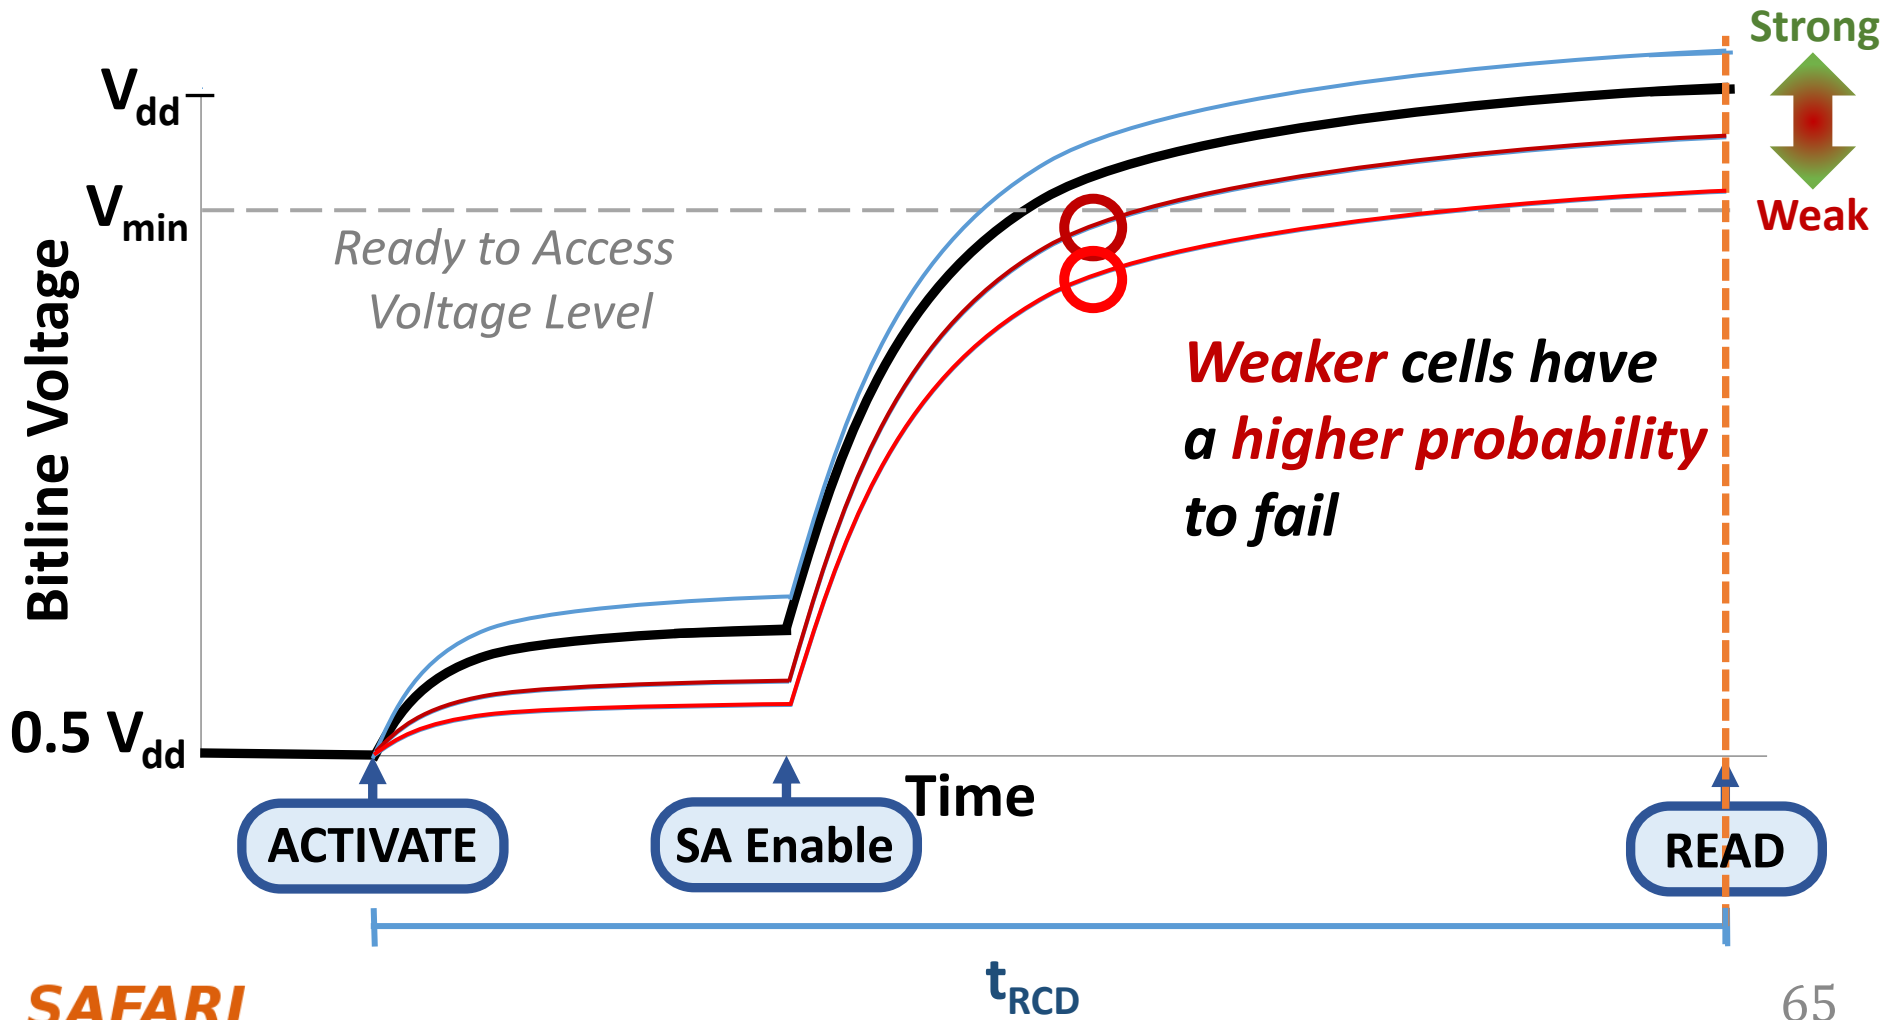
DRAM Latency Characterization of 282 LPDDR4 DRAM Devices

- Latency failures come from accessing DRAM with **reduced** timing parameters.
- **Key Observations:**
 1. A cell's **latency failure** probability is determined by **random process variation**
 2. Some cells fail **randomly**

DRAM Accesses and Failures



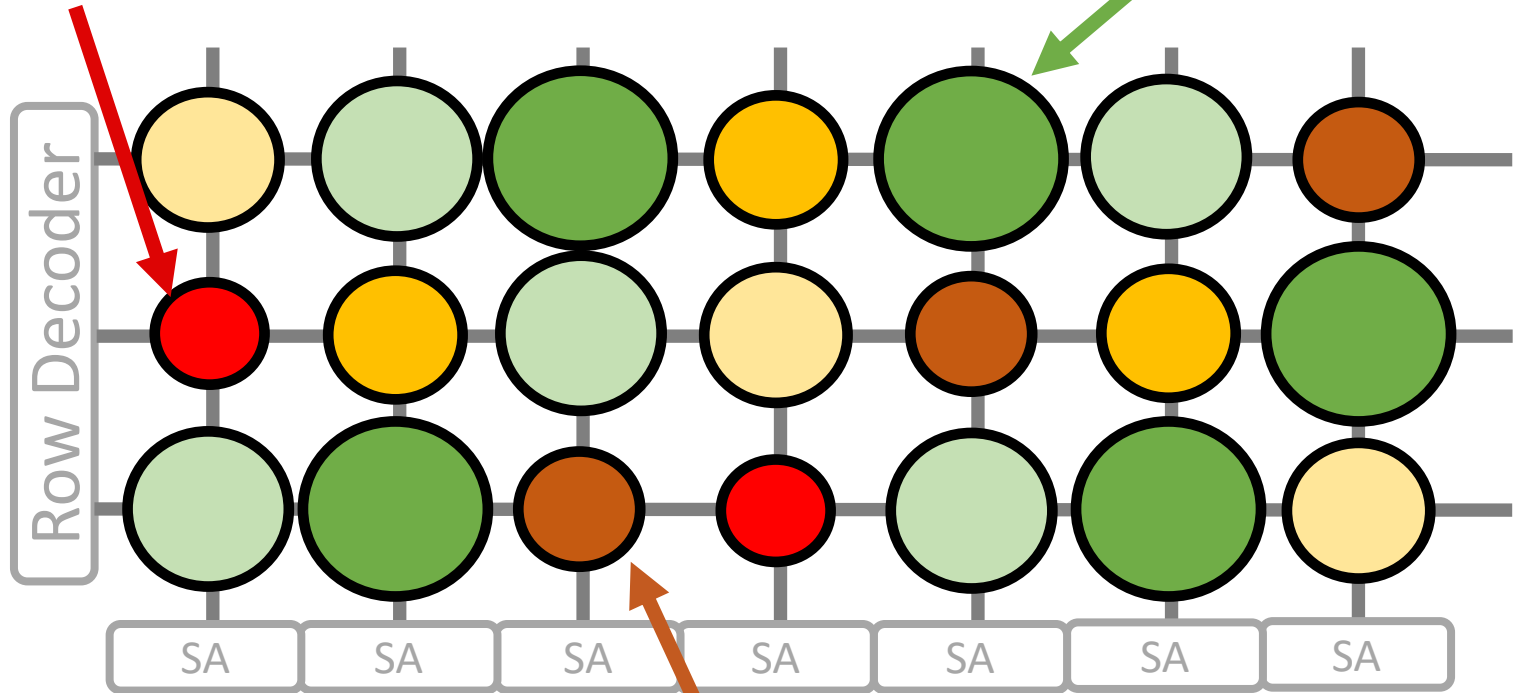
DRAM Accesses and Failures



D-RaNGe Key Idea

High % chance to fail
with reduced t_{RCD}

Low % chance to fail
with reduced t_{RCD}



Fails randomly
with reduced t_{RCD}

D-RaNGe Key Idea

High % chance to fail
with reduced t_{RCD}

Low % chance to fail
with reduced t_{RCD}

**We refer to cells that fail randomly
when accessed with a reduced t_{RCD}
as RNG cells**



Fails randomly
with reduced t_{RCD}

Our D-RaNGe Evaluation

- We generate **random values** by repeatedly accessing **RNG cells** and aggregating the data read
- The random data satisfies the NIST statistical test suite for randomness
- The **D-RaNGE** generates random numbers
 - **Throughput:** 717.4 Mb/s
 - **Latency:** 64 bits in <1us
 - **Power:** 4.4 nJ/bit

D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput

Jeremie S. Kim Minesh Patel

Hasan Hassan Lois Orosa Onur Mutlu

SAFARI

HPCA 2019

ETH zürich

Carnegie Mellon

More on D-RaNGe

- Jeremie S. Kim, Minesh Patel, Hasan Hassan, Lois Orosa, and Onur Mutlu,
"D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput"
Proceedings of the 25th International Symposium on High-Performance Computer Architecture (HPCA), Washington, DC, USA, February 2019.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Full Talk Video](#) (21 minutes)]
[[Full Talk Lecture Video](#) (27 minutes)]
Top Picks Honorable Mention by IEEE Micro.

D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput

Jeremie S. Kim^{‡§}

Minesh Patel[§]

Hasan Hassan[§]

Lois Orosa[§]

Onur Mutlu^{§‡}

[‡]Carnegie Mellon University

[§]ETH Zürich

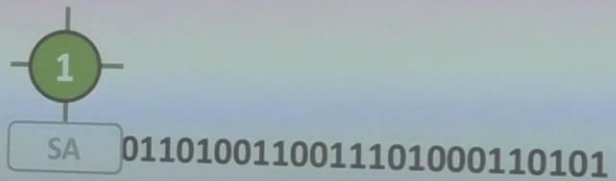
More on DRAM Latency TRNGs

D-RaNGe: Extracting Random Values

Identify all DRAM cells that fail randomly when accessed with a reduced t_{RCD} (**RNG Cell**)

- When accessing an RNG Cell with a reduced t_{RCD} , the values read will be truly random values

RNG Cell



SA 0110100110011101000110101

SAFARI 17/48

8:51 / 27:00 • D-RaNGe Key Idea >

ETH ZÜRICH

Computer Architecture - Lecture 11b: D-RaNGe: True Random Number Generation (ETH Zürich, Fall 2019)

449 views • Oct 31, 2019

Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED

In-DRAM True Random Number Generation

- Ataberk Olgun, Minesh Patel, A. Giray Yaglikci, Haocong Luo, Jeremie S. Kim, F. Nisa Bostanci, Nandita Vijaykumar, Oguz Ergin, and Onur Mutlu,
"QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAM Chips"
Proceedings of the 48th International Symposium on Computer Architecture (ISCA), Virtual, June 2021.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Short Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (25 minutes)]
[[SAFARI Live Seminar Video](#) (1 hr 26 mins)]

QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAM Chips

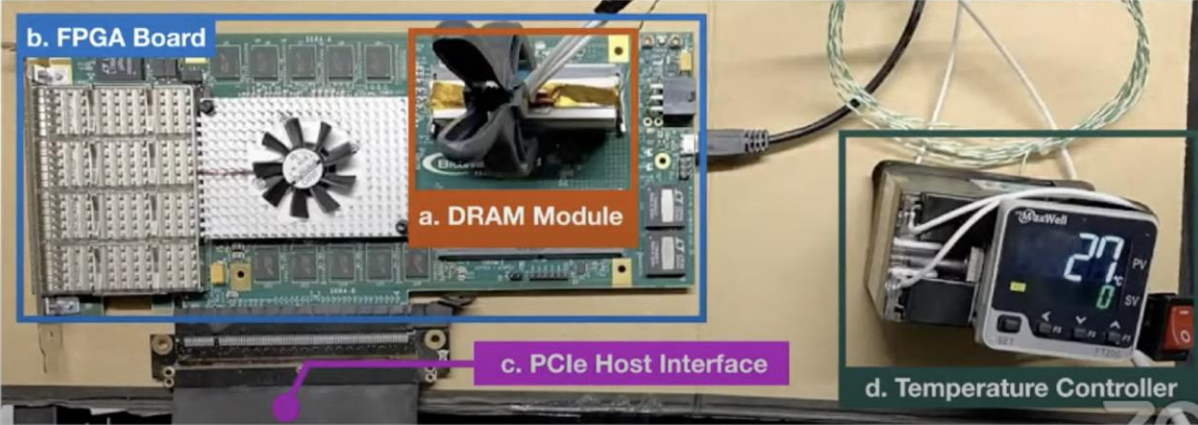
Ataberk Olgun^{§†} Minesh Patel[§] A. Giray Yağlıkçı[§] Haocong Luo[§]
Jeremie S. Kim[§] F. Nisa Bostancı^{§†} Nandita Vijaykumar^{§⊙} Oğuz Ergin[†] Onur Mutlu[§]
[§]ETH Zürich [†]TOBB University of Economics and Technology [⊙]University of Toronto

More on QUAC-TRNG

Real Chip Characterization

Experimentally study QUAC and QUAC-TRNG using 136 real DDR4 chips from SK Hynix

DDR4 SoftMC → DRAM Testing Infrastructure



Ataberk Olgun...

zoom

SAFARI kasirga [Hassan+ HPCA'17] <https://github.com/CMU-SAFARI/SoftMC> CC BY 3.0

SAFARI Live Seminar: High-Throughput TRNG Using Quadruple Row Activation in Commodity DRAM Chips

713 views • Streamed live on Sep 15, 2021

👍 27 💬 0 ➦ SHARE ⚙️ SAVE ...



Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED



Reducing Refresh Latency

Reducing Refresh Latency

- Anup Das, Hasan Hassan, and Onur Mutlu,
"VRL-DRAM: Improving DRAM Performance via Variable Refresh Latency"
*Proceedings of the 55th Design Automation Conference (**DAC**), San Francisco, CA, USA, June 2018.*

VRL-DRAM: Improving DRAM Performance via Variable Refresh Latency

Anup Das
Drexel University
Philadelphia, PA, USA
anup.das@drexel.edu

Hasan Hassan
ETH Zürich
Zürich, Switzerland
hhasan@ethz.ch

Onur Mutlu
ETH Zürich
Zürich, Switzerland
omutlu@gmail.com

Reducing Memory Latency by Exploiting Memory Access Patterns

ChargeCache: Exploiting Access Patterns

- Hasan Hassan, Gennady Pekhimenko, Nandita Vijaykumar, Vivek Seshadri, Donghyuk Lee, Oguz Ergin, and Onur Mutlu,
"ChargeCache: Reducing DRAM Latency by Exploiting Row Access Locality"
Proceedings of the 22nd International Symposium on High-Performance Computer Architecture (HPCA), Barcelona, Spain, March 2016.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Source Code](#)]

ChargeCache: Reducing DRAM Latency by Exploiting Row Access Locality

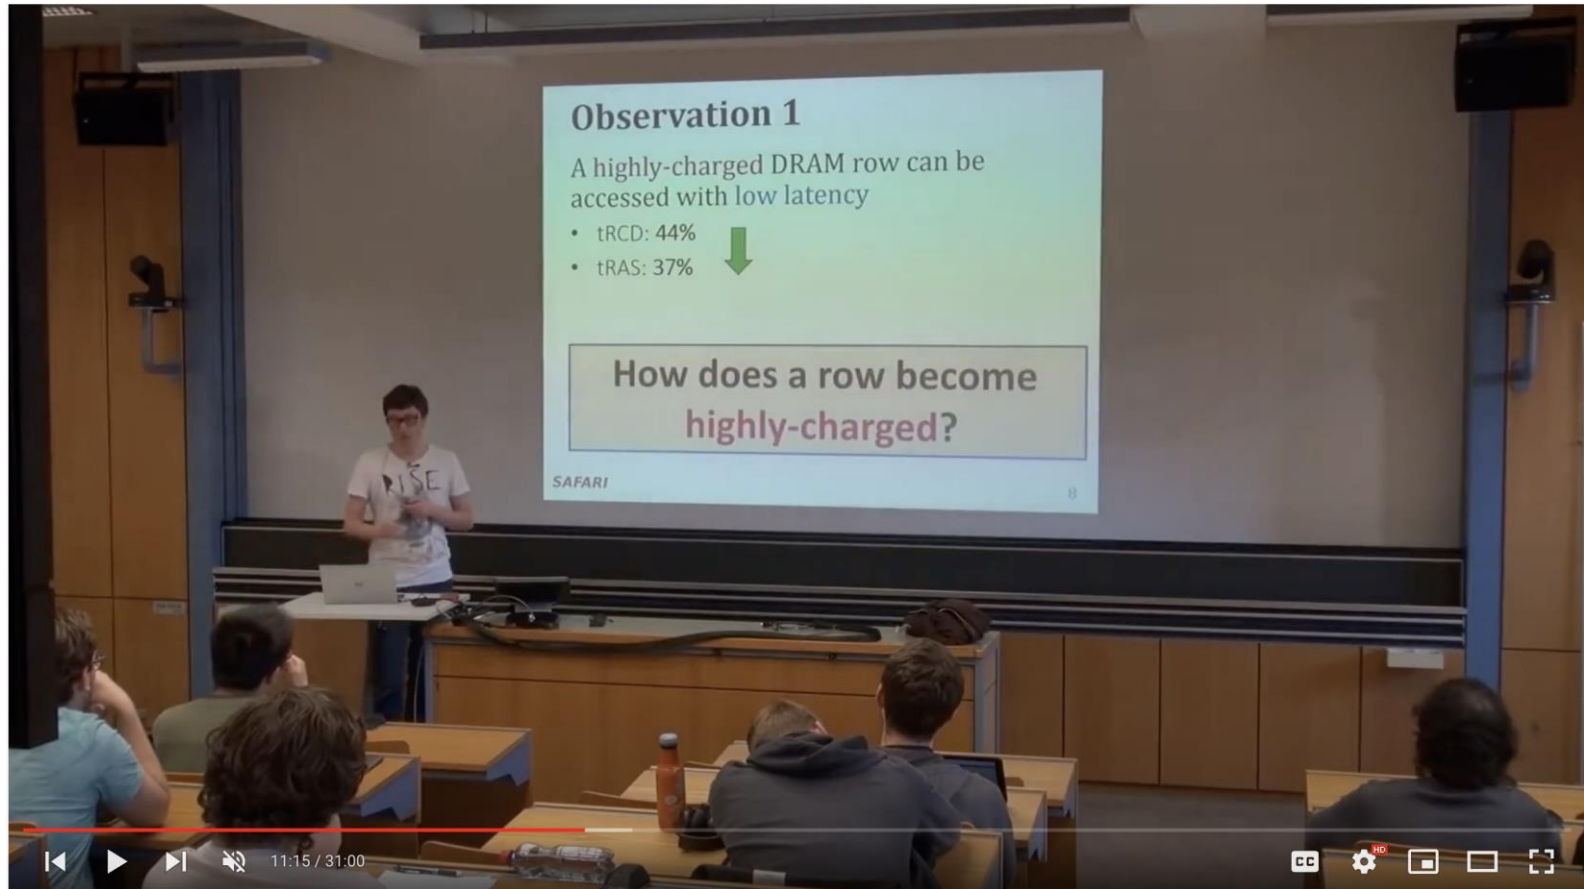
Hasan Hassan^{†*}, Gennady Pekhimenko[†], Nandita Vijaykumar[†]
Vivek Seshadri[†], Donghyuk Lee[†], Oguz Ergin^{*}, Onur Mutlu[†]

ChargeCache: Executive Summary

- **Goal**: Reduce average DRAM access latency with no modification to the existing DRAM chips
- **Observations**:
 - 1) A highly-charged DRAM row can be accessed with low latency
 - 2) A row's charge is restored when the row is accessed
 - 3) A recently-accessed row is likely to be accessed again:

Row Level Temporal Locality (RLTL)
- **Key Idea**: Track recently-accessed DRAM rows and use lower timing parameters if such rows are accessed again
- **ChargeCache**:
 - Low cost & no modifications to the DRAM
 - Higher performance (**8.6-10.6%** on average for 8-core)
 - Lower DRAM energy (**7.9%** on average)

More on ChargeCache



ETH ZÜRICH HAUPTGEBÄUDE

Computer Architecture - Lecture 6a: ChargeCache: Reducing DRAM Latency (ETH Zürich, Fall 2018)

519 views • Oct 10, 2018

9 0 SHARE SAVE ...



Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED



Partial Restoration of Cell Charge

- Yaohua Wang, Arash Tavakkol, Lois Orosa, Saugata Ghose, Nika Mansouri Ghiasi, Minesh Patel, Jeremie S. Kim, Hasan Hassan, Mohammad Sadrosadati, and Onur Mutlu,
"Reducing DRAM Latency via Charge-Level-Aware Look-Ahead Partial Restoration"
Proceedings of the 51st International Symposium on Microarchitecture (MICRO), Fukuoka, Japan, October 2018.

Reducing DRAM Latency via Charge-Level-Aware Look-Ahead Partial Restoration

Yaohua Wang^{†§} Arash Tavakkol[†] Lois Orosa^{†*} Saugata Ghose[‡] Nika Mansouri Ghiasi[†]
Minesh Patel[†] Jeremie S. Kim^{‡†} Hasan Hassan[†] Mohammad Sadrosadati[†] Onur Mutlu^{‡†}

[†]*ETH Zürich* [§]*National University of Defense Technology*

[‡]*Carnegie Mellon University* ^{*}*University of Campinas*

Parallelizing Refreshes and Accesses

- Kevin Chang, Donghyuk Lee, Zeshan Chishti, Alaa Alameldeen, Chris Wilkerson, Yoongu Kim, and Onur Mutlu,
"Improving DRAM Performance by Parallelizing Refreshes with Accesses"
Proceedings of the 20th International Symposium on High-Performance Computer Architecture (HPCA), Orlando, FL, February 2014.
[[Summary](#)] [[Slides \(pptx\)](#)] [[pdf](#)]

Reducing Performance Impact of DRAM Refresh by Parallelizing Refreshes with Accesses

Kevin Kai-Wei Chang Donghyuk Lee Zeshan Chishti[†]

Alaa R. Alameldeen[†] Chris Wilkerson[†] Yoongu Kim Onur Mutlu

Carnegie Mellon University [†]Intel Labs

On DRAM Power Consumption

VAMPIRE DRAM Power Model

- Saugata Ghose, A. Giray Yaglikci, Raghav Gupta, Donghyuk Lee, Kais Kudrolli, William X. Liu, Hasan Hassan, Kevin K. Chang, Niladrish Chatterjee, Aditya Agrawal, Mike O'Connor, and Onur Mutlu,

"What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study"

*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Irvine, CA, USA, June 2018.*

[[Abstract](#)]

[[POMACS Journal Version \(same content, different format\)](#)]

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

[[VAMPIRE DRAM Power Model](#)]

What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study

Saugata Ghose [†]	Abdullah Giray Yağlıkçı ^{‡†}	Raghav Gupta [†]	Donghyuk Lee [§]
Kais Kudrolli [†]	William X. Liu [†]	Hasan Hassan [‡]	Kevin K. Chang [†]
Niladrish Chatterjee [§]	Aditya Agrawal [§]	Mike O'Connor ^{§¶}	Onur Mutlu ^{‡†}

[†]Carnegie Mellon University

[‡]ETH Zürich

[§]NVIDIA

[¶]University of Texas at Austin

Power Measurement Platform

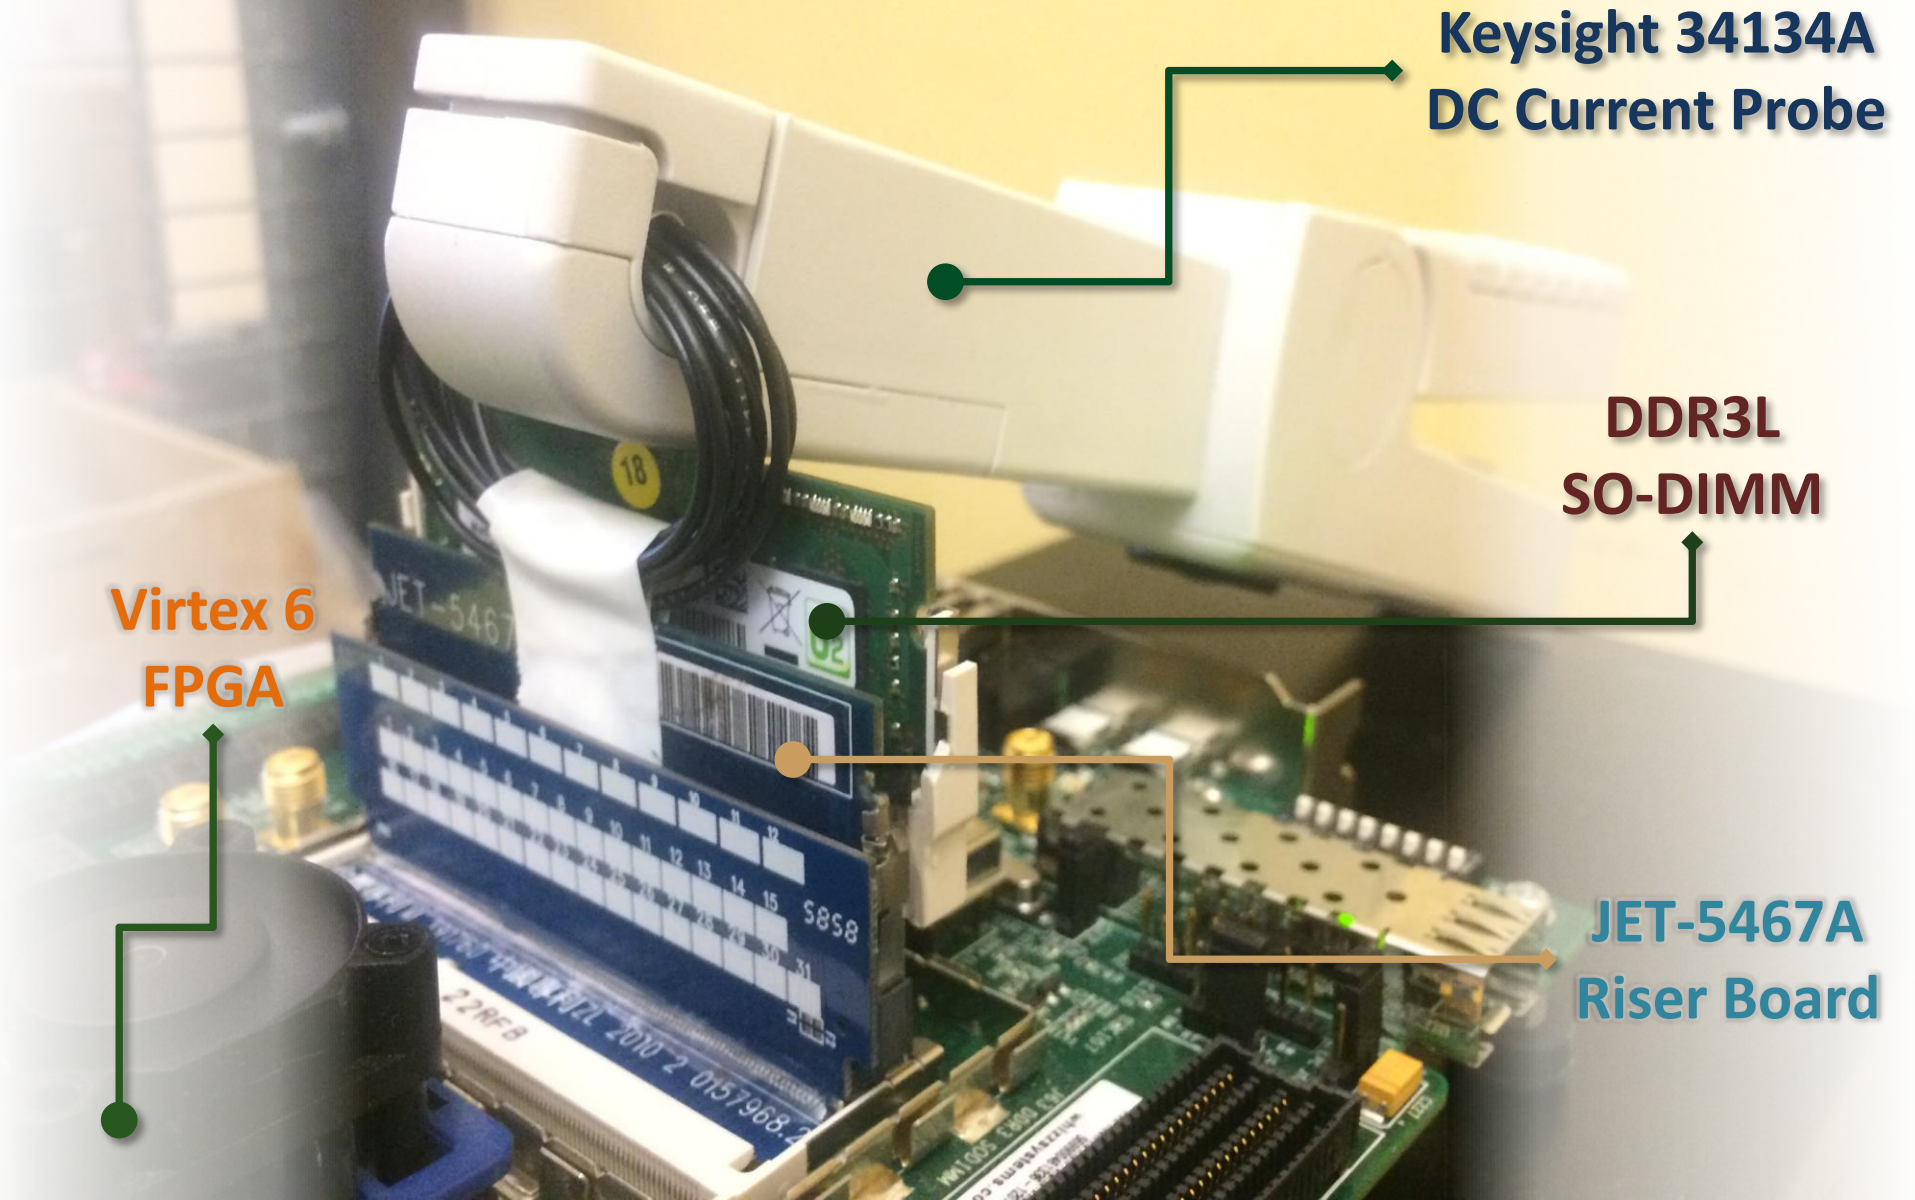
SAFARI

**Keysight 34134A
DC Current Probe**

**DDR3L
SO-DIMM**

**Virtex 6
FPGA**

**JET-5467A
Riser Board**



Summary: Low-Latency Memory

Fundamentally Low Latency Computing Architectures

Summary: Tackling Long Memory Latency

- Reason 1: Design of DRAM Micro-architecture
 - Goal: Maximize capacity/area, not minimize latency
- Reason 2: “One size fits all” approach to latency specification
 - Same latency parameters for all temperatures
 - Same latency parameters for all DRAM chips (e.g., rows)
 - Same latency parameters for all parts of a DRAM chip
 - Same latency parameters for all supply voltage levels
 - Same latency parameters for all application data
 - ...

We Can Reduce
Memory Latency
with Change of Mindset

Main Memory Needs
Intelligent Controllers
to Reduce Latency

Some Solution Principles

- Data-centric design
 - All components intelligent
 - Better cross-layer communication, better interfaces
 - Better-than-worst-case design
 - Heterogeneity
 - Flexibility, adaptability
- Open minds**

Four Key Current Directions

- Fundamentally **Secure/Reliable/Safe** Architectures
- Fundamentally **Energy-Efficient** Architectures
 - **Memory-centric** (Data-centric) Architectures
- Fundamentally **Low-Latency and Predictable** Architectures
- Architectures for **AI/ML, Genomics, Medicine, Health, ...**

Computer Architecture

Lecture 12a: Low-Latency Memory II

Prof. Onur Mutlu

ETH Zürich

Fall 2021

5 November 2021

Solar-DRAM

Solar-DRAM: Putting It Together

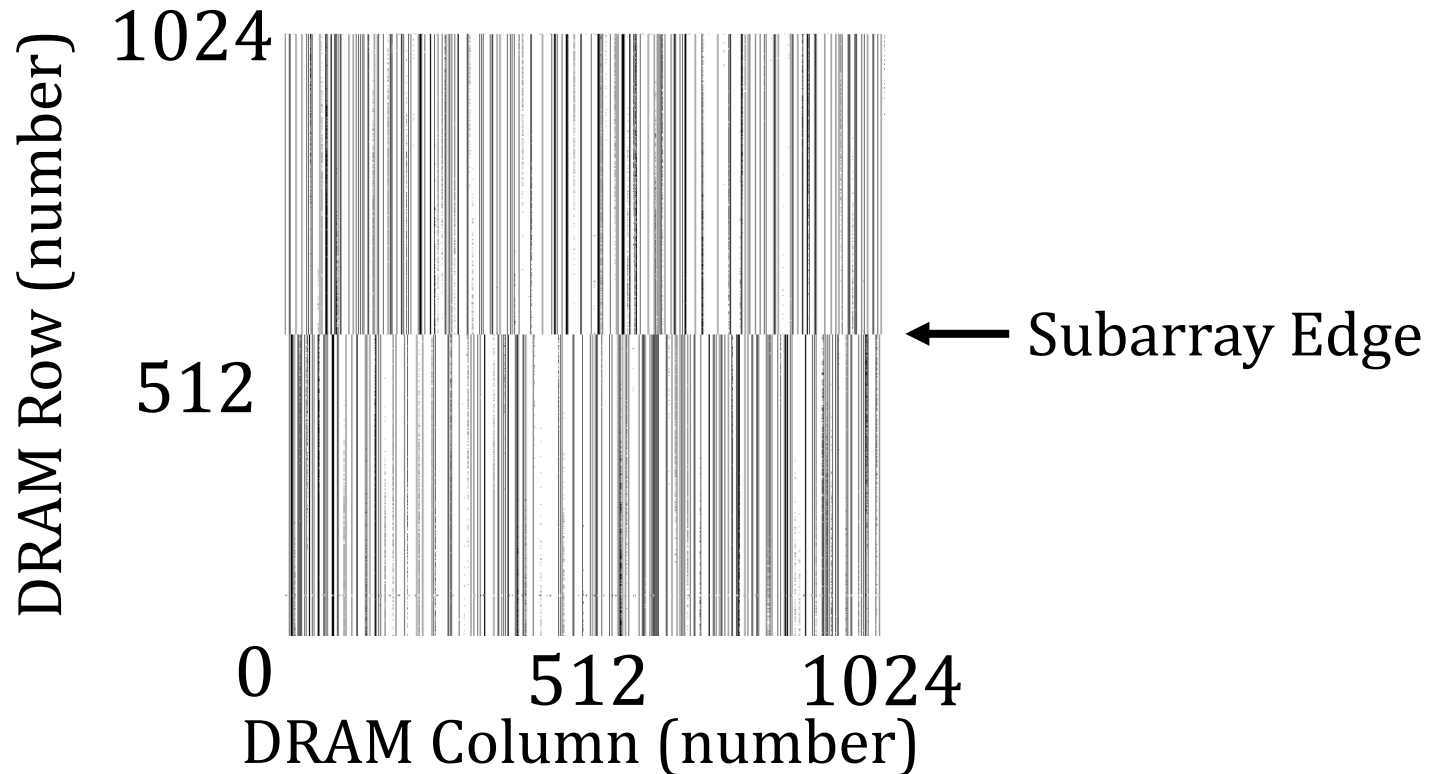
- Jeremie S. Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu,
"Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines"
Proceedings of the 36th IEEE International Conference on Computer Design (ICCD), Orlando, FL, USA, October 2018.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (16 minutes)]

Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines

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

Spatial Distribution of Failures

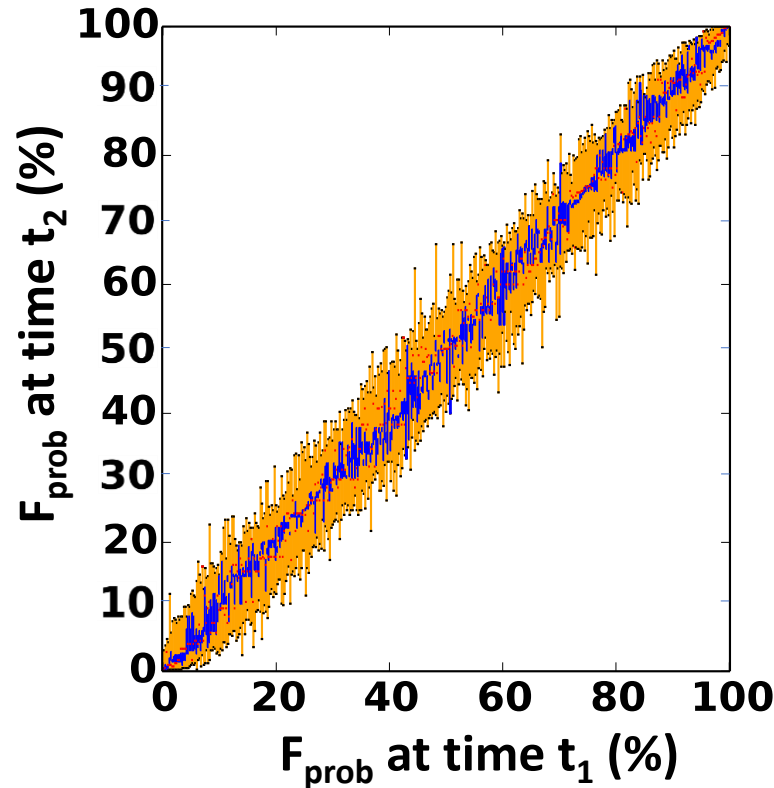
How are activation failures spatially distributed in DRAM?



Activation failures are **highly constrained**
to local bitlines

Short-term Variation

Does a bitline's probability of failure change over time?



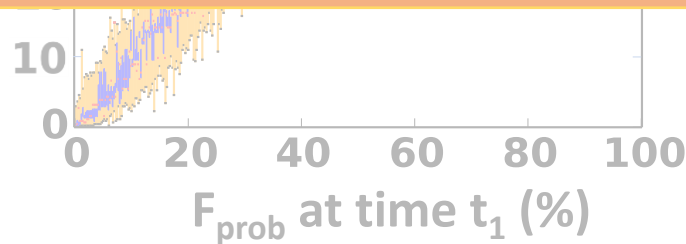
A **weak** bitline is likely to remain **weak** and
a **strong** bitline is likely to remain **strong** over time

Short-term Variation

Does a bitline's probability of failure change over time?



We can rely on a **static profile** of weak bitlines to determine whether an access will cause failures

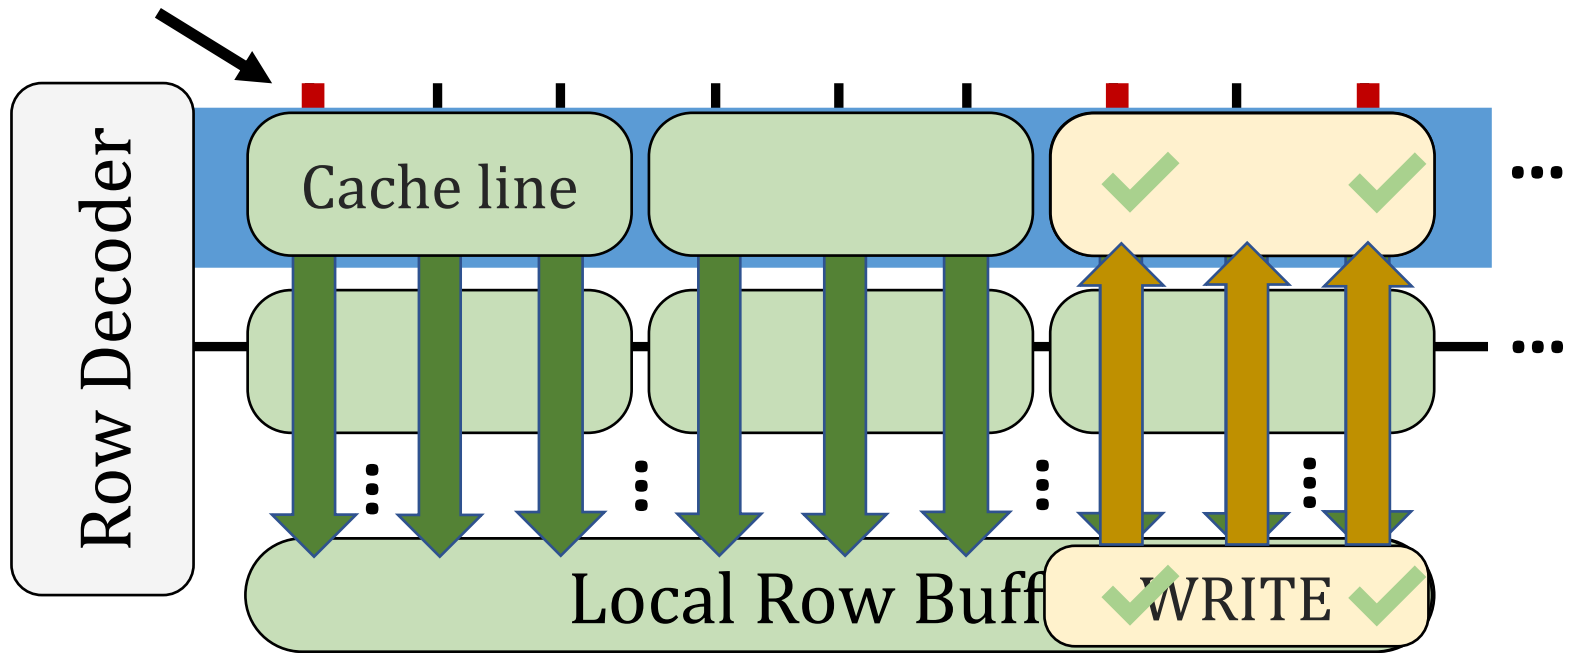


A **weak bitline** is likely to remain **weak** and a **strong bitline** is likely to remain **strong** over time

Write Operations

How are write operations affected by reduced t_{RCD} ?

Weak bitline



We can reliably issue write operations
with significantly reduced t_{RCD} (e.g., by 77%)

Solar-DRAM

Uses a **static profile of weak subarray columns**

- Identifies subarray columns as weak or strong
- Obtained in a one-time profiling step

Three Components

1. Variable-latency cache lines (VLC)
2. Reordered subarray columns (RSC)
3. Reduced latency for writes (RLW)

Solar-DRAM

Uses a **static profile of weak subarray columns**

- Identifies subarray columns as weak or strong
- Obtained in a one-time profiling step

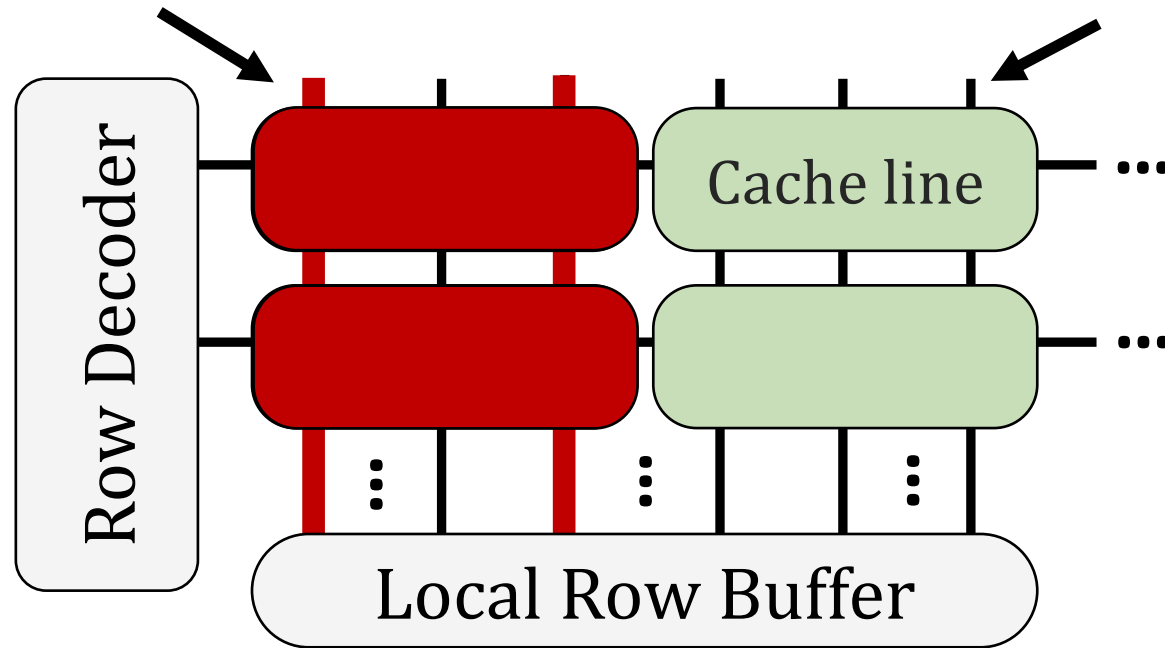
Three Components

1. Variable-latency cache lines (VLC)
2. Reordered subarray columns (RSC)
3. Reduced latency for writes (RLW)

Solar-DRAM: VLC (I)

Weak bitline

Strong bitline



Identify cache lines comprised of **strong bitlines**

Access such cache lines with a **reduced t_{RCD}**

Solar-DRAM

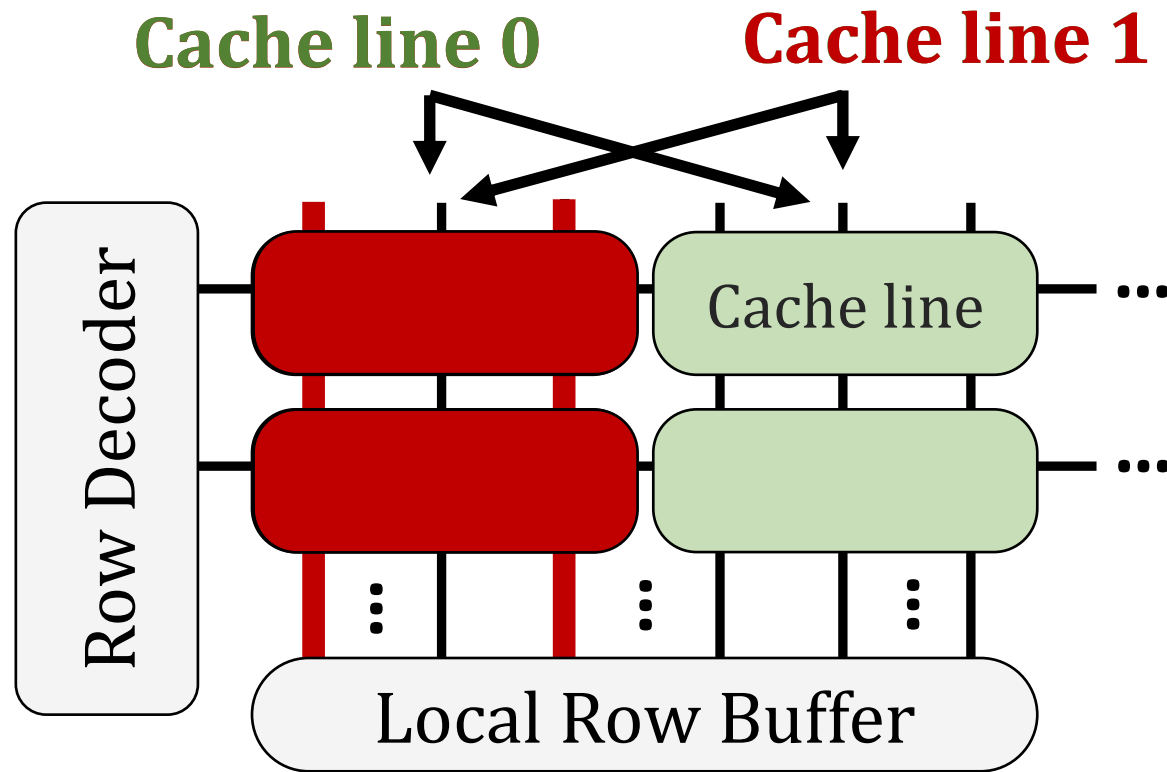
Uses a **static profile of weak subarray columns**

- Identifies subarray columns as weak or strong
- Obtained in a one-time profiling step

Three Components

1. Variable-latency cache lines (VLC)
2. Reordered subarray columns (RSC)
3. Reduced latency for writes (RLW)

Solar-DRAM: RSC (II)



Remap cache lines across DRAM at the memory controller level so cache line 0 will likely map to a **strong** cache line

Solar-DRAM

Uses a **static profile of weak subarray columns**

- Identifies subarray columns as weak or strong
- Obtained in a one-time profiling step

Three Components

1. Variable-latency cache lines (VLC)
2. Reordered subarray columns (RSC)
3. Reduced latency for writes (RLW)

Solar-DRAM: Putting It Together

- Jeremie S. Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu,
"Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines"
Proceedings of the 36th IEEE International Conference on Computer Design (ICCD), Orlando, FL, USA, October 2018.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (16 minutes)]

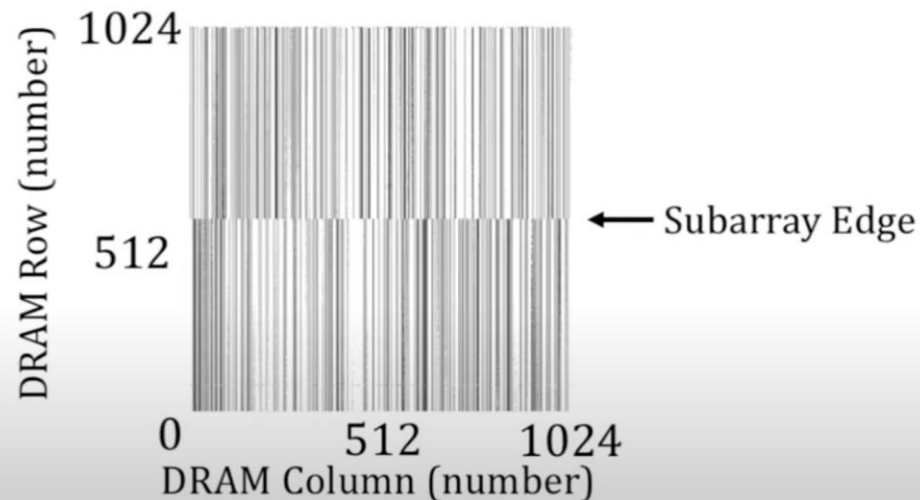
Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines

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

More on Solar DRAM

Spatial Distribution of Failures

How are activation failures spatially distributed in DRAM?



Activation failures are **highly constrained**
to local bitlines (i.e., subarrays)

Solar-DRAM: Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines - ICCD 2018

101 views • Oct 23, 2018

4 0 SHARE SAVE ...



Jeremie Kim
18 subscribers

SUBSCRIBE

Understanding & Exploiting the Voltage-Latency-Reliability Relationship

Analysis of Latency-Voltage in DRAM Chips

- Kevin Chang, A. Giray Yaglikci, Saugata Ghose, Aditya Agrawal, Niladrish Chatterjee, Abhijith Kashyap, Donghyuk Lee, Mike O'Connor, Hasan Hassan, and u,

"Understanding Reduced-Voltage Operation in Modern DRAM Devices: Experimental Characterization, Analysis, and Mechanisms"

*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Urbana-Champaign, IL, USA, June 2017.*

Understanding Reduced-Voltage Operation in Modern DRAM Chips: Characterization, Analysis, and Mechanisms

Kevin K. Chang[†] Abdullah Giray Yağlıkçı[†] Saugata Ghose[†] Aditya Agrawal[¶] Niladrish Chatterjee[¶]
Abhijith Kashyap[†] Donghyuk Lee[¶] Mike O'Connor^{¶,‡} Hasan Hassan[§] Onur Mutlu^{§,†}

[†]Carnegie Mellon University

[¶]NVIDIA

[‡]The University of Texas at Austin

[§]ETH Zürich

High DRAM Power Consumption

- Problem: High DRAM (memory) power in today's systems



>40% in POWER7 (Ware+, HPCA'10)



>40% in GPU (Paul+, ISCA'15)

Low-Voltage Memory

- Existing DRAM designs to help reduce DRAM power by lowering supply voltage conservatively
 - $Power \propto Voltage^2$
- DDR3L (low-voltage) reduces voltage from 1.5V to 1.35V (-10%)
- LPDDR4 (low-power) employs low-power I/O interface with 1.2V (lower bandwidth)

Can we reduce DRAM power and energy by further reducing supply voltage?

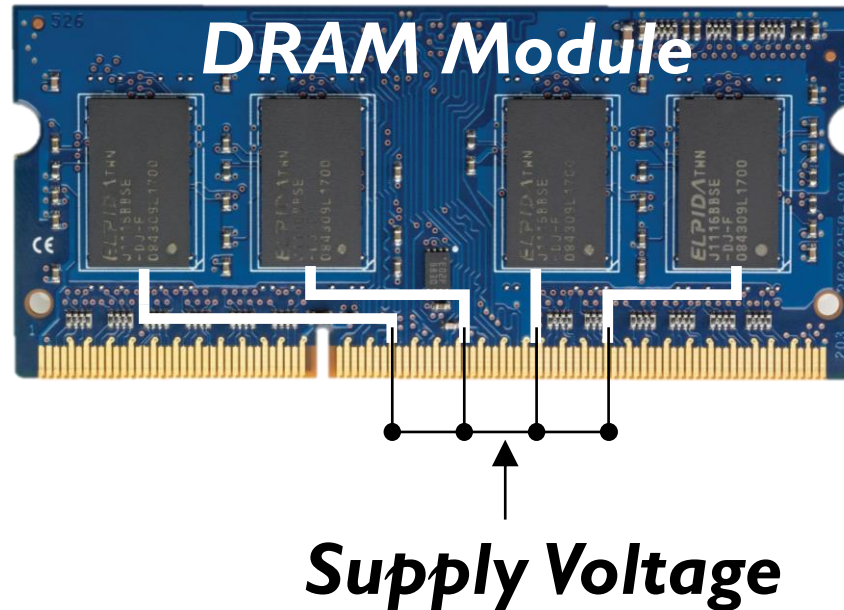
Goals

- 1 Understand and characterize the various characteristics of DRAM under **reduced voltage**
- 2 Develop a mechanism that reduces DRAM energy by **lowering voltage** while keeping performance loss within a target

Key Questions

- How does reducing voltage affect ***reliability*** (errors)?
- How does reducing voltage affect ***DRAM latency***?
- How do we design a new DRAM energy reduction mechanism?

Supply Voltage Control on DRAM



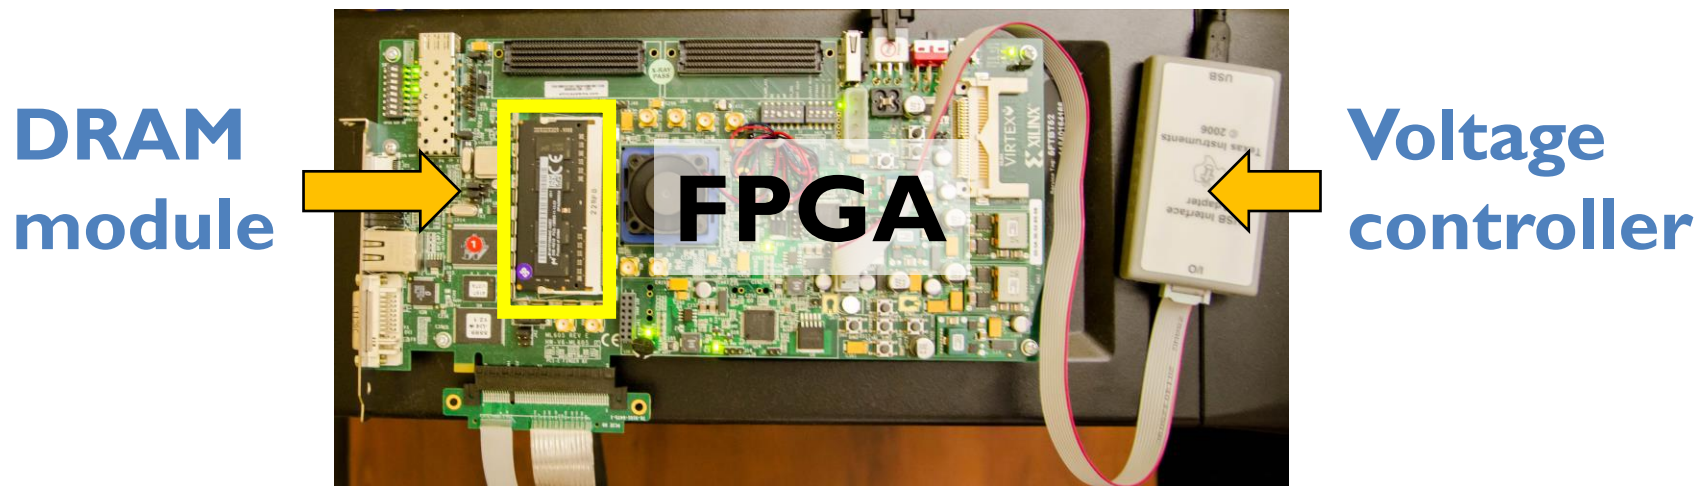
Adjust the *supply voltage* to every chip on the same module

Custom Testing Platform

SoftMC [Hassan+, HPCA'17]: FPGA testing platform to

- 1) Adjust supply voltage to DRAM modules
- 2) Schedule DRAM commands to DRAM modules

Existing systems: DRAM commands not exposed to users

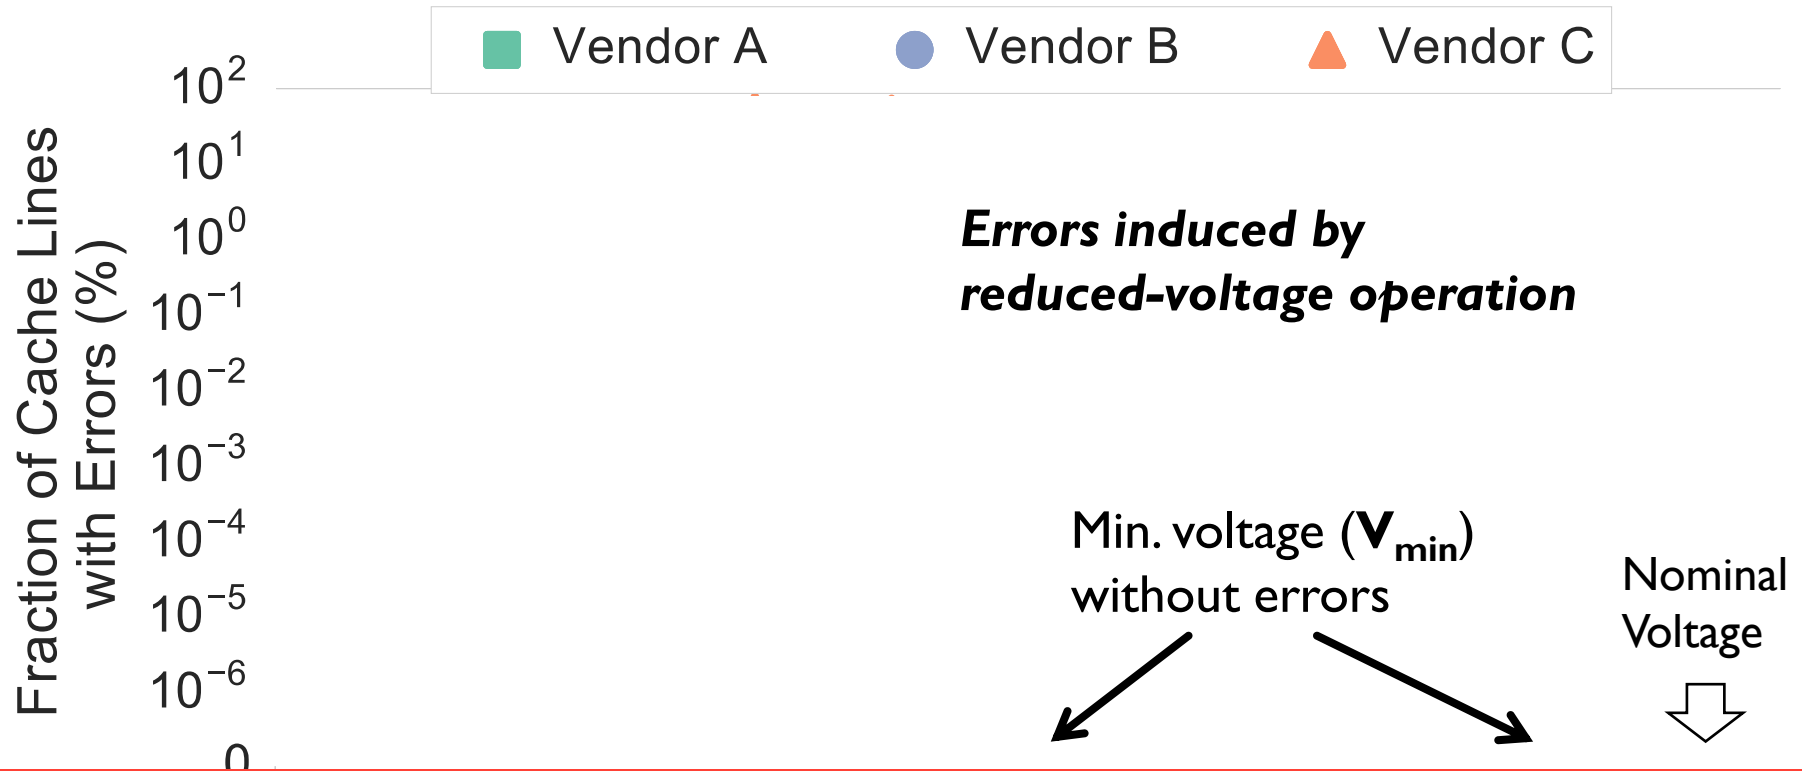


<https://github.com/CMU-SAFARI/DRAM-Voltage-Study>

Tested DRAM Modules

- **124 DDR3L** (low-voltage) DRAM chips
 - **31 SO-DIMMs**
 - **1.35V** (DDR3 uses 1.5V)
 - Density: 4Gb per chip
 - Three major vendors/manufacturers
 - Manufacturing dates: 2014-2016
- Iteratively read every bit in each 4Gb chip under a wide range of supply voltage levels: 1.35V to 1.0V (**-26%**)

Reliability Worsens with Lower Voltage



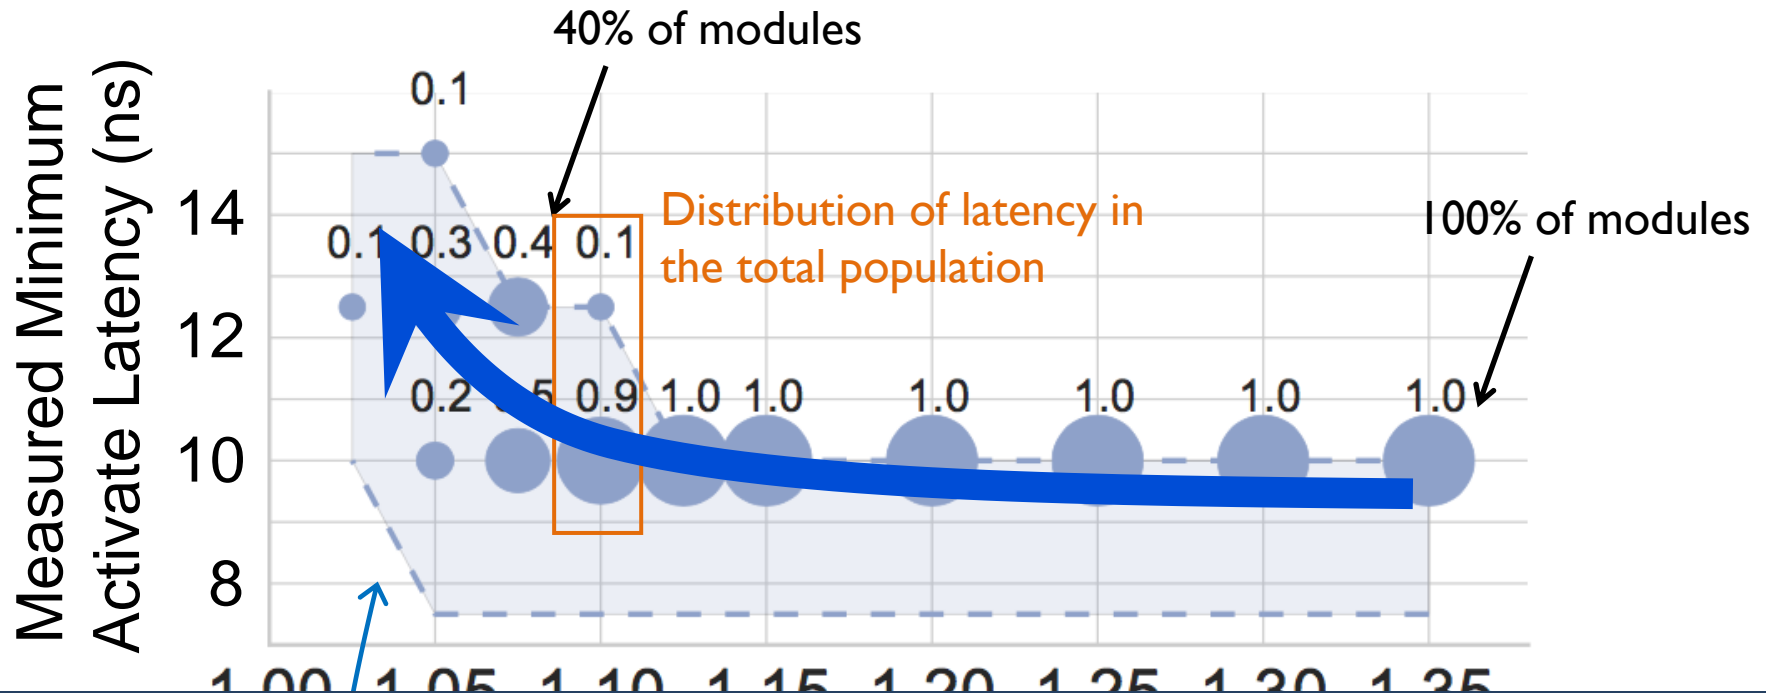
Reducing voltage below V_{\min} causes an increasing number of errors

Detailed circuit simulations (SPICE) of a DRAM cell array to model the behavior of DRAM operations



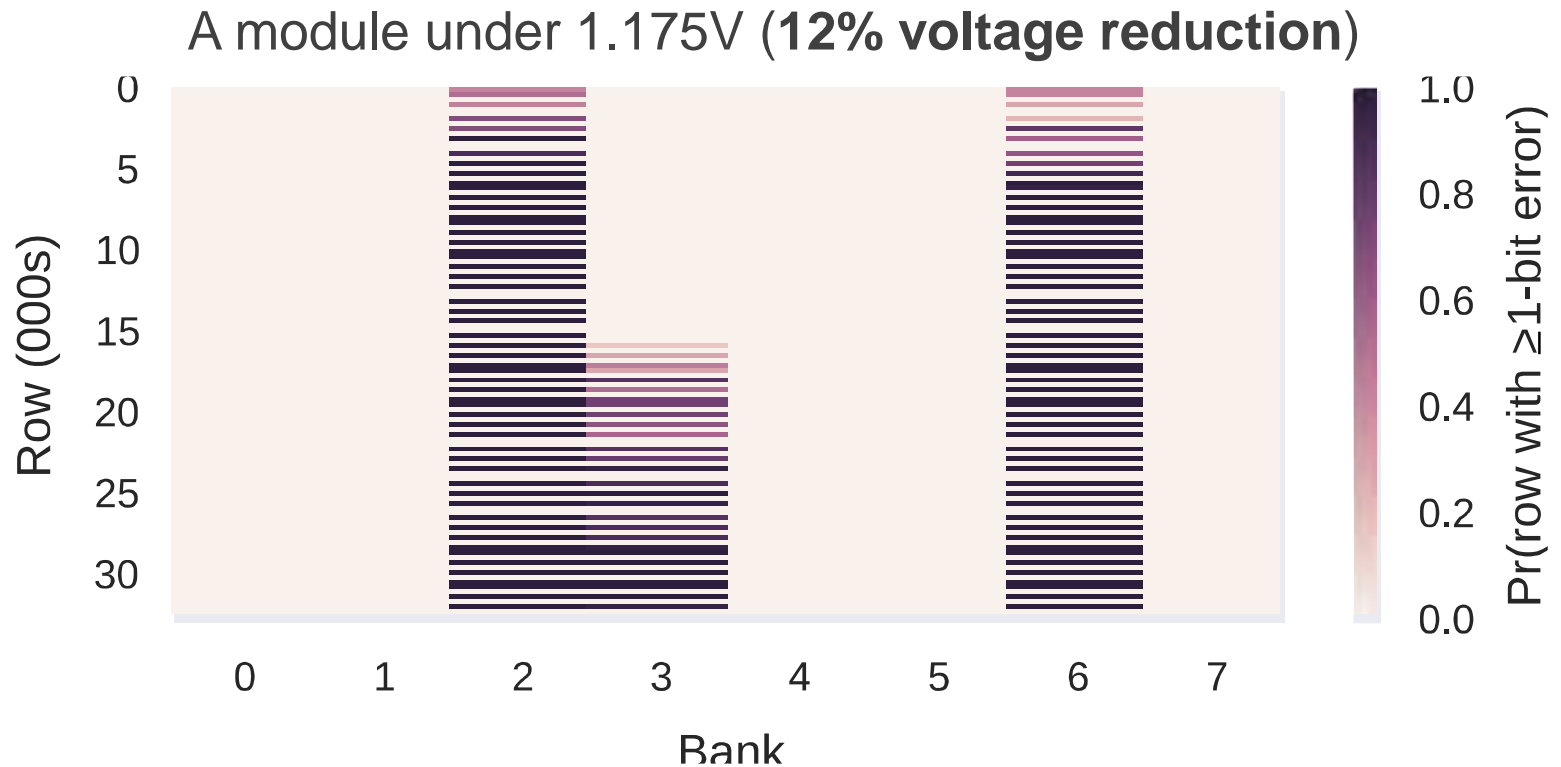
DIMMs Operating at Higher Latency

Measured minimum latency that *does not* cause errors in DRAM modules



DRAM requires longer latency to access data **without errors** at lower voltage

Spatial Locality of Errors



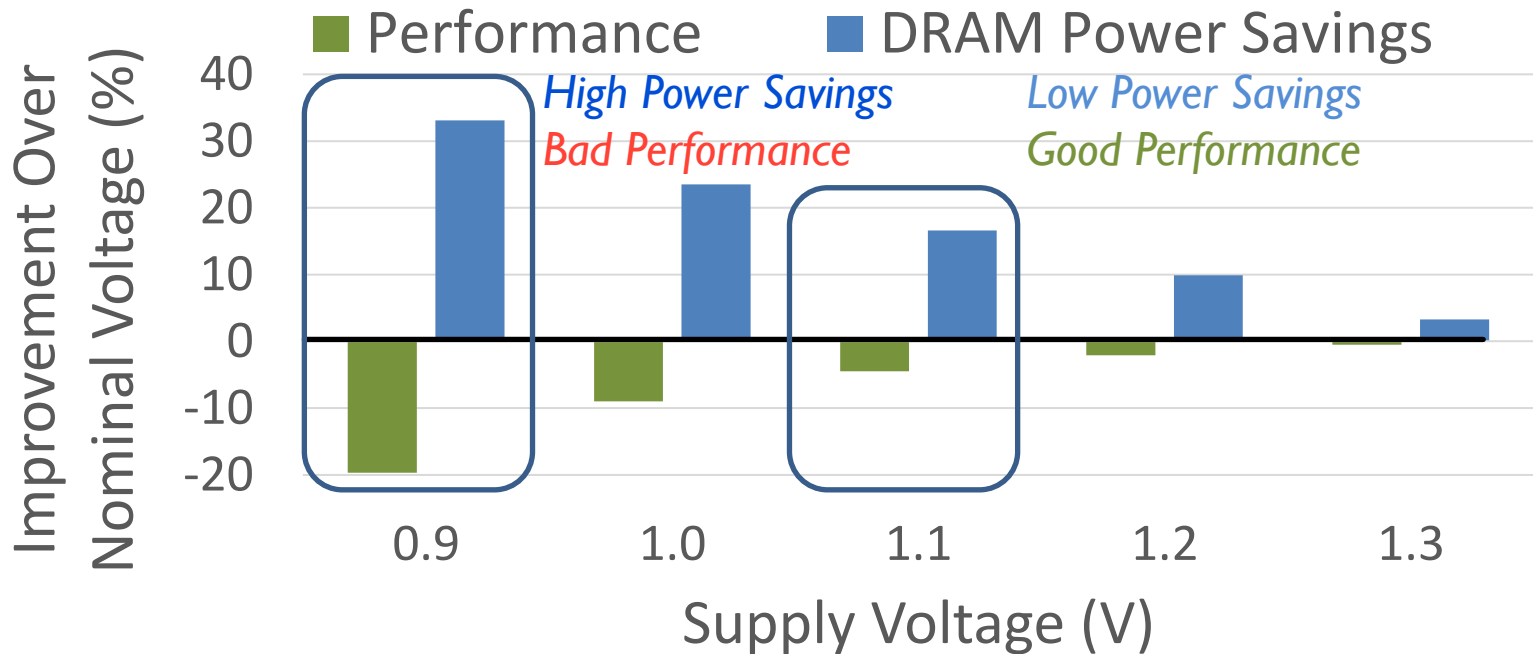
Errors concentrate in certain regions

Summary of Key Experimental Observations

- Voltage-induced errors increase as voltage reduces further below V_{\min}
- Errors exhibit spatial locality
- Increasing the latency of DRAM operations mitigates voltage-induced errors

DRAM Voltage Adjustment to Reduce Energy

- Goal: Exploit the trade-off between voltage and latency to reduce energy consumption
- Approach: Reduce DRAM voltage **reliably**
 - **Performance loss** due to increased latency at lower voltage

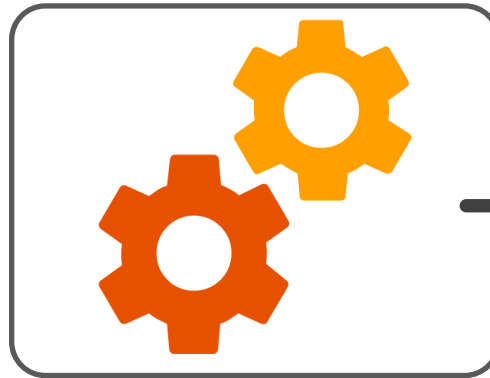


Voltron Overview

Voltron



User specifies the
performance loss target

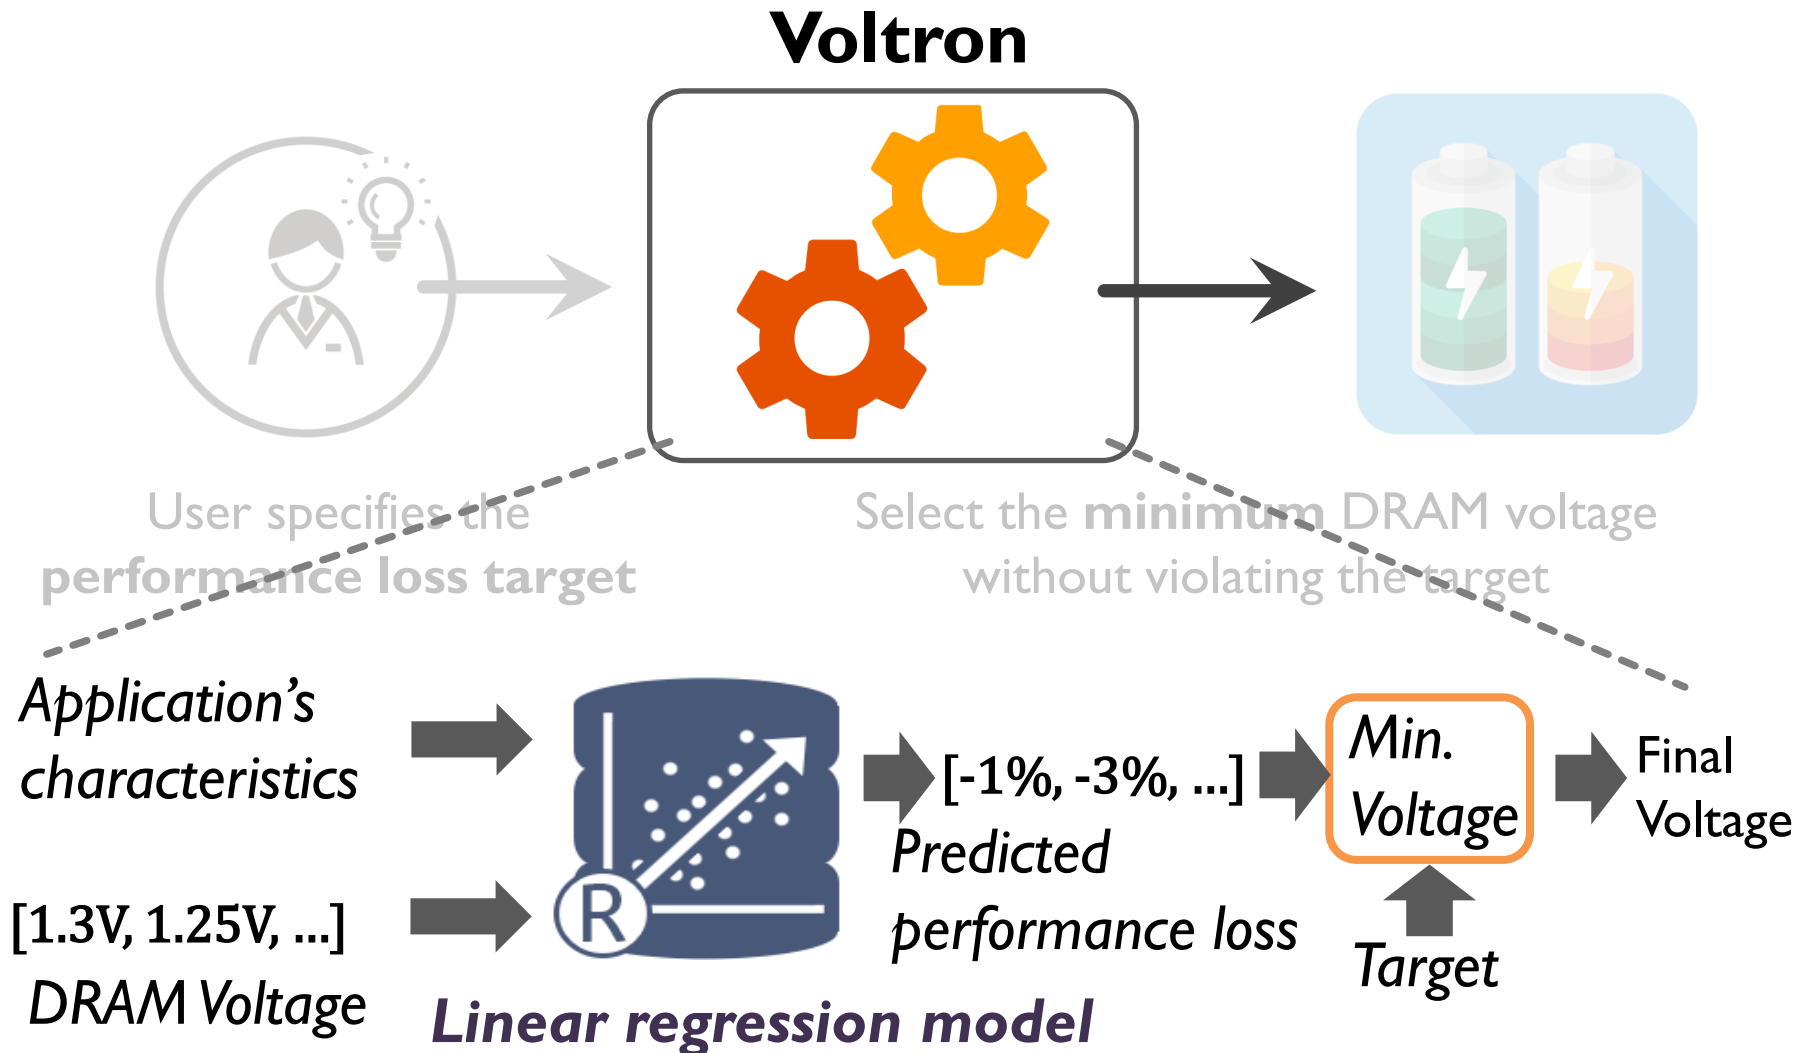


Select the **minimum** DRAM voltage
without violating the target



How do we predict performance loss due to increased latency under low DRAM voltage?

Linear Model to Predict Performance

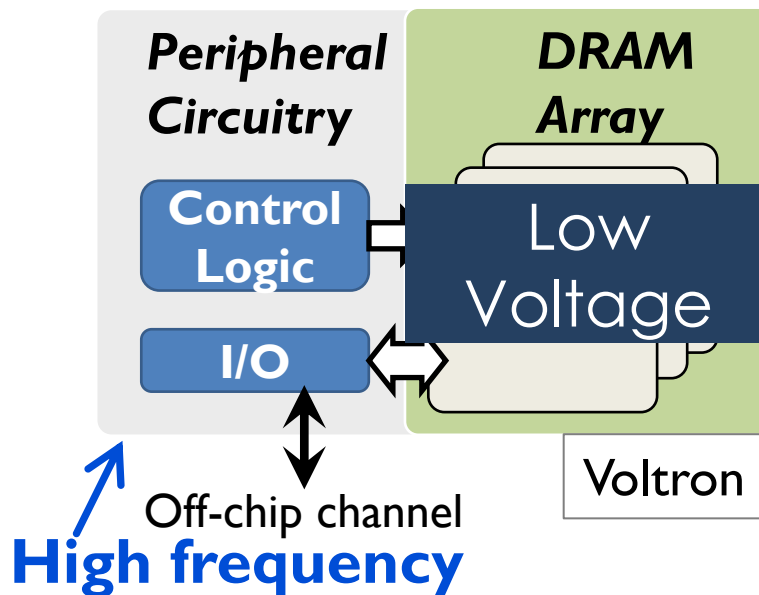
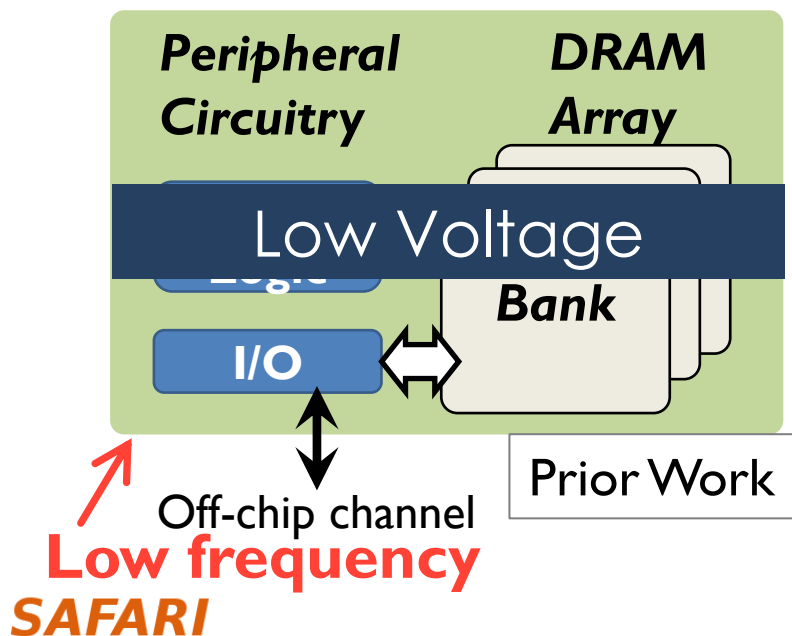


Regression Model to Predict Performance

- Application's characteristics for the model:
 - **Memory intensity**: Frequency of last-level cache misses
 - **Memory stall time**: Amount of time memory requests stall commit inside CPU
- Handling multiple applications:
 - Predict a performance loss for each application
 - Select the minimum voltage that satisfies the performance target for all applications

Comparison to Prior Work

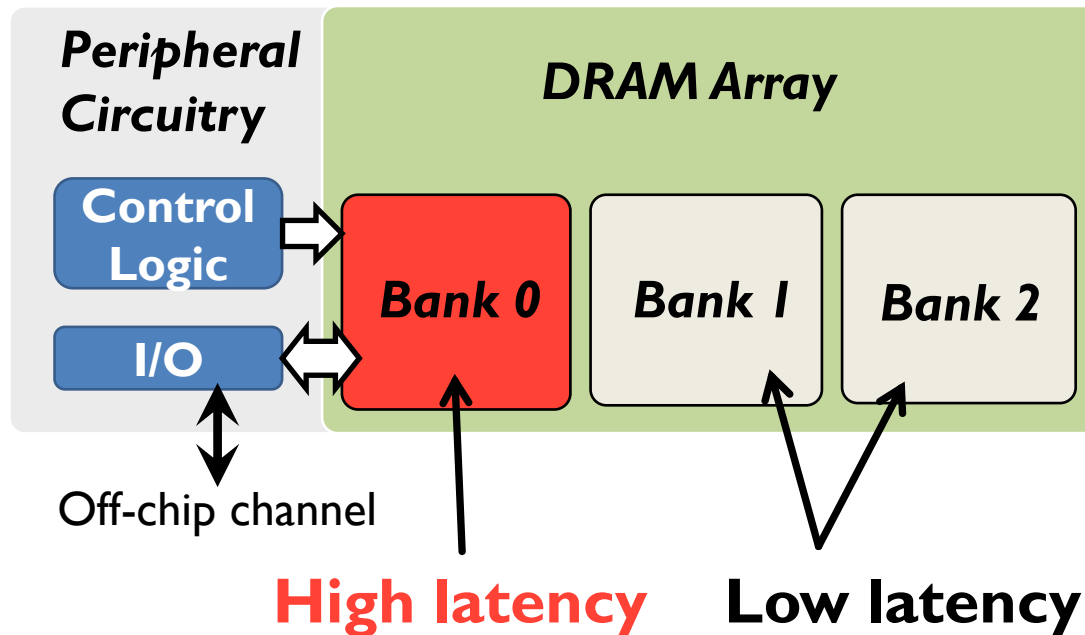
- Prior work: Dynamically scale *frequency and voltage* of the entire DRAM based on bandwidth demand [David+, ICAC'11]
 - Problem: Lowering voltage on the peripheral circuitry decreases channel frequency (memory data throughput)
- Voltron: Reduce voltage to only **DRAM array** without changing the voltage to peripheral circuitry



Exploiting Spatial Locality of Errors

Key idea: Increase the latency only for DRAM banks that observe errors under low voltage

- Benefit: Higher performance

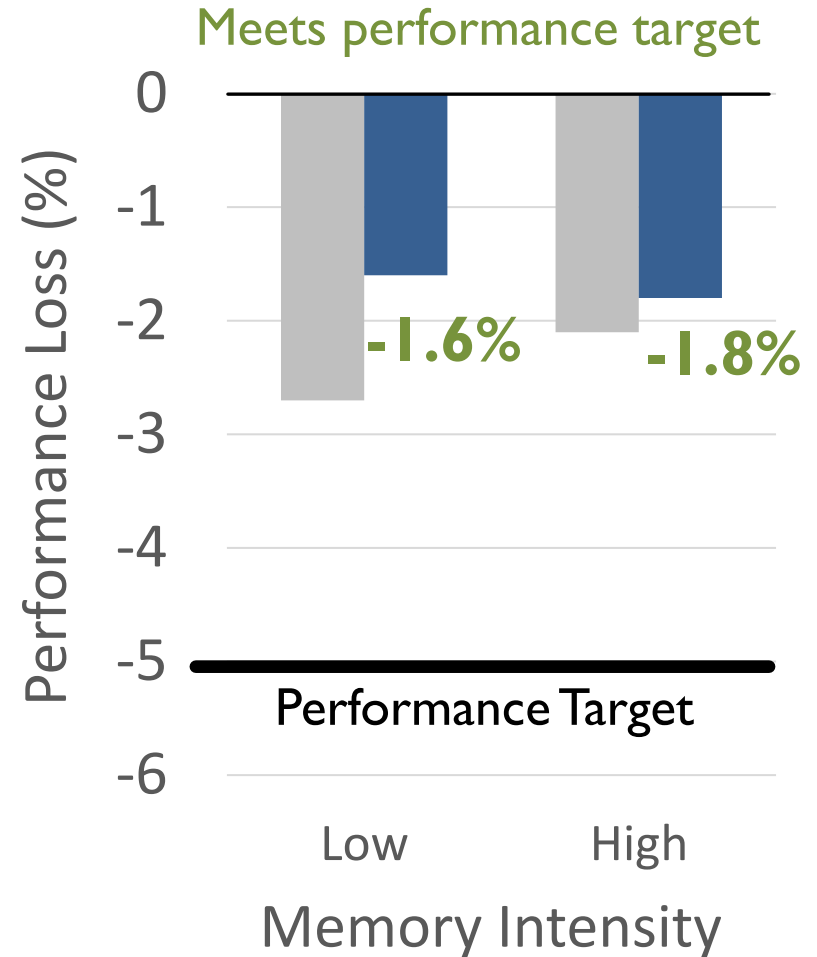
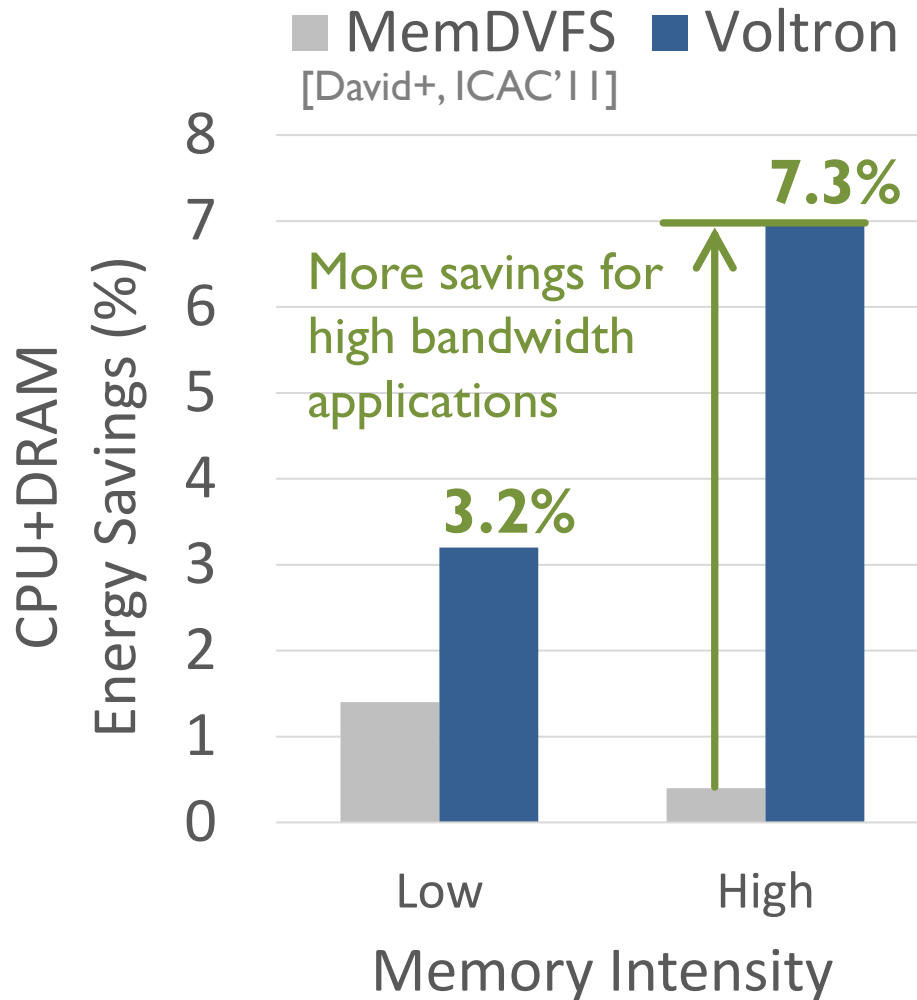


Voltron Evaluation Methodology

- **Cycle-level simulator:** Ramulator [CAL'15]
 - **McPAT** and **DRAMPower** for energy measurement

<https://github.com/CMU-SAFARI/ramulator>
- **4-core** system with DDR3L memory
- **Benchmarks:** SPEC2006, YCSB
- Comparison to prior work: **MemDVFS** [David+, ICAC'11]
 - Dynamic DRAM frequency and voltage scaling
 - Scaling based on the *memory bandwidth consumption*

Energy Savings with Bounded Performance



Voltron: Advantages & Disadvantages

■ Advantages

- + Can trade-off between voltage and latency to improve energy or performance
- + Can exploit the high voltage margin present in DRAM

■ Disadvantages

- Requires finding the reliable operating voltage for each chip → higher testing cost

Analysis of Latency-Voltage in DRAM Chips

- Kevin Chang, A. Giray Yaglikci, Saugata Ghose, Aditya Agrawal, Niladrish Chatterjee, Abhijith Kashyap, Donghyuk Lee, Mike O'Connor, Hasan Hassan, and Onur Mutlu,

"Understanding Reduced-Voltage Operation in Modern DRAM Devices: Experimental Characterization, Analysis, and Mechanisms"

*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Urbana-Champaign, IL, USA, June 2017.*

Understanding Reduced-Voltage Operation in Modern DRAM Chips: Characterization, Analysis, and Mechanisms

Kevin K. Chang[†] Abdullah Giray Yağlıkçı[†] Saugata Ghose[†] Aditya Agrawal[¶] Niladrish Chatterjee[¶]
Abhijith Kashyap[†] Donghyuk Lee[¶] Mike O'Connor^{¶,‡} Hasan Hassan[§] Onur Mutlu^{§,†}

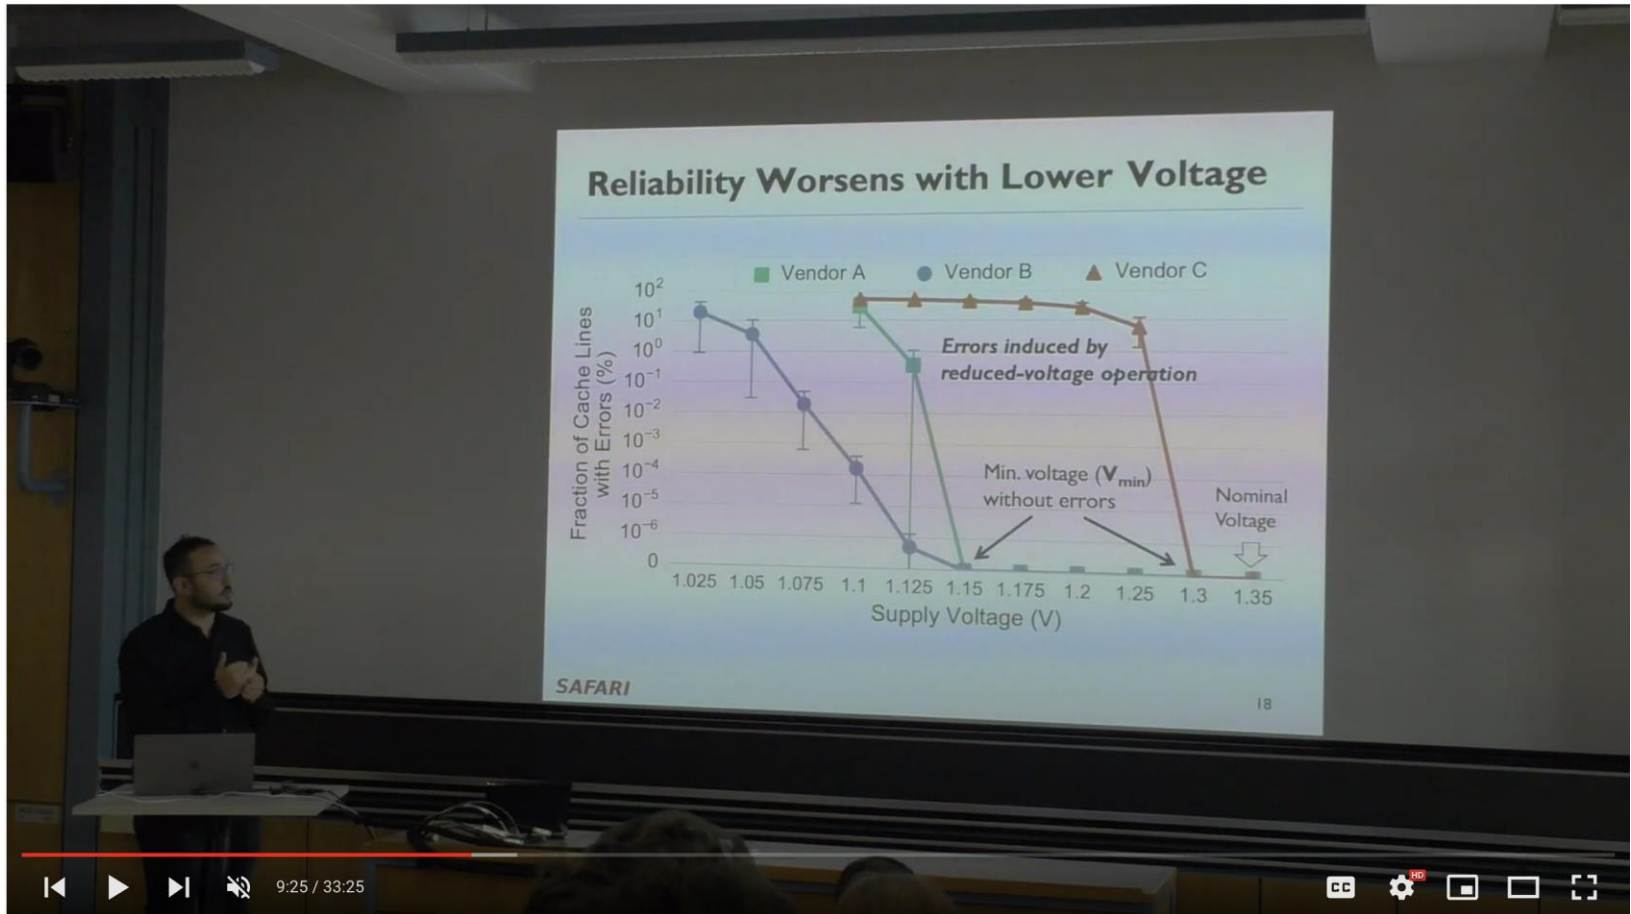
[†]Carnegie Mellon University

[¶]NVIDIA

[‡]The University of Texas at Austin

[§]ETH Zürich

More on Voltron



ETH ZÜRICH

Computer Architecture - Lecture 11c: Voltron: Reducing DRAM Energy (ETH Zürich, Fall 2019)

409 views • Oct 31, 2019

7 0 SHARE SAVE ...



Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED



Reducing Memory Latency to Support Security Primitives

Using Memory for Security

- Generating True Random Numbers (using DRAM)
 - Kim et al., HPCA 2019
 - Olgun et al., ISCA 2021
- Evaluating Physically Unclonable Functions (using DRAM)
 - Kim et al., HPCA 2018
- Quickly Destroying In-Memory Data (using DRAM)
 - Orosa et al., arxiv 2019 + ISCA 2021

D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput

Jeremie S. Kim Minesh Patel

Hasan Hassan Lois Orosa Onur Mutlu

HPCA 2019

SAFARI

ETH zürich

Carnegie Mellon

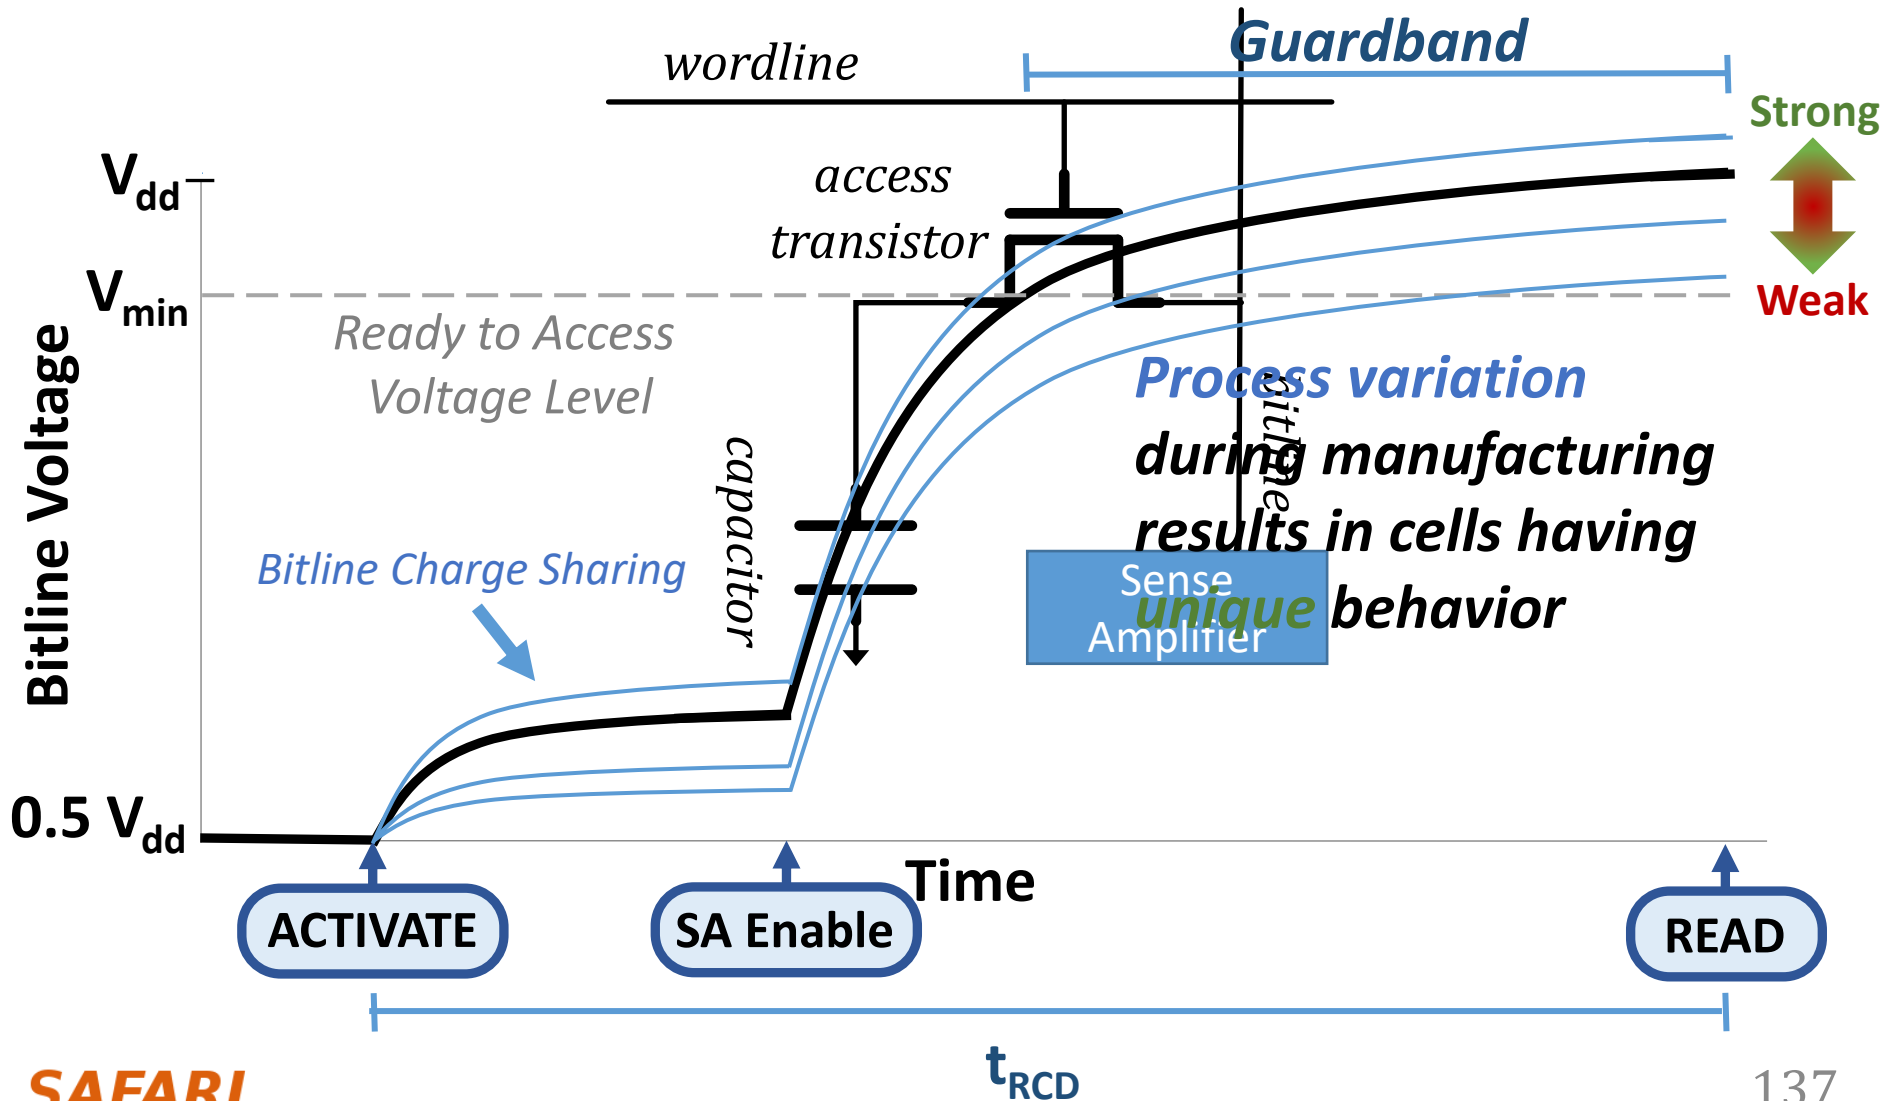
D-RaNGe Executive Summary

- **Motivation**: High-throughput true random numbers enable system security and various randomized algorithms.
 - Many systems (e.g., IoT, mobile, embedded) do not have dedicated **True Random Number Generator (TRNG)** hardware but have DRAM devices
- **Problem**: Current DRAM-based TRNGs either
 1. do **not** sample a fundamentally non-deterministic entropy source
 2. are **too slow** for continuous high-throughput operation
- **Goal**: A novel and effective TRNG that uses **existing** commodity DRAM to provide random values with 1) **high-throughput**, 2) **low latency** and 3) no adverse effect on concurrently running applications
- **D-RaNGe**: Reduce DRAM access latency **below reliable values** and exploit DRAM cells' failure probabilities to generate random values
- **Evaluation**:
 1. Experimentally characterize **282 real LPDDR4 DRAM devices**
 2. **D-RaNGe (717.4 Mb/s)** has significantly higher throughput (**211x**)
 3. **D-RaNGe (100ns)** has significantly lower latency (**180x**)

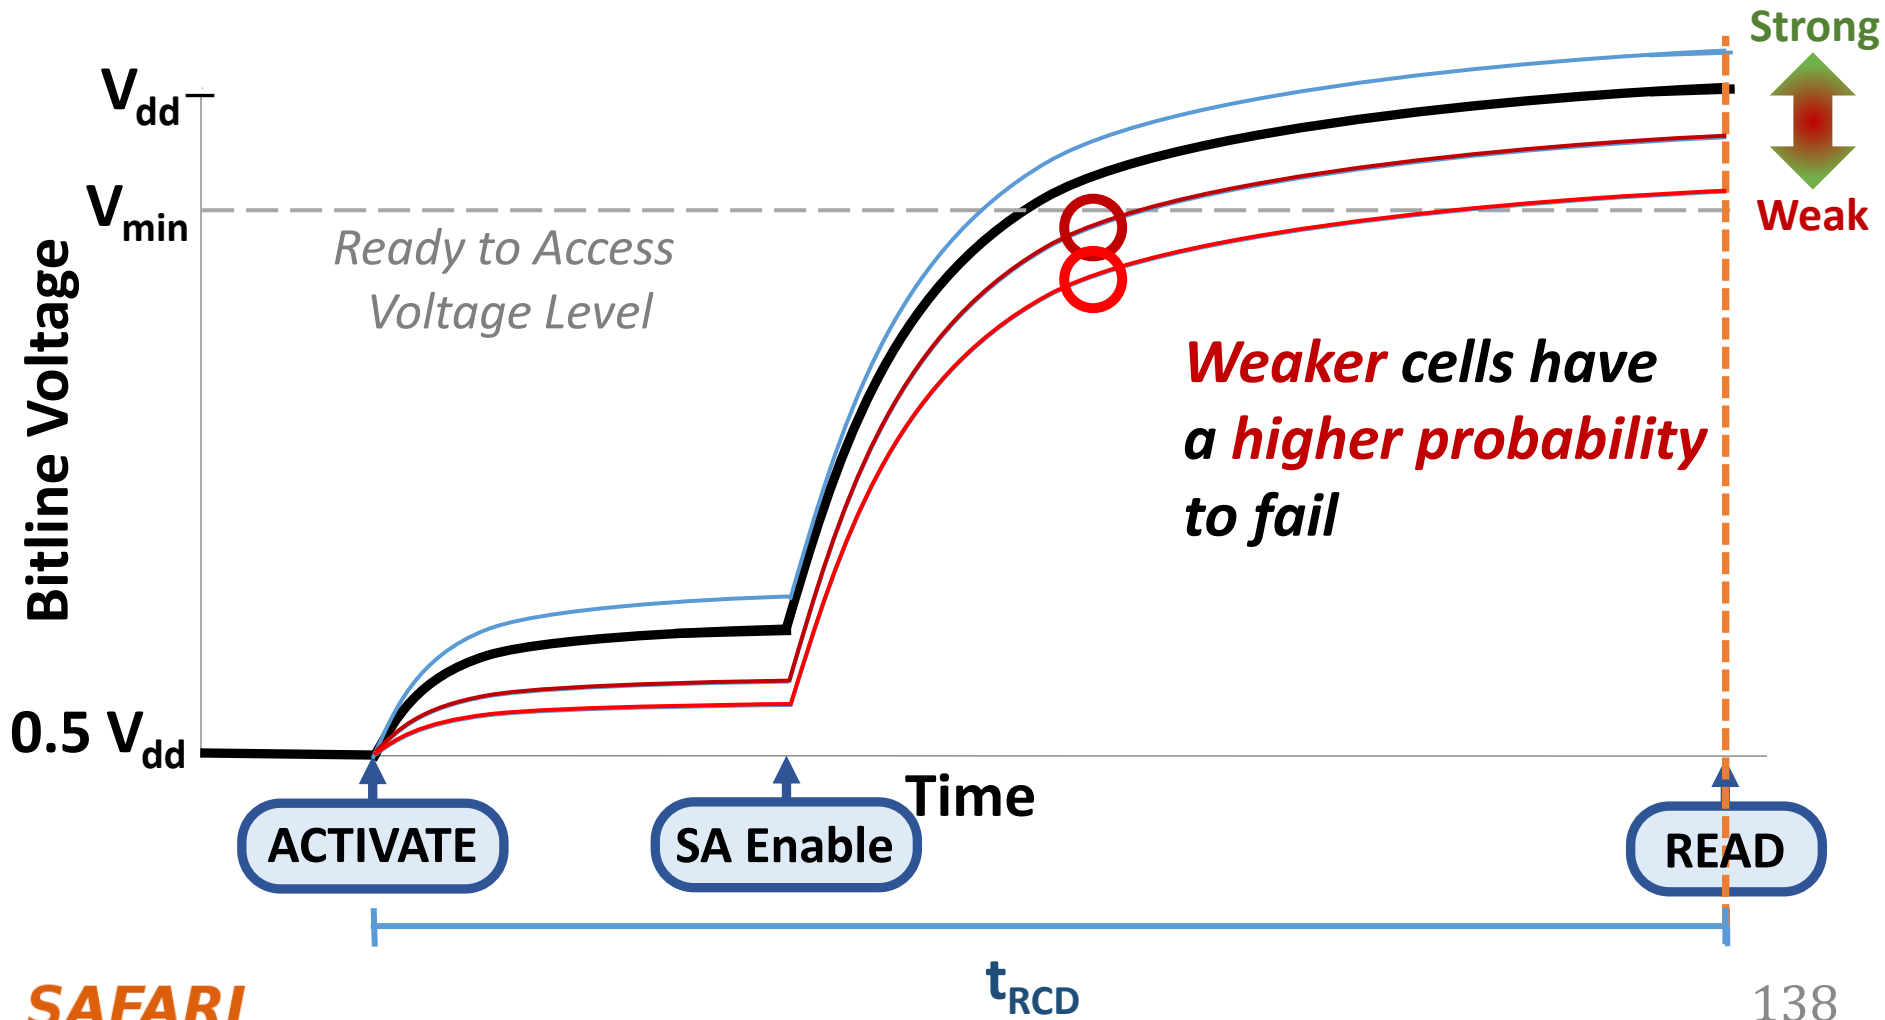
DRAM Latency Characterization of 282 LPDDR4 DRAM Devices

- Latency failures come from accessing DRAM with **reduced** timing parameters.
- **Key Observations:**
 1. A cell's **latency failure** probability is determined by **random process variation**
 2. Some cells fail **randomly**

DRAM Accesses and Failures



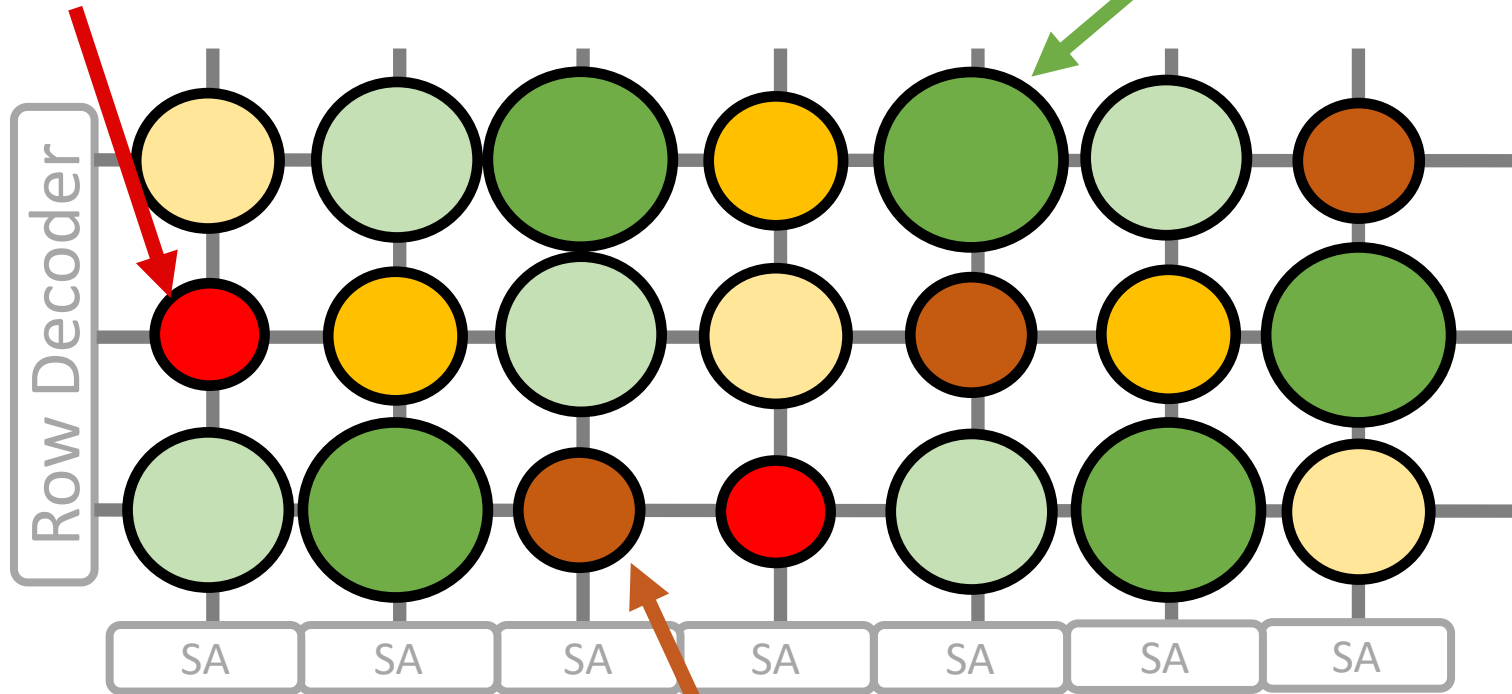
DRAM Accesses and Failures



D-RaNGe Key Idea

High % chance to fail
with reduced t_{RCD}

Low % chance to fail
with reduced t_{RCD}



Fails randomly
with reduced t_{RCD}

D-RaNGe Key Idea

High % chance to fail
with reduced t_{RCD}

Low % chance to fail
with reduced t_{RCD}

**We refer to cells that fail randomly
when accessed with a reduced t_{RCD}
as RNG cells**



Fails randomly
with reduced t_{RCD}

Our D-RaNGe Evaluation

- We generate **random values** by repeatedly accessing **RNG cells** and aggregating the data read
- The random data satisfies the NIST statistical test suite for randomness
- The **D-RaNGE** generates random numbers
 - **Throughput:** 717.4 Mb/s
 - **Latency:** 64 bits in <1us
 - **Power:** 4.4 nJ/bit

D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput

Jeremie S. Kim Minesh Patel

Hasan Hassan Lois Orosa Onur Mutlu

SAFARI

HPCA 2019

ETH zürich

Carnegie Mellon

More on D-RaNGe

- Jeremie S. Kim, Minesh Patel, Hasan Hassan, Lois Orosa, and Onur Mutlu,
"D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput"
Proceedings of the 25th International Symposium on High-Performance Computer Architecture (HPCA), Washington, DC, USA, February 2019.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Full Talk Video](#) (21 minutes)]
[[Full Talk Lecture Video](#) (27 minutes)]
Top Picks Honorable Mention by IEEE Micro.

D-RaNGe: Using Commodity DRAM Devices to Generate True Random Numbers with Low Latency and High Throughput

Jeremie S. Kim^{‡§}

Minesh Patel[§]

Hasan Hassan[§]

Lois Orosa[§]

Onur Mutlu^{§‡}

[‡]Carnegie Mellon University

[§]ETH Zürich

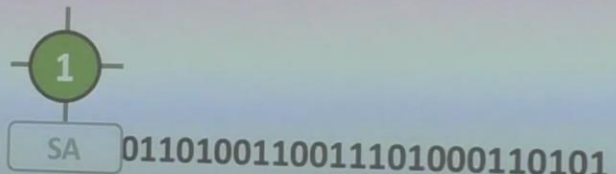
More on DRAM Latency TRNGs

D-RaNGe: Extracting Random Values

Identify all DRAM cells that fail randomly when accessed with a reduced t_{RCD} (**RNG Cell**)

- When accessing an RNG Cell with a reduced t_{RCD} , the values read will be truly random values

RNG Cell



SA 0110100110011101000110101

SAFARI 17/48

8:51 / 27:00 • D-RaNGe Key Idea >

ETH ZÜRICH

Computer Architecture - Lecture 11b: D-RaNGe: True Random Number Generation (ETH Zürich, Fall 2019)

449 views • Oct 31, 2019

Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED

Doing Better Than D-RaNGe

- Ataberk Olgun, Minesh Patel, A. Giray Yaglikci, Haocong Luo, Jeremie S. Kim, F. Nisa Bostanci, Nandita Vijaykumar, Oguz Ergin, and Onur Mutlu,
"QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAM Chips"
Proceedings of the 48th International Symposium on Computer Architecture (ISCA), Virtual, June 2021.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Short Talk Slides \(pptx\)](#)] [[pdf](#)]
[[Talk Video](#) (25 minutes)]
[[SAFARI Live Seminar Video](#) (1 hr 26 mins)]

QUAC-TRNG: High-Throughput True Random Number Generation Using Quadruple Row Activation in Commodity DRAM Chips

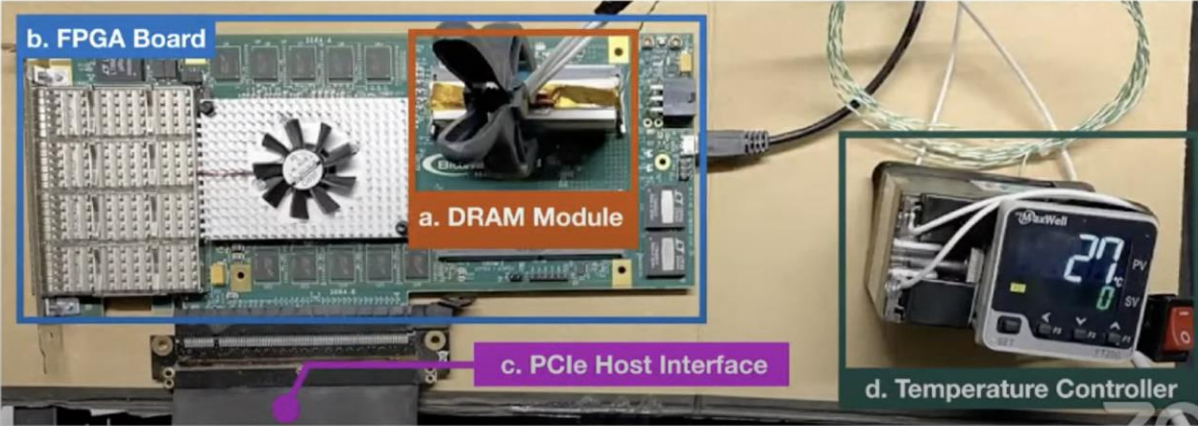
Ataberk Olgun^{§†} Minesh Patel[§] A. Giray Yağlıkçı[§] Haocong Luo[§]
Jeremie S. Kim[§] F. Nisa Bostancı^{§†} Nandita Vijaykumar^{§⊙} Oğuz Ergin[†] Onur Mutlu[§]
[§]ETH Zürich [†]TOBB University of Economics and Technology [⊙]University of Toronto

More on QUAC-TRNG

Real Chip Characterization

Experimentally study QUAC and QUAC-TRNG using 136 real DDR4 chips from SK Hynix

DDR4 SoftMC → DRAM Testing Infrastructure



Ataberk Olgun...

zoom

SAFARI kasirga [Hassan+ HPCA'17] <https://github.com/CMU-SAFARI/SoftMC> CC BY 3.0

SAFARI Live Seminar: High-Throughput TRNG Using Quadruple Row Activation in Commodity DRAM Chips

713 views • Streamed live on Sep 15, 2021

👍 27 💬 0 ➦ SHARE ⚙️ SAVE ...



Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED



DRAM Latency PUFs

- Jeremie S. Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu,
"The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern DRAM Devices"
Proceedings of the 24th International Symposium on High-Performance Computer Architecture (HPCA), Vienna, Austria, February 2018.
[[Lightning Talk Video](#)]
[[Slides \(pptx\)](#)] [[pdf](#)] [[Lightning Session Slides \(pptx\)](#)] [[pdf](#)]
[[Full Talk Lecture Video](#) (28 minutes)]

The DRAM Latency PUF:

Quickly Evaluating Physical Unclonable Functions

by Exploiting the Latency-Reliability Tradeoff in Modern Commodity DRAM Devices

Jeremie S. Kim^{†§}

Minesh Patel[§]

Hasan Hassan[§]

Onur Mutlu^{§†}

[†]Carnegie Mellon University

[§]ETH Zürich

The DRAM Latency PUF:

Quickly Evaluating Physical Unclonable Functions
by Exploiting the Latency-Reliability Tradeoff
in Modern Commodity DRAM Devices

Jeremie S. Kim Minesh Patel

Hasan Hassan Onur Mutlu



SAFARI

ETH zürich

Carnegie Mellon

DL-PUF: Executive Summary

- **Motivation:**

- We can authenticate a system via **unique signatures** if we can evaluate a **Physical Unclonable Function (PUF)** on it
- Signatures (**PUF response**) reflect inherent properties of a device
- DRAM is a promising substrate for PUFs because it is **widely** used

- **Problem:** Current DRAM PUFs are 1) very slow, 2) require a DRAM reboot, or 3) require additional custom hardware

- **Goal:** To develop a novel and effective PUF for **existing** commodity DRAM devices with **low-latency evaluation time** and **low system interference** across **all operating temperatures**

- **DRAM Latency PUF:** Reduce DRAM access latency **below reliable values** and exploit the resulting error patterns as **unique identifiers**

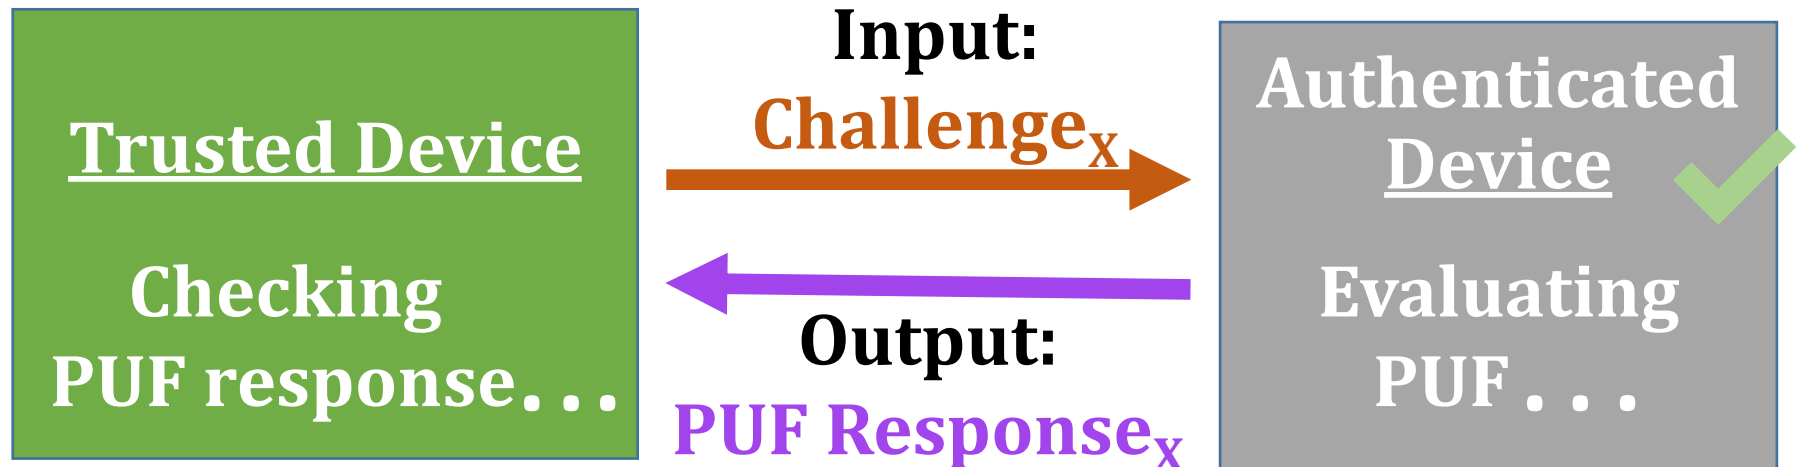
- **Evaluation:**

1. Experimentally characterize **223 real LPDDR4 DRAM devices**
2. **DRAM latency PUF** (88.2 ms) achieves a speedup of **102x/860x** at 70°C/55°C over prior DRAM PUF evaluation mechanisms

Motivation

We want a way to ensure that a system's components are not **compromised**

- **Physical Unclonable Function (PUF)**: a function we **evaluate** on a device to **generate** a **signature unique** to the device
- We refer to the unique signature as a **PUF response**
- Often used in a **Challenge-Response Protocol (CRP)**



Motivation

1. We want a **runtime-accessible** PUF
 - Should be evaluated **quickly** with **minimal** impact on concurrent applications
 - Can protect against **attacks that swap system components with malicious parts**
2. DRAM is a **promising substrate** for evaluating PUFs because it is **ubiquitous** in modern systems
 - Unfortunately, current DRAM PUFs are **slow** and get **exponentially slower** at lower temperatures

DRAM Latency Characterization of 223 LPDDR4 DRAM Devices

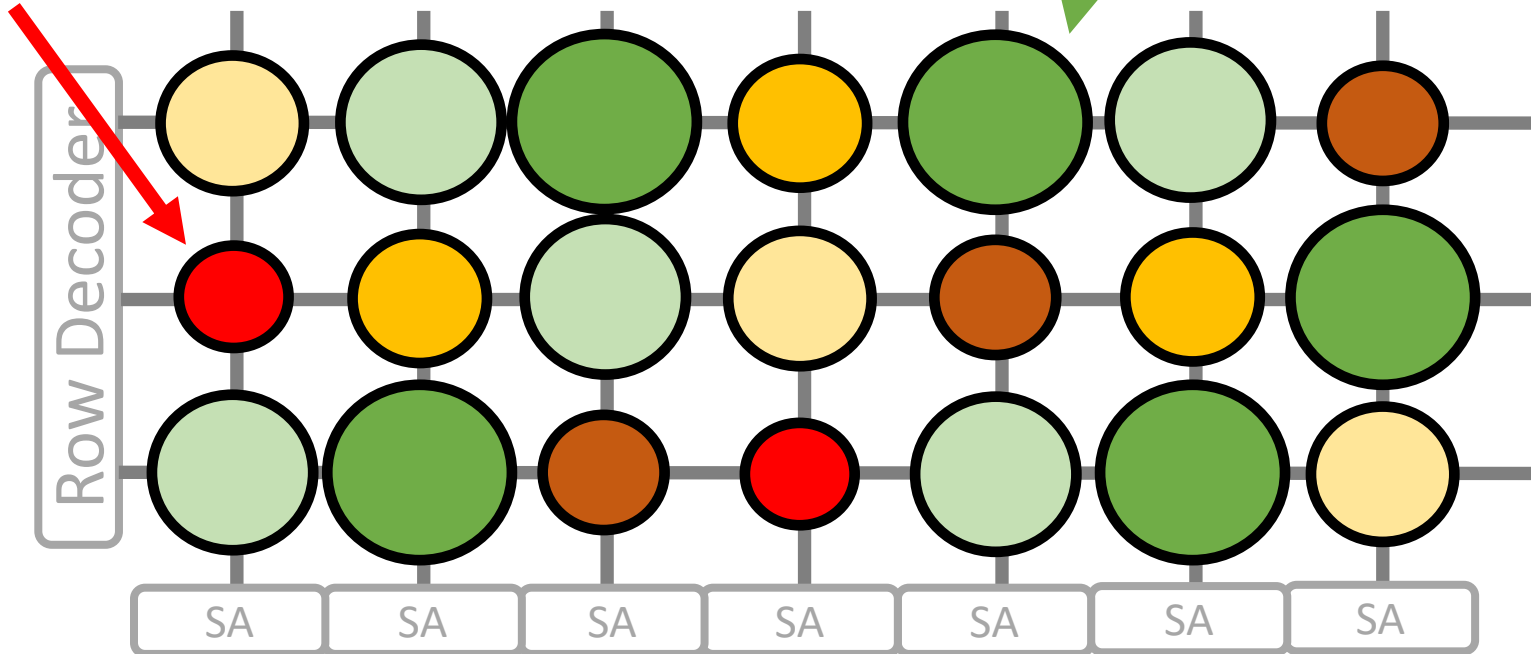
- Latency failures come from accessing DRAM with **reduced** timing parameters.
- **Key Observations:**
 1. A cell's **latency failure** probability is determined by **random process variation**
 2. Latency failure patterns are **repeatable and unique to a device**

DRAM Latency PUF Key Idea

- A cell's latency failure probability is inherently related to **random process variation** from manufacturing
- We can provide **repeatable and unique device signatures** using latency error patterns

High % chance to fail
with reduced t_{RCD}

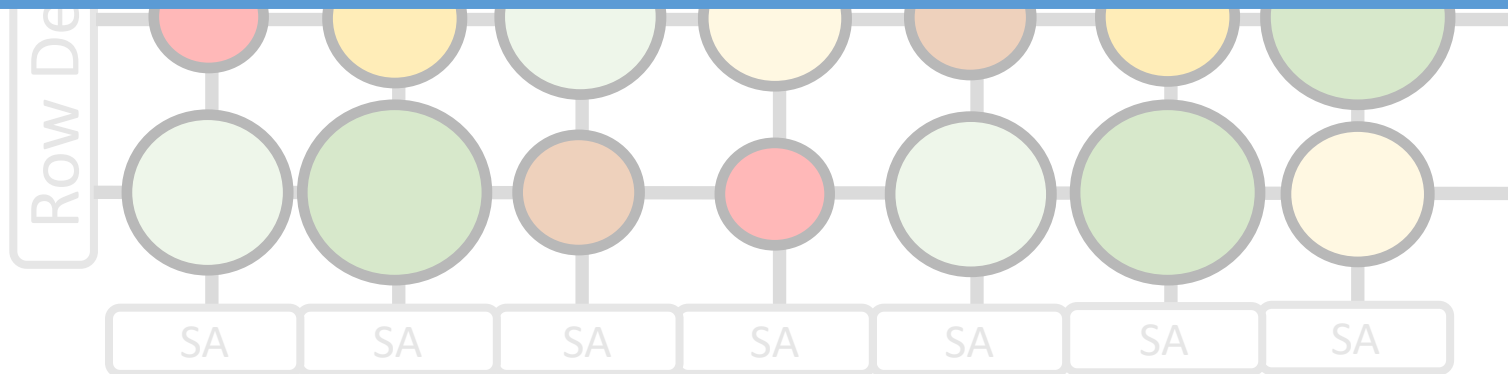
Low % chance to fail
with reduced t_{RCD}



DRAM Latency PUF Key Idea

- A cell's latency failure probability is inherently related to **random process variation** from manufacturing
- We can provide **repeatable and unique device**

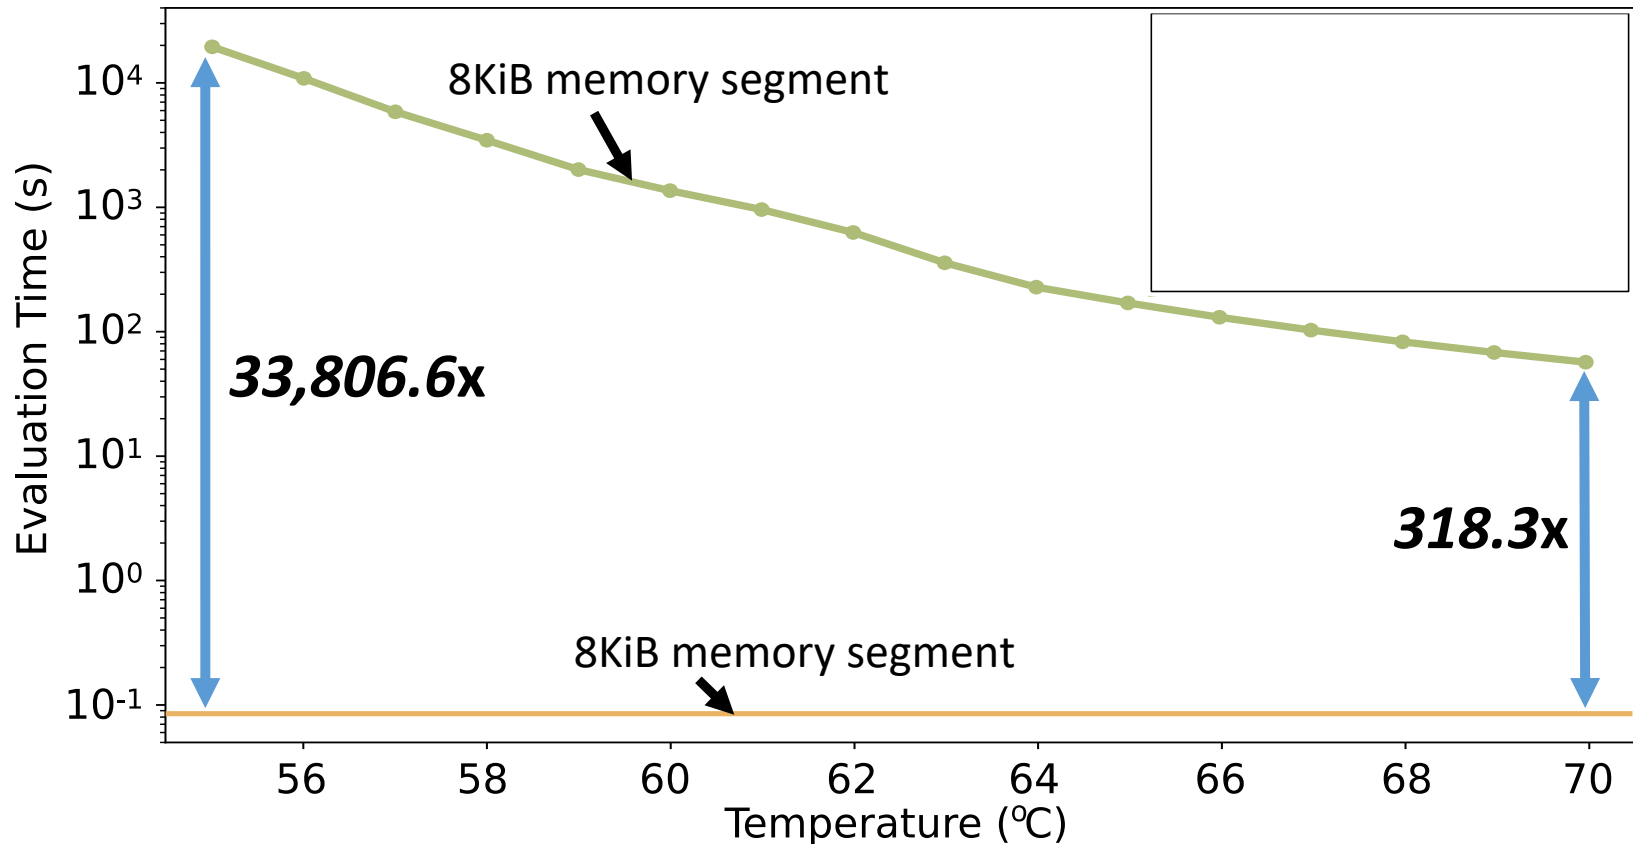
The **key idea** is to compose a PUF response using the DRAM cells that fail with **high probability**



The DRAM Latency PUF Evaluation

- We generate PUF responses using **latency errors** in a region of DRAM
- The latency error patterns **satisfy PUF requirements**
- The DRAM Latency PUF **generates PUF responses in 88.2ms**

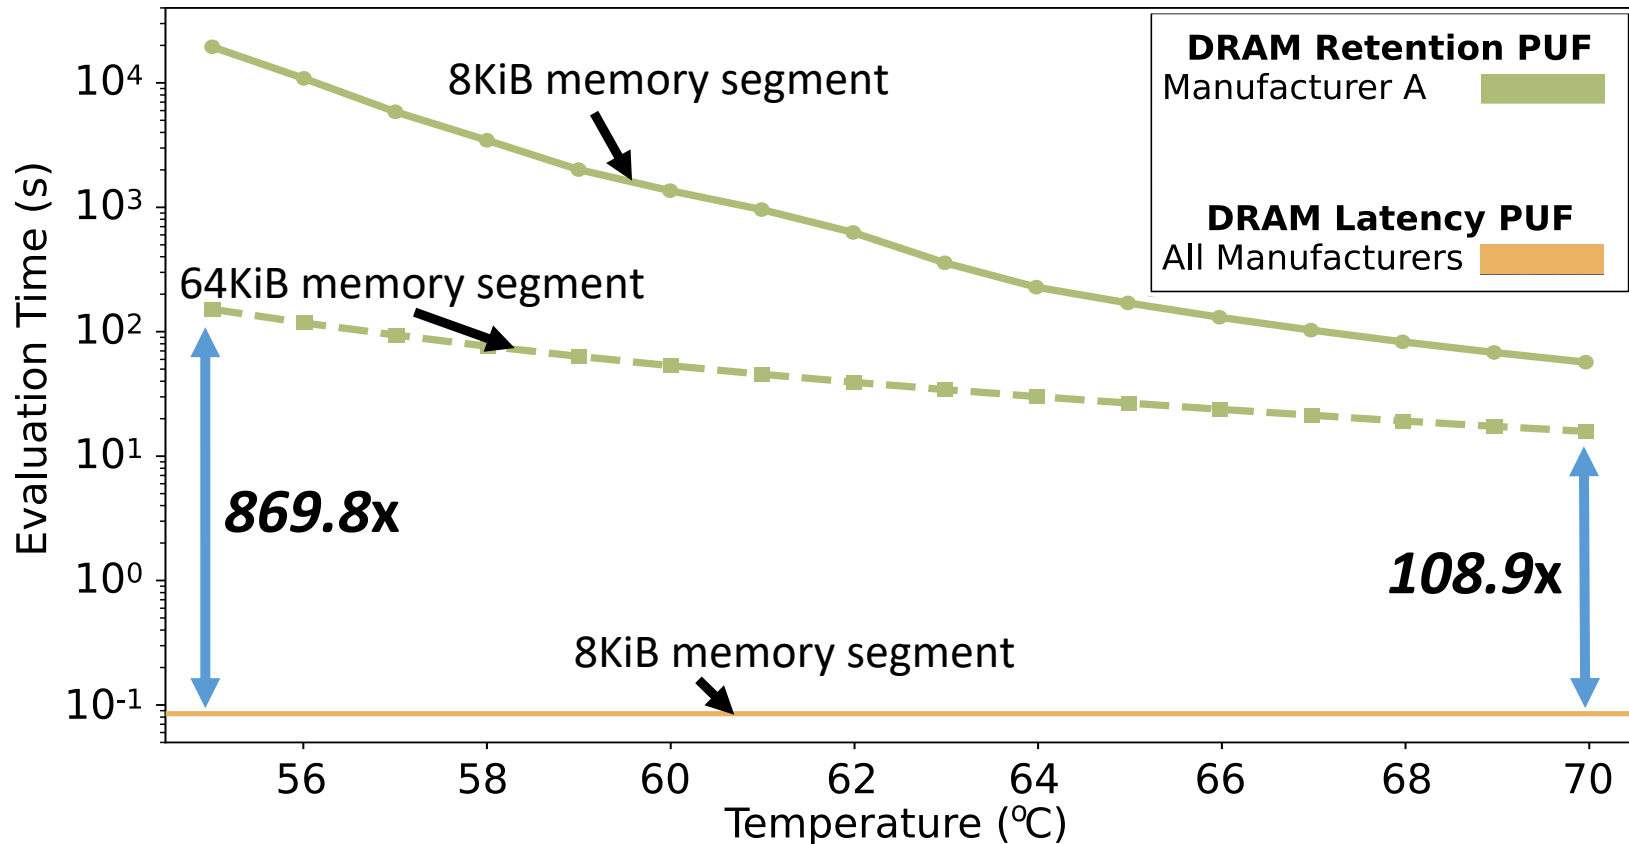
Results – PUF Evaluation Latency



DRAM latency PUF is

1. Fast and constant latency (**88.2ms**)

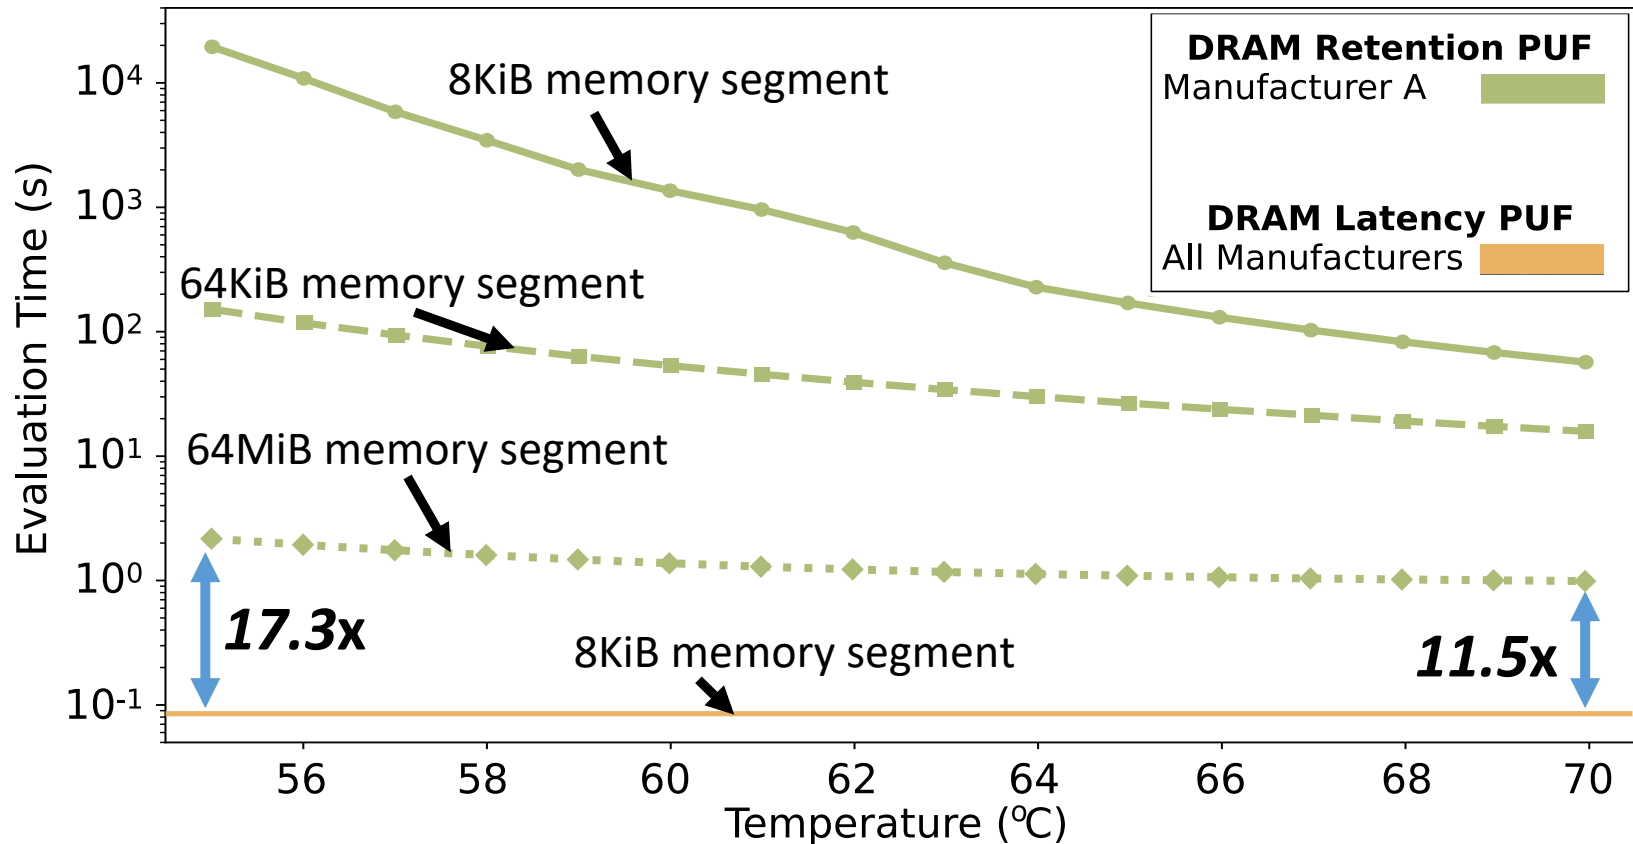
Results – PUF Evaluation Latency



DRAM latency PUF is

1. Fast and constant latency (88.2ms)

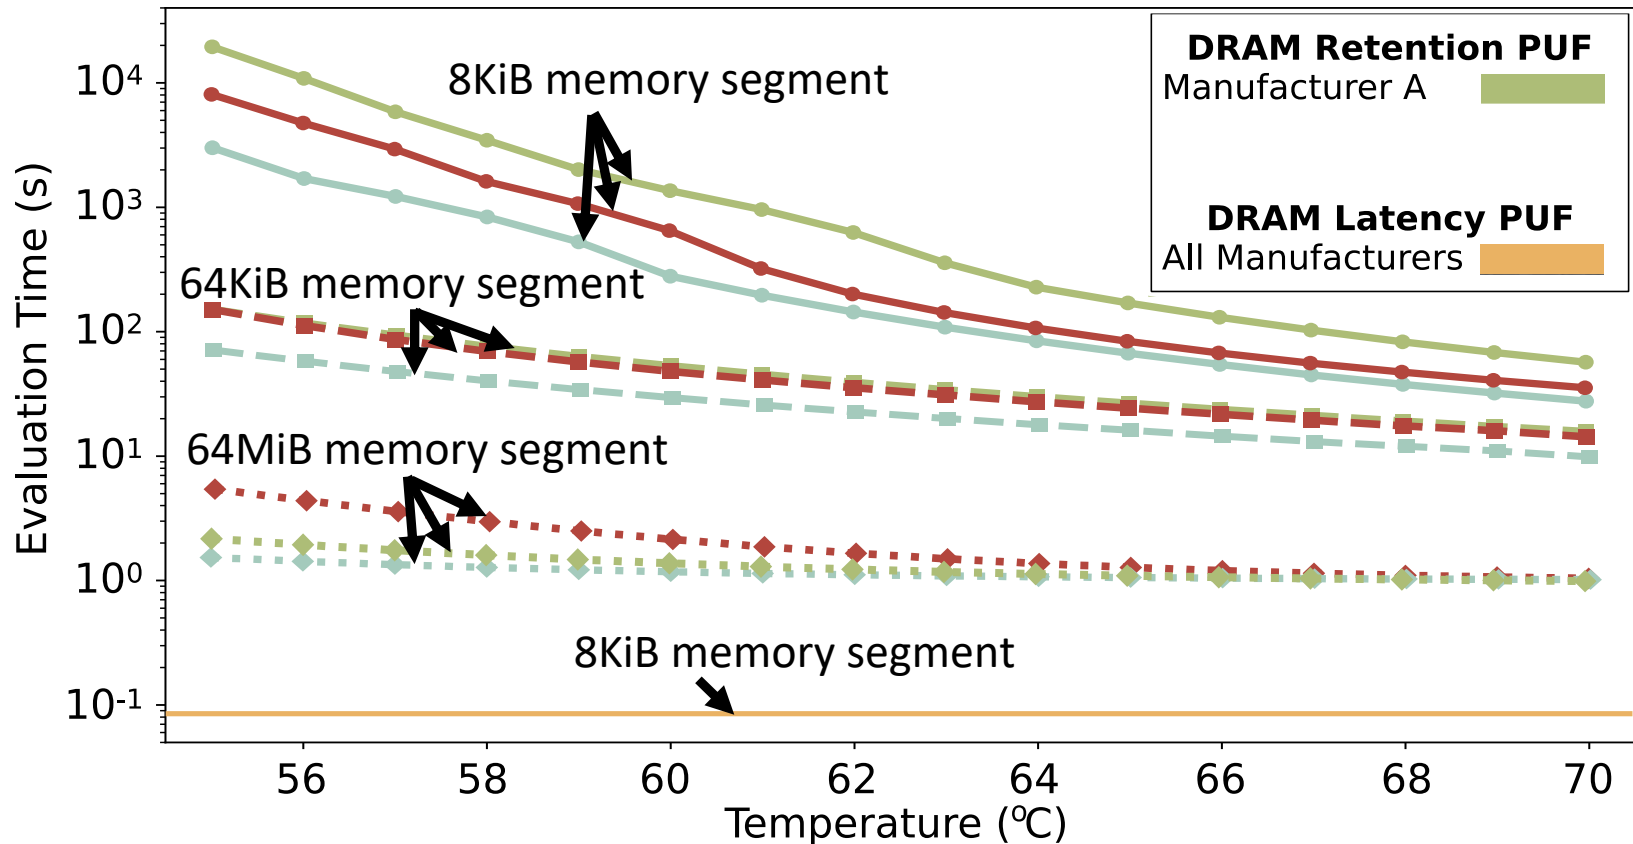
Results – PUF Evaluation Latency



DRAM latency PUF is

1. Fast and constant latency (88.2ms)

Results – PUF Evaluation Latency



DRAM latency PUF is

1. Fast and constant latency (**88.2ms**)
2. On average, **102x/860x** faster than the previous DRAM PUF with the same DRAM capacity overhead (64KiB)

Other Results in the Paper

- How the **DRAM latency PUF** meets the basic requirements for an effective PUF
- A **detailed** analysis on:
 - Devices of **the three major DRAM manufacturers**
 - The **evaluation time** of a PUF
- **Further discussion on:**
 - **Optimizing** retention PUFs
 - **System interference** of DRAM retention and latency PUFs
 - Algorithm to **quickly and reliably** evaluate DRAM latency PUF
 - **Design considerations** for a DRAM latency PUF
 - The DRAM Latency PUF overhead analysis

The DRAM Latency PUF:

Quickly Evaluating Physical Unclonable Functions
by Exploiting the Latency-Reliability Tradeoff
in Modern Commodity DRAM Devices

Jeremie S. Kim Minesh Patel

Hasan Hassan Onur Mutlu



QR Code for the paper

https://people.inf.ethz.ch/omutlu/pub/dram-latency-puf_hpca18.pdf

HPCA 2018

SAFARI



ETH zürich

Carnegie Mellon

More on DRAM Latency PUFs

- Jeremie S. Kim, Minesh Patel, Hasan Hassan, and Onur Mutlu, **"The DRAM Latency PUF: Quickly Evaluating Physical Unclonable Functions by Exploiting the Latency-Reliability Tradeoff in Modern DRAM Devices"**

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

[Lightning Talk Video]

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

The DRAM Latency PUF:

Quickly Evaluating Physical Unclonable Functions

by Exploiting the Latency-Reliability Tradeoff in Modern Commodity DRAM Devices

Jeremie S. Kim^{†§}

Minesh Patel[§]

Hasan Hassan[§]

Onur Mutlu^{§†}

[†]Carnegie Mellon University

[§]ETH Zürich

More on DRAM Latency PUFs

DRAM Latency PUF Key Idea

- A cell's latency failure probability is inherently related to **random process variation** from manufacturing
- We can provide **repeatable and unique device signatures** using latency error patterns

High % chance to fail with reduced t_{RCD}

SAFARI 19/45

ETH ZÜRICH

Computer Architecture - Lecture 11a: DRAM Latency PUF (ETH Zürich, Fall 2019)

449 views • Oct 31, 2019

6 0 SHARE SAVE ...



Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED



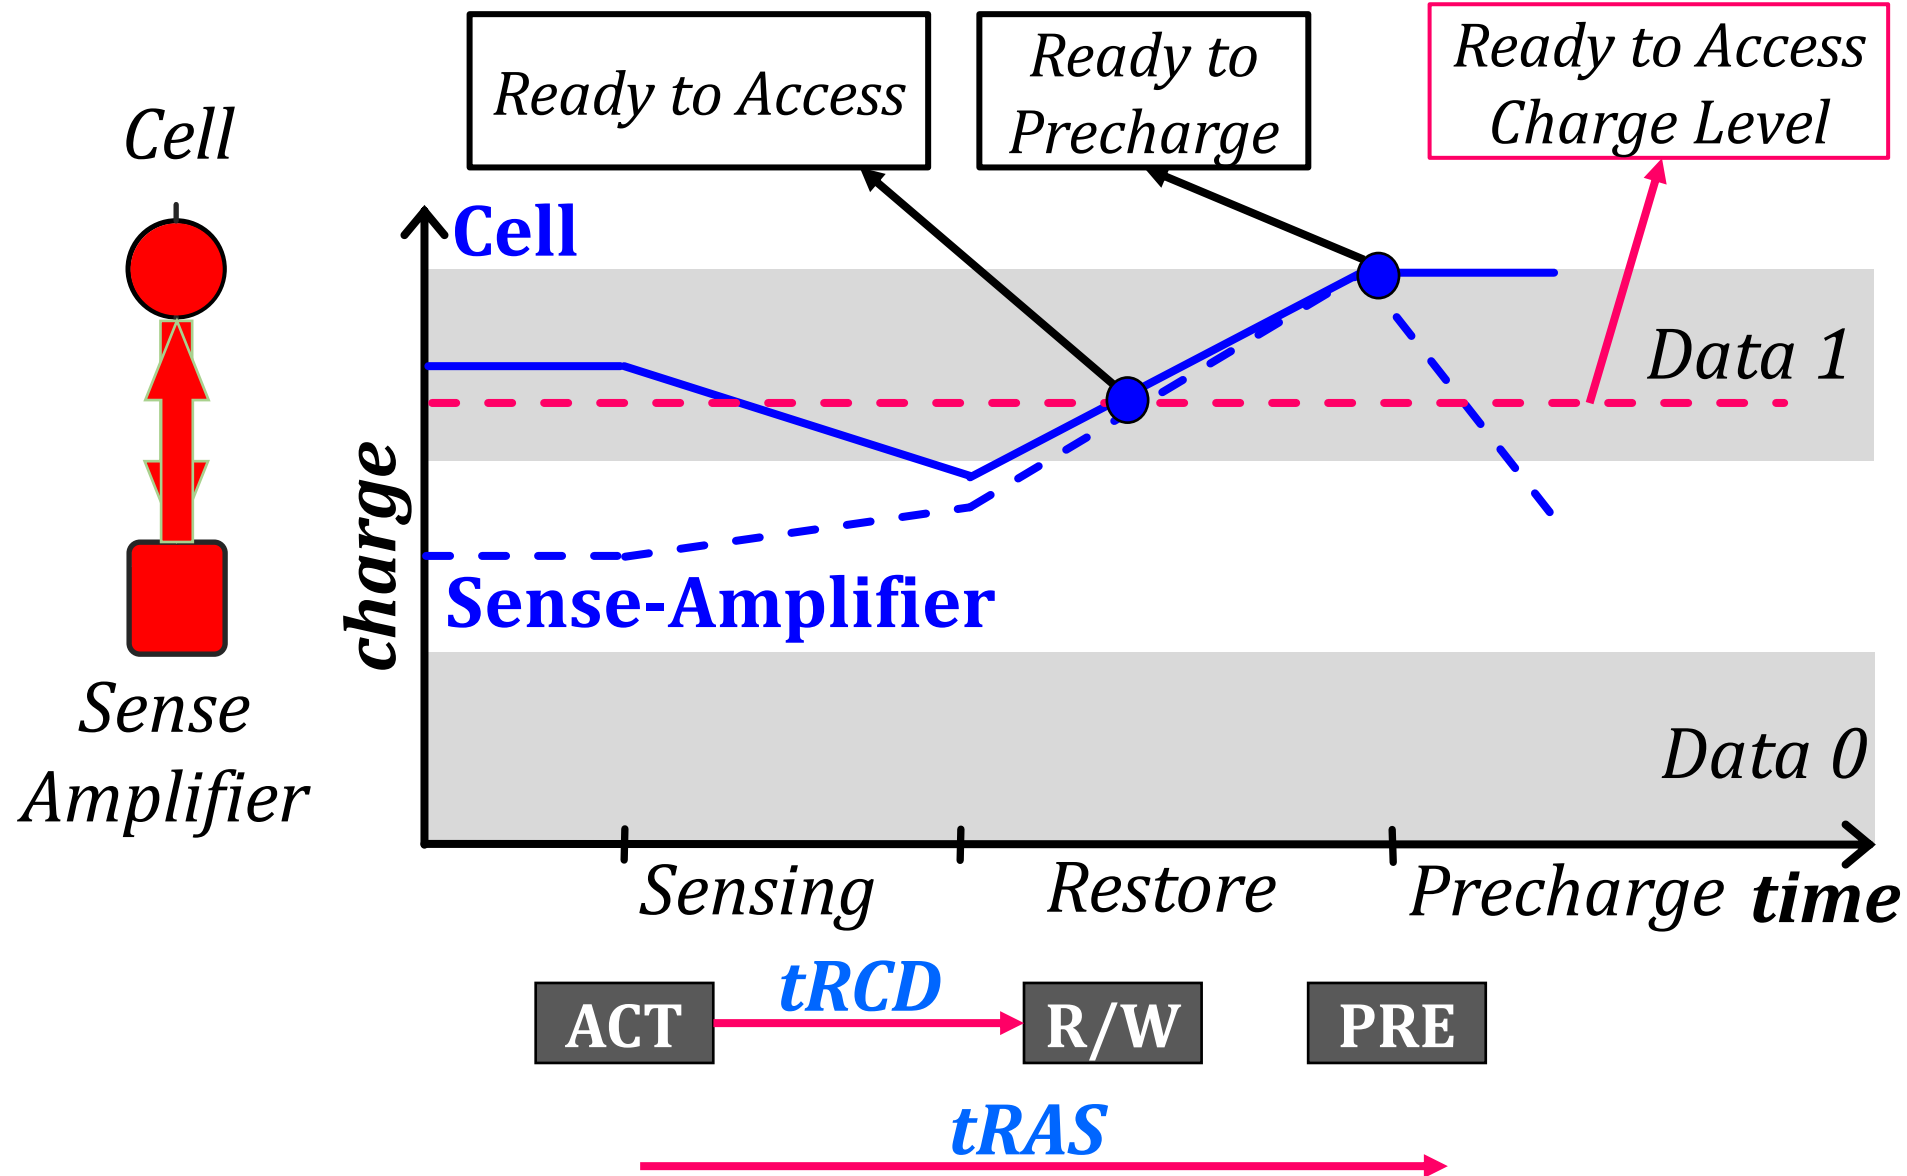
Reducing Memory Latency by Exploiting Memory Access Patterns

ChargeCache: Executive Summary

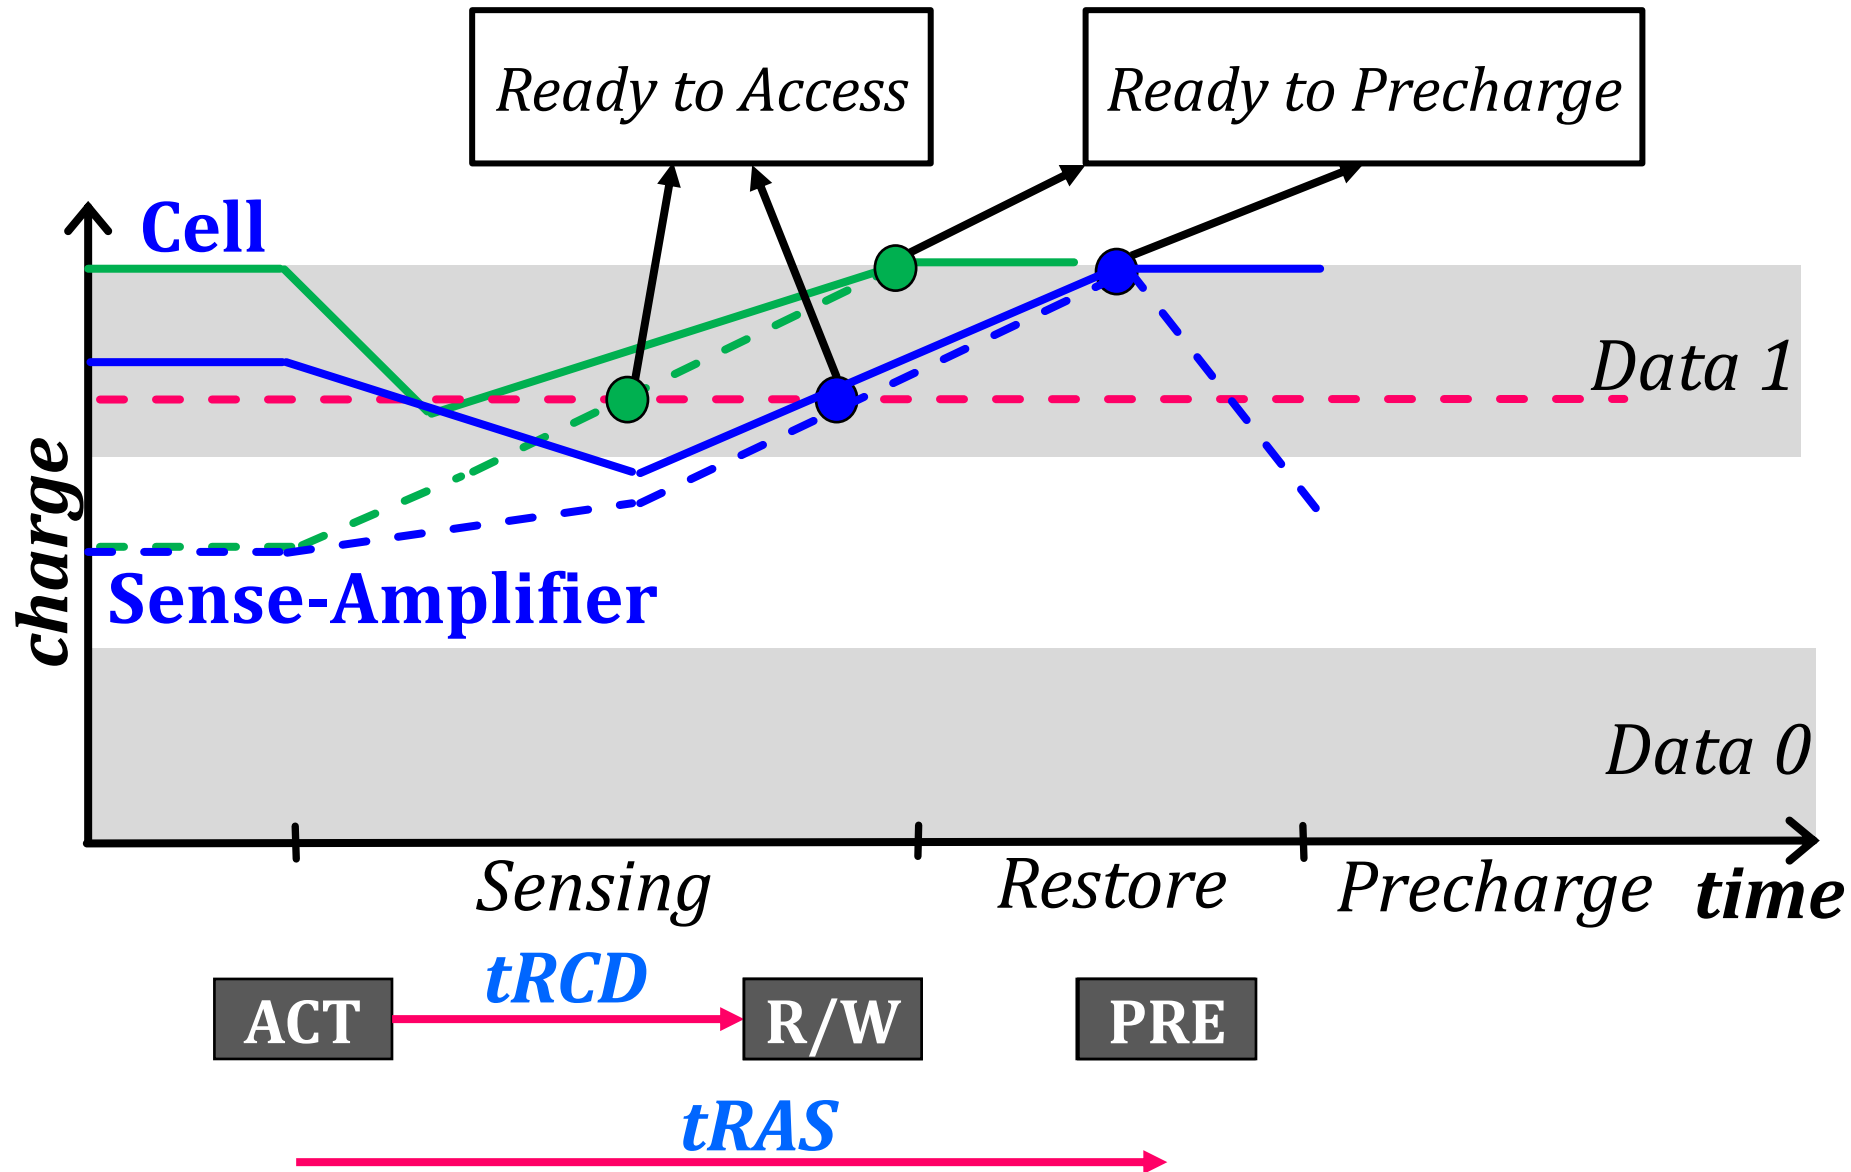
- **Goal**: Reduce average DRAM access latency with no modification to the existing DRAM chips
- **Observations**:
 - 1) A highly-charged DRAM row can be accessed with low latency
 - 2) A row's charge is restored when the row is accessed
 - 3) A recently-accessed row is likely to be accessed again:

Row Level Temporal Locality (RLTL)
- **Key Idea**: Track recently-accessed DRAM rows and use lower timing parameters if such rows are accessed again
- **ChargeCache**:
 - Low cost & no modifications to the DRAM
 - Higher performance (**8.6-10.6%** on average for 8-core)
 - Lower DRAM energy (**7.9%** on average)

DRAM Charge over Time



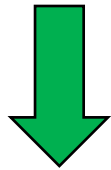
Accessing Highly-charged Rows



Observation 1

A **highly-charged** DRAM row can be accessed with **low latency**

- tRCD: 44%
- tRAS: 37%



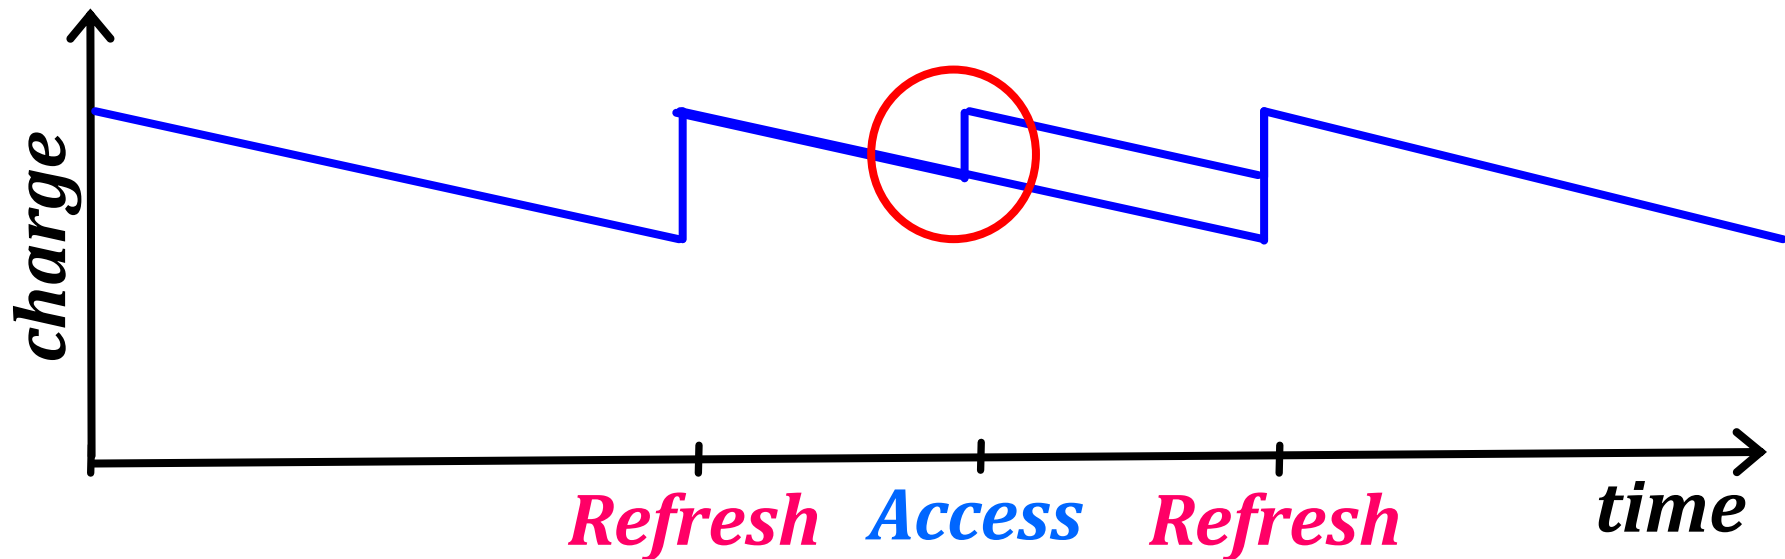
**How does a row become
highly-charged?**

How Does a Row Become Highly-Charged?

DRAM cells **lose charge** over time

Two ways of restoring a row's charge:

- Refresh Operation
- Access



Observation 2

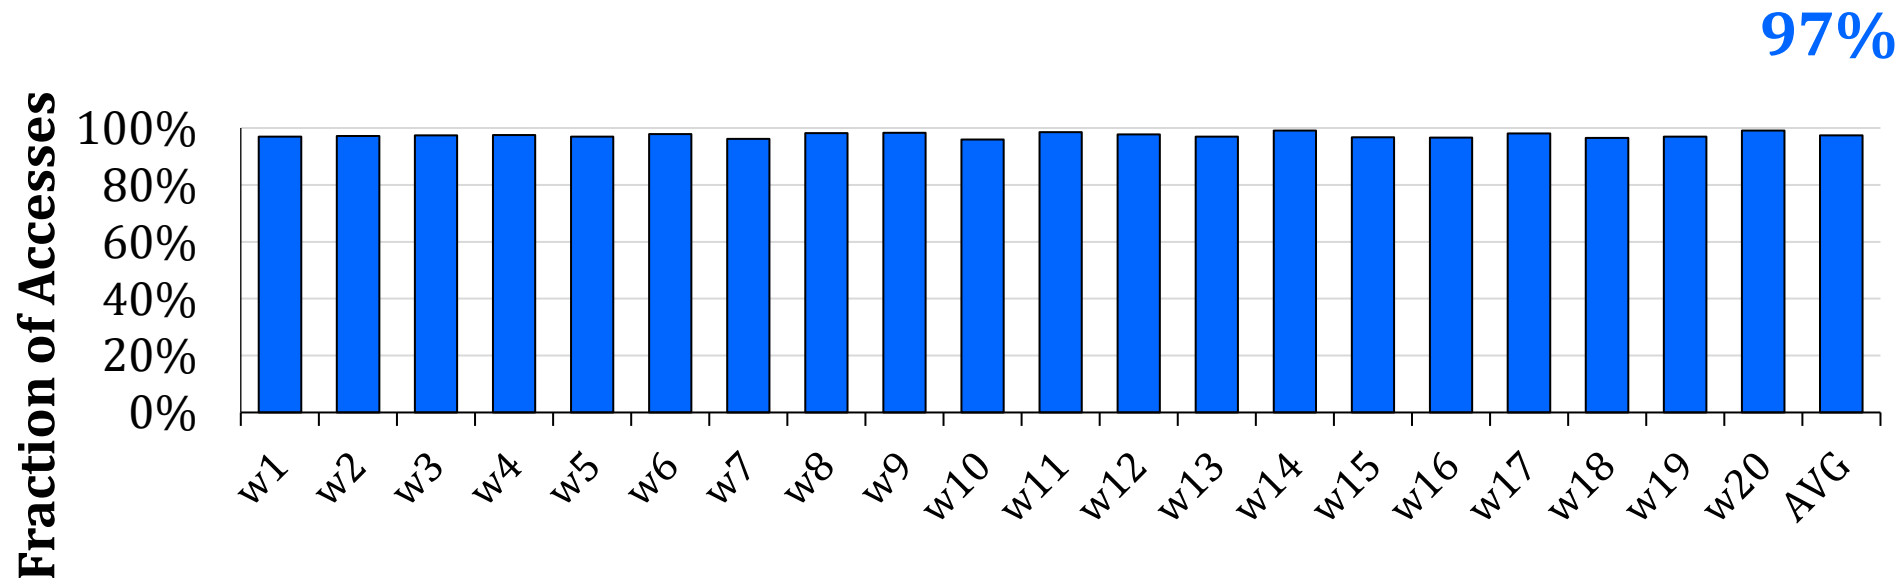
A row's charge is **restored** when the row is **accessed**

How likely is a **recently-accessed row to be accessed again?**

Row Level Temporal Locality (RLTL)

A **recently-accessed** DRAM row is likely to be accessed again.

- t -RLTL: Fraction of rows that are accessed within time t after their previous access

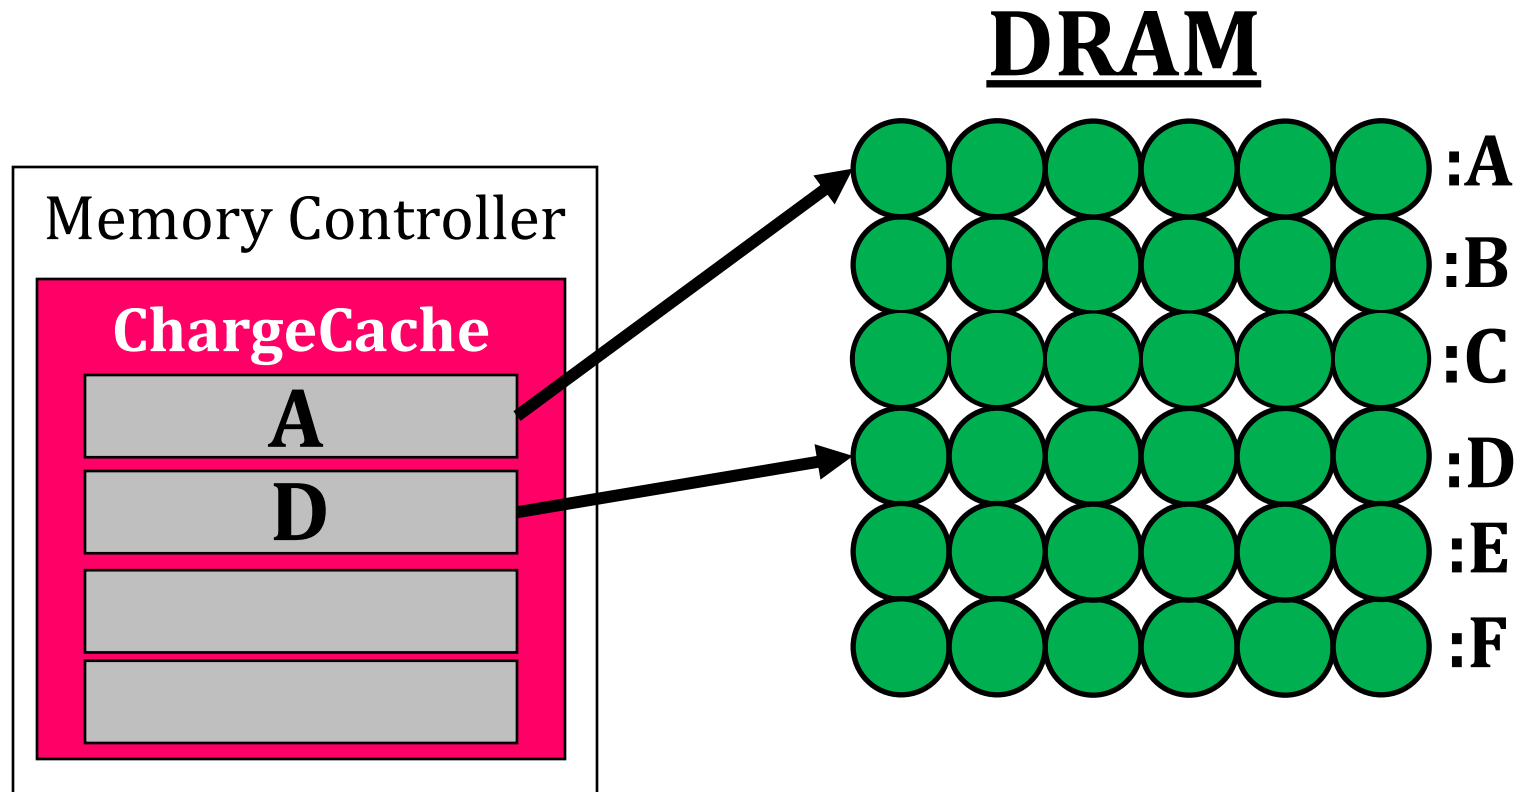


88ns — RLTL for eight-core workloads

Key Idea

Track **recently-accessed** DRAM rows and use **lower timing parameters** if such rows are accessed again

ChargeCache Overview



Requests: A D A 

ChargeCache Hits: Use Default Timings

Area and Power Overhead

- Modeled with CACTI

- Area

- ~5KB for 128-entry ChargeCache
- 0.24% of a 4MB Last Level Cache (LLC) area

- Power Consumption

- 0.15 mW on average (static + dynamic)
- 0.23% of the 4MB LLC power consumption

Methodology

- **Simulator**

- DRAM Simulator (Ramulator *[Kim+, CAL'15]*)
<https://github.com/CMU-SAFARI/ramulator>

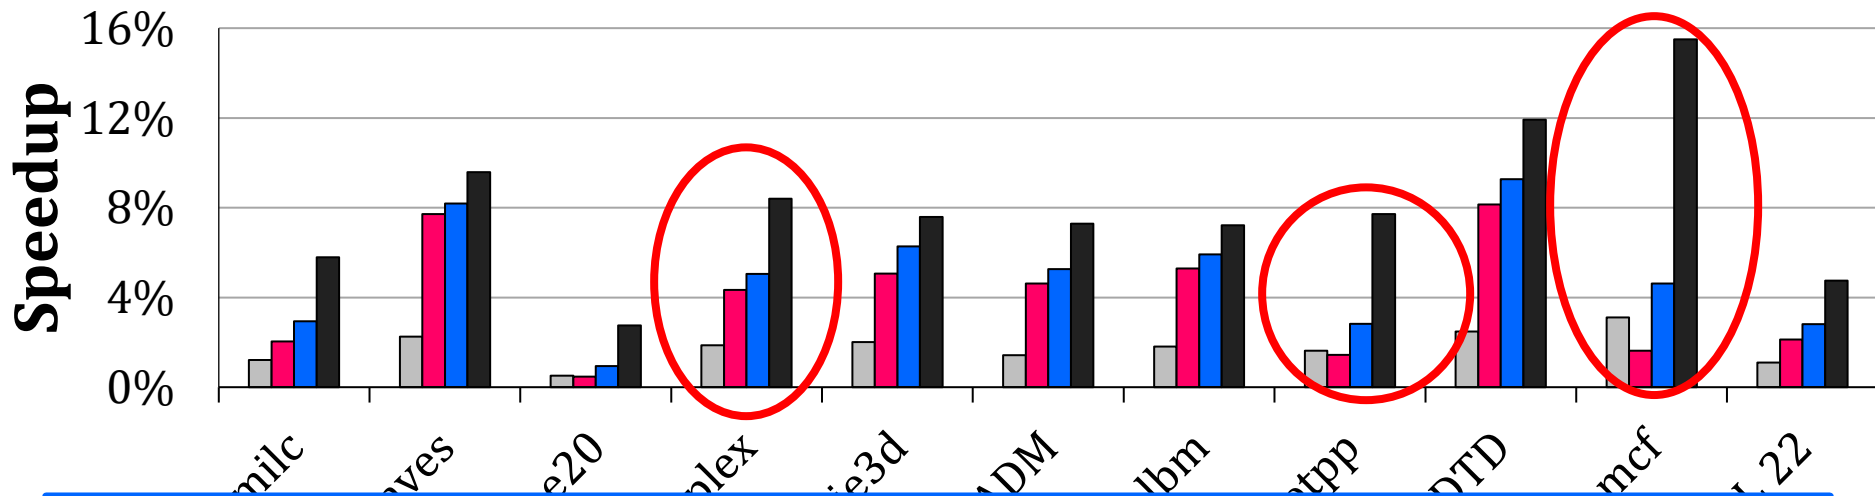
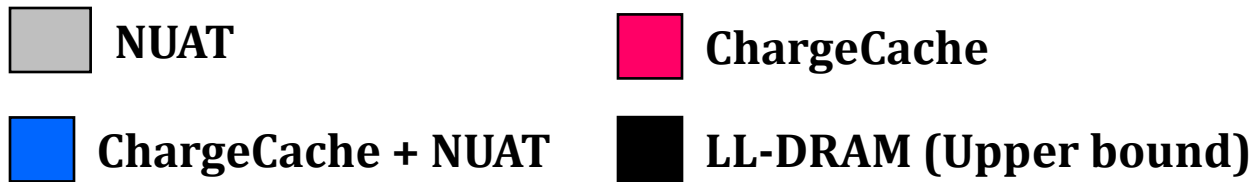
- **Workloads**

- 22 single-core workloads
 - SPEC CPU2006, TPC, STREAM
- 20 multi-programmed 8-core workloads
 - By randomly choosing from single-core workloads
- Execute at least 1 billion representative instructions per core (Pinpoints)

- **System Parameters**

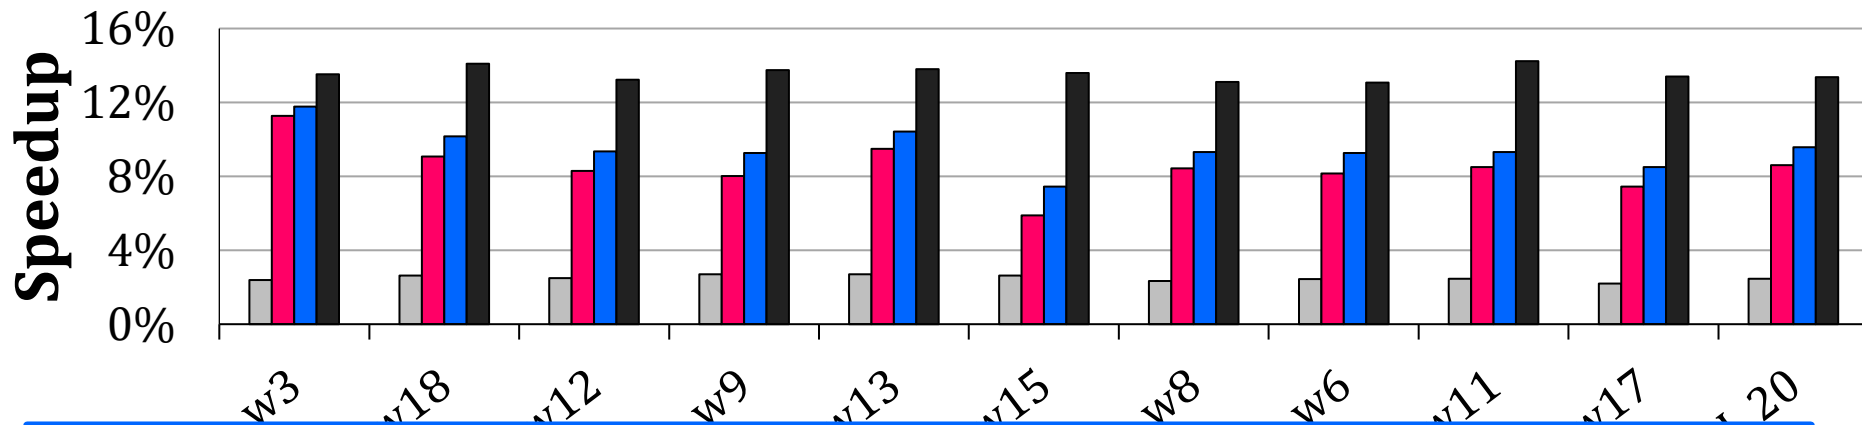
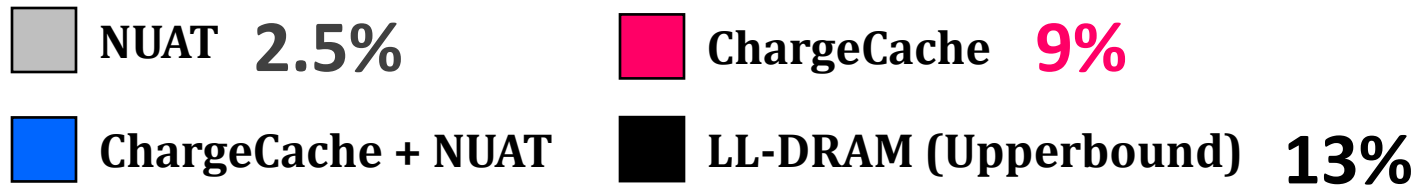
- 1/8 core system with 4MB LLC
- Default tRCD/tRAS of 11/28 cycles

Single-core Performance



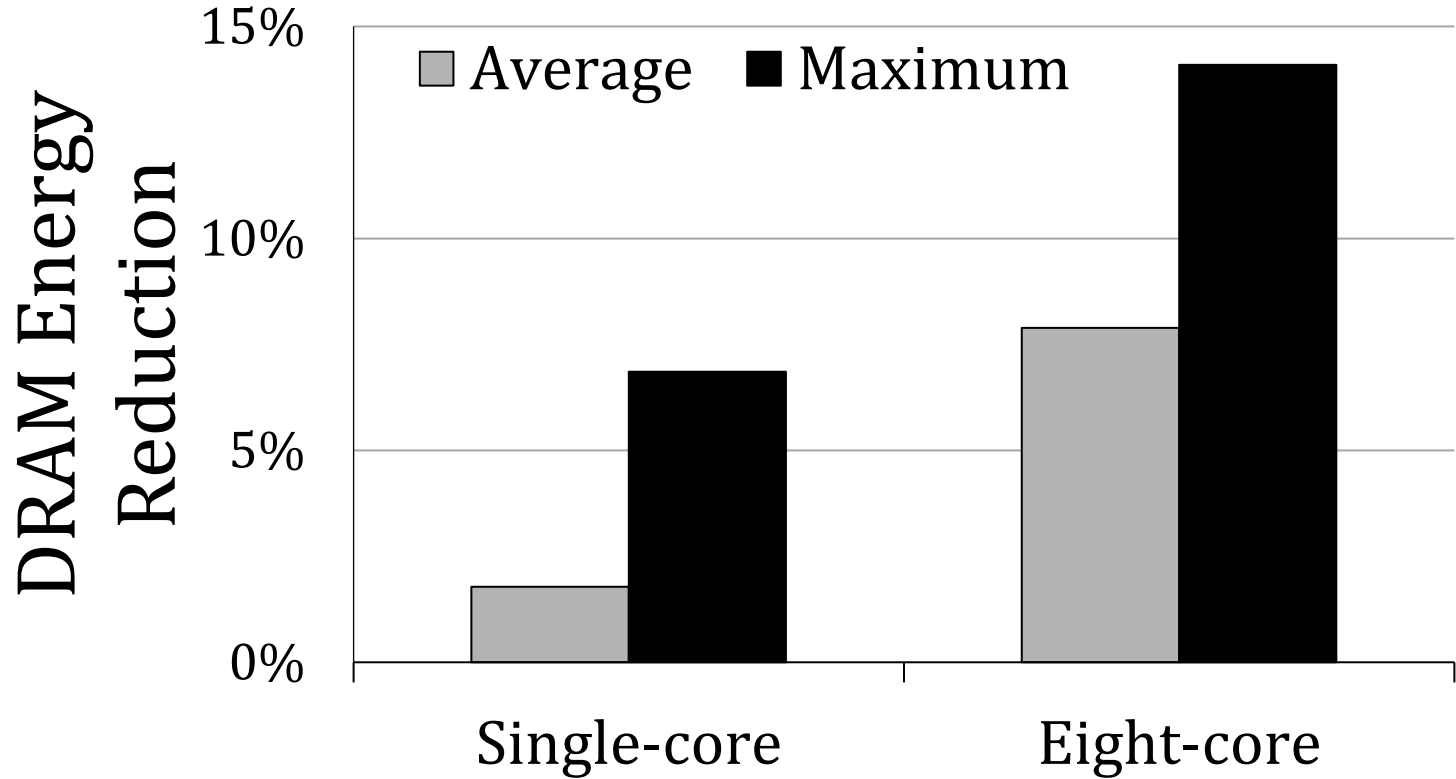
**ChargeCache improves
single-core performance**

Eight-core Performance



ChargeCache significantly improves multi-core performance

DRAM Energy Savings



ChargeCache reduces DRAM energy

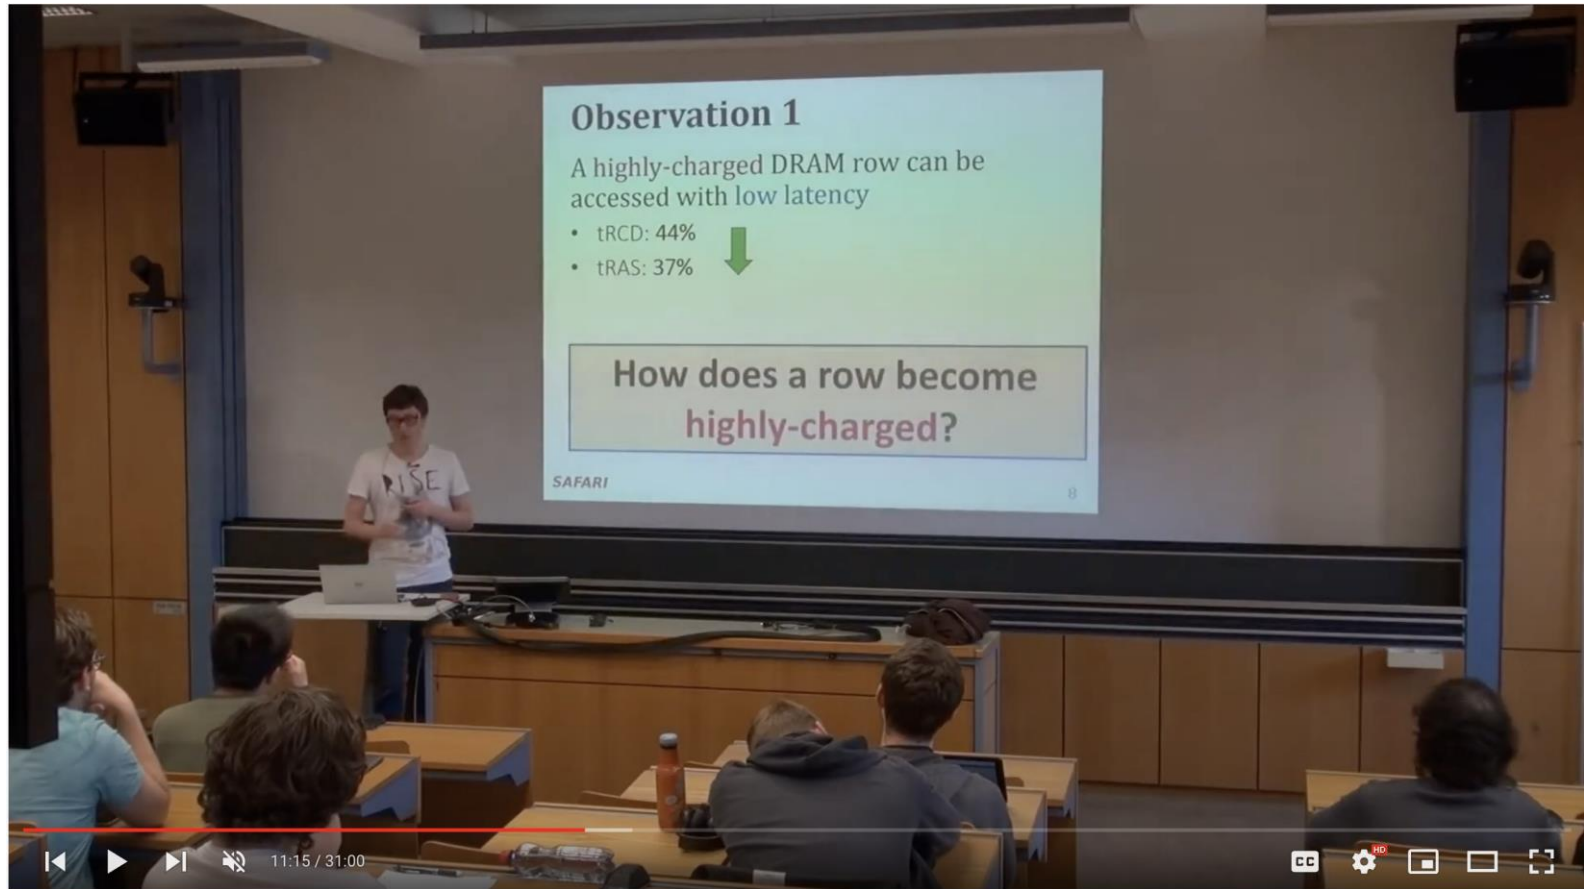
More on ChargeCache

- Hasan Hassan, Gennady Pekhimenko, Nandita Vijaykumar, Vivek Seshadri, Donghyuk Lee, Oguz Ergin, and Onur Mutlu,
"ChargeCache: Reducing DRAM Latency by Exploiting Row Access Locality"
Proceedings of the 22nd International Symposium on High-Performance Computer Architecture (HPCA), Barcelona, Spain, March 2016.
[[Slides \(pptx\)](#)] [[pdf](#)]
[[Source Code](#)]

ChargeCache: Reducing DRAM Latency by Exploiting Row Access Locality

Hasan Hassan^{†*}, Gennady Pekhimenko[†], Nandita Vijaykumar[†]
Vivek Seshadri[†], Donghyuk Lee[†], Oguz Ergin^{*}, Onur Mutlu[†]

More on ChargeCache



ETH ZÜRICH HAUPTGEBÄUDE

Computer Architecture - Lecture 6a: ChargeCache: Reducing DRAM Latency (ETH Zürich, Fall 2018)

519 views • Oct 10, 2018

9 0 SHARE SAVE ...



Onur Mutlu Lectures
19.7K subscribers

SUBSCRIBED



Partial Restoration of Cell Charge

- Yaohua Wang, Arash Tavakkol, Lois Orosa, Saugata Ghose, Nika Mansouri Ghiasi, Minesh Patel, Jeremie S. Kim, Hasan Hassan, Mohammad Sadrosadati, and Onur Mutlu,
"Reducing DRAM Latency via Charge-Level-Aware Look-Ahead Partial Restoration"
Proceedings of the 51st International Symposium on Microarchitecture (MICRO), Fukuoka, Japan, October 2018.

Reducing DRAM Latency via Charge-Level-Aware Look-Ahead Partial Restoration

Yaohua Wang^{†§} Arash Tavakkol[†] Lois Orosa^{†*} Saugata Ghose[‡] Nika Mansouri Ghiasi[†]
Minesh Patel[†] Jeremie S. Kim^{‡†} Hasan Hassan[†] Mohammad Sadrosadati[†] Onur Mutlu^{‡†}

[†]*ETH Zürich* [§]*National University of Defense Technology*

[‡]*Carnegie Mellon University* ^{*}*University of Campinas*

On DRAM Power Consumption

VAMPIRE DRAM Power Model

- Saugata Ghose, A. Giray Yaglikci, Raghav Gupta, Donghyuk Lee, Kais Kudrolli, William X. Liu, Hasan Hassan, Kevin K. Chang, Niladrish Chatterjee, Aditya Agrawal, Mike O'Connor, and Onur Mutlu,

"What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study"

*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Irvine, CA, USA, June 2018.*

[[Abstract](#)]

[[POMACS Journal Version \(same content, different format\)](#)]

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

[[VAMPIRE DRAM Power Model](#)]

What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study

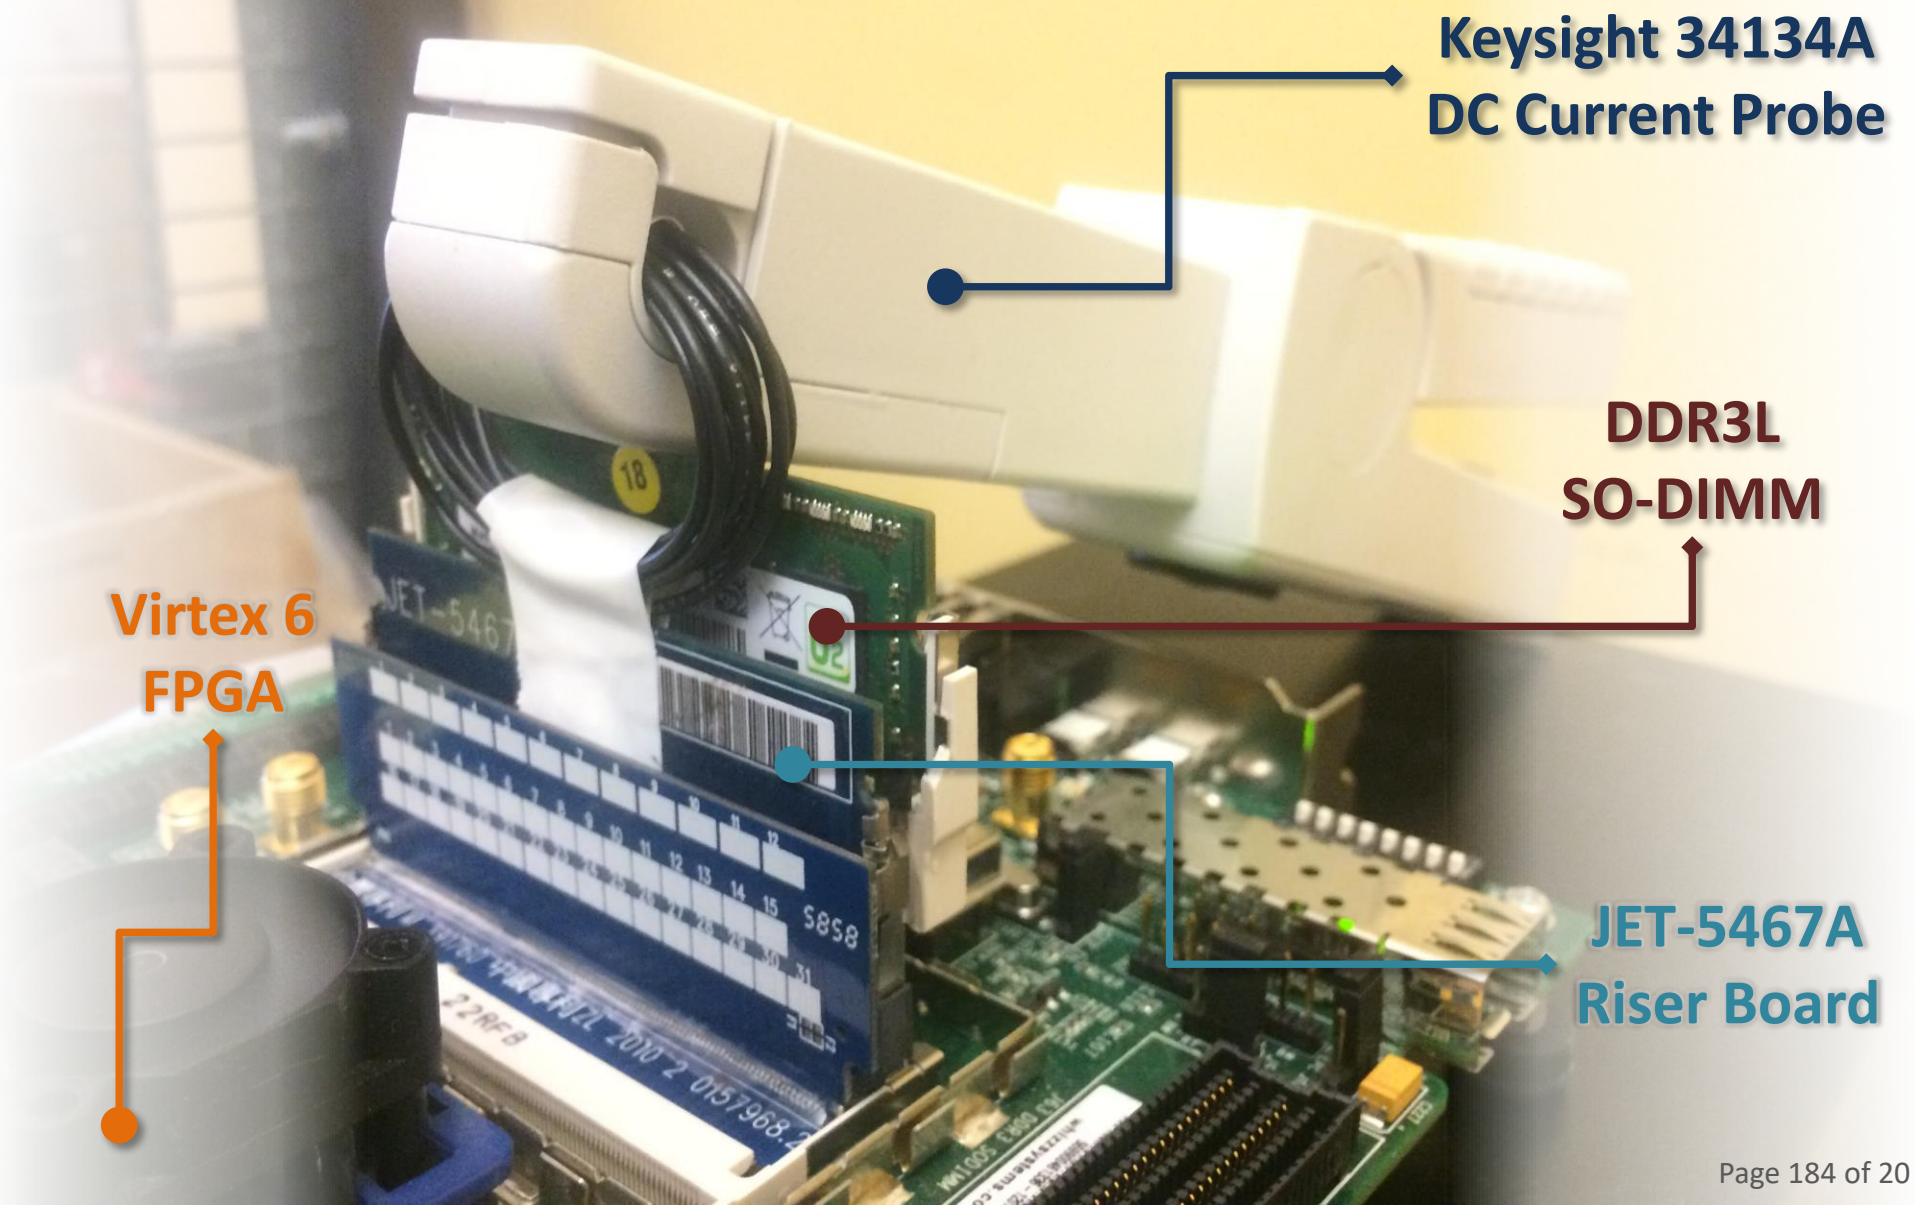
Saugata Ghose [†]	Abdullah Giray Yağlıkçı ^{‡†}	Raghav Gupta [†]	Donghyuk Lee [§]
Kais Kudrolli [†]	William X. Liu [†]	Hasan Hassan [‡]	Kevin K. Chang [†]
Niladrish Chatterjee [§]	Aditya Agrawal [§]	Mike O'Connor ^{§¶}	Onur Mutlu ^{‡†}

[†]Carnegie Mellon University

[‡]ETH Zürich

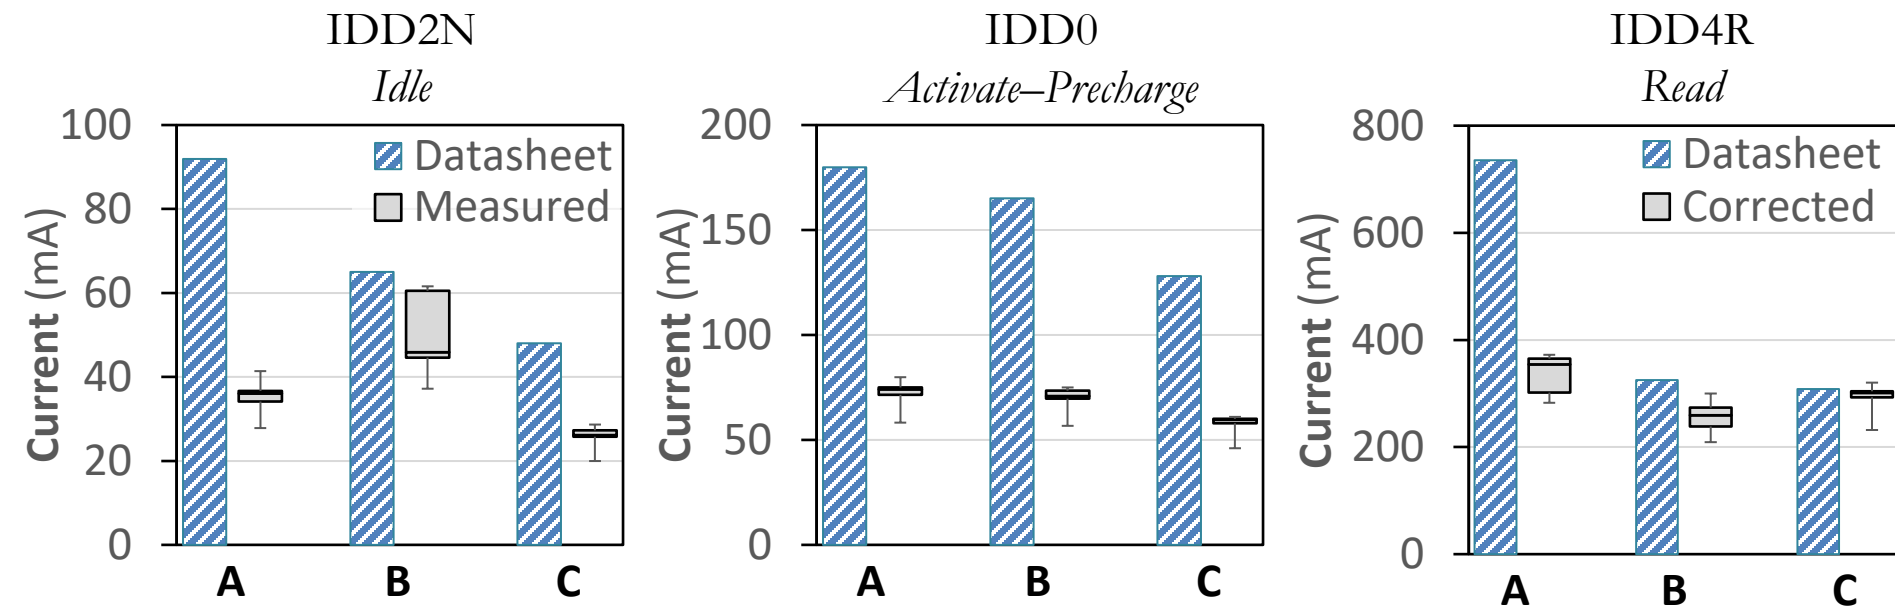
[§]NVIDIA

[¶]University of Texas at Austin



- **SoftMC: an FPGA-based memory controller** [Hassan+ HPCA '17]
 - Modified to repeatedly loop commands
 - Open-source: <https://github.com/CMU-SAFARI/SoftMC>
- **Measure current consumed by a module during a SoftMC test**
- **Tested 50 DDR3L DRAM modules** (200 DRAM chips)
 - Supply voltage: 1.35 V
 - **Three major vendors: A, B, C**
 - Manufactured between 2014 and 2016
- **For each experimental test that we perform**
 - 10 runs of each test per module
 - At least 10 current samples per run

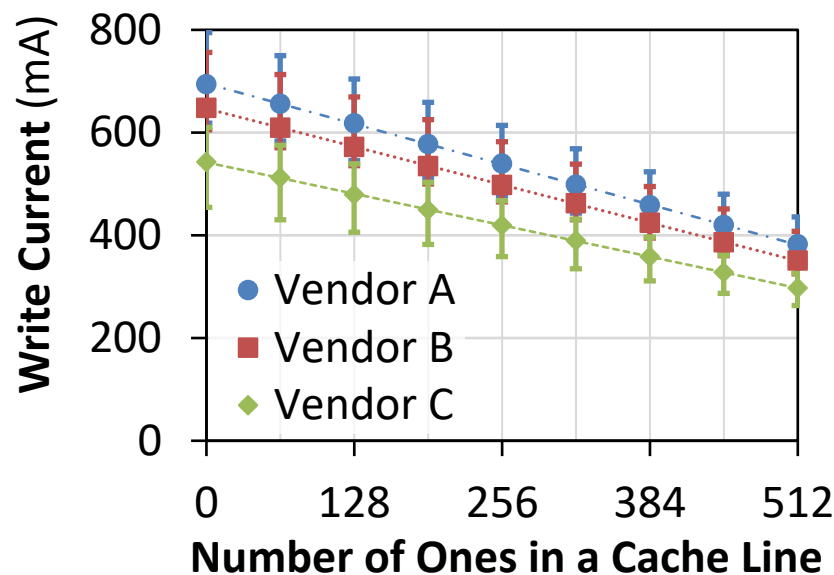
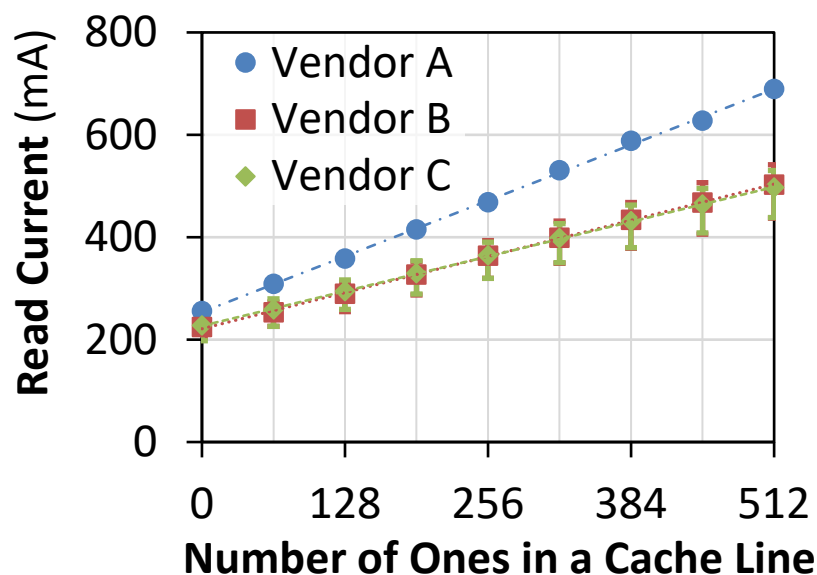
1. Real DRAM Power Varies Widely from IDD Values **SAFARI**



- Different vendors have very different margins (i.e., *guardbands*)
- Low variance among different modules from same vendor

Current consumed by real DRAM modules varies significantly for all IDD values that we measure

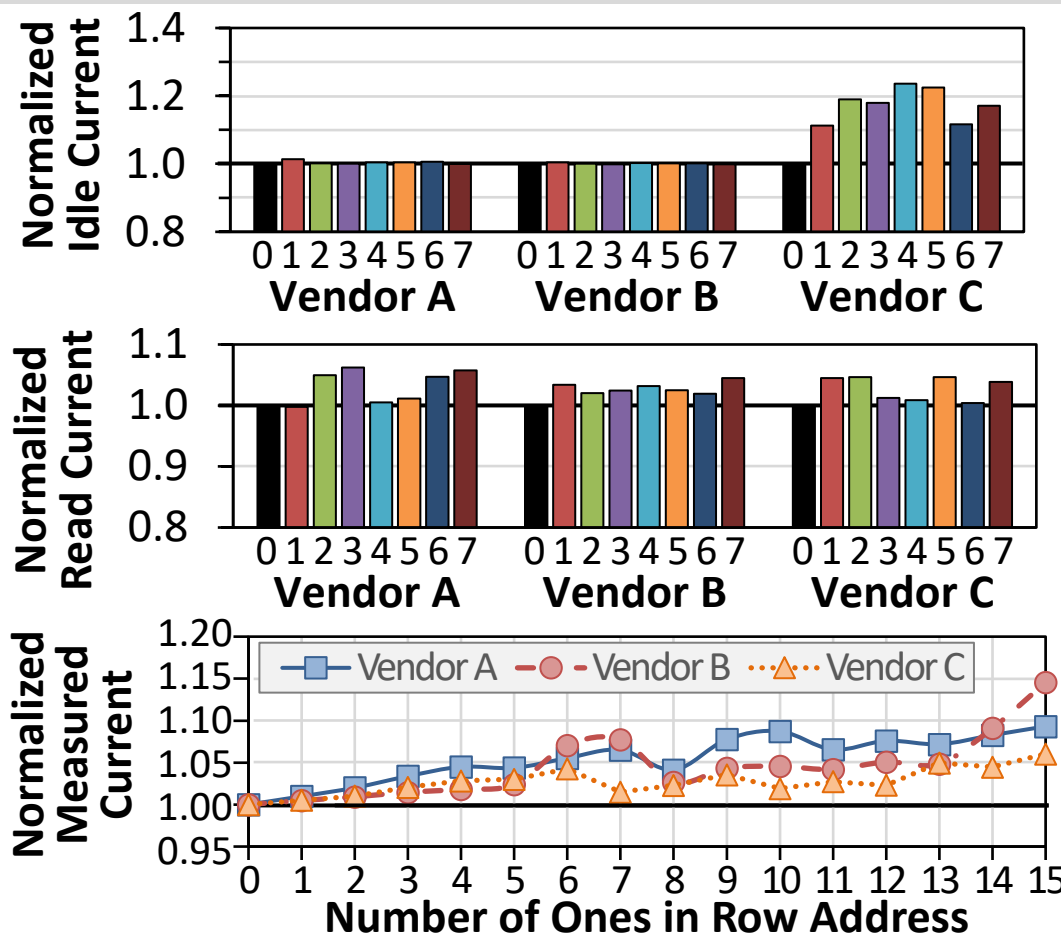
2. DRAM Power is Dependent on Data Values



- Some variation due to infrastructure – can be subtracted
- Without infrastructure variation: up to 230 mA of change
- Toggle affects power consumption, but < 0.15 mA per bit

DRAM power consumption depends *strongly* on the data value

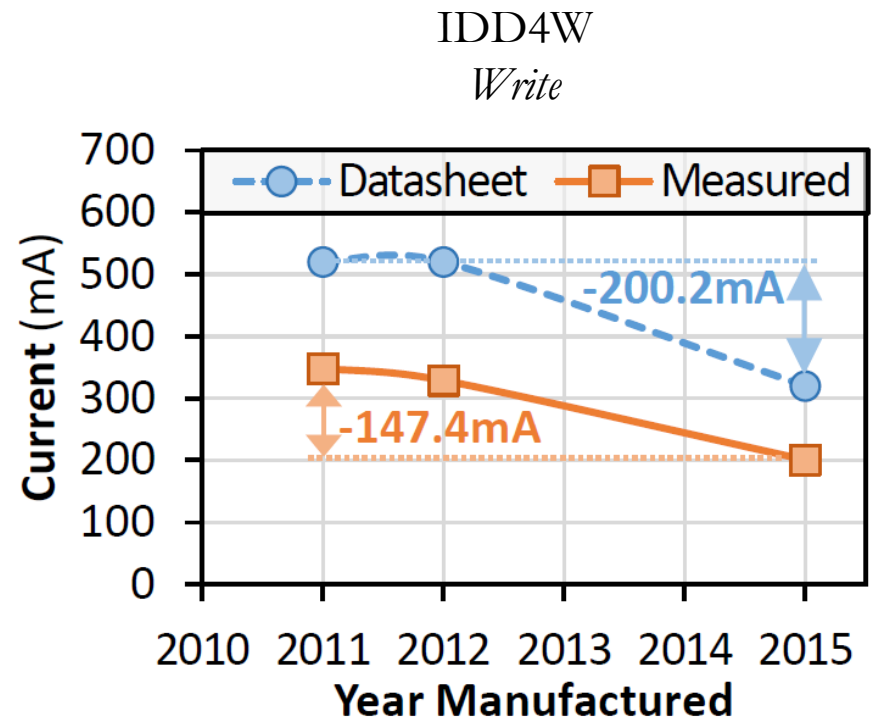
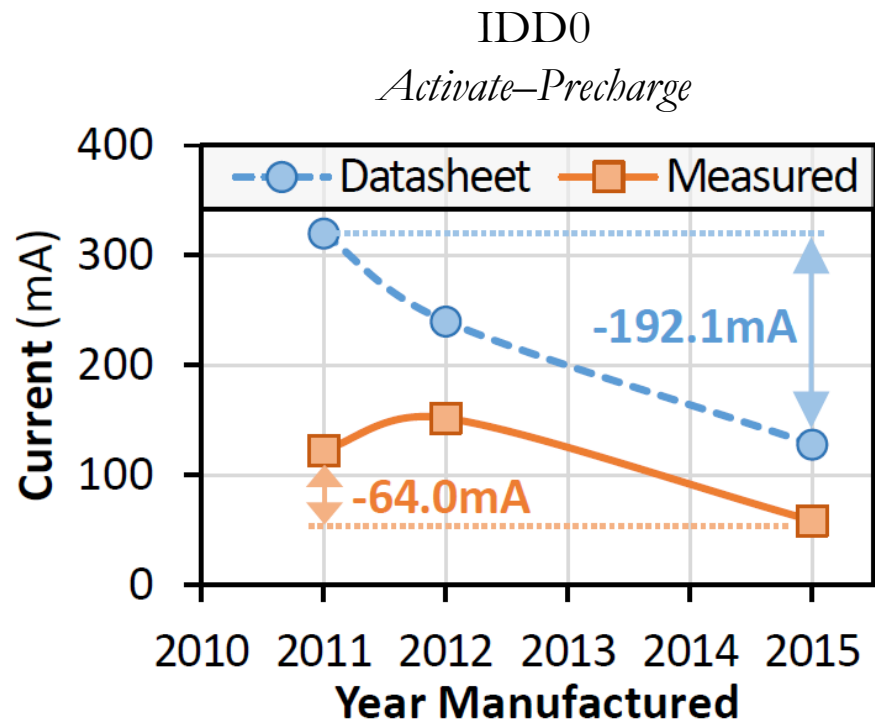
3. Structural Variation Affects DRAM Power Usage **SAFARI**



- Vendor C: variation in idle current across banks
- All vendors: variation in read current across banks
- All vendors: variation in activation based on

Significant structural variation:
DRAM power varies systematically by bank and row

4. Generational Savings Are Smaller Than Expected **SAFARI**



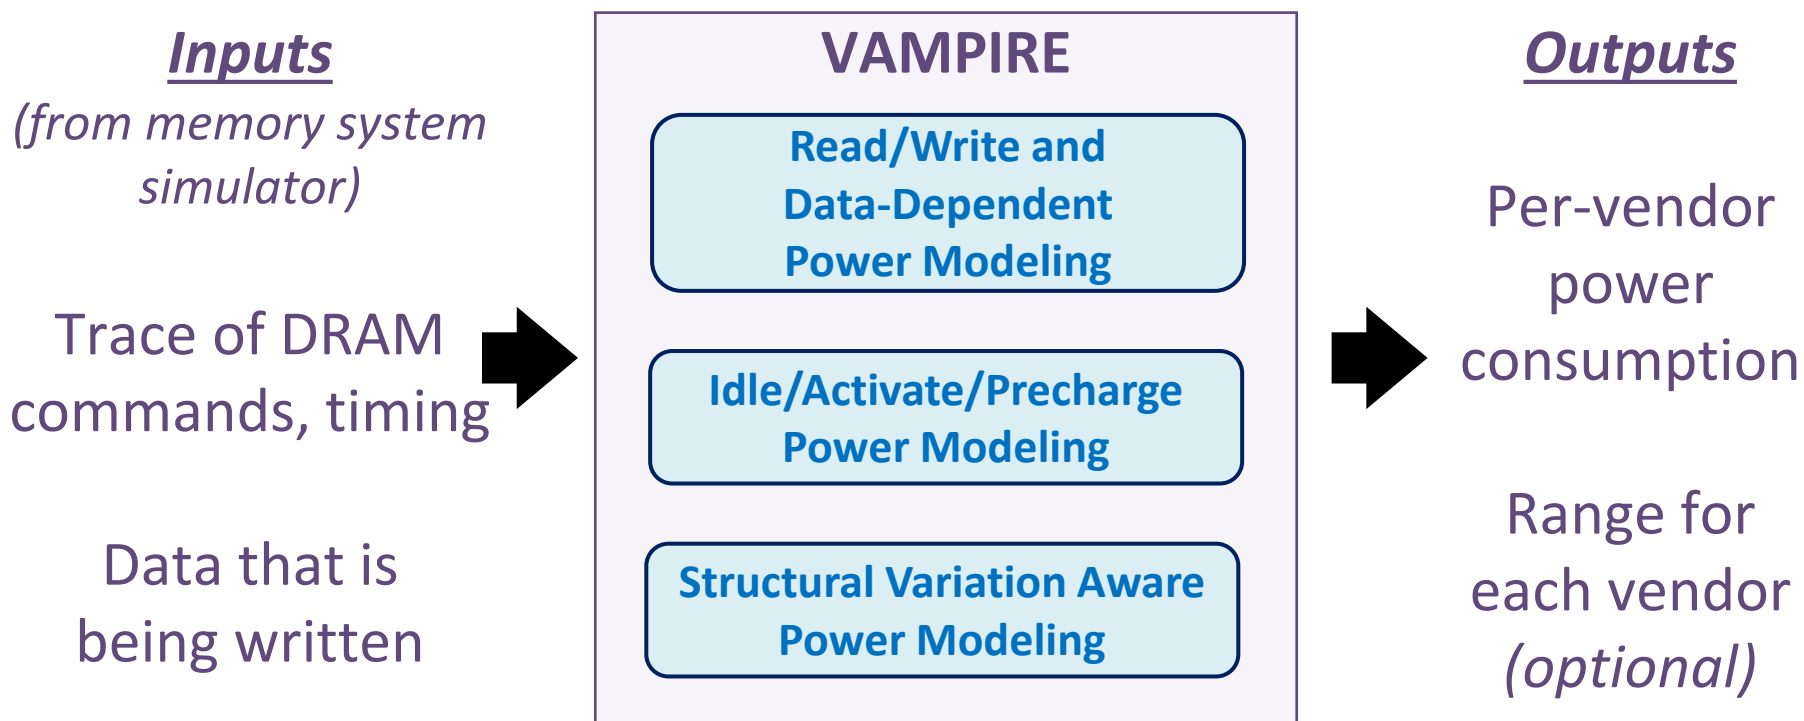
- Similar trends for idle and read currents

Actual power savings of newer DRAM is *much lower* than the savings indicated in the datasheets

1. Real DRAM modules often **consume less power** than vendor-provided IDD values state
2. DRAM power consumption is **dependent on the data value** that is read/written
3. Across banks and rows, **structural variation affects power consumption of DRAM**
4. **Newer DRAM modules save less power** than indicated in datasheets by vendors

Detailed observations and analyses in the paper

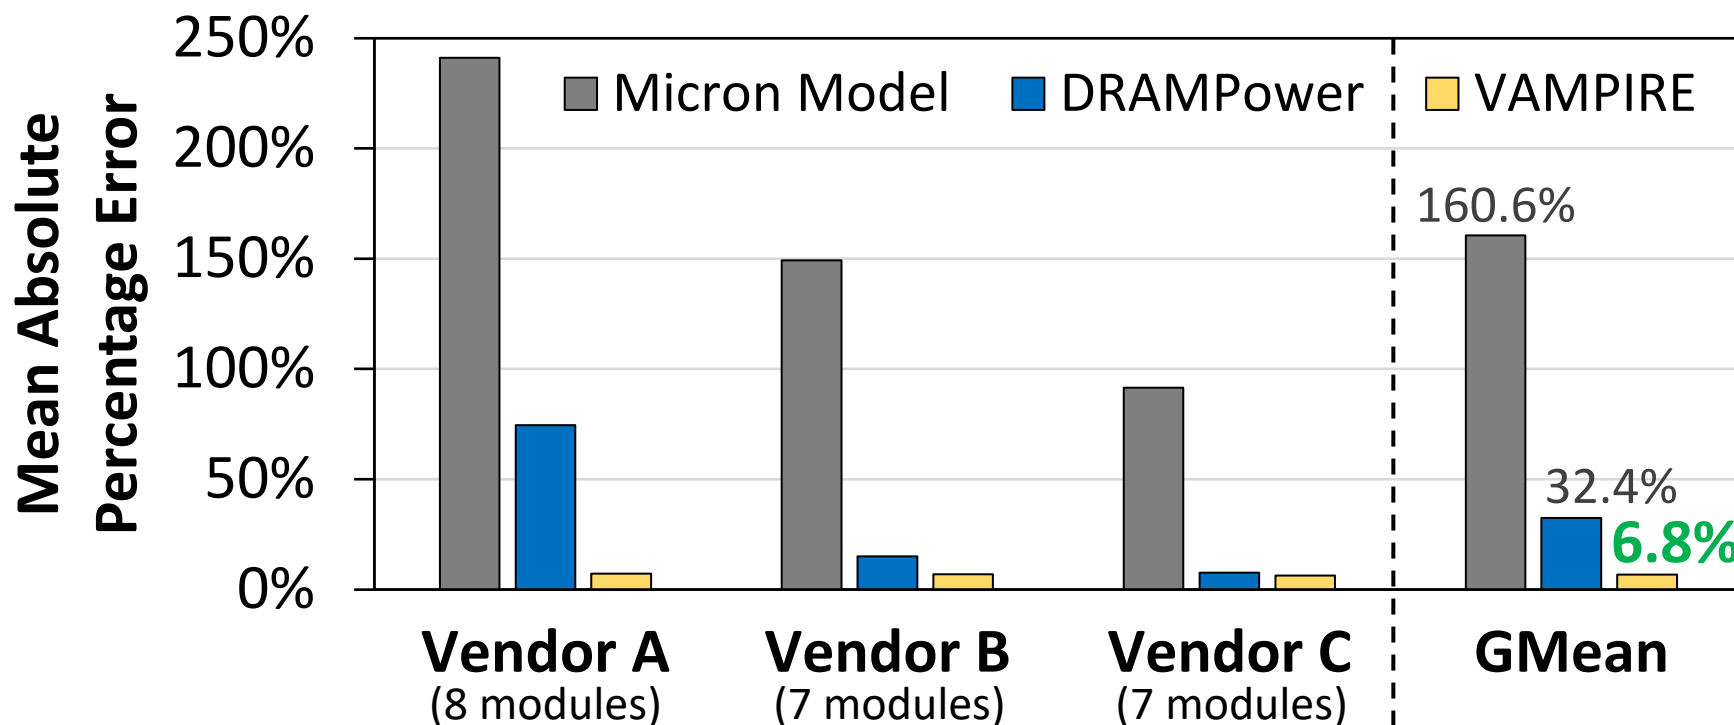
- **VAMPIRE: Variation-Aware model of Memory Power Informed by Real Experiments**



- VAMPIRE and raw characterization data are open-source:
<https://github.com/CMU-SAFARI/VAMPIRE>

VAMPIRE Has Lower Error Than Existing Models **SAFARI**

- Validated using new power measurements: details in the

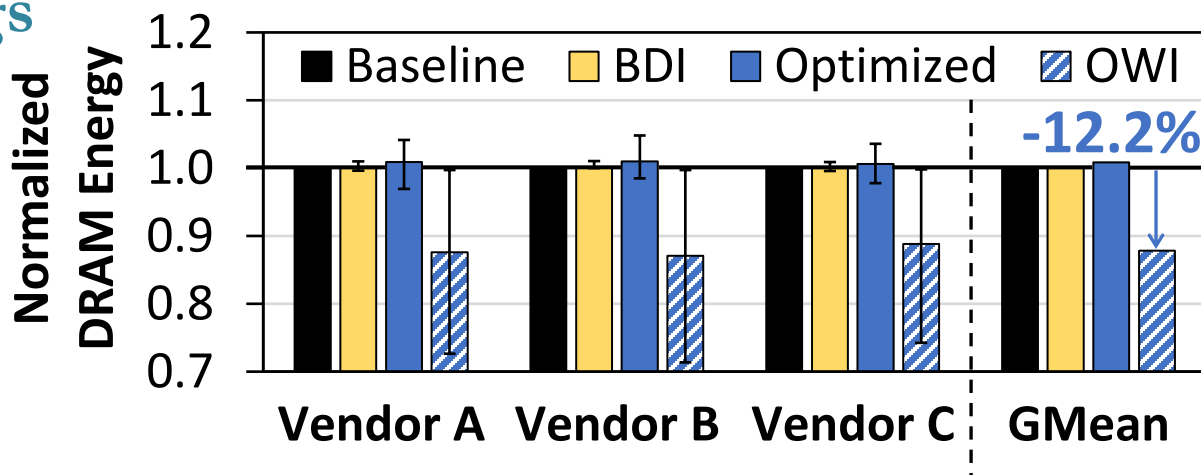


VAMPIRE has very low error for *all* vendors: 6.8%
Much more accurate than prior models

- Taking advantage of structural variation to perform **variation-aware physical page allocation** to reduce power
- Smarter DRAM **power-down scheduling**

- Reducing DRAM energy with **data-dependency-aware cache line encodings**

- 23 applications from the SPEC 2006 benchmark suite
- Traces collected using Pin and Ramulator



- We expect there to be many other new studies in the future

VAMPIRE DRAM Power Model

- Saugata Ghose, A. Giray Yaglikci, Raghav Gupta, Donghyuk Lee, Kais Kudrolli, William X. Liu, Hasan Hassan, Kevin K. Chang, Niladrish Chatterjee, Aditya Agrawal, Mike O'Connor, and Onur Mutlu,

"What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study"

*Proceedings of the ACM International Conference on Measurement and Modeling of Computer Systems (**SIGMETRICS**), Irvine, CA, USA, June 2018.*

[[Abstract](#)]

[[POMACS Journal Version \(same content, different format\)](#)]

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

[[VAMPIRE DRAM Power Model](#)]

What Your DRAM Power Models Are Not Telling You: Lessons from a Detailed Experimental Study

Saugata Ghose [†]	Abdullah Giray Yağlıkçı ^{‡†}	Raghav Gupta [†]	Donghyuk Lee [§]
Kais Kudrolli [†]	William X. Liu [†]	Hasan Hassan [‡]	Kevin K. Chang [†]
Niladrish Chatterjee [§]	Aditya Agrawal [§]	Mike O'Connor ^{§¶}	Onur Mutlu ^{‡†}

[†]Carnegie Mellon University

[‡]ETH Zürich

[§]NVIDIA

[¶]University of Texas at Austin