Variable Read Disturbance (VRD)
An Experimental Analysis of Temporal Variation in DRAM Read Disturbance

Ataberk Olgun, F. Nisa Bostancı, İsmail Emir Yüksel
Oğuzhan Canpolat, Haocong Luo, Geraldo F. Oliveira
A. Giray Yağlıkçı, Minesh Patel, Onur Mutlu

https://arxiv.org/pdf/2502.13075

ETH zürich SAFARI RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

A Typical DRAM-Based Computing System

DRAM Chip

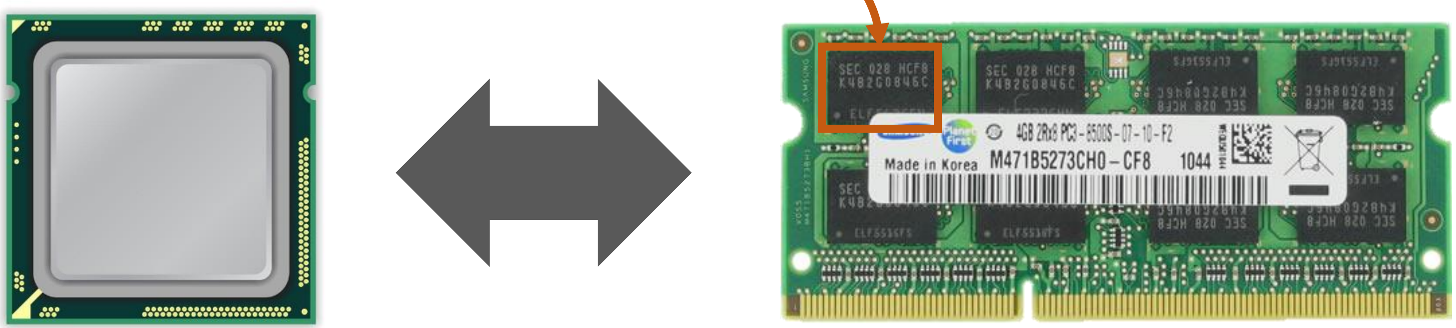

CPU

DRAM Module

DRAM Organization

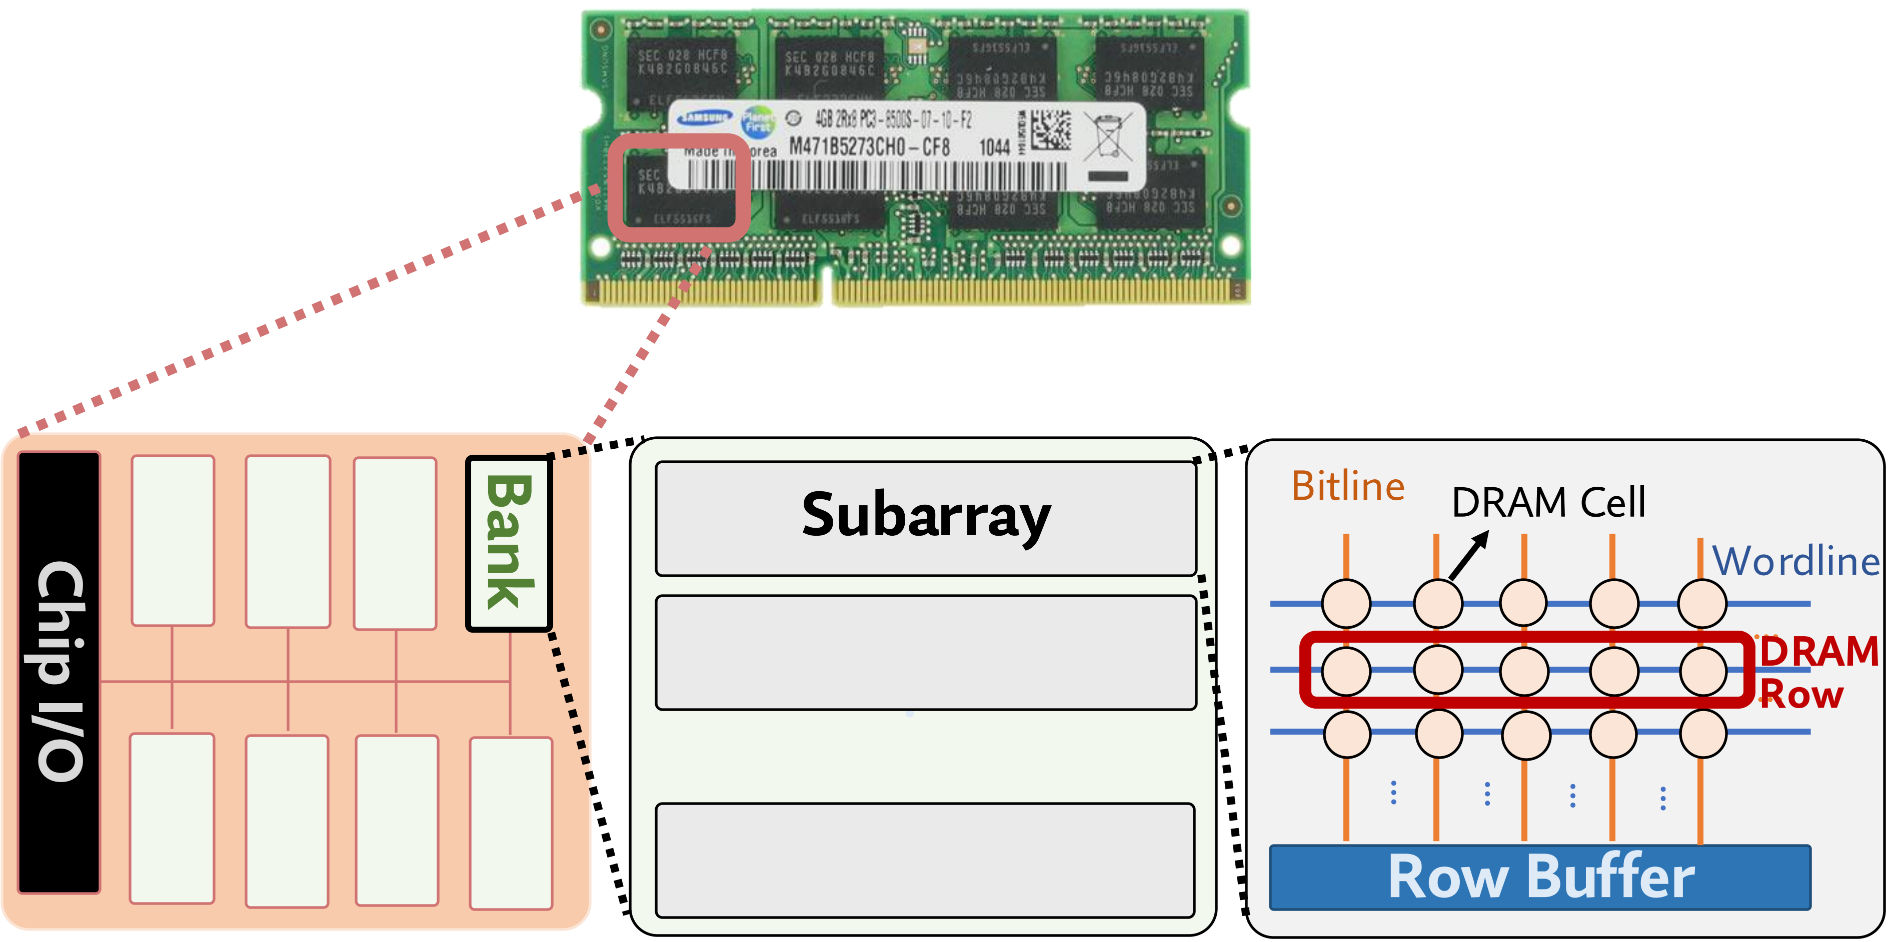

DRAM Chip DRAM Bank DRAM Subarray

Read Disturbance in DRAM (I)

- Read disturbance in DRAM breaks memory isolation
- Prominent example: RowHammer

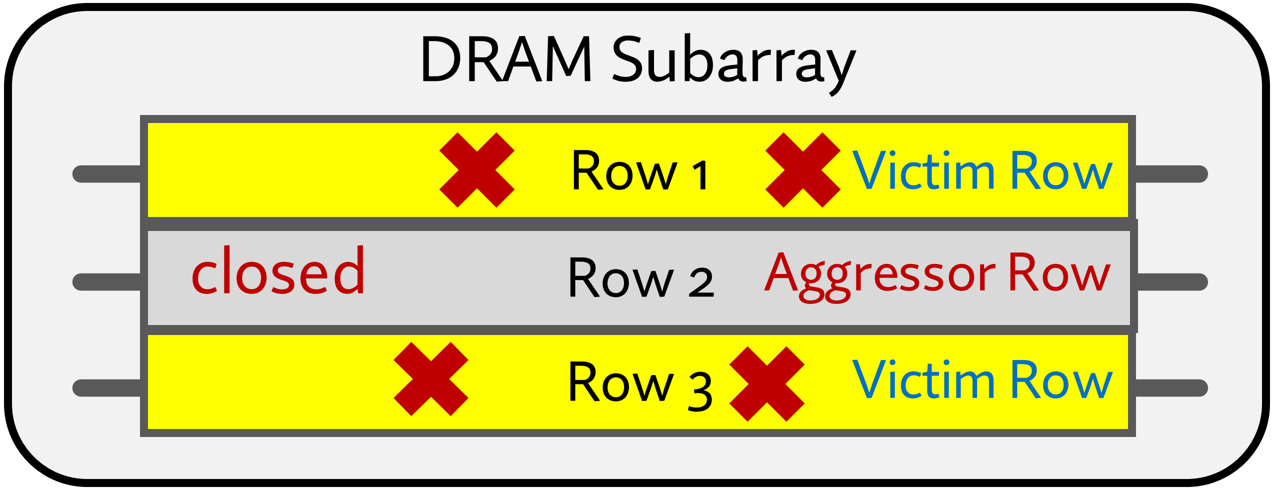

Repeatedly opening (activating) and closing a DRAM row many times causes RowHammer bitflips in adjacent rows

[Kim+, ISCA'14]

Read Disturbance in DRAM (II)

- Read disturbance in DRAM breaks memory isolation
- A new read disturbance phenomenon: RowPress

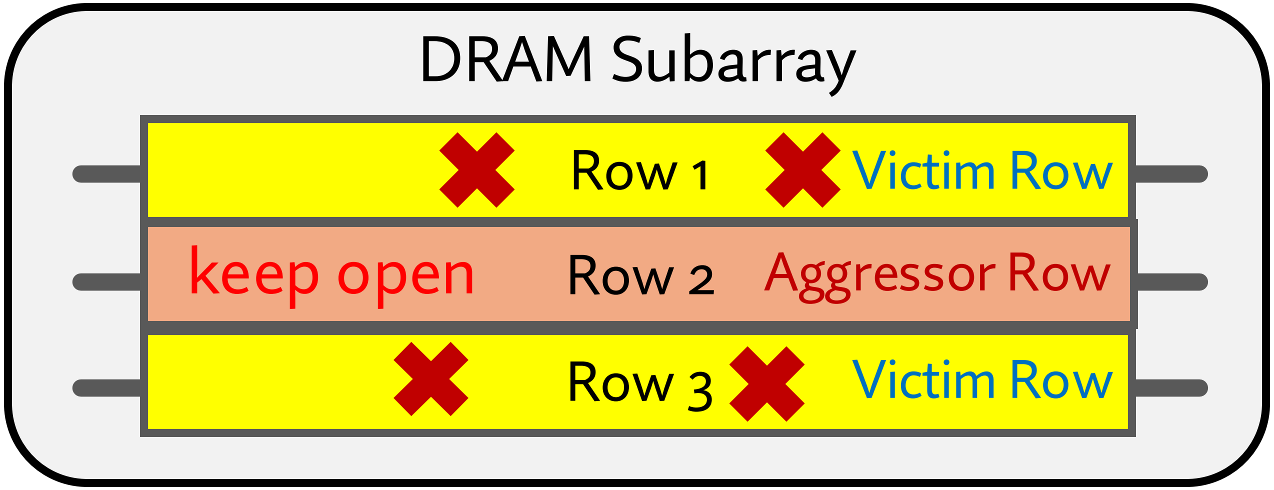

Keeping a DRAM row open for a long time causes bitflips in adjacent rows

[Luo+, ISCA'23]

Read Disturbance Solutions

There are many solutions to mitigate read disturbance bitflips

- More robust DRAM chips and/or error-correcting codes
- Increased refresh rate
- Physical isolation
- Row remapping
- Preventive refresh
- Proactive throttling

Each solution offers a different system design point
in reliability, performance, energy, and area tradeoff space

The Read Disturbance Threshold (RDT)

- Many secure read disturbance solutions take a preventive action before a bitflip manifests
  - E.g., preventively refresh a victim row

- Must **accurately quantify** the amount of disturbance that a row can withstand before experiencing a bitflip
  - Typically identified by testing for read disturbance failures

- Read Disturbance Threshold (RDT): The number of aggressor row activations needed to induce the first bitflip

Accurate Identification of Read Disturbance Threshold is Critical for System Security and Performance

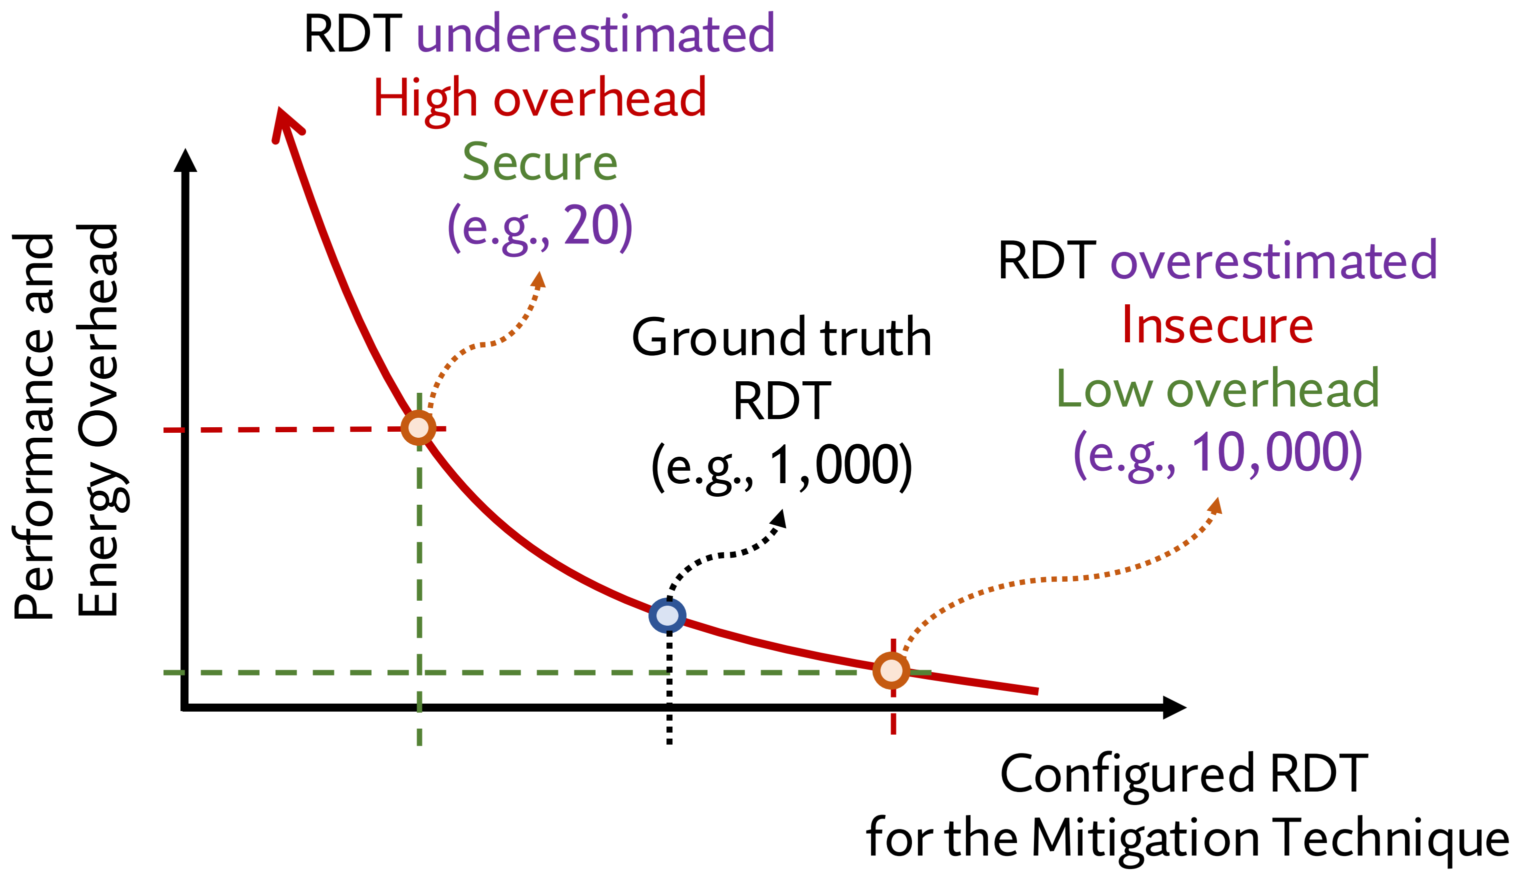

To securely prevent bitflips at low overhead
RDT must be **accurately identified** and carefully configured

Variable Read Disturbance (VRD) Summary

Research Question

- How accurately and efficiently can we measure the read disturbance threshold (RDT) of each DRAM row?

Experimental Characterization

- Record >100M RDT measurements across 3750 rows and many test parameters (e.g., temperature, data pattern) in 160 DDR4 and 4 HBM2 chips

Key Observations

- RDT changes significantly and unpredictably over time: VRD
- Maximum observed RDT is 3.5X higher than minimum (for a row)
- Smallest RDT (for a row) may appear after 94,467 measurements

Implications for System Security and Robustness

- RDT cannot be accurately identified quickly
- *Given our limited dataset*, guardbands (>10%) and ECC (SECDED or Chipkill) *may* prevent VRD-induced bitflips at significant performance cost
  - More data and analyses needed to make definitive conclusion
- Call for future work on understanding and efficiently mitigating VRD

Talk Outline

I. Motivation

II. Experimental Characterization Methodology

III. Foundational Results

IV. In-Depth Analysis of VRD

V. Implications for System Security and Robustness

VI. Conclusion

Talk Outline

I. **Motivation**

II. Experimental Characterization Methodology

III. Foundational Results

IV. In-Depth Analysis of VRD

V. Implications for System Security and Robustness

VI. Conclusion

Motivation

DRAM chips are increasingly more vulnerable to read disturbance with technology scaling

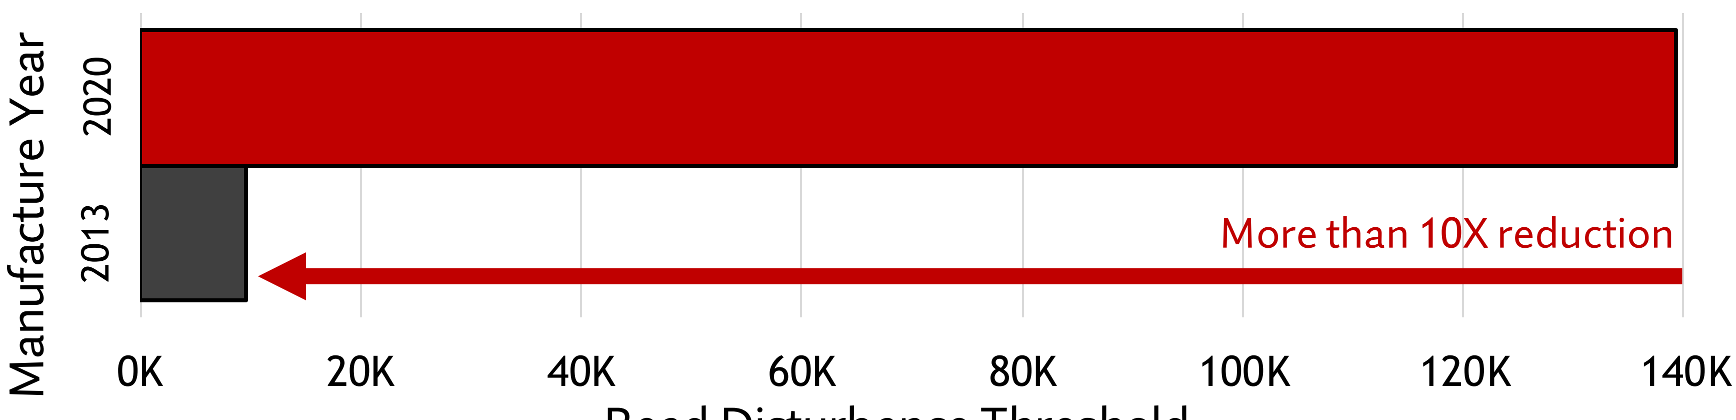

[Yaglikci+, HPCA'24]

Motivation

DRAM read disturbance worsens
as DRAM chip density increases

Existing solutions become more aggressive

Aggressive preventive actions make
existing solutions prohibitively expensive

Motivation

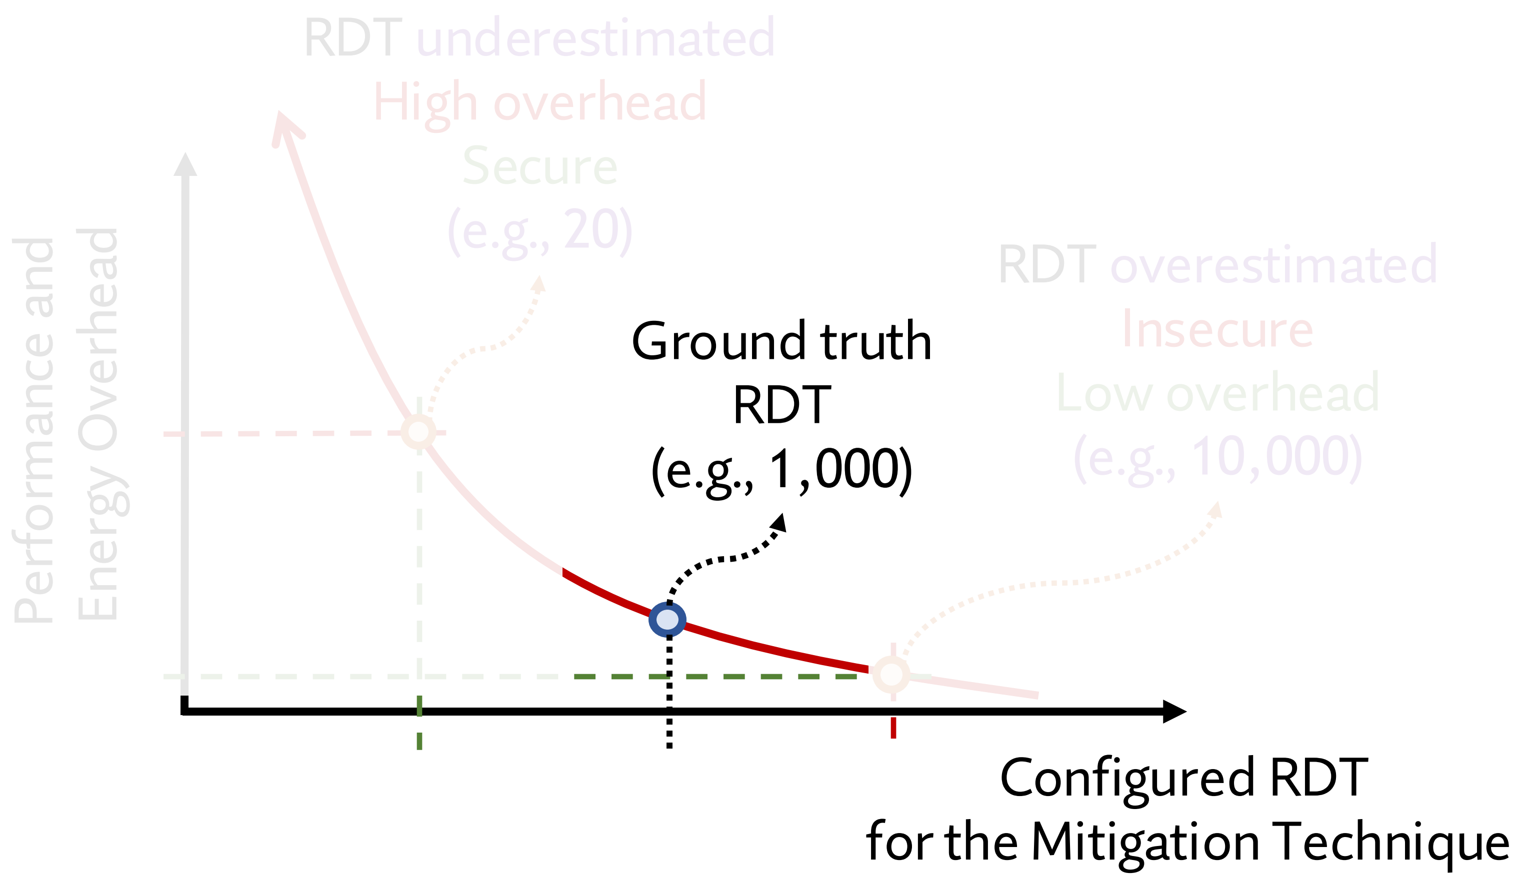

Prior works assume that the **ground truth**
Read Disturbance Threshold (RDT) **can be identified**

Problem

No prior work rigorously studies

temporal variation of

DRAM read disturbance threshold

&

implications for future solutions

Our Goal

Answer two research questions:

1. Does RDT change over time?

2. How reliably and efficiently can RDT be measured?

Analyze implications for read disturbance solutions

Talk Outline

I. Motivation

II. Experimental Characterization Methodology

III. Foundational Results

IV. In-Depth Analysis of VRD

V. Implications for System Security and Robustness

VI. Conclusion

DDR4 DRAM Testing Infrastructure

DRAM Bender on a Xilinx Virtex UltraScale+ XCU200

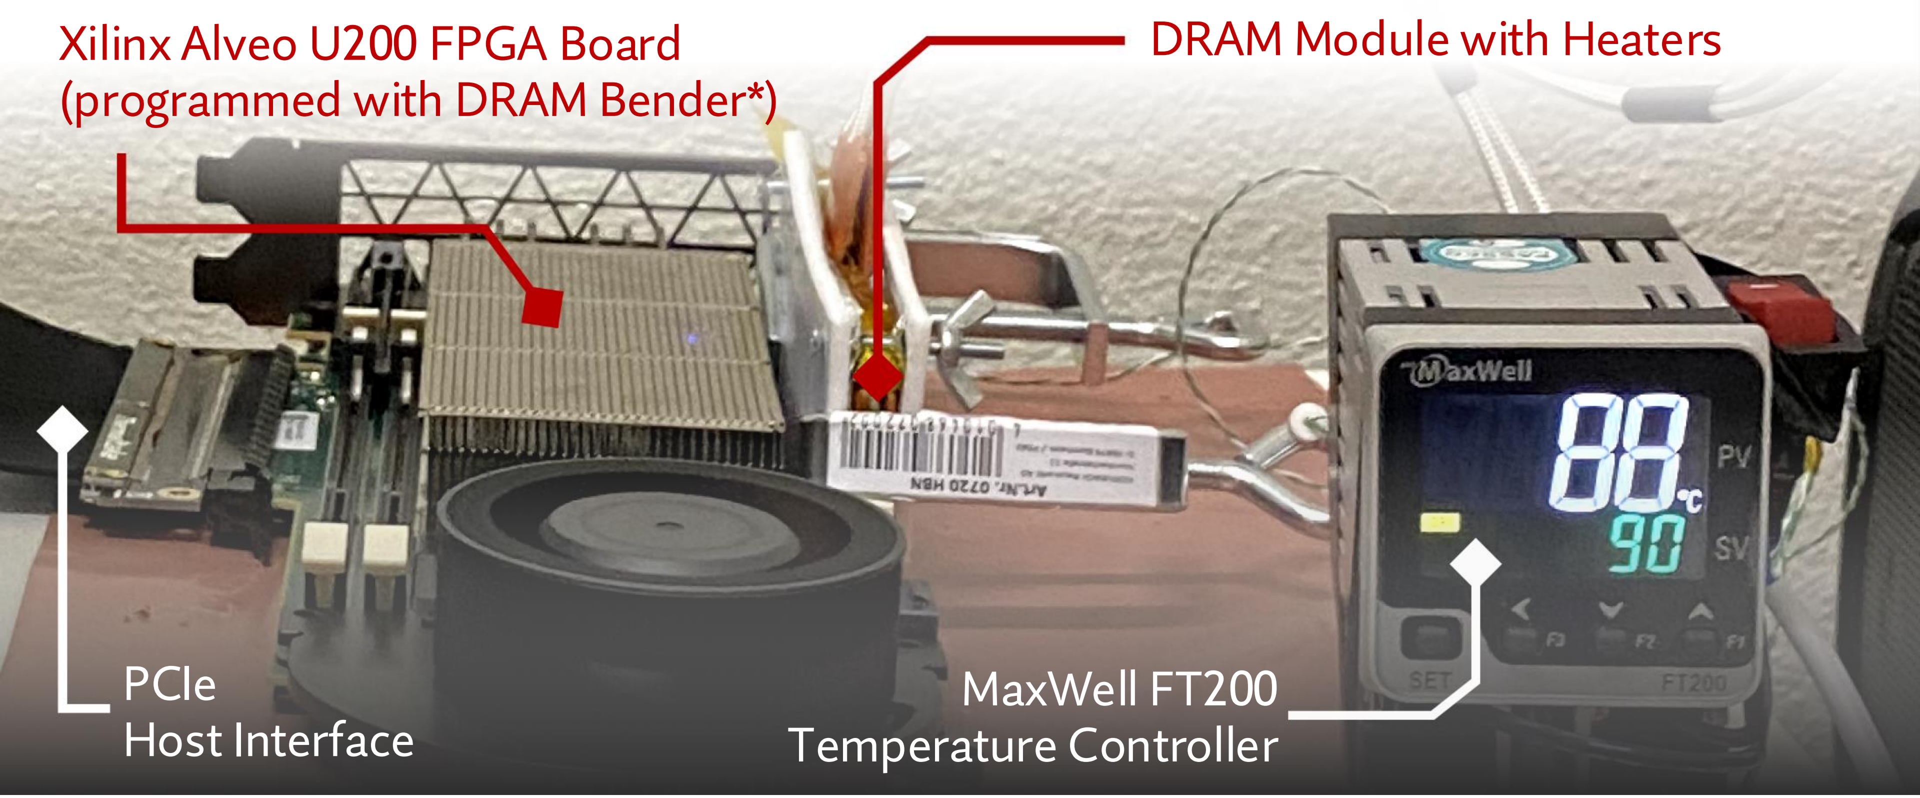

Fine-grained control over DRAM commands,
timing parameters (±1.5ns), and temperature (±0.5°C)

*Olgun et al., ["DRAM Bender: An Extensible and Versatile FPGA-Based Infrastructure to Easily Test State-of-the-Art DRAM Chips,"]() TCAD, 2023. [GitHub: https://github.com/CMU-SAFARI/DRAM-Bender]

HBM2 DRAM Testing Infrastructure

DRAM Bender on a Bittware XUPVVH

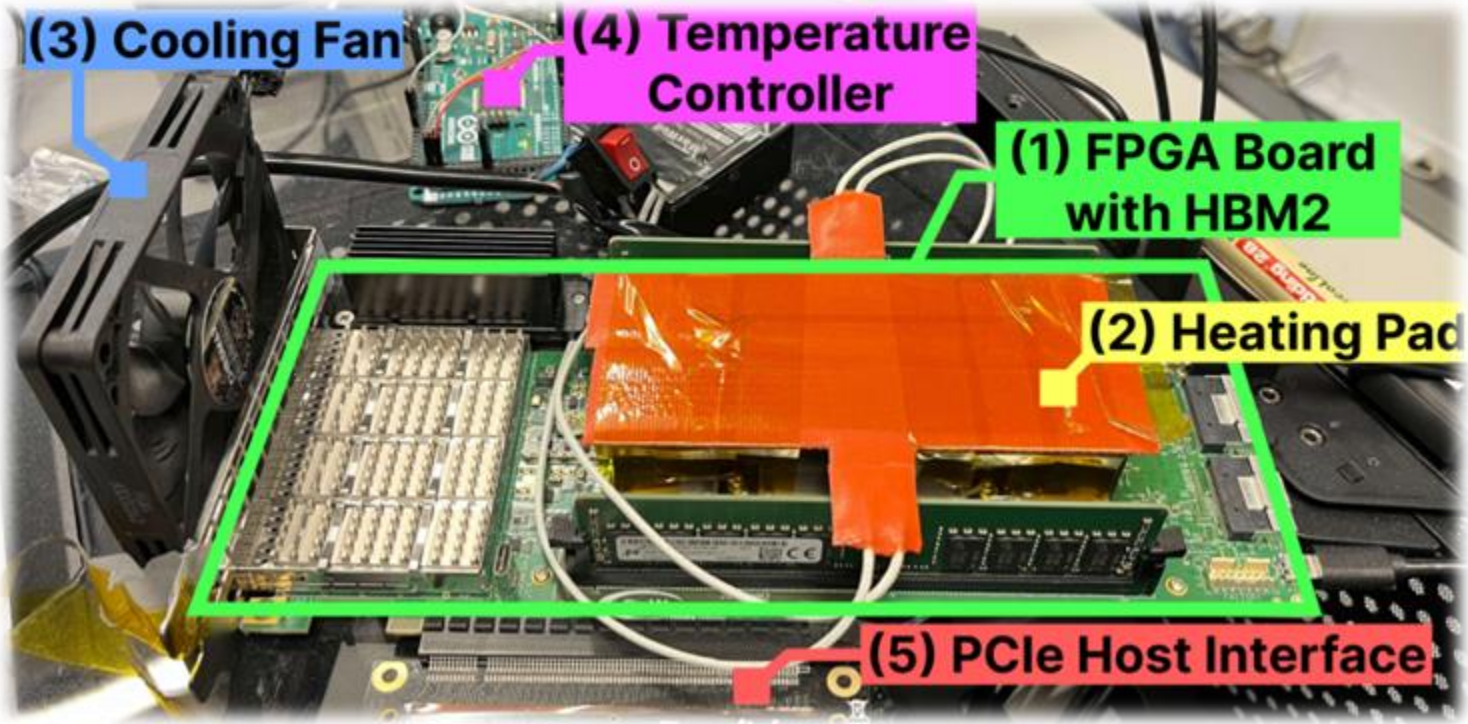

Fine-grained control over DRAM commands, timing parameters (±1.67ns), and temperature (±0.5°C)

*Olgun et al., ["DRAM Bender: An Extensible and Versatile FPGA-Based Infrastructure to Easily Test State-of-the-Art DRAM Chips,"]() TCAD, 2023. [GitHub: https://github.com/CMU-SAFARI/DRAM-Bender]

Tested DRAM Chips

160 DDR4 and 4 HBM2 Chips from SK Hynix, Micron, Samsung

Mfr.	DDR4 Module	# of Chips	Density Die Rev.	Chip Org.	Date (ww-yy)
Mfr. H (SK Hynix)	H0	8	8Gb – J	x8	N/A
	H1	8	16Gb – C	x8	36-21
	H2	8	8Gb – A	x8	43-18
	H3, H4	8	8Gb – D	x8	38-19
	H5, H6	8	8Gb – D	x8	24-20
Mfr. M (Micron)	M0	4	16Gb – E	x16	46-20
	M1	8	16Gb – F	x8	37-22
	M2	8	16Gb – F	x8	37-22
	M3, M4	8	8Gb – R	x8	12-24
	M5	8	8Gb – R	x8	10-24
	M6	8	16Gb – F	x8	12-24
Mfr. S (Samsung)	S0	8	8Gb – C	x8	N/A
	S1	8	8Gb – B	x8	53-20
	S2	8	8Gb – D	x8	10-21
	S3	8	16Gb – A	x8	20-23
	S4	4	4Gb – C	x16	19-19
	S5, S6	8	16Gb – B	x16	15-23
Mfr. S (Samsung)	HBM2 Chip Chip0 – Chip3	4	N/A	N/A	N/A

Testing Methodology

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

1. Prevent sources of interference during core test loop

- **No DRAM refresh**: to avoid refreshing victim row
- **No read disturbance mitigation mechanisms**: to observe circuit-level effects
- **No error correcting codes (ECC)**: to observe all bitflips
- Test for **less than a refresh window (32ms)** to avoid retention failures

2. Worst-case read disturbance access sequence

- We use **worst-case** read disturbance access sequence based on prior works' observations
- Double-sided read disturbance: **repeatedly access the two physically-adjacent rows**

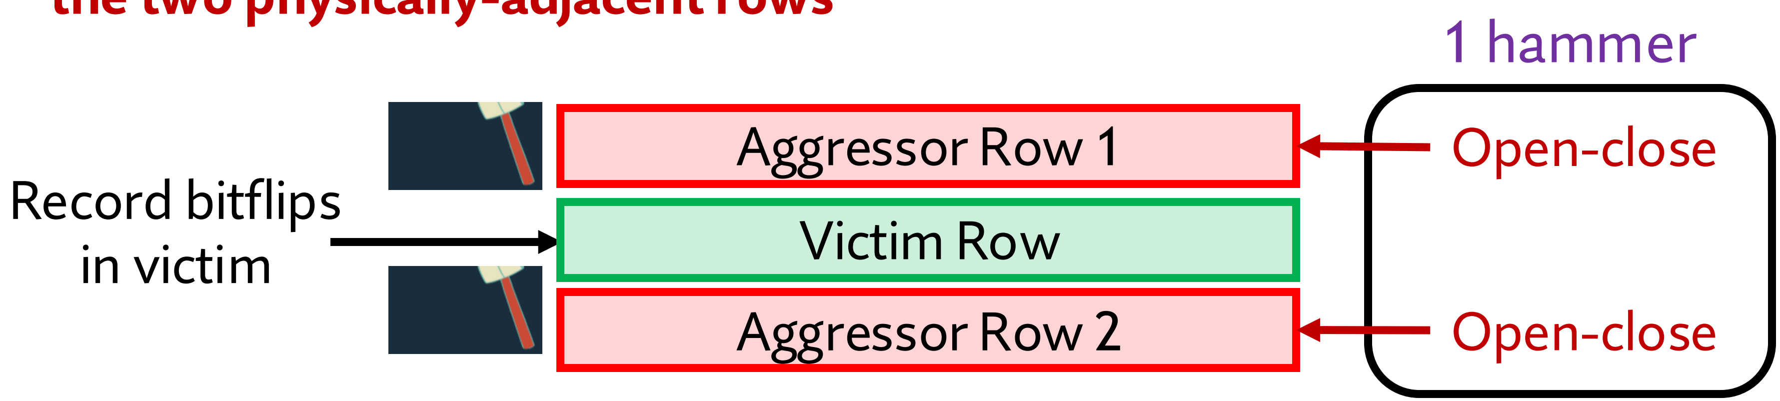

Talk Outline

Foundational Results: Key Takeaway

Key Takeaway

The Read Disturbance Threshold (RDT) of a row changes randomly and unpredictably over time

Accurately identifying RDT is challenging

Read Disturbance Threshold Changes Over Time

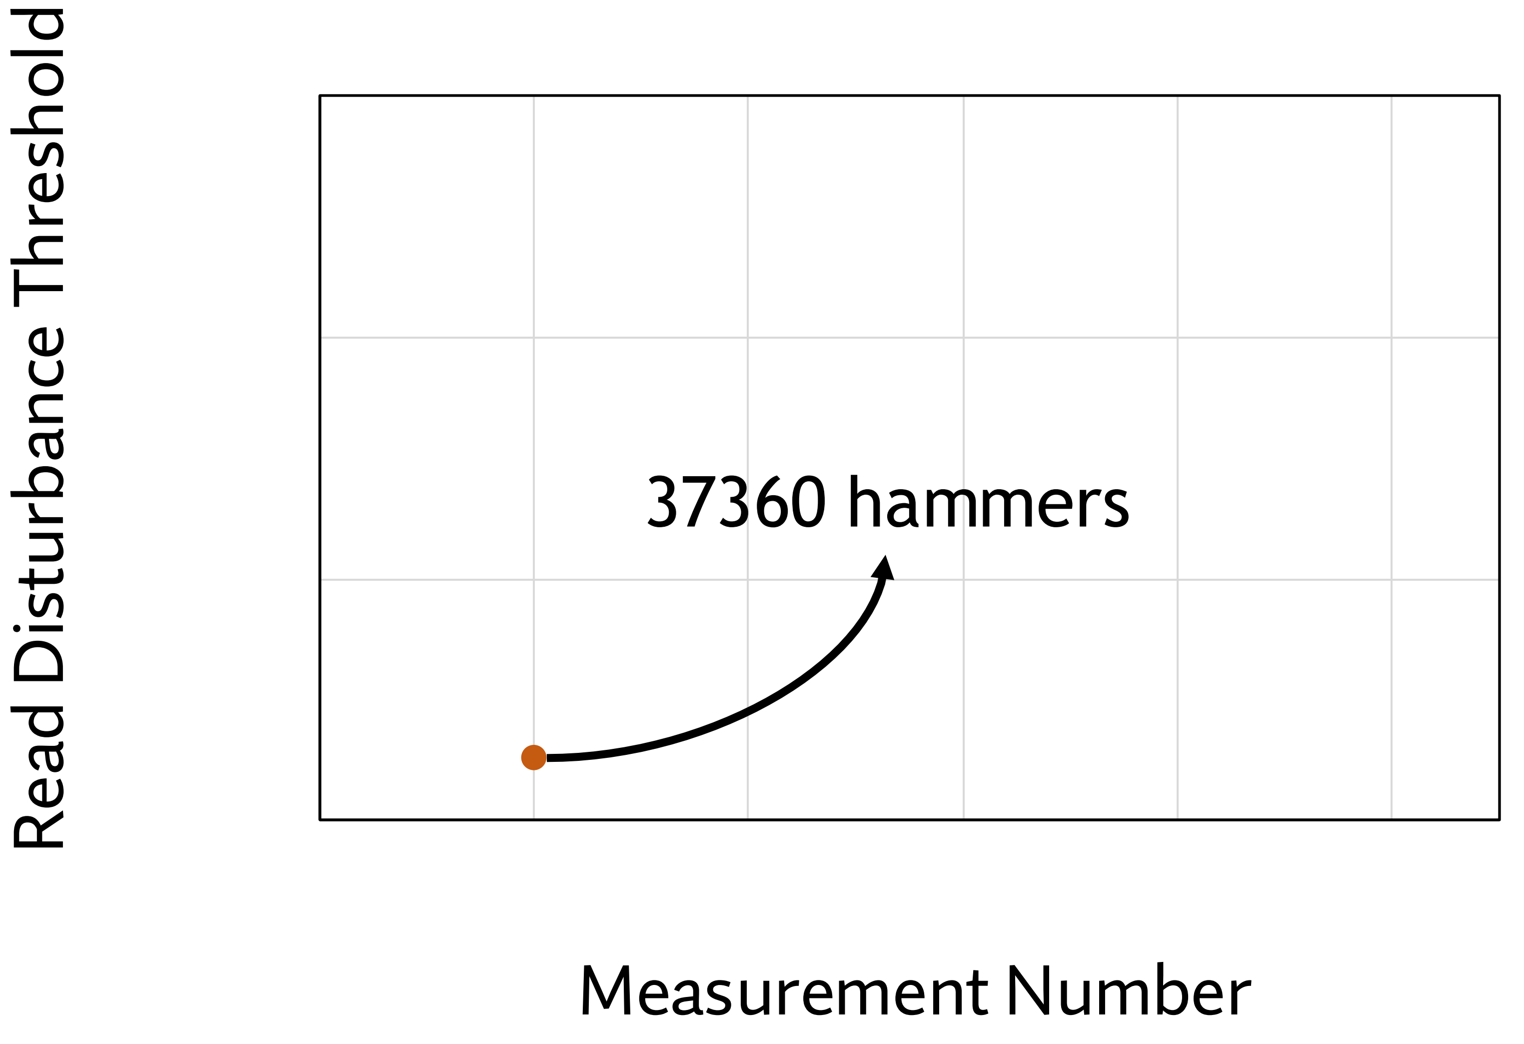

Read Disturbance Threshold Changes Over Time

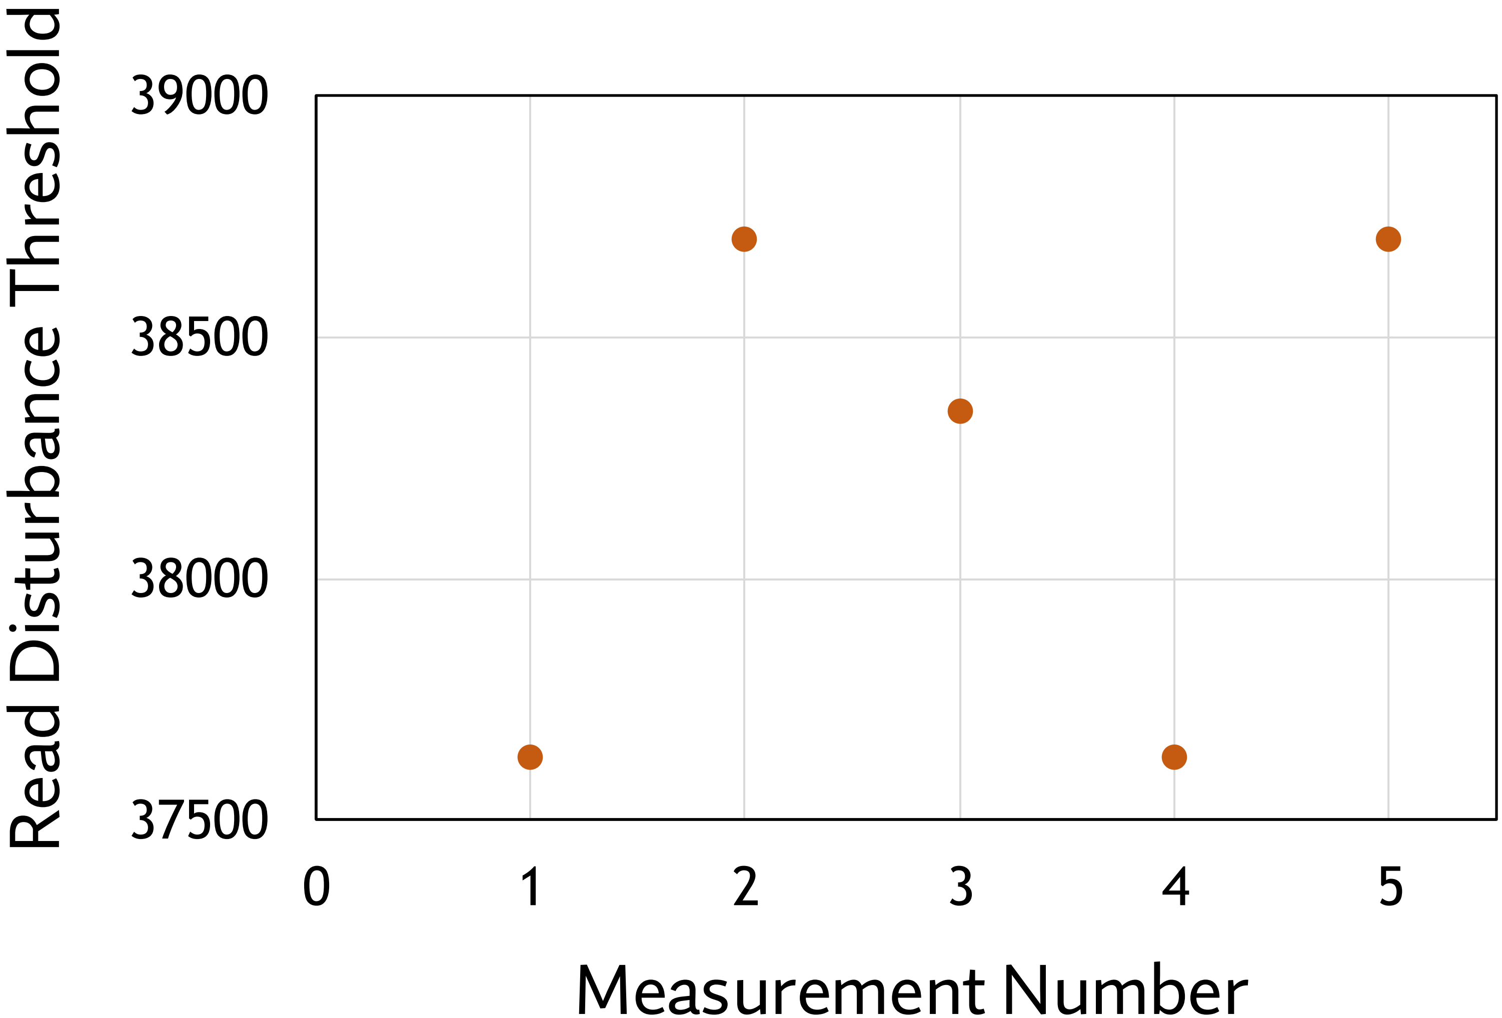

Read Disturbance Threshold Changes Over Time

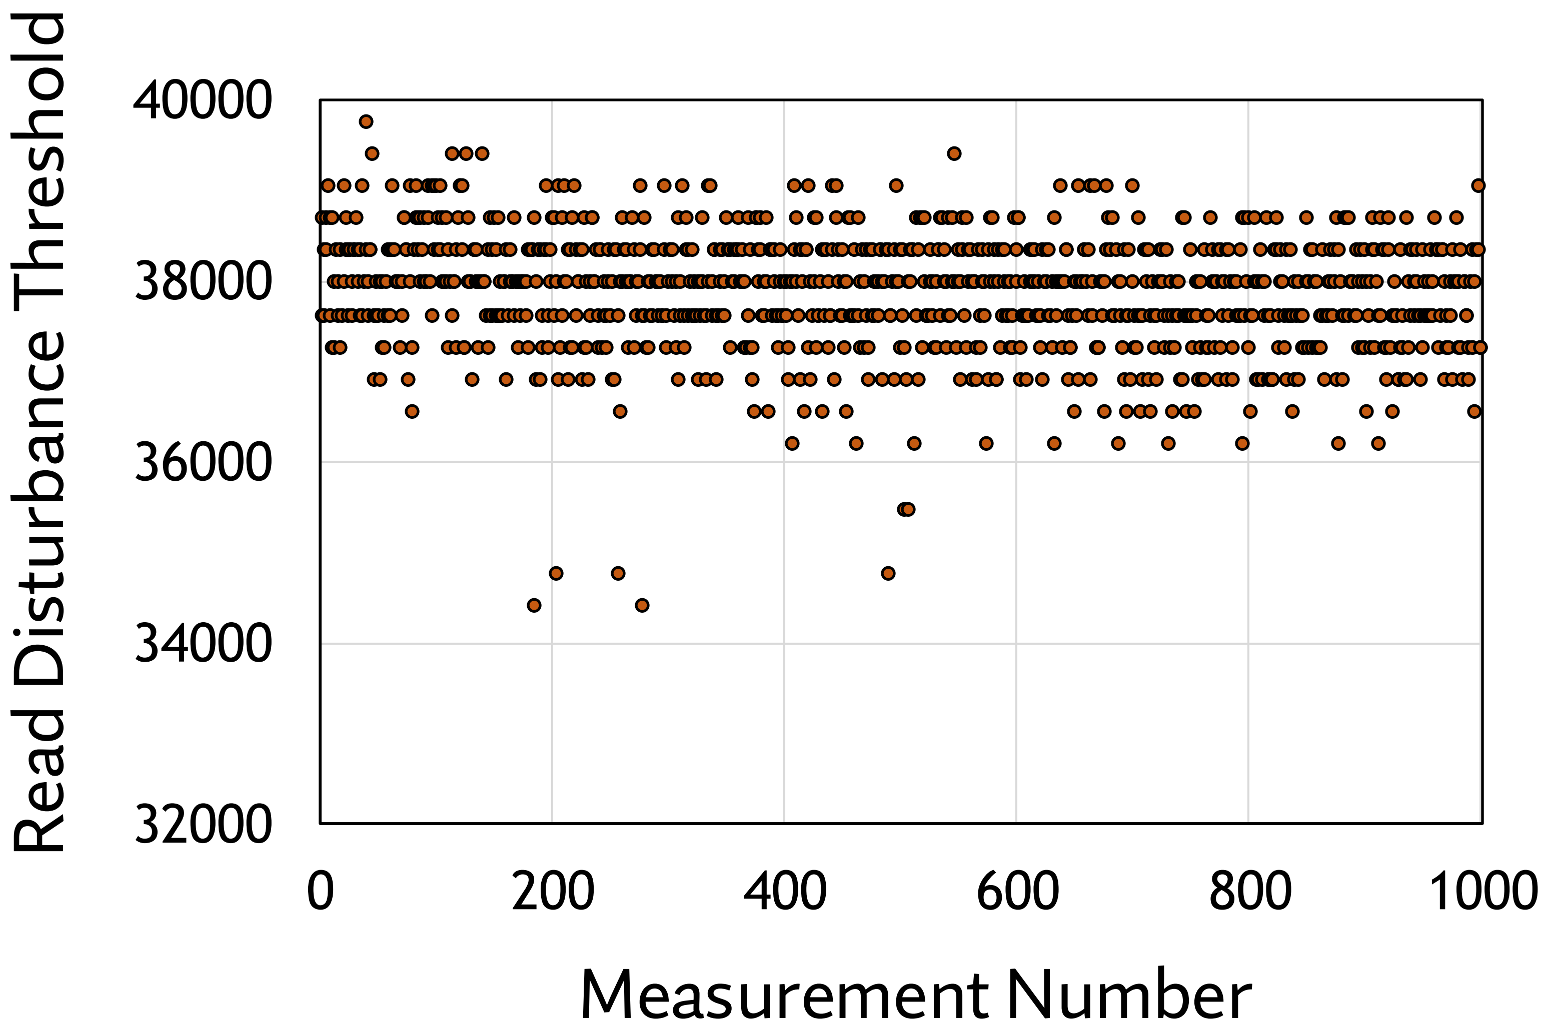

Read Disturbance Threshold Changes Over Time

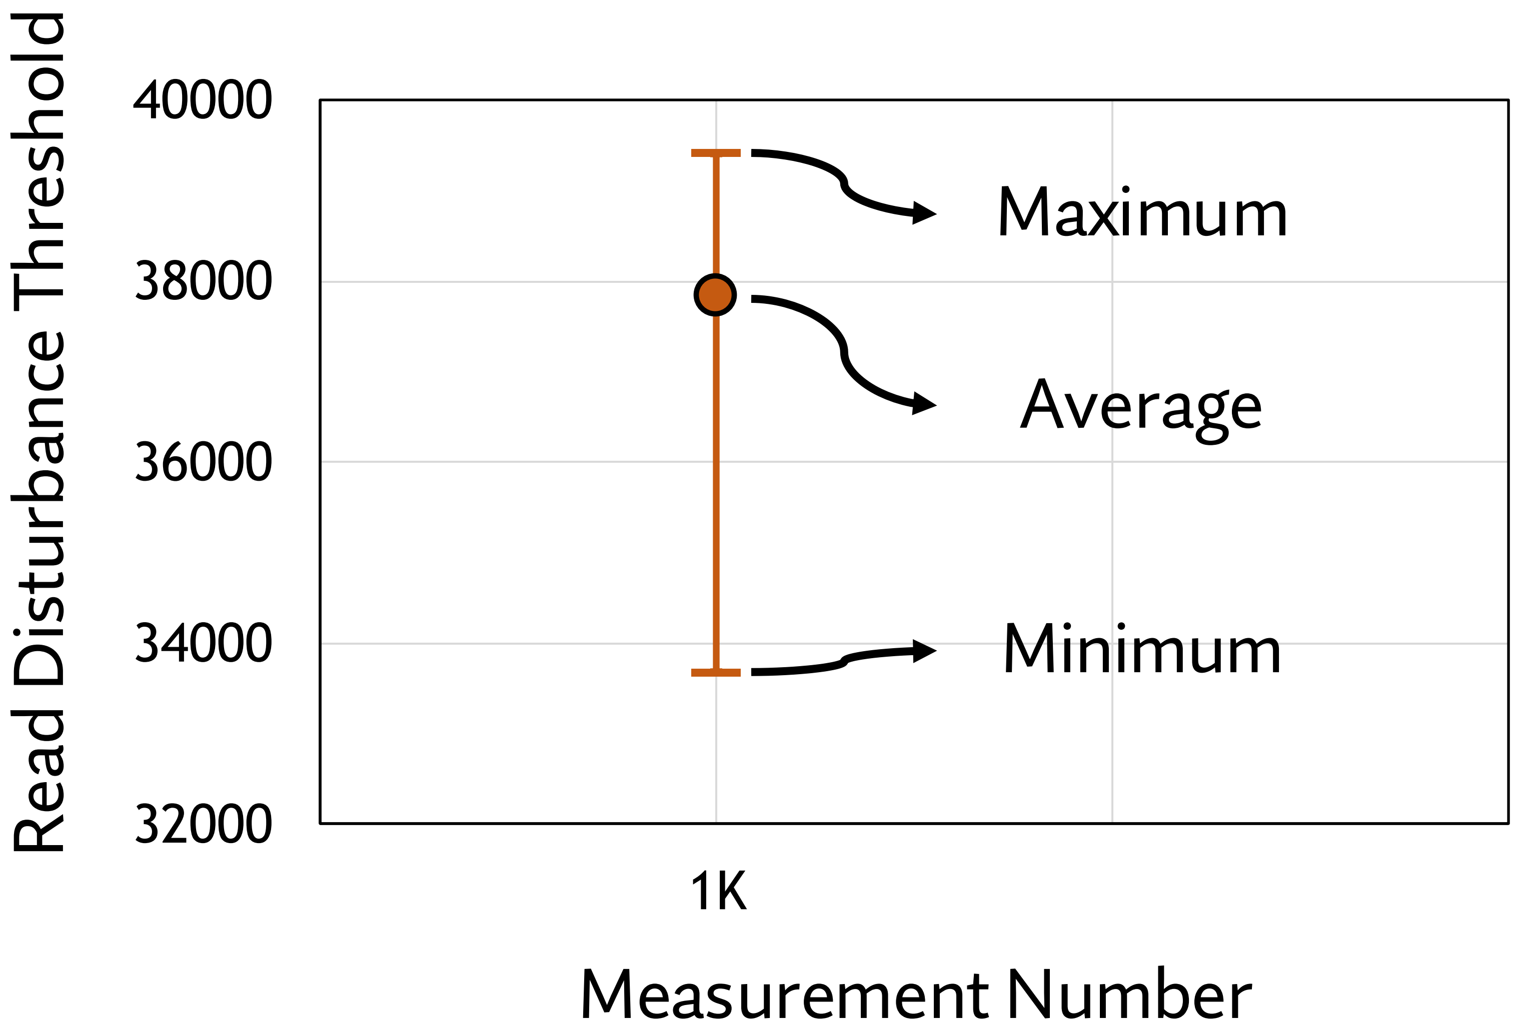

Read Disturbance Threshold Changes Over Time

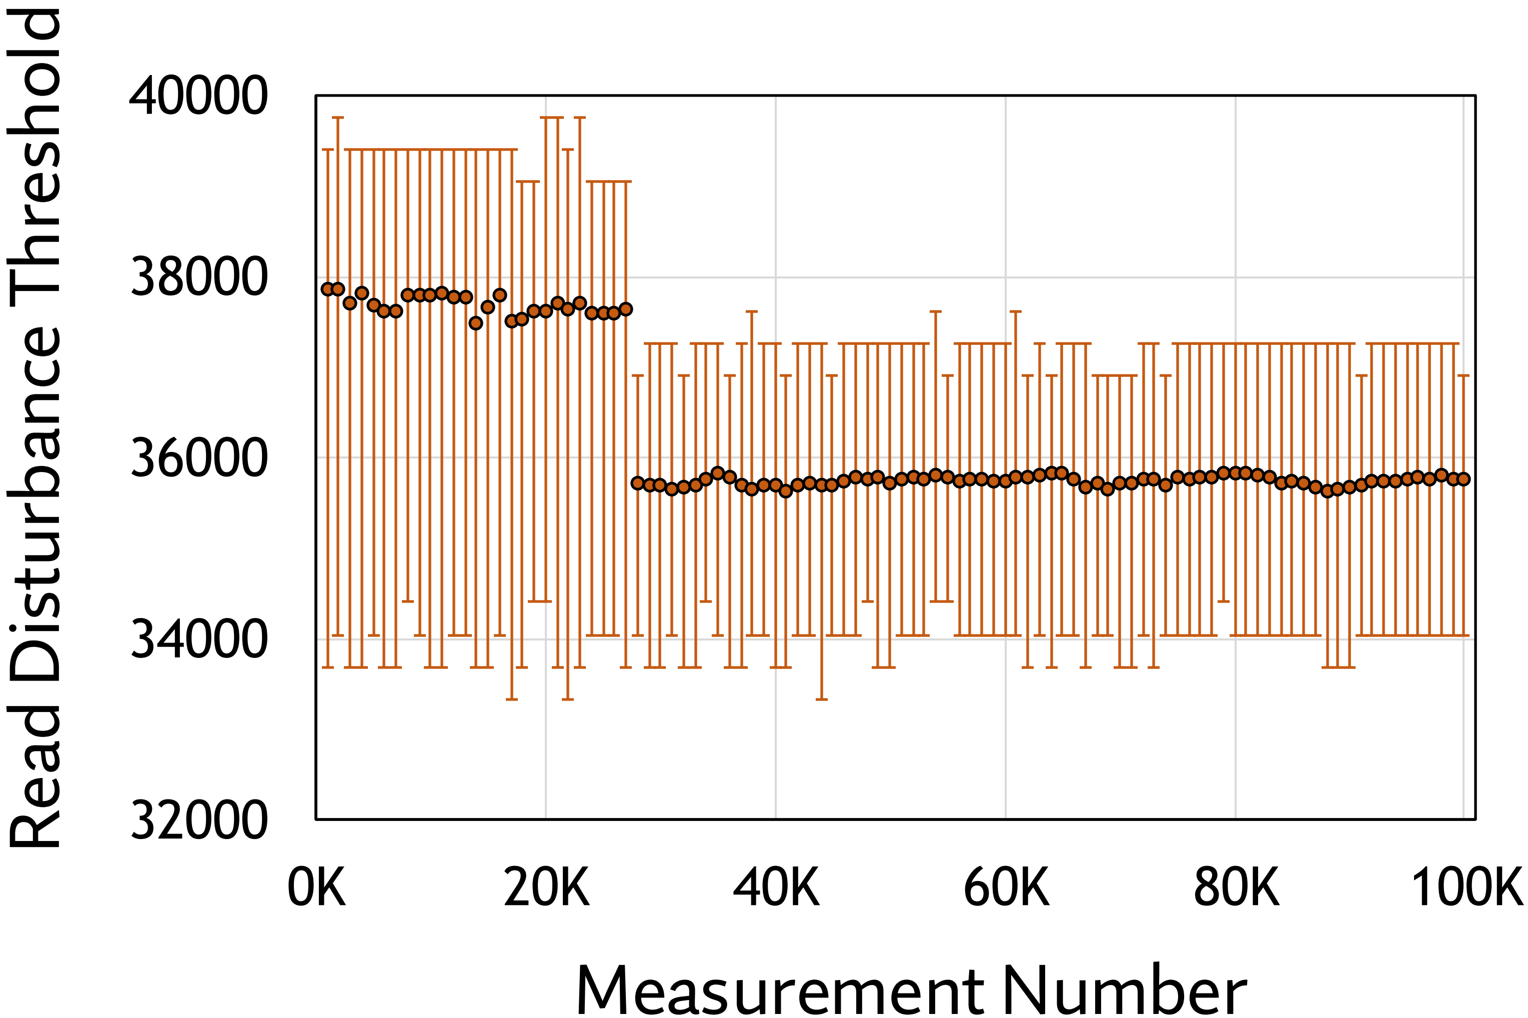

Read Disturbance Threshold Changes Over Time

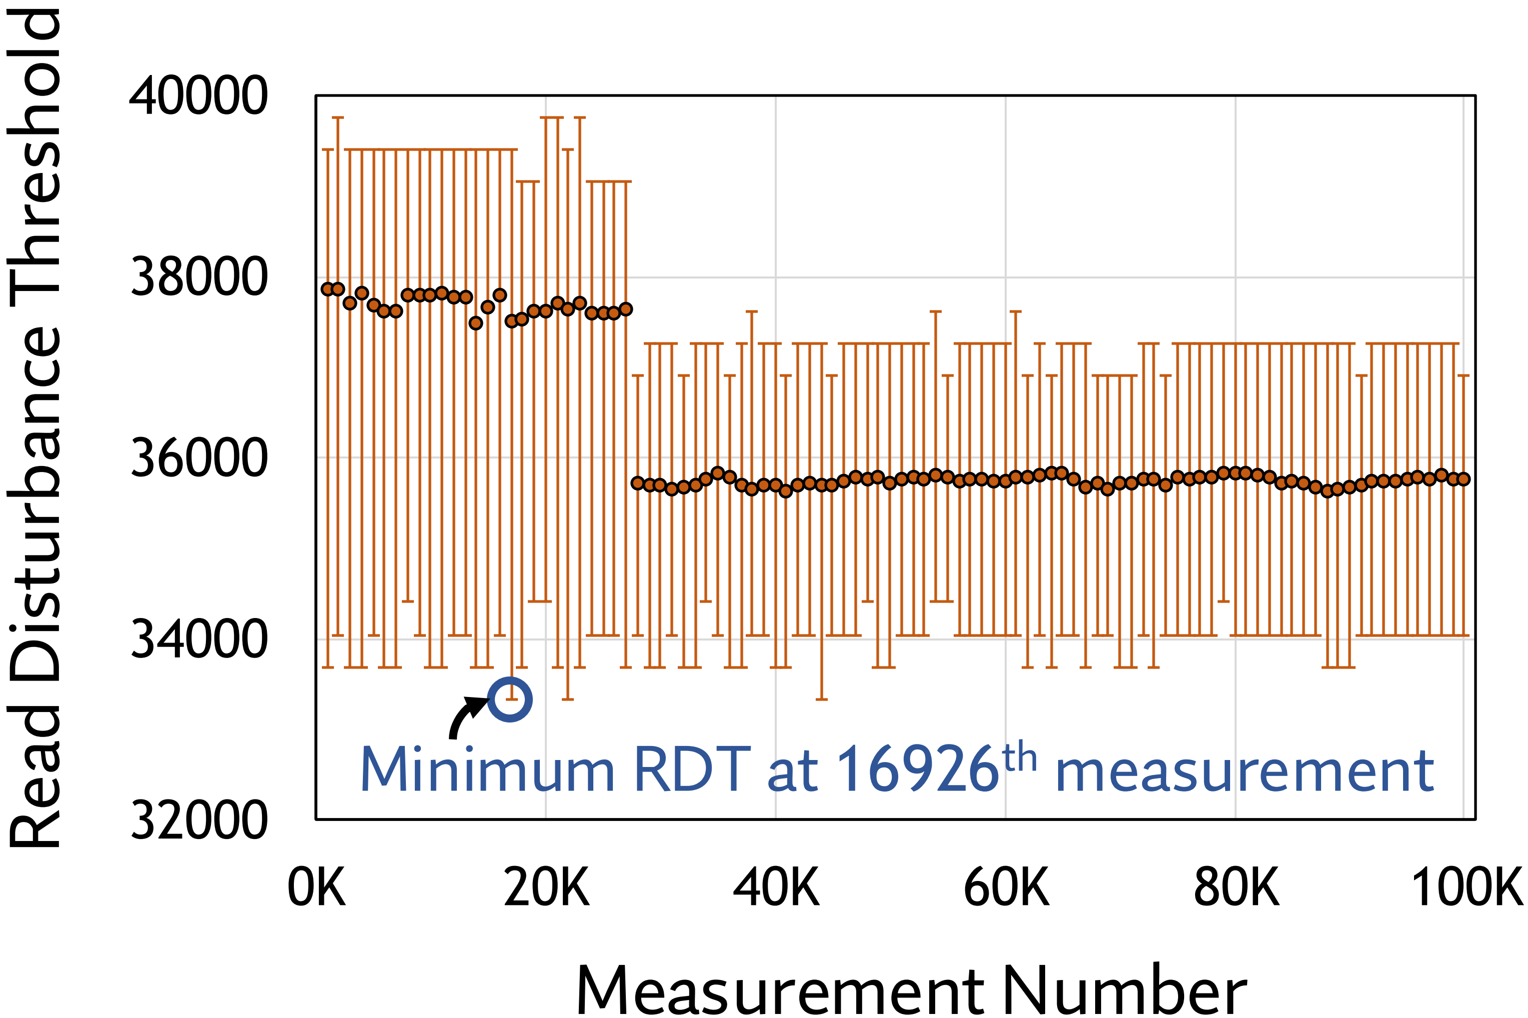

Read Disturbance Threshold Changes Over Time

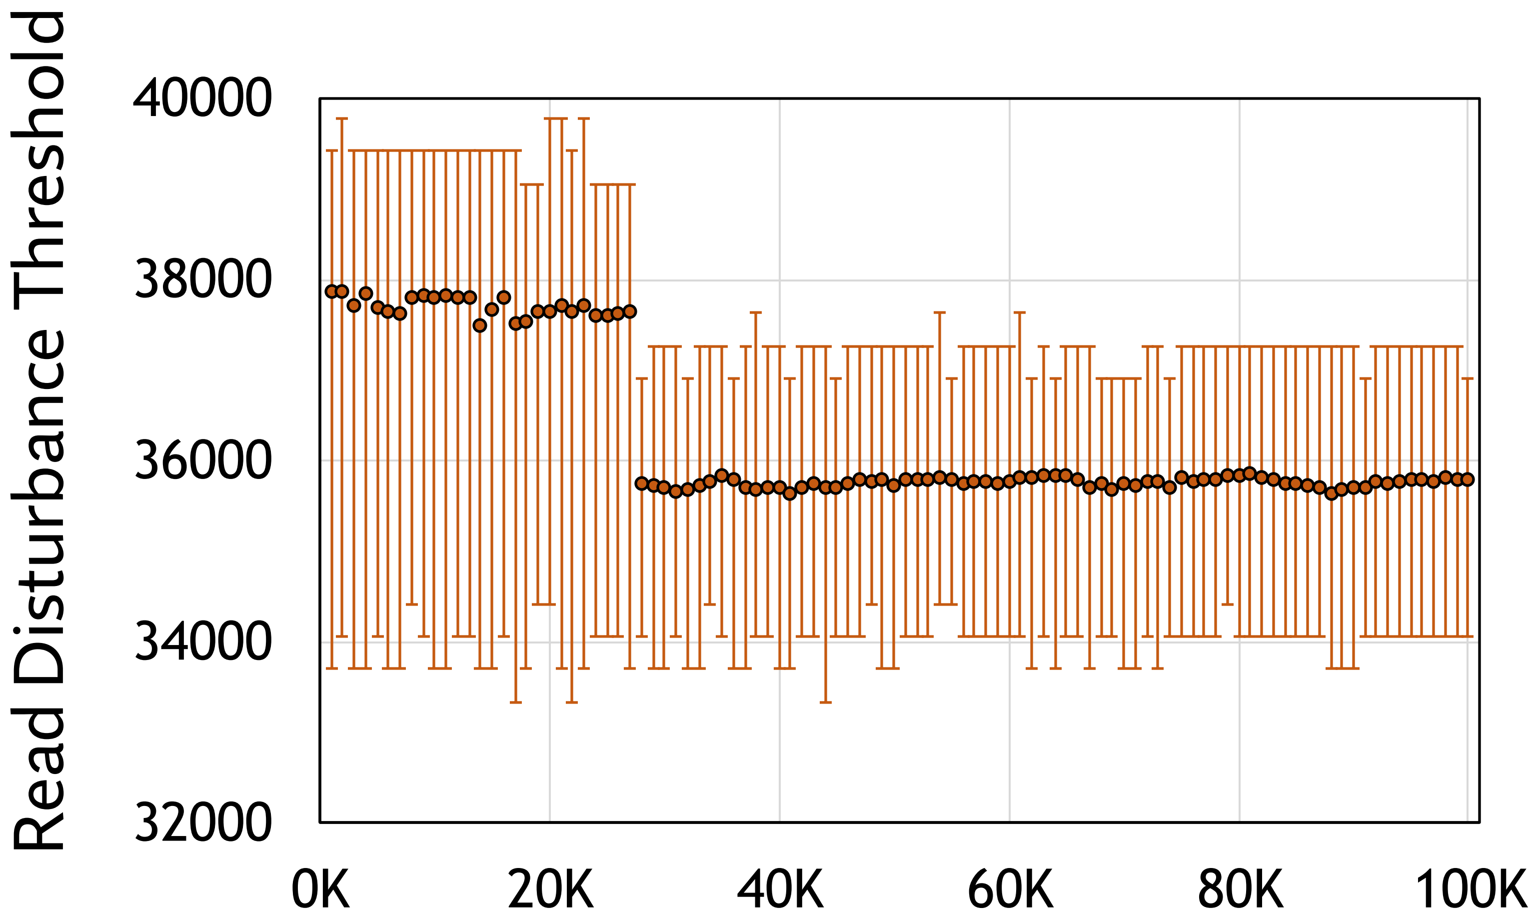

Read Disturbance Threshold of a DRAM row
varies over time: Variable Read Disturbance (VRD)

The RDT of a Row Has Multiple States

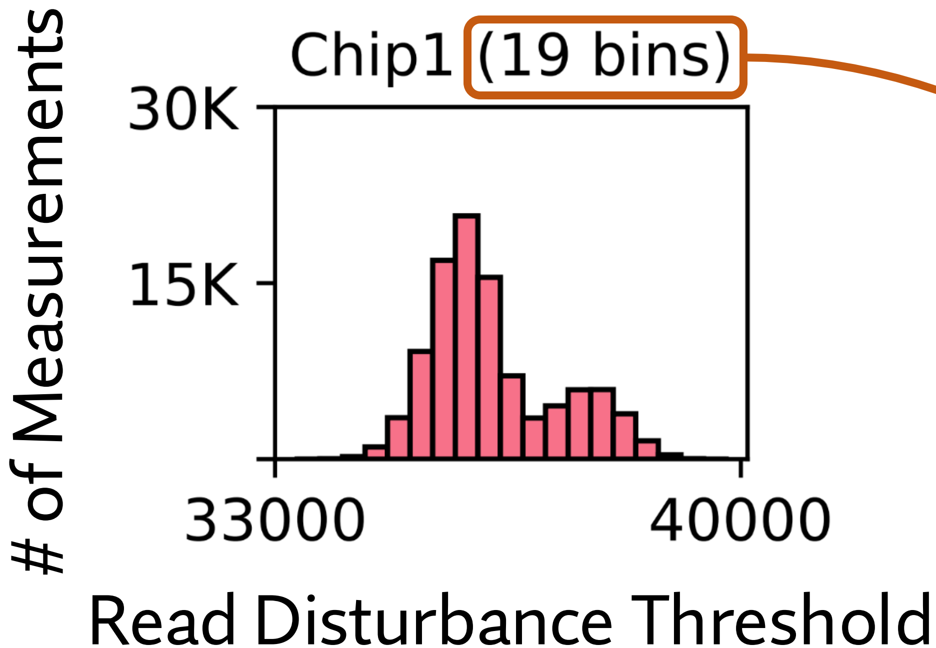

The RDT of a row takes various different values across 100,000 measurements

Variable Read Disturbance Across DRAM Chips

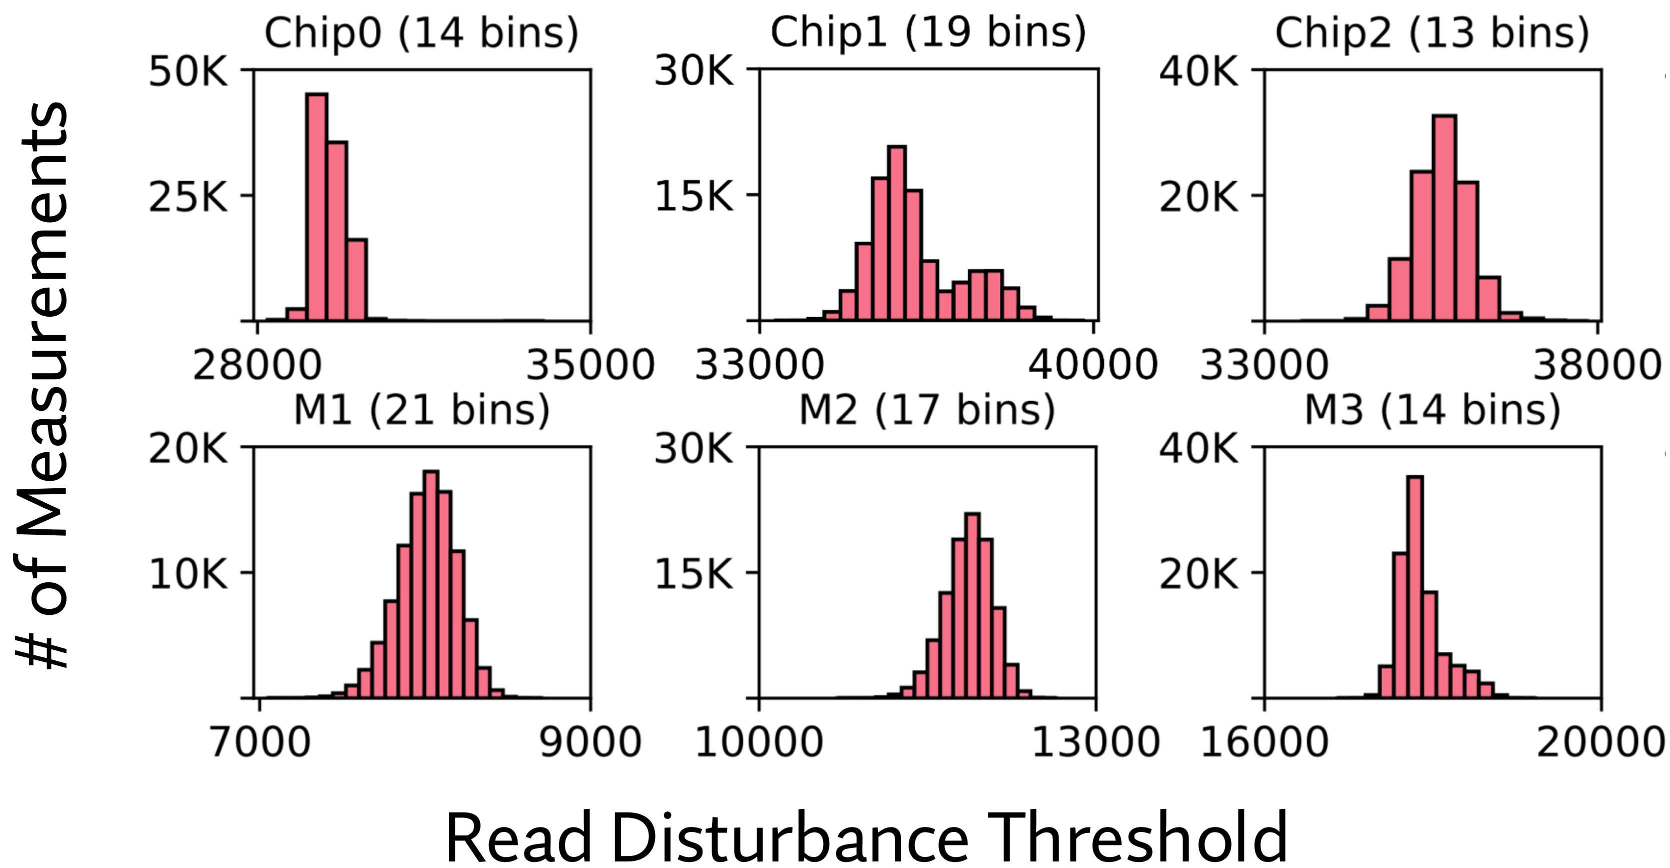

RDT consistently varies over time
across all tested DRAM chips

Variable Read Disturbance Across DRAM Chips

https://arxiv.org/pdf/2502.13075

Variable Read Disturbance:
An Experimental Analysis of Temporal Variation in DRAM Read Disturbance

Ataberk Olgun† F. Nisa Bostancı† İsmail Emir Yüksel† Oğuzhan Canpolat† Haocong Luo†
Geraldo F. Oliveira† A. Giray Yağlıkçı† Minesh Patel‡ Onur Mutlu†

ETH Zurich† *Rutgers University‡*

Modern DRAM chips are subject to read disturbance errors. These errors manifest as security-critical bitflips in a victim DRAM row *that is physically nearby a repeatedly activated (opened) aggressor row (RowHammer) or an aggressor row that is kept open for a long time (RowPress). State-of-the-art read disturbance mitigations rely on accurate and exhaustive characterization of the* read disturbance threshold *(RDT) (e.g., the number of aggressor row activations needed to induce the first RowHammer or RowPress bitflip) of every DRAM row (of which there are millions or billions in a modern system) to prevent read disturbance bitflips securely and with low overhead.*

We experimentally demonstrate for the first time that the RDT of a DRAM row significantly and unpredictably changes over time. We call this new phenomenon variable read disturbance (VRD). Our extensive experiments using 160 DDR4 chips and 4 HBM2 chips from three major manufacturers yield three key observations. First, it is very unlikely that relatively few RDT measurements can accurately identify the RDT of a DRAM row. The minimum RDT of a DRAM row appears after tens of thousands of measurements (e.g., up to 94,467), and the minimum RDT of a DRAM row is 3.5× smaller than the maximum RDT observed for that row. Second, the probability of accu-

row) *many times* (e.g., tens of thousands of times) induces *RowHammer bitflips* in physically nearby rows (i.e., victim rows) [1]. Keeping the aggressor row open for a long period of time amplifies the effects of read disturbance and induces *RowPress bitflips*, *without* requiring *many* repeated aggressor row activations [4].

A large body of work [1, 3, 26, 32, 39, 45, 69–141] proposes various techniques to mitigate DRAM read disturbance bitflips. Many high-performance and low-overhead mitigation techniques [1, 73, 74, 76, 79, 82–84, 86, 87, 91, 97, 133–135, 137–139, 142–146], including those that are used and standardized by industry [121, 126, 138, 139, 144], prevent read disturbance bitflips by *preventively* refreshing (i.e., opening and closing) a victim row *before* a bitflip manifests in that row.

To securely prevent read disturbance bitflips at low performance and energy overhead, it is important to *accurately* identify the amount of read disturbance that a victim row can withstand before experiencing a read disturbance bitflip. This amount is typically quantified using the *hammer count (the number of aggressor row activations) needed to induce the first read disturbance bitflip* in a victim row. We call this metric the *read disturbance threshold (RDT)* of the victim row.

VRD is (Likely) Unpredictable

- The outcome of the next read disturbance threshold (RDT) measurement cannot be predicted given past measurements

1 RDT histograms well resemble* random probability distributions e.g., normal distribution

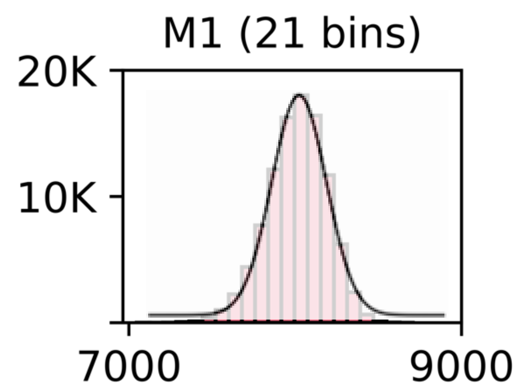

2 Analyze and find no repeating patterns in the series of consecutively measured RDT values using the autocorrelation function

** Resemblance quantified using statistical tests in the paper*

VRD is (Likely) Unpredictable

https://arxiv.org/pdf/2502.13075

Variable Read Disturbance:
An Experimental Analysis of Temporal Variation in DRAM Read Disturbance

Ataberk Olgun† F. Nisa Bostancı† İsmail Emir Yüksel† Oğuzhan Canpolat† Haocong Luo†
Geraldo F. Oliveira† A. Giray Yağlıkçı† Minesh Patel‡ Onur Mutlu†

ETH Zurich† *Rutgers University‡*

Modern DRAM chips are subject to read disturbance errors. These errors manifest as security-critical bitflips in a victim DRAM row *that is physically nearby a repeatedly activated (opened) aggressor row (RowHammer) or an aggressor row that is kept open for a long time (RowPress). State-of-the-art read disturbance mitigations rely on accurate and exhaustive characterization of the* read disturbance threshold *(RDT) (e.g., the number of aggressor row activations needed to induce the first RowHammer or RowPress bitflip) of every DRAM row (of which there are millions or billions in a modern system) to prevent read disturbance bitflips securely and with low overhead.*

We experimentally demonstrate for the first time that the RDT of a DRAM row significantly and unpredictably changes over time. We call this new phenomenon variable read disturbance (VRD). Our extensive experiments using 160 DDR4 chips and 4 HBM2 chips from three major manufacturers yield three key observations. First, it is very unlikely that relatively few RDT measurements can accurately identify the RDT of a DRAM row. The minimum RDT of a DRAM row appears after tens of thousands of measurements (e.g., up to 94,467), and the minimum RDT of a DRAM row is 3.5× smaller than the maximum RDT observed for that row. Second, the probability of accu-

row) *many times* (e.g., tens of thousands of times) induces *RowHammer bitflips* in physically nearby rows (i.e., victim rows) [1]. Keeping the aggressor row open for a long period of time amplifies the effects of read disturbance and induces *RowPress bitflips*, *without* requiring *many* repeated aggressor row activations [4].

A large body of work [1, 3, 26, 32, 39, 45, 69–141] proposes various techniques to mitigate DRAM read disturbance bitflips. Many high-performance and low-overhead mitigation techniques [1, 73, 74, 76, 79, 82–84, 86, 87, 91, 97, 133–135, 137–139, 142–146], including those that are used and standardized by industry [121, 126, 138, 139, 144], prevent read disturbance bitflips by *preventively* refreshing (i.e., opening and closing) a victim row *before* a bitflip manifests in that row.

To securely prevent read disturbance bitflips at low performance and energy overhead, it is important to *accurately* identify the amount of read disturbance that a victim row can withstand before experiencing a read disturbance bitflip. This amount is typically quantified using the *hammer count (the number of aggressor row activations) needed to induce the first read disturbance bitflip* in a victim row. We call this metric the *read disturbance threshold (RDT)* of the victim row.

Talk Outline

I. Motivation

II. Experimental Characterization Methodology

III. Foundational Results

IV. In-Depth Analysis of VRD

V. Implications for System Security and Robustness

VI. Conclusion

In-Depth Analysis: Parameter Space

- Four data patterns

Row Addresses	*Rowstripe0*	*Rowstripe1*	*Checkered0*	*Checkered1*
Victim (V)	0x00	0xFF	0x55	0xAA
Aggressors ($V \pm 1$)	0xFF	0x00	0xAA	0x55
$V \pm [2{:}8]$	0x00	0xFF	0x55	0xAA

- Three temperature levels: 50°C, 65°C, 80°C

- Three aggressor row on time values (RowPress):
  - Minimum t_{RAS} = ~35ns
  - Interval between two periodic refresh commands t_{REFI} = 7.8μs (DDR4)
  - Maximum interval between two refresh $9 \times t_{REFI}$ = 70.2μs (DDR4)

- Test 3750 rows and measure RDT 1000 times per row
  - Aside: what would happen if we measure >1M times?

In-Depth Analysis: Key Takeaways

Takeaway 1

All tested DRAM rows exhibit VRD

Takeaway 2

Relatively few (<500) RDT measurements are unlikely to yield the minimum RDT of a row

Takeaway 3

Data patterns, temperature, and aggressor row on time affect VRD

In-Depth Analysis: Key Takeaways

Takeaway 1

All tested DRAM rows exhibit VRD

Takeaway 2

Relatively few (<500) RDT measurements are unlikely to yield the minimum RDT of a row

Takeaway 3

Data patterns, temperature, and aggressor row on time affect VRD

VRD Across DRAM Rows

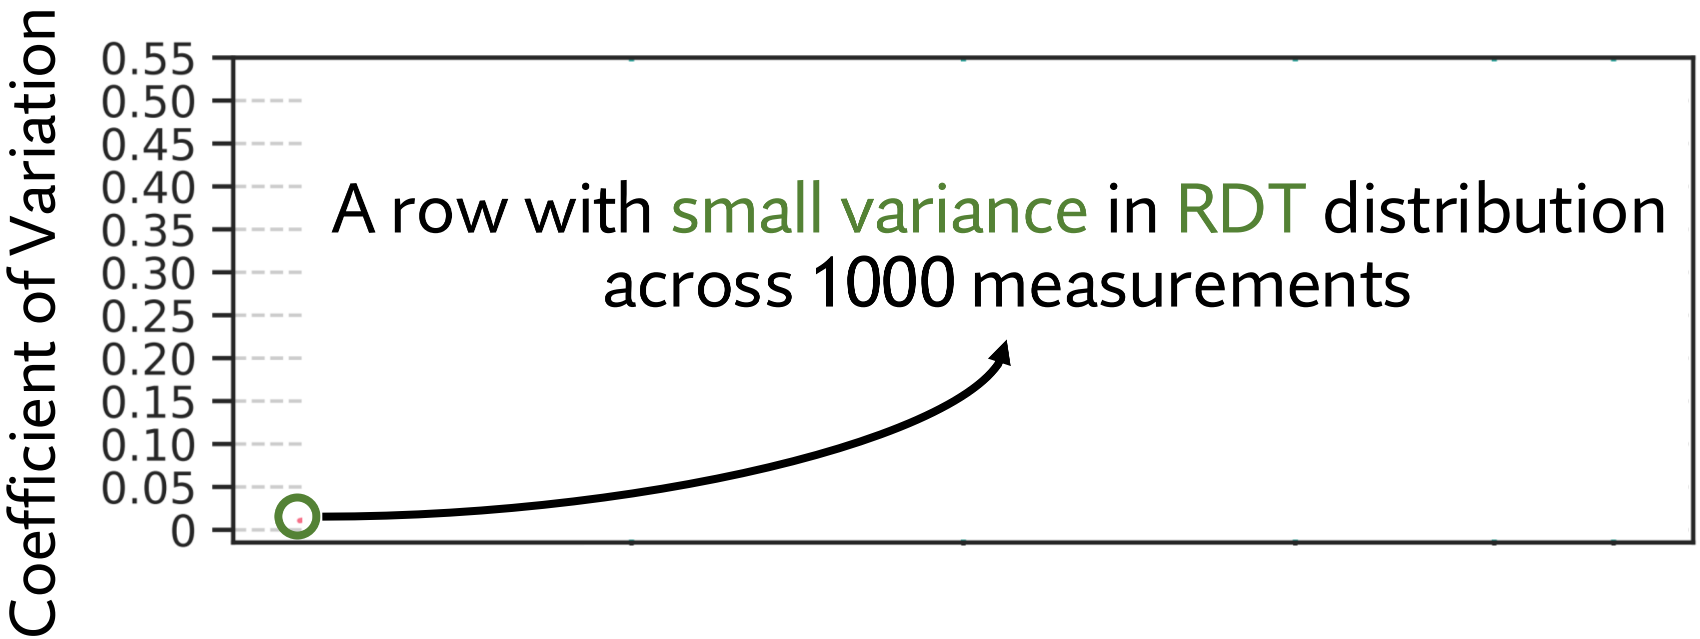

VRD Across DRAM Rows

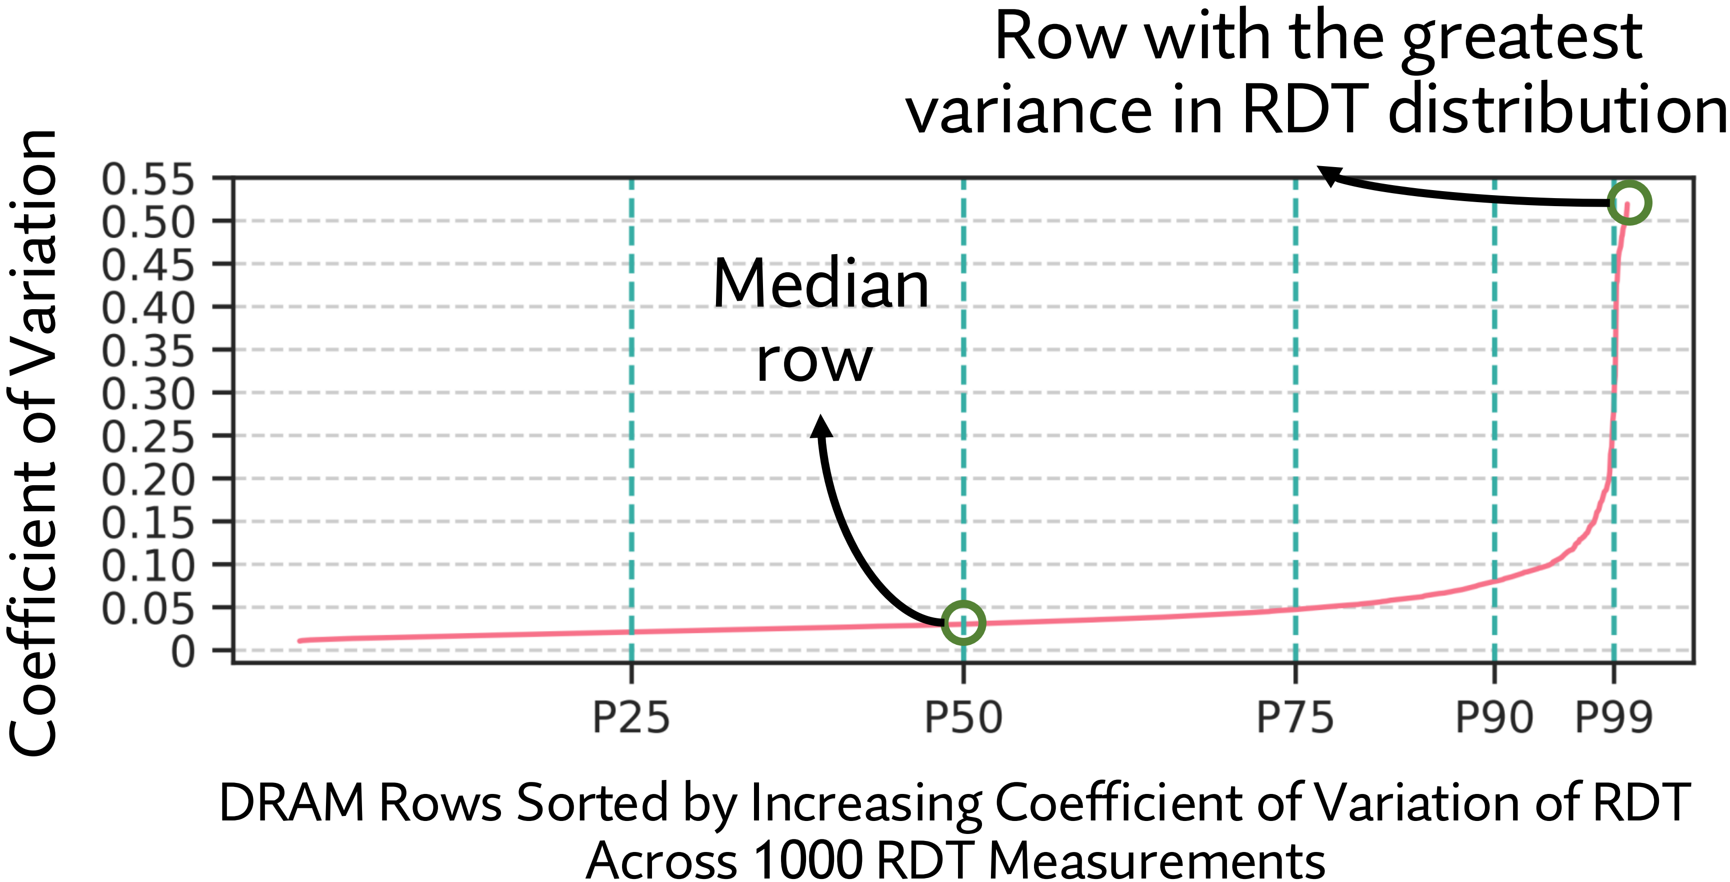

All tested rows exhibit VRD

VRD in Two Example Rows

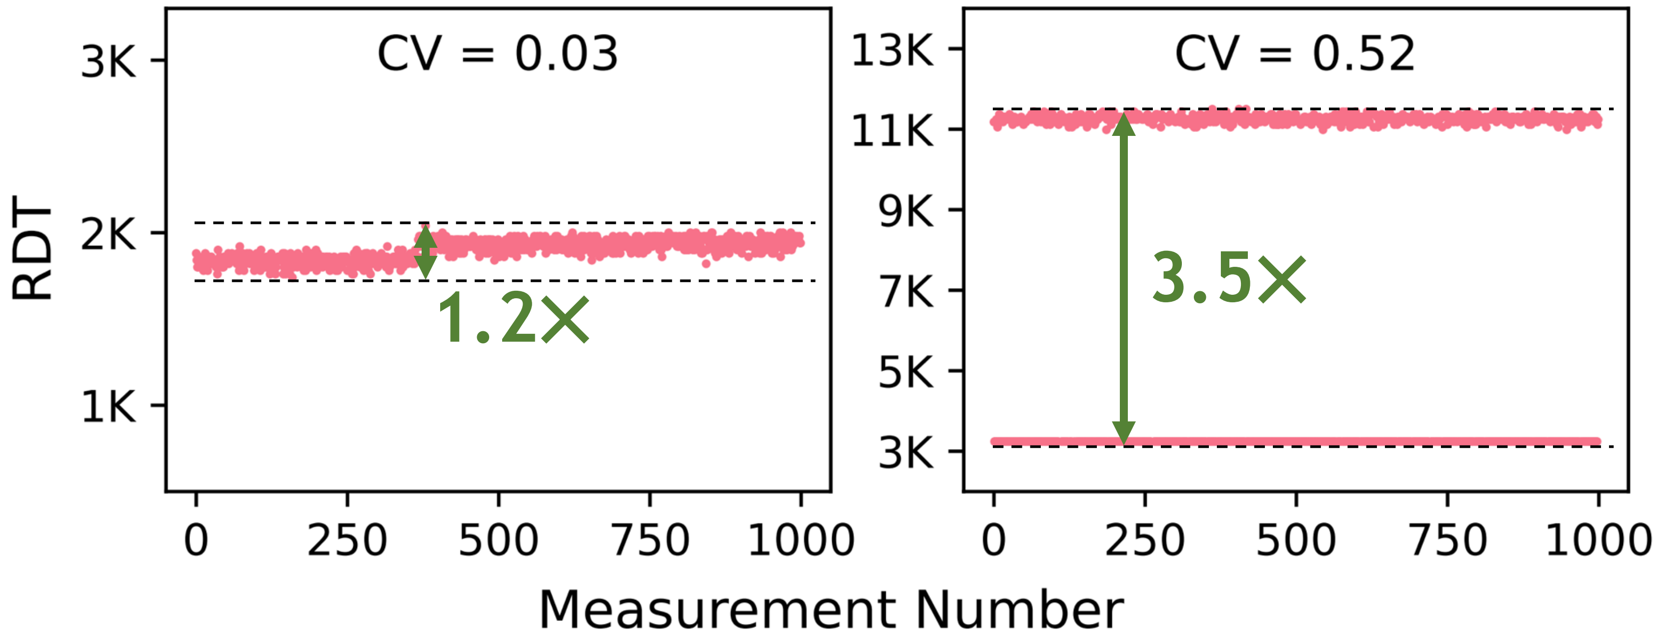

Variation in read disturbance threshold can reach 3.5×

In-Depth Analysis: Key Takeaways

Takeaway 1

All tested DRAM rows exhibit VRD

Takeaway 2

Relatively few (<500) RDT measurements are unlikely to yield the minimum RDT of a row

Takeaway 3

Data patterns, temperature, and aggressor row on time affect VRD

Probability of Identifying the Minimum RDT

- How likely is it that N < 1000 measurements yield the minimum RDT value across 1000 measurements?
- N = 1, 3, 5, 10, 50, and 500
- Monte Carlo simulations for 10K iterations

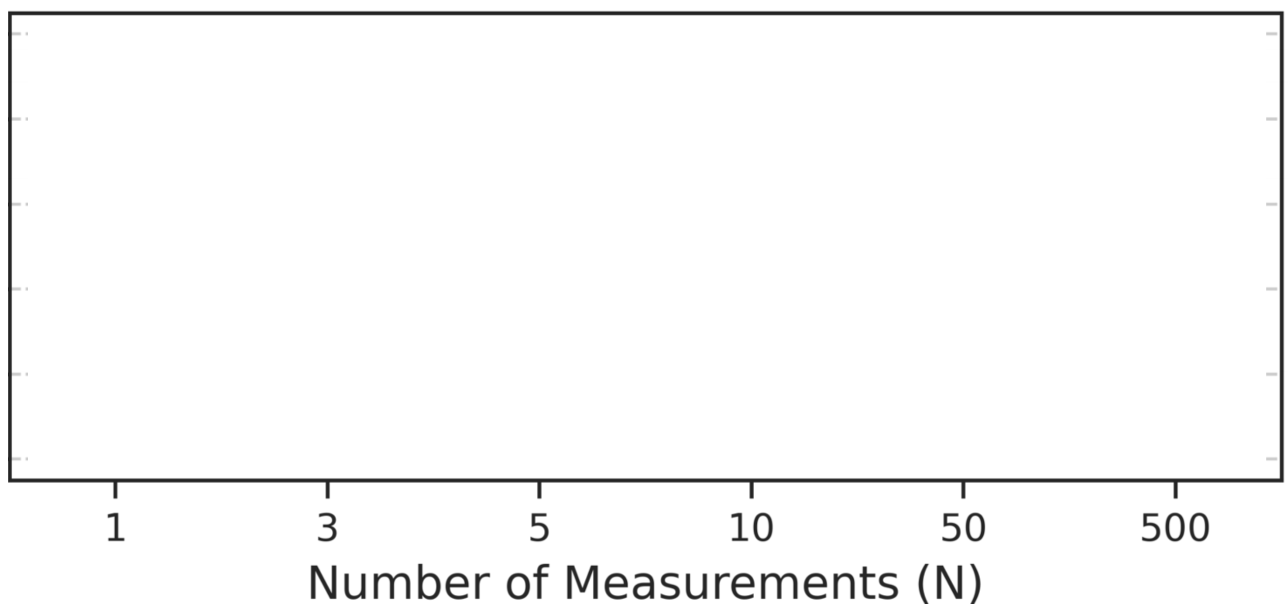

Probability of Identifying the Minimum RDT

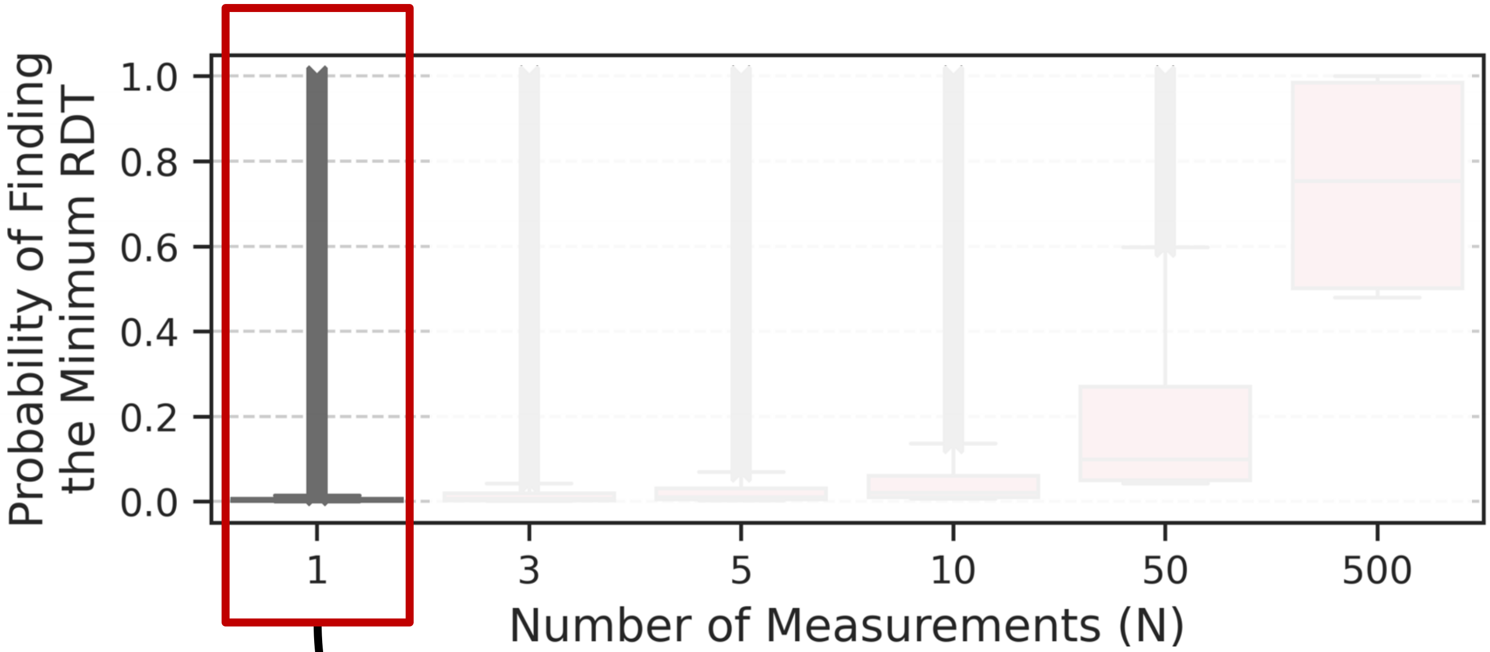

only 0.2% for the median row

Very unlikely to find the minimum RDT of a DRAM row with N = 1 measurement

Probability of Identifying the Minimum RDT

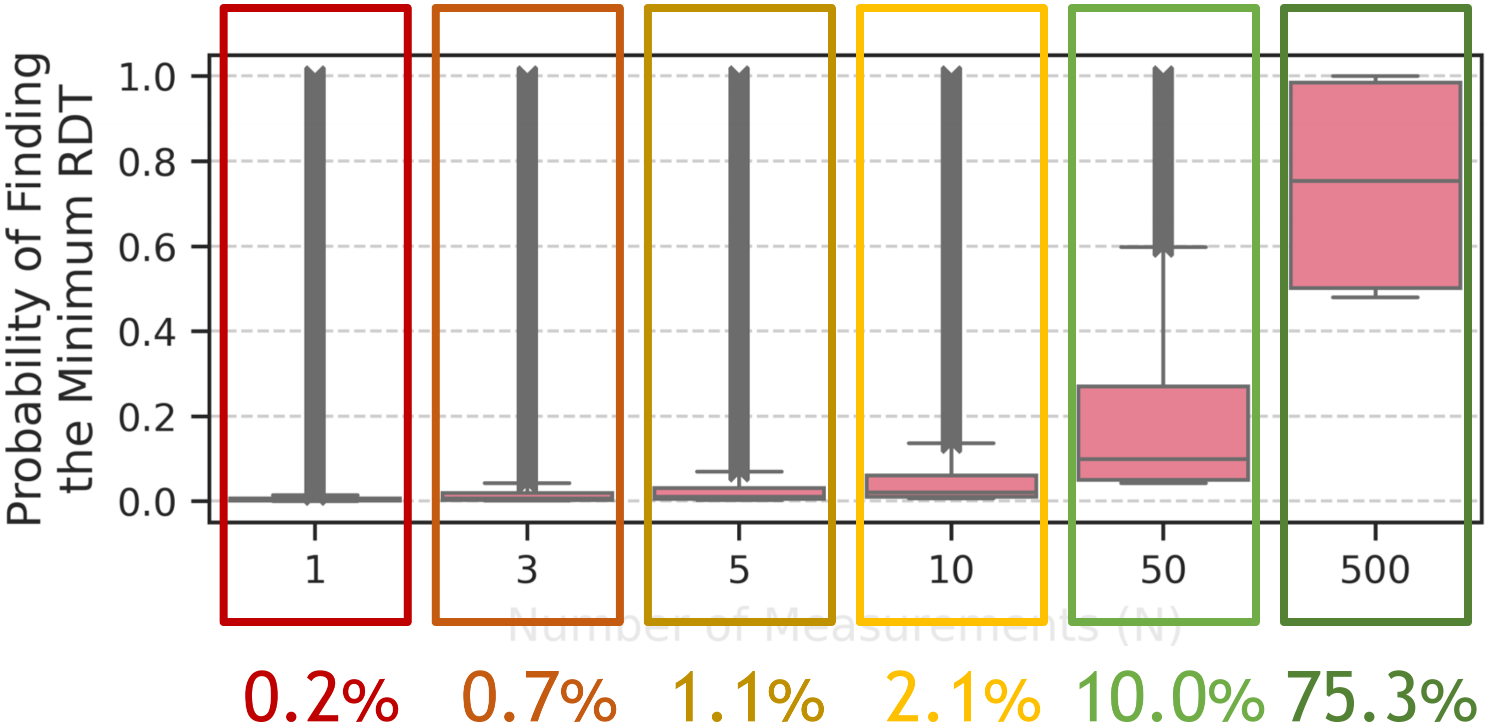

Probability values for the median row

Probability of finding the minimum read disturbance threshold increases with N (i.e., with more and more testing)

Expected Value of the Minimum RDT

- With only N < 1000 RDT measurements
  how far are we from the minimum RDT across 1000 measurements

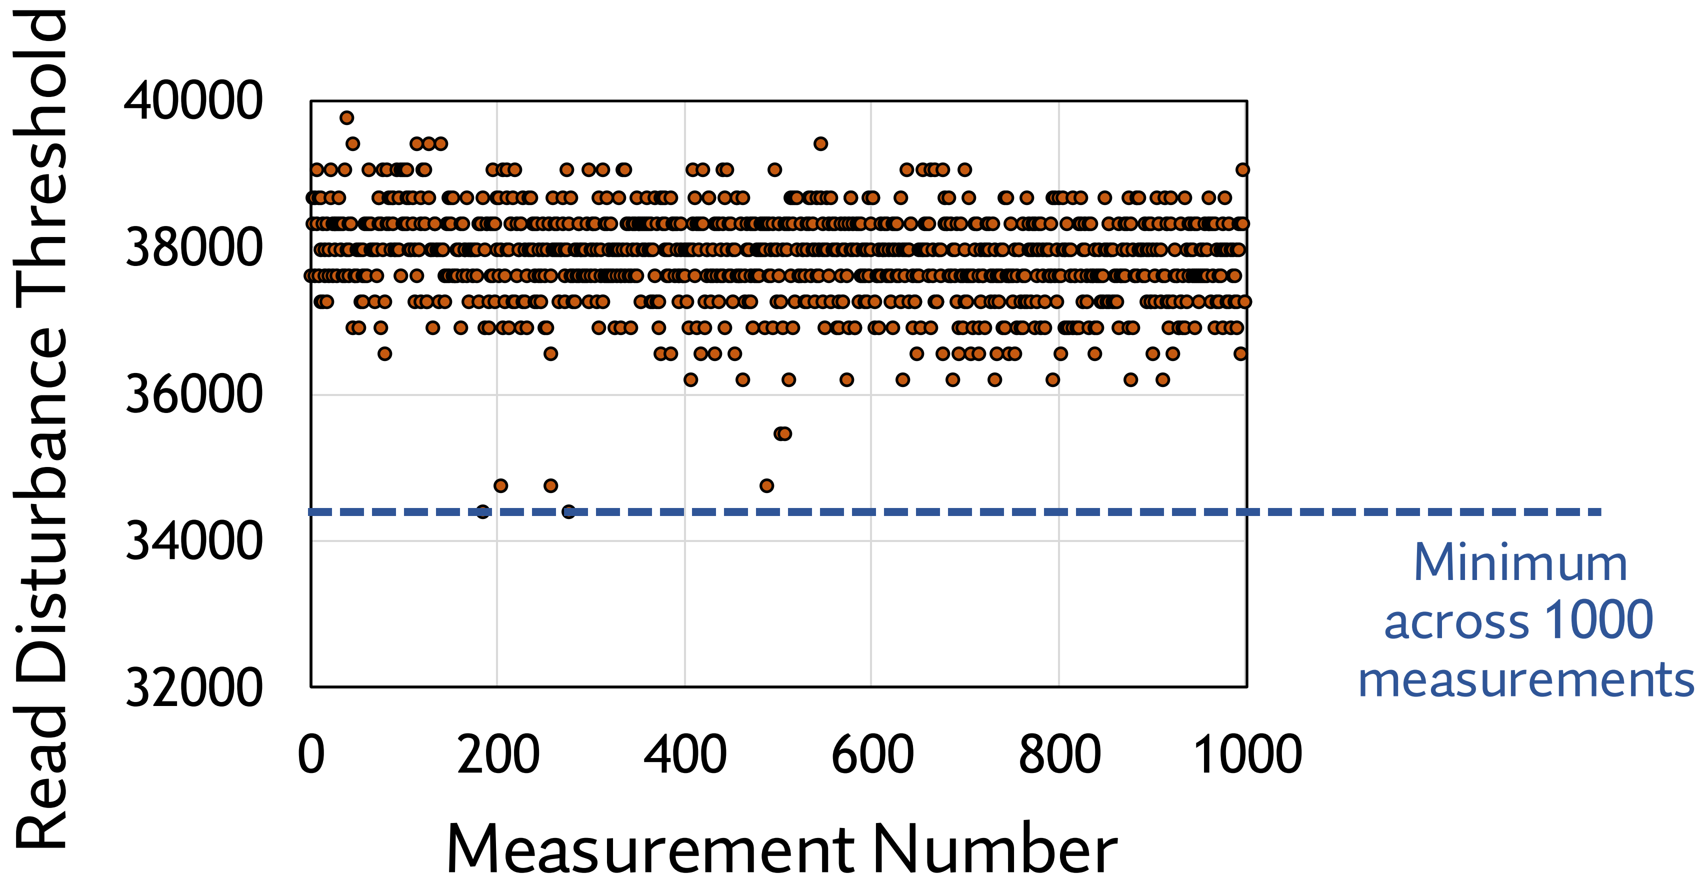

Expected Value of the Minimum RDT

- With only N < 1000 RDT measurements
  how far are we from the minimum RDT across 1000 measurements

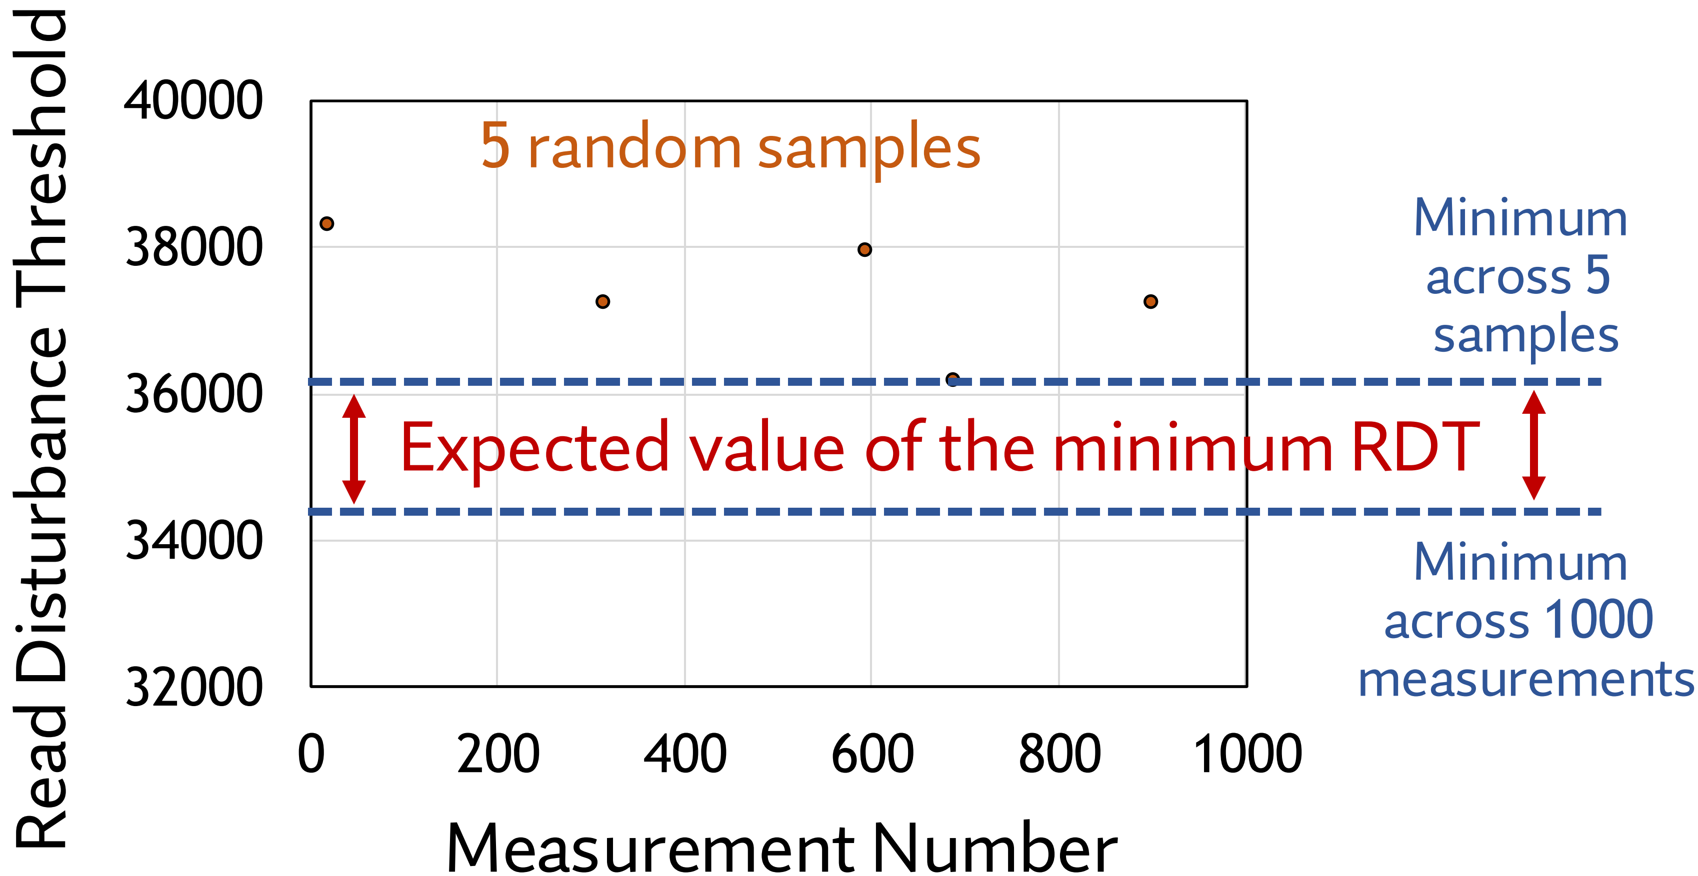

Expected Value of the Minimum RDT

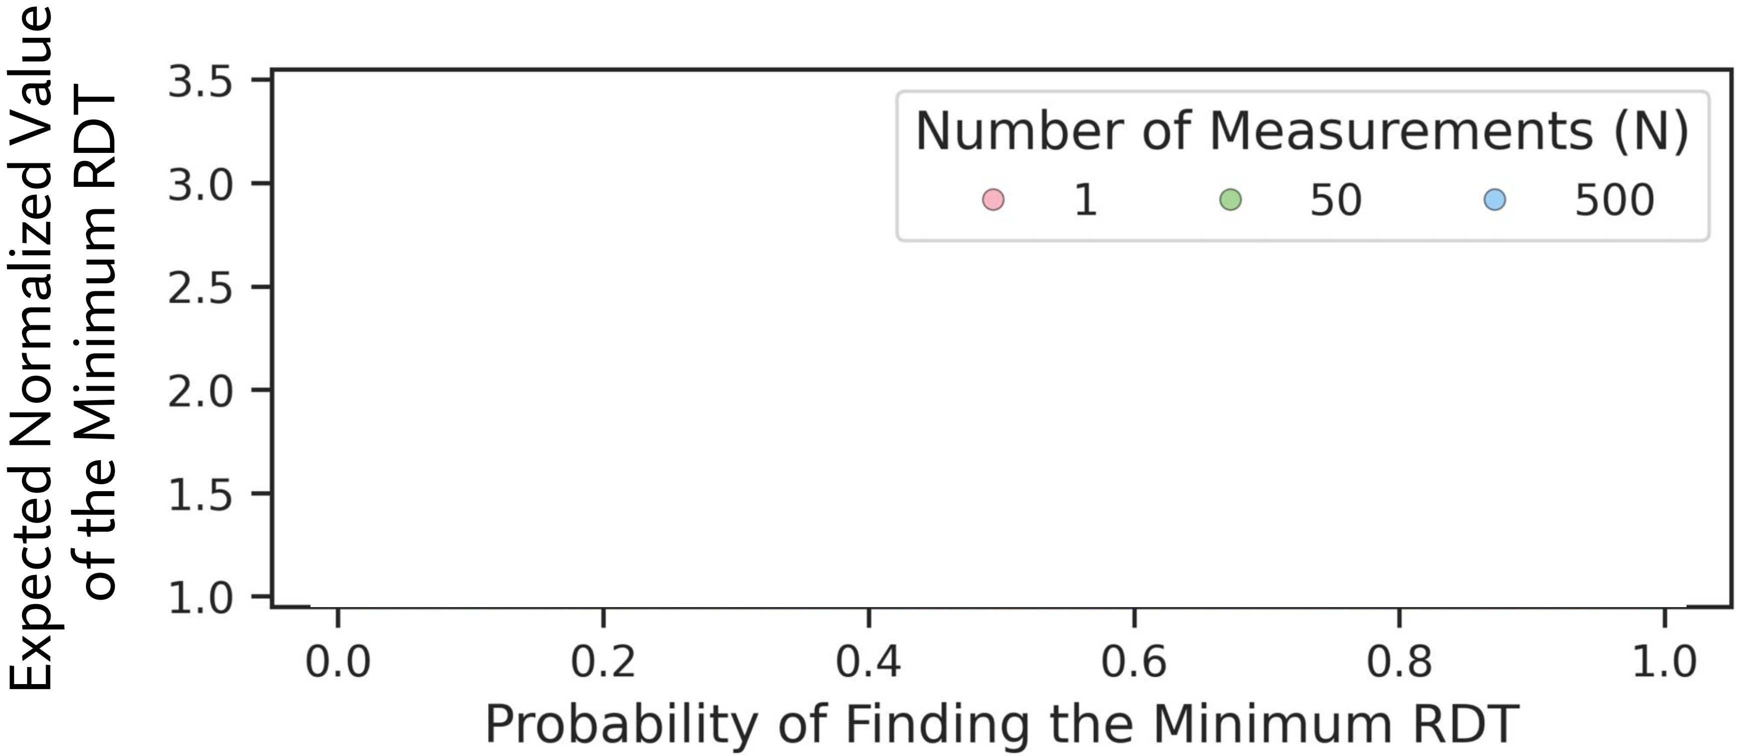

Expected Value of the Minimum RDT

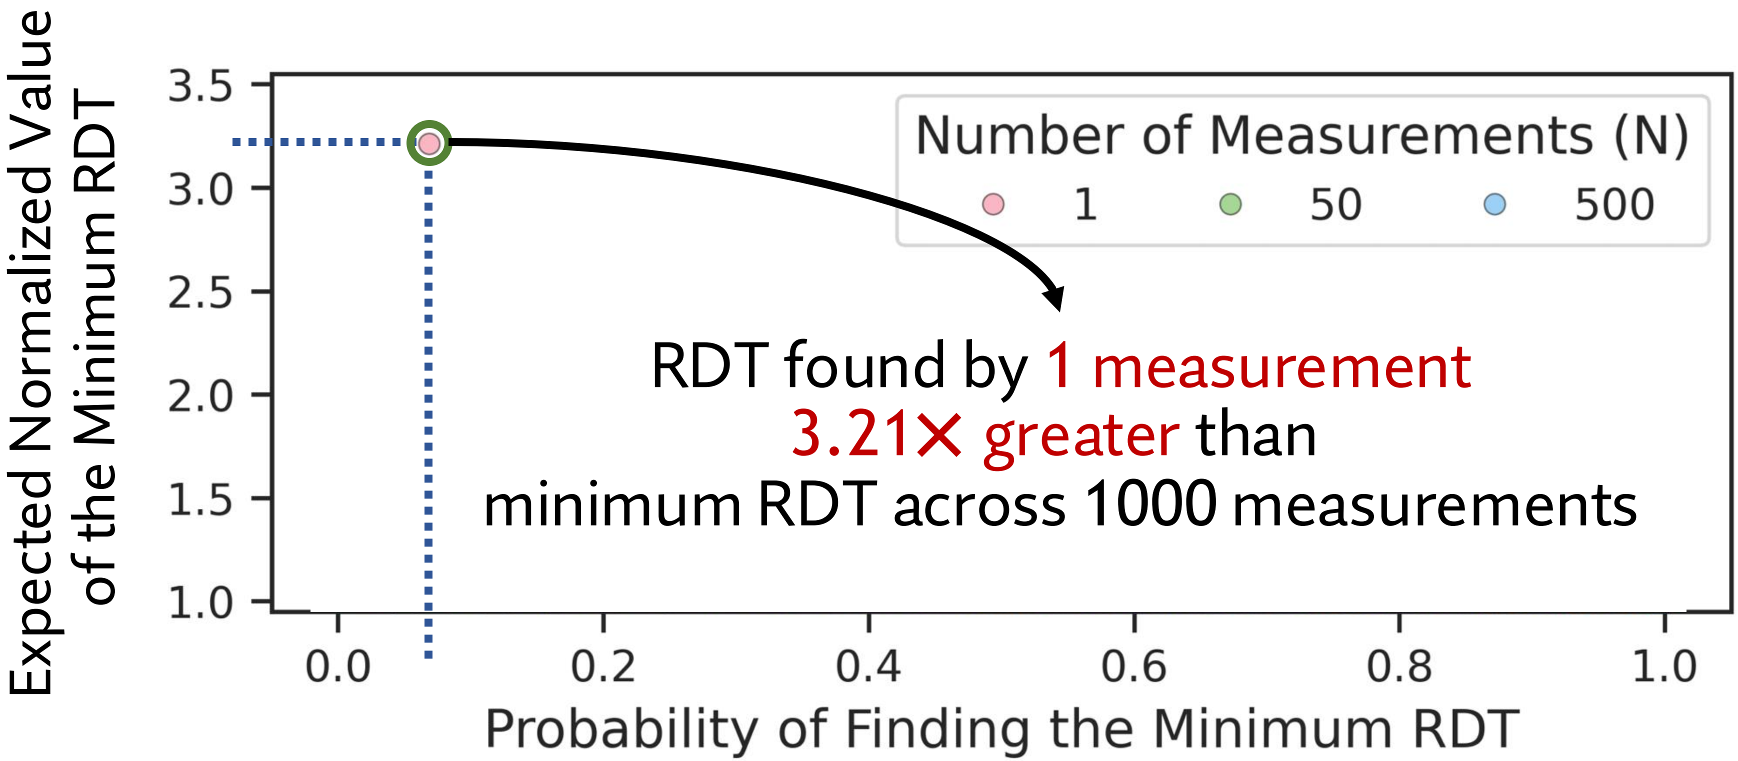

Expected Value of the Minimum RDT

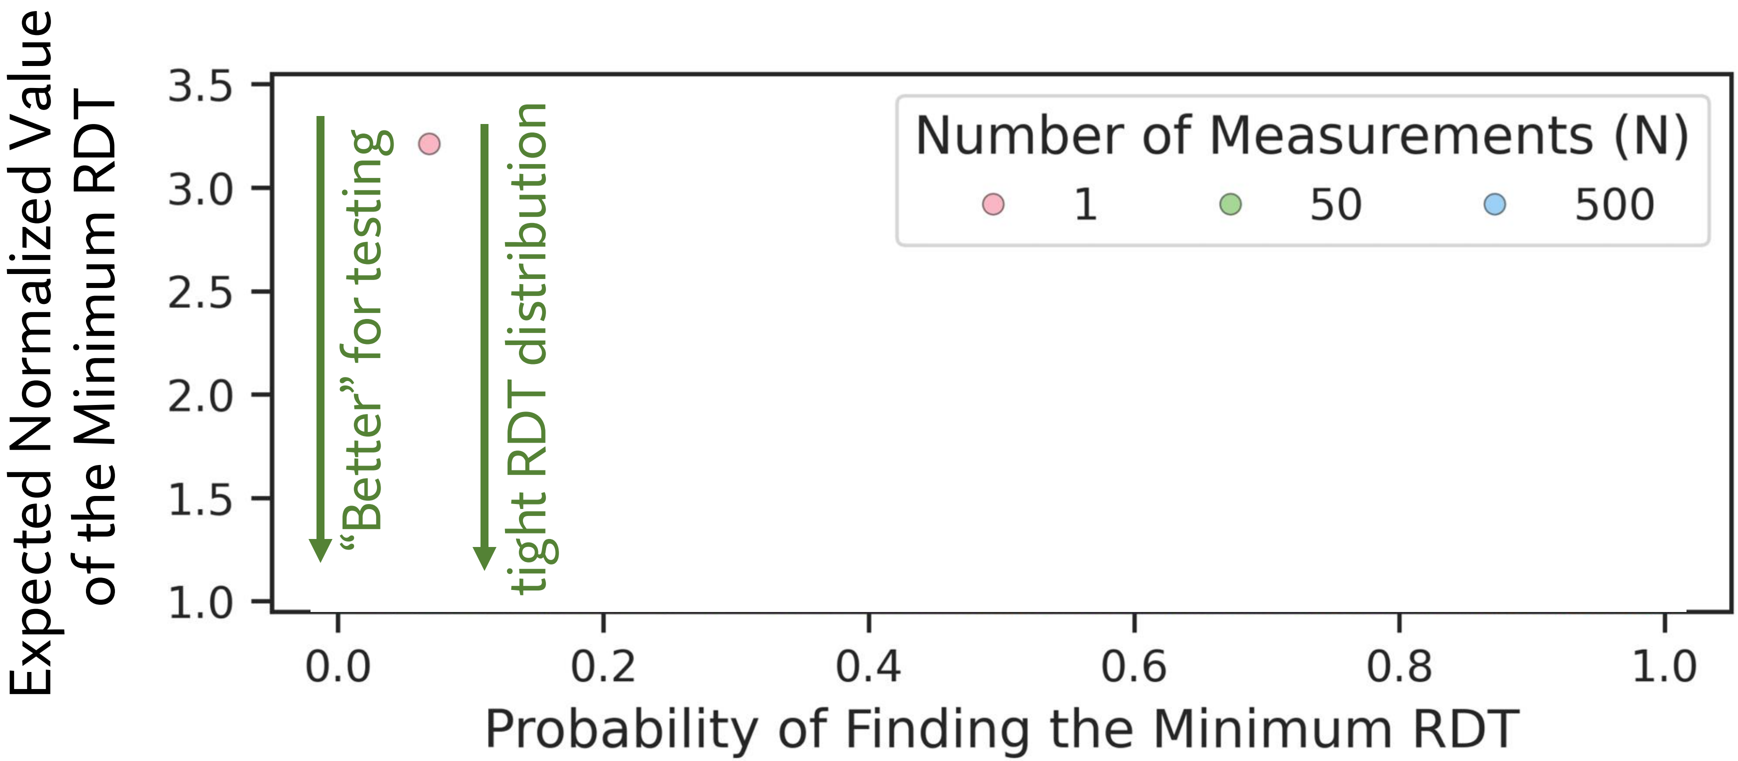

Plot interpretation

"Better" for testing == as tight an RDT distribution as possible

Expected Value of the Minimum RDT

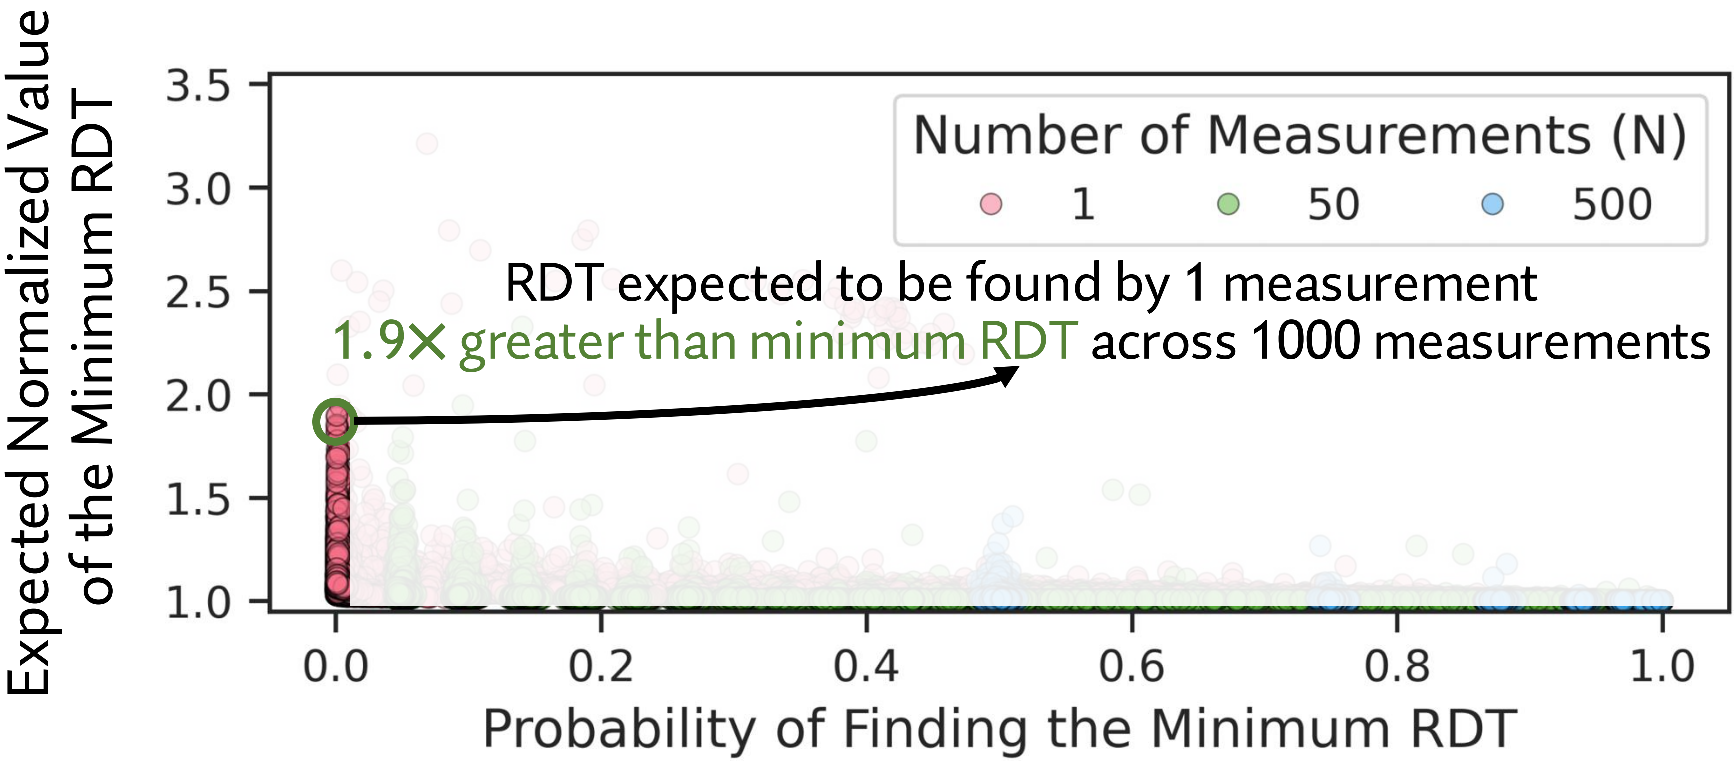

The minimum RDT is significantly smaller than the one expected to be found with N = 1 measurement

Expected Value of the Minimum RDT

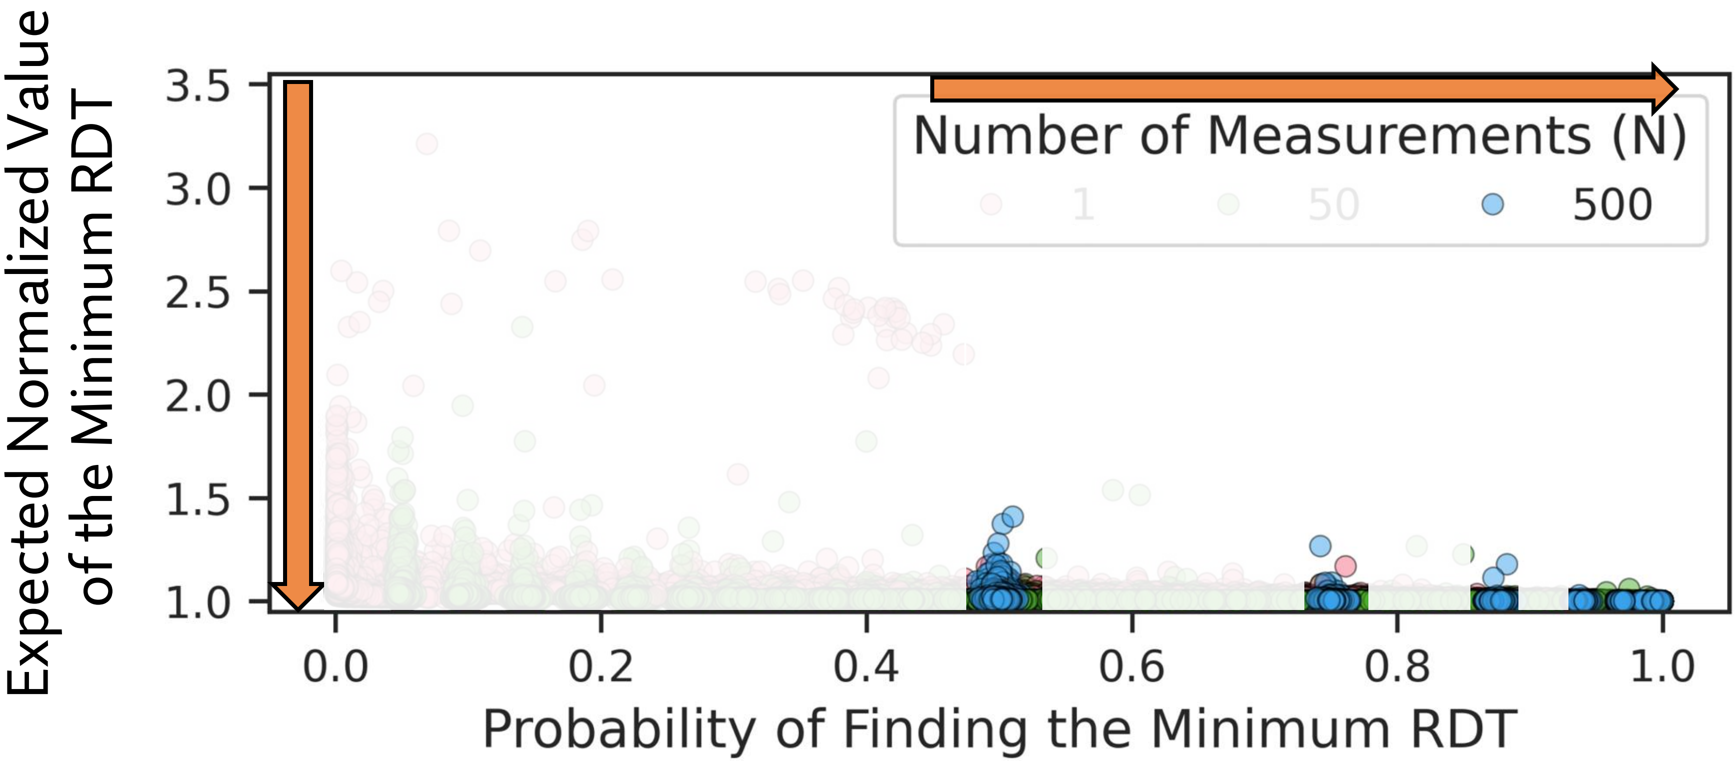

With increasing N (number of measurements) we expect to identify an RDT value closer to the minimum across 1000 measurements

In-Depth Analysis: Key Takeaways

Takeaway 1

All tested DRAM rows exhibit VRD

Takeaway 2

Relatively few (<500) RDT measurements are unlikely to yield the minimum RDT of a row

Takeaway 3

Data patterns, temperature, and aggressor row on time affect VRD

Talk Outline

I. Motivation

II. Experimental Characterization Methodology

III. Foundational Results

IV. In-Depth Analysis of VRD

V. Implications for System Security and Robustness

VI. Conclusion

Implications Summary

- Security guarantees provided by mitigation techniques rely on accurately identified minimum read disturbance threshold (RDT)
- Accurate identification of minimum RDT (for each row) is extremely challenging (even with 1000s measurements) because RDT unpredictably changes over time
- We analyze the use of a guardband for RDT and ECC
  - May prevent VRD-induced bitflips
  - Large guardbands induce performance overhead
- Call for future work on online RDT profiling and runtime configurable read disturbance mitigations

Important Caveat

- VRD solution analysis based on 1K or 10K read disturbance threshold measurements per row

- More measurements could yield **worse** results
  - Read disturbance threshold distribution tail could expand

- What results would millions or billions of RDT measurements yield?

Challenges of Accurately Identifying RDT

Variation in read disturbance threshold across 1000 measurements can reach 3.5× and may not be bounded

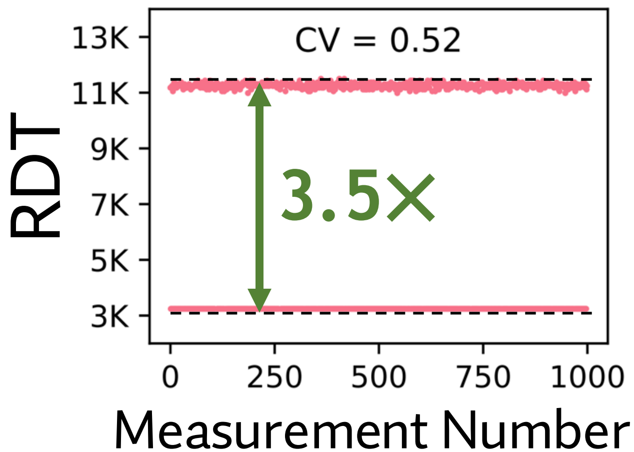

VRD is affected by data pattern, temperature, aggressor row on time

Comprehensive RDT profiling is **time-intensive**

Measuring RDT of *each row only once* with 8000 hammers using four data patterns, at three temperature levels takes 39 minutes in a bank of 256K rows

RDT Profiling is Time-Intensive

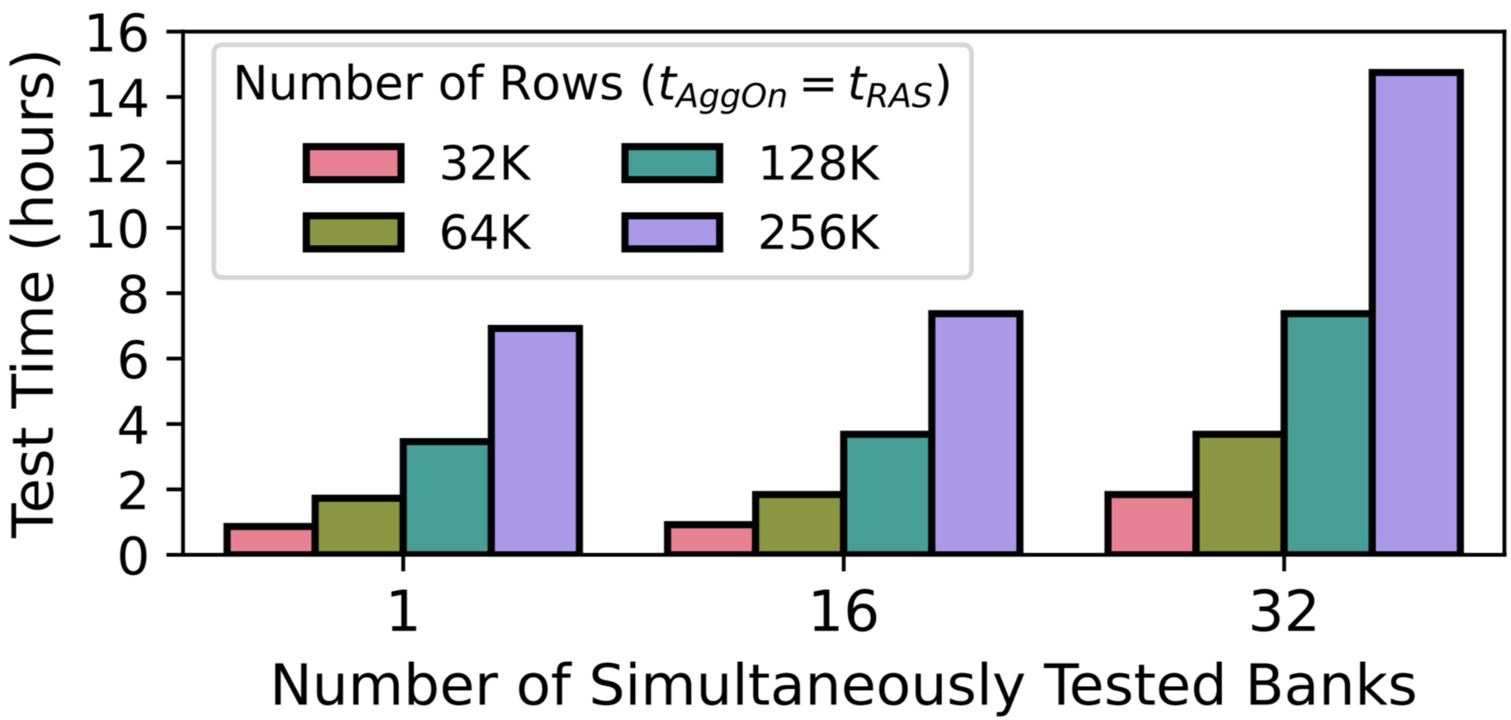

Comprehensive RDT testing can take tens of hours
(***only*** 1000 measurements, one data pattern,
one temperature level, one aggressor row on time)

RDT Profiling is Time-Intensive

https://arxiv.org/pdf/2502.13075

Variable Read Disturbance:
An Experimental Analysis of Temporal Variation in DRAM Read Disturbance

Ataberk Olgun† F. Nisa Bostancı† İsmail Emir Yüksel† Oğuzhan Canpolat† Haocong Luo†
Geraldo F. Oliveira† A. Giray Yağlıkçı† Minesh Patel‡ Onur Mutlu†

ETH Zurich† *Rutgers University‡*

Modern DRAM chips are subject to read disturbance errors. These errors manifest as security-critical bitflips in a victim DRAM row *that is physically nearby a repeatedly activated (opened) aggressor row (RowHammer) or an aggressor row that is kept open for a long time (RowPress). State-of-the-art read disturbance mitigations rely on accurate and exhaustive characterization of the* read disturbance threshold *(RDT) (e.g., the number of aggressor row activations needed to induce the first RowHammer or RowPress bitflip) of every DRAM row (of which there are millions or billions in a modern system) to prevent read disturbance bitflips securely and with low overhead.*

We experimentally demonstrate for the first time that the RDT of a DRAM row significantly and unpredictably changes over time. We call this new phenomenon variable read disturbance (VRD). Our extensive experiments using 160 DDR4 chips and 4 HBM2 chips from three major manufacturers yield three key observations. First, it is very unlikely that relatively few RDT measurements can accurately identify the RDT of a DRAM row. The minimum RDT of a DRAM row appears after tens of thousands of measurements (e.g., up to 94,467), and the minimum RDT of a DRAM row is 3.5× smaller than the maximum RDT observed for that row. Second, the probability of accu-

row) *many times* (e.g., tens of thousands of times) induces *RowHammer bitflips* in physically nearby rows (i.e., victim rows) [1]. Keeping the aggressor row open for a long period of time amplifies the effects of read disturbance and induces *RowPress bitflips*, *without* requiring *many* repeated aggressor row activations [4].

A large body of work [1, 3, 26, 32, 39, 45, 69–141] proposes various techniques to mitigate DRAM read disturbance bitflips. Many high-performance and low-overhead mitigation techniques [1, 73, 74, 76, 79, 82–84, 86, 87, 91, 97, 133–135, 137–139, 142–146], including those that are used and standardized by industry [121, 126, 138, 139, 144], prevent read disturbance bitflips by *preventively* refreshing (i.e., opening and closing) a victim row *before* a bitflip manifests in that row.

To securely prevent read disturbance bitflips at low performance and energy overhead, it is important to *accurately* identify the amount of read disturbance that a victim row can withstand before experiencing a read disturbance bitflip. This amount is typically quantified using the *hammer count (the number of aggressor row activations) needed to induce the first read disturbance bitflip* in a victim row. We call this metric the *read disturbance threshold (RDT)* of the victim row.

Making Do With Few RDT Measurements

- A system designer might measure RDT a few times
  and apply a safety margin (guardband) to the minimum observed value

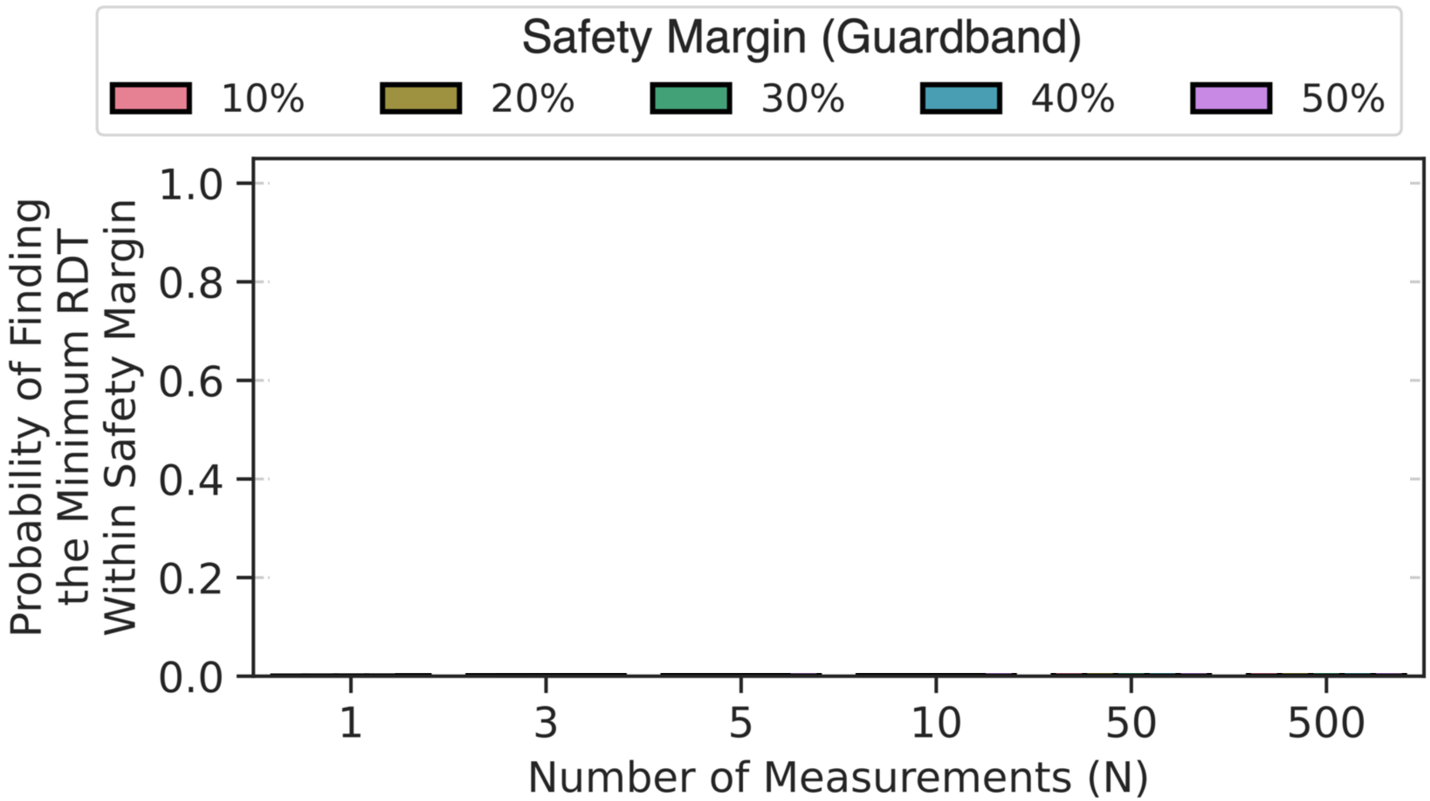

Making Do With Few RDT Measurements

- A system designer might measure RDT a few times and apply a safety margin (guardband) to the minimum observed value

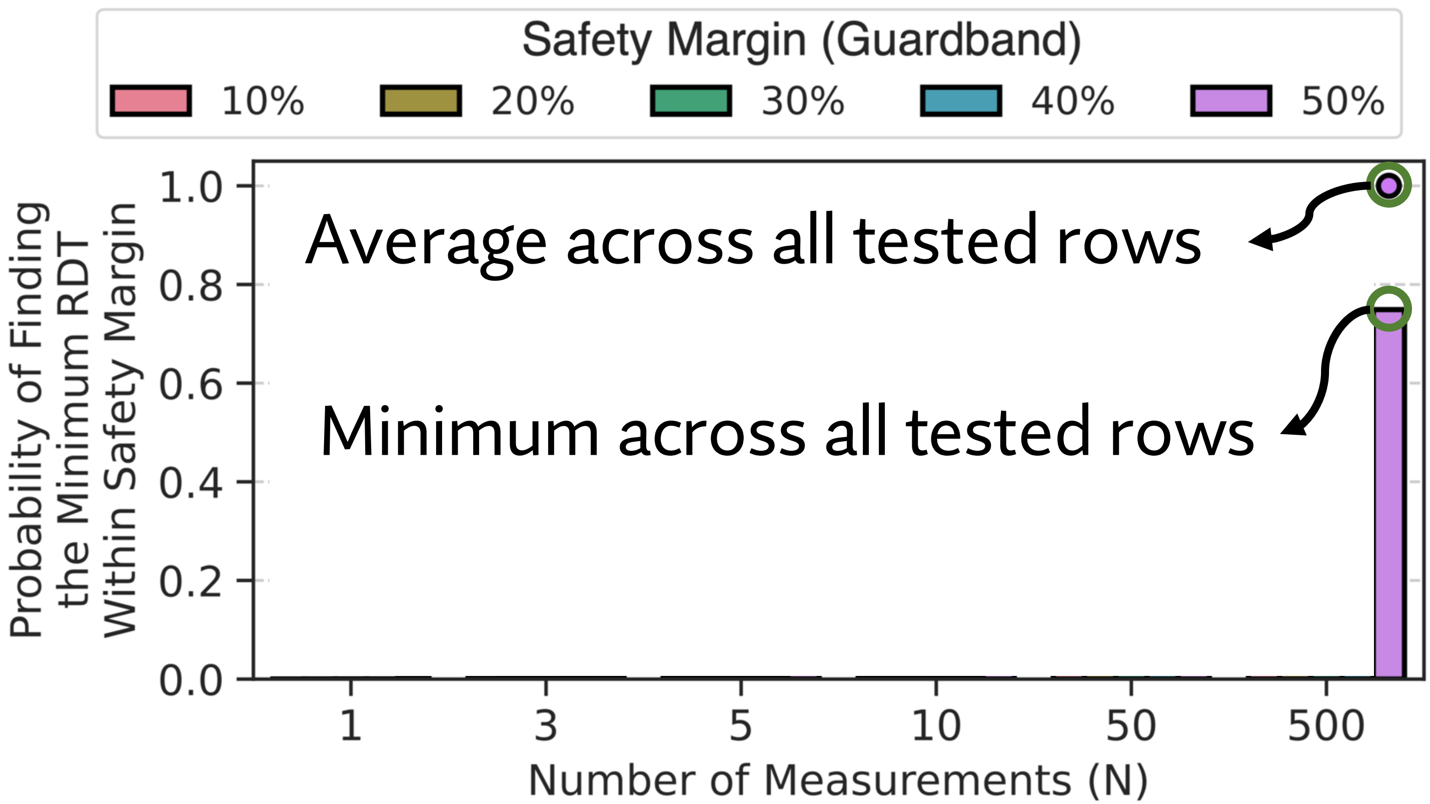

Making Do With Few RDT Measurements

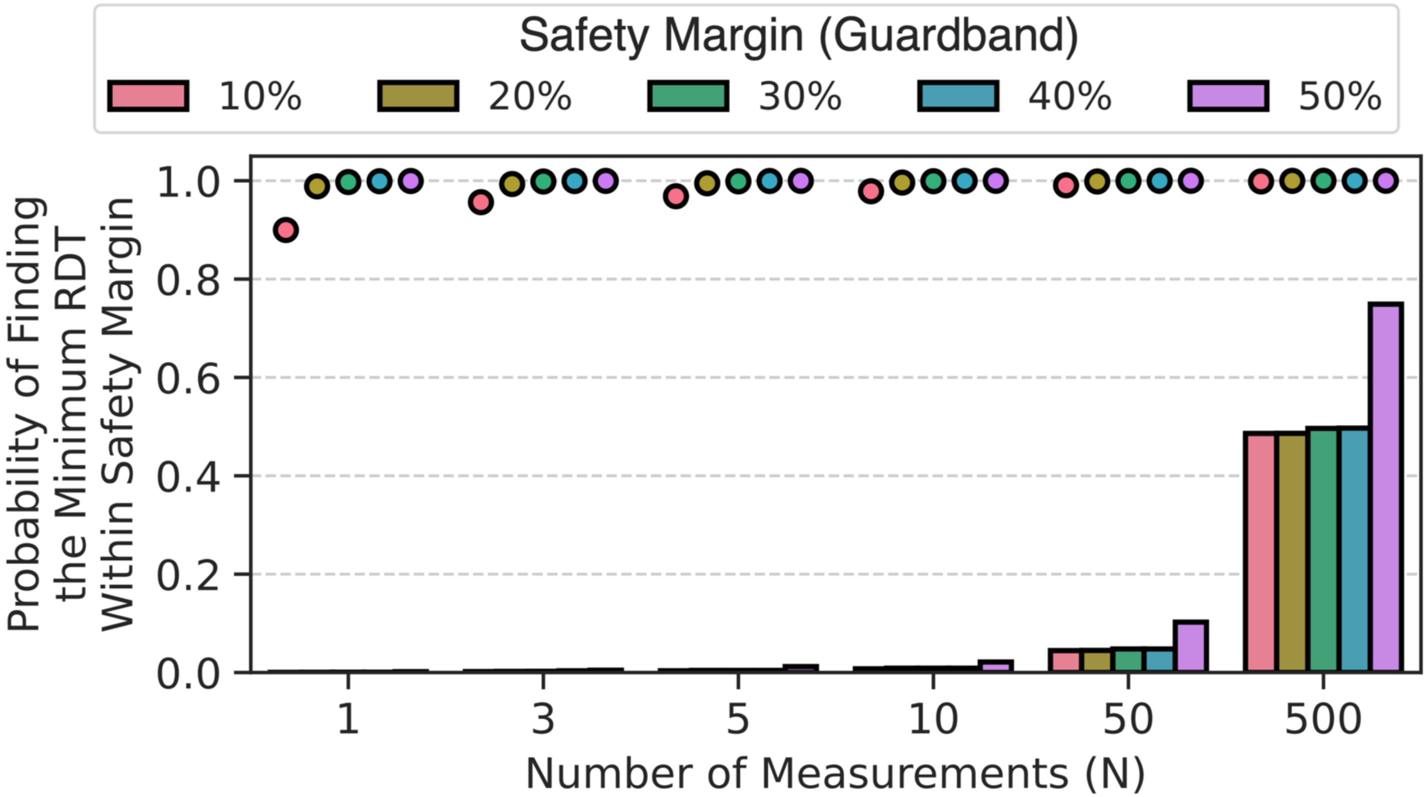

A large guardband does not guarantee that the minimum RDT is always identified

Using guardbands alone is likely not effective

RDT Guardband Increases Performance Overheads

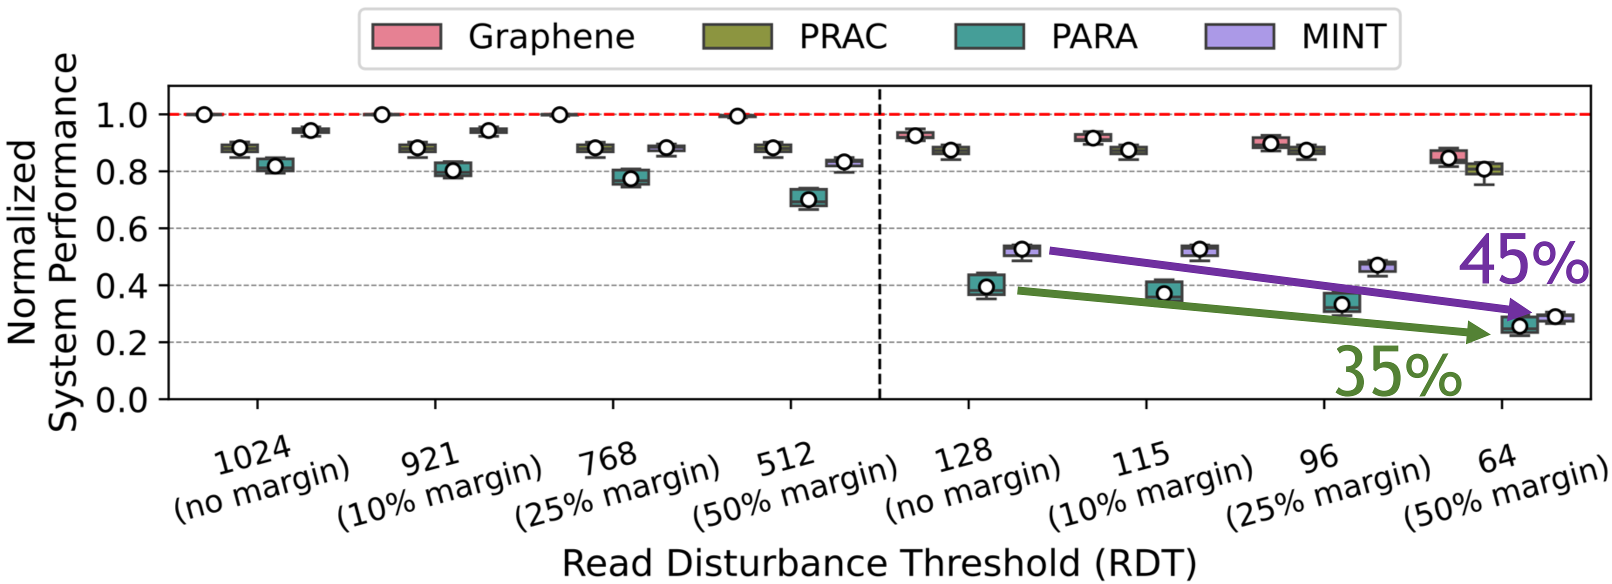

50% RDT safety margin can induce
45% additional overhead (over no margin)

Relying **solely** on guardbands **not** recommended

Combining ECC and Guardbands (I)

- Single-error correcting double-error detecting (SECDED) or Chipkill ECC combined with guardbands could mitigate VRD-induced bitflips

Unique bitflips when 10% RDT guardband applied

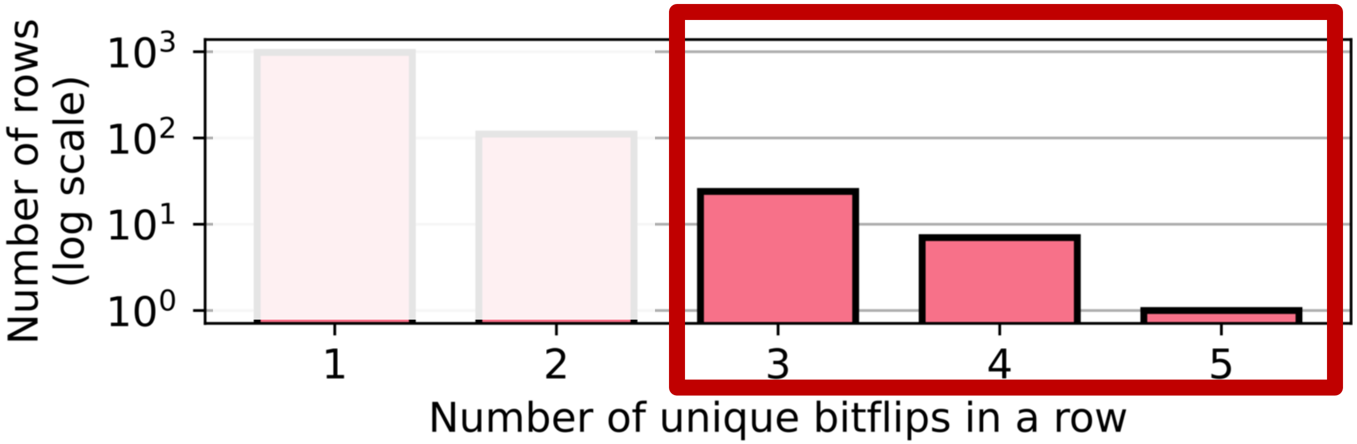

10% guardband combined w/ ECC is likely unsafe

Combining ECC and Guardbands (II)

RDT guardbands ≥20% yield 1 unique bitflip in a row

Given our limited measurement dataset (10K measurements)
RDT guardbands ≥20% combined with ECC
may prevent VRD-induced read disturbance bitflips

More detailed analysis (following a large-scale study)
needed to make a definitive conclusion

More in the Paper

- Hypothetical explanation for VRD
- Effect of True- and Anti-Cell Layout
  - Presence of true- and anti-cells in the victim row does not significantly affect the RDT distribution
- Read disturbance mitigation evaluation methodology
- Probability of errors at the worst observed bitflip rate for 10% RDT guardband
  - SEC, SECDED, and Chipkill-like (SSC)
- Read disturbance testing time and energy consumption
- Detailed information on tested modules and chips

More in the Paper

https://arxiv.org/pdf/2502.13075

Variable Read Disturbance:
An Experimental Analysis of Temporal Variation in DRAM Read Disturbance

Ataberk Olgun† F. Nisa Bostancı† İsmail Emir Yüksel† Oğuzhan Canpolat† Haocong Luo†
Geraldo F. Oliveira† A. Giray Yağlıkçı† Minesh Patel‡ Onur Mutlu†

ETH Zurich† *Rutgers University‡*

Modern DRAM chips are subject to read disturbance errors. These errors manifest as security-critical bitflips in a victim DRAM row *that is physically nearby a repeatedly activated (opened) aggressor row (RowHammer) or an aggressor row that is kept open for a long time (RowPress). State-of-the-art read disturbance mitigations rely on accurate and exhaustive characterization of the* read disturbance threshold *(RDT) (e.g., the number of aggressor row activations needed to induce the first RowHammer or RowPress bitflip) of every DRAM row (of which there are millions or billions in a modern system) to prevent read disturbance bitflips securely and with low overhead.*

We experimentally demonstrate for the first time that the RDT of a DRAM row significantly and unpredictably changes over time. We call this new phenomenon variable read disturbance (VRD). Our extensive experiments using 160 DDR4 chips and 4 HBM2 chips from three major manufacturers yield three key observations. First, it is very unlikely that relatively few RDT measurements can accurately identify the RDT of a DRAM row. The minimum RDT of a DRAM row appears after tens of thousands of measurements (e.g., up to 94,467), and the minimum RDT of a DRAM row is 3.5× smaller than the maximum RDT observed for that row. Second, the probability of accu-

row) *many times* (e.g., tens of thousands of times) induces *RowHammer bitflips* in physically nearby rows (i.e., victim rows) [1]. Keeping the aggressor row open for a long period of time amplifies the effects of read disturbance and induces *RowPress bitflips*, *without* requiring *many* repeated aggressor row activations [4].

A large body of work [1, 3, 26, 32, 39, 45, 69–141] proposes various techniques to mitigate DRAM read disturbance bitflips. Many high-performance and low-overhead mitigation techniques [1, 73, 74, 76, 79, 82–84, 86, 87, 91, 97, 133–135, 137–139, 142–146], including those that are used and standardized by industry [121, 126, 138, 139, 144], prevent read disturbance bitflips by *preventively* refreshing (i.e., opening and closing) a victim row *before* a bitflip manifests in that row.

To securely prevent read disturbance bitflips at low performance and energy overhead, it is important to *accurately* identify the amount of read disturbance that a victim row can withstand before experiencing a read disturbance bitflip. This amount is typically quantified using the *hammer count (the number of aggressor row activations) needed to induce the first read disturbance bitflip* in a victim row. We call this metric the *read disturbance threshold (RDT)* of the victim row.

Talk Outline

I. Motivation

II. Experimental Characterization Methodology

III. Foundational Results

IV. In-Depth Analysis of VRD

V. Implications for System Security and Robustness

VI. Conclusion

VRD Conclusion

Variable Read Disturbance (VRD)

The read disturbance threshold changes unpredictably over time

Minimum RDT (of a row) may appear after many measurements

RDT for a DRAM row can vary by 3.5X

Identifying the minimum RDT is challenging and time-intensive

Given our limited read disturbance bitflip dataset,
guardbands combined with error-correcting codes
may be a solution for VRD-induced bitflips.

More data and analyses needed to make definitive conclusion.

Future work could alleviate the shortcomings of existing mitigations & develop better understanding of inner workings of VRD

Extended Version on arXiv

https://arxiv.org/pdf/2502.13075

arXiv > cs > arXiv:2502.13075

Search... | All fields | Search

Help | Advanced Search

Computer Science > Hardware Architecture

[Submitted on 18 Feb 2025]

Variable Read Disturbance: An Experimental Analysis of Temporal Variation in DRAM Read Disturbance

Ataberk Olgun, F. Nisa Bostanci, Ismail Emir Yuksel, Oguzhan Canpolat, Haocong Luo, Geraldo F. Oliveira, A. Giray Yaglikci, Minesh Patel, Onur Mutlu

Modern DRAM chips are subject to read disturbance errors. State-of-the-art read disturbance mitigations rely on accurate and exhaustive characterization of the read disturbance threshold (RDT) (e.g., the number of aggressor row activations needed to induce the first RowHammer or RowPress bitflip) of every DRAM row (of which there are millions or billions in a modern system) to prevent read disturbance bitflips securely and with low overhead. We experimentally demonstrate for the first time that the RDT of a DRAM row significantly and unpredictably changes over time. We call this new phenomenon variable read disturbance (VRD). Our experiments using 160 DDR4 chips and 4 HBM2 chips from three major manufacturers yield two key observations. First, it is very unlikely that relatively few RDT measurements can accurately identify the RDT of a DRAM row. The minimum RDT of a DRAM row appears after tens of thousands of measurements (e.g., up to 94,467), and the minimum RDT of a DRAM row is 3.5X smaller than the maximum RDT observed for that row. Second, the probability of accurately identifying a row's RDT with a relatively small number of measurements reduces with increasing chip density or smaller technology node size. Our empirical results have implications for the security guarantees of read disturbance

Access Paper:

- View PDF
- TeX Source
- Other Formats

Current browse context:
cs.AR

< prev | next >

new | recent | 2025-02

Change to browse by:

cs

cs.CR

References & Citations

- NASA ADS
- Google Scholar
- Semantic Scholar

Export BibTeX Citation

Bookmark

Variable Read Disturbance (VRD)
An Experimental Analysis of Temporal Variation in DRAM Read Disturbance

Ataberk Olgun, F. Nisa Bostancı, İsmail Emir Yüksel
Oğuzhan Canpolat, Haocong Luo, Geraldo F. Oliveira
A. Giray Yağlıkçı, Minesh Patel, Onur Mutlu

https://arxiv.org/pdf/2502.13075

ETH zürich SAFARI RUTGERS THE STATE UNIVERSITY OF NEW JERSEY

Variable Read Disturbance

Backup Slides

Variable Read Disturbance (VRD) Summary

Research Question

- How accurately and efficiently can we measure the RDT of each DRAM row?

Experimental Characterization

- Record >100M RDT measurements across 3750 rows and many test parameters (e.g., temperature, data pattern) in 160 DDR4 and 4 HBM2 chips

Key Observations

- RDT changes significantly and unpredictably over time: VRD
  - Smallest RDT value (for a row) may appear after 94,467 measurements
  - Maximum observed RDT for a tested row can be 3.5X higher than minimum

Implications for System Security and Robustness

- RDT cannot be accurately identified quickly
- RDT guardbands (>10%) and ECC (SECDED or Chipkill) could prevent VRD-induced bitflips at significant performance cost

VRD Summary (I)

Motivation

- Read Disturbance Threshold (RDT) quantifies a DRAM row's read disturbance vulnerability
  - e.g., number of row activations needed to induce the first bitflip
- Read disturbance mitigation security depends on accurately-identified RDT for each DRAM row

Research Question

- How accurately and efficiently can we measure the RDT of each DRAM row?

Experimental Characterization

- Record >100M RDT measurements across 3750 rows and many test parameters (e.g., temperature, data pattern) in 160 DDR4 and 4 HBM2 chips

VRD Summary (II)

Key Observations

- RDT changes significantly and unpredictably over time: VRD
- The smallest RDT value (for a row) may appear after 94,467 successive RDT measurements
- The maximum observed RDT for a tested row is 3.5X higher than the minimum observed for that row

Implications for System Security and Robustness

- RDT cannot be accurately identified quickly
- RDT guardbands (>10%) and ECC (SECDED or Chipkill) could prevent VRD-induced bitflips at significant performance cost
  - 10% and 50% RDT guardbands respectively induce 6% and 45% performance overhead
- Call for future work on online profiling and runtime configurable read disturbance mitigation techniques

Mitigating Read Disturbance Bitflips

- Key Idea: Take an action before bitflip manifests

- Glass filled with water analogy
  - ACT -> fill with some water
  - ACT (keep open) -> fill with more water
  - Spill -> bitflip
  - Drink before spill -> no bitflip (and more room for water)

DRAM Read Disturbance Bitflips

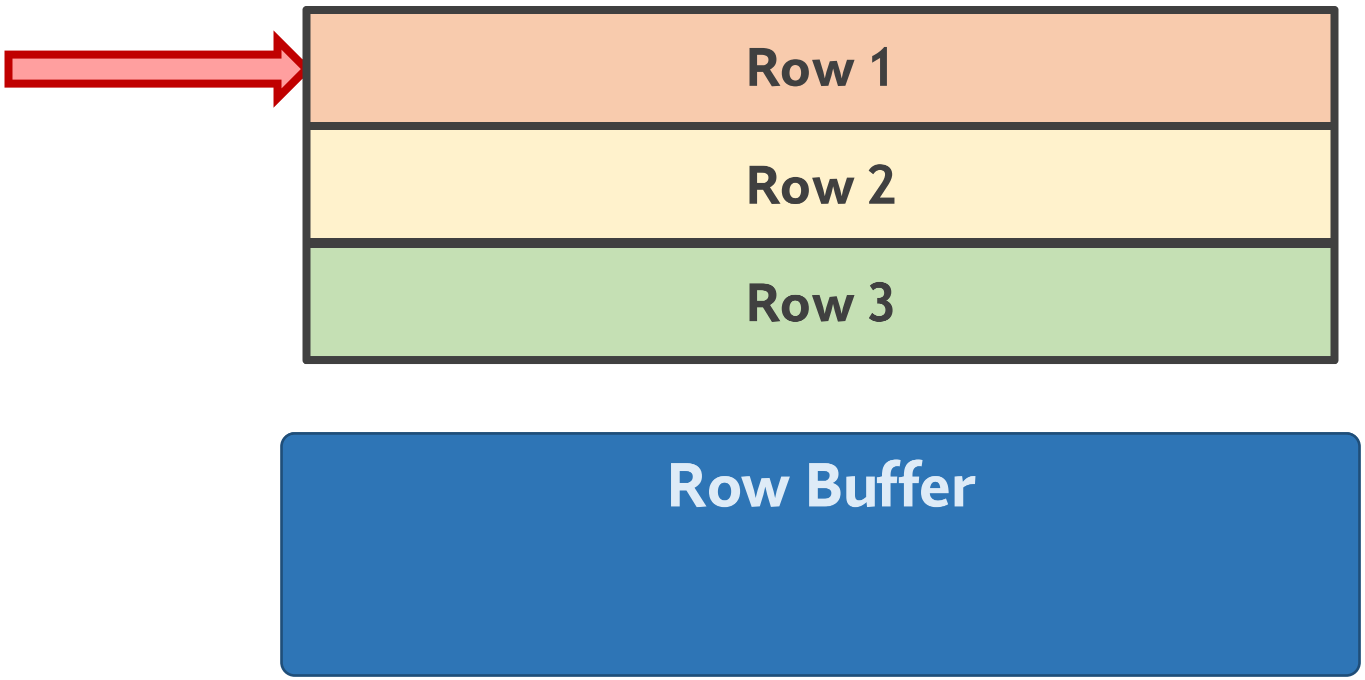

DRAM Read Disturbance Bitflips

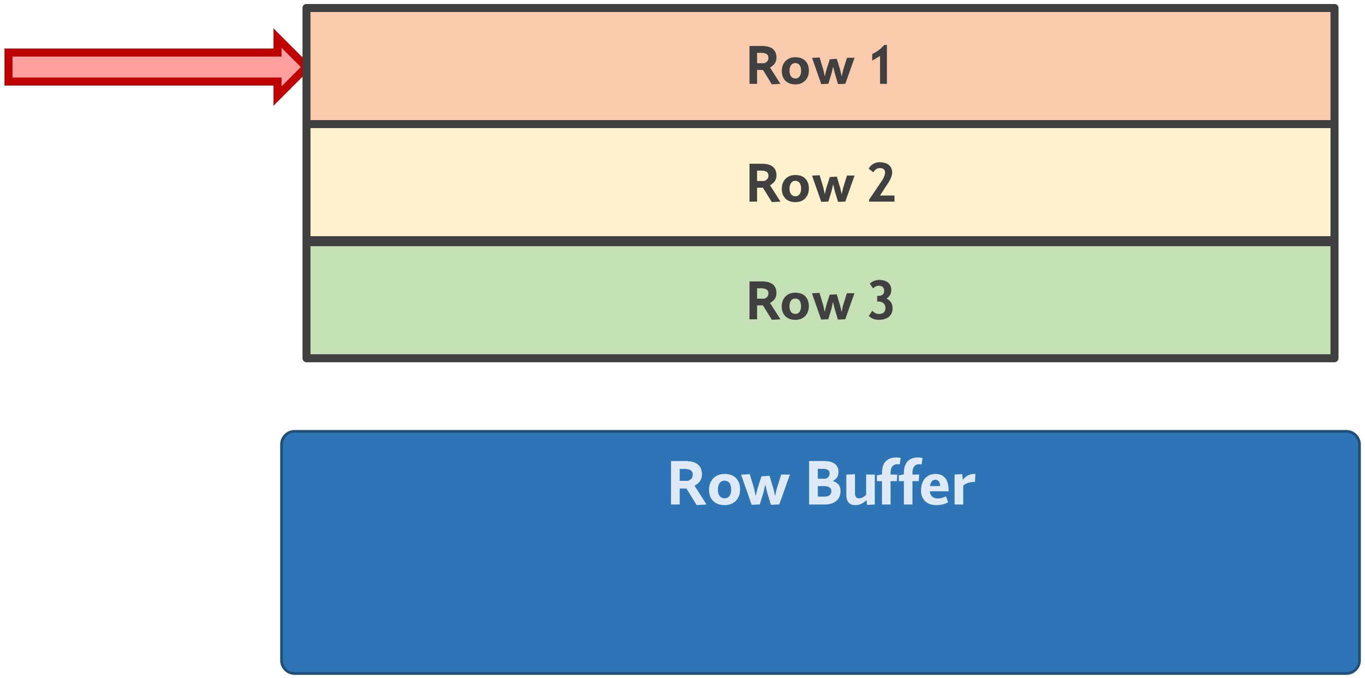

DRAM Read Disturbance Bitflips

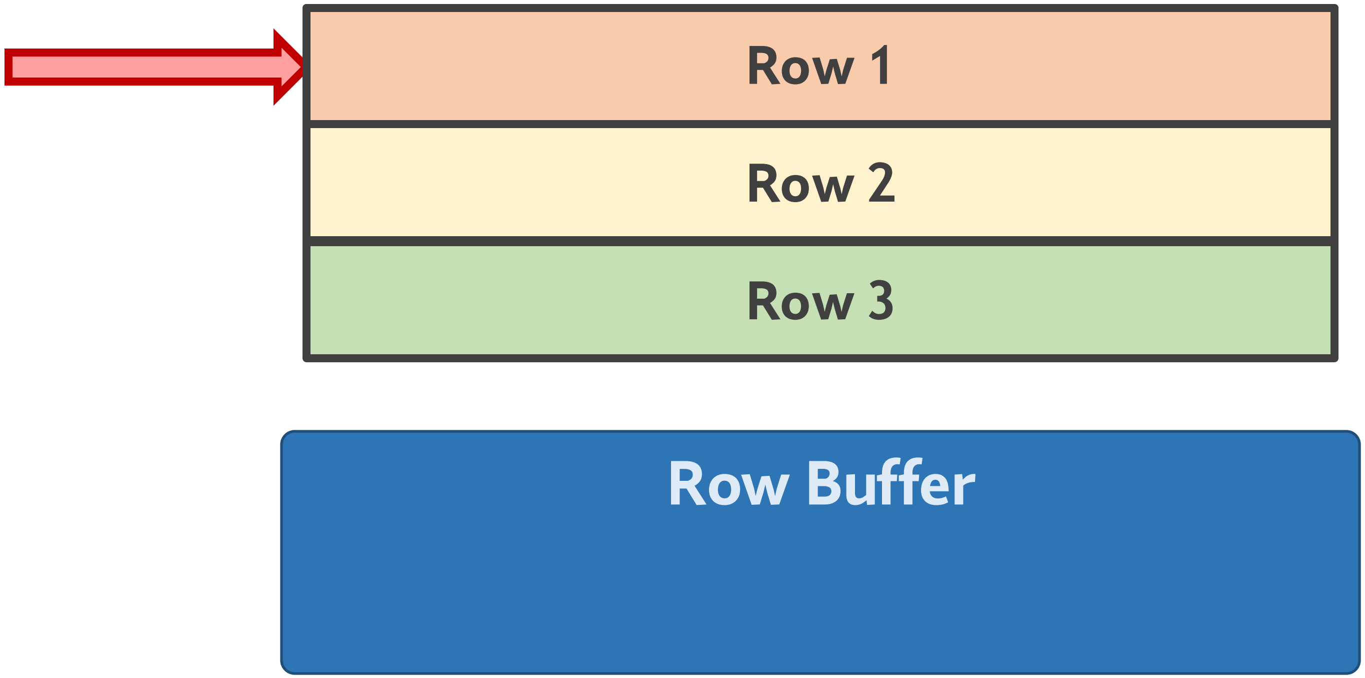

DRAM Read Disturbance Bitflips

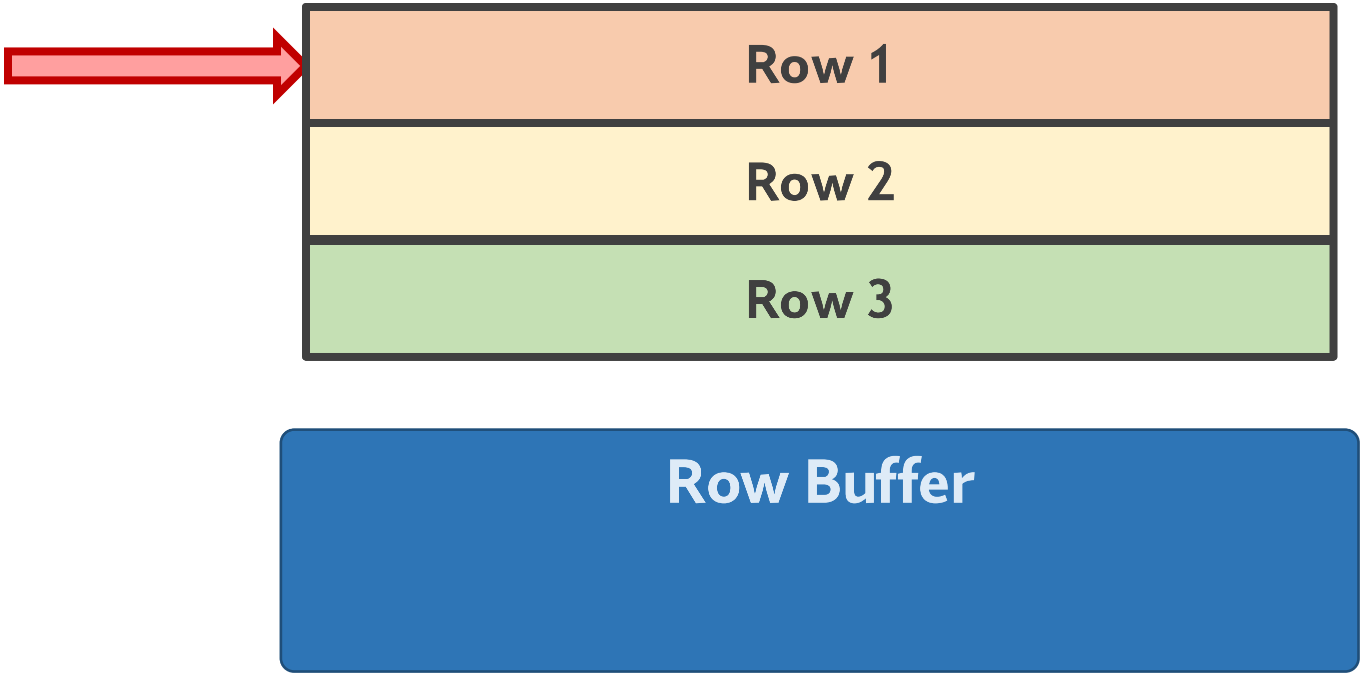

Expected Value of the Minimum RDT

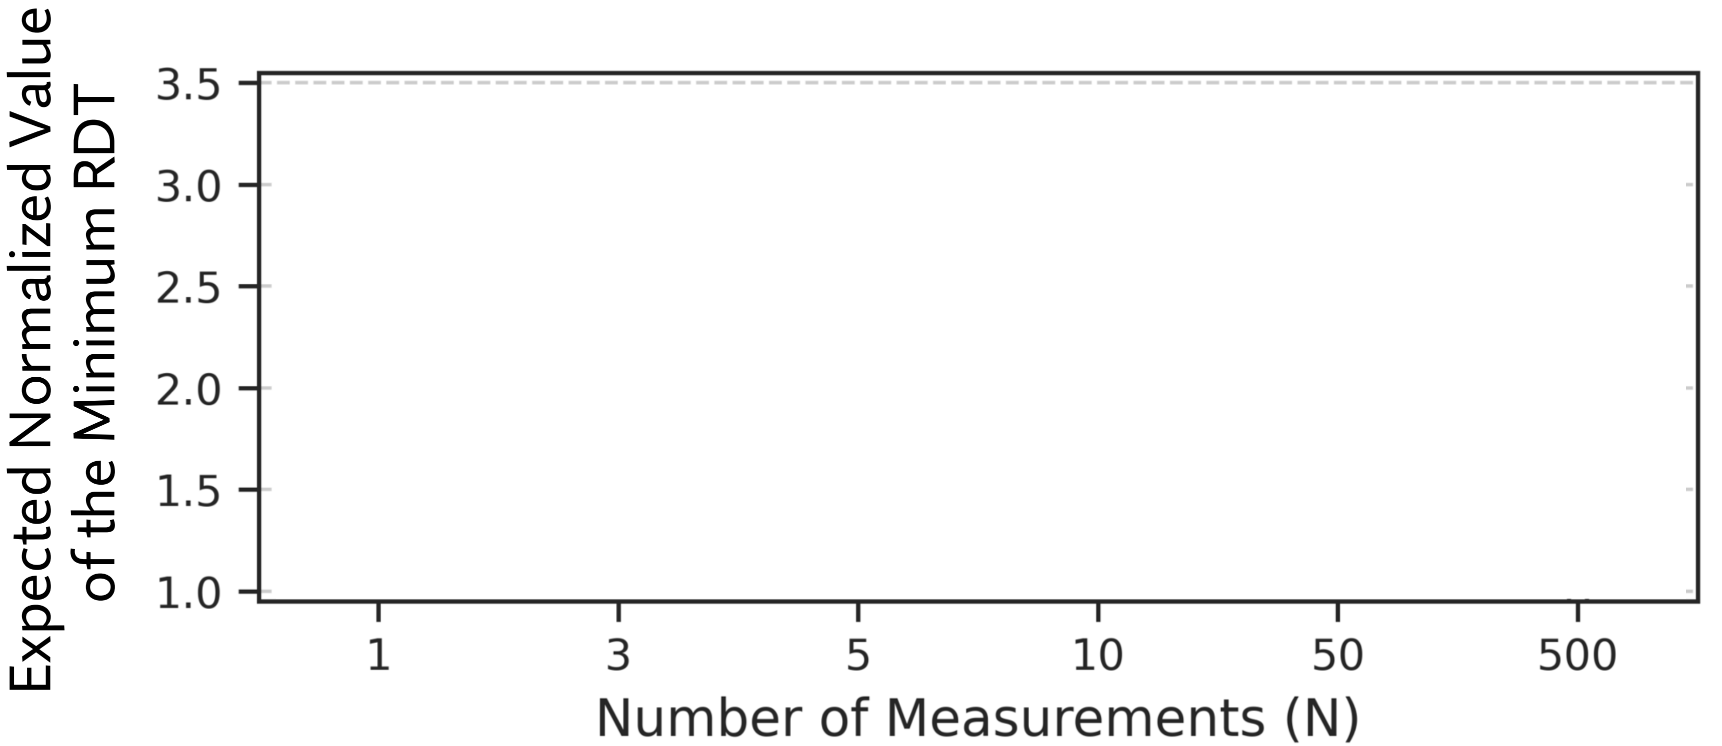

Expected Value of the Minimum RDT

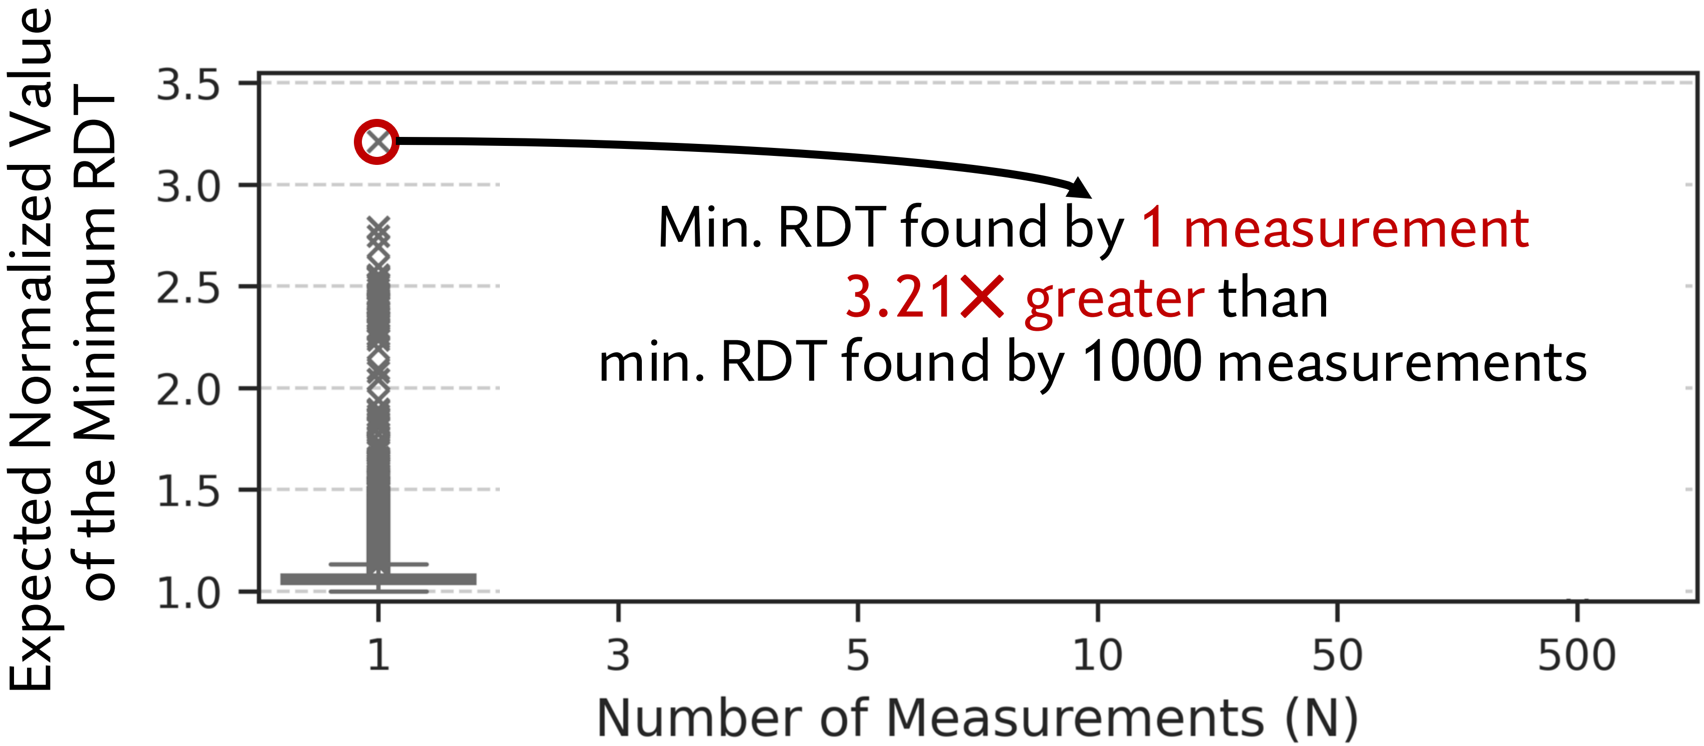

Expected Value of the Minimum RDT

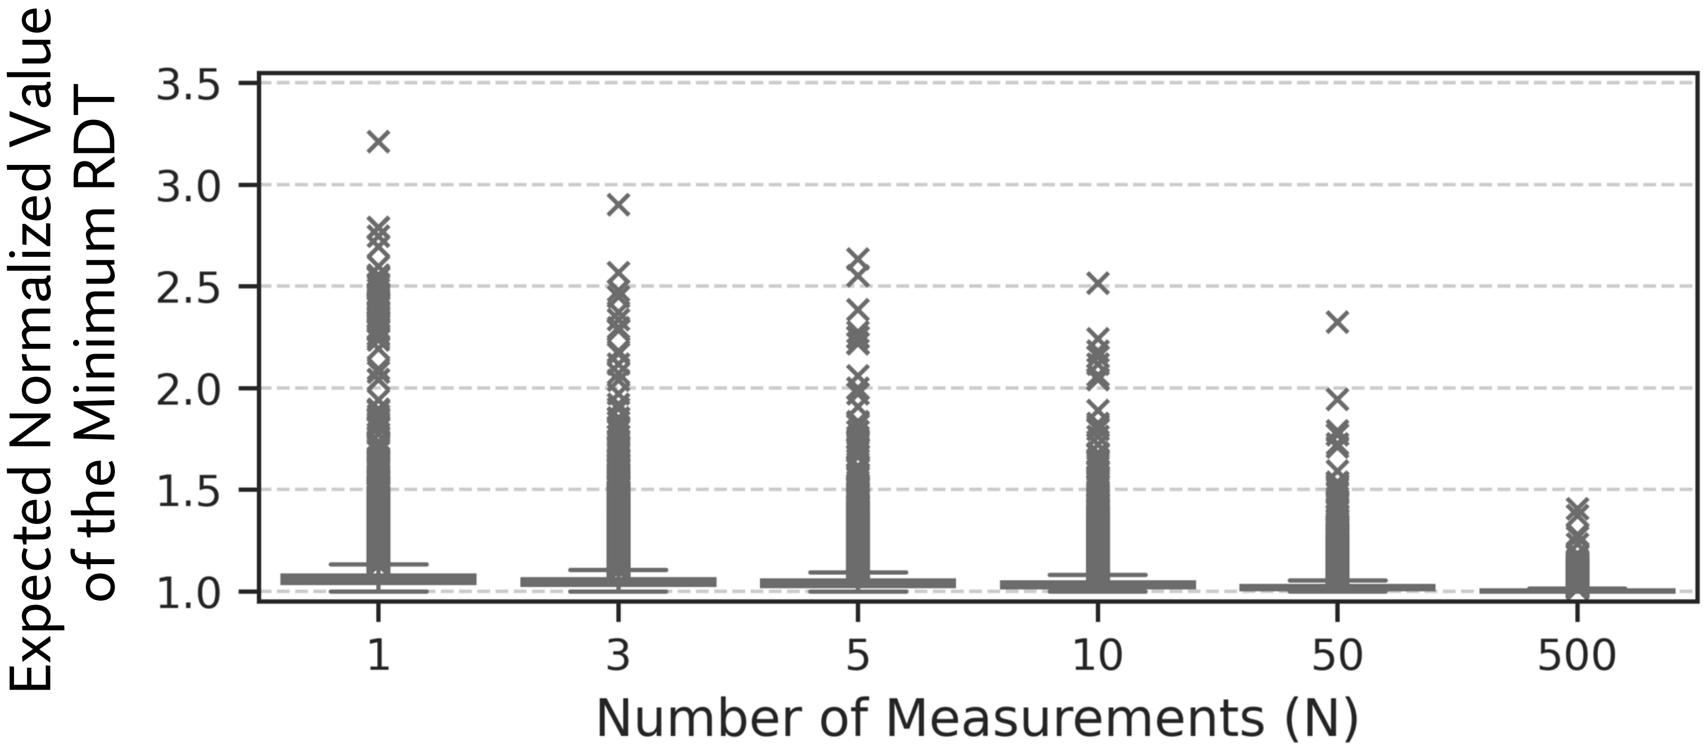

Hypothetical Explanation for VRD

- No device-level study shows temporal variations in read disturbance vulnerability

Effect of Die Density and Die Revision (I)

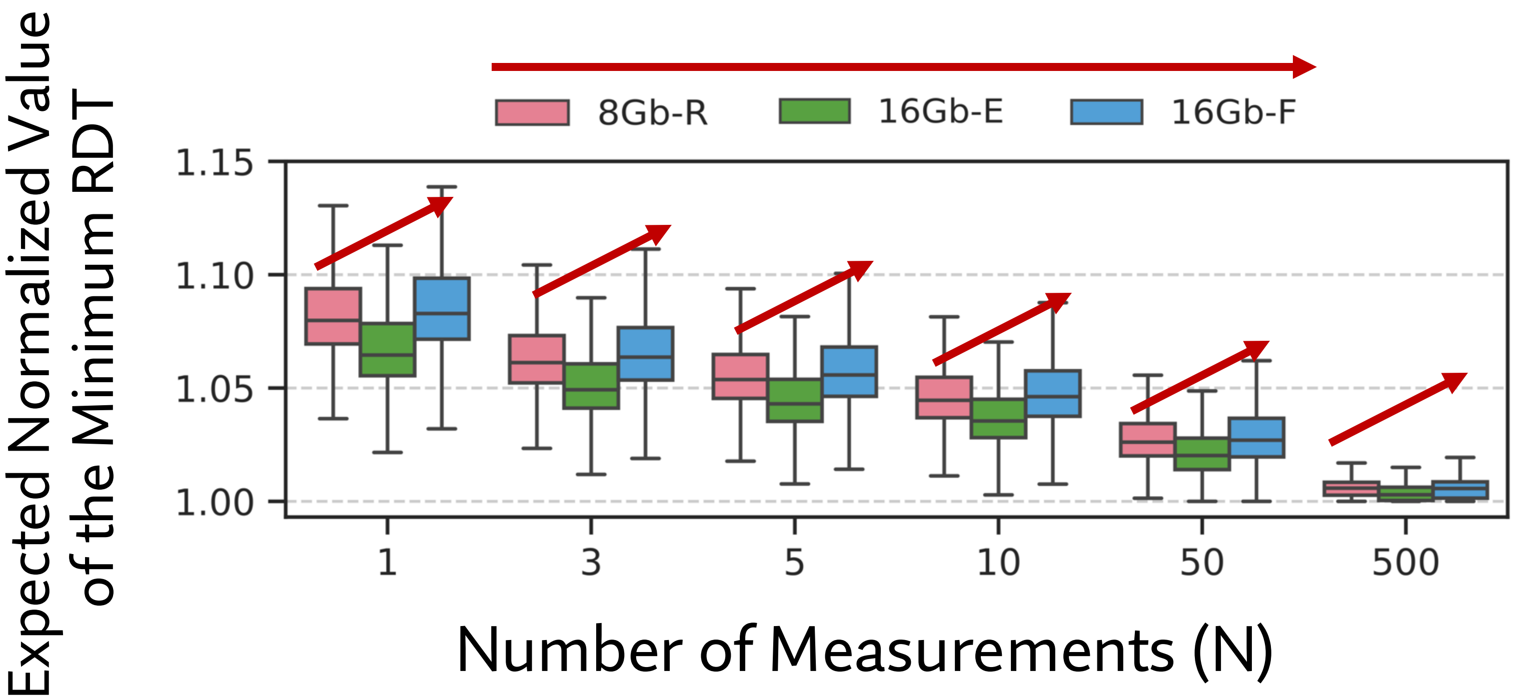

RDT distribution worsens with increasing die density and with advanced DRAM technology

Effect of Die Density and Die Revision (II)

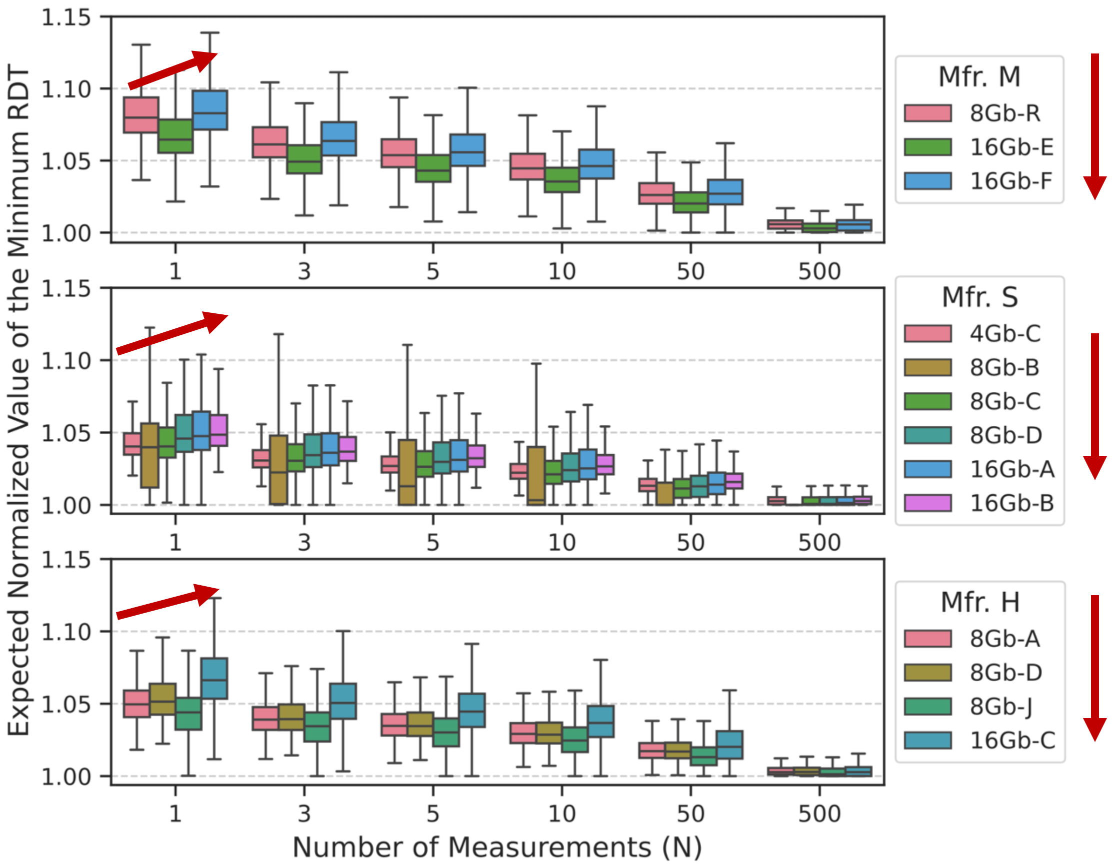

The effect of die density and die revision is consistent across all tested modules

Effect of Data Pattern (I)

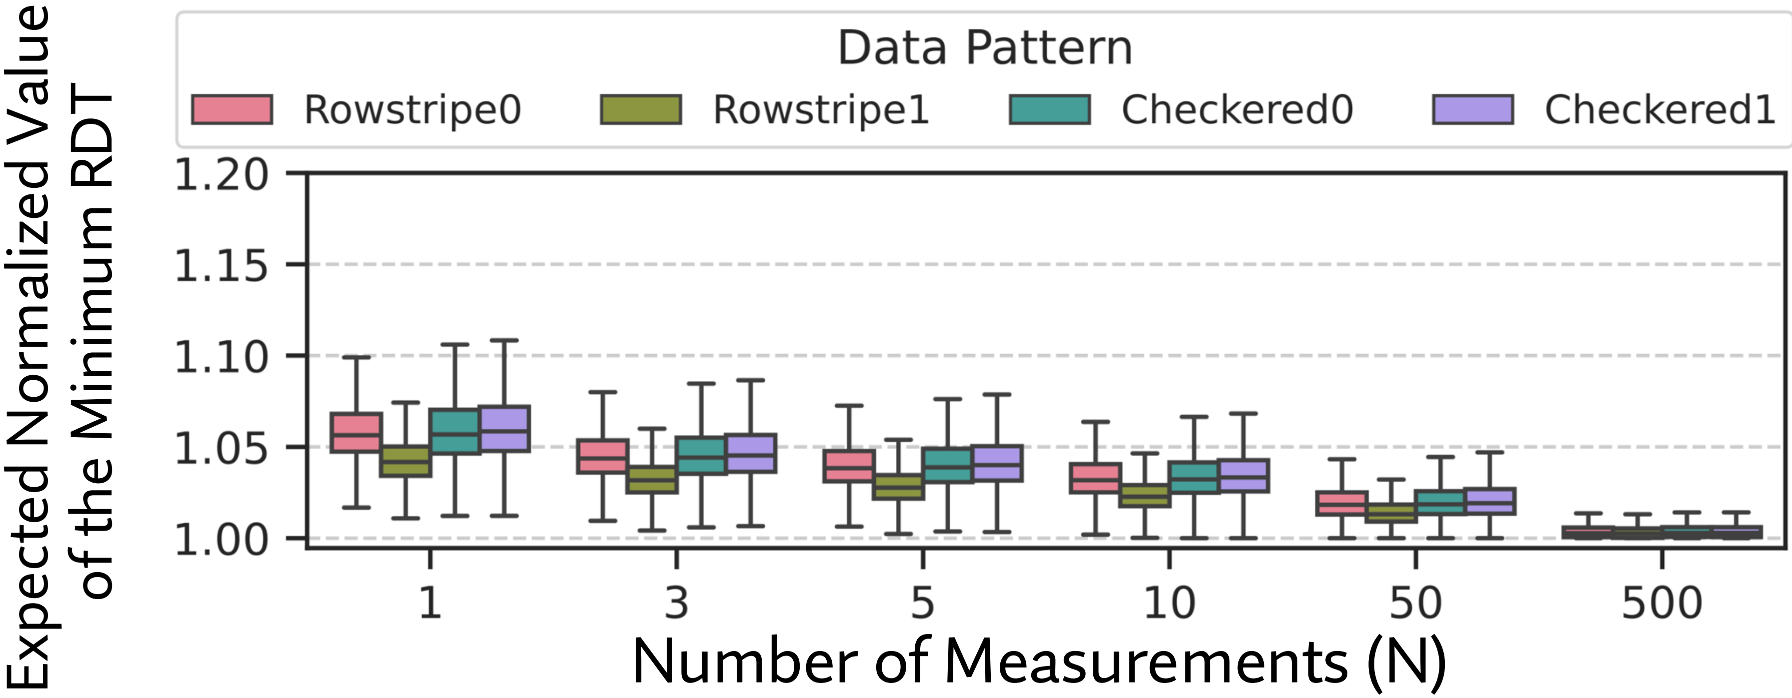

Effect of Data Pattern (I)

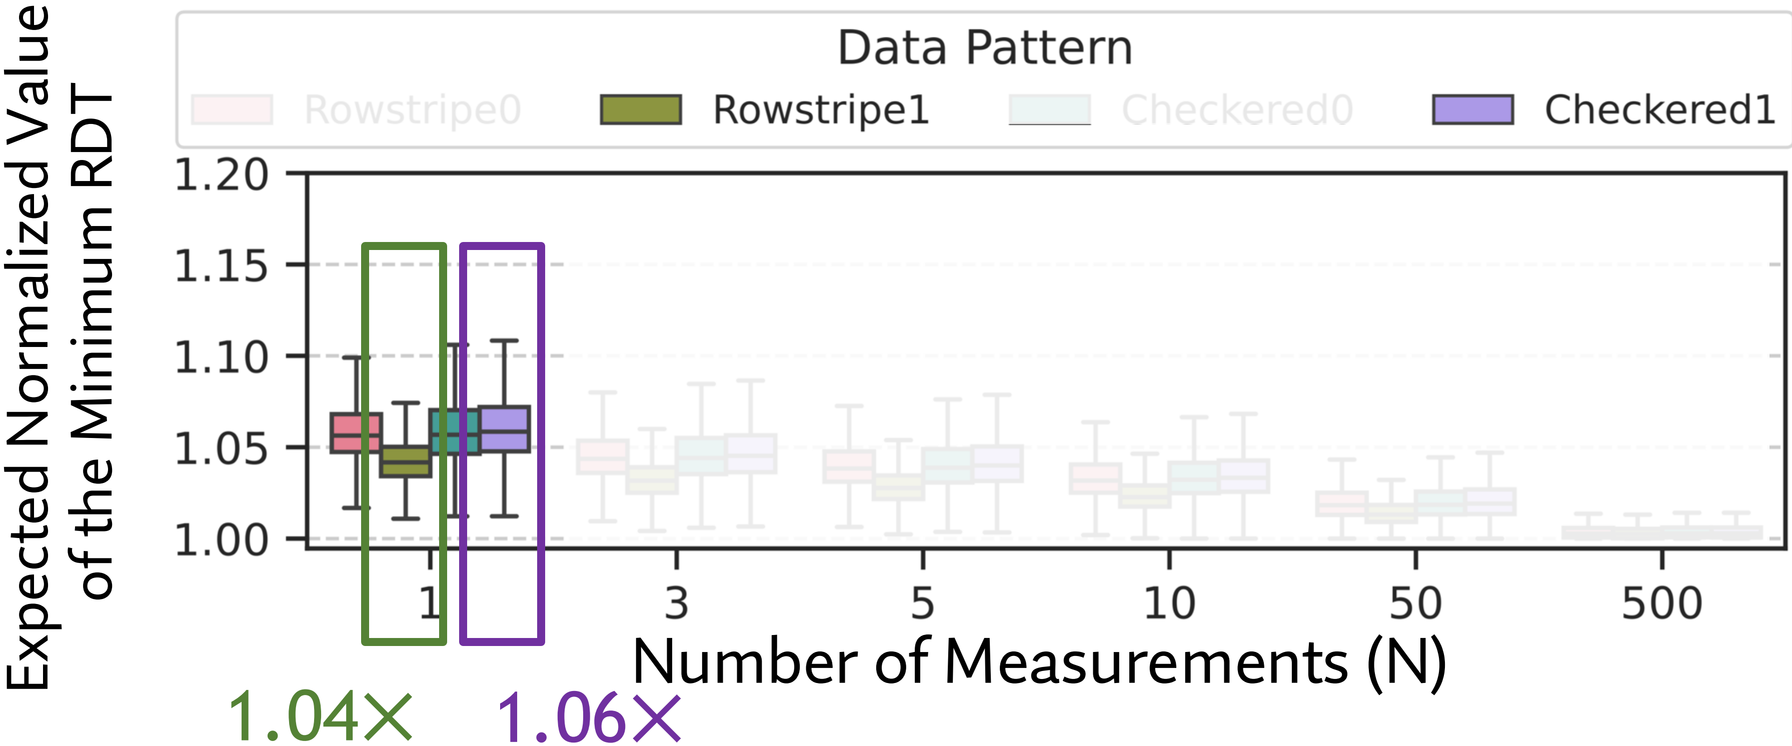

expected normalized value for the median row

RDT distribution changes with data pattern

Effect of Data Pattern (II)

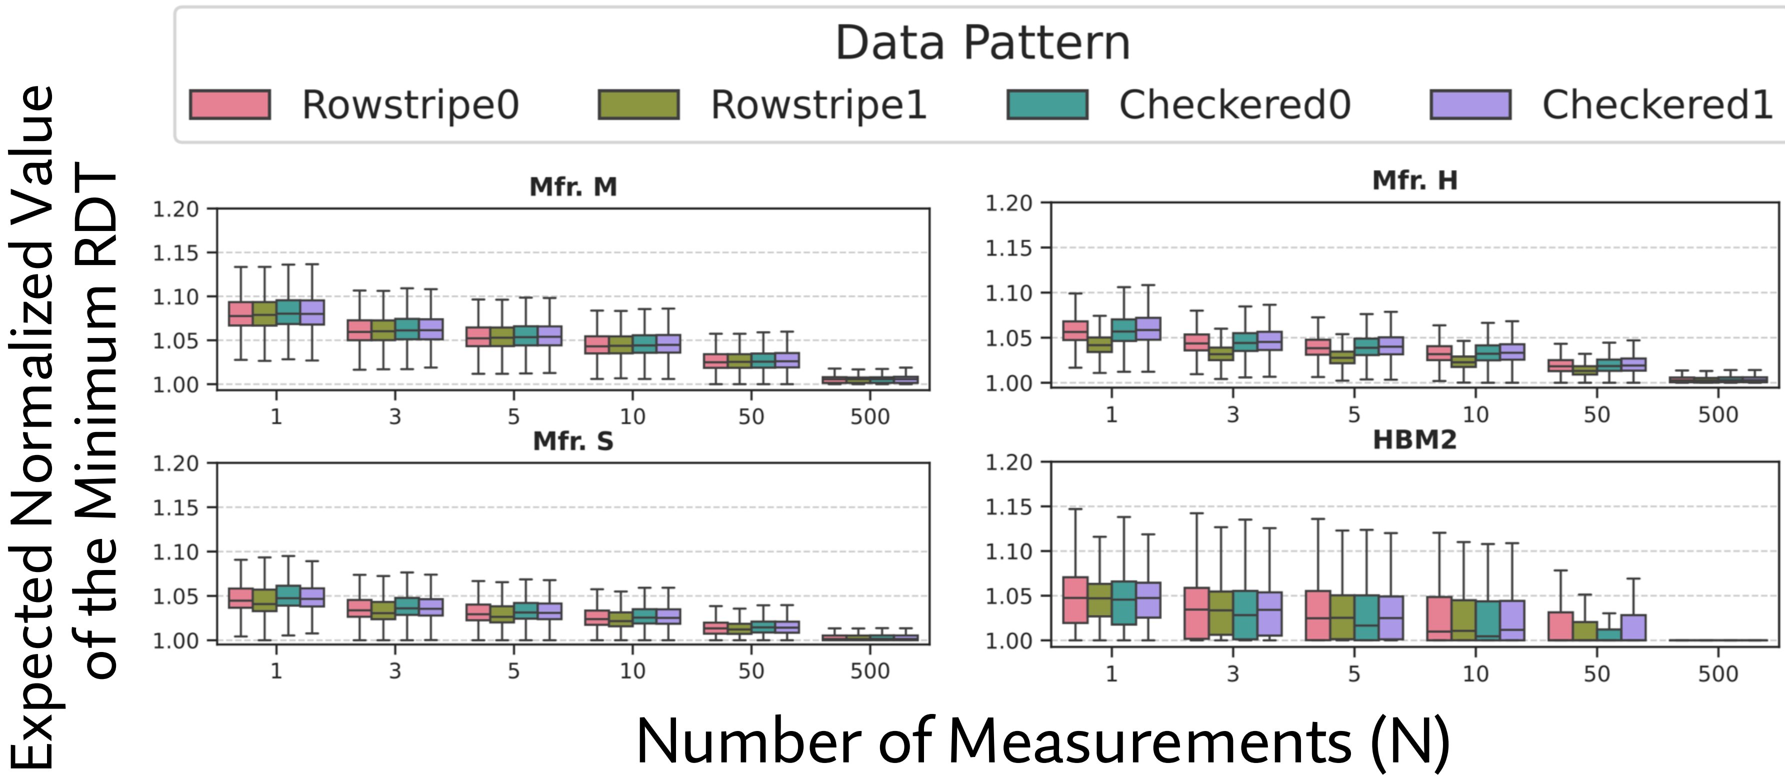

No single data pattern causes the worst RDT distribution across all tested DRAM chips

Effect of Aggressor Row On Time

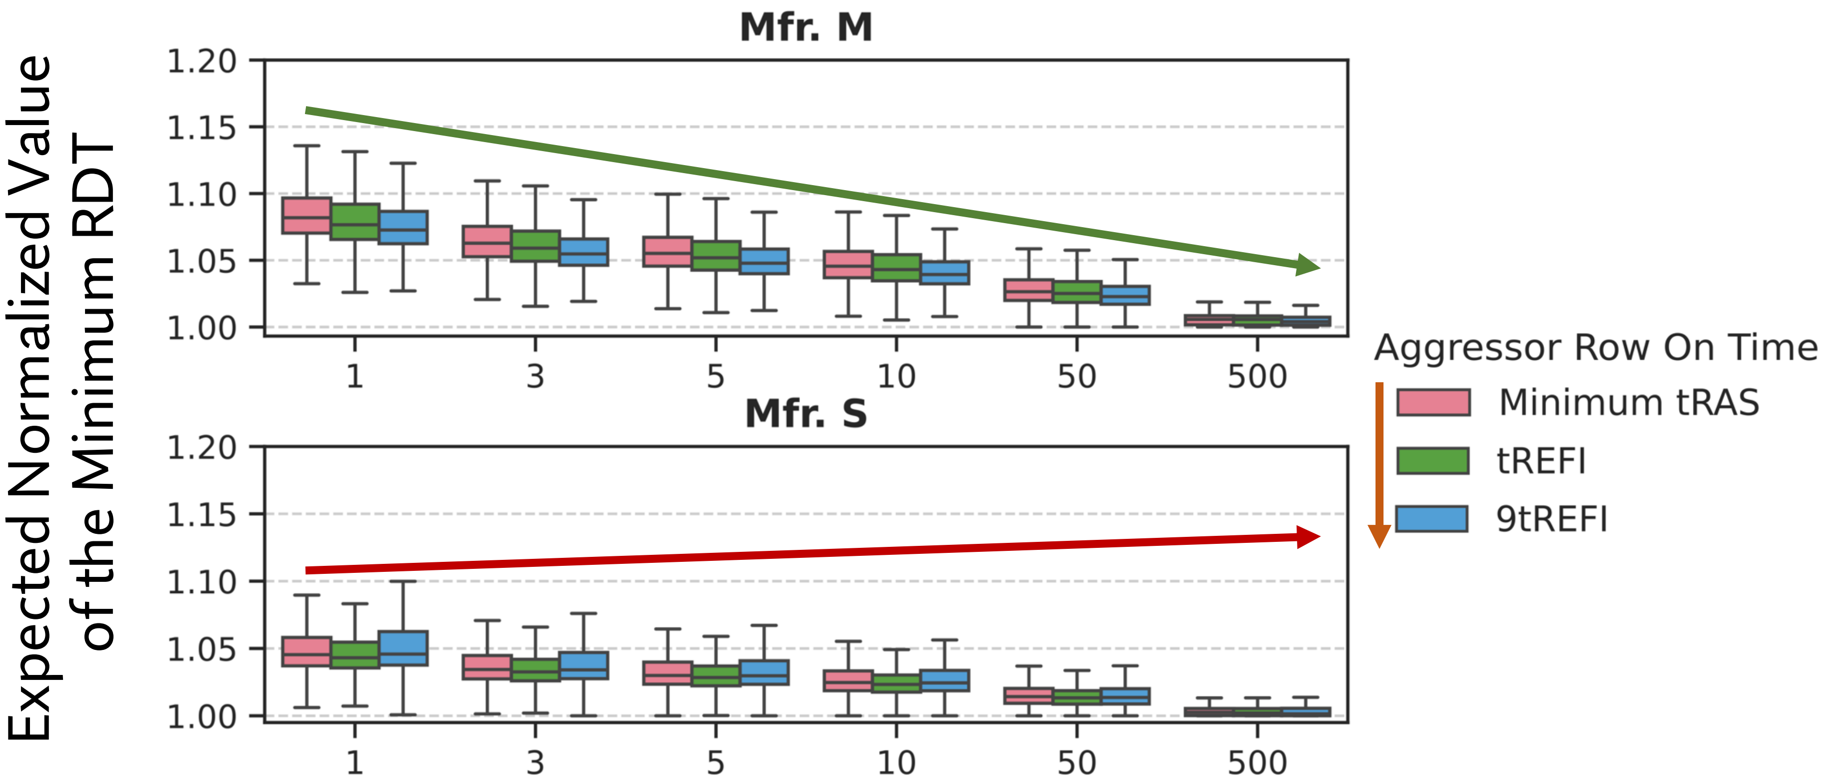

RDT distribution changes with aggressor row on time

RDT distribution can become better or worse with increasing row on time

Effect of Temperature

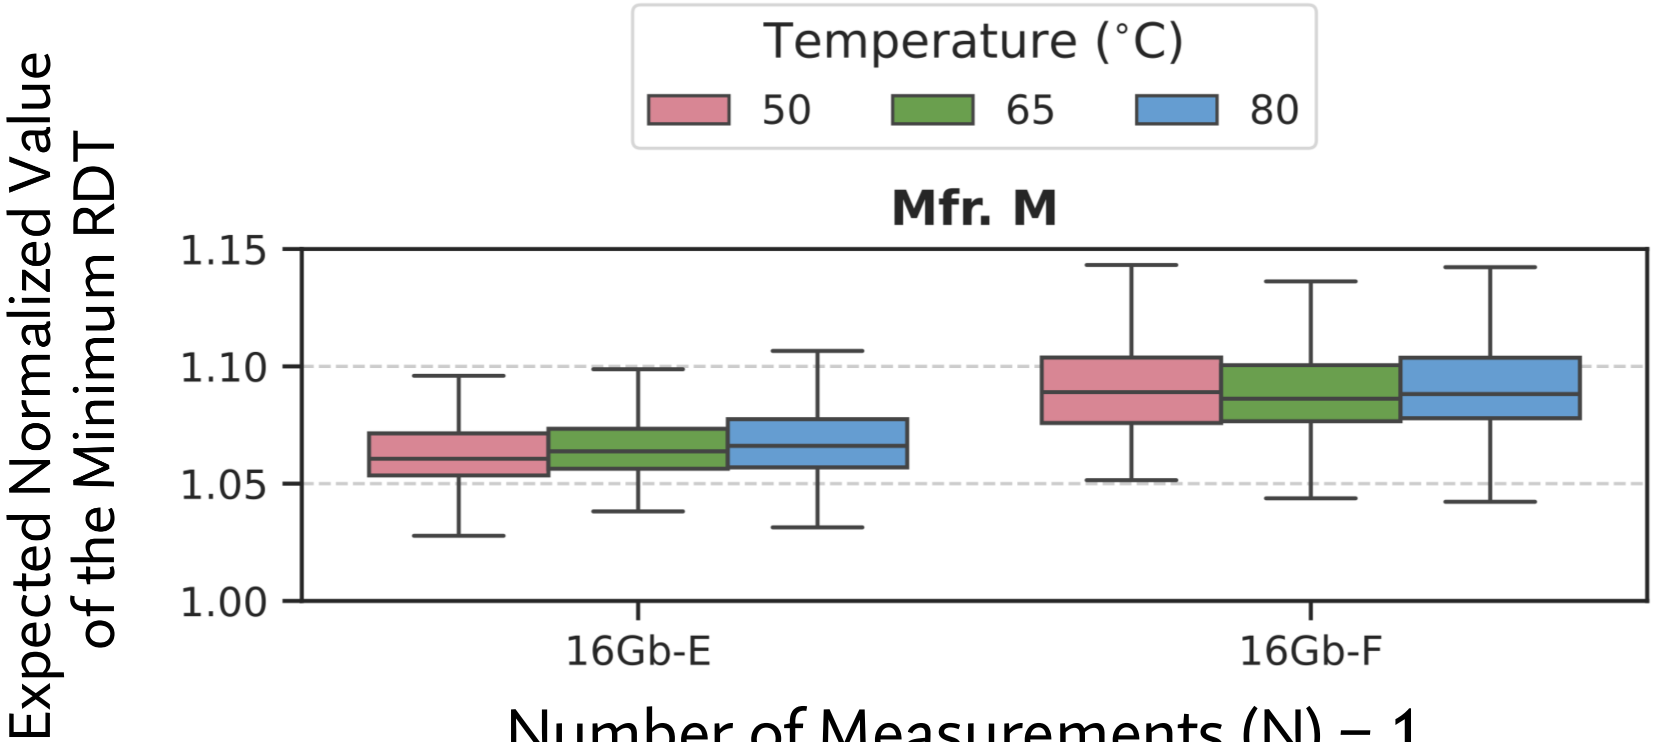

Number of Measurements (N) = 1

RDT distribution tends to change with temperature

DRAM Operation: Activate and Precharge

Access data in Row 1

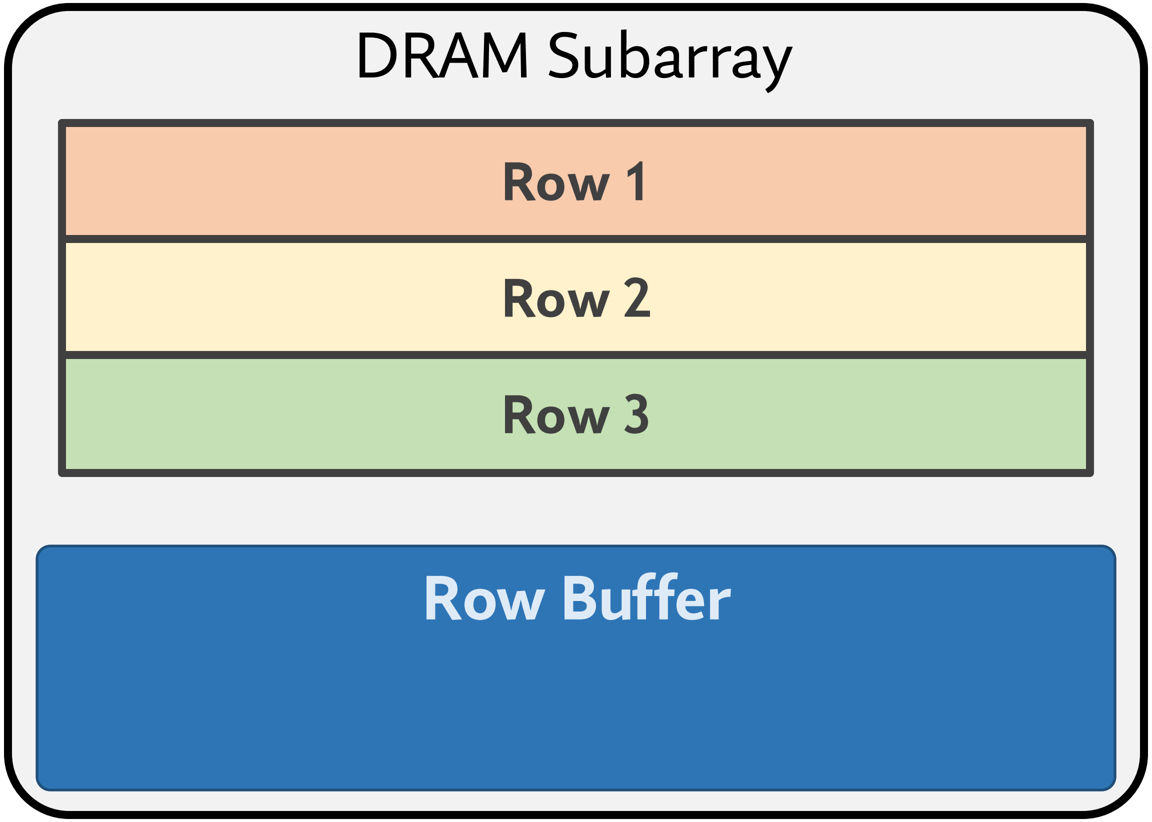

Row 1 is **closed**

DRAM Operation: Activate and Precharge

Access data in Row 1

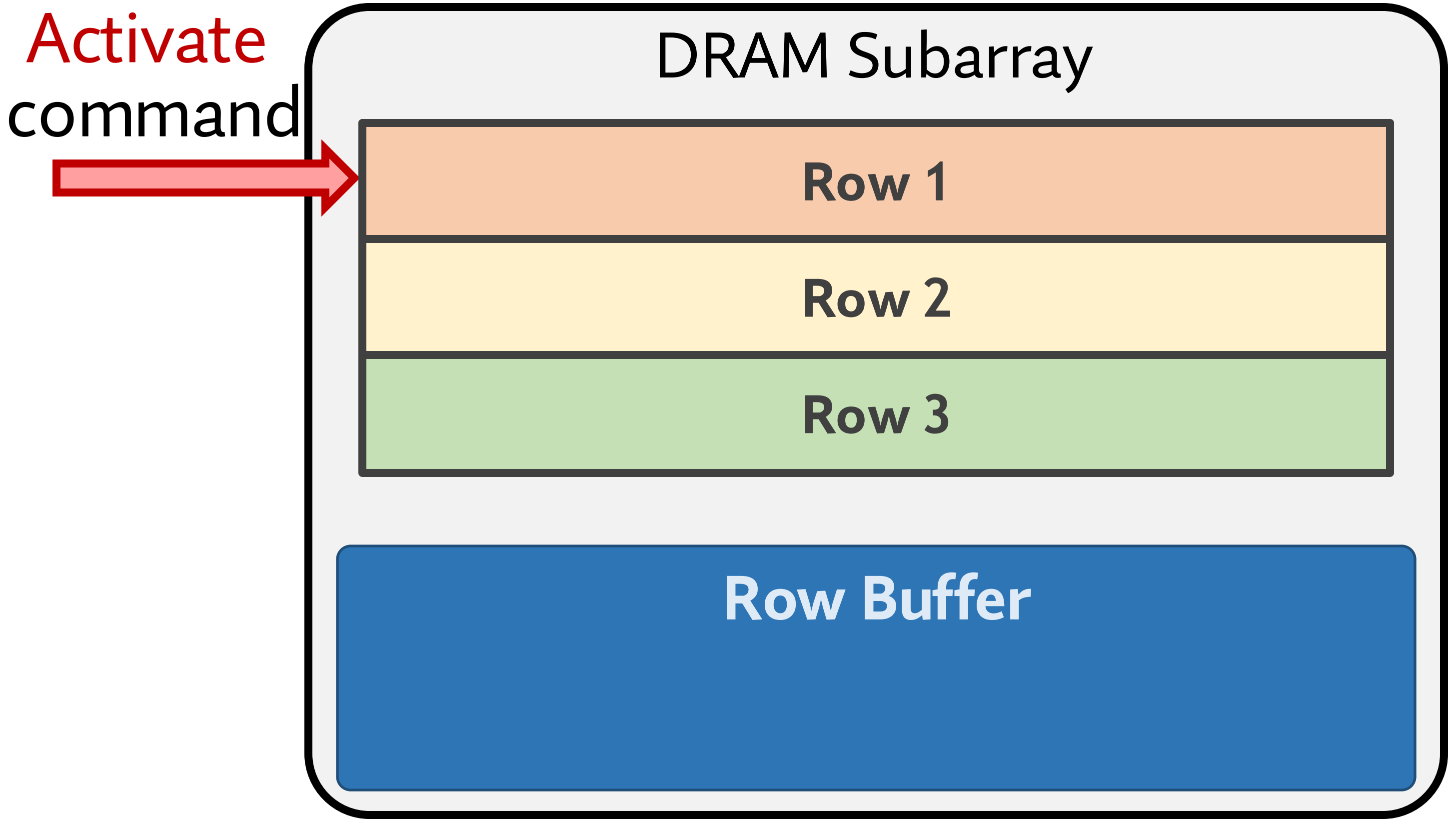

Row 1 is opened

Row 1 is closed

DRAM Operation: Activate and Precharge

Access data in Row 3

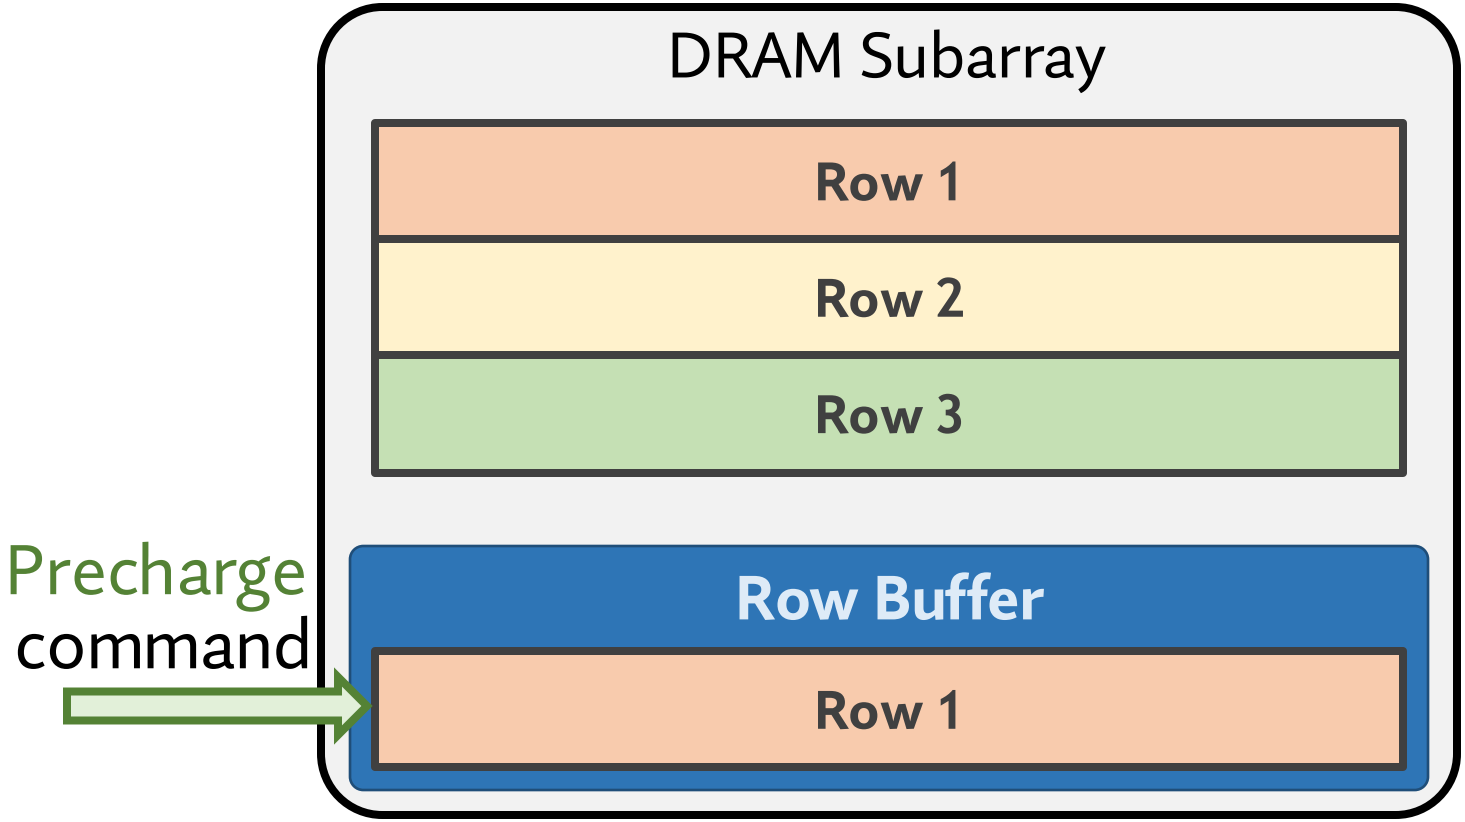

Row 3 is **closed**

DRAM Operation: Activate and Precharge

Access data in Row 3

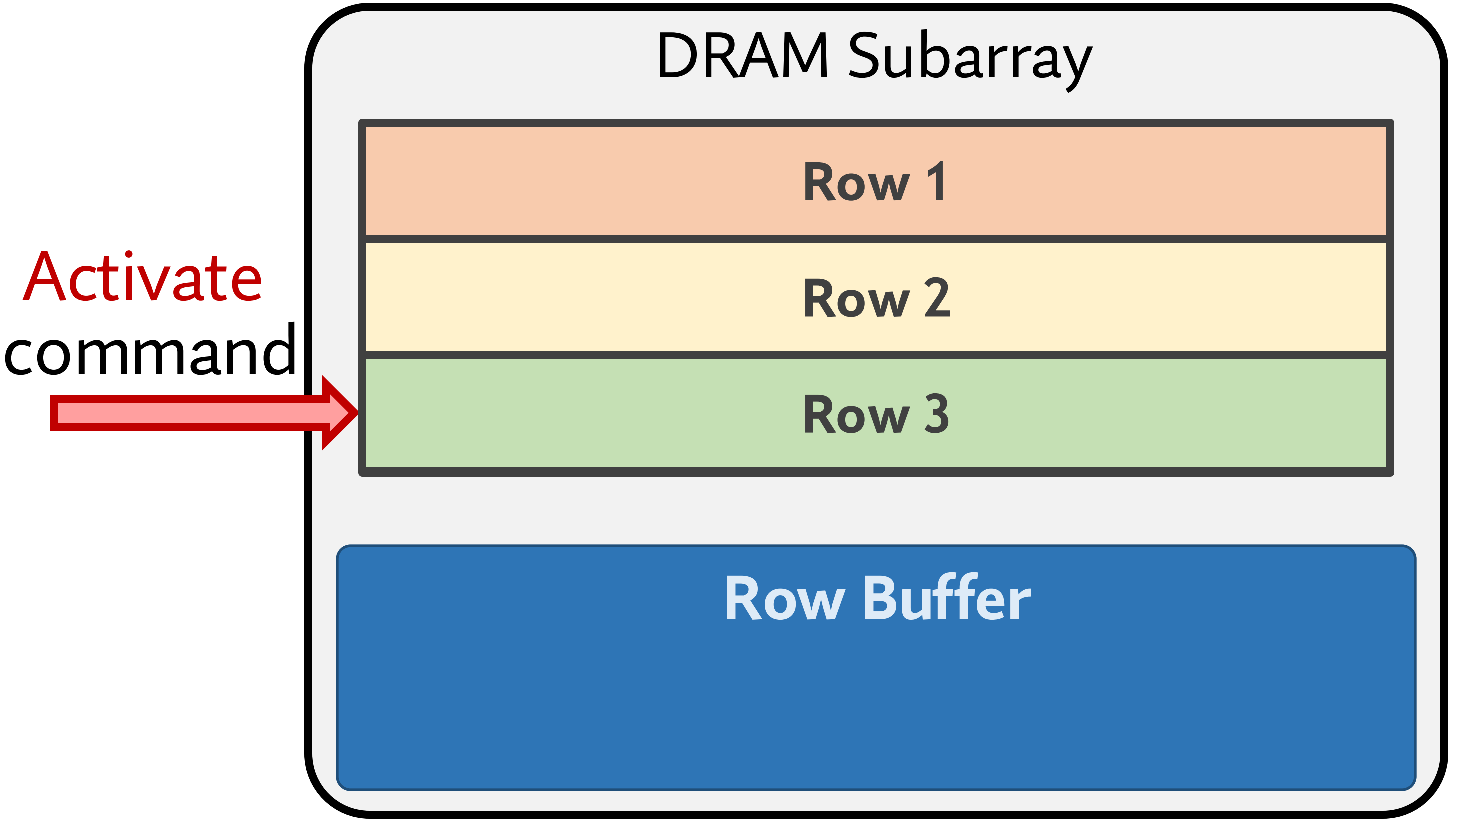

Row 3 is closed
Row 3 is opened

Implications Summary

- Security guarantees provided by mitigation techniques rely on accurately identified minimum RDT

- Accurate identification of minimum RDT is extremely challenging (even with 1000s measurements) because RDT unpredictably changes over time

- **Given our limited bitflip dataset** a ≥20% guardband for RDT combined with error-correcting codes (e.g., Chipkill) could prevent VRD-induced bitflips at performance cost

- Evaluate a short-term solution: combining a guardband for RDT and error-correcting codes
  - >10% guardband for the minimum observed RDT, combined with;
  - single-error correcting double-error detecting (SECDED) or Chipkill-like ECC
  - could prevent VRD-induced bitflips at performance cost

- Call for future work on online RDT profiling and runtime configurable read disturbance mitigations

Variable Read Disturbance (VRD) Summary

Research Question

- How accurately and efficiently can we measure the RDT of each DRAM row?

Experimental Characterization

- Record >100M RDT measurements across 3750 rows and many test parameters (e.g., temperature, data pattern) in 160 DDR4 and 4 HBM2 chips

Key Observations

- RDT changes significantly and unpredictably over time: VRD
  - Maximum observed RDT for a tested row can be 3.5X higher than minimum
  - Smallest RDT value (for a row) may appear after 94,467 measurements

Implications for System Security and Robustness

- RDT cannot be accurately identified quickly
- RDT guardbands (>10%) and ECC (SECDED or Chipkill) could prevent VRD-induced bitflips at significant performance cost

DRAM Cell Leakage

Each cell encodes information in **leaky** capacitors

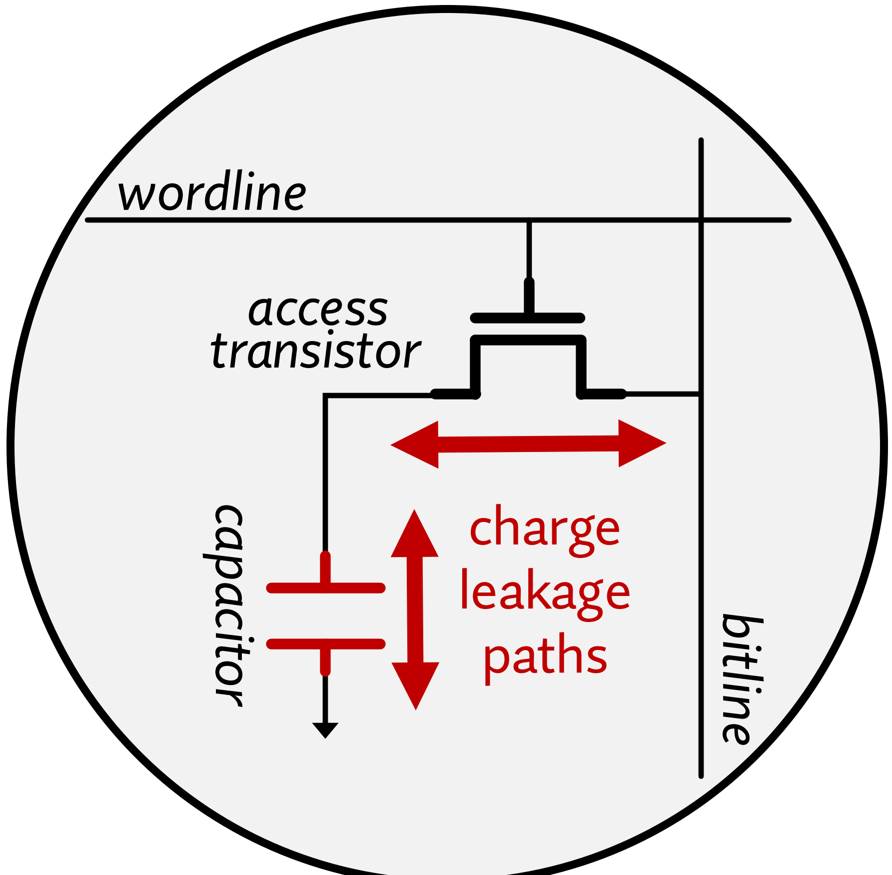

Stored data is **corrupted** if too much charge leaks (i.e., the capacitor voltage degrades too much)

[Patel+, ISCA'17]

DRAM Refresh

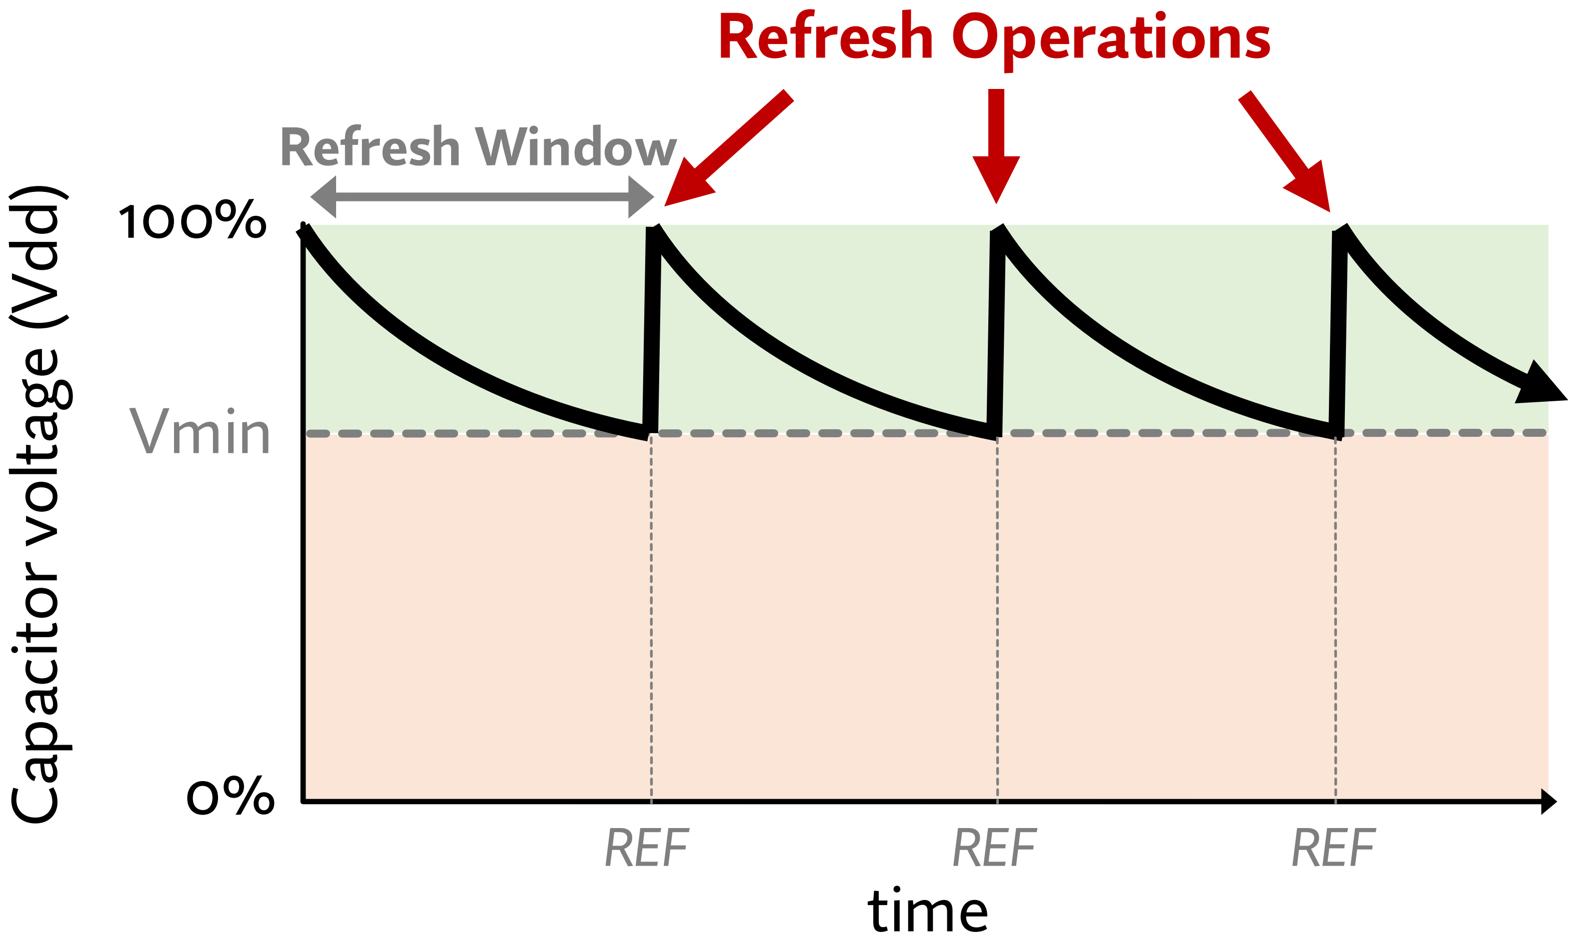

Periodic **refresh operations** preserve stored data

Read Disturbance Bitflips

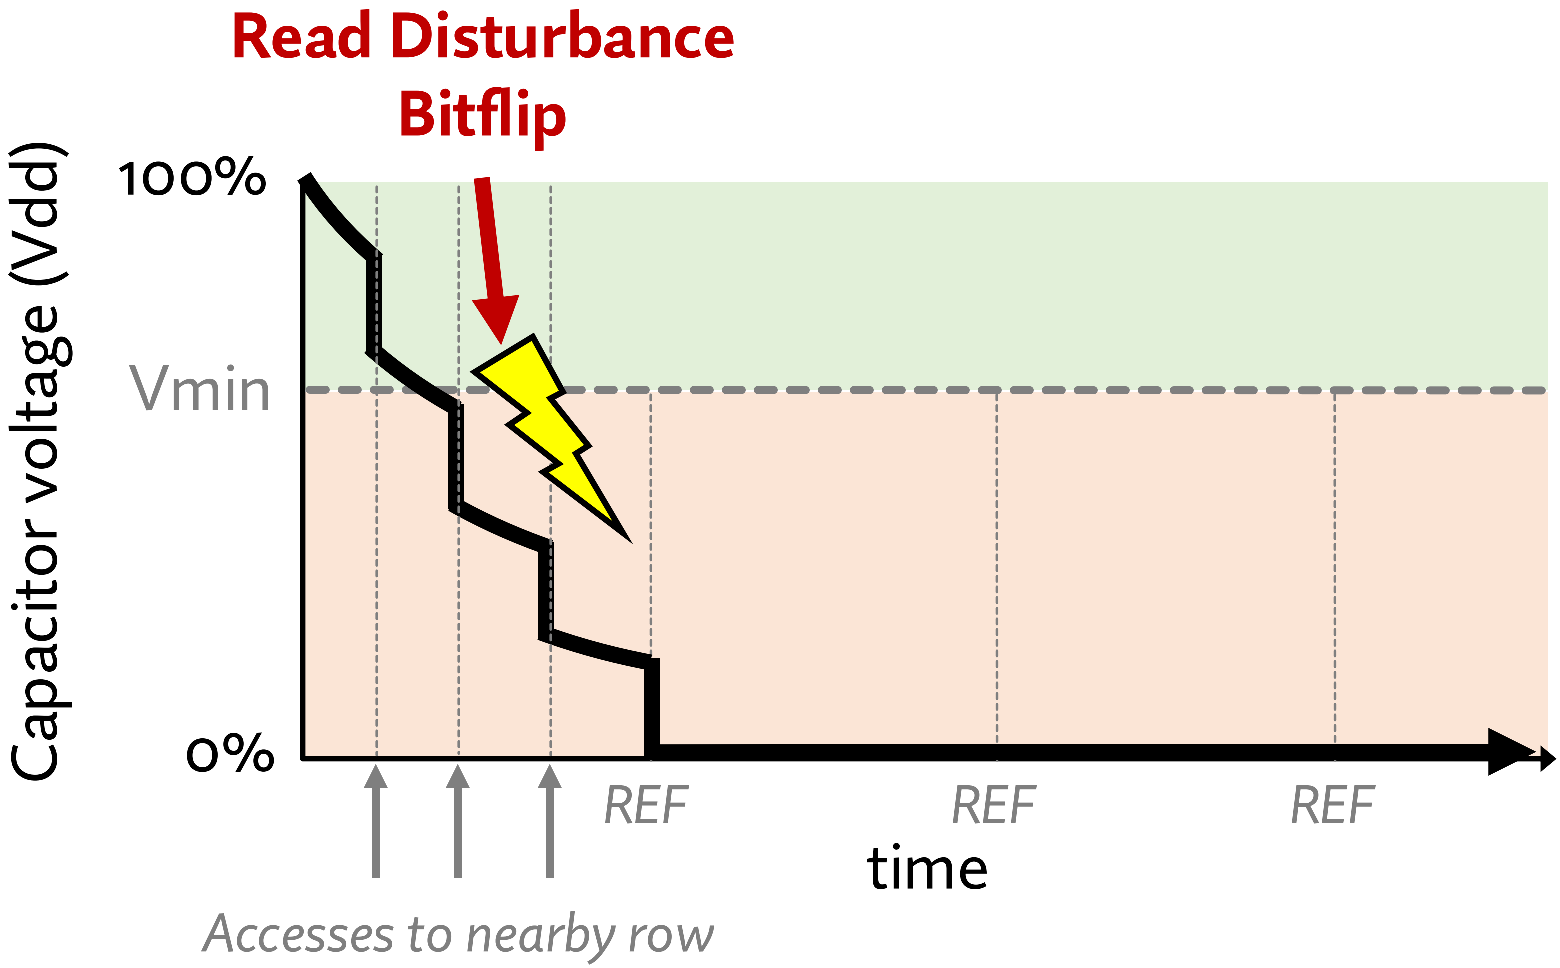

[Kim+, ISCA'20]