Optimizing DRAM Timing for the Common-Case

Adaptive-Latency DRAM

Donghyuk Lee

Yoongu Kim, Gennady Pekhimenko, Samira Khan,

Vivek Seshadri, Kevin Chang, Onur Mutlu

SAFARI

Carnegie Mellon

x86 CPU

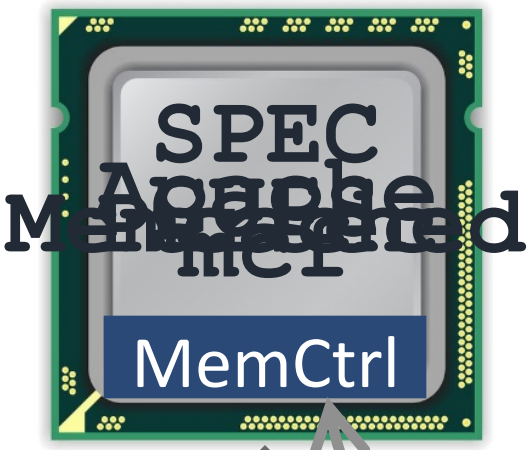

Runtime: 527min

Runtime: 477min

-10.5% (no error)

Timing Parameters

(11 – 11 – 28)

(8 – 8 – 19)

DRAM Module

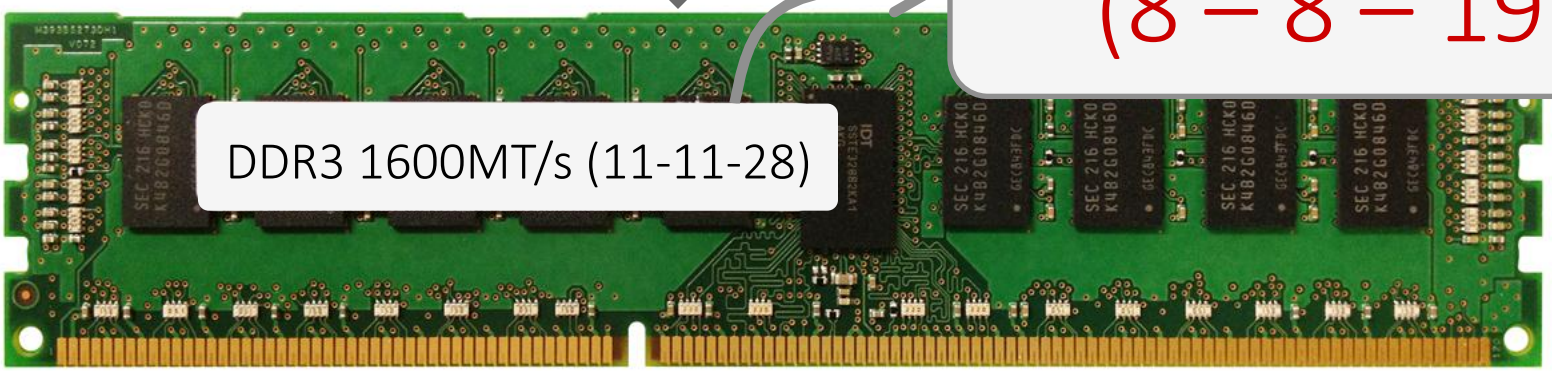

*Why can we **reduce DRAM timing parameters without any errors**?*

Executive Summary

- *Observations*
  - *DRAM timing parameters are dictated by the worst-case cell (smallest cell across all products at highest temperature)*
  - *DRAM operates at lower temperature than the worst case*
- *Idea: Adaptive-Latency DRAM*
  - *Optimizes DRAM timing parameters for the common case (typical DIMM operating at low temperatures)*
- *Analysis: Characterization of 115 DIMMs*
  - *Great potential to lower DRAM timing parameters (17 – 54%) without any errors*
- *Real System Performance Evaluation*
  - *Significant performance improvement (14% for memory-intensive workloads) without errors (33 days)*

1. DRAM Operation Basics

2. Reasons for Timing Margin in DRAM

3. Key Observations

4. Adaptive-Latency DRAM

5. DRAM Characterization

6. Real System Performance Evaluation

DRAM Stores Data as Charge

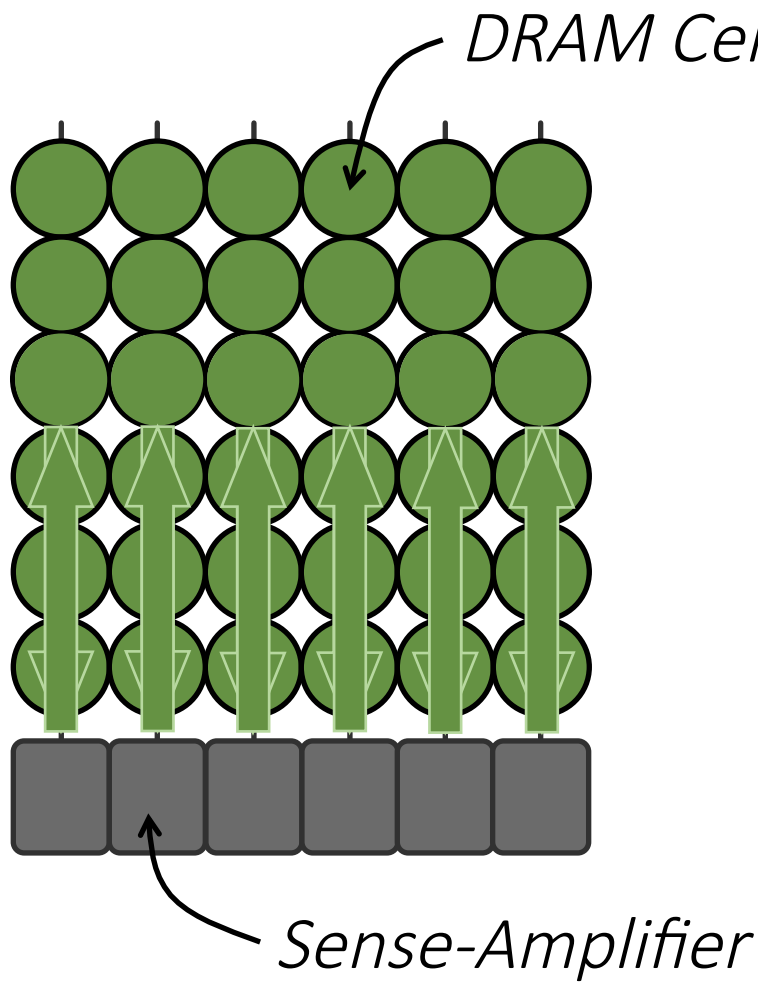

Three steps of charge movement

1. Sensing
2. Restore
3. Precharge

DRAM Charge over Time

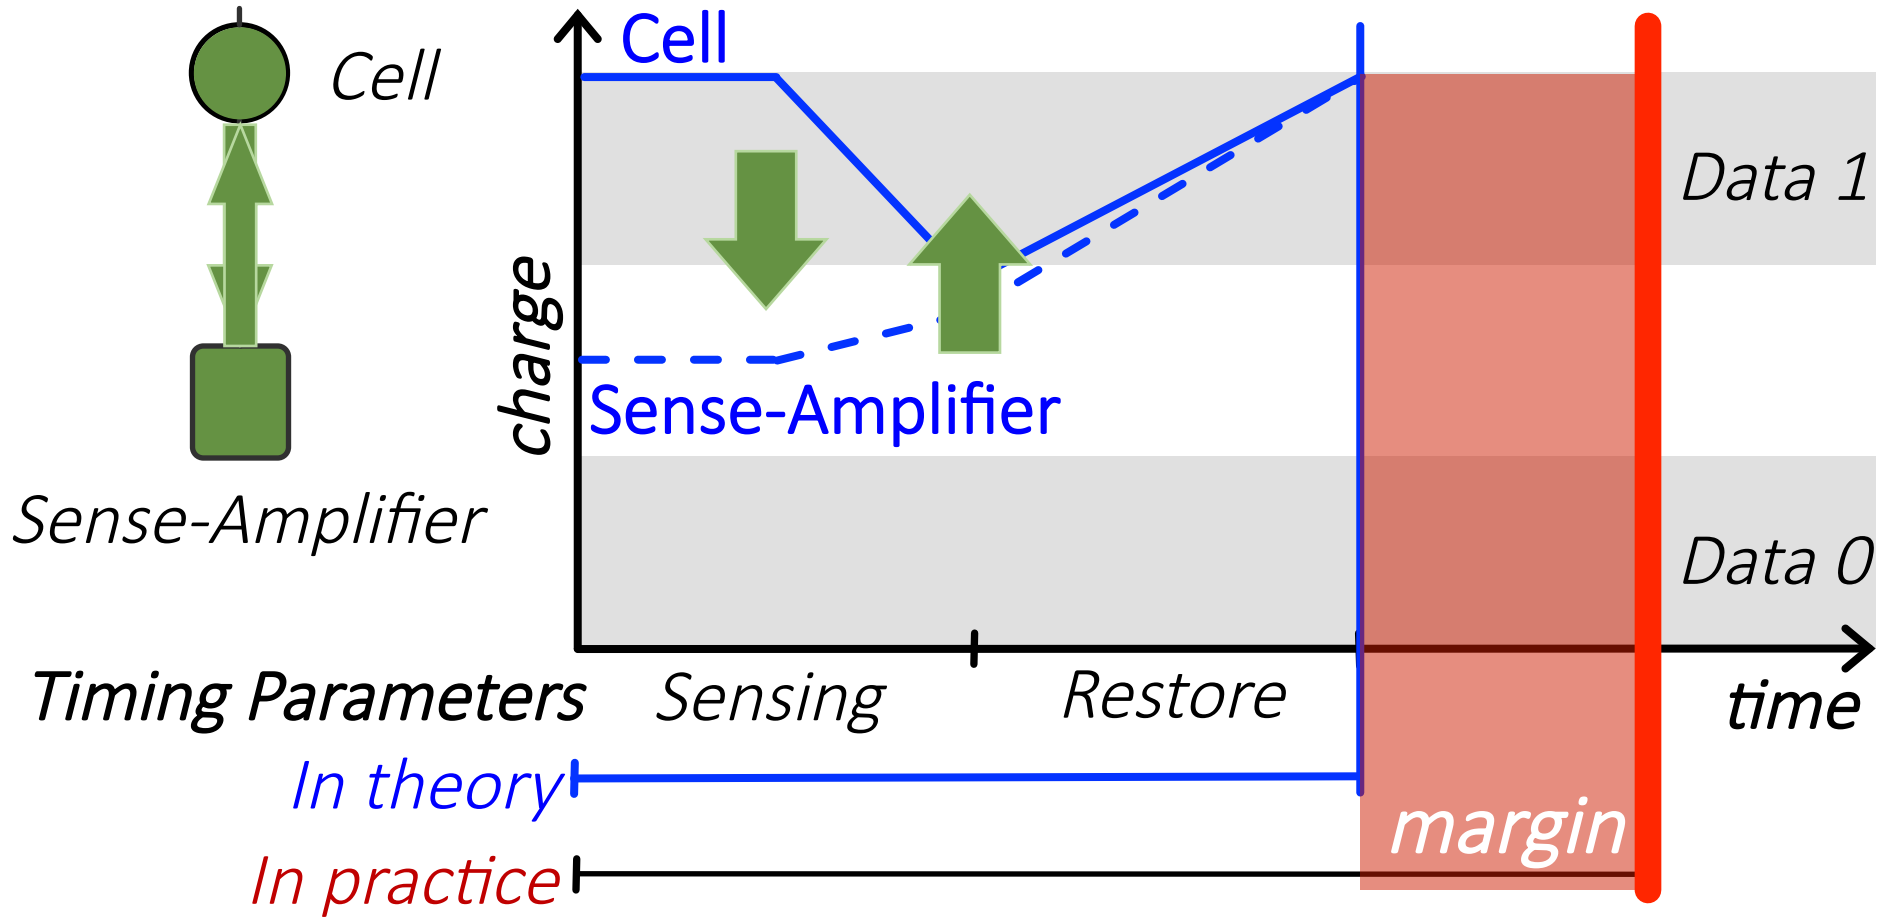

Why does DRAM need the extra timing margin?

1. DRAM Operation Basics

2. Reasons for Timing Margin in DRAM

3. Key Observations

4. Adaptive-Latency DRAM

5. DRAM Characterization

6. Real System Performance Evaluation

Two Reasons for Timing Margin

1. Process Variation

- DRAM cells are not equal
- Leads to extra timing margin for a cell that can store a large amount of charge

2. Temperature Dependence

DRAM Cells are Not Equal

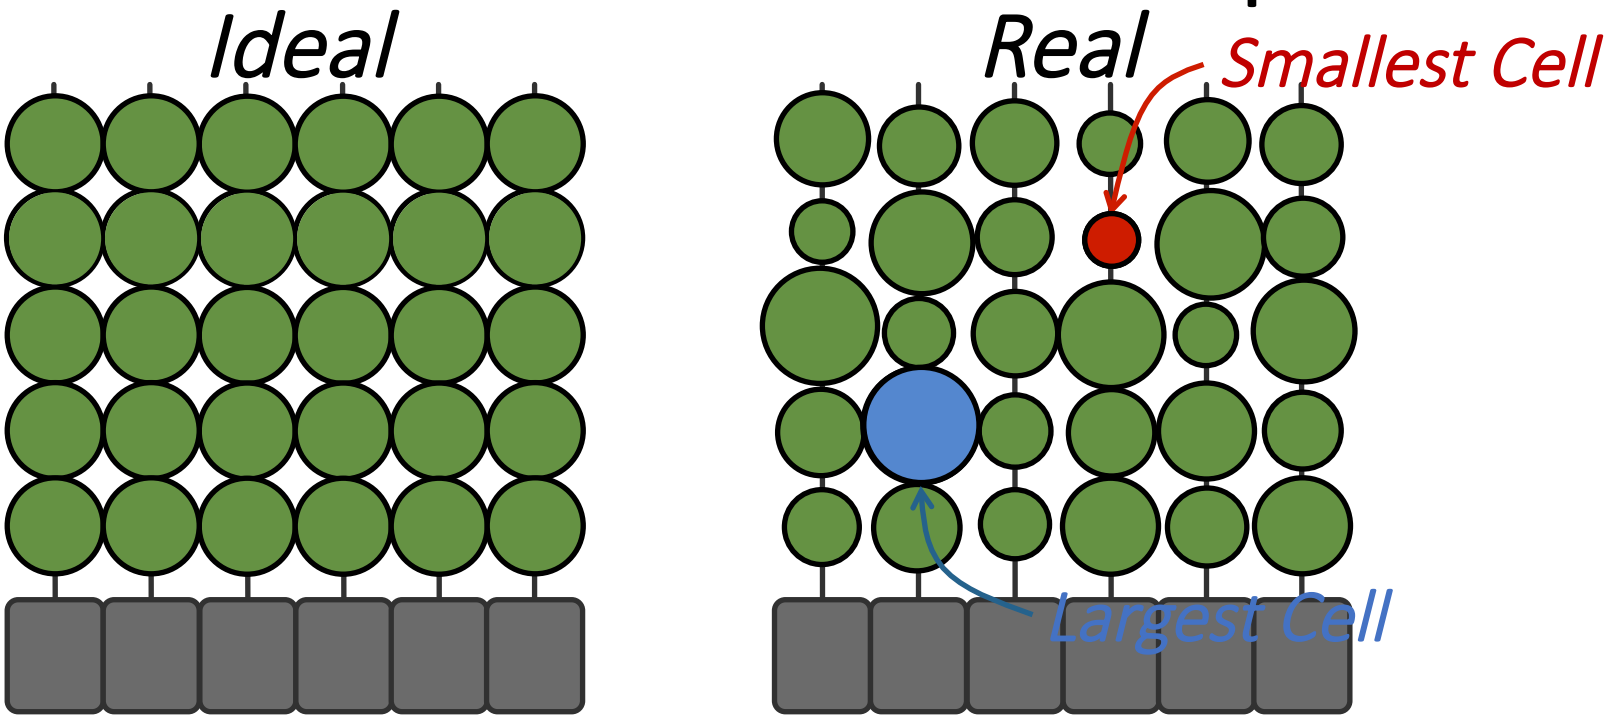

Large variation in cell size

Large variation in charge

Large variation in access latency

Process Variation

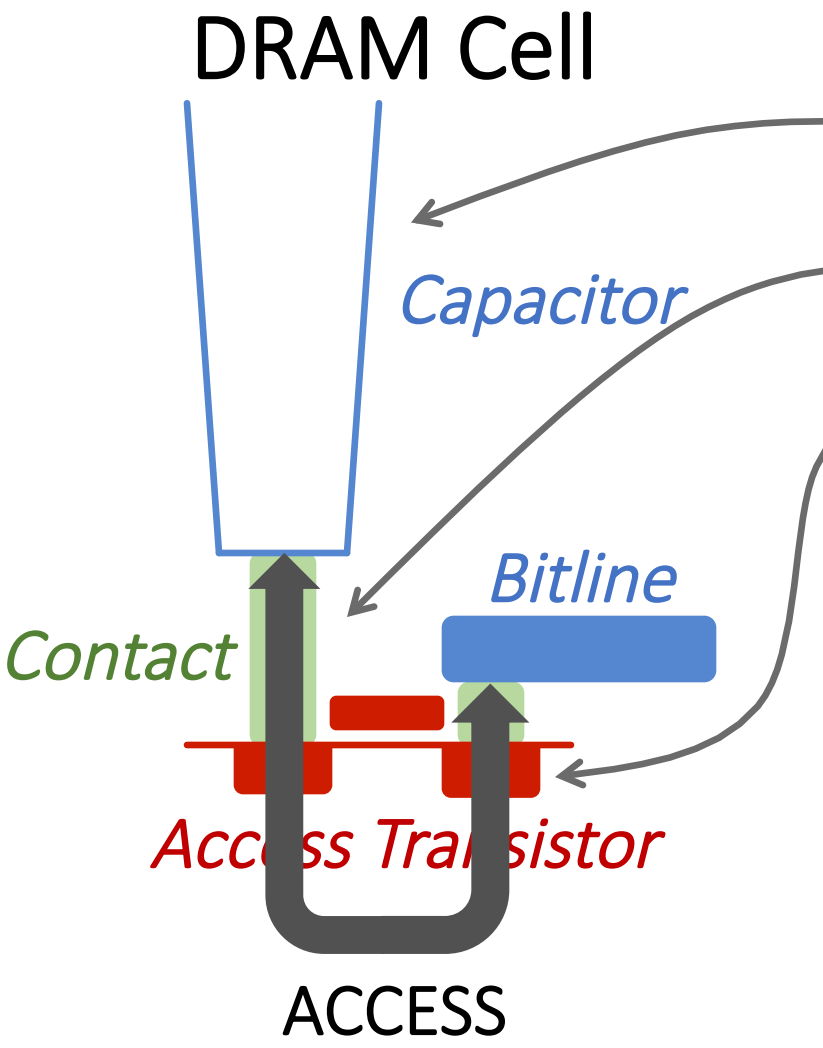

❶ Cell Capacitance

❷ Contact Resistance

❸ Transistor Performance

Small cell can store small charge

- *Small cell capacitance*
- *High contact resistance*
- *Slow access transistor*

➔ *High access latency*

Two Reasons for Timing Margin

1. Process Variation

- DRAM cells are not equal
- Leads to extra timing margin for a cell that can store a large amount of charge

2. Temperature Dependence

- DRAM leaks more charge at higher temperature
- Leads to extra timing margin for cells that operate at the low temperature

Charge Leakage ∝ Temperature

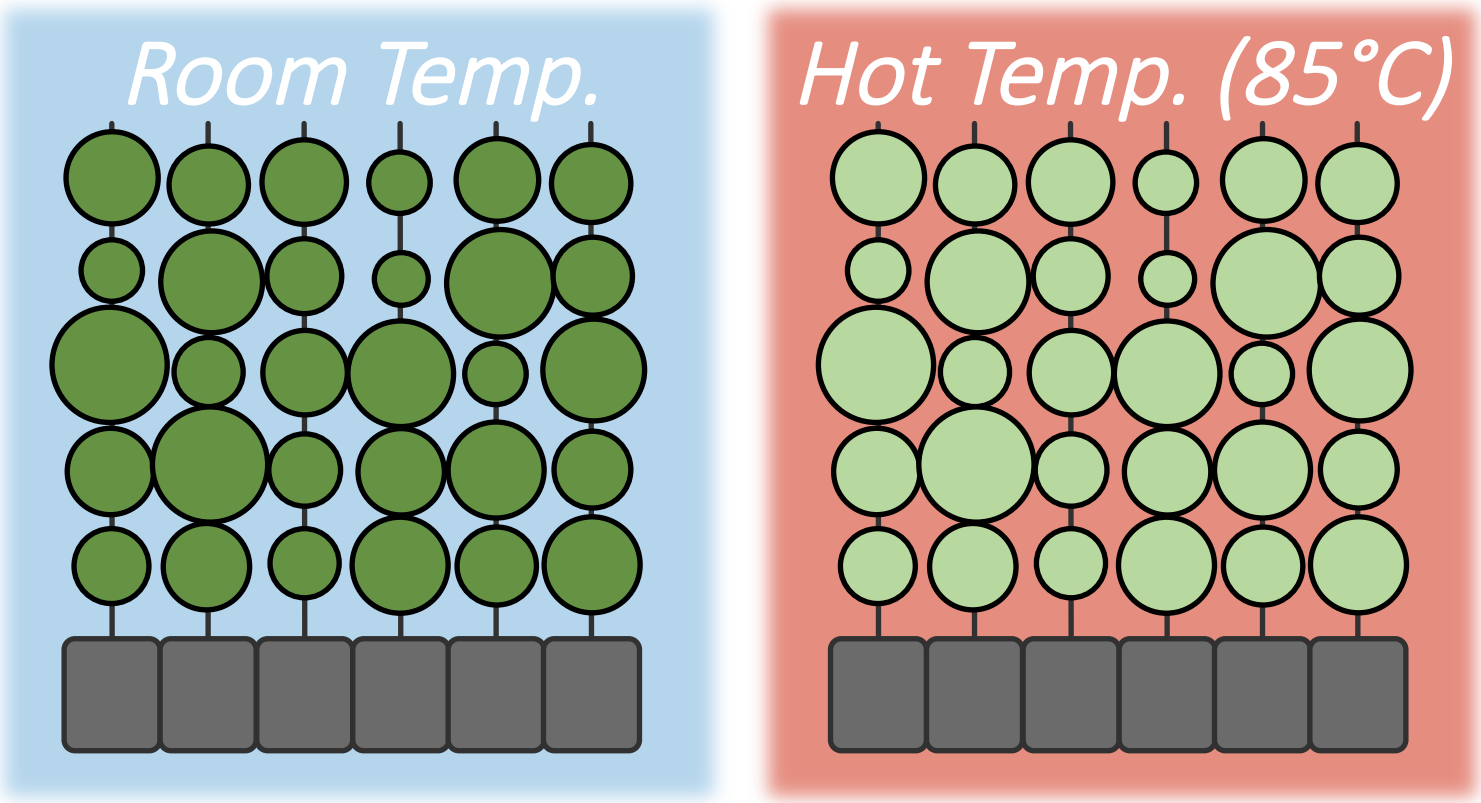

Cells store small charge at high temperature and large charge at low temperature

→ Large variation in access latency

DRAM Timing Parameters

- *DRAM timing parameters are dictated by **the worst-case***
  - The smallest cell with the smallest charge **in all DRAM products**
  - Operating at **the highest temperature**

- *Large timing margin for the common-case*

Our Approach

- *We optimize DRAM timing parameters for the common-case*
  - The smallest cell with the smallest charge in a DRAM module
  - Operating at the current temperature

- *Common-case cell has extra charge than the worst-case cell*

→ *Can lower latency for the common-case*

1. DRAM Operation Basics

2. Reasons for Timing Margin in DRAM

3. Key Observations

4. Adaptive-Latency DRAM

5. DRAM Characterization

6. Real System Performance Evaluation

Key Observations

1. Sensing

   Sense **cells with extra charge** faster

   → Lower sensing latency

2. Restore

   No need to fully restore **cells with extra charge**

   → Lower restore latency

3. Precharge

   No need to fully precharge bitlines for **cells with extra charge**

   → Lower precharge latency

Observation 1. Faster Sensing

Typical DIMM at Low Temperature

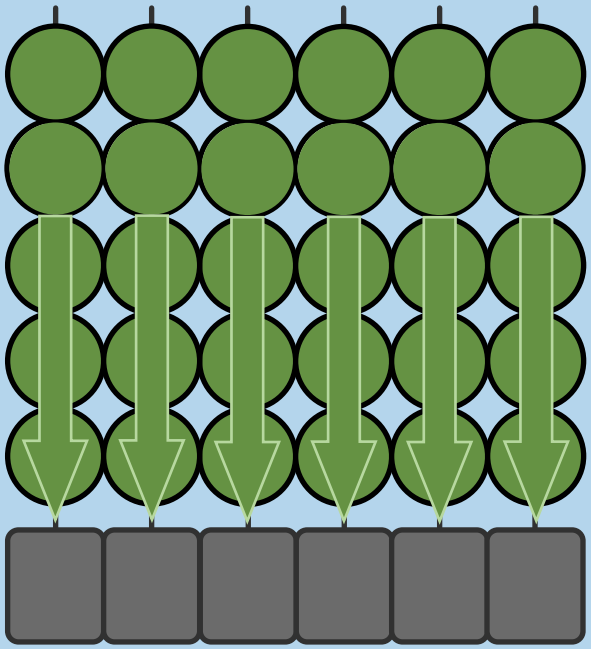

More Charge

Strong Charge Flow

Faster Sensing

115 DIMM Characterization

Timing (tRCD)

17% ↓

No Errors

Typical DIMM at Low Temperature

➔ *More charge* ➔ *Faster sensing*

Observation 2. Reducing Restore Time

Typical DIMM at Low Temperature

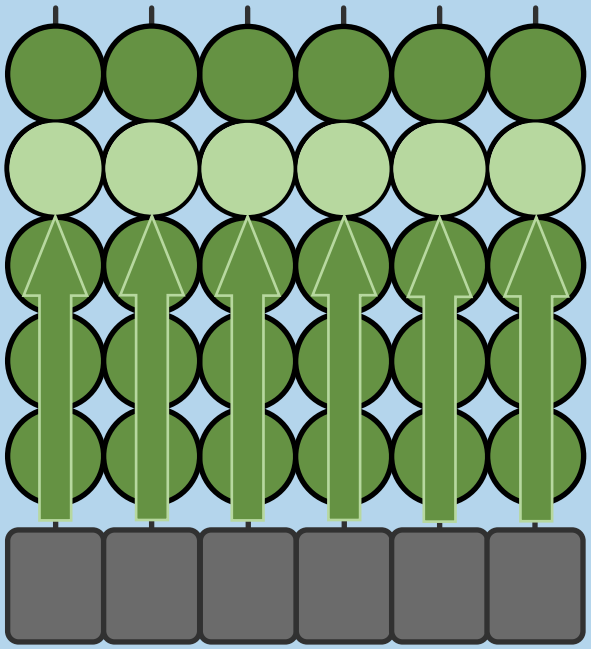

Larger Cell & Less Leakage ➔ Extra Charge

No Need to Fully Restore Charge

115 DIMM Characterization

Read (tRAS)

37% ↓

Write (tWR)

54% ↓

No Errors

Typical DIMM at lower temperature

➔ *More charge* ➔ *Restore time reduction*

Observation 3. Reducing Precharge Time

Typical DIMM at Lower Temperature

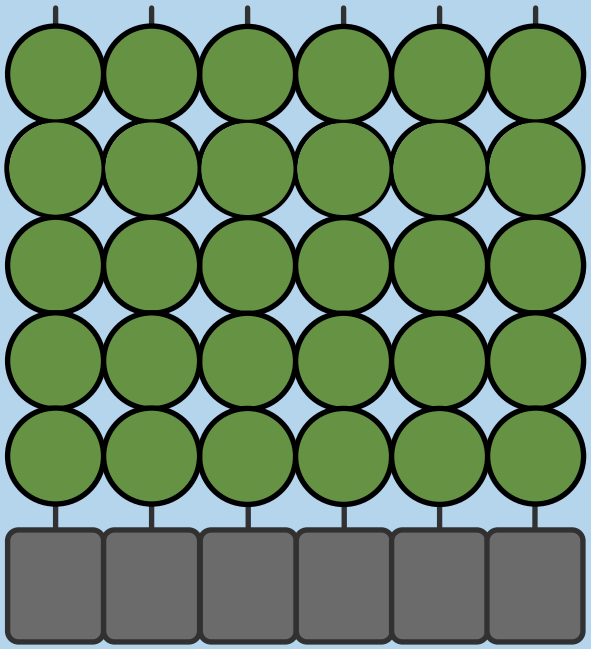

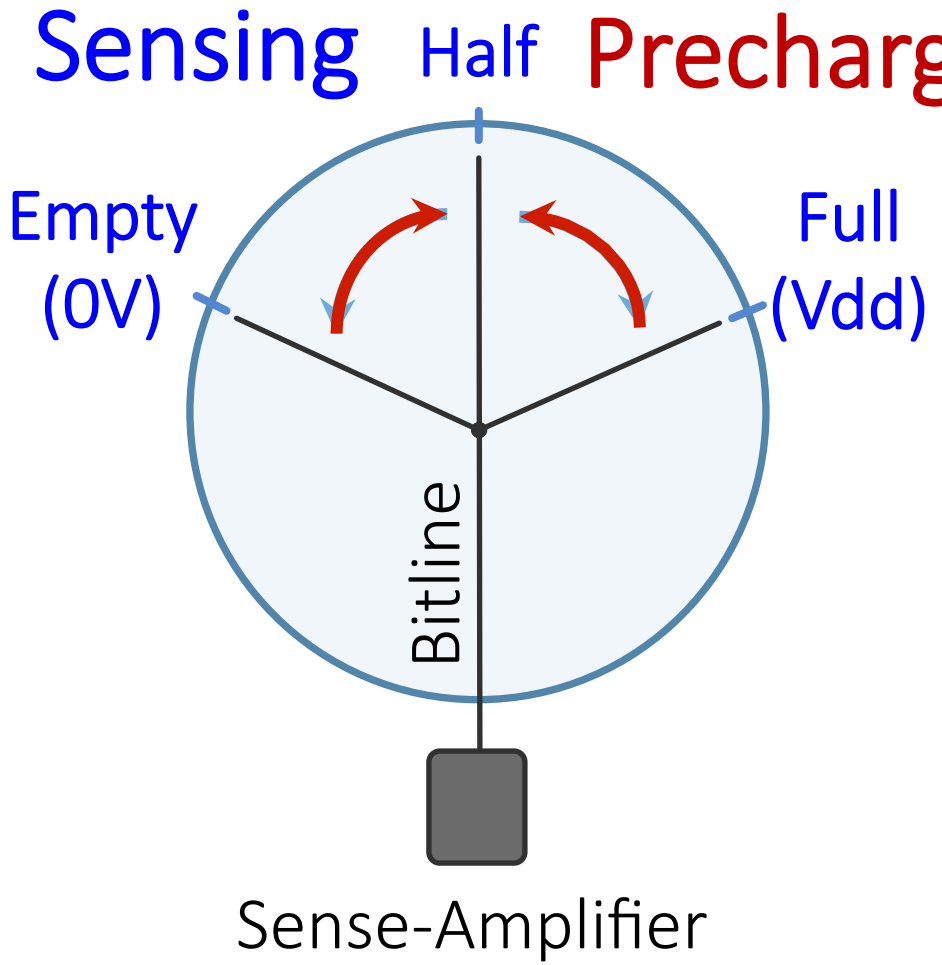

*Precharge ? – Setting bitline to **half-full charge***

Observation 3. Reducing Precharge Time

Access Empty Cell

Access Full Cell

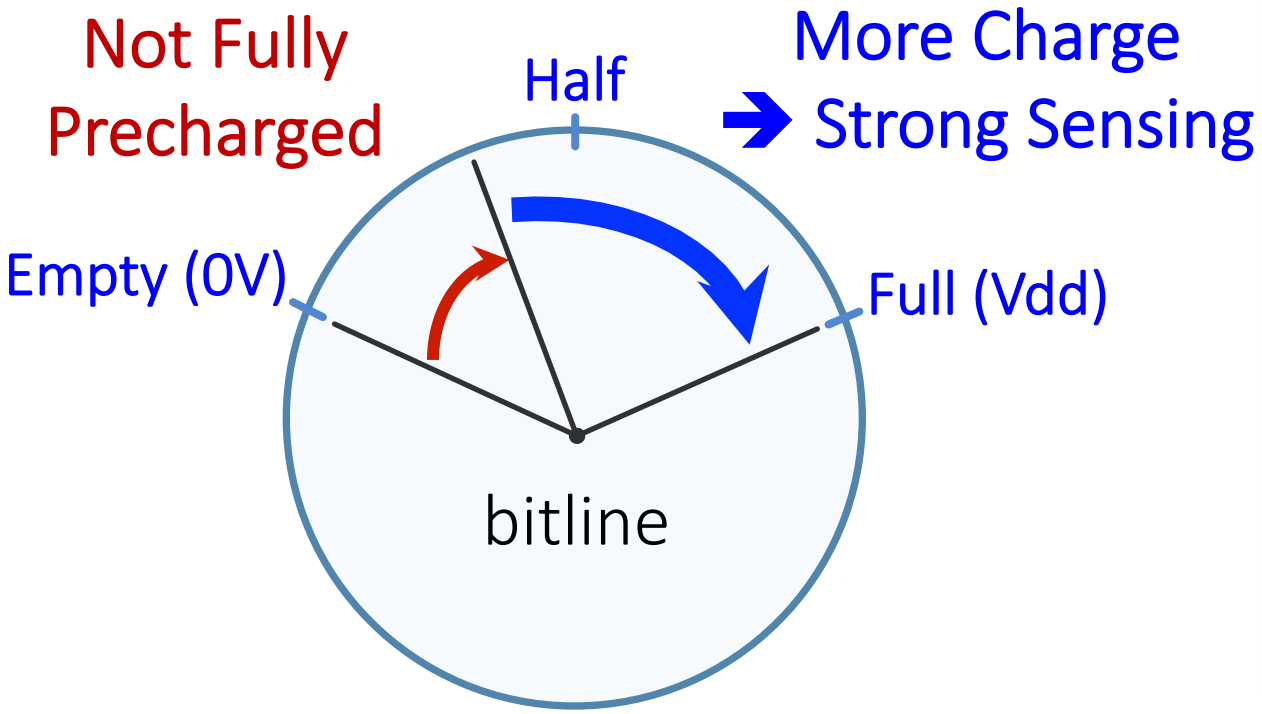

115 DIMM Characterization

Timing (tRP)

35% ↓

No Errors

Typical DIMM at Lower Temperature

➔ *More charge* ➔ *Precharge time reduction*

Key Observations

1. Sensing

   Sense **cells with extra charge** faster

   → Lower sensing latency

2. Restore

   No need to fully restore **cells with extra charge**

   → Lower restore latency

3. Precharge

   No need to fully precharge bitlines for **cells with extra charge**

   → Lower precharge latency

1. DRAM Operation Basics

2. Reasons for Timing Margin in DRAM

3. Key Observations

4. Adaptive-Latency DRAM

5. DRAM Characterization

6. Real System Performance Evaluation

Adaptive-Latency DRAM

- *Key idea*
  - **Optimize DRAM timing parameters online**
- *Two components*
  - DRAM manufacturer profiles multiple sets of reliable DRAM timing parameters at different temperatures for each DIMM
  - System monitors DRAM temperature & uses appropriate DRAM timing parameters

1. DRAM Operation Basics

2. Reasons for Timing Margin in DRAM

3. Key Observations

4. Adaptive-Latency DRAM

5. DRAM Characterization

6. Real System Performance Evaluation

DRAM Temperature

- *DRAM temperature measurement*
  - Server cluster: Operates at under 34°C
  - Desktop: Operates at under 50°C
  - *DRAM standard optimized for 85°C*

DRAM operates at low temperatures in the common-case

- *Previous works – Maintain DRAM temperature low*
  - David+ ICAC 2011
  - Liu+ ISCA 2007
  - Zhu+ ITHERM 2008

DRAM Testing Infrastructure

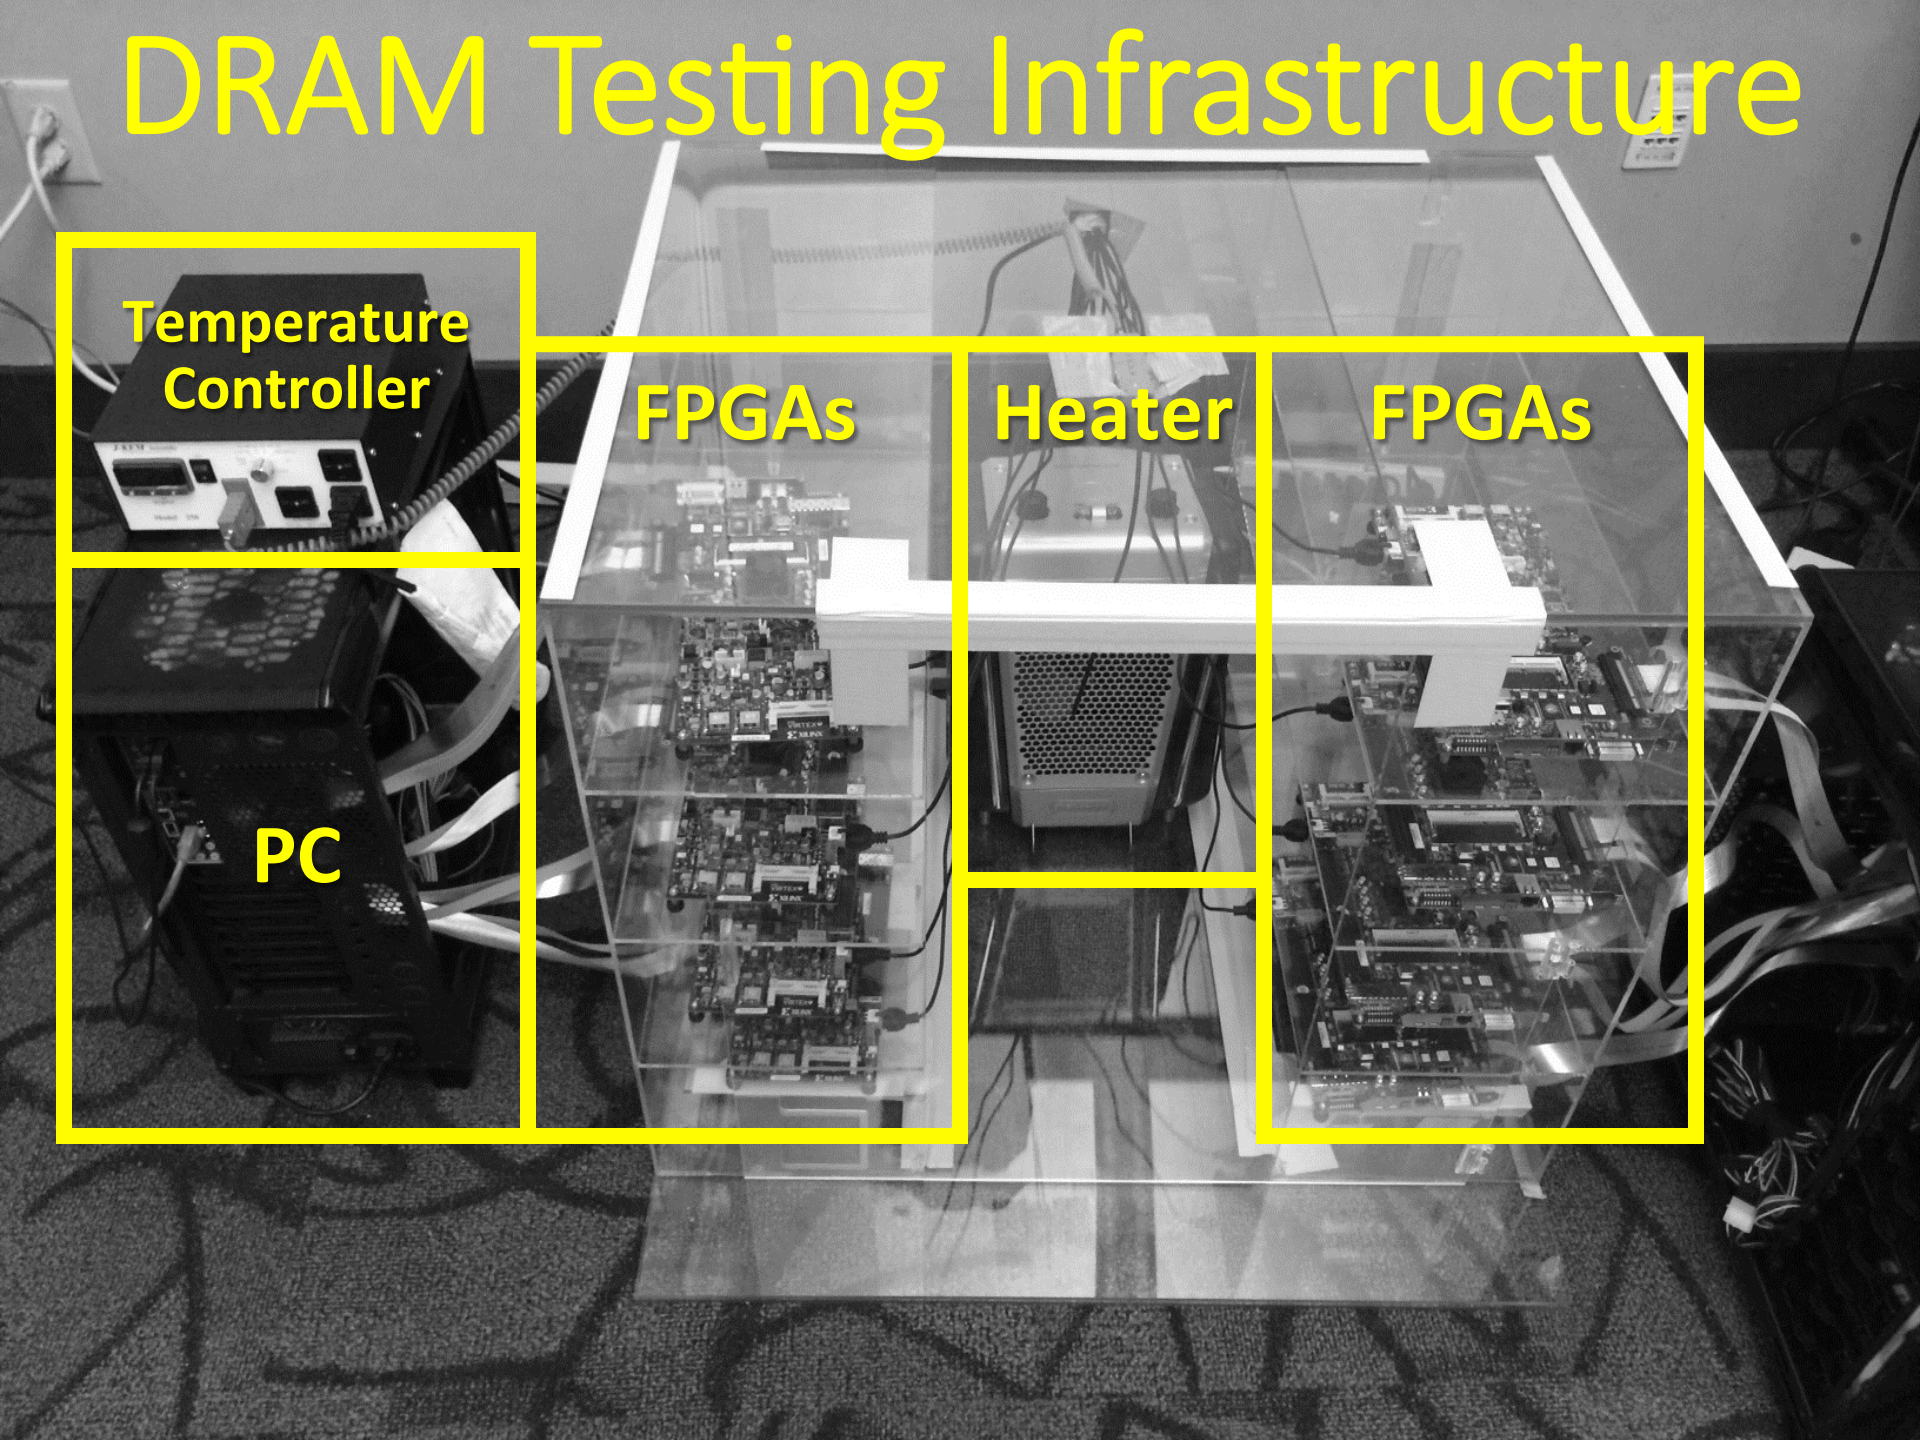

Test Pattern

- *Single cache line test (Read/Write)*

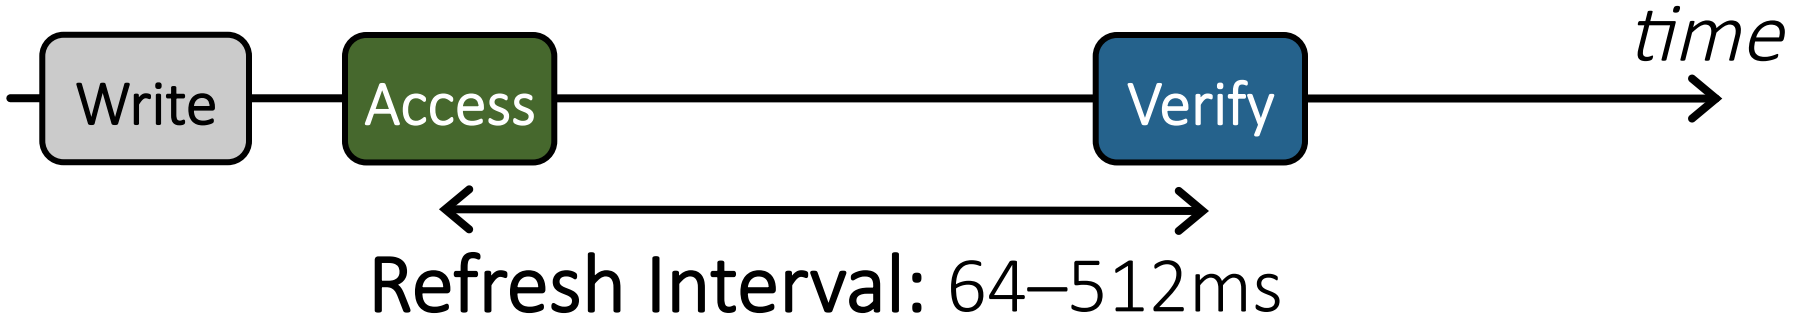

- *Overlapping multiple single cache line tests to simulate power noise and coupling*

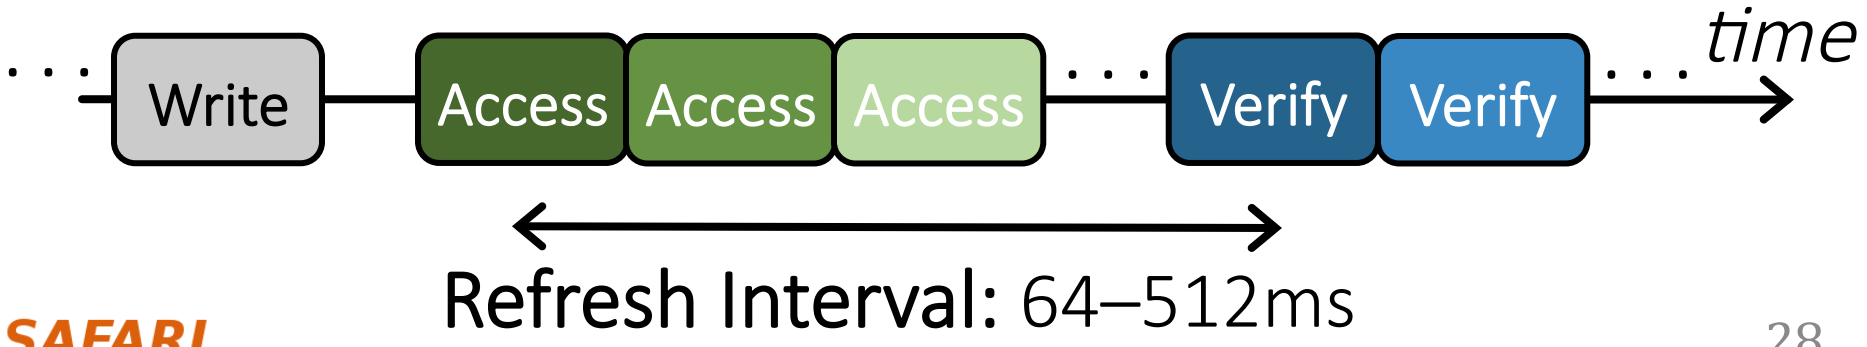

Control Factors

- *Timing parameters*
  - Sensing: tRCD
  - Restore: tRAS (read), tWR(write)
  - Precharge: tRP
- *Temperature: 55 – 85°C*
- *Refresh interval: 64 – 512ms*
  - Longer refresh interval leads to smaller charge
  - Standard refresh interval: 64ms

1. Timings ↔ Charge

Temperature: 85°C/Refresh Interval: 64, 128, 256, 512ms

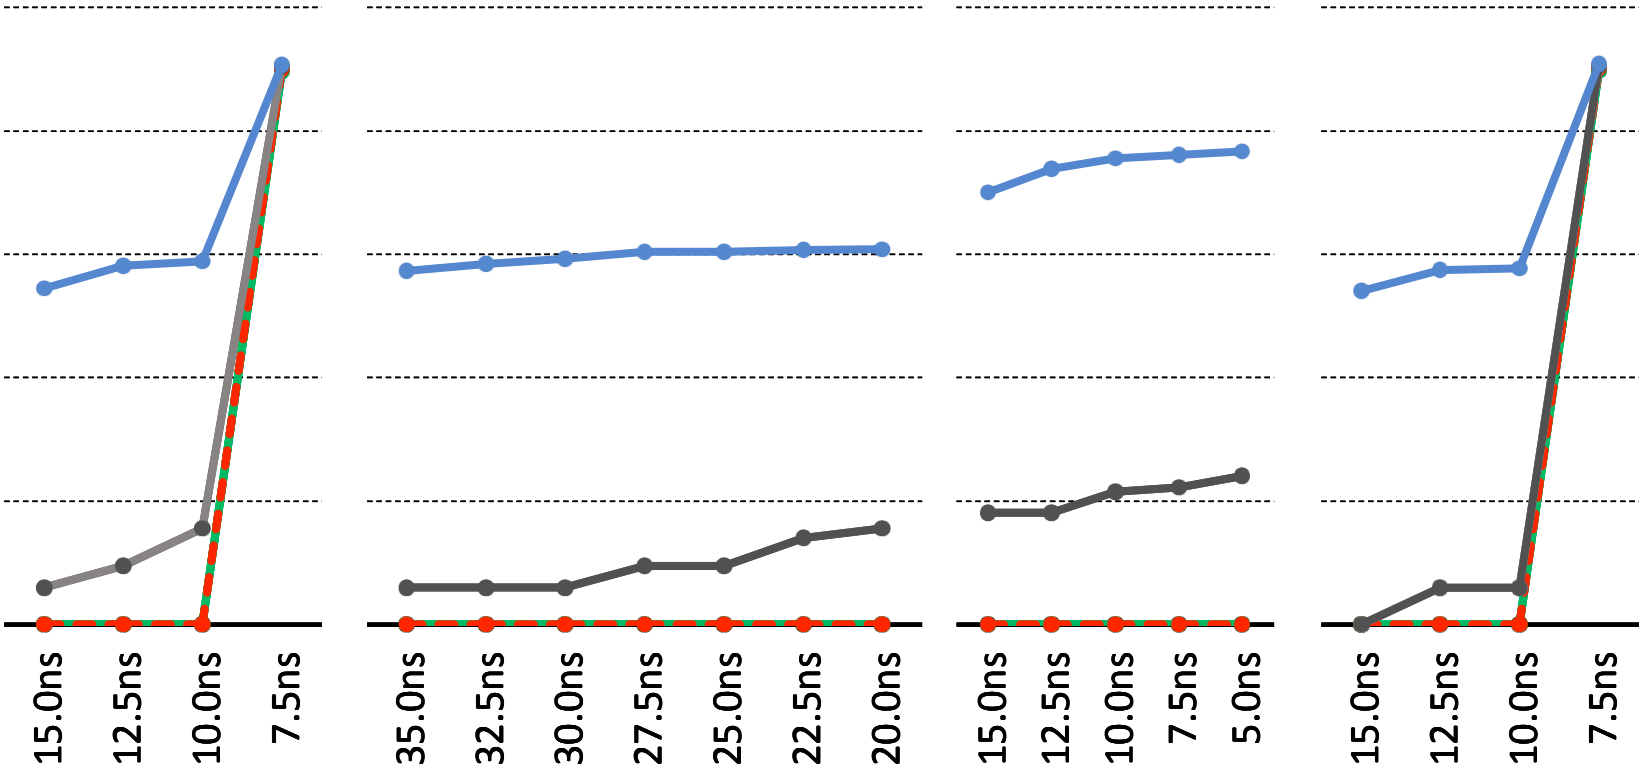

More charge enables

more timing parameter reduction

2. Timings ↔ Temperature

*Temperature: 55, 65, 75, 85°C/Refresh Interval: **512ms***

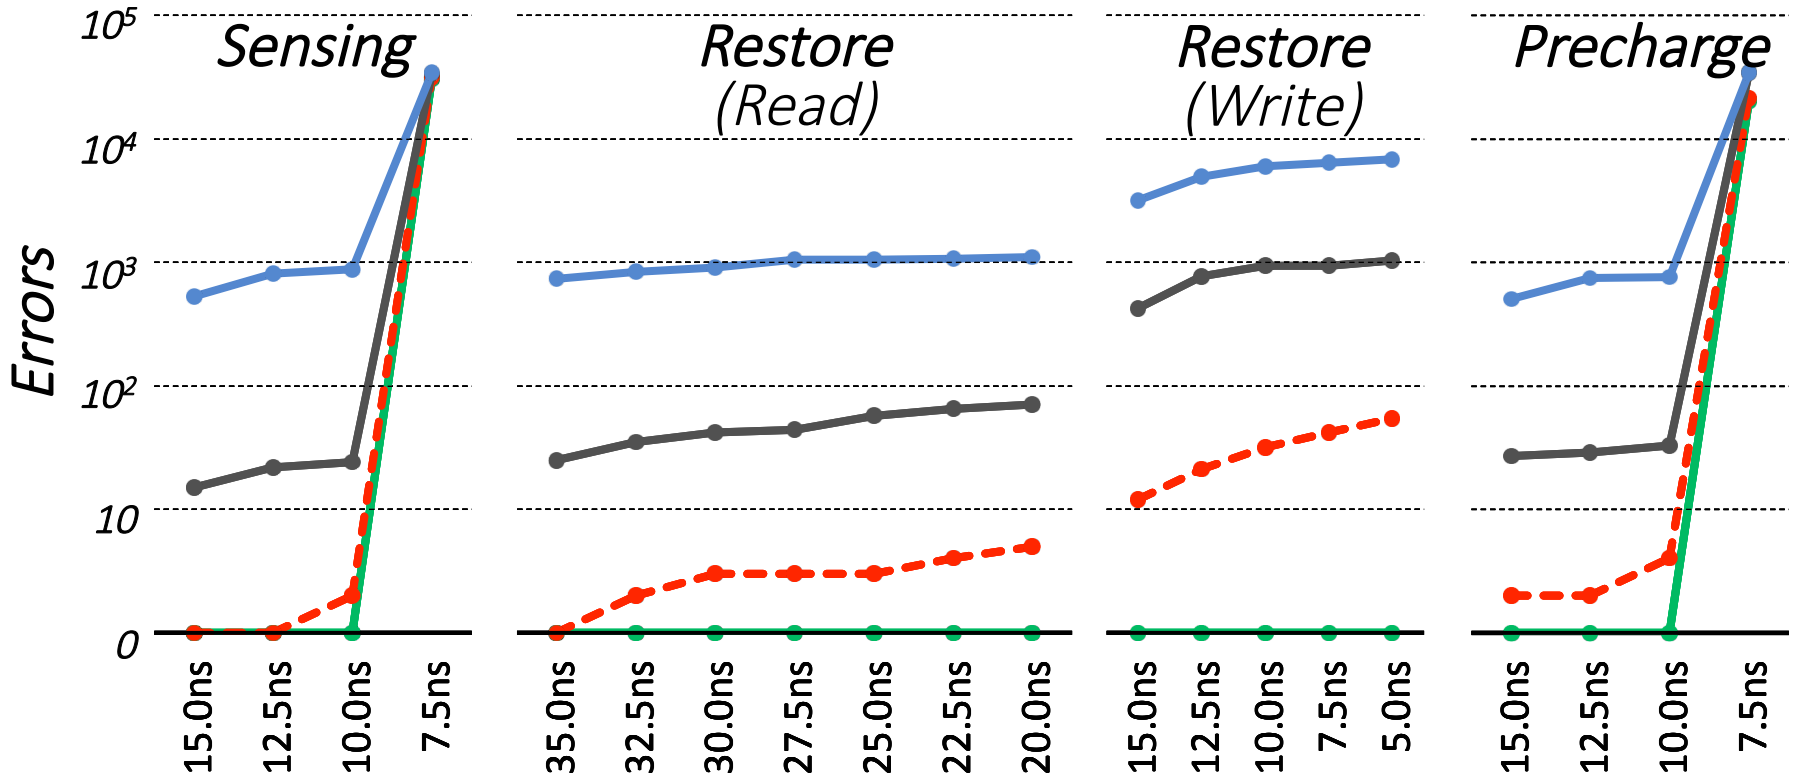

Lower temperature enables

more timing parameter reduction

3. Summary of 115 DIMMs

- *Latency reduction for read & write (55°C)*
  - *Read Latency: **32.7%***
  - *Write Latency: **55.1%***
- *Latency reduction for each timing parameter (55°C)*
  - *Sensing: **17.3%***
  - *Restore: **37.3%** (read), **54.8%** (write)*
  - *Precharge: **35.2%***

1. DRAM Operation Basics

2. Reasons for Timing Margin in DRAM

3. Key Observations

4. Adaptive-Latency DRAM

5. DRAM Characterization

6. Real System Performance Evaluation

Real System Evaluation Method

- *System*
  - *CPU: AMD 4386 (8 Cores, 3.1GHz, 8MB LLC)*

D18F2x200_dct[0]_mp[1:0] DDR3 DRAM Timing 0

Reset: 0F05_0505h. See 2.9.3 [DCT Configuration Registers].

Bits	Description
31:30	Reserved.
29:24	**Tras: row active strobe**. Read-write. BIOS: See 2.9.7.5 [SPD ROM-Based Configuration]. Specifies the minimum time in memory clock cycles from an activate command to a precharge command, both to the same chip select bank. Bits — Description 07h-00h — Reserved 2Ah-08h — <Tras> clocks 3Fh-2Bh — Reserved
23:21	Reserved.
20:16	**Trp: row precharge time**. Read-write. BIOS: See 2.9.7.5 [SPD ROM-Based Configuration]. Specifies the minimum time in memory clock cycles from a precharge command to an activate command or auto refresh command, both to the same bank.

Single-Core Evaluation

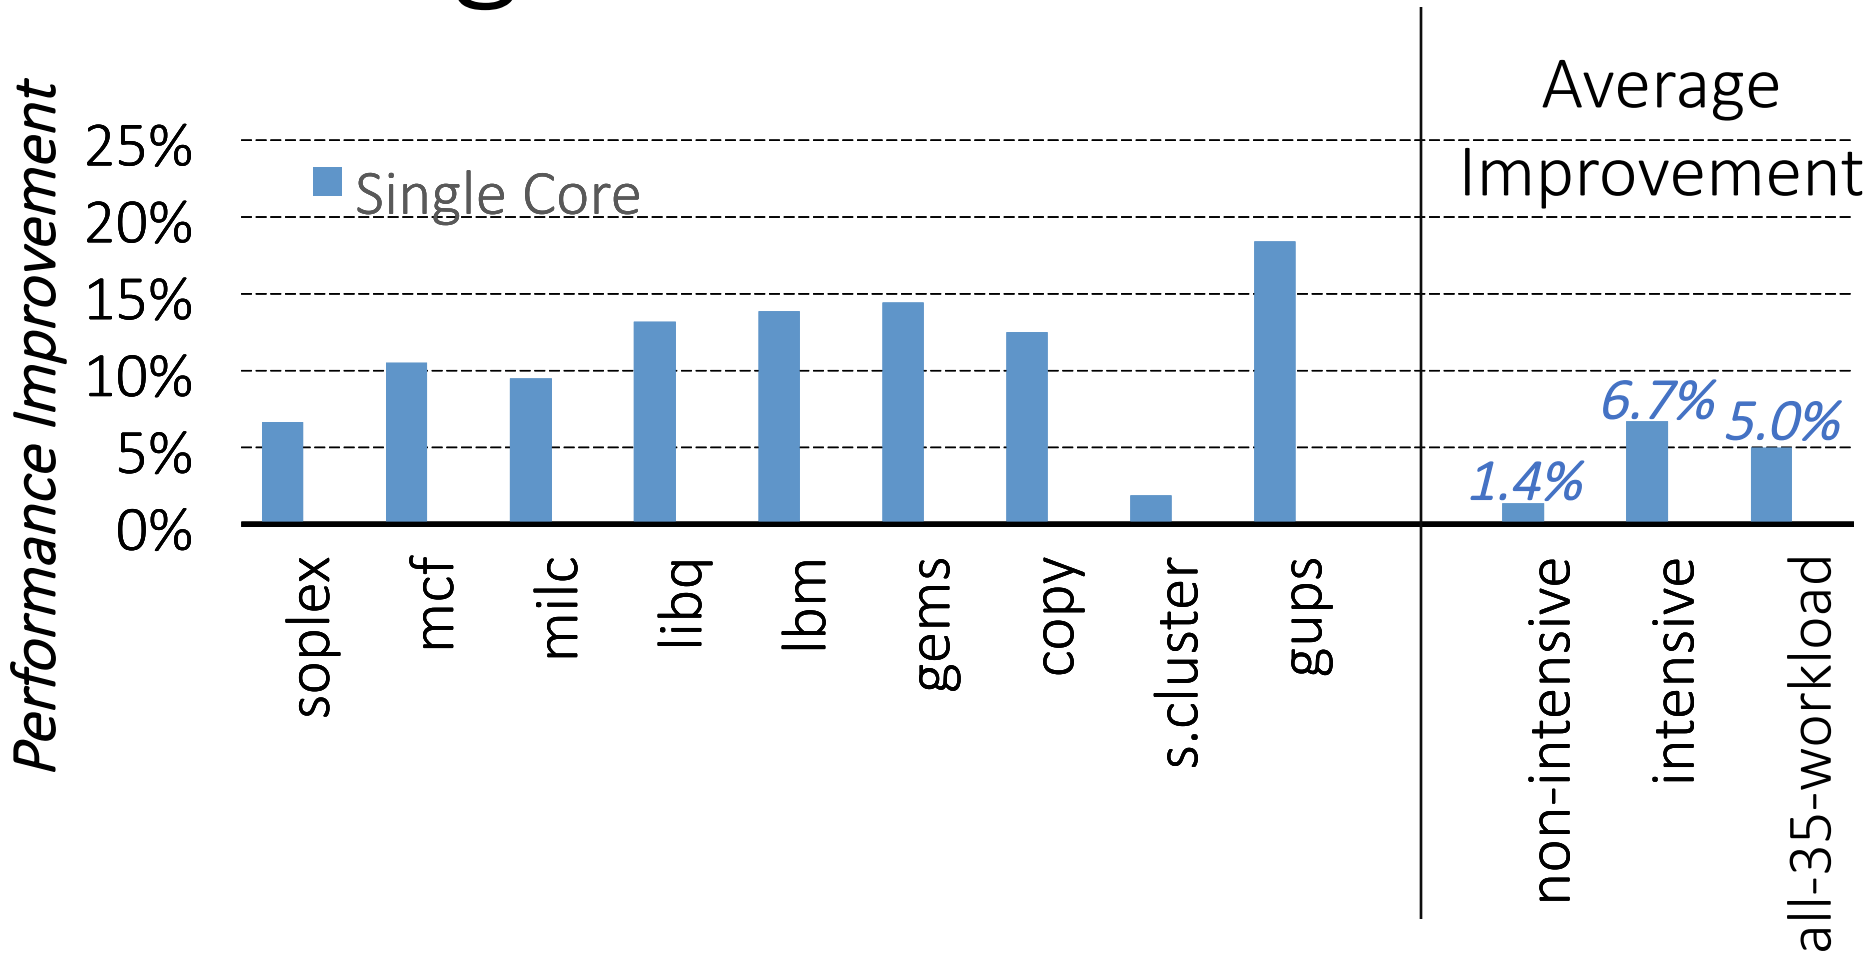

*AL-DRAM **improves performance** on a real system*

Multi-Core Evaluation

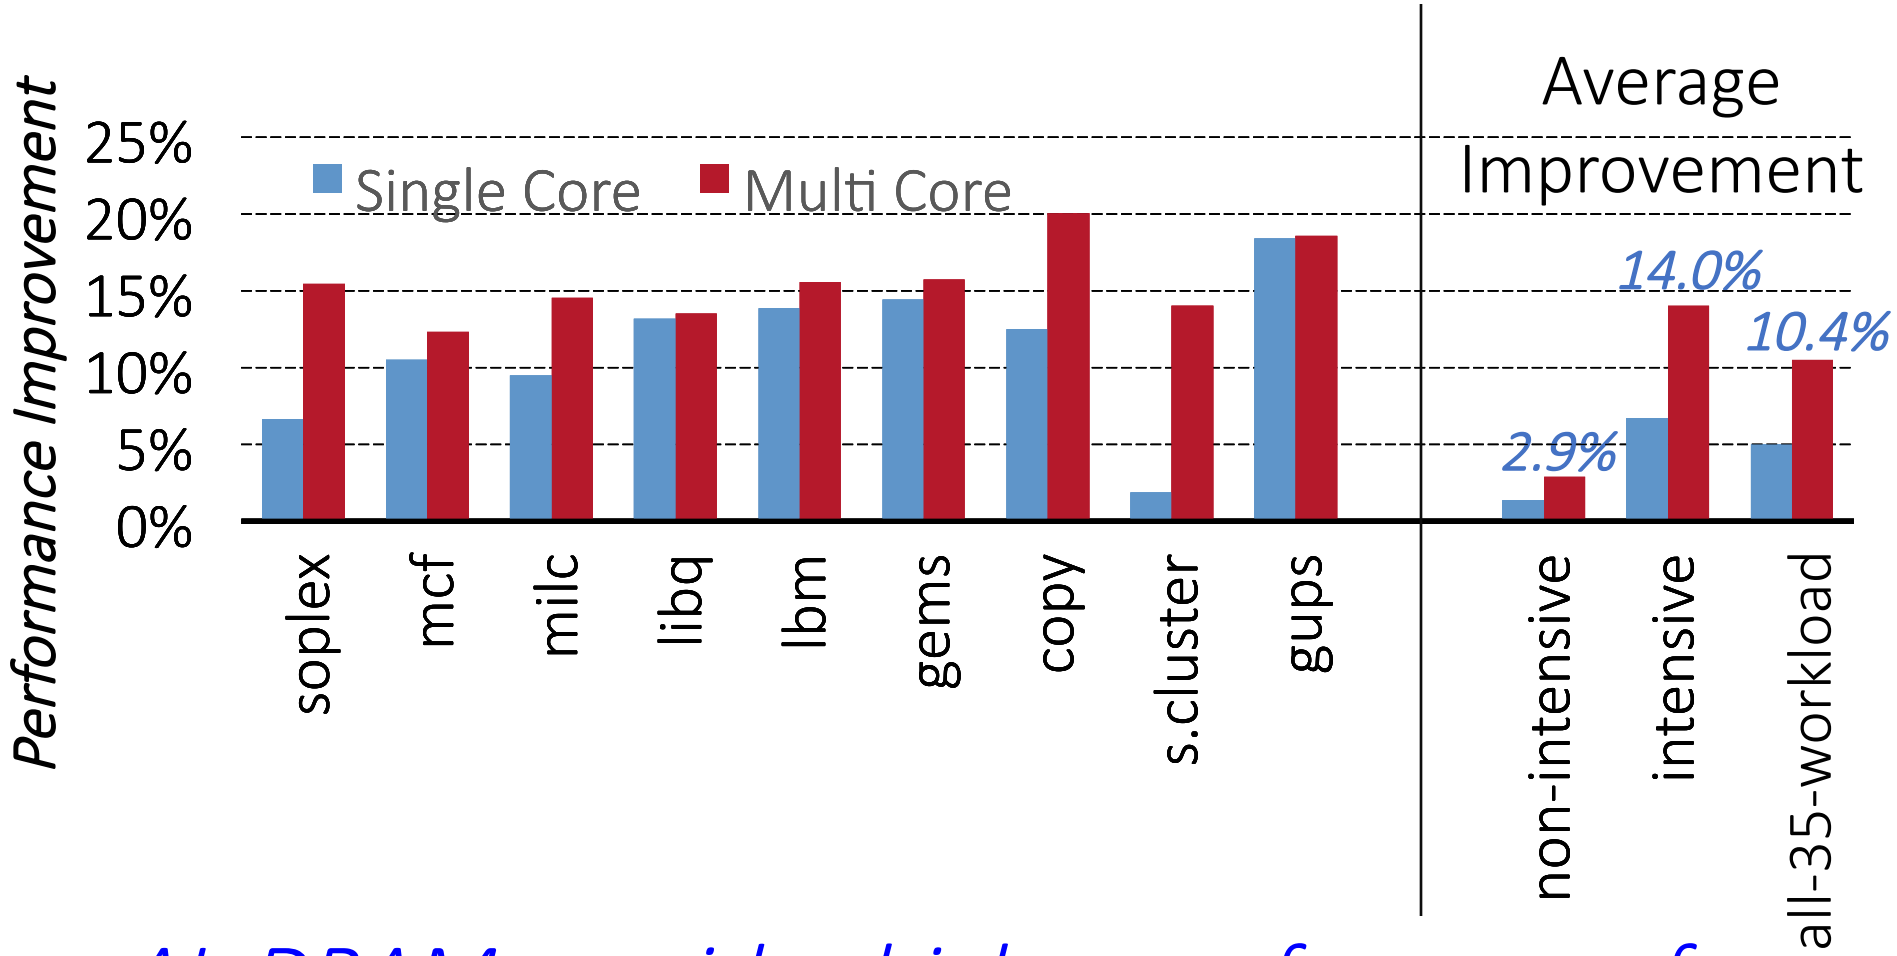

AL-DRAM provides higher performance for ***multi-programmed & multi-threaded workloads***

Conclusion

- *Observations*
  - *DRAM timing parameters are dictated by the worst-case cell (smallest cell across all products at highest temperature)*
  - *DRAM operates at lower temperature than the worst case*
- *Idea: Adaptive-Latency DRAM*
  - *Optimizes DRAM timing parameters for the common case (typical DIMM operating at low temperatures)*
- *Analysis: Characterization of 115 DIMMs*
  - *Great potential to lower DRAM timing parameters (17 – 54%) without any errors*
- *Real System Performance Evaluation*
  - *Significant performance improvement (14% for memory-intensive workloads) without errors (33 days)*

Optimizing DRAM Timing for the Common-Case

Adaptive-Latency DRAM

Donghyuk Lee

Yoongu Kim, Gennady Pekhimenko, Samira Khan,

Vivek Seshadri, Kevin Chang, Onur Mutlu

SAFARI

Carnegie Mellon

Backup Slides

Overhead

- *DRAM Manufacturer*
  - *Additional tests: can be integrated into existing test process (i.e., TCSR test)*
- *DRAM (DIMM)*
  - *Already have in-DRAM temperature sensor (i.e., Low Power DDR)*
  - *Multiple sets of timing parameters can be stored in SPD (Serial Presence Detect)*
- *System Support for AL-DRAM*
  - *Already have ability to change DRAM timing online*

Multiple Timing Parameters

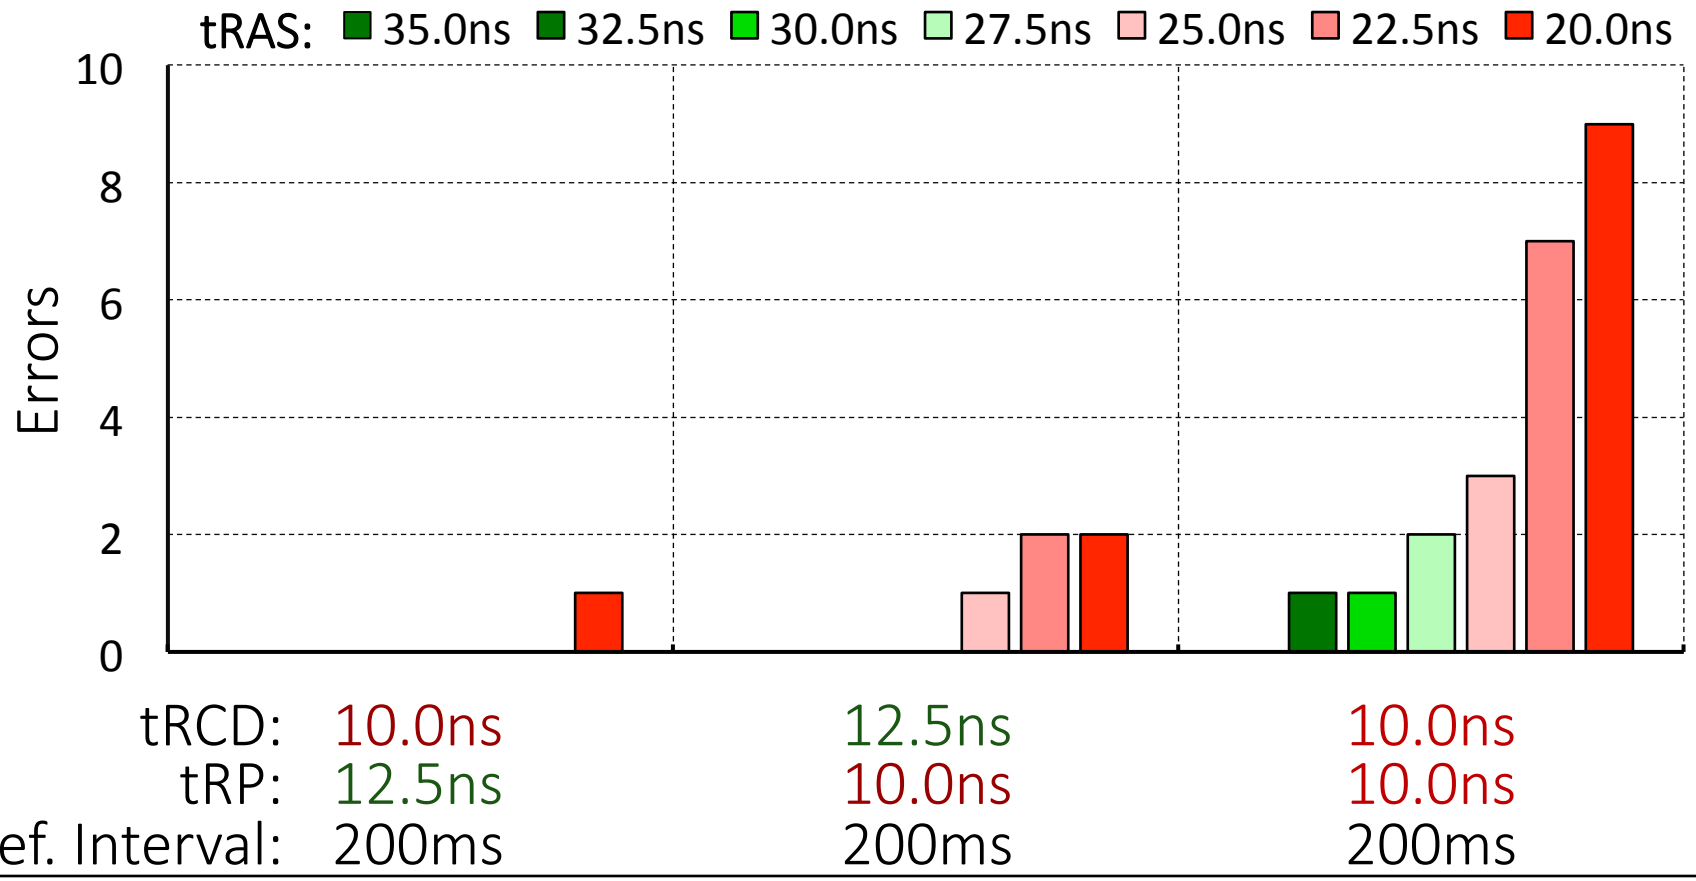

Reducing a timing parameter

➔ Reduces potential reduction of other parameters

Temperature ↔ Refresh Interval

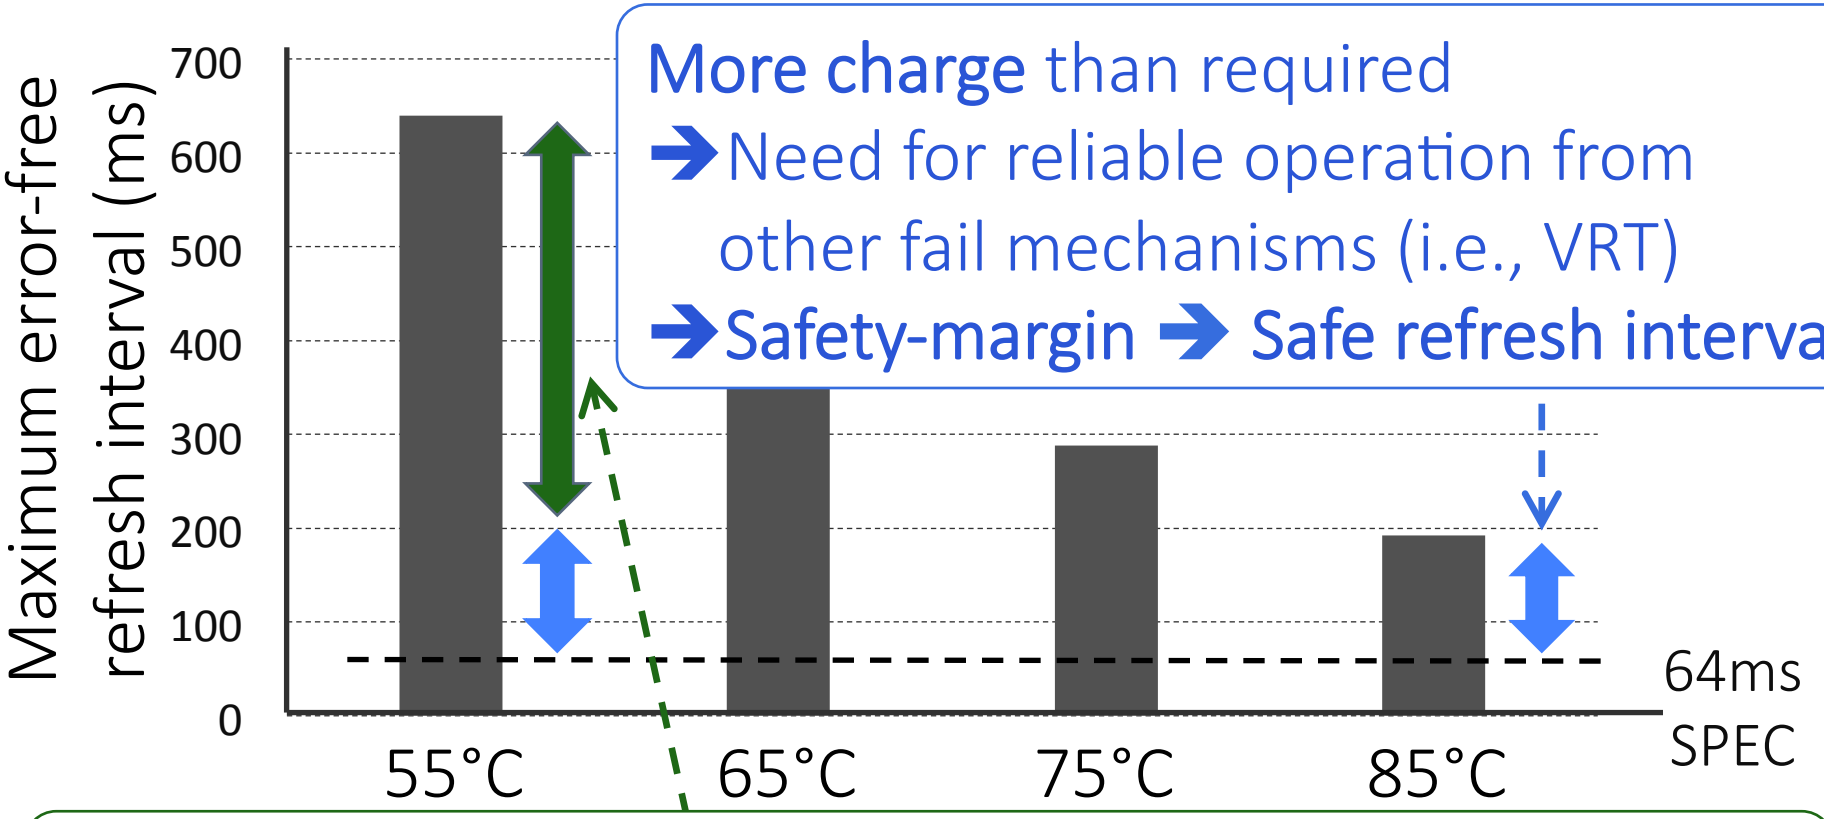

More charge than required
➔Need for reliable operation from other fail mechanisms (i.e., VRT)
➔**Safety-margin** ➔ **Safe refresh interval**

Extra charge that can be used for latency reduction

DRAM Cell Organization

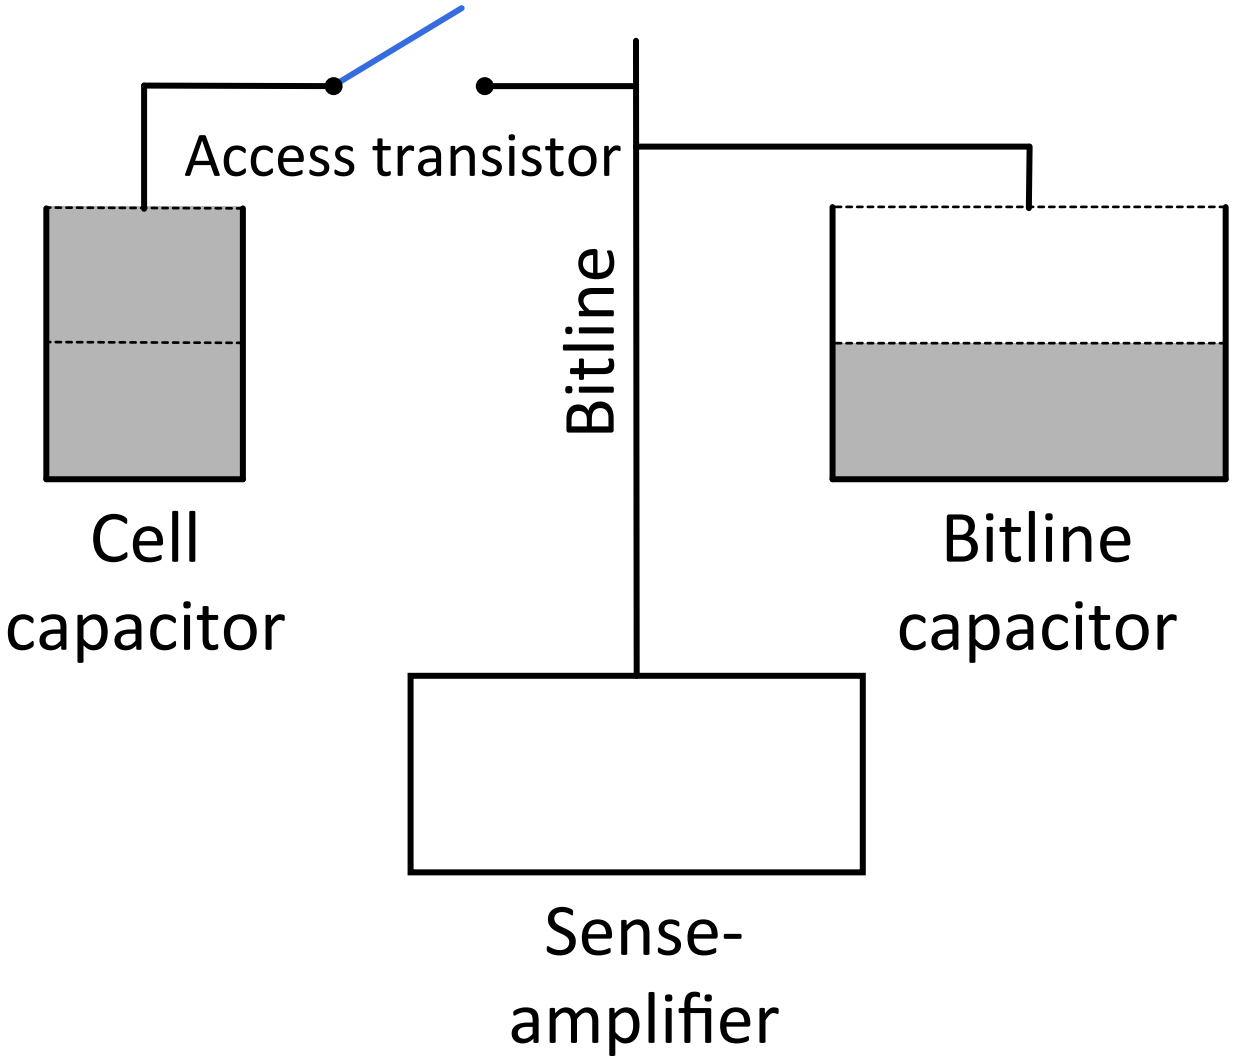

DRAM Cell Operation

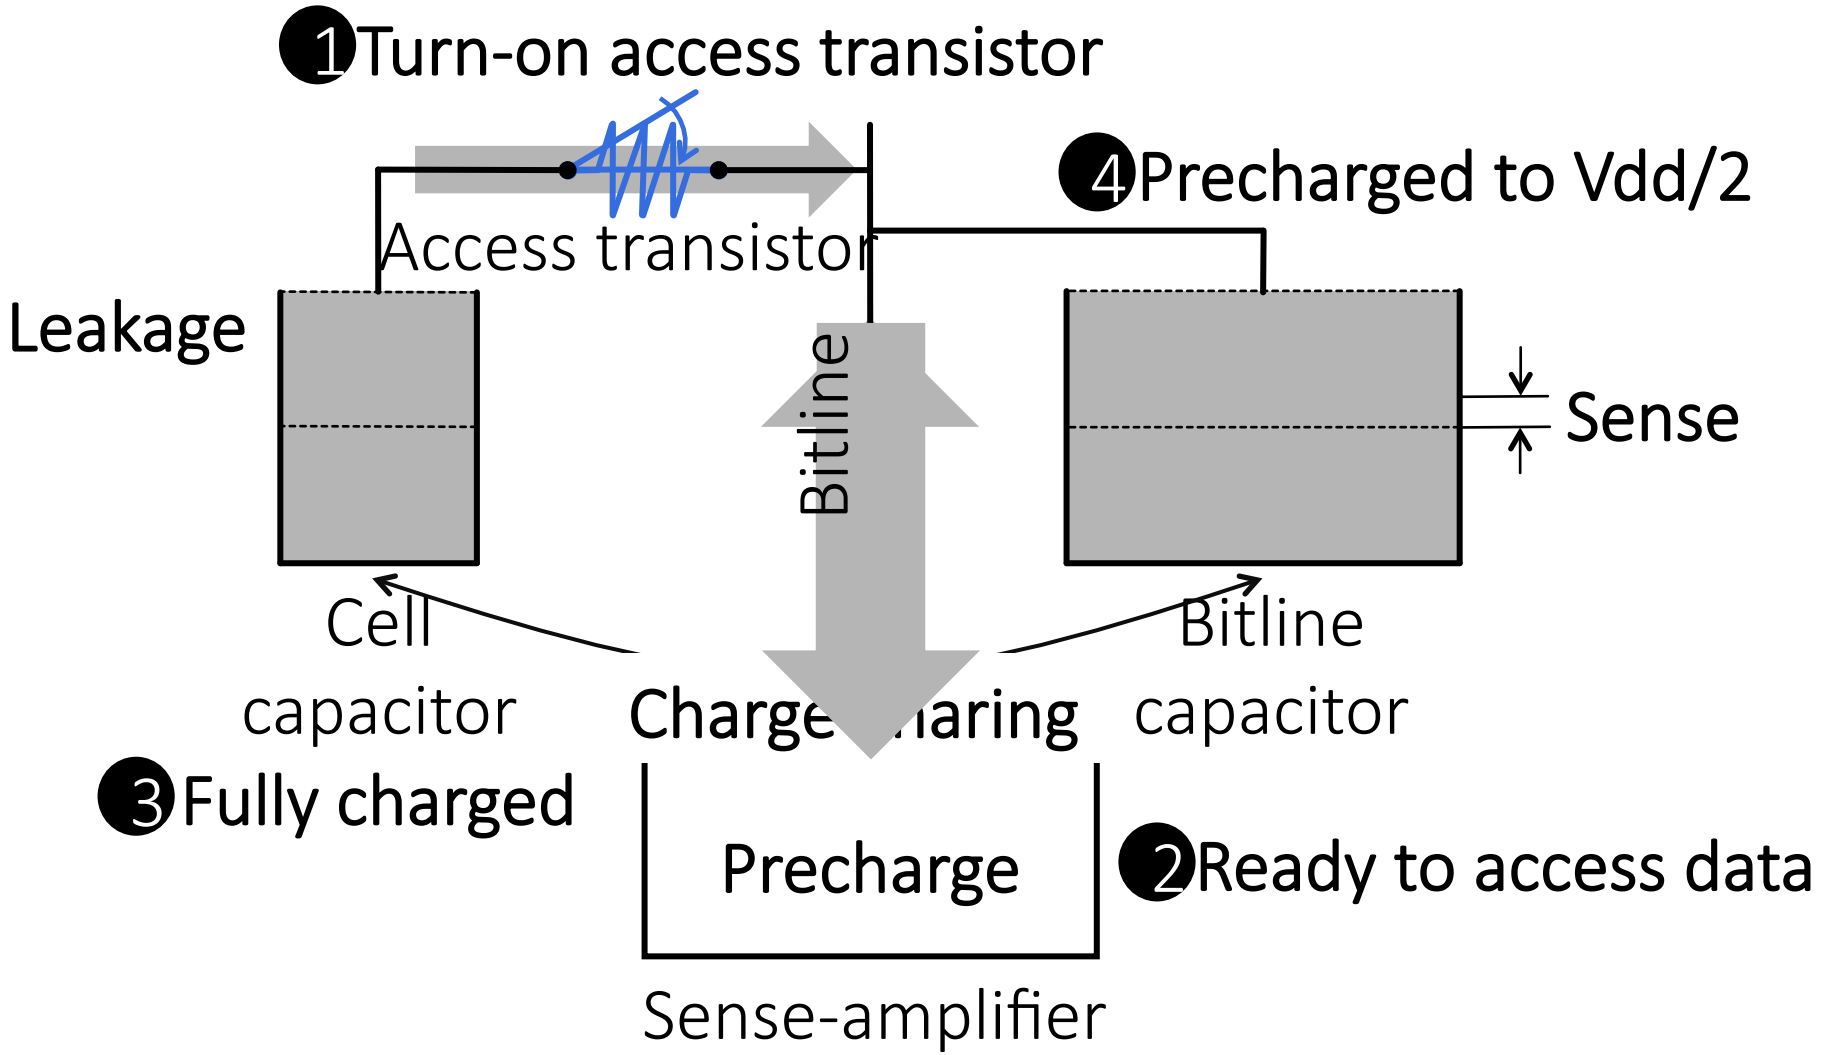

DRAM Cell Charge Variations