Solar-DRAM:

Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines

Jeremie S. Kim **Minesh Patel**

Hasan Hassan **Onur Mutlu**

Systems@ETH Zürich

SAFARI

ETH Zürich

Carnegie Mellon

Executive Summary

Motivation: DRAM latency is a **major performance bottleneck**

Problem: Many important workloads exhibit **bank conflicts** in DRAM, which result in even longer latencies

Goal:

1. Rigorously **characterize access latency** on LPDDR4 DRAM
2. Exploit findings to **robustly reduce DRAM access latency**

Solar-DRAM:

- Categorizes local bitlines as "*weak (slow)*" or "*strong (fast)*"
- Robustly **reduces DRAM access latency for reads *and* writes** to data contained in "*strong*" local bitlines.

Evaluation:

1. Experimentally characterize **282** real LPDDR4 DRAM chips
2. In simulation, **Solar-DRAM** provides **10.87%** system performance improvement over LPDDR4 DRAM

Solar-DRAM Outline

- Motivation and Goal
- DRAM Background
- Experimental Methodology
- Characterization Results
- Mechanism: Solar-DRAM
- Evaluation
- Conclusion

Solar-DRAM Outline

- **Motivation and Goal**
- DRAM Background
- Experimental Methodology
- Characterization Results
- Mechanism: Solar-DRAM
- Evaluation
- Conclusion

Motivation and Goal

- Many important workloads exhibit many bank conflicts
  - **Bank conflicts** result in an additional delay of $\mathbf{t_{RCD}}$
  - This negatively impacts overall system performance
- A prior work (FLY-DRAM) finds **weak (slow) cells** and uses variable $\mathbf{t_{RCD}}$ depending on cell strength, **however**
  - They do ***not*** show the **viability of static profile** of cell strength
  - They characterize an **older** generation (DDR3) of DRAM
- **Our goal** is to
  - **Rigorously characterize** *state-of-the-art* LPDDR4 DRAM
  - **Demonstrate** **viability of using static profile** of cell strength
  - **Devise** a mechanism to **exploit more activation failure ($\mathbf{t_{RCD}}$) characteristics** and **further reduce DRAM latency**

Solar-DRAM Outline

- Motivation and Goal
- DRAM Background
- Experimental Methodology
- Characterization Results
- Mechanism: Solar-DRAM
- Evaluation
- Conclusion

DRAM Background

Each DRAM cell is made of 1 capacitor and 1 transistor

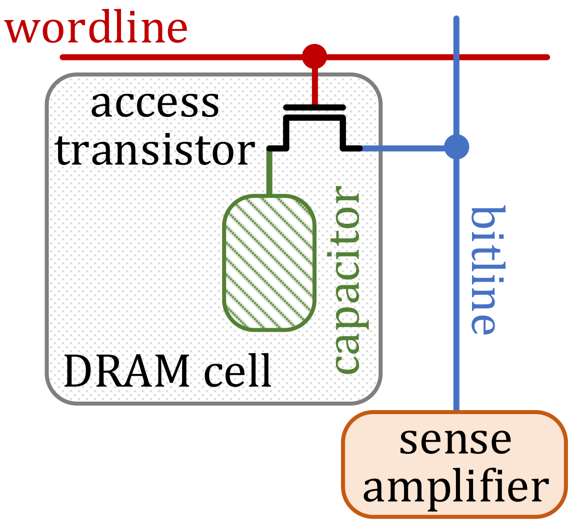

Wordline enables reading/writing data in the cell
Bitline moves data from cells to/from I/O circuitry

DRAM Background

A DRAM bank is organized hierarchically with **subarrays**

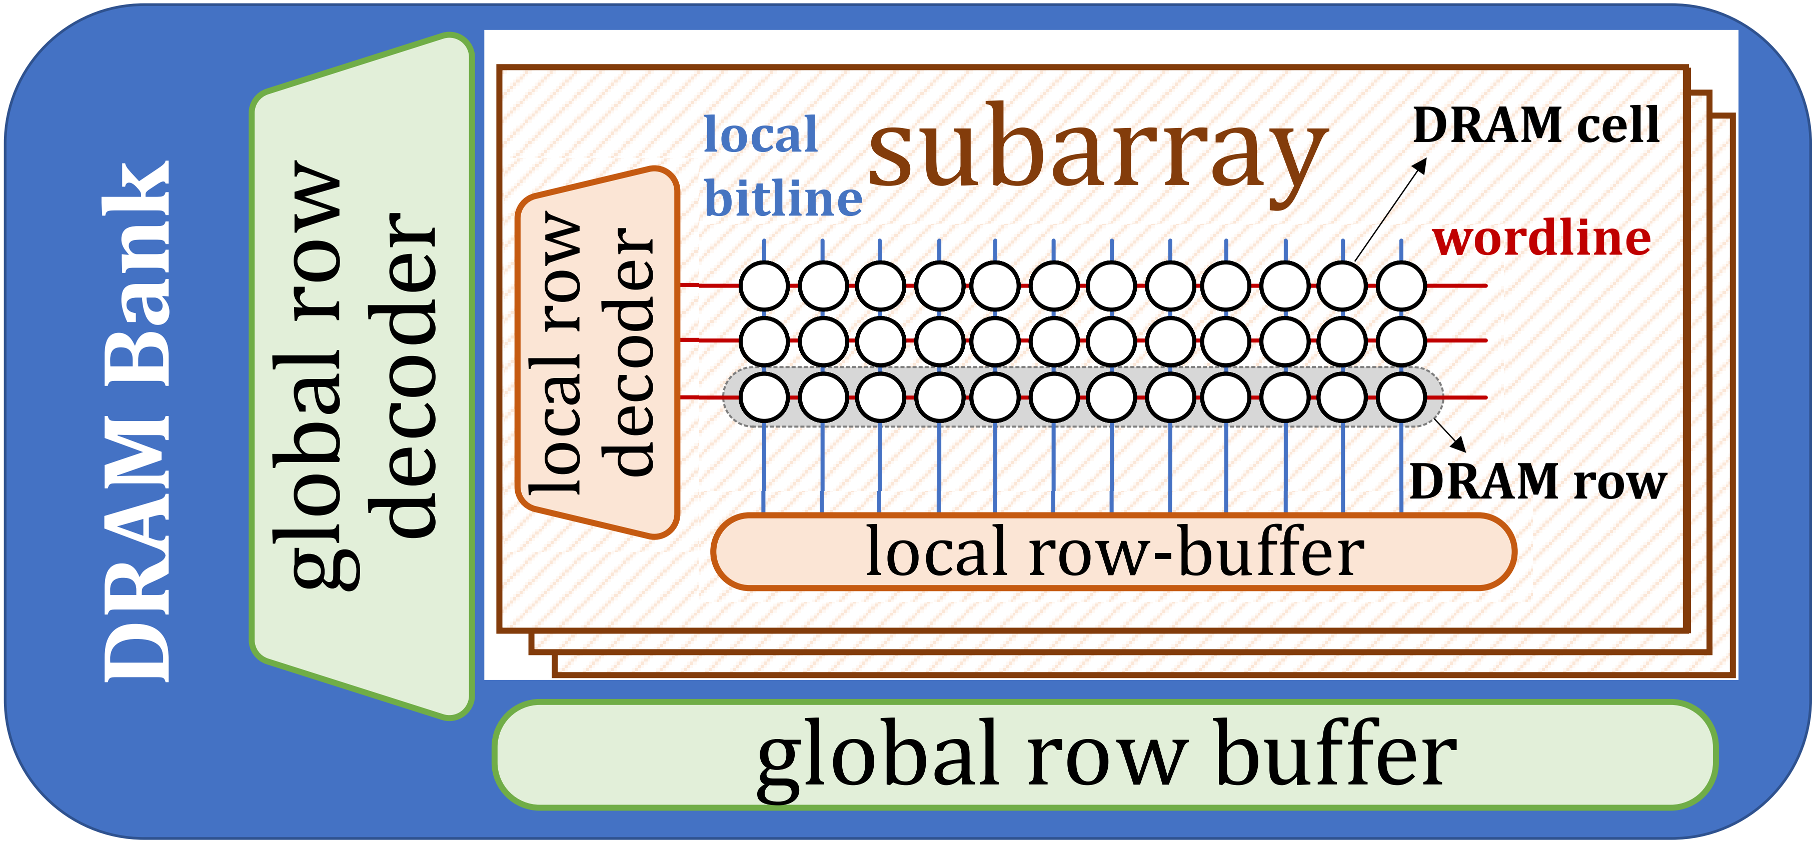

Columns of cells in subarrays share a **local bitline**

Rows of cells in a subarray share a **wordline**

DRAM Operation

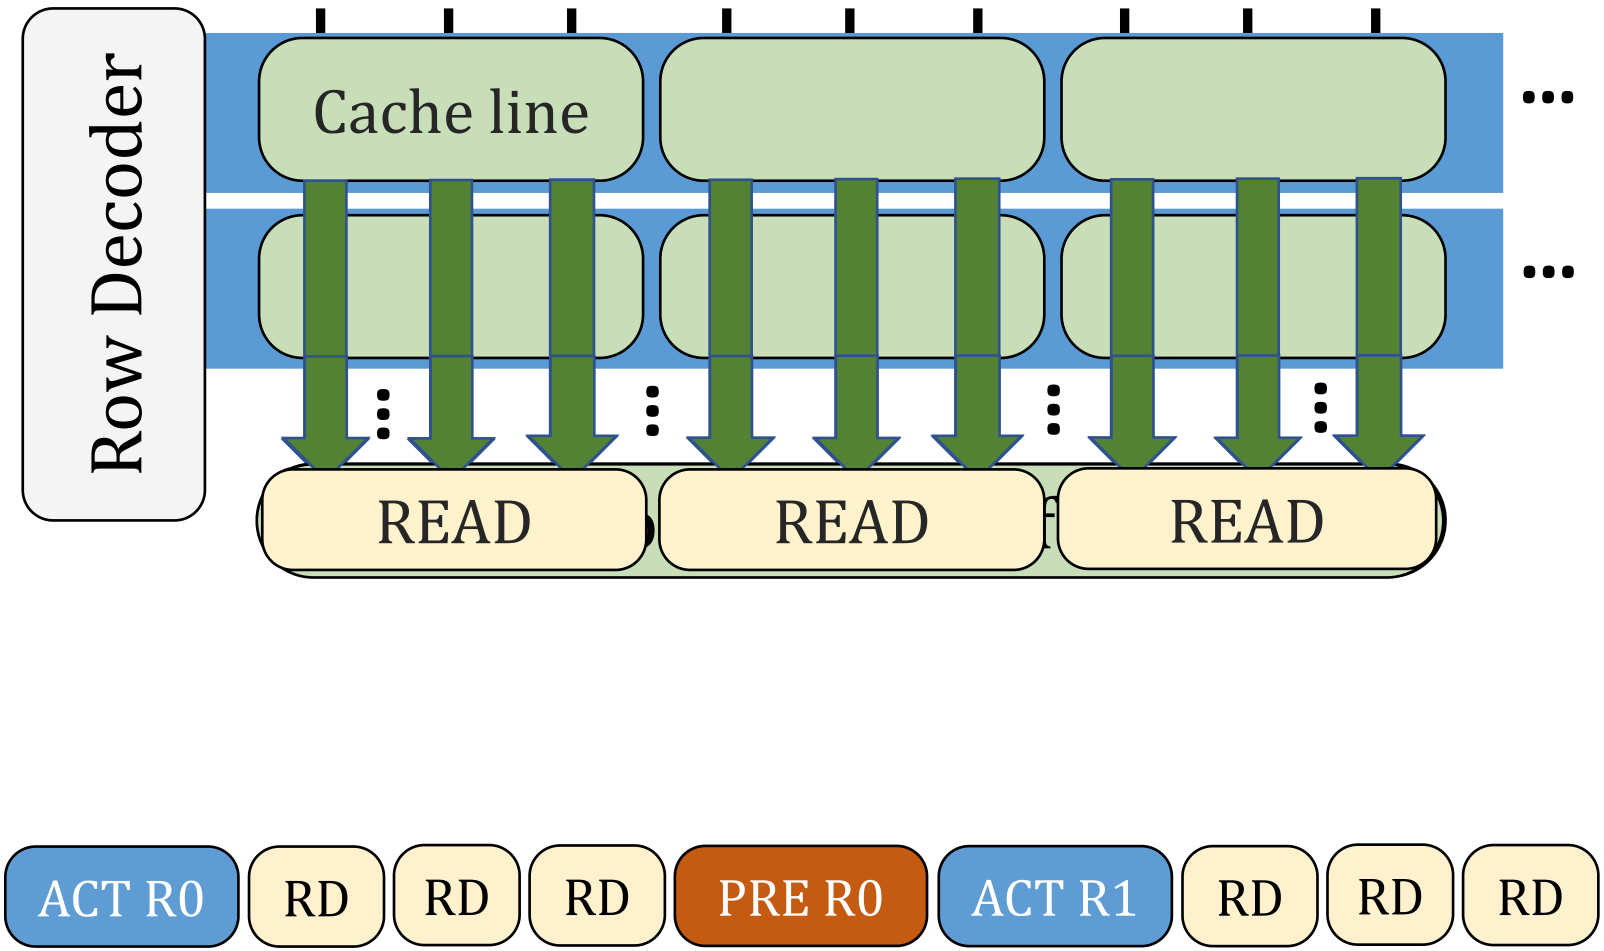

DRAM Accesses and Failures

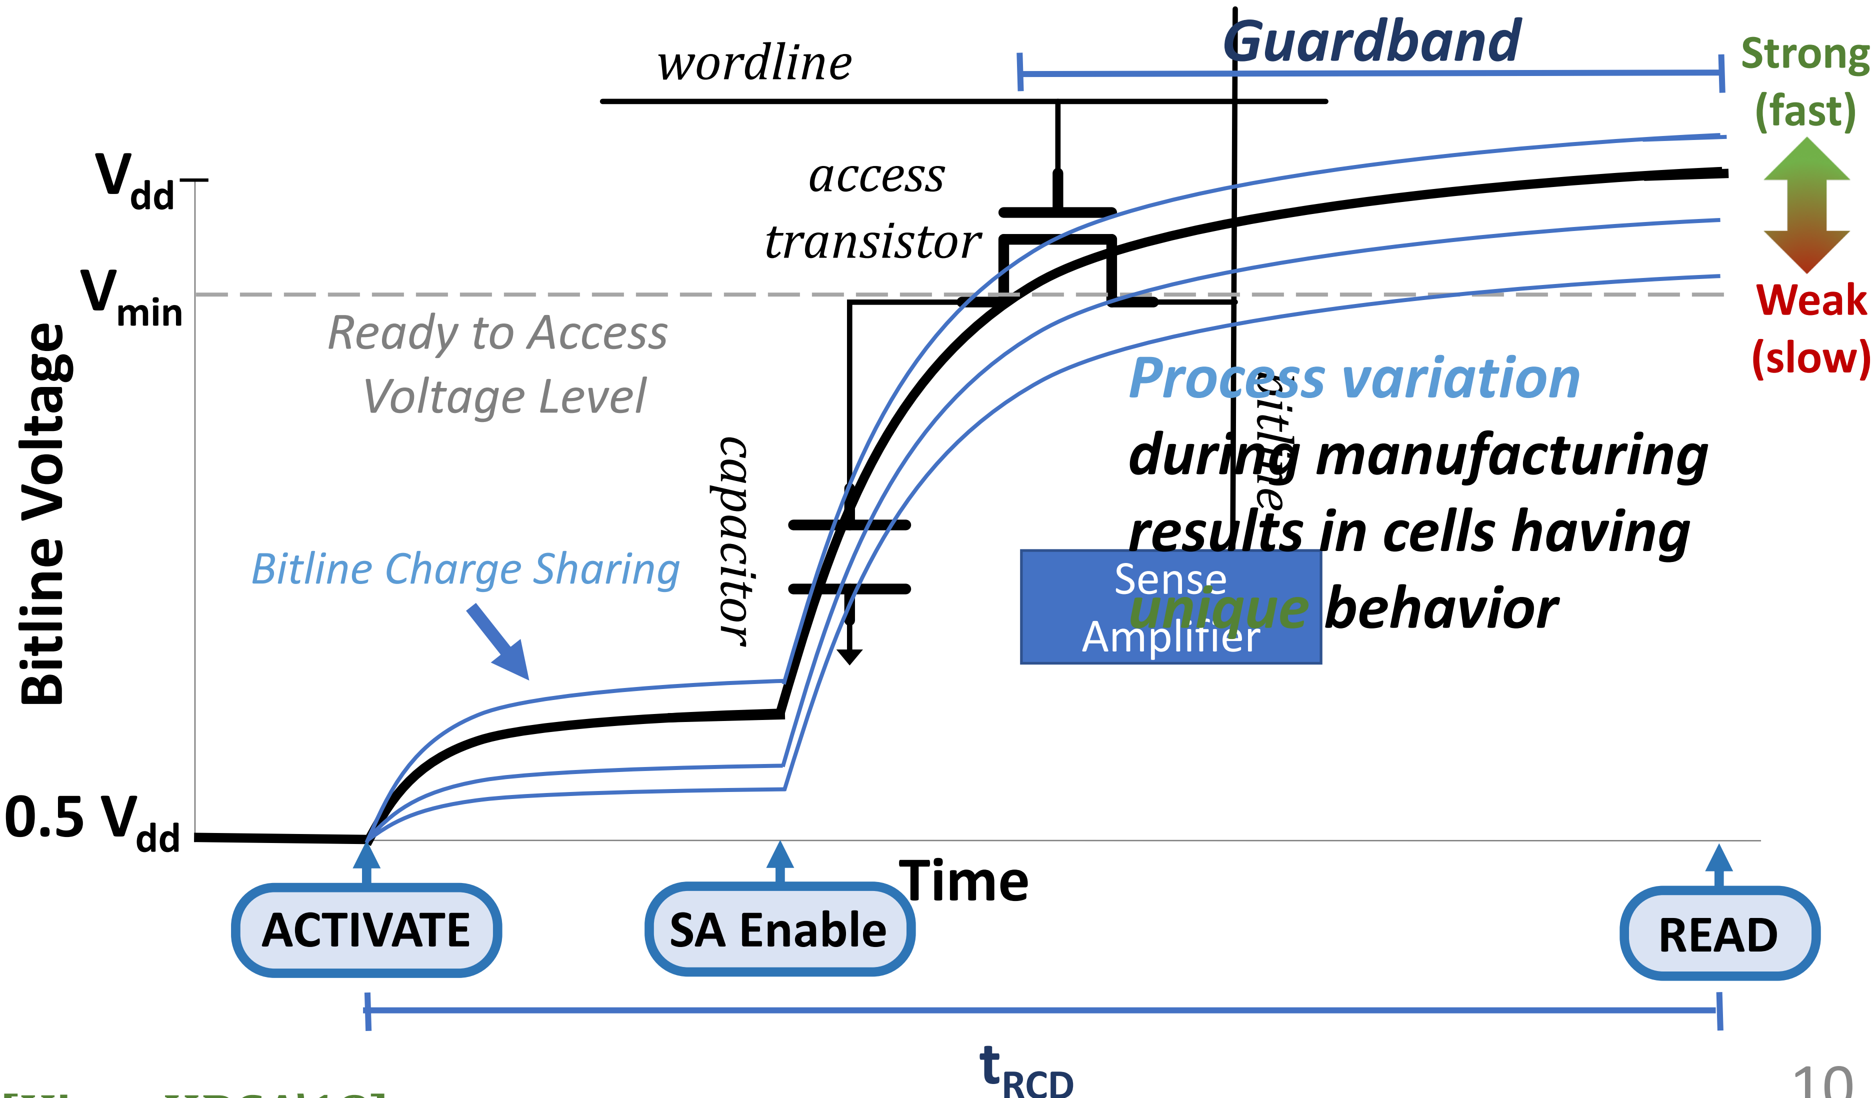

[Kim+ HPCA'18]

DRAM Accesses and Failures

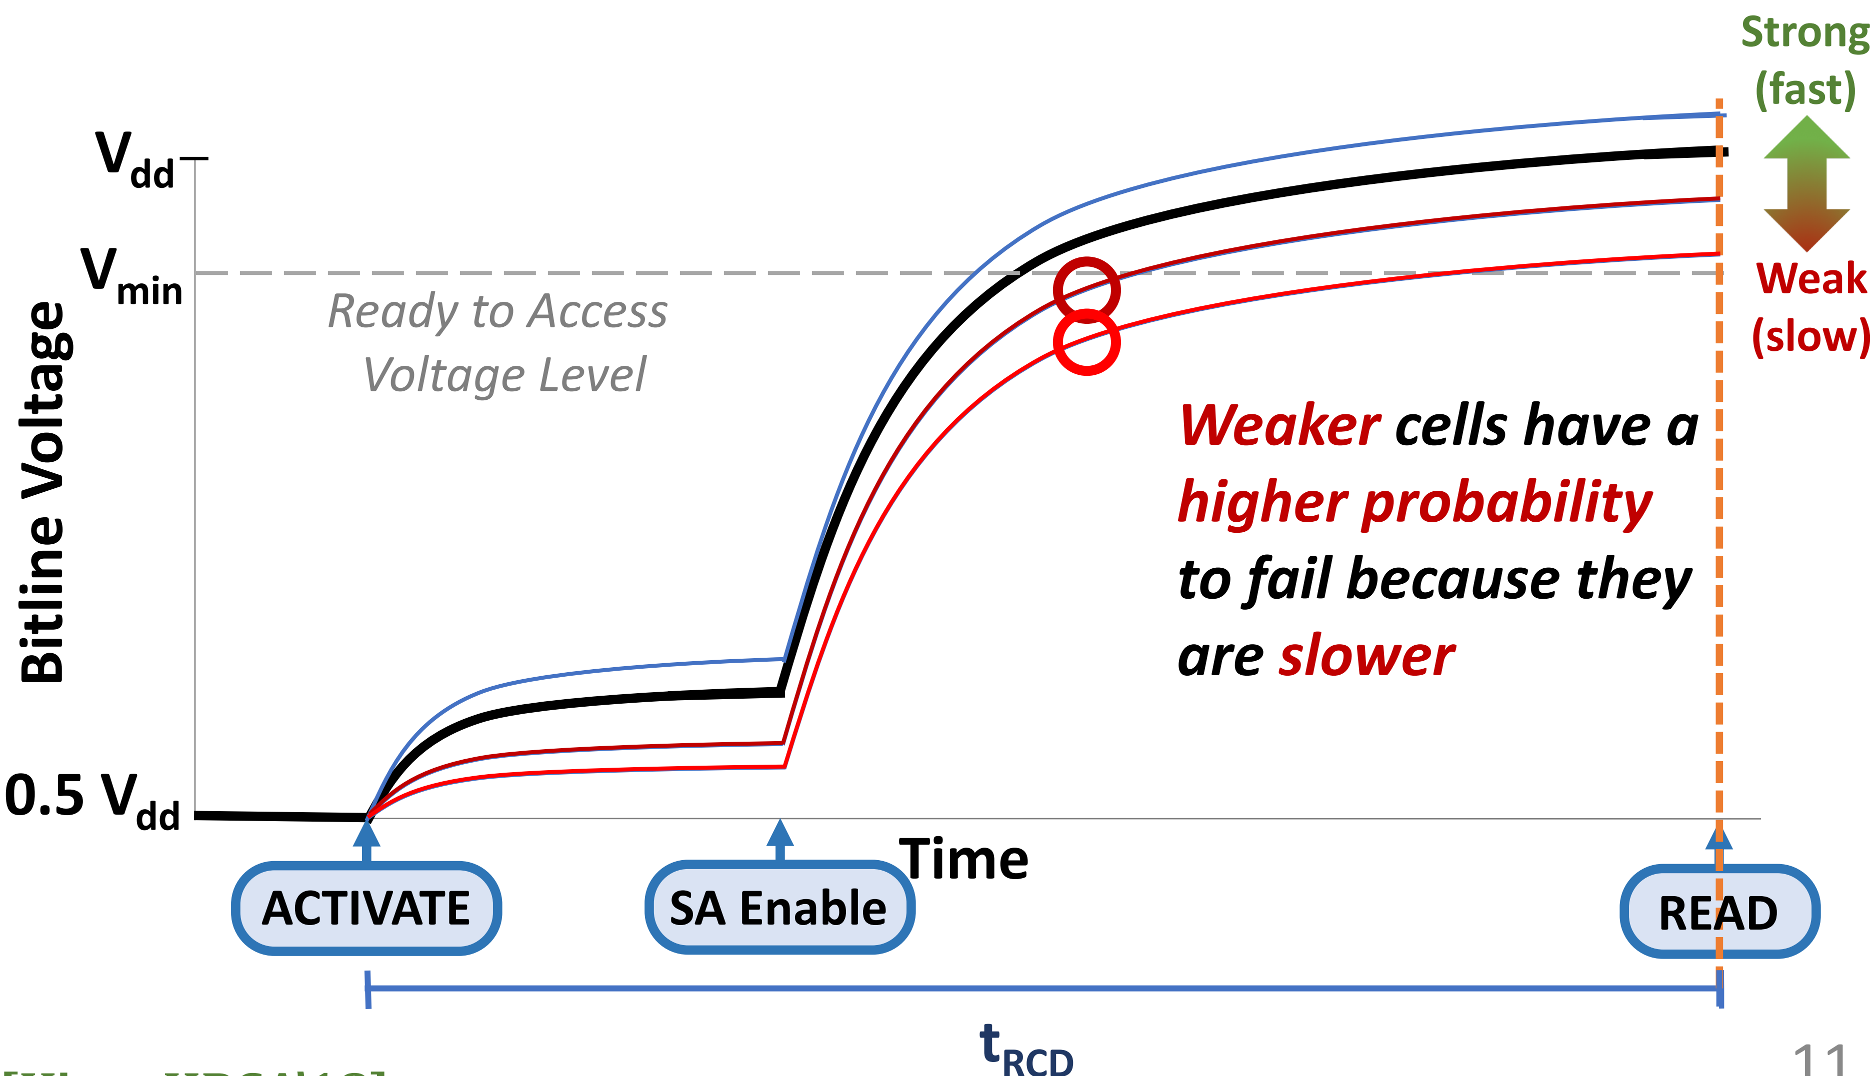

[Kim+ HPCA'18]

Recap of Goals

To identify the opportunity for reliably reducing **t_{RCD},** we want to:

1. **Rigorously characterize** *state-of-the-art* LPDDR4 DRAM
2. **Demonstrate** the **viability of using static profile** of cell strength
3. **Devise** a mechanism to **exploit more activation failure (t_{RCD}) characteristics** and **further reduce DRAM latency**

Solar-DRAM Outline

- Motivation and Goal
- DRAM Background
- **Experimental Methodology**
- Characterization Results
- Mechanism: Solar-DRAM
- Evaluation
- Conclusion

Experimental Methodology

- **282 2y-nm LPDDR4 DRAM modules**
  - **2GB** device size
  - From **3 major DRAM manufacturers**
- **Thermally controlled testing chamber**
  - Ambient temperature range: **{40°C – 55°C} ± 0.25°C**
  - DRAM temperature is held at 15°C above ambient
- **Precise control over DRAM commands and timing parameters**
  - Test reduced latency effects by **reducing t_{RCD} parameter**
- **Ramulator DRAM Simulator** [Kim+, CAL'15]
  - Access latency characterization in **real workloads**

Solar-DRAM Outline

Motivation and Goal

DRAM Background

Experimental Methodology

Characterization Results

Mechanism: Solar-DRAM

Evaluation

Conclusion

Characterization Results

1. **Spatial distribution of activation failures**

2. **Spatial locality of activation failures**

3. **Distribution of cache accesses in real workloads**

4. **Short-term variation of activation failure probability**

5. **Effects of reduced t_{RCD} on write operations**

Spatial Distribution of Failures

How are activation failures spatially distributed in DRAM?

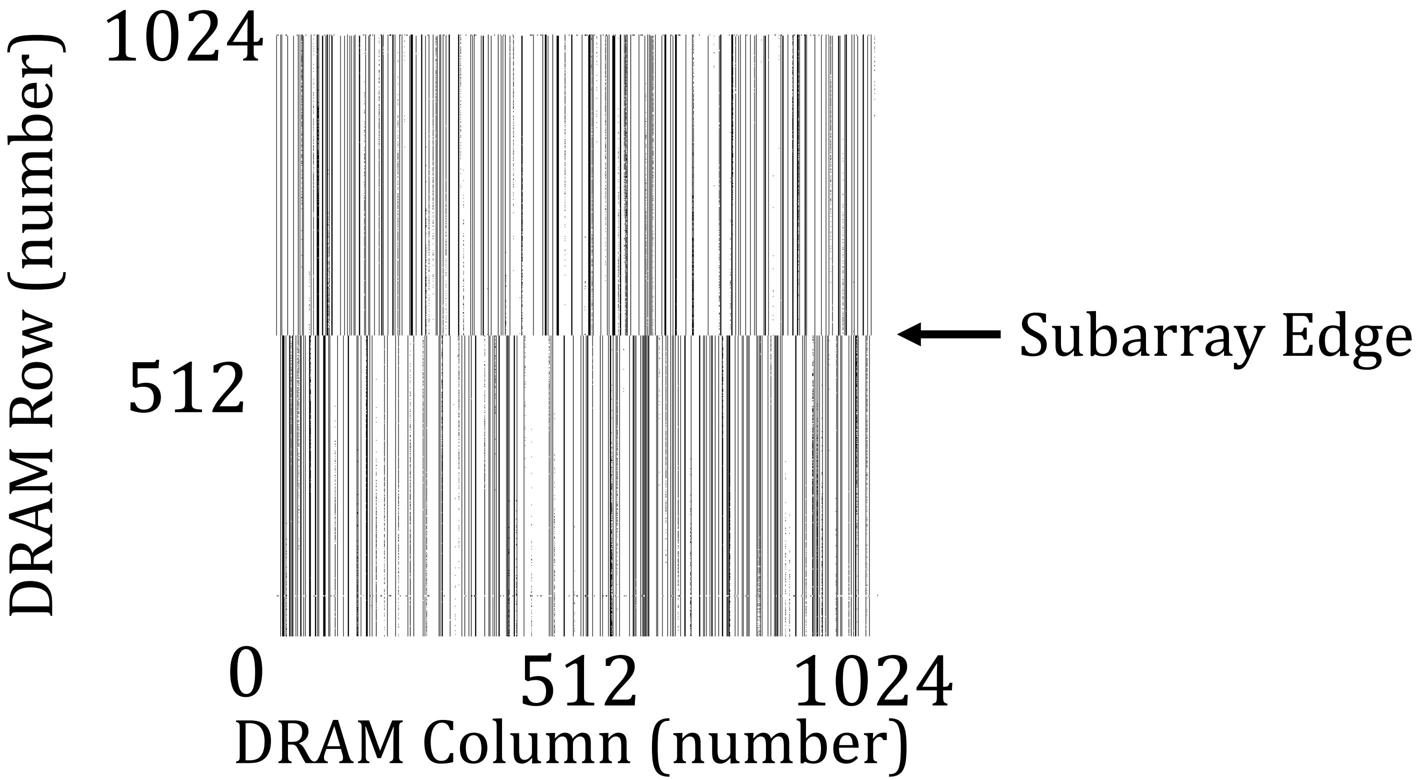

Activation failures are **highly constrained** to local bitlines (i.e., subarrays)

Spatial Locality of Failures

Where does a single access induce activation failures?

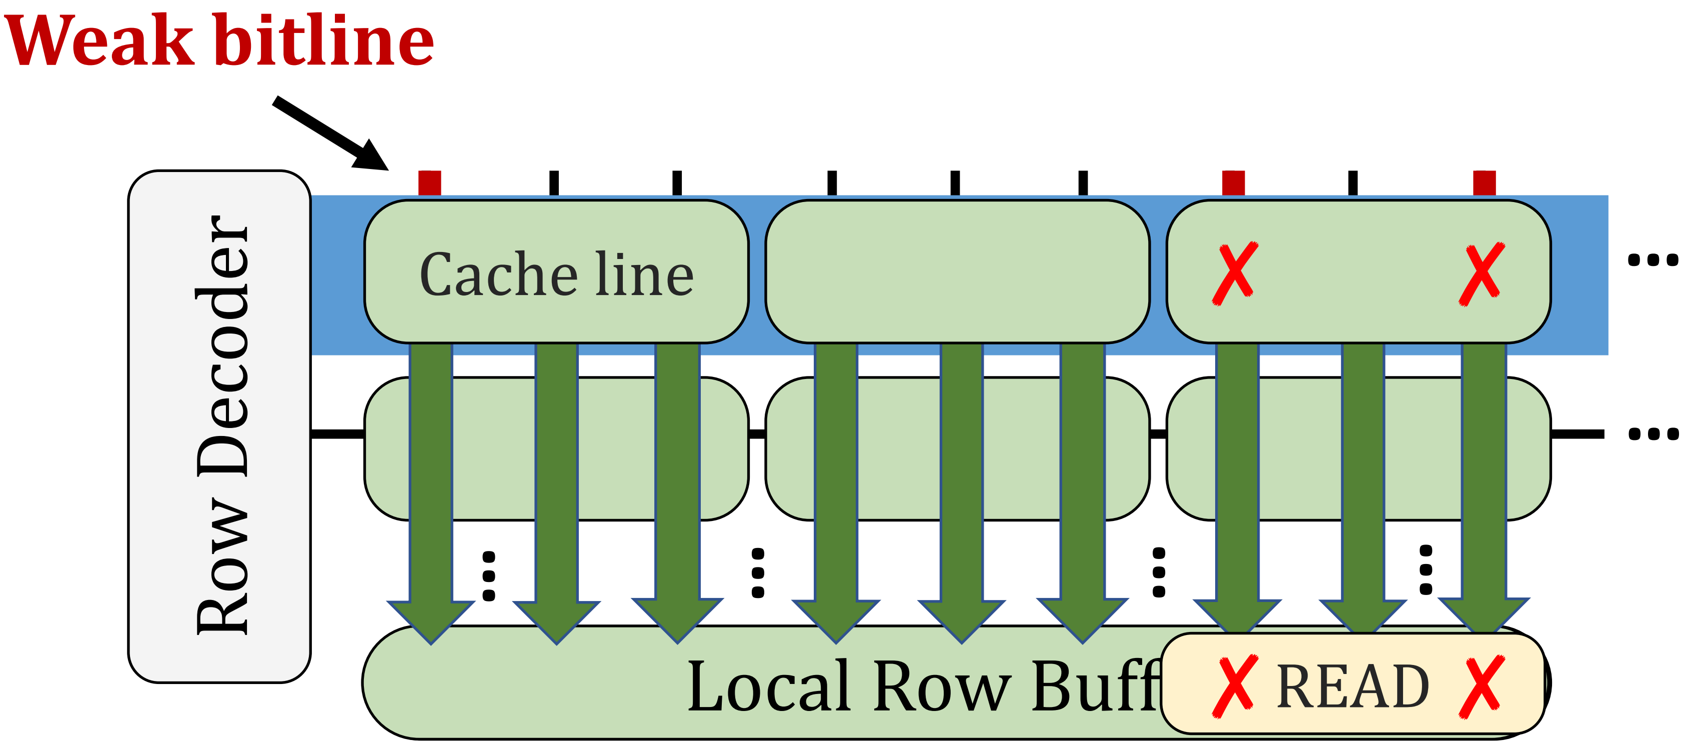

Activation failures are **constrained to the cache line** first accessed immediately following an activation

Spatial Locality of Failures

Where does a single access induce activation failures?

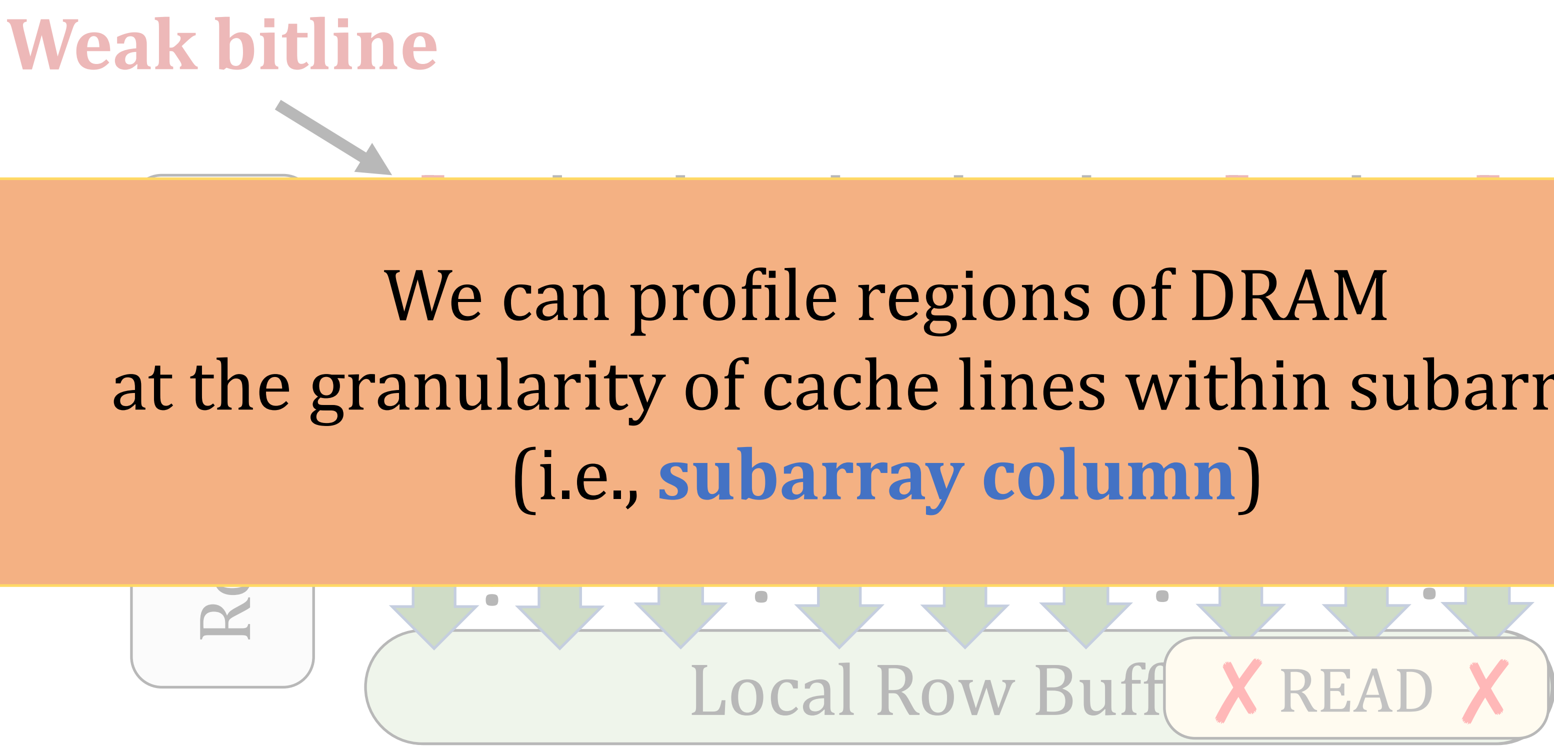

We can profile regions of DRAM
at the granularity of cache lines within subarrays
(i.e., **subarray column**)

Activation failures are **constrained to the cache line**
first accessed immediately following an activation

Distribution of Cache Accesses

Which cache line is most likely to be accessed first immediately following an activation?

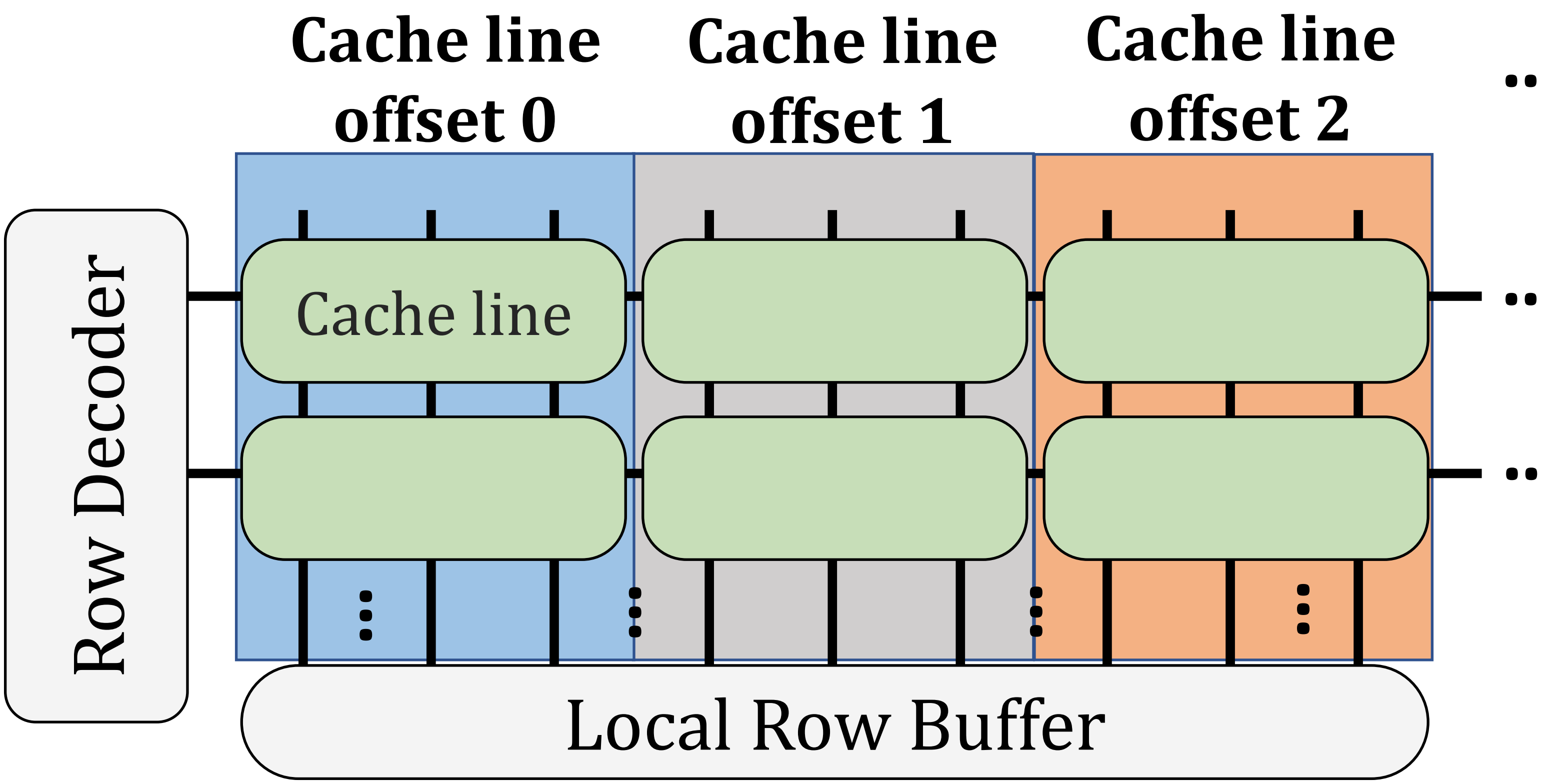

Distribution of Cache Accesses

Which cache line is most likely to be accessed first immediately following an activation?

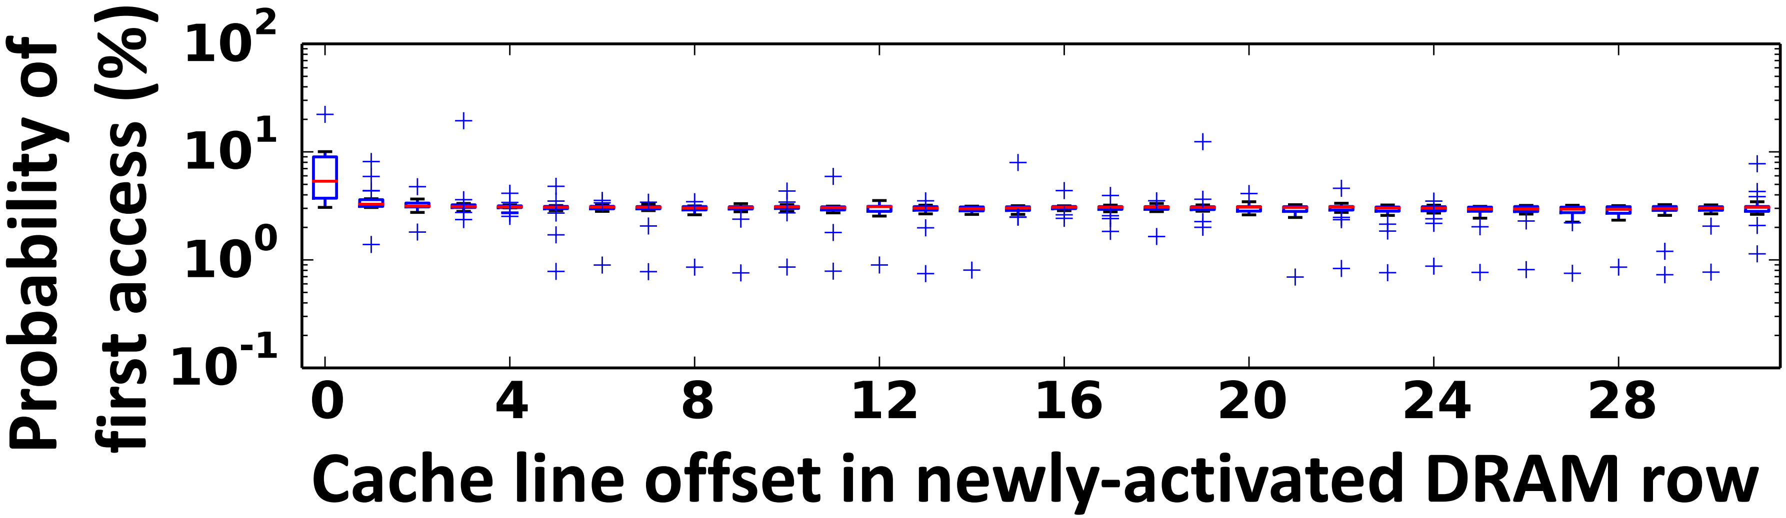

In some applications, up to **22.2%** of first accesses to a newly-activated DRAM row request **cache line 0** in the row

Distribution of Cache Accesses

Which cache line is most likely to be accessed first immediately following an activation?

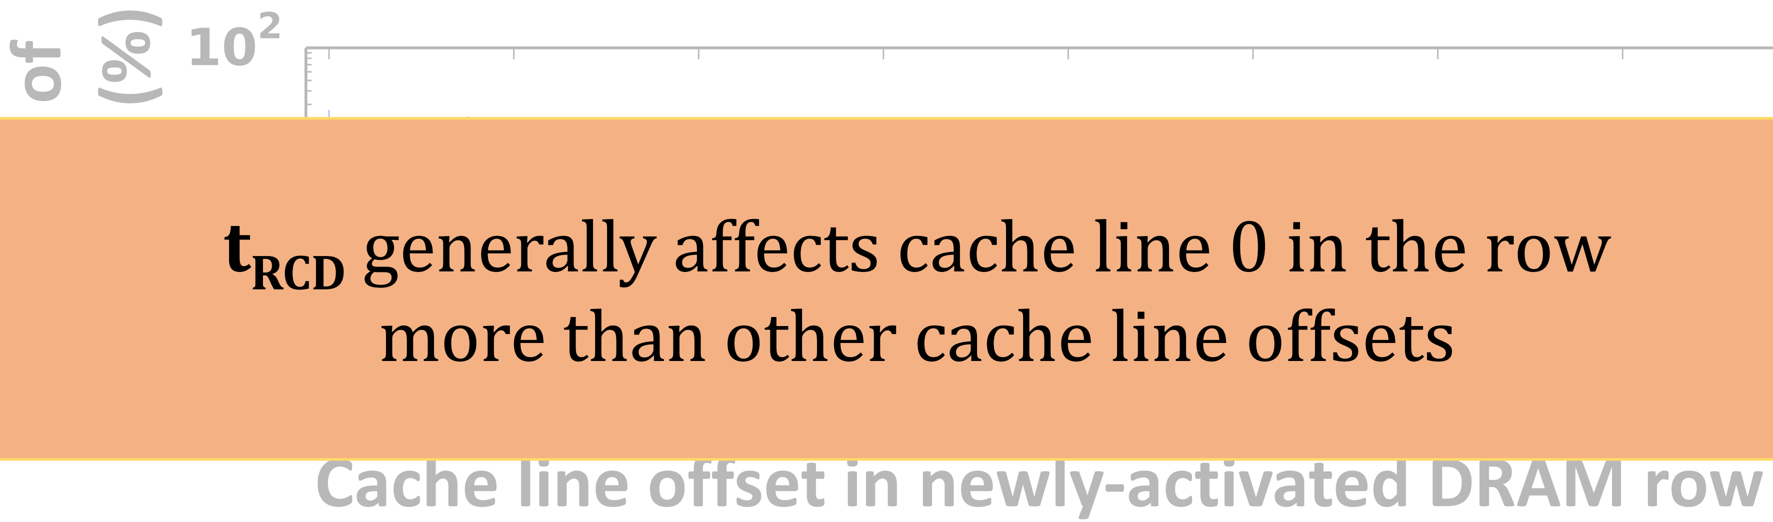

t_{RCD} generally affects cache line 0 in the row more than other cache line offsets

In some applications, up to **22.2%** of first accesses to a newly-activated DRAM row request **cache line 0** in the row

Short-term Variation

Does a bitline's probability of failure (i.e., latency characteristics) change over time?

$$F_{prob} = \sum_{n=1}^{cells\_in\_SA\_bitline} \frac{num\_iters\_failed_{cell_n}}{num\_iters \times cells\_in\_SA\_bitline}$$

cells_in_SA_bitline: number of cells in a local bitline

***num_iters*:** iterations we try to induce failures in each cell

$num\_iters\_failed_{cell_n}$**:** iterations $cell_n$ fails in

We sample F_{prob} many times over a long period and plot how F_{prob} varies across all samples

Short-term Variation

Does a bitline's probability of failure (i.e., latency characteristics) change over time?

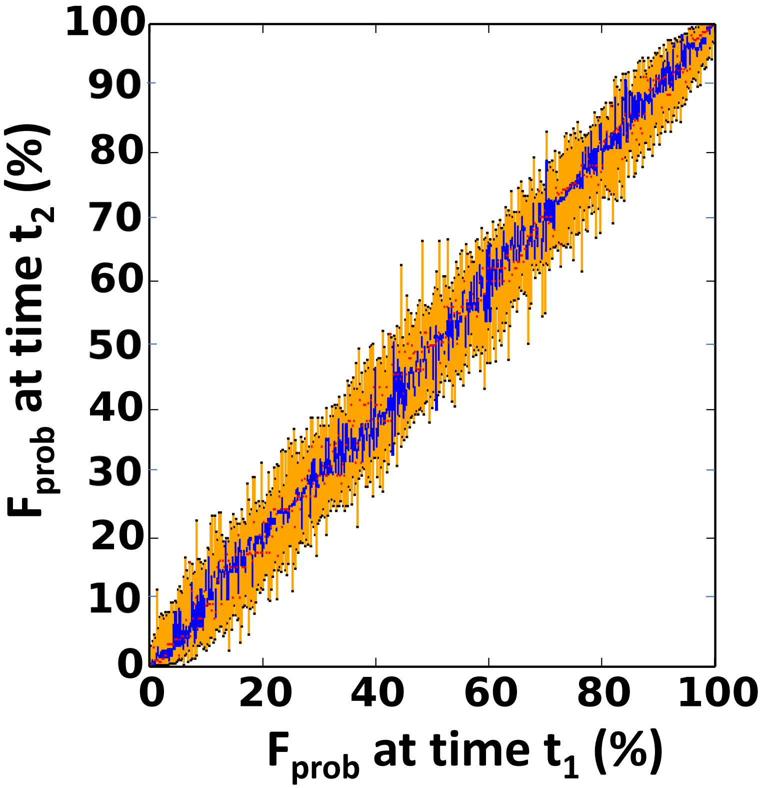

A **weak bitline** is likely to remain **weak** and a **strong bitline** is likely to remain **strong** over time

Short-term Variation

Does a bitline's probability of failure (i.e., latency characteristics) change over time?

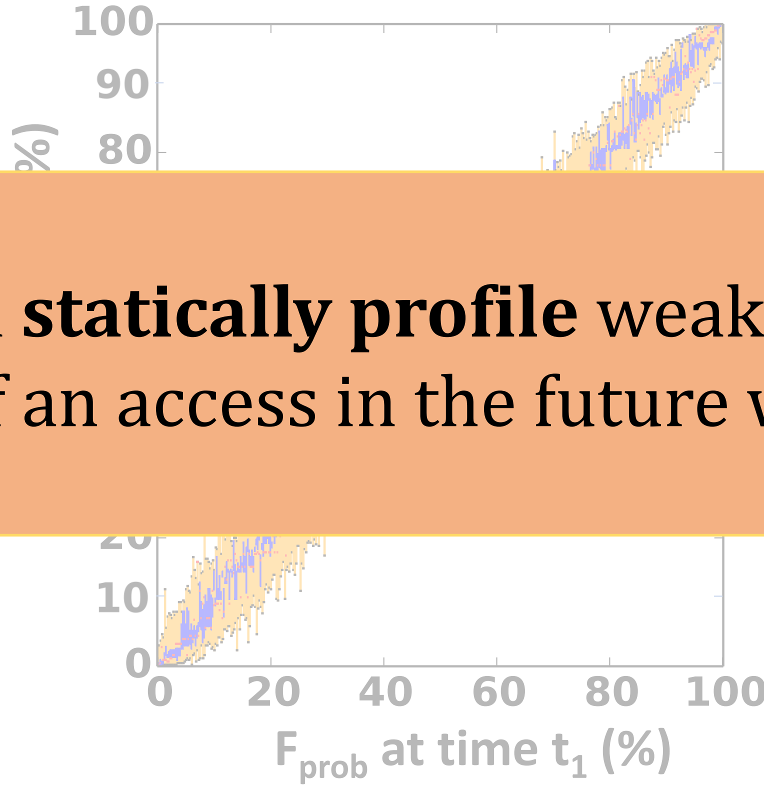

We can **statically profile** weak bitlines and determine if an access in the future will cause failures

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 $\mathbf{t_{RCD}}$?

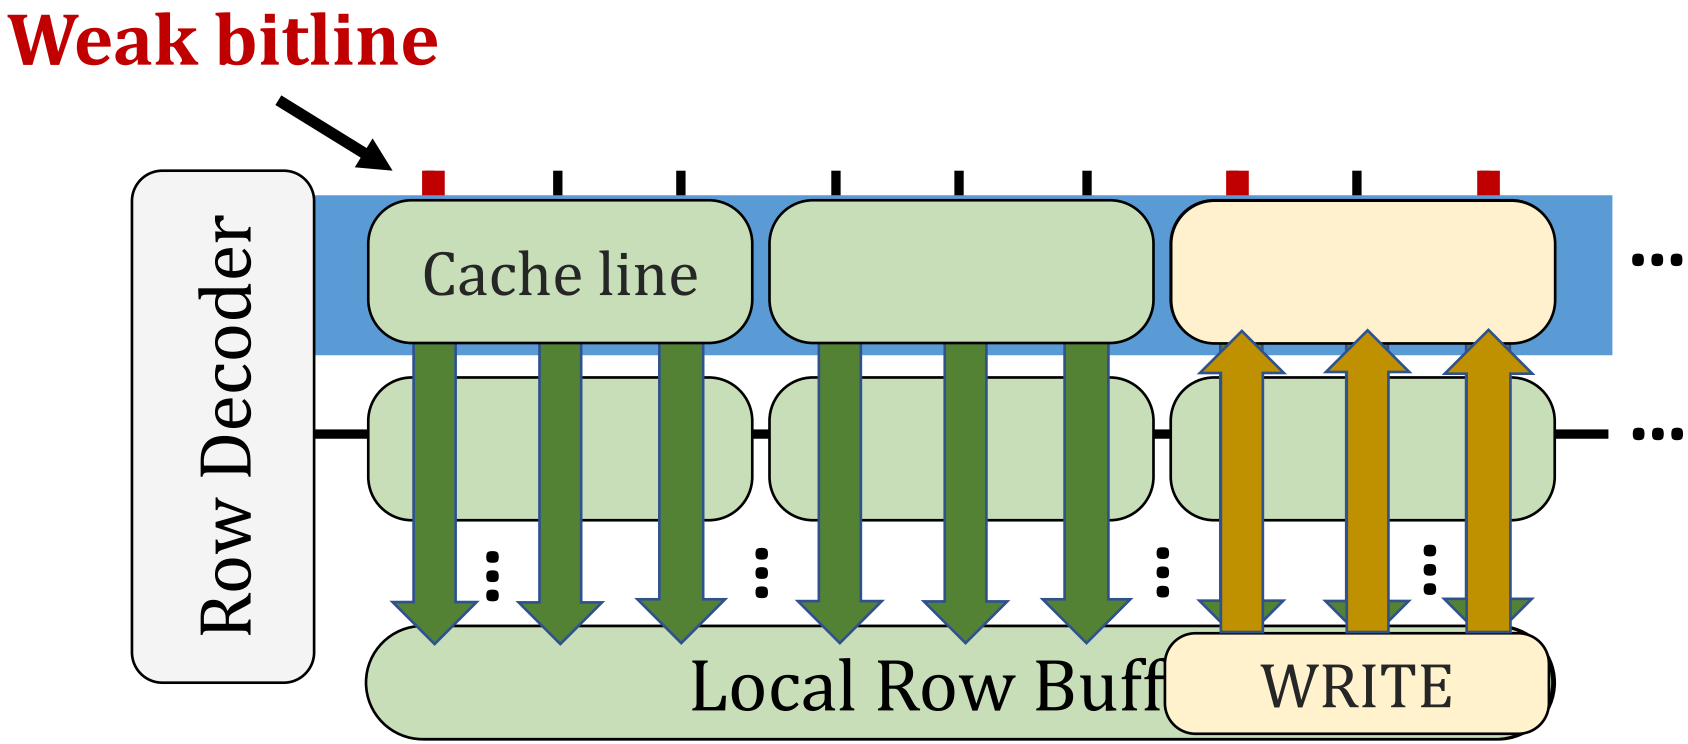

We can reliably issue write operations
with significantly reduced $\mathbf{t_{RCD}}$ (e.g., by 77%)

Solar-DRAM Outline

- Motivation and Goal
- DRAM Background
- Experimental Methodology
- Characterization Results
- **Mechanism: Solar-DRAM**
- Evaluation
- Conclusion

Solar-DRAM

Identifies subarray columns as "**weak (slow)**" or "**strong (fast)**" and accesses cache lines in strong subarray columns with reduced $\mathbf{t_{RCD}}$

Uses a static profile of weak subarray columns

- 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

Identifies subarray columns as "**weak (slow)**" or "**strong (fast)**" and accesses cache lines in strong subarray columns with reduced $\mathbf{t_{RCD}}$

Uses a static profile of weak subarray columns

- 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)

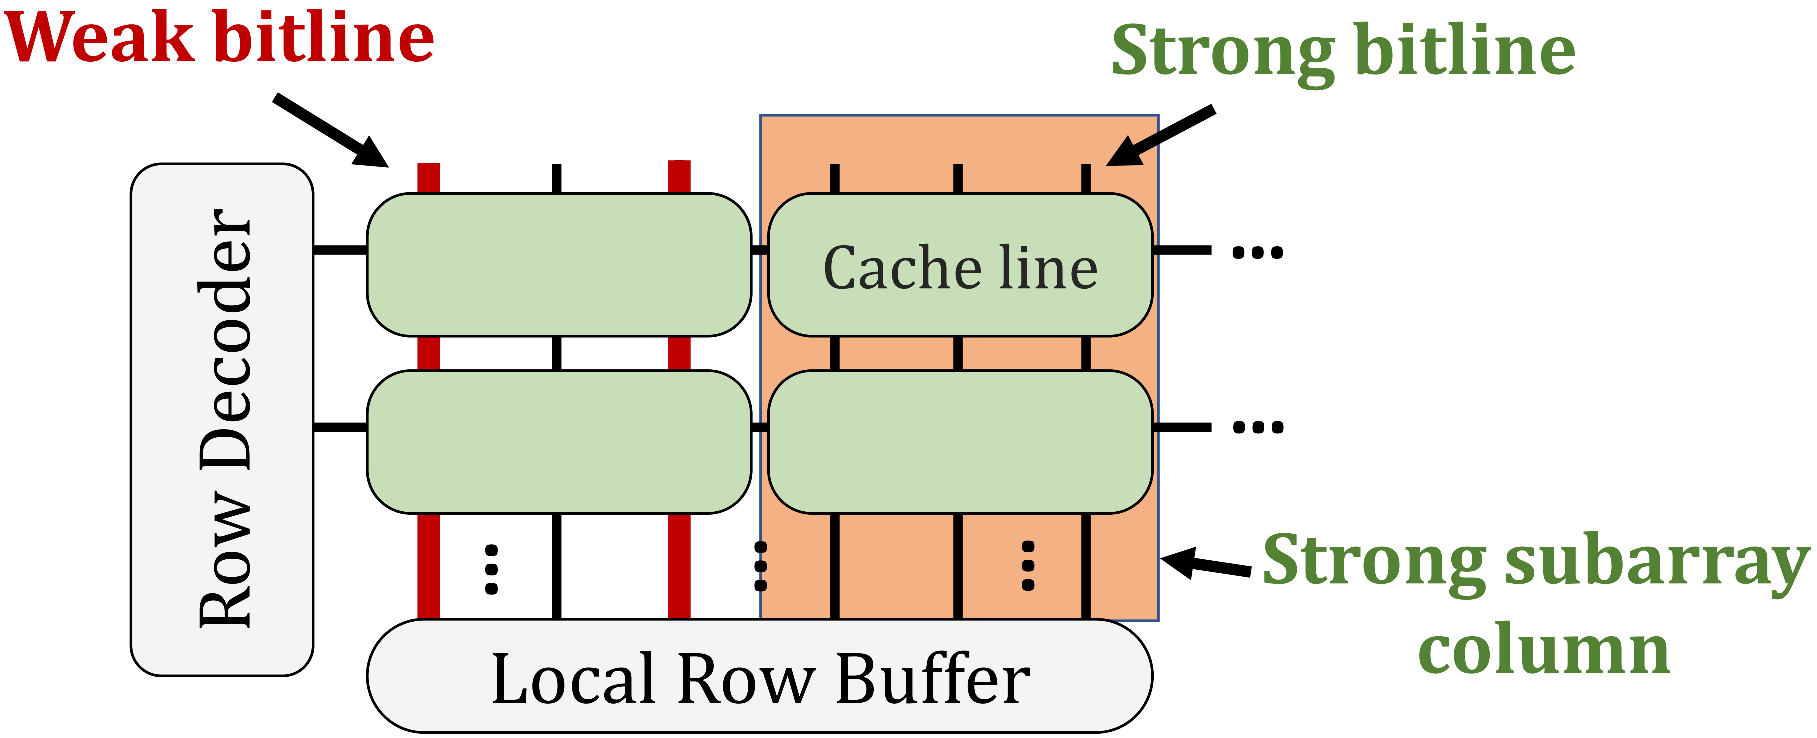

Identifies subarray columns comprised of **strong bitlines**

Access cache lines in strong subarray columns with a **reduced t_{RCD}**

Solar-DRAM

Identifies subarray columns as "**weak (slow)**" or "**strong (fast)**" and accesses cache lines in strong subarray columns with reduced $\mathbf{t_{RCD}}$

Uses a static profile of weak subarray columns

- 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)

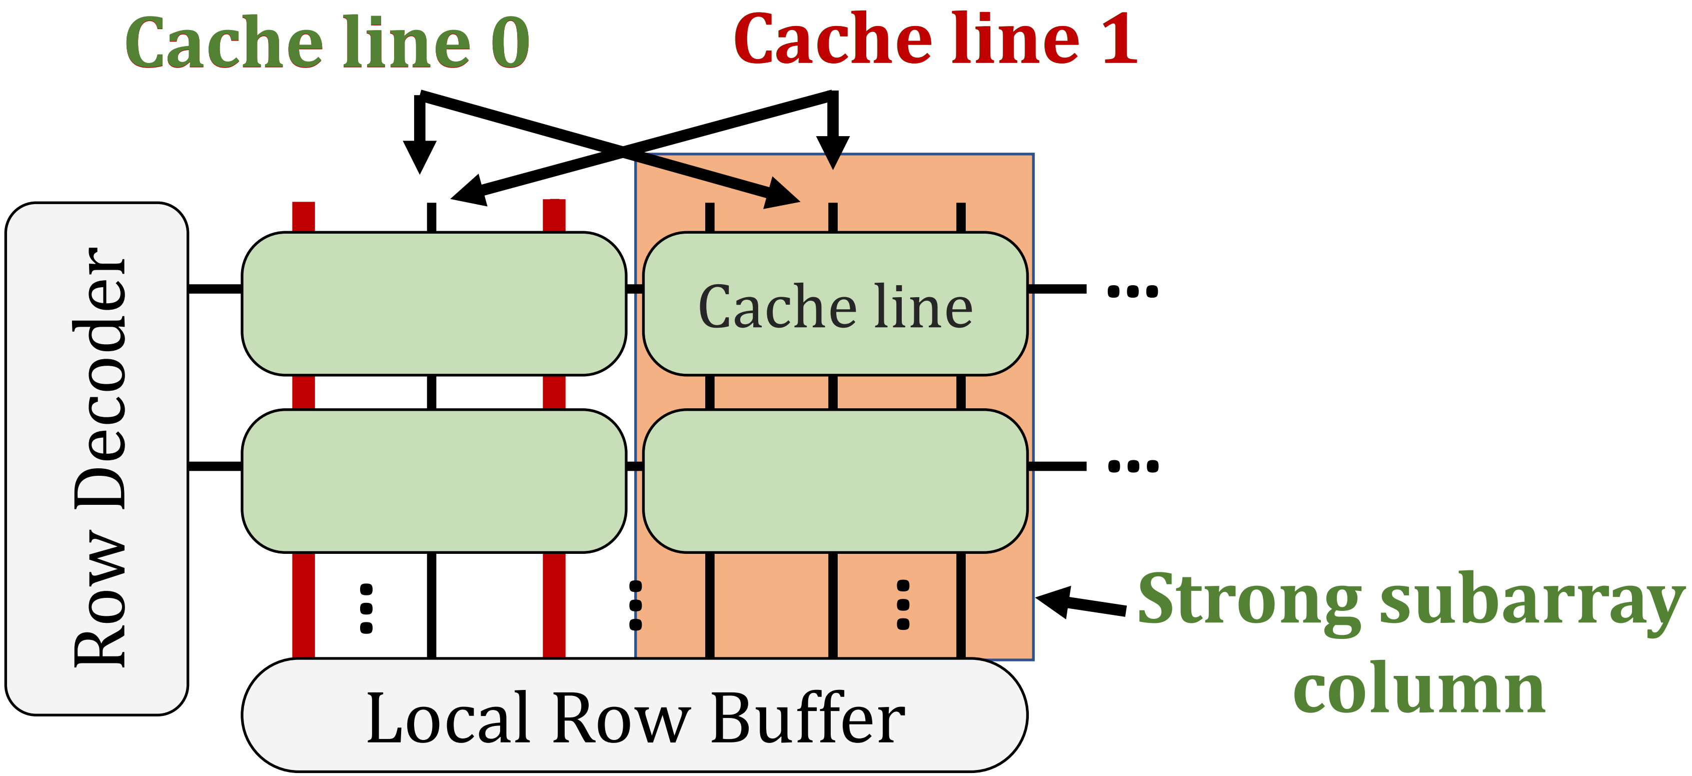

Remap cache lines across DRAM at the memory controller level so cache line 0 will likely map to a **strong** cache line

Solar-DRAM

Identifies subarray columns as "**weak (slow)**" or "**strong (fast)**" and accesses cache lines in strong subarray columns with reduced $\mathbf{t_{RCD}}$

Uses a static profile of weak subarray columns

- 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: RLW (III)

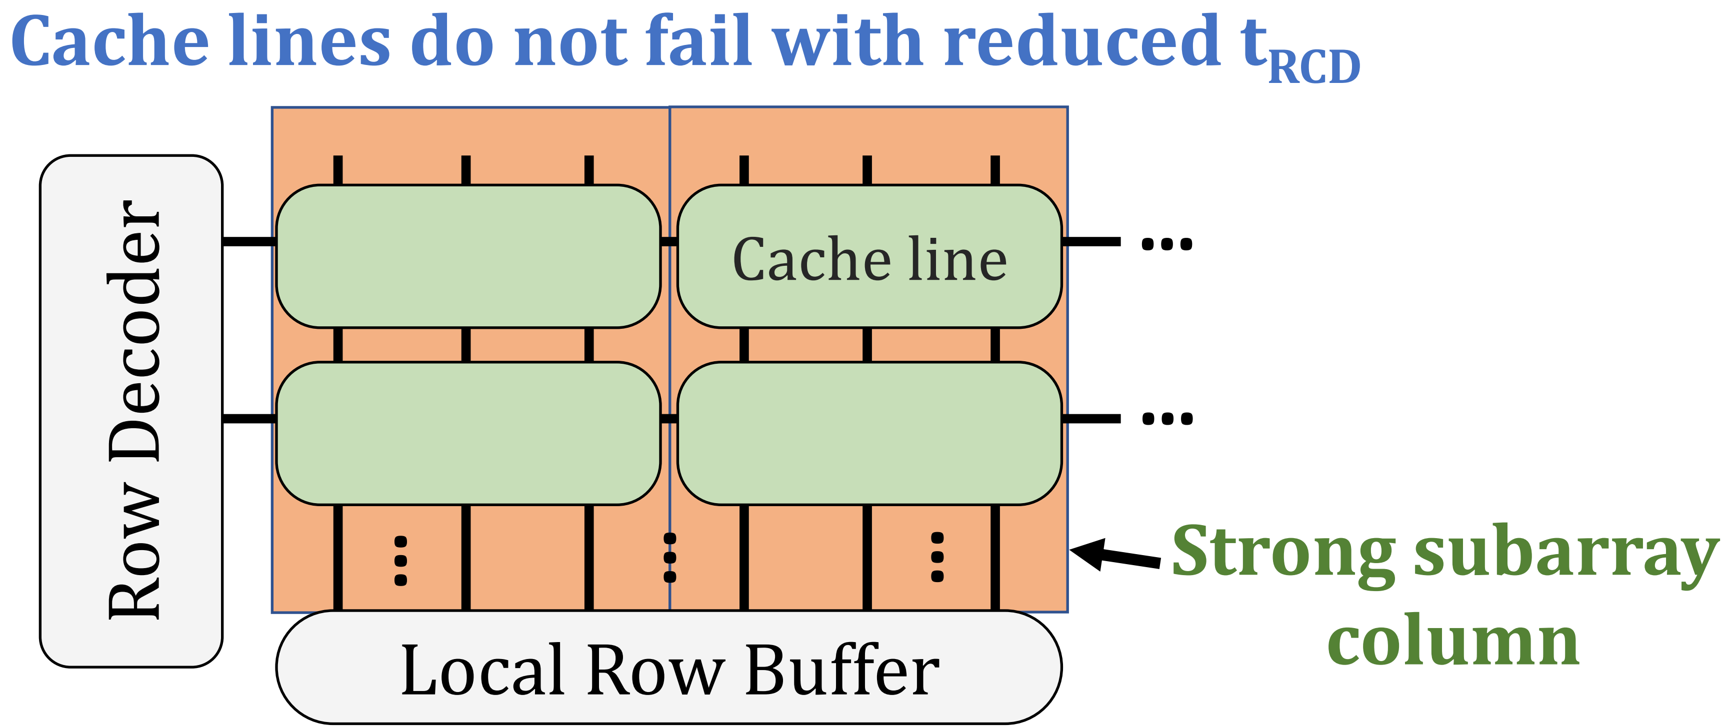

Write to all locations in DRAM with a significantly reduced $\mathbf{t_{RCD}}$ (e.g., by 77%)

Solar-DRAM: Putting it all Together

Each component increases the number of accesses that can be issued with a **reduced t_{RCD}**

They **combine** to further increase the number of cases where **t_{RCD} can be reduced**

Solar-DRAM utilizes each component (**VLC**, **RSC**, and **RLW**) in concert to reduce DRAM latency and **significantly improve system performance**

Solar-DRAM Outline

Evaluation Methodology

- **Cycle-level simulator**: Ramulator [Kim+, CAL'15]
  https://github.com/CMU-SAFARI/ramulator
- **4-core** system with LPDDR4-3200 memory
- **Benchmarks**: SPEC2006
  - 40 8-core workloads
- **Performance metric**: Weighted Speedup (WS)

Evaluation: Homogeneous workloads

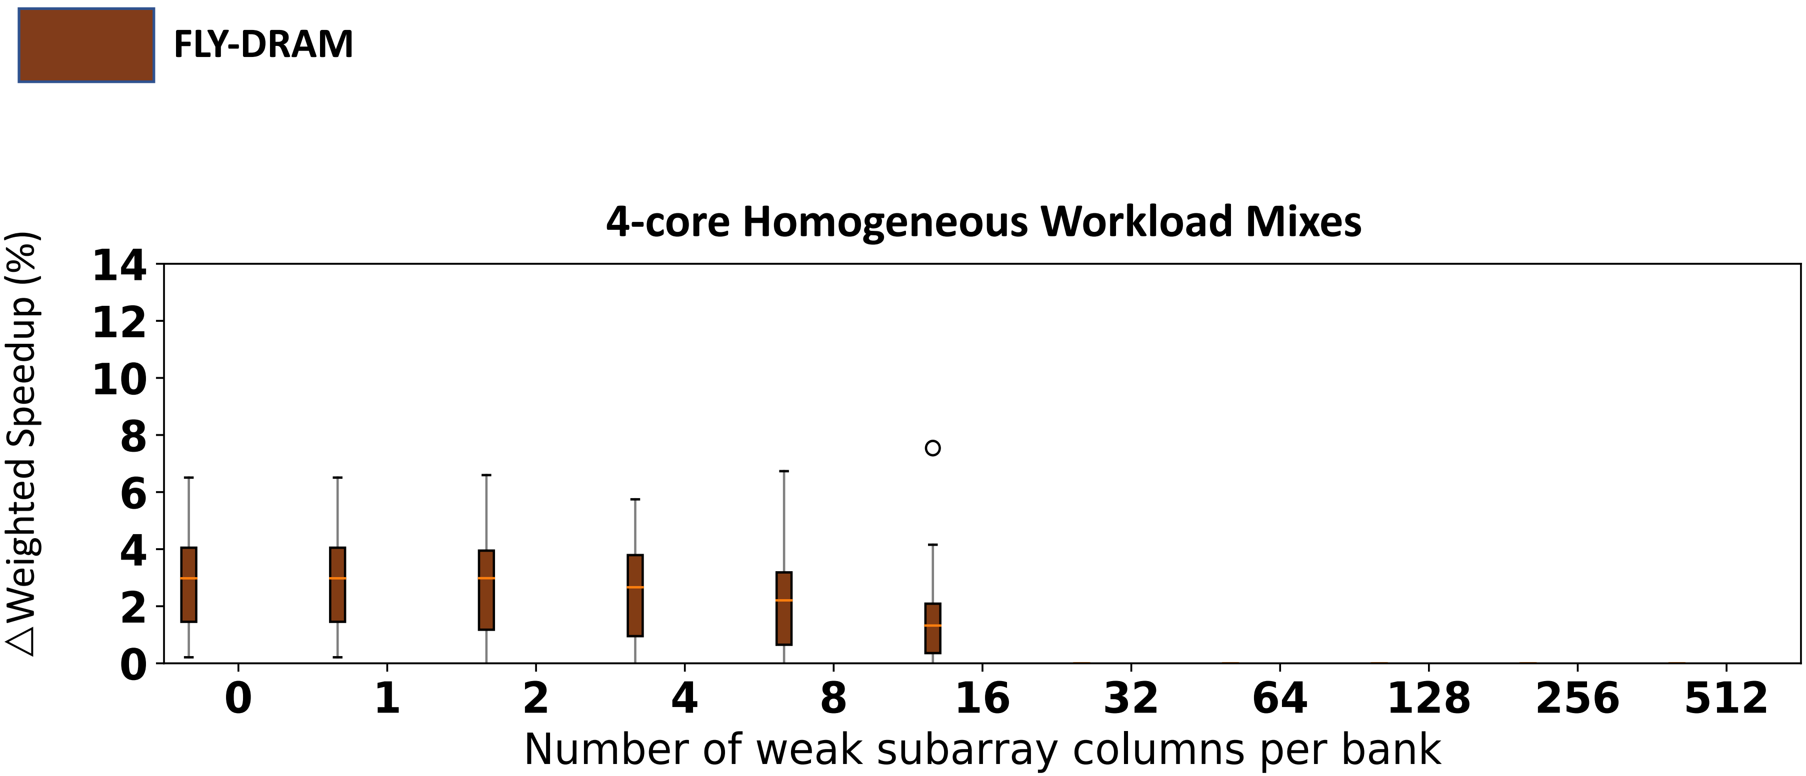

Evaluation: Homogeneous workloads

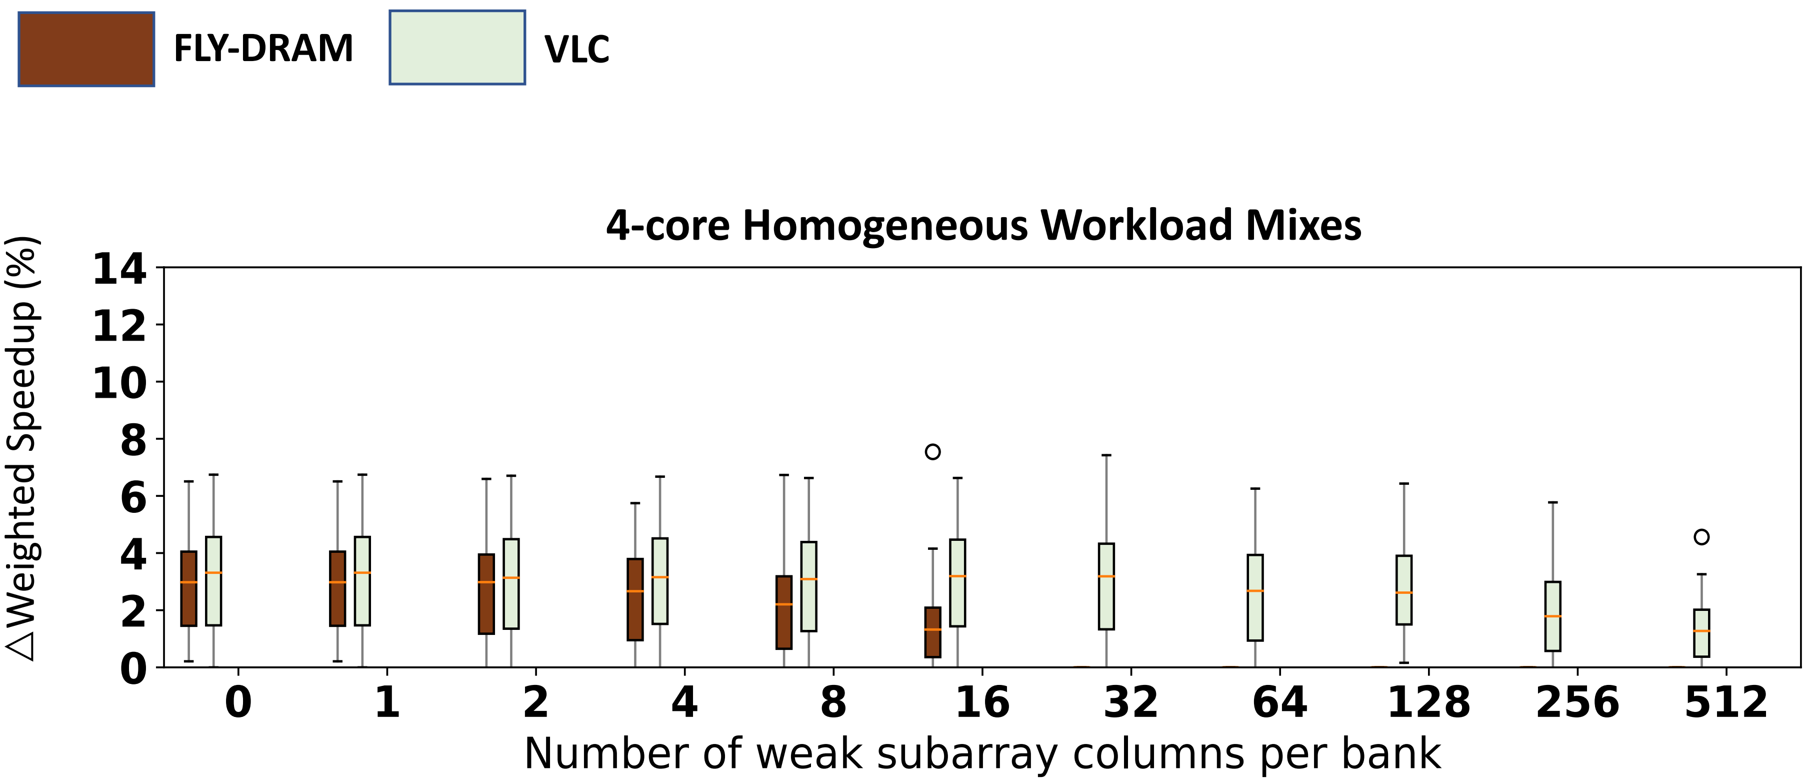

Evaluation: Homogeneous workloads

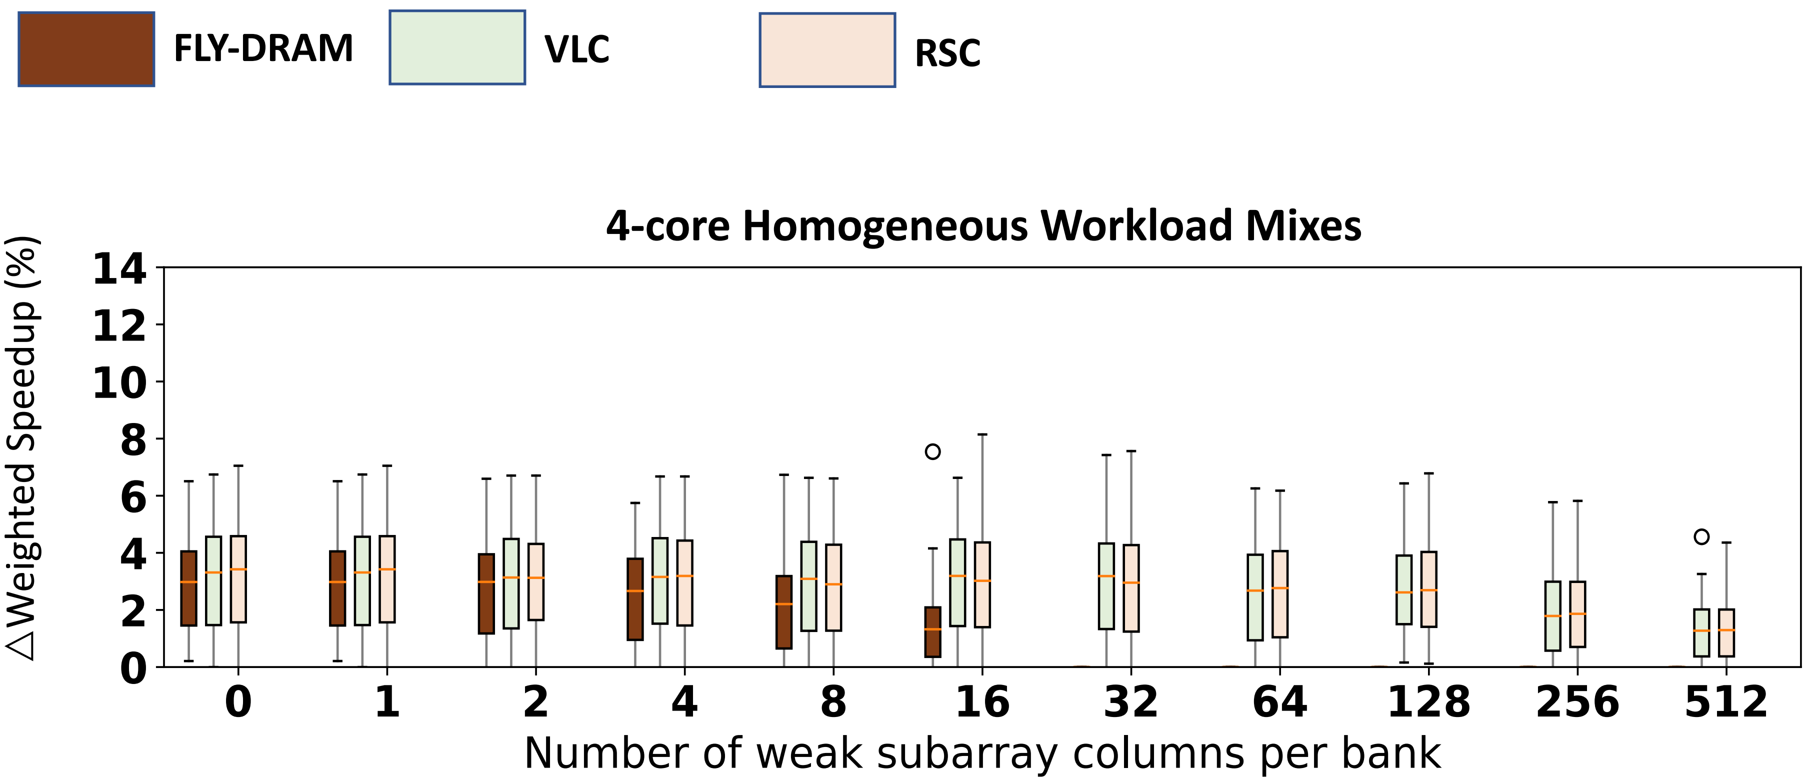

Evaluation: Homogeneous workloads

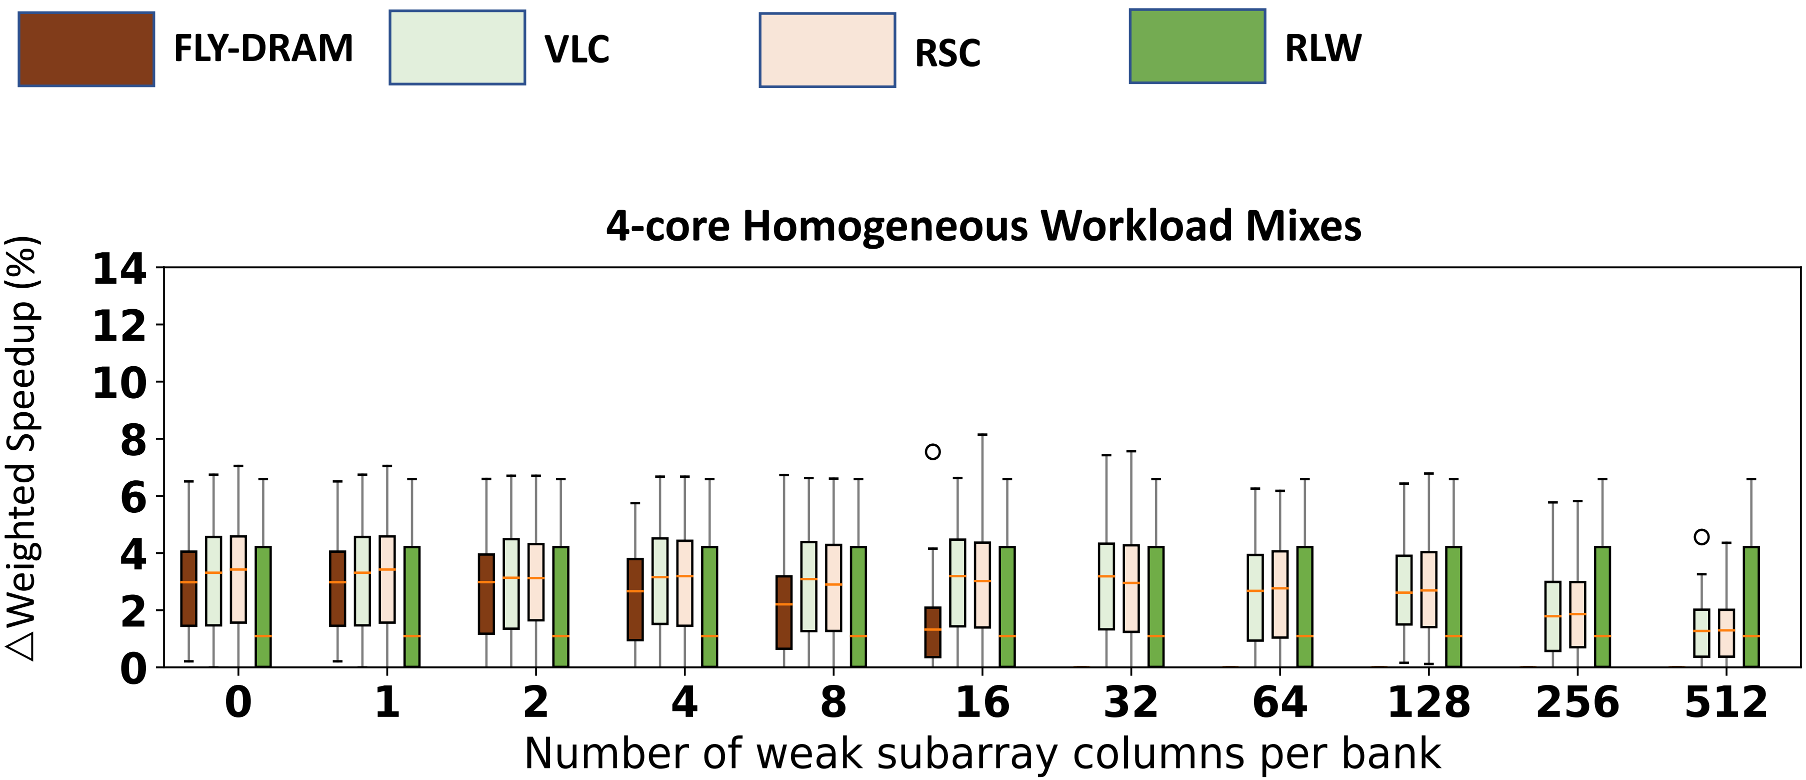

Evaluation: Homogeneous workloads

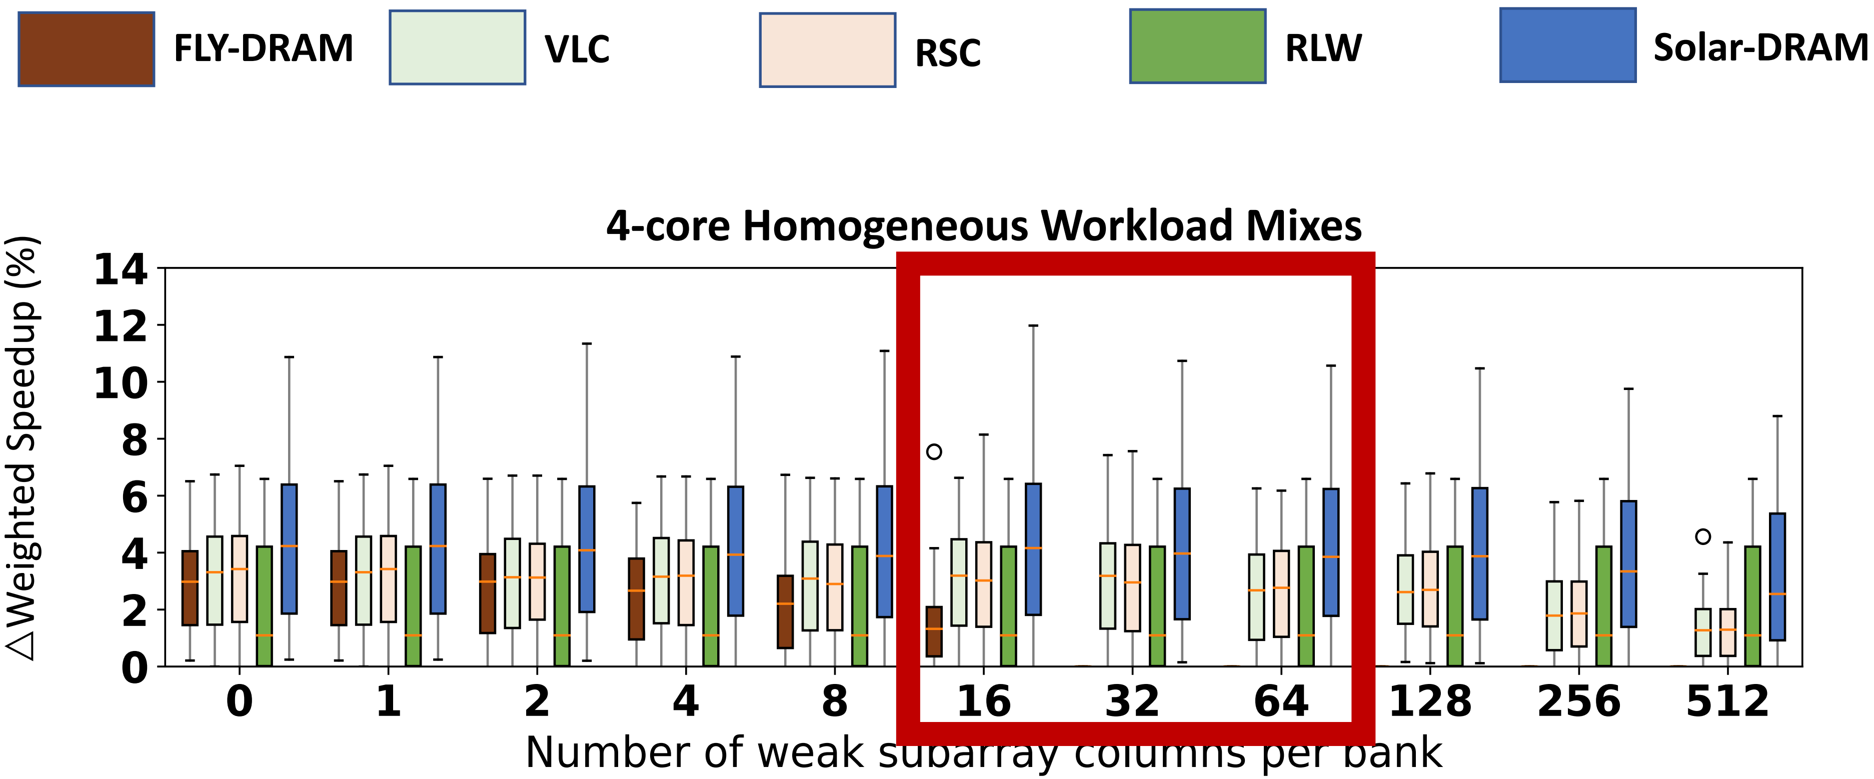

Solar-DRAM reduces $\mathbf{t_{RCD}}$ for more DRAM accesses and provides **10.87%** performance benefit

Other Results in the Paper

- **A detailed analysis on:**
  - Devices of **the three major DRAM manufacturers**
  - **Data Pattern Dependence** of activation failures
    - Random data pattern finds the highest coverage of weak bitlines
  - **Temperature effects** on activation failure probability
    - F_{prob} generally increases with higher temperatures
  - Evaluation with Heterogeneous workloads
    - Solar-DRAM provides **8.79%** performance benefit

- **Further discussion on:**
  - Implementation details
  - Finding a **comprehensive profile** of weak subarray columns

Solar-DRAM Outline

- Motivation and Goal
- DRAM Background
- Experimental Methodology
- Characterization Results
- Mechanism: Solar-DRAM
- Evaluation
- **Conclusion**

Executive Summary

Motivation: DRAM latency is a **major performance bottleneck**

Problem: Many important workloads exhibit **bank conflicts** in DRAM, which result in even longer latencies

Goal:

1. Rigorously **characterize access latency** on LPDDR4 DRAM
2. Exploit findings to **robustly reduce DRAM access latency**

Solar-DRAM:

- Categorizes local bitlines as "*weak (slow)*" or "*strong (fast)*"
- Robustly **reduces DRAM access latency for reads *and* writes** to data contained in "*strong*" local bitlines.

Evaluation:

1. Experimentally characterize **282** real LPDDR4 DRAM chips
2. In simulation, **Solar-DRAM** provides **10.87%** system performance improvement over LPDDR4 DRAM

Solar-DRAM:

Reducing DRAM Access Latency by Exploiting the Variation in Local Bitlines

Jeremie S. Kim **Minesh Patel**

Hasan Hassan **Onur Mutlu**

Systems@ETH zürich

SAFARI

ETH zürich

Carnegie Mellon

Evaluation: Heterogeneous workloads

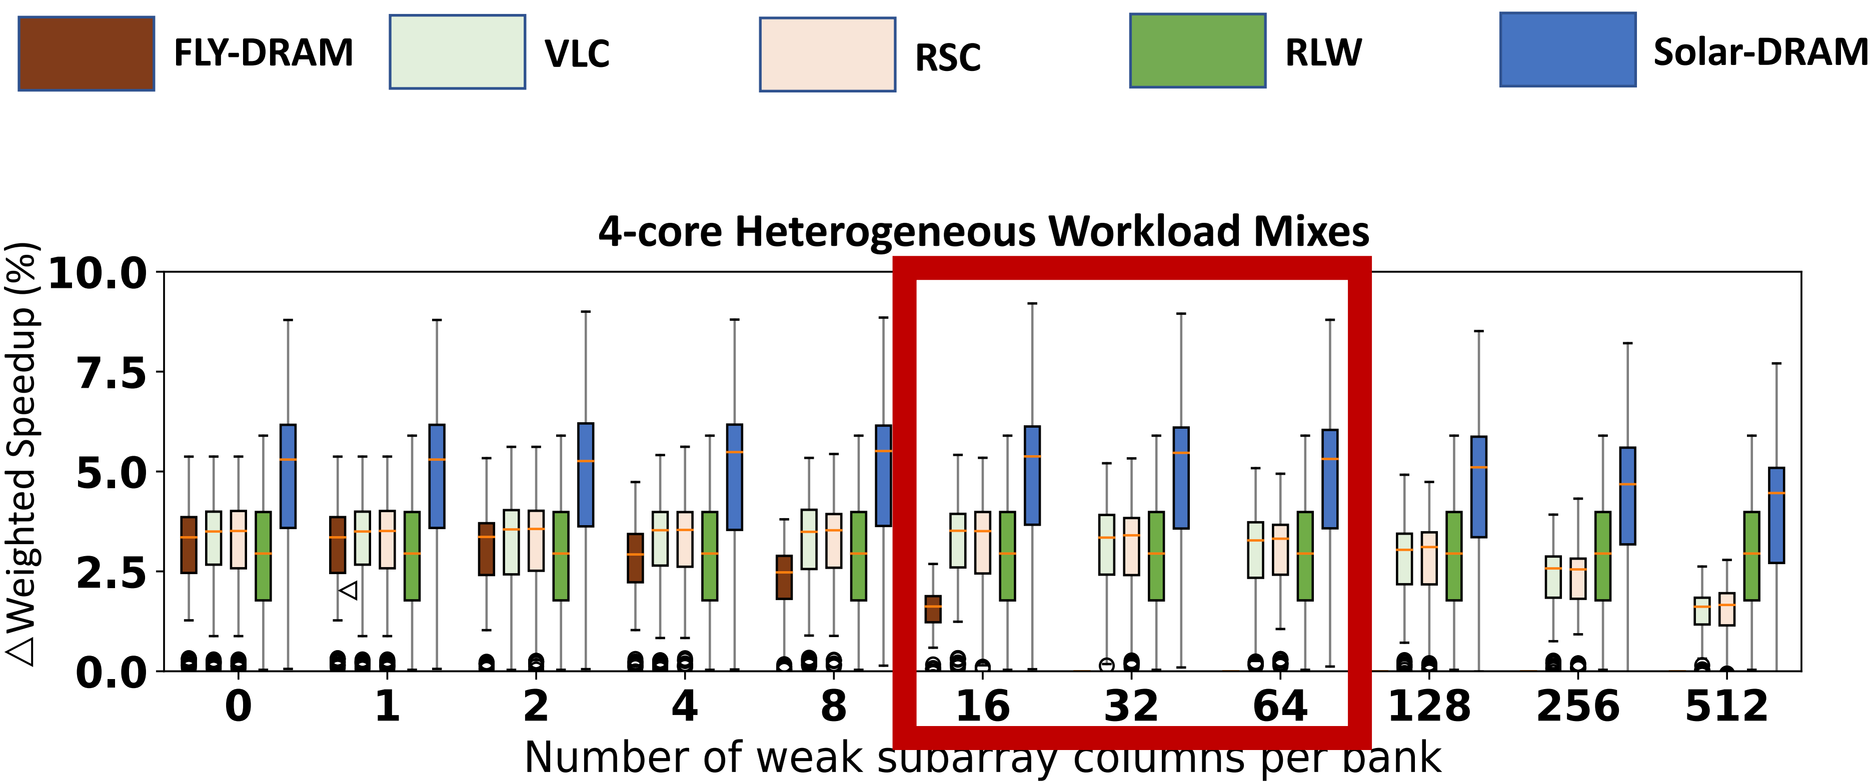

Solar-DRAM reduces $\mathbf{t_{RCD}}$ for more DRAM accesses and provides **8.79%** performance benefit

Temperature

We study the effects of changing temperature on Fprob. The x-axis shows the Fprob at a given temperature T, and the y-axis plots the distribution (box and whiskers plot) of Fprob at a higher temperature for the same bitline

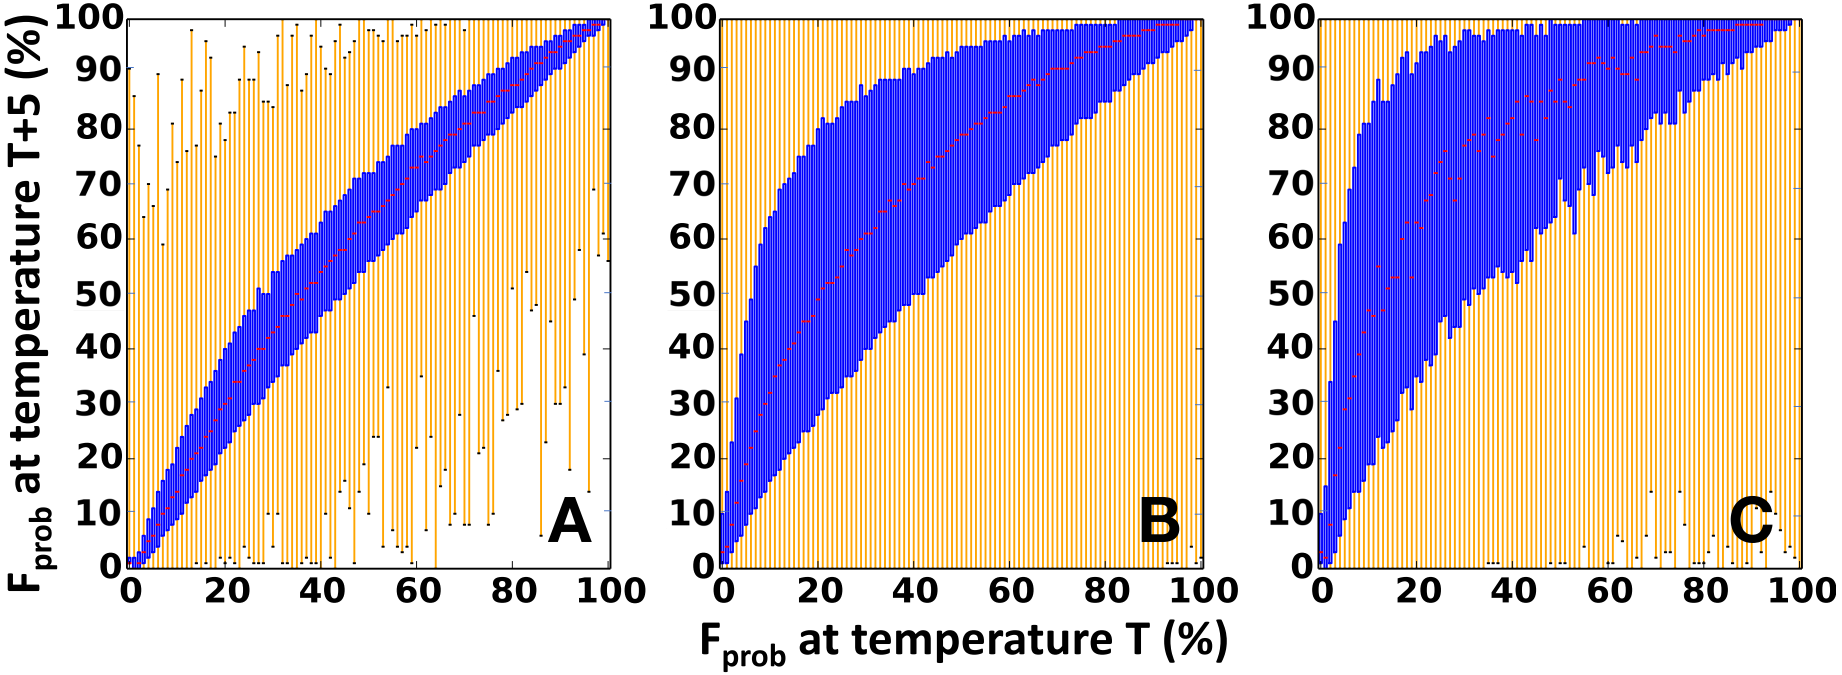

Since a majority of the data points are above the x=y line, Fprob generally increases with higher temperatures

Data Pattern Dependence

We study how using different data patterns affects the number of weak bitlines found over multiple iterations

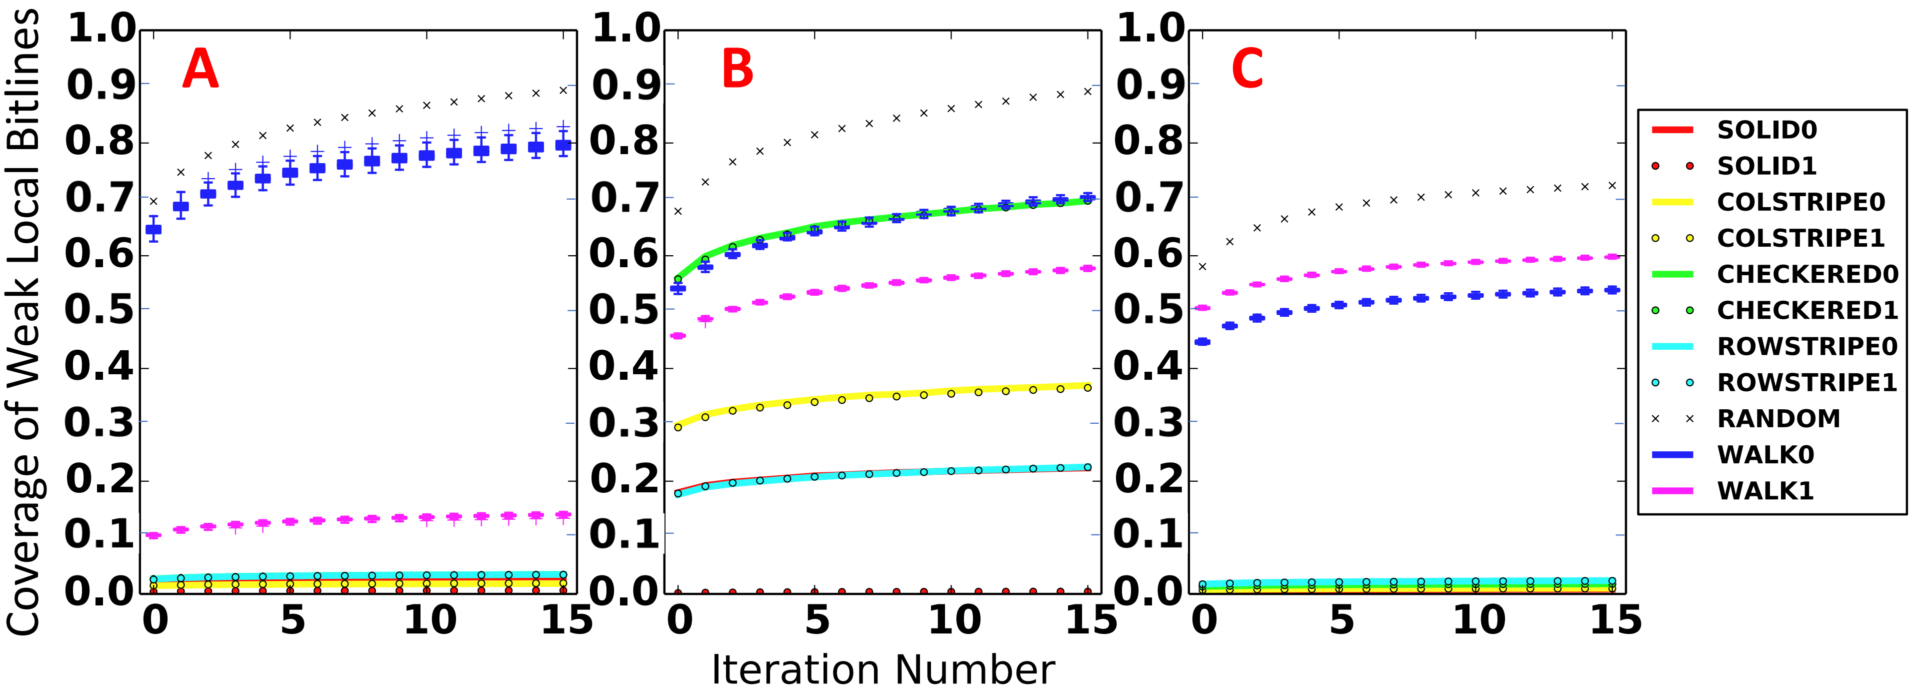

DRAM Background

DRAM chips are organized into DRAM ranks and modules.

The CPU interfaces with DRAM at the granularity of a module with a memory controller that has a 64-bit channel connection

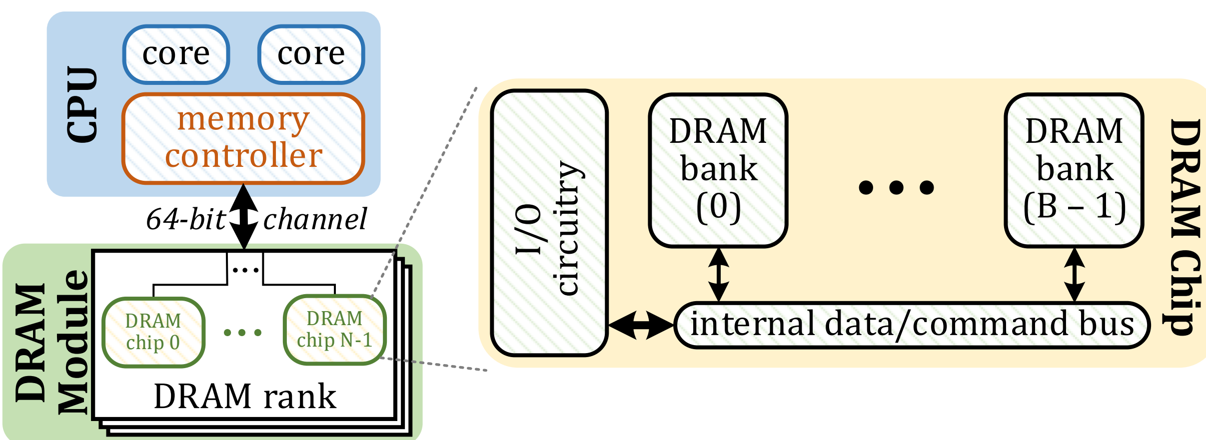

Evaluation Methodology

Processor	4 cores, 4 GHz, 4-wide issue, 8 MSHRs/core, OoO 128-entry window
LLC	8 *MiB* shared, 64B cache line, 8-way associative
Memory Controller	64-entry R/W queue, FR-FCFS [55, 74]
DRAM	LPDDR4-3200 [18], 2 channels, 1 rank/channel, 8 banks/rank, 64K rows/bank, 1024 rows/subarray, 8 *KiB* row-buffer, Baseline: $t_{RCD}/t_{RAS}/t_{WR}$ = 29/67/29 cycles (18.125/41.875/18.125 ns)
Solar-DRAM	reduced t_{RCD} for requests to strong cache lines: 18 cycles (11.25ns) reduced t_{RCD} for write requests: 7 cycles (4.375ns)

Table 1: Evaluated system configuration.

Testing Methodology

Algorithm 1: DRAM Activation Failure Testing

```
DRAM_ACT_fail_testing(data_pattern, reduced_tRCD):
    write data_pattern (e.g., solid 1s) into all DRAM cells
    foreach col in DRAM module:
        foreach row in DRAM module:
            refresh(row)                     // replenish cell voltage
            precharge(row)                   // ensure next access activates row
            read(col) with reduced_tRCD      // induce activation failures on col
            find and record activation failures
```

Implementation Overhead

[3]To store the lookup table for a DRAM channel, we require $num\_banks \times num\_subarrays\_per\_bank \times \frac{row\_size}{cacheline\_size}$ bits, where *num_subarrays_per_bank* is the number of subarrays in a bank, *row_size* is the size of a DRAM row in bits, and *cacheline_size* is the size of a cache line in bits. For a 4GB DRAM module with 8 banks, 64 subarrays per bank, 32-byte cache lines, and 2KB per row, the lookup table requires 4KB of storage.

The table is stored in the memory controller that interfaces with the DRAM channel