An Experimental Study of Data Retention Behavior in Modern DRAM Devices

Implications for Retention Time Profiling Mechanisms

Jamie Liu[1] Ben Jaiyen[1] Yoongu Kim[1]
Chris Wilkerson[2] Onur Mutlu[1]

[1] Carnegie Mellon University
[2] Intel Corporation

SAFARI

Carnegie Mellon

Talk Agenda

- DRAM Refresh: Background and Motivation
- Challenges and Our Goal
- DRAM Characterization Methodology
- Foundational Results
  - Temperature Dependence
  - Retention Time Distribution
- Data Pattern Dependence: Analysis and Implications
- Variable Retention Time: Analysis and Implications
- Conclusions

A DRAM Cell

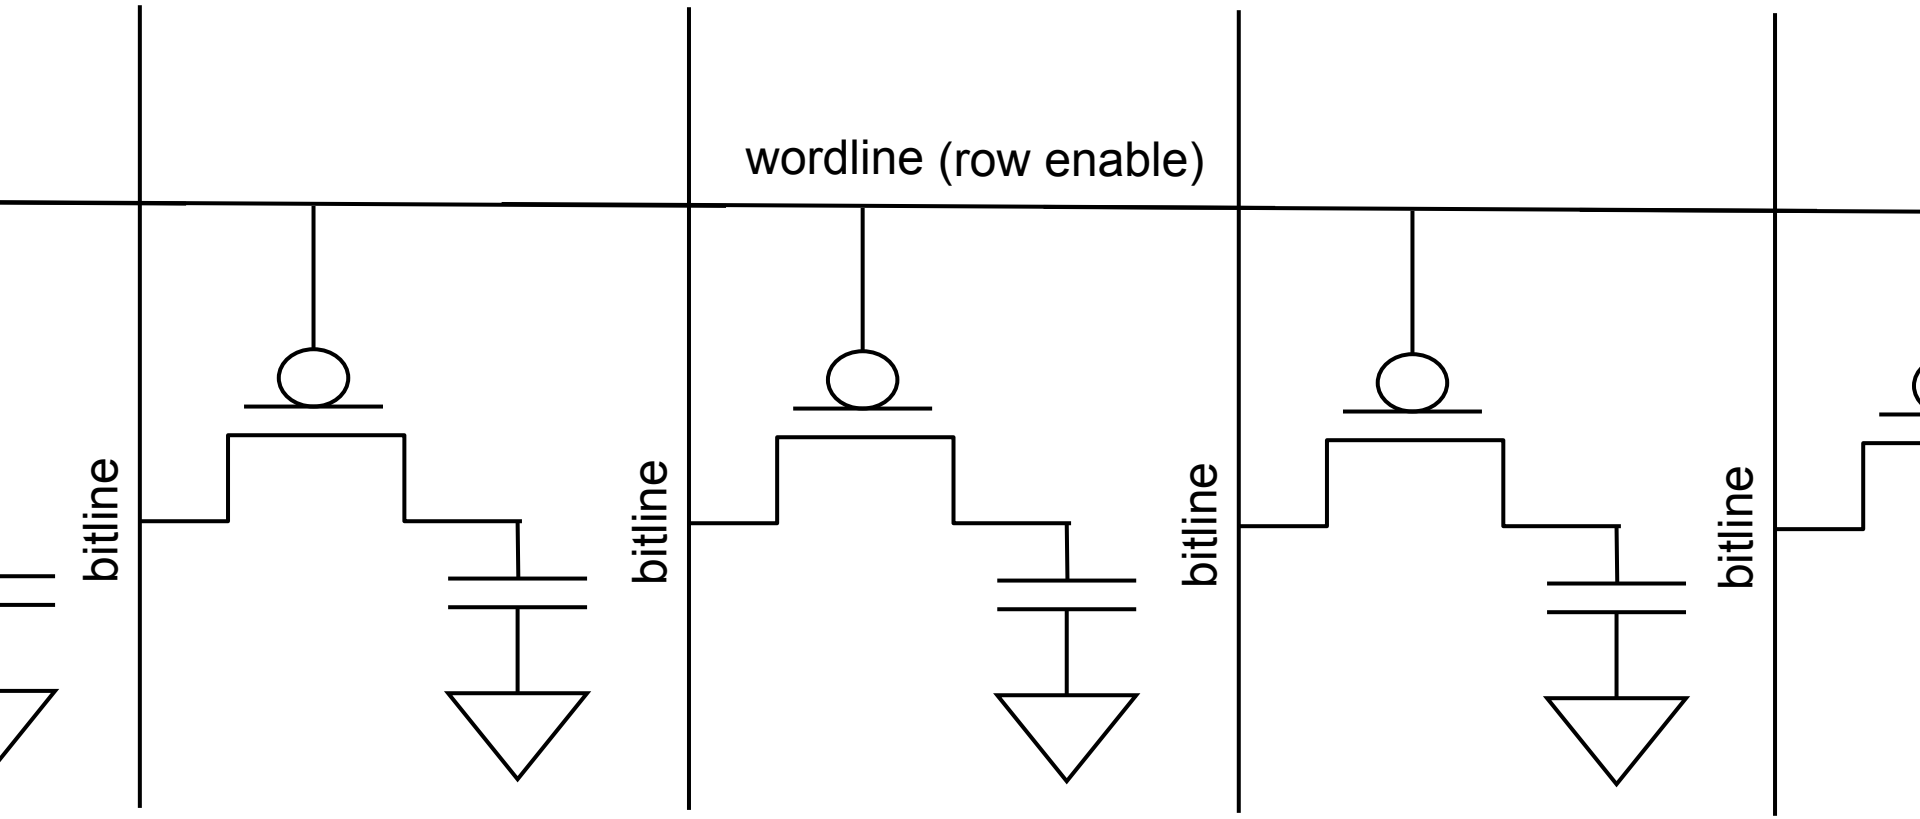

- A DRAM cell consists of a capacitor and an access transistor
- It stores data in terms of charge in the capacitor
- A DRAM chip consists of (10s of 1000s of) rows of such cells

DRAM Refresh

- DRAM capacitor charge leaks over time

- Each DRAM row is periodically refreshed to restore charge
  - Activate each row every N ms
  - Typical N = 64 ms

- Downsides of refresh

  -- Energy consumption: Each refresh consumes energy

  -- Performance degradation: DRAM rank/bank unavailable while refreshed

  -- QoS/predictability impact: (Long) pause times during refresh

  -- Refresh rate limits DRAM capacity scaling

Refresh Overhead: Performance

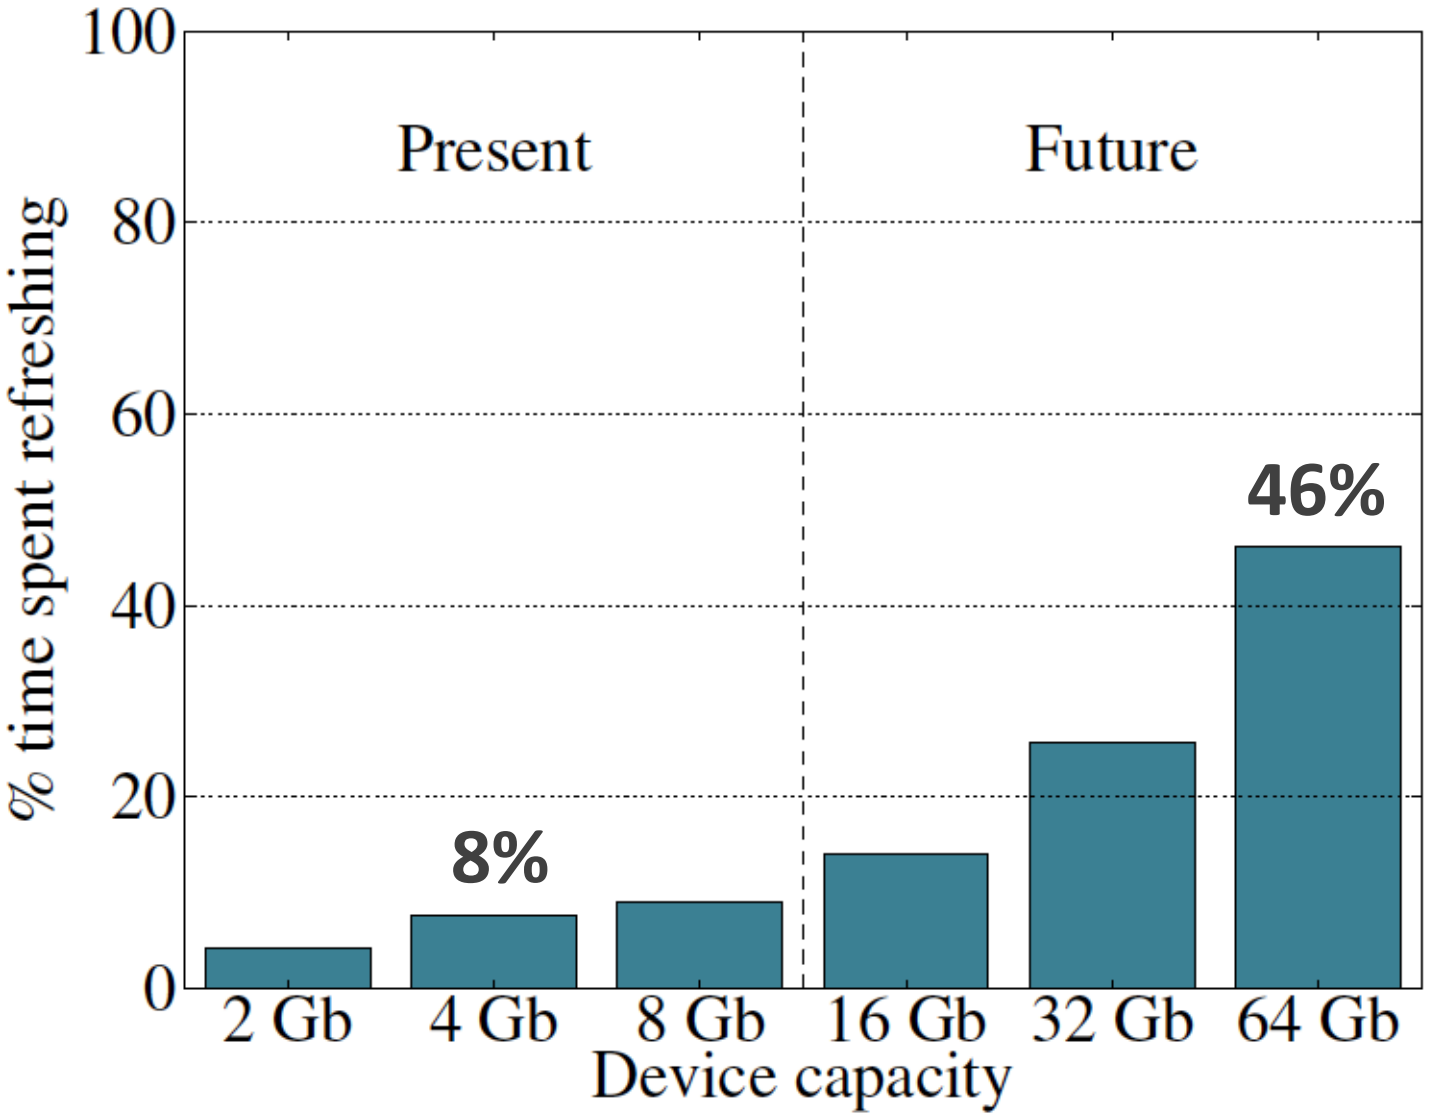

Liu et al., "RAIDR: Retention-Aware Intelligent DRAM Refresh," ISCA 2012.

Refresh Overhead: Energy

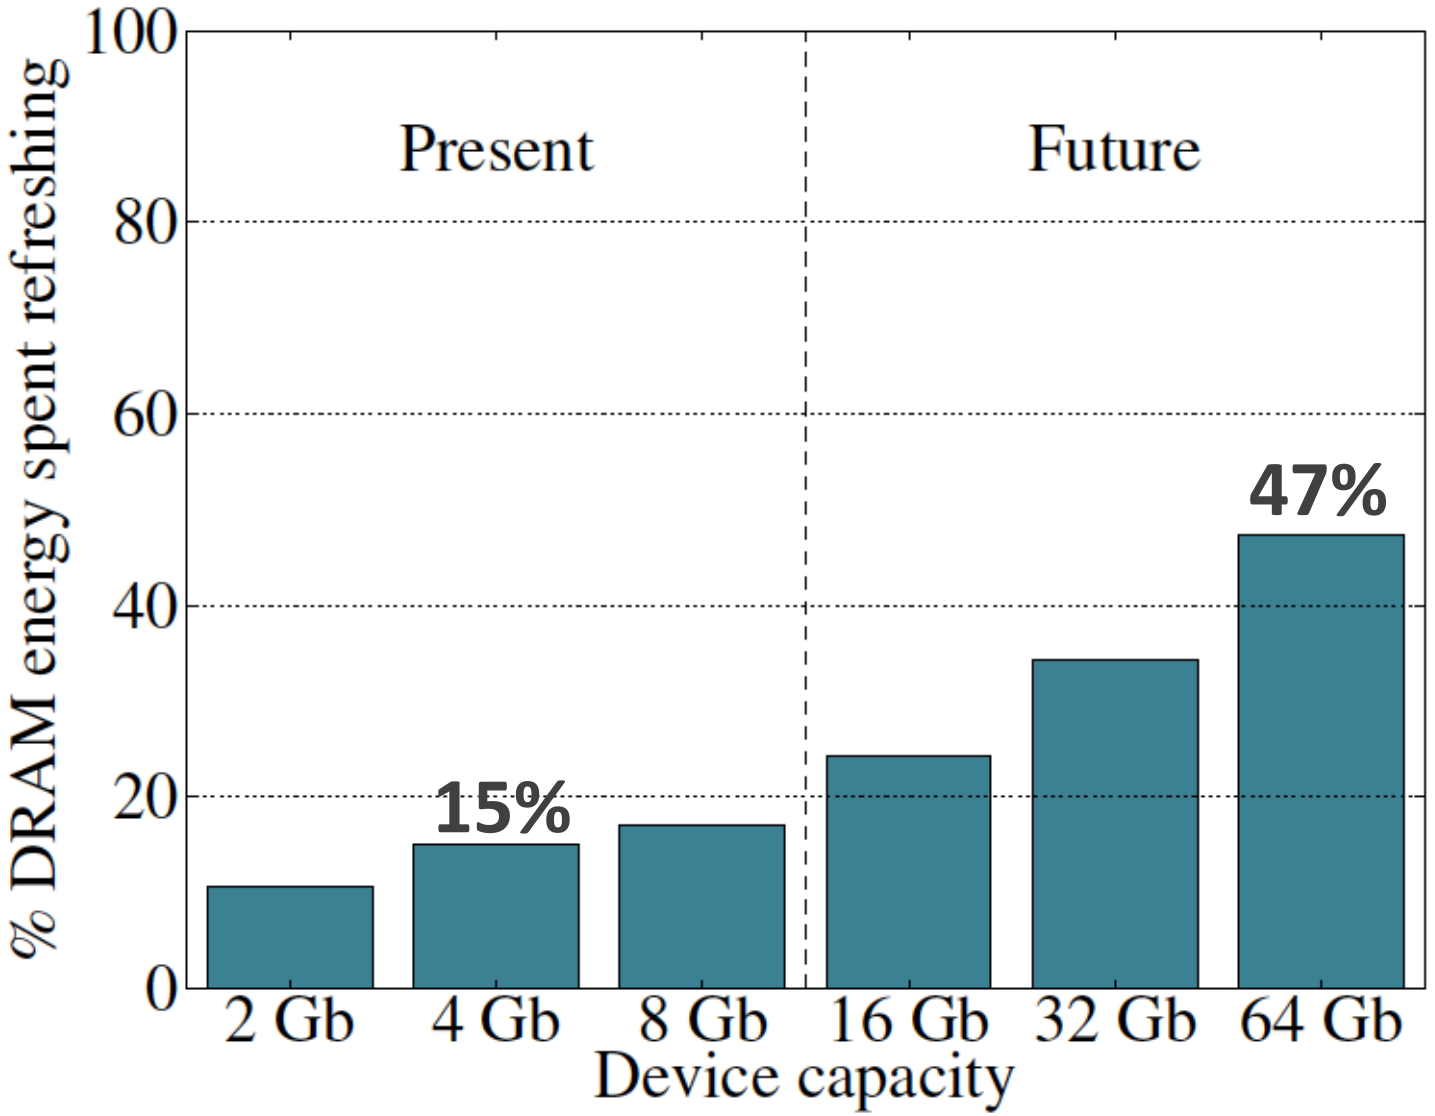

Liu et al., "RAIDR: Retention-Aware Intelligent DRAM Refresh," ISCA 2012.

Previous Work on Reducing Refreshes

- Observed significant variation in data retention times of DRAM cells (due to manufacturing process variation)
  - Retention time: maximum time a cell can go without being refreshed while maintaining its stored data

- Proposed methods to take advantage of widely varying retention times among DRAM rows
  - Reduce refresh rate for rows that can retain data for longer than 64 ms, e.g., **[Liu+ ISCA 2012]**
  - Disable rows that have low retention times, e.g., **[Venkatesan+ HPCA 2006]**

- Showed large benefits in energy and performance

An Example: RAIDR [Liu+, ISCA 2012]

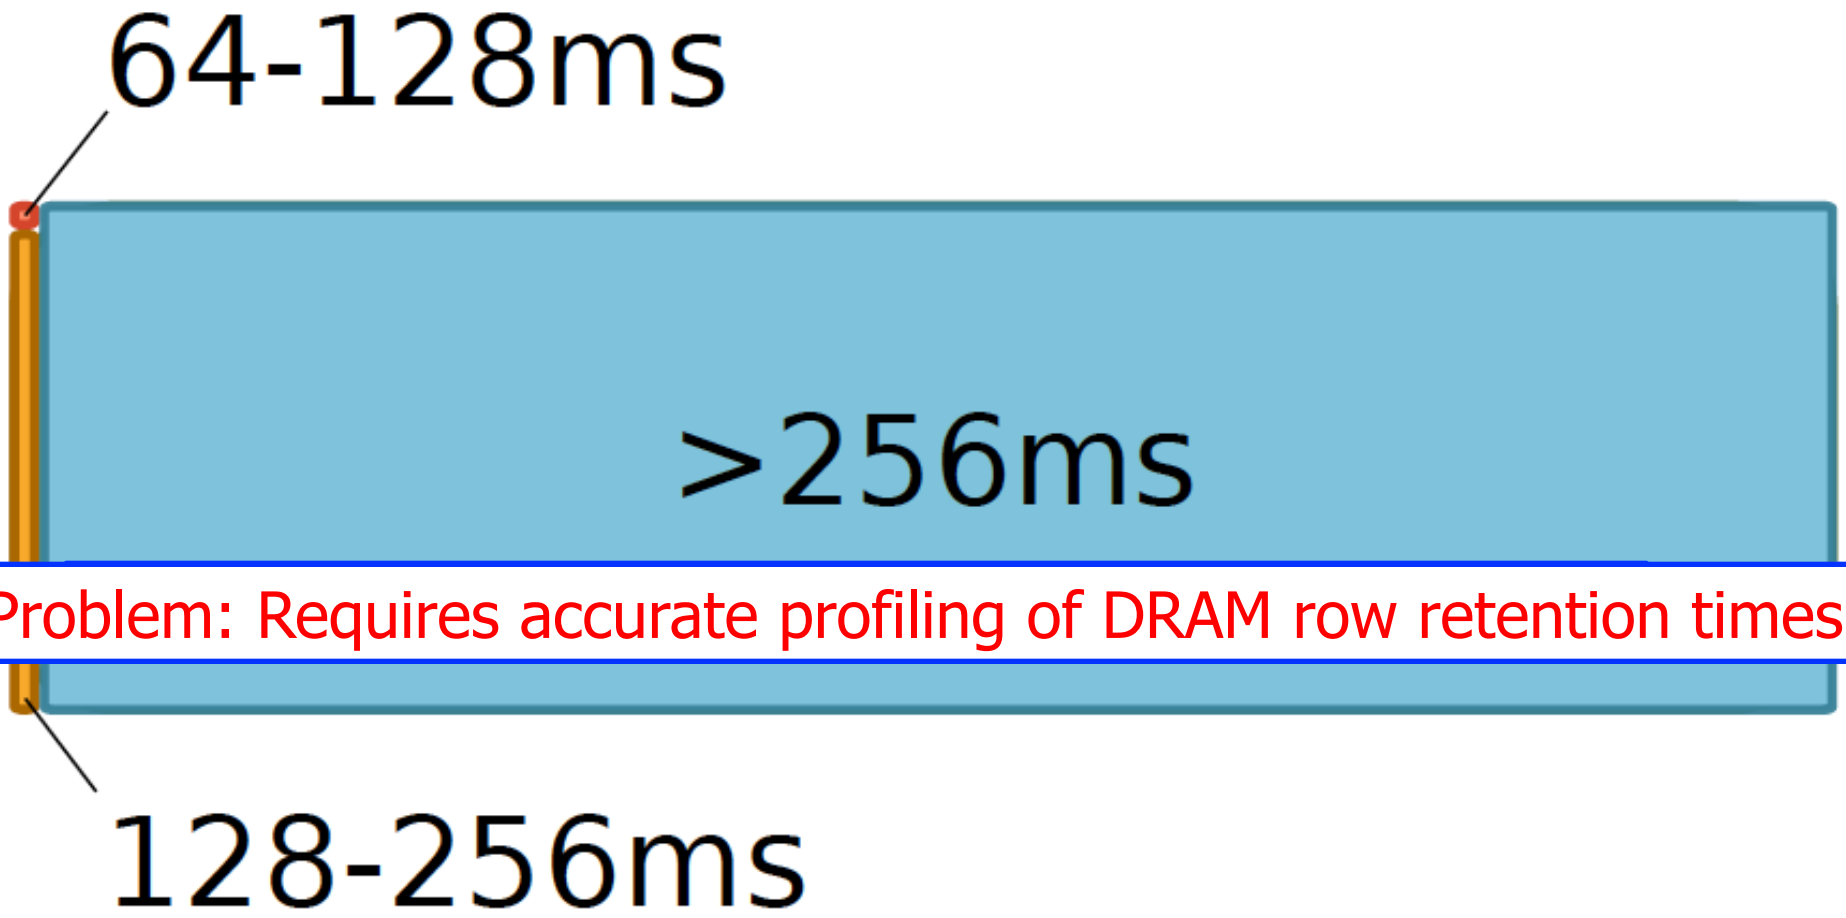

Problem: Requires accurate profiling of DRAM row retention times

Can reduce refreshes by ~75%
→ reduces energy consumption and improves performance

Motivation

- Past works require accurate and reliable measurement of retention time of each DRAM row
  - To maintain data integrity while reducing refreshes

- Assumption: worst-case retention time of each row can be determined and stays the same at a given temperature
  - Some works propose writing all 1's and 0's to a row, and measuring the time before data corruption

- Question:
  - Can we reliably and accurately determine retention times of all DRAM rows?

Talk Agenda

- DRAM Refresh: Background and Motivation
- Challenges and Our Goal
- DRAM Characterization Methodology
- Foundational Results
  - Temperature Dependence
  - Retention Time Distribution
- Data Pattern Dependence: Analysis and Implications
- Variable Retention Time: Analysis and Implications
- Conclusions

Two Challenges to Retention Time Profiling

- Data Pattern Dependence (DPD) of retention time

- Variable Retention Time (VRT) phenomenon

Two Challenges to Retention Time Profiling

- Challenge 1: Data Pattern Dependence (DPD)
  - Retention time of a DRAM cell depends on its value and the values of cells nearby it
  - When a row is activated, all bitlines are perturbed simultaneously

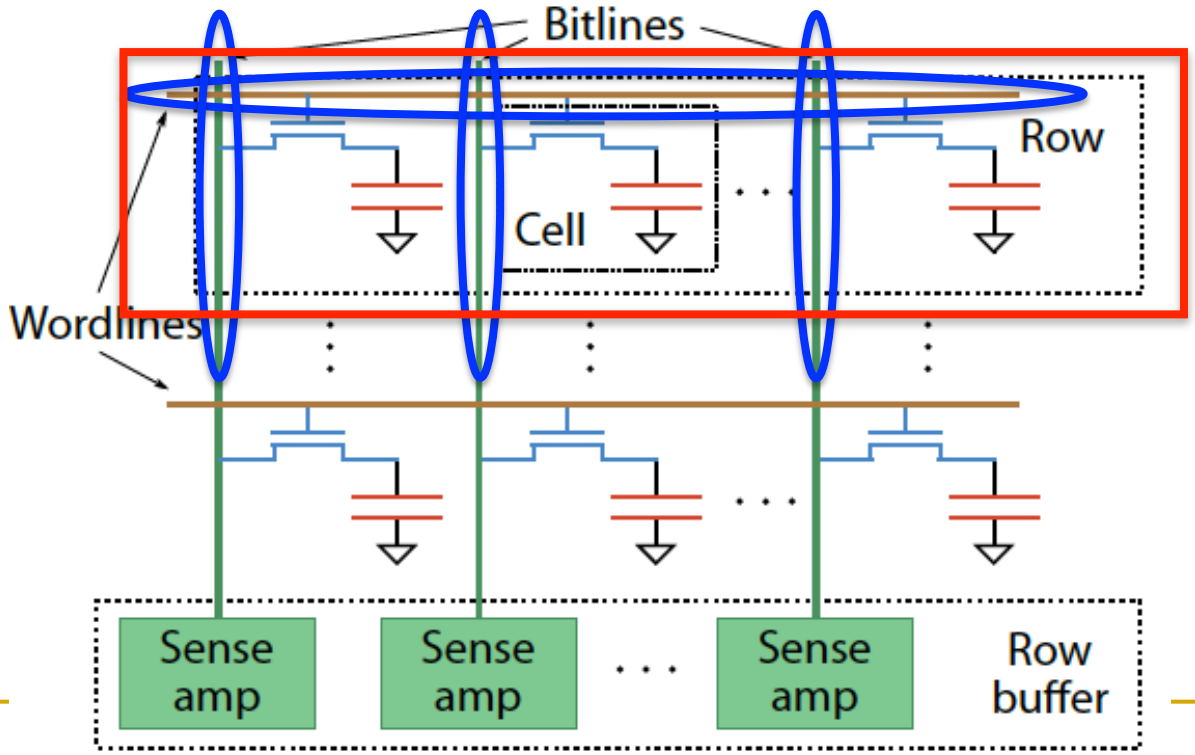

Data Pattern Dependence

- Electrical noise on the bitline affects reliable sensing of a DRAM cell
- The magnitude of this noise is affected by values of nearby cells via
  - Bitline-bitline coupling → electrical coupling between adjacent bitlines
  - Bitline-wordline coupling → electrical coupling between each bitline and the activated wordline

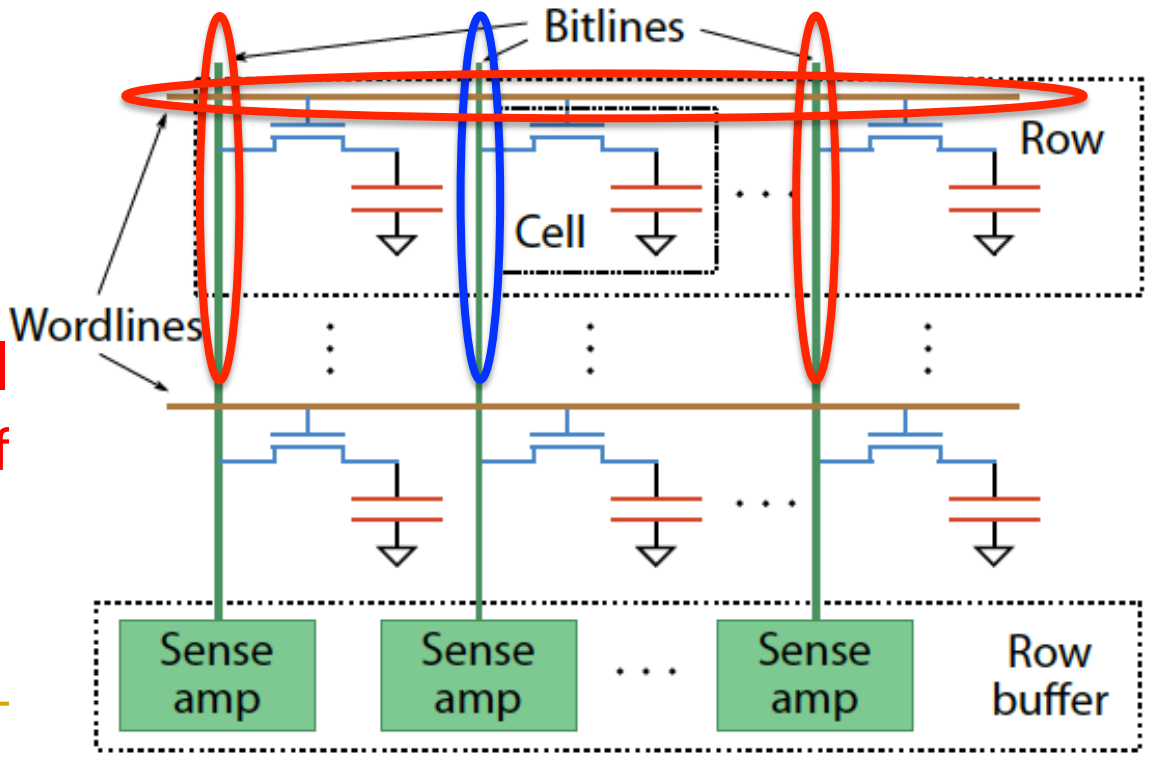

- Retention
  nearby cel
  → need to f

tored in

etention time

Two Challenges to Retention Time Profiling

- Challenge 2: Variable Retention Time (VRT)
  - Retention time of a DRAM cell changes randomly over time
    - a cell alternates between multiple retention time states
  - Leakage current of a cell changes sporadically due to a charge trap in the gate oxide of the DRAM cell access transistor
  - When the trap becomes occupied, charge leaks more readily from the transistor's drain, leading to a short retention time
    - Called *Trap-Assisted Gate-Induced Drain Leakage*
  - This process appears to be a random process ~~[Kim+ IEEE TED'11]~~
  - Worst-case retention time depends on a random process
    → need to find the worst case despite this

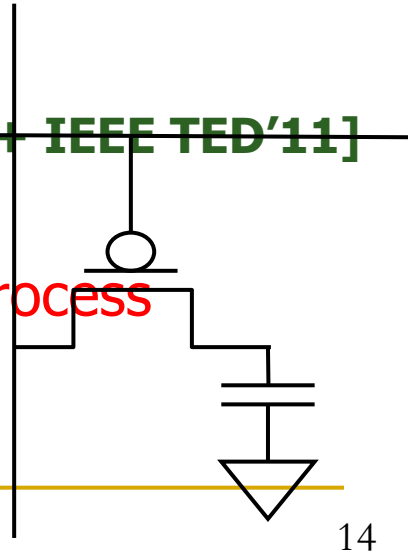

Our Goal

- Analyze the retention time behavior of DRAM cells in modern commodity DRAM devices
  - to aid the collection of accurate profile information

- Provide a comprehensive empirical investigation of two key challenges to retention time profiling
  - Data Pattern Dependence (DPD)
  - Variable Retention Time (VRT)

Talk Agenda

- DRAM Refresh: Background and Motivation
- Challenges and Our Goal
- DRAM Characterization Methodology
- Foundational Results
  - Temperature Dependence
  - Retention Time Distribution
- Data Pattern Dependence: Analysis and Implications
- Variable Retention Time: Analysis and Implications
- Conclusions

DRAM Testing Platform and Method

- Test platform: Developed a DDR3 DRAM testing platform using the Xilinx ML605 FPGA development board
  - Temperature controlled

- Tested DRAM chips: 248 commodity DRAM chips from five manufacturers (A,B,C,D,E)
- Seven families based on equal capacity per device:
  - A 1Gb, A 2Gb
  - B 2Gb
  - C 2Gb
  - D 1Gb, D 2Gb
  - E 2Gb

Experiment Design

- Each module tested for multiple ***rounds*** of ***tests***.
- Each test searches for the set of cells with a retention time less than a threshold value for a particular data pattern
- High-level structure of a test:
  - Write data pattern to rows in a DRAM bank
  - Prevent refresh for a period of time *tWAIT*, leave DRAM idle
  - Read stored data pattern, compare to written pattern and record corrupt cells as those with retention time < *tWAIT*
- Test details and important issues to pay attention to are discussed in paper

Experiment Structure

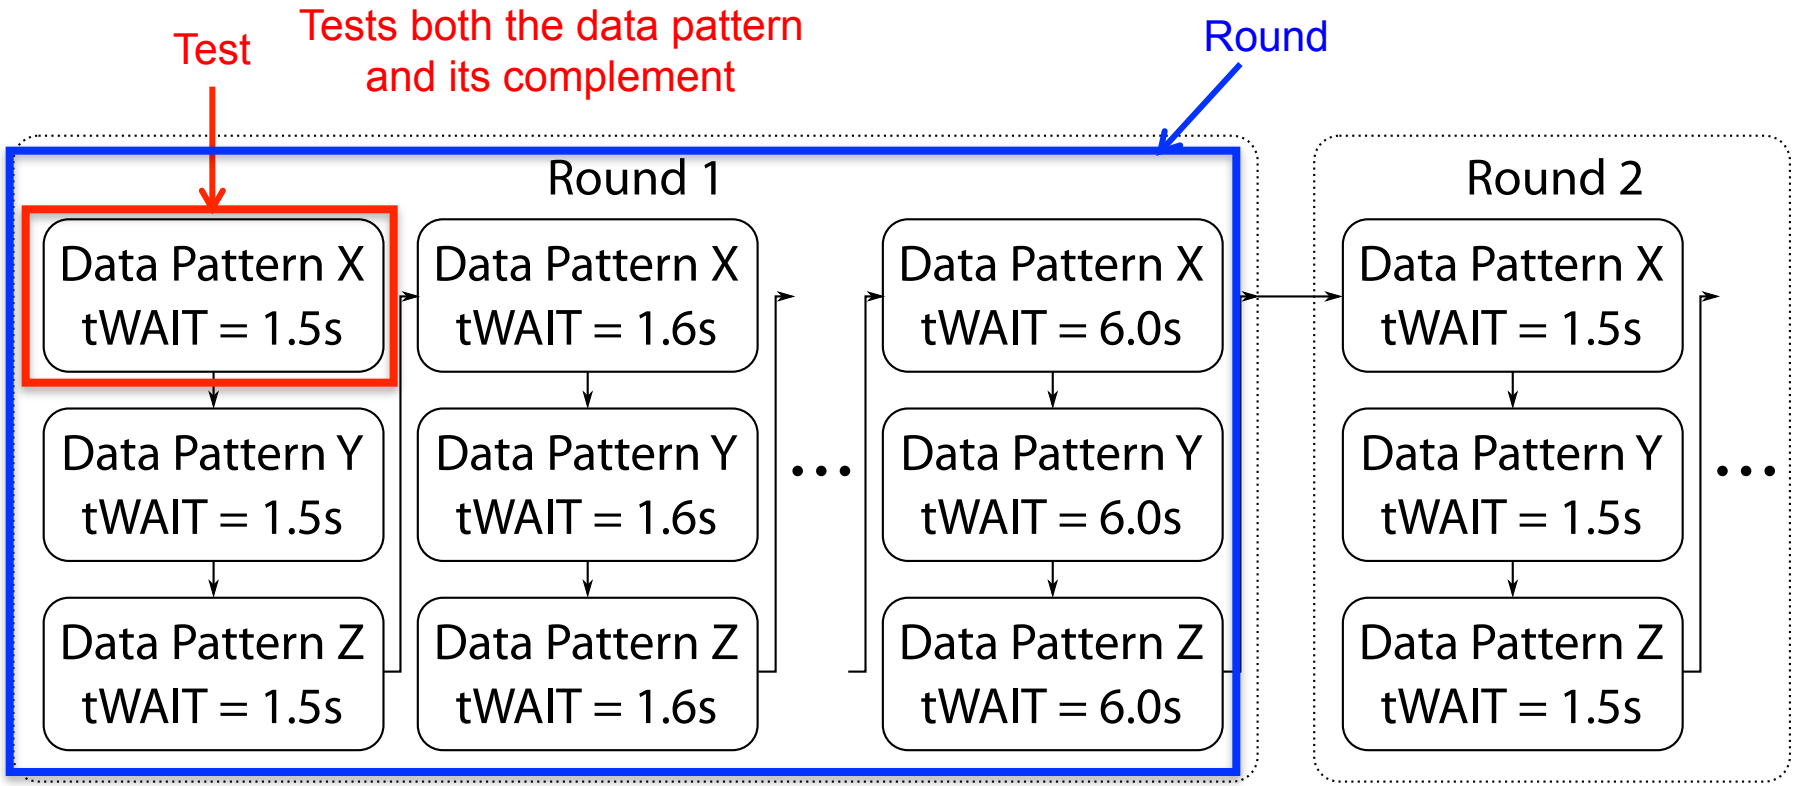

Experiment Parameters

- Most tests conducted at 45 degrees Celsius
- No cells observed to have a retention time less than 1.5 second at 45^{o}C
- Tested *tWAIT* in increments of 128ms from 1.5 to 6.1 seconds

Tested Data Patterns

Fixed patterns

- All 0s/1s: Value 0/1 is written to all bits
  - Previous work suggested this is sufficient
- Checkerboard: Consecutive bits alternate between 0 and 1
  - Coupling noise increases with voltage difference between the neighboring bitlines → May induce worst case data pattern (if adjacent bits mapped to adjacent cells)
- Walk: Attempts to ensure a single cell storing 1 is surrounded by cells storing 0
  - This may lead to even worse coupling noise and retention time due to coupling between *nearby* bitlines **[Li+ IEEE TCSI 2011]**
  - Walk pattern is permuted in each round to exercise different cells
- Random: Randomly generated data is written to each row
  - A new set of random data is generated for each round

Talk Agenda

- DRAM Refresh: Background and Motivation
- Challenges and Our Goal
- DRAM Characterization Methodology
- Foundational Results
  - Temperature Dependence
  - Retention Time Distribution
- Data Pattern Dependence: Analysis and Implications
- Variable Retention Time: Analysis and Implications
- Conclusions

Temperature Stability

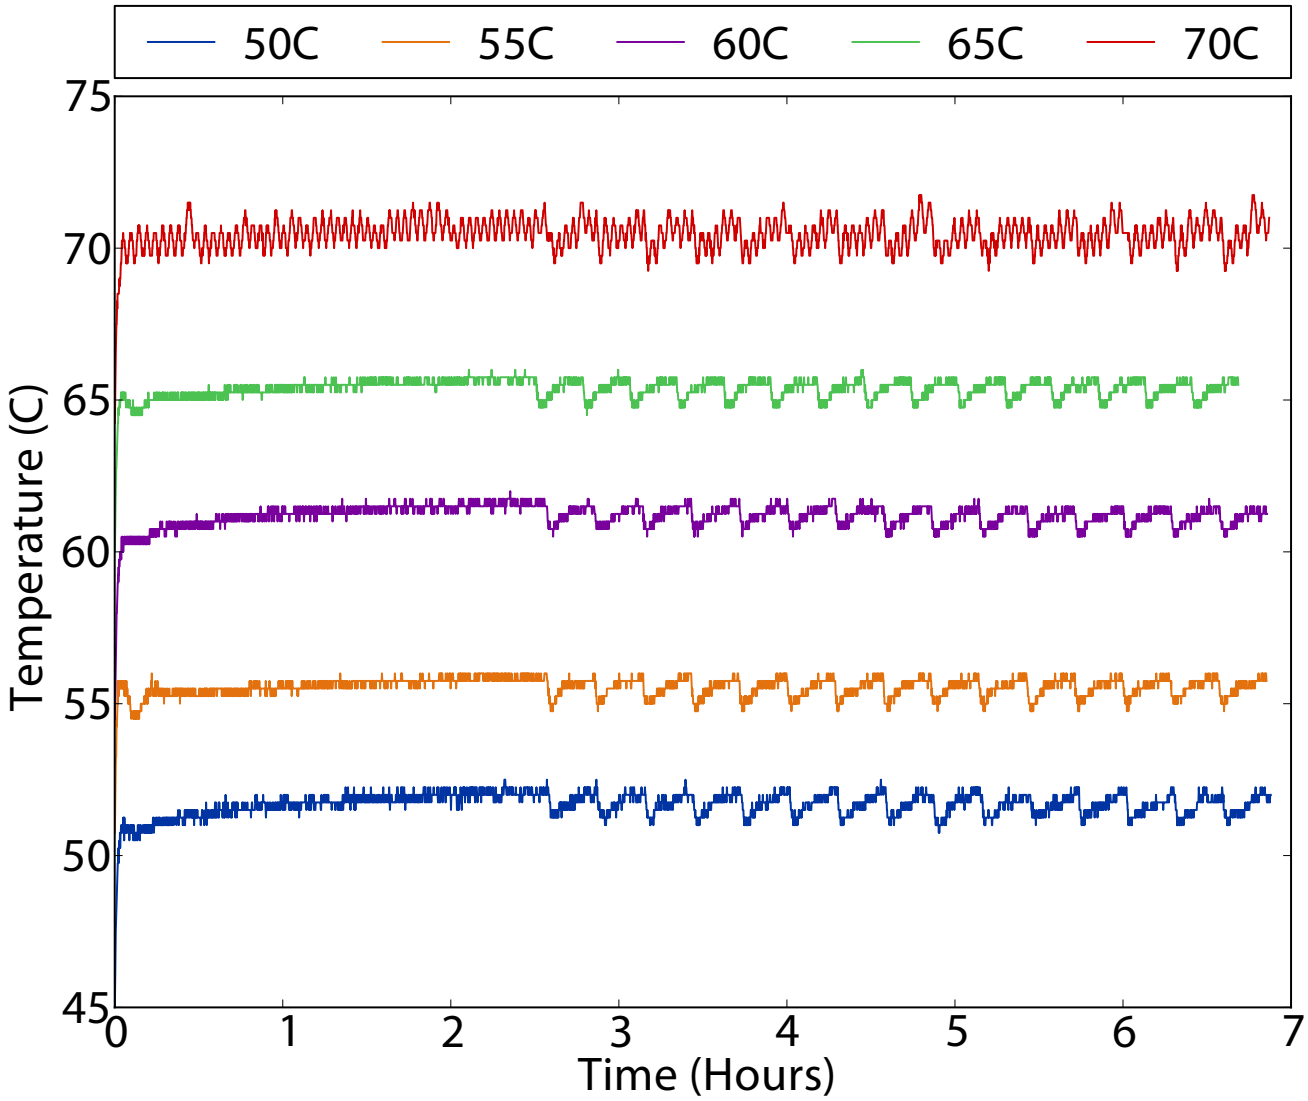

Tested chips at five different stable temperatures

Dependence of Retention Time on Temperature

Fraction of cells that exhibited retention time failure at any *tWAIT* for any data pattern at 50°C

Normalized retention times of the same cells at 55°C

Normalized retention times of the same cells At 70°C

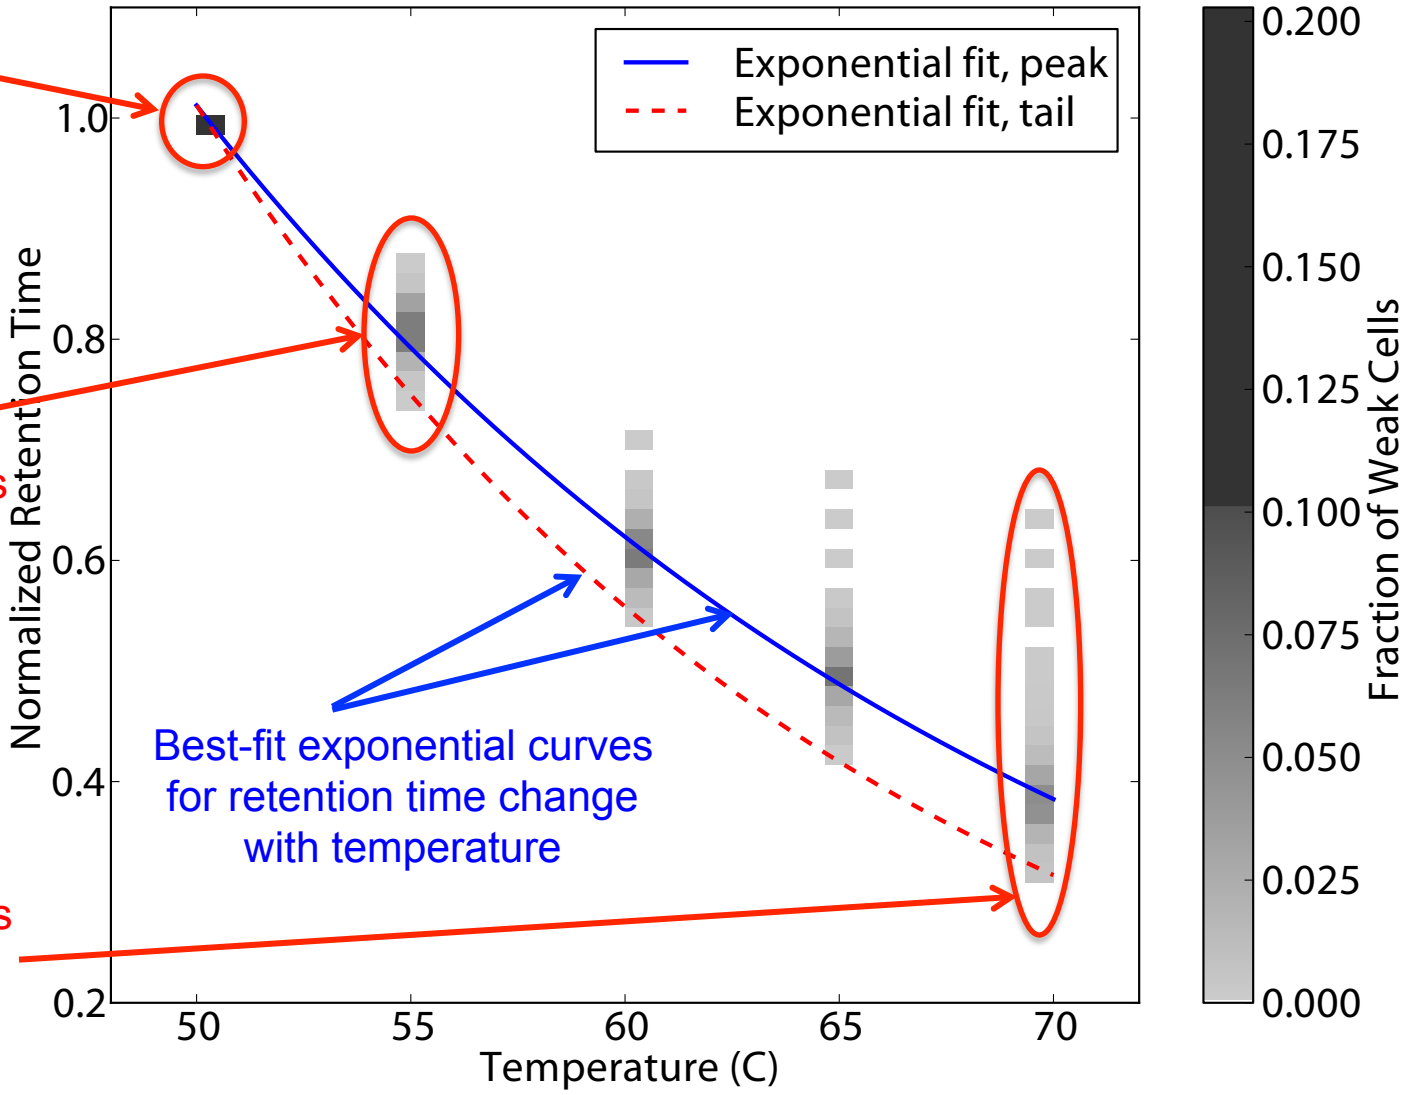

Dependence of Retention Time on Temperature

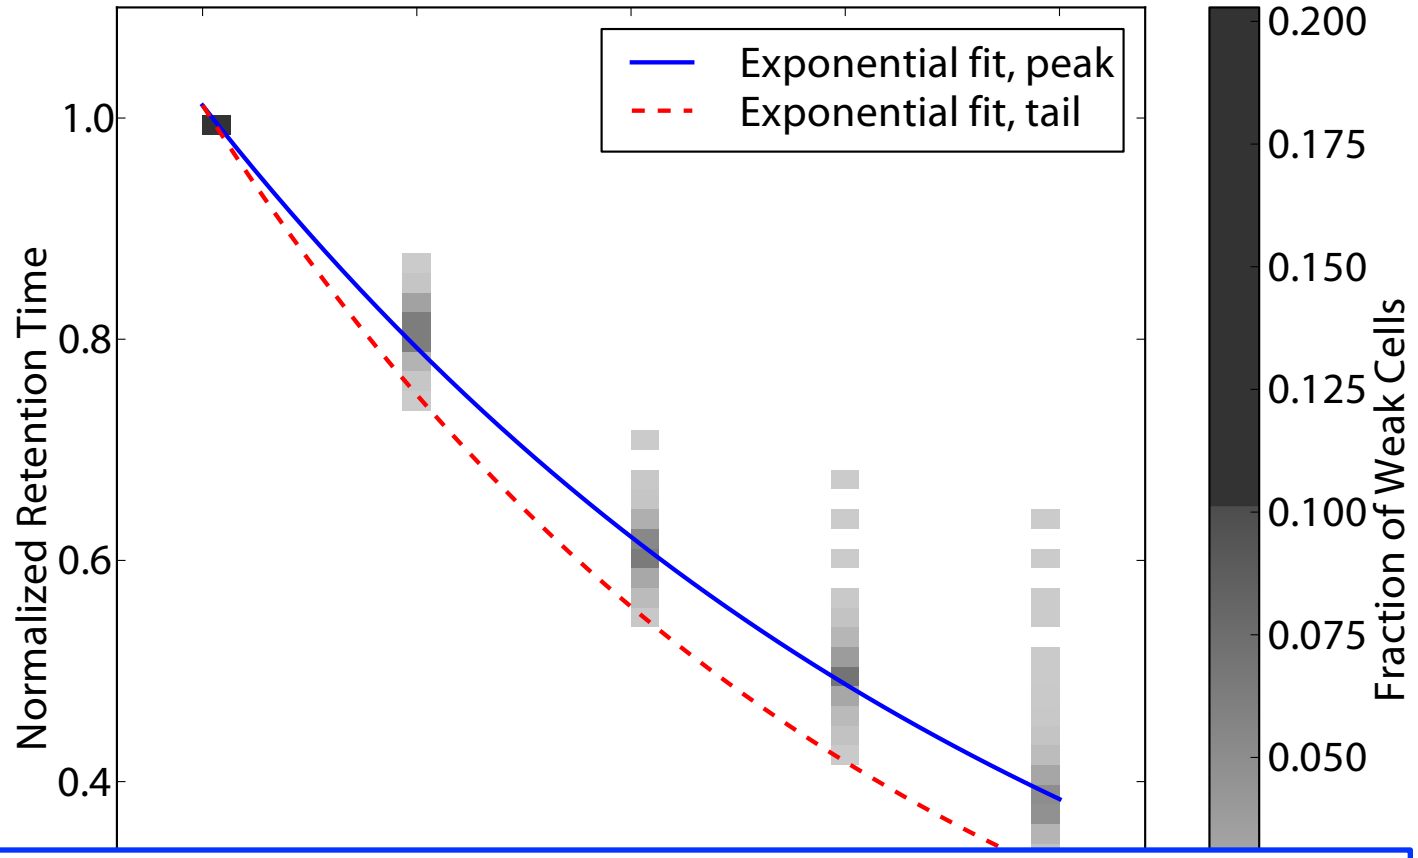

Relationship between retention time and temperature is consistently bounded (predictable) within a device

Every 10°C temperature increase
→ 46.5% reduction in retention time in the worst case

Retention Time Distribution

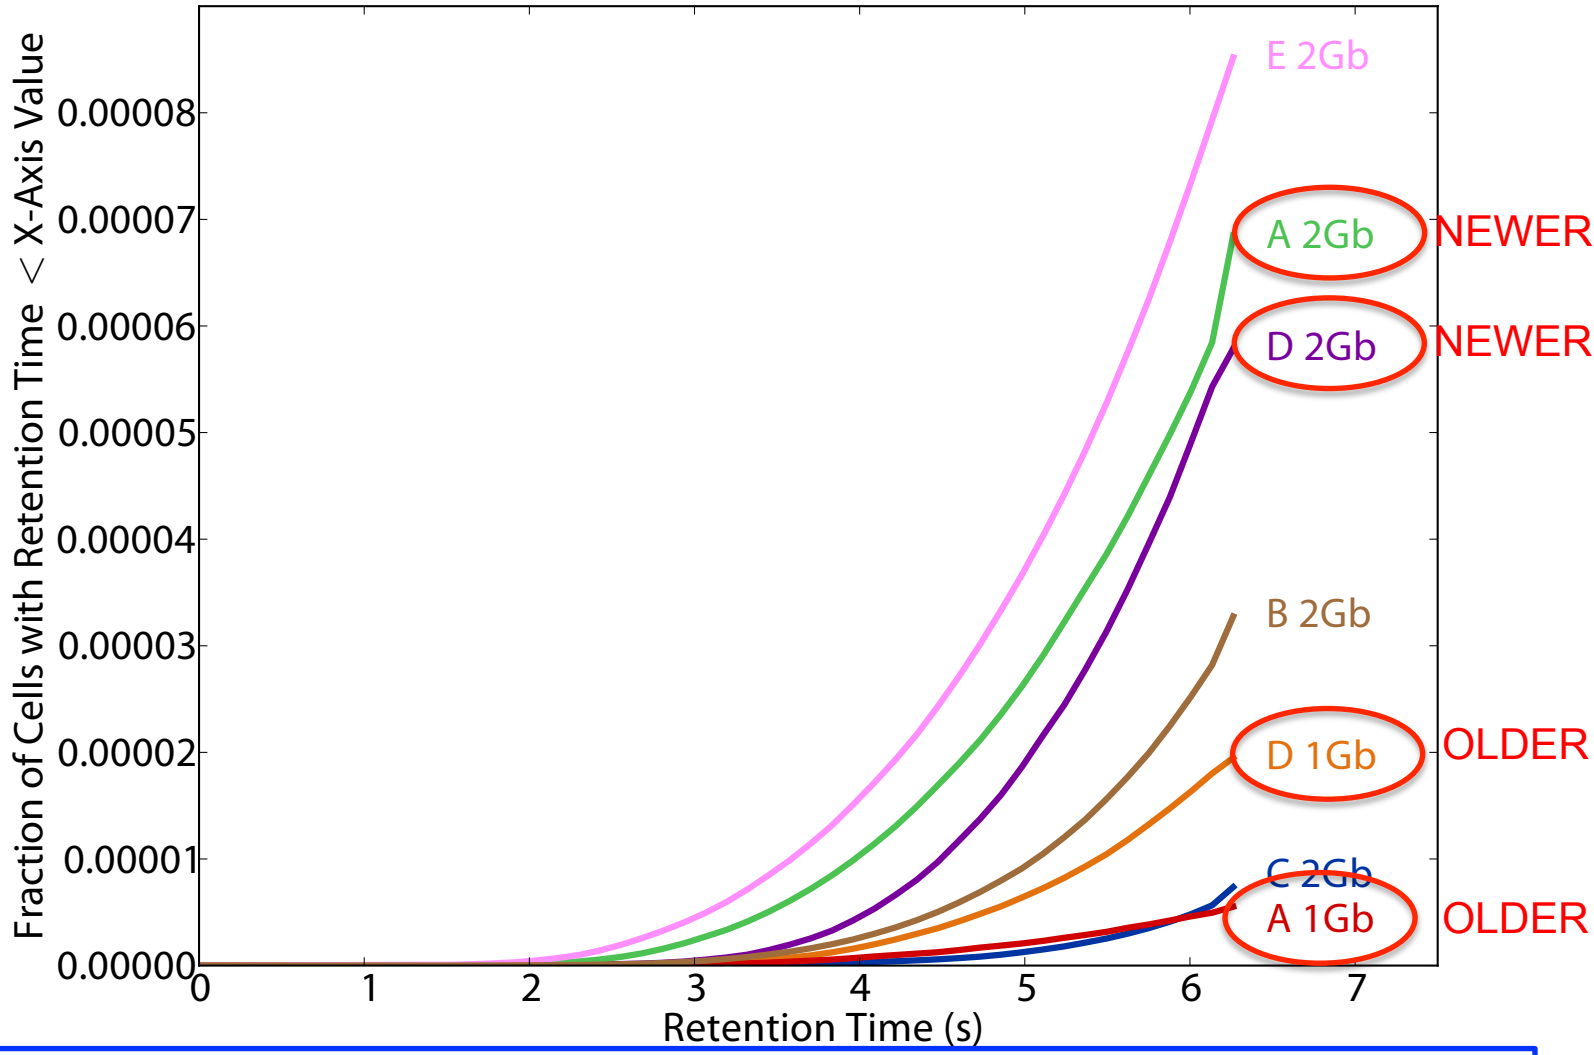

Newer device families have more weak cells than older ones
Likely a result of technology scaling

Talk Agenda

- DRAM Refresh: Background and Motivation
- Challenges and Our Goal
- DRAM Characterization Methodology
- Foundational Results
  - Temperature Dependence
  - Retention Time Distribution
- Data Pattern Dependence: Analysis and Implications
- Variable Retention Time: Analysis and Implications
- Conclusions

Some Terminology

- Failure population of cells with Retention Time X: The set of all cells that exhibit retention failure in any test with *any* data pattern at that retention time (*tWAIT*)

- Retention Failure Coverage of a Data Pattern DP: Fraction of cells with retention time X that exhibit retention failure with that *particular* data pattern DP

- If retention times are not dependent on data pattern stored in cells, we would expect
  - Coverage of any data pattern to be 100%
  - In other words, if one data pattern causes a retention failure, any other data pattern also would

Recall the Tested Data Patterns

Fixed patterns

- All 0s/1s: Value 0/1 is written to all bits
- Checkerboard: Consecutive bits alternate between 0 and 1
- Walk: Attempts to ensure a single cell storing 1 is surrounded by cells storing 0
- Random: Randomly generated data is written to each row

Retention Failure Coverage of Data Patterns

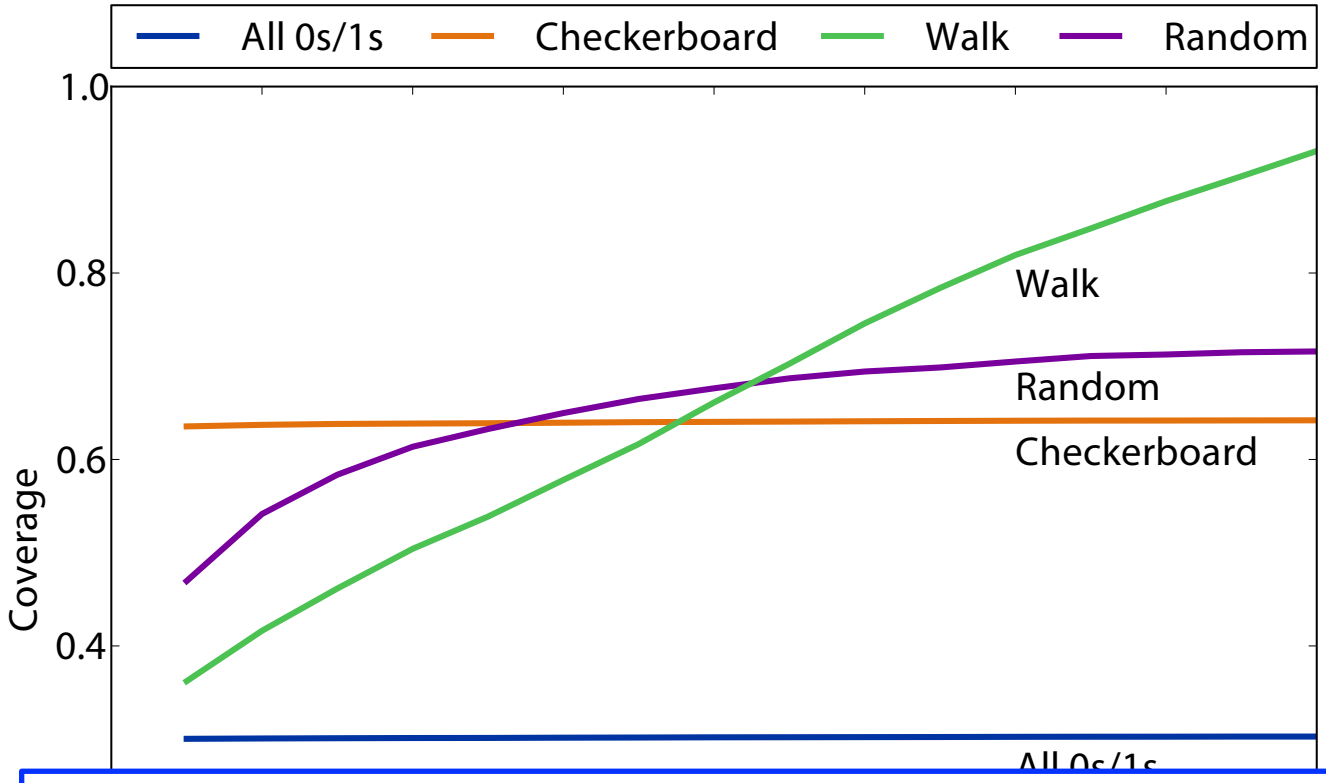

A 2Gb chip family

6.1s retention time

Different data patterns have widely different coverage: Data pattern dependence exists and is severe

Coverage of fixed patterns is low: ~30% for *All 0s/1s*

***Walk* is the most effective data pattern for this device**

No data pattern achieves 100% coverage

Retention Failure Coverage of Data Patterns

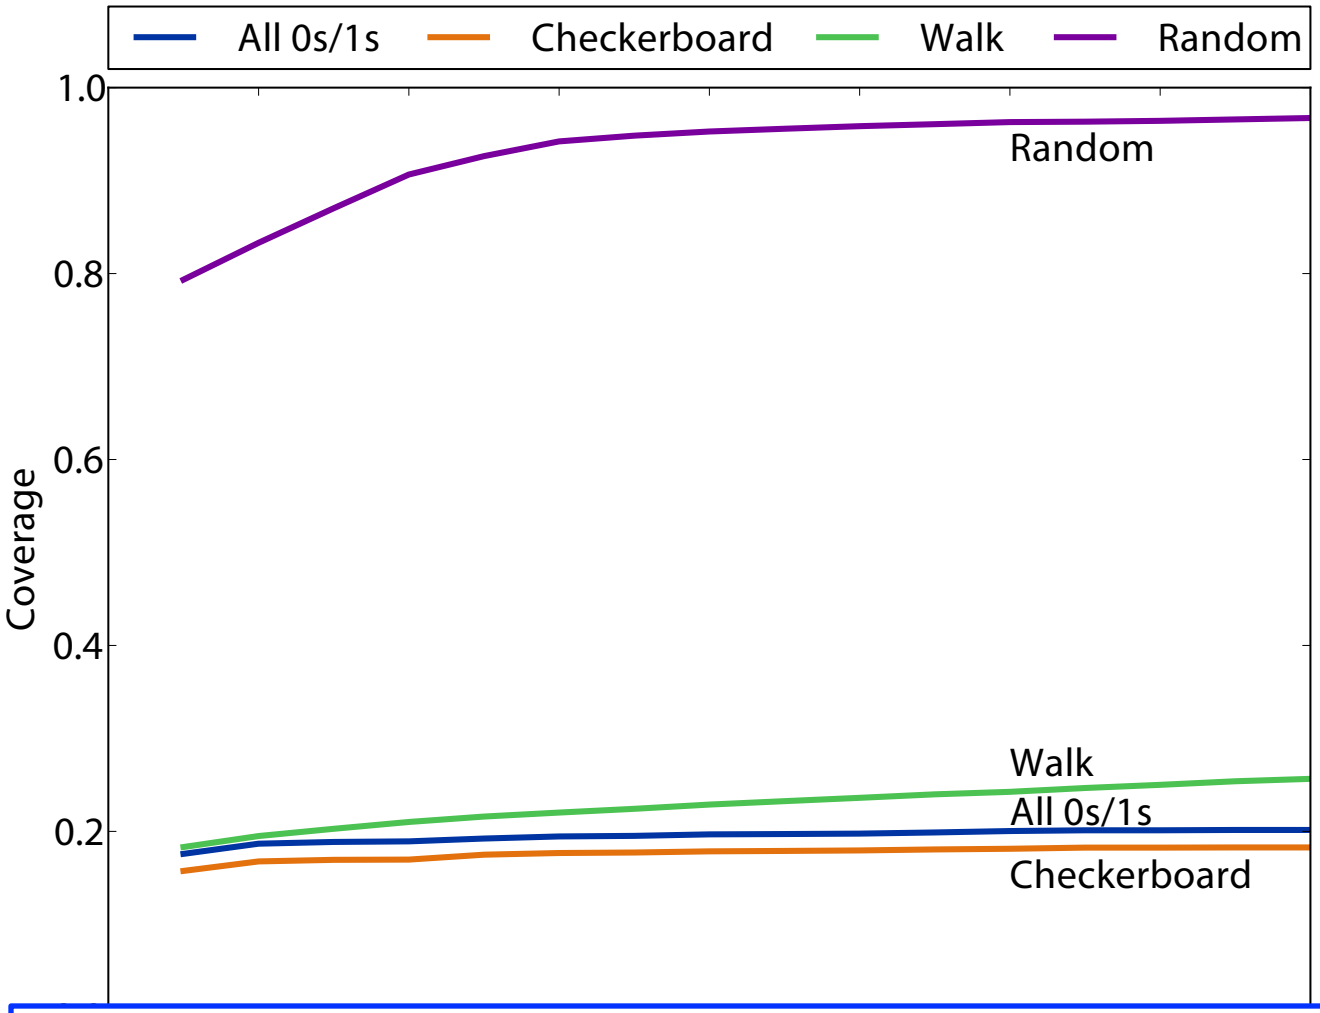

B 2Gb chip family

6.1s retention time

***Random* is the most effective data pattern for this device**

No data pattern achieves 100% coverage

Retention Failure Coverage of Data Patterns

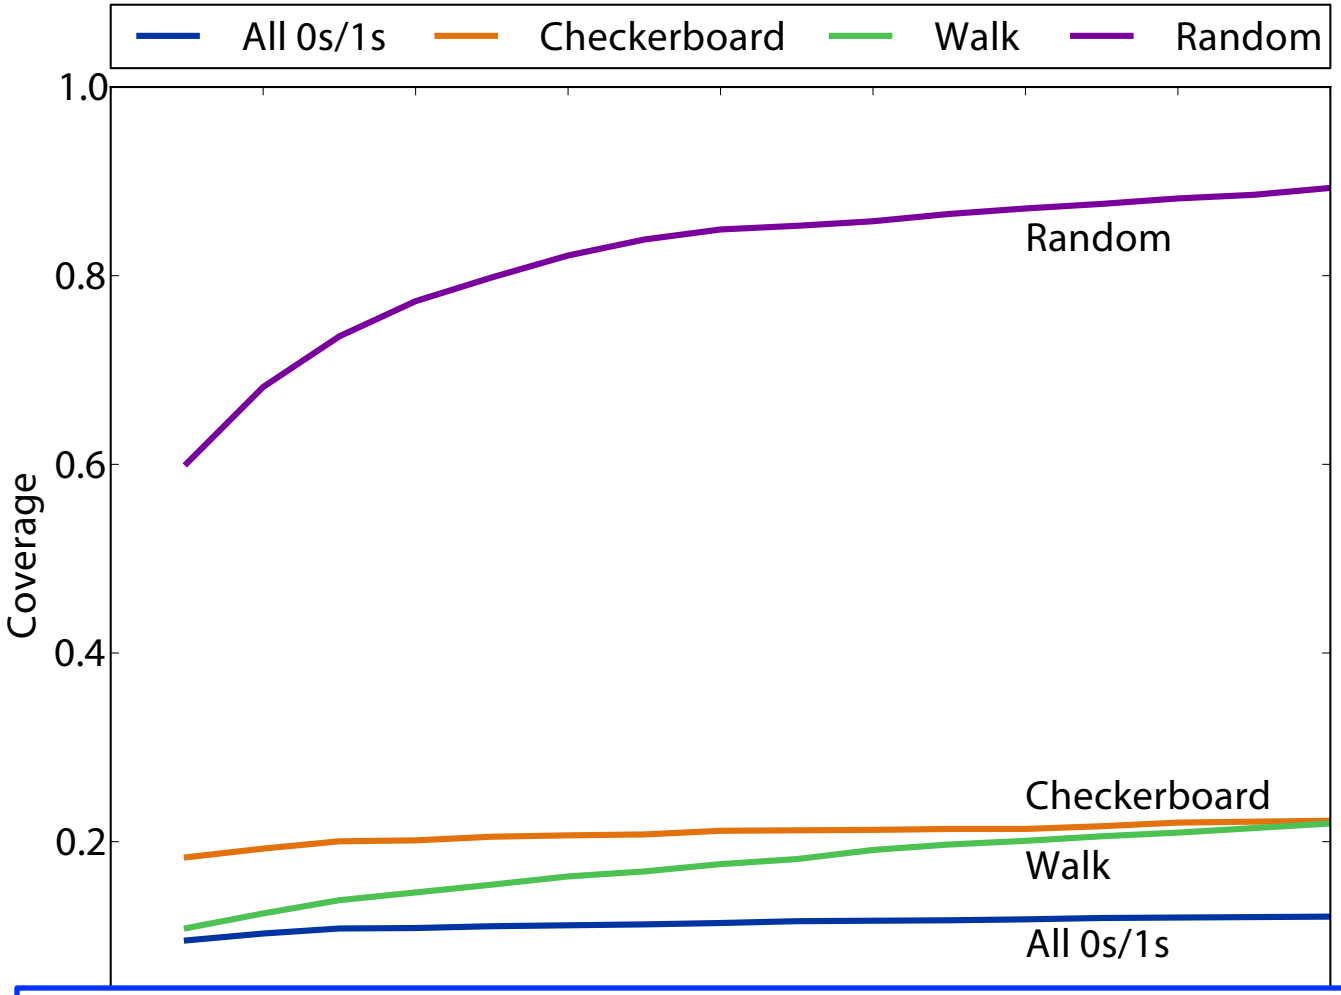

C 2Gb chip family

6.1s retention time

***Random* is the most effective data pattern for this device**

No data pattern achieves 100% coverage

Data Pattern Dependence: Observations (I)

- A cell's retention time is heavily influenced by data pattern stored in other cells
  - Pattern affects the coupling noise, which affects cell leakage

- No tested data pattern exercises the worst case retention time for all cells (no pattern has 100% coverage)
  - No pattern is able to induce the worst-case coupling noise for every cell
  - Problem: Underlying DRAM circuit organization is *not* known to the memory controller → very hard to construct a pattern that exercises the worst-case cell leakage
    - → Opaque mapping of addresses to physical DRAM geometry
    - → Internal remapping of addresses within DRAM to tolerate faults
    - → Second order coupling effects are very hard to determine

Data Pattern Dependence: Observations (II)

- Fixed, simple data patterns have low coverage
  - They do not exercise the worst-case coupling noise
- The effectiveness of each data pattern varies significantly between DRAM devices (of the same or different vendors)
  - Underlying DRAM circuit organization likely differs between different devices → patterns leading to worst coupling are different in different devices
- Technology scaling appears to increase the impact of data pattern dependence
  - Scaling reduces the physical distance between circuit elements, increasing the magnitude of coupling effects

Effect of Technology Scaling on DPD

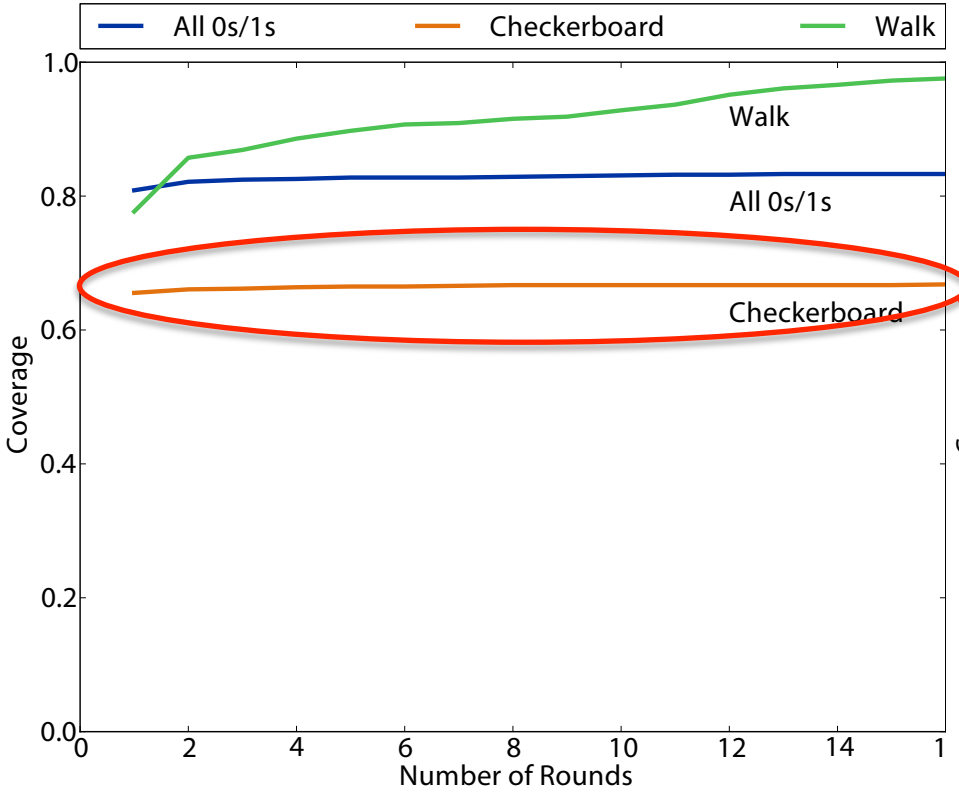

A 1Gb chip family

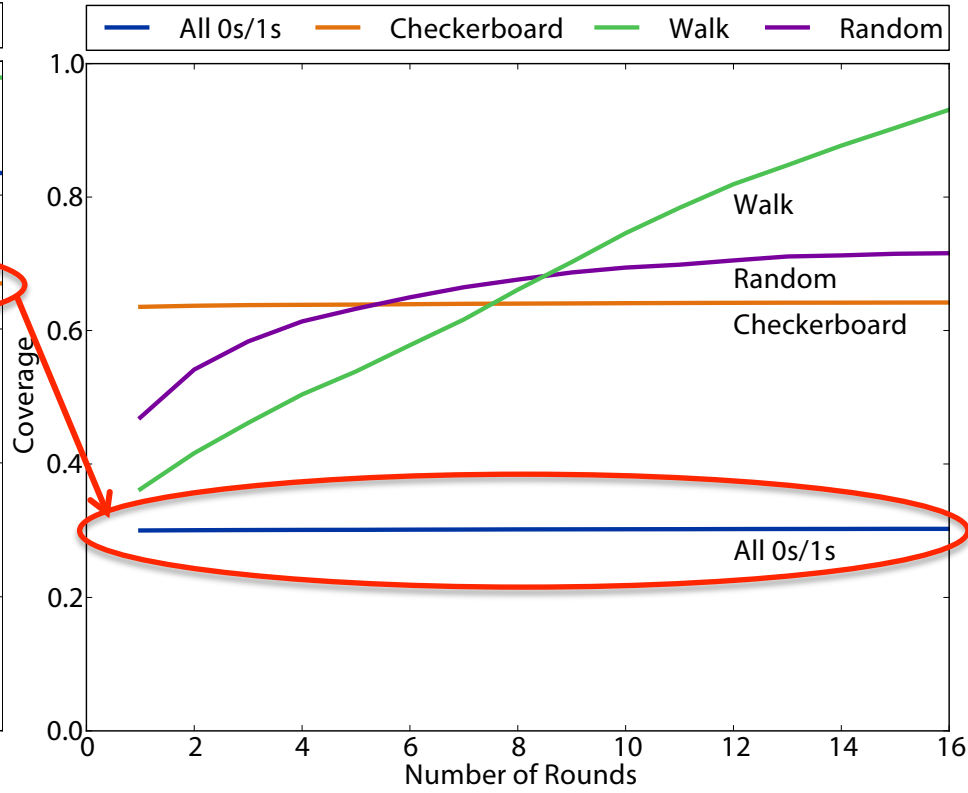

A 2Gb chip family

The lowest-coverage data pattern achieves much lower coverage for the smaller technology node

DPD: Implications on Profiling Mechanisms

- Any retention time profiling mechanism must handle data pattern dependence of retention time
- Intuitive approach: Identify the data pattern that induces the worst-case retention time for a particular cell or device

- Problem 1: Very hard to know at the memory controller which bits actually interfere with each other due to
  - Opaque mapping of addresses to physical DRAM geometry → logically consecutive bits may not be physically consecutive
  - Remapping of faulty bitlines/wordlines to redundant ones internally within DRAM

- Problem 2: Worst-case coupling noise is affected by non-obvious second order bitline coupling effects

DPD: Suggestions (for Future Work)

- A mechanism for identifying worst-case data pattern(s) likely requires support from DRAM device
  - DRAM manufacturers might be in a better position to do this
  - But, the ability of the manufacturer to identify and expose the entire retention time profile is limited due to VRT

- An alternative approach: Use random data patterns to increase coverage as much as possible; handle incorrect retention time estimates with ECC
  - Need to keep profiling time in check
  - Need to keep ECC overhead in check

Talk Agenda

- DRAM Refresh: Background and Motivation
- Challenges and Our Goal
- DRAM Characterization Methodology
- Foundational Results
  - Temperature Dependence
  - Retention Time Distribution
- Data Pattern Dependence: Analysis and Implications
- Variable Retention Time: Analysis and Implications
- Conclusions

Variable Retention Time

- Retention time of a cell can vary over time

- A cell can randomly switch between multiple leakage current states due to *Trap-Assisted Gate-Induced Drain Leakage,* which appears to be a random process

  [Yaney+ IEDM 1987, Restle+ IEDM 1992]

An Example VRT Cell

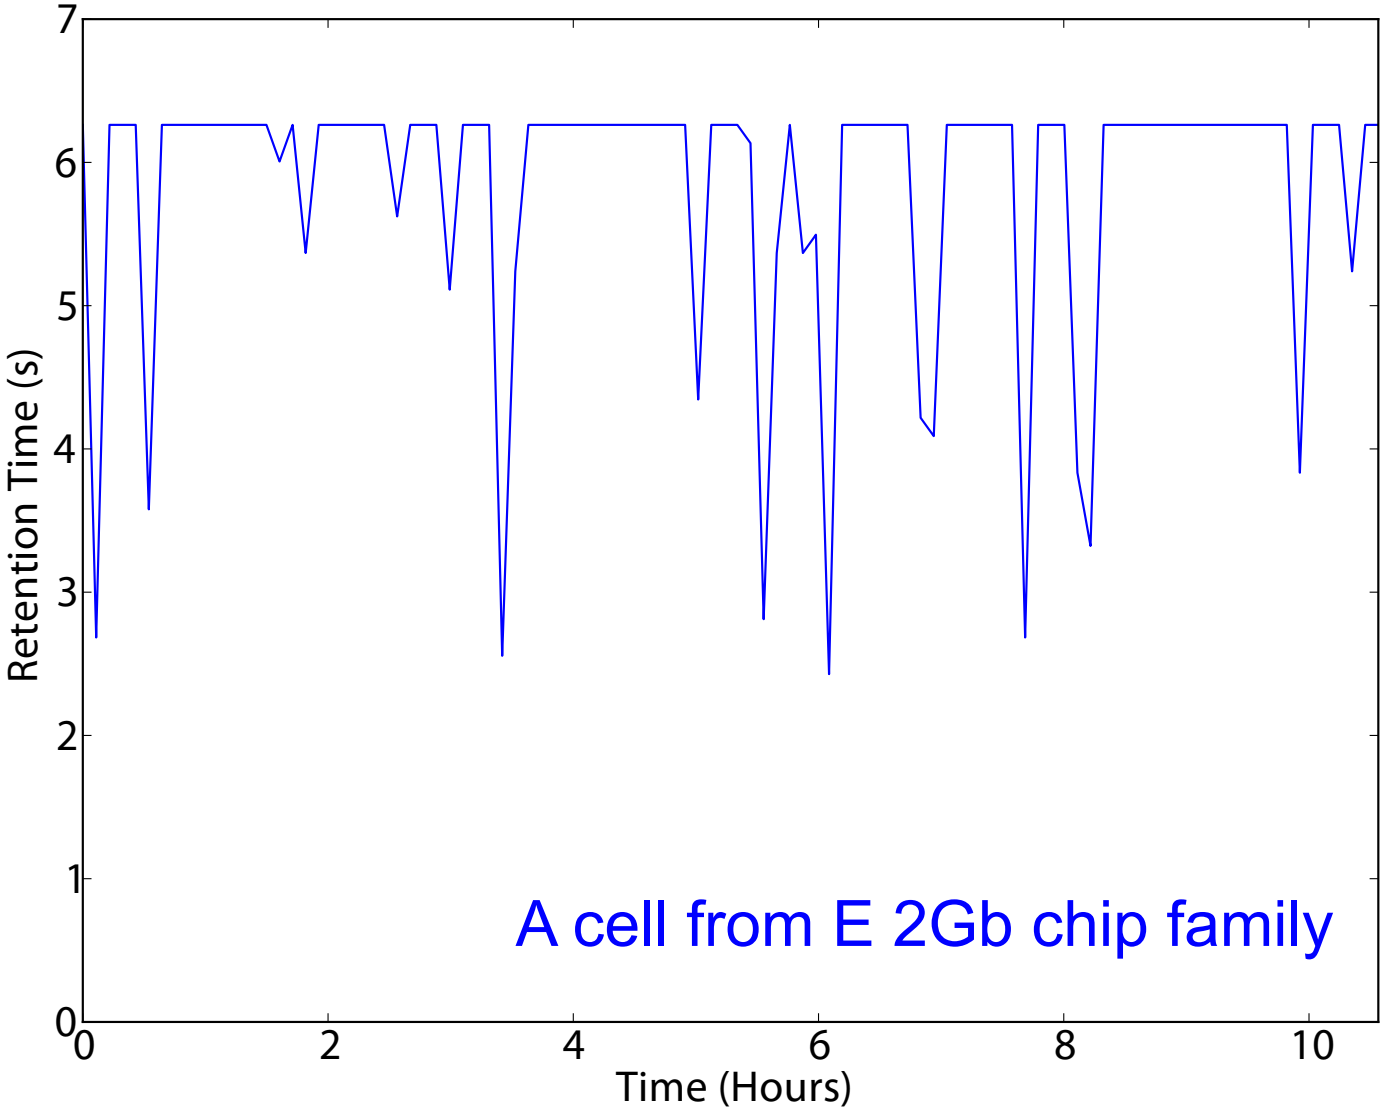

VRT: Questions and Methodology

- Key Questions
  - How prevalent is VRT in modern DRAM devices?
  - What is the timescale of observation of the lowest retention time state?
  - What are the implications on retention time profiling?

- Test Methodology
  - Each device was tested for at least 1024 rounds over 24 hours
  - Temperature fixed at 45^{o}C
  - Data pattern used is the most effective data pattern for each device
  - For each cell that fails at any retention time, we record the minimum and the maximum retention time observed

Variable Retention Time

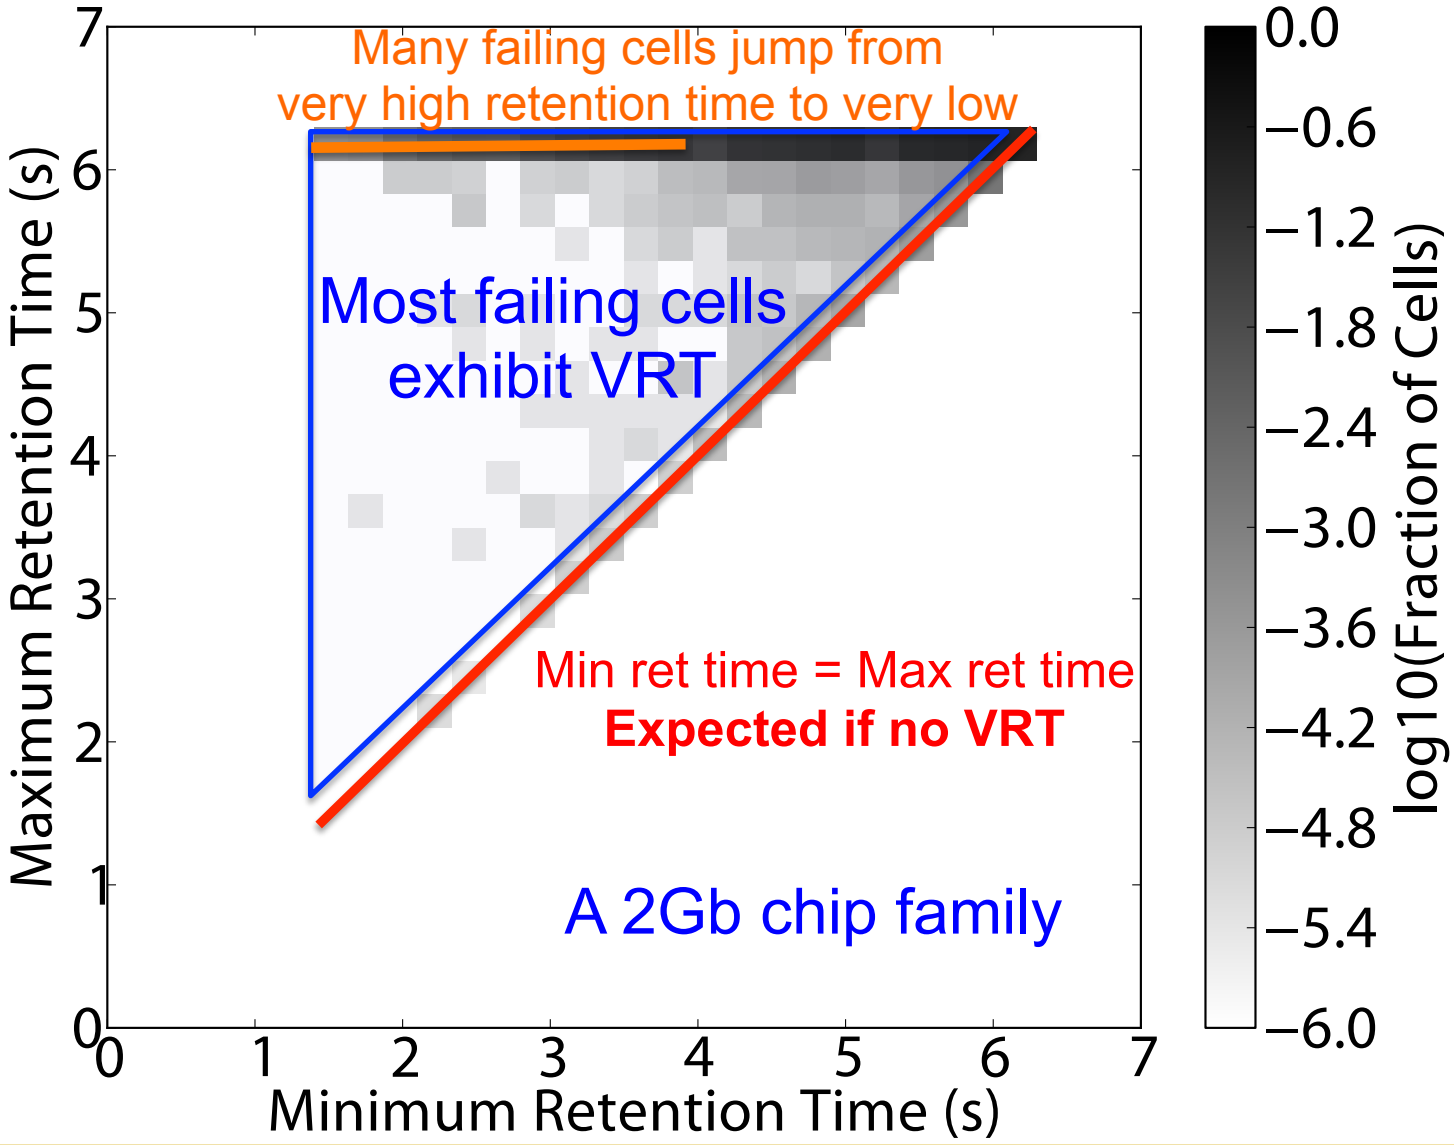

Variable Retention Time

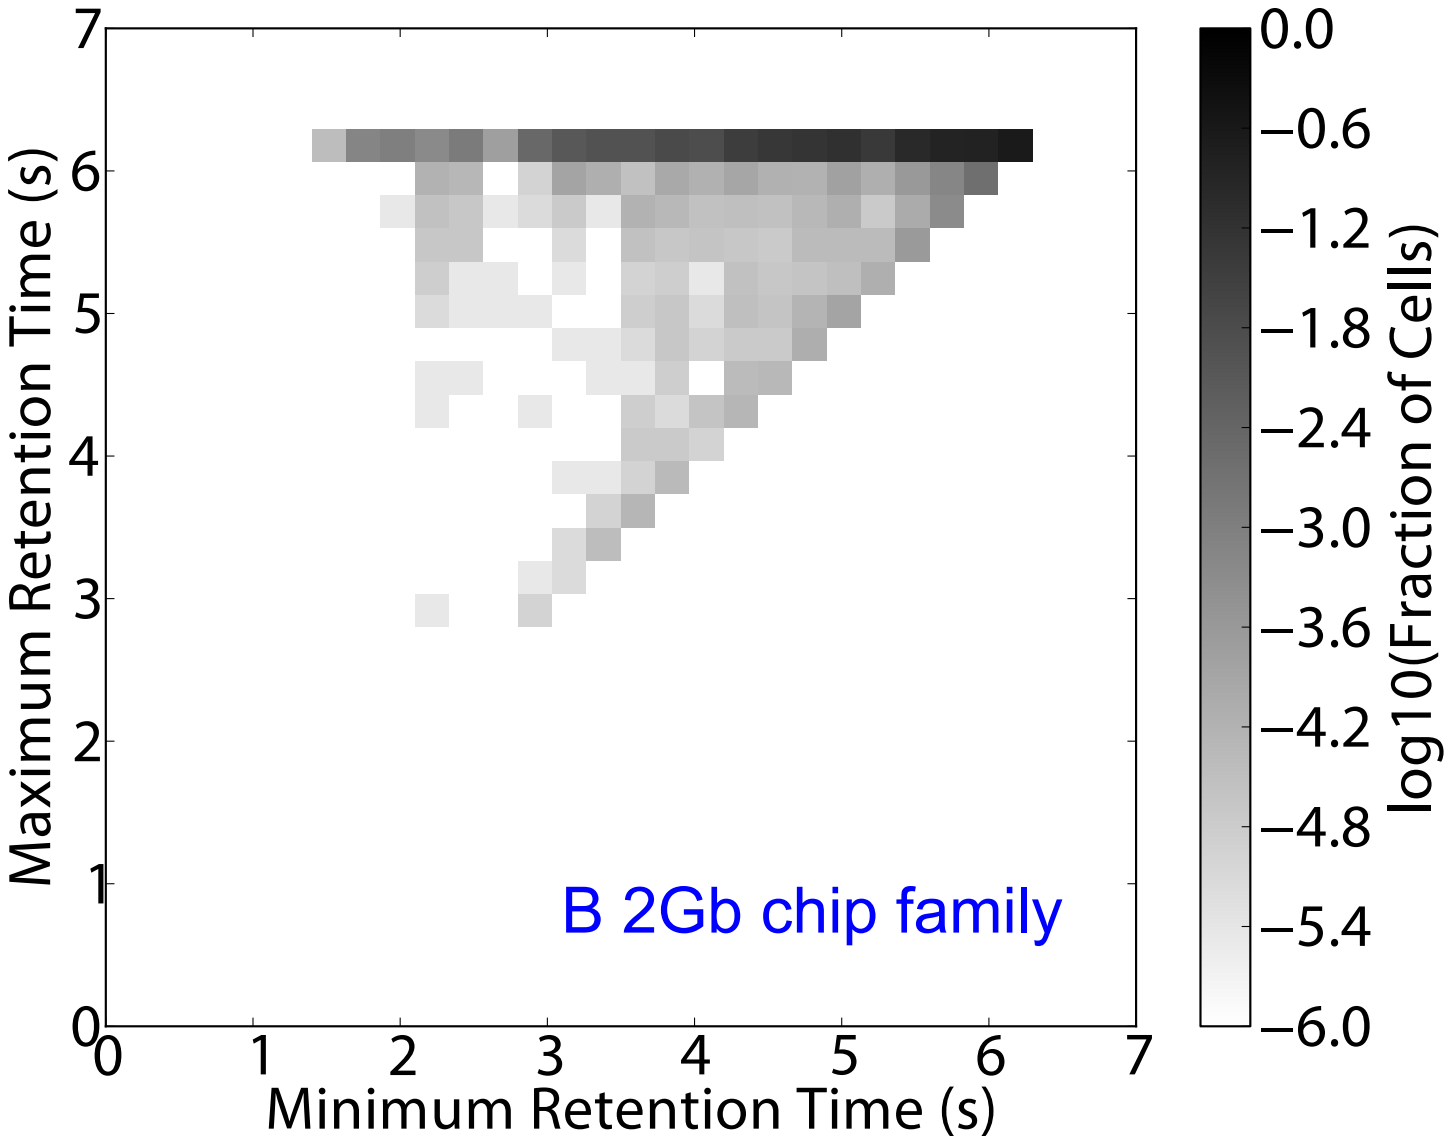

Variable Retention Time

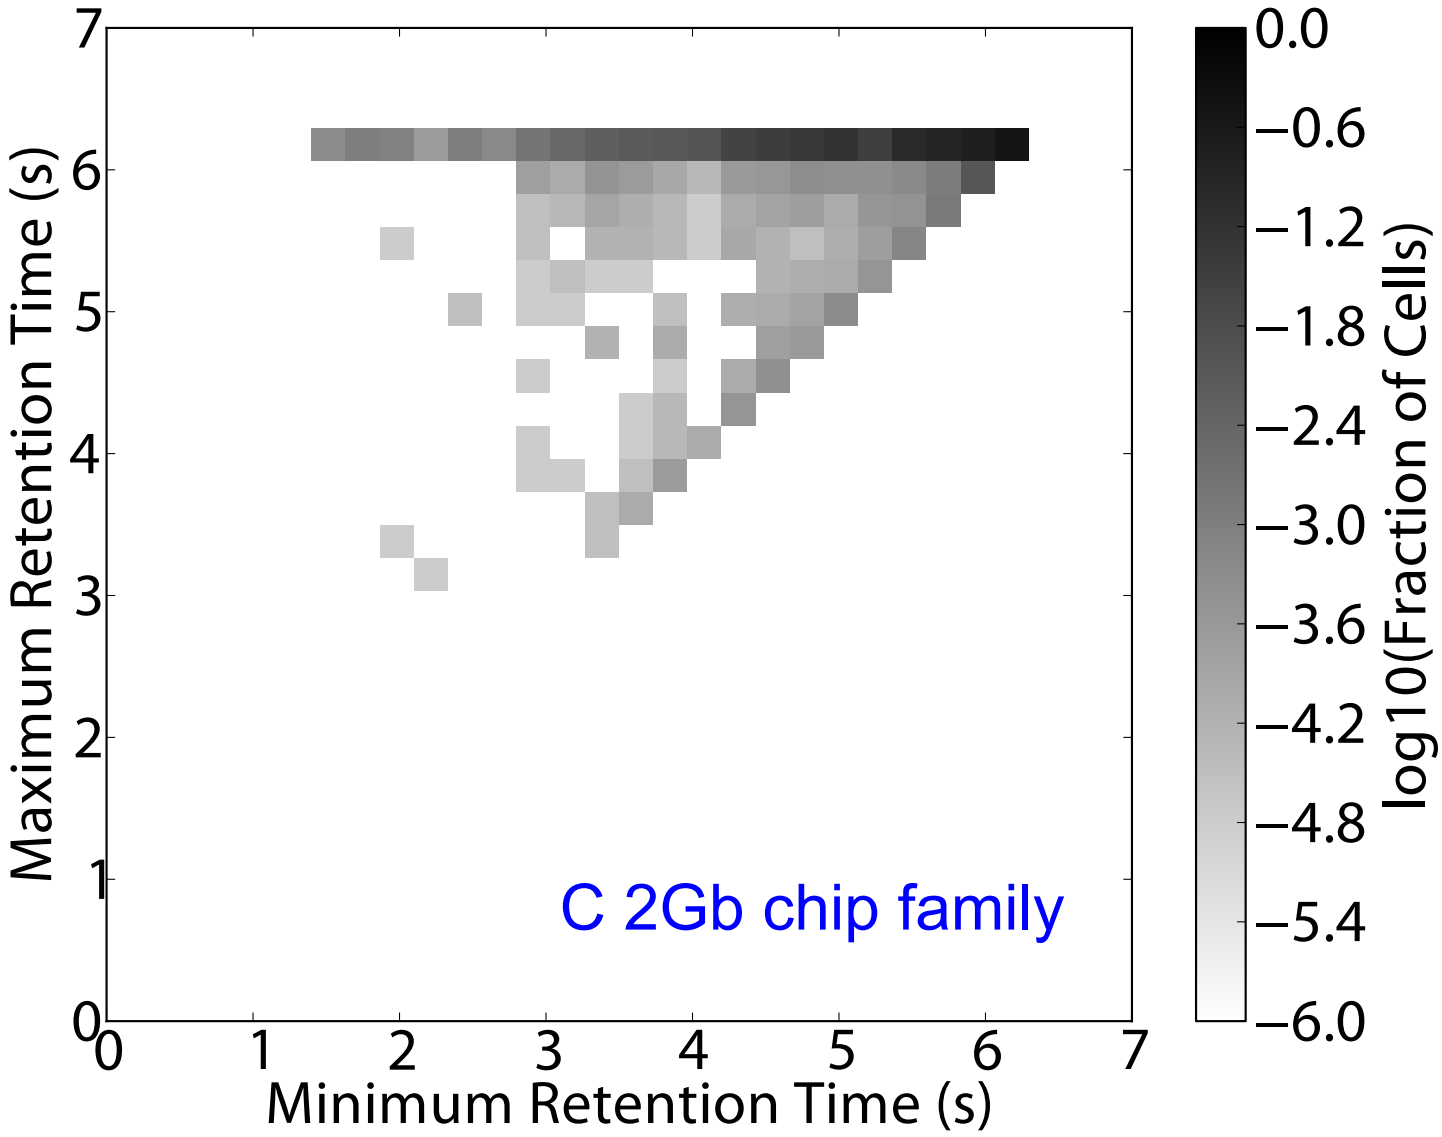

VRT: Observations So Far

- VRT is common among weak cells (i.e., those cells that experience low retention times)

- VRT can result in significant retention time changes
  - Difference between minimum and maximum retention times of a cell can be more than 4x, and may not be bounded
  - Implication: Finding *a* retention time for a cell and using a guardband to ensure minimum retention time is "covered" requires a large guardband or may not work

- Retention time profiling mechanisms must identify lowest retention time in the presence of VRT
  - Question: How long to profile a cell to find its lowest retention time state?

Time Between Retention Time State Changes

- How much time does a cell spend in a high retention state before switching to the minimum observed retention time state?

Time Spent in High Retention Time State

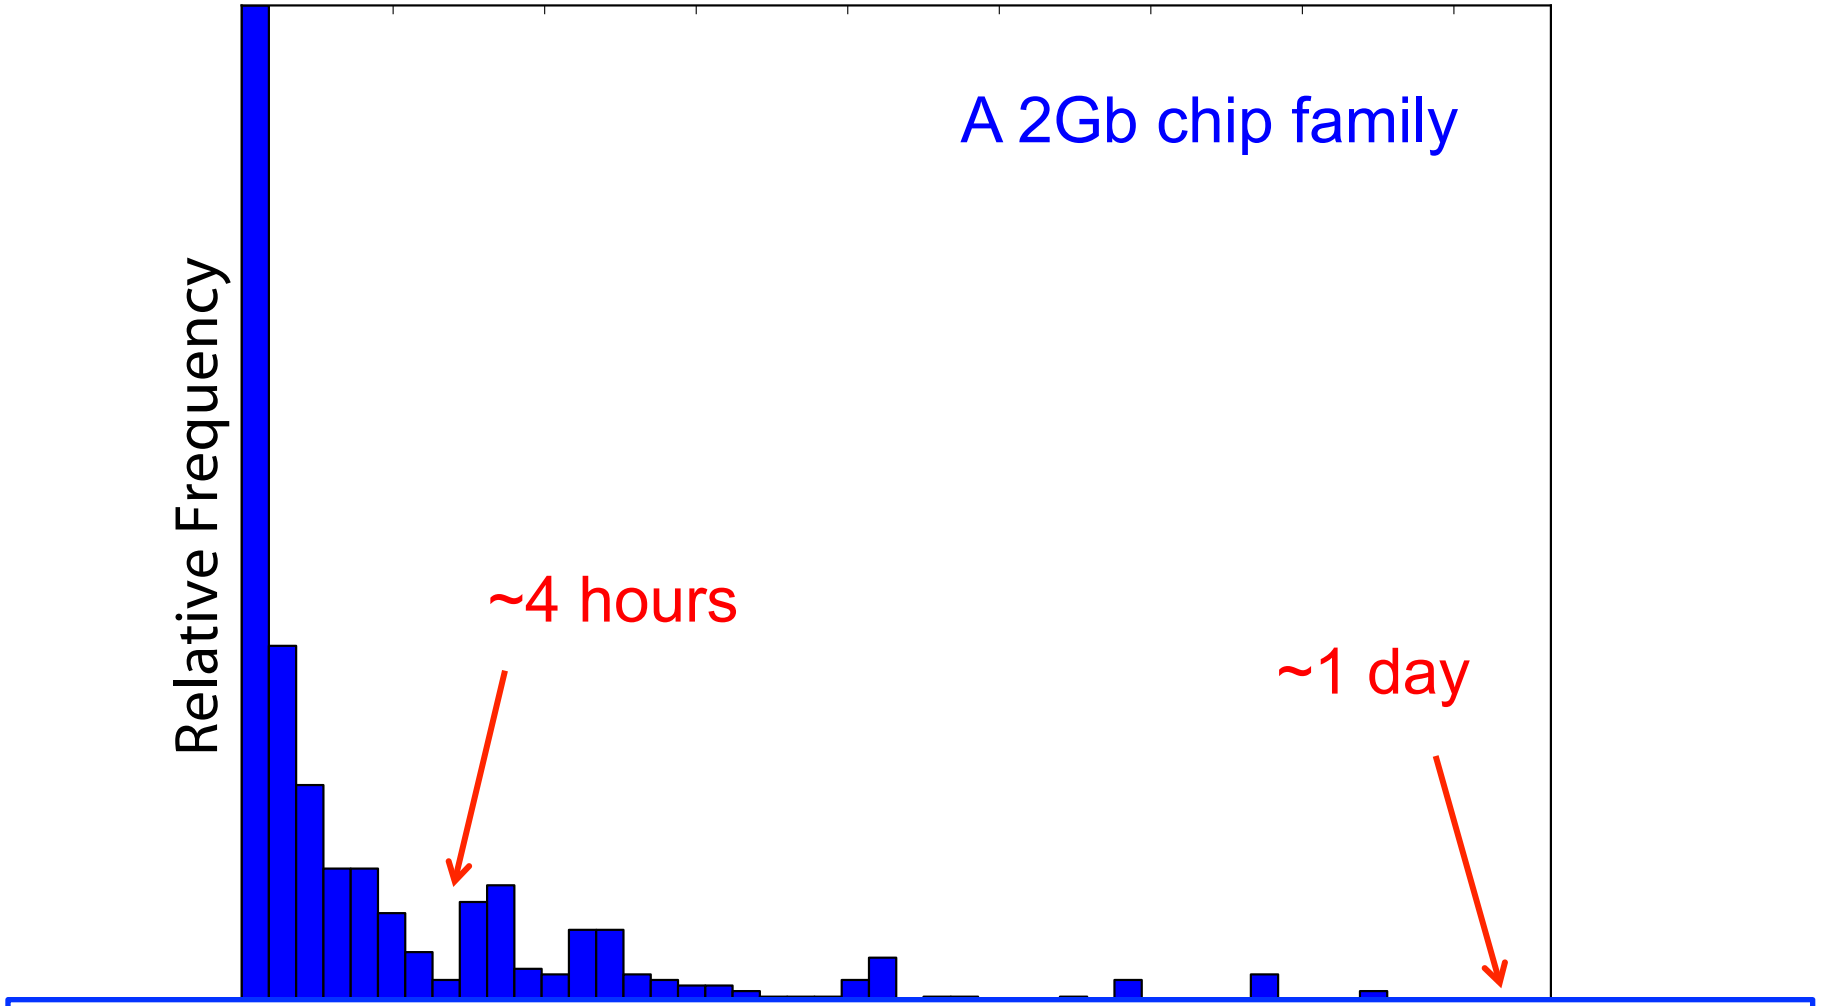

Need to profile for a long time to get to the minimum retention time state

Time Spent in High Retention Time State

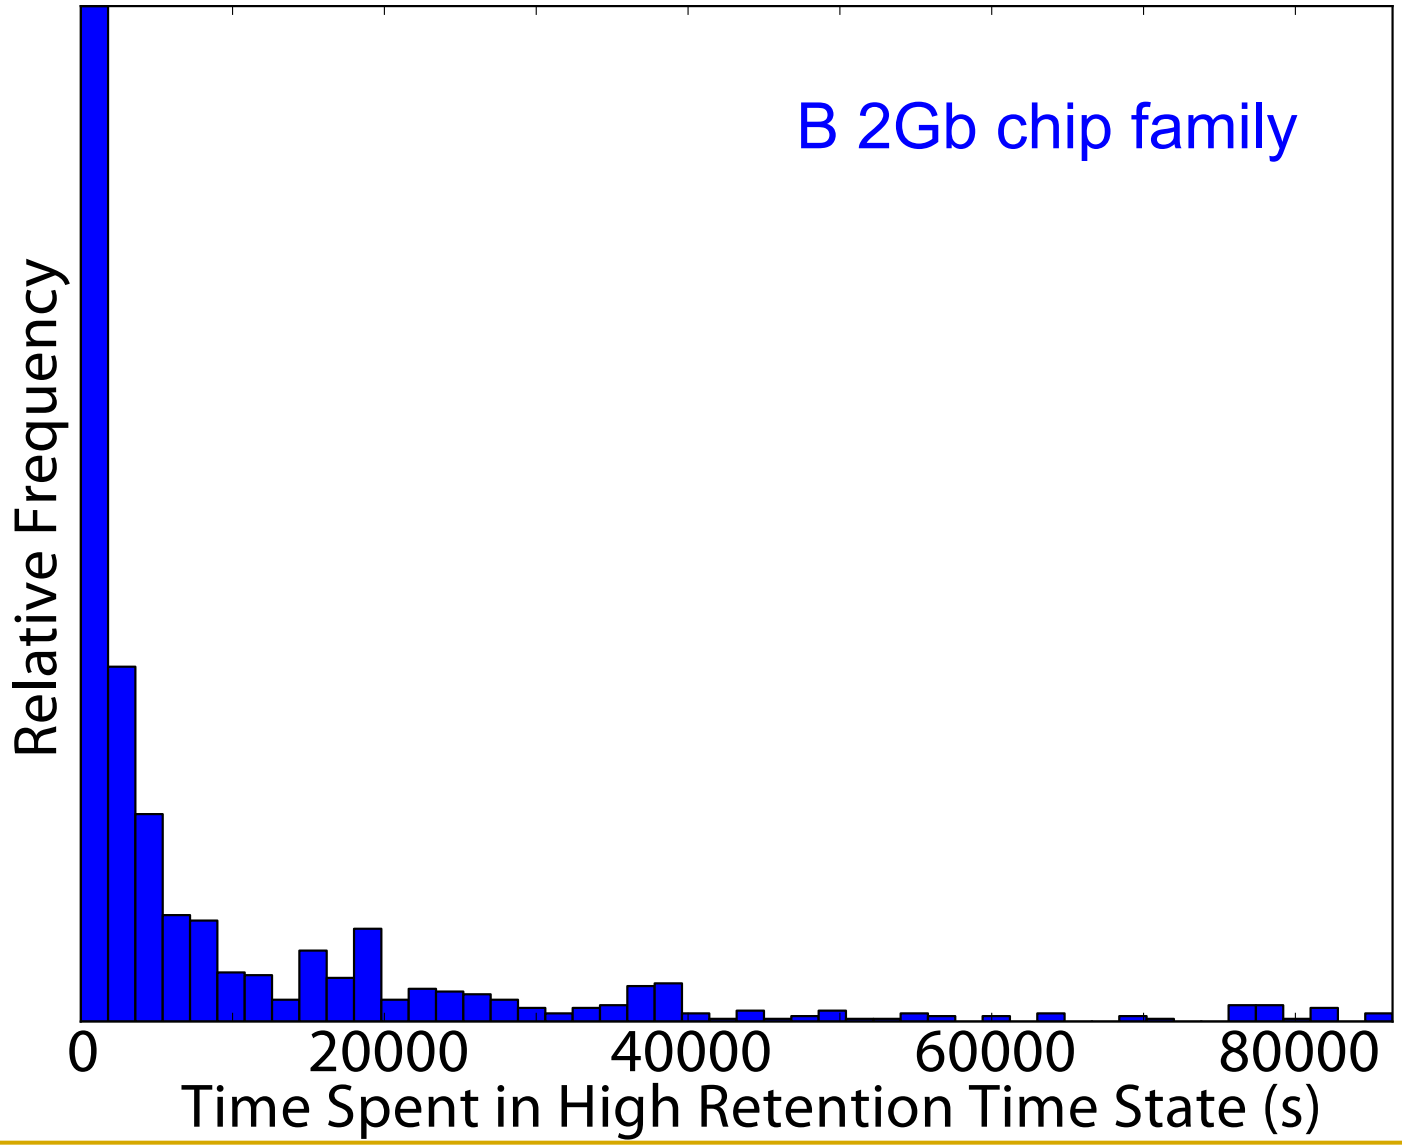

Time Spent in High Retention Time State

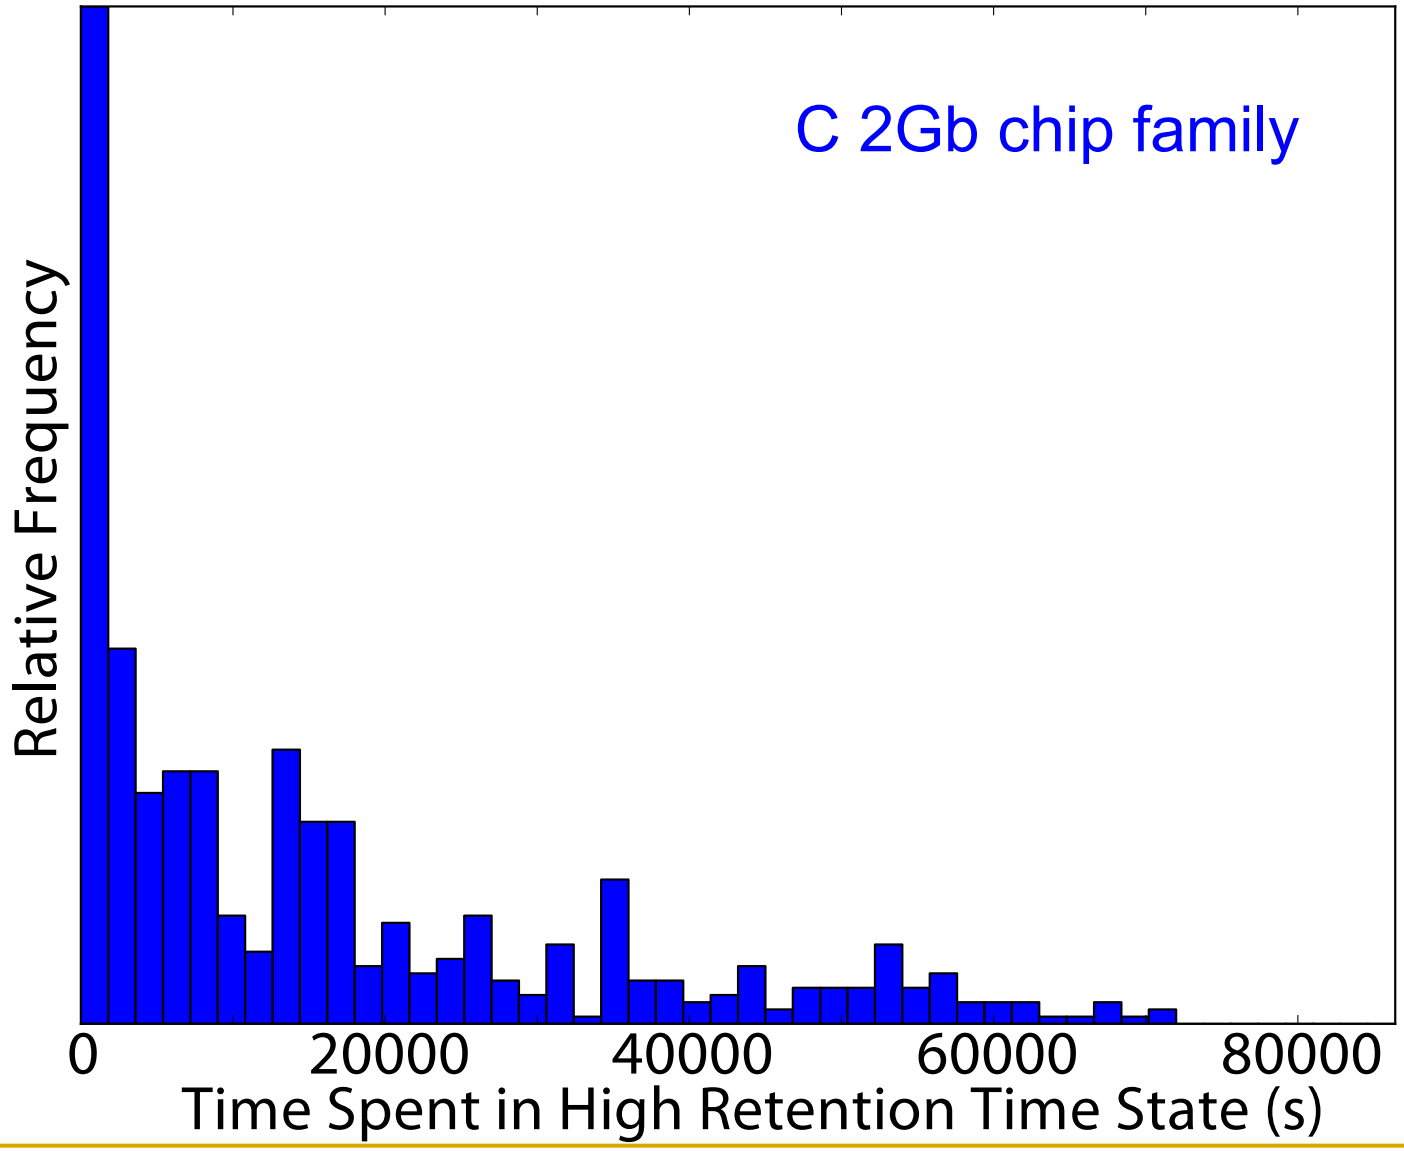

VRT: Implications on Profiling Mechanisms

- Problem 1: There does not seem to be a way of determining if a cell exhibits VRT without actually observing a cell exhibiting VRT
  - VRT is a memoryless random process **[Kim+ JJAP 2010]**
- Problem 2: VRT complicates retention time profiling by DRAM manufacturers
  - Exposure to very high temperatures can induce VRT in cells that were not previously susceptible

    → can happen during soldering of DRAM chips

    → manufacturer's retention time profile may not be accurate
- One option for future work: Use ECC to continuously profile DRAM online while aggressively reducing refresh rate
  - Need to keep ECC overhead in check

Talk Agenda

- DRAM Refresh: Background and Motivation
- Challenges and Our Goal
- DRAM Characterization Methodology
- Foundational Results
  - Temperature Dependence
  - Retention Time Distribution
- Data Pattern Dependence: Analysis and Implications
- Variable Retention Time: Analysis and Implications
- Conclusions

Summary and Conclusions

- DRAM refresh is a critical challenge in scaling DRAM technology efficiently to higher capacities and smaller feature sizes
- Understanding the retention time of modern DRAM devices can enable old or new methods to reduce the impact of refresh
  - Many mechanisms require accurate and reliable retention time profiles
- We presented the first work that comprehensively examines data retention behavior in modern commodity DRAM devices
  - Characterized 248 devices from five manufacturers
- Key findings: Retention time of a cell significantly depends on data pattern stored in other cells (data pattern dependence) and changes over time via a random process (variable retention time)
  - Discussed the underlying reasons and provided suggestions
- Future research on retention time profiling should solve the challenges posed by the DPD and VRT phenomena

An Experimental Study of Data Retention Behavior in Modern DRAM Devices

Implications for Retention Time Profiling Mechanisms

Jamie Liu[1] Ben Jaiyen[1] Yoongu Kim[1]
Chris Wilkerson[2] Onur Mutlu[1]

[1] Carnegie Mellon University
[2] Intel Corporation

SAFARI

Carnegie Mellon

Summary (I)

- DRAM requires periodic refresh to avoid data loss
  - Refresh wastes energy, reduces performance, limits DRAM density scaling
- Many past works observed that different DRAM cells can retain data for different times without being refreshed; proposed reducing refresh rate for strong DRAM cells
  - Problem: These techniques require an accurate profile of the retention time of all DRAM cells
- Our goal: To analyze the retention time behavior of DRAM cells in modern DRAM devices to aid the collection of accurate profile information
- Our experiments: We characterize 248 modern commodity DDR3 DRAM chips from 5 manufacturers using an FPGA based testing platform
- Two Key Issues:

  1. Data Pattern Dependence: A cell's retention time is heavily dependent on data values stored in itself and nearby cells, which cannot easily be controlled.

  2. Variable Retention Time: Retention time of some cells change unpredictably from high to low at large timescales.

Summary (II)

- Key findings on Data Pattern Dependence
  - There is no observed single data pattern that elicits the lowest retention times for a DRAM device → very hard to find this pattern
  - DPD varies between devices due to variation in DRAM array circuit design between manufacturers
  - DPD of retention time gets worse as DRAM scales to smaller feature sizes
- Key findings on Variable Retention Time
  - VRT is common in modern DRAM cells that are weak
  - The timescale at which VRT occurs is very large (e.g., a cell can stay in high retention time state for a day or longer) → finding minimum retention time can take very long
- Future work on retention time profiling must address these issues

Walk Data Pattern

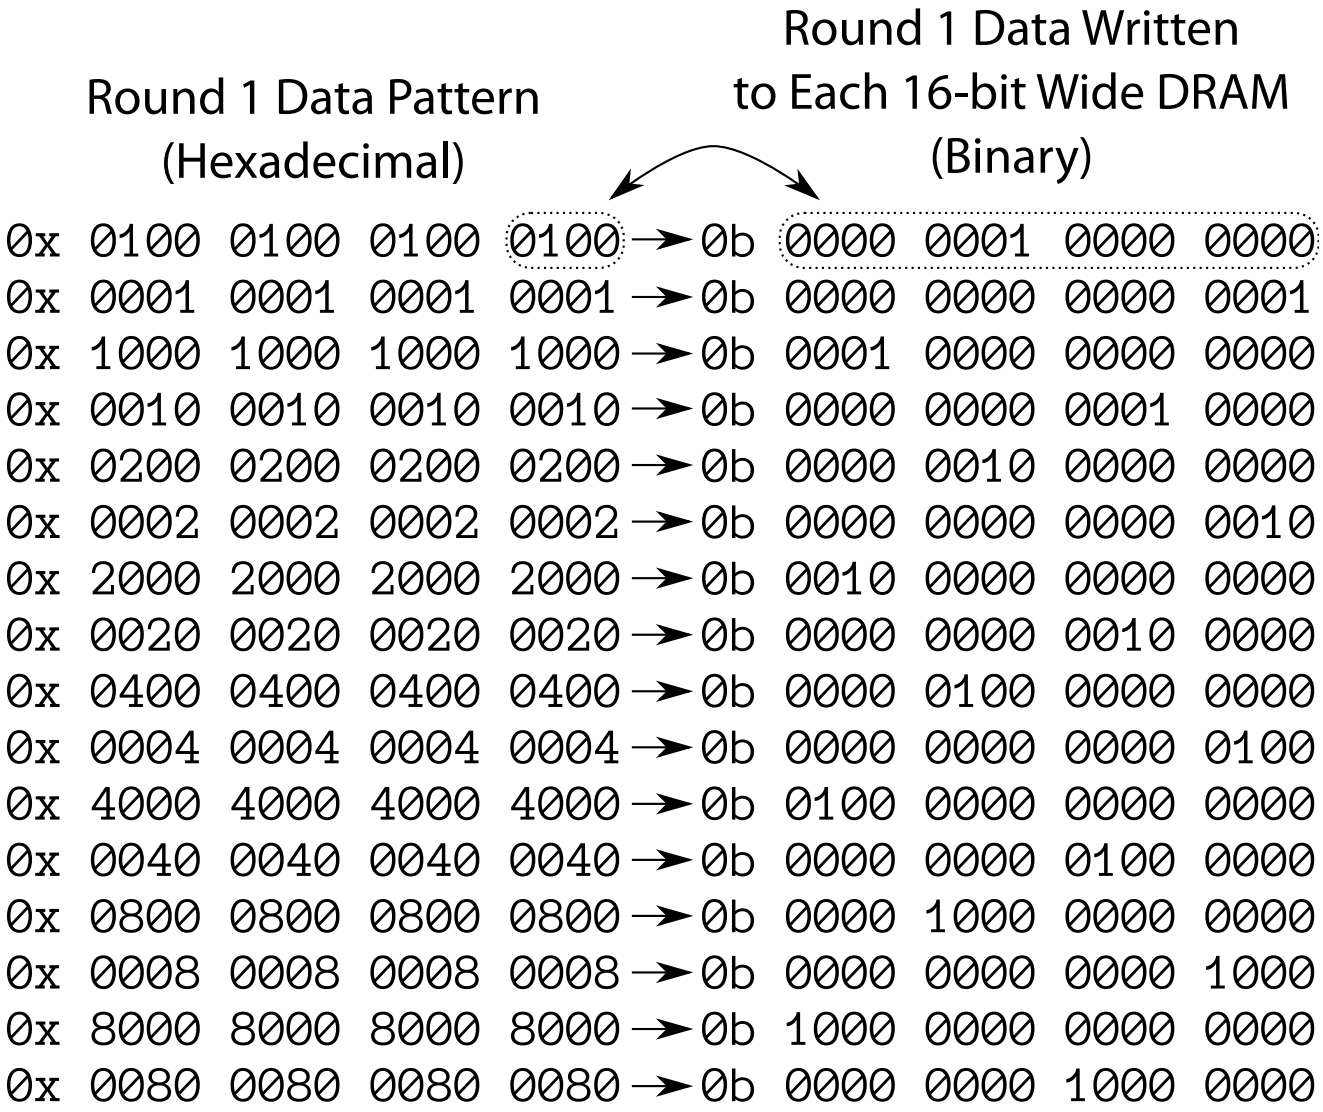

Walk Data Pattern After Round 1

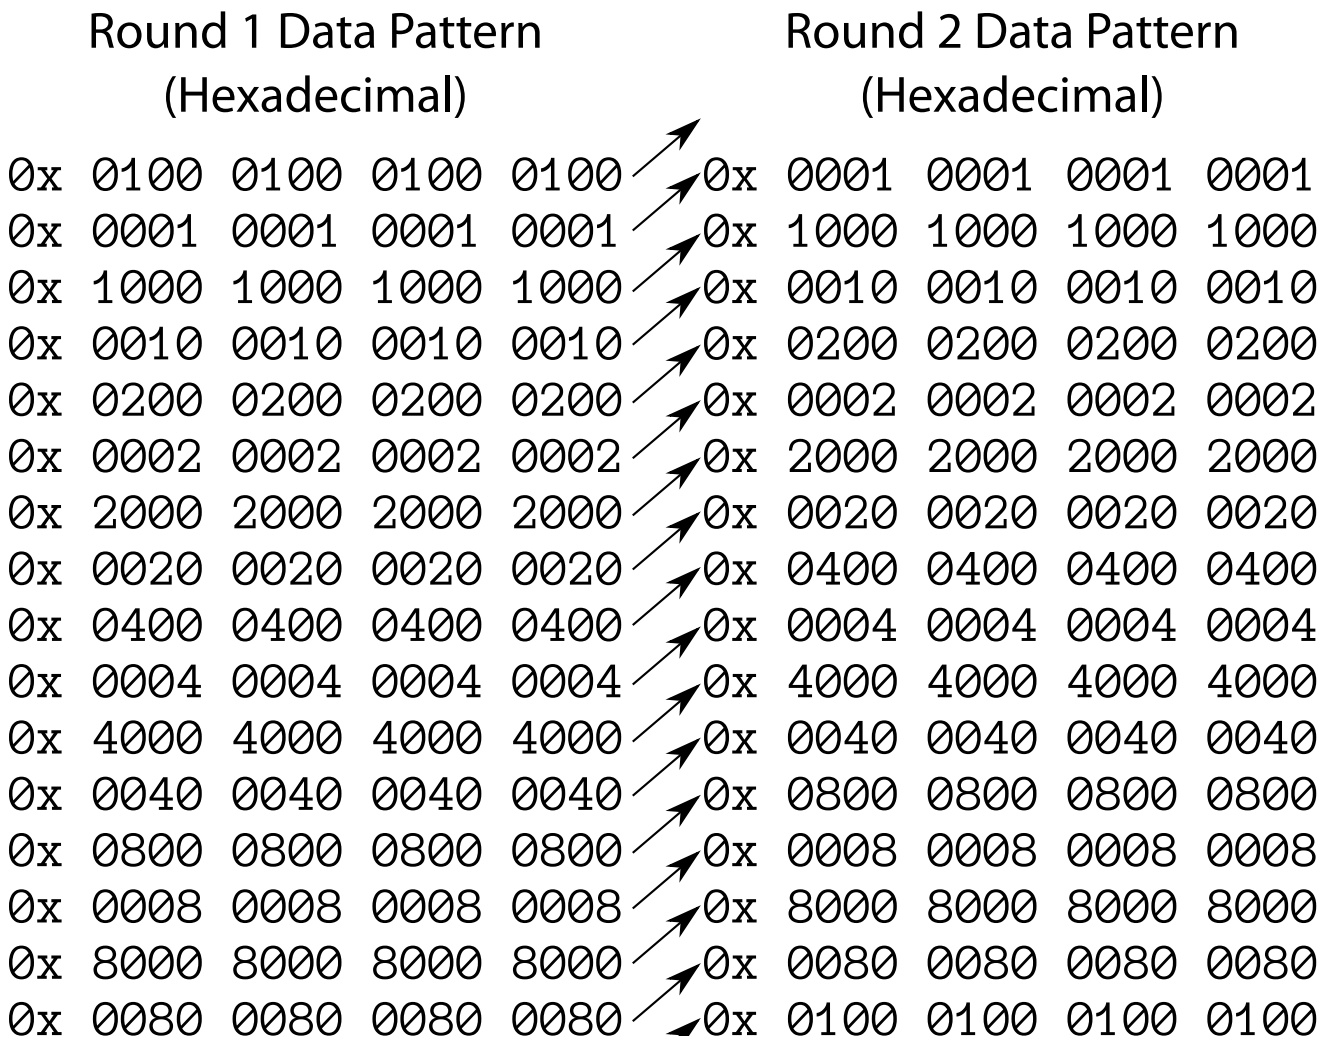

DRAM Activation

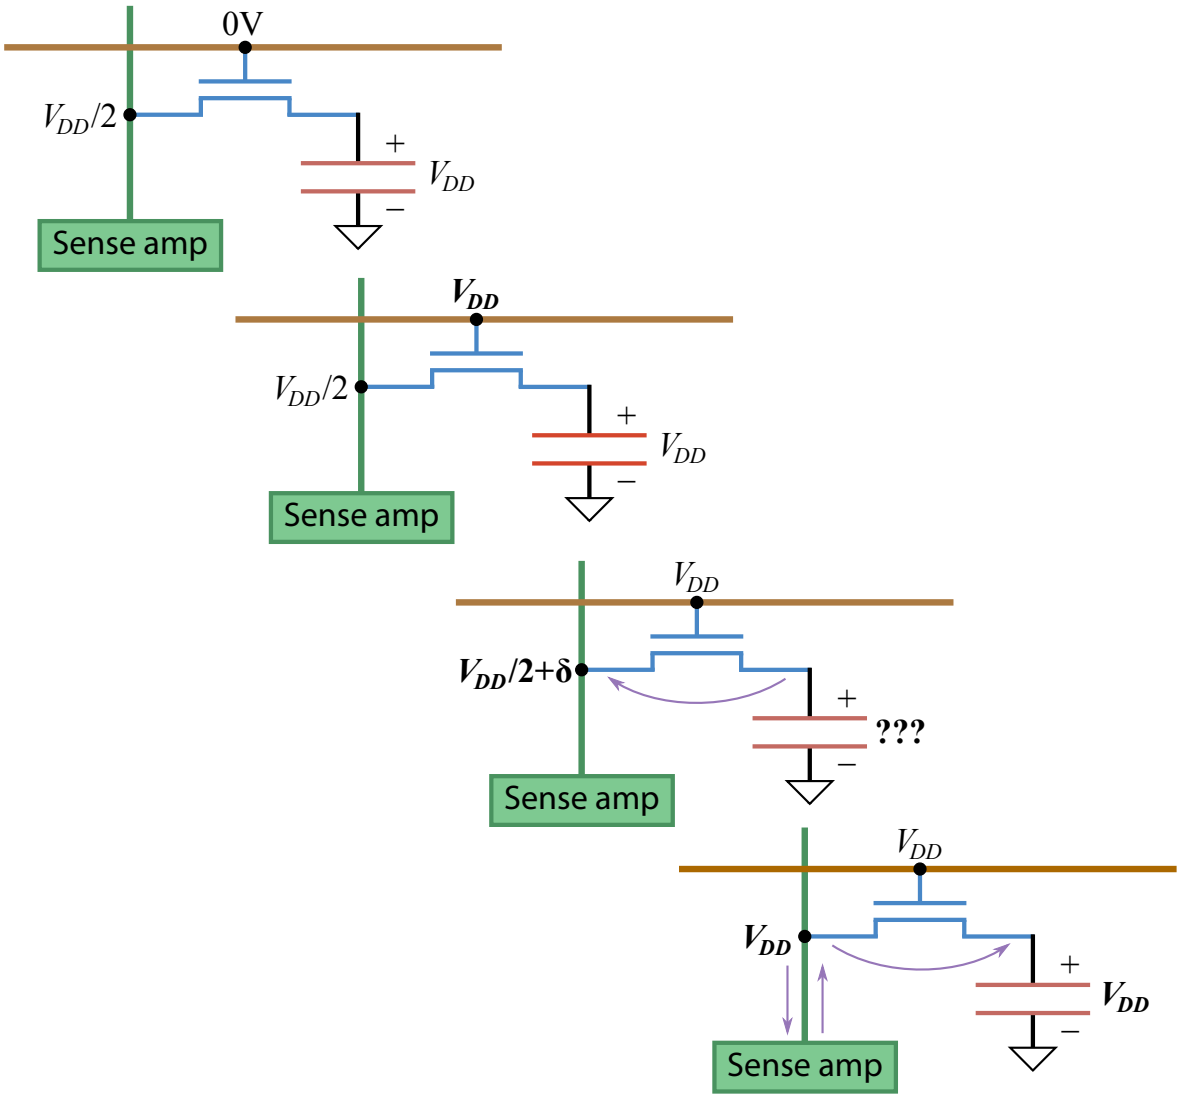

True Cell vs. Anti Cell

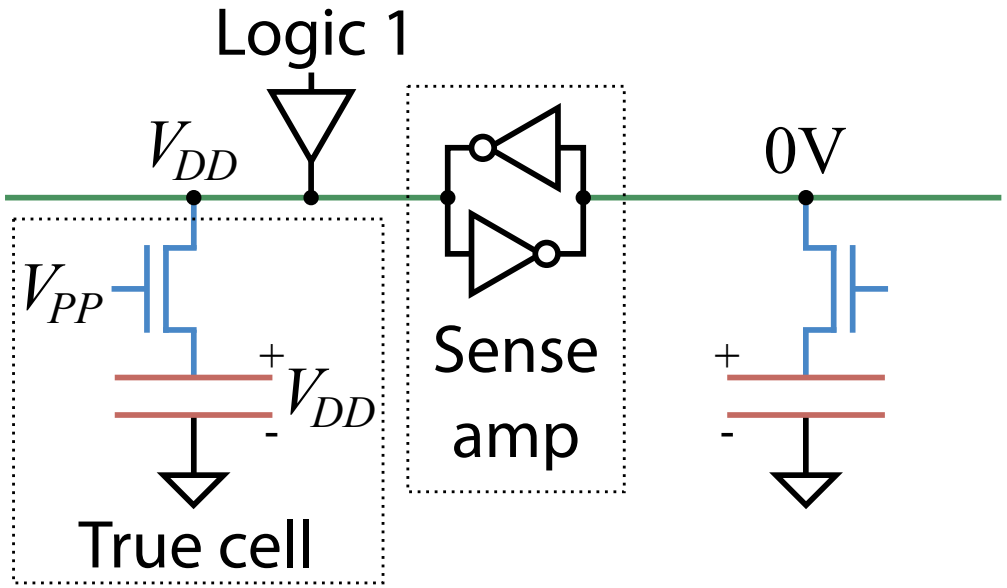

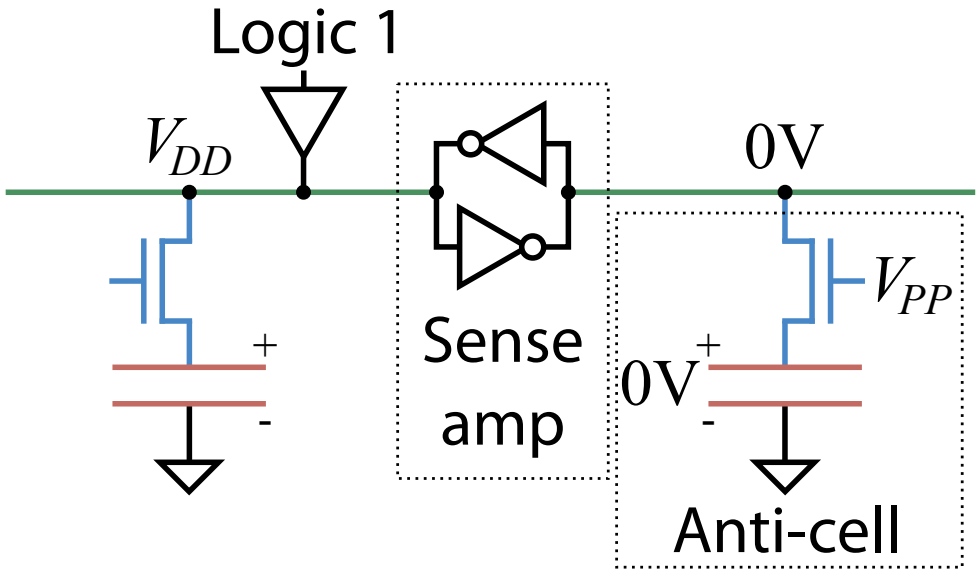

DRAM Organization

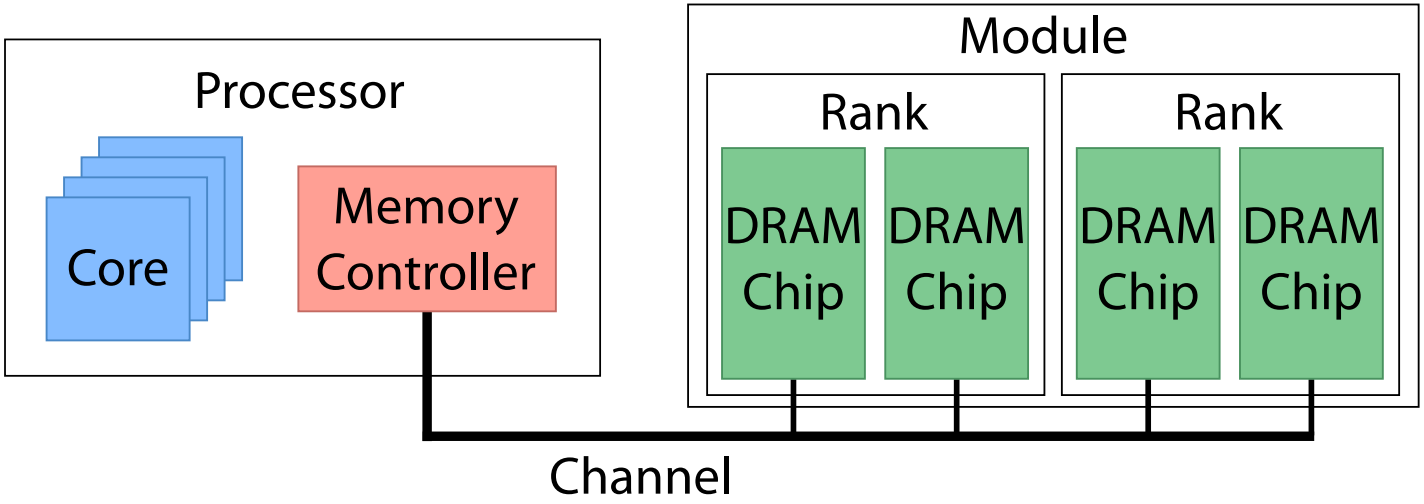

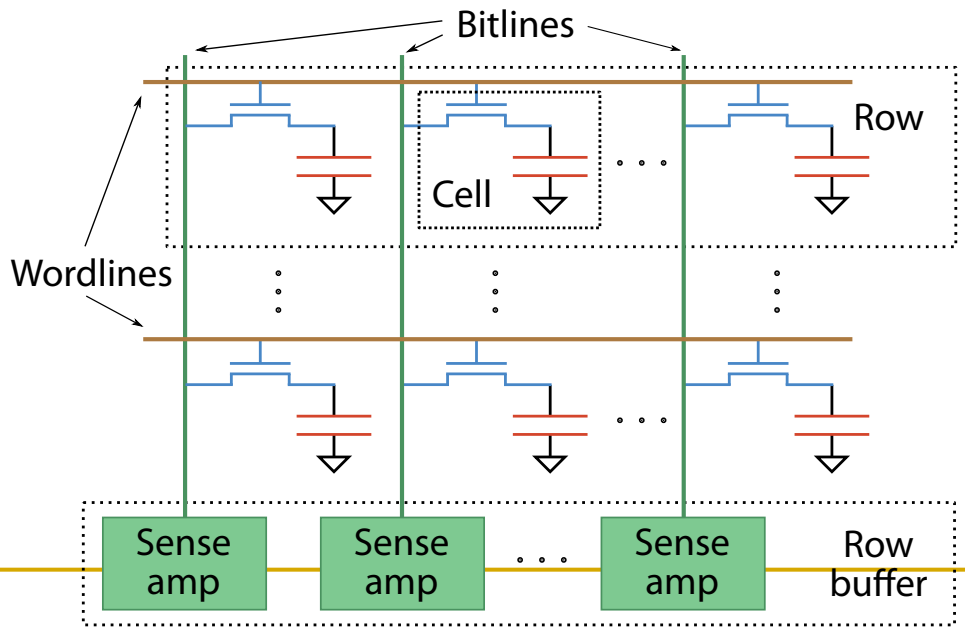

VRT as a Random Process

- Previous work has shown that each VRT cell spends an exponentially distributed amount of time in each state **[Restle+ IEDM 1992, Kim+ JJAP 2010]**, and that the distribution of time constants for these exponential distributions is itself exponentially distributed **[Kim+ IEEE TED 2011]**.

- The shape of our observed distributions appear to be consistent with this prior work.