RAIDR: Retention-Aware Intelligent DRAM Refresh

Jamie Liu Ben Jaiyen Richard Veras Onur Mutlu

Carnegie Mellon University

Presented by Robin Burkhard

14 November 2018

Background, Problem & Goal

Background: DRAM Refresh

- DRAM cells leak over time
  - Rows need to be refreshed periodically
- Cells have different retention times
  - Refresh each row every 64 ms
- Refresh degrades performance
  - Bank unavailable while refreshed
  - Memory accesses have to wait
- Each refresh consumes energy

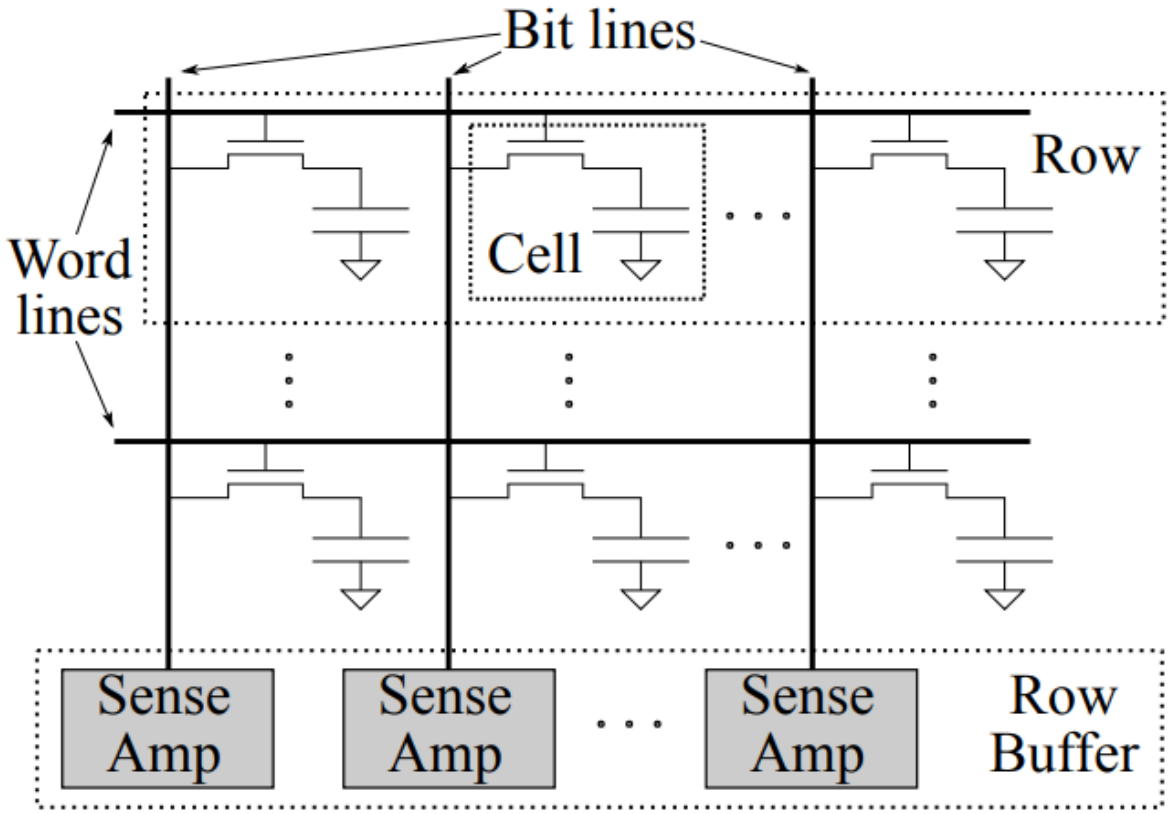

Background: DRAM Refresh

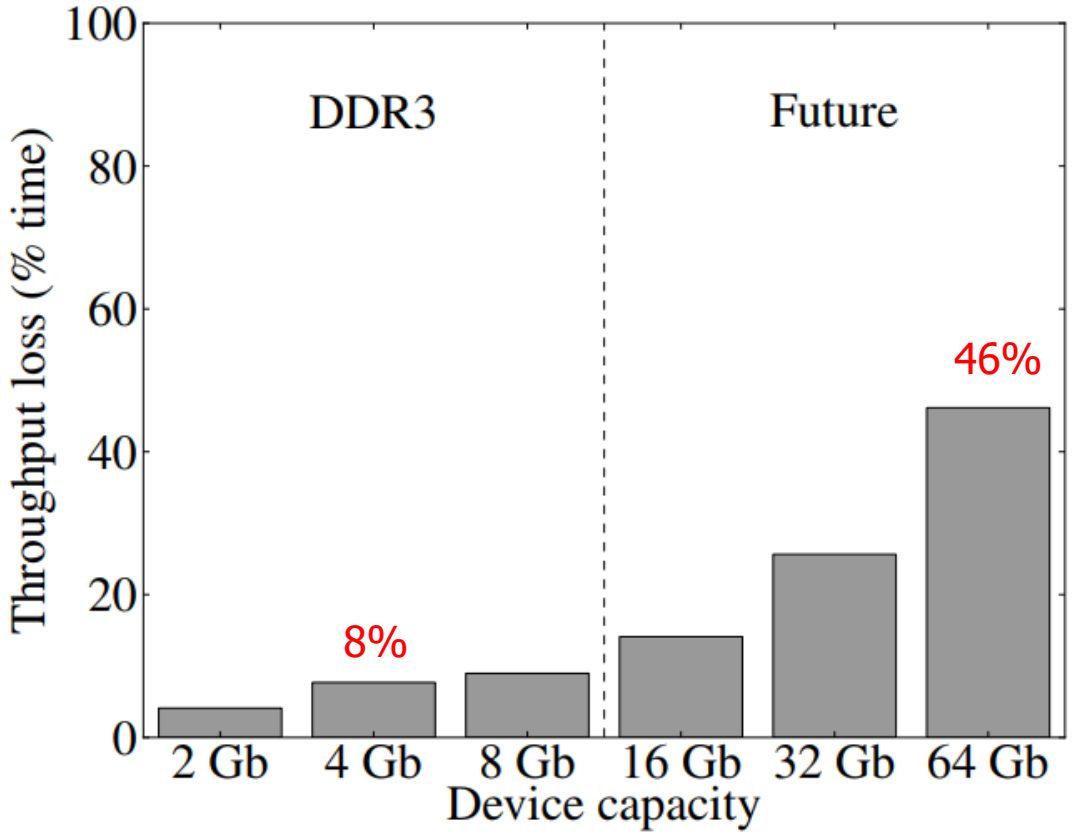

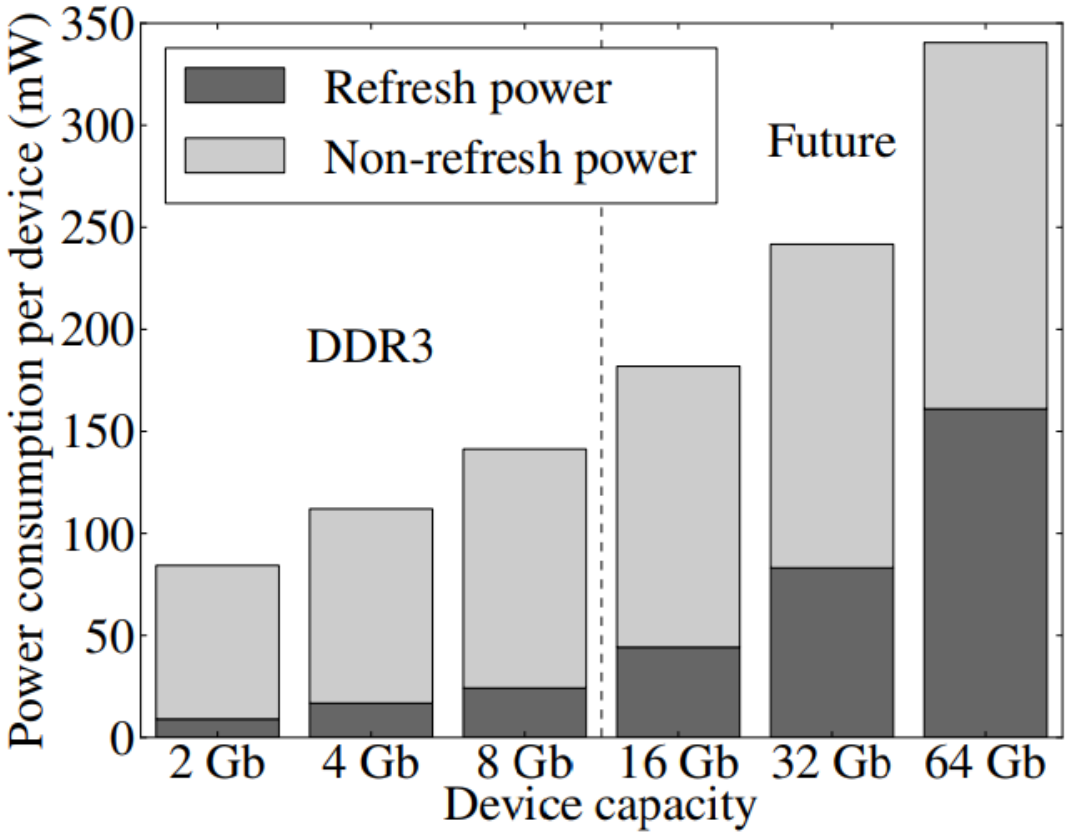

~ 50/50

Retention Time Distribution

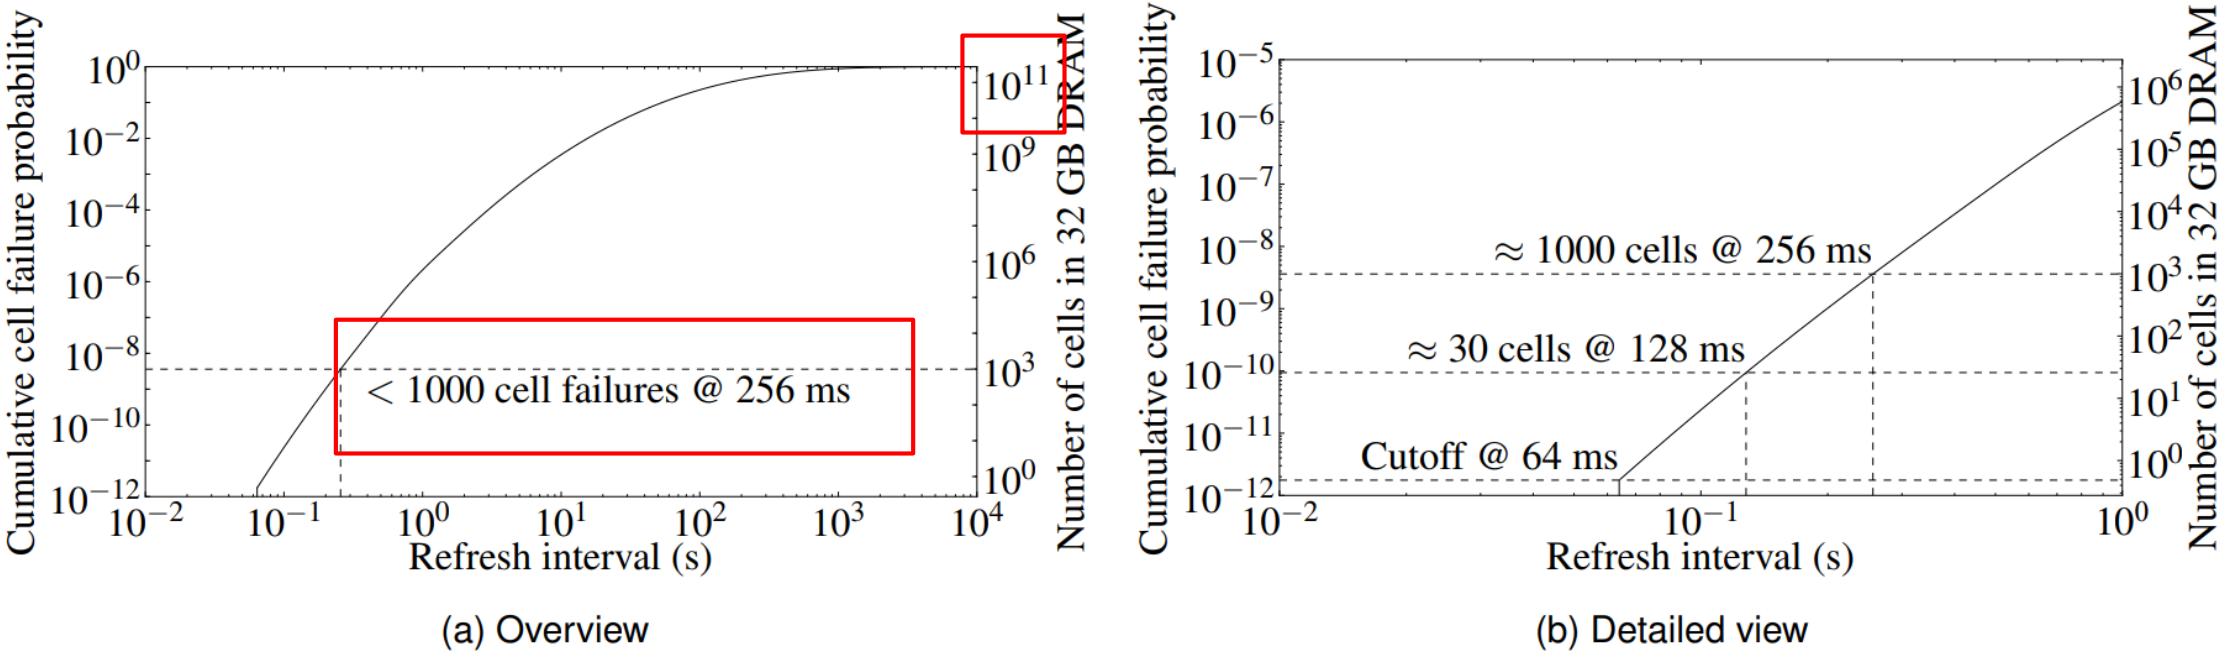

64ms refresh interval is a waste of energy and time caused by very few cells

Idea of RAIDR: Refresh weak rows at high frequency and all other rows at low frequency

Mechanism & Implementation

RAIDR Operation

1. Profiling: Get the retention time for each row

2. Binning: Store rows into different bins based on retention time

3. Refreshing: Use different refresh rates for different bins

RAIDR Profiling

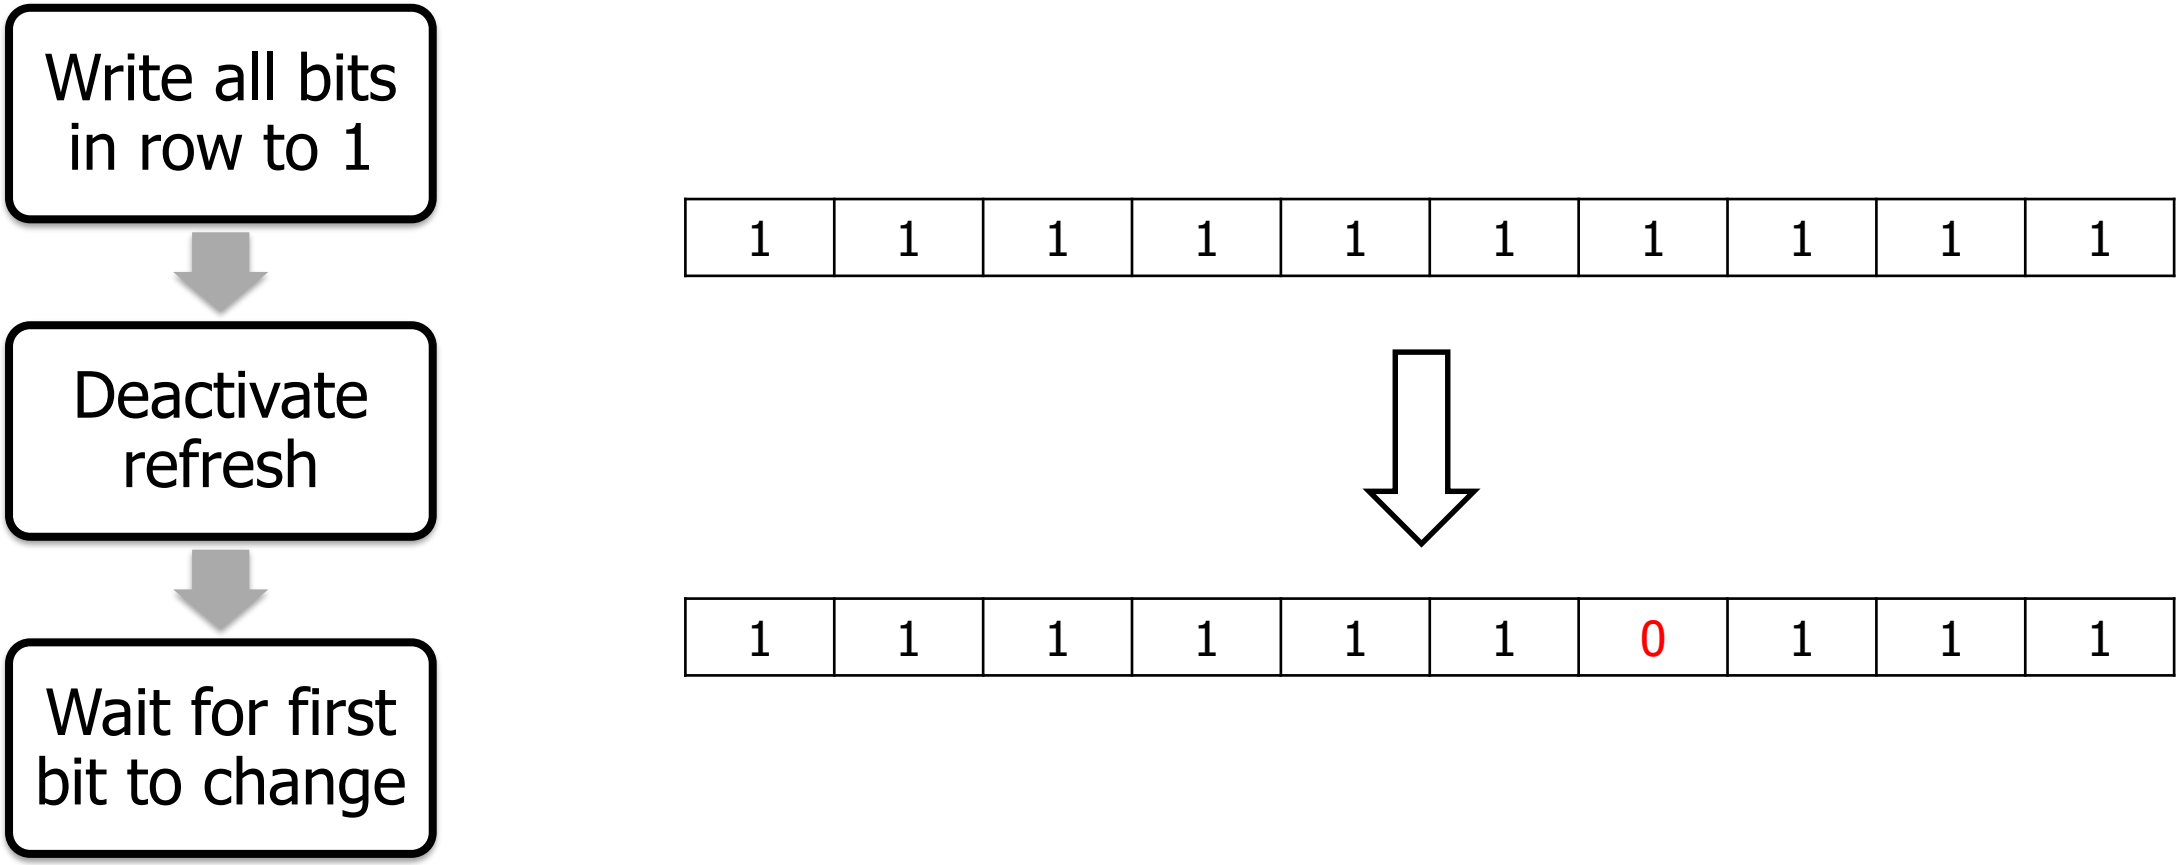

RAIDR Binning

Default Configuration of RAIDR for 32GB DRAM system:

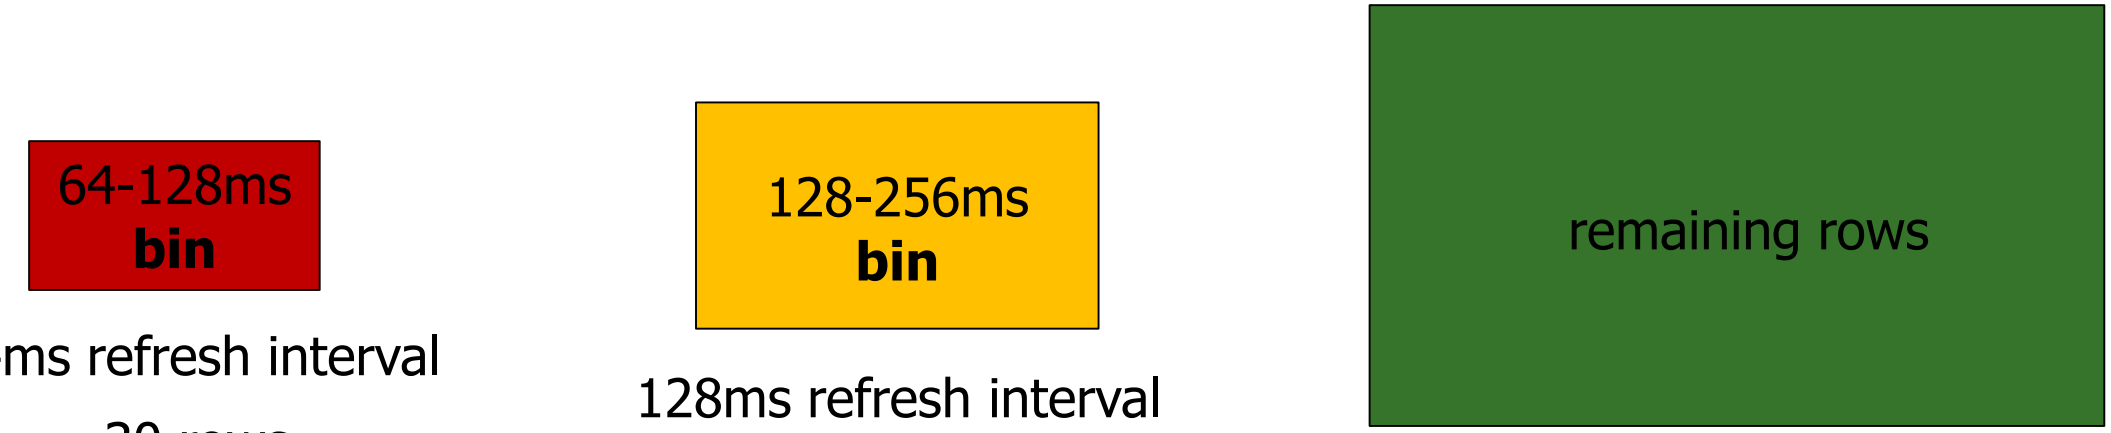

64ms refresh interval

~ 30 rows

128ms refresh interval

~ 1000 rows

256ms refresh interval

~ 4 Million rows

Space-efficient implementation using **Bloom Filters**

RAIDR Binning

16 bits size
3 hash functions

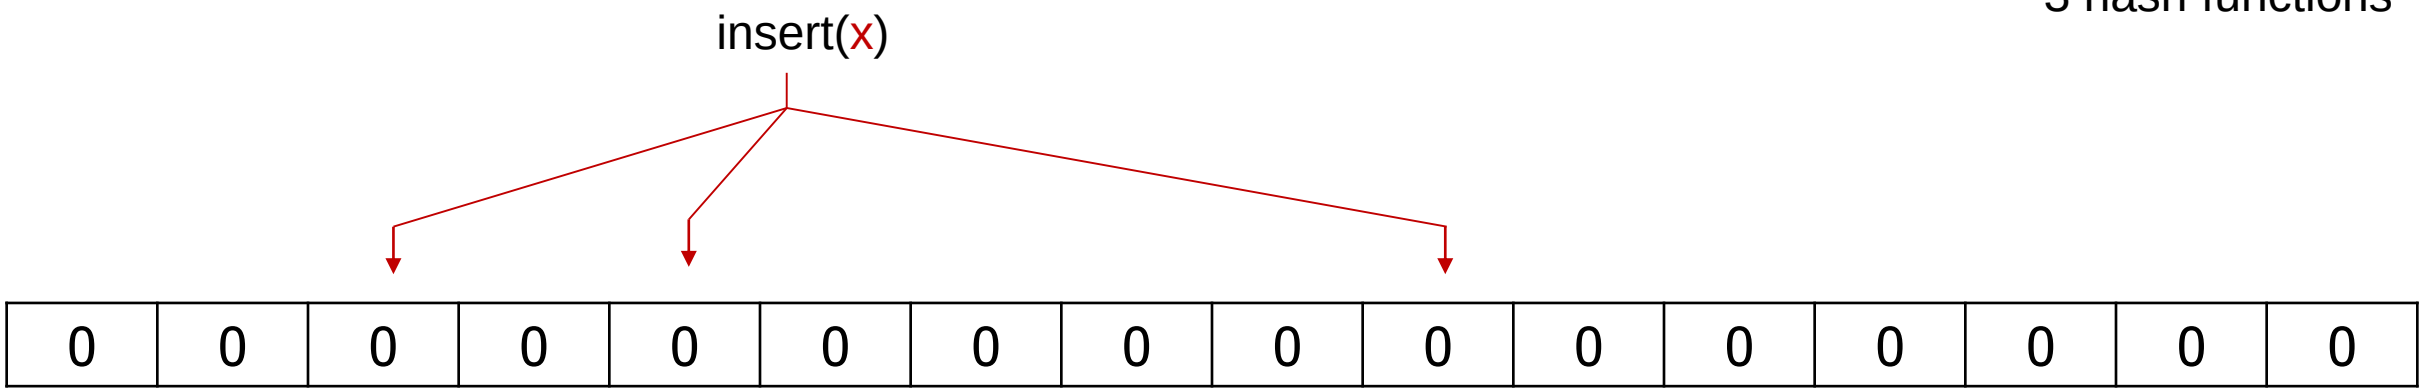

RAIDR Binning

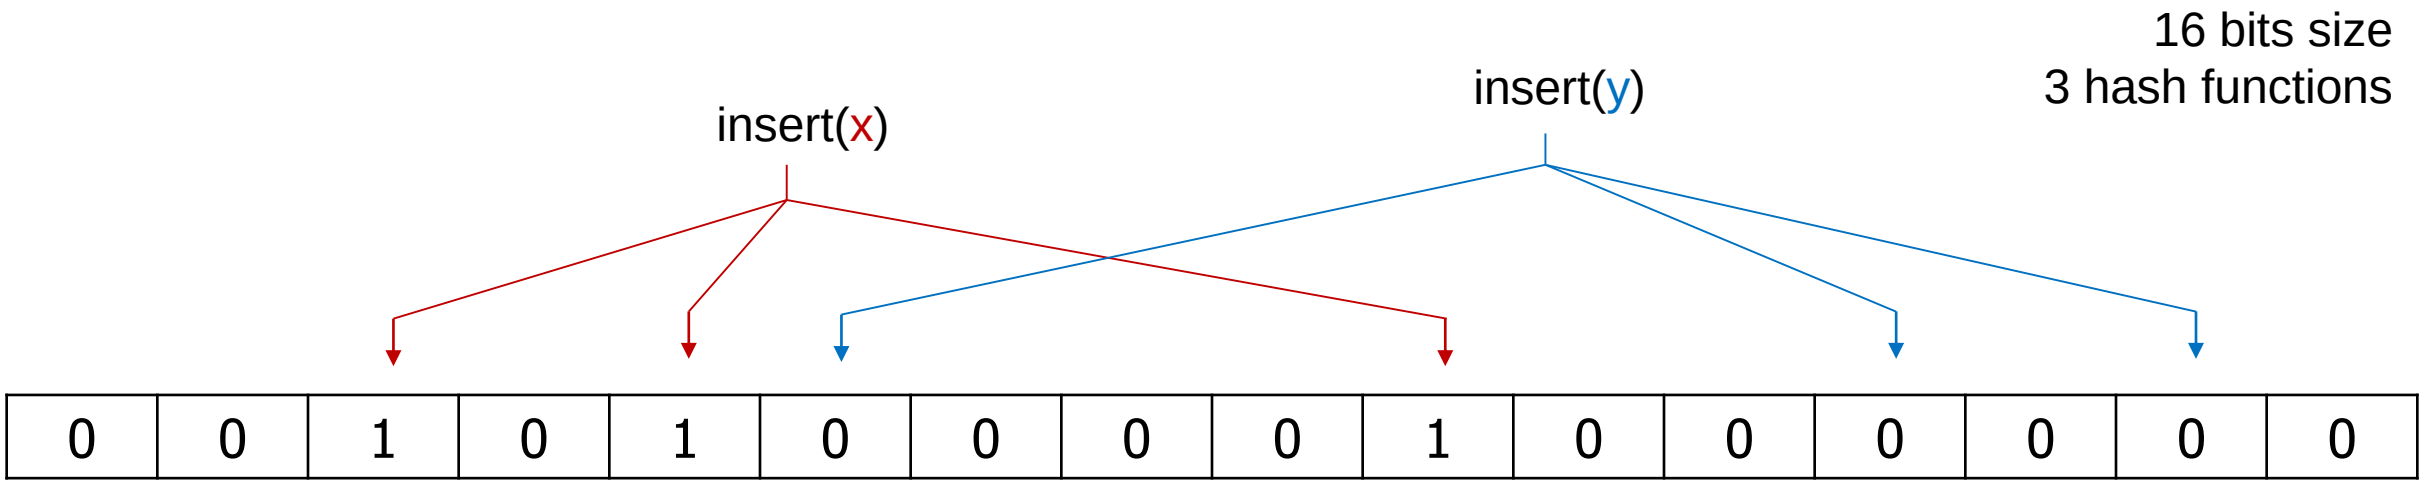

RAIDR Binning

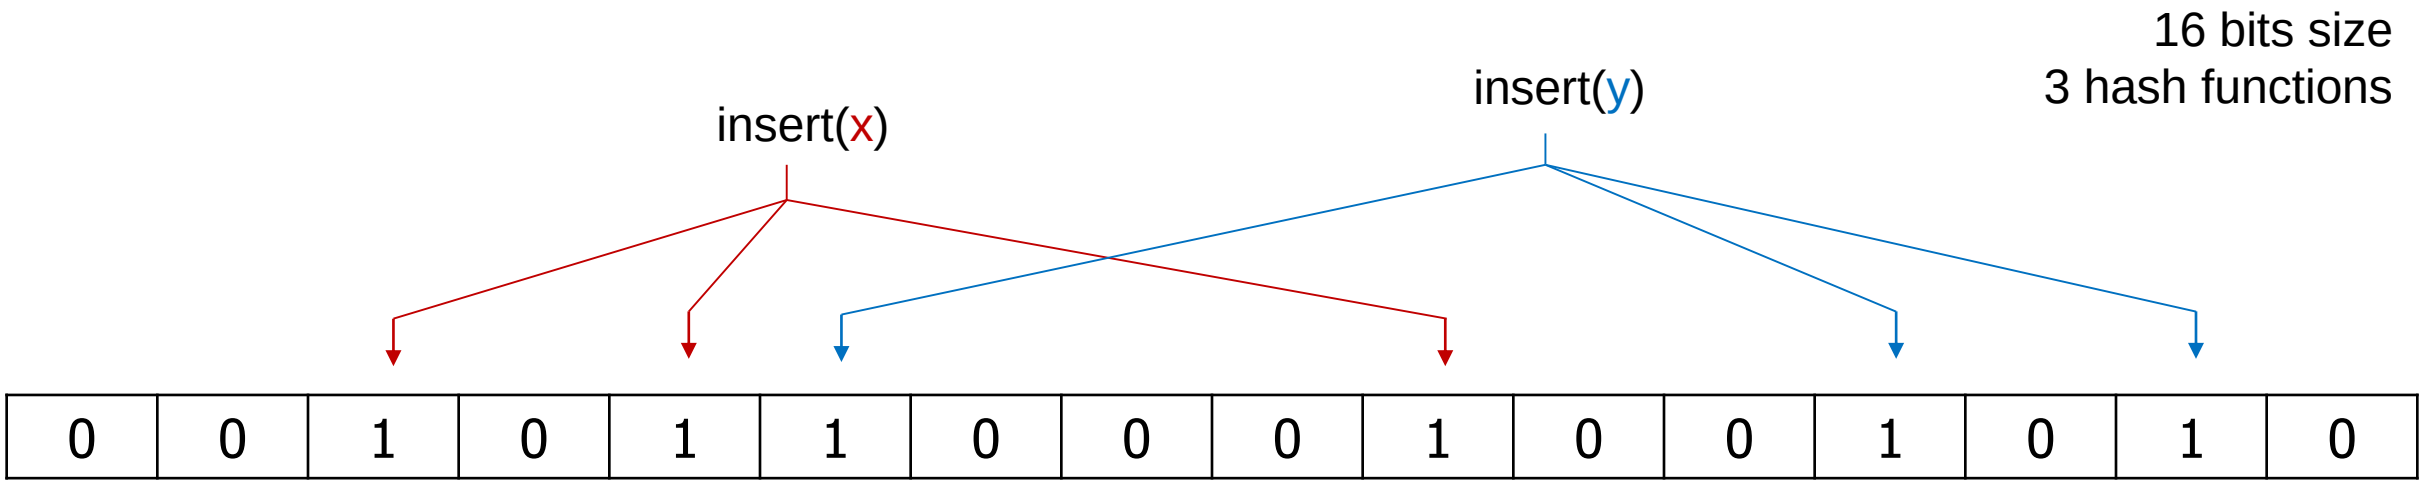

RAIDR Binning

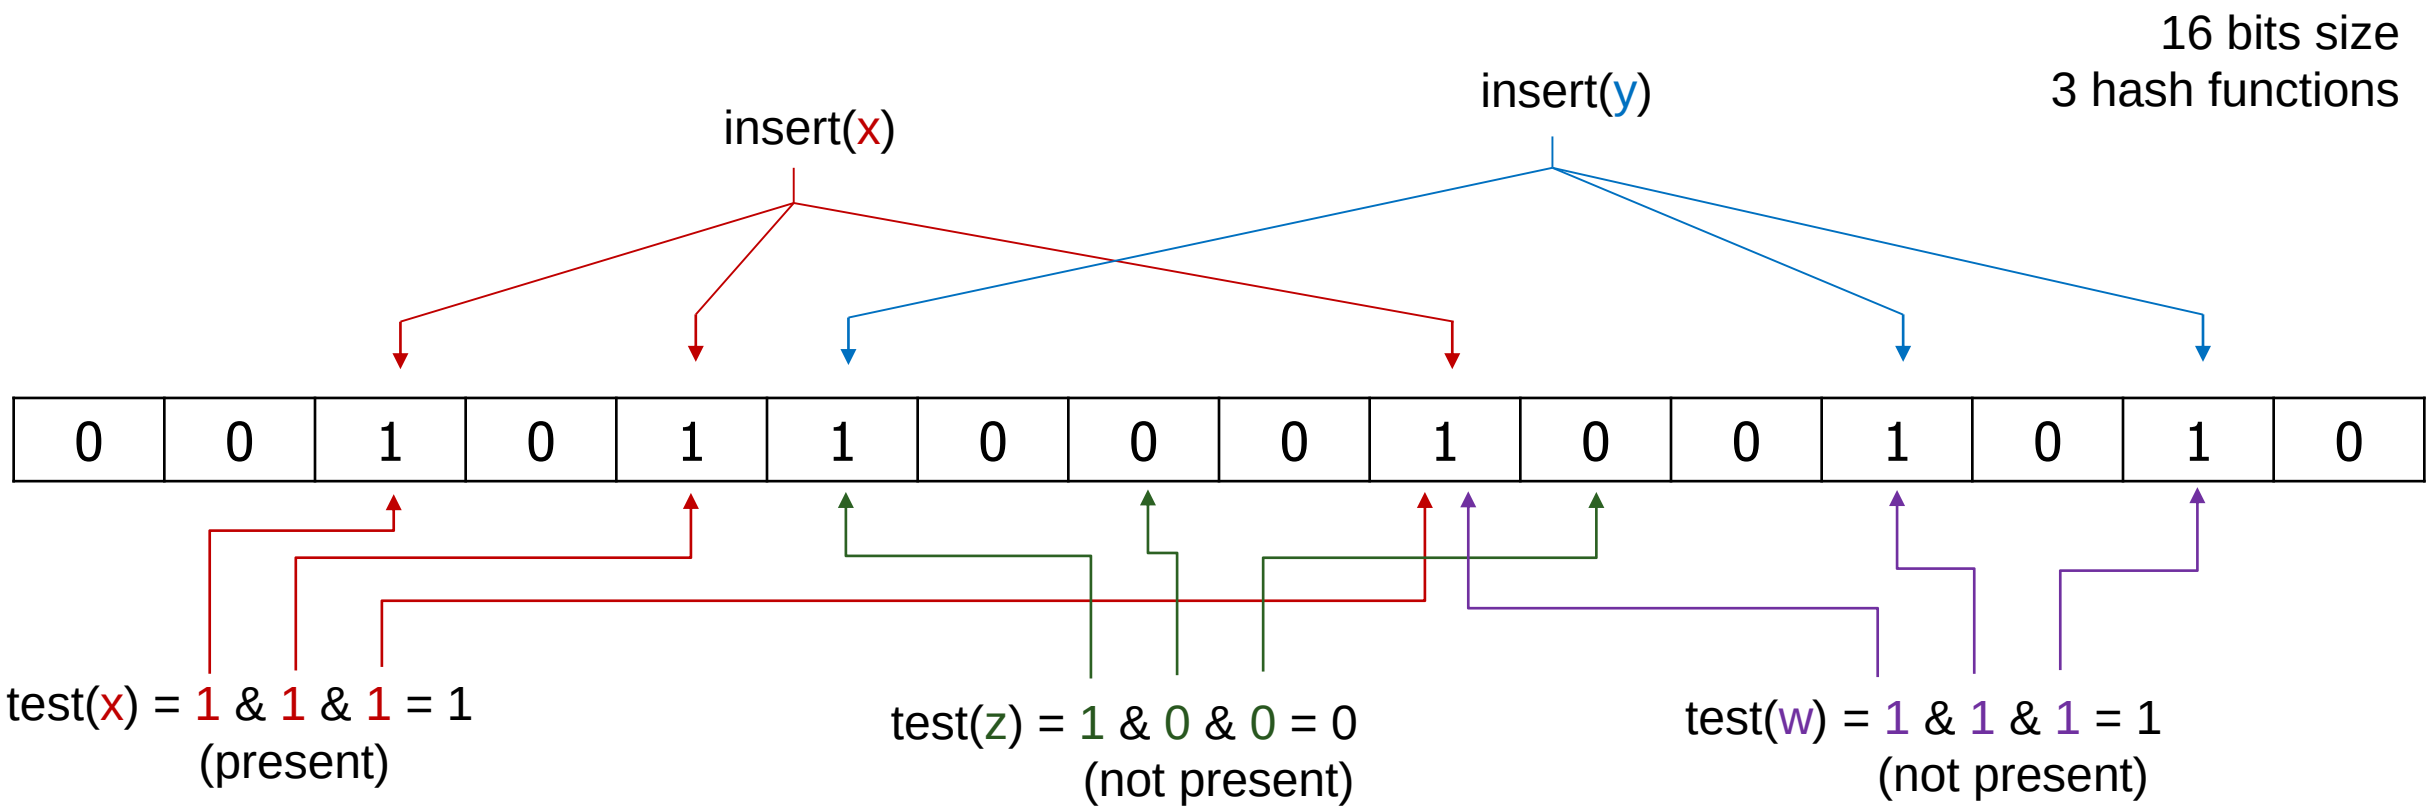

- no overflow property guaranteed
- false positives acceptable
- inability to remove elements acceptable

RAIDR Refreshing

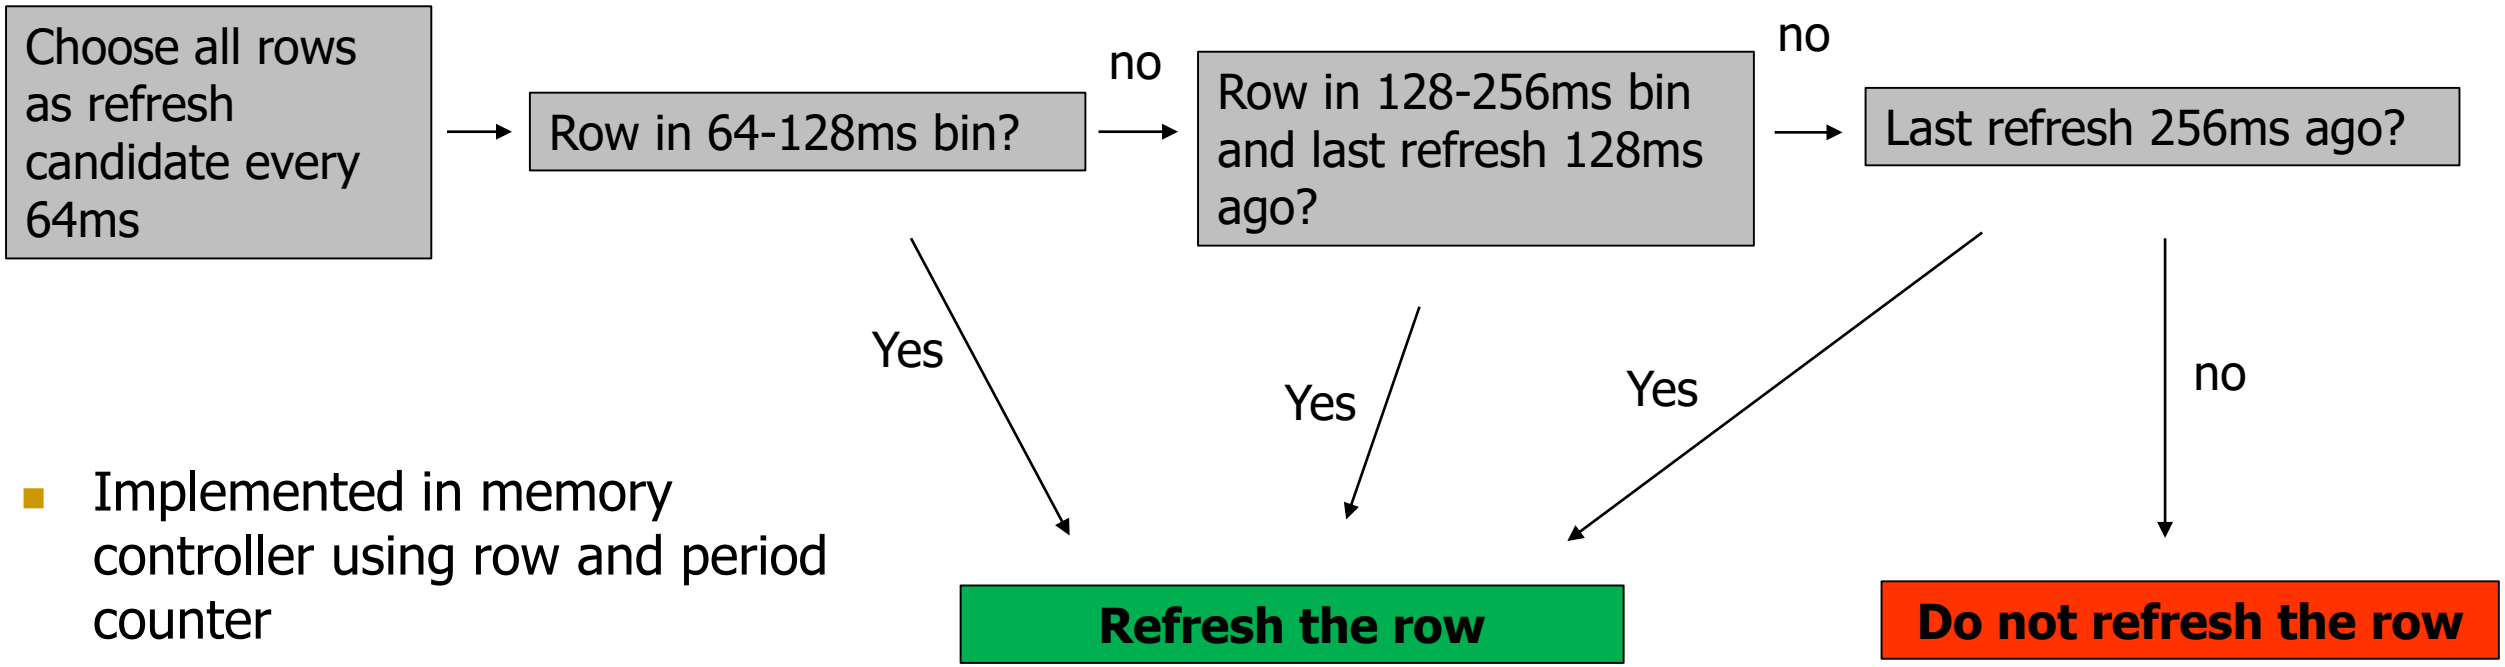

- Implemented in memory controller using row and period counter
- Rate Scaler increases refresh rate at extended temperature mode (>85° C)

Evaluation & Results

Evaluation System

Table 1: Evaluated system configuration

Component	Specifications
Processor	8-core, 4 GHz, 3-wide issue, 128-entry instruction window, 16 MSHRs per core
Per-core cache	512 KB, 16-way, 64 B cache line size
Memory controller	FR-FCFS scheduling [41, 54], line-interleaved mapping, open-page policy
DRAM organization	32 GB, 2 channels, 4 ranks/channel, 8 banks/rank, 64K rows/bank, 8 KB rows
DRAM device	64x Micron MT41J512M8RA-15E (DDR3-1333) [33]

Table 2: Bloom filter properties

Retention range	Bloom filter size m	Number of hash functions k	Rows in bin	False positive probability
64 ms – 128 ms	256 B	10	28	$1.16 \cdot 10^{-9}$
128 ms – 256 ms	1 KB	6	978	0.0179

1.25 KB in Memory Controller

Compared other Mechanisms

- **Auto Refresh**
  - Memory controller sends auto-refresh commands causing DRAM to refresh several rows in one rank automatically
  - Common mechanism in existing systems
- **Distributed Refresh**
  - Memory Controller sends address of row to be refreshed and can make use of bank-level parallelism
  - Same number of refreshes as Auto-Refresh
- **Smart Refresh**
  - Timeout counter for each row that is reset on access or refresh
  - Rows only get refreshed when the counter has expired
- **No Refresh**
  - Not used in practice

Results: Number of Refresh Operations

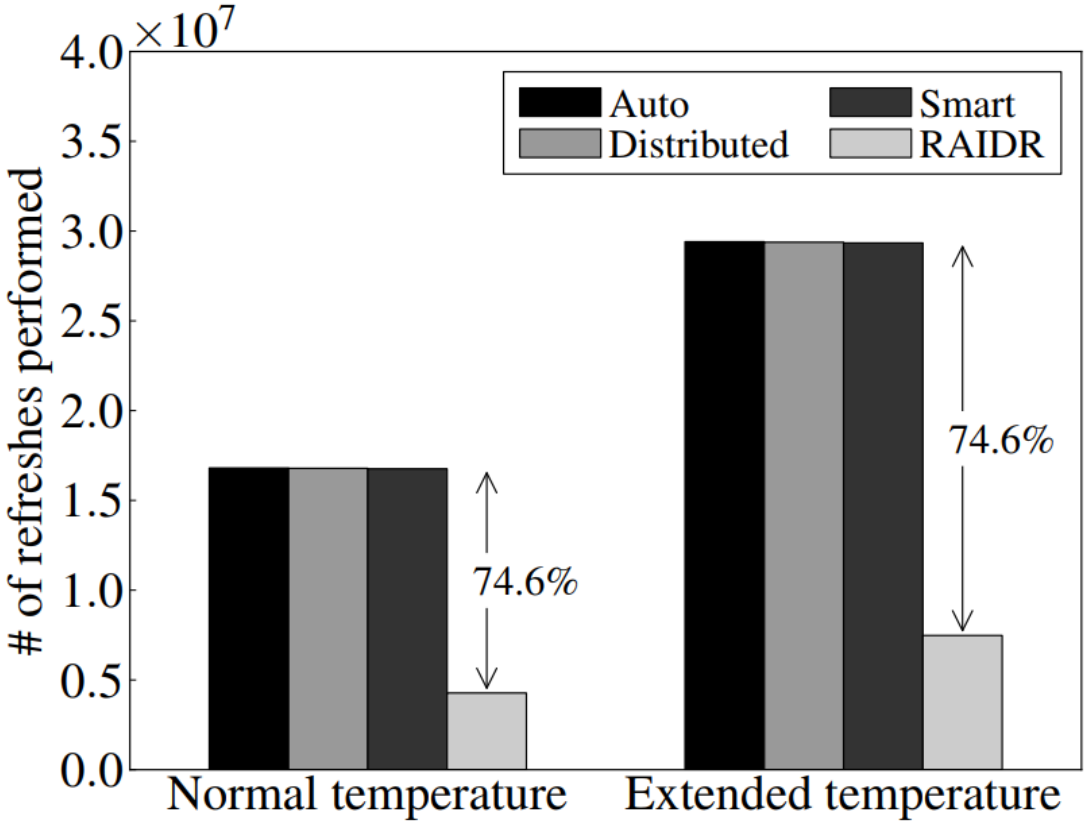

Results: System Performance

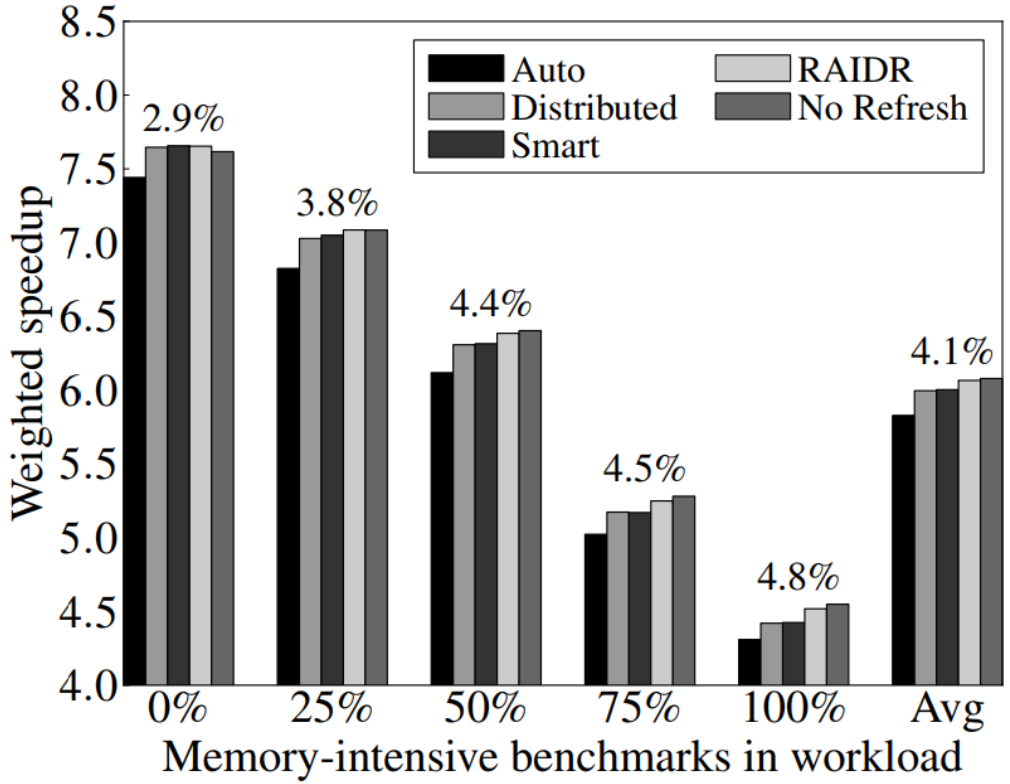

(a) Normal temperature range

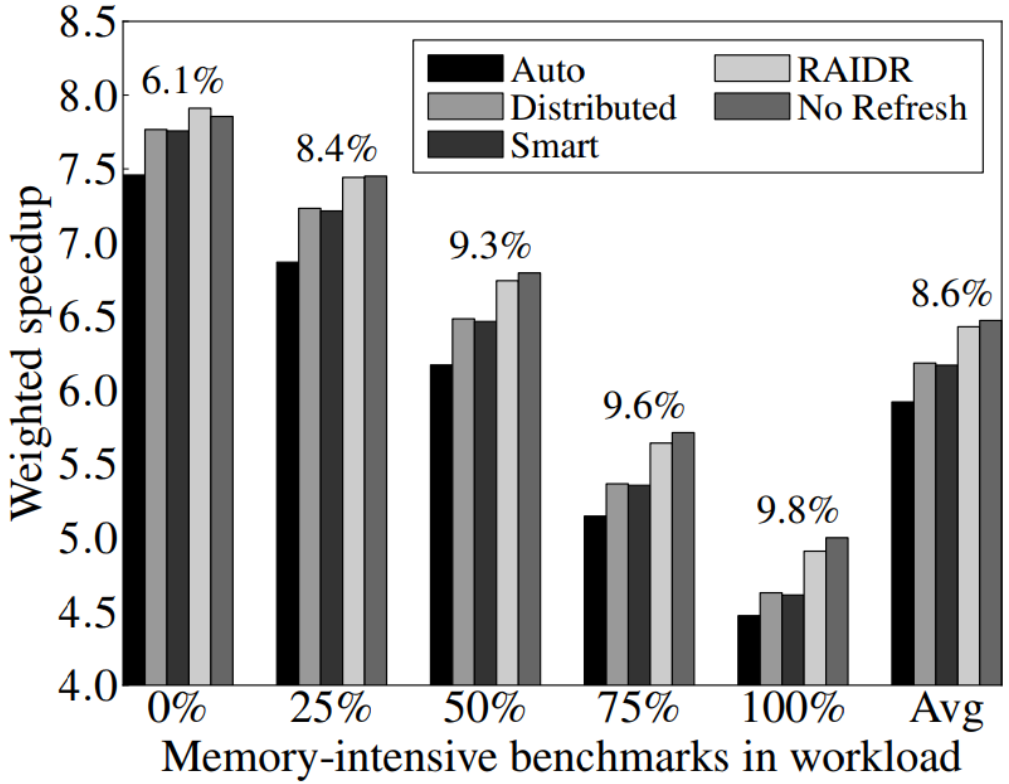

(b) Extended temperature range

Results: Energy Consumption

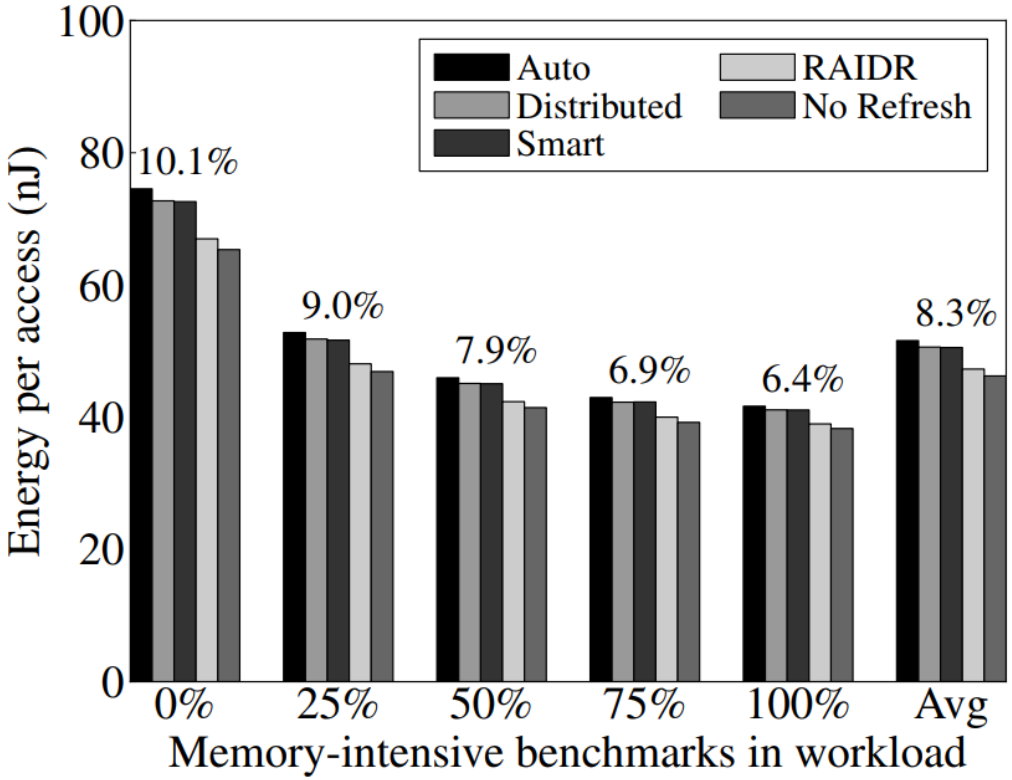

(a) Normal temperature range

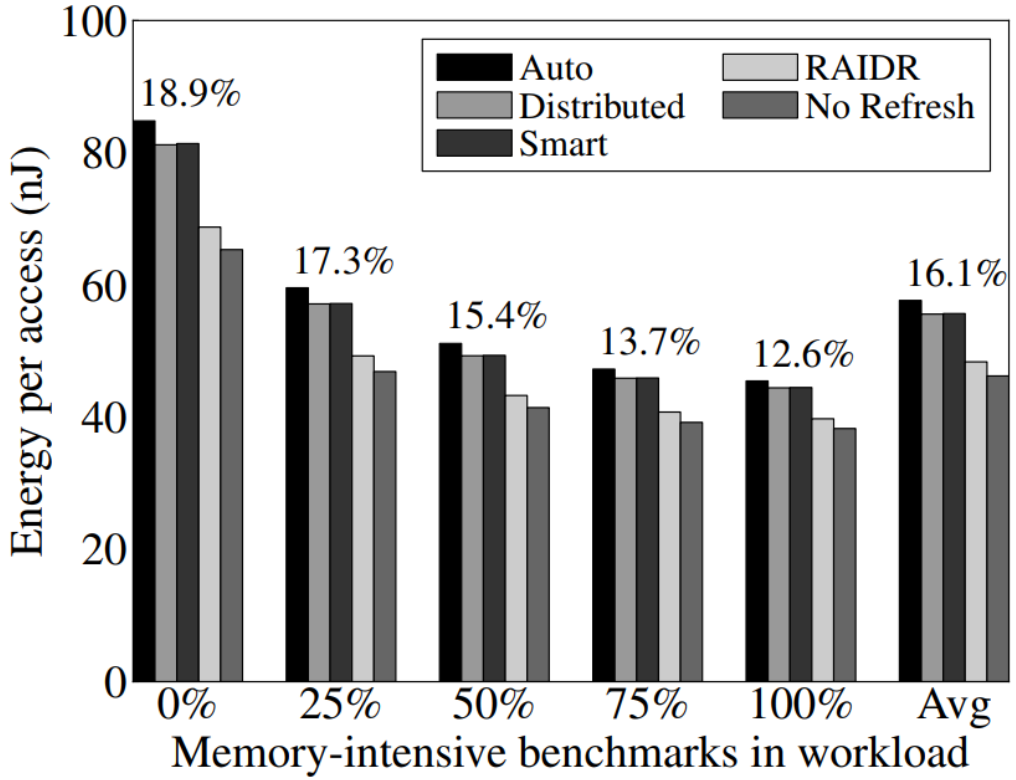

(b) Extended temperature range

Strengths & Weaknesses

Strengths

- Significant reduction in number of refreshes is achieved by very simple and clever idea
- Benefits of RAIDR increase with higher DRAM capacity
- Modifications to Memory Controller are small and not too complex
- No modifications to DRAM or software needed
- Applicable to many different types of DRAM (e.g. 3D-stacked DRAMs, eDRAMs)

Weaknesses

- RAIDR's profiling method might not identify all weak rows
  - Retention times of cells are dependent on values stored in nearby cells (Data Pattern Dependencies)
  - All 1 pattern used by RAIDR can cause cells to have higher retention times and be profiled wrong
- RAIDR is unable to adapt refresh rate for rows at runtime
  - Cells might shift randomly between different retention time states (Variable Retention Times)
  - Later change in retention time can cause a cell to be refreshed at too low rate
- DPD/VRT can cause retention errors which makes the system less reliable
- One weak cell still causes a whole row (64K cells) to be refreshed at high rate
- Setting default interval much higher than 256ms becomes inefficient again

Thoughts and Ideas

- Deactivating rows with lower retention time than 256ms ?
- Reduce refresh rate for everything and increase fault tolerance ?
- Skip refresh of rows that contain only ‘0’s ?
- Handle VRT failures at runtime and adapt refresh rate ?
- Replacing DRAM by devices that don’t need to be refreshed ?

Questions & Discussion

Backup Slides

DRAM System Organization

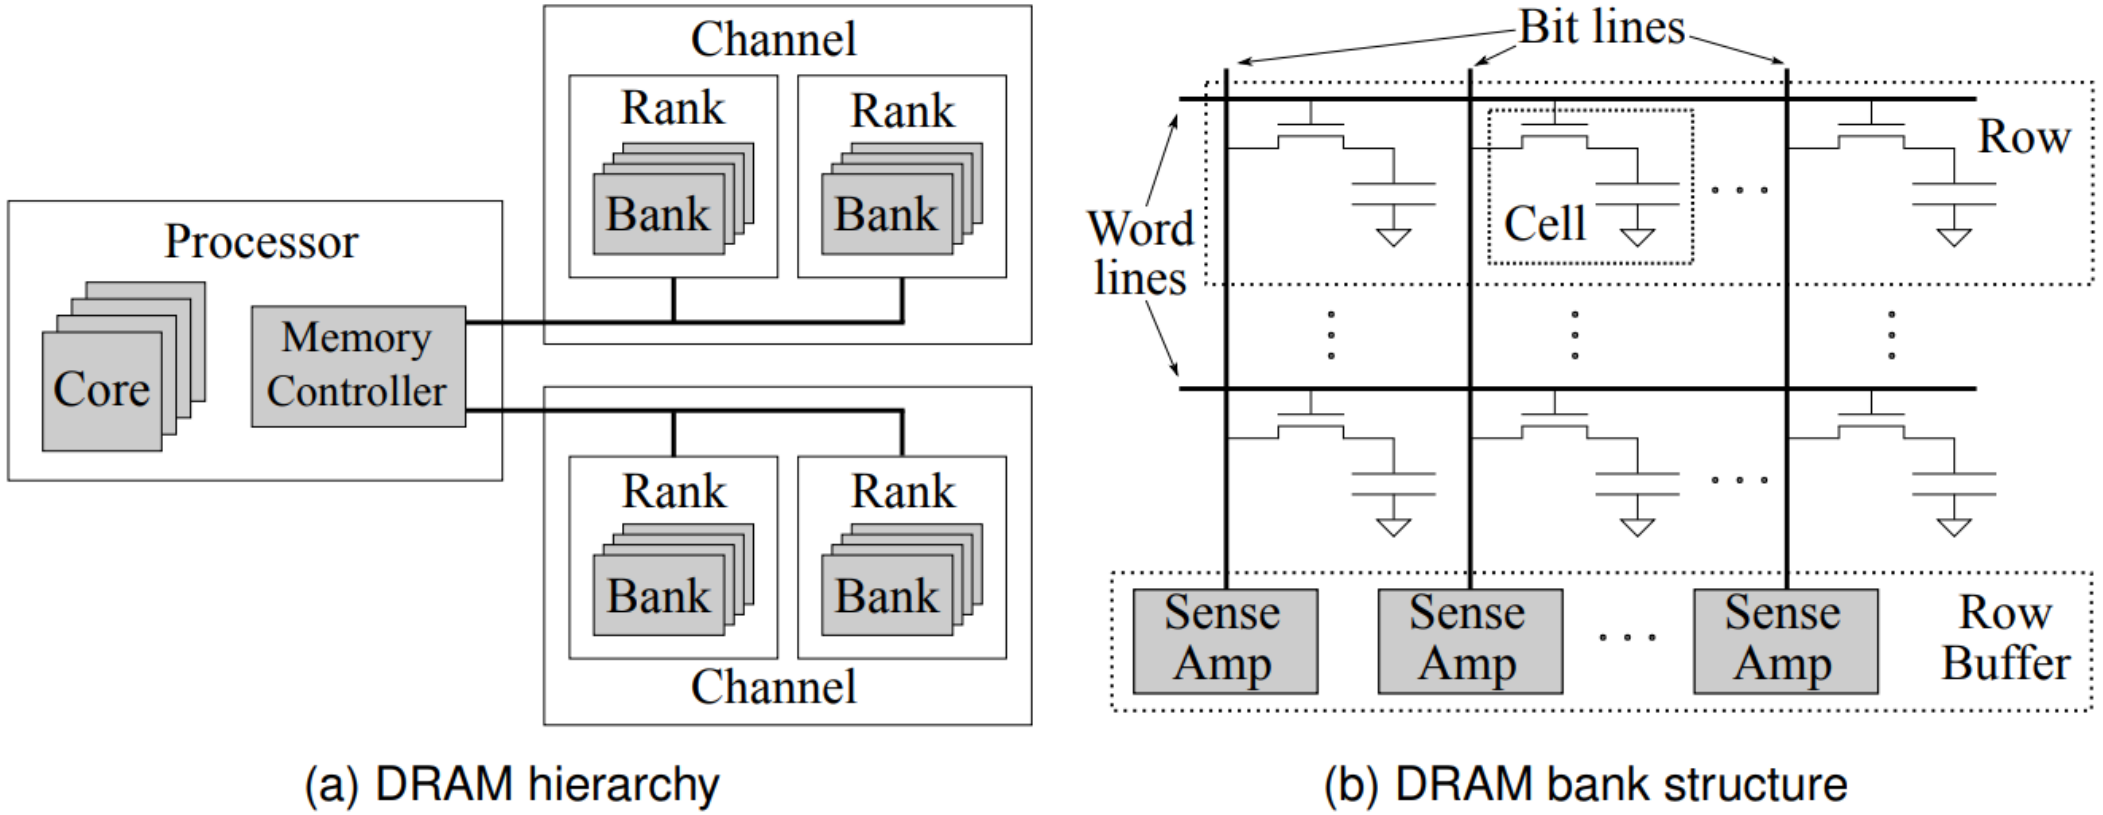

Figure 2: DRAM system organization

Idle Power Consumption

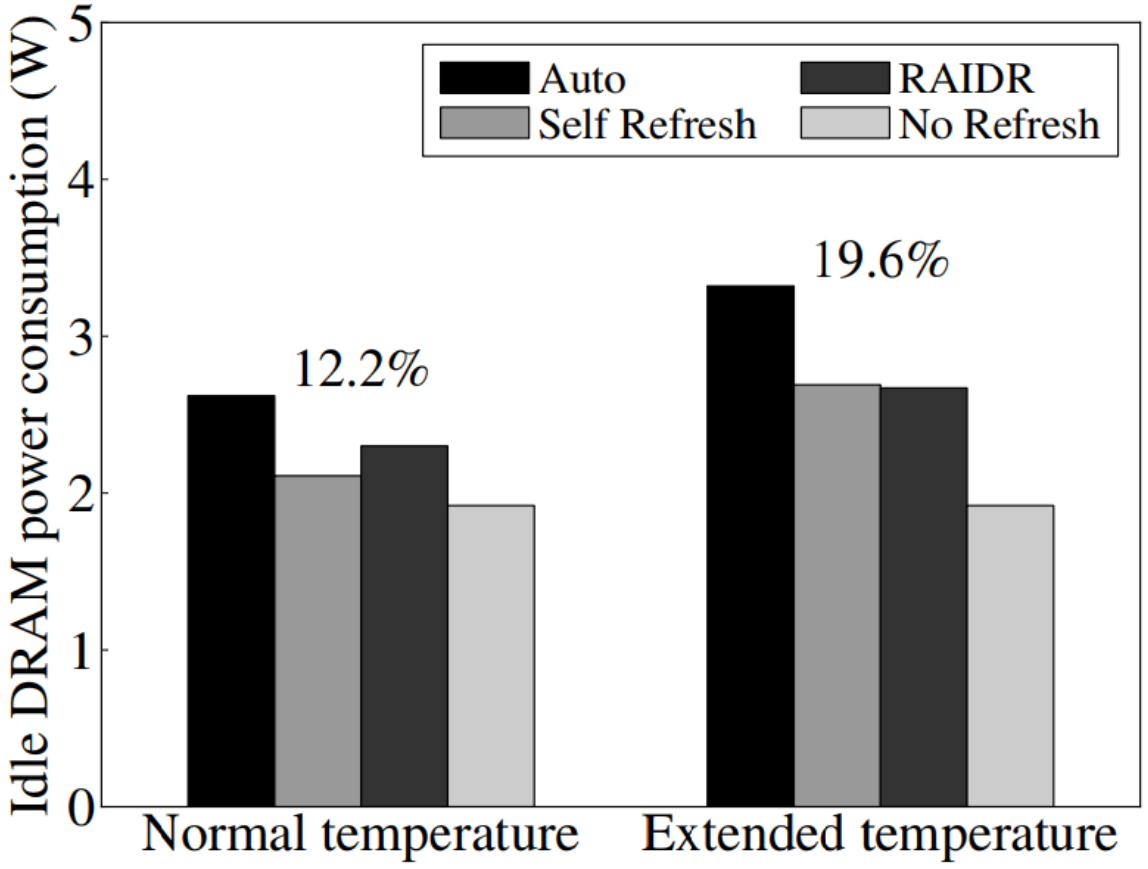

(c) Idle power consumption

RAIDR Configurations

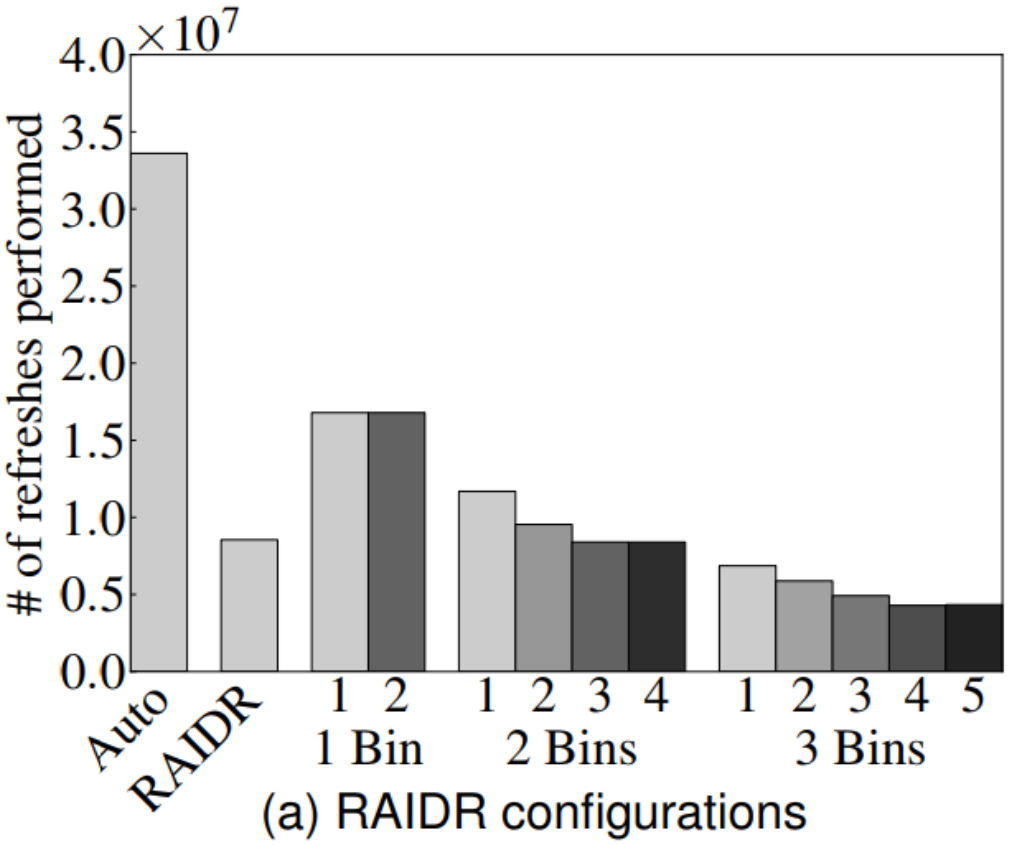

Table 3: Tested RAIDR configurations

Key	Description	Storage Overhead
Auto	Auto-refresh	N/A
RAIDR	Default RAIDR: 2 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 8192$)	1.25 KB
1 bin (1)	1 bin (64–128 ms, $m = 512$)	64 B
1 bin (2)	1 bin (64–128 ms, $m = 1024$)	128 B
2 bins (1)	2 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 2048$)	512 B
2 bins (2)	2 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 4096$)	768 B
2 bins (3)	2 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 16384$)	2.25 KB
2 bins (4)	2 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 32768$)	4.25 KB
3 bins (1)	3 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 8192$; 256–512 ms, $m = 32768$)	5.25 KB
3 bins (2)	3 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 8192$; 256–512 ms, $m = 65536$)	9.25 KB
3 bins (3)	3 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 8192$; 256–512 ms, $m = 131072$)	17.25 KB
3 bins (4)	3 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 8192$; 256–512 ms, $m = 262144$)	33.25 KB
3 bins (5)	3 bins (64–128 ms, $m = 2048$; 128–256 ms, $m = 8192$; 256–512 ms, $m = 524288$)	65.25 KB

Scalability

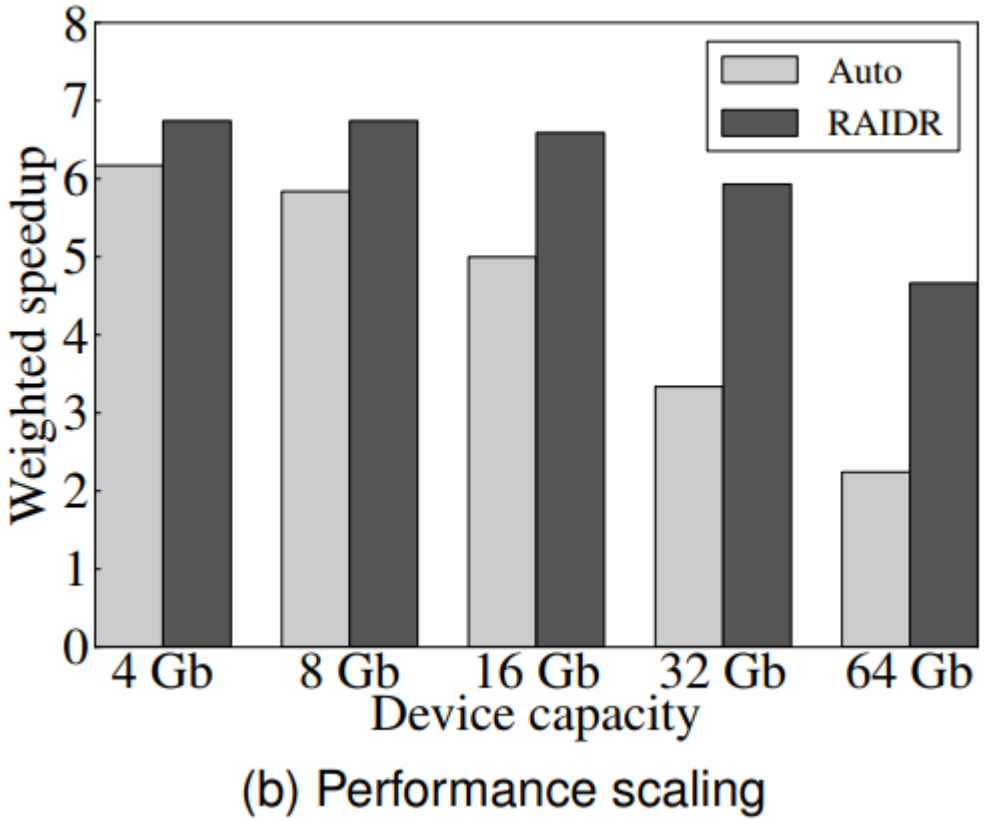

(b) Performance scaling

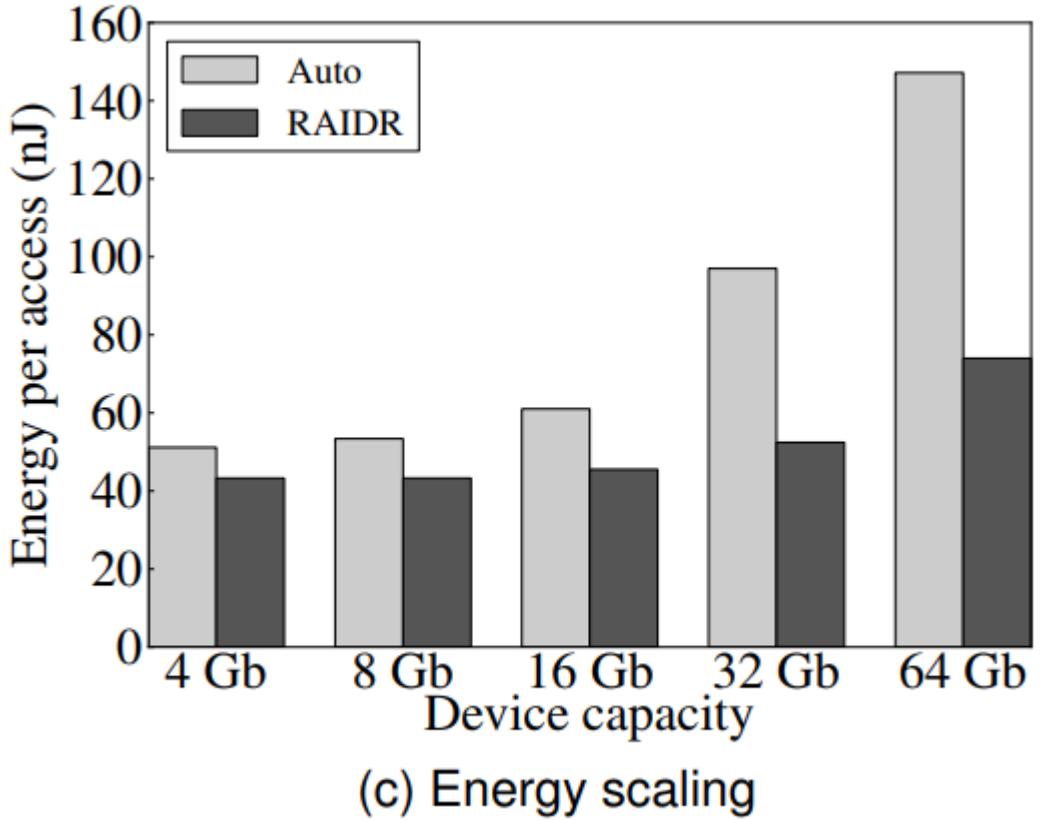

(c) Energy scaling

Period Counter

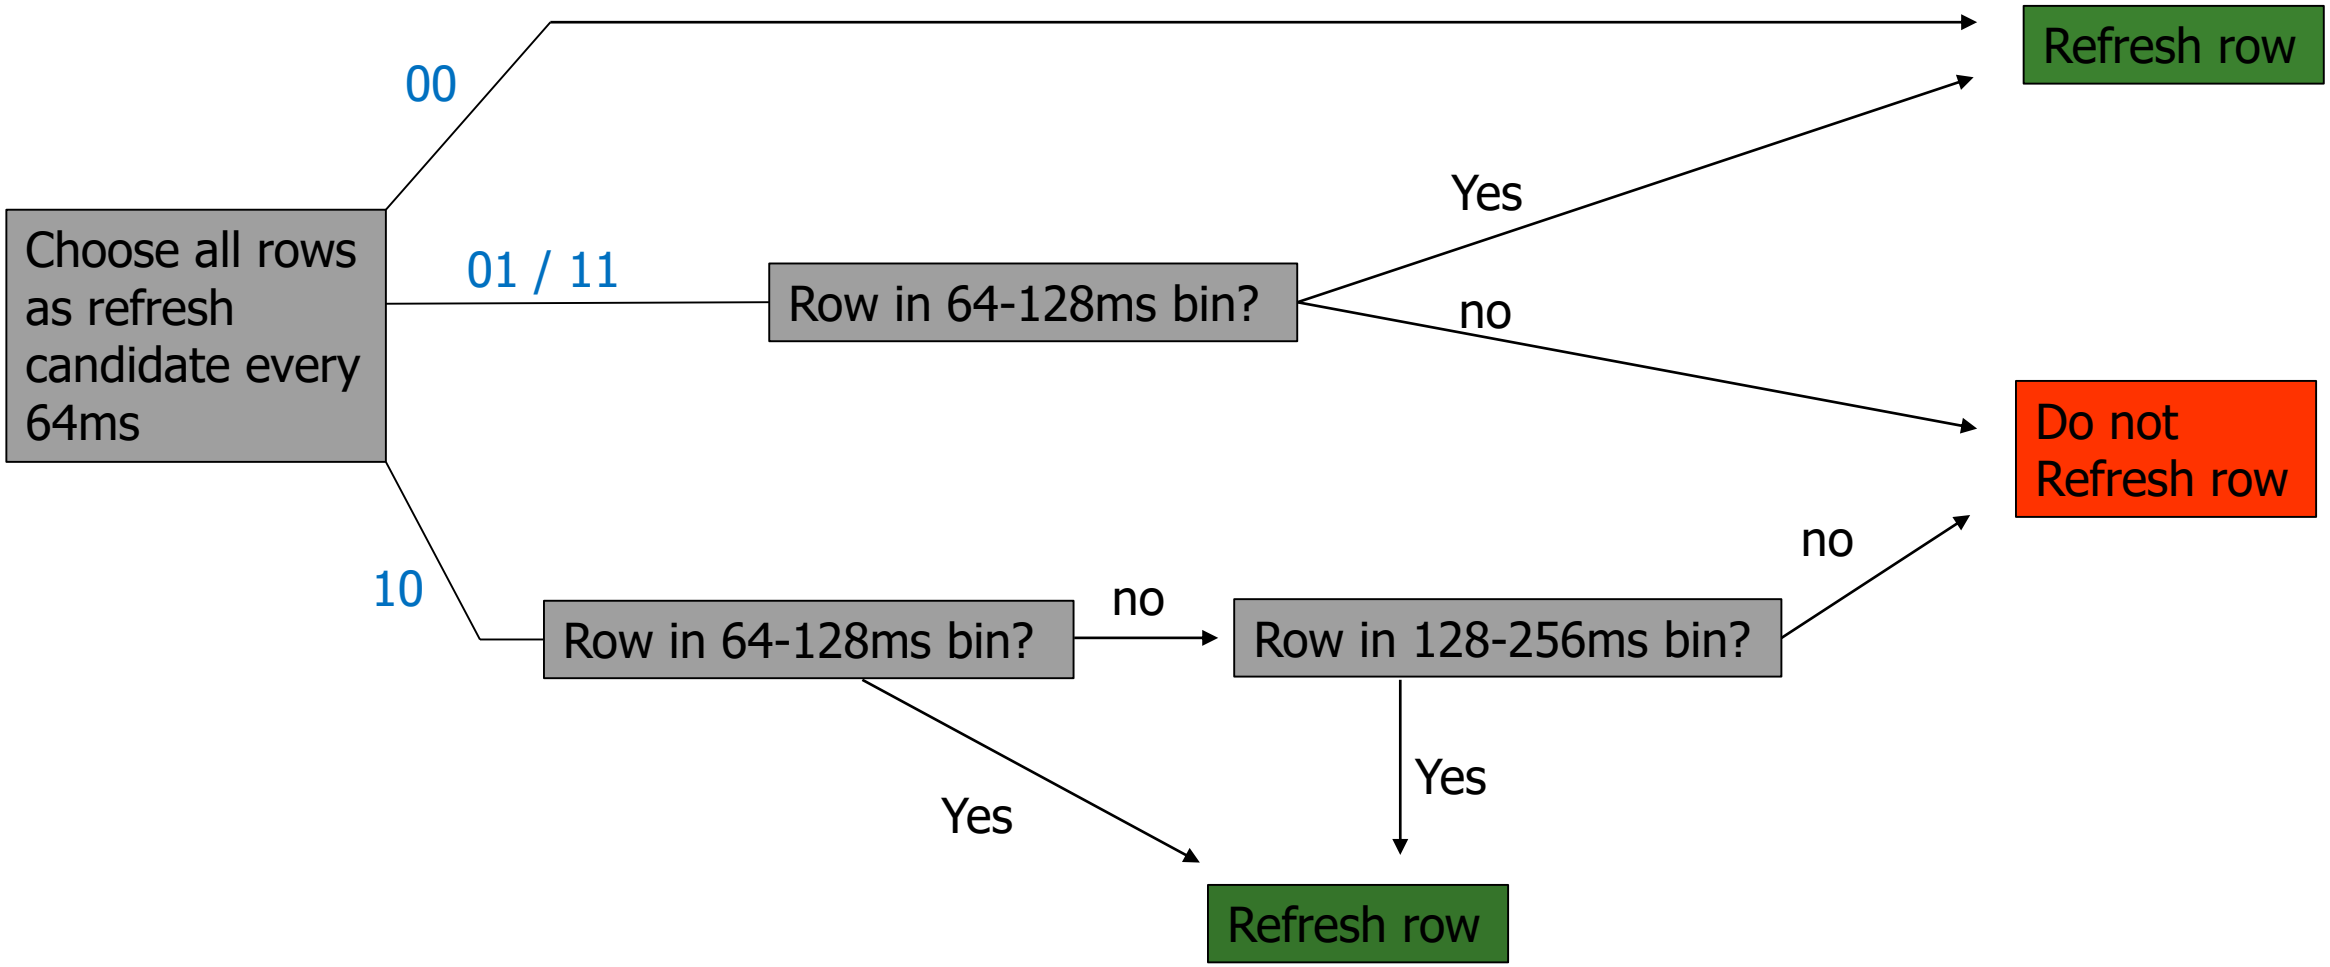

Papers

- https://users.ece.cmu.edu/~omutlu/pub/raidr-dram-refresh_isca12.pdf
- http://www.pdl.cmu.edu/PDL-FTP/NVM/dram-retention_isca13.pdf
- http://www.istc-cc.cmu.edu/publications/papers/2015/avatar-dram-refresh_dsn15.pdf
- http://www.xcg.cs.pitt.edu/papers/baek-tc13.pdf
- https://prashantnair.bitbucket.io/isca40/ArchShield.pdf