## BlockHammer: Preventing RowHammer at Low Cost by Blacklisting RapidlyAccessed DRAM Rows

First presented at: $27^{\text {th }}$ IEEE International Symposium on High-Performance Computer Architecture, 2021
Authors: A. Giray Yağlıkçı¹, Minesh Patelı ${ }^{1}$, Jeremie S. Kim ${ }^{1}$, Roknoddin Azizi¹, Ataberk Olgun ${ }^{1}$, Lois Oros ${ }^{1}$, Hasan Hassan ${ }^{1}$, Jisung Park ${ }^{1}$, Konstantinos Kanellopoulos ${ }^{1}$, Taha Shahroodi ${ }^{1}$, Saugata Ghose ${ }^{2}$, Onur Mutlu ${ }^{1}$
${ }^{1}$ ETH Zürich $\quad{ }^{2}$ University of Illinois at Urbana-Champaign

## Executive summary

## Problem:

- Memory density scaling of DRAM chips causes increasing vulnerability to RowHammer, but most solutions can't scale accordingly
- Current solutions often require knowledge of or modification to DRAM internals


## Goal:

- Find scalable and efficient way to prevent RowHammer without modifying DRAM chip


## Key idea:

- Selectively throttle memory accesses that can cause bit-flips


## Mechanism:

- Tracking all row activations and throttling RowHammer unsafe row accesses
- Identifying and throttling potential attacker threads


## Results:

- Hardware complexity: scalable
- Performance \& energy consumption: efficient \& scalable


## Overview



## Background,

 Problem \& Goal

## DRAM \& RowHammer



Cause: memory density scaling

$\downarrow$ DRAM cell size<br>$\downarrow$ cell-to-cell spacing

## DRAM \& RowHammer



以
Cause: memory density scaling

$\downarrow$ DRAM cell size<br>$\downarrow$ cell-to-cell spacing

RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer



Cause: memory density scaling $\square$ $\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows


## Victim rows

Aggressor row

## Victim rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## DRAM \& RowHammer

以
Cause: memory density scaling
$\downarrow$ DRAM cell size $\downarrow$ cell-to-cell spacing


RowHammer: rapidly activating (opening) and precharging (closing) DRAM row can cause bit-flips in nearby rows

## Current solutions to RowHammer



## Current solutions to RowHammer



## Increased refresh rate

What: refresh (all!) DRAM rows more often to reduce probability of successful bitflip

RowHammer ( RH ) is getting worse: cannot prevent RH without unacceptable performance loss and power consumption increase

## REFRESH



## REFRESH



## Increased refresh rate



## Current solutions to RowHammer



## Reactive refresh

What: observes activations and reacts by refreshing potential victim rows e.g., TWiCe, PARA, ProHIT, MRLoc, CAT, CBT, ...

Requires proprietary knowledge on DRAM internals: need to know which rows are adjacent to aggressor rows

- Faulty rows/cells/columns
- Differences in access latency of fastest \& slowest cell

Wang, Minghua, et al. "DRAMDig: a knowledge-assisted tool to uncover DRAM address mapping." 2020 57th ACM/IEEE Design Automation Conference (DAC). IEEE, 2020.

Some are probabilistic methods: do not prevent RowHammer completely

## REFRESH



## REFRESH

Proactive throttling


## Current solutions to RowHammer



## Physical isolation

## Already defeated!

What: separates physically sensitive da'
e.g., by adding buffer setween (ZebRAM)
e.g., by separating me ry of user and kernel mode (CATT)

RowHammer is getting worse: we need to provide greater isolation

- wastes memory capacity
- reduces fraction of cells we can protect from RH

Requires proprietary knowledge on DRAM internals: need to know which rows are adjacent to aggressor rows

- Faulty rows/cells/columns
- Differences in access latency of fastest \& slowest cell


## Physical isolation

## Buffer row



## Physical isolation

## Buffer rows

Buffer rows


## Current solutions to RowHammer



What: limit repeated access to the same row
e.g., by setting a minimum access delay
e.g., by limiting number of accesses to a row within refresh window

Challenge: performance overhead Will we delay every access?

## Challenge: area overhead

How do you track the number of row activations?

# Proactive throttling 

Countdown to next
row activation

Can I get access to row X ?

Countdown to next
row activation

0:00:00:005

Can I get access to row X ?

## Current solutions to RowHammer



## In search of a better solution

F
Efficient: low performance/area overhead

Scalable: we want things to work in the future

Implemented without knowledge of or modification to DRAM chip

## Key idea: selectively throttle RowHammer-like memory accesses by

## Tracking activation rates of all rows in an area-efficient way

Using tracking data to throttle RowHammer unsafe activations
$r$
Identifying and limiting row activation rates of potential attacker threads (minimizes performance degradation of benign threads)

# Mechanisms \& Implementation 

## BlockHammer =

a

RowBlocker

## AttackThrottler

## RowBlocker



## RowBlocker BL



## RowBlocker BL



Goal 1: Track which rows have been activated and how many times

Goal 2: Blacklist when activation rate exceeds blacklisting threshold

## How can we do this area-efficiently?

## Recap: Bloom filter



Question: does a set contain a certain element?


Main components: hash functions + bit array

Operations: insert, test, clear

## Recap: Bloom filter



Question: does a set contain a certain element?


## Recap: Bloom filter



## Recap: Bloom filter

## ! Insert 5

(i) Set $=\{5\}$


## Recap: Bloom filter

## ! Insert 7

(1) Set $=\{5\}$


## Recap: Bloom filter

## ! Insert 7

(i) Set $=\{5,7\}$


## Recap: Bloom filter

## ! Insert 9

(i) Set $=\{5,7\}$


## Recap: Bloom filter

## ! Insert 9

(i) Set $=\{5,7,9\}$


## Recap: Bloom filter

## ! Test 9

(i) Set $=\{5,7,9\}$


## Recap: Bloom filter

## ! Test 8

$$
\text { (i) Set }=\{5,7,9\}
$$

## False Positive!!



## Counting Bloom filter



Remember: we want to know how many times a row is activated (and blacklist it if activation rate > threshold)

Idea: Counting Bloom filter (CBF)
(tracks number of times an element is inserted into filter)

## Counting Bloom filter



## Counting Bloom filter



## Counting Bloom filter

## Insert 9

(i) Set $=\{5,7,9\}$


## Counting Bloom filter



Test 9
(i) Set $=\{5,7,9\}$

Here threshold $=0$


## Counting Bloom filter



Idea: Counting Bloom filter (CBF)
(tracks number of times an element is inserted into filter)

But Bloom filter is getting saturated

## Counting Bloom filter

## ! Delete 8

(i) Set $=\{5,7,9\}$


## Counting Bloom filter



## Counting Bloom filter


! Test 5
(i) Set $=\{5,7,9\}$

Here threshold $=0$


## Unified Bloom filter



Remember: we want to know how many times a row is activated (and blacklist it if activation rate > threshold)

But we can't prevent false negatives
(without compensating for it in terms of space)

Idea: Unified Bloom filter (UBF)
(tracks all elements inserted into filter during specific time window)

## Unified Bloom filter



Unified Bloom filter: active + passive Bloom filter

- Both insert all elements into filter
- Only active filter responds to test queries
- Active filter clears array at end of specified time interval (= epoch)
- Switch roles every epoch


## Guarantees no false negatives

when tested for elements inserted in the last two epochs

## Epoch 1

Epoch 2
Epoch 3

## Unified Bloom filter

| Filter $\mathbf{A}$ |  |  | $\ldots$ |
| :--- | :--- | :--- | :--- |
| Filter B |  |  | $\ldots$ |

Filter A


## Unified Bloom filter


! Insert 5


Set $=\{5\}$
Set $_{A}=\{5\}=$ Set $_{B}$
Filter A: active


## Unified Bloom filter


! Insert 7
(i)
Set $=\{5,7\} \quad \operatorname{Set}_{A}=\{5,7\}=$ Set $_{B}$

Filter A: active


## Unified Bloom filter



Clear A
(i)

Set $=\{5,7\} \quad \operatorname{Set}_{A}=\{ \}$, Set $_{B}=\{5,7\}$
Filter A: active
-


## Unified Bloom filter


! Insert 9
(i) Set $=\{5,7,9\} \operatorname{Set}_{A}=\{9\}, \operatorname{Set}_{B}=\{5,7,9\}$

Filter A: passive


Epoch 1
Epoch 2
Epoch 3

## Unified Bloom filter

Epoch 1

# (i) Set $=\{5,7,9\} \operatorname{Set}_{A}=\{9\}$, Set $_{B}=\{5,7,9\}$ 

Filter A: passive


## Dual counting Bloom filter



Remember: we want to know how many times a row is activated (and blacklist it if activation rate > threshold))

Idea: dual counting Bloom filter (D-CBF)
= unified Bloom filter + counting Bloom filter

- both filters use different hash functions
- hash functions of active filter are altered at end of epoch


## Dual counting Bloom filter

RowBlocker BL


Filter A: passive


## RowBlocker



## RowBlocker History Buffer (HB)

## RowBlocker HB

(per DRAM rank)

## RowBlocker HB



Goal 1: Track which rows were activated recentlyGoal 2: Test if current row is one of them

## RowBlocker HB



What: circular first-in-first-out (FIFO) queue (stores record of rows activated in last $t_{\text {delay }}$ time window)
$\stackrel{+7}{\times 7}$
Operations: insert, test, (update)

## RowBlocker HB

Row ID: rank-unique ID for all rows
© Timestamp: current time

Valid bit


Tail pointer (youngest entry)

## Update

Row ID: rank-unique ID for all rows

© Now - Timestamp >= telay

Valid bit: set to 0


Tail pointer (youngest entry)

## RowBlocker: is this row activation RH-safe?

## RowBlocker BL

(per DRAM bank)

Memory
Request
RowBlocker HB
(per DRAM rank)

## RowBlocker: is this row activation RH-safe?

## RowBlocker BL

(per DRAM bank)

Memory
Request
RowBlocker HB
(per DRAM rank)

## RowBlocker: is this row activation RH-safe?



## RowBlocker: is this row activation RH-safe?



## RowBlocker: is this row activation RH-safe?



## BlockHammer =

## 8

$+$

## RowBlocker

## AttackThrottler



Goal 1: identify potential attacker threads

Goal 2: limit their memory bandwidth usage

## AttackThrottler



Goal 1: identify potential attacker threads
Goal 2: limit their memory bandwidth usage

## 1. Identifying (potential) attacker threads

## 關 4

How: RowHammer Likelihood Index (RHLI)

$$
R H L I=\frac{\# \text { blacklisted row activations thread performs to DRAM bank }}{\text { max \# times blacklisted row can be activated in protected system }}
$$

RHLI $=\mathbf{0}$
(benign
threads)

More and more
likely to induce bit-flip

Quantifies similarity between a given thread's memory access pattern and a real RowHammer attack

## 1. Identifying (potential) attacker threads

Idea: 2 counters per <thread, bank> pair, used same time-interleaving mechanism of D-CBF

2 counters: active + passive counter

- Thread activates blacklisted row in bank $\rightarrow$ increment both counters
- Only active counter is used to calculate RHLI
- RowBlocker clears active filter in bank $\rightarrow$ AttackThrottler clears all active counters in bank and switches roles

Calculates RHLI from rows blacklisted in last two epochs

## AttackThrottler



Goal 1: identify potential attacker threads

## (C)

Goal 2: limit their memory bandwidth usage

## 2. Limiting memory bandwidth usage

How: by applying quota to thread's total in-flight memory requests

$$
\text { Quota } \sim \frac{1}{R H L I}
$$

Thread keeps activating blacklisted row: RHLI increases $\rightarrow$ quota decreases

Thread reaches quota: can't make new memory request (until ongoing request is completed)

Lessens memory bandwidth usage of attacker threads $\rightarrow$ frees up memory bandwidth for benign threads

## AttackThrottler: $3^{\text {rd }}$ goal?

Goal 1: identify potential attacker threads $\rightarrow$ RHLI

Goal 2: limit their memory bandwidth usage $\rightarrow$ quota

## 3. Share info with the Operating System

What: Share <thread, DRAM bank> RHLI values with OS

Goal: mitigate RH attack at software level
e.g., by killing or descheduling attacker thread

Results 사

## We compare BlockHammer with:



Baseline system: no RH mitigation

。
Three probabilistic mitigation mechanisms: PARA, ProHIT, MRLoc

囲
Three deterministic mitigation mechanisms: CBT, TWiCe, Graphene

## Results



Hardware complexity analysis
$\rightarrow$ scalable \& low cost


Performance \& energy consumption
$\rightarrow$ scalable \& efficient

## Results



Hardware complexity analysis
$\rightarrow$ scalable \& low cost
$\rightarrow$ scalable \& efficient

## 1. Hardware complexity analysis



## 1. Hardware complexity analysis

## RowHammer threshold 32K

- PARA, PRoHIT, MRLoc $\rightarrow$ extremely area-efficient (because probabilistic)
- Graphene << TWiCe, BlockHammer < CBT


## 1. Hardware complexity analysis

- PARA, PRoHIT, MRLoc $\rightarrow$ extremely area-efficient (because probabilistic)
- Graphene << TWiCe, BlockHammer < CBT


## RowHammer threshold 1K

- Graphene x28.5, TWiCE x34.5, CBT x19.7 $\leftrightarrow$ BlockHammer x11.2
- New order: Graphene < BlockHammer << TWiCE << CBT
- BlockHammer is catching up!


## 1. Hardware complexity analysis



Conclusion 1: BlockHammer is more scalable than other RowHammer mitigation mechanisms


Conclusion 2: Graphene mostly better than BlockHammer... for now at least...

- RowHammer will get worse $\rightarrow$ maybe < 1K? (currently at 9.6K)
- Graphene does not scale as well!


## Results

Hardware complexity analysis
$\rightarrow$ scalable \& low cost


Performance \& energy consumption
$\rightarrow$ scalable \& efficient

## 2. Performance \& energy consumption

- 

Single-core system performance

## 2. Performance \& energy consumption © Citc



Benign Application Groups

## 

Benign Application Groups
BlockHammer has no performance or energy overhead for singlecore benign applications

## 2. Performance \& energy consumption



## 2. Performance \& energy consumption



BlockHammer has competitive performance and energy consumption when no attack is present

## 2. Performance \& energy consumption

## 2. Performance \& energy consumption

## ○ <br> Single-core system performance

413
Scalability

## 2. Performance \& energy consumption



BlockHammer has negligible performance and energy consumption overheads and still does if RH worsens (when no attack is present)

## 2. Performance \& energy consumption



BlockHammer has significantly better performance and lower energy consumption as RH worsens (when attack is present)

## 2. Performance \& energy consumption



Conclusion 1: When the system is not under attack, BlockHammer is competitive with the other state-of-the-art mechanisms, also at the lowest RH thresholds

Conclusion 2: In the presence of a RH attack, BlockHammer has significantly better performance and lower energy consumption than all other state-of-the-art mechanisms, even at lower RH thresholds

## Summary <br> (3)

## Summary \& Conclusion

## Problem:

- Memory density scaling of DRAM chips causes increasing vulnerability to RowHammer, but most solutions can't scale accordingly
- Current solutions often require knowledge of or modification to DRAM internals


## Goal:

- Find a scalable and efficient way to prevent RowHammer, without knowledge of or modification to DRAM internals


## Mechanisms:

- RowBlocker: tracking all row activations efficiently (by using Bloom filters) and throttling RowHammer unsafe row accesses
- AttackThrottler: identifying (RHLI) and throttling (quota) potential attacker threads


## Results:

- Hardware complexity: most scalable solution (Graphene currently more efficient but not as scalable)
- Performance \& energy: 〔No RowHammer attack: competitive, even at lower RH thresholds RowHammer attack: significantly better than all other solutions


# Strengths \& Weaknesses 

## Strengths

- BlockHammer still scales well when DRAM chips are getting more vulnerable to RowHammer
- Implementation requires no knowledge of or modifications to DRAM internals (completely implemented in memory controller)
- Makes distinction between benign applications and potential attacks
- Introduces many new concepts and even more possible improvements
- Innovative idea $\rightarrow$ groundwork for new type of RowHammer mitigation: proactive throttling


## Weaknesses

- Completely implemented in memory controller $\rightarrow$ cannot be implemented in already manufactured processor chips
- Some empirically-determined parameters (e.g., Bloom filter size)
- Partially determines false positive rate $\rightarrow$ room for improvement!
- Evaluation is simulated on DDR4-based memory subsystem $\rightarrow$ what about LPDDR4?
- Results probably similar
- And hardware designers will redo it anyway...

Discussion 0

## Discussion

- Should we always aim for deterministic solutions or are probabilistic methods not that bad?
- Can we lower BlockHammer's hardware complexity by adopting a probabilistic approach? What would you change in BlockHammer to achieve that?
- Remember:
BlockHammer = RowBlocker (D-CBF + HB) + AttackThrottler (RHLI + quota)
- Is it a good idea to modify BlockHammer into a probabilistic mitigation mechanism? Why (not)?
- Are there other ways to reduce BlockHammer's hardware complexity?


## Discussion

- Once we can quickly reverse-engineer DRAM address mappings, will BlockHammer still be the best approach?
- What would be the ideal RowHammer mitigation mechanism and why?


## Discussion

- Do you think we can combine (parts of) BlockHammer with other mitigation mechanisms? What would be the (dis)advantages?
- Remember:

BlockHammer $=$ RowBlocker (D-CBF + HB) + AttackThrottler (RHLI + quota)

- Do you have any other ideas to improve BlockHammer?

| refreshing all DRAM rows = <br> - high performance loss \& energy consumption | Increased refresh rate | Reactive refresh | = victim row refresh <br> - challenge: finding victim rows <br> - some probabilistic methods |
| :---: | :---: | :---: | :---: |
| using buffer/isolation rows = <br> - challenge: finding victim rows <br> - RH gets worse $\rightarrow$ need more isolation | Physical isolation | Proactive throttling | $\leftarrow$ BlockHammer |

## Discussion

- What can we do with the RHLI at the software level?
- E.g. killing or descheduling a thread
- What problems would you encounter?

Backup Slides

## Insert

Row ID: rank-unique ID for all rows
© Timestamp: current time

Valid bit: set to 1


Tail pointer
(youngest entry
Timestamp
Valid bit

## Test: row recently activated?

Ríd Row ID $==$ to be accessed row
© Timestamp

## We want low latency!

Valid bit == 1
Head pointer
(oldest entry)

Tail pointer
(youngest entry)

## Test: row recently activated?

R Row ID == to be accessed row
T. Timestamp

## Store row addresses in CAM

Valid bit $==\mathbf{1}$
Head pointer
(oldest entry)

Tail pointer
(youngest entry)

## Comparison

- Compare BlockHammer with
- (Baseline system: no RH mitigation)
- 3 probabilistic mitigation mechanisms (errors still possible)
- PARA
- ProHIT
- MRLoc
- 3 deterministic mitigation mechanisms (usually area overhead)
- CBT
- TWiCe
- Graphene


## PARA: definition

- = Probabilistic Adjacent Row Activation
- Row gets activated $\rightarrow$ adjacent rows get activated (= refreshed) with probability p


## PARA: mechanism

- Remember: Reactive refresh



## PARA: mechanism

- Remember: Reactive refresh



## PARA: mechanism

- Remember: Reactive refresh



## PARA: mechanism

- Remember: Reactive refresh


## REFRESH

with probability $p$

## REFRESH

with probability $p$


## PARA: mechanism

- Remember: Reactive refresh


## REFRESH

with probability $p$

## REFRESH

with probability $p$


## PARA: mechanism

- Remember: Reactive refresh


## REFRESH

with probability $p$

## REFRESH

with probability $p$


## PARA: weaknesses

- Cannot prevent bit-flips with $100 \%$ certainty (probabilistic!)
- Performance $\rightarrow$ vulnerable to applications with mix of few frequently activated rows and many randomly activated ones (often the case in memory-intensive programs) $\rightarrow$ solution: ProHIT
- Knowledge on in-DRAM mapping needed


## ProHIT: definition

- Based on PARA
- Selects victim rows by considering the access patterns of applications (on top of probabilistic selection) $\rightarrow$ done by Probabilistic History Table
- Key operations: row activation $\rightarrow$
- Probabilistic table promotion (from cold to hot)
- Probabilistic promotion (from hot to hotter, i.e. higher priority)
- Probabilistic insertion (into highest priority cold table slot)
- Probabilistic eviction (one of the cold entries is evicted)


## ProHIT: mechanism

| Row A |
| :---: |
| Row B |
| Row C |
| Row D |
| Row E |
| Row F |
| Row G |
| Row H |



## ProHIT: mechanism



## ProHIT: mechanism



## ProHIT: mechanism



## ProHIT: mechanism



## ProHIT: mechanism

Activate row I


| Row A |
| :---: |
| Row B |
| Row C |
| Row D |
| Row J |
| Row E |
| Row F |
| Row H |

## ProHIT: mechanism



Promote to 'random' hot entry (with probability based on priority)

## ProHIT: mechanism



## After specific <br> ProHIT: mechanismime interval

 Invalidate entry + refresh highestpriority row| Row A |
| :---: |
| Row B |
| Row J |
| Row C |
| Row D |
| Row E |
| Row F |
| Row H |

## ProHIT: mechanism

 Invalidate entry + refresh highestpriority row|  |
| :---: |
| Row B |
| Row J |
| Row C |
| Row D |
| Row E |
| Row F |
| Row H |

## ProHIT: weaknesses

- Still cannot prevent bit-flips with $100 \%$ certainty (probabilistic!)
- But at least we have better performance!
- Knowledge on in-DRAM mapping still needed


## MRLoc: definition

- Based on PARA
- Mitigating Row-hammering based on memory Locality
- Optimizes refresh probability based on memory locality
- If a certain row has been accessed recently, a higher probability is assigned to its corresponding victim rows
- Victim rows are stored in queue


## MRLoc: mechanism



## MRLoc: mechanism



## MRLoc: mechanism



## MRLoc: weaknesses

- Cannot prevent bit-flips with $100 \%$ certainty (probabilistic!)
- Even worse performance now ...
- Knowledge on in-DRAM mapping needed


## Comparison

- Compare BlockHammer with
- (Baseline system: no RH mitigation)
- 3 probabilistic mitigation mechanisms
- PARA
- ProHIT
- MRLoc
- 3 deterministic mitigation mechanisms
- CBT
- TWiCe
- Graphene


## CBT: definition

## - = Counter-Based Tree

- Tree of counters that count row activations in disjoint memory regions
- Whenever parent node reaches certain threshold, memory region is halved (one half for each child)
- Predefined threshold for each level
- Leaf node reaches threshold: counter reset + refresh of respective memory region


## CBT: mechanism

CBT: mechanism

Activate row 1

CBT: mechanism

Activate row 1

CBT: mechanism

Activate row 4

CBT: mechanism

Activate row 4
2
Threshold = 2

## CBT: mechanism



Threshold = 2

Threshold = 5

## CBT: mechanism

Activate row 4


Threshold = 2

Threshold = 5

## CBT: mechanism

Activate row 4


## CBT: mechanism



Threshold = 2
And so on ...

CBT: mechanism

Reset \& Refresh!!


Threshold = 2

Threshold = 5

Threshold = 7

## CBT: mechanism

At end of refresh period (e.g. 64 ms )

## CBT: weaknesses

- Area vs. performance trade-off
- More levels means smaller memory region size and thus more correct refreshes (better performance), but at higher area cost
- Assumes rows are contiguous but might not be the case $\rightarrow$ DRAM remaps addresses internally


## TWiCe: definition

- = Time Window Counter based row refresh
- Maximum number of DRAM ACTs over $t_{\text {REFW }}$ is bounded
- Counter table: Valid bit | Row address | Activation count | Life
- Counter table + counter logic
- Activation count: records number of activations to the target row address
- Valid bit: is entry valid?
- Life: \# consecutive pruning intervals for which entry stays valid in the table


## TWiCe: mechanism

- Row activation
- Not in table $\rightarrow$ allocate entry


(1) Address not found.

New entry inserted.

2) Address found.
act_cnt incremented.

(3) thrн reached.

Victim rows refrest

(4) Table updated
during auto-refresh.

## TWiCe: mechanism

- Row activation
- Not in table $\rightarrow$ allocate entry
- In table $\rightarrow$ increment activation count


| valid | w_add | act_cnt | life |
| :---: | :---: | :---: | :---: |
| 1 | 0x50 | 32,767 | 3 |
| 1 | 0xC0 | 8 | 2 |
| 1 | 0xF0 | 1 | 1 |
| 0 |  |  |  |
| ACT/0x50 |  |  |  |


(3) thri reached.
Victim rows refreshed.

## TWiCe: mechanism

- Activation count reaches threshold $\rightarrow$ refresh victim rows \& set valid bit to 0


(1) Address not found. New entry inserted.


2) Address found.
act_cnt incremented.

(3) thrH reached.

Victim rows refreshed.

(4) Table updated during auto-refresh.

## TWiCe: mechanism

- After each pruning interval
- All entries with activation count $<$ th $_{\text {PI }} \times$ life $\rightarrow$ removed (NOT refreshed)
- Activation count $\geq$ th $_{\text {PI }} x$ life $\rightarrow$ increment life


(1) Address not found

New entry inserted.

(2) Address found.
act_cnt incremented.

(3) thRH reached.

Victim rows refreshed.

(4) Table updated during auto-refresh.

## TWiCe: weaknesses

- Relatively large area overhead as RH gets worse! (in comparison to BH and Graphene)
- Needs to identify victim rows $\rightarrow$ requires knowledge of DRAM internals!


## Graphene: definition

- Misra-Gries algorithm
- Solves frequent elements problem
- Find all elements in a (finite!) stream that occur more than a given fraction of the time
- Here: elements = memory requests


## Graphene: mechanism

- Activate row
- Row in table $\rightarrow$ increase count

| Row Address | Count |
| :---: | :---: |
| $0 \times 1010$ | 5 |
| $0 \times 2020$ | 7 |
| $0 \times 3030$ | 3 |
| Spillover Count | 2 |


|  | Row Address |  |
| :---: | :---: | :---: |
|  | Count |  |
|  | $0 \times 1010$ | 6 |
|  | $0 \times 2020$ | 7 |
|  | $0 \times 3030$ | 3 |
|  | Spillover Count | 2 |



## Graphene: mechanism

- Activate row
- Row not in table AND spillover count < count of all entries $\rightarrow$ increment spillover count


| Row Address | Count |  |
| :---: | :---: | :---: |
| $0 \times 1010$ | 6 |  |
| $0 \times 1010$ | 7 |  |
| $0 \times 3030$ | 3 |  |
|  | Spillover Count | 2 |


| Row Address | Count |  |
| :---: | :---: | :---: |
|  | $0 \times 1010$ | 6 |
|  | $0 \times 2020$ | 7 |
|  | $0 \times 3030$ | 3 |
|  | Spillover Count | 3 |



## Graphene: mechanism

- Activate row
- Row not in table AND spillover count $>=$ count of some entry $X \rightarrow$ replace entry $X$ with new row + increment count of that row


| Row Address |  | Count |  | Row Address |  | Count |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $0 \times 1010$ | 6 |  |  |  |  |  |
| 0 |  | $0 \times 1010$ | 6 |  |  |  |
| $0 \times 2020$ | 7 |  | $0 \times 2020$ | 7 |  |  |
| $0 \times 3030$ | 3 | $0 \times 5050$ |  | $0 \times 5050$ |  |  |
|  |  | 4 |  |  |  |  |
|  |  |  | Spillover Count | 3 |  |  |
|  |  |  |  |  |  |  |

## Graphene: mechanism

- Count $==$ (multiple of) threshold $\rightarrow$ refresh victim rows



## Graphene: weaknesses

- Needs to identify victim rows $\rightarrow$ requires knowledge of DRAM internals

Currently one of the best solutions (has good performance and low area overhead)

## 1. Hardware complexity analysis



## 1. Hardware complexity analysis



## 1. Hardware complexity analysis



## 1. Hardware complexity analysis

## RowHammer threshold 32K

- PARA, PRoHIT, MRLoc $\rightarrow$ extremely area-efficient (because probabilistic)
- Graphene << TWiCe, BlockHammer < CBT $\rightarrow$ still relatively area-efficient
- PARA, PRoHIT, MRLoc $\rightarrow$ extremely area-efficient (because probabilistic)
- Graphene << TWiCe, BlockHammer < CBT $\rightarrow$ still relatively area-efficient


## RowHammer threshold 1K

- Graphene x28.5, TWiCE x34.5, CBT x19.7 $\leftrightarrow$ BlockHammer x11.2
- New order: Graphene < BlockHammer << TWiCE << CBT
- BlockHammer is catching up!


## 1. Hardware complexity analysis



## 1. Hardware complexity analysis

RowHammer threshold 32 K

- PARA, PRoHIT, MRLoc $\rightarrow$ extremely area-efficient (because probabilistic)
- Graphene << TWiCe, BlockHammer < CBT $\rightarrow$ still relatively area-efficient


## RowHammer threshold 1K

- Graphene x23, TWiCE x35, CBT x20 $\leftrightarrow$ BlockHammer x10.7
- New order: Graphene < BlockHammer << TWiCE << CBT
- BlockHammer is catching up!


## 1. Hardware complexity analysis



## 1. Hardware complexity analysis

- PRoHIT, MRLoc $\rightarrow$ extremely efficient (because probabilistic)
- TWiCe < CBT << BlockHammer << Graphene $\rightarrow$ still relatively efficient


## RowHammer threshold 1K

- Graphene x22.6, TWiCE x15.6, CBT x14 $\leftrightarrow$ BlockHammer x4.9
- New order: BlockHammer <<< TWiCE, CBT << Graphene
- BlockHammer is most efficient!


## 1. Hardware complexity analysis

## 1. Hardware complexity analysis

- PRoHIT, MRLoc $\rightarrow$ extremely efficient (because probabilistic)
- Graphene << TWiCe, BlockHammer << CBT $\rightarrow$ still relatively efficient


## RowHammer threshold 1K

- Graphene x30.2, TWiCE x29.7, CBT x15.1 $\leftrightarrow$ BlockHammer x9.9
- New order: Graphene << BlockHammer <<< TWiCE << CBT
- BlockHammer is catching up!


## 2. Performance \& energy consumption

- Setup: DDR4 memory

| Processor | $3.2 \mathrm{GHz},\{1,8\}$ core, 4-wide issue, 128-entry instr. window |
| :--- | :--- |
| Last-Level Cache | 64-byte cache line, 8-way set-associative, 16 MB |
| Memory Controller | 64-entry each read and write request queues; Scheduling <br> policy: FR-FCFS [122, 164]; Address mapping: MOP [60] |
| Main Memory | DDR4, 1 channel, 1 rank, 4 bank groups, 4 banks/bank <br> group, 64K rows/bank |

Table 5: Simulated system configuration.

