



ABACuS: All-Bank Activation Counters for Scalable and Low Overhead RowHammer Mitigation

Ataberk Olgun, Yahya Can Tugrul, Nisa Bostanci, Ismail Emir Yuksel, Haocong Luo, Steve Rhyner, Abdullah Giray Yaglikci, Geraldo F. Oliveira, and Onur Mutlu, *ETH Zurich*

<https://www.usenix.org/conference/usenixsecurity24/presentation/olgun>

This paper is included in the Proceedings of the
33rd USENIX Security Symposium.

August 14-16, 2024 • Philadelphia, PA, USA

978-1-939133-44-1

Open access to the Proceedings of the
33rd USENIX Security Symposium
is sponsored by USENIX.

ABACuS: All-Bank Activation Counters for Scalable and Low Overhead RowHammer Mitigation

Ataberk Olgun Yahya Can Tugrul Nisa Bostanci Ismail Emir Yuksel Haocong Luo
Steve Rhyner Abdullah Giray Yaglikci Geraldo F. Oliveira Onur Mutlu

ETH Zurich

Abstract

We introduce ABACuS, a new low-cost hardware-counter-based RowHammer mitigation technique that performance-, energy-, and area-efficiently scales with worsening RowHammer vulnerability. We observe that both benign workloads and RowHammer attacks tend to access DRAM rows with the same row address in multiple DRAM banks at around the same time. Based on this observation, ABACuS’s key idea is to use a *single shared row activation counter* to track activations to the rows with the same row address in all DRAM banks. Unlike state-of-the-art RowHammer mitigation mechanisms that implement a separate row activation counter for each DRAM bank, ABACuS implements fewer counters (e.g., *only one*) to track an equal number of aggressor rows.

Our comprehensive evaluations show that ABACuS securely prevents RowHammer bitflips at low performance/energy overhead and low area cost. We compare ABACuS to four state-of-the-art mitigation mechanisms. At a near-future RowHammer threshold of 1000, ABACuS incurs *only* 0.58% (0.77%) performance and 1.66% (2.12%) DRAM energy overheads, averaged across 62 single-core (8-core) workloads, requiring *only* 9.47 KiB of storage per DRAM rank. At the RowHammer threshold of 1000, the best prior low-area-cost mitigation mechanism incurs 1.80% higher average performance overhead than ABACuS, while ABACuS requires 2.50× smaller chip area to implement. At a future RowHammer threshold of 125, ABACuS performs very similarly to (within 0.38% of the performance of) the best prior performance- and energy-efficient RowHammer mitigation mechanism while requiring 22.72× smaller chip area. We show that ABACuS’s performance scales well with the number of DRAM banks. At the RowHammer threshold of 125, ABACuS incurs 1.58%, 1.50%, and 2.60% performance overheads for 16-, 32-, and 64-bank systems across all single-core workloads, respectively. ABACuS is freely and openly available at <https://github.com/CMU-SAFARI/ABACuS>.

1 Introduction

Modern DRAM chips are vulnerable to RowHammer [1–13], where repeatedly opening and closing (i.e., activating and precharging, or simply *hammering*) a DRAM row (aggressor row) at a high enough rate can cause bitflips in physically nearby rows (victim rows). DRAM chips become more vulnerable to RowHammer as DRAM storage density increases across DRAM generations [2, 4, 14–19]. The minimum number of row activations needed to induce a RowHammer bitflip, i.e., the RowHammer threshold (N_{RH}), has reduced by more than an order of magnitude in less than a decade [14].¹ As many prior works demonstrate on real systems [1, 2, 4, 15, 20–83], RowHammer bitflips can lead to security exploits that 1) take over a system, 2) leak security-critical or private data, and 3) manipulate safety-critical applications’ behavior in undesirable ways. As a result, a large body of work [1, 15, 19, 38, 44, 55, 84–88, 88–135] proposes mitigation mechanisms to prevent RowHammer bitflips.

Key Problem. Many prior works (e.g., [1, 98, 102, 106, 107, 110, 112, 116, 117, 125, 134, 135]) propose using a set of counters to track the activation counts of potential aggressor rows (counter-based mechanisms). Using counters to determine rows that reach *close to RowHammer thresholds* and taking mitigating actions accordingly can prevent RowHammer bitflips at low performance and energy overheads. Unfortunately, mitigation mechanisms that rely on counters face two scalability challenges. First, they need to implement an increasingly large number of counters to track all potential aggressor rows as N_{RH} reduces. This is because an attacker can concurrently hammer *more* DRAM rows when N_{RH} is smaller. Second, the area overhead of these mechanisms linearly increases with the number of DRAM banks in the system, and modern systems continue to use more banks to scale up both DRAM capacity and bandwidth [136–152]. A small set of prior works (e.g., [1, 97, 100, 106, 116, 133]) aim to mitigate

¹For example, RowHammer threshold (N_{RH}) is *only* 4.8K and 10K for some newer LPDDR4 and DDR4 DRAM chips (manufactured in 2019–2020), which is 14.4× and 6.9× lower than the N_{RH} of 69.2K for some older DRAM chips (manufactured in 2010–2013) [14].

RowHammer at low area overhead. Unfortunately, to achieve low area overhead, these works cause (prohibitively) large performance overheads as DRAM chips become more vulnerable to RowHammer [2, 4, 14–19]. Therefore, it is important to provide a scalable RowHammer solution whose area overhead and performance overhead remain low as DRAM chips become more vulnerable to RowHammer.

Our goal is to prevent RowHammer bitflips at low performance, energy, and area overheads in modern and future DRAM-based systems with high RowHammer vulnerability. To this end, we propose a new low-cost and scalable counter-based RowHammer mitigation mechanism, *All-Bank Activation Counters for Scalable and low overhead RowHammer mitigation (ABACuS)*. ABACuS leverages our **key observation** on benign workloads’ and RowHammer attacks’ memory access patterns. Many workloads (both benign workloads and RowHammer attacks) tend to access DRAM rows with the *same* row address in *multiple* DRAM banks at *around the same time* because i) modern memory address mapping schemes interleave consecutive cache blocks across different banks (§2.1) and ii) workloads tend to access cache blocks in close proximity around the same time due to the spatial locality in their memory accesses (§3).

Leveraging this observation, ABACuS’s **key idea** is to use a *single shared activation counter* to track activations to the rows with the same row ID (i.e., same row address) in all DRAM banks. By doing so, ABACuS i) retains the performance- and energy-efficiency benefits of counter-based RowHammer mitigation mechanisms (§9), and ii) incurs low area cost, as it requires *only* a small number of counters to keep track of many aggressor rows (e.g., 2720 counters instead of 43520 [102] at an N_{RH} of 1000).

Key Mechanism. At a high level, ABACuS maps DRAM rows that have the same row ID in different banks (which we call *sibling rows*) to the same ABACuS counter. In each ABACuS counter, we store i) the *sibling activation vector* that contains as many bits as the number of banks (e.g., 16 bits if there are 16 banks), and ii) the *row activation count*. ABACuS tracks only the maximum (i.e., the worst) activation count of the activation counts of sibling rows. Before the activation count value reaches N_{RH} , ABACuS preventively refreshes all potential victim rows of each sibling row and thus prevents any potential RowHammer bitflips. ABACuS tracks the maximum activation count of sibling rows *without unnecessarily* incrementing the row activation count for each sibling row. This way, ABACuS reduces the number of unnecessary *preventive refresh* operations, lowering its performance and energy overheads. ABACuS is completely implemented inside the memory controller and therefore does *not* require any modifications to existing DRAM chips or software.

Key Results. We rigorously evaluate ABACuS’s i) impact on system performance and energy consumption using cycle-accurate memory system simulations (with Ramulator [153–156]), executing a diverse set of 62 single-core

and 62 8-core multi-programmed workloads from SPEC CPU2006, SPEC CPU2017, TPC, MediaBench, and YCSB benchmark suites and memory-intensive microbenchmarks, and ii) area overhead using CACTI [157]. We model ABACuS’s hardware design (RTL) in Verilog and evaluate its circuit area and latency overheads using modern ASIC design tools. We compare ABACuS to four state-of-the-art RowHammer mitigation mechanisms. We make four key observations from our evaluation. First, at a near-future RowHammer threshold of 1K, ABACuS incurs *only* 1) 0.58% average (32.00% maximum) performance and 1.66% average ($2.02\times$ maximum) DRAM energy overheads across 62 single-core workloads, and 2) 0.77% average (32.97% maximum) performance and 2.12% average ($2.17\times$ maximum) DRAM energy overheads across 62 8-core workload mixes compared to a system without any RowHammer protection, while requiring only 18.93 KiB of storage. Second, ABACuS scales well into the future for DRAM chips with extremely low RowHammer thresholds: e.g., at a RowHammer threshold of *only* 125, ABACuS’s performance and energy overheads are 1.45% and 1.27%, respectively, on average for single-core workloads, while requiring 151.41 KiB of storage. Third, at the N_{RH} of 125, ABACuS performs very similarly to the best prior performance- and energy-efficient RowHammer mitigation mechanism while requiring $22.72\times$ smaller chip area. Fourth, ABACuS scales well with the number of DRAM banks. At the N_{RH} of 125, ABACuS incurs 1.58%, 1.50%, and 2.60% performance overheads for 16-, 32-, and 64-bank systems across all 62 single-core workloads, respectively. Our evaluation of ABACuS’s circuit latency shows that ABACuS could be implemented off the critical path in the memory controller. ABACuS’s latency (1.22 ns) is easily overlapped with the latency (2.5 ns [158]) of issuing two successive DRAM row activation commands (*tRRD*). We open source our simulation infrastructure and all datasets at <https://github.com/CMU-SAFARI/ABACuS> to enable reproducibility and help future research. Detailed analyses and data, including memory intensity characteristics of each evaluated workload, per-workload performance and energy results, key configuration parameters of evaluated state-of-the-art RowHammer mitigations, and security analysis of some of the existing in-DRAM RowHammer mitigations, are in an extended version of this paper [159].

This work makes the following key contributions:

- We show that it is possible to leverage benign workload access patterns to prevent RowHammer bitflips at low overhead in terms of performance, energy, and area, even for DRAM chips with very high RowHammer vulnerability.
- We develop ABACuS, a new low-cost and scalable RowHammer mitigation mechanism. ABACuS prevents RowHammer bitflips with small average performance and energy overheads while using significantly smaller in-processor-chip storage compared to state-of-

the-art RowHammer mitigation mechanisms at very low RowHammer thresholds (i.e., 1K to 125).

- We evaluate the performance, energy, and area overheads of four state-of-the-art RowHammer mitigation mechanisms. We show that ABACuS performs very similarly to the best-performing mechanism at a much smaller (e.g., 22.72×) chip area overhead. We model ABACuS’s hardware design (RTL) in Verilog and evaluate its hardware complexity using modern ASIC design tools.

2 Background

2.1 DRAM Organization and Operation

DRAM Organization (Fig. 1a). A memory channel connects the processor (CPU) to a DRAM rank, a set of DRAM chips working in lockstep. A DRAM chip has multiple banks each of which containing a 2D DRAM cell array in the form of rows and columns. The DRAM cell stores one bit of data in the form of electrical charge in a capacitor. The access transistor, controlled by the wordline, connects the cell to the bitline which is connected to the row buffer.




Figure 1: DRAM organization (a), timing parameters (b), and RowHammer (c)

Operation. The memory controller issues DRAM commands, e.g., row activation (*ACT*), bank precharge (*PRE*), data read (*RD*), data write (*WR*), and refresh (*REF*) to serve memory requests. To do so, the memory controller first issues an *ACT* command with the bank address (i.e., bank ID) and row address (i.e., row ID) corresponding to the memory request’s address, activating the row. When a row is activated, its data is copied to the row buffer. The memory controller can read/write data at cache block (512 bits) granularity from/to the row buffer using a sequence of *RD/WR* commands. A subsequent access to the same row (i.e., a *row hit*) can be served quickly without issuing another *ACT*. To access another row (i.e., to serve a *row conflict*), the memory controller must issue a precharge command and close the open row.

Periodic Refresh. DRAM cells are inherently leaky and thus lose the stored electrical charge over time. To maintain data integrity, a DRAM cell is periodically refreshed every refresh window (t_{REFW}), which is typically 64 ms (e.g., [158, 160, 161]) or 32 ms (e.g., [162–164]). To timely refresh all cells, the memory controller periodically issues a refresh (*REF*) command every refresh interval (t_{REFI}), which is typically 7.8 μ s (e.g., [158, 160, 161]) or 3.9 μ s

(e.g., [162–164]). Upon receiving a *REF* command, the DRAM chip internally refreshes multiple DRAM rows for a refresh latency (t_{RFC}) amount of time.

DRAM Timing Parameters. The memory controller schedules DRAM commands according to certain timing parameters to guarantee correct operation [136, 143, 148, 158, 160, 162–166]. In addition to t_{REFW} , t_{REFI} , and t_{RFC} , two other timing parameters are relevant for this work: i) the minimum time between two consecutive row activations targeting the same bank (t_{RC}) and ii) the minimum time between two consecutive row activations targeting the same rank (t_{RRD}) (Fig. 1b).

Bank-Level Parallelism. Main memory accesses that target different banks can proceed concurrently [136]. Modern address mapping schemes (e.g., [167–171]) aim to interleave consequently addressed cache blocks across different banks to exploit *bank-level parallelism* [136, 138].

2.2 RowHammer Mitigation Techniques

To prevent RowHammer bitflips and protect computing systems against RowHammer attacks, prior works propose different RowHammer mitigation mechanisms [1, 15, 19, 38, 44, 55, 84–88, 88–135]. These works trigger their countermeasure (e.g., refreshing potential victim rows or throttling accesses to potential aggressor rows) based on either i) the result of a probabilistic procedure or ii) tracking the number of times DRAM rows are activated (i.e., row activation counts). While probabilistic procedures can be implemented at low chip area cost, they incur prohibitively large performance overheads when configured for sub-1K RowHammer threshold (N_{RH}) values [14, 104, 116]. Prior works propose several different row activation tracking mechanisms that detect the frequently-activated set of rows. Unfortunately, while providing better performance than the probabilistic mechanisms, the chip area overhead of these row-activation-count-tracking mechanisms significantly increases as DRAM chips become more vulnerable to RowHammer [13, 106].²

Frequent Item Counting. A naive, area-inefficient tracking method to detect possible aggressor rows is to store the activation count of each DRAM row in a dedicated counter. However, this method leads to impractical on-chip area overheads when used to protect modern, high-density DRAM modules. For example, 8-bit counters for a modern DDR4 rank with 2^{21} rows [158] would require 2 MiB on-chip storage and a newer and denser DDR5 rank with 2^{23} rows [163] would require an even larger 8 MiB on-chip storage. Fortunately, the problem of tracking the frequently activated DRAM rows can be interpreted as a frequent item counting problem and can be solved using more area-efficient algorithms. For example, the Misra-Gries algorithm [172] can be implemented in hardware to accurately track aggressor rows using a relatively small

²Hydra [106] is an exception in this classification as it incurs a low chip area overhead while tracking row activation counts. Hydra achieves this by storing the counters in the DRAM array and caching them in the memory controller. §3 discusses Hydra’s scalability limitations.

number of counters to detect potential aggressor rows, and its variants are adopted by several prior RowHammer mitigation mechanisms [102, 107, 110, 112, 117].

3 Motivation

Repeatedly activating and precharging (hammering) a DRAM row *at least* N_{RH} times in a refresh window induces one or multiple bitflips in that row. As DRAM chips become more vulnerable to RowHammer (i.e., the chip's rows have smaller N_{RH} values), fewer hammers can induce bitflips. Even though the number of activate and precharge commands that the memory controller can issue in a refresh window remains the same, more rows can be concurrently hammered N_{RH} times at a smaller N_{RH} . As RowHammer vulnerability increases, state-of-the-art counter-based RowHammer mitigation mechanisms need to track more DRAM rows and implement more activation counters.

A common method of increasing memory bandwidth and capacity is to increase the number of DRAM banks [136–141]. However, as the number of banks increases, counter-based mechanisms incur increasing chip area overhead.

Limitations of Prior Work. Several prior works [1, 98, 100, 105, 106, 126, 173] aim to mitigate RowHammer at low area overhead by implementing a limited set of row activation counters (i.e., fewer counters than there are rows in the DRAM module) at the cost of reduced tracking accuracy. However, a RowHammer mitigation countermeasure (e.g., preventively refreshing potential victim rows or throttling accesses to potential aggressor rows) fundamentally consumes memory bandwidth and reduced tracking accuracy exacerbates the number of countermeasures deployed by the mitigation mechanism. Thus, these mechanisms occupy a significant portion of main memory bandwidth and incur large system performance and DRAM energy overheads at small N_{RH} values. To provide more insight into this problem, we describe the key drawback of one such state-of-the-art mechanism, Hydra [106], as a concrete example. Hydra [106] maintains the activation count of each DRAM row in a physical location in main memory (i.e., in the DRAM chips). To minimize the overheads of fetching the counters from the main memory, Hydra implements a filtering and caching logic. The filtering logic groups a number of (e.g., 125) DRAM rows into row groups and assigns a counter to each row group called the *group counter*. DRAM row activations update *only* the corresponding group counters at the beginning of a refresh window. When a group counter exceeds a predetermined *group count threshold* (e.g., 400), the group counter's value is copied to the activation counters of rows in that group, such that Hydra can track each row's activation count individually and deploy its countermeasure (preventively refreshing victim rows) more accurately (e.g., instead of preventively refreshing *all* 125 rows in a group, Hydra can refresh one or several DRAM rows that are activated frequently in the group of 125 rows, depending on workload memory access patterns), and the group counter is no longer queried.

Hydra's mechanism has two *key drawbacks*. First, Hydra overestimates the activation counts of DRAM rows, causing a large number of unnecessary refresh operations. According to our system-level simulations (in §9), *approximately half* of Hydra's preventive refresh operations are unnecessary for 62 single-core workloads at a very low RowHammer threshold of 125. Hydra overestimates activation counts of DRAM rows because modern memory-intensive workloads can rapidly increase the group counter value to the group count threshold value with *only* a few activations to each DRAM row in the group. Such workloads can cause the activation counters to overestimate the actual activation count of each row in the group by up to 396 (in Hydra's default configuration for an N_{RH} of 1K). Therefore, Hydra often mistakenly refreshes the neighbors of many rows that will *not* be activated as many as N_{RH} times. Second, suppose the counter of an accessed row is not cached in the memory controller. In that case, Hydra needs to fetch the counter from the main memory, which incurs additional memory latency for writing back the evicted counter and fetching the new counter. Both of these drawbacks incur significant performance and energy overheads as Hydra increases the memory latency (i.e., the time it takes to serve a memory demand request) by 23.67% on average at $N_{RH} = 125$ (as we show in detail in §9).

3.1 Motivational Analysis for ABACuS

We investigate memory access patterns of modern memory-intensive workloads and existing RowHammer attacks. We observe that they activate DRAM rows with the same row address in multiple DRAM banks (i.e., sibling rows) at around the same time. This observation motivates us to design a performance-, energy, and area-efficient DRAM row activation count tracking mechanism by implementing *one shared activation counter* for all sibling rows. Implementing one shared activation counter reduces the number of counters required to track aggressor rows (and thus the area cost) by a factor of the number of banks (e.g., 16 in DDR4 [158]) compared to the aggressor row tracking mechanisms used by the state-of-the-art performance- and energy-efficient RowHammer mitigations (e.g., Graphene [102], Panopticon [134], and PRHT [135]). However, the shared counter may not accurately represent the activation counts of multiple sibling rows because the shared counter can store only one activation count. Misrepresenting the activation counts of sibling rows may induce performance and energy overheads due to unnecessary victim row refresh operations. In the remainder of this section, we show that 1) sibling rows are activated at around the same time, and 2) a shared activation counter can accurately represent the activation counts of multiple sibling rows.

We simulate 34 memory-intensive workloads, each having more than two row buffer misses per kilo instructions (RBMPKI), and three variations of the double-sided (ds) [14, 17, 28, 31, 32, 39, 41] and many-sided (ms) [15, 54–56, 58, 62] on a 32-bank system using the simulation methodology that

we explain in §8. We carefully create memory traces (load and store requests that arrive at main memory) of double- and many-sided attacks 1) without prefetching, 2) with a simple prefetcher that prefetches the *next cache line* (p1) [174, 175], 3) the *next eight cache lines* (p8), 4) and the *next 32 cache lines* (p32) for every load request. We name a RowHammer attack trace as the concatenation of its type (ds or ms) and the prefetcher configuration used in creating the trace (p1, p8, or p32). For example, *ds-p32* is the double-sided RowHammer attack with the next 32 cache line prefetcher. Fig. 2 shows how many sibling rows get activated before one of the sibling rows is activated *again*, on average across all DRAM row activations (y-axis) for each simulated workload (x-axis).³ Benign workloads are ordered from left to right in increasing memory intensity (in terms of row buffer misses per kilo instructions) in the figure. We highlight the highest possible y-axis value (31) with a red line on the plot. A workload with a bar closer to this line indicates that the workload accesses *all* sibling rows at around the same time.




Figure 2: Number of sibling rows activated before one sibling row is activated again

In the upper extreme case (at $y = 31$)⁴, the workload *always* activates *all* sibling rows once before activating any sibling row for the second time. This property of the workload makes it a good fit for using a single activation counter shared between sibling rows. For us to know when an aggressor sibling row has been activated N_{RH} times, the shared counter stores the highest activation count across all sibling rows. Because the workload activates all sibling rows once before activating any other for the second time, the shared counter accurately represents the activation count of every sibling row. In the lower extreme case (at $y = 0$), the workload *never* activates two or more sibling rows. For this type of a workload, the single shared activation counter’s value misrepresents almost all of the sibling rows’ activation counts (which are 0).

We make three key observations from Fig. 2. First, on average across all workloads, 12.8 sibling rows get activated before any sibling receives another activation. We observe that some workloads activate at least *three* sibling rows *once*,

³The box is lower-bounded by the first quartile (i.e., the median of the first half of the ordered set of data points) and upper-bounded by the third quartile (i.e., the median of the second half of the ordered set of data points). The inter-quartile range (*IQR*) is the distance between the first and third quartiles (i.e., box size). Whiskers show the central 90th percentile of the distribution.

⁴A DRAM row has 31 sibling rows in a system with 32 banks.

while some activate up to 25 sibling rows (out of 31), before activating any sibling row *again*. Second, the average sibling row activation count does *not* significantly correlate with the memory intensity of the workload. Third, the RowHammer attacks can activate up to 31 sibling rows before activating any sibling row again due to the prefetch requests generated by the simple next cache line prefetcher. From these observations, we conclude that after accessing a row with the address R in a bank, the rows at address R in other banks (i.e., sibling rows) are also likely to be activated. We hypothesize that this access pattern occurs due to two properties: i) the memory address mapping schemes that aim to increase memory bank-level parallelism to improve system performance (§2.1)⁵ and ii) the intrinsic spatial locality in workloads’ main memory accesses.

To strengthen our motivation for sharing an activation counter between sibling rows, we plot the distribution of the activation count of each sibling row when at least one sibling row gets activated N_{RH} times. In other words, one point in the distribution is an activation count of a row. One of this row’s siblings has been activated N_{RH} times. Fig. 3 shows how many times a sibling row gets activated (y-axis) before one of the sibling rows gets activated N_{RH} times for $N_{RH} = 500$, 250, and 125 (different subplots) across benign workloads and RowHammer attacks (x-axis). We highlight the highest possible y-axis values for each N_{RH} value.




Figure 3: The distribution of the number of activations a sibling row receives before any sibling row gets activated N_{RH} times

We make two observations from Fig. 3. First, a sibling row gets activated 99, 194, and 370 times for N_{RH} values of 125, 250, and 500, on average across all workloads. This indicates that when any sibling row gets activated N_{RH} times, the activation counts of all sibling rows are (very) close to N_{RH} . Second, as N_{RH} reduces, the gap between the average activation count of a sibling row and the sibling row with the highest activation count becomes smaller in proportion. For example, the *bfs_cm203* workload, on average, activates

⁵Our extended version [159] describes these schemes in more detail.

sibling rows 272 (54.4% of the N_{RH} of 500) and 108 (86.4% of the N_{RH} of 125) times for $N_{RH} = 500$ and 125, respectively.⁶

From our analysis, we conclude that a single shared activation counter, which stores the highest activation count among the activation counts of all sibling rows, can reasonably accurately represent the activation count of all sibling rows. This property of the shared activation counter becomes stronger as N_{RH} reduces from 500 to 125.

4 Mechanism

Overview. ABACuS is designed to prevent RowHammer bitflips at low performance, energy, and area overhead. Achieving low performance and energy overheads requires accurately identifying aggressor rows and preventively refreshing victim rows *only* when necessary. To this end, ABACuS adopts the Misra-Gries algorithm [172] (§2) to track aggressor rows, similar to prior work [102, 107, 110, 112, 117]. However, Misra-Gries alone *cannot* prevent RowHammer bitflips at low area cost (§7.1). Thus, ABACuS performs Misra-Gries tracking using *shared activation counters* to significantly reduce the area overhead of implementing many counters.

ABACuS’s key idea is to share a *row activation counter* among rows that have the same row ID across all banks (which we call *sibling rows*). The shared row activation counter tracks the maximum activation count among the sibling rows. ABACuS preventively refreshes all the neighboring rows of all the sibling rows tracked by the same row activation counter just enough *before* the row activation counter’s value reaches N_{RH} within a t_{REFW} to prevent RowHammer bitflips (i.e., none of the sibling rows’ activation count reaches N_{RH} within t_{REFW}).

While ABACuS tracks the maximum activation count among sibling rows, it does *not* unnecessarily increment the shared row activation counter with each sibling row activation. For example, when a workload activates multiple sibling rows *only once* (which is a common behavior we observe in §3), it is sufficient for ABACuS to increment the shared activation counter *by only one*. After a sibling row is activated and ABACuS increments the shared row activation counter, other sibling rows can be activated *at most once* before the shared counter is incremented again. To allow for multiple sibling rows to be activated *without* unnecessarily incrementing the shared row activation counter, ABACuS maintains a bit vector for the counter, *sibling activation vector*, that stores which sibling rows were activated *since the shared counter was last incremented*. ABACuS increments the shared row activation counter only when the bit corresponding to the activated sib-

⁶The gups workload *cannot* activate any same row 125 times because a workload is limited in the number of DRAM row activations it can issue to a DRAM chip by DRAM timing parameters (e.g., four row activation window, t_{FAW} [158]) and gups’ row activations are randomly and evenly distributed to all 128K DRAM rows. A workload that fully exercises the available DRAM row activation bandwidth can *only* issue 12’190’476 activate commands in a refresh window (64 ms) due to t_{FAW} (21 ns) timing constraint [158]. Therefore, randomly and evenly distributing 12’190’476 activate commands to 128K rows would activate a row at most 94 times.

ling row is already set (i.e., the sibling row was activated once since the shared counter was last incremented).

4.1 ABACuS’s Hardware Design

Fig. 4 presents ABACuS’s key components. ABACuS is placed inside the memory controller. The ABACuS counter table contains (❶) multiple ABACuS counters, each mapped to a row ID. There are exactly $N_{entries}$ ABACuS counters in the ABACuS counter table. An ABACuS counter (❷) consists of a row activation counter (RAC) of size S_{RAC} bits and a sibling activation (bit) vector (SAV) of size S_{SAV} bits. The ABACuS controller (❸) dynamically maps (not shown in the figure) a row ID to a counter during runtime and uses a *spillover counter* (❹) to track the maximum activation count of all DRAM rows that do *not* have an ABACuS counter assigned (§4.2). We explain how we determine the sizes of each key component (S_{RAC} , S_{SAV} , and $N_{entries}$) in §4.3.

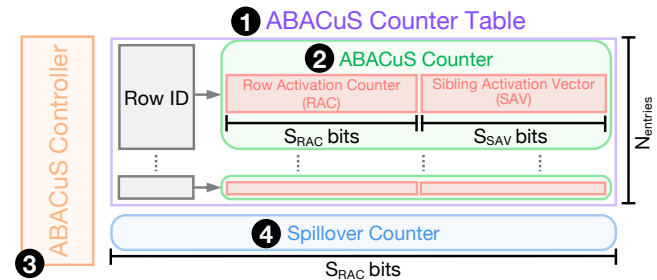


Figure 4: Key components of ABACuS

The row activation counter (RAC) in an ABACuS counter (❷) stores the maximum activation count across all sibling rows’ activation counts. The sibling activation vector (SAV) stores the sibling activation bits used by ABACuS to increment the RAC *only* when necessary (see §4). The ABACuS controller updates RAC and SAV to ensure the maximum activation count among all sibling rows is tracked in the RAC.

4.2 Operation of ABACuS

We describe ABACuS’s operation in five key steps.

(1) **Initialization and Reset.** Initially (at system power on) and after periodic ABACuS counter table reset (Step 5), no DRAM row is activated for the last t_{REFW} . Thus, row activation counters (RACs) and sibling activation vectors (SAVs) in all ABACuS counters and the spillover counter all store 0. (2) **ABACuS Counter Table Search.** The memory controller issues an *ACT* command to a row ID in a bank. Consequently, the ABACuS controller searches all row ID mappings to find if the activated row is already tracked by an ABACuS counter. No row is tracked immediately after initialization and reset. Thus, the ABACuS controller needs to map a row ID to an ABACuS counter. To find the counter that will be mapped to the activated row’s ID, the ABACuS controller looks for an ABACuS counter whose RAC stores the same value as the spillover counter’s value. If a RAC and the spillover counter have the same value, the RAC’s row ID is replaced with the activated row’s ID (Step 3). In contrast, if an ABACuS




Figure 5: ABACuS workflow using four activate (*ACT*) commands. We highlight state changes using black boxes and red text.

counter already tracks the row ID, the ABACuS controller updates the matching counter (Step 4). In case no RAC value equals the spillover counter’s value (i.e., the activated row’s ID *cannot* be mapped to an ABACuS counter) and there is no ABACuS counter that already tracks the activated row’s ID, the ABACuS controller increments the spillover counter value. When the spillover counter reaches a predefined value of the refresh cycle threshold (RCT), ABACuS issues t_{REFW}/t_{REFI} refresh (*REF*) commands to refresh all DRAM rows in the DRAM rank and resets all counters. We call the time when the memory controller refreshes all DRAM rows due to the spillover counter reaching the RCT a *refresh cycle*.

(3) ABACuS Counter Mapping and Replacement. The ABACuS controller maps the newly-activated row ID to the matching ABACuS counter. To correctly track the maximum activation count, the ABACuS controller i) initializes RAC with spillover counter value + 1, and ii) sets the bit in the SAV that corresponds to the bank ID of the activated row.

(4) ABACuS Counter Update. The ABACuS controller checks the SAV bit value corresponding to the activated row’s bank ID. If the SAV bit is *not* set (i.e., stores logic-0), i.e., the activated row is activated *for the first time* since the RAC was last incremented, the ABACuS controller sets the SAV bit. If the SAV bit is set (i.e., stores logic-1), i.e., the activated row is activated *again* since the RAC was last incremented, the ABACuS controller increments the RAC by 1. After incrementing the RAC, the ABACuS controller sets the SAV bit corresponding to the bank ID of the activated row and resets all other bits. This way, a set bit corresponding to the bank ID of the activated row in the SAV *still* indicates that the corresponding row was activated *once* since the RAC was last incremented. If the RAC’s value is a multiple of the *preventive refresh threshold (PRT)*, i.e., one of the sibling rows tracked by the RAC was activated PRT times since all the sibling rows’ victims were preventively refreshed, ABACuS performs preventive refresh operations to all the victim rows (the neighbors of the activated row) in all banks.

(5) Reset Period. It is sufficient for ABACuS to track the activation counts within a t_{REFW} , after which all DRAM rows are refreshed. Therefore, we reset the ABACuS counters and the spillover counter after every t_{REFW} (i.e., ABACuS’s reset period is t_{REFW}). After periodic reset, ABACuS’s state becomes as described in Step 1.

Determining the Preventive Refresh Threshold (PRT). ABACuS’s preventive refresh threshold is set accordingly to

prevent aggressor rows from being activated $N_{RH}/2$ times in a reset period. ABACuS is configured in this way because ABACuS does *not* precisely know when each row is periodically refreshed: A potential aggressor row might be hammered for $PRT - 1$ times before its neighbors are refreshed and ABACuS is reset. After ABACuS is reset, an attacker can hammer the same aggressor row for $2 * PRT$ times, accumulating a total activation count of $2 * PRT - 1$ on the aggressor row. Thus, we set PRT to $N_{RH}/2$.

4.3 Configuring ABACuS

Depending on the system’s RowHammer vulnerability (typically measured using N_{RH}), ABACuS has three key parameters that are configured at design time. First, the number of entries ($N_{entries}$) in the ABACuS counter table. Second, the size of each row activation counter (S_{RAC}). Third, the size of each sibling activation vector (S_{SAV}).

Configuring the Number of Entries ($N_{entries}$). We determine the number of entries based on how many rows can be hammered in one DRAM bank during one ABACuS reset period (64 ms) given i) the preventive refresh threshold (PRT), ii) t_{RC} , and iii) t_{RFC} as described in a prior work that adopts Misra-Gries tracking [102]. We first calculate how many *ACT* commands can be issued by the memory controller in a reset period when the bank is *not* unavailable due to periodic refresh (N_{ACT}) as $t_{REFW} * (1 - t_{RFC}/t_{REFI})/t_{RC}$. Thus, at most N_{ACT}/PRT rows can be activated PRT times in an ABACuS reset period in a bank. Setting the number of entries to N_{ACT}/PRT is sufficient to ensure that all hammered rows in an ABACuS reset period in one bank are tracked in one of the counters.

Configuring the Size of Row Activation Counters (S_{RAC}). The number of activations that the memory controller can issue in a t_{REFW} determines the row activation counter’s size. Thus, a row activation counter should be $S_{RAC} = \lceil \log_2((N_{ACT} * t_{RC})/t_{RRD}) \rceil$ bits large. However, we can reduce S_{RAC} to $\lceil \log_2(PRT) \rceil$ by adding an *overflow* bit [102] to each ABACuS counter. Using the overflow bit, we make sure that ABACuS preventively refreshes rows according to the Misra-Gries tracking algorithm [102]: When a RAC reaches PRT , ABACuS sets the overflow bit of the RAC and resets the RAC’s value. The set overflow bit in a RAC indicates that ABACuS should *not* replace the row ID tracked by this RAC, even if the RAC’s value equals the value of the spillover counter. The overflow bit is reset after each reset period.

Configuring the size of sibling activation vectors (S_{SAV}). A SAV contains as many bits as there are banks. Thus $S_{SAV} = N_{banks}$ bits. For example, there are 32 banks in a dual-rank DDR4-based system [158], making S_{SAV} 32-bit large.

4.4 Example ABACuS Workflow

Fig. 5 shows an example ABACuS workflow with three ABACuS counters and four sibling rows. We show how four *ACT* commands cause state changes in the three ABACuS counters. The first *ACT* command (❶) updates the sibling activation vector (SAV) of its ABACuS counter as the bank ID corresponds to a zero bit in the SAV. The second *ACT* command (❷) increments the row activation counter of its ABACuS counter as the bank ID corresponds to a non-zero bit in the SAV. The third *ACT* command (❸) replaces the second ABACuS counter with row ID 20 because i) row ID 20 is not tracked by any counter, and ii) the spillover counter value is equal to the second ABACuS counter’s RAC value. The fourth *ACT* command (❹) increments the spillover counter as i) row ID 7 is not tracked by any ABACuS counter, and ii) the spillover counter value is smaller than all RAC values.

5 Security Analysis

ABACuS preventively refreshes the victim rows of a potential aggressor row before the aggressor row is activated N_{RH} times in a t_{REFW} . Assuming ABACuS accurately stores the maximum activation count across all sibling rows, the Misra-Gries-based tracking technique guarantees that no aggressor row is activated more than the preventive refresh threshold in a t_{REFW} [102]. ABACuS accurately maintains the maximum activation count in the row activation counters because the row activation counter’s value is incremented 1) when *any* sibling row is activated for the first time, and 2) when a sibling row whose sibling activation vector bit is set (i.e., the sibling row was activated after the row activation counter’s value was last incremented) is activated. Appendix A formally analyzes and shows that the row activation counter always stores the maximum activation count.

6 Accounting for RowHammer Blast Radius

An aggressor row can cause bitflips in victim rows that are *not* physically adjacent [1, 14]. The impact of RowHammer on a victim row decreases and eventually disappears as the physical distance between the victim and the aggressor rows increases. To account for this characteristic, prior works define *blast radius* as the distance between an aggressor row and its furthest victim row [1, 14, 15, 17, 50, 54–56, 58, 102–104, 123, 126, 176]. A recent RowHammer attack, Half Double [58], exploits blast radius to induce bitflips with a significantly lower activation count. To account for blast radius and address Half Double, ABACuS 1) preventively refreshes all potential victim rows within the blast radius and 2) counts each preventive refresh as an additional activation. We configure ABACuS

and other state-of-the-art mechanisms with a blast radius of one in our performance and energy evaluation (§8).

7 Hardware Implementation

ABACuS is implemented in the memory controller. It does *not* require *any* modifications to existing DRAM chips.

Key Components. ABACuS’s hardware implementation consists of two components: i) the ABACuS counter table, and ii) the spillover counter. The ABACuS counter table contains: i) the Row ID Table (RIT), ii) the RAC Table (RACT), and iii) the SAV Table (SAVT). To efficiently track the number of activations, ABACuS searches and updates the RIT and the RACT. Thus, we implement RIT and RACT using content-addressable memory (CAM) arrays. We implement SAVT as an SRAM array since ABACuS does *not* search SAVT entries. A register stores the spillover counter’s value.

Performing Preventive Refresh. Since the standard *REF* command is row-address-agnostic in DRAM standards [147, 158, 163], ABACuS cannot use standard refresh commands to refresh detected victim rows. To remain compatible with existing DRAM chips and interface standards, ABACuS performs a preventive refresh operation by accessing a victim row once using *ACT* and *PRE* commands. When a tracked row’s RAC value reaches PRT, ABACuS performs preventive refresh operations to victim rows in all banks. ABACuS prioritizes preventive refreshes over other memory requests: the memory controller does *not* serve any memory request *to the same bank* until the victim rows are preventively refreshed.

7.1 Area Overhead

We evaluate ABACuS’ and four other state-of-the-art mitigation mechanisms’ [1, 102, 106, 177] (their configuration details are explained in §8) chip area, static power, and memory array access energy using CACTI [157]. Table 1 summarizes our results. We perform this analysis at N_{RH} values of 1000 and 125. Table 2 shows ABACuS’ key parameters.

All three ABACuS hardware structures (Row ID Table, RACT, and SAVT) contain $N_{entries}$ entries. At a near future N_{RH} of 1000, we estimate ABACuS’s overall area overhead to be $0.04mm^2$ per DRAM channel for a dual-rank system. ABACuS consumes approximately 0.02% of the chip area of a high-end Intel Xeon processor with four memory channels [178]. At a low N_{RH} of 125, ABACuS’s estimated chip area cost increases to $0.25mm^2$, taking up *only* approximately 0.11% of the same processor’s area.

Area Comparison. At an N_{RH} of 1000, ABACuS takes up $20.25\times$ and $2.50\times$ smaller chip area than Graphene [102] and Hydra [106], respectively. Graphene’s area overhead is larger than other mechanisms because it implements a large number of counters (e.g., 2720 per bank, 87040 in total for a dual-rank DDR4 system). REGA [177] takes 2.06% DRAM chip area to implement. Compared to ABACuS’s memory controller chip area requirement, REGA requires a larger DRAM chip

Table 1: Area, energy, power of ABACuS vs. state-of-the-art RowHammer mitigation mechanisms for a 2-rank memory system

Mitigation Mechanism	$N_{RH} = 1000$					Access Energy (pJ)	Static Power (mW)	$N_{RH} = 125$					Access Energy (pJ)	Static Power (mW)
	SRAM KB	CAM KB	mm ²	% CPU	% DRAM			SRAM KB	CAM KB	mm ²	% CPU	% DRAM		
ABACuS	10.63	8.30	0.04	0.02	-	24.36	12.19	85.00	66.41	0.25	0.11	-	36.87	50.39
Row ID Table	-	5.64	0.01	< 0.01	-	11.23	6.59	-	45.16	0.12	0.05	-	20.64	27.42
Row Activation Counter Table	-	2.66	0.02	< 0.01	-	11.13	4.66	-	21.25	0.06	0.03	-	11.66	15.53
Sibling Activation Vector	10.63	-	0.01	< 0.01	-	1.99	0.95	85.00	-	0.07	0.03	-	4.57	7.44
PARA [1]	-	-	-	< 0.01	-	-	-	-	-	-	< 0.01	-	-	-
Graphene [102]	-	286.51	0.81	0.35	-	876.85	188.64	-	2037.09	5.68	2.43	-	1042.49	1385.52
Hydra [106]	-	61.56	-	0.10	0.04	-	43.07	-	56.5	-	0.07	-	40.25	23.14
REGA [177]	-	-	-	-	2.06	-	-	-	-	-	-	2.06	-	-

Table 2: ABACuS Parameters

Term	Definition	Value			
N_{RH}	RowHammer threshold	1000	500	250	125
PRT	Preventive refresh threshold	500	250	125	62
RCT	Refresh cycle threshold	498	248	123	60
$N_{entries}$	Number of table entries	2720	5440	10880	21760
S_{SAV}	Bit-length of a SAV entry	32	32	32	32
S_{RID}	Bit-length of a Row ID entry	17	17	17	17
S_{RAC}	Bit-length of a RAC entry	10	9	8	7

area. PARA [1] does not maintain any state, thus it has *no* significant area overhead.

We repeat our area overhead analysis for future DRAM chips by scaling the RowHammer threshold down to 125. Although ABACuS’s area overhead increases as it implements a larger number of ABACuS counters at the lower N_{RH} , ABACuS still consumes a relatively small 0.11% processor chip area at an N_{RH} of 125. ABACuS requires $22.72\times$ less chip area to implement than Graphene. ABACuS’s area overhead at this very low N_{RH} is $3.57\times$ that of Hydra’s, however, Hydra incurs up to a very large 85.42% performance overhead for 8-core memory-intensive workloads at the same N_{RH} (see §9.1). Hydra’s chip area overhead reduces with decreasing N_{RH} as it requires counters with fewer bits of storage each. We conclude that ABACuS’s chip area requirement scales better than Graphene’s and that ABACuS’s area requirement at low N_{RH} is closer to the most area-efficient state-of-the-art mitigation mechanism, Hydra.

Energy and Static Power Comparison. For $N_{RH} = 1000$, ABACuS has $36.00\times$ and $1.77\times$ smaller access energy than Graphene and Hydra, respectively. ABACuS consumes $12.19mW$ of static power, which is $15.47\times$ and $1.99\times$ smaller than Graphene and Hydra’s static power consumptions. As N_{RH} reduces to 125, ABACuS’s static power and access energy scale more efficiently (similarly to Hydra) compared to Graphene, where Graphene has $28.27\times$ and $27.50\times$ the access energy and static power of ABACuS, respectively.

7.2 Latency Analysis

We implement ABACuS in Verilog and use Synopsys DC [179] to evaluate ABACuS’s latency impact on memory accesses. ABACuS needs 1.22 ns to update the ABACuS counter of an activated DRAM row. This latency overlaps with the latency of regular memory controller operations as it is smaller than t_{RRD} (e.g., 2.5 ns in DDR4 [158, 161]).

8 Evaluation Methodology

We evaluate ABACuS’s performance and energy consumption with Ramulator [153, 154], a cycle-accurate DRAM simulator, and DRAMPower [180]. We specify our simulated system’s configuration in Table 3.

Table 3: Simulated System Configuration

Processor	1 or 8 cores, 3.6GHz clock frequency, 4-wide issue, 128-entry instruction window
DRAM	DDR4, 1 channel, 2 rank/channel, 4 bank groups, 4 banks/bank group, 128K rows/bank, 3200 MT/s
Memory Ctrl.	64-entry read and write requests queues, Scheduling policy: FR-FCFS [181, 182] with a column cap of 16 [183], Address mapping: MOP [167, 169] 45 ns t_{RC} , 7.9 μ s t_{REFI} , 64 ms t_{REFW} 64 ms ABACuS reset period
Last-Level Cache	2 MiB per core

Address Mapping. Fig. 6 depicts our address mapping scheme. We use an address mapping scheme that interleaves consecutive cache blocks in the physical address space between different DRAM banks.




Figure 6: Simulated address mapping

Comparison Points. We compare ABACuS to a baseline system with no RowHammer mitigation and to four state-of-the-art RowHammer mitigation mechanisms: (1) Graphene [102] implements per bank counters to track the possible aggressor rows using Misra-Gries algorithm [172]. When a counter value exceeds a threshold value, Graphene issues preventive refreshes to the victim rows. (2) Hydra [106] implements a group count table to track activations for a group of rows and a row count table to track per row activations. We configure Hydra such that all rows in a row group have their row count table entries reside in two consecutive cache blocks (64 bytes each). The row count table is stored in the DRAM and cached in the memory controller. Hydra performs preventive refresh operations when a counter exceeds a threshold value. (3) REGA [177] augments DRAM design such that one or more victim rows can be refreshed when a DRAM row is activated. REGA tunes its protection guarantees by changing the default t_{RC} value. A smaller t_{RC} allows REGA to refresh

more rows concurrently with a DRAM row activation, at the cost of increased access latency. To simulate REGA, we modify t_{RC} as described in [177]. (4) PARA [1] protects against RowHammer by performing probabilistic adjacent row activation. When a row is closed (i.e., when the memory controller issues a precharge (PRE)), PARA issues preventive refreshes to the adjacent rows based on a probability threshold. We tune the probability threshold of PARA for a target failure probability of 10^{-15} within a 64 ms as in prior work [104].

Workloads. We evaluate 62 single-core and 62 homogeneous multi-programmed 8-core workloads from five benchmark suites: SPEC CPU2006 [184], SPEC CPU2017 [185], TPC [186], MediaBench [187], and YCSB [188]. Based on the row buffer misses-per-kilo-instruction (RBMPKI), we group the applications into three categories, which Table 4 describes: (1) L (low memory-intensity, $RBMPKI \in [0, 2)$), (2) M (medium memory-intensity, $RBMPKI \in [2, 10)$), (3) H (high memory-intensity, $RBMPKI \in [10, +)$). To do so, we obtain the RBMPKI values of the applications by analyzing each application’s SimPoint [189] traces (200M instructions). All of these traces are open-sourced [190].

Table 4: Evaluated single-core workloads

RBMPKI	Workloads
[10+] (High)	519.lbm, 459.GemsFDTD, 450.soplex, h264_decode, 520.omnetpp, 433.milc, 434.zeusmp, bfs_dblp, 429.mcf, 549.fotonik3d, 470.lbm, bfs_ny, bfs_cm2003, 437.leslie3d, gups
[2, 10) (Medium)	510.parest, 462.libquantum, tpch2, wc_8443, ycsb_aserver, 473.atar, jp2_decode, 436.cactusADM, 557.xz, ycsb_cserver, ycsb_eserver, 471.omnetpp, 483.xalancbmk, 505.mcf, wc_map0, jp2_encode, tpch17, ycsb_bserver, tpcc64, 482.sphinx3
[0, 2) (Low)	502.gcc, 544.nab, h264_encode, 507.cactuBSSN, 525.x264, ycsb_dserver, 531.deepsjeng, 526.blender, 435.gromacs, 523.xalancbmk, 447.dealIII, 508.namd, 538.imagick, 445.gobmk, 444.namd, 464.h264ref, ycsb_abgsave, 458.sjeng, 541.leela, tpch6, 511.povray, 456.hmmer, 481.wrf, grep_map0, 500.perlbench, 403.gcc, 401.bzip2

9 Evaluation

We 1) analyze ABACuS’s system performance and DRAM energy overheads and compare ABACuS’s system performance and DRAM energy overheads to state-of-the-art mitigation mechanisms (§9.1), 2) show the effect of the number of banks in the system on ABACuS’s performance, and 3) analyze ABACuS’s performance under RowHammer attacks.

9.1 System Performance and DRAM Energy

System Performance Overhead. Fig. 7 presents the performance (in instructions per cycle) of all single-core workloads (grouped into three categories and sorted based on RBMPKI; see Table 4) for four different near-future and very low RowHammer thresholds when executed on a system that uses ABACuS, normalized to a baseline system that does not employ any RowHammer mitigation mechanism.




Figure 7: Performance of single-core applications for four different RowHammer thresholds (higher is better)

We make two major observations from Fig. 7. First, ABACuS induces minor system performance overhead for all evaluated single-core workloads at a near-future N_{RH} of 1000. At such N_{RH} , ABACuS incurs *only* 0.58% (32.00%) slowdown on average (at maximum) across all workloads. ABACuS increases the average memory latency experienced by application memory requests by 1.87% on average across all workloads (not shown) due to preventive refreshes. Second, ABACuS can efficiently prevent RowHammer bitflips even at very low N_{RH} . At such a future N_{RH} of 125, ABACuS induces 1.45% (12.37%) slowdown on average (at maximum) across all workloads. At this N_{RH} , the average memory latency increases by 2.72% across all workloads on average due to preventive refresh operations. We attribute the increasing trend in ABACuS’s average slowdown to the larger number of preventive refresh operations performed by ABACuS at lower N_{RH} values. More workloads hammer more rows more times in a refresh window as N_{RH} reduces, which leads to both i) ABACuS row activation counters (RACs) incrementing faster and ii) more ABACuS RACs reaching the preventive refresh threshold earlier (and ABACuS performing costly preventive refresh operations). In contrast to the trend in the average slowdown, we observe that ABACuS induces a smaller maximum slowdown as N_{RH} reduces. This is because ABACuS implements more row activation counters at lower N_{RH} values and a very memory-intensive random access workload (e.g., gups) *cannot* quickly exhaust all ABACuS RACs and increases the spillover counter value to the refresh cycle threshold slower (than at higher N_{RH} values) such that ABACuS less frequently performs refresh cycles (§4.2).

DRAM Energy Overhead. Fig. 8 presents the DRAM energy consumption for all single-core workloads for four different RowHammer thresholds when executed on a system that uses ABACuS, normalized to a baseline system that does not employ any RowHammer mitigation mechanism.




Figure 8: DRAM energy for single-core applications for four different RowHammer thresholds (lower is better)

We make two key observations from Fig. 8. First, ABACuS induces minor DRAM energy overhead at $N_{RH} = 1000$. ABACuS increases DRAM energy consumption by *only* 1.66% (2.02 \times) on average (at maximum) across all evaluated workloads. Second, ABACuS increases DRAM energy consumption by 1.27% (30.46%) on average (maximum) across all workloads at $N_{RH}=125$. We attribute the DRAM energy overheads to i) increased DRAM activation, precharge, and command bus energy induced by the preventive refresh operations, and ii) increased DRAM background (standby) energy consumption due to increased execution time for applications.

Performance Comparison. Fig. 9 presents the performance impact of ABACuS and four state-of-the-art mechanisms on a single-core system for four different RowHammer thresholds, normalized to the baseline system.




Figure 9: Performance comparison of ABACuS vs. state-of-the-art mitigation techniques for single-core workloads at four different RowHammer thresholds

We make six key observations based on Fig. 9. First, ABACuS outperforms Hydra, REGA, and PARA at RowHammer thresholds below 1000 and performs similarly to Graphene at all tested RowHammer thresholds on average across all workloads. Second, ABACuS outperforms Hydra and PARA at $N_{RH} = 1000$. Third, REGA [177] at $N_{RH} = 1000$ does *not* incur any performance overhead. At $N_{RH} = 1000$, REGA can hide the latency of a preventive refresh behind the latency of a DRAM row access (i.e., preventive refresh happens concurrently with a DRAM row access at the nominal t_{RC} defined in the DDR4 standard [158]). However, REGA incurs increasingly higher overheads as N_{RH} reduces because REGA needs to perform multiple preventive refreshes on each DRAM row access. To perform 8 preventive refreshes on each DRAM row access at an N_{RH} of 125, REGA increases t_{RC} from the nominal value of 45.0 ns [158] to 167.5 ns, where REGA induces 16.65% performance overhead on average across all workloads as the average memory access latency increases.

Fourth, PARA [1] performs the worst among all evaluated mechanisms. PARA incurs 5.47% and 31.08% performance overheads on average across all workloads at $N_{RH} = 1000$ and 125, respectively, because it performs many unnecessary refresh operations [14, 104].

Fifth, Hydra [106] incurs 1.80%, 3.33%, 5.70%, and 9.75% higher performance overheads than ABACuS for N_{RH} of 1000, 500, 250, and 125, respectively, on average across all workloads. In addition to performing preventive refresh operations, Hydra also performs i) an *ACT* and a write (*WR*) command

when a counter in its row count cache (RCC) needs to be evicted to the row count table (RCT) in DRAM, and ii) an *ACT* and a read (*RD*) command when a counter needs to be retrieved from the RCT and placed in the RCC. These operations incur additional performance overheads due to i) row buffer misses that interfere with application memory requests, and ii) DRAM banks being unavailable during RCC eviction and RCT access operations, on top of the overheads caused by preventive refresh operations. For example, at an N_{RH} of 125, the Hydra-based system has i) a row buffer miss rate 6.22% larger than that of ABACuS and ii) an average memory latency experienced by application memory requests 20.94% higher than that of ABACuS on average across all workloads.

Sixth, Graphene [102] incurs slightly higher performance overhead than ABACuS on average across all workloads at an N_{RH} of 125. Even though ABACuS, compared to Graphene, performs 2.06 \times more preventive refresh operations as ABACuS’s shared activation counters reach the preventive refresh threshold faster, the amount of time where at least one DRAM bank is unavailable (for serving application memory requests) because of preventive refresh is an estimated 7.73 \times higher in Graphene compared to ABACuS. Once an ABACuS activation counter reaches N_{RH} , ABACuS performs 64 preventive refresh operations (to all 64 victim rows in 32 banks of the rank) in *quick succession*. The memory controller takes approximately 170 ns⁷ to issue all activate and precharge commands that make up a preventive refresh operation, leveraging bank-level parallelism. In contrast, issuing two preventive refresh operations to a single bank takes approximately 90 ns, an already large fraction of 170 ns. Keeping at least one DRAM bank unavailable for a longer total time increases the critical path for application memory requests in Graphene by a larger amount than in ABACuS. As such, the amount of time in which the processor *cannot* execute instructions due to the re-order buffer being full is higher by 1.87% in Graphene compared to ABACuS on average across all workloads.

Fig. 10 shows ABACuS and four state-of-the-art mechanisms’ performance impact in terms of weighted speedup [191–193] for four different N_{RH} values on an eight-core system, normalized to the baseline system.




Figure 10: Performance comparison for multi-programmed (8 core) workloads at four different RowHammer thresholds

⁷ Calculated as the time it takes to activate and precharge two rows in the same bank ($2 * t_{RC}$) plus the number of banks multiplied by the minimum time between two successive *ACT* commands to different banks in the same rank ($32 * t_{RRD}$).

We make three key observations from Fig. 10. First, ABACuS induces small system performance overhead across all evaluated workloads and RowHammer thresholds. ABACuS has 0.77%, 1.19%, 2.29%, and 4.48% performance overhead on average across all workloads for N_{RH} of 1000, 500, 250, and 125, respectively. Second, Hydra incurs 2.56% higher performance overhead than ABACuS at an N_{RH} of 1K. Third, ABACuS outperforms Hydra, REGA, and PARA at an extreme RowHammer threshold of 125. At such N_{RH} , ABACuS incurs *only* 4.48% performance overhead, whereas Hydra, REGA, and PARA incur 16.49%, 29.31%, and 40.16% performance overhead on average across all workloads.

Energy Comparison. Fig. 11 presents the DRAM energy consumption of ABACuS and four state-of-the-art mechanisms on a single-core system for four different RowHammer thresholds, normalized to the baseline system.




Figure 11: DRAM energy comparison for single-core workloads at four different RowHammer thresholds

From Fig. 11, we make two observations. First, ABACuS induces smaller DRAM energy overhead than other evaluated mitigation mechanisms (except Graphene) on average across all workloads for $N_{RH} < 1000$. Second, at an $N_{RH} = 1000$, ABACuS induces 1.34% and 1.36% smaller DRAM energy overheads than REGA and PARA, respectively, because REGA preventively refreshes one row with every DRAM row activation (at $N_{RH} = 1000$) and PARA performs many unnecessary refresh operations. ABACuS induces 1.66% average ($2.02 \times$ maximum) DRAM energy overhead at this N_{RH} , which is close to Hydra’s 0.73% average ($1.11 \times$ maximum) DRAM energy overhead.

Fig. 12 shows the DRAM energy consumption of ABACuS and four state-of-the-art mechanisms for four different N_{RH} on an eight-core system, normalized to the baseline system.




Figure 12: DRAM energy comparison for multi-programmed (8 core) workloads at four different RowHammer thresholds

From Fig. 12, we observe that ABACuS induces 2.12%, 2.44%, 3.25%, and 4.76% DRAM energy overhead at $N_{RH} =$

1000, 500, 250, and 125, respectively. At a very low $N_{RH} = 125$, ABACuS’s DRAM energy overhead is 19.64%, 70.41%, and 33.99% smaller than Hydra, REGA, and PARA, on average across all evaluated workloads. Graphene induces 3.95% lower DRAM energy overhead than ABACuS at $N_{RH} = 125$.

Summary. We conclude that ABACuS induces small system performance and DRAM energy overheads on average across all tested single-core and multi-core workloads for $N_{RH} = 1000, 500, 250,$ and 125 . ABACuS’s performance and DRAM energy overheads are closer to the most-performance-efficient state-of-the-art mechanism ([102]). ABACuS outperforms and consumes less DRAM energy than the most-area-efficient state-of-the-art (counter-based) mechanism ([106]).

9.2 Sensitivity to Number of Banks

We run 16-, 32-, and 64-bank (1-, 2-, and 4-rank) simulations using ABACuS and the baseline system. We observe that ABACuS can prevent RowHammer bitflips with low overhead in systems that use memory modules with different numbers of banks (ranks). At $N_{RH} = 125$, ABACuS incurs 1.58%, 1.50%, and 2.60% performance overheads for 16-, 32-, and 64-bank configurations, respectively, on average (geometric mean) across all evaluated single-core workloads.

9.3 Adversarial Workloads

ABACuS securely prevents bitflips under RowHammer attacks (§5). We demonstrate that, in a dual-core system, ABACuS incurs smaller performance overheads than Hydra, REGA, and PARA for the evaluated single-core workloads on average, while one core in the system executes a traditional RowHammer access pattern (RowHammer Attack) that repeatedly activates 32 rows in each bank in a bank-interleaved manner. We also develop two specialized RowHammer access patterns (which are open source [190]): Hydra-Adversarial and ABACuS-Adversarial. 1) Hydra-Adversarial exacerbates Hydra’s Row Count Cache eviction rate to significantly increase the throughput of Hydra’s DRAM read and write requests. 2) ABACuS-Adversarial rapidly increments the spillover counter value to cause frequent refresh cycles (§4.2). All three access patterns (RowHammer Attack, Hydra-Adversarial, and ABACuS-Adversarial) incur the same, substantially high rate of *ACT* commands in the memory controller. The memory controller issues an *ACT* command *every* 20 ns while executing each access pattern. Fig. 13 shows the performance impact of ABACuS and the state-of-the-art mechanisms on all evaluated single-core workloads in a dual-core system when the second core executes one of the three RowHammer access patterns.

We make three major observations. First, ABACuS induces *only* 0.88% performance overhead on average across all workloads when one core executes the RowHammer attack, whereas Graphene, Hydra, REGA, and PARA induce 0.61%, 3.03%, 4.43%, and 14.62%, respectively. Second, Hydra induces a large 73.96% average slowdown when one core




Figure 13: Performance comparison for single-core workloads with three different RowHammer access patterns ($N_{RH} = 500$)

executes the Hydra-Adversarial access pattern. We attribute this overhead to the high rate of Row Count Cache (RCC) evictions the Hydra-Adversarial access pattern incurs. Hydra evicts 1.13 RCC entries per last level cache miss on average across all workloads. The memory controller serves an RCC eviction by issuing high-priority *WR* and *RD* DRAM requests (i.e., *WR* and *RD* requests caused by RCC evictions are on the critical path of workload main memory requests). For the same access pattern, ABACuS incurs *only* 0.79% on average across all workloads. Third, ABACuS induces 9.20% performance overhead, on average across all workloads when one core executes the ABACuS-Adversarial access pattern. This is because the ABACuS-Adversarial access pattern triggers multiple ABACuS refresh cycles, during which no memory request can be serviced, while the single-core workload executes. The same access pattern incurs 48.08% performance overhead on average across all workloads for Hydra.

We conclude that ABACuS incurs almost-negligible additional performance overhead for benign workloads when another core executes a traditional RowHammer attack. Specialized adversarial access patterns can exacerbate such overheads by frequently triggering ABACuS refresh cycles.

9.3.1 Improving ABACuS’s Performance

A workload may, intentionally (e.g., ABACuS-Adversarial) or unintentionally (e.g., gups), rapidly increment the spillover counter’s value, frequently triggering *refresh cycles* where ABACuS issues a refresh command to each DRAM row ID in a rank, and cause substantial performance overheads in an ABACuS-based system. To prevent such overheads, a less-area-constrained version of ABACuS can remove the spillover counter and implement one shared activation counter (ABACuS counter) per DRAM row ID (i.e., ABACuS-Big’s $N_{entries}$ (Table 2) is equal to the number of rows in a DRAM bank). We design and evaluate ABACuS-Big, which implements one ABACuS counter per DRAM row ID. The ABACuS counter in ABACuS-Big is updated in the same way as in ABACuS. ABACuS-Big implements as many ABACuS counters as there are rows in a bank (i.e., there is a 1-1 mapping between ABACuS counters and DRAM row IDs) to keep *precise* track of *every sibling row’s* maximum activation count and ABACuS-Big does *not* need to use a spillover counter. Fig. 14 shows the performance impact of ABACuS and ABACuS-Big on all evaluated single-core workloads in a

dual-core system with the three RowHammer access patterns described earlier in this section.




Figure 14: Performance of ABACuS&-Big ($N_{RH} = 500$)

We observe that ABACuS-Big incurs *only* 0.28% performance overhead on average across all workloads for the ABACuS-Adversarial pattern, whereas ABACuS incurs 9.20% because ABACuS-Big does *not* perform any refresh cycles where the memory controller is busy rapidly issuing *REF* commands. We evaluate ABACuS-Big’s chip area using the methodology described in §7.1. ABACuS-Big requires 40 bits (8 bits for row activation counter, 32 bits for sibling activation vector) of storage per DRAM row, amounting to 640 KiB on-chip storage at an N_{RH} of 500 for 128K DRAM rows. ABACuS-Big takes up 0.48 mm² chip area (0.20% of a high-end Intel Xeon processor’s area [178]). We conclude that ABACuS-Big is a high-performance implementation of the ABACuS design, which system designers that have fewer chip area constraints can choose to implement, that improves system performance under adversarial workloads and some benign workloads (e.g., gups) compared to ABACuS.

10 Related Work

To our knowledge, ABACuS is the first work that mitigates RowHammer efficiently and scalably at very low RowHammer thresholds (e.g., 125) without incurring large area, performance, or energy overheads. Sections 7 and 9 already qualitatively and quantitatively compare ABACuS to the most relevant state-of-the-art mechanisms [1, 102, 106, 177]. This section discusses other RowHammer mitigation mechanisms. **Hardware-based Mitigation Mechanisms.** Many prior works [1, 87, 88, 97, 98, 100–102, 104–107, 110–112, 114, 116, 117, 122, 123, 126, 131, 133, 134, 173, 194–196] propose hardware-based mitigation mechanisms to prevent RowHammer bitflips. We classify these into three main categories. 1) Probabilistic preventive refresh (PPR) mechanisms [1, 97, 100, 114, 116, 122, 133, 194] preventively refresh victim rows based on a probability. PPR mechanisms incur impractical performance overheads at very low RowHammer thresholds as they perform many unnecessary preventive refresh operations. A recent work [195] proposes a new methodology for configuring PPRs. 2) Deterministic preventive refresh (DPR) mechanisms [88, 98, 101, 102, 105, 107, 110–112, 117, 123, 126, 131, 134, 173, 196] track activation counts of aggressor rows and preventively refresh victim rows. DPR mechanisms incur less performance overhead than PPR mechanisms (from fewer unnecessary preventive refresh operations) at the cost of larger chip area overhead to store aggressor row activation counters. and 3) Deterministic aggressor row access throttling (DAT) mechanisms [1, 87, 104] track activation counts of aggressor rows and preventively

block memory accesses to aggressor rows. *DAT mechanisms* incur average system performance and total chip area overheads similar to DPR mechanisms [104]. However, existing *DAT mechanisms* can incur delays in the order of microseconds on memory demand requests (e.g., load instructions) [104, 106, 117].

Software-based Mitigation Mechanisms. Many works [38, 44, 93, 108, 113, 118, 132] propose software-based mitigation mechanisms to avoid hardware modifications. Unfortunately, it is *not* possible for these mechanisms to monitor *all* memory requests, and thus most of these mechanisms have already been defeated by recent attacks [30, 37, 42, 46, 49, 52, 197].

Integrity-based Mitigation Mechanisms. Several works [115, 119, 127, 198] propose integrity check mechanisms to detect and correct bitflips that may have been induced by RowHammer. Unfortunately, it is either not possible, difficult, or prohibitively expensive to correct all possible RowHammer bitflips using these mechanisms. However, these mechanisms can be combined with ABACuS to improve overall system reliability and future work could demonstrate the benefits of combining them with ABACuS.

RowHammer Mitigation in Commodity Chips. DRAM manufacturers employ RowHammer mitigation mechanisms, commonly referred to as target row refresh (TRR), in commodity DRAM chips [158, 163] without publicly documenting their detailed designs. These mechanisms typically do *not* induce any performance overhead because they take action (e.g., refresh a victim row) when the DRAM chip is busy performing a periodic refresh operation (i.e., their victim row refresh latency is hidden by the latency of performing a periodic refresh operation). However, recent studies experimentally demonstrate that specialized adversarial access patterns can defeat some of these mechanisms [15, 29, 54–56, 128]. A recent work [121] develops a tool that can automatically infer parameters of TRR mechanisms. More recent works from industry design new in-DRAM RowHammer mitigation mechanisms [135, 199]. Unfortunately, these mechanisms cannot or are not proven to deterministically prevent all RowHammer bitflips.

Device-level Mechanisms for Mitigating RowHammer. Several prior works [8, 9, 130, 200] design new DRAM cells or arrays with improved RowHammer resilience. Unfortunately, these works alone *cannot* completely prevent RowHammer bitflips (but could effectively increase the RowHammer thresholds). They can be used with other hardware-/software-based mitigation techniques to mitigate RowHammer.

11 Conclusion

We introduced a new RowHammer mitigation mechanism that prevents RowHammer bitflips at low area, performance, and energy overheads for modern and future DRAM chips that are very vulnerable to RowHammer (e.g., with RowHammer thresholds as low as 125). Compared to existing RowHammer mitigation mechanisms, our technique, all-bank activation counters for scalable RowHammer mitigation (ABACuS)

technique incurs significantly smaller area, performance, and DRAM energy overheads for modern and future DRAM chips. Our technique achieves this by sharing activation counters of rows that has the same row ID in different banks. While ABACuS efficiently and securely prevents RowHammer bitflips, it also scales well with worsening RowHammer vulnerability down to RowHammer threshold (N_{RH}) = 125.

Acknowledgments

We thank our shepherd and the anonymous reviewers of USENIX Security 2023. We thank the SAFARI Research Group members for feedback and the stimulating intellectual environment. We acknowledge the generous gift funding provided by our industrial partners (especially Google, Huawei, Intel, Microsoft, VMware), which has been instrumental in enabling our decade-long research on read disturbance in DRAM and memory systems. This work was in part supported by a Google Security and Privacy Research Award and the Microsoft Swiss Joint Research Center.

References

- [1] Y. Kim et al. Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors. In *ISCA*, 2014.
- [2] Onur Mutlu. The RowHammer Problem and Other Issues We May Face as Memory Becomes Denser. In *DATE*, 2017.
- [3] Thomas Yang et al. Trap-Assisted DRAM Row Hammer Effect. *EDL*, 2019.
- [4] Onur Mutlu et al. RowHammer: A Retrospective. *TCAD*, 2019.
- [5] Kyungbae Park et al. Statistical Distributions of Row-Hammering Induced Failures in DDR3 Components. *Microelectronics Reliability*, 2016.
- [6] Kyungbae Park et al. Experiments and Root Cause Analysis for Active-Precharge Hammering Fault in DDR3 SDRAM under 3nm Technology. *Microelectronics Reliability*, 2016.
- [7] Andrew J. Walker et al. On DRAM RowHammer and the Physics on Insecurity. *IEEE TED*, 2021.
- [8] Seong-Wan Ryu et al. Overcoming the Reliability Limitation in the Ultimately Scaled DRAM using Silicon Migration Technique by Hydrogen Annealing. In *IEDM*, 2017.
- [9] Chia Yang et al. Suppression of RowHammer Effect by Doping Profile Modification in Saddle-Fin Array Devices for Sub-30-nm DRAM Technology. *TDMR*, 2016.
- [10] Chia-Ming Yang et al. Scanning Spreading Resistance Microscopy for Doping Profile in Saddle-Fin Devices. *IEEE Transactions on Nanotechnology*, 2017.
- [11] SK Gautam et al. Row Hammering Mitigation Using Metal Nanowire in Saddle Fin DRAM. *IEEE TED*, 2019.
- [12] Yichen Jiang et al. Quantifying RowHammer Vulnerability for DRAM Security. In *DAC*, 2021.
- [13] Onur Mutlu et al. Fundamentally Understanding and Solving RowHammer. In *ASP-DAC*, 2023.
- [14] Jeremie S. Kim et al. Revisiting RowHammer: An Experimental Analysis of Modern Devices and Mitigation Techniques. In *ISCA*, 2020.
- [15] Pietro Frigo et al. TRRespass: Exploiting the Many Sides of Target Row Refresh. In *S&P*, 2020.
- [16] A. Giray Yağlıkçı et al. Understanding RowHammer Under Reduced Wordline Voltage: An Experimental Study Using Real DRAM Devices. In *DSN*, 2022.

- [17] Lois Orosa et al. A Deeper Look into RowHammer’s Sensitivities: Experimental Analysis of Real DRAM Chips and Implications on Future Attacks and Defenses. In *MICRO*, 2021.
- [18] Onur Mutlu. RowHammer. Top Picks in Hardware and Embedded Security, 2018.
- [19] Onur Mutlu et al. Fundamentally Understanding and Solving RowHammer. *arXiv:2211.07613 [cs.CR]*, 2022.
- [20] Apostolos P Fournaris et al. Exploiting Hardware Vulnerabilities to Attack Embedded System Devices: A Survey of Potent Microarchitectural Attacks. *Electronics*, 2017.
- [21] Damian Poddebniak et al. Attacking Deterministic Signature Schemes using Fault Attacks. In *EuroS&P*, 2018.
- [22] Andrei Tatar et al. Throwhammer: Rowhammer Attacks Over the Network and Defenses. In *USENIX ATC*, 2018.
- [23] Sebastien Carre et al. OpenSSL Bellcore’s Protection Helps Fault Attack. In *DSD*, 2018.
- [24] Alessandro Barengi et al. Software-Only Reverse Engineering of Physical DRAM Mappings for Rowhammer Attacks. In *IVSW*, 2018.
- [25] Zhenkai Zhang et al. Triggering Rowhammer Hardware Faults on ARM: A Revisit. In *ASHES*, 2018.
- [26] Sarani Bhattacharya et al. Advanced Fault Attacks in Software: Exploiting the Rowhammer Bug. *Fault Tolerant Architectures for Cryptography and Hardware Security*, 2018.
- [27] SAFARI Research Group. RowHammer — GitHub Repository. <https://github.com/CMU-SAFARI/rowhammer>.
- [28] Mark Seaborn et al. Exploiting the DRAM Rowhammer Bug to Gain Kernel Privileges. <http://googleprojectzero.blogspot.com.tr/2015/03/exploiting-dram-rowhammer-bug-to-gain.html>, 2015.
- [29] Victor van der Veen et al. Drammer: Deterministic Rowhammer Attacks on Mobile Platforms. In *CCS*, 2016.
- [30] Daniel Gruss et al. Rowhammer.js: A Remote Software-Induced Fault Attack in Javascript. *arXiv:1507.06955 [cs.CR]*, 2016.
- [31] Kaveh Razavi et al. Flip Feng Shui: Hammering a Needle in the Software Stack. In *USENIX Security*, 2016.
- [32] Peter Pessl et al. DRAMA: Exploiting DRAM Addressing for Cross-CPU Attacks. In *USENIX Security*, 2016.
- [33] Yuan Xiao et al. One Bit Flips, One Cloud Flops: Cross-VM Row Hammer Attacks and Privilege Escalation. In *USENIX Security*, 2016.
- [34] Erik Bosman et al. Dedup Est Machina: Memory Deduplication as An Advanced Exploitation Vector. In *S&P*, 2016.
- [35] Sarani Bhattacharya et al. Curious Case of Rowhammer: Flipping Secret Exponent Bits Using Timing Analysis. In *CHES*, 2016.
- [36] Wayne Burleson et al. Invited: Who is the Major Threat to Tomorrow’s Security? You, the Hardware Designer. In *DAC*, 2016.
- [37] Rui Qiao et al. A New Approach for RowHammer Attacks. In *HOST*, 2016.
- [38] Ferdinand Brasser et al. Can’t Touch This: Software-Only Mitigation Against Rowhammer Attacks Targeting Kernel Memory. In *USENIX Security*, 2017.
- [39] Yeongjin Jang et al. SGX-Bomb: Locking Down the Processor via Rowhammer Attack. In *SOSP*, 2017.
- [40] Misiker Aga et al. When Good Protections Go Bad: Exploiting Anti-DoS Measures to Accelerate Rowhammer Attacks. In *HOST*, 2017.
- [41] Andrei Tatar et al. Defeating Software Mitigations Against Rowhammer: A Surgical Precision Hammer. In *RAID*, 2018.
- [42] Daniel Gruss et al. Another Flip in the Wall of Rowhammer Defenses. In *S&P*, 2018.
- [43] Moritz Lipp et al. Nethammer: Inducing Rowhammer Faults Through Network Requests. *arXiv:1805.04956 [cs.CR]*, 2018.
- [44] Victor van der Veen et al. GuardION: Practical Mitigation of DMA-Based Rowhammer Attacks on ARM. In *DIMVA*, 2018.
- [45] Pietro Frigo et al. Grand Pwning Unit: Accelerating Microarchitectural Attacks with the GPU. In *S&P*, 2018.
- [46] Lucian Cojocar et al. Exploiting Correcting Codes: On the Effectiveness of ECC Memory Against Rowhammer Attacks. In *S&P*, 2019.
- [47] Sangwoo Ji et al. Pinpoint Rowhammer: Suppressing Unwanted Bit Flips on Rowhammer Attacks. In *ASIACCS*, 2019.
- [48] Sanghyun Hong et al. Terminal Brain Damage: Exposing the Graceless Degradation in Deep Neural Networks Under Hardware Fault Attacks. In *USENIX Security*, 2019.
- [49] Andrew Kwong et al. RAMBleed: Reading Bits in Memory Without Accessing Them. In *S&P*, 2020.
- [50] Lucian Cojocar et al. Are We Susceptible to Rowhammer? An End-to-End Methodology for Cloud Providers. In *S&P*, 2020.
- [51] Zane Weissman et al. JackHammer: Efficient Rowhammer on Heterogeneous FPGA–CPU Platforms. *arXiv:1912.11523 [cs.CR]*, 2020.
- [52] Zhi Zhang et al. PTHammer: Cross-User-Kernel-Boundary Rowhammer Through Implicit Accesses. In *MICRO*, 2020.
- [53] Fan Yao et al. Deephammer: Depleting the Intelligence of Deep Neural Networks Through Targeted Chain of Bit Flips. In *USENIX Security*, 2020.
- [54] Finn de Ridder et al. SMASH: Synchronized Many-Sided Rowhammer Attacks from JavaScript. In *USENIX Security*, 2021.
- [55] Hasan Hassan et al. Uncovering In-DRAM RowHammer Protection Mechanisms: A New Methodology, Custom RowHammer Patterns, and Implications. In *MICRO*, 2021.
- [56] Patrick Jattke et al. Blacksmith: Scalable Rowhammering in the Frequency Domain. In *SP*, 2022.
- [57] M Caner Tol et al. Toward Realistic Backdoor Injection Attacks on DNNs using RowHammer. *arXiv:2110.07683v2 [cs.LG]*, 2022.
- [58] Andreas Kogler et al. Half-Double: Hammering From the Next Row Over. In *USENIX Security*, 2022.
- [59] Lois Orosa et al. SpyHammer: Using RowHammer to Remotely Spy on Temperature. *arXiv:2210.04084 [cs.CR]*, 2022.
- [60] Zhi Zhang et al. Implicit Hammer: Cross-Privilege-Boundary Rowhammer through Implicit Accesses. *IEEE Transactions on Dependable and Secure Computing*, 2022.
- [61] Liang Liu, Yanan Guo, Yueqiang Cheng, Youtao Zhang, and Jun Yang. Generating Robust DNN with Resistance to Bit-Flip based Adversarial Weight Attack. *IEEE Transactions on Computers*, 2022.
- [62] Yaakov Cohen et al. HammerScope: Observing DRAM Power Consumption Using Rowhammer. In *CCS*, 2022.
- [63] Mengxin Zheng et al. TrojViT: Trojan Insertion in Vision Transformers. *arXiv:2208.13049 [cs.LG]*, 2022.
- [64] Michael Fahr Jr et al. When Frodo Flips: End-to-End Key Recovery on FrodoKEM via Rowhammer. *CCS*, 2022.
- [65] Youssef Tobah et al. SpecHammer: Combining Spectre and Rowhammer for New Speculative Attacks. In *SP*, 2022.
- [66] Adnan Siraj Rakin et al. DeepSteal: Advanced Model Extractions Leveraging Efficient Weight Stealing in Memories. In *SP*, 2022.
- [67] Hakan Aydin et al. Cyber Security in Industrial Control Systems (ICS): A Survey of RowHammer Vulnerability. *Applied Computer Science*, 2022.
- [68] Koksal Mus et al. Jolt: Recovering TLS Signing Keys via Rowhammer Faults. *Cryptology ePrint Archive*, 2022.

- [69] Jianxin Wang et al. Research and Implementation of Rowhammer Attack Method based on Domestic NeoKylin Operating System. In *ICFTIC*, 2022.
- [70] Sam Lefforge. Reverse Engineering Post-Quantum Cryptography Schemes to Find Rowhammer Exploits. Bachelor's Thesis, University of Arkansas, 2023.
- [71] Michael Jacob Fahr. The Effects of Side-Channel Attacks on Post-Quantum Cryptography: Influencing FrodoKEM Key Generation Using the Rowhammer Exploit. Master's thesis, University of Arkansas, 2022.
- [72] Anandpreet Kaur et al. Work-in-Progress: DRAM-MaUT: DRAM Address Mapping Unveiling Tool for ARM Devices. In *CASES*, 2022.
- [73] Kunbei Cai et al. On the Feasibility of Training-time Trojan Attacks through Hardware-based Faults in Memory. In *HOST*, 2022.
- [74] Dawei Li et al. CyberRadar: A PUF-based Detecting and Mapping Framework for Physical Devices. arXiv:2201.07597, 2022.
- [75] Arman Roohi et al. Efficient Targeted Bit-Flip Attack Against the Local Binary Pattern Network. In *HOST*, 2022.
- [76] Felix Staudigl et al. NeuroHammer: Inducing Bit-Flips in Memristive Crossbar Memories. In *DATE*, 2022.
- [77] Li-Hsing Yang et al. Socially-Aware Collaborative Defense System against Bit-Flip Attack in Social Internet of Things and Its Online Assignment Optimization. In *ICCCN*, 2022.
- [78] Saad Islam et al. Signature Correction Attack on Dilithium Signature Scheme. In *Euro S&P*, 2022.
- [79] Anandpreet Kaur et al. Flipping Bits Like a Pro: Precise Rowhammering on Embedded Devices. *IEEE Embedded Systems Letters*, 2023.
- [80] Andrew J. Adiletta et al. Mayhem: Targeted Corruption of Register and Stack Variables. arXiv:2309.02545, 2023.
- [81] Jianshuo Dong et al. One-bit Flip is All You Need: When Bit-flip Attack Meets Model Training. In *ICCV*, 2023.
- [82] Fangxin Liu et al. HyperAttack: An Efficient Attack Framework for HyperDimensional Computing. In *DAC*, 2023.
- [83] M. Caner Tol et al. Don't Knock! Rowhammer at the Backdoor of DNN Models. In *DSN*, 2023.
- [84] Apple Inc. About the Security Content of Mac EFI Security Update 2015-001. <https://support.apple.com/en-us/HT204934>. June 2015.
- [85] Hewlett-Packard Enterprise. HP Moonshot Component Pack Version 2015.05.0, 2015.
- [86] Lenovo Group Ltd. Row Hammer Privilege Escalation, 2015.
- [87] Zvika Greenfield et al. Throttling Support for Row-Hammer Counters, 2016. U.S. Patent 9,251,885.
- [88] Dae-Hyun Kim et al. Architectural Support for Mitigating Row Hammering in DRAM Memories. *CAL*, 2015.
- [89] K.S. Bains and J.B. Halbert. Distributed Row Hammer Tracking. US Patent App. 13/631,781, April 3 2014.
- [90] K.S. Bains et al. Method, Apparatus and System for Providing a Memory Refresh. US Patent App. 13/625,741, March 27 2014.
- [91] K.S. Bains et al. Row Hammer Refresh Command. US Patent App. 13/539,415, January 2 2014.
- [92] K. Bains et al. Row Hammer Refresh Command. US Patent App. 14/068,677, February 27 2014.
- [93] Zelalem Birhanu Aweke et al. ANVIL: Software-Based Protection Against Next-Generation Rowhammer Attacks. In *ASPLOS*, 2016.
- [94] Kuljit Bains et al. Row Hammer Refresh Command, 2015. U.S. Patent 9,117,544.
- [95] Kuljit S Bains et al. Row Hammer Monitoring Based on Stored Row Hammer Threshold Value, 2016. U.S. Patent 9,384,821.
- [96] Kuljit S Bains et al. Distributed Row Hammer Tracking, 2016. U.S. Patent 9,299,400.
- [97] Mungyu Son et al. Making DRAM Stronger Against Row Hammering. In *DAC*, 2017.
- [98] S. M. Seyedzadeh et al. Mitigating Wordline Crosstalk Using Adaptive Trees of Counters. In *ISCA*, 2018.
- [99] Gorka Irazoqui et al. MASCAT: Stopping Microarchitectural Attacks Before Execution. *IACR Cryptology*, 2016.
- [100] Jung Min You et al. MRLoc: Mitigating Row-Hammering Based on Memory Locality. In *DAC*, 2019.
- [101] Eojin Lee et al. TWiCe: Preventing Row-Hammering by Exploiting Time Window Counters. In *ISCA*, 2019.
- [102] Yeonhong Park et al. Graphene: Strong yet Lightweight Row Hammer Protection. In *MICRO*, 2020.
- [103] A. Giray Yağlıkçı et al. Security Analysis of the Silver Bullet Technique for RowHammer Prevention. arXiv:2106.07084 [cs.CR], 2021.
- [104] A. Giray Yağlıkçı et al. BlockHammer: Preventing RowHammer at Low Cost by Blacklisting Rapidly-Accessed DRAM Rows. In *HPCA*, 2021.
- [105] Ingab Kang et al. CAT-TWO: Counter-Based Adaptive Tree, Time Window Optimized for DRAM Row-Hammer Prevention. *IEEE Access*, 2020.
- [106] Moinuddin Qureshi et al. Hydra: Enabling Low-Overhead Mitigation of Row-Hammer at Ultra-Low Thresholds via Hybrid Tracking. In *ISCA*, 2022.
- [107] Gururaj Saileshwar et al. Randomized Row-Swap: Mitigating Row Hammer by Breaking Spatial Correlation Between Aggressor and Victim Rows. In *ASPLOS*, 2022.
- [108] Radhesh Krishnan Konoth et al. ZebRAM: Comprehensive and Compatible Software Protection Against Rowhammer Attacks. In *OSDI*, 2018.
- [109] Saru Vig et al. Rapid Detection of Rowhammer Attacks Using Dynamic Skewed Hash Tree. In *HASP*, 2018.
- [110] Michael Jaemin Kim et al. Mithril: Cooperative Row Hammer Protection on Commodity DRAM Leveraging Managed Refresh. In *HPCA*, 2022.
- [111] Gyu-Hyeon Lee et al. CryoGuard: A Near Refresh-Free Robust DRAM Design for Cryogenic Computing. In *ISCA*, 2021.
- [112] Michele Marazzi et al. ProTRR: Principled yet Optimal In-DRAM Target Row Refresh. In *IEEE S&P*, 2022.
- [113] Zhi Zhang et al. SoftTRR: Protect Page Tables against Rowhammer Attacks using Software-Only Target Row Refresh. In *USENIX ATC*, 2022.
- [114] Biresh Kumar Joardar et al. Learning to Mitigate RowHammer Attacks. In *DATE*, 2022.
- [115] Jonas Juffinger et al. CSI: Rowhammer-Cryptographic Security and Integrity against Rowhammer. In *SP*, 2023.
- [116] A. Giray Yağlıkçı et al. HiRA: Hidden Row Activation for Reducing Refresh Latency of Off-the-Shelf DRAM Chips. In *MICRO*, 2022.
- [117] Anish Saxena et al. AQUA: Scalable Rowhammer Mitigation by Quarantining Aggressor Rows at Runtime. In *MICRO*, 2022.
- [118] Shuhei Enomoto et al. Efficient Protection Mechanism for CPU Cache Flush Instruction Based Attacks. *IEICE Transactions on Information and Systems*, 2022.
- [119] Evgeny Manzhosov et al. Revisiting Residue Codes for Modern Memories. In *MICRO*, 2022.

- [120] Samira Ajorpaz et al. EVAX: Towards a Practical, Pro-active & Adaptive Architecture for High Performance & Security. In *MICRO*, 2022.
- [121] Amir Naseredini et al. ALARM: Active LeArning of Rowhammer Mitigations. <https://users.sussex.ac.uk/~mfb21/rh-draft.pdf>, 2022.
- [122] Biresh Kumar Joardar et al. Machine Learning-Based Rowhammer Mitigation. *TCAD*, 2022.
- [123] Hasan Hassan et al. A Case for Self-Managing DRAM Chips: Improving Performance, Efficiency, Reliability, and Security via Autonomous in-DRAM Maintenance Operations. arXiv:2207.13358 [cs.AR], 2022.
- [124] Zhenkai Zhang et al. Leveraging EM Side-Channel Information to Detect Rowhammer Attacks. In *SP*, 2020.
- [125] Kevin Loughlin et al. Stop! Hammer Time: Rethinking Our Approach to Rowhammer Mitigations. In *HotOS*, 2021.
- [126] Fabrice Devaux et al. Method and Circuit for Protecting a DRAM Memory Device from the Row Hammer Effect, 2021. U.S. Patent 10,885,966.
- [127] Ali Fakhrzadehgan et al. SafeGuard: Reducing the Security Risk from Row-Hammer via Low-Cost Integrity Protection. In *HPCA*, 2022.
- [128] Stefan Saroiu et al. The Price of Secrecy: How Hiding Internal DRAM Topologies Hurts Rowhammer Defenses. In *IRPS*, 2022.
- [129] Kevin Loughlin et al. MOESI-Prime: Preventing Coherence-Induced Hammering in Commodity Workloads. In *ISCA*, 2022.
- [130] Jin Han et al. Surround Gate Transistor With Epitaxially Grown Si Pillar and Simulation Study on Soft Error and Rowhammer Tolerance for DRAM. *TED*, 2021.
- [131] Jeonghyun Woo et al. Scalable and Secure Row-Swap: Efficient and Safe Row Hammer Mitigation in Memory Systems. arXiv:2212.12613, 2022.
- [132] Carsten Bock et al. RIP-RH: Preventing Rowhammer-Based Inter-Process Attacks. In *ASIA CCS*, 2019.
- [133] Yicheng Wang et al. Discreet-PARA: Rowhammer Defense with Low Cost and High Efficiency. In *ICCD*, 2021.
- [134] Tanj Bennett et al. Panopticon: A Complete In-DRAM Rowhammer Mitigation. In *Workshop on DRAM Security (DRAMSec)*, 2021.
- [135] Woongrae Kim et al. A 1.1V 16Gb DDR5 DRAM with Probabilistic-Aggressor Tracking, Refresh-Management Functionality, Per-Row Hammer Tracking, a Multi-Step Precharge, and Core-Bias Modulation for Security and Reliability Enhancement. In *ISSCC*, 2023.
- [136] Yoongu Kim et al. A Case for Exploiting Subarray-Level Parallelism (SALP) in DRAM. In *ISCA*, 2012.
- [137] Kevin K Chang et al. Improving DRAM Performance by Parallelizing Refreshes with Accesses. In *HPCA*, 2014.
- [138] Onur Mutlu and Thomas Moscibroda. Parallelism-Aware Batch Scheduling: Enhancing Both Performance and Fairness of Shared DRAM Systems. In *ISCA*, 2008.
- [139] Lavanya Subramanian et al. BLISS: Balancing Performance, Fairness and Complexity in Memory Access Scheduling. *TPDS*, 2016.
- [140] Lavanya Subramanian et al. The Blacklisting Memory Scheduler: Achieving High Performance and Fairness at Low Cost. In *ICCD*, 2014.
- [141] Chang Joo Lee et al. Improving Memory Bank-Level Parallelism in the Presence of Prefetching. In *MICRO*, 2009.
- [142] Eiman Ebrahimi et al. Fairness via Source Throttling: A Configurable and High Performance Fairness Substrate for Multi Core Memory Systems. In *ASPLOS*, 2010.
- [143] Engin Ipek et al. Self-Optimizing Memory Controllers: A Reinforcement Learning Approach. In *ISCA*, 2008.
- [144] Yoongu Kim et al. ATLAS: A Scalable and High-Performance Scheduling Algorithm for Multiple Memory Controllers. In *HPCA*, 2010.
- [145] Yoongu Kim et al. Thread Cluster Memory Scheduling: Exploiting Differences in Memory Access Behavior. In *MICRO*, 2010.
- [146] Eiman Ebrahimi et al. Parallel Application Memory Scheduling. In *MICRO*, 2011.
- [147] Jamie Liu et al. RAIDR: Retention-Aware Intelligent DRAM Refresh. In *ISCA*, 2012.
- [148] Donghyuk Lee et al. Tiered-Latency DRAM: A Low Latency and Low Cost DRAM Architecture. In *HPCA*, 2013.
- [149] Onur Mutlu. Memory Scaling: A Systems Architecture Perspective. In *IMW*, 2013.
- [150] IBM. IBM Power S1014, S1022s, S1022, and S1024 Technical Overview and Introduction. <https://www.redbooks.ibm.com/abstracts/redp5675.html>, 2022.
- [151] AMD Inc. 4TH GEN AMD EPYC PROCESSOR ARCHITECTURE. <https://www.amd.com/system/files/documents/4th-gen-epyc-processor-architecture-white-paper.pdf>, 2023.
- [152] Intel Inc. Intel® Xeon® W-3400 & Intel® Xeon® W-2400 Processors and the Intel® W790 Chipset Workstation Platform Brief, 2023.
- [153] Yoongu Kim et al. Ramulator: A Fast and Extensible DRAM Simulator. *CAL*, 2016.
- [154] SAFARI Research Group. Ramulator — GitHub Repository. <https://github.com/CMU-SAFARI/ramulator>, 2021.
- [155] Haocong Luo et al. Ramulator 2.0: A Modern, Modular, and Extensible DRAM Simulator. arXiv:2308.11030 [cs.AR], 2023.
- [156] SAFARI Research Group. Ramulator 2.0 — GitHub Repository. <https://github.com/CMU-SAFARI/ramulator2>, 2023.
- [157] Rajeev Balasubramonian et al. CACTI 7: New Tools for Interconnect Exploration in Innovative Off-Chip Memories. *ACM TACO*, 2017.
- [158] JEDEC. *JESD79-4C: DDR4 SDRAM Standard*, 2020.
- [159] Ataberk Olgun et al. ABACuS: All-Bank Activation Counters for Scalable and Low Overhead RowHammer Mitigation. arXiv, 2023.
- [160] JEDEC. *JESD79-3: DDR3 SDRAM Standard*, 2012.
- [161] Micron Inc. SDRAM, 4Gb: x4, x8, x16 DDR4 SDRAM Features, 2014.
- [162] JEDEC. *JESD209-4B: Low Power Double Data Rate 4 (LPDDR4) Standard*, 2017.
- [163] JEDEC. *JESD79-5: DDR5 SDRAM Standard*, 2020.
- [164] JEDEC. *JESD209-5A: LPDDR5 SDRAM Standard*, 2020.
- [165] JEDEC. *JESD235C: High Bandwidth Memory (HBM) DRAM*, 2020.
- [166] JEDEC. *JESD79F: Double Data Rate (DDR) SDRAM Standard*, 2008.
- [167] Zhao Zhang et al. A Permutation-Based Page Interleaving Scheme to Reduce Row-Buffer Conflicts and Exploit Data Locality. In *MICRO*, 2000.
- [168] Jae Young Hur et al. Adaptive Linear Address Map for Bank Interleaving in DRAMs. *IEEE Access*, 2019.
- [169] Dimitris Kaseridis et al. Minimalist Open-Page: A DRAM Page-Mode Scheduling Policy for the Many-Core Era. In *MICRO*, 2011.
- [170] Y. Liu et al. Get Out of the Valley: Power-Efficient Address Mapping for GPUs. In *ISCA*, 2018.
- [171] Mohsen Ghasempour et al. DRAM: Dynamic Re-Arrangement of Address Mapping to Improve the Performance of DRAMs. In *MEMSYS*, 2016.

- [172] Jayadev Misra et al. Finding Repeated Elements. *Science of Computer Programming*, 1982.
- [173] Seyed Mohammad Seyedzadeh et al. Counter-Based Tree Structure for Row Hammering Mitigation in DRAM. *CAL*, 2017.
- [174] N.P. Jouppi. Improving Direct-Mapped Cache Performance by the Addition of a Small Fully-Associative Cache and Prefetch Buffers. In *ISCA*, 1990.
- [175] Alan Jay Smith. Cache Memories. *ACM Computing Surveys*, 1982.
- [176] Minesh Patel et al. A Case for Transparent Reliability in DRAM Systems. arXiv:2204.10378 [cs.AR], 2022.
- [177] M. Marazzi et al. REGA: Scalable Rowhammer Mitigation with Refresh-Generating Activations. In *IEEE S&P*, 2023.
- [178] WikiChip. Cascade Lake SP - Intel. https://en.wikichip.org/wiki/intel/cores/cascade_lake_sp.
- [179] Synopsys, Inc. Synopsys Design Compiler. <https://www.synopsys.com/support/training/rtl-synthesis/design-compiler-rtl-synthesis.html>.
- [180] Karthik Chandrasekar et al. DRAMPower: Open-Source DRAM Power & Energy Estimation Tool. <http://www.drampower.info/>.
- [181] Scott Rixner et al. Memory Access Scheduling. In *ISCA*, 2000.
- [182] William K Zuravleff et al. Controller for a Synchronous DRAM That Maximizes Throughput by Allowing Memory Requests and Commands to Be Issued Out of Order, 1997. U.S. Patent 5,630,096.
- [183] Onur Mutlu et al. Stall-Time Fair Memory Access Scheduling for Chip Multiprocessors. In *MICRO*, 2007.
- [184] Standard Performance Evaluation Corp. SPEC CPU 2006. <http://www.spec.org/cpu2006/>, 2006.
- [185] Standard Performance Evaluation Corp. SPEC CPU 2017. <http://www.spec.org/cpu2017/>, 2017.
- [186] Transaction Processing Performance Council. TPC Benchmarks. <http://tpc.org/>.
- [187] Jason E. Fritts et al. MediaBench II Video: Expediting the next Generation of Video Systems Research. *Microprocess. Microsyst.*, 2009.
- [188] Brian Cooper et al. Benchmarking Cloud Serving Systems with YCSB. In *SoCC*, 2010.
- [189] Greg Hamerly et al. SimPoint 3.0: Faster and More Flexible Program Phase Analysis. *Journal of Instruction-Level Parallelism*, 2005.
- [190] SAFARI Research Group. ABACuS — GitHub Repository. <https://github.com/CMU-SAFARI/ABACuS>, 2023.
- [191] Allan Snaveley et al. Symbiotic Job Scheduling for A Simultaneous Multithreaded Processor. In *ASPLOS*, 2000.
- [192] Stijn Eyerman et al. System-Level Performance Metrics for Multiprogram Workloads. *IEEE Micro*, 2008.
- [193] Pierre Michaud. Demystifying Multicore Throughput Metrics. *CAL*, 2012.
- [194] R. Zhou et al. LT-PIM: An LUT-Based Processing-in-DRAM Architecture With RowHammer Self-Tracking. *IEEE CAL*, 2022.
- [195] Stefan Saroiu et al. How to Configure Row-Sampling-Based Rowhammer Defenses. *DRAMSec*, 2022.
- [196] H. Hassan et al. CROW: A Low-Cost Substrate for Improving DRAM Performance, Energy Efficiency, and Reliability. In *ISCA*, 2019.
- [197] Zhi Zhang et al. TeleHammer: A Stealthy Cross-Boundary Rowhammer Technique. arXiv:1912.03076 [cs.CR], 2019.
- [198] Moinuddin Qureshi. Rethinking ECC in the Era of Row-Hammer. *DRAMSec*, 2021.
- [199] Seungki Hong et al. DSAC: Low-Cost Rowhammer Mitigation Using In-DRAM Stochastic and Approximate Counting Algorithm. arXiv:2302.03591, 2023.
- [200] H. Gomez et al. DRAM Row-Hammer Attack Reduction Using Dummy Cells. In *NORCAS*, 2016.

Appendix

A ABACuS Security Analysis

We explain how ABACuS maintains the maximum activation count among all sibling rows (§4) in RACs by showing that Invariant 1 holds after a RAC update upon a DRAM row activation.

Analysis Overview. Invariant 1 formally defines the property of the value stored in a row activation counter in terms of the *actual* activation count of sibling DRAM rows (i.e., each sibling row’s true activation count, regardless of the RAC’s value) in a refresh window (t_{REFW}). The spillover counter’s value is already greater than or equal to a row’s activation count for DRAM rows that are *not* tracked by any ABACuS counter [102].⁸

Invariant 1

Let $ACT_COUNT(bankID, rowID)$ denote the *actual* activation count of a row with row ID in a bank. If a row is tracked by an ABACuS counter, the row activation counter (RAC) corresponding to this row is always greater than or equal to the actual activation count of the row with the same row ID in *any* of the banks. That is, $\forall (b, r) \in (Banks, SiblingRows); RAC(r) \geq ACT_COUNT(b, r)$

Proof: By induction on the actual activation count of row r in bank b , tracked by $RAC(r)$.

Base Case: When the RAC starts tracking a row r in bank b for the first time, the following holds:

- $RAC(r) \geq ACT_COUNT(b, r)$
- $SAV(b, r)$ is set. Other SAV bits are zero.

Induction Hypothesis: Assume that invariant holds for any row r' in bank b' which are tracked by an ABACuS counter.

Step Case: Let r' be an arbitrary sibling of r in bank b' . Note that $RAC(r') = RAC(r)$ by definition of RAC. Assume that such an r' is activated. We distinguish between two cases.

Case 1: $SAV(b', r')$ is not set. In this subcase, before activating r' , we have $RAC(r') > ACT_COUNT(b', r')$ because $SAV(b', r')$ is not set. Therefore, after activating r' we have $RAC(r') \geq ACT_COUNT(b', r')$ and $SAV(b', r')$ is set.

Case 2: $SAV(b', r')$ is set. In this subcase, before activating r' , we have $RAC(r') \geq ACT_COUNT(b', r')$. Hence, after activating r' , the actual activation count of r' increases by one, and $RAC(r')$ is incremented. The $SAV(b', r')$ remains set, while other SAV bits are reset. Thus, $RAC(r') \geq ACT_COUNT(b', r')$ still holds, satisfying the invariant.

Based on the proof of the correctness of Invariant 1, we conclude that ABACuS accurately stores the maximum activation count across all sibling rows.

⁸Our extended version [159] provides a simple explanation for why the spillover counter’s value is greater than or equal to a non-tracked row’s activation count.