Assessing the Potential of Escalating RowHammer Attack Distance to Bypass Counter-based Defenses

Ranyang Zhou, *Student Member, IEEE,* Jacqueline Liu, *Student Member, IEEE,* Sabbir Ahmed, *Student Member, IEEE,* Nakul Kochar, *Student Member, IEEE,* Adnan Siraj Rakin, *Member, IEEE,* and Shaahin Angizi, *Senior Member, IEEE*

***Abstract*—This brief studies the impact of escalating DRAM RowHammer attack distance to potentially bypass well-developed counter-based defenses leveraging a multi-sided fault injection mechanism. By conducting systematic experimentation on 128 commercial DDR4 products, our results challenge recent research findings, showing that cells positioned at a greater physical distance from the target rows do not significantly affect performance across chips sourced from leading DRAM manufacturers. This implies such RowHammer models are unable to reliably bypass the latest counter-based defense mechanisms. We conduct an extensive attack design space exploration and compare the performance efficiency between this mechanism and the well-known double-sided attack.**

***Index Terms*—DRAM, RowHammer attack, attack distance.**

I. Introduction

THE far-reaching development of Deep Neural Networks (DNN) accuracy even with low-bit-width models has recently triggered various security-associated attacks in many applications [1], [2]. Recent studies show that an adversary can identify and manipulate a small number of vulnerable bits of off-the-shelf well-trained DNN weight parameters to significantly compromise the output accuracy [1]. As depicted in Fig. 1(a), such a Bit-Flip Attack (BFA) can degrade the accuracy of an 8-bit quantized ResNet-34 on ImageNet dataset from 73.1% to 0% by targeting 5 bits. BFAs have been enabled mainly due to a manifestation of a DRAM cell-to-cell interference and failure mechanism called RowHammer (RH) [3], [4]. RH attack is conducted when a malicious process activates and pre-charges a specific row (i.e., aggressor row) repeatedly to a certain threshold (T_{RH}) to induce bit-flips on immediate nearby rows (i.e., victim rows). Unfortunately, by scaling down the size of DRAM chips in the modern manufacturing process, DRAM becomes increasingly more vulnerable to RH bit-flips [5]. Fig. 1(b) shows that the T_{RH} has had a significant downward trend in recent years, e.g., the attacker needs $\sim$4.5$\times$ fewer Hammer Counts (HC) on LPDDR4 (new) as opposed to DDR3 (new) [6].

Addressing RH errors necessitates the implementation of more robust Error Correction Code (ECC) techniques, which come at the cost of excessive energy consumption, reduced performance, and capacity overhead [7]–[9]. The standard RH mitigation approach used by system manufacturers such as Apple [10] is to increase the refresh rate which imposes a humongous power consumption and can be easily compromised [7], [11]. Intel's pTRR [12] and several research works propose to proactively count the number of row activations (i.e., HC) by maintaining an array of counters in either the memory controller [13] or in the DRAM chips themselves [14]–[18]. Memory controller keeps the HC track and refreshes victim rows when the number of row activations issued to the DRAM exceeds Maximum Activate Count (MAC) threshold (T_{MAC}) which is typically saved on the Serial Presence Detect (SPD) chip within the DRAM module [17].

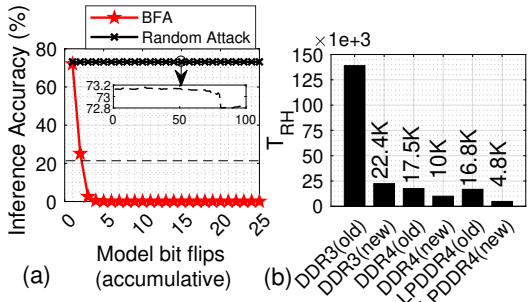

Fig. 1: (a) BFA vs. random bit flipping for an 8-bit quantized ResNet-34 on ImageNet, (b) RowHammer thresholds [6].

While conventional single- and double-sided RH attacks have been well-explored and can be potentially defended, Half-Double [19] presents a novel progression of RH attacks, showcasing its impact expanding beyond the immediate neighbors, where it flips bits in the victim rows by combining numerous accesses to a distant aggressor with just a few to a nearby aggressor. Target Row Refresh (TRR) [20] becomes a concern in such an attack as it inadvertently enables the Half-Double attack by transforming the initially refreshed row into a nearby aggressor, collaborating with the distant one that initially triggered the refresh. It has been demonstrated that the Half-Double can bypass counter-based mechanisms and grant an attacker arbitrary read and write access, e.g., on Chromebooks with ECC and TRR-protected LPDDR4x memory in an average runtime of just 45 minutes. TRRespass [17], U-TRR [21], Half-Double [19], BlackSmith [22], and SMASH [23] are the established research on multiple-row RH fault injection to create diverse RH patterns. They demonstrate such patterns can effectively induce bit-flips in DDR4 DRAM chips from all three major DRAM vendors.

The key question this work will investigate is that *Is it possible to identify an optimal set of HCs in a multi-sided RH fault injection scenario that is smaller than the HC requirement for a double-sided RH scenario to bypass the MAC?* The contributions of this paper are as follows: (1) We explore the impact of escalating RH beyond the adjacent row to unlock its attack potential. In this way, we consider a multi-sided fault injection model to fool the system, elevate the cost of counter-based defense, and in the end, overcome the established defense; and (2) We extensively analyze and experimentally

This work is supported in part by the National Science Foundation under Grant No. 2228028.

R. Zhou, N. Kochar, and S. Angizi are with the Department of Electrical and Computer Engineering, New Jersey Institute of Technology, Newark, NJ, USA. E-mail: {rz26,nk548,shaahin.angizi}@njit.edu.

J. Liu, S. Ahmed, and A. Siraj Rakin are with the Department of Computer Science, State University of New York at Binghamton, NY, USA. E-mail: {jliu28,sahmed9,arakin}@binghamton.edu.

verify the impact of such an attack model on 128 DDR4 DRAM chips across various manufacturers, namely Samsung, Micron, etc., with counter-based RH protection mechanisms enabled. Our findings challenge recent conclusions such as the one in [19], indicating that cells located at a greater physical distance from the target rows offer no or negligible improvement in generating bit-flips across various chips. In other words, it seems that only the adjacent rows play the key role in inducing bit-flips.

II. Background & Motivation

DRAM Organization. A DRAM chip comprises a two-dimensional matrix of memory cells, organized into sub-arrays called mats within each bank. Modern DRAM chips contain billions of these cells [21]. Each bit-cell includes a capacitor and an access transistor, with the capacitor's charge state representing binary data: full-"1" or empty-"0".

DRAM Timing Parameters. The most basic DRAM timing parameter is the clock cycle (t_{CK}). Row Active Time (t_{RAS}) encompasses the temporal window demarcating an activation (ACT) command and the subsequent precharge (PRE) command. During the prescribed t_{RAS} interval, the restoration of charge within the DRAM cells on the open DRAM row is effectuated to ensure optimal performance. Row Precharge Time (t_{RP}) signifies the temporal gap between the issuance of a PRE command and the subsequent ACT command. The imposition of t_{RP} is instrumental in closing the open WL and initiating the pre-charging of the DRAM BLs to the voltage level of $\frac{V_{DD}}{2}$. Retention time in DRAM refers to the duration for which a memory cell can hold its stored data without requiring a refresh operation. The Refresh Window (t_{REFW}) is essentially the interval within which all DRAM cells must be refreshed to prevent data loss or corruption.

RH Attack. Kim et al. [3] were the pioneers in conducting an extensive study on the characteristics of RH bit-flips in DDR3 modules. They observed that approximately 85% of the tested modules were susceptible to RH attack. Therefore, the majority of earlier RH research is centered on DDR3 systems [24]. With the prospect of having an RH-less landscape, DDR4 modules have been introduced. While there are documented instances of RH on DDR4 chips in previous studies [25], [26], these findings pertain to earlier generations of DDR4.

RH Defense. The hardware-based RH mitigation mechanisms can be classified into two categories, i.e., *victim-focused* mechanism with probabilistic refreshing (e.g., PRA [14], PARA [3]) and *aggressor-focused* mechanism by counting activations (e.g., U-TRR [21], Hydra [15], TWiCe [8], Graphene [27]). The system manufacturers tend to follow the mechanisms that explicitly detect RH conditions and intervene, such as increasing refresh rates and access counter-based approaches. However, such methods require add-on hardware to calculate rows' activation [14]–[17] and record it to other fast-read-memory (SRAM [8]/CAM [27]). The controller will then refresh the target row if the number reaches MAC [17]. The JEDEC standard outlines three potential configurations for the MAC value: (1) unlimited, if the DRAM module claims to be RH-free; (2) untested, if the DRAM module has not undergone post-production inspection; or (3) T_{MAC} indicating the specific number of ACTs the DRAM module can withstand

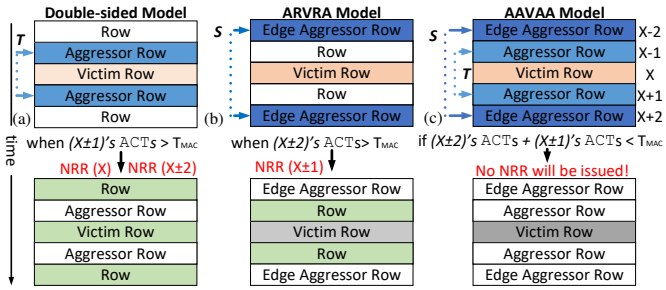

Fig. 2: RowHammer attack models (top) and refreshed rows after NRR commands (bottom): (a) Double-sided model, (b) under-test ARVRA model, (c) Proposed AAVAA model.

(e.g., 1M). It has been revealed in [17] that, irrespective of the DRAM manufacturer, the majority of DDR4 modules assert unlimited MAC value.

III. Proposed Model and Assessment

To enhance the defensive capabilities of DRAM modules, it is necessary to adopt an attacker's perspective, enabling a deeper comprehension of potential threats and more effective countermeasures. Existing counter-based RH prevention frameworks come with distinct challenges, specifically in their scope and thresholds. From an attacker's perspective, we can articulate three essential directions to defeating counter-based frameworks: (i) Broaden the attack area as extensively as possible to make the detection more complicated; (ii) Leverage various attack patterns such as side-kick aggressors or many-sided attacks [17], [19], [28]; and (iii) Reduce the HC if possible to fool the system by not being detected. Our objective here is to explore the third direction to elevate the cost of counter-based defense in DDR4 modules and, in the end, overcome established mitigation techniques.

Traditional fault injection models such as double-sided attacks can be effectively defended [12], [28] by counter-based frameworks. As shown in Fig. 2(a)-top, the double-sided RH model mainly affects the victim rows with two aggressors X±1. While there are three victim rows in this model, the primary focus of this approach is on victim row X, as both aggressor rows simultaneously exert a significant influence on it. Subsequent testing allows us to establish a range of aggressor rows' HCs, denoted by T, that effectively quantifies the vulnerability levels of the victim rows. The lower and higher boundaries of T correspond to the respective thresholds where the victim row first exhibits bit-flips and where the victim row is entirely flipped due to the attacks, respectively. Hence, defense mechanisms will easily identify anomalous rows that have been activated significantly more frequently than typical rows. As discussed, such defenses establish distinct thresholds depending on the manufacturer of the chips. If the defense mechanism properly detects that the row X±1 reaches the T_{MAC}, the Nearby Row Refresh command (NRR) will refresh row X and X±2 as shown in Fig. 2(a)-bottom. Fig. 3 shows the timing for such an RH attack. Assuming RH is implemented on the row 0x99. F is a flag used to decide whether to issue an NRR command or not. The memory controller issues an NRR for that row when HC surpasses MAC, which means HC$\leq \frac{T_{MAC}}{t_{RAS}}$. Common t_{RAS} values for DDR4 memory modules could range from

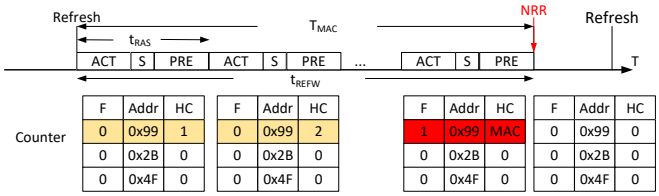

Fig. 3: Timing of RH attack.

around 36 to 48 t_{CK} [29], but these values can differ based on the module's speed rating (e.g., DDR4-2133, DDR4-2400, DDR4-3200, etc.). The duration of a clock cycle for DDR4-2400 memory can be calculated as $t_{RAS} = \frac{1}{2400MT/s}$. In our design, every t_{RAS} consists of three parts: `ACT`, `Sleep(S)`, and `PRE`, where `Sleep(S)` is set to $5{\times}t_{CK}$. Based on the existing conditions, we can give HC a limit value of 1M, which is applicable to DRAM chips of all frequencies on the market.

The multiple-row fault injection model herein represents a new concept rooted in the traditional double-sided RH attack model with an expanded set of attack vectors. Through the assessment, our objective is to identify potential HCs to enable a so-called soft-attack, i.e., weaker attack with fewer aggressor rows' `ACT`, as a means to circumvent counter-based defenses regardless of DRAM controller implementation details in various DRAM chips. The under-test *ARVRA RH model* depicted in Fig. 2(b)-top is a straightforward variant of the double-sided model, in which the two edge aggressor rows (X±2) that are one row apart from the targeted victim row, hammer it S times. Our working hypothesis is when (X±2)'s `ACTs` is greater than T_{MAC}, ARVRA model can victimize all three sandwiched rows and may effectively flip the bits in the targeted victim row X. By issuing the NRR command, X±1 rows will be refreshed where the victim row remains flipped as shown in Fig. 2(b)-bottom. The proposed *AAVAA RH model* aims to combine the double-sided model and ARVRA model to find a way to bypass the T_{MAC}. In this model, as shown in Fig. 2(c)-top, so-called edge aggressor rows (X±2) and typical aggressor rows (X±1) are soft-attacked/hammered S and T times, respectively. *Our working hypothesis is that DRAM modules exposed to AAVAA might be vulnerable to certain reduced hammering patterns by which (X±2)'s `ACTs` + (X±1)'s `ACTs` is less than T_{MAC}.* In this case, no counter-based technique will be able to figure out which row is victimized, and therefore no NRR command will be issued.

IV. Characterization and Observations

Framework Setup and Testing Infrastructure. We test the DRAM chips by modifying the DRAM-Bender [30] to have a versatile FPGA-based DRAM attack exploration framework for DDR4 with an in-DRAM compiler API installed on our host machine. Our testing infrastructure, as shown in Fig. 4, consists of the Alveo U200 Data Center Accelerator Card [31] as the FPGA that accepts DDR4 modules and runs the test programs by sending DDR4 command traces generated by the host machine. Besides, to have a fair comparison among various DRAM chips, the temperature is kept below 30°C with the INKBIRDPLUS 1800W controller.

Minimizing Interference. Before implementing the attack scenarios, DRAM refresh [20] and rank-level ECC are disabled to minimize their interference with RH bit-flips. Unlike proprietary RH protection techniques such as TRR [5], [19],

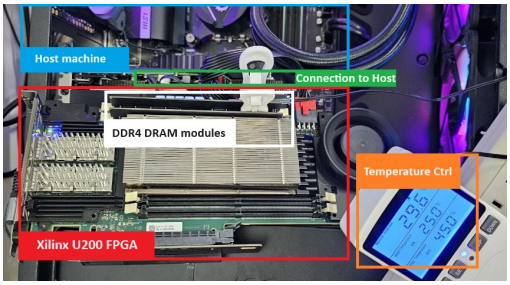

Fig. 4: Our testing infrastructure for DDR4 modules.

[32], which implement refresh operations, our experiment removes these protections to directly observe the behavior of bit-flips caused by RH.

Chips Tested. The experiments are conducted on a range of 128 commercialized DRAM chips from eight different manufacturers (mf.) in Table I with various die densities and die revisions. We tested 10-20 random rows from each bank of various chips. For instance, Samsung chips showed bit-flip variations up to 200, indicating consistent bit-flip generation within the same chip. We then selected one row for our experiment. To ensure accuracy and minimize data fluctuations, each activation count was repeated 10 times and averaged.

TABLE I: Under-test DRAM chips.

Vendor	#Chips	Freq (MHz)	Die rev.	Org.	Date
mf-A (Micron 16GB)	16	2133	B	x4	2126
mf-B (ATECH 16GB)	16	2933	A	x8	2597
mf-C (Crucial 16 GB)	16	3200	C	x8	N/A
mf-D (Kingston 16GB)	16	2666	G	x8	2152
mf-E (NEMIX 16GB)	16	2133	B	x4	1733
mf-F (SK Hynix 16GB)	16	2400	A	x8	1817
mf-G (Patriot Viper 16GB)	16	3600	C	x8	N/A
mf-H (Samsung 16GB)	16	2400	B	x8	2053

Results and Observations. To study the effectiveness of under-test models on read disturbance, we comprehensively analyze the AAVAA attack on various (S, T) configuration sets. Please note that ARVRA is a sub-set of AAVAA whereby T = 0. The characterization method remains consistent for DRAM modules from eight distinct manufacturers. It includes incrementing both S and T HCs to assess the effects of all conceivable combinations. The 3-D surface plots presented in Fig. 5 reveal distinct characteristics for each design. To facilitate the understanding of the results, we present a 2-D plane in Fig. 5(c). This figure shows a fixed number of bit-flips incorporated within different patterns of S and T.

We systematically examine 16 chips from every manufacturer through a rigorous testing process. Our findings reveal that the performance differentials among these chips are minimal, with variations consistently below the 5% threshold, attesting to their uniform quality and reliability. This implies that the test outcomes for individual chips indicate the overall impact of RH on each manufacturer's chips. Herein, as we acquire multiple samples of identical chips, we take the average to create a representative plot, ensuring a more accurate depiction of the chips' characteristics.

Obs.#1. The impact of the edge aggressor rows on the victim row is considerably lower than that of the standard aggressor rows.

Fig. 6(a) illustrates the differential impact of edge aggressor rows compared to standard aggressor rows on a victim row during an RH attack. The data from the AAVAA model of mf-C is plotted and labeled when $S = 0$ and $T = 0$. It is evident that when $S = 0$, the rate of bit-flips surpasses that

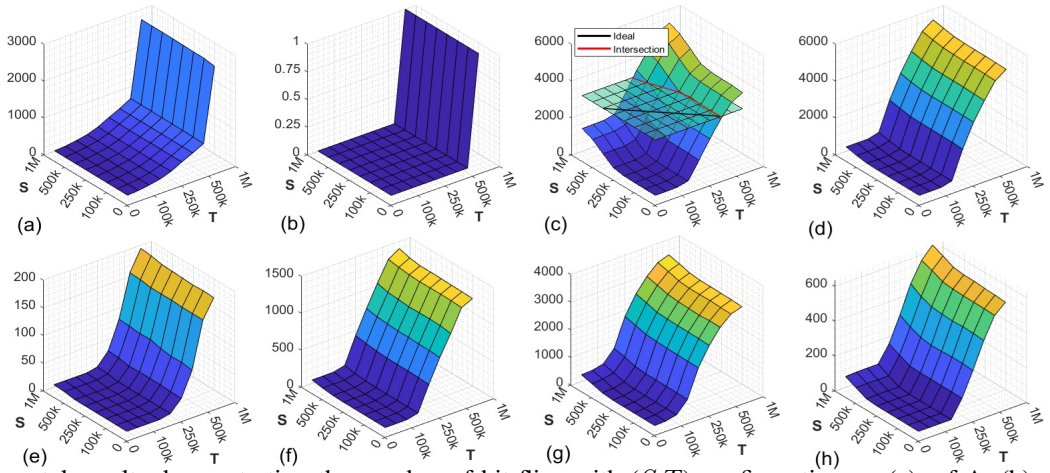

Fig. 5: Experimental results demonstrating the number of bit-flips with (S,T) configuration on (a) mf-A, (b) mf-B, (c) mf-C, (d) mf-D, (e) mf-E, (f) mf-F, (g) mf-G, (h) mf-H.

observed when $T = 0$. In other words, when S and T share identical values, the bit-flips for $S = 0$ exceed those for $T = 0$. This observation leads to the conclusion that the greater the distance, the smaller the impact on the victim row.

Obs.#2. Our assessment reveals that 75% of the modules demonstrate significantly higher resilience against aggressor rows from far distances.

For mf-A chips characterized in Fig. 5(a), as T increases, the number of bit-flips remains nearly constant, indicating that the fault injection technique is unable to decrease S by elevating T. The distinction lies in Fig. 5(b), wherewith the rise in T, the increment in the number of bit-flips surpasses that of increasing S. However, a notable issue with Fig. 5(b) is the extremely low total number of bit-flips which means mf-B chips are robust against any RH fault injection method. Although the bit-flips significantly increase with both S and T, the impact remains negligible at the chip level. Moreover, both mf-C and mf-H chips in Fig. 5(c)(h) can generate a substantial number of bit-flips, and T demonstrates a clear influence on bit-flips. However, both of them also reach a high threshold of S that affects bit-flips. The potential impact is that there may be a point where neither S nor T reaches the threshold of the tracking mechanisms as shown in Table II.

Obs.#3. The attacker cannot conduct a successful multi-sided RH attack on the under-test chips with a significantly smaller HC than the double-sided model.

The points where the planes intersect with the plot represent all HCs capable of producing that specific quantity of bit-flips. Taking mf-H characterized in Fig. 5(h) as an example, both the double-sided and the ARVRA models' HC reach 1M. We select the bit-flips equal to 3000 as the benchmark to draw the plane that crosses the plot. The intersection (red line) and the ideal pattern (black line) of the mesh graph in mf-C (Fig. 5(c)) is converted into a 2D plot in Fig. 6(b). We observe that the total number of HCs ($S + T$) in the AAVAA model is always higher than in the double-sided model., we can identify the optimal pattern along this intersection line. This pattern should satisfy the following criteria: (i) Both S and T in the pattern must be smaller than the T in the double-sided model or the S in the ARVRA model; (ii) The values of S and T should be as close or even equal as possible for soft-hammering; and (iii) The sum of S and T should be less than or equal to 700k which is represented as the ideal pattern in mf-C in Fig. 5(c). We observe that in Fig. 5, no chips can meet the three conditions at the same time. The optimal configuration set can be found at (700k, 250k) in mf-C chips (see Fig. 6(b)), in which there remains a disparity between the ideal patterns and the intersection. Here, bit-flips will only experience a notable increase when T surpasses S, and *this contradicts our initial hypothesis where we need neither T nor S to be higher than the HC required for the double-sided model otherwise, the attack will get detected.*

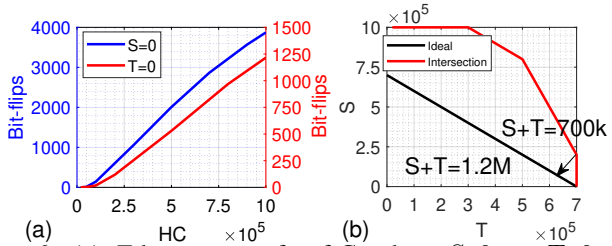

Fig. 6: (a) Edge curve of mf-C when S=0 or T=0, (b) Intersection curve and the ideal pattern.

Drawing from the aforementioned observations, the double-sided RH model remains the most direct and efficient attack method when analyzing chips sourced from leading DRAM manufacturers in the market. *Our analysis emphasizes that the majority of chips exhibit resistance to non-adjacent attack aggressors.* Therefore, merely augmenting the quantity of non-adjacent attack rows does not enhance the efficiency of RH. Hence, *to circumvent counter-based defense, the core strategy must rely on novel techniques capable of significantly diminishing the HCs such as RowPress attack [33].* RowPress indicates that bit-flips can occur when the activated row is not promptly closed and remains open for an extended period.

Practicality of Current Counter-based Mechanisms. In Table II, we summarize several previous counter-based defense mechanisms to underscore their robustness against the multi-sided attack model. Our observation reveals that Graphene [27] stands out as a reliable and practical method due to its

TABLE II: Generic RH mitigation frameworks.

Framework	capacity overhead	area overhead	defense threshold
Graphene [27]	0.53MB‡+1.12MB†	1 counter	50k
Hydra [15]	56KB†+4MB*	1 counter	500
TWiCE [8]	3.16MB†+1.6MB‡	1 counter	32,768
Counter-per-Row	32MB*	16384 counters	customized
Counter-Tree [16]	2MB*	1024 counters	customized

*The capacity overhead of DRAM. †The capacity overhead of SRAM. ‡The capacity overhead of CAM.

minimal capacity overhead and low RH detection. According to the results in Fig. 5, all the under-test chips do not generate bit-flips with the 50k threshold, which means Graphene can defend RH with minimum overhead. Given that counter-based mechanisms necessitate the incorporation of extra counters and space for storing the counting table, each such mechanism inherently incurs unavoidable overhead. Hence, improving performance entails minimizing the usage of counters and storage space. In addition, due to different algorithms, the number of HCs (defense threshold) that various mechanisms can defend are also different. Attackers need to activate aggressor rows many times, so the lower the HC threshold can be defended, the less possibility that attackers can flip the bits. Thus, as indicated in the table, it is evident that counter-per-row and counter-tree mechanisms [16] possess the capability to defend against attacks irrespective of the magnitude of the HC threshold. However, they necessitate a large number of counters and lack practical utility. Graphene [27], TWiCE [8], and Hydra [15] require only one counter. While Graphene consumes the smallest amount of storage space compared to the others, its defense threshold surpasses that of the other two mechanisms. At the same time, Hydra has the lowest defense threshold, but it takes up a lot of storage.

Discussions. Our study extends the hammering cycles beyond the typical refresh interval limitations. The total number of hammering cycles reached 4 million, which significantly exceeds the commonly accepted maximum of 1.6 million within a typical refresh window. However, this extension is justified by the need to build a robust bit-flip tendency model. Previous works, such as Half-Double [19] and DearDRAM [34], have demonstrated that extending the number of memory accesses (up to 10 million in Half-Double [19]) provides valuable insights into the behavior of bit flips under prolonged hammering. This approach ensures that our model is both comprehensive and accurate. To further refine our analysis, we differentiated between retention and RH-induced errors. Following the method used in RAIDR [35], we observed that only approximately 30 cells fail to tolerate a doubled refresh interval, and around 10^3 cells fail at four times the interval. The retention error rate remains around $\frac{1}{10^6}$, while the total bit flips are at the 10^3 level.

V. Summary

This study examines the effect of escalating DRAM RH attack distance to potentially bypass counter-based defenses leveraging a multi-sided fault injection mechanism. Through testing of 128 DDR4 chips from major manufacturers, we shed light on the resilience of DRAM chips showing that (i) The impact of the edge aggressor rows on the victim row is considerably lower than that of the standard aggressor rows, (ii) 75% of the DRAM modules demonstrate significantly higher resilience against aggressor rows from far distances, and (iii) The attacker cannot conduct a successful multi-sided attack on the under-test chips with a significantly smaller HC than the double-sided model. Overall, we conclude increasing the quantity of non-adjacent attack rows does not enhance the efficiency of RH. To overcome counter-based defense, one should choose novel techniques capable of significantly diminishing the HCs such as the RowPress attack [33].

References

[1] A. S. Rakin *et al.*, "Bit-flip attack: Crushing neural network with progressive bit search," in *ICCV*, 2019, pp. 1211–1220.
[2] V. Van Der Veen *et al.*, "Drammer: Deterministic rowhammer attacks on mobile platforms," in *CCS*, 2016, pp. 1675–1689.
[3] Y. Kim *et al.*, "Flipping bits in memory without accessing them: An experimental study of dram disturbance errors," *ACM SIGARCH Computer Architecture News*, vol. 42, no. 3, pp. 361–372, 2014.
[4] O. Mutlu *et al.*, "Fundamentally understanding and solving rowhammer," in *Proceedings of the 28th ASPDAC*, 2023, pp. 461–468.
[5] J. S. Kim *et al.*, "Revisiting rowhammer: An experimental analysis of modern dram devices and mitigation techniques," in *ISCA*. IEEE, 2020.
[6] J. Woo *et al.*, "Scalable and secure row-swap: Efficient and safe row hammer mitigation in memory systems," *arXiv:2212.12613*, 2022.
[7] O. Mutlu and J. S. Kim, "Rowhammer: A retrospective," *IEEE TCAD*, vol. 39, 2019.
[8] E. Lee *et al.*, "Twice: Preventing row-hammering by exploiting time window counters," in *ISCA*, 2019, pp. 385–396.
[9] R. Zhou *et al.*, "P-pim: A parallel processing-in-dram framework enabling rowhammer protection," 2023.
[10] (2015) Apple, inc. about the security content of mac efi security. [Online]. Available: https://support.apple.com/en-au/HT204934.
[11] A. S. Rakin *et al.*, "Deep-dup: An adversarial weight duplication attack framework to crush deep neural network in multi-tenant fpga," in *USENIX Security*, 2021, pp. 1919–1936.
[12] M. Kaczmarski, "Thoughts on intel xeon e5-2600 v2 product family performance optimisation–component selection guidelines," 2014.
[13] K. S. Bains *et al.*, "Method, apparatus and system for providing a memory refresh," 2015, uS Patent 9,030,903.
[14] D.-H. Kim *et al.*, "Architectural support for mitigating row hammering in dram memories," *IEEE CAL*, vol. 14, no. 1, pp. 9–12, 2014.
[15] M. Qureshi *et al.*, "Hydra: enabling low-overhead mitigation of row-hammer at ultra-low thresholds via hybrid tracking," in *ISCA*, 2022.
[16] S. M. Seyedzadeh *et al.*, "Counter-based tree structure for row hammering mitigation in dram," *CAL*, vol. 16, 2016.
[17] P. Frigo *et al.*, "Trrespass: Exploiting the many sides of target row refresh," in *SP*. IEEE, 2020, pp. 747–762.
[18] Y. Wang *et al.*, "Detect dram disturbance error by using disturbance bin counters," *IEEE CAL*, pp. 35–38, 2019.
[19] A. Kogler *et al.*, "{Half-Double}: Hammering from the next row over," in *USENIX Security*, 2022, pp. 3807–3824.
[20] (2020) Jesd79-4c: Ddr4 sdram standard. [Online]. Available: https://www.xilinx.com/products/boards-and-kits/alveo.html
[21] H. Hassan *et al.*, "Uncovering in-dram rowhammer protection mechanisms: A new methodology, custom rowhammer patterns, and implications," in *MICRO*, 2021, pp. 1198–1213.
[22] P. Jattke *et al.*, "Blacksmith: Scalable rowhammering in the frequency domain," in *2022 IEEE SP*. IEEE, 2022, pp. 716–734.
[23] F. de Ridder *et al.*, "{SMASH}: Synchronized many-sided rowhammer attacks from {JavaScript}," in *USENIX Security)*, 2021, pp. 1001–1018.
[24] M. Seaborn and T. Dullien, "Exploiting the dram rowhammer bug to gain kernel privileges," *Black Hat*, vol. 15, p. 71, 2015.
[25] M. Lipp *et al.*, "Nethammer: Inducing rowhammer faults through network requests," in *EuroS&PW*. IEEE, 2020, pp. 710–719.
[26] D. Gruss *et al.*, "Another flip in the wall of rowhammer defenses," in *SP*. IEEE, 2018, pp. 245–261.
[27] Y. Park *et al.*, "Graphene: Strong yet lightweight row hammer protection," in *MICRO*. IEEE, 2020, pp. 1–13.
[28] S. Saroiu *et al.*, "The price of secrecy: How hiding internal dram topologies hurts rowhammer defenses," in *IRPS*. IEEE, 2022, pp. 2C–3.
[29] H. Choi *et al.*, "Reducing dram refresh power consumption by runtime profiling of retention time and dual-row activation," *MICPRO*, 2020.
[30] A. Olgun *et al.*, "Dram bender: An extensible and versatile fpga-based infrastructure to easily test state-of-the-art dram chips," *TCAD*, 2023.
[31] (2021) Xilinx inc., xilinx alveo u200 fpga board. [Online]. Available: https://www.xilinx.com/products/boards-and-kits/alveo.html
[32] L. 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, pp. 1182–1197.
[33] H. Luo *et al.*, "Rowpress: Amplifying read disturbance in modern dram chips," in *Proceedings of the 50th ISCA*, 2023, pp. 1–18.
[34] X. Zhan *et al.*, "Deardram: Discard weak rows for reducing dram's refresh overhead," in *ACA*. Springer, 2018, pp. 109–124.
[35] J. Liu *et al.*, "Raidr: Retention-aware intelligent dram refresh," *ACM SIGARCH Computer Architecture News*, vol. 40, no. 3, pp. 1–12, 2012.