Rowhammer Vulnerability of DRAMs in 3-D Integration

Eduardo Ortega, Jonti Talukdar, Woohyun Paik, Tyler Bletsch, and Krishnendu Chakrabarty

***Abstract*— We investigate the vulnerability of 3-D-integrated dynamic random access memorys (DRAMs) [i.e., typically connected with silicon via (TSV), monolithic interconnect via (MIV)] to Rowhammer attacks. We have developed a SPICE framework to characterize Rowhammer attacks for the scenarios described. We utilize OPENROAD ASAP7 PDK for our simulation. We investigate horizontal (within the same tier) and vertical (across multiple tiers) variants of Rowhammer attacks. We show that horizontal Rowhammer vulnerability may be reduced through DRAM bank partitioning. In addition, we show that vertical parasitic capacitance in TSV 3D-DRAM is unlikely to lead to vertical Rowhammer attacks. However, vertical parasitic capacitance in MIV 3D-DRAM can make vertical Rowhammer attacks feasible.**

***Index Terms*— 3-D integration, dynamic random access memory (DRAM), Rowhammer.**

I. Introduction and Motivation

Rowhammer is a memory vulnerability that impacts system-level security across multiple generations of dynamic random access memory (DRAM) [1], [2], [4], [5], [7]. It occurs when a malicious actor accesses the same DRAM row multiple times. These multiple accesses cause disturbances that result in bit flips in physically adjacent DRAM rows (victim rows). The frequently accessed row(s) are known as the aggressor row(s). The minimum number of aggressor accesses to cause bit flips is called "hammer count first" (HC_{first}). HC_{first} has been on a historical downtrend since Rowhammer was first reported (i.e., HC_{first} decreased from 139 K in 2014 to 4.8 K in 2020, and 3.2 K in 2022) [1], [3], [7]. It is projected that HC_{first} will continue to decrease due to technology scaling. In this brief, we provide an analysis of Rowhammer attacks on next-generation DRAM memory. Typically, Rowhammer is known to flip bits within a DRAM bank that is layered horizontally (e.g., DDR4 Bank). However, 3D-DRAM has enabled higher-density DRAM DIMMs by stacking DRAM memory vertically. These vertical stacks of 3D-DRAM memory are typically connected with silicon vias (TSVs) [16]. This style of stacked memory with TSVs is vital for the heterogeneous integration of memory in next-generation packages. We refer to heterogeneous integration as HI hereafter. In addition, monolithic interconnect vias (MIVs) are another emerging packaging technology. MIVs provide greater packaging density than TSV-based packaging [17]. Hence, MIV-based 3D-DRAM is also integral for HI of memory. We simulate both scenarios to encapsulate the Rowhamer vulnerability of emerging DRAM chips.

Manuscript received 30 September 2023; revised 10 December 2023 and 8 February 2024; accepted 18 February 2024. Date of publication 27 February 2024; date of current version 26 April 2024. This work was supported in part by National Science Foundation (NSF) under Grant CNS-2011561 and in part by Semiconductor Research Corporation (SRC) under Grant 3199.001. *(Corresponding author: Eduardo Ortega.)*

Eduardo Ortega and Krishnendu Chakrabarty are with the School of Electrical, Computer, and Energy Engineering, Arizona State University, Tempe, AZ 85287 USA (e-mail: eeortega@asu.edu; krishnendu.chakrabarty@asu.edu).

Jonti Talukdar, Woohyun Paik, and Tyler Bletsch are with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: jonti.talukdar@duke.edu; woohyun.paik@duke.edu; tyler.bletsch@duke.edu).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TVLSI.2024.3368044.

Digital Object Identifier 10.1109/TVLSI.2024.3368044

We investigate the interlayer parasitic capacitance of TSV- and MIV-based 3D-DRAM. As packaging density is increasing, we hypothesize that the resultant parasitics from vertical interconnects (i.e., TSVs, MIVs, etc.) in 3D-DRAM may lead to a vertical evolution of the Rowhammer attack (Rowhammer across DRAM memory layers). For example, if a Rowhammer attack occurs at DRAM layer n then the vertical parasitic may induce bit flips at the DRAM layer $n \pm 1$. We refer to Rowhammer across DRAM layers as vertical Rowhammer, and Rowhammer within the same DRAM layer as horizontal Rowhammer. We simulate DRAM memory operations through SPICE with the OPENROAD ASAP7 PDK (7 nm node size) [12]. Our simulations consider parasitic based on interconnect geometry from ASAP7, analysis of real-silicon MIVs, and bonding between layers of TSV-based 3D-DRAM. Overall, this brief makes the following contributions.

1) We show that DRAM layout configurations (i.e., WL partitioning) can greatly impact the Rowhammer threshold (HC_{first}).
2) We provide evidence that vertical Rowhammer is unlikely for TSV-based 3D-DRAM in HI packages due to interlayer distance.
3) We provide evidence that vertical Rowhammer may be feasible for MIV-based 3D-DRAM in HI packages.

II. Prior Rowhammer Studies

The first DRAM study that showed the underlying disturbance (Rowhammer) was in 2014 [1]. It was shown that Rowhammer depends on aggressor row data patterns and the induced bit flips are repeatable for many consumer products [1]. In addition, some consumer products were biased to flip more bits from "1" to "0" than from "0" to "1" [1]. Another notable study showed the historical trend that HC_{first} was becoming smaller [7]. In conjunction, [7] also shows that Rowhammer can flip bits farther away than the immediate aggressor row, and newer DRAM chips are increasingly more vulnerable to Rowhammer.

Rowhammer vulnerability and the decreasing HC_{first} trend are shown through an evaluation using DRAM chips in [1] and [7], respectively. Orosa et al. [11] provide insights into Rowhammer through an experimental study of DDR4 and DDR3 DRAM chips. In addition, it has been observed that HC_{first} decreases when temperature increases [11]. Furthermore, Rowhammer vulnerability increases by up to 36% when the aggressor WL stays active longer or when the precharge occurs for a longer period [11].

Yang and Lin. [17] present a TCAD simulation that shows how charge pumping from interface traps is one of the underlying mechanisms that induce the Rowhammer vulnerability. Walker et al. [20] show how electron injection/capture and capacitive crosstalk induce Rowhammer vulnerability. Yang and Lin. [17] and Walker et al. [20] show that an increase in temperature can induce more carrier movement to the victim cells and victim interconnects (i.e., victim WL) and thus decrease HC_{first}. Hence, higher temperature and longer aggressor WL activation periods may help increase Rowhammer vulnerability, as corroborated in a chip study [11]. In addition, [19] shows that the underlying parasitic enables charge recombination (sub-threshold leakage) at victim cells due to neighboring aggressor cells and corresponding bitlines (BLs). Thus, longer precharging periods, also seen in [11], will lead to a greater impact on victim cells due to

TABLE I
INTERCONNECT GEOMETRY FOR ASAP7 PDK [12]

Metal Layer	Width	Thickness	Pitch
M1 through M3	18 nm	9 nm	36 nm
M4 and M5	24 nm	12 nm	48 nm
M6 and M7	32 nm	16 nm	64 nm
M8 and M9	40 nm	20 nm	80 nm

Rowhammer. These circuit-level analyses of Rowhammer vulnerability corroborate recent DRAM chip studies [11], [17], [19], [20].

Fabrication methods have been studied to alleviate the underlying causes of Rowhammer (e.g., charge interface traps and crosstalk) [18]. However, recent DRAM chip studies show the Rowhammer vulnerability is still present in commercial DRAM products and is projected to worsen [2], [7], [11].

III. SIMULATION SETUP AND METHODOLOGY

We carry out transient SPICE simulation of DRAM aggressor and victim cells through OPENROAD ASAP7 PDK [12]. In the ASAP7 PDK, we utilize the transistors that offer lower subthreshold leakage to represent real memory cell design (i.e., nMOS_sram) [12]. We also consider process variation of our devices (i.e., TT, FF, and SS) [12]. We insert parasitics into our SPICE simulation to represent coupling capacitance and wire resistance [15]. The inserted parasitics are lumped Π models based on the stated interconnect geometry [12], [15]. Equations (1) and (2) represent the capacitive crosstalk as discussed in prior circuit-level studies [17], [20]. The parameter A in (1) and (2) represents the area of the two facing metal interconnects, and the cross section of the metal interconnect, respectively. In (1), ϵ_{SiO_2} represents the permittivity of silicon dioxide ($3.9\epsilon_0$) [18]. In (2), ρ_{cu} represents the resistivity of Cu (1.7×10^{-8} Ωm) [18]

$$C_{coupling} = \varepsilon_{SiO_2} \left(A_{facing}/d\right) \tag{1}$$

$$R_{wire} = \rho_{cu} \left(l/A_{cross\,section}\right). \tag{2}$$

By inserting these computed resistance (R) and capacitance (C) values, we can incorporate the relative parasitic mechanisms that induce Rowhammer in our SPICE simulation. We assume that SiO_2, a common dielectric material, is used as the insulation material [18]. In addition, we assume Cu metal interconnects for the WL and BL, as specified by the ASAP7 PDK [12]. We utilize the interconnect geometry as specified by ASAP7 PDK (Table I) [12]. Note that the Fin and Gate Width (Pitch) are 7 nm (27 nm) and 20 nm (54 nm), respectively [12]. To accurately model practical DRAM architectures, we consider layout and architectural performance optimizations described in [9] and [10]. These DRAM performance optimizations make use of the partitioning of interconnects and peripheral circuits to reduce power and memory access time. They reduce memory access times by introducing subarray/subbank level parallelism for DRAM bank accesses [9]. This is done by using more peripheral circuit sense amplifiers (SAs) that divide the BL into smaller interconnects. Power optimizations are also considered in our study. To optimize power, the WL is partitioned to lower overall activation energy (fine-grain activation) [10]. By reducing overall interconnect WL length, DRAM architects can increase the overall performance and power efficiency of their memory [9], [10]. For our SPICE simulation, we consider subarrays/sub-banks partitioning to achieve smaller interconnects between DRAM cells (i.e., WL and BL). We divide the BL in our DRAM bank into partitions of 32 subarrays following [10]. In addition, we partition the WL of each subarray into 32 blocks for further granularity, also following [10]. Each partitioned WL is split further to represent fine-grain activation (2, 4, and 8 splits) [10]. The overall DRAM architecture is illustrated in Fig. 1.

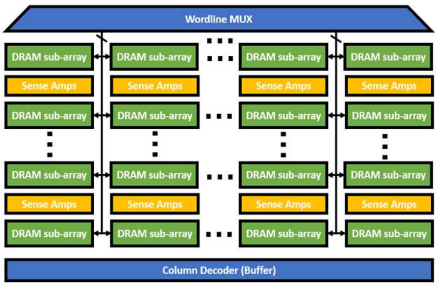

Fig. 1. DRAM Bank block-level layout example.

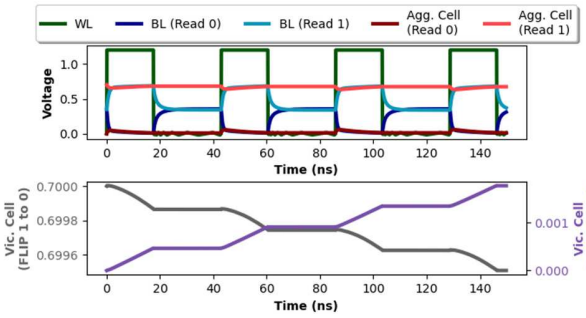

Fig. 2. Aggressor DRAM Read operation.

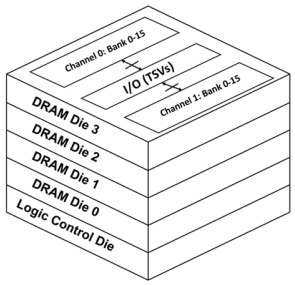

Fig. 3. TSV-based 3D-DRAM example.

The DRAM bank capacity for our SPICE simulation is modeled after MICROMETER MT40A2G4 (128 K rows and 1 K columns per bank group) [8]. The operating voltage of the DRAM is assumed to be 1.2 V ± 60 mV [8]. The devices provided by ASAP7 have an operating voltage (V_{DD}) of 0.7 V. Typically, in DRAM memory, WL is pulled higher than the operating voltage of the device to ensure that the memory cell is written into correctly [2]. In our simulation, we utilize the DRAM operating voltage as our WL voltage (1.2 V). We utilize the voltage variance of [8] to model the voltage variance of our simulation DRAM (i.e., V_{DD} ± 60 mV). To represent realistic DRAM cells and DRAM data access operation(s), we implement a SPICE version of a deep trench capacitor (DTC) and SA, respectively [13], [15]. The DTC SPICE model was selected because the design is based on real silicon-fabricated trench capacitors used for high-density capacitor applications such as DRAM [13]. The latch-based SA was chosen as this is a known peripheral circuit for DRAM memory operations [15]. Fig. 2 shows typical waveforms for reading a DRAM cell with an SA. The waveform represents a Read-based aggressor memory access pattern (i.e., read 0 and read 1) because multiple WL activations are occurring in the same row (Fig. 2). In addition, we show the read-based aggressor impact from hammering nearby victim cells for the considered bit flip scenarios, i.e., FLIP 0 to 1 and FLIP 1 to 0 (Fig. 2). We discuss the specifics of how certain Read operations may induce bit flips in Section IV.

There are several considerations for simulating 3D-DRAM and the vertical parasitics for vertical Rowhammer [6], [14], [21]. We consider memory placement and packaging density. TSV-based

3D-DRAM is packaged with 50 μm layers of DRAM chips (Fig. 3) [14]. For vertical integration, TSVs are utilized. However, TSVs introduce physical stress, therefore careful floorplanning is carried out to ensure that devices are placed a certain distance away from the TSVs (i.e., keep out zone) [14]. TSV interconnects act as a vertical bus for logic signals. These logic signals communicate the necessary stimuli to execute the DRAM memory operation via peripheral circuits (i.e., global WL, local WL, SA, etc.). However, these logic signals and related TSVs, are physically too far away to be related to the underlying mechanisms that induce Rowhammer. Thus, in our analysis, TSV-induced parasitics are not considered for vertical Rowhammer. However, the inter-layer bonding (μ-bumps) will lead to vertical parasitics. We follow the dimensions of our μ-bumps from a real 3D-DRAM design [6]: diameter is 20 μm, height is 10 μm, and pitch is 25 μm. In addition, we assume that the material for the μ-bumps is Cu_3Sn as in [6]. To characterize vertical Rowhammer, we utilize the bonding model from [6] to represent the parasitic from layer n-layer $n \pm 1$. We utilize this analysis to establish (3) and (4) [6]. These equations represent the resistance and capacitance of the bonding μ-bumps between chip layers, respectively. We utilize the following parameter values: $\rho_{Cu_3Sn} = 8.3 \times 10^{-8}$ Ωm [16], μ-bump height is H, μ-bump width is W, and μ-bump pitch is D

$$R_{\mu\text{-bump}} = \rho_{Cu_3Sn} H \Big/ \left(\pi\,(W/2)^2\right) \tag{3}$$

$$C_{\mu\text{-bump}} = \epsilon_{SiO_2} H W / D. \tag{4}$$

MIVs offer greater packaging density than TSVs for HI-orientated DRAM [16]. However, since MIVs are still an emerging technology, we utilize the current analysis of real-silicon MIVs for our MIV-based 3D-DRAM simulation [21]. The MIV dimensions given in [21] inform our simulated MIV-based 3D-DRAM (117 nm diameter and 1795 nm height). In our MIV-based 3D-DRAM simulations, we assume similar architecture and performance optimizations as described previously (Figs. 1 and 3). Note, that thermal issues are common for TSV- and MIV-based packages [14]. To capture the impact temperature has on Rowhammer in 3D-DRAM we sweep temperature according to [8] and [22]. In our simulation setup, we consider parasitic from horizontal and vertical interconnect geometry, process variation, voltage variation, temperature variation, layout optimizations, DTC SPICE model, DRAM SA analog circuits, as well as parasitics from HI packaging architectures. Overall, we simulate an aggressor cell's impact on a physically adjacent victim cell (horizontal Rowhammer) and on vertically packaged adjacent victim cells (vertical Rowhammer). An example of our SPICE simulation setup (for the TSV 3D-DRAM case) is shown in Fig. 4. Note that vertical Rowhammer for TSV-based 3D-DRAM integration considers the μ-bumps in conjunction with vertical crosstalk. Vertical Rowhammer for MIV-based 3D-DRAM integration considers vertical crosstalk and fringe capacitance coupling from relative WL and MIVs [12], [15], [21].

A Rowhammer attack may be conducted in different ways [23]. A single-sided Rowhammer occurs when a single aggressor row is accessed HC_{first} times [23]. A double-sided Rowhammer attack is when two aggressor rows access both sides of the victim [23]. Each aggressor row is access $HC_{first}/2$ times. These attack methods are used in Rowhammer frameworks such as TRRespass and BLACKSMITH [23]. TRRespass interlaces single-sided Rowhammer to attack multiple victim rows simultaneously [23]. BLACKSMITH launches aggressor access in a semi-random fashion to achieve Rowhammer [23]. Another attack, half-double, occurs when an attacker attempts to flip a bit (at row r) through a low-frequency (row $r+1$) and high-frequency (row $r+2$) aggressor access [23]. Feinting is a recent Rowhammer method that uses decoy aggressor access to build up enough access across to cause Rowhammer [3], [23].

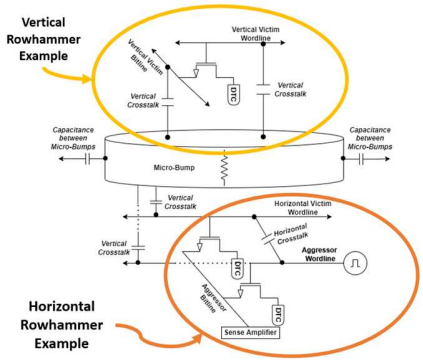

Fig. 4. Illustration of Rowhammer SPICE Simulation for TSV-based 3D-DRAM.

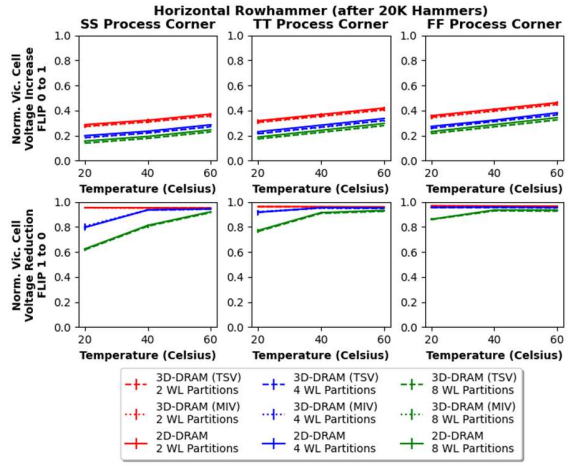

(a)

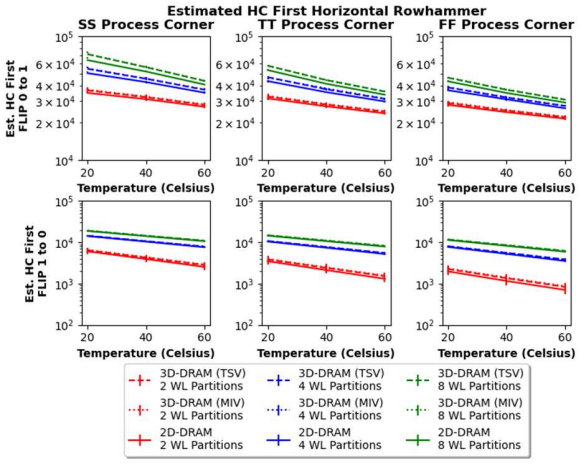

(b)

Fig. 5. Horizontal Rowhammer vulnerability. (a) Normalized victim cell voltage. (b) Impact of wordline partitions to HC_{first}.

While the distribution of access may differ for each attack method, all discussed attack methods require frequent access to the same aggressor row to achieve Rowhammer [23]. Thus, we utilize the single-sided attack to act as the Rowhammer base case. In future work, we aim to address other attacks and understand their relationship to HC_{first}.

IV. Results and Analysis

We show results for two sets of simulations. First, we show the impact of WL partitioning on horizontal Rowhammer. We present simulation results for vertical Rowhammer with crosstalk enabled by the HI bonding parasitics. Fig. 5 shows the results for the first set of simulations. Note, when simulating horizontal Rowhammer we

consider 2-D (i.e., DDR4) and 3-D packaging (i.e., TSV, MIVs) of DRAM. To ensure that we are simulating a bit flip from 0 to 1, we utilize a continuous "read 1" memory access pattern to the aggressor cell (the victim cell is initialized to 0 V). By executing a "read 1" memory access, we cause parasitic to impact the victim cell through the WL and BL interconnects, thereby inducing higher voltage on the victim cell. To simulate the flipping of a bit from 1 to 0, we execute a "read 0" memory access (the victim cell is initialized to V_{DD}). Hammering the victim WL enables subthreshold leakage, which expedites DRAM cell leakage. In addition, during a "read 0" memory access, the SA pulls the BL down to 0 V. This causes a voltage potential difference between the victim cell with a value of 1 (V_{DD}) and the BL (0 V). This voltage difference in combination with the hammered WL enables subthreshold leakage to accelerate the inherent DRAM leakage, thus resulting in a DRAM cell bit flip. Fig. 5 considers the previously described WL partitioning (2, 4, 8). The red/blue/green line represents the victim cell voltage with the WL partitioned into 2/4/8 segments, respectively. More WL partitions imply smaller interconnects and thus, lower parasitics that help reduce the Rowhammer vulnerability (Fig. 5). However, these partitioned interconnects require more local peripheral circuits to support DRAM memory accesses (e.g., more area usage because of local WL support circuits, etc.) [10]. The dotted, dashed, and solid line represents the victim cell voltage with 3-D TSV-based packaging, 3-D MIV-based packaging, and 2-D packaging. Note that the resultant impact of horizontal Rowhammer stayed consistent across all packaging schemes considered. For conciseness, we show the resultant normalized victim cell voltage changes after 1 ms of hammering (20 K hammers), and we utilize error bars to show the impact voltage variation has on the Rowhammer vulnerability. Note that flipping a bit (either 0 to 1 or 1 to 0) is a 0.5 victim cell voltage difference to the victim cell [Fig. 5(a)]. For flipping a bit from 1 to 0, we were able to cause bit flips across all considered test cases within 1 ms (20 K hammers). Regardless of packaging architecture (i.e., 2-D and 3-D), horizontal Rowhammer impact remained consistent. We show the relative HC_{first} for all considered test cases for flipping a bit 1 to 0 [Fig. 5(b)]. For flipping a bit 1 to 0, the estimated HC_{first} for Horizontal Rowhammer ranges from ~19 K hammers (8 WL partitions, SS process corner, 20 °C, $V_{DD} - 60$ mV) to ~500 (2 WL partitions, FF process corner, 60 °C, $V_{DD} + 60$ mV). For flipping a bit 0 to 1, we approach the necessary voltage change to induce a bit flip at higher temperatures [Fig. 5(a)]. For flipping a bit from 0 to 1, we use the voltage change after 1 ms of hammering to estimate HC_{first} using a linear function. The estimated HC_{first} for Horizontal Rowhammer ranges from ~65 K (8 WL partitions, SS process corner, 20 °C, $V_{DD} - 60$ mV) to ~21 K (2 WL partitions, FF process corner, 60 °C, $V_{DD} + 60$ mV). We show that WL partitions affect HC_{first}. The HC_{first} for bit flip "1" to "0" reduces faster than bit flip "0" to "1." DRAM cells are designed to be small. This design choice makes them susceptible to leakage and interference. It is easier to cause bit flips through loss of charge (1 to 0) than the injection of charge (0 to 1). Our simulation results corroborate real DRAM studies [1].

With our TSV-based (MIV-based) 3D-DRAM simulation, HI bonding material (vertical crosstalk) is considered. We repeat the simulations as prior for flipping bits from 0 to 1, and vice versa. Instead of a physically adjacent victim cell, we consider cells in the layer below and above the aggressor row. Fig. 6 shows the vertical Rowhammer event results. In Fig. 6, we see some induced voltage change in the victim cell with both TSV- and MIV-based 3D-DRAM vertical Rowhammer. It is important to consider benign subthreshold leakage for these simulations. For DRAM memory operations, the BL is initialized to $V_{DD}/2$. Regardless of whether the victim cell is 0 or 1, there will be subthreshold leakage due to DRAM cells' intrinsic characteristics [2], [15]. We compare the victim cell voltage changes during vertical Rowhammer versus a nonhammered cell that experiences only subthreshold leakage. We find that a vertically Rowhammer-ed cell, when TSVs are considered, is like a nonhammered cell that just experiences subthreshold leakage [Fig. 6(a)]. We show evidence that in TSV-based HI DRAM packages, vertical Rowhammer is unlikely due to the considerable interlayer distance. When MIVs are considered, the vertically Rowhammer-ed victim cell shows a notable induced voltage difference compared to subthreshold leakage for all test cases for flipping a bit from 0 to 1 [Fig. 6(a)]. For the inverse case (bit flip 1 to 0), the vertically Rowhammer-ed victim cell shows similar results but only for the cases associated with higher temperatures and the FF process corner. DRAM cells fabricated in the FF process corner inherently have higher subthreshold leakage. This increase in leakage makes DRAM cells more susceptible to losing charge (i.e., bit-flips occur) when they are hammered (Figs. 5 and 6). Fig. 6(b) shows that vertical Rowhammer is feasible when MIV-based packaging is considered. We show the estimated HC_{first} for flipping a bit 0 to 1 for MIV-based vertical Rowhammer [Fig. 6(b)]. The estimated HC_{first} for MIV-based Vertical Rowhammer (bit flip 0 to 1) ranges from ~414 K hammers (8 WL partitions, SS process corner, 20 °C, $V_{DD} - 60$ mV) to ~86 K hammers (2 WL partitions, FF process corner, 60 °C $V_{DD} + 60$ mV). For bit flips 1 to 0 for MIV-based Vertical Rowhammer, the HC_{first} estimate ranges from ~193 to ~190 K hammers for all considered WL partitions and voltage variance (FF, 60 °C). MIV-based vertical Rowhammer is feasible for both bit flips scenarios at higher temperatures due to process variation.

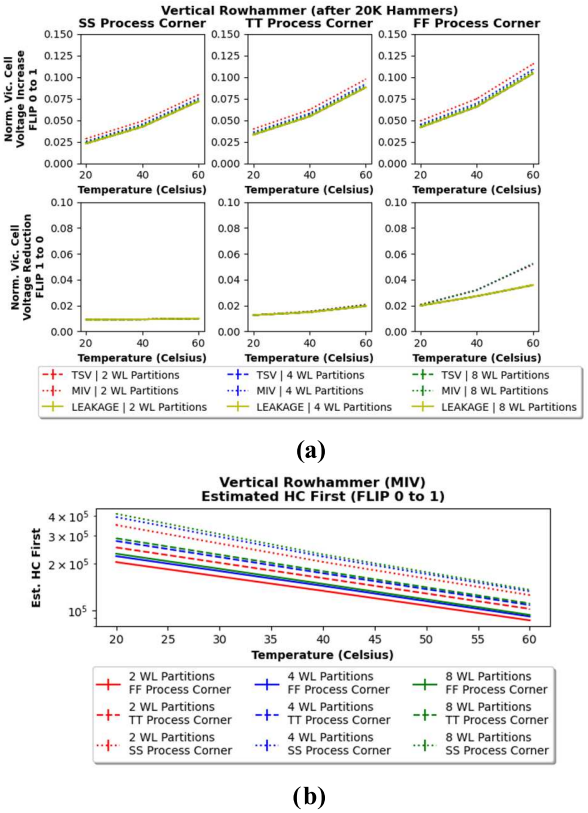

Fig. 6. Vertical Rowhammer in 3D-DRAMs. (a) Vertical Rowhammer versus leakage in 3D-DRAM. (b) Est. HC_{first} for vertical Rowhammer in 3D-DRAM.

V. Conclusion

We have studied Rowhammer vulnerability for heterogeneously integrated DRAM packages (i.e., 3D-DRAM). We have shown that

horizontal Rowhammer vulnerability may be reduced through known performance optimizations such as BL and WL interconnect partitioning in DRAM banks. We have also shown that vertical Rowhammer in TSV-based heterogeneously integrated DRAM is unlikely due to the larger distance between aggressor and victim layers. However, when MIV-based heterogeneous integration is considered, vertical Rowhammer is feasible due to vertical crosstalk, higher temperature, and process variation.

REFERENCES

[1] Y. Kim et al., "Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors," in *Proc. ACM/IEEE 41st Int. Symp. Comput. Archit. (ISCA)*, Jun. 2014, pp. 361–372.

[2] O. Mutlu and J. S. Kim, "RowHammer: A retrospective," *IEEE Trans. Comput.-Aided Design Integr. Circuits Syst.*, vol. 39, no. 8, pp. 1555–1571, Aug. 2020.

[3] M. Marazzi, P. Jattke, F. Solt, and K. Razavi, "ProTRR: Principled yet optimal in-DRAM target row refresh," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2022, pp. 735–753.

[4] A. Kwong, D. Genkin, D. Gruss, and Y. Yarom, "RAMBleed: Reading bits in memory without accessing them," in *Proc. IEEE Symp. Secur. Privacy (SP)*, May 2020, pp. 695–711.

[5] C. Evans. (2015). *Exploiting the DRAM Rowhammer Bug to Gain Kernel Privileges*. Accessed: Mar. 9, 2015. [Online]. Available: https://googleprojectzero.blogspot.com/2015/03/exploiting-dram-rowhammer-bug-to-gain.html

[6] *Analysis of Microbump Overheads for 2.5D Disintegrated Design, Technical Report CSE-TR-002-17*. Accessed: Sep. 15, 2023. [Online]. Available: https://www.eecs.umich.edu/techreports/cse/2017/CSE-TR-002-17.pdf

[7] J. S. Kim et al., "Revisiting RowHammer: An experimental analysis of modern DRAM devices and mitigation techniques," in *Proc. ACM/IEEE 47th Annu. Int. Symp. Comput. Archit. (ISCA)*, May 2020, pp. 638–651.

[8] Micron. (2021). *DDR4 SDRAM*. [Online]. Available: https://www.micron.com/-/media/client/global/documents/products/data-sheet/dram/ddr4/8gb_ddr4_sdram.pdf

[9] Y. Kim, V. Seshadri, D. Lee, J. Liu, and O. Mutlu, "A case for exploiting subarray-level parallelism (SALP) in DRAM," in *Proc. 39th Annu. Int. Symp. Comput. Archit. (ISCA)*, Jun. 2012, pp. 368–379.

[10] T. Zhang, K. Chen, C. Xu, G. Sun, T. Wang, and Y. Xie, "Half-DRAM: A high-bandwidth and low-power DRAM architecture from the rethinking of fine-grained activation," in *Proc. ACM/IEEE 41st Int. Symp. Comput. Archit. (ISCA)*, Jun. 2014, pp. 349–360.

[11] 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 *Proc. 54th Annu. IEEE/ACM Int. Symp. Microarchitecture*, Oct. 2021, pp. 1–13.

[12] L. T. Clark et al., "ASAP7: A 7-nm finFET predictive process design kit," *Microelectron. J.*, vol. 53, pp. 105–115, Jul. 2016.

[13] S. C. Shastri et al., "A distributed compact model for high-density, on-chip trench capacitors in high-frequency applications," in *Proc. NSTI Nanotech Conf*, 2005, pp. 119–122.

[14] T. Li and S. S. Sapatnekar, "Stress-aware performance evaluation of 3D-stacked wide I/O DRAMs," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Design (ICCAD)*, Nov. 2017, pp. 645–650.

[15] D. A. Hodges et al., "Additional topics in memory design," in *Analysis and Design of Digital Integrated Circuits: In Deep Submicron Technology*. New York, NY, USA: McGraw-Hill, 2005.

[16] C. C. Wei et al., "Electromigration in Sn–Cu intermetallic compounds," *J. Appl. Phys.*, vol. 105, no. 2, 2009, Art. no. 023715.

[17] T. Yang and X.-W. Lin, "Trap-assisted DRAM row hammer effect," *IEEE Electron Device Lett.*, vol. 40, no. 3, pp. 391–394, Mar. 2019.

[18] S.-W. Ryu et al., "Overcoming the reliability limitation in the ultimately scaled DRAM using silicon migration technique by hydrogen annealing," in *IEDM Tech. Dig.*, Dec. 2017, pp. 21.6.1–21.6.4.

[19] K. Park et al., "Experiments and root cause analysis for active precharge hammering fault in DDR3 SDRAM under 3× nm technology," *Microelectron. Rel.*, vol. 57, pp. 39–46, Feb. 2016.

[20] A. J. Walker, S. Lee, and D. Beery, "On DRAM rowhammer and the physics of insecurity," *IEEE Trans. Electron Devices*, vol. 68, no. 4, pp. 1400–1410, Apr. 2021.

[21] W. Wahby, A. Dembla, and M. Bakir, "Evaluation of 3DICs and fabrication of monolithic interlayer vias," in *Proc. IEEE Int. 3D Syst. Integr. Conf. (3DIC)*, Oct. 2013, pp. 1–6.

[22] F. Wang, T. Vogelsang, B. Haukness, and S. C. Magee, "DRAM retention at cryogenic temperatures," in *Proc. IEEE Int. Memory Workshop (IMW)*, May 2018, pp. 1–4.

[23] E. Ortega, T. Bletsch, B. Joardar, J. Talukdar, W. Paik, and K. Chakrabarty, "Simply-track-and-refresh: Efficient and scalable rowhammer mitigation," in *Proc. IEEE Int. Test Conf. (ITC)*, Oct. 2023, pp. 340–349.