Quantifying Rowhammer Vulnerability for DRAM Security

Yichen Jiang*§, Huifeng Zhu†§, Dean Sullivan*, Xiaolong Guo‡, Xuan Zhang† and Yier Jin*
*University of Florida, †Washington University in St. Louis, ‡Kansas State University
yichen.jiang@ufl.edu, zhuhuifeng@wustl.edu, deanms@ufl.edu, guoxiaolong@ksu.edu, xuan.zhang@wustl.edu, yier.jin@ece.ufl.edu

2021 58th ACM/IEEE Design Automation Conference (DAC) | 978-1-6654-3274-0/21/$31.00 ©2021 IEEE | DOI: 10.1109/DAC18074.2021.9586119

***Abstract*—Rowhammer is a memory-based attack that leverages capacitive-coupling to induce faults in modern dynamic random-access memory (DRAM). Over the last decade, a significant number of Rowhammer attacks have been presented to reveal that it is a severe security issue capable of causing privilege escalations, launching distributed denial-of-service (DDoS) attacks, and even runtime attack such as control flow hijacking. Moreover, the Rowhammer vulnerability has also been identified and validated in both cloud computing and data center environments, threatening data security and privacy at a large scale. Various solutions have been proposed to counter Rowhammer attacks but existing methods lack a circuit-level explanation of the capacitive-coupling phenomenon in modern DRAMs, the key cause of Rowhammer attacks.**

In this paper, we develop an analytical model of capacitive-coupling vulnerabilities in DRAMs. We thoroughly analyze all parameters in the mathematical model contributing to the Rowhammer vulnerability and quantify them through real DRAM measurements. We validate the model with different attributions on a wide range of DRAM brands from various manufacturers. Through our model we re-evaluate existing Rowhammer attacks on both DDR3 and DDR4 memory, including the recently developed TRRespass attack. Our analysis presents a new Rowhammer attack insight and will guide future research in this area.

I. Introduction

The Rowhammer attack was first proposed by Kim et al. [1] in 2014, in which it was demonstrated that a specially crafted workload could flip bits in DRAM memory cells without accessing them. As opposed to software-level attacks, the Rowhammer attack does not exploit software errors but rather low-level circuit side-effects. It was found that a parasitic capacitance could be induced to turn on the access transistor of victim cell by repeatedly accessing physically adjacent rows to a victim row. Researchers have used the Rowhammer vulnerability to launch advanced attacks such as privilege escalation [2], distributed denial-of-service (DDoS) attacks on Intel SGX [3], etc. The wide use of DRAM in computing systems makes Rowhammer attacks applicable to a wide range of devices and platforms. For example, the authors in [4] showed that the mobile phone was also vulnerable to such attacks. The authors in [5] and [6] exploited the Rowhammer attack in the cloud and high-performance computing systems.

Given the pervasive threat the Rowhammer attack poses to architecture security, various solutions have been proposed [7]–[9]. However, these countermeasures aim to prevent, not characterize, Rowhammer attacks. There lacks a comprehensive, quantitative circuit-model capable of providing insight into the susceptibility of a DRAM cell to Rowhammer. Statistical modelling has been introduced [10], but it is not efficient nor accurate when applied to a large set of DRAM chips despite massive amounts of measurement data.

In this paper, we will introduce an analytical model that simulates all parameters of the capacitive-coupling phenomenon responsible for inducing the Rowhammer attack at the circuit level. We abstract two new parameters, namely the *equivalent resistance of intrinsic leakage* and the *equivalent resistance of capacitive coupling leakage*, to quantitatively characterize a DRAM cell in response to the Rowhammer attack and evaluate cells' susceptibility of being leaked during such an attack. In our model, the *equivalent resistance of intrinsic leakage* (denoted as R_L) describes the retention time of the DRAM cells storage capacitor. The *equivalent resistance of capacitive coupling leakage* (denoted as R_{SW}) represents the DRAM cells resistance to repeated aggressor row activations. We experimentally determine these model parameters and how they can be used to characterize a DRAM's resilience against Rowhammer attacks. To the best of our knowledge, this is the first work to quantify the relationship between DRAM data retention and counts of aggressor row activations. We further show how the model helps re-examinate recent Rowhammer attacks and derives a new Rowhammer attack insight.

The main contributions of this paper are as follows.

- For the first time, we provide a quantitative model of the Rowhammer vulnerability at the granularity of a DRAM cell. We define two parameters, the *equivalent resistance of intrinsic leakage* and *equivalent resistance of capacitive coupling leakage* which allow us to use microarchitectural side-effects to accurately define an arbitrary DRAM cell's resilience to Rowhammer attacks. Indicated by the model, we show how the cell leakage time contributes to the Rowhammer attack.
- We derive two attributions from our model and use them to re-evaluate the state-of-the-art Rowhammer attacks, e.g., TRRespass attacks on DDR4 memory. New insights of Rowhammer attacks are derived guide by our model.

The rest of the paper is organized as follows: In Section II, we introduce background information related to our Rowhammer model. We introduce the abstracted circuit model and features in Section III. Section IV presents our model verification approach. Our re-examination of Rowhammer attacks is presented in Section V. We conclude our paper in Section VI.

II. Background

A. DRAM Circuit Structure

Figure 1(a) shows the circuit schematic of a DRAM chip composed mainly of three parts: the decoder, the row buffer, and the cell array. The schematic of *one* DRAM cell is illustrated in Figure 1(b). The decoder drives the cell array and reads or writes data to the corresponding DRAM cells with the aid of the row buffer. Each DRAM cell contains one capacitor that is connected to the bit line (BL) through an access transistor. The bit lines (also called digit line or column line) are arranged vertically and each line is shared by multiple rows of DRAM cells. The access transistors are controlled by a row decoder through word lines (also called row line), which is arranged horizontally and shared by columns.

When the word line is activated, the access transistor in Figure 1(b) is turned on, connecting the storage capacitor C_S to the bit line. The stored value is detected by the sense amplifier by comparing the charge stored on the capacitor with a reference voltage. In modern DRAM circuits, the charge stored on the capacitor is equal to $Q = 0$ or $Q = +V_{DD} \times C_S$. A high voltage is detected if $Q > +(V_{DD}/2) \times C_S$ while a low voltage is detected if

§The first two authors contributed equally to this work.

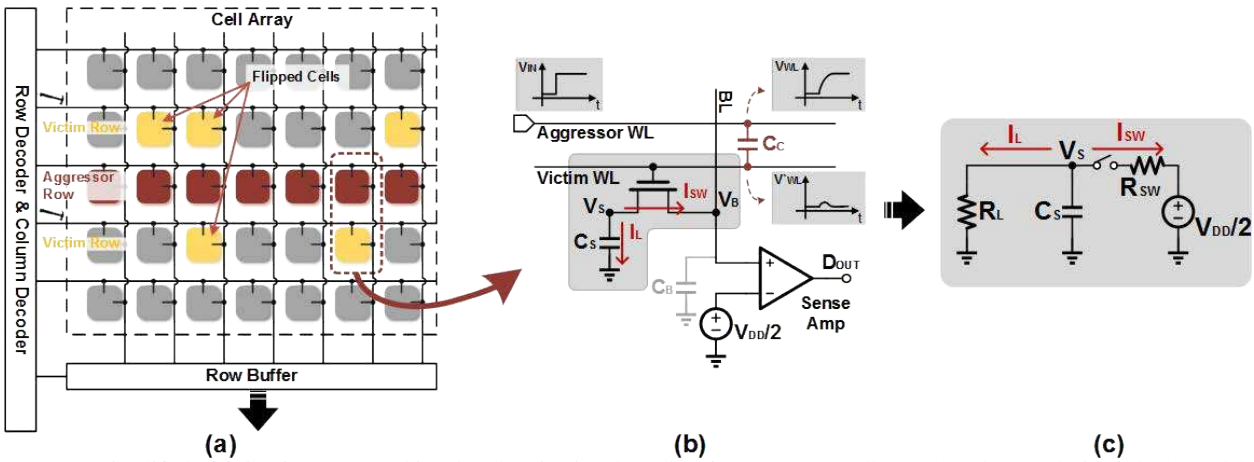

Figure 1: (a) Simplified circuit of a DRAM chip. (b) The circuit schematic of one DRAM cell. (c) The abstracted circuit-level model of the Rowhammer attack.

$Q < +(V_{DD}/2) \times C_S$.[1] The capacitor will slowly leak charge due to various reasons, but primarily its leakage current. Therefore, a periodic refresh operation is necessary to keep the stored data in the DRAM cell capacitor. The refresh interval is typically set to 64ms, i.e., all DRAM rows will be refreshed every 64ms in sequence.

B. Capacitive-Coupling Vulnerabilities

Capacitive-coupling vulnerabilities belongs to a newly developed charge-domain analog vulnerability [11]. This type of vulnerability utilizes analog behavior in low-level circuits to induce unexpected, often malicious, electrical charge transfers and/or redistribution for fault injections [12]. More specifically, the capacitive-coupling vulnerability relies on capacitive coupling effects associated with induced parasitic capacitors in digital circuits. The Rowhammer attack is one example of this type of vulnerability.

As shown in Figure 1(a) and (b), the matrix structure of DRAM cell arrays induces a parasitic capacitance between two long word lines which triggers the Rowhammer attack. When activating a specific word line, labeled the Aggressor WL in the figure, voltage fluctuations are induced on its adjacent word line, labeled the Victim WL in the figure, due to the parasitic coupling capacitance C_C between these two word lines. As a result, the access transistor(s) of the victim row can be partially opened to leak the charge stored in that cell. The leakage charge across the aggressor-to-victim wordline increases as the frequency of aggressor row activations increases. The victim cell's storage capacitor C_S will leak when the accumulated leakage is enough to turn on the victim cell's access transistor.

There are several ways to perform the Rowhammer attack against victim cells including single-sided, double-sided, and many-sided [1], [2], [13]. For the single-sided Rowhammer attack, the attacker reads one aggressor row at a high frequency to cause some bits of the neighboring victim rows to be flipped[2]. To further improve the attack speed, the double-sided Rowhammer attack was proposed which aggressively activates two adjacent aggressor rows of the victim row. The many-sided Rowhammer attack builds upon the previous techniques by using *many* aggressor rows to trigger Rowhammer and bypass a defense called Target Row Refresh implemented in latest DDR chips.

[1]It is also possible that the DRAM cell capacitor will be connected to $V_{DD}/2$ instead of ground. In this case, the stored charge in each cell will be $Q = \pm(V_{DD}/2) \times C_S$. Our framework covers both cases. Note that the real structure of commercial DRAMs is proprietary and undocumented.

[2]In real attack scenarios, the attacker needs to alternatively read from the aggressor row and another random row from the same bank to avoid the impact of the row buffer.

Parameter	Description
R_{SW}	Equivalent resistance of coupling leakage
R_L	Equivalent resistance of intrinsic leakage
I_{SW}	Coupling leakage current
I_L	Discharging current
V_{DD}	Power supply voltage
V_S	Voltage of the storage capacitor
C_S	Capacitance of the storage capacitor
N_{att}	Total toggling counts
t_I	Interval between the successive accessing

Table I: Parameters in the abstracted circuit model for a DRAM cell.

C. TRRespass Attack

One of Target Row Refresh (TRR) protection mechanism which monitors the access time of rows in one refresh interval is reported in [13]. If recorded access time of one row within 64ms is beyond the secure access number that stored in the Serial Presence Detect (SPD) of the DRAM, one or more extra refresh operation is generated to refresh the adjacent rows of that row. For this TRR mechanism, extra buffer (or sampler) inside the DRAM is required to store the information of the access time for the rows. Since the sampler size is limited, the row information recorded in the sampler is limited as well. If accessing a group of rows inside a refresh interval concurrently, it can overwhelm the sampler and bypass the TRR mechanism. Thus, the many-side Rowhammer (also called n-side Rowhammer where n denotes the number of the aggressive rows) is proposed. In many side row hammer attack, several aggressive rows are accessed in one refresh interval and the aggressive row and victim row follow the pattern as: AVAVAVA (A is for aggressive, and V is for victim). Due to the limited size of the sampler, some aggressive rows are not recorded in the sampler. Hence, the aggressive rows which is not recorded have the high possibility to flip the bits in their adjacent victim rows.

III. Rowhammer Circuit Model

A. Circuit-level Abstraction

The abstracted circuit model of the Rowhammer attack is shown in Figure 1(c), and the model parameters are listed in Table I.

Model Description. In our model, there are two dominant leakage current paths in a DRAM cell, the discharge current I_L due to the intrinsic leakage of the storage capacitor, and the coupling leakage current I_{SW} due to the partially closed access transistor during a Rowhammer attack. These two leakage paths are modeled as equivalent resistors, R_L and R_{SW}, respectively (see Table I). The intrinsic leakage, in the form of discharge current I_L, exists all the time. In fact, DRAM refresh is designed to compensate for the intrinsic leakage. The leakage caused by the access transistor is modeled as the equivalent resistor R_{SW} connected with a switch and a voltage source (the amplitude of the voltage source is $V_{DD}/2$ because the bit line is pre-charged to $V_{DD}/2$). When the Aggressor WL is activated (or charged), the access transistor of the Victim WL becomes partially closed for a short period of time due to the capacitive-coupling effects between the two word lines. To capture this effect in our model, the switch will be turned on for a short period of time with each activation of the Aggressor WL causing the charge to leak from the resistor R_{SW}. This state is called the aggressive situation.

Model Equations. Based on the above discussion, we derive the corresponding mathematical expressions describing the behavior under normal operation and during a Rowhammer attack. During normal operation of a DRAM cell, there only exists the intrinsic leakage and the charge of the storage capacitor leaks through the resistor R_L. Given the storage capacitor C_S, an RC discharging circuit's behavior can be expressed by the following equation.

$$V_S(t^1) = e^{-\frac{1}{R_L C_S}(t^1 - t^0)} V_S(t^0) \tag{1}$$

where $V_S(t^0)$ is the initial voltage of the storage capacitors.

In the aggressive situation, we need a more complex RC discharging circuit model since the charges leak from both R_L and R_{SW}. The corresponding equation can be presented as follows.

$$\frac{dV_S(t)}{dt} = \frac{1}{C_s}(I_L(t) + I_{SW}(t)) = -\frac{V_S(t)}{R' C_S} + \frac{V_{DD}/2}{R_{SW} C_S} \tag{2}$$

where $I_L(t) = -\frac{V_S(t)}{R_L}$, $I_{SW}(t) = -\frac{V_S(t) - V_{DD}/2}{R_{SW}}$, and $R' = R_L \parallel R_{SW}$. Given a differential equation with the form $\frac{dx(t)}{dt} = \lambda x(t) + f(t)$, the solution is $x(t) = e^{\lambda(t-t_0)}x(t_0) + \int_{t_0}^{t} e^{\lambda(t-\tau)} f(\tau) d\tau$. Using this we can derive the analytical expression for aggressive situation as follows.

$$V_S(t^2) = e^{-\frac{1}{R' C_S}(t^2 - t^1)} V_S(t^1) + \frac{V_{DD} R'}{2 R_{SW}} (1 - e^{-\frac{1}{R' C_S}(t^2 - t^1)}) \tag{3}$$

Rowhammer Attack Modeling. When the attacker launches the Rowhammer attack by repeatedly activating an aggressor row, the mode of charge leakage changes in both Equation (1) and Equation (3). Assume that the interval of two row activations is $t_I = (t^2 - t^1) + (t^1 - t^0) = B\Delta t + A\Delta t$, where the normal activation lasts for $A\Delta t$, and the aggressive activation lasts for $B\Delta t$. We can derive the storage capacitor voltage after one row activation by combining Equation (1) and Equation (3).

$$\begin{aligned} V_S(t^2) &= e^{-\frac{1}{R' C_S}(B\Delta t)} V_S(t^1) + \frac{V_{DD} R'}{2 R_{SW}} (1 - e^{-\frac{1}{R' C_S}(B\Delta t)}) \\ &= e^{-\frac{1}{R' C_S}(B\Delta t)} e^{-\frac{1}{R_L C_S}(A\Delta t)} V_S(t^0) \\ &+ \frac{V_{DD} R'}{2 R_{SW}} (1 - e^{-\frac{1}{R' C_S}(B\Delta t)}) \end{aligned} \tag{4}$$

The above equation can be viewed as a recursive formula $V_S(t_i) = f(V_S(t_{i-1}))$ where $V_S(t_i) := V_S(t^2)$ is the voltage after the i-th row activation and $V_S(t_{i-i}) := V_S(t^0)$ is the voltage after the $(i-1)$-th one. By iteratively applying Equation (4), we can derive the expression for V_S after N_{att} aggressor row activations.

$$\begin{aligned} V_S(t_N) &= e^{-N_{att} \frac{1}{C_S}(\frac{1}{R_L} + \frac{D}{R_{SW}}) t_I} V_{DD} \\ &+ \frac{V_{DD} R_L}{2(R_L + R_{SW})} (1 - e^{-N_{att} \frac{1}{C_S}(\frac{1}{R_L} + \frac{1}{R_{SW}}) D t_I}) \end{aligned} \tag{5}$$

where $t_N = \sum_{i=0}^{N_{att}} t_i^2 - t_i^0 = N_{att} \times t_I$, and $D = B/(A+B)$.

Evaluating Rowhammer Attack. $V_S(t_N)$ from Equation (5) represents the voltage of the victim cell under a Rowhammer attack. At any time, this voltage can be compared with the threshold of the sense amplifier ($V_{DD}/2$) to determine whether the charge on the storage capacitor C_S has leaked. We also incorporate the effects of charge sharing between the storage capacitor and the bitline in our comparison [14]. This is characterized with the following inequality whereby we decide whether a Rowhammer attack is successful or not.

$$\begin{cases} V_S(t_N) > \frac{C_S + C_B}{C_S} \times (\frac{V_{DD}}{2} + V_{SA}), & \text{unsuccessful attack} \\ V_S(t_N) < \frac{C_S + C_B}{C_S} \times (\frac{V_{DD}}{2} - V_{SA}), & \text{successful attack} \\ others, & \text{uncertain} \end{cases} \tag{6}$$

where C_B is the parasitic capacitance of the bit line, and V_{SA} is the resolution of the sense amplifier.

B. Model Analysis

In the capacitive-coupling model in Equation (5), N_{att} and t_I are parameters controlled by the attacker while V_{DD} and C_S are device-specific features which are available in the DRAM datasheet [15]. Therefore, in order to better understand the capacitive-coupling effect, we will mostly focus on the two abstracted resistance parameters, R_{SW} and R_L, in the model. Specifically, we address two attributes of the model (detailed justifications will be given in Section IV).

Attribute 1: *Both the leakage time of the cell and the activation time of the aggressive rows will effect Rowhammer attacks.*

Our model shows that the cell leakage time and activation time of aggressive row have influence to the Rowhammer attack and we use R_L and R_{SW} to quantity the effect respectively. In the experiments section, we will demonstrate how the leakage time contributes to the Rowhammer attack, contradicting to the argument from previous research that leakage time is not important to Rowhammer attacks.

Attribute 2: *If the cell's R_{SW} and R_L are fixed, the activation time for the adjacent aggressive row to induce bit flipping are determined.*

Previous research proves that R_L (the cell leakage time) is affected by different factors, e.g., temperature and data pattern [16]. However, if R_L and R_{SW} are fixed during the attack, as indicated by our model, the activation time of adjacent aggressive row to flip the cells is determined. As a result, our model have the conclusion that the aggressive activation time to induce bit flipping is fixed under the same R_{SW} and R_L regardless of the aggressive row access sequence.

IV. Model Verification

A. Experimental Setup

Our goal is to corroborate the abstracted Rowhammer circuit-level model against an arbitrary DRAM module. To this end, we developed experimental platforms with custom memory controllers that provide

Timings	Value	Unit	Description
tcke	5	ns	CKE mininum pulse width
tfaw	30	ns	Four Address Width
tras	35	ns	Active to Precharge command
trcd	13.75	ns	Active to Read or write delay
trefi	7.8	us	Average period refresh interval
trfc	110	ns	Refresh to Active/Refresh
trp	13.75	ns	Precharge command period
trrd	6	ns	Activate min. command period
trtp	7.5	ns	Read to Precharge delay
twtr	7.5	ns	Rank write to read delay

Table II: DRAM timing parameters used across all modules tested.

fine-grained control of DRAM address mapping, page policy, and refresh rate interval.

DDR3 Experimental Platform. We use the Xilinx Zynq-7000 ZC706 evaluation platform [17], a heterogeneous development platform containing a dual core ARM Cortex-A9 processor and reconfigurable logic on the same die. The ZC706 hosts a Kintex-7 FPGA with one DDR3 SODIMM slot. The time parameters for the experiment platform is listed in Table II. We run the SoC bare-metal, and experiment programs start executing in DRAM from the processor side.

DDR4 Experimental Platform. We use the Xilinx ZYNQ UltraScale+ ZCU104, a developmental platform containing a ARM Cortex-A53 and 16nm FinFET+ programmable logic (PL) [18]. The ZCU104 includes a 64-bit PL DDR4 SODIMM Connector. The testing process of DDR4 is the same as that in the DDR3 platform.

Memory Controller. We configure the FPGA as a DDR memory controller using the Xilinx 7 series memory interface generator [19]. We use the same parameters for each DRAM we evaluated to normalize the results. Physical addressing was configured in `row-bank-column` mode for single rank DDR. Bursts were configured to be handled sequentially. This memory controller uses a closed-page policy by default. When evaluating leakage times we disable DRAM cell refresh by configuring the `user_refresh` parameter to `TRUE`.

DRAM Chips. In table III, we list the DDR3 and DDR4 module we used for our experiment. The DRAM brands are from various manufacturers including Samsung, Kingston, Hynix, Micron, Axiom, Corsair, Crucial, TimeTec.

B. Computing R_L

In order to compute R_L for any given cell, we need to find the inherent leakage time of the cell[3]. Our analysis follows a similar procedure to prior research [16], outlined in Listing 1. For each DRAM, we iterate over every row within the DRAM. For each row, we perform the following steps.

- We activate the row and wait a `LEAKAGE_MIN` time;
- After waiting, the row is read back and checked for corruption;
- A cell that has leaked its charge is then reported to have a leakage time less than `LEAKAGE_MIN`.

We perform this operation for a `LEAKAGE_MAX` time. After the cell leakage time is measured, we take the value in Equation 1 to calculate the R_L.

[3]Note that the inherent leakage time is often referred to as DRAM cell's retention time in literature.

```
int find_leakage_time(target_addr) {
    /* Iterate over possible leakage times*/
    for (t = LEAKAGE_MIN -> LEAKAGE_MAX) {
        /* Set every cell in the target row */
        set_target_row(target_addr);
        sleep(t);
        /* Check any cell in the target row
        and exit if a cell has leaked */
        if (check_target_row(target_addr))
            return t;
}}
```

Listing 1 Pseudo-code used to find the leakage times within DRAM in computing R_L

C. Computing R_{SW}

In order to compute R_{SW} for a given cell in Equation 5, we need to find the activation time of the adjacent aggressor row to flip a bit in that cell. For a victim cell under observation, R_{SW} represents the likelihood of being discharged by repeated aggressor activation. That is, R_{SW} is the resistance to withstanding a Rowhammer attack. Our model provides a method for quantifying this resistance at a per cell granularity.

Listing 2 shows the pseudo-code used to find the number of aggressor row activations required to compute R_{SW}. For each DRAM, we iterate over every row within the DRAM. For each victim row we run a double-sided Rowhammer attack using the following steps.

- We first activate physically adjacent rows an `ACTVS_MIN` number of times;
- We then check the victim row for corruption;
- A cell that has flipped due to the attack is then reported to have activation count less of `ACTVS_MIN`.

```
int find_aggressor_count(victim_addr) {
    /* Set value in the victim row to 0xFF */
    set_target_row(victim_addr);
    /* Iteratively activate aggressor rows */
    for (num_actvs = ACTVS_MIN -> ACTVS_MAX){
        /* Double-sided Rowhammer attack */
        read_aggressor_row(victim_addr - 1);
        read_aggressor_row(victim_addr + 1);
        /* Check every cell in the victim row
        and break if a cell has leaked */
        if (check_victim_row(victim_addr))
            return num_actvs;
}}
```

Listing 2 Pseudo-code used to find the number of aggressor row activations for a DRAM cell in computing R_{SW}

D. Model Verification

In what follows, we verify the model with respect to the key model attributes presented in Section III.

Attribute 1: To verify Attribute 1, we take one 4GB HyperX DDR4 as the testing chip and calculate R_L and R_{SW} for all cells on this chip. We choose the cells with both R_{SW} and R_L value and separate them upon the different activation time of the aggressive row. The result is demonstrated in Figure 2 and Figure 3. From these figures, we found that both R_L and R_{SW} have certain range of value corresponding to each aggressive row activation time, e.g.,

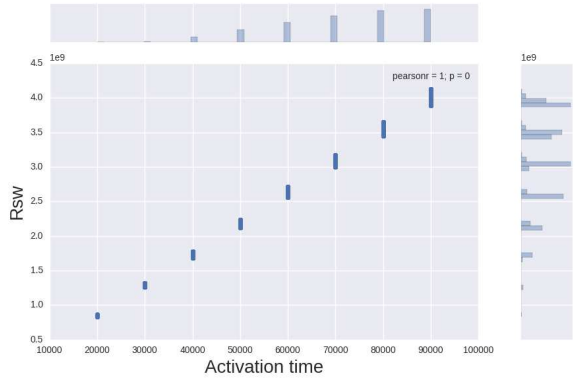

Figure 2: Activation time corresponding with R_{SW} on HyperX DDR4.

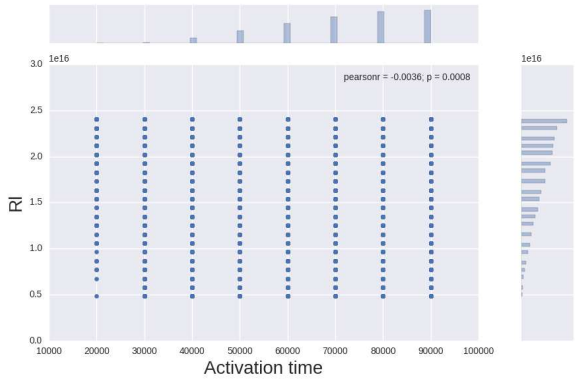

Figure 3: Activation time corresponding with R_L on HyperX DDR4.

R_{SW} varies from $3.4 * 10^9$ Ohm to $3.6 * 10^9$ Ohm and R_L ranges from $0.5{*}10^{16}$ Ohm to $2.5{*}10^{16}$ Ohm for 80,000 activation time. To investigate if the phenomenon is common situation across all DRAM modules. Next, we repeat the measurement for the different DRAMs listed in the Table III, and the range of R_L (the third column in Table III) and the range of R_{SW} (the fourth column in Table III) are recorded under each minimum activation time for flipping bits. The varied R_{SW} for all DRAMs demonstrate that the flipping cell has the varied resistance against the aggressive activation operation even under the same activation time. Since the higher resistance (the large value of the R_{SW}) of the cell obtained against the activation operation, a lower R_L requires for the cell to flip under the same activation time. Based on the model where the R_L describes the leakage time of the cell, we prove that the leakage time contributes to the Rowhammer attack along with the activating operation. It also delivers the result that for cells that are flippable, reducing either R_L or R_{SW} makes it more vulnerable.

Furthermore, we investigate the value of R_{SW} and R_L to find out if any correlation between these two parameters, e.g., the lower R_{SW} always has the lower R_L value. In Figure 4, we demonstrate the results for the determined R_{SW} and its corresponding R_L value from a HyperX DDR4 memory. From the figure, we notice that the R_L always has a large range for any R_{SW}. Thus, we conclude that the value of R_{SW} and R_L do not have observable correlations. Our finding also matches prior research that not always the cell with smallest leakage time is the most vulnerable cell to the Rowhammer attack [10]. Thus, we are confident that Attribute 1 is valid.

Attribute 2. For Attribute 2, we need to make sure that both R_L and R_{SW} are fixed. Thus, we use the same testing configuration and keep same temperature in each experiment. A random sequence

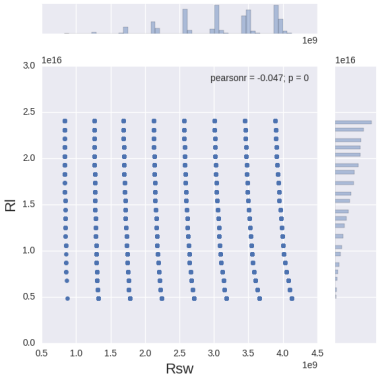

Figure 4: R_{SW} compared to R_L on HyperX DDR4.

DRAM	Type	R_L ($\times 10^{16}$Ohm)	R_{SW} ($\times 10^{9}$Ohm)
Axiom_1	DDR3	$6.24 - 8.17$	$81.4 - 83.9$
Corsair_1	DDR3	$2.40 - 9.13$	$51.9 - 58.9$
Corsair_2	DDR3	$1.43 - 7.69$	$52.9 - 62.9$
Crucial_1	DDR3	$5.76 - 9.13$	$95.4 - 99.9$
Hynix_1	DDR3	$1.92 - 7.69$	$81.9 - 95.9$
Hynix_2	DDR3	$3.84 - 5.76$	$83.9 - 87.9$
Kingston_1	DDR3	$1.43 - 9.61$	$155.0 - 202.0$
Kingston_2	DDR3	$1.43 - 8.65$	$157.0 - 202.0$
Micron_1	DDR3	$4.79 - 8.65$	$157.0 - 256.0$
Samsung_1	DDR3	$3.36 - 9.61$	$155.0 - 176.0$
Samsung_2	DDR3	$4.80 - 8.17$	$127.0 - 134.0$
TimeTec_1	DDR4	$1.82 - 2.40$	$1.68 - 1.69$
HyperX_1	DDR4	$0.48 - 2.40$	$0.82 - 0.86$

Table III: Tested DRAM modules

of aggressive row accessing is required. However, in double-side Rowhammer, it does not allow to change the memory access sequence since two aggressive rows require to access alternately to flush the row buffer. Hence, we insert a random number of interference row (the interference row is the row which stays in the same bank as the victim rows but different from aggressive row) between aggressive rows accessing to change the access sequence. We take an 8GB TimeTec DDR4 as the testing chip and record the cells addresses under different activation time within 128MB memory space. Next, we run the interference row inserted double-side Rowhammer attack to check if the different sequence of memory access influence the bit to flip. In Table IV, we shows the results for 20,000, 30,000 and 40,000 activation time with 5%, 10% interference rows inserted respectively. With the same bit flipping number against the different percent of interference row inserted, the results prove that the access sequence of aggressive row do not effect the bit to flip. We repeat the same experiment on all other DRAM chips and get similar results, a proof of the Attribute 2.

V. New Rowhammer Insights

A. Rowhammer on DDR4 Memory

In [13], many-side Rowhammer (aka n-side Rowhammer) is proposed to bypass the TRR mechanism and induce the bit flipping in DDR4 memory. Indicated by our model, the toggling count is determined if the certain parameters are given. That is, the activation time to induce the bit flipping should be same under the same platform configuration and the same testing environment. Thus, the n-side Rowhammer attack should have n aggressor rows to successfully

Activation time	Interference row percent	Bit flipping
20,000	5%	0
20,000	10%	0
30,000	5%	132
30,000	10%	132
40,000	5%	1028
40,000	10%	1028

Table IV: Different percent of the interference row inserted Rowhammer

3-side attack	4-side attack	5-side attack	6-side attack
1604	1594	1587	1590

Table V: many-side Rowhammer attack under varied aggressive row number

flip the bits if the TRR is bypassed. Specifically, if an n_1-side attack ($n_1 > 2$) can successful flip the bit, then for any n_2-side attack ($n_2 > n_1$) should also be successful to induce the bit flipping if the same activation time can be applied.

We repeat the n-side Rowhammer attack on various DDR4 memory to show the correctness of our findings. In the experiment, we use same experiment platform mentioned in experiment section and the 8GB TimeTec DDR4 is used as the testing memory chip. We first select the cells which is vulnerable to the 3-side Rowhammer attack with 50,000 activation time within 20MB memory space. We then increase the n-side attack from 3-side to 6-side but keep the activation time fixed to 50,000 for the aggressive row. In Table V, we demonstrate that the flipping bit count for different n-side Rowhammer attack. The result shows only a sightly difference between each experiments. The result matches that in [13] where they reported module A10 has the varied n from 4 to 32 to successful inducing the bit flipping.

B. New Rowhammer Attack Insights

m-gap Rowhammer Attack: *Relatively low R_L and R_{SW} in DDR4 memory may cause a new m-gap Rowhammer attack.*

Note that the new m-gap attack is totally different from the previous n-side attack. In n-side attack, n describes the total amount of aggressive rows. While in m-gap attack, only two aggressive rows are used and m denotes how many rows between each aggressive row and the victim row.

Our model shows that a lower R_{SW} and R_L makes the Rowhammer attack easier. Compared to DDR3, the R_{SW} of DDR4, shown in table III, is much smaller. Suggested by our model, we assume that DDR4 may be vulnerable to the m-gap Rowhammer attack. To validate that the new m-gap Rowhammer attack is possible, we perform the attack on 256MB memory space in a TimeTec DDR4 chip under 75 Celsius degrees (the high temperature will further reduce the R_L [16]). The auto-refresh is turned off during the experiment. Assume v_1 is the victim row number and m_1 is the number of rows between the aggressive row and the victim row. So the $v_1 + m_1$ and $v_1 - m_1$ rows are the aggressive rows for the m-gap Rowhammer attack. With $m_1 = 2$, bit flipping was successfully detected for the TimeTec DDR4 under 2-gap Rowhammer attack.

VI. Conclusion

In this paper, a Rowhammer vulnerability model is developed to evaluate the modern DRAM chips against Rowhammer attacks. For the first time, the equivalent resistances of coupling leakage and intrinsic leakage are defined at the circuit level to explain and evaluate the Rowhammer attack. We further show, for the first time, how the leakage time and activation time contribute to the Rowhammer attack separately and how our model quantify each of them. We validate our findings on a series of DDR3 and DDR4 from various manufacturers. Guide by our model, we present that our model's capability to help better understand the fundamental of the vulnerability of DDR memory. In the future, we will focus on developing countermeasures to protect DDR memory, especially the new DDR5 memory.

References

[1] Y. Kim, R. Daly, J. Kim, C. Fallin, J. H. Lee, D. Lee, C. Wilkerson, K. Lai, and O. Mutlu, "Flipping bits in memory without accessing them: An experimental study of dram disturbance errors," in *ACM SIGARCH Computer Architecture News*, vol. 42, no. 3. IEEE Press, 2014, pp. 361–372.

[2] M. Seaborn and T. Dullien, "Exploiting the dram rowhammer bug to gain kernel privileges," *Black Hat*, vol. 15, 2015.

[3] Y. Jang, J. Lee, S. Lee, and T. Kim, "Sgx-bomb: Locking down the processor via rowhammer attack," in *Proceedings of the 2nd Workshop on System Software for Trusted Execution*. ACM, 2017, p. 5.

[4] V. Van Der Veen, Y. Fratantonio, M. Lindorfer, D. Gruss, C. Maurice, G. Vigna, H. Bos, K. Razavi, and C. Giuffrida, "Drammer: Deterministic rowhammer attacks on mobile platforms," in *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*. ACM, 2016, pp. 1675–1689.

[5] Y. Xiao, X. Zhang, Y. Zhang, and R. Teodorescu, "One bit flips, one cloud flops: Cross-vm row hammer attacks and privilege escalation," in *25th {USENIX} Security Symposium ({USENIX} Security 16)*, 2016, pp. 19–35.

[6] V. Sridharan, J. Stearley, N. DeBardeleben, S. Blanchard, and S. Gurumurthi, "Feng shui of supercomputer memory positional effects in dram and sram faults," in *SC'13: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*. IEEE, 2013, pp. 1–11.

[7] Z. B. Aweke, S. F. Yitbarek, R. Qiao, R. Das, M. Hicks, Y. Oren, and T. Austin, "Anvil: Software-based protection against next-generation rowhammer attacks," *ACM SIGPLAN Notices*, vol. 51, no. 4, pp. 743–755, 2016.

[8] F. Brasser, L. Davi, D. Gens, C. Liebchen, and A.-R. Sadeghi, "Can't touch this: Software-only mitigation against rowhammer attacks targeting kernel memory," in *USENIX Security Symposium*, 2017.

[9] D. Gruss, M. Lipp, M. Schwarz, D. Genkin, J. Juffinger, S. O'Connell, W. Schoechl, and Y. Yarom, "Another flip in the wall of rowhammer defenses," in *2018 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2018, pp. 245–261.

[10] K. Park, D. Yun, and S. Baeg, "Statistical distributions of row-hammering induced failures in ddr3 components," *Microelectronics Reliability*, vol. 67, pp. 143–149, 2016.

[11] X. Guo, H. Zhu, Y. Jin, and X. Zhang, "When capacitors attack: Formal method driven design and detection of charge-domain trojans," in *2019 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 2019, pp. 1727–1732.

[12] K. Yang, M. Hicks, Q. Dong, T. Austin, and D. Sylvester, "A2: Analog malicious hardware," in *Security and Privacy (SP), 2016 IEEE Symposium on*. IEEE, 2016, pp. 18–37.

[13] P. Frigo, E. Vannacci, H. Hassan, V. van der Veen, O. Mutlu, C. Giuffrida, H. Bos, and K. Razavi, "Trrespass: Exploiting the many sides of target row refresh," *arXiv preprint arXiv:2004.01807*, 2020.

[14] Y. Li, H. Schneider, F. Schnabel, R. Thewes, and D. Schmitt-Landsiedel, "Dram yield analysis and optimization by a statistical design approach," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 58, no. 12, pp. 2906–2918, 2011.

[15] T. Instruments, "Ddr3 design requirements for keystone devices," 2014.

[16] J. Liu, B. Jaiyen, Y. Kim, C. Wilkerson, and O. Mutlu, "An experimental study of data retention behavior in modern dram devices: Implications for retention time profiling mechanisms," *ACM SIGARCH Computer Architecture News*, vol. 41, no. 3, pp. 60–71, 2013.

[17] Xilinx, "Zc706 evaluation board for the zynz-7000 xc7z045 soc, user guide," 2019.

[18] ——, "Vitis software platform: Embedded vision reference platforms user guide 2019.2 (ug1265)," 2019.

[19] ——, "7 series fpgas memory interface solutions: User guide," 2012.