

Quantifying Memory Cells Vulnerability for DRAM Security

Zilong Hu*, Hongming Fei*✉, Prosanta Gope†, Jack Miskelly‡, Owen Millwood§, Biplab Sikdar*

*National University of Singapore, Singapore
{huzilong, bsikdar}@nus.edu.sg, fei.hongming@u.nus.edu

†University of Sheffield, United Kingdom
p.gope@sheffield.ac.uk

‡Queen’s University Belfast, United Kingdom
J.Miskelly@qub.ac.uk

§Wiznet, Germany
owen.millwood@wiznet.eu

Abstract—Dynamic Random Access Memory (DRAM) is pervasive in computer systems. Cell vulnerabilities caused by unintended phenomena (forced retention failure, latency alteration, Rowhammer and Rowpress) lead to unintended bit flips in memory. These phenomena have been explored as attacks to violate data integrity and confidentiality during normal operation, but also exploited as a benefit in security systems as a method to generate random secret keys and unique device fingerprints (e.g. Physically Unclonable Functions). In both cases, attackers may wish to exploit knowledge of individual cell flip vulnerability to predict the current/future data contents of a set of cells, which can be utilised to break security systems. In this work, we develop a quantitative, cell-level circuit framework that models DRAM vulnerability directly from its physical charge leakage and disturbance pathways. By linking these device-layer behaviours to system-level security properties, our framework enables systematic evaluation of DRAM with respect to volatility (retention), integrity (disturbance-induced modification), and confidentiality (pattern-dependent leakage). We further demonstrate how the framework can be applied to well-known failure modes, revealing non-uniform and context-dependent vulnerability patterns. This work provides both theoretical foundations and practical evaluation tools for evaluating the suitability of DRAM use within security applications.

1. Introduction

Dynamic Random Access Memory (DRAM) is widely used as primary memory, from embedded IoT nodes to large-scale servers, and underpins almost every security-critical computation. During program execution, sensitive operating system metadata, cryptographic keys, user credentials, and privacy-relevant intermediate states inevitably reside in DRAM. The security posture of the overall system, therefore, depends not only on access-control policy and software isolation, but also on the physical behaviour of the memory array that stores these values. Physical, cell-level phenomena related to charge

properties can have significant impacts on the security of information in systems that use DRAM. Without an understanding of and ability to model those effects, any security analysis of such systems is incomplete. When it comes to these cell-level effects, two mechanisms are particularly influential. First, *retention loss* describes the natural decay of stored charge in a DRAM capacitor; without timely refresh, the storage node voltage drifts toward the sense threshold, eventually leading to a bit flip. Second, *read disturbance*—popularly known as Rowhammer and Rowpress [1], [2]—describes the effect where repeated or sustainable activations of one or more rows can induce bit flips in neighbouring rows even in the absence of write permissions. Both mechanisms have been observed in real systems and weaponized in attacks [3]–[5], however, both the security community and DRAM vendors mostly treat these failures empirically: attacks are tuned to device-specific behaviours, and mitigations are evaluated by ad-hoc testing. A comprehensive understanding of the cause is lacking, largely due to the absence of an **attacker-centric, quantitative model** for DRAM that links measurable cell-level parameters to concrete security properties (integrity, confidentiality, volatility).

To fulfill the gap and give a comprehensive analysis, we proposed a security-oriented framework: it aggregates physical leakage mechanisms into a small set of attacker-interpretable parameters (R_S , R_B , N), and connects these parameters to attack objectives such as precise bit flip control and pattern-based key inference to reveal the essence of the DRAM attacks from a theoretical to empirical way. This lets us reason about which cells and patterns are exploitable, rather than just how often errors occur. Moreover, we evaluate the **existing** attack surfaces systematically, then provide insights into potential attack surfaces.

1.1. Related Work

In recent years, the vulnerability of DRAM cells has been explored through empirical studies. Liu et al. [6] showed that data patterns significantly affect retention-induced bitflips due to parasitic coupling between adjacent wordlines and bitlines. Building on this, the Rowhammer

✉ Corresponding author

effect [1] demonstrated that repeated activation of aggressor rows can externally amplify cell fragility through disturbance-induced charge leakage, where the number of bitflips also depends on the aggressor-victim data patterns. *CheckerBoard* (i.e., alternating 010 between columns) is usually the most efficient pattern, while *Stripe* (same value for the entire row, such as a consistent 111 pattern) becomes weaker. Yağlıkçı et al. [7] demonstrate that tampering with wordline voltages can alter the rowhammer behaviour and increase vulnerability. Orosa et al. [8] demonstrate that the Rowhammer effect is more likely to occur when aggressor rows remain active for a longer duration. Furthermore, the RowPress effect [2] revealed that even static but prolonged wordline activation can amplify the vulnerability of the memory cells. Luo et al. [9] also reported inconsistency between previously proposed error mechanisms and pattern dependency effects with empirical results from 96 COTS DDR4 DRAM chips. Together, these mechanisms reveal that DRAM cell fragility arises not only from intrinsic retention loss but also from pattern-dependent disturbance and field-induced leakage for which modeling capabilities are currently limited.

While existing studies provide valuable empirical observations on cells’ vulnerability, many researchers [3]–[5], [10], [11] have also made use of these observations to perform real attacks. The most classic example is the Cold Boot Attack [10] which exploits the fact that DRAM, while nominally volatile, has a certain time before actual physical retention failure. That time is highly dependent on temperature, so rapidly cooling the DRAM after powering off allows partial reconstruction of sensitive data such as disk encryption keys once power is restored. This is an example of a *volatility* attack. Lin et al. [3] show that Rowhammer extends to discrete GPUs with GDDR6, where induced bit-flips can severely degrade machine learning model *integrity*, underscoring its system-level security impact. Lucian et al. [4] show that Rowhammer can be used to flip the targeted bits and keep the other bits unchanged by controlling the access patterns, another integrity attack. In addition to this kind of integrity violation, Kwong et al. [5] and Rakin et al. [12] further show that Rowhammer can be used in a *confidentiality* attack by observing flip patterns in attacker-controlled memory to infer victim data (e.g., cryptographic keys). These physical-level effects touch on all security properties of DRAM.

Jiang et al. [13] attempted to quantify Rowhammer at the circuit level by modelling the leakage mechanism. This study shows that when a victim cell’s storage resistance (R_{SW}) and leakage resistance (R_L) are fixed, the activation time of the adjacent aggressor row directly determines the likelihood of disturbance-induced bitflips. This provides a circuit-parameter-centric perspective on how Rowhammer arises from cell-level electrical constraints. On the other hand, Walker et al. [14] derive a mathematical and physics-based formulation of Rowhammer, offering a more analytical expression for how repeated activations couple into victim cells to induce charge leakage. Together, these two works illustrate complementary approaches: while Jiang et al. [13] emphasize empirical quantification through circuit parameters, Walker et al. [14] highlight analytical modeling grounded in physical

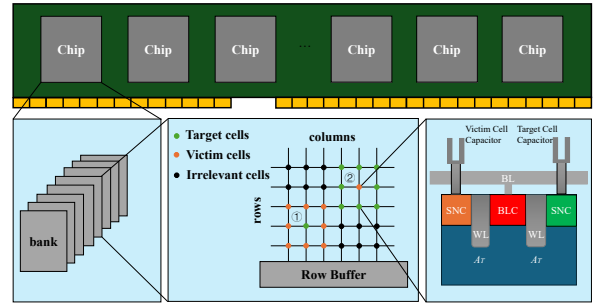


Figure 1: An overview of DRAM organization. Independent cell capacitors share substrate, BL, and WL, inducing pattern-dependent coupling. Scenario ① and ② represent two attack case studies as discussed in Section 3.

equations. Both perspectives converge to deepen our understanding of Rowhammer as a phenomenon governed by the joint properties of electrical devices and activation-induced disturbance physics.

1.2. Motivation

In summary, prior work has either (i) empirically demonstrated that DRAM failures can be weaponized in a number of attack surfaces, or (ii) built physical models to explain these failures, mostly for reliability purposes. To the best of our knowledge, no prior work has combined a cell-level physical model with a systematic, attacker-centric security evaluation, nor demonstrated that the underlying parameters can be reliably measured on commodity devices and used to compare the impact of different failure mechanisms on attack surfaces. This is precisely the gap our work aims to address.

To bridge this gap, we focus on a key phenomenon that naturally links physical cell behaviour with attacker capabilities: **pattern awareness**. Although adjacent DRAM cells store charge in independent capacitors (Figure 1), they are not isolated from each other. They share the same p-type substrate, bitlines (BL), and wordlines (WL), which introduce capacitive and subthreshold coupling. As a result, a cell’s flip probability is determined not only by its own device parameters and intrinsic leakage, but also by the *surrounding data pattern* encoded in its neighborhood.

From an attacker’s perspective, this pattern awareness immediately gives rise to two distinct classes of attack surfaces: (i) *pattern control*, where the adversary deliberately **sculpts** neighbour data so as to amplify or suppress a victim cell’s flip probability (targeted integrity violation); and (ii) *pattern inference*, where the adversary **measures** flip statistics on carefully chosen probe cells in order to infer unknown neighbor data (confidentiality violation). These behaviors have been repeatedly observed in empirical studies, but in the absence of a quantitative, cell-level model, it remains unclear how to systematically evaluate which cells, patterns, and devices are exploitable, or how different failure mechanisms (e.g., retention, Rowhammer, Rowpress) compare on the same footing.

This leads to a broad central question: From the attacker’s perspective, can we turn measurable cell-level DRAM parameters into quantitative indicators of how

exploitable a given device is for volatility-, integrity-, and confidentiality-based attacks? More specifically, we aim to answer the following research questions:

- RQ1:** How can we formalize a physical DRAM instance as an explainable model with well-defined, measurable parameters from the *attackers' perspective*?
- RQ2:** Can our model be exploited to accurately qualify or evaluate the performance of integrity and confidentiality attacks on the target device?
- RQ3:** If so, can the model provide new insights to existing attack surfaces not seen in empirical analysis?

1.3. Our Contributions

To answer the above research questions and bridge the gap between empirical DRAM characterization and *formal* security evaluation, we present a **model-driven, cell-level framework** that explains and quantifies DRAM physical behaviours from a security perspective. In Section 2, we introduce the background knowledge. After illustrating the threat model in Section 2.2, we present the proposed memory circuit model, two case studies and detailed attack qualification in Section 3 and make the following key contributions:

- 1 To answer the **RQ1**, We develop the *first* empirically verified cell circuit-level model that jointly characterizes DRAM volatility and read disturbance to **qualify** attacks from a *security* perspective. This model yields closed-form conditions for when a given cell is exploitable for *arbitrary and targeted bit flips* or *pattern inference*, as a function of R_S and R_B .
- 2 To answer the **RQ2**, we collect measurable electrical parameters on our FPGA-based DRAM testing platform (Xilinx Alveo U200 + DRAM Bender), and translate them into security-relevant objectives.¹ We show that these parameters can be robustly estimated on seven commodity DDR4 modules, and we characterize how exploitable vulnerabilities impact existing and potential attack surfaces.
- 3 To answer the **RQ3**, using our model-driven framework, we uncover potential Rowpress-based attack surfaces. Our evaluation reveals, for instance, that Rowpress offers significantly better selectivity for *targeted bit flip* and *pattern-based inference* than conventional Rowhammer on the tested devices, as demonstrated through theoretical analysis and chip-level experiments, suggesting high potential for new attack vectors and directions for mitigation.

The remainder of this paper is organized as follows. Section 2 provides background on DRAM architecture and summarizes known cell-level vulnerabilities that motivate our study. Section 2.2 introduces the threat model, defining the security properties of DRAM and outlining adversarial capabilities in common application scenarios. Section 3 presents our cell-level circuit model, which formalizes the physical mechanisms underlying retention and disturbance failures and connects them to concrete case studies and attack qualification. Section 4 reports

our experimental methodology and results on commodity DRAM devices, and demonstrates how the proposed parameters can be measured and adopted into the attack surfaces. Finally, Section 6 concludes the paper with a discussion of broader implications.

2. Preliminaries

In this section, we first introduce two classic disturbance-based attacks: the Rowhammer and Rowpress attacks. We then describe the underlying vulnerabilities in DRAM cells that enable these attacks and the resulting attack surfaces. Finally, we discuss the threat model and define the adversary capabilities used in our security evaluation.

2.1. DRAM Cell Vulnerabilities and Attack Surfaces

Rowhammer Attack. Rowhammer is a disturbance attack in which an adversary repeatedly activates and precharges one or more DRAM rows (the *aggressor rows*) at a high rate, causing charge to leak from physically adjacent *victim rows*. When the cumulative disturbance is large enough, victim cells experience accelerated charge loss and may flip their stored bits without any explicit write to the corresponding memory region, leading to a violation of memory integrity. Over the past decade, Rowhammer has been shown to corrupt page tables, escalate privileges, and compromise isolation boundaries on a wide range of commodity systems [4], [12], [15]–[17]. In all cases, the attack fundamentally exploits how repeated activations modulate the victim cell’s effective leakage paths, thereby shortening its time-to-failure.

Rowpress Attack. Rowpress is a closely related disturbance mechanism in which the adversary does not necessarily toggle the aggressor row at a high frequency, but instead keeps a row *open* (i.e., wordline asserted) for an extended period. This long-duration activation places sustained electrical stress on cells that share substrate and bitline structures with the pressed row, altering their leakage behaviour. Compared to Rowhammer, Rowpress can induce bitflips with fewer activation cycles and exhibits a stronger dependence on the stored data pattern in neighbouring rows, making it particularly relevant for *pattern-aware* integrity and confidentiality attacks [2]. In this work, we will treat Rowhammer and Rowpress as two distinct, but closely related, read-disturbance mechanisms acting on the same underlying cell-level vulnerabilities.

Existing work has demonstrated numerous feasible attacks that exploit memory cell vulnerabilities [3]–[5], [11]. Although these attacks are based on real empirical observations of cell behaviour, it remains unclear which physical leakage paths truly matter for each attack surface. In order to examine this more closely, we must establish the relevant leakage paths, the associated electrical properties, and their relation to the security properties of *volatility*, *integrity*, and *confidentiality*.

A DRAM cell typically consists of a single access transistor (A_T) and a storage capacitor (C_s), forming a 1T1C (one-transistor-one-capacitor) architecture. As illustrated in Fig. 2, one terminal of the storage capacitor is

1. To support reproducibility and foster further research, we make all artifacts of this work publicly available at Open Science Framework: https://osf.io/9gdyw/?view_only=eb552500a9f445d8a56611dc7ea36d25.

connected to a stable reference voltage, typically $V_{DD}/2$ or ground. In contrast, the other terminal is connected to the drain (or source) of A_T , whose source (or drain) is in turn connected to the bitline that interfaces with the sense amplifier.

We consider two primary leakage paths for the charge maintained in the storage capacitor C_s . These can be described in terms of two leakage currents, I_1 and I_2 . The first, I_1 , is the path from the storage node to **ground** through the p-well; the second, I_2 , is the path from the storage node to the **bitline**. Because cells along a column share the same bitline, the I_2 path is inherently influenced by neighbouring cells. This shared structure allows disturbances and side-channel information to propagate through the bitline-coupled leakage, making I_2 crucial for understanding pattern-dependent attack surfaces.

Throughout the rest of this paper, we will use a small set of effective parameters to summarize these two physical leakage paths. Specifically, we denote by R_S the effective resistance of the p-well leakage path associated with I_1 , and by R_B the baseline resistance of the bitline-coupled leakage path associated with I_2 under a neutral neighbour pattern. In addition, we introduce a pattern-dependent noise term $N(P)$ to capture how different neighbour patterns P modulate the bitline voltage and thus the effective leakage along the I_2 path. Intuitively, R_S primarily governs volatility and coarse integrity properties, while R_B and $N(P)$ are the key determinants of pattern-dependent integrity and confidentiality attacks. These quantities ($R_S, R_B, N(P)$) form the core parameters of our cell-level circuit model in Section 3 and will be explicitly measured and evaluated on real DRAM chips in Section 4.

2.1.1. P-well Leakage (I_1).

Intrinsic Leakage. Without cell refresh, the charge maintained in C_s will leak over time and eventually fail. The charge variation in the upper plate of the storage capacitor C_s determines the failure characteristics of the cell. Two major components determine the intrinsic leakage through I_1 . First, transistor leakage, mainly due to Gate-Induced Drain Leakage (GIDL), where a strong electric field near the drain junction under negative gate bias triggers band-to-band tunnelling (BTBT), generating unwanted current [18]. Second, dielectric leakage between the bitline contact (BLC) and storage node contact (SNC), which worsens with technology scaling as the dielectric spacing shrinks. Intrinsic leakage caused by these two behaviours forms the characteristic **volatility** of DRAM cells and does not require any external disturbance to occur.

Electron Injection Leakage. In read-disturbance scenarios, frequent activations of aggressor rows repeatedly toggle their wordlines, creating transient high electric fields in the substrate. These fields inject a small population of electrons into the p-well region beneath adjacent victim cells. Once injected, these electrons diffuse laterally and vertically, and some reach the victim cell’s storage node, effectively accelerating its charge leakage I_1 .

This mechanism differs fundamentally from retention failure, in which leakage is driven solely by passive decay paths, without external disturbance. Moreover, in contrast to the slow, steady nature of retention-based charge de-




Figure 2: DRAM cell leakage paths under retention and read disturbance.

cay, the injection-induced leakage current in Rowhammer scenarios can be significantly larger due to the strong, repetitive electric-field stress. This electron-injection leakage under Rowhammer disturbance gradually discharges the cell until it flips, and is a primary cause of DRAM **integrity** violations.

2.1.2. Bitline Leakage (I_2). The second pathway (I_2) involves charge sharing between the bitline and the storage capacitor, mediated by A_T .

Sensing Noise. During the sensing phase of a DRAM read, the bitline voltage changes only slightly, on the order of tens of millivolts, due to capacitive charge sharing with the storage capacitor [19]. This small swing must be amplified by the sense amplifier (SA). However, in this narrow amplification window, parasitic capacitive coupling from nearby structures can inject unwanted noise onto the bitline, a phenomenon known as **sensing noise**.

Under retention failure, a key feature of sensing noise is that it can induce **data-dependent** bitflips [6]. For instance, if a neighbouring bitline stores a 0, its falling transition during sensing may couple negatively onto an adjacent bitline storing a 1, suppressing its voltage rise. Such cross-bitline interference reduces the effective sensing margin and reduces the effective retention time of memory cells, thereby **amplifying** DRAM **volatility**.

Subthreshold Leakage. As shown in Fig. 2, when a victim row is under Rowhammer attack, the row’s wordline is partially activated by the crosstalk coupling induced by repeated aggressor row activations. The voltage between bitline and storage capacitor is V_{BS} , and this voltage difference leads to a subthreshold leakage current I_2 from the storage node to the bitline. Past a certain threshold of prolonged disturbance, V_S will decrease below $V_{DD}/2$ and the cell will flip from 1 to 0. However, the bitline voltage V_B fluctuation is determined by the values of the cells in the aggressor rows. If the cells in aggressor rows store a low voltage level, they pull V_B towards 0 and increase V_{SB} , thus amplifying the leakage current I_2 . In contrast, if the cells store a high voltage level, they pull V_B towards V_{DD} and decrease V_{SB} , thereby suppressing the leakage current I_2 . This subthreshold leakage, adjusted by the victim and aggressor data patterns, not only amplifies the impact of **integrity** violations, but also serves as a side channel that can be exploited to violate **confidentiality**.

2.2. Threat Model

Before discussing the circuit-level models, we first define the threat model that guides our interpretation of

DRAM cell parameters from a security perspective. Our goal is not to propose a new end-to-end exploit, but to provide a quantitative framework that characterizes how different DRAM devices behave under realistic adversarial capabilities.

Security Properties. Our analysis is organised around three core properties of memory:

- **Memory Integrity:** values in memory cannot be altered except through intended channels and with valid write access to the relevant memory region.
- **Memory Confidentiality:** values in memory cannot be learned except through intended channels and with valid read access to the relevant memory region.
- **Memory Volatility:** without active maintenance (refresh), the values in memory are irrecoverably lost.

Abstract Adversary Capabilities. We consider an *attacker-oriented* viewpoint in which the adversary interacts with DRAM only through architecturally allowed memory operations, but may have advanced control over how these operations are issued. In particular, we assume the adversary may possess some combination of the following high-level capabilities:

- **Pattern control:** the ability to place chosen data patterns P in selected memory regions (e.g., rows surrounding a victim cell) and to issue repeated accesses such as Rowhammer or Rowpress-style activations.
- **Observation:** the ability to repeatedly probe memory locations and observe whether bitflips occur, or whether previously written data can be partially recovered after a disturbance or power cycle.
- **Timing/refresh control:** in some scenarios, the ability to influence refresh behaviour or power cycles (e.g., cold-boot-like settings), thereby exposing retention-related failures.

These abstract capabilities capture a broad range of known DRAM-based attacks (e.g., Rowhammer, Rowpress, cold boot, RAMBleed) without committing to a particular exploit chain. The detailed mapping to concrete attacker models and system privileges is given in Appendix C.

Adversarial Objectives. Given these capabilities, we distinguish between *coarse* objectives, which do not exploit pattern dependence, and *pattern-aware* objectives, which do.

Coarse objectives.

- **OBJ-I-1 (Integrity-1):** alter *any* bits in a write-protected memory region, i.e., a coarse integrity violation.
- **OBJ-C-1 (Confidentiality-1):** learn *any* bits in a read-protected memory region, i.e., a coarse confidentiality violation.
- **OBJ-V-1 (Volatility-1):** partially retrieve bits of memory contents that should have been lost after refresh stops or after a power cycle, i.e., a coarse volatility violation.

Pattern-aware objectives. Knowledge of pattern dependence and the ability to control patterns allow the adversary to refine these into more targeted objectives:

- **OBJ-I-2 (Pattern-control integrity):** by writing specific data patterns in cells surrounding a victim cell, increase the victim’s flip probability $p_{\text{flip}}(P)$ and enable *pattern-sculpted* targeted modifications (e.g., steering bitflips to particular locations).




Figure 3: Equivalent Leakage Circuit.

- **OBJ-C-2 (Pattern-inference confidentiality):** by observing changes in a probe or victim cell’s flip probability across controlled disturbance trials, infer unknown neighbour data (the *surrounding pattern*).
- **OBJ-V-2 (Pattern-regeneration volatility):** by modeling retention failure characteristics and reading the start-up state of a specific victim cell after power-up, infer probable pre-power-down states of the stored data.

In summary, this threat model abstracts away system-specific exploit chains and instead focuses on *which cell-level parameters an adversary can, in principle, leverage to violate integrity, confidentiality, or volatility*. Our circuit model then provides a quantitative link between these parameters and the attacker objectives above, enabling a model-driven comparison of different DRAM devices and disturbance mechanisms.

3. Cell-Level Circuit Model and Attack Qualification

In this section, we start by presenting a unified cell-level leakage model that explains three classes of DRAM security-relevant behaviors: *volatility* (retention loss), *integrity violation* (disturbance-induced flips), and *confidentiality leakage* (pattern inference). As illustrated in Fig. 3, our cell-level circuit leakage model abstracts the physical cell circuit into an RC circuit. First, we characterize two primary leakage paths as two resistances R_S and R_B . Then, we assume the current I_1 leaks through R_S to the ground and I_2 leaks through R_B to the bitline. Under attack scenarios, bitline voltage will be modulated from 0 to V_{DD} due to the disturbance from the repeated read operations. Therefore, I_2 will be influenced by the data pattern in adjacent rows through the bitline. Using this model as a baseline we can delve into volatility, integrity, and confidentiality situations, providing a precise model of the cell vulnerability for each of them.

3.1. Volatility Model

From the perspective of volatility, we are interested in failures that occur *without any explicit disturbance*, purely due to intrinsic charge leakage. A DRAM cell is initially charged to V_{DD} (logic 1) and then left idle. A bitflip occurs when the held charge decays sufficiently that the wordline voltage is mischaracterized by the sense amplifier during read operations. During sensing, coupling noise from adjacent wordlines and bitlines can further reduce the margin of error, increasing the chance of a bitflip.

Accordingly, in the volatility model we consider only the leakage path I_1 through R_S and treat the bitline branch as effectively open ($R_B^{\text{eff}} \rightarrow \infty$). We then extend this to include pattern-induced sensing noise.

3.1.1. Noise-free Leakage Model. Retention failure occurs when the charge stored in the cell capacitor gradually leaks over time, causing the storage voltage to drop below the sensing threshold. We first consider all the neighboring cells hold the high voltage level so sensing noise can be neglected. During the idle period, the wordline is held low ($V_W = 0$), while the bitline is precharged to $V_{DD}/2$. The storage node voltage V_S is initialized to V_{DD} (logic 1). Considering only the intrinsic leakage I_1 in this setting, the behavior can be modeled by an equivalent RC circuit, where the capacitor C_S represents the storage element and the resistor R_S denotes the **intrinsic leakage resistance**.

The temporal evolution of the storage voltage follows a first-order RC discharge process:

$$\frac{dV_S(t)}{dt} = -\frac{1}{R_S C_S} V_S(t) \Rightarrow V_S(t) = V_{DD} \cdot e^{-t/(R_S C_S)}. \quad (1)$$

Here, the time constant $\tau = R_S C_S$ characterizes the retention capability of the cell, with reliable sensing typically guaranteed only within one time constant τ .

The instantaneous leakage current is given by

$$I_1(t) = C_S \cdot \frac{dV_S(t)}{dt} = -\frac{V_{DD}}{R_S} \cdot e^{-t/(R_S C_S)}. \quad (2)$$

This expression shows that the leakage current is inversely proportional to R_S and decreases exponentially over time. Cells with smaller R_S discharge more quickly, making them more vulnerable to retention failures.

Then the DRAM cell is accessed during a read operation, the stored charge in the capacitor is shared with the precharged bitline. This sharing produces a small voltage swing on the bitline given by

$$\Delta V_{BL} = \frac{C_S}{C_S + C_{BL}} \cdot \left(\frac{V_{DD}}{2} - V_S(t) \right), \quad (3)$$

where C_{BL} is the bitline capacitance. The ratio $\alpha = \frac{C_S}{C_S + C_{BL}}$ characterizes the fraction of the storage charge that is transferred to the bitline during sensing. The resulting signal $\Delta V_{BL}(t)$ is then amplified by the sense amplifier; if $|\Delta V_{BL}(t)|$ does not exceed the sensing threshold V_{SA} , the sense amplifier may misread the cell and the stored data is effectively lost.

3.1.2. Pattern-noised Leakage Model. As discussed in Section 2.1.2, when the surrounding cells are biased to a low-voltage state, they inject additional charge into the shared bitline through bitline and wordline capacitive coupling. We model the resulting *pattern-induced noise* as an effective voltage perturbation N on the bitline. Taking this noise into account, the readout condition becomes

$$\begin{cases} \Delta V_{BL} + N > V_{SA}, & \text{Correct read (still '1')} \\ \Delta V_{BL} + N < -V_{SA}, & \text{Flip error (misread as '0')} \\ |\Delta V_{BL} + N| < V_{SA}, & \text{Uncertain outcome} \end{cases} \quad (4)$$

In pure volatility, the bitline branch is effectively closed, $R_B^{\text{eff}} \rightarrow \infty$, so $\tau \approx C_S R_S$. Here integrity errors are time-pattern dependent through R_S and N : the flip probability increases with longer dwell time and larger $|N|$, and the attack knob is to *extend refresh interval* or *bias patterns* that reduce the effective R_S (coupling-assisted leakage). Formally, referring to equation 3, a flip occurs when

$$\frac{C_S}{C_S + C_{BL}} \cdot \left(\frac{V_{DD}}{2} - V_S(0) e^{-T/(C_S R_S)} \right) + N < -V_{SA}. \quad (5)$$

Thus, in the pattern-noised leakage model, both the intrinsic leakage resistance R_S and the external coupling noise N jointly determine the cell's data *volatility*. Larger noise magnitudes $|N|$ reduce the sensing margin and increase the probability that a cell becomes vulnerable to retention-induced bitflips.

3.2. Read Disturbance Model

Read-disturbance phenomena such as Rowhammer and Rowpress introduce failure modes that go beyond pure volatility. In the integrity and confidentiality model, we consider both leakage paths, I_1 and I_2 . As illustrated in Fig. 3, we customize their leakage resistances R_S and R_B in these two paths, where R_S and R_B jointly dominate the integrity loss of a cell. However, R_B is controlled by the neighboring *access pattern*, therefore, by observing or measuring the R_B we can infer *access pattern* inversely to violate confidentiality.

3.2.1. P-well Leakage (I_1). As discussed in Section 2.1.2, repeated wordline activations inject minority carriers into the cell-array p-well substrate. The transient voltage on the aggressor wordline locally increases the free-electron density in the p-well, and these carriers are subsequently captured by neighbouring cells. As a result, the intrinsic leakage path I_1 becomes more conductive, which we capture by an effective resistance R_S .

In contrast to pure retention, the relevant stress parameter is not elapsed time but the cumulative activation history of neighbouring rows. Each activation keeps the wordline high for an effective stress window t_{AggON} during which additional electrons are injected into the p-well. After hammering counts (HC) such activations, the victim cell has effectively experienced a total leakage exposure of $t_{\text{eff}} = HC \cdot t_{\text{AggON}}$.

Starting from the standard RC decay for the storage node, $V_S(t) = V_{DD} \exp(-t/(R_S C_S))$, we reparameterize the evolution in terms of the hammer count HC by substituting $t = t_{\text{eff}}$. This yields a count-domain leakage model

$$I_1(HC) = \frac{V_S(HC)}{R_S} = \frac{V_{DD}}{R_S} \cdot \exp\left(-\frac{HC \cdot t_{\text{AggON}}}{R_S C_S}\right). \quad (6)$$

3.2.2. Bitline Leakage (I_2). Bitline leakage I_2 indicates the subthreshold current between the source and drain of the access transistor when the channel is not totally opened. During an ACT command process, the aggressor

wordline rises from 0 to V_{PP} within T_{RCD} , producing a transient spike on the victim wordline:

$$\Delta V_{\text{victim}} \approx V_{PP} \cdot \frac{C_C}{C_C + C_{WL}}, \quad (7)$$

where C_C is the parasitic capacitance C_C couples adjacent wordlines and C_{WL} is the victim wordline capacitance to ground.

Although ΔV_{victim} does not exceed the threshold V_T , it biases the transistor into the subthreshold region, yielding a leakage current I_2

$$I_2(t) = \mu_n C_d v_{th}^2 \cdot \frac{W}{L} \cdot e^{\frac{\Delta V_{\text{victim}} - V_B - V_T}{n v_{th}}} \left(1 - e^{-\frac{V_S(t) - V_B}{v_{th}}} \right), \quad (8)$$

where μ_n is electron mobility, C_d the depletion capacitance, W/L the aspect ratio, v_{th} the thermal voltage, n the subthreshold swing coefficient, V_S the victim storage-node voltage, and V_B the bitline voltage. To further express effective resistance R_B :

$$R_B(V_S, V_B) = \frac{V_S - V_B}{A \left(1 - e^{-\frac{V_S - V_B}{v_{th}}} \right)}. \quad (9)$$

with the **Subthreshold Factor** A

$$A = \mu_n C_d v_{th}^2 \cdot \frac{W}{L} \cdot e^{\frac{\Delta V_{\text{victim}} - V_B - V_T}{n v_{th}}}, \quad (10)$$

As shown in Figure 3, R_B through the I_2 is controlled by the *access data pattern*. We will discuss the different patterns and how we model predicts as follows:

111 pattern case: If capacitors in the victim rows and aggressor rows store the same value (e.g., 111), during the read disturbance, V_B will be pulled up to V_{DD} by aggressor cells. Therefore, when $V_S \approx V_B$, it makes $|V_S - V_B| \ll v_{th}$, the exponential term can be linearized as $(1 - e^{-x}) \approx x$, leading to

$$R_B^{111} \approx \frac{v_{th}}{A \cdot (V_S - V_B)} \quad (11)$$

Consequently, the effective resistance R_B increases to a **very large value**, indicating that the leakage toward the bitline is practically *insignificant* in this case.

010 pattern case: In contrast, when the victim rows and aggressor rows store the opposite value (e.g., 010). V_B will be pulled down to 0 by aggressor rows under the read disturbance. In this case, $|V_S - V_B| \approx V_{DD} > v_{th}$, makes $(1 - e^{-x}) \approx 1$, leading to

$$R_B^{010} \approx \frac{v_{th}}{A} \quad (12)$$

Through this modeling, we find that under the 010 pattern effective resistance R_B becomes a reasonable value which depends on the **subthreshold voltage** v_{th} of access transistor and our defined **subthreshold factor** A . It also reveals the mechanisms of the widely observed phenomenon - why CheckerBoard pattern 010 is able to accelerate the charge loss of the cell then increase the likelihood of the bitflips.

3.3. Case Studies

Before detailed attack qualification, we present two representative scenarios to illustrate how the proposed (R_S, R_B) -based leakage model supports *decision-oriented* reasoning about integrity and confidentiality vulnerabilities. They include two attacks aligned with our model-driven adversarial objectives. **Case A** targets *integrity* via pattern-sculpted disturbance to edit a *targeted* sensitive bit; **Case B** targets *confidentiality* by inferring an unknown neighbour pattern from the flip statistics.

3.3.1. Case A: Pattern-Sculpted Positioned Edit (Integrity). This instantiates the precise integrity violation scenario (**OBJ-I-2**, Pattern Control Attacker) from Section 2.2. The goal is to flip a chosen bit in the DRAM arrays and leave other bits unchanged by increasing its flip likelihood through controlled neighbour access patterns. In the equivalent model, pattern control primarily acts through the bitline-side path: it reduces R_B under 010 and increases R_B under 111. Let the effective time constant under pattern pat be

$$\tau^{\text{pat}} = C_S (R_S \parallel R_B^{\text{pat}}), \quad (13)$$

so that, for a fixed stress window (e.g., t_{AggON} and activation count), the flip probability increases as τ decreases.

Selective flipping via R_S and R_B . For a target cell that flips only under 010, pattern control makes $R_B^{010} \ll R_B^{111}$. A sufficient condition for selectivity is the existence of a threshold τ^* (determined by the stress schedule) such that

$$\tau_{\text{target}}^{010} = C_S (R_S^{\text{tar}} \parallel R_B^{010}), \quad \tau_{\text{byst}}^{111} = C_S (R_S^{\text{byst}} \parallel R_B^{111}). \quad (14)$$

Selective-flip succeeds if there exists a stress level τ^* such that

$$\tau_{\text{target}}^{010} < \tau^* < \tau_{\text{byst}}^{111}$$

(e.g., choose $t_{\text{AggON}}/\#\text{activations}$ that corresponds to τ^*).

When $R_B^{010} \ll R_S^{\text{tar}}$, the target's time constant is dominated by the bitline path ($\tau_{\text{target}}^{010} \approx C_S R_B^{010}$), enabling a high flip likelihood. Conversely, suppressing bystanders relies not only on a large R_B^{111} but also on sufficiently large R_S^{byst} ; otherwise,

$$R_S^{\text{byst}} \ll R_B^{111} \Rightarrow \tau_{\text{byst}}^{111} \approx C_S R_S^{\text{byst}} \quad (\text{small}), \quad (15)$$

and unwanted flips may occur even under 111.

Practical implication.. Based on this analysis in the practical *positioned integrity* attack design, in order to increase the attack success rate, attackers should refer to the following metrics:

	Parameter push		Effect
Target (010)	$R_B^{010} \downarrow$	(stronger BL path)	$\tau_{\text{target}}^{010} \downarrow$
Bystanders (111)	$R_S^{\text{byst}} \uparrow$	(less self-leak)	$\tau_{\text{byst}}^{111} \uparrow$

Therefore, for practical attackers, their goal is to enlarge the gap $\Delta\tau = \tau_{\text{byst}}^{111} - \tau_{\text{target}}^{010}$ and the gap $\Delta R = R_S^{\text{byst}} - R_B^{010}$ to effectively launch the attacks.

3.3.2. Case B: Pattern-to-Flip Tomography (Confidentiality). This instantiates the precise confidentiality violation scenario (**OBJ-C-2**, Pattern Inference Attacker) from Section 2.2. The goal of the adversary is to infer the unknown data bit of a *neighbour* (victim) cell by observing the flip statistics of a *probe* cell whose neighbourhood includes that victim. We assume the *aggressor* row encodes the unknown bit $d \in \{0, 1\}$ to be inferred. For each victim cell (whose R_S and R_B^{pat} are known), we drive the neighbourhood with two patterns, $\text{pat} \in \{010, 111\}$, and read flips on the victim as a probe of d . The victim’s pattern-conditioned time constants are defined as

$$\tau^{\text{pat}} = C_S(R_S \parallel R_B^{\text{pat}}) \quad (16)$$

for a fixed stress window (e.g., t_{AggON} , activation count). The following box indicates the necessary conditions so that the cell flips under 010 and remains stable under 111.

**Unambiguous inference with one victim cell
requires a separable window**

$$\tau^{010} < \tau^* < \tau^{111}$$

When this holds, the decision rule
flip $\Rightarrow \hat{d} = 1$, no flip $\Rightarrow \hat{d} = 0$ is correct for that
cell.

Practical implication.. If there is not enough gap between the value of the R_B^{010} and R_S , the discharge time-constants under 010 and 111 become comparable, consequently the decision margin (likelihood ratio) collapses and the posterior $\Pr(\text{aggressor} = 0 | \text{obs})$ cannot be separated from $\Pr(\text{aggressor} = 1 | \text{obs})$ (i.e., stays close to 0.5).

Therefore, in a practical *confidentiality* attack, in order to develop the accuracy of the inference, attackers should refer to the following metrics:

Parameter push	Effect on τ^{010}, τ^{111}	Net impact
$R_B^{010} \downarrow$	$\tau^{010} \downarrow$	easier 010 flip
$R_S \uparrow$	both $\tau^{010}, \tau^{111} \uparrow$	harder 111 flip

Similarly, the goal of the *confidentiality* attackers is to enlarge the gap $\Delta\tau = \tau_{\text{probe}}^{111} - \tau_{\text{probe}}^{010}$ and the gap $\Delta R = R_S - R_B^{010}$ to effectively launch the attacks.

Taken together, Case A and Case B show how our (R_S, R_B) -based leakage model directly links device-level parameters to the success conditions of both targeted integrity edits and pattern-inference attacks. By reading off gaps such as $\Delta\tau$ and ΔR from our calibrated parameters, an adversary can systematically select the most exploitable cells and stress windows. We will give a more detailed qualification of attacks in the following section.

3.4. Attacks Qualification

After two case studies, we adopted the parameters in our model to precisely explain the integrity and confidentiality attacks quality in three classes of attack scenarios considered in the Threat Model (Section 2.2):

(i) **General integrity attacks (OBJ-I-1, OBJ-I-2, OBJ-V-1, OBJ-V-2).** The attacker’s goal is to flip stored bits without direct writes, i.e., to drive the storage node across the sense-amplifier decision threshold during

a read. In our model, the vulnerability of cells is governed by the effective leakage resistance R_S and R_B . For a given mechanism and pattern, the total leakage conductance can be written as

$$G_{\text{tot}} = \frac{1}{R_S} + \frac{1}{R_B} \quad (17)$$

where R_S captures the cell-to-well leakage and R_B aggregates any additional bitline-side leakage branch (if present). The corresponding time constant is $\tau \approx C_S/G_{\text{tot}}$, so a larger G_{tot} directly translates into shorter flip times and stronger integrity attacks.

Pure retention failures operate with $R_B \rightarrow \infty$, such that $G_{\text{tot}} \approx 1/R_S$. Rowhammer and Rowpress both add a finite bitline-side branch R_B and hence increase G_{tot} under the access pattern 010. Therefore, the quality of a general integrity attack can be determined by the total leakage conductance G_{tot} .

(ii) **Targeted integrity attacks (OBJ-I-2).** Beyond flipping bits as fast as possible, a more powerful adversary aims for *selective* integrity violations: inducing flips only under specific data patterns while keeping other patterns stable. In this case, the selectivity is not determined by the absolute leakage G_{tot} alone, but by the *pattern-induced difference* between leakage under two patterns, e.g., 010 vs. 111. For a victim cell i , we write

$$G_{111}(i) \approx \frac{1}{R_S(i)}, \quad G_{010}(i) = \frac{1}{R_S(i)} + \frac{1}{R_B(i)}, \quad (18)$$

so that the additional pattern lever is

$$\Delta G(i) \triangleq G_{010}(i) - G_{111}(i) = \frac{1}{R_B(i)} \quad (19)$$

A small R_B (large $1/R_B$) means that the 010 pattern decays faster than 111. However, true selectivity also requires that the baseline leakage under 111 remains small; otherwise both patterns may flip within the same stress window.

To capture this, we consider the *relative* pattern gap

$$\Delta G_{\text{rel}}(i) \triangleq \frac{G_{010}(i) - G_{111}(i)}{G_{111}(i)} = \frac{1/R_B(i)}{1/R_S(i)} = \frac{R_S(i)}{R_B(i)} \quad (20)$$

In this view, a large $R_S(i)$ suppresses leakage under 111 while a small $R_B(i)$ strongly enhances leakage under 010; both are jointly necessary for highly selective integrity attacks.

(iii) **Confidentiality attacks. (OBJ-C-2).** In confidentiality attacks, the adversary does not aim to corrupt data but to **infer** a hidden pattern (e.g., whether the neighborhood stores 0 or 1) from the flip statistics of a victim row. The attacker typically adopts a simple rule: if a flip is observed within the attack budget, guess the “vulnerable” pattern (e.g., 010, storing 0); if no flip is observed, guess the “stable” pattern (e.g., 111, storing 1).

Similar to the targeted integrity attack, it is the *gap* between the intrinsic p-well resistance R_S and the pattern-dependent bitline resistance R_B that drives confidentiality. For each bit i , the effective conductance under 111 and 010 patterns is already defined in Equation (18). For a fixed hammer budget, we abstract the stress as a conductance threshold θ : bits with $G_{010}(i) \geq \theta$ will flip with high probability under the 010 pattern, and bits with

$G_{111}(i) \geq \theta$ will flip with high probability under the 111 pattern. Let $p_{010}(\theta)$ and $p_{111}(\theta)$ denote the fractions of bits that flip under 010 and 111, respectively. Under the rule “flip \Rightarrow 0, no flip \Rightarrow 1” and assuming equal priors, the expected inference accuracy at stress level θ is

$$\widehat{\text{Acc}}(\theta) = \frac{1}{2} \left[p_{010}(\theta) + (1 - p_{111}(\theta)) \right]. \quad (21)$$

4. Experiments and Attack Evaluations

4.1. Experimental Objectives and Overview

The purpose of this section is to empirically validate that the cell-level leakage model in Section 3 provides *actionable, decision-oriented* insights for the security scenarios formalized in Section 2.2. In particular, this section addresses the three research questions (RQs) introduced in Section 1.2 from an empirical perspective.

RQ1 (Measurability of physical parameters). We first validate that the proposed circuit-level parameters, effective p-well resistance (R_S), effective bitline resistance (R_B), pattern-induced noise (N) can be reliably extracted from real DRAM modules. Through systematic retention profiling and read disturbance (Rowhammer and Rowpress) stress tests across multiple devices, we examine whether these quantities are measurable in practice. This confirms that our framework is not only theoretical but also experimentally grounded.

RQ2 (Security Analysis of attack surfaces). Given the extracted parameters, we then study how they can be used to analyse and qualify the performance of integrity and confidentiality attacks. We also compare the parameters across different DRAM chips as well as different read disturbance mechanisms (e.g., Rowhammer and Rowpress).

RQ3 (Insights on existing attack surfaces). Finally, we investigate whether Rowpress can strengthen or extend existing attack surfaces when compared to Rowhammer under a comparable “damage budget” (e.g., total flip count or stress window). We focus on two aspects: (i) selective integrity violation, where the attacker aims to steer flips towards specific victim cells while minimizing collateral damage, and (ii) pattern-dependent confidentiality leakage, where the attacker infers unknown data through disturbance-induced bitflips. Our experiments show that Rowpress exposes finer-grained, more pattern-aware control knobs, suggesting an expanded security risk surface that is not captured by prior empirical-only analyses.

4.2. Hardware Platform

In this section, we present the hardware platform of our test framework, including the FPGA-based experimental platform, the custom memory controller, and the diverse set of DRAM chips used for validation.

Experimental platform: We conduct our experiments on an Xilinx Alveo U200 FPGA accelerator card, which can work with high-bandwidth DDR4 DRAM modules. A custom memory controller is implemented on the FPGA, enabling fine-grained control over the memory under test.

Memory Controller: For this, we leverage *DRAM Bender* [20], a flexible open-source framework that enables direct manipulation of DRAM timing parameters

and access patterns. DRAM Bender allows us to bypass conventional memory controllers and issue raw DRAM commands (e.g., ACT, PRE, RD, WR) directly to the device. This low-level programmability makes it possible to construct customized access sequences, such as hammering patterns or controlled retention tests, and to precisely measure resulting bitflips. In addition, the customized memory controller sets timing configurations t_{RCD} and t_{RP} to 13.75 ns, while t_{RAS} was set to 32 ns. These timing parameters follow JEDEC specifications [21] and are not intentionally relaxed unless explicitly stated (e.g., Rowpress experiments).

DRAM Chips: To validate our cell-level circuit models, we selected a diverse set of DRAM modules that span multiple sources of variability. The detailed information of tested DRAM modules is shown in the Table 1. This set includes (i) multiple modules of the same model from a single vendor (Micron 16 GB) to capture device-to-device variation under nominally identical fabrication; (ii) modules with different density and scaling generations (Micron 16 GB, Samsung 16 GB, Lenovo 8 GB vs ADATA 4 GB) to reflect broader technology trends. (iii) modules from different third-party vendors (ADATA 4 GB, Innodisk 8 GB and Lenovo 8 GB) to compare vendor-specific circuit and process designs. Together, these modules provide a comprehensive and representative basis for validating the robustness and generality of our proposed models. It should be noted that our model focuses on the fundamental 1T1C cell architecture and its intrinsic leakage mechanisms (p-well and bitline leakage), which remain structurally consistent from DDR4 to DDR5. Of these two DDR4 is currently more practical for these kind of deep-dive empirical tests requiring custom memory controllers and varied test hardware. Thus, the experiments in this paper are based on DDR4 chips.

TABLE 1: Tested DDR4 DRAM modules

DIMM	Vendor	Type	Density	Organization	Speed	Year
D1	Micron	RDIMM	16GB	1Rx4	2400	2021
D2	Micron	RDIMM	16GB	1Rx4	2400	2021
D3	Micron	RDIMM	16GB	1Rx4	2400	2021
D4	Samsung	UDIMM	16GB	2Rx8	2400	2021
D5	Lenovo	SODIMM	8GB	1Rx4	2666	2024
D6	Lenovo	SODIMM	8GB	1Rx4	2666	2024
D7	Innodisk	UDIMM	8GB	1Rx4	2400	2017
D8	ADATA	UDIMM	4GB	1Rx4	2400	2018

4.3. Experimental Results

We conducted experiments on commodity DDR4 DRAM modules using our FPGA-based testing platform. For *retention profiling*, we extended the refresh-pausing window up to **3600 seconds**, for *Rowhammer*, we stressed the cells with double-sided maximum to **15M hammering counts (HC)**, and for *double-sided Rowpress*, we set the t_{AggON} to 1000ns, then stressed the cells with maximum to **1.5M hammering counts (HC)**. Both tests were with different written patterns (baseline a11-1 and checkerBoard) or aggressor-victim patterns (baseline 111 and 010) to capture pattern-dependent effects on the true cells. During these experiments, we collected both the total bitflip counts and the spatial distribution maps. By applying Algorithm 1 and Algorithm 2, we extracted the p-well effective resistance R_S , $R_S(N)$ (with noise) and the

Algorithm 1 Extraction of R_S^{RD} and R_B^{RD} from read disturbance (Rowpress and Rowhammer)

```

1: Input: cell  $v$ ; patterns  $P \in \{111, 010\}$ ; hammer rate  $f_{\text{RD}}$ ;
   target flips  $N_{\text{target}}$ ;  $C_S, V_{DD}, V_{\text{flip}}, V_{TH}$ 
2: Output:  $R_S^{\text{RD}}, R_B^{\text{RD}}$ 
3: function HCATARGETFLIPS( $v, P$ )
4:    $HC \leftarrow HC_{\text{init}}$ 
5:   while victim flips  $< N_{\text{target}}$  do
6:     Configure victim/aggressors for pattern  $P$ 
7:     Initialize  $v \leftarrow 1$ 
8:     Hammer aggressors with count  $HC$  at rate  $f_{\text{RD}}$ 
9:     Read victim and count flips
10:    if flips  $< N_{\text{target}}$  then
11:       $HC \leftarrow \alpha \cdot HC$  ▷ e.g.,  $\alpha = 2$ 
12:    return  $HC$ 
13:  $HC_{111} \leftarrow \text{HCATARGETFLIPS}(v, 111)$ 
14:  $HC_{010} \leftarrow \text{HCATARGETFLIPS}(v, 010)$ 
15:  $T_{111} \leftarrow HC_{111}/f_{\text{RD}}, T_{010} \leftarrow HC_{010}/f_{\text{RD}}$ 
16:  $R_{\text{eff}} \leftarrow \frac{T_{111}}{C_S \ln(\frac{V_{DD}}{V_{\text{flip}}})}$  ▷  $R_{\text{eff}} \approx R_S^{\text{RD}}$ 
17:  $E \leftarrow (\frac{V_{DD}}{V_{\text{flip}}})^{T_{010}/T_{111}}$ 
18:  $A \leftarrow \frac{1}{R_{\text{eff}}} \cdot \frac{V_{DD} - EV_{\text{flip}}}{E - 1}$ 
19:  $R_S^{\text{RD}} \leftarrow R_{\text{eff}}, R_B^{\text{RD}} \leftarrow \frac{V_{TH}}{A}$ 
20: return  $R_S^{\text{RD}}, R_B^{\text{RD}}$ 

```

Algorithm 2 Measurement of R_S and $R_S(N)$ with Pattern Noise

```

1: Input: Victim cell  $v$ , neighbor patterns  $\{N_0 =$ 
   all-1,  $N_1, \dots, N_k\}$ , time intervals  $\{t_1, \dots, t_n\}$ 
2: Output:  $R_S$  and  $R_S(N)$  with noise
3: for each pattern  $N$  do
4:   Write victim cell  $v \leftarrow 1$ 
5:   Program neighbors of  $v$  according to  $N$ 
6:   for each  $t_i$  do
7:     Disable refresh for duration  $t_i$ 
8:     Read cell  $v$ 
9:     if read result == 0 then
10:       $t_{\text{flip}} \leftarrow t_i$ , break
11:   Compute  $R_S(N) = \frac{t_{\text{flip}}}{C_S \cdot \ln(V_{DD}/V_{\text{flip}})}$ 
12: return  $\{R_S(N)\}$ 

```

bitline effective resistance R_B corresponding to *retention failures* and *read* disturbance conditions.

4.3.1. w.r.t. RQ1: Measurability of Physical Parameters. We first verify whether the proposed parameters R_S , R_B , and N can be reliably extracted from real DRAM modules, and whether their statistical properties are consistent with our models. Table 2 summarizes the tested DRAM modules and the extracted parameters of 5000 most vulnerable *true cells* in each row of each memory chip in our study. For each device, we report: (i) the p-well effective resistance R_S estimated under the all-1 retention pattern, (ii) the effective noise parameter $R_S(N)$ obtained from the checkerboard retention pattern, (iii) the p-well effective resistance R_S and bitline effective resistance R_B measured under Rowhammer and Rowpress with the 111 and 010 activation pattern.

To enable security reasoning in later sections, we summarise two empirically robust observations that directly

reflect the physical structure of DRAM cell leakage and determine which attack surfaces are practically viable.

Observation 1. *Rowpress-induced p-well effective resistance R_S^{RP} is consistently orders of magnitude greater than the Rowhammer-induced p-well effective resistance R_S^{RH} , revealing Rowpress induces much smaller leakage to the p-well than Rowhammer.*

We observe that for every tested Dual In-Line Memory Module (DIMM), the Rowpress p-well resistance range R_S^{RP} is shifted to significantly larger values than the corresponding Rowhammer range R_S^{RH} . For instance, on Micron₁ the Rowhammer R_S^{RH} spans 4.78–305 × 10¹⁰ Ω, whereas Rowpress reaches 36.1–1050 × 10¹⁰ Ω; similarly, Lenovo₁, Lenovo₂, Innodisk, ADATA and Team all show R_S^{RP} that are roughly one to two orders of magnitude higher than R_S^{RH} , confirming that Rowpress induces much weaker effective p-well leakage than Rowhammer in practice. This implies that, in the absence of any pattern-controlled bitline effect, Rowpress does aggravate the intrinsic cell vulnerability; however, it is in fact **less** damaging to the p-well leakage path than Rowhammer.

Observation 2. *Rowpress-induced bitline effective resistance R_B^{RP} is consistently orders of magnitude smaller than the Rowhammer-induced bitline effective resistance R_B^{RH} , revealing Rowpress induces much stronger leakage to the bitline than Rowhammer.*

From Table 2, we see that for all tested DIMMs, the Rowpress R_B^{RP} ranges (in 10⁸ Ω) are systematically shifted to much lower values than the corresponding Rowhammer R_B^{RH} ranges, often by one to two orders of magnitude. This implies that, once a suitable 010-like access pattern is applied, Rowpress can dominantly route the discharge through the bitline branch and thus offer attackers a substantially stronger handle for both targeted integrity edits and pattern-dependent confidentiality attacks than Rowhammer.


4.3.2. w.r.t. RQ2: Measured Parameters for Attack Qualification. So far, we have extracted the leakage parameters (R_S, R_B) from both retention-failure and read-disturbance experiments and reported several observations on how different mechanisms and patterns shape the attack surface. We now answer RQ2 by adopting these parameters to directly qualify the attacks by the metrics introduced in Section 3.4.

(i) General integrity attacks. Our **Observation 1** and **Observation 2** already implied that Rowhammer has stronger leakage than Rowpress through R_S and weaker leakage through R_B . Intuitively, it indicates that Rowhammer performs a stronger intrinsic leakage than Rowpress, but Rowpress induces more pattern-controlled leakage. For a more detailed quantification of their impacts, we can use Equation 4 to calculate and compare them.

For each victim cell i , we quantify the leakage conductance to evaluate direct integrity attacks under Rowpress relative to Rowhammer across different access patterns. In Figure 4, we aggregate $G_{\text{tot}}(i)$ over all test rows to obtain per-module medians across our test DRAM chips. Across modules, the RH-111 curve stays above RP-111, indicating stronger integrity capability when only R_S is active, while under the 010 pattern, the RP-010 curve is significantly elevated and always surpasses RH-010, showing that once bitline leakage is enabled, Rowpress

TABLE 2: Measured resistance ranges for DRAM modules

DIMM	Volatility		Rowhammer Test		Rowpress Test	
	$R_S (\times 10^{13} \Omega)$	$R_S(N) (\times 10^{13} \Omega)$	$R_S (\times 10^{10} \Omega)$	$R_B (\times 10^8 \Omega)$	$R_S (\times 10^{10} \Omega)$	$R_B (\times 10^8 \Omega)$
D1	10.6–5530	9.01–5244	4.78–305	57.9–909	36.1–1050	38.5–362
D2	9.2–4988	7.7–4620	5.31–306	37.2–905	25.5–2125	25.7–584
D3	13.2–6720	11.4–5420	5.31–240	76.7–829	55.3–1830	44.4–258
D4	7.3–5230	6.1–4240	6.33–5360	43.2–2350	47.8–84560	17.8–992
D5	8.50–8960	7.2–6980	4.54–31880	29.5–6470	102–638000	24.3–1170
D6	6.5–7350	5.5–6321	5.21–46300	36.2–8410	512–826000	15.6–2230
D7	14.4–5620	10.8–4770	8.77–442	44.3–486	25.5–1620	9.45–312
D8	15.6–6820	13.2–5183	4.31–85.0	37.8–643	31.9–424	26.7–393


 Figure 4: The total leakage conductance G_{tot} under Rowhammer vs Rowpress.


provides a much stronger direct integrity attack handle than Rowhammer.

Takeaway:

Rowhammer yields a larger attack capability than Rowpress under the 111 pattern, but once the 010 pattern enables a low- R_B bitline path, Rowpress overtakes Rowhammer and becomes the more powerful general integrity attack vector.

(ii) Targeted integrity attacks. Our two key observations also show, in general, Rowpress has a larger gap between R_S and R_B than Rowhammer. Therefore, this implies higher flipping selectivity if attackers can precisely control bit flips through access patterns. For further precise measurement of the attack quality, in section 3.4, we provide *relative* pattern lever $\Delta G_{\text{rel}}(i)$ (Equation 20) to measure the attacker’s capability of targeted integrity violation.

As illustrated in Figure 5, we aggregate $\Delta G_{\text{rel}}(i)$ over all test rows to obtain per-module medians across our test DRAM chips. It shows that, for every DIMM, Rowpress systematically yields much larger ΔG_{rel} than Rowhammer, often by more than an order of magnitude. It indicates that the pattern lever between vulnerable (010) and stable (111) configurations is significantly stronger under Rowpress and thus more favorable for targeted integrity attacks.


 Figure 5: Relative Pattern Lever ΔG_{rel} under Rowhammer vs Rowpress.

Takeaway:

Across all tested DIMMs, Rowpress consistently provides a much stronger capability for targeted integrity attacks than Rowhammer, enabling more selective and controllable bit flips under the same disturbance budget.

(iii) Confidentiality attacks. To capture the *practical* confidentiality attack capability on real DIMMs, we instantiate this model using the extracted per-row parameters (R_S, R_B) for both Rowhammer and Rowpress. For each mechanism and each DIMM, we rank all bits by their 010-pattern conductance and select the m_{010} most vulnerable bits such that, under a fixed hammer budget HC , all of them flip when the adjacent pattern is 010. Applying the same HC when the adjacent pattern is 111, we count how many of these m_{010} bits also flip and denote this number by m_{111} .

The attacker keeps the same decision rule: if a bit flips, guess 010; otherwise, guess 111. For the selected m_{010} bits, the success probability when the true pattern is 010 is 1 (they all flip by construction), while the success probability when the true pattern is 111 is $1 - m_{111}/m_{010}$. Assuming equal priors, the overall estimated inference accuracy becomes

$$\text{Acc} = \frac{1}{2} \left[1 + \left(1 - \frac{m_{111}}{m_{010}} \right) \right] = 1 - \frac{1}{2} \frac{m_{111}}{m_{010}}. \quad (22)$$

As illustrate in Figure 6, we set m_{010} to 5000 and estimate the average inference accuracy by Equation (22) across the test rows and devices. Across all tested DIMMs, Rowpress consistently yields a higher Acc_{max} for more aggressive targets (e.g., $> 90\%$ accuracy), while Rowhammer often fails to reach the same level within any stress




Figure 6: Estimated average confidentiality inference accuracy among the most vulnerable 5000 cells in each row across DIMMs.

window, highlighting the stronger confidentiality threat enabled by Rowpress.

Takeaway:

Rowpress consistently enables substantially stronger confidentiality attacks than Rowhammer: under the same hammer budget, it achieves higher estimated pattern-inference accuracy and exposes a larger fraction of bits whose values can be reliably inferred from flip behavior.

4.3.3. w.r.t. RQ3: Insights on Existing Attacks. While RQ2 demonstrates that the extracted (R_S, R_B) parameters accurately characterize integrity and confidentiality vulnerabilities, RQ3 asks a broader system-level question: *Does Rowpress merely replicate Rowhammer, or does it systematically amplify existing Rowhammer-style attacks under identical resource constraints?* Using our calibrated parameters, we compare the two mechanisms across all security-relevant dimensions.

First, for *direct integrity attacks*, we aggregate the total leakage conductance $G_{\text{tot}} = 1/R_S + 1/R_B$ for different mechanisms and patterns (Figure 4). Across all DIMMs, the RH-111 curve stays above RP-111, confirming that Rowhammer provides a stronger integrity handle when only the p-well path is active. However, once the 010 pattern enables a low- R_B bitline branch, the RP-010 curve is significantly elevated and often surpasses RH-010, showing that Rowpress becomes the more powerful direct integrity weapon when attackers can exploit pattern-dependent bitline leakage.

Second, for *selective integrity attacks*, we focus on the pattern-induced component $\Delta G = 1/R_B$ and its relative gap $\Delta G_{\text{rel}}(i) = R_S(i)/R_B(i)$. Figure 5 shows that Rowpress consistently yields much larger ΔG_{rel} than Rowhammer across our DIMMs. In other words, Rowpress not only accelerates flips, but also *sharpens* the contrast between vulnerable (010) and stable (111) patterns on the same physical row, making it substantially easier for an adversary to steer bitflips toward targeted data values while keeping bystanders intact.

Third, for *confidentiality attacks*, we instantiate the pattern-inference attacker using the measured (R_S, R_B) and the accuracy model in Equation (22). For each mechanism and DIMM, we select the m_{010} most vulnerable bits under the 010 pattern, choose a hammer budget HC such that all of them flip under 010, and then measure

m_{111} —how many of these bits also flip under 111 with the same HC . The resulting estimated inference accuracy $Acc = 1 - \frac{1}{2}m_{111}/m_{010}$ is summarized in Figure 6. Across all tested DIMMs, Rowpress consistently achieves higher Acc than Rowhammer, and is often the only mechanism that can reach aggressive targets (e.g., $> 90\%$ accuracy) within any stress window, highlighting a strictly stronger confidentiality threat.

Overall, the RQ3 analysis shows that Rowpress behaves as a systematic **amplifier** of existing Rowhammer-style attacks under realistic constraints. It increases G_{tot} when pattern-controlled bitline leakage is available (direct integrity), enlarges the pattern lever ΔG_{rel} (selective integrity), and enables significantly higher inference accuracy Acc for pattern-based key extraction (confidentiality) all without introducing any fundamentally new primitive beyond keeping a row activated for a longer dwell time (t_{AggON}). This suggests that future DRAM defences must explicitly reason about Rowpress-like attacks, rather than focusing solely on traditional Rowhammer activation patterns.

5. Impact of In-DRAM Mitigations

Modern DRAM devices may incorporate mitigation mechanisms such as Target Row Refresh (TRR) and on-die Error Correcting Codes (ECC) to reduce the effectiveness of disturbance-based attacks. These mechanisms can significantly affect the practical exploitability of disturbance-induced bit flips on a computer system.

TRR attempts to detect aggressive row activations and refresh nearby victim rows to suppress the accumulation of disturbance effects. As a result, it effectively limits the disturbance budget that an attacker can accumulate within a given time window. However, TRR does not modify the intrinsic electrical properties of the memory cells themselves. Consequently, the resistance parameters extracted in our work (R_S and R_B) still characterize the underlying physical vulnerability of the DRAM cells. In practice, TRR reduces the accessible disturbance strength rather than altering the resistance parameters governing bit-flip behaviour.

Similarly, on-die ECC is designed to correct a limited number of bit errors before they are exposed to the system. While this mechanism can mask some disturbance-induced errors at the system level, our experimental setup measures raw bit flips directly using an FPGA-based controller. Therefore, the measured resistance parameters reflect the intrinsic device behavior rather than the effectiveness of error correction.

Overall, these mitigation mechanisms primarily affect the feasibility of system-level exploitation by constraining disturbance accumulation or masking a subset of errors. However, our model focuses on the characterization of the underlying cell-level vulnerability and provides a quantitative framework for reasoning about disturbance feasibility under different disturbance budgets and mitigation constraints.

6. Conclusion

This work introduces the **first** empirically verified cell-circuit-level model that explicitly explains DRAM

retention and read-disturbance failures from a security perspective. By modeling the p-well and bitline leakage paths together with pattern-dependent coupling, we extract a small set of physical parameters (R_S, R_B) that directly govern different attack surfaces. Through case studies and attack qualification, we show how these parameters can be systematically adopted to analyze and construct both *integrity* and *confidentiality* attacks, rather than treating bitflips as opaque symptoms. Implemented and validated on an FPGA-based platform with commodity DRAM modules, our framework provides a quantitative basis for assessing the exploitable surface of real devices to attacks and for comparing mechanisms such as Rowhammer and Rowpress under a unified lens. It offers practical tools to evaluate, prioritize, and ultimately strengthen future DRAM-based designs and defenses.

Acknowledgment

This research is supported by A*STAR, CISCO Systems (USA) Pte. Ltd and the National University of Singapore under its Cisco-NUS Accelerated Digital Economy Corporate Laboratory (Award I21001E0002).

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," *ACM SIGARCH Computer Architecture News*, vol. 42, no. 3, pp. 361–372, 2014.
- [2] H. Luo, A. Olgun, A. G. Yağlıkcı, Y. C. Tuğrul, S. Rhyner, M. B. Cavlak, J. Lindegger, M. Sadrosadati, and O. Mutlu, "Rowpress: Amplifying read disturbance in modern dram chips," in *Proceedings of the 50th Annual International Symposium on Computer Architecture*, 2023, pp. 1–18.
- [3] C. S. Lin, J. Qu, and G. Saileshwar, "Gpuhammer: Rowhammer attacks on gpu memories are practical," *arXiv preprint arXiv:2507.08166*, 2025.
- [4] L. Cojocar, K. Razavi, C. Giuffrida, and H. Bos, "Exploiting correcting codes: On the effectiveness of ecc memory against rowhammer attacks," in *2019 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2019, pp. 55–71.
- [5] A. Kwong, D. Genkin, D. Gruss, and Y. Yarom, "Rambleed: Reading bits in memory without accessing them," in *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2020, pp. 695–711.
- [6] 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.
- [7] A. G. Yağlıkcı, H. Luo, G. F. De Oliveira, A. Olgun, M. Patel, J. Park, H. Hassan, J. S. Kim, L. Orosa, and O. Mutlu, "Understanding rowhammer under reduced wordline voltage: An experimental study using real dram devices," in *2022 52nd Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2022, pp. 475–487.
- [8] L. Orosa, A. G. Yağlıkcı, H. Luo, A. Olgun, J. Park, H. Hassan, M. Patel, J. S. Kim, and O. Mutlu, "A deeper look into rowhammer's sensitivities: Experimental analysis of real dram chips and implications on future attacks and defenses," in *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*, 2021, pp. 1182–1197.
- [9] H. Luo, İ. E. Yüksel, A. Olgun, A. G. Yağlıkcı, and O. Mutlu, "Revisiting dram read disturbance: Identifying inconsistencies between experimental characterization and device-level studies," in *2025 IEEE 43rd VLSI Test Symposium (VTS)*. IEEE, 2025, pp. 1–8.
- [10] J. A. Halderman, S. D. Schoen, N. Heninger, W. Clarkson, W. Paul, J. A. Calandrino, A. J. Feldman, J. Appelbaum, and E. W. Felten, "Lest we remember: cold-boot attacks on encryption keys," *Communications of the ACM*, vol. 52, no. 5, pp. 91–98, 2009.
- [11] J. Juffinger, S. R. Neela, M. Heckel, L. Schwarz, F. Adamsky, and D. Gruss, "Presshammer: Rowhammer and rowpress without physical address information," in *International Conference on Detection of Intrusions and Malware, and Vulnerability Assessment*. Springer, 2024, pp. 460–479.
- [12] A. S. Rakin, M. H. I. Chowdhury, F. Yao, and D. Fan, "Deepsteal: Advanced model extractions leveraging efficient weight stealing in memories," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 1157–1174.
- [13] Y. Jiang, H. Zhu, D. Sullivan, X. Guo, X. Zhang, and Y. Jin, "Quantifying rowhammer vulnerability for dram security," in *2021 58th ACM/IEEE Design Automation Conference (DAC)*. IEEE, 2021, pp. 73–78.
- [14] A. J. Walker, S. Lee, and D. Beery, "On dram rowhammer and the physics of insecurity," *IEEE Transactions on Electron Devices*, vol. 68, no. 4, pp. 1400–1410, 2021.
- [15] Y. Chen, Z. Liu, Y. Yuan, S. Hu, T. Li, and S. Wang, "Compiled models, built-in exploits: Uncovering pervasive bit-flip attack surfaces in dnn executables," *arXiv preprint arXiv:2309.06223*, 2023.
- [16] P. Jattke, V. Van Der Veen, P. Frigo, S. Gunter, and K. Razavi, "Blacksmith: Scalable rowhammering in the frequency domain," in *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, 2022, pp. 716–734.
- [17] A. Kwong, D. Genkin, D. Gruss, and Y. Yarom, "Rambleed: Reading bits in memory without accessing them," in *2020 IEEE Symposium on Security and Privacy (SP)*, 2020, pp. 695–711.
- [18] J.-H. Chen, S.-C. Wong, and Y.-H. Wang, "An analytic three-terminal band-to-band tunneling model on gidl in mosfet," *IEEE Transactions on electron devices*, vol. 48, no. 7, pp. 1400–1405, 2001.
- [19] T. Sekiguchi, K. Itoh, T. Takahashi, M. Sugaya, H. Fujisawa, M. Nakamura, K. Kajigaya, and K. Kimura, "A low-impedance open-bitline array for multigigabit dram," *IEEE Journal of Solid-State Circuits*, vol. 37, no. 4, pp. 487–498, 2002.
- [20] A. Olgun, H. Hassan, A. G. Yağlıkcı, Y. C. Tuğrul, L. Orosa, H. Luo, M. Patel, O. Ergin, and O. Mutlu, "Dram bender: An extensible and versatile fpga-based infrastructure to easily test state-of-the-art dram chips," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 42, no. 12, pp. 5098–5112, 2023.
- [21] JEDEC, "JESD79-4C: DDR4 SDRAM Standard," JEDEC, Tech. Rep. JESD79-4C, 2020.

Appendix A. DRAM Organization

Figure 1 shows an overview of how DRAM is structurally organised. DRAM consists at the lowest level as a set of independent matrices - each known as a bank - of capacitive bit-cells, whereby each bit-cell contains a measurable charge, denoting a logical one or zero as data stored there. Each bit-cell is a structure known as a 1T1C (one transistor, one capacitor) structure, and consists of five key components: a Cell Capacitor, an Access Transistor, a Storage Node Contact (SNC), a Wordline (WL), a Bitline (BL) and a Bitline Contact (BLC). The cell capacitor stores the actual data bit as electrical charge, with its top plate connected to the access transistor through the SNC, which facilitates the electrical pathway between the two components. The access transistor acts as a switch controlled by the WL signal at its gate terminal, with one source/drain terminal connected to the capacitor via the SNC and the other connected to the BL through

the BLC. When the WL is activated during a read or write operation, the access transistor activates, creating a conductive path that allows charge to flow between the capacitor and the BL - either sensing the stored charge (read) or storing a new charge (write). The BLC serves as the connection point between the transistor and the vertical BL that runs through all cells in a column, enabling analogue circuits called sense amplifiers to detect the small voltage changes during reads or to drive new values during writes, which in turn acts as a bridge to transfer data to/from the main memory controller to then be used by CPU. Analogue circuits known as row and column decoders select specific bit-cells to read/write to based on row and column addresses supplied by a main memory controller. Multiple banks (typically 8 to 16) are organised into a single DRAM chip. The smallest form factor seen on any given computer system would be an individual DRAM chip (typically the case with embedded systems), however in commodity/high-powered systems, multiple DRAM chips are found organised across a printed circuit board called a Dual In-Line Memory Module (DIMM). These enable simultaneous chip control across a set of ranks, where each share control signals, yet have separate data pathways. This parallel operation enables the memory controller to interleave read/write accesses across individual chips/banks, allowing DRAM to maintain a very high bandwidth and efficiency, even when individual banks are busy handling operations across multiple clock cycles.

Appendix B.

Parameter Characterization

B.1. R_S Measurement under Retention

To characterize a DRAM cell, we measure the effective leakage resistance R_S under different configurations. The procedure is based on retention profiling: the victim cell is initialized to logic 1, refresh is disabled for a sequence of intervals $\{t_1, t_2, \dots\}$, and the cell is read out to detect the earliest flip time t_{flip} . Given the storage capacitance C_S and the sense amplifier threshold V_{flip} , the effective resistance is computed as

$$R_S(N) = \frac{t_{\text{flip}}}{C_S \cdot \ln\left(\frac{V_{DD}}{V_{\text{flip}}}\right)}, \quad (23)$$

where N denotes the neighbor pattern applied to the eight adjacent cells. The baseline resistance $R_S(\text{all-1})$ is obtained when all neighbors are written with logic '1'. The complete procedure is summarized in Algorithm 2.

B.2. Measurement of R_S and R_B under Read Disturbance

To quantify the effect of subthreshold conduction under read disturb stress (Rowhammer/Rowpress), we adopt a two-pattern methodology that separates the baseline leakage path from its pattern-dependent extension. In the 111 configuration (victim and aggressors set to logic 1), the bitline is restored to V_{DD} after each read, leaving essentially no DC bias across the access transistor. Although wordline coupling may slightly perturb the victim,

the absence of a driving potential suppresses additional subthreshold current. Let HC_{111} be the minimal hammer count that produces the target number of flips (e.g., 5000) at an effective read disturb rate f_{RD} ; the corresponding flip time is $T_{111} = HC_{111}/f_{RD}$, and the baseline effective resistance is

$$R_{\text{eff}} \approx R_S^{\text{RD}} = \frac{T_{111}}{C_S \ln\left(\frac{V_{DD}}{V_{\text{flip}}}\right)}. \quad (24)$$

In contrast, the 010 configuration (victim at logic 1, aggressors at logic 0) establishes a non-zero bias between the storage node and bitline, biasing the access transistor into the subthreshold region and enabling an additional current component. Let HC_{010} be the hammer count needed to reach the same flip target, and $T_{010} = HC_{010}/f_{RD}$ the corresponding flip time. Defining

$$E \triangleq \left(\frac{V_{DD}}{V_{\text{flip}}}\right)^{T_{010}/T_{111}}, \quad (25)$$

and modelling the discharge as a combination of baseline leakage and a constant subthreshold term, we obtain the closed-form estimate of the subthreshold factor

$$A = \frac{1}{R_{\text{eff}}} \frac{V_{DD} - E V_{\text{flip}}}{E - 1}, \quad (26)$$

which depends only on measured hammer counts (HC_{111}, HC_{010}) and known device parameters ($f_{RD}, C_S, V_{DD}, V_{\text{flip}}$). The parameter A has units of current and quantifies the extra leakage due to subthreshold conduction; in our model, it directly leads to the bitline-side effective resistance $R_B^{\text{RD}} \approx V_{TH}/A$.

Open Science

To support reproducibility and foster further research, we make all artifacts of this work publicly available at https://osf.io/9gdyw/?view_only=eb552500a9f445d8a56611dc7ea36d25. This includes our FPGA-based measurement framework, implementation of the proposed models, and security evaluation scripts.

Ethical Considerations

Our experiments were performed on commodity DRAM modules in controlled lab settings, without involving production systems, user data, or human subjects. The attacks are evaluated solely for research purposes. The goal is to inform the community and vendors about security implications to enable more secure memory designs.