Reducing the Silicon Area Overhead of Counter-Based Rowhammer Mitigations

Loïc France, Florent Bruguier, David Novo, Maria Mushtaq, and Pascal Benoit

Abstract

Modern computer memories have shown to have reliability issues. The main memory is the target of a security threat called Rowhammer, which causes bit flips in adjacent victim cells of repeatedly activated aggressor rows [1]. This issue is becoming more important as DRAM technology scales down, with the required aggressor activations to corrupt a victim going from 130k for DDR3 [1] to around 10k for the most recent LPDDR4 memories [2]. Numerous countermeasures have been proposed, implemented either in software [3], [4] or in hardware [1], [5]–[10]. Among the hardware-based proposals, some rely on probability, randomly refreshing neighbors of activated rows [1], [5], [6], while others rely on row activation counters to detect aggressor rows before acting to prevent the corruption [7]–[9]. Counter-based hardware mitigation proposals offer the lowest performance overhead, as the mechanism only acts when an aggressor is detected and does not disturb the system for harmless applications. However, they require a lot of counters to track row activations. Considering the unrealistic amount of counters needed to track every rows, those mitigation exploit different most-frequent-elements detection algorithms to reduce the number of counters needed while keeping a complete protection and a minimal false positive rate. Most of them offer a bank-level attack detection, with a separate set of counters for each bank. In this talk, We will show you that by changing the counting granularity from bank-level to rank-level, we can further reduce the total required number for counters from 20% for DDR3 to 70% for DDR5, thus reducing the silicon area and energy overheads of such mitigations.

Index Terms

Security, Rowhammer, DRAM.

References

[1] Y. Kim *et al.*, “Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors,” in *ISCA*, 2014.
[2] J. S. Kim *et al.*, “Revisiting rowhammer: An experimental analysis of modern dram devices and mitigation techniques,” in *ISCA*, 2020.
[3] A. Chakraborty *et al.*, “Deep learning based diagnostics for rowhammer protection of DRAM chips,” in *ATS*, 2019.
[4] M. Alam *et al.*, “Performance counters to rescue: A machine learning based safeguard against micro-architectural side-channel-attacks.” *IACR Cryptology ePrint Archive*, 2017.
[5] J. M. You and J.-S. Yang, “MRLoc: Mitigating row-hammering based on memory locality,” in *DAC*, 2019.
[6] M. Son *et al.*, “Making dram stronger against row hammering,” in *DAC*, 2017.
[7] Y. Park *et al.*, “Graphene: Strong yet lightweight row hammer protection,” in *MICRO*, 2020.
[8] A. G. Yağlikçi *et al.*, “Blockhammer: Preventing rowhammer at low cost by blacklisting rapidly-accessed dram rows,” in *HPCA*, 2021.
[9] E. Lee *et al.*, “Twice: Preventing row-hammering by exploiting time window counters,” in *ISCA*, 2019.
[10] L. France *et al.*, “Vulnerability assessment of the rowhammer attack using machine learning and the gem5 simulator - work in progress,” in *SaT-CPS*, 2021.

- *Loïc France, Florent Bruguier, David Novo, and Pascal Benoit are with LIRMM, University of Montpellier, CNRS, France.*
  E-mail: {firstname}.{lastname}@lirmm.fr
- *Maria Mushtaq is with LTCI, Télécom Paris, Institut Technologique de Paris, France.*
  E-mail: maria.mushtaq@telecom-paris.fr