

SPOILER-GUARD: Gating Latency Effects of Memory Accesses through Randomized Dependency Prediction

Gayathri Subramanian, Girinath P, Nitya Ranganathan, Kamakoti Veezhinathan, Gopalakrishnan Srinivasan

Department of Computer Science and Engineering, Indian Institute of Technology Madras

{gayathri, cs22s021, kama, sgopal}@cse.iitm.ac.in, nitya_rise@icsrps.iitm.ac.in

Abstract—Modern microprocessors depend on speculative execution, creating vulnerabilities that enable transient execution attacks. Prior defenses target speculative data leakage but overlook false dependencies from partial address aliasing, where repeated squash and reissue events increase the load-store latency, which is exploited by the SPOILER attack. We present SPOILER-GUARD, a hardware defense that obfuscates speculative dependency resolution by dynamically randomizing the physical address bits used for load-store comparisons and tagging store entries to prevent latency-amplifying misspeculations. Implemented in gem5 and evaluated with SPEC 2017, SPOILER-GUARD reduces misspeculation to 0.0004 percent and improves integer and floating-point performance by 2.12 and 2.87 percent. HDL synthesis with Synopsys Design Compiler at 14 nm node demonstrates minimal overheads - 69 ps latency in critical path, 0.064 square millimeter in area, and 5.863 mW in power.

Index Terms—Secure Architecture, Memory Dependence Prediction, SPOILER, Address Aliasing, Transient Execution Attack

I. INTRODUCTION

Modern out-of-order microprocessors use speculative execution to hide latency by predicting branch outcomes and memory dependencies. Misspeculation rolls back architectural state, but microarchitectural state such as cache contents remain altered and are typically exploited by transient execution attacks [1]–[3]. Spectre [4] leverages branch misprediction while Meltdown [5] exploits deferred exception handling. Speculation also affects memory operations, where loads may bypass preceding stores if memory-dependence predictors within the pipeline do not predict the dependency [6]–[8].

SPOILER [9] exploits false dependencies induced via partially aliased accesses, revealing fine-grained physical address (PA) information. This accelerates virtual-to-physical reverse-engineering by $256\times$ and eviction-set construction by $4096\times$, amplifying attacks such as Prime+Probe [10] and Rowhammer [11]. Partial address checks in the Memory Order Buffer [12]–[14] can trigger repeated squash-and-reissue cycles, and the partial physical address bits stored in the Store Address Buffer (SAB) make aliasing-induced latency observable.

In this work, we assume an unprivileged attacker on the same physical core, who is capable of allocating memory, crafting load-store sequences, and performing fine-grained timing via `rdtsc` instruction [15]. Our key observation is that SPOILER succeeds because the dependence predictor performs static physical address comparisons, producing repeatable squash-and-reissue cycles. Existing defenses [16]–[21] target speculative data leakage by blocking or sanitizing unsafe speculative loads. However, mechanisms like NDA [22] and SSBD [23] incur high overhead, up to 125% in strict modes and 6.6–10.7% in practical ones. SPOILER-ALERT [24] embeds a cuckoo filter

in the Store Buffer to track store addresses and detect load-store aliasing with 99.99% accuracy, but cannot block leakage. Intel guidelines [25] ignore structural leakage whereas speculative-interference studies [26], [27] fail to address contention effects, leaving SPOILER exploitable. SPOILER-GUARD obfuscates dependency resolution by randomizing the physical address bits used for speculative checks and tagging store entries to prevent repeated replays. The key contributions of our work are:

- SPOILER-GUARD breaks the correlation between partial address aliasing and microarchitectural timing by dynamically randomizing address bits and tagging misspeculated stores to avoid further misspeculation.
- We implement and evaluate SPOILER-GUARD in gem5 using SPOILER binaries and SPEC 2017 benchmark suite to demonstrate security and performance.
- SPOILER-GUARD lowers misspeculation to 0.0004% and improves integer and floating-point performance by 2.12% and 2.87% respectively, with negligible area (0.064 mm²), power (5.863 mW), and latency (69 ps) overhead.

II. PROPOSED DESIGN




Fig. 1. Overview of the proposed SPOILER-GUARD defense.

In SPOILER, the attacker fills the Store Buffer with aliased stores and issues probe loads with 1 MB aliasing, causing repeated squash-and-reissue cycles. SPOILER-GUARD breaks this correlation by randomizing the partial address bits used for dependency resolution. The fixed 8-bit comparison is replaced by comparison of dynamically selected 12-bits (determined by the *mask*), reducing the attacker’s probability of inferring the aliasing pattern from $1/256$ to $1/4096$. In simulation, the mask is generated using a Mersenne Twister Pseudo Random Number Generator (or PRNG) seeded with high-entropy operating system sources [28]. In hardware, equivalent randomness can be realized using a True Random Number Generator (TRNG) with entropy conditioning and expansion through a cryptographically secure PRNG [29].

SPOILER-GUARD adds three microarchitectural changes: (1) SAB entries gain a PC tag and spoiler-vulnerability flag, (2) partial physical address field is widened to 12 bits, and (3) LSQ integrates mask-generation logic with a remasking controller. An initial mask is set at system setup, with remasking triggered

Gayathri Subramanian is the corresponding author.




Fig. 2. Average latency of the speculatively issued malicious load for M1, M2, and M3 configurations.

on misspeculation. When a load is executed and encounters a loosenet hit but a finenet miss, the dependency predictor flags a potential dependency and performs a masked 12-bit partial physical address comparison. A partial hit triggers speculative store-to-load forwarding, and the corresponding SQ entry is tagged with the load’s PC. After resolution of the full physical address, the predictor detects misspeculation, sets the vulnerability bit, triggers remasking, and squashes the load, thus preventing repeated aliasing and making any residual latency indistinguishable from normal speculation (see Fig. 1).

III. EVALUATION

A. Setup and Evaluation

We use the gem5 architectural simulator [30] to model a single-core and single-threaded microarchitecture (downsizing Intel i7-8700), allowing predictable speculative dependency behavior and detailed observation of LSQ interactions. All simulations are run in a full-system Ubuntu Linux x86 environment, capturing realistic system-level interactions to evaluate defense effectiveness. We compare three configurations: M1 - baseline with virtual address (VA) forwarding, M2 - baseline with SPOILER vulnerability, and M3 - SPOILER-GUARD.

B. Security Analysis

The three models are evaluated by their obfuscation of partially aliased load latency and their frequency of false positives from speculative forwarding. A standard attack binary, based on [9], allocates a 4KB buffer across 1024 pages, filling the store buffer to induce 4KB and 1MB aliasing, and measuring average latency over 100 rounds per page. This reliably reproduced the attack on Intel i7-8700 and newer processors, including the i9 systems, revealing higher-order physical address bits.

In M1, speculative load latency is uniform with minimal 4KB aliasing delays. In M2, 1 MB aliasing triggers repeated squash-and-reissue (about 255,941 SPOILER-violations and 255,924 attacker-induced stalls), resulting in the high observable latency for these loads. In M3, SPOILER-violations fall to 14 and attacker-induced stalls to 1, with latency of target load matching M1 aside from the unavoidable 4 KB aliasing (see Fig. 2).

In M3, the attacker’s aliasing probability is $P_{M3}(alias) = 2^{-12}$ per trial, and remasking on each misspeculation makes attempts statistically independent, collapsing SPOILER’s required temporal correlation. This neutralizes partial physical-address aliasing effects, making speculation indistinguishable from purely VA-based behavior, thereby eliminating the leakage channel and making the defense secure as well as practical.

C. Performance Analysis

We evaluated performance using the SPEC CPU2017 suite (compiled using GCC 7.4.0), simulating an average of one billion committed instructions per benchmark. Performance was measured using cycles-per-instruction (CPI), with speedup computed relative to baseline models M1 and M2. Memory and compute intensive workloads were analyzed separately to assess differential responsiveness. SPOILER-GUARD shows overall speedups of 2.87% for floating-point workloads and 2.12% for integer workloads. Memory-bound floating-point workloads gain 3.23%, while memory-bound integer workloads gain 1.13%. The performance of compute-bound benchmarks remained near baseline. Non-classified workloads benefit more substantially, achieving 7.03% improvement for floating-point and 3.46% for integer benchmarks (see Fig. 3).




Fig. 3. Speedup achieved by SPOILER-GUARD for SPEC2017 integer and floating-point benchmark suite.

D. Storage Overhead

The enhanced SAB expands the partial physical address field from 8 to 12 bits and introduces two metadata components: a 48-bit load PC tracker and a 1-bit spoiler-vulnerability flag. This adds 53 bits per entry, or 2968 bits (≈ 0.36 KB) for a 56-entry SAB. Relative to a typical 32 KB L1 data cache, the SAB incurs less than 0.2% area overhead while providing significant microarchitectural security enhancements.

E. Power, Area, and Timing Overhead

To evaluate SPOILER-GUARD’s hardware impact, the partial address comparison and SAB remasking logic were implemented in RTL and synthesized using the Synopsys Design Compiler (DC) with a 14 nm standard-cell library. Compared to M2, M3 incurs minimal overheads of 0.064 mm^2 (less than 0.8% of a Skylake core’s 8.73 mm^2 footprint synthesized in a comparable 14 nm process [31]), 5.863 mW in power, and 69 ps along the critical path (negligible impact), demonstrating that SPOILER-GUARD adds minimal hardware cost while securing speculative dependency resolution.

IV. CONCLUSION AND FUTURE WORK

We proposed SPOILER-GUARD, a hardware defense that mitigates SPOILER by obfuscating partial address aliasing in LSQ. Applicable to all affected Intel Core microarchitectures, it provides robust security and speedups for memory-intensive workloads. Logic synthesis confirmed minimal area, power, and timing overheads, thereby demonstrating an efficient and practical mitigation. Future work may extend dynamic dependency resolution and store tagging to other speculative attacks.

REFERENCES

- [1] R. J. Colvin and K. Winter, "An Abstract Semantics of Speculative Execution for Reasoning About Security Vulnerabilities," In *Formal Methods. FM 2019 International Workshops, Lecture Notes in Computer Science*, vol. 12233, 2020, pp. 323-341.
- [2] Y. Jin, P. Qiu, C. Wang, Y. Yang, D. Wang, and G. Qu, "Timing the Transient Execution: A New Side-Channel Attack on Intel CPUs," *arXiv preprint arXiv:2304.10877*, 2023. [Online]. Available: <https://arxiv.org/abs/2304.10877>
- [3] J. Van Bulck, M. Minkin, O. Weisse, D. Genkin, B. Kasikci, F. Piessens, M. Silberstein, T. F. Wenisch, Y. Yarom, and R. Strackx, "Foreshadow: Extracting the Keys to the Intel SGX Kingdom with Transient Out-of-Order Execution," In *Proc. USENIX Security Symp.*, 2018, pp. 991-1008.
- [4] P. Kocher et al., "Spectre attacks: Exploiting speculative execution," In *Proc. IEEE Symp. Security Privacy (SP)*, 2019, pp. 1-19.
- [5] M. Lipp et al., "Meltdown: Reading kernel memory from user space," In *Proc. USENIX Security Symp.*, 2018, pp. 973-990.
- [6] G. Z. Chrysos and J. S. Emer, "Memory dependence prediction using store sets," In *Proc. 25th Int. Symp. Comput. Archit. (ISCA)*, 1998, pp. 142-153.
- [7] R. E. Kessler, "The Alpha 21264 microprocessor," *IEEE Micro*, vol. 19, no. 2, Apr., pp. 24-36, 1999.
- [8] A. Yoaz, M. Erez, R. Ronen, and S. Jourdan, "Speculation techniques for improving load related instruction scheduling," in *Proc. 26th Int. Symp. Comput. Archit. (ISCA)*, 1999, pp. 42-53.
- [9] S. Islam et al., "SPOILER: Speculative load hazards boost Rowhammer and cache attacks," In *Proc. USENIX Security Symp.*, 2019, pp. 621-637.
- [10] D. A. Osvik, A. Shamir, and E. Tromer, "Cache Attacks and Countermeasures: The Case of AES," In *Lecture Notes in Computer Science*, vol. 3897, Feb. 2006, pp. 1-20.
- [11] Y. Kim et al., "Flipping bits in memory without accessing them: An experimental study of DRAM disturbance errors," *SIGARCH Comput. Archit. News*, vol. 42, no. 3, Jun., pp. 361-372, 2014.
- [12] J. M. Abramson et al., "Method and apparatus for dispatching and executing a load operation to memory," U. S. Patent 5,717,882, Feb. 10, 1998.
- [13] S. Hilv, Z. Zhang, and P. Hammarlund, "Resolving false dependencies of speculative load instructions," U. S. Patent 7,603,527 B2, Oct. 13, 2009.
- [14] S. Kosinski et al., "Store forwarding for data caches," U. S. Patent 9,507,725 B2, Nov. 29, 2016.
- [15] Intel, "Intel 64 and IA-32 architectures software developer's manual, vol. 3B: System programming guide, part 2," Intel, Santa Clara, CA, USA, 2023. [Online]. Available: <https://www.intel.com/content/www/us/en/developer/articles/technical/intel-sdm.html>
- [16] K. N. Khasawneh et al., "SafeSpec: Banishing the Spectre of a Meltdown with Leakage-Free Speculation," In *56th ACM/IEEE Design Automation Conference (DAC)*, 2019, pp. 1-6.
- [17] J. Yu, M. Yan, A. Khyzha, A. Morrison, J. Torrellas and C. W. Fletcher, "Speculative taint tracking (STT): A comprehensive protection for speculatively accessed data," In *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchit.*, 2019, pp. 954-968.
- [18] M. Yan, J. Choi, D. Skarlatos, A. Morrison, C. W. Fletcher, and J. Torrellas, "Invisispec: Making speculative execution invisible in the cache hierarchy," In *MICRO*, 2018, pp. 428-441.
- [19] J. Fustos, F. Farschi, and H. Yun, "Spectreguard: An efficient data-centric defense mechanism against spectre attacks," In *DAC*, 2019, pp. 1-6.
- [20] O. Oleksenko, B. Trach, T. Reiher, M. Silberstein, and C. Fetzer, "You Shall Not Bypass: Employing data dependencies to prevent Bounds Check Bypass," *arXiv.org*, 2018. [Online]. Available: <https://arxiv.org/abs/1805.08506>
- [21] K. Barber, A. Bacha, L. Zhou, Y. Zhang, and R. Teodorescu, "SpecShield: Shielding speculative data from microarchitectural covert channels during speculation," In *Proc. Int. Conf. Parallel Archit. Compilation Tech. (PACT)*, 2019, pp. 11-23.
- [22] O. Weisse, I. Neal, K. Loughlin, T. F. Wenisch, and B. Kasikci, "NDA: Preventing speculative execution attacks at their source," In *Proc. 52nd Annu. IEEE/ACM Int. Symp. Microarchit. (MICRO)*, 2019, pp. 572-586.
- [23] Intel, "Speculative store bypass / CVE-2018-3639 / INTEL-SA-00115," Intel Security Advisory, May 21, 2018. [Online]. Available: <https://www.intel.com/content/www/us/en/security-center/advisory/intel-sa-00115.html>
- [24] J. Cui, Y. Yin, C. Chen, and J. Zhang, "SPOILER-ALERT: Detecting SPOILER Attacks Using a Cuckoo Filter," In *Proc. Design, Autom. Test Eur. Conf. Exhib. (DATE)*, 2023, pp. 1-6.
- [25] Intel, "Speculative Execution Side Channel Mitigations," Intel Developer Guidance, Jan. 2018. [Online]. Available: <https://www.intel.com/content/www/us/en/developer/articles/technical/software-security-guidance/technical-documentation.html>
- [26] L. Culbertson, "Addressing new research for side-channel analysis," Intel Newsroom, Mar. 2018. [Online]. Available: <https://newsroom.intel.com/editorials/addressing-new-research-for-side-channel-analysis/>
- [27] M. Behnia et al., "Speculative interference attacks: Breaking invisible speculation schemes," In *Proc. 26th ACM Int. Conf. Archit. Support Program. Lang. Operating Syst. (ASPLOS)*, 2021, pp. 1046-1060.
- [28] R. Ostertág and M. Stanek, "Entropy assessment of Windows performance counters for PRNG seeding," *arXiv:1311.3139*, Nov. 2013. [Online]. Available: <https://arxiv.org/abs/1311.3139>
- [29] V. Piscopo, A. Dolmeta, M. Mirigaldi, M. Martina, and G. Masera, "A high-entropy true random number generator with Keccak conditioning for FPGA," *Sensors*, vol. 25, no. 6, Art. no. 1678, Mar., pp. 1678, 2025.
- [30] N. L. Binkert et al., "The gem5 simulator," *SIGARCH Comput. Archit. News*, vol. 39, no. 2, May, pp. 1-7, 2011.
- [31] "Skylake (Client) - Microarchitectures - Intel," WikiChip. [Online]. Available: [https://en.wikichip.org/wiki/intel/microarchitectures/skylake_\(client\)](https://en.wikichip.org/wiki/intel/microarchitectures/skylake_(client)). [Accessed: Sep. 2, 2025].