Microarchitectural signals analysis platform for the implementation of Hardware Security Counters

Lucas Georget[1,2], Vincent Migliore[1], Vincent Nicomette[1], Frédéric Silvi[2] and Arthur Villard[2]*

[1]LAAS-CNRS, Toulouse, France
[2]EDF R&D, Paris-Saclay, France

Abstract

Detecting malicious software or hardware behavior during the operation of a computer system requires observables from one or more abstraction layers of the system. This abstraction, however, tends to limit the ability to detect behavioral deviations, especially for attack classes that exploit vulnerabilities very close to the target hardware. Conversely, too low a level of abstraction tends to significantly increase the complexity of the system model, and therefore poses a number of difficulties for the extraction and selection of relevant observables for a given class of attack.

Hardware performance counters in particular have been used as an indirect means of observing microarchitecture behavior and detecting software attempting to exploit hardware vulnerabilities. In order to improve the various detection methods, we propose the construction of hardware metrics designed from the outset for security, by studying the correlation between signals from the microarchitecture and the various classes of attack in the literature, targeting both conventional IT and industrial OT systems. By extension, this work aims to detect attacks originating from hardware Trojans, the latter having the effect of changing the behavior of a given microarchitecture.

Extended abstract

The detection of low-level attacks, especially hardware malware or software malware targeting microarchitectural vulnerabilities, is quite complex. Hardware Performance counters (HPC) for example, can be used to trace hardware behavior and divert it for security purposes. However, the further down we go, the more complex it becomes to use observables. No hardware metrics were originally designed for security purpose. This is why it is essential to study the impact of attacks on microarchitectural signals in order to build specific counters that could be used to reference the internal state of our machines and detect attacks at microarchitectural level, as well as hardware Trojans, on both traditional and industrial equipment.

HPCs have been used in different research works for security purposes. In the context of a fleet of IoT devices executing the same software, early work [1] aimed to identify deviations in the behavior of one or more devices compared to the others. For that purpose, an intrusion system was designed, based on the analysis through machine learning algorithms [2] of the HPC collected on all the devices in order to exhibit outliers. Other methods based on learning algorithms have been used on HPCs to detect temporal attacks on processor caches, as reported in [3]. Other work [4] has led to the development of a monitoring and tracing system for lightweight systems against radio attacks.

However, to reach a finer level of granularity and detect even more subtle attacks, we argue that it is necessary to analyze various hardware signals and try to identify which signals are relevant to detect some specific class of attacks. Some specific platforms, mainly based on FPGA, are necessary to carry out such experiments. But, to the best of our knowledge, few solutions are currently available for that purpose. Only debugging solutions such as Xilinx ChipScope and Intel SignalTap are available on the market. Mao et al [5] proposed a RocketChip instrumentation in Scala to provide such a solution, with strong constraints on the bandwidth for data collection.

The purpose of our research work is thus to design and implement a hardware platform, generic enough so that we can observe and exports various microarchitectural signals from the board (processor and peripherals) to perform a concrete and complete analysis and correlate this data.

In this work, we propose a flexible solution that captures a vector of internal signals from a System on Chip's microarchitecture with no impact on the running software. The solution requires:

- A reconfigurable target system with an integrated logic analyzer for extracting microarchitectural signals
- A host system to collect these data, with good storage capacity and bandwidth with the target
- A high-performance system (perhaps the same as the host) to further process and analyze the

*Corresponding author: lucas.georget@laas.fr

Received: February 7, 2025, **Published:** March 7, 2025

data.

The solution has been implemented on the high-end FPGA-based plateform Alveo U50-DD connected to a host computer through PCIe, for high data transfer speed capacity. For a fast and portable SoC deployment into the programmable logic, we used the emerging framework LiteX [6].

For the embedded FPGA logic analyzer part, LiteScope [7] was integrated to provide an initial observation of microarchitectural signals from, for example, the instruction and data buses in the CPU. It is a small footprint and configurable tool able to capture signals in real time, with limited resources and without any perturbation of the system.

To finalize the platform for our use cases, we still need to :

- Use custom OS images to run benign and malicious program benchmarks
- Modify LiteScope for continuous monitoring, as it currently only fills a predefined storage space and then terminates, to cover 1 nano second of system usage
- Segmentation and correlation of signal readings by process, so that we can extract only the signals associated to a particular process (as part of a single-core, uninterruptible architecture)

For security purposes, we have identified three possible use cases:

- Software attacks such as Spectre, Meltdown and Rowhammer, as well as Cache Side-Channel and Return-Oriented Programming Attacks
- Hardware attacks: detection of automatic Trojan insertion at processor and peripheral levels
- Reverse engineering of CPU behavior at microarchitectural level

All the metrics collected during these different use cases will be stored and analyzed via Machine Learning techniques. The objective is to exhibit common detection criteria that will be used to design relevant hardware security counters.

References

[1] Malcolm Bourdon et al. "Hardware-Performance-Counters-based anomaly detection in massively deployed smart industrial devices". In: *2020 IEEE 19th International Symposium on Network Computing and Applications (NCA)*. 2020, pp. 1–8. DOI: 10.1109/NCA51143.2020.9306726.

[2] Nikolaos Foivos Polychronou et al. "MaDMAN: Detection of Software Attacks Targeting Hardware Vulnerabilities". In: *2021 24th Euromicro Conference on Digital System Design (DSD)*. 2021, pp. 355–362. DOI: 10.1109/DSD53832.2021.00060.

[3] Maria Mushtaq. "Software-based Detection and Mitigation of Microarchitectural Attacks on Intel's x86 Architecture". Theses. Université de Bretagne Sud, Sept. 2019. URL: https://theses.hal.science/tel-02988980.

[4] Mohamed El-Bouazzati. "A Lightweight Host-based Intrusion Detection System using a Hardware-Assisted Monitor to detect Wireless Attacks Targeting Constrained IoT Devices". Theses. Université de Bretagne Sud, Dec. 2023. URL: https://cnrs.hal.science/tel-04612764.

[5] Yuxiao Mao, Vincent Migliore, and Vincent Nicomette. "MATANA: A Reconfigurable Framework for Runtime Attack Detection Based on the Analysis of Microarchitectural Signals". In: *Applied Sciences* 12.3 (2022). ISSN: 2076-3417. DOI: 10.3390/app12031452. URL: https://www.mdpi.com/2076-3417/12/3/1452.

[6] Florent Kermarrec et al. *LiteX: an open-source SoC builder and library based on Migen Python DSL*. 2020. arXiv: 2005.02506 [cs.AR]. URL: https://arxiv.org/abs/2005.02506.

[7] EnjoyDigital. *LiteScope - A small footprint and configurable embedded FPGA logic analyzer*. 2015. URL: https://github.com/enjoy-digital/litescope.