



**HAL**  
open science

# Detecting Hardware Trojans in Microprocessors via Hardware Error Correction Code-based Modules

Alessandro Palumbo, Ruben Salvador

## ► To cite this version:

Alessandro Palumbo, Ruben Salvador. Detecting Hardware Trojans in Microprocessors via Hardware Error Correction Code-based Modules. 2025 IEEE 31st International Symposium on On-Line Testing and Robust System Design (IOLTS), Jul 2025, Naples, Italy. pp.1-7, <10.1109/IOLTS65288.2025.11116960>. <hal-05055204>

**HAL Id: hal-05055204**

**<https://hal.science/hal-05055204v1>**

Submitted on 4 May 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

# Detecting Hardware Trojans in Microprocessors via Hardware Error Correction Code-based Modules

Alessandro Palumbo

CentraleSupélec, Inria, CNRS, IRISA, France

alessandro.palumbo@inria.fr

Ruben Salvador

CentraleSupélec, Inria, CNRS, IRISA, France

ruben.salvador@inria.fr

**Abstract**—Software-exploitable Hardware Trojans (HTs) enable attackers to execute unauthorized software or gain illicit access to privileged operations. This manuscript introduces a hardware-based methodology for detecting runtime HT activations using Error Correction Codes (ECCs) on a RISC-V microprocessor. Specifically, it focuses on HTs that inject malicious instructions, disrupting the normal execution flow by triggering unauthorized programs. To counter this threat, the manuscript introduces a Hardware Security Checker (HSC) leveraging Hamming Single Error Correction (HSEC) architectures for effective HT detection. Experimental results demonstrate that the proposed solution achieves a 100% detection rate for potential HT activations, with no false positives or undetected attacks. The implementation incurs minimal overhead, requiring only 72 #LUTs, 24 #FFs, and 0.5 #BRAM while maintaining the microprocessor’s original operating frequency and introducing no additional time delay.

**Index Terms**—Error Correction Codes, Hardware Security, Hardware Trojans, Microprocessor-based System, RISC-V.

## I. INTRODUCTION AND RELATED WORK

The ongoing drive to lower production costs and shorter time-to-market has led to the globalization of design and manufacturing processes for modern integrated circuits (ICs) [1]. Design tasks for components and subsystems are often outsourced, third-party intellectual property cores (3PIPs) are commonly purchased, and the final chips are fabricated by external foundries [2]. While this approach has effectively reduced both design time and cost, it has also led to a significant reduction in trust regarding the integrity of the produced ICs [3].

Ensuring trust among all participants in a globalized supply chain has become an increasingly complex challenge that is now virtually unattainable. As a result, various threats may emerge, including the overproduction or counterfeiting of integrated circuits (ICs), the improper use of licenses, and the introduction of Software-Exploitable Hardware Trojans (HTs) [4], [5].

HTs are subtle, hard-to-detect alterations in a system designed to remain hidden for most of the time, only activating under certain (often rare) conditions to disrupt system behaviour by, e.g., performing attacks by accessing protected memory locations or stealing sensitive data [6], [7]. Malicious actors can introduce such modifications, whether they are third-party IP providers, employees, CAD tools, mask suppliers, or silicon foundries.

Initially, HTs were considered more of an academic problem due to the challenges associated with their real-world implementation, which limited their potential use by attackers. However, recent research has demonstrated that they can be embedded in commercial microprocessors, enabling attackers to execute their own malicious software, modify running code, or gain root privileges [8]–[10]. For instance, a few years ago, the *Rosenbridge* backdoor was discovered in a Via Technologies C3 processor [11]. Attackers could activate the backdoor through a specific instruction sequence, granting access to supervisor mode.

Several techniques have been developed for detecting HTs during the design phase, focusing on analyzing the system at the circuit level before deployment. These methods include logic testing [12], program run code obfuscation [13], [14], formal property verification [15], side-channel analysis [16], structural and behavioral analysis [17], [18], and machine learning [19], [20]. However, due to the inherently stealthy nature of HTs, detecting them before deployment is extremely challenging. This has led to the emergence of the *Design-for-Trust* paradigm, which emphasizes *system-level* techniques to construct trusted systems from potentially untrusted components [21], [22]. Additionally, HTs can be countered by enabling trusted software execution on systems that include untrusted microprocessor-based components [23], [24].

This paper proposes a system-level solution for detecting HT activations at runtime in microprocessor-based systems, which could compel the microprocessor to execute malicious code. The presented approach integrates a Hardware Security Checker (HSC) composed of two Hardware Security Modules (HSMs) between the microprocessor and the main memory to monitor fetching activity. The HSC is *configured* during the program installation in memory, utilizing information about the program’s instructions and memory locations.

During the program execution, at runtime, the HSC verifies that the correct instructions are fetched from the correct memory addresses. This mechanism allows the HSC to detect HT activation at runtime, potentially compromising the microprocessor, the memory, or the bus. The proposed solution operates transparently without interrupting or interfering with the normal execution of the microprocessor. The HSC functions passively, ensuring the integrity of the protected program while remaining non-intrusive.

The HSC has been tested on a case study system based

on a RISC-V microprocessor implemented on an FPGA, running a set of software benchmarks. The RISC-V ISA could implement the Physical Memory Protection (PMP) mechanism to control memory access and protect against unauthorized instruction execution or data manipulation [25], [26]. However, PMP is not immune to attacks that exploit and modify its configuration [25], [27]. The results showed that the proposed HSC can detect 100% potential HT activations without false alarms or undetected attacks. We observed an area overhead of less than 1% in terms of #LUTs and #FFs, with 0.5 #BRAM required and no reduction in operating frequency.

Some system-level design-for-trust methodologies are reported in [23], [24], [28], [29]. Some solutions [23], [24] assume the microprocessor is untrusted while the memory is trusted. In [23], the checker monitors the validity of opcodes and control signals and the number of clock cycles required to execute an instruction. Meanwhile, the solution in [24] tracks the microprocessor’s liveness and its privilege mode. However, neither of these solutions is designed to detect HTs that force the CPU to execute legitimate instructions without altering the privilege mode. Specifically, they do not detect scenarios where the microprocessor is compelled to execute an unintended program or access unauthorized memory locations, which our solution addresses.

Similar works [28]–[30] to this paper propose checkers to detect HT activations affecting the microprocessor, bus, and memory. The solution in [28] incurs significant area overhead and still experiences a non-zero false negative rate. The solution in [29] improves the overhead but leaves vulnerabilities when the accessed address is correct, yet the fetched instruction (opcode or operands) has been maliciously altered. In [30], the authors improved the HT detections by putting in parallel to the HSC a Hamming Single Error Correction (HSEC) circuitry. In [31], the authors presented a hardware-based methodology that monitors the microprocessors fetched instruction to detect Side-Channel Attacks (SCA). The approach presented in this manuscript integrates insights from these works by monitoring addresses and fetched instructions while adopting and evaluating HSEC codes to detect HT activations. This is achieved while maintaining minimal area and power overhead and without reducing the operating frequency.

This paper is organized as follows: Section II presents the HT models targeted by our proposal; Section III presents the proposed HT detection methodology and the details of the HSMs on which it relies; Section IV discusses the results of the proposed HSC, while Section V discusses the security analysis; Section VI concludes the paper.

## II. THE CONSIDERED THREAT MODELS

This manuscript focuses on attacks that compel the CPU to execute unauthorized code. An HT could be embedded in the microprocessor, forcing the fetch unit to access instruction memory locations where malicious code is stored. Similarly, HTs might compromise the instruction memory or system bus, altering the referenced memory locations and thus facilitating the launch of a malicious program.

According to the classical HTs classification [6], no assumptions have been made about the triggering mechanism of the referred HTs. The proposed approach considers both triggered and always-on HTs. The presented methodology assumes that these HTs are introduced by a malicious Intellectual Property (IP) provider while the design team and foundry involved in the microprocessor’s development are trusted. Therefore, the HSC is assumed to be reliable. Given that the attacker is the IP provider, we assume they have full knowledge of the hardware platform when inserting the HT.

The possible attack scenarios addressed by this proposal include:

- 1) An HT in the bus or main memory forcing the fetching from a particular memory address, which in turn would contain an instruction forced by the attacker;
- 2) An HT in the bus or main memory that forces the fetching of an instruction from the legitimate program but at a different point in time than intended at design time.

## III. THE PROPOSED SECURITY SOLUTION

The proposed approach involves introducing an HSC between the microprocessor and the instruction memory (as illustrated in Figure 1) to detect HTs attempting to execute malicious code. The HSC is configured during the installation phase of the program(s) the system will execute. At runtime, the HSC monitors the instruction fetch phase to detect and signal any HT activation<sup>1</sup>.




Figure 1: The proposed protection architecture

When the program is being installed into the main memory, the HSC operates in *configuration* mode. During this phase, it is programmed with the parity bits from the program instructions and the corresponding memory addresses. Once the program is installed, the HSC switches to *query* mode. During this runtime phase, in parallel to each instruction fetch, the HSC checks if the accessed memory address and the fetched instruction parity bits match the previously configured data. Specifically, the HSC verifies that the accessed address belongs to the memory space of the running program and that the fetched instruction parity bits correspond to what was originally stored at that address during the installation phase.

<sup>1</sup>The handling of the warning, such as via a non-maskable interrupt by the operating system, is outside the scope of this paper.




Figure 2: The methodology of the proposed Hardware Security Module

### A. The Architecture of the Hardware Security Modules

The architecture of the proposed methodology is illustrated in Figure 2: the HSMs implementing the HSC receive as input a memory address, an instruction, and the CONFIGURE/QUERY signal (which indicates whether the HSMs are operating in configure or query mode) and outputs a warning signal. In configure mode, the user provides both the addresses and instructions, meaning the program is being installed into the system’s memory (Figure 3a). In query mode, the addresses are the ones required by the microprocessor during the instruction fetch phase, while the instructions come from the main memory output (Figure 3b).

The input addresses and instructions are divided into chunks (in a manner detailed in Subsection III-B) and used to point to a set of memories within the HSMs during both the configure and query modes. During the configuration phase, the ECC sequences (i.e., the checking bits) are calculated for each instruction chunk. These checking bits are stored in dedicated memories, where each instruction chunk has a corresponding address. During the query phase, the checking bits of the fetched instruction are recalculated and compared with those stored previously during the configuration phase at the memory address from which the microprocessor is fetching the current instruction. By incorporating these checking bits, the system can detect errors (instruction bit flips) and HT activation using the same hardware blocks, ensuring efficient and robust security checking. In query mode, the HSMs raise a warning if the calculated checking bits do not match the stored ones.

Table I: ECC Types and Parameters

ECC Type	$n/k$	$k$	$p$
HSEC32	32	1	6
HSEC16	16	2	5
HSEC8	8	4	4

### B. Configuration and Usage of the Hardware Security Module

The HSMs receive an address and an instruction as inputs. These two inputs are concatenated and then divided into a series of vector chunks (INSTR\_CHUNK<sub>1</sub> to INSTR\_CHUNK<sub>k</sub> in Figure 2). Specifically, the number of memories in the HSMs is  $k$ , the *fragmentation factor*, while  $n$  represents the bit-size of addresses and instructions in the architecture under consideration. INSTR\_CHUNK<sub>1</sub> is defined as the chunk containing the first  $n/k$  bits of the address coupled with the first  $n/k$  bits of the instruction. Likewise, INSTR\_CHUNK<sub>2</sub> includes the second  $n/k$  bits of the address and the second  $n/k$  bits of the instruction, and this pattern continues. These chunks are then inputted to ECC architectures, providing parity bits that will be stored in ECC’s corresponding memories.

When the HSMs operate in configure mode, the user space provides both addresses and instructions while setting up the program. Figure 3 illustrates an example of the configuration process for the program’s two initial instructions within a system with 32-bit addresses and instructions and an HSM with  $n = 32$  and  $k = 4$ . It is important to note that different instructions may lead one or more vector chunks to reference the same parity bits. Consequently, different ECC memory locations can contain the same parity bits.

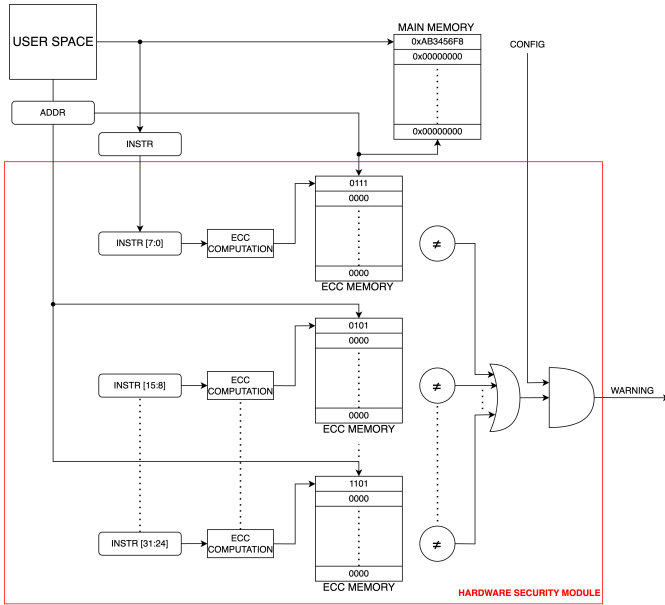
When the HSMs are in query mode, the address is provided by the microprocessor, and after retrieving it from the main memory, the corresponding instruction is also obtained. The  $k$  vector chunks are generated in the same manner as previously described; however, in this phase, the contents of the ECC Memory are read (as opposed to being written in configure mode). The HSMs compute the parity bits for the instruction retrieved from the main memory and compare them with the ones previously stored in the ECC memory at the address from the microprocessor. If this comparison indicates a discrepancy, the HSM raises an alarm. Figure 4 provides an example query process for the program’s first instruction.

## IV. EXPERIMENTAL RESULTS

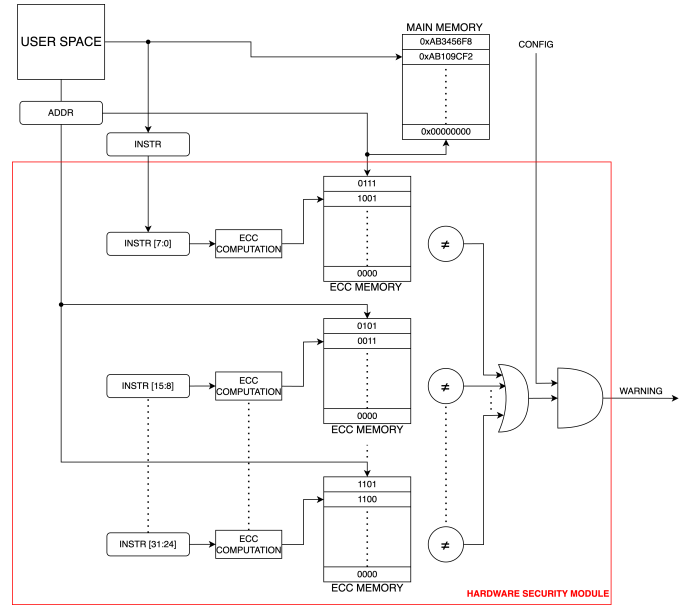
This section describes the hardware platform and benchmark programs used, presents the results obtained from the experiments, and provides a detailed security analysis. Three Hamming Single Error Correction (HSEC) functions were assessed regarding detection rates. Their architectural details are summarized in Table I, where  $n/k$  indicates the number of bits in each chunk,  $k$  denotes the fragmentation factor (representing the number of chunks managed by the HSMs), and  $p$  is the number of bits per ECC memory entry.

### A. Experimental setup

For benchmarking, the programs listed in Table II were selected, along with their respective counts of assembly instructions. The experimental analysis utilized the RI5CY [32] core from the PULPINO architecture, a 32-bit ultra-low-power processing platform designed for Internet of Things applications [33]. When implemented on a Xilinx Artix XC7A35T FPGA [34], the RI5CY core requires 15,314 LUTs and 9,881 FFs, achieving a clock frequency of approximately 50 MHz.



(a) Writing the first program instruction



(b) Writing the second program instruction

Figure 3: Configuration Methodology of the Hardware Security Module

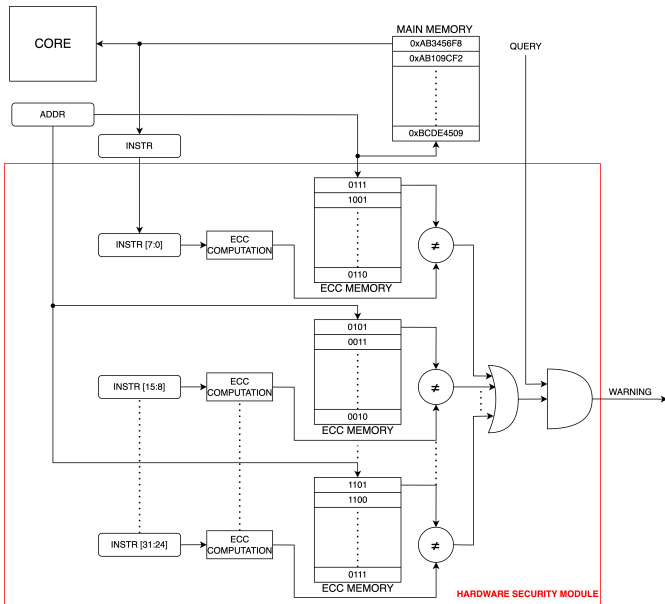


Figure 4: Querying Methodology of the Hardware Security Module

Using the same board and core, Table VII summarizes the resources utilized by the HSC and the percentage overheads associated with different ECC architectures. These results are compared with a previous HSC proposal [30], which employs ECCs to enhance the detection of HT activations. The HTs were emulated by randomly modifying the instruction memory address from which the microprocessor fetches instructions at runtime. Specifically, HTs capable of hijacking the execution flow toward a new program were considered. These HTs

Table II: The considered benchmark programs

Benchmarks	#32-bit Instruction (I)
Coremark (CM)	1288
Matrix Multiplication (MM)	216
Quick Sort (QS)	1023
RSort (RS)	4466
Sha (SHA)	516

force the selected memory address to fall outside the memory space of the program currently under execution. A total of 10,000 randomly generated HT activations were run, and the false negative (FN) rate was calculated as the ratio of runs in which the checker failed to raise an alarm to the total number of runs. Similarly, another 10,000 runs were performed without activating any HTs, and the false positive (FP) rate was measured as the ratio of runs in which the HSMs raised an alarm to the total number of runs. FPs, corresponding to false alarms, are theoretically and mathematically impossible due to the inherent properties of the error detection mechanism. Specifically, the HSMs approach relies on deterministic parity checks that uniquely map to the expected instruction execution flow. Since the ECC mechanism is designed to detect only deviations from this predefined structure, it cannot erroneously flag a correctly executed instruction sequence as an anomaly. In other words, given that the error-checking codes are computed based on a fixed and predictable transformation of the instruction data, any modification required to produce a false positive would necessitate an alteration that coincidentally results in a valid parity match—an event with negligible probability. Therefore, under ideal conditions without external sources of hardware faults or transient bit flips unrelated to the attack model, FPs cannot occur. The results of this experiment

Table III: FP and FN rates when the HT modifies the accessed instruction memory location (Threat model 1 reported in Section II)

Benchmarks	HSEC32		HSEC16		HSEC8	
	FP	FN	FP	FN	FP	FN
CM	0%	1.66%	0%	3.20%	0%	0%
MM	0%	1.52%	0%	3.28%	0%	0%
QS	0%	1.40%	0%	3.15%	0%	0%
RS	0%	1.55%	0%	3.00%	0%	0%
SHA	0%	1.61%	0%	2.97%	0%	0%
AVG	0%	1.548%	0%	3.12%	0%	0%

Table IV: FP and FN rates when the HT modifies the fetched instruction (Threat model 2 reported in Section II)

Benchmarks	HSEC32		HSEC16		HSEC8	
	FP	FN	FP	FN	FP	FN
CM	0%	1.64%	0%	3.55%	0%	0.12%
MM	0%	1.72%	0%	2.52%	0%	0.51%
QS	0%	1.61%	0%	2.93%	0%	0.07%
RS	0%	1.60%	0%	3.12%	0%	0.02%
SHA	0%	1.71%	0%	3.21%	0%	0.14%
AVG	0%	1.656%	0%	3.066%	0%	0.172%

are summarized in Tables III and IV. Such tables seem to reveal an anomalous trend in the False Negative (FN) rates for the HSEC16 implementation compared to HSEC32 and HSEC8. Theoretically, it is expected that: HSEC32, calculated on a single 32-bit block, would exhibit the highest FN rate because it evaluates a larger data chunk and provides a more robust parity check; HSEC16, calculated on two separate 16-bit blocks, would show an intermediate FP rate; HSEC8, calculated on four separate 8-bit blocks, would have the lowest FN rate due to finer granularity and an increased number of independent calculations. However, the results indicate that the FN rate for HSEC16 is higher than for HSEC32 and HSEC8. Such a trend has been confirmed by similar tests running this time CRC algorithms on the same chunk divisions (CRC32, CRC16, CRC8), as shown in Tables VI and V.

In these tests, the FN rates followed the same trend, suggesting that the discrepancies are not specific to the ECC type (Hamming or CRC) but are instead related to intrinsic factors such as block segmentation and the statistical distribution of instructions. The HSEC16 (and the CRC16) configuration relies on two 16-bit blocks for checking bit calculations. After these experiments, we can conclude that the data (i.e.,  $INSTR\_CHUNKS_1$  and  $INSTR\_CHUNKS_2$ ) in these blocks exhibit a high degree of correlation or overlap during the benchmark execution, reducing the effectiveness of checking bits, leading to a higher FN rate.

Given these findings, the proposed HSC combines HSEC32- and HSEC8-based HSMs to optimize detection accuracy while minimizing FN. The proposed HSC is illustrated in Figure 5.

Finally, targeting an FPGA implementation, the overhead introduced by the proposed HSC in terms of used resources and working frequency has been evaluated. Table VII reports the number of LUTs, FFs, and BRAM blocks required, as well as the maximum working frequencies, by the proposed HSC

Table V: FP and FN rates when the HT modifies the accessed instruction memory location (Threat model 1 reported in Section II)

Benchmarks	CRC32		CRC16		CRC8	
	FP	FN	FP	FN	FP	FN
CM	0%	0.09%	0%	0.83%	0%	0.09%
MM	0%	0.44%	0%	0.44%	0%	0.44%
QS	0%	0.01%	0%	0.01%	0%	0.01%
RS	0%	0.03%	0%	0.34%	0%	0.03%
SHA	0%	0.18%	0%	1.11%	0%	0.18%
AVG	0%	0.015%	0%	0.546%	0%	0.15%

Table VI: FP and FN rates when the HT modifies the fetched instruction (Threat model 2 reported in Section II)

Benchmarks	CRC32		CRC16		CRC8	
	FP	FN	FP	FN	FP	FN
CM	0%	1.66%	0%	3.20%	0%	0%
MM	0%	1.52%	0%	3.28%	0%	0.01%
QS	0%	1.40%	0%	3.15%	0%	0%
RS	0%	1.55%	0%	3.00%	0%	0%
SHA	0%	1.61%	0%	2.97%	0%	0%
AVG	0%	1.548%	0%	3.12%	0%	0.002%

and compares them with previous work [30]. The overhead in terms of additional FFs and LUTs is negligible between the two solutions. In addition, both of them do not slow down the microprocessor that works at 50MHz. On the other hand, our solution requires much less BRAMs.

We also compare undetected and false alarm rates. Table VIII reports the rate of the threat model 2 reported in Section II (both the solutions do not have any FP or FN on the threat model 1 reported in Section II). **One of the key contributions of this work is demonstrating that ECCs are capable of detecting and correcting errors and fully capable of detecting runtime security attacks, such as HT activations.** This is a significant advancement compared to previous work [30], which showed that combining ECCs with additional hash functions could enhance attack detection. Here, **we demonstrate that ECCs alone can effectively detect attacks**, eliminating the need for additional hardware complexity. This finding could have profound implications for microprocessor-based systems:

- **Reuse of Existing ECC Blocks:** Modern microprocessors often include ECC-based error detection and correction mechanisms to ensure data integrity. Our results show that these existing ECC blocks can also be repurposed or extended to detect security threats, providing a dual-use capability without significant additional resource overhead.
- **Streamlined Design:** By relying solely on ECCs, our approach simplifies the hardware implementation, making it more cost-effective and easier to integrate into existing microprocessor architectures. This is especially valuable for low-power or resource-constrained environments like IoT devices or embedded systems.
- **Comprehensive Security Coverage:** Unlike the previous solution, which required hash functions to complement

Table VII: Resource occupation and working frequency

Benchmarks	Proposed solution				Solution in [30]			
	#LUTs	#FFs	#BRAMs	F. (MHz)	#LUTs	#FFs	#BRAMs	F. (MHz)
CM	72 (0.47%)	24 (0.24%)	0.5	560 MHz	82 (0.53%)	31 (0.31%)	8.5	275 MHz
MM	72 (0.47%)	24 (0.24%)	0.5	560 MHz	82 (0.53%)	31 (0.31%)	8.5	275 MHz
QS	72 (0.47%)	24 (0.24%)	0.5	560 MHz	82 (0.53%)	31 (0.31%)	8.5	275 MHz
RS	72 (0.47%)	24 (0.24%)	0.5	560 MHz	82 (0.53%)	31 (0.31%)	9.5	275 MHz
SHA	72 (0.47%)	24 (0.24%)	0.5	560 MHz	-	-	-	-

Table VIII: FP and FN rates (Threat model 2 reported in Section II)

Benchmarks	Proposed solution		Solution in [30]	
	FP	FN	FP	FN
CM	0%	0.12%	0%	0.11%
MM	0%	0.44%	0%	0.34%
QS	0%	0.01%	0%	0.08%
RS	0%	0.12%	0%	0.05%
SHA	0%	0.14%	-	-
AVG	0%	0.166%	0%	0.145%

ECCs for effective attack detection, we demonstrate that ECCs are sufficient to detect attacks with high accuracy. This fully exploits the potential of ECCs, turning them into a robust tool for runtime security monitoring.

- **Scalability and Adaptability:** The ability to detect errors and attacks using the same ECC-based hardware ensures that the solution can be easily adapted to various microprocessor architectures without redesigning core security components.




Figure 5: The structure of the proposed HSC

## V. SECURITY ANALYSIS

The presented experimental results demonstrate that the proposed HSC allows the detection of 100% of the runtime activations of HTs that try to force the CPU to execute malicious programs installed in instruction memory locations outside the memory space of the running program as well as 100% of HTs that try to force the CPU to execute a legit instruction but in a different moment respect to the designed one. Furthermore, the proposed HSC never incurs false alarms. It is worth mentioning that, as it has already been discussed, the effectiveness of the proposed solution is independent of the triggering mechanism of the HT, i.e., combinational/sequentially triggered, externally activated, time-bombs and always-on, and of the design stage during which the HT has been inserted.

The proposed solution could be defeated by denial-of-service HTs that modify the execution flow of the legal program. We identified two possible scenarios: i) HTs that halt the system by maliciously making the CPU fetch always the same legal instruction (or sequence of legal instructions) from legal memory locations, and ii) HTs that halt the system by making it crash by fetching a legal instruction from a memory location belonging to the authorized program but at the wrong time or in the wrong order, e.g., fetching a jump instruction too early during the execution flow.

## VI. CONCLUSION

We presented a security architecture to protect microprocessor-based systems against HT attacks. We integrated our proposal within a system featuring a RISC-V processor implemented on an FPGA device and running a set of software benchmarks. Our proposal can detect 100% of possible Hardware Trojan (HT) activations without false or undetected alarms. We measured a LUT and FF overhead of less than 1%, with 0.5 #BRAMs required and no working frequency reduction.

## REFERENCES

- [1] DIGITIMES, "Trends in the global IC design service market." <http://www.digitimes.com/news/a20120313RS400.html?chid=2>.
- [2] M. Rostami, F. Koushanfar, J. Rajendran, and R. Karri, "Hardware security: Threat models and metrics," in *Proc. Int. Conf. Computer-Aided Design*, pp. 819–823, 2013.
- [3] Mohammad Tehranipoor and Cliff Wang, *Introduction to Hardware Security and Trust*. Springer-Verlag New York, 2012.
- [4] L. Cassano, S. D. Mascio, A. Palumbo, A. Menicucci, G. Furano, G. Bianchi, and M. Ottavi, "Is risc-v ready for space? a security perspective," in *2022 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, pp. 1–6, 2022.
- [5] P. R. Nikiema, A. Palumbo, A. Aasma, L. Cassano, A. Kritikakou, A. Kulmala, J. Lukkarila, M. Ottavi, R. Psiakis, and M. Traiola, "Towards dependable risc-v cores for edge computing devices," in *2023 IEEE 29th International Symposium on On-Line Testing and Robust System Design (IOLTS)*, pp. 1–7, 2023.
- [6] M. Tehranipoor and F. Koushanfar, "A survey of hardware trojan taxonomy and detection," *IEEE Design & Test of Computers*, vol. 27, no. 1, 2010.
- [7] M. Xue, C. Gu, W. Liu, S. Yu, and M. O'Neill, "Ten years of hardware trojans: a survey from the attacker's perspective," *IET Computers & Digital Techniques*, vol. 14, no. 6, pp. 231–246, 2020.
- [8] Y. Jin, M. Maniatakos, and Y. Makris, "Exposing vulnerabilities of untrusted computing platforms," in *Proc. Int. Conf. Computer Design*, pp. 131–134, 2012.
- [9] N. G. Tsoutsos and M. Maniatakos, "Fabrication attacks: Zero-overhead malicious modifications enabling modern microprocessor privilege escalation," *IEEE Trans. Emerging Topics in Computing*, vol. 2, no. 1, pp. 81–93, 2014.
- [10] X. Wang, T. Mal-Sarkar, A. Krishna, S. Narasimhan, and S. Bhunia, "Software exploitable hardware trojans in embedded processor," in *2012 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, pp. 55–58, IEEE, 2012.

- [11] C. Domas, "Hardware backdoors in x86 cpus." <https://i.blackhat.com/us-18/Thu-August-9/us-18-Domas-God-Mode-Unlocked-Hardware-Backdoors-In-x86-CPUs-wp.pdf>, 2018.
- [12] X. Chuan, Y. Yan, and Y. Zhang, "An efficient triggering method of hardware Trojan in AES cryptographic circuit," in *Proc. Int. Conf. Integrated Circuits and Microsystems*, pp. 91–95, 2017.
- [13] L. Cassano, M. Iamundo, T. A. Lopez, A. Nazzari, and G. Di Natale, "Deton: Defeating hardware trojan horses in microprocessors through software obfuscation," *Journal of Systems Architecture*, vol. 129, p. 102592, 2022.
- [14] A. Palumbo, M. Ottavi, and L. Cassano, "Built-in software obfuscation for protecting microprocessors against hardware trojan horses," in *2023 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, pp. 1–6, IEEE, 2023.
- [15] J. Zhang, F. Yuan, L. Wei, Y. Liu, and Q. Xu, "Veritrust: Verification for hardware trust," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, vol. 34, no. 7, pp. 1148–1161, 2015.
- [16] Y. Liu, Y. Zhao, J. He, A. Liu, and R. Xin, "Scca: Side-channel correlation analysis for detecting hardware trojan," in *Proc. Int. Conf. Anti-counterfeiting, Security, and Identification*, pp. 196–200, 2017.
- [17] H. Salmani and M. Tehranipoor, "Analyzing circuit vulnerability to hardware trojan insertion at the behavioral level," in *Proc. Int. Symp. Defect and Fault Tolerance in VLSI and Nanotechnology Systems*, pp. 190–195, 2013.
- [18] H. Salmani and M. Tehranipoor, "Layout-aware switching activity localization to enhance hardware trojan detection," *IEEE Trans. Information Forensics and Security*, vol. 7, no. 1, pp. 76–87, 2012.
- [19] A. Palumbo, L. Cassano, B. Luzzi, J. A. Hernández, P. Reviriego, G. Bianchi, and M. Ottavi, "Is your fpga bitstream hardware trojan-free? machine learning can provide an answer," *Journal of Systems Architecture*, vol. 128, p. 102543, 2022.
- [20] S. Ribes, F. Malatesta, G. Garzo, and A. Palumbo, "Machine learning-based classification of hardware trojans in fpgas implementing risc-v cores," in *ICISSP*, pp. 717–724, 2024.
- [21] D. Šišejković, F. Merchant, R. Leupers, G. Ascheid, and S. Kegreiss, "Control-lock: Securing processor cores against software-controlled hardware trojans," in *Proceedings of the 2019 on Great Lakes Symposium on VLSI, GLSVLSI '19*, pp. 27–32, 2019.
- [22] A. Basak, S. Bhunia, T. Tkacik, and S. Ray, "Security assurance for system-on-chip designs with untrusted ips," *IEEE Transactions on Information Forensics and Security*, vol. 12, no. 7, pp. 1515–1528, 2017.
- [23] J. Dubeuf, D. Hély, and R. Karri, "Run-time detection of hardware trojans: The processor protection unit," in *2013 18th IEEE European Test Symposium (ETS)*, pp. 1–6, 2013.
- [24] G. Bloom, B. Narahari, and R. Simha, "Os support for detecting trojan circuit attacks," in *2009 IEEE International Workshop on Hardware-Oriented Security and Trust*, pp. 100–103, 2009.
- [25] K. Cheang, C. Rasmussen, D. Lee, D. W. Kohlbrenner, K. Asanović, and S. A. Seshia, "Verifying risc-v physical memory protection," *arXiv preprint arXiv:2211.02179*, 2022.
- [26] J. H. Ng, C. H. Ang, and H. C. Law, "A realization of io physical memory protection for risc-v systems," in *2022 IEEE 15th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MC-SoC)*, pp. 375–380, IEEE, 2022.
- [27] C. Shepherd, K. Markantonakis, and G.-A. Jaloyan, "Lira-v: Lightweight remote attestation for constrained risc-v devices," in *2021 IEEE Security and Privacy Workshops (SPW)*, pp. 221–227, IEEE, 2021.
- [28] A. Bolat, L. Cassano, P. Reviriego, O. Ergin, and M. Ottavi, "A micro-processor protection architecture against hardware trojans in memories," in *2020 15th Design Technology of Integrated Systems in Nanoscale Era (DTIS)*, pp. 1–6, 2020.
- [29] A. Palumbo, L. Cassano, P. Reviriego, G. Bianchi, and M. Ottavi, "A lightweight security checking module to protect microprocessors against hardware trojan horses," in *2021 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, pp. 1–6, 2021.
- [30] A. Palumbo, L. Cassano, P. Reviriego, and M. Ottavi, "Improving the detection of hardware trojan horses in microprocessors via hamming codes," in *2023 IEEE International Symposium on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFT)*, pp. 1–6, IEEE, 2023.
- [31] K. Arkan, A. Palumbo, L. Cassano, P. Reviriego, S. Pontarelli, G. Bianchi, O. Ergin, and M. Ottavi, "Processor security: Detecting microarchitectural attacks via count-min sketches," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 30, no. 7, pp. 938–951, 2022.
- [32] M. Gautschi, P. D. Schiavone, A. Traber, I. Loi, A. Pullini, D. Rossi, E. Flamand, F. K. Gurkaynak, and L. Benini, "Near-threshold risc-v core with dsp extensions for scalable iot endpoint devices," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 25, pp. 2700–2713, Oct 2017.
- [33] A. Traber, F. Zaruba, S. Stucki, A. Pullini, G. Haugou, E. Flamand, F. K. Gurkaynak, and L. Benini, "Pulpino: A small single-core risc-v soc," in *3rd RISCv Workshop*, 2016.
- [34] R. Höller, D. Haselberger, D. Ballek, P. Rössler, M. Krapfenbauer, and M. Linauer, "Open-source risc-v processor ip cores for fpgas — overview and evaluation," in *2019 8th Mediterranean Conference on Embedded Computing (MECO)*, pp. 1–6, June 2019.