Exploring the Security of an Accelerator integrated with Core-V eXtention InterFace (CV-X-IF)

Alessandra Dolmeta[1]*, Behnam Farnaghinejad[2], Davide Bellizia[3], Guido Masera[1], Maurizio Martina[1], Stefano Di Carlo[2], Alessandro Savino[2], and Ernesto Sanchez[2]

[1]DET, Politecnico di Torino, Torino, Italy, [2]DAUIN, Politecnico di Torino, Torino, Italy, [3] Telsy S.p.A, Rome, Italy

Abstract

The adoption of hardware accelerators in embedded systems enhances performance but raises security concerns, particularly regarding side-channel leakage. The Core-V eXtension Interface (CV-X-IF) simplifies accelerator integration in RISC-V, yet its security remains unexplored. This study presents the first side-channel analysis of CV-X-IF, comparing a native XOR instruction to a custom accelerator using Test Vector Leakage Assessment (TVLA). Power trace analysis reveals measurable differences, highlighting potential vulnerabilities. These findings underscore the need for secure accelerator design in RISC-V microcontrollers.

Introduction

As embedded systems become global across IoT devices, automotive systems, and industrial control networks, the demand for robust microcontroller security has reached critical importance. The increase in connected devices handling sensitive data has intensified the exploration of side-channel vulnerabilities, where attackers exploit physical leakage (e.g., power consumption, electromagnetic emissions) to extract secrets. Currently, the push for higher performance and energy efficiency has driven the adoption of specialized hardware accelerators, which offload computationally intensive tasks from the main processor core. While these accelerators enhance functionality, their integration into microcontroller architectures introduces new attack surfaces, requiring careful evaluation of both performance gains and security trade-offs.

The modularity of the open-source RISC-V architecture has driven innovation in secure processor design, enabling customizable extensions without proprietary constraints. However, traditional methods for integrating accelerators—such as modifying the core's instruction decoder—often introduce complexity, increase vulnerability to design flaws, and hinder scalability. The Core-V eXtension Interface (CV-X-IF)[1] addresses these challenges by providing a standardized framework for attaching tightly coupled co-processors directly to the pipeline, bypassing invasive changes to the core's decode logic. By streamlining accelerator integration, CV-X-IF promises to enhance performance and design flexibility while maintaining compliance with RISC-V's security-focused ecosystem. Yet, the security implications of this interface remain unexplored, particularly its resilience against side-channel attacks—a critical oversight given its growing adoption in security-critical applications.

This study presents the first security analysis of the CV-X-IF interface, focusing on its potential to introduce or amplify side-channel leakage. We compare two implementations of a fundamental operation (XOR): one using the native RISC-V `xor` instruction and another executed via a custom accelerator integrated through CV-X-IF. Employing Test Vector Leakage Assessment (TVLA), a standardized methodology for detecting information leakage, we evaluate whether the accelerator interface introduces measurable side-channel vulnerabilities with respect to the conventional instruction-based approach. Our findings aim to establish foundational insights into the risks and mitigations associated with CV-X-IF.

Methodologies

To evaluate the side-channel resilience of the CV-X-IF-integrated accelerator, we employed a Test Vector Leakage Assessment (TVLA) framework. Instead of the standard fixed-vs-random approach, we analyzed power traces while executing the same XOR operation on identical input data across two implementations. This method isolates differences in leakage purely due to the execution path rather than input variations.

The experiment compared two XOR implementations:

- **Native RISC-V Instruction**: The baseline xor instruction executed directly on the CVA6 core (Ariane RISC-V CPU[2]).
- **Custom Accelerator**: A hardware-optimized XOR module implemented as a tightly coupled co-processor via the CV-X-IF, attached to the CVA6 pipeline.

*Corresponding author: `alessandra.dolmeta@polito.it`

[1] `https://github.com/openhwgroup/core-v-xif/tree/main`

[2] `https://github.com/openhwgroup/cva6`

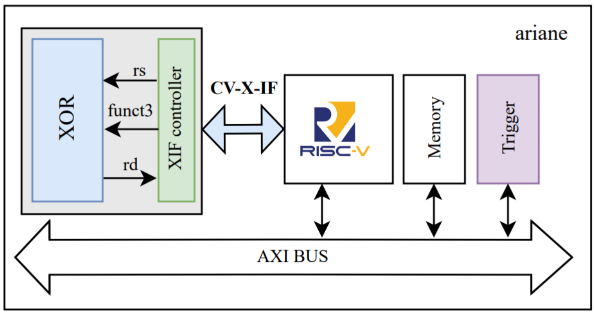

Figure 1: *Architecture of the tested RISC-V SoC.*

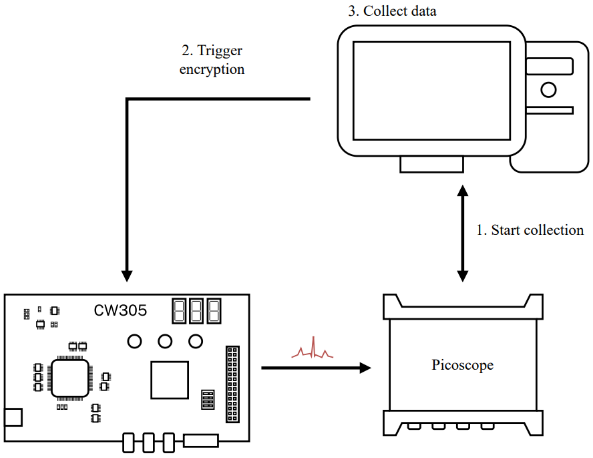

Figure 2: *Setup environment.*

Both implementations were tested under identical environmental conditions to eliminate external variables. The test platform, shown in Fig.2, comprised:

- **Target Hardware**: the CVA6 core with the CV-X-IF integrated XOR accelerator implemented on an AMD Xilinx Artix-7 FPGA and hosted on a CW305 board [3]. The core is running at 5 MHz and powered at 1 V.
- **Measurement Tools:** a Picoscope 6404D oscilloscope for capturing power traces at 1.25 GS/s, synchronized via a trigger signal.
- **Control Interface**: a UART-based communication protocol between the FPGA and a host PC, enabling:
  - PRNG Initialization: A cryptographically secure pseudorandom number generator (AES-CBC with a 128-bit seed) to randomize the selection between the two XOR implementations, preventing timing or execution-pattern biases.
  - Trace Configuration: Dynamic adjustment of the number of traces (10,000–50,000 per test) to ensure statistical significance.

A GPIO pin is asserted at the start of each XOR operation provided a precise trigger signal, isolating the relevant execution window (10 ns) in power traces. For each trial, the PRNG selected either the native xor instruction or the CV-X-IF accelerator, ensuring unbiased alternation. Identical input data and cryptographic keys were used for both implementations to isolate leakage differences attributable to the CV-X-IF interface. We have leveraged on the Test Vector Leakage Assessment methodology (TVLA) [1] to evaluate the presence of a meaningful statistical difference between the two aforementioned execution. Results are shown in Fig.3. The value of the t-score widely exceeds the +/-4.5 threshold normally used in the context of side-channel analysis to remark a meaningful statistical difference in the execution of the XOR instruction on the native execution stage of the core and the same instruction running in the CV-X-IF-based accelerator.

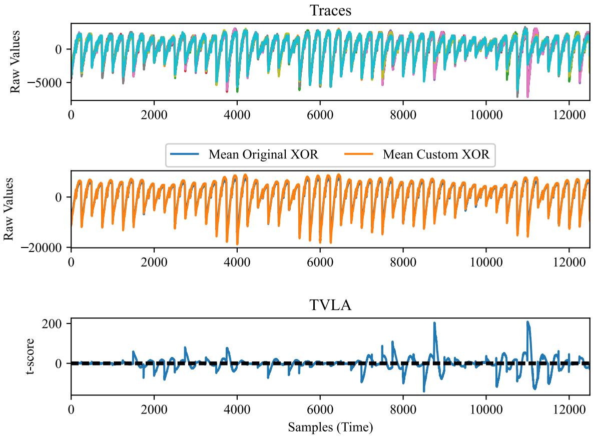

Figure 3: *Test results.*

Conclusions

The results of the TVLA tests demonstrate a significant difference between the two XOR implementations. This discrepancy is expected due to the distinct signal paths and interactions introduced by the CV-X-IF interface. However, the findings underscore an essential starting point for a more comprehensive security analysis of this new extension mechanism. Understanding and mitigating the side effects of the CV-X-IF interface will be crucial in designing secure and efficient cryptographic accelerators within RISC-V-based systems. Future work will extend this analysis to enhance the security of a cryptographic coprocessor integrated with the CV-X-IF interface.

References

[1] George Becker et al. “Test vector leakage assessment (TVLA) methodology in practice”. In: *International Cryptographic Module Conference*. Vol. 1001. sn. 2013, p. 13.

[3] https://www.newae.com/products/nae-cw305