

COVERT: Trojan Detection in COTS Hardware via Statistical Activation of Microarchitectural Events

Mahmudul Hasan
Dept. of EECS
University of Kansas
m.hasan@ku.edu

Sudipta Paria
Dept. of ECE
University of Florida
sudiptaparia@ufl.edu

Swarup Bhunia
Dept. of ECE
University of Florida
swarup@ece.ufl.edu

Tamzidul Hoque
Dept. of EECS
University of Kansas
hoque@ku.edu

Abstract—Commercial Off-The-Shelf (COTS) hardware, such as microprocessors, are widely adopted in system design due to their ability to reduce development time and cost compared to custom solutions. However, supply chain entities involved in the design and fabrication of COTS components are considered untrusted from the consumer’s standpoint due to the potential insertion of hidden malicious logic or hardware Trojans (HTs). Existing solutions to detect Trojans are largely inapplicable for COTS components due to their black-box nature and lack of access to a golden model. A few existing studies rely on expensive equipment, lack scalability, and are applicable only to a limited class of Trojans. In this work, we present a novel golden-free trust verification framework, **COVERT**, for COTS microprocessors, which can efficiently test the presence of hardware Trojan implants by identifying microarchitectural rare events and transferring activation knowledge from existing processor designs to trigger highly susceptible internal nodes. **COVERT** leverages Large Language Models to automatically generate test programs that trigger rare microarchitectural events, which may be exploited to develop Trojan trigger conditions. By deriving these events from publicly available Register Transfer Level implementations, **COVERT** can verify a wide variety of COTS microprocessors that inherit the same Instruction Set Architecture. We have evaluated **COVERT** on open-source RISC-V COTS microprocessors and demonstrated its effectiveness in activating combinational and sequential Trojan triggers with high coverage, highlighting the efficiency of the trust verification. By pruning rare microarchitectural events from mor1kx Cappuccino OpenRISC processor design, **COVERT** has been able to achieve more than 80% trigger coverage for the rarest 5% of events in or1k Marocchino and PicoRV32 as COTS processors.

Index Terms—Trojan Detection, Test Generation, Microarchitectural Events, COTS, Combinational and Sequential Trojans.

I. INTRODUCTION

Commercial off-the-shelf (COTS) components offer a compelling system design paradigm due to reduced development time, lower hardware costs, and market availability compared to custom design approach. As a result, COTS microprocessors have been widely adopted in military, avionics, finance, and commercial sectors. According to a 2022 report, around 98% of microelectronic components used in defense applications are COTS [1]. The increasing reliance on COTS components has simultaneously introduced significant security concerns. Suppliers of COTS components distribute design, manufacturing, and testing across various domestic and foreign untrusted entities. Any untrusted entities with access to the design could introduce hidden malicious logic or hardware Trojans capable of causing functional failures or leakage of sensitive information such as encryption keys [2]. Many real-world cyber attacks indicate the rising threat of hardware Trojans in untrusted components [3]. Most existing hardware Trojan research primarily focuses on

insertion threats either from untrusted foundry or untrusted IP vendor [4], [5]. Most of these countermeasures rely on one or more of the following: (i) access to golden (Trojan-free) chips, (ii) white-box access to the design, or (iii) design-time modifications to integrate countermeasures [2]. None of these assumptions hold for COTS components, rendering existing methods ineffective. For example, logic testing approaches depend on design access to generate effective test vectors [6], [7], while most side-channel analysis (SCA) methods require golden chips to establish reference signatures and design access for test generation to ensure switching activity in Trojans [8]. A few studies have developed runtime approaches applicable to COTS hardware to detect, tolerate, or prevent Trojan activation in the field [9]–[11]. While runtime solutions introduce an additional layer of security, trust verification approach is generally more desirable as the primary defense mechanism. Side-channel assisted golden-chip free trust-verification methods have been developed that can be applied to COTS hardware, but these techniques also assume the presence of design-aware test-generation methods that can ensure activation of Trojan nets [12], [13]. Therefore, there is a critical unmet need for high-confidence black-box trust verification of COTS hardware.

In this paper, we propose **COVERT** (COTS VERification Framework for Trojan Detection), a novel test generation approach for hardware Trojan detection in COTS microprocessors through statistical activation of microarchitectural events. Unlike arbitrary rare nets, rare microarchitectural events represent controllable and architecturally meaningful behaviors that can be deliberately exercised through program execution, enabling attackers to achieve stealthy yet predictable Trojan activation. **COVERT** leverages the software code generation capabilities of Large Language Models (LLMs) to generate targeted test programs for activating the rare internal events that can form the triggers of potential Trojans. While the golden design of a COTS processor may not be available, a procurer of the COTS component typically receives comprehensive documentation of the instruction set architecture (ISA) and architectural features, including the list of interrupts, events, and limited hardware components details. **COVERT** utilizes this publicly available ISA together with an open-source RTL implementation of that ISA to explore rare microarchitectural events. The majority of these events are implementation agnostic, meaning they are applicable to COTS processors of the same ISA. Next, **COVERT** construct targeted test programs that exercise those rare microarchitectural events, which are common targets for crafting hard-to-activate Trojan trigger conditions.

COVERT leverages the relationship across design abstractions to connect low-level Trojan nets with high-level events.

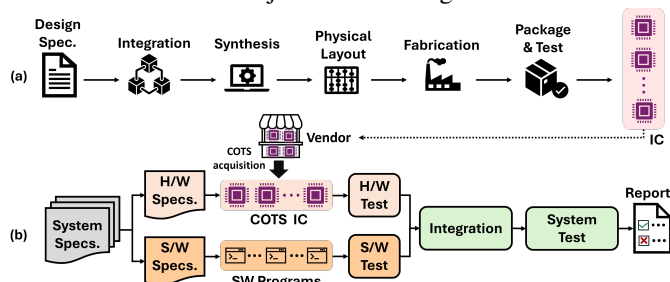


Fig. 1. (a) Modern IC design flow; (b) COTS IC integration flow [14], [15].

Microarchitectural events, such as pipeline hazards, cache behaviors, or exception conditions are directly derived from the RTL design and represent exact functional scenarios within the processor. If an event is rarely exercised at the microarchitectural level, the individual registers and logic cones that implement it at the gate level will likewise experience rare switching activity. COVERT prunes the rare microarchitectural event space for a given ISA based on publicly available information and open-source implementations and develops tests to thoroughly explore the potential Trojan design space based on the events. The key contributions of this paper include:

- We propose a novel framework, COVERT, that enables Trojan detection in COTS microprocessors without access to their implementation by transferring activation knowledge from ISA and existing processor designs.
- We present a systematic process for extracting rare microarchitectural events from known RTL implementations of microprocessors. To activate these events, we have developed an LLM-assisted program generation framework that consists of an agent-based workflow in LangGraph (built on LangChain) with iterative mitigation of program errors through automated feedback.
- We have extensively evaluated COVERT by applying test programs generated from the OpenRISC mor1kx Cappuccino reference design to two black-box COTS processors: OpenRISC Marocchino, which shares the same ISA, and PicoRV32, which follows a similar RISC design philosophy. Although PicoRV32 differs in ISA and instruction encoding from OpenRISC, both PicoRV32 and Marocchino follow core RISC principles such as a load-store model and simple in-order execution. We observed over 80% trigger coverage for the rarest event groups in both. We also analyzed coverage for complex Trojans with multiple excitation of these events. We have open-sourced all the related artifacts from our experiments in this link [16].

II. BACKGROUND

A. Trust Issues in COTS Microprocessors

The design process of a custom IC goes through several steps starting from the specification to all the way to the layout, as shown in Fig. 1(a). Alternatively, if a COTS IC is available in the market, a system developer can avoid these steps and directly obtain a COTS component from the market. For example, a COTS processor based on desired specifications, performance,

TABLE I
COMPARATIVE ANALYSIS OF EXISTING VERIFICATION TECHNIQUES VS. THE PROPOSED COVERT FRAMEWORK.

Proposed Solution	Goal	Task	Trojan Det.?	#events / #prog.
LLM4DV [18]	Functional Verification	Stimulus generation	No	N/A / N/A
Xiao et al. [19]	Processor Verification	Test program generation	No	N/A / 164
LLM-TG [20]	Processor Verification	Test generation	No	N/A / 132
Paiya et al. [15]	COTS Verification	Test program generation	No	9-11 / ~10
COVERT (this work)	COTS Verification + Trojan Detection	Test program generation	Yes	368 / ~500

and ISA can be procured, as shown in Fig. 1(b). However, the use of COTS hardware introduces trust issues for the procurer due to limited visibility into the design and manufacturing processes, where any of these entities can introduce malicious modification or hardware Trojans.

B. Existing Verification Techniques

Existing hardware Trojan detection methods are applicable under one or more of the assumptions: i) access to a golden (Trojan-free) design as reference, ii) white-box accessibility to the chip design, iii) design modification to facilitate verification. None of these requirements can, however, be met in the context of COTS components [2], [4], [9], [17]. For instance, logic testing methods such as MERO [6] aim to generate test vectors to ensure activation of hard-to-activate rare nets within the design. However, existing test generation methods focusing on rare net activation rely on access to the netlist [6]–[8]. A few existing verification techniques leverage side channels (e.g., power or delay). Authors in [12] employ a laser-assisted side channel to observe suspicious hard-to-activate flip-flops that are not part of the scan chain. In [13], the authors propose unsupervised clustering of power side-channels to identify suspicious chips without a golden signature. Both studies, however, require awareness of the design information to create test patterns that activate the rare nets to their rare values.

C. LLM-assisted Processor Verification

LLM-based test program generation methods have been developed for functional verification and bug detection [18]–[20]. These techniques have shown potential in achieving high code/functional coverage without access to the processor design. However, they target functional verification for uncovering unintentional design bugs rather than trust verification, and they do not generate programs that exercise complex microarchitectural events or activate the rare nets that often serve as Trojan triggers. They also cannot be readily extended to Trojan detection because they lack a security-oriented coverage plan, they assume a golden model for correctness, they do not redefine coverage bins around security targets, and they require access to internal RTL coverage in some cases, which is infeasible for COTS. Table I provides a comparative analysis of the proposed COVERT framework against existing verification techniques applicable to processors, highlighting their limitations in detecting HTs in COTS hardware.

III. METHODOLOGY

Fig. 2 illustrates the proposed COVERT framework. The process begins by identifying the ISA of the target COTS microprocessor and obtaining an open-source RTL design that implements the same ISA or follows the same architecture philosophy. From this RTL, we extract detailed microarchitectural information to define high-level events. Signal-probability analysis is then used to rank these events by rarity. Finally, a




Fig. 2. COVERT framework overview for Trojan detection in COTS hardware. state-driven agent generates test programs to trigger the rare events and reports their trigger coverage.

A. Microarchitectural Event Exploration

A *microarchitectural event* refers to an observable internal condition or transition in the processor’s microarchitecture that encapsulates a meaningful high-level behavior resulting from the activation of specific low-level nets within the microprocessor. For example, if a rare net is influenced by the instruction page-fault flag, privilege mode bits, and TLB miss status, it can be abstracted as the microarchitectural event: “instruction page fault in user mode with pending TLB refill.” Such event abstractions effectively map low-level signals to architecturally invariant state, enabling precise targeting through test programs. Here **architecturally invariant means the event remains valid and maintains identical semantics across processors that implement the same ISA, regardless of internal design.** The systematic flow of event exploration is illustrated in Fig. 3.

The process begins with ① **identifying the ISA of the target COTS microprocessor** and selecting representative RTL implementations that conform to the same ISA and specification. One ② **Verilog RTL design** is chosen as the **reference** implementation to facilitate rare event identification and subsequent targeted test generation. We chose the RTL design for event generation because it exposes the internal control logic and dependencies needed to locate hard-to-reach conditions that public ISA manuals and datasheets do not reveal. The selected ③ **RTL is then parsed and transformed into an Abstract Syntax Tree (AST)** enabling both structural and semantic analysis of the design. A custom script ④ **recursively traces each identified rare signal backward** through its fan-in logic in the AST, following both intra-module and cross-module signal dependencies up to a predefined depth. This backward tracing is typically bounded at the decoder stage, which is ideal for revealing the specific instructions responsible for activating the rare signal. The resulting hierarchical trace provides a crucial link between low-level net activity and corresponding high-level architectural behavior, stored in a JSON database. The trace is further abstracted into simplified textual representations of procedural and control logic blocks, providing the necessary ⑤ **context for the LLM to interpret structural dependencies** in natural language and **infer corresponding microarchitectural events**. To maintain semantic accuracy and prevent hallucinations, the LLM is provided with carefully crafted system prompt specifying its role, contextual information, analysis rules, and the expected output format. LLM generates event names with high-level descriptions linked to the traced signals, architectural summaries, test generation guidance, and relevant instruction categories, serving as the foundation for producing targeted test programs.

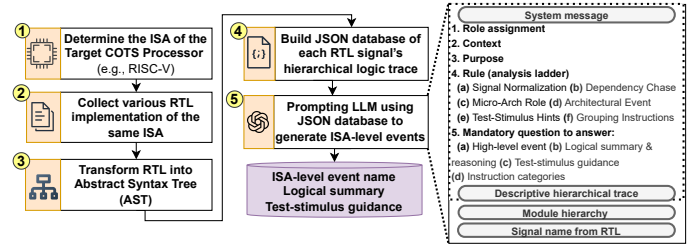


Fig. 3. Overview of the event exploration process, where it maps RTL signals into microarchitectural events by hierarchical logic unrolling through AST.

B. Rare Event Identification

We perform rarity analysis to identify microarchitectural events that occur infrequently and serve as Trojan triggers. The reference RTL is stimulated with diverse benchmark programs that generate realistic switching activity. These benchmarks are chosen to cover a wide range of instruction types, data patterns, and control-flow behaviors, providing a realistic view of real-world applications and enabling extraction of rare events. During simulation, Value Change Dump (VCD) traces are generated, capturing signal transitions over time. From the VCD text files, Signals are classified into *single-bit nets* and *multi-bit buses*, each requiring a distinct activity metric for rarity analysis. For **single-bit nets**, the probabilities of the signal being at logic high (\hat{p}_1) and logic low (\hat{p}_0) are estimated over the simulation cycles. The rarity metric is defined as $\theta = \min(\hat{p}_0, \hat{p}_1)$. For **multi-bit buses**, activity is measured using the normalized toggle rate defined as:

$$\theta = \frac{\text{Number of value changes}}{\text{Total number of consecutive cycle pairs}}$$

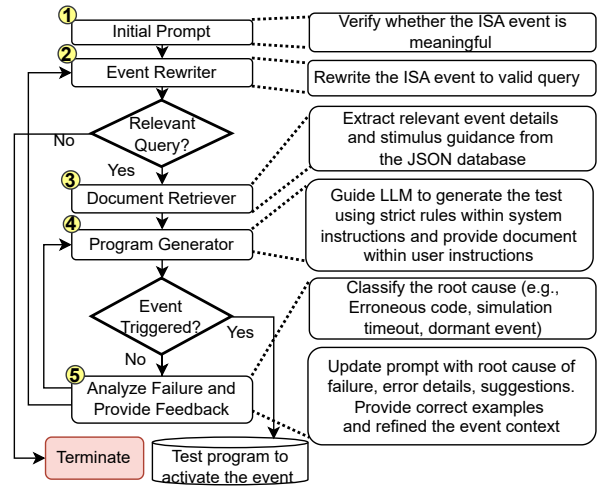


Fig. 4. Proposed agent-driven workflow in COVERT for test program generation leveraging LLM to activate microarchitectural rare events.

C. Test Program Generation by leveraging LLM

COVERT integrates the automated software code generation capabilities of LLMs to generate targeted test programs in C/C++ or inline assembly. An LLM-based state-driven agentic workflow, where each state performs a specific, sequenced task, for program generation is illustrated in Fig. 4. The following subsections describe each state briefly.

1) *Event Rewriter and Topic Classifier*: COVERT starts with processing the user-provided event query to ① validate its correctness, ensuring the event name is valid, consistent, and

free from linguistic errors. If errors are detected, the agent ② rewrites the query while preserving its intended semantics. The revised query is then passed to a topic classifier, which verifies whether it maps to a valid microarchitectural event under the constraints of the target ISA.

2) *Document Retriever*: After validating the event query, the agent proceeds to ③ retrieve microarchitectural knowledge associated with the event. This step is critical for guiding the LLM to generate meaningful programs beyond surface-level event names by providing the underlying logical reasoning. The knowledge base for this retrieval is constructed during the event generation phase (Section III-A) and stored in a structured JSON file. For each event, the database contains three key metadata fields in addition to the event name: (a) a logical summary describing the reasoning chain that leads to the event, (b) stimulus-generation guidance, and (c) relevant instruction categories indicating which ISA operations are most applicable.

3) *Program Generator*: COVERT leverages LLM to ④ generate test programs using a systematic prompt template. The process begins with an iteration check, with no feedback applied in the first iteration due to the absence of prior output. The state then iterates until the generated program meets the trigger condition. If the generated program fails, the agent refines the prompt by incorporating failure feedback and tightening constraints, progressively increasing the likelihood of success. The prompt is organized using LLM APIs into two layers: system and user instructions, allowing dynamic updates to the user layer while keeping the system layer tasks stable. The system prompt defines the model’s role, enforces global rules (e.g., single code block, no fabricated opcodes), and outlines a step-by-step analysis procedure. It also includes k-shot learning through working examples: one in each standard C and inline assembly, to demonstrate the expected output format. The user prompt includes four key components: (a) event-specific metadata from previous state (see Section III-C2), (b) user-defined constraints and objectives, (c) a pre-generation self-checklist, and (d) a strictly defined output structure.

Another challenge is that LLMs often struggle to consistently extract correct ISA semantics from large, unstructured PDF manuals. To mitigate this, the ISA documentation is restructured into two parts: (a) JSON representation capturing valid opcodes, operands, and encoding rules; and (b) the original PDF content for high-level architectural context (e.g., memory models, exceptions, and special-purpose registers). This dual-format knowledge base, combined with an automated retrieval-augmentation mechanism, significantly improves code correctness and validity and also reduces the number of required iterations for successful event triggering.

4) *Detection of Event Trigger*: In this state, the agent compiles and simulates the test program generated in the previous stage and produces the VCD trace. This trace is then compared against a golden reference trace, generated by executing an empty main function (i.e., `int main(){ return 0; }`), to detect microarchitectural transition differences. An increase in transitions indicates that the test program successfully triggered the targeted event and reached the termination state. If the

event is not triggered, the framework attributes the failure to one of three primary causes: (a) compilation error, (b) simulation timeout, or (c) absence of event activation. The agent logs the failure reason and associated outputs, which are forwarded to the feedback stage to guide the next iteration of test generation.

5) *Failure Analyzer and Feedback*: The primary objective is to establish a structured ⑤ feedback loop that enables LLM to iteratively refine its output based on the root causes of failure. The agent analyzes failed outputs and incorporates targeted feedback into subsequent prompt iterations. For compiler errors, the agent parses the compiler log into a structured JSON format detailing the error type, location, and cause. For simulation timeouts, the system prompt is updated to instruct the LLM to detect potential infinite loops, add termination conditions, or revise problematic inline assembly segments. When a program compiles and executes but fails to trigger the target event, the agent supplements the prompt with previously successful, event-relevant examples to improve context. Throughout all iterations, the system prompt enforces repair strategies and contractual rules, while the user prompt includes structured error reports and a history of prior code attempts. For generating valid inline assembly, the agent adds a list of frequently used ISA instructions to the user prompt, reducing the model’s need to search the JSON database for targeted instructions. This guided feedback ensures that LLM receives precise, context-aware information, significantly improving convergence toward valid and effective test programs.

IV. RESULTS

We have used the mor1kx Cappuccino core [21] from OpenRISC [22] architecture as the sample design to extract microarchitectural events. For verification, we selected the or1k Marocchino [23] processor, which implements the exact ISA but follows an out-of-order pipeline, and the PicoRV32 [24] processor, which follows a similar ISA and design philosophy. We applied the CHStone [25] and MiBench [26] benchmark suites under random testing using Verilator [27] to obtain cycle-accurate simulations, then computed signal probabilities to identify rare nets. We acknowledge that no fixed benchmark suite is exhaustive; therefore, COVERT does not claim completeness but instead provides a scalable mechanism to maximize the probability of activating rare events. PyVerilog [28] was used to parse the RTL and generate an AST for mapping these rare nets to architectural events. We completed the framework using LangGraph from LangChain [29] to build a state-driven agent for automated test generation. The agent connected to the OpenAI Assistant API [30], using its file assistant to retrieve ISA documentation information. Inside the API, we selected the OpenAI 4.1 model (mini), which provided strong coding ability and offered better cost-performance efficiency compared to larger models. All experiments were performed on a 24-core Intel Core i7-13700F system (2.1 GHz base, 64 GB RAM) running Ubuntu 22.04.5. We have open-sourced all related artifacts for community use (see [16]).

A. Mapping Signals to Architectural Events

Table II summarizes the available signals across the five pipeline modules of the mor1kx Cappuccino core. In principle,

TABLE II
SIGNALS DISTRIBUTION ACROSS PIPELINE MODULES OF MORI-KX CAPPUCCINO WITH
EXAMPLES OF RTL SIGNALS MAPPED TO THEIR HIGH-LEVEL EVENTS.

Stage	Total Signals	Total Events	Rare Events	Example Signal	Corresponding High-Level Event
Decode	56	45	11	decode_op_- movhi_o	Pipeline stall on instruction decode due to dependency on a high immediate value.
Fetch	237	149	65	icache_refill_- done_o	Cache refill completes, allowing instruction fetch to resume.
ALU	93	76	16	overflow_set_o	Detection of signed overflow or division by zero.
LSU	278	173	37	align_err_short	Misaligned memory access.
Ctrl	194	90	25	ctrl_op_- mtps_r_i	Move-to-SPR instruction transfers data from General Purpose Register (PR) to Special PR.

microarchitectural events may be derived from meaningful combinations of multiple signals. However, in our experiments, each signal is treated as a potential architectural event, while excluding signals that never toggled during simulation or failed during AST abstraction. We also reported the number of rare events when rarity threshold $\theta \leq 0.05$ and $\theta \in [0, 1]$. Table II illustrates five representative examples, one from each pipeline module, where RTL signals are mapped to architectural events.

B. Test Program Generation for COTS

To produce valid high-level and bare-metal programs by modern LLMs that reliably trigger specific events requires precise knowledge of register addresses and bit-level configurations, which an LLM cannot infer from ISA documentation. To overcome these limitations, we progressively refined our prompting strategies. Fig. 5 shows how each refinement, starting from baseline prompting to iterative feedback, increased the percentage of events triggered in the sample design.

T1. Baseline Prompting: In the initial stage, only the ISA documentation in PDF format was provided to the LLM (GPT 4.1 mini), with minimal rule-setting restricted to the user role. The event name and a brief architectural summary were supplied as input. Under this setting, we observed the lowest rate of correct test program generation. Most outputs either failed to compile or resulted in simulation timeouts, leading to a lower number of events successfully triggered across pipeline modules shown in Fig. 5.




Fig. 5. Percentage of events triggered in sample design across pipeline modules using four progressively refined prompting strategies (T1–T4).

T2. Role-Guided Prompting: After identifying the common issues, we addressed them in two approaches: (i) we converted the ISA’s table-based descriptions from the PDF into a JSON schema `{id, syntax, encoding, example_hex, description}`, and retaining the PDF only for memory and Special Purpose Register (SPR) details; (ii) we enforced strict coding rules and a repair loop that feeds the compiler and

simulation errors back to the LLM to fix the current program rather than regenerate. We also incorporated compiler header definitions that provide predefined macros and access functions for SPRs. As a result, we were able to reduce errors, which mostly came from invalid opcode/operand choices and SPR bit configurations. Fig. 5 shows that the number of events triggered increased significantly: Fetch from 24% to 47%, LSU from 43% to 68% and Ctrl from 28% to 62%, while Decode and ALU improved slightly since most of their events already required less low-level programming effort.

T3. Incorporating Event Guidance: For events that still failed to trigger, we added higher-level context, i.e., logical summary of the signal’s behavior, guidance on suitable test stimuli, and categories of instructions likely to influence the event. This abstraction allowed the LLM to connect low-level activity with ISA-level outcomes while remaining architecture-invariant. Building on T2, these additions improved the percentage of events triggered across all pipeline modules.

T4. Per Run Role Definition with Exemplars: Finally, T3 was extended with iterative feedback loops where system and user instructions were adjusted based on the current state of program generation. When an event failed to trigger, the framework retrieved a small set of previously successful programs (e.g., 3–5) from the existing pool and supplied them as exemplars alongside the updated instructions. This approach led to further improvements in event triggering, shown in Fig. 5. Even after applying T4, coverage stayed low for the fetch module because the LLM generated small, loop-based tests that tried to trigger memory events, such as instruction cache misses, page faults, cache invalidation via special registers, and instruction bus fetches. These tests did not create the necessary setup, including large and varied code, jumps across far-apart code locations, and long run time that changes cache and memory settings.

C. Event Coverage for COTS

To evaluate the effectiveness of COVERT on unknown COTS processors, we executed the test programs generated from the sample implementation. For COTS, direct one-to-one matching of events generated from signals with the sample design is challenging because of implementation differences. Therefore, we reran the same CHStone and MiBench benchmarks on Marocchino and PicoRV32 to identify rare events. For Marocchino, we simulated module-specific test programs. For PicoRV32 (organized as a single top-level module without distinct pipeline stages), we executed the entire set of test programs. OpenRISC inline assembly was automatically converted to the RV32IMC instruction set for PicoRV32 using LLM. Table III summarizes the results. With a rarity threshold (see Section III-B) of 5%, Marocchino achieved nearly 80% event coverage in the Fetch and LSU modules, while Decode and ALU exceeded 90%, and Ctrl reached over 86%. PicoRV32 showed similar activation rates. As the rarity threshold was relaxed, coverage increased steadily across all modules. Decode and ALU depend less on microarchitectural features, whereas Fetch and LSU involve cache, MMU, and SPR interactions, explaining their slightly lower coverage in Marocchino. PicoRV32, a size-optimized core lacking caches,

TABLE III

EVALUATION OF COVERT USING TEST PROGRAMS FROM THE SAMPLE MORIKX CAPPUCCINO RTL ON TWO TARGETS: MODULE-WISE RARE-EVENT COVERAGE FOR THE OPENRISC MAROCCHINO PROCESSOR AND OVERALL RARE-EVENT COVERAGE FOR THE SINGLE-MODULE RISC-V PICORV32.

module	OpenRISC Marocchino															RISC-V PICORV32IMC		
	Fetch			Decode			ALU			LSU			Ctrl			Combined		
	Rare events	Rare events triggered	% of triggered	Rare events	Rare events triggered	% of triggered	Rare events	Rare events triggered	% of triggered	Rare events	Rare events triggered	% of triggered	Rare events	Rare events triggered	% of triggered	Rare events	Rare events triggered	% of triggered
θ	Test programs simulated = 80			Test programs simulated = 41			Test programs simulated = 70			Test programs simulated = 91			Test programs simulated = 77			Test programs simulated = 364		
0.05	78	62	79.49	21	21	100	12	11	91.67	107	85	79.44	86	74	86.05	101	86	85.15
0.15	110	94	85.45	34	34	100	26	25	96.15	195	172	88.21	95	83	87.37	134	119	88.81
0.25	120	102	85.00	41	41	100	42	41	97.62	220	197	89.55	99	87	87.88	147	132	89.80
0.35	125	106	84.80	50	50	100	56	55	98.21	247	224	90.69	103	91	88.35	154	139	90.26
0.45	130	111	85.38	71	71	100	63	62	98.41	253	230	90.91	108	96	88.89	167	152	91.02
0.55	165	143	86.67	90	90	100	65	64	98.46	258	235	91.09	116	104	89.66	170	155	91.18
0.65	176	153	86.93	94	94	100	67	66	98.51	260	237	91.15	120	108	90.00	176	161	91.48
0.75	180	157	87.22	99	99	100	70	69	98.57	260	237	91.15	120	108	90.00	179	164	91.62
0.85	183	160	87.43	101	101	100	73	72	98.63	261	238	91.19	125	113	90.40	179	164	91.62
1	199	175	87.94	104	104	100	75	74	98.67	266	243	91.35	135	123	91.11	182	167	91.76

memory management units, and access mode, gained only about 2% additional coverage even after running the converted assembly programs. **These results demonstrate that microarchitecturally derived programs from the sample core design remain effective across different implementations, providing high coverage despite the absence of internal design knowledge.** COTS microprocessors that follow a similar ISA will always exhibit overlapping rare events, while unique events arise from implementation specific features. COVERT's event-extraction and activation methodology can scale by adapting to the event space for target architectures. By continually generating microarchitectural events from diverse open-source RTL designs, a procurer can progressively cover both the shared and unique rare events, increase rare-node activation, and trust. A Trojan trigger can be combinational or sequential. Combinational Trojan is triggered by a single event; therefore, this analysis can be considered combinational trigger coverage.




Fig. 6. Percentage of sequential trigger coverage across pipeline modules with joint rarity threshold θ in the OpenRISC Marocchino microprocessor.

D. Complex Trigger Coverage for COTS Marocchino

While an occurrence of one or more distinct events could be considered as combinational triggers, a sequential trigger could be designed using a sequence of multiple events. We extended the analysis by constructing sequential Trojan instances through joint activation of individual events. To keep activation probability low but realistic [31], we keep event combinations within a realistic rarity threshold range shown in Fig. 6. We observed that sequential trigger coverage across all modules of the OpenRISC Marocchino ranged from about 60% in Fetch to nearly 100% in Decode and ALU, with combined joint θ values spanning approximately 4.6×10^{-16} to 0.24.

E. Mathematical Analysis for Coverage

Coverage is calculated as the fraction of Trojans triggered during testing out of the total number of Trojans.

(i) For individual test programs, each Trojan j is targeted by a single test program t_j that triggers it with probability p_j . Assuming all T test programs are applied independently, the expected number of Trojans triggered is:

$$\mathbb{E}[\text{Triggers}] = \sum_{j=1}^T p_j \text{ and Exp. coverage: } \mathbf{Coverage}_{indv} = \frac{1}{N} \sum_{j=1}^N p_j \quad (1)$$

(ii) For combining test programs, the probability that event e_j is triggered by at least one test program is:

$$P_{\text{trigger},j} = 1 - \prod_{t_i \in T_j} (1 - p_{ij}) \quad (2)$$

The expected number of triggered Trojans across all N events is then:

$$\mathbb{E}[\text{Triggers}] = \sum_{j=1}^N \left(1 - \prod_{t_i \in T_j} (1 - p_{ij}) \right) \quad (3)$$

where, $\{p_{ij}\}_{i=1}^n$ be a set of independent probabilities in $[0, 1]$.

$$\text{Exp. coverage : } \mathbf{Coverage}_{comb} = \frac{1}{N} \sum_{j=1}^N \left(1 - \prod_{t_i \in T_j} (1 - p_{ij}) \right) \quad (4)$$

It can be shown that $1 - \prod_{i=1}^n (1 - p_{ij}) \geq \max_i p_{ij}$ with strict inequality if more than one $p_{ij} > 0$. This implies $\mathbf{Coverage}_{comb} \geq \mathbf{Coverage}_{indv}$. All of the above trends were observed in our experiments.

V. DISCUSSION AND CONCLUSION

We have presented COVERT, a practical solution to verification of COTS hardware trust while being scalable (with respect to design size) and flexible (with respect to microarchitectural variations). The central idea behind COVERT, a golden-free COTS trust verification framework, is the use of microarchitectural events, which are implementation invariant, and triggering them multiple times using a set of test programs to statistically maximize the probability of triggering Trojans. To our knowledge, this is the first instance of trust verification of COTS microprocessors, which can be used by Original Equipment Manufacturers (OEMs) to verify the trust of untrusted COTS components before integrating them into systems.

While COVERT is promising, we note that there are significant opportunities to improve coverage, test efficiency, and scalability. It includes (1) accounting for trigger coverage drop even under the same ISA when the target COTS includes unseen implementation features such as memory protection mechanisms or error-correction code, and (2) increased diversity of the test programs in terms of instruction mix. There are several methods to achieve this. For instance, adding a random prefix code to a valid test program generated by COVERT can alter both the instruction mix and execution trace, thereby statistically enhancing the probability of activating a random trigger condition. Our future work will investigate these opportunities.

VI. ACKNOWLEDGMENT

The authors acknowledge support from the Purdue Center for Secure Microelectronics Ecosystem – CSME#210205.

REFERENCES

- [1] S. Shivakumar and C. Wessner. (2022, June) Semiconductors and national defense: What are the stakes? Center for Strategic and International Studies. [Online]. Available: https://www.csis.org/analysis/semiconductors-and-national-defense-what-are-stakes?utm_source=chatgpt.com
- [2] K. Xiao, D. Forte, Y. Jin, R. Karri, S. Bhunia, and M. Tehranipoor, "Hardware Trojans: Lessons Learned after One Decade of Rresearch," *ACM Transactions on Design Automation of Electronic Systems (TODAES)*, vol. 22, no. 1, p. 6, 2016.
- [3] P. P. Sarker, U. Das, N. Varshney, S. Shi, A. Kulkarni, F. Farahmandi, and M. Tehranipoor, "When everyday devices become weapons: A closer look at the pager and walkie-talkie attacks," *arXiv preprint arXiv:2501.17405*, 2025.
- [4] A. Jain, Z. Zhou, and U. Guin, "Survey of recent developments for hardware trojan detection," in *2021 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE, 2021, pp. 1–5.
- [5] J. Cruz, A. Gaikwad, and S. Bhunia, "Automatic hardware trojan insertion using machine learning to generate training datasets," in *2022 IEEE International Symposium on Hardware Oriented Security and Trust (HOST)*. IEEE, 2022, pp. 140–145.
- [6] R. S. Chakraborty, F. Wolff, S. Paul, C. Papachristou, and S. Bhunia, "MERO: A statistical Approach for Hardware Trojan Detection," in *Cryptographic Hardware and Embedded Systems-CHES 2009*. Springer, 2009, pp. 396–410.
- [7] S. Paria, P. Gaikwad, A. Dasgupta, and S. Bhunia, "Latent: Leveraging automated test pattern generation for hardware trojan detection," in *2024 IEEE 33rd Asian Test Symposium (ATS)*, 2024, pp. 1–6.
- [8] Y. Huang, S. Bhunia, and P. Mishra, "Scalable test generation for trojan detection using side channel analysis," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 11, pp. 2746–2760, 2018.
- [9] M. Hasan, J. Cruz, P. Chakraborty, S. Bhunia, and T. Hoque, "Trojan resilient computing in cots processors under zero trust," *IEEE Transactions on Very Large Scale Integration (VLSI) Systems*, vol. 30, no. 10, pp. 1412–1424, 2022.
- [10] 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.
- [11] M. Beaumont, B. Hopkins, and T. Newby, "Safer Path: Security Architecture using Fragmented Execution and Replication for Protection against Trojaned Hardware," in *Proceedings of the Conference on Design, Automation and Test in Europe*. EDA Consortium, 2012, pp. 1000–1005.
- [12] A. Stern, D. Mehta, S. Tajik, F. Farahmandi, and M. Tehranipoor, "Sparta: A laser probing approach for trojan detection," in *2020 IEEE International Test Conference (ITC)*. IEEE, 2020, pp. 1–10.
- [13] S. Yang, P. Chakraborty, P. SLPSK, and S. Bhunia, "Trusted electronic systems with untrusted cots," in *2021 22nd International Symposium on Quality Electronic Design (ISQED)*. IEEE, 2021, pp. 198–203.
- [14] T. Hoque, P. SLPSK, and S. Bhunia, "Trust issues in cots: The challenges and emerging solution," in *Proceedings of the 2020 on Great Lakes Symposium on VLSI*, ser. GLSVLSI '20. New York, NY, USA: Association for Computing Machinery, 2020, p. 211–216.
- [15] S. Paria, A. Dasgupta, and S. Bhunia, "Towards automated verification of ip and cots: Leveraging llms in pre- and post-silicon stages," in *2025 IEEE 43rd VLSI Test Symposium (VTS)*, 2025, pp. 1–5.
- [16] M. Hasan and S. Paria, "COVERT Artifact," 2025. [Online]. Available: <https://github.com/UF-Nelms-IoT-Git-Projects/COVERT>
- [17] J. Cruz and J. Hamlet, "A survey on the design, detection, and prevention of pre-silicon hardware trojans," *IEEE Access*, 2025.
- [18] Z. Zhang, G. Chadwick, H. McNally, Y. Zhao, and R. Mullins, "Llm4dv: Using large language models for hardware test stimuli generation," 2023.
- [19] C. Xiao, Y. Deng, Z. Yang, R. Chen, H. Wang, J. Zhao, H. Dai, L. Wang, Y. Tang, and W. Xu, "Llm-based processor verification: A case study for neuromorphic processor," in *2024 Design, Automation & Test in Europe Conference & Exhibition (DATE)*, 2024, pp. 1–6.
- [20] Y. Deng, R. Chen, C. Xiao, Z. Yang, Y. Luo, J. Zhao, N. Li, Z. Wan, Y. Ai, H. Dai, and L. Wang, "Llm - tg: Towards automated test case generation for processors using large language models," in *2024 IEEE 42nd International Conference on Computer Design (ICCD)*, 2024, pp. 389–396.
- [21] OpenRISC, "mor1kx: An openrisc 1000 processor ip core," <https://github.com/openrisc/mor1kx>, version 5.1.1.
- [22] OpenRISC, "Openrisc project overview," <https://openrisc.io>, accessed: 2025-09-14.
- [23] OpenRISC, "or1k_marocchino: Openrisc processor ip core based on tomasulo algorithm," https://github.com/openrisc/or1k_marocchino.
- [24] YosysHQ, "Picorv32: A size-optimized risc-v cpu," <https://github.com/YosysHQ/picorv32>, 2019, version 1.0.
- [25] Y. Hara, H. Tomiyama, S. Honda, and H. Takada, "Proposal and quantitative analysis of the CHStone benchmark program suite for practical c-based high-level synthesis," in *Proceedings of the IEEE International Symposium on Circuits and Systems (ISCAS)*, 2009, pp. 1–4.
- [26] M. R. Guthaus, J. S. Ringenber, D. Ernst, T. M. Austin, T. Mudge, and R. B. Brown, "MiBench: A free, commercially representative embedded benchmark suite," in *Proceedings of the IEEE International Workshop on Workload Characterization (WWC)*, 2001, pp. 3–14.
- [27] W. Snyder, "Verilator: Fast and open-source verilog/systemverilog simulator," <https://www.veripool.org/verilator/>, accessed: 2025-09-14.
- [28] S. Takamaeda-Yamazaki, "PyVerilog: A python-based hardware design processing toolkit for verilog hdl," in *Proceedings of the 11th IEEE International Conference on Open Source Systems and Technologies (ICOSST)*, 2015, pp. 45–51.
- [29] L. Team, "Langgraph: State-driven workflows for langchain," <https://docs.langchain.com/docs/langgraph/>, accessed: 2025-09-14.
- [30] OpenAI, "Openai assistants api," <https://platform.openai.com/docs/assistants>, accessed: 2025-09-14.
- [31] B. Shakya, T. He, H. Salmani, D. Forte, S. Bhunia, and M. Tehranipoor, "Benchmarking of hardware trojans and maliciously affected circuits," *Journal of Hardware and Systems Security*, vol. 1, no. 1, pp. 85–102, 2017.