

IMS: Intelligent Hardware Monitoring System for Secure SoCs

Wadid Foudhaili^{1*}, Aykut Rencher^{2*}, Anouar Nechi¹, Rainer Buchty¹, Mladen Berekovic¹,
Andres Gomez^{2*}, and Saleh Mulhem^{1*}

¹*Institute of Computer Engineering, Universität zu Lübeck, Lübeck, Germany*

²*Institute of Computer and Network Engineering, TU Braunschweig, Braunschweig, Germany*
wadid.foudhaili@uni-luebeck.de

Abstract—In the modern Systems-on-Chip (SoC), the Advanced eXtensible Interface (AXI) protocol exhibits security vulnerabilities, enabling partial or complete denial-of-service (DoS) through protocol-violation attacks. The recent countermeasures lack a dedicated real-time protocol semantic analysis and evade protocol compliance checks. This paper tackles this AXI vulnerability issue and presents an intelligent hardware monitoring system (IMS) for real-time detection of AXI protocol violations. IMS is a hardware module leveraging neural networks to achieve high detection accuracy. For model training, we perform DoS attacks through header-field manipulation and systematic malicious operations, while recording AXI transactions to build a training dataset. We then deploy a quantization-optimized neural network, achieving 98.7% detection accuracy with $\leq 3\%$ latency overhead, and throughput of > 2.5 million inferences/s. We subsequently integrate this IMS into a RISC-V SoC as a memory-mapped IP core to monitor its AXI bus. For demonstration and initial assessment for later ASIC integration, we implemented this IMS on an AMD Zynq UltraScale+ MPSoC ZCU104 board, showing an overall small hardware footprint (9.04% look-up-tables (LUTs), 0.23% DSP slices, and 0.70% flip-flops) and negligible impact on the overall design’s achievable frequency. This demonstrates the feasibility of lightweight, security monitoring for resource-constrained edge environments.

Index Terms—AXI protocol security, hardware security monitoring, ML-based monitoring, SoC security, protocol-level attacks, denial-of-service, RISC-V.

I. INTRODUCTION

Contemporary electronic devices have become ubiquitous, from smartphones to automotive systems. These devices are mainly powered by system-on-chip (SoC) architectures. Within these SoCs, individual components – so-called Intellectual Property (IP) cores – are interconnected and communicate using on-chip buses. Modern SoC architectures rely on the Advanced eXtensible Interface (AXI) protocol for high-performance communication between IP cores [1]. However, AXI’s design prioritizes performance and flexibility over security, which inadvertently leads to security vulnerabilities and flaws.

Recent security analysis has revealed implementation flaws in AXI interconnects that enable new attack vectors targeting protocol-level vulnerabilities [2], [3]. Analysis tools such as

XRAY have identified 41 distinct implementation vulnerabilities in certain AXI interconnects [2], [4], demonstrating that even protocol-compliant traffic can be exploited to bypass conventional protection mechanisms [2]. These vulnerabilities come from implementation flaws rather than protocol specification issues, yet enable sophisticated attacks [2], [3].

A. SoC Denial of Service

Denial-of-service (DoS) attacks represent a critical attack vector where malicious or compromised components exploit implementation weaknesses and protocol characteristics to disrupt communication among SoC components [5]. Here, *partial blocking* increases the unavailability of some IPs with a possibility of SoC malfunction, while *complete blocking* causes immediate SoC malfunction. Fig. 1 illustrates a typical SoC architecture where the host CPU connects to peripherals and IP cores via the AXI bus. Both malicious and buggy legitimate cores can trigger protocol violations, resulting in a partial or complete DoS. This can be performed or happens through malformed header fields, such as illegal burst lengths, transaction ID reuse, or signal flooding. Such protocol-level malicious operations remain undetectable to external monitoring systems because they occur within the SoC’s internal communication fabric. Consequently, the SoC may continue operating while performance degrades, potentially violating service-level agreements or causing complete DoS.

The current security countermeasures and mechanisms focus primarily on access control, preventing unauthorized memory access through memory protection units [6], [7] and interconnect policies [5], [8]. However, these approaches cannot detect protocol violations where malicious components use legitimate access patterns to violate protocol semantics through header field manipulation. This highlights a new class of SoC security challenges and vulnerabilities:

- C1 New attack vectors that exploit protocol semantics of on-chip buses.
- C2 The inability of existing protection mechanisms to detect or mitigate such low-level real-time threats.

B. Paper Contributions

This paper presents a novel approach to detect and mitigate DoS attacks on SoC AXI buses using Machine Learning

* These authors contributed equally

This work has been partially funded by the German Federal Ministry of Research, Technology and Space (BMFTR) through the project RILKOSAN.

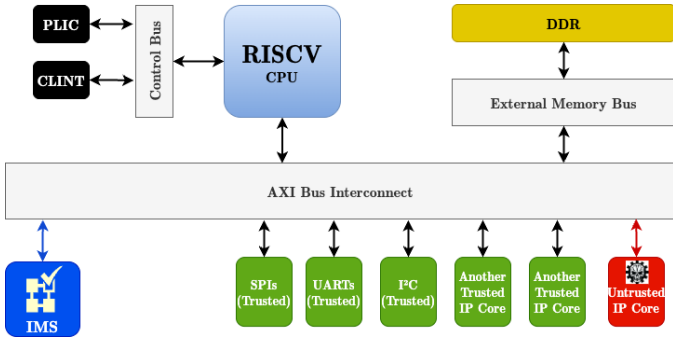


Fig. 1: AXI-Bus Hardware Monitoring System Concept

(ML). We advance the state-of-the-art with the following key contributions:

- We introduce a new security countermeasure called the intelligent hardware monitoring system (IMS) against critical AXI protocol-violation attacks, causing partial or complete DoS of the SoC.
- We show how to build IMS as a machine learning model. Therefore, we start with the generation of a learning dataset, and we propose a thoroughly optimized hardware-friendly ML model by applying several ML optimization techniques.
- Our proposed IMS achieves a highly accurate detection rate of 98.7%, making it suitable for real-world deployment. To demonstrate, we implement IMS in an RISC-V SoC and show the required hardware overhead for IMS integration.

Our dataset for attack detection in AXI bus headers is made available for the public on our repository to allow reproducibility of the presented results and to address the lack of public datasets and benchmarks for protocol-level security research in the SoC environment.

II. BACKGROUND & MOTIVATION

This section introduces our threat model and highlights its related state-of-the-art countermeasures. Then, we motivate our proposed solution.

A. Threat Model

SoC security validation addresses the following security principles: Confidentiality, Integrity, Availability, and Authenticity [5]. This results in three main threat categories [5]: (i) **Availability Violations:** Malicious or malfunctioning components/IP cores may make shared resources unavailable to legitimate users, (ii) **Confidentiality Breaches:** It mainly covers unauthorized access to sensitive data, and (iii) **Integrity Compromises:** Untrusted components try to modify a critical system state.

In this work, we focus on availability-violation threats, which represent a critical and immediate concern in SoC designs as they can cause immediate system-wide failures and are readily observable through performance degradation [2], [4]. While confidentiality and integrity violations are equally

important in comprehensive security frameworks, availability attacks often serve as first indicators of system compromise and directly impact operational functionality.

Therefore, the proposed threat model considers an adversary who can be a malicious IP core or can exploit a malfunctioning IP to disrupt the availability of shared resources on the AXI bus infrastructure. This threat model is an extension of the established model in [8], which focuses on the transaction level of the AXI bus protocol [5], [8].

B. Hardware-Based Security Monitoring

Current hardware security monitoring approaches for SoCs employ two primary strategies: **Memory Protection Unit (MPU)** and **Access Control Policy (ACP)**. MPU-based monitoring approaches enforce access boundaries to specific memory regions, ensuring only authorized components can read or write to them [6], while ACP-based monitoring approaches define communication rules among components/IP blocks using primarily address-based filtering [5], [8]. However, these solutions focus on preventing unauthorized access rather than detecting protocol-level semantic violations. Therefore, they exhibit fundamental limitations, as they rely on static threat models, cannot adapt to evolving attack patterns, and lack the intelligence to distinguish between legitimate transactions and protocol-compliant malicious behavior [9], [10].

C. Research Gap and Motivation

Existing solutions focus on access control and memory protection but lack dedicated real-time protocol semantic analysis [3]. While static verification tools can identify design-time vulnerabilities, they cannot address runtime attacks that evade protocol compliance checks. This indicates a critical security gap, which current countermeasures cannot overcome. Therefore, we introduce an intelligent hardware monitoring system (IMS), deploying an ML model to monitor the AXI. ML has been intensively explored and investigated to monitor network transaction [11], device operation, and CPU execution [12]. Our work addresses this gap by introducing protocol-aware ML models deployed as a lightweight hardware engine for continuous on-chip monitoring. To our knowledge, no existing approach employs ML for real-time AXI protocol monitoring in resource-constrained environments.

III. IMS DESIGN METHODOLOGY AND REALIZATION

This section presents the methodology for developing our intelligent IMS for real-time AXI protocol monitoring and security analysis. Our approach includes dataset generation, preprocessing pipeline design, ML model optimization and evaluation metrics.

A. IMS Design Steps

The key idea of our proposed work is designing an Intelligent Hardware Monitoring System (IMS) for AXI bus traffic. The IMS operates as a passive hardware monitor strategically positioned within the AXI-bus interconnect to observe transaction patterns between system components without disrupting

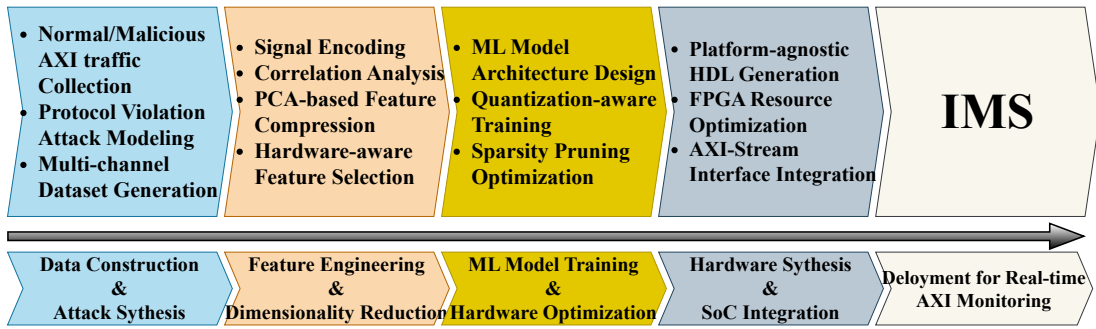


Fig. 2: Overview of the IMS Design Steps

normal operation. The system architecture centers around a SoC interfacing with system resources through a hierarchical bus infrastructure and multiple devices (IP cores) competing for bus access via the AXI protocol, and hence worth use the AXI interconnect. Untrustworthy components or IP cores, such as a malicious master or a malfunctioning device, can disrupt bus operation, potentially leading to system-wide stalls, and partial or complete DoS. Therefore, our target is to design an IMS that can monitor the AXI-bus traffic in real time and analyze the AXI protocol. Fig. 2 illustrates the IMS design steps. First, we build a training dataset by recording the transactions of the AXI bus traffic. The dataset includes normal and malicious data. Indeed, malicious data indicates any abnormal behavior through header-level violations of the AXI protocol or resource starvation attacks. Then, we apply feature engineering methods on the dataset to use the most representative features during the training. We perform ML model optimization techniques and generate the RTL code for this ML model to serve as an IMS for any SoC.

B. Dataset construction and Attack Synthesis

AXI has five channels: Write Address (AW), Write Data (W), Write Response (B), Read Address (AR), and Read Data (R), where anomalies can occur in any channel. The absence of publicly available AXI security datasets necessitated the creation of a comprehensive training corpus. We developed a dual-mode (**Normal** and **Malicious operations**) data collection strategy using Chipyard Platform [13] for RISC-V SoC with standard AXI4 interconnect capabilities. This SoC includes multiple controllers (processors, DMA controllers, hardware accelerators) that act as AXI masters (initiators), competing for access to shared peripherals (e.g., memory, I/O), coordinated by the AXI interconnect’s arbitration logic.

Normal Operation Capture: We collected 16,383 legitimate transactions during standard Linux OS runtime, ensuring comprehensive coverage of typical system behavior across all five AXI channels (AW, W, R, B, AR). This baseline captures the natural transaction patterns expected in production SoC environments. The normal transactions are captured during the System runtime, where nothing disturbs the transactions in any AXI channels.

Malicious Operation Synthesis: To record malicious operations, we introduce malicious behavior to the AXI protocol, covering several attack scenarios where a malicious or malfunctioning master (initiator) IP/component blocks the bus by not completing a write transaction, causing system-wide stalls and hardware-level DoS. We systematically generated 3,242 attack samples representing three critical attack scenarios:

- 1) *Illegal Burst Configurations* [Fig. 3a]: AWLEN values exceeding 15 to force bus re-initialization
- 2) *Transaction ID Exploits* [Fig. 3b]: Duplicate ARID values inducing cache incoherence
- 3) *QoS Signal Flooding* [Fig. 3c]: AWQoS saturation (0xF) to starve low-priority traffic

These attacks lead to DoS of the targeted SoC.

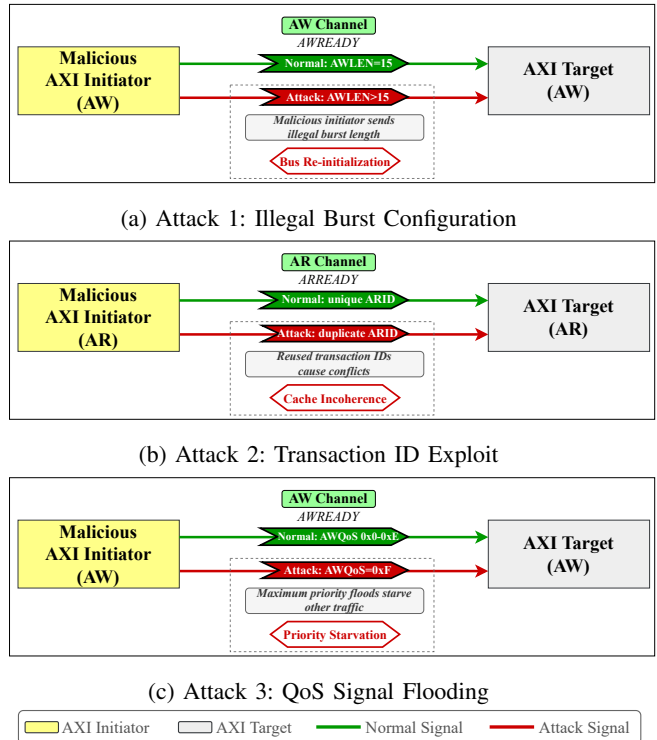


Fig. 3: AXI Protocol Example of Security Vulnerabilities and Attacks Vectors

For instance, a malicious transaction can be injected as follows: After the initiator receives AWVALID in the AW channel, we hold WDATA in the W channel. These attack vectors simulate hardware-level scenarios where malicious initiators block bus transactions, causing system-wide stalls and a DoS.

Our instrumentation employs Vivado’s Integrated Logic Analyzer (ILA) [14] with strategically placed probes at RISC-V initiator and AXI target/device interfaces to mirror all relevant AXI signals. We captured 19,625 transactions (raw samples) divided into 16,383 normal and 3,242 malicious by performing the proposed three attack scenarios. This provides sufficient data for robust ML training. We export the captured data in three formats: raw, VCD, and CSV files, where the raw format is used for hardware instrumentation, VCD for dataset generation, and CSV format for ML-model training. The raw dump file contains 290 signals corresponding to 57 features, represented in their simplest form as bits.

C. Pre-processing Pipeline and Feature Engineering

Our feature engineering methodology transforms raw protocol signals into machine learning-ready representations through a systematic three-stage pipeline designed to maximize information content while minimizing computational overhead.

(1) **Stage 1 - Signal Encoding:** We decode binary and hexadecimal header fields to decimal representation using standardized AXI protocol specifications. After removing 5 ILA-specific debug signals not part of the AXI4 standard, 52 protocol-relevant features remain for subsequent processing.

(2) **Stage 2 - Correlation Analysis:** We apply statistical correlation analysis to identify and eliminate redundant features exhibiting high inter-correlation or constant values across samples. This process reduces the feature space from 52 to 22 dimensions, achieving a 58% reduction while preserving essential discriminative information.

(3) **Stage 3 - Dimensionality Optimization:** Principal Component Analysis (PCA) compresses the remaining features into lower-dimensional representations while retaining statistical significance. Our iterative approach achieves substantial compression: a combination of the original features as 4 components retains 90% variance, 6 components preserve 95%, and 8 components maintain 97% of original information content.

The iterative application of correlation analysis and PCA continues until optimal feature reduction is achieved while maintaining the minimum acceptable variance threshold (90-97%), ensuring that essential attack signatures remain detectable in the compressed feature space.

D. ML Model Architecture and Training

We select a supervised Multilayer Perceptron (MLP) architecture based on its effectiveness with tabular protocol data and compatibility with hardware synthesis frameworks. Our network architecture comprises two hidden layers containing 32 neurons each with ReLU activation functions, followed by

TABLE I: Dataset Statistics and Preprocessing Results

Metric	Normal	Malicious	Mixed Data
Original Features	52	52	52
Post-Correlation Analysis (Features)	22	22	22
Post-PCA (90% variance) (Components)	4	4	4
Post-PCA (95% variance) (Components)	6	6	6
Post-PCA (97% variance) (Components)	8	8	8

a sigmoid output layer optimized for binary classification. The model parameters are in Float32 data format.

Training Configuration: Data partitioning follows standard machine learning practices with 80% allocated for training (augmented via SMOTE for class balance) and 20% reserved for testing. We employ the ADAM optimizer with binary cross-entropy loss and L2 regularization ($\lambda=0.0001$) to prevent overfitting while maintaining generalization capability.

Model Optimization: Hardware deployment requirements necessitated aggressive optimization techniques. We implement QKeras-based [15] quantization-aware training, reducing weight precision to $\langle 8,5 \rangle$ format while applying 80% sparsity pruning during training. These optimizations achieve substantial resource reduction without sacrificing detection accuracy, as validated through comprehensive performance metrics including accuracy, precision, recall, F1-score, and *Area Under the Receiver Operating Characteristic Curve* (AUC-ROC) analysis.

IV. HARDWARE IMPLEMENTATION, OPTIMIZATION AND DISCUSSION

This section presents the results of our experiments, including the performance of the model at different quantization levels, the results of the hardware implementation, and the performance of the attack detection. We also discuss the significance of these findings in the context of AXI protocol security.

A. IMS Implementation and Optimization

To implement the developed MLP as IMS on the targeted Chipyard SoC Platform [13], we convert the trained MLP model into a hardware description language (HDL) using High-Level Synthesis (HLS) tools. We use the HLS4ML [16] framework to convert our trained model into HLS, resulting in a synthesized FPGA bitstream using the AXI-Stream interface. The resulting IP core integrates seamlessly with existing SoC architectures through standardized AXI-Stream interfaces, enabling deployment as a memory-mapped peripheral or dedicated security coprocessor.

To enable efficient hardware mapping, further optimization is needed. We perform a QKeras-based [15] weight quantization down to $\langle 8,5 \rangle$ precision to minimize memory footprint. During training, we apply a constant sparsity pruning (target 80%), forcing a fraction of zero weights to further reduce hardware usage. Table II shows the impact of quantization levels on ML performance. The $\langle 8,5 \rangle$ quantization level achieves optimal performance, even slightly outperforming the

baseline Float32 model due to regularization effects. Performance remains acceptable until extreme quantization ($\langle 2,0 \rangle$), where accuracy drops significantly. For this application, this shows the practical limits of weight compression. We chose the $\langle 8,5 \rangle$ quantization (QKeras [15]) in the experiments, as its recall score of 100% shows a high detection rate. We test its AUC-ROC, the result shows that the $\langle 8,5 \rangle$ quantization (QKeras [15]) has AUC-ROC values exceeding 99%.

TABLE II: Quantization Impact on Model Performance

Quantization Level	Acc(%)	Pr(%)	R(%)	F1(%)
Float32	98.9	94.9	99.99	97.4
$\langle 8,5 \rangle$	99.1	95.8	99.99	97.9
$\langle 8,3 \rangle$	98.7	94.2	99.99	97.0
$\langle 8,1 \rangle$	97.3	91.8	99.99	95.7
$\langle 2,0 \rangle$	85.2	78.9	92.1	85.0

Note: Acc (Accuracy), Pr (Precision), R (Recall), F1 (F1 Score)

In our mapping of the quantized and pruned MLP to FPGA hardware (ZCU104), we aim to cover edge-device requirements with constrained resources. Therefore, we focus on the following metrics:

- **Latency**, measured in milliseconds, quantifies the inference time per sample.
- **Throughput** is expressed by the number of inferences per second (inference/s).
- **Resource Utilization** results from the synthesis and elaboration of the hardware design, reporting the number of DSP slices, LUTs, flip-flops, and block RAM used.

Table III shows the hardware implementations on FPGA ZCU104. The quantized model achieves a noticeable reduction in resources, particularly in DSP slices (99.5% reduction) and overall very low hardware resource usage on a ZCU104, making it suitable for resource-constrained SoC deployments. The low utilization percentages allow significant headroom for additional security features or multiple detector instances.

TABLE III: Detailed FPGA Resource Utilization (ZCU104)

Resource Type	Baseline Model	Quantized Model	Reduction (%)	HW Resource	Usage (%)
DSP Slices	799	4	99.5	1,728	0.23
LUTs	45,238	20,841	53.9	230,400	9.04
Flip-Flops	8,450	3,224	61.8	460,800	0.70
Block RAM	12	8	33.3	312	2.56
Clock Frequency	250 MHz	250 MHz	0	-	-

To evaluate the performance of IMS, we increase AXI bus loads (10%, 25%, etc) while we measure the latency and the throughput of IMS. Table. IV summarizes the IMS performance results, which show a stable IMS performance for different AXI bus loads.

B. IMS Detectability Analysis

To perform such an analysis, we start by investigating and evaluating the overall IMS detection rate, and then we closely look at each attack.

TABLE IV: IMS Performance Metrics vs. System Load

System Load (%)	Inference Latency (ms)	Throughput (inferences/s)
10	1.523	2,567,891
25	1.544	2,542,103
50	1.566	2,509,578
75	1.589	2,478,920
100	1.612	2,445,682

1) *IMS Evaluation*: Table. V compares the performance of the baseline and quantized models. The high recall (R) score indicates that the IMS correctly detects the malicious operations. The high AUC-ROC score means that IMS can reliably separate malicious from normal operations. These results show an outstanding performance of IMS and its high level of detectability.

TABLE V: IMS Performance Evaluation

Model	Acc(%)	P(%)	R(%)	F1(%)	AUC-ROC
IMS (Baseline)	98.9	94.9	99.99	97.4	0.993
IMS (Quantized $\langle 8,5 \rangle$)	99.1	95.8	99.99	97.9	0.993

Note: Acc (Accuracy), Pr (Precision), R (Recall), F1 (F1 Score)

2) *Attack-Specific Detection Evaluation*: Table. VI shows that the proposed IMS demonstrates consistently high detection rates across all attack types, with AWLEN-overflow attacks achieving perfect detection. The performed attack-specific detection evaluation results in the following observations:

- **Burst-Length Exploits**: 100% detection rate for illegal AWLEN values (>15) causing bus-stall attacks
- **QoS Flooding**: 97.3% precision in detecting priority-inversion pattern attacks via abnormal AWQOS signals.
- **Transaction ID Reuse**: 98.1% recall for cache-incoherence scenarios (ARIDS duplication identification)

Mixed attack patterns, representing real-world sophisticated attacks, maintain detection rates above 98%. The detection rate across the attacks validates the model’s robustness against complex threat scenarios.

TABLE VI: Attack-Specific Detection Performance Evaluation

Attack Vector	Samples	T.P.	F.N.	D.R. (%)	Pr (%)
AWLEN Overflow (> 15)	642	642	0	100.0	98.7
ARID Duplication	558	547	11	98.0	96.2
AWQOS Flooding (0xF)	423	411	12	97.2	94.8
AWSIZE Invalid	389	381	8	97.9	95.1
ARRPROT Violation	345	339	6	98.3	96.7
Mixed Attack Pattern	885	873	12	98.6	97.1
Overall	3,242	3,193	49	98.5	96.4

Note: T.P. (True Positives), F.N. (False Negatives), D.R. (Detection Rate), Pr (Precision)

Our findings indicate that the IMS exhibits perfect detection (100% rate) against very critical protocol violation attacks (AWLEN overflow) while maintaining high detection rates ($>97%$ precision) across all attack vectors, including sophisticated mixed attack patterns and a false positive rate far below 1% (at 99% recall). This demonstrates our model’s robustness and reliability in real-world scenarios.

TABLE VII: Comparison of AXI Security & Monitoring Solutions

Reference	Target Prot.	HW/ SW	Timing Metrics	Transac. Level	Phase Level	Protocol Check	Header Process.*	Perf. Metrics	M.O. Supp.†	Scalability
XRAY [2]	AXI	SW	–	✓	✓	✓	Static	–	✓	✓
eXpect [3]	AXI	SW	–	✓	✓	✓	Static	–	✓	✓
TMU Tiny [17]	AXI4	HW	✓	✓	–	–	Basic (error logs)	✓	✓	✓
TMU Full [17]	AXI4	HW	✓	✓	✓	–	Basic	✓	✓	✓
DD-MPU [6]	generic	HW	–	✓	–	–	Addr-only	–	–	✓
VAST IFT [18]	VP Level	SW	–	–	–	–	N/A	–	–	✓
Trust MU [19]	AHB/ Wishbone	HW	✓	✓	–	–	Basic behavior	–	–	–
This work	AXI4	HW	✓	✓	✓	✓	ML-based	DSP, Latency, Throughput	✓	✓ (ML)

***Header Process.** indicates the depth of protocol header field analysis: *Static* = design-time verification only; *Basic* = simple rule-based validation; *Addr-only* = address-based filtering; *ML-based* = intelligent runtime header analysis that detects malformed headers appearing as valid transactions.

†M.O. Supp. denotes Multiple Outstanding transaction (initiated by not completed) support.

V. RELATED WORK

Recently, AXI monitoring has become a crucial task, especially for anomaly and malicious operation detection purposes. Such monitoring increases the SoC availability against AXI DoS and arouses interest in the context of security.

To show the advantages of our proposed IMS, we systematically compare it with the state-of-the-art solutions, shown in Table. VII. Although XRAY [2] and eXpect [3] serve as robust static analysis frameworks, offering significant protocol coverage at both transaction and phase levels during the design phase. Nonetheless, these software-centric strategies present notable operational limitations, particularly their inability to address runtime threats or adapt to the dynamic attack patterns. This static nature prevents the detection of sophisticated runtime exploits that manifest through protocol-compliant malicious operations. The Transaction Monitoring Unit (TMU) variants proposed [17] exhibit hardware-based real-time monitoring capabilities with TMU Full [17] providing comprehensive coverage across transaction and phase-level monitoring. However, these monitoring solutions fundamentally rely on simple rule-based validation methods that lack a semantic comprehension of protocol violations. TMU Tiny [17] achieves minimal overhead through simplified error logging, yet it fails to recognize intricate attack patterns. Conversely, while TMU Full [17] offers broader coverage, it remains fundamentally limited in its ability to differentiate between legitimate transactions and semantically malicious behaviors that comply with protocol standards. The DD-Memory Protection Units (DD-MPU) [6] illustrates memory-centric security strategies that focus solely on address-based filtering mechanisms. DD-MPU achieves efficiency and scalability in hardware implementation, however, it operates independently of the protocol-level semantic violations, which are critical attack vectors. The address-only processing paradigm can not detect changes (manipulations) in header fields and shows a limited security countermeasure against illegal burst configurations or QoS flooding attacks. Vast IFT [18] and Trust MU [19] are specialized frameworks that address distinct protocol contexts. These solutions demonstrate efficiency within their targeted domains but cannot address the challenges of AXI protocol monitoring. Our proposed IMS uniquely processes AXI protocol-header features using an ML-based approach.

Unlike static verification frameworks which operate exclusively at design time, or basic rule-based hardware monitors, our IMS intelligently detects sophisticated attack patterns appearing as legitimate transactions to conventional validation systems.

VI. CONCLUSION

The protocol-level violations against AXI bus used in the system-on-chip (SoC) leads to a system-on-chip’s (SoC) partial or complete denial of service. This work presents the first comprehensive approach for detecting protocol-level malicious operations by using a machine learning (ML) model integrated into RISC-V-based SoCs as a memory-mapped IP core. We address the absence of public datasets for low-level SoC-bus violations by constructing a novel dataset built from captured normal AXI traffic and from systematically injected DoS, both collected via a RISC-V SoC on a Xilinx ZCU104 FPGA platform. The optimized ML model achieves high accuracy (up to 99.11% AUC-ROC) and robust performance in precision, recall, and F1 metrics. The model is deployed as an IP core on the FPGA and demonstrates real-time inference with minimal resource usage (0.23% DSP and 0.70% Flip-flops utilization on a ZCU104), low latency (1.566 ms), and high throughput (>2.5 million inferences/s). This work demonstrates that protocol-aware, ML-based detection of malicious operations can be effectively realized in resource-constrained edge environments, providing a practical path for securing SoC interconnects against sophisticated hardware attacks.

As future work, we aim to investigate the following avenues:

- Broader attack coverage: Extend the attack scenarios to cover additional AXI channels and include protocol violations beyond DoS scenarios
- ML Model diversity: Investigate alternative ML architectures (e.g., LSTM, GNN) for zero-day attack detection.
- Integration with SoC security frameworks: Explore co-design with existing SoC firewalls, hypervisors, or runtime monitors to provide layered, defense-in-depth protection.
- ASIC integration using state-of-the-art technology libraries.

REFERENCES

- [1] R. Patil and P. Sangamkar, "A review of system-on-chip bus protocols," *International Journal of Advanced Research in Electrical, Electronics and Instrumentation Engineering*, vol. 04, pp. 271–281, 01 2015.
- [2] M. Zonta, N. Hinderling, and S. Shinde, "Xray: Detecting and exploiting vulnerabilities in arm axi interconnects," in *2025 Design, Automation and Test in Europe Conference (DATE)*, pp. 1–7, 2025.
- [3] M. Zonta-Roudes, A. Meza, N. Hinderling, L. Deutschmann, F. Restuccia, R. Kastner, and S. Shinde, *eXpect: On the Security Implications of Violations in AXI Implementations*, pp. 183–191. New York, NY, USA: Association for Computing Machinery, 2025.
- [4] Xilinx, "Axi protocol violations," Technical Report PG247, AMD Xilinx, 2022. Accessed: 07 July 2025.
- [5] E. N. D. Coşkun, S. Ahmadi-Pour, M. Hassan, and R. Drechsler, "Security coverage metrics for information flow at the system level," in *2024 29th Asia and South Pacific Design Automation Conference (ASP-DAC)*, pp. 945–950, IEEE, 2024.
- [6] C. Heinz and A. Koch, "Dd-mpu: Dynamic and distributed memory protection unit for embedded system-on-chips," in *International Conference on Embedded Computer Systems*, pp. 285–295, Springer, 2023.
- [7] C. Ewert, A. Neskovic, C. Heinz, F. Muuss, A. Treff, M. Gourjon, R. Buchty, T. Eisenbarth, A. Koch, M. Berekovic, *et al.*, "Lightweight authenticated integration and in-field secure operation of system-in-package," *ACM Transactions on Design Automation of Electronic Systems*, vol. 30, no. 5, pp. 1–23, 2025.
- [8] M. Goli and R. Drechsler, "Early socs information flow policies validation using systemc-based virtual prototypes at the esl," *ACM Trans. Embed. Comput. Syst.*, vol. 23, Aug. 2024.
- [9] M. Ayache, E. Rama, S. Mulhem, M. Berekovic, and M. Korb, "Holistic framework for evaluating the trustworthiness of integrated circuits," in *2024 IFIP/IEEE 32nd International Conference on Very Large Scale Integration (VLSI-SoC)*, pp. 1–4, 2024.
- [10] E. Rama, M. Ayache, R. Buchty, B. Bauer, M. Korb, M. Berekovic, and S. Mulhem, "Trustworthy integrated circuits: From safety to security and beyond," *IEEE Access*, vol. 12, pp. 69603–69632, 2024.
- [11] W. Foudhaili, A. Nechi, C. Thermann, M. Al Johmani, R. Buchty, M. Berekovic, and S. Mulhem, "Reconfigurable edge hardware for intelligent ids: Systematic approach," in *2024 IFIP/IEEE 32nd International Conference on Very Large Scale Integration (VLSI-SoC)*, (Berlin, Heidelberg), pp. 1–6, Springer-Verlag, 2024.
- [12] D. Hirsch, F. Hoffmann, A. Neskovic, C. Thermann, R. Buchty, M. Berekovic, and S. Mulhem, *Efficient AI-based Attack Detection Methods for Sensitive Edge Devices and Systems*, pp. 177–196. Rives Publishers, 02 2024.
- [13] A. Amid, D. Biancolin, A. Gonzalez, D. Grubb, S. Karandikar, H. Liew, A. Magyar, H. Mao, A. Ou, N. Pemberton, P. Rigge, C. Schmidt, J. Wright, J. Zhao, Y. S. Shao, K. Asanović, and B. Nikolić, "Chipyard: Integrated design, simulation, and implementation framework for custom socs," *IEEE Micro*, vol. 40, no. 4, pp. 10–21, 2020.
- [14] I. Xilinx / Advanced Micro Devices, *Vivado Design Suite User Guide: Programming and Debugging (UG908)*, 2025. Version 2025.1; see section "ILA – Integrated Logic Analyzer", especially pages 24–30.
- [15] C. N. Coelho, A. Kuusela, S. Li, H. Zhuang, J. Ngadiuba, T. K. Aarrestad, V. Loncar, M. Pierini, A. A. Pol, and S. Summers, "Automatic heterogeneous quantization of deep neural networks for low-latency inference on the edge for particle detectors," *Nature Machine Intelligence*, vol. 3, p. 675–686, June 2021.
- [16] FastML Team, "fastmachinelearning/hls4ml," 2025.
- [17] C. Liang, T. Benz, A. Ottaviano, A. Garofalo, L. Benini, and D. Rossi, "Towards reliable systems: A scalable approach to axi4 transaction monitoring," in *2025 Design, Automation & Test in Europe Conference (DATE)*, pp. 1–7, 2025.
- [18] E. N. D. Coşkun, M. Hassan, M. Goli, and R. Drechsler, "Vast: Validation of vp-based heterogeneous systems against availability security properties using static information flow tracking," in *2023 24th International Symposium on Quality Electronic Design (ISQED)*, pp. 1–8, IEEE, 2023.
- [19] M. Flaskkamp, C. Klarhorst, and J. Hagemeyer, "Trustworthy system-on-chip by monitoring system behavior at runtime," in *Proceedings of the 1st Safety and Security in Heterogeneous Open System-on-Chip Platforms Workshop (SSH-SoC 2023)*, 2023.