PRIVATEER: Secure FPGA Acceleration for 6G AI Edge Analytics

Aimilios Leftheriotis[†], Ilias Papalamprou[*], Apostolos Garos[‡], Georgios Gardikis[‡], Maria Christopoulou[§], George Xilouris[§], Georgios Livanos[¶], Nikolaos Papadakis[¶], Dimosthenis Masouros[*], George Theodoridis[†], Dimitrios Soudris[*]
[*]National Technical University of Athens, Greece [†]University of Patras, Greece [‡]R&D Department, Space Hellas S.A., Greece
[§]NCSR "Demokritos" Institute of Informatics and Telecommunications, Greece [¶]Infili Technologies S.A., Greece

2025 IEEE Computer Society Annual Symposium on VLSI (ISVLSI) | 979-8-3315-3477-6/25/$31.00 ©2025 IEEE | DOI: 10.1109/ISVLSI65124.2025.11130215

***Abstract*—The progression towards 6G networks promises enhanced performance but introduces significant security and privacy vulnerabilities, particularly at the network edge. This work presents advancements in PRIVATEER, focusing on AI-driven anomaly detection accelerated on FPGAs and robust security countermeasures for FPGA deployments. We detail an Attention-Autoencoder model for detecting DDoS attacks and evaluate its high-performance, energy-efficient FPGA implementation, achieving >8x/>9x latency reduction and >130x/>30x energy savings compared to CPU/GPU baselines, respectively, without accuracy loss. Additionally, we discuss security mechanisms including remote attestation, Physical Unclonable Functions (PUFs), and side-channel mitigation, demonstrating their efficacy with low overhead. These results showcase viable solutions for secure, hardware-accelerated AI analytics in future 6G edge systems.**

***Index Terms*—6G, Security, Privacy, AI, FPGA, Hardware Acceleration, Anomaly Detection, Attestation**

I. Introduction

The relentless evolution of mobile communication technologies is rapidly propelling us from the current 5G era towards the horizon of 6G networks. This next generation of wireless systems is envisioned to deliver transformative capabilities, including unprecedented data rates, near-zero latency, and ubiquitous connectivity, thereby enabling a new wave of innovative applications such as immersive augmented and virtual reality, massive Internet of Things deployments, and sophisticated AI-driven services [1], [2]. While the potential benefits are immense, the increased complexity, decentralization, and reliance on heterogeneous technologies inherent in 6G also introduce significant and novel security and privacy challenges [3]. The vastly expanded attack surface, coupled with the sensitive nature of data processed by 6G applications, necessitates a fundamental rethinking of security paradigms.

The PRIVATEER project[1] is dedicated to addressing these emerging threats by developing a comprehensive suite of privacy-centric security enablers specifically tailored for the demanding 6G environments [4], [5]. The project's philosophy is rooted in the "privacy-by-design" principle, ensuring that data protection and security are integral components of the network architecture from its inception, in alignment with stringent EU standards such as GDPR. PRIVATEER aims to foster trust across a multi-stakeholder ecosystem, encompassing end-users, infrastructure providers, network operators, and service providers.

A cornerstone of 6G networks will be the pervasive integration of Artificial Intelligence (AI) and Machine Learning (ML) [3]. AI/ML algorithms are poised to revolutionize network management, automate complex orchestration tasks, and, crucially, power advanced security mechanisms. One such critical application is anomaly detection, particularly for identifying and mitigating threats like Distributed Denial-of-Service (DDoS) attacks, which can cripple network services by overwhelming them with malicious traffic. However, the sophisticated AI models required for real-time, effective anomaly detection are computationally intensive, often exceeding the processing capabilities of general-purpose CPUs, especially under the strict low-latency and energy-efficiency constraints of edge computing environments [6].

To overcome these computational bottlenecks, the adoption of a heterogeneous computing continuum, incorporating hardware accelerators, is paramount. Among various accelerator options, Field-Programmable Gate Arrays (FPGAs) have emerged as a compelling technology for 6G edge deployments [7]. FPGAs offer a unique blend of reconfigurability, high parallelism, and energy efficiency, making them ideal for accelerating AI inference tasks directly at the network edge, where data is generated and immediate responses are critical.

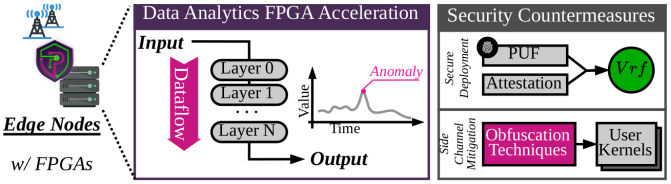

Fig. 1: PRIVATEER concept: FPGA-based AI acceleration and security countermeasures at the 6G edge.

This work has been partially funded by the PRIVATEER project. PRIVATEER has received funding from the Smart Networks and Services Joint Undertaking (SNS JU) under the European Union's Horizon Europe research and innovation programme under Grant Agreement No. 101096110. Views and opinions expressed are however those of the author(s) only and do not necessarily reflect those of the European Union or SNS JU. Neither the European Union nor the granting authority can be held responsible for them.

[1]https://www.privateer-project.eu

While FPGAs provide significant performance advantages, their integration into complex, and potentially untrusted, edge infrastructures introduces new security vulnerabilities. These programmable hardware devices can become targets for a range of attacks, including intellectual property (IP) theft through bitstream reverse engineering, malicious hardware Trojan insertions, unauthorized modifications of the deployed accelerator (kernel), and physical attacks such as side-channel analysis to extract sensitive information (e.g., cryptographic keys or model parameters) [8], [9]. Ensuring the integrity, authenticity, and confidentiality of FPGA-based accelerators is therefore crucial for the secure operation of 6G networks.

This paper presents significant advancements within the PRIVATEER project, focusing on the secure deployment and acceleration of AI-powered data analytics on FPGAs at the 6G network edge. Fig. 1 illustrates the core concept addressed in this work, highlighting the synergy between FPGA-based AI acceleration for network analytics and the essential security countermeasures required to protect these deployments.

The primary contributions of this paper are threefold:

- **Advanced AI Model for Anomaly Detection:** We detail an Attention-Autoencoder (Att-AE) ML model, designed for detecting DDoS attacks by analyzing multivariate time-series data from 6G network monitoring functions.
- **High-Performance FPGA Acceleration:** We present a custom FPGA architecture for the Att-AE model, optimized for low-latency inference and energy efficiency, demonstrating its superiority over CPU and GPU implementations for edge deployment scenarios.
- **Robust FPGA Security Countermeasures:** We describe and evaluate a suite of security mechanisms for FPGA-based accelerators, including remote attestation protocols for integrity verification, the use of Physical Unclonable Functions (PUFs) for secure device identification and key management, and techniques for mitigating side-channel attacks.

The remainder of this paper is structured as follows: Section II elaborates on the AI-driven hardware-accelerated security analytics and the secure deployment mechanisms for FPGA-based kernels developed within PRIVATEER. Section III presents a comprehensive evaluation of these enablers, quantifying the performance of the AI model, the efficiency of the FPGA accelerator, and the effectiveness of the security countermeasures. Finally, Section IV summarizes our findings and outlines future research directions. The developed components and methodologies are being made available through PRIVATEER's open-source repositories[2].

II. PRIVATEER's Security Enablers

This section delves into the technical specifics of the key security enablers developed by PRIVATEER, focusing on AI-driven anomaly detection and the mechanisms for its secure hardware acceleration on FPGAs at the network edge.

[2]https://github.com/privateer-project

A. AI-Driven Hardware Accelerated Security Analytics

The proliferation of connected devices and the increasing complexity of network traffic in 6G environments necessitate sophisticated and timely threat detection capabilities. Maintaining the reliability and security of 6G networks hinges on the ability to efficiently identify and respond to malicious activities. DDoS attacks, in particular, pose a significant threat to 6G infrastructure, capable of disrupting critical services by inundating network resources [10].

PRIVATEER addresses this challenge by developing advanced AI-based security analytics. These analytics leverage data streams from the Network Data Analytics Function (NWDAF), a key component in 5G and future 6G architectures responsible for collecting and analyzing network data to provide insights for network automation and optimization. By processing NWDAF data, our AI models can learn patterns of normal network behavior and detect deviations indicative of attacks. To meet the stringent sub-millisecond latency requirements often encountered at the network edge [6], these AI models are deployed on FPGAs, enabling low-latency and energy-efficient execution.

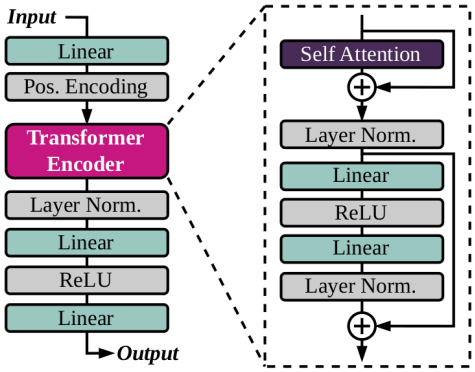

Fig. 2: Att-AE model architecture

1) Attention-Autoencoder (Att-AE) for Anomaly Detection: This work advances PRIVATEER's anomaly detection capability using an Attention-Autoencoder (Att-AE) architecture, illustrated in Fig. 2. The fundamental principle of an autoencoder is to learn an efficient representation (encoding) of input data and then reconstruct the original input from this representation (decoding). When trained exclusively on normal network behavior data from NWDAF streams, the Att-AE model becomes adept at reconstructing such typical patterns with low error. Conversely, when anomalous sequences, such as those indicative of a DDoS attack, are fed into the trained model during inference, they typically result in a significantly higher reconstruction error, thereby enabling their detection.

The encoder component of our Att-AE model is specifically designed to capture intricate temporal dependencies within the sequential network data. Input sequences first pass through a linear layer for embedding, followed by the addition of positional encoding to provide the model with information about the order of data points. The core of the encoder is

a Transformer encoder layer [11], which utilizes the self-attention mechanism. This mechanism allows the model to weigh the importance of different parts of the input, proving highly effective for learning complex patterns in time-series data and generating a rich latent representation. The decoder component then takes this latent representation and aims to reconstruct the original input sequence. In our Att-AE architecture, the decoder consists of a series of linear layers that progressively transform the latent representation back to the dimensionality of the input data.

2) FPGA Acceleration of the Att-AE Model: To facilitate the deployment of the computationally intensive Att-AE model at the network edge while meeting stringent latency and energy constraints, we have developed a dedicated, custom FPGA hardware architecture.

Our FPGA implementation leverages a dataflow paradigm, structuring the computation as a pipeline of concurrent hardware modules communicating through FIFO buffers. This approach enables coarse-grained parallelism at the model layer level, where different layers of the Att-AE (e.g., linear layers, attention mechanism, normalization) operate concurrently on different segments of the data as it streams through the pipeline. This deep pipelining helps achieve high clock frequencies and efficient resource utilization by breaking down complex computations into smaller, manageable stages.

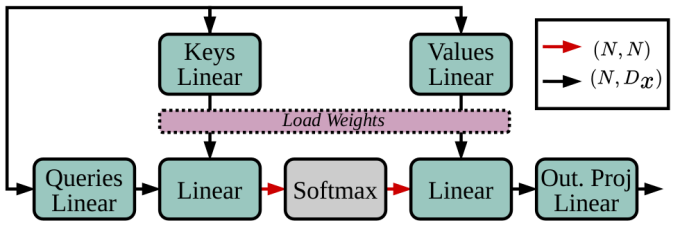

Fig. 3: Representation of Self-Attention Layer mapping to Linear Layers

A detailed analysis of the Att-AE model's computational graph revealed that linear (fully connected) layers constitute the most significant computational bottleneck. The model comprises several explicit linear layers, and critically, the computationally intensive self-attention mechanism can be effectively decomposed into multiple linear layer operations, as conceptually illustrated in Fig. 3. This decomposition highlights that optimizing linear layer execution is paramount for overall accelerator performance.

Listing 1: HLS Pseudo-code representation of the linear layer computation.

```
for(int t = 0; t < Nts; t++){
    accums[Fout] = bias
    for(int x = 0; x < Fin; x++)
        input = input_stream.read()
        for(int y = 0; y < Fout; y++)
            #pragma HLS UNROLL
            accums[y] += input * w[y,x];
        }
    }
    output_stream.write(accums);
}
```

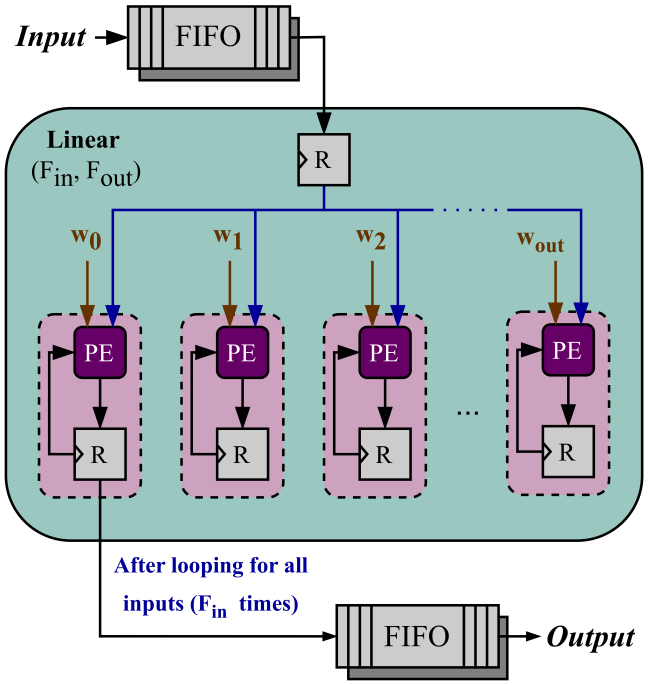

Fig. 4: High-level overview of implemented Linear layer accelerator. Color-coded to match the Pseudo-code in Listing 1.

Consequently, we designed a highly parameterized and optimized hardware module specifically for accelerating these linear layer computations. This module is architected for efficient processing by maximizing parallelism in multiply-accumulate operations and by carefully managing data movement. Fig. 4 provides a high-level schematic of this linear accelerator core, while Listing 1 presents a pseudo-code representation of its computational flow, with color coding to aid understanding.

The operational principle of the linear layer module is as follows: Input activation values (F_{in}). are streamed into the module. As depicted in Fig. 4, each incoming activation value is broadcast in parallel to an array of Processing Elements (PEs). Concurrently with receiving the broadcast activation, each PE accesses its assigned local weight. The PE then performs a multiplication between the broadcast input activation and its local weight, adding the result to an internal accumulator register dedicated to a specific output feature. This multiply-and-accumulate step occurs in parallel across the PEs. This entire sequence is iterated for all F_{in} input activations of the current input. Once all F_{in} activations have been processed, the final accumulated sums which constitute the F_{out} dimensional output vector, are written to an output FIFO for consumption by the subsequent layer in the model. The configurability of this linear layer module, allowing adjustments, via compile-time parameters, to the number of PEs, internal data widths, and tiling strategies , facilitates Design Space Exploration (DSE) to trade-off resource utilization against throughput and latency for different FPGA targets.

Beyond the linear layers, dedicated hardware modules were also developed for other essential components of the Att-

AE model, including Positional Encoding, Layer Normalization, and the Softmax activation function. These modules are designed to seamlessly integrate into the overall dataflow architecture, consuming data from input FIFOs and producing results onto output FIFOs, thereby maintaining pipeline concurrency and maximizing performance.

Approximation techniques, such as quantization to fixed-point arithmetic (e.g., Q8.24 format, providing 24 fractional bits for precision and 7 integer bits plus sign for dynamic range), are carefully applied to balance computational reduction with the preservation of anomaly detection accuracy. The primary goal is to ensure that any approximations do not significantly degrade the model's ability to distinguish anomalous traffic, while yielding substantial benefits in terms of hardware efficiency.

B. Secure Deployment of FPGA-based Kernels

PRIVATEER employs multiple mechanisms to ensure the secure deployment of services, including accelerated data analytics. The proposed security countermeasures focus on verifying the integrity and authenticity of all components provided by multiple stakeholders in a 6G network. At its core, the system uses custom *remote attestation* protocol [12], where an external attestation server communicates with individual edge nodes to verify received values, that include cryptographic checksums, digital signatures etc. The process includes authentication of the underlying infrastructure, as well as verification of the deployed accelerated kernel. To enhance system security, PRIVATEER leverages hardware components on the FPGA, providing the necessary security functions.

Apart from the convential operations, e.g., AES encryption, a fundamental requirement for securely configuring any device is reliable access to cryptographic keys. Traditionally, these keys are stored in non-volatile memory (NVM), which makes them vulnerable to various attacks, including side-channel exploits. To mitigate this risk, *Physical Unclonable Functions (PUFs)* have emerged as a promising solution. PUFs exploit inherent, uncontrollable physical variations in semiconductor manufacturing to generate unique, device-specific cryptographic keys. Unlike stored keys, PUF-based keys are derived dynamically and never persistently stored, significantly reducing the risk of extraction by malicious users. This makes PUFs a lightweight and tamper-resistant solution for authentication and key generation in secure FPGA configuration.

PRIVATEER enhances the reliability and uniqueness of PUF responses by carefully selecting and tuning the PUF architecture implemented in the Programmable Logic (PL) of the FPGA. The custom PUFs developed for this system exploit manufacturing variability specifically in core FPGA resources, such as Addressable Shift Registers (ASRs) and carry chains. [13]. In contrast to traditional Ring Oscillator-based PUFs [14], the chosen PUF architecture has advantages, in terms of reliability. The selected PUF architecture intentional glitches are generated in the ASRs. In each PUF cell, the output register is set to '1' or '0' based on which glitch arrives first. Due to inherent process variations during FPGA manufacturing, the arrival order is unpredictable, ensuring that each device produces a unique and random PUF response. Furthermore, to enhance the reliability of PUF responses, post-processing techniques are applied. In particular, we use majority voting to improve the stability of the PUF output. For each user input, the PUF is challenged multiple times and for every bit position, the final response is determined by selecting the value that appears most frequently across the responses.

Additionally, PRIVATEER incorporates *obfuscation strategies*, including power masking techniques, to mitigate side-channel attacks aimed at extracting sensitive data, such as encryption keys [15]. Our proposed solution relies on obfuscating the power samples collected by malicious users, using power waster circuits that are deployed on the FPGA.

III. Evaluation of Security Enablers

A. Anomaly Detection Model Evaluation

To evaluate the efficacy of PRIVATEER's Att-AE model in detecting network anomalies, we utilized the NCSRD-DS-5GDDoS v3.0 dataset [16]. This dataset is specifically curated for research in 5G security and comprises multivariate time-series data, capturing a rich set of metrics related to both 5G radio access network and core network components during normal operation and synthetically generated DDoS attacks.

The Att-AE model configuration evaluated here processes input sequences of 12 timesteps, where each timestep consists of 8 distinct features. The encoder part of the model first embeds these input sequences into a 32-dimensional space. This is followed by a Transformer encoder layer, which leverages its self-attention mechanism to capture complex temporal dependencies within the sequence. The decoder then reconstructs the original 8-dimensional input features per timestep using two fully connected layers.

TABLE I: Anomaly Detection Metrics

ROC-AUC	Precision	Recall	F1 Score
0.9543	0.6190	0.9008	0.6634

The performance of the Att-AE model is summarized in Table I. The model demonstrates strong discriminative power in identifying anomalies, achieving a Receiver Operating Characteristic Area Under Curve (ROC-AUC) of 0.9543. This metric indicates a high probability that the model will rank a randomly chosen positive instance (anomaly) higher than a randomly chosen negative instance (normal traffic). The precision, which measures the proportion of correctly identified anomalies among all instances flagged as anomalous, is 0.6190. While this value suggests some false positives, the recall of 0.9008 is notably high, signifying that the model successfully identifies 90.08% of all actual DDoS attacks present in the dataset. High recall is often critical in security applications to minimize missed threats. The F1 Score, which is the harmonic mean of precision and recall, stands at 0.6634, providing a balanced measure of the model's overall detection effectiveness against this challenging dataset.

B. FPGA Acceleration Performance

The performance of our custom FPGA-accelerated Att-AE model was evaluated against conventional software implementations on diverse hardware platforms. The experimental testbed included:

- `CPU`: An Intel Xeon Gold 6530 processor.
- `GPU`: An NVIDIA V100 graphics processing unit.
- `FPGA-A`: The Att-AE accelerator implemented on an AMD ALVEO U280 FPGA card.
- `FPGA-B`: The Att-AE accelerator implemented on an AMD ZCU104 MPSoC FPGA.

For the FPGA implementations, numerical precision was carefully managed using 32-bit fixed-point arithmetic (Q8.24 format) for both weights and activations. This choice provided sufficient precision to ensure functionally lossless execution, meaning the FPGA accelerators produce anomaly detection results identical to the floating-point CPU and GPU baselines, without any accuracy degradation due to quantization.

TABLE II: FPGA Resource Utilization (%) for the Att-AE Accelerator

Device	LUT (%)	FF (%)	BRAM (%)	DSP (%)
`FPGA-A`	26.11	12.87	39.74	34.72
`FPGA-B`	43.04	18.52	77.08	18.06

Table II presents the FPGA resource utilization for the Att-AE accelerator on both target platforms. On the `FPGA-A` the design utilizes 26.11% of Look-Up Tables (LUTs), 12.87% of Flip-Flops (FFs), 39.74% of Block RAM (BRAM), and 34.72% of Digital Signal Processors (DSPs). For the `FPGA-B` the corresponding utilization figures are 43.04% for LUTs, 18.52% for FFs, 77.08% for BRAMs, and 18.06% for DSPs. These figures demonstrate an efficient mapping of the Att-AE model onto the FPGA fabric. While `FPGA-B` shows a higher percentage utilization for BRAM, reflecting its relatively smaller on-chip memory capacity compared to the Alveo U280, both platforms retain considerable headroom. This unused resource capacity is highly advantageous, as it allows for the potential co-location of other critical functions, such as security modules for encryption or attestation, or even scaling the analytics capability by instantiating multiple accelerator cores if higher throughput is demanded by specific 6G edge scenarios.

Fig. 5 illustrates the average inference latency (Fig. 5a) and energy consumption per inference (Fig. 5b) over 1000 inferences. The `FPGA-B` implementation achieves the lowest average inference latency at a mere 0.054 ms. This is a substantial improvement, translating to an 8.3x speedup over the `CPU` and a 9.9x speedup over the `GPU`. Such ultra-low latency is critical for enabling real-time responses to detected threats at the network edge. The `FPGA-A` also delivers impressive acceleration, with an average latency of 0.076 ms. The moderately higher latency of `FPGA-A` compared to `FPGA-B` is primarily attributed to the communication overheads inherent in the PCIe-based ALVEO platform, whereas the ZCU104

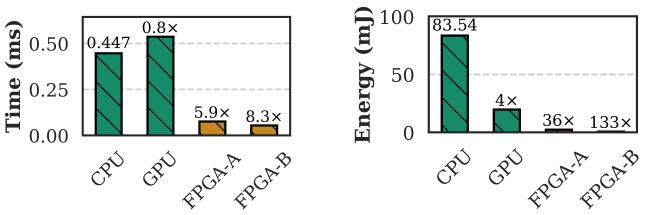

(a) Latency per inference (ms) (b) Energy per inference (mJ)

Fig. 5: Latency and energy-per-inference comparison for the Att-AE model across different hardware platforms. Results are averaged over 1000 inferences.

benefits from a more tightly integrated Arm processing system and programmable logic.

In terms of energy efficiency, as shown in Fig. 5b, the `FPGA-B` platform demonstrates remarkable performance, consuming only 0.627 mJ per inference. This represents a drastic reduction in energy usage: 133.2x less energy than the `CPU` (consuming approx. 83.5 mJ) and 31.4x less energy than the `GPU` (consuming approx. 19.7 mJ). The `FPGA-A` implementation also provides considerable energy savings, consuming approximately 2.3 mJ per inference, still significantly outperforming both CPU and GPU. The superior energy efficiency of `FPGA-B` is substantially influenced by the lower overall power envelope of the ZCU104 MPSoC board (typically around 10-15W under load) compared to the ALVEO U280 card, which can draw significantly more.

Crucially, as mentioned earlier, these significant gains in latency and energy efficiency are achieved without any degradation in the anomaly detection accuracy. The careful use of 32-bit fixed-point (Q8.24) arithmetic in the FPGA design ensures that the precision is sufficient to mirror the floating-point calculations of the CPU/GPU baselines, thus maintaining the model's effectiveness as reported in Section III-A. This combination of high performance, low energy consumption, and preserved accuracy validates the FPGA acceleration approach for deploying sophisticated AI analytics in future 6G edge systems.

C. Secure FPGA Configuration

Remote Attestation: The overhead introduced by the proposed secure deployment methods for FPGA-based services, such as the accelerated AE, is evaluated against a baseline unsecured deployment. The execution time overhead of the remote attestation is minimal, under 7 sec. in both edge and far-edge FPGA devices.

PUF: We evaluate the proposed PUF module on am AMD ZCU104 MPSoC FPGA board. To interpret the collected experimental results, we focus on two key metrics: *Reliability*, which measures the consistency of PUF responses under repeated evaluations, and *Uniformity*, which assesses the balance of '1's and '0's in the output, with ideal values close to 50%, indicating a well-distributed response. As shown in Fig. 6a the PUF achieves reliability above 96% across the majority of

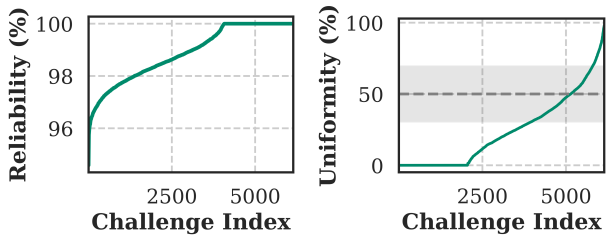

(a) Reliability (b) Uniformity

Fig. 6: PUF Evaluation over different metrics.

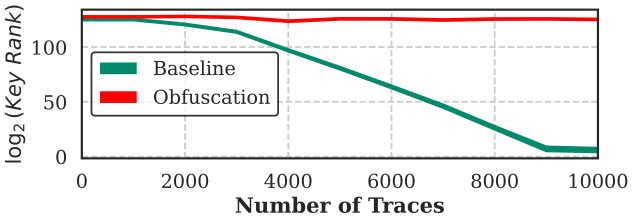

Fig. 7: Comparison of AES encryption key recovery with and without the proposed obfuscation technique.

input challenges, which is considered sufficiently high for our application. Furthermore, Fig. 6b demonstrates that a large number of PUF responses (over 1500) achieve acceptable levels of uniformity. This characteristic is crucial for the overall system, as a high number of valid PUF challenge-response pairs strengthens resistance to replay attacks, which is critical in secure service deployments that rely on remote attestation protocols. In terms of execution time, the overhead for generating a single PUF response is approximately 110 ms, with an encryption key width of 256 bits.

Obfuscation Techniques: We evaluate the effectiveness of the proposed power obfuscation solution by measuring the attacker's ability to extract the AES encryption key. As a proof of concept, we use the side channel attack setup from [15], in a Nexys A7 (100T) FPGA board. The collected power samples from the on-board sensors are analyzed using Correlation Power Analysis (CPA). To assess the effectiveness of the attack in extracting the encryption key, we use the key rank metric. This metric leverages the correlation values obtained through CPA to estimate how much effort remains for an attacker to identify the correct key. As the key rank converges to zero, it indicates that the secret key has been successfully recovered. As illustrated in Fig. 7, an attacker using 10,000 power traces can successfully recover the key in the unprotected scenario (curve colored green). In contrast, with our proposed power obfuscation countermeasure in place, the same number of traces is insufficient to retrieve the key, as higher key ranks reflect increased uncertainty and reduced attack effectiveness (curve colored red).

IV. Conclusion & Future Work

This work demonstrated PRIVATEER's secure, hardware-accelerated AI for 6G edge analytics. We detailed an Attention-Autoencoder model for effective DDoS detection and demonstrated its efficient implementation on FPGA platforms, achieving significant latency and energy improvements while maintaining accuracy. Furthermore, we outlined and evaluated key security countermeasures for FPGAs, confirming their effectiveness with minimal overhead. Future work targets orchestrator integration and broader threat detection.

References

[1] A. Dogra, R. K. Jha, and S. Jain, "A survey on beyond 5g network with the advent of 6g: Architecture and emerging technologies," *IEEE access*, vol. 9, pp. 67512–67547, 2020.

[2] V.-L. Nguyen, P.-C. Lin, B.-C. Cheng, R.-H. Hwang, and Y.-D. Lin, "Security and privacy for 6g: A survey on prospective technologies and challenges," *IEEE Communications Surveys & Tutorials*, vol. 23, no. 4, pp. 2384–2428, 2021.

[3] W. Jiang, B. Han, M. A. Habibi, and H. D. Schotten, "The road towards 6g: A comprehensive survey," *IEEE Open Journal of the Communications Society*, vol. 2, pp. 334–366, 2021.

[4] D. Masouros, D. Soudris, G. Gardikis, V. Katsarou, M. Christopoulou, G. Xilouris, H. Ramón, A. Pastor, F. Scaglione, C. Petrollini, *et al.*, "Towards privacy-first security enablers for 6g networks: the privateer approach," in *International Conference on Embedded Computer Systems*, pp. 379–391, Springer, 2023.

[5] I. Papalamprou, A. Leftheriotis, A. Garos, G. Gardikis, M. Christopoulou, G. Xilouris, L. Argyriou, A. Karamatskou, N. Papadakis, E. Kalotychos, *et al.*, "Multi-partner project: Secure hardware accelerated data analytics for 6g networks: The privateer approach," in *2025 Design, Automation & Test in Europe Conference (DATE)*, pp. 1–4, IEEE, 2025.

[6] M. A. Ferrag, O. Friha, B. Kantarci, N. Tihanyi, L. Cordeiro, M. Debbah, D. Hamouda, M. Al-Hawawreh, and K.-K. R. Choo, "Edge learning for 6g-enabled internet of things: A comprehensive survey of vulnerabilities, datasets, and defenses," *IEEE Communications Surveys & Tutorials*, 2023.

[7] V. Ziegler, H. Viswanathan, H. Flinck, M. Hoffmann, V. Räisänen, and K. Hätönen, "6g architecture to connect the worlds," *IEEE Access*, vol. 8, pp. 173508–173520, 2020.

[8] J. Zhang and G. Qu, "Recent attacks and defenses on fpga-based systems," *ACM Transactions on Reconfigurable Technology and Systems (TRETS)*, vol. 12, no. 3, pp. 1–24, 2019.

[9] W. Hu, C.-H. Chang, A. Sengupta, S. Bhunia, R. Kastner, and H. Li, "An overview of hardware security and trust: Threats, countermeasures, and design tools," *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, vol. 40, no. 6, pp. 1010–1038, 2020.

[10] S. B. Prathiba, G. Raja, S. Anbalagan, K. Arikumar, S. Gurumoorthy, and K. Dev, "A hybrid deep sensor anomaly detection for autonomous vehicles in 6g-v2x environment," *IEEE Transactions on Network Science and Engineering*, vol. 10, no. 3, pp. 1246–1255, 2022.

[11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin, "Attention is all you need," *Advances in neural information processing systems*, vol. 30, 2017.

[12] I. Papalamprou, N. Fotos, N. Chatzivasileiadis, A. Angelogianni, D. Masouros, and D. Soudris, "Post-quantum and blockchain-based attestation for trusted fpgas in b5g networks," in *2025 IEEE International Symposium on Circuits and Systems (ISCAS)*, pp. 1–5, IEEE, 2025.

[13] F.-J. Streit, P. Krüger, A. Becher, J. Schlumberger, S. Wildermann, and J. Teich, "Choice–a tunable puf-design for fpgas," in *2021 31st International Conference on Field-Programmable Logic and Applications (FPL)*, pp. 38–44, IEEE, 2021.

[14] M. T. Rahman, D. Forte, J. Fahrny, and M. Tehranipoor, "Aro-puf: An aging-resistant ring oscillator puf design," in *2014 design, automation & test in Europe conference & exhibition (DATE)*, pp. 1–6, IEEE, 2014.

[15] D. Spielmann, O. Glamočanin, and M. Stojilović, "Rds: Fpga routing delay sensors for effective remote power analysis attacks," *IACR Transactions on Cryptographic Hardware and Embedded Systems*, pp. 543–567, 2023.

[16] N. C. of Scientific Research "Demokritos" and S. H. (Greece), "NCSRD-DS-5GDDoS: 5G Radio and Core metrics containing sporadic DDoS attacks," Oct. 2024. https://doi.org/10.5281/zenodo.13900057.