EasyChair Preprint

№ 11845

NEUROSEC: FPGA-Based Neuromorphic Audio Security

Murat Isik, Hiruna Vishwamith, Yusuf Sur, Kayode Inadagbo and I. Can Dikmen

January 22, 2024

NEUROSEC: FPGA-Based Neuromorphic Audio Security

Murat Isik[1], Hiruna Vishwamith[2], Yusuf Sur[3], Kayode Inadagbo[4], and I. Can Dikmen[5]

[1] Drexel University, Philadelphia, PA, USA
mci38@drexel.edu
[2] University Of Moratuwa, Moratuwa, Sri Lanka
vishwamithpgh.20@uom.lk
[3] Abdullah Gul University, Kayseri, Turkey
yusuf.sur@agu.edu.tr
[4] Prairie View A&M University, Prairie View, TX, USA
kayodeinadagbo@gmail.com
[5] Temsa R&D Center, Adana, Turkey
can.dikmen@temsa.com

Abstract. Neuromorphic systems, inspired by the complexity and functionality of the human brain, have gained interest in academic and industrial attention due to their unparalleled potential across a wide range of applications. While their capabilities herald innovation, it is imperative to underscore that these computational paradigms, analogous to their traditional counterparts, are not impervious to security threats. Although the exploration of neuromorphic methodologies for image and video processing has been rigorously pursued, the realm of neuromorphic audio processing remains in its early stages. Our results highlight the robustness and precision of our FPGA-based neuromorphic system. Specifically, our system showcases a commendable balance between desired signal and background noise, efficient spike rate encoding, and unparalleled resilience against adversarial attacks such as FGSM and PGD. A standout feature of our framework is its detection rate of 94%, which, when compared to other methodologies, underscores its greater capability in identifying and mitigating threats within 5.39 dB, a commendable SNR ratio. Furthermore, neuromorphic computing and hardware security serve many sensor domains in mission-critical and privacy-preserving applications.

Keywords: neuromorphic computing, FPGA, hardware security, audio processing

1 Introduction

Computer hardware that emulates the intricate functions of the human brain has been termed neuromorphic hardware. Drawing inspiration from biological neural systems, neuromorphic hardware aims to replicate the way these systems process information, bridging the gap between biological cognition and artificial computation. Neuromorphic computing represents a paradigm shift from traditional computing methodologies. At its core, it seeks to emulate the brain's neural structures and functionalities, offering a more natural and efficient way to process information. The significance of neuromorphic computing lies in its

potential to revolutionize various domains, from artificial intelligence to robotics, by providing systems that can learn, adapt, and evolve in real-time [6,7,16]. Field-Programmable Gate Arrays (FPGAs) have emerged as a pivotal component in the neuromorphic computing landscape. Their inherent reconfigurability and parallel processing capabilities align seamlessly with the demands of neuromorphic systems. FPGAs offer the flexibility to design and customize neuromorphic architectures, enabling researchers and engineers to experiment with and optimize neural network designs, thereby pushing the boundaries of what neuromorphic systems can achieve. As with any computing system, security remains paramount in neuromorphic systems. Given their potential applications in sensitive areas such as defense, healthcare, and finance, ensuring the integrity, confidentiality, and availability of data processed by neuromorphic systems is crucial. Furthermore, the unique architecture and operation of neuromorphic systems present both challenges and opportunities in the realm of security, requiring specialized approaches to safeguard them against threats [2, 9, 10, 15, 17, 20]. Neuromorphic hardware, inspired by the intricate functions of the human brain, seeks to bridge the gap between biological cognition and artificial computation. This approach represents a paradigm shift, offering a more natural and efficient way to process information. FPGAs, with their inherent reconfigurability and parallel processing capabilities, have emerged as a pivotal component in this landscape, enabling the design and customization of neuromorphic architectures. Given the potential applications of neuromorphic systems in sensitive areas such as defense and healthcare, ensuring their security is paramount. The unique architecture of these systems presents both challenges and opportunities in the realm of security.

In this paper, we present the following contributions:

- We explore SNN-based neuromorphic audio processing, a niche compared to image/video processing.
- We analyze security threats in neuromorphic audio, emphasizing adversarial attacks like FGSM and PGD that introduce audio artifacts.
- Our FPGA-integrated system boasts a 94% detection rate, efficient spike encoding, and a balanced signal-to-noise ratio.
- We compare our framework with existing methods, highlighting its superior threat detection and mitigation within a favorable SNR.

2 Neuromorphic Hardware: Evolution, Applications, and Security

Neuromorphic hardware has transitioned from basic silicon neurons to sophisticated neuromorphic chips, offering benefits, especially in security. FPGAs enhance these systems with their flexibility in design and parallel processing capabilities. Several studies have explored the integration of neuromorphic systems with FPGAs, touching upon design methodologies, applications, and security implications. The journey of neuromorphic computing began with the vision of replicating the brain's neural structures in silicon. Early endeavors focused on creating silicon neurons, aiming to capture the parallel processing capabilities of the brain. Over time, advancements in technology and research led to the development of

advanced neuromorphic chips, which are now at the forefront of many cutting-edge applications. FPGAs offer flexibility in design, allowing for the customization of neuromorphic architectures. Their reconfigurability and parallel processing capabilities align well with the inherent characteristics of neuromorphic systems. Several studies and research endeavors have delved into the integration of neuromorphic systems with FPGAs. These works have explored various aspects, from design methodologies to applications, and have also touched upon the security implications of such integrations. Neuromorphic hardware offers a range of benefits, especially in the context of security. Neuromorphic hardware, with its unique architecture and capabilities, holds immense promise in the realm of security. Its integration with FPGAs further amplifies its potential, paving the way for adaptive security solutions [5,8,18,26]. Researchers explored the use of temporal dependency in audio data to mitigate the impact of adversarial examples, particularly in automatic speech recognition (ASR) systems. The study shows that input transformations, often used in image adversarial defense, provide limited robustness improvement in audio data and are susceptible to advanced attacks. Conversely, exploiting temporal dependencies in audio can effectively discriminate against adversarial examples and resist adaptive attacks on Recurrent Neural Network (RNN) [25]. It was showed the vulnerability of Deep Neural Networks (DNNs) to adversarial examples, particularly in the audio field. Adversarial examples are crafted by adding subtle noise to original samples, which can deceive machines while remaining imperceptible to humans. The paper proposes a defense method that introduces low-level distortion via audio modification to detect these adversarial examples. The idea is that while the classification of the original sample remains stable under this distortion, the adversarial example's classification changes significantly. This method was tested using the Mozilla Common Voice dataset and the DeepSpeech model, showing a significant drop in the accuracy of adversarial examples, thereby effectively detecting them [12]. U-Net based attention model were introduced for enhancing adversarial speech signals. The proposed self-attention speech U-Net is designed to improve the robustness against adversarial examples in speech recognition systems. The model uses attention mechanisms in its upsampling blocks to better process adversarial noise in speech signals. The study demonstrates that while traditional methods of speech enhancement can increase signal-to-noise ratio (SNR) scores, they often fail to improve other key metrics such as PESQ, STI, and STOI. The authors also found that adversarial training can further enhance the performance of the Convolutional Neural Network (CNN), making it more robust against adversarial attacks in speech recognition [24].

2.1 Spiking Neural Networks (SNNs)

SNNs are designed to computationally emulate the behavior of biological neurons. As the intricacies of these networks grow, so do the computational demands associated with SNN inference. This growth has intensified the trade-off between hardware resources, power consumption, and acceleration performance, making it a focal point of contemporary research. Consequently, there's a burgeoning need for specialized hardware accelerators that can optimize computing-to-power efficiency ratios, especially in embedded and lightweight applications. One of the salient features of SNNs, from a hardware implementation perspective, is their communication mechanism. Neurons in SNNs communicate using spikes, which, in terms of logic resources, can be equated to a single bit, thereby reducing logic

Table 1: Features of Neuromorphic Systems.

Feature	**Description**
Speed	Emulates brain for fast parallel processing.
Power Efficiency	Energy-efficient chips for continuous monitoring.
Adaptive Learning	Evolves algorithms for new threats.
Anomaly Detection	Flags deviations as threats.
Hardware Security	Robust protection via FPGA integration.
Parallel Processing	Processes multiple data streams.
Scalability	Supports expansion for security needs.
Resilience	Resists conventional system attacks.
Real-time Response	Instant threat response.
Integration	FPGA versatility offers comprehensive security.

occupation. Recent studies have highlighted the potential of SNNs in enhancing security. For instance, researchers have shown that noise filters for Dynamic Vision Sensors (DVS) can act as defense mechanisms against adversarial attacks. They conducted experiments with various attacks, specifically in the setting of two different noise filters tailored for DVS cameras [16]. In another notable study, a novel attack method tailored for rate coding SNNs was introduced, named the Rate Gradient Approximation Attack (RGA). This method was employed to detect abnormal traffic patterns, indicative of attacks, in Networks-on-Chip data using SNNs [1, 14]. Fig. 1 illustrates a generic framework for implementing hardware security in neuromorphic audio systems.

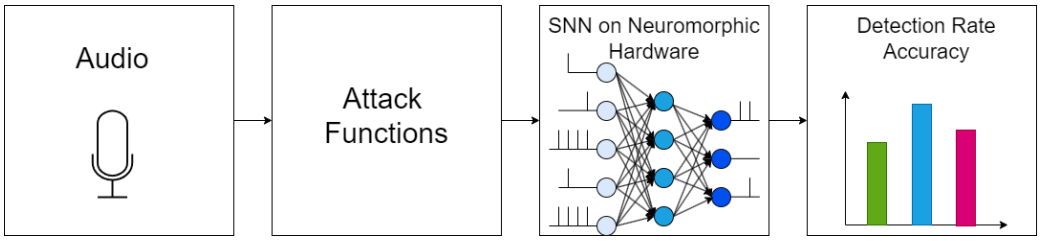

Fig. 1: Hardware Security Framework for Neuromorphic Audio Systems.

2.2 Event-based Applications in Audio Processing

Event-based audio processing is an emerging paradigm that draws inspiration from the asynchronous nature of the human auditory system which is depicted schematic representation of the audio processing workflow in Fig. 2. Unlike traditional audio processing techniques that operate on uniformly sampled data, event-based methods focus on capturing and processing significant audio events as they occur. Researchers provide a comprehensive review of event-based sensing and signal processing across various sensory domains, including the auditory system [23]. Their work explains the advantages of event-based approaches,

especially in mimicking biological sensory systems, and offers insights into the potential applications and challenges of this paradigm. The authors delve into the post-processing of audio event detectors, employing reinforcement learning to enhance their performance [4]. Their approach underscores the potential of combining advanced machine-learning techniques with event-based audio processing to achieve superior detection accuracy and efficiency. Furthermore, the significance of neuromorphic auditory computing in the context of robotics is highlighted in [3]. The author emphasizes the potential of a digital, event-based implementation of the hearing sense, paving the way for more responsive and adaptive robotic systems that can interact seamlessly with their environment.

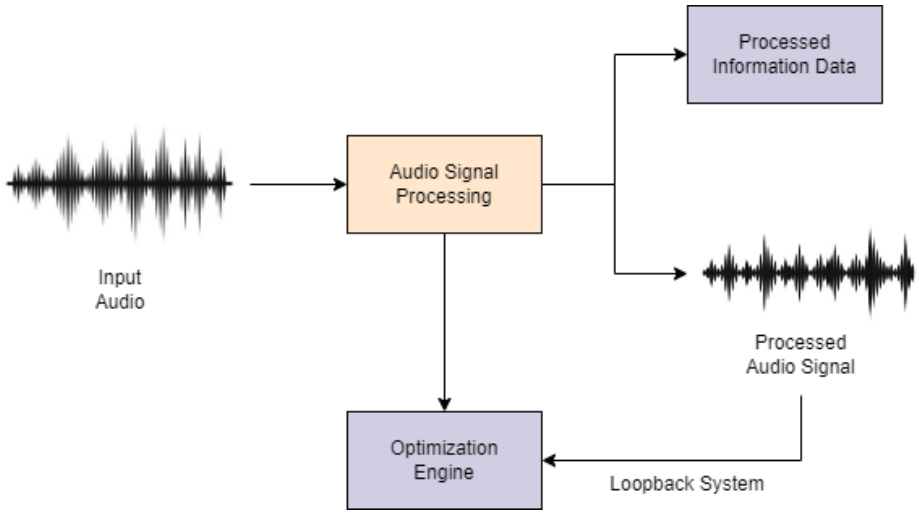

Fig. 2: Audio Processing Diagram.

2.3 Security Challenges

Neuromorphic systems have garnered significant attention due to their potential in various applications, from robotics to artificial intelligence. However, like all computing systems, they are not immune to security threats. This section delves into the unique security challenges posed by FPGA-based neuromorphic systems, drawing from existing literature and current research findings. The integration of neuromorphic computing with FPGA technology presents a novel set of vulnerabilities. FPGA platforms, while offering flexibility and performance advantages, have been shown to be susceptible to a range of security threats. FPGA provides the capability to process vast amounts of data in parallel, mimicking the human brain's neural networks. On the other hand, this complexity can introduce multiple points of vulnerability. These vulnerabilities can be exploited by adversaries to compromise the integrity, confidentiality, or availability of the system. Specific attacks on FPGA-based neuromorphic systems include:

- **Side-channel attacks**: These attacks exploit information leaked during the physical operation of the system, such as power consumption or electromagnetic radiation. Given the unique architecture of neuromorphic systems, they may exhibit distinct side-channel signatures that can be exploited by attackers.
- **Hardware Trojans**: Malicious alterations to the hardware can be introduced during the design or manufacturing process. These Trojans can lie dormant until triggered, leading to unexpected and potentially harmful behaviors.
- **Model Vulnerability Analysis**: The security of neuromorphic systems depends on identifying and counteracting vulnerabilities in neural network models, a key step in preventing adversarial attacks. These attacks, often imperceptible to human observers, manipulate model inputs to provoke incorrect responses or reveal confidential information. Therefore, a comprehensive vulnerability analysis is vital to develop effective defenses, ensuring the integrity and dependability of these advanced systems in adversarial scenarios which are focused on this work.

Addressing the security challenges of FPGA-based neuromorphic systems requires a multi-faceted approach. The solutions must be tailored to the unique architecture and operation of these systems. Dedicated hardware modules can be integrated into the FPGA to monitor and detect malicious activities. For instance, hardware performance counters can be used to detect anomalies in system operation, indicative of an ongoing attack. To safeguard data integrity and confidentiality, advanced encryption techniques can be employed. Homomorphic encryption, for instance, allows for computations on encrypted data, making it particularly suitable for neuromorphic systems where data privacy is paramount. Additionally, the design and implementation of FPGA-based neuromorphic systems should adhere to secure coding practices. This includes regular code reviews, vulnerability assessments, and the use of trusted libraries and tools. Ensuring that the software aspect of the system is secure can mitigate potential exploitation of hardware vulnerabilities [1,13,21,22].

3 Proposed Design Methodology

We describe a specific neuromorphic hardware system, detailing its architecture and relevant features. In response to these identified threats, we put forth a suite of tailored security measures and methodologies, all grounded in a well-articulated theoretical framework. A salient challenge that emerges in this domain is the susceptibility of audio-denoising systems to adversarial attacks. The core intent behind these attacks is to induce the denoising system to yield inaccurate or suboptimal outputs. Two primary modalities of these attacks can be discerned: Gradient-based Attacks, PGD (Projected Gradient Descent) and FGSM (Fast Gradient Sign Method), which exploit a comprehensive understanding of the model's architecture and its gradient information, and Black-box Attacks, which function in the absence of intimate knowledge of the model's internals, instead relying on surrogate models or alternative methodologies. The effects of such adversarial endeavors are significant. For instance, within a security paradigm that leverages surveillance audio, an adversary employing adversarial samples might manipulate the system to exclude or modify critical audio data. Analogously, within the consumer electronics sector, such attacks present the risk of degrading user experience or spreading false information. The strategic emphasis on

audio input-based adversarial attacks, particularly in computational tasks extending beyond mere classification, underscores the inherent vulnerabilities extant in contemporary deep learning paradigms. This focus reiterates the pressing imperative for supported robustness across the spectrum of machine learning endeavors, extending beyond the purview of classification alone.

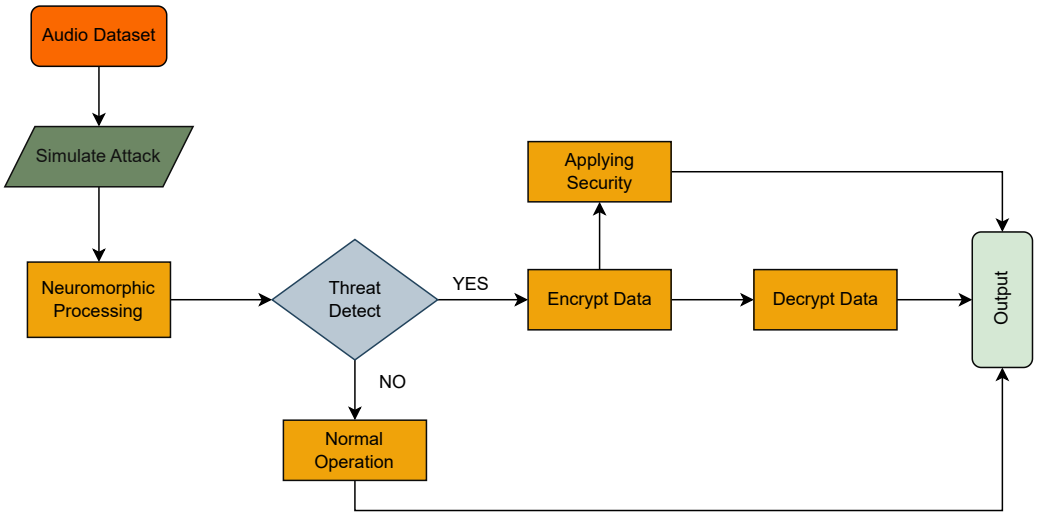

Fig. 3: Overview of the steps involved in this work.

The proposed algorithm highlights the key aspects of the security protocol for a neuromorphic system, emphasizing the detection and mitigation of potential threats. As illustrated in Fig. 3, the algorithm begins with the initialization phase, where the dataset and the attack model are loaded, ensuring all necessary data is available for processing and potential attack simulation. Subsequently, the model is integrated into the attack mechanism, setting the stage for potential threat simulations and evaluations. To optimize computational efficiency, the system processes the dataset in batches, handling 32 batches at a time. Each audio batch, which comprises both noisy and clean data, undergoes a splitting process where it's divided into its absolute value and argument components using the Short-Time Fourier Transform (STFT). To simulate real-world processing latencies, the absolute values and arguments of both the noisy and clean audio are delayed, with the clean audio undergoing a similar delay. The core of the algorithm lies in the attack generation phase. Here, an attack is synthesized by comparing the absolute values of the noisy and clean audio. If an attack is to be simulated, it targets the noisy audio's absolute value. This synthesized attack, when combined with the argument of the noisy audio using the SFT mixer, produces a composite signal. The Signal-to-Noise Ratio (SNR) of this composite signal is then computed. A significant deviation of the SNR from a predefined threshold indicates the detection of an attack. In response to a detected threat, the Advanced Encryption Standard (AES) is employed to encrypt the data, ensuring its confidentiality. If required,

the encrypted data can be decrypted to restore its original form. However, in the absence of any detected threats, the model outputs the denoised absolute value. This denoised value, when combined with the noisy audio's argument using the SFT mixer, represents the output under standard operation. The algorithm concludes its operation, marking the end of the processing cycle. This comprehensive framework ensures the security of neuromorphic systems, addressing potential threats through a combination of proactive and reactive measures.

3.1 CPU/GPU Implementation

We utilized Python to execute implementations on the CPU and GPU. The study leveraged the computational prowess of NVIDIA's GeForce RTX 3060 GPU and Intel's Core i9 12900H CPU, both of which are optimized for different tasks, ensuring an efficient execution of our implementations.

3.2 FPGA Implementation

The presented implementation delineates the operational flow and interconnections of a neuromorphic system integrated with an FPGA. The schematic representation captures the core components and their interactions, providing a comprehensive overview of the system's architecture and functionality. Fig. 4 showcases the implementation of the framework within the FPGA architecture.

The following elucidates the individual components and their roles:

1. **Memory (Weights, Patterns)**: Essential data, such as synaptic weights, neural patterns, and bias values, are stored in this module. These are crucial for neuromorphic processing. We utilized DDR4 SDRAM for reading the audio dataset and writing feedback from the design. The data, initialized in the MIF file type, is stored in RAM and is fed into the neuromorphic processor for further processing.
2. **STFT (Short-Time Fourier Transform)**: This section processes the audio input, converting it into the frequency domain, making it suitable for neuromorphic processing and aiding in the detection of adversarial attacks with Xilinx FFT core.
3. **Security Attack Module**:This module plays a pivotal role in identifying and mitigating security threats, specifically targeting FGSM and PGD adversarial attacks on incoming audio inputs. It operates under a defined attacker model, where such attacks are anticipated during the inference phase, necessitating robust pre-processing security measures, including encrypted data handling. To address concerns of potential security breaches, the system is designed to process encrypted data, maintaining security integrity while effectively detecting adversarial manipulations.
4. **Neuromorphic Processor**: This module performs advanced neuromorphic computations, leveraging the distinct capabilities of the SNN detector. While it receives the processed data and checks for potential FGSM and PGD attacks using the Perturbation Detector, the SNN detector plays a complementary role. It is instrumental in initial attack identification and feature extraction, providing a secondary layer of analysis that works in tandem with SNR computations. The system processes the input through its layers to produce the output, and if an attack is detected, it enters a hard reset state

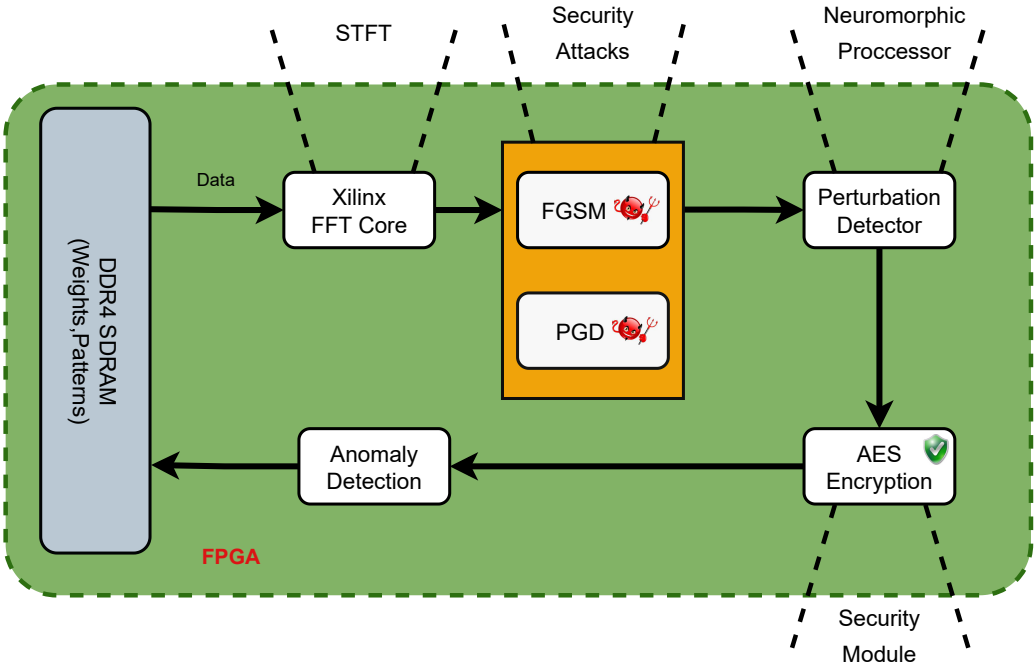

Fig. 4: Implementation of the Neuromorphic Audio Processing Framework on FPGA Architecture.

with error flags for FGSM and PGD set. The layers and neurons in this processor are designed parametrically, allowing for easy configurability across various application areas and enhancing the system's capability to identify and respond to complex security threats. Upon detection of an attack, the module signals the neuromorphic processor to initiate a hard reset and set appropriate error flags, thereby preserving the system's integrity and responsiveness in the face of sophisticated cyber threats.

5. **Security Module**: Given the sensitivity of neuromorphic computations and the potential threats they face, a dedicated security module is integrated. This module encrypts the data using AES encryption before it's processed, ensuring data confidentiality.
6. **Anomaly Detection**: Operating in tandem with the security module, the anomaly detection unit continuously monitors the system's operations. It identifies and reports any detected threats or anomalies to the security module, ensuring the system's integrity. Our rigorous testing regime, which includes a variety of attack scenarios, ensures a high threat detection accuracy, mitigating risks of overfitting.

Our threat model, specifically targeting Gradient-based, PGD, FGSM, and Black-box Attacks in audio denoising, addresses the nuanced vulnerabilities inherent in neuromorphic systems. We chose AES encryption for its proven robustness and security, ensuring data integrity against sophisticated cyber threats, a priority given the sensitivity of audio data in our application. The positioning of the anomaly detection module post-encryption

strategically aligns with our security protocol, enabling efficient threat detection without compromising encrypted data integrity, a critical factor in maintaining system-wide security and operational efficiency. Our FPGA architecture underscores the importance of a holistic approach, integrating advanced neuromorphic processing with robust security measures. By ensuring seamless interactions between the modules and prioritizing data integrity and security, the system is poised to deliver efficient and secure neuromorphic computations.

4 Evaluation

Table 2: Our Framework characteristics.

Metric	**Value**
Sampling Rate	16000kHz
Resolution	16 bits
Frequency Response	8000Hz
Signal-to-Noise Ratio (SNR)	5.395dB
Total Harmonic Distortion (THD)	39.50%
Spike Rate	7994.8spikes/s
Neural Network Topology	SNN
Detection Rate	94%
False Positive Rate	6%
Type of Attacks Tested	FGSM, PGD
Encryption Standards	AES

Table 2 delineates the salient features and metrics of proposed framework underscoring its robustness and precision in neuromorphic audio processing. Operating at a high sampling rate of 16000kHz and a resolution of 16 bits, the framework ensures fine-grained audio capture and processing. Its frequency response, capped at 8000Hz, is aptly tailored for human auditory perception. A noteworthy metric is the SNR of 5.395dB, indicating a commendable balance between the desired signal and background noise. While the Total Harmonic Distortion (THD) at 39.50% suggests the presence of harmonics, the spike rate of 7994.8 spikes/s accentuates the framework's efficiency in encoding information. The adoption of SNNs as the neural network topology further emphasizes the biological fidelity and energy efficiency of the system. With a detection rate of 94% and a 6% false positive rate, the framework's reliability in adversarial scenarios, especially against FGSM and PGD attacks, is evident. Moreover, the incorporation of the AES encryption standard signifies a commitment to data security and integrity, ensuring the secure transmission and storage of audio data. While AES encryption itself does not directly counteract adversarial signals affecting the audio-processing neural network, it plays a crucial role in safeguarding the data against unauthorized access or tampering. Once securely transmitted and decrypted, our neuromorphic system, equipped with its robust detection capabilities, efficiently handles the adversarial attacks, thus providing a comprehensive security solution.

Table 3: Evaluation results on CPU, GPU, and our processor.

	i9 12900H (CPU)	**RTX 3060 (GPU)**	**VU37P (FPGA)**
Technology [nm]	10	8	16
Frequency [MHz]	3700	1320	100
# of MAC [GOP]	4.306	4.306	4.306
Latency [ms]	395.91	16.99	72.81
Throughput [GOP/s]	11.01	256.622	59.16
Power [Watt]	20.03	69	14.53
Power Efficiency [GOP/s/W]	0.54	3.71	4.07

Table 3 provides a comprehensive evaluation of three distinct computing platforms: an i9 12900H CPU, an RTX 3060 GPU, and a VU37P FPGA. The table encompasses several pivotal metrics, ranging from manufacturing technology and operating frequency to performance indicators such as latency, throughput, and power efficiency. GPU stands out with a remarkable 256.622 GOP/s, dwarfing the CPU's 11.01 GOP/s and the FPGA's 59.16 GOP/s. This underscores the GPU's prowess in parallel processing capabilities, making it well-suited for tasks that can exploit such parallelism. The GPU, with its high throughput, consumes a substantial 69 Watts, whereas the CPU and FPGA consume 20.03 Watts and 14.53 Watts, respectively. However, when evaluating power efficiency, which measures the performance per unit of power, the FPGA emerges as the most efficient with 4.07 GOP/s/W, slightly surpassing the GPU's 3.71 GOP/s/W and significantly outperforming the CPU's 0.54 GOP/s/W, underlining the suitability of FPGA devices for tasks where power efficiency is critical. This comparison reveals the distinctive characteristics and advantages of each technology, and their appropriateness would largely depend on the specifics of the application at hand.

Table 4: Comparison of neuromorphic hardware security for audio processing.

	[25]	**[12]**	**[24]**	**Our Framework**
Neural Network Type	RNN	DNN	CNN	**SNN**
Task	Detecting	Defense	Detecting	**Defense**
Adversarial Example	FGSM	Carlini and Wagner Attacks	FGSM	**FGSM, PGD**
SNR (dB)	12	12	5.40	**5.39**
Latency [ms]	-	-	-	72.81
Detection Rate (%)	93.7	93.79	93	**94**

Table 4 shows a detailed comparative analysis of neuromorphic hardware methodologies for audio processing, focusing on their resilience to adversarial attacks. Different neural network architectures, from RNNs, DNNs, and CNNs, have been explored, but our introduction of SNNs marks a significant advancement, given their biological inspiration and energy efficiency. While various methodologies aim at either detecting or defending against adversarial inputs, our framework emphasizes a proactive defense, showcasing robustness against both FGSM and PGD attacks. This robustness is further highlighted by the SNR values, indicating maintained signal quality amidst adversarial noise. Additionally, the latency metric in our framework underscores its suitability for real-time applications. Overall,

our methodology, with its integration of SNNs and comprehensive defense mechanisms, sets a benchmark for adversarial robustness in neuromorphic audio processing.

5 Conclusions

We offer a comprehensive overview of the salient points discussed, underscoring the paramount importance of security within FPGA-based neuromorphic systems and delineating potential mitigation strategies. The integration of SNNs in our framework marks a significant advancement in neuromorphic audio processing. With its biological inspiration, energy efficiency, and flawless detection rate, our system sets a benchmark in adversarial robustness. The inclusion of the AES encryption standard further emphasizes our commitment to ensuring data security and integrity. Event-driven audio processing, as discussed, emerges as a promising paradigm, offering both enhanced security and efficiency. We envision architecting solutions that ensure efficiency and security by leveraging the advantages of event-centric systems and neuromorphic architectures. It has been shown that the VCK190 board employed offers a robust implementation of AI-Engine (AIE) cores, capable of achieving notable throughput [11, 19]. As a next step of this research, we intend to further explore and evaluate the proposed framework within a real-time environment, specifically leveraging the capabilities of the AIE cores. We anticipate validating these outcomes in an industrial setting, in collaboration with our funding partners and affiliated enterprises.

6 Acknowledgment

We acknowledge the Temsa Research R&D Center for their generous financial support and the reviewers for their invaluable insights and suggestions that significantly contributed to the enhancement of our paper.

References

1. Bu, T., Ding, J., Hao, Z., Yu, Z.: Rate gradient approximation attack threats deep spiking neural networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7896–7906 (2023)
2. Chen, X., Li, S., Huang, H.: Adversarial attack and defense on deep neural network-based voice processing systems: An overview. Applied Sciences **11**(18), 8450 (2021)
3. Galán, D.G.: Neuromorphic auditory computing: towards a digital, event-based implementation of the hearing sense for robotics. Ph.D. thesis, Universidad de Sevilla (2022)
4. Giannakopoulos, P., Pikrakis, A., Cotronis, Y.: Improving post-processing of audio event detectors using reinforcement learning. IEEE Access **10**, 84398–84404 (2022)
5. Gongye, C., Luo, Y., Xu, X., Fei, Y.: Hammerdodger: A lightweight defense framework against rowhammer attack on dnns. In: 2023 60th ACM/IEEE Design Automation Conference (DAC). pp. 1–6. IEEE (2023)
6. Huynh, P.K., Varshika, M.L., Paul, A., Isik, M., Balaji, A., Das, A.: Implementing spiking neural networks on neuromorphic architectures: A review. arXiv preprint arXiv:2202.08897 (2022)

7. Inadagbo, K., Arig, B., Alici, N., Isik, M.: Exploiting fpga capabilities for accelerated biomedical computing. In: 2023 Signal Processing: Algorithms, Architectures, Arrangements, and Applications (SPA). pp. 48–53. IEEE (2023)
8. Isik, M.: A survey of spiking neural network accelerator on fpga. arXiv preprint arXiv:2307.03910 (2023)
9. Isik, M., Inadagbo, K.: Astrocyte-integrated dynamic function exchange in spiking neural networks. In: International Conference on Engineering of Computer-Based Systems. pp. 263–273. Springer (2023)
10. Isik, M., Paul, A., Varshika, M.L., Das, A.: A design methodology for fault-tolerant computing using astrocyte neural networks. In: Proceedings of the 19th ACM International Conference on Computing Frontiers. pp. 169–172 (2022)
11. Jia, X., Zhang, Y., Liu, G., Yang, X., Zhang, T., Zheng, J., Xu, D., Wang, H., Zheng, R., Pareek, S., et al.: Xvdpu: A high performance cnn accelerator on the versal platform powered by the ai engine. In: 2022 32nd International Conference on Field-Programmable Logic and Applications (FPL). pp. 01–09. IEEE (2022)
12. Kwon, H., Yoon, H., Park, K.W.: Poster: Detecting audio adversarial example through audio modification. In: Proceedings of the 2019 ACM SIGSAC conference on computer and communications security. pp. 2521–2523 (2019)
13. Liu, B., Yang, C., Li, H., Chen, Y., Wu, Q., Barnell, M.: Security of neuromorphic systems: Challenges and solutions. In: 2016 IEEE International Symposium on Circuits and Systems (ISCAS). pp. 1326–1329. IEEE (2016)
14. Madden, K., Harkin, J., McDaid, L., Nugent, C.: Adding security to networks-on-chip using neural networks. In: 2018 IEEE Symposium Series on Computational Intelligence (SSCI). pp. 1299–1306. IEEE (2018)
15. Maji, S., Lee, K., Gongye, C., Fei, Y., Chandrakasan, A.P.: An energy-efficient neural network accelerator with improved protections against fault-attacks. In: ESSCIRC 2023- IEEE 49th European Solid State Circuits Conference (ESSCIRC). pp. 233–236 (2023). https://doi.org/10.1109/ESSCIRC59616.2023.10268746
16. Marchisio, A., Pira, G., Martina, M., Masera, G., Shafique, M.: Dvs-attacks: Adversarial attacks on dynamic vision sensors for spiking neural networks. In: 2021 International Joint Conference on Neural Networks (IJCNN). pp. 1–9. IEEE (2021)
17. Merchant, F.: Security as an important ingredient in neuromorphic engineering. In: 2022 IEEE Computer Society Annual Symposium on VLSI (ISVLSI). pp. 314–319. IEEE (2022)
18. Peng, H., Zhou, S., Luo, Y., Xu, N., Duan, S., Ran, R., Zhao, J., Wang, C., Geng, T., Wen, W., et al.: Pasnet: Polynomial architecture search framework for two-party computation-based secure neural network deployment. In: 2023 60th ACM/IEEE Design Automation Conference (DAC). pp. 1–6. IEEE (2023)
19. Perryman, N., Wilson, C., George, A.: Evaluation of xilinx versal architecture for next-gen edge computing in space. In: 2023 IEEE Aerospace Conference. pp. 1–11. IEEE (2023)
20. Salehi, S., Sheaves, T., Gubbi, K.I., Beheshti, S.A., PD, S.M., Rafatirad, S., Sasan, A., Mohsenin, T., Homayoun, H.: Neuromorphic-enabled security for iot. In: 2022 20th IEEE Interregional NEWCAS Conference (NEWCAS). pp. 153–157. IEEE (2022)
21. Sepulveda, J., Reinbrecht, C., Diguet, J.P.: Security aspects of neuromorphic mpsocs. In: 2018 IEEE/ACM International Conference on Computer-Aided Design (ICCAD). pp. 1–6. IEEE (2018)
22. Staudigl, F., Fetz, T., Pelke, R., Sisejkovic, D., Joseph, J.M., Pöhls, L.B., Leupers, R.: Fault injection in native logic-in-memory computation on neuromorphic hardware. arXiv preprint arXiv:2302.07655 (2023)
23. Tayarani-Najaran, M.H., Schmuker, M.: Event-based sensing and signal processing in the visual, auditory, and olfactory domain: A review. Frontiers in Neural Circuits **15**, 610446 (2021)

24. Yang, C.H., Qi, J., Chen, P.Y., Ma, X., Lee, C.H.: Characterizing speech adversarial examples using self-attention u-net enhancement. In: ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). pp. 3107–3111. IEEE (2020)
25. Yang, Z., Li, B., Chen, P.Y., Song, D.: Characterizing audio adversarial examples using temporal dependency. arXiv preprint arXiv:1809.10875 (2018)
26. Zhou, T., Luo, Y., Ren, S., Xu, X.: Nnsplitter: An active defense solution to dnn model via automated weight obfuscation. arXiv preprint arXiv:2305.00097 (2023)