Energy-Efficient ECG Classification via Event-Based Sampling and In-Memory Computing

Varun Raghuraman*, Ryan Brinson*, Rhimesh Lwagun* Xiaoyue Hu†, Weiling Li‡ Yan Fang*
*Department of Electrical and Computer Engineering, Kennesaw State University, Marietta, GA, USA
†Department of Computer Science, Kennesaw State University, Marietta, GA, USA
‡School of Computer Science and Technology, Dongguan University of Technology, Dongguan, China

***Abstract*—This paper presents a hybrid event-driven framework for low-power cardiac arrhythmia classification that integrates a Level-Crossing Analog-to-Spike Converter (LC-ASC) with a lightweight Event-driven Convolutional Neural Network (ECNN) classifier mapped onto a resistive compute-in-memory (CIM) system-on-chip (SoC). The proposed preprocessing converts sparse LC-ASC outputs into compact spatio-temporal spike-based frames, enabling efficient feature extraction and robust inference. Evaluated via 5-fold stratified cross-validation on the highly imbalanced MIT-BIH Arrhythmia Database, the proposed framework achieves a superior mean overall accuracy of 98.61% and a Macro-F1 score of 92.68%, significantly outperforming deep learning baselines particularly in detecting rare arrhythmias. The ultra-lightweight architecture requires only ≈857 kFLOPs per heartbeat inference. When evaluated on a 25mW 40nm RRAM-based CIM SoC, the system completes end-to-end inference within 5 µs using 100–150 µJ, demonstrating orders-of-magnitude improvements in latency and energy compared to standard desktop GPUs and low-power edge deployments. These results validate an effective hardware–software co-design for near-sensor edge intelligence in next-generation wearable health monitors.**

***Index Terms*—electrocardiogram (ECG) classification, event-based processing, System on Chip (SoC), neuromorphic sensing, Resistive RAM, Compute-in-Memory**

I. Introduction

Cardiovascular diseases remain a leading cause of global mortality, making the electrocardiogram (ECG) an indispensable tool for diagnosis and real-time monitoring. The expansion of wearable health devices has driven the demand for continuous, autonomous arrhythmia detection systems that can operate for extended periods on a limited power budget. A fundamental challenge for these battery-powered devices is the high energy consumption of conventional data acquisition systems. These systems typically rely on uniform Nyquist-rate sampling, which is inefficient for sparse biomedical signals like the ECG, as significant power is wasted sampling redundant data during inactive or "silent" periods. This issue, often termed the "Sampling Wall," is a primary barrier to creating long-endurance health monitors [1].

To overcome this limitation, event-driven sensing has emerged as a promising low-power alternative. A Level-Crossing Analog-to-Spike Converter (LC-ASC) was proposed to reduce data volume by generating sparse spike trains that encode signal activity [1]. However, while fully neuromorphic LC-ASC+SNN pipelines offer the highest theoretical efficiency, they introduce complexity in model training and hardware design [2]. Hybrid LC-ASC+ANN architectures present a more practical solution but often depend on simple feature extraction that underutilizes the temporal information in spike data [3]. Recent advances emphasize near-sensor intelligence for reduced power and real-time processing, aligning with the event-driven paradigm for wearable biomedical systems [4]. This paper proposes a hybrid, event-driven sensing and representation framework that bridges neuromorphic signal acquisition with efficient deep-learning inference on in-memory computing hardware. Sparse LC-ASC spike outputs are reframed into compact 2D spatio-temporal "spike images," allowing an Event-driven Convolutional Neural Network (ECNN) to achieve high classification accuracy with greatly reduced computation compared to standard frame-based approaches. The resulting model is mapped to a resistive compute-in-memory (CIM) SoC to achieve low-latency, energy-efficient operation suitable for edge deployment in wearable health devices.

The main contributions of this work are summarized as follows:

- **Efficient Event-based Encoding:** Utilization of the LC-ASC architecture to capture essential ECG dynamics in a sparse spike representation. Profiling on real patient data demonstrates this preprocessing step introduces negligible computational overhead (<3.5 kFLOPs per heartbeat), ensuring real-time feasibility.
- **Robust Lightweight Inference:** A compact ECNN that extracts spatio-temporal features of ECG events, achieving competitive Macro-F1 scores. Aided by inverse-frequency class weighting, the model surpassing standard Nyquist CNN baselines in identifying rare but critical arrhythmia classes (Supraventricular and Unclassified beats) while maintaining a minimal computational footprint.
- **Hardware/Software Co-design:** Adaptation of the proposed algorithm on an RRAM-based CIM processor and standard edge accelerators (e.g., NVIDIA Jetson Nano). The system achieves ultra-low latency validating its suitability for long-term wearable monitoring.
- **Real-Time Capability:** End-to-end system evaluation demonstrates real-time operation with approximately

This work was supported by NSF Grant CCF-2153440. Corresponding authors: Yan Fang (yfang@kennesaw.edu) and Weiling Li (liweiling@dgut.edu.cn).

0.81 ms preprocessing latency and 5 μs CIM inference latency per heartbeat, resulting in a total processing time below 1 ms. This provides over two orders of magnitude timing margin relative to the minimum physiological heartbeat interval (300 ms at 200 bpm), validating suitability for real-time wearable cardiac monitoring.

II. Related Work

Automated electrocardiogram (ECG) classification for arrhythmia detection has been implemented through various processing pipelines [2]. The conventional Nyquist+ANN pipeline achieves high accuracy but suffers from excessive power consumption due to dense, high-rate sampling. A Nyquist+SNN configuration improves computational efficiency but retains the same inefficient front-end. The fully event-driven LC-ASC+SNN architecture offers the highest theoretical power efficiency but introduces significant complexity in model training and hardware realization [2]. Similar event-based style information processing has also been successfully applied to computer vision with event camera and SNN, focusing on motion and temporal visual information [5]–[8].

This work focuses on the hybrid LC-ASC+ANN pipeline, previously explored in [3]. Although that study achieved high accuracy, its preprocessing compressed spike data into a 1D feature vector for a simple Multi-Layer Perceptron (MLP), thereby discarding valuable spatio-temporal information in the multi-channel spike signals. In contrast, our goal is to design a simpler yet more expressive framework that achieves comparable or superior accuracy to the LC-ASC+SNN model, which reached 95.34% in [2].

III. Method

This project seeks to design an event-driven pipeline that combines the low-power benefits of neuromorphic signal acquisition with the robust feature extraction capability of convolutional neural networks (CNNs). The proposed architecture leverages Level-Crossing Analog-to-Spike Converter (LC-ASC) circuits [1] to generate event-based spike signals from ECG data. Using this event-driven model, classification only occurs when meaningful spike information is available, allowing deep neural processing to remain quiet during periods of inactivity. The workflow is divided into three major stages: dataset preparation, spike-based preprocessing, and neural classifier development.

A. Dataset

The foundation for this work is the MIT-BIH Arrhythmia Database [9]. This clinically validated benchmark consists of 48 half-hour, two-channel ambulatory ECG recordings from 47 subjects, with sampling at 360 Hz. The database contains thoroughly annotated heartbeat events, amounting to over 100,000 labeled beats.

To ensure standardization and resolve inconsistencies between specific annotation labels and broad arrhythmia categories, we adopt the Association for the Advancement of Medical Instrumentation (AAMI) standard. We map the original MIT-BIH annotations (e.g., Left bundle branch block 'L', Right bundle branch block 'R', atrial premature 'A', etc.) into five distinct superclasses: Normal (N), Supraventricular (S), Ventricular (V), Fusion (F), and Unclassified (Q). The dataset exhibits a highly imbalanced class distribution, with the Normal (N) class comprising roughly 75% of all beats, while the S, V, F, and Q classes represent much smaller proportions. These classes are crucial for arrhythmia detection and classification. The same dataset and mapping scheme were used in prior works such as [1], [2], and [3].

B. LC-ASC and Preprocessing

The Level-Crossing Analog-to-Spike Converter (LC-ASC) serves as the event-driven front-end of the proposed architecture, enabling asynchronous sample generation that significantly reduces data volume compared to conventional uniform sampling. As illustrated in Fig. 1a, the LC-ASC quantizes the analog ECG waveform based on a dynamic threshold defined by:

$$LSB = \frac{A_{FS}}{2^M} \tag{1}$$

where A_{FS} is the signal amplitude full-scale range and M is the quantization bit resolution (set to $M = 6$ in this work) [1]. The system triggers a discrete event only when the signal amplitude changes by at least one LSB compared to the previous reference level. An increasing signal generates an 'UP' spike, while a decreasing signal generates a 'DN' spike. This mechanism compresses the raw time-series voltage data into a sparse, timestamped spike sequence.

Following the acquisition, the raw spike data is processed into a format suitable for training. The preprocessing pipeline first generates a "whitespace-free" event list (Fig. 2d), removing inactive temporal segments to maximize information density. For each annotated heartbeat, a fixed temporal window spanning 361 samples (capturing the 180 time steps immediately preceding and following the central beat annotation inclusive) is extracted and mapped onto a standardized 4 × 361 "spike image" frame (Fig. 2e). This frame separates the UP and DN spikes for both channels into distinct spatial rows, effectively converting temporal jitter into spatial features compatible with Convolutional Neural Networks [2]. It is important to note that these spike events serve as an event-driven signal encoding rather than neuronal firing events, and no spiking neuron dynamics are modeled in the subsequent neural network inference stage.

To validate the efficiency of this specific encoding step, we profiled the preprocessing algorithm on the full dataset. Experimental results demonstrate that converting a full heartbeat window (361 samples) into the sparse spike representation requires approximately 3.5 kFLOPs and executes in 0.81 ms on a standard CPU. Given that a typical physiological heartbeat spans approximately one second (1000 ms), this processing latency represents a speedup factor of over 1,200× compared to real-time acquisition, confirming that the LC-ASC encoding step introduces negligible delay.

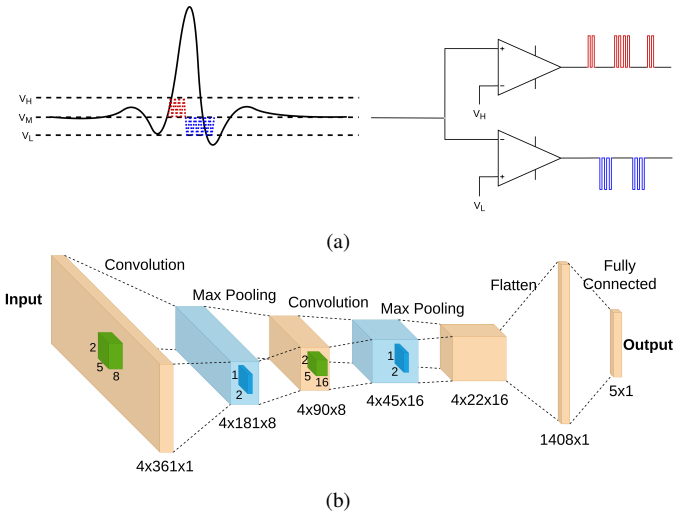

Fig. 1: (a) Event-based sampling that transforms ECG waveform into spike trains based on the increment and decrement of signal. (b)Architecture of ECNN that takes the ECG events performs classification .

C. ECNN Architecture

All standardized spike image frames are compiled for neural analysis, forming the training and evaluation sets for subsequent classification. Arrhythmia classification is handled by an event-based convolutional neural network (ECNN) (Fig. 1b) that takes as its input a $4 \times 361 \times 1$ matrix represented by 4 output signals from the LC-ASC and a frame that extends 180 elements before and after the current event inclusive. This input is then passed through the first convolution that has 8 filters of size of 2×5, a stride of 1 and 2 respectively, with padding set to maintain input dimensions using the formula $\lceil I/S \rceil$ [10], where I is the input dimension and S is the stride. After the first convolution, a max pooling with a frame of 1×2 and a stride of $[1\ 2]$ to output a $4 \times 45 \times 16$ matrix. This matrix is then convolved again with the same frame and stride size but an increased number of filters from 8 to 16. Again a max pooling is performed with the same hyperparameters as before, this time producing a $4 \times 22 \times 16$ matrix. The ECNN is then finished off by creating a fully connected (FC) that outputs the 5 possible classifications according AAMI standards.

D. Experimental Setup and Training Methodology

We implemented the ECNN models using PyTorch on a CUDA-enabled GPU. To ensure robust evaluation against class imbalance,We employed 5-fold stratified cross-validation (intra-patient beat shuffling) to ensure robust evaluation against class imbalance. Training utilized the Adam optimizer (initial learning rate 10^{-3}) and categorical cross-entropy loss for a maximum of 15 epochs with a batch size of 128 samples.

E. Compute-In-Memory Hardware

For hardware acceleration and energy efficiency, we map the proposed classification process on a 40nm 25mW CIM microprocessor composed of a 2.25MB RRAM-based accelerator with power gating and an embedded Cortex M3 core [11].

1) CIM on RRAM Macro: CIM processors contain crossbar arrays of RRAM devices that can be programmed into high/low resistance states (HRS/LRS). Vector-Matrix Multiplication (VMM) $\vec{y} = W\vec{s}$ can be carried out with such a configuration, as shown in the bottom of Fig. 3. The weight matrix W can be programmed on the resistive devices. The activation input vector s is applied to the wordlines (WL). The current flowing through a bitline provides the results of the multiply and accumulate (MAC) operation. The current is read by an analog-to-digital converter (ADC) for further processing in the digital domain. As binary states can be reliably programmed on RRAMs, a single weight element $W[i, j]$ is programmed across multiple columns. Simultaneously, the input ($\vec{s}[i]$) is applied in a binarized fashion. Therefore, for a quantization of W, shift and add operations are carried out in the digital domain.

2) Hybrid Digital/CIM RRAM SoC Architecture: We map the computation of each layer of the neural network on an RRAM CIM + digital processing enabled SoC by matching the required operation with high-speed and low-power matrix multiplication kernels. The structure of CIM SoC utilized for evaluation is shown in Fig. 3. This SoC integrates 288 RRAM modules (green) and each module storing 8kB of RRAM as 256×256 1T1R crossbar array. An embedded processor

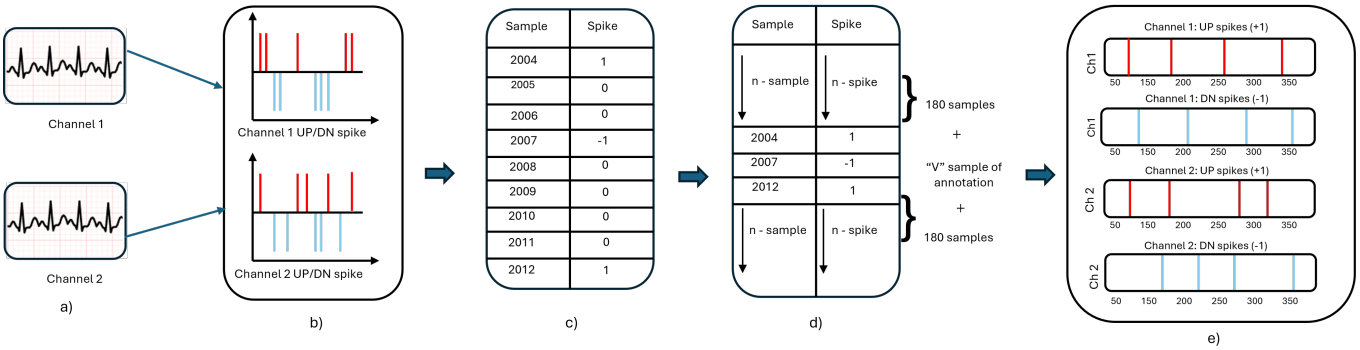

Fig. 2: Data preprocessing pipeline. (a) Two-channel analog ECG signals from the MIT-BIH database are (b) converted into four event-driven spike trains (UP/DN for each channel) using a virtual LC-ASC. (c) These spike trains are then represented as a compact 'whitespace-free' list of events. (d) For each annotated heartbeat, a fixed temporal window is extracted from the event list and (e) mapped onto a 4x361 2D 'spike image' frame, which forms a spatial-temporal representation input to the ECNN.

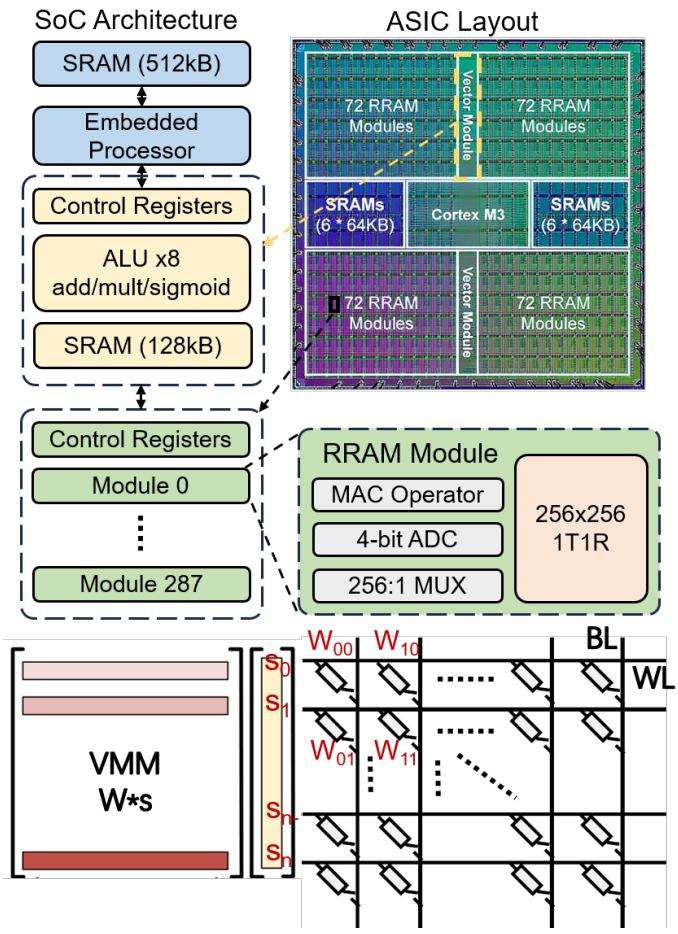

Fig. 3: General layout and die micrograph of the chip illustrated in SoC architecture with 288 RRAM CIM modules and detail structures.

(blue) controls an SRAM cache and the RRAM modules carrying out matrix-vector multiplications. Two vector modules (yellow) communicate with RRAM modules to transfer VMM results and perform supportive calculations, including activation functions. The control registers accessible by the processor are used to select and read/write the data. All modules can operate in parallel. As a single CIM module provides limited throughput, multiple independent and parallel operating modules are demanded for our computation task in this work. Fig. 3 also shows the die micrograph with different sections annotating the regions in the floor plan.

Different modules are programmed to implement different elements of the weight matrix. The large footprint and power consumption of ADCs require a small number of ADCs to be shared among the bitlines. Therefore, a multiplexer is added to determine the BL to be read. The output current is read using a 4-bit ADC in every module allowing reliable reading of 8 WLs in parallel. Each module also contains a dedicated shift and add unit. The reading happens sequentially. Shift and add operations are carried out in the digital domain.

3) Adapting ECNN on CIM processor : To efficiently deploy the proposed ECNN on the RRAM-based compute-in-memory (CIM) processor, the network architecture and dataflow are adapted to match the constraints and execution model of the crossbar arrays. Each convolution and fully connected layer is decomposed into a sequence of vector–matrix multiplication (VMM) operations, which are directly mapped onto the 256×256 RRAM crossbar arrays. Kernel weights are quantized to 8-bit fixed-point precision and stored as conductance values in the crossbar cells, while input activations are binarized to enable voltage-mode driving of the array rows.

During inference, convolutional operations are executed by unfolding the input feature maps into sliding window vectors, which are applied to the rows of the crossbar arrays. For each window, multiple output channels are computed in parallel by activating different columns within the same array. Max-pooling and feature map reshaping are implemented in the digital peripheral logic between successive CIM operations, avoiding additional analog computation overhead.

To ensure compatibility with binary input driving, the ReLU activation function is replaced by a step function during inference, while sigmoid activation is retained during training to preserve gradient stability. This modification enables efficient threshold-based activation without introducing spiking neuron dynamics or event-driven inference. Overall, the ECNN is executed as a sequence of deterministic and highly parallel CIM operations, allowing predictable latency and energy consumption while fully leveraging the inherent parallelism of the RRAM crossbar architecture.

IV. Results

A. Algorithm Evaluation

The classification performance of the proposed LC-ASC ECNN architectures was evaluated based on the 5-fold stratified cross-validation experimental setup described previously in III-D.

Table I summarizes the comparative performance metrics averaged across all five folds. To establish a benchmark, we compared our approach against a deep ResNet CNN trained on standard Nyquist-sampled ECG data [12]. While this baseline achieves a high overall accuracy of 93.4%, its Macro-F1 score is limited to 67.5% (calculated from the reported confusion matrix in [12]), reflecting poor generalizability to minority classes due to the severe dataset imbalance.

The proposed LC-ECNN models, operating on sparse, event-based spike frames, significantly outperform this deep baseline in both metrics. The ECNN variant utilizing the **ReLU** activation function achieved the highest performance, reaching a mean overall accuracy of 98.60% and a Macro-F1 score of 92.68%. The ECNN employing the **sigmoid** activation function also performed well, surpassing the baseline with an accuracy of 98.16% and a Macro-F1 score of 90.45%.

A critical challenge in arrhythmia diagnosis is accurately classifying minority classes alongside dominant ones. Figure 4 presents the aggregated confusion matrices across all five folds for both models. Both variants demonstrate exceptional performance in classifying the clinically significant Normal (N) and Ventricular ectopic (V) classes.

Most notably, as detailed in the per-class metrics (Fig. 4), the proposed framework shows marked improvement in detect-

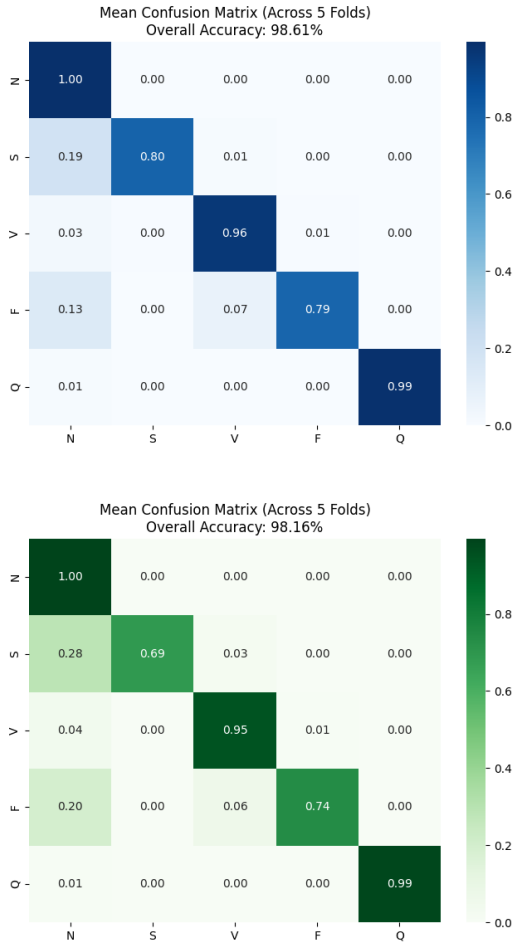

Fig. 4: Confusion matrix of the ECNN with two different activation functions, ReLU (left) and sigmoid/step-function (right)

ing severely underrepresented classes compared to standard approaches. The **ReLU** model successfully identified these minority classes, achieving a recall of 79.76% and an F1-score of 84.15% for the challenging S class, as well as an F1-score of 99.32% for the rare Q class. The **Sigmoid** model also demonstrated detection capabilities for these classes, achieving an F1-score of 78.39% for the S class and 98.94% for the Q class.

B. Hardware Evaluation

Our proposed lightweight ECNN architecture demonstrates high computational efficiency, requiring only approximately 857 kFLOPs for a single heartbeat inference operation.With Matlab and an NVIDIA RTX 4050 GPU running on a laptop in 80W, the computation latency of one ECNN inference is about $40\mu s$, resulting in a total energy consumption around 3.2mJ for the classification of one ECG event. To further evaluate performance in an embedded edge-computing scenario, inference was also conducted on an NVIDIA Jetson Orin Nano 8GB module operating at 5 W. In this case, the inference of one ECG event took $2.177, ms$, with an energy consumption of approximately $12.5mJ$.

TABLE I: Comparison of Classification Performance with Related Work

	Input Paradigm	**Overall Accuracy**	**Macro-F1**
ResNet CNN [12]	Nyquist Sampling (Raw Data)	93.4%	67.5%
Neuromorphic SNN [2]	Event-driven (LC Sampling)	97.8%	NR
Event-Driven ANN [3]	Event-driven (LC Sampling)	96.9%	NR
ECNN with ReLu (this paper)	**Event-driven (Spike Frames)**	**98.61%**	**92.68%**
ECNN with Sigmoid (this paper)	**Event-driven (Spike Frames)**	**98.16%**	**90.45%**

TABLE II: Comparison of Hardware Implementation Efficiency across Platforms

Hardware Platform	**Operating Voltage**	**Energy Consumption**	**Latency**
Nvidia RTX 4050 GPU	80W	3.2 mJ	40 μs
Nvidia Jetson Nano	5W	12.5 mJ	2.177ms
Proposed RRAM CIM Processor	25mW	100-150 μJ	5 μs

As summarized in Table II, Compared to the GPU, our evaluation of the hybrid CIM processor demonstrates superior efficiency in both computational and energy consumption. The end-to-end computation latency of one inference is around $5\mu s$ with the full power mode (25mW) of the processor. The energy consumption ranges from 100 to $150\mu J$, around 1/20 of the energy consumed by the NVIDIA RTX 4050 GPU and and about 1/80 of that consumed by the NVIDIA Jetson Orin Nano.

To better attribute the observed efficiency gains, the end-to-end system performance can be decomposed into three stages: event-driven sensing and encoding, ECNN inference, and memory movement. The LC-ASC front-end and spike-frame construction reduce the input data volume by eliminating redundant Nyquist samples, contributing to lower preprocessing and memory access overhead. The lightweight ECNN architecture further minimizes computational demand, requiring only 857 kFLOPs per heartbeat. Finally, mapping vector–matrix multiplications to the RRAM-based CIM arrays avoids repeated data transfers between memory and compute units, which dominate energy consumption in conventional von Neumann architectures. Together, these factors jointly enable the sub-millisecond latency and microjoule-level energy consumption observed in the proposed system.

These results validate the efficiency of the proposed algorithm model for deployment on edge hardware, particularly in in-memory compute architectures that accelerate VMM or MAC operations.

C. Enable Real-Time Processing

Although the spike-frame representation is described using fixed temporal windows for clarity, the proposed system operates in a fully streaming manner. The LC-ASC generates spike events asynchronously, which are processed incrementally as they arrive. A spike frame is constructed only when a heartbeat annotation is detected, triggering ECNN inference without requiring batch storage of long ECG segments. This event-driven, on-demand processing model enables continuous real-time monitoring while maintaining the sub-millisecond end-to-end latency reported in this section. A critical requirement

for wearable cardiac monitors is real-time operation, where the total system latency must be significantly less than the interval between heartbeats to prevent data backlog. Even for a severe condition like tachycardia at 200 bpm, the interval between beats is 300 ms.

Our proposed hybrid framework easily meets this constraint. As detailed in Section III-B, the software-based preprocessing step executes in approximately 0.81 ms. The subsequent hardware inference on the RRAM-based CIM SoC is completed within a negligible 5 μs. Therefore, the total processing latency per heartbeat is comfortably below 1 ms. This provides a safety margin of over two orders of magnitude compared to the tightest physiological timing constraints, guaranteeing that classification results are available virtually instantaneously for real-time clinical alert systems.

V. Discussion

A. Event-Driven Sparsity

In this work, sparsity is primarily exploited at the sensing and data-representation stages. The LC-ASC front-end suppresses redundant samples by emitting events only upon significant signal changes, while the subsequent whitespace-free event list removes inactive temporal regions prior to frame construction. As a result, the ECNN operates on a compact and information-dense spatio-temporal representation (4×361), which substantially reduces memory movement and total MAC operations compared to Nyquist-sampled pipelines. This design choice favors predictable, highly parallel execution that is well matched to compute-in-memory (CIM) architectures.

B. Robustness to Class Imbalance

A significant finding of this work is the model's ability to identify the 'S' (Supraventricular) and 'Q' (Unclassified) classes, which are often missed by standard classifiers due to their low prevalence in the MIT-BIH dataset. The combination of event-driven features—which highlight temporal changes rather than absolute amplitude—and inverse-frequency class weighting allows the ECNN to focus on the morphological distinctiveness of these rare beats. Unlike Nyquist-based CNNs that may overfit to the dominant 'Normal' rhythm, the sparse spike representation naturally filters out baseline wander and high-frequency noise, preserving only the essential structural dynamics of the arrhythmia.

VI. Conclusion

We proposed a hybrid event-driven ECG classification framework that integrates Level-Crossing Analog-to-Spike Converter (LC-ASC) sensing with a lightweight Convolutional Neural Network (ECNN) mapped onto an RRAM-based compute-in-memory (CIM) SoC. By transforming sparse LC-ASC outputs into compact spatio-temporal "spike-image" frames, the proposed approach effectively bridges neuromorphic signal acquisition and conventional deep-learning inference.

Crucially, the system achieves superior classification performance compared to Nyquist-based deep learning baselines, demonstrating exceptional robustness in detecting rare, clinically significant arrhythmia classes. Furthermore, the lightweight network architecture drastically minimizes computational complexity. Evaluation on the hybrid RRAM-based CIM SoC exhibits orders-of-magnitude reductions in computation latency and energy consumption compared to standard GPU implementations.

In summary, our work establishes a computationally and energy-efficient pathway for near-sensor intelligence in biomedical monitoring. By leveraging data-sparse event sampling, a lightweight neural model, and in-memory computing co-design, the LC-ASC + ECNN framework provides a practical blueprint for future low-power, real-time, and intelligent wearable healthcare systems. The complete pipeline achieves sub-millisecond end-to-end latency per heartbeat (0.81 ms preprocessing and 5 μs inference), providing over two orders of magnitude timing margin relative to physiological constraints.

References

[1] J. Chen, H. Wu, F. Tian, Q. Hou, S. Lin, J. Yang, and M. Sawan, "A low-power level-crossing analog-to-spike converter intended for neuromorphic biomedical applications," in *2024 IEEE International Symposium on Circuits and Systems (ISCAS)*, pp. 1–5, 2024.

[2] H. Chu, H. Jia, Y. Yan, Y. Jin, L. Qian, L. Gan, Y. Huan, L. Zheng, and Z. Zou, "A neuromorphic processing system for low-power wearable ecg classification," in *2021 IEEE Biomedical Circuits and Systems Conference (BioCAS)*, pp. 1–5, 2021.

[3] Y. Zhao, Z. Shang, and Y. Lian, "A 13.34μw event-driven patient-specific ann cardiac arrhythmia classifier for wearable ecg sensors," *IEEE Transactions on Biomedical Circuits and Systems*, vol. 14, no. 2, pp. 186–197, 2020.

[4] Y. Zhao and Y. Lian, "Event-driven circuits and systems: A promising low power technique for intelligent sensors in aiot era," *IEEE Transactions on Circuits and Systems II: Express Briefs*, vol. 69, no. 7, pp. 3122–3128, 2022.

[5] G. Gallego, T. Delbrück, G. Orchard, C. Bartolozzi, B. Taba, A. Censi, S. Leutenegger, A. J. Davison, J. Conradt, K. Daniilidis, *et al.*, "Event-based vision: A survey," *IEEE transactions on pattern analysis and machine intelligence*, vol. 44, no. 1, pp. 154–180, 2020.

[6] L. Yu, H. Chen, Z. Wang, S. Zhan, J. Shao, Q. Liu, and S. Xu, "Spikingvit: A multiscale spiking vision transformer model for event-based object detection," *IEEE Transactions on Cognitive and Developmental Systems*, vol. 17, no. 1, pp. 130–146, 2024.

[7] A. S. Lele, Y. Fang, A. Anwar, and A. Raychowdhury, "Bio-mimetic high-speed target localization with fused frame and event vision for edge application," *Frontiers in Neuroscience*, vol. 16, p. 1010302, 2022.

[8] A. Renner, L. Supic, A. Danielescu, G. Indiveri, E. P. Frady, F. T. Sommer, and Y. Sandamirskaya, "Visual odometry with neuromorphic resonator networks," *Nature Machine Intelligence*, vol. 6, no. 6, pp. 653–663, 2024.

[9] A. L. Goldberger, L. A. N. Amaral, L. Glass, J. M. Hausdorff, P. C. Ivanov, R. G. Mark, and H. E. ... Stanley, "PhysioBank, PhysioToolkit, and PhysioNet: Components of a new research resource for complex physiologic signals," 2000.

[10] Mathworks, "Convolution2dlayer - 2-d convolutional layer," Mar 2016.

[11] M. Chang, S. D. Spetalnick, B. Crafton, W.-S. Khwa, Y.-D. Chih, M.-F. Chang, and A. Raychowdhury, "A 40nm 60.64 tops/w ecc-capable compute-in-memory/digital 2.25 mb/768kb rram/sram system with embedded cortex m3 microprocessor for edge recommendation systems," in *2022 IEEE International Solid-State Circuits Conference (ISSCC)*, vol. 65, pp. 1–3, IEEE, 2022.

[12] M. Kachuee, S. Fazeli, and M. Sarrafzadeh, "Ecg heartbeat classification: A deep transferable representation learning framework," *IEEE Journal of Biomedical and Health Informatics*, vol. 22, no. 6, pp. 1721–1731, 2018.