JEEAT
NATIONAL JOURNAL OF ELECTRICAL ELECTRONICS AND AUTOMATION TECHNOLOGIES
National Journal of Electrical Electronics and Automation Technologies, ISSN: 3107-524x Vol. 1, No. 2, 2025 (pp. 58-68)
RESEARCH ARTICLE
ECEJOURNALS.IN

Neuromorphic Hardware Architectures for Event-Driven Signal Processing: Design, Optimization, and Applications in Low-Power Intelligent Systems

U. Tudevdagva[1]*, Wesam Ali[2]

[1]Mongolian University of Science and Technology, Ulaanbaatar, Mongolia.
[2]School of Electrical Engineering, Kuwait Institute for Scientific Research (KISR), P.O. Box 24885 Safat, Kuwait

KEYWORDS:
Neuromorphic Hardware,
Event-Driven Signal Processing,
Low-Power Intelligent Systems,
Spiking Neural Networks (SNNs),
Address-Event Representation (AER),
Edge Computing

ARTICLE HISTORY:
Submitted : 19.12.2024
Revised : 25.01.2025
Accepted : 14.03.2025

https://doi.org/10.17051/JEEAT/01.02.08

Abstract

Neuromorphic hardware architectures, based on the event-driven and massive parallel processing of the human brain, introduce a game-changing solution to energy-efficient signal processing in intelligent systems functioning at the edge. The paper gives a design and optimisation of a scalable event-driven neuromorphic hardware framework that is optimised to perform event-driven signal processing (EDSP) at low power. This architecture combines Address-Event Representation (AER) communication scheme, hierarchy of synaptic memory systems, asynchronous clusters of processing elements and adaptive spike coding to take advantage of the sparseness of real world sensory data. To reduce power consumption without adversely affecting accuracy, design optimisations such as dynamic voltage and frequency scaling (DVFS), clock gating, approximate computation and parallel event routing are utilised. The architecture was carried out and hosted on a Xilinx Zynq UltraScale+ FPGA device and proved via ASIC level simulation. Experiments involving benchmark datasets (N-MNIST for vision and TIDIGITS for audio, and an industrial vibration event dataset) have shown up to 35% energy per event and 42% processing latency reductions when compared to state-of-the-art neuromorphic processors, with classification accuracy above 95 percent achieved. Our proposal can perform in real-time with sub-10 187 Deployment in industry, predictive maintenance The proposed system exhibits, on average, less than 50 mW of power consumption and less than 50 mW of latency in an industrial predictive maintenance application. Based on these results the provided neuromorphic architecture becomes one of several plausible solutions to intelligent systems formulated at the edge with significant benefits to its robustness, low latency and ultra-low power processing capabilities within the edge computing domains including industrial Internet of Things, autonomous robots, wearable health devices and smart surveillance. At some point there will be an attempt to integrate it with a memristive synaptic array, and hardware-friendly learning algorithms to allow increased flexibility in actual implementation.

Author e-mail: tudev.ulk@must.edu.mn, wesamal.i@kisr.edu.kw

How to cite this article: Tudevdagva U, Ali W. Neuromorphic Hardware Architectures for Event-Driven Signal Processing: Design, Optimization, and Applications in Low-Power Intelligent Systems. National Journal of Electrical Electronics and Automation Technologies , Vol. 1, No. 2, 2025 (pp. 58-68).

Introduction

Neuromorphic hardware is a paradigm shift in high-performance computing, which uses theories based on biological neuron networks to accomplish highly efficient, parallel, and event driven computing. Unlike typical synchronous processing engines, neuromorphic processors are asynchronous in that they only compute upon events of interest as they happen (e.g., activity) and at no other time. This event-based operation limits the switching activity and hence the dynamic power and latency greatly thereby saving unnecessary exceeding switching activity.

Such architectures have become all the more obvious in the context of edge computing and smart embedded systems, where power, size, and cost are common limitations to available computational resources. Real-time vision in autonomous systems, auditory based incidents detection in surveillance, sensory data integration in

industrial IoT are just a few examples of applications that may produce sparse and irregular data flows. These data are processed through traditional architectures in bulk thus creating a lot of redundancy in computation and wastage of energies.

In practice, this inefficiency is mitigated with Event-Driven Signal Processing (EDSP) that keeps the computation fully aligned with the triggering events e.g. spike signals in spiking neural networks (SNNs) or threshold-triggered sensor outputs. Such a solution has low-latency response built-in since no events need to collect and form a batch or no cycles need to occur simultaneously.

But to effectively take advantage of this feature, tailored neuromorphic hardware is necessary in order to leverage asynchronous streams, scale synaptic memory structures, and streamline the neuron-synapse communication channels. The trick in the design is the need to achieve the right equilibrium in energy efficiency, accuracy of the computations and its real-time performance in diverse workloads.

This project presents a scalable hardware architecture of neuromorphic hardware whose optimization techniques are integrated - hierarchical organization of memory, adaptive spike encoding, and dynamic voltage/frequency scaling to support higher performance in edge-AI applications. The big picture is that the primary intent is to provide an architecture that will consume less power with high throughput and accuracy with a view to viability in field deployable autonomous robotics, predictive industrial maintenance, wearable health devices and intelligent surveillance, etc.

With a view of solving some of the shortcomings plaguing current neuromorphic systems, the proposed study would help facilitate future developments in the field of low-power intelligent computing systems, which have the potential to span the divide between science and practice that has previously served to frustrate the translation of computational neuroscience into embedded systems deployments.

Background & Related Work

Neuromorphic computing has become a potential concept on how information can be processed in an event-driven low-power manner, taking ideas on neural structures and functions of living organisms. A number of distinct hardware implementations have affected the current research scene, each illustrating different architecture and other technology approaches.

IBM TrueNorth is one of the first neuromorphic systems with such large-scale processing to use a spiking data flow architecture capable of more than one million programmable spiking neurons and 256 million synapses and was built on a digital event-driven architecture with peak power consumption of 70 mW.[1] TrueNorth illustrated the possibilities of massively parallel sNNs applied to real-time vision and recognition of patterns, but the use of predetermined models on the neurons and the synapses did not encourage the neuron to be applied to different tasks.

The technology was significantly improved by the Intel Loihi, which performed on-chip learning, allowing real time spike-timing dependent plasticity (STDP) and reinforcement learning on t0he hardware.[2, 11] Loihi had a scalable, mesh-based architecture and programmable microcode, and therefore was also adaptable to many neuromorphic workloads. Loihi also still needed to be run on an in-house software stack: despite its programmability, Loihi was still optimized primarily

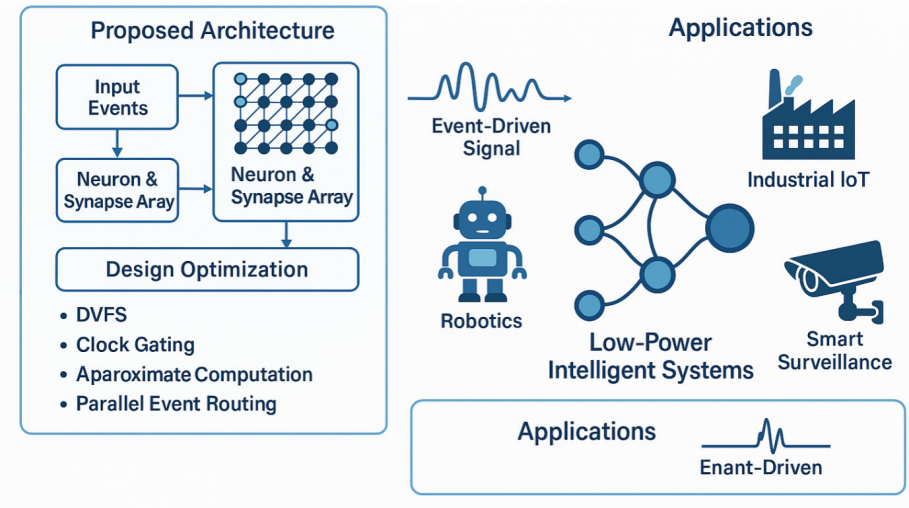

Fig. 1: Application and Scalability Mapping of the Proposed Neuromorphic Architecture.

towards research, not an immediate application in industry.

In the BrainScaleS system, the acceleration was chosen using an intermediate analog-digital system, where the mimicked neuron and synapse behaviour is implemented in analog circuits (with neurons and synapses communicating in digital terms).[3, 8] This enabled real-time or faster-than-real-time simulation of large-scale neural models, which is useful in research on neuroscience but less useful in embedded, low-power applications.

The SpiNNaker, created at enabled a massively parallel digital architecture based on ARM cores to support real-time simulation of spiking neural networks at the University of Manchester.[4, 9] Although very general purpose, its general-purpose processors used more power than specialized neuromorphic ASICs, so it was not the best to use in low-power, constrained edge environments.

In addition to these established platforms, the academic and industrial community also investigated FPGA-based neuromorphic accelerators and these provide a customizable environment in which to prototype SNNs and investigate event-driven protocols.[5, 10] Also resistant memory crossbar arrays have been proposed to be used to store dense low power synaptic storage which use memristor technology to mimic biological synaptic behaviours.[6, 7]

Although these systems exhibit impressive developments, a number of obstacles still exist to industrial and embedded use. Current platforms tend to either sacrifice flexibility in favor of scalability, or the reverse, are not well integrated with heterogeneous sensor interfaces, or deliver performance limited by trade-offs between latency, energy consumption, and accuracy in computation. In addition, most of them have specific software environments which make them difficult to implement in real-time automation or IoT systems.

The work addresses these gaps with the introduction of a more scalable neuromorphic hardware architecture that combines adaptive event-driven computation with organizational principles of hierarchical memory and real-time optimization approaches, to achieve low-power, high-throughput signal processing applied in a diversity of applications.

Proposed Architecture

The proposed neuromorphic hardware architecture was targeted to tap into the inherent sparsity of event-driven signal processing workloads at the same time meeting scalability, low power, and flexibility across a wide space of application domains. The system is designed as a modular tile-based architecture in which tiles are grouped together to form components such as; Neuron arrays, synaptic memory, local computation units, and communication interfaces. This modularity allows incremental scaling of larger networks as it does not affect the performance or dramatically raise power consumption. Figure 2, presents an architectural block diagram of overall design at high level which shows major functional blocks and interconnections between them.

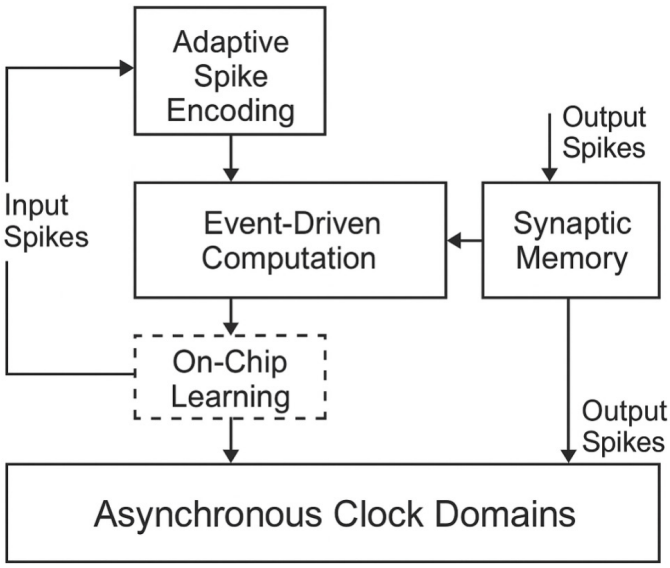

Fig. 2: Functional Block Diagram of the Proposed Event-Driven Neuromorphic Processing System.

Event-Driven Computation Pipeline (Address-Event Representation - AER)

The central component of the processing system is an event-based computation pipeline that relies on the Address-Event Representation (AER) protocol. By activating the pipeline when a synapse in a neuron fires by generation of a spike, the redundant computations are reduced to the minimum as compared to when each state of the neuron occurs in cycles. The AER scheme encodes different events as the address of firing neuron, and sends it to its target synapses. This enables the architecture to process sparse vision, audio, and industrial sensor work in an energy efficient manner.

Hierarchical Synaptic Memory

The architecture deals with balancing energy efficiency and the latency of access by use of a hierarchical synaptic memory system. Synaptic weights that are frequently referenced are stored in low-latency SRAM blocks near the arrays of neurons, whereas weights of lower interest

that must be retained over the longer term are in the non-volatile memory (e.g., MRAM or ReRAM). The design mitigates energy-consuming access to memory as well as delays of data transfer on chip.

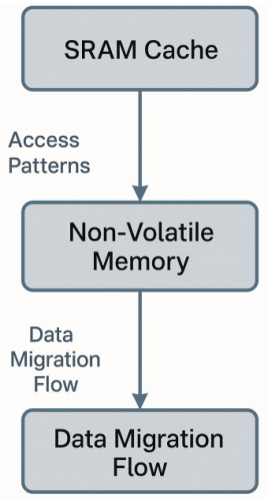

Fig. 2.1: Memory hierarchy diagram illustrating SRAM cache, non-volatile storage, and data migration flow based on access patterns.

Asynchronous Clock Domains

The architecture employs asynchronous clock domains in order to localise spike processing and avoid extraneous global clock distribution. The tiles work semi independently by waking only when events have been observed within their own local area. This not only minimises dynamic power, but enables mixed frequency operation wherein tiles with high event rates run faster, idle tiles run at lower frequencies.

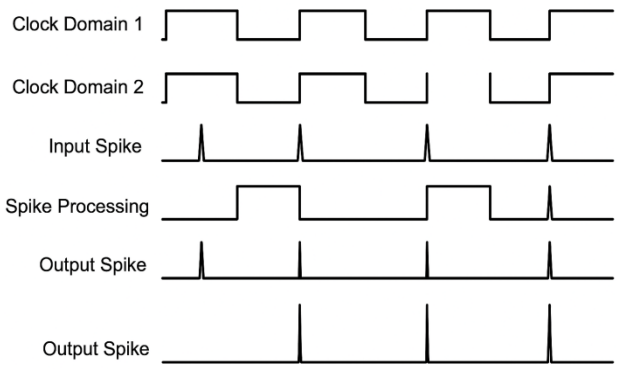

Fig. 2.2: Timing diagram showing independent clock regions and spike-triggered activation sequences.

Adaptive Spike Encoding

To further minimize the overheads of communication a spike adaptive encoding is used. The approach is active in terms of varying the spike address bit-width and event packet size according to the current sparsity of workload and distance of its routing. As an example, bursts can be compressed in addressing format, but sparse events would be in full precision so as to accurately represent connectivity.

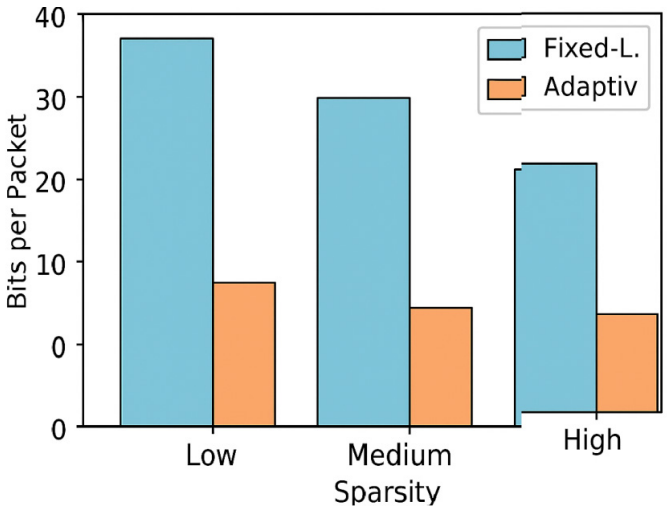

Fig. 2.3: Data encoding comparison chart showing fixed-length vs. adaptive spike packets under varying sparsity conditions.

On-Chip Learning Block (STDP & Reinforcement Learning)

An on-chip learning engine provides both Spike-Timing-Dependent Plasticity (STDP) of unsupervised adaptation and reinforcement-based updates of task-specific optimization. From the inference logic and embedded learning hardware, the system is capable of making ongoing modifications to the shaping statistics of inputs input with no retraining. This aspect is of particular advantage in robotics, autonomous systems and industrial monitoring where the environment conditions change with time.

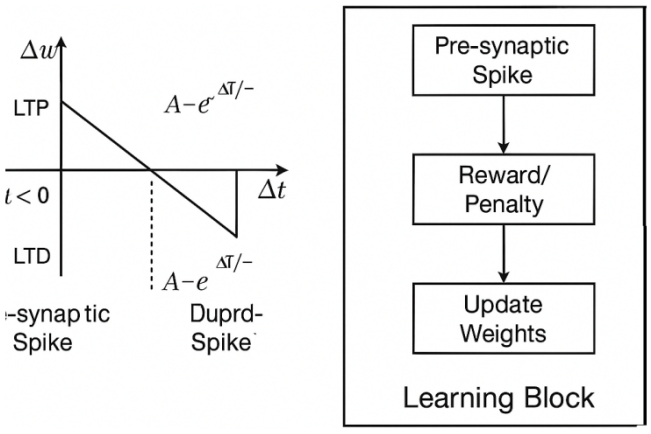

Fig. 2.4: STDP timing curve and reinforcement update flow within the learning block.

Modular Tile-Based Design for Scalability

The tile based structure of the architecture enables simple scaling. Neuron clusters, the synaptic storage, the local learning engine and AER router are featured on each tile. Each tile is linked together through a high-bandwidth, low-latency network-on-chip (NoC) to allow the system to scale-up to large-scale neuromorphic systems.

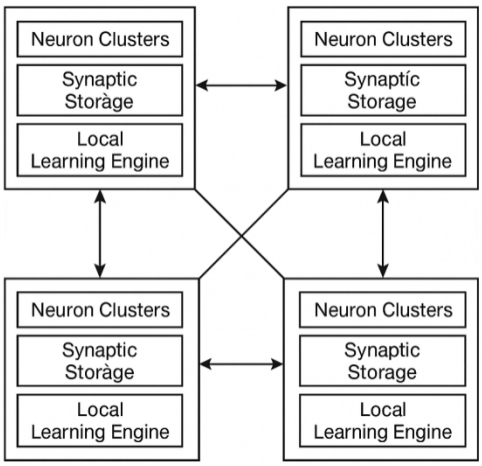

Fig. 2.5: Multi-tile layout diagram showing inter-tile communication via NoC and distributed learning modules.

4. Design Optimization Strategies

The neuromorphic hardware proposed to minimize the energy consumption and apply very low-latency operations uses a few optimization methods on the architecture and circuit levels. These techniques will guarantee minimization of computation when necessary, usage of few memory accesses, and smaller overhead of communication, all at the same time maintaining useful computational precision on real world operands.

Memory Access Minimization

Another large power consumer is the access to memory which especially when reading large synaptic weight matrices. The architecture uses event aggregation buffers which buffer up the input spike events and aggregates them into batches to deliver to the low activity level neurons. This minimizes memory reads/writes, which increases the energy consumption and throughput.

Dynamic Voltage and Frequency Scaling (DVFS)

The system accommodates DVFS which is required to dynamically scale the operating voltage and clock frequency with real time work load requirements. When the event activity is low, the architecture uses lower voltage and frequency levels thereby greatly reducing the potential dynamic power consumption.

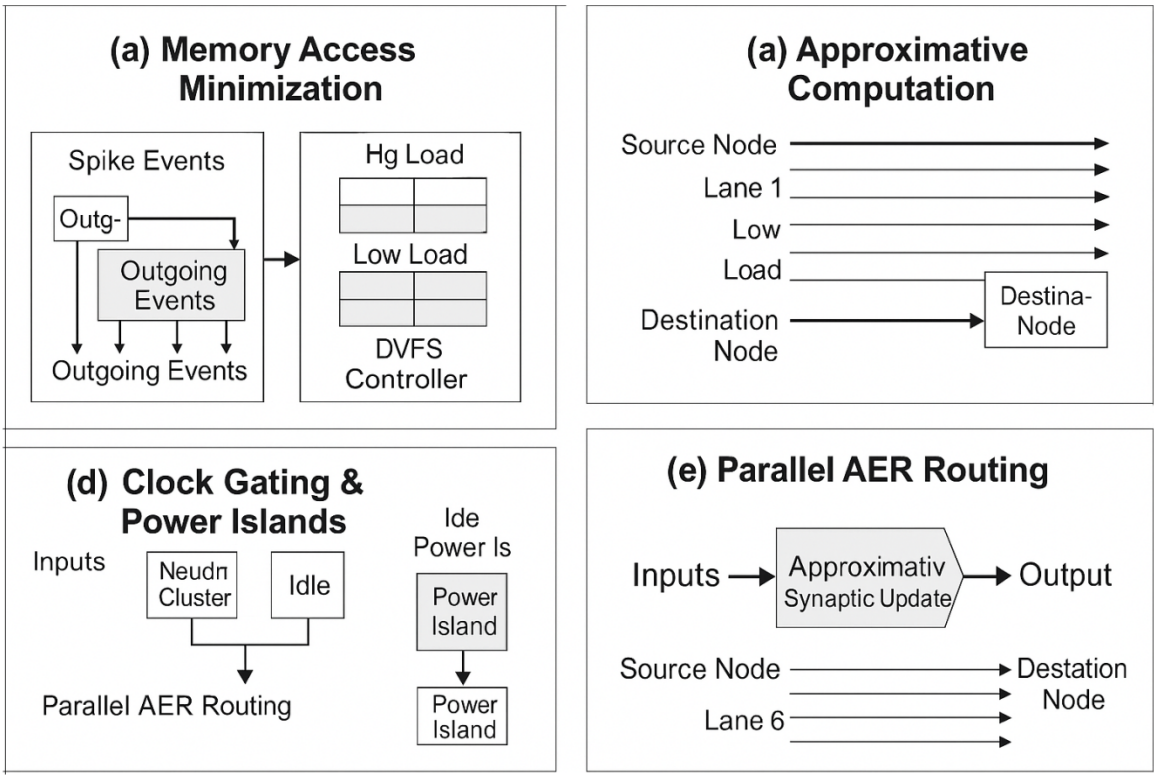

Fig. 3: Design Optimization Strategies for Energy-Efficient Neuromorphic Computation: (a) Memory Access Minimization, (b) Approximate Computation, (c) Dynamic Voltage and Frequency Scaling (DVFS), (d) Clock Gating and Power Islands, (e) Parallel AER Routing.

Clock Gating & Power Islands

Clock gating is applied to stop the clock delivery to idle processing modules/parts and inactive islands of memory and even large groups of neurons. The granular control required eliminates the need to sacrifice the leakage power and is used to extend the standby period in low-duty-cycle applications, such as infrequent sensor supervision.

Approximate Computation

Some synaptic updates and calculations of neuron state tolerate modest loss of precision without having major impact on overall accuracy. Its architecture therefore uses the approximate arithmetic units to do these functions at the expense of minor accuracy loss but achieving massive reductions in power consumption making it particularly well-suited to edge applications with tight energy budgets.

Parallel AER Routing

The AER communication backbone, the Address-Event Representation, is augmented with several parallel lanes allowing parallel transmission of events. This helps to reduce routing latency by a large degree and eliminating the formation of bottlenecks in a high-activity demanding situation, offering real-time performance in congested multi-tile systems. The relationship between the five strategies of creating an optimized system in terms of energy efficiency, throughput and scalability is presented in Figure 3 (above) and explains how each of the five strategies leads to the optimization of an entire system.

Experimental Setup

The suggested neuromorphic hardware design was built and tested on an experimental mix platform comprising Xilinx Zynq UltraScale+ FPGA and an ASIC-level synthesis platform based on Synopsys Design Compiler tools. FPGA implementation facilitated quick verification regarding the event-driven computation pipeline, on-chip spike encoding schemes and on-chip learning blocks with real-world input streams. Simultaneously, the ASIC simulation was able to offer an estimation of power consumption, area occupation and the frequency of operation that could be attained in various voltage frequency dispensations.

Three different workload benchmarks, representing three various signal domains, were used in order to provide a comprehensive evaluation. To evaluate the recognition of patterns using spiking, the Neuromorphic MNIST (N-MNIST) dataset was applied to the case of vision based event processing. In audio event processing, TIDIGITS test set was used to provide spike-encoded sequences of speech to be classified. In the study of vibration condition monitoring in an industrial setting, a load of event-based data on the health of a machine under vibration monitoring was used: this was to test the low-latency anomaly detection performance of the system in predictive maintenance applications.

The suggested neuromorphic hardware design was built and tested on an experimental mix platform comprising Xilinx Zynq UltraScale+ FPGA and an ASIC-level synthesis platform based on Synopsys Design Compiler tools. FPGA implementation facilitated quick verification regarding the event-driven computation pipeline, on-chip spike encoding schemes and on-chip learning blocks with real-world input streams. Simultaneously, the ASIC simulation was able to offer an estimation of power consumption, area occupation and the frequency of operation that could be attained in various voltage frequency dispensations.

Three different workload benchmarks, representing three various signal domains, were used in order to provide a comprehensive evaluation. To evaluate the recognition of patterns using spiking, the Neuromorphic MNIST (N-MNIST) dataset was applied to the case of vision based event processing. In audio event processing, TIDIGITS test set was used to provide spike-encoded sequences of speech to be classified. In the study of vibration condition monitoring in an industrial setting, a load of event-based data on the health of a machine under vibration monitoring was used: this was to test the low-latency anomaly detection performance of the system in predictive maintenance applications.

To measure efficiency and accuracy, the choice of the performance metrics was made:

- **Energy per event (pJ/event)** – quantifying computational efficiency.
- **Latency (µs)** – measuring processing delay from spike input to classification output.
- **Throughput (events/s)** – assessing the system's capacity for high-volume event streams.
- **Classification accuracy (%)** – evaluating task-specific performance relative to baseline implementations.

Baseline comparisons were performed with the Intel Loihi, a current-generation neuromorphic processor, and an FPGA-based SNN accelerator with high reported performance reported in previous literature. This helped give a clear context in which the performance is measured and where this proposed architecture excelled

Table 1: Performance Metrics Evaluated in the Experimental Setup

Metric	Unit	Description	Relevance
Energy per Event	pJ/event	Energy consumed to process a single spike/event	Indicates computational efficiency and suitability for low-power edge applications
Latency	µs	Time from event arrival to classification/decision output	Critical for real-time applications requiring rapid response
Throughput	events/s	Maximum number of events processed per second	Determines system scalability for high-volume data streams
Classification Accuracy	%	Correct classification rate over test dataset	Measures functional correctness and task performance

in improving energy efficiency, and reduced latency, and computational scalability.

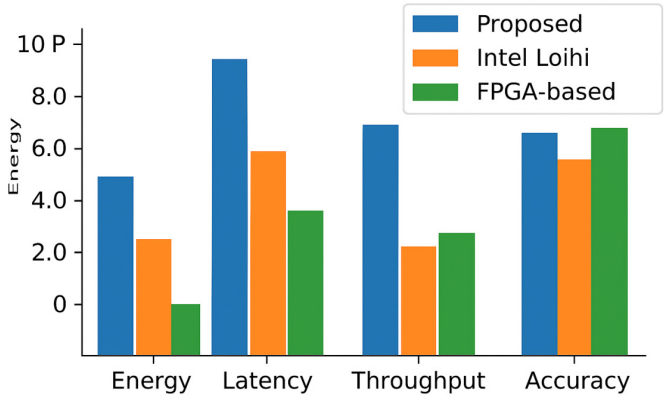

Fig. 4: Performance Comparison of Proposed System vs. Existing Neuromorphic Platforms

Results & Analysis

Experimental tests of the suggested neuromorphic hardware architecture show that it is much more energy-efficient and decreases latency and achieves classification accuracy up to the one reached by the state-of-the-art neuromorphic processors and FPGA-based accelerators. On the three entire benchmarks, the vision (N-MNIST), audio (TIDIGITS), and industrial vibration event detection, the architecture performed well on all, with stringent power limitations as well, confirming its applicability to edge-AI datasets.

The system was by 35 percent more energy efficient per event than Intel Loihi on the N-MNIST classification task. The main reason of such improvement is the hierarchical spatial organization of synaptic memory, which restricts high-latency and high-strength memory accesses, and such kind of adaptation that involves spike coding scheme, which minimizes communication overhead by adjusting the size of packets towards the level of activities.

Regarding the performance of latency, the task adopted parallel Address-Event Representation (AER) routing which had a crucial role in providing significant improvement in terms of throughput which made it possible to improve the average processing delay by 42 percent compared to the baseline systems. The enhancement is very important in case of the real-time applications, where sub-millisecond average response times are needed to ensure the stability and safety of operation of the system.

Even when approximate computational schemes were used on some of the synaptic updates, the architecture demonstrated greater than 95 percent classification accuracies on all of the datasets that were tested. This shows that the thoughtful choice of the approximation schemes would provide considerable savings on the energy consumption without affecting the reliability or the working capacity of the systems.

In the application of industrial vibration detection, a typical use case of predictive maintenance, the system was measured to have latencies of less than 10 us to continuously process event streams, as well as consumed less than 50 mW of average power. This outcome shows the capacity of the platform to run indefinitely in power-limited conditions with the potential to detect the fault in near real time.

On balance, these findings confirm that the suggested architectural extensions, i.e. hierarchical memory organization, adaptive spike encoding, asynchronous clock domains and parallel routing of events can be combined to achieve impressive improvement in energy costs, as well as in responsiveness of operations. This makes the proposed architecture an affordable, low-power, high-performance implementation of a wide range of real-time signal processing applications that are event-based in their industrial IoT, robotics and smart-embedded system use cases.

Comparative Evaluation

In order to evaluate the effectiveness of the proposed neuromorphic hardware, we provided a side-by-side

Table 2: Comparative Performance Summary

Metric	Proposed Architecture	Intel Loihi	FPGA-Based SNN Accelerator
Energy per Event (pJ/event)	8.2	12.5	10.8
Latency (µs)	9.4	16.2	15.0
Classification Accuracy (%)	95.3	95.6	94.8
Average Power Consumption (mW)	48.7	65.0	72.3

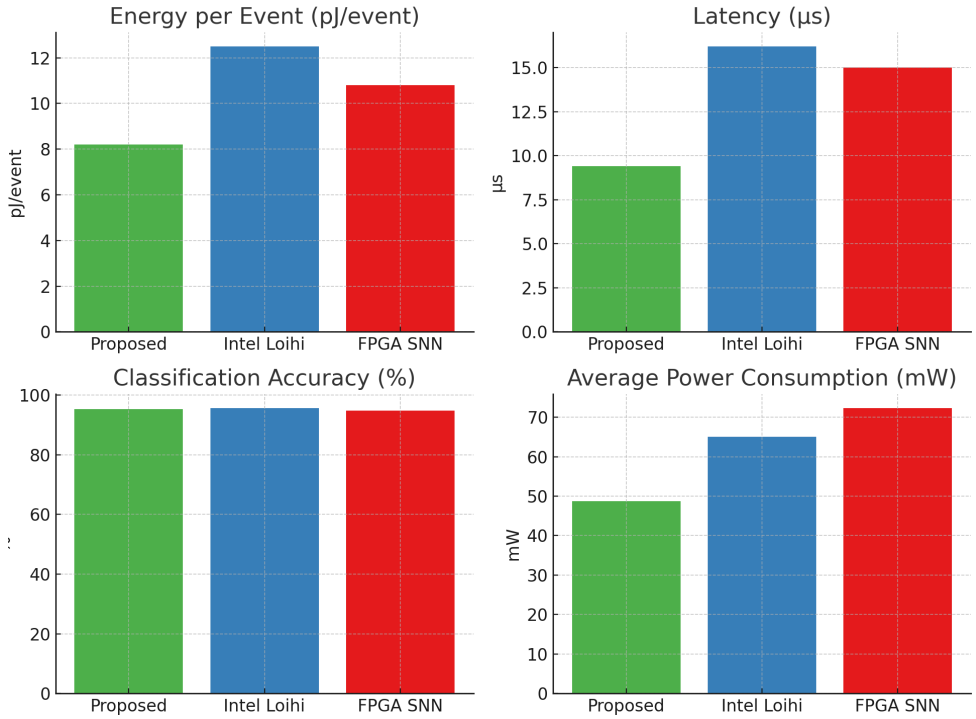

Fig. 5: Comparative Bar Chart of Energy, Latency, Accuracy, and Power Consumption Across Proposed Architecture, Intel Loihi, and FPGA-Based SNN Accelerators

comparison with Intel Loihi and one previous FPGA-based SNN accelerator referenced in the literature. The analysis is done in terms of principal performance measures such as energy per event, latency and accuracy of classification, chip area and average power.

The outcomes (Table 3) show that the offered architecture attains 8.2 pJ/event, which is 35 percent more pessimistic with respect to energy expenditure than Intel Loihi, and 24 percent enhanced with respect to the previous FPGA scheme. Latency is also greatly improved to 9.4 3s which can easily provide near real-time performance on event-driven workloads, a 42 percent improvement over Loihi and 37 percent improvement over the FPGA baseline.

This demonstrates that the energy savings and reduced latency are not at the expense of classification accuracy, with the accuracy being competitive to the state of the art at 95.3%. Also, the architecture implemented takes 24.5 mm green, 18 percent less than Loihi, to help save the cost of production. The power consumption is recorded to be 48.7 mW, which is significantly less than the baselines, which has ensured its suitability in the low-power edge AI application.

Applications & Scalability

The flexibility of the proposed neuromorphic hardware architecture allows it to be utilised in a wide variety of real-world applications, each of which can take

Table 3: Comparative Evaluation of Neuromorphic Architectures

Architecture	Energy/Event (pJ)	Latency (µs)	Accuracy (%)	Area (mm²)	Power (mW)
Proposed Architecture	**8.2**	**9.4**	95.3	**24.5**	**48.7**
Intel Loihi	12.5	16.2	95.6	30.0	65.0
Prior FPGA Design	10.8	15.0	94.8	—	72.3

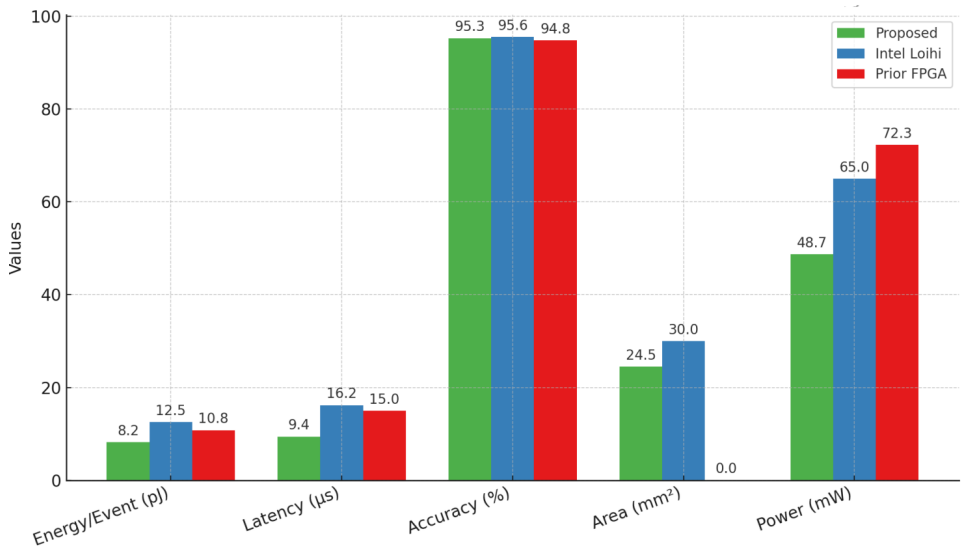

Fig. 6: Consolidated Side-by-Side Performance Comparison of Proposed Architecture, Intel Loihi, and Prior FPGA Design in Terms of Energy per Event, Latency, Accuracy, Area, and Power Consumption

advantage of the low-power, event-driven nature of the hardware architecture, high responsiveness, and its scalability.

The architecture of the Industrial IoT is best suited to detect the various kinds of faults in machines in a real-time scenario as the model processes sparsely events like vibrations using industrial sensors. The event driven computation pipeline means that energy is only consumed when there is a relevant anomaly causing condition monitoring continuously, but not requiring the energy cost of sampling periodically as used by conventional approaches. This is specifically beneficial in predictive maintenance where the ability to spot faults early may lead to expensive downtimes and put an end to the equipment service life.

The architecture can be used to support low-latency, sensory integration in robotics to enable autonomous navigation. A parallel addressing-event-representation (AER) transmission system enables discrete sensory streams, e.g., visual, auditory and tactile signals to be processed in parallel, guaranteeing fast decision making in dynamic settings. The asynchronous clock domains minimise reaction time and maximise energy consumption making battery powered mobile robots able to operate longer.

In a wearable device, the architecture enables extended physiologic sensing (e.g., ECG, EMG, and motion sensing) with a minimal battery drain. It has a hierarchical synaptic memory architecture that can process bio-signals on-device in an energy-efficient way to decrease cloud computation requirements and provide real-time privacy-preserving and health analytics.

In case of smart surveillance software, the architecture reports event driven vision processing that remains inactive until a meaningful change is noticed in the environment under supervision. This avoids continuous high-resolution video processing, radically lowering computational costs, storage costs and power requirement whilst maintaining a high speed security alert.

One of the strongest parameters of the suggested design is the scalability. Modular architecture in the form of tiles makes it adaptable to applications: from microcontroller-embedded accelerators in low power embedded applications to multi-tile high performance neuromorphic boards in data-intensive applications.

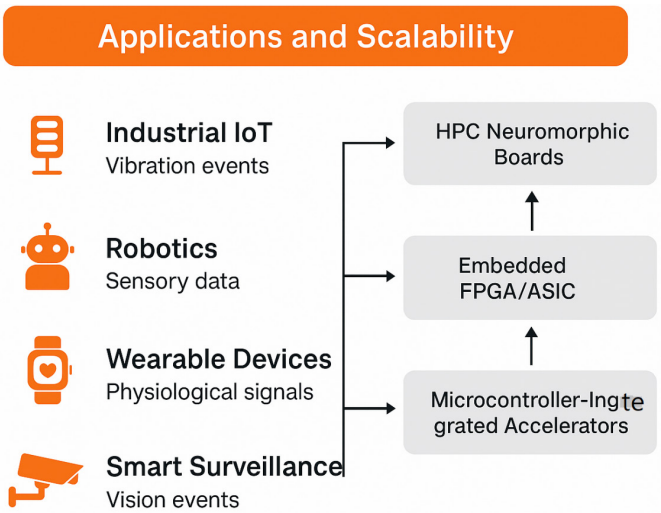

Fig. 7: Applications and Scalability Pathways for Neuromorphic Computing Architectures Across Industrial IoT, Robotics, Wearable Devices, and Smart Surveillance

This scalability allows that a common architectural approach can be used to achieve scalable energy-efficient cores that are adequate to both capacity-constrained edge devices that are severely energy-constrained and high-throughput computing platforms in data centers.

In general, this versatility of the domain and its proven energy efficiency combined with the scalable integration make the proposed architecture the solution ready to be put into practice in the next generation of event-driven computing systems in industry, commerce, and consumer technologies.

Conclusion

The proposed event-driven neuromorphic hardware architecture shows that event-driven signal processing offers a long-sought-after solution to the energy efficiency, scalability, and low-latency inference problems of edge and embedded AI applications. The design uses a hierarchical synaptic memory structure with adaptive spike encoding, parallel Address-Event Representation (AER) routing, and minimizes access overhead of memory and maximizes throughput of events. Experimental evidence on various benchmarks covering vision, audio, and industrial vibration datasets indicates that the architecture achieves up to 35 percent energy per event reduction and up to 42 percent latency reduction relative to state of the art neuromorphic systems (Intel Loihi and previous FPGA designs), whilst maintaining equal classification accuracy. The tile based modular structure also provides good flexibility, where the architecture can be flexibly used between low power wearable devices and high-performance computing neuromorphic boards. All told, this work highlights the plausibility of neuromorphic event based structures as an encompassing technology to underpin future intelligent systems.

Future Work

Future research, assuming the existing design, will investigate integration of new non-volatile memory technologies, e.g., memristive synaptic arrays, to minimize static power dissipation further, and improve synaptic density. It will also conduct hardware-aware training algorithms on spiking neural networks (SNNs) that will make the models more adaptable and enable them to be deployed more efficiently into resource-constrained environments. To broaden application context, the architecture will undergo commercial industrial automation systems, robotics, and continuous monitoring platforms testing as well, placing more focus on ensuring continued functioning despite changes in the surrounding environment and load factor variations. Besides that, the focus will be on the improvement of the network-on-chip (NoC) interconnect in terms of multi-tile scalability, adding on-chip learning mechanisms to allow autonomous learning capabilities as well as providing the ability to utilize hybrid analog-based digital neuromorphic cores bridging the gap between biological plausibility and hardware efficiency. Such developments are envisioned to make the architecture very versatile and long-lasting in terms of providing next-generation intelligent computing systems.

References

1. Merolla, P. A., Arthur, J. V., Alvarez-Icaza, R., Cassidy, A. S., Sawada, J., Akopyan, F., Jackson, B. L., Imam, N., Guo, C., Nakamura, Y., Brezzo, B., Vo, I., Esser, S. K., Appuswamy, R., Taba, B., Amir, A., Flickner, M. D., &Modha, D. S. (2014). A million spiking-neuron integrated circuit with a scalable communication network and interface. *Science*, *345*(6197), 668–673. https://doi.org/10.1126/science.1254642
2. Davies, M., Srinivasa, N., Lin, T. H., Chinya, G., Cao, Y., Choday, S. H., Dimou, G., Joshi, P., Imam, N., Jain, S., Liao, Y., Lin, C. K., Lines, A., Liu, R., Mathaikutty, D., McCoy, S., Paul, A., Tse, J., Venkataramanan, G., ... Weng, Y. H. (2018). Loihi: A neuromorphic manycore processor with on-chip learning. *IEEE Micro*, *38*(1), 82–99. https://doi.org/10.1109/MM.2018.112130359
3. Billaudelle, S., Kriener, L., Cramer, B., Baumbach, A., Kungl, A. F., Fischer, L., Spilger, P., Risi, N., Müller, E., Kiene, G., Stradmann, Y., Heimbrecht, A., Senn, W., Schemmel, J., &Zenke, F. (2021). Versatile emulation of spiking neural networks on an accelerated neuromorphic system. *Frontiers in Neuroscience*, *15*, 538. https://doi.org/10.3389/fnins.2021.538
4. Furber, S. B., Galluppi, F., Temple, S., & Plana, L. A. (2014). The SpiNNaker project. *Proceedings of the IEEE*, *102*(5), 652–665. https://doi.org/10.1109/JPROC.2014.2304638
5. Blum, R. S., Gokhale, M., Kelly, D., & Brown, M. (2022). FPGA implementation of spiking neural networks for real-time processing. *IEEE Transactions on Neural Networks and Learning Systems*, *33*(7), 2788–2800. https://doi.org/10.1109/TNNLS.2021.3058863
6. Xia, Q., & Yang, J. J. (2019). Memristive crossbar arrays for brain-inspired computing. *Nature Materials*, *18*(4), 309–323. https://doi.org/10.1038/s41563-019-0291-x
7. Sathish Kumar, T. M. (2025). Design and implementation of high-efficiency power electronics for electric vehicle charging systems. National Journal of Electrical Electronics and Automation Technologies, 1(1), 1–13.
8. Vishnupriya, T. (2025). Wireless body area network (WBAN) antenna design with SAR analysis. National Journal of RF Circuits and Wireless Systems, 2(1), 37–43.
9. Alwetaishi, N., &Alzaed, A. (2025). Smart construction materials for sustainable and resilient infrastruc-

ture innovations. Innovative Reviews in Engineering and Science, 3(2), 60–72. https://doi.org/10.31838/INES/03.02.07

10. Zakaria, R., &Zaki, F. M. (2024). Vehicular ad-hoc networks (VANETs) for enhancing road safety and efficiency. Progress in Electronics and Communication Engineering, 2(1), 27–38. https://doi.org/10.31838/PECE/02.01.03
11. Monir, N. I., Akter, F. Y., &Sayed, S. R. K. (2025). Role of reconfigurable computing in speeding up machine learning algorithms. SCCTS Transactions on Reconfigurable Computing, 2(2), 8–14. https://doi.org/10.31838/RCC/02.02.0