Accepted at ACM International Conference on Neuromorphic Systems (ICONS) 2025

Neuromorphic Cybersecurity with Semi-supervised Lifelong Learning

Md Zesun Ahmed Mia
School of EECS
The Pennsylvania State University
University Park, PA, USA
zesun.ahmed@psu.edu

Malyaban Bal
School of EECS
The Pennsylvania State University
University Park, PA, USA
mjb7906@psu.edu

Sen Lu
Dept. of EECS
University of Michigan
Ann Arbor, MI, USA
senlu@umich.edu

George M. Nishibuchi
Quantum Ventura, Inc.
San Jose, CA, USA
max@quantumventura.com

Suhas Chelian
University of Texas at Arlington Research Institute
Fort Worth, TX, USA
suhas.chelian@uta.edu

Srini Vasan
Quantum Ventura, Inc.
San Jose, CA, USA
srini@quantumventura.com

Abhronil Sengupta
School of EECS
The Pennsylvania State University
University Park, PA, USA
sengupta@psu.edu

***Abstract*—Inspired by the brain's hierarchical processing and energy efficiency, this paper presents a Spiking Neural Network (SNN) architecture for lifelong Network Intrusion Detection System (NIDS). The proposed system first employs an efficient static SNN to identify potential intrusions, which then activates an adaptive dynamic SNN responsible for classifying the specific attack type. Mimicking biological adaptation, the dynamic classifier utilizes Grow When Required (GWR)-inspired structural plasticity and a novel Adaptive Spike-Timing-Dependent Plasticity (Ad-STDP) learning rule. These bio-plausible mechanisms enable the network to learn new threats incrementally while preserving existing knowledge. Tested on the UNSW-NB15 benchmark in a continual learning setting, the architecture demonstrates robust adaptation, reduced catastrophic forgetting, and achieves 85.3% overall accuracy. Furthermore, simulations using the Intel Lava framework confirm high operational sparsity, highlighting the potential for low-power deployment on neuromorphic hardware.**

***Index Terms*—Spiking Neural Network (SNN), Lifelong learning, Hierarchical architecture**

I. Introduction

NETWORK Intrusion Detection Systems (NIDS) face significant challenges in scalability and energy efficiency, especially in high-throughput environments. The inherent nature of network attacks—often sparse events within a continuous temporal data stream—makes SNNs a compelling, bio-inspired alternative [1] to tackle such an application driver. SNNs offer potential for energy savings and real-time processing due to their event-driven, sparse computation. Furthermore, real-world cybersecurity scenarios often lack comprehensive labeled data, necessitating unsupervised or semi-supervised learning approaches. While Spike-Timing-Dependent Plasticity (STDP) is a common unsupervised learning rule in SNNs, its performance can be limited in complex classification tasks [2]. To enhance the efficacy and efficiency of STDP-based SNNs for NIDS, we explore a **hierarchical architecture** inspired from the brain's architectural organization [3]. This involves an initial, lightweight detection phase filtering benign traffic, followed by a more detailed classification phase, optimizing resource usage and potentially improving learning focus. Beyond static threats, NIDS must contend with the dynamic nature of cyberattacks, requiring adaptation to novel threats over time—a challenge known as **lifelong learning**. Conventional networks, and often static SNNs, suffer from *catastrophic forgetting*, losing old knowledge when learning new patterns [4]. Addressing this requires mechanisms for incremental learning.

To tackle both the STDP performance limitations and the lifelong learning challenge, this paper proposes and evaluates a **Hierarchical Dynamic Spiking Neural Network (D-SNN)**. This architecture combines the efficiency of the hierarchical structure with mechanisms for continuous adaptation. Learning employs our novel **Adaptive STDP (Ad-STDP)** rule, using a neuron's activity history ('firing factor') to modulate plasticity, stabilizing memories while learning new patterns [2]. We test this D-SNN on the UNSW-NB15 NIDS dataset, analyzing its adaptation, knowledge retention, efficiency, and neuromorphic hardware suitability via Lava simulations [5].

II. Related Works and Main Contributions

Neuromorphic approaches, particularly SNNs, are being explored for cybersecurity due to their potential for low-power, real-time processing. Prior work in neuromorphic NIDS includes systems demonstrating significant speed/energy improvements over conventional methods [6] and implementations on hardware like Intel's Loihi or using ANN-to-SNN conversion techniques [7]. Unsupervised methods using autoencoders on neuromorphic simulators or hardware have also shown promise [8]. However, many existing neuromorphic NIDS often rely on supervised training paradigms or do not explicitly tackle the challenge of continuously adapting to new, unseen threats without forgetting past ones. Life-

This material is based upon work supported in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, under Award Number #DE-SC0021562.

long learning, or continual learning, aims to address this adaptation challenge, but mitigating catastrophic forgetting remains difficult, especially in SNNs [4]. Strategies explored in SNNs include leveraging synaptic plasticity rules like STDP [2], developing adaptive plasticity mechanisms [9], controlling forgetting through neuromodulation or structural plasticity [10], and employing dynamic architectures inspired by concepts like GWR networks [11]. While these approaches show promise, integrating them effectively into a practical, efficient, and hierarchical NIDS framework remains an area ripe for investigation. Our work builds upon these foundations, combining a hierarchical SNN structure (justified in Sec. I) with dynamic adaptation and adaptive learning rules—specifically, our novel Ad-STDP incorporating a **neuron-specific firing factor to modulate plasticity for lifelong learning, differing from prior adaptive mechanisms**—for robust lifelong NIDS, explicitly tackling limitations of prior works, particularly their tendency to overlook hierarchical scalability or rigorous lifelong learning evaluation in this context. The main contributions of this paper are:

- **Hierarchical SNN with Dynamic Lifelong Learning Classifier:** Design and application of a novel, bio-plausible hierarchical SNN architecture for NIDS. This features an efficient static SNN detector (using standard STDP) followed by an adaptive dynamic SNN classifier. The dynamic classifier employs GWR-inspired structural plasticity combined with Adaptive STDP (Ad-STDP) learning to enable continuous learning of new attack types while mitigating catastrophic forgetting.
- **Semi-Supervised Lifelong Evaluation:** Demonstration of the semi-supervised D-SNN's effectiveness on the UNSW-NB15 benchmark [12] in a lifelong learning scenario, showcasing adaptation to new attacks while retaining prior knowledge compared to static counterparts.

III. Proposed Hierarchical D-SNN Methodology

Our proposed approach utilizes a **Hierarchical D-SNN** architecture specifically designed for adaptive and efficient Network Intrusion Detection. The motivation for a hierarchical structure stems from several observations relevant to NIDS. Firstly, cyberattacks are often sparse events compared to the high volume of benign network traffic. A single, complex classifier processing all traffic is inefficient. Secondly, real-world network data is inherently imbalanced. A hierarchical design allows for efficient filtering and helps mitigate the challenges posed by this imbalance. Inspired by biological processing pathways [3] and the need for efficiency, our architecture decomposes the intrusion detection task into two cascaded SNN modules (illustrated in the inset of Fig. 1):

Phase 1 (Attack Detection): A lightweight SNN module acts as an initial filter. It employs a **static architecture with a fixed size of 100 neurons**. It processes incoming network features (encoded as Poisson spike trains [2]) to make a coarse determination: is the traffic potentially malicious or benign?

Phase 2 (Attack Classification): This module is activated only when Phase 1 flags potential malicious activity. It utilizes a **dynamic structure** capable of adaptation. It receives the original input features concatenated with the activity state (average spiking rate) of the Phase 1 excitatory neurons. Its task is to classify the specific type of attack detected.

Both modules are SNNs built with Leaky Integrate-and-Fire (LIF) neurons, leveraging their event-driven nature for potential energy savings [1]. Lateral inhibition and homeostatic adaptive thresholds are used within each module's excitatory layer to promote neuron specialization and prevent dominance [2]. Beyond the hierarchical structure, the core innovation lies in the network's dynamic nature, enabling adaptation and lifelong learning. Inspired by GWR principles and the need to combat catastrophic forgetting, the Phase 2 D-SNN dynamically adjusts its structure and synaptic plasticity. The complete workflow, including the hierarchical structure, dynamic adaptation, and learning, is depicted in Fig. 1.

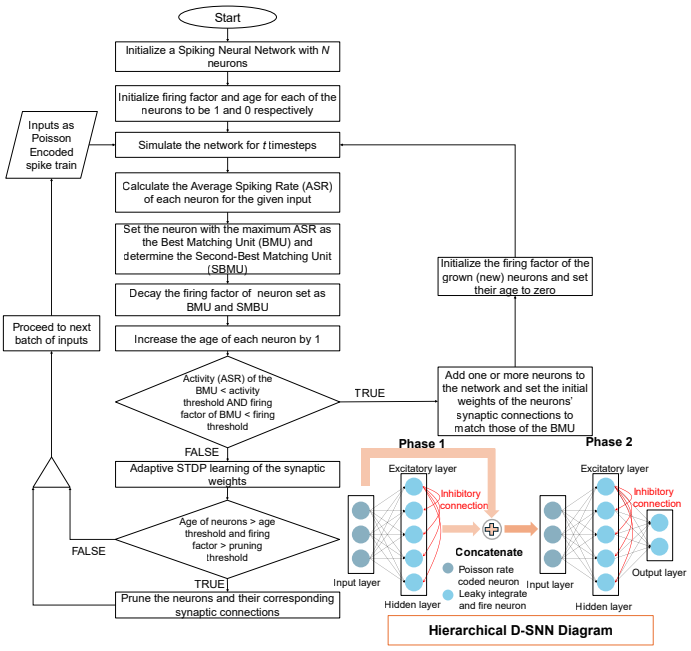

Fig. 1. Flowchart illustrating the core D-SNN algorithm, including dynamic structural plasticity (growth/pruning) and Ad-STDP learning, as applied in the Phase 2 module. The inset shows the two-phase hierarchical SNN architecture (Phase 1 is static, Phase 2 is dynamic).

A. Key Concepts for Dynamic Adaptation

Several key metrics and concepts govern the dynamic behavior of the D-SNN module and adaptive learning -

Average Spiking Rate (ASR): A measure of a neuron's recent firing activity over a sliding time window, indicating its response to current input stimuli.

Best Matching Unit (BMU) & Second-Best Matching Unit (SBMU): For a given input, the BMU is the excitatory neuron with the highest ASR, representing the closest match in the network. The SBMU is the neuron with the second-highest ASR [11].

Firing Factor (f_i): To balance plasticity and stability during lifelong learning, we introduce a neuron-specific firing factor. Intuitively, new neurons need to be highly adaptable to learn new patterns, while neurons that have already specialized

in representing certain inputs should become more stable to retain that knowledge. This concept is inspired by the biological principle of *habituation* (where neurons become less responsive to repeated stimuli) and GWR's habituation counter [11]. Our proposed firing factor (f_i) implements this idea by tracking a neuron's activity history and influencing its plasticity. It starts high (e.g., 1) for newly added neurons, promoting learning. As a neuron is frequently selected as the BMU or SBMU for inputs, indicating its successful integration and specialization, its firing factor decays over time according to $f_i = 1 - \frac{1}{\alpha_i}\left(1 - e^{-(\alpha_i \cdot n)/\tau_{ff}}\right)$, where n tracks the number of times the neuron was selected as BMU or SBMU, τ_{ff} is the decay time constant, and α_i is a neuron-specific rate constant. The SBMU's decay rate (α_{sbmu}) is further scaled by its ASR relative to the BMU ($\alpha_{SBMU} = \alpha_{BMU} \cdot \frac{ASR_{SBMU}}{ASR_{BMU}}$).

Neuron Age: A counter associated with each neuron, incremented over time (e.g., per mini-batch). It is used in conjunction with the firing factor for the pruning mechanism in Phase 2.

B. Dynamic Structural Plasticity (Phase 2)

The Phase 2 SNN module adapts its structure through neuron growth and pruning:

Network Growth: To accommodate new patterns without causing catastrophic forgetting, new excitatory neurons are added strategically. Growth is triggered when the BMU responds weakly (ASR $< a_{th}$, an activity threshold) to an input pattern that it *should* recognize, indicated by its low, decayed firing factor ($f_{BMU} < f_{th}$, a firing threshold). A low firing factor signifies that the BMU has already specialized in learning previous patterns; forcing it to learn this new, poorly matched pattern could overwrite its existing knowledge. Therefore, the dual condition (low ASR and low f_{BMU}) identifies the need for a *new* neuron to handle the novel pattern. This new neuron inherits weights similar to the BMU for rapid integration but starts with a high firing factor ($f_i = 1$) and zero age, maximizing its initial plasticity specifically for learning the new pattern.

Network Pruning: To maintain efficiency and remove redundant units, neurons are pruned based on their age and activity history. If a neuron's age exceeds a maximum threshold ($age > age_{max}$) *and* its firing factor remains consistently high ($f_i > p_{th}$, a pruning threshold), it suggests the neuron has failed to specialize or contribute meaningfully to pattern representation. Such neurons, along with their connections, are removed from the network.

C. Learning Rules: STDP and Ad-STDP

Synaptic weights are updated using different STDP-based rules in each phase. The static Phase 1 module employs standard STDP, where weight changes depend only on spike timing [2]. The dynamic Phase 2 module uses our novel Adaptive STDP (Ad-STDP) rule, a key contribution of this work. This rule introduces and incorporates the presynaptic neuron's firing factor (f_i) to modulate plasticity (Eq. 1), balancing stability and adaptability.

$$\Delta w_{ij} = \begin{cases} A_+ \cdot f_i \cdot \exp(-\Delta t/\tau_{pre}) & \text{if } \Delta t > 0 \text{ (LTP)} \\ A_- \cdot f_i \cdot \exp(+\Delta t/\tau_{post}) & \text{if } \Delta t < 0 \text{ (LTD)} \end{cases} \quad (1)$$

Here, $\Delta t = t_{post} - t_{pre}$ is the relative timing difference between postsynaptic (t_{post}) and presynaptic (t_{pre}) spikes. A_+ and A_- represent the maximum amplitudes for Long-Term Potentiation (LTP) and Long-Term Depression (LTD), respectively. τ_{pre} and τ_{post} are the time constants governing the STDP window for potentiation and depression. **Our work's novelty lies in modulating these updates with f_i:** High f_i of new or inactive neurons allows larger weight updates, facilitating rapid learning. As a neuron becomes established (f_i decays), the magnitude of weight updates decreases. This stabilizes learned representations and prevents new learning from drastically overwriting existing knowledge, effectively mitigating catastrophic forgetting. If f_i drops very low (approaching a habituated state), plasticity is significantly reduced for that neuron's outgoing synapses.

D. Semi-Supervised Labeling

Following the unsupervised phase involving structural plasticity and **STDP/Ad-STDP learning**, a small amount of labeled data is used to assign functional labels to the excitatory neurons [2]. In Phase 1, neurons are labeled as 'Attack' or 'Benign' based on their maximal ASR response to corresponding labeled inputs. In Phase 2, neurons are assigned specific attack type labels (e.g., 'DOS', 'DDOS', etc.) using the same principle. This semi-supervised approach leverages the network's self-organization while minimizing the requirement for extensively labeled datasets, making it suitable for real-world scenarios where labeled data may be scarce.

IV. Experimental Setup and Results

A. Experimental Setup

We evaluate our proposed **Hierarchical D-SNN** against a baseline Static Hierarchical SNN on the UNSW-NB15 NIDS dataset [12]. Our version focuses on lifelong learning across six distinct attack classes selected from the original nine; the remaining three classes (e.g., Worms, Shellcode, Analysis) are excluded due to having too few samples to form meaningful sequential tasks in our lifelong learning protocol. The data is preprocessed using standard cleaning, scaling, and random forest feature selection (42 features), with an 8:1:1 train/validation/test split. To assess adaptation, we simulate a task-incremental **lifelong learning** scenario: the network is trained sequentially on distinct tasks, each introducing benign traffic and a new, disjoint set of attack types (e.g., Task 1: DOS/Scanning; Task 2: Backdoor/DDOS, etc.), without revisiting prior task data. This protocol mimics real-world adaptation needs without full retraining. We use Python-based simulation framework with BindsNET [13]. Efficiency analysis use Intel Lava framework [5] simulations.

B. Results

Compared to the static baseline which suffers significant performance degradation, the proposed Hierarchical D-SNN effectively adapts to new tasks and mitigates catastrophic

forgetting, demonstrating the benefits of its dynamic structure and adaptive learning for lifelong operation. The mechanism enabling this improved adaptation is visualized in Fig. 2, tracking the structural evolution during the lifelong learning process. The network begins with few neurons and dynamically increases its size (growing to approx. 90 neurons) as it encounters new information corresponding to different attack classes. This contrasts sharply with the static baseline's fixed capacity.

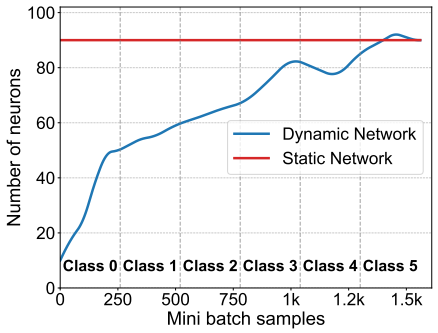

Fig. 2. Neuron count evolution in the D-SNN during lifelong learning.

The performance benefit of dynamic adaptation is reflected in Fig. 3. Critically for NIDS, the dynamic network shows substantially higher recall for most attack classes (0, 1, 4, 5), indicating superior pattern learning and knowledge retention essential for lifelong learning. While precision varies and static recall is higher for Class 3, the overall recall trend supports the dynamic approach's adaptability. While recent static, supervised deep learning models report high multi-class classification accuracies on UNSW-NB15 (often exceeding 95% [14]), our **Hierarchical D-SNN addresses the distinct challenges of lifelong learning using a semi-supervised SNN approach and is the first to report performance of a neuromorphic algorithm in this domain.** Based on the Phase 1 detection accuracy (94.3%) and Phase 2 classification accuracy (66.3%), weighted by the proportion of benign (72.5%) and attack (28.5%) traffic, the estimated overall system accuracy is approximately 85.3%. This significantly outperforms the static SNN baseline, whose overall accuracy under the same conditions is estimated at 80.0% (using a Phase 2 static accuracy of 46.6%).

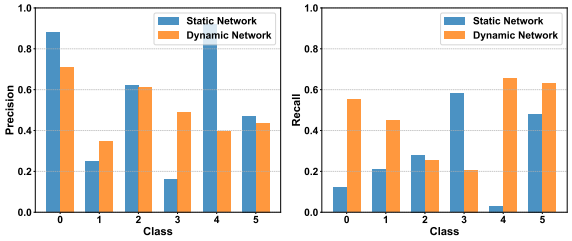

Fig. 3. Precision and Recall comparison per NIDS attack class for static vs. dynamic SNNs.

Efficiency analysis using the Intel Lava framework highlights the benefits of our approach. The Hierarchical D-SNN operates with high sparsity, indicated by very low average inference spike rates per neuron (~ 0.0008 for the Phase 1 filter, $\sim 0.001 - 0.002$ for the dynamic classifier over 200 timesteps). This inherent sparsity, particularly in the initial detection stage, is significantly greater than typical ANN-SNN conversion techniques [1] and directly contributes to potential energy savings, as computation is primarily event-driven.

V. Conclusions

Our Hierarchical D-SNN integrates structural plasticity with adaptive Ad-STDP learning, enabling lifelong NIDS capabilities that mitigate catastrophic forgetting and improve pattern learning over static SNNs. Its advantages include hierarchical efficiency, inherent sparsity suitable for neuromorphic hardware, and reduced label dependency through semi-supervised learning. As the first hierarchical D-SNN combining these techniques for semi-supervised continual learning in NIDS, this architecture offers a promising direction for robust, energy-efficient cybersecurity, despite overheads associated with dynamic network growth/pruning. Future work will focus on on-chip learning, Ad-STDP refinement, and evaluation on more complex datasets.

References

[1] A. Sengupta, Y. Ye, R. Wang, C. Liu, and K. Roy, "Going deeper in spiking neural networks: VGG and residual architectures," *Frontiers in neuroscience*, vol. 13, p. 95, 2019.

[2] P. U. Diehl and M. Cook, "Unsupervised learning of digit recognition using spike-timing-dependent plasticity," *Frontiers in computational neuroscience*, vol. 9, p. 99, 2015.

[3] J. LeDoux, "Rethinking the emotional brain," *Neuron*, vol. 73, no. 4, pp. 653–676, 2012.

[4] R. Kemker, M. McClure, A. Abitino, T. Hayes, and C. Kanan, "Measuring catastrophic forgetting in neural networks," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 32, 2018.

[5] M. Davies, A. Wild, G. Orchard, Y. Sandamirskaya, G. A. F. Guerra, P. Joshi, P. Plank, and S. R. Risbud, "Advancing neuromorphic computing with loihi: A survey of results and outlook," *Proceedings of the IEEE*, vol. 109, no. 5, pp. 911–934, 2021.

[6] D. R. Follett, D. Townsend, P. L. Follett, G. D. Karpman, J. H. Naegle, R. A. Suppona, J. B. Aimone, and C. D. James, "Neuromorphic data microscope," in *Proceedings of the Neuromorphic Computing Symposium*, 2017, pp. 1–5.

[7] W. Zahm, T. Stern, M. Bal, A. Sengupta, A. Jose, S. Chelian, and S. Vasan, "Cyber-neuro RT: Real-time Neuromorphic Cybersecurity," *Procedia Computer Science*, vol. 213, pp. 536–545, 2022.

[8] M. Z. Alom and T. M. Taha, "Network intrusion detection for cyber security on neuromorphic computing system," in *2017 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2017, pp. 3830–3837.

[9] P. Panda, J. M. Allred, S. Ramanathan, and K. Roy, "Asp: Learning to forget with adaptive synaptic plasticity in spiking neural networks," *IEEE Journal on Emerging and Selected Topics in Circuits and Systems*, vol. 8, no. 1, pp. 51–64, 2017.

[10] J. M. Allred and K. Roy, "Controlled forgetting: Targeted stimulation and dopaminergic plasticity modulation for unsupervised lifelong learning in spiking neural networks," *Frontiers in neuroscience*, vol. 14, p. 492718, 2020.

[11] S. Marsland, J. Shapiro, and U. Nehmzow, "A self-organising network that grows when required," *Neural networks*, vol. 15, no. 8-9, pp. 1041–1058, 2002.

[12] N. Moustafa and J. Slay, "Unsw-nb15: a comprehensive data set for network intrusion detection systems (unsw-nb15 network data set)," in *2015 military communications and information systems conference (MilCIS)*. IEEE, 2015, pp. 1–6.

[13] H. Hazan, D. J. Saunders, H. Khan, D. Patel, D. T. Sanghavi, H. T. Siegelmann, and R. Kozma, "Bindsnet: A machine learning-oriented spiking neural networks library in python," *Frontiers in neuroinformatics*, vol. 12, p. 89, 2018.

[14] M. Kassem, A. H. Al-Hamami, and A. Kassem, “Machine learning-based network intrusion detection for big and imbalanced data using oversampling, stacking feature embedding and feature extraction,” 2024.