Performance Evaluation of Neuromorphic Hardware for Onboard Satellite Communication Applications

Eva Lagunas, *Senior Member, IEEE,* Flor G. Ortiz, *Member, IEEE,* Geoffrey Eappen, *Member, IEEE,* Saed Daoud, *Member, IEEE,* Wallace A. Martins, *Senior Member IEEE*, Jorge Querol, *Member, IEEE,* Symeon Chatzinotas, *Fellow IEEE,* Nicolas Skatchkovsky, *Member, IEEE,* Bipin Rajendran, *Senior Member IEEE*, and Osvaldo Simeone, *Fellow IEEE.*

***Abstract*—Spiking neural networks (SNNs) implemented on neuromorphic processors (NPs) can enhance the energy efficiency of deployments of artificial intelligence (AI) for specific workloads. As such, NP represents an interesting opportunity for implementing AI tasks on board power-limited satellite communication spacecraft. In this article, we disseminate the findings of a recently completed study which targeted the comparison in terms of performance and power-consumption of different satellite communication use cases implemented on standard AI accelerators and on NPs. In particular, the article describes three prominent use cases, namely payload resource optimization, onboard interference detection and classification, and dynamic receive beamforming; and compare the performance of conventional convolutional neural networks (CNNs) implemented on Xilinx's VCK5000 Versal development card and SNNs on Intel's neuromorphic chip Loihi 2.**

***Index Terms*—Neuromorphic Processors, Machine Learning, Satellite Communications.**

I. Introduction

ARTIFICIAL intelligence (AI) has been identified as an essential ingredient of the next generation of wireless communications. AI opens up exciting opportunities, including autonomous network reconfiguration to changing environments, improved system performance, and enhanced customer experience. Given the benefits of AI in the terrestrial wireless environment, the satellite communication (SatCom) research community has also started exploring AI for various applications.

Traditionally, SatCom payloads have been built based on hardware components generally conceived for dedicated tasks. Software-defined radio (SDR) based technology has proven to be beneficial to the satellite industry [1]. By replacing hardware components with software, SDR allows a more flexible satellite payload, offering system reconfiguration in real time.

Existing works in the open literature have analyzed the adoption of machine learning (ML) algorithms for SDR-based SatCom use cases at the level of concepts and simulations [2], [3]. While testing in these works relies on software models, their successful transition to real deployments must hinge on their efficient implementation on chipsets that are able to perform the task within reasonable accuracy, computation time, and power budget. This is particularly relevant for onboard AI implementation, where the AI chipset is to be integrated within the satellite payload.

Neuromorphic processors (NPs) represent a potentially revolutionary solution for efficiently deploying AI models targeting specific workloads. NPs implement energy-efficient spiking neural networks (SNNs) that mimic the operation of the human brain. Unlike conventional Von-Neumann computing, NPs do not need to fetch their instructions and data from memory, and then process them in a sequential step. NPs process data in real-time, overcoming the bottleneck of the memory-processor bus, which is significantly power-hungry. Given the stringent onboard power limitations of satellite platforms, NPs represent a major opportunity to unlock the potential benefits of AI and ML solutions for SatCom systems, thanks to their energy efficiency. Power must be generated onboard, typically exploiting the sunlight using solar panels. However, the size of such panels dictate the size of the satellite and determines the launching cost, being one of the most critical design aspects of space missions. Furthermore, the capability to dynamically reconfigure the satellite payload and antenna pattern is becoming a must for future broadband satellite systems. Adapting transmission parameters according to varying environment conditions and spatio-temporal traffic demand variations allows to satellite operators to maximize the amount of system capacity that is actually used (i.e. sold) while ensuring quality of service.

Industry and academia are currently working on the development of specific AI processors that provide better energy and latency performance for computationally intensive AI algorithms. Some of these AI processors are already available on the market, while many are still in testing and design phase. In this work, we first discuss the methodology for shortlisting the most promising SatCom use cases, and provide a description of the three selected ones, which are: (1) Resource optimization in flexible satellite payloads; (2) Onboard interference detection and classification; and (3) Dynamic digital beamforming for fast-moving users. Subsequently, we discuss and analyze the representative SatCom use cases using two promising alternative commercial options, namely Xilinx's VCK5000 Versal development card for classical AI deployments and Intel's Loihi 2 chipset for neuromorphic platforms. The first is an acceleration chip from the Versal family already commercialized by Xilinx, which was specifically designed to implement high throughput AI inference and signal processing computing performance. The latter corresponds to the second generation of Intel's neuromorphic research chips. We quantify gains in terms of computational time and power consumption, as well as performance on the specific communications-related task. The research activity presented in this manuscript was carried out in collaboration with the European Space Agency (ESA) in the context of the ESA NeuroSat project [4].

This article is organized as follows. First, we describe the selected use cases in Section II. Next, we discuss and compare the two selected AI chipsets in Section III, and discuss the data encoding methodologies for neuromophic implementation in Section IV. Finally, we present the performance comparisons in Section V, followed by the conclusion in Section VI.

II. Selected SatCom Use Cases

In recent years, AI has been applied for a plethora of SatCom use cases [2], [3]. Such widespread interest is a natural response to the emergence of a new technology, and it highlights the need for a clear identification of the most practically relevant use cases.

In the context of the ESA project NeuroSat [4], a workshop including both commercial AI-chip vendors and SatCom service providers, as well as leading European research institutions, was held in July 2022 with the aim of collecting feedback from experts regarding the most promising and feasible scenarios and use cases for the implementation of learning-based methods. Participants ranged from space companies like Hispasat, Thales Alenia Space, and SES; neuromorphic chipset manufacturers like IBM and Intel; as well as national space agencies like NASA and ESA; and few academic institutions. Details on the workshop are available on the project website [4]. The participants were asked to vote for the most promising scenario based on their expertise. The outcomes of such voting are detailed in Fig. 1. Notably, SatCom service providers identified use cases such as flexible payload operations and interference detection and mitigation as critical, given their current needs and the recent advances in network reconfiguration and the increasing spectrum congestion. In contrast, they tended to give less importance to well-studied problems for which they already have efficient solutions based on classical optimization such as precoding matrix calculation. For NP experts, having the possibility of a good temporal encoding for the corresponding AI-model inputs was found to be crucial to obtain gains via NP deployments.

Based on the preferences of stakeholders and considering the potential gains that AI could offer, the following use cases were chosen for further evaluation:

1) resource optimization in flexible satellite payloads.
2) onboard interference detection and classification.
3) dynamic digital beamforming for fast-moving users.

In the following subsections, we present the details of the shortlisted use cases.

A. Resource Optimization in Flexible Satellite Payloads

Dynamically adapting the radio resource allocation to match the spatiotemporal on-ground traffic demand variations is a challenging task for the new generation of software-defined satellite payloads [5]. Power and bandwidth assignment per beam are the two elements typically considered as degrees of freedom to achieve a good match between the users' on-ground demand and the actual capacity offered by the communication satellite.

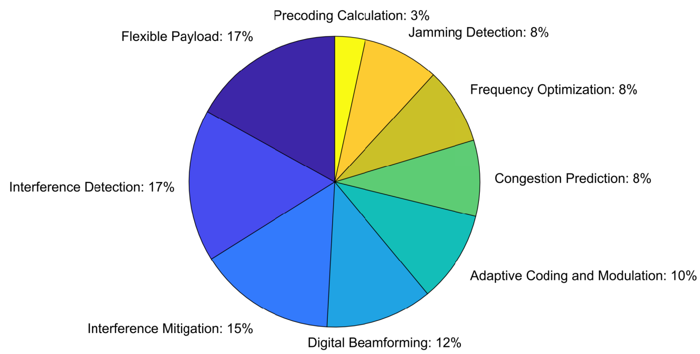

Fig. 1. Outcome of use case relevance for AI application from survey conducted on industry and academia representatives. [4]

The input of the resource optimization is given by the on-ground traffic demand. Such input considers a geographical map of the service area, divided into a grid of latitude and longitude points, forming a matrix. This matrix represents the traffic demand at each point on the map. Each element $r_{i,j}$ in this matrix represents the required data traffic (in Mbps) at a specific geographic location indexed by i and j. Each location within the satellite service area provides information about its data traffic needs, while areas outside the service zone are irrelevant and are thus assigned zero values.

The objective is to determine the optimal combination of power and bandwidth settings for each satellite beam within our service area. We create "payload configurations" consisting of power and bandwidth pairs for each beam. A particular payload configuration selected for a specific time τ can be denoted as $\left[(P_\tau^1, \mathrm{W}_\tau^1), (P_\tau^2, \mathrm{W}_\tau^2), \ldots, (P_\tau^K, \mathrm{W}_\tau^K)\right]$, where $k \in \{1, \ldots, K\}$ denotes the index of satellite beams. The selected power and bandwidth belong to a discrete set, i.e., $P_\tau^k \in \mathcal{P}_\tau$ and $\mathrm{W}_\tau^k \in \mathcal{W}_\tau$. However, it is crucial to recognize that not all combinations might be feasible due to practical constraints on total power and bandwidth.

The proposed approach is a classification problem, where a convolutional neural network (CNN) takes as input the geographical traffic demand distribution and selects the appropriate payload configuration. To minimize complexity in the number of outputs (i.e., possible payload configurations), instead of considering all possible combinations of power and bandwidth, we make a simplification by: (1) Considering only the combinations of power and bandwidth options that comply with the system constraints in terms of maximum power and bandwidth; and (2) Pre-training the system and removing the payload configurations from the candidate pool that are rarely selected for accommodating the traffic demand patterns expected by the system. Once the output possibilities are narrowed down, supervised learning comes into play to match traffic demand with the optimal payload configuration.

The optimization task for flexible satellite payload is summarized in Table I.

TABLE I
SATELLITE PAYLOAD CONFIGURATION OPTIMIZATION

Input	On-ground traffic demand in Mbps per latitude and longitude grid point.
Output	A selection from a limited set of power-per-beam and bandwidth-per-beam configurations.
Objective	Minimize the difference between on-ground traffic demand and offered capacity.
Type of Problem	Classification, i.e., matching each input with one of the potential outputs.

B. Onboard Interference Detection and Classification

The intended or unintended interference experienced in space, either generated by ground transmitters or by other SatCom systems, represents one of the major problems of the satellite industry, which is being aggravated by the trends of launching more and more in-orbit satellite systems, particularly the popular low-Earth orbit (LEO) constellations. Interference events may occur due to intentional jamming, or more likely due to equipment and/or human errors (e.g., antenna's misalignment and cross-polarization effects). For a satellite operator, the first step is to detect the interference event and characterize the received interference. In our approach, the interference detection and classification problem is reduced to a classification problem with the following output classes:

- The received signal is free of interference (class 0).
- The received signal is interfered at subband $\mathcal{F}_i$, $i \in \{1,\ldots,L\}$ (classes $1,2,\ldots,L$).

The onboard interference detection task is summarized in Table II.

TABLE II
ONBOARD INTERFERENCE DETECTION AND CLASSIFICATION

Input	Frequency and/or time domain representation of the received signal.
Output	The frequency band of the detected interference (out from a limited set of options).
Objective	Detect the interference and classify the spectrum of interest.
Type of Problem	Classification, i.e., matching each input with one of the potential output.

C. Dynamic Digital Beamforming for Fast-Moving Users

Receiver beamforming in the user uplink can significantly improve the received signal quality. However, it might be challenging to properly steer the beams toward users that move at high speeds (e.g., aircraft), especially when it comes to non-geostationary orbits. For such high-speed users, the location information received from the gateway might be substantially outdated in scenarios with small spot-beams and, therefore, render beam-pointing errors. In this context, the problem boils down to selecting the optimum beamforming coefficients that allow to point the beam toward the fast-moving user. To address this problem, we use a least absolute shrinkage and selection operator (LASSO) regression to obtain the targeted beamforming vector, where the sparsity term is applied to the real and imaginary parts of the beamforming vector in an attempt to maximize the zero components and, thus, switch off as many radio frequency (RF) chains as possible. The beamforming task is summarized in Table III.

TABLE III
DYNAMIC DIGITAL BEAMFORMING FOR FAST-MOVING USERS

Input	Received signal at the satellite composed of a single uplink from a specific fast-moving aircrafts.
Output	Onboard beamforming weights, which could in principle assume any complex value.
Objective	A design for receive digital beamforming which can be implemented turning off as many RF chains as possible.
Type of Problem	LASSO regression optimization.

III. SELECTED AI CHIPSETS: VCK5000 VS. LOIHI 2

Although the field of AI has evolved rapidly in the past decades, computing processors on which AI algorithms run have not developed at the same pace, which may limit or delay their anticipated benefits. Existing computers are predominantly based on the classical von Neumann architecture, where computation and memory are implemented as separate elements connected via a common data bus, resulting in significant overheads, both in terms of latency and energy for shuttling data back and forth. This bottleneck is critical for AI workloads that involve constant fetching of big amounts of data, impacting power consumption and peak performance [6]. As a response, the community has been intensively investigating AI hardware accelerators, i.e., dedicated and customized processors for AI-specific tasks.

A detailed overview of commercial off-the-shelf (COTS) AI-capable chipsets was presented in [7]. The Xilinx Versal family showed the best trade-off in terms of performance and computation efficiency, while keeping an easy-to-use software flow. Based on that, and on the current availability and lead times, we selected the Xilinx VCK5000 for the evaluation of the machine learning models described in this article.

The Xilinx Versal VCK5000 is an ML-based development card built on Xilinx FPGAs and adaptive compute acceleration platforms. It is designed for high-efficiency AI acceleration with optimized deep learning processor unit cores for solving problems in 5G and beyond communication, signal processing, radar, and satellite-based applications. In the present scenario, VCK5000 is a fully developed AI chip capable of running CNNs, recurrent neural networks, and natural language processing-based models for cloud and edge applications. The training of a neural network model on VCK5000 employs three primary steps, which are quantization, compilation, and deployment. The insights obtained from training the CNN model on Versal VCK5000 is that the VCK5000 is a domain-specific architecture. The AI developer on VCK5000 should be specific about the deep learning processor unit, as it is related to the employed AI models. The training of CNN models on VCK5000 has high computational accuracy. This results in higher energy consumption as a trade-off.

Recently, the concept of NP has gathered significant attention as an alternative architecture imitating the biological brains by operating in an event-driven fashion. NPs may therefore represent a suitable approach to unlocking the potential benefits of AI. To select the NP for this study, we considered the most cutting-edge hardware options available in 2022. By that time, Intel's had recently announced Loihi 2 chipset, a research neuromorphic chip that uses asynchronous spiking neurons to implement fine-grained, event-driven, adaptive, self-modifying, parallel computations. The advanced technology of Loihi 2 combined with the open source framework associated to it (i.e. "Lava"), provided the best tool and necessary abstractions to develop applications for neural computation. Intel's Lava simulator is a programming framework developed by Intel designed to make neuromorphic systems accessible to non-experts developers [8]. To carry out our study, we joined the Intel Neuromorphic Research Community (INRC), which provides access to the latest NP technology from Intel for research purposes. Loihi 2, the next-generation iteration of Loihi, serves as a follow-up to the original neuromorphic research test chip, Loihi. Within Loihi 2, multiple neuromorphic cores, each housing numerous artificial neurons, are interconnected, and they receive spikes from other parts of the network. When these received spikes accumulate over a specific time interval and surpass a predetermined threshold, the respective core initiates its own spike transmission to connected neurons. Previous spikes strengthen existing neuronal connections, whereas subsequent spikes inhibit these connections, gradually reducing connectivity until all activity ceases. Loihi 2 is constructed using the Intel 4 process and boasts a total of 1 million artificial neurons per chip and 120 million synapses per chip. In addition to the 128 neuromorphic cores, the chip incorporates 6 processor cores [9].

IV. Data Encoding for Neuromorphic Implementation

While an SNN can theoretically receive data in the form of an analog input current, NPs can typically only handle data in the form of binary streams of spikes as inputs. Consequently, the natural signals we consider throughout the aforementioned SatCom use cases need to be encoded into binary spikes for processing via Intel's Loihi 2 chip.

The general encoding procedure is as follows. Considering a general feature matrix $\boldsymbol{x} \in \mathbb{R}^{n \times m}$, we first obtain an $nm \times 1$ vector $[x_{1,1} \ldots x_{n,m}]^{\mathrm{T}}$, before performing the encoding process into a spiking signal $\boldsymbol{X} \in \{0,1\}^{nm \times T}$, with T being the number of encoding time-steps.

Among the different neural coding schemes [10], we focus on two of the most predominant ones: rate coding and temporal coding, both illustrated in Fig. 2 and briefly explained below:

a) rate coding: Rate coding utilizes spiking rates to represent information, i.e., larger real-valued input generates a large number of spikes within a fixed encoding window time.

b) temporal coding: Temporal coding utilizes precise spike arrival times to encode information.

Unfortunately, the selection of coding technique is difficult, as some might be better suited than others depending on the type of data and the structure of the spiking neural network.

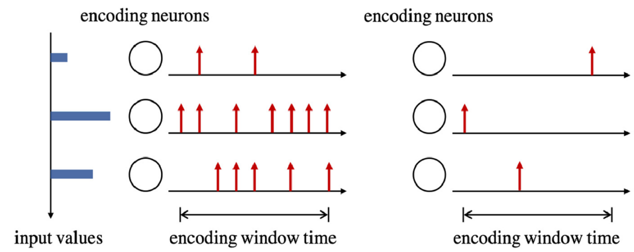

Fig. 2. Illustration of encoding schemes: rate coding (left) and temporal coding (right).

Therefore, there is no "one-coding-fits-all" solution. We now describe the details of signal encoding used for the selected use cases.

A. Resource Optimization in Flexible Satellite Payloads

The inputs to be encoded in the resource optimization use case described in Section II-A correspond to the on-ground traffic demands in latitude and longitude, at different times of the day. Since the temporal aspect is intrinsically related to the input signal, we propose to use temporal coding for the on-ground traffic demand of each geographical location. Therefore, we end up with one time series for each geographical position. More precisely, we consider a temporal model based on leaky integrate-and-fire (LIF) neurons, whereby a spike is emitted when the "voltage" of the neuron (i.e., the traffic demand) crosses a predefined threshold.

B. Onboard Interference Detection and Classification

The use case scenario of the interference detection employs discrete Fourier transform (DFT), implemented via a fast Fourier transform (FFT) algorithm, on the input samples. The magnitude of the FFT signal can be encoded either with rate or temporal coding without problem. However, the preprocessing of the temporal signal to compute its spectral components via an auxiliary processor may incur a significant latency in a neuromorphic system. Recent works have demonstrated that resonate-and-fire (R&F) neurons can compute the short-term Fourier transform (STFT) directly in the spiking domain [11]. An R&F neuron is an oscillatory system that operates by maintaining a complex-valued internal variable that spikes when its real part crosses a predetermined threshold and its argument is zero. For comparison purposes, we checked the computation of the STFT on an auxiliary processor and followed by temporal coding.

C. Dynamic Digital Beamforming for Fast-Moving Users

The LASSO beamforming solution does not involve a learning task but rather an optimization procedure. In this case, we selected the spiking locally competitive algorithm (S-LCA) to solve the problem. Locally competitive algorithm (LCA) was first introduced in [12], as the first neural network model aimed at solving LASSO-type of problems. The S-LCA is essentially the same but adapted to spiking networks, providing the way to implement LASSO problems in SNN in an efficient and effective manner. The details of the S-LCA can be found in [13].

V. Performance Comparison

With the aim of having a fair comparison in terms of energy consumption and processing time, both non-neuromorphic and neuromorphic models were fed with the same input data, and they were given the same output data. The same training dataset was given to use-cases in Section IV-A and IV-B, which were designed aiming at an accuracy on the training and on the validation dataset of $>97\%$.

The graph in Fig. 3 displays the results for the first two use cases, namely the resource optimization for flexible payloads (denoted by RRM) and the interference detection and classification (denoted by ID). Fig. 3 offers a comparison in a single plot of the two chipsets that have been studied, i.e., Xilinx's VCK5000 and Intel's Loihi 2, in terms of energy consumption and time taken to converge to a solution (delay).

To make sense of the graph, we follow [14] and adopt the energy-to-delay ratio (EDP) as a benchmark. The EDP metric has been selected because it captures two key metrics in a single figure of merit: (i) the total energy consumption, and (ii) the delay or amount of time for executing the inference. In particular, this ratio serves as a reference point and is interpreted as follows:

- **Points above the EDP:** Any data points located further to the right of the EDP line represent instances where Intel's Loihi 2 outperforms the Xilinx's VCK5000 in the sense that Loihi 2 is more efficient in terms of power and time.
- **Points below the EDP:** Conversely, data points in the lower-left portion of the graph indicate scenarios where the reference architecture Xilinx's VCK5000 perform better than Loihi 2. These are situations in which the reference architectures are more efficient.

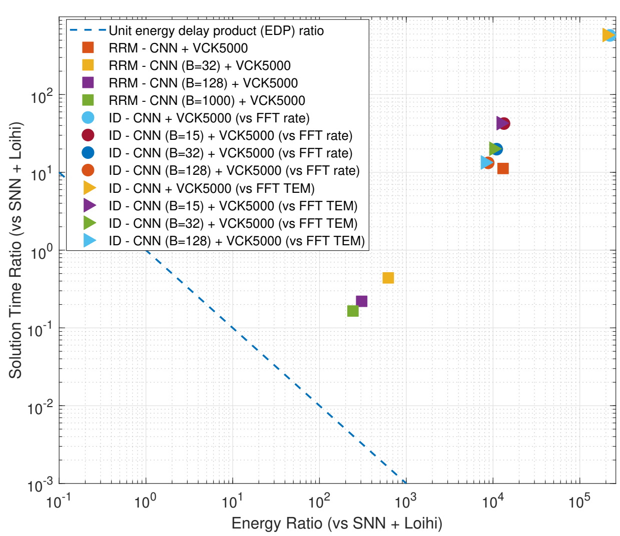

Fig. 3. Comparison in terms of energy ratio and delay ratio comparing the two chipsets (Xilinx's VCK5000 and Intel's Loihi 2) for different applications.

Fig. 3 shows results for different batch sizes, denoted by B, i.e., the number of input samples given to the chip. The terms "FFT rate" and "FFT TEM" in the legend of Fig. 3 refer to the use of either the rate or temporal encoding applied to the frequency representation of the signal, respectively. Below we summarize the conclusions that can be extracted from Fig. 3:

- **Superiority of SNN and Loihi 2:** The results highlight that, in all cases, Intel's Loihi 2 performs better than the CNN implemented in Xilinx's VCK5000. However, it is worth noting that, as the batch size increases, the advantage of Loihi became less pronounced as the performance points move closer to the EDP line, although it still outperforms the Xilinx's VCK5000 implementation. The latter confirms that NP is a better fit for processing small batch size data, and thus a better fit for real-time decision making.
- **Interference detection as a promising use case:** Among the considered scenarios, it seems that the interference detection and classification benefited the most from the implementation on Intel's Loihi chipset. Even though the time ratio remained generally higher than one, the energy savings achieved with Loihi were significant. In some cases, the energy ratio reached values as high as 10^5, which is remarkable.
- **SNN encoding:** Fig. 3 compares the impact of FFT Rate and FFT TEM encoding for the ID scenario. Interestingly, the type of coding used did not have a significant impact on this comparison.
- **RRM's Performance:** Although RRM did not achieve as pronounced energy savings as in the ID scenario, all RRM experiments have shown an energy ratio greater than 10^2. However the energy ratio is significantly higher for batch size 1 (i.e. energy ratio greater than 10^4, with time ratio around 10). The latter was expected, as NPs are designed to work with data on the fly without the need of batch-processing. As a consequence, batch sizes greater than 1 are shown to reduce both energy and time ratio.

Regarding the digital beamforming performance for fast-moving users, we compared the conventional LASSO solution provided by CVX [15] (a well-know Matlab-based software toolbox for convex optimization) with the solution of S-LCA on Intel's Lava simulator. Note that this use-case was run on a simulator, not in the chipset itself. This is because S-LCA was not yet available in Loihi 2. Firstly, it is worth highlighting that the proposed beamforming formulation yielded sparse beamforming vectors, with both solutions being able to turn off up to 60% of the RF chains without compromising the resulting beampatterns. Regarding performance comparisons between the two solutions, both generated satisfactory beampatterns with the main lobe pointing toward the aircraft. For a numerical comparison, the beamformer's average output power was considered as key performance indicator to assess the beamforming capabilities to mitigate the effects of noise and interference while focusing on the desired signal direction. In this context, the S-LCA solution was able to reach lower levels of beamformer's average output power, around 19% below the value reached by the CVX solution, but with a much higher spreading of values, around 4 times higher than the CVX solution, when comparing the lower and upper quartiles of beamformer's average output power.

Coming back to the results of the RRM and interference detection use-cases (see Section IV-A and IV-B, respectively), the results presented in this article were conducted with Loihi 2 but using the remote access from Intel's research cloud. Although Intel offers the possibility of shipping physical chips to INRC partners premises, at the moment of developing these results the primary access to Loihi 2 was through the Intel's neuromorphic research cloud. Obviously, the remote access introduced some additional limitations as it was not possible to control for concurrent usage by other users, which could lead to delays and increased power consumption. Additionally, the specific interface used by Intel on their cloud was not disclosed, potentially resulting in differences when conducting measurements with future releases of Loihi 2. Furthermore, due to the runtime limitation of 20 minutes for jobs on Intel's cloud, on-chip training was not possible, restricting the exploration of this important capability of Loihi.

The cloud interface plays a key role, as it impacts the transfer of spiking signals to the chipset. High input size may span long per-step processing time. For example, in the flexible payload use case, the execution time per example increased from approximately 5 ms to 100 ms when the input size went from 252 neurons to 299 neurons.

VI. Conclusion

While we enter the era of AI, it becomes evident that energy consumption is a limiting factor when training and implementing neural networks with significant number of neurons. This issue becomes particularly relevant for nonterrestrial communication devices, such as satellite payloads, which are in need of more efficient hardware components in order to benefit from the potential of AI techniques.

In this work, we analyzed 3 on-board satellite use-cases and we identified significant power savings when utilizing the neuromorphic chipset. Furthermore, the performance gap between the NP and the benchmark AI accelerator increases in favor of NP when processing individual data samples and are, therefore, a better fit for use cases where real world data arrives to the chip and it needs to be processed right away. In this article, we verified these hypothesis using real standard processor solutions, such as the Xilinx's VCK5000 and the Intel's Loihi 2 chipset.

Acknowledgments

This work has been supported by the European Space Agency (ESA) funded under Contract No. 4000137378/22/UK/ND - The Application of Neuromorphic Processors to Satcom Applications. Please note that the views of the authors of this paper do not necessarily reflect the views of ESA. Furthermore, this work was partially supported by the Luxembourg National Research Fund (FNR) under the project SmartSpace (C21/IS/16193290).

References

[1] Marco Giordani and Michele Zorzi, "Non-terrestrial networks in the 6G era: Challenges and opportunities," *IEEE Network*, vol. 35, no. 2, pp. 244–251, 2021.

[2] M.A. Vazquez, P. Henarejos, I. Pappalardo, E. Grechi, J. Fort, J.C. Gil, and R.M. Lancellotti, "Machine Learning for Satellite Communications Operations," *IEEE Communications Magazine*, vol. 59, no. 2, pp. 22–27, 2 2021.

[3] G. Fontanesi et al., "Artificial intelligence for satellite communication and non-terrestrial networks: A survey," https://arxiv.org/abs/2304.13008, 2023.

[4] ESA FPE NeuroSat Project, "The Application of Neuromorphic Processors to Satcom Applications," https://sites.google.com/view/evalagunas/home/projects/esa-neurosat.

[5] S. Kisseleff, E. Lagunas, T.S. Abdu, S. Chatzinotas, and B. Ottersten, "Radio Resource Management Techniques for Multibeam Satellite Systems," *IEEE Communications Letters*, vol. 25, no. 8, pp. 2448–2452, 2021.

[6] G. Indiveri and Y. Sandamirskaya, "The Importance of Space and Time for Signal Processing in Neuromorphic Agents: The Challenge of Developing Low-Power, Autonomous Agents That Interact With the Environment," *IEEE Signal Processing Magazine*, vol. 36, no. 6, pp. 16–28, 2019.

[7] F. Ortiz et al., "Onboard Processing in Satellite Communications Using AI Accelerators," *MDPI Aerospace*, vol. 10, no. 2, 2023.

[8] Intel, "Lava Software Framework," [Online]. Available from: https://lava-nc.org/, 2021, Last accessed 16 April 2024.

[9] M. Davies et al., "Taking neuromorphic computing to the next level with Loihi 2," *Intel Labs' Loihi*, vol. 2, pp. 1–7, 2021.

[10] D. Auge, J. Hille, E. Mueller, and A. Knoll, "A Survey of Encoding Techniques for Signal Processing in Spiking Neural Networks," *Neural Process Letters*, vol. 53, pp. 4693–4710, 2021.

[11] G. Orchard et al., "Efficient Neuromorphic Signal Processing with Loihi 2," in *IEEE Workshop on Signal Processing Systems (SiPS)*, 2021, pp. 254–259.

[12] C.J. Rozell, D.H. Johnson, R.G. Baraniuk, and B.A. Olshausen, "Sparse Coding via Thresholding and Local Competition in Neural Circuits," May Oct. 2008, Neural Comput.

[13] Ping Tak Peter Tang, Tsung-Han Lin, and Mike Davies, "Sparse coding by spiking neural networks: Convergence theory and computational results," 2017.

[14] M. Davies, A. Wild, G. Orchard, Y. Sandamirskaya, G.A. Guerra-Fonseca, 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.

[15] M. Grant and S. Boyd, "CVX: Matlab Software for Disciplined Convex Programming, version 2.1," http://cvxr.com/cvx, Mar. 2014.

VII. Biography Section

Eva Lagunas received the M.Sc. and Ph.D. degrees in telecommunications engineering from the Polytechnic University of Catalonia (UPC), Barcelona, Spain, in 2010 and 2014, respectively. She has held positions at UPC, Centre Tecnologic de Telecomunicacions de Catalunya (CTTC), University of Pisa, Italy; and the Center for Advanced Communications (CAC), Villanova University, PA, USA. In 2014, she joined the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg, where she currently holds a Research Scientist position. Her research interests include terrestrial and satellite system optimization, spectrum sharing, resource management and machine learning.

Flor G. Ortiz received her B.S. degree in telecommunications engineering and her M.S. degree in electrical engineering-telecommunications from the Universidad Nacional Autónoma de México (UNAM), Mexico City, Mexico, in 2015 and 2016, respectively. Flor obtained her Ph.D. degree in Telecommunication Engineering (September 2021) at the Universidad Politecnica de Madrid (UPM), Madrid, Spain. During her PhD, she performed a research period at the University of Bologna, in Bologna, Italy. She started a close collaboration between UPM and the University of Bologna, opening a new research line for both groups on applying Machine Learning for radio resource management. She is joined as a Research Associate at the Interdisciplinary centre for Security, Reliability, and Trust (SnT) at University of Luxembourg. Her research interests are focused on implementing cutting-edge Machine Learning techniques including Continual Learning and and Neuromorphic Computing for operations in Satellite Communications systems.

Geoffrey Eappen (Member, IEEE), received the Ph.D. degree from the Vellore Institute of Technology (VIT), Vellore, and Brunel University London. He was a Senior Research Fellow with the Council of Scientific and Industrial Research India, Department of Wireless Communication (SENSE), VIT. He is currently working as a Research Associate with the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg. His research interest include cognitive radio network, spectrum sensing methodologies, metaheuristic optimization schemes, artificial neural network, 5G beamforming, and satellite communication. He is the recipient of the U.K. Commonwealth Fellowship.

Saed Daoud Saed Daoud received his M.Sc. degree from Jordan University of Science and Technology, Jordan in 2010, and his PhD degree from Concordia University, Montreal, Canada in 2015, both in Electrical Engineering with specialization in Wireless Communication. He was a postdoctoral researcher at Ecole Polytechnique de Montreal, and currently a research associate with the Interdisciplinary Centre for Security, Reliability and Trust (SnT), University of Luxembourg. His research interests broadly cover physical-layer transceiver design and performance evaluation for underwater acoustic (UWA), terrestrial and satellite communication systems, for technologies such as OFDM, SC-FDE, cognitive radio, spread spectrum, MIMO, and precoding.

Wallace A. Martins received his D.Sc. degree in Electrical Engineering from the Federal University of Rio de Janeiro (UFRJ), Brazil, in 2011. He is currently a Full Professor with ISAE-SUPAERO, Université de Toulouse, France. His research interests encompass digital signal processing and telecommunications, focusing on future wireless and satellite networks.

Jorge Querol Jorge Querol (S'15–M'18) was born in Forcall, Castelló, Spain, in 1987. He received the B.Sc. (+5) degree in telecommunication engineering, the M.Sc. degree in electronics engineering, the M.Sc. degree in photonics, and the Ph.D. degree (Cum Laude) in signal processing and communications from the Universitat Politècnica de Catalunya - BarcelonaTech (UPC), Barcelona, Spain, in 2011, 2012, 2013 and 2018 respectively. His Ph.D. thesis was devoted to the development of novel anti-jamming and counter-interference systems for Global Navigation Satellite Systems (GNSS), GNSS-Reflectometry, and microwave radiometry. One of his outstanding achievements was the development of a real-time standalone pre-correlation mitigation system for GNSS, named FENIX, in a customized Software Defined Radio (SDR) platform. FENIX was patented, licensed and commercialized by MITIC Solutions, a UPC spin-off company. Since 2018, he is with the SIGCOM research group of the Interdisciplinary Centre for Security, Reliability, and Trust (SnT) of the University of Luxembourg, Luxembourg and head of the 6GSPACE Laboratory. He is involved in several ESA and Luxembourgish national research projects dealing with signal processing and satellite communications. His research interests include SDR, real-time signal processing, satellite communications, 5G non-terrestrial networks, satellite navigation, and remote sensing. He received the best academic record award of the year in Electronics Engineering at UPC in 2012, the first prize of the European Satellite Navigation Competition (ESNC) Barcelona Challenge from the European GNSS Agency (GSA) in 2015, the best innovative project of the Market Assessment Program (MAP) of EADA business school in 2016, the award Isabel P. Trabal from Fundació Caixa d'Enginyers for its quality research during his Ph.D. in 2017, and the best Ph.D. thesis award in remote sensing in Spain from the IEEE Geoscience and Remote Sensing (GRSS) Spanish Chapter in 2019.

Symeon Chatzinotas (MEng, MSc, PhD, FIEEE) is currently Full Professor / Chief Scientist I and Head of the research group SIGCOM in the Interdisciplinary Centre for Security, Reliability and Trust, University of Luxembourg. In parallel, he is an Adjunct Professor in the Department of Electronic Systems, Norwegian University of Science and Technology and a Collaborating Scholar of the Institute of Informatics & Telecommunications, National Center for Scientific Research "Demokritos". In the past, he has lectured as Visiting Professor at the University of Parma, Italy and contributed in numerous R&D projects for the Institute of Telematics and Informatics, Center of Research and Technology Hellas and Mobile Communications Research Group, Center of Communication Systems Research, University of Surrey. He has received the M.Eng. in Telecommunications from Aristotle University of Thessaloniki, Greece and the M.Sc. and Ph.D. in Electronic Engineering from University of Surrey, UK in 2003, 2006 and 2009 respectively. He has authored more than 800 technical papers in refereed international journals, conferences and scientific books and has received numerous awards and recognitions, including the IEEE Fellowship and an IEEE Distinguished Contributions Award. He is currently in the editorial board of the IEEE Transactions on Communications, IEEE Open Journal of Vehicular Technology and the International Journal of Satellite Communications and Networking.

Nicolas Skatchkovsky received his Ph.D. in Electrical Engineering from King's College London in 2022, under the supervision of Prof. Osvaldo Simeone. Before that, he received his M.Eng (with Honours) and B.Eng in Electrical Engineering from CentraleSupélec, Gif-sur-Yvette, France. He was a post-doctoral research associate at King's College London from 2022 to 2023, where he worked on Bayesian learning and applications of neuromorphic learning to space communications. He is now a post-doctoral research associate at the Francis Crick Institute, working on the modelling of attention in mamallian brains using machine learning techniques. His current research interests include applications of AI to neurosciences under an information-theoretic perspective.

Bipin Rajendran is a Professor of Intelligent Computing Systems in the Department of Engineering, King's College London, where he co-directs the Centre for Intelligent Information Processing (CIIPS). He received a B. Tech degree from I.I.T. Kharagpur in 2000, and M.S. and Ph.D. degrees in Electrical Engineering from Stanford University in 2003 and 2006, respectively. He was a Master Inventor and Research Staff Member at IBM T. J. Watson Research Center in New York during 2006-'12 and has held faculty positions in India and the US. His research focuses on building algorithms, devices, and systems for intelligent computing systems. He has co-authored over 95 papers in peer-reviewed journals and conferences, one monograph, one edited book, and 59 issued U.S. patents. He is a recipient of the IBM Faculty Award (2019), IBM Research Division Award (2012), and IBM Technical Accomplishment Award (2010). He was elected a senior member of the US National Academy of Inventors in 2019. His research has been supported by the Engineering and Physical Sciences Research Council (EPSRC), the US National Science Foundation (NSF), the European Commission, the European Space Agency, Semiconductor Research Corporation as well as Intel, IBM, and Cisco. In 2022, he was awarded an Open Fellowship of the EPSRC.

Osvaldo Simeone is a Professor of Information Engineering. He co-directs the Centre for Intelligent Information Processing Systems within the Department of Engineering of King's College London, where he also runs the King's Communications, Learning and Information Processing lab. He received an M.Sc. degree (with honors) and a Ph.D. degree in information engineering from Politecnico di Milano, Milan, Italy, in 2001 and 2005, respectively. From 2006 to 2017, he was a faculty member of the Electrical and Computer Engineering (ECE) Department at New Jersey Institute of Technology (NJIT), where he was affiliated with the Center for Wireless Information Processing (CWiP). His research interests include information theory, machine learning, wireless communications, neuromorphic computing, and quantum machine learning. Dr Simeone is a co-recipient of the 2022 IEEE Communications Society Outstanding Paper Award, the 2021 IEEE Vehicular Technology Society Jack Neubauer Memorial Award, the 2019 IEEE Communication Society Best Tutorial Paper Award, the 2018 IEEE Signal Processing Best Paper Award, the 2017 JCN Best Paper Award, the 2015 IEEE Communication Society Best Tutorial Paper Award and of the Best Paper Awards of IEEE SPAWC 2007 and IEEE WRECOM 2007. He was awarded an Open Fellowship by the EPSRC in 2022 and a Consolidator grant by the European Research Council (ERC) in 2016. His research has been also supported by the U.S. National Science Foundation, the European Commission, the European Research Council, the Vienna Science and Technology Fund, the European Space Agency, as well as by a number of industrial collaborations including with Intel Labs and InterDigital. He was the Chair of the Signal Processing for Communications and Networking Technical Committee of the IEEE Signal Processing Society in 2022, as well as of the UK & Ireland Chapter of the IEEE Information Theory Society from 2017 to 2022. He was a Distinguished Lecturer of the IEEE Communications Society in 2021 and 2022, and he was a Distinguished Lecturer of the IEEE Information Theory Society in 2017 and 2018. Prof. Simeone is the author of the textbook "Machine Learning for Engineers" published by Cambridge University Press, four monographs, two edited books, and more than 200 research journal and magazine papers. He is a Fellow of the IET, EPSRC, and IEEE.