



# Brain-Inspired Hardware for Artificial Intelligence: Accelerated Learning in a Physical-Model Spiking Neural Network

Timo Wunderlich<sup>1</sup>(✉), Akos F. Kungl<sup>1</sup>(✉), Eric Müller<sup>1</sup>, Johannes Schemmel<sup>1</sup>,  
and Mihai Petrovici<sup>1,2</sup>

<sup>1</sup> Kirchhoff Institute for Physics, Heidelberg University, Heidelberg, Germany  
{timo.wunderlich, fkungl}@kip.uni-heidelberg.de

<sup>2</sup> Department of Physiology, University of Bern, Bern, Switzerland

**Abstract.** Future developments in artificial intelligence will profit from the existence of novel, non-traditional substrates for brain-inspired computing. Neuromorphic computers aim to provide such a substrate that reproduces the brain's capabilities in terms of adaptive, low-power information processing. We present results from a prototype chip of the BrainScaleS-2 mixed-signal neuromorphic system that adopts a physical-model approach with a 1000-fold acceleration of spiking neural network dynamics relative to biological real time. Using the embedded plasticity processor, we both simulate the Pong arcade video game and implement a local plasticity rule that enables reinforcement learning, allowing the on-chip neural network to learn to play the game. The experiment demonstrates key aspects of the employed approach, such as accelerated and flexible learning, high energy efficiency and resilience to noise.

**Keywords:** Neuromorphic computing · Spiking networks · Plasticity

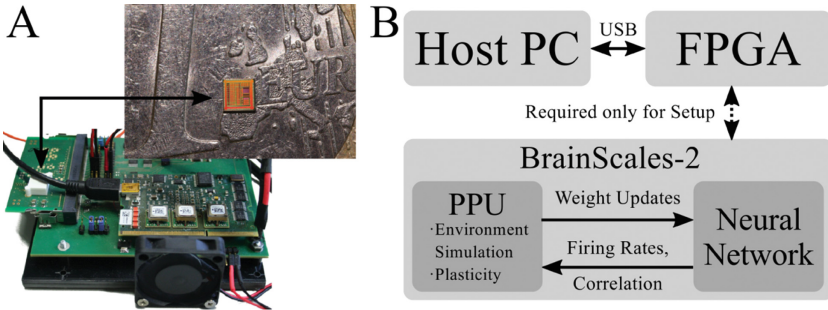
## 1 Introduction

Many breakthrough advances in artificial intelligence incorporate methods and algorithms that are inspired by the brain. For instance, the artificial neural networks employed in deep learning are inspired by the architecture of biological neural networks [1]. Very often, however, these brain-inspired algorithms are run on classical von Neumann devices that instantiate a computational architecture remarkably different from the one of the brain. It is therefore a widely held view that the future of artificial intelligence will depend critically on the deployment of novel computational substrates [2]. Neuromorphic computers represent an attempt to move beyond brain-inspired software by building hardware that structurally and functionally mimics the brain [3].

In this work, we use a prototype of the BrainScaleS-2 (BSS2) neuromorphic system [4]. The employed physical-model approach enables the accelerated (1000-fold with respect to biology) and energy-efficient emulation of spiking neural networks (SNNs). Beyond SNN emulation, the system contains an embedded

plasticity processing unit (PPU) that provides facilities for the flexible implementation of learning rules. Our prototype chip (see Fig. 1A) contains 32 physical-model neurons with 32 synapses each, totalling 1024 synapses. The neurons are an electronic circuit implementation of the leaky integrate-and-fire neuron model. We use this prototype to demonstrate key advantages of our employed approach, such as the 1000-fold speed-up of neuronal dynamics, on-chip learning, high energy-efficiency and robustness to noise in a closed-loop reinforcement learning experiment.

## 2 Experiment




**Fig. 1.** A: prototype chip and evaluation board. B: schematic of experimental setup. The chip runs the experiment fully autonomously, with the PPU simulating the virtual environment (Pong) and calculating reward-based weight updates that are applied to the on-chip analog neural network. The neural network receives the discretized ball position as input and outputs the target paddle position. B taken from [5].

The experiment represents the first demonstration of on-chip closed-loop learning in an accelerated physical-model neuromorphic system [5]. It takes place on the chip fully autonomously, with external communication only required for initial configuration (see Fig. 1B). We use the embedded plasticity processor both to simulate a simplified version of the Pong video game (opponent is a solid wall) and to implement a reward-modulated spike-timing-dependent plasticity (R-STDP) learning rule [6] of the form  $\Delta w_{ij} \propto (R - \bar{R})e_{ij}$ , where  $R$  is the reward,  $\bar{R}$  is a running average of the reward and  $e_{ij}$  is an STDP-like eligibility trace. Each synapse locally records the STDP-like eligibility trace and stores it as an analog value (a voltage), to be digitized and used by the plasticity processor [4].

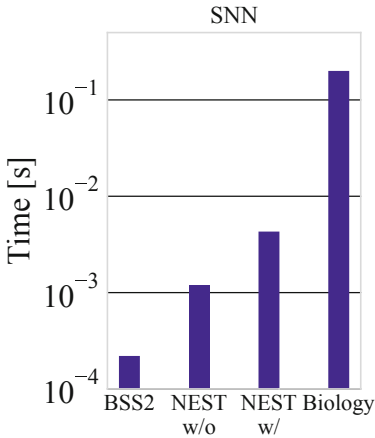
The two-layer neural network receives the ball position along one axis as input and dictates the target paddle position, to which the paddle moves with constant velocity, using the neuronal firing rates. It receives reward depending on its aiming accuracy (i.e., how close it aims the paddle to the center of the ball), with  $R = 1$  for perfect aiming,  $R = 0$  for not aiming under the ball and graded steps in between. By correlating reward and synaptic activity via the given learning rule, the SNN on the chip learns to trace the ball with high fidelity.

### 3 Results



**Fig. 2.** Screenshot of experiment interface for live demonstration. Left: playing field and color-coded neuronal firing rates. Middle: color-coded synaptic weight matrix. Right: performance in playing Pong.

A screen recording of a live demonstration of the experiment (see Fig. 2) is available at <https://www.youtube.com/watch?v=LW0Y5SSIQU4>. The recording allows the viewer to follow the game dynamics, neuronal firing rates, synaptic weight dynamics and learning progress. The learning progress is quantified by measuring the relative number of ball positions for which the paddle is able to catch the ball. The learned weight matrix is diagonally dominant: this expresses a correct mapping of states to actions in the reinforcement learning paradigm.



**Fig. 3.** Comparison of experiment durations. w/o: without. w/: with. Adapted from [5].

v2.14.0) on an Intel processor (i7-4771), when considering only the numerical state propagation, is at least an order of magnitude slower than our neuromorphic emulation (see Fig. 3 and [5]). Besides this, the emulation on our prototype is at least 1000 times more energy-efficient than the software simulation

Importantly, neuronal firing rates vary from trial to trial due to noise in the analog chip components. Used appropriately, this can become an asset rather than a nuisance: in our reinforcement learning scenario, such variability endows the neural network with the ability to explore the action space and thereby with a necessary prerequisite for trial-and-error learning. We also found that neuronal parameter variability due to fixed-pattern noise, which is inevitable in analog neuromorphic hardware, is implicitly compensated by the chosen learning paradigm, leading to a correlation of learned weights and neuronal properties [4].

The accelerated nature of our substrate represents a key advantage. We found that a software simulation (NEST

(23  $\mu$ J vs. 106 mJ per iteration). This evinces the considerable benefit of using the BSS2 platform for emulating spiking networks and hints towards its decisive advantages when scaling the emulated networks to larger sizes.

## 4 Conclusions

These experiments demonstrate, for the first time, functional on-chip closed-loop learning on an accelerated physical-model neuromorphic system. The employed approach carries the potential to both enable researchers with the ability to investigate learning processes with a 1000-fold speed-up and to enable novel, energy-efficient and fast solutions for brain-inspired edge computing. While digital neuromorphic solutions and supercomputers generally achieve at most real-time simulation speed in large-scale neural networks [7–9], the speed-up of BrainScaleS-2 is independent of network size and will become a critical asset in future work on the full-scale BrainScaleS-2 system.

## References

1. Hassabis, D., et al.: Neuroscience-inspired artificial intelligence. *Neuron* **95**, 245–258 (2017). <https://doi.org/10.1016/j.neuron.2017.06.011>
2. Dean, J., et al.: A new golden age in computer architecture: empowering the machine-learning revolution. *IEEE Micro* **38**(2), 21–29 (2018). <https://doi.org/10.1109/MM.2018.112130030>. ISSN 0272-1732
3. Schuman, C.D., et al.: A survey of neuromorphic computing and neural networks in hardware. *CoRR* abs/1705.06963 (2017)
4. Friedmann, S., et al.: Demonstrating hybrid learning in a flexible neuromorphic hardware system. *IEEE Trans. Biomed. Circ. Syst.* **11**(1), 128–142 (2017). <https://doi.org/10.1109/TBCAS.2016.2579164>
5. Wunderlich, T., et al.: Demonstrating advantages of neuromorphic computation: a pilot study. *Front. Neurosci.* **13**, 260 (2019). <https://doi.org/10.3389/fnins.2019.00260>
6. Frémaux, N., et al.: Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Front. Neural Circ.* **9**, 85 (2015). <https://doi.org/10.3389/fncir.2015.00085>. ISSN 1662-5110
7. Jordan, J., et al.: Extremely scalable spiking neuronal network simulation code: from laptops to exascale computers. *Front. Neu roinform.* **12**, 2 (2018). <https://doi.org/10.3389/fninf.2018.00002>. ISSN 1662-5196
8. van Albada, S.J., et al.: Performance comparison of the digital neuromorphic hardware SpiNNaker and the neural network simulation software NEST for a full-scale cortical microcircuit model. *Front. Neurosci.* **12**, 291 (2018). <https://doi.org/10.3389/fnins.2018.00291>. ISSN 1662-453X
9. Mikaitis, M., et al.: Neuromodulated synaptic plasticity on the SpiNNaker neuromorphic system. *Front. Neurosci.* **12**, 105 (2018). <https://doi.org/10.3389/fnins.2018.00105>. ISSN 1662-453X