

Closed-loop experiments on the BrainScaleS-2 architecture

K. Schreiber,
T. C. Wunderlich, C. Pehle
Kirchhoff-Institute for Physics, Heidelberg University
Heidelberg, Germany

M. A. Petrovici
Department of Physiology, University of Bern
Bern, Switzerland

J. Schemmel, K. Meier
Kirchhoff-Institute for Physics, Heidelberg University
Heidelberg, Germany

ABSTRACT

The evolution of biological brains has always been contingent on their embodiment within their respective environments, in which survival required appropriate navigation and manipulation skills. Studying such interactions thus represents an important aspect of computational neuroscience and, by extension, a topic of interest for neuromorphic engineering. Here, we present three examples of embodiment on the BrainScaleS-2 architecture, in which dynamical timescales of both agents and environment are accelerated by several orders of magnitude with respect to their biological archetypes.

CCS CONCEPTS

• **Hardware** → Analog and mixed-signal circuits.

KEYWORDS

Closed-loop, neuromorphic hardware, path integration, reinforcement learning, neurorobotics

ACM Reference Format:

K. Schreiber, T. C. Wunderlich, C. Pehle, M. A. Petrovici, and J. Schemmel, K. Meier. 2020. Closed-loop experiments on the BrainScaleS-2 architecture. In *Neuro-inspired Computational Elements Workshop (NICE '20)*, March 17–20, 2020, Heidelberg, Germany. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3381755.3381776>

1 INTRODUCTION

Neuromorphic engineering aims at overcoming certain limitations of traditional computer architectures by reproducing particular aspects of structure and function of biological neural networks in VLSI. Within this context, it is important to note that all biological brains are part of a body that interacts with its environment in various ways: information continuously flows from a diverse range of sensory organs or cells into the nervous system. The nervous system processes this information and in turn provides signals to organs or cells that are involved in motion or communication, resulting in actions within the environment. This interplay of information exchange and processing appears to be inseparably linked to the working principles of biological brains and forms a closed loop of environment, body, and brain. These important aspects of so-called embodied cognition have received increasing attention in cognitive science and related fields [12], evidently pertaining to both biological and artificial brains. For versions of the latter running on analog neuromorphic hardware, embodiment carries particular

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

NICE '20, March 17–20, 2020, Heidelberg, Germany

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7718-8/20/03.

<https://doi.org/10.1145/3381755.3381776>

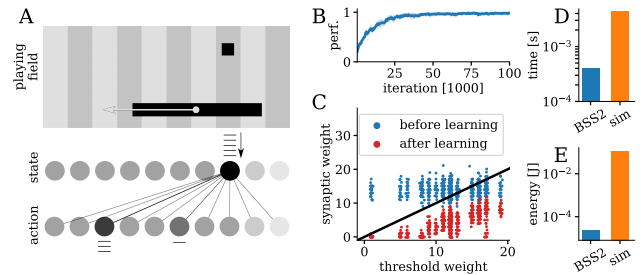


Figure 1: Reinforcement learning with reward-modulated STDP. A) The plasticity processing unit (PPU) simulates a simplified version of Pong. The horizontal position of the ball serves as input for a 2-layer neural network, with the resulting output dictating the target paddle position. The network receives reward based on its aiming accuracy. B) Playing performance during learning. C) Synaptic depression automatically adapts to the excitability of neurons. D, E) Wall-clock duration and power consumption of a single iteration on BrainScaleS-2 (blue) and an equivalent software simulation using NEST (orange).

constraints, as the continuous-time dynamics of physical analog neurons and synapses cannot be accelerated, slowed-down or even halted as easily as for their digital counterparts.

Here, we discuss three different implementations of closed-loop experiments on the BrainScaleS-2 system [1–3]: a neural agent playing the game of pong, a robotic application in which the neural network is connected to a pantographic robotic system, and an insectoid agent that performs a path integration task in a 2D environment. By allowing continuous access to the spike I/O of the analog neural core, as well as digital online updates of network parameters, the hybrid architecture of BrainScaleS-2 lends itself particularly well to the study of neural agents that behave and learn within an interactive environment.

2 REINFORCEMENT LEARNING

Reinforcement learning represents a natural fit for closed-loop experiments in which an agent tries to maximize a reward signal based on its actions within an environment. It was recently shown that BrainScaleS-2 can be used to implement closed-loop reinforcement learning, where all components of the loop, including the virtual environment simulation, are computed on-chip, thus creating a fully autonomous setup [13].

Our referenced work uses a three-factor learning rule called R-STDP [5, 6]. This rule multiplies a scalar reward signal, akin to a global neuromodulator such as dopamine, with an STDP-like synaptic eligibility trace. In our case, the latter is computed locally at each synapse and in an analog fashion [7].

The virtual environment is a simplified version of the Pong video game, as shown in Fig. 1A. The spiking neural network dynamics are emulated by the neuromorphic substrate while the embedded plasticity processor takes on a dual role: it both simulates the game




Figure 2: PlayPen: a pantographic robotic system. A) Arbitrary optical content of a screen is captured by a sensor S and translated into a continuous afferent stream of spikes. This input is, in turn, used by spiking networks emulated on the BrainScaleS-2 system to control the position of the sensor by emitting spike trains that drive the motors MX and MY . B) Example screen contents for four possible experimental tasks: stay within a potential well, follow a gradient, follow a trace, follow a dot.

dynamics and computes synaptic weight updates using R-STDP, creating a fully autonomous setup.

During training, the network learns to keep the paddle approximately centered under the ball (Fig. 1B). Implicitly, the experiment also demonstrates how learning can compensate fixed-pattern noise in the analog neuro-synaptic circuits (Fig. 1C): we found that while the excitability of uncalibrated neurons varied significantly due to mismatch effects, synapses that would negatively impact correct tracking of the ball were systematically depressed below the individual firing threshold of their postsynaptic neuron. Furthermore, this setup demonstrates the speed and power advantages of the BrainScaleS-2 architecture compared to software simulations (Fig. 1D/E).

3 ACCELERATED ROBOTICS

By implementing a fast low-latency link for peripheral spike input and output, we realized an experimental platform that provides BrainScaleS-2 with access to real-world robotic actuators and sensors. Unlike common neurorobotics, the inherent thousandfold speed-up of BrainScaleS-2 demands exceptionally quick actuation and sensory feedback. We therefore needed to reduce the weight – and thus also the complexity – of the mechanical components as much as possible. The setup translates efferent spike trains into impulses that control two motors of a pantographic actuator that moves an optical sensor over a two-dimensional surface. The data gathered by this sensor is in turn translated into an afferent spike train that is transmitted back to the neuromorphic core. This simple construction, which we call the "PlayPen", provides an intuitive and easily configurable setup for the study of spike-based learning in the context of real-world robotic control. The setup is depicted in Fig. 2 together with four example experiments.

4 INSECT NAVIGATION

Recent developments in biological imaging have facilitated unprecedented insight into numerous functional aspects of insect brains [4, 10, 11], such as their navigational capabilities [8]. Based on physiological data from the bee's brain and following [9], we emulated a network for path integration (Fig. 3A) that reproduces




Figure 3: Virtual insectoid agent performing path integration on BrainScaleS-2. A) Network schematic and activity histogram. The information flows from the sensory layer at the top through an integration and a steering layer to the motor neurons at the bottom. R and L indicate the right and left side, respectively. B) A typical trajectory of the virtual insect which turns to random looping around the home position upon reaching it. C) Overlay of 100 trajectories like in B), each with a different random outbound journey.

bees' ability to return to their nest's location after exploring the environment for sources of food.

At the beginning of each experiment, a virtual insect performed a random walk starting from a certain starting position. During this phase, the modeled network had no effect on the insect's motion, but was provided with sensory data of the absolute head orientation and the optical flow field from the insect's eyes. The insect's head orientation was encoded by four spike sources that each represented a cardinal direction similar to a compass. The optical flow field was similarly represented by two spike generators that fired with a rate proportional to the optical flow as derived from the left and right eye (FL and FR). In the second part of the experiment, the return phase, the insect's motion was determined by the two motor neurons in the network (ML and MR), which steered the insect by providing propulsion on the left or right hand side, similar to a tank drive. Across multiple experiments, the emulated navigation network was able to reliably guide the insect back to its starting position (Fig. 3C). As with Pong (Sec. 2), the plasticity processor was used to handle multiple tasks: the processing of synaptic modulations for the integrator neurons, the simulation of the environment, an emulation of all sensors including the corresponding spike stimuli, the translation of neuronal data into actions of motion, and the entire experiment control. The experiment can thus run entirely self-contained on the BrainScaleS-2 system.

5 CONCLUSION

We outlined three agent-environment interaction scenarios using spiking neural networks emulated on the BrainScaleS-2 architecture. In two of these, the sensors and actuators of an embodied neural agent, along with its environment, were also emulated on-chip using the embedded vector processing unit. On the other hand, the high-speed PlayPen setup represents a physical embodiment and environment that facilitates a large variety of experimental tasks on accelerated time scales. These experiments demonstrate the capabilities of the BrainScaleS-2 architecture for studying embodied cognition and provide many of the necessary tools required for this endeavor.

ACKNOWLEDGMENTS

The author wishes to thank all present and former members of the Electronic Vision(s) research group contributing to the BSS-1 and BSS-2 hardware as well as software. We gratefully acknowledge funding from the European Union under grant agreements 604102, 720270, 785907 (HBP) and the Manfred Stärk Foundation.

REFERENCES

- [1] S. A. Aamir, P. Müller, G. Kiene, L. Kriener, Y. Stradmann, A. Grübl, J. Schemmel, and K. Meier. 2018. A Mixed-Signal Structured AdEx Neuron for Accelerated Neuromorphic Cores. *IEEE Transactions on Biomedical Circuits and Systems* 12, 5 (Oct 2018), 1027–1037. <https://doi.org/10.1109/TBCAS.2018.2848203>
- [2] S. A. Aamir, Y. Stradmann, P. Müller, C. Pehle, A. Hartel, A. Grübl, J. Schemmel, and K. Meier. 2018. An Accelerated LIF Neuronal Network Array for a Large-Scale Mixed-Signal Neuromorphic Architecture. *IEEE Transactions on Circuits and Systems I: Regular Papers* 65, 12 (Dec 2018), 4299–4312. <https://doi.org/10.1109/TCSI.2018.2840718>
- [3] Sebastian Billaudelle, Yannik Stradmann, Korbinian Schreiber, Benjamin Cramer, Andreas Baumbach, Dominik Dold, Julian Goltz, Akos F Kungl, Timo C Wunderlich, Andreas Hartel, Eric Müller, Oliver J Breitwieser, Christian Mauch, Mitja Kleider, Andreas Grübl, David Stöckel, Christian Pehle, Arthur Heimbrecht, Philipp Spilger, Gerd Kiene, Vitali Karasenko, Walter Senn, Karlheinz Meier, Johannes Schemmel, and Mihai A Petrovici. 2020. Versatile emulation of spiking neural networks on an accelerated neuromorphic substrate. In *IEEE International Symposium on Circuits and Systems, ISCAS 2020, Sevilla, Spain, May 17-20, 2020*. IEEE. in preparation.
- [4] Ann-Shyn Chiang, Chih-Yung Lin, Chao-Chun Chuang, Hsiu-Ming Chang, Chang-Huain Hsieh, Chang-Wei Yeh, Chi-Tin Shih, Jian-Jheng Wu, Guo-Tzau Wang, Yung-Chang Chen, et al. 2011. Three-dimensional reconstruction of brain-wide wiring networks in *Drosophila* at single-cell resolution. *Current Biology* 21, 1 (2011), 1–11.
- [5] Nicolas Frémaux and Wulfram Gerstner. 2016. Neuromodulated spike-timing-dependent plasticity, and theory of three-factor learning rules. *Frontiers in neural circuits* 9 (2016), 85.
- [6] Nicolas Frémaux, Henning Sprekeler, and Wulfram Gerstner. 2010. Functional requirements for reward-modulated spike-timing-dependent plasticity. *Journal of Neuroscience* 30, 40 (2010), 13326–13337.
- [7] S. Friedmann, J. Schemmel, A. Grübl, A. Hartel, M. Hock, and K. Meier. 2017. Demonstrating Hybrid Learning in a Flexible Neuromorphic Hardware System. *IEEE Transactions on Biomedical Circuits and Systems* 11, 1 (2017), 128–142. <https://doi.org/10.1109/TBCAS.2016.2579164>
- [8] Kirsia Neuser, Tilman Triphan, Markus Mronz, Burkhard Poeck, and Roland Strauß. 2008. Analysis of a spatial orientation memory in *Drosophila*. *Nature* 453, 7199 (2008), 1244.
- [9] Thomas Stone, Barbara Webb, Andrea Adden, Nicolai Ben Weddig, Anna Honkannen, Rachel Templin, William Wcislo, Luca Scimeca, Eric Warrant, and Stanley Heinze. 2017. An anatomically constrained model for path integration in the bee brain. *Current Biology* 27, 20 (2017), 3069–3085.
- [10] Shinya Takemura, Arjun Bharioke, Zhiyuan Lu, Aljoscha Nern, Shiv Vitaladevuni, Patricia K Rivlin, William T Katz, Donald J Olbris, Stephen M Plaza, Philip Winston, et al. 2013. A visual motion detection circuit suggested by *Drosophila* connectomics. *Nature* 500, 7461 (2013), 175.
- [11] Shin-ya Takemura, Yoshinori Aso, Toshihide Hige, Allan Wong, Zhiyuan Lu, C Shan Xu, Patricia K Rivlin, Harald Hess, Ting Zhao, Toufiq Parag, et al. 2017. A connectome of a learning and memory center in the adult *Drosophila* brain. *Elife* 6 (2017), e26975.
- [12] Margaret Wilson. 2002. Six views of embodied cognition. *Psychonomic bulletin & review* 9, 4 (2002), 625–636.
- [13] Timo Wunderlich, Akos Ferenc Kungl, Eric Müller, Andreas Hartel, Yannik Stradmann, Syed Ahmed Aamir, Andreas Grübl, Arthur Heimbrecht, Korbinian Schreiber, David Stöckel, et al. 2019. Demonstrating advantages of neuromorphic computation: a pilot study. *Frontiers in Neuroscience* 13 (2019), 260.