Brain-Inspired Hardware for Artificial Intelligence

Electronic Vision(s)

Timo C. Wunderlich[1], Akos F. Kungl[1], Eric Müller[1], Johannes Schemmel[1], Mihai A. Petrovici[1,2]

1: Kirchhoff Institute for Physics, Heidelberg University, Germany; 2: Department of Physiology, University of Bern, Switzerland

Correspondence: timo.wunderlich@kip.uni-heidelberg.de

Summary

We showcase a mixed-signal neuromorphic chip with an acceleration factor of 1000 relative to biology and novel facilities for synaptic plasticity. A fully on-chip closed–loop reinforcement learning experiment allows us to demonstrate the speed and energy-efficiency of our neuromorphic emulation as compared to a digital simulation on an off-the-shelf CPU. Besides this, we demonstrate that temporal noise can be used as a computational resource and that fixed-pattern noise can be implicitly compensated by learning. For detailed results, we refer to [1].

Background

Physical-model neuromorphic hardware emulates the dynamics of neuron and synapse models using analog electronic circuits.

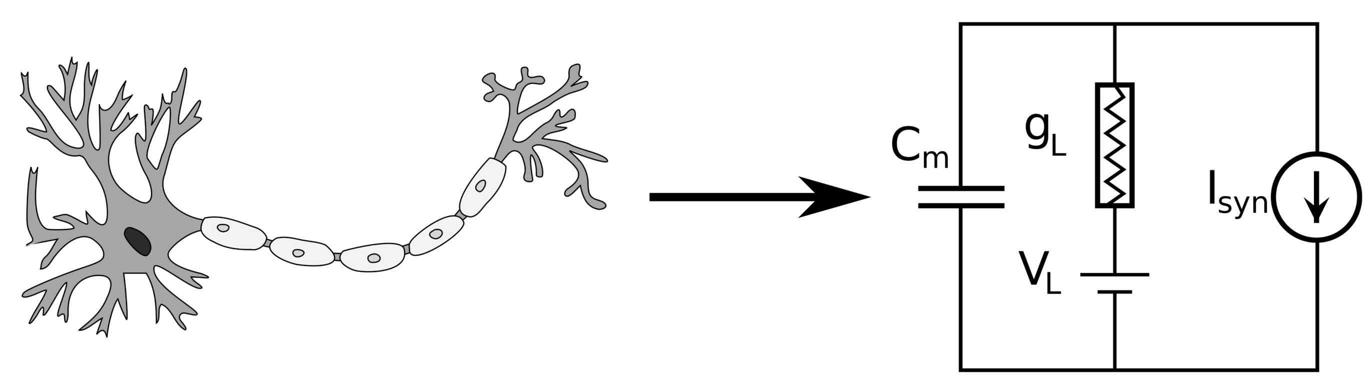

Such neuromorphic chips that contain neural networks made up of electronic neurons and synapses can be produced using CMOS technology.

BrainScaleS 2

BrainScaleS 2 (BSS2) is a neuromorphic platform developed in Heidelberg [2]. It features

- configurable analog neurons based on the AdEx model and synapses with correlation sensors and short-term plasticity,
- a 1000-fold acceleration of neuronal dynamics relative to biology,
- an embedded plasticity processor for flexible synaptic plasticity.

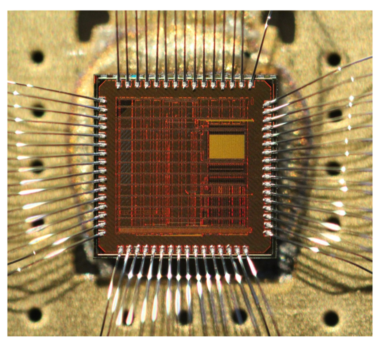

Single-chip prototype (65nm CMOS by TSMC)

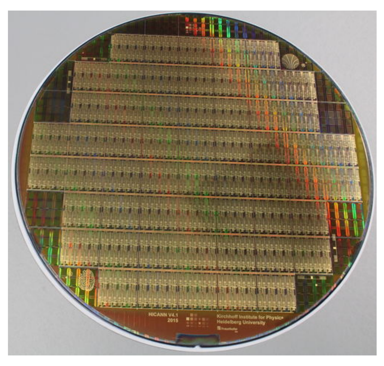

BrainScaleS 1 wafer-scale system [3]

Materials and Methods

We use a single-chip prototype of BrainScaleS 2 with 32 analog leaky integrate-and-fire neurons with 32 synapses each.

The experiment simulates the Pong video game using the embedded processor and uses the on-chip spiking neural network for control and reinforcement learning [1].

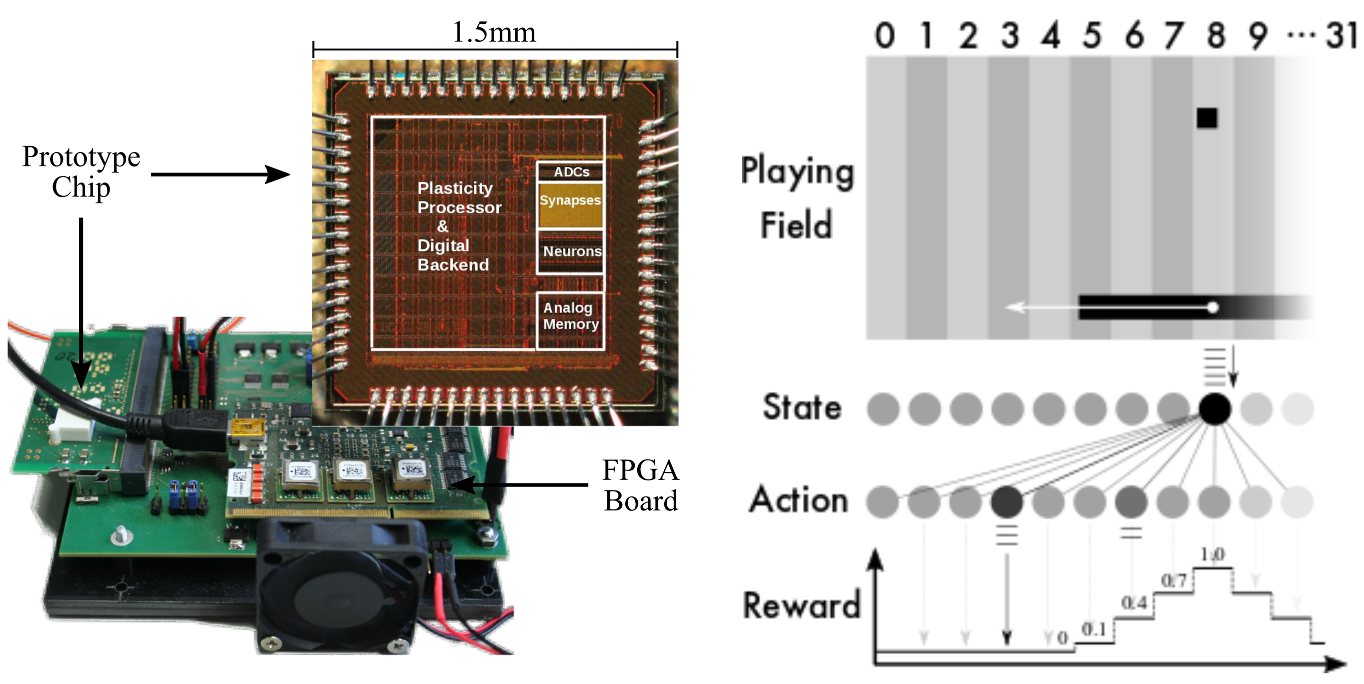

The chip learns to map states to rewarded actions using Reward-modulated Spike-Timing-Dependent Plasticity (R-STDP) [4]:

$$\Delta w_{ij} = \gamma \cdot e_{ij} \cdot (R - \bar{R})$$

Live-On-Tape Demonstration:

bit.ly/pongdemo

Results

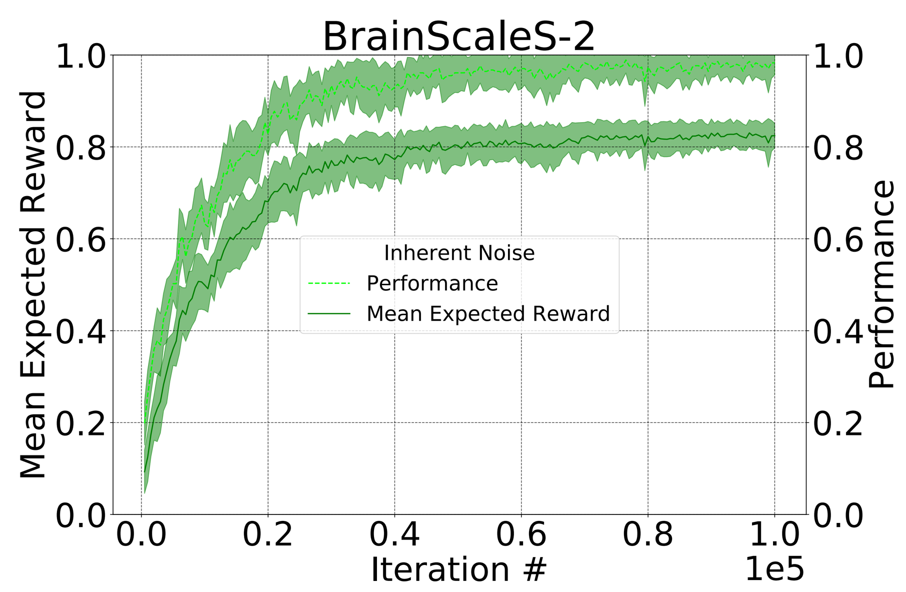

Temporal noise on the neuromorphic chip causes trial-to-trial variability of firing rates, enabling action exploration and hence, learning.

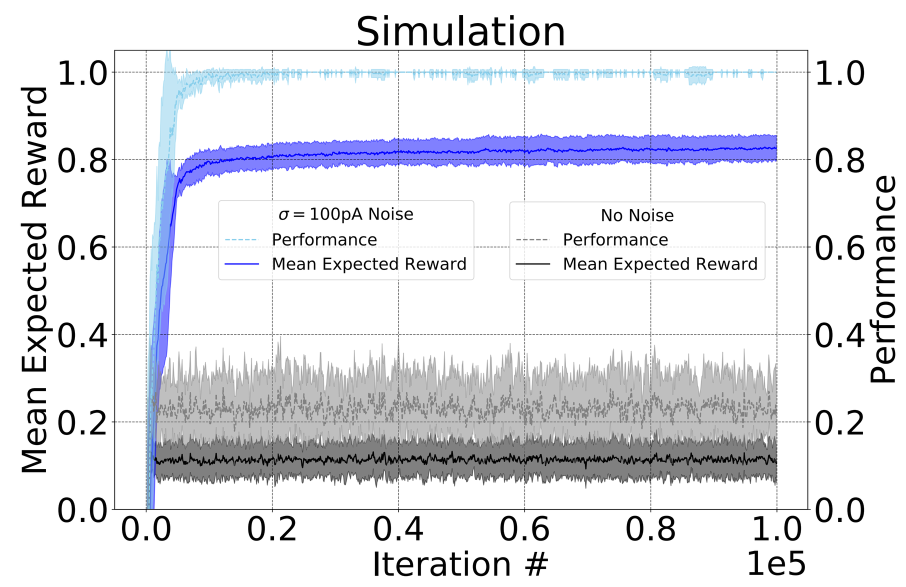

An equivalent software simulation using NEST running on an Intel CPU requires the injection of noise for learning.

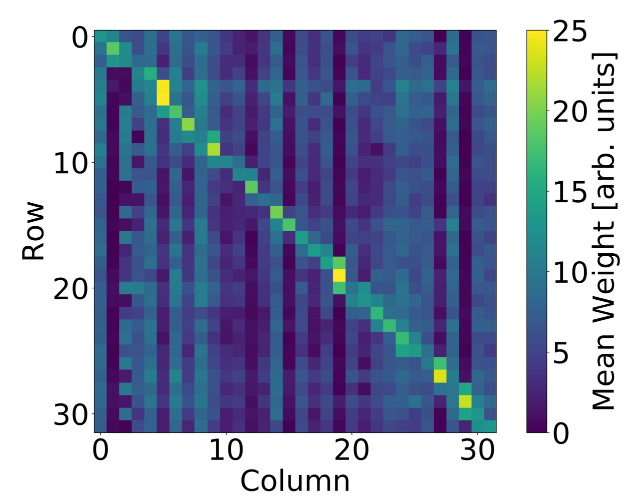

The learning process turns a uniform into a diagonally-dominant weight matrix.

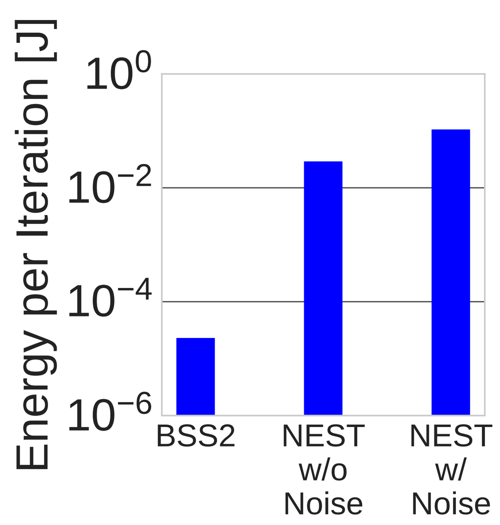

Compared to the digital simulation, BSS2 consumes at least three orders of magnitude less energy per iteration.

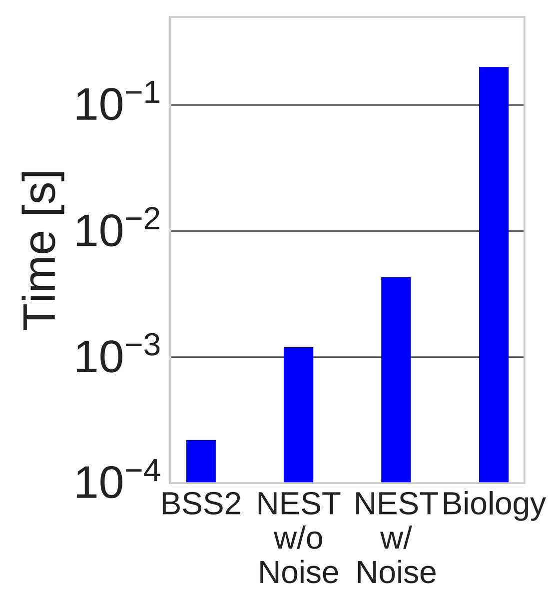

Compared to the digital simulation, BSS2 is at least an order of magnitude faster.

[1] Timo Wunderlich et al. "Demonstrating Advantages of Neuromorphic Computation: A Pilot Study". In: Frontiers in Neuroscience (2019). doi:10.3389/fnins.2019.00260
[2] Simon Friedmann et al. "Demonstrating Hybrid Learning in a Flexible Neuromorphic Hardware System". In: IEEE Transactions on Biomedical Circuits and Systems 11.1 (Feb. 2017), pp. 128–142. doi:10.1109/TBCAS.2016.2579164.
[3] Schemmel et al. "A wafer-scale neuromorphic hardware system for large-scale neural modeling". In: Proceedings of 2010 IEEE International Symposium on Circuits and Systems (2010). doi: 10.1109/ISCAS.2010.5536970.
[4] Nicolas Frémaux et al. "Neuromodulated Spike-Timing-Dependent Plasticity, and Theory of Three-Factor Learning Rules." In: Frontiers in neural circuits 9 (2015), p. 85. issn: 1662-5110. doi: 10.3389/fncir.2015.00085