Internship report

Spike-timing dependent plasticity on BrainScaleS-2

Felix Jaeckle
University of Heidelberg, Electronic Visions Group
Supervisors: Amani Atoui, Jakob Kaiser, Eric Müller

January 2026

Abstract

Each hardware synapse on the BrainScaleS-2 chip has a correlation sensor that detects the temporal correlation of presynaptic and postsynaptic spikes. As a first step towards emulating SNNs with spike-timing dependent plasticity on the BrainScaleS-2 chip, the correlation sensors were characterized. To this end, an experiment was designed that allows a correlation sensor to be stimulated with full control over the pre- and post-spiketimes and its data to be read out. The various parameters were introduced as variables to investigate the behavior of the correlation sensor under arbitrary conditions. The results show the desired behavior with realistic deviation due to hardware implementation. The experiments designed during this internship can be used to test individual correlation sensors for special cases by selecting suitable experiment parameters. However, measuring one correlation sensor is not representative of the behavior of several sensors working in parallel.

Contents

1. Introduction

Spiking neural networks are computational models inspired by biological neuronal dynamics, designed to capture key temporal aspects of neural information processing.

Of particular interest here is neural plasticity, which describes the system's ability to adapt its dynamics and structure in response to internal and external events [1]. It is shown that synaptic plasticity plays a major role in this and is involved in learning and memory processes [3].

The exact relationship between microscopic synaptic properties and macroscopic functional consequences has only been partially researched to date. Phenomenological models in the form of theoretical learning rules are therefore of great importance for the analysis, simulation and emulation of neural plasticity.

In the last decade of the twentieth century, a study demonstrated for the first time the relevance of the relative timing between action potentials in relation to plasticity [8]. The group of learning rules that emerges from this concept is referred to as spike-timing dependent plasticity (STDP) [2]. Since the underlying biological mechanisms in vivo are complex and, in some cases, still the subject of research, it is plausible to investigate the principle in vitro and on spiking neuron models.

The BrainScaleS-2 platform offers the opportunity to emulate spiking neural networks. The aim of the underlying internship is therefore to take the first steps towards STDP on the BrainScaleS-2 platform, especially focusing at the hardware implementation.

1.1 Spike-timing dependent plasticity (STDP)

The principle of spike time dependent plasticity is based on the change in synaptic weights depending on the temporal correlation between pre- and postsynaptic spike behavior. More specifically, it depends on spike pairs and the size of their inter-spike intervals (ISIs). If the presynaptic spike is immediately followed by a postsynaptic spike, this causal correlation should result in a positive weight update (long-term potentiation (LTP)). Conversely, an anticausal correlation (post before pre) should cause a negative weight update (long-term depression (LTD)) [7]. The relationship between the weight change and the amount of the inter-spike interval is modeled by using an exponential function:

$$\Delta w = \begin{cases} A_+ \exp\left(-\frac{\Delta t}{\tau_+}\right) & \text{for } \Delta t > 0 \\ -A_- \exp\left(-\frac{\Delta t}{\tau_+}\right) & \text{for } \Delta t \leq 0\,. \end{cases} \tag{1}$$

where Δw is the weight update, $A_\pm$ some constants, $\tau_\pm$ the time constants and Δt the ISI length. The range of inter-spike intervals considered by the learning rule defines the so called learning window.
The general form of STDP described by Equation 1 can be extended and varied depending on the model, so that the weights do not diverge or only certain pre-post-events are considered. The exact form of STDP enabled by the BrainScaleS-2 chip is described in subsection 1.4.
The standard STDP curve, which has become the icon of theoretical STDP research [3], can be seen in Figure 1.

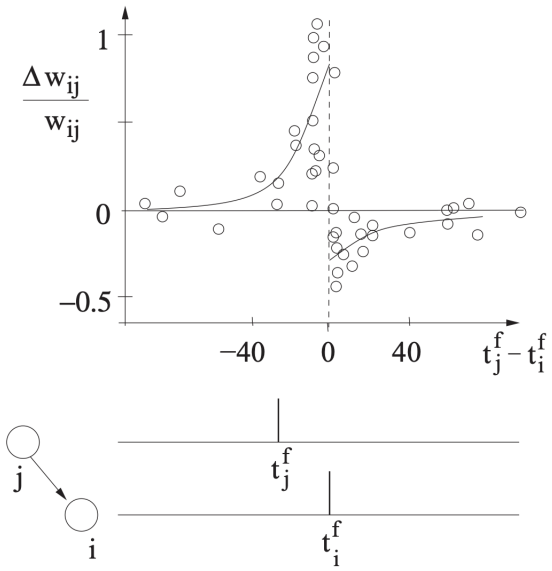

Figure 1: Experimentally measured weight changes (circles) as function of the differences between pre- an postsynaptic spike timing $t_j^f - t_i^f$ in milliseconds overlayed on a schematic two-phase learning window (solid line). Data point redrawn after experiments of Bi and Poo (1998)[5], (figure taken from: [7]).

1.2 BrainScaleS-2

The BrainScaleS-2 (BSS-2) platform enables fast and energy-efficient emulation of spiking neural networks (SNNs) on analog hardware. The BSS-2 chip is designed as an accelerated neuromorphic system and is approximately 1000 times faster than biological timescales. This means that one day in biological time can be emulated within two minutes. The BSS-2 chip is a mixed signal chip that has 512 neurons (256 per hemisphere) and two crossbars of 256x128 synapses [4]. Each of these synapses, in turn, has a correlation sensor circuit that records the temporal correlation of pre- and post-synaptic events and processes them analogously according to the principle described in subsection 1.1. Differences and similarities between the general model and the implemented model type on the BSS-2 platform will be discussed in subsection 1.4. The column-wise analog-to-digital converter (CADC) allows all correlation sensors to be read out in parallel, and the digitized values to be passed on to the plasticity processing unit (PPU) with an 8-bit resolution. The PPU is part of the plasticity sub-system calculates the new synaptic weights so that the 6-bit static random-access memory (SRAM) weights can be updated during the experiment [6]. The frequency with which the PPU updates the weights based on a plasticity rule is determined by a PPU timer.

1.3 Correlation Sensors

The schematic structure of the correlation sensors can be seen in Figure 2. The timing control receives pre and post signals and uses them to classify the internal timing as causal or anti-causal events. This is done based on the last input. If the last spike was a post spike, the event will charge the causal capacitor and if the last spike was a pre spike, the event will charge the anti-causal capacitor. A time-to-voltage conversion generates a voltage that represents the amount of the last inter-spike interval. The calibration of the time-to-voltage conversion effectively corresponds to the scaling of the inter-spike interval Δt and can thus be used to modulate the time constant τ.

The voltage is scaled by the storage gain control and used as an argument in an exponential function. The calibration of the storage gain control effectively corresponds to the scaling of the exponential function and can thus be used to modulate the maximum weight change A. Depending on the causal or anti-causal classification, the resulting voltage is added to either the causal storage or the anti-causal storage. Causal storage and anti-causal storage can then be read out in parallel for all synapses using the CADC and passed on to the PPU for calculation of the underlying learning rule. The time to voltage control and the storage gain control each have a 2-bit digital calibration. Thus, the correlation has a total of 4 bits SRAM for calibration. Additionally, the time constant for the time to voltage conversion and the storage gain control can be globally set for one row of synapses by a control voltage.

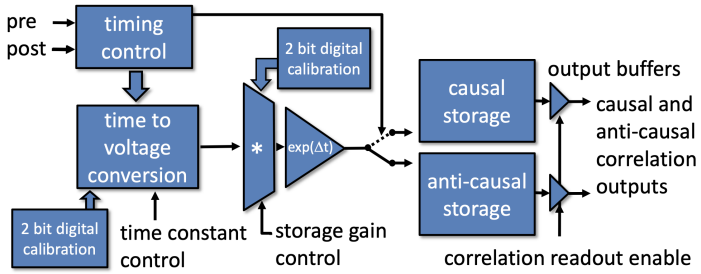

Figure 2: Block diagram of the correlation sensor circuit (Reference: [6]).

The exact circuit is described in detail in [4]. To understand how it works, it is sufficient to consider only the key dynamics. There is a causal capacitor and an anti-causal capacitor. Figure 3 shows how the voltage traces of the causal and anti-causal capacitors behave according to pre and post spikes. When a pre spike arrives, the causal capacitor is discharged. Until the post-spike arrives, the capacitor is charged with a constant current. The voltage of the capacitor at the end of the inter-spike intervals is therefore linearly related to the duration of the interval. When a post-spike arrives, exactly the same discharge/charge process occurs in the anti-causal capacitor. Two or more pre or post spikes in a row would only restart the discharge/charge process without changing the storage capacitors [6]. The causal or anti-causal capacitor is then connected with a transfer capacitor, so that a voltage is established that is proportional to the previous voltage on the causal or anti-causal capacitor.

To apply the exponential dynamics to the voltage on the transfer capacitor, a metal–oxide–semiconductor field-effect transistor (MOSFET) is used, which can be interpreted as a "voltage-dependent current source". The MOSFET effectively relates to an exponential function, that takes the voltage on the transfer capacitor as an argument and charges the storage capacitor by a charge that corresponds to the exponential output.

The accumulated and exponential-weighted ISIs are stored by the storage capacitors. The storage capacitors can then be readout via the CADC or get reset.

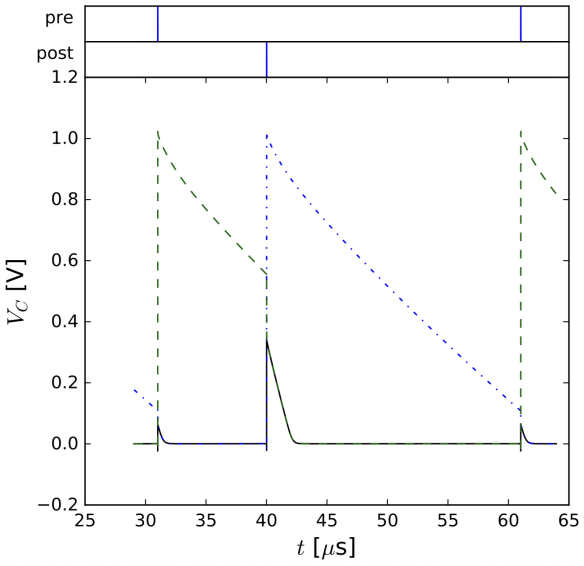

Figure 3: Simulation showing the charging of C_{transfer} (solid trace) after a post pulse (at $t = 40\mu s$). The dashed trace shows the voltage on C_{causal} and the dashed-dotted trace the voltage on $C_{\text{anti-causal}}$. During the charging process, the appropriate C_{storage} capacitor is discharged. (Reference: [4]).

1.4 STDP on BrainScales2

Due to the physical circuits on the BSS-2 chip, its functional flexibility is limited, e.g. the chip will not support arbitrary time constants or pairing schemes. Thus, the variability of the models that can be emulated with the BSS-2 chip is limited. It is therefore important to consider for which specific type of STDP the hardware of the BSS-2 chip was designed. This section should reconstruct the theoretical STDP model on the BSS-2 chip, referring to the general theoretical model from subsection 1.1. It is important to mention that real experiment results will also differ from the BSS-2 specific theoretical model due to fixed pattern noise and other effects of the physical circuits.

The local 4-bit calibration of the time constant and maximum weight change applies to both causal and anti-causal events, which means that at the hardware level, $\tau_- = \tau_+ = \tau$ and $A_- = A_+ = A$ apply. Because weight changes on the BSS-2 chip are stored at different addresses for causal and anti-causal events, these cases are described separately. Using the operating principle from [6], the following equation applies as the BSS-2 implemented variant of Equation 1:

$$c_{\text{causal}} = \begin{cases} -A \exp\left(-\dfrac{\Delta t}{\tau}\right), & \text{for } \Delta t > 0 \\ 0, & \text{for } \Delta t < 0 \end{cases} \tag{2}$$

$$c_{\text{anti-causal}} = \begin{cases} 0, & \text{for } \Delta t > 0 \\ -A \exp\left(-\dfrac{\Delta t}{\tau}\right), & \text{for } \Delta t < 0 \end{cases} \tag{3}$$

where c_{causal} and $c_{\text{anti-causal}}$ are the correlations for causal and anti-causal events. Sticking with the Notation of Δw would be misleading here because the correlation data is not necessarily the weight change. The correlation data is only passed to the PPU, where the weight change

can be calculated based on the given plasticity rule.

It is not only important how the correlation data depend on the interval lengths, but also which intervals are considered. For this, there are several pairing schemes. Three different schemes are shown in Figure 4.

The operating principle of the correlation sensor was described in subsection 1.3, and the voltages of the causal and anti-causal capacitors are shown in Figure 3. This shows that the BSS-2 chip implements next-neighbor pairing, which corresponds to case C in Figure 4. This pairing scheme has the advantage that the interval lengths are more limited due to the reduced combinatorics, which makes the cut-off of the learning window, as provided for in the theoretical model, superfluous on the BSS-2 chip.

There are several effects that are not intended in the theoretical model, but can occur in analog hardware due to the real circuit structure.

- One such effect is the drift effect that every capacitor has. This could cause a loss of information because the storage capacitors drift significantly before they are read out.
- Also relevant are the saturation effects of the capacitors, which are expected especially during long readout periods or high activity.
- Last but not least, there are time windows between the reading and resetting of the correlation sensors during which the correlation sensor is blind. This loss of information can be minimized by resetting directly after reading in the PPU code. Since the reset only applies to causal and anticausal storage, there is no loss of information for spike pairs that are separated by the end of a PPU period. It is assumed that the effect can be neglected when minimizing the time window in the PPU code.

Whether and to what extent these effects actually play a significant role needs to be investigated. At least in the case of systematic deviations from the expected behavior, it might be useful to take these effects into account.

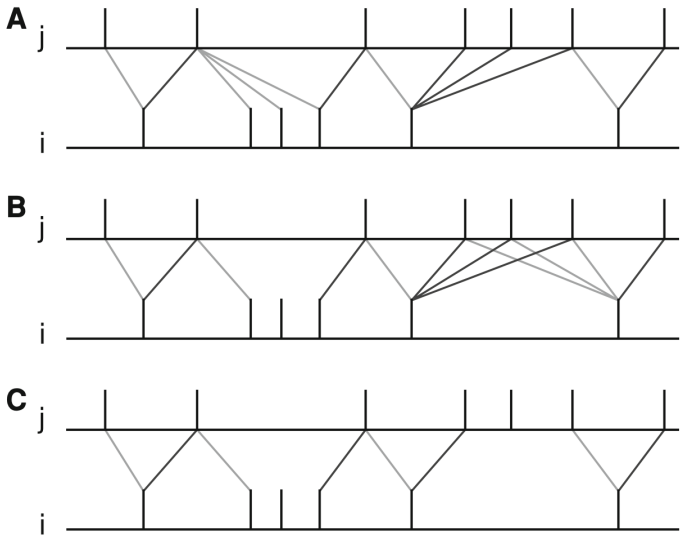

Figure 4: Examples of nearest neighbor spike pairing schemes for a presynaptic neuron j and a postsynaptic neuron i. In each case, the dark gray indicate which pairings contribute toward depression of a synapse, and light gray indicate which pairings contribute toward potentiation. **A** Symmetric interpretation: each presynaptic spike is paired with the last postsynaptic spike, and each postsynaptic spike is paired with the last presynaptic spike. **B** Presynaptic centered interpretation: each presynaptic spike is paired with the last postsynaptic spike and the next postsynaptic spike. **C** Reduced symmetric interpretation: as in B but only for immediate pairings (Reference: [3]).

2. Methods

In order to characterize the correlation sensors, an experiment is designed that records the correlation data for different ISIs so that the values can be compared with the theoretical model from subsection 1.1. Ideally, this data should reproduce the functions from Equation 2 and Equation 3. In addition, the drift effect of the correlation sensors is investigated qualitatively by reading out the correlation data after different drift times.

2.1 Experiment setup

A schematic graphic of the experiment setup is shown in Figure 5. The setup consists of two neurons: a presynaptic neuron and a postsynaptic neuron. The spike times of the presynaptic and postsynaptic neurons are each controlled via a spike source with a bypass projection. The bypass projection causes a strong pulse to be sent to the bypass neuron for each spike from the spike source, so that the pre- and post neuron are spiking without accumulation. In order to capture the correlation data, a "main projection" is also created from the pre-neuron to the post-neuron. Since the spike times are to be controlled only by the spike source arrays, the weight of the synapse of the main projection is set to zero. The synapse circuit of the main projection contains the correlation sensor, through which the correlation data is read out during the experiment. The correlation data is read out via the PPU. The PPU is equipped with a timer that can be used to define the PPU period, which determines the frequency at which the PPU code is executed by the PPU. To control the PPU, a plasticity rule is created for the synapse of the main projection. The plasticity rule is implemented in a PPU kernel in which a C++ code is given to the PPU as a string. The code instructs the PPU to measure a baseline within the first PPU period. In the later periods, the PPU reads-out the correlation sensors, subtracts the baselines, stores the data, and resets the correlation sensors for the next PPU period.

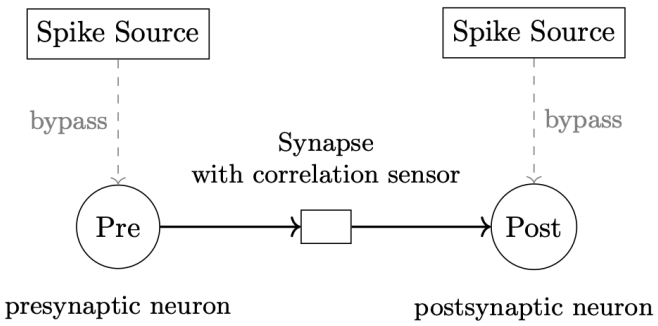

Figure 5: Experiment setup consisting a pre and post neuron connected with a synapse including the correlation sensor. The pre and post neuron are connected to the spikesources via a bypass projection in order to fully control the spike times.

2.2 Calibration

The correlation sensors are calibrated using an automated Calix calibration, which is performed every week. The time constant and amplitude can be calibrated globally for one synapse row containing 128 synapses/correlation sensors. To reduce fixed pattern noise, the two parameters are additionally locally calibrated via the 4-bit SRAM. For the used calix calibration the time constant is calibrated to $\tau = 30\mu s$ and the amplitude to $A = 1.5$ [CADC a.u.].

The choice of these parameters also seems relatively plausible from a biological point of view, as most experiments are done with repetitions of 50–60 pairs of spikes, whereas a single pair has no effect [3][9].

2.3 Spike protocols

Because the correlation sensors are reset after each PPU period, the spike pairs should be arranged within the PPU period so that each period generates one data point.

First, a spike protocol is created to record the learning window. A schematic sketch of the spike protocol can be seen in Figure 6. In each PPU period, the correlation sensor is probed by spike pairs with a variable interval length Δt (Figure 6a). The correlation data read out after the period can then be plotted against Δt to create the desired learning window plot. Due to the small amplitude of 1.5, it is necessary to inject several spike pairs for one data point (Figure 6b). To prevent detecting a correlation between the spike pairs, an inter-spike interval is introduced so that the correlation between the spike pairs is negligible. The time interval between the last spike of the spike pair sequence and the reading of the correlation sensor should be fixed at a value, which is defined as the "drift gap". To prevent the PPU from taking longer to execute the code than the PPU period allows, a "PPU-gap" is also introduced. The PPU-gap determines the minimum time interval in each PPU period before the arrival of the spike pairs and thus effectively extends the PPU period without changing the relevant dynamics. A function takes the parameters: `num_pairs`, `num_datapoints`, `delta_max`, `inter_pair_interval`, `ppu_gap` and `drift_gap` (Figure 6c). As output, it generates the spike trains, which are then injected into the corresponding neurons via the spike sources and the bypass.

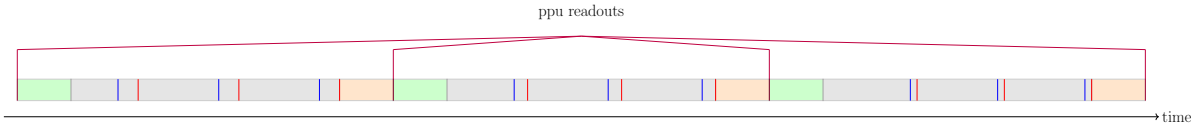

(a) Zoomed-out view of the spike protocol across multiple PPU periods. For each datapoint, an additional PPU period is appended with adjusted intra-pair interval Δt. After finishing for causal pairs the same appends for anti-causal measurements, within the same spiketrain.

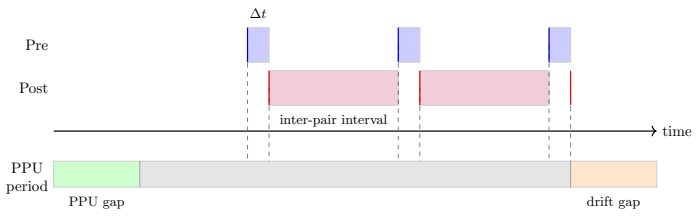

(b) Spike protocol used to generate the learning window measurement. The schematic shows a single PPU period with multiple spike pairs with corresponding Δt. The spikepairs are seperated by a long inter-pair interval to avoid correlations between spikepairs.

```python
def complete_spiketrain(
        num_pairs: int,
        num_datapoints: int,
        delta_max: float,
        inter_pair_interval: float,
        ppu_gap: float,
        drift_gap: float):
        ...
        return pre_spikes, post_spikes
```

(c) Corresponding code structure implementing the spike protocol and the systematic temporal adjusting of intra-pair intervals across PPU periods.

Figure 6: Spike protocol and implementation used for learning window characterization. (a) Temporal extension of the protocol over multiple PPU periods to generate a full learning window. (b) Schematic representation of a single PPU period with spikepair constraints. (c) Code-level implementation of the protocol.

Another spike protocol is created to investigate the drift effect mentioned in subsection 1.4. A schematic sketch of the spike protocol can be seen in Figure 7. The drift gap Δt_{drift} is not constant for this and increases with each PPU period (Figure 7a). The correlation data read out after each period in the experiment can then be plotted against Δt_{drift} to create the desired drift plot. This time, two identical equidistant spike trains are shifted into each other so that a pre-post and post-pre interval of equal length always follows each other alternately (Figure 7b). Here, too, several spike pairs must be used due to the small amplitude. A function takes the following parameters as input: `num_pairs`, `num_datapoints`, `intra_pair_interval`, `ppu_gap` `min_drift_gap`, and `max_drift_gap` (Figure 7c). As output, it generates the spike trains again, which are then injected into the corresponding neurons via the spike sources.

(a) Zoomed-out view of the spike protocol across multiple PPU periods. For each datapoint, an additional PPU period is appended with shifted spike trains until the first pre-synaptic spike reaches the maximum drift gap boundary.

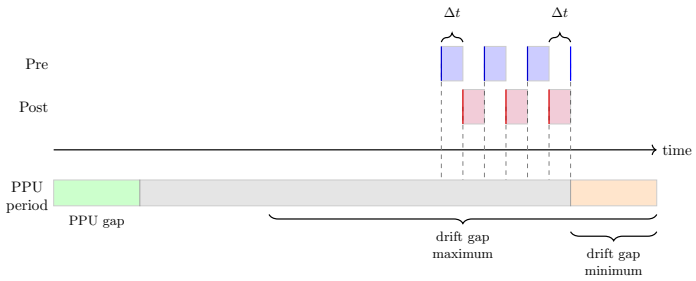

(b) Spike protocol used to generate the drift measurement. The schematic shows a single PPU period with pre- and post-synaptic spike pairs and the associated drift gap boundaries.

```python
def complete_spiketrain(
        num_pairs: int,
        num_datapoints: int,
        intra_pair_interval: float,
        ppu_gap: float,
        min_drift_gap: float,
        max_drift_gap: float):
        ...
        return pre_spikes, post_spikes
```

(c) Corresponding code structure implementing the spike protocol and the systematic temporal shifting across PPU periods.

Figure 7: Spike protocol and implementation used for drift characterization. (a) Temporal extension of the protocol over multiple PPU periods to generate a full drift curve. (b) Schematic representation of a single PPU period with drift gap constraints. (c) Code-level implementation of the protocol.

3. Results

The experiment measures one correlation sensor at a time. The following results show an exemplary chip. Other correlation sensors that have been used and have shown similar behavior are: hxcube1fpga0chip76_1, hxcube1fpga3chip83_1, hxcube6fpga0chip65_1, hxcube6fpga3chip66_1, hxcube8fpga0chip93_1, hxcube14fpga3chip63_1.

3.1 Test of the learning window

The learning window is plotted here using the spike protocol from Figure 6 as an example for the W74F0 chip (identifier: hxcube14fpga0chip90_1). The parameters are stored in a configuration file (config.yaml). The following parameters were used for this experiment:

```
num_pairs: 40
num_datapoints: 20
delta_max: 0.05
inter_pair_interval: 0.5
ppu_gap: 2
drift_gap: 0.01
```

The result of the plot can be seen in Figure 8. In addition, a fit was performed for the exponential ranges using the theoretical functions from Equation 2 and Equation 3 to check the deviation from the calibration.

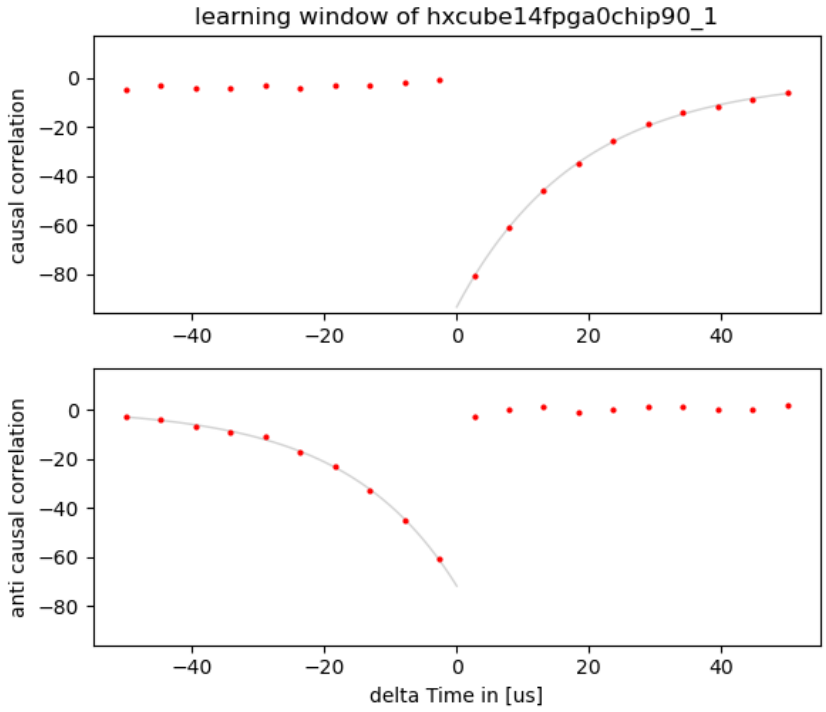

Figure 8: Causal and anti-causal correlation functions derived from the sensor readout as a function of Δt. The upper panel shows the causal correlation, and the lower panel the anti-causal correlation, both plotted versus the delta Time (in milliseconds). Discrete data points (red markers) represent the measured correlations, while the solid light-gray lines indicate the corresponding fitted model functions. The fitted parameters (time constant τ and amplitude) for both causal and anti-causal correlations are reported within the figure.

3.2 Test of the drift effect

The drift is plotted using the spike protocol from Figure 7. The same chip is used, and the code is executed after the learning window experiment to base the configuration values on these results. The parameters are again stored in a configuration file. The following parameters were used for this experiment:

```
num_pairs : 40
num_datapoints : 20
intra_pair_interval : 0.005
ppu_gap : 2
min_drift_gap : 0.01
max_drift_gap : 1000
```

The result can be seen in Figure 9. Although the drift lasts up to 1000 ms, almost no drift can be detected here.

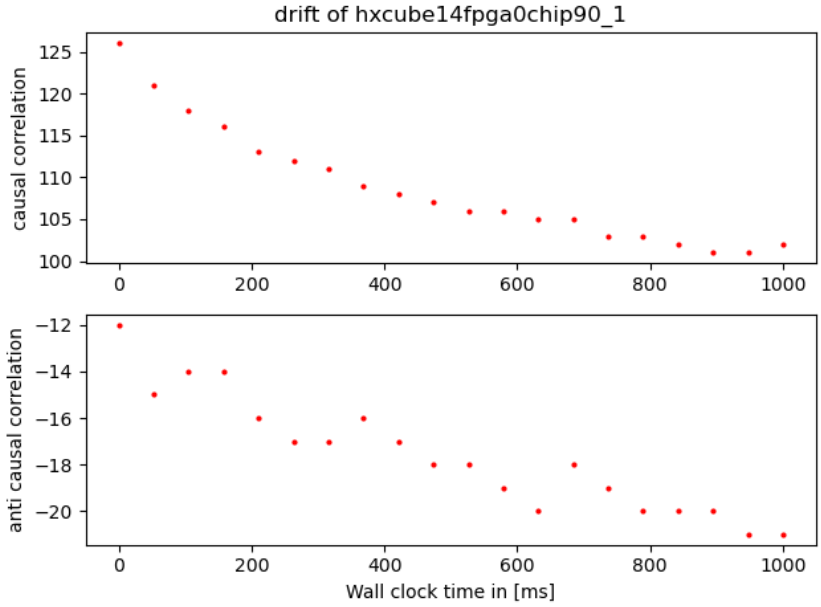

Figure 9: Temporal drift of causal and anti-causal correlations over wall-clock time. The upper panel displays the evolution of the causal correlation, and the lower panel the anti-causal correlation, plotted as a function of wall-clock time (in milliseconds). Each marker corresponds to a single measurement point, illustrating the long-term stability and drift behavior.

3.3 Approximations for suitable plasticity parameters

In order for the STDP learning rule to take effect in later experiments, the input spikes must lie within a suitable frequency range. A rough calculation is made to estimate a suitable firing rate for this chip. For this purpose, a simplified mean ISI length $\langle \text{ISI} \rangle = \bar{t}$ is assumed for uncorrelated pre and post spikes. This refers only to intervals that are considered by the next-neighbor scheme. For reasons of symmetry, the value is the same for causal and anti-causal ISIs. This value should be close to the half-way time, or rather below it, depending on the variance of the intervals. The half-way times $T_{\frac{1}{2}}$ are approximated from the fit parameters of the time constants:

$$\begin{aligned} T_{\frac{1}{2},\text{causal}} &= \frac{\tau_{\text{causal fit}}}{\ln(2)} \\ T_{\frac{1}{2},\text{anti-causal}} &= \frac{\tau_{\text{anti-causal fit}}}{\ln(2)} \,. \end{aligned} \tag{4}$$

For an appropriate ratio of the PPU period T_{PPU} to the mean ISI length $\bar{t}$, at least the following inequality should apply in order to avoid saturation with the calibration used:

$$\frac{T_{\text{PPU}}}{\bar{t}} \cdot A_{\text{calibration}} \exp\left(-\frac{\bar{t}}{\tau_{\text{calibration}}}\right) < 255 \tag{5}$$

Because an upward deviation has a much greater effect due to the exponential function, the value should be well below the CADC maximum of 255.

To avoid saturation due to the hardware capacitors, the drift experiment described above can be used to test whether the capacitors are in saturation. For this, the number of spike pairs in the drift protocol can be varied. If the correlation data does not increase despite increasing the spike pairs, the capacitor is in saturation.

The saturation problems mentioned can be mitigated in both cases by appropriate calibration.

4. Discussion

The plots for the chip shown here demonstrate that the correlation sensors function very well, and repetitions with other chips confirm this. While this is not surprising, it is nevertheless very important for upcoming STDP experiments that build on it. Credit for this achievement goes to the developers of the correlation sensors (S. Friedmann, J. Schemmel et al., 2017) [6].

The plots shown in Figure 8 and Figure 9 characterize only one synapse on a chip and are therefore not representative. In fact, the W60F0 chip shown as an example in this report demonstrates very poor agreement with the expected calibration values compared to other chips. Normally, both the amplitudes and the time constants are closer to the values adjusted by the calibration.

The strength of the experiments designed here lies in the variability of the parameters, which allows the correlation sensor to be tested under experimental conditions. Other experiments already exist for the averaged characterization of all correlation sensors on a chip.

5. Outlook

Since the correlation sensors on the BSS-2 chip work well, as shown, the logical next step is to use them in a plasticity experiment and to solve a task designed for this purpose based on a STDP learning rule. The STDP learning rule can still be modified for future experiments, e.g. so that the weights are limited by upper and lower bounds. In addition, the STDP learning rule can also be combined with other learning rules, e.g. from the areas of structural plasticity or homeostatic plasticity. When implementing such an experiment, it may be helpful to refer to the spike protocols designed in subsection 2.3.

After this project internship, I am highly motivated to devote myself to the tasks at hand and am therefore delighted to have the opportunity to do so in my upcoming bachelor's thesis.

References

[1] Mateos-Aparicio P, Rodríguez-Moreno A (2019). *The Impact of Studying Brain Plasticity.* Front Cell Neurosci. https://pmc.ncbi.nlm.nih.gov/articles/PMC6400842/

[2] Markram H, Gerstner W, Sjöström PJ (2011). *A history of spike-timing-dependent plasticity.* Front Cell Neurosci. https://pmc.ncbi.nlm.nih.gov/articles/PMC3187646/#sec12

[3] Morrison A, Diesmann M, Gerstner W (2008). *Phenomenological models of synaptic plasticity based on spike timing.* Biol Cybern. https://pmc.ncbi.nlm.nih.gov/articles/PMC2799003/

[4] Pehle C. et al. (2022). *The brainscales-2 accelerated neuromorphic system with hybrid plasticity.* Frontiers in Neuroscience. https://www.frontiersin.org/journals/neuroscience/articles/10.3389/fnins.2022.795876/full

[5] Bi GQ, Poo MM. (1998). *Synaptic modifications in cultured hippocampal neurons: dependence on spike timing, synaptic strength, and postsynaptic cell type.* J Neurosci. https://www.jneurosci.org/content/18/24/10464

[6] Friedmann S, Schemmel J, Gruebl A, Hartel A, Hock M, Meier K (2016). *Demonstrating Hybrid Learning in a Flexible Neuromorphic Hardware System.* arXiv. https://arxiv.org/pdf/1604.05080

[7] Gerstner W, Kistler WM (2002). *Spiking Neuron Models: Single Neurons, Populations, Plasticity.* Cambridge University Press. https://lcnwww.epfl.ch/gerstner/PUBLICATIONS/SpikingNeuronM-extracts.pdf

[8] Markram H, Lübke J, Frotscher M, Sakmann B (1997). *Regulation of synaptic efficacy by coincidence of postsynaptic APs and EPSPs.* Science. https://pubmed.ncbi.nlm.nih.gov/8985014/

[9] Sjöström PJ, Turrigiano GG, Nelson SB (2001). *Rate, timing, and cooperativity jointly determine cortical synaptic plasticity.* Neuron. https://pubmed.ncbi.nlm.nih.gov/11754844/

List of Acronyms

BSS-2 BrainScaleS-2. 2, 4, 5, 13

CADC column-wise analog-to-digital converter. 2, 3, 12

ISI inter-spike interval. 1, 3, 7, 11, 12

LTD long-term depression. 1

LTP long-term potentiation. 1

MOSFET metal–oxide–semiconductor field-effect transistor. 3

PPU plasticity processing unit. 2–5, 7–9, 12

SNN spiking neural network. 2

SRAM static random-access memory. 2, 3, 7

STDP spike-timing dependent plasticity. 1, 4, 11, 13