Department of Physics and Astronomy
University of Heidelberg

Master Thesis in Physics
submitted by
Laura Kriener
born in Offenburg (Germany)
March 2017

Characterization of Single-Neuron Dynamics in the Development of Neuromorphic Hardware

This Master Thesis has been carried out by
Laura Kriener
at the
Kirchhoff-Institute for Physics
Heidelberg University
under the supervision of
Prof. Dr. Karlheinz Meier
and
Prof. Dr. Michael Hausmann

Abstract

In the development of accelerated analog neuron circuits, a trade-off between energy efficiency and area consumption on the one hand and precision and configurability of the emulated neuron model on the other hand has to be made. It is therefore essential to verify that the designed circuits capture the most relevant characteristics of their ideal mathematical description.
To this end, we present a simulation-based, application-oriented characterization methodology for analog neuron circuit dynamics which was applied during the design phase of the HICANN DLS 3 prototype chip. As high-level test cases for the neuron circuits we choose three different biologically inspired single-neuron experiments, each focused on a different aspect of neuron functionality. The application of these test cases in transistor-level simulations of the neuron circuits is preceded by a detailed verification and characterization which is focused on the practical applicability and usability of the neuron circuits. The combination of high- and low-level investigations of the full neuron circuit yielded valuable information ensuring the correct implementation during the development as well as a set of improvements guiding the design of future circuit generations.

Zusammenfassung

Während der Entwicklung von beschleunigten analogen Neuronschaltungen erfolgt eine Gratwanderung zwischen Energieeffizienz und Flächenverbrauch einerseits und Präzision und Konfigurierbarkeit des Neurons andererseits. Daher ist es unumgänglich sicherzustellen, dass die entwickelten Schaltungen die wichtigsten Charakteristika ihres mathematischen Modells abbilden. Zu diesem Zweck präsentieren wir eine simulations-basierte, anwendungsorientierte Methodik um analoge Neuronschaltungen zu charakterisieren, die während der Entwicklungsphase des HICANN DLS 3 Prototyp Chips angewandt wurde. Als umfassende Tests für die Neuronschaltung verwenden wir drei biologisch inspirierte Einzelneuron-Experimente. Jedes testet dabei einen anderen Aspekt der Neuronfunktionalität. Ein detaillierter, auf praktische Anwendbarkeit fokussierter Verifikations- und Charakterisierungsprozess geht der Anwendung dieser Experimente in Schaltungssimulationen voraus. Die Kombination aus umfassenden und spezialisierten Untersuchungen des gesamten Neuronschaltkreises lieferte wertvolle Informationen, welche sowohl die korrekte Implementation des Neurons in der Designphase ermöglichten als auch zu Verbesserungsvorschlägen für zukünftige Chipgenerationen führten.

Contents

1 Introduction

In March 2016 the neural network AlphaGo beat Lee Sedol, who ranked fourth in the international ranking list, in four out of five matches (DeepMind web page, 2017; BayesElo-rating, 2017). Before the appearance of AlphaGo a computer program mastering the game of Go was thought to be at least a decade away (Silver et al., 2016). Impressive results such as this demonstrate the potential of artificial neural networks.

However, although these networks are called *neural* networks, they are heavily simplified compared to their biological counterparts. Understanding the mechanisms that allow the brain to perform and especially learn tasks that still pose a challenge for artificial neural networks could lead to even more powerful computing technologies.

The field of computational neuroscience tries to obtain an understanding of the mechanisms in the brain by performing simulations on different levels of detail and abstraction. To this end neurons and their interactions are described by mathematical equations that can be solved numerically. Typically, very detailed models that capture a multitude of biological features, such as the different ion channels governing the neuron dynamics, are only used for small-scale or even single neuron simulations. For the investigation of large networks, these models are too computationally intensive and therefore more abstract and thus simple neuron models are used. However, even this becomes unfeasible in terms of power consumption and simulation time, for increasing network sizes or studies of neuron behavior over longer time scales than a few seconds in biology. Using the full K supercomputer, simulating a network of 1.86×10^9 neurons and 1.1×10^{13} synapses for one biological second required forty minutes simulation time (Kunkel et al., 2014).

The approach of accelerated, analog neuromorphic hardware can provide an alternative to numerical simulations (Schemmel et al., 2010). In analog neuromorpic hardware, neurons and synapses consist of electrical circuits that are designed to emulate the dynamics of mathematical neuron models. In contrast to a numerical simulation the states of the dynamic variables in those physical models are not calculated at discrete time steps but they evolve in continuous time as they are represented by electrical quantities in the circuits. The membrane potential of a neuron, for example, is a physical voltage which is stored on a capacitor in the neuron circuit. By designing the size of the whole neuron circuit to be approximately $200 \times 12\,\mu\text{m}$ (Aamir et al., 2016), the neuron and synapse time constants are reduced to micro seconds in comparison to milli seconds in biology, which results in an acceleration factor compared to biology of approximately one thousand.

This acceleration, however, comes at a cost: in the design of the neuromorphic neuron and synapse circuits we sacrifice configurability and precision of the implemented mathematical neuron model for energy efficiency and emulation speed. Therefore, it must be ensured that in spite of this trade-off the essential features of the emulated neuron model are still captured.

In this thesis we describe a simulation-based, application-oriented characterization process for the neuron circuit dynamics which was applied in the design phase of the HICANN DLS 3 prototype chip developed in the Electronic Vision(s) Group at the Kirchhoff-Institute for Physics in Heidelberg. The characterization is based on three different biologically

inspired single neuron experiments that are used as high-level test cases for the neuron circuit: we present an extension to the single neuron modeling competition by Jolivet et al. (2008b) making it suitable for the application to analog neuromorphic hardware. As a second test case we choose the firing patterns exhibited by an adaptive exponential integrate-and-fire neuron (Naud et al., 2008) and finally use the backpropagation-activated calcium spike firing mechanism measured described by Larkum et al. (1999) as test for the newly implemented multi-compartment and plateau potential features on the chip. The execution of these high-level experiments is preceded by a low-level verification of the general functionality and an application-focused characterization of the neuron circuits.
As these evaluations of the neuron circuits take place during the design phase of the chip, all investigations are performed in simulation. We use a combination of a fast, ideal neuron simulation in the NEST simulator (Gewaltig and Diesmann, 2007) and detailed transistor-level simulations of the circuit.
The applied combination of high- and low-level investigations of the full neuron circuit yields valuable information such as the detection of errors in the circuits during the design process and a set of improvement suggestions for future chip generations.

2 Modeling Neurons

2.1 Mathematical Neuron Models

2.1.1 LIF and AdEx

A simple and widely used mathematical neuron model is the leaky integrate-and-fire (LIF) neuron (Dayan and Abbott, 2001). The membrane voltage of a LIF neuron is described by

$$C_{\mathrm{m}}\frac{dV}{dt} = -g_{\mathrm{L}} \cdot (V - E_{\mathrm{L}}) + I_{\mathrm{syn}} + I_{\mathrm{ext}} \tag{2.1}$$

where C_{m} describes the membrane capacitance, I_{ext} is an external- and I_{syn} the synaptic current onto the membrane capacitance, g_{L} the leakage conductance and E_{L} the leakage potential. The membrane time constant τ_{m} can be calculated as $\tau_{\mathrm{m}} = \frac{C_{\mathrm{m}}}{g_{\mathrm{L}}}$. If the membrane voltage crosses a threshold V_{th}, i.e. if the neuron emits a spike t_0, V is reset to V_{reset} and kept there for the time τ_{refrac}:

$$V(t_0) = V_{\mathrm{th}} \tag{2.2}$$

$$V(t) = V_{\mathrm{reset}} \quad \forall t \in (t_0, t_0 + \tau_{\mathrm{refrac}}] \tag{2.3}$$

In equation (2.1) all synaptic input is comprised in I_{syn}. For a LIF neuron with current based, exponential synapses the synaptic input, separated in excitatory (exc) and inhibitory (inh) input, can be written as

$$I_{\mathrm{syn,\,exc}}(t) = \sum_i \sum_j w_i \cdot \exp\left(\frac{t_{ij} - t}{\tau_{\mathrm{syn,\,exc}}}\right) \cdot \Theta(t - t_{ij}) \tag{2.4}$$

$$I_{\mathrm{syn,\,inh}}(t) = -\sum_k \sum_l w_k \cdot \exp\left(\frac{t_{kl} - t}{\tau_{\mathrm{syn,\,inh}}}\right) \cdot \Theta(t - t_{kl}) \tag{2.5}$$

$$I_{\mathrm{syn}} = I_{\mathrm{syn,\,exc}} + I_{\mathrm{syn,\,inh}}. \tag{2.6}$$

The sum over i and k adds all presynaptic partners while j and l run over all spike times of one presynaptic partner. $\tau_{\mathrm{syn,\,exc}}$ and $\tau_{\mathrm{syn,\,inh}}$ are the synaptic time constants and w_i and w_k denote the synaptic strength. An alternative model with conductance based synapses includes a dependency of the synaptic current on the membrane potential:

$$I_{\mathrm{syn,\,exc}}(t) = \sum_i \sum_j w_i \cdot \exp\left(\frac{t_{ij} - t}{\tau_{\mathrm{syn,\,exc}}}\right) \cdot \Theta(t - t_{ij}) \cdot \big(E_{\mathrm{rev,\,exc}} - V(t)\big) \tag{2.7}$$

$E_{\mathrm{rev,exc}}$ is the excitatory reversal potential. The synaptic current for the inhibitory case can be calculated accordingly using the inhibitory reversal potential $E_{\mathrm{rev,inh}}$ and $\tau_{\mathrm{syn,inh}}$.
To model more biological features the LIF model can be extended to the adaptive exponential integrate-and-fire (AdEx) model (Brette and Gerstner, 2005). It includes an exponential term which causes a strong rise of the membrane voltage if a soft threshold V_{T} is crossed. Additionally, the adaptation variable w is added. The adaptation allows a change in the behavior of the neuron (e.g. the spike frequency) while being stimulated with the same

stimulus all the time (see for example figure 18). The differential equations for the AdEx neuron model are

$$C_{\mathrm{m}}\frac{dV}{dt} = -g_{\mathrm{L}} \cdot (V - E_{\mathrm{L}}) + g_{\mathrm{L}} \cdot \Delta_{\mathrm{T}} \exp\left(\frac{V - V_{\mathrm{T}}}{\Delta_{\mathrm{T}}}\right) + I_{\mathrm{syn}} + I_{\mathrm{ext}} - w \tag{2.8}$$

$$\tau_{\mathrm{w}}\frac{dw}{dt} = a \cdot (V - E_{\mathrm{L}}) - w \tag{2.9}$$

with Δ_{T} and V_{T} as voltage parameters for the exponential term and the adaptation conductance a. Both variables V and w are reset if the threshold V_{th} is crossed:

$$V(t_0) = V_{\mathrm{th}} \tag{2.10}$$

$$V(t) = V_{\mathrm{reset}} \quad \forall t \in (t_0, t_0 + \tau_{\mathrm{refrac}}] \tag{2.11}$$

$$w \to w + b \tag{2.12}$$

The increasing of w by a fixed amount of b is called spike-triggered adaptation.
Figure 1 shows a simulation of an exemplary voltage trace that demonstrates the behavior of an AdEx neuron. The neuron receives synaptic input at irregularly spaced intervals. For each arriving input a post synaptic potential (PSP) is visible on the membrane voltage. If the inputs arrive shortly after each other, the PSPs integrate. When the soft threshold voltage of $V_{\mathrm{T}} = -50\,\mathrm{mV}$ is crossed a sharp increase in the membrane voltage caused by the exponential term is visible. The neuron then reaches the spike threshold and is reset. After the reset follows a small downswing of the membrane potential which is caused by the spike-triggered adaptation.

2.1.2 Multi-compartment neurons

The neuron models covered in section 2.1.1 are called point neuron models. This means that the spatial structure of a biological neuron as shown in figure 2 is reduced to a single point. A biological neuron has multiple dendrites and an axon connecting to the cell body called the soma. The synapses which connect the neuron to its presynaptic partners are located at the dendrites. Arriving signals travel from the dendrites to the soma. If the accumulated input is sufficiently strong, a spike is generated in the axon and travels down the axon to the synapses which connect the neuron to its postsynaptic partners.
A simple way to add spatial structure to the neuron models described in section 2.1.1 is to connect point neurons with an intercompartment conductance g_{ic} as shown in figure 3 (Gerstner and Kistler, 2002). This couples the differential equations of the neuron models by adding an additional current I_{ic}. For a compartment that is connected to n other compartments the sum of all intercompartment currents is

$$I_{\mathrm{ic}} = \sum_{i=0}^{n} g_{\mathrm{ic,\,i}} \cdot \left(V_{\mathrm{comp,\,i}} - V\right) \tag{2.13}$$

which depends on the voltage differences between the membrane voltages of the neuron V and the other compartments $V_{\mathrm{comp,\,i}}$. By that the former point neurons become compartments of a larger neuron which interact passively via the intercompartment conductances. Depending on the location in the compartment tree and parameter choices, the individual compartment can mimic the roles of dendrites, soma or axon on a basic level.

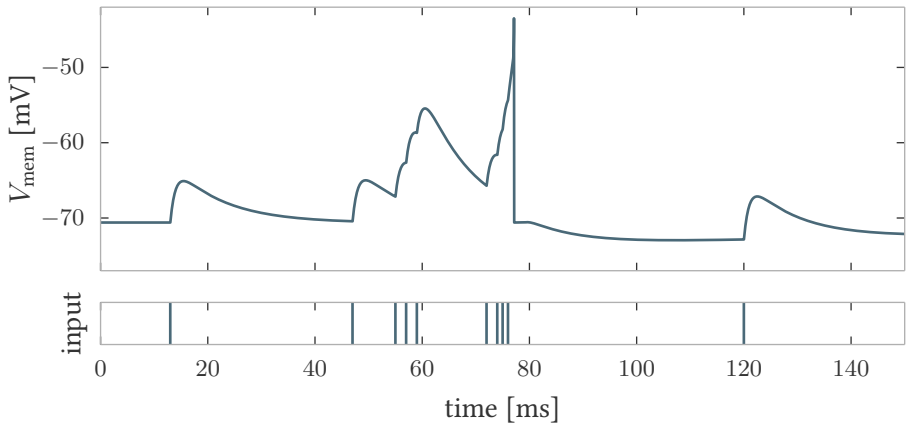

Figure 1: Top: Software simulation of an exemplary AdEx neuron using the PyNN simulator (Davison et al., 2009) with NEST as back-end (Gewaltig and Diesmann, 2007). The neuron is stimulated by spike input (bottom). When the inputs arrive in short succession the PSPs sum up. When the membrane voltage reaches $V_{\mathrm{T}} = -50\,\mathrm{mV}$ which triggers a sharp upswing of the membrane potential caused by the exponential term. After reaching the threshold the membrane voltage is reset and kept at the reset voltage for the duration of the refractory period. The spike-triggered adaptation causes the small downswing of the voltage after the reset.

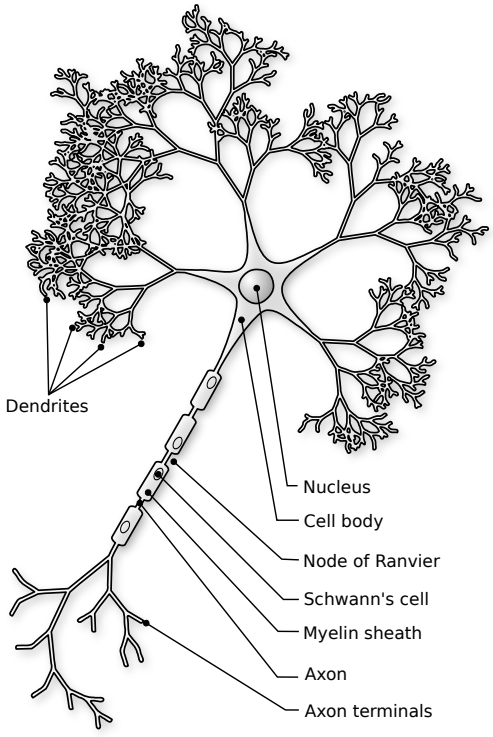

Figure 2: Drawing of a biological neuron by Rougier (2007). It shows the basic structure of a neuron. There are three main parts, the dendrites, the cell body (soma) and the axon.

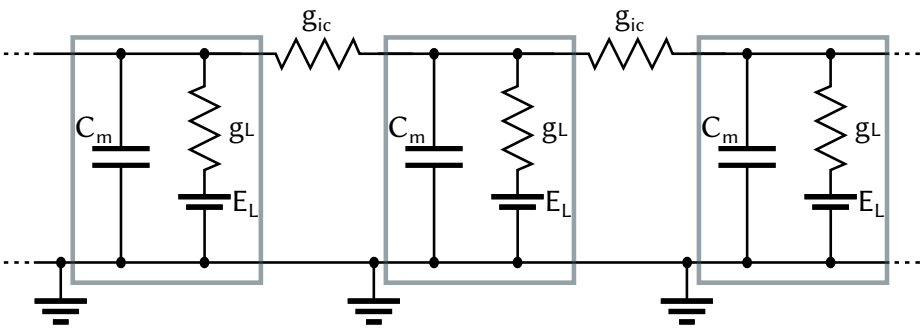

Figure 3: Illustration of a simple multi-compartment neuron consisting of multiple LIF neurons (gray boxes picturing membrane capacitance C_{m}, leakage conductance g_{L} and the leakage potential E_{L}) connected to a chain via intercompartment conductances g_{ic}.

2.1.3 Dendritic plateau potentials

Physiological measurements have shown an important communication mechanism between dendrites and soma, which passively connected compartments lack, the dendritic spikes (Antic et al., 2010). Figure 4 shows a collection of spike types, their typical location of occurrence and their waveform in a cortical pyramidal neuron.
The waveform of the action potential, which is initiated in the axon consists of a sharp and short rise in the membrane voltage. This is the signal which allows interneuron communication. Its characteristic sharp rise and short duration are modeled by the exponential term and the threshold in the AdEx equations.
The calcium and N-methyl-D-aspartate (NMDA) plateau potentials are initiated in the dendrites and the apical trunk. The apical trunk is, as shown figure 4 B, the connection between the tuft dendrites and the soma. Their shapes differ strongly from the shape of the action potential. Typically, they show a plateau shape of a length in the order of 50 to 100 ms. The addition of active and nonlinear components is believed to play a significant role in cortical information processing Antic et al. (2010); Larkum (2013).

2.2 Accelerated Mixed-Signal Neuromorphic Hardware

The neuromorphic hardware discussed in this thesis uses a mix of analog and digital circuits (mixed-signal design). The neurons and synapses consist of analog circuits that are designed to mimic the dynamics of mathematic neuron model e.g. the AdEx neuron. In contrast to a numerical simulation the state of the variables is not calculated at discrete points in time but they evolve in continuous time since they are represented by electrical quantities in the circuit. The membrane potential of a neuron, for example, is a physical voltage which is stored on a capacitor in the neuron circuit. The components of the neuron and synapse circuits are extremely small, which allows for small capacitances and high conductances. Therefore, the neuron time constants such as the membrane time constant $\tau_{\mathrm{m}} = \frac{C_{\mathrm{m}}}{g_{\mathrm{L}}}$ are up to a factor of ten thousand times shorter than in biology. This allows to reduce the duration experiments which would take years in biological real time to hours. But the high acceleration factor also increases the difficulty of digital communication on

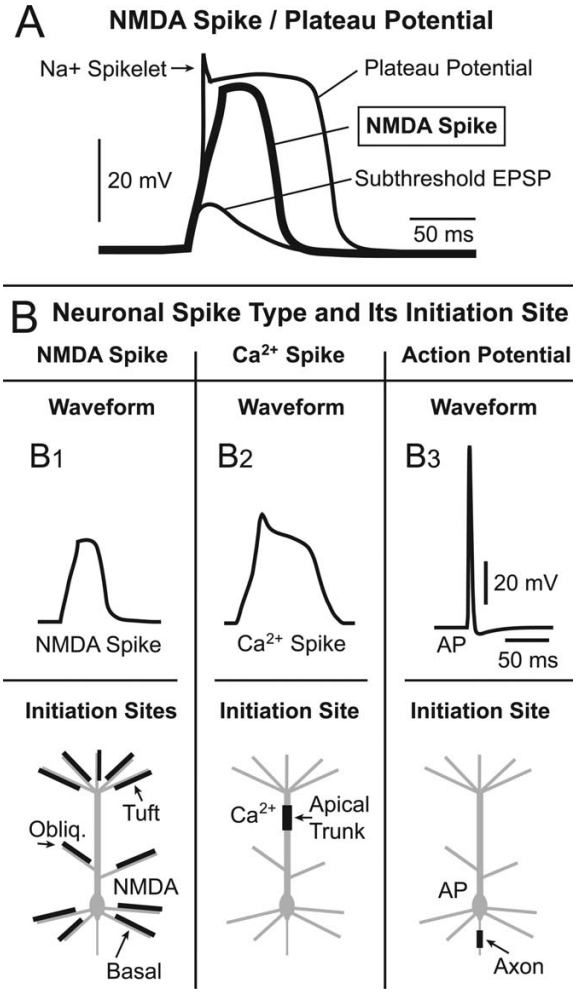

Figure 4: Overview of spike types in a cortical pyramidal neuron, taken from Antic et al. (2010). **A** shows an overlay of the shapes of an excitatory post synaptic potential (PSP), an NMDA spike and a calcium plateau potential. In **B** the waveforms of calcium plateau potential (initiation in the apical trunk) and NMDA spike (initiation in the dendrites) are compared to an action potential initiated in the axon.

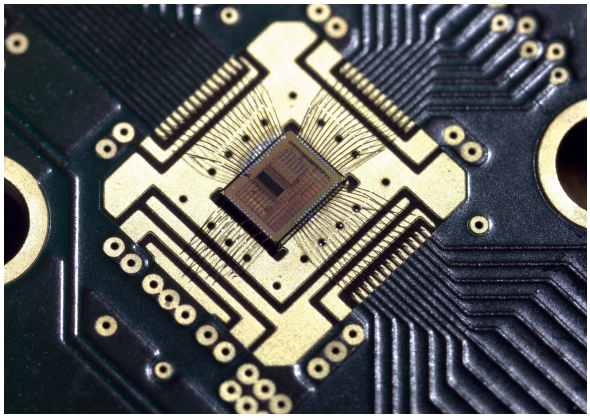

Figure 5: Photograph of the HICANN DLS 2 prototype chip, taken by Matthias Hock.

the hardware itself, e.g. spikes, or between the hardware and host computer, since high bandwidths and precision are required.

2.2.1 HICANN Wafer System

The High Input Count Analog Neural Network (HICANN) chip (Millner, 2012) is the basis of the BrainScaleS system which is developed in the scope of the BrainScaleS project (BrainScaleS, 2012) and the Human Brain Project (HBP SP9 partners, 2014). The HICANN contains 512 neurons and 512×224 synapse circuits that are fabricated in 180 nm CMOS process. The circuits are designed to mimic the dynamics of AdEx neurons with conductance based synapses. In order to be able to emulate networks with more than 512 neurons multiple chips must be used simultaneously. To this end, the wafer of HICANN chips is not cut into single chips after production but instead a post-processing layer is added that allows communication between the chips. A full wafer contains 200000 neurons and 44 million synapses that can be connected to a network (Schemmel et al., 2010, 2008).

2.2.2 HICANN DLS

The area consumption of the circuit components can be reduced by using the 65 nm fabrication process. This allows the addition of more complex circuitry as for example the plasticity processing unit (PPU) (Friedmann et al., 2016). New versions of the HICANN chips fabricated in 65 nm process and with added PPU are called HICANN Digital Learning System (HICANN DLS). HICANN DLS wafers will in the future replace the HICANN wafers currently in use. Before the final wafers are being produced a row of test chips were (HICANN DLS 1, 2 and 3) and will be fabricated. Figure 5 shows a photograph of the second test chip, the HICANN DLS 2.

The HICANN DLS 3 is the newest test chip. It adds several new features to HICANN DLS 2 (Aamir et al., 2016). In the following we introduce additions to the neuron circuit that will be treated in the following chapters.

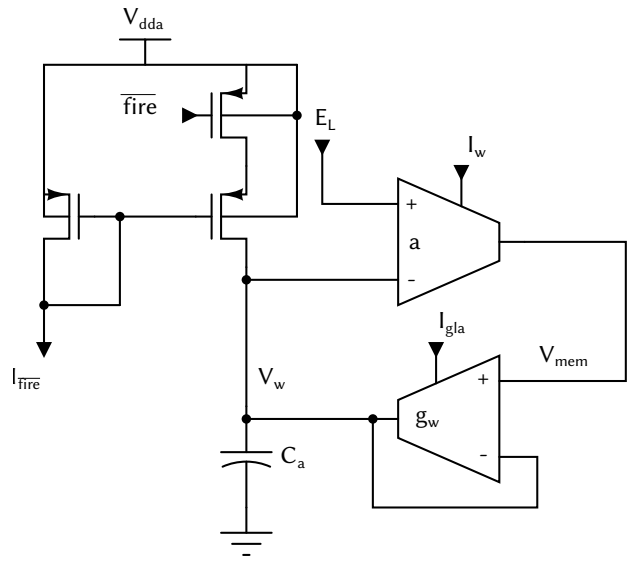

Figure 6: Schematic of the adaptation term on HICANN, taken from (Millner, 2012). On the top-left the mechanism for spike-triggered adaptation (enabled by the fire signal of the neuron) is shown. The adaptation variable V_w is stored on the capacitor C_a. The decay of V_w with the time constant $\tau_\mathrm{w} = \frac{C_\mathrm{a}}{g_\mathrm{w}}$ is modeled with the OTA on the bottom right. The adaptation current onto the membrane is produced by the OTA on the top right.

Switchable Adaptation The circuit in figure 6 shows the functional principle of the adaptation term on the HICANN chip (Millner, 2012). In hardware the adaptation current w is modeled as a voltage V_w which is stored on the capacitor C_a. The transformation between hardware and AdEx adaptation (see equation (2.9)) is

$$w = a \cdot (V_\mathrm{w} - E_\mathrm{L}) \tag{2.14}$$

which changes the differential equation of the adaptation variable to

$$\frac{dV_\mathrm{w}}{dt} = \frac{1}{\tau_\mathrm{w}} \cdot (V - V_\mathrm{w}) \tag{2.15}$$

where $\tau_\mathrm{w} = \frac{C_\mathrm{a}}{g_\mathrm{w}}$. In contrast to the AdEx model the spike-triggered adaptation value b is not added directly to the adaptation current w flowing onto the membrane, but rather as a voltage step to V_w. The resulting adaptation current onto the membrane is $I_\mathrm{adapt} = a \cdot (V_\mathrm{w} - E_\mathrm{L})$. As the circuit for adding the spike-triggered adaptation b to V_w only allows for a positive b and the conductance a, realized by an operational transconductance amplifier (OTA), is also always positive, the hardware neuron realizes only one out of four possible sign combinations of a and b.

Compared to the adaptation circuit on the HICANN chip the new circuit for HICANN DLS 3 (designed by Syed Ahmed Aamir) allows more flexibility: The spike-triggered adaptation b can now be modeled by a positive or a negative current onto C_a. The sign switch is controlled by the digital parameter `en_pos_vw`. Additionally, the input into the OTA that emulates the adaptation conductance a is switchable (see figure 7) to allow both signs for a. The sign switch is realized by the digital parameter `en_neg_va`.

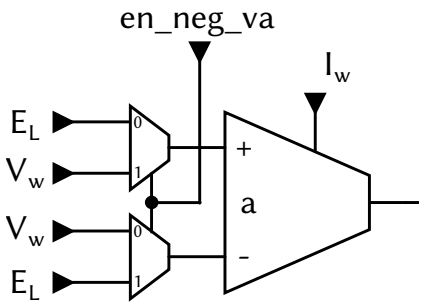

Figure 7: Schematic of new circuit for the production of the adaptation current. The inputs of the OTA are connected to two multiplexers (MUX) which choose, depending on their parameter `en_neg_va` the upper or lower of their inputs. This effectively switches the sign of the adaptation conductance a modeled by the OTA.

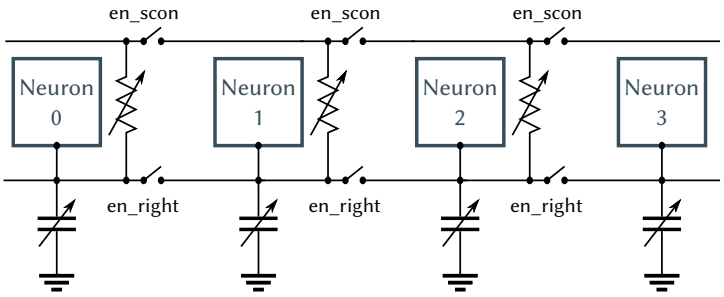

Figure 8: Schematic drawing of multi-compartment circuit on HICANN DLS 3. The neuron circuits except for the membrane capacitance are summarized in the gray boxes. To connect the neurons to one larger single-compartment neuron the `en_right` switches can be enabled. The `en_scon` switches connect the individual neurons via configurable conductances to a multi-compartment neuron.

Multi-compartment Circuits In addition to the AdEx features, multi-compartment circuits by Johannes Schemmel have been added to HICANN DLS 3 (Schemmel et al., 2017). Figure 8 shows a schematic drawing of the circuit. There are two possibilities of connecting the membrane capacitances of two neurons: First, they can be connected via the `en_right` switches, which connect the capacitances directly and merge the neurons into one neuron with a larger membrane capacitance. This can be useful to achieve for example longer membrane time constants τ_{m} and allows a higher number of synaptic inputs per neuron. Secondly the neurons can be connected to the soma line via a configurable conductance using the `en_scon` switches. This couples the individual neurons with a conductance and connects them to a multi-compartment neuron as described in section 2.1.2.

Plateau Potentials Another new feature of the HICANN DLS 3 neuron is the possibility to emulate plateau potentials as described in section 2.1.3 (Schemmel et al., 2017). A plateau potential is realized using the redesigned reset mechanism of the neuron.

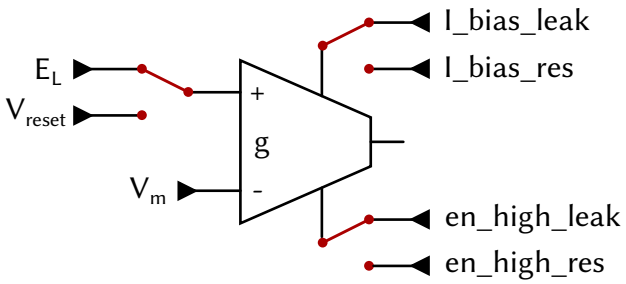

Figure 9: Schematic drawing of the analog part of the reset circuits. The newly designed leak OTA allows the switching between two modes, the normal leak mode and the second mode, which is used for the reset. The reset signal from the digital reset circuit triggers the switching between leak and reset (red switches). The digital parameters highs_leak and highs_res allow to increase the conductance of the OTA by a factor of 8 to 10.

Figure 9 shows the analog part of the new reset. The leak OTA of the neuron was redesigned for HICANN DLS 3 in order to allow a switching between two modes, the normal leak mode and the reset mode. Both modes effectively function the same way, but they can have different parameters and inputs. The reset signal from the digital reset circuit switches from the parameter set and inputs of the leak to the parameters and input of the reset and back. A digital reset signal also has the advantage that it allows a more precisely timed refractory period compared to an analog circuit as in the HICANN. Typically, the reset mode is configured to be significantly stronger than the leak mode.
By setting the reset voltage to a value above the threshold voltage and the refractory time to a high value, a plateau potential similar to the NMDA and Calcium spikes in biology is realized. Figure 10 shows an exemplary simulation of the circuits (for details on simulation method see section 4.2). The neuron is stimulated with synaptic input, reaches the threshold and is reset to a voltage above the threshold voltage.
The possibility to set the reset voltage above the threshold requires the addition of a holdoff time t_{holdoff} during which the reset mechanism is already switched off but the neuron can not fire again. This is necessary to prevent the neuron from firing continuously as soon it reaches the threshold once. The holdoff time allows the membrane voltage to decay below the threshold towards the leak potential before it is allowed to fire again. The refractory time is then defined as

$$\tau_{\text{refrac}} = t_{\text{reset}} + t_{\text{holdoff}} \tag{2.16}$$

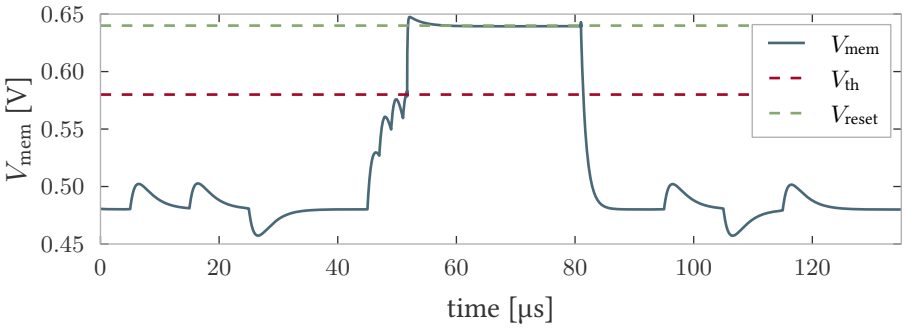

Figure 10: Exemplary transistor level simulation of a plateau potential, adapted from Schemmel et al. (2017). The plot shows the result of a circuit simulation, therefore the voltage and time are given in hardware domain. The neuron is stimulated with synaptic input (excitatory and inhibitory). When it reaches the threshold (lower dashed line) a spike is triggered and the reset mechanism pulls up the membrane voltage to the reset potential (upper dashed line). After a refractory time of approximately $\tau_{\text{refrac}} \approx 30\,\mu\text{s}$ the reset is switched of and the membrane potential decays back to the resting potential.

3 Single-Neuron Experiments

In the development of low-power accelerated neuromorphic hardware a trade-off between energy efficiency and speed on the one hand and precision and configurability on the other hand has to be made. It is therefore essential to assess how much the implemented circuits deviate from the mathematical model and whether the relevant characteristics of the model are captured by the circuits. This can be done using high-level test cases that are as similar as possible to the intended use cases of the hardware. However, during the development of new circuits, a detailed simulation (usually on transistor level) of more than very few neuron circuits is not feasible due to long simulation times. Therefore, this chapter presents a set of single-neuron experiments that allows to evaluate the performance of neuron circuits with respect to their accordance with the mathematical model they emulate and their biological plausibility.

3.1 Single-Neuron Modeling Competition

The quantitative single-neuron modeling competition was developed by Renaud Jolivet, Felix Schürmann, Thomas Berger, Richard Naud and others in an attempt to come up with a standardized test to quantify the performance of single-neuron models reproducing biological behavior. The competition took place in the years of 2007 to 2009 and the setup and results of 2007 and 2008 are extensively described in Jolivet et al. (2008b) and Jolivet et al. (2008a). Each year, four separate challenges were posed, approximately half of them for multi-compartment neurons the other half for single-compartment neurons. A challenge consists of test and training data. The training data contains input for the tested neuron model and output recordings of a biological neuron which serve as reference. The test data only contains the input data. In this thesis, we focus on challenge A of the year 2007 as it deals with single compartment neurons and tests the neuron models' ability to reliably reproduce biological spike times. The reliability of spike timing is important, as the spikes are the only means of communication of neurons in a network on the chip.
The challenge A of 2007 contains 12 data sets (cha, 2017). Eight of them are training sets and contain an input current for the neuron as well as four membrane voltage and spike time recordings of the same biological neuron which was stimulated with the given current. During the competition the test sets contained only the input currents, however, later, the target spike times of the biological neuron were made public. Figure 11 shows the current input and the four recordings of the target spike output per data set. The currents are given by an Ornstein-Uhlenbeck process

$$I(t+dt) = I(t) - \frac{I(t)}{\tau_\mathrm{I}}\,dt + m_\mathrm{I}\,dt + s_\mathrm{I}\xi(t)\sqrt{dt} \tag{3.1}$$

with a correlation length $\tau_\mathrm{I} = 1\,\mathrm{ms}$ and a time step of $0.2\,\mathrm{ms}$ (Jolivet et al., 2008a). $\xi(t)$ is a Gaussian random variable with zero mean and variance one. The mean μ_I of the resulting current distribution is $\mu_\mathrm{I} = m_\mathrm{I} \cdot \tau_\mathrm{I}$ and the standard deviation can be calculated as $\sigma_\mathrm{I}^2 = s_\mathrm{I}^2 \frac{\tau_\mathrm{I}}{2}$. While μ_I and σ_I vary over the different data sets, the same random sequence was used for the simulation of the stochastic process $\xi(t)$. Each current stimulus is $6.8\,\mathrm{s}$

long. The goal of the challenge is to reproduce the spike times of the biological neuron (also shown in figure 11) as closely as possible.

3.1.1 Γ-Measure

As method to determine how well the spike train produced by the neuron model matches the target spike train of the biological neuron the Γ-measure is chosen by Jolivet et al. (2008b). It is originally described in Kistler et al. (1997) where it is used to compare the spike times of different neuron models. The measure is calculated as

$$\Gamma = \frac{N_{\text{coinc}} - \langle N_{\text{coinc}} \rangle}{\frac{1}{2} \cdot \left(N_{\text{target}} + N_{\text{model}}\right)} \cdot \alpha \quad \text{with} \quad \langle N_{\text{coinc}} \rangle = \frac{2\Delta}{T} N_{\text{target}} N_{\text{model}} \tag{3.2}$$

where N_{coinc} is the number of coincident spikes of model and target, $\pm\Delta$ is the time window within which spikes are treated as coincident, T is the duration of the experiment, N_{target} and N_{model} are the number of spikes in the target and model spike train respectively and $\langle N_{\text{coinc}} \rangle$ is the number of coincident spikes that is expected if the model was a Poisson process. α is a normalization factor.
A score of $\Gamma = 1$ denotes a perfect match, $\Gamma = 0$ means that the model shows as many coincidences with the target as a random spike train of the same frequency would and a negative Γ indicates anti-correlation between model and target spike trains.
The normalization factor α is defined in Kistler et al. (1997) and Jolivet et al. (2008b) as

$$\alpha = \frac{1}{1 - \frac{2\Delta}{T} N_{\text{target}}}, \tag{3.3}$$

where Δ is the so-called coincidence window and T is the duration of the experiment. However, the normalization is defined differently in Jolivet et al. (2008a) and in the MATLAB script for the calculation of Γ provided on the web page of the competition (mat, 2016).

$$\alpha = \frac{1}{1 - \frac{2\Delta}{T} N_{\text{model}}} \tag{3.4}$$

In addition to the differences in the definition of α, the definition of a coincidence also varies in different sources. Figure 12 shows a collection of the different coincidence definitions and the resulting number of coincidences N_{coinc} they produce in an example. Kistler et al. (1997) defines coincidence by stating "In so doing, we have accepted a spike of the spike response model to be coincident with the corresponding spike of the Hodgkin-Huxley model if it arrives within a temporal precision of ± 2 ms." Instructions for the participants of the year 2009 (ins, 2016) state however "To evaluate this quantity, we calculate the number of coincidences N_{coinc} between the spikes in the data spike train one repetition at a time (target) and the spike train of the model submitted by a participant. This number is calculated by counting the number of target spikes for which we can find at least one model spike within ± 4 ms." Finally, the MATLAB script provided for the participants of competition in 2007 (mat, 2016) behaves differently from the versions described before (see figure 12). The script counts exactly one coincidence if multiple model spikes are within

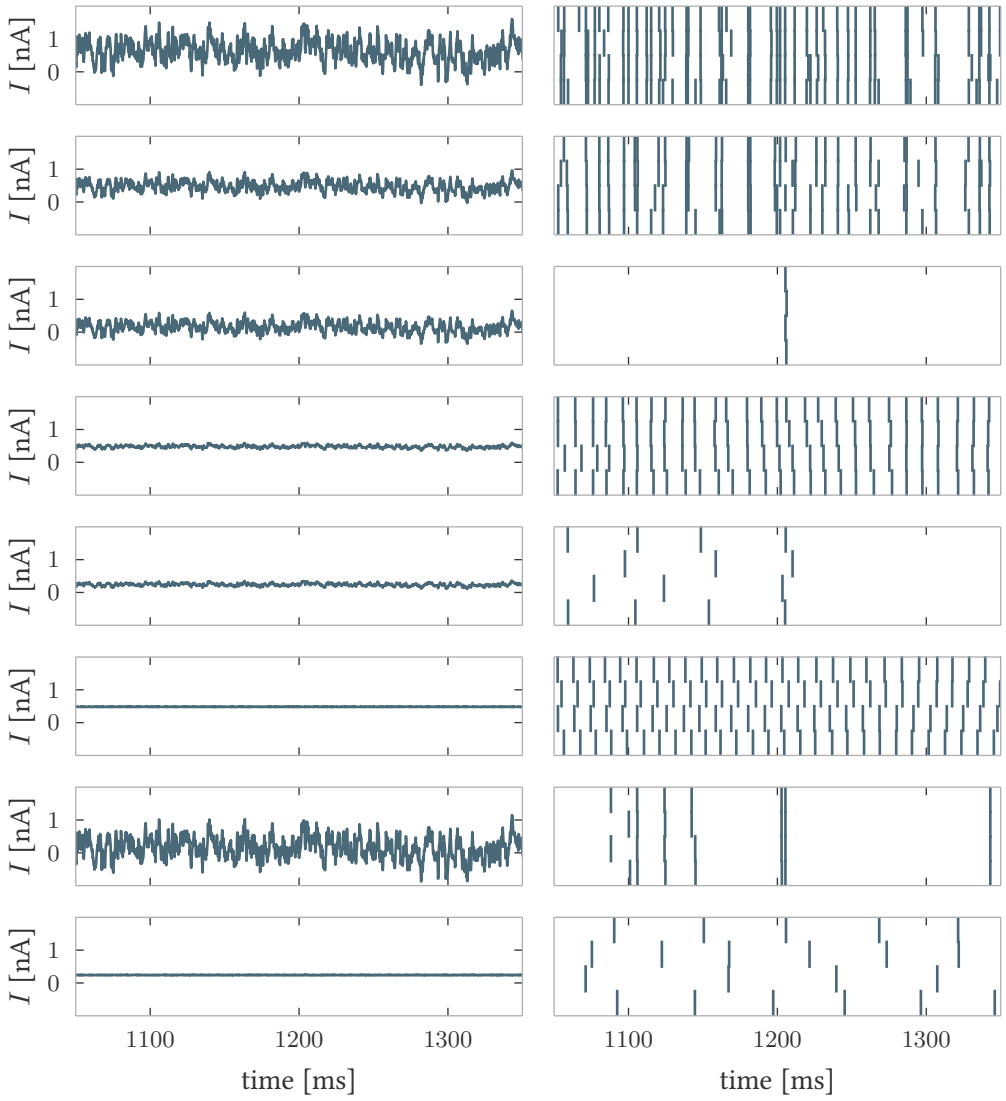

Figure 11: Left: Sections of the input currents in the training sets of challenge A. The means and standard deviations vary over the sets. Right: Recorded spike times of the biological neurons when stimulated with the corresponding current on the left. The recordings were repeated four times. Data taken from cha (2017).

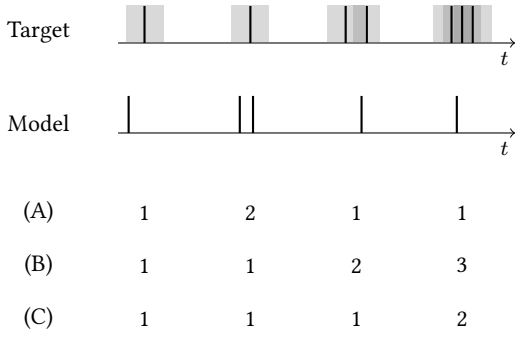

Figure 12: Different definitions of coincidences. The coincidence windows Δ are shaded gray around the target spikes. For the third and fourth group of target spikes the Δ windows overlap. The numbers below indicate the number of coincidences counted for the corresponding group of spikes depending on the used definition. The coincidences in (A) are calculated using the definition by Kistler et al. (1997). (B) follows the instructions for participants of the year 2009 (ins, 2016). (C) uses the algorithm of the provided matlab script (mat, 2016).

the Δ windows of one target spike. If there are is one model spike within the window of two target spikes, also only one coincidence is counted. However, if the Δ windows of $n > 2$ target spikes overlap one model spike is counted as $n - 1$ coincidences. This is most likely due to an implementation mistake. In the following we will, always count one coincidence if multiple target spike are met by one model spike and vice versa.
Equation (3.2) shows that the value of the Γ measure depends on the coincidence window Δ. As visible in figure 13, there are three regions of that dependency. For very small Δ, spikes must happen at nearly the same time to be recognized as a coincidence. This is rarely the case and therefore Γ is very low. This region is followed by a regime where Γ is nearly constant due to the fact that all coincidences are recognized and a larger Δ does not increase N_{coinc}. In the third region the denominator of the normalization factor $\alpha = \frac{1}{1-\frac{2\Delta}{T}N_{\text{target}}}$ approaches zero and causes a pole in Γ. The Γ measure is not valid in this region. Unfortunately, the third test data set is, with the $\Delta = 2\,\text{ms}$ given in Jolivet et al. (2008a) and (Jolivet et al., 2008b), in the invalid region.

3.1.2 Intrinsic Reliability and Performance

In order to estimate how well the biological neuron itself reliably reproduces its own spike times when stimulated with the same stimulus all recordings were repeated 4 times. As a measure for the intrinsic reliability Γ_{int} is used. It is calculated as the mean of the Γ-values between different repetitions of the recordings b_i, b_j of data set s:

$$\Gamma_{\text{int, s}} = \frac{2}{N_{\text{reps}}(N_{\text{reps}}-1)} \sum_{i=1}^{N_{\text{reps}}} \sum_{j=i+1}^{N_{\text{reps}}} \Gamma(b_i, b_j). \tag{3.5}$$

N_{reps} denotes the number of recording repetitions (in our case $N_{\text{reps}} = 4$). Figure 14 shows that the intrinsic reliability of a data set depends on the variability of the input

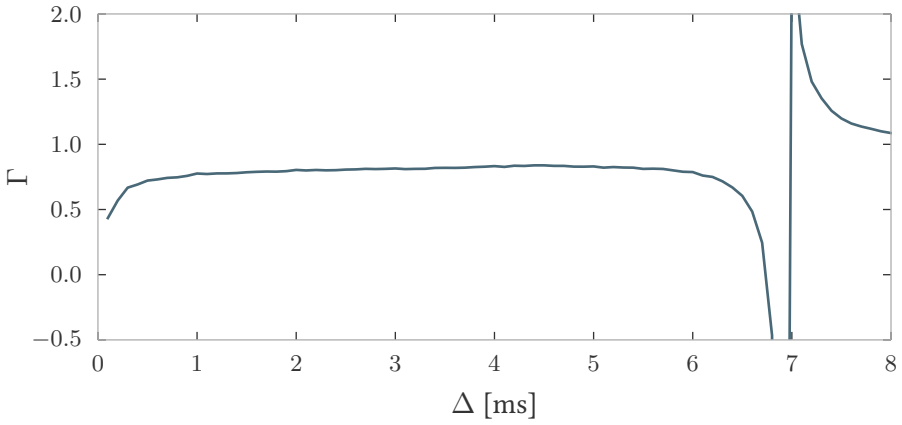

Figure 13: Exemplary evaluation of the Γ measure using different coincidence windows Δ. The measure is nearly independent of Δ over a wide range. However, for large Δ, Γ shows a pole, which renders the measure unstable and invalid in this region.

current. Irregular spike trains caused by an input current with high variation tend to have a higher intrinsic reliability than regular spike trains. This tendency can also be observed in figure 11.

If the biological neuron reproduces its spike times with a low reliability the neuron model can not be expected to perform well in reproducing the spikes in the 4 recordings. To take this into consideration, the overall performance of a neuron in the challenge is defined as

$$P = \frac{1}{N_{\text{sets}}} \frac{1}{N_{\text{reps}}} \sum_{s=1}^{N_{\text{sets}}} \sum_{i=1}^{N_{\text{reps}}} \frac{\Gamma(b_{si}, m_s)}{\Gamma_{\text{int}}(b_s)} \tag{3.6}$$

where b_{si} are the biological spike times in the i-th repetition of the s-th data set, and m_s are the spike times of the tested model for this data set.

Equation (3.6) shows that the intrinsic reliabilities of the biological measurements must be known in order to calculate the performance. Jolivet et al. (2008a) give the values for the four test sets $\Gamma_{\text{int}} = 0.22$, 0.76, 0.85 and 0.99. However, the values for Γ_{int} calculated from the test data published after the competition do not match the given values. We calculate $\Gamma_{\text{int}} = 0.99$, 0.93, 0.72 and 0.80. As it is not clear how the given values were calculated and which of them corresponds to which test set, the performance values published in Jolivet et al. (2008a) can not be used for further considerations in this work.

3.1.3 Parameter Fits

In order to test a neuron model in the single-neuron modeling competition the neuron parameters that lead to the closest match need to be determined. This is done by fitting the neuron parameters to the data given in the training data sets. Jolivet et al. (2008a) describes the method used for the AdEx neuron. An optimal parameter set was determined using a genetic algorithm (Vanier and Bower, 1999) using the Γ measure as fitness function. The

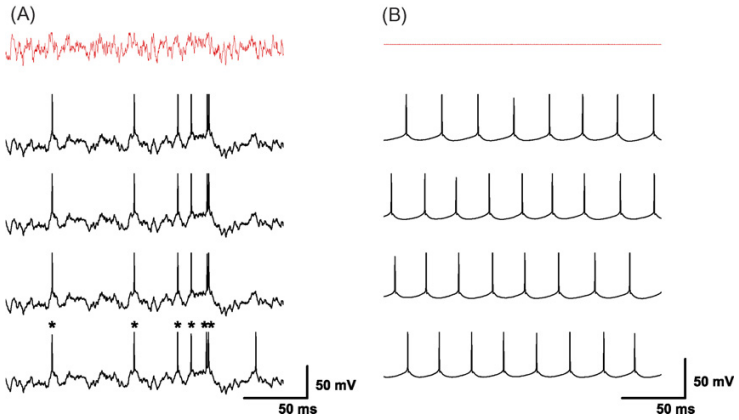

Figure 14: Taken from Jolivet et al. (2008a). Comparison of the reliability in spike timing for strongly varying input (A) and a nearly constant input current (B). For the input with high variability the spike timing is more reliable.

resulting parameters are given as $C_\mathrm{m} = 72\,\mathrm{pF}$, $g_\mathrm{L} = -13\,\mathrm{nS}$, $E_\mathrm{L} = -60\,\mathrm{mV}$, $V_\mathrm{T} = -38\,\mathrm{mV}$, $\Delta_\mathrm{T} = 0.006\,\mathrm{mV}$, $a = -0.5\,\mathrm{nS}$, $b = 36\,\mathrm{pA}$ and $\tau_\mathrm{w} = -25\,\mathrm{ms}$.

The membrane time constant and the adaptation constant are negative which is most likely due to a different sign definition in differential equations of the neuron model in the simulator. Additionally, the values for V_reset and V_th are missing. Due to the very low value of Δ_T the exact value of V_th is not relevant. V_reset, however, needs to be determined, as the parameter set can not be used without it. In order to determine the missing reset voltage a sweep over V_reset was performed, where, for each value, the neuron's score in the challenge test sets was calculated and compared to the test score given in Jolivet et al. (2008a).

Figure 15 shows the results of the sweep. As it is not certain which definition of the normalization factor α in the Γ measure was used in the paper, Γ was calculated using both possible definitions. We see that for low reset voltages the resulting Γ is very similar for both normalizations. For larger voltages the normalization in Jolivet et al. (2008a) shows a pole at approximately $V_\mathrm{reset} = -48\,\mathrm{mV}$ which renders the measure invalid for all higher reset values. Unfortunately, the only reset voltage, which produces a $\Gamma = 0.7$ as given in the paper, is in the invalid regime.

As the full optimal parameter set could not be determined with certainty, we determine an own optimal parameter set. Similar to the genetic algorithm used by Jolivet et al. (2008a) we use an evolutionary algorithm as well. However, in order to reduce simulation time, we narrow down the possible ranges for some of the neuron parameters beforehand.

By minimizing the squared distance between the simulated neuron model and the biological recordings $(V_\mathrm{simulated} - V_\mathrm{bio})^2$ in a region without spikes, the parameters are determined. For the fit nine patches of the voltage traces in the third training data set were used. Each patch was at least 200 ms long. The fit allowed to narrow down the parameter ranges of the membrane capacitance, membrane time constant and the leakage reversal

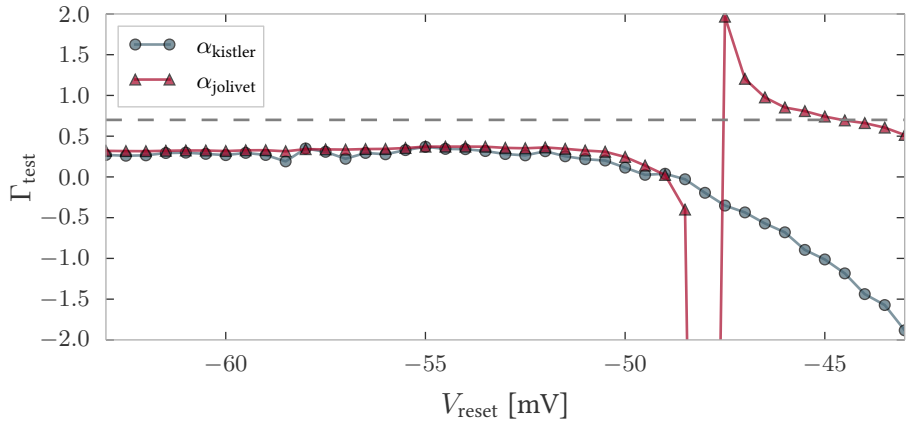

Figure 15: Sweep over V_{reset} in order to determine the reset voltage that produces the score $\Gamma = 0.7$ given in the paper (dashed line). The scores are calculated with both of the possible normalization factors α_{kistler} (blue circles) and α_{jolivet} (red triangles). The Γ calculated with α_{jolivet} shows a pole, which renders all the scores produced by larger reset voltages invalid. The score mentioned in the paper can only be reached in the invalid regime.

Method	Simulated Annealing	Genetic Algorithm	Differential Evolution
Final relative error	6.72×10^{-4}	2.022×10^{-3}	1.42×10^{-4}
Duration	6 d 6 h	1 h 52 min	57 min

Table 1: Comparison of optimization algorithms used for neuron parameter fitting. Values taken from Buhry et al. (2011).

potential

$$C_{\text{m}} = 43.9 \pm 6.8\,\text{pF} \tag{3.7}$$

$$\tau_{\text{m}} = 2.40 \pm 0.69\,\text{ms} \tag{3.8}$$

$$E_{\text{L}} = -58.5 \pm 1.0\,\text{mV} \tag{3.9}$$

The values are the averages of the fit results on the individual patches. The errors denote the standard deviation. Figure 16 shows the fit result on an exemplary patch of the used voltage traces.

The fit results can be used as a starting point for an evolutionary algorithm. We use the differential evolution algorithm (Storn and Price, 1997; Price et al., 2006) which was compared to other widely used optimization algorithms in the task of neuron parameter fitting (Buhry et al., 2011). It outperformed simulated annealing and a genetic algorithm (Vanier and Bower, 1999) in speed as well as precision (see table 1).

Similar to a genetic algorithm, differential evolution operates on a population that consists of N_{pop} vectors $\vec{x_i}$ in parameter space with $i \in [1, N_{\text{pop}}]$. From these vectors an offspring population $\vec{o_i}$ of the same size is created by drawing randomly 3 vectors of the original

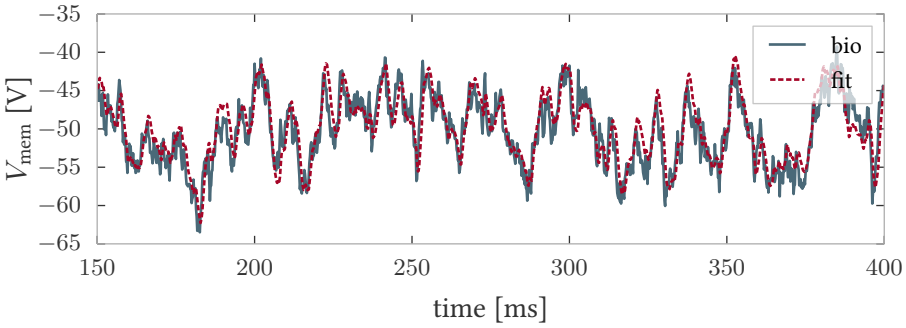

Figure 16: Simulation of neuron using parameters resulting from fit (dashed line) and recorded biological voltage (solid line) on an exemplary patch of voltages used for the least squares fit.

population and adding the scaled difference of two to the third:

$$\vec{o_i} = \vec{x_j} + F \cdot (\vec{x_k} - \vec{x_l}) \quad \text{with} \quad j, k, l \in [1, N_{\text{pop}}] \tag{3.10}$$

The scale factor F typically is smaller than 1. The offspring vectors $\vec{o_i}$ are crossed over with the parent vectors $\vec{x_i}$ by replacing, with a crossover probability p_c, the j-th component of the offspring vector j-th component of its parent. Then the fitness of all offspring and parent vectors is evaluated and for each i the offspring or the parent is chosen depending on their fitness value.

In our case, each component of the vectors $\vec{x_i}$ denotes one neuron parameter. For each parameter we choose a value range within which it can evolve. The ranges of the parameters which were predetermined in the least squares fit were set to the fit result $\pm$ two times the standard deviation. The ranges for the other parameters can be found in section D.1.

At first, we chose the Γ measure as fitness function. In contrast to the evaluation of the challenge, we set the coincidence window Δ to $\Delta = 1\,\text{ms}$ during the evolution, because the measure is not able to distinguish the quality of a perfect match and, for example, a spike train that shows deviations from the correct spike times within the coincidence window. With a smaller coincidence window, the inability to distinguish results appears in a later stage of the evolution.

This shortcoming of the Γ measure was also the reason to repeat the evolution with another fitness function: The reduced Gauss measure (Born, 2012) is inspired by the Schreiber similarity measure (Schreiber et al., 2003) which convolves the spike trains with gaussian functions and then calculates the cross correlation. The normalization of the Schreiber measure fails if multiple model spikes are close to one single target spike. This problem is solved by the reduced Gauss measure which first deletes all spikes from the model and the target spike train that do not have a corresponding partner in the other spike train. To be corresponding, the spikes must lie in the ± 1 ms interval of each other. The reduced spike trains are then convolved with a Gaussian, which results in the functions g_{model} and g_{target}

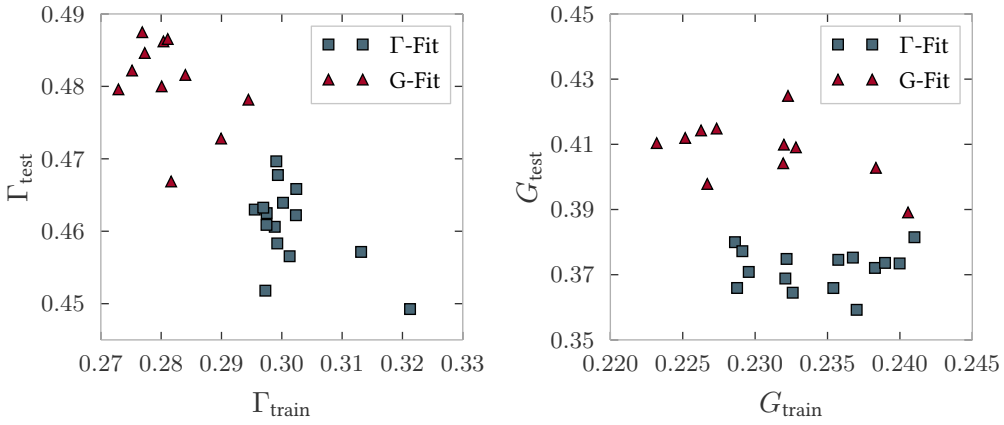

Figure 17: Results of the differential evolution. Left: Test and training scores of the resulting parameter sets calculated with the Γ measure, triangles denote parameter sets where the Gauss measure was used in the evolution, squares denote parameter sets where the Γ measure was used as fitness function. Right: Test and training scores of the same parameter sets but calculated with the reduced Gauss measure.

which are used to calculate the value G of the measure:

$$G = \frac{2}{N_{\text{model}} + N_{\text{target}}} \cdot \int_0^\infty g(t)_{\text{model}} \cdot g(t)_{\text{target}}\, dt \tag{3.11}$$

For the normalization, the original numbers of spikes in the spike trains N_{model} and N_{target} are used. This penalizes spike trains with too few or too many spikes.
Figure 17 shows the results of the parameter fits performed using differential evolution with the Γ measure as fitness function in comparison to the reduced Gauss measure. The evolution algorithm was started multiple times and evolved into a similar fitness regime for each repetition. We see that, although the results produced with the reduced Gauss measure show worse training scores when evaluated with the Γ measure, the test scores are better than the one where Γ was used in the evolution. This might be due to an overfitting when the Γ measure is used in the evolution.
From the collection of results, we choose three parameter sets for all further investigations. They can be found in table 2.

3.1.4 Extension to Spike Input

On the current HICANN and HICANN DLS chips, there is no possibility to stimulate the neuron with a varying current. Therefore, we extend the data sets of the single-neuron modeling competition by data sets that use spike input. In order to mimic the challenge data as closely as possible, the spike input is designed in a way, that the synaptic currents produced by the spikes show means and variances in the same range as the competition currents. For one excitatory and one inhibitory Poisson spike train with rates ν_{exc} and ν_{inh} the mean and variation of the resulting current can be calculated using the Siegert

Parameter	AdEx 0	AdEx 1	AdEx 2
C_m [pF]	0.044	0.044	0.044
τ_m [ms]	2.84	2.63	2.49
E_L [mV]	-59.05	-59.18	-58.81
V_T [mV]	-37.13	-38.60	-39.36
V_th [mV]	-36.53	-37.43	-37.92
V_reset [mV]	-57.55	-60.56	-60.85
Δ_T [mV]	2.51	1.74	1.86
τ_w [ms]	89.4	62.3	91.6
a [nS]	1.67	1.09	2.49
b [nA]	0.012	0.016	0.010

Table 2: Parameter sets determined by the differential evolution algorithm.

approximation (Siegert, 1951; Stromatias et al., 2015)

$$\mu_\mathrm{I} = \tau_\mathrm{syn} \cdot (w_\mathrm{exc} \cdot \nu_\mathrm{exc} + w_\mathrm{inh} \cdot \nu_\mathrm{inh}) \tag{3.12}$$

$$\sigma_\mathrm{I}^2 = \frac{\tau_\mathrm{syn}}{2} \cdot \left(w_\mathrm{exc}^2 \cdot \nu_\mathrm{exc} + w_\mathrm{inh}^2 \cdot \nu_\mathrm{inh} \right) \tag{3.13}$$

where w_exc, w_inh are the excitatory and inhibitory synaptic weights and τ_syn is the synaptic time constant. The competition currents cover a range of $0 < \mu_\mathrm{I} < 650\,\mathrm{pA}$ and $0 < \sigma_\mathrm{I} < 350\,\mathrm{pA}$. A choice of the Poisson rate ν and the synaptic time constant then allows the calculation of the synaptic weights for a given mean and variation, if excitatory and inhibitory input rates and time constants are chosen to be equal. However, a solution only exists if

$$\frac{\sigma_\mathrm{I}^2 \tau \nu}{\mu_\mathrm{I}^2} > \frac{1}{4} \tag{3.14}$$

is fulfilled. This may for example not be the case for a high mean and low variance. If the solution does not exist the data set can nevertheless be created by stimulating the neuron with a constant current $I_\mathrm{offset} = \mu_\mathrm{I}$ in addition to the spike input. Then the excitatory and inhibitory synaptic weights are symmetric $w_\mathrm{exc} = -w_\mathrm{inh}$.
We created data sets for the Poisson frequencies $\nu = 100\,\mathrm{Hz}$, $500\,\mathrm{Hz}$ and $2\,\mathrm{kHz}$ and the synaptic time constants $\tau_\mathrm{syn} = 1\,\mathrm{ms}$ and $5\,\mathrm{ms}$. For each combination of τ_syn and ν the range of μ_I and σ_I is covered by 11 different input spike trains and their corresponding weights. As references we use simulated AdEx neurons with the optimal neuron parameter sets determined in section 3.1.3.

3.2 AdEx Firing Patterns

The collection of firing patterns that can be achieved with an AdEx neuron, which is stimulated with a step current, described in Naud et al. (2008) can serve as a test case for analog implementations of the AdEx model. To reproduce the patterns, the neuron parameters,

especially the adaptation parameters, are required to be precisely tunable over a wide value range. Figure 18 shows a software simulation of the eight different firing patterns using an adapted version of the neuron parameter sets given in (Naud et al., 2008) by Paul Müller. A table with all neuron parameters can be found in section B.1.
In figure 18 the voltage traces of the patterns *tonic spiking*, *adaptation*, *initial bursting*, *regular bursting*, *delayed accelerating*, *delayed regular bursting* and *irregular firing* are shown. Additionally, figure 19 shows the phase space plots of the patterns together with the nullclines of the differential equations of the AdEx neuron. The V-nullcline can be calculated as

$$w = -g_{\mathrm{L}} \cdot (V - E_{\mathrm{L}}) + g_{\mathrm{L}} \cdot \Delta_{\mathrm{T}} \exp\left(\frac{V - V_{\mathrm{T}}}{\Delta_{\mathrm{T}}}\right) + I_{\mathrm{ext}} \tag{3.15}$$

and the w-nullcline as

$$w = a \cdot (V - E_{\mathrm{L}})\,. \tag{3.16}$$

Tonic spiking is characterized by equidistant spikes followed by exclusively sharp resets or exclusively broad resets. After a sharp reset, the membrane potential rises immediately after the end of the refractory period, while after a broad reset the membrane potential first decreases below the reset potential and afterwards rises more slowly. A broad reset occurs when the neuron is reset above the V-nullcline and it must return below the nullcline before being able to spike again. The tonic spiking in figure 18 shows sharp resets since the neuron is (as visible in figure 19) far below the V-nullcline.
The adaptation pattern is characterized by increasing interspike intervals (ISI) and exclusively sharp resets. The phase space plot shows that a strong spike-triggered adaptation is necessary for this pattern: When the neuron reaches the threshold, the membrane voltage is reset to V_{reset} but the adaptation variable is increased by a large amount. As $-w$ is the current onto the membrane, a larger w increases the time until the neuron reaches the threshold again.
The initial bursting consists of a burst at the beginning, which is characterized by multiple spikes with short ISI and sharp resets, followed by single spikes with broad resets in between. During the burst all resets are below the V-nullcline, the burst ends when the nullcline is crossed. After the initial burst the neuron is close enough to the nullcline to cross it with every spike. Because of that, all following spikes are followed by a broad reset.
The regular bursting pattern is similar to initial bursting except for the fact that the first burst is followed by other bursts instead of single spikes with broad resets in between. This is due to the fact that in contrast to the initial bursting a single spike is not enough to reset the neuron above the V-nullcline. Initial bursting and regular bursting require a fine tuning of the adaptation parameters a and b.
Delayed accelerating is characterized by a long time difference between the onset of the stimulating step current and the first spike of the neuron. Additionally, the ISIs decrease in this pattern. In contrast to the previous patterns, a is negative here. This causes the neuron to move faster in the phase space the further away it is from the w-nullcline, i.e.

the higher V is, which causes the slow onset of the firing. The delayed regular bursting is similar to the regular bursting but shows a delayed onset such as delayed accelerating, which is due to the negative a.
A transient spiking pattern shows only one single spike followed by the membrane potential settling on a constant value. This is caused by the fact that the nullclines cross in the phase space, which produces a fixed point. Once the neuron reaches this fixed point after the first spike it stays there until the current stimulus ends and the fixed point disappears.
The irregular firing pattern shows chaotic firing behavior of the neuron. The ISIs change without any periodicity and there are sharp and broad resets. This pattern appears only for a set of parameters that seems not to be connected in the parameter space although mostly it appears for negative a, large positive spike-triggered adaptation b and high resets (Naud et al., 2008).
In order to reproduce all firing patterns on a neuromorphic implementation of the AdEx model, the parameter ranges especially for the adaptation but also for example for the leak- and exponential term must be chosen appropriately and need to be precisely tunable.

3.3 Backpropagation-Activated Calcium Spike Firing

Backpropagation-activated calcium spike (BAC) firing is a spike mechanism measured in layer 5 pyramidal neurons (Larkum et al., 1999). It allows a coincidence detection between the basal and apical input of a neuron. The mechanism is hypothesized to be the foundation of cortical associations on a cellular level (Larkum, 2013).
Figure 20 shows the BAC firing mechanism as measured in a layer 5 pyramidal neuron. Stimulating the dendrite of the neuron with a current mimicking synaptic input produces a PSP in the dendrite which travels passively into the apical trunk and the soma. On the way, the PSP is strongly attenuated and results in a marginal effect at the soma. When the soma is stimulated with a step current a somatic spike is triggered. The action potential propagates back into the apical trunk and the dendrites. A combined stimulation of soma and dendrites triggers first a spike in the soma, which propagates into apical trunk and dendrites and triggers, in combination with the dendritic stimulation, a calcium spike and an NMDA plateau potential. These, in turn, increase the membrane potential in the soma and cause the soma to emit a burst of spikes.
The BAC firing mechanism can serve as a biologically inspired test case for the novel multi-compartment and plateau potential features of the HICANN DLS 3 as it requires configurable intercompartment connections as well as different plateau potential durations and heights.

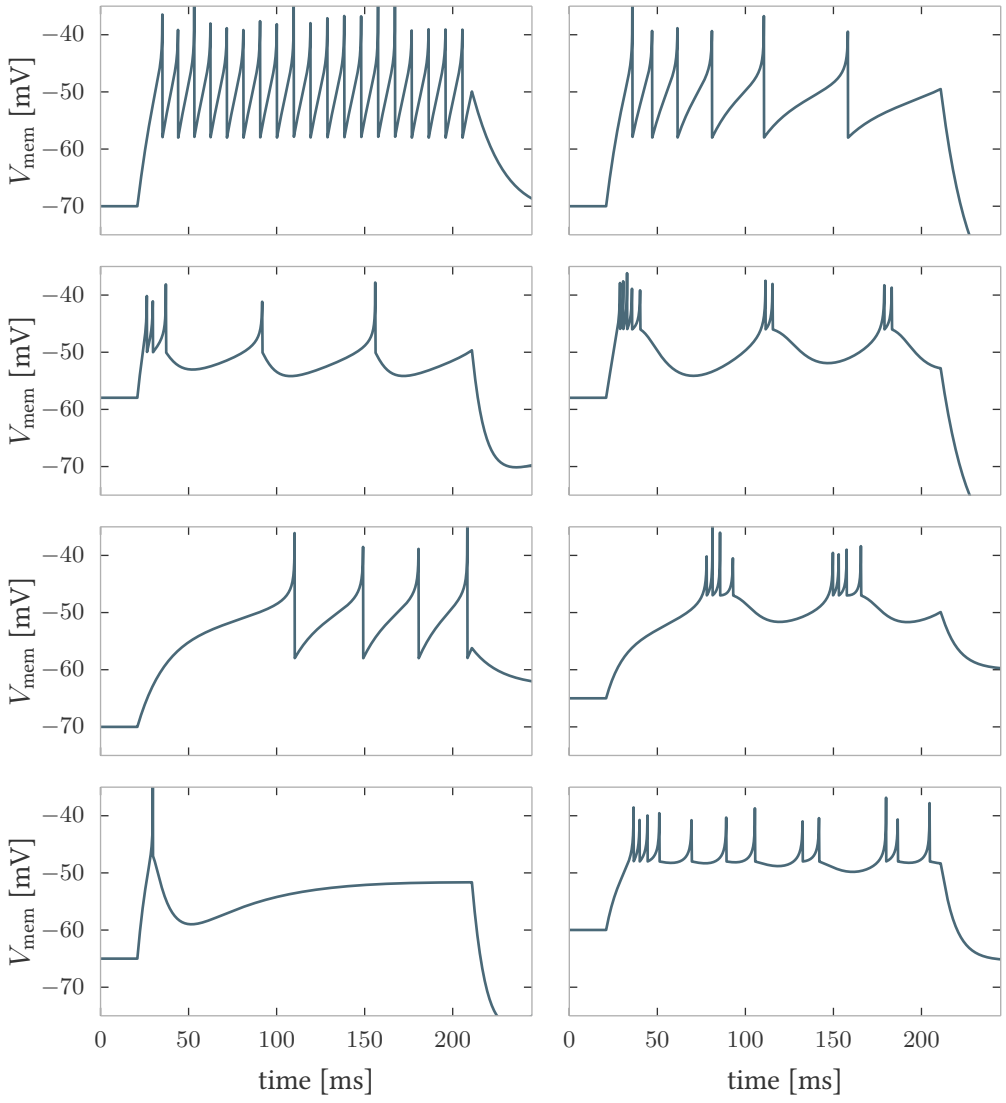

Figure 18: AdEx firing patterns as described in (Naud et al., 2008) simulated using PyNN. From top-left to bottom right the plots show tonic spiking, adaptation, initial bursting, regular bursting, delayed accelerating, delayed regular bursting, transient spiking and irregular firing.

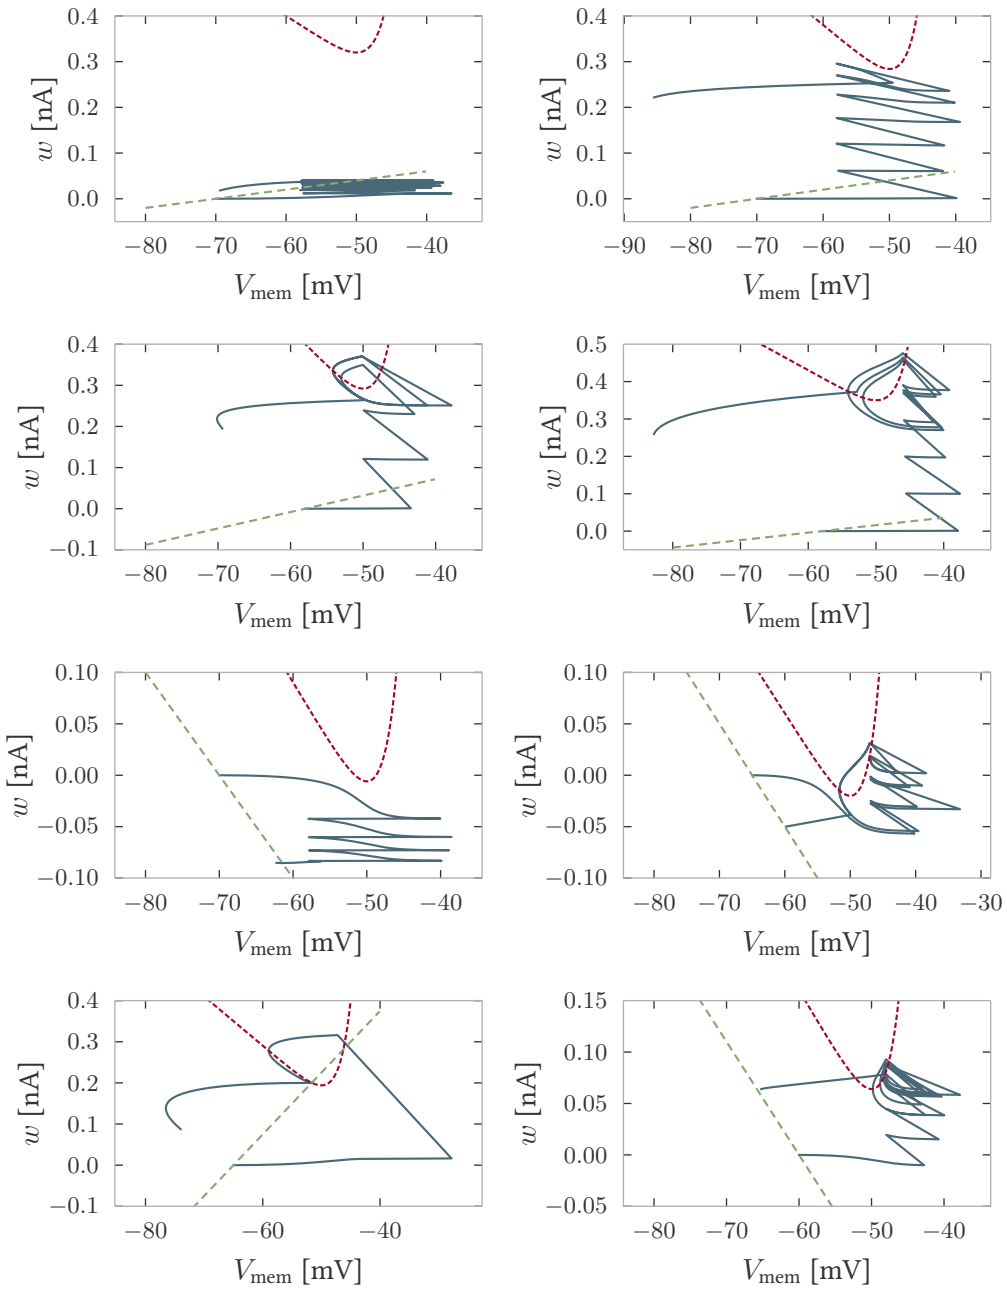

Figure 19: The trajectories in the phase space of the firing patters shown in figure 18 (blue). From top-left to bottom right the plots show tonic spiking, adaptation, initial bursting, regular bursting, delayed accelerating, delayed regular bursting, transient spiking and irregular firing. The nullclines are drawn with dashed lines, the V-nullcline in red and the w-nullcline in green.

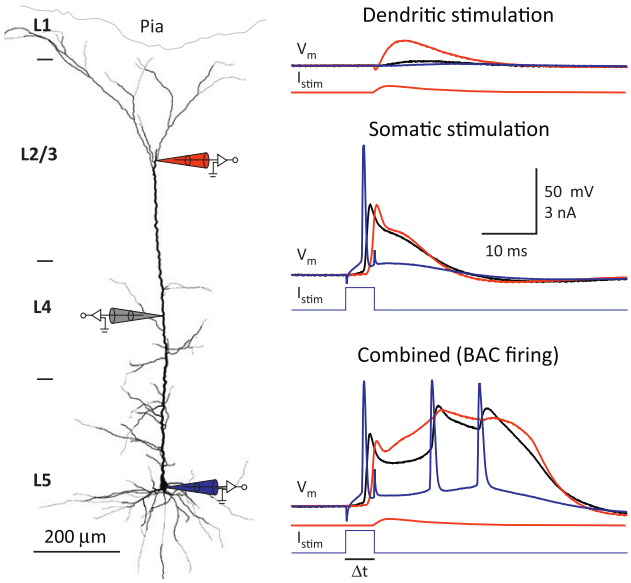

Figure 20: BAC firing mechanism measured in a layer 5 pyramidal neuron, taken from (Larkum, 2013). Three different stimulation patterns are applied to the neuron: A current stimulus mimicking synaptic input is applied to the dendrites and produces a PSP in the dendrite (top right, red). Stimulating the soma with a step current produces a single somatic spike (middle right, blue) that travels into the apical trunk and the dendrites (middle right, black and red). The combination of both stimuli triggers a somatic burst (bottom right, blue) and a plateau potentials in the apical trunk and dendrites.

4 Simulation Back-Ends

In this chapter we present two back-ends that are used for the analyses performed in the following chapters. First, the Extended Executable System Specification (Exess), a hardware system simulation tool and secondly the transistor-level circuit simulations for the HICANN DLS 3 chip.

4.1 Exess

The development of the hardware platform is accompanied by the implementation of an executable system specification, the successor of the ESS for the BrainScaleS system (Petrovici et al., 2014). The new, NEST-based simulator Exess (implemented by Oliver Breitwieser) is intended to simulate hardware platform restrictions such as bandwidth limitations and synapse loss. These simulations can provide a basis for hardware design decisions and give an insight into the influence of hardware restrictions on the performance of networks. The used neuron model (implemented by Paul Müller) is based on the implementation of the AdEx neuron in NEST but models the circuits of the HICANN DLS 3 neuron assuming ideal electrical components. This includes for example modeling the adaptation variable as a voltage (see section 2.2.2) instead of a current and implementing ideal models of electrical circuits. The leakage current for example is produced by a function modeling an ideal OTA, with output current I_{leak}

$$I_{\text{leak}} = \frac{I_{\text{bias,leak}}}{f_{\text{ota,leak}}} \cdot (V - E_{\text{L}}) \tag{4.1}$$

where $I_{\text{bias,leak}}$ is the bias current of the OTA and $f_{\text{ota,leak}}$ an internal scaling voltage. The Exess neuron also contains additional features compared to the HICANN DLS 3 neuron such as conductance based input and different adaptation circuits. As the Exess neuron is supposed to mimic the behavior of a hardware neuron, it operates in hardware time and voltage domain. This means that the time constants are 10^3 times shorter than in biology and the membrane voltage for example can have values in the range from 0 to 1.2 V. To compare Exess simulations to a PyNN simulation or biological recordings hardware times and voltages must be transformed into biological time and voltage domain:

$$t_{\text{hw}} = \frac{t_{\text{bio}}}{\alpha_{\text{t}}} \tag{4.2}$$

$$V_{\text{hw}} = V_{\text{bio}} \cdot \alpha_{\text{v}} + \omega_{\text{v}} \tag{4.3}$$

Typically, α_{v} is on the order of 10 and ω_{v} is on the order of 1 V. The acceleration factor α_{t} is 10^3 for the HICANN DLS 3.

The AdEx neuron in the PyNN simulator and the Exess neuron are parameterized differently, as shown in table 3. In the following a set of transformations that allows to translate PyNN parameters into the parameters for Exess neuron are introduced. The reversal potentials, the threshold voltages and the reset voltage can be transformed according to equation (4.3). The transformation parameters α_{v} and ω_{v} can be calculated by mapping the

PyNN inhibitory reversal potential to a hardware voltage of 0.2 V and the spike threshold to 1.2 V.

$$\alpha_{\mathrm{v}} = \frac{1}{V_{\mathrm{th,bio}} - E_{\mathrm{rev,bio}}} \tag{4.4}$$

$$\omega_{\mathrm{v}} = 1.2 - V_{\mathrm{th,bio}} \cdot \alpha_{\mathrm{v}} \tag{4.5}$$

The resistances controlling the synaptic time constants are calculated as

$$R_{\mathrm{syn}} = \frac{\tau_{\mathrm{syn}}}{\alpha_{\mathrm{t}} \cdot C_{\mathrm{syn}}}. \tag{4.6}$$

Conductance based weights are transformed as

$$w = \frac{C_{\mathrm{m,hw}}}{C_{\mathrm{m,bio}}} \cdot \alpha_{\mathrm{t}} \cdot w_{\mathrm{bio}} \cdot \frac{f_{\mathrm{ota,syn}} \cdot f_{\mathrm{syn,res}}}{I_{\mathrm{bias,cond}}} \cdot C_{\mathrm{syn}} \tag{4.7}$$

with fixed value for $I_{\mathrm{bias,cond}} = 1\,\mu\mathrm{A}$. The membrane time constant is controlled by

$$I_{\mathrm{bias,leak}} = f_{\mathrm{ota,leak}} \cdot \frac{C_{\mathrm{m,hw}}}{\tau_{\mathrm{m,bio}}} \cdot \alpha_{\mathrm{t}}. \tag{4.8}$$

The adaptation parameters are:

$$I_{\mathrm{bias,adapt}} = a_{\mathrm{bio}} \cdot f_{\mathrm{ota,adapt}} \cdot \frac{C_{\mathrm{m,hw}}}{C_{\mathrm{m,bio}}} \cdot \alpha_{\mathrm{t}} \tag{4.9}$$

$$I_{\mathrm{tau,w}} = f_{\mathrm{ota,adapt}} \cdot \frac{C_{\mathrm{w}}}{\tau_{\mathrm{w}}} \cdot \alpha_{\mathrm{t}} \tag{4.10}$$

$$b = \frac{b_{\mathrm{bio}}}{a_{\mathrm{bio}}} \cdot \alpha_{\mathrm{v}} \tag{4.11}$$

The exponential parameters are transformed as

$$\Delta_{\mathrm{T}} = \alpha_{\mathrm{v}} \cdot \Delta_{\mathrm{T,bio}} \tag{4.12}$$

$$I_{\mathrm{exp0}} = \frac{C_{\mathrm{m,hw}} \cdot \alpha_{\mathrm{t}} \cdot \Delta_{\mathrm{T}}}{\tau_{\mathrm{m,bio}}}. \tag{4.13}$$

Figure 21 shows the result of a simulation in PyNN and a simulation in Exess with transformed parameters. In order to make the traces comparable the PyNN trace was translated into hardware time and voltage domain. The untransformed trace can be found in figure 1. The novel multi-compartment and plateau potential features of the HICANN DLS 3 (Schemmel et al., 2017) are supported in the Exess neuron as well. Similar to the chip the plateau potential are realized by a high reset voltage. The connection of neurons via an inter-compartment conductance is realized by the NEST feature of gap-junctions (Hahne et al., 2015), which are voltage dependent channels between neurons.

4.2 HICANN DLS 3 Simulation

4.2.1 Transistor-level simulations

Transistor-level simulations are a valuable verification tool during the design phase of a chip. The simulations allow a detailed investigation of a circuits behavior for example

Neuron parameter	PyNN	PyNN unit	Exess	Exess unit
Membrane capacitance	cm	nF	C_m	F
Membrane time constant	tau_m	ms	I_bias_leak	A
Leakage reversal potential	v_rest	mV	E_L	V
Reset potential	v_reset	mV	V_reset	V
Spike threshold	v_spike	mV	V_peak	V
Exponential threshold	v_thresh	mV	V_th	V
Synaptic reversal potential	e_rev_E,I	mV	E_rev_x,i	V
Synaptic time constant	tau_syn_E,I	ms	R_syn_x,i	Ω
Synaptic weight	weight	µS	weight	C
Slope of exponential term	delta_T	mV	Delta_T	V
Adaptation conductance	a	nS	I_bias_adapt	A
Spike-triggered adaptation	b	nA	b	V
Adaptation time constant	tau_w	ms	I_tau_w	A
Refractory time	tau_refrac	ms	t_ref	s
Scaling factor exponential term	-	-	I_exp0	A
Capacitance synaptic input	-	-	C_syn	A
Capacitance adaptation term	-	-	C_w	A
Scaling voltage leak OTA	-	-	F_ota_leak	V
Scaling voltage synaptic OTA	-	-	F_ota_syn	V
Scaling voltage synaptic resistor	-	-	F_ota_syn_res	V
Scaling voltage adaptation OTA	-	-	F_ota_adapt	V

Table 3: Top: Collection of the names and units of the AdEx parameters in PyNN and Exess. Bottom: Additional parameters only used in Exess controlling circuit specific quantities, e.g. the size of a capacitor in the synaptic input circuit. The Exess specific parameters (scaling voltages and additional capacitances) are fixed to characteristic circuit values (we use $C_{\mathrm{w}} = 1\,\mathrm{pF}$, $C_{\mathrm{syn}} = 1\,\mathrm{pF}$, F_ota_leak = F_ota_syn = F_ota_adapt = 0.5 and F_ota_syn_res = 0.15).

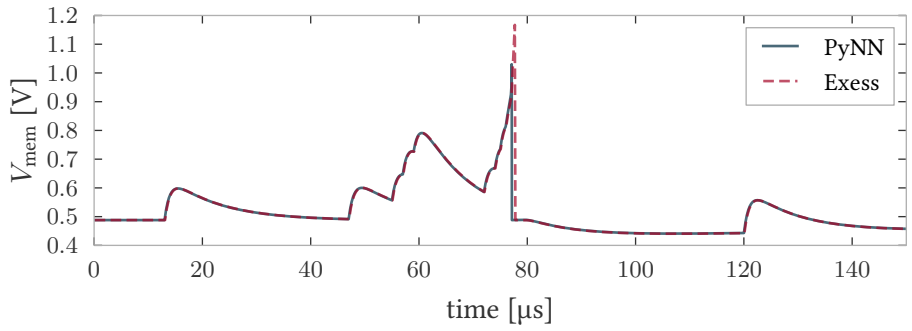

Figure 21: Comparison of voltage traces recorded using PyNN and, after parameter transformation, using Exess. The trace is shown in hardware time- and voltage domain. The PyNN result transformed according to equation (4.3).

under different parameter settings, varying input signals or changing temperature. Each individual transistor is parameterized using detailed physical characteristics provided by the manufacturer (in the case of HICANN DLS3 TSMC, Taiwan Semiconductor Manufacturing Company, Ltd., Hsinchu, Taiwan). The physical models include process variations and mismatch, see for example (Drennan and McAndrew, 2003). The circuits are simulated using the Spectre-simulator (Cadence Design Systems, Inc., San Jose, CA, USA).
Process variations describe the fact that variations in doping or film thicknesses can occur from one wafer to another. These variations are described by the so called process corners *tt, ss, ff, fs* and *sf*. The *s* stands for the slow, the *f* for the fast and *t* for the typical case. In the fs corner for example the nMOS transistors in the circuit are fast and the pMOS transistors are slow. Fast and slow describe mobility of the charge carriers, the higher the mobility, the faster the transistor.
In a corner simulation a transistor is either typical, slow or fast and two transistors that lie in the same corner and with the same parameters (e.g. width and length), are exactly identical. However, this does not account for the fact, that due to imperfections in the production process, all transistors on a chip are different, an effect called mismatch. This can be simulated in mismatch Monte Carlo simulations, where the transistor parameters are changed randomly. Many repetitions of Monte Carlo simulations allow to gain statistical insight about the circuit behavior under parameter variations that will occur during production.

4.2.2 HICANN DLS 3 Simulation Setup

The transistor-level simulation setup for the HICANN DLS 3 neuron uses the Cadence interface *Teststand* implemented by Sebastian Billaudelle. Using Teststand, simulations of a circuit can be parameterized and executed using a Python API which makes parameter sweeps and the analysis of the simulation data more convenient.
With the DLS 3 neuron testbench it is possible to simulate up to 4 neuron circuits with up to 8 synapses each. As the simulation time increases with the number of components, only the circuits that are necessary for verification are simulated in full detail. To reduce simulation time, the digital components are replaced by behavioral models. In addition to the increasing number of components it is more time-consuming to simulate digital circuits than analog ones. This is due the fact that the simulator uses adaptive time steps. If all signals change slowly in time, the time step is chosen to be large and the simulation time is short. However, digital circuits such as the digital reset mechanism have a clock signal, which changes quickly at every clock cycle. Therefore, there is always at least one fast signal which prevents the increasing of time steps.
Figure 23 shows a detailed view of the simulated components. All neuron circuits except for the digital reset mechanism are fully simulated on transistor level. The circuit which pulls the membrane towards the reset potential is analog and part of the simulation, the digital counters which determine the length and which trigger the beginning and end of the refractory period are implemented as behaviorals. The neurons receive their voltage and current parameters, e.g. the bias currents for the OTAs, from the capacitive memory (Hock, 2014). It is replaced by a behavioral, designed by Matthias Hock, which is shown in

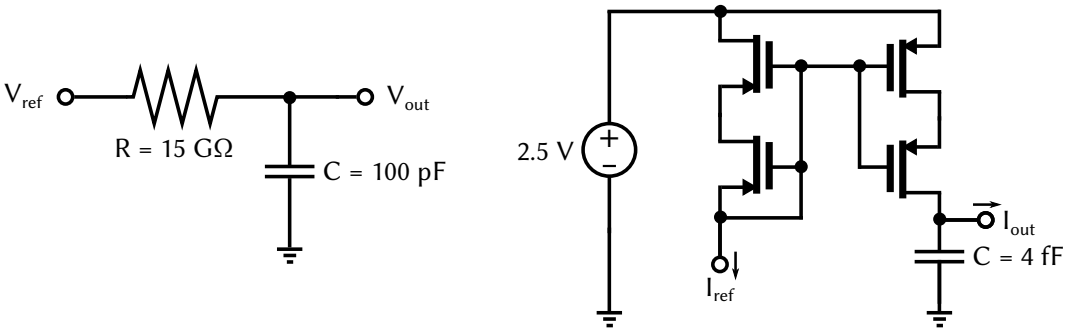

Figure 22: Behavioral circuits for the capacitive memory, designed by Matthias Hock. The circuit on the left is used to replace the voltage cells, the circuit on the right models the current cells.

figure 22. The voltage cells are modeled by an ideal capacitor and is charged over an ideal resistor, which stands in place for the periodic update of the capacitance by the refresh mechanism. The current cells are modeled as a current mirror with the output transistors of the same dimensions as in the original circuits (Hock, 2015). The SRAM, where the digital parameters for the circuits are stored, is initialized in the beginning of the simulations. The process of writing the SRAM is omitted.

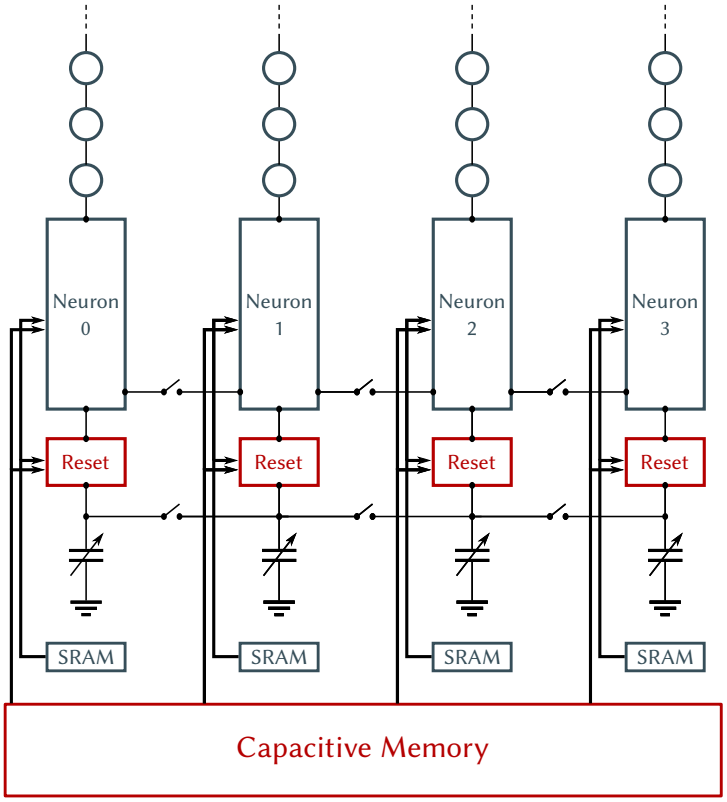

Figure 23: Simulation setup for the HICANN DLS 3 neuron. 4 Neuron circuits (gray boxes) including the multi-compartment features, with up to 8 synapses (gray circles) each can be simulated. The digital reset mechanism and the capacitive memory are replaced by behavioral to reduce simulation time (red boxes) 8 synapses (gray circles) each can be simulated. The digital reset mechanism, the capacitive memory and the SRAM are replaced by behavioral to reduce simulation time (red boxes). Adapted from Kriener et al. (2017).

5 Circuit Verification

Before high-level experiments, such as the ones described in section 3, can be performed, a general verification of the circuit functionality is necessary. The neuron circuit simulations described in section 4.2 are a valuable tool for this task. During the design process of individual components like the adaptation term, usually only the subcircuit in question is simulated. This might allow interdependencies of one subcircuit onto another stay undiscovered. A simulation of the full neuron design can uncover such effects while testing the circuits in situations similar to the intended use cases. In the following, the verification results of the new circuits in HICANN DLS 3, in particular, the discovered errors and the implemented solutions, are presented.

5.1 Transmission Gates

Many subcircuits of the HICANN DLS 3 neuron can be individually enabled or disabled. This is done by disconnecting them from the membrane capacitor via transmission gates that are controlled by digital neuron parameters.

From a high-level point of view, a transmission gate is a switch. The digital parameter controlling the switch is stored in the neuron SRAM. Compared to a single pass transistor, a transmission gate offers a high conductance over the whole voltage range. This is useful it is expected that voltage differences over the whole voltage range will appear.

A drawing of the transmission gate circuit is shown in figure 24. The transmission gate consists of a pMOS and an nMOS transistor that are connected in parallel. The nMOS transistor is conducting if its gate is high (1.2 V), while the pMOS transistor is conducting when its gate is low (0 V). The output voltage of the SRAM cell is connected to the nMOS gate and the inverted output is supplied to the pMOS transistor. By setting the SRAM cell to a high value, the transmission gate is opened, while it closed in the other case.

In the neuron circuit, some voltages can rise above 1.2 V, which is used to supply the thin-oxide transistors in the transmission gates. If this happens, the gate-source voltage of the pMOS transistor in the transmission gate becomes positive, leading to an increased conductance. This results in a leakage current that can disturb the neuron behavior and the emulation results.

For some experiments, it might be necessary to have longer time constants than the leak OTA settings allow. This can then be achieved by merging the adaptation capacitance with the membrane capacitance. This configurable connection between both capacitors was implemented with a 1.2 V transmission gate in the first version of the circuit. However, both voltages can rise above 1.2 V therefore potentially resulting in a leakage current between the adaptation circuit and the membrane as shown in figure 25. This is especially problematic as such behavior can occur during normal operation of the adaptation circuit and the only option to prevent it with certainty is to switch off adaptation.

Another site where the leakage occurs is at the voltage readout. In HICANN DLS 3 the membrane, synaptic input, or the adaptation voltage can be read out. The voltages that are not read out are separated from the readout circuits by transmission gates. The leakage is

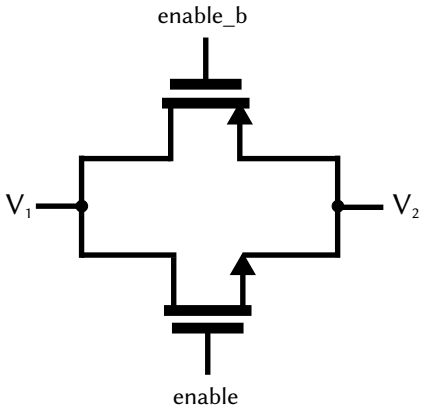

Figure 24: Circuit of a transmission gate which is controlled by the digital parameter `enable`. The transmission gate separates the voltages V_1 and V_2. The pMOS transistor at the top is conducting if its gate voltage (here `enable_b`, which is the inverted enable signal) is low. The nMOS transistor is conducting if its gate is high. Therefore, the transmission gate is open, if `enable` is high, and closed if it is low.

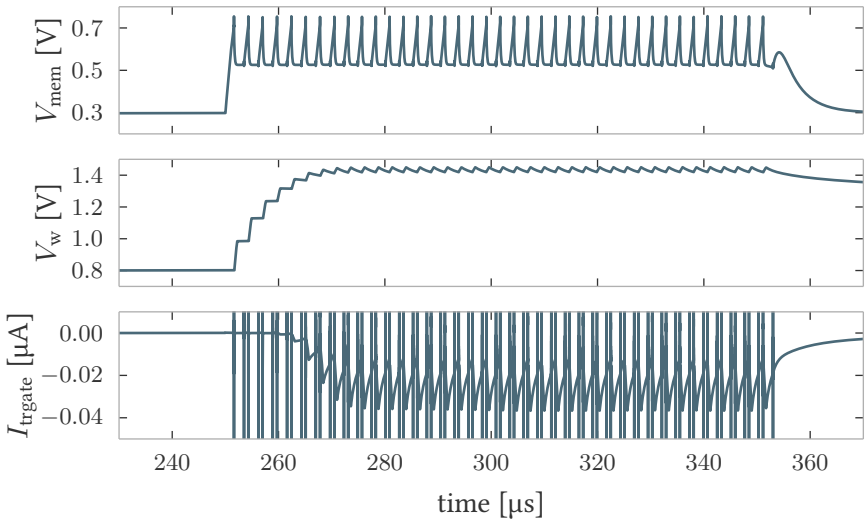

Figure 25: Demonstration of the leakage through the transmission gate connecting the membrane capacitance and the adaptation capacitance. The circuits are configured to cause a "worst-case" scenario, by switching off the adaptation but enabling the spike-triggered adaptation. The spike-triggered adaptation causes V_{w} to rise above 1.2 V. As soon as this threshold is reached, a current I_{trgate} flows through the closed transmission gate. The short, sharp peaks on I_{trgate} are caused by the switching between leak and reset mode of the leak/reset OTA when the neuron spikes.

undesirable as it distorts the measurement results. This effect can also be observed in chip measurements on HICANN DLS 2 (Stradmann, 2017).
As the problem was observed in the neuron simulations before the tapeout of the chip, the transmission gates were changed in order to prevent the leakage. At the sites where the voltage on either side of a transmission gate can rise above 1.2 V the transistors in the transmission gates were replaced by thick-oxide transistors with a logic level of 2.5 V. The thick-oxide transmission gates require that the high value of the digital parameters is changed to 2.5 V as well. This necessitates the implementation of additional level shifter circuits. As no voltage in the circuit can rise above the highest supply voltage (2.5 V), the leakage can no longer occur.

5.2 Adaptation

5.2.1 Connection to Capacitive Memory

The adaptation circuit can be configured using multiple analog parameters provided by the capacitive memory. One of the parameters shown in figure 7 is the leak potential E_L which is used as one of the inputs for the adaptation OTA. However, it is not supplied by the same storage cell in the capacitive memory as the parameter which is used as leak potential in the other parts of the neuron circuit. Therefore it will be called $V_\mathrm{leak,adapt}$ in the following. It is convenient to have an additional parameter for the leak potential in the adaptation term, as it allows, for example, to compensate the offset of the adaptation OTA by adding an offset to $V_\mathrm{leak,adapt}$. As figure 7 shows, the adaptation OTA was extended from the circuit implemented in the HICANN chip by a sign switching mechanism. The mechanism works by switching the two inputs between the negative and positive input of the OTA. Each of the two multiplexer circuits consists of two transmission gates. They, therefore, show the same leakage problem as described in section 5.1. In this case, the leakage causes the parameter stored in the capacitive memory to change. This is due to the fact, that the voltage parameter is stored on a capacitor and the local parameter cells of the capacitive memory are not buffered. A leakage current onto that capacitor therefore changes the voltage.
Figure 26 shows the effect of the leakage. If the adaptation voltage V_w rises above 1.2 V the parameter $V_\mathrm{leak,adapt}$ starts to change its value. This changes the behavior of the circuit severely as the adaptation current is calculated as

$$I_\mathrm{adapt} = a \cdot \left(V_\mathrm{leak,adapt} - V_\mathrm{w}\right) \tag{5.1}$$

which means that a $V_\mathrm{leak,adapt}$ that approaches the adaptation voltage causes a breakdown in the adaptation current. The problem was solved by introducing thick-oxide transmission gates, similar to the solution described in section 5.1.

5.2.2 Spike-Triggered Adaptation

Spike-triggered adaptation is realized by a current pulse onto the adaptation capacitance C_a, which causes a stepwise increase in the adaptation voltage. In the original design,

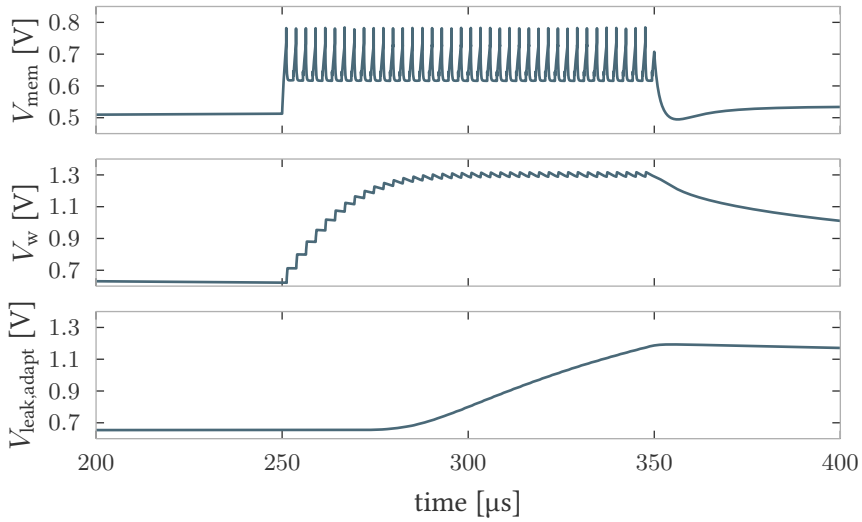

Figure 26: Top: Membrane voltage of a neuron which is stimulated with a constant current. Center: Adaptation variable V_{w} of the neuron above. As spike-triggered adaptation is enabled, the voltage is increased with every spike. At approximately $t = 280\,\mu\mathrm{s}$ the voltage rises above 1.2 V. Bottom: Voltage parameter for the leakage potential provided by the capacitive memory. When V_{w} crosses 1.2 V the current leaks onto the capacitance in the capacitive memory, where the voltage parameter is stored. This causes the $V_{\mathrm{leak,adapt}}$ to follow the voltage of V_{w}.

it was planned to use, in analogy to the HICANN design, the fire pulse of the neuron to time the current pulse of the spike-triggered adaptation. However, the fire pulse is only $t_{\mathrm{pulse}} = 8\,\mathrm{ns}$ long. The strongest current that can flow onto the adaptation capacitance during the fire pulse is, by design, $I_{\mathrm{adapt,w}} = 1\,\mu\mathrm{A}$. This yields a maximum increase of V_{w} of

$$\Delta V_{\mathrm{w}} = \frac{I_{\mathrm{adapt,w}} \cdot t_{\mathrm{pulse}}}{C_{\mathrm{a}}} = 4\,\mathrm{mV} \tag{5.2}$$

with an adaptation capacitance of $C_{\mathrm{a}} = 2\,\mathrm{pF}$. In a typical experiment, the value of V_{w} can vary over a range of 0.5 V. A voltage step of 4 mV is therefore too small, as it is much smaller than other changes that occur during an experiment (for a more detailed treatment of the appropriate range for ΔV_{w} see section 6.2.2). To allow larger ΔV_{w} the pulse length was chosen to be dependent on another digital signal. The adaptation pulse is, as the fire pulse, generated by the digital neuron back-end (developed by Gerd Kiene) and can have pulse lengths of up to 14 µs. However, it can not be longer than the refractory time of the neuron. With the longer adaptation pulse, ΔV_{w} that cover the whole range of V_{w} are theoretically possible (although not used in practice for the reasons described in section 6.2.2).

The left plot in figure 27 shows that for $I_{\mathrm{adapt,w}} > 0.2\,\mu\mathrm{A}$ the steps ΔV_{w} that can be calculated with equation (5.2) do not match the measured steps in ΔV_{w}. This can be explained by the input part of the spike-triggered adaptation circuit depicted in figure 28. The topmost

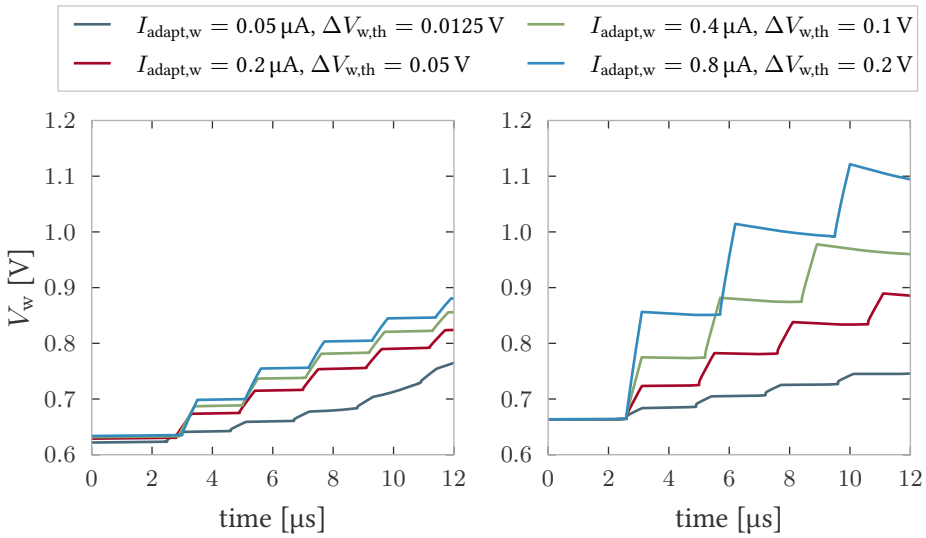

Figure 27: Spike-triggered adaptation before (left) and after (right) extending the range of possible current pulse strengths. For large $I_{adapt,w}$ the increasing steps in V_w are much smaller than the expected values of $V_{w,th}$ which are calculated from the pulse length and the set value of $I_{adapt,w}$. In the improved circuit on the right, the parameter $I_{adapt,w}$ controls the strength of the spike-triggered adaptation more reliably and the expected values for $\Delta V_{w,th}$ approximately match the observed steps in V_w.

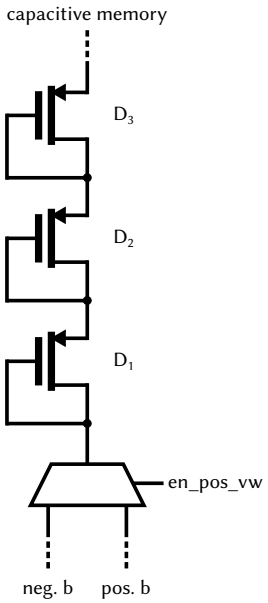

Figure 28: Part of the spike-triggered adaptation circuit. The three transistors operating as diodes D_1 to D_3 separate the circuit from the capacitive memory which provides the parameter $I_{adapt,w}$. The MUX at the bottom chooses, depending on `en_pos_vw` whether the circuit for positive or negative spike-triggered adaptation is provided with $I_{adapt,w}$.

transistor is connected to the capacitive memory which provides the parameter $I_{\text{adapt,w}}$. In order to prevent the adaptation voltage to rise up to values close to 2.5 V the three transistors labeled D_1, D_2 and D_3 at the top are connected as diodes. A diode is a device that blocks current in one direction and allows current to flow in the other direction as long as a specific voltage drop is ensured. For silicon-based semiconductors this drop is approximately 0.7 V. If the voltage drop is below the threshold only minor leakage currents can flow through the diode. The output stage of the capacitive memory is a thick-oxide pMOS transistor with its source connected to 2.5 V. For V_{w} being larger than 0.4 V, which is usually the case, a voltage drop of three times 0.7 V can therefore not occur. This limits the current $I_{\text{adapt,w}}$ that can be provided to a voltage-dependent leakage current.
This problem was observed before the tapeout of the chip and the circuit was adapted. The topmost diode was removed, which reduces the summed voltage drop over the diodes and allows more than a leakage current to flow. The right plot in figure 27 shows the same experiment as before, but with the improved circuit. We see that for high desired values of $I_{\text{adapt,w}}$ the voltage drop on V_{w} matches the theoretically calculated values much better than before. Although, for rising V_{w} the measured ΔV_{w} still decreases, as with decreasing difference between V_{w} and 2.5 V the current that can flow through the diodes decreases.

5.2.3 Sign Switches in Adaptation

In comparison to the HICANN circuit, the adaptation on HICANN DLS 3 allows for a switchable sign for the adaptation parameters a and for the spike-triggered adaptation b. The sign of the adaptation parameters is configured by the digital parameters `en_neg_va`, which switches the sign of a, and `en_pos_vw`, which allows to choose the sign of the spike-triggered adaptation. As described in section 2.2.2 the spike-triggered adaption is realized by a voltage step on V_{w} which is then converted in a current step using the adaptation OTA. Therefore, the configuration of the `en_pos_vw` switch for a chosen sign of the spike-triggered adaptation must depend on the sign which was chosen for the adaptation by the setting of `en_neg_va`.
Figure 29 shows a demonstration of the effect of the parameters mentioned above. In each of the four experiments, the neuron is stimulated with a constant current and the membrane voltage, the adaptation voltage and the adaptation current are recorded. For `en_neg_va=False` and `en_pos_vw=True` (top left) we see positive adaptation and positive spike-triggered adaptation. The neuron spikes in the beginning, then the spike frequency decreases due to the rising adaptation current and the neuron stops firing. For `en_neg_va=False` and `en_pos_vw=False` we see that the positive adaptation and the negative spike-triggered adaptation counteract each other. This is also the case for `en_neg_va=True` and `en_pos_vw=False` with negative adaptation but positive spike-triggered adaptation (bottom right). In the recordings at the bottom left the digital parameters are set to `en_neg_va=True` and `en_pos_vw=True`. The adaptation is negative, which causes a positive feedback when the neuron rises above its leak potential. The increasing adaptation current allows the neuron to fire towards the end of the stimulating current step. As the spike-triggered adaptation is negative as well, the spike frequency

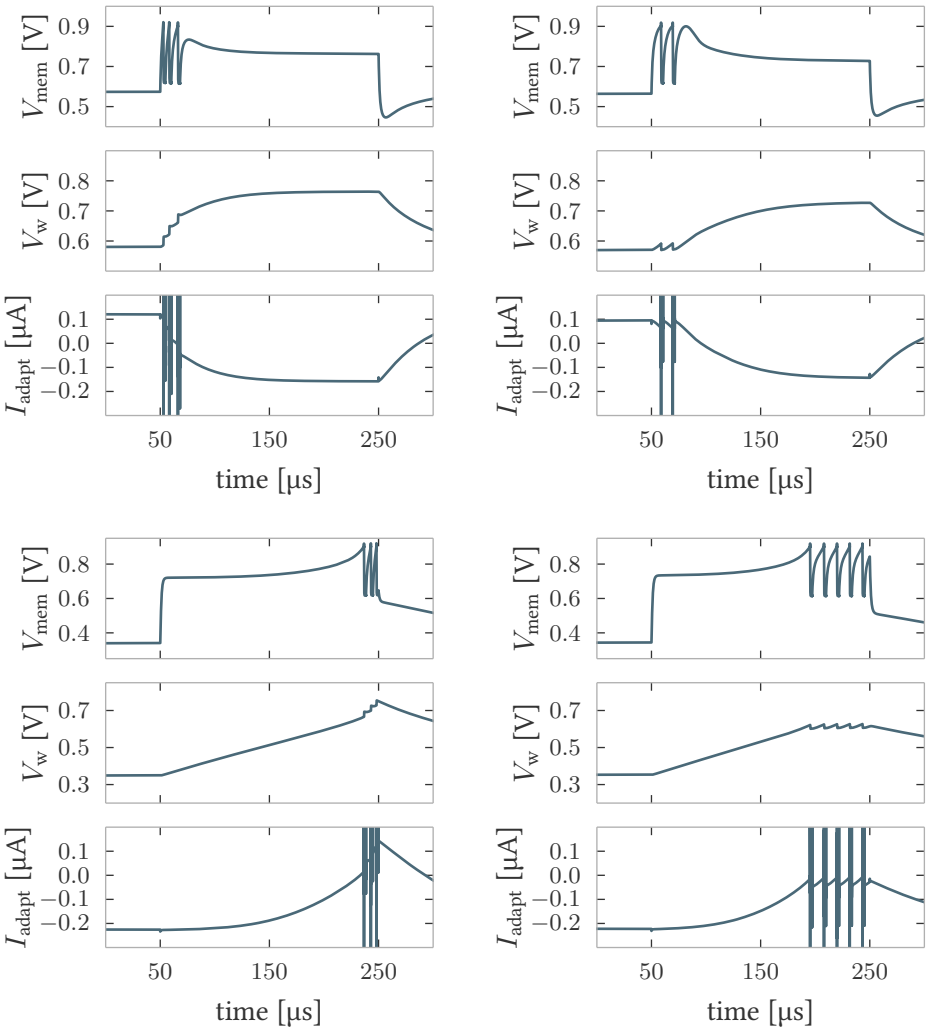

Figure 29: Simulations demonstrating the functionality of adaptation sign switches. A neuron is stimulated with a constant current and the membrane potential, the adaptation voltage and the adaptation current are measured.
Top left: The digital parameters controlling the sign of the adaptation are set to `en_neg_va=False` and `en_pos_vw=True`. This enables positive spike-triggered adaptation and positive adaptation. Top right: Here the parameters are set to `en_neg_va=False` and `en_pos_vw=False`. The spike-triggered adaptation is negative while the adaptation is positive. Bottom left: The digital parameters are set to `en_neg_va=True` and `en_pos_vw=True`. The adaptation, as well as the spike-triggered adaptation, are negative. Bottom right: The digital parameters are set to `en_neg_va=False` and `en_pos_vw=False`. This setting allows negative adaptation with positive spike-triggered adaptation. The short, sharp peaks on I_{adapt} are caused by the switching between leak and reset mode of the leak/reset OTA when the neuron spikes.

increases. The simulation indicates that combinations of the digital sign switches show the intended behavior.

5.3 Exponential Term

In addition to the sign switches in the adaptation circuit, an exponential term circuit with a three bit digital weight parameter `exp_weight_b` extends the LIF neuron of HICANN DLS 2. To test and illustrate the functionality of the exponential term circuit the neuron is stimulated with multiple step currents of different amplitude. Figure 30 shows the output current of the exponential circuit as a function of the membrane voltage. We see that for membrane voltages below $V_{\text{mem}} = 1.0\,\text{V}$, I_{exp} is approximately an exponential function of the membrane voltage (for detailed evaluation see section 6.3). For larger voltages, the exponential current drops because the transistor in the output stage of the circuit leaves the saturation as its drain-source voltage decreases towards zero. An exponential current of nearly zero for a membrane voltage of $V_{\text{mem}} \geq 1.2\,\text{V}$ is intended, to prevent the membrane voltage from rising far above 1.2 V.
The measurement also shows that the weight parameters modulate the strength of the exponential current. The weight is given as a tuple of the negated boolean values and therefore the setting `exp_weight_b=(True, True, True)` is the weakest setting, which switches off the term, while `exp_weight_b=(False, False, False)` yields maximum strength. The first value in the tuple is the least significant bit. For a more detailed investigation of the influence of the weight parameter see section 6.3.

5.4 Bypass Mode

The bypass mode is a debug and test feature that was introduced in HICANN DLS 1. It allows to bypass the neuron circuits and produce a fire signal of the neuron every time an input arrives at the neuron. This allows, for example, to test spike routing from neuron to neuron, external spike input and spike readout independently from the neuron circuits.
Using the neuron simulation setup two errors in the bypass circuit were discovered: At first, the output signal of the bypass was inverted, due to changes introduced in the interface between the neuron and the new digital reset mechanism. This error could be easily solved by adding an additional inverter in the signal path.
The second problem arose because the bypass circuit did not change from HICANN DLS 2 to HICANN DLS 3 although the spike signals did. The synaptic input, as well as the bypass of a neuron, see an incoming spike event as a current pulse. On HICANN DLS 2 the pulse length was determined by the FPGA, which was used to route spike signals, and was at least 10 ns long. On HICANN DLS 3 the synapse drivers produce pulses of the length of maximally 4 ns. The signal with the reduced pulse length is too weak to trigger a response of the bypass circuit. Figure 31 shows that this can be circumvented to a certain degree by triggering the bypass using bursts of input spikes with the minimal interspike interval that can be handled by the synapse drivers of 8 ns. We see that in the typical case (*tt* corner) the bypass requires four spikes in short succession to trigger one fireout signal. In the

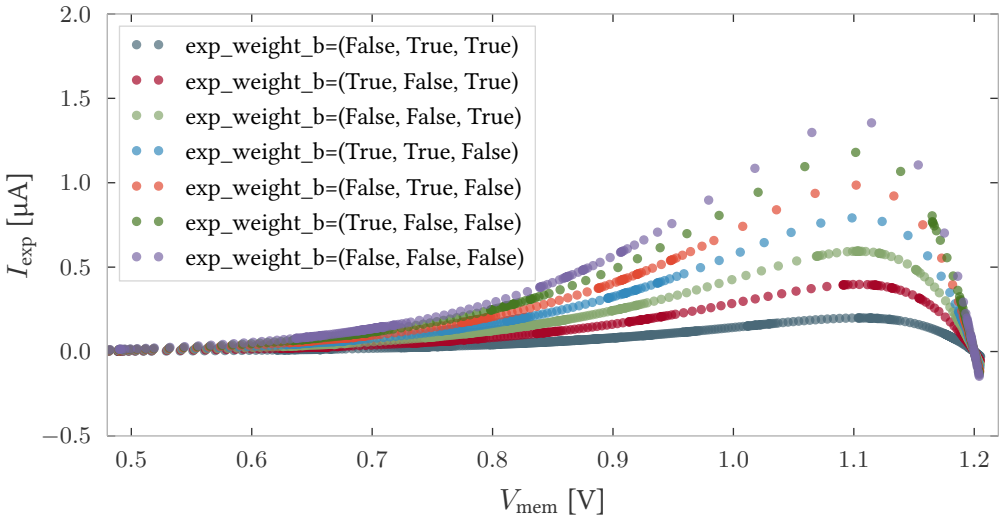

Figure 30: Demonstration of the scaling parameter `exp_weight_b` of the exponential term. For every possible parameter setting (except `(True, True, True)` which is equivalent to switched off) the neuron was stimulated with a series of step currents of different strengths and the membrane voltage and exponential currents were recorded. The plot of I_{exp} over V_{mem} shows that for all settings the current increases with the membrane voltage until approximately $V_{\text{mem}} = 1.0\,\text{V}$ the current starts to drop.

ss corner seven spikes per burst are needed, in the *ff* corner two spikes are sufficient to trigger the bypass.

This issue was observed shortly before the tapeout of the chip and there was not enough time for the circuits to be changed.

5.5 Plateau Potentials and Multi-Compartment Circuits

The multi-compartment features described in section 2.2.2 allow connecting the neuron circuits to a multi-compartment neuron using intercompartment conductances. The left part of figure 32 shows simulation results of a simple multi-compartment neuron consisting of two compartments. One compartment receives synaptic input and shows a PSP and, for stronger input, a spike. The membrane voltage of the second compartment passively follows the membrane of the first compartment. As the multi-compartment circuit acts as a low-pass filter, the fast features of the membrane of the active compartment are not visible on the second compartment.

The right part of figure 32 shows a test of the combination of the plateau potential mechanism in combination with the multi-compartment connections. The active compartment is configured to produce a plateau potential if a spike is triggered. The plot shows three different experiments where the length and the amplitude of the plateau potential as well as the strength of the intercompartment conductance are varied.

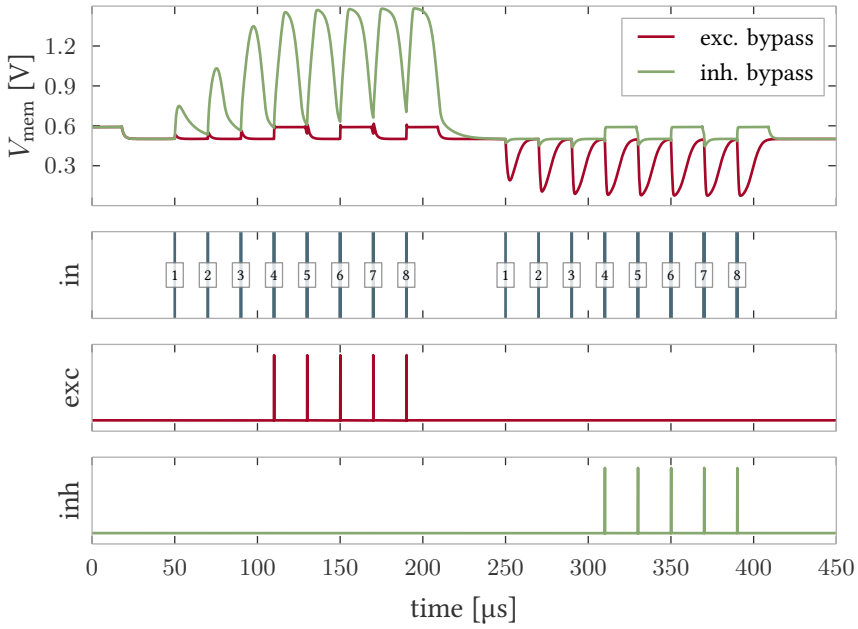

Figure 31: Simulation of bypass functionality. A neuron is stimulated with spike bursts of increasing size first via the excitatory input then via the inhibitory input. The membrane voltage (top) and the fire signals (bottom) are recorded with excitatory bypass enabled (red) and inhibitory bypass enabled (green). The input spikes are shown in blue (center). We see that the bypass circuits require a burst consisting of four spikes to emit a fire signal.

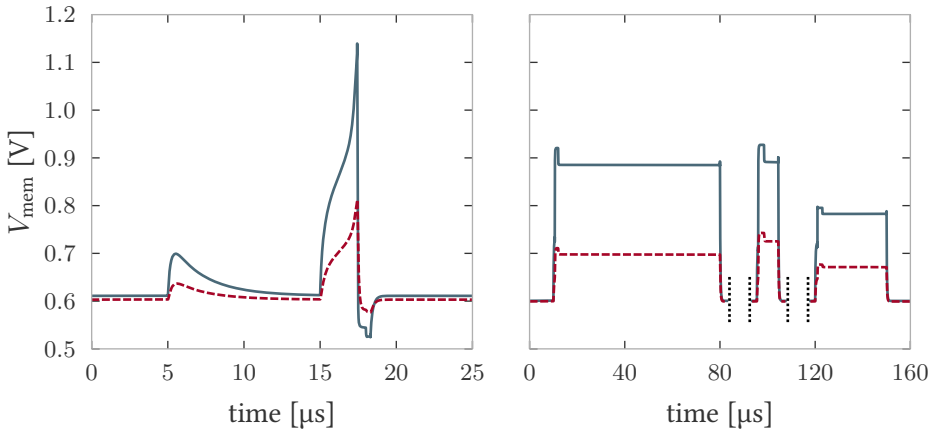

Figure 32: Demonstration of multi-compartment and plateau potential functionality. Left: One neuron circuit is stimulated with spike input. The second input is strong enough to trigger a spike. A second neuron circuit is connected to the first using the multi-compartment circuits. Its membrane voltage passively follows the voltage of the first neuron circuit. Right: The reset of the first neuron is now configured for plateau potentials. The experiment is repeated 3 times with different plateau potential lengths and heights. Adapted from Schemmel et al. (2017).

These tests show that the multi-compartment circuits are functional and that a variation of the parameters, controlling the height of the plateau potential and the intercompartment conductance, produce the expected change in behavior.

6 Characterization

In addition to the verification process, which, in general, tests whether circuits are functional and behave as intended, the characterization process addresses more detailed questions. How do the circuits react to changes in parameterization? In what range can the subcircuits of the neuron be configured? How does the parameterization correspond to parameters of the mathematical neuron model? Are the available parameter ranges in accordance with intended use cases? Do the results above change in a critical way in the process corners? In the following, characterization results for the new reset mechanism, the adaptation and the exponential term are presented.

6.1 Reset Current

For the emulation of plateau potentials using a long reset with a high reset voltage the implemented reset mechanism changed in comparison to previous chip versions. The previous requirements for the reset mechanism however are still valid. The membrane voltage must be pulled towards the reset voltage and kept at this voltage for the duration of the refractory period. This requires at least that the current produced by the reset mechanism must be larger than the sum of all other currents pulling the membrane voltage away from the desired reset potential. These other currents can be the adaptation current, the exponential current, intercompartmental currents and synaptic input currents. If the reset current does not exceed the other currents, it is not able to pull the membrane voltage towards the reset potential and it can happen that the neuron stays above the threshold, which then results in continuous firing of the neuron. If the reset current is only slightly larger than the sum of the opposing currents, the membrane potential moves slowly towards the reset potential. This can be a problem if the refractory time is too short to reach the desired voltage, because then effectively the reset potential of the neuron is changed.
The first of the possible outcomes of an insufficiently strong reset, the unintended continuous firing, is a severe problem, as it completely changes the behavior of the neuron. In practice it is therefore desirable to have the highest possible reset current that can be produced with the available settings. The right plot in figure 33 shows the maximum of the available reset current for different combinations of reset potentials and bias parameter settings for the leak/reset OTA. For each data point we determine the reset current at $V_{\text{mem}} = 1.4\,\text{V}$ because, as figure 33 shows, the current in this region is nearly independent of the membrane voltage. We see that the bias current, for which the highest current is produced, changes with the reset potential. Additionally, for $V_{\text{reset}} > 0.89\,\text{V}$, the highest possible current is below $3\,\mu\text{A}$. This is can lead to a non-functional reset, as the exponential term alone is able to produce currents of more than $1\,\mu\text{A}$. As this only occurs, if the reset potential is high, it is mainly relevant only for the plateau potential use case. The easiest way to avoid continuous firing in this case is to not use the exponential term and, if necessary, limit the maximal current produced by the synaptic input.
Figure 33 is simulated for the *tt* case, figure 34 shows the same simulation but in the other process corners. We see that in the *ss* and *fs* corner the currents are significantly weaker than in the typical case, while in the other corners, the currents are much higher. The slow

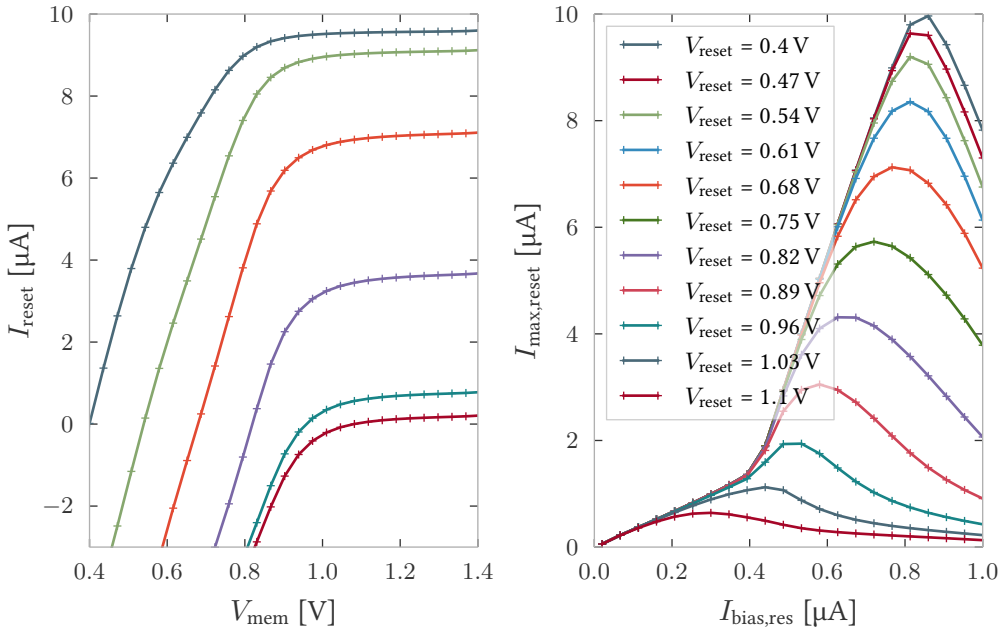

Figure 33: Left: Reset current over membrane voltage for different settings of V_{reset}. The bias parameters are set to $I_{\text{bias,res}} = 0.8\,\mu\text{A}$ and $I_{\text{bias,sd,res}} = 1.0\,\mu\text{A}$. The colors correspond to the legend on the right. Right: Maximum of the reset current that can be provided by the leak/reset OTA depending on the setting of $I_{\text{bias,res}}$. The source degeneration bias $I_{\text{bias,sd,res}}$ is set to 1 µA. The current values at $I_{\text{bias,res}} = 0.8\,\mu\text{A}$ correspond to the values at $V_{\text{mem}} = 1.4\,\text{V}$ in the left plot. For rising values of the reset voltage V_{reset}, the highest achievable current drops below 1 µA.
Data provided by Paul Müller.

pMOS corners show a critically low reset current for reset potentials of approximately $V_{\text{reset}} > 0.7\,\text{V}$. This can not only affect the plateau potential use case but also the "normal" reset behavior.

These simulations showed that not only the amplitude of the highest achievable reset current but also the parameter i_bias_res, for which it is produced, is different for every neuron. To avoid the effects of an insufficient reset a calibration mechanism must be implemented that ensures the optimal reset setting for each neuron. Additionally, a circuit change could be implemented that allows to disable the exponential term during the reset. This significantly reduces the current that has to be produced by the reset OTA.

6.2 Adaptation

6.2.1 Adaptation Conductance

The adaptation circuit on the HICANN DLS 3 (see figure 6 and figure 7) is controlled by the digital parameters described in section 5.2.3 and by the current parameters i_adapt_w controlling the spike-triggered adaptation, i_bias_adapt, i_bias_adapt_sd and i_bias_ad-

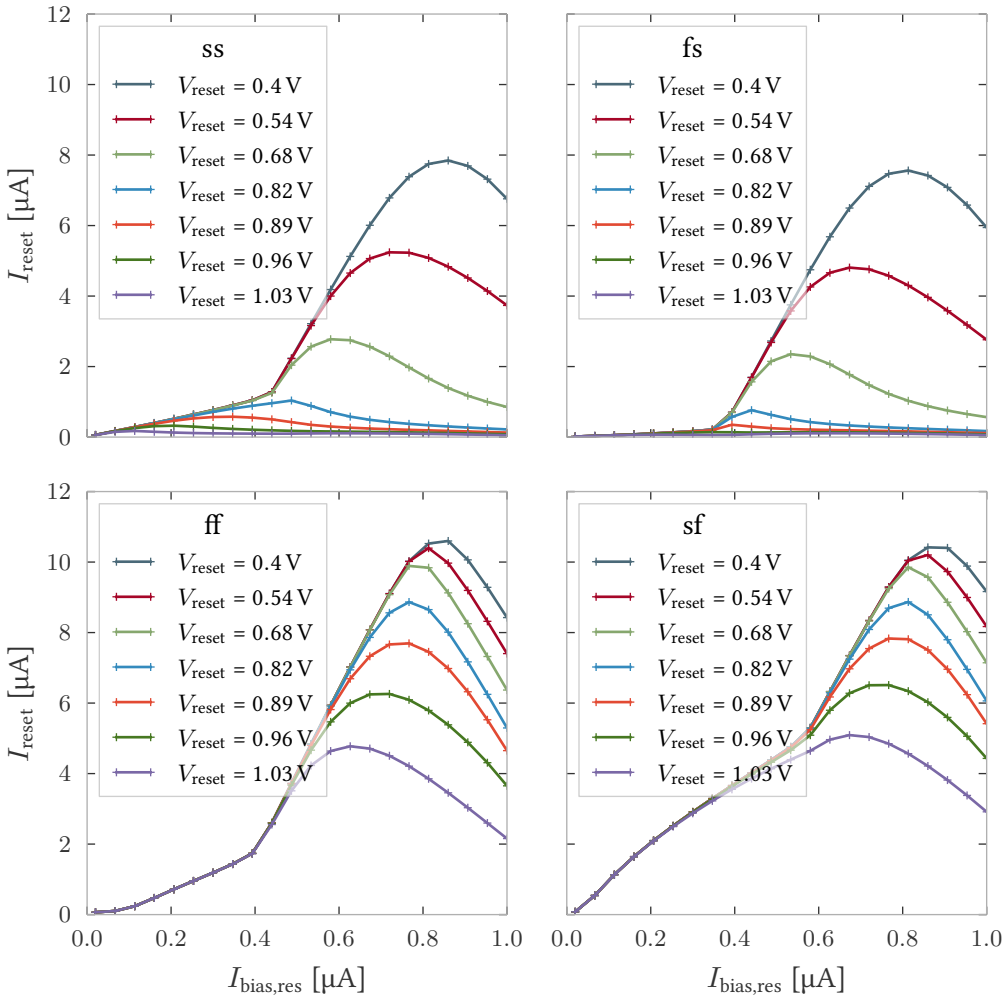

Figure 34: Maximum of the reset current that can be provided by the leak/reset OTA depending on the setting of $I_{bias,res}$ in the process corners. The source degeneration bias $I_{bias,sd,res}$ is set to 1 µA. For rising values of the reset voltage V_{reset}, the highest achievable current drops below 0.5 µA in the slow pMOS corners.
Data provided by Paul Müller.

apt_res. i_bias_adapt_res controls the adaptation time constant and is treated in section 6.2.3. i_bias_adapt and i_bias_adapt_sd are the bias currents for the adaptation OTA, which models the adaptation conductance a. The source degeneration bias can be configured per neuron while the bias current is a global parameter, i.e. it is shared between all neurons.

Figure 35 shows the dependency of a on the source degeneration bias of the adaptation OTA. In the simulation a neuron was stimulated with a step current starting at t_1 and ending at t_2 then the experiment was repeated with the adaptation switched off. The signals $V_{\text{mem, a}}$, $V_{\text{w, a}}$ and $I_{\text{adapt,a}}$ for the simulation with enabled adaptation and V_{mem}, V_{w} and I_{adapt} for disabled adaptation were recorded. As a is a conductance it can not be measured directly in the simulation. To calculate its value from the recorded current and voltage signals two different methods were used.

The first method derives a formula for a assuming an ideal AdEx neuron (equation (2.9)) that requires only the signals $V_{\text{mem, a}}$ and V_{mem}. We use the fact that, if the applied current step is long enough, the membrane voltage reaches a steady value during its application. It was ensured that the steady state is reached during the current step, by choosing a short adaptation time constant. This means that the applied current stimulus I_{stim} and the leak and adaptation current cancel each other.

$$I_{\text{stim}} = -g_{\text{L}} \cdot (V_{\text{mem,a}} - E_{\text{L}}) - a \cdot (V_{\text{w,a}} - E_{\text{L}}) \tag{6.1}$$

As the adaptation voltage slowly follows the membrane voltage we can assume

$$(V_{\text{mem,a}} - E_{\text{L}}) = (V_{\text{w,a}} - E_{\text{L}}) = \Delta_a. \tag{6.2}$$

When the adaptation is switched off, the stimulating current is canceled only by the leak current.

$$I_{\text{stim}} = -g_{\text{L}} \cdot (V - E_{\text{L}}) \tag{6.3}$$

We write $(V - E_{\text{L}}) = \Delta$. As the stimulating current is the same in both simulations we can calculate a from the leak conductance g_{L} and Δ, Δ_a:

$$a = g_{\text{L}} \cdot \left(\frac{\Delta}{\Delta_a} - 1 \right) \tag{6.4}$$

The second method estimates the transconductance of the adaptation OTA with a difference approximation using the adaptation current and adaptation voltage before the current stimulus $I_{\text{adapt,1}}$, $V_{\text{w,1}}$ and during the current stimulus $I_{\text{adapt,2}}$, $V_{\text{w,2}}$

$$I_{\text{adapt,1}} = a \cdot (V_{\text{w,1}} - E_{\text{L}}) \tag{6.5}$$

$$I_{\text{adapt,2}} = a \cdot (V_{\text{w,2}} - E_{\text{L}}) \tag{6.6}$$

$$I_{\text{adapt,1}} - I_{\text{adpat,2}} = a \cdot (V_{\text{w,1}} - V_{\text{w,2}}) \tag{6.7}$$

$$\tag{6.8}$$

which yields

$$a = \frac{I_{\text{adapt,1}} - I_{\text{adapt,2}}}{V_{\text{w,1}} - V_{\text{w,2}}}. \tag{6.9}$$

In this simulation $V_{\text{w,1}} = 0.62\,\text{V}$ and, depending on the setting of i_bias_adapt_sd, $0.63\,\text{V} \leq V_{\text{w,2}} \leq 0.8\,\text{V}$.
As we see in figure 35, both methods yield results that differ by approximately 10 %. This can be explained by the fact, that the first method assumes an ideal AdEx neuron with a constant leak conductance. The transconductance of the leak OTA however, is not perfectly constant over the range of the used membrane voltages, which influences the results calculated with the first method. The first method has the additional disadvantage of requiring two measurements, one with and one without adaptation. However, it is still better suited for actual chip measurements since the second method requires a simultaneous recording of I_{adapt} and V_{w} while the first method relies on the membrane potential only.
The results for the parameter range of a on HICANN DLS 3 can be transformed into the biological domain by

$$a_{\text{bio}} = \frac{C_{\text{bio}}}{C_{\text{hw}}} \cdot \frac{1}{\alpha_{\text{t}}} \cdot a_{\text{hw}} \tag{6.10}$$

with $\alpha_{\text{t}} = 1000$. Assuming $C_{\text{bio}} = 200\,\text{pF}$ which is a value used in the AdEx firing patterns (see section B.1) and the full membrane capacitance on hardware $C_{\text{hw}} = 2.36\,\text{pF}$ we obtain a maximum biological adaptation conductance of approximately $a_{\text{bio}} \approx 805\,\mu\text{S}$. This is nearly a factor of 50 larger than the largest conductance required for the AdEx firing patterns. However, figure 35 also shows that much smaller values for a, as for example needed for the firing patterns, can be achieved by using i_bias_adapt_sd < 250 nA.
Theoretically, the adaptation current is independent of the membrane voltage V_{mem}. However, as figure 36 shows, this is not the case for all situations. For this simulation the membrane and adaptation voltages were clamped to fixed values and the adaptation current was measured. The simulation was performed with a comparatively low adaptation conductance using a source degeneration bias of 250 nA. We see that for low membrane voltages the output current of the OTA is not only dependent on its inputs but also on the output voltage V_{mem}. Additionally, we notice that for adaptation voltages of $V_{\text{w}} > 1.0\,\text{V}$ the adaptation current no longer rises with increasing V_{w} which is expected due to the saturation of the OTA.
Both effects observed here can alter the behavior of a neuron and should to be avoided if possible. However, the membrane voltage, below which the adaptation conductance depends on V_{mem} depends on i_bias_adapt and i_bias_adapt_sd. The same is the case for the adaptation voltages, where the adaptation OTA is in saturation. Nevertheless, figure 36 shows that, in general, low membrane voltages as well as high adaptation voltages should be avoided to stay close to the mathematical model.

6.2.2 Spike-Triggered Adaptation

The mechanism of the spike-triggered adaptation circuit shows a fundamental difference from the adaptation in the AdEx model. In the model, the spike-triggered adaptation b is completely independent of the adaptation conductance a. In the hardware implementation, however, the spike-triggered adaptation is added to the adaptation voltage rather

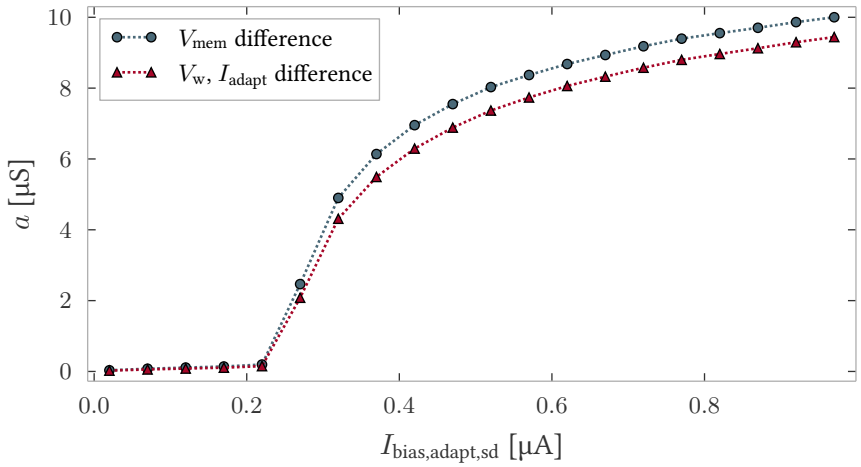

Figure 35: Dependency of the adaptation conductance a on the OTA parameter $I_{\text{bias,adapt,sd}}$. The global parameter $I_{\text{bias,adapt}}$ is set to 1 µA. To determine a, the membrane is stimulated with a step current (once with adaptation enabled and once without). Then a is determined using the differences in the membrane voltage before and during the current (blue circles) or using the increase in adaptation current and adaptation voltage before and during the current (red triangles). We see that the two methods yield increasingly different results for larger bias currents.

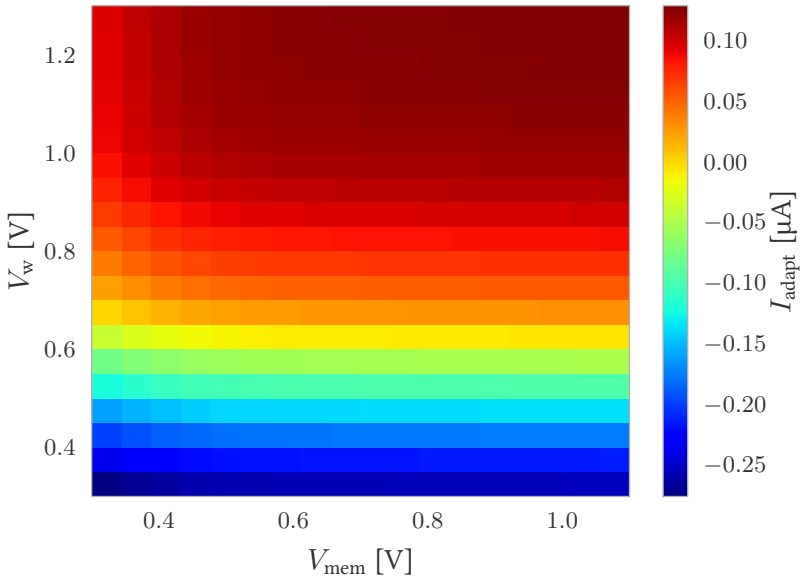

Figure 36: Two-dimensional sweep over membrane and adaptation voltage. The resulting adaptation current is depicted as color. The adaptation current should be independent of the membrane voltage, which is the case for most combinations of V_{mem} and V_{w} except for small membrane voltages. The sweep was performed with $I_{\text{bias,adapt}} = 1\,\mu\text{A}$ and $I_{\text{bias,adapt,sd}} = 0.25\,\mu\text{A}$.

than directly to the adaptation current. As the adaptation current is calculated from the adaptation voltage using the adaptation conductance, the voltage step ΔV_{w} and a_{hw} are not independent given a fixed value of b_{bio}. Assuming a certain combination of biological parameters a_{bio} and b_{bio} need to be realized on the hardware, the required hardware adaptation conductance a_{hw} can be calculated using equation (6.10). The required size of the voltage step for the spike-triggered adaptation then is

$$\Delta V_{\mathrm{w}} = \frac{b_{\mathrm{bio}}}{a_{\mathrm{bio}}} \cdot \alpha_v \tag{6.11}$$

with the transformation factor α_v typically in the range of 5 to 15. This introduces a problem for combinations of small adaptation conductances and large spike-triggered adaptation. The smaller a is, the larger the voltage step ΔV_{w} needs to be to realize the same value of b. As V_{w} has a finite range and can not rise above approximately 1.4 V it can happen, in the case of a large ΔV_{w}, that after very few spikes of the neuron, spike-triggered adaptation has pushed V_{w} to its maximum and further adaptation is not possible. For the regular bursting pattern for example $a = 2$ nS and $b = 0.1$ nA with a transformation factor of $\alpha_v = 10$ this yields a required voltage step of $\Delta V_{\mathrm{w}} = 0.5$ V. Such large steps allow at most two spikes until the maximum of V_{w} is reached.
A possibility to avoid this situation is to chose a smaller α_v. This however also scales all other voltages, which might be unwanted and additionally reduces the range of the membrane voltage. As long spike-triggered adaptation and a are implemented this way, the dependency of ΔV_{w} can not be fully avoided.

6.2.3 Adaptation Time Constant

The adaptation time constant τ_{w} can be transformed from biological into hardware domain by

$$\tau_{\mathrm{w,hw}} = \tau_{\mathrm{w,\,bio}} \cdot \frac{1}{\alpha_t}. \tag{6.12}$$

On hardware, the time constant is controlled by the parameter i_bias_adapt_res. Figure 37 shows the dependency of τ_{w} on this parameter. For this simulation the membrane potential was first clamped to the leak potential and then increased in a step of 200 mV. The adaptation voltage V_{w} follows this step on an exponential curve with the time constant τ_{w}. To determine τ_{w}, an exponential function is fitted to the measured curve (see figure 37 for an example with one setting of the bias parameter). A sweep over the range of the bias current shows that the available range of τ_{w} values goes from below 10 µs to above 200 µs (measured for a bias current of 70 nA). The data point for the lowest bias current of 20 ns is not shown here, because the increase of V_{w} is so slow for this bias setting, that V_{w} rises only 100 mV over the whole simulation time. For a reliable fit result a significant increase would be required. Although the exact value of τ_{w} could not be determined, the fact that the rise of V_{w} was so slow, shows that significantly larger time constants than $\tau_{\mathrm{w}} = 200$ µs are possible.

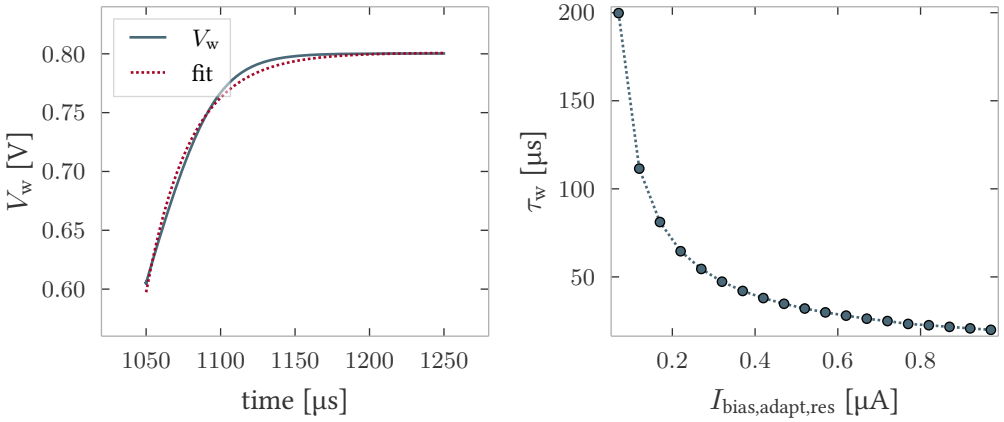

Figure 37: Measurement of τ_w depending on the parameter $I_{bias,adapt,res}$. Left: Exponential fit on the rise of the adaptation voltage V_w which is caused by a voltage step in the membrane voltage. The time constant of the rise is the adaptation time constant τ_w. We see that the rise of the adaptation voltage only approximately follows an exponential curve. Right: Sweep over the available range of $I_{bias,adapt,res}$ shows that the possible range of τ_w. The measurement result for $I_{bias,adapt,res} = 20\,\text{nA}$ is not shown here because the fit does not work reliably for the very large time constant.

As 1 µs on hardware translates to 1 ms in biology the parameter range measured here fits well to the range of adaptation time constants used in the AdEx firing patters, which range from 30 ms to 300 ms.

6.3 Exponential Term

The exponential term in the AdEx model is described by

$$I_{exp} = g_L \cdot \Delta_T \exp\left(\frac{V - V_T}{\Delta_T}\right) \tag{6.13}$$

with the two exponential term parameters Δ_T and V_T and the leak conductance g_L. The circuit on the HICANN DLS 3 has no dependency of the exponential current on the leak conductance and can be configured to eight different strengths using a 3 bit digital parameter `exp_weight_b`. The weight is given as tuple of booleans where `(True, True, True)` is the weakest setting and corresponds to a switched off exponential term. The first value of the tuple is the least significant bit.

To obtain an estimate for the parameters V_T and Δ_T, which can not be set directly, we perform a simulation as described in section 5.3 and fit equation (6.13) to the exponential current using a value for g_L which was determined in simulation before. Figure 38 shows the current produced by the exponential term plotted over the membrane voltage for the different weight values. We see that the dependency is only exponential for membrane voltages below 1.0 V. Therefore, the exponential fit is only performed over a voltage range of 0.55 V to 1.0 V. Figure 39 shows the residuals of the fits shown in figure 38. They show

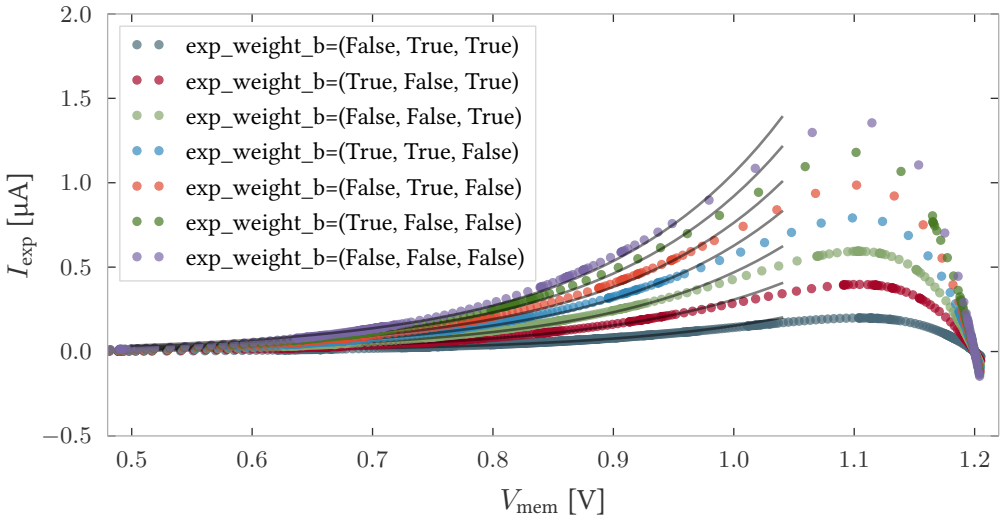

Figure 38: Exponential fit (black) to determine the AdEx parameters Δ_T and V_T in the exponential term circuit. Only data points in the voltage range between 0.55 V and 1.0 V were used as the exponential current no longer follows an exponential curve for large voltages. In this plot the fit curve intentionally has been drawn over a larger voltage range to show the increasing deviation from an exponential curve. In the chosen voltage range the dependency of I_{exp} on V_{mem} is approximately exponential, which allows to gain an approximate value for the model parameters Δ_T and V_T.

that, although the dependency of I_{exp} on V_{mem} is only approximately exponential, the fits are good enough to obtain an estimate for V_T and Δ_T.

The simulations for the different weight values were performed using the same series of step currents injected into the membrane. To investigate the influence of changes in the stimulation protocol to the fit result, the simulation is repeated 10 times, each time with a different series of step currents. The fit results for all 10 runs can be found in figure 40. In general, we see that there is a clear dependency of the parameter V_T on the weight parameter. Strong weights (low values) lead to a low V_T. This relation is nearly independent from the exact measurement protocol (for all stimulation protocols we see low V_T for strong weights and high V_T for weak weights). This is different for Δ_T. We see that Δ_T changes with the weight, but the dependency between Δ_T and the weight is different for different stimulation protocols.

We therefore conclude that there exists a relation between V_T and the exponential weight and we can estimate the relation with the described fit method. For Δ_T we can only determine a rough estimate which is not dependent on the setting of `exp_weight_b`.

6.3.1 Temperature Variations

Transistor-level simulations allow to simulate a circuit's behavior at different temperatures. This is important, since a circuit should be designed in a way, that changes in temperature of $\pm 10\,^{\circ}\mathrm{C}$ do not significantly alter its behavior. Typically, circuits are simulated

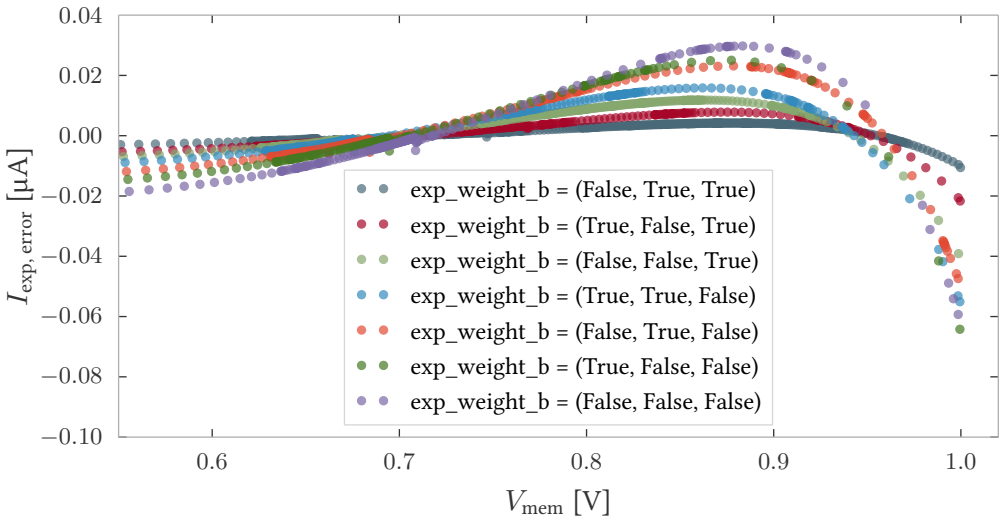

Figure 39: Residual of the exponential fit to the current produced by the exponential term performed in figure 38. For high voltages the observed residuals are below 10 % of the exponential current. At the lower end of the voltage range, for very small exponential currents, the relative deviations increase.

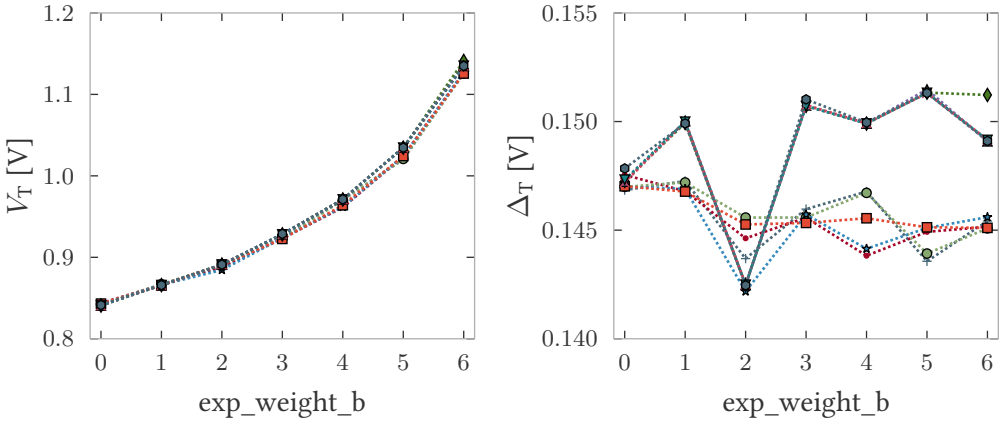

Figure 40: Estimation of the influence of the stimulation procedure. The AdEx parameters Δ_{T} and V_{T} are plotted against their weight parameter `exp_weight_b`. The tuple of booleans in `exp_weight_b` was translated into an integer (0 stands for (False, False, False), i.e. the strongest weight). The stimulation of the neuron with several current steps and the fit was repeated 10 times with different combinations of current steps to estimate the influence of changes in the stimulation protocol. We see that the exact protocol has very little influence on the resulting dependency of V_{T} on `exp_weight_b`, the relation between Δ_{T} and the weight however varies strongly.

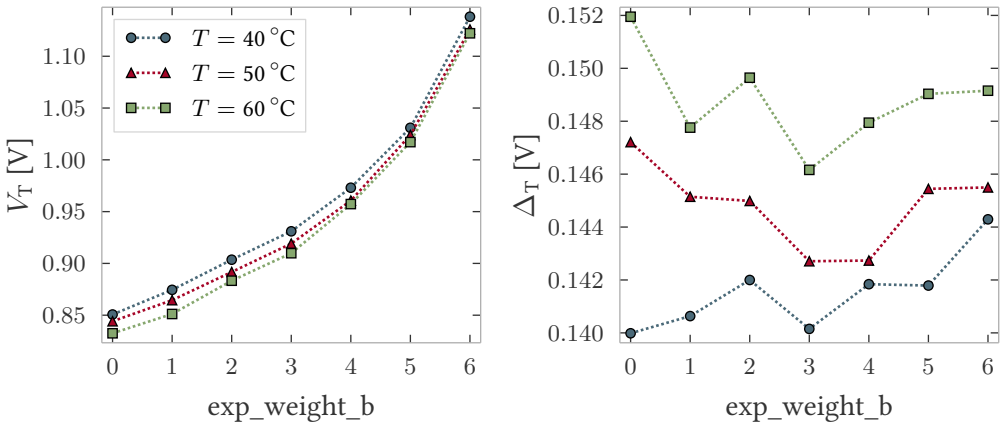

Figure 41: Investigation of temperature dependency of the exponential term circuit. The neuron was simulated at $T = 40\,°C, 50\,°C, 60\,°C$. Left: The parameter V_T depends only weakly on the temperature. Still, for lower temperatures the value of V_T at the same weight is consistently higher than for higher temperatures. Right: Even though the variation in Δ_T for different temperatures are not much larger than the ones observed in figure 40, the values for Δ_T seem to increase with rising temperature.

at 50 °C as expected operation temperature. Figure 41 shows a comparison of simulation results at the temperatures 40 °C, 50 °C and 60 °C. We see only a decrease of approximately 10 mV per 10 °C for V_T which is approximately 1 % of the absolute value. Although the values for Δ_T are only an estimate, we can nevertheless get a sense for the temperature dependence of Δ_T and see an increase in Δ_T of around 4 % for a change of 10 °C. The effect of the described temperature dependencies in an application will be treated in section 7.3.2 (see figure 62).

6.3.2 Process Corners and Mismatch

In addition to temperature sensibility, the influence of process variations and transistor mismatch was investigated in simulations of the circuit. The upper panels of figure 42 show the behavior of the exponential circuit in the process corners. We see that V_T as well as Δ_T vary strongly between different corners, the difference between the *sf* and *fs* corners is greater than 250 mV for V_t (roughly 30 %) and nearly 20 mV for Δ_T (roughly 15 %). Figure 42 top right, suggests that the value for Δ_T only significantly depends on the speed of the nMOS transistors but not on the pMOS transistors as the *ff* and *fs* corner on the one hand and the *ss* and *sf* corner on the other hand show approximately equal values. This is due to the fact that the single transistor which is responsible for the exponential dependency on V_{mem} is an nMOS transistor.
The bottom panels of figure 42 show the same experiment repeated with 15 Monte-Carlo samples. We observe a similar spread in V_T and Δ_T as in the corner simulation. Figure 43 shows the effect of the spread in V_T and Δ_T by comparing the current produced by the exponential term at a membrane voltage of 1.0 V. We see that for the 15 Monte-Carlo

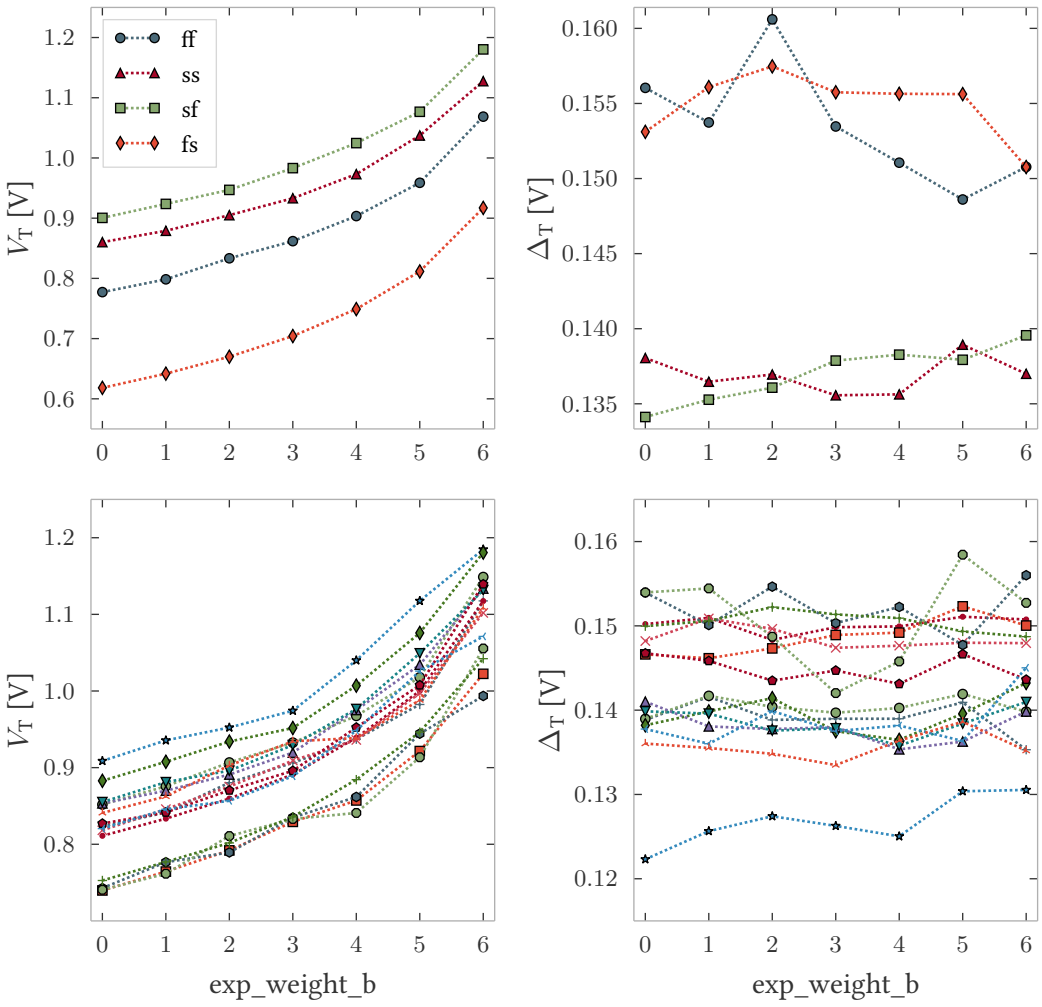

Figure 42: Top: Investigation of behavior of the exponential parameters in the process corners. The left plot shows a strong difference in the V_T parameter. The right plot shows that a fast nMOS corner produces a significantly higher Δ_T than a slow nMOS corner. The pMOS corner seems to have no influence on Δ_T.
Bottom: Monte-Carlo simulations with 15 different random seeds. We see a spread of approximately 200 mV for V_T and close to 30 mV for Δ_T.

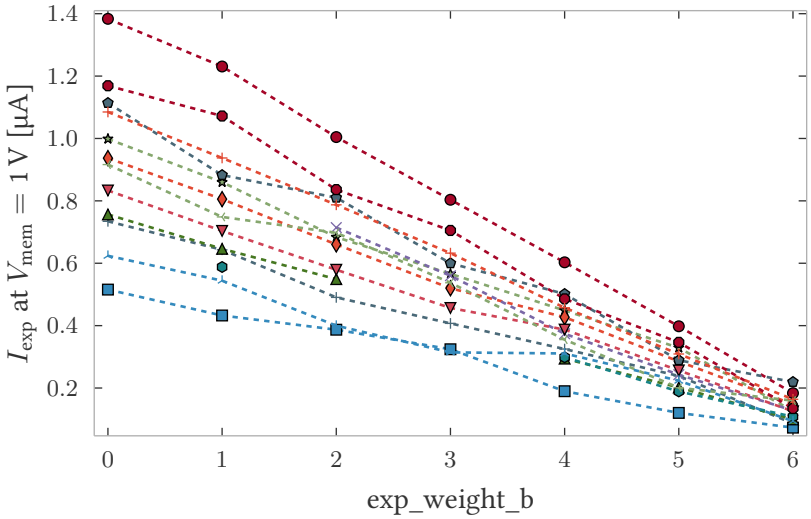

Figure 43: Exponential current at a membrane voltage of $V_{mem} = 1.0\,V$ over the exponential weight, for the 15 Monte-Carlo simulations shown in figure 42. The simulations show that the produced current can vary by a factor of 3 between different Monte-Carlo samples.

samples the current can differ by a factor of 3. These differences can lead to significantly different behavior of two neurons with the same parameter settings.

6.3.3 Calibration

The results in section 6.3.2 and figure 43 show the need for a calibration mechanism for the exponential term that allows to set the digital weight parameter that produces the V_T that is closest to the target value. Without that, two neurons with the same parameter settings might significantly differ in behavior due to strong differences in the strength of the exponential term.
To decide which hardware setting is required to achieve a certain set of AdEx model parameters $V_{T,bio}$ and $\Delta_{T,bio}$ we need to consider that in the model, the exponential current depends on the leak conductance but on the hardware it does not. Additionally, there is no parameter that controls $\Delta_{T,hw}$ which is fixed to approximately 0.15 V. The transformation between hardware and biological domain is

$$\Delta_{T,hw} = \alpha_v \cdot \Delta_{T,bio} \tag{6.14}$$

which yields $\Delta_{T,hw} = 15\,mV$ for $\alpha_v = 10$. This is much higher as for example the $\Delta_{T,bio}$ in the AdEx firing patterns, which is 2 mV. The only way to set a certain biological value is to adjust the transformation factor accordingly. With an $\alpha_v = 75$ the Δ_T of the firing patterns would be realized in the circuit. However, as α_v scales all voltages, the difference between leak and threshold potential would be larger than 1.2 V and therefore larger than the available voltage range on the hardware. With the current circuit is therefore not possible to set the correct Δ_T for the AdEx firing patterns.

In the circuit V_T is controlled by the `exp_weight_b` settings. The simulations in figure 40 show the dependency of V_T on that weight parameter. However, the simulations and fits were performed with one setting for the leak conductance. For a different setting the values for V_T change. As it is not feasible to determine the dependency of V_T on `exp_weight_b` for every possible leak conductance, we need a way to determine $V_\mathrm{T}(w_\mathrm{exp})$ once for one leak conductance and calculate the relationship for all other leak conductances:
The biological parameter is transformed into hardware domain by

$$V_\mathrm{T,hw} = V_\mathrm{T,bio} \cdot \alpha_v + \omega_v \tag{6.15}$$

which is the standard voltage transformation. Assuming we measured the dependency of $V_\mathrm{T,hw}$ on the weight parameter at leak conductance of g_L1 and now want to set $V_\mathrm{T,hw}$ for an experiment with a different g_L2, we need to translate the values for V_T1 measured with the first leak conductance into the equivalents we would have measured if the leak conductance would have been g_L2. To translate from the measured voltages to the domain of the target parameter we use that the exponential current on the hardware is independent of the leak conductance. With

$$I_\mathrm{exp,1} = \Delta_\mathrm{T} g_\mathrm{L1} \cdot \exp\left(\frac{V - V_\mathrm{T,1}}{\Delta_\mathrm{T}}\right) \tag{6.16}$$

$$I_\mathrm{exp,2} = \Delta_\mathrm{T} g_\mathrm{L2} \cdot \exp\left(\frac{V - V_\mathrm{T,2}}{\Delta_\mathrm{T}}\right) \tag{6.17}$$

$$I_\mathrm{exp,1} = I_\mathrm{exp,2} \tag{6.18}$$

we get

$$\frac{g_\mathrm{L,1}}{g_\mathrm{L,2}} = \exp\left(\frac{V - V_\mathrm{T,2} - V + V_\mathrm{T,1}}{\Delta_\mathrm{T}}\right) \tag{6.19}$$

$$V_\mathrm{T,1} = \Delta_\mathrm{T} \cdot \ln\left(\frac{g_\mathrm{L,1}}{g_\mathrm{L,2}}\right) + V_\mathrm{T,2} \tag{6.20}$$

where $V_\mathrm{T,1}$ is the value we need look up in a measurement like in figure 42 which was performed with $g_\mathrm{L,1}$ to get a value of $V_\mathrm{T,2}$ if the leak conductance $g_\mathrm{L,2}$ is used.
This method allows to determine the dependency of V_T on `exp_weight_b` once for one leak conductance and calculate the needed parameter for all other leak conductances.

6.4 V_mem, V_w - Phase space

To compare the mathematical AdEx model with the AdEx neuron realized on the HICANN DLS 3 we can compare their phase spaces. The phase space has two dimensions, which are the membrane voltage and the adaptation variable. In the case of the AdEx model the adaptation variable is the current w, for the hardware it is V_w.
A collection of phase space plots of the mathematical AdEx model can be found in figure 19. The shape of the V-nullcline depends on the leak conductance and on the exponential term. For low membrane voltages the leak dominates, for higher voltages the exponential current is dominant. The w-nullcline is a straight line with positive slope if the adaptation conductance is positive and a negative slope otherwise.

The nullclines for the AdEx circuit can not be calculated analytically, but they can be made visible by plotting the vector field in the state space. To measure the vector field we record the adaptation current and all other currents flowing onto the membrane capacitance while the membrane voltage and the adaptation voltage are pinned to points on a grid of values in the phase space. The current onto the adaptation capacitance is proportional to the derivative of the adaptation variable, while the sum of all currents onto the membrane (consisting of leak current, exponential current and adaptation current) is proportional to the derivative of the membrane voltage. We can therefore plot the vector field by drawing a vector for every grid point. The vectors are calculated as

$$\begin{pmatrix} \dot{V}_{\text{mem}} \\ \dot{V}_{\text{w}} \end{pmatrix} = \begin{pmatrix} \frac{I_{\text{mem}}}{C_{\text{mem}}} \\ \frac{I_{\text{w}}}{C_{\text{a}}} \end{pmatrix} \tag{6.21}$$

where I_{mem} are all currents onto the membrane capacitance C_{m} and I_{w} is the current onto the adaptation capacitance C_{a}.

In reality the phase space of the neuron circuit has more than two dimensions as all currents and voltages in the circuit are dynamic variables of the neuron. However, as figure 55 will show in section 7.2 a reduction to the two dimensions of V_{mem} and V_{w} still captures the relevant characteristics of the neuron.

Figure 44 shows the resulting phase space for positive adaptation and the weakest setting for the exponential weight (`exp_weight_b = 6`). The adaptation nullcline, i.e. the points where the sign of the V_{w}-component of the vectors switch, is denoted by a color change. As the V_{mem} component is much larger than the V_{w}-component, which is to be expected, as V_{mem} changes much faster than V_{w}, the adaptation nullcline would barely be visible otherwise. We see that the adaptation nullcline shows the expected shape of a straight line with positive slope. The bent at low membrane voltages is an artefact of the representation, as the long arrows from very low voltages cover the other arrows and hide the color change.

The V_{mem}-nullcline shows a different shape than the model in figure 19. In the model the nullcline is close to a straight line for low voltages, while for the hardware the nullcline is curved. This is caused by the fact, that in this region of the membrane voltage, the circuit implementation deviates from the mathematical AdEx model. Figure 36 shows, that for low membrane voltages the adaptation conductance is not constant. Additionally, the leak conductance is dependent on the membrane voltage as well.

Figure 45 shows the repetition of the simulation for a stronger setting of the exponential weight (`exp_weight_b = 5`). We see that the increase in the strength of the exponential current moves the rising and falling slopes of the V_{mem}-nullcline closer together.

Section 7.2 will show the phase space plots described here can made be useful in the task of qualitatively reproducing behavioral patterns of the mathematical neuron model the hardware emulates.

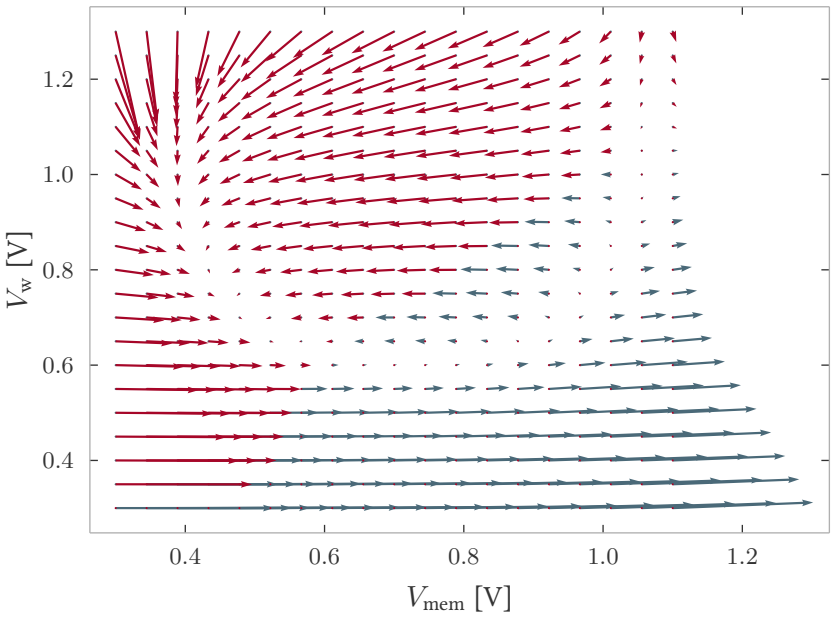

Figure 44: Phase space of the HICANN DLS 3 neuron with weak exponential term. The V_{mem}-nullcline is where the horizontal component of the arrows is zero. The V_w-nullcline is marked by the color change.

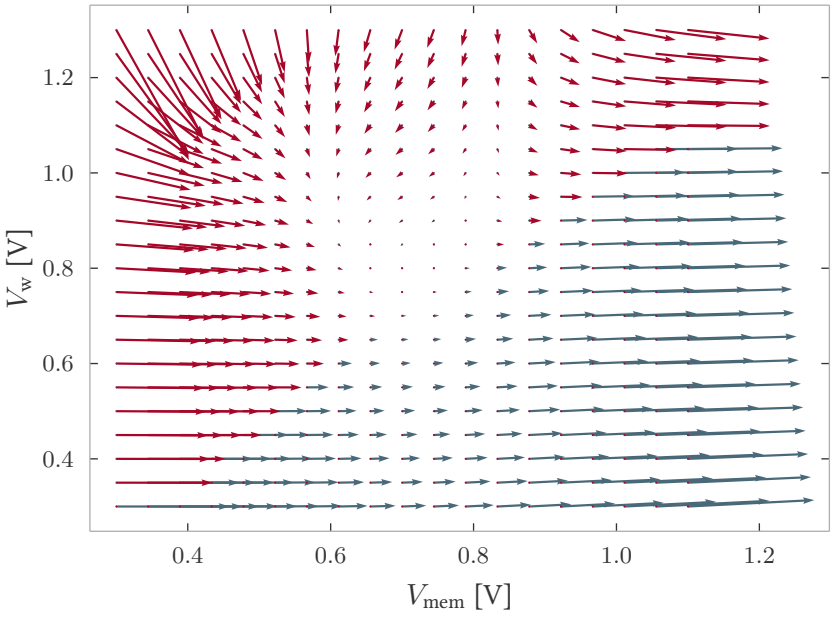

Figure 45: Phase space of the HICANN DLS 3 neuron with stronger exponential term. The V_{mem}-nullcline is where the horizontal component of the arrows is zero. The V_w-nullcline is marked by the color change. The rising and falling part of the V_{mem}-nullcline move closer together compared to figure 44.

7 High-Level Results

7.1 Single-Neuron Modeling Competition

7.1.1 Exess

The Exess neuron is designed to be an ideal version of the HICANN DLS 3 AdEx neuron and therefore needs to behave exactly as an PyNN or nest AdEx neuron which is transformed into hardware time and voltage domain. To verify this we use the data sets of the spike-based extension of the single-neuron modeling competition described in section 3.1.4.

Figure 46 shows a comparison between the Exess neuron and the PyNN.nest reference neuron for an exemplary data set. The parameters of the PyNN neuron have been transformed to Exess parameters using the transformations given in section 4.1. The plots show the traces of membrane voltage, adaptation variable and the synaptic input conductances in biological domain. We see that in this case, after a transformation into the biological domain, the traces of the Exess neuron perfectly match the PyNN.nest reference.

There are however data sets where the traces do not match perfectly. This does not necessarily mean that the deviations observed in these cases are caused by errors in the implementation of the Exess neuron. Different simulators use different algorithms to solve the differential equations of the neurons, can round of values at different stages of the simulations and by that cause different simulation results (Henker et al., 2012). To estimate the extent of simulator dependent deviations we run all data sets of the spike-based extension to the single-neuron modeling competition on 3 different simulators: the PyNN.nest reference, on PyNN with the NEURON simulator as back-end (Carnevale and Hines, 2006) and with the Exess neuron (which uses nest directly without the PyNN front-end).

Figure 47 shows a comparison of the results. To estimate the differences we compare the number of spikes produced by the simulators for each data set and plot histograms of the occurring differences. In this case it is not useful to use the Γ measure as indicator for the amount of deviation, as the differences are very small and all Γ values therefore are very close to 1. We see that the changing of the PyNN back-end can cause a change in the number of produced spikes in one data set by up to 7 spikes. This happens in situations where the neuron is very close to the threshold and for one simulator crosses the threshold while for the other stays just below the threshold. Figure 47 shows that the deviations between the Exess neuron and the PyNN.nest reference are smaller than the deviations between PyNN.neuron and the reference. We therefore conclude that the observed differences between the Exess neuron and the reference simulations are not caused by implementation errors, but by simulator internal differences.

In contrast to a software simulation a neuron on hardware can only stimulated by a spike at certain points in time. This is due to the fact, that spikes are sent into the chip as digital events. Therefore, they are arranged on a binned time axis before being sent to the neuron. The details of this protocol, e.g. the size of the time bins, depend on the specific implementation and vary between different chip versions.

In order to investigate the influence of the binning of the input spikes we use the Exess neuron and evaluate its performance in the spike-based modeling competition with pre-

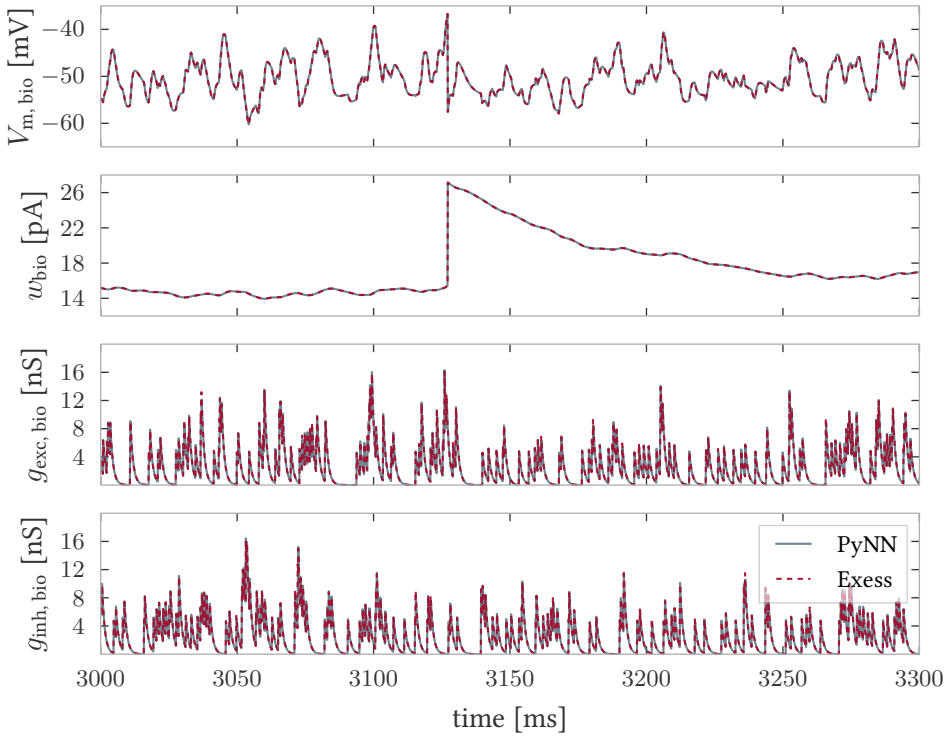

Figure 46: Comparison between PyNN with NEST as a back-end and an Exess simulation to verify implementation of the Exess neuron. As input a data set of the spike input extension of the single-neuron modeling competition is used. From top to bottom we see the membrane potential, the adaptation variable, the excitatory and inhibitory input conductances. The Exess traces were transformed into biological domain in order to compare them to the PyNN result. All traces show a perfect match.

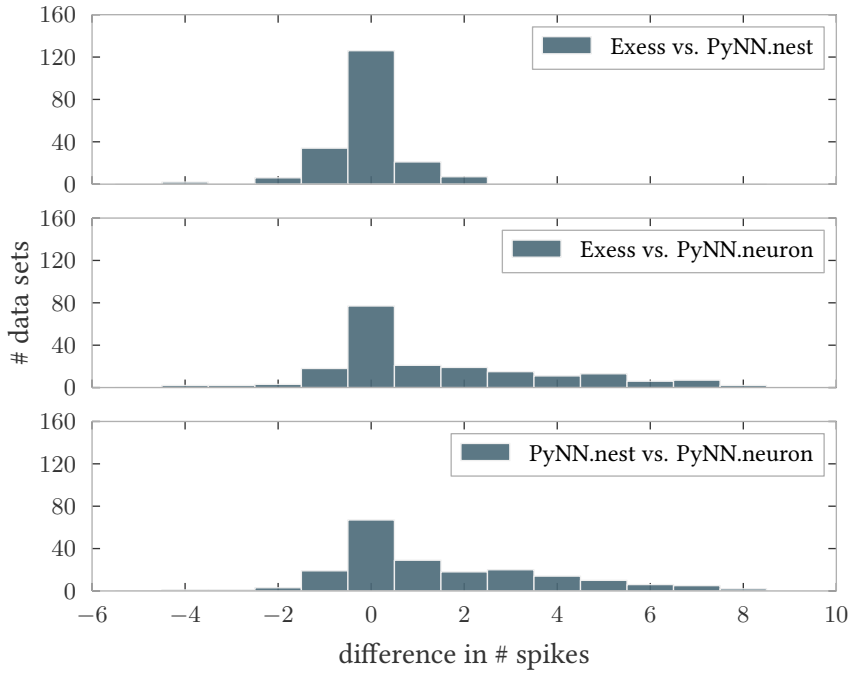

Figure 47: Comparison of the results in the spike-based competition of different simulators to the reference which was generated using PyNN.nest. As all simulators produce very high Γ values ($\Gamma > 0.9$), the comparison in the number of spikes produced in each data set shows differences more clearly. Since the Exess neuron is simulated using nest, it is not surprising, that the differences between PyNN.nest and Exess are smaller than the differences between Exess and PyNN.neuron or PyNN.nest and PyNN.neuron.

processed, i.e. time binned input. The Exess neuron allows to investigate the influence of time binning separately from other disturbances that can occur on the hardware, as we can be sure, that without the binned input, the Exess neuron reaches a nearly perfect score.

There are two possible ways of binning the original spike input. The first possibility is to choose a time axis with time step dt and move all spikes within the interval $t_0 \pm \frac{dt}{2}$ to t_0. If multiple spikes are within the same time bin, all except for one are dropped. In the second case multiple spikes within one bin are not dropped. All second spikes in a bin are sent to the neuron using a second synapse, all third spikes via a third synapse and so on. It can however not be guaranteed that this is always possible on hardware, as especially for large dt the number of spikes in a time bin might exceed the number of available synapses.

Figure 48 shows the results achieved by the Exess neuron in all data sets for all 3 reference AdEx neuron parameters depending on the size of the time bin dt (as reference data the PyNN.nest data is used). In this case the simple time binning without multiple synapses is used. All performance values Γ are calculated with a coincidence of $\Delta = 1.0\,\text{ms}$ in biological time domain. We see that for high input frequencies the performance measured by the Γ value drops significantly for $dt > 0.05\,\text{ms}$. The data sets with a lower input frequency

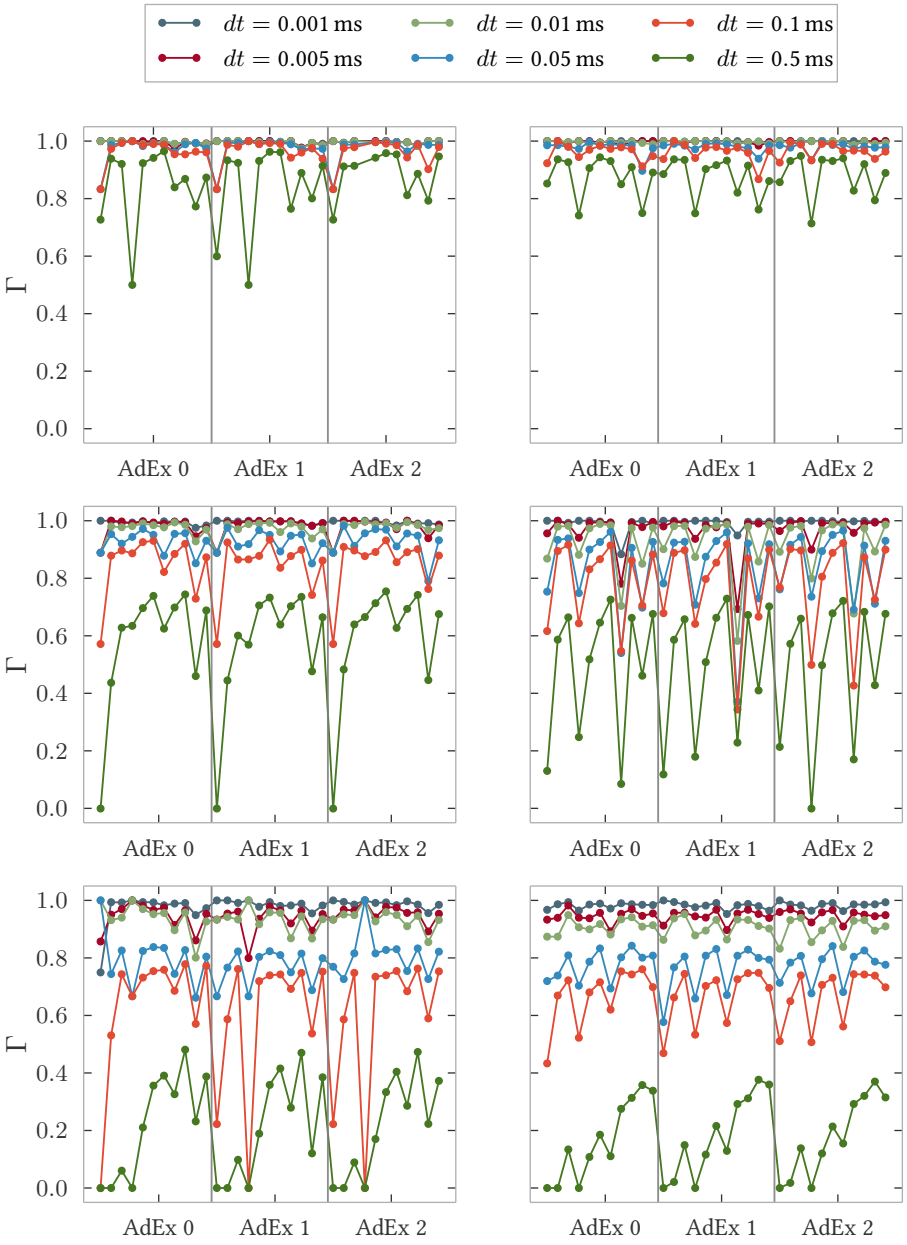

Figure 48: Results of time binning in the input without using multiple synapses. The plots from top left to bottom right show the Γ value the Exess neuron with time binned input reaches in the data sets with $\nu_{\text{bio}} = 100\,\text{Hz}$ and $\tau_{\text{syn,bio}} = 1\,\text{ms}$, $\nu_{\text{bio}} = 100\,\text{Hz}$ and $\tau_{\text{syn,bio}} = 5\,\text{ms}$, $\nu_{\text{bio}} = 500\,\text{Hz}$ and $\tau_{\text{syn,bio}} = 1\,\text{ms}$, $\nu_{\text{bio}} = 500\,\text{Hz}$ and $\tau_{\text{syn,bio}} = 5\,\text{ms}$, $\nu_{\text{bio}} = 2\,\text{kHz}$ and $\tau_{\text{syn,bio}} = 1\,\text{ms}$, $\nu_{\text{bio}} = 2\,\text{kHz}$ and $\tau_{\text{syn,bio}} = 5\,\text{ms}$. The dt denotes the size of the time bin. In each plot the results for all 3 AdEx reference parameters are shown. The higher input frequencies show a stronger performance drop for increasing time bins.

can tolerate a larger bin size of up to 0.1 ms. The size of the time bin is given in biological time domain. For the Exess simulation the time binned input has been transformed into the hardware domain.

This dependency on the input frequency indicates, that the drop in performance in mainly caused by the deleting of input spikes that are in the same time bin and not by the binning itself. For low frequencies fewer spikes are in the same time bin, therefore fewer spikes are deleted and the neuron performs better. This is indeed confirmed by figure 49 which shows the performance of the Exess neuron if the second binning procedure with multiple synapses is used. We see that the performance is nearly perfect, i.e. $\Gamma \approx 1.0$, for $dt < 0.5\,\text{ms}$. For larger time bins, the performance starts to decrease. In this case no clear dependence of the performance on the input frequency is visible, instead for very large time bins $dt > 1.0\,\text{ms}$ the performance seems to depend on the synaptic time constant. For the data sets with $\tau_{\text{syn, bio}} = 5.0\,\text{ms}$ (column on the right) on average larger time bins can be tolerated. This could be explained by the fact that with a larger synaptic time constant the input spike affects the neuron over a longer time, which makes the exact timing of the start of the stimulus less important.

Figure 50 illustrates the influence of the time binning on the performance using the voltage trace of an exemplary data set (all values in this figure are given in hardware domain). In this case the binning with multiple synapses is used. For small dt the neuron scores $\Gamma = 1.0$ which indicates a perfect reproduction of the reference spike times within the coincidence window. For $dt = 0.5\,\mu\text{s}$ the neuron produces an additional spike ($t \approx 300\,\mu\text{s}$) and does not spike at the required time ($t \approx 1650\,\mu\text{s}$). The additional spike is caused by the fact, that multiple spikes, that were in different time bins before, are merged into the same bin and therefore arrive at exactly the same time, which drives the neuron over the threshold. The missing spike could be explained by two spikes in adjacent time bins. For smaller bins they arrive at the neuron shortly after each other and are strong enough to cause a spike, for increasing time bins, they are further apart, which makes the summed up input for the neuron weaker. If the bin size is increased further they are moved into the same bin, arrive at the neuron at exactly the same time, and the neuron again produces an output spike (see $dt = 1.0\,\mu\text{s}$). For $dt = 5.0\,\mu\text{s}$ we see that the similarity to the spike train is completely lost, due to 3 additional and 1 shifted spike compared to the reference.

These evaluations show that a neuron can achieve good scores in these tests even if the spike input needs to binned. Depending on the binning mechanism (i.e. if multiple synapses can be used) bin sizes of $dt_{\text{bio}} = 1.0\,\text{ms}$ still allow an overall score of $\Gamma \approx 0.9$.

7.1.2 HICANN DLS 3 Simulation

A disadvantage of the transistor-level simulations, namely long simulation times, clearly shows, when we execute the spike-based extension of the single-neuron modeling competition. It takes close to two hours to run the full 5 s of one data set ($\nu = 2000\,\text{Hz}$, $\tau_{\text{syn}} = 1\,\text{ms}$). As this is not feasible for detailed investigations, we restrict ourselves to the first 1.5 s of all data sets and additionally limit the number of data sets by using only the first AdEx neuron as reference and only data sets with $\nu = 500\,\text{Hz}$, $\tau_{\text{syn}} = 1\,\text{ms}$. The input

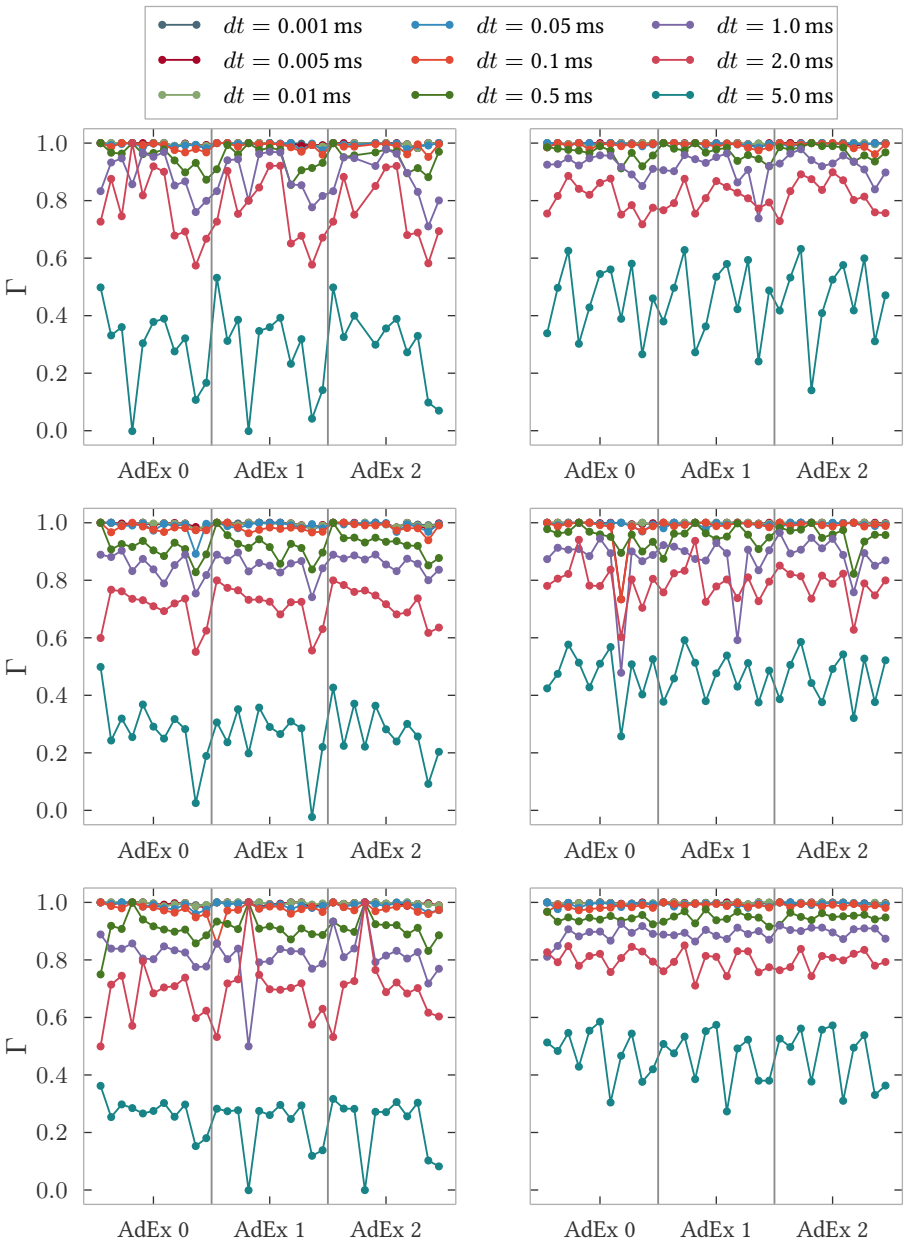

Figure 49: Results of time binning in the input using multiple synapses. The plots from top left to bottom right show the Γ value the Exess neuron with time binned input achieves in the data sets with $\nu_{\text{bio}} = 100\,\text{Hz}$ and $\tau_{\text{syn,bio}} = 1\,\text{ms}$, $\nu_{\text{bio}} = 100\,\text{Hz}$ and $\tau_{\text{syn,bio}} = 5\,\text{ms}$, $\nu_{\text{bio}} = 500\,\text{Hz}$ and $\tau_{\text{syn,bio}} = 1\,\text{ms}$, $\nu_{\text{bio}} = 500\,\text{Hz}$ and $\tau_{\text{syn,bio}} = 5\,\text{ms}$, $\nu_{\text{bio}} = 2\,\text{kHz}$ and $\tau_{\text{syn,bio}} = 1\,\text{ms}$, $\nu_{\text{bio}} = 2\,\text{kHz}$ and $\tau_{\text{syn,bio}} = 5\,\text{ms}$. The dt denotes the size of the time bin. In each plot the results for all 3 AdEx reference parameters are shown. Compared to figure 48 much larger time bins are possible before a strong decrease in performance is visible.

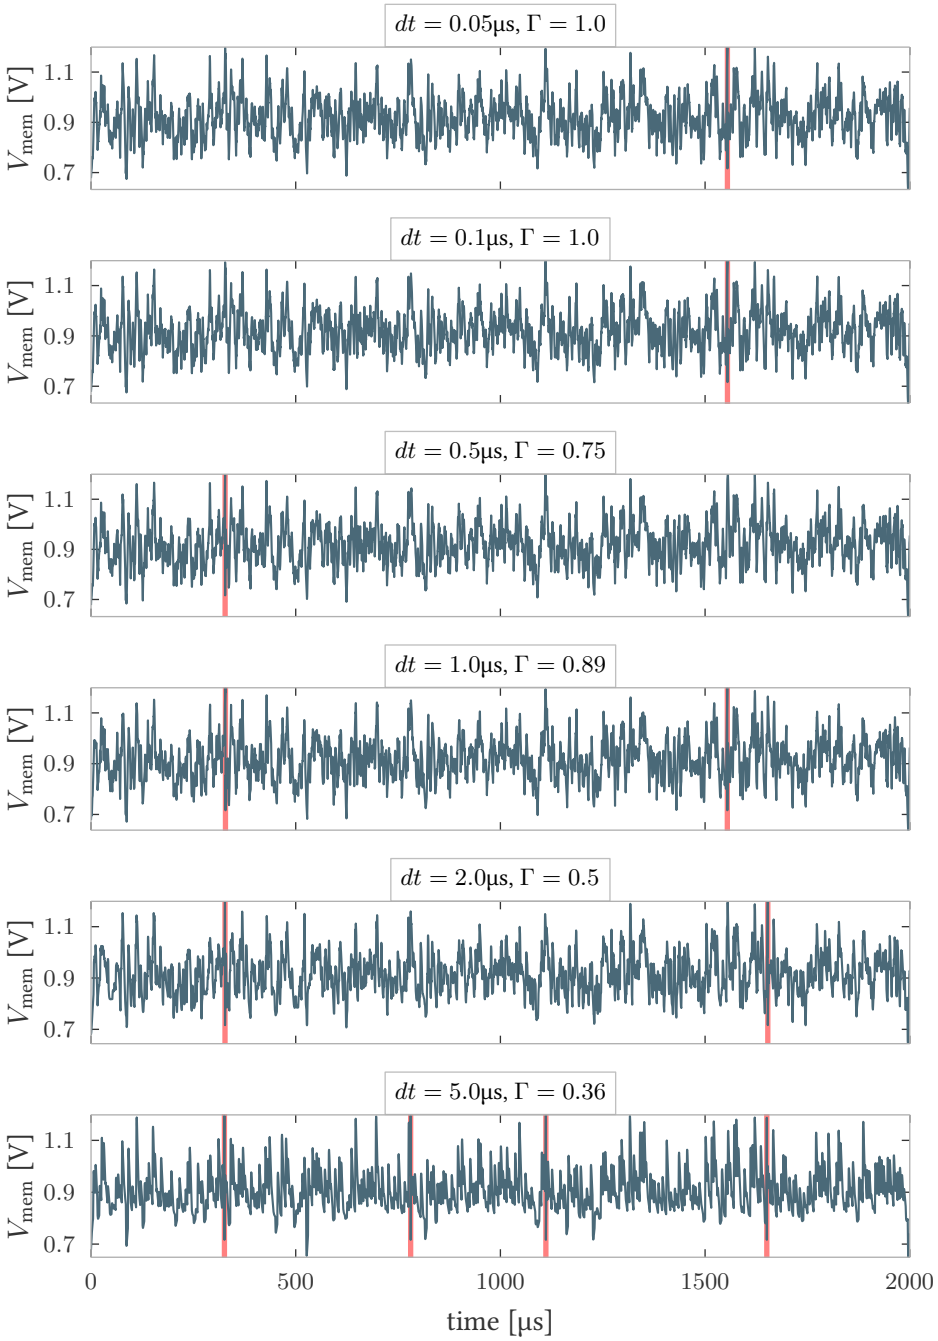

Figure 50: Resulting voltage traces for the Exess neuron with time binned input with multiple synapses in data set number 0 with $\nu_{bio} = 2\,\text{kHz}$ and $\tau_{syn,bio} = 1\,\text{ms}$ (bottom left in figure 49) for different sizes of time bins. The voltages and times in this plot are in hardware domain. The spikes of the neuron are marked in red. For increasing time bins not only the exact spike times are lost, but also additional spikes occur.

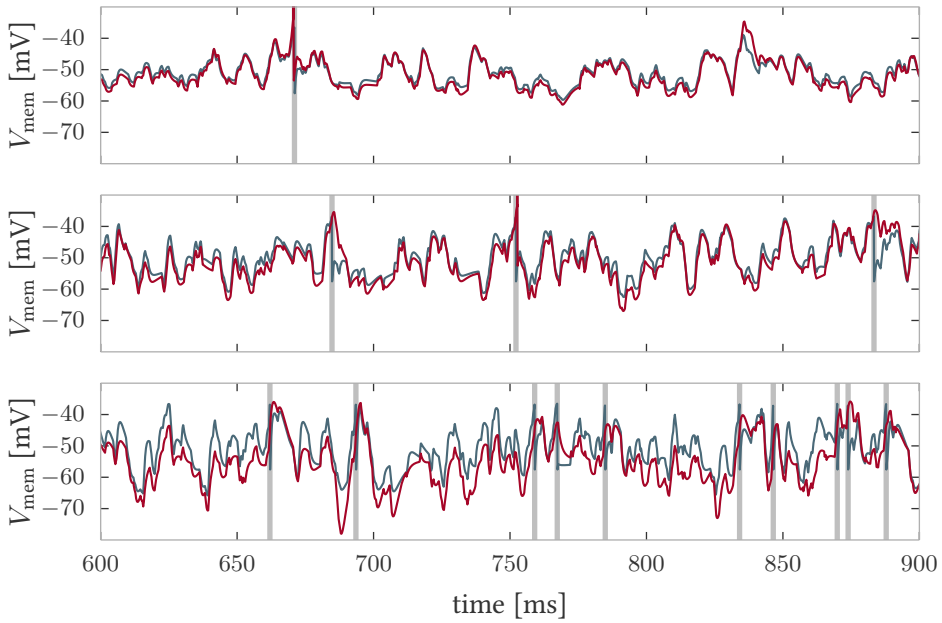

Figure 51: Voltage traces of the HICANN DLS 3 simulation (red) compared to the PyNN reference (blue) for 3 different data sets plotted in biological time- and voltage domain. All sets have an input frequency of 500 Hz (in biological time domain) and a synaptic time constant of $\tau_{\text{syn, bio}} = 1\,\text{ms}$. The topmost plot shows the result for data set 0, the plot in the center for data set 1 and the plot at the bottom for data set 2. In each plot, the reference spikes are marked in gray. We see that the reference of data set 0 is reproduced very well. For data set 1 more deviations occur, and for data set 2 no reference spike is matched by a spike of the DLS 3 neuron.

parameters for the used data sets can be found in section D.2. Time binning the input also speeds up the simulation, we therefore set a bin size of $dt_{\text{bio}} = 0.1\,\text{ms}$.

To obtain the parameter settings for the neuron, we use a calibration algorithm implemented by Paul Müller. We observed that after the calibration the exponential current was far too large compared to the reference. This is probably due to the fact, that Δ_T can not be configured. The exponential term was therefore switched off and is not used in the following simulations. Figure 51 shows the voltage traces of the HICANN DLS 3 neuron in three of the used data sets. The used neuron configuration can be found in section D.3.

We see that for the first data set the HICANN DLS 3 neuron reproduces the voltage trace very well. Also, the produced spike is coincident with the reference spike. For the second data set, the DLS neuron reproduces the reference trace rather closely, although the deviations are larger than in the topmost plot. Only 1 of 3 reference spikes is matched by a spike of the DLS neuron. We see that most of the time, the membrane voltage of the reference neuron is slightly higher, than the membrane voltage of the DLS neuron. This effect can be observed in a stronger fashion for the last data set. Here, reference and DLS

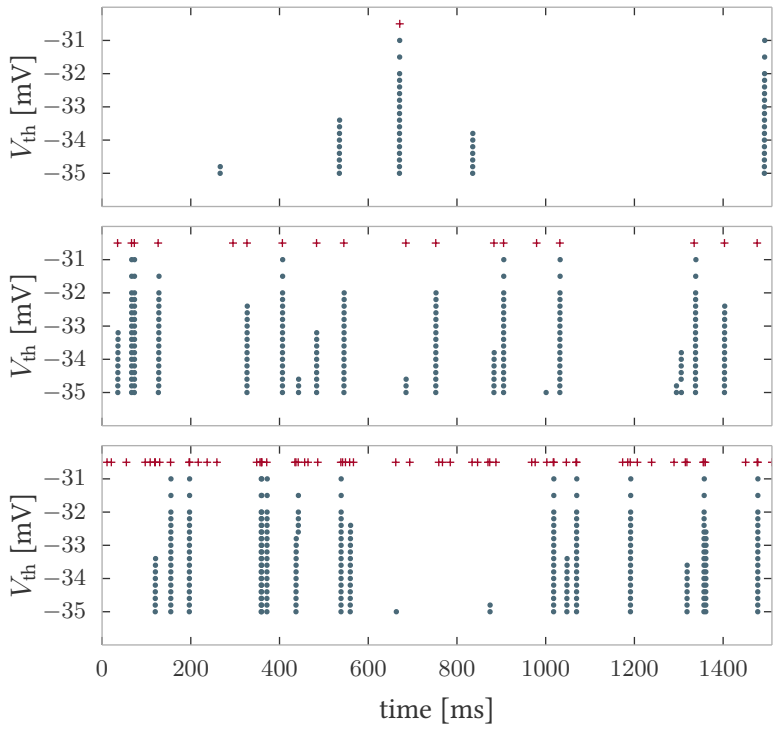

Figure 52: Results of a sweep over V_{th} for the data sets shown in figure 51. For each of the V_{th} values the spike times are shown (blue points). The reference spike times are shown as red + signs. While data set 0 shows too many spikes, the data sets 1 and 2 show too few.

trace match poorly and none of the reference spikes is reproduced. It is clearly visible that there is an offset between reference and the simulated DLS trace.

This is caused by too high excitatory weights in the data sets. On HICANN DLS the synaptic weight is controlled by a digital 6 bit parameter. The calibration algorithm determines the weight value that would be required for a certain target weight. If the required value exceeds the available range, the weight is configured to 63, which is the highest possible value.

In the case of figure 51 an excitatory weight of 77 for the panel in the center and of 112 for the bottom plot would have been required. For the simulation the weights in both cases were set to 63. A simple solution for this is to use two synapses with half of the weight each. However, this was not possible, as the simulation setup only allows for four excitatory synapses. Three of these four available synapses are already used for the time binning of the input spike as described in section 7.1.1.

In order to determine whether the observed offsets can be compensated for, we sweep the threshold voltage. The results for the same data sets as in figure 51 (recorded with $V_{\text{th,bio}} = -32.4\,\text{mV}$) can be found in figure 52. We see that, as expected, the lowering of the threshold does not improve the performance in the first data set, as the available

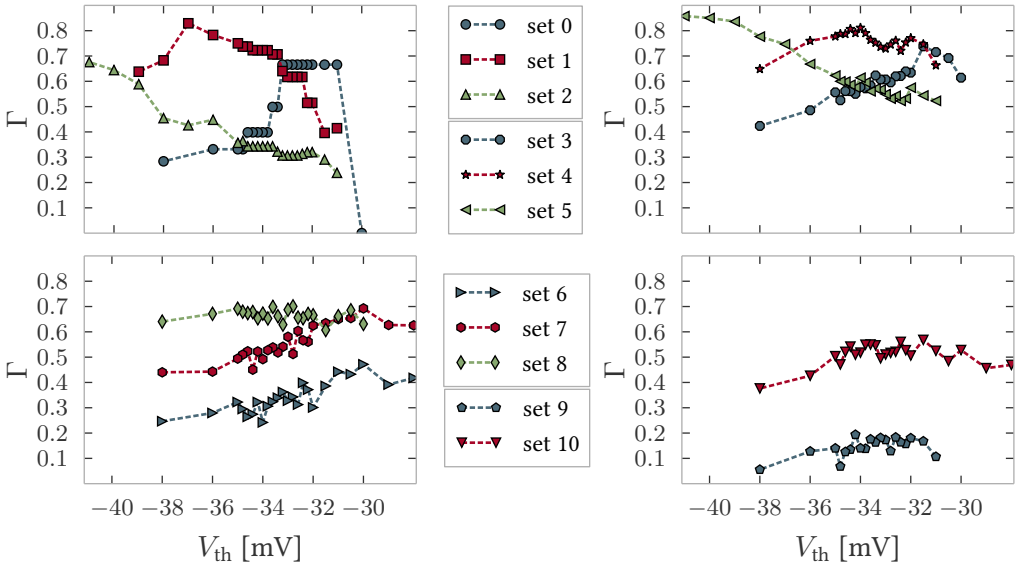

Figure 53: Summary of the results of the sweep over the threshold voltage. For each threshold the neuron's performance in all data sets was evaluated using the Γ measure. The data sets are split into different plots to improve readability.

excitatory synaptic weight was sufficient there. For the other data sets we find that for a lower threshold more spikes (with the correct timing) are produced.

Figure 53 shows a summary of the results obtained in the sweep over V_{th}. For each threshold voltage all data sets were simulated and the performance determined using the Γ measure. We see that there are several data sets for which the performance improves with a decreasing threshold voltage (sets 1, 2, 5). As explained above, this is due to the limitation of the excitatory synaptic weight.

For most of the other data sets, generally, higher threshold voltages produce a better performance. This is problematic, because the reference for all data sets was produced using the same AdEx parameters. It should therefore not be necessary to choose different DLS parameters for the data sets.

Figure 54 shows, for the example of data set 7, why this is the case. In the top panel on the left we see that the slope of the dependency of the leak current on the membrane voltage is not constant as in the AdEx reference. Additionally, the slope, which corresponds to the leak conductance g_{L} is, for voltages between 0.76 V and 0.8 V significantly lower than in the reference. The plot of the membrane voltage over time on the right, shows that the neuron is in this voltage region for most of the simulation.

The center left panel shows that in the beginning of the simulation the adaptation current of the circuit and the reference match, but then the adaptation OTA saturates and therefore the current in the circuit simulation can no longer follow the further increasing reference. The center right panel shows, for the excitatory case, that the calibration of the synaptic weights and the synaptic time constants worked well, as the current produced by the synaptic input circuit and the reference match well. The visible offset is intentional, as it is used to cancel the offset of the leak OTA. The bottom panel shows a comparison

between the membrane voltages of the circuits and the reference in biological time and voltage domain. We see that with increasing time, they differ more strongly. The voltage of the AdEx reference is below the voltage of the DLS 3 circuits most of the time. This can be explained by the differences in the leak and adaptation currents, as discussed before. A weaker adaptation as well as leak lead, for the same threshold voltage, to a higher spike rate. This explains, why the increase of the threshold improves the performance, as it partly compensates for the decreased leak and adaptation. Figure 64 and Figure 65 in section A show the same plot for the fourth and ninth data sets. The fourth data set shows a good performance of $\Gamma \approx 0.8$ while the performance for the ninth set is the worst with $\Gamma \approx 0.2$. We see that indeed the effects impairing the performance are the same for these data sets as well and that only the magnitude of the disturbance differs.
For the described investigations only a subset of the available data sets could be used due to long simulation times. To allow for more detailed evaluations the simulations must be sped up. Additionally, the problem of the truncated synaptic weights (in this case for the data set 1, 2 and 5) needs to be solved, for example by increasing the number of available synapses in the simulation. As long as the performance in these data sets is impaired by a decreased input strength a detailed comparison to the other data sets is not meaningful. The comparison of membrane currents to th reference for the other data sets showed that the main effects lowering the performance are the saturation of the adaptation current and the decreased leak conductance. The observed mismatch between the leak conductance in the AdEx reference and the DLS 3 simulation (figure 54) can be circumvented using the linear part of the relation between the leak current and the membrane voltage. This may improve the results shown in figure 53.

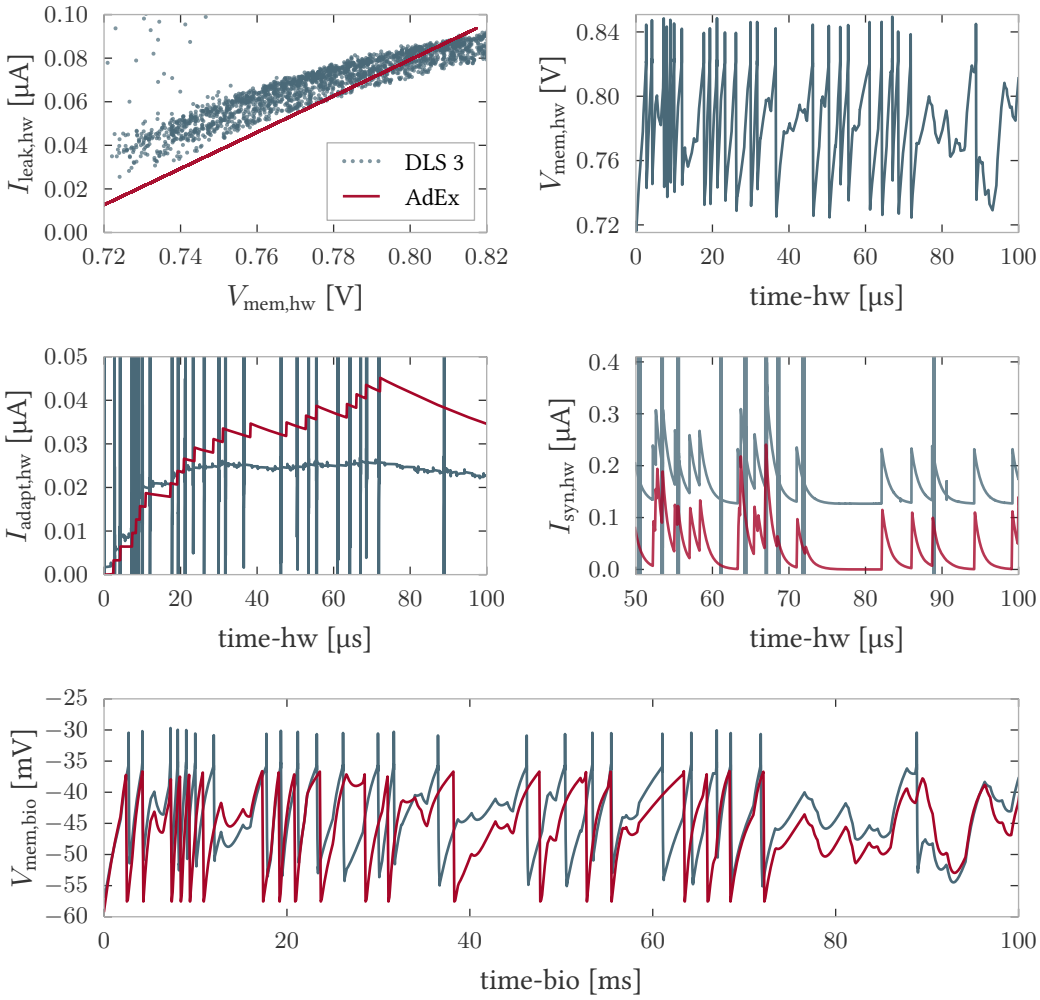

Figure 54: Comparison between the circuit simulation of HICANN DLS 3 and the AdEx reference in the seventh data set. Top: On the left a comparison of the leak current over the membrane voltage in hardware domain is shown. In contrast to the reference (red) leak conductance (slope) is not constant for the circuit (blue). The data is extracted from a transient simulation and thus noisy. The right plot shows the membrane voltage in the circuit simulation in hardware domain. Center: The left plot shows a comparison of the adaptation currents (colors as before). The adaptation current produced by the adaptation circuit saturates at $t \approx 30\,\mu s$. The right plot shows that the excitatory synaptic weight and the synaptic time constant match the reference well. Bottom: Comparison of the membrane potentials over time. The short, sharp peaks on I_{syn} and I_{adapt} are caused by the switching between leak and reset mode of the leak/reset OTA when the neuron spikes.

7.2 AdEx firing Patterns

7.2.1 HICANN DLS 3 Simulation

Seven out of eight of the firing patterns described in section 3.2 could be reproduced qualitatively in transistor-level simulations of the HICANN DLS 3. As starting point, the parameters in biological domain summarized in section B.1 were translated into the hardware domain. Then the measurements described in section 6 were used to determine the hardware parameter values that approximately realize the target parameters. For an ideal circuit this would suffice to reproduce the firing patterns, however with real components the effects of limited configurability, parameter ranges and saturation effects need to be taken into account.

In the case of the HICANN DLS 3 circuits the value for Δ_T can not be configured and is larger than the value which is required for the firing patterns. Therefore, the other parameters need to be adjusted accordingly. Additionally, the leak as well as the adaptation conductance are not constant over the complete voltage range of V_mem and V_w. Because of that a simple parameter transformation from biological to hardware domain is not feasible here either.

Using the naively transformed parameters as a starting point the parameters were fine-tuned by hand. Comparing the trajectories in the phase space to the original phase space paths in figure 19 allowed to ensure that the circuits qualitatively reproduce the correct AdEx features in the pattern. Figure 55 shows the HICANN DLS 3 trajectory in the phase space for the regular bursting pattern (the voltage trace can be found in figure 56). We see that the mechanism producing the pattern is as in the AdEx reference: The neuron spikes several times until it is reset above the V_mem-nullcline, which causes a longer pause between the spikes with a downswing of the membrane, as the neuron needs to return below the nullcline before it can spike again.

Figure 56 shows the voltage traces of the firing patterns in the HICANN DLS 3 simulation. The parameter settings producing the firing patterns can be found in section B.2. The space for the irregular firing pattern has been left blank, as it could not be reproduced. This might be due to the fact that as described in Naud et al. (2008) it occurs only for scattered points in the parameter space, which makes it difficult to obtain by hand tuning. It is not excluded that the pattern can be produced with the circuits, but the required settings were not found.

In addition to the difficulties in reproducing the irregular firing the delayed regular bursting pattern is instable in the circuit simulation. This is due to the fact that it requires a negative adaptation conductance of the same absolute value as the leak conductance $a = -g_\mathrm{L}$. Even in a numerical simulation this is unstable and very sensitive to parameter variations, because, if $|a| > g_\mathrm{L}$ and the membrane voltage is above the leak potential there is a positive feedback current onto the membrane and the neuron starts to fire without any current input. In the circuits it is practically impossible to have a perfectly equal adaptation and leak conductance as the parameters can only be set with a fixed precision and leak as well as adaptation conductance depend on the membrane and adaptation voltage. It is therefore not surprising that in the circuit simulation of the delayed regular bursting pattern we see the instability shown in figure 57, where the neuron starts to fire long

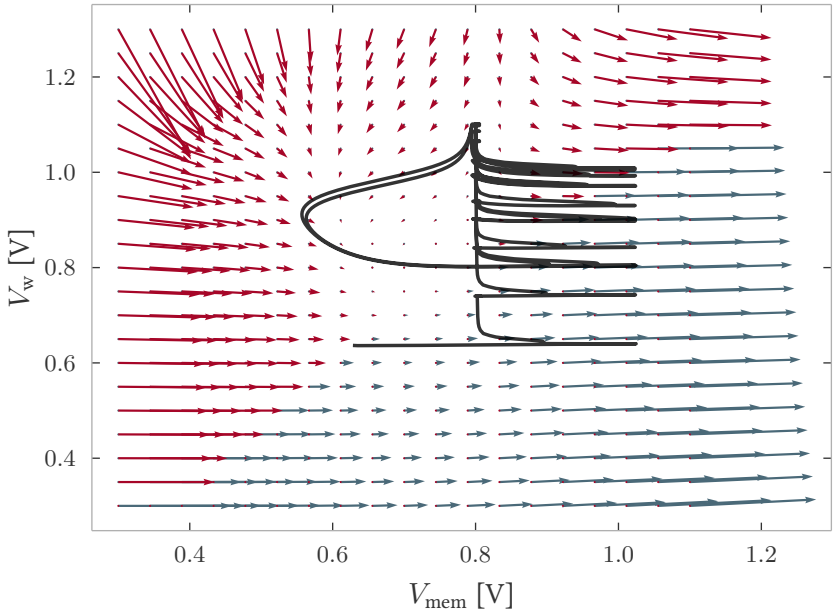

Figure 55: Phase space trajectory of the HICANN DLS 3 reproducing the regular bursting pattern. The neuron spikes several times and is reset to a higher V_w value due to the spike-triggered adaptation. When the neuron is reset above the V_{mem}-nullcline, a broad reset with a longer spike pause follows. To plot the trajectory the same data as for the regular bursting pattern in figure 56 is used.

before the current stimulus is applied. In figure 56 the time interval between the start of the simulation and the start of the current stimulus is chosen to be short, which starts the current stimulus before the neuron can fire by itself. Setting the time of the start of the stimulus to a small value hides the instability, but does not prevent the effect.
The other patterns however do not show any instabilities and reliably reproduce the key characteristics of their model references.

7.3 BAC Firing

7.3.1 Exess

The Exess neuron allows to simulate multi-compartment neuron consisting of ideal hardware neurons. This is done by using the gap junction features of the NEST simulator (Hahne et al., 2015). In NEST gap junctions are modeled as electrical synapses, that produce a current

$$I_{gap} = g_{ij} \cdot (V_i - V_j) \tag{7.1}$$

which shows the same voltage dependence as an intercompartment conductance in a simple multi-compartment neuron model. A reset mechanism which allows plateau potentials was implemented as well.
These features allow to set up a 3-compartment neuron, which is configured to reproduce the key features of the BAC firing mechanism described in section 3.3. Figure 58 shows the

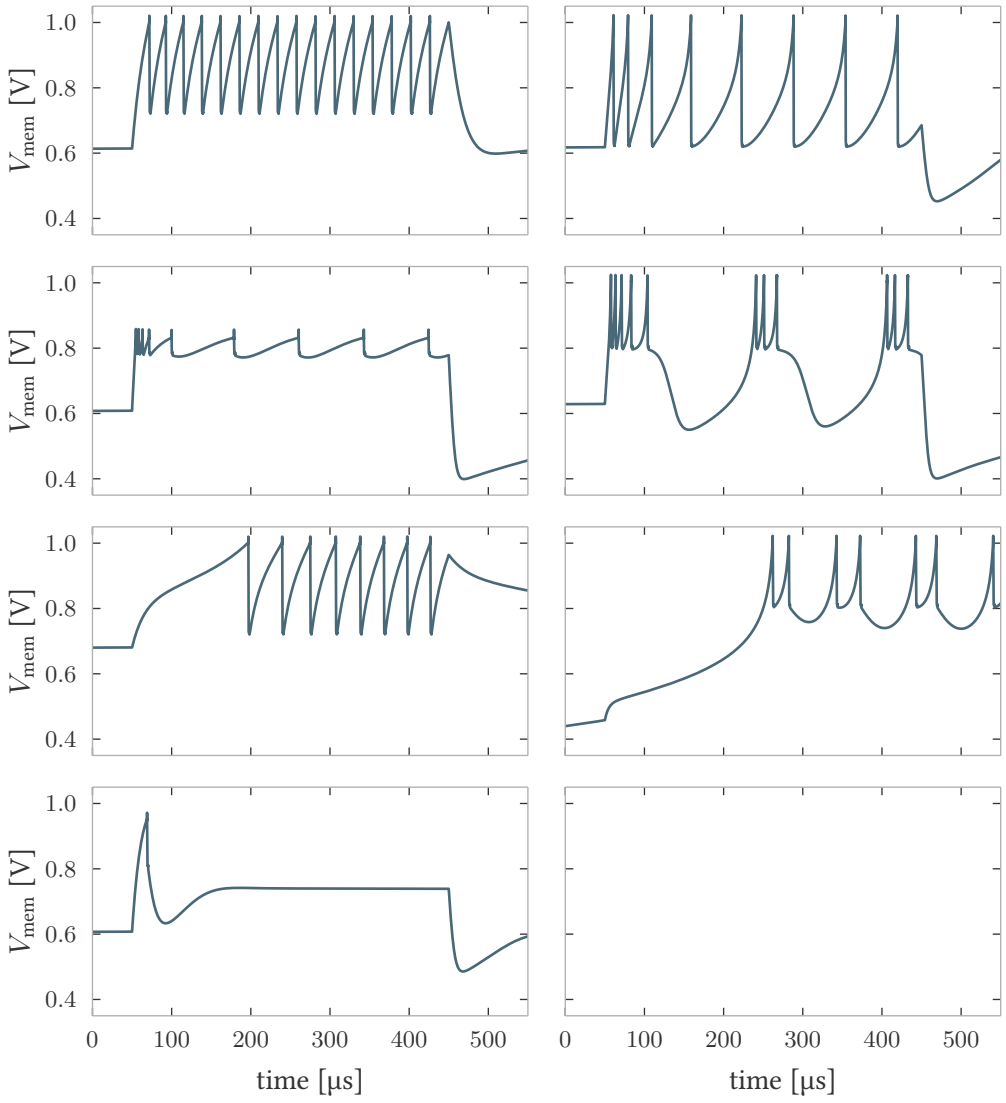

Figure 56: Simulations of the AdEx firing patterns on the HICANN DLS 3. It was possible to tune the neuron parameters to qualitatively reproduce 7 of the 8 patterns shown in figure 18. From top-left to bottom right the plots show tonic spiking, adaptation, initial bursting, regular bursting, delayed accelerating, delayed regular bursting and transient spiking.

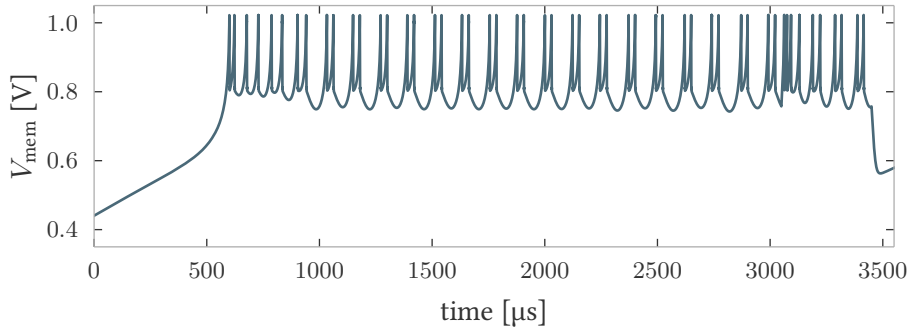

Figure 57: Demonstration of the instability of the delayed regular bursting pattern. In comparison to figure 56 only the time before the start of the simulation was varied. The stimulus starts at $t = 3000$ µs however the neuron starts to fire much earlier. This is caused by the fact that in the model, the leak and adaptation conductance are equal for this pattern. On the hardware however they vary effectively with changing membrane voltage. This, in combination with the negative adaptation, causes a positive feedback and the neuron fires without additional current input.

3 sub-experiments of the BAC firing mechanism, the dendritic stimulation that produces a PSP in the dendrite (labeled with NMDA in the figure), the somatic stimulation which produces a spike in the soma and the combination of both inputs that produces plateau potential modeling NMDA and calcium spikes and a burst in the soma. The neuron parameters for this experiment can be found in section C.1.
The influence of the gap junctions functioning as intercompartment conductances can, for example, be seen in the dendritic stimulation, where the PSP in the dendritic compartment propagates into the calcium compartment and then into the soma and is attenuated on the way. The low-pass feature of this connection can be seen in the somatic stimulation, where the spike of the soma propagates into the other compartments. We see that the signal is not only attenuated but fast changes are filtered out as well. This shows that the gap junction implementation in the NEST simulator indeed can be used to model intercompartment conductances.

7.3.2 HICANN DLS 3 Simulation

The BAC firing mechanism can serve as a high-level test for the new plateau potential and multi-compartment features of the HICANN DLS 3. The mechanism can be modeled, as shown in section 7.3.1 using a 3 compartment neuron where the compartments form a chain. Figure 59 shows a schematic drawing of the realization of this structure on the chip. Due to the design of the multi-compartment connections it is necessary to merge two neuron circuits into one compartment, as the chain structure otherwise can not be realized.
With this setup the result shown in figure 60 can be obtained (Schemmel et al., 2017). As at the time of this simulation a full neuron calibration was not available all neuron parameters and the input strengths are hand tuned. We see that the compartments in figure 60 show

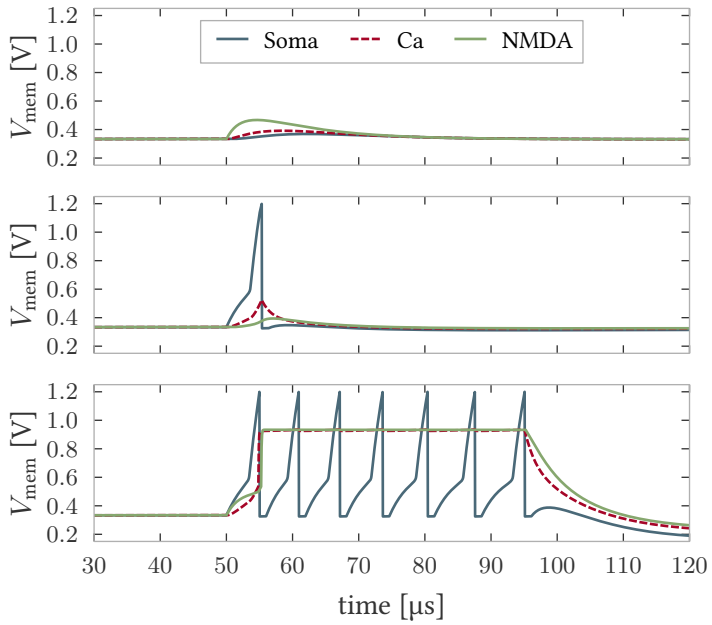

Figure 58: BAC firing mechanism simulated with a three-compartment Exess neuron using the NEST implementation of gap junctions as intercompartment conductances. The mechanism is modeled correctly: Dendritic stimulation (top) causes a PSP in the dendritic compartment (green), somatic stimulation (center) causes a spike in the soma (blue) and the combination of both stimuli (bottom) causes a plateau potential in the dendritic compartment and in the calcium compartment (red) and a burst in the soma.

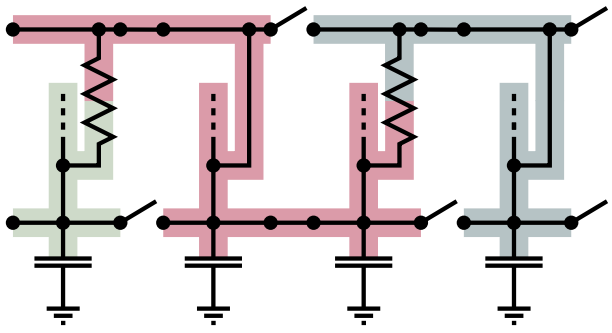

Figure 59: Multi-compartment switch configuration for BAC firing. The two neuron circuits in the middle are merged into one compartment in order to allow a chain topology. The colors correspond to the trace colors in figure 60. Drawing by Paul Müller.

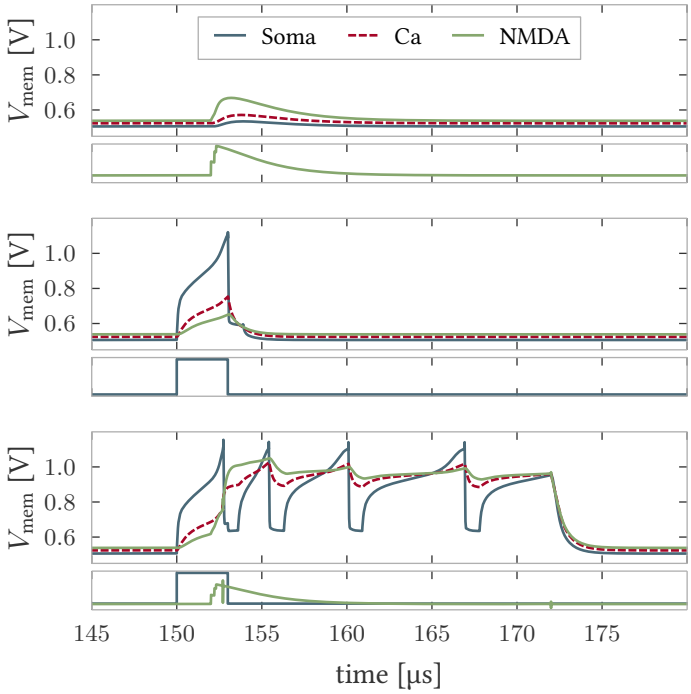

Figure 60: Transistor-level simulation of the HICANN DLS 3 qualitatively reproducing the BAC firing mechanism shown in figure 20, adapted from (Schemmel et al., 2017). Dendritic stimulation using synaptic input induces a PSP (top), somatic stimulation induces a somatic spike (center) and the combination of both inputs induce plateau potentials in the Ca and NMDA compartments as well as a spike burst in the soma (bottom). The neuron parameters for this application can be found in section C.2.

the expected behavior described in section 3.3. As the strength of a single synapse was not sufficient to provide a sufficiently strong PSP in the dendritic compartment, 3 spike inputs in short succession were used (see lower plot in top of figure 60). The parameters used to obtain this result can be found in section C.2.

The BAC firing mechanism is fragile and can break easily if parameters like for example the leak conductance or firing thresholds are changed without adjustments for the other parameters. It is therefore interesting to investigate the influence of possible distortions that can happen on the chip.

First we investigate the influence of temperature changes. As measure for the difference between functional reproductions of the BAC firing mechanism we use the number of soma spikes in the burst occurring with the combined stimulation. To have a more meaningful measure we increase the length of the plateau potentials in order to allow more spikes. A functional reproduction of the mechanism must include: no spikes for the dendritic stimulation, a single somatic spike for the somatic stimulation, no plateau potentials

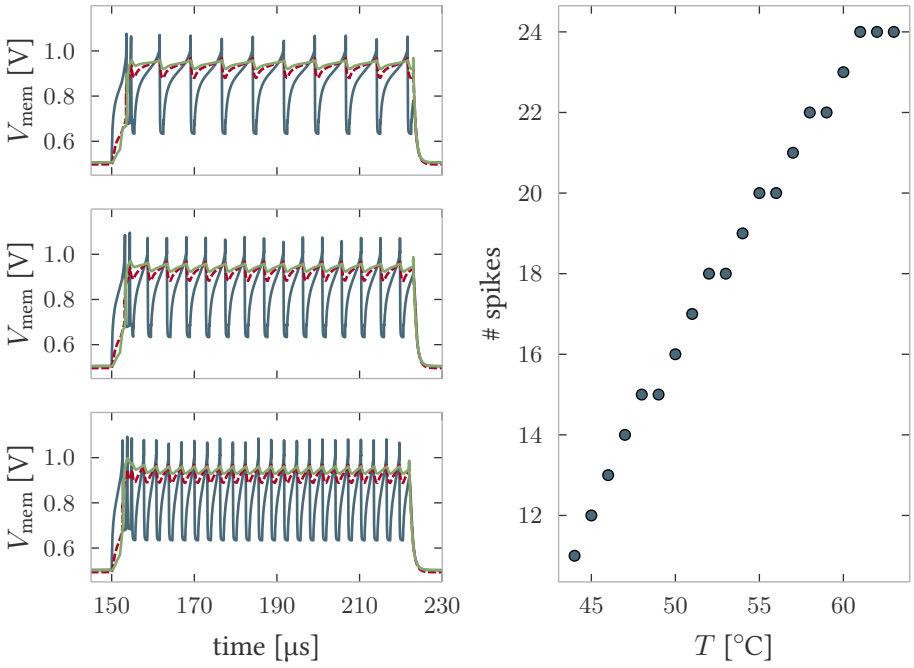

Figure 61: Left: Voltage traces of combined stimulation for increasing temperatures (from top to bottom 44 °C, 50 °C, 63 °C). The increasing number of spikes during the burst is clearly visible. Right: Number of somatic spikes in the burst depending on the circuit temperature. Only temperatures that correctly reproduce the firing behavior are included.

for the somatic stimulation, more than one somatic spike for the combined stimulation and plateau potentials in the calcium and the NMDA compartment.

Figure 61 shows the variation of the number of spikes produced in the somatic burst depending on the temperature of the circuits. Only the temperature was varied, no neuron or input parameters were adjusted. For temperatures below 44 °C and above 63 °C the firing mechanism can not be reproduced. We see that over a range of 19 °C the number of spikes during the burst increases by more than a factor of two.

The simulations in figure 61 were performed with an earlier version of the schematic compared to the simulations in figure 60, which used the final schematic. However, the old simulations are still valid, as the later changes in the schematic either influenced the adaptation circuit, which is not used here, or are compensated for by neuron parameter changes (e.g. the parameterization of C_{m} changed, which is compensated by choosing different settings for the `en_mem_cap` parameter and adjusting other neuron parameters accordingly).

To investigate the source of the temperature dependency, the current contributions to the total current onto the membrane were analyzed separately. Figure 62 shows the comparison of these currents for the minimal and maximal temperature for which the correct firing behavior is produced as well as for the standard simulation temperature of 50 °C. We see that out of leak, synaptic input, intercompartment currents and exponential current only

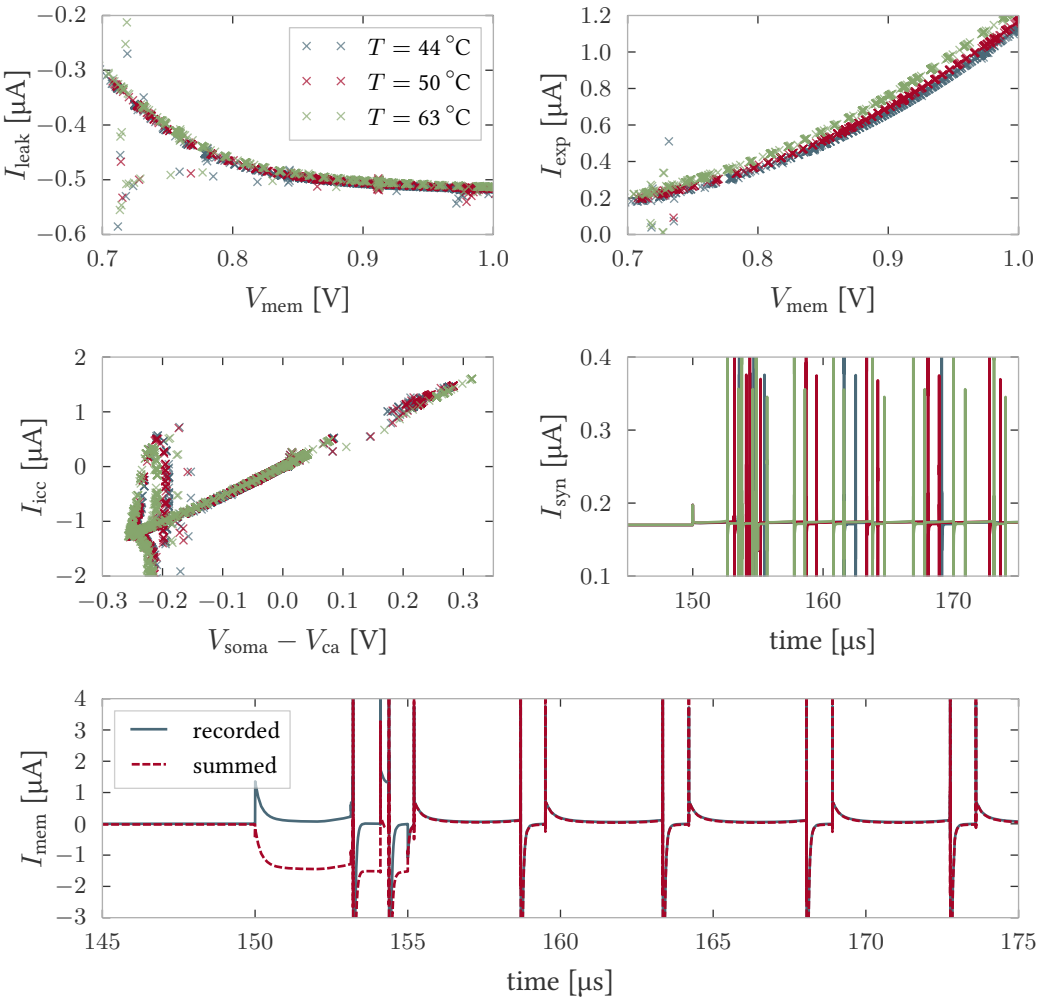

Figure 62: Investigation of temperature dependency of spike frequency. Top and center: Comparison of currents flowing onto the membrane for different temperatures. The data is extracted from a transient simulation and thus noisy. The leak current is plotted over the membrane voltage (top left) and does not show a temperature dependence. The exponential current is plotted over the membrane voltage but here a temperature dependency is visible. For increasing temperatures the exponential current rises as well. The plots in the center shows that neither the intercompartment current (left) nor the synaptic input current (right) show a strong temperature dependency. Bottom: Currents shown above summed up and plotted together with measures current flowing onto the membrane. The only visible difference is at $T = 150 - 155\,\mu\mathrm{s}$ which is the current stimulus into the soma. This shows that all currents that affect the evolution of the membrane are considered above. The short, sharp peaks on I_{syn} and I_{mem} are caused by the switching between leak and reset mode of the leak/reset OTA when the neuron spikes.

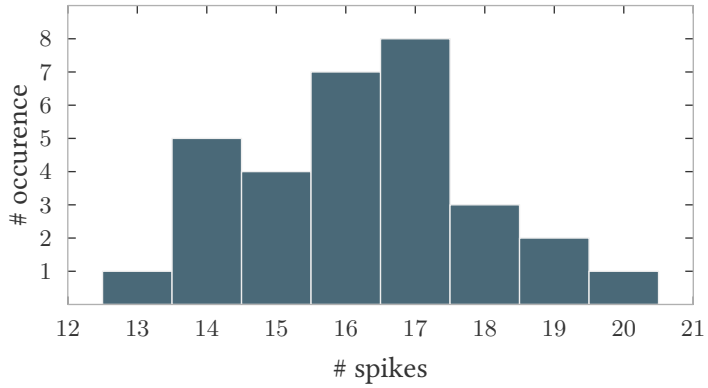

Figure 63: Distribution of number of spikes per burst measured with the same settings as figure 61 but added parameter variations of ± 1 LSB to the analog neuron parameters provided by the capacitive memory. The number of spikes can vary up to a factor of 1.5.

the exponential term shows a visible temperature dependency. As the bottom panel of figure 62 shows, the sum of above currents is equal to the complete current flowing onto the membrane, which was recorded directly at the membrane capacitance. This indicates, that the there is no additional relevant current, that could cause the observed temperature dependency, which therefore is most likely caused by the exponential term.

The investigation of temperature stability is essential as temperature fluctuations on the chip can be caused by the PPU. At the moment there are no measurements on the possible temperature changes caused by the PPU but, depending on cooling, chip size and the executed code, they can vary strongly. A possible scenario where these fluctuations pose a problem would be, for example, that the PPU is used to tune a network of neurons to show the BAC firing behavior with a certain firing rate in the burst. When the parameter tuning is done, the PPU is switched off to save energy, which causes the temperature on the chip to decrease. The temperature drop then can, as figure 61 has shown, either break the BAC firing mechanism or change the firing rate significantly.

In addition to temperature variations, the effect of parameter variation caused by the capacitative memory (Hock, 2014) were investigated. The analog parameters stored in the capacitive memory can be set with 10 bit precision, i.e. to 1024 different values. With a maximum current value of 1 µA this yields a step size (or least significant bit, LSB) of approximately 1 nA. For the voltage parameters a maximum of 1.8 V is given in Hock (2015). This yields a voltage step size of 1.8 mV. Discretizing the desired target parameters of a neuron configuration can, in the worst case, lead to an error of ± 0.5 LSB.

Figure 63 shows an estimation of the influence of capacitive memory induced parameter variations on the BAC firing behavior. As measure for the difference of measurements again the number of spikes in the burst is used. To estimate the influence 31 experiments were performed. For each measurement random numbers within a uniform distribution between ± 1 LSB were added to all current and voltage parameters of the neuron. In all 31 tries the neuron still showed the correct BAC firing behavior, however the number of

spikes varied up to a factor of 1.5. We chose to add parameter noise of 1 LSB in order to not only estimate the effect of discretization, as those are not the only disturbances that can occur in the capacitive memory (e.g. crosstalk as described in section 2.2 of Stradmann (2016))

This shows that the parameter fluctuations caused by discretization of the analog neuron parameters are unlikely to break the BAC firing mechanism of a neuron but have a strong effect on the detailed behavior of the neuron such as the firing rate during a burst. The investigations can also suggest that the hand tuning of the neuron parameters did not lead to an instable point in the parameter space, where small changes already lead to a completely different behavior, but instead yielded a set of neuron parameters, that reliably produce similar behavior under small disturbances.

8 Discussion

Summary In this thesis we introduced three different single-neuron experiments that can be used to characterize neuron circuits in their development phase. Each experiment is focused on a different aspect of the neuron's functionality.

First, the single-neuron modeling competition initiated by Jolivet et al. (2008b) which poses the challenge for mathematical neuron models to predict the precise spike times of a biological neuron given an input current. Here, we employed an optimization algorithm to determine neuron parameters for the Adaptive Exponential Integrate-and-Fire neuron model in this competition, which complements the findings and solves some of the problems in Jolivet et al. (2008a). Additionally, we extended the competition by spike-based data sets in order to make it suitable for the testing of accelerated neuromorphic hardware neurons.

Secondly, the data set of the competition are supplemented with a set of biologically relevant firing patterns introduced by Naud et al. (2008) which can be used to test neuron circuits for their adaptation and exponential parameter ranges and precise tuneability of the circuits.

Finally, the backpropagation-activated calcium spike (BAC) firing mechanism, measured by Larkum et al. (1999), was used as an exemplary test case for multi-compartment and plateau potential circuits.

The collection of single-neuron experiments was used to characterize the neuron circuits in the development phase of HICANN DLS 3. To this end, we used two complementary simulation back-ends, the NEST-based Exess neuron and transistor-level simulations of the neuron circuits. The former allows fast, idealized simulations while the latter provides the basis for detailed investigations.

Before the application of the single-neuron experiments a general verification of the functionality and characterization of the neuron circuits was performed. The verification process of the new circuits uncovered errors that would have severely impaired the usability of the neuron if they had stayed unnoticed.

We found that, due to the possibility of the adaptation and membrane voltage rising above 1.2 V, thin-oxide transmission gates do not electrically separate these subcircuits in a reliable way. Additionally, it was discovered that the usage of thin-oxide transmission gates in the multiplexers at the input of the adaptation OTA is harmful, as the occurring leakage current flows onto the capacitive memory and alters the stored value for the parameter $V_{\text{leak,adapt}}$. These issues were addressed by replacing the problematic thin-oxide transmission gates with thick-oxide transmission gates, which show no significant leakage for all possibly occurring voltages.

In the original design the voltage steps applied to V_{w} to model the spike-triggered adaptation were significantly smaller than the values required for practical applications. This was due to a wrongly chosen pulse length of the signal controlling the spike-triggered adaptation and was solved by introducing another signal. However, the observed voltage steps ΔV_{w} were still smaller than theoretically expected, which was due to a too large

voltage drop over diodes separating the adaptation circuits from the capacitive memory. By removing one of these diodes, the voltage steps increased.
Finally, it was found that the bypass signal was inverted, which was corrected by the insertion of an additional inverter into the signal path. All circuit modifications were performed by the designers of the neuron circuits.
In addition to the uncovered design errors, the verification process allowed to confirm the basic functionality of the sign-switching mechanism in the adaptation circuit, the exponential term and its weight scaling parameter as well as the multi-compartment connections and the plateau potential mechanism.
In contrast to the verification process, the characterization of the circuits included more detailed investigations which are more focused on practical applicability of the circuits.
We found that the new reset current produced by the combined leak-reset OTA depends strongly on the bias settings of the OTA and the chosen reset potential. Especially in the slow pMOS corners the highest achievable reset current is critically low, which can lead to an effectively changed reset potential and in the worst case to non-functional reset with a continuously spiking neuron. As the reset current depends on the bias currents of the OTA this also shows the need for a calibration of the reset circuit, as it is critical to choose the bias settings which allow a stable and robust operation.
A detailed review of the adaptation circuit showed that the adaptation conductances that can be achieved with the adaptation OTA range from comparatively low values as required for example for the AdEx firing patterns to approximately 50 times the highest value used in the firing patterns. It was found that combinations of low adaptation conductances with strong spike-triggered adaptation can not be realized well with the current design. The adaptation time constant can cover a wide range from below 10 ms to far above 200 ms which fits well to range required by the firing patterns.
The exponential term was characterized by an investigation of the realized parameters of the AdEx model V_T and Δ_T. Δ_T can not be configured on HICANN DLS 3. Its value was determined as approximately $\Delta_\mathrm{T,hw} \approx 150\,\mathrm{mV}$. V_T can be configured to seven different values using the digital exponential weight parameter. We found that the circuit is strongly susceptible to production variations and requires a calibration.
The performing of the spike-based extension of the single-neuron modeling competition using the Exess neuron allowed on the one hand to verify the implementation of the Exess neuron. On the other hand an investigations on the impact of discretized spike times, as they occur in hardware due to digital input signals, in the input of the competition could be performed. We investigated two possible discretization schemes, one where spikes were dropped, if multiple spikes at the same time step occurred, and one where multiple input synapses were used to ensure, that no spike loss occurred. The simulations showed that if no input spikes are dropped due to the binning, for the evaluated data sets and neuron parameters time steps of up to $dt_\mathrm{bio} \leq 1\,\mathrm{ms}$ still allow a $\Gamma \geq 0.9$.
The transistor-level simulations of the full neuron circuit are computationally intensive, which prohibits the usage of the full competition data sets due to long simulation times. On a subset of the competition data sets we found that for some data sets the synaptic input strength is too low which can be partly compensated by the lowering of the firing

threshold. For the remaining data sets the performance is increased for a higher threshold as this partly compensates for the observed reduced leak conductance and saturated adaptation current.

With the new adaptation and exponential circuits it was possible to qualitatively reproduce 7 out of 8 of the firing patterns introduced in (Naud et al., 2008). In particular, the newly implemented sign-switching mechanism allowed the reproduction of two patterns that were not possible in previous hardware generations (Tran, 2013). As Δ_T can not be configured a straight forward parameter translation from the given biological parameters to the hardware domain is not possible and therefore the parameters for the HICANN DLS 3 simulation needed to be tuned individually. A parameter set that produces the irregular firing pattern was not found, as the corresponding volume in the parameter space is small and the long simulation durations are prohibitive for extensive parameter sweeps.

The last of the performed single-neuron experiments, the BAC firing mechanism was simulated using both the Exess neuron and the HICANN DLS 3 simulation. The simulation of the multi-compartment neuron was done by connecting Exess neurons with the readily available gap junctions implemented in NEST. This allowed the reproduction of the BAC firing mechanism with a three-compartment Exess neuron. This shows that the modeling of intercompartment conductances using the gap junction model in NEST as well as the implemented plateau potential mechanism for the Exess neuron function as expected.

Finally, four neuron circuits in the HICANN DLS 3 simulation were connected to form a three-compartment neuron that was able to qualitatively reproduce the BAC firing mechanism as well. Additional simulations showed that the mechanism is susceptible to temperature variations smaller than 10 K. It was also shown, by adding parameter variations of the magnitude expected from the capacitive memory, that the tuning of the neuron parameters did not lead to an unstable point in parameter space, but that the found behavior is reliably reproduced under small parameter fluctuations.

Conclusion The described improvements to the HICANN DLS 3 neuron circuits demonstrate the value of full-neuron-circuit simulations during the development phase of analog neuron circuits. These simulations accompany the verification of the subcircuits usually performed by the designers and allow to uncover harmful interdependencies between subcircuits that can otherwise stay unnoticed.

Additionally, the goal of performing a small-scale application already during the design phase enables a verification and characterization focused on functionality as well as configurability and usability. An example for this is the spike-triggered adaptation, where the calculation of the required values for ΔV_w for the AdEx firing patterns showed the mismatch between achievable and required parameter range. The application of the firing patterns also revealed that the present implementation is not suited for the emulation of small adaptation conductances in combination with strong spike-triggered adaptation. The fact that this information on necessary improvements is available now, before the first measurement on the chip is performed, allows to include them in the planning of experiments as well as in the design of future chip versions.

The investigation of possibly harmful effects such as temperature variations in the context of an application, in this case the BAC firing mechanism, yields valuable information in comparison to a simple investigation of the temperature stability of a circuit. In the case of the exponential term the simulations of the effect of temperature changes revealed a variation in the parameters of below 5 % for a temperature change of 10 K. However, the demonstration in the BAC firing mechanism showed that, although the variation is low, the impact on the neuron behavior can still be significant.

The extension of the single-neuron modeling competition by a spike-based data set and a set of AdEx neuron references made it applicable to the development of our analog neuron circuits. However, as at the moment the simulation times are too long for practical usage of the full data set, a way for faster execution has to be found. Already with the limited data sets, important information could be gained: The simulation setup for the HICANN DLS 3 needs to be extended to allow for more synapses. Additionally, we found that the saturation of the adaptation OTA, which was observed in the characterization of the circuit, indeed impairs the performance of the neuron and an improvement of the circuit should be considered.

The strategy of using functional experiments with simulated neuron circuits as component of the hardware design process is validated by the results outlined in this thesis. The continuation of this strategy is encouraged for future hardware generations.

9 Outlook

With the HICANN DLS 3 chip soon being available for the first measurements, work can continue by reproducing the single-neuron experiments described in this thesis in hardware measurements. As in simulation, this should be preceded by general functionality tests similar to the ones described in section 5 and section 6. The measurements to be performed can confirm the simulation results or, if errors are uncovered that were not visible in the transistor-level simulations, point to necessary improvements in the simulation setup. In parallel to the reproduction of the simulation results, already confirmed issues, as for example the bypass mode which can only be triggered by multiple input events in short succession, can be improved for the next circuit generation.

The execution of the spike-based extension of the single-neuron modeling competition on HICANN DLS 3 must be preceded by an implementation of calibration algorithms for the neuron circuits. During that process the data sets can serve as a test setup for the calibration. Due to the accelerated nature of the circuits, the long execution times observed in the transistor-level simulations do not pose a problem on hardware. This should allow to run the full data sets. Additionally, the availability of 32 synapses per neuron allows to avoid the issue of truncated synaptic weights by using multiple synapses. However, in comparison to transistor-level simulations, it is more difficult in hardware measurements to find the causes of decreased performance such as the saturating adaptation current. In these cases, the hardware measurements can be accompanied by fast software simulations of the Exess neuron, where the impairing effects can be investigated in an isolated fashion. To this end, the implementation of the Exess neuron needs to be extended to allow for the modeling of hardware restrictions. One aspect, a more realistic modeling of OTA circuits, which allows for example to simulate their saturation, has already been implemented by Maria Susanna Fuhrmann (Fuhrmann, in prep.).

For the reproduction of the AdEx firing patterns on chip the circuit parameters found in the transistor-level simulation can be used as a starting point, however, as the simulations were performed for a typical neuron, the parameters will most likely not directly produce the desired patterns without additional tuning. Due to fast emulation times the parameter tuning and the necessary parameter space explorations are significantly less time-consuming and might even allow to find a parameter set for the irregular firing pattern, which was not possible in simulation due to prohibitively long simulation times.

If the newly designed multi-compartment and plateau potential circuits are functional as suggested by the performed simulations, the BAC firing mechanism is a natural test case of moderate complexity. As the BAC firing mechanism showed to be sensitive to temperature variations in simulation, a detailed measurements on the magnitude of possible temperature effects and especially their dependence on PPU operation, on the chip should be performed. Additionally, it should be evaluated to what extend the circuit behavior is stable within the temperature range that can be measured during normal operation.

The next step is to include the BAC firing mechanism in a practical application example that uses the capability of a single neuron to act as a coincidence detector between its dendritic and somatic input. This can yield information on the necessity of possible circuit

extensions such as, for example, plateau potential lengths that vary depending on the intensity of glutamatergic stimulation (Antic et al., 2010). A high-level description of a use case can be found in Larkum (2013), where it is hypothesized that the BAC firing mechanism can be the basis for associative computations in the cortex.
A key challenge for accelerated spiking neuromorphic hardware is to realize efficient learning mechanisms for computational tasks. A mapping of a backpropagation algorithm to the HICANN chip has been demonstrated by Schmitt et al. (2017) with the weight updates being calculated off-chip. To fully exploit the speedup of the hardware system the calculation of the weight updates must be performed on chip, as has been demonstrated by Friedmann et al. (2016).
One possible candidate for a hardware experiment can be found in Schiess et al. (2016), where the authors describe a plasticity rule for supervised and reinforcement learning that makes use of a multi-compartment neuron model with active dendrites. A mapping of this learning rule onto the HICANN DLS can allow to explore the computational potential of the newly implemented multi-compartment and plateau potential circuits.

Appendices

A Additional Figures

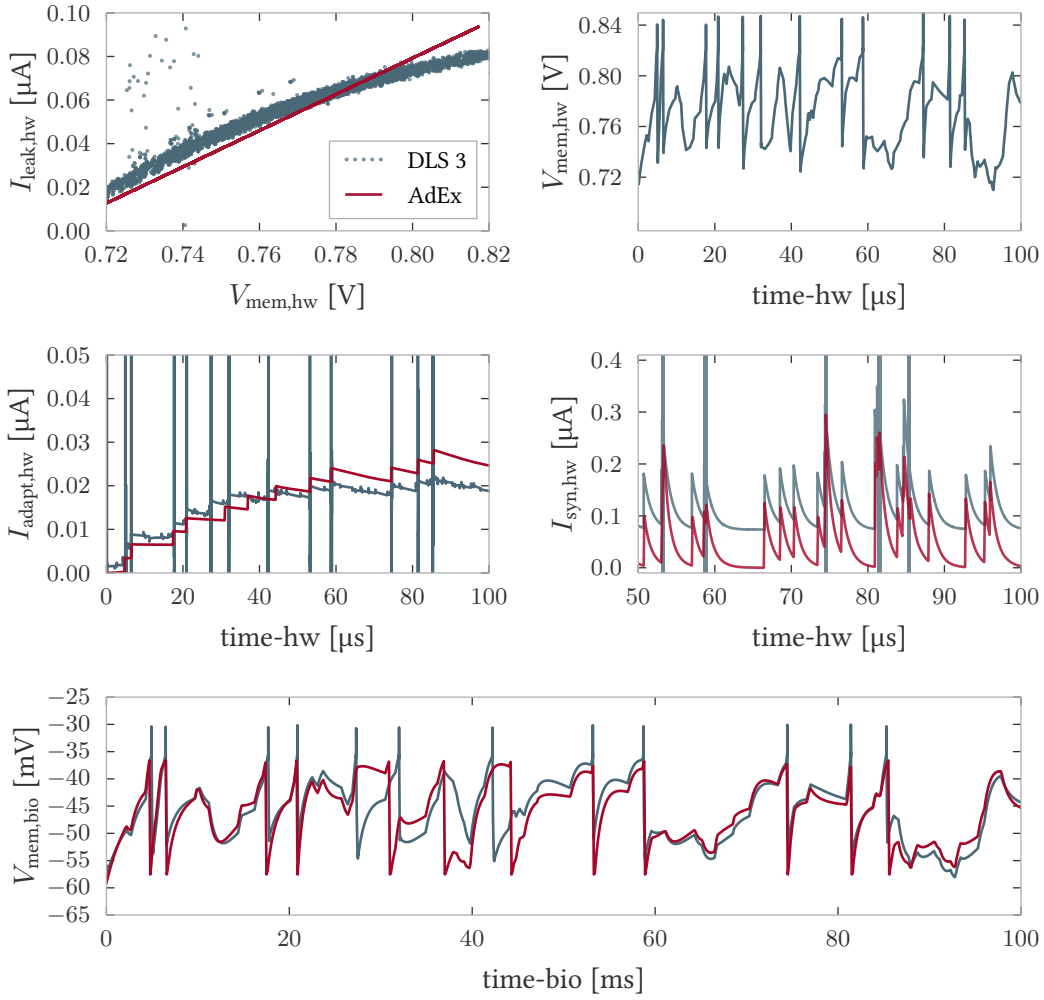

Figure 64: Comparison between the circuit simulation of HICANN DLS 3 and the AdEx reference in the fourth data set, which shows a high performance. Top: On the left a comparison of the leak current over the membrane voltage in hardware domain is shown. In contrast to the reference (red) leak conductance (slope) is not constant for the circuit (blue). The data is extracted from a transient simulation and thus noisy. The right plot shows the membrane voltage in the circuit simulation in hardware domain. Center: The left plot shows a comparison of the adaptation currents (colors as before). The right plot shows that the excitatory synaptic weight and the synaptic time constant match the reference well. Bottom: Comparison of the membrane potentials over time.

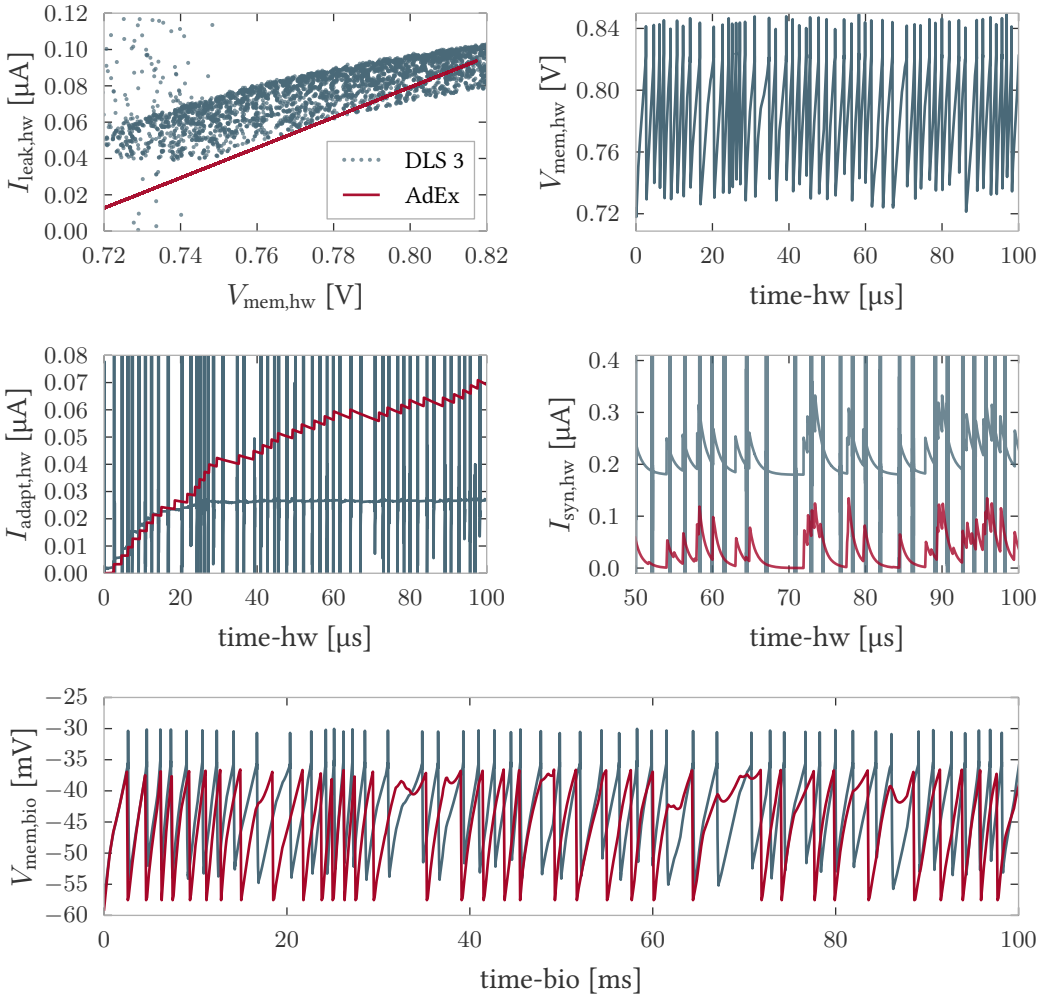

Figure 65: Comparison between the circuit simulation of HICANN DLS 3 and the AdEx reference in the ninth data set, which shows the lowest performance. Top: On the left a comparison of the leak current over the membrane voltage in hardware domain is shown. In contrast to the reference (red) leak conductance (slope) is not constant for the circuit (blue). The data is extracted from a transient simulation and thus noisy. The right plot shows the membrane voltage in the circuit simulation in hardware domain. Center: The left plot shows a comparison of the adaptation currents (colors as before). The right plot shows that the excitatory synaptic weight and the synaptic time constant match the reference well. Bottom: Comparison of the membrane potentials over time.

B Neuron parameters for AdEx firing patterns

B.1 PyNN

PyNN parameter	PyNN unit	A	B	C	D	E	F	G	H
cm	nF	0.2	0.2	0.13	0.2	0.2	0.1	0.1	0.1
tau_m	ms	5	5	5	5	5	5	5	5
v_rest	mV	-70	-70	-58	-58	-70	-65	-65	-60
v_reset	mV	-58	-58	-50	-46	-58	-47	-47	-48
v_thresh	mV	-50	-50	-50	-50	-50	-50	-50	-50
delta_T	mV	2	2	2	2	2	2	2	2
a	nS	2	2	4	2	-10	-10	15	-11
b	nA	0.0	0.06	0.12	0.1	0.0	0.03	0.3	0.03
tau_w	ms	30	300	150	120	300	90	90	130
tau_refrac	ms	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
I_stim	nA	0.5	0.5	0.4	0.41	0.21	0.11	0.35	0.16

Table 4: Collection of PyNN parameters for the AdEx firing patterns, originally from Naud et al. (2008), modified to match figures in Naud et al. (2008) by Paul Müller. A: tonic spiking, B: adaptation, C: initial bursting, D: regular bursting, E: delayed accelerated, F: delayed regular bursting, G: transient spiking, H: irregular firing.

B.2 HICANN DLS 3

DLS 3 parameter	Unit	A	B	C	D	E	F	G
v_leak	V	0.60	0.60	0.60	0.60	0.60	0.60	0.60
v_leak_adapt	V	0.60	0.60	0.60	0.60	0.60	0.60	0.60
v_reset	V	0.72	0.62	0.78	0.80	0.72	0.806	0.806
v_thresh	V	1.00	1.00	0.83	1.00	1.00	1.00	0.95
i_bias_leak	µA	1.00	1.00	1.00	1.00	1.00	1.00	1.00
i_bias_leak_sd	µA	0.74	0.76	0.76	0.77	0.76	0.77	0.77
i_bias_res	µA	1.00	1.00	1.00	0.70	1.00	1.00	1.00
i_bias_res_sd	µA	1.00	1.00	1.00	1.00	1.00	1.00	1.00
i_bias_adapt	µA	1.00	1.00	1.00	1.00	1.00	1.00	1.00
i_bias_adapt_sd	µA	0.05	0.25	0.25	0.25	0.21	0.22	0.25
i_bias_adapt_res	µA	0.50	0.11	1.00	0.05	0.08	0.15	0.15
i_bias_adapt_w	µA	0.02	0.05	0.08	0.14	0.02	0.14	0.35
i_stim	µA	0.071	0.077	0.133	0.080	0.025	0.030	0.094
en_neg_va	digital	False	False	False	False	True	True	False
en_pos_vw	digital	True	True	True	True	True	False	True
en_exp	digital	False	True	True	True	True	True	True
exp_weight_b	digital	7	6	7	5	7	6	7
en_mem_cap	digital	(True, True, True, True, True, True)						
en_adapt	digital	(True, True)						
en_ana_in	digital	True						
highs_leak	digital	False						
highs_res	digital	True						
en_syn_i_exc	digital	False						
en_syn_i_inh	digital	False						
i_bias_syn_gm_exc	µA	0.02						
i_bias_syn_gm_inh	µA	0.02						
i_bias_syn_res_exc	µA	0.8						
i_bias_syn_res_inh	µA	0.8						
holdoff_time	clk cycles	1						
adaptation_time	clk cycles	13						
refrac_time	clk cycles	15						
refrac_clk_freq	MHz	10						

Table 5: Collection of HICANN DLS 3 parameters for the AdEx firing patterns. A: tonic spiking, B: adaptation, C: initial bursting, D: regular bursting, E: delayed accelerated, F: delayed regular bursting, G: transient spiking. As these experiments are performed with a single neuron, all multi-compartment switch parameters are set to False.

C Neuron parameters for BAC firing

C.1 Exess

PyNN parameter	PyNN unit	NMDA	Ca	Soma
cm	nF	0.2	0.2	0.2
tau_m	ms	10.0	10.0	10.0
v_rest	mV	-70.0	-70.0	-70.0
v_reset	mV	-25.0	-25.0	-70.6
v_thresh	mV	-55.0	-55.0	-55.1
v_spike	mV	-54.9	-54.9	-5.0
delta_T	mV	1.0	1.0	1.0
a	nS	4.0	4.0	4.0
b	nA	0.08	0.08	0.08
tau_w	ms	150.0	150.0	150.0
tau_refrac	ms	60.0	60.0	1.0
tau_syn_E	ms	5.0	5.0	5.0
E_rev_E	mV	0.0	0.0	0.0
w_exc	µS	0.02	0.0	0.0

Table 6: Collection of PyNN parameters for the BAC firing mechanism, which were then translated into exess parameters using the transformations in section 4.1. For the transformation $\alpha_v = 13.33$ and $\omega_v = 1.27$ was used. The intercompartment conductances in hardware domain were set to $g_{\text{icc}} = 0.5\,\mu\text{S}$ and the current stimulus for the soma was $I_{\text{stim}} = 0.25\,\mu\text{A}$.

C.2 HICANN DLS 3

DLS 3 parameter	Unit	NMDA	Ca_1	Ca_2	Soma
v_leak	V	0.65	0.65	0.65	0.60
v_leak_adapt	V	0.65	0.65	0.65	0.60
v_reset	V	1.20	1.20	1.20	0.60
v_thresh	V	0.75	0.75	0.75	1.10
i_bias_leak	µA	1.00	0.80	0.80	0.80
i_bias_leak_sd	µA	1.00	0.80	0.80	0.80
i_bias_res	µA	1.00	1.00	1.00	1.00
i_bias_res_sd	µA	1.00	1.00	1.00	1.00
i_bias_adapt	µA	1.00	1.00	1.00	1.00
i_bias_adapt_sd	µA	0.02	0.02	0.02	0.02
i_bias_adapt_res	µA	0.02	0.02	0.02	0.02
i_bias_adapt_w	µA	0.00	0.00	0.00	0.00
i_bias_nmda	µA	0.70	1.00	0.90	1.00
i_stim	µA	0.00	0.00	0.00	1.50
en_neg_va	digital	False	False	False	False
en_pos_vw	digital	True	True	True	True
en_exp	digital	False	False	False	True
en_nmda	digital	True	False	True	False
en_scon	digital	True	False	True	False
en_soma	digital	False	True	False	True
en_bot	digital	False	False	False	False
en_right	digital	False	True	False	False
en_ana_in	digital	False	False	False	True
exp_weight_b	digital	0	0	0	0
holdoff_time	clk cycles	10	10	10	1
refrac_time	clk cycles	30	30	30	10
refrac_clk_freq	MHz	1	1	1	10
en_mem_cap (soma)	digital	(True, True, False, True, False, False)			
en_mem_cap (others)	digital	(False, False, False, False, True, True)			
en_adapt	digital	(False, False)			
highs_leak	digital	False			
highs_res	digital	True			
en_syn_i_exc	digital	True			
en_syn_i_inh	digital	True			
ib_nmda_mul4	digital	True			
i_bias_syn_gm_exc	µA	1.00			
i_bias_syn_gm_inh	µA	1.00			
i_bias_syn_res_exc	µA	0.1			
i_bias_syn_res_inh	µA	0.1			
adaptation_time	clk cycles	1			

Table 7: Collection of HICANN DLS 3 parameters for BAC firing.

D Single-Neuron Modeling Competition

D.1 Differential Evolution

Parameter	AdEx 0	AdEx 1	AdEx 2
τ_{m} [ms]	2.84	2.63	2.49
E_{L} [mV]	-59.05	-59.18	-58.81
V_{T} [mV]	-37.13	-38.60	-39.36
V_{th} [mV]	-36.53	-37.43	-37.92
V_{reset} [mV]	-57.55	-60.56	-60.85
Δ_{T} [mV]	2.51	1.74	1.86
τ_{w} [ms]	89.4	62.3	91.6
a [nS]	1.67	1.09	2.49
b [nA]	0.012	0.016	0.010
C_{m} [pF]		0.044	
$E_{\mathrm{rev,exc}}$ [mV]		0.0	
$E_{\mathrm{rev,inh}}$ [mV]		-80.0	
τ_{refrac} [ms]		0.1	

Table 8: Parameter sets determined by the differential evolution algorithm. The parameters in the lower part were fixed during evolution ($E_{\mathrm{rev,exc/inh}}$ as no spike inputs is used in simulation, C_{m} as only the leak conductance needs to be determined which is $g_{\mathrm{L}} = \frac{\tau_{\mathrm{m}}}{C_{\mathrm{m}}}$. As τ_{m} evolves the membrane capacitance can be kept constant).

Parameter	Min	Max
V_{T} [mV]	-41.0	-37.0
V_{th} [mV]	-45.0	-15.0
V_{reset} [mV]	-65.0	-45.0
Δ_{T} [mV]	0.0	5.0
τ_{w} [ms]	50.0	200.0
a [nS]	-10.0	10.0
b [nA]	0.0	0.15

Table 9: Parameter ranges within which the neuron parameters can evolve. Wherever possible the parameter ranges were narrowed down by comparison to the voltage recordings of the biological neuron (e.g. spike threshold and reset voltage).

D.2 Input Parameters for Spike-Based Extension

Parameter	0	1	2	3	4	5
w_{exc} [µS]	5.0e-3	8.8e-3	1.3e-2	5.0e-3	8.8e-3	1.3e-2
w_{inh} [µS]	5.0e-3	8.8e-3	1.3e-2	5.0e-3	8.8e-3	1.3e-2
I_{offset} [nA]	0.1	0.1	0.1	0.3	0.3	0.3
Parameter	6	7	8	9	10	
w_{exc} [µS]	5.0e-3	8.8e-3	1.3e-2	5.0e-3	1.3e-2	
w_{inh} [µS]	5.0e-3	8.8e-3	1.3e-2	5.0e-3	1.3e-2	
I_{offset} [nA]	0.5	0.5	0.5	0.7	0.7	

Table 10: Input parameters for spike-based extension for the data sets with input frequency $\nu = 500\,\text{Hz}$ and $\tau_{\text{syn}} = 1\,\text{ms}$.

D.3 HICANN DLS 3 Neuron parameters

DLS 3 parameter	Unit	AdEx0
v_leak	V	0.705
v_leak_adapt	V	0.705
v_reset	V	0.713
v_thresh	V	0.831
i_bias_leak	µA	1.000
i_bias_leak_sd	µA	0.827
i_bias_res	µA	0.777
i_bias_res_sd	µA	1.000
i_bias_adapt	µA	1.000
i_bias_adapt_sd	µA	0.142
i_bias_adapt_res	µA	0.085
i_bias_adapt_w	µA	0.699
en_neg_va	digital	False
en_pos_vw	digital	True
en_exp	digital	False
exp_weight_b	digital	7
en_mem_cap	digital	(True, True, True, True, True, True)
en_adapt	digital	(True, True)
en_ana_in	digital	False
highs_leak	digital	False
highs_res	digital	True
en_syn_i_exc	digital	False
en_syn_i_inh	digital	False
i_bias_syn_gm_exc	µA	1.000
i_bias_syn_gm_inh	µA	1.000
i_bias_syn_res_exc	µA	0.550
i_bias_syn_res_inh	µA	0.565
holdoff_time	clk cycles	1
adaptation_time	clk cycles	1
refrac_time	clk cycles	3
refrac_clk_freq	MHz	10

Table 11: Neuron parameters for the typical neuron used in the spike-based extension of the single-neuron modeling competition. They were determined from the first AdEx parameters in table 2 using the calibration algorithms.

E Software Versions and HICANN DLS 3 Schematic

Experiment	hicann-dls-fc	hicann-dls-testbench	single-neuron-benchmarks
Figure 10	298ba44b8d	832095da5c	e2fedee2193
Figure 25	a6b81712b9a	46641365ab2	e2fedee2193
Figure 26	a6b81712b9a	46641365ab2	e2fedee2193
Figure 27 (left)	a6b81712b9a	4051733eded	c5953085720
Figure 27 (right)	d771a338707	984d3a5639f	e8043311465
Figure 29	d771a338707	984d3a5639f	e8043311465
Figure 30	d771a338707	984d3a5639f	e8043311465
Figure 31	298ba44b8d	832095da5c	e2fedee2193
Figure 32	298ba44b8d	832095da5c	e2fedee2193
Figure 33 - Figure 45	d771a338707	984d3a5639f	e8043311465
Figure 51 - Figure 57	d771a338707	984d3a5639f	e8043311465
Figure 60	d771a338707	984d3a5639f	e8043311465
Figure 61 - Figure 63	15f87b990ca	99666e699c8	c814220c02a

Table 12: Git commit ID for the used repositories to create corresponding figures. The neuron schematic is in the hicann-dls-fc repository, the neuron testbench and calibration algorithms are in hicann-dls-testbench and the experiments in single-neuron-benchmarks.

References

Instructions for participants 2009. http://incf.org/community/competitions/archive/spike-time-prediction/2009/InstructionForParticipants2009.pdf/view, 2016. Accessed: 2016-02-27.

Script for gamma measure. http://icwww.epfl.ch/~gerstner/QuantNeuronMod2007/GamCoincFac.m, 2016. Accessed: 2016-02-27.

Challenge data 2007. http://icwww.epfl.ch/~gerstner/QuantNeuronMod2007/challenge.html, 2017. Accessed: 2017-01-26.

Syed Ahmed Aamir, Paul Müller, Andreas Hartel, Johannes Schemmel, and Karlheinz Meier. A highly tunable 65-nm cmos lif neuron for a large scale neuromorphic system. In *European Solid-State Circuits Conference, ESSCIRC Conference 2016: 42nd*, pages 71–74. IEEE, 2016.

Srdjan D Antic, Wen-Liang Zhou, Anna R Moore, Shaina M Short, and Katerina D Ikonomu. The decade of the dendritic nmda spike. *Journal of neuroscience research*, 88(14):2991–3001, 2010.

BayesElo-rating. Woldranking go. https://www.goratings.org/, 2017. Accessed: 2017-03-06.

Jonathan Born. Distanzmaße für spiketrains im rahmen der optimierung von neuronparametern. Bachelor thesis (German), University of Heidelberg, HD-KIP 12-18, 2012.

BrainScaleS. Research project. http://brainscales.kip.uni-heidelberg.de/public/index.html, 2012.

Romain Brette and Wulfram Gerstner. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *Journal of neurophysiology*, 94(5):3637–3642, 2005.

Laure Buhry, Filippo Grassia, Audrey Giremus, Eric Grivel, Sylvie Renaud, and Sylvain Saïghi. Automated parameter estimation of the hodgkin-huxley model using the differential evolution algorithm: application to neuromimetic analog integrated circuits. *Neural computation*, 23(10):2599–2625, 2011.

Nicholas T Carnevale and Michael L Hines. *The NEURON book*. Cambridge University Press, 2006.

Andrew Davison, Daniel Brüderle, Jens Kremkow, Eilif Muller, Dejan Pecevski, Laurent Perrinet, and Pierre Yger. Pynn: a common interface for neuronal network simulators. 2009.

Peter Dayan and Laurence F Abbott. *Theoretical neuroscience*, volume 806. Cambridge, MA: MIT Press, 2001.

DeepMind web page. Google deepmind - alphago. https://deepmind.com/research/alphago/, 2017. Accessed: 2017-03-05.

Patrick G Drennan and Colin C McAndrew. Understanding mosfet mismatch for analog design. *IEEE Journal of Solid-State Circuits*, 38(3):450–456, 2003.

Simon Friedmann, Johannes Schemmel, Andreas Grübl, Andreas Hartel, Matthias Hock, and Karlheinz Meier. Demonstrating hybrid learning in a flexible neuromorphic hardware system. *IEEE Transactions on Biomedical Circuits and Systems*, 2016.

Maria Susanna Fuhrmann. Influence of limited otas and noise on the hicann-dls-nn3 nest model. Internship Report, University of Heidelberg, in prep.

Wulfram Gerstner and Werner M Kistler. *Spiking neuron models: Single neurons, populations, plasticity*. Cambridge university press, 2002.

Marc-Oliver Gewaltig and Markus Diesmann. Nest (neural simulation tool). *Scholarpedia*, 2(4):1430, 2007.

Jan Hahne, Moritz Helias, Susanne Kunkel, Jun Igarashi, Matthias Bolten, Andreas Frommer, and Markus Diesmann. A unified framework for spiking and gap-junction interactions in distributed neuronal network simulations. *Frontiers in neuroinformatics*, 9, 2015.

Andreas Hartel. *Implementation and Characterization of Mixed-Signal Neuromorphic ASICs*. PhD thesis, Universität Heidelberg, 2016.

HBP SP9 partners. *Neuromorphic Platform Specification*. Human Brain Project, March 2014.

Stephan Henker, Johannes Partzsch, and René Schüffny. Accuracy evaluation of numerical methods used in state-of-the-art simulators for spiking neural networks. *Journal of computational neuroscience*, 32(2):309–326, 2012.

Matthias Hock. *Modern semiconductor technologies for neuromorphic hardware*. PhD thesis, Universität Heidelberg, July 2014.

Matthias Hock. Capacitive memory specification. https://brainscales-r.kip.uni-heidelberg.de/projects/cm-spec, 2015. Accessed: 2017-02-20.

Renaud Jolivet, Ryota Kobayashi, Alexander Rauch, Richard Naud, Shigeru Shinomoto, and Wulfram Gerstner. A benchmark test for a quantitative assessment of simple neuron models. *Journal of neuroscience methods*, 169(2):417–424, 2008a.

Renaud Jolivet, Felix Schürmann, Thomas K Berger, Richard Naud, Wulfram Gerstner, and Arnd Roth. The quantitative single-neuron modeling competition. *Biological cybernetics*, 99(4-5):417–426, 2008b.

Werner M Kistler, Wulfram Gerstner, and J Hemmen. Reduction of the hodgkin-huxley equations to a single-variable threshold model. *Neural Computation*, 9(5):1015–1045, 1997.

Laura Kriener, Paul Müller, Sebastian Billaudelle, Syed Ahmed Aamir, Johannes Schemmel, and Karlheinz Meier. Neuronal dynamics: Silicon vs. biology. 1st HBP Student Conference, 2017.

Susanne Kunkel, Maximilian Schmidt, Jochen M Eppler, Hans E Plesser, Gen Masumoto, Jun Igarashi, Shin Ishii, Tomoki Fukai, Abigail Morrison, Markus Diesmann, et al. Spiking network simulation code for petascale computers. *Frontiers in neuroinformatics*, 8: 78, 2014.

Matthew Larkum. A cellular mechanism for cortical associations: an organizing principle for the cerebral cortex. *Trends in neurosciences*, 36(3):141–151, 2013.

Matthew E Larkum, J Julius Zhu, and Bert Sakmann. A new cellular mechanism for coupling inputs arriving at different cortical layers. *Nature*, 398(6725):338–341, 1999.

Sebastian Millner. Development of a multi-compartment neuron model emulation. 2012.

Richard Naud, Nicolas Marcille, Claudia Clopath, and Wulfram Gerstner. Firing patterns in the adaptive exponential integrate-and-fire model. *Biological cybernetics*, 99(4-5): 335–347, 2008.

Mihai A Petrovici, Bernhard Vogginger, Paul Müller, Oliver Breitwieser, Mikael Lundqvist, Lyle Muller, Matthias Ehrlich, Alain Destexhe, Anders Lansner, René Schüffny, et al. Characterization and compensation of network-level anomalies in mixed-signal neuromorphic modeling platforms. *PloS one*, 9(10):e108590, 2014.

Kenneth Price, Rainer M Storn, and Jouni A Lampinen. *Differential evolution: a practical approach to global optimization*. Springer Science & Business Media, 2006.

Nicolas Rougier. Biological neuron schema. https://commons.wikimedia.org/wiki/File:Neuron-figure.svg, June 2007. URL `https://commons.wikimedia.org/wiki/File:Neuron-figure.svg`. Accessed: 2017-02-23.

Johannes Schemmel, Johannes Fieres, and Karlheinz Meier. Wafer-scale integration of analog neural networks. In *Neural Networks, 2008. IJCNN 2008.(IEEE World Congress on Computational Intelligence). IEEE International Joint Conference on*, pages 431–438. IEEE, 2008.

Johannes Schemmel, Daniel Brüderle, Andreas Grübl, Matthias Hock, Karlheinz Meier, and Sebastian Millner. A wafer-scale neuromorphic hardware system for large-scale neural modeling. *Proceedings of the 2010 IEEE International Symposium on Circuits and Systems (ISCAS"10)*, pages 1947–1950, 2010.

Johannes Schemmel, Laura Kriener, Paul Müller, and Karlheinz Meier. An accelerated analog neuromophic hardware system emulating NMDA- and calcium-based non-linear dendrites. In *The 2017 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2017. To be published.

Mathieu Schiess, Robert Urbanczik, and Walter Senn. Somato-dendritic synaptic plasticity and error-backpropagation in active dendrites. *PLoS Comput Biol*, 12(2):e1004638, 2016.

Sebastian Schmitt, Johann Klähn, Guillaume Bellec, Andreas Grübl, Maurice Güttler, Andreas Hartel, et al. Neuromorphic hardware in the loop: Training a deep spiking network on the brainscales wafer-scale system. In *The 2017 International Joint Conference on Neural Networks (IJCNN)*. IEEE, 2017. To be published.

Susanne Schreiber, Jean-Marc Fellous, D Whitmer, P Tiesinga, and Terrence J Sejnowski. A new correlation-based measure of spike timing reliability. *Neurocomputing*, 52:925–931, 2003.

Arnold JF Siegert. On the first passage time probability problem. *Physical Review*, 81(4): 617, 1951.

David Silver, Aja Huang, Chris J Maddison, Arthur Guez, Laurent Sifre, George Van Den Driessche, Julian Schrittwieser, Ioannis Antonoglou, Veda Panneershelvam, Marc Lanctot, et al. Mastering the game of go with deep neural networks and tree search. *Nature*, 529(7587):484–489, 2016.

Rainer Storn and Kenneth Price. Differential evolution–a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization*, 11(4): 341–359, 1997.

Yannick Stradmann. personal communication, March 4th, 2017.

Yannik Stradmann. Characterization and calibration of a mixed-signal leaky integrate and fire neuron on hicann-dls. Bachelor thesis, 2016.

Evangelos Stromatias, Daniel Neil, Michael Pfeiffer, Francesco Galluppi, Steve B Furber, and Shih-Chii Liu. Robustness of spiking deep belief networks to noise and reduced bit precision of neuro-inspired hardware platforms. *Frontiers in neuroscience*, 9, 2015.

Binh Tran. Demonstrationsexperimente auf neuromorpher hardware. Bachelor thesis (German), University of Heidelberg, HD-KIP 13-55, 2013.

Michael C Vanier and James M Bower. A comparative survey of automated parameter-search methods for compartmental neural models. *Journal of computational neuroscience*, 7(2):149–171, 1999.

Acknowledgments

I would like to thank

- Prof. Dr. Karlheinz Meier and Dr. Johannes Schemmel for the possibility to write my thesis in this amazing research group and for the chance to experience the unique atmosphere at the workshops in Manchester, Capo Caccia and Vienna.
- Prof. Dr. Michael Hausmann for being the second referee for this thesis.
- Paul Müller for being the best advisor for my master thesis I could have wished for. Thank you!
- Gerd, Korbi and Sebastian B. for patient electronics explanations.
- The proofreaders Agnes, Ákos, Andi H., Benjamin, Mihai, Paul, Sebastian B., Sebastian S. and Yannick.
- "The entire Electronic Vision(s) Group for general support, being great coworkers and such an amazing bunch of crazy people." (Hock, 2014; Hartel, 2016)
- The tea and coffee guys (among others Andi H., Eric, Johannes B., Mihai, Paul, Sebastian B. and Vitali) for nice breaks with interesting discussions and/or relaxing chats.
- The bouldering guys (among others Alex, Andi H., Mihai, Paul, Sebastian B., Sebastian J. and Vitali) for challenging sessions at the Boulderhaus and in Fontainebleau.
- The board- and card game guys (among others Albert, Christian M., Joscha, Mihai, Paul and Vitali) for some heated discussions and amazingly fun evenings.
- My friends at home, especially Julia, for keeping up with my periodical disappearance and reappearance correlated with my work load.
- My family for always having my back.

Statement of Originality (Erklärung):

I certify that this thesis, and the research to which it refers, are the product of my own work. Any ideas or quotations from the work of other people, published or otherwise, are fully acknowledged in accordance with the standard referencing practices of the discipline.

Ich versichere, dass ich diese Arbeit selbständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt habe.

Heidelberg, March 07, 2017 .....

(signature)