

# ACCELERATED ANALOG NEUROMORPHIC COMPUTING

Johannes Schemmel, Sebastian Billaudelle, Phillip Dauer, Johannes Weis  
Heidelberg University, Heidelberg, Germany  
{schemmel}@kip.uni-heidelberg.de

## Abstract

This paper presents the concepts behind the BrainScales (BSS) accelerated analog neuromorphic computing architecture. It describes the second-generation BrainScales-2 (BSS-2) version and its most recent in-silico realization, the HICANN-X Application Specific Integrated Circuit (ASIC), as it has been developed as part of the neuromorphic computing activities within the European Human Brain Project (HBP). While the first generation is implemented in an 180 nm process, the second generation uses 65 nm technology. This allows the integration of a digital plasticity processing unit, a highly-parallel micro processor specially built for the computational needs of learning in an accelerated analog neuromorphic systems.

The presented architecture is based upon a continuous-time, analog, physical model implementation of neurons and synapses, resembling an analog neuromorphic accelerator attached to build-in digital compute cores. While the analog part emulates the spike-based dynamics of the neural network in continuous-time, the latter simulates biological processes happening on a slower time-scale, like structural and parameter changes. Compared to biological time-scales, the emulation is highly accelerated, i.e. all time-constants are several orders of magnitude smaller than in biology. Programmable ion channel emulation and inter-compartmental conductances allow the modeling of non-linear dendrites, back-propagating action-potentials as well as NMDA and Calcium plateau potentials. To extend the usability of the analog accelerator, it also supports vector-matrix multiplication. Thereby, BSS-2 supports inference of deep convolutional networks as well as local-learning with complex ensembles of spiking neurons within the same substrate. A prerequisite to successful training is the calibratability of the underlying analog circuits across the full range of process variations. For this purpose a custom software toolbox has been developed, that facilitates complex calibrated Monte-Carlo simulations.

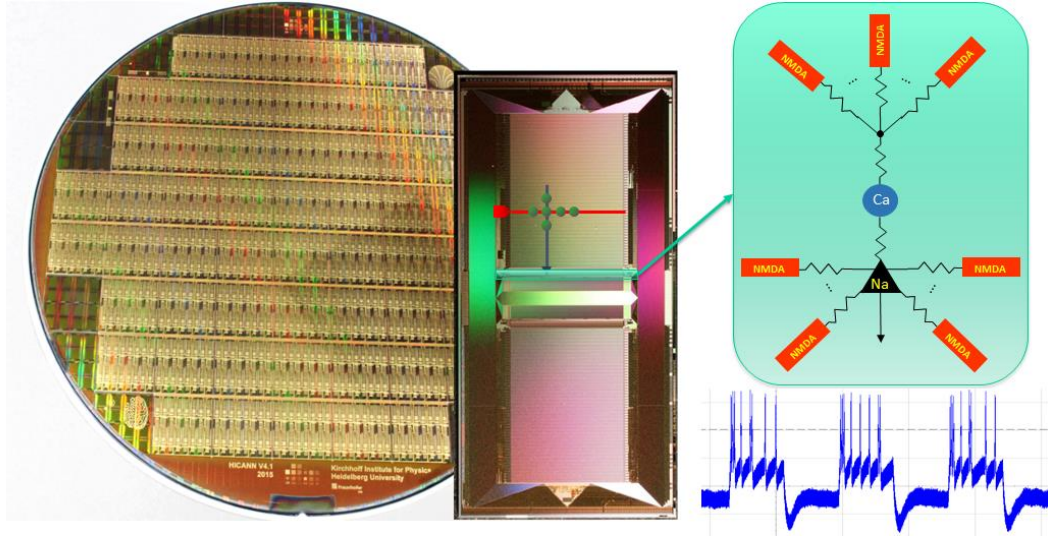


Fig. 1: Basic elements of the BrainScaleS architecture: wafer, BSS-1 ASIC, BSS-2 neuron and exemplary membrane voltage trace.

## 1. Introduction

The basic concept of the BrainScaleS systems is the emulation of biologically-inspired neural networks with physical models[1]. It differs from comparable neuromorphic approaches based on continuous-time analog circuits[2, 3, 4] in many aspects, like the high acceleration factor[5, 6], usage of wafer-scale integration[7], calibratability towards biologically-sound neuron parameters[8, 9], a software-interface based on the simulator-agnostic description language PyNN[10, 11], support for non-linear dendrites and structured neurons[12] as well as on-chip support for complex plasticity rules based on a combination of analog measurements internal analog-to-digital conversion and build-in microprocessors.

The first generation, BrainScaleS 1, has been completed[13] and is used mostly for research of connectivity aspects of large accelerated analog neural networks and the further development of wafer-scale integration technology. The main short-coming of the BrainScaleS 1 system is the rather inflexible implementation of long-term plasticity based solely on Spike-timing dependent plasticity (STDP), which has been taken over from its predecessor[6]. Already at the very beginning of the BrainScaleS project this was considered a conceptual weakness and an upgrade path was devised to implement the more flexible hybrid plasticity[14] scheme in future revisions. Due to the process technology used within BrainScaleS 1, 180 nm, it was not feasible to integrate the necessary standard cell logic without sacrificing too much area to digital circuits in relation to the analog neurons and synapses. Therefore, the decision was made to develop a second Brain-

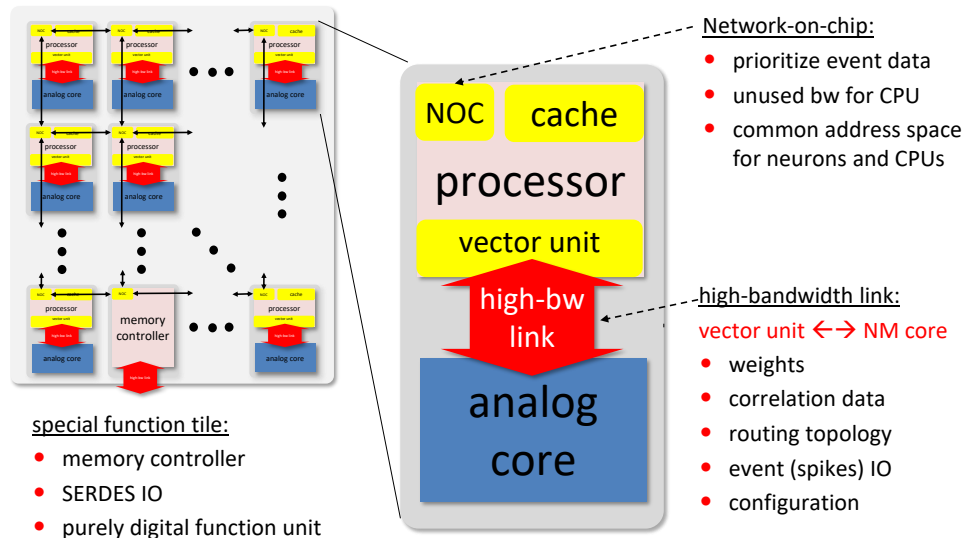


Fig. 2: A neuromorphic SOC consisting of a multitude of digital CPU cores with special vector units attached to analog Neuromorphic accelerators.

ScaleS generation, BrainScaleS 2, which is based from the beginning on a smaller process technology, namely 65 nm. Fig. 1 shows the main elements of the BSS architecture. At the very left, a BSS-1 wafer containing approx. 500 interconnected ASICs is shown. To its right, a BSS chip illustrates the characteristic layout of BSS neuromorphic chips: a central neuron area surrounded by two large synapse blocks. The sketched overlay shows the rectangular orientation of input (pre-synaptic) and output (post-synaptic) signals: the input is routed horizontally through the synapse array, while the output of the synapses connects them vertically to the neurons in the center. Next to it the graphical representation of an emulated structured neuron is shown above a measured voltage-trace from the membrane capacitor of a neuron.

One major improvement is the inclusion of a digital plasticity processor in the BSS-2 ASIC[14]. This specialized highly-parallel Single Instruction Multiple Data (SIMD) microprocessor adds an additional layer of modeling capabilities, covering all aspects of structural and parameter changes during network operation. By including the necessary logic directly within the analog network core, a communication bottleneck to the host system is avoided. This allows to scale-up all novel plasticity features for wafer-scale integration within the BrainScaleS 2 system. In the finale multi-wafer version of the BrainScaleS 2 system, which is planned to be capable of extending experiments across several hundreds of wafers, the distributed local compute capability will be even more essential. It will not only perform all levels of plasticity calculations, but also the initialization and calibration of the numerous analog mixed-signal circuits within the ASIC. The

role of the analog neural network block changes by the transition from BSS-1 to BSS-2. The analog part becomes an attachment to the CPU cores, similar to a complex accelerator. Fig. 2 illustrates this architecture.

The remainder of this publication is organized as follows: Section 2 gives an overview of the BSS-2 architecture. Section 3 presents the current prototype, the single-chip variant of BSS-2, called HICANN-X (HICANN-X). Section 4 shows some examples of the complex calibrated Monte-Carlo simulations used to verify that the analog neurons circuits are always capable of correctly emulating their biological counterparts, i.e. their calibratability under all process and device variations. The paper closes with a conclusion in Section 5.

## **2. Overview of the BSS neuromorphic architecture**

As shown in Fig. 2, the BSS architecture is based on the close interaction of digital and analog circuit blocks. Because of their primary intended function, the digital processor cores are called Plasticity Processing Units (PPUs). As the main neuromorphic component, the analog core contains synapse and neuron circuits [15, 16], analog parameter memories, PPU interfaces as well as all event related interface components.

The PPU is an embedded microprocessor core with a highly parallel SIMD unit optimized for the calculation of plasticity rules in conjunction with the analog core[17]. In the current incarnation of the BSS architecture, BSS-2, two PPUs share an analog core. This allows the most efficient arrangement of the neuron circuits in the center of the analog core. Fig. 3 depicts the individual function blocks located within the ANNCORE:

### **synapse arrays**

The total number of synapses are split up in four equally sized blocks to keep the vertical and horizontal lines traversing the sub-arrays as short as possible, thereby reducing their parasitic capacitances (see [17, 16]). Each synapse array resembles a block of static memory, with 16 memory cells located in each synapse, organized in two words of eight bits each. A synapse array also contains the sense amplifiers, precharge and write control circuits as well as word-line decoders and buffers. Thereby it can be connected directly to the digital, standard cell based parts of the chip. Two PPUs connect to the static memory interfaces of the two adjacent synapse arrays, using a fully parallel connection to the  $8 \times 256$  data lines.

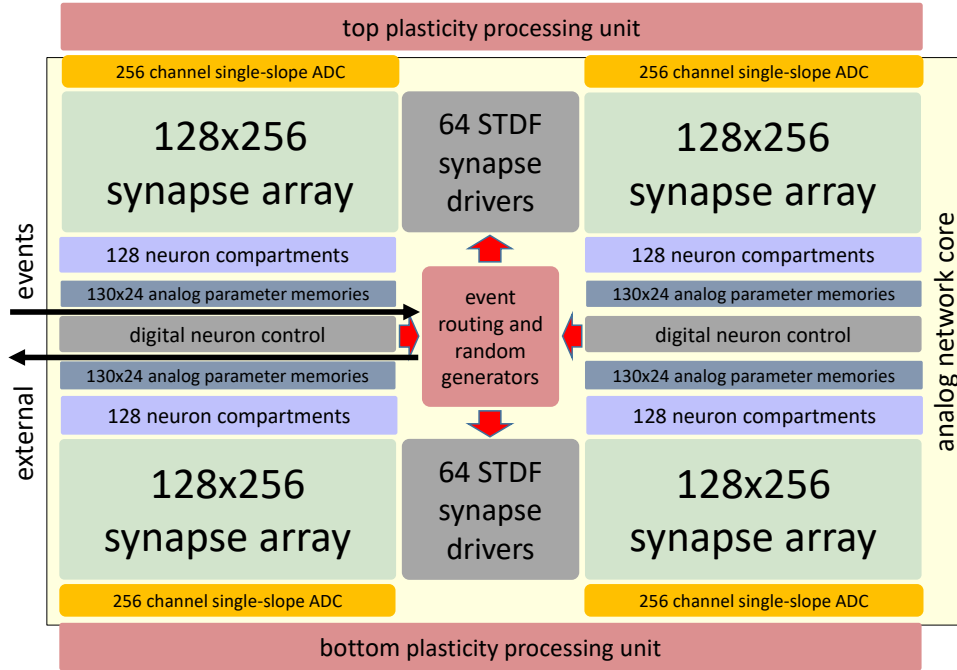


Fig. 3. Block diagram of the Analog Network Core (ANNCORE).

### neuron compartment circuits

Four rows of neuron compartment circuits are located at the edges of the synapse blocks. Each pair of dendritic input lines of a neuron compartment is connected to a column of 256 synapses. The neuron compartment implements the Adaptive-Exponential Integrate-and-Fire (AdEx) neuron model. They can be connected to form larger neurons, emulating either point or structured neurons. See [12] for more details about the multi-compartment capabilities.

### analog parameter memories

Adjacent to each row of neuron compartments is a row of analog parameter storages. These capacitive memories [18] store 24 analog values per neuron and an additional 48 global parameters. They are auto-refreshed from values stored digitally inside the memory block.

### digital neuron control

Two neuron rows share a digital neuron control block which synchronizes neural events to the digital system clock of 125 MHz and serializes them onto digital output buses.

### synapse drivers with short term plasticity

The pre-synaptic events are fed into the array via the synapse drivers. Besides timing control and buffering they contain short-term plasticity circuits emulating a simplified Tsodyks-Markram model [19, 6]. The synapse drivers

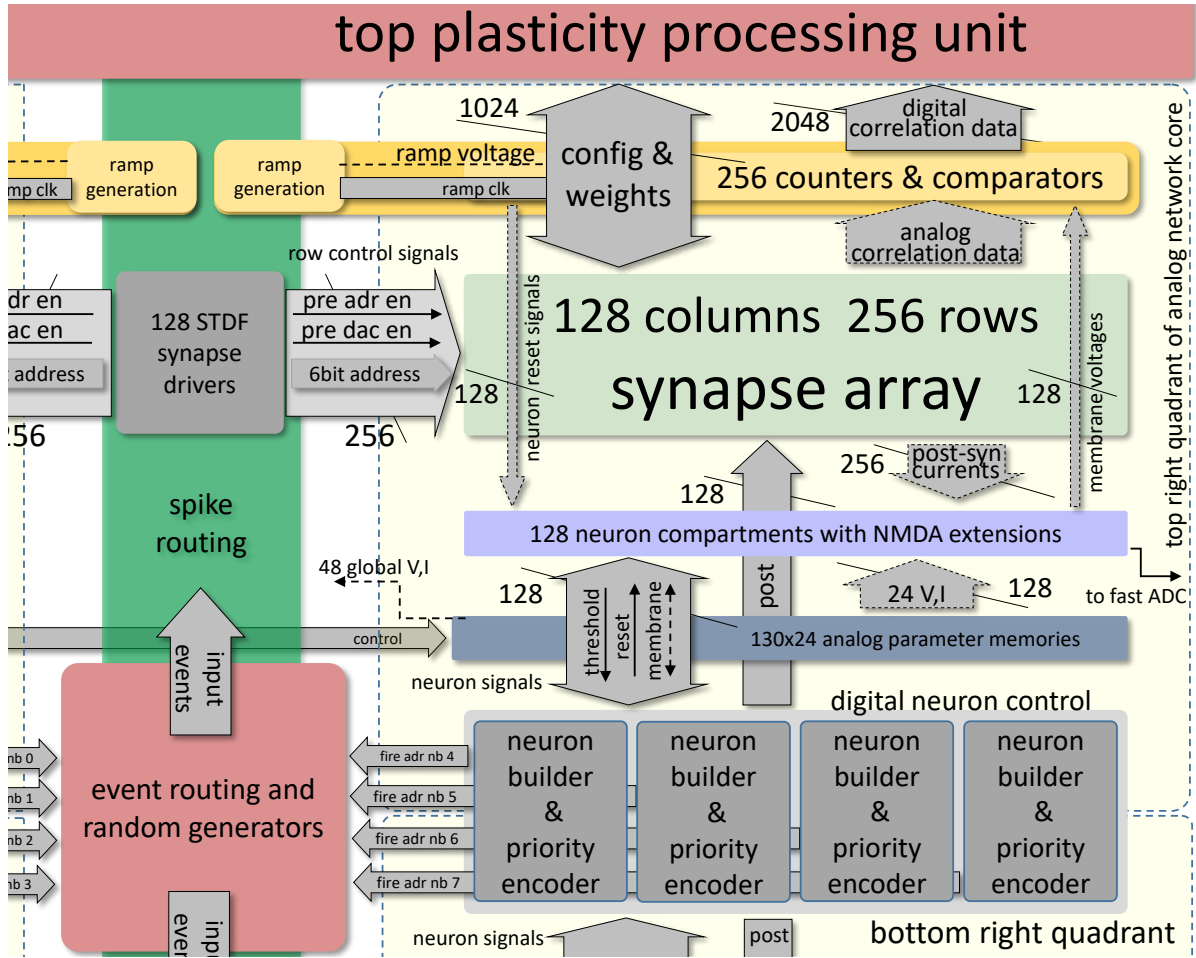


Fig. 4. Detailed block diagram of the ANNCORE's upper right quadrant.

can handle single- or multi-valued input signals, depending on the current operation mode of the synapse row, which may be either rate or spike based.

### random event generators

The random generators produce random background events fed directly into the synapse array via the synapse drivers, strongly reducing the external bandwidth usage when stochastic models [20, 21] are used.

### correlation Analog to Digital Converters (ADCs)

The top and bottom edges of the ANNCORE are lined by the SIMD units of the top and bottom PPUs. A column-parallel ADCs converts the analog data from the synapse arrays as well as selected analog signals from the neurons into the digital representations needed by the PPUs .

Fig. 4 shows a zoom-in into the upper right quadrant of the ANNCORE. For compatibility with BSS-1, the synapse drivers and digital neuron control circuits are arranged in a similar substructure as they have been previously: one synapse driver controls two rows of synapses in both adjacent blocks and the digital neu-

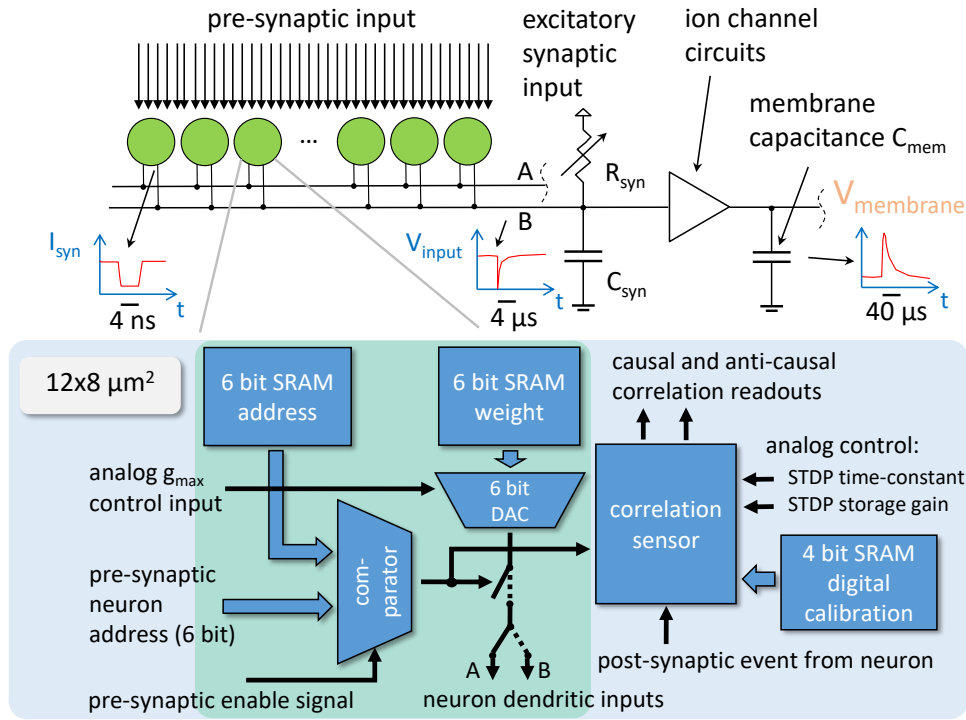


Fig. 5: Top: Operating principle and basic timing relationships of an accelerated BrainScaleS spiking neuron. Bottom: Block-diagram of a synapse.

ron control is split in eight blocks controlling 64 neuron compartments each. Four blocks are located in the left and four in the right half of ANNCORE. Each block contains the so-called neuron builder logic, which allows to interconnect analog membrane and digital spike output signals from neuron compartments being either vertically or horizontally adjacent to each other. To serialize the up-to 64 spike outputs each digital neuron control block contains priority encoder circuits that arbitrate the access to the output bus. It also contains a  $8 \times 64$  neuron source address memory[22].

The pre-synaptic input for the synapse drivers of one chip half comes from a set of local event input buses driven by the central event router. The event router within the ANNCORE mixes global, local and random event sources. In Fig. 4 the synapses are arranged in a two-dimensional array between the PPU and the neuron compartment circuits. Pre-synaptic input enters the synapse array at the left edge. For each row, a set of signal buffers transmit the pre-synaptic pulses to all synapses in the row. The post-synaptic side of the synapses, i.e. the equivalent of the dendritic membrane of the target neuron, is formed by wires running vertically through each column of synapses. At each intersection between pre- and post-synaptic wires, a synapse is located. To avoid that all neuron compartments share the same set of pre-synaptic inputs, each pre-synaptic input line transmits - in a time-multiplexed fashion - the pre-synaptic signals of up to 64 different pre-



synaptic neurons. Each synapse stores a pre-synaptic address that determines the pre-synaptic neuron it responds to.

Fig. 5 illustrates the basic operation of the BrainScaleS accelerated analog neuron and its associated synapses. Due to space limitations the dendritic column is rotated by  $90^\circ$  in the figure. The bottom half of the figure shows a block diagram of the synapse circuit. The main functional blocks are the address comparator, the Digital to Analog Converter (DAC) and the correlation sensor. Each of these circuits has its associated memory block. The address comparator receives a 6 bit address and a pre-synaptic enable signal from the periphery of the synapse array as well as a locally stored 6 bit neuron number. If the address matches the programmed neuron number, the comparator circuit generates a pre-synaptic enable signal local to the synapse (*pre*), which is subsequently used in the DAC and correlation sensor circuits. Each time the DAC circuit receives a *pre* signal, it generates a current pulse. The height of this pulse is proportional to the stored weight, while the pulse width is typically 4 ns. This matches the maximum pre-synaptic input rate of the whole synapse row which is limited to 125 MHz. The remaining 4 ns are necessary to change the pre-synaptic address. The current pulse can be shortened below the 4 ns maximum pulse length to emulate short-term synaptic plasticity [6, 23].

Each neuron compartment has two inputs, labeled A and B in Fig. 5. Usually, the neuron compartment uses A as excitatory and B as inhibitory input. Each row of synapses is statically switched to either input A or B, meaning that all pre-synaptic neurons connected to this row act either as excitatory or inhibitory inputs to their target neurons. Due to the address width of 6 bit the maximum number of different pre-synaptic neurons is 64[24]. The output currents of all synapses discharge the synaptic input capacitance  $C_{\text{syn}}$ , which is realized predominantly by the shielding capacitance of the long synaptic input wires. An adjustable Metal-Oxid Semiconductor (MOS) resistor,  $R_{\text{syn}}$ , restores the charge. Due to the short time-constant of the synaptic input pulse compared to the time constant of the synaptic input line  $\tau_{\text{input}} = C_{\text{syn}}R_{\text{syn}}$ , which is three orders of magnitude longer, the voltage trace  $V_{\text{input}}(t)$  is a single exponential.

The ion-channel circuits in BrainScaleS should implement the full AdEx neuron model, as it is the case in the BSS-1 system.[25, 26, 16]. In BSS-2 some terms are still under development at the time of this writing. The minimum configuration available in all prototype versions of BSS-2 is a set of two current-based inputs, one for inhibitory synaptic input, connected to input A in Fig. 5 and one for excitatory (input B), in combination with a leak circuit and spike and reset generation[15]. Therefore the membrane voltage is given by the standard Integrate-and-Fire (I&F) neuron model[27]. Typically, the membrane time



constant set by the leakage term is another order of magnitude above the time-constant of the synaptic input. These temporal relationships are visualized in the small timing diagram inserts in Fig.5.

The remaining functional block of the synapse shown in Fig. 5 is the correlation sensor. Its task is the measurement of the time difference between pre- and post-synaptic spikes. To determine the time of the pre-synaptic spike it is connected to the *pre* signal. The post-synaptic spike-time is determined by a dedicated signaling line running from each neuron compartment vertically through the synapse array to connect to all synapses projecting to inputs A or B of the compartment[17].

### 3. The HICANN-X chip

Although the target of the BSS architecture is wafer-scale integration, which offers a cost-effective possibility to build brain-size spiking neural network models, smaller solutions based upon single ASICs are needed to develop and debug the final design. They also shorten the time to first experiments, of which a significant proportion does need only hundreds up to a few thousand neurons and therefore does not necessarily rely on wafer-scale integration. Depending of the complexity of the neuron model they utilize, a few tens of interconnected BSS ASICs might be sufficient. To support these goals, an intermediate version of the second-generation BSS technology has been developed: suited for single- or multi-chip operation, but simultaneously prepared for later wafer-scale integration. This section will introduce said single-chip version of BSS-2, called HICANN-X, in more detail. Fig. 6 shows a block diagram of HICANN-X. In total, the HICANN-X chip uses 16 differential Low Voltage Differential Signalling (LVDS) lines for the host communication. A single chip has the same bandwidth as the full BSS-1 reticle build from eight individual chips. Using this link arrangement, the HICANN-X chip can be directly connected to one communication module of the Brain-ScaleS system, providing an easy upgrade path[28]. The layout and photograph of the chip are shown in Fig. 7.

#### 3.1. Event-routing within HICANN-X

HICANN-X uses the same two-level communication infrastructure as the first BSS generation[29]: a real-time address-event layer without handshake, called Event Link Layer 1 (Layer1), and a second layer using time-stamped event pack-

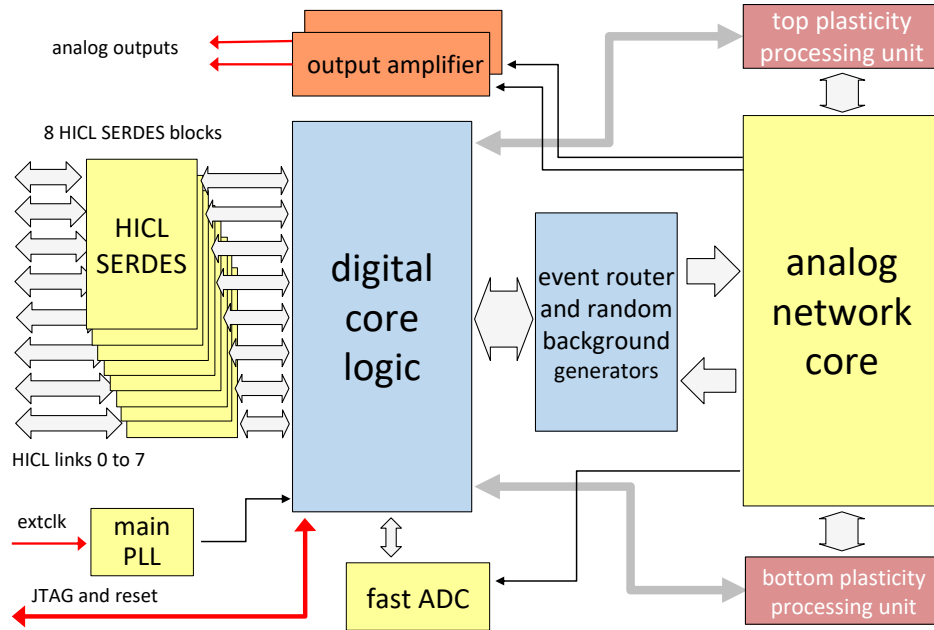


Fig. 6. Block diagram of the HICANN-X ASIC.

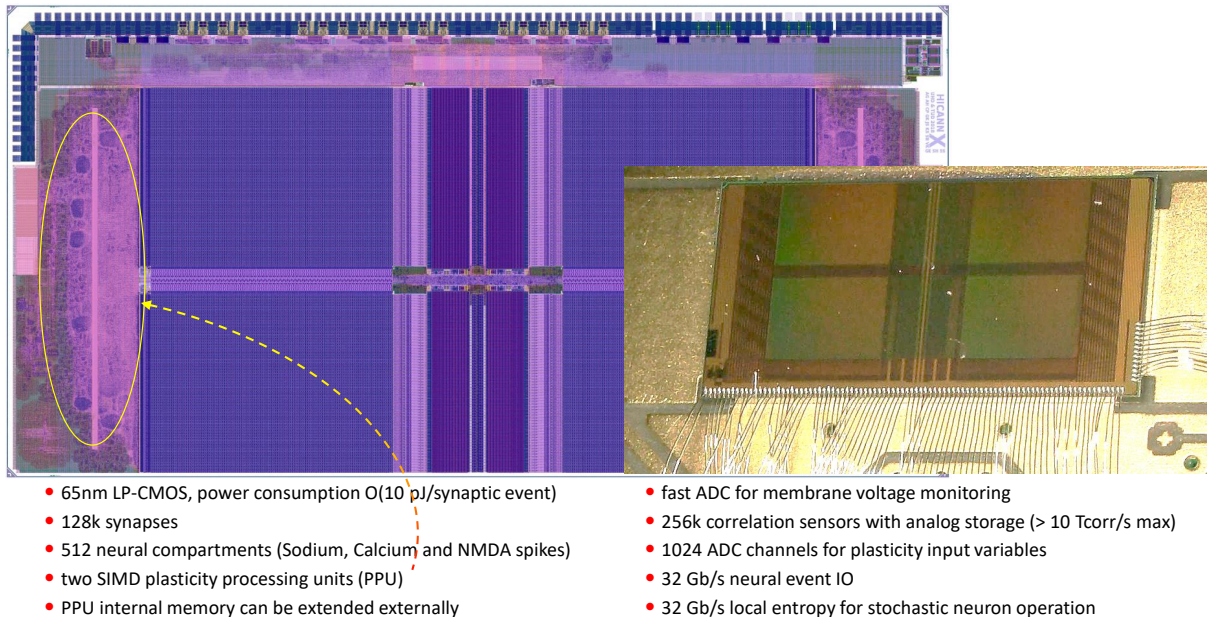


Fig. 7: Top: Layout drawing and chip photograph of the HICANN-X ASIC. Bottom: key features of HICANN-X.

	syndriver top				syndriver bottom				L1→L2			
	0	1	2	3	0	1	2	3	0	1	2	3
4 neuron output channels left half of anncore	0	X				X				X		
	1		X				X				X	
	2			X				X				X
	3				X				X			X
4 neuron output channels right half of anncore	0	X				X				X		
	1		X				X				X	
	2			X				X				X
	3				X				X			X
4 L2 → L1	0	X	X	X	X	X	X	X	X	X		
	1	X	X	X	X	X	X	X	X		X	
	2	X	X	X	X	X	X	X	X			X
	3	X	X	X	X	X	X	X	X			X
8 background generators	0	X								X		
	1		X								X	
	2			X								X
	3				X							X
	4					X				X		
	5						X				X	
	6							X				X
	7								X			X

Fig. 8: Conceptual view of the internal digital event routing matrix of the HICANN-X chip. All 20 sources are shown vertically on the left, while the 12 output channels are listed at the top. At each position marked with a cross a programmable routing element is located.

ets<sup>1</sup>. Fig. 8 shows the implementation of the central Layer1 digital event routing network. There are two main sources and sinks for event data: the analog network core, which has eight input event and eight output event buses, as well as the Event Link Layer 2 (Layer2)→Layer1 converter, which provides four links in each direction. With the exception of the analog core input buses, each link can handle one event per clock cycle of 4 ns. The ANNCORE input buses are limited to one event every two cycles.

All eight High Input Count Analog Neural Network (HICANN) compatible links are used for Layer2 based event transport. An event is encoded as a combination of neuron address and time stamp. The conversion between time-stamped Layer2 data and real-time Layer1 data is performed inside the Layer2→Layer1 converter located in the digital core logic. It uses a globally synchronized system time counter for this purpose. The routing of all Layer1 events is done within the router matrix. Inside this module are several columns of buffered n-to-1 event merger stages allowing to combine the data of a set of inputs into one Layer1 output channel. All eight physical links of the chip can be simultaneously used for neuron event data (Layer2), slow control and PPU global memory accesses. The number of active links might be statically programmed to be any number between one and the maximum of eight. This is useful if several chips should be connected to a single host with a limited number of available links. All events transferred via

<sup>1</sup>The Layer1 data format codes a neural event as a parallel bit-field containing the neuron address and a valid bit. It is real-time data with a temporal resolution of the system clock, which is 250 MHz in HICANN-X.

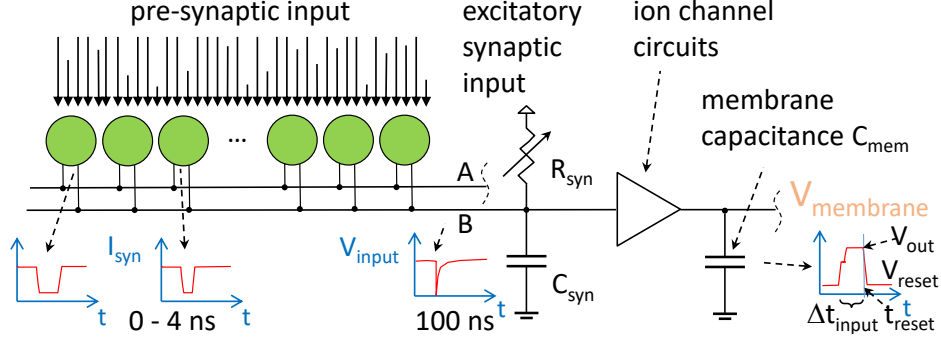


Fig. 9. Operating principle of the HAGEN extensions in the HICANN-X chip.

Layer2 are protected against undetected bit-errors by Cyclic Redundancy Check (CRC) fields.

### 3.2. Analog Inference: Rate-based Extension of HICANN-X

One of the first neuromorphic systems build in Heidelberg was Heidelberg Analog Evolvable Neural network (HAGEN), a fast analog Perceptron-based network chip optimized for hardware-in-the-loop training[30]. Caused by parallel activities withing the Heidelberg Electronic Vision(s) research group[31] it was mainly trained by evolutionary algorithms[32], explaining the acronym. Nevertheless, it was perfectly usable for other hardware-in-the-loop based algorithms, similar to the deep-learning results that have been more recently achieved by other neural network chips used in a Perceptron-like fashion[33]. Although the HICANN architecture has been successfully used to implement deep multi-layer networks using rate-based spiking models[34] and back-propagation based training, it loses some of its power-efficiency by emulating a Perceptron model. Encoding the activation in the time between spikes can enhance the efficiency significantly[35]. In all spiking solutions the network operates in continuous time and therefore the size of the network is limited to the number of neurons and synapses available on the chip. The HAGEN extension, which is part of the HICANN-X chip, allows a seamless mixture of spiking and non-spiking operation within a single chip. Since this rate-based operation is based on discrete-time analog vector-matrix multiplication, a time-multiplexing scheme can be employed, similar to digital accelerators for deep convolutional networks[36]. In this case the size of the network is limited only by the size of any external memory.

Fig. 9 visualizes the differences between standard spiking mode and HAGEN

mode, which eliminates all temporal dynamics from the neuron. By disabling the leakage term of the neuron the membrane just sums up the synaptic input. The excitatory input is added with a positive and the inhibitory input with a negative sign. All input is applied during the time interval  $\Delta t_{\text{input}}$ , after which the membrane voltage is digitized by the Correlation-readout ADC (CADC) and the neuron is set to the reset voltage  $V_{\text{reset}}$  by a reset signal from the PPU.  $\Delta t_{\text{input}}$  can be as short as 100 ns. It depends on the bandwidth of the synaptic input and the number of synaptic rows used, i.e. the total time required to transfer all input events to the synapses. Since the minimum time is at least a few synaptic time constants and nothing is gained by setting the integration time shorter than the conversion time of the CADC, a typical value for  $\Delta t_{\text{input}}$  is about 500 ns. Thereby, the network can evaluate  $2 \cdot 10^6 \times 256 \times 512 = 2.62 \cdot 10^{11}$  multiply-accumulate operations per second. By shortening the conversion time of the CADC further speed improvements are possible. Since the reset voltage of the neuron membrane can be aligned with the lower bound of the CADC conversion range the neuron acts like a ReLU unit in this setting[37].

A standard synapse within BrainScaleS reacts to a pre-synaptic event in a digital fashion: the arrival of a pre-synaptic event generates a fixed current pulse. By enabling short-term facilitation or depression[23] the synaptic strength depends on the pre-synaptic firing history. This is achieved by modulating the pulse length generated by the synapse. Instead of using the firing history, in HAGEN mode the pulse length is transmitted together with the pre-synaptic spike and converted into variable length pulses by the existing Short Time Plasticity (STP) pulse-length modulation circuits. The digital pulse length information is transmitted by reusing the 5 lower address bits of the Layer1 event data, since in the HAGEN mode the network structure is much more regular and not all pre-synaptic address bits are needed.

Fig.10 shows some early results using the activity-based Perceptron mode from HICANN-X for analog vector-matrix multiplication. In the left part of the figure, 127 neurons are measured simultaneously. Their synaptic weights increase linearly from -63 to 63, i.e. all synapses connected to a single neuron are set to the same weight while the weights increase from neuron to neuron. All synapses receive the same input: 0, 3 or 7 for the black, red and blue traces respectively. The outputs of all 127 neurons are digitized simultaneously by the CADC and the digital values are plotted over the weight values of the neurons. Although the neuron circuits are calibrated, some fixed-pattern noise remains visible. The temporal variations are caused by a well-understood circuit flaw, that will be removed in future iterations.

The chip has been subsequently used to perform inference on the MNIST dataset[38].

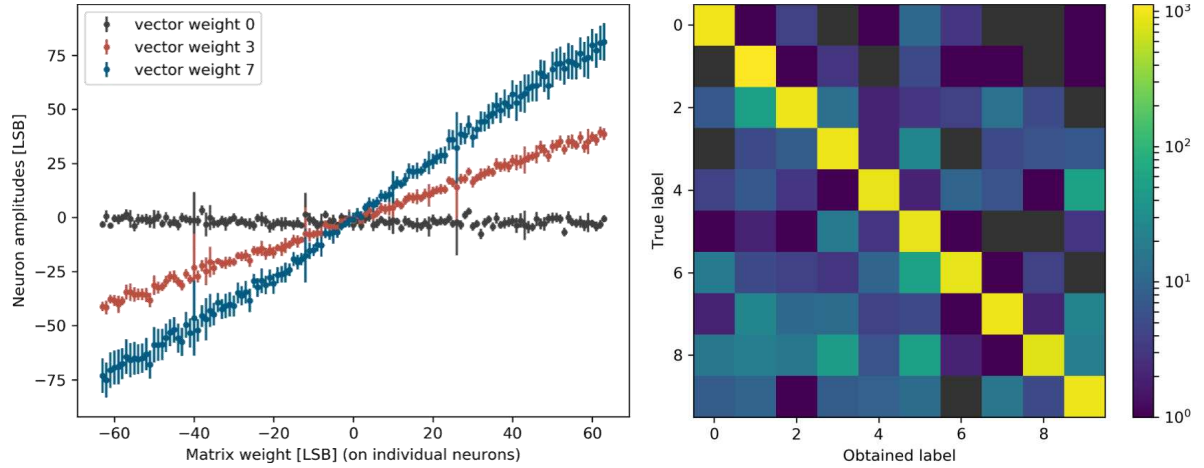


Fig. 10: Left: Results for analog vector-matrix multiplication. Right: Confusion matrix for MNIST.

A three-layer network has been trained in Tensorflow[39] to reach a classification rate of 97.43%. The weights and input activations of this network have been quantized to 6 bit weight and 5 bit input resolution, to fit the trained network to the dynamic range of the analog circuits. The inference on the test data set has been repeated using the HICANN-X chip. The resulting classification accuracy was 92.48%. The corresponding confusion matrix is shown in the right panel of Fig. 10. The deterioration is most likely caused by the remaining fixed-pattern noise. In the future we will include the hardware in the forward-path of the training loop, similar to the approach followed in [34], which will most likely improve the accuracy significantly.

#### 4. Analog Verification of Complex Neuron Circuits

The BrainScaleS systems feature complex mixed-signal circuits to emulate the rich properties of their biological counterparts. Our neuron circuits, implementing the AdEx equations [25], possess a multitude of individual subcomponents, such as a leak or adaptation term. Each of these units is parameterized through a number of digital controls as well as analog voltage and current biases. Designed to support a variety of different tasks, ranging from biologically realistic firing patterns to analog matrix multiplication, these circuits have to be operated at widely different operating points. The correct behavior has to be ensured prior to fabrication. Individual components can often be unit-tested in isolation, making use of conventional simulation strategies. The assessability of a complete design is, however, limited due to error propagation and inter-dependencies of parameters.

A suite of benchmark tasks, evaluated on comprehensive testbenches, is required



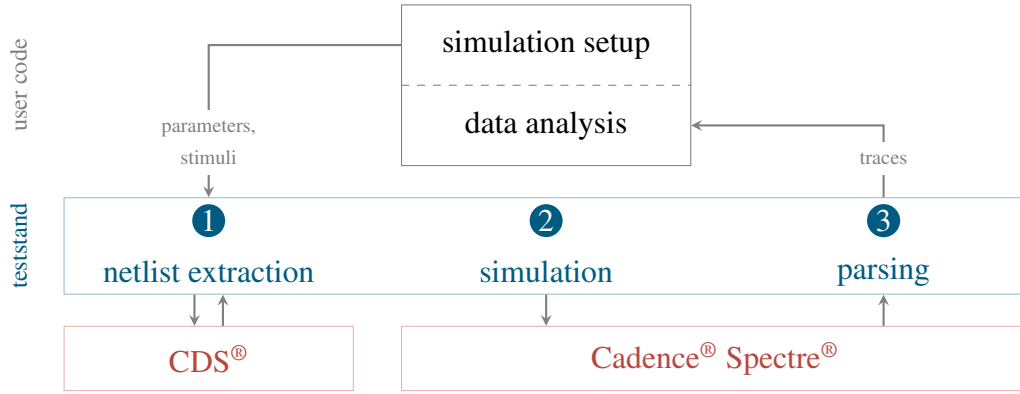


Fig. 11: Structure of a teststand-based simulation highlighting the interaction with the Cadence Design Suite. Image taken from [40].

for pre-tapeout verification. To ensure the required degree of precision over larger arrays of analog circuits, mismatch effects introduced through imperfections in the production process, have to be covered through Monte Carlo (MC) simulations. Different incarnations of a circuit can be obtained by individually fixing the MC seed. These virtual instances can then be characterized, very similar to the fabricated siblings. Similarly, the worst case behavior can be characterized for the process corners.

In the following paragraphs, we present our simulation strategy and a custom library to aid software-driven simulations within the rich ecosystem of the Python programming language. We will guide through our benchmarking flow for our current generation of AdEx neurons. Similar approaches have successfully been taken for the verification of plasticity circuits and vector-matrix multiplication circuits.

#### 4.1. Interfacing analog simulations from Python

Our custom Python module *teststand* provides a tight integration between analog circuit simulations and the ecosystem of the programming language[40]. It mainly consists of a software layer to interface with the Cadence Spectre simulator and other tools from the Cadence Design Suite.

Teststand extracts the testbench’s netlist directly from the target cell view as available in the design library. The data is accessed by querying the database via an OCEAN script executed as a child process. Teststand then reads the netlist and modifies it according to the user’s specification. In addition to the schematic description, Spectre netlists also contain simulator instructions. Teststand generates



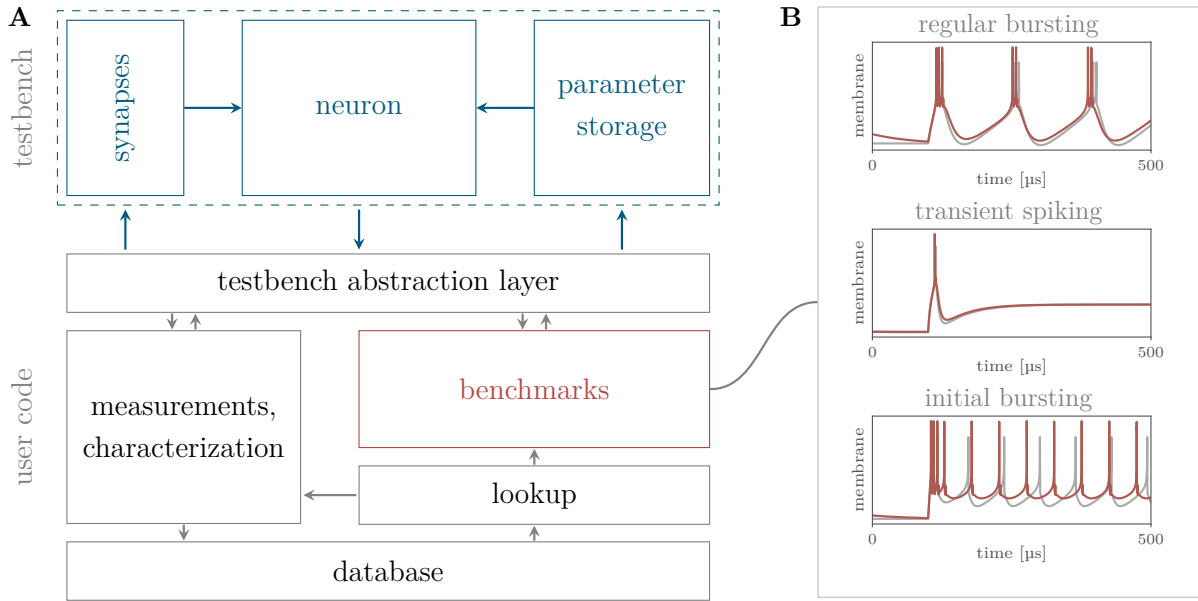


Fig. 12: MC calibration workflow of an AdEx neuron circuit using teststand. **A** Testbench and overview on the software stack for characterization and calibration. **B** Membrane traces (red) of a neuron circuit simulation configured for regular bursting, transient spiking, and initial bursting. The results of a numerical integration of the AdEx equations is shown as a reference (gray).

these statements according to the user's Python code. Specifically, the user can define analyses to be performed by the simulator, such as DC, AC, and transient simulations. MC analyses are supported as well and play an important role in the verification strategies presented below. Teststand can be easily extended to support all features provided by the backend.

All circuit parameters, stimuli, and nodes to be recorded are specified using an object-oriented interface that resembles Spectre simulation instructions.

```
cell = ('mylib', 'mycell', 'schematic')
nets = ['I0.mynet']
```

```
teststand = Teststand(cds_lib, cell)
tran = TransientAnalysis('tran', 1e-3)
simulation = Simulation(
    [tran], params, save=nets)
result = teststand.simulate(simulation)
```

The `simulate()`-call executes Spectre as a child process. Basic parallelization features are natively provided via the *multiprocessing* library. Scheduling can

be trivially extended to support custom compute environments. The simulation log is parsed and potential error messages are presented to the user as Python exceptions.

Results are read and provided to the user as structured NumPy arrays. This allows to resort to the vast amount of data processing libraries available in the Python ecosystem to process and evaluate recorded data. Most notably, this includes NumPy [41], SciPy [42], and Matplotlib [43]. As a side effect, the latter allows to directly generate rich publication-ready figures from analog circuit simulations.

## 4.2. Monte Carlo calibration of AdEx neuron circuits

As shown in Fig. 12A, we used the teststand library, *inter alia*, for the verification of our AdEx design. The model equations feature a high-dimensional parameter space, allowing for a wide range of behaviors. Our circuit, on the other hand, is parameterized through 24 individual analog bias sources and a set of digital controls. Starting from first-order models of the utilized subcomponents, we characterized the circuit's dynamics through a set of measurements on the full neuron circuit. With the results stored in a database, we established a transformation between the circuits's and the models's parameter spaces. The influence of mismatch effects manifests itself in deviations in these calibration curves for individual neuron instances. We applied the above framework for a large number of neuron incarnations, obtained by fixing the respective MC seeds.

The circuit was benchmarked against multiple firing patterns, such as *transient spiking*, *regular bursting*, and *initial bursting* [44]. For each of these targets, a set of biases, corresponding to the respective parameter set from literature, was determined through a reverse lookup based on the above transformations. Exemplary results for a single neuron simulation are shown in Fig. 12 B.

The presented approach enforces the development of calibration algorithms before tape-out. Especially for circuits with large parameter spaces, there might occur multi-dimensional dependencies which can be hard to resolve. The strategy might also reveal an insufficient parametrization not necessarily apparent from individual unit tests. In order to uncover potential regressions due to modifications to a circuit, simulations based on teststand can easily be automated and allow continuous integration testing for full-custom designs.

## 5. Conclusion

The development and implementation of the presented second generation Brain-ScaleS architecture will hopefully continue during the next years. The outcome we hope for is a multi-wafer system, constructed from hundreds of 30 cm silicon wafers, each one directly embedded in a printed circuit board (PCB) and all of them interconnected to form a novel large-scale analog neuromorphic platform. A system capable of answering questions about learning and development in large scale, biologically realistic neural networks.

Utilizing standard Complementary Metal-Oxid-Semiconductor (CMOS) technology to build large-scale analog accelerated neuromorphic hardware systems places our approach in the middle between the two major research directions for AI circuits: digital accelerators and novel persistent memory devices. It presents a complementary option to these technologies. Compared to systems based on novel device technology it has advantages, like the high operational speed, low energy requirements for learning, the possibility to use any standard CMOS process without regards to back-of-the-line compatibility and the capability to replicate relevant biological structures more easily. In comparison to digital implementations, like Loihi or SpiNNaker [45, 46], the fully analog implementation of complex neural structures combined with true in-memory computing allows for time-continuous emulation of neural dynamics and much higher emulation speed at similar energy efficiencies. Most importantly, analog CMOS implementations might be the essential step to uncover the learning rules needed to cope with substrate variations. In our systems the local learning rules do not only train the system to perform a certain task, but simultaneously adjust the operating point of the circuits and compensate fixed pattern noise[23]. This will be an essential property for future novel computing systems based on advanced device technologies as well, since they all are expected to have substantially increased device-to-device variations. We hope that our BSS platform will help to gain insight in the necessary algorithms in the upcoming future.

In the short term the BSS system allows the combination of energy- and cost-efficient analog inference with local learning rules for a multitude of practical applications, scaling from small systems for edge computing up to high-performance neuromorphic cloud computing.

## 6. Acknowledgments

The authors wish to express their gratitude to Andreas Grübl, Yannik Stradmann, Vitali Karasenko, Korbinian Schreiber, Christian Pehle, Ralf Achenbach, Markus Dorn and Aron Leibfried for their invaluable help and active contributions in the development of the BrainScaleS 2 ASICs and systems.

They are not forgetting the important role their former colleagues Andreas Hartel, Syed Aamir, Gerd Kiene, Matthias Hock, Simon Friedmann, Paul Müller, Laura Kriener and Timo Wunderlich had in these endeavors.

They also want to thank their collaborators Sebastian Höppner from TU Dresden and Tugba Demirci from EPFL Lausanne for their contributions to the BrainScaleS 2 prototype ASIC.

Very special thanks go to Eric Müller and his team for leading the software development as well as Mihai Petrovici, Sebastian Schmitt and the late Karlheinz Meier for their invaluable advice.

This work has received funding from the European Union Seventh Framework Programme ([FP7/2007-2013]) under grant agreement no 604102 (HBP rampup), 269921 (BrainScaleS), 243914 (Brain-i-Nets), the Horizon 2020 Framework Programme ([H2020/2014-2020]) under grant agreement 720270 and 785907 (HBP SGA1 and SGA2) as well as from the Manfred Stärk Foundation.

## 7. Author Contribution

J.S. created the concept, has been the lead architect of the BSS systems and wrote the manuscript except for Section 4, which was written by S.B. S.B. also created the teststand software and conceived the simulations jointly with P.D, who performed the simulations and prepared the results. J.W. performed the measurements for the HAGEN mode and created figure 10. All authors edited the manuscript together.

## References

- [1] J. Schemmel, D. Brüderle, A. Grübl, M. Hock, K. Meier, and S. Millner, “A wafer-scale neuromorphic hardware system for large-scale neural modeling,” in *Proceedings of the 2010 IEEE International Symposium on Circuits and Systems (ISCAS)*, 2010, pp. 1947–1950.

- [2] G. Indiveri, B. Linares-Barranco, T. J. Hamilton, A. van Schaik, R. Etienne-Cummings, T. Delbruck, S.-C. Liu, P. Dudek, P. Häfliger, S. Renaud, J. Schemmel, G. Cauwenberghs, J. Arthur, K. Hynna, F. Folowosele, S. Saighi, T. Serrano-Gotarredona, J. Wijekoon, Y. Wang, and K. Boahen, “Neuromorphic silicon neuron circuits,” *Frontiers in Neuroscience*, vol. 5, no. 0, 2011. [Online]. Available: [http://www.frontiersin.org/Journal/Abstract.aspx?s=755&name=neuromorphicengineering&ART\\_Doi=10.3389/fnins.2011.00073](http://www.frontiersin.org/Journal/Abstract.aspx?s=755&name=neuromorphicengineering&ART_Doi=10.3389/fnins.2011.00073)
- [3] B. V. Benjamin, P. Gao, E. McQuinn, S. Choudhary, A. R. Chandrasekaran, J.-M. Bussat, R. Alvarez-Icaza, J. V. Arthur, P. A. Merolla, and K. Boahen, “Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations,” *Proceedings of the IEEE*, vol. 102, no. 5, pp. 699–716, 2014.
- [4] R. Douglas, M. Mahowald, and C. Mead, “Neuromorphic analogue VLSI,” *Annu. Rev. Neurosci.*, vol. 18, pp. 255–281, 1995.
- [5] J. Schemmel, A. Grübl, K. Meier, and E. Muller, “Implementing synaptic plasticity in a VLSI spiking neural network model,” in *Proceedings of the 2006 International Joint Conference on Neural Networks (IJCNN)*. IEEE Press, 2006.
- [6] J. Schemmel, D. Brüderle, K. Meier, and B. Ostendorf, “Modeling synaptic plasticity within networks of highly accelerated I&F neurons,” in *Proceedings of the 2007 IEEE International Symposium on Circuits and Systems (ISCAS)*. IEEE Press, 2007, pp. 3367–3370.
- [7] K. Zoschke, M. Güttler, L. Böttcher, A. Grübl, D. Husmann, J. Schemmel, K. Meier, and O. Ehrmann, “Full wafer redistribution and wafer embedding as key technologies for a multi-scale neuromorphic hardware cluster,” in *2017 IEEE 19th Electronics Packaging Technology Conference (EPTC)*. IEEE, 2017, pp. 1–8.
- [8] S. Millner, A. Grübl, K. Meier, J. Schemmel, and M.-O. Schwartz, “A VLSI implementation of the adaptive exponential integrate-and-fire neuron model,” in *Advances in Neural Information Processing Systems 23*, J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, R. Zemel, and A. Culotta, Eds., 2010, pp. 1642–1650.
- [9] T. Pfeil, A. Grübl, S. Jeltsch, E. Müller, P. Müller, M. A. Petrovici, M. Schmücker, D. Brüderle, J. Schemmel, and K. Meier, “Six networks on a universal neuromorphic computing substrate,” *Frontiers in Neuroscience*, vol. 7, p. 11, 2013. [Online]. Available: [http://www.frontiersin.org/neuromorphic\\_engineering/10.3389/fnins.2013.00011/abstract](http://www.frontiersin.org/neuromorphic_engineering/10.3389/fnins.2013.00011/abstract)

- [10] A. P. Davison, D. Brüderle, J. Eppler, J. Kremkow, E. Muller, D. Pecevski, L. Perrinet, and P. Yger, “PyNN: a common interface for neuronal network simulators,” *Front. Neuroinform.*, vol. 2, no. 11, 2008.
- [11] D. Brüderle, M. A. Petrovici, B. Vogginger, M. Ehrlich, T. Pfeil, S. Millner, A. Grübl, K. Wendt, E. Müller, M.-O. Schwartz, D. de Oliveira, S. Jeltsch, J. Fieres, M. Schilling, P. Müller, O. Breitwieser, V. Petkov, L. Muller, A. Davison, P. Krishnamurthy, J. Kremkow, M. Lundqvist, E. Muller, J. Partzsch, S. Scholze, L. Zühl, C. Mayr, A. Destexhe, M. Diesmann, T. Potjans, A. Lansner, R. Schüffny, J. Schemmel, and K. Meier, “A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems,” *Biological Cybernetics*, vol. 104, pp. 263–296, 2011. [Online]. Available: <http://dx.doi.org/10.1007/s00422-011-0435-9>
- [12] J. Schemmel, L. Kriener, P. Müller, and K. Meier, “An accelerated analog neuromorphic hardware system emulating NMDA-and calcium-based non-linear dendrites,” *arXiv preprint arXiv:1703.07286*, 2017.
- [13] C. S. Thakur, J. L. Molin, G. Cauwenberghs, G. Indiveri, K. Kumar, N. Qiao, J. Schemmel, R. Wang, E. Chicca, J. Olson Hasler *et al.*, “Large-scale neuromorphic spiking array processors: A quest to mimic the brain,” *Frontiers in neuroscience*, vol. 12, p. 891, 2018.
- [14] S. Friedmann, J. Schemmel, A. Grübl, A. Hartel, M. Hock, and K. Meier, “Demonstrating hybrid learning in a flexible neuromorphic hardware system,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 11, no. 1, pp. 128–142, 2017.
- [15] S. A. Aamir, P. Müller, A. Hartel, J. Schemmel, and K. Meier, “A highly tunable 65-nm CMOS LIF neuron for a large-scale neuromorphic system,” in *Proceedings of IEEE European Solid-State Circuits Conference (ESSCIRC)*, 2016.
- [16] S. A. Aamir, Y. Stradmann, P. Müller, C. Pehle, A. Hartel, A. Grübl, J. Schemmel, and K. Meier, “An accelerated lif neuronal network array for a large-scale mixed-signal neuromorphic architecture,” *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 65, no. 12, pp. 4299–4312, 2018.
- [17] S. Friedmann, J. Schemmel, A. Grübl, A. Hartel, M. Hock, and K. Meier, “Demonstrating hybrid learning in a flexible neuromorphic hardware system,” *IEEE Transactions on Biomedical Circuits and Systems*, vol. 11, no. 1, pp. 128–142, 2017.

- [18] M. Hock, A. Hartel, J. Schemmel, and K. Meier, “An analog dynamic memory array for neuromorphic hardware,” in *Circuit Theory and Design (ECTD), 2013 European Conference on*, Sep. 2013, pp. 1–4.
- [19] M. Tsodyks and H. Markram, “The neural code between neocortical pyramidal neurons depends on neurotransmitter release probability,” *Proceedings of the national academy of science USA*, vol. 94, pp. 719–723, Jan. 1997.
- [20] T. Pfeil, J. Jordan, T. Tetzlaff, A. Grübl, J. Schemmel, M. Diesmann, and K. Meier, “The effect of heterogeneity on decorrelation mechanisms in spiking neural networks: a neuromorphic-hardware study,” *arXiv preprint arXiv:1411.7916*, 2014.
- [21] J. Jordan, M. A. Petrovici, O. Breitwieser, J. Schemmel, K. Meier, M. Diesmann, and T. Tetzlaff, “Deterministic networks for probabilistic computing,” *Scientific reports*, vol. 9, no. 1, pp. 1–17, 2019.
- [22] G. Kiene, “Mixed-signal neuron and readout circuits for a neuromorphic system,” Masterthesis, Universität Heidelberg, 2017.
- [23] S. Billaudelle, “Design and implementation of a short term plasticity circuit for a 65 nm neuromorphic hardware system,” Masterarbeit, Universität Heidelberg, 2017.
- [24] S. Billaudelle, B. Cramer, M. A. Petrovici, K. Schreiber, D. Kappel, J. Schemmel, and K. Meier, “Structural plasticity on an accelerated analog neuromorphic hardware system,” *arXiv preprint arXiv:1912.12047*, 2019.
- [25] R. Brette and W. Gerstner, “Adaptive exponential integrate-and-fire model as an effective description of neuronal activity,” *J. Neurophysiol.*, vol. 94, pp. 3637 – 3642, 2005.
- [26] S. Millner, “Development of a multi-compartment neuron model emulation,” 2012.
- [27] R. Jolivet, T. J. Lewis, and W. Gerstner, “Generalized integrate-and-fire models of neuronal activity approximate spike trains of a detailed model to a high degree of accuracy,” *Journal of neurophysiology*, vol. 92, no. 2, pp. 959–976, 2004.
- [28] V. Thanasoulis, J. Partzsch, S. Hartmann, C. Mayr, and R. Schüffny, “Dedicated fpga communication architecture and design for a large-scale neuromorphic system,” in *2012 19th IEEE International Conference on Electronics, Circuits, and Systems (ICECS 2012)*. IEEE, 2012, pp. 877–880.
- [29] J. Schemmel, J. Fieres, and K. Meier, “Wafer-scale integration of analog neural networks,” in *Proceedings of the 2008 International Joint Conference on Neural Networks (IJCNN)*, 2008.



- [30] J. Schemmel, S. Hohmann, K. Meier, and F. Schürmann, “A mixed-mode analog neural network using current-steering synapses,” *Analog Integrated Circuits and Signal Processing*, vol. 38, no. 2-3, pp. 233–244, 2004.
- [31] J. Langeheine, M. Trefzer, D. Brüderle, K. Meier, and J. Schemmel, “On the evolution of analog electronic circuits using building blocks on a CMOS FPTA,” in *Proceedings of the Genetic and Evolutionary Computation Conference (GECCO2004)*, 2004.
- [32] S. Hohmann, J. Fieres, K. Meier, J. Schemmel, T. Schmitz, and F. Schürmann, “Training fast mixed-signal neural networks for data classification,” in *Proceedings of the 2004 International Joint Conference on Neural Networks (IJCNN’04)*. IEEE Press, 2004, pp. 2647–2652.
- [33] E. Nurse, B. S. Mashford, A. J. Yepes, I. Kiral-Kornek, S. Harrer, and D. R. Freestone, “Decoding eeg and lfp signals using deep learning: heading truenorth,” in *Proceedings of the ACM International Conference on Computing Frontiers*, 2016, pp. 259–266.
- [34] S. Schmitt, J. Klähn, G. Bellec, A. Grübl, M. Güttler, A. Hartel, S. Hartmann, D. Husmann, K. Husmann, S. Jeltsch, V. Karasenko, M. Kleider, C. Koke, A. Kononov, C. Mauch, E. Müller, P. Müller, J. Partzsch, M. A. Petrovici, B. Vogginger, S. Schiefer, S. Scholze, V. Thanasoulis, J. Schemmel, R. Legenstein, W. Maass, C. Mayr, and K. Meier, “Classification with deep neural networks on an accelerated analog neuromorphic system,” *arXiv*, 2016.
- [35] J. Göltz, A. Baumbach, S. Billaudelle, O. Breitwieser, D. Dold, L. Kriener, A. F. Kungl, W. Senn, J. Schemmel, K. Meier *et al.*, “Fast and deep neuromorphic learning with time-to-first-spike coding,” *arXiv preprint arXiv:1912.11443*, 2019.
- [36] A. Shawahna, S. M. Sait, and A. El-Maleh, “Fpga-based accelerators of deep learning networks for learning and classification: A review,” *IEEE Access*, vol. 7, pp. 7823–7859, 2018.
- [37] P. Sharma and A. Singh, “Era of deep neural networks: A review,” in *2017 8th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*. IEEE, 2017, pp. 1–5.
- [38] Y. LeCun and C. Cortes, “The mnist database of handwritten digits,” 1998.
- [39] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar,

- P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng, “TensorFlow: Large-scale machine learning on heterogeneous distributed systems,” 2015. [Online]. Available: <http://download.tensorflow.org/paper/whitepaper2015.pdf>
- [40] A. Grübl, S. Billaudelle, B. Cramer, V. Karasenko, and J. Schemmel, “Verification and design methods for the brainscales neuromorphic hardware system,” *arXiv preprint*, 2020. [Online]. Available: <http://arxiv.org/abs/2003.11455>
- [41] T. E. Oliphant, *A guide to NumPy*. Trelgol Publishing USA, 2006, vol. 1.
- [42] E. Jones, T. Oliphant, and P. Peterson, “SciPy: Open source scientific tools for Python,” 2001. [Online]. Available: <http://www.scipy.org/>
- [43] J. D. Hunter, “Matplotlib: A 2d graphics environment,” *Computing in Science Engineering*, vol. 9, no. 3, pp. 90–95, May 2007.
- [44] R. Naud, N. Marcille, C. Clopath, and W. Gerstner, “Firing patterns in the adaptive exponential integrate-and-fire model,” *Biological Cybernetics*, vol. 99, no. 4, pp. 335–347, Nov 2008. [Online]. Available: <http://dx.doi.org/10.1007/s00422-008-0264-7>
- [45] M. Davies, N. Srinivasa, T.-H. Lin, G. Chinya, Y. Cao, S. H. Choday, G. Dimou, P. Joshi, N. Imam, S. Jain *et al.*, “Loihi: A neuromorphic manycore processor with on-chip learning,” *IEEE Micro*, vol. 38, no. 1, pp. 82–99, 2018.
- [46] S. B. Furber, F. Galluppi, S. Temple, and L. A. Plana, “The spinnaker project,” *Proceedings of the IEEE*, vol. 102, no. 5, pp. 652–665, 2014.