

The Operating System of the Neuromorphic BrainScaleS-1 System

Eric Müller^{*||}, Sebastian Schmitt^{*||}, Christian Mauch^{*||},
Sebastian Billaudelle^{||}, Andreas Grübl^{||}, Maurice Güttler^{||}, Dan Husmann^{||}, Joscha Ilmberger^{||}, Sebastian Jeltsch^{||},
Jakob Kaiser^{||}, Johann Klähn^{||}, Mitja Kleider^{||}, Christoph Koke^{||}, José Montes^{||}, Paul Müller^{||},
Johannes Partzsch[§], Felix Passenberg^{||}, Hartmut Schmidt^{||}, Bernhard Vogginger[§], Jonas Weidner^{||},
Christian Mayr[§] and Johannes Schemmel^{||}

^{*}contributed equally

^{||} Kirchhoff-Institute for Physics

Ruprecht-Karls-Universität Heidelberg, Germany

Email: {mueller,sschmitt,cmauch}@kip.uni-heidelberg.de

[§] Chair of Highly-Parallel VLSI-Systems and Neuro-Microelectronics
Technische Universität Dresden, Germany

Abstract—BrainScaleS-1 is a wafer-scale mixed-signal accelerated neuromorphic system targeted for research in the fields of computational neuroscience and beyond-von-Neumann computing. The BrainScaleS Operating System (BrainScaleS OS) is a software stack giving users the possibility to emulate networks described in the high-level network description language PyNN with minimal knowledge of the system. At the same time, expert usage is facilitated by allowing to hook into the system at any depth of the stack. We present operation and development methodologies implemented for the BrainScaleS-1 neuromorphic architecture and walk through the individual components of BrainScaleS OS constituting the software stack for BrainScaleS-1 platform operation.

I. INTRODUCTION

State-of-the-art neuromorphic architectures pose many requirements in terms of system control, data preprocessing, data exchange and data analysis. In all these areas, software is involved in satisfying these requirements. Several neuromorphic systems are directly used by individual researchers in collaborations, e.g., [1–4]. In addition, some systems are operated as experiment platforms providing access for external users [1–3, 5].

Especially systems open for a broader range of users require clear and concise interfaces. Neuromorphic platform operators have additional requirements in resource management, runtime control and –depending on data volumes– “grid-computing”-like data processing capabilities. At the same time, usability and experiment reproducibility are crucial properties of all experiment platforms, including neuromorphic systems.

Modern software engineering techniques such as code review, continuous integration as well as continuous deployment can help to increase platform robustness and ensure experiment reproducibility. Long-term hardware development roadmaps and experiment collaborations draw

attention to platform sustainability. Technical decisions need to be evaluated for potential future impact; containing and reducing technical debt is a key objective during planning as well as development. Regardless of being software-driven simulations/emulations, or being physical experiments, modern experiment setups more and more depend on these additional tools and skills in order to enable reproducible, correct and successful scientific research.

This paper describes the results of a ten-year project delivering the software environment and platform operation tools for the BrainScaleS-1 neuromorphic system. The following sections describe the hardware substrate and give a general overview. Section II introduces the methods and software tools we employ. In section III, the scopes and implementation details of the main software layers and libraries are explained, followed by an overview over the operation of the platform in section IV. Section V exemplifies the usage of the BrainScaleS Operating System on a simple experiment and describes larger experiments carried out in the past. We close in section VI with an overview over future developments and discuss our endeavor and the lessons learned in section VII.

A. The BrainScaleS-1 Neuromorphic System

Classical neuromorphic systems make use of VLSI to implement electronic analog circuits mimicking neuro-biological architectures in the nervous system [6]. Contemporary systems also employ mixed-signal techniques to enable flexible system connectivity based on conventional digital interfaces [7]. Recently, purely digital systems emerged [1, 3, 4]. Compared to the analog approach typical advantages of such systems are: deterministic behavior and arbitrarily programmable neuron dynamics. However, when comparing




Fig. 1: (a) 3D-schematic of a BrainScaleS Wafer Module (dimensions: 50 cm \times 50 cm \times 15 cm) hosting the wafer (A) and 48 FPGAs (B). The positioning mask (C) is used to align elastomeric connectors that link the wafer to the large main PCB (D). Support PCBs provide power supply (E & F) for the on-wafer circuits as well as access (G) to analog dynamic variables such as neuron membrane voltages. The connectors for inter-wafer and off-wafer/host connectivity (48 \times Gigabit-Ethernet) are distributed over all four edges (H) of the main PCB. Mechanical stability is provided by an aluminum frame (I). (b) Photograph of a fully assembled wafer module.

at the same technology node the energy as well as area consumption is increased.

Based on the ideas of a single-chip implementation called *Spikey* [8], BSS-1 is a mixed-signal architecture providing accelerated Adaptive Exponential Integrate-and-Fire (AdEx) neuron dynamics and plastic synapses [5, 9, 10]. While many neuromorphic systems target biological real-time execution (i.e. model time constants in the same order as their biological counterpart) [11], BSS-1 evolves in continuous time, typically at a speed-up factor of 1000–10 000 faster than biological real time. Consequently, real-time interfacing to, e.g., sensors and robotic applications are not the main goal of the architecture. The design focuses on fast model dynamics, controllable model parameters and system scalability, thereby allowing for time-compressed emulations of longer experiment time scales. Plasticity and learning processes can therefore be investigated in manageable time frames.

Figure 1 depicts a BSS-1 wafer module. The main constituent is a silicon wafer, manufactured in 180 nm Complementary Metal Oxide Semiconductor (CMOS) technology, carrying 384 HICANN chips that are interconnected via an on-wafer bus network. Each chip hosts up to 512 AdEx neurons and 113k plastic synapses. 48 Xilinx Kintex-7 Field Programmable Gate Arrays (FPGA) provide an I/O interface for configuration, stimulus and recorded data. The connection between these FPGAs and the control cluster network is established via 1-Gigabit and 10-Gigabit Ethernet, cf. Schmitt et al. [12].

B. Performing Experiments

Providing access to and experimenting with neuromorphic systems is an active field of research [13–16]. In 2017 Schuman et al. [17] highlighted that:

Supporting software will be a vital component in order for neuromorphic systems to be truly successful and accepted both within and outside the computing community.

We second this statement and this paper details our approach on how to tackle the software challenge in the context of BSS-1 in particular. However, many of the developed solutions are of general scope and should be applicable to other systems that have similar features and targets.

Different approaches to user interfaces are viable, e.g., the interface for *Spikey* —a previous chip developed by the Heidelberg group— mainly used *Python* for configuration and experiment description [18]. In contrast, BrainScaleS OS provides only thin *Python* wrappers for all user-facing Application Programming Interfaces (API) while the core software is written in *C++*, cf. section II-B2. *Spikey* focused on *PyNN* as the experiment description language [19], see section II-B1. As of today, common spiking neural network simulators and a few hardware emulators support *PyNN* [13, 19]. Further work investigated the typical workflow when porting experiments from pure software simulations to neuromorphic platforms [20].

BSS-1 has a fairly large parameter space, $O(50 \text{ MiB per wafer})$ of static configuration data, and the analog characteristics of the system require expert knowledge when configuring the system on a low abstraction level. Thus, an easy-to-use interface and the support of non-expert users are the main objectives of the BSS-1 software development effort. In addition to the neuroscientific API, the current BSS-1 software stack provides access a multitude of interfaces to manipulate the experiment description and system configuration on lower level, cf. sections III-A and III-B. For example, configuring certain hardware entities in a manual fashion but still being able to rely on the automated process in other areas facilitates the commissioning of new systems. This approach allows for exposing expert-level manipulation of the complex neuromorphic substrate without giving up all benefits of automation. Finally, increasing numbers of platform users, the parallel operation of production-type systems as well as commissioning future hardware generations pose challenges for the development and operation methodologies: platform robustness, experiment reproducibility and, to a lesser extent, the reduction of turnaround times between hardware revisions are essential.

C. Experiment Platform

Providing the research community with access to neuromorphic systems has been an ongoing effort pushed by large-scale research projects, such as the Human Brain Project. In the case at hand, a large-scale neuromorphic system is operated in a multi-user setting. From a high-performance computing (HPC) perspective, “fast” neuromorphic systems resemble spiking neural network accelerators. To efficiently utilize the available hardware systems requires solutions that are also common in HPC centers: resource management, time-sharing, fairness, accounting, monitoring and visualization. However, providing access to external researchers

also increases the need for robust operation, experiment reproducibility and support.

II. METHODS AND TOOLS

Performing well-controlled experiments on the BSS-1 system is the main task of BrainScaleS OS. In broad terms, experiments are defined by *what* and *when*, representing the data and the control flow. APIs for spiking neural network (SNN) descriptions, e.g., *PyNEST* or simulator-agnostic *PyNN* typically focus on the initial experiment setup, i.e. network topology, model parameters, plasticity rules, recording settings and the stimulus definition. Similarly, neuromorphic hardware requires an initial configuration that is typically performed before any stimulus is connected to the SNN.

Neuromorphic hardware is also different, as it often needs additional—more technical—settings compared to simulators which numerically calculate the time evolution of differential equations as, e.g., described by Einevoll et al. [21]. In our case, APIs that solely concentrate on aspects of neuron and synapse models, network topology and stimulus are insufficient, especially during commissioning. Usability for both experts and non-experts is the key feature of the software stack. The main points of usability are: 1) encapsulation of domain knowledge into software layers; 2) validity checking of settable parameters; 3) error reporting and explicit error handling; 4) consistency in the API layers concepts, and their representation of hardware entities; 5) availability of tested settings and configuration protocols to the user; 6) possibility to inject customized behavior at all levels of the software stack.

The development of the BrainScaleS-1 platform started already in 2008. Though many hardware and software components are now over a decade old, updates and improvements are continuously being made. In this section, we shortly describe the development methodology and foundations for the BrainScaleS OS.

A. Methodology

Compared to previous efforts made when developing the software stack for Spikey [18], large-scale neuromorphic systems introduce additional complexity. For example, multi-chip setups require more automation and robustness in all parts of configuration and runtime control. Hence, more people collaborate on different aspects of the system which, in turn, introduces friction in the development and commissioning process. When the development of BrainScaleS OS started, we employed version control, personal interaction and test frameworks. Within the first four years of development, a chat system [22]—also utilized for users support—, continuous integration [23] and formalized code review [24] were added to the development process. At the time of writing, over 10 000 changes were submitted for discussion. We do not adopt a strict development process framework, e.g., like *scrum*, however, we do include ourselves within the agile movement. Weekly meetings provide the scope for structured, long-term development, whereas our chat and code review systems encourage technical discussions of

details. The scenario is similar to CERN-style development where the developers are also to a large extent the users [25]. Over 120 individuals contributed across various projects. In the following paragraphs, we introduce the key concepts.

1) *Open Source*: Open Source software is a vital part in almost all fields of research. If not stated otherwise, the developed libraries and tools are published at <https://github.com/electronicvisions> under the *LGPL v2.1* [26] license. We also actively report bugs and push features upstream to third-party libraries.

2) *Software Design*: The long-term software and hardware development roadmaps are aligned to each other. Weekly software development meetings form the basis of the collaborative development. Problems and feature requests are discussed as well as medium-term development planning is performed. If needed, smaller teams are formed to come up with proposals that are then discussed in the plenum. However, the process is not fully democratic and at the end the maintainers take the final decisions.

3) *Review*: The BrainScaleS hardware and software development teams adopted an explicit review-based development scheme. Tracking of the development history and the current state of all components is handled by a set of version-controlled repositories. Developers propose changes to aspects of this state which are subsequently reviewed by other developers. At the end, the automated verification of each change and an iterative review process result in a final version which is then applied to the repository and becomes the new current state. This enables a rolling release scheme.

4) *Verification*: Based on the ideas of the continuous integration development methodology, the BSS verification methodology consists not only of software tests but also of hardware-based as well as simulation-based tests. For each proposed change the test result is fed back into the review system. When changes to software components are applied to the current state, the modified software is automatically deployed. Nightly tests serve as a measure for hardware platform health. The same experiment protocol can be used on hardware and in a combined FPGA and digital chip simulation. The latter is used for pre-tapeout verification during chip development.

5) *Software Environment*: Reuse of existing software packages reduces development costs but also introduces technical debt in the form of dependencies on external software packages [27]. BrainScaleS uses a containerized and explicit software dependency tracking system based on singularity [28] and spack [29]. Updating the software environment is based on the same review and verification system as before: developers propose a change to the dependency list, a testing container image is built and all tests are executed using this container image. If the container has build and code review as well as all tests have passed, the proposed change can be applied and the modified container image becomes the new default container image.

B. Foundations

1) *PyNN*: *PyNN* [19] is a simulator-agnostic domain-specific language for describing spiking neural network

models. Rooted in computational neuroscience, it focuses on the initial network topology, model parameters and plasticity rules, definition of input stimulus and “recording” settings.

Matching our goal of an backend-independent experiment description language for spiking neural networks, we adopted *PyNN* as our high-level API. However, BSS-1 is not as flexible as a software simulator. For example, it only supports a fixed neuron model, limited-resolution synapses and a sparse connectivity matrix. Transforming a user-defined *PyNN* experiment into a similar, well-fitting hardware configuration is challenging task. In particular, it is a matter of neuron and synapse placement, spike routing and model parameter translation. Due of imperfections of the analog substrate and limited resources, like bandwidth, there will always be differences to the user-defined target. For a detailed study, see Brüderle et al. [20].

2) *Programming Environment: GNU/Linux* is a flexible and well-supported host environment when developing custom hardware. In addition, computational neuroscience relies heavily on libraries and tools that are available for *NIX-like operating systems. Therefore, we only target *Linux*.

All core libraries are written in *C++* with the exception of parts of the transport layer that need a tight coupling to the *Linux* kernel and are therefore written in *C*.

We chose *C++* because of several reasons that are not unique to an neuromorphic operating system but apply to requirements of large software suites that have at least a modest need for performance and robustness. It is a multi-paradigm strongly-typed compiled language which leads to the discovery of many problems at compile time instead of runtime. *C++* supports many low-level manipulations that are essential when directly communicating with a custom hardware system. For example, in lower software layers the in-memory layout of data structures is required to match the formatting expected by the system. We always use the latest language standard that typical open-source compilers, e.g., *GCC* [30] and *LLVM* [31] support.

3) *Python Wrapping*: In experimental usage settings, scripting languages offer large advantages compared to compiled languages. For example, the read-eval-print loop (REPL) allows for iterative testing of the hardware and also for exploration of the software itself. Integration with the broad *Python* ecosystem of scientific libraries, e.g., *numpy* [32] or *matplotlib* [33], is an advantage in scientific efficiency. Therefore, we support *Python* in addition to *C++*. To link the *Python* and the *C++* world, we adopted a fully automated wrapper code generation scheme based on *py++* and *pygccxml* [34]. The generated wrapper code uses *boost::python* [35]. Customizations of the wrapping process have been collected in a library.

4) *Serialization*: Serialization describes the process of transforming in-memory data structures or object states into a store and loadable format. This format can be written to, e.g., disk and loaded to restore the in-memory data structures at a later point in time. Together with a framework for remote procedure calls, such as *RCF* [36], this allows for inter-process communication of higher-level data structures.

Though *C++* does not offer built-in support for serialization, it can be made available through several third party libraries. For BrainScaleS OS, we decided to use *boost::serialization*.

Listing 1: Example for *boost* serialization.

```
class Spike
{
// ...
private:
    template<typename Archiver>
    void serialize(Archiver& ar,
                  unsigned int const /*version*/)
    {
        using boost::serialization::make_nvp;
        ar & make_nvp("addr", m_addr)
          & make_nvp("time", m_time);
    }

    addr_type m_addr;
    time_type m_time;
};
```

Listing 1 exemplifies the serialization of the two member variables of the *Spike* class. More complex serialization functions are needed, e.g., when references and pointers are involved or different versions should be considered to support long-term compatibility for pre-existing data sets. However, *boost::serialization* has excellent support for all these scenarios. The on-disk format ranges from binary to text-based, such as *JSON* (custom extension) and *XML*.

5) Utility Libraries:

a) *Ranged enumeration types*: In *C++* numeric types do not have built-in support for range checks. Yet, it is beneficial to have such concepts, since over and under runs can be a threat to both correctness and security. The *rant* [37] library provides ranged integers and provides compile-time as well as runtime checking of ranges. In case of compile-time statements, violations produce compile errors; runtime errors raise exceptions. Checks are implemented to be lightweight enough to be included in production code. If required, ranged types can be replaced by their native counterparts via a compile flag to get rid of any remaining overhead. Ranged integers are heavily used for the coordinates described in the next section. Listing 2 demonstrates the use on the example of a ranged integer type.

Listing 2: Example for a ranged type.

```
rant::integral_range<int, 0, 5> ranged_integer;

// e.g. 6 => throws
ranged_integer = get_large_int();

// fails to compile (constexpr)
ranged_integer = -1;

// works
ranged_integer = 3;
```

b) *Python-style C++ convenience library: pythonic++* [38] brings some *Python*-style programming to *C++* for ease and more expressive code. Listing 3 exemplifies this idea on the enumeration during the iterations over a vector.

Listing 3: The `pythonic::enumerate` function can be used to count the iterations over an `STL` conform container.

```
using namespace pythonic;

typedef std::vector<int> vec;

for (auto v : enumerate(vec{0, -1337, 42}))
{
    std::cout << v.first << " " << v.second << '\n';
}
```




Fig. 2: The software stack as layers of abstraction. The core is the neuromorphic hardware BSS-1. It is followed by physical and software communication layers, i.e. the FPGAs and communication layer, followed by hardware abstraction and core functionality like mapping and routing. Next comes expert-level software used for blacklisting and calibration and the layer providing the `PyNN` abstraction. The last layer are the user-level applications and experiments.

c) *Bit Manipulation Library*: *bitter* [39] provides a common interface for bit operations on integral types as well as `std::bitset`. Operations like reversal or cropping to ranges are implemented.

III. IMPLEMENTATION

The BrainScaleS OS consists of several software component categories which are described in the following sections, see fig. 2 for an overview. At the end, the user is enabled to describe and execute a neuromorphic experiment without detailed knowledge of the underlying parts.

A. Configuration

The correct configuration of any hardware is a non-trivial task, e.g., one needs to ensure to write the correctly formatted command to the correct memory location, while taking into account other subtleties as, e.g., configuration order. In addition, there can be a mismatch between the addressing of individual circuit instances and their physical or logical placement. Another aspect is configuration timing as some

entities require settling times that have to be taken into account. To address the issues raised, the BrainScaleS OS comprises of several libraries that allow for an easy and correct configuration and control flow.

1) *Coordinate System*: In natural sciences, the proper choice of a coordinate system often strongly contributes to a simple and clean solution for a problem. We argue that this applies also to the usage of hardware in general and in particular to neuromorphic computing. The myriad of components on a wafer lead to a configuration space in the order of 50 MiB [20]. The memory for 4-bit weight and 4-bit address filter of the 384 chips \times 113k synapses per chip alone amounts to 41 MiB per wafer. One needs a representation of those components in software [40].

Many symmetries in chip layout combined with wafer-scale integration naturally lead to abstraction on different scales. Figure 3 gives an overview of a BSS-1 wafer and the structure of its components. Framed in blue one can see the layout of a single HICANN chip and its high degree of self-similarity. In the background one silicon wafer containing 384 of those chips is shown. This translational symmetry is reflected in a hierarchical structure of the coordinates. For each component we define a coordinate with the smallest granularity which then can be combined to define a higher hierarchical layer. We will illustrate this exemplarily with the coordinate for neuron circuits. First, we represent one neuron circuit on a chip: *NeuronOnHICANN*. This can then be combined with *HICANNOnWafer* resulting in *NeuronOnWafer* representing a specific neuron circuit on a wafer. Finally, *NeuronGlobal* can be composed from *NeuronOnWafer* and *Wafer* to uniquely identify one neuron circuit in the whole BSS-1 system. It is also possible to cast down to lower levels of representation, e.g., *NeuronOnWafer::toNeuronOnHICANN()*. Besides “lateral” conversions between hierarchical layers it is also possible to translate “horizontally” among coordinates on the same level. For example *SynapseOnHICANN::toNeuronOnHICANN()* yields the matching *Neuron* of a *Synapse*. See listing 4 for additional examples.

Another important feature is the possibility to create two dimensional grids that also have a notion of orientation, e.g., north and south. *SynapseOnHICANN* for example is structured in a grid of neurons per chip hemisphere and synapses per neuron. Grid coordinates also provide enumeration which is done in row-major order as shown in orange in fig. 3. Enumeration enables iteration of all coordinates which is supported in both, *C++* and *Python*, cf. listing 5 and listing 5. A string serialization exists that serves as both, convenient short format for logging and for argument parsing, cf. section III-C1. An example for this functionality can be found in listing 7. The consistence of this hierarchical structure is essential for a descriptive, reliable and maintainable low-level code base. About 80 distinct coordinate types are used to describe elements of a wafer module.

2) *Bit Formatting*: Typically, to configure a hardware entity a pair of an address and its content is needed.

Listing 4: Example coordinate conversion.

```
nrn = NeuronOnWafer(NeuronOnHICANN(Enum(5)),
    ↪ HICANNOnWafer(Enum(5)))
nrn.toNeuronOnHICANN()
nrn.toHICANNOnWafer()
```

Listing 5: Example coordinate iteration in C++.

```
for(auto nrn : iter_all<NeuronOnHICANN>()) {
    std::cout << nrn << '\n';
}
```

Listing 6: Example coordinate iteration in Python.

```
for nrn in iter_all(NeuronOnHICANN):
    print(nrn)
```

Listing 7: Example coordinate short formatting.

```
h = HICANNGlobal(HICANNOnWafer(Enum(5)), Wafer(6))
print(short_format(h))
# W006H005
print(from_string("W3"))
# Wafer(3)
```




Fig. 3: Chip structure and coordinate system of a BSS-1 wafer: the background shows a silicon wafer with highlighted structures of chips grouped in units of 4-by-2. The zoom-in shows one single HICANN chip layout. Framed in white are various component categories. Black lines illustrate the structure of intra- and inter-chip event buses. The row-major ordering scheme of a two-dimensional coordinate is shown in orange over a synapse array.

Addresses are used to identity writable and/or readable memory locations. The formatting of the content may depend on various aspects, e.g., the entity’s physical location on the chip. Therefore, functionality to address hardware entities and to format bits is an essential part of the software [41]. We use coordinates, section III-A1, to logically represent addresses. The content is represented as data structures encapsulating functionality of the underlying entity, cf. listing 8.

During development, for commission, expert use and debugging, an iterative & interactive usage is facilitated by *Python* bindings for the lower-level configuration functions.

Listing 8: Excerpt of the data structure *SynapseDriver*.

```
struct SynapseDriver {
public:
    bool is_enabled() const { return enable; }
    ...
private:
    bool enable = false;
    ...
};
```

Listing 9: Example for a pair of getter and setter functions of the stateless hardware abstraction layer.

```
void set_synapse_driver(
    Handle::HICANN& handle,
    SynapseDriverOnHICANN const& sdoh,
    SynapseDriver const& driver);

SynapseDriver get_synapse_driver(
    Handle::HICANN& handle,
    SynapseDriverOnHICANN const& sdoh);
```

This allows to, e.g., change parts of the configuration —also out of order w.r.t. the canonical flow— and directly observe the effects.

Listing 9 gives an example for a pair of getter and setter functions. The *handle* represents the backend, either accessing the hardware, the simulation backend described in section III-B3, or other debug facilities. The coordinate *SynapseDriverOnHICANN* specifies for which synapse driver the settings in the *SynapseDriver* container *driver* should be applied. The getter function is the same in reverse. Here, only the handle and the coordinate are passed. The bits read back from the hardware are decoded into and returned as a *SynapseDriver* object.

3) *High-level Configuration*: The core principal of the configuration of the neuromorphic hardware is that the user first specifies the desired state to which then the hardware is configured to [42]. Then, the hardware is configured to this state. To facilitate this user-driven configuration, all configurable settings have functional names, e.g., the neuron configuration. By this, a viable level of “self-documentation” is achieved. The user-facing configuration does not necessarily reflect the exact granularity in which the hardware can be configured, however, it does reflect, as stated above, an achievable final state which then also allows validity checks.

Listing 10: Example for the stateful hardware abstraction layer.

```
sthal::Wafer wafer;
auto& hicann = wafer[HICANNOnWafer(Enum(5))];
hicann.synapses[SynapseOnHICANN(Enum(123))].weight
    ↪ = SynapseWeight(3);
```

Listing 11: Example for per-FPGA parallelism via *OpenMP*.

```
#pragma omp parallel for schedule(dynamic)
for (size_t fpga_enum = 0; fpga_enum <
    ↪ FPGAOnWafer::end; ++fpga_enum) {
    ...
}
```

Listing 10 demonstrates how to set the weight of a single synapse. The needed objects and bookkeeping structures are created on the fly. Also, checks on the availability database,

see section III-C1, are performed and raise exceptions if the requested resources are not available.

The configuration is carried out with the maximum parallelism supported by the system, e.g., on a per-FPGA-basis with the help of *OpenMP*, see listing 11.

B. Control

1) *Experiment Control Flow*: The BSS-1 platform supports two distinct operation modes, both relying on FPGAs for data I/O and for control flow. Figure 4 illustrates the control flow for the primarily used mode, the *batch* mode, which suits independent pre-defined experiments.

In either case, the first step is configuring the neuromorphic hardware. Many configuration register accesses on the chip use a non-blocking access scheme requiring correct timing. This is implemented by inserting wait instructions between configuration commands; the time intervals use a static worst-case timing model. It is also important to configure hardware entities in a valid order. This is especially relevant when entities are configured in parallel, e.g., per FPGA. Synchronization barriers must then be added to the configuration flow so that only when all entities have reached a certain stage, configuration is continued. Another point to take care of is enabling triggers for, e.g., the recording of analog membrane traces which are supposed to start with the experiment; it is the point in time when the stimulus begins and recording of spike events is enabled.

Timed spike event release as well as recording is handled by 48 FPGAs on each wafer module. Each FPGA has accesses to 1.25 GiB of DRAM providing buffer memory for, e.g., input stimulus and recorded data. In case of the batch mode, the complete input to the network is predefined on the host, then sent to the FPGA and released upon experiment start. Simultaneously, spikes generated by the neuromorphic chips are recorded. At the end of the experiment, the recorded spikes and analog membrane traces are sent back to host. One typical application of this mode are deep neural networks where synaptic weights are optimized by an offline learning algorithm. If the hardware is allocated for a longer time period, the experiment framework also supports selectable automatic differential configuration reducing configuration overhead in later iterations.

In the other so called *hybrid* operation mode parts of the network or a virtual environment are simulated on the control cluster and interact with the spiking neural network running on BSS-1. This mode of operation is also known as *real-time closed-loop*. Control flow differs compared to the aforementioned mode as spike events sent by the host are not pre-buffered and timed by the FPGA but instead, they are directly injected into the chip upon arrival. Vice versa, emitted spikes from the network are directly sent to host and reacted upon. The challenge is to match the acceleration factor in both realms and keep the latency of the network communication as low as possible, see section III-B2.

2) *Communication Layer*: There are two main categories of data which need to be transferred between the neuromorphic hardware and conventional compute nodes. On the one hand

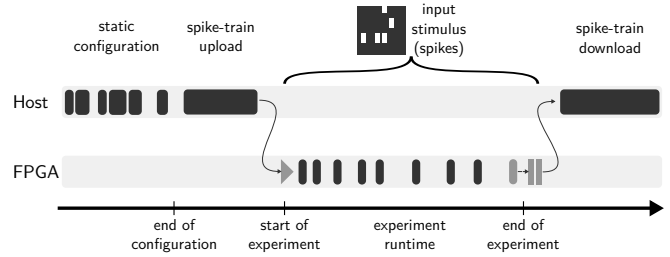


Fig. 4: Control flow of a typical experiment in batch mode. Black boxes indicate activity of host and FPGA during the different steps.

there is configuration data, e.g., neuron parameters, network topology, and on the other hand the activity of the network, i.e. spike events. Due to the accelerated operation of the BSS-1 system typical neuron activities of $O(100 \text{ Hz})$ result in on-wafer event rates exceeding Tera-Events per second. This demands high-throughput data exchange between BSS-1 and cluster control nodes. Ethernet was chosen as conventional data network equipment is readily available and, at the time of writing, commercial hardware supports bandwidths of up to 100 GiB/s . On BSS-1, the external connectivity is provided by 48 1-Gigabit Ethernet links. However, this bandwidth is still not sufficient to completely monitor the aforementioned on-wafer activity. The FPGAs provide an additional buffer stage for input and output data, but filtering and selecting in- and outputs is still needed. In the case of a deep neural network this would for example simply be the in- and output layers.

Furthermore, transfer of data, especially configuration data, needs to be robust. For Ethernet-based communication the Transmission Control Protocol (TCP) on top of the Internet Protocol (IP) is most commonly used as a reliable secure transport layer protocol. At the time of development there were no open source FPGA implementations of TCP available and even now available solutions are very resource demanding [43]. Hence, the BSS-1 FPGAs implement a custom sliding-window protocol with an automatic resend mechanism (ARQ) on top of the unreliable User Datagram Protocol (UDP) over IP. The software implementation has been open-sourced in the past [44]. Additional features analogous to congestion control, like roundtrip-time estimation as well as the slow start algorithm have been implemented in both, software and hardware.

However, the hybrid operation mode, cf. section III-B1, demands low-latency and low-jitter transport of spike events; configuration data is still transmitted via the reliable custom transport layer protocol but spike events are transferred best-effort facilitated by memory-mapped zero-copy receive and transmit ring buffers based on *PACKET_MMAP* [45].

Additional measures like setting CPU core affinity are taken to reduce jitter to a minimum on the host side.

3) *Hardware Simulation*: The so-called “executable system specification” [46] (ESS) is a hardware simulator of the BSS-1 system implemented in *C++/SystemC*. It contains behavioral, timing-accurate models of the digital components

and functional models of the analog neural components, e.g., the hardware neurons are numerically simulated AdEx neurons. Offering the same configuration interface as the real hardware and being fully executable, the ESS has been essential for the hardware-software co-design [20] and still serves as a validation tool for the software stack, especially for the mapping, configuration and experiment execution steps. In addition, the ESS allows to evaluate the effect of BSS-1 design-specific constraints (e.g., limited stimulation and recording bandwidth, spike time jitter, reduced parameter resolution) on neuromorphic experiments in isolation from distortions due to the mismatch of the mixed-signal circuits. For a detailed study see Petrovici et al. [47].

C. Conditions Support

Wafer-scale hardware operates under the assumption that individual components can be switched off and circumvented. In addition, the analog nature makes it necessary to, at least, apply a working point calibration. For this, conditions support libraries are put in place and described below.

1) *Availability Database*: Errors during the manufacturing process and the assembly of the wafer lead to varying conditions of individual components. Moreover, modifying hardware parameters may lead to a change in the response of these components. Disregarded, they might either distort simulation results or make the execution of experiments impossible in the first place. Consequently, it is mandatory to be aware of the state of the components and handle it dynamically. Therefore, the availability database was developed [48].

Combined with digital tests, cf. section IV-C1, this allows for storing and handling of the used components. It is implemented in C++ and uses XML with *boost::serialization* as the storage backend. Based on the coordinate system it stores a sparse representation of the flagged components without the notion of reasons as a whitelist or a blacklist. By this, the natural hierarchy of the system is mapped to the database. Thus, e.g., *HICANNOOnWafer* flags the full chip and *NeuronOnHICANN* flags only a single neuron circuit of the chip. Subsequently, using the database, other parts of the software can simply avoid the flagged components.

This also allows for the second use case of the availability database. Components can be marked artificially as not available to manipulate the hardware resources of an experiment without the requirement of an additional interface. *Python* bindings allow to construct convenience tools like a command line interface, cf. listing 12. As a result, a per experiment set of usable components can be generated and adapted dynamically.

Listing 12: Availability database command line interface.

```
redman_cli.py . W33H0 has neuron 0
# True

redman_cli.py . W33H0 disable neuron 1
```

2) *Parameter Translation and Calibration Database*: Microelectronics' manufacture deals with non-uniformities in the circuits produced across a silicon wafer. These transistor mismatches result in varying response from neuron to neuron circuit. To compensate, a calibration framework [49] maps high-level parameters to the hardware parameter space, homogenizing the response of neuron circuits.

First, the biological units are converted to the hardware compatible range. For time constants, the acceleration factor $\alpha = 1000 \dots 10\,000$ is taken into account:

$$\tau_{\text{hardware}} = \alpha \cdot \tau_{\text{biology}}.$$

Voltages also need to be scaled by s and shifted by o , respectively:

$$V_{\text{hardware}} = s \cdot V_{\text{biology}} + o,$$

with typical values $s = 10$ and $o = 1.2\text{ V}$. Similar conversions are needed for synaptic weights.

Now that the desired hardware values are known in physical units, the conversion to the digital domain can happen. This step does two things. The translation from physical units to digital units while at the same time taking into account variations from circuit to circuit, i.e. it applies calibration data. For this, the calibration database allows to store parameters for a set of pre-defined functions, e.g., polynomials. In addition, the transformation classes provide a numerical function inversion.

Also, the input values can be checked to lie within a given range of validity. The returned value can then be either clipped, an exception can be thrown or the validity range can be ignored.

Listing 13 demonstrates the usage of the library on the example of a linear function.

Listing 13: Example for a linear calibration function.

```
// linear transformation from, e.g., 0 - 1.8 V to 0
// - 1023 DAC
Polynomial linear({0.0, 1023./1.8}, 0.0, 1.8);
linear.apply(0.9);
// 511.5
linear.reverseApply(256);
// 0.450
linear.apply(2); // defaults to clipping
// 1023
```

D. Network Description

1) *PyNN Interface*: We implement the *PyNN*-API as a thin C++ library for which *Python* bindings are generated [50]. Compared to a *Python*-based implementation, this allows for a memory-efficient handling of larger data sets such as weight matrices of large neural networks, stimulus or recorded data. For the user, however, it appears like any other *PyNN* implementation, e.g., *PyNN.nest*. Internally, it translates from the *PyNN*'s imperative experiment description to an object-oriented description in the underlying C++ layer. The individual elements, e.g., populations and projections, are similar in their structure to *PyNN*. However, if found necessary, we restructure the data to our liking as

it is decoupled from the user-facing API. Also, pure *C++* usage is supported in a structured way.

2) *Map & Route*: Mapping and routing of neural networks described in *PyNN* to the neuromorphic hardware is a non-trivial task. The complexity and scope of the problem is similar to the synthesis of FPGA bitfiles. Therefore, the process is only described briefly. The full implementation can be found at [51]. Also, the map & route implementation undergoes substantial changes as new features and improvements are being developed.

In its simplest form, we implement a greedy strategy without back tracking. First, we place neurons from *PyNN* populations to hardware neuron circuits. Hereby, different neurons may be represented by a different number of neuron circuits on the hardware. Insertion points for spike input from the FPGAs are placed as well. The user has the option to constrain the automatic placement of neurons and spike sources. User parametrization is facilitated by a custom class that works aside from *PyNN*. Listing 14 shows how a user can restrict the placement of a population to a certain HICANN or to a list of allowed options.

Listing 14: Example for constraining placement.

```
pop = pynn.Population(...)
marocco.manual_placement.on_hicann(pop,
    ↳ HICANNOnWafer(Enum(42)))

pop2 = pynn.Population(...)
marocco.manual_placement.on_hicann(pop2,
    ↳ [HICANNOnWafer(Enum(23)),
    ↳ HICANNOnWafer(Enum(24))])
```

Listing 15: Example for querying a mapping result, where the hardware neurons corresponding to *PyNN* neurons are retrieved.

```
pop = pynn.Population(5, ...)
for pynn_neuron in enumerate(pop):
    items =
    ↳ runtime.results().placement.find(pynn_neuron)
    for item in items:
        for hardware_neuron in item.logical_neuron():
            ...
```

After the placement of neurons, the *PyNN* projections are transformed into synapses and on-chip routes on the hardware. This is the most time-consuming step as several hardware constraints must be taken into account, e.g., the limited number of allowed switches per route.

Listing 15 shows how to retrieve information on the allocated hardware after the placement: the hardware neuron circuits are looked up for all neurons of a *PyNN* population. This is useful for, e.g., directly manipulating low-level hardware parameters. The link between *PyNN* and the result of the map & route step is stored into an intermediate representation format, cf. fig. 5.

The on-chip bus network is represented as a graph using the *boost::graph* [52] library, where bus lines are vertices and switches are edges. During the creation of the graph, hardware availability data is already taken into account, i.e. hardware components that should not be used are not included in the graph representation. On-wafer routes can




Fig. 5: The transformation of *PyNN* to a hardware configuration (container) makes use of intermediate representations (IR). The IR also links *PyNN* and hardware entities and allows for look-ups in both directions.




Fig. 6: Screenshot of the web-based visualization. Chips are colored with increasing opacity proportional to the number of placed neurons. On-chip routes are also colored and can be click-selected to reveal more details.

be found by custom traversal algorithms of the graph or by using graph search algorithms like Dijkstra.

Being able to have a visual representation of the found hardware configuration is important for both, debugging and understanding possible improvements of manual or automatic placement. For this, a web-based visualization has been developed. Based on previous efforts to build visualization tools, the main requirement is to not replicate any code paths that are already part of the software stack. Another requirement is the possibility to run the tool locally and standalone, i.e. without the need for a server and the availability of the full software stack. This is possible by transpiling parts of the BrainScaleS OS *C++* libraries to *JavaScript*, including classes representing the map & route intermediate representation and its serialization implementation. Now, only the transpiled *JavaScript* libraries and the output of the mapping must be at hand. The top-level code is written in *TypeScript* [53].

We rely on *PixiJS* [54] for a fast 2D graphics engine supporting WebGL. It is capable of rendering large networks with many details of the hardware configuration. The tool offers different levels of details where, e.g., by zooming in all used synapses and neurons become visible. An example is shown in fig. 6.

IV. OPERATION

A. Resource Management and System Access

BrainScaleS neuromorphic platform resources are time-shared and partitioned between multiple experiments and/or users. In contrast to typical digital systems, analog neuromorphic hardware substrates are not homogeneous. Users need to be able to request specific hardware instances when running experiments. We use *SLURM* [55] —a HPC job scheduler— to handle resource requests for hardware components.

SLURM was extended utilizing its plugin API to handle various requirements related to inhomogeneous hardware resources. Being a mixed-signal neuromorphic system, individual BrainScaleS systems behave slightly differently which is why experimenters need a way to explicitly specify individual resource instances. The coordinate system described in section III-A1 is used to provide a familiar interface to the user. Different hardware components have varying degrees of granularity that can overlap and have interdependencies. We allocate the smallest needed subset of resources inferred from the user request. In principle, the Ethernet-based communication described in section III-B2 allows access to each FPGA from any conventional compute node in the same network. To prevent accidental clashes between concurrent experiment runs we separate individual BSS-1 modules into IPv4 subnets and deny access based on default firewall settings. When a user specifies a hardware resource for a *SLURM* job a firewall rule to accept traffic is automatically added during job runtime. Experiment software also compares its own resources with the allocated *SLURM* resources to detect possible mismatch.

On top of the direct access to the system as explained above, we provide access via HBP’s collab infrastructure [56]. Jobs are fetched from the HBP neuromorphic platform queuing service with the help of [57] and passed on to *SLURM*. Every few seconds, the job states of our scheduler and the upstream queue manager are synchronized.

B. Monitoring

Managing a large complex hardware system is unfeasible without extensive monitoring, as malfunction of any individual component can be fatal for operation. Monitoring can generally be split into three steps: aggregation, storage and visualization. Likewise there are two different types of data to be gathered, time-series data (e.g., voltage, temperature) and event data (e.g., powering off components).

The general flow for monitoring data of a wafer module is shown in fig. 7. There are around 1200 time-series data sources within one wafer module. Important sensors like wafer temperature are read out every few seconds. Additionally, events for powering parts on/off or alerts are generated. The data aggregation is performed on a Raspberry Pi via a software daemon handling several communications channels, e.g., I2C. On the Raspberry Pi a first data analysis is done in order to have a quick response to a dangerous system state. For example the temperatures are checked to be in an allowed range, above a given threshold the system

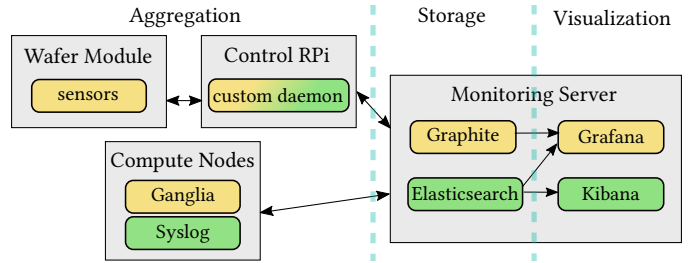


Fig. 7: Flow of monitoring data from aggregation to storage and visualization. Grey boxes represent involved devices on which the corresponding software libraries run. Software responsible for time-series data is shown in yellow and for event data in green. Arrows illustrate connectivity between the different components.

is turned off. Furthermore, all microcontrollers providing data to the Raspberry Pi perform local data evaluation tests and, therefore, detect false states faster than the Raspberry Pi. Time-series data is stored on a central *Carbon* [58] server outside the Raspberry Pi. For the conventional compute nodes we use *Ganglia* [59] for data aggregation which also feeds into the *Graphite* data base. *Graphite* uses a round-robin database for automatic data compression after certain intervals. Event aggregation is done utilizing *syslog* [60] which is parsed by *Logstash* [61]. Filtered events are stored in an *Elasticsearch* [62] database. *Grafana* [63] is used to visualize time-series data. It allows the creation of dashboards which give insight to the state of the system on various levels of detail. This facilitates getting a quick overview of relevant data from wafer modules and the state of the conventional compute cluster while simultaneously allowing to drill down for more details. Additionally, events like powering on components can also be shown in *Grafana* to easily link events and changes in time-series data. In general we use *Kibana* [64] to visualize event data.

C. Commissioning

1) *Digital memory tests*: In large complex hardware systems, variations of individual components are inevitable. As already mentioned in section III-C1 the behavior of these components change due to varying hardware parameters such as supply voltage and might disturb the execution of experiments. As a result, it is important to keep track of the state of the components and be aware of it during experiment execution. This is achieved by digital memory tests that are executed after assembly as well as periodically. Here, for a specific hardware configuration given by, e.g., supply voltage and clock frequency, each digital memory of every HICANN is read/write-tested with random values. The results are compared and if a malfunctioning component is found it is flagged in the availability management database. The database reflects the hierarchical structure of the hardware, so that always the largest functional unit that exclusively depends on the malfunctioning components is flagged, shown in fig. 8. The information can then be extended individually for each experiment and stored by serializing the updated database to disk, which is typically an XML-based file format, cf. section III-C1. During experiment execution this data is




Fig. 8: Digital test of malfunctioning components highlighted in grey. After the test (right side) these components are marked as not available using the hierarchy of the system. As a result, individual components up to large functional units, consisting of many components, are marked as not available.

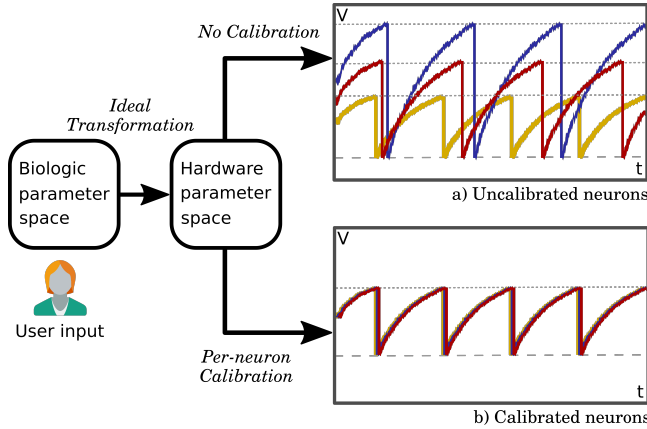


Fig. 9: Three neurons set for continuous spiking activity exhibit non-uniform threshold voltage under the same floating gate configuration. After applying the per-neuron calibration, parameters like $V_{threshold}$ can be set accurately across different neuron circuits (Traces were hand-drawn for illustrative purposes, with attention only to $V_{threshold}$).

then deserialized, which allows for skipping the unavailable components. Besides the experiment execution, the digital memory tests are also used in continuous integration to monitor and store the state of the hardware.

2) *Calibration*: The one-time circuit characterization [65] runs sequences of experiments that sweep the neuron parameters (stored as 10-bit values in analog floating gates on the HICANN), measures the changes' impacts, and employs different fits depending on the parameter effect's response. A calibration database is then filled with the transformation data for its utilization on routine hardware usage. The effect of applying such parameter mapping to the neuron configuration is exemplified in fig. 9.

V. APPLICATIONS

The previous sections motivated and detailed the status of the BrainScaleS OS. In the following, first a minimal experiment is demonstrated with the key concepts in action. It is followed by examples for more complex "full" experiments.

A. A Minimal Experiment

Listing 16: Example for an experiment.

```
import numpy as np

# BrainScaleS OS imports
import pyhmf as pyhmf
import pyhalco_hicann_v2 as C
from pyhalco_common import Enum
import pyhalbe
import pysthal
from pysthal.command_line_util import init_logger
from pymarocco import PyMarocco
from pymarocco.runtime import Runtime
from pymarocco.results import Marocco

init_logger("WARN", [])

marocco = PyMarocco()
runtime = Runtime(C.Wafer(33))

pyhmf.setup(marocco=marocco, marocco_runtime=runtime)

experiment_duration = 1000 # ms

neuron_parameters = {
    'cm': 0.2, # nF
    'v_reset': -20., # mV
    'v_rest': -20., # mV
    'v_thresh': 100, # mV
    'e_rev_I': -20., # mV
    'e_rev_E': 0., # mV
    'tau_m': 10., # ms
    'tau_refrac': 0.1, # ms
    'tau_syn_E': 2., # ms
    'tau_syn_I': 5., # ms
}

pop = pyhmf.Population(2, pyhmf.IF_cond_exp,
    ↪ neuron_parameters)
stimulus = pyhmf.Population(1, pyhmf.SpikeSourceArray, {
    'spike_times': [0, 5, 10]})

# choose a specific HICANN for given pyhmf population
marocco.manual_placement.on_hicann(pop,
    ↪ C.HICANNOnWafer(Enum(0)))

# record both, neuron membrane trace and spikes
pop.record()
pyhmf.PopulationView(pop).record_v()

# connect the stimulus
proj = pyhmf.Projection(stimulus, pop,
    ↪ pyhmf.AllToAllConnector(weights=0.001),
    ↪ target='excitatory')

# perform mapping but do not execute on hardware
marocco.backend = PyMarocco.None
pyhmf.run(experiment_duration)

# hardware-level settings
synapses_result =
    ↪ runtime.results().synapse_routing.synapses()
proj_items = synapses_result.find(proj)
for proj_item in proj_items:
    syn_c = proj_item.hardware_synapse()
    hicann = runtime.wafer()[syn_c.toHICANNOnWafer()]
    synapse = hicann.synapses[syn_c.toSynapseOnHICANN()]
    synapse.weight = pyhalbe.HICANN.SynapseWeight(3)

# run on hardware
marocco.skip_mapping = True
marocco.backend = PyMarocco.Hardware
pyhmf.run(experiment_duration)

np.savetxt("membrane.txt", pop.get_v())
np.savetxt("spikes.txt", pop.getSpikes())
```


Listing 16 shows an example experiment. It demonstrates all software features discussed in the previous sections. First, a couple of *Python* modules are imported. The *marocco* object is instantiated that allows for custom configurations that are not part of the *PyNN* API. The *runtime* object holds, amongst others, the *sthal* representation of the wafer configuration that will be used for low-level re-configuration. Next, parameters like the experiment duration and neuron parameters are set as variables. A population of neurons as well as a stimulus are created with two and one neuron, respectively. The population of neurons is placed on a specific HICANN. If manual placement is not given, the mapping software will find a location depending on the chosen mapping algorithm. No mapping hint is given for the stimulus. It will be inserted as close as possible to the mapped neuron population while adhering to bandwidth limitations as good as possible. The population is then asked to record both, its spikes and membrane potential. Next, a projection is drawn between the stimulus and the neurons. The projection is stored in a variable for later lookup.

Now that the network is completely setup, the mapping can be carried out, but it is not yet executed (*backend=None*). By this, the user can look up the hardware synapse between the stimulus and the neurons. Doing so we set manually a digital weight of 3. Then we skip the mapping, set the backend to hardware and execute. After *pynn.run* the resulting membrane trace and spikes can be read out.

Listing 17: Example for an experiment invocation.

```
# allocate the full module
srun -p experiment --wafer 33 experiment_example.py

# allocate only HICANN 0 (with analog readout and
  ↳ trigger by default)
srun -p experiment --wafer 33 --hicann 0
  ↳ experiment_example.py
```

The network execution is then invoked by calling listing 16 with a *SLURM* command like *srun*, see listing 17. The system on which the experiment is conducted is given as well as the partition which is used for accounting and priority.

B. Examples for Full Experiments

In the simple example explained above, the split between mapping and execution is only necessary if low-level access is wished. However, an important application are chip-in-the-loop experiments where it is crucial that an iterative re-configuration is possible to, e.g., compensate for trial-to-trial variations.

An experiment where this was used is detailed in [10] for training of a deep network for digit classification. Figure 10 shows the concept. After training an artificial neural network, the weights of a matching hardware network are set accordingly. However, due to both, trial-to-trial variations and differing responses of the artificial w.r.t. the hardware neurons, the classification performance is diminished. By continuing the training in the loop, the performance can be restored. For this, the response of all neurons in the

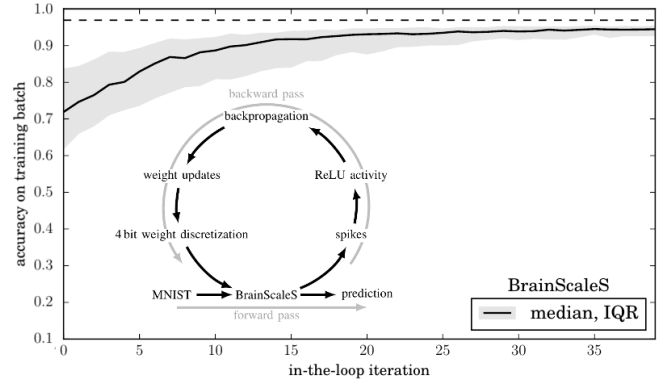


Fig. 10: Each iteration of in-the-loop training consists of two passes. In the forward pass, the output firing rates of the LIF network are measured in hardware. In the backward pass, these rates are used to update the synaptic weights of the LIF network by computing the corresponding weight updates in the ReLU network and mapping them back to the hardware. Adapted from Schmitt et al. [10].

hardware network is fed back into the training loop of the artificial network.

Accelerated physical emulation of Bayesian inference in spiking neural networks was demonstrated in Kungl et al. [66] where the full BrainScaleS OS was used as well. A network of spiking neurons was set up to sample from a Boltzmann distribution. The network was also trained iteratively, however, not in companion with an artificial network. Classification and pattern completion were demonstrated on two datasets.

Another example is Göltz et al. [67] that demonstrates classification based on spike timing only. Again, the hardware in-the-loop approach was used to train a network classifying images on BSS-1.

VI. FUTURE DEVELOPMENTS

A. Separation of Experiment Configuration and Execution

The experiment demonstrated in section V-A executed the mapping and the neuromorphic emulation in one process. Most importantly, the requested hardware resources had to be specified prior to the mapping. This is an unfortunate order as it does not allow for, e.g., choosing the system dynamically or specifying only a subset of the wafer for running experiments in parallel. However, the necessary ingredients to overcome this problem are in place and work is currently carried out to implement a solution based on the serialization capabilities of the data structures, see section II-B4.

B. Next-generation Python Binding Generation

The design and development of the BrainScaleS OS started in 2009. In the meantime, several external dependencies have been deprecated. In particular, our auto-generated *Python* wrapping depends on *gccxml* where development stopped in 2015. It depends on *gcc* $\leq 4.9.3$ blocking the usage of the latest *C++* features from the 14, 17 and 20 standards in header files. We evaluated several approaches, including the usage of *LLVM*'s *castxml*, but resorted to developing a new

wrapper code generator —*genpybind* [68]— which is based on *LLVM* libraries. The transition from the *py++*-based to *genpybind* is now in progress. In addition, *genpybind* also attacks binding generation from a different angle with a more fine-grained and explicit approach.

C. Towards BrainScaleS-2

Software development for BSS-2 started in 2016 and builds upon the results —the BrainScaleS OS— presented in this work. We try to re-use and adapt as much as possible from the existing code base. Especially the coordinate system, cf. section III-A1, has proven beneficial. However, early in the design phase we decided to improve the hardware abstraction layers by introducing structured types encapsulating all on-chip and on-FPGA configurable hardware entities. These types also provide explicit implementations for encoding to and decoding from hardware configuration bitstreams. In conjunction with a timer-based execution flow on the FPGA, this allows for experiments being described as an timed event sequence. The *C++* API makes heavy use of *std::future*-like interfaces to expose an asynchronous interface to the experiment control flow. Additionally, a fast experiment scheduler has been developed for BSS-2. It allows for approximately ten experiments per second where —due to the neuromorphic speed-up factor— each experiment represents up to 100 s of emulated time in the biological model. Based on this, restructuring work on BSS-1 started and a corresponding implementation was developed. Similarly, the *genpybind* tool was created during BSS-2 low-level software development. See Müller et al. [69] for a detailed description.

VII. DISCUSSION

This work describes the latest version of BrainScaleS OS, the software stack operating the BSS-1 platform. It allows to accomplish the main goal of the wafer-scale mixed-signal neuromorphic system BrainScaleS-1: designing and running wafer-scale experiments. The software stack aims for non-expert usage, e.g., by neuroscientists, while maintaining access to all other abstraction levels for expert users. We give a detailed overview of the individual software components and describe different aspects. From the hardware configuration, over the interaction with the system, e.g., setup, runtime control and result read out. We describe the transformation of user-defined experiments into a valid hardware configuration, as well as the necessary resource management and monitoring.

BrainScaleS adopted development methodologies and tools originating in software engineering to improve platform robustness and experiment reproducibility. BSS-1 is operated as a platform which is available for the research community.

Several experiments [12, 66, 67] demonstrate that BrainScaleS OS is a viable basis for using BSS-1. In addition to the publications, several thesis in our group made use of it for conducting neuromorphic experiments and commissioning work.

VIII. CONTRIBUTIONS

E. Müller is the lead developer and architect of the BrainScaleS software stack. S. Schmitt contributed to the calibration, the stateful configuration layer and the general usage flow. C. Mauch contributed to the system configuration layers as well as system operation. S. Billaudelle is a main contributor to the BSS-1 *PyNN* API implementation. A. Gröbl contributed software for low-level configuration as well as the system simulation backend. M. Güttler is the main developer of the system-level operation software, e.g., system monitoring and controlling, and contributed to low-level firmware. D. Husmann is the main developer of the system-level test software suite and contributed to system-level operation software. J. Ilmberger contributed to the communication layer and the analog readout framework. S. Jeltsch is the main developer of the map & route layer. J. Kaiser contributed to speed-up the synapse configuration. J. Klähn contributed to the map & route layer. M. Kleider contributed to the calibration of the system. C. Koke is the main developer of the stateful configuration layer and the calibration framework. J. Montes contributed to calibration scalability. P. Müller evaluated the performance of the BSS-1 neuromorphic circuit implementation. J. Partzsch contributed software for the low-level system configuration. F. Passenberg optimized the map & route algorithms to enable successful topology mapping to wafer modules with non-ideal blacklisting state. H. Schmidt contributed to digital blacklisting. B. Vogginger is a main developer to the simulation backend and contributed to the map & route layer. J. Weidner contributed to the web-based configuration visualization and acquired configuration results for the Jülich cortical column network. C. Mayr contributed to the system design (hardware and software) of the off-wafer communication stack. J. Schemmel is the lead designer and architect of the BrainScaleS-1 neuromorphic system. All authors discussed and contributed to the manuscript.

ACKNOWLEDGMENTS

The authors wish to thank all present and former members of the Electronic Vision(s) research group contributing to the BSS-1 hardware system, development and operation methodologies, as well as software development. The authors express their special gratitude towards: 1) Daniel Barley for his contribution to the parallel ADC-readout software; 2) Richard Boell for his contribution to the web-based visualization tool; 3) Patrick Häussermann for his contribution to experiment isolation in the scheduler; 4) Kai Husmann for his contribution to the low-level system control environment; 5) Lukas Pilz for his contribution to evaluate support for iterative configuration; 6) Vitali Karasenko for his contribution to the graphics; 7) Alexander Kononov for his effort when commissioning the HICANN chip; 8) Daniel Kutny for his contribution to the monitoring solution; 9) Dominik Schmidt for his contribution to the calibration framework; 10) Moritz Schilling for his contribution to the initial implementation of the secure transport layer protocol; 11) Andreas Baumbach, 12) Oliver Breitwieser and 13) Yannik

Stradmann for their work on continuous integration and structured deployment of the software environment. We especially express our gratefulness to the late Karlheinz Meier who initiated and led the project for most of its time.

This work has received funding from the EU ([FP7/2007-2013], [H2020/2014-2020]) under grant agreements 604102 (HBP), 269921 (BrainScaleS), 243914 (Brain-i-Nets), 720270 (HBP) and 785907 (HBP) as well as from the Manfred Stärk Foundation.

REFERENCES

- [1] Steve B. Furber, David R. Lester, Luis A. Plana, et al. "Overview of the SpiNNaker System Architecture". In: *IEEE Transactions on Computers* 99.PrePrints (2012). ISSN: 0018-9340. DOI: <http://doi.ieeecomputersociety.org/10.1109/TC.2012.142>.
- [2] Thomas Pfeil, Andreas Grübl, Sebastian Jeltsch, et al. "Six networks on a universal neuromorphic computing substrate". In: *Frontiers in Neuroscience* 7 (2013), p. 11. ISSN: 1662-453X. DOI: http://www.frontiersin.org/neuromorphic%5C_engineering/10.3389/fnins.2013.00011/abstract.
- [3] Mike Davies, Narayan Srinivasa, Tsung-Han Lin, et al. "Loihi: A neuromorphic manycore processor with on-chip learning". In: *IEEE Micro* 38.1 (2018), pp. 82–99.
- [4] Paul A Merolla, John V Arthur, Rodrigo Alvarez-Icaza, et al. "A million spiking-neuron integrated circuit with a scalable communication network and interface". In: *Science* 345.6197 (2014), pp. 668–673.
- [5] Johannes Schemmel, Daniel Brüderle, Andreas Grübl, et al. "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.
- [6] C. A. Mead. "Neuromorphic Electronic Systems". In: *Proceedings of the IEEE* 78 (1990), pp. 1629–1636.
- [7] S. Moradi and G. Indiveri. "An Event-Based Neural Network Architecture With an Asynchronous Programmable Synaptic Memory". In: *IEEE Transactions on Biomedical Circuits and Systems* 8.1 (Feb. 2014), pp. 98–107. ISSN: 1940-9990. DOI: [10.1109/TBCAS.2013.2255873](http://dx.doi.org/10.1109/TBCAS.2013.2255873).
- [8] J. Schemmel, A. Grübl, K. Meier, et al. "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.
- [9] Sebastian Millner, Andreas Grübl, Karlheinz Meier, et al. "A VLSI Implementation of the Adaptive Exponential Integrate-and-Fire Neuron Model". In: *Advances in Neural Information Processing Systems* 23. Ed. by J. Lafferty, C. K. I. Williams, J. Shawe-Taylor, et al. 2010, pp. 1642–1650.
- [10] Sebastian Schmitt, Johann Klähn, Guillaume Bellec, et al. "Classification With Deep Neural Networks on an Accelerated Analog Neuromorphic System". In: *Proceedings of the 2017 IEEE International Joint Conference on Neural Networks* (2017). DOI: [10.1109/IJCNN.2017.7966125](http://dx.doi.org/10.1109/IJCNN.2017.7966125). URL: <http://ieeexplore.ieee.org/document/7966125/>.
- [11] Chetan Singh Thakur, Jamal Lottier Molin, Gert Cauwenberghs, et al. "Large-Scale Neuromorphic Spiking Array Processors: A Quest to Mimic the Brain". In: *Frontiers in Neuroscience* 12 (2018), p. 891. ISSN: 1662-453X. DOI: [10.3389/fnins.2018.00891](http://dx.doi.org/10.3389/fnins.2018.00891). URL: <https://www.frontiersin.org/article/10.3389/fnins.2018.00891>.
- [12] Sebastian Schmitt, Johann Klähn, Guillaume Bellec, et al. "Classification With Deep Neural Networks on an Accelerated Analog Neuromorphic System". In: *Proceedings of the 2017 IEEE International Joint Conference on Neural Networks* (2017). DOI: [10.1109/IJCNN.2017.7966125](http://dx.doi.org/10.1109/IJCNN.2017.7966125). URL: <http://ieeexplore.ieee.org/document/7966125/>.
- [13] Oliver Rhodes, Petruț A. Bogdan, Christian Brennkmeijer, et al. "sPyNNaker: A Software Package for Running PyNN Simulations on SpiNNaker". In: *Frontiers in Neuroscience* 12 (2018), p. 816. ISSN: 1662-453X. DOI: [10.3389/fnins.2018.00816](http://dx.doi.org/10.3389/fnins.2018.00816). URL: <https://www.frontiersin.org/article/10.3389/fnins.2018.00816>.
- [14] Andrew G. D. Rowley, Christian Brenninkmeijer, Simon Davidson, et al. "SpiNNTools: The Execution Engine for the SpiNNaker Platform". In: *Frontiers in Neuroscience* 13 (2019), p. 231. ISSN: 1662-453X. DOI: [10.3389/fnins.2019.00231](http://dx.doi.org/10.3389/fnins.2019.00231). URL: <https://www.frontiersin.org/article/10.3389/fnins.2019.00231>.
- [15] Chit-Kwan Lin, Andreas Wild, Gautham N Chinya, et al. "Programming Spiking Neural Networks on Intel's Loihi". In: *Computer* 51.3 (2018), pp. 52–61.
- [16] Arnon Amir, Pallab Datta, William P Risk, et al. "Cognitive Computing Programming Paradigm: A Corelet Language for Composing Networks of Neurosynaptic Cores". In: *The 2013 International Joint Conference on Neural Networks (IJCNN)*. IEEE. 2013, pp. 1–10.
- [17] Catherine D. Schuman, Thomas E. Potok, Robert M. Patton, et al. *A Survey of Neuromorphic Computing and Neural Networks in Hardware*. 2017. eprint: [arXiv: 1705.06963](https://arxiv.org/abs/1705.06963).
- [18] Daniel Brüderle, Eric Müller, Andrew Davison, et al. "Establishing a novel modeling tool: a python-based interface for a neuromorphic hardware system". In: *Frontiers in Neuroinformatics* 3 (2009), p. 17. ISSN: 1662-5196. DOI: [10.3389/neuro.11.017.2009](http://dx.doi.org/10.3389/neuro.11.017.2009). URL: <https://www.frontiersin.org/article/10.3389/neuro.11.017.2009>.
- [19] A. P. Davison, D. Brüderle, J. Eppler, et al. "PyNN: a common interface for neuronal network simulators". In: *Front. Neuroinform.* 2.11 (2009). DOI: [3389/neuro.11.011.2008](http://dx.doi.org/10.3389/neuro.11.011.2008).
- [20] Daniel Brüderle, Mihai A. Petrovici, Bernhard Vogginger, et al. "A comprehensive workflow for general-purpose neural modeling with highly configurable neuromorphic hardware systems". In: *Biological Cybernetics* 104 (4 2011), pp. 263–296. ISSN: 0340-1200. URL: <http://dx.doi.org/10.1007/s00422-011-0435-9>.

- [21] Gaute T. Einevoll, Alain Destexhe, Markus Diesmann, et al. "The Scientific Case for Brain Simulations". In: *Neuron* 102.4 (May 2019), pp. 735–744. ISSN: 0896-6273. DOI: [10.1016/j.neuron.2019.03.027](https://doi.org/10.1016/j.neuron.2019.03.027). URL: <https://doi.org/10.1016/j.neuron.2019.03.027>.
- [22] Mattermost, Inc. *Mattermost: Open Source, Self-hosted Slack Alternative*. URL: <https://mattermost.com>.
- [23] Valentina Armenise. "Continuous Delivery with Jenkins: Jenkins Solutions to Implement Continuous Delivery". In: *Proceedings of the Third International Workshop on Release Engineering*. RELENG '15. Florence, Italy: IEEE Press, 2015, pp. 24–27.
- [24] *Gerrit Code Review*. <https://www.gerritcodereview.com/>. accessed March 11, 2020. 2020.
- [25] René Brun, Federico Carminati, and Giuliana Galli Carminati, eds. *From the Web to the Grid and Beyond*. Springer, Berlin, Heidelberg, 2012.
- [26] *GNU Lesser General Public License*. Version 2.1. Free Software Foundation. URL: <http://www.gnu.org/licenses/gpl.html>.
- [27] Russ Cox. "Surviving Software Dependencies". In: *Commun. ACM* 62.9 (Aug. 2019), pp. 36–43. ISSN: 0001-0782. DOI: [10.1145/3347446](https://doi.org/10.1145/3347446). URL: <https://doi.org/10.1145/3347446>.
- [28] Gregory M. Kurtzer, Vanessa Sochat, and Michael W. Bauer. "Singularity: Scientific containers for mobility of compute". In: *PLOS ONE* 12.5 (May 2017), pp. 1–20. DOI: [10.1371/journal.pone.0177459](https://doi.org/10.1371/journal.pone.0177459).
- [29] Todd Gamblin, Matthew LeGendre, Michael R. Collette, et al. "The Spack Package Manager: Bringing Order to HPC Software Chaos". In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. SC '15. Austin, Texas: ACM, 2015, 40:1–40:12. ISBN: 978-1-4503-3723-6. DOI: [10.1145/2807591.2807623](https://doi.org/10.1145/2807591.2807623).
- [30] B. Gough and R.M. Stallman. *An Introduction to GCC: For the GNU Compilers Gcc and G++*. Network theory manual. Network Theory, 2005. ISBN: 9780954161798. URL: <https://books.google.de/books?id=yIGKQAAACAAJ>.
- [31] Chris Lattner and Vikram Adve. "LLVM: A Compilation Framework for Lifelong Program Analysis and Transformation". In: San Jose, CA, USA, Mar. 2004, pp. 75–88.
- [32] S. van der Walt, S. C. Colbert, and G. Varoquaux. "The NumPy Array: A Structure for Efficient Numerical Computation". In: *Computing in Science Engineering* 13.2 (Mar. 2011), pp. 22–30. ISSN: 1558-366X. DOI: [10.1109/MCSE.2011.37](https://doi.org/10.1109/MCSE.2011.37).
- [33] John D. Hunter. "Matplotlib: A 2D Graphics Environment". In: *IEEE Computing in Science and Engineering* 9.3 (2007), pp. 90–95.
- [34] Roman Yakovenko. *pygccxml/py++*. URL: <https://sourceforge.net/projects/pygccxml>.
- [35] Boost.Python. *Version 1.71.0 Website*. http://www.boost.org/doc/libs/1_71_0/libs/python. 2019.
- [36] Delta V Software. *Remote Call Framework*. URL: www.deltavsoft.co.
- [43] Mario Ruiz, David Sidler, Gustavo Sutter, et al. "Limago: an FPGA-based Open-source 100 GbE TCP/IP Stack". In: Sept. 2019. DOI: [10.1109/FPL.2019.00053](https://doi.org/10.1109/FPL.2019.00053).
- [45] *PACKET(7) Linux Programmer's Manual*. Feb. 2020. URL: <http://man7.org/linux/man-pages/man7/packet.7.html>.
- [47] Mihai A Petrovici, Bernhard Vogginger, Paul Müller, et al. "Characterization and Compensation of Network-Level Anomalies in Mixed-Signal Neuromorphic Modeling Platforms". In: *PLOS ONE* 9.10 (2014), e108590.
- [52] Boost.Graph. *Version 1.71.0 Website*. http://www.boost.org/doc/libs/1_71_0/libs/graph. 2019.
- [53] Microsoft. *TypeScript: JavaScript For Any Scale*. URL: <https://www.typescriptlang.org/>.
- [54] *Pixijs 5*. 2019. URL: <https://www.pixijs.com/>.
- [55] Andy B Yoo, Morris A Jette, and Mark Grondona. "Slurm: Simple linux utility for resource management". In: *Workshop on Job Scheduling Strategies for Parallel Processing*. Springer. 2003, pp. 44–60.
- [56] Katrin Amunts, Christoph Ebell, Jeff Muller, et al. "The Human Brain Project: Creating a European Research Infrastructure to Decode the Human Brain". In: *Neuron* 92.3 (Nov. 2016), pp. 574–581. ISSN: 0896-6273. DOI: [10.1016/j.neuron.2016.10.046](https://doi.org/10.1016/j.neuron.2016.10.046). URL: <https://doi.org/10.1016/j.neuron.2016.10.046>.
- [57] Human Brain Project. *Python client for the Human Brain Project Neuromorphic Computing Platform*. URL: <https://github.com/HumanBrainProject/hbp-neuromorphic-client>.
- [58] Graphite Project. *Carbon*. URL: <https://github.com/graphite-project/carbon>.
- [59] M. Massie, B. Li, B. Nicholes, et al. *Monitoring with Ganglia*. O'Reilly and Associate Series. O'Reilly Media, Incorporated, 2012. ISBN: 9781449329709. URL: <http://books.google.de/books?id=w4LLpXeVCbkC>.
- [60] R. Gerhards. *The Syslog Protocol*. RFC 5424. RFC Editor, Oct. 2009. URL: <https://www.rfc-editor.org/rfc/rfc5424.txt>.
- [61] elastic. *Logstash: Collect, Parse, Transform Logs*. URL: <https://www.elastic.co/logstash>.
- [62] elastic. *Elasticsearch: The Official Distributed Search & Analytics Engine*. URL: <https://www.elastic.co/elasticsearch>.
- [63] Grafana Labs. *Grafana: The open observability platform*. URL: <https://grafana.com>.
- [64] elastic. *Kibana: Explore, Visualize, Discover Data*. URL: <https://www.elastic.co/kibana>.
- [66] Akos F. Kungl, Sebastian Schmitt, Johann Klähn, et al. "Accelerated Physical Emulation of Bayesian Inference in Spiking Neural Networks". In: *Frontiers in Neuroscience* 13 (2019), p. 1201. ISSN: 1662-453X. DOI: [10.3389/fnins.2019.01201](https://doi.org/10.3389/fnins.2019.01201). URL: <https://www.frontiersin.org/article/10.3389/fnins.2019.01201>.

- [67] Julian Göltz, Andreas Baumbach, Sebastian Billaudelle, et al. *Fast and deep neuromorphic learning with time-to-first-spike coding*. 2019. eprint: [arXiv:1912.11443](https://arxiv.org/abs/1912.11443).
 - [68] Johann Klähn. *genpybind software v0.2.0*. 2020. DOI: [10.5281/zenodo.372674](https://doi.org/10.5281/zenodo.372674). URL: <https://github.com/kljohann/genpybind>.
 - [69] Eric Müller, Christian Mauch, Philipp Spilger, et al. “Extending BrainScaleS OS for BrainScaleS-2”. In: *arXiv preprint* (Mar. 2020). URL: [TODO](#).
- OWN SOFTWARE
- [37] Sebastian Jeltsch. *rant*. URL: <https://github.com/ignatz/rant>.
 - [38] Electronic Visions(s), Heidelberg University. *pythonic*. URL: <https://github.com/electronicvisions/pythonic>.
 - [39] Sebastian Jeltsch. *bitter*. URL: <https://github.com/ignatz/bitter>.
 - [40] Electronic Visions(s), Heidelberg University. *halco*. URL: <https://github.com/electronicvisions/halco>.
 - [41] Electronic Visions(s), Heidelberg University. *halbe*. URL: <https://github.com/electronicvisions/halbe>.
 - [42] Electronic Visions(s), Heidelberg University. *sthal*. URL: <https://github.com/electronicvisions/sthal>.
 - [44] Eric Müller, Moritz Schilling, and Christian Mauch. *HostARQ Slow Control Transport Protocol*. Apr. 2018. URL: <https://github.com/electronicvisions/sctrltp>.
 - [46] UHEI, TUD. *ESS*. URL: <https://github.com/electronicvisions/systemsim-stage2>.
 - [48] Electronic Visions(s), Heidelberg University. *redman*. URL: <https://github.com/electronicvisions/redman>.
 - [49] Electronic Visions(s), Heidelberg University. *calibtic*. URL: <https://github.com/electronicvisions/calibtic>.
 - [50] Electronic Visions(s), Heidelberg University. *pyhmf*. URL: <https://github.com/electronicvisions/pyhmf>.
 - [51] Electronic Visions(s), Heidelberg University. *marocco*. URL: <https://github.com/electronicvisions/marocco>.
 - [65] Electronic Visions(s), Heidelberg University. *cake*. URL: <https://github.com/electronicvisions/cake>.