Sustainable operation of research infrastructure for novel computing

Yannik Stradmann*, Joscha Ilmberger*, Eric Müller, and Johannes Schemmel

***Abstract*—Novel compute systems are an emerging research topic, aiming towards building next-generation compute platforms. For these systems to thrive, they need to be provided as research infrastructure to allow acceptance and usage by a large community. By the example of the neuromorphic BrainScaleS-2 system, we showcase the transformation from a laboratory setup to a sustainable, publicly available platform. It is embedded into a purpose-built institute, tightly coupling a conventional cluster with novel compute hardware. The network infrastructure is optimized for robust operation, even in the case of unintended behavior of individual devices. The systems themselves are packaged into 19-inch compatible units to allow for easy maintenance and extension. We operate the platform using modern CI/CD techniques and continuously assert its health using automated system monitoring. Finally, we share our lessons learned during the decade-long endeavor of operating analog neuromorphic systems as a publicly available research platform.**

***Index Terms*—analog computing, platform operation, continuous integration**

I. Introduction

Neuromorphic hardware describes a variety of novel computing approaches aiming to mimic data processing strategies found in neurobiology. Taking inspiration from the sparsity in communication found in such spiking neural systems, many of these devices target ultra-low-power edge applications [1]–[8]. As such, they are usually developed and deployed as standalone system-on-chips (SoCs). In contrast, large-scale digital systems have been put forward to accelerate compute-intensive workloads and facilitate research in cognitive neurosciences [6], [9]–[13]. Combining both approaches, the BrainScaleS-2 (BSS-2) platform embeds an analog neural network core into digital periphery. It targets biologically inspired multi-timescale online-learning experiments and therefore emulates neuronal dynamics with a speed-up factor of 1 000 compared to biological real-time [14].

BSS-2 is envisioned as a versatile research platform for the interdisciplinary fields of cognitive neurosciences and machine learning, requiring continuous operation and remote accessibility. It combines a user-facing software stack and services [15], with physical infrastructure to host the neuromorphic platform. In contrast to many systems based on novel materials, the BSS-2 ASIC uses analog CMOS devices as computational units and therefore does not require a specialized operational environment. Combined with the maturity of the platform, this allows us to adopt practices established by data center operation.

Figure 1. BrainScaleS-2 platform in the machine hall of the European Institute for Neuromorphic Computing. Two rack cabinets house up to 16 systems, each containing two ASICs. They are placed on retractable drawers to allow convenient access for maintenance. Each of these contains all components for operation, requiring only mains supply and a single Ethernet uplink. Both are supplied using local cable passages through the floor to the server room beneath (not visible in the picture).

This manuscript describes the physical infrastructure for operating BSS-2, one of the world's largest analog compute systems (Fig. 1). We will introduce the Ethernet-based network topology, the composition of the setup, continuous integration workflows, and health monitoring. Finally, we describe lessons learned during the multi-year efforts of operating this infrastructure in an academic setting.

II. Network infrastructure at EINC

The European Institute for Neuromorphic Computing (EINC) is purpose-built to host research of novel compute systems at large scale. Across four floors, it houses a state-of-the-art server room, a machine hall, and laboratories: In its maximum

* These authors contributed equally to this work, listed in random order.
Yannik Stradmann, Joscha Ilmberger and Eric Müller are with the Kirchhoff-Institute for Physics, Heidelberg University. Johannes Schemmel is with the Institute of Computer Engineering, Heidelberg University.

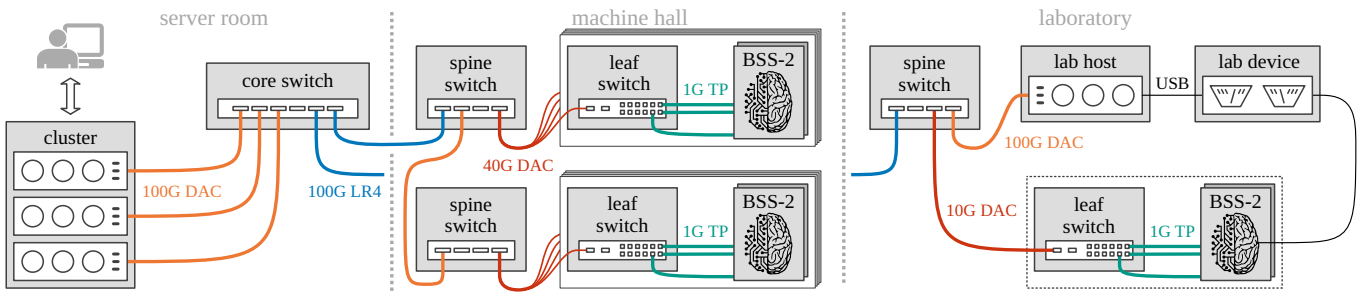

Figure 2. Network infrastructure of the BrainScaleS-2 (BSS-2) systems, as installed at the European Institute for Neuromorphic Computing (EINC). The installation is distributed across three floors: Users execute experiment code on a compute cluster in the basement, where all nodes are equipped with a dedicated 100 GbE uplink for experiment traffic. They connect to a 100 GbE core switch, which in turn provides connectivity to the machine hall (production systems) and laboratories (test systems) via fiber-optic connections. In both cases, each rack has a small 100 GbE spine switch, splitting to multiple 10 GbE leaf switches. Up until this point, all network traffic consists of multiple VLANs dedicated to experiment and management data. The leaf switch finally provides untagged 1 GbE connections to the BrainScaleS-2 setups.

designed configuration, the server room can accommodate 63 racks with 1.5 MW "indirect free cooling" capacity. The conventional compute cluster in operation consists of 26 nodes spread across three racks with a total installed energy budget of 30 kW. Users execute neuromorphic experiments on the BSS-2 platform via this cluster (Fig. 2). It orchestrates hardware access and transfers data to and from the systems, each utilizing a 1 GbE link. A core switch (HPE FM3032Q) interconnects all cluster nodes with the neuromorphic hardware in the machine hall and laboratories via 100 GbE. This dedicated physical network is segmented in multiple IEEE 802.1Q VLANs for experiment and different categories of management data:

1) Infrastructure management (power distribution, network)
2) Cluster node management (IPMI)
3) BSS-2 system management
4) Experiment data to/from BSS-2 systems

Together with queing disciplines providing minimum bandwidth guarantees, this separation ensures management access in case of unintended behavior of any experiment device, such as denial-of-service.

Network traffic from the server room is distributed to the machine hall (production systems) and laboratories (test systems and maintenance) using fiber-optic connections. In both cases, it is further distributed through 100 GbE spine switches (Mikrotik CRS504-4XQ-IN) to leaf switches (Mikrotik CRS326-24G-2S) via 40 GbE-to-4×10 GbE breakout cables. Here, the VLANs terminate into untagged 1 GbE ports, to which individual BSS-2 systems connect. They are comprised of one management controller (VLAN 3) and two ASICs, which communicate experiment data via FPGAs (VLAN 4).

In addition to this common infrastructure, the laboratories are equipped with remote cluster nodes that accommodate for measurement equipment without networking capabilities, such as USB devices. This setup enables a location-agnostic operation of the BSS-2 systems in production and during maintenance.

The chosen network hierarchy supports concurrent operation of all neuromorphic systems at line speed, with headroom for future extensions.

III. Systems

We house all BSS-2 systems in conventional 19-inch rack cabinets, which allows easy integration with off-the-shelf components for power distribution and networking. Since the neuromorphic systems have originally been developed for laboratory use, their physical form factor does not comply with this standard. We therefore mount them on retractable drawers, which allows easy access and removal for maintenance. Figure 3 shows such a unit, where—in addition to two BSS-2 systems—all components necessary for operation from mains power and a single network uplink are placed. Specifically, we provide 12 V DC power through a single AC adapter per drawer, which feeds all FPGAs, digital periphery, system controllers and cooling fans. The ASIC supplies are derived from a separate 6 V AC/DC adapter per system to prevent coupling of digital noise into the analog circuits. Data connectivity is established through a single 10 GbE uplink, which provides multiple separate VLANs. The managed Ethernet leaf switch exposes these virtual networks as untagged ports for the experiment systems and their ARM-based system controllers.

This controller has been inherited from the previous generation of BrainScaleS [16]. It provides I^2C and JTAG connections to the associated neuromorphic system for power management and configuration. In addition, it continuously collects 29 health metrics, such as power supply status and temperatures. This data is streamed to a database over a dedicated Ethernet connection (see Section V). Over the same interface, users can remotely access system management functionality and JTAG connections to all FPGAs.

IV. Platform Integration into CI/CD

To ensure robust operation of the BSS-2 platform, we employ a workflow based on continuous integration and -delivery (CI/CD). In addition to human review, any change to software or hardware components proposed by a developer is automatically verified in simulation and on the hardware systems. As such, the BSS-2 platform itself is tightly integrated into the CI/CD workflow. We demonstrate the steps involved by the example of an in-review software modification that requires an update of the system's FPGA design (Fig. 4)

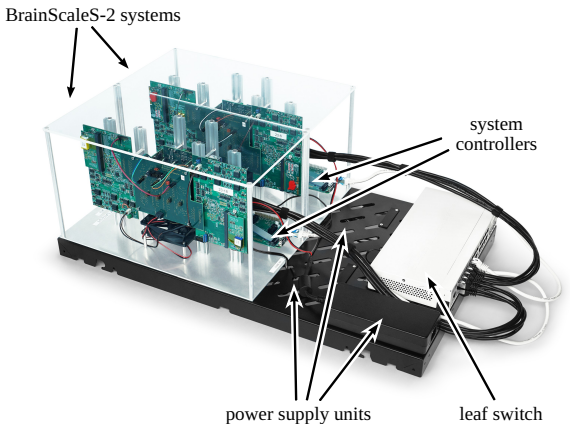

Figure 3. A single drawer of the BrainScaleS-2 platform, containing two systems[2] and all necessary periphery for operation from mains power and a single 10 GbE uplink. This includes ARM-based system controllers, three power supply units and the leaf switch also shown in Fig. 2.

1) The developer modifies the RTL code of the FPGA locally. They indicate the co-dependent software change in the commit message and push the patch to the Gerrit code review system.
2) Gerrit notifies a security-hardened Jenkins instance with access to commercial EDA tools about this proposed changeset. This triggers multiple jobs in parallel:
   a) RTL-based simulations of the BSS-2 system (ASIC and modified FPGA). The executed testsuite includes all referenced software modifications.
   b) Build of the FPGA design, resulting in a binary hardware configuration. This file, together with metadata, is distributed to a cluster-wide staging area.
3) Once step 2b has completed, a second—less restricted—Jenkins instance is triggered to perform downstream verification on the physical hardware systems:
   a) It builds the BSS-2 software stack [19] together with test suites ranging from transport-layer validation to full-stack experiments—again, including the referenced software modifications.
   b) A BSS-2 hardware setup is allocated and the staged configuration from step 2b is loaded into the FPGA. Subsequently, the prepared test suites are executed on the hardware system.
4) All test results are aggregated and reported to the Gerrit instance in form of a mandatory vote.
5) After the changeset has additionally been approved by a human reviewer and submitted by the author, a release build is run and delivered as new stable version.

The strong isolation of individual system components described in Section II as well as the runtime reconfiguration of the FPGA allows for verification of work-in-progress changesets that may introduce unintended behavior.

[2]The shown BrainScaleS-2 systems have been developed in collaboration with the Chair of Highly-Parallel VLSI Systems and Neuro-Microelectronics at Technische Universität Dresden.

The tight integration of the BSS-2 hardware platform into the development workflow ensures robust operation. Additionally, this standardized workflow enables efficient collaboration of researchers throughout all levels of seniority.

V. Operational Monitoring

The health of the BSS-2 platform is continuously monitored on multiple levels ranging from high-level experiment fidelity through network availability to power supply status. We collect any time series data in a Graphite database, event-like data in InfluxDB and use Grafana for visualization and alerting. In addition, system operators annotate events, such as power outages or maintenance, directly within the monitoring data.

We collect information about the system state on multiple timescales: Nightly Jenkins jobs assess the fidelity of experiments, such as training of spiking neural networks for classification tasks. Similarly, we run bihourly high-level health checks on all unoccupied systems. These validate the overall system state, including tests for SRAM access, high-speed communication channels, and ASIC power supplies. In addition, we read out and store the currently configured FPGA design revision for future reference.

Data that has to be available on a finer timescale is collected using dedicated monitoring services: Once every minute, we execute ICMP echo- and ARP requests from a central host to the FPGAs of all BSS-2 systems and evaluate the results in terms of response ratio and -time. Physical monitoring data, specifically supply voltages and system temperature, is collected by the system controller introduced in Section III every second.

Quantities that need to be observed on shorter timescales, such as the ASIC's power consumption, are not considered to be part of operational monitoring and rather available to users directly within the experiment result structures.

This multi-level approach allows us to maintain high system availability through early notifications in case of failures and degradation. The collected data has proven vital for understanding unexpected system behavior by allowing for comparison to the past.

VI. Lessons learned

The presented system represents the current iteration of a decade-long endeavor to provide a reliable platform for neuromorphic computing from within a university-based research initiative. While their applicability to other environments may vary, we nevertheless want to share lessons we have learned during this process:

a) Strict resource management: BSS-2 systems can be used concurrently by different users in an interleaved fashion, with individual runtimes ranging from seconds to days. All user-accessible devices must therefore be managed and any interference with unallocated resources must be prevented through automation. While malicious intent usually is not part of our threat model, we have made the experience that unintentional access can lead to significant disturbance of system availability. We therefore integrate all available hardware resources into the SLURM resource manager [20]

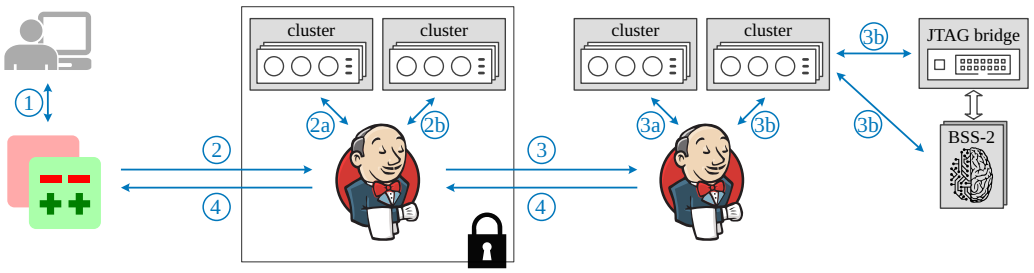

Figure 4. CI/CD workflow for proposed changes to the BrainScaleS-2 software and hardware configuration. Developers upload their patches to a Gerrit instance for code review [17]. Each upload triggers a build on a security-hardened Jenkins [18] instance with access to servers containing commercial EDA tools. These run RTL simulations of the combined system and build a bitfile. Once the latter is available in binary form, a second Jenkins instance executes downstream tests using this staging bitfile. To do so, it builds the BrainScaleS-2 software stack—including changes, the original patch might require—, configures the staging bitfile on a test setup and executes the tests. The results will be propagated back to the original patch submitter and be available next to human code-review comments.

and dynamically adjust firewall rules upon system allocation. Without an active allocation, no package will be routed to the respective system.

b) Independence of hardware resources: In line with the strict resource management through software mentioned above, hardware should be built in a way that two independent systems cannot influence each other (e.g., through adjustments of a shared power supply).

c) Network separation: While Ethernet has proven to be a very robust protocol for our use case, unintended behavior of a single device may stall operation through denial-of-service. For such cases, sufficient bandwidth must be reserved for management traffic to shut down the misbehaving device. While separate physical networks are ideal for this purpose, they pose significant financial and spatial constraints. We therefore utilize queuing disciplines in our switch infrastructure to guarantee minimum bandwidth allocations for specific VLANs.

d) Prefer Ethernet over USB: Previous iterations of the BSS-2 platform utilized deep trees of USB devices for managing the individual systems. Since then, we have migrated to Ethernet where possible—specifically, for the SoC-based system controller described in Section III—and also prefer this protocol for new measurement equipment: In contrast to USB, Ethernet poses fewer constraints on tree size and -topology [21] and allows galvanic isolation between devices. USB additionally requires a fixed pairing of devices to a single host system, which in practice poses strong requirements on the locality of nodes. Ethernet-based networking, on the contrary, allows multiple—remote—hosts to access the same equipment and thereby allows for fault-tolerance and load-balancing.

e) Monitoring with long retention periods: A reliable monitoring infrastructure is crucial for any sustainable platform operation. Here, we'd like to additionally highlight that retention periods of high-resolution data must be longer than it takes users to notice typical fail cases. For the specific application described in this manuscript, the intrinsic stochasticity of most experiments can lead to multi-month periods in which users accumulate errors until they report them.

f) Physical access control: Even in a university-based context, physical access to production systems must be limited to a small group of maintainers. While this statement itself is trivial, it creates additional constraints on the system—especially for research platforms: All functionality must be controllable remotely, including controlled access to different reset domains, power management and data acquisition.

VII. Discussion

We have shown how we transformed a laboratory system for exploring novel compute to a publicly available research platform operated in an academic environment.

While we base our approach on data centers, the presented implementation poses certain shortcomings: Most importantly, we do not target high availability of the systems and therefore omit redundancy for most power supplies and networking. Especially with academic personnel—instead of dedicated technical staff—operating the system, maintenance periods may be elongated. The building itself does not come with a continual power system, making the infrastructure susceptible to power outages.

Other fields of novel computing have shown similar approaches of integrating laboratory systems into compute cabinets [22]. While the CMOS-based compute substrate of BSS-2 does not require similarly complex auxiliary components, its acceleration factor still poses strong constraints on network bandwidth and -latency. We therefore base our core network infrastructure on 100 GbE. It allows—due to the strong separation based on VLANs—the described hardware-enabled CI/CD workflow to be executed on the production systems. The infrastructure is prepared to accommodate additional scaled-up systems in active development.

The presented neuromorphic platform—BrainScaleS-2—has been publicly available since 2017, initially through the Human Brain Project and later via the European research infrastructure EBRAINS[3]. Together with its predecessor system, it was moved to its current location in the EINC in 2023, with currently 13 systems in productive operation. Between January 2024 and June 2025, more than 170 individual researchers conducted an average of over 5 700 experiments per day.

Acknowledgements

The authors would like to thank Christian Mauch for his contributions to hardware resource management and the system

[3] Access to BrainScaleS-2 can be requested via https://ebrains.eu/nmc

controller, Julian Göltz for spearheading automated system health assessment when training spiking neural networks on the platform, and the Chair of Highly-Parallel VLSI Systems and Neuro-Microelectronics at Technische Universität Dresden for their contributions to BrainScaleS-2. Finally, we thank all past and present members of the Electronic Visions group for their contributions to the BrainScaleS platforms.

Funding

The presented work has received funding from the EC Horizon 2020 Framework Programme under grant agreements 720270 (HBP SGA1), 785907 (HBP SGA2), 945539 (HBP SGA3) and Horizon Europe grant agreement 101147319 (EBRAINS 2.0), and the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy EXC 2181/1-390900948 (Heidelberg STRUCTURES Excellence Cluster).

References

[1] C. Frenkel, J.-D. Legat, and D. Bol, "MorphIC: A 65-nm 738k-synapse/mm^2 quad-core binary-weight digital neuromorphic processor with stochastic spike-driven online learning," *IEEE transactions on biomedical circuits and systems*, vol. 13, no. 5, pp. 999–1010, 2019. DOI: 10.1109/ISCAS.2019.8702793.

[2] S. Moradi, N. Qiao, F. Stefanini, and G. Indiveri, "A scalable multicore architecture with heterogeneous memory structures for dynamic neuromorphic asynchronous processors (DYNAPs)," *IEEE Trans. Biomed. Circuits Syst.*, vol. 12, no. 1, pp. 106–122, 2018. DOI: 10.1109/TBCAS.2017.2759700.

[3] O. Richter *et al.*, "DYNAP-SE2: A scalable multi-core dynamic neuromorphic asynchronous spiking neural network processor," *Neuromorphic Computing and Engineering*, vol. 4, no. 1, p. 014 003, Jan. 2024, ISSN: 2634-4386. DOI: 10.1088/2634-4386/ad1cd7.

[4] M. Yao *et al.*, "Spike-based dynamic computing with asynchronous sensing-computing neuromorphic chip," *Nature Communications*, vol. 15, no. 1, May 2024, ISSN: 2041-1723. DOI: 10.1038/s41467-024-47811-6.

[5] J. Pei *et al.*, "Towards artificial general intelligence with hybrid tianjic chip architecture," en, *Nature*, vol. 572, no. 7767, pp. 106–111, Aug. 2019.

[6] M. Davies *et al.*, "Loihi: A neuromorphic manycore processor with on-chip learning," *IEEE Micro*, vol. 38, no. 1, pp. 82–99, 2018. DOI: 10.1109/MM.2018.112130359.

[7] W. Wan *et al.*, "A compute-in-memory chip based on resistive random-access memory," *Nature*, vol. 608, no. 7923, pp. 504–512, Aug. 2022, ISSN: 1476-4687. DOI: 10.1038/s41586-022-04992-8.

[8] A. Neckar *et al.*, "Braindrop: A mixed-signal neuromorphic architecture with a dynamical systems-based programming model," *Proceedings of the IEEE*, vol. 107, no. 1, pp. 144–164, 2018.

[9] M. Khan *et al.*, "SpiNNaker: Mapping neural networks onto a massively-parallel chip multiprocessor," in *IEEE International Joint Conference on Neural Networks (IEEE World Congress on Computational Intelligence)*, IEEE, 2008, pp. 2849–2856. DOI: 10.1109/IJCNN.2008.4634199.

[10] H. A. Gonzalez *et al.*, *SpiNNaker2: A large-scale neuromorphic system for event-based and asynchronous machine learning*, 2024. arXiv: 2401.04491 [cs.ET].

[11] P. A. Merolla *et al.*, "A million spiking-neuron integrated circuit with a scalable communication network and interface," *Science*, vol. 345, no. 6197, pp. 668–673, 2014. DOI: 10.1126/science.1254642.

[12] B. V. Benjamin *et al.*, "Neurogrid: A mixed-analog-digital multichip system for large-scale neural simulations," *Proceedings of the IEEE*, vol. 102, no. 5, pp. 699–716, 2014.

[13] J. Park, T. Yu, S. Joshi, C. Maier, and G. Cauwenberghs, "Hierarchical address event routing for reconfigurable large-scale neuromorphic systems," *IEEE Transactions on Neural Networks and Learning Systems*, vol. 28, no. 10, pp. 2408–2422, 2017. DOI: 10.1109/TNNLS.2016.2572164.

[14] C. Pehle *et al.*, "The BrainScaleS-2 accelerated neuromorphic system with hybrid plasticity," *Front. Neurosci.*, vol. 16, 2022, ISSN: 1662-453X. DOI: 10.3389/fnins.2022.795876.

[15] E. Müller *et al.*, "A scalable approach to modeling on accelerated neuromorphic hardware," *Front. Neurosci.*, vol. 16, 2022, ISSN: 1662-453X. DOI: 10.3389/fnins.2022.884128.

[16] H. Schmidt *et al.*, "From clean room to machine room: Commissioning of the first-generation BrainScaleS wafer-scale neuromorphic system," *Neuromorphic comput. eng.*, vol. 3, no. 3, p. 034 013, 2023. DOI: 10.1088/2634-4386/acf7e4.

[17] *Gerrit code review*. [Online]. Available: https://www.gerritcodereview.com (visited on 02/27/2025).

[18] *The Jenkins project*. [Online]. Available: https://www.jenkins.io (visited on 02/27/2025).

[19] E. Müller *et al.*, "A scalable approach to modeling on accelerated neuromorphic hardware," *Front. Neurosci.*, vol. 16, 2022, ISSN: 1662-453X. DOI: 10.3389/fnins.2022.884128.

[20] A. B. Yoo, M. A. Jette, and M. Grondona, "Slurm: Simple linux utility for resource management," in *Workshop on Job Scheduling Strategies for Parallel Processing*, Springer, 2003, pp. 44–60.

[21] Compaq Computer Corporation *et al.*, *Universal serial bus specification*, Revision 2.0, 2000.

[22] I. Pogorelov *et al.*, "Compact ion-trap quantum computing demonstrator," *PRX Quantum*, vol. 2, no. 2, Jun. 2021, ISSN: 2691-3399. DOI: 10.1103/prxquantum.2.020343.