Grant Number:	720270	Grant Title:	Human Brain Project SGA1

Deliverable Title:	D9.5.2 (D53.2, D47) SP9 Neuromorphic Computing Platform - Results for SGA1 Period 2
Contractual Number and type:	SGA1 D9.5.2 (D53.2 D47) Report
Dissemination Level:	PU (= PUBLIC)
Version / Date:	Resubmission of revised V1.2 on 24 July 2018 after initial acceptance in review report and subsequent rejection by EU, accepted 14 Sep 2018

Abstract:	This deliverable is the annual compound of HBP deliveries and results (outputs and outcomes) from Sub-Project SP9 – Neuromorphic Computing – here for year 2 of SGA1 = from April 2017 – March 2018. The SP9 Key Results are: • BrainScaleS-1:1st generation machine operation • BrainScaleS-2: 2nd generation prototype development • SpiNNaker-1: 1st generation machine operation • SpiNNaker-2: 2nd generation prototype development • Neuromorphic Platform software development • Computational principles: Theory for computation in neuromorphic systems
Keywords:	Neuromorphic Computing, BrainScaleS, SpiNNaker, Computational Principles, Software, Collaboratory

SP9
Neuromorphic Computing Platform

SGA1
M24 Report

Targeted users/readers	Scientific Community, General Public, Funders
Contributing Work-Package(s):	SGA1 WPs 9.1 – 9.5
Initially Planned Delivery Date:	SGA1 M24 = 31 March 2018

Authors:	SP9: SP and WP leaders
Compiling Editors:	Karlheinz MEIER, Björn KINDLER UHEI (P47)
Contributors:	SP9 WP and task leaders
SciTechCoord Review:	EPFL (P1) Martin TELEFONT
Editorial Review:	EPFL (P1) Annemieke MICHELS

Table of Contents

List of Figures

List of Tables

1. Introduction

Neuromorphic computing (NMC) transfers aspects of structure and function observed in biological neural circuits to custom designed hardware, mostly to silicon-based CMOS substrates. The potential benefits of NMC comprise energy efficiency, resilience against resource loss and, most importantly, the ability to self-organise and learn from external data.

The on-going success story of artificial neural networks (ANNs) as the main conceptual basis for state-of-the-art artificial intelligence (AI) has also generated a renewed interest in special hardware approaches that deviate substantially from the established von Neumann architecture. However, there is still a very significant performance gap between current algorithms and hardware for AI and the biological brain. The need for very large labelled data sets, the separation of training and inference, the long training times and the lack of energy efficiency are only a few examples for this gap. There is a general consensus in the AI community that a more biologically grounded approach is urgently needed to move the field and the emerging applications forward.

Here, the HBP offers a worldwide unique opportunity to transfer biological principles to novel hardware architectures. From the planning phase on, the path from biological knowledge to NMC has been a key element of research in the HBP. NMC in the HBP is based on 2 complementary and internationally recognised approaches that have been conceived around 2005 and initially supported by the FET integrated projects FACETS and BrainScaleS, as well as by the *Engineering and Physical Sciences Research Counc*il (EPSRC) in the UK.

SpiNNaker is a custom designed neuromorphic many-core chip based on the well-known ARM architecture. Special emphasis has been placed on the development of a packet-based spike communication network optimised for spike transmission that enables construction and real-time operation of very large-scale networks. The 1st generation SpiNNaker machine (2.3) currently offers 0.5 Million cores and is located in Manchester (UK).

BrainScaleS is a physical model system of neural circuits with local analog images of neurons and synapses and binary spike communication in continuous time. The analog parameters have been selected to operate the machine in an accelerated mode 10,000-fold faster than real-time. The 1st generation BrainScaleS machine (2.1) currently offers 20 wafer modules with a maximum total of 4 Million neurons and 1 Billion synapses; it is located in Heidelberg (Germany).

Organisation of work in SP9	
BrainScaleS-1 1st generation machine operation	SpiNNaker-1 1st generation machine operation
BrainScaleS-2 2nd generation prototype development	SpiNNaker-2 2nd generation prototype development
Software development Configure and operate the neuromorphic platform	
Computational principles Theory development for computation in neuromorphic systems	

Table 1: Work in SP9 is organised around the 6 key results reported in this document.

Both first generation machines are now fully integrated in the HBP infrastructure. They are remotely accessible through the Collaboratory or through direct batch job submission at the machine locations. During the reporting period, the number of users and number of jobs submitted has shown a steady increase (see Figure 19 and Figure 20), demonstrating a good uptake of this novel computing approach. As those 1st generation machines are not commercial purchases like HPC systems, substantial efforts went into maintaining and further commissioning the 1st generation. Still, today these machines represent the only remotely accessible large-scale neuromorphic computing systems in the world and are as such a unique capability offered only by the HBP.

The second major activity of the SP9 group is the design and evaluation of the 2nd generation prototype chips: SpiNNaker-2 and BrainScaleS-2. Those represent major upgrades with respect to the 1st generation systems. The novel features are enabled by the availability of very advanced CMOS process nodes as well as by the co-design process in the HBP. SpiNNaker-2 will have a tenfold increase of computational capability with a constant power budget. BrainScaleS-2 will offer on-chip hybrid plasticity and structured analog neurons with active dendrites. These developments have the potential to secure Europe's excellent position in the field of neuromorphic computing. A joint presentation with Intel's new Loihi chip at the international conference on neuro-inspired -computational-elements (NICE) in March 2018 has clearly confirmed this. The next step of realising large-scale systems based on the 2nd generation chips during the early SGA3 period will be very important to achieve this goal.

SP9 also delivers software support for the NMC systems ranging from low-level, hardware specific software to the high-level integration into the HBP Platform infrastructure which will eventually enable the NMC to be part of the HBP joint Platform. The SP9 software represents a unique ecosystem for neuromorphic computing. The user-accessible software is open source and managed in public GitHub repositories.

Finally, but of crucial importance is the very close collaboration with theory which has successfully enabled the co-design process of 2nd generation neuromorphic hardware driven by theoretical principles of brain information processing. The most important principle is learning in neural systems but also other ideas have

contributed to the design on the 2nd generation NMC systems. Theory enters SP9 through an internal Work Package as well as through SP4 and CDP5.

In this context, we quote the corresponding argument given in the CDP5 deliverable: "Besides brain science, another branch of science is interested in uncovering the principles of learning: artificial intelligence (AI). However, the field of AI has a different scope, as it is not confined by biological boundaries. Nevertheless, as Nature writes in its editorial from Feb. 8, 2018, on the Hardware Upgrade "Artificial intelligence is driving the next wave of innovations in the semiconductor industry". The editorial concludes with "We welcome papers that will enable computing architectures beyond von Neumann, such as components for neuromorphic chips and in memory processing. Scientists across many fields are waiting for the result …." By working on unravelling the learning ability of the mammalian brain, CDP5 implicitly addresses the advancement of AI and its desire to build fast, energy-efficient, massively parallel hardware. A key partner is therefore SP9, which is devoted to hardware implementations of learning circuitry, with the multi-core SpiNNaker system in Manchester and the physical-model BrainScaleS Platform in Heidelberg. The core of CDP5 connects the computational theories on learning and plasticity, developed in SP4 (WP4.3) and in part in SP9 itself (WP9.4), to the neuromorphic hardware developed in SP9."

All key results presented in this report are aligned with the work organisation at shown in Table 1.

2. Key Results

In this chapter we present the key results obtained by SP9 during the reporting period, as well as their impacts, mostly through publications. Publication lists are just providing important examples; complete lists are available in the periodic reports. For each Key Result we also provide a short description of the challenges encountered during the work and an outlook to future work. As general statement relating to all Key Results one can say that the work on the SGA1 funding period proceeded very well and that all challenges have been met with excellent success.

2.1 BrainScaleS-1 (1st generation machine)

2.1.1 Results

During the reporting period, the BrainScaleS-1 machine (Figure 1 and Figure 2) was operated at the Heidelberg location in a 24/7 mode. Job submission can be performed through a local batch queue (SLURM) and through the HBP Collaboratory, also labelled NMPI (Neuromorphic Platform interface).

Figure 1: Wafer module of BrainScaleS-2 machine

Figure 2: BrainScaleS machine. 20 wafer modules controlled by local compute cluster

As the machine is still in a commissioning phase, subsets of the complete system have been available to users at any given time. Part of the on-going commissioning work during the reporting period went into investigating a remaining connectivity problem between the silicon wafer and the printed circuit board. Together with the Fraunhofer IZM group the problem was traced back to local damages occurring during the wafer post-processing in some of the wafers. A new batch of wafers has been submitted for post-processing within the reporting period and their installation is expected to take place in April 2018.

Figure 3 shows the evolution of the Platform use. Batch submission (i.e. the non-NMPI user jobs) dominates by far the statistics, because of its ease of use, speed and the advanced capabilities to run scripted parameter scans. Continuous integration through the Jenkins tool ensures quality control and an operational system with regular updates in hardware and software.

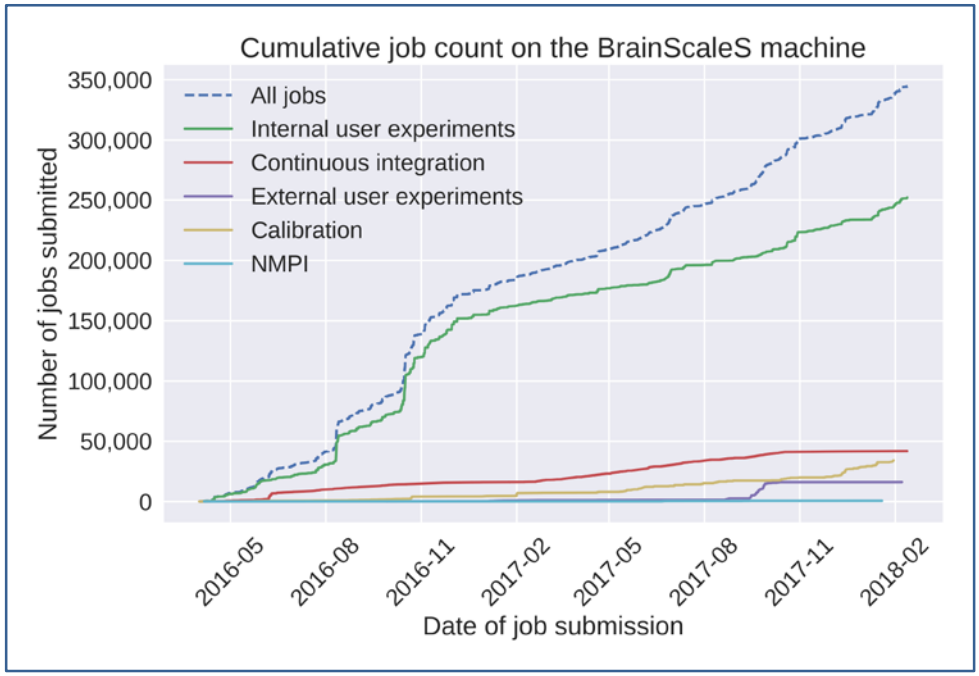

Figure 3: Cumulative job submission numbers to the BrainScaleS machine from March 2015 to February 2018

Calibration and understanding of the system took substantial efforts by SP9 scientists. This part of the work was mostly dedicated to a detailed calibration of analog cell parameters and setting-up of the on-wafer routing. Figure 4 shows the example of a routed chain of spiking neurons on a hardware wafer crossing boundaries between many chips and reticles.

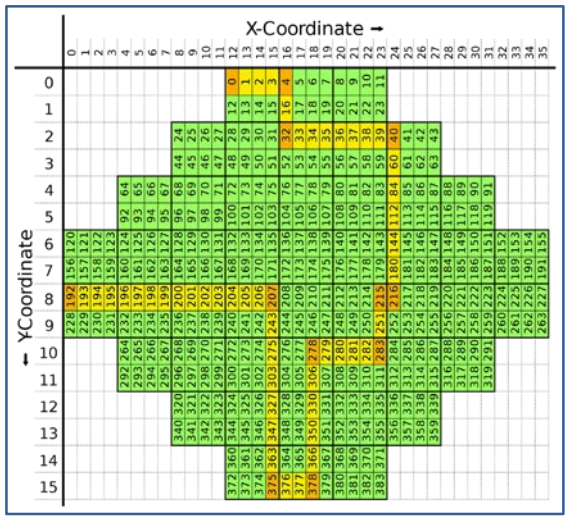

Figure 4: Example of on-wafer routing on the BrainScaleS system

Several state-of-the-art experiments have been carried out with the BrainScaleS machine. Here, we report on 3 fundamentally different approaches, which demonstrate the universality of the system, and interesting features of spike-based, accelerated physical model implementations of neuromorphic computing.

A physical model spiking deep networks solving the MNIST machine learning benchmark

Here, it was demonstrated how iterative training of a hardware-emulated network can compensate for anomalies induced by the analog substrate. We first convert a deep neural network trained in software to a spiking network on the BrainScaleS

wafer-scale neuromorphic system, thereby enabling an acceleration factor of 10,000 compared to the biological time domain. This mapping is followed by the in-the-loop training, where in each training step, the network activity is first recorded in hardware and then used to compute the parameter updates in the software via backpropagation. An essential finding is that the parameter updates do not have to be precise, but only need to approximately follow the correct gradient, which simplifies the computation of updates. Using this approach, after only several tens of iterations, the spiking network shows an accuracy close to the ideal software-emulated prototype. These techniques show that deep spiking networks emulated on analog neuromorphic devices can attain a good computational performance despite the inherent variations of the analog substrate (text taken from abstract of Ref 1 Schmitt *et al*.).

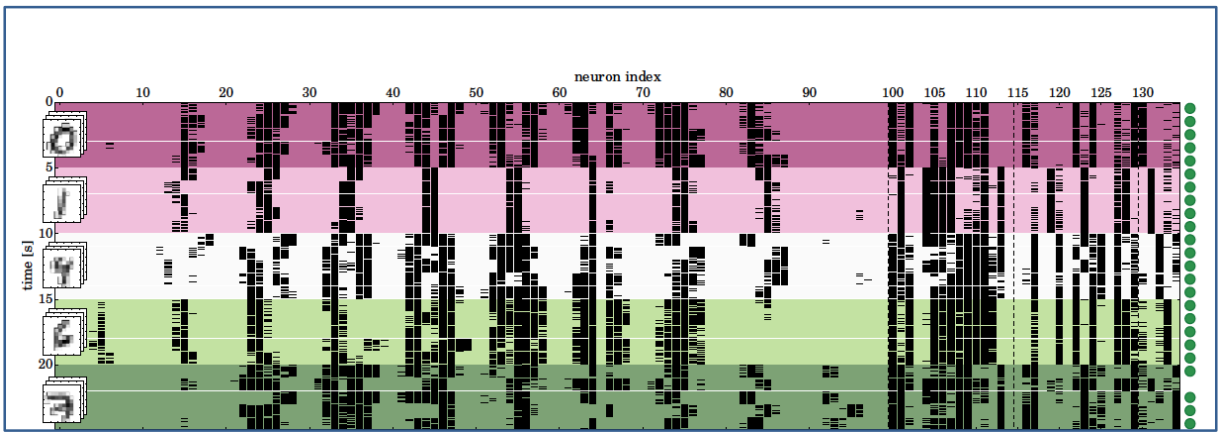

Figure 5: Spike pattern in network layers (horizontal) vs. pattern input (vertical). Green dots in the label layer (right) show correct results.

Figure 5 shows the recorded spike activity after presentation of a sequence of hand-written numbers to the trained network of the BrainScaleS system.

Neural sampling with spiking neurons.

During the SGA1 funding period, SP9 members have pioneered a formal theory of neural sampling with spiking LIF neurons (Ref 2). Neural sampling is based on stochastic neurons that require noise input from either internal or external sources. Sampling neurons can either generate patterns from learned data, or recognise patterns from incomplete or noisy inputs. During the reporting period, neural sampling has been implemented for the first time on the BrainScaleS machine, using the activity of an adjacent, on-wafer functional network as a noise source (Figure 6).

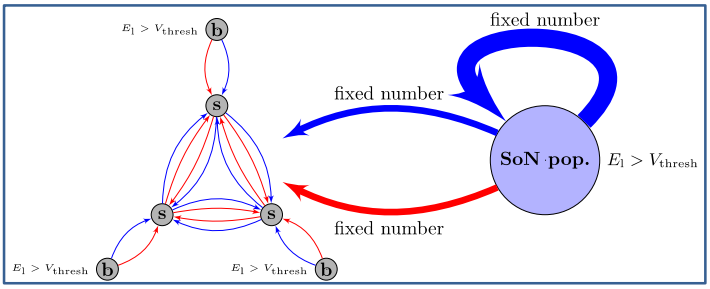

Figure 6: Interconnected (left) and noise network (right) on the BrainScaleS machine

This “sea-of-noise-concept” was originally developed in the BrainScaleS project (Ref 3), it represents a promising concept to produce *in vivo*-like conditions for information processing with neuromorphic hardware without bandwidth penalties.

Constraint satisfaction problems

Another example of a recent category of applications running on the BrainScaleS system are constraint satisfaction problems (CSF), i.e. problems in which a whole set of constraints must be satisfied for a given state. A well-known benchmark for a CSF is the Sudoku problem which was applied as a first use case with a 4x4 cell configuration. For a maximum number of 12 free cells, the percentage of correctly solved Sudoku set-ups could be increased from below 20% to 50% for 100 randomly drawn Sudoku set-ups. At 8 free cells, the solving accuracy with training is at 100% while it declines below 70% without training (from Ref 4).

2.1.2 Impact

BrainScaleS-1 is the worldwide first and only large-scale, remotely accessible physical model neuromorphic computing system. The following references point to some key publications.

Ref 1: Sebastian Schmitt, Johann Klähn, Guillaume Bellec, Andreas Grübl, Maurice Güttler, Andreas Hartel, Stephan Hartmann, Dan Husmann,Kai Husmann, Sebastian Jeltsch, Vitali Karasenko, Mitja Kleider, Christoph Koke, Alexander Kononov, Christian Mauch, Eric Müller, Paul Müller, Johannes Partzsch, Mihai A. Petrovici, Stefan Schiefer, Stefan Scholze, Vasilis Thanasoulis, Bernhard Vogginger, Robert Legenstein, Wolfgang Maass, Christian Mayr, René Schüffny, Johannes Schemmel, Karlheinz Meier (2017, May). Neuromorphic hardware in the loop: Training a deep spiking network on the brainscales wafer-scale system. In Neural Networks (IJCNN), 2017 International Joint Conference on (pp. 2227-2234). IEEE. https://doi.org/10.1109/IJCNN.2017.7966125

Ref 2: Petrovici, M. A., Bill, J., Bytschok, I., Schemmel, J., & Meier, K. (2016). Stochastic inference with spiking neurons in the high-conductance state. *Physical Review E*, *94*(4), 042312. https://doi.org/10.1103/PhysRevE.94.042312

Ref 3: Jordan, J., Petrovici, M. A., Breitwieser, O., Schemmel, J., Meier, K., Diesmann, M., & Tetzlaff, T. (2017). Stochastic neural computation without noise. *arXiv preprint arXiv:1710.04931*.

Ref 4: Alexander Kugele, Solving the Constraint Satisfaction Problem Sudoku on Neuromorphic Hardware, MSc thesis in Physics, Heidelberg University, 2018.

2.1.3 Challenges

The main challenge of this Key Result was to maintain and further develop a very large-scale hardware system in a stable and reliable mode of operation. This includes a continuous monitoring of hardware and software, a support of established and new users, and a continuous upgrade and replacement of hardware components like improved and newly reprocessed wafers and analog readout boards. The SP9 large-scale systems are sometimes compared to commercial supercomputing systems which represent investments of several 100 Million Euros and are simply purchases from major industry players. Investments into SP9 large-scale systems are about 2 orders of magnitude less in funding and in-house constructions. Today, they

represent worldwide unique facilities with very remarkable results as described in the Key Results in this report. The challenge was met by highly competent and motivated scientists and engineers in SP9, which represent a powerful workforce that today represents one of the most important assets of the HBP.

2.1.4 Future Work

The BrainScaleS large-scale system will now be continuousy operated until the next generation machine is available with the size as contracted in the FPA. Efforts will now be focused on user support and software development. Hardware replacements and minor upgrades will still be carried out wherever neccessry and useful for the user requirements. A comprehensive list of large-scale experiments is currently being prepared. It is expected that in the year 2020 (end of SGA2) the existing machine will be moved into the new European Institute for Neuromorphci Computing (EINC) currently under construction at the University of Heidelberg.

2.1.5 Component Dependencies

Component ID	Component Name	HBP Internal	Comment
1	BrainScaleS-1 Neuromorphic Computing System	No	This Key Result

2.2 BrainScaleS-2 (2nd generation prototype development)

2.2.1 Results

The 2nd generation BrainScaleS-2 neuromorphic chip was developed within the HBP, in close collaboration with theoretical neuroscientists; it is an example of successful co-design.

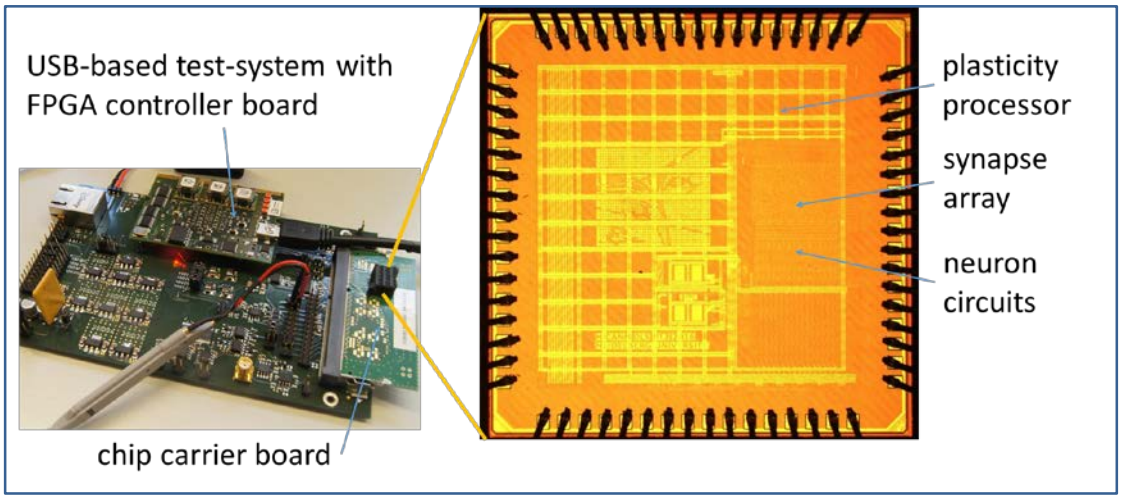

Figure 7: Test-system (left) and chip photograph (right) of a BrainScaleS-2 prototype test chip

A major outcome of the reporting period is the development of the first full scale 2nd generation BrainScaleS-2 ASIC (application specific integrated circuit) that is submitted for production. This first full-size chip is based on results obtained with previous prototype test chips (see Figure 7 for an example) which each tested different aspects of the final BrainScaleS 2nd generation neuromorphic hardware system. Most notably, during the reporting period, the first prototype that implemented non-linear dendrites and multi-compartmental modelling capabilities, was successfully tested (Figure 8).

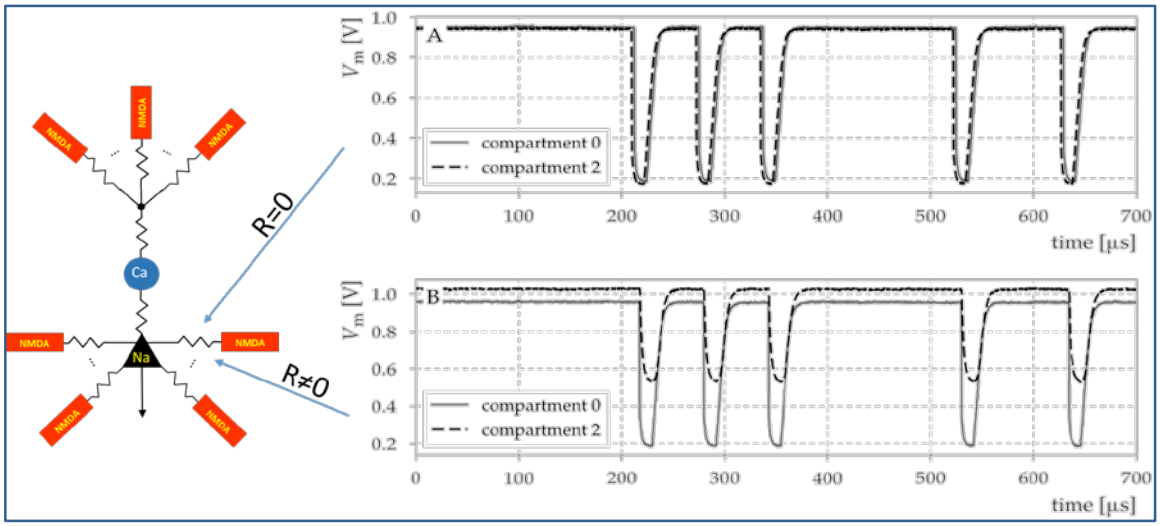

Figure 8: Measurements of multi-compartment operation. Top trace: large inter-compartmental conductance (i.c.c.), Bottom: small i.c.c.

The main purpose of the first full-size 2nd generation chip (internal name HICANN-X) will not primarily be for testing, but rather to serve already as the first new experimental platform of the BrainScaleS-2 system available to external users.

The basic neural network layout of the HICANN-X chip is similar to the BrainScaleS 1st generation ASIC, to ease the migration of software and models. It is manufactured in a 65nm low-power and low-leakage CMOS technology. 512 neuron compartment circuits are located in a horizontal orientation in the centre of the analog network core while two synapse arrays with 256x256 synapses form the top and bottom halves of the structure. A routing logic allows flexible programming of network topologies which can be extended across several chips by using the built-in eight 2 Gbit/s full-duplex links. The links are compatible to the FPGA boards of the existing BrainScaleS system, which saves considerable resources. The integrated fast analog-to-digital converter allows the readout of membrane voltages.

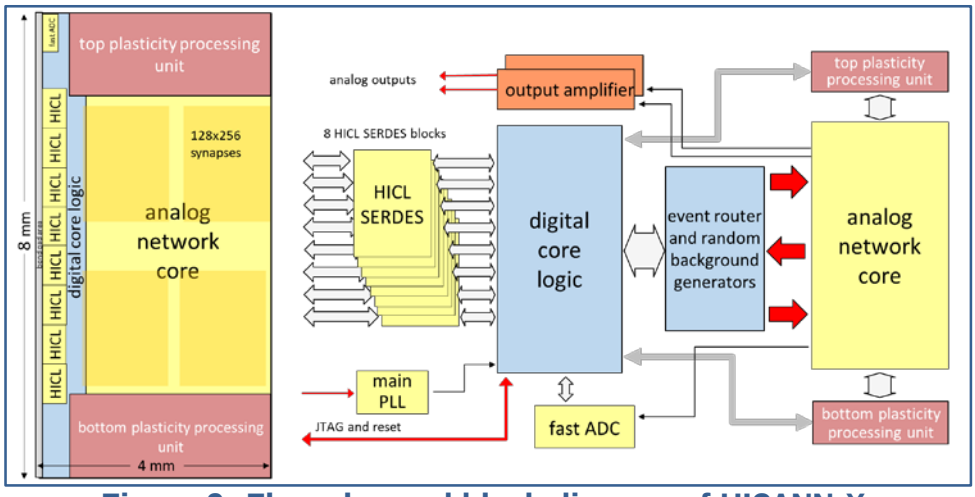

Figure 9: Floorplan and block diagram of HICANN-X.
Colours indicate technologies: Yellow - mixed signal, blue or red – standard cell, orange - analog

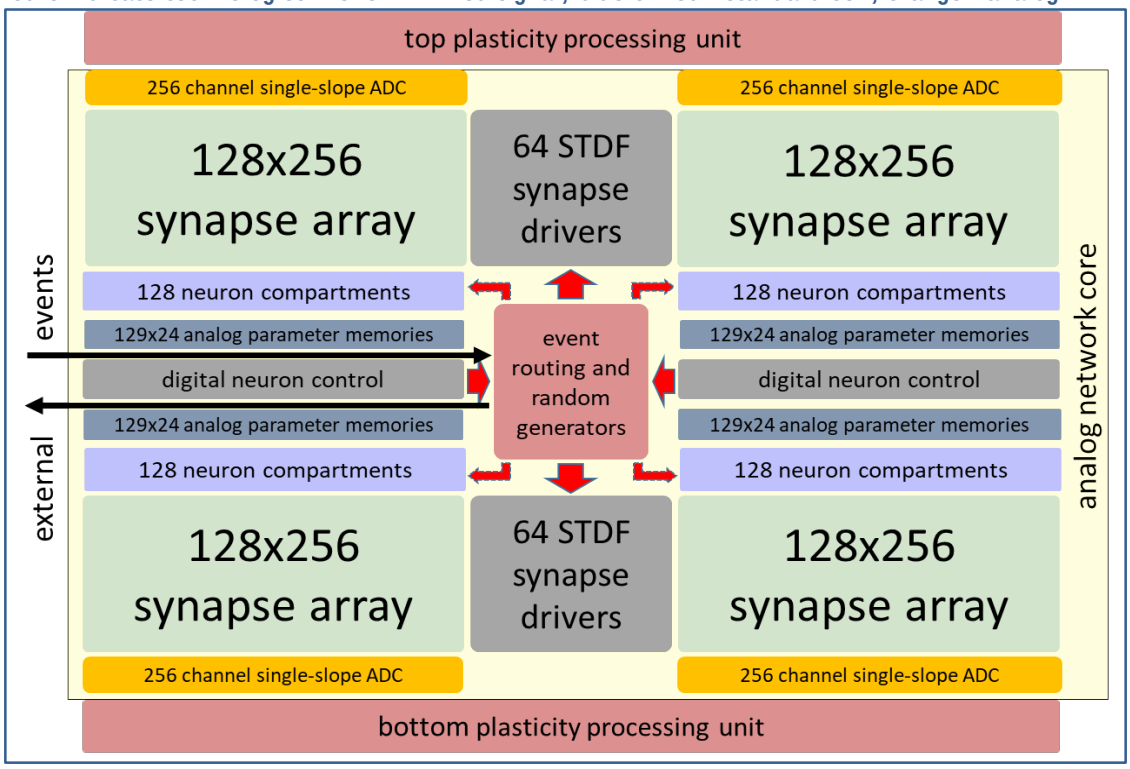

Figure 10: Block diagram of analog network core of HICANN-X.
Digital event routing and random generators are physically located in the centre of the core.

HICANN-X includes a full implementation of hybrid plasticity. Two build-in SIMD plasticity processing units, each adjacent to one of the synapse arrays, allow the execution of software-defined plasticity algorithms in the models' real-time, which is in BrainScaleS 2 1000 times faster than its biological counterpart. Besides the spike time information, it is possible to include membrane voltages, average firing rates and neuro-modulators in the plasticity calculations. The multi-compartment extensions are also fully supported by the plasticity processing units. Structural plasticity concepts, like synaptic sampling, are part of the planned usage of the hybrid plasticity technology as well.

The neuron model contains adaptation and exponential terms according to the Adaptive-Integrate-and-Fire model, similar to BrainScaleS 1. In addition, a digital refractory circuit enhances precision and provides a very large dynamic range, suitable for the emulation of effects like NMDA plateau potentials.

2.2.2 Impact

The 2nd generation BrainScaleS chip represents the first neuromorphic chip worldwide with an embedded processor for local plasticity and learning, as well as the first analog CMOS implementation of structured neurons with active dendrites. The first prototype chips are already remotely accessible and well used (Figure 11). The large increase of usage is related to the learning-to-learn experiments which make heavy use of fast parameter scans.

Ref 5: Schemmel, J., Kriener, L., Müller, P., & Meier, K. (2017, May). An accelerated analog neuromorphic hardware system emulating NMDA-and calcium-based non-linear dendrites. In *Neural Networks (IJCNN), 2017 International Joint Conference on* (pp. 2217-2226). IEEE, https://doi.org/10.1109/IJCNN.2017.7966124

Ref 6: Friedmann, S., Schemmel, J., Grübl, A., Hartel, A., Hock, M., & Meier, K. (2017). Demonstrating hybrid learning in a flexible neuromorphic hardware system. *IEEE transactions on biomedical circuits and systems*, *11*(1), 128-142, http://doi.org/10.1109/TBCAS.2016.2579164

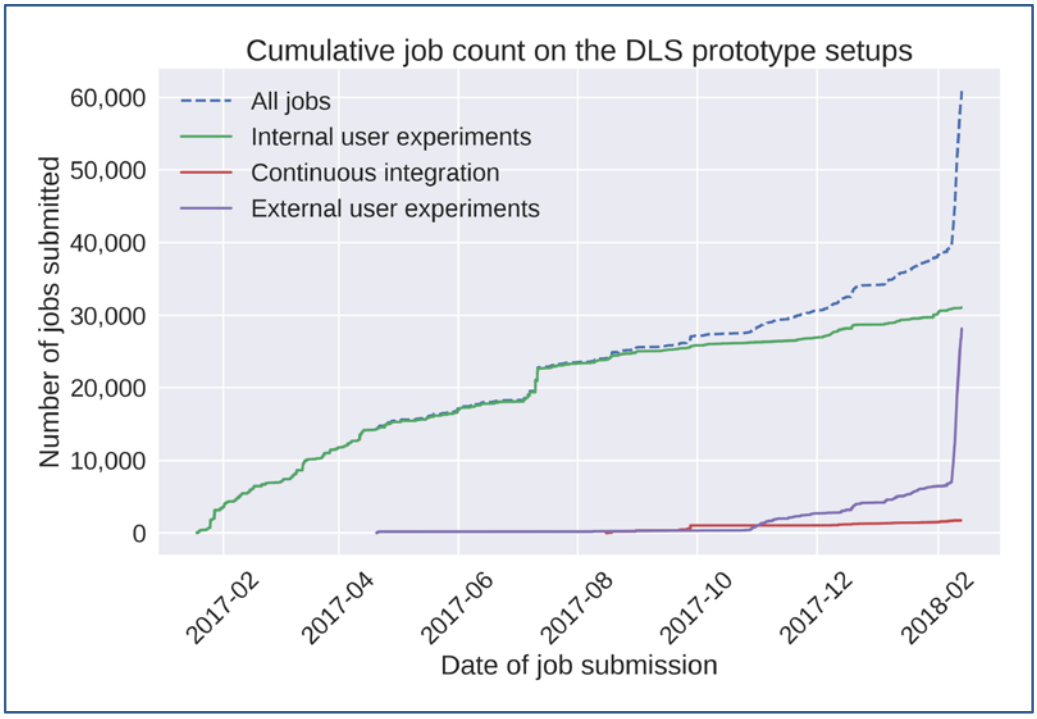

Figure 11: Cumulative job count on the DLS prototype setups

2.2.3 Challenges

There have been 2 main challenges for the second generation BrainScaleS system.

Scientifically the system implements very recent state-of-the-art results obtained by the neuroscience groups in the HBP. This includes novel principles of dendritic integration, the multicompartment structure of complex neurons and several principles for plasticity, learning and development. The very successful transfer of these results into neuromorphic hardware was mostly enabled by a very good interdisciplinary communication. Key enablers for this communication have been CDP5, the EITN in Paris, the Fürberg workshops organized by SP9 and first and foremost a highly motivated group of Post-Docs, PhD and Ma students, who enthusiastically worked very hard towards the goal of producing a truly unique and novel neuromorphic chip.

In terms of project planning the very short 2 year SGA funding periods without any flexibility of budget transfers pose a major problem and risk for chip design and production, which has to match the rigorous schedules of commercial manufacturing plants. SP9 managed to handle this problem by in-kind contributions and shifting funds between staff and consumables, but the situation is clearly not satisfactory and substantially worse than in other large scientific projects like elementary particle physics with similar challenges but a more flexible and adaptive funding scheme.

2.2.4 Future Work

The future work of the second generation chips is specified and described in detail in the SGA2 proposal. All features are defined and tested. A full size prototype has been submitted for production. The chip wil be evaluated during SGA2 and the design step for full size wafer production will be carried out, so that construction of the

next generation full size sytems can proceed in SGA3 as planned from the beginning of the HBP and described and contracted in the FPA.

2.2.5 Component Dependencies

Component ID	Component Name	HBP Internal	Comment
457	BrainScaleS 2 Neuromorphic Computing System	No	This Key Result
1	BrainScaleS-1 Neuromorphic Computing System	No	BrainScaleS-1 version – the starting point for the development of BrainScaleS-2

2.3 SpiNNaker-1 (1st generation machine)

2.3.1 Results

SpiNNaker-1 is a large-scale many-core Neuromorphic Platform incorporating 500,000 ARM processor cores (Figure 12). It is built using the SpiNNaker-1 chip, each of which incorporates 18 ARM968 processor cores and a novel multicast packet routing engine optimised for conveying very large numbers of very small data packets, where typically each spike in a spiking neural network simulation is conveyed on one packet. The SpiNNaker chips are packaged with memory chips, assembled onto circuit boards, and mounted in 19" rack cabinets where each cabinet contains 100,000 processors on 120 circuit boards. Neuron and synapse models are implemented in software on SpiNNaker, offering considerable flexibility in the choice of equations and learning rules. The machine typically executes spiking networks of any size up to its maximum capacity (around 100 million neurons) in biological real time.

The 500,000-core HBP Platform is shown in the photo below (Figure 12). Five racks hold the SpiNNaker boards, and the sixth rack holds the server.

Figure 12: The 0.5M core SpiNNaker-1 machine located in Manchester

The SpiNNaker-1 machine was designed and developed using UK (EPSRC) funding. HBP supports ongoing software development and support, hardware maintenance, and the building of further small machines for dissemination and wider use. Over 90 SpiNNaker-1 boards are now in the hands of research groups all around the world (Figure 13).

Figure 13: Global distribution of small-scale SpiNNaker machines

The 500,000-core SpiNNaker-1 machine at Manchester is available through the HBP Collaboratory since the start of HBP SGA1 and has attracted a range of users over the subsequent period. Jobs submitted through the HBP Collaboratory can be written as PyNN 0.7, 0.8 or 0.9 spiking neural networks. They are compiled and run on the machine using a software tool flow running on a server at Manchester:

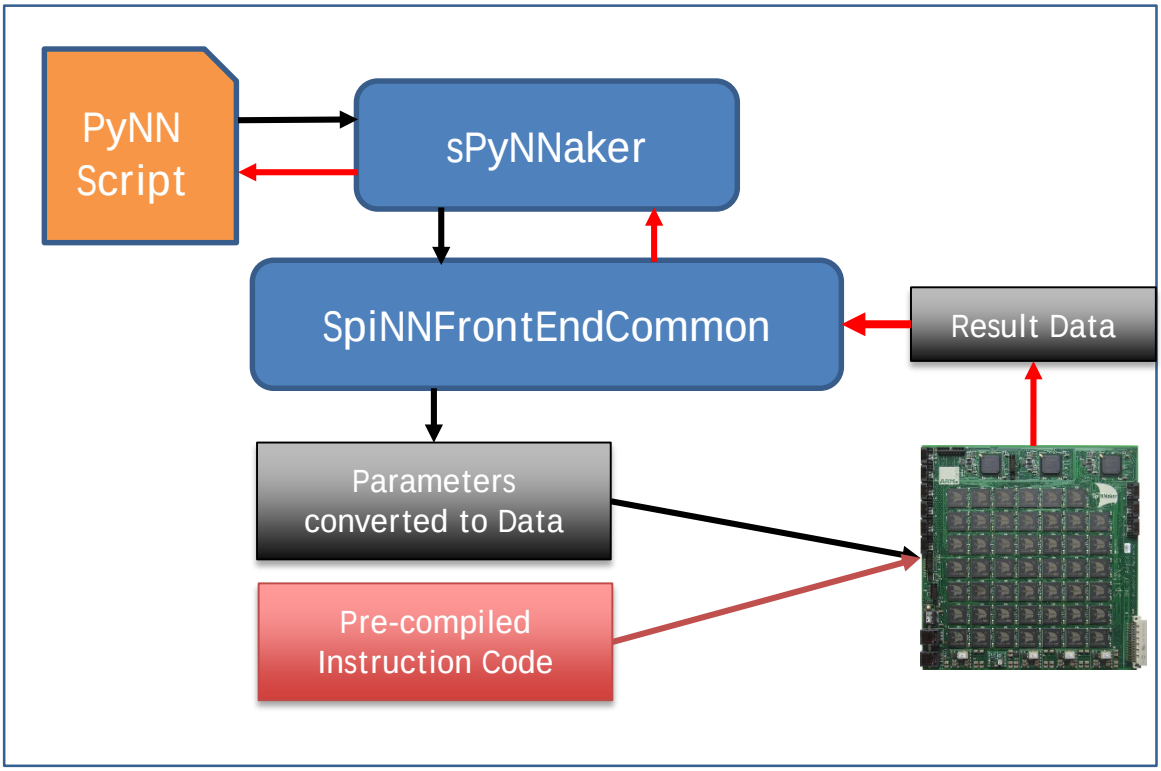

Figure 14: SpiNNaker software framework

As at January 11, 2018 there were 64 remote HBP portal SpiNNaker-1 users and 1,306 jobs had been run through the HBP Collaboratory. The software tools are hosted on GitHub, and users are supported through a SpiNNaker Users Mailing list (which attracted 44 posts in December 2017). Open issues are logged on GitHub and addressed according to priorities that are reviewed regularly.

Other status highlights include:

- SpiNNaker-1 uses fixed-point arithmetic. This delivers comparable numerical accuracy to that achieved with double-precision floating-point arithmetic used in simulations, as shown in a comparison of the results of the machine running PyNN with those from a conventional HPC running NEST on a cortical micro column model developed at Jülich.
- The current SpiNNaker-1 machine has the capability of running stochastic spiking neural networks to solve constraint satisfaction problems [Ref 7].
- An alternative way to map problems onto SpiNNaker-1 has been developed [Ref 8].
- Hardware optimisations have been investigated [Ref 9].
- A spiking vision benchmarking suite is available [Ref 10].
- Biological models [Ref 11] and learning algorithms [Ref 12] have been developed on SpiNNaker-1.
- A review of large-scale neuromorphic platforms is available [Ref 13].

2.3.2 Impact

SpiNNaker-1 is the worldwide first and only large-scale, remotely accessible many-core neuromorphic computing system. The following references point to some key publications.

Ref 7: G.A. Fonseca Guerra, S.B. Furber, "Using Stochastic Spiking Neural Networks on SpiNNaker to Solve Constraint Satisfaction Problems", Frontiers in Neuroscience – Neuromorphic Engineering, Dec 2017. https://doi.org/10.3389/fnins.2017.00714

Ref 8: Knight, James C. and Furber, Steve B., "Synapse-Centric Mapping of Cortical Models to the SpiNNaker Neuromorphic Architecture", Frontiers in Neuroscience 10, p.420, 14 September 2016. https://doi.org/10.3389/fnins.2016.00420

Ref 9: A. Yousefzadeh, M. Jabłon ▫ski, T. Iakymchuk, A. Linares-Barranco, A. Rosado, L.A. Plana, S. Temple, T. Serrano-Gotarredona, S.B. Furber and B Linares-Barranco, "On Multiple AER Handshaking Channels Over High-Speed Bit-Serial Bidirectional LVDS Links With Flow-Control and Clock-Correction on Commercial FPGAs for Scalable Neuromorphic Systems", in IEEE Trans Biomedical Circuits and Systems 11(5), pp. 1133-1147, Oct. 2017.DOI: 10.1109/TBCAS.2017.2717341, https://doi.org/10.1109/TBCAS.2017.2717341

Ref 10: Qian Liu, Garibaldi Pineda-García, Evangelos Stromatias, Teresa Serrano-Gotarredona, Steve B. Furber, "Benchmarking Spike-Based Visual Recognition: A Dataset and Evaluation", Frontiers in Neuroscience, Vol.10, 2016, https://doi.org/10.3389/fnins.2016.00496

Ref 11: B. Sen-Bhattacharya, T. Serrano-Gotarredona, L. Balassa, A. Bhattacharya, A. Stokes, A. Rowley, I. Sugiato, S. Furber, "A spiking neural network model of the Lateral Geniculate Nucleus on the SpiNNaker machine", Frontiers in Neuroscience, Aug 2017. https://doi.org/10.3389/fnins.2017.00454

Ref 12: James C. Knight, Philip J. Tully, Bernhard A. Kaplan, Anders Lansner and Steve B. Furber, "Large-Scale Simulations of Plastic Neural Networks on

Neuromorphic Hardware", Frontiers in Neuroanatomy 10(37), 7 April 2016. https://doi.org/10.3389/fnana.2016.00037

Ref 13: Steve Furber, "Large-scale Neuromorphic Computing Systems", Journal of Neural Engineering 13(5), 2016, pp. 1-14. https://doi.org/10.1088/1741-2560/13/5/051001

2.3.3 Challenges

The challenges for the large-scale SpiNNaker systems are similar to those for the BrainScaleS systems as both are lab-made large-scale hardware facilities built from a very limited budget. As SpiNNaker is based on ARM based many-core chip directly mounted on printed-circuits-boards, there is less commissioning work required compared to the BrainScaleS wafer-based system. Challenges were mostly on the software side, in particular the challenge of mapping very large-scale networks which exploit a substantial fraction of the machines computational capability. As the results presented in this report show, this challenge has been met. Again, the main reason is the ability of SP9 groups to attract highly motivated staff that recognises neuromorphic computing as a rewarding and forward looking technology for the future of brain-science and a more biologically driven machine learning.

2.3.4 Future Work

Also here, the future work for the large-scale SpiNNaker system is very similar compared to BrainScaleS. The main focus is usability, software support and the running of large-scale computational tasks suitable for a comparison with conventional computers. The machine will also be upgraded from 0.5 Million to 1 Million ARM cores.

2.3.5 Component Dependencies

Component ID	Component Name	HBP Internal	Comment
2	SpiNNaker Neuromorphic Computing System	No	This Key Result

2.4 SpiNNaker-2 (2nd generation prototype development)

2.4.1 Results

As small-scale prototype of the 2nd generation SpiNNaker many-core system, a test chip (Santos) (Figure 15) has been implemented in 28nm SLP CMOS technology. It contains 4 processing elements (PE) based on ARM M4F processors. The PEs are connected via a network-on-chip. A DRAM interface is available of off-chip memory, e.g. used for synaptic weight storage. A demonstrator PCB hosting 4 chip modules has been developed and is used for evaluation.

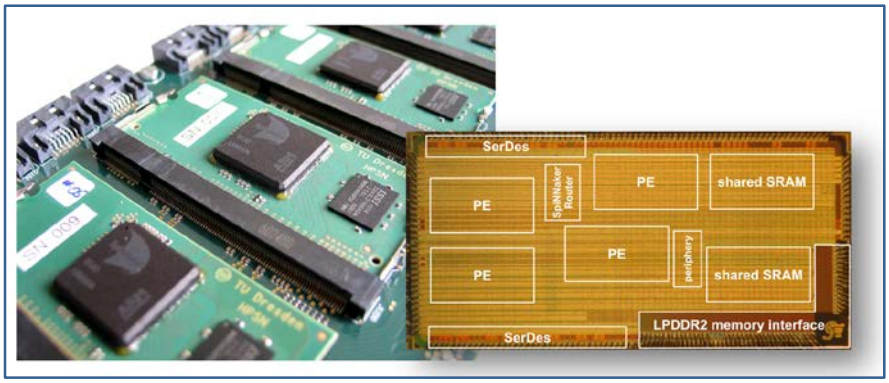

Figure 15: Santos prototype PCB and chip photo.

The chip contains novel hardware features to enhance the efficiency of neuromorphic real-time simulations on the next generation SpiNNaker system. These new features include:

- Neuromorphic Power Management: Fast switching of supply voltage level (from 0.70V to 1.0V) and processor clock frequency (from 100MHz to 500MHz) within <100ns for fast adaption of the compute performance to the temporal workload of the experiment. The performance level can be adjusted individually per core and simulation time step (e.g. 1ms), leading to up to 50% reduction of energy per synaptic event (down to 0.83nJ) and up to 85% reduction of total power consumption.
- Exponential Function Acceleration: Hardware unit offload, the exp() function calculation from the processor, beneficial for neuromorphic compute problems such as STDP or BCPNN. Achieving 250Mexp/s throughput at 0.44nJ/exp for nominal supply (1.0V), or 0.21nJ/exp at 0.7V supply and 77Mexp/s, demonstrating a throughput multiplication of almost 50 and 98% energy reduction at 2% area overhead.
- True Random Number Generation: Circuits to extract true random numbers from the clock generator jitter with minimum energy (1.55 pJ/Bit) and chip area overhead for the entropy sources, post processing units and additional PRNG accelerators. Random numbers pass NIST and DIEHARD benchmarks.

The new hardware features are evaluated by benchmarks: Examples include:

- Synfire Chain Network and Bursting Network Simulations: Using neuromorphic power management, showing the benefit of fast performance level scaling, only around 1% of the simulations time steps need to be executed on the highest performance level, 90% are executed at the lowest level for maximum energy efficiency (Figure 16).

- Reward-Based-Model (with TU Graz): Implementation of a reward-based synaptic sampling network using the exp() function and the random number generators. Showing up to 50% reduction of the computation time per stochastic synapse update (Figure 17).

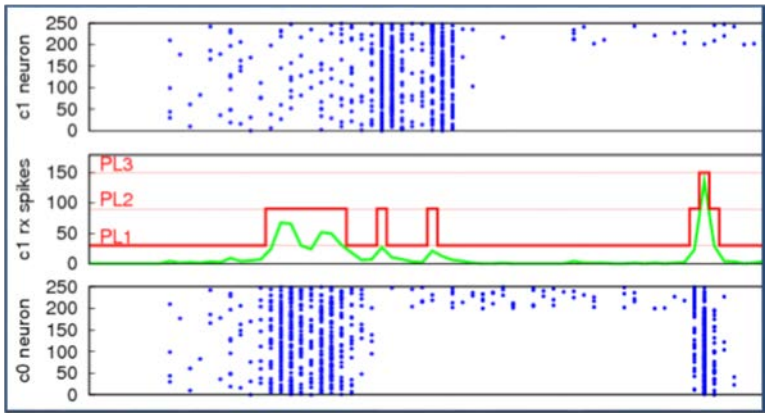

Figure 16: Synfire chain simulation with dynamic power management, processors are clocked to higher performance level (PL) if high computational load is required.

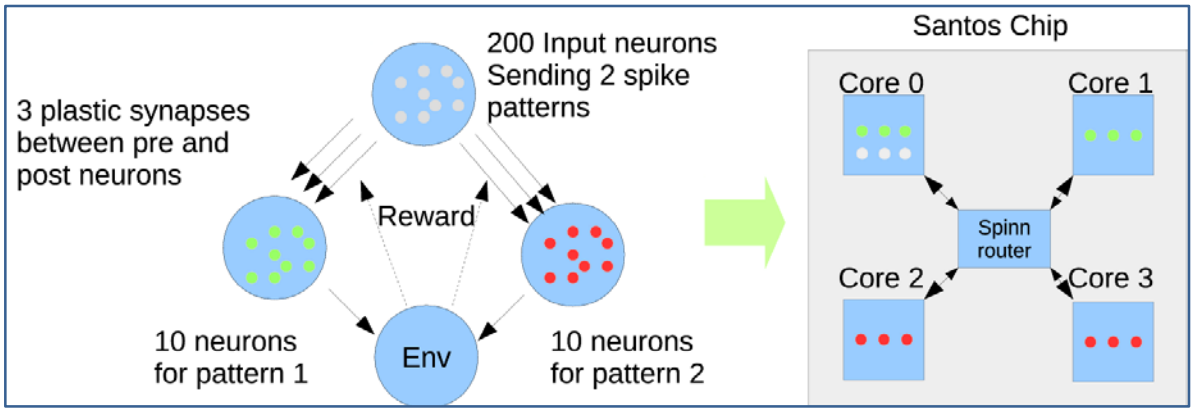

Figure 17: Reward-based synaptic sampling network mapped to the SpiNNaker-2 prototype.

Additional Benchmarks under evaluation include the simulation of BCPNN networks (using the exp() unit), Dynamic Vision Sensor (DVS) data processing (using power management) and applications for Spike detection and sorting (real-time biological data processing).

2.4.2 Impact

Ref 14 Sebastian Höppner, Yexin Yan, Bernhard Vogginger, Andreas Dixius, Johannes Partzsch, Felix Neumärker, Stephan Hartmann, Stefan Schiefer, Stefan Scholze, Georg Ellguth, Love Cederstroem, Matthias Eberlein, Christian Mayr, Steve Temple, Luis Plana, Jim Garside, Simon Davison, David R Lester, Steve Furber: Dynamic voltage and frequency scaling for neuromorphic many-core systems, 2017 IEEE International Symposium on Circuits and Systems (ISCAS), https://doi.org/10.1109/ISCAS.2017.8050656

Ref 15 Johannes Partzsch, Sebastian Höppner, Matthias Eberlein, Rene Schüffny, Christian Mayr, David R Lester, Steve Furber: A fixed-point exponential function accelerator for a neuromorphic many-core system, 2017 IEEE International

Symposium on Circuits and Systems (ISCAS), https://doi.org/10.1109/ISCAS.2017.8050528

Ref 16 Felix Neumarker, Sebastian Höppner, Andreas Dixius, Christian Mayr: True random number generation from bang-bang ADPLL jitter, 2016 IEEE Nordic Circuits and Systems Conference (NORCAS), https://doi.org/10.1109/NORCHIP.2016.7792875

2.4.3 Challenges

The next generation SpiNNaker system implements digital neuromorphic computing in an almost state-of-the-art 28nm process node. It has several novel features with as main goal to improve the efficiency of the system in terms of computational capability for a given amount of energy.

Also here, the challenges are twofold. Having access to and making efficient use of such state-of-the-art technology requires an excellent understanding of the toolsets for simulation, design and layout. The Manchester and Dresden groups are of exceptional quality and do have the required knowledge to design a test system and bring it into operation during the very short SGA1 funding period. This is the main outcome of this Key Result.

However, like in the case of BrainScaleS also here, the very short SGA periods are a substantial problem for efficient work. Workarounds have been found by shifting tasks and budget allocations with in the SGA but the situation is far from ideal.

2.4.4 Future Work

The future work of the second generation SpiNNaker chips is specified and described in detail in the SGA2 proposal. All features are defined and tested. A full size prototype has been submitted for production. The chip will be evaluated during SGA2 and the design step for full size system production will be carried out, so that construction of the next generation full size systems can proceed in SGA3 as planned from the beginning of the HBP and described and contracted in the FPA.

2.4.5 Component Dependencies

Component ID	Component Name	HBP Internal	Comment
467	SpiNNaker small-scale NM-MC System	No	SpiNNaker-2 evaluation
2	SpiNNaker Neuromorphic Computing System	No	SpiNNaker-1 – the basis for the SpiNNaker-2 development

2.5 Neuromorphic Computing Platform (Software development)

2.5.1 Results

The Neuromorphic Computing Platform offers access to two large-scale, multi-user facilities, the BrainScaleS 20-wafer system in Heidelberg and the SpiNNaker 500,000 core machine in Manchester. Recently, also 2nd generation prototype systems have

been added. Our goal is to make the use of these machines as easy as possible, for neuroscientists, machine-learning researchers, and students. This implies providing i) remote access, through both a graphical interface (for less technical users, and for result visualisation) and a scriptable interface (for power-users), and ii) an easy-to-learn, flexible and powerful programming interface (PyNN).

A further requirement, given that neuromorphic simulations in general form only part of a scientific workflow, is that data produced on the Neuromorphic Platform can be transferred to traditional computing facilities, like the user's local computer, Jupyter notebooks in the Collaboratory, or the HBP HPAC Platform. An example of a job output from a BrainScaleS experiment, accessed in the HBP Collaboratory, is shown in Figure 18. On a more fine-grained level, we also wish to support run-time communication of data between neuromorphic systems and external software, for example simulated virtual environments.

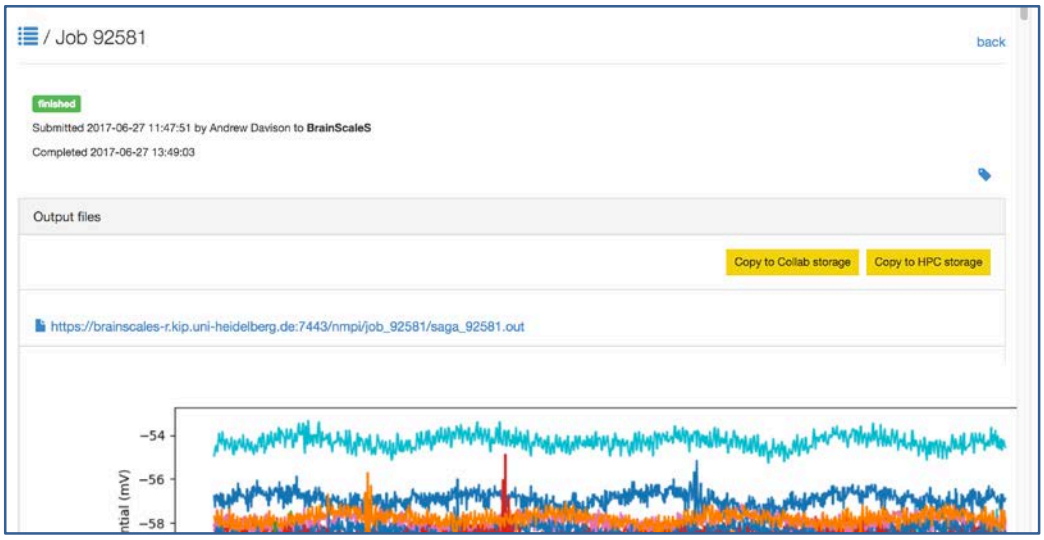

Figure 18: Job output example from a BrainScaleS experiment, accessed in the HBP Collaboratory

Remote access

The components that implement remote access to the Platform, listed in the table below, were first released at the end of the Ramp-Up Phase of HBP, in March 2016. Since then, the components have been continuously improved, each new feature being released as soon as it passes quality control.

In the first year of SGA1 (April 2016-March 2017), improvements and new features were developed and released in several areas. In the Collaboratory Job Manager app, the option was added to edit and resubmit an existing job, making iterative development of models much less time consuming. "Deep-links" to individual jobs were added, making it possible to share results just by passing around a URL. The Python client was extended to allow copying of simulation results to long-term storage or to the HPAC Platform, and querying of quota usage for the different systems. The robustness of the back end web service against user and hardware errors was improved. Quality control of the back end web service and the Python client was improved by writing unit tests for all components, and deploying continuous integration using Travis CI. The source code for all components was released under open-source licences through Github.

In the second year of SGA1 (April 2017-March 2018), we introduced support in both the Job Manager app and the Python client for tagging jobs, commenting on jobs,

and search/filtering of jobs. These features make it easier to find results of interest in a long-lived project, and to collaborate effectively with others. Based on user feedback, we added options to provide simulation code from files or folders in Collaboratory storage, and to use a Jupyter notebook as simulation input; the latter features make it simpler to develop a model using a software simulator such as NEST or NEURON, and then to switch to a neuromorphic simulation when the model is ready. To enhance reproducibility, limited provenance information was added to the job results, including, for example, the precise versions of libraries used on the hardware systems. In addition to bug fixes and performance improvements, quality control was further improved by developing unit tests for the front end Javascript code underlying the Collaboratory apps.

Programming interface

The programming language for using both the SpiNNaker and BrainScaleS systems is Python, using the PyNN API. Importantly, PyNN can also be used to program the NEST and NEURON simulators used in the HBP Brain Simulation and Neurorobotics Platforms, which makes it relatively straightforward to cross-check results, and to transfer models between platforms. PyNN, and the related interface modules for SpiNNaker (sPyNNaker) and BrainScales (pyhmf) were originally developed prior to the Human Brain Project. At the end of the HBP Ramp-Up Phase, version 0.8.0 of PyNN, with support for NEST v2.8 and NEURON v7.3, had recently been released. pyhmf and sPyNNaker implemented most features of version 0.7 of the PyNN API, but did not yet support PyNN 0.8.

In the first year of SGA1, three minor versions of PyNN (0.8.1 – 0.8.3) were released, supporting NEST versions 2.10 and 2.12, NeuroML v2, and new spike source models. In the BrainScaleS implementation of the PyNN API and underlying software, coverage of the API was improved, logging and monitoring of hardware statistics was introduced, and new data structures representing the Map & Route results were added, giving faster reconfiguration of the system for in-the-loop experiments. Other performance improvements included optimised data handling and parallel chip configuration for time-consuming operations. For the BrainScaleS 2 prototype system, software support for the second prototype chip was completed, and initial support for the third prototype implemented. The SpiNNaker implementation of the PyNN API (« sPyNNaker ») saw multiple performance improvements, including support for higher-firing rates through fewer dropped packets, support for running multiple simulations at the same time, improved debug support, faster mapping of model networks to the hardware, and support for parameter changes without re-mapping.

In the second year of SGA1, PyNN 0.9.0 was released, with support for the new, simplified Neo object model (Neo v0.5+; cf the report for SP5, where the impact of Neo for experimental datasets is described). Two minor releases (0.9.1, 0.9.2) introduced new stochastic synapse models, recording of injected currents, and support for NEST v2.14 and NEURON v7.5. For the BrainScaleS system we optimised re-configuration speed for in-the-loop experiments (e.g. reconfiguration of all synapse weights in O(1s)), improved detection for faulty hardware components in user experiments, and implemented a browser-based visualisation tool for mapping results. For the BrainScaleS Prototype System, support for the ADC was made available on the prototype system PCBs, and 5 setups are now available for remote users. sPyNNaker added support for versions 0.8 and 0.9 of the PyNN API while maintaining support for version 0.7. Quality control was improved by deploying

continuous integration. New features included improved support for external devices, focussing on the TUM PushBot, improved partitioning speed through estimation of synaptic matrix size, support for estimating power usage of scripts, support for alpha synapses, lazy list support to improve host memory usage, the ability to record only parts of a population and/or at a reduced sampling rate, and faster extraction of recorded data.

Benchmarking

For the development of platform-agnostic benchmark programmes, an additional abstraction layer with an accompanying software framework has been developed and extended during this project phase. The framework SNAB Suite (Spiking Neural Architecture Benchmark Suite) provides a uniform interface for collections of individual benchmarks and supports automated evaluation runs for platform comparisons and performance regression detection, as well as parameter sweeps. The included set of benchmarks consists of building blocks for neuromorphic applications [Ref 17] and explorations of platform parameters in terms of the application domain.

Interaction of neuromorphic systems with external software

MUSIC is an API and library for on-line communication of spike events and other data during simulation. By providing a standard interface, it helps to connect software and hardware in a modular way. MUSIC itself was developed prior to HBP, while further improvements and supporting software have been developed during SGA1.

During the first year of SGA1, we added a tool chain for real-time communication with ROS (the “Robot Operating System”, used in the HBP Neurorobotics Platform.

During the second year, we developed a new MUSIC prototype supporting dynamic port creation and deletion, for use by the Neurorobotics Platform, a new MUSIC API, supporting spike communication at accelerated time scales with the BrainScaleS system, and a MUSIC-SpiNNaker interface.

Component ID	Component Name	HBP Internal	Comment
349 (Table 10)	PyNN	Partially	Python library
376	Software for BrainScaleS Systems	Yes	Python and C++ software
375	SpiNNaker software stack	Yes	Python, C++, Java software
343	Neuromorphic Job Manager app	Yes	Collaboratory app
344	Neuromorphic Job Queue service	Yes	Web service

Component ID	Component Name	HBP Internal	Comment
345	Python client for the Neuromorphic Computing Platform	Yes	Python library
368	Neuromorphic jobs database	Yes	Database
369	Neuromorphic Quotas service	Yes	Web service
370	Neuromorphic Resource Manager app	Yes	Collaboratory app
371	Neuromorphic Dashboard app	Yes	Collaboratory app
2462 (Table 7)	SNABSuite benchmarks and framework	Partially	Python and C++ software

Table 2: List of software components

Usage of the Platform

Since the end of the Ramp-Up Phase, the number of people using the Neuromorphic Platform through the Collaboratory and/or Python client has increased from 13 to 66, while the total number of jobs submitted through these interfaces has increased from 780 to 2433, see Figure 19 and Figure 20 below. Note that these figures do not include developers and/or power users with direct or ssh access to the systems.

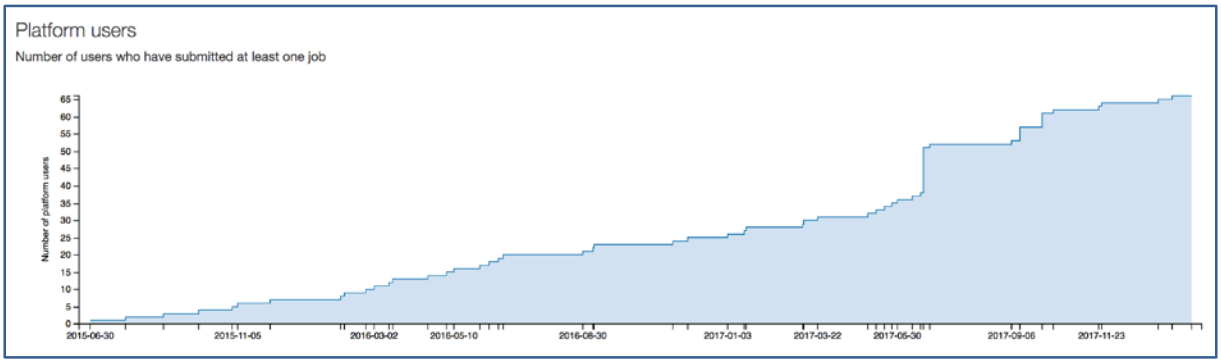

Figure 19: Number of users who submitted at least one job to the SP9 Platform

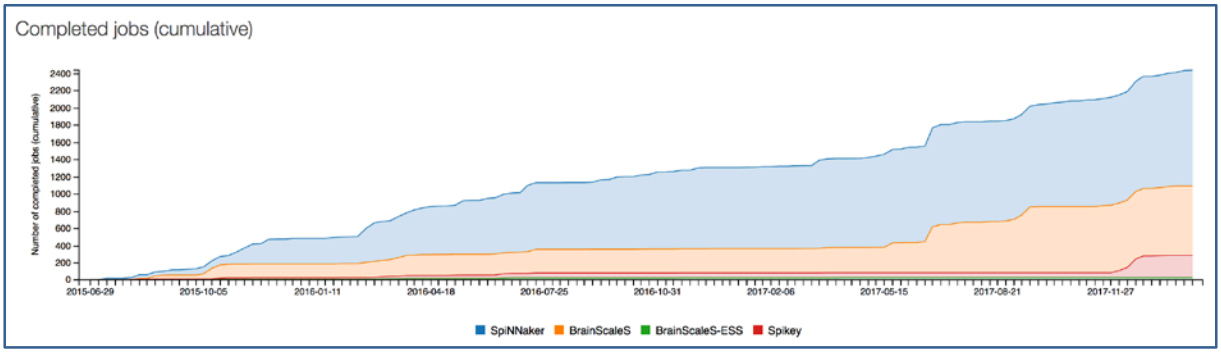

Figure 20: Cumulative number of jobs submitted to the SP9 Platform

2.5.2 Impact

During the reporting period, software development has made major advances towards a unique ecosystem for neuromorphic computing in SP9. All SpiNNaker, BrainScaleS-1, BrainScaleS-2 and Spikey interfaces, PyNN as well as all remote access components, are open source and publicly available on GitHub repositories: https://github.com/SpiNNakerManchester, https://github.com/electronicvisions, https://github.com/HumanBrainProject/hbp-neuromorphic-client, https://github.com/NeuralEnsemble/).

Intel announced to support PyNN for the new Loihi chip.

Ref 17 A. Stöckel, C. Jenzen, M. Thies, U. Rückert, "Binary Associative Memories as a Benchmark for Spiking Neuromorphic Hardware", Frontiers in Computational Neuroscience 11:71, August 2017. https://doi.org/10.3389/fncom.2017.00071

2.5.3 Challenges

Providing neuromorphic computing with an attractive software ecosystem that can be used by non-hardware experts is a key challenge for the field. HBP has been a pioneer in this approach. Tools like PyNN and the Collaboratory are well established and even used by new contributors like Intel, which announced the use of PyNN. The challenge for PyNN is the integration of new HBP specific hardware capabilities like various learning algorithms, structured neurons and the integration with supercomputers towards hybrid systems. SP9 has a strong software group that at the same time works in direct cooperation with the hardware designers. It is this direct cooperation that allows to meet the challenge of building a powerful integrated framework for neuromorphic computing.

2.5.4 Future Work

Future work on the software side of neuromorphic computing will have a strong focus on the second-generation systems, both from the BrainScaleS and the SpiNNaker approach. Both systems offer very unique capabilities in particular related to learning. HBP is in a leading position to secure a strong European standing in this field and the software will be crucial to bring the new hardware capabilities to the users.

2.5.5 Component Dependencies

Component ID	Component Name	HBP Internal	Comment
1	BrainScaleS-1 Neuromorphic Computing System	No	BrainScaleS-1 version – the starting point for the development of BrainScaleS-2
2	SpiNNaker Neuromorphic Computing System	No	SpiNNaker-1 – the basis for the SpiNNaker-2 development
457	BrainScaleS 2 Neuromorphic Computing System	No	BrainScaleS-2 evaulation
467	SpiNNaker small-scale NM-MC System	No	SpiNNaker-2 evaluation
	Components listed in Table 2: List of software components above		

2.6 Computational principles (Theory development for neuromorphic systems)

During the reporting period partners in SP9 have tackled important questions to foster the potential of neuromorphic hardware for powerful computation.

2.6.1 Results

A model for function-oriented rewiring of networks of neurons in the brain

Recent experimental data in neuroscience point to a structural difference between the commonly considered connectivity of artificial neural networks (ANNs) and digital computers on one hand, and networks of neurons in the brain (BNNs) on the other hand: BNNs continuously undergo rewiring, especially during learning. Previous models for learning in BNNs did not include this important effect. In [Ref 18], partner TUGRAZ proposed a model for reward-based network plasticity that integrates rewiring with reward-based STDP in a principled manner. This model can reproduce quite a number of experimental data for BNNs. However, it also turns out to provide new algorithmic ideas for integrating rewiring into network plasticity to reduce the memory footprints of ANNs. While memory is a critical bottleneck to reduce the size and energy consumption of neuromorphic hardware, TUGRAZ has shown in [Ref 19] that the memory requirement of machine learning benchmarks is reduced by a factor of 20 to 100 when using rewiring. Through a collaboration with TU Dresden, preliminary results have confirmed that rewiring makes it possible to implement realistic BNNs and performing ANNs on the next generation of neuromorphic hardware.

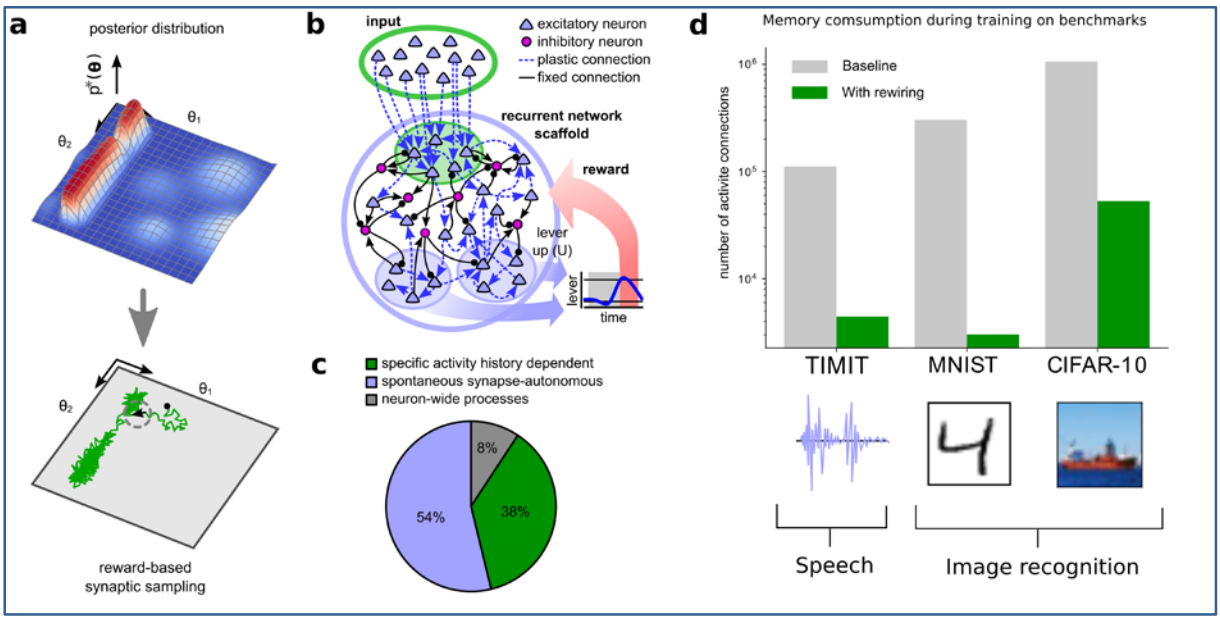

Figure 21: Self-configuration and rewiring of recurrent networks.

Figure 21 shows (a) Synaptic sampling model (schema): The main idea is that trajectories of network parameters θ (bottom) represent a distribution p*(θ) over network configurations (top). This distribution is shaped to plasticity rules: good configurations will have a high probability. (b) Recurrent network scaffold and task scheme. Two pools D and U were randomly selected to control lever movement. Inset: lever movement to receive a reward. (c) Synaptic parameter changes in the model and in experimental data show random components. Contributions of history-dependent activity, synapse-autonomous and neuron-wide processes to synaptic dynamics (d) In artificial neural networks synaptic sampling is used to reduce the memory footprint by continuously rewiring the network.

Methods to reproduce computations essential for higher level cognitive processing in the human brain.

Numerous experimental data from neuroscience highlight the importance of an intermediate level of network organisation and neural coding in BNNs, that lies between the level of single neurons and that of the whole network: neural coding of concepts and other tokens of higher-level brain computations through assemblies of neurons. However, it turned out to be quite nontrivial to reproduce the emergence of neural assemblies in networks of spiking neurons. TUGRAZ has characterised conditions under which assemblies emerge as codes for frequently occurring network inputs, both theoretically [Ref 20, Ref 21] and in simulations of biologically quite realistic recurrent networks of spiking neurons with STDP [Ref 22, Ref 23]. In addition, it was investigated both theoretically [Ref 21] and in computer simulations [Ref 22] under which conditions associations between two such external stimuli can be encoded in the network, as reported by [Ison *et al*., 2015[1]] for the human brain. Furthermore, such assemblies can function as basic tokens for higher-level cognitive abilities [Ref 24]. These results provide a solid basis for emulating a significant class of cognitive computations in large-scale networks of spiking neurons, in particular in neuromorphic hardware systems.

Learning-to-Learn (L2L) applicable for research in computational neuroscience and neuromorphic engineering.

L2L has recently become a powerful tool in machine learning, in particular for enabling transfer learning, i.e. for enabling fast learning of a new task after previous learning of related tasks. However, L2L so far has not been applied to biologically more realistic neural network models, in particular not to networks of spiking neurons. However, the potential advantage of using L2L in computational neuroscience is quite large, it also includes aspects of automatic parameter tuning of models for BNNs in computers and neuromorphic hardware. Therefore, a scalable software framework for applying L2L methods to networks of spiking neurons using gradient free optimisation algorithms was created at TUGRAZ both for digital simulations, including large-scale simulations on a supercomputer, and neuromorphic hardware. This software framework, written in Python, is designed with a general purpose API to facilitate adoption, it uses distributed scheduling to use all available computational resources, and it contains implementations of various optimisation algorithms. Proofs of concept for the viability of this new research strategy in modelling and in understanding learning in BNNs have already been produced. It was shown that L2L can create and optimise new algorithms for supervised learning in recurrent networks of spiking neurons, similarly as demonstrated by Hochreiter *et al* (2001) for artificial neural networks. This opens new doors for exploring and understanding a large variety of possible learning methods for spiking neural networks. In addition, a proof of concept was given in a collaboration of TUGRAZ and UHEI for transfer learning in networks of spiking neurons, so that a new task can be learnt much faster when the network had previously learned somewhat related tasks. In particular, with this method we could enable transfer learning of the neuromorphic HICANN-DLS chip that has been developed in Heidelberg, and which will provide the basis for the next generation BrainScaleS Neuromorphic Platform. Hence, we expect that L2L will significantly

[1] MJ Ison, RQ Quiroga and I Fried. Rapid encoding of new memories by individual neurons in the human brain. Neuron, 87(1), 220-230

improve the capability for efficient autonomous learning of this neuromorphic system.

A theoretical framework for stochastic inference in networks of integrate-and-fire neurons ported to accelerated analog neuromorphic substrates

The ability to perform probabilistic (Bayesian) inference is a hallmark of mammalian cognition and a coveted feature for embedded AI. Recent developments in machine learning have tried to capture this kind of computation with so-called "deep" architectures, but the analogy to biology remains superficial. UHEI, UBERN and TUGRAZ have developed a framework for cognitive computation with spiking neurons that narrows the gap between biological and artificial deep networks, while employing well-documented aspects of cortical dynamics such as spike-based communication, operation in a high-conductance state, short-term plasticity and background-driven stochasticity (Component 2474; Ref 2). A particularly intriguing finding was that spike-triggered short-term plasticity enables a form of local tempering, allowing the network to more easily jump between low-energy modes, thereby significantly speeding up inference [Ref 25].

As many other models in computational neuroscience, these networks rely on the existence of background noise to provide the necessary stochasticity. However, more than just assuming the availability of such uncorrelated noise – a feature that is difficult to reconcile with cortical architecture and dynamics – we have shown how our functional networks can be embedded in larger networks of recurrently connected neurons in order to achieve the required input statistics (Component 2475; Ref 3). These "sea-of-noise networks" also provide the basis for large-scale hardware implementation.

In addition to building a bridge to biology, the above-mentioned features are part of the BrainScaleS Neuromorphic Platform design specifications. Component variability was tackled by a combination of hierarchical networks' intrinsic robustness with respect to some of these sources of distortion [Ref 26] and in-the-loop training [Ref 1]. This allowed the implementation and accelerated emulation of hierarchical spiking networks able to perform probabilistic inference in high-dimensional data spaces (Component 2476, 2477; Ref 27).

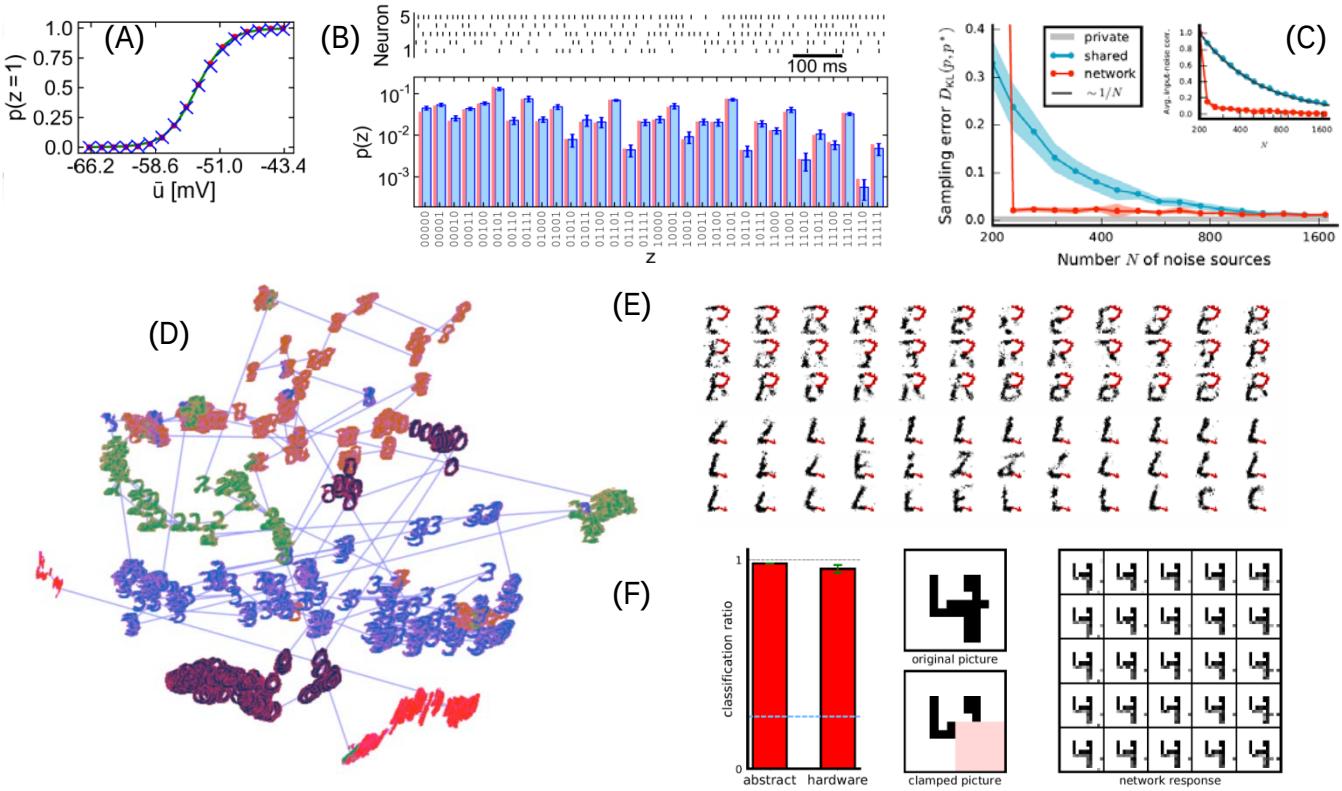

Figure 22: Stochastic inference in networks of integrate-and-fire neurons

(A) Simulated activation function of an integrate-and-fire neuron in the high-conductance state (blue) and theoretical prediction (red). The fast membrane enables a more symmetric shape, which enables the often-used approximation by a logistic function (green). (B) Spiking network sampling (blue) from a target distribution (red). (C) The stochastic output of cortical networks is driven by stochastic spiking input, generated by the cortex itself. An average connectivity dominated by inhibition allows even small, deterministic networks (red) to supply functional networks with uncorrelated noise, unlike a limited pool of independent Poisson sources (blue). This observation enables the emulation of stochastic networks in neuromorphic hardware, even when harsh limits on external bandwidth are present (panel F). (D) Visible layer activity of a hierarchical spiking network. While classical neural nets tend to become stuck in local energy minima, spiking networks endowed with short-term plasticity are able to overcome energy barriers, enabling them to simultaneously be good generative and discriminative models of real-world data. (E) In a large enough ensemble of stochastic networks, external noise can even become superfluous. The two networks shown here, performing pattern completion on the ambiguous red input, only receive background input from (and provide it to) other functional spiking networks in the ensemble. (F) Hierarchical sampling network trained on handwritten digits performing classification and pattern completion on the BrainScaleS Neuromorphic Platform.

2.6.2 Impact

The impact of the theory work in SP9 is exemplified by the following list of references:

Ref 1: Sebastian Schmitt, Johann Klähn, Guillaume Bellec, Andreas Grübl, Maurice Güttler, Andreas Hartel, Stephan Hartmann, Dan Husmann,Kai Husmann, Sebastian Jeltsch, Vitali Karasenko, Mitja Kleider, Christoph Koke, Alexander Kononov, Christian Mauch, Eric Müller, Paul Müller, Johannes Partzsch, Mihai A. Petrovici, Stefan Schiefer, Stefan Scholze, Vasilis Thanasoulis, Bernhard Vogginger, Robert Legenstein, Wolfgang Maass, Christian Mayr, René Schüffny, Johannes Schemmel, Karlheinz Meier (2017, May). Neuromorphic hardware in the loop: Training a deep spiking network on the brainscales wafer-scale system. In Neural Networks (IJCNN), 2017 International Joint Conference on (pp. 2227-2234). IEEE. https://doi.org/10.1109/IJCNN.2017.7966125

Ref 2: Petrovici, M. A., Bill, J., Bytschok, I., Schemmel, J., & Meier, K. (2016). Stochastic inference with spiking neurons in the high-conductance state. *Physical Review E*, *94*(4), 042312.

Ref 3: Jordan, J., Petrovici, M. A., Breitwieser, O., Schemmel, J., Meier, K., Diesmann, M., & Tetzlaff, T. (2017). Stochastic neural computation without noise. *arXiv preprint arXiv:1710.04931*.

Ref 18[Kappel *et al*. 2018] D Kappel, R Legenstein, S Habenschuss, Ml Hsieh and W Maass; A dynamic connectome supports the emergence of stable computational function of neural circuits through reward-based learning. arXiv:1704.04238 [q-bio.NC]

Ref 19[Bellec *et al*. 2018] G Bellec, D Kappel, W Maass and R Legenstein. Deep Rewiring: Training very sparse deep network.Proceedings of the International Conference on Learning Representations (ICLR)

Ref 20 [Legenstein *et al*., 2017] R Legenstein, Z Jonke, S Habenschuss and W. Maass. A probabilistic model for learning in cortical microcircuit motifs with data-based divisive inhibition. arXiv:1707.05182

Ref 21 [Legenstein *et al*., 2018] R Legenstein, W Maass, CH Papadimitriou and SS Vempala. Long term memory and the densest K-subgraph problem. In Proc. of Innovations in Theoretical Computer Science (ITCS)

Ref 22 [Pokorny *et al*., 2017] . Pokorny, MJ Ison, A Rao, R Legenstein, C Papadimitriou and W Maass.Associations between memory traces emerge in a generic neural circuit model through STDP. bioRxiv:188938 https://doi.org/10.1101/188938

Ref 23 [Jonke *et al*., 2017] Z Jonke, R Legenstein, S Habenschuss and W Maass. Feedback inhibition shapes emergent computational properties of cortical microcircuit motifs. Journal of Neuroscience, 37(35):8511-8523, https://doi.org/10.1523/JNEUROSCI.2078-16.2017

Ref 24 [Legenstein *et al*., 2016] R Legenstein, CH Papadimitriou, S Vempala and W Maass. Assembly pointers for variable binding in networks of spiking neurons. arXiv preprint arXiv:1611.03698

Ref 25 [Leng *et al*., 2017] L Leng, R Martel, O Breitwieser, I Bytschok, W Senn, J Schemmel, K Meier, MA Petrovici. Spiking neurons with short-term synaptic plasticity form superior generative networks. arXiv:1709.08166

Ref 26 [Petrovici *et al*., 2017a] MA Petrovici, A Schroeder, O Breitwieser, A Grübl, J Schemmel, K Meier. Robustness from structure: Inference with hierarchical spiking networks on analog neuromorphic hardware. Proceedings of the 2017 IEEE International Joint Conference on Neural Networks

Ref 27 [Petrovici *et al*., 2017b] MA Petrovici, S Schmitt, J Klähn, D Stöckel, A Schroeder, G Bellec, J Bill, O Breitwieser, I Bytschok, A Grübl, M Güttler, A Hartel, S Hartmann, D Husmann, K Husmann, S Jeltsch, V Karasenko, M Kleider, C Koke, A Kononov, C Mauch, E Müller, P Müller, J Partzsch, T Pfeil, S Schiefer, S Scholze, A Subramoney, V Thanasoulis, B Vogginger, R Legenstein, W Maass, R Schüffny, C Mayr, J Schemmel, K Meier. Pattern representation and recognition with accelerated analog neuromorphic systems. Proceedings of the 2017 IEEE International Symposium on Circuits and Systems (ISCAS) https://doi.org/10.1109/ISCAS.2017.8050530

2.6.3 Challenges

In almost all neuromorphic computing projects worldwide the main problem is a lack of integration between hardware engineering, software development and grounding in neuroscience, in particular mathematical and theoretical neuroscience. In HBP, and SP9 in particular this has been solved in a way that is recognised as a model case for many upcoming activities worldwide, like in China and the US.

Several instruments have been implemented in the HBP to meet the integration challenge. Within SP9 an internal workpackage 4 develops principles of neuromorphic computing and has implemented the highly successful Fürberg workshop series. Co-design project CDP5 has been strongly supported by SP9 during the second half of SGA1 and is now a very strong element to implement learning as a key principle in the SP9 systems. The EITN in Paris is an internationally recognised contact point to enable collaboration especially also outside the HBP.

2.6.4 Future Work

The connection between theory and neuromorphic computing has identified several novel principles that will be further pursued- Among those are: Neural spike-based sampling, learning-to-learn and several novel methods of biological deep learning pursued in CDP5. Reviewers have identified this connection to be one of the most promising aspects of the entire HBP.

3. The Industry Angle

Computer technology has progressed spectacularly since the world's first electronic stored program computer successfully ran its first program in Manchester on June 21st 1948. Today's machines are a hundred billion times more energy efficient (measured in terms of the energy they consume to execute an instruction), much smaller and much cheaper than those early machines, and this has enabled computer technology to penetrate every area of human activity. As an example of the degree of penetration, shipments of ARM processors have passed 75 billion, representing 10 computers for every human on the planet.

However, computer technology is now at a crossroads.

The major means whereby computer technology has made these advances has been the ever-shrinking transistor used on integrated circuits. This process, typically understood as being driven by Moore's Law, has been driven by formidable investments in manufacturing technology to exploit the fact that historically, as transistors are made smaller, they become faster, cheaper, and more energy efficient - a win-win scenario. However, although this process has driven exponential progress (with a doubling time constant of around two years) for half a century, it cannot continue forever, and today's transistors are approaching physical limits determined primarily by the size of a silicon atom. As these limits are approached the win-win effects have begun to weaken, and the industry is looking at a future which will be very different from the last half-century.

The first signs of the limits to exponential progress arrived a decade ago. During the 1990s the microprocessor business was characterised by ever increasing clock rates, but soon after the turn of the century all of the major players threw in the towel on this approach. The reason was power. Although smaller transistors are more energy-efficient, they can also be packed closer together, and the increased density dominates the efficiency improvements, so chips get hotter until heat limits their operating speed. So, although Moore's Law continued to deliver more transistors on a chip, the operating clock had to be constrained to avoid overheating, and the performance of an individual microprocessor ceased improving. Instead, the extra transistors were used to put several microprocessors onto the same chip, and the era of multicore computing had arrived. Multicore computing is an efficient way to deliver more processing power, but it requires a different programming paradigm, and effective parallel programming had been the holy grail of computer science for half a century! Still, there was now no other way forward.

Today multi- and many-core computing is everywhere, from the mobile phone and tablet through to the high-performance computer. There are still issues with developing software to make effective use of these resources, but on the whole we have learnt to live with the problem, and the technology works well.

However, the gains from Moore's Law have almost dried up, and new directions must be found if we are to come close to maintaining historic rates of progress in computer performance. As with the transition to multicore computing, energy-efficiency considerations dominate. We can now put so many computational resources onto a chip that we face the prospect of "dark silicon" - microchips where we cannot afford to turn on more than a fraction of the chip's resources at any time, otherwise the chip will overheat and fail. New, more efficient models of computation are needed urgently!

A major feature of the standard model of computation is the need for precision. Ideally, computations should be exact, deterministic, reliable and repeatable. This requirement is clearly vital when the computer is used for banking, for example, but increasingly computers are used in less critical applications, and in many cases in applications where exact, repeatable results are not necessary nor, in some cases, possible. Examples of such applications include computer vision, which is of increasing importance in car driver assistance technology and in driverless cars and robotics, and speech recognition systems such as those widely used on mobile phones. Today, these applications use advanced machine learning techniques such as deep neural networks and convolutional neural networks, where very low precision parameters can be used without affecting performance. Indeed, using very accurate arithmetic such as the standard double-precision floating-point hardware on most microprocessors leads to at least an order of magnitude increase in energy consumption compared with using small integers, for no benefit in terms of recognition performance.

Neuromorphic Computing is one example of a more efficient model of computation. Neuromorphic Computing takes its inspiration from our (partial) knowledge of how the brain works, a similar approach, though closer to the biology, to that which led to the deep learning and convolution neural networks mentioned above. In Neuromorphic Computing the brain-like algorithms are not simply modelled on a conventional computer, they are ingrained into the design of the computational hardware itself, to a greater or lesser degree, in order to deliver more of the efficiency that is exemplified by the biological system itself.

Visibility and impact of neuromorphic computing have increased substantially during the last year. The reasons for this development are twofold.

- The first reason is connected to the huge success of deep learning (DL). The more this approach is accessible to everybody through open and easy to use software, the more the limitations become evident to the users. DL requires very large labelled data sets for a training process that works fundamentally different from the biological brain. The quest for learning principles derived from biology becomes more and more urgent as biological systems are capable of continuous learning and learning with small (even single shot) data sets. Transferring biological learning to artificial systems would be a game changer to AI with very substantial impacts of economy and society.
- The second reason is the very strong involvement of major industry players in alternative approaches to computing like quantum and neuromorphic computing. The fact that Intel has invested substantial resources to develop and support a spiking neural network chip very much along the lines of the HBP systems has also generated an enormous visibility for the research and development carried out in the HBP.

How will these 2 reasons affect our future work and the impact it will have on the industrial landscape in Europe?

The move towards biologically inspired learning mechanisms requires close, even day-by-day collaboration with theoretical neuroscience. The HBP is the only project worldwide that offers this collaboration between neuroscience and neuromorphic engineering. If this is properly exploited, the HBP could well develop into a very major player in the field of AI by co-designing novel hardware according to computational principles found in the brain.

The large investment by companies like Intel may appear like a strong competition that a European project with very limited funding cannot match. By comparing the 3 approaches (SpiNNaker, BrainScaleS and Loihi by Intel) it is however clear, that there is a very good complementarity, which offers the opportunity to explore different application spaces with the 3 approaches. It will also be essential that the 3 approaches collaborate in the development of support aspects like software. For the HBP systems it will be essential to go beyond the pure research aspect very soon and to receive substantial support for the development of systems that demonstrate a learning performance beyond the one reachable with conventional computing.

In an approach to address the industrial angle we present here 3 outlooks at different time scales

A short term outlook

A number of demonstrations of the benefits of neuromorphic technology are beginning to emerge, and more can be expected in the short to medium term. Various start-up companies are emerging, in the USA and elsewhere, to exploit the prospective advantages of neuromorphic and similar technologies in these new machine learning application domains.

An outlook toward the end of the HBP

In the medium term we may expect neuromorphic technologies to deliver a range of applications more efficiently than conventional computers, for example to deliver speech and image recognition capabilities in smartphones. (Currently such capabilities are available only using powerful cloud resources to implement the recognition algorithms.) These will require small-scale neuromorphic accelerators integrated with the application processor, using a fraction of the resources of a single chip.

An outlook into the far future

In the long term there is the prospect of using neuromorphic technology to integrate energy-efficient intelligent cognitive functions into a wide range of consumer and business products, from driverless cars to domestic robots. While human-level “strong” artificial intelligence remains a mystery, and indeed may depend on the emergence of an understanding of information processing in the biological brain (through initiatives such as the Human Brain Project) before it becomes a practical reality, there are many useful applications that can benefit from more modest cognitive capabilities. The technology is relatively young, and there is much uncertainty as to where it will find its place in the wider world, but it clearly meets a need in the rapidly changing world of computing.

4. Conclusion and Outlook

The end of the SGA1 funding period represents a major milestone in the entire 10-year HBP endeavour. Essentially, half of the project time has elapsed and about 40% of the planned EC resources have been invested in research and technology development. What have taxpayers received for their substantial investment until today?

SP9 has reached 2 major milestones that were expected at this stage of the project:

1) SP9 operates the only large-scale, remotely accessible neuromorphic computing system worldwide.
2) SP9 delivered proven designs and prototypes for the next generation neuromorphic systems, developed in close collaboration with theoretical neuroscientists in the HBP.

Noteworthy, the field of neural computation has undergone rapid developments since the early days of the HBP planning in 2009 (pilot phase) and 2013 (start of the main Project). Brain-inspired computing concepts became a major driver of what is called today Artificial Intelligence (AI). Traditional deep learning approaches are limited by their need for very large sets of labelled data and extremely power hungry hardware for the slow learning processes.

There is general consensus, that brain inspired concepts of learning represent the only way out of this dilemma. Biological brains learn on very small datasets, even on single events. They learn continuously and do not separate between learning and inference. And, of course, their energy efficiency is still far better than the one of conventional computer.

Here, the HBP is in a unique position. It is the only project worldwide that has established a research driven link between neuroscience and computing.

The new 2nd generation designs have generated substantial visibility in the international community, including major industry players. During the recent NICE conference at Intel Labs in the US, BrainScaleS-2 and SpiNNaker-2 played in the same league as Intel's new Loihi chip.

To exploit the leading edge Europe has gained in the field of brain inspired computing, it is essential to make the important step from 2nd generation prototype chips to large-scale neuromorphic systems that exhibit superior performance, compared to their competitors based on traditional computing architectures. SGA2, and even more SGA3, will show us whether Europe and European research funding are able to live up to these expectations.

.

5. Literature

Ref 1: Sebastian Schmitt, Johann Klähn, Guillaume Bellec, Andreas Grübl, Maurice Güttler, Andreas Hartel, Stephan Hartmann, Dan Husmann,Kai Husmann, Sebastian Jeltsch, Vitali Karasenko, Mitja Kleider, Christoph Koke, Alexander Kononov, Christian Mauch, Eric Müller, Paul Müller, Johannes Partzsch, Mihai A. Petrovici, Stefan Schiefer, Stefan Scholze, Vasilis Thanasoulis, Bernhard Vogginger, Robert Legenstein, Wolfgang Maass, Christian Mayr, René Schüffny, Johannes Schemmel, Karlheinz Meier (2017, May). Neuromorphic hardware in the loop: Training a deep spiking network on the brainscales wafer-scale system. In Neural Networks (IJCNN), 2017 International Joint Conference on (pp. 2227-2234). IEEE. https://doi.org/10.1109/IJCNN.2017.7966125 12

Ref 2: Petrovici, M. A., Bill, J., Bytschok, I., Schemmel, J., & Meier, K. (2016). Stochastic inference with spiking neurons in the high-conductance state. *Physical Review E*, *94*(4), 042312. https://doi.org/10.1103/PhysRevE.94.042312 12

Ref 3: Jordan, J., Petrovici, M. A., Breitwieser, O., Schemmel, J., Meier, K., Diesmann, M., & Tetzlaff, T. (2017). Stochastic neural computation without noise. *arXiv preprint arXiv:1710.04931*.....12

Ref 4: Alexander Kugele, Solving the Constraint Satisfaction Problem Sudoku on Neuromorphic Hard 12

Ref 5: Schemmel, J., Kriener, L., Müller, P., & Meier, K. (2017, May). An accelerated analog neuromorphic hardware system emulating NMDA-and calcium-based non-linear dendrites. In *Neural Networks (IJCNN), 2017 International Joint Conference on* (pp. 2217-2226). IEEE, https://doi.org/10.1109/IJCNN.2017.7966124 16

Ref 6: Friedmann, S., Schemmel, J., Grübl, A., Hartel, A., Hock, M., & Meier, K. (2017). Demonstrating hybrid learning in a flexible neuromorphic hardware system. *IEEE transactions on biomedical circuits and systems*, *11*(1), 128-142, http://doi.org/10.1109/TBCAS.2016.2579164.....16

Ref 7: G.A. Fonseca Guerra, S.B. Furber, "Using Stochastic Spiking Neural Networks on SpiNNaker to Solve Constraint Satisfaction Problems", Frontiers in Neuroscience – Neuromorphic Engineering, Dec 2017. https://doi.org/10.3389/fnins.2017.00714 21

Ref 8: Knight, James C. and Furber, Steve B., "Synapse-Centric Mapping of Cortical Models to the SpiNNaker Neuromorphic Architecture", Frontiers in Neuroscience 10, p.420, 14 September 2016. https://doi.org/10.3389/fnins.2016.00420 21

Ref 9: A. Yousefzadeh, M. Jabłon ▫ski, T. Iakymchuk, A. Linares-Barranco, A. Rosado, L.A. Plana, S. Temple, T. Serrano-Gotarredona, S.B. Furber and B Linares-Barranco, "On Multiple AER Handshaking Channels Over High-Speed Bit-Serial Bidirectional LVDS Links With Flow-Control and Clock-Correction on Commercial FPGAs for Scalable Neuromorphic Systems", in IEEE Trans Biomedical Circuits and Systems 11(5), pp. 1133-1147, Oct. 2017.DOI: 10.1109/TBCAS.2017.2717341, https://doi.org/10.1109/TBCAS.2017.271734121

Ref 10: Qian Liu, Garibaldi Pineda-García, Evangelos Stromatias, Teresa Serrano-Gotarredona, Steve B. Furber, "Benchmarking Spike-Based Visual Recognition:

6. Component Details

The following is a list of the newly released internal Components for this deliverable.

6.1 BrainScaleS-1 Neuromorphic Computing System

Field Name	Field Content	Additional Information
ID	1	
Component Type	Hardware	20 wafer-module system + local cluster in Heidelberg
Contact	MEIER, Karlheinz, SCHEMMEL, Johannes	
Latest Release	Latest wafer version is V4.1	3 of the 20 modules are on wafer version 4.1 as of March 2018
TRL	5	
Location	D-69118 Heidelberg, Im Neuenheimer Feld 227b	
Maintenance	Continuous operation and maintenance by UHEI team in WP9.2	
Curation Status	NA	
Validation - QC	Yes	Continuous integration tests for software. Nightly Jenkins test jobs for hardware and software, benchmarking
Validation – Users	Yes	Usage via the collab is reported in the dashboard: https://collab.humanbrainproject.eu/#/collab/609/nav/7820 Usage dominated by direct local use
Validation – Publications	Yes	Need to add ref to paper here
Privacy Constraints	No privacy constraints	
Sharing	publically authenticated	Accessible with HBP community account after signing the Platform User Agreement form
Licence	-	
Component Access URL	Getting started info: https://www.humanbrainproject.eu/en/silicon-brains/neuromorphic-computing-platform/	

Field Name	Field Content	Additional Information
	Collaboratory: Neuromorphic Computing Platform https://collab.humanbrainproject.eu/#/collab/51/nav/244	
Technical documentation URL	D9.7.1_Neuromorphic_Platform_Specification_-_public_version.pdf (a "living document")	
Usage documentation URL	HBP Neuromorphic Computing Platform Guidebook, also available as downloadable .pdf version. Both documents linked from the public getting-started page https://www.humanbrainproject.eu/en/silicon-brains/neuromorphic-computing-platform/	
Component Dissemination Material URL	"Silicon brains" part of the HBP public website. Videos (of talks, trainings): https://www.humanbrainproject.eu/en/silicon-brains/media/videos/ Events: https://www.humanbrainproject.eu/en/silicon-brains/events/	

Table 3: Component 1: BrainScaleS-1 Neuromorphic Computing System

6.2 SP9 BrainScaleS 2 Neuromorphic Computing System

Field Name	Field Content	Additional Information
ID	457	
Component Type	Hardware	BrainScaleS-2 prototype
Contact	SCHEMMEL, Johannes and GRÜBL, Andreas	
Latest Release	HICANN-X	HICANN-X to be taped out early April 2018
TRL	3	
Location	Heidelberg	
Maintenance	Active development WP9.2	

Field Name	Field Content	Additional Information
Curation Status		
Validation - QC	yes	Test chips for checking and evaluating several aspects of the system are in use and test.
Validation – Users	yes	Internal first-test usage of small scale feature test chips ; TUGraz and UBERN users of the DLS systems
Validation – Publications	no	
Privacy Constraints	no privacy constraints	
Sharing	-	
Licence	-	
Component Access URL	Remotely accessible, but not yet integrated in the NMPI frontend to access the machines.	
Technical documentation URL	NA	
Usage documentation URL	NA	
Component Dissemination Material URL	NA	

Table 4: Component 457: BrainScaleS 2 Neuromorphic Computing System

6.3 SP9 SpiNNaker Neuromorphic Computing System

Field Name	Field Content	Additional Information
ID	2	
Component Type	hardware	
Contact	FURBER, Steve	
Latest Release	SpiNNaker 1	500.000 core system accessible
TRL	TRL 5 - Prototype Integration	
Location	Manchester	
Maintenance	Continuous operation and maintenance by UMAN team in WP9.3	

Field Name	Field Content	Additional Information
Curation Status	NA	
Validation - QC	Yes	Benchmarking, code testing
Validation – Users	Yes	Usage info via the dashboard: https://collab.humanbrainproject.eu/#/collab/609/nav/7820
Validation – Publications	Yes	
Privacy Constraints	no	
Sharing	publically authenticated	Accessible with HBP community account after signing the Platform User Agreement form
Licence		
Component Access URL	Getting started info: https://www.humanbrainproject.eu/en/silicon-brains/neuromorphic-computing-platform/ Collaboratory: Neuromorphic Computing Platform https://collab.humanbrainproject.eu/#/collab/51/nav/244	
Technical documentation URL	D9.7.1_Neuromorphic_Platform_Specification_-_public_version.pdf (a "living document")	
Usage documentation URL	HBP Neuromorphic Computing Platform Guidebook, also available as downloadable .pdf version. Both documents linked from the public getting-started page https://www.humanbrainproject.eu/en/silicon-brains/neuromorphic-computing-platform/	
Component Dissemination Material URL	"Silicon brains" part of the HBP public website. Videos (of talks, trainings): https://www.humanbrainproject.eu/en/silicon-brains/media/videos/ Events: https://www.humanbrainpr	

Field Name	Field Content	Additional Information
	oject.eu/en/silicon-brains/events/	

Table 5: Component 2: SP9 SpiNNaker Neuromorphic Computing System (hardware)

6.4 SpiNNaker 2 small-scale NM-MC System

Field Name	Field Content	Additional Information
ID	467	
Component Type	hardware	
Contact	HÖPPNER, Sebastian	
Latest Release		
TRL	3	
Location	Manchester	
Maintenance	Software development in GIT repository, Hardware desing sources in SVN	
Curation Status	NA	
Validation - QC		Validation of the hardware by multiple users from TUD and UMAN Execution of low level tests (ARM software execution) and high level benchmarks (e.g. synfire chain, bursing network, dynamic vision sensor interface, BCPNN, reward based stochastic synaptic sampling)
Validation – Users		Usage by partners UMAN, TUGRAZ and TUD
Validation – Publications		
Privacy Constraints	None	
Sharing	consortium - share with any consortium members	
Licence		
Component Access URL	(not yet remote accessible)	
Technical documentation URL	https://hpsn.et.tu-dresden.de/svn/p_cfaed/s_santos28/v0.0.1/doc/Santos28_Documentation/	Internal only URL

	Santos28_Testchip_Documentation.doc	
Usage documentation URL	NA	
Component Dissemination Material URL	NA	

Table 6: Component 467: SpiNNaker small-scale NM-MC System (hardware)

6.5 SNABSuite

	Main Meta Data	Comment/title
Component	2462	SNABSuite
Component Type	Software	Python and C++ library
Contact	THIES, Michael	
Latest Release	February 2018	
TRL	5	
Location	hosted by other non-HBP 3rd party: Github	
Maintenance	Active development and maintenance in WP9.1	
Curation Status	NA	
Validation - QC	Pass	Continuous Integration
Validation – Already existing users	Yes	
Validation – Use in publications	Yes	
Access privacy	No Privacy Constraint	
Access sharing	anonymous	
Access licence	GNU GENERAL PUBLIC LICENSE, Version 3	
URL to access component	https://github.com/hbp-unibi/SNABSuite	
URL to component documentation	https://hbp-unibi.github.io/SNABSuite/index.html	

	Main Meta Data	Comment/title
URL to component usage documentation	https://github.com/hbp-unibi/SNABSuite/blob/master/README.md	
URL to dissemination material highlighting component	https://doi.org/10.3389/fncom.2017.00071	

Table 7: Component 2462: SNABSuite

6.6 Neuromorphic Job Manager App

	Main Meta Data	Comment/title
Component	343	
Component Type	Software	
Contact	DAVISON, Andrew	
Latest Release	2017-12-01	
TRL	7	
Location	Data hosted by collaborator storage	
Maintenance	Active development and maintenance in WP9.1	
Curation Status	NA	
Validation - QC	Pass	Agile Quality Assurance
Validation – Already existing users	Yes	See Dashboard - https://nmpi.hbpneuromorphic.eu/dashboard/
Validation – Use in publications	No	
Access privacy	No Privacy Constraint	
Access sharing	public authenticated - share with authenticated non-consortium members e.g. public collab	
Access licence	Apache v2 license	
URL to access component	https://collab.humanbrainproject.eu/#/collab/19/nav/2108?state=software,Neuromorphic-Computing-Platform-Job-Manager	

	Main Meta Data	Comment/title
URL to component documentation	https://collab.humanbrainproject.eu/#/collab/51/nav/1069	
URL to component usage documentation	https://collab.humanbrainproject.eu/#/collab/51/nav/1069	
URL to dissemination material highlighting component	NA	

Table 8: Component 343: Neuromorphic Job Manager App

6.7 Neuromorphic Job Queue Service

	Main Meta Data	Comment/title
Component	344	
Component Type	Service	
Contact	DAVISON, Andrew	
Latest Release	2017-12-01	
TRL	7	
Location	data hosted by other non-HBP 3rd party	Code: Github Hosting: Digital Ocean AMS2 data centre, Amsterdam
Maintenance	Active development and maintenance in WP9.1	
Curation Status	NA	
Validation - QC	Pass	Agile Quality Assurance
Validation – Already existing users	Yes	See Dashboard - https://nmpi.hbpneuromorphic.eu/dashboard/
Validation – Use in publications	NA	
Access privacy	No Privacy Constraint	
Access sharing	public authenticated - share with authenticated non-consortium members e.g. public collab	
Access licence	Apache v2 license	

	Main Meta Data	Comment/title
URL to access component	The job queue service is the backend for the experiments submitted via the Collab	
URL to component documentation	https://collab.humanbrainproject.eu/#/collab/51/nav/1069	
URL to component usage documentation	https://collab.humanbrainproject.eu/#/collab/51/nav/1069	
URL to dissemination material highlighting component	NA	

Table 9: Component 344: Neuromorphic Job Queue Service

6.8 PyNN

	Main Meta Data	Comment/title
Component	349	
Component Type	Software	
Contact	DAVISON, Andrew	
Latest Release	February 2018	
TRL	5	
Location	data hosted by other non-HBP 3rd party	Github
Maintenance	Active development and maintenance in WP9.1	
Curation Status	NA	
Validation - QC	Pass	Continuous Integration
Validation – Already existing users	Yes	
Validation – Use in publications	Yes	
Access privacy	No Privacy Constraint	
Access sharing	anonymous	
Access licence	GNU GENERAL PUBLIC LICENSE, Version 3	

	Main Meta Data	Comment/title
URL to access component	https://github.com/hbp-unibi/SNABSuite	
URL to component documentation	https://hbp-unibi.github.io/SNABSuite/index.html	
URL to component usage documentation	https://github.com/hbp-unibi/SNABSuite/blob/master/README.md	
URL to dissemination material highlighting component	https://doi.org/10.3389/fncom.2017.00071	

Table 10: Component 349: PyNN

6.9 SP9 Neuromorphic Python client

	Main Meta Data	Comment/title
Component	349	
Component Type	Software	
Contact	DAVISON, Andrew	
Latest Release		
TRL	7	
Location	data hosted by other non-HBP 3rd party	Github
Maintenance	Active development and maintenance in WP9.1	
Curation Status	NA	
Validation - QC	Pass	Agile Quality Assurance
Validation – Already existing users	Yes	https://github.com/HumanBrainProject/hbp-neuromorphic-client/stargazers
Validation – Use in publications	*No*	
Access privacy	No Privacy Constraint	
Access sharing	anonymous	
Access licence	Apache v2 license	

	Main Meta Data	Comment/title
URL to access component	https://github.com/HumanBrainProject/hbp-neuromorphic-client	
URL to component documentation	https://collab.humanbrainproject.eu/#/collab/51/nav/1069	
URL to component usage documentation	https://collab.humanbrainproject.eu/#/collab/51/nav/1069	
URL to dissemination material highlighting component		

Table 11: Component 349: SP9 Neuromorphic Python client

6.10 Software model of sampling with LIF neurons

Field Name	Field Content	Additional Information
ID	2474	
Component Type	Model	
Contact	PETROVICI, Mihai	
Latest Release	-	
TRL	6	
Location	data hosted by other non-HBP 3rd party	[Petrovici et al., 2016], [Leng et al., 2017]
Maintenance	Active development WP9.4, CDP5	
Validation - QC	-	
Validation – Users	-	
Validation – Publications	-	

Table 12: Component 2474 Software model of sampling with LIF neurons

6.11 Software model of noise generation by a balanced random spiking network

Field Name	Field Content	Additional Information
ID	2475	
Component Type	Model	
Contact	PETROVICI, Mihai	
Latest Release	-	
TRL	5	
Location	data hosted by other non-HBP 3rd party	arXiv preprint arXiv:1710.04931 [Jordan et al. 2017]
Maintenance	Active development WP9.4	
Validation – Users	-	
Validation – Publications	-	

Table 13: Component 2475 Software model of noise generation by a balanced random spiking network

6.12 Hardware emulation of LIF sampling with subnetwork modules

Field Name	Field Content	Additional Information
ID	2476	
Component Type	Model	
Contact	PETROVICI, Mihai	
Latest Release	-	
TRL	5	
Location	data hosted by other non-HBP 3rd party	[Petrovici et al., 2017b] https://doi.org/10.1109/ISCAS.2017.8050530
Maintenance	Active development WP9.4	
Validation - QC	-	
Validation – Users	-	
Validation – Publications	-	

Table 14: Component 2476 Hardware emulation of LIF sampling with subnetwork modules

6.13 Hardware emulation of hierarchical sampling networks

Field Name	Field Content	Additional Information
ID	2477	
Component Type	Model	
Contact	PETROVICI, Mihai	
Latest Release	-	
TRL	3	
Location	data hosted by other non-HBP 3rd party	[Petrovici et al., 2017a], [Schmitt et al., 2017]
Maintenance	Active development WP9.4, CDP5	
Validation - QC	-	
Validation – Users	-	
Validation – Publications	-	

Table 15: Component 2477 Hardware emulation of hierarchical sampling networks