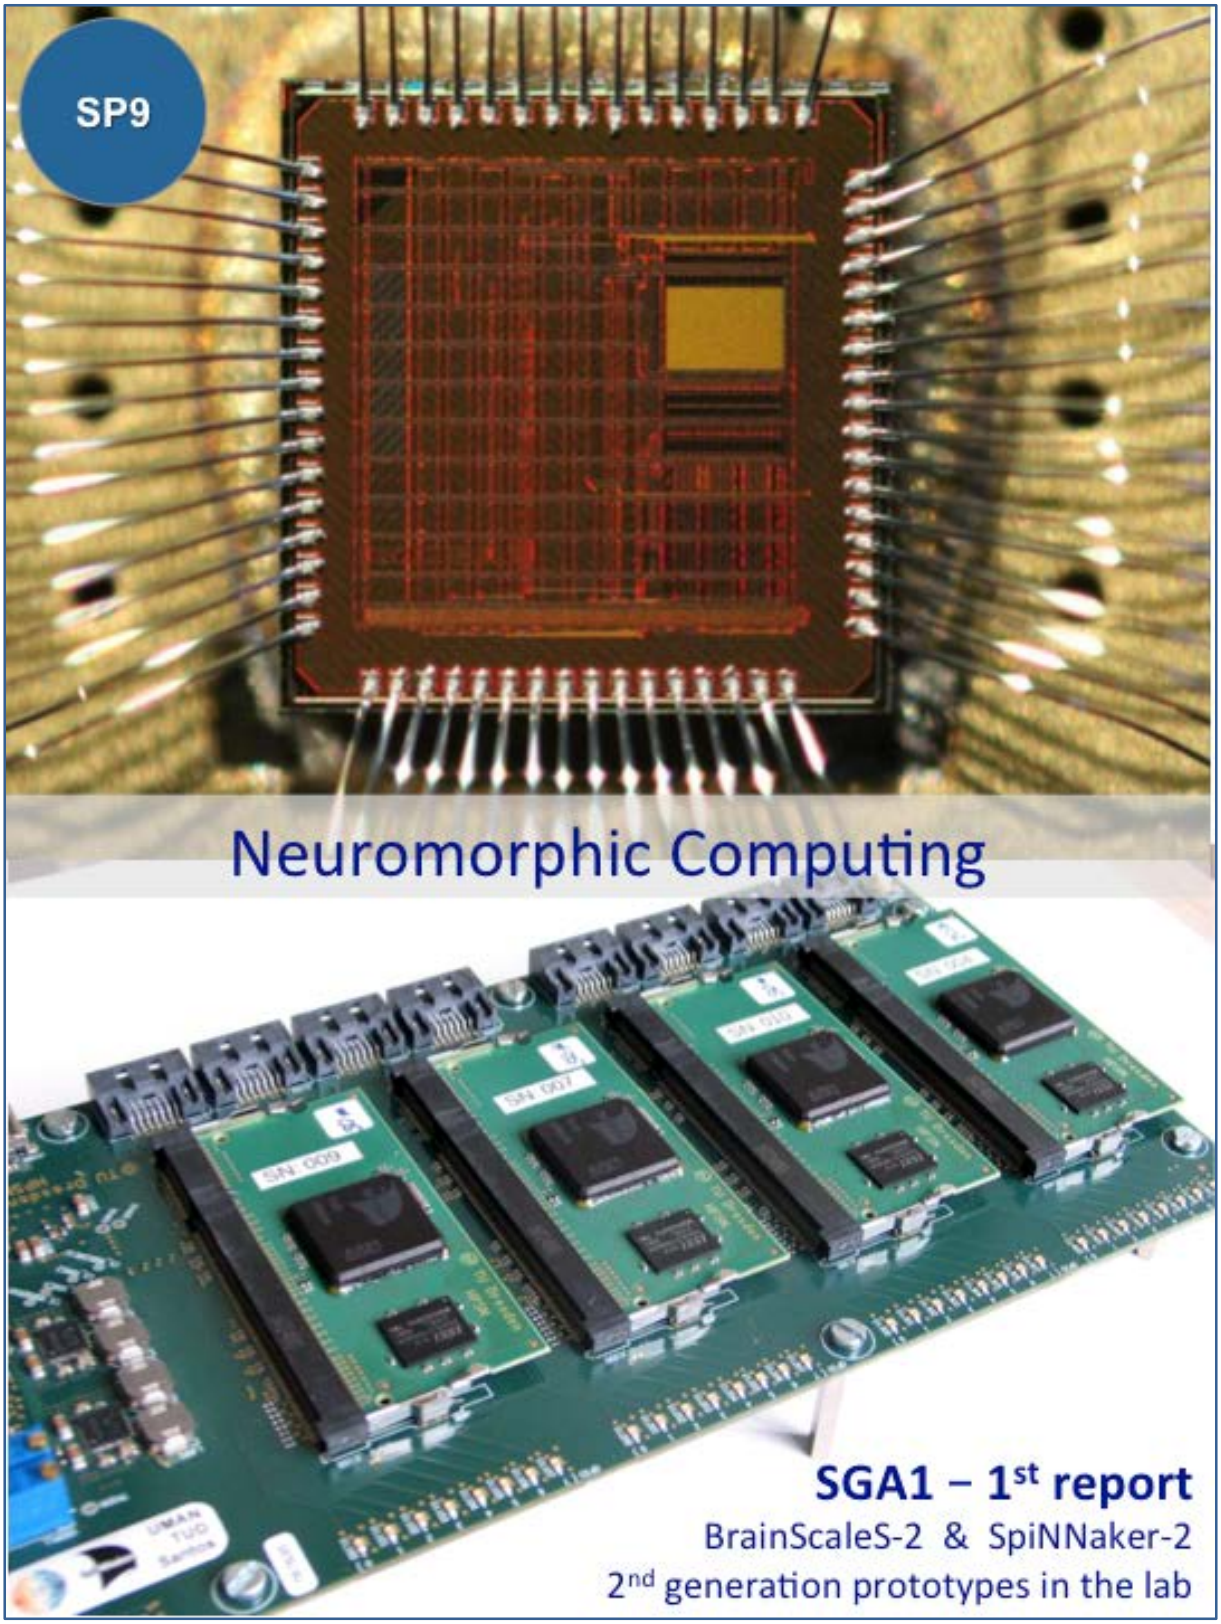
SP9
Neuromorphic Computing
SGA1 – 1st report
BrainScaleS-2 & SpiNNaker-2
2nd generation prototypes in the lab

Project Number:	720270	Project Title:	Human Brain Project SGA1

Document Title:	D9.5.1: SP9 Neuromorphic Computing Platform - Results for SGA1 Period 1
Document Filename:	SP9 D9.5.1 FINAL
Deliverable Number:	SGA1 D9.5.1
Deliverable Type:	Report
Work Package(s):	WP9.1, WP9.2, WP9.3, WP9.4, WP9.5
Dissemination Level:	PU (= public)
Planned Delivery Date:	SGA1 M12 / 31 March 2017
Actual Delivery Date:	SGA1 M12 / 27 March 2017 with an update May 2017 (Delivery Date 23 May 2017)

Authors:	SP9 Task-, WP- and SP-leaders
Compiling Editors:	
Contributors:	SP9 Task-, WP- and SP-members
SciTechCoord Review:	EPFL (P1): Jeff MULLER, Martin TELEFONT, Marie-Elisabeth COLIN UHEI (P47): Martina SCHMALHOLZ, Sabine SCHNEIDER
Editorial Review:	EPFL (P1): Guy WILLIS, Colin McKINNON

Abstract:	Progress report from the first HBP SGA1 year (1 April 2016 – 31 March 2017) by the Neuromorphic Computing Platform teams
Keywords:	Neuromorphic computing, SP9, SpiNNaker, BrainScaleS, physical model, many-core, computational principles, applications, training, innovation

Table of Contents

List of Figures

List of Tables

1. SP Leader's Overview

1.1 Key Personnel

Subproject Leader: Karlheinz MEIER (UHEI)

Subproject Deputy Leader: Steve FURBER (UMAN)

Subproject Manager: Björn KINDLER (UHEI)

1.2 Progress

The HBP Neuromorphic Computing Platform (SP9) offers access to a unique set of 2 complementary neuromorphic machines for modelling neural microcircuits and applying brain-like principles in machine learning and cognitive computing. For this purpose, it is developing fundamental design principles, neuromorphic chips, systems and the required support software. Users of the Platform receive support through training and access to software tools and hardware systems.

The Subproject is structured into 5 Work Packages.

SP9 made very good progress according to the SGA1 work plan in all areas of planned work (software, hardware, platform operation and maintenance, computational theory and applications).

First generation platform machines

The neuromorphic platform machines in Manchester (SpiNNaker-1) and Heidelberg (BrainScaleS-1) are readily accessible through the Collaboratory as well as through direct local job submission. Commissioning of the 1st generation BrainScaleS-1 system has progressed well and the SpiNNaker-1 system with 0.5 Million ARM cores is fully operational. Both systems are currently used for experiments.

Software integration into the HBP Platform infrastructure (Collaboratory) was successfully carried forward according to plan.

Second generation prototype chips and systems

SP9 differs from all other HBP Platforms in that is has very significant activity in hardware development, including custom chip design. An essential advantage of the HBP over other projects is the close collaboration between computer architects and neuroscientists. Here, CDP5 and the dedicated theory Work Package within SP9 have been of key importance. A prototype chip for the BrainScaleS-2 system has been submitted to manufacturing. It includes a multitude of innovative circuits such as flexible hybrid plasticity, multi-compartment neurons with non-linear dendrites and an improved technology for parameter storage. Testing of the next generation SpiNNaker-2 prototype chip (Santos) for the next generation SpiNNaker system has progressed well and delivered the expected results. The finalisation of first 2nd generation single chip systems by the end of SGA1 is well under way and will open a new way towards neuromorphic computing.

The work towards full-scale next-generation systems is fully in line with the chip and system design roadmap defined in the HBP FPA as shown in all recent project reviews. Building a neuromorphic system from the chip upwards is a 10-year effort that has been carefully planned at the beginning of HBP. No technological hurdles have so far been identified that would impede this development. It will now be critical to provide the required financial resources during SGA2, in particular for the hardware construction.

Computational principles

An important aspect of SP9 is the integration of research on theoretical principles for neuromorphic computation into the subproject proper. This turned out to be a success model with an easy and efficient exchange of ideas in a non-bureaucratic way. The regular Fürberg

workshops going back to the BrainScaleS project tradition are now established as an exchange forum especially for students who receive substantial inspiration from there. Neural sampling with spiking neurons made it from a theoretical principle to a real hardware implementation (published in this reporting period) with substantial catalysing effects by the Fürberg workshops.

Applications and Innovation

Integration of former SP11 (application) groups who joined HBP through an open call has worked well and a variety of applications and benchmark developments are under way. Together with the FET unit SP) has organized an innovation workshop in Brussels to pave the way for joint projects of industry and academia.

Strategy document for neuromorphic computing

Towards the next phase of HBP (SGA2), SP9 has delivered a comprehensive strategy document for neuromorphic computing in HBP. That document is attached as an Appendix.

Outreach and dissemination

In the first project year, SP9 partners conducted 77 disseminations for different audiences (including scientific audiences, industry, media and the general public). Invited presentations at industry include companies like apple, google and SAP.

1.3 Deviations

Some SP9 groups reported delays in hiring staff because of the very late formal approval of SGA1 funding 6 months after project start. For legal reasons many European universities cannot sign contracts without a formally operational and funded project. It appears that no significant delays have been accumulated by this problem but discontinuities in the funding pose a substantial risk for the Subproject and have to be avoided by all means.

1.4 Impact of work done to date

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.

With the 2nd generation prototype chips submitted and likely to be successfully tested during Y2 of SGA1, European neuromorphic computing made a major step forward to maintain its leading global role. The currently operational Platform machines are still based on chip designs from 2005. The new 2nd generation chips are genuine HBP work based on collaboration within the Project. A recent conference in the US (NICE2017) has shown that the HBP prototypes are among very few new neuromorphic systems under production at this point in time. It should be noted that DARPA (US) has just called for proposals for systems carrying out “life-long-learning” as a strong competitor for the European activities. In conclusion, the impact is high but there is growing competition. Here again, the close integration of chip design with theory through WP9.4 has been and will be an important ingredient to remain competitive on the European side. It will be essential to provide the necessary and well defined funding to realise the next generation machines as a major contribution to the HBP Platform infrastructure.

The transfer of research and development results obtained in the HBP has been explored in a dedicated workshop organised by the FET unit in Brussels on February 3rd, 2017. The involvement of other funding schemes like ECSEL (Electronic Components and Systems for

European Leadership) have been identified as a promising way to achieve an impact of neuromorphic computing on the industrial scale.

1.5 Priorities for the remainder of the SGA1 phase

SP9 has a clearly defined priority for the remainder of the SGA1 phase:

Integrating fully functional 2nd generation prototype chips into systems ready for external user evaluation and smaller-scale experiments and applications.

The next generation chip will have several new features that make them much more powerful than the existing ones and highly competitive in the international comparison:

- *BrainScaleS-2* will feature a flexible on-chip plasticity-processor for various synaptic and neuronal local learning mechanisms, the capability for on-the-fly network reconfiguration, multi-compartment neurons and non-linear, active dendrites.
- *SpiNNaker-2* will employ around 32 quad-core Quad Processing Elements (QPEs), delivering 25 GIPS at 1Watt on a single die, with floating point precision and an on-chip true random generator.

It will be essential to give small-scale but fully functional systems based on the new chips to users by the end of SGA1 (March 2018). **This is certainly the most important goal for SP9.** Reaching this goal in SGA1 will allow for the next steps towards 2nd generation large-scale machines currently planned for construction in SGA2 and SGA3.

In addition the further development and upgrades for the existing machines including software tools for their operation will also be an important objective for the next phase of SGA1.

2. WP 9.1: Platform Software and Operations

2.1 Key Personnel

Work Package Leader: Andrew P. DAVISON (CNRS)

2.2 WP Leader's Overview

What went particularly well?

The first year of SGA1 was primarily a period of consolidation, with the tools developed during the Ramp-Up Phase being adopted by an increasing number of users (e.g. number of users accessing the Platform remotely increased from 13 to 30) and new developers beginning to contribute to the Platform. A particular highlight was the adoption of the MUSIC library into tools developed by other SPs, such as the Closed-Loop Engine developed in SP10 and the Monsteer visualisation tool in SP7. Considerable progress has also been made on a framework for performance benchmarking of the neuromorphic systems.

What didn't go according to plan?

Hiring of new personnel was delayed by the extended contract negotiations. This reduced the velocity of development in the first 12 months; however, we expect this delay to be caught up by Month 24.

Impact of work done

By operating and improving the public interface to the Platform and its software underpinnings we have made two cutting-edge neuromorphic computing systems available to a general scientific public.

2.3 Priorities for the remainder of the phase

In the remainder of SGA1 our priorities are

- to operate and maintain the Platform, with a particular emphasis on improved status monitoring, extensive benchmarking of system performance and on provenance tracking to support reproducible research.
- to lay the groundwork for the next generation of neuromorphic systems, with a particular emphasis on support for modelling of synaptic plasticity and on interoperability with the Brain Simulation and Neurorobotics platforms.

2.4 T9.1.1 Platform Operation and Maintenance

2.4.1 Key Personnel

Task Leader: Andrew P. DAVISON (CNRS)
Other Researcher: Ulrich RÜCKERT (UNIBI)
Other Researcher: Eric MÜLLER (UHEI)
Other Researcher: David R. LESTER (UMAN)
Other Researcher: Christian MAYR (TUD)

2.4.2 SGA1 DoA Goals

Operate and maintain the Neuromorphic Computing Platform hardware and software constructed in the HBP Ramp-Up Phase.

2.4.3 Component Progress

2.4.3.1 Neuromorphic Job Queue Service

The Neuromorphic Job Queue Service is a web service providing a REST (representational state transfer) API for submitting jobs to the BrainScaleS and SpiNNaker systems and for retrieving job results. It is developed by CNRS.

Progress on Component:

The service was operated continuously with almost no downtime. A number of bugs, identified by users and by extension of the test suite, were fixed. The source code was moved to a public repository at:

https://github.com/HumanBrainProject/hbp_neuromorphic_platform

Both planned formal releases occurred on time: M5 update of authorisation to support SGA1 accreditation; M12 augmented notifications, with notification e-mails including a link to the relevant Collab and the last few lines of the logs produced by the Neuromorphic systems. In addition, several new REST endpoints were added to support the functionality of the Neuromorphic Dashboard app (see below).

Quality Control:

- Upstream Component "HBP Identity Service"
  Task responsible: T11.3.2
  Status/quality: stable and reliable
- Upstream Component "Collaboratory Service"
  Task responsible: T11.3.2
  Status/quality: stable and reliable
- Upstream Component "Collaboratory Storage Service"
  Task responsible: T11.3.2
  Status/quality: stable and reliable
- Upstream Components "Neuromorphic jobs database", "Neuromorphic Quotas service",
  Task responsible: this Task (T9.1.1)
  Status/quality: see detailed reports below
- Upstream Component "PyNN",
  Task responsible: T9.1.4
  Status/quality: see detailed report under Task T9.1.4 below
- Upstream Component "Upload caches & services"
  Tasks responsible: T7.5.2, T7.5.3, T7.5.6
  Status/quality: the UNICORE REST API allows data to be copied from the Neuromorphic systems to HPAC machines for which the user has an account. This works reliably.
- Upstream Component "SpiNNaker Neuromorphic Computing System",
  Tasks responsible: T9.3.1, T9.3.2, T9.3.3, T9.1.1
  Status/quality: the SpiNNaker system support software regularly polls the Job Queue server, and posts back the results of completed jobs. This subcomponent has been reliable, and the occasional bug rapidly fixed.
- Upstream Component "BrainScaleS 1 Neuromorphic Computing System",
  Tasks responsible: T9.2.1, T9.2.3, T9.1.1
  Status/quality: the BrainScaleS system support software regularly polls the Job Queue server, and posts back the results of completed jobs. This subcomponent has been reliable, and the occasional bug rapidly fixed, except for one period of a few weeks when the support systems were being upgraded and the service was unavailable.

- Downstream Components: “Bio-mimetic classifier” (T9.5.4), “Neuromorphic Benchmarks service” (T9.1.3), “Neuromorphic Dashboard app” (T9.1.1), “Python client for the Neuromorphic Computing Platform” (T9.1.1), “Neuromorphic Job Manager app”.

  Status: the Neuromorphic Job Queue Service was provided to the above components during the Ramp-Up Phase and has been continuously available, with improvements, since then.

2.4.3.2 Neuromorphic Job Manager app

The Neuromorphic Job Manager app is a Collaboratory app providing a graphical interface for submitting jobs to the BrainScaleS and SpiNNaker systems and for retrieving job results. It is a graphical front-end to the Neuromorphic Job Queue Service. It is developed by CNRS.

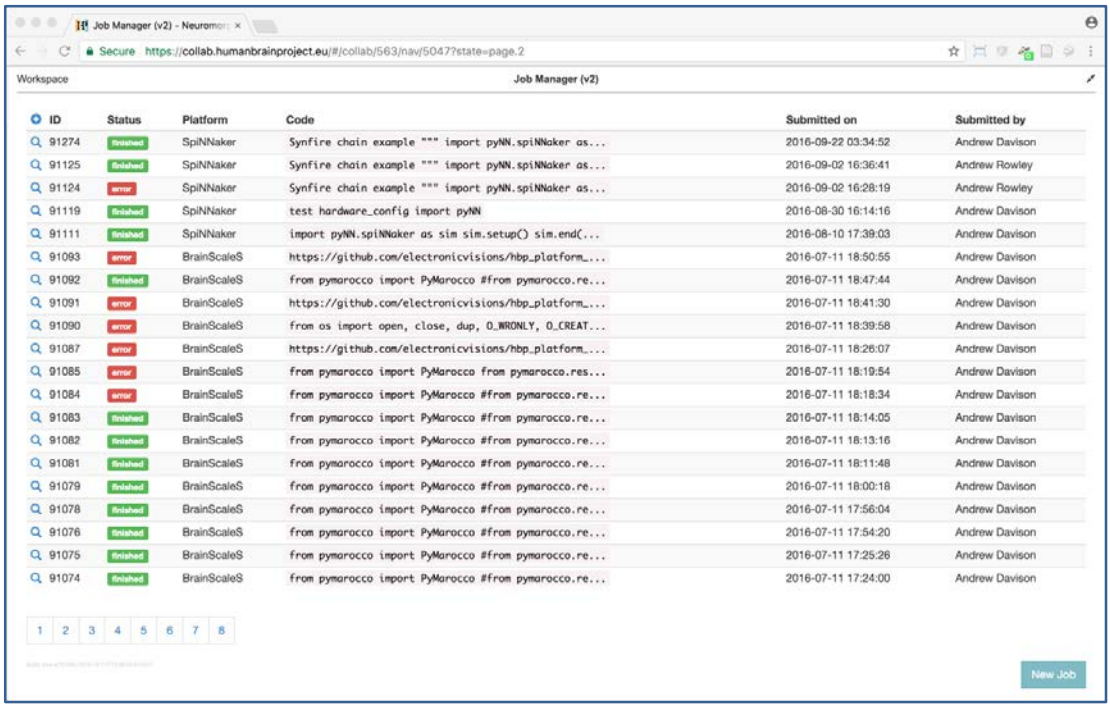

Figure 1: Neuromorphic Job Manager app

Progress on Component:

The service was operated continuously with almost no downtime. A number of bugs, identified by users and by extension of the test suite, were fixed. The app was updated to keep track with improvements to the underlying JavaScript libraries. The source code was moved to the same public repository as the Job Queue Service (see above).

Both planned formal releases occurred on time: M6 - added deep linking, so that state is not lost on page refresh, and job detail pages can be bookmarked. M12 - allow scripts to be loaded from the Collaboratory Storage Service (previously, scripts could only be pasted in directly or obtained from a public Git repository). Additional improvements were rolled out as they became available, including the ability to delete jobs and to edit/resubmit failed jobs.

Quality Control:

- Upstream Component “Collaboratory Service”
  Task responsible: T11.3.2
  Status/quality: stable and reliable
- Upstream Component “Collaboratory Storage Service”
  Task responsible: T11.3.2
  Status/quality: stable and reliable

- Upstream Component "Provenance viewer for the Neuromorphic Job Manager app"
  Task responsible: T9.1.1, T9.1.2
  Status: Work on this Component is planned to start in the second half of SGA1.
- Upstream Component "Neuromorphic Job Queue Service"
  Task responsible: T9.1.1
  Status/quality: see detailed report above.
- Downstream Components: "Workflow for comparison of spiking neuron data obtained from NEST and SpiNNaker simulation", "SP9 Platform Training events".

  Status: the Neuromorphic Job Manager app was provided to the above Components during the Ramp-Up Phase and has been continuously available, with improvements, since then.

2.4.3.3 Neuromorphic Dashboard app

The Neuromorphic Dashboard app is a Collaboratory app providing graphs of statistics about the Neuromorphic Computing Platform, such as the length of the job queues, the number of completed jobs for each system, and the distribution of job durations. It is developed by CNRS.

Progress on Component:

The planned M12 release took place considerably ahead of schedule, during M2, and added job duration histograms and a graph of queue-lengths.

Quality Control:

- Upstream component "Neuromorphic Job Queue Service"
  Task responsible: T9.1.1
  Status/quality: see detailed report above.
- Downstream Components: None. The app is intended for direct use by platform users.

Figure 2: Neuromorphic Dashboard app

2.4.3.4 Python client for the Neuromorphic Computing Platform

The Python client allows programmatic interaction with the Neuromorphic Job Queue Service and the Neuromorphic Quotas Service. It can be used in Jupyter notebooks, Python scripts or in a Python console. The client is distributed via the Python Package Index (PyPI)

with the package name "hbp_neuromorphic_platform", under the Apache License. The source code is available at https://github.com/HumanBrainProject/hbp-neuromorphic-client. It is developed by CNRS and UHEI.

Progress on Component:

Two releases were made during the reporting period:

- Version 0.4.3 – 17/06/2016 – added the ability to copy data produced by neuromorphic jobs to Collaboratory storage or to the HPAC Platform.
- Version 0.5.0 – 14/03/2017 - added methods for working with the Quotas service, e.g. making a request for a resource allocation, listing current quota usage.

In addition, continuous integration was set up, so that unit tests are automatically run on each commit.

Quality Control:

- Upstream Component "Neuromorphic Job Queue Service"
  Task responsible: T9.1.1
  Status/quality: see detailed report above.
- Downstream Component "Neuromorphic Benchmarks service" (T9.1.3)
  Status: the Python client was already available when work on this Component began. See evaluation under T9.1.3 below.
- Downstream Component "BrainScaleS frontend software (T9.2.3)
  Status: these two Components were integrated during the Ramp-Up Phase.

2.4.3.5 Neuromorphic Quotas service

The Neuromorphic Quotas Service is a web service providing a REST API for managing requests for computing resources on the BrainScaleS and SpiNNaker systems, and tracking the usage of granted resources. It is developed by CNRS.

Progress on Component:

The service was operated continuously with almost no downtime. A number of bugs, identified by users and by extension of the test suite, were fixed. The source code was moved to a public repository at

https://github.com/HumanBrainProject/hbp_neuromorphic_platform

The planned formal release occurred on time: M12 - notification of user when an access request is accepted or a quota added.

Quality Control:

- Upstream Component "HBP Identity Service"
  Task responsible: T11.3.2
  Status/quality: stable and reliable
- Upstream Component "Collaboratory Service"
  Task responsible: T11.3.2
  Status/quality: stable and reliable
- Upstream Components "Neuromorphic jobs database"
  Task responsible: this Task (T9.1.1)
  Status/quality: see detailed report below
- Downstream Components: "Neuromorphic Job Queue Service" (T9.1.1), "Neuromorphic Resource Manager app" (T9.1.1).
  Status: the Neuromorphic Quotas Service was provided to the above Components during the Ramp-Up Phase and has been continuously available, with improvements, since then.

2.4.3.6 Neuromorphic Resource Manager app

The Neuromorphic Resource Manager app is a Collaboratory app providing a graphical interface for managing requests for computing resources on the BrainScaleS and SpiNNaker systems, and tracking the usage of granted resources. It is a graphical front-end to the Neuromorphic Quotas Service. It is developed by CNRS.

Progress on Component:

The service was operated continuously with almost no downtime. A number of bugs, identified by users and by extension of the test suite, were fixed. The app was updated to keep track with improvements to the underlying JavaScript libraries. The source code was moved to the same public repository as the Job Queue Service (see above).

No formal releases were planned in the first 12 months of SGA1.

Quality Control:

- Upstream Component "Neuromorphic Quotas service"
  Task responsible: this Task (T9.1.1)
  Status/quality: see detailed report above.
- Downstream Components: None. The app is intended for direct use by platform users.

2.4.3.7 Neuromorphic jobs database

The Neuromorphic Job Queue and Quotas services use a common object-relational PostgreSQL database, backed-up hourly. Restoring from backup is semi-automated, and is checked regularly. CNRS maintains it.

Progress on Component:

The database operated continuously with almost no downtime.

Quality Control:

- Downstream Components: "Neuromorphic Job Queue Service", "Neuromorphic Quotas Service" (T9.1.1).
  Status: the jobs database was provided to the above Components during the Ramp-Up Phase and has been continuously available since then.

2.4.3.8 BrainScaleS-1 Neuromorphic Computing System

The neuromorphic computing hardware physical model system ("BrainScaleS") in Heidelberg (big system). System version 1 (=NM-PM1 systems, 20 cm wafers). Task T9.1.1 is responsible for the ongoing operation of the first generation system. The system is operated by UHEI.

Progress on Component:

See report under T9.2.3 below.

2.4.3.9 BrainScaleS system software

See report under T9.2.3 below.

2.4.3.10 BrainScaleS frontend software

See report under T9.2.3 below.

2.4.3.11 User Guidebook for the Neuromorphic Computing Platform

See report under T9.5.2 below.

2.5 T9.1.2 Platform Integration with HBP Collaboratory and other Platforms

2.5.1 Key Personnel

Task Leader: Andrew P. DAVISON (CNRS)
Other Researcher: Mikael DJURFELDT (KTH)
Other Researcher: Eric MÜLLER (UHEI)
Other Researcher: David R. LESTER (UMAN)

2.5.2 SGA1 DoA Goals

1) To extend the existing Platform interface to ensure that the Neuromorphic Computing Platform is fully integrated with the HBP Collaboratory.
2) To implement integration with the Neuroinformatics Platform to ensure complete provenance tracking and correct registration of models and datasets.
3) To implement real-time communication protocol(s) and support software using the MUSIC library to enable spike-based communication between the neuromorphic systems and external software, in particular the software of the Neurorobotics Platform.

2.5.3 Component Progress

2.5.3.1 Components described under Task 9.1.1

A number of components receive contributions from both Tasks T9.1.1 and T9.1.2. The contributions of Task T9.1.2 relate specifically to features that involve interaction with the other HBP Platforms. The following components are described above under Task T9.1.1:

- "Neuromorphic Resource Manager app"
- "Neuromorphic Quotas service"
- "Neuromorphic Job Manager app"
- "Neuromorphic Job Queue Service".

2.5.3.2 Provenance viewer for the Neuromorphic Job Manager app

This Component adds a viewer for provenance information (i.e. software and hardware versions, hardware operating conditions, source of input data, etc.) to the Neuromorphic Job Manager app.

Progress on Component:

Work on this Component is scheduled to begin in the second year of SGA1.

2.5.3.3 KnowledgeGraph integration for neuromorphic simulations

Enable users of the Neuromorphic Computing Platform to register their models and results with the Neuroinformatics Platform KnowledgeGraph.

Progress on Component:

Work on this Component is scheduled to begin in the second year of SGA1.

2.5.3.4 MUSIC library

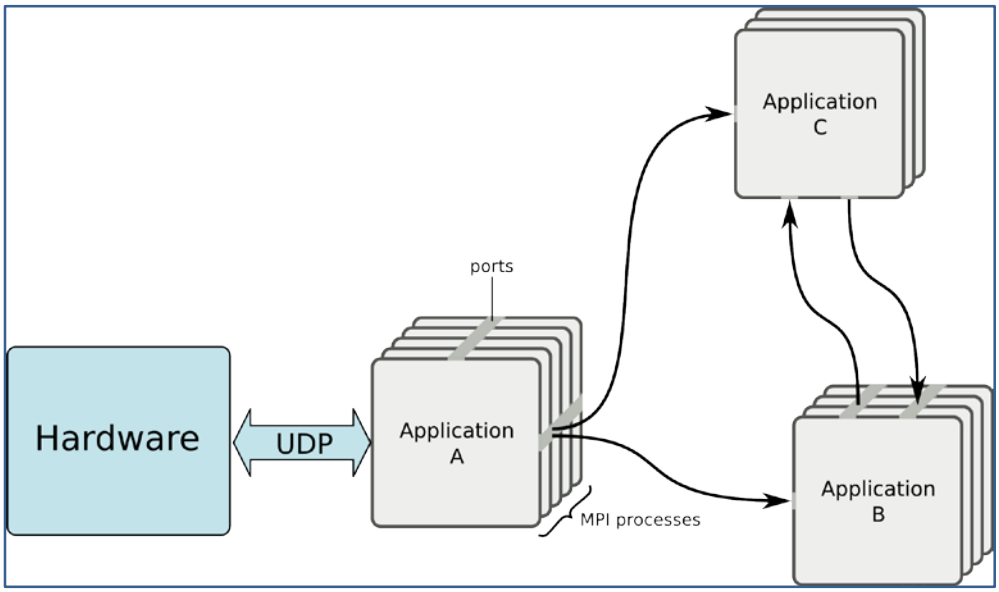

Figure 3: Three parallel applications A, B and C communicate through MUSIC ports.

Each gray square represents one process. Application A communicates with hardware using UDP.

MUSIC (Djurfeldt et al. 2010) is a software framework for on-line communication of spike events and other types of data between simulation tools. It also supports modularity in that different tools and models can easily be connected together to form larger systems without the need to know details of connected components. Figure 3 shows an example setup where three parallel applications are executed together while one of them communicates with external hardware. The applications could be neuronal network simulators responsible for different modules of a large-scale brain simulation while the hardware could be a robot. Figure 4 shows further use cases for MUSIC.

In HBP, the CLE (SGA1 task 10.5.2; Closed-Loop Engine) of the Neurorobotics Platform uses MUSIC for communication with simulators (Figure 4D) and will use it for communication with neuromorphic hardware (Figure 4C). MUSIC is also suitable for interactive visualisation (Figure 4B; SGA1 tasks T7.3.1, T7.3.2; Monsteer), especially in combination with large-scale simulations, and in studies of closed-loop systems. Work on combining neurons from different simulators in a single multi-scale network (Figure 4A) is planned in SGA2 T4.4.3.

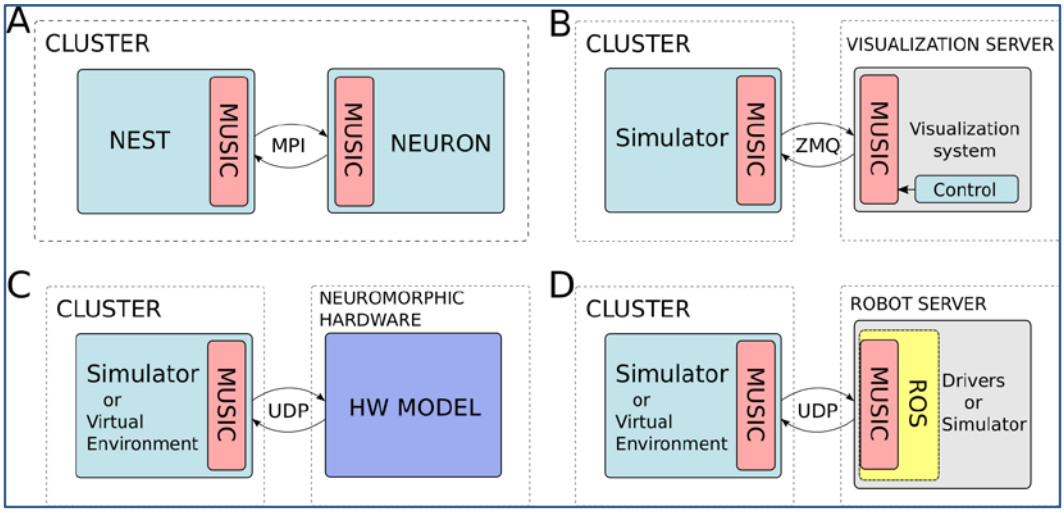

Figure 4: Example usage scenarios for MUSIC.

A. Neurons simulated by different simulators (here NEST and NEURON) can form a single network. B. Data shipped out of simulator for interactive visualization. C. Model or virtual environment connects to

hardware. D. Model or virtual environment connects to robot. (MPI presently used instead of ZMQ and UDP in B and D.)

Progress on Component:

During this reporting period, the KTH partner has assisted members of SP10 (T10.5.2) in incorporating MUSIC as a middleware in the CLE and have collaborated further with Philipp Weidel at Simlab Neuroscience, F.Z. Juelich, in developing and benchmarking a MUSIC-based toolchain which allows MUSIC-aware simulators to talk to simulated or real robots through ROS (Weidel et al. 2016). As part of the work for MS9.1.3, a prototype MUSIC-SpiNNaker interface, which allows populations of spiking neurons in the NM-MC-1 system to be abstracted as MUSIC ports, has been developed. The next step is to integrate it into the SpiNNaker software stack such that it will become available to HBP users. During the period, the KTH partner has also continued MUSIC development, including further work on a new communication algorithm as well as benchmarking.

MUSIC currently ships data between simulation tools which is sufficient for integration of spiking networks and for streaming data out of a simulation. However, in some multi-scale co-simulations the situation arises where variables solved by numerical solvers needs to be connected into a larger system. As an example, a solver for compartmentalized neurons might provide a Ca-signal to a simulator for reaction-diffusion systems, which might, in turn, send a signal back about the phosphorylation state of ion channels. It turns out that doing this naively can lead to several problems such as unexpected stiffness and instability.

As an in-kind contribution, the KTH partner has worked on theory of numerical integration of multi-scale systems with the aim to extend MUSIC to allow for correct integration of numerical solvers. Two papers on this work has been published (Brocke et al. 2016a, 2016b).

Quality Control:

- Downstream Component: “Closed-loop engine”
  Task responsible: T10.5.2
  Status: MUSIC is stable/reliable but is in need of better documentation and more user-friendly error reporting. Also, for the purposes of the CLE, a new MUSIC API is needed to make MUSIC more flexible and to allow for starting/stopping simulations and dynamic reconfiguration of the system.
- Downstream Component “SP9 SpiNNaker MUSIC integration”
  Task responsible: T9.3.2
  Status: Software is in prototype state and needs to be integrated into the SpiNNaker software stack.
- Downstream Component “Monsteer”
  Task responsible: T7.3.1, T7.3.2
  Status: MUSIC is stable/reliable. Monsteer is an adapter between MUSIC-aware simulators, such as NEST, and socket based communication, suitable for shipping spikes out of a supercomputer. For the purposes of visualisation, it would be good if a future MUSIC version could directly support, e.g., ZMQ.
- Downstream Component “PyNN”
  Task responsible: T9.1.4
  Status: The PyNN/MUSIC interface, which makes it possible to set up MUSIC co-simulations in an easy and friendly way from within PyNN, currently has a bug. This needs to be resolved and the interface integrated into the PyNN master branch.

2.6 T9.1.3 Platform Performance Benchmarks

2.6.1 Key Personnel

Task Leader: Ulrich RÜCKERT (UNIBI)
Other Researcher: Andrew P. DAVISON (CNRS)
Other Researcher: Eric MÜLLER (UHEI)
Other Researcher: David R. LESTER (UMAN)

2.6.2 SGA1 DoA Goals

1) The development and maintenance of performance benchmarks, based on Use Cases from neuroscience and machine learning.
2) Implementation and operation of a continuous integration system providing regular runs of the entire benchmark suite.

2.6.3 Component Progress

During the first semester (M1 to M6) of SGA1 we have developed the software abstraction layer Cypress that hides the remaining differences between the PyNN-based software interfaces to the different neuromorphic hardware and simulation platforms. These differences are expected to become smaller, once the feature set of platform-specific software support stacks stabilises, but will not vanish completely, because important universally available PyNN functions need to be supplied with platform-specific parameters or configuration data. The Cypress layer has been successfully used to run our previously developed associative memory benchmarks unchanged on the NM-PM1, NM-MC1, and Spikey hardware platforms.

Upon this foundation, we have built the modular benchmark framework SNABSuite, shown in Figure 5. This framework functions as a harness where one instance automatically runs a set of benchmarks on a given target platform and collects all results in a format suitable for import into the benchmarks result database (densely dotted arrows in Figure 5). The individual benchmarks inside this framework use a shared implementation across all target platforms to render benchmark results comparable, not only with former results obtained on the same target platform, but also with performance results for other platforms, to track changes in relative performance differences.

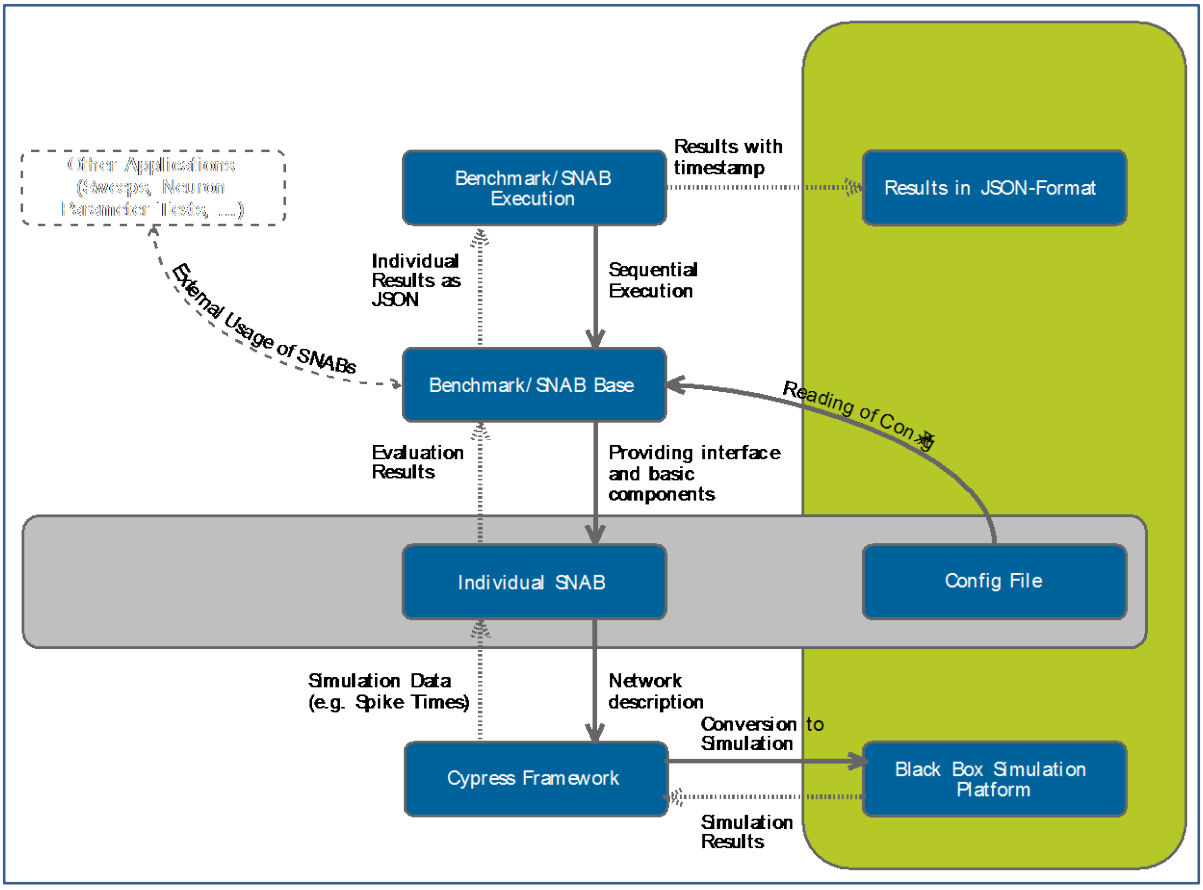

Figure 5: Flow chart of the benchmark framework SNABSuite

The dashed parts in the upper-left corner are currently under development. Addition of a new benchmark requires only the two components shown in the grey box in the centre.

The SNABSuite framework supports platform-specific parameterisation of benchmarks (solid arrows in Figure 5), for example to express valid ranges for hardware parameters involved in parameter sweeps. The same mechanism can be used to limit (or even exclude) execution of specific benchmarks on certain platforms so that a continuous automated tracking of platform performance is not hampered by temporary hardware or software problems or by known defects. On the other hand, scalability of suitable benchmarks will also be achieved by using SNABSuite parameters to express, for example, neuron population sizes, multiplicity, or connection degrees. In this way, small-scale as well as large-scale platforms can each be exercised up to their maximum neuron or synapse capacity. For composite platforms, intermediate parameter sets are added that match the capacity of hierarchical substructures of the platform. This helps to identify performance differences caused by communication mechanisms between those substructures.

From a software-technical perspective, the SNABSuite framework moves typical administrative parts surrounding the benchmark code, which would otherwise be repeated in every individual benchmark, into a common superclass and provides a library of reusable support functions. This reduces the cognitive load and the effort required by benchmark authors to add new benchmarks to the suite and ensures consistent interaction between SNABSuite clients, like the benchmarks service, and the individual benchmarks contained in the suite.

As of M10, we are currently integrating our associative memory benchmarks into the SNABSuite framework and have in the process validated the design of the framework as well as fine-tuned the division of responsibilities between the shared framework code and the individual benchmarks. The next steps are to integrate further realistic application benchmarks to the benchmark suite. These are based on essential parts of (currently often

platform-specific) applications developed in SP9 and on typical computational tasks that occur in many neuromorphic applications.

The parameter mechanism of SNABSuite has been successfully used to integrate an initial set of synthetic benchmarks that check platform performance under workloads with extreme properties or measure platform-/machine-specific parameters. Beyond uncovering performance bottlenecks and stress-testing specific components of hardware platforms, the results of such benchmarks will help prospective users of neuromorphic computing platforms to pick the best-suited platform for the specific requirements of the user application.

2.6.3.1 Neuromorphic Benchmarks repository

Description: Python scripts defining a set of PyNN network models and benchmark measurements aiming to measure the performance of the PM1 and MC1 hardware designs.

Progress on Component: The release of the first set of benchmarks is imminent as the primary Component of the M12 Milestone MS9.1.1. The benchmark framework has been developed by partner UNIBI, as well as the synthetic benchmarks for characterising hardware platforms and the BiNAM associative memory benchmarks. Key components of application benchmarks and platform-specific development support has been provided by partners UHEI and UMAN.

Quality Control:

- Upstream Component "Bio-mimetic classifier"
  Task responsible: T9.5.4
  Status: no release has been received yet (matches planning)
- Upstream Component "PyNN"
  Task responsible: T9.1.4, T6.4.3
  Status: intermediate release has been received
  Quality of upstream Component: near-release quality with very few remaining bugs and some opportunities for improvements to the programming interface (API)
- Upstream Component "SP9 BrainScaleS frontend software"
  Task responsible: T9.2.3, T9.2.2
  Status: intermediate release has been received
  Quality of upstream Component: usable for prototyping and testing purposes with some remaining bugs and known limitations in currently implemented functionality
- Upstream Component "SP9 SpiNNaker software stack"
  Task responsible: T9.3.2
  Status: intermediate release has been received
  Quality of upstream Component: fully usable for prototyping and testing purposes with only few remaining bugs and some known limitations in currently implemented functionality
- Downstream Component "SP9 BrainScaleS frontend software"
  Task responsible: T9.2.3, T9.2.2
  Status: bug fixes/patches or requests/suggestions for enhancements have been provided
- Downstream Component "SP9 SpiNNaker software stack"
  Task responsible: T9.3.2
  Status: bug fixes/patches or requests/suggestions for enhancements have been provided
- Downstream Component "SP9 Neuromorphic Benchmarks results database"
  Task responsible: T9.1.3
  Status: results for the currently implemented set of benchmarks are routinely generated in the format required for importing into the database
- Downstream Component "SP9 Neuromorphic Benchmarks app"

Task responsible: T9.1.3

Status: example benchmarks have been integrated into the app for testing

2.6.3.2 Neuromorphic Benchmarks service

Description: Continuous integration system providing regular runs of the entire neuromorphic benchmark suite to track changes of platform performance over time.

CDP Contributions: none

Progress on Component: This component will be released as part of MS9.1.2 in M12. It forms the foundation for the benchmarks app that is currently undergoing testing. The service itself has been developed by partner CNRS with inputs, testing feedback and platform-specific support provided by partners UNIBI, UHEI and UMAN.

Quality Control:

- Upstream Component "SP9 BrainScaleS standalone next generation single chip physical model system"
  Task responsible: T9.2.1, T9.2.4
  Status: no release has been received yet (matches planning)
- Upstream Component "SP9 SpiNNaker next generation (NM-MC2, SGA1) chip"
  Task responsible: T9.3.1
  Status: no release has been received yet (matches planning)
- Upstream Component "SP9 BrainScaleS USB-Spikey (standalone physical model hardware system)"
  Task responsible: T9.2.4
  Status: final release has been received
  Quality of upstream Component: usable for validation of benchmark implementations and comparison/reference purposes with some remaining bugs in the associated software stack and known limitations concerning sensitiveness to operating environment (e.g. temperature, supply voltage)
- Upstream Component "Python client for the Neuromorphic Computing Platform"
  Task responsible: T9.1.1
  Status: intermediate release has been received
  Quality of upstream Component: release-quality software with some enhancements planned to improve the software interface for reproducible/comparable experiment workflows
- Upstream Component "SP9 specific provenance for the BrainScaleS system"
  Task responsible: T9.2.3
  Status: service is available in intermediate form
  Quality of upstream Component: usable for prototyping and testing purposes with known limitations affecting monitoring of certain parameters (like energy consumption) in the current form
- Upstream Component "SP9 specific provenance for the SpiNNaker system"
  Task responsible: T9.3.2
  Status: service is available in intermediate form
  Quality of upstream Component: usable for prototyping and testing purposes with some known limitations affecting extensive monitoring of certain machine parameters in current form
- Upstream Component "SP9 Neuromorphic Job Queue Service"
  Task responsible: T9.1.2, T9.1.1
  Status: service is available in intermediate form

Quality of upstream Component: fully usable for prototyping and testing purposes with planned enhancements concerning scalability to a larger user base and management of experiment workflows

- Upstream Component "SP9 Neuromorphic Benchmarks results database"
  Task responsible: T9.1.3
  Status: intermediate release has been received
  Quality of upstream Component: small data sets usable for prototyping and testing are available
- Downstream Component "SP9 Neuromorphic Benchmarks app"
  Task responsible: T9.1.3
  Status: for benchmarks app integration, the benchmarks service has been made available as a simplified façade implementation, which adheres to the final software interface

2.6.3.3 Neuromorphic Benchmarks results database

Description: performance results collected during regular runs of the entire neuromorphic benchmark suite to track changes of platform performance over time.

Progress on Component: The reference results for the first collection of neuromorphic benchmarks running on the currently available versions of the neuromorphic computing hardware and simulation platforms will be documented as part of the imminent Milestone MS9.1.1 in M12. It is expected that results for energy consumption will initially only be available on the small-scale hardware platforms, as actual measurements are not currently supported on the large-scale hardware platforms. Results will be measured by partner UNIBI inside the infrastructure provided by partner CNRS and platform-specific support provided by partners UHEI and UMAN.

Quality Control:

- Upstream Component "SP9 Neuromorphic Benchmarks repository"
  Task responsible: T9.1.3
  Status: intermediate release has been received
  Quality of upstream Component: the available sample of benchmarks with associated exemplary results suffices to validate the underlying database schema and to test the access interfaces
- Downstream Component "SP9 Neuromorphic Benchmarks service"
  Task responsible: T9.1.3
  Status: an interface description with an initial, simple implementation has been provided for integration with the service
- Downstream Component "Bio-mimetic classifier"
  Task responsible: T9.5.4
  Status: no task-specific benchmark results have been released yet (matching planning); first suitable benchmarks are scheduled for inclusion in MS9.1.1
- Downstream Component "SP9 SpiNNaker next generation (NM-MC2, SGA1)"
  Task responsible: T9.3.1
  Status: some relevant observations have been reported; comprehensive data will be made available after the release of the MS9.1.1 set of benchmarks
- Downstream Component "SP9 SpiNNaker software stack"
  Task responsible: T9.3.2
  Status: feedback concerning usability, robustness, and performance of the software stack has been and is routinely provided while porting and evaluating new benchmarks on the SpiNNaker platform
- Downstream Component "SP9 SpiNNaker next generation system (NM-MC2) (integrating hardware and software)"

Task responsible: T9.3.3

Status: no task-specific benchmark results have been released yet (matching planning); larger application benchmarks, required for expressive task-relevant results, are expected as part of the work on the MS9.1.5 benchmark set due in M24

- Downstream Component "SP9 BrainScaleS frontend software"

  Task responsible: T9.2.3, T9.2.2

  Status: feedback concerning usability, robustness, and performance of the software stack has been and is actively exchanged while porting and evaluating new benchmarks on the BrainScaleS platform

- Downstream Component "SP9 BrainScaleS 2 Neuromorphic Computing System (version 2 = NM-PM2) -- system under development"

  Task responsible: T9.2.1

  Status: no task-specific benchmark results have been released yet (matching planning); while first relevant results are expected when evaluating the full MS9.1.1 set of benchmarks, more influential application benchmarks will be included in the MS9.1.5 benchmark set due in M24

- Downstream Component "SP9 BrainScaleS standalone next generation single chip physical model system"

  Task responsible: T9.2.1, T9.2.4

  Status: some conclusions concerning system performance have already been drawn from the currently available subset of benchmarks; a more comprehensive analysis of platform traits will become possible with future extended benchmark sets in MS9.1.1 and MS9.1.5

2.6.3.4 Neuromorphic Benchmarks app

Description: web-based interface to trigger runs of the neuromorphic benchmark suite and to access the collected performance results.

Progress on Component: This Component constitutes the user-visible interface part of the benchmarks service that will be released as part of MS9.1.2 in M12. The Component is currently undergoing internal testing. The app itself has been developed primarily by partner CNRS with testing support and feedback provided by partners UNIBI, UHEI, and UMAN.

Quality Control:

- Upstream Component "SP9 Neuromorphic Benchmarks repository"

  Task responsible: T9.1.3

  Status: intermediate release has been received

  Quality of upstream Component: the available subset of the planned first benchmark collection is sufficient for validation of the concept of the benchmarks app and initial testing except for potential scalability issues.

- Upstream omponent "SP9 Neuromorphic Benchmarks service"

  Task responsible: T9.1.3

  Status: intermediate release has been received

  Quality of upstream Component: the received façade implementation is sufficient for prototyping the benchmarks app and conducting integration tests with the other software components of the system

2.7 T 9.1.4 Model and Experiment Descriptions

2.7.1 Key Personnel

Task Leader: Andrew P. DAVISON (CNRS)
Other Researcher: Ulrich RÜCKERT (UNIBI)

Other Researcher: Eric MÜLLER (UHEI)
Other Researcher: David R. LESTER (UMAN)
Other Researcher: Christian MAYR (TUD)
Other Researcher: Mikael DJURFELDT (KTH)

2.7.2 SGA1 DoA Goals

1) Continued development and maintenance of the PyNN model description API, and its implementation for the NEST and NEURON simulators and for the NM-PM-1 and NM-MC-1 neuromorphic hardware systems.
2) Investigation of the feasibility of code generation for the SpiNNaker system and the BrainScaleS 2 plasticity processor from model descriptions in the NineML and/or LEMS model description languages.
3) Improved interoperability with the model and experiment description languages used in the Brain Simulation Platform including support for import and export of NeuroML model descriptions.

2.7.3 Component Progress

2.7.3.1 PyNN

PyNN (pronounced 'pine') is both a simulator-independent Python API for building neuronal network models and a Python package that implements this API for the NEST, NEURON and Brian simulators. It is developed using an open-source community model. CNRS (P10) is one of the major contributors and the project administrator.

Progress on Component:

Two API versions are currently maintained, v0.7 and v0.8, since the production BrainScaleS and SpiNNaker systems do not yet support 0.8.

Three minor releases were made during the reporting period:

- Version 0.8.1 – 25/05/2016 – support for NEST 2.10.0, bug fixes. Release notes: http://neuralensemble.org/docs/PyNN/releases/0.8.1.html
- Version 0.8.2 – 06/12/2016 – new spike source models, minor enhancements, bug fixes. Release notes: http://neuralensemble.org/docs/PyNN/releases/0.8.2.html
- Version 0.8.3 – 08/03/2017 – support for NEST 2.12.0, NeuroML export (see following component). Release notes: http://neuralensemble.org/docs/PyNN/releases/0.8.3.html

During this period, we began tracking the NEST master branch in a specific development branch of PyNN, which allowed us to reduce the time needed to release a compatible PyNN release following a NEST release from six months with NEST 2.10 to seven days with NEST 2.12.

CDP to which Component contributes: CDP 5

Quality Control:

- Upstream Component "NEURON"

Task responsible: T6.3.4

Status/quality: NEURON is a well-established, widely used and high-quality simulator developed primarily outside HBP. The PyNN unit tests are run with NEURON version 7.4.

- Upstream Component "NEST - The Neural Simulation Tool"

Tasks responsible: T7.1.1, T7.1.3, T7.1.4, T7.4.1

Status: NEST 2.10.0 was released shortly before the end of the Ramp-Up Phase (31/12/2015). NEST 2.12.0 was released on 01/03/2017.

Quality: NEURON is a well-established, widely used and high-quality simulator with a large development community both inside and outside HBP, which is highly responsive to bug reports.

- Upstream Component "pyNN.neuroml"

See detailed report below.

- Upstream Component: Plasticity models: SP 4

Tasks responsible: T4.3.1, T4.3.2

Status/Quality: The plasticity models developed in WP 4.3 will be used when reviewing the PyNN API for synaptic plasticity modelling, in the second year of SGA1.

- Upstream Component: "MUSIC library"

Task responsible: T9.1.2

Status/quality: see detailed report above

- Upstream Component: "Neo"

Task responsible: T5.7.1, T5.7.2

Status: Neo 0.4.0 was released on 07/07/2016, Neo 0.4.1 was released on 03/08/2016.

Quality: Neo is a stable and well-tested Python package with a medium-to-large development community both inside and outside HBP (Over the past twelve months, 16 developers contributed to Neo [source: https://www.openhub.net/p/python-neo])

- Upstream Component "BrainScaleS frontend software"

See detailed report under T9.2.3 below.

- Upstream Component "SpiNNaker software stack"

See detailed report under T9.3.2 below.

- Downstream Components "Neuromorphic Visual motor coordination" (T10.2.3, T10.4.3), "SP9 Neuromorphic Benchmarks repository" (T9.1.3), "Sensory-motor maps" (T10.1.2), "Workflow for comparison of spiking neuron data obtained from NEST and SpiNNaker simulation" (T9.1.5), "Sensory models" (T10.1.2), "Robot Demonstrator – Mobile platform" (T10.4.3), "Biomimetic learning control architecture for modular robots" (T10.4.4), "Laminart with segmentation and retina into the NRP" (T10.2.1), "Model of spinal cord using reservoir computing for real-time control" (T10.4.5), "SP9 Neuromorphic Job Queue Service" (T9.1.1), "pyNN.neuroml" (T9.1.4)

  Status: PyNN 0.8.0 has been available to downstream Components since the beginning of SGA1. Three further releases of PyNN have taken place during the first 12 months of SGA1.

2.7.3.2 pyNN.neuroml

A PyNN module for import/export and simulation of NeuroML/LEMS models. This component was developed in collaboration with the NeuroML development community, in particular Dr. Padraig GLEESON (UCL).

Progress on Component:

During this period, a new version of the pyNN.neuroml module, supporting NeuroML v2, was developed, primarily by Dr. Gleeson. This enables a large subset of PyNN models to be exported as NeuroML, and subsequently simulated and/or visualised in any of the tools supporting NeuroML v2 (see https://neuroml.org/tool_support).

Quality Control:

Upstream Component "PyNN"

Note that pyNN.neuroml is a sub-module of PyNN, and uses part of the PyNN "common" module. It is therefore both an upstream and downstream Component of PyNN.

Upstream Component "NEST code with abstracted neuron model representations"

Task responsible: T7.1.3

Status: Work is underway, but not yet released (expected Month 24).

Upstream Component "Model representations for synaptic plasticity"

Task responsible: T6.4.3

Status: Work will begin in the second year of SGA1

2.7.3.3 Neuromorphic code generation feasibility study

Investigation of the feasibility of code generation for the SpiNNaker system and BrainScaleS-2 plasticity processor from model descriptions in the NineML and/or LEMS model description languages.

Progress on component:

Initial discussions have taken place; the majority of this work is planned for the second year of SGA1.

CDP to which Component contributes: CDP 5

Quality control:

Upstream Component "Plasticity models"

Task responsible: T4.3.1

Status: The plasticity models developed in WP 4.3 will be used as part of the feasibility study, in the second year of SGA1

Upstream Component "Model representations for synaptic plasticity"

Task responsible: T6.4.3

Status: Work will begin in the second year of SGA1.

2.7.3.4 The connection-set algebra library

The CSA (connection-set algebra; Djurfeldt 2012) is an algebra and two parallelized implementations thereof (one in Python and one in C++) which can succinctly express complex neuronal network connection patterns as mathematical expressions with operators on sets. The ConnectionGenerator interface (Djurfeldt et al 2014) is a generic interface for connecting connection generating libraries to neuronal network simulators. It is currently supported by NEST.

Progress on component:

As an in-kind contribution from the KTH partner, the C++ CSA library (libcsa) has been extended with an XML parser for CSA expressions which makes it possible to move expressions between the two CSA implementations and also to import expressions from external tools. The KTH partner has also developed a new version of the ConnectionGenerator interface which is an improvement over the first version both with regard to efficiency and functionality. The new interface supports MPI- and OpenMP-aware connection generators. In addition, the KTH partner has been involved in work on the description of connectivity as well as benchmarking of connectivity primitives in a collaboration with JUELICH and NMBU, Norway. This work will eventually result in a paper.

2.8 T9.1.5 Model Simplification and Validation

2.8.1 Key Personnel

Task Leader: Michael DENKER (JUELICH)

Other Researcher: Andrew DAVISON (CNRS)

Other Researcher: David LESTER (UMAN)

2.8.2 SGA1 DoA Goals

1) Develop, improve and implement algorithms/methods for simplifying models so as to be able to run them on neuromorphic hardware

2) Develop systematic workflows for comparison of simulation results from models at different levels of simplification and running on different hardware systems, including maintenance and

3) Further development of the Elephant data analysis toolkit

2.8.3 Component Progress

2.8.3.1 Model simplification service (neuromorphic)

This Component will provide a service to simplify neuronal network models by modifying the underlying model descriptions of network components (e.g., neurons, synapses,..) based on the PyNN framework. This Component is related to DoA goal 1 of T9.1.5.

Progress on Component:

Work on the simplification workflow has continued, in collaboration with SP6:

- Added 'GIF_cond_exp' model (used in the SP6 Model simplification pipeline) to PyNN, with NEST and NEURON implementations (CNRS).
- Added stochastic synapses to PyNN (both static and Tsodyks-Markram plasticity); NEURON implementation complete, NEST implementation in testing (CNRS).
- The SpiNNaker team (UMAN) are investigating the implementation of the GIF model and of stochastic synapses.

Quality Control:

- Upstream Component: "Simplified brain models (model)"

  Task responsible: T6.2.7

  Status: no inputs yet received beyond the work done in the Ramp-Up Phase.

- Upstream Component: "Fitting Generalized Integrate-and-Fire models (model)"

  Task responsible: T4.1.2

  Status: The GIF model was delivered during the Ramp-Up Phase and has now been incorporated into PyNN.

2.8.3.2 Workflow for comparison of spiking neuron data obtained from NEST and SpiNNaker simulation

Workflow for comparison of spiking neuron data obtained from NEST and SpiNNaker simulation. This Component is related to DoA goal 2 of T9.1.5

Progress on Component:

In M01-M12, JUELICH performed further work on the initial prototype of the validation workflow of simulation results of NEST and SpiNNaker using the Elephant framework (available at https://collab.humanbrainproject.eu/#/collab/507/nav/6326). This collaborative effort between SPs 5, 6, 7, and 9 was prominently disseminated to Project members during the HBP Summit, and in particular was featured as the primary demonstration on how to perform collaborative work in the context of the session on Collaboratory use. JUELICH demonstrated this workflow at the HBP student conference in 2017. Furthermore, JUELICH has presented the work at the JARA HPAC Conference in Aachen (Oct. 4-5, 2016), and has published the result as the first article describing a real-world collaborative workflow based on multiple software tools and integrated by the Collaboratory (Senk, J., Yegenoglu, A., Amblet, O., Brukau, Y., Davison, A., Lester, D.R., Lührs, A., Quaglio, P., Rostami, V., Rowley, A., et al. (2017). A Collaborative Simulation-Analysis Workflow for Computational Neuroscience Using HPC. In High-Performance Scientific Computing, E. Di Napoli, M.-A. Hermanns, H. Iliev, A. Lintermann, and A. Peyser, eds. (Cham: Springer International Publishing), pp. 243–256.). In order to implement the next version of validation tests based on the existing NEST-SpiNNaker-Elephant-workflow, and to facilitate collaboration with

T4.5.1, an additional student (Robert GUTZEN, JUELICH) has been hired, compensating also for relocation of the Elephant component of this Task to SP5 during the DPIT process. In this context, we started to investigate (i) measures that quantify the difference in observables from the two simulation environments and (ii) use the eigenvalue decomposition of the correlation matrix from network simulations to extract measures that characterise the structure of correlation exhibited by the simulation, going beyond descriptions on the level of the distributions of pairwise correlation coefficients. This and future sets of enhanced NEST-SpiNNaker validations are made available in the Collaboratory at https://collab.humanbrainproject.eu/#/collab/2366/nav/20166. More details are available in the description of the Milestone MS9.1.2 "Workflow for systematic comparison of simulation results obtained by models running on traditional HPC resources and on neuromorphic hardware, using the Elephant data analysis toolkit."

Quality Control:

- Upstream Component: "PyNN (software)"

  Task responsible: T6.4.3/T9.1.4

  Status: Operational for this component.

- Upstream Component: "UNICORE (software)"

  Task responsible: T7.5.6

  Status: Operational for this component.

- Upstream Component: "SP9 Neuromorphic Job Manager app (service)"

  Task responsible: T9.1.1/9.1.2

  Status: In development, current version included in the workflow.

- Upstream Component: "Collaboratory Task Service (service)"

  Task responsible: T11.3.1

  Status: currently not used in the component. Added value for this Task in the scope of future work: coupling of the Collaboratory task service to UNICORE.

- Upstream Component: "Collaboratory Provenance Service (service)"

  Task responsible: T11.3.2

  Status: basic functionality working; currently sufficient for this Component.

- Upstream Component: "Collaboratory Storage Service (service)"

  Task responsible: T11.3.2

  Status: Operational for this Component.

- Upstream Component: "Collaboratory Jupyter Notebook (service)"

  Task responsible: T11.3.2

  Status: Basic functionality required is operational.

- Upstream Component: "HPC systems at JSC (hardware)"

  Task responsible: T7.5.1

  Status: Operational for this component.

- Upstream Component: "ViSTA (software)"

  Task responsible: RUP T7.3.4

  Status: currently not used by this Component.

- Upstream Component: “SP9 SpiNNaker Neuromorphic Computing System (hardware)”

  Task responsible: no task

  Status: Operational for this Component.

- Upstream Component: “NEST – The Neural Simulation Tool (software)”

  Task responsible: T7.4.1/T7.1.4/T7.1.3/T7.1.1

  Status: Operational for this component.

- Upstream Component: “Elephant (software)”

  Task responsible: T5.7.1 (pre-DPIT: T9.1.5)

  Status: Successfully used in the validation workflow, future work includes the addition of analysis functionality developed for validation tests into the Elephant Component (see eigenvalue decomposition above).

- Downstream Component: “SP6-T6.4.4-SGA1-Validation result service (service)”

  Task responsible: T6.4.4

  Status: Components not yet linked.

2.8.3.3 Elephant

Elephant is a toolbox for the analysis of electrophysiological data based on the Neo framework. This component is related to DoA goal 3 of T9.1.5. The Elephant development has been moved from T9.1.5 to T5.7.1 during the DPIT process. See T5.7.1 for a cumulative description of the component progress for M1-M12.

3. WP 9.2: Next Generation Physical Model Implementation

3.1 Key Personnel

Work Package Leader: Johannes SCHEMMEL (UHEI)

3.2 WP Leader's Overview

What went particularly well?

During the reporting period the consortium reached two important achievements. First, the third prototype ASIC for the second-generation BrainScaleS-2 physical model system has been sent to manufacturing, including a multitude of novel circuits like short-time plasticity, neuronal adaptation and non-linear dendrites. Second, as part of the commissioning of the first generation BrainScaleS-1 system, a classification task based on the MNIST database has been successfully implemented on the system.

What didn't go according to plan?

Commissioning of the first-generation system has progressed well but is still ongoing.

Impact of work done

Three accepted peer-reviewed publications (see list of publications).

3.3 Priorities for the remainder of the phase

Regarding the development of the second-generation system, the third prototype ASIC will be tested and evaluated. Using the results from these measurements together with the already existing silicon components, the first full-size implementation of the second-generation neural network chip, called HICANN-DSL-SR, will be designed and manufactured.

Simultaneously, the commissioning of the first-generation system will be an ongoing activity, together with some substantial enhancements of the first-generation platform hardware. The available number of analog readout channels will be increased by a factor of 32. This is achieved by replacing the current membrane voltage readout systems by custom build highly-integrated multi-channel ADC boards directly integrated into the BrainScaleS wafer modules.

3.4 T9.2.1 Next Generation System Development (Chips)

3.4.1 Key Personnel

Task Leader: Johannes SCHEMMEL (UHEI)

Other Researcher: Sebastian HÖPPNER (TUD)

Other Researcher: Yusuf LEBLEBICI (EPFL)

3.4.2 SGA1 DoA Goals

Task T9.2.1 will develop all VLSI circuits necessary for the implementation of the NM-PM roadmap. It will produce prototype chips to verify their operation. It will design the necessary hardware for testing the prototype chips and develop the soft- and firmware associated with these test systems.

3.4.3 Component Progress

During the reporting period the Task developed a multitude of novel circuits for the first implementation of the neural network ASIC for the second-generation BrainScaleS system. This ASIC will be a central element of the following component: single–chip system. It will

be also the basis of the neural network wafer used in the component: second-generation BrainScaleS system.

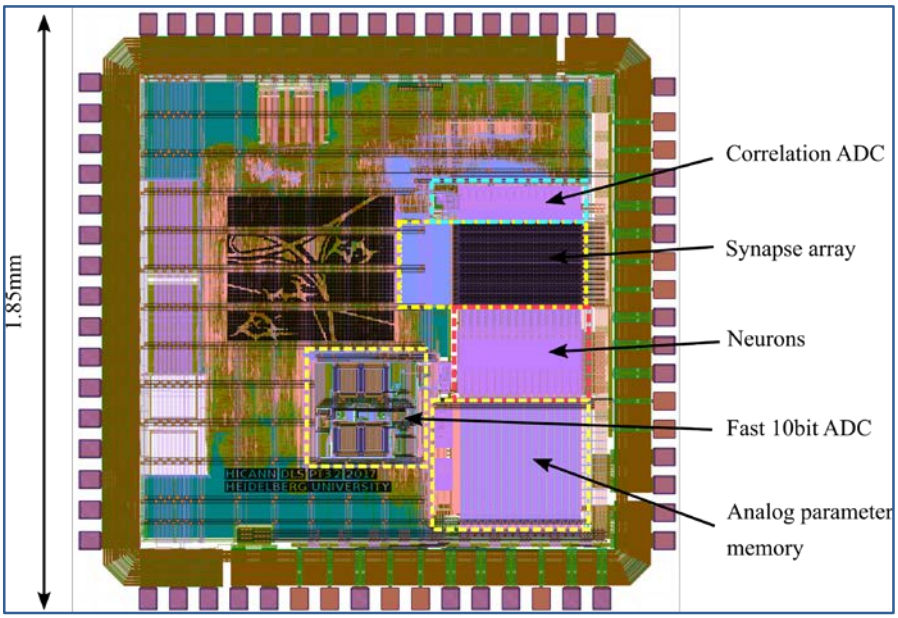

Figure 6: Layout drawing of the third HICANN-DLS prototype ASIC.

Analog neuromorphic blocks are labelled according to their respective function. The remainder of the chip area contains digital logic that comprises the plasticity processor (PPU), control logic for the analog blocks and interface logic for external communication.

The following circuits have been developed:

- Random number vector load for the plasticity processing unit

An array of pseudo random number generators has been added to the vector part of the PPU. Their output is multiplexed with the synapse data bus. This allows fully parallel loading of random numbers into a vector register.

- Short-term plasticity

At the input of each synapse row, a novel full-custom circuit has been inserted which allows a modulation of the synapse strength depending on the short-term spiking history of all 64 possible pre-synaptic neurons projecting into each synapse row. It allows either depression or facilitation according to the Tsodyks-Markram model.

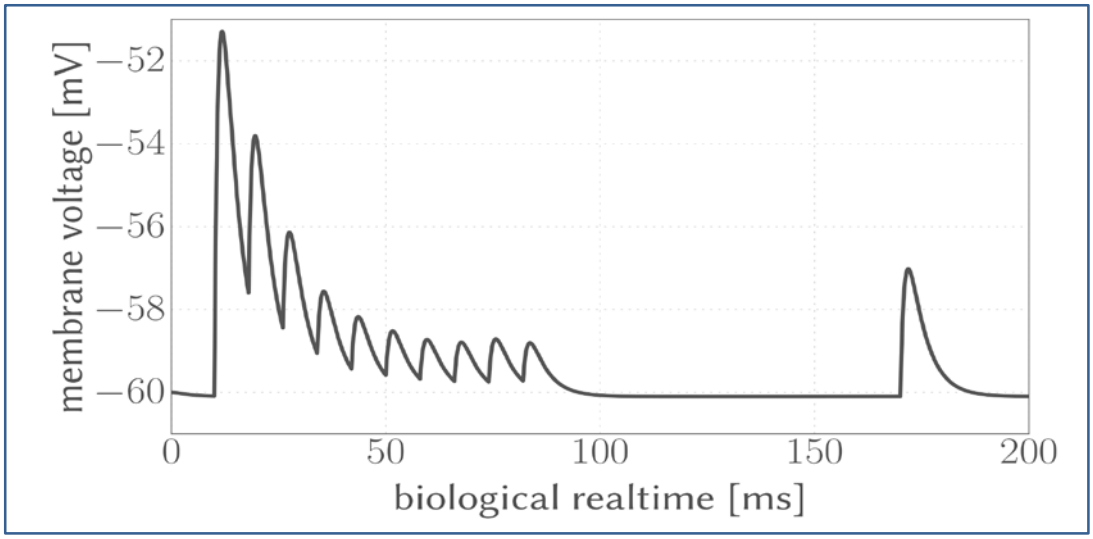

Figure 7: Demonstration of short-term plasticity.

A series of spikes with inter-spike-intervals of 8 ms was used to depress the target synapse. After 90 ms another event was injected, showing the recovery response. The synaptic current was taken from a transistor-level simulation of a synapse driver and a synapse circuit. The membrane dynamics were simulated in software.

- Digital neuron back-end

A novel circuit comprising of spike-detection, spike-registration, spike-synchronisation, refractory-time generation, spike output bus arbitration and spike output bus has been developed and implemented, allowing the capture and transmission of up to 2 GSpikes/s in the HICANN-DLS-SR chip.

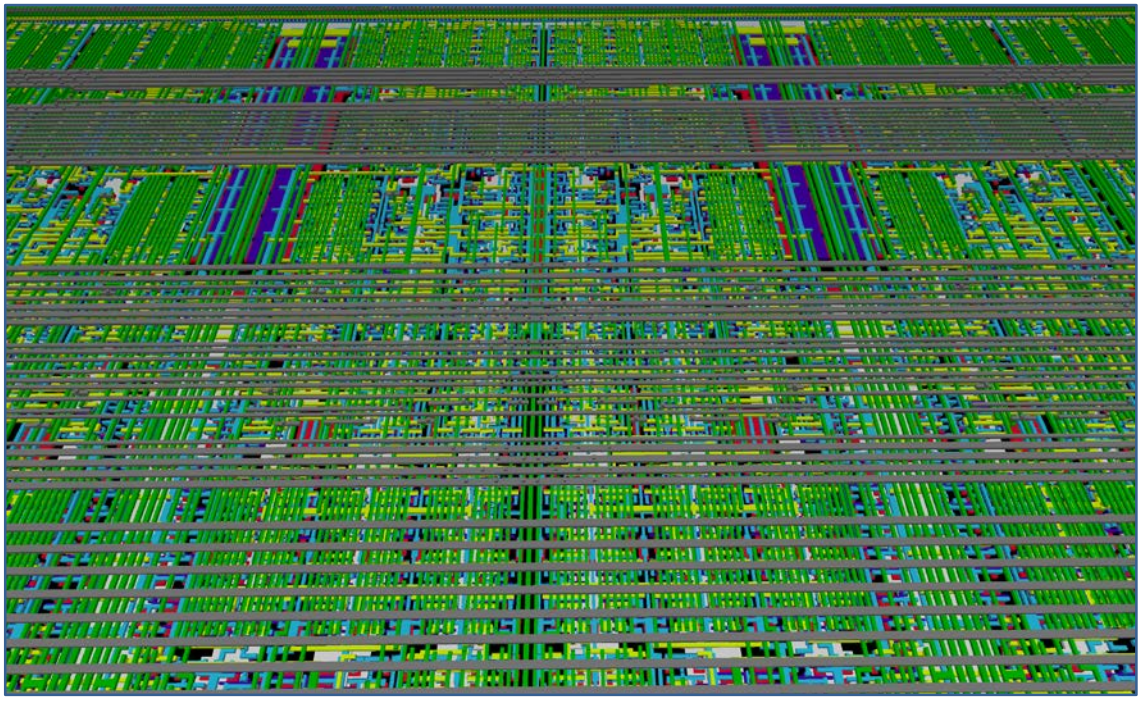

Figure 8: 3D rendering of the digital neuron backend layout.

The digital neuron backend provides serialization and bus access conflict resolution for access to the digital bus. Also it produces control pulses for the neuron facilitating extremely long refractory times and widening the dynamic range of adaptation in the AdEx neuron.

- Non-linear dendrites

A novel circuit implementing three types of non-linear dendrites has been developed. It allows the emulation of NMDA, calcium and sodium spikes.

- Multi-Compartment extensions

The BrainScaleS generation one flexible neuron sizing scheme has been extended to support programmable conductances between the individual neuronal compartments. This allows, among other things, the emulation of the dendritic structure of pyramidal neurons.

- Upgraded correlation ADC

A second revision of the correlation ADC has been implemented, resulting in true 8 bit DNL and INL at a total sampling rate of 1 GSample/s in the HICANN-DLS-SR chip.

- Digital neuron read-out chain

An embedded fast 10bit ADC together with the necessary amplifiers and multiplexers has been developed and integrated into the HICANN-DLS prototype. It allows the simultaneous sampling of two analog sources, i.e. two membrane voltages, at more than 30 MSamples/s each.

- Adaptive Exponential neuron circuits

Progress has been made to complete the circuits emulating the full Adaptive Exponential neuron model. A first version of an implementation of the adaptation and exponential terms has been developed.

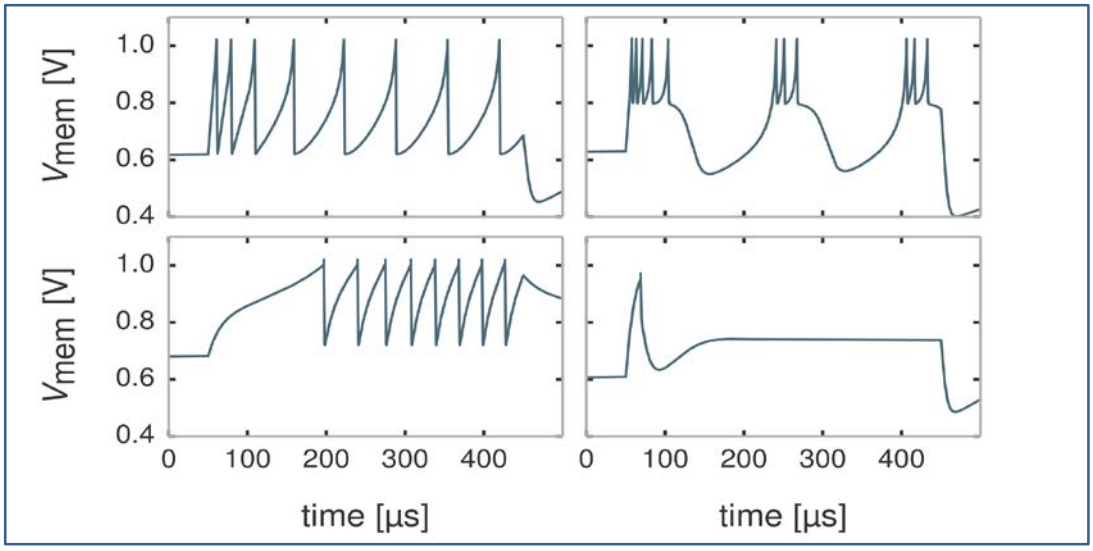

Figure 9: Firing patterns of the Ad-Ex neuron observed during a step current stimulation.

The patterns show adaptation (top left), regular bursting (top right), delayed acceleration (bottom left) and transient spiking (bottom right). Taken from L. Kriener, P. Müller, S. Billaudelle, S. A. Aamir, J. Schemmel and K. Meier, "Neuronal Dynamics: Silicon vs. Biology", 1st HBP Student Conference, 2017.

All novel circuits together with the results from the Ramp-Up Phase have been integrated into the third HICANN-DLS prototype ASIC and sent to manufacturing during the reporting period.

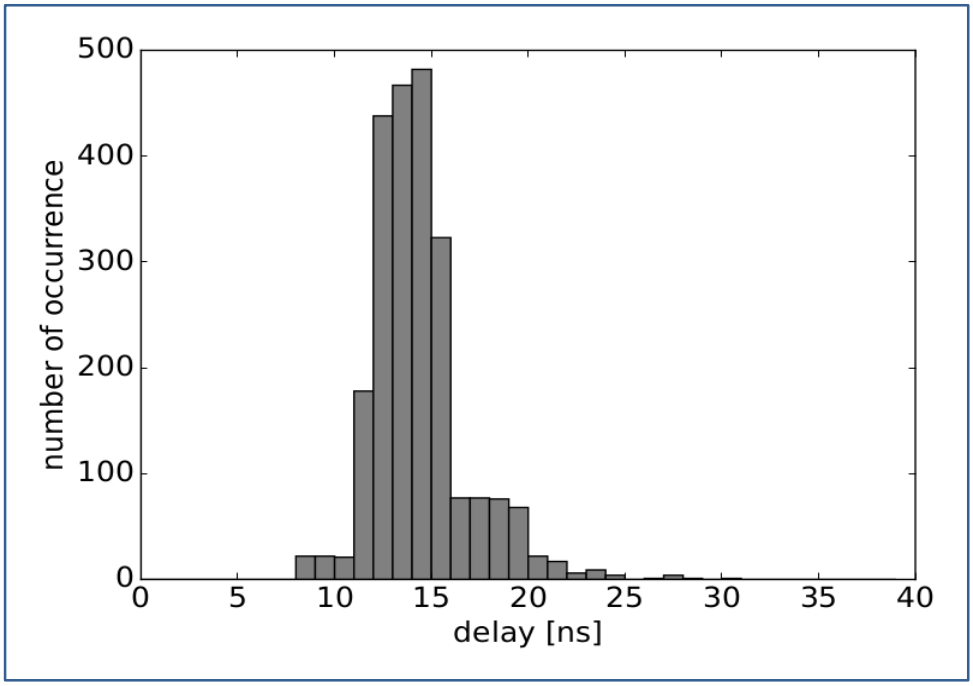

Figure 10: Delay distribution introduced by the digital neuron backend resulting from 1MHz input from 32 neurons over 200µs.

Translated to biological timescales: 1kHz input over the course of 200ms. Even when all neurons fire at such high rates, the firing distribution is narrow and the mean biological delay introduced is just 15µs, hence not measureable against typical time constants.

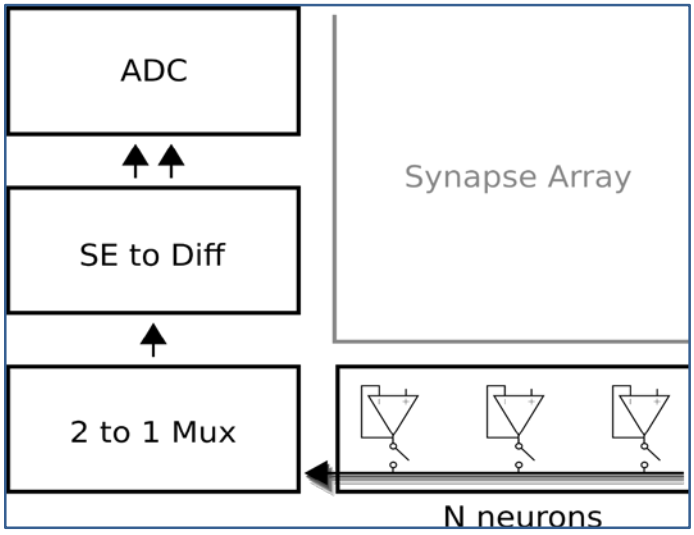

Figure 11: Fully differential analog readout chain employing an analog 2 to 1 active mux to allow for readout of simultaneously observable quantities in every of the currently 32 neurons.

The concept is scalable and employs analog tristate buffers and fully differential signaling for the most sensitive parts of the readout chain. Not only the neuron membrane voltage, but also the other state variables like the synaptic voltage can be read out.

3.5 T9.2.2 Next Generation System Development (Boards, Firmware, Software)

3.5.1 Key Personnel

Task Leader: Dan HUSMANN (UHEI)

Other Researcher: Sebastian HÖPPNER (TUD)

Other Researcher: Oswin EHRMANN (IZM/FG)

3.5.2 SGA1 DoA Goals

Task T9.2.2 will develop the main PCBs and all necessary auxiliary components to prototype the complete, operational wafer-modules. It will develop novel packaging technologies according to the NM-PM Roadmap. It will provide the technical documentation for operation and manufacture of the wafer-modules. It will engineer the compute and network infrastructure for the NM-PM systems. It will develop the system management and communication soft- and firmware.

3.5.3 Component Progress

During the reporting period the planned improvements of the Tasks could be achieved. A new cooling system for the wafer modules with adjustable FANs was built and integrated into the system software. System control and system test software was improved and unified.

See the following detailed descriptions of component progress.

SP9 BrainScaleS Embedded Wafer Prototype

Novel packaging technologies have been developed by SP9 partner Fraunhofer IZM, according to the NM-PM Roadmap.

Progress on Component:

IZM: Two functional post-processed HICANN wafers have been thinned, prepared and laminated into printed circuit boards.

UHEI (M. Güttler): A 2-layer printed circuit board for the lamination process has been designed. The board includes test structures to verify the divergent interconnect levels.

A complex test-equipment for the HICANNs of the laminated wafer has been developed, see Figure 12.

The connectors showed no problems and created a reliable connection between the boards.

Tests of almost all connected elements of the laminated wafer have been performed and the results showed no structural problems.

No differences between a standard packaged HICANN chip and an embedded chip were observable.

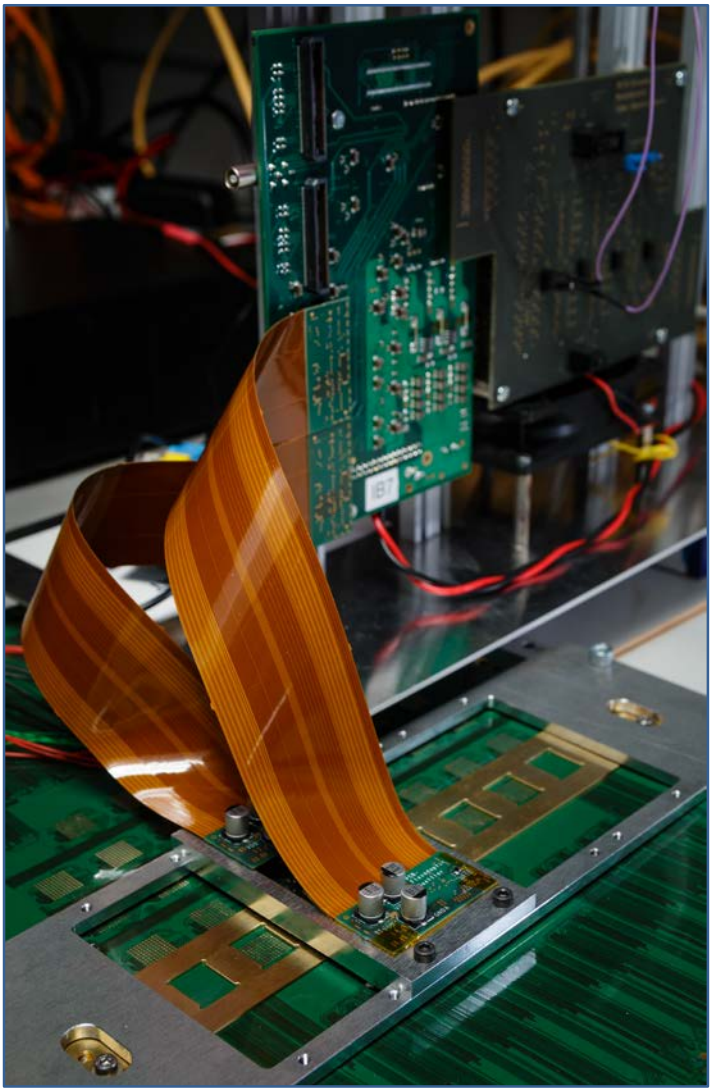

Figure 12: Test equipment for the embedded wafer testing.

At the bottom there is the printed circuit board with the HICANN wafer inside. On top there are the mechanical components for the fixation of the additional printed circuit boards, which establish the connection between the wafer and the HICANN test device (in the background). For the HICANN test device there is no difference between a standard HICANN test chip and an embedded HICANN chip. This reuse of existing hardware led to a significant shorter development time.

For a solderless connection between embedded wafer printed circuit board and the test equipment Samtec ZA1 connectors were used, see Figure 13.

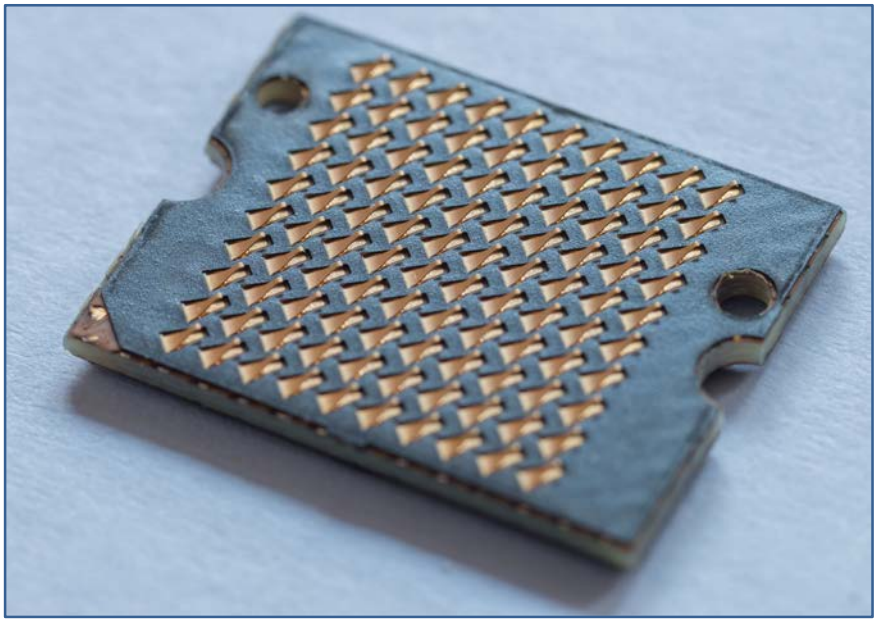

Figure 13: The additional printed circuit boards for the testing are connected to the printed circuit board with the embedded wafer using Samtec ZA1 connectors.

They have spring contacts on both sides and create under compression a reliable vertical connection. Due to the coarse pitch of 1 mm between the pins no special adjustment device is needed for the alignment.

3.5.4 SP9 BrainScaleS prototype hardware for upgrade of the ADC-subsystem

Custom analog readout system for improved robustness and calibration capabilities.

Progress on Component:

UHEI (J. Ilmberger): A prototype board (see Figure 14) of the ADC-subsystem upgrade has been developed. The features of this prototype are:

- fully parallel readout of all analog signals using 12bit 32.5Msps ADCs
- fully parallel low-speed precision DC measurement capability for all analog signals
- attached FPGA unit for data compression and gigabit Ethernet network connectivity
- interface between Wafer Module and System Management Control Unit for system power switching and monitoring and system temperature measurements

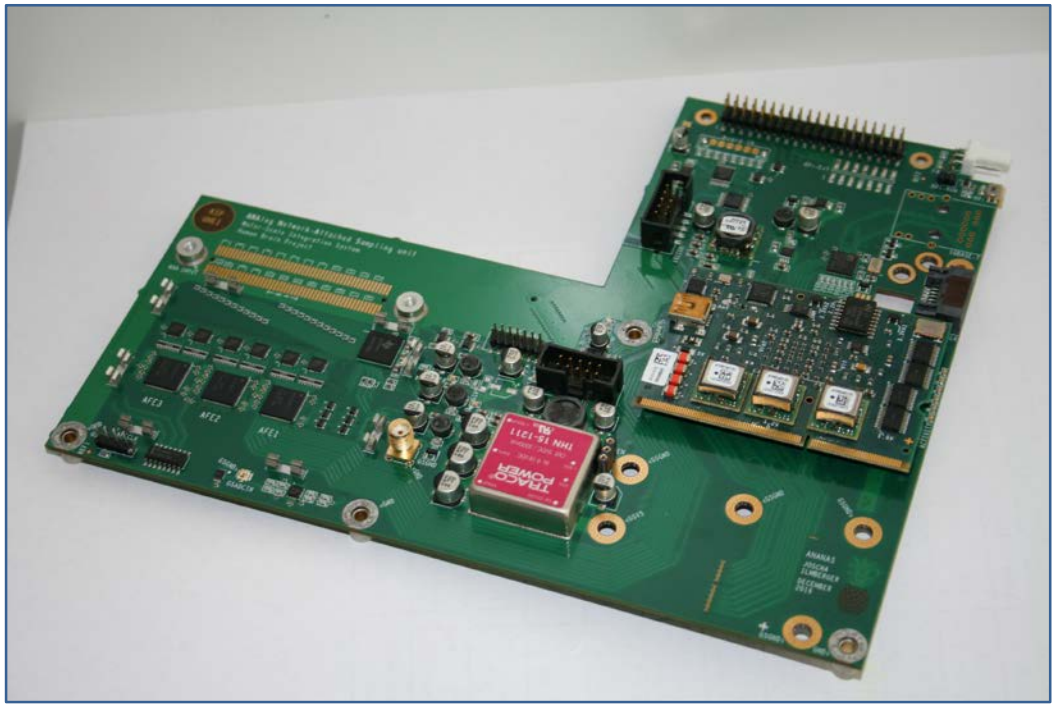

Figure 14: Top view of the ADC-subsystem upgrade board (ANANAS) for fully parallel readout of all analog signals with additional precision DC measurement capability for calibration purposes.

SP9 BrainScaleS firmware for the communication subsystem

Firmware for the FPGAs of the communication subsystem.

Important extensions of the functionality have been added to the FPGA firmware:

- FPGA reprogramming via Ethernet
- Support for global system time reference
- Extensions for faster and more reliable experiment execution
- More fine-grained status monitoring and experiment control via FPGA

CDP5: Accelerated neuromorphic implementation of action- perception and learning loops for spatial navigation on the BrainScaleS-1/2 systems

Quality Control:

Upstream Components:

- SP9 BrainScaleS code repositories: E. MÜLLER; T9.1.1, T9.2.3
- SP9 Neuromorphic Benchmarks repository: A. DAVISON; T9.1.3, T9.3.4
- SP9 Neuromorphic Benchmarks result database: A. DAVISON; T9.1.3

Downstream Components

- PyNN: A. DAVISON; T6.4.10, T9.1.3, T6.4.3, T9.1.4, T9.5.4
- SP9 Neuromorphic Benchmark repository: A. DAVISON; T9.1.3, T9.3.4
- Neuromorphic Direct Access service: A. DAVISON; T9.1.4

3.6 T9.2.3 Next Generation System Integration

3.6.1 Key Personnel

Task Leader: Eric MÜLLER (UHEI)

3.6.2 SGA1 DoA Goals

Task T9.2.3 manages the production and operation of the NM-PM Platform hardware. It is responsible for the large-scale manufacturing and test of all system components according to the NM-PM Roadmap and the results from Tasks T9.2.1 and T9.2.2. It will operate the NM-PM Platform hardware during the duration of the project as part of the Neuromorphic Computing Platform. It will provide technical personnel to provide technical assistance to remote users.

3.6.3 Component Progress

3.6.4 SP9 BrainScaleS Frontend Software

Operation software of the physical model neuromorphic compute systems (BrainScaleS systems). Component "*SP9 BrainScaleS System Software*" tracks software layers that are exposed to users of the BrainScaleS systems.

CDP to which Component contributes: CDP5 "Plasticity, Learning and Development", Use Cases 1, 2, 3, 6.

3.6.4.1 Resource Management

Scheduling of experiments (or low-level tasks) on the BrainScaleS system uses the SLURM scheduler software. To improve usability and robustness of the software, we have been implementing a custom job submit plugin to manage the hardware-related resource scheduling. Single hardware components, such as FPGAs and ADCs, can be requested by the user. Interdependencies between hardware components are automatically resolved and a convenient translation between different "coordinate systems" is provided. In addition to fine-grained sub-wafer-level access, a full-wafer allocation mechanism has been implemented. All hardware information is acquired from a central hardware database which tracks the current status of the system (e.g. connected FPGAs, analog readout lines, and availability information for other hardware Components).

3.6.5 SP9 BrainScaleS System Software

User frontend software for the BrainScaleS neuromorphic computing system, including the PyNN implementations for the BrainScaleS systems.

CDP to which Component contributes: CDP5 "Plasticity, Learning and Development", Use ases 1, 2, 3, 6.

3.6.5.1 Monitoring

A central database for monitoring and logging of hardware statistic data has been implemented: all available values (e.g., temperatures, supply voltages, currents, fan speeds) are stored into a carbon-based backend, visualization is provided by graphite (Figure 15); both tools are open-source software. The data is organized in a hierarchical structure: wafer modules provide groupings for FPGAs, supply power PCBs, fans and reticles. Based on the monitored data, an alert mechanism has been implemented using Shinken.

Figure 15: Visualisation of logged monitoring values (here temperature graphs) by graphite

3.6.5.2 Intermediate representation of mapping result

During the last year, the mapping result storage was modified to provide a structured API to access mapping results. Previously (Figure 16, panel A), the mapping software directly wrote to a data container representing the low-level configuration (configuration bits) of the BrainScaleS hardware system. The new "intermediate representation" (internal name: N1C3) provides a solution. Users can now access the low-level data via PyNN objects (e.g. PyNN "Neurons" can be used as an "index" to the hardware container) and modify the low-level configuration of the associated hardware data structure.

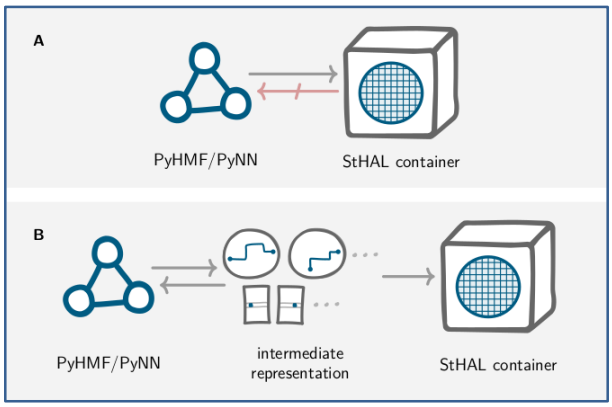

Figure 16: New structured API to access mapping results

3.6.5.3 Low-level Software

3.6.5.3.1 DLS

- The low-level software layer for the future DLS-based systems (e.g. the HICANN-DLS-SR system) has been started. As a first step, an abstract coordinate system (comparable to the one existing for the BrainScaleS System) has been defined and implemented in C++. The completion of a container structure for the full hardware configuration is almost complete. Full test coverage of the coordinate system and the container has been reached.
- Spike routing with the FPGA has been integrated.

- Support for NUX (the embedded processor on the HICANN-DLS) cross-compilation flow has been integrated into the BrainScaleS build system (waf-based).
- Preliminary support for the NUX vector unit has been integrated into GCC:
  - support for C-style vector parallelization with common vector type and intrinsic functions
  - assembler inline coding supports automatic variable handling
  - optimisation is enabled for vector code
- Due to the improved GCC support, overall usability for NUX programming has improved and a simple PPU testing framework has been implemented. Several tests now make use of the upgraded GCC cross-compiler for the PPU.
- Software and hardware tests verifying DLS functionality have been implemented.
- Experiments: Figure 17

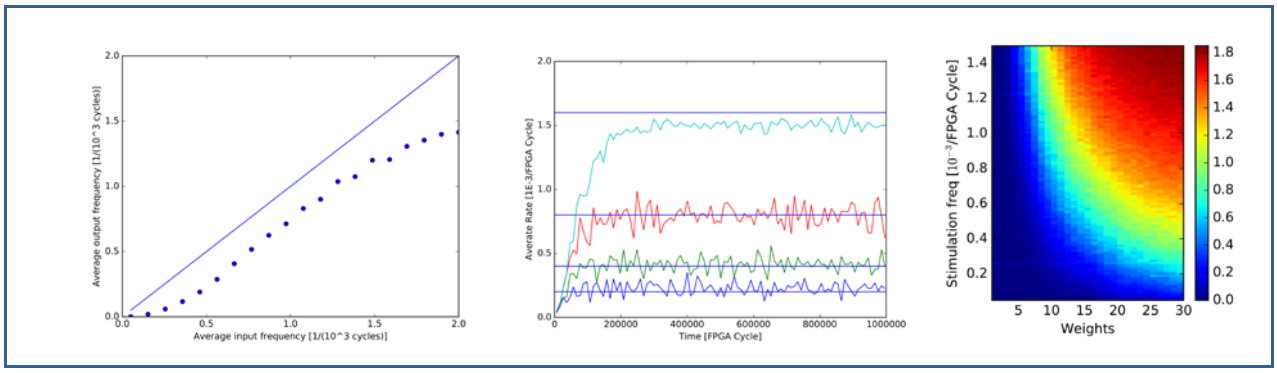

Figure 17: Experiments on DLS.

Left: Fast measurement of network gain function (average output frequency vs. average input frequency).; Middle: Use PPU for sweeping weights during emulation; Right: Use rate counter for homeostasis mechanism for rate stabilisation

3.6.5.3.2 Synapse Calibration

For the usage of the neuron circuits as LIF model emulation, the synaptic and membrane time-constants are characteristic parameters of the interaction between neurons. For these new methods for calibration were developed based on voltage recordings of PSP curves, which are caused by incoming spike events. The time-constants are extracted from a recorded PSP by fitting a parameterized model of its shape. Because recording of the neuron membrane in HICANN is afflicted by noise, multiple PSPs are averaged to improve the precision and verify the quality of fits using various tests.

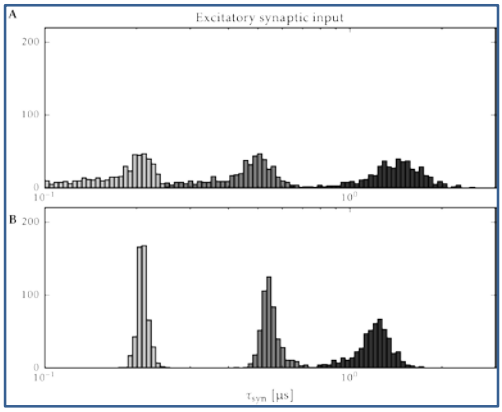

Figure 18: Evaluation of the excitatory synaptic input.

A Shows the raw data obtained during the calibration. The resulting mean time-constant µ and the standard deviation are noted in the table. B Shows the evaluation measurement using N = 30 repetitions. For each neuron its mean value is plotted. We compare the width σ m of the resulting distribution, with the mean trial-to-trial variability ⟨σ_t⟩ over all neurons. For both lower values, σ_m is dominated by the trial-to-trial variations. For τ_{syne} = 2 µs, σ m is twice as large as ⟨σ_t⟩. This are caused by trial-to-trial variations occurring during the calibration, that are not perfectly averaged out.

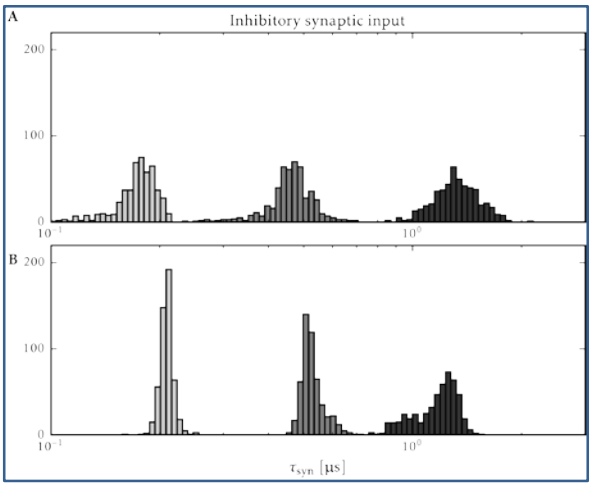

Figure 19: Evaluation of the inhibitory synaptic input.

It confirms that the behaviour of both inputs is indeed identical. A Shows the raw data obtained during the calibration. The resulting mean time-constant µ and the standard deviation are noted in the table. B Shows the evaluation measurement using N = 30 repetitions. For each neuron its mean value is plotted. We compare the width σ m of the resulting distribution, with the mean trial-to-trial variability ⟨σ_t⟩ over all neurons. For both lower values σ m is dominated by the trial-to-trial variations. For τ_{syni} = 2 µs σ_m is twice as large as ⟨σ_t⟩. This is caused by trial-to-trial variations occurring during the calibration, that are not perfectly averaged out.

Figure 18 and Figure 19 show the result of those measurements for a single HICANN using all of its 512 neurons. Here, we find that the achieved range of time-constants varies greatly between neurons. The excitatory time-constant of a single neuron can be between 0.14 ± 0.13 µs and 2.3 ± 1.1 µs, where the error indicates the standard deviation of the neuron-to-neuron variability. The common range, which can be covered by 99.7 % (corresponding to a 3σ interval of a normal distribution) of the neurons, extends from 0.18 to 1.1 µs. Respectively, the inhibitory time-constant for a single neuron lies between 0.11 ± 0.12 µs and 4.3 ± 1.7 µs and that the common range extends from 0.14 to 2.0 µs. The differences between excitatory and inhibitory time-constants are mainly caused by the asymmetry of the inhibitory and excitatory input on hardware. The precision of the calibration achieves the limit given by the trial-to-trial variability of the analog parameters. The same method

can be employed to calibrate the membrane time-constant. However, here we considered only the case of a scaling of 1:3 of the control parameter of the leakage conductance, so we analysed only a subset of the available parameter space. We find the possible range for the individual neuron reaching from 0.47±0.09µs to 3.3±1.9µs and the common range extends from 0.65 to 1.14µs. Here we can also reach a precision close to the trial-to-trial variability for the whole parameter range, but observe the same systematic error as for the synaptic time-constants.

In addition, a model for the strength of the synapse circuits on HICANN was developed. The strength of the synaptic event dependents on its digital weight w and divisor gdiv as well as the analog reference current Vgmax. Since those are not accessible to direct measurements, it is based on transistor-level simulations. We find that the synapses have a weight-dependent offset, which we incorporated into the model. Monte Carlo simulations were used to estimate the expected variation in strength. We find that the variations mostly depend on the chosen divisor gdiv and that for gdiv < 10 a relative error below 10 % can be expected.

These calibrations improve the ability to run neural network experiments on the BrainScaleS System by increasing precision of the LIF neuron model in the hardware.

3.7 T9.2.4 Small-Scale NM-PM System

3.7.1 Key Personnel

Task Leader: Andreas GRÜBL (UHEI)

3.7.2 SGA1 DoA Goals

Task T9.2.4 manages the development and test of a single-chip version of the NM-PM2 system. It will build upon the prototype design developed for the wafer-scale system within the scope of Task T9.2.1 and will extend it to provide the necessary communication bandwidth and network size to be useful outside of the wafer-module context. It will manufacture said modified NM-PM2 chip and develop and produce a test platform for its evaluation. The expected result of Task T9.2.4 is a small stand-alone NM-PM2 platform providing an implementation of the hybrid plasticity technology developed for NM-PM2 during the Ramp-Up Phase. The network size will be comparable to that of a single HICANN in NM-PM1.

3.7.3 Component Progress

3.7.3.1 SP9 BrainScaleS standalone next generation single chip physical model system

CDP to which Component contributes: CDP5: “Plasticity, Learning and Development”, use case: “Accelerated neuromorphic implementation of action- perception and learning loops for spatial navigation on the BrainScaleS-1/2 systems (SP9, SP3, SP4)”

Component progress:

The single chip system will reuse the existing Xilinx Kintex7-based FPGA boards from the NM-PM1 (BrainScaleS-1) system. Therefore, adaption of the NM-PM1 FPGAs’ firmware has been started with an initial implementation of the low-level communication layer between FPGA and the planned NM-PM2 (BrainScaleS-2) prototype ASIC. In a first version, this communication layer will rely on the existing implementation in the NM-PM1 circuits and extend this to increase communication bandwidth for pulse and configuration data transport.

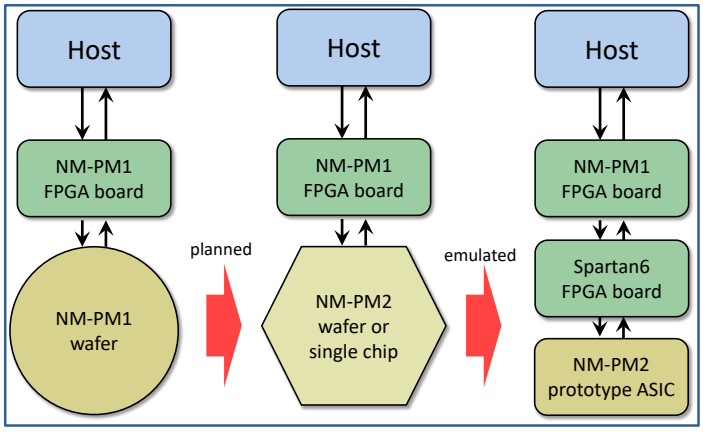

Figure 20: Illustration of the setup that will be used to emulate the communication between the existing host computers and NM-PM1 FPGA boards on the one hand and the next generation single chip version of the NM-PM2 wafer.

In a first step, all digital communication will be emulated on a Spartan6 based FPGA board that is already in use for prototype testing.

Since the NM-PM2 prototype circuits developed in Task 9.2.1 do not yet contain this communication layer, it has been implemented in a Xilinx Spartan6 FPGA that is already being used to control the current NM-PM2 prototype ASICs in their according test setup (see Figure 20 for an illustration of this setup). This way, testing the new communication layer is possible, already with the current NM-PM2 prototype ASICs. Communication between both FPGAs was successfully simulated and the designs were synthesised without timing or FPGA resource issues. Initial tests of the physical communication between the FPGAs have been completed successfully.

Quality Control:

- The Task uses the existing (and known-to-be-good) NM-PM1 FPGA-Boards and the Prototype ASIC that is to be developed in task T9.2.1. The ASIC has not been finished which is according to the work plan.
- Delivery to downstream Components will be possible with the implemented single-chip system at hand.

4. WP 9.3: Next Generation Many Core Implementation

4.1 Key Personnel

Work Package Leader: Steve FURBER (UMAN)

4.2 WP Leader's Overview

What went particularly well?

The testing of Santos (RUP next generation SpiNNaker chip prototype) has shown that we have functional hardware, and that TUD/UMAN teams can work constructively together.

What didn't go according to plan?

Hiring of new personnel was badly delayed by the delayed contract negotiations.

Impact of work done:

The many-core platform is now significantly more stable than it was at the end of the RUP. The evaluation of the Santos chip has shown that three of the hardware accelerators will work, and given a solid basis on which to estimate power consumption for the final chip.

4.3 Priorities for the remainder of the phase

There are two priorities in SGA-1: designing and taping out the Quad-Processor Unit (QPE), which is an M24 deliverable, and enhancing the current Many-core platform (SpiNNaker-1), which is also an M24 deliverable.

4.4 T9.3.1 Next Generation System Development (Hardware)

4.4.1 Key Personnel

Task Leader: Christian MAYR (TUD)

Other Researchers: Sebastian HÖPPNER (TUD), Steve FURBER (UMAN)

4.4.2 SGA1 DoA Goals

This Task is focused on the development of the NM-MC2 chip. This activity includes development and refinement of the chip architecture including a holistic verification approach by partner UMAN. This includes the development of an FPGA or other prototype/simulation of the NM-MC2 chip for design verification and early stage firmware and software development.

Key circuit components of the final NM-MC2 chip are to be designed and prototyped in the target 28nm SLP CMOS technology on a test chip by partner TUD. This includes serial-off chip interfaces and a memory interface to connect Hybrid Memory Cube modules (HMC). The NM-MC2 chips are to be assembled into a compact 2.5D or 3D silicon interposer package for close integration with the HMC memory module for a small mechanical footprint.

Within this Task, the concepts of this integration scheme are to be developed and prototyped by Partner FG. Furthermore, the work includes the development of power supply concepts for the NM-MC2 system by Partner EPFL. Provision is made in this Task for the full documentation via a data-sheet. Associated Tasks are the NM-MC2 software development in T9.3.2 and the realisation of a small scale NM-MC2 system in T9.3.4. Hardware development will be carried out in close collaboration with WP4.1, WP4.2, WP4.3 and WP4.4.

4.4.3 Component Progress

SP9 SpiNNaker next generation (NM-MC2, SGA1) chip

Development and refinement of the chip architecture including a holistic verification approach by partner UMAN. This includes the development of an FPGA or other prototype/simulation of the NM-MC2 chip for design verification and early stage firmware and software development.

Key circuit components of the final NM-MC2 chip are to be designed and prototyped in the target 28nm SLP CMOS technology on a test chip by partner TUD. This includes serial-off chip interfaces and a memory interface to connect Hybrid Memory Cube modules (HMC). The NM-MC2 chips are to be assembled into a compact 2.5D or 3D silicon interposer package for close integration with the HMC memory module for a small mechanical footprint.

Progress on Component:

The SpiNNaker2 chip top-level architecture has been developed as shown in Figure 21. Key component of the processor array is the quad-processing-element (QPE), which consists of 4 ARM processors with local memory and a network-on-chip router for packet based communication from/to the SpiNNaker router, off-chip memory and other (Q)Pes on the SpiNNaker2 chip. The work included architecture refinement with respect to the power management architecture (fine-grained dynamic voltage and frequency scaling per core, as evaluated in the silicon prototype of NM-MC2 from the HBP Ramp-Up-Phase) and the hierarchical RTL2GDS implementation approach for feasible layout realisation of this multi-million-gate system-on-chip (*contributions by TUD and UMAN*). The work has been documented in a draft-version of the SpiNNaker2 chip datasheet (internal documentation). In 09/2016 a workshop on SpiNNaker 2 chip architecture development was held in Manchester (attendants from UMAN and TUD).

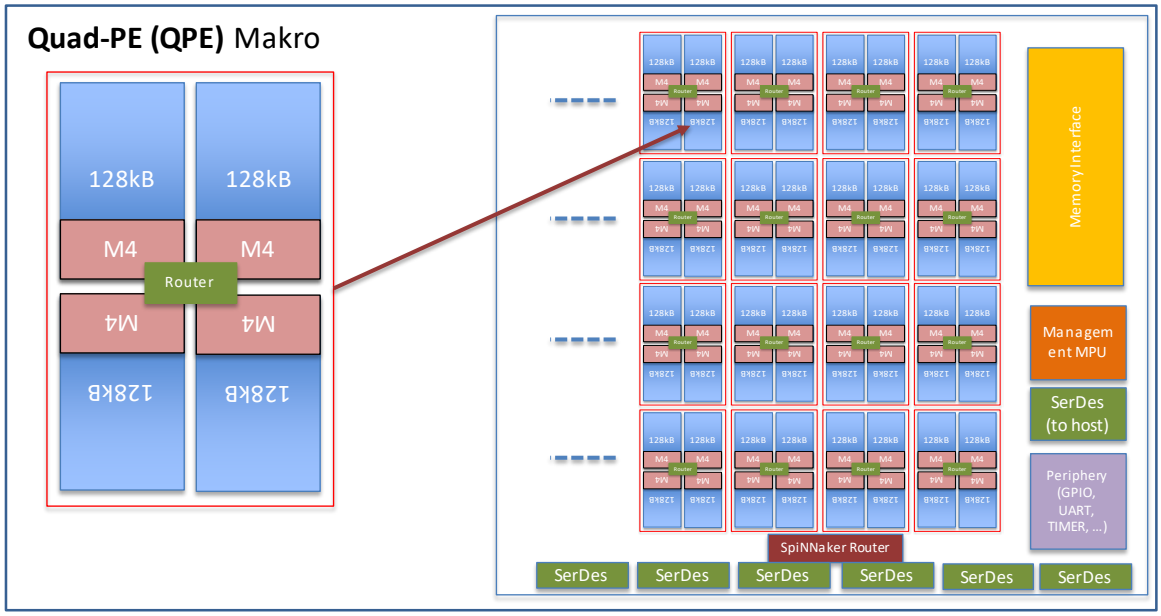

Figure 21: SpiNNaker2 chip toplevel architecture block diagram (simplified)

Key circuits for on-chip communication and memory access have been specified, partly implemented in synthesisable Verilog RTL code and verified by simulations. This includes the NoC routers for system configuration and communication (*contribution by TUD*) and the QPE internal busses (*contribution by UMAN*).

An initial version of QPE has been implemented in synthesisable Verilog RTL code and is ready for verification by simulations and execution of low level software. Therefore, a software built-flow for usage within circuit verification, FPGA prototyping and later chip lab evaluation has been set up (*contribution by TUD*).

The exponential function co-processor which is targeted for the acceleration of synaptic computations for learning (e.g. STDP) has been evaluated on the NM-MC2 silicon prototype

for the HBP Ramp-Up-Phase. Results have been submitted for publication (*contributions by TUD and UMAN*):

A Fixed Point Exponential Function Accelerator in 28nm CMOS for a Digital Neuromorphic System, Johannes Partzsch, Sebastian Höppner, Matthias Eberlein, Christian Mayr, David R Lester, Stephen B Furber, *submitted* to IEEE International Symposium on Circuits and Systems, 2017, ISCAS2017

Based on detailed estimations of the performance (data-rate) of the Hybrid Memory Cube (HMC) off chip memory modules and its power consumption, this memory solution is now being considered as not suitable for the targeted energy efficiency of the SpiNNaker2 system, mainly due to its high idle power consumption. Therefore, an LPDDR4 memory interface solution has been evaluated. This includes re-planning of the packaging concept for close integration of the SpiNNaker2 chips and their off-chip memory. Currently a Package-on-Package (PoP) solution is preferred. (*contributions by TUD and UMAN*).

Quality Control:

Upstream Component:

- SP9 model: Principles for Brain-Like Computation [important]

  Received a benchmark for stochastic neuromorphic computing from TUG, a software model has been delivered, and ported to the ARM M4F cores for NM-MC2. A 1-week workshop on the bring-up of this benchmark on the NM-MC2 prototype was held in 10/2016 at TUD (attendants from TUG and TUD).

Downstream Component:

- SP9 SpiNNaker small-scale NM-MC System SGA1 [essential]

  The QPE synthesisable Verilog RTL, its verification environment and the software built flow has been provided to T9.3.4. (TUD) as intermediate release.

- SP9 SpiNNaker next generation system (NM-MC2) (integrating hardware and software) [essential]

  Provided nothing (SpiNNaker2 chip not yet manufactured)

4.5 T9.3.2 Next Generation System Development (Software)

4.5.1 Key Personnel

Task Leader: Dave LESTER (UMAN)

Other Researcher: Andrea ACQUAVIVA (Polito)

4.5.2 SGA1 DoA Goals

This Task is responsible for the PyNN-SpiNNaker compiler software. It will be responsible for the development, maintenance and upgrades of the NM-MC-1 and NM-MC-2 compiler software. Our goal is to make this sustainable software, (i.e. maintainable by the community after HBP finishes).

4.5.3 Component Progress

4.5.3.1 SP9 SpiNNaker software stack (software)

Description of Component (from PLA)

The Distributed Compiler and run-time system for the SpiNNaker neuromorphic computing system.

CDP to which Component contributes: CDP 1, 2, and 5

Progress on Component, including: Release 2016.001 “Another Fine Product..” released 28 April 2016, 2016.001.001 released 5 May 2016, and 3.0.0 released 2 September 2016. Collaboration Manchester/Polito.

Quality Control:

- Upstream Components
  - SP9 Neuromorphic Benchmarks repository [added value]

    Support of this activity provided us with helpful alpha, beta testers
  - SP9 Neuromorphic Benchmarks results database [added value]

    Support of this activity provided us with helpful alpha, beta testers
- Downstream Components
  - T3.3.3 Decoded spike patterns of neural ensembles in cortex and hippocampus during multimodal scene representation [added value]
  - T3.3.3 Multi-area recordings from visual and somatosensory cortices, perirhinal and entorhinal cortex and hippocampal CA1 [added value]
  - Neuromorphic Direct Access service [essential]

    Intermediate Releases provided (2016.001, *.001, 3.0.0).
  - SP9 Neuromorphic Benchmarks repository [essential]

    Intermediate Releases provided (2016.001, *.001, 3.0.0).
  - Robot Demonstrator – Robot arm [important]

    Intermediate Releases provided (2016.001, *.001, 3.0.0).
  - Robot Demonstrator – Mobile platform [essential]

    Intermediate Releases provided (2016.001, *.001, 3.0.0).
  - Bio-mimetic classifier [essential]
  - Modelling of network-level mechanisms from T3.6.3a [important]
  - Analysis of network-level mechanisms constraining the in vivo implementation of learning rules and implementing integration, encoding and recall of multisensory memories [important]
  - SP9 SpiNNaker small-scale NM-MC System SGA1 [essential]

    Intermediate Releases provided (2016.001, *.001, 3.0.0).
  - PyNN [essential]

    Intermediate Releases provided (2016.001, *.001, 3.0.0).

4.5.3.2 Nengo-NEF for SpiNNaker

Description of Component (from PLA)

This software component provides support for the Nengo/NEF modelling tool developed by Eliasmith and co at the University of Waterloo. The SpiNNaker version of this software is a joint development by UMAN and UWaterloo. (Repository: https://github.com/project-rig/nengo_spinnaker/)

Progress on Component, including: Nengo-NEF is released. There are still issues to do with communication/routing for the full SPAUN model.

4.6 T9.3.3 Next Generation System Integration

4.6.1 Key Personnel

Task Leader: Steve FURBER (UMAN)

Other Researcher: Sebastian HÖPPNER (TUD)

4.6.2 SGA1 DoA Goals

This Task is responsible for the porting of the lowest level software of the SpiNNaker software stack onto the hardware. In short it provides the integration between hardware and software.

4.6.3 Component Progress

4.6.3.1 SP9 SpiNNaker next generation system (NM-MC2) (integrating hardware and software) (service)

Description of Component (from PLA): Integration between the SpiNNaker next generation system hardware and the software. This service includes fast I/O ports (e.g. SATA) and external hardware devices e.g. silicon retina, silicon cochlea.

Progress on Component, including: A prototype loading system is in place, permitting the loading of test code onto Santos. An attempt will be made in April/May 2017 (M13-24) to link Santos to other neuromorphic devices.

For service-like activities, we currently advise using the more numerous and better supported SpiNNaker-1 platform – either the 48-node boards or the 4-node boards.

Quality Control:

- Upstream component: Prototype only.
- Downstream component: Prototype only.

4.7 T9.3.4 Small-Scale NM-MC System

4.7.1 Key Personnel

Task Leader: Sebastian HÖPPNER (TUD)

Other Researcher: Dave LESTER (UMAN)

4.7.2 SGA1 DoA Goals

Development of a small-scale version of the NM-MC2 neuromorphic many core system.

4.7.3 Component Progress

SpiNNaker small-scale NM-MC System SGA1

Description of Component

A small-scale version of the NM-MC2 neuromorphic many-core system. This includes the realisation of a small-scale multi-processor macro block with all required functionality for energy efficient neuromorphic computing (e.g. power management, hardware accelerators for neuromorphic computing) and communication fabric interfaces. This is intended for later scaling towards the final NM-MC2 chip by multiple instantiations of these macros and interconnection over a network-on-chip.

Progress on Component, including: planned releases achieved / which partner did what

A small scale NM-MC2 prototype system based on the Santos 28nm chips from the HBP Ramp-Up-Phase have been assembled and ramped-up for lab evaluation (Figure 22). The chips contain 4 ARM M4F cores with fine grained neuromorphic power management and hardware accelerators (exponential function) for neuromorphic computing (*contribution by TUD*).

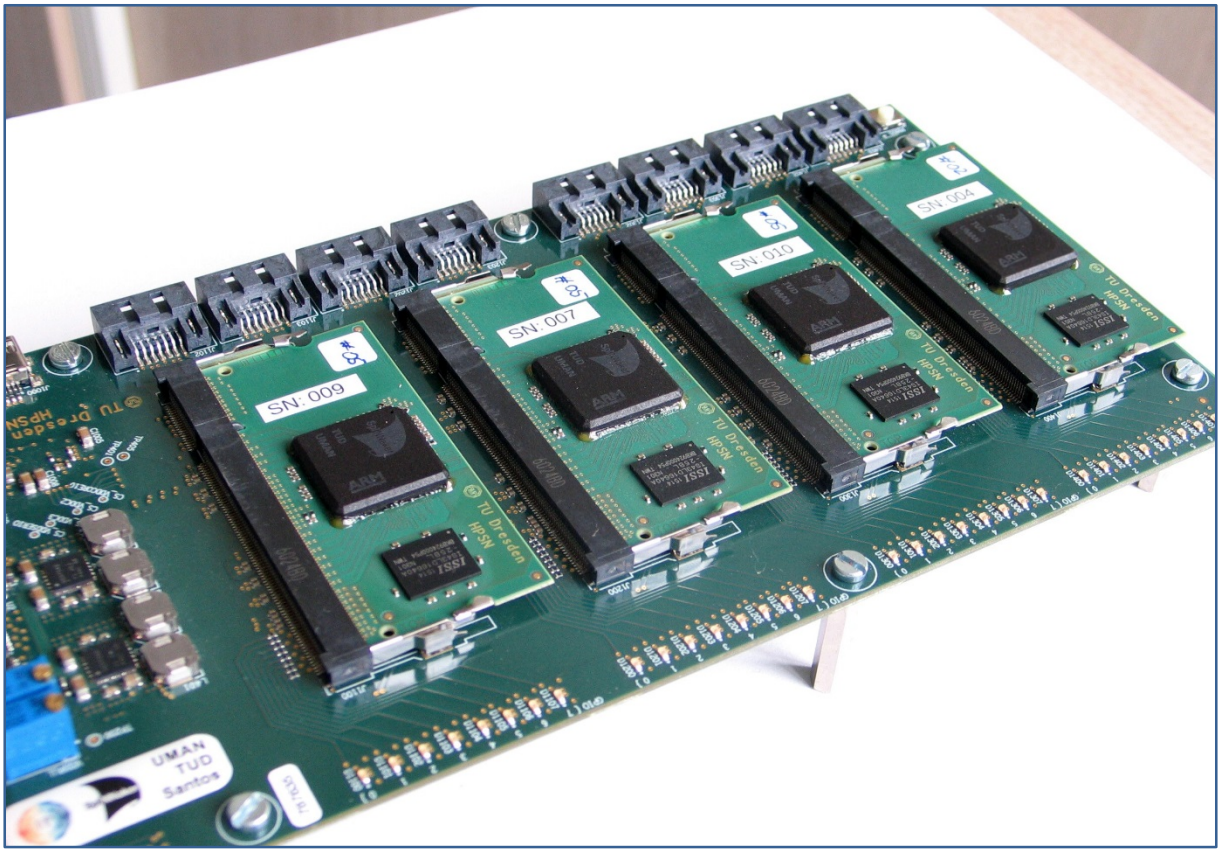

Figure 22: Prototype SpiNNaker-2 test chip and small-scale system (Santos)

Lab evaluation has been started with initial focus on the power management architecture for neuromorphic computing. A synfire chain benchmark has been implemented on the system and the effectiveness of DVFS has been demonstrated successfully. It has been shown that typically only approx. 1% of the simulation time steps have to be processed at the highest performance level (1.0V supply, 500MHz clock frequency) to maintain the biological real time performance (1ms time step) of the systems, whereas all other time steps can be processed at lower voltage and frequency levels, down 0.7V supply voltage operation at 125MHz clock frequency. By this approach the total system power consumption can be reduced by 45% without loss off simulation performance. A live demonstration setup for neuromorphic DVFS on this platform has been developed.

5. WP 9.4: Computational Principles

5.1 Key Personnel

Work Package Leader: Wolfgang MAASS (TUGRAZ)

5.2 WP Leader's Overview

What went particularly well?

Substantial progress was made in the implementation of paradigms for stochastic computations in neuromorphic hardware. Furthermore, new models for stochastic computation with assemblies of spiking neurons, that are based on data on network activity in the brain, where explored and prepared for implementing cognitive computations in neuromorphic hardware. Finally, a new model for network learning that integrates rewiring with goal oriented (reward based) synaptic plasticity was established. In contrast to previous models for network plasticity this model also takes into account the surprisingly strong stochastic component of rewiring and synaptic plasticity that has been reported in neurobiological experiments ("synaptic sampling"). Work has already started on the integration of this new model for network learning into neuromorphic hardware.

What didn't go according to plan?

No deviations.

Impact of work done

The HBP Platforms have already become international leaders in the implementation of stochastic computations. This strength has been enhanced, and complemented by new paradigms for network learning that are now also for the first time inherently stochastic. These advances have raised substantial international interest during invited presentations at the recent 2017 Conference on Neuro Inspired Computing Elements at the IBM Research Center in Almaden, USA.

5.3 Priorities for the remainder of the phase

Work on the implementation of synaptic sampling in both neuromorphic hardware platforms, BrainScaleS and SpiNNaker, will be continued, and applied to demanding learning tasks. The implementation of these new learning methods will be enhanced through the creation of an outer optimisation loop that optimises hyperparameters of the network for a variety of learning tasks ("learning-to-learn"). We need to find out which optimisation methods work based for these novel applications of the machine-learning inspired learning-to-learn paradigm to networks of spiking neurons in software and hardware. In addition, the new paradigms for cognitive computation, based on data about the role of assemblies of concept cells and their flexible associations in the human brain, that we have started to explore in the first year of SGA1 will be scaled up to larger networks, memory capacity, and more complex webs of associations, and implemented on SpiNNaker. This work will have to supported by theoretical analyses of the resulting new models for associations between memories. Other priorities for further theoretical work is the development of a better understanding of stochastic computation and learning in hierarchically structured networks of spiking neurons. Furthermore, combinations of unsupervised learning and reinforcement learning will be explored, taking experimental data from neuroscience as well as new ideas from machine learning into account.

5.4 T9.4.1 Principles for Brain-Like Computation

5.4.1 Key Personnel

Task Leader: Mihai A. PETROVICI (UHEI)

Other Researchers: David KAPPEL, Anand SUBRAMONEY, Guillaume BELLEC, Wolfgang MAASS (TUGraz)

Other Researchers: Sebastian SCHMITT, Eric MÜLLER, Oliver BREITWIESER, Mihai A. PETROVICI (UHEI)

Other Researchers: Jakob JORDAN, Markus DIESMANN, Tom TETZLAFF (Juelich)

5.4.2 SGA1 DoA Goals

1) Understand how the brain computes reliably with unreliable elements, and develop principles for porting this capability into neuromorphic hardware;
2) Investigate which properties of stochasticity are functionally relevant and develop methods for implementing these efficiently in neuromorphic hardware. In particular, random number generators will be analysed on test chips of TUD, and their applicability for stochastic neuromorphic computation will be investigated;
3) Understand the functional role of stereotypical spatio-temporal firing patterns, how different computational functions can be multiplexed, and how processing rules can be implemented in neural networks in order to support cognitive computing.

5.4.3 Component Progress

In work towards goals 1 and 2, UHEI showed how neurons in the high-conductance state can achieve an approximately logistic activation function. This endows them with an essential feature required for neural sampling as described in Büsing et al. 2011 for an abstract neuron model.

In [Petrovici 2016a], UHEI extended the previously developed theory to networks of spiking neurons (see also SP9 model component “Software model of sampling with LIF neurons”). With an appropriate parametrisation, such networks can be made to sample from arbitrary Boltzmann distributions, thereby allowing the implementation of a powerful machine learning model into networks of spiking neurons that are compatible with a large number of existing neuromorphic platforms.

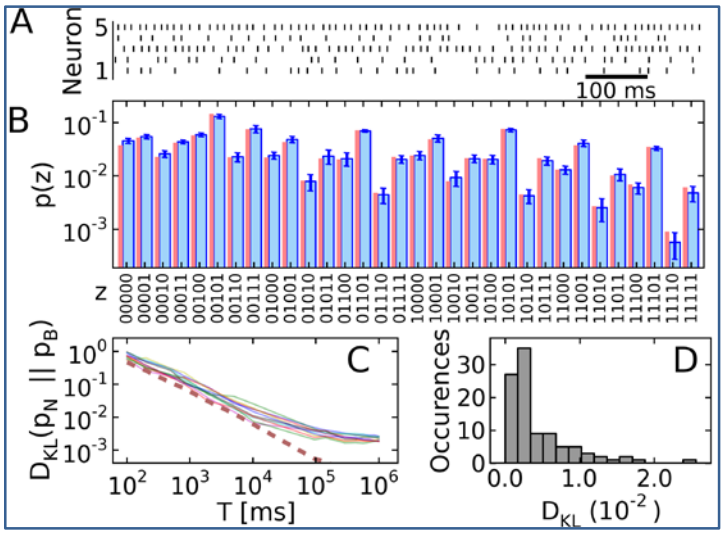

Figure 23: Software simulations of sampling with LIF neurons.

(A) Spike pattern of a recurrent network of 5 LIF neurons during sampling from a randomly generated Boltzmann distribution. (B) Sampled distribution of network states (blue) and target distribution (red). The sampled distribution was estimated from ten 10 s simulation runs (error bars: std. deviation between

runs). (C) Distance between sampled and target distribution as a function of integration time T for 10 trials. The red dotted line shows convergence for the theoretically optimal abstract model. (D) Distance between sampled and target distribution when sampling for T = 10^6 ms from 100 different randomly generated target distributions.

The neural sampling framework assumes the presence of an independent source of stochasticity for each neuron. On any finite-sized back-end – including neuromorphic hardware – a limited pool of independent noise sources necessarily leads to noise correlations that impair the performance of embedded stochastic networks. In previous work, FZJ and UHEI discussed how a small network of laterally connected neurons can actively decorrelate the input for a sampling network and fully restore its functionality. This has been extended with a more detailed investigation of how the nature of the sampled distribution influences the requirements for the provided noise (see also SP9 model component "Software model of noise generation by a balanced random spiking network").

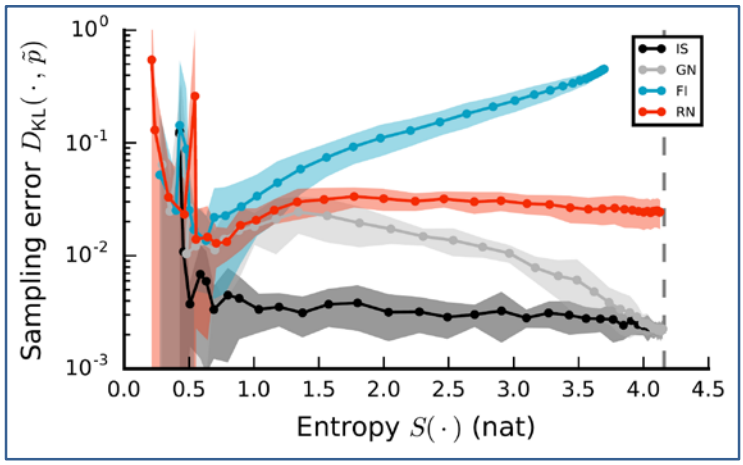

Figure 24: Dependence of the sampling error on the entropy of the state distribution.

IS: Intrinsically stochastic units. GN: Deterministic units with additive Gaussian Noise. FI: Deterministic units with input from a finite population of independent stochastic units. RN: Deterministic units with input from a finite recurrent network of deterministic units. Vertical dashed dark gray line indicates entropy of uniform distribution. Sampling-network size: 100 units. Number of noise sources for FI and RN: 222. Sampling time: 105ms.

Previously, UHEI ported the LIF sampling framework to an accelerated neuromorphic back-end. In order to compensate for various hardware-inherent effects that would otherwise distort the sampled distribution, a small subnetwork was designed to take the role of a sampling unit. The emulated model was shown to take full advantage of the accelerated hardware, reaching the desired approximation of the target distribution at a speed-up of 10,000 compared to an equivalent biological network (see also SP9 model Component "Hardware emulation of LIF sampling with subnetwork modules").

Depending on the function of the emulated stochastic network, distortive effects may play a more or less critical role. In [Petrovici et al. 2017a] UHEI showed both in software simulations and in hardware emulations that a hierarchical stochastic network emulating an RBM can be inherently robust to certain types of hardware-induced distortions (see also SP9 model component "Hardware emulation of hierarchical sampling networks").

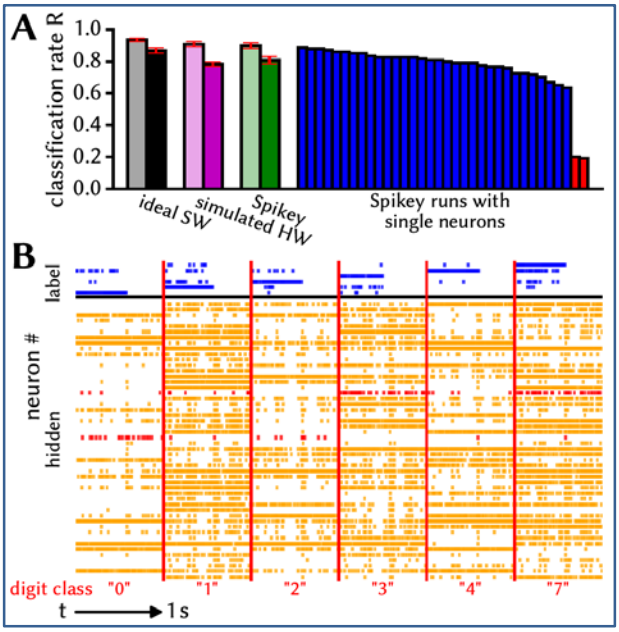

Figure 25: Study of a direct-to-hardware mapping of a hierarchical LIF network.

(A) Classification performance. Black: Software simulation of a distortion-free LIF network. Purple: Software simulation of the LIF network with all distortion mechanisms being present simultaneously, with amplitudes and variances as measured on Spikey. Green: Hardware emulation of the hidden layer with software evaluation of the label layer. Blue/red: Emulation of the hidden layer with repeated use of single neurons, followed by software evaluation of the label layer. The two "bad" neurons marked in red were not well configurable and therefore performed at chance level. (B) Exemplary spike trains of a subset of neurons in the LIF network with the hidden layer running on hardware. Spike trains belonging to the two "bad" neurons from A are marked in red.

This and further work done in WP9.4 has been the subject of two reviews: [Petrovici et al. 2016b] and [Petrovici et al. 2017b].

TUD and TUGRAZ have also carried out joint work towards goal 2 (part "analysis of random number generators on test chips of TUD"). A prototype of the reward-gated synaptic sampling model from [Kappel et al. 2016] was ported to the new SpiNNaker hardware that is developed at TUD. Due to the extensive use of random numbers, this model is ideal for evaluating the new hardware design and identifying possible performance bottlenecks. A first reward-based classification experiment was conducted and successfully realised on the hardware platform. The SpiNNaker implementation of the synaptic sampling model is currently being benchmarked and further optimised.

Regarding goal 3 of this task, Anand SUBRAMONEY and Guillaume BELLEC of TUGRAZ carried out research on the role of stereotypical spatio-temporal firing patterns in networks of spiking neurons for cognitive tasks that require a working memory. This work supports a new model for working memory that has been suggested by more recent experimental data based on Ca-imaging in rodents. This model for working memory can readily be implemented in neuromorphic hardware. The work on this model will be completed and two publications will be submitted later this year.

Finally, an analysis of experimental data that elucidate how the brain is able to compute reliably with unreliable elements was carried out by TUGRAZ in [Maass 2016] (see goal 1).

References:

[Petrovici et al. 2016a] Petrovici, M.A., Bill, J., Bytschok, I., Schemmel, J., Meier, K. Stochastic inference with spiking neurons in the high-conductance state. Physical Review E 94, 042312. 2016.

[Petrovici et al. 2017a] Petrovici, M.A., Schroeder, A., Breitwieser, O., Grübl, A., Schemmel, J., Meier, K. Robustness from structure: Inference with hierarchical spiking networks on analog neuromorphic hardware. Accepted at IJCNN 2017.

[Petrovici et al. 2016b] Petrovici, M.A., Leng, L., Breitwieser, O., Stöckel, D., Bytschok, I., Martel, R., Bill, J., Schemmel, J., Meier, K. Stochastic inference with spiking neural networks. BMC Neuroscience 2016, 17(Suppl 1):P96.

[Petrovici et al. 2017b] Petrovici, M.A., Schmitt, S., Klähn, J., Stöckel, D., Schroeder, A., Bellec, G., Bill, J., Breitwieser, O., Bytschok, I., Grübl, A., Güttler, M., Hartel, A., Hartmann, S., Husmann, D., Husmann, K., Karasenko, V., Kleider, M., Koke, C., Mauch, C., Müller, P., Partzsch, J., Pfeil, T., Schiefer, S., Scholze, S., Thanasoulis, V., Vogginger, B., Legenstein, R., Maass, W., Schüffny, R., Mayr, C., Schemmel, J., Meier, K. Pattern representation and recognition with accelerated analog neuromorphic systems. Accepted at ISCAS 2017.

[Maass 2016] Maass, W. Searching for principles of brain computation. Current Opinion in Behavioral Sciences (Special Issue on Computational Modelling), 11:81-92, 2016

5.5 T9.4.2 Emergence of Computational Capabilities through Learning

5.5.1 Key Personnel

Task Leader: Wolfgang MAASS (TUGraz)

Other Researcher: D. KAPPEL, R. LEGENSTEIN, T. LIMBACHER, G. BELLEC (TUGraz)

Other Researcher: Eric MÜLLER, Mihai A. PETROVICI, Sebastian SCHMITT (UHEI)

Other Researcher: Jamie KNIGHT (UMAN)

5.5.2 SGA1 DoA Goals

1) Identify plasticity mechanisms and principles for the large-scale organisation of network learning that enable users to have a variety of cognitive functions emerge in simulated neural systems through learning.
2) Develop principles for organising the plasticity of diverse types of neurons and synapses (including adaptive rewiring) in such a way that stable learning and network function is achieved in spite of numerous disturbances and device mismatches.

5.5.3 Component Progress

In work towards goals 1 and 2 a principle was developed for organising rewiring and synaptic plasticity in a generic recurrent network of spiking neurons in such a way that stable learning and network function is achieved in spite of disturbances. This principle takes the experimentally reported strong stochastic component of rewiring and synaptic plasticity into account, and uses it for enhancing the capabilities of network plasticity ("synaptic sampling"). The results are written down in a report, that will be submitted for publication shortly:

[Kappel et al. 2016] Kappel, D., Regenstein, R., Habenschuss, S., Hsieh, M., Maass, W. Reward-based self-configuration of neural circuits. Internal report, 2017.

In addition, it was shown that generic molecular mechanism in the postsynaptic density are likely to enhance the resulting learning model. These results were written down in a preprint;

Yu, Z., Kappel, D., Legenstein, R., Song, S., Chen, F., Maass, W. CaMKII activation supports reward-based neural network optimization through Hamiltonian sampling. arXiv:1606.00157, 2016

and accepted for presentation in the Workshop "Brains and Bits" at NIPS 2016.

These results suggest a new perspective of network learning:

1. We arrive at a Bayesian model for network plasticity, where a prior (encoding e.g. structural constraints, innate knowledge, previously learned information...) modulates network plasticity.

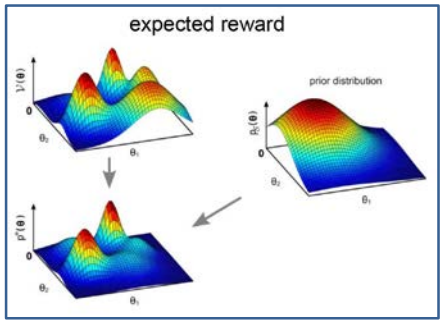

2. Gradient ascent in network fitness is replaced by stochastic sampling from a posterior distribution
3. On the abstract level of reinforcement learning theory our model proposes to replace policy gradient by continuous Bayesian policy sampling

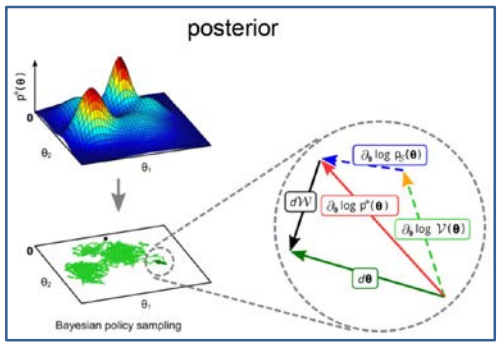

4. This continuous sampling aspect provides automatic compensation for changes in the network or task. Plotted below is the evolution of the vector of network parameters -- projected onto the first 3 principal components--for a motor control task. The colour of the curve indicates the current performance level (red= fast movement completion = good performance). After 24h or learning the motor task was changed, but the network automatically searched for a solution of this new task, that was found hour 28.

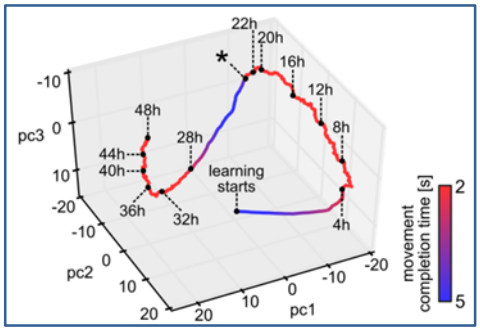

TUGraz also implemented the reward-gated synaptic sampling model in the NEST neural simulator. In collaboration between KIT and TUG a library was developed that allows the model neurons and synapses to interact with the robotics software framework ROS (Robot Operating System) and the Gazebo simulator. The library builds on the existing NEST/MUSIC interface for communicating spike events and analog signal traces. This approach allows us to perform closed-loop simulations efficiently with very small loop delays. The system can

be run in real-time, which is crucial for later applications with real robots. In a first test this system successfully was used together with a Gazebo model of an event-based camera that was developed by KIT to learn a closed-loop balancing task.

Another option for learning in networks of spiking neuron proposed by researchers at IBM was reviewed in:

Maass, W. Energy-efficient neural network chips approach human recognition capabilities. PNAS, 113(40):doi/10.1073/pnas.1614109113, 2016.

In addition, TUGRAZ has investigated options for a hybrid training paradigm for neuromorphic hardware. In this paradigm, neural network weight changes are computed offline in high precision on a digital computer, but these calculations are based on the activations of neurons in the spiking neuromorphic hardware. This training paradigm enables us to utilise the powerful backpropagation training algorithm for neuromorphic hardware training. In one version of this approach, a rate-based coding scheme was employed. Here, TUGRAZ investigated a complementary approach where information is encoded by single spikes. This promises much faster operation under much lower power consumption.

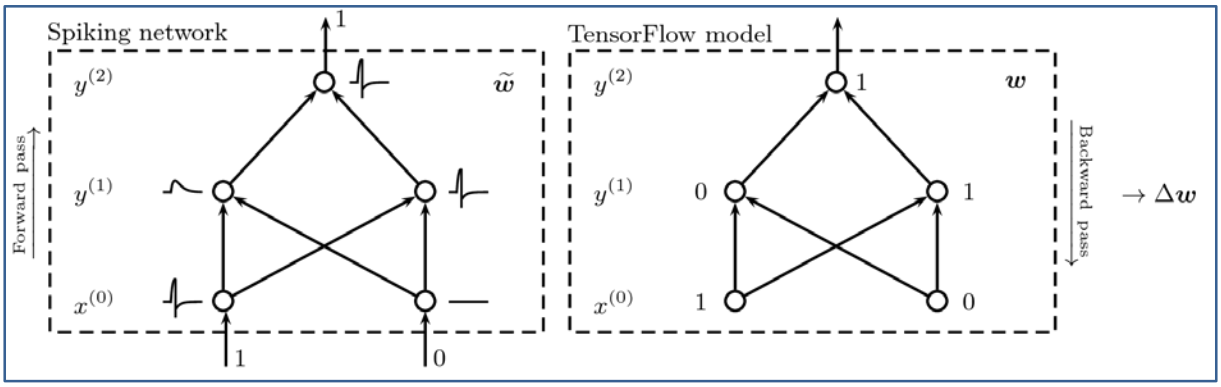

Figure 26: Schemata of the neuromorphic spiking network (left) and the digital computer model (right).

The input layer, hidden layer, and the output layer are denoted by $x^{(0)}$, $y^{(1)}$, and $y^{(2)}$, respectively. The input spikes are presented to the neuromorphic hardware which implements the forward pass. The spike events of the hidden neurons and the spike events of the output neurons are then extracted and interpreted as output of the corresponding artificial neuron in the digital model. The backpropagation step is then performed in the digital model based on the imposed hidden and output layers' outputs. Finally, Δw is used to update the weights in both networks.

The results have been reported in a Master's Thesis:

Limbacher, T. Supervised Learning Algorithms for Spiking Neuromorphic Hardware. Master's Thesis, Graz University of Technology, 2017.

At UMAN a novel mapping of large-scale neural networks onto SpiNNaker was explored:

Knight, J.C., Furber, S.B. Synapse-Centric Mapping of Cortical Models to the SpiNNaker Neuromorphic Architecture. Frontiers in Neuroscience, vol.10, p.420, 14 September 2016

In work towards goal 1, UHEI and TUGRAZ have successfully implemented a deep spiking network architecture on the BrainScaleS wafer-scale system. In [Schmitt et al. 2017], a deep neural network trained in software (using Tensorflow) was converted to a spiking network on the BrainScaleS system. This process was followed by "in-the-loop" training, where in each training step, the network activity was first recorded in hardware and then used to

compute the parameter updates in software via backpropagation. An essential finding was 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 showed an accuracy close to the ideal software-emulated prototype. The presented techniques demonstrate that deep spiking networks emulated on analogue neuromorphic devices can attain good computational performance despite the inherent variations of the analogue substrate.

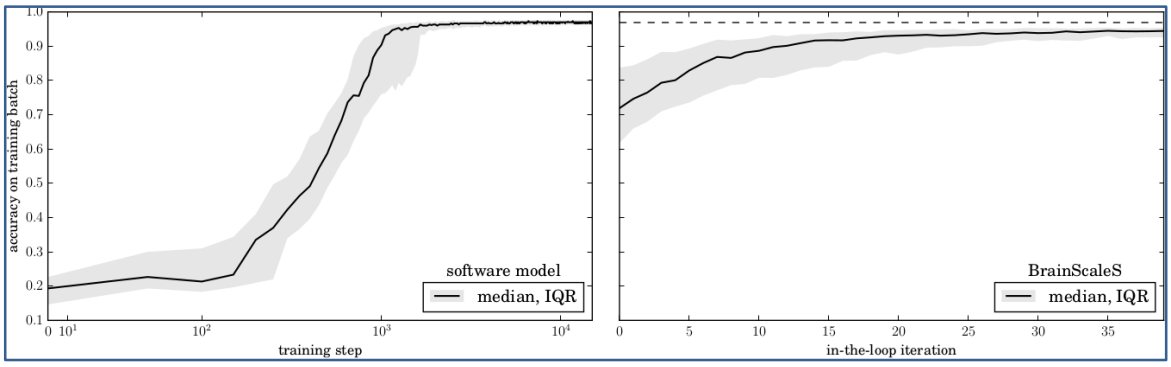

Figure 27: Evolution of the accuracy per training batch for both the software model and the in-the-loop training of the hardware for 130 different sets of hardware neurons and initial weights of the software model.

The total classification accuracy is computed as the sum of correctly classified patterns divided by the total number of patterns in the test set. After 15000 training steps, the accuracy of the software model is 97% with a negligible uncertainty arising from the choice of initial weights. Directly after converting the artificial network to the network of spiking neurons, the accuracy is reduced to 72^{+12}_{-10}%. It increases to 95^{+1}_{-2}% at the end of the in-the-loop training, being close to the performance of the software model with the uncertainty given by the interquartile range (IQR).

[Schmitt et al. 2017] Schmitt, S., Klähn, J., Bellec, G., Grübl, A., Güttler, M., Hartel, A., Hartmann, S., Husmann, D., Husmann, K., Jeltsch, S., Karasenko, V., Kleider, M., Koke, C., Kononov, A., Mauch, C., Müller, E., Müller, P., Partzsch, J., Petrovici, M.A., Schiefer, S., Scholze, S., Thanasoulis, V., Vogginger, B., Legenstein, R., Maass, W., Mayr, C., Schüffny, R., Schemmel, J., Meier, K. Neuromorphic Hardware In The Loop: Training a Deep Spiking Network on the BrainScaleS Wafer-Scale System. Accepted at IJCNN 2017.

5.6 T9.4.3 Large-Scale Organisation of Cognitive Computation

5.6.1 Key Personnel

Task Leader: Wolfgang MAASS (TUGraz)

Other Researchers: Arjun RAO, Christoph POKORNY (TUGraz)

Other Researcher: Eric MÜLLER, Sebastian SCHMITT, Mihai A. PETROVICI (UHEI)

Other Researcher: Gabriel Fonseca GUERRA (UMAN)

5.6.2 SGA1 DoA Goals

1) Understand how distributed brain computations are organized into task-specific functional networks of brain areas that are formed on the fly

2) How is computational load and knowledge shared between brain areas, and what aspects of brain dynamics and architecture support coherent perception and cognition in very large distributed systems;

3) Develop principles that enable users to port these features into artificial cognitive computations.

4) This Task will integrate results achieved in SP3 on Models of Cognitive Processes.

5.6.3 Component Progress

In work towards goals 2 and 3, a spiking neural network model was derived for the emergence of assemblies of "concept cells" in the human MTL, and for the modulation of these assemblies when associations between concepts are formed. In this model, assemblies of concept cells emerge rapidly through synaptic plasticity. Further, when two concepts (e.g. images, or more general patterns) become associated through simultaneous presentation as input to the network, the corresponding two assemblies extend into each other, so that they get a significant overlap. This network feature reproduces recent experimental from the human MTL (see Ison et al., Neuron, 87(1), 2015). The draft of our modelling results is contained in an internal report:

Pokorny, C., Ison, M., Legenstein, R., Papadimitriou, C., Venkatesh, S., Maass, W. Memory-specific assemblies and associations between assemblies emerge in a cortical microcircuit models through STDP. Internal Report, 2017.

A corresponding paper will be submitted for publication within the next few months. Furthermore the feasibility of porting this model onto SpiNNaker was analysed.

In addition, a new model for variable binding in cognitive computations was developed. The model is based on recent fMRI recordings from the human brain, which suggest that subregions of the temporal cortex are dedicated to the representation of specific roles (e.g., subject or object) of concepts in a sentence or visually presented episode. We propose that quickly recruited assemblies of neurons in these subregions (termed neural spaces) represent variables and act as pointers to previously created assemblies in some content space that represent concepts. We provide a proof of principle that the resulting model for binding through assembly pointers can be implemented in networks of spiking neurons, and supports basic operations of brain computations, such as structured information retrieval and copying of information.

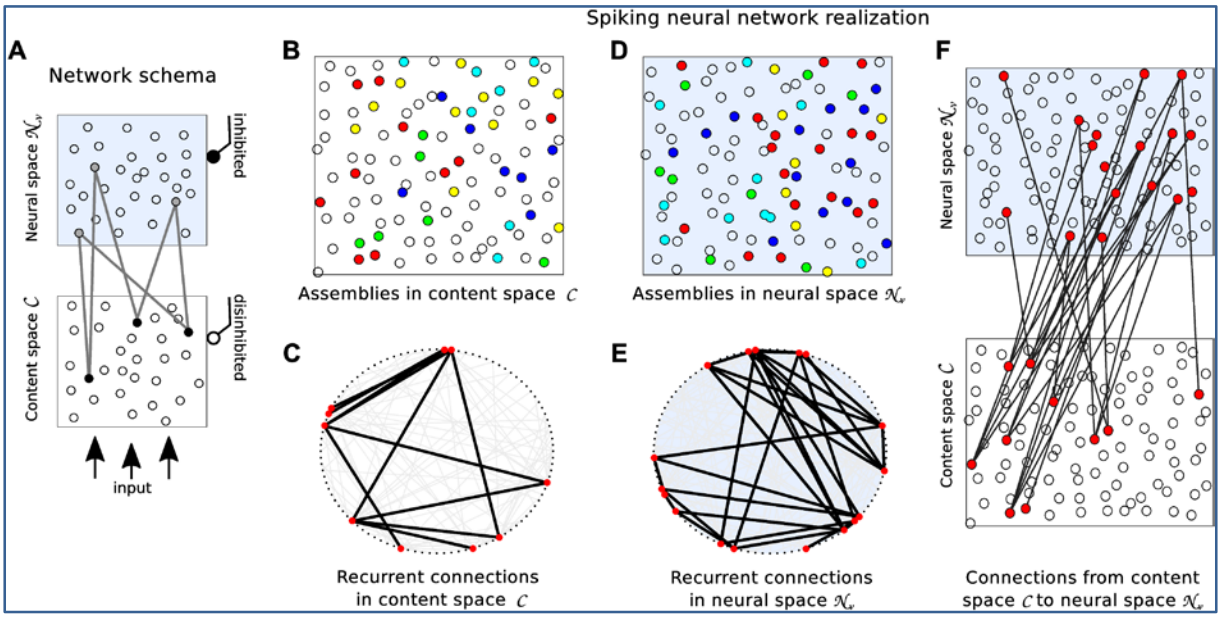

Figure 28: Neural spaces and assembly pointers.

A) Network structure. Rectangles indicate the content space *C* (white) and the neural space N_v for variable *v* (light blue shading). Circles denote neurons (open: inactive; filled black: active; filled gray: potentially active but inhibited). Spaces *C* and N_v also include sparse recurrent connections which are not shown for clarity. Circle on top right of each space indicates disinhibition (filled black circle: inhibited; open circle: disinhibited). Concepts are encoded in content space through neural assemblies. Initially, the neural space N_v is inhibited. Filled gray circles indicate neurons with connections from active neurons in *C*. These neurons may become active when the neural space is disinhibited to constitute an assembly pointer B,D) Assembly code in content space *C* (B) and neural space N_v (D) after induction of assemblies through STDP in a spiking neural network (SNN) model. Filling color indicates assembly identity. C,E) Assembly formation in content space (C) and neural space (E) of SNN. The 100 neurons shown in panel (B) and (D) are rearranged on a circle (black and red dots). Red dots denote neurons of the red assembly and thick black (light gray) lines strong (weak) connections between neurons F) Connections from content space *C* to neural space N_v after an assembly pointer to the red assembly in (B) has been created through STDP.

The model is described in a preliminary report:

Legenstein, R., Papadimitriou, C.H., Vempala, S., Maass, W. Assembly pointers for variable binding in networks of spiking neurons. arXiv preprint arXiv:1611.03698, 2016.

This paper was accepted for presentation at the Workshop "Cognitive Computation: Integrating Neural and Symbolic Approaches" at NIPS 2016. A journal publication based on these results is in preparation. This resulting model can be used for implementing variable binding for cognitive computations on the SpiNNaker hardware.

The role of the two previously sketched models for a new understanding of memory-based cognitive computations that also provides a basis for implementation in neuromorphic devices was described by TUGRAZ and collaborators in the paper:

Maass, W., Papadimitriou, C.H., Vempala, S., and Legenstein, R. Brain computation: A computer science perspective. Invited contribution to Springer Lecture Notes in Computer Science, vol. 10 000, 2017.

This approach will be implemented and tested in a large-scale model in SpiNNaker at UMAN. Preliminary discussions were carried out in order to facilitate porting of the NEST software from TUGRAZ to SpiNNaker.

In addition, UMAN has developed a stochastic neural network that solves Boolean satisfiability problems on SpiNNaker, starting with a solver for Sudoku that reliably finds solutions to the hardest class of Sudoku problems:

Fonseca Guerra, G.A., Furber, S. Spiking Neural Solver for Boolean Satisfiability Problems on SpiNNaker. Frontiers journal, in preparation.

In work towards goals 1 and 2, UHEI has implemented changes to the operating software that offer improved manual control over the mapping process and provide a convenient interface to access the hardware components using the high-level description of the network. This allows more precise characterisation and calibration of single neurons in order to reproduce the logistic activation functions required by LIF sampling.

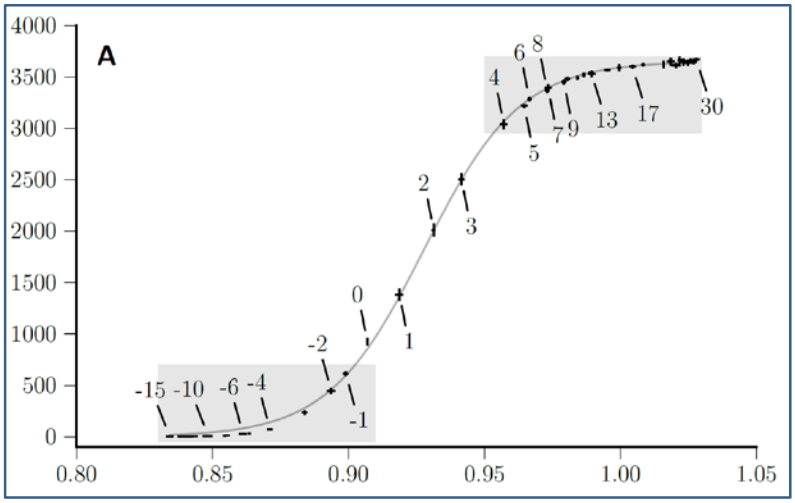

Figure 29: Exemplary activation function (spike rate vs. mean free membrane potential) of a single neuron on the BrainScaleS system.

The sweep was performed by modifying the bias input, which itself was controlled by digital parameters, allowing a high degree of precise control over the neuron.

6. WP 9.5: Platform Training and Coordination

6.1 Key Personnel

Work Package Leader: Karlheinz MEIER (UHEI): also DIR and SIB member

Other Researchers: group leaders from all SP9 partners

Additional roles: Andrew DAVISON (CNRS, Software and Infrastructure groups), David LESTER (UMAN) and Christian MAYR (TUD) as Ethics rapporteurs, Björn KINDLER (UHEI, SP-manager, science coordination group member, community coordinator group member)

6.2 WP leaders overview

What went particularly well?

Several training events have been carried out and introduced newcomers into the principles of neuromorphic computing. The first ever regular university lab tutorial in neuromorphic computing was implemented at the university of Heidelberg and subsequently transferred to the University of Zürich and TU Munich.

The coordination of the SP went very well and created a functioning Subproject with good working communication.

A comprehensive strategy document for neuromorphic computing in HBP has been produced.

What didn't go according to plan?

The regular HBP code jam has been deferred to Y2 of SGA1.

Impact of work done

SP9 is the only project structure worldwide in which operation and development of neuromorphic computing systems is closely integrated with experimental and theoretical neuroscience.

6.3 Priorities for the remainder of the phase

The Training and Coordination WP will continue to provide SP9 internal and inter-SP coordination and to provide training for using the two neuromorphic hardware system types. The integration of the next generation systems into the training and documentation will be of particular importance. In June 2017 the 2nd HBP School for computing will be held under the directorship of the SP9 leader. The school will offer lectures and hands-on training with HBP neuromorphic systems.

6.4 T 9.5.1 Platform Training

6.4.1 SGA1 DoA Goals

The objective of this Task is to deliver documentation and provide training for the use of the neuromorphic platform. Specifically, to carry out events like code jams and other training events.

6.4.2 Component Progress

Training events offered for the BrainScaleS system

We offered a hands-on training event during the "Cognitive Neuromorphic Engineering Workshop" at CapoCaccia (Italy) in spring 2016. Users were introduced to the neuromorphic system and received support and advise for the realisation of their own project ideas on chip. The training event lasted for a full week and was offered every afternoon.

A tutorial for students was developed and set-up at Heidelberg University. The tutorial makes use of the portable system and involves a basis introduction to neuromorphic computing as well as 7 introductory experiments ranging from single neuron response to plasticity in networks. The documentation for the lab class can be found here:

http://www.physi.uni-heidelberg.de/Einrichtungen/FP/anleitungen/F09-10.pdf

The on-going use of the experiment id monitored on this website (in German):

http://www.physi.uni-heidelberg.de/Einrichtungen/FP/buchungsuebersicht.php

The response of about 10 groups that carried out the experiment so far is very positive. The experiment has been exported to the University of Zürich and the TU Munich.

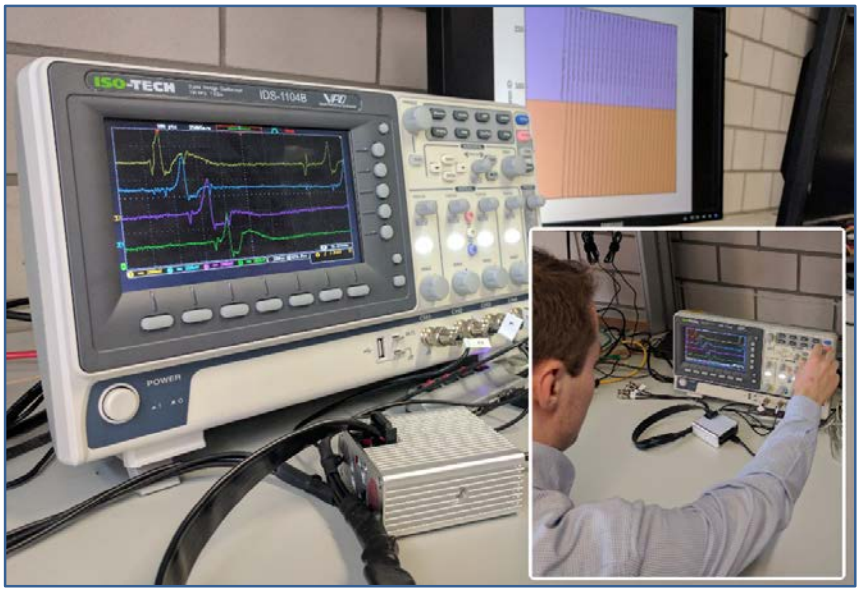

Figure 30: In the Heidelberg-based teaching laboratory, additional readout connections allow students to inspect multiple internal voltages of the analog circuitry simultaneously.

Training events offered for the SpiNNaker system

Two training events were performed, an introduction to SpiNNaker (5-7 September 2016), and an advanced course (8/9 September 2016). There were a total of 21 participants, 2 stated that they were HBP funded and 19 stated that they were not.

6.5 T 9.5.2 Platform User Support

6.5.1 SGA1 DoA Goals

The goal of this task is to provide support for current and interested users of the neuromorphic platform. In addition to extensive e-mail support a living document (neuromorphic guidebook) has been developed and maintained from the beginning of the HBP.

6.5.2 Component Progress

The neuromorphic guidebook received substantial new contributions during Y1 of the SGA1.

Web based version: https://electronicvisions.github.io/hbp-sp9-guidebook/index.html

Printable pdf version: https://flagship.kip.uni-heidelberg.de/SU/p/NM_Guidebook_pdf

https://electronicvisions.github.io/hbp-sp9-guidebook/getting_help.html is pointing users of the systems to the appropriate support channels. Specifically, these are the NeuralEnsemble forum, the Github issue tracker for PyNN bug reports, the mailing list for BrainScaleS wafer-scale hardware and ESS users and the SpiNNaker Users Google group.

In addition, SP9 provided comprehensive support to external users of the single-chip system via email and chat, as well as through personal communication during workshops and visits. The number of external users has grown to 25 scientists and students with remote access to the system. In addition, several external researcher groups possess local Spikey systems.

6.6 T 9.5.3 Platform Coordination

6.6.1 SGA1 DoA Goals

The goal of this Task is to ensure that the Subproject works towards the contractual project goals, delivers results in time and of high quality, perform medium and long term planning and ensure internal and external communication.

6.6.2 Component progress

SP9 has performed all contractual work according to the SGA1 plan. Reporting was delivered complete and in time.

The Subproject has carried out monthly meetings with all partners to coordinate the ongoing work. Quarterly in-person meetings at changing partner locations are complemented by videoconferences in-between the in-person-meeting months. The frequent meeting schedule is an important factor for the information flow and the good communication in the SP.

In-person meetings in SP9 have a structure made of a half-day meeting reviewing scientific progress in SP9 and conveying the political developments in HBP. The other half day is dedicated to a selected scheme on interest. Recent schemes included the SP9 strategy document, theoretical principles with special emphasis on learning and benchmarking.

SP9 went through a major effort to review its scientific and technological goals. The main motivation behind this work was the planning of SGA2 but it also served as a general discussion on the progress in the field and the HBP contributions. The work resulted in a comprehensive strategy document on neuromorphic computing in HBP. The document is appended to this report.

Finally, as part of this Work Package the web services of SP9 are operated and further developed. The SP runs an internal website with a document server, meeting planner and many other functions (https://flagship.kip.uni-heidelberg.de). Also, the general HBP website receives regular updates from SP9 as part of the coordination work.

6.7 T9.5.4 Platform Applications

6.7.1 Key Personnel

Task Leader: Michael SCHMUKER (HERTS)

Other Researcher: Chris HUYCK (MU), Thomas NOWOTNY (UoS)

6.7.2 SGA1 DoA Goals

1) Implement a bio-mimetic classifier on the NM-PM-1 wafer scale system to classify selected high-dimensional data sets, e.g. high dimensional chemical sensor data. This will involve on-board plasticity mechanisms.

2) Assess classifier function against existing SpiNNaker classifier benchmarks (developed during the Ramp-Up Phase) on MNIST and the selected data sets and to Compare classifier systems on large SpiNNaker systems, NM-PM-1 wafer scale, GPU and classical serial CPU with respect to speed and energy consumption (UoS).
3) Extend existing virtual agents developed in the Ramp-Up Phase. The goal is to develop the mechanisms that will enable others to extend these virtual agents on the HBP platforms.
4) Disseminate expertise on building agents.
5) Contribute Use-Case reports to training workshops and participate actively on the user mailing lists for both Platforms.

6.7.3 Component Progress

The work on the classifiers is planned to start in M13.

For the period M1-12 work was performed by Chris Huyck on the virtual agents. The progress is described in the Milestones reports to MS9.5.1 and MS9.5.2

6.7.3.1 MS9.5.1 Open loop virtual agent using HiCANN

From the MS9.5.1 description:

The open loop agent running on the BrainScaleS hardware accounts correctly for three commands. These commands are turn right, turn left, and turn toward the pyramid. Turning toward the pyramid is context sensitive. The neuron connections are all to all.

This system was run 10 times on each of the four commands. In each case, the appropriate neurons fired, and only the appropriate neurons fired.

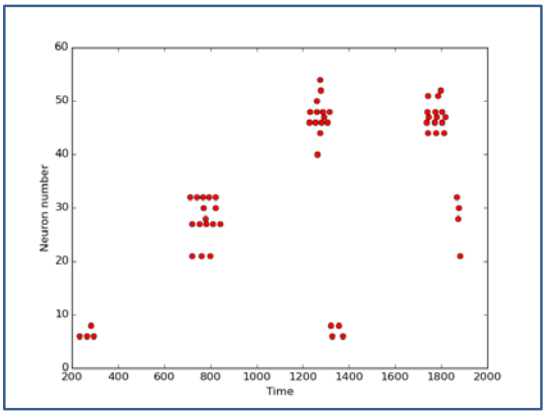

Figure 31: Raster Plot of Neurons on all Four Commands. "Time" is simulated time in ms.

Figure 31 shows a raster plot of the firing behaviour of the neurons on the system running on all four conditions. The tests are left, right, turn to the pyramid with one on the left, and turn to the pyramid with one on the right. The left neurons are 0-19, the right neurons are 20-39, and the pyramid neurons are 40-59. All spikes are correct.

The code can be found at: http://www.cwa.mdx.ac.uk/NEAL/code/cabot1/cabot1A.txt

6.7.3.2 MS9.5.2 Closed loop virtual agent

From the description of Milestone MS9.5.2:

Several agents were developed on the Neurorobotics Platform. Four have been made available ranging from simple to complex. The most complex responds to five user commands, turn left, turn right, move forward, move pyramid and stop. The move pyramid command is context sensitive, needing to find a blue object. In this agent, all processing is

done in simulated neurons. The particular virtual environment is minor modification of an existing NRP environment with a box added.

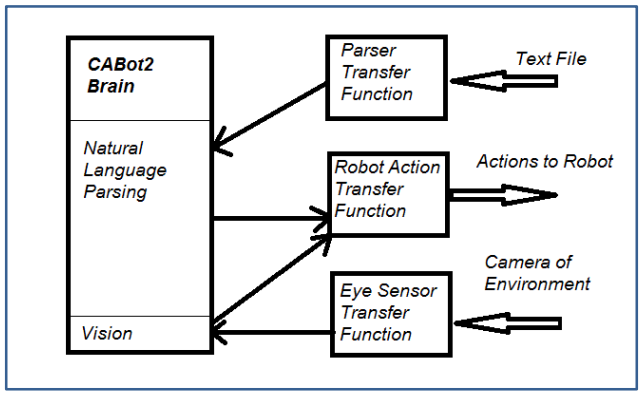

Figure 32 NRP CABot2 System Description

Figure 32 is a representation of the agent. The Neural portion is represented on the left by the CABot2 Brain. Most of the neurons are for text processing, though the final state neurons also act as goals. Vision is managed by two neurons. The input to the "Brain" are managed by two NRP transfer functions: the Parser Transfer Function and the Eye Sensor Transfer Function. Parsing checks the text file to see if there it has been modified; if so, it starts a parse by activating the input neurons in the "Brain". Similarly, the Robot Action Transfer Function polls the goal neurons. If one of them is above an activation value, the goal is set. The Forward, Left and Right goals all go directly to the Husky Robot. The Move Pyramid goal is modulated by the vision neurons.

The Eye Sensor processes the Picture, and if the left side of the picture has blue in it, it will stimulate the left neuron; if the right side has blue, the sensor will stimulate the left neuron. When the goal is move pyramid, the Robot Action Transfer Function polls these neurons. If one is high and the other low, it turns that way. If both are high, it moves forward. If blue is nearby, it moves to the pyramid. If there is no blue, it turns left.

The code can be found at:

http://www.cwa.mdx.ac.uk/NEAL/code/cabot2/NEALCABot2NRPMarch2017.tar

7. Milestones

Table 1: Milestones for SP9

MS No.	Milestone Name	Leader	Task(s) involved	Expected Month	Achieved Month	Comments
MS 9.1.1	Neuromorphic benchmarks 1	UNIBI	T9.1.3	12	12	
MS 9.1.2	Neuromorphic-HPC comparison workflow	Juelich	T9.1.5	12	12	
MS 9.4.1	Computational principles to platforms, part 1	TUGraz	T9.4.1, T9.4.2, T9.4.3	12	12	
MS 9.5.1	Open loop virtual agent using HiCANN	MU	T9.5.4	1	10	
MS 9.5.2	Closed loop virtual agent	MU	T9.5.4	5	12	
MS 9.5.3	First year coordination meetings done	UHEI	T9.5.3	12	12	
MS 9.5.4	Trainings done year 1	UHEI	T9.5.1	12	12	

(Additional details regarding MS9.5.1 and MS9.5.2 are available in the reports stored in the Milestone report repository https://flagship.kip.uni-heidelberg.de/jss/FileExchange?hbpUsID=166)

8. Co-Design Project 5: Plasticity, Learning and Development: Modelling the Dynamic Brain

8.1 Key Personnel

CDP Science Leader: Walter SENN (UBERN)

CDP Implementation Leader: David LESTER (UMAN)

8.2 CDP Leader's Overview (Dave Lester)

The purpose of this CDP is to connect the various strands of research and the various researchers themselves, who have an interest in Plasticity, Learning and Development. To this end, the primary output of this activity is meetings, workshops and scientific papers. The actual implementation activity constitutes just 20% of the overall CDP budget and is shared by LESTER (UMAN), MORRISON (FZJ), and DAVISON (CNRS). In addition Andre GRUENING (USurrey) will undertake significant development work.

- **What went particularly well?** The initial kick-off meeting took place at the EITN in Paris (12/13 May 2016), and was attended by CDP representatives, interested HBP partners, and non-members of HBP. In addition, other CDP-related meetings have taken place:
  - Workshop on Biological Deep Learning, July 16/17, 2016, Bern
  - Joint SP4/SP9 meeting on Dendritic computation in 26/27 Sept 2016 at EITN.
  - CDP5 presentation at HBP Review 11-14 October, 2016 Florence. (M12 deliverable)
  - CodeGen Workshop at Jülich 7-9 December, 2016.
- **What didn't go according to plan?** The cancellation of the work of CDP5 going forwards into SGA2 is disappointing.
- **Impact of work done:** It is critical that the integration of various aspects of the work of the Project as a whole continues.

8.3 Use Case Progress

8.3.1 CDP5-P1 Suite of Benchmark Learning Tasks

Use Case: As a roadmap for the CDP5 activities, we will define network learning tasks that can eventually be implemented on the various Platforms. These benchmark tasks will be designed to test different aspects of learning and memory (i.e. temporally separated and/or multi-modal inputs, reward learning, memory extinction), while enforcing biological constraints. In selecting the learning tasks it will be acknowledged that the brain shows a diversity of synapses and plasticity rules, depending on brain area and neuron type. The network structures will be chosen based on the available knowledge of their biology, plasticity and memory function.

Use Case Leader: David LESTER (UMAN)

Contributing Tasks & Components:

Progress summary: Current progress on this Task is – as intended at this stage – 'piecemeal'; the full, intended, integration and portability is not yet present at M12, i.e. plasticity and learning tasks that perform on SpiNNaker do not yet seamlessly port to BrainScaleS and *vice versa.* The combination into a portable benchmark suite is was always intended as an M24 deliverable.

8.3.2 CDP5-P2 Cellular and Cognitive Determinants of Functional Plasticity

Use Case Description: Various experimental results on synaptic plasticity have been identified in the recent years within and outside the HBP Consortium that potentially improve plasticity rules. Among such results are the modulation of synaptic plasticity by dendritic NMDA-, calcium- and sodium-spikes, the semi-global modulation by astrocytes and glia cells, the strategic co-localisation of synapses at spines, dendritic shafts and the axonal hillock, different types of plasticity rules on the various excitatory and inhibitory neurons, an STDP window that is broadened to the behavioural time scale, modulation by attention, novelty and cognitive engagement, and so forth. These results need to be taken into account, evaluated in a model, and judged for the suitability of porting them to the large-scale systems.

Use Case Leader: Cyriel PENNARTZ

Contributing Tasks & Components:

Progress summary: The progress of the new SP3 has been primarily focused on the SP3 Deliverables, which include a number of components required by this use-case. It is expected that a closer integration of these activities will be demonstrated in the M24 Deliverables.

8.3.3 CDP5-P3 Guiding Platform Design on Functional Plasticity

Description: The process of 'copying' a well-studied plasticity rule to a platform is itself a task that requires a systematic approach, with forth- and back-iterations between theorising about implementational details and small-size test simulations. Product 3 supports the implementation of the benchmark learning tasks (Product 1), using the enhanced learning rules (Product 2), to the Platforms. It requires a re-shaping of the learning rules to the specific implementational details with a re-testing of the functionality on a small-scale system.

Use Case Leader: Andre GRUENING, (USurrey)

Contributing Tasks & Components: T4.3.3, T9.3.2

Progress: As reported in T4.3.3. In short, progress was made with this Use Case in conjunction with the SpiNNaker team. The conclusion reached is that a prototype Use Case is achievable within the time-scales of SGA-1.

8.3.4 CDP5-P4 Concept Showcases in Big Systems

Use Case Description: The final target of the CDP5 activities are the large-scale applications of learning. Within the 2-year time horizon of SGA1, however, it would be too ambitious to expect the up-scaling of the small-system learning tests (Product 3) to showcase the large-scale systems. For the Heidelberg Physical Model and the Cortical Column Model these large-scale demos are scheduled for the SGA2 phase. Nevertheless, simulations with the Multi Core Hard/Software (UMAN) that are merely based on the classical STDP type of plasticity are within reach. The system will also be combined with NEST, the large-scale simulator of spiking in Jülich (SP7).

Specifically, we intend here to demonstrate that the neuromorphic platforms are useful simulators for large-scale learning tasks.

Use Case Leader: Andre GRUENING (USurrey)

Contributing Tasks & Components: T4.3.3, T9.2.2, T9.3.2

Progress summary: As these activities rely on a substantially complete system of plasticity being implemented on each of the Platforms, this Use Case is intended to be completed by M24.

8.4 Priorities for the remainder of the phase

For the remainder of the current funding phase (SGA1) the CDP will continue to operate as it was intended to: running workshops and meetings, and continuing to deliver successfully on all of the commitments we made in the SGA1 proposal.

9. Publications

In the following a list of 28 publications from SP9 is collected.

By partner (publications can be from multiple partners): CNRS: 1, EPFL: 1, JUELICH: 5, KTH: 2, MU: 1, POLITO: 2, TUD: 3, TUGRAZ: 7, UHEI: 10, UMAN: 7

- Yousefzadeh A, Plana LA, Temple S, Serrano-Gotarredona T, Furber SB, Linares-Barranco B. 18 February 2016. Fast Predictive Handshaking in Synchronous FPGAs for Fully Asynchronous Multisymbol Chip Links: Application to SpiNNaker 2-of-7 Links: IEEE Transactions on Circuits and Systems II: Express Briefs (Volume: 63, Issue: 8, Aug. 2016), DOI 10.1109/TCSII.2016.2531092. Partner: UMAN
- Pfeil T, Jordan J, Tetzlaff T, Grübl A, Schemmel J, Diesmann M, Meier K. 18 May 2016. Effect of Heterogeneity on Decorrelation Mechanisms in Spiking Neural Networks: A Neuromorphic-Hardware Study. Physical Review X, DOI 10.1103/PhysRevX.6.021023. Partner: JUELICH, UHEI, RWTH
- Yu Z, Kappel D, Legenstein R, Song S, Chen F, Maass W. I June 2016. CaMKII activation supports reward-based neural network optimization through Hamiltonian sampling. Computational Science - Neural and Evolutionary Computing, https://arxiv.org/abs/1606.00157. Partner: TUGRAZ
- Urgese G. 15 June 2016. Optimizing Network Traffic for Spiking Neural Network Simulations on Densely Interconnected Many-Core Neuromorphic Platforms. IEEE Transactions, DOI 10.1109/TETC.2016.2579605. http://ieeexplore.ieee.org/document/7492182/. Partner: POLITO
- Torre E, Canova C, Denker M, Gerstein G, Helias M, Grün S. 15 July 2016. ASSET: Analysis of Sequences of Synchronous Events in Massively Parallel Spike Trains. PLOS Computational Biology, DOI 10.1371/journal.pcbi.1004939. Partner: JUELICH
- Weidel P, Djurfeldt M, Duarte RC, Morrison A. 3 August 2016. Closed Loop Interactions between Spiking Neural Network and Robotic Simulators Based on MUSIC and ROS. Frontiers in Neuroinformatics, DOI 10.3389/fninf.2016.00031. Partner: KTH, JUELICH
- Furber S. 16 August 2016. Large-scale neuromorphic computing systems (Topical Review). J Neural Eng. 2016 Oct;13(5):051001, DOI 10.1088/1741-2560/13/5/051001. Partner: UMAN
- Petrovici MA, Leng L, Breitwieser O, Stöckel D, Bytschok I, Martel R, Bill J, Schemmel J, Meier K. 18 August 2016. Stochastic inference with spiking neural networks. BMC Neuroscience 2016, 17(Suppl 1):P96, DOI 10.1186/s12868-016-0283-6. Partner: UHEI
- Brocke E, Bhalla US, Djurfeldt M, Hellgren Kotaleski J, Hanke M. 12 September 2016. Efficient Integration of Coupled Electrical-Chemical Systems in Multiscale Neuronal Simulations. Frontiers in Computational Neuroscience, DOI 10.3389/fncom.2016.00097, http://journal.frontiersin.org/article/10.3389/fncom.2016.00097/full. Partner: KTH, KI
- Aamir SA, Müller P, Hartel A, Schemmel J, Meier K. 12 September 2016. A highly tunable 65-nm CMOS LIF neuron for a large scale neuromorphic system. European Solid-State Circuits Conference, ESSCIRC Conference 2016: 42nd, DOI 10.1109/ESSCIRC.2016.7598245. Partner: UHEI
- Knight JC, Furber SB. 14 September 2016. Synapse-Centric Mapping of Cortical Models to the SpiNNaker Neuromorphic Architecture. Frontiers in Neuroscience, DOI 10.3389/fnins.2016.00420, http://journal.frontiersin.org/article/10.3389/fnins.2016.00420/full. Partner: UMAN
- Pfeil T, Jordan J, Tetzlaff T, Grübl A, Schemmel J, Diesmann M, Meier K. 22 September 2016. Effect of Heterogeneity on Decorrelation Mechanisms in Spiking Neural Networks:

A Neuromorphic-Hardware Study. Proceedings of the Bernstein Conference 2016, DOI 10.12751/nncn.bc2016.0166. Partner: JUELICH, UHEI

- Liu Q, Furber S. 30 September 2016. Noisy Softplus: A Biology Inspired Activation Function. Neural Information Processing. ICONIP 2016. Lecture Notes in Computer Science, vol 9950., DOI 10.1007/978-3-319-46681-1_49. Partner: UMAN
- Maass W. 1 October 2016. Searching for principles of brain computation. Current Opinion in Behavioral Sciences (Special Issue on Computational Modelling), DOI 10.1016/j.cobeha.2016.06.003, http://www.sciencedirect.com/science/article/pii/S235215461630119X. Partner: TUGRAZ
- Wolfgang Maass. 4 October 2016. Energy-efficient neural network chips approach human recognition capabilities. PNAS October 11, 2016 vol. 113 no. 41 11387-11389, DOI 10.1073/pnas.1614109113, http://www.pnas.org/content/113/41/11387.full. Partner: TUGRAZ
- Petrovici MA, Bill J, Bytschok I, Schemmel J, Meier K. 20 October 2016. Stochastic inference with spiking neurons in the high-conductance state. Physical Review E, DOI 10.1103/PhysRevE.94.042312, http://journals.aps.org/pre/abstract/10.1103/PhysRevE.94.042312. Partner: UHEI, TUGRAZ
- Liu Q, Pineda-García G, Stromatias E, Serrano-Gotarredona T, Furber SB. 2 November 2016. Benchmarking Spike-Based Visual Recognition: A Dataset and Evaluation. Frontiers in Neuroscience, DOI 10.3389/fnins.2016.00496, http://journal.frontiersin.org/article/10.3389/fnins.2016.00496/full. Partner: UMAN
- Legenstein R, Papadimitriou CH, Vempala S, Maass W. 11 September 2016. Assembly pointers for variable binding in networks of spiking neurons. accepted for presentation at the Workshop "Cognitive Computation: Integrating Neural and Symbolic Approaches", https://arxiv.org/abs/1611.03698. Partner: TUGRAZ
- Siino A, Barchi F, Davies S, Urgese G, Acquaviva A. 8 December 2016. Data and Commands Communication Protocol for Neuromorphic Platform Configuration. 2016 IEEE 10th International Symposium on Embedded Multicore/Many-core Systems-on-Chip (MCSOC), DOI 10.1109/MCSoC.2016.41. Partner: POLITO
- Singh N, Huyck C, Gandhi V, Jones A. 1 February 2017. Neuron Based Control Mechansims for a Robotic Arm and Hand. International Journal of Computer, Electrical, Automation, Control and Information Engineering, http://waset.org/publications/10006430/neuron-based-control-mechanisms-for-a-robotic-arm-and-hand. Partner: MU
- Senk J, Yegenoglu A, Amblet O, Brukau Y, Davison A, Lester DR, Lührs A, Quaglio P, Rostami V, Rowley A, Schuller B, Stokes AB, van Albada SJ, Zielasko D, Diesmann M, Weyers B, Denker M, Grün S. 19 February 2017. A Collaborative Simulation-Analysis Workflow for Computational Neuroscience Using HPC. High-Performance Scientific Computing. JHPCS 2016. Lecture Notes in Computer Science, DOI 10.1007/978-3-319-53862-4_21. Partner: CNRS, JUELICH, UMAN, EPFL, RWTH
- Friedmann S, Schemmel J, Grübl A, Hartel A, Hock M, Meier K. 20 February 2017. Demonstrating Hybrid Learning in a Flexible Neuromorphic Hardware System. IEEE Transactions on Biomedical Circuits and Systems, DOI 10.1109/TBCAS.2016.2579164. Partner: UHEI
- Schmitt S, Klaehn J, Bellec G, Gruebl A, Guettler M, Hartel A, Hartmann S, Husmann D, Husmann K, Karasenko V, Kleider M, Koke C, Mauch C, Mueller E, Mueller P, Partzsch J, Petrovici MA, Schiefer S, Scholze S, Vogginger B, Legenstein R, Maass W, Mayr C, Schemmel J, Meier K. 6 March 2017. Neuromorphic Hardware In The Loop: Training a

Deep Spiking Network on the BrainScaleS Wafer-Scale System. (arXiv.org pre-print of:) Proceedings of the 2017 IEEE International Joint Conference on Neural Networks, https://arxiv.org/abs/1703.01909. Partner: UHEI, TUD, TUGRAZ, UBERN

- Petrovici MA, Schroeder A, Breitwieser O, Grübl A, Schemmel J, Meier K. 12 March 2017. Robustness from structure: Inference with hierarchical spiking networks on analog neuromorphic hardware. IJCNN 2017,https://arxiv.org/abs/1703.04145. Partner: UHEI, UBERN
- Furber S. 15 March 2017. Microprocessors: the engines of the digital age. Proc. R. Soc. A 473.2199 (2017): 20160893, DOI 10.1098/rspa.2016.0893. Partner: UMAN
- Petrovici MA, Schmitt S, Klähn J, Stöckel D, Schroeder A, Bellec G, Bill J, Breitwieser O, Bytschok I, Grübl A, Güttler M, Hartel A, Hartmann S, Husmann D, Husmann K, Jeltsch S, Karasenko V, Kleider M, Koke C, Kononov A, Mauch C, Müller P, Partzsch J, ..., Vogginger B, Legenstein R, Maass W, Schüffny R, Mayr C, Schemmel J, Meier K. 17 March 2017. Pattern representation and recognition with accelerated analog neuromorphic systems. (arXiv.org pre-print of:) Proceedings of the 2017 IEEE International Symposium on Circuits and Syst, https://arxiv.org/abs/1703.06043. Partner: UHEI, TUD, UBERN, TUGRAZ
- Schemmel J, Kriener L, Müller P, Meier K. 21 March 2017. An Accelerated Analog Neuromorphic Hardware System Emulating NMDA- and Calcium-Based Non-Linear Dendrites. (arXiv.org preprint of:) IJCNN 2017, https://arxiv.org/abs/1703.07286. Partner: UHEI
- Neumarker F, Höppner S, Dixius A, Mayr C. 27 March 2017. True random number generation from bang-bang ADPLL jitter. 2016 IEEE Nordic Circuits and Systems Conference (NORCAS), DOI 10.1109/NORCHIP.2016.7792875. Partner: TUD

10. Dissemination

SP9 is very active in disseminating concept and results of neuromorphic computing to scientists and to the public. The following list contains corresponding events in the first year of SGA1.

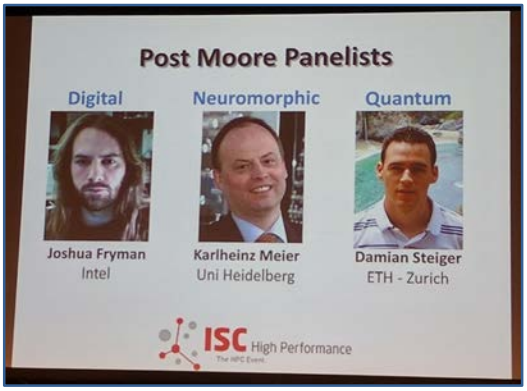

Figure 33: Slide shown by the organisers of the ISC 2017 Post Moore session

10.1.1 Disseminations to the general public

- **Understanding Minds, Brains and Artificial Intelligence**: Talk for the Christians in Science "Engage Science' Events 2017" series. 13 March 2017, Wilmslow, UK. Addressed: United Kingdom. Partner: UMAN. Audience: general public
- **The (past &) future of Computing**: Alumni event. 9 March 2017, San Francisco, USA. Addressed: USA. Audience size: 60. Partner: UMAN. Audience: general public
- **Building Brains**: Talk "The 2017 Gareth Roberts Lecture" at Durham University. https://www.dur.ac.uk/physics/newsandevents/lectures/roberts/lecture2017/ 1 March 2017, Durham, UK. Partner: UMAN. Audience: general public
- **Computers and Brains**: School talk. 25 January 2017, Bolton, UK. Addressed: United Kingdom. Audience size: 30. Partner: UMAN. Audience: general public
- **Computers and Brains**: Keynote at school visit day. 30 November 2016, Manchester, UK. Addressed: United Kingdom. Audience size: 50. Partner: UMAN. Audience: general public
- **Artificial Intelligence**: ARTE TV interview. 7 November 2016, Manchester, UK. Addressed: France & Germany. Partner: UMAN. Audiences: general public, media
- **Computers like brains – A new paradigm in information technology**: Talk at the DE-CIX summit 2016. 13 October 2016, Munich. Addressed: international. Partner: UHEI. Audiences: general public, industry, scientific community
- **Building Brains**: Talk to Cardiff Scientific Society. 5 October 2016, Cardiff, UK. Partner: UMAN. Audience: general public
- **Komplexität und Information - Ist unser Gehirn ein Computer?**: Vorlesung Herbstschule, evening lecture. 8 September 2016, Maria Laach, Germany. Addressed: national. Partner: UHEI. Audiences: general public, scientific community
- **Cognitive Computing**: Lecture at a Manchester University alumni event. 26 July 2016, Manchester, UK. Addressed: United Kingdom. Audience size: 100. Partner: UMAN. Audience: general public
- **Pi with the Prof event**: EuroScience Open Forum ESOF 2016. 25 July 2016, Manchester, UK. Addressed: United Kingdom. Partner: UMAN. Audience: general public

- **Building Brains**: Talk at Bluedot Festival 2016. 23 July 2016, Jodrell Bank, UK. Addressed: United Kingdom. Partner: UMAN. Audience: general public
- **Tech Pioneer, Steve Furber launches Summer Festival**: Viva Computer! festival (http://www.computinghistory.org.uk/news/42101/Tech-Pioneer-Steve-Furber-launches-Summer-Festival/). 26 June 2016, Cambridge, UK. Audience size: 50. Partner: UMAN. Audience: general public
- **SpiNNaker and the Human Brain Project**: Talk to Manchester IET section. Video available at: https://communities.theiet.org/groups/blogpost/view/124/181/3845. 15 June 2016, Manchester, UK. Addressed: United Kingdom. Partner: UMAN. Audience: general public
- **Artificial Intelligence – Episode 35 – The Oxford Comment Podcast**: Podcast. 2 June 2016, https://blog.oup.com/2016/06/artificial-intelligence-oxford-comment/. Addressed: International. Partner: UMAN. Audiences: general public, media, scientific community, civil society
- **HBP**: Camera Interview Deutsche Welle, Hr. Bernert, Berlin. 9 May 2016, Heidelberg. Addressed: national. Partner: UHEI. Audience: general public
- **TV interview**: Interview for Turkish TV TRT World (http://beta.trtworld.com/) Unknown whether it has been used. 29 April 2016. Addressed: Turkey. Partner: UMAN. Audiences: general public, media

10.1.2 Disseminations to the media

- **"The BrainScaleS physical model machine - From commissioning to real world problem solving"**: workshop NICE 2017 conference in Almaden, USA. 6 March 2017, Almaden USA. Addressed: international. Partner: UHEI. Audiences: industry, media, scientific community
- **SP9 exhibition at the science market in the European Parliament** : Exhibition at the STOA (Science and Technology Options Assessment) workshop "UNDERSTANDING THE HUMAN BRAIN - A new era of big neuroscience" on 29-30 November 2016. URL: http://www.europarl.europa.eu/stoa/cms/home/workshops/neuroscience2016. 29 November 2016, European Parliament, Brussels, Belgium. Addressed: Europa. Partner: UHEI, UMAN. Audiences: media, scientific community, policy makers
- **Artificial Intelligence**: ARTE TV interview. 7 November 2016, Manchester, UK. Addressed: France & Germany. Partner: UMAN. Audiences: general public, media
- **Artificial Intelligence – Episode 35 – The Oxford Comment Podcast**: Podcast. 2 June 2016, https://blog.oup.com/2016/06/artificial-intelligence-oxford-comment/. Addressed: International. Partner: UMAN. Audiences: general public, media, scientific community, civil society
- **TV interview**: Interview for Turkish TV TRT World (http://beta.trtworld.com/) Unknown whether it has been used. 29 April 2016. Addressed: Turkey. Partner: UMAN. Audiences: general public, media

10.1.3 Disseminations specifically for policy makers

- **The SpiNNaker Project**: Talk at UK/Japan AI mission in the British Embassy. 31 March 2017, Tokyo, Japan. Addressed: UK & Japan. Audience size: 30. Partner: UMAN. Audiences: industry, policy makers
- **NICE Workshop 2017: Poster "Benchmarking Neuromorphic Hardware"**: Poster Presentation. 6 March 2017, San Jose, CA, United States. Addressed: International. Audience size: 300. Partner: UNIBI. Audiences: industry, scientific community, policy makers

- **SP9 exhibition at the science market in the European Parliament**: Exhibition at the STOA (Science and Technology Options Assessment) workshop "UNDERSTANDING THE HUMAN BRAIN - A new era of big neuroscience" on 29-30 November 2016. URL: http://www.europarl.europa.eu/stoa/cms/home/workshops/neuroscience2016. 29 November 2016, European Parliament, Brussels, Belgium. Addressed: Europa. Partner: UHEI, UMAN. Audiences: media, scientific community, policy makers
- **Presentation of the Neuromorphic Computing Platform (booth) during ESFRI visit to Biotech Campus - Science Market event**: Exhibition. 2 June 2016, Geneva, Switzerland. Addressed: International. Partner: UHEI, CNRS, UMAN. Audiences: scientific community, policy makers

10.1.4 Disseminations to industry

- **The SpiNNaker Project**: Talk at UK/Japan AI mission in the British Embassy. 31 March 2017, Tokyo, Japan. Addressed: UK & Japan. Audience size: 30. Partner: UMAN. Audiences: industry, policy makers
- **"Computer und Gehirn - eine spannende Liaison"**: public evening lecture. 13 March 2017, Nuertingen, Germany. Addressed: national. Partner: UHEI. Audiences: industry, scientific community
- **From ARMs to Brains**: Industry seminar at Google. 9 March 2017, Mountain View, USA. Addressed: USA. Audience size: 25. Partner: UMAN. Audience: industry
- **From ARMs to Brains**: Industry seminar at Apple. 9 March 2017, Cupertino, USA. Addressed: USA. Audience size: 40. Partner: UMAN. Audience: industry
- **"The BrainScaleS physical model machine - From commissioning to real world problem solving"**: workshop NICE 2017 conference in Almaden, USA. 6 March 2017, Almaden USA. Addressed: international. Partner: UHEI. Audiences: industry, media, scientific community
- **NICE Workshop 2017: Poster "Benchmarking Neuromorphic Hardware"**: Poster Presentation. 6 March 2017, San Jose, CA, United States. Addressed: International. Audience size: 300. Partner: UNIBI. Audiences: industry, scientific community, policy makers
- **Brain-inspired computing in the European Human Brain Project**: seminar at Lawrence Berkeley National Laboratory, Berkeley, CA. 3 March 2017, USA, Berkeley. Addressed: international. Partner: UHEI. Audiences: industry, scientific community
- **Workshop on the Exploitation of Neuromorphic Computing Technologies**: Participation to a workshop. URL: https://ec.europa.eu/digital-single-market/en/news/workshop-exploitation-neuromorphic-computing-technologies. 3 February 2017, Brussels. Audience size: 70. Partner: UHEI, TUD, UMAN. Audiences: industry, scientific community, other
- **The ARM Story**: Talk at the Industry Fellows Networking Event (November 2016) of the The Royal Society. 10 November 2016, London, UK. Partner: UMAN. Audience: industry
- **Computers like brains – A new paradigm in information technology**: Talk at the DE-CIX summit 2016. 13 October 2016, Munich. Addressed: international. Partner: UHEI. Audiences: general public, industry, scientific community
- **The brain, the universe and the need for integrated infrastructures**: Keynote presentation for the HBP at DI4R, Krakow. 28 September 2016, Krakow, Poland. Addressed: international. Partner: UHEI, Audiences: industry, scientific community
- **ISC Session: Beyond von Neumann**: ISC2016, Tutorials, Conference & Exhibition. 22 June 2016, Frankfurt am Main, Germany. Addressed: international. Partner: UHEI. Audiences: industry, scientific community

- **SpiNNaker - an architecture for real-time brain modelling**: Talk at the ChipEx2016 (largest annual event of the Israeli microelectronics industry). Slides: http://www.chipex.co.il/_Uploads/dbsAttachedFiles/SteveFurberplenarysession.pdf. 8 May 2016, Tel Aviv, Israel. Addressed: International. Partner: UMAN. Audiences: industry, scientific community
- **"Brain Inspired Computing"**: SAP Kolloquium, Rethinking Business and IT. 15 April 2016, SAP, Walldorf, Germany. Addressed: national. Audience size: 200. Partner: UHEI. Audiences: industry, scientific community
- **"From ARMs to Brains" - Evening plenary talk at the AISB Convention 2016, organised by the Society for the Study of Artificial Intelligence and Simulation of Behaviour**: Talk. 4 April 2016, Sheffield, UK. Addressed: UK, international. Partner: UMAN. Audiences: industry, scientific community

10.1.5 Disseminations to scientific audiences

- **BrainScaleS 2 capabilities poster at L2M proposers day**: Poster with details of the BrainScaleS 2 capabilities at the L2M proposers day, 30 March 2017, Washington DC (USA). Addressed: international. Audience size: 120. Partner: UHEI. Audience: scientific community
- **Modelling the Brain and the Human Brain Project**: Talk at Royal College of Anaesthetists NW branch 25th anniversary meeting. 16 March 2017, Warrington, UK. Addressed: United Kingdom. Audience size: 60. Partner: UMAN. Audience: scientific community
- **Neuromorphic computing - Principles, achievements and potentials**: Keynote at BVM 2017, Heidelberg Bildverarbeitung. 14 March 2017, Heidelberg, DKFZ. Addressed: international. Partner: UHEI. Audience: scientific community
- **"Computer und Gehirn - eine spannende Liaison"**: public evening lecture. 13 March 2017, Nuertingen, Germany. Addressed: national. Partner: UHEI. Audiences: industry, scientific community
- **From carbon to silicon - How real are physical models?**: talk at workshop "Consciousness in humans, animals & robots I: concepts, theories, and neural mechanisms". 10 March 2017, EITN, Paris. Addressed: international. Partner: UHEI. Audience: scientific community
- **"The BrainScaleS physical model machine - From commissioning to real world problem solving"**: workshop NICE 2017 conference in Almaden, USA. 6 March 2017, Almaden USA. Addressed: international. Partner: UHEI Audiences: industry, media, scientific community
- **NICE Workshop 2017: Poster "Benchmarking Neuromorphic Hardware"**: Poster Presentation. 6 March 2017, San Jose, CA, United States. Addressed: International. Audience size: 300. Partner: UNIBI. Audiences: industry, scientific community, policy makers
- **Abstract Presentation SpiNNaker: Large-scale Real-time Neural Simulation**: Talk at the NICE 2017 conference, https://www.src.org/calendar/e006125/. 6 March 2017, Almaden, USA. Addressed: International. Partner: UMAN. Audience: scientific community
- **Brain-inspired computing in the European Human Brain Project**: seminar at Lawrence Berkeley National Laboratory, Berkeley, CA. 3 March 2017, USA, Berkeley. Addressed: international. Partner: UHEI. Audiences: industry, scientific community
- **«Künstliche neuronale Systeme - Eine stille Revolution»**: public lecture. 1 March 2017, Ludwigshafen, Germany. Addressed: national. Partner: UHEI. Audiences: scientific community, civil society

- **"Will there ever be a standard model of the brain?"**: Colloquium at Durham University, Calman Learning Centre. 14 February 2017, Durham, UK. Addressed: international. Partner: UHEI. Audience: scientific community
- **Workshop on the Exploitation of Neuromorphic Computing Technologies**: Participation to a workshop. URL: https://ec.europa.eu/digital-single-market/en/news/workshop-exploitation-neuromorphic-computing-technologies. 3 February 2017, Brussels. Audience size: 70. Partner: UHEI, TUD, UMAN. Audiences: industry, scientific community, other
- **Manycore, past, present and future**: Talk at the kick-off meeting of the MaRIONet (Manycore Research, Innovation, and Opportunities Network), http://manycore.org.uk/London_jan17.html. 11 January 2017, London, UK. Partner: UMAN. Audience: scientific community
- **Neuromorphic Computing (SP9)**: Talk and exhibition at the HBP booth at the US-Brain Meeting. 12 December 2016, Washington (USA). Partner: UHEI. Audience: scientific community
- **Computer und Gehirn - eine spannende Liasion**: Talk and discussion "Wilhelm-Ostwald-Gespräche". 10 October 2016, University Leipzig, Germany. Addressed: national. Partner: UHEI. Audience: scientific community
- **Towards Machine Intelligence**: Lecture in Artificial Cognition research forum. 2 December 2016, Tampere, Finland. Addressed: Finland. Audience size: 30. Partner: UMAN. Audience: scientific community
- **The SpiNNaker Project**: Lecture in Artificial Cognition research forum. 1 December 2016, Tampere, Finland. Addressed: Finland. Audience size: 30. Partner: UMAN. Audience: scientific community
- **SP9 exhibition at the science market in the European Parliament** : Exhibition at the STOA (Science and Technology Options Assessment) workshop "UNDERSTANDING THE HUMAN BRAIN - A new era of big neuroscience" on 29-30 November 2016. URL: http://www.europarl.europa.eu/stoa/cms/home/workshops/neuroscience2016. 29 November 2016, European Parliament, Brussels, Belgium. Addressed: Europa. Partner: UHEI, UMAN. Audiences: media, scientific community, policy makers
- **Building Brains**: Keynote at York Doctoral Symposium (https://www.cs.york.ac.uk/yds/yds2016//programme/). 17 November 2016, York, UK. Addressed: United Kingdom. Audience size: 100. Partner: UMAN. Audience: scientific community
- **The SpiNNaker Project**: Seminar at ETH Zurich. 14 November 2016, Zurich, Switzerland. Addressed: Switzerland. Audience size: 50. Partner: UMAN. Audience: scientific community
- **Human Brain Project Satellite Symposium, Society for Neuroscience Meeting 2016: "Collaborative Neuroscience and Enabling Infrastructure"**: Organisation of the workshop and presentations. 11 November 2016, San Diego, USA. Addressed: Worldwide. Partner: EPFL, CNRS, UMAN. Audience: scientific community
- **Talk "Exploring bio-inspired pattern recognition on neuromorphic hardware" at Human Brain Project Satellite Symposium @ SfN 2016.** : Participation to a conference. 11 November 2016, San Diego. Addressed: worldwide. Audience size: 50. Partner: HERTS, UoS. Audience: scientific community
- **Brain inspired hardware architectures - Can they be used for particle physics?**: Detector Seminar at CERN. 4 November 2016, CERN, Switzerland. Partner: UHEI. Audience: scientific community

- **Anders als Turing - Physikalische Modellsysteme neuronaler Schaltkreise**: Festvortrag 10 Jahre jDPG Dresden. 30 October 2016, Dresden, Germany. Addressed: national. Partner: UHEI. Audience: scientific community
- **The SpiNNaker Project**: Keynote talk at the INNS Conference on Big Data. 24 October 2016, Thessaloniki, Greece. Partner: UMAN. Audience: scientific community
- **Neuromorphic Computing in Odour Space.**: Other. 19 October 2016, University of Hertforshire. Addressed: United Kingdom, Audience size: 30. Partner: HERTS. Audience: scientific community
- **neuromorphic computing**: Interview with Rachel Courtland, IEEE Spectrum. 17 October 2016, Heidelberg, phone. Addressed: international. Partner: UHEI. Audience: scientific community
- **Computers like brains – A new paradigm in information technology**: Talk at the DE-CIX summit 2016. 13 October 2016, Munich. Addressed: international. Partner: UHEI. Audiences: general public, industry, scientific community
- **Panel on Emerging Computer Paradigms**: Paneldiscussion at the Symposium on Emerging Trends in Computing 2016. 11 October 2016, Montreux. Addressed: international. Partner: UHEI. Audience: scientific community
- **Brain derived computer architectures - How much biology do we need?**: Talk at the Academie des sciences, Paris. 4 October 2016, Paris. Addressed: international. Partner: UHEI. Audience: scientific community
- **The man who is building a brain**: Feature in the Winter 2016 issue of The University of Manchester Magazine, https://www.manchester.ac.uk/discover/magazine/features/the-man-who-is-building-a-brain/ 3 October 2016, Manchester, UK. Addressed: United Kingdom. Partner: UMAN. Audience: scientific community
- **Neural hardware architectures for spatio-temporal data processing**: Talk at TWEPP2016. 29 November 2016, KIT, Karlsruhe, Germany. Addressed: international. Partner: UHEI. Audience: scientific community
- **The brain, the universe and the need for integrated infrastructures**: Keynote presentation for the HBP at DI4R, Krakow. 28 November 2016, Krakow, Poland. Addressed: international. Partner: UHEI. Audiences: industry, scientific community
- **Bernstein Conference 2016 workshop on HPC in Neuroscience, talk "Neuromorphic solutions vs. GPUs: practical insights from a bio-inspired pattern recognition task"**: Participation to a conference. 21 October 2016, Berlin, Germany. Addressed: worldwide. Audience size: 50. Partner: HERTS, UoS. Audience: scientific community
- **SpiNNaker - Massively-Paralled Brain Modelling with ARMs**: Talk at the ARM Research Summit 2016 "Computing for the next decade" (agenda https://developer.arm.com/research/summit/previous-summits/2016/agenda). 16 September 2016, Cambridge, UK. Addressed: International. Partner: UMAN. Audience: scientific community
- **Komplexität und Information - Ist unser Gehirn ein Computer?**: Vorlesung Herbstschule, evening lecture. 8 September 2016, Maria Laach, Germany. Addressed: national. Partner: UHEI. Audiences: general public, scientific community
- **The Spinnaker Project**: Key Note Talk at the 16th UK Workshop on Computational Intelligence (UKCI 2016), http://wp.lancs.ac.uk/ukci2016/key-note-talk-2/. 8 September 2016, Lancaster, UK. Addressed: International. Partner: UMAN. Audience: scientific community

- **Simulation des menschlichen Gehirns**: Interview with WIRED Germany, Bernd Skischally. 24 August 2016, Heidelberg, phone. Addressed: international. Partner: UHEI. Audience: scientific community
- **Building Brains: bio-inspired computing**: Panel session ESOF 2016 (Karlheinz Meier, Steve Furber, ...). http://manchester2016.esof.eu/en/the-programme/event-information/building-brains-bio-inspired-computing.html. 27 July 2016, Manchester, UK. Addressed: international. Partner: UHEI, UMAN. Audience: scientific community
- **Talk in the "Department Rankers and Rankings: Truths and Consequences" session at the CRA 2016 Conference at Snowbird**: Talk. 19 July 2016, Snow Bird, USA. Addressed: USA. Partner: UMAN. Audience: scientific community
- **The SpiNNaker Project**: Keynote talk at UCNC'16 (Unconventional Computation and Natural Computation). 15 July 2016, Manchester, UK. Addressed: International. Audience size: 40. Partner: UMAN. Audience: scientific community
- **The SpiNNaker Project**: Talk at Teratec 2016, Slides at: http://www.teratec.eu/library/pdf/forum/2016/Presentations/A4_01_Forum_TERATEC_2016_FURBER_UNIV_MANCHESTER.pdf. 29 June 2016, Paris, France. Addressed: France. Partner: UMAN. Audience: scientific community
- **SpiNNaker Project**: Talk in the "Latest ARM developments in HPC" session during the "Developing Next-gen HPC Architectures - A Hardware Prototyping Workshop" at ISC'2016. 23 June 2016, Frankfurt, Germany. Addressed: International. Partner: UMAN. Audience: scientific community
- **ISC Session: Beyond von Neumann**: ISC2016, Tutorials, Conference & Exhibition. 22 June 2016, Frankfurt am Main, Germany. Addressed: international. Partner: UHEI. Audiences: industry, scientific community
- **Human Brain Project's neuromorphic computing platform**: Interview with Debora MacKenzie, New Scientist. 19 June 2016, Heidelberg, phone. Addressed: international. Partner: UHEI. Audience: scientific community
- **Neuromorphic computing - What are the principles?**: Talk at the Visionary Seminar. 7 June 2016, Imec, Leuven, Belgium. Addressed: international. Partner: UHEI. Audience: scientific community
- **Presentation of the Neuromorphic Computing Platform (booth) during ESFRI visit to Biotech Campus - Science Market event**: Exhibition. 2 June 2016, Geneva, Switzerland. Addressed: International. Partner: UHEI, CNRS, UMAN. Audiences: scientific community, policy makers
- **Artificial Intelligence – Episode 35 – The Oxford Comment Podcast**: Podcast. 2 June 2016, https://blog.oup.com/2016/06/artificial-intelligence-oxford-comment/. Addressed: International. Partner: UMAN. Audiences: general public, media, scientific community, civil society
- **Seminar Talk on "Neuromorphic computing in odor space", Computer-assisted drug design group, ETH Zürich**: Seminar talk. 27 May 2016, ETH Zürich. Addressed: Switzerland. Audience size: 20. Partner: UoS. Audience: scientific community
- **Neuromorphes Computing**: Public talk. 24 May 2016, Heidelberg university, Germany. Addressed: national. Partner: UHEI. Audience: scientific community
- **Brain-inspired computing beyond von Neumann**: Talk at INC12. 12 May 2016, Imec, Leuven, Belgium. Addressed: international. Partner: UHEI. Audience: scientific community
- **SpiNNaker - an architecture for real-time brain modelling**: Talk at the ChipEx2016 (largest annual event of the Israeli microelectronics industry). Slides:

http://www.chipex.co.il/_Uploads/dbsAttachedFiles/SteveFurberplenarysession.pdf. 8 May 2016, Tel Aviv, Israel. Addressed: International. Partner: UMAN. Audiences: industry, scientific community

- **Neuromorphic computing - Extreme approaches to weak and strong scaling**: Keynote at EASC 2016. 27 April 2016, Stockholm, Sweden. Addressed: international. Partner: UHEI. Audience: scientific community
- **After dinner speech**: Talk at the 2016 CPHC (The Council of Professors and Heads of Computing) Conference. 25 April 2016, Manchester, UK. Addressed: United Kingdom. Partner: UMAN. Audience: scientific community
- **"Brain Inspired Computing"**: SAP Kolloquium, Rethinking Business and IT. 15 April 2016, SAP, Walldorf, Germany. Addressed: national. Audience size: 200. Partner: UHEI. Audiences: industry, scientific community
- **Seminar talk @ MPI for Chemical Ecology Jena: "Integrating physical, chemical, and computational aspects of Odour Space."**: Seminar. 13 April 2016, Max-Planck-Institute for Chemical Ecology, Jena. Addressed: Germany. Audience size: 40. Partner: UoS. Audience: scientific community
- **"From ARMs to Brains" - Evening plenary talk at the AISB Convention 2016, organized by the Society for the Study of Artificial Intelligence and Simulation of Behaviour**: Talk. 4 April 2016, Sheffield, UK. Addressed: UK. international. Partner: UMAN, Audiences: industry, scientific community

10.1.6 Disseminations in other categories

- **NM Computing (SP9) exhibition at the Swiss Embassy reception during the US-Brain meeting**: Exhibition. 13 December 2016, Swiss Embassy in Washington (USA). Audience size: 50. Partner: UHEI. Audience: other
- **Driverless Cars**: Talk to barristers at Exchange Chambers. 27 July 2016, Manchester, UK. Addressed: United Kingdom. Audience size: 20. Partner: UMAN Audience: other

11. Education

David Lester (UMAN), Steve Furber (UMAN) and Andreas Grübl (UHEI) contributed lectures to the "ICT for non-specialists" HBP education programme. David Lester also served as course director.

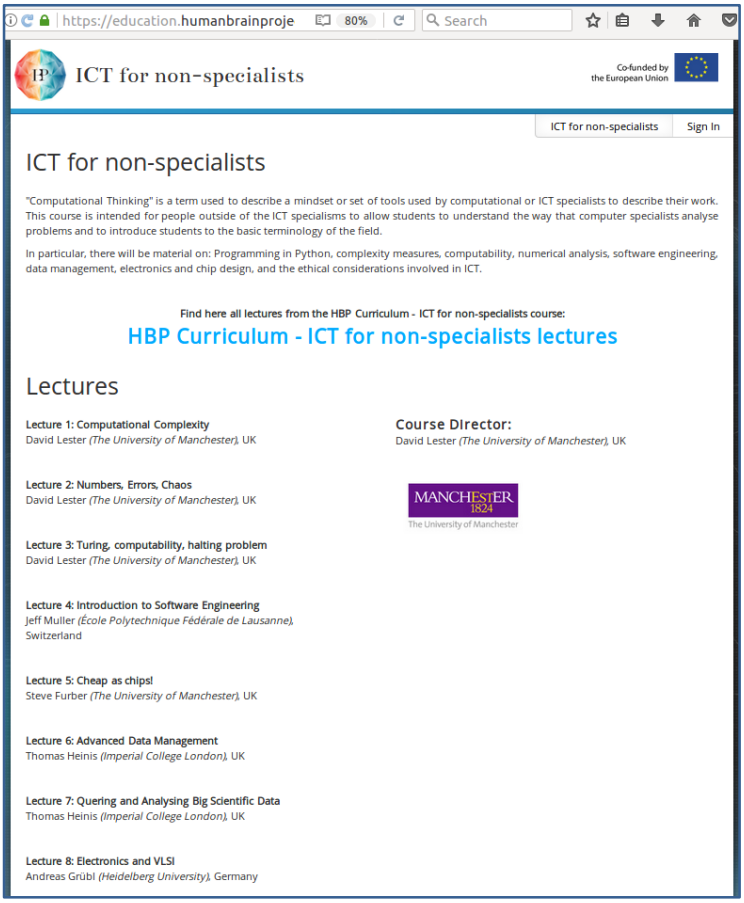

Figure 34: The HBP Education Programme 'ICT for Non-Specialists

The lecture videos are publicly accessible at the URL:

https://education.humanbrainproject.eu/de/web/ict-for-non-specialists

12. Ethics

No specific activities.

13. Innovation

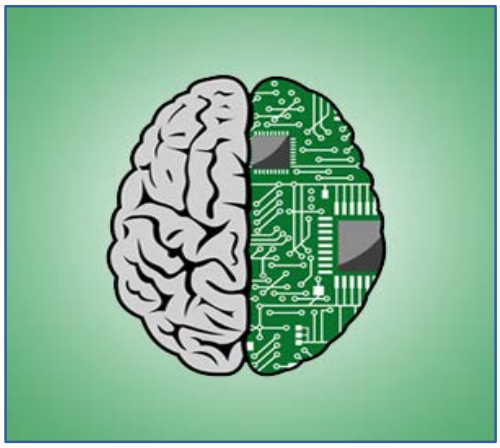

An innovation workshop "Workshop on the Exploitation of Neuromorphic Computing Technologies" has been organised by the FET unit. It took place in Brussels on 3 February 2017: "The workshop presented the neuromorphic technologies developments in the Future and Emerging Technologies (FET) Programme, and in particular in the Human Brain Project. The final report is now available.

The public report has been largely written by Christian Mayr and Karlheinz Meier.

The report is available online:

https://ec.europa.eu/digital-single-market/en/news/workshop-exploitation-neuromorphic-computing-technologies

14. Appendix: The HBP SP9 Strategy document for neuromorphic computing

The attached pdf is the HBP SP9 Strategy document for neuromorphic computing in the version 1.15 from 20 March 2017

The Human Brain Project

SP9 - Neuromorphic Computing strategy

- V0.1-30.08.2016-initial version
- V0.2-01.09.2016-WP9.1 included + changes
- V1.0-09.10.2016-WP9.4 included, WP9.5 updated-first complete version
- V1.1-10.10.2016-modify tasks 9,1,5, 9.5.4, 9.5.5
- V1.2-15.10.2016-modify some task names, swap task leaders of 9.2.2 and 9.2.3, added placeholder for open call partners
- V1.3-20.Jan2017: include changes from documents sent to Karlheinz end of 2016:
    - Michael Schumker document, sent to Karlheinz 20 December 2016
    - Wolfgang Maass document, sent to Karlheinz 21 December 2016. This removes the forth task from WP 9.4
- V1.5 8 Feb 2017: Added the changes to Task 9.1.3 and 9.1.4 provided by Andrew
- V1.6 9 Feb 2017: putting in the VARIATION SUGGESTION 1: All direct costs (but travel, a bit of conference fees and a bit of training) which are not HW related: **removed**, FG added 200.000 Euro HW cost, all other HW cost kept, **PM drastically reduced** to fit to about 92% of SGA1 cost *without* the HW cost. Original numbers left in the text, **task description not adapted to the lower PM numbers**
- V1.7: in-kind contributions of UHEI added. 1FTE move in WP9.4 added. WP9.5 renamed (taken Benchmarks out from the WP title)
- V1.8: Re-added the benchmarking task 9.5.4 (sorry for the confusion)
- V1.9: in-kind contributions added: UMAN, TUGraz. Moved a part of TUGraz contribution from 9.41. to 9.4.2, Deletions in T9.3.3 provided by David Lester. Added the general WP9.3 introduction (provided by Steve). Moved 0.5 FTE in T9.1.2 from MU (now 0 FTE) to KTH (now 0,70 FTE (different cost rates). Added some TUD requested changes. Not included yet is a price change from 1.4 Mio Euro to 2 Mio Euro for the SpiNNaker 2 chips.
- V1.10: 5000 Euro direct cost in 9.1.1
- V1.11: reduced direct HW cost: T9.2.2: 150.000 (instead of 200.000) FG, T9.2.1: 1.150.000 (instead of 1.165.000) UHEI
- V1.12: Taken out questions in the text. Added the changes in WP9.4 (including CNRS not in 9.4.3)
- V1.13 Changes by Michael Schmuker in WP9.5.1
- V1.15 last modification 20 March 2017

SP9 Strategy and Planning Document for SGA2
Version **1.15 (20 March 2017)**

Abstract : This document describes the strategy and planning of subproject 9 on neuromorphic computing in the Human Brain Project (HBP) for the time period from 1.4.2018 to 30.3.2020 (SGA2). It provides a detailed description of the planned work and the required resources to achieve the subproject goals. The work is organized in 6 workpackages.

Introduction

Neuromorphic computing implements aspects of structure and function of neural circuits on electronic substrates. Specifically, SP9 has been following two complementary approaches from the beginning of the Human Brain Project. : The physical model approach and the many-core approach. Both approaches have delivered operational large-scale machines and small-scale evaluation systems by the end of the HBP ramp-up phase (RUP) in March 2016.

The on-going SGA1 phase currently focuses on the operation of the available machines as well as on the development of small-scale demonstrators for the next generation implementation of neuromorphic computing in the HBP. In addition SP9 has integrated a new workpackage on computational principles as well as demonstrator applications as part of the coordination workpackage. The project is also strongly involved in the co-design project 5 (CDP5) on plasticity and learning.

The workplan for the SGA2 presented in this document contains both, a continuation and consolidation of current platform operation as well as a major new step towards the next generation machines. The new and cost-intensive work will be the design and fabrication of the first full-size next generation chip and wafer systems. This development is based on the preparatory work carried out in the HBP during the first 4.5 years (RUP and SGA1) and is a necessary step to maintain Europe's significant role in neuromorphic hardware development and operation. It is therefore considered to be of highest priority for the future work of SP9.

Structure of the subproject

Subproject 9 is organized in 6 workpackages:

Compared to the SGA1 phase a new workpackage (WP 9.5, Application and benchmarks) has been added combing application development and benchmarking that have so far been run under platform operation and coordination.

General information on resources

In this document capital costs are given as EU contributions in units of year 2016 € including 25% overheads. Personnel efforts are given as EU contributions in units of full-time equivalents (FTE). Travel costs are calculated as a lump sum / FTE. Workshops including invitations of external speakers will be covered by workpackage 9.6.1.

Workpackage 9.1: Platform software and operations – WP leader: Andrew Davison

Summary of workpackage objectives

The objectives of WP 9.1 are to operate and maintain the HBP Neuromorphic Platform, maintain and improve the underlying software, and extend the Platform with new functionality. New functionality planned for SGA2 includes:

- Support for the second generation hardware systems
- Robust integration with the Neurorobotics Platform
- Code generation for user-defined neuron, synapse and plasticity models
- Higher-level APIs for cognitive modelling and non-neuroscience applications
- Richer user-interfaces, including visualization and graphical model-building tools

The workpackage is led by Andrew Davison (CNRS)

Task 9.1.1 Software operation and maintenance – Task leader: Andrew Davison

Task 9.1.1 will operate the front-end services of the neuromorphic platform, maintain the functionality developed during this and previous phases of the project, and adapt the platform software as necessary to deal with (i) changes to other HBP platforms and base infrastructure (e.g. Collaboratory services); (ii) increased demand from users. (Note that operation of the hardware systems will be carried out in Tasks 9.2.3 and 9.3.3.)

To support larger numbers of users, and more and larger jobs, we will implement a strategy of platform hardening, including improved monitoring, load-balancing, fail-overs, and demand-driven scaling, with the aim of reaching TRL 8 by the end of SGA2.

The task includes provision of technical support and documentation, in collaboration with Task 9.6.3.

Resources:

CNRS

- ~~1.5 FTE~~ 1,4 FTE
- ~~30 k€ for cloud computing (web, database and monitoring servers for the platform services with redundancy and support for demand-driven scaling).~~ 5000 Euro for commercial hosting of the current microservices architecture frontend services

Components :
From Ramp-up phase and SGA1 :

Neuromorphic jobs database, Neuromorphic Job Queue Service, Neuromorphic Job Manager app, Python client for the Neuromorphic Computing Platform, Provenance viewer for the Neuromorphic Job Manager app, Neuromorphic Quotas service, Neuromorphic Resource Manager app, Neuromorphic Dashboard app, Model simplification service (neuromorphic), Knowledge graph integration for neuromorphic simulations, Neuromorphic Benchmarks results database, Neuromorphic Benchmarks service, Neuromorphic Benchmarks app

New:
Neuromorphic Platform monitoring service, load-balancing/scaling service

Task 9.1.2 Hybrid simulations and robotics – Task leader: Mikael Djurfeldt

Task 9.1.2 will ensure integration of the Neuromorphic Computing Platform with the Neurorobotics Platform and further develop, test and maintain software layers and example use cases, for the interaction between neuromorphic hardware and virtual or real environments.

Neuronal networks interact with environments. They are both studied and used in applications in such a setting. Neuronal networks can be simulated on a PC or supercomputer, or emulated in neuromorphic hardware. Environments can be simulated on a PC or supercomputer, or be constituted of neurorobotic hardware.

This leads to problems and complexity arising when shipping spikes, or other information, between neuronal populations and environments. Some of the problems concern communication protocols, synchronization and timing, avoiding lockups, bandwidth, latency, and organization and structure of sources and targets.

MUSIC is a software framework abstracting some of these complexities and supporting modularity in neuronal data communication. In HBP, the CLE of the Neurorobotics platform uses MUSIC for communication with simulators and neuromorphic hardware. MUSIC is used for visualization (SGA1 task 7.3.2) and in studies of closed-loop systems.

In this Task, the KTH partner will, in collaboration with the MU, UMAN, UHEI and CNRS partners, and the Neurorobotics Platform, further develop MUSIC and associated software layers. A new major release of MUSIC will be released, supporting features that are currently requested from users in the HBP including connecting to a running simulation, dynamic MUSIC port creation and deletion, and online update of connectivity (to support structural plasticity).

The~~ MU and~~ KTH partner~~s~~ will develop and test example use cases/demonstrators based on the full neuromorphic software stack, MUSIC and the CLE from the Neurorobotics platform. This will provide examples for users to build upon, ensure the integrity of the full software stack as well as identify problems and limitations during its development.

Resources:

- KTH – ~~1.5 FTE~~ 0,7 FTE

- MU – ~~1 FTE~~ -0 FTE
- UMAN – ~~0.25 FTE~~ 0,1 FTE

Components :

From Ramp-up phase and SGA1 :
MUSIC, ROS-MUSIC interface

New:
Dynamic MUSIC port creation/deletion, online connectivity update

Task 9.1.3 Neuromorphic software development tools – Task leader: Andrew Davison

From a user perspective, "programming" neuromorphic hardware consists of specifying neuronal and synapse parameters, network connectivity, plasticity models, inputs and outputs. For the HBP Neuromorphic Computing Platform, this specification takes the form of either (i) a Python script using an implementation of the PyNN API; (ii) an XML file using the NeuroML or NineML standards, which is then interpreted by PyNN.

For the SpiNNaker hardware, in addition, a feasibility study in SGA1 demonstrated that custom neuron/synapse/plasticity mechanism models expressed in NineML or LEMS format can be used to generate low-level SpiNNaker code. The same feasibility study demonstrated that NineML plasticity models can be mapped to the NM-PM-2 plasticity processor.

The abstraction level of the PyNN API is well suited for use by neuroscientists, particularly for models at the scale of brain regions and smaller. For larger scale neuroscience models, cognitive models, and non-neuroscience applications (e.g. machine learning), higher-level abstractions will be needed to increase user productivity.

Task 9.1.3 will therefore have four main sub-tasks:

(1) continued development and maintenance of the PyNN API, implementations of the API for the NEST and NEURON simulators (the implementations for the two neuromorphic hardware systems will be developed and maintained by the workpackages dedicated to those systems), and support for NeuroML and NineML network descriptions. This will include further development of libcsa, the C++ implementation of the Connection Set Algebra (CSA), used for advanced network connectivity specifications.

(2) full integration of NineML and LEMS code generation support into the Platform user workflow

(3) exploration of potential higher-level APIs for larger scale neuroscience models, cognitive models, and non-neuroscience applications, driven by user consultations and in collaboration with WPs 9.4 and 9.5.

(4) building a library of modular sub-networks which can be composed into larger applications. Currently, network models exist independently of each other, and must be assembled manually and trained separately. A modular architecture could greatly reduce the required effort, allowing less experienced users to assemble brain-like architectures at the abstraction level of a flow-graph. This sub-task requires development of an interoperability layer to interface spiking neural networks, which unifies the existing assortment of spike train data-encodings.

Resources:

- CNRS ~~2 FTE~~ 1,5 FTE
- KTH ~~0.5 FTE~~ 0,25 FTE
- UNIBI ~~0.5 FTE~~ -0,5 FTE

Components:

From Ramp-up phase and SGA1:
PyNN, libNeuroML, libNineML, Neo, pyNN.neuroml, pyNN.nineml, libcsa

New:
code generation software for SpiNNaker; code generation software for BrainScaleS plasticity processor, prototype API for cognitive modelling, prototype API for a non-neuroscience application, library of modular and composable sub-networks

Task 9.1.4 Extended user interfaces – Task leader: Ulrich Rückert

The principal user interface for the platform at the end of SGA1 is a job submission system, accessible both through a web browser and through scripting (e.g. in a Jupyter notebook), and using Python source code or XML documents as user input. Task 9.1.4 will extend this user interface in three ways:

Tighter integration with Brain Simulation and HPAC platforms.

A typical project using HBP infrastructure might involve HPC simulation using the Brain Simulation and HPAC platforms, neuromorphic simulation/emulation using the Neuromorphic Computing Platform, and data analysis using the HPAC platform. A convenient and productive user workflow will require a single workspace to interact with all the results. This sub-task will collaborate with SP5, SP6 and SP7 to develop such an integrated workspace as a Collaboratory app. [Such an approach will require a project-wide, universal provenance store, preumably the SP5 Knowledge Graph] This will also enable the use of visualization tools developed in SP6 and SP7 for neuromorphic simulations.

Graphical tools for model building.

While power users will most likely continue to prefer Python scripting as a method for building models and describing computational experiments, students and non-computational scientists may be more productive with a graphical interface. This sub-task will develop a number of such interfaces and deploy them as Collaboratory apps.

The first app will allow building networks at the level of the PyNN API (populations of neurons, and patterns of connections between populations).

A second groups of apps will provide existing applications developed in previous project phases to non-expert users. Therefore, existing applications must be adapted for use in a common software framework. In its final form, a user should be able to choose a computational task (e.g. classification, association, pattern completion) and upload training and test data. The software will then automatically prepare the network and execute it on the BrainScaleS or SpiNNaker system, depending on the user's choice. Afterwards, the results will be available to the user, for example as an interactive visualisation.

The final app will allow graphical assembly of the modular sub-networks, which are developed in Task 9.1.3. In this way, networks for compound neuroscience tasks can be specified in a notation similar to domain-specific planning or overview diagrams. The generated PyNN implementation of the network can be made available as a starting point for further refinements of the implementation by advanced users, if desired. We expect synergies through collaboration with Work Package 9.4 (Theory) and aim for a proof-of-concept tool with representative example applications.

Direct reserved access to the neuromorphic hardware.

For some applications, particularly those involving real-world robotics, direct interactive or semi-interactive access to neuromorphic hardware is required. This sub-task will make such direct access available to HBP members, and implement the required extensions to the quota system to ensure equitable access to resources.

Resources:

- UNIBI ~~1.5 FTE~~ 1FTE
- CNRS – ~~1.5 FTE~~ 1 FTE
- UMAN - ~~0.25 FTE~~ 0,1 FTE
- UHEI – ~~0.25 FTE~~ 0,2 FTE

Components

From Ramp-up phase and SGA1:
Neuromorphic Quotas service, Neuromorphic Resource Manager app

New:
Integrated modelling workflow app, graphical neuromorphic model-building app,
Neuromorphic computing direct access service.

WP9.1 in kind contributions

- UHEI: 2FTE

Workpackage 9.2: BrainScaleS systems – WP leader: Johannes Schemmel

The goal of WP 9.2 is to develop, prototype, manufacture, assemble, test and operate next generation hardware systems to implement massively parallel, physical models of brain cells, circuits and networks. The workpackage has 3 tasks.

Summary of the WP's objectives

WP 9.2 has 3 major objectives .

First, wafer-scale integration of the HICANN-DLS second generation neuromorphic microchip and its integration in the BrainScaleS platform, in particular :

- complete L1 wafer-level communication circuits
- tape-out MPW prototype of HICANN-DLS,
- create post-processing masks
- create tape-out data for wafer-version of HICANN-DLS (no bond-pad version)
- manufacture wafer, post-process and cut, test and mount in existing NM-PM1 platform
- integrate HICANN-DLS in platform software

Second, development of all necessary components for the 30cm wafer-module for the subsequent construction of the final BrainScaleS systems during SGA3, particular :

- finalize routing protocol
- prototype routing ASICs (miniASIC and FPGA)
- enhance mapping to include new routing features
- ~~develop high-speed serdes (5Gb/s) with direct wafer-to-wafer capability and fast power-down states~~

Third, development of the 30cm wafer embedding technology, in particular :

- build NM-PM2 based test modules
- create and evaluate a small test system (10 wafers) as proof of concept

The corresponding workpackage structure consisting of 3 tasks is similar to the one currently implemented in SGA1, with the exception of task 4, which was a start-up effort linked to CDP5.

Task 9.2.1 – Circuit development – Task leader: Johannes Schemmel

Task 9.2.1 will develop all VLSI circuits necessary for the implementation of the NM-PM roadmap as described in the approved FPA. It will produce full-size prototype chips to verify their operation. It will design the necessary hardware for testing the prototype chips and develop the soft- and firmware associated with these test systems. Chip development will be carried out in close collaboration with the relevant other tasks of the workpackage.

Task-leader: Johannes Schemmel

Use of resources in task 9.2.1, direct costs:

UHEI (615 ~~1165~~ k€):

- two prototype systems for testing and further developing the NM-PM2 circuits: 35 k€ ASIC manufacturing, 5 k€ test setup
- wafer version of NM-PM2 system: 121 k€ ASIC manufacturing and 4 k€ test setup
- mask-set and initial wafer production: 450 k$ (date of offer: 3/2017)

TUD (23 k€):

- one prototype system for testing the NM-PM2 high-speed serdes: 18 k€ ASIC manufacturing, 5 k€ test setup (including chip-to-chip high-speed tests simulating wire and connector structure of a >1000 wafer machine)

Use of resources in task 9.2.1, personnel :

UHEI (~~6 FTE~~ 3 FTE):

- chip development HICANN-DLS
  - complete L1 wafer-level communication circuits
  - tape-out MPW prototype of HICANN-DLS
  - production, post-processing and test of the full set of NM-PM2 wafers for the platform upgrade (create tape-out data for wafer-version of HICANN-DLS (no bond-pad version) and post-processing masks, wafer test and integration)

- chip development 30cm wafer and routing
  - finalize routing protocol (specification of the wafer-to-wafer protocol for the PM3 wafer event router circuits
  - prototype routing ASICs (miniASIC and FPGA)

TUD (0.5 FTE)

- chip development and SerDes IP Integration
- lab evaluation of serial link testchip from SGA1, serial link IP integration into HICANN DLS, contribution to serial link verification for HICANN DLS.

Task 9.2.2 – System development – Task leader: Andreas Grübl

Task 9.2.2 will develop the main printed circuit boards (PCBs) and all necessary auxiliary components to prototype complete, operational wafer-modules. It will develop novel packaging technologies according to the NM-PM roadmap. It will provide the technical documentation for operation and manufacture of the wafer-modules.

It will engineer the compute and network infrastructure for the NM-PM systems. It will develop the system management and communication soft- and firmware.

It will provide the documentation of all relevant APIs for integration of the NM-PM system in the SP9 software infrastructure as it will be defined by WP9.1.

Task-leader: Andreas Grübl, UHEI

Use of resources in task 9.2.2, direct costs:

UHEI (20k€):

- post-processing for the NM-PM2 wafers for the platform upgrade from NM-PM1 to NM-PM2: *costs in the Fraunhofer direct cost*
- prototype hardware for the platform upgrade: 20 k€

TUD:

- ~~development of communication module upgrade for NM-PM-2~~

FhG:

- prototype manufacturing for 30cm wafer embedding plus post-processing for NM-PM2: 150.000 *Euro total*

Use of resources in task 9.2.2, personnel :

FhG (1 FTE):

- extend direct wafer embedding technology to 30cm wafers, using NM-PM2 wafers as test vehicles

TUD (0.5 FTE):

- update the firmware of the communication subsystem for NM-PM2
- ~~develop an upgrade of the communication module~~

UHEI (~~6 FTE~~: 3 FTE):

Platform hardware development:

- hicann-dls wafer-module integration
  - upgrade of the platform hardware (power supply, system monitoring etc.) to accommodate the NM-PM2 wafers
  - update a wafer-module for testing the NM-PM2 wafers
  - update the firmware of the communication subsystem for NM-PM2
  - specify the communication module upgrade

- manufacture and test of two prototype modules
- develop and test printed circuit board layouts for the development of 30cm direct wafer embedding technology

Platform software development

- NM-PM2 specific software for integration into HBP framework
  - hicann-dls integration
  - hicann-dls mapping integration
  - hicann-dls ESS

- enhance mapping to include new routing features for routing ASICs

Task 9.2.3 – Production, operation and maintenance– Task leader: Eric Müller

Task 9.2.3 manages the production and operation of the NM-PM platform hardware. It is responsible for the large-scale manufacturing and test of all system components according to the NM-PM roadmap and the results from Tasks 9.2.1 and 9.2.2.

It will operate the NM-PM platform hardware during the duration of the project as part of the Neuromorphic Computing Platform. It will provide technical personnel to provide technical assistance to remote users (made explicit in the coordination workpackage 9.6).

Task-leader: Eric Müller, UHEI

Use of resources in task 9.2.3, direct costs:

UHEI (10 k€):

- hardware maintenance of NM-PM1 platform : 10 k€

Use of resources in task 9.2.3, personnel:

UHEI (~~6 FTE~~ 3 FTE):

- platform Technology Readiness Level (TRL) improvement
  - calibration
  - software stack
  - applications (user demos)
  - software development and maintenance in close interaction with user feedback
  - applications (user demos) HICANN-DLS

- platform hardware maintenance
  - repair and replacement
  - installation of the NM-PM2 wafers
  - technology hardening

- platform user support
  - user education for NM-PM2

WP9.2 in kind contributions

- UHEI: 3 FTE
- TUD:
  - 0.5FTE: design flow setup, operation and support, lab infrastructure and operation (e.g. for serial link characterization);
  - 0.5FTE: Neuromorphic benchmarking for NM-PM1/PM2 system application

Workpackage 9.3: SpiNNaker systems – WP leader: Steve Furber

The goal of WP 9.3 is to develop the next generation many-core hardware microchip to implement massively parallel, many-core models of brain cells, circuits and networks, and to extend the software support for the current many-core system.

Summary of the WP's objectives
WP 9.3 has 3 major objectives .

- First, circuit development of the second generation SpiNNaker2 many-core microchip:
  - RTL design & verification
  - physical design, synthesis and P&R
  - package and board-level concept development & design
  - serial link and memory interface integration
  - tape-out and testing of prototype device and tape-out of final chip
- Second, architecture development and verification for the second generation chip :
  - virtual (FPGA) prototype
  - architecture and verification
- Third, maintain the current platform and extend the capabilities of the software support :
  - mapping software, including support for CSA
  - extended support for accurate numerics, analytics & plasticity, MUSIC
  - system maintenance and support

The corresponding workpackage structure consisting of 3 tasks is similar to the one currently implemented in SGA1.

Task 9.3.1 - Circuit development – Task leader: Sebastian Höppner

In this task the circuit design and SpiNNaker2 chip implementation for the NM-MC2 system are executed. The test-chip from SGA1 is integrated in a small scale evaluation and verification platform for SpiNNaker2. This allows for low-level firmware development and integration of SpiNNaker2 into the high-level software stack. Another small-scale test-chip containing the final processor component is implemented. Special emphasis is put on serial links for spike communication, the memory interface integration and the on-chip communication fabric (NoC). By this, all custom circuit components will be silicon-proven before the tape-out of the SpiNNaker2 chip. Finally, the SpiNNaker2 chip, containing a large number of processor cores, is implemented.

The work includes RTL design entry and generation, verification by digital circuit simulation, logic synthesis, place & route, timing and power optimization and verification, IR-drop analysis and physical sign-off (DRC, LVS). A design-for-test (DFT) concept is developed and implemented which includes architectures for efficient test coverage of the produced chips, such as built-in-self-test (BIST) circuits for memories, data transmission and clock generators. Approaches for redundancy and built-in-self-repair for yield enhancement are evaluated.

The chip is to be taped-out as full mask set (full wafer). Jointly to the chip design activities a packaging concept is developed for close processor and memory integration for the final SpiNNaker2 platform PCBs and modules.

Task 9.3.1 resources:

Personnel

- 1.25 FTE RTL design & verification (TUD)
- 1.25 FTE physical design, synthesis, P&R (TUD)
- ~~1 FTE package concept development, board/package design, small-scale system, Nanolink_2 (quad-core) deployment (TUD)~~
- 0,75 FTE serial link and memory interface integration (TUD)

Capital expenditure

- Nanolink28_3 tape-out, packaging and PCB setup 140 k€
- SpiNNaker2 tape-out 1400 k€

Total resources: ~~5 FTE~~ + 1.54 M€ 3,25 FTE

Task 9.3.2 - System development – Task leader: Steve Furber

This task is concerned with the top-level architectural design of the SpiNNaker2 chip (carrying this task over from SGA1), and the extensive verification work that is required to establish that the detailed design is correct prior to tape-out.

The primary verification methodology is to develop prototype application software which can be run on the SpiNNaker2 Verilog model within the TUD design environment. The verification tests will be done by UMAN and run remotely on the TUD system. This is highly demanding and time-consuming work as the Verilog models will run very slowly, so the tests must be designed with great care to deliver useful results within a practical time. Where problems are identified extensive detailed debugging will be required, and appropriate changes made to the Verilog model.

Aspects of this work can be accelerated by developing a virtual prototype of the SpiNNaker2 chip on an FPGA platform. This can be synthesized from the same Verilog source and will run much faster than the software Verilog model, but it will capture only part of the system for verification. The FPGA prototype is maintained by TUD and low-level software is developed, which uses the new hardware features with focus on the neuromorphic application, such as dynamic power management and hardware acceleration of specific functions (e.g. fixed-point exponential).

The SpiNNaker2 chip will have in the region of a billion transistors, all of which must be tested for correct functionality prior to tape-out; hence the requirement for 8 person-years of verification effort.

Task 9.3.2 resources:

Resources

- TUD: ~~1 FTE~~ 0,75 FTE virtual (FPGA) prototype and low-level software (TUD)
- UMAN: ~~3 FTE~~ 1,8 FTE architecture & verification (UMAN)

Total resources: ~~4 FTE~~ 2,25 FTE

Task 9.3.3 Production, operation and maintenance– Task leader: David Lester

This task provides the hardware maintenance and operations support for the current SpiNNaker1 platform. In addition software will be maintained for the current platform and ported to the new SpiNNaker2 platform which is being developed in WP 9.3.

The primary software support provided by this task is the software that takes a PyNN, nineML or Nego script and translates the script program into a form capable of being run on SpiNNaker1 or SpiNNaker2 as appropriate. The software will therefore distribute a compressed version of the PyNN script to each core, where it will be expanded into the required form for execution. A facility to snap-shot the state of the computation will be required, and it is possible – given the size of the data-sets envisaged – that on chip data analysis will need to be performed.

In amongst the software maintenance tasks we will need to provide in-house numerical analysis expertise, as best efficiency with SpiNNaker occurs when calculations are performed in fixed-point arithmetic rather than using floating-point arithmetic. Better load and execution performance will be obtained by better place & route algorithms, so provision is made for these to be more fully explored than is done currently. If MUSIC and Connection Set Algebra (CSA) support has not already been incorporated into the tool chain by the start of SGA-2, then this too will need to be budgeted for; it will vastly improve the load-performance of the system for probabilistically specified projections between populations of neurons.

Finally, we need the ability to make required changes to the software to support cross-SP activity. Currently, we envisage this to include providing support for Neurorobotics (as part of CDP-1), providing the links into SpiNNaker for use of SP7 visualization tools, and of course supporting the requirements generated by CDP-5.

An approximate budget is:

- ~~Porting of SpiNNaker-1 software to SpiNNaker-2 (24PM)~~
- ~~Embedded micro O/S (12PM)~~
- Improved Place and Route (24PM)
- Numerics (12PM)
- Probability and Data Analytics (12PM)
- CSA (= Connection Set Algebra 12PM)
- Plasticity (12PM)
- General Support (24PM)
- Server (= host) (12PM)
- MUSIC (12PM)
- ~~Visualization (hooks for links to SP7 12PM)~~
- ~~Robotics (hooks for links to SP10 12PM)~~
- ~~Graphical machine status/utilization (12PM)~~

Task 9.3.3 resources:

Resources:
- UMAN: ~~7 FTE~~ 4,2 FTE software development & support (UMAN)
- POLITO: 1 FTE software development & support (POLITO)

Capital expenditure:

- SpiNNaker1 maintenance 100 k€
- SpiNNaker1&2 dissemination boards 100 k€
- SpiNNaker host system & network 100 k€

Total resources: ~~8 FTE~~ 5,2 FTE + €300k

WP9.3 in kind contributions

- UMAN: 5 FTE (PhD students) = 120 PM
- TUD:
  - 0.5FTE: design flow setup, operation and support, MPSoC infrastructure integration, technical administration of foundry contact and MPW handling.
  - 0.5FTE: Neuromorphic benchmarking for NM-MC2 system application

Workpackage 9.4: Computational Principles – WP .leader: Wolfgang Maass

The goal of WP9.4 is to develop computational principles and learning methods that support the use of large neuromorphic systems, that help to exploit their specific advantages, and which can guide the design of the final version of neuromorphic systems for SGA3. Experts for theory and neuromorphic hardware will closely collaborate in this WP. This research will build on the work during SGA1, make use of insight from other SPs, and integrate new results from Neuroscience, Cognitive Science, and Machine Learning. Resulting efficient simulations of large brain networks on neuromorphic hardware will enable us to evaluate different hypotheses about the organization of computation and learning in brain systems, and thereby support other SPs

Main objectives of this workpackage:

First, development of principles of stochastic computation and probabilistic inference that support computational uses of the neuromorphic systems produced in WP 9.2, 9.3, and 9.5 during SGA2, and which can guide the design of updated neuromorphic systems for SGA3. Task 9.4.1 (Stochastic neural computation) is dedicated to achieving these objectives

Second, the development of learning methods that enable large neuromorphic systems to learn a variety of computational and cognitive functions. We want to enable neuromorphic systems to learn not only with the help of a teacher, but also from its own observations, play, internal and external rewards, and to transfer knowledge from previously learnt tasks. Task 9.4.2 (Learning) is dedicated to these objectives.

Third, development of functional architectures that enable large neuromorphic systems to act as artificial brains, encompassing higher cognitive functions such as the formation of memory and associations between memory items, use of this memory for solving problems, abstract reasoning, and planning. The corresponding Task 9.4.3 (Architectures for neuromorphic intelligent systems) will be based on results achieved in Tasks 9.4.1 and 9.4.2, and will prepare the ground for the design of an artificial neuromorphic brain during SGA3.

Task 9.4.1 Stochastic neural computation – Task leader: Mihai Petrovici

Concrete points to be investigated:

- Development of a method for soft neural sampling and computation with spatio-temporal patterns, where random variables are represented by assemblies of neurons or spatio-temporal patterns, rather than single neurons.

This coding method is likely to support the robustness of stochastic neural computation and probabilistic inference with regard to failures of neurons and mismatch of physical neuron models. Biological data suggest that this is also the way how random variables

are represented in biological networks of neurons. In particular tasks that require computations on dynamic inputs, i.e., high-dimensional time series, will be addressed.

- Development of methods for hierarchical neural sampling in networks of spiking neurons

The goal is here to enable higher layers of the network to guide and speed up sampling from salient parts of the underlying joint distribution, both for static and dynamic network inputs. Hierarchical network topologies further appear to play an important role in the robustness of information processing in spiking networks embedded in imperfect physical substrates, be they biological or neuromorphic.

- Investigation of the functional role of dendrites in stochastic neural computations

Experimental data from neuroscience suggest that top-down inputs arrive in the brain to a large extent in the dendritic tuft of pyramidal cells on layers 2/3 and 5/6, and that they bias these cells to burst if they also receive inputs from other pathways. We will study the viability of this hypothesis in particular in the context of models for hierarchical stochastic computations..

- Development of a neural network implementation of probabilistic programs

Substantial work in cognitive science and machine learning (especially in the group of Josh Tenenbaum at MIT) has focused on probabilistic programs as putative representations for higher level concepts and behaviors in the human brain. But an implementation in neural networks is missing. Besides possible implementations via neural sampling we will also explore possible combinations of neural sampling and other potential neural network mechanisms for probabilistic inference.

- ~~Investigation of the computational role of criticality in dynamical systems~~

~~Numerous experimental studies have shown that brain network are typically in a critical regime, based on the analysis of the statistics of so-called avalanches of activity and irregularity of spike trains. However, there is a lack of insight into how these or other features of criticality could support network computation or learning in neural networks or others types of dynamical system. Additionally, criticality appears to depend strongly on the exact temporal shape of synaptic interactions – a feature that is often not accounted for in abstract models of spiking networks. We will characterize the dependence of critical behavior (phase transitions, critical exponents) on the exact shape of neuro-synaptic dynamics and investigate its role in computation.~~

- ~~Investigation of possible links between concepts from quantum computation and neural computation.~~

~~A number of recent experimental data from neuroscience suggest that brain networks are typically in one of a discrete set of states (becoming visible through the activation of one of a discrete set of assemblies of neurons in the network). One also has observed in the brain a different type of network activity when it has to search for a new behavioural rule or strategy, that could be related to complex superpositions of discrete stereotypical patterns (assemblies). We will explore to what extent concepts from quantum computing may help to understand these phenomena, and whether this link provides new ideas for computation in neuromorphic systems or brains.~~

Task leader: Mihai Petrovici

Direct costs:
TUGraz: ~~20k Euro for computing~~

Personnel:

- UHEI: ~~2.5 FTE~~ 0,3 FTE
- UMAN: ~~0.5 FTE~~ 0 FTE
- TUGraz: ~~2.5 FTE~~ 0,85 FTE

Task 9.4.2 Learning – Task leader: Robert Legenstein

Concrete points to be investigated:

- Development of a sequence of challenging benchmark learning tasks that can drive learning and emergence of complex computational and cognitive functions in large neuromorphic systems

In order to leverage the capability of large neuromorphic systems to engage hierarchical networks we will include learning tasks that require the formation of internal representations simultaneously on several spatial and temporal scales, and on several levels of abstraction. We will set up -- in collaboration with the international research community --a repository of such tasks that are especially suited for training complex neuromorphic systems or other types of functional brain models.

- ~~Online learning~~

~~Many learning algorithms, especially those that are applied in deep learning, only work well in a learning mode where the training data is drawn from a static data distribution that describes a single task. We want to enable neuromorphic systems to learn like brains in an online manner, where new examples, tasks, and challenges can be efficiently integrated into previously learnt knowledge.~~

- ~~Develop methods for simultaneous learning on several levels of abstraction in hierarchical neural systems, where abstract representations support fast learning of new tasks.~~

~~Methods that could achieve this have already been proposed on abstract levels in machine learning and cognitive science. We will develop neural network implementations of such methods.~~

- Develop Learning-to-Learn methods that enable networks of spiking neurons to learn many tasks from few examples

Learning-to Learn has become a powerful new tool in Machine Learning. We will modify and apply it in order to produce networks of spiking neurons that are able to learn a variety of tasks very efficiently, because they can encode and transfer knowledge from previously learnt tasks. In addition, this method promises to provide new insight into the functional role of specific features of neurons, dendrites, synapses, synaptic plasticity, and network architectures in the brain.

- Develop Learning-to-Learn methods that enable neural circuits of the Brainscales hardware to carry out a variety of computational and learning tasks in a reliable and efficient manner, in spite of mismatch and other noise sources in the hardware

We apply here the new Learning-to-Learn technology to networks of neurons in the Brainscales hardware. One goal is to replace work-intense calibration of hardware neurons by an automated procedure, that directly optimizes their parameters to optimally support a variety of network functions. But this work will also provide new insight into the distribution of regions in high-dimensional parameter spaces that support stable and efficient function of the Brainscales hardware, and other novel energy-efficient computing hardware. It will capitalize on this insight for the design of algorithms that efficiently search for these parameter regions.

- Develop methods for deep learning that can be implemented in neuromorphic hardware

A number of new deep learning methods have recently been developed that support supervised learning in energy-efficient neuromoprhic hardware. We will develop variations of deep learning that are especially suited for the neuromorphic hardware in SP9, including new methods for unsupervised learning and reward-based learning in deep neural networks. We will also explore biologically more realistic alternatives to LSTM cells in artifical neural networks, and to backpropagation through time, which underlies most recent successes in emulating cognitive computation in artificial neural networks.

- Develop methods for deep learning that exploit the abundance of lateral connections on each level of a brain hierarchy.

We will explore in particular the hypothesis that lateral connections within an area of a hierarchical brain network (e.g., the visual system) support the comparison of bottom-up and top-down information with stereotypical patterns that have been learnt and are represented in that area.

- Long-term learning

Like Learning-to-Learn, this research will exploit that learning processes are sped up – in comparison with biological time – by a factor 10 000 in the Brainscales NM hardware. Hence year-long learning processes can be simulated in about an hour. This provides an ideal research tool for studying longterm learning in NM models of neural circuits and brain areas.

Task Leader: Robert Legenstein

Direct costs:
TUGraz: ~~25k Euro for computing~~

Personnel:
- UHEI: ~~2 FTE~~ 0,3 FTE
- UMAN: ~~1 FTE~~ 0,5 FTE
- TUGraz: ~~3 FTE~~ 1,5 FTE

Task 9.4.3 Architectures for neuromorphic intelligent systems – Task leader: Wolfgang Maass

Concrete research goals:

- Definition of a sequence of challenging benchmark tasks that allow us to test the levels of cognitive capabilities that neuromorphic systems have achieved.

Several task sets have already been proposed in the literature, e.g. tasks that were used to test functional capabilities of the SPAUN model from the Lab of Eliasmith. Other benchmark tasks have recently been proposed to test cognitive capabilites of digital computers, such as those published by Facebook (bAbI tasks, MazeBase tasks). A suitable sequence of tasks for challenging and testing neuromorphic artificial brains will be developed in collaboration with the international research community.

- Develop methods that enable networks of spiking neurons to create variables, bind variables to concrete meanings, and support abstract reasoning and rule application with the help of this level of abstraction.

This work will build on ideas and paradigms that are currently developed at TUGraz. One overriding goal is to develop methods that are based on recent experimental data about corresponding processes in the brain.

- ~~Develop methods for carrying out cognitive tasks in networks of spiking neurons without copying of information~~

~~Many abstract algorithms for solving cognitive tasks involve copying of information from one part to another part of a larger network. The brain appears to be able to carry out these tasks without copying of information. In fact, since information is largely encoded in the brain by the set of neurons that currently fire, and it would be very difficult to copy information that is encoded in this way. The development of new algorithms that do not require copying of information is of particular interest also from the general perspective of designing energy-efficient new computer hardware, since a large fraction of the energy that is spent in current computing hardware is spent for copying of information. One such new algorithm has been proposed in (Legenstein and Maass, 2016).~~

- Develop a model for human memory and associations between memory items in networks of spiking neurons, based on experimental data from humans.

Electrode recordings from epileptic human patients have produced a wealth of insights and hints about neural codes for higher level concepts and memories in the medial temporal lobe, and how these neural codes change when an association between two concepts or memory items is created. Corresponding small scale models are currently developed at TUGraz, and will be scaled up to almost real-size models on the hardware developed in WP 9.3. This will result in a new paradigm for memory in neural networks that integrates more recent brain data, and which also allows us to study the formation of the web of associations between memory items that is essential for the function of the human brain. Large-scale implementations in neuromorphic hardware will allow us to also investigate scaling properties of the resulting new model for memory in neural networks.

- ~~Develop a neuromorphic model for the formation and functional use of episodic memory in networks of spiking neurons, based on experimental data from the human brain.~~

~~Small scale models, based on recent experimental data, are currently developed at TUGraz. We will scale up this model to neuromorphic hardware with at least a million neurons. Functional tasks will require to make use of this memory network (including episodic memory) for solving tasks. Experimental data and preliminary feasibility~~

~~studies at TUGraz suggest that these capabilities can be attained through suitable network architectures in combination with learning.~~

- ~~Develop software that enables users to design and carry out higher level cognitive operations on neuromorphic hardware without having to worry about details of the implementation~~

~~This software will be developed in collaboration with WP 9.1. It will support the dissemination and use of results from SP9.~~

Task Leader: Wolfgang Maass
Direct costs:
TUGraz: ~~25k Euro for computing~~
Personnel:

- CNRS: ~~1 FTE~~ 0 FTE
- UMAN: ~~1.5 FTE~~ 0,5 FTE
- UHEI: ~~0.5 FTE~~ 0,2 FTE
- TUGraz: ~~3 FTE~~ 1.5 FTE

SP9.4 in kind contributions

- UHEI: 2.5 FTE
- TUGraz: 50% of the EC paid FTE contributed on top as in-kind contribution

Workpackage 9.5: Applications and Benchmarks – WP leader: Michael Schmuker

Applications in WP 9.5 are thought of as "demonstrators" that display the use of the neuromorphic platforms to obtain new scientific insight, and exhibit novel technological solutions for future computing and robotics. WP9.5 will also provide a framework to facilitate the development of larger neuromorphic software systems within the subproject, the overall HBP, and beyond.

The first task will demonstrate the capability of the SpiNNaker platform to interface in real time with sensors and actuators. Task 2 will display the proficiency of spiking networks to exploit temporal stimulus features for object recognition, and makes use of the high acceleration of NM-PM-1 for "lifetime learning". Task 3 gives a demonstration of cognitive agents implemented on spiking hardware platforms.

Finally, in task 4 we provide Benchmarks that evaluate the performance of the hardware systems, including the software interface that is exposed to external users. Benchmarks thus provide a quantifiable indicator of progress in making the platforms available to the scientific community and commercial users.

Task 9.5.1: Neuronal signal processing – Task leader: Michael Schmuker

In this task, we will develop a demonstrator for signal processing and pattern recognition with neuromorphic hardware for application in mobile robots.
The sensory systems of animals are highly efficient in processing and classifying external stimuli. Brain-inspired methods for processing visual and olfactory information, and for pattern recognition have already been demonstrated during the ramp-up phase within HBP, on neuromorphic hardware. The aim of this task is to use this technology to endow mobile robots with decision making capabilities that draw information on real-time sensor input, which is processed using neuromorphic hardware.

We will investigate processing visual and olfactory information using small, portable neuromorphic systems, to make navigational decisions. These decisions will be made using neuromorphic pattern recognition algorithms based on visual and olfactory input, provided by a camera and an electronic nose. A particularly novel aspect is the use of spatio-temporal dynamics in gas concentration to estimate gas source proximity, which has been explored by PI Schmuker in a recent publication (Sens Act B: Chemical, 2016). If successful, this application will therefore not only demonstrate the fitness of the HBP's neuromorphic hardware for practical use in robotics, but also explore a novel approach in the yet heavily underexplored domain of gas-based robot navigation.

Naturally, this project will benefit from cross-SP collaboration with the Neurorobotics platform, since the simulation environment provided therein will greatly facilitate rapid prototyping and development of an neuromorphic application that interacts with the real world.

Resources:

- HERTS: ~~1,2 FTE~~ 1 FTE (research assistant)
    - Develop a signal processing pipeline on neuromorphic hardware that can implement various temporal filters.
    - Test signal processing capabilities using real sensor data.
    - Quantify benefits of neuromorphic signal processing for neuromorphic pattern recognition.
    - Evaluate signal processing performance and compare to conventional approaches.
    - Active contribution to supporting platform users via mailing list.

Direct costs: 0 Euro

Task 9.5.2: Agent and Agent Components – Task leader: Chris Huyck

This task has two primary goals: firstly, to create agents in simulated neurons; secondly, to provide components of these agents that can be reused.

Two new agents will be created, the first a virtual agent, and the second a physical robotic agent. These agents will make use of existing virtual agents that run on SpiNNaker/NM-MC-1, HiCANN/NM-PM-1, and Nest. The existing agents work in virtual environments and have vision, natural language processing (NLP), and planning components.

An agent in a virtual environment will add an associative memory component. This will enable the agent to learn relations between objects in the environment and will enable the agent to improve its performance. A cognitive model will also be developed.

Extra motor behaviour will be needed for a robotic agent. Existing spiking neuron based (bioloid) robotic work will be extended, and integrated with existing virtual agent components. If time allows, the associative memory component will also be added. This may be extended to other robot platforms.

In collaboration with WP 9.1 independent components for vision, NLP, planning, motor action, and associative memory will be developed. These will be combined to form the virtual and robotic agents. Different versions of the components will be evaluated within these agents. The agents will also provide an API for integration for other components (such as spatial cognitive mapping and olfaction). These tasks will obviously involve the neuro-robotics subproject. We will use their virtual environment, and hope to support their physical robotic work.

Moreover, we will collaborate with the systems and cognitive science subproject; one of their neural systems will be integrated into an agent. Learning across components will be considered in collaboration with the Theory subproject.

- Developing agents and agent components, and a cognitive model for learning relations in an environment.
- Active contribution to supporting platform users via mailing list.

Resources:

- MU: ~~1.2 FTE~~ 1 FTE

Task 9.5.3: Object recognition -- Task leader: Thomas Nowotny

It has been established in animals that stimuli are interpreted, and objects recognised, based on stimulus onset synchrony or temporal correlations in stimulus occurrence. For example, in vision, dots that move in a correlated fashion on a background of randomly moving dots are interpreted as an object. In hearing, mixtures of pure tones are conceptualised as a single sound, if they have a synchronous onset, but as separate sounds if not. In olfaction, mixtures of odours that co-fluctuate are distinguished from those where onsets of stimuli are asynchronous.

In this task we will investigate neuromorphic solutions to object recognition based on neural networks from neuroscience models of olfaction and vision. In addition to the temporal binding of stimuli to objects described above, this work will also investigate the formation of stimulus primitives through "lifetime learning", i.e. the ability to identify particular complex stimulus mixtures as one elemental stimulus, e.g. the smell of coffee, or a complex Chinese character, based on familiarity acquired from a large number of such inputs and over a long time period. The 10,000X speedup of NM-PM will be of particular value here.

The work will combine the refinement of networks (neuromorphic algorithms) on standard PCs and GPUs with implementing and investigating them on the hardware systems. Some basic benchmark measurements will be taken as in our earlier work in SP11.3.6 during the HBP ramp-up phase.

The Task will be complemented by related work in the "Odor Objects" project (biological basis of the neuromorphic algorithms investigated here from honeybees, HFSP 2015-2017) and the "Brains on-Board" project (bee-inspired neuromorphic algorithms on GPUs for use in autonomous drones, EPSRC 2017-2022).

Staff costs:
UoS: ~~1.2 FTE (1.0 FTE for researcher and 0.2 FTE for PI)~~ 1 FTE

- Develop a neuromorphic solution to object recognition that exploits temporal correlation of sensory inputs.
- Investigate the formation of elemental stimulus prototypes in very long simulations, exploiting the speedup factor of the BrainScaleS hardware.
- Quantify network performance and compare to GPU implementations.
- Active contribution to supporting platform users via mailing list.

In-kind contributions:
UoS: Direct costs:
CPU time on Sussex HPC: 1152 EUR
GPU time on Sussex HPC GPU server: 500 EUR

Task. 9.5.4: Benchmarking – **Task leader:** Ulrich Rückert

Benchmarking provides a method for comparing the performance of systems or subsystems with respect to real world tasks or synthetic programs. Benchmarks are

designed to mimic a particular type of workload on a hardware or software component of a system. Synthetic benchmarks do this by specially created programs that impose a designed artificial workload on the component. Functional benchmarks run application programs on the system for performance evaluation. While functional benchmarks usually give a much better measure of application performance on a given system, synthetic benchmarks are useful for characterizing the role of individual components within the system. Benchmarks help system designers to select the optimal system components (hardware as well as software) and to understand the performance and energy characteristics of their systems. Benchmarking is not easy and often involves several iterative design-and-execution cycles in order to arrive at predictable, useful conclusions. Furthermore, interpretation of benchmarking data is also extraordinarily difficult.

In this task, we aim at developing benchmark suits as well as tools for analysing the benchmark results in an efficient and unbiased manner. Two benchmark sets have to be developed and maintained: First, a representative set of neural network architectures studied in HBP; b) a synthetic set of specialized networks for evaluating bottlenecks within the neuromorphic computational architectures of SP9 (NM-PM, NM-MC).
In order to identify a representative set of network architectures in HBP to extend the functional benchmark set implemented during SGA1, we will adapt and integrate the neuromorphic architectures from tasks 9.5.1, 9.5.2, and 9.5.3 and WP 9.4. In addition, we will adapt selected neural architectures from SP4 and SP10 to the neuromorphic hardware in order to show the benefits of these hardware platforms. This also underlines the important role of benchmarking in stimulating cross-SP collaboration in HBP.

The second important benchmarking aspect is the search for architectural bottlenecks in the neuromorphic systems, in particular with respect to large-scale experiments involving very high neuron counts and long biological run times. Such experiments are a hallmark of the hardware systems developed in SP 9, and the identification of bottlenecks and architectural weaknesses is very important for the design of future versions of the hardware platforms and for optimizing the software framework. Hence, each HW platform requires an own specific benchmark set which has to be specified in close cooperation with the system architects of both platforms (WP 9.2, WP 9.3).

Last but not least, tools for analysing the benchmark runs will be developed. We aim at a software package for building, running and comparing large-scale benchmarks on neuromorphic hardware. The software package will be divided in a frontend for building platform-independent large-scale benchmark runs, and a backend for running the benchmark on a selected neuromorphic hardware. Comparisons will be done automatically by an analysis tool based on a flexible set of performance measures for multi-objective optimization.

Staff costs:
UNIBI: 1 FTE

- Develop a representative benchmark suite for the hardware systems in SP9.
- Cross-SP collaboration to increase diversity of benchmark applications.
- Develop analysis tools for benchmarks.
- Develop a frontend for running the benchmarks.
- Active contribution to supporting platform users via mailing list.

Direct costs:
UNIBI: None.

Task 9.5.5 (Placeholder task for a potential open-call work on machine learning)

WP9.5 in kind contributions

- MU
  - 0.7 FTE (0.2 FTE for PI and 0.5 FTE for PhD student)
  - Direct costs: MDX: 10,000 EUR for 1 SpiNNaker board.
- UNIBI: 1 FTE
- UoS: 0.2 FTE + CPU+GPU time on Sussex HBP / HPC GPU: 1600 Euro

Workpackage 9.6: Coordination and Infrastructure – WP leader: Karlheinz Meier

Workpackage 9.6 ensures the SP management, outreach and communication as well as training and user support of the neuromorphic platform. It is in charge of the integration of SP9 as part of the HBP infrastructure and enables the interaction with other SPs and the CDPs. The workpackage has 4 tasks.

The WP is led by Karlheinz Meier

Task 9.6.1 Management – Task leader: Björn Kindler

Management is led by Björn Kindler. It carries out the following tasks :

- Organize monthly subproject video conferences
- Organize 4 face-to-face subproject meeting per year
- Organize 4 project internal workshops led by SP9 (example : Fürberg workshop)
- Organize joint meetings with other SPs and CDPs
- Maintain and further develop SP9 webspace
- Ensure proper reporting of milestones and deliverables
- Prepare SP9 contribution to project reviews
- Provide input to requests to the ethics SP12

Task 9.6.2 Outreach and communication – Task leader: Björn Kindler

Outreach and communication is led by Björn Kindler. It carries out the following tasks

- Maintain and further develop SP9 webspace
- Present SP9 in social media
- Produce media material
- Organize and carry out public SP9 presentations
- Coordinate SP9 conference contributions
- Maintain and develop SP9 support for the NICE conference series
- Collect and database SP9 publications
- Maintain database of SP9 students
- Provide hardware and software for school and universities
- Maintain and establish industry relations
- Organize 2 application workshops per year
- Maintain and further develop innovation roadmap of SP9

Task 9.6.3 NM platform training and user support (PLAN/MAYBE: move the education part to SP13)– Task leader: David Lester/Sebastian Schmitt

Training and user support is led by David Lester. It carries out the following tasks

- Provide support of platform users together with WPs 9.1.1, 9.1.2 and 9.1.3
- Organize 2 platform training workshops per year
- Organize 2 Codejams per year

Maintain and develop support for Capo Caccia Workshop

Task 9.6.4 Integration and Infrastructure – Task leader: Karlheinz Meier

Integration and infrastructure is led by Karlheinz Meier. It carries out the following tasks

- Integrate the NM platform into the HBP infrastructure
- Enable interaction with other SPs
- Enable interaction with the CDPs

Use of resources in WP 9.6

- UHEI 1.5 FTE
- Each partner 0,025 FTE for management.
- No EC paid PM planned here for training (in kind contributions)

~~100 k€ for workshop and training support (room and equipment rent, travel for external speakers, meeting materials)~~

Direct cost: Total of 20.000 Euro supporting training and SP9 meetings

WP 9.6 in kind contributions

- UHEI: 1FTE