

D9.3.1: PyNN 0.9 on SpiNNaker

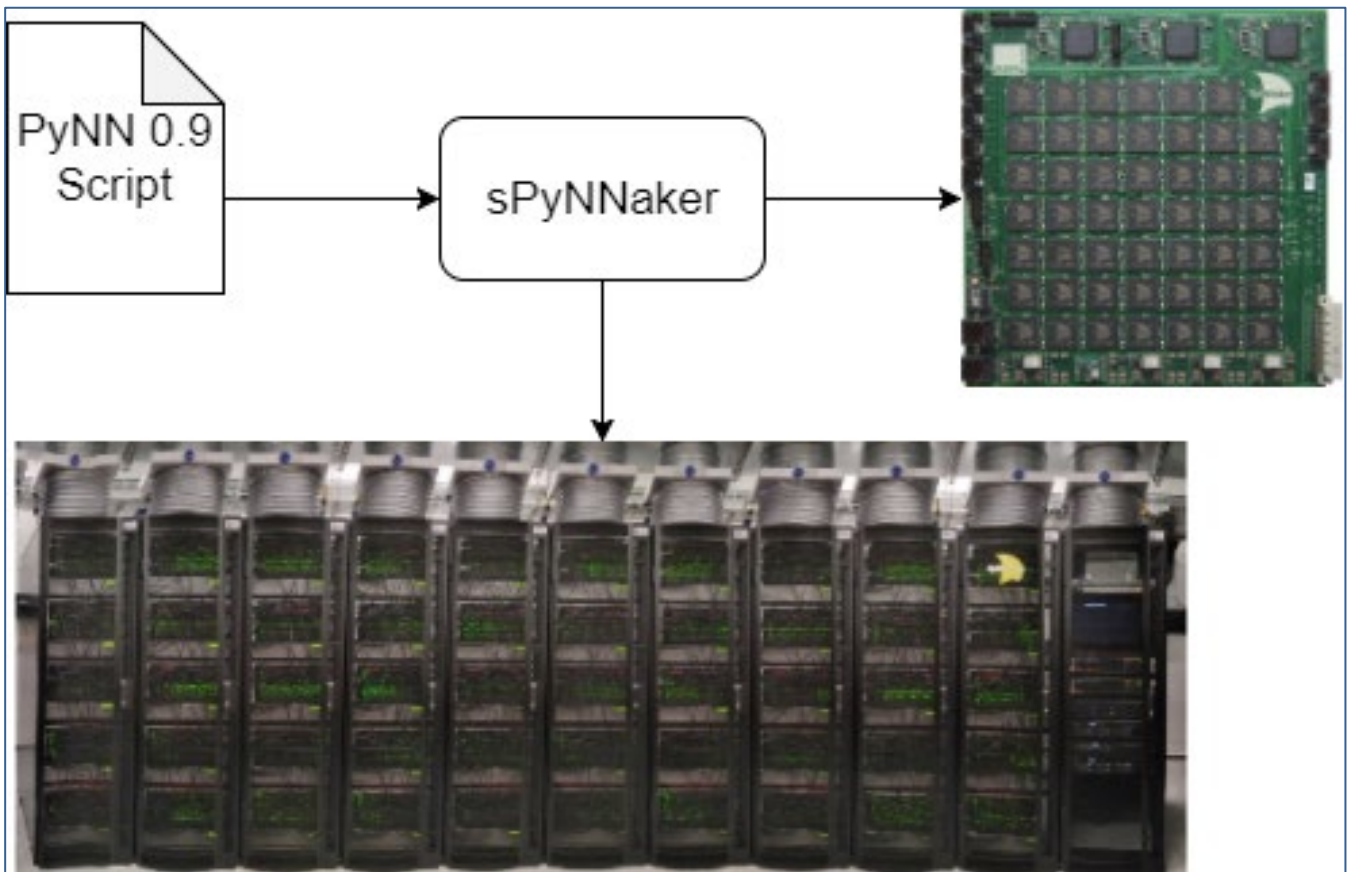


Figure 1: A PyNN 0.9 Script running on SpiNNaker.

The script is executed by the sPyNNaker software, which translates the script into data for executables to be run on the SpiNNaker machine.

Project Number:	785907	Project Title:	Human Brain Project SGA2
Document Title:	D9.3.1: PyNN 0.9 on SpiNNaker		
Document Filename:	D9.3.1 (D59.1 D92) SGA2 M12 ACCEPTED 190723		
Deliverable Number:	SGA2 D9.3.1 (D59.1, D92)		
Deliverable Type:	Other		
Work Packages:	WP9.3		
Dissemination Level:	PU = Public		
Planned Delivery Date:	SGA2 M12 / 31 March 2019		
Actual Delivery Date:	SGA2 M12 / 1 March 2019; ACCEPTED 23 Jul 2019		
Author(s):	Andrew ROWLEY, UMAN (P63)		
Compiled by:	Andrew ROWLEY, UMAN (P63)		
Contributor(s):	Andrew ROWLEY, UMAN (P63)		
SciTechCoord Review:	Marc MORGAN, EPFL (P1)		
Editorial Review:	Guy WILLIS, EPFL (P1)		
Description in GA:	As complete a PyNN 0.9 implementation on SpiNNaker as possible (Task T9.3.3)		
Abstract:	<p>PyNN is the language used to describe spiking neural networks (SNNs) of point neurons which can then be simulated on various SNN simulators, including neuromorphic hardware. PyNN version 0.9 has been available since April 2017, but the SpiNNaker implementation requires special handling of some features, and so tends not to be released on the same schedule. The SpiNNaker backend has now been updated to comply with PyNN 0.9 as much as is possible, with the exception of a few less-used features that are more difficult to implement on neuromorphic hardware.</p>		
Keywords:	PyNN, SpiNNaker, Neuromorphic Computing		
Target Users/Readers:	Computational Neuroscientists, Neuromorphic Computing Community, Platform Users		

Table of Contents

1. Introduction	4
2. PyNN 0.9 Implementation, as of 27 Feb 2019	4
3. Testing.....	5
4. Documentation	5
5. Conclusion	6

1. Introduction

PyNN¹ is a Python-based API used to describe and execute simulations of spiking neural networks. The API allows users to implement their network in a platform-agnostic manner, and then execute the same network on many platforms. This ability relies on the backend platforms implementing the same version of the API, as well as all the features of the API; PyNN does not require the latter in its entirety, with the only requirement being that at least two platforms implement any given feature for it to appear in the PyNN API.

A PyNN backend is available for the SpiNNaker platform in the form of the sPyNNaker software and associated libraries. This implementation is done at the API level as shown in Figure 1, meaning that there is little tie-in with the actual PyNN library itself. This is because some of the implementation details of the PyNN API can be implemented more efficiently, both in time and space, when it is known that SpiNNaker is the target platform. The sPyNNaker implementation does not fully implement all the features of the PyNN API, as some features are more difficult to implement on SpiNNaker, such as the full Hodgkin-Huxley model.

PyNN version 0.9 has been available since April 2017 but, until now, sPyNNaker had only implemented versions 0.7 and 0.8.

2. PyNN 0.9 Implementation, as of 27 Feb 2019

The software has now been updated to support version 0.9 of the PyNN API (PyNN 0.9.3, which is the latest release as of the date of this document). Although not all features of the API have been implemented, the most commonly used and important features have.

As of Feb 2019, the features that are missing include:

- PyNN Assembly objects². These are groupings of Population, PopulationView and other Assembly objects into a single item with a common index. We have not had any users requesting these objects to date.
- Projection³ between PopulationView objects. A PopulationView contains a subset of the neurons in a Population. Projections between PopulationView objects would allow the user to better specify connectivity between these subsets. However, it is possible for users to specify this connectivity using a FromListConnector, which is supported, or else by separating the Population objects into several sub-populations in advance, where this makes sense. Although there has been some interest in this feature, users have so far managed without it.
- TsodyksMarkramSynapse⁴. This is used to specify short-term plasticity. An initial implementation of this is available but this still requires testing.
- Various neuron models. Some of the neuron models are hard to implement within the limited code space of the SpiNNaker cores, and without floating point support. These include Integrate and Fire models with adaptation (with the exception of the Izhikevich model, which is implemented), and the Hodgkin-Huxley model.
- SpikeSourceInhGamma⁵. We have so far had no requests for this to be implemented.
- Current Sources⁶. All the neuron models support the use of `i_offset`, which is the equivalent of a DCSource. This can also be varied through the ability to call `run` a number of times.

¹ <http://neuralensemble.org/docs/PyNN/>

² <http://neuralensemble.org/docs/PyNN/reference/populations.html#assemblies>

³ <http://neuralensemble.org/docs/PyNN/reference/projections.html#projections>

⁴ <http://neuralensemble.org/docs/PyNN/reference/plasticitymodels.html#short-term-plasticity-mechanisms>

⁵ <http://neuralensemble.org/docs/PyNN/reference/neuronmodels.html#spike-sources>

⁶ <http://neuralensemble.org/docs/PyNN/reference/electrodes.html#current-sources>

Please note that, in spite of these features not being implemented, we nevertheless consider the implementation to date to be complete with respect to the API.

3. Testing

Passed - 173

> test_get_gsyn - p8_integration_tests.test_0_1_time_steps.test_gsyn_allow_violate.test_synfire_Odot1_timestep_test_get_gsyn.TestGsyn	42s
> test_get_gsyn - p8_integration_tests.test_0_1_time_steps.test_gsyn_allow_violate.test_synfire_Odot1_timestep_test_print_gsyn.TestPrintGsyn	49s
> test_get_gsyn - p8_integration_tests.test_0_1_time_steps.test_gsyn_do_not_allow_violate.test_synfire_Odot1_timestep_test_get_gsyn.TestGsyn	<1s
> test_script - p8_integration_tests.test_0_1_time_steps.test_malloc_key_allocator.test_synfire_if_curr_exp_with_malloc_key_allocator.TestMallocKeyAllocatorWithSynfire	44s
> test_script - p8_integration_tests.test_0_1_time_steps.test_overload_system_to_check_recovery.test_read_spike_file_Odot1_time_step.TestReadingSpikeArrayDataAndBigSlices	58s
> test_get_spikes - p8_integration_tests.test_0_1_time_steps.test_spikes.test_synfire_Odot1_timestep_test_get_spikes.TestGetSpikesAt0_1msTimeStep	46s
> test_print_spikes - p8_integration_tests.test_0_1_time_steps.test_spikes.test_synfire_Odot1_timestep_test_print_spikes.TestPrintSpikes	42s
> test_get_voltage - p8_integration_tests.test_0_1_time_steps.test_v.test_synfire_Odot1_timestep_test_get_v.TestGetVoltage	46s
> test_print_voltage - p8_integration_tests.test_0_1_time_steps.test_v.test_synfire_Odot1_timestep_test_print_v.TestPrintVoltage	1m 31s
> test_va_benchmark - p8_integration_tests.test_0_1_time_steps.test_va_bench_mark_tests.test_va_benchmark.TestVABenchmarkSpikes	54s
> test_get_gsyn - p8_integration_tests.test_1_0_time_steps.test_gsyn.test_synfire_1dot0_timestep_test_get_gsyn.TestGetGsyn	41s
> test_get_gsyn - p8_integration_tests.test_1_0_time_steps.test_gsyn.test_synfire_1dot0_timestep_test_print_gsyn.TestPrintGsyn	42s
> test_end_before_print - p8_integration_tests.test_1_0_time_steps.test_malloc_key_allocator.test_synfire_if_curr_exp_with_malloc_key_allocator.TestMallocKeyAllocatorWithSynfire	42s
> test_script - p8_integration_tests.test_1_0_time_steps.test_malloc_key_allocator.test_synfire_if_curr_exp_with_malloc_key_allocator.TestMallocKeyAllocatorWithSynfire	42s
> test_single_neuron - p8_integration_tests.test_1_0_time_steps.test_single_neuron_tests.test_single_if_curr_exp_neuron.TestIfCurrExpSingleNeuron	48s
> test_script - p8_integration_tests.test_1_0_time_steps.test_spike_array_from_read_in_spikes_and_big_slices.test_read_spike_file_1dot0_time_step.TestReadingSpikeArrayDataAndBigSlices	51s
> test_get_spikes - p8_integration_tests.test_1_0_time_steps.test_spikes.test_synfire_1dot0_timestep_test_get_spikes.TestGetSpikesAt0_1msTimeStep	42s
> test_print_voltage - p8_integration_tests.test_1_0_time_steps.test_spikes.test_synfire_1dot0_timestep_test_print_spikes.TestPrintVoltage	47s
> test_print_voltage - p8_integration_tests.test_1_0_time_steps.test_v.test_synfire_1dot0_timestep_test_get_v.TestPrintVoltage	47s
> test_print_voltage - p8_integration_tests.test_1_0_time_steps.test_v.test_synfire_1dot0_timestep_test_print_v.TestPrintVoltage	47s
> test_va_benchmark - p8_integration_tests.test_1_0_time_steps.test_va_bench_mark_tests.test_va_benchmark.TestVABenchmarkSpikes	43s
> test_run - p8_integration_tests.test_array_connector.test_array_connector.ArrayConnectorTest	49s
> test_more_runs - p8_integration_tests.test_auto_pause_and_resume_tests.test_shorter_run.test_synfire_1_very_low_sdram.TestVeryLow	49s
> test_too_low - p8_integration_tests.test_auto_pause_and_resume_tests.test_too_little_memory.test_synfire_1_too_low_sdram.TestTooLow	19s
> test_run - p8_integration_tests.test_backwards_compatibility_issues.test_ssa.SynfireIfCurrExp	<1s
> test_run - p8_integration_tests.test_buffer_manager.test_if_curr_exp_with_all_recording.test_synfire_if_curr_exp.SynfireIfCurrExp	50s
> test_run - p8_integration_tests.test_buffer_manager.test_if_curr_exp_with_no_recording.test_synfire_if_curr_exp.SynfireIfCurrExp	48s
> test_run - p8_integration_tests.test_buffer_manager.test_if_curr_exp_with_v_recording.test_synfire_if_curr_exp.SynfireIfCurrExp	49s
> test_run - p8_integration_tests.test_buffer_manager.test_poisson_with_recording.test_pynnBrunnelBrianNestSpinnaker.PynnBrunnelBrianNestSpinnaker	49s
> test_run - p8_integration_tests.test_buffer_manager.test_poisson_without_recording.test_pynnBrunnelBrianNestSpinnaker.PynnBrunnelBrianNestSpinnaker	48s
> test_no_split - p8_integration_tests.test_change_neuron_parameters_between_runs.test_change_parameter.TestChangeParameter	52s
> test_split - p8_integration_tests.test_change_neuron_parameters_between_runs.test_change_parameter.TestChangeParameter	53s
> test_no_split - p8_integration_tests.test_change_neuron_parameters_between_runs.test_change_spike_source_poisson.TestChangeParameter	47s
> test_one_core - p8_integration_tests.test_change_neuron_parameters_between_runs.test_set_toffset.TestSetOffset	45s
> test_three_cores - p8_integration_tests.test_change_neuron_parameters_between_runs.test_set_toffset.TestSetOffset	45s
> test_one_core - p8_integration_tests.test_change_neuron_parameters_between_runs.test_set_toffset_izk.TestSetOffset	45s
> test_three_cores - p8_integration_tests.test_change_neuron_parameters_between_runs.test_set_toffset_izk.TestSetOffset	45s
> test_get_gsyn - p8_integration_tests.test_change_neuron_parameters_between_runs.test_synfire_if_curr_exp_parameter.TestGetGsyn	59s
> test_synfire_poisson_if_curr_exp_parameter - p8_integration_tests.test_change_neuron_parameters_between_runs.test_synfire_poisson_if_curr_exp_parameter.TestSynfirePoissonIfCurrExpParameter	1m 4s
> test_synfire_poisson_if_curr_exp_parameter_test_second_none - p8_integration_tests.test_change_neuron_parameters_between_runs.test_synfire_poisson_if_curr_exp_parameter_test_second_none.TestSynfirePoissonIfCurrExpParameterTestSecondNone	0s
> test_run - p8_integration_tests.test_csa_connectors.test_csa_connectors.CSAConnectorTest	49s

Figure 2: Output of Jenkins, the integration testing environment used to test the PyNN API compatibility.

A combination of unit tests and integration tests are set up to be performed on the software using the Jenkins CI continuous integration suite. These tests were used to verify the compatibility with the PyNN API, and continue to be executed when changes are made to ensure that those changes do not break the software or the compatibility with the PyNN API. The tests performed can be seen at https://github.com/SpiNNakerManchester/sPyNNaker8/tree/master/p8_integration_tests. An example of the output of these tests run on the Jenkins system is shown in Figure 2.

4. Documentation

The documentation for this software consists of many elements, including:



- The HBP neuromorphic guidebook⁷. This is online documentation that evolves with the software, and includes links to the PyNN API and example code.
- The sPyNNaker code-level documentation⁸. This details the modules available in the software, which includes PyNN extensions.
- The SpiNNaker software pages⁹. This includes installation guides and general usage instructions and limitations.
- The SpiNNaker Workshop slides and tutorials¹⁰. This includes training material with slides and lab manuals to help users learn to use PyNN with SpiNNaker.

5. Conclusion

The PyNN 0.9 API has been implemented on SpiNNaker. There are a few limitations to this interface, but they can be addressed through the ongoing maintenance of the software, as the need arises. The testing framework and documentation have been updated to match the implementation.

The changes made will make it easier for users to use SpiNNaker. With these changes, users now have the ability to reduce the data extracted from the machine through the implementation of the PopulationView object which enables selective recording, which will make the execution of scripts faster. The fact that the latest API has been implemented means that users can execute the same scripts on the SpiNNaker platform as they can on other platforms that support the API, with only minor changes required should their scripts happen to rely on the unimplemented features.

⁷ <https://collab.humanbrainproject.eu/#/collab/51/nav/1069>

⁸ <https://spynaker8.readthedocs.io/en/latest/>

⁹ <http://spinnakermanchester.github.io/>

¹⁰ <http://spinnakermanchester.github.io/workshops/eighth.html>