Sub realtime simulation of a full density microcircuit model on a single compute node

Anno Kurth[1,2], Justin Finnerty[1], Dennis Terhorst[1], Jari Pronold[1,2], Johanna Senk[1], Markus Diesmann[1,3,4]

[1] Institute of Neuroscience and Medicine (INM-6) and Institute for Advanced Simulation (IAS-6) and JARA-Institute Brain Structure-Function Relationships (INM-10), Jülich Research Centre, Jülich, Germany
[2] RWTH Aachen University, Aachen, Germany
[3] Department of Psychiatry, Psychotherapy and Psychosomatics, School of Medicine, RWTH Aachen University, Aachen, Germany
[4] Department of Physics, Faculty 1, RWTH Aachen University, Aachen, Germany

Email: a.kurth@fz-juelich.de

The cortical microcircuit is a building block of the mammalian brain. In a model of the network below a 1 mm^2 patch of cortical surface [1] the spatial structure is replaced by cell-type specific random connectivity. Each layer is represented by an excitatory and an inhibitory population of integrate-and-fire model neurons. The network model is a benchmark for neuromorphic systems [2, 3, 4].

This contribution shows performance data for the microcircuit model on two AMD EPYC Rome 128 core compute nodes coupled by a direct Infiniband interconnect and running NEST 2.14 [5] (with fix 726f9b04bbd47c). On a single node we observe sub realtime performance, on two the simulation is 1.7 times faster than realtime. Our study of the aged 4g kernel serves as a reference for present optimizations, exposes bottlenecks, and guides the design of future computing systems.

For the single node the energy per synaptic event is 0.26 μJ, and for the fastest configuration using two nodes 0.39 μJ. These values are in the same order of magnitude as the lowest reported so far. The findings confirm a non-trivial relationship [2] between the resources in use and the energy required.

At the poster we demonstrate how power measurements with a contemporary PDU can be aligned with benchmark timers to obtain a reliable time course of power consumption.

Acknowledgements

Partially supported by EU Horizon 2020 945539 (HBP SGA3) and Helmholtz IVF SO-092 (ACA).

References

1. Potjans TC & Diesmann M (2014) **The cell-type specific cortical microcircuit: relating structure and activity in a full-scale spiking network model**. *Cerebral Cortex* 24:785–806. doi: 10.1093/cercor/bhs358
2. van Albada SJ, et al. (2018) **Performance comparison of the digital neuromorphic hardware SpiNNaker and the neural network simulation software NEST for a full-scale cortical microcircuit model**. *Front Neurosci 12:291*. doi: 10.3389/fnins.2018.00291
3. Knight JC & Nowotny T (2018) **GPUs outperform current HPC and neuromorphic solutions in terms of speed and energy when simulating a highly-connected cortical model**. *Front Neurosci* 12:941. doi: 10.3389/fnins.2018.00941
4. Rhodes O, et al. (2019) **Real-time cortical simulation on neuromorphic hardware**. *Phil Trans R Soc A* 378:20190160. doi: 10.1098/rsta.2019.0160
5. Peyser A, et al. (2017) **NEST 2.14** *Zenodo*. doi: 10.5281/zenodo.882971