

Simulating macroscale brain circuits with microscale resolution

Abigail Morrison¹, Tobias C Potjans^{2,3}, Susanne Kunkel⁴,
Hans-Ekkehard Plesser⁵, Markus Diesmann^{1,3,4}

1 RIKEN Brain Science Institute, Wako City, Saitama, Japan

2 Brain and Neural Systems Team, RIKEN Computational Science Research Program, Wako City,
Saitama, Japan

3 Institute of Neurosciences and Medicine, Research Center Juelich, Juelich, Germany

4 Bernstein Center for Computational Neuroscience, Albert-Ludwigs University, Freiburg, Germany

5 Norwegian University of Life Sciences, Aas, Norway

E-mail: abigail@brain.riken.jp

Abstract : In order to understand higher brain functions, it is necessary to simulate macroscale circuits involving multiple areas of the brain. Such brain-scale networks represent a challenge to current simulation technology, as design solutions that are appropriate for simulating local cortical networks lead to unreasonable serial memory overhead when the network is scaled up by one or two orders of magnitude. We quantify these limitations and present strategies for overcoming them on computer architectures with up to 10k processors.

1 Linking microscale to macroscale brain networks


Thanks to distributed computing techniques, today it is possible to routinely simulate local cortical networks of around 10^5 neurons with up to 10^9 synapses on clusters and multiple processor shared memory machines. Simulations of this type carried out with NEST [1] scale well up to at least 1000 processors. However, simulations of microscale networks corresponding to approximately 1mm^3 of the cortex are limited in their explanatory power. To understand the functions of the brain, we need to simulate macroscale circuits involving multiple interacting areas [2]. One approach is to develop brain-scale networks in which the individual nodes are realized by microcircuits at the resolution of point neurons and synapses such as the layered local cortical network model we recently developed [3] (see figure A). These

networks will be one or two orders of magnitude larger than the previously studied models, not only in terms of numbers of neurons and synapses but also in terms of computational load.



2 Simulation technology for brain-scale spiking networks

This presents a number of challenges to current simulation technology. Firstly, as the number of processors increases, the memory overhead due to serial data structures eventually dominates the total memory usage and so limits the parallelization. An example of this is the usage of proxies representing remote neurons on each machine and providing an interface to the local neurons. Although such representations of remote neurons require much less memory than local neurons, their proportion of the total neuronal memory usage

approaches one as the number of machines increases. Secondly, the size of the maximal synaptic delay determines the size of the spike buffers each neuron uses to queue incoming spikes [4]. For local networks these buffers are small, as the synaptic delays are in the order of milliseconds. If interacting brain areas are to be considered, the synaptic delays



increase by an order of magnitude, which entails a corresponding increase of memory usage for the spike buffers. We quantify the effects of these memory limitations in our application up to the order of 10k processors and present strategies for addressing these problems.



3 Performance example of NEST on Blue Gene/P supercomputers

Having implemented these strategies for reducing memory consumption, we tested the performance of NEST on a Blue Gene/P supercomputer. Figure B shows the time required to simulate one second of biological time as a function of the number of cores. The blue circles correspond to a network with 10^6 neurons and 10^4 synapses per neuron.

Previously only networks in the order of 10^5 neurons were possible. The simulation shows good scaling up to 8192 processors and still improves when using 16384 processors. An even larger network with 6 million neurons and full connectivity (10^4 synapses per neuron, red squares) exhibits good scaling up to 32768 processors. Figure C shows the scaled speed-up: computing time as a function of the number of cores when simultaneously increasing the problem size. In this case, a horizontal line corresponds to optimal performance. Red squares show the data for 550 neurons per core, blue circles for 1100 neurons per core (all simulations with 10^4 synapses per neuron). Both cases show only a slight increase in the computing time when simulating a network five times the original problem size on an equivalently enlarged computer system.

Future work concentrates on the further reduction of memory consumption.

Acknowledgements

Partially funded by Next-Generation Supercomputer Project of MEXT, Japan, EU Grant 15879 (FACETS), BMBF Grant 01GQ0420 to BCCN Freiburg, and the Helmholtz Alliance on Systems Biology.

References

- [1] Gewaltig M-O, Diesmann M: NEST (neural simulation tool). Scholarpedia 2007, 2(4), 1430
- [2] Sporns O, Tononi G and Koetter R: The Human Connectome: A Structural Description of the Human Brain, PLoS Comput Biol 2005, 1(4):e42
- [3] Potjans TC, Diesmann M: Consistency of in vitro and in vivo connectivity estimates: statistical assessment and application to cortical network modeling. 38th Soc for Neurosci Meeting 2008, 16.1
- [4] Morrison A, Mehring C, Geisel T, Aertsen A, Diesmann M: Advancing the Boundaries of High-Connectivity Network Simulation with Distributed Computing, Neural Comput 2005, 17(8):1776-801