A Configurable Simulation Environment for the Efficient Simulation of Large-Scale Spiking Neural Networks on Graphics Processors

Jayram Moorkanikara Nageswaran, Nikil Dutt, Jeffrey L Krichmar*,
Alex Nicolau, Alex Veidenbaum

Donald Bren School of Information and Computer Science,
University of California, Irvine, CA 92697 USA
* Department of Cognitive Sciences, School of Social Science,
University of California, Irvine, USA 92697

May 4, 2009

Abstract

Neural network simulators that take into account the spiking behavior of neurons are useful for studying brain mechanisms and for various neural engineering applications. Spiking Neural Network (SNN) simulators have been traditionally simulated on large-scale clusters, super-computers, or on dedicated hardware architectures. Alternatively, Compute Unified Device Architecture (CUDA) Graphics Processing Units (GPUs) can provide a low-cost, programmable, and high-performance computing platform for simulation of SNNs. In this paper we demonstrate an efficient, biologically realistic, large-scale SNN simulator that runs on a single GPU. The SNN model includes Izhikevich spiking neurons, detailed models of synaptic plasticity and variable axonal delay. We allow user-defined configuration of the GPU-SNN model by means of a high-level programming interface written in C++ but similar to the PyNN programming interface specification. PyNN is a common programming interface developed by the neuronal simulation community to allow a single script to run on various simulators. The GPU implementation (on NVIDIA GTX-280 with 1GB of memory) is up to 26 times faster than a CPU version for the simulation of 100K neurons with 50 Million synaptic connections, firing at an average rate of 7Hz. For simulation of 10 Million synaptic connections and 100K neurons, the GPU SNN model is only 1.5 times slower than real-time. Further, we present a collection of new techniques related to parallelism extraction, mapping of irregular communication, and network representation for effective simulation of SNNs on GPUs. The fidelity of the simulation results was validated on CPU simulations using firing rate, synaptic weight distribution, and inter-spike interval analysis. Our simulator is publicly available to the modeling community so that researchers will have easy access to large-scale SNN simulations.