Abstract Submitted
for the MAR17 Meeting of
The American Physical Society

Toward Petascale Biologically Plausible Neural Networks LYLE LONG, Pennsylvania State Univ — This talk will describe an approach to achieving petascale neural networks. Artificial intelligence has been oversold for many decades. Computers in the beginning could only do about 16,000 operations per second. Computer processing power, however, has been doubling every two years thanks to Moore's law, and growing even faster due to massively parallel architectures. Finally, 60 years after the first AI conference we have computers on the order of the performance of the human brain (10^{16} operations per second). The main issues now are algorithms, software, and learning. We have excellent models of neurons, such as the Hodgkin-Huxley model, but we do not know how the human neurons are wired together. With careful attention to efficient parallel computing, event-driven programming, table lookups, and memory minimization massive scale simulations can be performed. The code that will be described was written in C++ and uses the Message Passing Interface (MPI). It uses the full Hodgkin-Huxley neuron model, not a simplified model. It also allows arbitrary network structures (deep, recurrent, convolutional, all-to-all, etc.). The code is scalable, and has, so far, been tested on up to 2,048 processor cores using 10^7 neurons and 10^9 synapses.

Lyle Long
Pennsylvania State Univ

Date submitted: 10 Nov 2016 Electronic form version 1.4