SAND2011-8576C

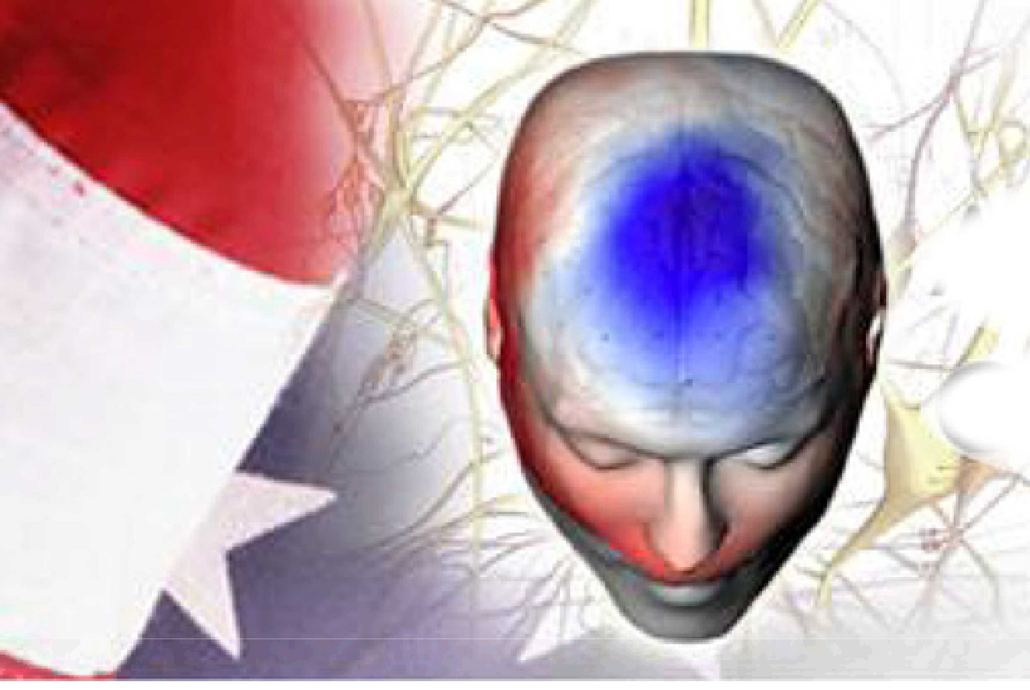

Advanced partitioning and integration techniques to improve parallel performance of densely connected neuron simulations

Richard Schiek[1], Christina Warrender[2], Heidi Thornquist[1], Ting Mei[1], Eric Keiter[1], Tom Russo[1], Eric Rankin[1]

Sandia National Laboratories, Albuquerque, NM

[1]Electrical Systems Modeling, [2]Cognitive Modeling

Introduction

The goal of this work is to develop computational tools to enable efficient parallel simulation of networks of detailed neuron models. Algorithmic advances in parallel partitioning, preconditioning and time integration can significantly improve parallel performance. This work examines linear, graph and hyper-graph partitioning schemes, complete and incomplete preconditioning and BDF vs. trapezoidal time integration techniques.

We have implemented standard Hodgkin-Huxley and cable equations within a parallel circuit simulator, Xyce (xyce.sandia.gov).

Potential advantages of Xyce:

- separation of model topology and solver methods
- simulator automatically handles parallelization
- arbitrary splitting of neurons across processors

Some of these potential advantages come at the expense of optimizations used by existing neural simulators, so we use Neuron simulations as a control and to illustrate areas where improvements are needed.

Approach

To simulate systems of graded complexity, we generated random networks of 1000 branched neurons connected on average at 10, 20, 50 or 100 synapses per neuron.

Neural complexity (number of branches) was proportional to the average #presynaptic connections for each neuron. Example of increasing connectivity and neural complexity:

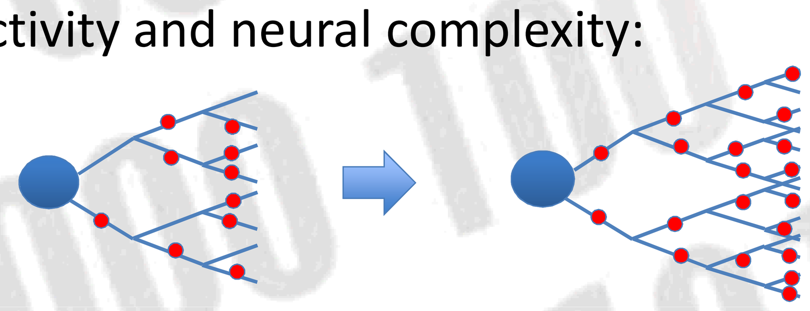

Scaling synaptic density as we have done here geometrically increases the network complexity as synapses are added because neuronal branching is simultaneously increased. Thus, as the networks get more highly connected, they also get significantly denser, adding to the numerical complexity.

Results

Scaling

Synapses per Neuron	Num. Unknowns
10	86,000
50	366,000
100	736,000
1000	6,116,000
10,000	60,498,000

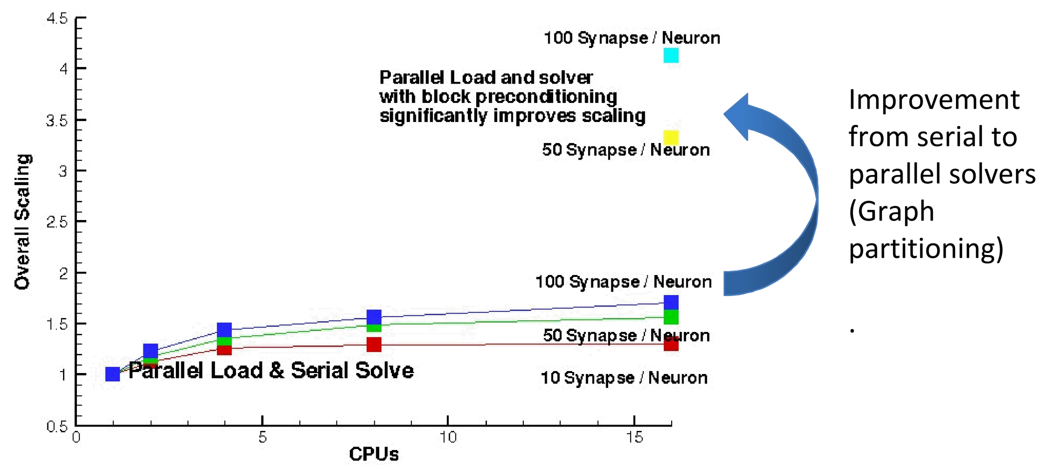

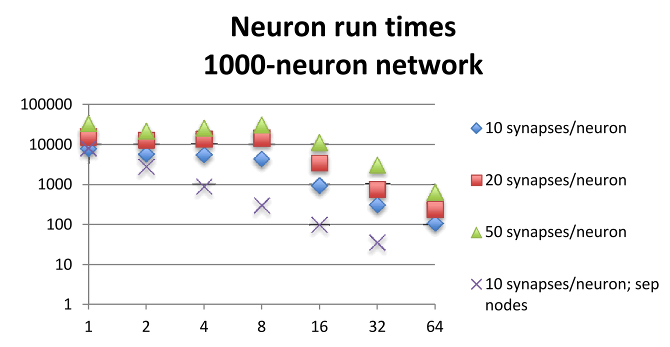

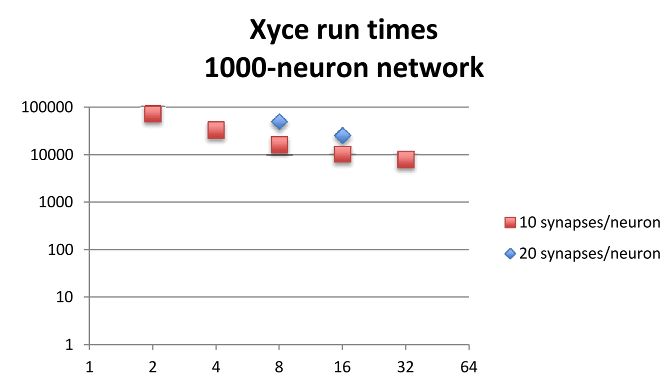

Xyce takes larger time steps on average, but is slower for this problem size. Note, however, that the modeler does not have to manage problem distribution and load balancing.

Switching from serial to parallel solvers pays off when the system size is over around 250,000.

Accuracy

Parallel scale up may change the underlying problem. Differences can be due to solver problems, time integration problems and/or model sensitivities.

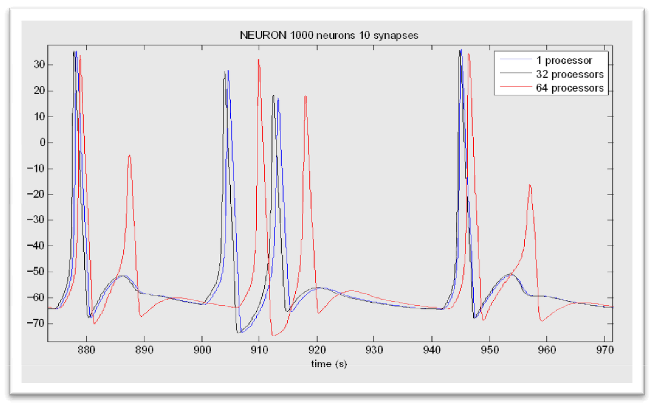

Neuron's solution scheme scales well, but can lead to different activity patterns for different numbers of processors. Whether these are functionally equivalent is a question for the researcher.

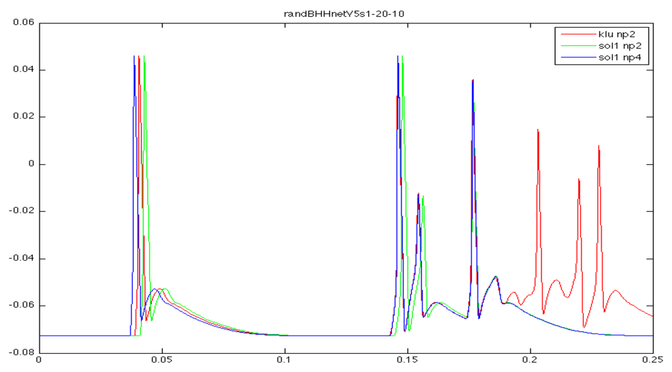

Xyce's serial and parallel solvers yielded similar activity for most of the simulation, but diverged near the end. This could be a numerical instability of the neuron or synapse model, since the more stable serial method showed continued activity whereas the iterative solver failed to find it.

Discussion

In other problem domains, solving very large simulations on parallel machines has required use of iterative solvers and appropriate preconditioning methods. Xyce offers a platform for evaluating such techniques for neural network simulation. While Neuron uses Gaussian elimination to solve the underlying system, we use KLU in serial for small networks and a parallel iterative solver for the larger simulations (gmres from AztecOO). The use of parallel, iterative solvers can be quite efficient if the system is well posed via preconditioning. We used both simple global, block triangular form (BTF) and hypergraph partitioning, and found that BTF did better than hypergraph. The Gaussian Elimination used by Neuron works very well but can be limited to topologies in which there is minimal splitting of neurons across processors. The simulations in Xyce show that new techniques in linear algebra can lead to significant parallel scaling independent of the network topology, with arbitrary splitting of neurons and synapses across processor boundaries. It is also significant and useful that in Xyce one can try multiple solution techniques to find what works best for the underlying problem.

NNSA National Nuclear Security Administration

Sandia National Laboratories is a multi-program laboratory managed and operated by Sandia Corporation, a wholly owned subsidiary of Lockheed Martin Corporation, for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-AC04-94AL85000.

Sandia National Laboratories

Xyce PARALLEL ELECTRONIC SIMULATOR