BMC Neuroscience

Poster presentation

Fully implicit parallel simulation of single neurons
Michael Hines*[1] and Felix Schuermann[2]

Address: [1]Department of Computer Science, Yale University, New Haven, CT 06520, USA and [2]Brain Mind Institute, EPFL, Lausanne, Switzerland

Email: Michael Hines* - michael.hines@yale.edu

* Corresponding author

When tree topology matrices are divided into subtrees where each subtree is on a different cpu and with the constraint that other subtrees are not connected to a given subtree at more than two distinct points (defining a backbone path on that subtree), the entire system remains amenable to direct gaussian elimination. The complexity increase is twice the number of divisions and four times the number of multiplications normally required along the backbones due to the necessity, during the triangularization phase, of transforming the tridiagonal backbone into an N topology matrix. In addition, each subtree is required to send its root diagonal and right hand side element, or, in the case of a subtree with a backbone, the 2 × 2 matrix and right hand sides of the backbone end points, to one of the cpus where that information is added together to form a reduced tree matrix of rank equal to the number of split points on the cell. The reduced tree matrix equation is solved, giving the voltages at the split points, and this information is sent back to the appropriate subtrees on the other cpus. Those subtrees with backbones can then use the N topology to quickly compute the voltages along the backbone and everyone can complete the back substitution phase of their gaussian elimination. Accuracy is the same as with standard gaussian elimination on a single cpu and any quantitative differences are attributed to accumulated round off error due to different ordering of subtrees containing backbones.

With this method, it is often feasible to divide a 3-d reconstructed neuron model into a dozen or so pieces and experience almost linear speedup. We have used the method for purposes of load balance in network simulations when some cells are very much larger than the average cell and there are more cpus than cells. The method is available in the current standard distribution of the NEURON simulation environment.