

---

Edmond Chow  
**Reducing communication costs in distributed relaxation  
methods**

School of Computational Science and Engineering  
Georgia Institute of Technology  
Atlanta  
GA  
USA  
echow@cc.gatech.edu  
Jordi Wolfson-Pou

Inter-node communication has high energy cost compared to computation. In exascale computing, reducing communication and thus energy cost may become even more important than reducing overall time. We apply this extreme objective to relaxation-type iterative methods and propose a new, “parallel Southwell” iteration for solving sparse linear systems. Standard Southwell iteration relaxes one row at a time, corresponding to the row with the highest residual, and is a sequential algorithm by definition. Southwell generally converges faster than Gauss-Seidel in terms of the number of relaxations. Since each relaxation is associated with communication, communication cost is also reduced. However, to achieve this faster convergence, Southwell needs global communication after every relaxation step to determine the row with the highest residual.

In parallel Southwell, we approximate the standard Southwell iteration by 1) avoiding global communication, and 2) allowing some rows to be relaxed simultaneously. To describe the method, assume one processor node assigned to one row or grid point. At each iteration, each node determines if it is the node with the largest residual among its neighbors in the grid. If so, then this node relaxes its row and sends updated values to its neighbors. Nodes that have recently relaxed do not relax again for a given number of iterations, to prevent too many nodes relaxing simultaneously. Open questions include how each node maintains an estimate of the residual on neighboring nodes (this process does not need to be exact). Also, the algorithm is ideally implemented as an asynchronous method, and thus we discuss the challenges of implementing this method using MPI one-sided communication (remote memory access). The goal is to develop a method to solve problems with less communication than with Gauss-Seidel or Jacobi with the possibility of also being faster on extreme scale distributed systems.