

Parallel Solution of the Boundary Integral Equations for Elasticity ^{† ‡}

Tan, E. S., Gray, L. J., and Phan, A.-V.
Computer Science and Mathematics Division
Oak Ridge National Laboratory, Oak Ridge TN 37831-6367.

The primary motivation for this work is to have the ability to carry out crack propagation simulations in materials at the meso-scale level. These analyses will incorporate multiple cracks, voids, and grain boundaries, and this small scale structure naturally leads to a large scale calculation. As a consequence a PC cluster implementation has been developed.

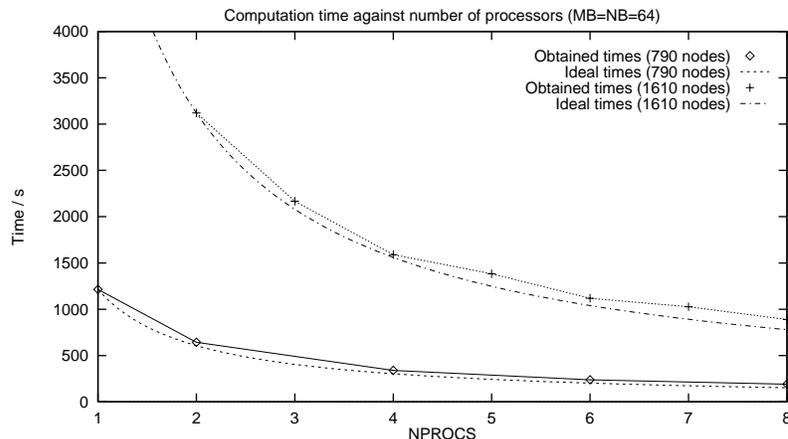
The two-dimensional linear elastic analysis is achieved by solving the boundary integral equations for surface displacement and surface traction:

$$u(P) + \int_{\Gamma} T(P, Q)u(Q) dQ = \int_{\Gamma} U(P, Q)\tau(Q) dQ$$
$$\tau(P) + \int_{\Gamma} W(P, Q)u(Q) dQ = \int_{\Gamma} S(P, Q)\tau(Q) dQ.$$

These equations are numerically approximated using the symmetric-Galerkin approximation employing quadratic elements. The matrix element generation and the solution to the system of linear equations $A \cdot x = b$, the two most processor intensive routines, are made to execute in parallel. The I/O routines are left to run in serial to minimise overhead costs in inter-processes communications. The machines in the cluster were equipped with PVM on the message passing tier with ScaLAPACK (and its dependencies) built on it.

In calling the ScaLAPACK PDGESV linear system solver, depending on the data set, different setups for the cluster may yield a superior performance; thus the program accepts the following parameters as inputs: (i) number of process rows in the process grid, $NPROW$, (ii) number of process columns in the process grid, $NPCOL$, and (iii) global row block size for partitioning in the global matrix, MB (equals the global row column size, NB). First, the processes are aligned in a one-dimensional block row-distribution of $(NPROW \times NPCOL)$ rows for the matrix element generation. Upon the completion of the matrix element generation, each process stores row blocks of the global matrices. The matrix elements are then redistributed to a two-dimensional block-cyclic distribution by the PDGEMR2D routine to prepare the global matrix for the PDGESV solver.

For large matrices ($n \gtrsim 400$) with k processors in the cluster, the original computation time T , can be reduced to about 10–20% of T/k .



[†] This research was supported by the Applied Mathematical Sciences Research Program of the Office of Mathematical, Information, and Computational Sciences, U.S. Department of Energy under contract DE-AC05-96OR22464 with Lockheed Martin Energy Research Corp.

[‡] Submitted to European Congress on Computation Methods in Applied Sciences and Engineering, 11–14 September 2000, at Barcelona, Spain. <http://www.cimne.upc.es/eccomas/>. Session title: Parallel Computing on PC Cluster.