Highly Scalable Dense Linear Solver Based on High Performance Linpack (HPL)

R. Barrett*, T. Chan†, **E. D'Azevedo***, E. Jaeger‡, K. Wong§, and R. Wong†

This poster describes the effort for enhancing the performance of parallel linear solver in the SciDAC fusion code AORSA2D by adapting the High Performance Linpack (HPL) benchmark, commonly used in ranking the top 500 computers. The modified HPL solver is more scalable than the ScaLAPACK library. The original HPL software was configured as a benchmark to solve a randomly generated double precision matrix. HPL was modified using scripts in a semi-automatic manner to solve a double complex system and provide a software interface that is compatible with the ScaLAPACK library. The result is a significant increase in performance, achieving 87.5TFLOPS on over 20,000 processors on the Cray XT4 at ORNL.

AORSA2D (All Orders Spectral Algorithm) simulation program is developed within the SciDAC *Numerical Computation of Wave Plasma-Interactions in Multi-dimensional Systems* for modeling the radio frequency (RF) heating of plasma in a Tokamak device. AORSA2D uses Fourier basis to represent the electric field and uses collocation on a rectangular grid to generate a large dense complex linear system. A 500×500 modes problem will lead to a system size of 524,474. Thus the solution of the complex linear system is a major computational kernel.

The algorithm used in HPL is a right-looking variant of the block LU factorization with row partial pivoting featuring multiple look-ahead depths, recursive panel factorization with pivot search and column broadcast combined, various virtual panel broadcast topologies, bandwidth reducing swap-broadcast algorithm, and a ring broadcast algorithm in the backward substitution phase.

Our current effort is in developing mixed precision iterative solver where the LU factorization is computed with reduced precision in HPL. For matrices that are not too ill-conditioned, this approach has been shown to produce significant performance benefits.

*Computer Science and Mathematics Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831. Research supported by the Office of Scientific Computing Research, U.S. Department of Energy. The work was performed at the Oak Ridge National Laboratory, which is managed by UT-Battelle, LLC under Contract No. DE-AC05-00OR22725. The research used resources of the Center for Computational Sciences at the Oak Ridge National Laboratory, which is supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC05-00OR22725.

†Chinese University of Hong Kong, Hong Kong, China. Summer Internships for T. Chan and R. Wong were supported by the Department of Mathematics, Chinese University of Hong Kong. Internship opportunity was provided by the Joint Institute for Computational Sciences, the University of Tennessee and the Oak Ridge National Laboratory.

‡Fusion Energy Division, Oak Ridge National Laboratory, Oak Ridge, TN 37831.

§Joint Institute for Computational Sciences, University of Tennessee, TN, USA.

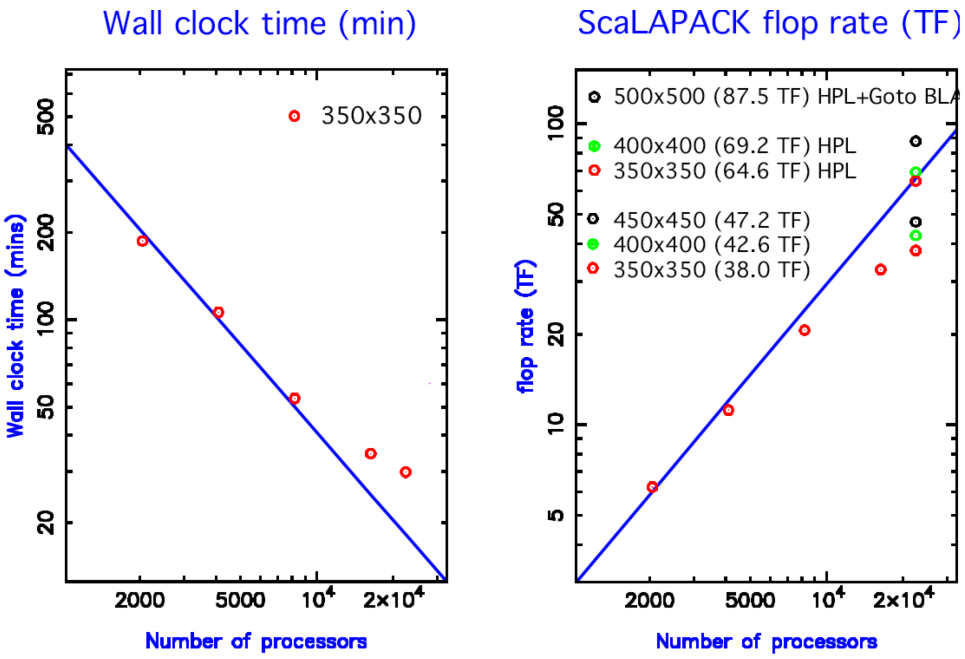

Figure 1: Performance of AORSA2D on Jaguar using ScaLAPACK and HPL.