Abstract Submitted
for the DPP12 Meeting of
The American Physical Society

Scaling Optimization of the SIESTA MHD Code[1] SUDIP SEAL, STEVEN HIRSHMAN, KALYAN PERUMALLA, Oak Ridge National Laboratory — SIESTA is capable of computing three-dimensional plasma equilibria with magnetic islands at high spatial resolutions for toroidally confined plasmas. Originally, it was developed to exploit the small-scale parallelism offered by shared-memory machines. Algorithms designed for shared-memory machines do not necessarily scale to large number of processors on large distributed-memory machines. Mainly, this is because communication overheads in distributed-memory architectures can easily and very quickly dominate the gain in computation time as the number of processors is increased. Here, we report the results of a scaling effort that increases both the speed and resolution of the SIESTA magnetohydrodynamic equilibrium code on large-scale tightly coupled distributed-memory computing platforms. We investigate the performance profile of the original SIESTA code to identify scale-dependent bottlenecks and develop scalable alternative functionality. This improves both its runtime speed (on the same number of processors) as well as its scalability (across larger number of processors) by an order of magnitude. The net outcome allows SIESTA to utilize a few thousand processors and to simulate high spatial-resolution scenarios in under an hour for the first time.

[1]Work supported by U.S. DOE under Contract DE-AC05-00OR22725 with UT-Battelle, LLC.

Steven Hirshman
Oak Ridge National Laboratory

Date submitted: 17 Jul 2012 Electronic form version 1.4