Abstract Submitted
for the DPP13 Meeting of
The American Physical Society

Graphics Processing Unit Acceleration of Gyrokinetic Turbulence Simulations BENJAMIN HAUSE, SCOTT PARKER, YANG CHEN, University of Colorado, Boulder — We find a substantial increase in on-node performance using Graphics Processing Unit (GPU) acceleration in gyrokinetic delta-f particle-in-cell simulation. Optimization is performed on a two-dimensional slab gyrokinetic particle simulation using the Portland Group Fortran compiler with the OpenACC compiler directives and Fortran CUDA. Mixed implementation of both Open-ACC and CUDA is demonstrated. CUDA is required for optimizing the particle deposition algorithm. We have implemented the GPU acceleration on a third generation Core I7 gaming PC with two NVIDIA GTX 680 GPUs. We find comparable, or better, acceleration relative to the NERSC DIRAC cluster with the NVIDIA Tesla C2050 computing processor. The Tesla C 2050 is about 2.6 times more expensive than the GTX 580 gaming GPU. We also see enormous speedups (10 or more) on the Titan supercomputer at Oak Ridge with Kepler K20 GPUs. Results show speed-ups comparable or better than that of OpenMP models utilizing multiple cores. The use of hybrid OpenACC, CUDA Fortran, and MPI models across many nodes will also be discussed. Optimization strategies will be presented. We will discuss progress on optimizing the comprehensive three dimensional general geometry GEM code.

Scott Parker
University of Colorado, Boulder

Date submitted: 10 Jul 2013 Electronic form version 1.4