



Contribution ID: 422

Type: **Oral presentation to parallel session**

Synergia-CUDA: GPU Accelerated Accelerator Modeling Package (video conference)

Monday, 14 October 2013 14:16 (22 minutes)

Synergia is a parallel, 3-dimensional space-charge particle-in-cell code that is widely used by the accelerator modeling community. We present our work of porting the pure MPI-based code to a hybrid of CPU and GPU computing kernels. The hybrid code uses the CUDA platform, in the same framework as the pure MPI solution. We have implemented a lock-free collaborative charge-deposition algorithm for the GPU, as well as other optimizations, including local communication avoidance for GPUs, customized FFT, and fine-tuned memory access patterns. On a small GPU cluster (up to 4 Tesla C1070 GPUs), our benchmarks exhibit both superior peak performance and better scaling, when compared to a CPU cluster with 16 nodes and 128 cores. We have further compared the code performance on different GPU architectures, including C1070 Tesla, M2070 Fermi, and K20 Kepler. We show 10 to 20% performance increases with optimizations addressing each specific hardware architectures.

Summary

Primary author: LU, Qiming (Fermi National Accelerator Laboratory)

Co-author: AMUNDSON, James (Fermi National Accelerator Laboratory)

Presenter: LU, Qiming (Fermi National Accelerator Laboratory)

Session Classification: Software Engineering, Parallelism & Multi-Core

Track Classification: Software Engineering, Parallelism & Multi-Core