

Abstract Submitted
for the DPP10 Meeting of
The American Physical Society

Adaptable Particle-in-Cell Algorithms for Graphical Processing Units¹ VIKTOR DECYK, TAJENDRA SINGH, UCLA — Emerging computer architectures consist of an increasing number of shared memory computing cores in a chip, often with vector (SIMD) co-processors. Future exascale high performance systems will consist of a hierarchy of such nodes, which will require different algorithms at different levels. Since no one knows exactly how the future will evolve, we have begun development of an adaptable Particle-in-Cell (PIC) code, whose parameters can match different hardware configurations. The data structures reflect three levels of parallelism, contiguous vectors and non-contiguous blocks of vectors, which can share memory, and groups of blocks which do not. Particles are kept ordered at each time step, and the size of a sorting cell is an adjustable parameter. We have implemented a simple 2D electrostatic skeleton code whose inner loop (containing 6 subroutines) runs entirely on the NVIDIA Tesla C1060. We obtained speedups of about 16-25 compared to a 2.66 GHz Intel i7 (Nehalem), depending on the plasma temperature, with an asymptotic limit of 40 for a frozen plasma. We expect speedups of about 70 for an 2D electromagnetic code and about 100 for a 3D electromagnetic code, which have higher computational intensities (more flops/memory access).

¹Work supported by USDOE.

Viktor Decyk
UCLA

Date submitted: 26 Jul 2010

Electronic form version 1.4