

Abstract Submitted
for the DPP11 Meeting of
The American Physical Society

Efficient GPU implementation for implicit particle-in-cell simulations G. CHEN, L. CHACÓN, ORNL, D.C. BARNES, Coronado Consulting — A recent proof-of-principle study of an energy- and charge-conserving, fully implicit particle-in-cell (PIC) algorithm¹ demonstrated that accurate and efficient PIC simulations with very large time steps are possible. A key component of the algorithm is the enslavement of particle orbits to the field equations. With particle enslavement, orbit integration is a segregated operation, which is perfectly suited for emerging heterogeneous architectures that combine CPUs with GPUs. The use of GPUs is promising on implicit PIC, as it is naturally data parallel (thus suited for extreme multi-threading), and it is compute-bounded (vs. explicit schemes, typically memory-bounded). However, the particle mover in [1] is adaptive, and particles have to stop at cell-boundaries to conserve charge locally. This creates load imbalances and dynamic control flows, which poses a challenge to utilize fully the GPU computing power. This work demonstrates that a highly efficient GPU implementation of the implicit particle mover (using CUDA) is possible. We obtain 300 to 400 GOps/s (counting floating, integer and special function operations) using single precision. This is about 20% to 25% of the peak performance of the GPU (GeForce GTX580), and about 200 to 300 times faster than a single CPU (Xeon@3.16GHz) implementation.

¹G. Chen, L. Chacón, and D.C. Barnes, *J. Comput. Phys.* **18**(2011).

Guangye Chen

Date submitted: 19 Jul 2011

Electronic form version 1.4