

Abstract Submitted
for the DPP15 Meeting of
The American Physical Society

GPU acceleration of particle-in-cell methods¹ BENJAMIN COWAN, JOHN CARY, DOMINIC MEISER, Tech-X Corporation — Graphics processing units (GPUs) have become key components in many supercomputing systems, as they can provide more computations relative to their cost and power consumption than conventional processors. However, to take full advantage of this capability, they require a strict programming model which involves single-instruction multiple-data execution as well as significant constraints on memory accesses. To bring the full power of GPUs to bear on plasma physics problems, we must adapt the computational methods to this new programming model. We have developed a GPU implementation of the particle-in-cell (PIC) method, one of the mainstays of plasma physics simulation. This framework is highly general and enables advanced PIC features such as high order particles and absorbing boundary conditions. The main elements of the PIC loop, including field interpolation and particle deposition, are designed to optimize memory access. We describe the performance of these algorithms and discuss some of the methods used.

¹Work supported by DARPA contract W31P4Q-15-C-0061 (SBIR).

Benjamin Cowan
Tech-X Corporation

Date submitted: 24 Jul 2015

Electronic form version 1.4