Abstract Submitted
for the DPP17 Meeting of
The American Physical Society

GPU Acceleration of Particle-In-Cell Methods[1] BENJAMIN COWAN, SERGEY AVERKIN, Tech-X Corporation, JOHN CARY, Tech-X Corporation and University of Colorado, JARROD LEDDY, SCOTT SIDES, Tech-X Corporation, GREGORY WERNER, University of Colorado — 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 access. 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 recent progress in these algorithms, including arbitrary grid types and multiple GPUs per node.

[1]Work supported by DARPA Contract No. W31P4Q-16-C-0009.

Benjamin Cowan
Tech-X Corporation

Date submitted: 14 Jul 2017 Electronic form version 1.4