Abstract Submitted
for the GEC18 Meeting of
The American Physical Society

Efficient use of GPUs for Particle-in-Cell codes PETER MESSMER, NVIDIA — The physical accuracy of Particle-in-Cell codes is strongly related to the available computing power. It is therefore no surprise that PIC codes are often among the first to embrace novel hardware architectures like GPU accelerated hybrid systems. With the recently introduced Summit supercomputer at Oak Ridge National Laboratory, largest-scale GPU-accelerated systems have been taken to the next level. And while the fundamental architecture is still unchanged, GPUs have seen tremendous advances over the past decade, both in terms of performance and usability, affecting hardware, software and interconnect. In this presentation, I will review some of the latest architectural changes in the light of core algorithms of PIC codes: How can we optimally deposit charge on latest generation GPUs? How to use multiple GPUs for Poisson solvers? What programming models to use for what parts of the code? etc. In addition, we will also look at hardware features targeting domains like machine learning or computer graphics and how they could be used for kinetic plasma simulations.

Peter Messmer
NVIDIA

Date submitted: 18 Jun 2018 Electronic form version 1.4