

MAR14-2013-020150

Abstract for an Invited Paper
for the MAR14 Meeting of
the American Physical Society

Particle-in-Cell Simulations on Emerging Architectures¹

VIKTOR DECYK, University of California, Los Angeles

Emerging High-Performance Computing (HPC) Architectures consist of multiple accelerators with multiple vector processors. Such platforms require programming with multiple levels of parallelism and pose considerable challenges for performing simulations of plasmas. In this talk we focus on lessons learned from our experience with implementations of Particle-in-Cell (PIC) codes on GPUs. We will discuss aspects of mixed shared memory/distributed memory algorithms, streaming, blocking (tiling), and vector , data structures, latency and load balancing. Many of these lessons are of course familiar from other architectures, but will be seen here in a new perspective. We will discuss strategies for the development of algorithms for PIC codes that are expected to work on a variety of emerging HPC architectures.

¹Supported by USDOE, NSF, NASA, and UCLA IDRE.