

Abstract Submitted
for the DPP19 Meeting of
The American Physical Society

VPIC on GPU¹ ROBERT BIRD, PATRICK KILLIAN, BRIAN ALBRIGHT, Los Alamos National Laboratory — Efficient operation of Particle-in-Cell codes on Graphics Processing Units (GPUs) has been a coveted goal since they were adopted by High Performance Computing platforms years ago. While a variety of research exist on this topic, many of the worlds highest performing PIC codes have yet to demonstrate their ability to effectively use large GPU machines at extreme scale. In this work we demonstrate the effort of Los Alamos National Laboratory to port VPIC to run on large-scale DoE GPU super computers. We demonstrate the codes ability to scale to thousands of GPUs, directly compare code performance to previous systems, and present lessons learnt from porting to the performance-portable code framework Kokkos. We document our strategy for data management in the context of the limited memory regime presented by GPUs, and also demonstrate the strategies we employ to minimize data copies between host and device.

¹Work performed under the auspices of the U.S. DOE by Triad National Security, LLC, LANL.

Robert Bird
Los Alamos National Laboratory

Date submitted: 02 Jul 2019

Electronic form version 1.4