Abstract Submitted
for the DPP19 Meeting of
The American Physical Society

Scalable, Performance-Portable Particle-in-Cell Simulations and PByte-Scale Data-Challenges A. HUEBL, Lawrence Berkeley National Laboratory, Helmholtz-Zentrum-Dresden-Rossendorf, R. WIDERA, Helmholtz-Zentrum Dresden-Rossendorf, M. GARTEN, Helmholtz-Zentrum Dresden-Rossendorf, Technical University Dresden, R. PAUSCH, K. STEINIGER, S. BASTRAKOV, A. DEBUS, T. KLUGE, S. EHRIG, F. MEYER, M. WERNER, Helmholtz-Zentrum Dresden-Rossendorf, B. WORPITZ, LogMeIn Inc., A. MATTHES, F. POESCHEL, Helmholtz-Zentrum Dresden-Rossendorf, Technical University Dresden, S. STARKE, M. BUSSMANN, Helmholtz-Zentrum Dresden-Rossendorf — We present the architecture, abstractions, novel developments, and workflows that enable high-resolution, fast turn-around computations on contemporary, leadership-scale supercomputers powered by both GPUs and CPUs from various vendors and on top of a generalized programming model (Alpaka). From the experience developing the open-source community code PIConGPU, strategies for handling PByte-scale data flows from thousands of computing devices for analysis with in situ processing and open data formats (openPMD) are presented. Furthermore, simulation control via a lightweight Python Jupyter interface as well as recent research towards just-in-time kernel generation for C++ with Cling-CUDA are shown as a mean for fast turn-around, close-to-experiment simulations.

Ruth Teferi
APS

Date submitted: 10 Jul 2019 Electronic form version 1.4