

PIConGPU

A scalable implementation of the Particle-in-Cell Algorithm for GPU Clusters

Michael Bussmann, Heiko Burau, Florian Berninger, Thomas Kluge,
Alexander Debus, Ulrich Schramm, Thomas E. Cowan
Helmholtz-Zentrum Dresden-Rossendorf, Dresden, Germany

Felix Schmitt, René Widera, Wolfgang Hönig, Guido Juckeland, Wolfgang Nagel
Zentrum für Informationsdienste und Hochleistungsrechnen

We present PIConGPU, a performant implementation of the particle-in-cell algorithm for GPUs that is scalable on GPU clusters. PIConGPU is used for fast-response simulations of laser-plasma interaction, including laser wakefield acceleration using the sliding-window technique. We discuss lessons learned from going from the initial two-dimensional PIC-implementation to a full 3D implementation, focusing on data storage on the GPU and data communication between GPU nodes in a cluster.

We show how communication and data storage can be efficiently hidden from users who want to extend the code by adding new physics so that users can assume to be working in a single-data single-instruction environment without deeper knowledge of GPU programming. As an example we show how the far field of relativistic electrons performing betatron oscillations in a laser-driven wakefield can be calculated from macro-particle trajectories on the GPU and subsequently stored in CPU.

First simulation results obtained with PIConGPU are shown to illustrate the advantage of fast response simulations for large parameter scans.