High-Performance Implementation Techniques of CUDA-based 1D and 2D Particle-in-Cell/MCC Plasma Simulations

Zoltán Juhász[1], Péter Hartmann[2], Zoltán Donkó[2]

[1] Department of Electrical Engineering and Information Systems, University of Pannonia, Veszprém, Hungary

[2] Department of Complex Fluids, Institute for Solid State Physics and Optics, Wigner Research Centre for Physics, Hungarian Academy of Sciences, Budapest, Hungary

The goal of our work is to create high-performance GPU implementations for Particle-in-Cell plasma simulations. The motivation lies in the fact that the simulation of low-pressure electrostatic plasma systems with Monte-Carlo collisions results in extreme execution times (ranging from several days to months). Speeding up simulations with parallelism using GPUs seemed to be the most promising implementation option. Over the past two years we have developed CUDA GPU implementations for one and two-dimensional plasma systems achieving two orders of magnitude speedup. Results of this work (approach to parallelization, achieved accuracy and performance) were already presented in the GPU Day 2017 and 2018 events.

In this talk, we would like to concentrate on the performance optimization aspects of this work. We describe the approach we took to port the original sequential implementation to GPU and in a step-by-step refinement fashion we demonstrate the performance benefits of various optimization methods, such as algorithmic modifications, CUDA kernel fusion, the use of shared memory, memory access optimizations, concurrent kernels and streams, high-performance atomic and warp operations.

Finally, we outline possible future improvement opportunities, including the use of multi-GPU systems.

Contact address: juhasz@virt.uni-pannon.hu