Highly Parallel GPU-based Particle-in-Cell/MCC Plasma Simulation

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

[1]Dept. of Electrical Engineering and Information Systems, University of Pannonia, Veszprem, Hungary

[2]Dept. of Complex Fluids, Institute for Solid State Physics and Optics, Wigner Research Centre for Physics, Budapest, Hungary

Simulation is a common tool in studying the behaviour of complex physical systems. It allows researchers to explore the effect of various parameter changes and assists the understanding of the details of physical mechanisms. Our interest lies in the simulation of low-temperature plasmas, which represent a special case of N-body particle simulations in which particles affect one another by electrostatic forces. In these systems, particle collisions are crucial to maintaining the plasma state via ionization processes. Due to the high number of particles required for accurate simulations and the low speed of convergence, plasma simulation is a time-consuming process on sequential computers (runtimes often range from days to weeks). As a result, parallel computing has been a long-term aid of particle simulations. For collision-free plasma simulation, several parallel implementations exist that show good performance on clusters and supercomputers alike. Parallel simulation of highly collisional plasmas that are of our interest, on the other hand, is a far greater challenge.

The aim of our project is to create a parallel implementation of a validated and trusted sequential onedimensional plasma simulation code written in C at the Wigner Research Centre. Initial experiments with an OpenMP-based approach identified problems (e.g. non-uniform and unpredictable memory access patterns) that prove simple CPU-based parallelisation approaches incapable of providing acceptable speedup levels. Our current focus is thus on GPU architecture and the CUDA programming framework. We plan to explore the possibilities as well as limitations of this technology in implementing Monte-Carlo collision Particle-in-Cell plasma simulations hoping that significant speedup is achievable with this approach.

In the case of one-dimensional simulations, the number of particles is typically in the same order as the number of GPU threads required for efficient execution on state-of-the-art GPU cards. This presents a significant challenge in algorithm design and implementation, and questions the applicability of tried-and-tested existing multi-CPU algorithms.

The talk will give a brief overview of the basics of plasma simulation and approaches to parallelisation. Afterwards, we will describe the details of our parallel implementation in its current state focusing on design decisions, mapping the problem to several thousand cores, the used optimization techniques, and parallel data structure management. We finish with discussing the current performance of the program and the lessons learned from the project.