

A Task-Based Particle-in-Cell Method with Automatic Load-Balancing using the AllScale Environment

Roman Iakymchuk¹, Herbert Jordan², Philipp Gschwandtner², Thomas Heller³, Peter Thoman², Xavier Aguilar¹, Thomas Fahringer², Erwin Laure¹, Stefano Markidis¹

¹*KTH Royal Institute of Technology, Sweden* {riakymch,xaguilar,markidis,erwin}@kth.se

²*University of Innsbruck, Austria* {herbert,philipp,petert,tf}@dps.uibk.ac.at

³*University of Erlangen-Nürnberg, Germany* thomas.heller@fau.de

The Particle-in-Cell (PIC) method is one of the most common and powerful numerical techniques for the simulation of fusion, astrophysical, and space plasmas. For instance, PIC simulations are used to study the interaction of the Earth's electromagnetic field with the hot plasma emanated by the sun, the so-called *solar wind*. To note, high energy plasma in space can damage spacecrafts and put at risk the life of astronauts. Thus, it is important to enable efficient large-scale PIC simulations that are capable to predict different phenomena in space.

The PIC method solves the kinetic equation of plasmas by first sampling plasma distribution functions with computational particles and then following their trajectories by solving the equation of motion for each particle. The electromagnetic field determining the particle trajectory is calculated by computing the Maxwell's equations on a grid. The coupling between particle and both electric and magnetic fields is provided by the so-called *interpolation* functions (aka gather and scatter). Typically, parallel PIC simulations divide the simulation box in several equal-in-size domains with initially the same number of particles [1]. Each domain is assigned to a process that carries out the computation for the particles in the domain. When a particle exits the domain, it is communicated to a different domain. Because of the non-uniform configuration of the electromagnetic field in space, computational particles concentrate in relatively small spatial regions while few particles cover other spatial regions, resulting in load imbalance. *Workload-imbalance* is the most severe and limiting problem in large-scale PIC simulations, resulting in up to 60% of the process imbalance.

In this talk, we present a new formulation of the PIC method to provide automatic load-balancing using the AllScale toolchain¹. The AllScale approach, which is based on task-centric nested recursive parallelism, aims to provide a unified programming system for the effective development of highly scalable, resilient and performance-portable parallel applications for Exascale systems. In the spirit of this approach, we redesign the PIC method: particles are placed inside individual cells rather than both structures stored separately, requiring extra mappings. Moreover, the field solver implements an explicit scheme which allows to eliminate global communications and synchronizations. This facilitates the accommodation of the nested, recursive, and asynchronous task parallelism that is in the core of the AllScale toolchain. Furthermore, our method employs dynamic load balancing: the simulation is recursively divided into smaller domains, forming a set of tasks to be computed. The granularity and distribution of those tasks is thereby controlled by the scheduler provided by the AllScale runtime system, leading to an even load balance where each processing unit spends approximately the same amount of time on processing its assigned share of particles. We demonstrate our new approach using a case study with severe load imbalance, the Earth's Van Allen radiation belts.

References

- [1] S. Markidis et al. Multi-scale simulations of plasma with iPIC3D. *Mathematics and Computers in Simulation*, 80(7):1509–1519, 2010.

¹The AllScale project has received funding from EC H2020 programme (No. 671603). allscale.eu