Scalable and efficient algorithms for kinetic plasma physics simulations

Lukas Einkemmer

University of Innsbruck

Plasma physics provides a number of challenging computational problems. This is, in particular, true for the kinetic equations considered in this talk. Kinetic equations are posed in an up to six-dimensional phase space and their discretization by standard methods leads to a severe CFL restriction. For such problems, semi-Lagrangian methods based on spline interpolation are often employed in practice. However, this approach introduces non-local data communication into the numerical algorithm. This significantly limits the scalability of the algorithm on clusters and poses significant challenges for their efficient implementation on graphic processing units (GPUs). Thus, it is of great interest to construct numerical methods that do not suffer from these problems and to demonstrate their efficiency on modern computer systems.

In this talk, we will consider the more recently developed semi-Lagrangian discontinuous Galerkin method, which besides being competitive from an algorithmic standpoint, can be implemented with only local data accesses. We demonstrate its utility by conducting a performance comparison (on the Vienna Scientific Cluster) of our semi-Lagrangian discontinuous Galerkin code `https://bitbucket.org/leinkemmer/sldg` with the SeLaLib library (which is based on cubic spline interpolation). The results show that the scalability is significantly improved as no global communication is necessary. Incidentally, this feature also improves the performance on a single node.

The developed algorithm is particularly suitable for GPUs, which require a large number of threads with fine grained parallelism. Some challenges that sets this problem apart from more traditional stencil codes is the fact that not all degrees of freedom within the discretization are treated equally, the high-dimensionality of the problem, and that the memory access pattern is local but not easily predictable. We show how these challenges are overcome in our implementation and highlight the excellent performance results that are achieved. A further advantage of the semi-Lagrangian discontinuous Galerkin approach is that it lends itself to mixed precision computations. This is possible since the degrees of freedom used in this method are naturally ordered. This can be exploited in order to further increase the performance of our algorithm on modern computer systems (in particular on accelerators).

References

[1] L. Einkemmer, Journal of Computationa Physics **376**, 937–951 (2019).

[2] L. Einkemmer, High Performance Computing and Simulation (HPCS), 2016 International Conference on.

[3] L. Einkemmer, Computer Physics Communications **202**, 326-336 (2016).