GPU acceleration of DEMO particle exhaust simulations

S. Varoutis, C. Tantos, Chr. Day

Karlsruhe Institute of Technology, Institute for Technical Physics, 76344, Eggenstein-Leopoldshafen, Germany

Simulating the DEMO particle exhaust using the DIVGAS (Divertor Gas Simulator) code, is in general a challenging computational task due to its numerical complexity and the computing cost. The DIVGAS code developed at the Karlsruhe Institute of Technology (KIT) is a verified numerical tool for investigation of the neutral gas dynamics in the subdivertor region for different divertor configurations under steady-state operation. The DIVGAS code currently is based on the stochastic Direct Simulation Monte Carlo (DSMC) method, in which the solution of the Boltzmann kinetic equation is reproduced.

Currently, the lowest computing costs are obtained when applying a graphic processing unit (GPU) originally developed for speeding up graphic processing. In this work we present an implementation of accelerating the neutral gas flow in a Single-Null DEMO divertor configuration on a GPU, using the DIVGAS code. Although the DIVGAS code provides accurate predictions for the case of two-dimensional neutral gas flows, any future extension to more complex flows (gas mixtures, polyatomic gases, three dimensional geometries) is essentially related to the possibility of reducing the memory demand and exploit the different types of memory found on GPUs, as for instance shared, local, global, constant etc. In the present work all these parallelization techniques are utilized.

For comparison purposes, two types of GPUs will be used, namely the NVIDIA Tesla K80, 24 MB and NVIDIA Tesla V100, 32 GB. Furthermore, the computation accuracy of the DIVGAS code on both described GPUs has been validated with the corresponding CPU-based benchmark case. To evaluate the performance gains, the computing time on each GPU against its sequential CPU counterpart has been compared. The measured speedups show that the GPU can accelerate the execution of the DIVGAS code by a factor 5–10, depending on the chosen particle number of simulation particles.

The parallelization approach presented here significantly reduces the cost of DIVGAS simulations and has the potential to scale to large CPU/GPU clusters, which could enable future applications, which focus on even more complex 3D neutral flow problems. The accelerated version of the DIVGAS code on GPUs is considered to be a major breakthrough in the reduction of the needed computational time for fusion related applications.