Contribution ID: **163** Type: **Sectional reports**

Characteristics of Nvidia CUDA and AMD ROCm Platforms Affecting Performance Portability

Friday, 9 July 2021 12:15 (15 minutes)

The development and popularization of the AMD ROCm platform with HIP technology allows one to create code that is not locked to a specific vendor maintaining a high level of performance. A lot of legacy but still supported codes is originally written in CUDA, and now it is getting ROCm HIP support as well. In a recent paper [1], the performance of popular molecular dynamics packages with GPU support was discussed in detail. The research includes the LAMMPS package providing backends for CUDA, OpenCL, and HIP. Based on this package, we can compare and define in detail the platform properties and performance impact of real parallel code. Differences can be found in the characteristics of the target hardware, the operation of the software environment and drivers, and even in the logic of the application code itself. In continuation of the study, the work of computational GPU kernels in the application using several MPI processes for each GPU is considered.

1. Kondratyuk N, Nikolskiy V, Pavlov D, Stegailov V. GPU-accelerated molecular dynamics: State-of-art soft

Summary

Primary author: NIKOLSKIY, Vsevolod (HSE)

Presenter: NIKOLSKIY, Vsevolod (HSE)

Session Classification: HPC

Track Classification: 5. High Performance Computing