Performance and portability of state-of-art molecular dynamics software on modern GPUs

Evgeny Kuznetsov[1], Nikolay Kondratyuk[1,2,3], Mikhail Logunov[2,3],
Vsevolod Nikolskiy[1,2], Vladimir Stegailov[1,2,3]
[1]National Research University Higher School of Economics, Moscow, Russia
[2]Joint Institute for High Temperatures of RAS, Moscow, Russia
[3]Moscow Institute of Physics and Technology, Dolgoprudny, Russia
v.stegailov@hse.ru

Classical molecular dynamics (MD) calculations represent a significant part of utilization time of high performance computing systems. As usual, efficiency of such calculations is based on an interplay of software and hardware that are nowadays moving to hybrid GPU-based technologies. Several well-developed MD packages focused on GPUs differ both in their data management capabilities and in performance. In this paper, we present our results for the porting of the CUDA backend of LAMMPS to ROCm HIP that shows considerable benefits for AMD GPUs comparatively to the existing OpenCL backend. We consider the efficiency of solving the same physical models using different software and hardware combinations. We analyze the performance of LAMMPS, HOOMD, GROMACS and OpenMM MD packages with different GPU back-ends on modern Nvidia Volta and AMD Vega20 GPUs.

Keywords: LAMMPS, HOOMD, GROMACS, OpenMM, OpenCL, Nvidia CUDA, AMD ROCm HIP.