

Суперкомпьютерные дни в России

международная научная конференция

26-27 сентября 2022 г.

State-of-the-Art Molecular Dynamics Packages for GPU Computations: Performance, Scalability and Limitations


Vsevolod Nikolskiy, Daniil Pavlov, Vladimir Stegailov



Contents

- Introduction
- Hardware used in the study
- LAMMPS
 - GPU package
 - Kokkos
- OpenMM
- Benchmarks: scaling comparison
- Current limitations

CLASSIC MOLECULAR DYNAMICS



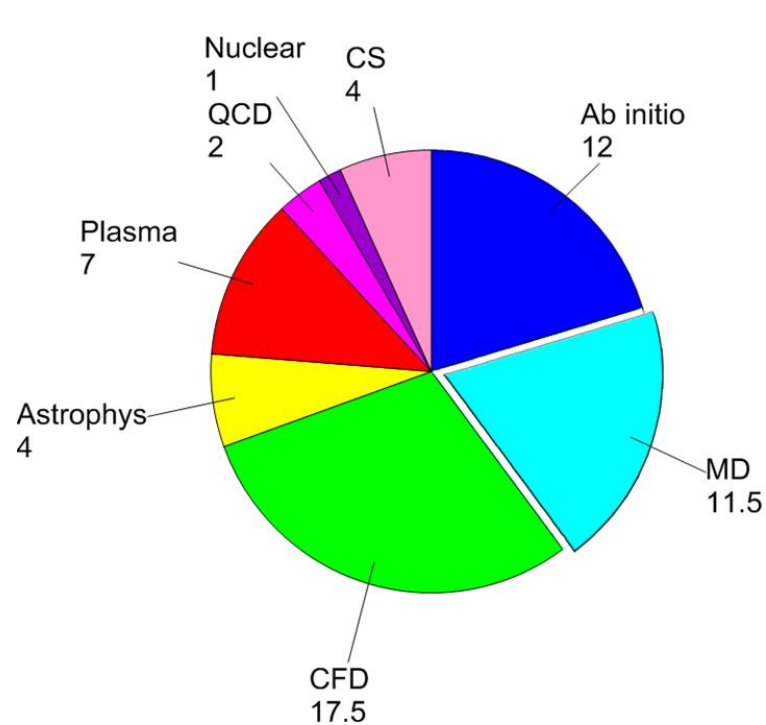
The simulation method based on numerical integration of the equations of motion for a system of a large number of particles with a given interaction potential

$$m_i \frac{d\vec{v}_i}{dt} = - \frac{\delta U(\{\vec{r}_k\})}{\delta \vec{r}_i}$$


МД – ЭТО СУПЕРКОМПЬЮТЕРНАЯ ЗАДАЧА

СОСТАВЛЯЕТ ЗНАЧИТЕЛЬНУЮ ДОЛЮ НАГРУЗКИ НА БОЛЬШИХ СУПЕРКОМПЬЮТЕРАХ

Кол-во проектов:

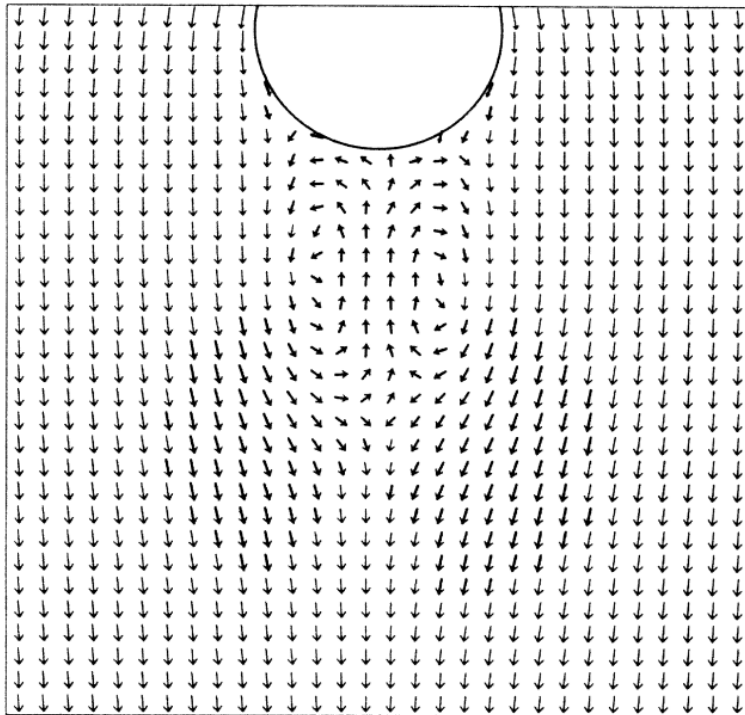


Выделенное время



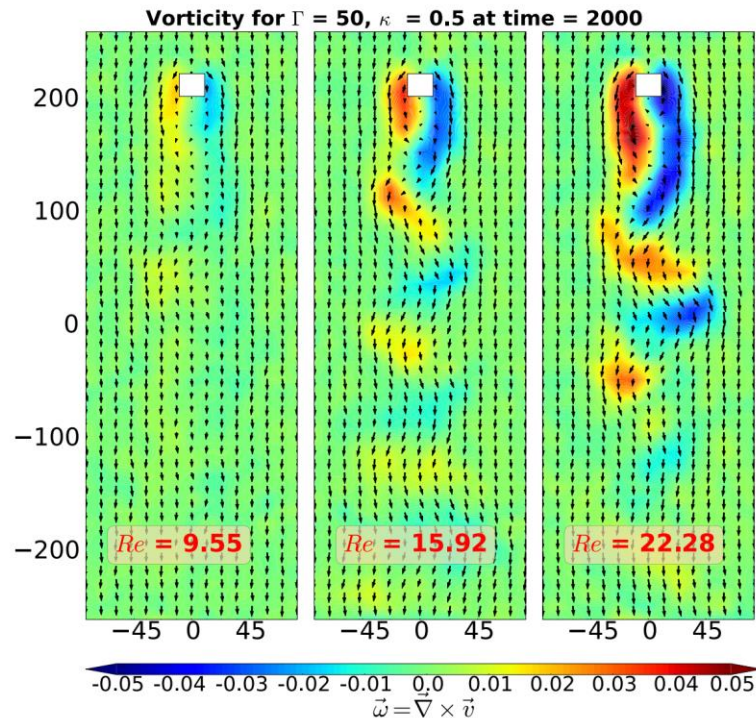
Innovate and Novel Computational Impact on Theory and Experiment (INCITE)

Molecular dynamics




Rapaport (Phys. Rev. Lett., 1986)

$N \approx 10^4$
2D



Charan&Ganesh
(Physics of Plasmas, 2016)

$N \approx 6 \cdot 10^4$
2D




Smith (Physics of Fluids, 2015)

$N \approx 3 \cdot 10^8$
3D

Hardware


	Memory	Bandwidth	Performance
AMD MI50	32 GB	1 TB/s	13.3 TFlops/sec SP, 6.6 TFlops/sec DP
Nvidia A30	24 GB	0.9 TB/s	10.3 TFlops/sec SP, 5.2 TFlops/sec DP
Nvidia V100	32 GB	0.9 GB/s	14 TFlops/sec SP, 7.1 TFlops/sec DP
Nvidia A100	80 GB	2 TB/s	19.5 TFlops/sec SP, 9.7 TFlops/sec DP



Steven J. Plimpton, Sandia

**Fast Parallel Algorithms for Short-Range
Molecular Dynamics**, S. Plimpton, J Comp Phys,
117, 1-19 (1995).

**LAMMPS - a flexible simulation tool for particle-based materials modeling at the
atomic, meso, and continuum scales**, A. P. Thompson, H. M. Aktulga, R. Berger, D. S.
Bolintineanu, W. M. Brown, P. S. Crozier, P. J. in 't Veld, A. Kohlmeyer, S. G. Moore, T. D.
Nguyen, R. Shan, M. J. Stevens, J. Tranchida, C. Trott, S. J. Plimpton, Comp Phys Comm,
271 (2022) 10817.



Extremely scalable:

Full load of the petascale Titan machine with about 16000 nodes at Oak Ridge National Laboratory

Nguyen, T.D. et al. Rupture mechanism of liquid crystal thin films realized by large-scale molecular simulations. *Nanoscale* 6(6), 3083–3096 (2014)

One of the longest simulations in the world to the 2017 at LLNL on the Sequoia Blue Gene Q machine with petascale performance


Zepeda-Ruiz et al. Probing the limits of metal plasticity with molecular dynamics simulations. *Nature* 550(7677), 492–495 (2017)

Simulation using over 100,000 processor at Los Alamos National Laboratory on the Trinity machine

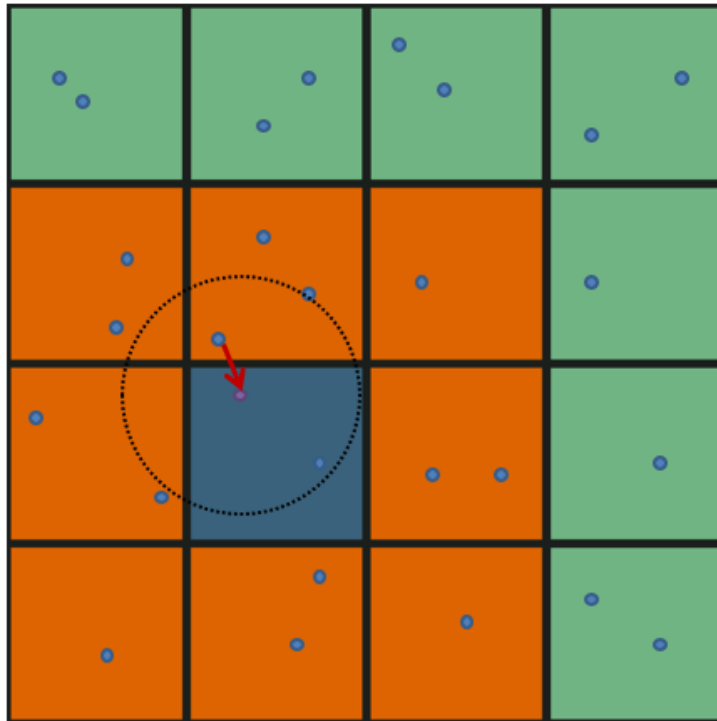
Wood et al. Multiscale modeling of shock wave localization in porous energetic material. *Physical Review B* 97(1), 014109 (2018)

Simulation using 4,650 nodes (27,900 NVIDIA V100 GPUs) at Oak Ridge National Laboratory on the Summit machine that was #1 supercomputer in 2018-2020

Nguyen-Cong et al. Billion atom molecular dynamics simulations of carbon at extreme conditions and experimental time and length scales (2021)




Domain decomposition




a east/west exchanges



b north/south exchanges




c

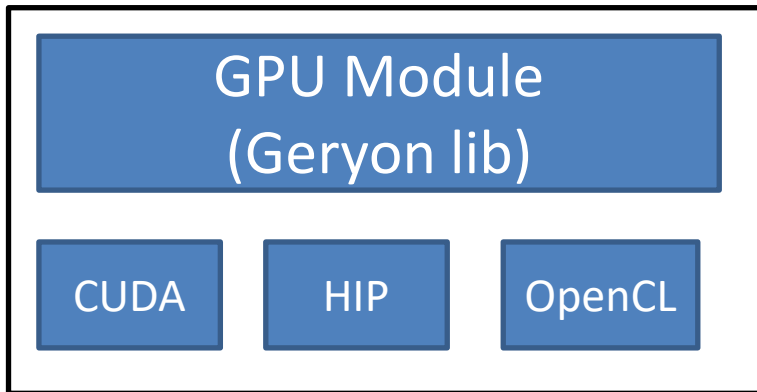


up/down exchanges

FIG. 8. Method by which a processor acquires nearby atom positions in the spatial-decomposition algorithm. In six data exchanges all atom positions in adjacent boxes in the (a) east/west, (b) north/south, and (c) up/down directions can be communicated.




Acceleration backends:



Сравнение доступности компонентов, используемых в модели воды SPC/E

LAMMPS component	GPU	KOKKOS
Neighbor Lists	calc_neigh_list_cell	NPairPotential
Pair Style	k_lj_coul_long	PairLJCutCoulLong
PPPM	interp, make_rho, particle_map	PPPM
Fix		FixShake
Integrator		FixNVEInitialIntegrate FixNVEFinalIntegrate
Compute		ComputeTemp
Bond Style		BondHarmonic
Angle Style		AngleHarmonic
Other	zero, etc.	View, Fill, Zero, etc.


GPU with 4 MPI ranks per GPU



KOKKOS with 4 MPI ranks per GPU

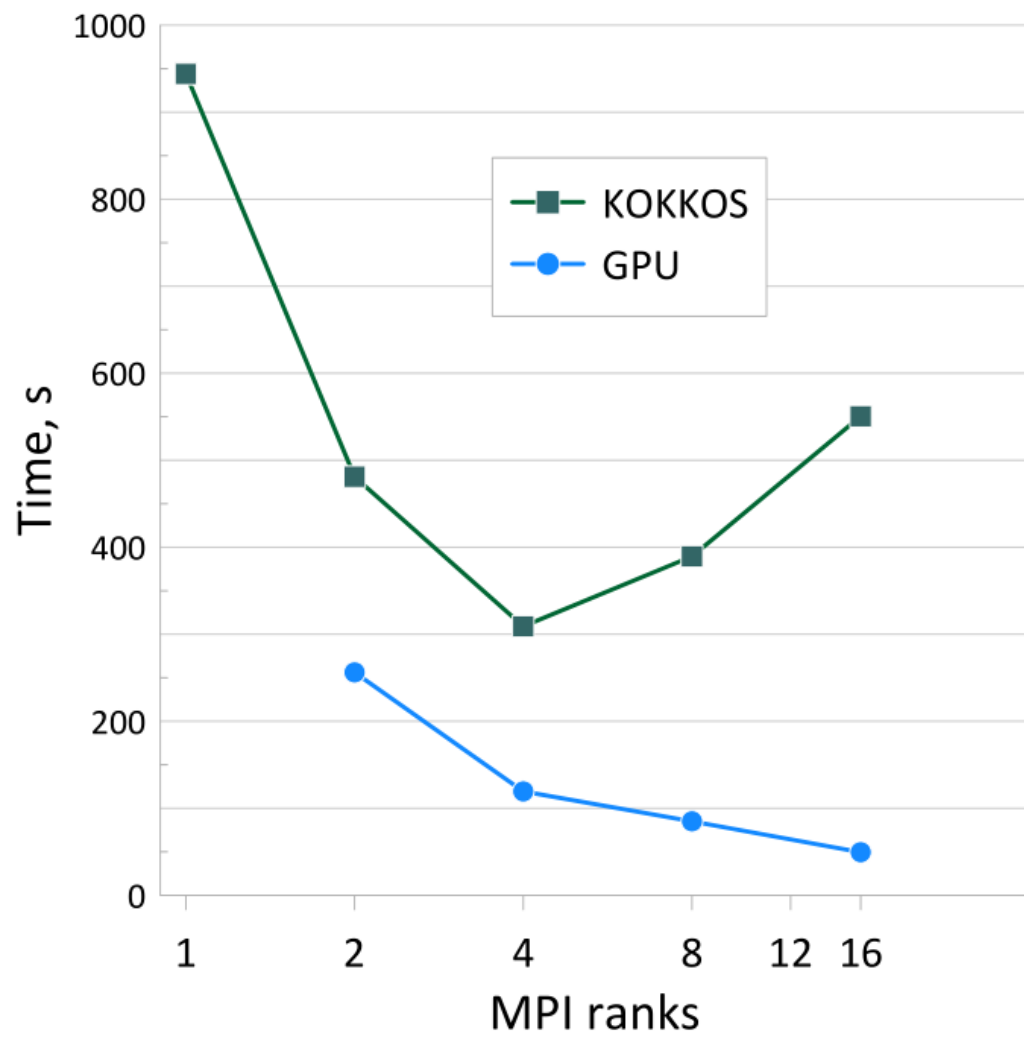



KOKKOS with 1 MPI rank per GPU



Обмены между ГПУ и основной памятью

	GPU 4 rank / GPU	Kokkos 1 rank / GPU	Kokkos 4 ranks / GPU
Invocations	3288	13759	89615
Total GB	78	29	572
Avg. MB per Invocation	23.7	2.1	6.38





OpenMM

High performance, customizable molecular simulation.



Dr. Peter Eastman, Stanford University

- High performance molecular dynamics library
- It is able to efficiently utilize GPUs
- Convenient C++ and Python APIs
- Easy to customize and extend



OpenMM

High performance, customizable molecular simulation.

Reference

CPU

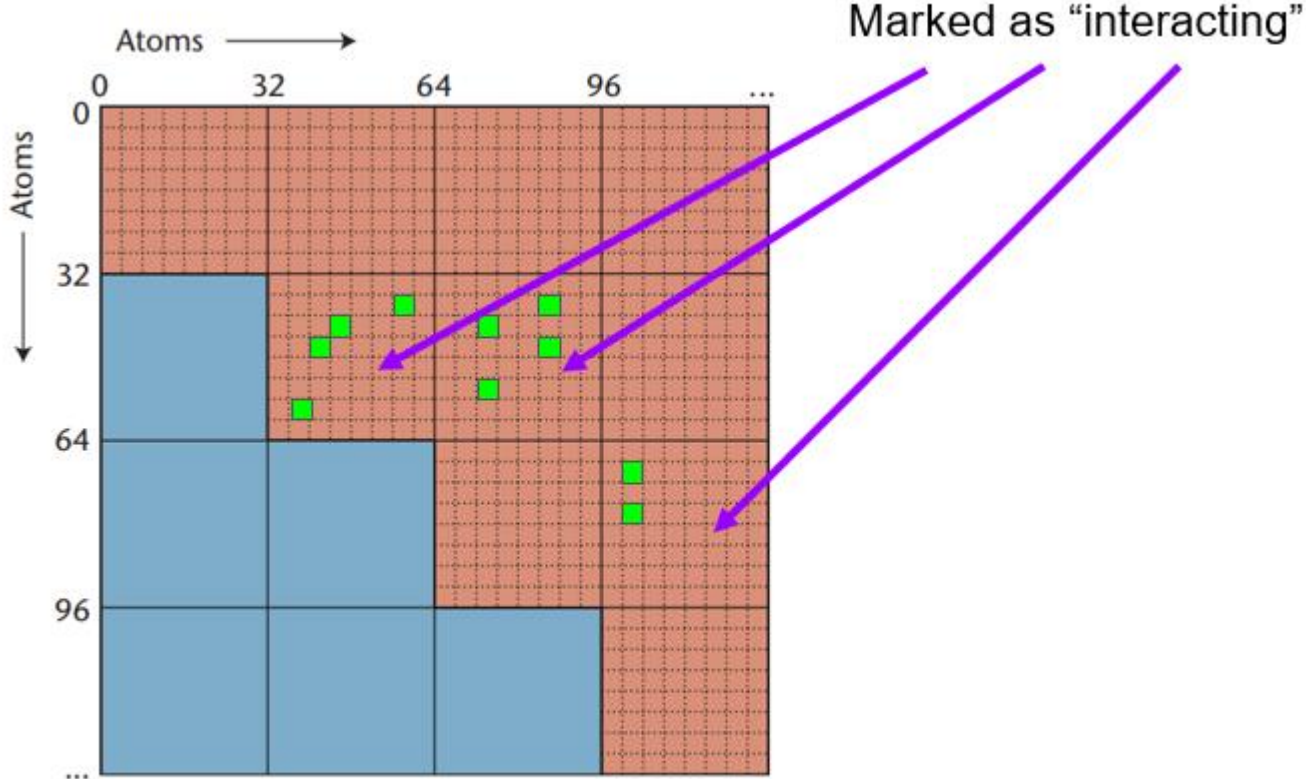
CUDA

OpenCL

HIP

Common Compute

Force decomposition



Eastman, Pande
(Computing in science and engineering, 2010)

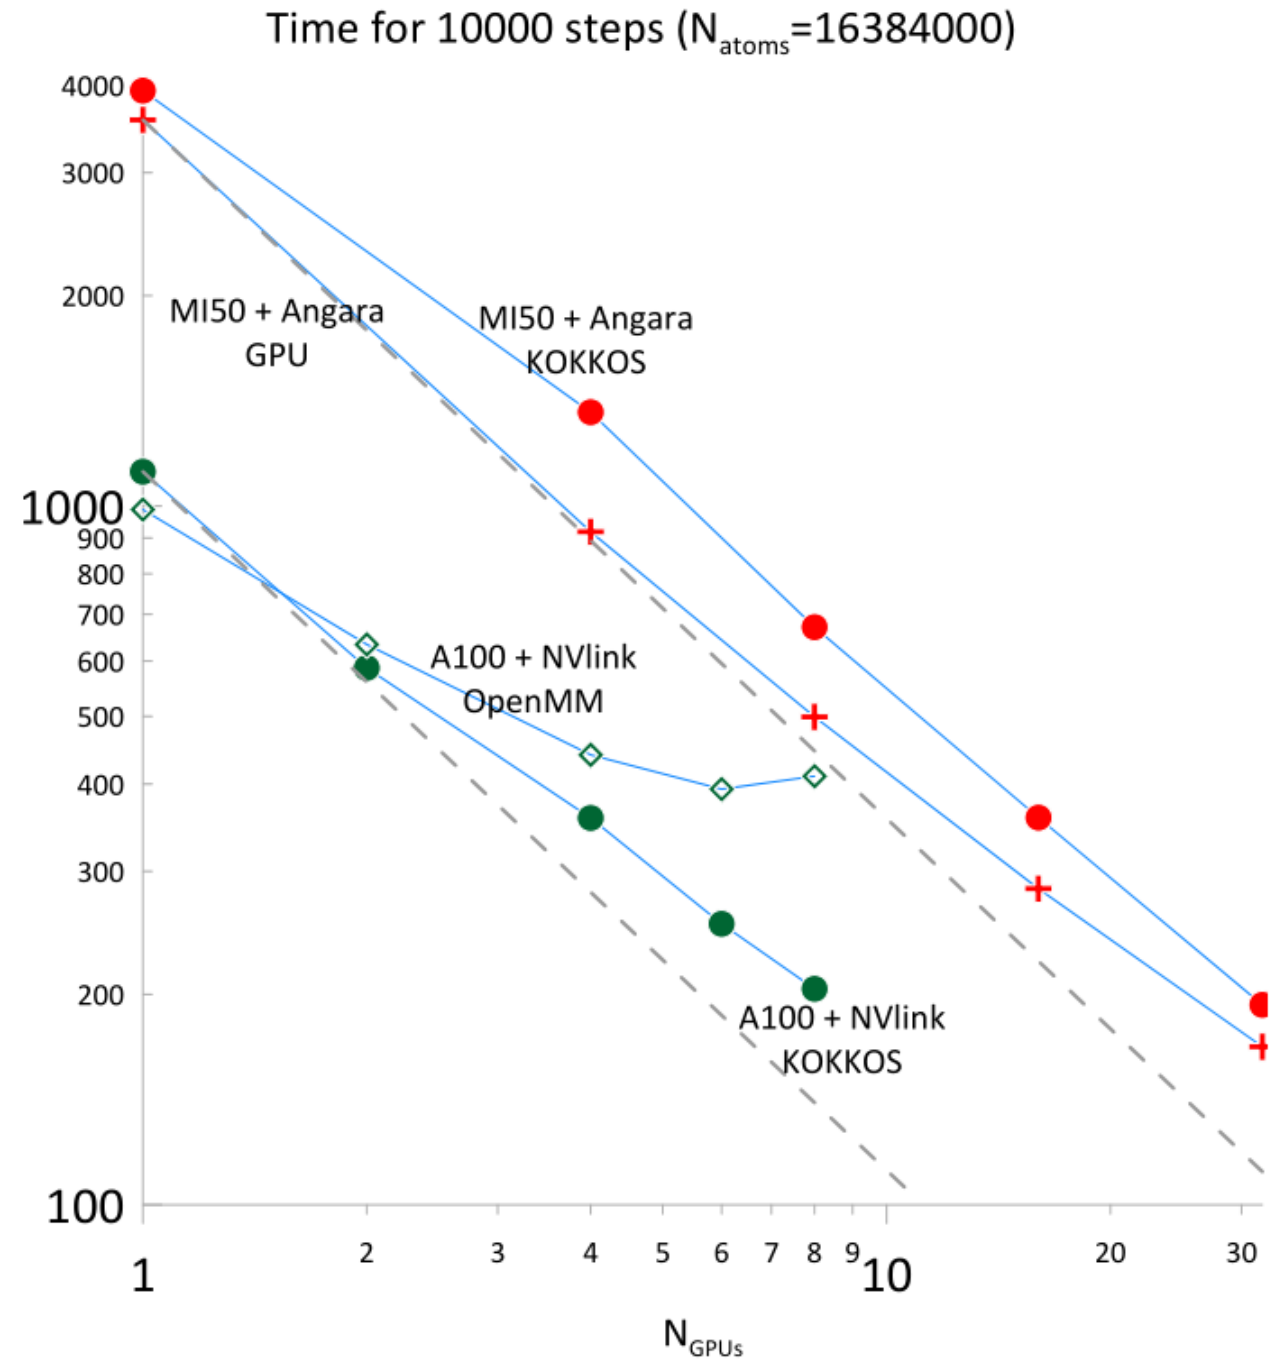
Strong scaling

Lennard-Jones liquid benchmark with 16384000 atoms (the largest size that fits into 32 GB of one AMD MI50)

Average number of neighbors is ~180




Stegailov, V., Dlinnova, E., Ismagilov, T., Khalilov, M., Kondratyuk, N., Makagon, D., ... & Timofeev, A. (2019). Angara interconnect makes GPU-based Desmos supercomputer an efficient tool for molecular dynamics calculations. *The International Journal of High Performance Computing Applications*, 33(3), 507-521.



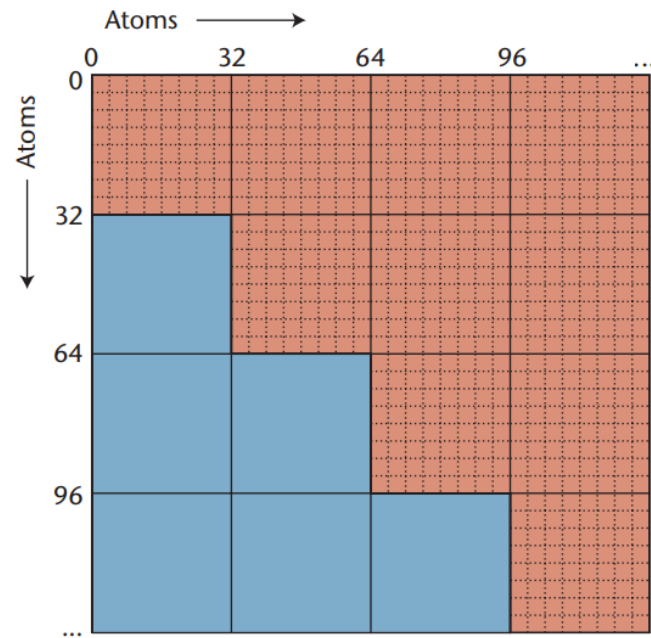
Weak scaling

Same particle density but different number of atoms

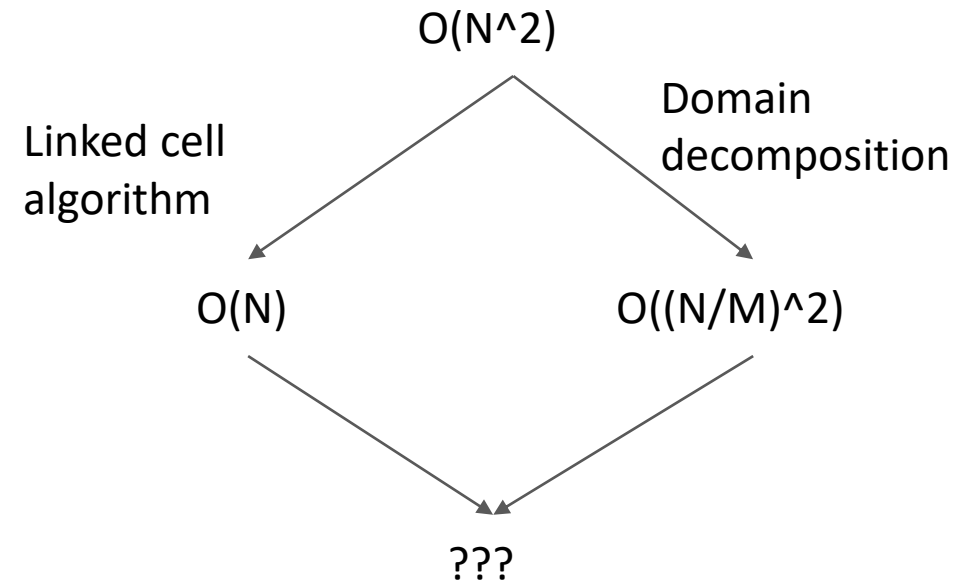
Average number of neighbors is ~ 180



OpenMM internal algorithm



Eastman, Pande
(Computing in science and engineering, 2010)



- ❖ Представлены результаты тестов LAMMPS с использованием пакетов GPU и KOKKOS. Показано отличие этих пакетов при сопоставимом общем уровне производительности.
- ❖ Количество передаваемых данных в KOKKOS существенно зависит от настроек запуска на одном и том же оборудовании. Использование нескольких процессов на один ускоритель в случае KOKKOS может существенно повышать количество передаваемых данных, в то время как для пакета GPU желательно использовать большое количество процессов на ускоритель.
- ❖ Сильное масштабирование LAMMPS при запуске на нескольких узлах с сетью Ангара не хуже, чем при использовании платформы Nvidia Delta с быстрым интерконнектом Nvlink через NVSwitch
- ❖ При близком уровне пиковой производительности ускоритель AMD MI50 показывает вдвое меньшую производительность при использовании KOKKOS на реальной задаче, чем Nvidia A30.
- ❖ OpenMM существенно превосходит по скорости счета LAMMPS на небольших системах, но его масштабируемость на несколько GPU значительно ограничена в OpenMM из-за $O(N^2)$ алгоритма.