

MORE BANG FOR YOUR BUCK!

IMPROVED USE OF GPU NODES FOR GROMACS 2018

C Kutzner,^a S Páll,^b M Fechner,^a A Esztermann,^a BL de Groot,^a H Grubmüller^a

^aTheoretical and Computational Biophysics, Max Planck Institute for Biophysical Chemistry, Göttingen
^bCenter for High Performance Computing, KTH Royal Institute of Technology, Stockholm

INTRO

The molecular dynamics (MD) simulation package GROMACS runs efficiently on a wide variety of hardware from commodity Linux workstations to HPC clusters.^{1,2} GROMACS exploits the underlying hardware on all possible levels with a combination of SIMD parallelism, OpenMP multi-threading, and MPI, whereas graphics processing units (GPUs) can be used as accelerators to "offload" calculations from the CPU.^{2,3} Here,⁴ we determine which hardware generates the maximum amount of MD trajectory for a fixed budget with GROMACS 2018.

SUMMARY OF OUR 2014 INVESTIGATION

From our earlier investigation³ using GROMACS 4.6 on 2014 hardware, we concluded that single nodes with GeForce consumer GPUs outperform CPU nodes and nodes with professional Tesla GPUs by a factor 2–3 in performance to price (P/P) ratio (Fig. 1).

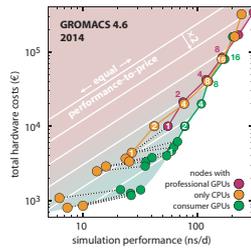


FIG. 1 Summary of our 2014 investigation.³ The performance to price (P/P) ratio for nodes with consumer GPUs (green) is 2–3x higher compared to other node types.

HARDWARE DEVELOPMENTS 2014-2018

Over the past 5 years, the nominal (FLOP-based) GPU processing power has increased more than threefold. In combination with microarchitectural improvements that made GPUs better suited for general-purpose compute, this led to an up to 6x performance increase for the GROMACS non-bonded pair interaction GPU kernels (Fig. 2). In contrast, CPU-side application performance has seen only modest gains during the same period. Whereas professional Tesla GPUs can compete with consumer GPUs in terms of performance (Fig. 2), they are lagging far behind in terms of the P/P ratio (Fig. 3).

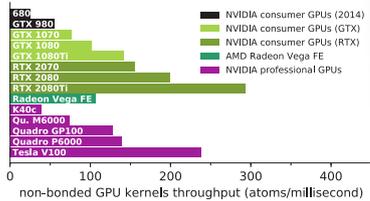


FIG. 2 Throughput of the GPU-offloaded GROMACS computation on various GPU devices.

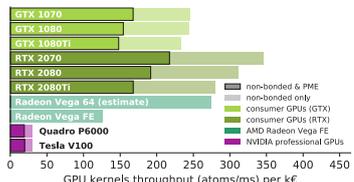


FIG. 3 Proxy metric for the P/P ratio of selected devices, computed as GROMACS GPU kernel performance divided by the average net price of that GPU in 2018.

SOFTWARE DEVELOPMENTS

The growing CPU/GPU performance imbalance motivated the implementation of tailored software features (Fig. 4): The dual pair list algorithm significantly reduces the CPU time spent for pair search (panel C) and PME offloading (panel D) enables higher simulation performances as well as significantly higher P/P ratios for cheap GPUs.

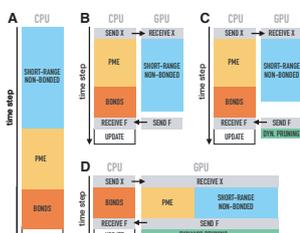


FIG. 4 Comparison of different offloading schemes employed by GROMACS 4.6, 5.x, and 2016 (B), and GROMACS 2018 (C, D).

RESULTS

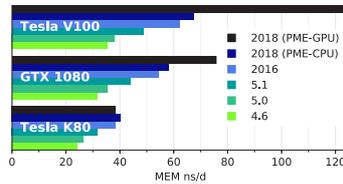


FIG. 5 GROMACS performance evolution (here on GPU nodes with 2x 8 core Xeon CPUs). Short-range non-bonded interactions are offloaded in all cases. For version 2018, also the PME mesh is offloaded (black bars).

GROMACS PERFORMANCE DEVELOPMENT

GROMACS performance on GPU nodes has increased continuously over the releases, but the most pronounced increase (given enough GPU power) is gained from PME offloading (Fig. 5). Additionally, with PME offloading, far less (4–6) CPU cores are needed to reach >80% of peak simulation performance (gained at 16 cores in Fig. 6). With a mid- to high end GPU, 10–15 "core-GHz" are generally sufficient to reach close to peak performance.

WHICH HARDWARE IS OPTIMAL FOR MD?

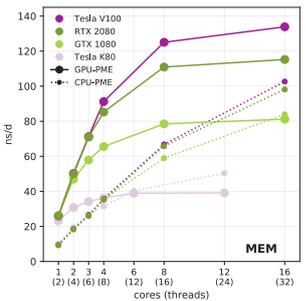


FIG. 6 GROMACS 2018 performance as a function of CPU cores used per GPU.

Fig. 7 (see left) show node costs versus GROMACS 2018 performance for selected node types. Our examined consumer GPU nodes have a P/P ratio that is 3–6x larger than their counterparts without GPUs. The best P/P ratio with current pricing is offered by combining a E5-2630v4 or Ryzen 1950X CPU with 2+ RTX 2080 GPUs, starting a ~3,000 EUR net. The best aggregate performance for consumer GPU nodes was identified for the AMD Epyc 7401P node combined with 4x 1080TI GPUs.

FREE LUNCH! GPU UPGRADES

An even greater benefit can be drawn from upgrading old nodes with state-of-the-art GPU models, as illustrated by the yellow and orange data points in Fig. 7 (see legend).

ENERGY EFFICIENCY

From the hardware costs, the measured power consumption, and the performance, we determine the total costs of producing MD trajectory on selected node types. (Fig. 8). Trajectory costs are highest on nodes without GPUs. With GROMACS 2018 on current hardware, trajectory costs on optimized GPU nodes are only 0.3x as high as on their CPU-only counterparts.

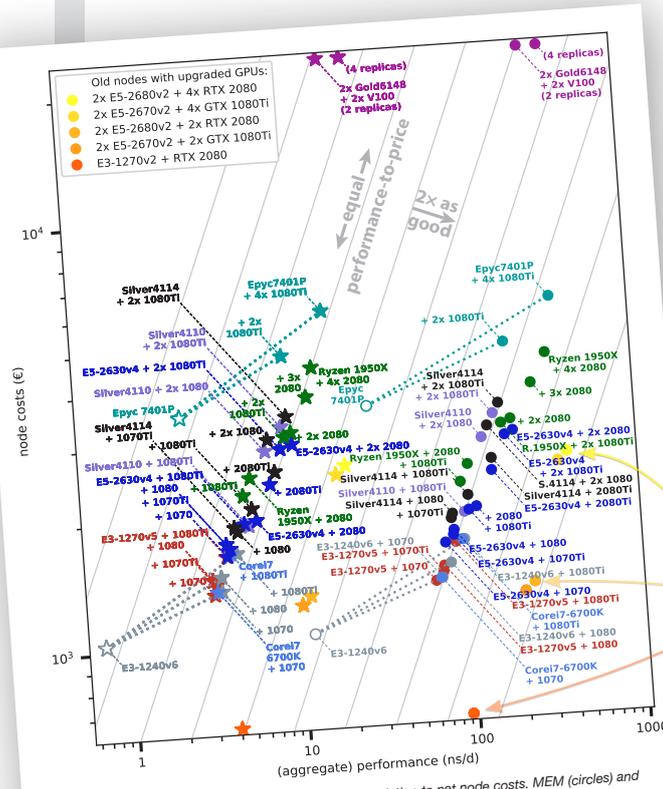


FIG. 7 (Aggregate) simulation performance in relation to net node costs. MEM (circles) and nodes (stars) symbols are colored depending on CPU type. Symbols with white fill denote nodes without GPUs; dotted lines connect GPU nodes with their CPU-only counterparts. Grey: iso-lines of equal P/P as in Fig. 1, with superior configurations to the lower right.

CONCLUSIONS

BAM! Consumer GPU nodes have 3–6x higher trajectory output per invested €, in contrast to nodes with just CPUs or with professional Tesla GPUs, when accounting the raw node price. If energy costs are included, the factor is 3.

BAM! Since GROMACS 2018, PME can be offloaded to the GPU. This feature has two ramifications:

BAM! A drastic move of the optimal hardware balance towards the GPU side: Now 10–15 "core-GHz" per GPU (1080TI or 2080) are sufficient!

BAM! A cheap upgrade of old nodes once tailored for older GROMACS versions is possible, yielding nodes that are in no way inferior as brand-new hardware!

OUTLOOK

- Our results are valid as well for GROMACS 2019.
- GROMACS 2019 allows PME offloading also for AMD GPUs.
- Benchmark your own hardware with our .tpr: <https://www.mpibpc.mpg.de/grubmueller/bench>



REFERENCES

- 1 B Hess, C Kutzner, D van der Spoel, E Lindahl. GROMACS 4: Algorithms for highly efficient, load-balanced, and scalable molecular simulation. JCTC, 4:435–447 (2008).
- 2 S Páll, M Abraham, C Kutzner, B Hess, E Lindahl. Tackling Exascale Software Challenges in Molecular Dynamics Simulations with GROMACS. EASCC 2014, LNCS 8759, pp. 3–27 (2015)
- 3 C Kutzner, S Páll, M Fechner, A Esztermann, BL de Groot, H Grubmüller. Best bang for your buck: GPU nodes for GROMACS biomolecular simulations. JCC 36 (26), pp. 1990–2008 (2015)
- 4 C Kutzner, S Páll, M Fechner, A Esztermann, BL de Groot, H Grubmüller. More bang for your buck: Improved use of GPU nodes for GROMACS 2018. Manuscript submitted to JCC (2019)

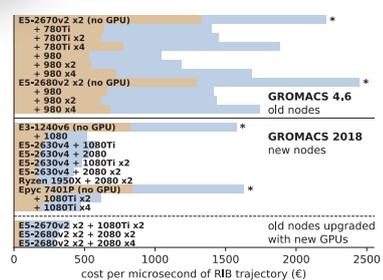


FIG. 8 RIB trajectory costs assuming five years of operation, including costs of 0.2 EUR per kWh for energy and cooling. Nodes without GPUs (*) are most expensive in terms of trajectory production.

METHODS

We use the same two representative benchmark systems as in our previous study³ MEM is a 81k atoms MD system, a membrane protein embedded in a lipid bilayer surrounded by water and ions (2 fs time step, 1 nm cutoffs, 0.12 nm PME grid spacing). RIB is a ribosome solvated in water and ions comprising altogether 2.1M atoms (4 fs time step, 1 nm cutoffs, 0.135 nm PME grid spacing). On each node, the highest GROMACS 2018 performance is determined. On nodes with N GPUs, we determine the aggregate performance of running N simulations in parallel using the -multidir functionality.

