INTERNATIONAL JOURNAL OF ENGINEERING
AND APPLIED SCIENCE
ISSN 2709-8117
DEC. 2021 VOL. 2, NO. 4

Comparative Performance Analysis of Molecular Dynamics Software, Taking the Popular Gromacs Package as an Example

Danil C. Fedorov, Vladimir G. Kholina

Federal Scientific and Clinical Center of Specialized Types of Medical Care and Medical Technologies, Federal Medical-Biological Agency of Russia, Moscow, Russia

ABSTRACT

Molecular dynamics has proved itself as a powerful computer simulation method to study dynamics, conformational changes, and interactions of biological macromolecules and their complexes. In order to achieve the best performance and efficiency, it is crucial to benchmark various hardware platforms for the simulations of realistic biomolecular systems with different size and timescale. Here, we compare performance and scalability of a number of commercially available computing architectures using all-atom and coarse-grained molecular dynamics simulations of water and the Ndc80-microtubule protein complex in the GROMACS-2019.4 package. We report typical single-node performance of various combinations of modern CPUs and GPUs, as well as multiple-node performance of the "Lomonosov-2" supercomputer. These data can be used as the practical guidelines for choosing optimal hardware for molecular dynamics simulations.

KEYWORDS

Molecular dynamics; Coarse grain; Tubulin; Microtubule; Ndc80.

1. Introduction

Over the last decades, molecular dynamics (MD) simulations have become a powerful tool for investigating molecular systems (including protein assemblies) with the exceptionally high temporal and spatial resolutions unattainable so far using experimental techniques. Molecular systems of biological interest typically consist of up to millions of atoms, so MD simulations of biomacromolecules still represent a major computational challenge. The GROMACS package [1] is among the most efficient and popular engines for MD simulations as it runs efficiently on a wide variety of hardware from desktop workstations to supercomputers. Here, we evaluate which hardware combinations show optimal performance. This work continues the systematic comparison of the multiple currently available hardware architectures in terms of their MD simulation performance, which has been initiated in [3] for a number of biologically relevant molecular systems. In the present study, we use two types of molecular systems as computing benchmarks: (i) water boxes (WB) of different size and (ii) a biomolecular system consisting of the Ndc80 protein complex with a microtubule (MT)

fragment [4]. Moreover, we extend our benchmark to coarse-grained (CG) MD simulations using MARTINI force field [7], a popular CG model for biomolecular simulations [8], using the same protein complex for testing. The MARTINI model is based on a four-to-one mapping, i.e., on average four heavy atoms plus associated hydrogens are represented by a single interaction center (called a bead). The overall aim of the coarse-graining approach is to provide a simple model that is computationally fast and easy to use, yet flexible enough to be applicable to a large range of biomolecular systems.

The paper is organized as follows. In Section 1, we introduce the biomolecular systems used for benchmarking and describe the simulation setups employed for all-atom and coarse-grained simulations. In Section 2, we provide a comprehensive overview of performance achieved on various computational platforms. Finally, in Conclusion, we summarize the results and outline possible directions for further work.

2. Methods

All calculations were performed using the GROMACS-2019.4 version, which allows parallel computing on hybrid architectures. All benchmarks were run for 15 minutes. All-atom (AA) simulations were run in the explicit solvent using the TIP3P water model and the CHARMM27 force field for proteins. The production simulation runs were carried out in the NPT ensemble at 300K, using the Parrinello-Rahman algorithm [10] and the V-rescale thermostat for a duration of 1μs each.

The structure of Ndc80 in complex with a MT fragment was obtained from the Protein Data Bank (PDB id 2VE7). The size of the virtual cell was chosen in such a way that the distance from the protein surface to the nearest box boundary was not less than two nanometers. For AA simulations the particle mesh Ewald (PME) method was used for the long-range electrostatics. Here, we used the interpolation order of 4 for PME, which equals cubic interpolation and should give electrostatic energies accurate to about 5· 10−3. The mass rescaling approach (i.e., partial transfer of mass from heavy atoms to the hydrogens bound to them) [5] allowed us to use 4-fs time step for AA MD simulations of the Ndc80 system instead of 1-fs time step used for the WB simulations and thus to accelerate them. Further details about the utilized MD protocol can be found in [4].

Table 1. Molecular dynamics systems used in the benchmark. Water box size is shortcutted in the system name, where the number stands for the thousands of particles in the system. Ndc80 is an acronym of a kinetochore protein

MD systems	MD system name	Number of particles	Box type	System size (nm)	Time step (fs)
Water box (WB)	WB-10	10206	cube	4.7x 4.7x 4.7	1
	WB-80	80232	cube	9.3x9.3x9.3	1
	WB-120	121527	cube	10.7x10.7x10.7	1
	WB-160	159780	cube	11.7x11.7x11.7	1
	WB-200	203415	cube	12.7x12.7x12.7	1
Ndc80 complex with microtubule	Ndc80 AA	750295	cube	22.1x17.2x20.1	4
	Ndc80 CG W	133371	cube	27x22x27	20
	Ndc80 CG PW	386294	cube	27x22x27	20

Coarse-grained simulations were run using the most recent version 2 of the MARTINI force field and the yet unreleased version 3 of this force field. The simulations with MARTINI 2 were run in combination with the polarizable water model (PW) [12], which allows for proper screening of interactions and other polarization effects. For these simulations, we utilized either

The PME or the reaction-field (RF) approach for the long-range electrostatics (with the cut-off rc = 1.1 nm, the dielectric constant beyond the cut-off was set to infinity [2]). The MARTINI beta version 3 is currently lacking an appropriately parameterized polarizable water model, so the simulations using this version of the CG force field were run with the standard water model (W) corresponding simply to a van der Waals particle. The RF approach was used for the longrange electrostatics in this case. The detailed protocols can be found elsewhere [6]. Specifications of all MD systems used for benchmarking are summarized in Tab. 1.

3. Results and Discussion

In order to examine the performance of MD simulations as a function of particle count both in homogeneous systems and in realistic systems, we have carried out MD simulations for a series of water boxes of the increasing size (see Tab. 1) and for the Ndc80-MT complex in the explicit solvent. In contrast to the previously reported benchmark [3], the updated results suggest that an increase of the system size leads to a commensurable decrease of computer performance in the explored range of system size (10,000-200,000 atoms), i.e., 20-fold increase in the number of atoms results in approximately equal decrease of performance for the CPU-only architecture. For GPU-accelerated simulations, we have found out even slower decrease of performance, which, for instance, scales down by the factor of 12.5 with the 20-fold increase of the system size for the RTX3080/Intel Core i9-9940X combination, see Tab. 2. However, for the largest system, which we have tested in the present benchmark, the all-atom Ndc80-MT complex consisting of over 750,000 atoms, the performance drop becomes disproportional. Overall, it implies an extremely high potential of the single-node hybrid architectures for the simulations of molecular systems with up to 100,000‑200,000 atoms particularly emanating from the recent adaptations of the MD software for such platforms [9]

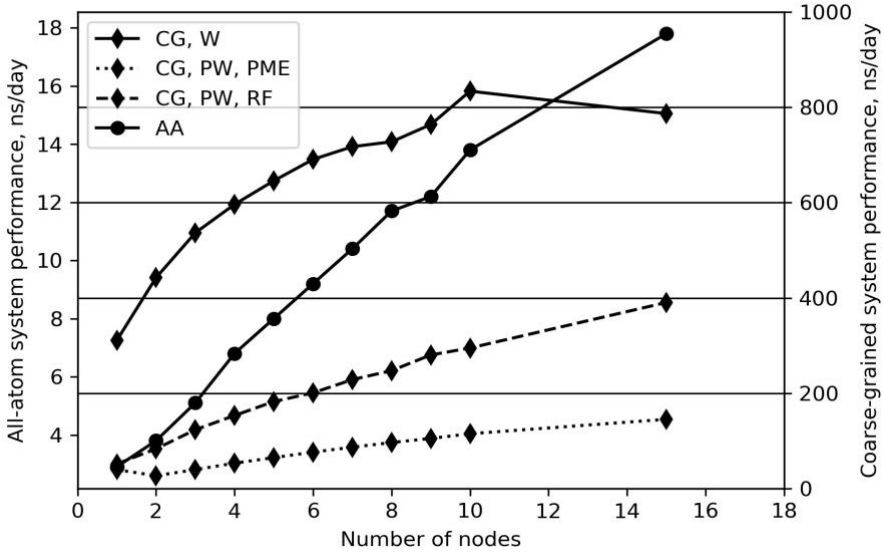

Figure 1. Performance of "Lomonosov-2" supercomputer (ns/day), depending on the number of computing nodes for all-atom (AA) and coarse-grained (CG) MD simulations of the Ndc80 complex with a fragment of MT

We have also addressed the question of scalability of MD simulations in GROMACS by estimating the dependence of the "Lomonosov-2" supercomputer [11] performance on the number of computer nodes used. For the all-atom Ndc80 system, as well as for the two CG systems with polarizable water (PW), the performance grows almost linearly as a function of the number of supercomputer nodes (Fig. 1) roughly following Amdahls law. However, for the smallest coarsegrained system (with the standard water model, W), the performance reaches the plateau at 10 nodes.

Finally, we have assessed the performance of two alternative schemes commonly used to treat the long-range electrostatics in CG simulations: PME and RF. For all the tested systems, RF outperforms PME suggesting the former as the best option. However, the difference is the strongest for the no-GPU platforms where it can be as large as 7-fold. For the GPU-accelerated platforms, a much lower handicap of 1.2-3.5 is observed

Table 2. Single-node performance (ns/day) for systems of different size depending on various combinations of CPUs and GPUs

GPU	MD system								
	WB-10	WB-80	WB-120	WB-160	WB-200	Ndc80 AA	Ndc80 CG W	Ndc80 CG PW PME	Ndc80 CG PW RF
no GPU*	100.3	14.1	9.0	6.6	5.2	4.2	910.8	25.9	153.8
RTX 2070 Super*	290.8	48.5	32.3	25.8	20.9	15.1	1657.5	152.2	260.9
RTX 2080 Ti*	349.9	62.6	43.8	34.5	26.8	19.5	1750.2	187.5	286.6
RTX 3080*	365.8	67.9	47.5	37.3	29.4	23.5	1817.6	252.4	330.4
no GPU**	113.6	16.5	10.5	8.1	6.3	5.7	1127.8	24.4	181.4
RTX 2080 Ti**	338.5	45.5	29.1	22.2	16.5	14.3	1993.1	207.7	331.4
2 RTX 2080 Ti**	207.4	49.6	30.0	22.6	17.5	14.7	2232.4	88.2	279.3

* Nodes with the Intel Core i9-7900X CPU.
** Nodes with the Intel Core i9-9940X CPU.

Our comparative performance analysis of molecular dynamics software (using the popular GROMACS package as an example) provides the guidelines for selection of the best-performing GPU-based architectures for both all-atom and coarse-grained MD simulations of realistic molecular systems. It also outlines certain limitations of the single-node workstations and highlights the importance of the HPC platforms (e.g., "Lomonosov-2" supercomputer) for the simulations of large systems exceeding ca. 200,000 particles on the relevant time scale.

In the case of the AA MD simulations, the presence of a modern graphics accelerator speeds up the calculations by about 5 times both for a 100 thousand atom system, and a million atom system. Moreover, as the size of the system grows, the acceleration increases slightly. However, in the case of CG models, the increase in performance with a modern GPU is not so impressive—only about 2 times. In this case, the multicore and multiprocessor computing architectures are also very important.

When using a parallel supercomputer 15 nodes (or more) are optimal for both AA and CG calculations. However, for CG systems with non-polarizable water, 10 nodes is the optimal choice, since with a further increase in the number of nodes, performance begins to decline.

Update on Performance Analysis of Different Computational Architectures: Molecular...

4. Conclusion

Our comparative performance analysis of molecular dynamics software (using the popular GROMACS package as an example) provides the guidelines for selection of the best-performing GPU-based architectures for both all-atom and coarse-grained MD simulations of realistic molecular systems. It also outlines certain limitations of the single-node workstations and highlights the importance of the HPC platforms (e.g., "Lomonosov-2" supercomputer) for the simulations of large systems exceeding ca. 200,000 particles on the relevant time scale.

In the case of the AA MD simulations, the presence of a modern graphics accelerator speeds up the calculations by about 5 times both for a 100 thousand atom system, and a million atom system. Moreover, as the size of the system grows, the acceleration increases slightly. However, in the case of CG models, the increase in performance with a modern GPU is

not so impressive—only about 2 times. In this case, the multicore and multiprocessor computing architectures are also very important.

When using a parallel supercomputer 15 nodes (or more) are optimal for both AA and CG calculations. However, for CG systems with non-polarizable water, 10 nodes is the optimal choice, since with a further increase in the number of nodes, performance begins to decline.

Update on Performance Analysis of Different Computational Architectures: Molecular...

Number equations consecutively with equation numbers in parentheses flush with the right margin, as in (1). First use the equation editor to create the equation. Then select the "Equation" markup style. Press the tab key and write the equation number in parentheses.

References

Abraham, M. J., Murtola, T., Schulz, R., et al.: GROMACS: "High performance molecular simulations through multi-level parallelism from laptops to supercomputers." SoftwareX, pp. 1-2, 19-25, 2015. DOI: 10.1016/j.softx.2015.06.001

Abraham, M., van der Spoel, D., Lindahl, E., et al.: "The GROMACS development team GROMACS User Manual," version 2019. http://www.gromacs.org, 2019.

Fedorov, V.A., Kholina, E.G., Kovalenko, I.B., et al.: "Performance analysis of different computational architectures: Molecular dynamics in application to protein assemblies, illustrated by microtubule and electron transfer proteins." Supercomputing Frontiers and Innovations, 5(4), pp. 111-114, 2018. DOI: 10.14529/jsfi180414

Fedorov, V.A., Orekhov, P.S., Kholina, E.G., et al.: "Mechanical properties of tubulin intraand inter-dimer interfaces and their implications for microtubule dynamic instability." PLoS computational biology, 15(8), e1007327, 2019. DOI: 10.1371/journal.pcbi.1007327

Feenstra, K.A., Hess, B., Berendsen, H.J.C.: "Improving efficiency of large time-scale molecular dynamics simulations of hydrogen-rich systems." Journal of Computational Chemistry, 20(8), pp. 786‑798, 1999. DOI: 10.1002/(SICI)1096-987X(199906)20:8¡786::AIDJCC5¿3.0.CO;2-B

Kholina, E.G., Kovalenko, I.B., Bozdaganyan, M.E., et al.: "Cationic antiseptics facilitate pore formation in model bacterial membranes." The Journal of Physical Chemistry, B, 124(39), pp. 8593-8600, 2020. DOI: 10.1021/acs.jpcb.0c07212

Marrink, S.J., Risselada, H.J., Yefimov, S., et al.: "The MARTINI force field: coarse grained model for biomolecular simulations." The journal of physical chemistry, B, 111(27), pp. 7812-7824, 2007. DOI: 10.1021/jp071097f

Monticelli, L., Kandasamy, S.K., Periole, X., et al.: "The MARTINI coarse-grained force field: extension to proteins." Journal of chemical theory and computation, 4(5), pp. 819-834, 2008. DOI: 10.1021/ct700324x

P ´ all, S., Zhmurov, A., Bauer, P., et al.: "Heterogeneous parallelization and acceleration of molecular dynamics simulations in GROMACS." The Journal of Chemical Physics, 153(13), pp. 134110, 2020. DOI: 10.1063/5.0018516

Parrinello, M., Rahman, A.: "Polymorphic transitions in single crystals: A new molecular dynamics method." Journal of Applied physics, 52(12), pp. 7182-7190, 1981. DOI: 10.1063/1.328693

Voevodin, V.V., Antonov, A.S., Nikitenko, D.A., et al.: "Supercomputer Lomonosov-2: large scale, deep monitoring and fine analytics for the user community." Supercomputing Frontiers and Innovations, 6(2), pp. 4-11, 2019. DOI: 10.14529/jsfi190201

Yesylevskyy, S.O., Sch ¨ afer, L.V., Sengupta, D., et al.: "Polarizable water model for the coarsegrained MARTINI force field." PLoS Comput Biol, 6(6), pp. e1000810, 2010. DOI: 10.1371/journal.pcbi.1000810