DOI: 10.14529/jsfi260102

High-Performance Computing in the Molecular Dynamics of Tubulin Cytoskeleton Polymers

Ilya B. Kovalenko[1], *Vladimir A. Fedorov*[1], *Ekaterina P. Vasyuchenko*[1], *Ekaterina G. Kholina*[1], *Svetlana Yu. Kovalenko*[1], *Andrey B. Rubin*[1]

High-performance computing is one of the most essential tools fueling the advancement of computational biology. The article discusses the application of the full-atom molecular dynamics (MD) method to study the dynamic behavior of filaments formed by the protein tubulin, and presents the results of testing the calculation performance depending on the latest models of central processors and video accelerators. Our comparative performance analysis of GPU-based computing architectures for all-atom MD simulations of biomolecular systems not only provides guidance on choosing the best computing solution in terms of price-performance ratio, but also shows the maximum potential computational performance that modern CPUs and GPUs can provide. For example, MD of the biomolecular system containing a tubulin protofilament in an explicitly specified solvent consisting of more than 300 thousand atoms can be studied with performance of 232 ns/day at time step 2 fs when using single-node computer with the latest CPU and GPU generation architecture. Constantly evolving computing resources coupled with modern software enable us to solve increasingly complex problems in life sciences.

Keywords: molecular dynamics, tubulin, microtubule, CPU, GPU, computing performance.

Introduction

Tubulins and tubulin-like proteins form cellular structures that play a crucial role in numerous cellular processes, including cell division. For example, tubulins form microtubules, which are involved in chromosome search and separation. This is possible due to the unique property of microtubules to spontaneously polymerize and depolymerize, a property known as dynamic instability [6]. Another protein, FtsZ (Filamenting temperature-sensitive mutant Z), is a tubulin-like protein that forms filaments with a repeating arrangement of subunits. These filaments form a ring (so-called Z-ring) around the longitudinal midpoint, or septum, of a bacterial cell [9]. FtsZ is essential for cell division in almost all bacteria and in many but not all archaea [8].

Although eukaryotic microtubules are a well-studied subject, the specific molecular mechanisms underlying and governing their dynamic instability remain unclear. One of the questions concerning microtubules is whether individual microtubule protofilaments assume a straight or curved shape in solution. We know even less about the specific molecular mechanisms underlying the dynamic behavior of FtsZ filaments that form the Z-ring. In any case, it is known that FtsZ protofilaments have polarity and move in one direction by treadmilling [9].

To study the dynamic behavior of both individual protofilaments and entire microtubules formed by tubulin, as well as filaments of the FtsZ protein, it is convenient to use methods of molecular computer modeling, in particular, the classical molecular dynamics (MD) method. However, studying such large molecular systems as microtubules and even individual tubulin protofilaments or FtsZ filaments requires extremely high computational resources. The situation is further complicated by the fact that both tubulins and FtsZ possess long, unstructured C-terminal regions or tails (in the case of FtsZ, this unstructured tail can reach a hundred amino acid residues in length), the role of which in dynamic behavior remains unclear. Because these

[1]Lomonosov Moscow State University, Moscow, Russia

regions lack a defined structure and their dynamic behavior is highly variable, the building of a molecular model requires increasing the reaction volume and obtaining long MD trajectories.

High-performance computing (HPC) is one of the most essential tools that fuels the advancement of computational biology. Since our latest overview of HPC capabilities, benchmarked on simulations of tubulin dynamics in 2023 [4], new version of CPU and GPU became commercially available. This article discusses the issue of choosing the optimal computational architecture for studying large biological systems using the MD method and the gains in computing speed that are provided by the latest models of central and graphic processors. Since 2018, when NVIDIA released the first RTX series GPUs, four generations of these accelerators have been released, and each new generation increased the speed of calculations. New, more powerful central processing units (CPUs) have also been designed by Intel, and although their contribution to the increase in MD calculation performance is less than that of GPUs, we also discuss them in this article, including Intel Ultra series CPU.

The article is organized as follows. Section 1 is devoted to the computational methods used in the paper. In Section 2 we provide the results of molecular dynamics performance tests using different computer architectures. The conclusion summarizes the study.

1. Methods

All tests were performed using the all-atom explicit solvent MD. The calculations were executed using the GROMACS 2022.5 or GROMACS 2025.2 software package [2] which facilitates parallel computing on hybrid architectures and incorporates the CHARMM27 force field [10]. Each benchmark simulation was run for a duration of 30 minutes, employing the TIP3P water model. The tubulin tetramer structure was obtained from the Protein Data Bank (PDB id 5SYF [7]). The dimensions of the simulation volume were selected to ensure that the distance from the protein's surface to the nearest boundary of the simulation box was no less than two nanometers. Long-range electrostatic interactions were taken into account using the particle mesh Ewald method [3]. Both Coulomb and Lennard-Jones cutoffs were configured to 1.25 nm. Molecular dynamics (md) integrator was used. For the water box systems, the time step was 1 fs and no restraints were used. For the tubulin protofilament, the time step was 2 fs and constraints were imposed on the bonds of atoms with hydrogens. The "GPU-resident" option was tested, which updates coordinates on the GPU when all force and coordinate data remain resident on the GPU for a number of steps [1]. Specifications of MD systems used for benchmarking are summarized in Tab. 1.

Table 1. Specification of molecular dynamics systems used in the benchmarks

MD systems	MD system name	Number of atoms	System size, nm
Water box (WB)	WB-10	10206	$4.7 \times 4.7 \times 4.7$
	WB-80	80232	$9.3 \times 9.3 \times 9.3$
	WB-200	203415	$12.7 \times 12.7 \times 12.7$
	WB-500	500076	$17.2 \times 17.2 \times 17.2$
	WB-1000	1000005	$21.7 \times 21.7 \times 21.7$
Tubulin tetramer	Tub-4	307453	$11.3 \times 12.5 \times 22.2$

2. Results and Discussion

2.1. Performance Tests

We have tested the performance of all-atom MD calculations for water box systems of different sizes as well as tubulin tetramer in water for both considered Intel CPUs (Core i9-14900KF and Core Ultra 9 285K) and a variety of NVIDIA RTX graphic accelerators. The table with all performance tests is available via this link. The results of the performance tests are also presented selectively in Fig. 1 and Tables 2–4.

2.2. Dependence of Performance on System Size

To understand how the performance of MD calculations depends on the size of the molecular system, we tested five water box systems of different sizes (10, 80, 200, 500, and 1000 thousands of atoms) for both considered Intel CPUs with and without NVIDIA RTX5090 graphic accelerator. We normalized the performance value to the system size and the performance of the Core i9-14900KF CPU only, Fig. 1. The graph shows that performance normalized to the number of atoms is virtually independent of system size when calculations are performed solely on the CPU, without a GPU. However, when using the top-of-the-line NVIDIA RTX5090 GPU for calculations, the situation changes greatly. First, performance with a GPU is significantly higher than without it, reaching its maximum (more than 20 times faster than with a CPU only) on large molecular systems (half a million and a million atoms). Second, on smaller molecular systems, normalized performance drops significantly. For example, for a system size of 80000 atoms, it drops by a factor of 2 on an Intel Ultra 9 285K processor, and by approximately 20% on a Core i9-14900KF processor. These results will be discussed in more detail in the following sections.

The cause of the relatively poor normalized performance for a small system (less than a hundred thousand atoms) may be the ineffective use of GPU computing resources. However, small systems are computationally fast, the computing performance of a system of 10000 atoms reaches 1 μs/day. In addition, biological systems of interest often contain a much larger number of atoms.

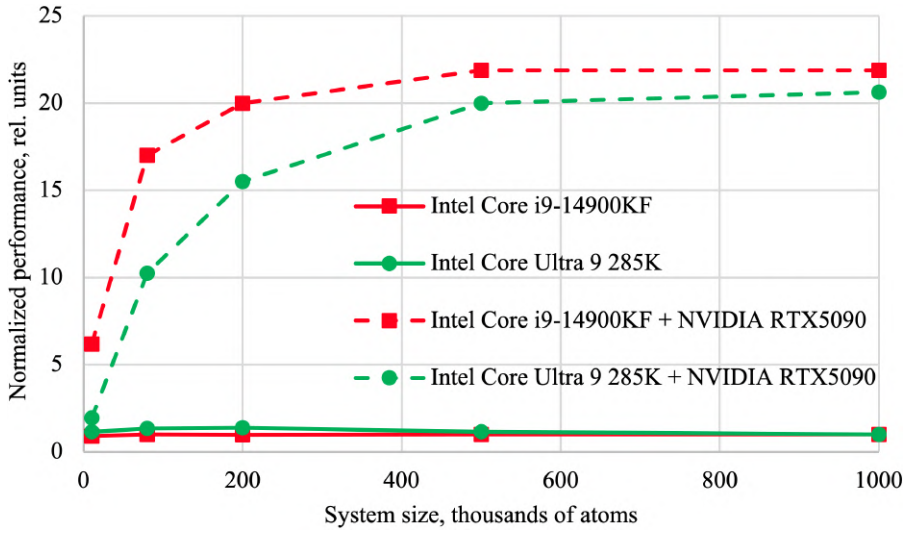

Figure 1. Normalized performance of MD calculation for systems of different sizes both for CPU only and CPU+GPU computer systems

2.3. Comparison of the Latest Intel Ultra 9 285K and i9-14900KF Processors in CPU-only Mode and in Combination with the Latest NVIDIA RTX5090 Graphics Accelerator

If the GPU is not used for molecular dynamics calculations, but only the CPU, the latest Intel Ultra 9 285K processor demonstrates almost 1.5 times higher performance than the Intel i9-14900KF on small systems of approximately 10000 atoms (Tab. 2, the third and fourth columns). However, the larger the molecular system, the smaller the performance gain. For a system consisting of a million atoms, no gain is observed, and these two CPUs demonstrate identical performance. It should be noted that the two versions of Gromacs compared, using only the CPU, produce virtually identical results, especially for the Intel i9-14900KF processor.

For large molecular systems with 200000 or more atoms, the latest NVIDIA RTX5090 graphics accelerator provides a 20-fold increase in performance when running MD simulations on the latest version of Gromacs 2025.2, or when using version 2022.5 with the GPU-resident option (Tab. 2). This observation roughly corresponds to data obtained three years ago (2022) on a previous-generation computer with an Intel i9-13900K processor and an NVIDIA RTX4090 graphics accelerator: back then, the graphics accelerator increased the computing speed by 14–22 times, depending on the system size [4]. Of course, a 2025 computer system with the graphics accelerator is about 1.3 times faster than the previous one anyway. It is worth adding that the new Intel i9-14900KF CPU delivers slightly faster, more stable and predictable GPU-based MD performance than the latest Intel Ultra 9 285K, especially for large molecular systems (Tab. 2).

Table 2. Single-node performance (ns/day) for water box systems of different size depending on various combinations of CPUs and RTX5090 GPU and version of Gromacs MD software

Molecular system	Intel Core CPU model	2022.5 No GPU	2025.2 No GPU	2022.5 GPU	2022.5 GPU-resident	2025.2 GPU	2025.2 GPU-resident
WB-10	i9-14900KF	142	143	768	1053	990	1016
WB-80	i9-14900KF	20	20	197	365	340	350
WB-200	i9-14900KF	7.8	7.8	85	160	160	160
WB-500	i9-14900KF	3.2	3.2	35	69	70	70
WB-1000	i9-14900KF	1.6	1.6	14	35	35	35
WB-10	Ultra 9 285K	205	185	336	348	310	319
WB-80	Ultra 9 285K	30	27	28	218	205	205
WB-200	Ultra 9 285K	12	11	28	127	124	123
WB-500	Ultra 9 285K	3.9	3.7	24	65	64	64
WB-1000	Ultra 9 285K	1.7	1.6	15	35	34	33

2.4. Notes on Using the GPU-resident Option in the Latest Version of Gromacs

In our previous paper [4], we argued that the GPU-resident option significantly (up to 2.3 times) improved the performance of MD calculations. Using the latest version of Gromacs, this option no longer provides any performance gain, at least on large systems over 200000 atoms (Tab. 2, the last two columns).

Table 3. Single-node Gromacs 2025.2 molecular dynamics performance (ns/day) for tubulin tetramer on various GPUs for Intel i9-14900KF CPU (results with GPU-resident option on and off are equal)

Test number	GPU	Series name	Release date	Performance, ns/day
1	No GPU	—	—	11
2	RTX2080ti	Turing	Sep 2018	45
3	RTX3080	Ampere	Sep 2020	65
4	RTX3090	Ampere	Sep 2020	75
5	RTX4080	Ada Lovelace	Nov 2022	121
6	RTX4090	Ada Lovelace	Oct 2022	167
7	RTX5080	Blackwell	Jan 2025	146
8	RTX5090	Blackwell	Jan 2025	232

2.5. Contribution of the Graphics Accelerator to the Performance of MD Calculations

We tested the performance of MD calculations on the latest version of Gromacs and the new Intel i9-14900KF processor, depending on the model of the NVIDIA RTX (Ray Tracing shader eXtreme) graphics accelerator, starting with the most advanced RTX20 models (Turing series), RTX2080ti, and ending with the latest model RTX5090 (Blackwell series), see Tab. 3. The performance test was conducted on a full-atom model of a real biological system, which is a tetramer of the protein tubulin in explicit water.

Immediately striking is the significant increase in MD performance starting with the 40 series (Ada Lovelace). For example, RTX4080 with its 121 ns/day is more than 1.5 times faster than the previous top-end model, RTX3090 (78 ns/day). Moreover, the top-end version of the 40 series (167 ns/day) is actually more than twice as fast as the top-end version of the 30 series.

As for the latest 50 series of NVIDIA graphics accelerators, they do not show such impressive results compared to the 40 series, although the top version of the 50 series RTX5090 with its fantastic 232 ns/day is nevertheless 1.4 times faster than RTX4090. The younger model of the 50 series, RTX5080, shows significantly worse performance even than the top card of the previous 40 series. Overall, in the past 7 years, starting in 2018, the MD performance of top-end RTX series graphics cards has increased more than 5 times. If we compare the performance of a computer with a top-end graphics accelerator versus one without a graphics accelerator at all, the performance gain is more than 20 times.

2.6. Performance of the Latest Intel Ultra 9 Processor in MD Calculations When Using the New Top-End NVIDIA RTX Graphics Accelerators

We tested the performance of MD calculations on the latest version of Gromacs and the newest Ultra series of Intel desktop CPUs, the Ultra 9 285K processor, depending on the NVIDIA RTX graphics accelerator model (Tab. 4). Performance tests were produced on the same biological system as for Intel i9-14900KF processor – a tubulin tetramer in explicit water. Since the test results were quite unexpected, we also performed the test on a larger molecular system, namely WB-1000.

Using NVIDIA RTX30-series graphics accelerators, the computational performance of both types of the latest Intel processors (Ultra 9 285K and i9-14900KF) is virtually identical. The same situation is repeated with the lower-end graphics cards of the next generations – RTX4080 and RTX5080 – although in these cases the Ultra processor demonstrates slightly better results. However, the test results become completely unpredictable and different when it comes to the top-end GPUs, namely RTX4090 and RTX5090. With an Intel Ultra processor, these graphics cards demonstrate significantly lower performance even compared to their lower-end counterparts, not to mention the fantastic speed they demonstrated with the Intel i9-14900KF. Indeed, for RTX5090, the MD computation performance with the Ultra 9 285K processor dropped more than 2-fold compared to i9-14900KF! This seems incredible and may be explained by the significantly different internal architecture of Intel's new Ultra series CPU, in which the processor consists of multiple dies. Instead of a single monolithic die, Intel Core Ultra uses a disaggregated chiplet design and features a neural processing unit (NPU) for AI acceleration. This architecture separates components like the CPU, GPU, and NPU into "tiles" that are manufactured on optimal processes before being combined, leading to power efficiency.

However, for the WB-1000 molecular system, three times larger and consisting of a million atoms, the Ultra 9 285K processor again outperforms the i9-14900KF when using the 5090 graphics card (38 ns/day versus 33 ns/day). The result regarding the computational slowdown when switching from the 80th to the 90th GPU also does not hold for the WB-1000 system (Tab. 4, last column). Indeed, the calculation speed of the 5090 graphics card is twice that of the 5080.

The Intel Core Ultra 9 processor performs slower with memory than the i9-14900K, as the latter has higher clock speeds and better multi-core performance, which is critical for memory performance in applications. Indeed, the Ultra series processor has 24 threads, while the i9-14900K has 32. At the same time, the smaller the size of the molecular system, the more frequent data exchange with the video card occurs. This may explain why the Ultra 9 285K performance drops more sharply compared to the i9-14900K as the biological system size decreases (Tab. 2). Note that performance drops only when using the top-end 40- and 50-series graphics accelerators. This is likely due to their higher memory bus widths (384 and 512 bits, respectively) and larger memory capacities (24 GB and 32 GB, respectively) than their lower-end counterparts. It can be concluded that Ultra 9 series processors, despite their powerful integrated graphics, energy efficiency, and AI capabilities, are focused on different tasks than high-performance desktop processors like the i9-14900K, which is aimed at maximum performance in traditional computing tasks.

Conclusion

Over the past seven years, since the introduction of the NVIDIA RTX series of graphics accelerators in 2018 and the release of the latest versions of Gromacs, the speed of all-atom MD calculations has increased more than fivefold. Moreover, since 2014, MD performance has increased by almost 35 times when comparing a top-of-the-line gaming CPU Intel Core i7-4790K and a top video accelerator NVIDIA GTX 980, both released in 2014 [5]. All this has become possible thanks to the rapid and continuous development of the computing capabilities of central and, in particular, graphic processors.

Our comparative performance analysis of GPU-based computing architectures for all-atom MD simulations of biomolecular systems not only provides guidance on choosing the best computing solution in terms of the price-performance ratio, but also shows the maximum potential

Table 4. Single-node Gromacs 2025.2 molecular dynamics performance (ns/day) for tubulin tetramer and WB-1000 molecular systems on various GPUs for Intel Ultra 9 285K CPU (results with GPU-resident option on and off are equal)

Test number	GPU	Series name	Release date	Performance for tubulin tetramer, ns/day	Performance for WB-1000, ns/day
1	No GPU	—	—	15	1.6
2	RTX3080	Ampere	Sep 2020	65	11
3	RTX3090	Ampere	Sep 2020	77	12
4	RTX4080	Ada Lovelace	Nov 2022	126	18
5	RTX4090	Ada Lovelace	Oct 2022	99	25
6	RTX5080	Blackwell	Jan 2025	136	17
7	RTX5090	Blackwell	Jan 2025	104	33

computational performance that modern CPUs and GPUs can provide. In addition, we show how the latest software versions and computational options can improve the performance of the calculations. MD of the biomolecular system containing a tubulin protofilament in an explicitly specified solvent consisting of more than 300 thousand atoms can be studied with a performance of 232 ns/day at time step 2 fs when using a single-node computer with the latest CPU and GPU generation architecture (Intel Core i9-14900KF and Nvidia RTX5090, respectively).

Acknowledgments

This study was supported by the Russian Science Foundation, project No. 24-74-00002, https://rscf.ru/en/project/24-74-00002/.

References

1. Abraham, M., Alekseenko, A., Bergh, C., *et al.*: GROMACS 2023.2 Manual (Jul 2023). https://doi.org/10.5281/zenodo.8134388

2. Abraham, M.J., Murtola, T., Schulz, R., *et al.*: GROMACS: High performance molecular simulations through multi-level parallelism from laptops to supercomputers. SoftwareX 1, 19–25 (2015). https://doi.org/10.1016/j.softx.2015.06.001

3. Ewald, P.P.: Die Berechnung optischer und elektrostatischer Gitterpotentiale. Annalen der Physik 369(3), 253–287 (1921). https://doi.org/10.1002/andp.19213690304

4. Fedorov, V.A., Kholina, E.G., Gudimchuk, N.B., Kovalenko, I.B.: High-performance computing of microtubule protofilament dynamics by means of all-atom molecular modeling. Supercomputing Frontiers and Innovations 10(4), 62–68 (2023). https://doi.org/10.14529/jsfi230406

5. Fedorov, V.A., Kholina, E.G., Kovalenko, I.B., Gudimchuk, N.B.: 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), 111–114 (2018). https://doi.org/10.14529/jsfi180414

6. Gudimchuk, N.B., McIntosh, J.R.: Regulation of microtubule dynamics, mechanics and function through the growing tip. Nature Reviews Molecular Cell Biology 22(12), 777–795 (2021). https://doi.org/10.1038/s41580-021-00399-x

7. Kellogg, E.H., Hejab, N.M., Howes, S., *et al.*: Insights into the distinct mechanisms of action of taxane and non-taxane microtubule stabilizers from cryo-EM structures. Journal of Molecular Biology 429(5), 633–646 (2017). https://doi.org/10.1016/j.jmb.2017.01.001

8. Liao, Y., Ithurbide, S., Evenhuis, C., *et al.*: Cell division in the archaeon Haloferax volcanii relies on two FtsZ proteins with distinct functions in division ring assembly and constriction. Nature Microbiology 6, 594–605 (2021). https://doi.org/10.1038/s41564-021-00894-z

9. Loose, M., Mitchison, T.J.: The bacterial cell division proteins FtsA and FtsZ self-organize into dynamic cytoskeletal patterns. Nature Cell Biology 16(1), 38–46 (2014). https://doi.org/10.1038/ncb2885

10. MacKerell Jr, A.D., Feig, M., Brooks, C.L.: Improved treatment of the protein backbone in empirical force fields. Journal of the American Chemical Society 126(3), 698–699 (2004). https://doi.org/10.1021/ja036959e