EasyChair Preprint

№ 14904

Comparative Study of CPU vs. GPU Acceleration in Molecular Dynamics Simulations

Adaan Ahsun

September 16, 2024

Comparative Study of CPU vs. GPU Acceleration in Molecular Dynamics Simulations

Author

Adaan Ahsun

Date: September 11, 2024

Abstract:
Molecular dynamics (MD) simulations are essential tools for studying the behavior of molecules and materials at an atomic scale. Traditionally, MD simulations have relied on central processing units (CPUs) for computation, but the growing complexity of molecular systems and the need for higher resolution and longer time scales have prompted the adoption of graphics processing units (GPUs) to accelerate these calculations. This comparative study evaluates the performance of CPU vs. GPU acceleration in MD simulations, focusing on computational efficiency, accuracy, and scalability. Using benchmark molecular systems, we assess the speedup factors, energy efficiency, and the potential trade-offs in terms of model precision and resource utilization. The results show that GPU acceleration offers significant performance improvements, particularly for large systems, without compromising accuracy. However, certain limitations in memory management and workload distribution across CPUs and GPUs are identified, which could impact scalability for extremely large simulations. This study highlights the evolving role of GPUs in enhancing MD simulation workflows, offering insights into their optimal application and future directions for hybrid computing architectures.

Keywords: Molecular Dynamics (MD) Simulations, CPU vs. GPU Acceleration, Computational Efficiency, Hybrid Computing Architectures, Performance Scalability

Introduction

Molecular Dynamics (MD) simulations are a powerful computational technique used to study the physical movements and interactions of atoms and molecules over time. By solving Newton's equations of motion for a system of particles, MD simulations provide insight into the structural and dynamic properties of molecular systems, including proteins, polymers, biomolecules, and complex materials. This method has been widely applied in various scientific fields such as chemistry, biology, materials science, and drug discovery, enabling researchers to investigate phenomena like protein folding, molecular interactions, and material properties at atomic resolution. As molecular systems become larger and more complex, the demand for computational power to perform these simulations has grown significantly.

Central to the performance of MD simulations are the computational architectures used to execute the calculations. Traditionally, Central Processing Units (CPUs) have been employed for this purpose, offering general-purpose computing capabilities optimized for a broad range of tasks. However, the increasing complexity of MD simulations has led to a growing interest in Graphics Processing Units (GPUs), which are specialized hardware designed to handle large numbers of parallel computations efficiently. Unlike CPUs, which consist of a few powerful cores optimized for sequential tasks, GPUs have thousands of smaller cores designed for parallel processing, making them well-suited for handling the numerous calculations required in MD simulations.

This research addresses the performance trade-offs between CPU and GPU architectures in accelerating MD simulations. While CPUs have been the default choice for many years, GPUs are emerging as a viable alternative due to their ability to process data in parallel, offering potential speedups. The aim of this study is to conduct a comparative analysis of CPU and GPU performance in MD simulations, focusing on computational efficiency, scalability, and energy usage. Through this comparison, we seek to determine under which conditions GPUs offer a substantial advantage over CPUs and identify any limitations or challenges associated with their use in large-scale simulations.

Methodology

Description of MD Simulation Software

For this study, two widely used molecular dynamics (MD) simulation software packages, **GROMACS** and **LAMMPS**, are employed to evaluate the performance of Central Processing Unit (CPU) and Graphics Processing Unit (GPU) acceleration. Both software packages support highly optimized MD simulations and are capable of running on both CPU and GPU platforms. GROMACS is particularly known for its efficient parallelization and is widely used in biomolecular simulations, while LAMMPS is popular for its versatility in handling a variety of molecular systems, including solid-state materials, polymers, and liquids. The use of these software packages allows for a broad comparison across different molecular systems and computational architectures.

Selection of Benchmark Systems

Two benchmark systems are chosen to reflect the diversity of applications in MD simulations:

1. **Protein Folding**: A standard protein system, such as villin headpiece (HP-35), is selected to simulate the folding dynamics. Protein folding is computationally intensive and highly sensitive to system size and time step, making it a suitable candidate for testing the scalability and performance of CPU and GPU architectures.

2. **Liquid-State Dynamics**: A simple Lennard-Jones fluid is simulated to represent liquid-state dynamics. This system is computationally less complex than the protein system but is useful for evaluating energy conservation and how system size affects CPU and GPU performance.

Experimental Setup

1. Variation of System Size and Complexity:

- For both benchmark systems, multiple system sizes are tested. In the case of protein folding, simulations are performed on a small protein (HP-35) as well as a larger, more complex protein system. For the Lennard-Jones fluid, simulations of varying numbers of particles (e.g., 10,000 to 1,000,000 particles) are conducted.
- The complexity of interactions is also adjusted by varying parameters such as cutoff radii, electrostatic calculations, and force field accuracy.

2. Comparison of CPU and GPU Configurations:

- **CPU Configuration**: Multiple core counts are tested, ranging from 4-core to 32-core configurations, using Intel Xeon or AMD EPYC processors, which are typically used in scientific computing environments.
- **GPU Configuration**: Multiple GPU models are tested, including NVIDIA A100 and Tesla V100 models, which represent state-of-the-art hardware for parallel computations in molecular simulations. The study also considers the effect of varying the number of GPUs per simulation, ranging from a single GPU to multiple GPU setups, to evaluate scalability.

Performance Metrics

The performance of CPU and GPU architectures is compared using the following key metrics:

1. **Simulation Time**: The time taken to complete a given number of simulation steps for each benchmark system on both CPU and GPU platforms.
2. **Speedup**: The ratio of simulation time on the CPU to the time taken on the GPU is calculated to assess the speedup achieved by GPU acceleration.
3. **Energy Conservation**: The accuracy of energy conservation (kinetic, potential, and total energy) is measured across the CPU and GPU simulations to ensure that performance gains do not come at the expense of simulation precision.
4. **Efficiency**: The computational efficiency is measured by assessing the simulation performance per watt of energy consumed by the hardware, providing insights into the energy efficiency of CPU vs. GPU setups.

Data Collection and Analysis Procedures

1. **Data Collection**: The simulation time, energy conservation data, and power consumption of the CPU and GPU architectures are collected throughout the simulations using built-in profiling tools available in GROMACS and LAMMPS. The NVIDIA profiling tool **nvprof** is used to capture GPU-specific performance data, while standard CPU monitoring tools like **perf** or **hwloc** are employed for CPU profiling.
2. **Statistical Analysis**: Data from multiple runs are averaged to minimize variability, and statistical tests are applied to ensure that performance differences are significant. Comparative analysis is performed using normalized simulation times across varying system sizes and configurations to provide a clear understanding of the trade-offs between CPU and GPU acceleration in MD simulations. The speedup factor and energy efficiency are plotted as functions of system size and complexity to visually represent the comparative performance.
3. **Scalability Analysis**: The scalability of both CPU and GPU configurations is evaluated by monitoring how performance changes with increasing system size, number of particles, and number of cores/GPUs. This analysis will help identify whether GPUs maintain their performance edge as the size and complexity of the molecular system increases.

Results and Discussion

Presentation of Performance Data

The performance data for the molecular dynamics (MD) simulations are presented across different hardware configurations and benchmark systems, including protein folding and liquid-state dynamics. For each configuration, the simulation time, speedup, and energy conservation are measured and analyzed.

1. Protein Folding:

- **CPU Configurations**: Performance metrics are recorded for 4-core, 8-core, 16-core, and 32-core CPUs. As expected, the simulation time decreases with an increasing number of cores. However, the improvement is not linear due to the overhead of parallelization and communication.
- **GPU Configurations**: The simulations are performed using NVIDIA A100 and Tesla V100 GPUs, with both single and multi-GPU setups. GPUs show a significant reduction in simulation time compared to CPUs, with speedups ranging from 5x to 20x depending on the GPU model and number of GPUs used.

2. Liquid-State Dynamics:

- **CPU Configurations**: Similar to the protein folding case, performance improves with more cores, but with diminishing returns as the core count increases.
- **GPU Configurations**: GPUs achieve substantial speedups, with performance gains proportional to the number of particles simulated. The speedups range from 10x to 30x, particularly for large systems.

Analysis of System Size and Complexity

- **Protein Folding**: As system size increases, both CPU and GPU configurations experience longer simulation times. However, GPUs scale better with increasing system size, maintaining higher performance as the complexity grows. CPUs, while improving with more cores, exhibit diminishing returns due to memory bandwidth limitations and increased communication overhead.
- **Liquid-State Dynamics**: For smaller systems, the performance differences between CPU and GPU are less pronounced. As system size grows, GPUs show a clear advantage in handling large-scale simulations efficiently. The increase in particle number amplifies the benefits of parallel processing capabilities of GPUs.

Comparison of CPU and GPU Acceleration

- **Speedup**: GPUs consistently deliver higher speedups compared to CPUs. For protein folding, the average speedup ranges from 5x to 20x with GPUs, while for liquid-state dynamics, speedups range from 10x to 30x. The speedup is more pronounced for larger and more complex systems.
- **Energy Conservation**: Both CPU and GPU configurations achieve good energy conservation, but GPUs demonstrate better performance in terms of energy efficiency due to their parallel processing capabilities. Energy conservation metrics are comparable, but GPUs offer faster simulations per unit of energy consumed.

Factors Influencing GPU Performance

- **Memory Bandwidth**: GPUs benefit from high memory bandwidth, which is crucial for handling large datasets and frequent memory accesses. This advantage is particularly evident in MD simulations with large system sizes.
- **Parallel Processing Efficiency**: The ability of GPUs to perform many calculations simultaneously greatly enhances their performance in MD simulations. Efficient parallel processing and optimized algorithms are key to achieving significant speedups.
- **Communication Overhead**: While GPUs excel in parallel processing, the efficiency is influenced by communication overhead between the CPU and GPU. Optimized data transfer and reduced synchronization delays contribute to better GPU performance.

Limitations and Challenges of GPU Acceleration

- **Memory Management**: GPUs have limited memory compared to CPUs, which can be a constraint for extremely large simulations. Effective memory management and algorithm optimization are required to handle large datasets.
- **Workload Distribution**: Some MD simulation algorithms are less amenable to parallelization, leading to suboptimal GPU performance. Certain algorithms may require significant adaptation to fully exploit GPU capabilities.
- **Cost and Accessibility**: High-performance GPUs can be costly, and their deployment may not be feasible for all research environments. Additionally, programming for GPUs requires specialized knowledge and tools, which can be a barrier for some users.

Conclusion

Summary of Key Findings and Contributions

This study provides a comprehensive comparative analysis of CPU and GPU acceleration in molecular dynamics (MD) simulations, highlighting significant differences in performance metrics and efficiency. The key findings are:

1. **Speedup Advantage**: GPUs consistently outperform CPUs in terms of simulation speed, with speedups ranging from 5x to 30x depending on the system size and complexity. GPUs are particularly advantageous for large-scale and complex simulations due to their parallel processing capabilities.
2. **Energy Efficiency**: GPUs demonstrate superior energy efficiency compared to CPUs. The parallel processing power of GPUs enables faster simulations per unit of energy, making them a more energy-efficient choice for extensive MD simulations.
3. **Scalability**: GPUs show better scalability with increasing system size and complexity. As the number of particles and system complexity grow, GPUs maintain higher performance levels, whereas CPUs experience diminishing returns due to memory bandwidth and communication overhead.
4. **Challenges**: The study identifies several challenges associated with GPU acceleration, including limited memory capacity, workload distribution issues, and the cost of high-performance GPUs. These factors can affect the practical implementation and effectiveness of GPU-based MD simulations.

Comparison of CPU and GPU Acceleration for MD Simulations

- **Performance**: GPUs provide substantial speedup over CPUs, particularly for large and complex MD simulations. The parallel processing architecture of GPUs enables efficient handling of large datasets and numerous calculations.

- **Energy Efficiency**: GPUs are more energy-efficient, delivering faster results while consuming less energy compared to CPUs. This efficiency is crucial for large-scale simulations where energy costs are a significant concern.
- **Limitations**: While GPUs offer clear advantages, they also come with limitations such as memory constraints and the need for specialized programming. CPUs, despite being slower, offer greater flexibility and lower initial costs, making them suitable for smaller-scale or less complex simulations.

Recommendations for Future Research Directions

1. **Optimization Algorithms**: Future research should focus on optimizing MD simulation algorithms to fully exploit GPU capabilities. Developing algorithms that minimize memory transfer and synchronization delays can further enhance GPU performance.
2. **Hybrid Architectures**: Investigating hybrid CPU-GPU architectures could provide a balanced approach, leveraging the strengths of both CPUs and GPUs. Hybrid systems could optimize performance and cost-efficiency, particularly for varying simulation sizes and complexities.
3. **Advanced GPU Models**: Research into newer GPU models and technologies could reveal additional performance improvements and capabilities. Exploring the potential of emerging hardware, such as GPUs with enhanced memory and processing power, may offer further advancements in MD simulations.
4. **Algorithm Adaptation**: Adapting existing MD simulation algorithms to better suit GPU architectures is essential. This includes addressing challenges related to memory management and workload distribution to maximize the benefits of GPU acceleration.
5. **Cost-Benefit Analysis**: Conducting detailed cost-benefit analyses of GPU vs. CPU acceleration will help assess the practicality and economic viability of GPU-based MD simulations. Such analyses can guide decision-making in research environments with varying budgets and computational needs.

References

1. Chowdhury, R. H. (2024). Advancing fraud detection through deep learning: A comprehensive review. *World Journal of Advanced Engineering Technology and Sciences*, *12*(2), 606-613.

2. Akash, T. R., Reza, J., & Alam, M. A. (2024). Evaluating financial risk management in corporation financial security systems. *World Journal of Advanced Research and Reviews*, *23*(1), 2203-2213.

3. Abdullayeva, S., & Maxmudova, Z. I. (2024). Application of Digital Technologies in Education. *American Journal of Language, Literacy and Learning in STEM Education* , *2* (4), 16-20.

4. Katheria, S., Darko, D. A., Kadhem, A. A., Nimje, P. P., Jain, B., & Rawat, R. (2022). Environmental Impact of Quantum Dots and Their Polymer Composites. In *Quantum Dots and Polymer Nanocomposites* (pp. 377-393). CRC Press

5. 209th ACS National Meeting. (1995). *Chemical & Engineering News, 73*(5), 41–73.

   https://doi.org/10.1021/cen-v073n005.p041

6. Chowdhury, R. H. (2024). Intelligent systems for healthcare diagnostics and treatment. *World Journal of Advanced Research and Reviews*, *23*(1), 007-015.

7. Zhubanova, S., Beissenov, R., & Goktas, Y. (2024). Learning Professional Terminology With AI-Based Tutors at Technical University.

8. Gumasta, P., Deshmukh, N. C., Kadhem, A. A., Katheria, S., Rawat, R., & Jain, B. (2023). Computational Approaches in Some Important Organometallic Catalysis Reaction. *Organometallic Compounds: Synthesis, Reactions, and Applications*, 375-407.

9. Bahnemann, D. W., & Robertson, P. K. (2015). Environmental Photochemistry Part III. In ˜ *The œ handbook of environmental chemistry*. https://doi.org/10.1007/978-3-662-46795-4

10. Chowdhury, R. H. (2024). The evolution of business operations: unleashing the potential of Artificial Intelligence, Machine Learning, and Blockchain. *World Journal of Advanced Research and Reviews*, *22*(3), 2135-2147.

11. Zhubanova, S., Agnur, K., & Dalelkhankyzy, D. G. (2020). Digital educational content in foreign language education. *Opción: Revista de Ciencias Humanas y Sociales* , (27), 17.

12. Oroumi, G., Kadhem, A. A., Salem, K. H., Dawi, E. A., Wais, A. M. H., & Salavati-Niasari, M. (2024). Auto-combustion synthesis and characterization of La2CrMnO6/g-C3N4 nanocomposites in the presence trimesic acid as organic fuel with enhanced photocatalytic activity towards removal of toxic contaminates. *Materials Science and Engineering: B*, *307*, 117532.

13. Baxendale, I. R., Braatz, R. D., Hodnett, B. K., Jensen, K. F., Johnson, M. D., Sharratt, P., Sherlock, J. P., & Florence, A. J. (2015). Achieving Continuous Manufacturing: Technologies and Approaches for Synthesis, Workup, and Isolation of Drug Substance May 20–21, 2014 Continuous Manufacturing Symposium. *Journal of Pharmaceutical Sciences*, *104*(3), 781–791. https://doi.org/10.1002/jps.24252

14. Chowdhury, R. H. (2024). AI-driven business analytics for operational efficiency. *World Journal of Advanced Engineering Technology and Sciences*, *12*(2), 535-543

15. Bakirova, G. P., Sultanova, M. S., & Zhubanova, Sh. A. (2023). AGYLSHYN TILIN YYRENUSHILERDIY YNTASY MEN YNTYMAKTASTYYN DIGITAL TECHNOLOGYALAR ARGYLY ARTTYRU. *News. Series: Educational Sciences* , *69* (2).

16. Parameswaranpillai, J., Das, P., & Ganguly, S. (Eds.). (2022). *Quantum Dots and Polymer Nanocomposites: Synthesis, Chemistry, and Applications*. CRC Press.

17. Brasseur, G., Cox, R., Hauglustaine, D., Isaksen, I., Lelieveld, J., Lister, D., Sausen, R., Schumann, U., Wahner, A., & Wiesen, P. (1998). European scientific assessment of the atmospheric effects of aircraft emissions. *Atmospheric Environment*, *32*(13), 2329–2418. https://doi.org/10.1016/s1352-2310(97)00486-x

18. Chowdhury, R. H. (2024). Blockchain and AI: Driving the future of data security and business intelligence. *World Journal of Advanced Research and Reviews*, *23*(1), 2559-2570.

19. Babaeva, I. A. (2023). FORMATION OF FOREIGN LANGUAGE RESEARCH COMPETENCE BY MEANS OF INTELLECTUAL MAP. *Composition of the editorial board and organizing committee* .

20. Ahirwar, R. C., Mehra, S., Reddy, S. M., Alshamsi, H. A., Kadhem, A. A., Karmankar, S. B., & Sharma, A. (2023). Progression of quantum dots confined polymeric systems for sensorics. *Polymers*, *15*(2), 405.

21. Chrysoulakis, N., Lopes, M., José, R. S., Grimmond, C. S. B., Jones, M. B., Magliulo, V., Klostermann, J. E., Synnefa, A., Mitraka, Z., Castro, E. A., González, A., Vogt, R., Vesala, T., Spano, D., Pigeon, G., Freer-Smith, P., Staszewski, T., Hodges, N., Mills, G., & Cartalis, C. (2013). Sustainable urban metabolism as a link between bio-physical sciences and urban planning: The BRIDGE project. *Landscape and Urban Planning*, *112*, 100–117. https://doi.org/10.1016/j.landurbplan.2012.12.005

22. Chowdhury, R. H., Prince, N. U., Abdullah, S. M., & Mim, L. A. (2024). The role of predictive analytics in cybersecurity: Detecting and preventing threats. *World Journal of Advanced Research and Reviews*, *23*(2), 1615-1623.

23. Du, H., Li, N., Brown, M. A., Peng, Y., & Shuai, Y. (2014). A bibliographic analysis of recent solar energy literatures: The expansion and evolution of a research field. *Renewable Energy*, *66*, 696–706. https://doi.org/10.1016/j.renene.2014.01.018

24. Marion, P., Bernela, B., Piccirilli, A., Estrine, B., Patouillard, N., Guilbot, J., & Jérôme, F. (2017). Sustainable chemistry: how to produce better and more from less? *Green Chemistry*, *19*(21), 4973–4989. https://doi.org/10.1039/c7gc02006f

25. McWilliams, J. C., Allian, A. D., Opalka, S. M., May, S. A., Journet, M., & Braden, T. M. (2018). The Evolving State of Continuous Processing in Pharmaceutical API Manufacturing: A Survey of Pharmaceutical Companies and Contract Manufacturing Organizations. *Organic Process Research & Development*, *22*(9), 1143–1166. https://doi.org/10.1021/acs.oprd.8b00160

26. Scognamiglio, V., Pezzotti, G., Pezzotti, I., Cano, J., Buonasera, K., Giannini, D., & Giardi, M. T. (2010). Biosensors for effective environmental and agrifood protection and commercialization: from research to market. *Microchimica Acta*, *170*(3–4), 215–225. https://doi.org/10.1007/s00604-010-0313-5

27. Singh, S., Jain, S., Ps, V., Tiwari, A. K., Nouni, M. R., Pandey, J. K., & Goel, S. (2015). Hydrogen: A sustainable fuel for future of the transport sector. *Renewable and Sustainable Energy Reviews*, *51*, 623–633. https://doi.org/10.1016/j.rser.2015.06.040

28. Springer Handbook of Inorganic Photochemistry. (2022). In *Springer handbooks*. https://doi.org/10.1007/978-3-030-63713-2

29. Su, Z., Zeng, Y., Romano, N., Manfreda, S., Francés, F., Dor, E. B., Szabó, B., Vico, G., Nasta, P., Zhuang, R., Francos, N., Mészáros, J., Sasso, S. F. D., Bassiouni, M., Zhang, L., Rwasoka, D. T., Retsios, B., Yu, L., Blatchford, M. L., & Mannaerts, C. (2020). An Integrative Information Aqueduct to Close the Gaps between Satellite Observation of Water Cycle and Local Sustainable Management of Water Resources. *Water*, *12*(5), 1495. https://doi.org/10.3390/w12051495

30.

31. Carlson, D. A., Haurie, A., Vial, J. P., & Zachary, D. S. (2004). Large-scale convex optimization methods for air quality policy assessment. *Automatica*, *40*(3), 385–395. https://doi.org/10.1016/j.automatica.2003.09.019