

Powering Discovery: Leveraging NVIDIA H100 for Advanced Molecular Dynamics Simulations

Author: BenchChem Technical Support Team. **Date:** December 2025

Compound of Interest

Compound Name: H100

Cat. No.: B15585327

[Get Quote](#)

Application Note and Protocols for Researchers, Scientists, and Drug Development Professionals

The NVIDIA **H100** Tensor Core GPU, built on the Hopper architecture, represents a significant leap forward in computational power, offering unprecedented performance for molecular dynamics (MD) simulations. This document provides detailed application notes and protocols for harnessing the capabilities of the **H100** to accelerate research in drug discovery, materials science, and structural biology. By leveraging the architectural advancements of the **H100**, researchers can tackle larger, more complex biological systems over longer timescales, leading to deeper insights into molecular mechanisms.

Introduction to H100 for Molecular Dynamics

The NVIDIA **H100** GPU introduces several key features that are highly beneficial for MD simulations, including fourth-generation Tensor Cores, a significant increase in FP64 and FP32 performance, and HBM3 memory with higher bandwidth compared to its predecessor, the A100.^{[1][2][3][4]} These improvements translate to substantial performance gains in widely-used MD software packages such as GROMACS, AMBER, NAMD, and OpenMM.

The **H100**'s enhanced memory capacity and bandwidth allow for the simulation of larger molecular systems, while its superior computational throughput accelerates the time-to-solution for complex calculations like non-bonded interactions and long-range electrostatics.^{[1][2]}

Performance Benchmarks

The following tables summarize the performance of the NVIDIA **H100** GPU in various MD simulation benchmarks, compared to other relevant GPUs. The performance metric is typically measured in nanoseconds of simulation per day (ns/day), where a higher value indicates better performance.

Table 1: GROMACS Performance Benchmarks

System Description	Number of Atoms	GPU	Performance (ns/day)
R-143a in hexane	20,248	H100	Data not available in search results
Short RNA piece in explicit water	31,889	H100	Data not available in search results
Protein in membrane (explicit water)	80,289	H100	Data not available in search results
Protein in explicit water	170,320	H100	Data not available in search results
Protein membrane channel (explicit water)	615,924	H100	Data not available in search results
Huge virus protein (STMV)	1,066,628	H100	Data not available in search results

Note: While specific ns/day values for **H100** with GROMACS were not found in the provided search results, the general consensus is that the **H100** significantly outperforms previous generations. One source indicated that for large models, consumer GPUs like the RTX 4090 can sometimes match or exceed the performance of datacenter GPUs like the **H100** when using a limited number of OpenMP threads.[\[5\]](#)

Table 2: AMBER Performance Benchmarks

System Description	Number of Atoms	GPU	Performance (ns/day)
DHFR (JAC Prod.) NVE 2fs	23,558	H100 (PCIe)	~1400
FactorIX NVE 2fs	90,906	H100 (PCIe)	~400
Cellulose NVE 2fs	408,609	H100 (PCIe)	~150
STMV Production NPT 4fs	1,067,095	H100 (PCIe)	~50

Data is estimated from benchmarks presented by Exxact Corp.[6] It's important to note that some benchmarks show high-end consumer cards like the RTX 4090 outperforming the **H100** for certain AMBER simulations, suggesting that the **H100**'s strengths are more pronounced in AI and mixed-precision workloads.[7][8]

Table 3: NAMD Performance Benchmarks

System Description	Number of Atoms	GPU	Performance (ns/day)
ApoA1	92,224	H100 (PCIe)	Data not available in search results
STMV	1,066,628	H100 (PCIe)	17.06

The NAMD benchmark for the **H100** PCIe on the STMV system showed a performance of 17.06 ns/day.[9] It was noted that for this specific workload, the cost-performance of the **H100** might not be optimal compared to other GPUs.[9]

Table 4: OpenMM Performance Benchmarks

System Description	Number of Atoms	GPU	Throughput (μs/day)
DHFR (8 simulations with MPS)	23,558	H100	Approaches 5
Cellulose (with MPS)	409,000	H100	~20% increase over single simulation

With OpenMM, using NVIDIA's Multi-Process Service (MPS) can significantly improve throughput by allowing multiple simulations to run concurrently on the same GPU, which is particularly effective for smaller systems that do not fully saturate the GPU's resources.^[10] For some systems, this can more than double the total throughput.^[10]

Experimental Protocols

This section provides detailed methodologies for setting up and running MD simulations on systems equipped with NVIDIA **H100** GPUs.

System Preparation and Software Installation

- Hardware and Drivers:
 - Ensure the system has one or more NVIDIA **H100** GPUs installed.
 - Install the latest NVIDIA drivers for the **H100** GPU to ensure optimal performance and compatibility.
 - A sufficient amount of system RAM (at least 256GB is recommended for complex simulations) should be available.^[11]
- CUDA Toolkit:
 - Download and install the latest version of the NVIDIA CUDA Toolkit (version 12.x or later is recommended).^{[11][12]} The CUDA Toolkit is essential for the GPU to be utilized by the MD software.^[12]
- MD Software Installation:

- GROMACS: Download the latest version of GROMACS and compile it with CUDA support. Ensure that the GROMACS build is optimized for the Hopper architecture.
- AMBER: Install the latest version of Amber and AmberTools. The configuration script should automatically detect the **H100** and compile the necessary CUDA code.
- NAMD: Use a version of NAMD that is compiled with CUDA support.[\[12\]](#)
- OpenMM: Install OpenMM, which is designed for high performance on GPUs.

General Molecular Dynamics Simulation Protocol

The following protocol outlines the general steps for running a biomolecular simulation.

- System Setup:
 - Obtain Initial Structure: Start with a high-resolution crystal structure from the Protein Data Bank (PDB) or a model generated by homology modeling.
 - Prepare the Protein: Clean the PDB file by removing any unwanted molecules (e.g., crystal waters, ligands not under study). Add missing hydrogen atoms.
 - Choose a Force Field: Select an appropriate force field (e.g., AMBER, CHARMM, OPLS) based on the nature of the biomolecule.
 - Solvation: Place the biomolecule in a periodic box of a chosen shape (e.g., cubic, dodecahedron) and solvate it with an explicit water model (e.g., TIP3P, SPC/E).
 - Add Ions: Add ions to neutralize the system and to mimic physiological salt concentration.
- Energy Minimization:
 - Perform energy minimization to relax the system and remove any steric clashes or unfavorable geometries introduced during the setup. This is typically done in two stages: first minimizing the positions of water and ions with the protein heavy atoms restrained, and then minimizing the entire system without restraints.
- Equilibration:

- Gradually heat the system to the desired temperature (e.g., 300 K) under the NVT (canonical) ensemble, with position restraints on the protein heavy atoms. This allows the solvent to equilibrate around the protein.
- Switch to the NPT (isothermal-isobaric) ensemble to equilibrate the pressure and density of the system, while gradually reducing the position restraints on the protein.
- Production Run:
 - Run the production MD simulation for the desired length of time (nanoseconds to microseconds) under the NPT ensemble without any restraints. This is the main data-gathering phase of the simulation.
- Analysis:
 - Analyze the trajectory to study the dynamics, structure, and interactions of the biomolecular system. This can include calculating root-mean-square deviation (RMSD), root-mean-square fluctuation (RMSF), radius of gyration, hydrogen bonds, and performing principal component analysis (PCA).


Optimizing Performance on H100

To maximize the performance of MD simulations on the **H100**, consider the following settings:

- Precision: Utilize mixed-precision (FP32/FP64) modes where supported by the simulation software to balance accuracy and speed.[\[11\]](#) The **H100**'s Tensor Cores can accelerate FP64 calculations.[\[11\]](#)
- Parallelization: For large-scale simulations, leverage multi-GPU support. The **H100**'s NVLink technology provides high-speed communication between GPUs.[\[1\]](#)[\[11\]](#)
- Memory Management: Efficiently allocate GPU memory by adjusting buffer sizes in the MD software to minimize data transfers.[\[11\]](#)
- NVIDIA Multi-Process Service (MPS): For running multiple smaller simulations, especially with OpenMM, using MPS can significantly increase overall throughput by allowing concurrent execution on a single GPU.[\[10\]](#) This is achieved by eliminating context-switching overhead.[\[10\]](#)


Visualizations

The following diagrams illustrate key workflows and relationships in utilizing the **H100** for molecular dynamics simulations.



[Click to download full resolution via product page](#)

Caption: General workflow for a molecular dynamics simulation.



[Click to download full resolution via product page](#)

Caption: Hardware and software stack for MD simulations on **H100**.

Need Custom Synthesis?

BenchChem offers custom synthesis for rare earth carbides and specific isotopic labeling.

Email: info@benchchem.com or [Request Quote Online](#).

References

- 1. What are the key differences between NVIDIA A100 and H100 GPUs for running large-scale molecular dynamics simulations? - Massed Compute [massedcompute.com]
- 2. Can you compare the performance of A100 SXM4 and H100 SXM5 GPUs in specific HPC applications? - Massed Compute [massedcompute.com]

- 3. Comparing NVIDIA H100 vs A100 GPUs for AI Workloads | OpenMetal IaaS [openmetal.io]
- 4. trgdatacenters.com [trgdatacenters.com]
- 5. blog.salad.com [blog.salad.com]
- 6. NVIDIA Ampere GPU Benchmarks for AMBER 22 [exxactcorp.com]
- 7. [AMBER] H100 & RTX4090 Benchmarks from Ross Walker via AMBER on 2022-12-03 (Amber Archive Dec 2022) [archive.ambermd.org]
- 8. Which GPU for MD Simulations? | SabrePC Blog [sabrepc.com]
- 9. NAMD GPU Benchmarks and Hardware Recommendations | Exxact Blog [exxactcorp.com]
- 10. Maximizing OpenMM Molecular Dynamics Throughput with NVIDIA Multi-Process Service | NVIDIA Technical Blog [developer.nvidia.com]
- 11. What are the recommended settings for NVIDIA H100 GPU acceleration in molecular dynamics simulations? - Massed Compute [massedcompute.com]
- 12. researchgate.net [researchgate.net]
- To cite this document: BenchChem. [Powering Discovery: Leveraging NVIDIA H100 for Advanced Molecular Dynamics Simulations]. BenchChem, [2025]. [Online PDF]. Available at: [https://www.benchchem.com/product/b15585327#using-h100-for-molecular-dynamics-simulations]

Disclaimer & Data Validity:

The information provided in this document is for Research Use Only (RUO) and is strictly not intended for diagnostic or therapeutic procedures. While BenchChem strives to provide accurate protocols, we make no warranties, express or implied, regarding the fitness of this product for every specific experimental setup.

Technical Support: The protocols provided are for reference purposes. Unsure if this reagent suits your experiment? [[Contact our Ph.D. Support Team for a compatibility check](#)]

Need Industrial/Bulk Grade? [Request Custom Synthesis Quote](#)

BenchChem

Our mission is to be the trusted global source of essential and advanced chemicals, empowering scientists and researchers to drive progress in science and industry.

Contact

Address: 3281 E Guasti Rd

Ontario, CA 91761, United States

Phone: (601) 213-4426

Email: info@benchchem.com