HPC Comes to Life with AMD Instinct™ GPUs and NAMD

November, 2022

Application:
NAMD code for molecular dynamics modeling

Challenge:
In order to run larger and more complex simulations, researchers using NAMD need access to faster, higher-precision processing capable of processing parallel code more efficiently.

Solution:
AMD Instinct™ MI200 GPU

Benefits:
AMD Instinct MI200 GPUs provide organizations with the flexibility to meet any research requirement using NAMD with a range of options that can deliver massive system-level performance and high density, or superior per GPU performance.

The molecular dynamics code package NAMD is used to simulate how biomolecules move over time. By creating computer models of what is occurring at the molecular level, NAMD makes it possible for researchers to gain a deeper knowledge of biological processes. NAMD has outstanding scalability and can be applied to thousands of CPU cores and GPUs, making it uniquely capable of modeling large, complex biomolecular systems such as complete cells and, recently, the entire Covid virion. AMD Instinct accelerators provide breakthrough performance well-suited to the demands of NAMD-based simulation.

AMD Instinct™ MI200 GPUs were designed using the 2nd Generation AMD CNDA™ Architecture, an innovative platform that delivers a quantum leap in HPC and AI performance over competitive data center GPUs. Industry-leading double-precision performance for HPC workloads enables scientists and researchers across the globe to process HPC parallel codes more efficiently. As a result, AMD Instinct™ MI200 GPUs provide the tremendous performance needed to make the most of the latest generation solutions such as NAMD. The AMD Instinct MI200 series accelerator is even the compute engine selected to support the first U.S. Exascale supercomputer, Oak Ridge National Laboratory's Frontier system.

AMD innovations in architecture, packaging, and integration push the boundaries of computing by unifying the most important processors in the data center, the CPU and the GPU accelerator. With industry-first multi-chip GPU modules and 3rd Generation AMD Infinity Architecture, AMD delivers performance, efficiency, and overall system throughput for HPC and AI.

Guided by a philosophy of choice without restriction, the AMD Instinct family makes it possible to choose the GPU that best meets your needs, confident that you'll experience the breakthrough performance you need. Instinct GPUs provide a range of options to fit any accelerator requirement, with configurations that range from one GPU (MI210 PCIe CEM card) to multiple GPUs per system (MI250 OAM or up to 8xMI210).

AMD Infinity Fabric™ Link Technology	Ultra-Fast HBM2e Memory
AMD Instinct MI200 series accelerators with advanced peer-to-peer I/O connectivity through a maximum of eight AMD Infinity Fabric™ links deliver up to 800 GB/s I/O bandwidth performance. WiUltra-Fast HBM2e Memory.	Up to 128GB High bandwidth HBM2e memory with ECC support at a clock rate of 1.6 GHz. and deliver an ultra-high 3.2 TB/s of memory bandwidth to help support your largest data sets and eliminate bottlenecks in moving data in and out of memory.

AMD INSTINCT MI200 ACCELERATORS DELIVER PERFORMANCE, EFFICIENCY, AND FLEXIBILITY TO MAXIMIZE HPC AND AI WORKLOADS

PERFORMANCE. 2nd Gen AMD CDNA™ architecture brings enhanced Matrix Core Technology with 880 Matrix Cores in AMD Instinct™ MI200 OAM accelerators boosting computational capabilities and throughput. AMD CDNA™ 2 supports a wider range of data types and applications, bringing full rate double precision and new FP64 Matrix operations for HPC workloads.

EFFICIENCY. AMD CDNA is the architectural foundation of MI200 series, designed to leverage the performance, power and efficiency advantages of the 6nm fabrication process, helping to improve code performance.

FLEXIBILITY. ROCm™ 5.0 extends AMD's open platform for HPC and AI, optimizing for you at the device level with a freely shared codebase that can evolve as systems and their parts evolve.

AMD INSTINCT FOR MOLECULAR DYNAMICS

Core IPC (instructions per cycle) is a critical factor in optimizing the performance of NAMD. AMD Instinct GPUs employing the revolutionary 2nd Generation AMD CNDA™ architecture ensure that you get the most out of your system, minimizing execution time and increasing overall utilization of your deployment.

Ultra-Fast Memory. AMD Instinct MI200 series accelerators offer up to 128GB high-bandwidth HBM2e memory with ECC support at a clock rate of 1.6 GHz. An ultra-high 3.2 TB/s of memory bandwidth helps support even the largest data sets and eliminate bottlenecks when moving data in and out of memory.

Many High-Performance Compute (HPC) workloads require you to balance performance vs. per-core license costs to manage your overall cost. AMD Instinct GPUs offer consistent features across the product line, allowing users to optimize their workloads without sacrificing features, memory channels, memory capacity, or I/O lanes.

AMD Infinity Fabric™ Link Technology. As workloads demand more data processing capabilities, the communications between Instinct GPUs become critical to efficiently solving the complex problems customers face. AMD Instinct MI200 series accelerators feature advanced peer-to-peer I/O connectivity through a maximum of eight AMD Infinity Fabric™ links creating a system capable of delivering up to 800 GB/s I/O bandwidth performance.

MI200 PRODUCT OPTIONS

	Instinct MI100	Instinct MI210	Instinct MI250
GPU Architecture	CDNA	CDNA2	CDNA2
Compute Units	120	104	208
Total Board Power (TBP)	300W Peak	300W Peak	500W | 560W Peak
Dedicated Memory Size	32 GB	64 GB	128 GB
Dedicated Memory Type	HBM2	HBM2e	HBM2e
Memory Clock	1.2 GHz	1.6 GHz	1.6 GHz
Peak Memory Bandwidth	Up to 1228.8 GB/s	Up to 1638.4 GB/s	Up to 3276.8 GB/s
Form Factor	PCIe® Add-in Card	PCIe® Add-in Card	
Bus Type	PCIe® 4.0 x16, PCIe® 3.0 x16	PCIe® 4.0 x16	OAM Module
Infinity Fabric™ Links	3	3	8
Peak Infinity Fabric™ Link Bandwidth	92 GB/s	100 GB/s	100 GB/s
Cooling	Passive	Passive	Passive OAM

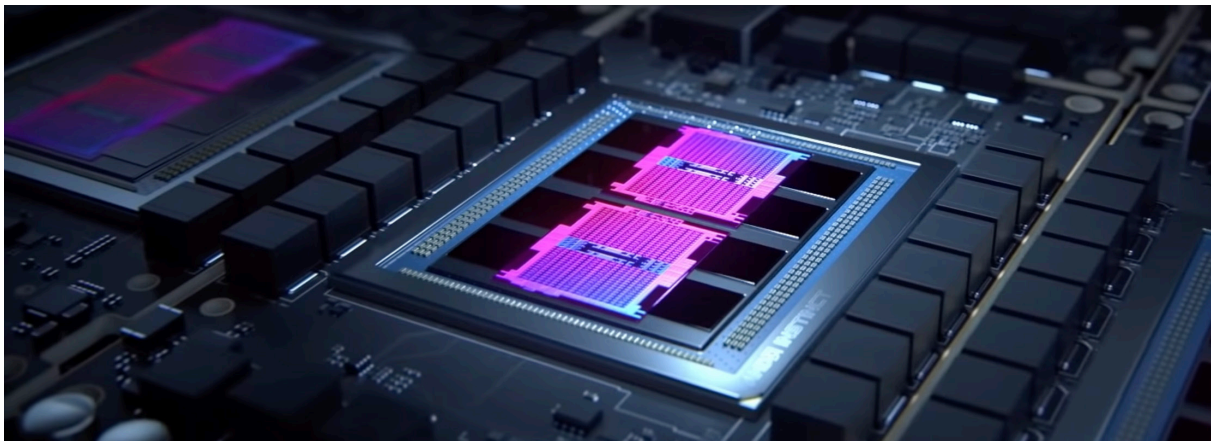

NAMD

Based on Charm++ parallel objects, NAMD scales to hundreds of cores for typical simulations and beyond 500,000 cores for the largest simulations. NAMD uses the popular molecular graphics program VMD for simulation setup and trajectory analysis but is also file-compatible with AMBER, CHARMM, and X-PLOR. NAMD is distributed free of charge with source code. You can build NAMD yourself or download binaries for a wide variety of platforms.

NAMD BENCHMARKS

The industry standard benchmarks for NAMD include STMV, ApoA1, and f1atpase. These benchmarks are well-established NAMD workloads that allow users to compare the performance of different hardware solutions to determine which solutions are best for their needs.

PERFORMANCE BENCHMARKS AND TESTING

NAMD benchmarks provide a basis of evaluating hardware performance. Standard models are provided that represent typical usage. The benchmarks used were APOA1 NVE and STMV NVE benchmark STMV, APOA1.

	Instinct MI100	Instinct MI210	Instinct MI250
CPU		2P AMD EPYC™ 7763	2P EPYC™ 7763
Cores		64	64
Memory			2 TiB (16 DIMMs 3200 mts)
Storage			3.5T NVMe drive
GPU		MI210 GPU (PCIe® 64GB 300W)	4x AMD Instinct™ MI250 (128 GB HBM2e) 560W
OS / SW Settings		Ubuntu® 21.04.4 LTS with Linux kernel 5.13.0-44-generic, System BIOS M12, ROCm 5.1.3	Ubuntu® 21.04.4 LTS with Linux kernel 5.13.0-44-generic, System BIOS M12, ROCm 5.1.3

Application SW	
APOA1 NVE	AMDih/NAMD3:3.0a9 container
STMV NV	AMDih/NAMD3:3.0a9 container

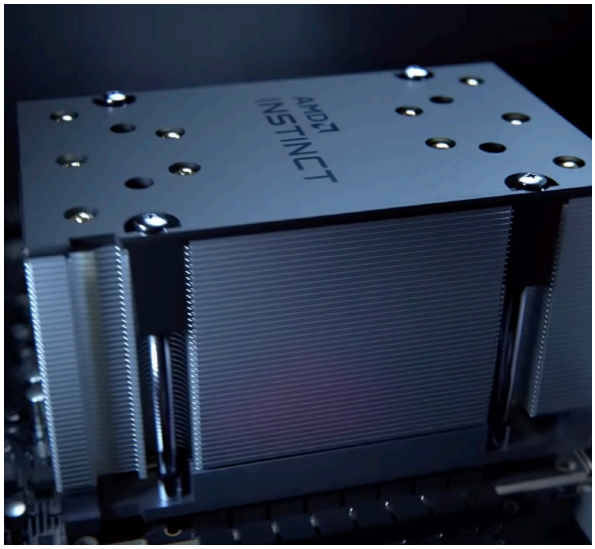

NAMD Performance and Scaling: Single Node Performance Single node performance was compared between the M100, MI210 and MI250 systems. The STMV benchmark is one of the most common benchmarks used to compare NAMD performance across various platforms. STMV is useful for demonstrating scaling across accelerators.

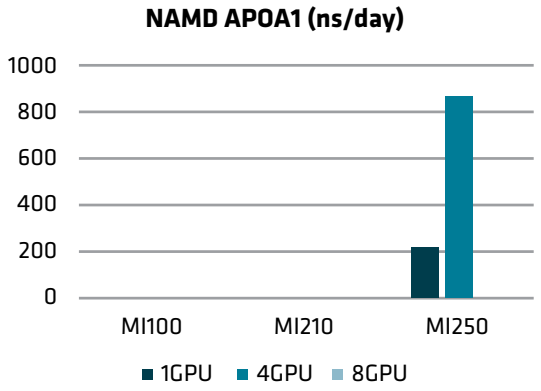

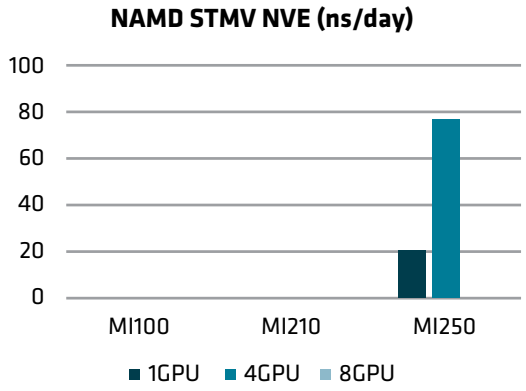

SUMMARY

NAMD benchmarks were conducted on single, two-socket systems running AMD EPYC 7763 processors with AMD Instinct MI100, MI210, and MI250 GPUs. NAMD scales very well across all benchmarks to a large number of cores and GPUs.

CONCLUSION

Scale-out testing on the AMD Instinct cluster shows impressive results on these benchmarks. Whether you need the dominating system-level performance and density or the equally dominating per GPU performance of the AMD Instinct, all products offer exceptional performance that will significantly benefit your organization.

Testing conducted by AMD performance lab 8.2.22 running NAMD v3 APOA1 NVE and STMV NVE.