Enhancing LAMMPS Simulations with AMD Instinct™ Accelerators: Unleashing Performance and Scalability

June, 2023

Application:
LAMMPS code for molecular dynamics modeling

Challenge:
To simulate increasingly larger and more complex materials and systems in LAMMPS, researchers require access to high-performance computing resources that can process parallel code more efficiently.

Solution:
AMD Instinct™ MI200 GPU

Benefits:
AMD Instinct MI200 accelerators versus the comparative NVIDIA offering provide enhanced performance, efficiency, and adaptability to optimize high-performance computing and artificial intelligence workloads when utilizing LAMMPS.

LAMMPS is used to mimic the long-term behavior of atomic and molecular systems using molecular dynamics. LAMMPS helps scientists gain a better understanding of a wide range of material and biological processes by creating computer models that represent interactions at the molecular level. Since LAMMPS can take advantage of hundreds of CPU cores and GPUs, it is well-suited for modeling complex molecular systems like biomolecules and cutting-edge materials. When it comes to LAMMPS-based simulations, the revolutionary performance offered by AMD Instinct accelerators is what's needed.

To provide a significant improvement in HPC and AI performance over other data center GPUs, AMD Instinct™ MI200 GPUs are constructed on the 2nd Generation AMD CDNA™ Architecture. Scientists and researchers all over the world can now process HPC parallel codes with more efficiency thanks to the double-precision performance for HPC workloads. Therefore, cutting-edge solutions like LAMMPS benefit greatly from the outstanding performance provided by AMD Instinct™ MI200 GPUs. As an example, the Frontier system at Oak Ridge National Laboratory, which is the first U.S. Exascale supercomputer, is powered by AMD Instinct MI200 series accelerators.

By unifying the CPU (central processing unit) and the GPU (graphics processing unit) accelerator, AMD's latest innovations in architecture, packaging, and integration pave the way for revolutionary computing. With its third-generation AMD Infinity Architecture and groundbreaking multi-chip GPU modules, AMD delivers on performance, efficiency, and overall system throughput for high-performance computing (HPC) and artificial intelligence (AI) applications.

The AMD Instinct™ accelerator family, which adheres to a concept of freedom of choice, lets you pick the GPU that's a perfect fit for your needs, guaranteeing you'll get the kind of ground-breaking performance you're hoping for. AMD Instinct GPUs come in a wide range of configurations, from a single GPU (the AMD Instinct MI210 PCIe CEM card) to as many as eight GPUs in a single system (the MI250 OAM or up to eight MI210 GPUs).

simulations and obtain deeper insights into the characteristics and behavior of materials at the atomic and molecular level by exploiting these powerful computing tools.

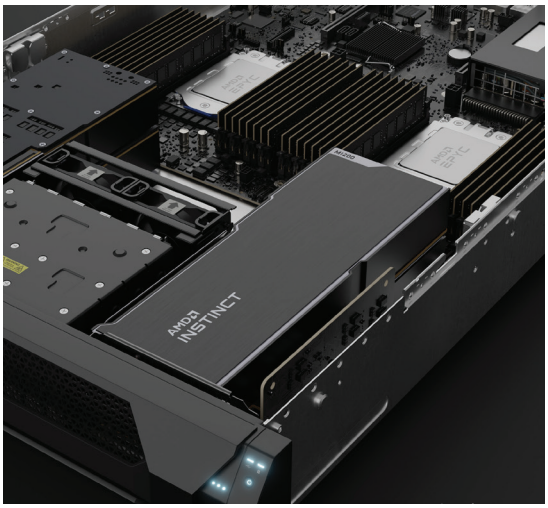

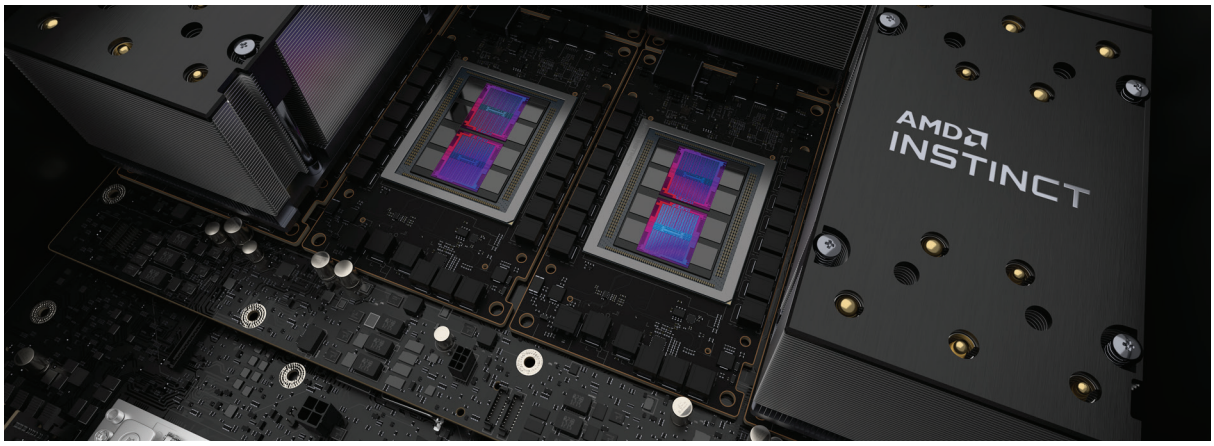

AMD Infinity Fabric™ Link Technology	Ultra-Fast HBM2e Memory
AMD Instinct MI200 series accelerators with advanced peer-to-peer I/O connectivity through a maximum of eight AMD Infinity Fabric™ links deliver up to 800 GB/s I/O bandwidth performance. WiUltra-Fast HBM2e Memory.	Up to 128GB High bandwidth 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.

USERS OF LAMMPS REQUIRE HIGH-PERFORMANCE COMPUTING RESOURCES THAT PROVIDE THE BEST PERFORMANCE, EFFICIENCY, AND FLEXIBILITY

PERFORMANCE. The AMD Instinct™ MI200 OAM accelerators are powered by the 2nd Gen AMD CDNA™ architecture, which provides extreme computational capabilities and throughput with 880 Matrix Cores and improved Matrix Core Technology. This architecture also allows a broad range of data types and applications, such as full-rate double precision and new FP64 Matrix operations, which are very beneficial for HPC workloads.

EFFICIENCY. The 2nd Gen. AMD CDNA™ 2 architecture is designed to take use of the 6nm fabrication process's performance, power, and efficiency gains, resulting in Increased code performance for LAMMPS simulations.

FLEXIBILITY. AMD ROCm™ 5.0 platform expands AMD's open platform for HPC and AI by enabling device-level optimization and providing a freely shared codebase that can change alongside LAMMPS simulations as systems and their components improve. LAMMPS users may effectively execute complicated simulations and get deeper insights into the characteristics and behavior of materials at the atomic and molecular level by exploiting these powerful computing technologies.

LAMMPS ACCELERATION WITH AMD INSTINCT FOR MOLECULAR DYNAMICS

The performance of LAMMPS simulations is determined by a number of elements, including Core IPC (instructions per cycle) and memory bandwidth. The AMD Instinct GPUs, which are driven by the groundbreaking 2nd Generation AMD CDNA™ architecture, will help you get the most out of your LAMMPS simulations by giving low execution time and enhanced overall deployment utilization. The AMD Instinct MI200 series accelerators feature ultra-fast memory that can handle even the largest data sets and eliminate bottlenecks when moving data in and out of memory, with up to 128GB of high-bandwidth HBM2e memory and a clock rate of 1.6 GHz.

Balancing performance with per-core license prices can be difficult for many high-performance computation (HPC) workloads while managing the overall cost of your deployment. AMD Instinct GPUs have consistent features across the product line, enabling users to optimize workloads without losing features, memory channels, memory capacity, or I/O lanes.

Advanced peer-to-peer I/O connectivity is also required for LAMMPS simulations to solve complicated issues quickly. AMD Instinct MI200 series accelerators use AMD Infinity Fabric™ Link Technology, which delivers up to eight AMD Infinity Fabric™ links, resulting in a system with up to 800 GB/s I/O bandwidth performance. LAMMPS users may efficiently execute complicated simulations and obtain deeper insights into the characteristics and behavior of materials at the atomic and molecular level by exploiting these powerful computing tools.

AMD INSTINCT MI200 SERIES ACCELERATOR PRODUCT OPTIONS

	Instinct MI100	Instinct MI210	Instinct MI250
GPU Architecture	AMD CDNA 2	AMD CDNA 2	AMD CDNA 2
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	OAM Module
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

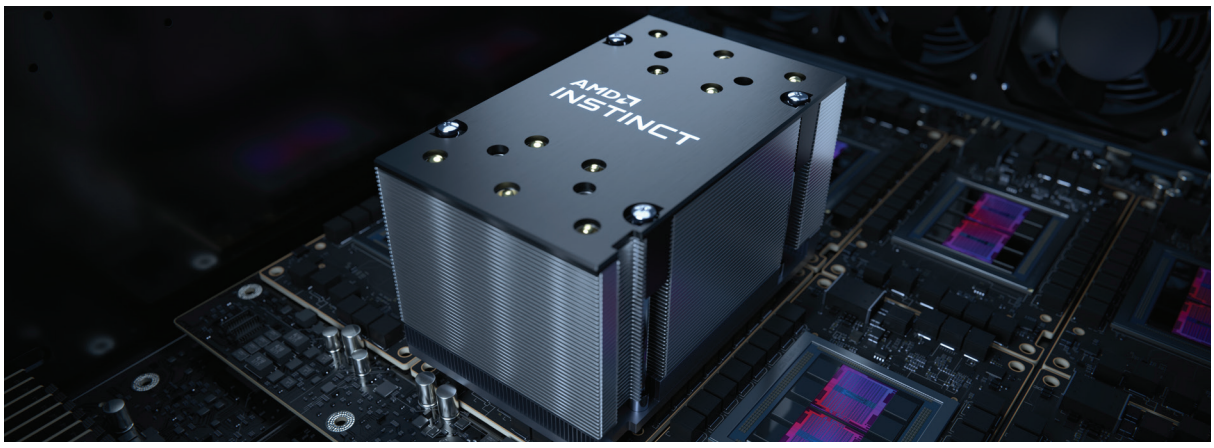

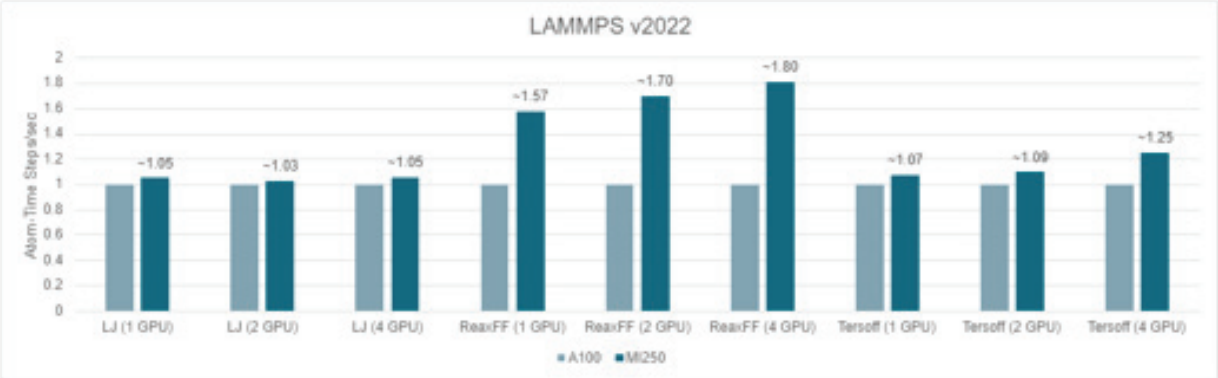

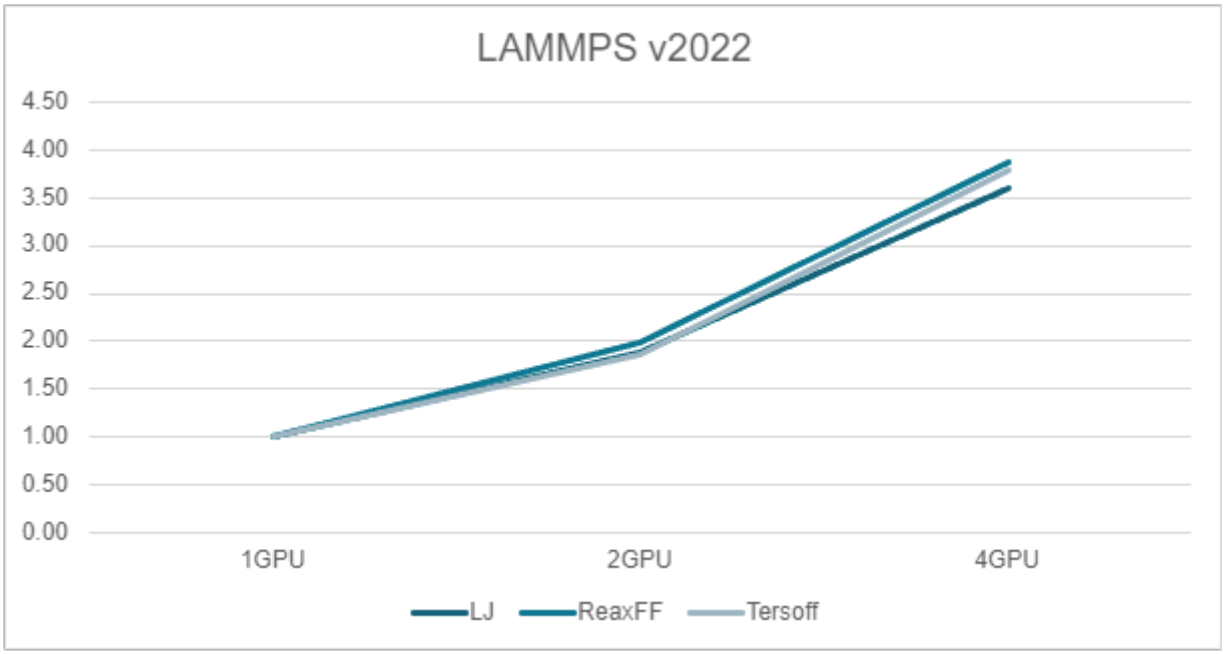

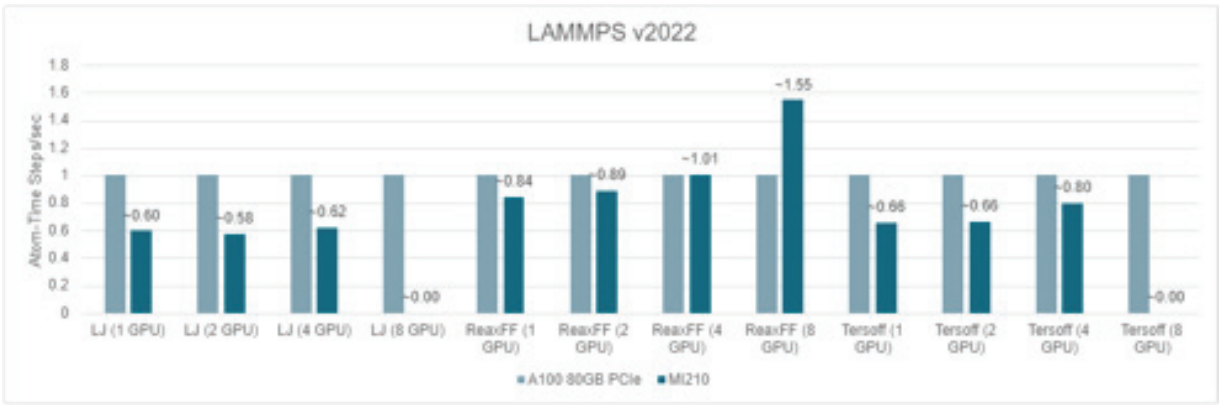

[1] MI200-77: Testing Conducted by AMD performance lab 109.0322.22 using LAMMPS: EAM, LJ, ReaxFF and Tersoff comparing two systems: 2P EPYC™ 7763 powered server with 4x AMD Instinct™ MI250 (128GB HBM2e) 560W GPU, ROCm 5.2.0, LAMMPS container 2021.5.14_121amdih/lammps:2022.5.04_130 vs. Nvidia public claims http://web.archive.org/web/20220718053400/https://developer.nvidia.com/hpc-application-performance. (stable 23Jun2022 update 1). Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations.

[2] MI200-47A: Testing Conducted on 10.03.2022 by AMD performance lab on a 2P socket AMD EPYC™ '7763 CPU Supermicro 4124 with 4x and 8x AMD Instinct™ MI210 GPUs (PCIe® 64GB,300W) with AMD Infinity Fabric™ technology enabled. SBIOS2.2, Ubuntu® 18.04.6 LTS, host ROCm™ 5.2.0. LAMMPS container amdih-2022.5.04_130 (ROCm 5.1) versus Nvidia public claims for 4x and 8x A100 PCIe 80GB GPUs http://web.archive.org/web/20220718053400/https://developer.nvidia.com/hpc-application-performance. Server manufacturers may vary configurations, yielding different results. Performance may vary based on use of latest drivers and optimizations.