Breaking Barriers in Plasma Physics with PIConGPU and AMD Instinct™ MI250 GPU

June, 2023

Application:
PIConGPU for plasma physics simulations

Challenge:
PIConGPU requires high-performance processing and large memory capacity to model complex plasma and laser interactions accurately. Efficient data handling and energy usage are also crucial for running extensive and complex simulations.

Solution:
AMD Instinct™ MI250 GPU

Benefits:
The AMD Instinct MI250 GPU, with its high computational power and large memory capacity, enables PIConGPU to run complex plasma simulations efficiently. Its energy-efficient design and compatibility with various software APIs and frameworks enhance its suitability for PIConGPU.

PIConGPU, a pioneering plasma physics simulation system, is at the forefront of exploring the intricate dynamics of plasma and laser interactions. This fully relativistic, many GPGPU, 3D3V particle-in-cell (PIC) code is optimized for GPU acceleration, enabling it to run complex simulations at an unprecedented scale. The AMD Instinct MI250 GPU has significantly driven forward the system's capabilities, a computational powerhouse that has proven instrumental in driving groundbreaking research in plasma physics.

The AMD Instinct MI250 GPU, part of the innovative AMD Instinct MI200 family of accelerators, is designed with high-performance computing in mind. Its large 128GB memory capacity, exceptional performance, and energy-efficient design make it ideally suited to handle the extensive and complex simulations run by PIConGPU. Built on AMD CDNA™ 2 architecture, the MI250 delivers a peak half-precision performance of 362.1 TFLOPs, providing the computational muscle required for PIConGPU's massive and highly-detailed simulations.

The MI250 GPU's capabilities extend beyond raw power. It excels in handling the large amounts of data generated by PIConGPU simulations, thanks to its dedicated HBM2e memory and high memory bandwidth. This efficient data handling is critical in enabling PIConGPU to push the boundaries of our understanding of plasma and laser interactions.

The AMD Instinct family, including the MI250, offers a range of GPUs to cater to diverse accelerator requirements. Whether the need is for a single GPU or multiple GPUs per system, AMD Instinct accelerators provide the flexibility and performance necessary for groundbreaking exploration in plasma physics. The MI250, with its extreme performance and robust architecture, is more than just a tool; it's a partner that enables the PIConGPU team to pioneer new frontiers in plasma physics.

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.	Up to 128GB High bandwidth HBM2e memory with ECC support at a clock rate of 1.6 GHz. and deliver an ultrahigh 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.

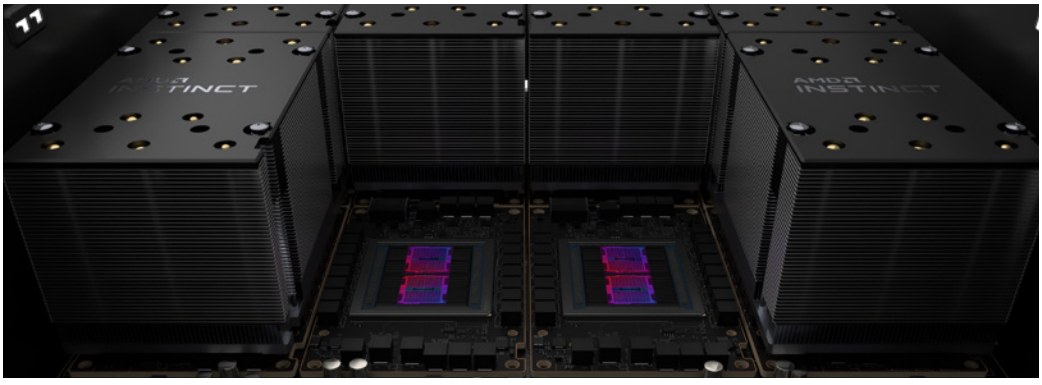

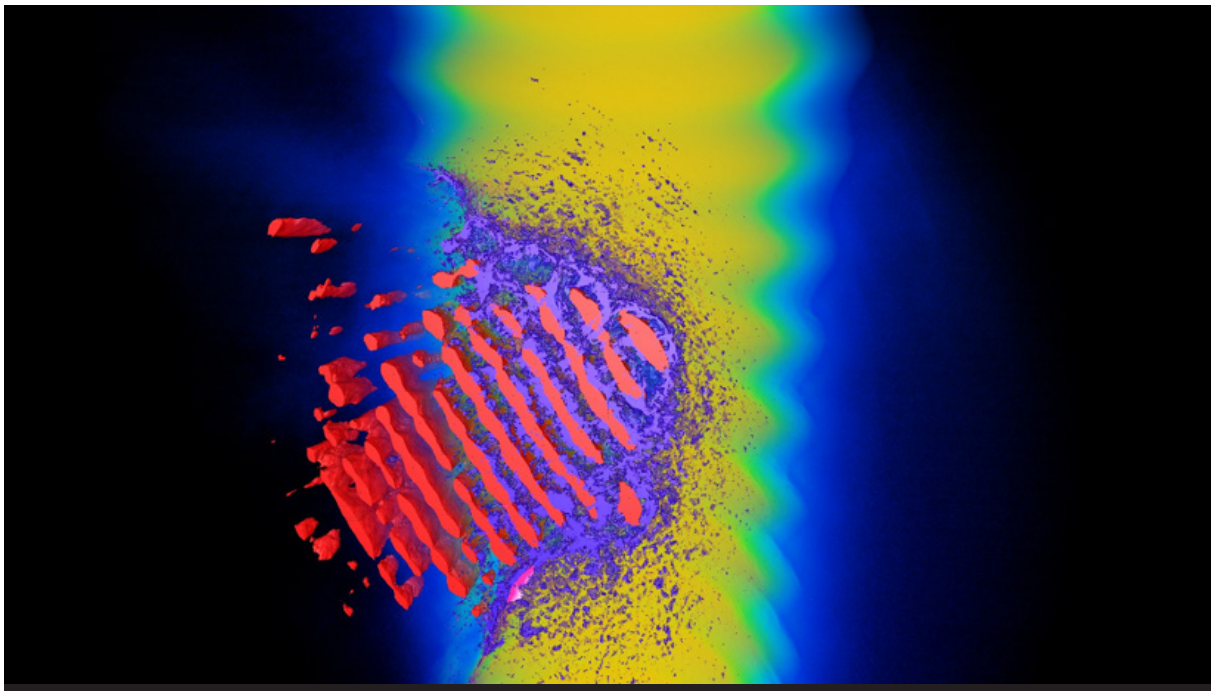

Image and video credit: Richard Pausch and René Widera @ (HZDR)

PIConGPU simulation of a cryogenic hydrogen jet interacting with a high-intensity laser beam, a model of an experiment recently performed at Helmholtz-Zentrum Dresden-Rossendorf (HZDR) at the DRACO laser system.

AMD INSTINCT MI250 GPU BRINGS OUTSTANDING PERFORMANCE, EFFICIENCY, AND FLEXIBILITY TO DRIVE PICONGPU SIMULATION

PERFORMANCE. The AMD Instinct MI250 GPU, built on the 2nd Gen AMD CDNA™ architecture, features 13,312 stream processors and 208 compute units, boosting computational capabilities and throughput. The architecture supports a wide range of data types and applications, offering full-rate double precision and new FP64 Matrix operations for HPC workloads. This high performance is crucial for PIConGPU simulations, which require extensive computational power to model the intricate dynamics of plasma and laser interactions accurately.

EFFICIENCY. The MI250 GPU, part of the AMD Instinct accelerator family, is designed on the AMD CDNA architecture, leveraging the performance, power, and efficiency advantages of the 6nm fabrication process. This efficiency is critical for PIConGPU, as the system needs to simultaneously utilize substantial amounts of compute resources without drawing excessive power.

FLEXIBILITY. The ROCm™ 5.0 open platform, operating on top of AMD's HIP and ROCm, extends AMD's commitment to HPC and AI. It optimizes at the device level with a freely shared codebase that can evolve as systems and their parts evolve. This flexibility is essential for PIConGPU, which uses the Alpaka backend and the Plasma-in-Cell (PIC) algorithm for its science case simulations, requiring a flexible and evolving software ecosystem. The combination of the AMD MI250 GPUs, Alpaka, and the ISAAC plug-in enables real-time steering of PIConGPU simulations, facilitating remote control from anywhere in the world.

AMD INSTINCT ACCELERATORS FOR PLASMA PHYSICS SIMULATIONS

Core IPC (instructions per cycle) is a critical factor in optimizing the performance of PIConGPU. AMD Instinct MI250 GPUs, built on the revolutionary 2nd Generation AMD CDNA™ architecture, ensure that you get the most out of your system, minimizing execution time and increasing overall utilization of your deployment. This optimization is crucial for PIConGPU, which requires high performance to model complex plasma and laser interactions accurately.

Ultra-Fast Memory. The AMD Instinct MI250 GPU offers a large memory capacity of 128GB of high-bandwidth HBM2e memory with ECC support. This ultra-high memory capacity is essential for PIConGPU, which handles extensive and complex simulations that generate large amounts of data.

Balancing performance vs. per-core license costs is a common challenge in High-Performance Compute (HPC) workloads. 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. This consistency benefits PIConGPU, which requires a stable and reliable software ecosystem for its science case simulations.

AMD Infinity Fabric™ Link Technology. As workloads demand more data processing capabilities, the communications between Instinct GPUs become critical to efficiently solving customers' complex problems. AMD Instinct MI250 GPUs feature advanced peer-to-peer I/O connectivity through AMD Infinity Fabric™ links, creating a system capable of delivering high I/O bandwidth performance. This advanced connectivity is crucial for PIConGPU, which engages in data movement and communication as part of its simulation process.

Energy Efficiency: PIConGPU simulations require substantial computational resources, which must be used efficiently without drawing excessive power. The AMD Instinct MI250 GPU is designed with energy efficiency in mind, which is a critical factor for modern HPC systems like PIConGPU.

Software Compatibility: PIConGPU uses the Alpaka backend and operates on top of AMD's HIP and ROCm. The AMD Instinct MI250 GPU's compatibility with these software APIs and frameworks allows the PIConGPU team to leverage a wide range of tools and technologies in their research.

AMD INSTINCT MI200 SERIES ACCELERATOR PRODUCT OPTIONS

	AMD Instinct MI100	AMD Instinct MI210	AMD Instinct MI250
GPU Architecture	AMD CDNA	AMD CDNA2	AMD 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

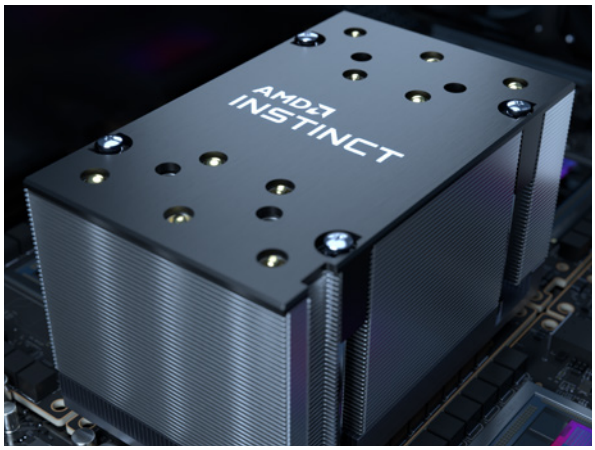

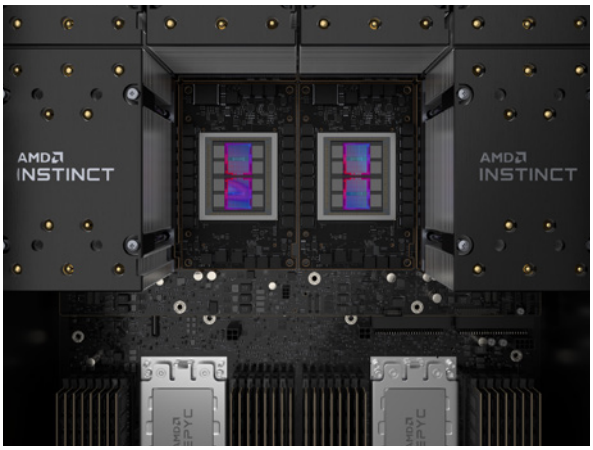

MORE ABOUT PICONGPU

PIConGPU is a state-of-the-art plasma physics simulation system leveraging the power of GPU acceleration to model complex plasma and laser interactions. It utilizes the Alpaka backend and the Plasma-in-Cell (PIC) algorithm, ensuring high performance and portability across various accelerators. Operating on AMD's HIP and ROCm platforms, PIConGPU provides an open-source ecosystem that encourages collaboration and innovation. This robust tool is freely distributed with source code, fostering continuous development and adaptation to meet the evolving needs of the scientific community.

© 2023 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, AMD CDNA, AMD Infinity Fabric, AMD Instinct, ROCm, and combinations thereof are trademarks of Advanced Micro Devices, Inc. Other product names used in this publication are for identification purposes only and may be trademarks of their respective owners.

[Insert third party trademark attribution here per AMD's Guidelines on Using Trademark Notice and Attribution.]