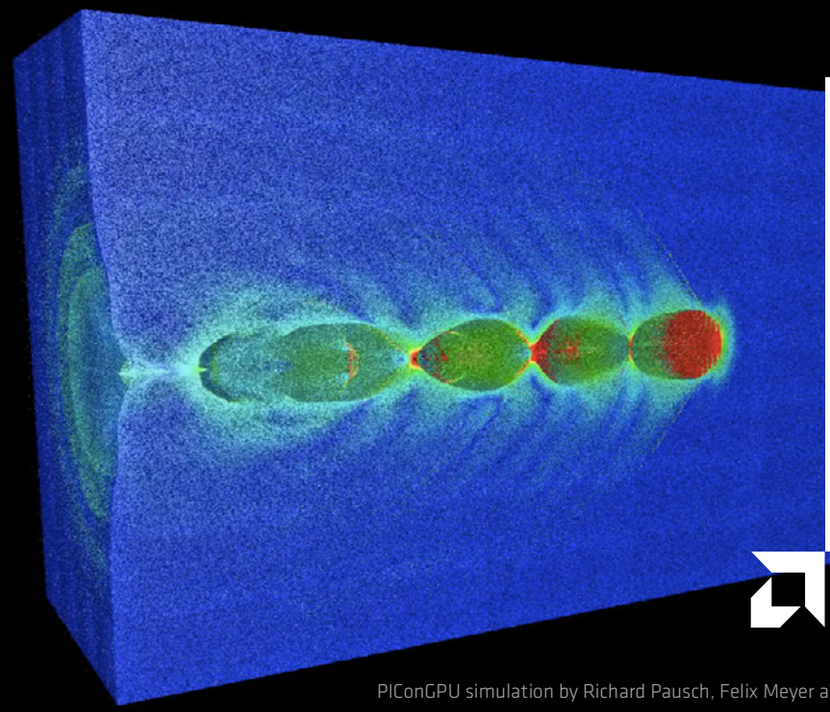

PIConGPU simulation by Richard Pausch, Felix Meyer and team from HZDR and Benjamin Hernandez (ORNL) on Summit

ORNL expands possibilities for plasma physics with open and portable AMD ROCm

Increased fidelity and longer simulations using GPU-accelerated software from AMD

AMD EPYC | INSTINCT

CUSTOMER

OAK RIDGE National Laboratory

INDUSTRY

Scientific research

CHALLENGES

Being able to run larger, longer simulations with PIConGPU plasma physics simulations when Frontier supercomputer comes online

SOLUTION

Deploy AMD Instinct™ MI100 GPUs with ROCm™ open software to prepare for Frontier supercomputer

RESULTS

Easy porting to AMD Instinct GPUs with ROCm and HIP

AMD TECHNOLOGY AT A GLANCE

AMD Instinct MI100 GPUs

TECHNOLOGY PARTNER

Cray

Oak Ridge National Laboratory (ORNL) has been home to some of the most powerful computational resources used in science for decades. It currently hosts the second most powerful supercomputer in the world, Summit. ORNL has commission a new supercomputer. Powered by AMD CPUs and GPUs, Frontier is expected to be one of the fastest systems in the world when it arrives in 2021. ORNL is readying several key scientific applications that can take advantage of the exascale power the new system will deliver via its Center for Accelerated Application Readiness (CAAR).

Amongst these applications is Particle-In-Cell GPU (PIConGPU). "We are one of the eight teams across the States chosen for the CAAR project to prepare the software stack for Frontier," says Sunita Chandrasekaran, Assistant Professor of Computer and Information Sciences at the University of Delaware. She is the Principal Investigator of an international collaboration with HZDR and CASUS from Germany, making PIConGPU, a plasma physics simulation system, ready for Frontier. Its main relevance is for applications in the radiotherapy of cancer, as well as in probing the structure of matter via X-rays in material and life sciences. "We want to be able to generate fast and predictive simulations for the next generation user plasma accelerators." The huge number of super-fast next-generation AMD GPUs in Frontier promise to deliver on this endeavor and advance the boundaries of science by another order of magnitude.

Increasing the kinetic energy of particle acceleration

GPUs have been essential for the computational capability required by PIConGPU, with every advance in computing performance and memory throughput providing a step forward. "For more realistic science cases, higher fidelity models, and more density simulations, we need more compute power per GPU," says Ronnie Chatterjee, Oak Ridge Leadership Computing Facility (OLCF) liaison for PIConGPU. This is because the research team wants to support particle accelerators capable of producing a greater level of kinetic energy than current instruments.

"Frontier's AMD Instinct GPU computing power will enable us to find answers to questions not accessible before."

Sunita Chandrasekaran, Assistant Professor of Computer and Information Sciences at the University of Delaware

"We want to get to the next generation of compact accelerators," says Dr Alexander Debus, Institute of Radiation Physics, Researcher at Helmholtz Zentrum Dresden Rossendorf. "The single-stage is a couple of centimeters long and is limited to 10 GeV. We want to get beyond that. Ideally, we want to reach 100s or 1000s of GeV. We have a couple of sophisticated approaches that can be used, which require a lot of computational power, because it's not just a large simulation, but it's also a long simulation."

In readiness for the huge boost in computational capabilities to come with Frontier, the PIConGPU team has been porting its code to from CUDA to run on AMD GPUs using the HIP programing model. But this wasn't the group's first experience with HIP.

"We have been using HIP developer tools, even prior to starting off with CAAR project, because PIConGPU has been written using an affordable software stack called Alpaka, which targets HIP," says Chandrasekaran. "The developers have therefore already been exposed to HIP, prior to the CAAR project, and the ROCm stack from AMD."

Easy performance improvements with AMD HIP

The HIP open source implementation fits perfectly with scientific workloads, so that researchers can more easily contribute improvements to the code. "The entire software stack for PIConGPU is open source," says Chandrasekaran. "This is important because then we can work with both the computer scientists as well as the domain science community. We are very appreciative of the fact that AMD has the ROCm™ open source software stack with the HIP programing model. In contrast, NVIDIA's CUDA is a black box, so we don't know what's happening under the hood, and that's not very welcoming to researchers like us. We want to dig into the bottom of it."

"We are very appreciative of the fact that AMD has the ROCm open source software stack with the HIP programing model."

Sunita Chandrasekaran, Assistant Professor of Computer and Information Sciences at the University of Delaware

"People from the outside are not doing bug fixes in CUDA," adds Chatterjee. "Having the opportunity to fix bugs with HIP gives not just us, but also AMD, the ability to grow as we are working together closely towards Frontier. It feeds back into the ROCm stack for the next releases.. There are specific changes that have happened as a result. It's not just about us proposing fixes. We're also getting these accepted by the AMD side."

"When you're looking at trying to access half the largest machines in the whole world, this sort of exercise is paramount," says Michael Bussmann from the Center for Advanced Systems Understanding at HZDR responsible for assuring performance portability for the software stack underlying PIConGPU. "The project is not just improving PIConGPU, or the seven other apps. It helps other scientific codes that will start using Frontier. This portability is helping us fit PIConGPU on Summit and Frontier, without trying to redo the entire software stack. Our thanks to the AMD team for making this possible."

Faster perfomance on AMD Instinct GPUs, ready for Frontier

As part of the readying process, PIConGPU code has been performance-tested on ORNL's Summit supercomputer and Poplar, a Cray Center of Excellence system configured for CAAR. Summit runs NVIDIA Tesla V100 GPUs exclusively, but Poplar offers both these and AMD GPUs, including AMD Radeon Instinct™ MI50 and AMD Instinct™ MI100 accelerators. Not only was compatibility with AMD GPUs demonstrated on both GPU vendor's hardware, but also significant improvements in performance with the AMD GPUs. "We ran the full PIConGPU on Poplar," says Chatterjee. "But we ran a smaller science data set, because we don't have the full system. We did not see any challenges. We are currently seeing about 1.4X speed up from the MI60s to the MI100s." The PIConGPU teams saw virtually no difference in performance when running the PIConGPU via HIP on NVIDIA Tesla V100 GPUs when compared to its CUDA counterpart on the same GPUs, using the Summit supercomputer, showing the efficiency of the software stack.

This shows great promise for when Frontier comes online in 2021. "We believe with Frontier we will be able to run advanced laser-plasma accelerator simulations that investigate how to achieve high electron energies, such as the Traveling-wave Electron accelerator scheme (TWEAC)," says Chandrasekaran. "These simulations in PIConGPU are not only large, but also run much longer, which at the required scale wouldn't be possible with Summit. Frontier's AMD Instinct GPU computing power will enable us to find answers to questions not accessible before."

About Oak Ridge National Laboratory (ORNL)

Oak Ridge National Laboratory is the largest US Department of Energy science and energy laboratory, conducting basic and applied research to deliver transformative solutions to compelling problems in energy and security. ORNL's diverse capabilities span a broad range of scientific and engineering disciplines, enabling the Laboratory to explore fundamental science challenges and to carry out the research needed to accelerate the delivery of solutions to the marketplace. Within ORNL, the Oak Ridge Leadership Computing Facility was established to accelerate scientific discovery and engineering progress by providing world-leading computational performance and advanced data infrastructure, including Summit, the second most powerful supercomputer in the world. For more information, visit ornl.gov.

About AMD

For 50 years AMD has driven innovation in high-performance computing, graphics, and visualization technologies—the building blocks for gaming, immersive platforms, and the data center. Hundreds of millions of consumers, leading Fortune 500 businesses, and cutting-edge scientific research facilities around the world rely on AMD technology daily to improve how they live, work, and play. AMD employees around the world are focused on building great products that push the boundaries of what is possible. For more information about how AMD is enabling today and inspiring tomorrow visit amd.com/Instinct.