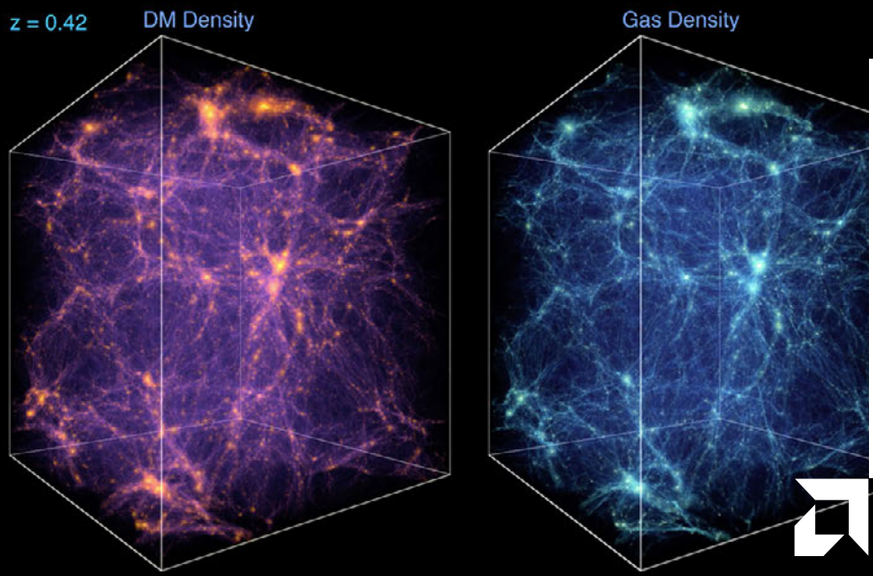

ORNL advances galaxy simulations with AMD Instinct™ MI100 GPUs

Accelerating galaxy formation research at unprecedented scale and resolution using AMD GPU-powered supercomputing

AMD EPYC | INSTINCT

CUSTOMER

OAK RIDGE National Laboratory

INDUSTRY

Scientific research

CHALLENGES

Increasing scale and resolution with numerical simulation of galactic evolution with CHOLLA software

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

Supercomputing is revolutionizing science, with the fastest systems in the world providing unparalleled path to discovery. Oak Ridge National Laboratory (ORNL) is at the forefront of this trend, and already hosts one of the fastest supercomputer in the world. But ORNL has commissioned a new supercomputer for its Leadership Computing Facility (OLCF). Called Frontier, it's expected to be the fastest open science supercomputer in the world when it arrives in 2021, and one of the first to offer exascale computing power of 1 exaFLOPS or more. It will also have both AMD CPUs and AMD GPUs at its heart.

"Having access to this exascale machine is a game changer for the kinds of problems that we can simulate."

Evan Schneider, Assistant Professor of Physics and Astronomy at the University of Pittsburgh

ORNL is preparing eight key scientific applications for Frontier, and one of them is CHOLLA, which investigates astrophysics and galaxy formations, one of the first workloads in its class to run on GPUs. The researchers behind this code have ported it from CUDA to run on AMD GPUs, in preparation for Frontier. They have been able to do this with ease, and are already finding considerable performance improvements, with much more promised when Frontier comes online. The potential benefits for our understanding of the universe are immense.

Harnessing GPUs to explore galaxy formation

"Astronomy is an observational science," says Evan Schneider, Assistant Professor of Physics and Astronomy at the University of Pittsburgh, and one of the chief architects of CHOLLA. "But a revolution in astronomy over the last 40 years has been our ability to use numerical simulations to try to understand how the universe is evolving. Unlike most physical sciences where you can conduct experiments, and the time scales of the experiments happen in relevant human lifetimes, in astronomy, things change on much longer timescales. The only way we can get a moving picture of things is by conducting numerical simulations."

CHOLLA was created to provide this time-based analysis, particularly focusing on galaxy formation, with its processing accelerated by GPUs. "The universe is mostly composed of gas and dark matter," says Schneider. "So the two most relevant pieces of physics are fluid dynamics, which is how we understand the evolution of gas, and gravity, which is how we understand the evolution of dark matter. Those are the two pieces of physics that our simulation software needs to include."

CHOLLA started off with simulating fluid dynamics. "CHOLLA stands for Computational Hydrodynamics on Parallel Architecture," says Reuben Budiardja, Computational Scientist at ORNL. "It has now been extended to have self-gravity and particle-based gravity, with magnetic field also in development. The idea for this project is to run simulations of a Milky Way-like galaxy at a resolution that's sufficiently high so

that we can also incorporate star formation." This is where GPU power provides the boost required for truly revelatory science. "The more physics we put in, the more computationally intensive it is," says Schneider. "It's really a nice match between GPUs and our end goals since GPUs are happy with more computational work."

Solving the problem of scale with AMD Instinct™ GPUs

"GPUs are good at massively parallel local problems," explains Schneider. "The fundamental problem in computational astrophysics is one of scale. The big picture evolution of a galaxy like the Milky Way depends on processes that are happening at very small scale. In order to understand how that evolves over time, we need to account for the formation of stars and the explosion of supernova within the galaxy and those explosions take place on much smaller scales than the whole galaxy. One of the things that we would like to do is have simulations with higher resolution."

"We got most of the porting [of CHOLLA] to HIP to run on AMD hardware done in a few hours."

Reuben Budiardja, Computational Scientist at ORNL

"If you don't have enough resolution, you just have to gloss over the star formation," says Budiardja. "If we want to simulate the Milky Way, its diameter is about 50,000 parsecs. Then typically star clusters have about a few parsecs as their diameters. If you divide these two numbers, then you get about 10,000 cells. Then you have three dimensions, so 10,000 cubed is the total number of cells in the computational domain. So now we are trying to evolve 10,000 cubed cells over 500 million years of evolution for the whole galaxy to form. Our timestep resolves to one kiloyear giving you a million total timesteps. Then the question becomes, how fast can you do one timestep to update about 10,000 cubed cells? That gives you a sense of how long the simulation will need to run."

Porting the CHOLLA code over from existing NVIDIA CUDA platform to AMD ROCm™ open software platform and both AMD and Radeon Instinct™ GPUs proved easy and provided immediate performance gains. "We got most of the porting to HIP to run on AMD hardware done in a few hours," says Budiardja. "With the AMD Radeon Instinct MI50, we've got relatively similar performance to the NVIDIA Tesla V100. With the AMD Instinct MI100, we realized about 1.4x speed up without doing anything at all. We just changed to compile to ROCm to get the code to run. There is a lot of benefit in HIP using similar function calls. If you're already familiar with CUDA, you can take one look at the HIP function and realize what it's doing."

Exascale GPU supercomputing for revolutionary discovery

With the further performance that will come with the next-generation AMD GPUs in Frontier, Budiardja expects the necessary scale for CHOLLA to be well within reach. "Coming back to the original target problem with 10,000 cubed cells of Milky Way galaxy, we are very optimistic that we can do that when Frontier is ready," he says. "We are now even exploring possibilities to run something larger. This will be the first time-simulation of this scale on a Milky Way-like galaxy ever done. It will be unprecedented to be able to run this size of simulation with this resolution and with all the physics that's required for the fidelity we need."

"Observational surveys of the Milky Way now have better resolution than any of our simulations," says Schneider. "One of our ways of advancing the field is to compare observational data to simulation data and if there is a mismatch, then you understand that you are missing some piece of physics that's important. But if you don't have simulations at the resolution required to compare to the observational data, that's a hard comparison to make." Schneider expects the AMD Instinct GPUs in the ORNL Frontier supercomputer to fully deliver what's needed, based on experience with the MI100. "The hardware upgrades are just immediate. If you can make a hardware that runs as twice as fast, then my simulation code runs twice as fast. That means that I can run higher resolution simulations. Having access to this exascale machine is a game changer for the kinds of problems that we can simulate."

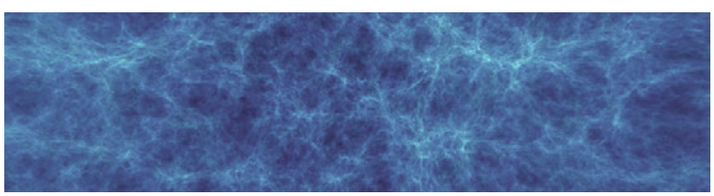

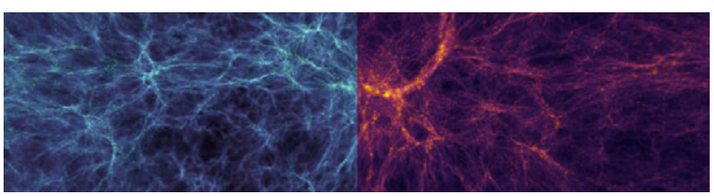

About ORNL

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.