



ORNL CoMet team explores complex biological systems with AMD Instinct™ GPUs

AMD GPU-powered supercomputing enables study of biological systems from the molecular level to planetary scale.



CUSTOMER



INDUSTRY

Scientific Research

CHALLENGES

Power some of the most computationally intensive algorithms to pursue study of complex biological systems at the highest possible resolution

SOLUTION

Deploy the newest AMD Instinct GPUs with HIP API

RESULTS

Based on early testing, the projection is that Frontier may provide a four or fivefold increase in what could be achieved using in the previous architecture

AMD TECHNOLOGY AT A GLANCE

AMD Instinct MI200 accelerator

TECHNOLOGY PARTNER



**Hewlett Packard
Enterprise**

High-performance computing is the engine that powers modern scientific discovery. Now one of the world's newest exascale supercomputers, the AMD GPU-powered Frontier system at the Oak Ridge Leadership Computing Facility (OLCF) established at Oak Ridge National Laboratory (ORNL), is poised to come online. In preparation, scientists, application developers, and their vendor partners are hard at work optimizing code and preparing workloads that will make the most of the new system as they drive innovation in science and technology.

"Bigger, faster machines like Frontier [with AMD Instinct GPUs] allow us to pursue higher resolution and answer more sophisticated questions."

*Daniel Jacobson,
Computational Systems
Biologist, ORNL*

Among those efforts are eight research projects selected by the Frontier Center for Accelerated Application Readiness (CAAR) program. One of those projects is a comparative genomics application called CoMet, short for Combinatorial Metrics, and it is led by computational systems biologist Daniel Jacobson.

Staying ahead of the biggest questions

CoMet first emerged on the Titan and Summit supercomputers, and it is now being prepared for the next-generation Frontier system. This suite of algorithms based on comparative genetics code was designed to analyze biological challenges at scales that have never been tried before. This includes looking at populations of genomes at every point of land on the planet and trying to compare them all to each other. Such analysis creates challenging questions that require the evaluation of extremely long vectors.

"We made a big effort many years ago to port higher-order, triple vector versions of these algorithms to run on GPUs in order to accelerate these sorts of calculations," said Jacobson. More recently, "the Advanced Readiness program allows us to get early access to nodes based on the latest AMD Instinct GPUs before Frontier even exists so that we can port CoMet to that hardware, put it through its paces, and be ready to hit the ground running."

According to Jacobson, CoMet has historically been one of the most demanding algorithms run on previous generation supercomputers. This makes CoMet particularly useful to OLCF as it seeks to test the capabilities of Frontier. "CoMet computations are massive and intense. We soak up as much computer power as is available. We deal with multi-scale looks at systems going from atomistic and sometimes even electron level detail inside cells up through ecosystems and planetary-scale calculations. Bigger, faster machines like Frontier allow us to pursue higher resolution and answer more sophisticated questions."

The Frontier supercomputer was designed by a collaboration between the U.S. Department of Energy, Oak Ridge National Laboratory, HPE, and AMD. Frontier is expected to deliver more than 1.5 exaflops of peak processing power. It leverages the AMD Infinity Architecture for high-speed data transfer between GPUs, CPUs, and memory to help ensure low latency at the scale of this unique system.

Open source provides future-ready portability

Jacobson's team has already ported and optimized CoMet code to run on AMD Instinct GPUs using the open-source Heterogeneous Computing Interface for Portability, or HIP. One of the benefits of converting to HIP is that unlike previous CUDA versions of code that could only run on Nvidia GPUs, the same source code is now portable between GPUs. HIP is also the native format for AMD ROCm™, the open, programming language-independent, exascale-class platform for accelerated computing. ROCm provides coding flexibility that enables developers to select and employ the best tools and run-time language for specific applications while making the most of the hardware performance to meet incredibly demanding workloads.

Breakthroughs that impact everybody

Recently the team has studied SARS-CoV-2, the virus behind the Covid-19 pandemic. While most of the scientific community has focused on the spike protein that defines coronaviruses, Jacobson points out, "there are 27 other proteins in that viral genome, all with important functions. We've modeled every protein in the viral genome. Using molecular dynamics and machine learning approaches, we can then understand the specific changes in the virus and their functional impact on viral biology and variants of concern. We can start to explain how they're behaving differently by those structural changes. Such molecular dynamics is computationally expensive and intense, and something that works really well on GPUs."

"Based on our early testing of the newest AMD Instinct GPUs, the projection is that Frontier will be another step function in power. We may come close to a four or fivefold increase in what we could do before in the previous architecture."

Daniel Jacobson, Computational Systems Biologist, ORNL

That provides a clue as to the promise Frontier and its massive GPU computational capacity holds. "Each machine architecture that we've forwarded code to has been a step function in compute power that allows us to tackle these problems," says Jacobson.

"Based on our early testing of the newest AMD Instinct GPUs, the projection is that Frontier will be another step function in power. We may come close to a four or fivefold increase in what we could do before in the previous architecture. That means a lot when you're tackling these sorts of problems. Increasing compute power by hundreds of percent means the sort of problems we can tackle become larger and more sophisticated and better biology coming out of them."

"Scientifically, the end goal is to really understand biological systems from incredibly fine levels of detail up through how they interact with the environment," says Jacobson. "We want to solve problems in agriculture. We want to solve problems in precision health. And we want to solve problems in pandemic prevention. We're pushing hard to build new scientific programs that are very compute hungry. They will integrate molecular level detail and viruses, and cells, up through ecosystem and planetary-scale information. We want to include remote sensing, satellite platforms, every level of detail across the board, so we can learn such things as the rules that lead to zoonotic spillovers, that lead to outbreaks, that lead to pandemics. We're seeking to solve some of the world's biggest challenges. We know that Frontier powered by AMD Instinct GPUs will play a critical role in helping us do that."



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](https://www.ornl.gov).

About AMD

For over 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.

All performance and cost savings claims are provided by ORNL and have not been independently verified by AMD. Performance and cost benefits are impacted by a variety of variables. Results herein are specific to ORNL and may not be typical. GD-181

©2021 Advanced Micro Devices, Inc. All rights reserved. AMD, the AMD Arrow logo, EPYC, 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 companies.