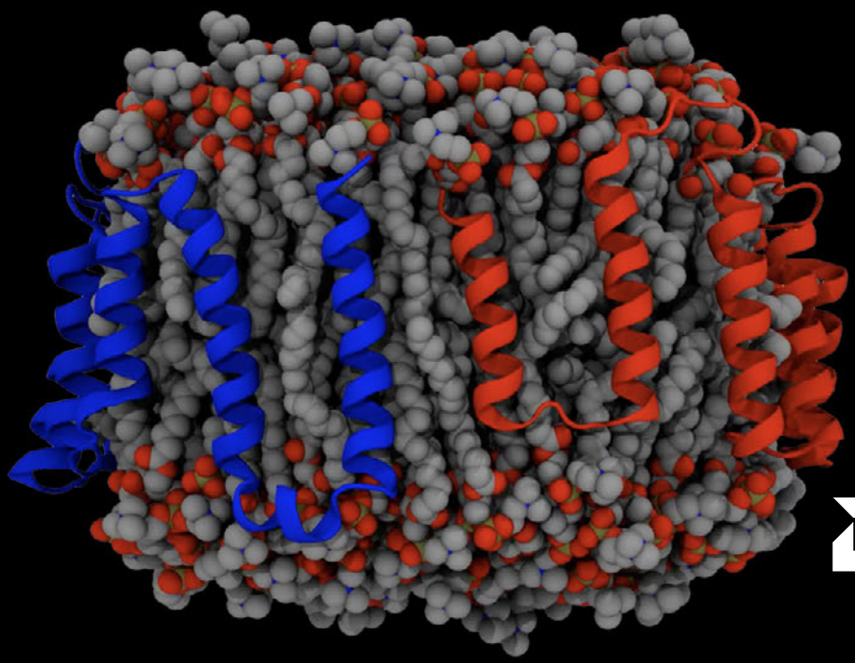

ORNL accelerates molecular dynamics with AMD Instinct™ MI100 GPUs

Large-scale simulation of entire cells and viruses within reach with NAMD and AMD GPU-powered supercomputing

AMD EPYC | INSTINCT

CUSTOMER

OAK RIDGE National Laboratory

INDUSTRY

Scientific research

CHALLENGES

Increase the size of molecules over greater durations of change that can be simulated

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) is getting ready to host the Frontier supercomputer, expected to be one of the first exascale machines, capable of over 1.5 exaFLOPS processing power when it comes online in 2021. Frontier's unprecedented level of computing will push the limits of scale in today's scientific applications, so ORNL is preparing key workloads to take advantage of the new capabilities via its Center for Accelerated Application Readiness (CAAR).

One of the eight workloads chosen by CAAR is NAMD, which stands for Nanoscale Molecular Dynamics. However, while NAMD has been targeting GPUs for over a decade, its focus has been to build with CUDA. Yet Frontier will be built on AMD CPUs, GPUs, and the ROCm™ software platform, not on the proprietary CUDA software stack. The NAMD team together with ORNL scientists has been porting their code to run on AMD GPUs, with some very promising results for future performance on Frontier.

"You can get a threefold increase in performance for NAMD 2.14 with the AMD Instinct™ MI100 compared to the NVIDIA Tesla VI100."

Josh Vermaas, Computational Scientist at ORNL

Pioneering molecular dynamics simulation with GPUs

"NAMD is a molecular dynamics simulation engine," explains Josh Vermaas, Computational Scientist at ORNL. "It is a computational microscope. Experimentally this is hard to do because you can't look at the individual atoms and molecules. The technology doesn't exist, physics prevents you from seeing it, but by making computational models of what's going on at the molecular level instead, we can have better understanding of the processes that are underlying biology and that's what NAMD lets you do."

This has been particularly important in the recent pandemic. "NAMD is really good at simulating large systems at impressive scales, such as whole cells, or the recent work that's being done with simulating the entire COVID virion," continues Vermaas. Getting NAMD running on the exascale computing power of Frontier will pay huge dividends, thanks to its huge complement of AMD GPUs. "For molecular dynamics, GPUs have increasingly become the de facto architecture, achieving impressive sampling speeds," adds Julio Maia, Research Programmer at the University of Illinois. "NAMD has been GPU enabled since 2007, so it was one of the earliest to use GPU acceleration." So far NAMD has employed a combination of GPUs and CPUs, offloading the computationally intensive parts to the GPUs. "That scheme allows us to scale across thousands and thousands of GPUs right now," continues Maia.

However, in order to take advantage of the scale of Frontier, adjustment will be required. "The way that NAMD performs overall on heterogeneous architectures has become limited by CPUs, so we've been working on moving everything to the GPU," says David Hardy, Senior Research Programmer at the University of Illinois. "This is important because NAMD has a couple of different ways that it could scale to use these parallel resources. It can scale a large simulation across multiple nodes, or run multiple copies of a smaller simulation."

"If you're trying to figure out what is a population distribution between two states of a protein, those are the kinds of things that multiple copies could be helpful with, by stringing together the transition path between the two states," says Vermaas. "A single protein is never going to scale to all of Frontier, but doing one of those calculations where you have one protein over many, many replicas, that would. One really big simulation of a whole cell is also something that's Frontier sized. NAMD has to be able to support both because they're both scientifically useful to answer questions that biophysicists are interested in."

Porting NAMD from CUDA to HIP

The NAMD team has found the AMD HIP environment extremely useful to prepare the application for AMD GPUs and the massive scale of Frontier. "HIP has helped us out tremendously," says Hardy. "We have been a CUDA-enabled code since 2007 when CUDA first came out. HIP is very important for us to be able to port the NAMD kernels to make use of AMD GPUs. So far, we're impressed by the performance of the AMD GPUs."

"HIP is very important for us to be able to port the NAMD kernels to make use of AMD GPUs. So far, we're impressed by the performance of the AMD GPUs."

David Hardy, Senior Research Programmer at the University of Illinois

"HIP implements a lot of the features that CUDA has and does so in a syntax that's very similar to CUDA," says Vermaas. "Porting across is a very straightforward process."
The NAMD team has been able to compare performance across NVIDIA and AMD GPUs using the Cray Center of Excellence Poplar system, which offers both vendors' hardware on the same platform. Benchmark tests were run using simulations of the Apolipoprotein A1 (APOA1) protein and Satellite Tobacco Mosaic Virus (STMV) using a single Poplar node with four GPUs, either four NVIDIA Tesla V100s, four AMD Radeon Instinct MI50s or four AMD Instinct™ MI100s. "Memory bandwidth has substantially improved going from a V100 to an MI50 or MI100," says Vermaas. "You can get a threefold increase in performance for NAMD 2.14 with the AMD Instinct MI100 compared to the NVIDIA Tesla VI100."

Enabling larger-scale cell simulation with AMD Instinct GPUs

This performance improvement per node, combined with Frontier's massive number of nodes, promises an increase in the dynamic time periods that can be simulated with NAMD, even compared to ORNL's Summit, the second fastest supercomputer in the world. "For the biggest, billion-atom systems right now, you might get a nano second, or a couple of nanoseconds, a day," says Vermaas. "Even if you're using a lot of Summit, it is functionally impossible to get meaningful results because you can't accumulate enough of a trajectory time to have a real statistical average of anything. If we're just comparing the number of GPUs on Summit to what we're expecting to have of Frontier, that already multiplies the processing power. And if each of those GPUs is also 3-4x more powerful than the ones on Summit, now you're talking about a substantial increase in how much simulation you can do. So you can work on bigger problems than you could before. Whether that's more replicas of a small system, or a bigger system that you couldn't physically do with previous hardware."

"We are trying to build an atomic scale model of a whole cell and we could not do it because the scales we were operating at were so enormous that you could not get a long enough time scale to get something scientifically meaningful," adds Vermaas. "A very small bacterial cell with a surrounding environment is about a billion atoms. That's at the very edge of what's doable right now. Being able to do that more routinely, who knows what that's going to enable, and going from a small bacterial cell to larger bacterial cell, you could investigate how antibiotics enter bacterial cells, for instance, or figure out how diffusion impacts photosynthetic efficiency. There are a lot of things about life science at the small scale that you could use molecular dynamics to attack, and Frontier with AMD Instinct GPUs will enable us to investigate them."

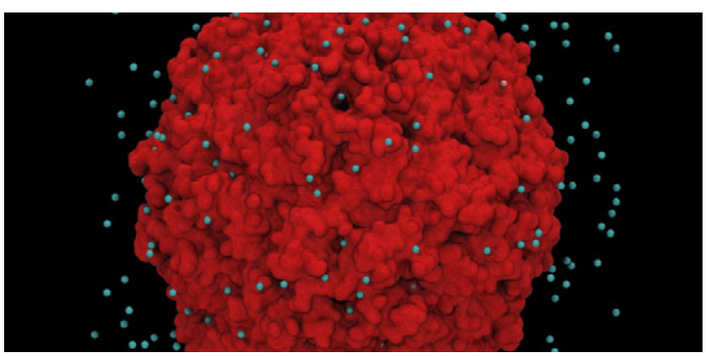

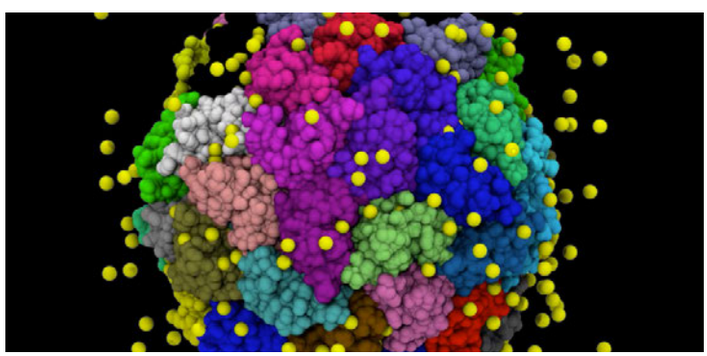

About ORNL

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.