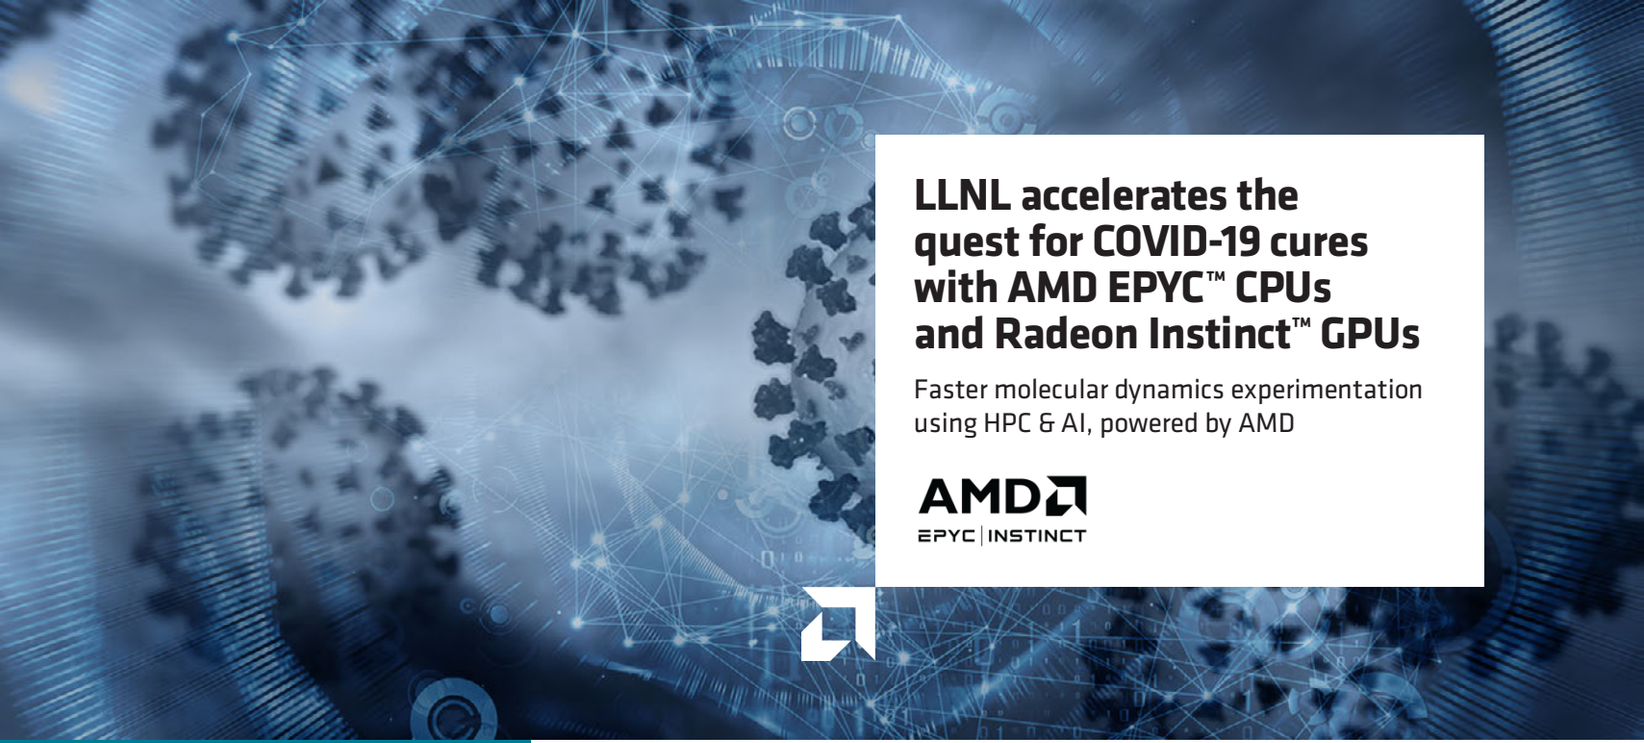

LLNL accelerates the quest for COVID-19 cures with AMD EPYC™ CPUs and Radeon Instinct™ GPUs

Faster molecular dynamics experimentation using HPC & AI, powered by AMD

AMD
EPYC | INSTINCT

CUSTOMER

COMPUTATION

Lawrence Livermore National Laboratory

INDUSTRY

Scientific research

CHALLENGES

Accelerating molecular dynamics simulations to provide faster experimental results

SOLUTION

Deploy AMD Radeon Instinct™ MI50 and MI60 GPUs

RESULTS

Accelerating the discovery of lead compounds for potential therapeutics, with the aim of reducing the time taken from three years to six months

AMD TECHNOLOGY AT A GLANCE

1st Gen AMD EPYC™ CPUs
2nd Gen AMD EPYC™ CPUs
AMD Radeon Instinct™ MI50 GPUs

TECHNOLOGY PARTNER

Penguin Computing

Lawrence Livermore National Laboratory (LLNL) has pioneered the use of computing in scientific discovery for nearly 70 years. It purchased one of the first UNIVAC computers in 1953, and now hosts the third most powerful supercomputer in the world, Sierra. But the research facility is always looking for new ways to increase the potency of the computing it can offer the scientific community across its significant portfolio of platforms. Recently, this has shifted the focus towards the acceleration that can be provided by GPUs, particularly Radeon Instinct™ GPUs.

The power of Radeon Instinct GPUs came into focus when LLNL turned its attention to molecular dynamics simulations that could potentially provide solutions to the Coronavirus pandemic. The calculations involve huge data sets and millions of simulations, for which GPU power is particularly well suited. Using 1,640 Radeon Instinct GPUs, the LLNL compute platform Corona (named after the light ring around a solar eclipse, rather than the virus), has unleashed unprecedented levels of performance for this research.

Discovering potential therapeutics with machine learning

The LLNL team has been employing machine learning (ML) to facilitate its work. "Every atom is a particle, and we move the particle over time," explains Felice Lightstone, Senior Scientist at LLNL.

"AMD GPUs have been exceptionally helpful. Performance with AMD Radeon Instinct using OpenMM has been great."

Edmond Lau, Staff Scientist at LLNL

"This is very useful to understand how the proteins function and how they interact with small molecules and other proteins. We've developed ML algorithms to learn the different aspects of drug design, so that we could get through not hundreds of millions, but billions of compounds when we screen them."

"The training does take a long time, but once it's trained, it's just a burst," says Lightstone. "What you really need to prepare is the small molecule library for the virtual screening. Not long ago, we thought a billion compounds was really the top of our limit. But with the pandemic, a lot of chemical companies have expanded their databases to be over a billion compounds, so we've actually prepared 1.6 billion compounds now for the small molecule screening."

"There are so many possibilities," says Lightstone. "That's why it has to go through this active learning loop." This is where ML is key to making this process effective. Same is true for the therapeutic antibody work. "We can't try every possible change, so it's targeted," adds Edmond Lau, Staff Scientist at LLNL. "The ML helps to figure out what are the interesting possibilities that we should try experimentally." The simulations that the LLNL team run are on the whole protein. "We have to chop it down, so they tend to be in chunks of 90,000 to 100,000 particles normally. This still requires about 40 nodes, which is 160 GPUs, to go through each amino acid mutation."

Unleashing the power of the GPU for molecular dynamics

Performing lots of similar calculations with incremental changes in this way is exactly where the forte of GPUs lies. "I've been using a program called OpenMM, which does calculations exclusively on the GPUs," says Lau. "Having the speed of the GPUs has helped a lot. You just can't do these calculations in a useful amount of time on CPUs. So just having the availability of MI50s and MI60s makes these free energy calculations possible. AMD has helped us out a lot by also placing one of their engineers to help convert OpenMM to our system. Performance with Radeon™ Instinct using OpenMM has been great."

The aim has been to design neutralizing antibodies that can defend cells from the Coronavirus pathogen. "This is the initial way that SARS-CoV-2 enters the human cells, so if we can actually block that with antibodies, we can stop the infection," says Lau. "We've been doing single-point mutations on residues. They take up a ton of GPUs, and AMD GPUs have been exceptionally helpful. Using the Instinct MI60s, we've been doing free energy simulations on Corona. Using the expansion with the new MI50 cards we recently received, I'll be able to do many more free energy calculations."

"Previous to the COVID-19 pandemic, we trimmed these simulations down to the 100,000-atom level so that they would be manageable," continues Lightstone. "One of the trimmings is to not include so many water molecules, because that just takes up a lot of compute time, and we don't gain a lot of information. But we're unable to trim the protein, because we are just doing these single amino acid changes, and we are looking for exact interactions with the viral spike RBD protein. Everything we've done is in anticipation that we have to run a lot of simulations to get the answers we need. And the answers only come through the entire ensemble of the calculations."

The Corona system that LLNL has been using for this work has 170 nodes with 1st Gen AMD EPYC™ processors, and 121 nodes with 2nd Gen AMD EPYC processors, offering 48 cores per node. Of these nodes, 82 have four AMD Radeon Instinct™ MI50 GPUs each, 82 have four MI60s, and 123 have eight MI50s, for a grand total of 1,640 GPUs, providing a peak 10,948 TFLOPS of performance and 52,480 GB of GPU memory.

The eight-GPU nodes take advantage of the huge allocation of PCI Express® lanes available on the AMD EPYC platform.

Conquering chemical space with GPU computing

The research being performed by LLNL could help with progress towards a cure for COVID-19 in several ways. "The first one is vaccines," says Lightstone. "But we have already made strides in the second way, of designing a therapeutic antibody that is more targeted to SARS-CoV-2. The third option is via small molecule drug therapies, which would be a pill or an IV-drug. For all these options, the time scale for development is normally very long. The average time to go from identifying your protein that you want to target to an investigative new drug that you can file with the FDA is six to seven years. What we've been trying to do is to crush down the discovery phase of that from three years to six months."

"It's a huge matrix of information that needs to be discovered. Computing is the only way we'll be able to conquer that in the future."

Felice Lightstone, Senior Scientist at LLNL

"Screening throws away the garbage," continues Lightstone. "Then there's testing where we iteratively improve any of the hits to try to make them into candidates. Traditionally, when you get to clinical safety trials, the failure rate is about 90% for a new therapy. We hope that we're making a huge impact by accelerating the discovery of lead compounds that could potentially become therapeutics for COVID-19. We're trying to accelerate therapeutic discovery using high-performance computing."

The abilities that LLNL has to offer are set to increase even further when its new supercomputer, El Capitan, comes online. This is scheduled to be installed at LLNL in early 2023, using AMD EPYC CPUs and Radeon Instinct GPUs. The Exascale level of computing El Capitan promises will provide an even greater level of performance for solving the molecular dynamics problems required in finding cures for diseases like COVID-19.

"It is a scale problem," says Lightstone. "Chemical space, the number of all possible molecules and compounds, is estimated to be 10 to the power of 60. When you're only looking at a billion compounds, and then consider all the possibilities, you realize it's still a very small proportion. And then you think about all human proteins and all the viral proteins and all the bacterial proteins that we have to cover in that space, it's a huge matrix of information that needs to be discovered. Computing is the only way we'll be able to conquer that in the future."

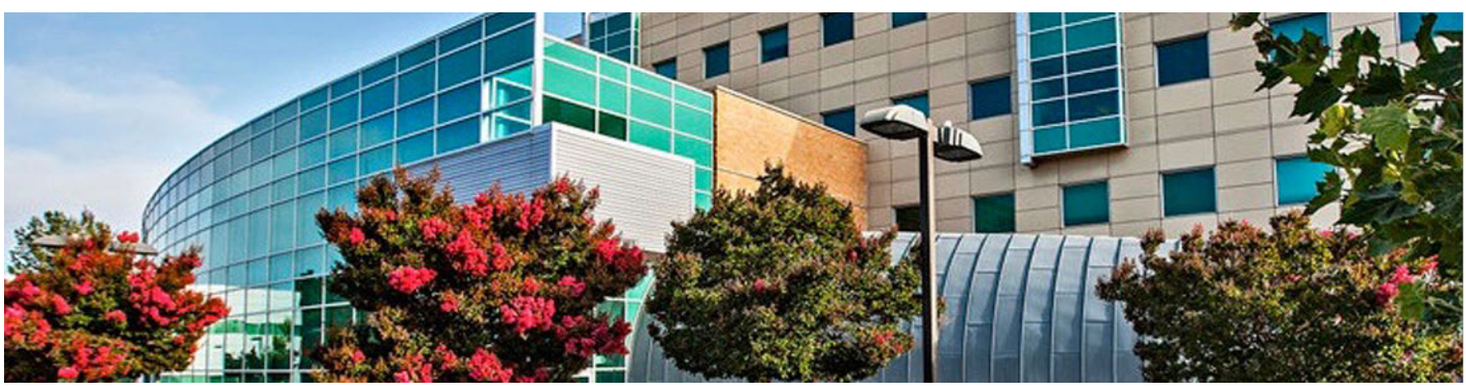

About LLNL

Lawrence Livermore National Laboratory is a US federal research facility based in Livermore, California. It is a premier research and development institution for science and technology applied to national security. It also applies its special expertise towards multidisciplinary capabilities including energy and environmental needs, scientific research and outreach, and economic competitiveness. Throughout its history, LLNL has been a leader in the use of computing for scientific discovery, starting with the purchase of one of the first UNIVAC computers. Its Sierra system is currently the third most powerful supercomputer in the world. For more information, visit llnl.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.