EXASCALE COMPUTING PROJECT

ECP Application Development Summaries

This is a compilation of summaries of the scientific, engineering, and machine learning application codes that were funded under the Exascale Computing Project. The research, development, and deployment activities of these code teams produced simulation software capable of running on the world's fastest supercomputers to address the nation's most critical challenges in scientific discovery, energy assurance, economic competitiveness, and national security.

July 2024

This research was supported by the Exascale Computing Project (17-SC-20-SC), a joint project of the U.S. Department of Energy's Office of Science and National Nuclear Security Administration, responsible for delivering a capable exascale ecosystem, including software, applications, and hardware technology, to support the nation's exascale computing imperative.

U.S. DEPARTMENT OF ENERGY | Office of Science

NNSA National Nuclear Security Administration

CONTENTS

The CANDLE Project: Using Exascale Computing to Improve Cancer Research and Treatment Outcomes

Cancer is the second most common cause of death in the United States. According to the NIH, roughly two million Americans were diagnosed with cancer in 2023, and more than 600,000 of those people died from the illness. In addition to cancer's staggering incidence and mortality rate, it imposes severe economic burdens on affected families, totaling more than $170 billion in direct medical costs alone in 2020. High performance computers and advanced software solutions such as artificial intelligence and machine learning are among the most critical new tools being employed in the fight against cancer. These new technologies can greatly accelerate the pace of research in drug design, treatment optimization and analysis, and fundamental cancer biology, all with the goal of improving patient prognoses and quality of life.

The Cancer Distributed Learning Environment (CANDLE) application is a software platform sponsored by the Exascale Computing Project (ECP) and built to combine and apply high performance computing and machine learning to key areas of cancer research and care. The application uses a deep neural network, a type of machine learning algorithm capable of recognizing patterns and classifying information with minimal human supervision, to rapidly complete large and complex data analysis and modeling tasks when accelerated with the power of modern supercomputers.

Using CANDLE, researchers can quickly create and train large numbers of computational models for three key tasks: predicting drug interactions with various cancer types to optimize treatment and drive precision medicine tailored to individual patients; understanding the molecular dynamics in the RAS/RAF pathway, a ubiquitous network of cellular interaction involved in 30% of cancers; and automating the analysis of millions of patient treatment records to find patterns which are not recognizable to human analysts across factors such as patient lifestyles, cancer types, and environmental exposures.

Machine learning cannot be applied to cancer research and clinical intervention without tools to quickly train effective computational models for highly specialized research and clinical tasks. Creating and training the millions of unique machine learning models required to adequately address sweeping research questions--such as interactions between thousands of different drugs and cancer types--requires an enormous amount of computational power, as a single model may take several hours to train given traditional methods. This computational demand requires improvements in hardware capabilities and innovation in machine learning training approaches.

The CANDLE application team has leveraged the power of exascale computation to dramatically improve researchers' ability to create machine learning models for cancer research. Using the CANDLE tool, researchers can create models at unprecedented speeds—almost 500 times more quickly than on previous petascale machines—and with greatly improved versatility. The application boasts features which allow researchers to view models' confidence in the validity their predictions and to screen predictions for minimum acceptable accuracy. Furthermore, CANDLE's versatility in training machine learning models allows the application to be implemented in research on other illnesses, such as COVID-19.

Exascale-accelerated machine learning is positioned to transform cancer research and care in the US. With support from the CANDLE application, researchers will accelerate development of new cancer treatments with improved clinical outcomes, optimize healthcare providers' decision-making by providing largescale analysis across nation-wide databases, and support basic research by elucidating the fundamental networks of cellular interaction which cause cancers.

The Combustion-Pele Project: Using Exascale Computing to Improve Engine Efficiency and Enable Sustainable Fossil Fuel Alternatives

Diesel and gas-turbine engines drive much of the world's industry and transportation. However, burning fossil fuels to power these engines strengthens the climate-warming greenhouse effect and emits harmful pollutants—which, according to the World Health Organization, kill more than 13 people every minute due to lung cancer, heart disease, and stroke. Scientists have spent the past half-century in search of safe and economically viable fuels, but a poor understanding of the complex physics inside combustion chambers has hindered their development efforts. High-performance computing is a key tool in overcoming this obstacle; researchers can begin to understand the best conditions for burning new and existing types of fuel by simulating mixing and combustion processes such as fuel-air mixture formation and ignition processes in extreme detail. Information from these simulations will be used to develop and test new fuel types and combustion conditions, ultimately leading to low-emission and high-efficiency combustion engines and a more sustainable standard in global transportation and power generation.

The Exascale Computing Project's Combustion-Pele application is designed to apply the power of the world's fastest supercomputers to recreate complex combustion environments in unprecedented detail. The Combustion-Pele team focused on developing new software architecture for exascale modeling of key scientific unknowns such as low-temperature combustion dependent on mixture formation, and the effects of multistage ignition. Understanding these processes will lay the groundwork to build new and much more efficient compression ignition combustion engines and will support the development and implementation of viable alternatives to fossil fuels.

The physics of ordinary combustion engines are already well understood, but new designs require far more precision to ensure that they function safely and consistently. These designs are difficult to implement without understanding the unique fluid and combustion dynamics that result from altering factors such as fuel types which affect, ignition temperatures and fuel injection timing. Researchers cannot efficiently address these unknowns without extremely granular models of the dynamic processes within combustion chambers, and creating these models requires previously unattainable computational power combined with software developments to greatly improve computational efficiency.

The Combustion-Pele application addresses these needs, delivering a system capable of the most precise physical models of combustion processes to date. The application boasts several innovative software improvements—including adaptively refined simulation mesh for improved computational efficiency and increased fidelity in critical regions; more realistic physics models for multi-component sprays, soot formation, gas dynamics, and thermal radiation effects; and a new portable code that functions on a wide range of computing hardware and scales up to more than three-quarters of Frontier, the world's largest exascale computer. The Combustion-Pele team has used these improvements to simulate the interaction of two unique fuels with varying reactivity under a multipulse injection strategy. This work provides a functional baseline for future simulations to test the most promising biologically derived fuels and hydrogen fuel blends for use in key sectors such as aviation and power generation.

Future simulations using Combustion-Pele will capture the complex physical and chemical processes in realistic engine environments, augmenting the limited measurements that can be made in laboratory-scale engine configurations. Researchers can use this high-fidelity data to train artificial intelligence systems for more complex and efficient simulations in the future and to develop predictive models for engineering simulations. The impending rapid improvement in combustion simulation capabilities will allow scientists and engineers to develop much more efficient fossil fuel combustion engines that can utilize sustainable fossil fuel alternatives, thereby greatly reducing the negative impact of transportation and power generation on our environment.

The Energy Exascale Earth System Model: Largescale Climate Simulations for Water Cycle and Weather Prediction

Climate change is influencing the global water cycle, resulting in droughts, changes to cloud formation and rainfall, as well as more frequent and powerful storms. These changes may negatively impact regional energy and water supplies, agriculture, and human and environmental health if not adequately addressed. Predicting changes in the water cycle gives critical information to regional policy makers in multiple sectors of the US and global economies, and advanced climate models are a key resource for providing this information. Researchers must consider complexities such as the microscale chemistry and physics of cloud formation and the impacts of anthropogenic climate change on cloud formation, and must incorporate these factors into simulations which routinely span decades while maintaining computational and economic feasibility. High performance computers are an essential part of this process, as they allow for greatly enhanced fidelity across longer timescales, leading to much more accurate and cost-effective predictions.

The Exascale Computing Project's Energy Exascale Earth System Model (E3SM) has developed a computational model capable of running on the world's fastest supercomputers, simulating climate-induced changes to cloud formation and weather events with unprecedented speed and accuracy. The E3SM application runs multidecade climate simulations at resolutions that accurately capture cloud formation through multiphysics integration and eddy capturing ocean models.

Accurately simulating the key processes in cloud formation requires minimum resolutions of 1 km in the atmosphere. Previous petascale computing systems are capable of such resolution, but only at great expense and for very short durations—measured in days rather than years. Running conventional climate models at the resolutions and durations needed for accurate and applicable cloud resolution simulations requires a 5,000× increase in computing resources. While modern exascale computers boast a more than 1000x improvement in processing speed compared with previous state-of-the-art machines, innovative software approaches were still needed to meet this benchmark.

The E3SM application team met and exceeded this computational challenge, running the first ever earth system model with a fully weather-resolving atmosphere and embedded cloud resolving model. The application uses superparameterization to resolve the convective processes responsible for cloud formation within 1 km grid boxes, with the rest of the model processes being resolved in 25 km boxes. These approaches, combined with efficient use of GPUs on exascale computers, improved E3SM's simulation rate from initial speeds of .011 years per day to 5 simulated years per day.

The greatly improved simulation speed and fidelity generated by E3SM on exascale machines will give the scientific community the much-needed ability to predict, assess, and respond to the challenges imposed by local variations in the water cycle caused by global climate change in the years to come. These simulations will provide key information to policy makers, enabling the mitigation of potentially disastrous changes in weather patterns and supporting long-term stability in energy generation, agricultural production, and water quality and supply.

The EQSIM Project: Using Exascale Computing to Improve Risk and Hazard Assessment for Earthquake Disaster Preparedness

The costs associated with earthquake damage are staggering, ranging from direct property damage to emergency response, loss of wages, rebuilding costs, interruption of transportation and daily life, not to mention injury and the loss of life. According to the U.S. Geological Survey and the Federal Emergency Management Agency (April 2023), earthquakes cost the nation an estimated $14.7 billion annually in building damage and associated losses alone. Computational models of ground motion during an earthquake help engineers and seismologists identify the areas most at risk of experiencing a severe quake. This information can be used to create disaster response plans, improve building specifications, and inform decisions about where construction should take place.

The Exascale Computing Project's Earthquake Simulation (EQSIM) application utilizes high performance computing to improve our understanding of the physics of earthquakes. The application enables an accurate end-to-end earthquake simulation, from the beginning of a fault rupture to the eventual impact of surface ground motions on buildings across a wide range of sizes and material compositions. These simulations can be coupled to engineering codes to identify and mitigate the vulnerabilities of key structures such as apartment buildings, medical facilities, and power plants. To evaluate and improve its simulations, the EQSIM team created a regional-scale model of the San Francisco Bay Area. This area contains the Hayward fault line, which is considered the most dangerous fault in the country. The Hayward fault line has not ruptured since 1868, and has ruptured roughly every 150 years on average.

Previous earthquake simulations relied on rough estimates of ground motions compiled from many earthquakes in different locations. This method results in uniform ground motion simulations regardless of location, making site-specific predictions extremely difficult. Conventional simulation techniques also lack the computational capacity to model ground motions at a sufficient frequency and fidelity to accurately determine the effects of earthquakes on infrastructure in more than one dimension, significantly limiting their applicability to real-world structures and engineering specifications.

The EQSIM team has leveraged exascale computing to address these shortcomings. The EQSIM software suite contains invaluable tools for modeling factors such as soil composition and surface topography. These new tools allow researchers to model ground motions using location-specific information, greatly improving the simulation's ability to accurately predict the resulting effects on structures. Furthermore, the software suite can simulate ground motions at sufficient frequency to study earthquake effects on large structures of varying material compositions in three dimensions, improving researchers' ability to model how specific structures will be affected.

These achievements greatly improve our understanding of how earthquakes affect key infrastructure. Using EQSIM, researchers, engineers, and planners can now understand the location-specific impacts of variable-strength earthquakes, which will inform city planning, improve disaster preparedness, and expand on foundational research in seismology and beyond.

Using Exascale Computing to Improve Molecular Dynamics Simulations and Develop New Materials and Technologies

Molecular dynamics is a core research technique in the 21st century. Scientists use the technique to create physically accurate simulations of the motion and interaction of atoms and molecules. Molecular dynamics simulations offer an extremely precise view of how nature works at the smallest scales, and their results help us understand how to solve key problems in fields ranging from nuclear fission and fusion, to biophysics, to materials science. These simulations require a huge amount of computational resources, and even the world's fastest computers are currently limited to relatively tiny physical spaces and timescales. Exascale computing represents a substantial improvement in these capabilities, offering researchers the tools they need to create new materials and develop next generation technologies.

The Exascale Computing Project's Exascale Atomistic Capability for Accuracy, Length, and Time (EXAALT) application was built to utilize exascale computing to solve the world's hardest problems in molecular dynamics. EXAALT integrates and enhances several widely used codes to improve three key aspects of molecular dynamics simulations: the physical accuracy of a simulation, the size of the simulation area, and the simulation's duration. The application boasts new physics models for describing materials' properties from first principals, and offers unprecedented physical accuracy for simulations of materials in highly dynamic situations such as nuclear fusion and fission conditions.

Improved molecular dynamics simulations will enable several key scientific advances which were previously out of reach. For example, molecular dynamics simulations currently lack the size and accuracy to model the formation of nanoscale "fuzz" in the tungsten walls of fusion reactors during use—a defect which destabilizes fusion plasma and reduces reactor efficiency. Successfully modeling this process requires simulation of hundreds of thousands of atoms with physical accuracy down to the quantum scale. Challenges such as this one will remain intractable without improved computational resources, hampering scientists' understanding of key issues and stalling the pace of technological innovation.

The EXAALT team has successfully developed an integrated code base which allows users to solve problems of unprecedented complexity, completing molecular dynamics simulations more quickly, accurately, and at larger scales than ever before. The application boasts performance enhancements of nearly 400x compared with 2016 benchmarks, and offers a range of physics models as well as advanced task management options which can greatly improve performance across multiple simulations. These improvements allowed the EXAALT team to effectively model nanoscale fuzz formation as well as other key problems, such as predicting the of buildup of radiation and gas byproducts during nuclear fission for more efficient fuel use and reduced nuclear waste.

EXAALT's improvements to the size, accuracy, and duration of molecular dynamics simulations will allow scientists to explore much larger and more diverse problem spaces than ever before. Increased computing power and innovative software approaches will give researchers the tools they need to answer questions about how biomolecules form and mis-form in our bodies and in the outside world, how to create resilient materials for our most advanced technologies in fields ranging from energy production to aerospace engineering, and beyond.

Exascale 3D printing simulations enable accurate predictions for manufacturing metal parts

Additive manufacturing, also known as 3D printing, plays a critical role in transforming the manufacturing sector by providing the ability to make parts with very complex geometries. While 3D printers for home use can easily make plastic components, industrial use of 3D printing often seeks to make metal components. For many applications, such as aerospace and automotive, the strength of these metal components is very important. The strength of a metal component is determined by the component's microstructure, or the pattern that the metal crystals make as they solidify. Predicting the strength of a metal component – critical for industry that relies on them – is necessary and requires great computing power.

The biggest influence on metal microstructure is the thermal history that the component experiences as it is printed layer by layer, with a laser that melts the appropriate pattern of metal powder in each layer. Every location in one of these parts can have a different thermal history, so every location can have a different microstructure and different strength properties. Different component geometries can have very different thermal histories, even if the 3D printer settings were exactly the same. Thus, predicting the ultimate performance of a 3D printed metal part is extremely challenging.

ExaAM is an Exascale Computing Project effort aimed at producing a suite of simulation codes capable of predicting the strength properties of a 3D printed metal part. This is done by predicting the thermal history of every location in the part, and then predicting the effect that these thermal histories have on the microstructure and the resulting strength properties. Previous modeling efforts had successfully demonstrated similar abilities on a very small scale, such as a 1 millimeter section of part, but to enable this capability on a large scale of an entire part requires very large computational resources, such as the Frontier supercomputer at ORNL.

The ExaAM project spent significant time and effort making the various codes in the suite capable of running efficiently on GPUs, since the vast majority of the capabilities of modern supercomputers comes from GPUs. The resulting suite, which is freely available as open source software, successfully ran on more than 8,000 nodes on the Frontier supercomputer (more than 32,000 GPUs!) and successfully predicted the strength properties of a National Institute of Standards and Technology additive manufacturing benchmark component.

Though the ExaAM projected ended in December 2023, the key components of the suite have proven so valuable that other Department of Energy projects – such as the Digital Factory at ORNL's Manufacturing Demonstration Facility and the Advanced Materials and Manufacturing Technology effort within DOE's Nuclear Energy program – continue to fund the development of these simulation tools. Thus the ExaAM supercomputing project promises to move additive manufacturing forward, as it continues to pay benefits for the nation and world.

The ExaBiome Project: Applying Exascale Algorithms and Software to Improve the Quality of Life

Accelerating foundational biological research and biotechnology development are critical steps in addressing the ecological problems of the 21st century, from the spread of antibiotic-resistant diseases to the accelerating formation of deserts and ocean dead zones due to pollution and climate change. The CDC estimates that antibiotic-resistant diseases infect more than 2.8 million U.S. citizens and cost more than $55 billion in treatment and lost productivity annually, and the collapse of ocean and terrestrial ecosystems has been identified as a catalyst to starvation and water shortage, the spread of disease, and mass migration from affected areas.

Historically, identification and reverse engineering of microbes—the most widespread and diverse form of life on Earth—has yielded powerful technologies such as uniquely effective antibiotics, revolutionary tools for genome editing such as CRISPR-Cas, and advanced biomanufacturing techniques for vaccines and agricultural products. Applying high performance computing to this field will greatly improve the speed with which microorganisms can be identified and reverse engineered for new biotechnologies, providing new solutions in medicine, agriculture, and environmental science.

The ExaBiome team within the Department of Energy's Exascale Computing Project has created a software platform which enables unprecedented insight into a staple biological research technique called metagenomics—the reconstruction and comparison of genomic information from entire communities of microbes found in soil, water, or tissue samples. With the application of exascale computation, researchers can more quickly identify new species of microbes and viruses, map population changes in a community or environment over time, and understand the function of unique cellular machinery by comparing huge numbers of similar genes and proteins.

Without exascale computing researchers cannot fully analyze complex microbial datasets—which are often dozens of terabytes in size—in a reasonable timeframe, and must resort to analyzing and computationally aggregating subsamples of large datasets. This approach limits researchers' ability to discover and characterize uncommon species, reduces the accuracy of assembled genomes, and obscures the function of genes and proteins.

To address these issues, the ExaBiome team has created a scalable metagenome assembler, protein analysis, and similarity search tool. These tools allow researchers to take raw data and form genome sequences for individual species, cluster proteins and accurately identify their function, and analyze multiple metagenomes—which can show how an environment has changed with time. The ExaBiome team has used these tools to analyze over 400 million protein sequences in less than 4 hours, a process that would take weeks using previous methods. The team has also used metagenomic techniques to assemble the largest environmental dataset to date, allowing for the discovery of new species and insight into the composition of microbial communities that cannot be replicated with subsampling.

Exascale delivers unprecedented performance and fidelity to move research forward at a faster pace. These new capabilities will dramatically accelerate the development of new biotechnologies by unlocking the potential of as-yet poorly understood biological systems. As an example, these technologies can be implemented in the development of next-generation fertilizers for increased agricultural yield and improved soil health, bioremediation tools to stabilize critical ecosystems and keep our environment livable, new methods to analyze and treat infection and diseases of the human microbiome, and beyond.

The ExaFEL Application: Using Exascale Computing to Analyze Light Source Experiments in Real Time

Light source experiments using x-ray free-electron lasers (XFEL) allow scientists to observe the structure and dynamics of individual atoms and molecules in high resolution. Researchers use the results from these experiments to predict how various biomolecules and engineered materials interact and, ultimately, to create new technologies in the biological, materials, and energy sciences. However, light source experiments can be difficult to conduct due to high construction and operating costs for facilities, limited operating time on XFEL machines, and difficulty processing the large amount of data that these experiments generate. High-performance computing can significantly improve research efficiency at light source facilities by enabling rapid data analysis and enhanced resolution in diffraction pattern models. These improvements will help researchers more thoroughly understand fundamental chemical processes such as bonding and catalysis and help them apply this information for technological developments throughout the sciences.

The Exascale Computing Project's ExaFEL application is built to enable near real–time analysis of data from XFEL experiments to maximize the scientific output at billion-dollar facilities such as the SLAC National Accelerator Laboratory's Linac Coherent Light Source (LCLS). ExaFEL supports efficient methods to rapidly identify molecular structures and reconstruct them in 3D so researchers can visualize how a structure is changing over time. These capabilities reveal structures and molecular interactions with atomic-level detail and enable real-time feedback and control over experiments.

LCLS users require an integrated combination of data processing and scientific interpretation, both of which demand intensive computational analysis. Traditional analysis methods typically take weeks to process data from light source experiments, thereby limiting researchers' ability to iterate their experiments and extract the most value from scarce beam time. Furthermore, data throughput at light source facilities is increasing rapidly as these facilities become more sophisticated, exacerbating the need for improved analysis capabilities. For example, new additions to the SLAC LCLS facility will increase its data throughput by three orders of magnitude by 2025, rendering current analysis methods untenable.

The ExaFEL team has addressed this urgent computational need by creating an exascale-based data analysis workflow that reduces analysis times from weeks to minutes. The team designed new GPU-accelerated reconstruction algorithms that have improved calculation speeds by 1,000× since 2016 while also improving image fidelity. Furthermore, the ExaFEL application is designed to scale with the increasing computational demands of next-generation light source facilities, ensuring rapid data analysis and improved outcomes in future light source experiments.

ExaFEL enables researchers to modulate their experimental parameters during runs on advanced light source machines. This new capability will improve research efficiency at light source facilities and will accelerate scientific progress as a result. Using ExaFEL, researchers at facilities such as the upgraded LCLS will gain unprecedented insight into key unknowns such as the real-time functionality of biomolecules, fundamental interactions in quantum and nanoscale material dynamics, catalysis and photocatalysis for new chemical transformation and solar energy conversion processes, and beyond.

The ExaSGD Project: Enabling a Resilient and Efficient Power Grid through Exascale Computational Modeling

The U.S. electric grid is a massive and intricately connected network capable of generating more than 1 million megawatts across over 600,000 miles of transmission lines. As renewable sources of energy such as wind and solar power become more significant elements of the U.S. power supply, grid operators will have to develop strategies to balance fluctuations of the resource supply against energy needs. Pressures from increasingly frequent and powerful weather events, as well as cyber-attacks and random equipment failures, requires detailed models of power grids in which multiple failure types and locations can be quickly and accurately responded to. Without these capabilities, power grids and the citizens who rely on them are vulnerable to uncontrolled blackouts at potentially critical locations and times.

The Exascale Computing Project's Stochastic Grid Dynamics (ExaSGD) application gives controllers the tools to successfully manage power grids under changing conditions, minimizing outages and throttling, while augmenting the grid's ability to utilize energy from renewable sources and strengthening its resilience to external threats. This is achieved by creating and analyzing accurate computational models of the energy grid under numerous potential scenarios.

Previous efforts to model and optimize the power grid relied on simplified models which supplied rough optimization attempts based on known past events. These models are often inefficient and are increasingly liable to fail under pressure from changing weather patterns and a diversifying energy grid.

ExaSGD can simulate power grids with tens of thousands of components under tens of thousands of operational scenarios. The large-scale, nonlinear power grid optimization runs on exascale supercomputers like Frontier allows for "whole grid" simulations in minutes or hours instead of days or weeks. The new methods and algorithms in ExaSGD can give a similar boost in awareness of much smaller scenarios that are important to grid operators, which can be run on computers as accessible as laptops.

This new tool will provide critical modeling and optimization capabilities to improve the strength and efficiency of the U.S. power grid and, with the power of exascale computers, deliver simulations in near-real time to inform grid operators' response to changes in power availability. This improved flexibility will help keep homes and critical infrastructure like hospitals powered through crisis conditions such as natural disasters, extreme temperatures, or cyber-attacks, which may overwhelm the present-day grid.

Exascale computing sheds light on the dark structure of the universe, simulating its large-scale structure and the processes that drive its evolution

The universe is filled with dark unknowns. Dark energy, dark matter, the mysterious neutrinos, and how objects like massive black holes form. Scientists think about 85 percent of the mass in the universe is comprised of dark matter that has conventional gravitational interactions, but they do not know what the dark matter is made of. Dark energy, the mysterious agent causing the current acceleration of the rate of expansion of the universe, is even less understood.

Since scientists are unable to experiment with the universe, they turn to computers to run experiments. In this case, a suite of software codes solve the underlying complex equations that best model the current knowledge and generate catalogues of objects such as galaxies and galaxy clusters that live within dark matter clumps. The ExaSky project's codes are designed to run on the world's fastest supercomputers, like Frontier and Aurora, exascale computers that can do a billion billion calculations per second. The codes create simulations of the universe in several different scenarios, requiring calculations of massive amounts of data, aiming to produce the most extensive synthetic sky maps ever created.

ExaSky drew on two major sets of computer codes that simulate a large range of physical processes, such as how billions of galaxies formed and arranged themselves in what is known as the cosmic web, from star birth to supernova death throes and everything in between. The challenge was to account for sizes up to several gigaparsecs in scale – length scales that are probed by current and future cosmic surveys, with galaxy formation-related physics modeled down to kiloparsecs—a dynamic range of one part in a million. The codes predict details of the structure and properties of individual galaxies as well as how they interact with other galaxies and with dark matter, via gravity. The ExaSky team updated these codes to best use exascale computers' capabilities, by rewriting the solvers and by adding new physical processes, a process that took seven years and many contributors to make it happen. One set of codes, called Nyx, was initially written for CPUs and the team had to rewrite much of it to run on GPUs, the computational technology that exascale computers use. The other main code, HACC, was already written for GPUs, but was re-optimized and made much more powerful by adding a number of astrophysical modeling capabilities. Both ExaSky codes, HACC and Nyx, although using different methods, have been shown to produce simulation results that agree at the 1% level, a major improvement over the previous state-of-the-art.

The ExaSky simulation program is now shedding light on some of the biggest questions in cosmology today: the mystery of dark matter and dark energy; the nature of primordial fluctuations, or how the universe evolved from a relatively smooth almost wrinkle-free vacuum to the wiggles and bumps that gave way to galaxies; and determining the mass of the mysterious neutrinos.

Cosmological simulations can represent a huge range of scales, from the size of the smallest galaxies to a distance of less than a fifth of the way to the edge of the observable universe – a huge scale from the smallest to the largest scale of observation. They are also allowing scientists to model what happens inside the central core of energetic galaxies, where radiation is detected but is not coming from stars. If it Is the work of supermassive black holes, what are the processes forming them? Supercomputer models such as the ExaSky simulations can help scientists figure that out.

Analysis of all this data can be used to interpret sky survey observations by both terrestrial and space-borne telescopes. The needed simulations of the past, present, and perhaps future of the universe are not possible without the power of exascale computing. The ExaSky project codes can produce levels of performance that is an order of magnitude greater than previous efforts. The scale of improvement for HACC is a factor of 270 times over from when the project started. Coupled with next-generation sky surveys, the ExaSky simulations have already and will continue to improve scientists' understanding of the large-scale physical processes that drive the evolution of structure in the universe.

Exascale simulations pave the way for new generation of nuclear power reactor design

Nuclear power reactors are an important part of a strategy to reduce greenhouse gases and achieve a carbon-neutral society, but they require extensive time and money to design and build. High-performance computing, however, can help accelerate the design of future reactors.

The next generation of supercomputers can help predict behavior inside the next generation of reactors, helping engineers and reactor designers see how they would behave, how much fuel they would need, how long the cycles last, and more. Such simulation can provide valuable information in months that otherwise would take many years of experimentation.

The Exascale Computing Project's Small Modular Reactor (ExaSMR) project has developed the first of its kind simulations of advanced nuclear reactors, including small modular reactors and others, such as molten salt reactors and microreactors. This advanced software suite developed for exascale computing over the past seven years may help create the next generation of nuclear power generation in the United States.

With high confidence predictions and high resolution simulations, reactor companies don't have to run as many experiments, which are time consuming and costly. The Exascale codes provide the means to see what the experiments would show, but more quickly and more accurately, than what has been available previously. This can reduce the lead time from concept to construction of a prototype, which also saves money. Previous simulation campaigns have been limited by the large computing power necessary to perform detailed reactor simulations. The ExaSMR code team had to rewrite new algorithms that enable it to run efficiently on the Frontier supercomputer. The result is that these codes can help reduce dependence on physical experiments because they are more accurate and predictive.

The project performed simulations involving both neutron transport and fluid dynamics. Neutron transport describes how neutrons move around a reactor and cause fission events, producing the heat that drives a power reactor. Fluid dynamics describes how the coolant moves through the reactor, transferring heat from the solid fuel to the liquid coolant. The ExaSMR team performed the largest reactor simulation ever, involving the simulation of more than 1 trillion neutron histories within the neutron transport solver and solving for more than 1 billion spatial regions in the fluid solver. The demonstration calculation achieved 100 times the computational performance that was demonstrated for the neutron transport solver just seven years earlier, and more than 400 times the performance for the fluid dynamics calculation. For this simulation, the code developers were named finalists for the Association for Computing Machinery's Gordon Bell Prize for outstanding achievement in high-performance computing.

The tools enabled by the supercomputer are for demonstration purposes, but private companies are looking at the ExaSMR results for possible use in design of new reactors. The code could be useful for a variety of reactor types, including molten salt and pebble bed reactors that the U.S. Department of Energy and Nuclear Regulatory Commission have under consideration. Exascale computing could allow them to optimize designs to accelerate the advanced reactor concept-to-design-to-build cycle that has constrained the nuclear energy industry for decades.

Exascale simulations help explain how a star explodes and creates heavy elements and particles throughout the universe

Exploding stars, known as supernova, that ripped through the universe before our sun was a star are regarded as a source of most of the heavy elements in the known universe by ejecting them into space with the violent force of nuclear blasts. Yet the process of how massive stars collapse after burning all their fuel until they explode is still largely unknown, and there remain many open questions of fundamental physics. These include questions about how ultra-dense matter, strong gravity, and energetic neutrinos behave under extreme conditions. But the next generation of supercomputing, exascale computing, is providing a new way to probe physics at these extremes.

An exploding star brings together a lot of complicated physics into one problem: Fluid mechanics, particle physics, and the mystery of how neutrinos drive the explosion and accompanying shock wave all have to come together to solve sets of equations at the same time to make a simulation that captures all the violence and turbulence of an exploding star. That challenge, which has taken seven years on the world's fastest supercomputer, is showing how stardust becomes heavy elements like iron and titanium and particles that populate Earth, its inhabitants, and the known universe.

ExaStar, as the simulation program is called, required major overhauls in developing codes of new physics modules that solve problems related to Einstein's theory of general relativity (rest assured, it remains intact), and many of the known equations of classic physics. Among the largest challenges was integrating all the code into GPU architecture, as that is what the Frontier supercomputer uses at the exascale. Developing new physics codes and having them run on GPUs was a major step forward: The team successfully implemented GPU code for almost all multiphysics modules, including neutrino transport, general relativistic gravity, nuclear kinetics, and hydrodynamics.

The simulation encloses the helium shell of the star, which has a radius around 10,000 km, with a maximum spatial resolution finer than 1 km in the innermost regions. At least 20 energy groups were used to resolve the spectra of neutrinos of all "flavors" (i.e., electron, mu, tau, and their antiparticles) from 0 to 300 MeV. It produced a detailed set of neutrino-matter interaction rates that are integrated with a high-density equation of state that provides pressures, entropies, and other thermodynamic variables. Through a better understanding of the locations in which the heaviest elements are made, ExaStar can help focus experimental efforts on the Facility for Rare Isotope Beams and other science facilities that are now a priority of the next decade of nuclear physics.

ExaStar now has multiple 3D simulations running of an exploding star. The star's core collapses into a little nugget, but the rest of the star is gigantic, so the algorithms are quite complex and the equations complicated. The goal is to have an entire suite of simulations that physicists can use to understand how such a star explodes, and perhaps answer longstanding questions of how the heaviest elements in the universe were born and thrust into space. The complication is that there are many different types of stars, and the code hopes to simulate how they turn into different types of supernova and produce the different elements that swarm the universe. It turns out it takes a lot of work to build codes with that kind of capability. The software must simulate capture behavior of matter at very high density and calculate the radiation transport of neutrinos. Predicting signals of such major events requires massive complexity in the software, which the team has completed. Thus exascale computing through the ExaStar suite is poised to answer some of the most fundamental questions of the universe, and perhaps the most basic of all – where did we come from?

The ExaWind Project: Optimizing Wind Energy Production with New Turbine and Wind Farm Designs

The efficiency and reliability of renewable energy will continue to grow in importance as the U.S. energy grid transitions from reliance on fossil fuels. Wind power is the largest source of renewable energy in the U.S., and one of the fastest growing energy sources in the country: According to the U.S. Office of Efficiency and Renewable Energy, wind energy made up more than 10% of U.S. power in 2022, and supplied more than 40% of new U.S. energy capacity in 2020, and more than 20% of new capacity in 2022. High performance computing is poised to greatly improve the efficiency of wind power generation by using advanced multiphysics simulations to model and optimize key operational parameters such as airflow and turbine structural dynamics.

The Exascale Computing Project's ExaWind application provides the software architecture needed to apply the power of next-generation supercomputers to the burgeoning wind energy sector. The project simulates the complex physics of an entire windfarm under various geographic and weather conditions. These simulations help determine the impacts of turbine placement, wind farm location, and blade design, driving improvements in efficiency and reducing the cost per megawatt hour of energy produced.

High performance computers have been used to improve wind turbine efficiency for more than a decade, but simulation fidelity and duration has been heavily limited by processing speed, resulting in simulations which could only yield general trends in efficiency with low degrees of confidence. To improve the accuracy of turbine and wind farm simulations, researchers needed to capture the physical airflow processes at the micron-scale boundary layer around turbine blades and to couple these processes to larger interactions occurring at the scales of turbines and entire wind farms. This coupling spans roughly eight orders of magnitude, ballooning the computational cost and forcing less powerful computational systems to sacrifice accuracy for simulation speed.

The ExaWind team has addressed these shortcomings for exascale machines, successfully integrating a whole-wind-farm simulation at scales ranging from microns to kilometers, and has reduced the duration of these simulations from days to hours. The integrated simulations span up to 38 billion grid points and simultaneously model multiple turbines at high fidelity, allowing for high-confidence numerical predictions of efficiency and energy loads, as well as accurate models of structural stability and deformation under operational and storm conditions.

Using ExaWind's large scale, high fidelity multiphysics simulations, researchers can model new turbine and wind farm design approaches without the slow and cost-intensive processes of real-world development and testing. This new tool will greatly improve the speed and cost-efficiency of wind energy innovation while supporting "longshot" designs which were previously not financially viable, and which may rapidly accelerate progress in the field.

The GAMESS Application: Improving Computational Chemistry Simulations for New Materials and Biotechnologies

Computational chemistry is a key research technique that enables scientists to use precise simulations of atomic and molecular interactions to understand the function of novel materials and biomolecules, design new drugs, and create technologies in fields ranging from environmental remediation to medicine. However, these simulations require significant computational resources owing to the complex chemical and quantum mechanical processes involved in chemical bonding, and even modern supercomputers can struggle with models of large atomic systems such as heterogenous catalysts and complex biomolecules.

The Exascale Computing Project's GAMESS application was created to address this shortcoming. The application significantly expands capabilities in computational chemistry, allowing for simulations of far larger and more complex systems in shorter timescales. Researchers can use GAMESS to design and test aptamers—DNA and RNA fragments—for detecting novel pathogens, explore how various materials can be used to remove CO2 from the atmosphere and transform it into useful compounds, and design more effective catalysts for chemical manufacturing processes.

The extreme complexity of atomic and molecular interactions has limited the scope of legacy computational chemistry simulations. Traditional systems can typically model interactions between hundreds of atoms before being forced to sacrifice accuracy to reduce simulation durations to feasible levels. However, simulations of this size prohibit the direct study of critical systems, including biomolecular interactions and most catalytic reactions. To compensate for this deficit, researchers have resorted to using smaller models that approximate complete systems. Unfortunately, these approximations are often inaccurate and result in incomplete or erroneous results when compared with information gathered from real-world experiments.

The GAMESS team has addressed this shortcoming by delivering an application that leverages exascale computing to extend the size of physically accurate computational chemistry simulations by orders of magnitude. Using GAMESS, researchers can now model systems of up to tens of thousands of atoms without sacrificing accuracy, which allows for the direct study of a vast number of molecular systems that were not previously accessible. The application team demonstrated this capability by modeling the dynamics and reaction mechanism of mesoporous silica nanoparticles—a ubiquitous and highly effective catalyst used in chemical synthesis—using tens of thousands of atoms.

As a key research tool, GAMESS software is already used by more than 150,000 researchers worldwide. The advances that this application brings to computational chemistry will significantly accelerate technological development in active areas of research such as climate mitigation, materials design, and drug synthesis. These developments will pave the way for new carbon sequestration technologies to reduce the impact of the greenhouse effect, methods for generating useful compounds such as ethanol, and more effective medications and more sensitive screens for pathogenic illnesses. As computational chemistry continues to advance our fundamental understanding of chemical interactions, GAMESS ensures that current and future high-performance computing systems can be fully leveraged for these simulations.

The LatticeQCD Project: Using exascale computing to simulate subatomic interactions to advance fundamental physics research

Almost all the mass in the visible universe arises from interactions between quarks and gluons, which make up seven of the sixteen elementary particles in the Standard Model of physics. Exploring the ways in which these elementary particles interact is the central focus of modern particle physics, and provides valuable information for nuclear physics and astrophysics researchers. High performance computers have become an integral part of this process by interpreting information from high energy physics experiments at particle accelerators and providing powerful simulations of quark-gluon interactions to guide future experiments.

Quantum Chromodynamics (QCD) is the mathematical framework used to understand subatomic particle interactions within a nucleus, but applying the model to real-world interactions requires huge computational resources. The Exascale Computing Project's Lattice Quantum Chromodynamics (LatticeQCD) team uses next-generation supercomputers to model subatomic interactions at unprecedented resolutions and timescales. The team's applications are used by hundreds of researchers around the world to generate new experimental designs and to analyze and validate data more quickly than ever before.

Physicists have already used QCD to determine some properties of particles such as protons and neutrons, computing their masses and decay properties, but until recently they lacked the computational resources to study more fundamental particle interactions at sufficient spatial and temporal fidelity. Exascale computing has enabled the study of these interactions by guiding experimental research on current and future particle accelerators such as the Large Hadron Collider (LHC) at CERN and the Electron-Ion Collider (EIC) under construction at Brookhaven National Laboratory.

While exascale computing is now a reality, it could not be exploited fully to solve QCD problems without major advances in algorithms to take advantage of exascale architectures. The LatticeQCD team has created and optimized software for next-generation exascale computers, improving simulation speeds by a factor of more than fifty since 2016 and enabling more extensive gauge-field simulations—a critical tool for understanding the dynamics of quarks and gluons. These improvements will enhance the return on investment in the multi-billion-dollar LHC and EIC, leading to deeper understanding of how nature works at the tiniest length scales ever explored.

As the electric, atomic, and quantum revolutions have illustrated, advances in understanding of the fundamental laws of nature precipitate dramatic technological innovation. By enabling physicists to explore how elementary particles interact, LatticeQCD will contribute to the development of the Standard Model of physics, refine the theoretical underpinnings of nuclear and astrophysicists, and open the door to a new understanding of nature.

The MFIX-Exa Application: Using Exascale Computing to Advance Carbon Capture and Storage Technologies

The United States is adding renewable energy sources to its power grid at an increasing pace, but fossil fuels still generate about 60% of the grid's power. Although that number is decreasing, the US Energy Information Administration projects that it will still be more than 40% in 2050. Carbon capture and storage technologies such as chemical looping reactors (CLRs) will be critical in reducing the impact of greenhouse gases released from burning fossil fuels. However, current designs are still in laboratories and must be scaled to industrial proportions. Building and testing large- and intermediate-scale reactors are expensive and time-consuming tasks, so researchers are using high-performance computers to simulate various designs and identify the most promising candidates. Using this approach, carbon capture technologies can be developed at less expense and with greater assurance that reactors will meet necessary safety and performance standards.

The Exascale Computing Project (ECP) MFIX-Exa application was built to model industrial reactors for carbon capture applications with the highest possible fidelity The application models the coupled motion of particles and fluids within traditional power plants and industrial processes based on fossil fuels. These models will enable chemical looping combustion and point source capture, two key technologies used to capture CO_2 before it is released into the atmosphere. MFIX-Exa simulations accelerate the development of functional reactors that operate with low carbon emissions and minimal cost and energy penalties.

Modeling even small-scale reactors at a sufficient fidelity to track individual particles within a fluid flow requires enormous computational resources. Even a small 50-kilowatt thermal unit contains approximately 5 billion particles, and commercialized systems could easily contain trillions. High-fidelity physical simulations at these scales require billions of computational cells to represent fluid and particle dynamics, and this computational load is prohibitive without advanced software solutions to increase computational efficiency and access to the most powerful computers available.

The MFIX-Exa team enabled physically accurate CLR simulations on Frontier, the world's fastest supercomputers. The team collaborated with other ECP development efforts to enable new simulation methods such as adaptive mesh refinement and block-structured numerical algorithms, which improve computational efficiency without sacrificing simulation fidelity in critical regions. Furthermore, the MFIX-Exa application boasts extended physical modeling capabilities such as heat and mass transfer and chemical reactions for improved predictive power. These innovations coupled with the power of exascale computation helped the team to successfully model the National Energy Technology Laboratory's 50-kilowatt thermal CLR with almost a billion computational cells. This achievement marks the first simulation of a large-scale gas-solids chemical reactor with individual particle tracking.

Exascale computing with MFIX-Exa will be used to evaluate and improve emerging carbon capture technologies. These high-fidelity simulations will dramatically reduce development costs and will ensure that reactor designs scaled up from the lab will meet performance and safety targets. These accelerated technologies will be critical for reaching US decarbonization goals and will contribute to an energy sector free from carbon pollution in the decades to come.

The QMCPACK Application: Using Exascale Computing to Improve Quantum Simulations for Materials Design

Advanced materials design is critical for developing new technologies in fields ranging from energy science and manufacturing to transportation and aerospace. However, engineering lighter and stronger materials or discovering materials with completely new behaviors requires highly accurate models and immensely complex calculations to describe how atomic and molecular systems will interact. Researchers are increasingly turning to high performance computing to understand and accurately predict the quantum mechanical interactions which determine the properties and behavior of a material, taking advantage of new and more powerful machines along with a slew of novel theoretical and computational approaches. Quantum Monte Carlo (QMC) methods are among the most accurate tools for designing accurate materials simulations based on first principles, albeit at a high computational cost. QMC allows for highly accurate calculations and enables feasible testing of any approximations, rendering computational models of this type truer to nature than other methods.

The Exascale Computing Project's QMCPACK application is designed to apply the power of the world's fastest supercomputers to meet the enormous computational costs of advanced QMC simulations. Researchers use the QMCPACK software suite to accurately and reliably simulate both novel and known materials with unparalleled scale and efficiency. These simulations reveal material structures and interactions with high confidence and no artificial bias in results, leading to advances in both fundamental research and the design of new materials-based technologies.

QMC provides a solid theoretical basis for accurate and testable materials simulations, but the computational load involved in these simulations has proven prohibitive for all but the simplest of systems, such as materials in their bulk pure form. Without the power of exascale computing, large systems with complex characteristics—such as defects and intentionally added impurities or structural asymmetries—cannot be studied. Because these properties are often key to a substance's physical properties, exascale-enabled simulations are critical for broadly applying QMC to materials research.

The QMCPACK development team has met the exascale computational challenge by delivering an application that, while running on Frontier and Aurora, is capable of the largest and most accurate QMC simulations to date. The application supports highly accurate simulations of up to around 1000 atoms while accommodating material complexities such as asymmetries and impurities. The QMCPACK team has demonstrated these improved capabilities by simulating a range of scientifically relevant materials such as transition metal oxides—which are commonly used as catalysts and display properties including superconductivity—as well as novel nanoscale materials for optical applications and potential materials candidates for next generation electronic devices.

The advances to materials simulations afforded by QMCPACK will provide important predictive capabilities for materials where previous methods were unreliable. Furthermore, exascale-enabled QMCPACK simulations will allow researchers to perform highly accurate benchmark calculations on a wide range of materials, providing much-needed reference and calibration data for cheaper and less powerful computational methods. These data will be useful across the materials science, nanoscience, and physics communities, particularly where experimental data are costly or difficult to obtain, and can be used directly in machine learning and artificial intelligence applications.

Exascale Subsurface Project: Applying exascale computing for safe and reliable subsurface infrastructure

Underground infrastructure—such as wellbores for natural resource extraction, storage sites for hazardous materials, and portions of geothermal powerplants—is critical for supplying the U.S. with electricity and resources like natural gas. According to the Energy Information Administration, natural gas made up nearly 40% of U.S. electricity generation in 2022, and the EPA estimates that at least 1% of natural gas produced each year—about 6.5 million metric tons, is lost to leakage. Subsurface leakage is also a key factor in the safety and reliability of nuclear waste containment, geothermal energy production, and beyond. Applying high performance computing to model subsurface fracture formation will allow engineers to test different materials in various subsurface conditions as they are subjected to chemical and mechanical stresses, greatly improving the durability of subsurface infrastructure and supporting the safety and efficiency of resource extraction, power generation, and hazardous waste containment.

The Exascale Computing Project's Subsurface team has created a software suite for modeling subsurface fractures in extreme detail. The suite enables detailed analyses at scales which can accurately capture the chemical and mechanical stresses leading to the formation fractures and pillars—the microscale structures which keep fractures open. These simulations will allow researchers to compare the physical and chemical resilience of various materials, and ultimately to select the materials and locations most suited to the task at hand. The Subsurface code suite can also be used to determine the risk of leakage in preexisting subsurface infrastructure such as wellbores, which are built to variable safety standards.

Previous subsurface simulations were modeled on large scales from 100-1000 meters and operated on timescales up to 10 years. This approach is useful for modeling leakage in subsurface infrastructure like wellbores, but lacks the resolution and appropriate physical parameters to accurately determine the cause and location of subsurface fractures. Without more accurate models of fracture formation, and the features like pillars and asperities that keep the fractures open, researchers and engineers can only mitigate the effects of leakage rather than predicting and preempting it.

To address this issue, the Subsurface team used the power of exascale computing to integrate a single high resolution multiphysics simulation across scales ranging from kilometers to microns. The Subsurface code models fracture sites at micron scales while coupling information from a coarser simulation of the changes to materials surrounding the fracture, enabling detailed study of the causes and effects of fracture formation. The code suite also incorporates variable chemical compositions in subsurface infrastructure and surrounding host rock, allowing for comparative analysis between potential locations and materials.

Exascale-enabled models will improve our ability to select safe and reliable sites and materials to minimize fracture formation in subsurface infrastructure projects, and to understand the controls of what governs whether fractures open or close due to some combination of chemical and mechanical stresses. These improvements will enable several critical developments, including improved efficiency and energy generation from advanced geothermal power plants, more feasible and secure storage of nuclear waste and CO2, and safer and more efficient extraction of natural resources.

WarpX: Using Exascale Plasma Acceleration Simulations to Develop Compact Particle Accelerators

Particle accelerators such as the Large Hadron Collider at CERN are a foundational tool for research in modern physics and chemistry. Since 1939, more than 30% of all Nobel prizes in physics and 4 of the past 14 prizes in chemistry have been awarded for work with particle accelerators. These tools are now used to treat cancer and produce medical supplies, support research in applied fields such as pharmaceuticals, create microcircuits, and sterilize food—and scientists are continually experimenting with and implementing more uses. Researchers are motivated to create novel particle accelerator designs to discover new applications, solve more complex problems, and reduce the difficulty in constructing the large facilities needed to house accelerators built using legacy technology—which are often dozens of kilometers long. Plasma-based accelerators are among the most promising new design options because they support new imaging capabilities and can be deployed in far smaller facilities. However, significant unknowns exist in the physics of transport and acceleration of particle beams in the long chains of plasma channels used in plasma accelerators. Resolving these questions will require extensive calculations derived from computational models of plasma accelerators in unprecedented number and fidelity.

The Exascale Computing Project (ECP) WarpX application was created to deliver a software suite that can capture the full complexity of the acceleration processes within plasma-based designs. Scientists are using WarpX simulations to run virtual experiments on the functionality of various plasma accelerator prototypes to answer outstanding questions within key areas of research such as beam stability and quality. These simulations will increase the pace of accelerator development and will greatly reduce the costs associated with planning, constructing, and iterating on plasma-based collider designs.

Validating plasma-based particle accelerator designs will require a significant investment of computational resources. Researchers must conduct dozens or even hundreds of intricate simulations to measure key properties such as stability and energetic output during the tens to hundreds of plasma acceleration stages needed to achieve high-energy collisions—while accounting for a range of changing conditions and accelerator designs. These simulations require a combination of large scale and extreme fidelity, and legacy computing systems such as petascale supercomputers cannot complete such simulations in a feasible time frame. Practical plasma acceleration simulations will rely on improved computational output and innovative software approaches that improve efficiency while maintaining high fidelity.

The WarpX application team met this computational challenge by delivering an exascale-enabled system capable of the most precise plasma acceleration simulations to date. WarpX has boasted a 500× performance improvement since the project's inception in 2016 and can now simulate up to 20 consecutive stages of laser-driven plasma in a prototype multistage accelerator. WarpX uses methods such as adaptive mesh refinement to improve computational efficiency on exascale systems, thereby greatly enhancing simulation fidelity without ballooning real-time experimental durations.

Plasma-based accelerators promise advances to particle acceleration technology in basic research and applied fields. Furthermore, these machines will be reduced in size from kilometers to meters, thus greatly reducing construction, time, and costs and providing researchers and industrialists far easier access to accelerators. By validating plasma-based accelerators via WarpX exascale modeling, the ECP supports the rapid and cost-efficient development of tens of thousands of new particle accelerators for various applications that improve our lives and understanding of the world around us.