

Scaling Up Energy Innovation



Computer scientists Pythagoras Watson (left) and Teresa Kamakea work on Sierra, a high-performance computing (HPC) system that can process more than 260 trillion floating-point operations per second. Project teams supported by the hpc4energy incubator, a yearlong pilot program at Livermore, received time on this workhorse machine to accelerate development efforts on energy technologies.

through **Advanced Computing**

Collaborations through the hpc4energy incubator demonstrate how high-performance computing can help energy companies accelerate their research.

HIGH-PERFORMANCE computing (HPC) systems and software have already changed how science is done. Industry is its next frontier. With HPC, computing systems with millions of processors run calculations simultaneously—in parallel—instead of sequentially, often slashing run times from weeks to hours or minutes. This massive increase in computational power and speed has allowed scientists to explore research problems that, because of scale or complexity, have previously been impossible to model.

Despite the promise of HPC for innovation, dedicated and experienced industrial users remain a relatively small group. Large communities of researchers and technology developers within both the public and private sectors use desktop workstations exclusively, and as a result, they have complex problems that remain unsolved. Many segments of the energy sector, for instance, rely solely on experiments and physical prototypes for product development. (See *S&TR*, December 2011, pp. 4–11.) Trillions of calculations per second may sound impressive, but companies want further proof that HPC modeling and simulation will benefit their business model before they move beyond what commercial software and a modest in-house computer cluster can offer.

Lawrence Livermore's hpc4energy incubator (see *S&TR*, June 2012, pp. 24–25) has demonstrated how national laboratories and companies might partner to expand

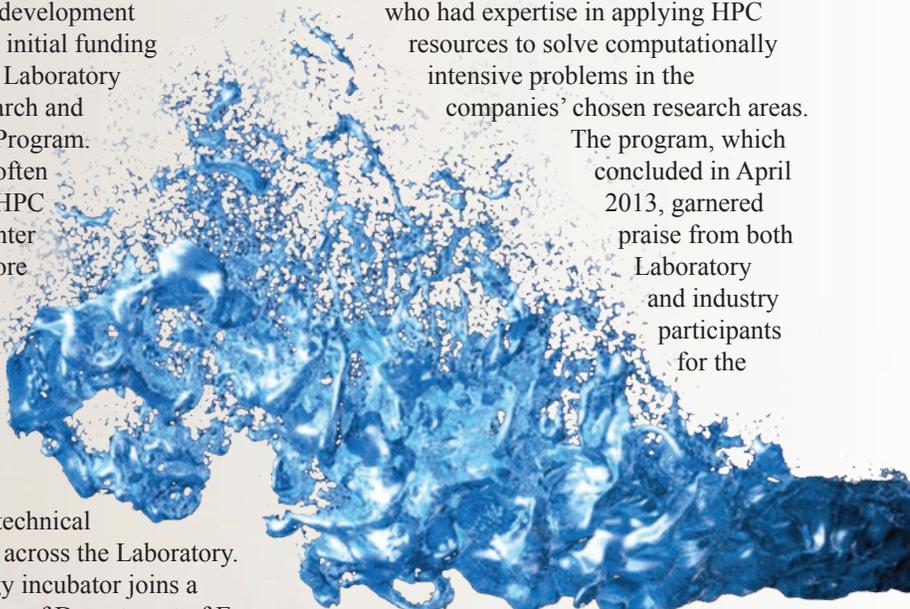
the adoption of HPC, a tool capable of addressing such national grand challenges as energy security, economic prosperity, and scientific leadership. The one-year pilot program was designed to exhibit the benefits of incorporating HPC modeling and simulation into energy technology development.

The hpc4energy incubator is part of Livermore's broader industrial outreach and economic development initiatives, and initial funding came from the Laboratory Directed Research and Development Program. Project teams often worked at the HPC Innovation Center on the Livermore Valley Open Campus, which provides industrial clients with access to computing resources and technical expertise from across the Laboratory. The hpc4energy incubator joins a growing number of Department of Energy–sponsored partnerships that are helping to connect energy businesses with HPC resources and meet the nation's carbon emission and energy security goals.

The six industrial participants, selected through a competition, differ in their level of HPC experience and the problems they want to examine, but they share one

commonality. "These are all forward-thinking companies willing to embrace new technologies and research tools," says former Livermore engineer Clara Smith, who managed the incubator program. Each company was granted time on Sierra, a supercomputing system that can perform more than 260 trillion floating-point operations per second. Project teams were matched with Livermore staff members who had expertise in applying HPC resources to solve computationally intensive problems in the companies' chosen research areas.

The program, which concluded in April 2013, garnered praise from both Laboratory and industry participants for the



GE Global Research and Livermore coupled sophisticated numerical methods and high-performance computing (HPC) to create a high-resolution simulation of liquid fuel spray. This research effort focused on reducing the number of design iterations needed to create advanced fuel injectors. (Courtesy of Computational Thermo-Fluids Laboratory, Cornell University.)

exposure to new research tools, methods, and collaborators.

John Grosh, deputy associate director for Computation Programs, says, “I was pleased to see how enthusiastic and committed the six companies were to our collaboration. They knew the program was a unique opportunity, and they were intellectually engaged and responsive.”

Whether the goal was designing a more efficient combustion engine, simulating complex energy networks to improve planning and scheduling, refining a new geothermal drilling technology, or analyzing energy use in buildings, the partners concurred that Livermore supercomputers and expertise accelerated the cycle for product and service development. The incubator project encouraged them to approach research problems in a new way. “When you have computing capabilities you didn’t have before, you think differently about problems,” says Eugene Litvinov, the senior director for Business Architecture and Technology at ISO New England. “You can ask questions you didn’t think of asking before.”

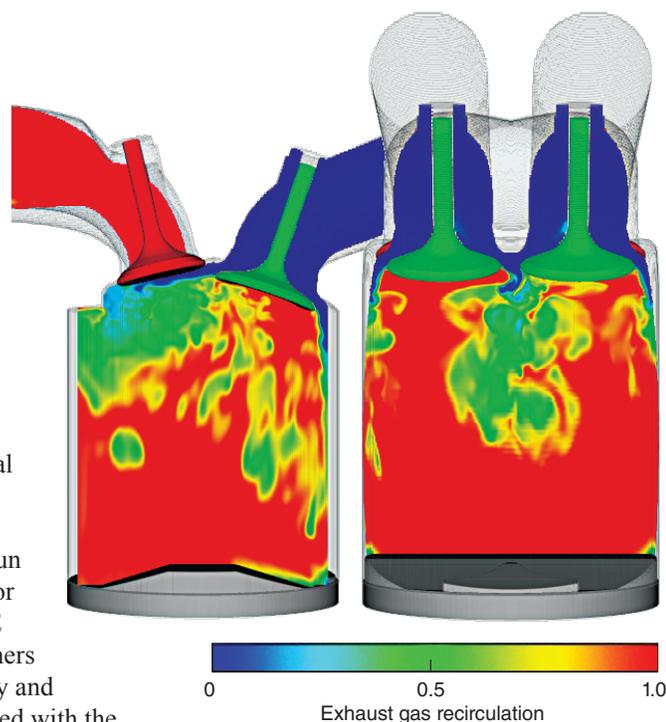
Injecting HPC into Fuel Simulation

Aircraft and automobile manufacturers are interested in developing more-efficient engines that last longer and generate less pollution, but the current engine optimization cycle of prototype design, fabrication, and laboratory testing is onerous. Designers cannot see inside or measure many engine systems during laboratory experiments. Conducting a series of tests that replicate the full range of potential operating conditions is impractical. As a result, the current prototyping and testing process can be expensive and inefficient. Two hpc4energy participants, GE Global Research and Robert Bosch, LLC, are augmenting laboratory measurements with modeling and simulation to better understand combustion, turbulence, and other processes that affect engine performance.

An engine’s overall performance hinges in part on how well the injected liquid fuel disperses into a turbulent spray of droplets that mix with an oxidizer and burn. Scientific understanding of turbulence has improved in recent years, largely because of new high-fidelity numerical techniques, but industry researchers often do not have the computational power to run the software. For the incubator effort, jet engine designer GE Global Research and researchers from Arizona State University and Cornell University collaborated with the Laboratory to deploy numerical methods developed by the universities on an HPC machine. Using these codes on Sierra, team members simulated turbulent spray breakup in three dimensions and, with those results, evaluated two designs for fuel-injector engines.

For the test simulation, the incubator team examined liquid fuel entering the combustion chamber to determine how the shape of the opening affects spray breakup as fuel droplets intersect with a turbulent cross-flowing current of air. Both numerical codes simulated the problem at a range of resolutions, in some cases with droplets as small as 20 micrometers in diameter. These simulations used nearly 16 million core-hours of computing time on Sierra, by far the largest demand for HPC resources among the incubator projects. A typical run took more than 3 days of continuous computing on more than 11,000 cores.

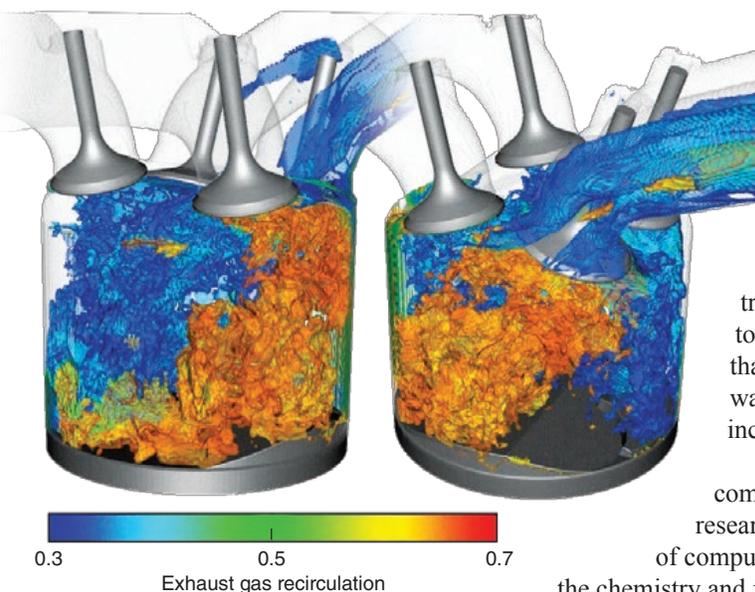
Postprocessing analysis of the simulation data sets and comparison with existing experimental measurements are ongoing, but the initial findings have been encouraging. For instance, results showed that Cornell’s NGA code effectively captures the shape and distribution of fuel droplets as the fuel and air mix. The simulations also



Researchers at Robert Bosch, LLC, have developed a three-dimensional numerical code to simulate how exhaust gas recirculates in the combustion chamber of a new engine design. Calculations run on the Bosch in-house computing cluster show fresh air (blue) entering the combustion chamber and interacting with leftover combustion products from the previous cycle.

indicated that experimentally drawn conclusions accurately describe how the injector pipe’s shape affects fuel droplet size, speed, distribution, and shape. These factors are important because they influence the efficiency, durability, reliability, and safety of engine operation.

Livermore computational physicist Gregory Burton observes, “The fidelity of the simulations has been astonishing. The work provides an unprecedented opportunity to delve deeply into the physical processes governing droplet formation and ultimately learn how to control these processes to develop higher-performance engines.” Gary Leonard, the global technology director at GE Global Research, adds that through the incubator effort, “We’ve started to learn about some



In their hpc4energy project, Bosch researchers worked with Laboratory computer scientists to scale the Bosch code to run efficiently on Livermore's Sierra supercomputer. They reduced the calculation time for each engine cycle by 70 percent and improved the model resolution, allowing them to examine combustion in greater detail than the Bosch computer system can produce.

of the physics going on in our jet engines that we didn't know about. And we've been building jet engines for 60 years."

Combustion Switches Gears

Certain advanced combustion technologies in automobile engines could reduce fuel consumption by more than 30 percent in low- to moderate-load driving conditions, such as while idling or cruising at a steady speed. Unfortunately, these technologies cannot provide the same performance level as conventional combustion in high-load circumstances, for example, when merging onto a highway or starting from a complete stop. Engines currently under development combine the two combustion modes to optimize performance and fuel efficiency. Switching between modes, however, will require

advanced control algorithms. Before those calculations can be developed, researchers must better understand the complex chemistry and

physics involved in the transition from one mode to the other—a problem that Bosch collaborators wanted to evaluate in their incubator project.

Using their in-house computer cluster, Bosch researchers needed two weeks of computational time to calculate the chemistry and flow through one engine cycle. Understanding how a set of timing parameters affects the transition from conventional to high-efficiency combustion requires modeling 10 four-stroke engine cycles: two cycles in conventional mode, followed by eight in high-efficiency mode. Processing those calculations on the Bosch system would require 20 weeks of uninterrupted computational time—much too slow to make meaningful progress.

For the incubator project, the collaborators scaled the sophisticated Bosch codes to run efficiently on the Laboratory's computing system. The team reduced the calculation time for each cycle by 70 percent, making it possible for the first time to simulate the full transition process. Expedited processing times also enabled the team to complete multiple sets of test cycles and vary parameters such as the fuel-injection method to more thoroughly validate the combustion model.

Livermore mechanical engineer Dan Flowers says, "Combustion is one of the hardest processes to model because events are occurring at a wide range of size and timescales simultaneously." Capturing the fine-scale effects of combustion was the most important factor in this study, so the team focused its effort on computationally expensive simulations at extremely high

resolutions. The researchers observed some unusual and possibly significant physical behaviors, such as high-frequency pressure wave effects, that would rarely if ever appear in the results of less-detailed simulations.

Through the hpc4energy collaborations, GE Global Research and Bosch had access to far more computing power than is available to either group internally, allowing their researchers to perform multiple tests and study complex physical phenomena with more precision than they could have otherwise. The insights these teams gained into spray breakup and combustion modes will help the companies improve the models they use for engine design and testing.

Livermore researchers emerged with a better understanding of the engineering and simulation problems that most concern these two participants and similar energy companies. Laboratory scientists also gained first-hand experience with two advanced numerical techniques developed by GE Global Research to examine spray breakup problems. In addition, the GE codes will now be available for use in other research projects through the new Turbulence Analysis and Simulation Center at Livermore.

Anticipating Changes in the Grid

In its 2003 book *A Century of Innovation: Twenty Engineering Achievements That Transformed Our Lives*, the National Academy of Engineering identified the U.S. electric power grid as the greatest engineering achievement of the 20th century. Modernizing the grid's infrastructure and tools to meet 21st-century needs poses a significant challenge. New, "smart" grids incorporate more sensors and automated controls to operate efficiently and better prevent interruptions.

However, many organizations involved in grid planning and scheduling have found that the accompanying increases in data volume and system complexity are straining available processor performance.

In addition, integrating intermittent power sources such as wind and solar energy into the grid can confound planning efforts. To address these issues, two incubator teams examined how HPC could improve energy-grid modeling and planning for complex networks.

In many parts of the world, electricity is now a dependable and vital resource. If a power line is damaged by a falling tree or a connection is offline for maintenance, electricity must keep flowing or quickly be restored to customers. Utility planners often use the GE Positive Sequence Load Flow (PSLF) software to predict how a system’s events might affect energy transmission. PSLF analyzes the systemwide effects of failed components for a given grid configuration and pinpoints configurations that will continue to operate in the event of such failures.

When calculations are made on a desktop computer, the contingency analyses, or what-if scenarios, must be examined consecutively. Completing the calculations for larger grid systems can take hours or days. To reduce the turnaround time, planners may lean on their expertise. For example, they might identify the most-likely failures and simulate only those scenarios. However, this type of selective testing increases the chance of overlooking a critical failure point.

The PSLF developer, GE Energy Consulting, wanted to scale up the software model while reducing the computational load. Through the incubator program, GE Energy could run the software on more powerful computing systems and simulate extremely large networks, many times the size of those previously simulated on the company’s computing cluster. The collaboration also enabled GE Energy scientists to perform more contingency analyses than they could process in-house. Livermore computer scientist Steve Smith says, “By running a more comprehensive simulation on our machines, the GE team increased confidence in the results produced by PSLF.”

Laboratory researchers optimized the PSLF code to run in parallel on Sierra, thereby reducing the run time for all contingencies to the time required to solve the longest running contingency. In a study of 4,217 contingencies, the total calculation time on Sierra was only 23 minutes. Processing this number of analyses consecutively would have required an estimated 23.5 days. Devin Van Zandt, software products manager for GE Energy Consulting, says, “Working with the Laboratory team to improve the performance of our code was a great opportunity. Their experience in solving problems similar to ours coupled with the

HPC infrastructure at Livermore helped us understand the potential of our code base.”

After the successful Sierra test, GE Energy worked with Livermore experts to speed up the processing on individual contingency analyses. GE Energy researchers are now studying how to incorporate these scaling and performance improvements into a new version of PSLF.

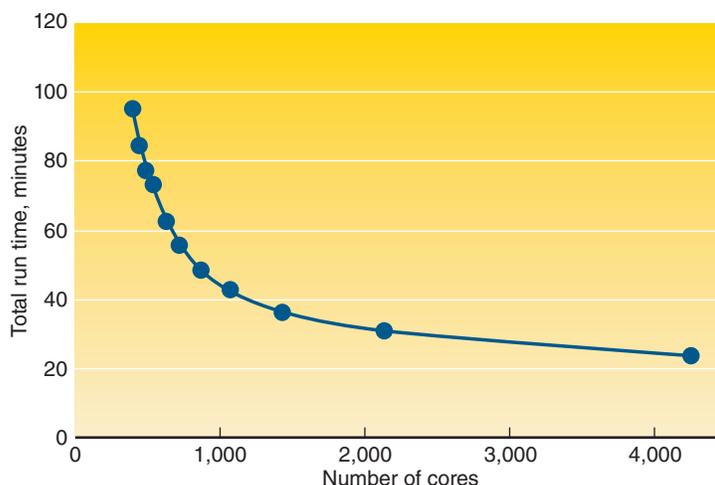
Tomorrow’s Forecast for Energy

In scheduling electric grid operations, system operators evaluate the supply and demand, or load conditions, forecast for the following day and, through a process called unit commitment, determine which power generators will be used. If the day’s forecast is inaccurate—say, for example, the amount of wind energy available is significantly below projections—operators may not be able to start up the uncommitted power sources quickly enough to alleviate the shortfall. Traditional models for unit commitment are thus designed to generate conservative estimates, an approach that is neither very cost effective nor efficient in its use of renewable resources.

ISO New England, the company responsible for power generation and transmission throughout that region, enlisted Livermore’s help in comparing this point forecast approach with a new method that accommodates a range of

uncertainty. A preliminary experiment by ISO New England indicated that the newer technique, termed robust unit commitment, could substantially reduce dispatch costs and improve reliability. Assessing the operational and economic benefits of robust commitment required generating day-ahead forecast ranges for projected electric load and renewable power generation. These ranges were then used to identify an optimal unit commitment schedule for

GE Energy performs contingency analysis to determine the stability of an electric grid when a connection is removed from the system, for example, for required maintenance or because of a damaged power line. With the increased computational power provided by Sierra, GE Energy could simulate all of the contingencies for a large network simultaneously rather than in sequence, producing more rapid and thorough results.



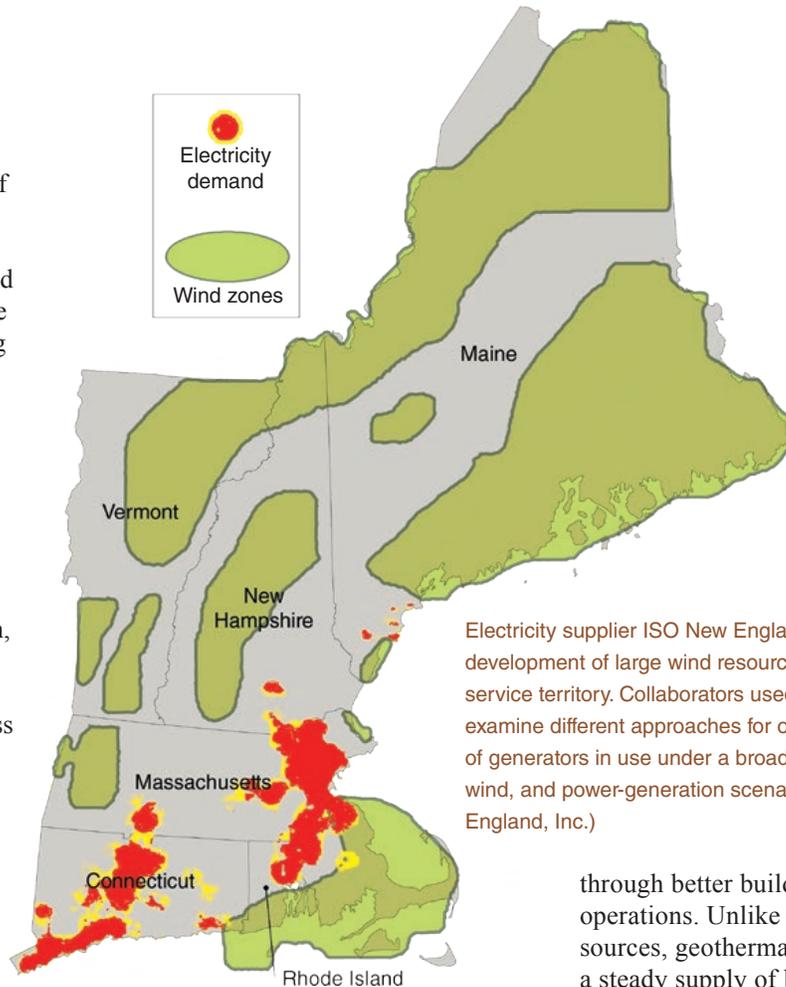
the day. The final step was testing the system dispatch over that day for each of many possible realizations of load and renewable power generation.

Since each unit on the grid is scheduled to be either on or off, optimizing schedule scenarios is a computationally demanding mathematical problem that leads to millions of possible combinations. For a set of 1,600 robust unit commitment configurations, for example, solving each scenario takes about 30 minutes, or 800 hours of processing time on one desktop computer. Up to 10,000 dispatch problems, each one taking 15 seconds, must be answered for every configuration, for a total dispatch-solving time of 67,000 hours on a desktop.

By parallelizing the software to process the robust problem set and running it on 1,600 core processors, the hpc4energy collaboration reduced the calculation time from 800 hours to 90 minutes. After demonstrating that the optimized code could efficiently process such a large number of scenarios, the team completed more than 10,000 simulations using Monte Carlo sampling based on historical load and wind-generation data combined with a Livermore-developed statistical model of this behavior. The team also evaluated the optimal size of the uncertainty range for the robust approach.

Processing thousands of occurrences at once provided researchers with a more comprehensive evaluation of the next day's schedule and allowed them to assess the effectiveness of the robust unit commitment approach under a broad range of energy use, wind, and power-generation scenarios. "We can't explore all of the possibilities because it would cause what is appropriately called combinatorial explosion," says Livermore computer scientist Barry Rountree. "But we can look at thousands more scenarios than ISO New England could with their in-house computing resources."

The Livermore team's greatest contribution in the ISO New England



Electricity supplier ISO New England is studying future development of large wind resources throughout the company's service territory. Collaborators used the hpc4energy project to examine different approaches for optimizing the daily schedule of generators in use under a broad range of energy use, wind, and power-generation scenarios. (Courtesy of ISO New England, Inc.)

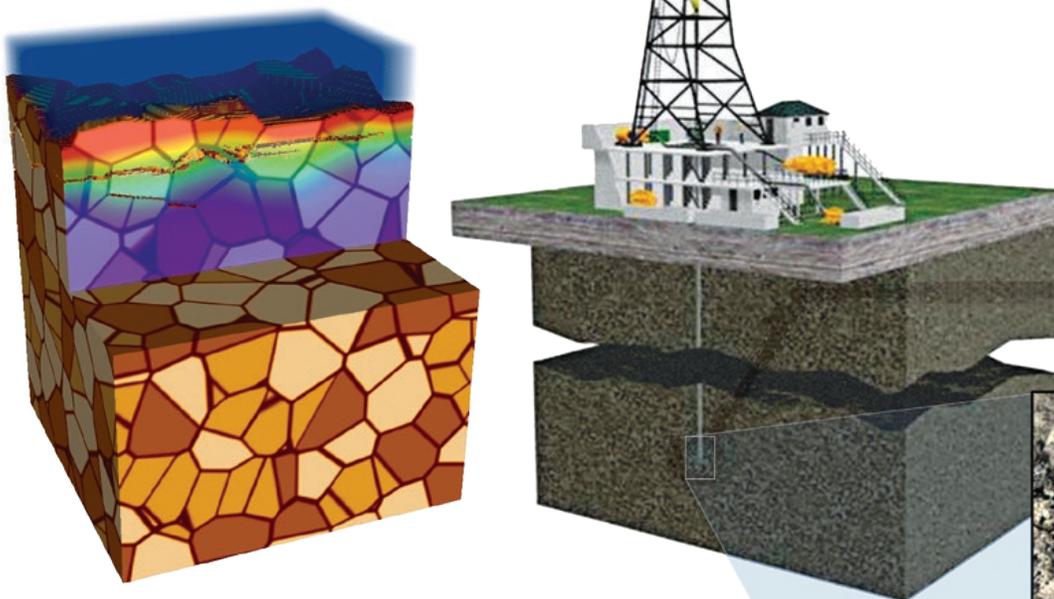
project was in statistical modeling and visualization. Laboratory researchers helped frame the problem and prepared a statistical model to simulate wind and load for the Sierra calculations. In the GE Energy effort, the Laboratory team provided a third-party evaluation of the GE application and demonstrated how larger-scale HPC resources could benefit that work. Using Sierra to efficiently run thousands of simulations, the two energy companies gathered statistically significant results that they can use to evaluate and develop sophisticated scheduling and contingency analysis software.

Modeling a New Drilling Technique

The final two incubator efforts focused on making a clean energy source more cost competitive and on saving energy

through better building design and operations. Unlike other renewable power sources, geothermal resources generate a steady supply of base-load energy without the need for storage or power-grid modifications. However, capital costs and potential economic risk have deterred investors. Geothermal wells are drilled up to 6 kilometers underground in hard rock formations such as granite, which have a high heat capacity. Because of the depths involved and the difficulties associated with penetrating hard rock, mechanically drilling the wells is expensive.

Potter Drilling, a small start-up company, has developed a technology called hydrothermal spallation drilling that could reduce drilling time and cost and make existing wells twice as productive. Thermal spallation uses a jet of superheated water to penetrate granite at two or more times the rate of conventional technologies, without the drill bit ever contacting the rock. Livermore mathematician Stuart Walsh notes that once a well location is determined, "The economics of drilling a well is the time



Potter Drilling worked with Laboratory researchers to create a grain-scale model (shown at far left on a 1-centimeter cube of granite slab) of hydrothermal spallation drilling, a new technique (below) that efficiently removes rock from a well bore. Hydrothermal spallation could reduce the time and cost for drilling deep geothermal wells and could make existing wells more productive.

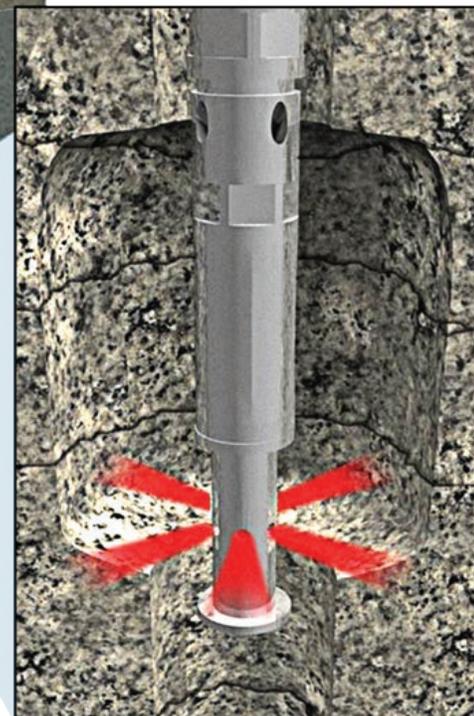
involved. If you have a technology that allows you to drill faster or replace drill bits less frequently, then you gain a competitive advantage.”

Before its collaboration with Lawrence Livermore, Potter Drilling relied on laboratory experiments and field testing. Unfortunately, most of the relevant work occurs in deep boreholes where it is impossible to monitor the process directly, making system optimization difficult, slow, and expensive. In the two years prior to the incubator project, Potter engineers completed 15 field trials and 30 design and process changes. Through the hpc4energy collaboration, the company conceptually tested operations under a broad range of conditions, such as encountering different rock types or drilling in deeper wells with higher pressure than could be experimentally simulated.

The incubator team developed a model to study hydrothermal spallation drilling at the mineral grain scale and combined it with Livermore’s GEODYN and PSUADE codes to determine how about two dozen model parameters affect the drilling process. Even with the computational support provided through the incubator, completing three-dimensional simulations

for the full range of circumstances would have been computationally prohibitive. Instead, the researchers first ran a set of over 7,000 less-taxing two-dimensional simulations, each using 72 computer processors. With those results, they constrained the number of three-dimensional simulations required, each of which would run on 1,020 processors. Parameters and boundary conditions for the modeling studies were derived from Potter’s past experiments, and the results were compared to data acquired in field tests of the new technique.

Jared Potter, cofounder of Potter Drilling, says, “Initial results of the Livermore modeling studies have shown a good replication of the key features thought to occur during the spallation process, including the initiation of spalls and movement of the damaged zone into the rock with time. Future collaborative work could lead to predictive capabilities and the ability to change drilling tool operating conditions, depending on such parameters as rock type or drilling depth.” Working with the 19-employee company was a valuable experience for Livermore participants, highlighting the needs and concerns of start-up companies, whose



core business can rapidly evolve. The collaboration also produced software tools that will be valuable to future energy and drilling research efforts at the Laboratory.

Factoring in Uncertainty

Commercial and residential buildings account for more than 40 percent of U.S. energy consumption. Making buildings more energy efficient could thus generate substantial savings in operational costs and energy use while reducing carbon emissions. Predicting a building’s efficiency requires engineers and planners

to quantify uncertainties in the parameters they use to estimate energy consumption for a building's lifecycle, from design and construction to operations and maintenance. The actual energy usage for a building is often up to 30 percent greater than originally projected because of variability in building materials, occupant behavior, weather, and maintenance. If the primary factors causing uncertainty can be identified and analyzed through simulation, an accurate risk assessment can be made prior to investing in a building retrofit, making energy-efficient designs and remodeling more attractive to investors.

Through the incubator program, United Technologies Research Center (UTRC) identified key operational variables for predicting building energy performance and estimated some effects of such uncertainty. UTRC researchers will use this information to validate and calibrate the company's models and to advise builders and building operators on a facility's use patterns and areas to target for improvement. "For existing facilities, the operations staff can determine a sustainable or wasteful path for building use," says Livermore engineer Noah Goldstein. "The people who run these buildings are really starting to understand the value HPC can add to operations."

UTRC researchers had incorporated parallel computing into their analysis process prior to the incubator project, reducing the turnaround time for whole-building analysis from 2 weeks to less than 24 hours. But their analysis was somewhat limited in scope. With the Livermore collaboration, they expanded their optimization effort to run the simulations in a massively parallel form. Working with an energy use model of an office building located in the Philadelphia Navy Yard, the incubator team tested variations in more than 900 operating parameters, including such factors as the performance of heating and air-conditioning equipment, energy used by lighting fixtures and appliances, insulation properties of different walls and



- Other, 2 percent
- Maximum flow rate, 7 percent
- Pressure rise, 28 percent
- Fan efficiency, 28 percent
- Minimum intake flow fraction, 35 percent

Researchers from United Technologies Research Center used Livermore's HPC resources to analyze which parameters most affect the energy efficiency of an office building. In reviewing air-handling fan units—one of the largest energy "consumers" in a building—the collaboration found that only 3 of the 917 operating parameters examined significantly alter a unit's electricity consumption. This kind of analysis provides vital information because it reveals precisely how retrofit designs will affect a building's energy performance.

windows, weather changes, and occupant activities in the building. On Sierra, operating with 1,000 computer cores per simulation, calculation time was reduced by a factor of 60, permitting UTRC to efficiently complete 10,000 simulations and accelerate uncertainty modeling. Analysis turnaround time was reduced from days to only hours.

Goldstein notes that collaborating with UTRC helped Laboratory scientists better understand how to simulate the energy consumption of a building. It also allowed them to compare approaches of uncertainty quantification to Livermore-developed methods.

The company is now comparing the model results and sensitivity levels from the incubator project with a year's worth of data on energy usage and occupant behavior from the Philadelphia building. The UTRC analysis indicates that, for simulating a typical building, only a small number of parameters has the greatest effect when modeling energy use. Variations in all other parameters have

a negligible effect. One such example is the electricity consumption of a building air-handling system. Results from the HPC models suggest that the system's performance depends primarily on only three parameters, while the combined effect of all other parameters is less than 10 percent. This finding, if proven accurate, would simplify a building retrofit design, because it gives clear guidance to the engineers regarding where to focus their effort and resources.

In the project's final months, the incubator team began designing a more efficient refrigeration system to meet the needs of certain commercial building users. HPC has rarely been used to consider heating and air-conditioning designs because such computations cost more than the retail price of the system. The incubator collaborations are thus helping address a wider range of engineering problems.

Although UTRC is already an HPC user, the company benefited from working with researchers who had an established HPC-based workflow for building analysis

and were experienced in quantifying uncertainty. Both the Potter Drilling and UTRC efforts show how advanced computing can enable companies to quickly narrow the list of significant variables, freeing up time and resources to focus on optimizing the factors that will likely yield the greatest improvements.

Problems on a Whole New Scale

The Laboratory's hpc4energy incubator helped demonstrate that businesses of all sizes can benefit from applying HPC to reduce risk and optimize designs. Through the pilot program, energy companies observed new details about the physical processes involved, modeled problems with extremely fine resolution, and captured the variability in the systems they were evaluating. Livermore participants learned about the energy industry's requirements and goals. In the process, they helped solve new problems, worked with unfamiliar algorithms and methods, and expanded their computational skills.

Initiatives such as the hpc4energy incubator support the Laboratory's mandate to transfer technology and expertise to the private sector. At a November 2012 workshop, Laboratory Director Penrose (Parney) Albright noted,

"The Lab's culture is about bringing together multidisciplinary teams to work on problems of national importance. Now, we want to take some of the capabilities we've developed and put them in the service of the U.S. economy."

In addition, fostering wider adoption of and demand for HPC drives down hardware and software costs, which benefits Lawrence Livermore and its mission work. "Broadening the HPC user base makes it easier and more affordable to push the cutting edge," says Livermore computer scientist Rob Neely. "Everyone wins."

With an effective template for industry engagement in place, the Laboratory's HPC leaders are considering which other research areas to "incubate." Several industrial sectors with urgent research problems have yet to fully embrace HPC and could benefit from Livermore's computational expertise. For instance, an HPC incubator for bioinformatics and pharmaceutical design could expand the traditional "number-crunching" capabilities of HPC simulation by marrying it with the Laboratory's strengths in molecular dynamics and big data analysis. (See *S&TR*, January/February 2013, pp. 4–11.) Such collaborations have the potential to

produce highly precise results and reduce the need for laboratory testing.

The competitive nature of business often makes companies reluctant to share their HPC success stories. An effort such as hpc4energy can strengthen the case for investing in this technology and can teach prospective users to think differently about problem solving. "Because of limited computational resources, many researchers have gotten used to constraining their models in resolution and scale," says Goldstein. "We've shown them that they can model a whole system, do full-scenario planning, or even do real-time modeling. The hpc4energy incubator has helped crack open the box of what simulation and modeling can do. It's not just faster—it's about thinking on different scales."

—Rose Hansen

Key Words: clean energy technology, contingency analysis, fuel-injection engine, high-efficiency combustion, high-performance computing (HPC), hpc4energy incubator, HPC Innovation Center, hydrothermal spallation drilling, parallel computing, smart electric grid, turbulence, uncertainty quantification, unit commitment.

For further information contact John Grosh (925) 424-6520 (grosh1@llnl.gov).