

Accelerating Computational Science Symposium 2012

Hosted by:

The Oak Ridge Leadership Computing Facility at Oak Ridge National Laboratory

The National Center for Supercomputing Applications at the University of Illinois at Urbana-Champaign

The Swiss National Supercomputing Centre at ETH Zurich



Scientists Assess a New Supercomputing Approach and Judge It a Win

Graphics processing units speed up research codes on hybrid computing platforms

Can scientists and engineers benefit from extreme-scale supercomputers that use application-code accelerators called GPUs (for graphics processing units)? Comparing GPU accelerators with today's fastest central processing units, early results from diverse areas of research show 1.5- to 3-fold speedups for most codes, which in turn increase the realism of simulations and decrease time to results. With the availability of new, higher-performance GPUs later this year, such as the Kepler GPU chip to be installed in the 20-petaflop Titan supercomputer, application speedups are expected to be even more substantial.

An innovation disrupts the marketplace when it confers an undeniable advantage over earlier technologies. Developing a disruptive innovation—such as the telephone that so long ago replaced the telegraph—can be a high-risk, high-rewards venture. Early in a product's life, it may be unclear whether customers will derive enough gain from a new technology to justify its developers' pain. Then unprecedented success compels a community to embrace the innovation. With supercomputers, which must balance performance and energy consumption, hybrid architectures incorporate high-performance, energy-efficient scientific-application-code accelerators called graphics processing units (GPUs) as well as traditional central processing units (CPUs). For scientists who must beat competitors to discoveries in a publish-or-perish world or engineers who must invent and improve products for the global marketplace, the ability to more fully exploit massive parallelism is proving advantageous enough for GPU/CPU hybrids to displace CPU-only architectures.

The Accelerating Computational Science Symposium 2012 (ACSS 2012), held March 28–30 in Washington, D.C., brought together about 100 experts in science, engineering, and computing from around the world to share transformational advances enabled by hybrid supercomputers in diverse domains including chemistry, combustion, biology, nuclear fusion and fission, and seismology. Participants also discussed the new breadth and scope of research possible as the performance of petascale systems, which execute quadrillions of calculations per second, continues to increase. At the start of the symposium, neither the scientific community nor the government officials charged with providing supercomputing facilities to meet the needs of researchers knew whether hybrid computing architectures conferred incontestable benefits to the scientists and engineers using leadership-class computing systems to conduct digital experiments. To understand the impact hybrid architectures are making on science, attendees heard from two dozen researchers who use a combination of traditional and accelerated proces-



Jack Wells, ORNL

sors to speed discoveries in fields from seismology to space and aid innovations such as next-generation catalysts, drugs, materials, engines, and reactors. The overwhelming consensus: GPUs accelerate a broad range of computationally intensive applications exploring the natural world, from subatomic particles to the vast cosmos, and the engineered world, from turbines to advanced fuels.

“The symposium [was] motivated by society’s great need for advances in energy technologies and by the demonstrated achievements and tremendous potential for computational science and engineering,” said meeting program chair Jack Wells, director of science at the Oak Ridge Leadership Computing Facility (OLCF), which co-hosted ACSS 2012 with the National Center for Supercomputing Applications (NCSA) and the Swiss National Supercomputing Centre (CSCS). Supercomputer vendor Cray Inc. and GPU inventor NVIDIA helped sponsor the event. “Computational science on extreme-scale hybrid computing architectures will advance research and development in this decade, increase our understanding of the natural world, accelerate innovation, and as a result, increase economic opportunity,” Wells said.

Concurring was James Hack, director of the National Center for Computational Sciences, which houses the U.S. Department of Energy’s (DOE’s) Jaguar, the National Science Foundation’s (NSF’s) Kraken, and the National Oceanic and Atmospheric Administration’s (NOAA’s) Gaea supercomputers. “As a community we’re really at a point where we need to embrace this disruptive approach to supercomputing,” he said. “This meeting was an opportunity to catalyze a broad discussion of how extreme-scale hybrid architectures are already accelerating progress in computationally focused science research.” At its conclusion, attendees had a better idea of where the community stood in exploiting hybrid architectures.

The ACSS 2012 researchers presented results obtained with four high-performance systems employing different computing strategies. First, CSCS’s 402-teraflop XE6 Cray supercomputer, called Monte Rosa, at the Swiss Federal Institute of Technology in Zurich (ETH Zurich) uses two modern 16-core CPUs from AMD named Interlagos per compute node. Second, a two-cabinet XK6 supercomputer called Tödi, the first hybrid



Monte Rosa, the Swiss National Supercomputing Centre’s 402-teraflop XE6 Cray supercomputer, uses only CPUs. The entire computational-science community benefits from efforts to optimize CPU performance. Image credit: Michele De Lorenzi/CSCS



Tödi, the Swiss National Supercomputing Centre’s 137-teraflop XK6 Cray supercomputer, is the first GPU/CPU hybrid supercomputing system. Image credit: Michele De Lorenzi/CSCS

system from Cray with Interlagos CPUs and NVIDIA Fermi GPUs, has been in operation at CSCS since fall 2011. Third, NCSA’s Blue Waters is a hybrid Cray system under construction at the University of Illinois at Urbana-Champaign that gives more weight to CPUs than GPUs. It will have only the Interlagos CPUs in 235 XE6 cabinets and both CPUs and the NVIDIA Fermi GPUs in 30 XK6 cabinets. A testbed of 48 XE6 cabinets has been operational since mid-March. Fourth, the OLCF’s 200-cabinet Jaguar supercomputer, America’s fastest, is undergoing an upgrade at its Oak Ridge National Laboratory (ORNL) home that will transform it into a mostly hybrid

computational giant called Titan. So far the upgrade has converted the Cray XT5 system, an all-CPU model with an older, less-capable version of the AMD processor, into a Cray XK6 system with Interlagos CPUs, twice the memory of its predecessor, and a more capable network. After the cabinets have been equipped with NVIDIA’s newest Kepler GPUs by the end of 2012, the peak performance will be at least 20 petaflops. In the meantime, ten of Jaguar’s cabinets have already gone hybrid to create a testbed system called TitanDev.



The National Center for Supercomputing Applications' Blue Waters, under construction at the University of Illinois at Urbana-Champaign, is a hybrid Cray system that gives more weight to CPUs than GPUs. Image courtesy of the National Center for Supercomputing Applications and the Board of Trustees of the University of Illinois


At the start of ACSS 2012, managers of extreme-scale supercomputing centers and program managers from DOE's Office of Science wanted to know if researchers would rather use an all-CPU XE6 or a hybrid GPU/CPU XK6. The results presented at the meeting laid that question to rest. Most scientific application codes showed a 1.5- to 3-fold speedup on the GPU-equipped XK6 compared to the CPU-only XE6, with evidence of a 6-fold speedup for one code. The greater computing speed makes possible simulations of greater complexity and realism and quicker solutions to computationally demanding problems. "The Fermi [GPU from NVIDIA] delivers higher performance than Interlagos [CPU from AMD]," said meeting co-chair Thomas Schulthess, who is head of CSCS and a professor at ETH Zurich. "And with Fermi we're talking about a GPU that is 1 or 2 years old compared to Interlagos, which is a very new processor that just came out."

Compute Node Architectures of Select Cray Supercomputers

	XT5	XE6	XK6 (w/o GPU)	XK6	Next-Gen XK
AMD Istanbul CPU (6 cores)	2				
AMD Interlagos CPU (16 cores)		2	1	1	1
NVIDIA Fermi				1	
NVIDIA Kepler					1

Added Schulthess: "Given the applications that we've seen and the benchmarks that were representative of real science runs, I haven't seen the case where the GPU doesn't provide a performance increase. Really, it makes no sense not to have the GPU."

The scientific results presented at the symposium were way ahead of what the high-performance computing community had expected, Schulthess said. Scientists from diverse fields of study were able to demonstrate impressive results using hybrid architectures.



Thomas Schulthess, ETH Zurich and CSCS

"The XK architecture is looking like it will be a big success," Schulthess opined. "It's going to be interesting to see in the next few years if there's going to be a small avalanche, or is

a big avalanche coming that's really going to revolutionize computational science. I think for these reasons we really want to invest in hybrid computing from an application developer point of view."

Work remains to determine if the payoff for other applications and algorithms is great enough to develop and port codes to hybrid, multicore systems. "What we're watching is a bunch of pioneers, people who have taken the first steps in embracing and adopting what amounts to an architectural sea change," Hack said. "Compared to the past when we made major architectural changes, this one is different because the vendors are part of the solution. The partnering that's going on in terms of allowing early adaptation or adoption of this new architecture includes working with the vendors, something I hope will continue. The kind of work that's going on now is identifying where the holes are and where the opportunities are for making progress."

NCSA Deputy Director Rob Pennington, a meeting co-chair, added that the partnership between domain and computational scientists is also crucial to providing whole solutions. "You're solving the problems now that are going to save your successors a lot of work," he said.

While some scientific application codes have been fully developed to run on GPUs, others need additional attention from computational scientists, computer scientists, software developers, and vendors combining their talents.


The success of the ACSS 2012 scientists helps decision makers for supercomputing centers weigh whether it makes more sense to fill empty slots with the fastest CPUs available or embrace a hybrid system employing both CPUs, which excel at serial processing, and GPUs, which excel at parallel processing.

Caffeinated codes, transformational results

The symposium highlighted advances in combustion, fusion energy, seismology, nuclear physics, biology, medicine, chemistry, and more made possible through GPU computing.

Combustion research. Combustion scientist Jacqueline Chen of Sandia National Laboratories, the first speaker, gave an overview of direct numerical simulations of turbulence-chemistry interactions. Combustion of fossil fuels accounts for 83 percent of U.S. energy use, with transportation alone accounting for two-thirds of petroleum use and one-quarter of carbon dioxide emissions. America aims to reduce petroleum usage 25 percent by 2020 and greenhouse gas emissions 80 percent by 2050. Meeting these goals will require a new generation of combustion systems that can achieve high efficiency and low emissions using diverse fuels—from conventional fuels to new fuels obtained from heavy hydrocarbons (oil from oil sands and oil shale as well as liquid fuels derived from coal) and renewable fuels (ethanol and biodiesel).

The combustion-engine community has identified combustion of dilute fuels at lower temperatures and higher pressures than achieved in today's engines as a promising way to increase fuel efficiency and decrease pollutant emissions. Chen detailed



Researchers use the S3D combustion code to investigate factors underlying the clean burning of fuels in next-generation engines. Shown are a lifted ethylene-air jet flame computed from direct numerical simulation (DNS) and tracer particle trajectories. The simulation explored 22 species in a reduced mechanism, used 1.3 billion grid points, and had a jet Reynolds number of 10,000, which indicates the simulation encompassed a great range of scales. Image credit: Sandia National Laboratories mechanical engineers Chun Sang Yoo and Jacqueline Chen performed the DNS. Hongfeng Yu of Sandia; Ray Grout of the National Renewable Energy Laboratory performed volume rendering

homogeneous charge compression ignition (HCCI) as a technology of great promise, but its control remains a challenge. “HCCI is really a chemical engine,” Chen said. Whereas the initiation of a spark or fuel-injection timing controls combustion in conventional engines, HCCI combustion starts when a piston compresses a charge to the point of ignition. Because combustion occurs by volumetric autoignition, thermal or chemical-species stratification is required to control how fast fuel burns. Too rapid a burning rate can lead to engine knock.

Hundreds of molecules with different chemical and physical characteristics have been proposed as alternative biofuels. How can engineers assess which are worth pursuing to optimize engine designs like HCCI? Chen's detailed combustion simulations on petascale supercomputers provide fundamental physical insights into the nuances of microscale turbulent mixing and chemical reactions in engine-relevant conditions. This knowledge is used to develop and assess predictive engineering models to facilitate an optimal combined engine-fuel system.


Supercomputers calculate the ability of turbulent mixing to optimize thermal and chemical stratification so that spontaneous combustion through ignition reactions occurs in proper phase with the piston motion. Chen is working with Ray Grout (National Renewable Energy Laboratory), John Levesque (Cray), Ramanan Sankaran (ORNL), and Cliff Woolley (NVIDIA) on revising the S3D scientific application code to make use of Titan's multicore hybrid architecture to enable simulations of greater chemical kinetic complexity. With Jaguar scientists could simulate hydrogen and small hydrocarbon fuels. Using Titan they aspire to simulate the larger hydrocarbons of practical transportation fuels such as heptane, octane, and biodiesel.

Grout followed Chen with a technical talk on efforts to “hybridize” the S3D legacy code. The message-passing interface (MPI), open message passing (OpenMP), and open accelerator (OpenACC) are community-developed, standards-based programming models for parallel computers that may be used separately or in concert

to enhance performance and portability between different and evolving computer architectures. “Movement of outer loops to the highest level in the code facilitates hybrid MPI-OpenMP usage that is necessary for performance on the machines coming on line now and allows an elegant path to GPU-enabled computing using OpenACC,” Grout said. “Setting up our algorithms this way, we get to keep a maintainable code yet still have acceptable performance across the range of machines we anticipate.”

The approach works: the hybridized version of the code performed twice as fast as the original code with each running on the now replaced Cray XT5 system. And running this new, improved code on today’s state-of-the-art technologies, the GPU-accelerated XK6 hardware outperforms the dual-CPU XE6 by a factor of 1.4. Grout said such performance levels should bring a target HCCI problem incorporating a realistic heptane or octane fuel and nearly 2 billion grid points into the realm of feasibility under the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program, the major means by which the scientific community gains access to the petascale computing resources of the OLCF.

Nuclear fusion energy. Bill Tang and Stephane Ethier of Princeton Plasma Physics Laboratory next spoke of extreme-scale computing as an increasingly vital tool for accelerating progress in developing nuclear fusion for clean, abundant energy. A fusion reactor magnetically confines deuterium and tritium fuel and releases energy when the two hydrogen isotopes fuse to form helium and a fast neutron—with a total energy multiplication of 450 times per reaction. The process depends on keeping the plasma hot, but turbulence causes heat losses. It is the balance between such losses and the fusion reaction heating rates that will ultimately determine the size and cost of a reactor. Understanding the associated complex mechanisms requires extreme-scale computing



The goal of ITER is to get as close as possible to achieving an ignited and burning, or self-sustaining, plasma in a tokamak that produces 500 megawatts of power for 400 seconds, getting 10 times more energy out than will be put in to ignite the plasma. Image credit: Jamison Daniel, ORNL

to simulate conditions in ITER, a multibillion-dollar experimental reactor in France with the collaboration of seven governments representing more than half the world’s population. Ultimately achieving commercial fusion energy depends on success at a dramatic next step—burning plasma in ITER to generate 500 MW of power for 400 seconds, getting ten times more energy out than will be put in. Current devices generate 10 MW and put out about as much energy as they consume.

The rapid advances in supercomputing power combined with the emergence of effective new algorithms and computational methodologies helps enable corresponding increases in the physics fidelity and the performance of the scientific codes used to model complex physical systems. “If properly validated against experimental measurements

and verified with mathematical tests and computational benchmarks, these codes can provide reliable predictive capability for the behavior of fusion-energy-relevant high-temperature plasmas,” Tang said.

“If properly validated against experimental measurements and verified with mathematical tests and computational benchmarks, these codes can provide reliable predictive capability for the behavior of fusion-energy-relevant high-temperature plasmas,” Tang said.

The fusion community has made excellent progress in developing advanced codes for which computer run time and problem size scale well with the number of processors on massively parallel supercomputers. For example, scaling is excellent for GTC, a global particle-in-cell kinetic code modeling the motion of charged particles and

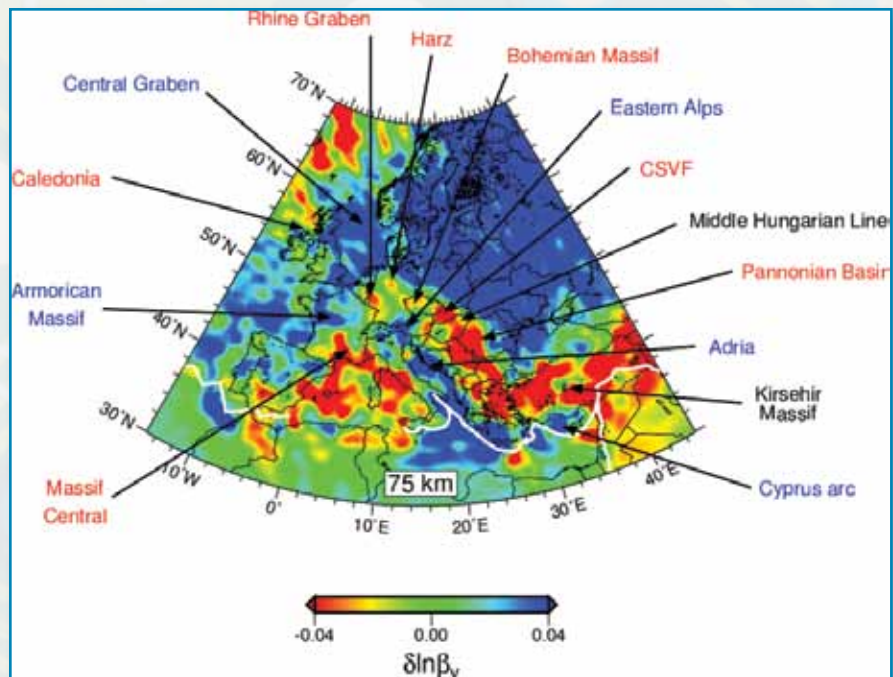
their interactions with the surrounding electromagnetic field. Several publications demonstrate that faster machines enable new physics insights with GTC, which was the first fusion code to deliver production simulations at petaflops in 2009. Showing data from China’s Tianhe-1A, currently the world’s second-fastest super-

computer, and America's TitanDev, Ethier said GTC on GPUs gives a 3-fold speedup over CPUs alone in some cases.

"We talked a lot about the GPU on the Cray, but don't forget there's also a brand new network and powerful CPUs in the XK6 platform," Ethier said of TitanDev. "Things that we couldn't do before, now we can do with the new network. So it's a full machine that we can optimize."

Seismology. Today petascale computing is used to create detailed "shake movies" such as M8, a 2010 simulation of a hypothetical magnitude-8 quake roiling Southern California that employed nearly all of Jaguar's approximately 220,000 processing cores and sustained 300 teraflops to run the anelastic wave propagation code AWP-ODC. With GPUs taking computing closer to exaflop performance levels—calculation speeds a thousand times faster than petaflops—Jeroen Tromp of Princeton University and Olaf Schenk of the University of Lugano said large-scale seismic inverse modeling and imaging are possible that will allow researchers to work backward after a big shake to determine in three dimensions where it originated and what phenomena shaped it. While earthquake prediction is not possible, inverse modeling of virtual quakes may help engineers assess potential hazards and build structures to withstand them.

Much as medical tomography uses waves to image sections of the human body, seismic imaging uses waves generated by earthquakes to reveal the structure of the solid Earth. Numerically solving the equations that govern motion of diverse geological solids challenges even the most advanced high-performance computing systems. Tromp used inversion to analyze 190 European quakes and revealed seismic hotspots in the upper mantle. The imaging detailed the subduction of Africa, volcanism in the Czech Republic, a "hole" under Bulgaria, and Italy's counterclockwise rotation over the past 6 million years.




Princeton's Jeroen Tromp used seismic inverse modeling and imaging to analyze 190 European quakes to reveal seismic hotspots (red indicates fast seismic shear wave speeds whereas blue indicates slow wave speeds). Imaging of the upper mantle revealed volcanism in the Czech Republic, a "hole" under Bulgaria, and Italy's counterclockwise rotation over the past 6 million years. An allocation of 739 million core hours on Titan would facilitate a global inversion using data from 5,000 earthquakes, allowing unprecedented understanding of forces that agitate Earth. Image courtesy of Jeroen Tromp

"Our ultimate goal is to move toward adjoint tomography of the entire planet," Tromp said. Getting started on this ambitious goal means analysis of a minimum of 250 earthquakes worldwide, requiring 14 million CPU core hours. But completing the task will require analysis of 5,000 earthquakes worldwide, using an estimated 739 million core hours. This is a huge amount of computing for today's systems—equivalent to running on Jaguar nonstop for 6 months. However, with GPU performance acceleration and Titan on the way, Tromp is working to make such an ambitious campaign a reality.

"Our ultimate goal is to move toward adjoint tomography of the entire planet,"
Tromp said.

"These seismic inverse problems are known as PDE [partial differential equation]-constrained optimization problems and are significantly more

difficult to solve than the PDE-forward problems," Schenk said. With the spectral element wave propagation code SPECFEM, one global-scale simulation of the forward problem might take 10,000 processor hours, but the inverse problem would take 300 million processor hours, he said. Optimization on tens of thousands of processor cores offers challenges, but solving them is worth the effort in terms of what the computational power doing so would provide the Earth science community. With GPUs on Tödi and TitanDev, Schenk showed much greater performance for the spectral-element wave propagation solver SPECFEM over that of CPU-only machines—approximately four times faster than XK6 with no accelerator (which has one CPU per node) and two times faster than Monte Rosa XE6 (which has two CPUs per node). TitanDev and Tödi have NVIDIA's Fermi GPUs. But Titan will have NVIDIA's even faster next-generation Kepler GPUs.



Researchers use the Denovo application code to model and simulate radiation transport in nuclear reactors for nuclear forensics, radiation detection, and analysis of fuel cores and radiation shielding. Without GPUs, using the 3.3-petaflop Jaguar system, Denovo can simulate a fuel rod through one cycle of use in a reactor core employing 60 hours of wall clock time. With GPU acceleration in the 20-petaflop Titan system, this simulation would take only 13 hours. Image credit: ORNL

Nuclear fission energy. Speaking next were ORNL's Doug Kothe and Tom Evans, who use modeling and simulation to improve the safety and performance of nuclear power plants, which supply 20 percent of America's electricity today and are projected to supply 25 percent by 2030. Kothe and Evans are partners in the multi-institutional Consortium for Advanced Simulation of Light Water Reactors (CASL), DOE's first Energy Innovation Hub, established in 2010 and headquartered at ORNL. CASL researchers model and simulate reactors to address three critical performance areas for nuclear power plants: reducing capital and operating costs per unit of energy by enabling power uprates and lifetime extension for existing plants and by increasing the rated powers and lifetimes of next-generation plants, reducing nuclear waste by enabling higher fuel

burnup, and enhancing nuclear safety by predicting component performance through the onset of failure.

Evans leads development of the Denovo code, which charts radiation transport in a reactor. "To match the fidelity of existing 2D industry calculations with consistent 3D codes, we need to solve a minimum of 10,000 times more unknowns than we have achieved to this point," he said. Denovo's most challenging calculation, an algorithm that sweeps through the virtual reactor to track the location of particles, was consuming 80 to 95 percent of the code's run time on ORNL's 3.3-petaflop Jaguar, which uses only CPUs. Using hybrid TitanDev to pass this calculation off to GPUs, researchers saw Denovo's speed increase 3.5 fold. "Getting a factor of 3 to 4 from GPUs means that a 30- to 40 petaflop machine could

allow fully consistent, 3D neutronics simulations that could be used to address CASL challenge problems," Evans said.

Fundamental physics. What is the nature of the nuclear force that binds protons and neutrons into stable nuclei and rare isotopes? What is the origin of simple patterns in complex nuclei? What is the origin of elements in the cosmos? These are a few of the unanswered physics questions under investigation by ORNL's David Dean, who reported on the present status and future prospects for computing nuclei, which comprise 99.9 percent of baryonic matter in the universe and are the fuel that burns in stars. A comprehensive description of nuclei and their reactions requires extensive computational capabilities. Dean described applications of nuclear coupled-cluster techniques for descriptions of isotopic chains (recent work focused on oxygen and calcium). Today scientists use 150,000 cores to compute one nucleus. Recent work by Dean collaborators Hai Ah Nam and Gaute Hagen indicated that GPUs may triple the speed of the calculation. Dean also detailed applications of nuclear density functional theory, which today use all of Jaguar to simulate 20,000 nuclei in 12 million configurations in an effort to determine the limits of existence of nuclei. "In both cases we are benefiting from the application of petascale computing and will benefit from hybrid architectures in the future," Dean said.

Scientists theorize that the four fundamental forces of nature—electromagnetism, gravity, the strong force (which holds an atom's nucleus together), and the weak force (which is responsible for radioactive decay)—are related. Balint Joo of Jefferson Lab described investigations in the field of quantum chromodynamics, which aims to understand the strong interaction in terms of quarks and gluons. He used the TitanDev testbed to show strong scaling to 768 GPUs (more than three-quarters of TitanDev). He and NVIDIA's Mike Clark compared

the performance of GPU-based solvers from the QUDA library with their CPU-based counterparts and found a 3- to 10-fold speed boost, depending on problem size. The calculation had to solve 679 million unknowns and did so at a speed of 100 teraflops. With Titan, work to better understand the strong force may enable scientists to compute the properties of never-before-seen exotic mesons prior to their planned production in GlueX, a flagship nuclear physics experiment that will run around 2015 in Jefferson Lab's accelerator after an upgrade to double the energy of its electron beam from 6 billion to 12 billion electron volts.

Climate science. Hack, who heads the Oak Ridge Climate Change Science Institute, introduced a discussion of the challenges in accelerating coupled models such as the state-of-the-art Community Earth System Model used by more than 300 climate scientists. This mega-model couples component models of atmosphere, land, ocean, and ice to reflect their complex interactions. For example, Phil Jones of Los Alamos National Laboratory and a large team of collaborators couple high-resolution


global ocean models to high-resolution atmosphere models (25-kilometer-wide grid cells for atmosphere and 10-kilometer resolution for ocean) to more accurately simulate the climate system and its extremes by including ocean eddies, tropical storms, and other fine-scale processes. "The current state is 2 to 3 simulated years per CPU day," he said. "A speedup of 3 to 10 times permits real science using limited ensembles and longer time integrations."

"Since atmospheric cores used for numerical weather prediction and climate prediction can be very similar, this result is strong evidence that the desired 3- to 10-fold speedup can be achieved," Henderson said.

Matt Norman from ORNL spoke of porting an important atmospheric code to hybrid platforms. "By porting to the Cray XK6 [1 Fermi GPU and 1 Interlagos CPU per node], we achieved a 1.5-fold speedup over Cray

XE6 nodes for the Community Atmosphere Model-Spectral Element code using 900 nodes on each machine," he said. "This currently ported portion of the code, or subroutine, accounts for 80 percent of the run time for our target science problems, and most of the remaining run time will be ported to GPUs soon." Similarly, Tom Henderson from NOAA shared his group's achievement of a 5-fold speedup of an atmospheric dynamical core used for numerical weather prediction (NWP). "Since atmospheric cores used for NWP and climate prediction can be very similar, this result is strong evidence that the desired 3- to 10-fold speedup can be achieved for climate codes," Henderson said.

Math libraries. To address the complex challenges of hybrid environments, optimized software solutions will have to fully exploit both CPUs and GPUs. The OLCF's Jack Wells introduced a panel discussion on math libraries aimed at accomplishing this feat. Jack Dongarra of the University of Tennessee-Knoxville talked about the Matrix Algebra on GPU and Multicore Architectures, or MAGMA, project to develop a dense linear algebra library for hybrid architectures. ORNL's Chris Baker next discussed Trilinos efforts to develop parallel solver algorithms and libraries within an object-oriented software framework for large-scale multiphysics applications. Finally, Ujval Kapasi of NVIDIA gave an overview of CUDA, NVIDIA's parallel computing platform and programming model that can send C, C++, and Fortran code straight to GPUs with no assembly language required. The sequential part of a workload still runs on the CPU, but parallel processes are sent to the GPU to accelerate the computation. CUDA has found use in applications from simulating air traffic to visualizing molecular motions. The panel made apparent the vital role of early partnerships between technology developers and scientific-application users in optimizing libraries that will eventually empower the entire computational science community.



Accelerated computing may benefit simulations of ocean eddies, tropical storms, and other fine-scale processes. Ocean simulations that include ecosystems are important for determining how much carbon oceans can sequester in current and future climates. Shown in green is chlorophyll, a measure of biological activity that can be compared with observational data obtained with satellites. Image credit: Matthew Maltrud, Los Alamos National Laboratory

Medicine and biology. The second day's lectures kicked off with Jeremy Smith of the University of Tennessee–Knoxville and ORNL speaking on supercomputing in biology and medicine. He described large-scale docking calculations used in drug discovery, a process in which thousands of compounds are screened to identify hundreds of candidates that are then winnowed to dozens of effective agents. A few go on to preclinical development, the stage at which most drugs fail, and fewer still make it to become the subjects of expensive and time-consuming clinical trials. Supercomputers can speed analysis and lower costs. High-throughput screening of drug libraries can quickly identify promising candidates and generate copious initial data about how drugs bind to their protein targets, decreasing the likelihood that the drugs will later prove toxic due to undesired binding, or cross-reactivity. “If drug candidates are going to fail, you want them to fail fast, fail cheap,” Smith said. In one day a petaflop computer can analyze 10 million drugs binding to a target, Smith said, whereas a 10-petaflop machine can analyze a million drugs interacting with 100 targets. A future exaflop machine, in comparison, could in one day assess the interactions of 10 million drugs with 1,000 different targets—one representative from every class of proteins in the human body.

“If drug candidates are going to fail, you want them to fail fast, fail cheap,” Smith said.

Smith also highlighted molecular dynamics simulations of cellulosic ethanol systems under investigation at DOE's Bioenergy Science Center. The goal is to better understand two steps necessary to generate “gasoline” from cellulose, a complex carbohydrate made of sugar units that forms cell walls in stalks, trunks, stems, and leaves. Hydrolysis breaks cellulose into glucose and enzymes, and fermentation turns the glucose

into alcohol. Petascale simulations revealed new knowledge about hydrolysis, an expensive step that must be optimized to make cellulosic ethanol commercially viable. Smith shared a capability-class, 24-million-atom simulation that detailed enzymatic binding to lignocellulose that had been pretreated with heat and dilute acid. With a more powerful 10-petaflop supercomputer, he said, a 300-million-atom simulation could shed light on how specialized microbes break down biomass. An even more powerful exaflop supercomputer could, in principle, simulate the 10 billion atoms in a living cell, although considerable technical barriers to such an achievement still remain.

“Over the last decade molecular dynamics simulation has evolved from a severely limited esoteric method into a cornerstone of many fields, in particular structural biology, where it is now just as established as NMR [nuclear magnetic resonance] or X-ray crystallography,” said Erik Lindahl of Stockholm University, who next discussed scaling the GROMACS code for execution on hybrid, accelerated architectures. “A central challenge for the entire community is that we frequently achieve scaling by going to increasingly larger model systems, reaching millions or even hundreds of millions of particles. While this has resulted in some impressive showcases, the problem for structural biology is that most molecules of biological interest (for example, membrane proteins) only require maybe 250,000 atoms.” The biggest challenge for the future is rather that scientists would like to reach timescales many orders of magnitude larger than what is possible today. He discussed using Copernicus, an open framework that parallelizes aspects of distributed computing, to run molecular dynamics on high-performance computers with fast interconnects. Copernicus treats each problem like an ensemble in which free energy is calculated for thousands of molecular shapes and the computation converges on the conformation

with the lowest energy. The result is larger-scale simulations than would be possible with distributed computing projects, such as Folding@Home, and reduced time to solution. In 46 hours a supercomputer can determine the 3D structure of a small protein from scratch, given its 2D amino-acid sequence; GPU-accelerated nodes will reduce this time by a factor of 3 to 4, Lindahl said.

Jim Phillips of NCSA at University of Illinois at Urbana-Champaign shared his work on scalable molecular dynamics using the NAMD code, which has 51,000 users worldwide to attest to its wide, deep impact. Projects driving the need for scalable molecular dynamics include investigations of ribosome dysfunctions leading to neurodegenerative disease and infection processes of the viruses that lead to polio and AIDS. His talk detailed how supercomputers probe the behavior of biosensors, whole cells, blood coagulants, integrin [a receptor mediating attachment between a cell and surrounding tissues], and membrane transporters and included a status update on scaling and acceleration of NAMD and Charm++, an adaptive, parallel run-time system. Acceleration with GPUs on Japan's Tsubame supercomputer allowed a large simulation of a 20-million-atom photosynthetic membrane patch. Cray's Gemini interconnect doubled the usable nodes for strong scaling (i.e., problems of a fixed size, in which increasing the number of processors shortens the time to solution) in a more moderately sized simulation of 92,000-atom apolipoprotein A1 on a Blue Waters prototype. On a 100-million-atom benchmark molecular dynamics calculation run on 768 nodes, GPU-accelerated XK6 hardware sped the calculation above and beyond what was possible with CPUs—a 2.6-fold acceleration on the XK6 without acceleration and a 1.4-fold boost on the dual-CPU XE6. These results show it would be more advantageous to fill a second slot of a compute node with a GPU than another CPU.

Chemical and materials

sciences. Christopher Mundy of Pacific Northwest National Laboratory discussed chemistry at interfaces, which has far-reaching implications for fuel cells, biological pumps, chemical catalysts, and chemical reactions in the atmosphere and on surfaces. He and his team performed research using first-principles interaction potentials in conjunction with statistical mechanics methods on leadership-class computers through the INCITE program. He performed the investigations under the auspices of two DOE programs that aim to elucidate complex chemical phenomena in heterogeneous environments: the Basic Energy Sciences' Condensed Phase and Interfacial Science program, which improves the molecular understanding of processes at interfaces, and the Energy Frontiers Research Center on Molecular Electrocatalysis, which supports studies of the conversion of electrical energy into chemical bonds and vice versa. His team used the CP2K computer code—nicknamed “the Swiss army knife of molecular simulation” because of its wide range of tools—to explore the role of electrons in explaining the behavior of molecules and materials. Parallel computations of free-energy profiles of individual chemical reactions used ensemble methods to statistically explore important configurations that give rise to novel phenomena. Mundy described a recent controversy—whether the surface of liquid water is acidic or basic. His work, using high-quality interaction potentials and based in quantum mechanics, suggests that the hydroxide anion is only slightly stabilized at the interface and that no strong propensity exists for the surface to be either acidic or basic.

How Effective Are GPUs on Scalable Applications? Code Performance Measurements on TitanDev

Application	XK6 (with GPU) vs XE6 Performance Ratio TitanDev : Monte Rosa
S3D Turbulent combustion	1.4
Denovo 3D neutron transport for nuclear reactors	3.3
LAMMPS High-performance molecular dynamics	3.2
WL-LSMS Statistical mechanics of magnetic materials	1.6
CAM-SE Community atmosphere model	1.5
NAMD High-performance molecular dynamics	1.4
Chroma High-energy nuclear physics	6.1
QMCPACK Electronic structure of materials	3.0
SPECFEM-3D Seismology	2.5
GTC Plasma physics for fusion energy	1.6
CP2K Chemical physics	1.5

Cray XK6: Fermi GPU plus Interlagos CPU | Cray XE6: Dual Interlagos and No GPU

As water is necessary for life on Earth, a deeper understanding of it could have a profound impact.

Joost VandeVondele of ETH Zurich spoke of nanoscale simulations using large-scale and hybrid computing with CP2K, which aims to describe at an atomic level the chemical and electronic processes in complex systems. The work may aid understanding of novel materials, precisely tuned

chemical reactions, and nanoscale processes that are crucial to addressing challenges in energy, environment, and health. “The next-generation hardware will accelerate this research,” he said. “However, this will require that domain scientists, in collaboration with system experts, design novel and powerful algorithms that are carefully implemented to exploit the strengths of the new hardware.” He presented recent efforts by the CP2K team to

simulate systems of ten thousand to millions of atoms and increase the accuracy for condensed-phase system simulations. He discussed early results using accelerators in a massively parallel context, showing GPU-accelerated XK6 performance more than doubled compared to CPUs alone in the XE6.

David Ceperley of the University of Illinois at Urbana-Champaign, who highlighted simulations of condensed-matter systems, began his talk with a Paul Dirac quote from 1929: “The general theory of quantum mechanics is now almost complete. The underlying physical laws necessary for the mathematical theory of a large part of physics and the whole of chemistry are thus completely known, and the difficulty is only that the exact application of these laws leads to equations much too complicated to be soluble.” Today’s supercomputers make these equations solvable, Ceperley said. “Recently, we have developed methods to couple the simulation of the quantum electron degrees of freedom of a microscopic system with the ionic degrees of freedom without evoking the approximations involved in density functional theory or semi-empirical potentials,” he said. “This allows an unprecedented description of all condensed-matter systems.” He spoke of simulations providing new insights into dense hydrogen, important for understanding the giant planets, and water, the key to many biophysical processes. He also discussed supercomputers speeding the design of materials with specific properties. Computation for each compound was a doctoral thesis 15 years ago, he said. Now, supercomputers can search 10,000 compounds to find the optimal material. “There’s plenty of work at the petascale and exascale,” Ceperley said.

Jeongnim Kim of ORNL is the principal author of QMCPACK, a code widely used by physical scientists for large-scale simulations of molecules, solids, and nanostructures. The code uses a technique called Quantum

Monte Carlo (QMC) that has enabled accurate, many-body predictions of electronic structures. “The resources at DOE leadership computing facilities have allowed QMC to seek fundamental understanding of materials properties, from atoms to liquids, at unprecedented accuracy and time-to-solution,” she said. “We have developed QMCPACK to enable QMC studies of the electronic structure of realistic systems, while fully taking advantage of such theoretical development and ever-growing computing capabilities and capacity.” She reviewed progress in QMC algorithms and their implementations in QMCPACK, including accelerations using GPUs. In recent work she simulated 256 electrons in a 64-atom piece of graphite. GPUs on TitanDev sped the calculation about 3.8-fold compared to the CPU-only portion of the Jaguar XK6 and about 3.0-fold over Monte Rosa’s dual-CPU XE6 nodes. Kim said a hybrid supercomputer of 10 to 100 petaflops could support the Materials Genome Initiative, which aims to accelerate understanding of the fundamentals of materials science, providing practical information that entrepreneurs and innovators can use to develop new products and processes.

To 20 petaflops and beyond

Arthur Bland, director of Jaguar’s upgrade and transformation project to create Titan, gave the meeting attendees an overview of that project, which by the end of 2012 will deliver an XK6 system with 600 terabytes of memory, Cray’s new high-performance Gemini network, and 18,688 compute nodes each containing one 16-core AMD Opteron CPU. While TitanDev now has Fermi GPUs, in fall 2012 these will be removed to allow most of the compute nodes in TitanDev and Jaguar to be upgraded with NVIDIA’s Kepler GPUs. Notably, the project plan is for 4,096 XK6 compute nodes to remain without accelerators. Titan users will have access to the Spider file system with more than 10 petabytes of storage capacity and an initial data bandwidth of 240 gigabytes



Hybrid architecture is the foundation of ORNL’s 200-cabinet Titan supercomputer, which will be a groundbreaking new tool for scientists to leverage the massive power of hybrid supercomputing for transformational research and discovery. Image credit: Andy Sproles, ORNL

per second and a planned upgrade of up to 1 terabyte per second. Bland announced that during the week of March 26, the DOE Office of Science approved the build-out of Titan to at least 20 petaflops.

Pennington followed with an update on Blue Waters, which will have more than 25,000 compute nodes. All XE6 nodes are expected to be in place by mid-year, and all XK components should be in place by the end of the year. Blue Waters will have more than 235 XE6 cabinets and more than 30 XK6 cabinets equipped with NVIDIA’s Kepler GPUs. It will have 1.5 petabytes of memory, and its disk storage will exceed 25 petabytes, with near-line storage that can go to 500 petabytes.

The major means of accessing Titan will be through the INCITE program, which awarded 1.7 billion core hours in 2012 and will allocate nearly 5 billion core hours in 2013 on the 10-petaflop IBM BG/Q Mira supercomputer at Argonne National Laboratory and Titan at ORNL. The Blue Waters system is accessed through the NSF’s Petascale Computing Resource Allocations process. —by Dawn Levy

Contact information:

Jack Wells

Oak Ridge Leadership Computing Facility
Oak Ridge National Laboratory
Phone: 865-241-2853
Email: wellsjc@ornl.gov
www.olcf.ornl.gov

Thomas Schulthess

Swiss National Supercomputing Center (CSCS)
Swiss Federal Institute of Technology in Zurich (ETH Zurich)
Phone: +41 91 610 82 01
Email: schulthess@cscs.ch
<http://www.cscs.ch>

Rob Pennington

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign
Phone: 217-244-1052
Email: robp@ncsa.illinois.edu
<http://www.ncsa.illinois.edu/>

