NERSC NEWS

NATIONAL ENERGY RESEARCH SCIENTIFIC COMPUTING CENTER

MARCH 2009

Green Flash Prototype Has a Successful First Run

Berkeley Lab's Green Flash project, which is exploring the feasibility of building a new class of energy-efficient supercomputers for climate modeling, has successfully reached its first milestone by running the atmospheric model of a full climate code on a logical prototype of a Green Flash processor.

"The logical prototype simulates the entire circuit design of the proposed processor," says John Shalf of the NERSC Division, principal investigator of Green Flash.

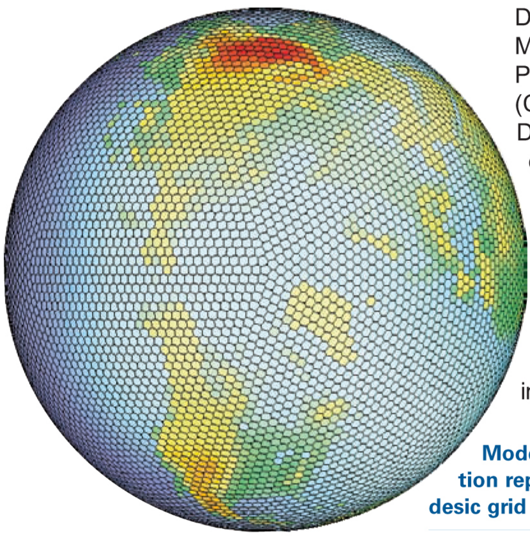

Modeling with a Geodesic Grid: Surface elevation represented on a 10,242-cell spherical geodesic grid in a GCRM model.

The prototype was designed in collaboration with Tensilica, Inc., using Tensilica's Xtensa LX extensible processor core as the basic building block, and was run on a RAMP (Research Accelerator for Multiple Processors) BEE3 hardware emulator, which is used for computer architecture research.

The climate code that ran on the prototype was a next-generation, limited area model version of the geodesic Global Cloud-Resolving Model (GCRM) developed by the research group of David Randall, Director of the Center for Multiscale Modeling of Atmospheric Processes at Colorado State University (CSU) and principal investigator of the DOE SciDAC project "Design and Testing of a Global Cloud-Resolving Model." Randall, Ross Heikes, and Celal Konor of CSU are collaborating with the Green Flash team.

David Donofrio of Berkeley Lab's Computational Research Division (CRD), who works on the hardware design of Green Flash, ran the prototype in a demonstration at the SC08 conference in Austin, Texas in November 2008.

Green Flash was first proposed publicly in the paper "Towards Ultra-High Resolution Models of Climate and Weather," written by Michael Wehner and Lenny Oliker of CRD and Shalf of NERSC. The paper was published in the May 2008 issue of the International Journal of High Performance Computing Applications.

Addressing Three Problems

The Green Flash project addresses three research problems simultaneously — a climate science problem, a computer architecture/hardware problem, and a software problem.

The climate science problem stems from current low resolution climate models that cannot capture the behavior of cumulus convective cloud systems. Instead, researchers must use statistical patterns to approximate this important mechanism of heat and moisture transport. Direct numerical simulation of individual clouds systems has been proposed as a solution to replace these approximations with a more rigorous formulation but would require horizontal grid resolutions approaching 1 km. Randall's research group is gradually working toward that goal; their current SciDAC

continued on page 3

What's New with HPSS?

HPSS Upgrades Simplifies User Passwords

Researchers storing valuable scientific data on NERSC's High Performance Storage System (HPSS) can now instantly access or manage their HPSS account from an internet browser anywhere in the world, 24 hours a day, seven days a week, and 365 days a year.

This kind of on-demand access to massive archival datasets is unprecedented, and is only now possible because of a series of custom software upgrades to HPSS at NERSC, which makes the system compatible with NERSC's Information Management (NIM) system.

"If you think of HPSS as a safety deposit box that stores valuable scientific data, NIM is like an online banking feature which allows you to set up your account and track your account activity, HPSS being one of your accounts," says Jason Hick, head of NERSC's Mass Storage Group.

He notes that the new software upgrade replaces an ad hoc authentication system where users were required to apply for multiple passwords to use NERSC systems — a NIM password to track their activity on supercomputers like Franklin and Bassi, and a separate password to retrieve archival data stored on HPSS. With the upgrade, HPSS is now integrated with NIM to enable users to manage their HPSS accounts and passwords along

continued on page 4

Kathy Yelick Appointed to the California Council on Science and Technology

NERSC Division Director Kathy Yelick has been appointed as a council member to the California Council on Science and Technology (CCST), a nonpartisan, nonprofit organization that offers expert advice on science and technology related issues to state government.

"California has a tremendous science and technology brain trust in its universities, research laboratories, and industry, and I look forward to working with the Council on how the state can expand that resource and use it to address some of the challenges and opportunities facing

continued on page 4

User Services Group Hires Two Consultants

Alice Koniges


Alice Koniges

Alice Koniges may be the newest high performance computing consultant in the User Services Group, but she is no stranger to NERSC. Soon after she became the first woman ever to earn a PhD in Applied and Computational Mathematics at Princeton University, Koniges began her career as a member of NERSC's Computational Physics Group in 1984. Back then, NERSC was known as the Magnetic Fusion Energy Computing Center and was based at the Lawrence Livermore National Laboratory, parallel computing was the frontier, and the first Cray-2 machine, complete with four processors, had just been installed at the center. A relative newbie, Koniges achieved the first successful parallel code run on the machine.

"Four processors were a really big deal back then, and one of the biggest challenges was getting a parallel code to demonstrate overlap when it was open to other users.… I essentially got the machine to show it was capable of computing in parallel by running my own physics code and letting the memory get bigger and bigger until most of the other projects running on the machine were kicked off, which allowed me to grab all the processors and demonstrate overlap. This meant the machine worked as we hoped it would," says Koniges.

She began her career researching parallel computing and computational plasma physics, eventually achieving luminary status in both fields. Her expertise in the transition from vector to parallel computing culminated in her textbook *Industrial Strength Parallel Computing*, published by Morgan Kaufmann in January 2000. Nowadays, her research interests include cloud computing, benchmarking, the multi-core revolution and programming languages associated with multi-core computing, as well as her mainstay of application supercomputing.

Originally from Poway, California, a small town outside of San Diego, Koniges always harbored an interest in science. As a child, she bonded with her father in family amateur astronomy classes. At the age of 15, her high school science teacher noted her interest and drove her several days a week after school to work with a research astronomy professor at San Diego State University. This work resulted in her attending the International Science Fair, where she won top prize with a project on binary star systems. A paper resulting from this project was published in the Astrophysical Journal.

While other teens took on paper routes or retail sales positions, Koniges' International Science Fair fame landed her a summer job at the Naval Research Laboratory, where she was introduced to computer programming. She eventually made her way through undergraduate studies at the University of California, San Diego, and graduate and doctoral studies at Princeton University, teaching programming to others along the way. A natural educator, Koniges has taught technical seminars on a variety of topics at every Supercomputing Conference since 1998.

She currently resides in Livermore with her youngest daughter and husband, a physicist at LLNL. In her free time, Koniges enjoys horseback riding, playing piccolo in local orchestras, and biking. She is currently training for the Cinderella Classic, a 100 km all-women's bicycle tour of the San Francisco area's East Bay.

Woo-Sun Yang

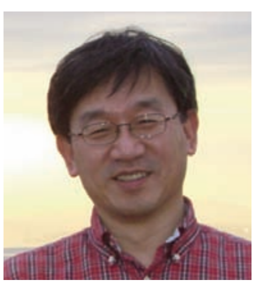
Woo-Sun Yang

As a new scientific consultant with NERSC's User Services Group, Woo-Sun Yang will combine an interest in numerical modeling with a knack for enhancing computer productivity to further the frontiers of science.

Yang admits that his road to NERSC was not entirely straightforward. As a child in South Korea, Yang dreamed of becoming a scientist and even received Bachelor and Master of Science degrees in geology at the Seoul National University. However, upon arriving at the University of Illinois, Urbana-Champaign, he decided to change the focus of his PhD thesis to geodynamics after reading an inspiring paper describing how simple mathematical models could help researchers better understand the fluid dynamics in the Earth's mantle.

"The elegant approach in solving the proposed fluid dynamic problems by using a simple mathematical model was very fascinating. I felt as though that was the field I had been preparing to study by then," says Yang.

Yang's first experience with supercomputers stemmed from a curiosity about whether the fluid dynamics of the Earth's mantle affected the plate tectonics, movement of plates on the planet's crust. He notes that the only way to study the inner workings of the mantle is with supercomputer models. From then on, he was hooked.

"Many people are naturally introduced to and step into computational science through creating models for scientific research," says Yang. "At this point, I was very much into numerical modeling and I decided to switch to computational science."

Shortly after receiving his PhD, Yang began a postdoctoral position at CRD, working on a parallel IO library for a climate computer model. This transitioned nicely into his second career post as an onsite engineer for Cray Korea Inc., based at the Korea Meteorological Administration. Here he helped scientists and technical staff forecast weather by porting, debugging, and parallelizing weather codes on the 1042-PE Cray X1E vector supercomputer. He also taught classes on using Cray supercomputers and the Fortran language. He moved back to the United States in 2007 as a Cray onsite engineer at NERSC, where he advised on Franklin-related efficiency issues.

"The work that I did at Cray is similar to what I do as a NERSC user consultant, but the current task requires interactions with more diverse and wider audience. I am excited that now I have research opportunities to work with great staff members," says Yang.

He currently lives in Moraga with his family and a sizeable collection of LP records, which his wife has collected over the years. In his spare time he enjoys listening to classical and Korean folk music on old LPs and traveling across the American Southwest.

Green Flash Project Runs First Prototype Successfully

continued from page 1

project aims to develop a cloud model with 3 km resolution.

To develop a 1 km cloud model, scientists would need a supercomputer that is 1,000 times more powerful than what is available today. But building a supercomputer that powerful with conventional microprocessors (the kind used to build personal computers) would cost about $1 billion and would require 200 megawatts of electricity to operate — enough energy to power a city of 100,000 residents. That constitutes the computer architecture problem. In fact, the energy consumption of conventional computers is now recognized as a major problem not just for climate science, but for all large-scale computing.

Shalf, Wehner, and Oliker see a possible solution to these challenges — achieving high performance with a limited power budget and with economic viability — in the low-power embedded microprocessors found in cell phones, iPods, and other electronic devices. Unlike the general-purpose processors found in personal computers and most supercomputers, where versatility comes at a high cost in power consumption and heat generation, embedded processors are designed to perform only what is required for specific applications, so their power needs are much lower. The embedded processor market also offers a robust set of design tools and a well-established economic model for developing application-specific integrated circuits (ASICs) that achieve power efficiency by tailoring the design to the requirements of the application. Chuck McParland of CRD has been examining issues of manufacturability and cost projections for the Green Flash design to demonstrate the cost-effectiveness of this approach.

Meeting the performance target for the climate model using this technology approach will require on the order of 20 million processors. Conventional approaches to programming are unable to scale to such massive concurrency. The software problem addressed by the Green Flash project involves developing new programming models that are designed with million-way concurrency in mind, and exploiting auto-tuning technology to automate the optimization of the software design to operate efficiently on such a massively parallel system.

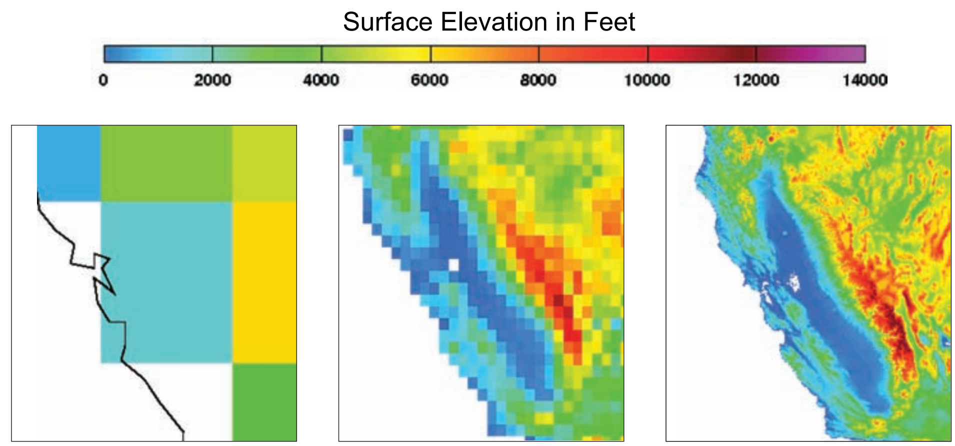

Increasing Model Resolution: Topography of California and Nevada at three different model resolutions. The left panel shows the relatively low resolution typical of the models used for the International Panel on Climate Change's Fourth Assessment Report, published in 2006. The center panel shows the upper limit of current climate models with statistical approximations of cloud systems. The right panel shows the resolution needed for direct numerical simulation of individual cloud systems.

To meet this challenge, Tony Drummond of CRD and Norm Miller of the Earth Sciences Division are working on analyzing the code requirements; and Shoaib Kamil, a graduate student in computer science at the University of California, Berkeley who is working at NERSC, has been developing an auto-tuning framework for the climate code. This framework automatically extracts sections of the Fortran source code of the climate model and optimizes them for Green Flash and a variety of other architectures, including multicore processors and graphics processors.

Developing Hardware and Software Together

An innovative aspect of the Green Flash research is the hardware/software co-design process, in which early versions of both the processor design and the application code are developed and tested simultaneously. The RAMP emulation platform allows scientists to run the climate code on different hardware configurations and evaluate those designs while they are still on the drawing board. Members of the RAMP consortium on the UC Berkeley campus, including John Wawrzyneck and Krste Asanovic (both of whom have joint appointments at NERSC), Greg Gibling, and Dan Burke, have been working closely with David Donofrio of CRD and the Green Flash hardware team throughout the development process.

At the same time, auto-tuning tools for code generation test different software implementations on each hardware configuration to increase performance, scalability, and power efficiency. Marghoob Mohiyuddin, another UCB graduate student at NERSC, has been working on automating the hardware/software co-design process and has recently demonstrated a 30–50% improvement in power efficiency over conventional approaches to design space exploration. The result will be a combination of hardware and software optimized to solve the cloud modeling problem.

The researchers estimate that the proposed Green Flash supercomputer, using about 20 million embedded microprocessors, would deliver the 1 km cloud model results and cost perhaps $75 million to construct (a more precise figure is one of the project goals). This computer would consume less than 4 megawatts of power and achieve a peak performance of 200 petaflops.

The hardware goals for the prototype research are fairly simple: produce a hardware prototype of a single Green Flash processor by fall 2009, and an entire node of the system (64 to 128 processors) by fall 2010. For more information on the project, go to http://www.lbl.gov/CS/html/greenflash.html.

What's New with HPSS?

continued from page 1

with every other system at NERSC.

"The goal is to streamline the process and make it easier for users," says Hick. "Before this upgrade, many of the calls that came through the user support center were related to obtaining a HPSS password and problems with authentication. Since we implemented the new software there has been a significant drop in user support calls."

In addition to making this system compatible with NIM, members of the Mass Storage Group also added a real-time monitoring capability that allows the center's administrative staff to more easily detect and solve internal HPSS problems. Besides the new features, the HPSS v6.2 system that is now in production has already proven to require less regular maintenance, which will be a big benefit to users of HPSS at NERSC.

Planned HPSS Hardware Upgrade Expands Scientific Archive

HPSS is software that efficiently manages petabytes of data (1,000,000 gigabytes) by moving information from high-cost, high-speed disks to low-cost tapes, and protecting that data over time. NERSC currently runs two different HPSS systems: an archival HPSS system that stores user files, provides high transfer rates with lower response times, among other things; and a backup HPSS system that retains supercomputer and NERSC global file system backups and provides for restoration of user data in the event of a disaster.

Since its inception 11 years ago, the amount of data stored on NERSC's HPSS has doubled in size every two years. To keep up with the demand for archive space, the center will replace its current 20 and 200 gigabyte tapes with 75 and 1,000 gigabyte tapes respectively. In addition, the center will increase the number of recordable cartridges in its library from 44,000 to 59,000.

"Users most likely will not notice the upgrade or immediately recognize the increased storage capabilities because user transfers predominantly involve HPSS disk resources. However, the benefit to science is enormous in that it enables us to keep pace with increased storage demand," says Hick. "In January 2009, HPSS experienced two days of record data transfers. The upgraded cartridges will allow us to keep up with the 50% daily demand increase we have seen since Franklin went into production."

The oldest science file saved in HPSS dates back to January 2, 1976. As of January 27, 2009, NERSC's archive HPSS contained over 66 million files, adding up to over 3.9 petabytes of data. Meanwhile, the backup HPSS holds over 12 million files, which add up to 2.8 petabytes of data.

Cost Efficient Archive: HPSS at NERSC archives petabytes of data on systems consisting of spinning disks and tapes.

Kathy Yelick Appointed to Council

continued from page 1

the state," says Yelick.

The CCST Council is an assembly of corporate CEOs, academicians, scientists, and scholars. The organization brings together those who create knowledge with those who create wealth, in cooperation with those who make policy, to utilize science and technology for the economic and social well being for California. Currently, seventy-eight members and fellows are also members of the National Academies, five are Nobel laureates, eight are National Medal of Science recipients and two are recipients of the National Medal of Technology.

Kathy Yelick

Yelick is a professor of computer science at the University of California, Berkeley, in addition to heading NERSC.

NERSC News, produced as a pdf since February 2004, is evolving into Computing Sciences News, a monthly newsletter highlighting the latest news from NERSC, ESnet and CRD in html format. For more information contact editor Linda Vu at LVu@lbl.gov or 510.495.2402

Workshop to Improve Next Generation HDF Held at Oakland Science Facility

Scientists and application developers from academia, the Department of Energy (DOE) and industry all met at NERSC from January 20–31 to develop strategies for usability improvements to version 5 of the Hierarchical Data Format (HDF5) software suite.

HDF5 is the primary scientific data format and digital library designed to handle terabyte-sized datasets and parallel file processing. The software allows users to efficiently store and access large files, as well as save a diverse set of files in one container. It can also grow to accommodate new types of objects and their metadata.

“The workshop primarily looked at how HDF might be made faster and more scalable for scientists,” says David Skinner, who heads NERSC’s Software Integration Group.

Scientific users at DOE and NASA represent most of the user base for HDF5, which was developed by engineers at the National Center for Supercomputing Applications (NCSA) at the University of Illinois Urbana-Champaign.

For more information on the workshop’s discussions, please visit the SciDAC Outreach Center: http://outreach.scidac.gov/hdf/

NERSC Employees Inspire Future Computer Scientists on Job Shadow Day

NERSC employees Katie Antypas and David Paul (at right) introduce four Albany High School students to the intricacies of parallel computing with parts from the center’s decommissioned Seaborg supercomputer at a job shadow tour of the Berkeley Lab’s Oakland Science Facility on March 4,2009.

Paul holds a Seaborg 16 GB memory module in his left hand, and points to a Seaborg four-processor CPU module with his right hand. At the end of the day, the students even went away with Seaborg souvenirs — a high-speed switch module, frame ID number magnet and IBM logo.

“When I walked into the gargantuan room the held the supercomputers, I expected to be gazing upon a nightmare of cables, engineers, commotion, and disgruntled looking scientists. Fortunately, that was not the case.… Everything was very orderly and awe-inspiring,” says Kevin Gumina, a junior classman at Albany High School. “Overall I enjoyed the day and hope one day I can work in this field.”

“At the NERSC job shadow day, I learned that those who work alongside supercomputers are actually quite genial, friendly and sociable — I had previously imagined that it would be some uptight group of nerds,” says Christian Pedersen, also an Albany High School junior. “After my trip to NERSC, I want to be a computer scientist more than ever.”

“The students were great, really interested in our topics, amazed by our scale and they asked intelligent engaging questions,” says Paul. “I truly enjoyed participating with them.”

From left, students are Kevin Gumina, Patrick Leng, Cong Chen and Christian Pederson. Overall 11 juniors from Albany High School shadowed lab staff on February 25 and March 4, as part of the school’s annual Job Shadow Day. Students were mentored by employees in ALS, AFRD, Computing Sciences, NERSC, CSO, Earth Sciences and Health Services.

Seaborg was NERSC’s IBM SP RS/6000 system, which was retired in 2008 after seven years of service.