

YOUNG TALENTS

Computing Sciences will host summer students in research and administrative support projects

Four graduate students who are members of the DOE Computational Science Graduate Fellowship program have joined Computing Sciences this summer. As part of their participation in the fellowship program, the students spend a summer conducting research at a DOE national lab. During the past seven years, 27 CSGF fellows have selected LBNL as the site for their summer practicum. These fellows are some of the brightest minds in fields that include chemistry, engineering, computer sciences, mathematics and life sciences.

Meet the four fellows:



Julianne Chung

Julianne Chung hails from Emory University and studies computational mathematics. She received a bachelor's degree in mathematics from Emory in 2004. Chung, who also minored in dance and movement

studies, is working with Chao Yang in CRD's Scientific Computing Group. Chung's Ph.D. research focuses on developing numerical methods and software for obtaining more clear and detailed images from biomedical imaging devices. To achieve higher-resolution images requires algorithms that can solve large systems of equations. Chung said she is particularly interested in developing iterative methods that can be used to solve large systems arising in ill-posed inverse problems.

During her summer at LBNL, she will be implementing a massively parallel data distribution scheme for the cryo-EM reconstruction project. Due to memory limitations, current parallel reconstruction algorithms cannot handle large-volume structures such as viruses and ribosomes sampled with a small pixel size. With Chung's work, structural biologists and researchers will be able to reconstruct 3D macromolecules and other large structures like the adenovirus.

Peter Norgaard, who studies computational plasma dynamics at Princeton University, is working with Phil Colella in

CRD's Applied Numerical Algorithms Group. Norgaard earned a bachelor's degree in aerospace engineering from the University of Washington in 2004. His Ph.D. research involves theoretical and numerical investigations of fast ion redistribution during periods of large-scale magnetohydrodynamic activities inside a tokamak, a magnetic confinement device and leading reactor design concept for producing fusion energy. Norgaard models the fast ion dynamics by numerically integrating the guiding center equations for representative test particles as they orbit the tokamak in the presence of oscillatory perturbation modes, and then extracts dynamical properties of the system from the numerical data.

With Colella, Norgaard is working on an extension of the Chombo adaptive mesh refinement software package to particle-in-cell methods. This will include developing a suite of test problems to verify that the interaction between the particles and the grid-based fields is self-consistent, particularly near the refinement boundaries.

Michael Sekora studies continuum mechanics, partial differential equations and numerical analysis at Princeton University. Sekora, who received bachelor's degrees in mathematics and physics from MIT in 2005, also is working with Phil Colella while at Berkeley Lab. In his Ph.D. research,



Michael Sekora

Sekora aims to develop the mathematical analysis of three-dimensional numerical methods for solving radiation hydrodynamic problems and answer questions about star formation, supernovae and accretion disks around black holes. He's particularly interested in extending Godunov, Piecewise-Parabolic and adaptive mesh refinement methods to resolve the problems.

This summer, Sekora said he wants to



Peter Norgaard

develop a numerical method for solving three-dimensional gas dynamical problems that has fourth-order spatial accuracy. It is conceivable that such a numerical method can be constructed by combining a technique for computing fourth-order face quadratures and the fourth-order spatial differencing of the unsplit piecewise parabolic method.

Benjamin Smith comes to CRD from Harvard University, where he studies experimental high energy physics. Smith, who received a bachelor's degree in physics and mathematics from Harvard in 2005, is working with Ernest Szeto in the Biological Data Management and Technology Center (BDMTC) within CRD. Smith's Ph.D. research involves hunting for the elusive Higgs Boson particle, whose existence can prove a current theory on how fundamental particles such as electrons and quarks acquire mass. The theory, based on the Standard Model, contends that these fundamental particles acquire mass while interacting with a scalar field called the Higgs field.



Benjamin Smith

Recent advances in biological data acquisition, such as rapid sequencing of microbial genomes, has resulted in a dramatic increase in the amount of data that must be stored, distributed and analyzed by biologists. The Integrated Microbial Genome (IMG) system, developed by the BDMTC, uses cutting-edge techniques in computer science to allow direct searches and comparisons of microbial genome sequences. Smith will be working on developing new techniques for efficiently indexing new genome sequences using the high performance computing resources at LBNL.

A list of the current fellows and more information about their academic background can be found at http://www.krellinst.org/csgf/fellows/listing/fellows_alph.cgi.

In addition to the CSGF fellows, Dong
continued on page 5

Quantum Leap *continued from page 1*

research ideas.

Speaking to a packed room, Rose described the basic architecture of Orion and how the quantum processor, built with superconducting niobium and chilled to near minus 273-degree Celsius in liquid helium, can carry out calculations. As a 16-qubit analog computer, Orion currently lacks the computational power of a conventional desktop computer.

Rose said when the system launches commercially next May, Orion will be 100 times faster than the best systems used in industries such as banking, pharmaceutical and defense—a problem that would take a conventional computer 12 hours to solve would take Orion just seven minutes. The business plan calls for leasing time on Orion through a web-based service, and Rose said tasks sent by customers could use common programming languages such as SQL.


Many scientists in the audience were less interested in D-Wave's business plan, but cared more about whether Orion can perform and scale as promised. During the Q&A session and private meetings with company officials, CRD and NERSC researchers probed for more details and debated whether Orion meets the definition of quantum computing, such as achieving quantum mechanical phenomena such as entanglement. After the meetings, the scientists said many issues remained unanswered.

"A big question is if quantum coherence is actually achieved across all 16 qubits, and if it could be achieved across thousands of qubits for long enough to sustain a quantum computation," said Erich Strohmaier in CRD's Future Technologies Group.

Orion is an adiabatic quantum computer developed specifically to provide more precise answers to NP-complete problems, which stipulate that all possible solutions must be checked out before arriving at the best one. Such specific design makes it attractive to businesses in areas such as engineering, finance and logistics.

But it also limits its ability to compete with more conventional supercomputers in many other tasks. At least for now.

"From the quantum simulation/material science perspective, much larger machines are necessary for any potential



Source: D-Wave Systems

The optical image shows the processor that D-Wave used during its quantum computing demo in February at the Computer History Museum in the Silicon Valley. The chip contains 16 qubits (the quasi-circular loops arranged in a 4x4 array). Each of the qubits is coupled to its nearest neighbors (N, S, E, W) and next-nearest neighbors (NW, NE, SW, SE) via a tunable flux transformer, giving a total of 42 of these couplers.

applications," said Joshua Schrier in CRD's Scientific Computing Group. "The current device (with only the Ising-type couplings) would not be suitable for more general quantum model Hamiltonian simulations."

For the commercial market, D-Wave sees its system as a good fit for nanotechnology research, such as drug discovery, given a quantum computer's potential to do a better job than conventional supercomputers in simulating the behavior of objects at the molecular level.

Orion's aim at solving tricky NP-complete problems also has generated interesting research questions for mathematicians in CRD. Esmond Ng, leader of the Scientific Computing Group, plans to exchange ideas with D-Wave engineers on the possibility of tackling combinatorial problems in sparse matrix computing using quantum computers.

Ali Pinar, another scientist in the group, would like to collaborate with D-

Wave on the design of decomposition methods, which would solve problems more quickly by grouping variables into sets. With the addition of a front-end server and the use of decomposition methods, Orion's capabilities can be broadened to tackle a broader set of problems, Pinar said.

"The success of quantum computing, and Orion in particular, depends not only on hardware, but also on algorithms that will best utilize the underlying infrastructure," Pinar said. "Orion is fundamentally different than anything we used before, so shall the algorithms be."

While many questions about quantum computing remains open, Berkeley Lab scientists say they will closely follow D-Wave's progress in advancing the field.

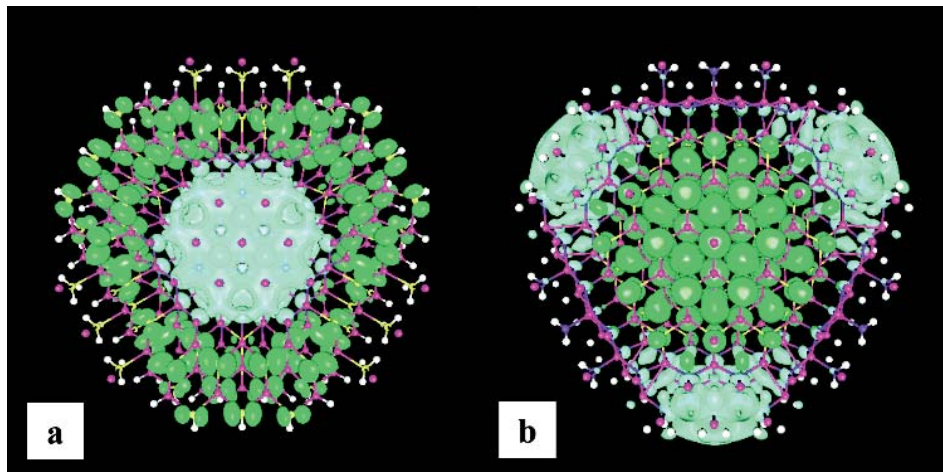
A blog by Rose that describes Orion can be found at <http://dwave.wordpress.com/2007/01/19/quantum-computing-demo-announcement>.

Nanowire *continued from page 1*

The 4-nanometer wire, modeled using a NERSC supercomputer, proved to be good at keeping electrons away from the holes and minimize energy loss. Unlike a conventional coaxial cable, which has an insulator to separate the central copper core from the braided copper conductor in the outer shell, the nanowire designed by Wang and his colleagues doesn't require an insulating material to keep the charges separate between the core and the shell.

Using gallium nitride and gallium phosphide together also creates a structure with a smaller band gap than they would otherwise have separately, leading to better solar cell efficiency. Having a smaller energy gap between the conduction and valence bands means the device can utilize a wider spectrum of the sunlight.

However, since the output voltage also depends on the conduction and valence band gap, there is an optimal band gap for solar cell application. The core-shell structure allows the scientists to manipulate the band gap by changing the thickness ratio between the core and the shell.



Cross-section view of charge distributions of the conduction band minimum's electron state (cyan) and valence band maximum's hole state (green) in a nanowire. (a) A nanowire with a gallium nitride core and gallium phosphide shell. (b) A nanowire with a gallium phosphide core and gallium nitride shell.

The scientists also alternated the use of gallium nitride and gallium phosphide for the core and the shell and found both models can produce similar band gaps, around 1.2 eV, which is similar to the band gap of silicon, and ideal for solar cell application.

Wang's work was published in the April 5 online edition of *Nano Letters* by

the American Chemical Society.

The research is far from finished.

"We'll look for different materials and different geometries, and how we can manipulate them," Wang said.

Wang is involved in various nanomaterial research projects. Check out his work at <http://hpcrd.lbl.gov/~linwang>.

Hall of Fame *continued from page 1*

researcher in CRD's Scientific Computing Group, and her project co-leader from France, plan to use the money to further their work on developing novel algorithms for solving large sparse linear systems, used in many time-consuming computations in fields such as fluid dynamics, structural analysis and nanomaterial designs.

Both the Berkeley and the French teams include members who have developed some of the most widely used sparse solvers in the last decade. Li's parallel SuperLU solver, for example, was the second-most downloaded code from the Berkeley Lab site in fiscal 2006. The solver has

been a key in carrying out simulations for the design of International Linear Accelerator, the study of magnetohydrodynamics in fusion energy and the understanding of quantum scattering theory.

Current solvers are capable of tackling large sparse linear systems with more than one million equations. But the growing demand for high-resolution simulations will require new algorithms to solve systems with hundreds of millions of equations.

Iain Duff, project leader of Parallel Algorithms Group at CERFACS in Toulouse, is heading the French team. The 1-year project is scheduled to begin in January next year.

Defining Future Research

The third and final town hall meeting to brainstorm for novel computational science challenges ended at Argonne National Laboratory recently, with CRD scientists playing key roles in moderating discussions and reporting their findings.

The three town hall meetings, held first at Berkeley Lab and then at Oak Ridge National Laboratory last month, are part of an effort by ASCR to generate ideas for an innovative program for tackling new computational sciences and expand ASCR's contributions to

continued on page 5

Summer Students *continued from page 2*



Dong Ryeol Lee

Ryeol Lee, a Ph.D. candidate from Georgia Tech University, has joined CRD as a fellow for the Department of Homeland Security Scholars and Fellows Program. The program pays for each fellow's tuition and other fees

and an annual stipend. Each appointment can last up to three years.

Lee graduated from Carnegie Mellon University with dual degrees in computer science and mathematics in 2005. He then enrolled in the Ph.D. program in computer science at Georgia Institute of Technology. His Ph.D. research focuses on developing scalable machine learning and data mining algorithms for enabling the usage of models that are computationally intensive but more expressive.

During this summer, he is working with Deb Agarwal, head of the Distributed Systems Department in CRD, on a KDD workflow project. He will be working on algorithms to resolve conflicts, implement control structure and merge search results within the KDD workflow system.

Computing Sciences also will host other high school and college students through various programs. The high school students are Dimitry Kisliyuk, Aaron Marshall and Michael Marshall.

The college students are Elena Caraba (Louisiana State University); Jonathan Chu (UC Berkeley); Oliver Dario (UC Davis); Christopher Creighton (Contra Costa Community College); Daniel Schreiber (University of Illinois at Urbana-Champaign); Louise Kutten (UC Berkeley); Andrew Rodriguez (UC San Diego); Sahak Margossian (UC Berkeley); Karen McVey (Ohlone College); Matt Doiron (Big Bend Community College); Matthew Smith (Big Bend Community College); Zachariah Tanko (Big Bend Community College); Annabelle Carter (Diablo Valley College); Cobber Lam (California State University at East Bay) and Pietro Cicotti (UC San Diego).

In addition, four international students from ETH Zurich will participate: Mauro Calderara, Jonathan Coles, Fracois Gognat and Denis Nordmann.

Hall of Fame *continued from page 4*

the DOE Office of Science mission over the next decade. The program is currently called Exsacle for Energy and Environment (E3).

During the first meeting at Berkeley Lab, scientists broke into nine groups to discuss subjects such as renewable energy, distributed computing, astrophysics and biology. Ideas from those discussions were further refined in the two subsequent meetings. Many researchers attended all three meetings.

More than a dozen CRD researchers are involved in drafting the final reports. When the town hall meeting took place at Berkeley Lab, Lab Director Steve Chu presented a lunch talk about LBNL's increasing emphasis on energy efficiency and new, renewable energy sources. Chu underscored his case for increased investment by citing both economic and environmental factors, which show that current energy practices are too costly to sustain on a global scale.

SciDAC Show and Tell

CRD and NERSC scientists spoke and presented posters this month at SciDAC 2007, the conference that brought together recipients of the Scientific Discovery through Advanced Computing Program from DOE's Office of Science. The program funds and brings together computational scientists, applied mathematicians and computer scientists from U.S. universities and national labs to work on algorithms and software development.

Phil Colella, John Shalf and Horst Simon were featured speakers during the five-day conference in Boston. Colella, head of the Applied Numerical Algorithms Group in CRD, talked about "Performance and

Scaling of Locally Structured Grid Methods for Partial Differential Equations." Shalf, the head of the Science-Driven System Architecture Team at NERSC, presented "The New Landscape of Parallel Computing Architecture." Simon, head of CRD and NERSC, presented a review of the three town-hall meetings on develop Exsacle for Energy and Environment (E3), a program for tackling novel computational sciences in the next decade.

Several CRD scientists also presented posters. They were John Bell, Andrew Canning, Lin-Wang Wang, Xiaoye Sherry Li, Parry Husbands, Brian Tierney and Arie Shoshani.

The conference took place in Boston this year. Berkeley Lab will host SciDAC 2008 in Seattle. More information about this year's SciDAC conference can be found at <http://www.scidac.gov/Conference2007>.

Large Applications, Distributing Computing

Keith Jackson and Dan Gunter were invited speakers at the conference, Challenges of Large



Keith Jackson

Applications in Distributed Environments (GLADE), which took place this month in Monterey Bay, California.

Jackson, head of the Advanced Application Interface Technologies Group in CRD, spoke about the remote analysis of large-scale data. Gunter, in the Collaborative Computing Technologies Group,




Dan Gunter

continued on page 6

Hall of Fame *continued from page 5*

spoke about troubleshooting data movement in his presentation.

Both scientists also took part in a panel discussion about data challenges for running large applications in distributed environments. More information about the conference can be found at <http://www-unix.mcs.anl.gov/%7Eschopf/CLADE2007/program.html>.



A New Book Out

David Bailey, CRD's chief technologist, is the co-author of "Experimental Mathematics in Action," a newly published 322-page book which grew out of a two-day course

Bailey helped lead at the 2006 annual joint meeting of the American Mathematical Society and the Mathematical Association of America.

The eight lectures presented during the course are analogous to the eight chapters of the book, which is co-authored by Jonathan M. Borwein, Neil J. Calkin, Roland Girgensohn, D. Russell Luke and Victor H. Moll. A draft version of the book had actually been written before the course, with the contents then revised and polished in light of the authors' experiences at the meeting, which drew about 80 participants.

The goal of this book is to present a variety of accessible examples of modern mathematics where intelligent computing plays a significant role. The book highlights some of the key algorithms and teaches key experimental approaches. A.K. Peters Ltd. of Massachusetts published the book.

In 2004, Bailey and Borwein co-authored "Mathematics by Experiment: Plausible Reasoning in the 21st

Century," and with Girgensohn they wrote "Experimentation in Mathematics: Computational Paths to Discovery."

Multicore Computing



Kathy Yelick

Kathy Yelick, head of the Future Technologies Group, was one of seven speakers at the conference "Multicore —the New Face of Computing: Promises and Challenges," that took place at Stanford University this month.

The conference, co-sponsored by the Santa Clara Chapter of the IEEE Computer Society and the North America Taiwanese Engineers' Association (NATEA), is part the "New Frontiers in Computing Technology" series sponsored by the two groups and featured speakers from both academia and industry.

Yelick presented a talk called "The Berkeley View: Applications-Driven Research in Parallel Program-ming Models and Architectures."

The conference provided a technical overview of the hardware and software issues involved in making effective use of multicore technology, including compilation, scalability graphics co-processing and potential applications in the new domain. For more information, see http://www.natea.org/sv/conferences/nfic/2007/nfic_2007.php.

Career Tips

Juan Meza, head of the CRD High Performance Computing Research Department, shared his insights into getting a research position with nearly



Juan Meza

250 students in a talk presented by the Rice-Houston AGEP (Alliance for Graduate Education and the Professoriate) this month.

The Rice-Houston AGEP unites the city's universities and community colleges in the mission to increase the number of minority students earning the Ph.D. and position them to become leaders in science. Meza earned his Ph.D. and master's degrees in computational mathematics from Rice University.

In his talk titled "Research Job Interviews: The Unwritten Rules," Meza combined practical advice and humor, such as reminding students to "READ you're CV double check spelling and grammar." Another tip: avoid hot-button issues such as politics, religion and American Idol during dinner conversations.

Meza's presentation is posted at <http://hpcrd.lbl.gov/~meza/presentations.html>.

About CRD Report

CRD Report, which publishes every other month, highlights the cutting-edge research conducted by staff scientists in areas including turbulent combustion, nano materials, climate change, distributed computing, high-speed networks, astrophysics, biological data management and visualization. CRD Report Editor Uclia Wang can be reached at 510 495-2402 or Uwang@lbl.gov. Find previous CRD Report articles at <http://crd.lbl.gov/html/news/CRDreport.html>.

DISCLAIMER

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California. Ernest Orlando Lawrence Berkeley National Laboratory is an equal opportunity employer.