SCIENCE, Computation, and Collaboration at BERKELEY LAB

R. KALTSCHMIDT, LBNL

LBNL Director Dr. Steve Chu discusses Berkeley Lab, high-performance computing, and collaborations for solving important science problems in the future.

SciDAC Review: *Several years ago, you left Stanford to accept the position of Lab Director at Lawrence Berkeley National Laboratory (LBNL, or Berkeley Lab). What led you to that decision and how do you like heading up the "father" of DOE's national laboratories?*

Dr. Chu: It's been quite an experience, and a wonderful opportunity to be part of what I consider to be great national assets. We have a very diverse set of talents which can be brought to bear on many different things. Our strengths are in many, many areas, from our roots in physics and chemistry to biology, computation, engineering, life sciences, and earth sciences. These represent a unique opportunity—with many scientific challenges in the future, you want "bench strength" in many areas. For example, with the Helios project, we will contribute in the areas of biofuels, nanotechnology for solar energy, and artificial photosynthesis/catalysis, among others. I think we can make similar contributions in studying climate change.

A number of our national user facilities—the National Energy Research Scientific Computing (NERSC) Center, the Advanced Light Source, the Molecular Foundry, and a part of the Joint Genome Institute—are truly national assets that are used by teams of scientists from around the country.

I think that one of the things that differentiate this lab from the other labs is that we have much stronger ties to a great research university, the University of California (UC)–Berkeley, and this has been a tradition since Lawrence founded his lab. The core of this has been historical ties to physics and the College of Chemistry. We also have great opportunities with Electrical Engineering and Computer Sciences, Mechanical Engineering, and Earth Sciences, but we haven't taken advantage of those relationships as much as I would like. There are a number of distinguished faculty members, and building stronger ties with them is one of my goals for the laboratory.

This is an exciting time for Berkeley Lab, and while we are building more connections, we need to keep our focus on being a national lab. Ideally, we will be able to leverage the intellectual strength not only of UC–Berkeley, but with the entire San Francisco Bay Area, which is a founding place for many, many things. The computer revolution and the information technology revolution, for example, grew out of Silicon Valley. Biotech has its roots here. And I think we will find our region at the heart of the Green Tech Revolution, with Berkeley Lab poised to be in the center of it all.

If you look at the industries in the Bay Area, you'll see that many of them grew out of the research institutions, with universities feeding repeatedly into the defining companies of these technology revolutions. There is an entrepreneurial, risk-taking spirit here, and I think that is accompanied by higher self-expectations and higher performance among many of the people who live and work here.

For decades after its founding in 1931, Berkeley Lab was known as a physics lab. In the 1970s, that began to change with the growth of energy efficiency research and scientific computing. In the 1990s, LBNL was catapulted to a leadership position in computational science when DOE moved both NERSC and the Energy Sciences Network (ESnet) to Berkeley. Ten years on, how would you measure the results of this experiment?

The Lab's strengths grew out of high-energy physics and nuclear chemistry, but I'd like to backtrack a little bit. We have real strength in chemistry, which began with Glenn Seaborg and Edwin McMillan in the early days of LBNL. In fact, of our 11 Nobel laureates, four of them—Seaborg, McMillan, Melvin Calvin, and Yuan Lee—were recognized for their work in chemistry. They started their work as laboratory-scale experiments. And now we have new generations of scientists. Then there's biochemistry and astrophysics, including the research led by George Smoot and Saul Perlmutter. And of course computational science is growing.

Biography in Brief: Dr. Steve Chu

Dr. Steve Chu became the sixth director of Lawrence Berkeley National Laboratory (LBNL) on August 1, 2004. His appointment marked a return to Berkeley for the physicist, who began his career in laboratory research as a postdoctoral fellow in physics at the University of California's Berkeley campus from 1976–1978, during which time he also utilized the facilities of Berkeley Lab. He is currently also a Professor of Physics and Professor of Molecular and Cellular Biology at the University of California–Berkeley.

Dr. Chu's first career appointment was as a member of the technical staff at AT&T Bell Laboratories from 1978–1987. He spent many years there as the Head of the Quantum Electronics Department, during which time he began his groundbreaking work in cooling and trapping atoms by using laser light.

In 1987, he became a Professor in the Physics and Applied Physics Departments at Stanford University. His work eventually led to the **Nobel Prize in Physics** in 1997, an honor he shared with Claude Cohen-Tannoudji of France and United States colleague William D. Phillips. Their discoveries, focusing on the so-called "optical tweezers" laser trap, were instrumental in the study of fundamental phenomena and in measuring important physical quantities with unprecedented precision.

Dr. Chu was the Theodore and Francis Geballe Professor of Physics and Applied Physics at Stanford University, where he remained for 17 years as highly decorated scientist, teacher, and administrator. At Stanford, he helped start Bio–X, a multidisciplinary initiative linking the physical and biological sciences with engineering and medicine. He has become active in the energy problem and is co-chairing an international InterAcademy Council (IAC) study, **"Transitioning to Sustainable Energy."** The IAC represents over 90 national academics of science around the world.

He has held numerous visiting lectureships that include Harvard University, the JILA Institute, Collège de France, the University of Oxford, and the University of Cambridge. He is a member of the National Academy of Sciences (NAS), the American Philosophical Society, the American Academy of Arts and Sciences, the Academia Sinica, and is a foreign member of the Chinese Academy of Sciences and of the Korean Academy of Science and Engineering.

He serves on the Boards of the Hewlett Foundation, the University of Rochester, and NVIDIA, and on the scientific boards of the Moore Foundation, NABsys and Helicos. He has served on a number of committees, including the Augustine Committee that produced the report **"Rising Above the Gathering Storm"** in 2006, advisory committees to the Director of the National Institutes of Health and the National Nuclear Security Agency, and the Executive Committee of the NAS Board on Physics and Astronomy.

Born in St. Louis and raised in New York, Dr. Chu earned an A.B. in mathematics and a B.S. in physics from the University of Rochester, a Ph.D. in physics from the University of California–Berkeley, and eight honorary degrees. He maintains a vigorous research program and directly supervises a team of graduate students and postdoctoral fellows. He is author or co-author of more than 200 articles and professional papers, and over two dozen former members of his group are now professors at leading research universities around the world.

NERSC has been a huge boost for the Laboratory. Here again, we have an amazing opportunity going into the future, especially when you consider the ecosystem we are embedded in. Both UC and Stanford have outstanding computer science programs and we are in the center of industry-leading companies in the areas of computing and technology. We should be taking full advantage of this ecosystem.

I think one of the things we're most proud of with NERSC is that we don't just run a machine that provides flops for users, but we are helping those users solve scientific problems. In the future we certainly want to continue to provide this service, but also, we'd like to take more advantage of our ecosystem. In the future, if you look at the trajectories of computational science, you could think of a Moore's Law of computational capability on the horizon; as supercomputers incorporate many more processors, there is an issue of having a wide range of scientific users who can take advantage of this capability. In this regard, you might think that NERSC could be located anywhere, as long as you have high-speed connections, that being tied to the Bay Area doesn't matter. But somehow, it does seem to matter, although I don't fully understand why, given that the Office of Science computers are accessible by the broader intellectual community.

I know that NERSC has helped raise awareness of the role of computational science at Berkeley Lab and among our collaborators. As first practiced, science was purely theoretical, or what was then called "philosophy." And then experiment emerged as a scientific tool, and it was called "natural philosophy." By toggling back and forth between theory and experiment, we created what we now call modern science, about 400–500 years ago.

There are certain things you can compute; in the early days we limited this to problems with analytical solutions. But with some problems we're interested in today, it's hard to get tractable analytic solutions. Simulation can provide that third leg where we don't have an opportunity to conduct experiments. A good example is climate research—we're doing the experiment right now. But with simulations, we can get some assurance of what the outcomes might be for different scenarios.

> As supercomputers incorporate many more processors, there is an issue of having a wide range of scientific users who can take advantage of this capability.

In areas where we can do experiments, computation lets you look at a level where our instrumentation doesn't allow us to look, or at the quantum level, which we could not measure directly, or in a molecular dynamics experiment, where we still don't have the ability to look at atomistic details, but simulation and modeling can help us see what is happening.

There continues to be a Moore's Law of computational capability, but as large systems incorporate more and more processors, there is an issue as to whether a wide range of scientific users will be able to take advantage of these technologies. But even at the cluster level, we have a great opportunity to think about new architectures. The capability for providing semi-automatic parallelization has to be brought into the hardware design process with the software engineers, not with the hardware designed first and handed off to the software guys. One of the goals I would like to see NERSC achieve is how users can get this parallelization of their codes without being one of the few selected users for petascale systems. We need to make these systems available to more of the average scientific users to solve their problems. NERSC and members of our Computational Research Division (CRD) have opportunities to partner with our local ecosystem and design how the next generation of computing will be done.

Since you arrived at LBNL in 2004, you have led—and continue to lead—a number of high-profile research initiatives that continue to change the Lab's scientific landscape. These include the Helios partnership with UC–Berkeley, BP, and the University of Illinois, and the Joint BioEnergy Institute. Which disciplines are involved and how do you see them coming together?

As I said, we have strengths in many, many areas. In the area of biofuels, biology, chemical engineering, and microbiology will be used. Another part of Helios is harnessing solar energy to create fuel through artificial photosynthesis, but that's in the longer term. Within 10–15 years, we want to have plants, such as algae, convert solar energy to fuel. The goal is to develop better plants for converting this energy, not to use conventional food crops. Developing other plants is more sustainable in terms of land use, higher yields, reducing the net carbon dioxide, and greatly reducing the fossil fuel inputs. It's much better environmentally.

Artificial photosynthesis marries bioscience with inorganic sciences—catalysis and chemistry. There is great crossing over there. With artificial photosynthesis, if we get it right, we will be able to use less water to produce a higher fraction of fuel. I've seen some estimates that biofuels could yield up to 10% of the world's energy budget, with the limiting factor being water. Biofuels could provide up to two-thirds of the energy used in the U.S.

With energy efficiency, I see crossing over between Materials Science and other divisions, such as the technology side of our Environmental Energy Technologies Division (EETD). And I would like to see the policy side of EETD draw on the scientific and technical resources—including computing—of the rest of the Lab as well.

Our intent is not to turn the Lab into an applied technology center. Instead, we want to take the strengths we have and form outstanding teams and address every problem we can, such as climate prediction, the flow of water. We certainly don't want to throttle off basic sciences—those are our basic core strengths, and they lead to strengths in other areas.

R. KALTSCHMIDT, LBNL

Figure 1. *Director of LBNL, Dr. Steve Chu.*

I see lots of opportunities by crossing over boundaries between both scientific disciplines and organization charts. That's what makes this Lab so special, and we should take full advantage of these strengths.

Continuing on the energy theme, one of the major topics of discussion in the field of computing these days is improving energy efficiency both at the systems level and for data centers. You have been promoting energy efficiency in many areas. What is the special importance of lowering power consumption for high-performance computers? What role should DOE and LBNL play here?

I don't know if most people realize it, but if you leave your PC on most of the time at home, it is the largest single largest consumer of electricity in your house. If you have a high-end desktop at home, it's heating up your house. If you look at the huge server stations for Google and Yahoo, they are consuming over 100 megawatts, 24/7. This is not a sustainable path in terms of energy consumption.

It all starts with the chip and goes up to the packaging, the installation on racks, and so on. For years, we have been using commercial chips in systems, and then trying to improve efficiency by building more efficient buildings to house the computer centers. But if we start with the design of the chips, we can create a lot of better options. If you look at current chips, especially multi-core, where each processor is the same, there is not an optimal use of real estate. The manufacturers know it and so do we. Each side is beginning to focus on the issue. There's no reason why the processors have to be the same on a multi-core chip.

One of the arguments against improving efficiency from the chip up is the current availability of relatively cheap power in some areas. Sitting your hungry power users near cheap power sources might appear to be good business sense in a micro-sense. But we need to look at the total amount of power we use in the United States. As we begin to develop a truly national, very high-voltage and highly efficient power grid, this will become less important. This grid will make it easier to store and transport electricity, which will also boost renewable energy sources.

Energy conservation also has national security implications. As currently configured, the locations of the country's largest high-performance computing (HPC) centers are easily detectable by the power lines feeding into them and the amount of heat they generate.

As a major user of HPC, DOE is well-positioned to help lead the move to more efficient computing, but it starts with the chip. A group from LBNL and UC–Berkeley is currently pursuing this approach by working with a manufacturer of the small, highly efficient processors used in consumer electronics and a computer manufacturer to prototype a large-scale system that is both less expensive to build and less expensive to power.

Another major undertaking is the construction of the new Computational Research and Theory facility, which will bring together NERSC, computer science, and applied math programs, and provide infrastructure for increased collaboration with UC–Berkeley. What do you see as the return on this major investment?

I see greater connections between Computing Sciences at Berkeley Lab and Electrical Engineering and Computer Sciences (EECS) on campus, connections which will benefit research communities beyond our Lab and campus. Historically, we have not had as close of a connection as I would like. Many times, a unit of research consists of a professor and his group of students. With some projects, such as tools for large-scale simulations, the work may require a longer-term concerted effort, one that outlives a group of graduate students and exceeds the incremental improvements that often result.

> Petascale systems will give us a lot of detailed information from which we can make choices.

The focal point of the new center will be to help scientists throughout the country by developing tools to help them, and specifically to help research groups where the primary members are not computing experts. They want to solve problems, but the group may not have the staying power to develop the tools themselves. The goal is to help researchers take full advantage of scientific computing. The best way to do that is to understand the problems they are trying to solve.

Ideally, we would have more types like John Bell and Phil Colella, who seek a full understanding of the problem they are tackling. Then they don't just develop a one-off solution, but tools that can be handed off and used by others too. This gives them a bigger reach, far beyond what they would be capable of as a single group studying one problem.

I would love to see NERSC and CRD develop methods to help thousands, tens of thousands of researchers use HPC.

The computing community developed through town hall meetings in the last few months a prospectus for exaflops computing, in particular for applications in energy and the environment. What are, in your opinion, the biggest scientific opportunities for computing in the age of petaflops and beyond?

Detailed climate prediction in regional areas comes to mind first. While climate averages are important, what will move policymakers are detailed simulations of what will happen in their own areas. Which areas will become deserts? Where will vegetation become weakened by environmental changes and become susceptible to destruction by parasites? Who will face water shortages? Climate change is a global problem, but what will drive us to make changes is what will happen locally.

Then there are the bigger climate questions. We don't know at what temperature carbon dioxide and methane frozen in tundra regions will be released into the atmosphere. The polar ice regions are melting faster than expected.

There are a lot of tipping points and unknowns out there. We need to better understand what's going on and what the risks are. Petascale systems will give us a lot of detailed information from which we can make choices.

And petascale simulations of new materials may help us make better choices. For example, with fossil fuel power plants, if we can burn fuel at higher temperatures, we can increase the efficiency by up to 50%. If we can burn coal in a more oxygen-rich atmosphere, we increase the efficiency and make it easier to capture the carbon emissions.

But we don't have the steels able to withstand these higher temperatures. Thinking out loud, I wonder if detailed simulations can give us guidance on solving these metallurgy problems. If we can predict the properties without doing experiments, that would be a huge advance. There are many new technologies that can be introduced by improved materials science—safer, more efficient cars, photovoltaic materials for tapping solar energy, developing more energy-efficient buildings…

The other side of petascale systems is providing more computing capacity to researchers who may not be able to quickly scale up to tens of thousands of processors. By developing better methods, such as the adaptive mesh refinement work by Bell and Colella and their groups, scientists can improve the performance of their codes by several orders of magnitude while running on the same number or processors. Then when you combine these high-level programming frameworks with petascale computing, you can solve problems that are otherwise inaccessible to traditional methods, even when those methods are run on petascale platforms. And, as I mentioned, providing users with automated high-level programmability will also make both the researchers and the systems they use more effective and efficient.

Thank you for taking the time to answer our questions.