

# COMPUTING and DISCOVERY



Dr. Michael Strayer is the Associate Director of Science for Advanced Scientific Computing Research at the U.S. Department of Energy (DOE), and also serves as SciDAC Program Director.

**SciDAC Review:** *First of all, congratulations on the first petascale computers for science that the Advanced Scientific Computing Research (ASCR) program manages at Oak Ridge National Laboratory (ORNL), Argonne National Laboratory (ANL), and Lawrence Berkeley National Laboratory (LBNL). These are arguably the most powerful computing capability for science and engineering. What does this responsibility mean to you and for the Office of Science?*

**Dr. Michael Strayer:** It's a great responsibility because the country is becoming dependent on, that is, becoming increasingly based on, computation and large-scale computational science as the basis for scientific competitiveness. ASCR and the Office of Science have the vanguard responsibility to make sure that the U.S. has the most advanced computing capability for open science. The three laboratories are a mechanism for doing that. It's a great responsibility, and we are very lucky to have great laboratories to work with to carry out this responsibility.

*The landscape of science, discovery, and computing does not stand still; and in this environment, growth needs to be measured in terms of applications and achievements. Can you comment on the major science breakthroughs from these leadership systems?*

There have been many breakthroughs across many domains in science ranging from computational astrophysics, to computational biology, to nuclear energy, and to materials science.

Recently a committee from across the facilities reviewed the program ("Top Breakthroughs in Computational Science" p 32). Here are just a few of the projects highlighted in that review:

- **Modeling the Molecular Basis of Parkinson's Disease (life sciences).** A group from the University of California–San Diego is researching a treatment for Parkinson's disease. This group is using the Blue Gene/P to simulate how proteins called alpha-synucleins damage neurons.
- **Discovery of the Standing Accretion Shock Instability and Pulsar Birth Mechanism in a Core-Collapse Supernova Evolution and Explosion (astrophysics).** Core-collapse supernovae are the dominant source of elements in the Universe, including all elements between oxygen and iron and half the elements heavier than iron; life would not exist without these elements. This research, being done on the Cray XT4, is critical to realistic modeling of the neutrino shock reheating believed to be central to the supernova explosion mechanism.
- **Prediction and Design of Macromolecular Structures and Functions (computational proteomics).** A group from the University of Washington is using the Blue Gene/P to develop tools to help experimentalists identify the structure of biologically important proteins for which experimental X-ray phases are not available or are hard to obtain. The tools will also be used to refine early-stage nuclear magnetic resonance (NMR) structures to significantly speed NMR structure determination. This will help decipher which portion of the protein surface is responsible for a particular function, since many proteins carry out multiple complex cellular functions.
- **Understanding How Lifted Flame Stabilizes in a Hot Coflow (combustion).** A group from Sandia is using the Cray XT4 for direct numerical simulation and large eddy simulation to capture the complex aero–thermo–chemical interactions of flames. The goal is to understand the effects of fuel variability.
- **New Insights from LCF-enabled Advanced Kinetic Simulations of Global Turbulence in Fusion Systems (plasma physics).** A group at the Princeton Plasma Physics Laboratory is using an advanced version of the Gyrokinetic Toroidal Code to study plasma micro-turbulence in magnetically confined high-temperature plasmas. Simulations on the Blue Gene/P will enable a realistic examination of the influence of collisions on long-time, steady-state plasma transport behavior.

*Nearly five years after the Leadership Computing program was launched by the department, are you pleased with the progress that has been made?* I am very pleased with the progress. Back in 2004 when the initiative was launched, the systems that were available to the community were measured in the low tens of teraflops. Now we

---

are planning a petaflop capability. The next-generation systems will be in the tens of petaflops, and we are looking forward to exascale. In the past five years it's been amazing growth, and it's been a wonderful opportunity for the community to have impressive new capabilities. Whether the next five years can be as productive as the current five years will depend on the support we get from Congress and the community.

*What do you think about supercomputers as major scientific facilities? What do we need to do for high-performance computing to truly emerge as the third leg of scientific discovery, together with theory and experiment?*

Two things make computing at the moment different from other large experimental facilities:

One is that computers—large-scale, high-performance computers—have a lifetime that is measured in only a few years, maybe four or five years, whereas other large-scale facilities have lifetimes that are measured in decades. There has to be constant renewal for scientific computing facilities. That means you must constantly justify the reinvestment needed to keep them at state of the art. That is one major difference.

Another difference is the structure of scientific disciplines. In most fields the same discipline program office funds the facility and funds the science. It may be the case that over the next five years we will have to re-examine how to ensure that the resources are going into human capital and algorithm development needed to take advantage of large-scale systems. That may mean really embracing computational science as a discipline in and of itself that becomes more tightly associated with math and computer science and the facilities.

*As computers have become more powerful, they have stretched the infrastructure requirements. Are you surprised that space, power, and cooling have become critical factors in deploying leading-edge systems?*

I don't think that we are surprised that they have become incredibly important factors. Right now the systems require significant machine rooms and significant sources of power, but as we look forward, these requirements can't grow unbounded. They have to be managed in a way that the hardware that is developed is practical, not something suitable just for a handful of sites to use for science, but something that ultimately is deployable in industry and universities at some scale. Therefore, these power and infrastructure requirements are a major challenge for the vendors. They will really have to make big strides in power efficiency going forward.

*The processor technology is driving us to heterogeneous, many-core processors. What are you doing as the Associate Director of ASCR to prepare the scientific community?*

The first thing we are doing is developing a roadmap for computer science research that will go out about five years. Some of the high-level tasks that the roadmap will have to encompass are how to engage the research community in developing new algorithms and tools for multicore processors, and how to make a really large push toward new programming models—maybe, ultimately, programming models that are scale invariant and are not tightly coupled to the actual scale of the underlying hardware. We need breakthroughs in computer science to meet the challenges of these multicore opportunities.

*SciDAC has added a new dimension to science and its practice and has driven both pure and applied sciences to new plateaus of achievement.*

*How can we increase the penetration of SciDAC more broadly across all programs, particularly the technology programs?*

SciDAC has been an incredible success through partnerships where the Office of Science computing program partners with different scientific domains, application domains, and application program offices to fund joint projects. The challenge for expanding these projects is to find opportunities, particularly in engineering and technology where sustained high-end computing is going to have a revolutionary impact, and then to find a way to build long-term partnerships to co-invest in those. Some areas where progress is being made are nuclear energy and energy efficiency.

*You were quoted as saying: "The 21st century should pave the way to a millennium that excels in science, technology, and the way in which these disciplines interface with society. Advanced computing and computational science will be indispensable parts of the new ethos, and SciDAC will help lead the way." What are the top two or three impacts where you feel SciDAC played a role in shaping high-end computing over the years?*

SciDAC has had an enormous impact on scaling of climate model codes. Climate simulations on large-scale computing platforms are still a major challenge going forward, but there has been enormous progress. It is an application that is incredibly important to society. SciDAC has also made investments in fundamental physics, such as understanding the dynamics of supernovae. This may not sound like it has enormous impact on society, but it is about improving our understanding of the Universe and working out how computing can be as powerful a tool of understanding as observational science for astronomy. A third area of enormous impact is

---

## The next-generation systems will be in the tens of petaflops, and we are looking forward to exascale.

---

computational biology. We now have a large number of tools that will run at scale on large-scale platforms starting to be used for research, for example, on diseases such as Parkinson's.

But I think that beyond these science accomplishments, SciDAC also has fundamentally changed the way we as a community go about making progress in computational science. SciDAC was the first federal program ever that supported collaborations of mathematicians and computer scientists with applications scientists. When we started with SciDAC in 2001 this was a risky proposition, and it was by no means clear that this model would succeed. Now, eight years later the success is firmly established. Other agencies and several foreign countries are interested in joining SciDAC or creating similar programs of their own. I have had discussions with European and Japanese program managers about how to build such a program. Imitation is the best indication that we have created a paradigm shift in collaborative multidisciplinary science.

*Over the years, you have worked with a great many pioneers in the field of computational science. What are the most important partnerships that you have had?*

Of course I had many important relationships with colleagues in scientific applications, but there are two other important partnerships that I want to mention. In 2006 both the National Science Foundation (NSF) and the National Nuclear Security Administration (NNSA) joined ASCR in SciDAC. This was the first formal step in what I would hope will be a long-lasting collaboration among our agencies in promoting computational science. Based on the exascale town hall meetings in 2007, we know that tremendous challenges are ahead of us to move to the next level of computational science. I think we must work together with all available resources to reach this level. Therefore, I believe that partnerships with Dimitri Kusnetsov at NNSA and Ed Seidel at NSF will become increasingly important.

Another partnership that is critical to our success is the partnership with vendors. Our current petascale platforms are the results of collaborations between our labs and vendors. Early engagement with Cray in the Red Storm project and with IBM in the BG/L project led to the high level of performance that we see today in our facilities. These partnerships are critical to the success of our program, and I am very grateful that we had the support of IBM and Cray. In the future we are looking forward to strengthening and deepening these relationships and extending them to other vendors. As I said above, the challenges at the exascale are so daunting that we need full community involvement and support.

*As a computational scientist, what do you see as the major differences in the top issues with high-end computing from your early days to today?*

We have several challenges. One is to identify scientific and engineering targets that, if solved, will have a dramatic impact on developing new sources in energy, making energy efficiency more feasible, and power plants safer. Another challenge is energy applications that help us understand the impact of energy technologies on the environment, such as climate models, land use models, and water models that relate environment to energy in terms of energy production. A third area is defining really fundamental problems in basic science. Once we have defined the problems, we need to work out what are the software challenges we must solve in order to make these problems feasible to run on advanced computers. We will need progress in programming models, algorithms, libraries, and tools for management of data, visualization, and the systems in general. And finally we will need sustained progress in the next-generation architectures. To have the same impact in the next 10 years we have had in the past 10 years, we will need to grow computational capability by a factor of a thousand. That means developing new systems that will target exascale performance, making them usable enough so they can be effective tools for actually attacking the science problems. We need to get the science right, we need to get the software story right, and we need to bring the hardware along.

*What are the key challenges for the next 10 years?*

In the early days we were essentially focused on finding an algorithm to solve a problem. The space of possible algorithms

was still being researched by a lot of people, and the technical challenges of running the code were much less important. The machines, of course, were much smaller and memories were smaller; the principal challenge was the algorithm.

Now with massive parallelism, it is not sufficient to have an algorithm; we need to have sufficient parallel algorithms—that is a major difference. We also have enormous requirements for

---

SciDAC has fundamentally changed the way we as a community go about making progress in computational science.

---

tools and infrastructure at the leading edge because we are producing much more data, requiring substantial capabilities to move data around and to visualize it. Fifteen, twenty, thirty years ago one person could write a code, run it, and get a nice result. Now it takes a team and a lot of infrastructure to push the leading edge. It has become a team sport, a team venture.

*How has the DOE's view of high-end computing changed over the years?*

I think DOE's view has tracked the community. DOE was a leader in creating several things. DOE promoted the high-performance networking that enabled laboratories and DOE-supported university investigators to share data and to have access to computing facilities. DOE created the first really viable open supercomputer center, which later became the National Energy Research Scientific Computing (NERSC) Center. On the defense program side, DOE pioneered the modern supercomputer at the weapons labs. These things were essentially DOE creations. NSF later played a huge role in expanding things out into the university community and adding to the scale of the community, but DOE has really been the lead. I think that from this idea that we need only one center, or one center at a lab, now we need multiple centers that play synergetic roles with each other. We need comprehensive networking infrastructure, and we need to tie modeling and simulation aggressively to experimental activities in DOE that support projects like the Large Hadron Collider (LHC) and ITER internationally. It becomes a rich ecosystem as we move forward, and that's really the difference.

*What are the most important advances that are needed in HPC to solve science and engineering problems of global significance?*

What we have to do is make it possible for more scientists and engineers to make use of high-performance computing platforms effectively. That means we have to reduce the barriers for algorithm development, tools development, and modeling, through programming models, libraries, and abstractions. We have to build an international community that is able to freely share software and share access to open facilities. We have to really look at how we make balanced investments, where we invest in both the hardware and the computer science, and the

challenges of the software and also in building strong application co-teams. These co-teams are of substantial scale, so it will mean in some sense rethinking how we fund computational science at the high end and doing that in a national context.

*We spend a lot of time discussing hardware and computer design these days. However, system software plays a very large role in the usability of supercomputers. What is your vision for the future for the software stack, specifically operating systems and compilers, given that the coming generations of computer users are going to be facing systems with millions of cores?*

Clearly we are going to have to make it possible for vendors to build and support systems with millions of cores. The systems software strategy I believe should be open source, based on common requirements from across communities and across vendors, so that we can leverage the resources that we have to do the systems software. The software stack should be very open, and ideally we should have multiple solutions to each element of the system we need to deploy. We should reach out to partnerships to build this software. Scalable systems software is incredibly important; these systems are not going to be usable without it. It is one of the three things we need to make dramatic progress on, and we have made progress. There is a plan under way to build an international consortium to attack this problem that will have substantial investment from both DOE and other agencies and international developers.

*In computational science, as in the other sciences, the arena of international collaboration is the forum for science for the 21st century. How do we effectively foster international collaborations?*

Right now the U.S. is in a very dominant position in high-performance computing, and we should use that position to convene international partners to lay out a plan for the next decade and find ways of involving all the resources, both human resources and investment, necessary to do this. Right now is actually the time for international cooperation in this area. I think DOE is well positioned, along with NSF, to help foster these international programs. DOE has a lot of experience in doing this in other areas such as high-energy physics and magnetic fusion, so we have the framework for doing it. Sometimes there can be extra work in getting started, but it has been shown that these international programs have a way of sustaining operation over a long time. This is something that is really timely, and our office is going to make a high priority.

*Basic and applied sciences influence and empower society in multiple ways. How will advanced computing and computational science help lead the way?*

Computational science at the high end can be used in many ways to empower progress in society. First of all we can use large-scale computations to demonstrate a path forward for the first time, making it possible to solve a problem for the first time, and making it understood that computational approaches actually have value. This will result in more groups, whether it will be industry or other academic groups, to take the lead from the high end and make it more feasible over time to put computational science into practice. In many ways high-

performance computing can be used as a time machine to find out what is possible and to help to create paths for the future where more production oriented computing can follow. It also can be used as a way of actually bringing people together. It takes a team of people typically to build a code that is competitive at the very high end. These teams have a positive impact on the institutions in which they sit and the disciplines in which they sit, so it has both a direct effect and indirect effect.

*What advice do you have for young scientists just entering the R&D field in high-end computing?*

First of all my advice is they should learn how to be fearless. They should have a grounding in mathematics, physics, or the underlying discipline, and they should be willing to do whatever is necessary to play an important role in these teams as they are building up. Someone who is going to have a long-term impact on computational science and high-end computing needs to be able to understand a multitude of issues, from algorithms, to architecture, to software engineering, and to the details of their discipline, whether it is biology or physics. We are looking for

---

The challenges at the exascale are so daunting that we need full community involvement and support.

---

people who have this broad kind of perspective and are good team builders. These are some of the attributes that we need from the next generation.

*You are an ardent supporter of science education, including your personal involvement in such programs as Adventures in Supercomputing and Computational Science Education Partnership. What should we do to educate the next generation of computational scientists and engineers?*

We have to try many things; we have to do many things. We need programs that get kids interested in computing at an early age. We need programs that support math and science education more generally. We need opportunities for young people to get involved in internships and activities at national labs and universities. We have had enormous success with projects like the Computational Science Graduate Fellowship program, which aims at graduate-level support. We have also had success in programs like Adventures in Supercomputing that gets kids exposed to computational science and modeling and simulation at an early age. Kids like to see these big computers; they are fascinated with them. So, not only do we need to capture their imagination and make it interesting for them, but we also need to grow the field so that when the young people go into the field, they have lots of opportunities to become professionals in their field over time. We also need to have outreach and engagement at a young age and be diligent in building out the communities to expand the space of opportunities for people as they go forward.

*Thank you for taking the time to answer our questions.*