

Leading the Push Toward Exascale



Dr. Thom Mason has served as Laboratory Director at Oak Ridge National Laboratory since July, 2007.

SciDAC Review: *Prior to becoming Lab Director at ORNL in 2007, you were the Director of ORNL's Spallation Neutron Source (SNS), now the world's most capable neutron source and a major user facility for neutron science. Given your perspective on big experimental science at a national laboratory, how do you view DOE's supercomputing enterprise in context?*

Dr. Thom Mason: Supercomputers are similar to big accelerator facilities, but they are inherently multipurpose (that is, Turing machines). They have more than one potential strategic outcome. Accelerators, like synchrotron sources and neutron sources, are mission specific and therefore are built by program offices, and these offices pay for the requisite scientific instrumentation and own the outcomes. For example, the Tevatron at FermiLab and the SNS at ORNL are both proton accelerators, but the work they do is very different. A consequence of supercomputing having more than one potential strategic outcome is the potential for a lack of clarity about high-level outcomes and who owns them, as well as who pays for development and maintenance of the scientific software that is the requisite "scientific instrumentation" for these facilities. This is why programs like Scientific Discovery for Advanced Computing (SciDAC) are important in giving the programs that rely on high-performance computing a window into the facilities.

Most people would agree that having a facility with the capability to solve more than one strategic problem is a good thing. Are there examples within the DOE R&D enterprise where the ownership of high-level outcomes from computational science is vague?

The National Nuclear Security Administration (NNSA), with its

strong focus on nuclear security, supports its own computing hardware, software, and application development infrastructure. I would say that NNSA is very focused on strategic outcomes. DOE's energy technology offices do not currently focus major support for computing infrastructure or use because their use is still not as well developed in the applied sciences. DOE's Office of Science supports world-leading computing infrastructure, including hardware and software, for its missions of discovery and use-inspired science. The Office of Science mission is appropriately very broad, and high-performance computing has the potential to accelerate progress in every problem on which the Office is working. But because of this breadth of programmatic and mission impact, ownership of the high-level outcomes can be unclear.

So, what are the consequences of this situation?

The consequence for delivering DOE mission impact in energy and science programs is the existence of gaps, or programmatic "white space," that have constrained the contribution computational science can make to our mission impact. As I mentioned, to first order DOE's energy technology programs do not fund computing or its application. For example, there exists no equivalent of the SciDAC program for energy technologies. However, we know – the report on the exascale Town Hall meetings is one recent source of information – that energy technology research is rich with opportunities for computing. I should mention that the recently announced competition that will occur this year for the Nuclear Energy Modeling and Simulation Hub is a welcome opportunity to demonstrate the role that supercomputing can play in an important energy technology program. Within the Office of Science, some programs tend not to focus on the development of applications for high-end computational science. The SciDAC program is not a complete solution – for example, not all Office of Science programs have adopted the SciDAC funding model. There is sporadic support for software sustainability or "application development." As a consequence, we are constantly at risk of losing the utility of generations of codes developed with DOE funding for the reason that funding to modernize these codes – for example, the insertion of better physical models and numerical algorithms – does not exist. Moreover, there is very strong international competition (from Europe, Japan, China, India, and others) for leadership in scientific software development within DOE's mission space. The leadership of important topical areas that DOE pioneered now resides outside the United States as the result of the absence of programmatic focus to maintain leadership in these capabilities. In addition, the appropriate

translation from basic research in computer science and math to computational science impact is broken in places.

So, the result is that national laboratories must develop creative solutions for computational science impact on DOE missions. For example, at ORNL we have prioritized discretionary investment toward building interdisciplinary teams of domain scientists, computer scientists, and mathematicians to develop new scientific applications (“codes”) in our priority areas – for example materials science, nanoscience, nuclear energy, and climate science – which can take advantage of our state-of-the-art supercomputers.

DOE's big synchrotron and neutron facilities address a broad range of science topics in materials research, from biology and chemistry to structural and engineering materials. Do the organization and structure of these facilities offer insight into possible ways to address the gaps in computational science research and produce a healthy software ecosystem?

ORNL's proposal to DOE in 2004 for the Leadership Computing Facility drew directly on the model established for DOE's large-scale experimental user facilities. At the SNS and other user facilities, scientists and engineers make use of “end stations” – best-in-class instruments supported by instrument specialists. By organizing software specialists and community codes into “computational end stations” in climate science, chemistry, and materials science, ORNL and other national laboratories have been able to offer the research community access to best-in-class scientific application codes and world-class computational specialists, building highly productive collaborations like the one that recently received the Gordon Bell Prize for analyzing the magnetic properties of materials. However, the computational end station concept, appropriately differentiated from the experimental analogs, is still not widely embraced. At ORNL, our successes have typically been initiated through our discretionary research program and more at a grassroots level.

DOE's SciDAC and Innovative and Novel Computational Impact on Theory and Experiment (INCITE) programs have played a major role in focusing the Department's high-performance computing resources on grand challenges in science and engineering and in attracting both DOE researchers and scientists and engineers from industry and the academic community to these resources, again in parallel with the experience at user facilities.

These are manageable challenges, and we think we have a model that is effective in building the requisite interdisciplinary teams and community software codes. But questions remain as to how the programmatic gaps should be filled. Should there be a SciDAC-like program for use of the NNSA and Office of Science computing resources by the energy technology programs? Should the Office of Science facilities be used to accelerate progress on DOE's energy mission? Should the energy technology programs build their own computing facilities? In Office of Science programs, should additional science scope be added to the computing programs? Should computational end stations be more common? What are the best practices in these areas on which to build? Right now these are open questions.

Could you discuss how HPC relates to other scientific tools?

Since the 1980s, we've been hearing that modeling and simulation are the “third leg of science,” along with theory and experiment. We're now beginning to see concrete examples of

Bio in Brief: Dr. Thomas E. Mason

Thom Mason is originally from Halifax, Nova Scotia, Canada. He majored in physics at Dalhousie University in Halifax, where he met his wife, Jennifer McGillivray, and received his Ph.D. from McMaster University in Toronto.

Thom came to the United States as a postdoctoral fellow at AT&T Bell Laboratories in New Jersey in 1990. He moved to Denmark in 1992 to work as a senior scientist at the Risø National Laboratory. In 1993, he returned to Canada as a professor in the Department of Physics at the University of Toronto in Ontario.

Thom joined Oak Ridge National Laboratory in 1998 to participate in the Spallation Neutron Source (SNS) project. He was named Associate Laboratory Director for the SNS in 2001, leading this \$1.4 billion project to completion on time, on budget, and on scope with an outstanding safety record. In 2006 with the SNS in operation, Thom was named Associate Laboratory Director for Neutron Sciences. In July 2007 he became director of ORNL and president and CEO of UT-Battelle, LLC.

As director of ORNL, Thom is responsible for the management of DOE's largest multipurpose science and energy laboratory, with 4,250 employees and an annual budget of more than \$1 billion. Under his direction, the Laboratory carries out a broad program of research and development that spans advanced materials, neutron science, high-performance computing, systems biology, energy, and national and homeland security.

Thom has authored or co-authored more than 100 papers in the open scientific literature. He is a Fellow of the American Association for the Advancement of Science and the American Physical Society. In 2005, he received the Postma Young Professional Award from the East Tennessee Economic Council, and in 2008 he received the Distinguished Alumni Award for the Sciences from McMaster University.

Thom and Jennifer live in Oak Ridge with their two sons, William and Simon. Thom serves as Chair of the Oak Ridge Public Schools Foundation, which helped to raise \$55 million for the renovation of Oak Ridge High School. He also chairs the Board of Directors of Innovation Valley Inc. and is a member of the ETEC Board of Directors.

this. For instance, a team of researchers at ORNL has used the Bio-SANS instrument at the High Flux Isotope Reactor and the OLCF supercomputers to investigate biomass samples from our BioEnergy Science Center, with the aim of obtaining a molecular-level understanding of the breakdown of plant cell wall structures during biomass processing. By comparing best-fit models of the lignocellulose molecules derived from experimental neutron scattering profiles to the results of computer simulations, they are gaining a detailed understanding of the behavior of water and of molecular order in these systems. This will move us toward cost-effective production of ethanol and other liquid fuels from renewable biomass, supporting DOE's goals for energy security.

Climate change science provides another example with a different set of challenges – here we need HPC not only to perform the high-resolution, long-duration simulations that are required

to understand the Earth's climate system, but also to manage the increasingly massive and complex datasets that result from these simulations and to integrate them with observational and experimental data to create useful and usable resources for the research community.

Finally, integration of experiment, theory, and simulation has been essential to advances in fusion energy science; we can trace the roots of DOE's scientific supercomputing centers back to the Controlled Thermonuclear Research Computer Center, which evolved into the Magnetic Fusion Energy Computer Center and then into the National Energy Research Supercomputer Center. With ITER on the horizon, the computational fusion community is focusing on obtaining more accurate predictions of performance and providing the modeling and control tools that will be needed to make the most of this multibillion-dollar investment in the next step toward fusion power.

ORNL is executing computational science programs for multiple agencies – DOE, the National Science Foundation (NSF), and, in the near future, The National Oceanic and Atmospheric Administration (NOAA). What additional value is there in a national laboratory executing such a multi-programmatic strategy that might not be available from a single program sponsor?

National laboratories are experienced in the development and deployment of the large multidisciplinary teams needed to answer complex science questions. We can also leverage both our experience and expertise in computing and the extensive computing infrastructure that DOE has put in place to meet the complementary needs of multiple customers. DOE has made the specialized capabilities of the national laboratories available to other federal agencies for many years through the Work for Others program. The cost and demands associated with delivering and operating exascale computers mean that only a few institutions will be able to offer these systems, and we can provide the best return to the taxpayer by making them available to the appropriate teams to solve problems of national importance.

For example, the technical issues of operating these massive computing machines and of scaling complex applications codes to run effectively on hundreds of thousands of processors are common to the work we perform for DOE, NSF, and now NOAA. This is expertise that we have and are investing our discretionary funds to further develop.

ORNL's multiagency success has resulted in a significant increase in demand at ORNL for highly-skilled computational and computer scientists with distinctive and sometimes unique experience in very large-scale parallel computing. Building and operating a stable computing machine with over one hundred thousand processing cores and getting real science applications to scale to the full size of such a machine is no small task. How are your recruiting efforts going? Are there enough appropriately trained people out there to staff your workforce?

High-performance computing is still a relatively small community, so we have to work hard to recruit qualified technical staff. We have been relatively successful in attracting good people, and our recruiters have been essential and diligent in this effort. But we need more! We think that ORNL is a great place to work and East Tennessee is a great place to live. Our facili-

ties require world-class staff to operate, and our leadership computing hardware and multi-programmatic environment helps to attract these folks. There is much to be gained by fostering interactions among a critical mass of experts focused on a common set of reliability and scalability issues. And these issues are common to the DOE, NSF, and NOAA facilities. It is nice to be in the position of needing to find good people.

Over the long term, we expect to see an increasing demand for scientists, computer scientists, and engineers trained to work at the very large scale. We need to work with the leading universities to establish training and graduate education programs to attract and prepare students to enter this field. For example, in January 2010, Tennessee Governor Phil Bredesen announced a new graduate education department of the University of Tennessee to be co-located on the ORNL campus. Over the next few years, we expect this initiative to double the number of graduate students working at ORNL.

We expect to see an increasing demand for scientists, computer scientists, and engineers trained to work at the very large scale.

Looking to the future, advancing computational science from the petascale today to the exascale in a few short years will be a challenging task. From your perspective, why is this investment of resources and talent a good idea? How will it pay off?

It is well established from an economic perspective that research in general, and computing and computational science in particular, are drivers for growth. An excellent example is the innovation that has occurred in the design and manufacturing of advanced airplanes and the corresponding decrease in the number of large wind tunnels needed by this important industry. Computing speeds design, lowers cost, and enhances our competitiveness. Boeing, in designing its new 787 Dreamliner, used first-rate computational science and our Office of Science computing facilities to solve a complex wing design problem and incorporate advanced carbon fiber materials into its designs. In the 1980s, when it was designing the 767, Boeing built 77 wings for testing. In designing the 787 in recent years, Boeing built only seven wings for testing. This is really what we mean we say that leadership computing is an important contributor to innovation and competitiveness. It is also an excellent example of how the tools that were initially motivated by problems in the science domain find application in areas of industrial relevance just as has been the case with accelerators and reactors. Over the long term, leading the push toward exascale computational science, I believe, will mean the solution of the nation's top science, energy, and security challenges, and this will translate into an improved quality of life for all of us.

Thank you for taking the time to answer our questions.