Advanced Scientific Computing Advisory Committee

February 2007 ASCAC Meeting Highlights

The primary business at the DOE Office of Science (SC) Advanced Scientific Computing Advisory Committee (ASCAC) 27-28 FEB 2007 meeting was to review and accept the Petascale Metrics Panel Report as an ASCAC Report. This Report, presented to the ASCAC by Petascale Metrics Panel Chair, ASCAC member Gordon Bell, provides the ASCAC response to the March 2006 Charge from Dr. Raymond Orbach, Under Secretary of Energy for Science, "... to weigh and review the approach to performance measurement and assessment at [the Department of Energy Office of Science High Performance Computing Centers and related projects], the appropriateness and comprehensiveness of the measures, the [computational science component] of the science accomplishments and their effects on the Office of Science's science programs, ... [and] the role and computational needs over the next 3–5 years" ("Highlights of the ASCAC March 2006 Meeting," *SciDAC Review*, Fall 2006, p4).

The Petascale Metrics Panel Report describes four aspects of science accomplishment that the Panel recommends that the ASCR High Performance Computing Centers should report and track "to assist in the measurement of scientific output from its projects: Publications, Code & Datasets, People, and Technology Transfer…; Project Milestones versus the Proposal's Project Plan …; Exploiting Parallelism and/or Improved Efficiency: aka Code Improvement …; and Break-throughs, an immeasurable goal. The Panel could not identify metrics or any method that could be used to anticipate discoveries that occur on the leading edge of fundamental science."

The ASCAC unanimously endorsed the Petascale Metrics Panel Report, noting that the difficult problem of measuring scientific understanding and accomplishment needs additional study. As observed by ASCAC Chair, Dr. Jill Dahlburg, in a paper titled "The S&T Innovation Conundrum," there are at least two distinct phases in the history of a major scientific accomplishment. "There is an early, searching phase that is evocative of prospecting. It is characterized by a few discrete but high-impact events. There is little functional capability produced during this phase and the individuals contributing to the discrete events, while generally sure that they are involved in profoundly exciting research and technology, often have no idea what the ultimate functional capability will be. Continuing with the analogy, this early phase is followed by a later, more predictable phase that is much like mining. This latter phase seems to be dominated by continuous improvement in functional capability as characterized by a larger number of lower-impact innovations than occurred in the early phase. During the mining phase, the capability produced can usually be related to the funding applied and to the inherent potential of the technology being exploited, resulting in desirable features such as measurable and predictable return on investment." Metrics, while in some sense essential for the mining phase, are generally not applicable to the quantification of success in the prospecting phase of a major scientific breakthrough.

Dr. Ellen Stechel, ASCAC member and Chair of the ASCAC Networking Subcommitee, provided an update about that status of the report to address Dr. Orbach's 10 MAR 2006 Charge to examine the role and efficiency of networking and networking research within the DOE SC. In the report, which is due in November 2007, the Subcommittee hopes to produce useful guidance on networking research and the priority of networking needs.

On February 28, two new Charges from Dr. Orbach were provided to the ASCAC. The first is a Committee of Visitors (COV) Review of SciDAC, which will be led by ASCAC member Dr. Robert Voigt. The second is a Joint Committee with the Biological and Environmental Research Advisory Committee (BERAC) to examine the issue of computational models for Genomes to Life (GTL). ASCAC member Dr. Rick Stevens has agreed to serve as ASCAC Co-Chair of this Joint Committee, the report for which is due at the August 2007 ASCAC meeting.

Contributor: Written by Dr. Jill Dahlburg of the Naval Research Laboratory, these highlights were largely summarized from the ASCAC Petascale Metrics Report, The S&T Innovation Conundrum report, and the February 2007 ASCAC Meeting Minutes (Fred O'Hara, recording secretary)

SciDAC PI Meeting

Expanding SciDAC's Collaborative Network

In February, nearly 80 SciDAC-supported researchers, project managers, and support staff gathered for a day-and-a-half meeting in Atlanta, GA. The purpose of the meeting was to reinforce existing partnerships and discuss potential new areas of collaboration. In addition to presentations about successful coll- aborations between scientists, applied mathematicians, and computer scientists under the first round of SciDAC projects, the Atlanta meeting also featured two poster sessions designed to catalyze discussions among the representatives of the new SciDAC Science Applications, Centers, and Institutes.

"Without the posters, it would have been very hard, if not impossible, for us to communicate with one another about some of the key concepts, challenges, and capabilities," said Dr. Kwan-Liu Ma, leader of the SciDAC Institute for Ultrascale Visualization. "The posters also enabled us to quickly identify potential collaborations, and limiting the meeting size made it possible for me to get to know and also talk to almost all the PIs. It was also helpful to hear the answers of many questions asked by others … [questions] I did not think of."

One interactive poster featured a matrix in which PIs were asked to identify known and potential collaborations. The data collected from this poster will be used to create a Web-based display of SciDAC partnerships, which are an integral component of the program's strategy for achieving computational science breakthroughs. More information can be found at the SciDAC Outreach Center website: http:// outreach.scidac.gov/

Conference

SciDAC 2007

The 2007 DOE Scientific Discovery through Advanced Computing (SciDAC) conference is scheduled for June 24–28 in Boston, MA. SciDAC 2007 will showcase the results of DOE-sponsored computational science across a range of disciplines, with an emphasis on science at scale. The conference will also feature computational science from outside of the SciDAC Program and will foster communication between SciDAC researchers and those funded by other agencies. The Conference Chair for 2007 is Dr. David Keyes of Columbia University.

More information about the meeting can be found online at: http://www.scidac.org/Conference2007/

Advanced Scientific Computing Research

Workshops Support Key Areas of Research

In December 2006 and January 2007 the Advanced Scientific Computing Research (ASCR) program held several workshops to identify research opportunities in support of collaborations across the department: Mathematical Research Challenges in Optimization of Complex Systems; Computational Subsurface Sciences; and Cybersecurity Research and Development (R&D) Challenge for Open Science.

In December 2006, a diverse group of distinguished mathematical scientists gathered in Bethesda, MD to consider Mathematical Research Challenges in Optimization of Complex Systems. The goal of this workshop was to identify opportunities for mathematical research relevant to DOE applied science and technology programs in areas that are not already a major part of the Office of Science (SC) applied mathematics portfolio. Domain experts gave presentations on four applications of great interest to the DOE: advanced fossil fuel power generation; the nuclear fuel lifecycle; power grid control and optimization; and risk assessment for cybersecurity.

To determine the basic research needs for computational subsurface sciences, the Computational Subsurface Sciences Workshop was held in Bethesda, MD during January 9–12, 2007. Collaborating DOE offices were SC, the Office of Environmental Management (EM), the Office of Fossil Energy (FE), and the Office of Civilian Radioactive Waste Management (RW). The purpose of the workshop was to obtain community input on computational science research needs and opportunities in the subsurface sciences and related areas, with a focus on developing a next generation of numerical models of subsurface flow and process simulation. Highlighted areas included potential terascale (and future petascale) computational algorithms to enable high-fidelity subsurface simulation models that fully couple key physical, chemical, geological, and biological processes with new capabilities to quantify and reduce model uncertainty.

A cybersecurity workshop, the Cybersecurity R&D Challenge for Open Science, was held on January 24–26, 2007. This workshop brought together cybersecurity production personnel, cybersecurity researchers, and scientific application researchers from across the DOE laboratories and DOE headquarters. The purpose of the workshop was to identify the research needs and opportunities associated with cybersecurity for open science. The meeting focused especially on those needs associated with DOE supercomputing, user facilities, high-speed networks, laboratories, and other open collaborative science stakeholders. A discussion of how open science cybersecurity differs from general cybersecurity explored the implications these differences may have for cybersecurity research activities.

Full reports from these three workshops can be accessed online, at: www.sc.doe.gov/ascr/Misc/WorkshopsConferences.html

Hardware

Upgrade at ORNL

ORNL

Figure 1. *In April 2007, the ORNL Leadership Computing Facility increased the computing power of Jaguar, a Cray XT-3 computing system, to 119 teraflops, making it the most powerful open scientific system in the world.*

The Oak Ridge Leadership Computing Facility (LCF) recently doubled the performance of the Cray XT-3 supercomputer, nicknamed Jaguar (figure 1), at the National Center for Computational Sciences. Work on doubling the size of the system was a substantial undertaking, involving 124 cabinets and 11,708 dual-core AMD Opteron processors. The system has a rating of 119 teraflops of peak performance, which translates to 119 trillion mathematical calculations per second. The enhanced computer also features 46 terabytes of memory and 750 terabytes of disk storage.

The LCF at ORNL is on a path to reach 250 teraflops by late-2007 as dual-core processors will be replaced with AMD quad-core processors. The memory will also be doubled and the operating system will migrate to a version of Linux on the computer nodes.

FY 2008 Budget Request

Increases Drive Science Towards the Petascale Era

The President's fiscal year (FY) 2008 budget request was delivered to Congress in February 2007. This budget request furthered the President's commitment to the American Competitiveness Initiative (sidebar "American Competitiveness Initiative," *SciDAC Review*, Fall 2006, p53) by providing significant increases throughout the Office of Science, including Advanced Scientific Computing Research (ASCR). Of particular note were increases in facilities that will initiate the era of petascale computing for open science at the ORNL Leadership Computing Facility.

In addition, there were significant increases in core research in applied mathematics to support critical long-term mathematical research issues relevant to petascale science (+$2,000,000), research in optimization control and risk analysis in complex systems (+$1,900,000), and increased research in multiscale mathematics (+$2,505,000). Total support for multiscale mathematics will be $11,000,000 in FY 2008.

Increases in core research in computer science will have a focus on strengthening long-term research in computer science to ensure that scientific applications can fully exploit future hardware performance. In the software environment area, additional funding (+$4,137,000) will focus on developing software and tools to enable both experienced and new researchers to make effective use of petascale systems at the national Leadership Computing Facilities and at the supercomputer facilities. These new enabling tools will include integrated/intelligent software development environments and new generation debugging and performance analysis tools providing improved ease-of-use at the petascale. Budget increases in data analyses, management, and visualization (+$1,000,000) will focus on accelerated visualization, uncertainty, and user interface environments.