ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE

MEETING MINUTES

American Geophysical Union
2000 Florida Avenue, NW – Washington, DC

March 27 – 28, 2012

PARTICIPANTS

Advanced Scientific Computing Advisory Committee (ASCAC) members present:
Roscoe Giles, Chair
Marsha Berger
Marjory Blumenthal
Vincent Chan
Jackie Chen
Barbara Chapman
Jack Dongarra
Susan Graham
James Hack
Wendy Huntoon
Sharon Glotzer (by telephone)
John Negele (by telephone)
Linda Petzold
Vivek Sarkar (by telephone)
William Tang
Victoria White
Dean Williams

Members of the U.S. Department of Energy (DOE) present for all or part of the meeting:
William Brinkman, Director, Office of Science
Daniel Hitchcock, Acting Associate Director, Office of Advanced Scientific Computing Research (ASCR), Office of Science
Christine Chalk, ASCAC Designated Federal Officer
Steve Binkley
Laura Biven
Richard Carlson
Vince Dattoria
Susan Gregurick
William Harrod
Barbara Helland
Marcos Huertz
Sandy Landsberg
Randall Laviolette
Thomas Ndousse-Fetter
Lucille Nowell
Karen Pao
Brian Plesser
Sonia Sachs
Jeff Salmon
Yukiko Sekine
Ceren Susut

Others present for all or part of the meeting:
S.F. Ashby, Pacific Northwest National Laboratory
Bill Bartolone, Intel
Mitra Basu, National Science Foundation
Jonathan Carter, Lawrence Berkeley National Laboratory
Alok Choudhary, Northwestern University
Steve Collis, Sandia National Laboratory
Dona Crawford, Lawrence Livermore National Laboratory
Lori Diachin, Lawrence Livermore National Laboratory
Paul Ducette, Battelle
Stephen Elsby, Research Councils UK
Ian Foster, Argonne National Laboratory
Ray Grout, National Renewable Energy Laboratory
Benjamin Grover, Lawrence Berkeley National Laboratory
Mark Guiton, Cray Computers
Charles Hayes, HCS
Jeff Hittinger, Lawrence Livermore National Laboratory
Vasant Honavar, National Science Foundation
Paul Houland, Argonne National Laboratory
Fred Johnson, SAIC
Gary Johnson, Computational Science Solutions
Brad Keelor, British Embassy
Moe Khaleel, Pacific Northwest National Laboratory
Omar Knio, Duke University
Alejandro Lopez-Bezanilla, Oak Ridge National Laboratory
Arthur Macek, Oak Ridge National Laboratory
Reinhold Mann, Lawrence Berkeley National Laboratory
John May, Lawrence Livermore National Laboratory
Michael McGuigan, Lawrence Berkeley National Laboratory
Paul Messina, Argonne National Laboratory
Jeff Nichols, Oak Ridge National Laboratory
Esmond Ng, Lawrence Berkeley National Laboratory
Jim Rogers, Oak Ridge National Laboratory
Paul Runci, Pacific Northwest National Laboratory
Nagiza Samatova, North Carolina State University
Lee Schroeder, Tech-Source, Inc.
Ray Stults, National Renewable Energy Laboratory
Peter Swart, Los Alamos National Laboratory
David Turek, IBM
Jack Wells, Oak Ridge National Laboratory
Julia White, Oak Ridge National Laboratory
Victoria White, Fermi National Accelerator Laboratory
Kathy Yelick, Lawrence Berkeley National Laboratory

MEETING MINUTES

Tuesday, March 27, 2012

The meeting of the Advanced Scientific Computing Advisory Committee (ASCAC) was convened at 9:00 a.m. EST by ASCAC Chair Roscoe Giles. Dr. Giles welcomed new ASCAC members Marjory Blumenthal, Vincent Chan, Wendy Huntoon and Dean Williams.

Presentation: View from Washington, Office of Science

Dr. William Brinkman, Director, Office of Science (SC), gave an update on the FY 2013 budget request for the SC. The request of 4.992M is 2.4 percent higher than the FY12 appropriation of 4.873M. Dr. Brinkman believes that funding will stay constant for the next five to 10 years.

SC achievements are at the frontiers of science with more than 100 Nobel Prizes given to SC-supported researchers over the past 60 years. SC provides 45 percent of all Federal support for physical and energy-related sciences and research in biology and computing. SC supports more than 25,000 Ph.D. scientists, students, and support staff at more than 300 institutions. SC's portfolio of science user facilities is the largest in the world and supports 26,500 users.

Research support can foster innovation and economic growth, and reflects the Administration's interest in science and technology, and continuing interest in basic research with an emphasis on energy. This motivates DOE's focus when it comes to budget decisions.

Science underpins America's technological and energy future, and computing is a major component. As an example, quantum mechanics knowledge has lead to other innovations. A highly-trained workforce and science, technology, engineering and mathematics (STEM) education is also vital to economic development. Dr. Carl Wieman of the White House Office of Science and Technology Policy (OSTP) is concerned about this, particularly attrition in undergraduate engineering programs.

The DOE laboratories and ASCAC efforts can impact American society. The work of the labs is not fully known. Their charters have changed in the last 50 years from pure and basic research to providing unique tools and diverse activities. SC needs to focus on whether labs are doing the right things for the country and their influence on its future.

Clean energy is one area impacted by DOE centers and new centers. Proposals for the new battery hub center are due on May 1 with reviews and an award announcement to take place over the summer. Dr. Brinkman highlighted work by the Office of Energy Efficiency and Renewable Energy (EERE) and in critical materials with the hope that a new funding opportunity announcement will appear in the next few months.

Themes driving the SC are: 1) Materials science, including materials analysis, 2) Biosystems and work to find new microbes and genetic codes, and 3) Modeling and simulation. Basic science underlies this and DOE must stay at the forefront of high-performance computing (HPC).

Dr. Brinkman highlighted Materials and Chemistry by Design, in particular efforts to identify desired points in the spectrum by color-coding them in order to try to synthesize them. This will lead to an online software tool available for public use. He also emphasized work being done in Biosystems by Design and a focus on design engineering, genomics and tools.

Despite some budget decreases, Dr. Brinkman noted the 3.3 percent increase for Advanced Scientific Computing Research (ASCR), a 2.6 percent increase for Biological and Environmental Research (BER), and a 6.6 percent increase for Basic Energy Science (BES).

ASCR is constructing and managing next generation computers for exascale computing. There are new peta-flop (PF) machines at the Argonne and Oak Ridge national laboratories. The exascale movement needs a roadmap and plan for issues such as power use and limiting statistical errors while computing. A high point is DOE facility use by industry and DOE software development.

DOE's INCITE is partnering with industry. GE, Ramgen, Boeing, General Motors and United Technologies are all customers. Dr. Brinkman described modeling studies for long-haul trucks that draw on HPC and are up to 15 percent more efficient due to aerodynamic add-ons.

The BES unit is meeting clean energy goals with research in science-based chemical and material discovery. This starts with theory and modeling and that work needs to be sustained. DOE is finishing the NSLS-II (National Synchrotron Light Source) at Brookhaven. DOE will also continue to leverage x-ray laser capabilities that for the first time positioned a laser in the x-ray region. This work amplifies an x-ray line to achieve a factor of a 40 decrease in inline width.

The 46 Energy Frontier Research Centers (EFRC) are part of BES. The Centers are under review with 50 percent done at this time. A specific location for the future battery hub is undetermined.

The Materials by Design component is advancing work in light absorption in wire arrays and researchers are developing new approaches for solar cells. Silicon rods are placed on a substrate and the space between is filled with aluminum oxide. Light bounces around and that gets absorbed and generates electricity rather than heat. This achieves 28 percent greater efficiency.

Dr. Brinkman mentioned the Fuels from Sunlight hub and the partnership between the University of California-Berkeley and CalTech.

DOE research is advancing light source facility upgrades and new instrumentation. BES needs to upgrade and enhance facilities, such as the LCLS (Linac Coherent Light Source), the NSLS, and the APS (Advanced Photon Source).

Forty-four of the Fortune 500 companies are BES facility users. GE uses cyclotrons to characterize sodium sulfide batteries for energy storage on a grid. The Eli Lilly Corporation has a beam line at Argonne for advanced drug development and the examination of different proteins and drug interactions.

The BER has reviewed three bioenergy centers. Each has worked to understand the basics of breaking down plants for fuel. The BioEnergy Science Center studies the lignin aspect of biofuels to change DNA and prepare to make fuels. The Joint BioEnergy Institute is focused on microbes that take cellulose and produce fuels. And, the Great Lakes Bioenergy Research Center is examining plants and soils and the production of biofuels without impacting the food supply.

BER's Joint Genome Institute is advancing metagenomics to support microbe culturing and the complicated interactions among communities of microbes. This helps climate research and climate modeling. DOE works with the U.S. Global Change Research Program to coordinate inter-agency climate research.

Climate studies have lead to DOE participation in developing a computer earth model and understanding the aerosol and cloud interaction. An important thing is DOE's contribution to work being done at Lawrence Livermore National Laboratories (LLNL). This impacts the work of the Intergovernmental Panel on Climate Change.

BER is working with other agencies to get climate modeling efforts out of isolation. This research is coordinated and somewhat motivated by attacks by anti-climate opponents.

Fusion Energy Sciences is supporting the International Thermonuclear Experimental Reactor (ITER). The project budget is around $25B. The European Union will provide $1B Euros this year and the next. The U.S. will provide $145M and deliver in-kind pieces such as the water cooling system and central solenoid. The U.S. sends little money overseas with around 80 percent of U.S. funds being spent here. The ITER facility's foundation is based on pedestals with a shock-absorbing cushion. This will help the structure withstand earthquakes and vibrations. General Atomic's DIII-D Facility has informed the construction of ITER and shown that there are ways that plasma can go unstable. There is a coil combination that prevents this.

The FES portfolio also includes an upgrade to the National Spherical Torus Experiment (NSTX) fusion reactor at Princeton to make a smaller central solenoid in the next year.

Within SC Nuclear Physics, there are three user facilities. One is the Continuous Electron Beam Accelerator Facility (CEBAF). The machine is going through an upgrade to 12GeV to increase the energy of proteins per cubic length. This seeks to understand the nature of protons and neutrons and how spin is distributed. There is debate in nuclear physics about the origins of spin.

The Facility for Rare Isotope Beams (FRIB) is able to make neutron-heavy nuclei. This is useful for nuclear physics and astrophysics. The postulation is that a supernova explosion will create neutron gas. This hits a cloud of other atoms and nuclei and generates very neutron-heavy nuclei.

The Nuclear Physics facilities are expensive to run. Among the existing facilities and FRIB and new work to begin there, budget considerations will have to be worked out.

Nuclear Physics is also taking responsibility for isotope supply and demand. SC is working with other agencies to synchronize those who need isotopes and supply. Dr. Brinkman proposed that industry may be able to develop facilities to produce isotopes.

SC's High Energy Physics is undergoing changes. Future direction is being planned along with understanding what the U.S. is doing. A recent study suggested building a large argon-based detector. It has been suggested that there is a need for ultra-pure argon to help collect data with wires that are meters apart and support the charged particles created to avoid decay away over the distance. Design costs presented in January 2012 are around $2B. Dr. Brinkman has asked Argonne to think of other approaches.

Within High Energy Physics, SC celebrated Saul Perlmutter, the 2011 Nobel Prize recipient in physics for his work on dark energy and the expanding universe. Work is also in process in Neutron Physics and work being done with China to support a neutron reactor near Hong Kong. High Energy Physics is also developing an experiment called LUX (Liquid Underground Xenon) at the Homestake Mine in South Dakota. The experiment will require the installation of a water tank and uses experimental work done in the 1970s.

Dr. Brinkman concluded by sharing that the Tevatron at Fermi National Accelerator Laboratory successfully completed its final run in September 2011. The data shows that there is a bump in the region of the spectrum. Similar findings came from the Large Hadron Collider (LHC). The LHC will be shutdown in all of 2013 for upgrades.

Discussion

Dr. Giles talked about leveraging all that the science community is producing and wonders to what degree Congress understands these advancements. Dr. Brinkman said there are many strong supporters who understand basic research's role. While this can vary, he has worked with committee members who demonstrate personal interest. Some have been on facility tours.

When asked about educational investments, Dr. Brinkman noted an error in using American Recovery and Reinvestment Act funds to propose a large graduate fellowship program. This encroached upon the National Science Foundation's mission. There is less resistance to programs that address specific communities. This effort will get back on track in the next year or two.

Dr. Brinkman addressed industry engagement, noting that the DOE talent pool is drawn into doing simulations for industry. He proposed an industry internship at DOE facilities so that private sector personnel can have improved capabilities and can be trained to do simulations themselves. Dr. Hitchcock responded that the Small Business Innovation Research Office and the Scientific Discovery through Advanced Computing (SciDAC) program are expanding talks with industry. He noted that small- and medium-sized businesses never have enough people to support simulation and modeling. A thought is to build an incentivized layer within industry for this work, as has been done in the seismic realm.

When asked for his perspective on the budget, Dr. Brinkman said that a fusion study to be completed in early April 2012 will determine budget direction. Some in Congress do not like the changes being considered and SC has to decide how to go forward. In nuclear physics, SC must work to figure out how to accommodate things being proposed.

Dr. Brinkman was asked about facility closures and the impact on the student and postdoctoral talent pool and maintaining their interest. SC is trying to keep programs balanced and ensure that there are as many students in the pipeline as SC can handle. The facility closure at MIT was unfortunate yet 300 students in fusion are still being funded. Dr. Brinkman commented that a bad thing is that ITER was building interest and excitement in this field and that move did not help.

Dr. Petzold asked about the interaction of large data and exascale computation science. Dr. Brinkman does not see any interference between the two. There is a question as to whether there will be an application where exascale accumulates large amounts of data. There do not need to be big investments on large data and those who have issues with this have addressed these problems. In the biological world, large data is a challenge and one that is not unfamiliar to this field. There are databases for genome sequencing, so a lot of money is not needed for this.

When asked about international collaboration with countries like China and areas of high interest such as fusion, Dr. Brinkman sees an opportunity to more actively engage others. At first, it was recognized that there was a real culture lag in the DOE laboratories. Previously, the attitude in high energy physics was to help China but now there is a more balanced program. He is going to China in a few weeks to negotiate various programs, but sees the need for care. Congress passed a bill that prevents NASA and the OSTP from talking with China at all, and Dr. Brinkman has to make an argument for engagement and defend the SC.

Public comment

None

Presentation: View from Germantown, Office of Science for ASCR

Dr. Daniel Hitchcock, Acting Director of ASCR, presented challenges and highlights from FY2013. The ASCR budget request for FY13 is $14.7M higher than FY12 and research support and facilities are down slightly.

The ASCR has an opening for Facilities Division Director. The position closes on April 9, 2012, and additional openings in computer science, computational mathematics, and facilities management are forthcoming.

Exascale computing is advancing and there is international recognition of this field. China has plans for large computers for military and industrial use, and to build better aircraft to avoid buying directly from the U.S.

The DOE is furthering a partnership with the National Nuclear Security Administration (NNSA). An MOU was signed on April 12, 2011, and DOE has supported joint meetings. A partnership between SciDAC and NNSA lead to computational work and facility use at the National Ignition Facility (NIF).

The need for ASCR to research ways to manage data is pushed by the move to exascale. Data movement requires power. If this and other operations succeed, the exascale investment can achieve broad effects for the SC. High-energy physics and other SC areas are creating great volumes of data. LCLS is an example and is a petabyte per year facility. Hardware changes are a challenge and experimentation is needed to determine that the computers the SC is considering are not just fiction. In moving to a data rich future, the way that scientists work has to change. Computing workflow must be figured out to make data useful to science.

The President's Council of Advisors on Science and Technology (PCAST) recognize the importance of this work. ASCR facilities give an opportunity to test new ideas, to make applications real, and to avoid conflict with research, as both are needed. ASCR is part of the SC and is working to impact the Office to enable scientific discovery.

Leadership Computing Facilities (LCF) are expanding. Oak Ridge LCF (OLCF) is upgrading and hopes to achieve 10 to 20 PF by November. Argonne LCF (ALCF) will upgrade to 10 PF and the IBM BlueGene/Q by this summer. LCFs will also allow more hours to people next year.

Dr. Hitchcock announced that the INCITE proposal writing process will open soon. Proposal writing webinars have been conducted (See www.olcf.orl.ogv/event/2013-incite-proposal-writing-webinar). The effort will award compute time on the Cray XK6 (Titan) system at the OLCF and the IBM Blue Gene/P (Intrepid) and IBM Blue Gene/Q (Mira) systems at the ALCF. Awards will be based on peer-review.

Work by researchers at LBNL was acknowledged. The Jaguar supercomputer at OLCF helps capture sunlight. If researchers can get nanoscience-based solar-based detectors to be more efficient, then more solar cells can be produced and made more available for mass use.

Dr. Hitchcock echoed Dr. Brinkman's message about industry engagement, specifically highlighting work with GE. Work with INCITE has helped them produce 15 percent greater fuel efficiency. Boeing is working with the ASCR to make landing gear quieter during landings.

The National Energy Research Scientific Computing Facility (NERSC) will move to the LBNL campus in 2013. A power upgrade is being planned. Some costs are reduced through LBNL's discovery that winds coming across the bay are cool enough to run the facility's chillers. NERSC will also use 14 Kv and less copper. This comes after learning that using 48 Kv at the Lawrence Berkeley National Laboratory (LBNL) would not work.

Important examples of NERSC-supported work are a study of solar materials and understanding how these materials work, and the use of compressed sensing in mathematics. NERSC used this to learn when things turn on and off in DNA. This does not require understanding everything that happens with folding. This combines mathematics with computing and different disciplines to give new ways of looking at these problems.

Dr. Hitchcock described the ESnet upgrade to achieve 100 gigabits per wavelength service. Several states have bought equipment and hooked up to Internet 2. With reasonably moderated and in-ground fiber, 10 times the data can be carried by simply using different features. A demonstration at Supercomputing 2011 showed data visualization at 10Gpbs and at 100 Gbps.

One could perform visualizations and get this much detail to users without them being there. NERSC and ESnet were essential in solving the puzzle of the neutrino.

Moving to research and evaluation prototypes, Dr. Hitchcock described this as ASCR's way of engaging industry. Despite a reduction in FY13 from $30M in FY12 to $22.5M, work here will enable researchers' use of next generation computers. Funding is set for Blue Gene architecture at LLNL and ASCR will do more with NNSA and move toward exsascale computing.

ASCR is continuing to enable the competitiveness of the U.S. IT sector. Dr. Hitchcock pointed out a number of companies that are using software developed by ASCR. Specifically, without the MPICH message passing library, the design for the F135 jet engine would not have been possible. Many of ASCR's contributions have lead to entirely new industries.

In the area of applied mathematics, new investments are planned that will integrate things and help to deal with large data. The root problems are how to envision high-dimensional data and how to effectively move it around. ASCR can work with others to create an end-to-end solution and make investments across its portfolio to address these bottlenecks.

The DOE Computational Science Graduate Fellowship Program is subject to discussion by the administration about the role of such investments in the mission agencies. Changes to the 2013 budget can mitigate any risks should the administration decide that mission agencies should not support these programs. ASCR believes that this type of program is important and decoupling it from the mission would be unfortunate. Actions have been taken to ensure that Fellows would not be terminated now but that may start in 2013.

Investments within Computer Science and in exascale are proceeding. While building an exascale computer may be optional, managing energy use with building cores is not. The hardware is advancing and ASCR does not have a choice in that regard.

Within the area of Computational Partnerships, SciDAC has been reengineered. Moving to the next generation of computers, complexity is being allocated from applications to hardware and this will make computer science easier. ASCR is working with four SciDAC codesign centers to identify needed applications and hardware configurations. The centers can respond to prior SciDAC concerns about who to talk to. Centers are responding to partners and helping manage data generation. They are also looking to the math community to identify useful things.

ASCR has changed the way it partners with other programs. This evolution of SciDAC used discussions with other associate directors to learn of their interests and what would help their programs in the next three to five years. A joint solicitation issued at end of the last fiscal year has lead to five FOAs and proposals out for joint review. A solicitation is out in Basic Science and that will be co-funded with ASCR to leverage other programs. High Energy Physics is focused on accelerator design. A barrier to next generation facilities is the number of volts one can achieve in the next facility; having shorter accelerators is important.

Going forward, SciDAC will inform future research through applications and delivery. SciDAC has discussed with industry ways to get their people to be more HPC literate in order to understand opportunities. One company that has done this the best is Proctor and Gamble.

Work is continuing in the Next Generation Networking for Science area. A solicitation is out for the distribution of data-intensive science. This represents tactics to deal with the challenges brought about by increased bandwidth.

ASCR is supporting a range of workshops. In particular, the high energy physics workshop in May 2012 will discuss what to do with high energy event software. Currently, they send an event to the core, wait for an answer, and it goes back into a queue when done. Starting to implement this on a new machine with 100 cores per processor will prompt the need to think about what to do in future. There is also modeling and simulation work being done with partners at the NNSA. There is a need for appropriate language and planning for exascale computing. This is the basis for a joint plan from NNSA and SC that is going through the concurrence process.

Discussion

Dr. Giles asked about delivery of the joint exascale report, noting that the ASCAC would be eager to comment on this and possibly have a Web-based or phone discussions before its next meeting in August. Dr. Hitchcock reported that the goal was to have this available to Congress by March 20, 2012. Concurrence has taken longer than expected, yet he believes the report will be available in a few weeks. He agreed that it would be nice for ASCAC to have this discussion and that SC can figure out how to do that.
Dr. Dongarra noted that exascale was originally planned for 2018 or 2020, but wonders about this timeline due to budget cuts. Dr. Hitchcock responded that there are things that will go out soon that will determine what is possible. Predicting the Federal budget outlook and the next 10 years of this field is challenging. Starting with less funding adds risk to the scope and schedule. A current dilemma is how to drive down energy consumption, and consideration of how to trade off paying for the investment now and pay off energy costs later. By August 2012, SC will know more about this. In the run-up to exascale, he believes that networks will not get much bigger as fiber power plants can connect things such as 700 cabinets. Most exascale challenges will occur in a cabinet and in a node. The rate of increase in power consumption is challenging. Google spends $300M per year on electrical power. It is possible to think about a 300mw computer but that is an unattractive design point.

In response to Dr. Blumenthal’s question about uniting people who are working in the SC, Dr. Hitchcock noted that PI meetings accomplish this by engaging up to 300 people. Other meetings are more open to new ideas and people at different times of the year. This summer, more data intensive meetings will examine different techniques to meet that type of challenge.

When asked about SciDAC’s further discussions and progress in light of budget problems, Dr. Hitchcock shared that progress is deeply discussed with partners on a monthly and sometimes weekly basis. The group is trying to winnow down challenges to generate focus on vital issues that also meet everyone’s schedules.

Dr. Hitchcock responded to interest in analogous exascale efforts going on across DOE, sharing that there are discussions with applied programs to improve the coupling of efforts. For instance, energy is meeting with applied mathematics and applied energy delivery personnel. Discussions have shown that all have complex systems and lack ideas on the underlying set of mathematics. It is important to determine how the pieces fit together, and to include wind, building and combustion experts to educate them and determine their needs and impediments that could prevent working together.

Dr. Tang inquired about international engagement the cross-fertilization of ideas. There is extensive discussion with the EU, said Dr. Hitchcock. Dr. Harrod is going to Japan to examine opportunities, especially on the software side. China is a more complicated issue but there have been discussions. SC is working out IP issues to find feasible approaches and balance.

Public comment

Mr. Jack Wells expressed concern about the “data intensive science” theme and wondered if this should be framed as a policy question about who owns the data. In an operational model, user scientists receive the data but it seems as though there is a limit on the level of analysis with large data sets. Dr. Hitchcock responded that SC has a data working group that is addressing policy issues and data ownership. Storing data and assuming that it is available for later analysis requires understanding of the accompanying metadata and what should be kept to compare results in an honest way. The SC’s Laura Biven chair’s the SC working group on digital data. Mr. Dave Turek of IBM commented that an explicit target date is needed to establish the exascale initiative relative to budget fluctuation. He also offered to give an industry viewpoint on exascale and its relationship to data.

Presentation: Understanding Climate Change: A Data-Driven Approach

Dr. Alok Choudhary and Dr. Nagiza Samatova reported on joint research supported by the Office of Computing Research. Work is now being leveraged by the National Science Foundation. The project is exploring a fourth paradigm of data-driven science that drives use cases and the extraction of information for decision-making, policy, health care and other applications. This project is analyzing data collected over many years to better predict extreme events such as weather-influence meningitis outbreaks, and hurricane and cyclone events.

Climate system complexity results from factors such as temporal-spatial scales, the range of time, and the number of variables. The challenge is connecting elements in complex systems to gain knowledge from data. Processing involves things such as eliminating certain correlations and understanding others. The next step is building networks to work with decadal and other data to help identify communities of data that can be used to predict extreme events. The algorithms used are applicable across many domains such as biology and social media applications.

An analysis of changes to the Sahara is a product of this work. Adjustments are made to account for seasonality, and stable clusters of data lead to “climate communities.” These show changes over time.

This application can help define fundamental computer science research and enable the understanding of problems in computer science. An example is relationship mining wherein there are few algorithms that adequately address this. Predictive modeling is another area that depends on relationship mining to help model typical and extreme behavior. This leads to HPC and achieving efficient analytics on future generation exascale HPC platforms with complex memory hierarchies. New algorithms are needed that can evolve to new systems.

The Global Cloud Resolving Model is an example of a new system. The simulation produces 1.4 petabytes of data and the fundamental techniques have helped reduce I/O modeling efforts.

Climate science is an application that can be enabled by a data-driven approach that leverages HPC. Components are high spatial or temporal resolution, higher data dimensionality, greater complexity per data point, and a shorter response time. There is a definite case and the case is not parallel to other things but is composed computation. There are also examples in different types of mathematics such as building highly-optimized, scalable kernels.

Along with developing these techniques, the movement of data, data amounts, and energy consumption must be considered. Approximations using power-aware analytics allow for a reduction in energy use and the complexity of energy use, and optimization efforts to avoid errors. This gives a different way of thinking about analytics.

Dr. Samatova presented examples and results of the application. Multiple connections make climate systems complex. The challenge is how to discover system components that cross and to define the functional response of these variables coming together. An example is the connections that are formed with El Nino.

Dr. Choudhary and Dr. Samatova's work has pointed out rainfall anomalies that can lead to meningitis outbreaks in Africa. This results from the complexity of conditions in the Northern Atlantic in its cold phase and the Southern Atlantic in its warm phase. Another example is hurricanes that start in the Western Sahara and impact North America. The interconnection of these factors in the West African climate has been studied for more than 25 years. Dr. Choudhary and Dr. Samatova seek to show how these pieces link, and to identify new pieces and how they can magnify or suppress system response.

One outcome could be theorizing and forecasting hurricanes, made possible through collaboration in the machine-learning and data-mining communities. This research is modeling systems as complex climate networks, and modeling systems when in extreme and normal phases. The work also looks at connections between geographic locations and other phenomena to understand the changing and dynamic nature of climate systems. Data mining will lead to machine-based models to enable forecasting and the ability to push the forecast accuracy and predictability to up to 90 percent and determine both the start and end-state of hurricanes. An outcome could also be predicting hurricane activity 10 to 15 days in advance of the end-state.

Specific to modeling and complex systems, biology uses hierarchical modularity in which a system is decomposed into smaller subsystems. Machine-learning inspires the use of modularity and gives understanding of the relationship features within each model. This goes one step further to understand the cross-talk between systems, and when mining these networks, shows how they are cross-talking and the existence of causality relationships.

State-of-the-art groups are using cross-fertilization to achieve 64 percent accuracy using regression modeling. A Georgia Tech team uses hybrid modeling to achieve a 65.5 percent rate.

This data-driven approach will help forge foundational inter-relationships in climate systems in fundamental ways. Discovering knowledge from massive amounts of data is the next HPC frontier.

Discussion

This research has examined other applications, starting with biology and then climate. It has also been used in medicine to predict mortality indices using data from thousands of patients. A next step is to bring in genomic data and other types of data. There are also graphic algorithms to analyze massive networks from online social networks. Dr. Choudhary and Dr. Samatova want to work with scientists in these domains.

When asked by Dr. Chen about getting kernels at the right level of granularity to work in-situ and with solvers, Drs. Choudhary and Samatova see these as two aspects not yet understood. In the climate context, 20 years worth of data is combined with real-time data. A lot depends on the complexity of the analytics. Running data into systems and scaling algorithms is challenging.

Dr. Giles commented that network analytics seem to work well for predictive work, but wondered if it is working better that Dr. Choudhary and Dr. Samatova thought it would. He also asked if there are circumstances where it would perform well and those where it will not, and how prediction reliability is understood. Dr. Choudhary noted that network construction is not a problem. Starting with exploratory analytics brings in another dimension of HPC. This makes the situation more complex. By doing various comparative studies using varying statistical methods, this can constrain the space. Dr. Choudhary and Dr. Samatova identify some plausible seed points from which to start to mine networks, identify the constraints of the search space, and have observed that there are new ways to construct the climate indices.

Prompted by Dr. Hack’s question about bringing this to deterministic models, it was noted that there should be a dialogue between model developers and knowing what network matrices mean and why some physical models are not as good as they should be. Dr. Choudhary also wondered about shifting from modeling to prediction. Data is plentiful and there is a nice predictive model that is better than the statistical methods of other groups, but it is still not as good as needed.

In response to Dr. Tang’s question about developing new local I/O capabilities that would generate high impacts, it was noted that this would be important for many reasons. If given the option from an analytics perspective of having more of that by fewer flops, Dr. Choudhary would pick more storage for many reasons. For instance, with this type of algorithm the analysis does not generate much locality. This allows more transactions and that would be important. Data should be kept near to where it is being computed and it can have a multiplicative effect. Dr. Samatova commented that some are more worried about getting data out of the machine versus how climatologists view data. Things need to be optimized for heterogeneous reading and from an in-situ view the level of preparatory analytics that can be done when data is available.

Public comment

None

Presentation: Density Functional Theory and Renormalization Techniques in Condensed Matter

Dr. Alejandro Lopez-Bezanilla described the research he is conducting as an American Recovery and Reinvestment Act (ARRA) Postdoctoral Associate Fellow at Oak Ridge National Laboratory's Center for Nanophase Materials Science. The Center supports computational and experimental studies to explore structure-function-transport relations of materials.

A challenge in computational science is coupling the atomic world with measurements done in materials experimentation. Changes in the quantum, nanoscopic and microscopic worlds are transferring to the atomic world, as are advances in measurement capabilities. Modeling is enabling the ability to predict, validate and innovate. Some examples are work in carbon nanotubes and the integration of atoms, engineering graphene, identifying new features in boron nitride, and spectroscopy modeling to improve technologies for nanotechnology research.

An area showing the integration of this is the modification of bio and photo sensors at the atomic level and understanding physics at the quantum level. An example is a double-walled nanotube. Molecules can be placed in the outer wall. The voltage is increased between the two and a range of current is observable because the inner wall allows for conduction. The group testing this may not understand what is happening to the inner tube. Dr. Lopez-Bezanilla's work simulates the movement of the electron and the probability transmission case with a certain number of molecules attached on the surface. The capabilities of the inner wall are not precisely observed. These results need to be coupled with microscopic measurements which will require a longer nanotube.

Two steps are involved. The first is building an efficient tool for electronic structure calculations. The second is enacting the transmission and reflexion probabilities of the electron. The computational tool used is the SIESTA scaling method.

The mathematical approach uses a system divided into three parts. First, electrons enter from the right and left sides with a central channel. There are coupling matrices. In principle, the right and left leads are nearly infinite. The center has hundreds of thousands of atoms. Using decimation to achieve a three-site model allows for changing the complex system into something more achievable and of a reduced size. This allows moving coupling from one system to the next to build a longer system. Building this linear approach allows shifting from a finite set of content that can be put into the computer and calculated. Transport formalism uses Green's Function for implementation. The technique that defines this code allows for sending this calculation to one of the nodes.

The power of this technique can be applied to the arrangement of hybridized boron nitride and graphene domains and it conducts electricity. With a lot of boron nitrate, the current is zero. The goal is to give further insight on the mechanism that allows the prediction of graphene in the boron nitride application.

Thousands of computations are needed to deal with impurities in the systems and to deal with how they scatter to understand how electrons go through a channel. For the scatter, one to two thousand configurations are needed over 20 different systems. This is fed into Jaguar, and has been done with simple models at this point. However, exascale resources will allow for generating precise data to achieve more accurate descriptions.

In materials by design, Dr. Lopez-Bezanilla's work has lead to the development of a plasticized material that was then made metallic by placing electrons along the edges. This process - VASP optimization - has brought about new features in the material and it is more useful. This code is now 10 times faster and the valuation will be available shortly. An outcome is learning how to deal with codes to predict behaviors in next generation materials.

Experimentalists are also trying to measure the nanoscopic world and that they need insight to improve their capabilities. This can come through Dr. Lopez-Bezanilla's and others' work with polymers.

Work will continue on parallelization methods and techniques to use new hybrid BN-C nanostructures. Dr. Lopez-Bezanilla's work will also calculate Green's Function with more powerful techniques. He is considering implementing techniques based on new codes that use CPU/GPUs, and doing other things to achieve a more accurate description of the transportation system.

Discussion

Dr. Chan asked about the boundary conditions in the use of graphene ribbons used in simulations. Dr. Lopez-Bezanilla explained that when a piece of a system is long enough, at the end of a ribbon an impurity is not seen in the neighboring cell. The impurities do not have an interaction. For state-of-the-art graphene, additional dimensions are unnecessary.

In response to Dr. Giles question about geometries and the methods that will expand upon current methods, Dr. Lopez-Bezanilla responded that when calculating the electronic structure of a system, certain conditions are not important. Movement to other systems allows for calculating with worrying about the conditions.

Public comment:

None

Presentation: Update on ASCR Small Business Innovate Research

Dr. Rich Carlson, Program Manager in ASCR, described the Small Business Innovative Research / Small Business Technology Transfer (SBIR/STTR) that operates across the SC and other agencies. The program works differently at each agency. Direction for DOE is set by ASCR. Funding in FY12 totaled $2.59M. Funds have come from a Congressional set-aside and last year Congress added more funding. The total for SBIR/STTR was $12.4M in FY11 and will grow to $16.1M in FY17 assuming a flat budget. ASCR wants to leverage this program.

Small businesses are eligible based on a Congressional definition and can receive R&D support. The program is split into three phases. Phase one covers science and engineering topics, phase two is for prototype development, and phase three is for commercialization. Those funded in an earlier phase can apply for the next phase and more funds. PIs from research institutions can qualify which allows labs and institutional PIs to participate.

Throughout ASCR, all program groups have received proposals. Two years ago, approximately 230 proposals were received. DOE had difficulty tracking this quantity and the impacts. ASCR has moved to technology transfer to transition research to commercial applications. Things that fit well are network operations and services, and HPC usage by engineering and manufacturing communities to help them leverage DOE research and advance their own product development.

A major theme in 2012 is advanced network technologies and services. This can involve researchers doing advanced networking, and those managing larger networks such as commodity networks or local service providers who can take advantage of DOE’s operation of networks. Those involved in hardware technologies for optical networking are also of interest.

Another theme is HPC and exploring turnkey solutions or leveraging codes that have been developed for years. An outcome is that an engineer could leverage HPC resources without having to learn and install these things from scratch. The HPC community was ready for this theme. One ASCR staff member started reaching out to this community and found this to be appealing to multiple consortia. The groups gave input on the tools and things that would enable engineering advances. Two articles on this appeared in the Digital Manufacturing Report.

Eighty proposals were received in 2012 with 56 in HPC for manufacturing and engineering. Fifty reviewers helped with proposal evaluation and identified 25 proposals for funding at $3.5M. Six proposals were funded from the networking side with19 from HPC.

ASCR plans to do this again and reach the same meetings as in 2011. Outreach efforts will emphasis what ASCR was able to accomplish through community input. ASCR is also considering partnerships with Applied Program offices.

Topic areas are expanding based on emerging needs. One area is videoconferencing and seeing if this is a tool that needs to be improved for the science community. There may be a desire to avoid using a single vendor product and to identify ways to work on a single desktop with an integrated solution across the scientific communities. ESnet is a foundation for this capability, so ASCR is looking to them as experts and with a focus on the impacts that small businesses could make. ASCR is also looking at software as a service model to start selling products and tools that DOE has developed. On the HPC side, ASCR is looking at 3-D and additive manufacturing.

ASCR will continue to engage communities and convey that ASCR is not the end consumer. Laboratories, universities and contractors who run network infrastructures are the ones who need the tools that would be developed and provide a basis for this program. Planned outreach for network operations includes engaging the HPC community and increasing the awareness of the network tools topic. Greg Bell is integrating ASCR's message into a slide set to enable the ASCAC and others to share this message with communities. Dr. Carlson asked the ASCAC members to engage their communities and identify tools and services needed to help manage and operate networks, and to engage independent software vendors and small businesses.

Discussion

Dr. Carlson described the SBIR/STTR phases. Phase one is a new call and solicitation and generates a large response. Those successful in fulfilling phase one can submit for phase two. Most of the funding supports phase two and awards of $1M. Phase three does not offer funding, but connects the prototype and developers to the venture capital and private sector communities to produce a commercial product. ASCR has not monitored phase three activities, but the Federal SBIR program office does track this and is pushing for concept commercialization.

Dr. Blumenthal asked about the involvement of venture capital, if there is repeat interest, and if there is more is being done work with small businesses. Dr. Carlson noted that most awards are for some small companies that can write very good SBIR proposals. The program office has a new division director who is focused on commercialization. This can also be a threshold for rejecting proposals. Each SBIR program office is responsible for reaching out to the venture capital side. ASCR is working to ensure that proposals are technically sound.

Dr. Carlson described ASCR's engagement with awardees. An ASCR program manager is assigned to the program. Kick-off meetings are held by teleconference, leading to working relationships with the PIs. Managers are invited to PI meetings and sought for participation in other venues.

Public comment

None

Presentation: Extreme-Scale Solvers Workshop

Dr. Karen Pao described the Extreme-Scale Solvers Workshop held on March 8 – 9, 2012.[1] The workshop sought to determine research areas for extreme-scale algorithms and software to effectively use the 100 PF systems and prepare for exascale systems. Approximately 40 researchers attended from DOE labs, universities, industry, and ASCR. Most were self-identified as numerical algorithm designers, involved in numerical methods deployment in HPC, and as HPC numerical library developers. A draft workshop report will be available on April 10th.

[1] http://orau.gov/extremesolvers2012/

There are preparations for a paradigm shift that moves to 100 ~ 300 PF systems and ~15MW systems sometime around 2020. A few other machines are needed prior to the shift. These challenges have been apparent for some time and were cited at a 2010 workshop that resulted in the report “Scientific Grand Challenges: Cross-Cutting Technology for Computing at the Exascale.” This exposed the need to study numerical analysis issues moving away from bulk-synchronous programming and developing tools to aid understanding energy and data use.

The workshop objective was to “determine research areas needed for extreme-scale algorithms and software to utilize effectively the 100PF systems and prepare for the exascale systems.” The 100 PF machine was chosen in order to have some idea of what a machine might look like. A future machine will likely be a natural evolution of today’s machines. It is expected to be heterogeneous with CPUs and GPUs in order to deal with power constraints. Immediate solutions are helpful yet the emphasis is not on a “Band-Aid” but rather an impactful solution with relevance and impact beyond the machine.

New algorithms are needed to deal with extreme parallelism, data movement, resilience, and potentially, heterogeneity. There is also an opportunity for influencing the design of future extreme-scale computers.

The opening talk given by Bob Lucas from the University of Southern California Information Sciences Institute detailed challenges inherent at the extreme scale. In addressing resiliency, he offered that there is an economic argument to addressing this as more hardware is developed. Solvers will have to change what they do to help achieve greater use of code.

Bill Gropp from the University of Illinois talked about rethinking due to exascale directions that include power constraints, the detection and management of faults, and the architectural forces that force a confrontation with architectural realities.

Mathematicians were asked to consider solver issues from different angles. People expressed worry about managing memory hierarchy and dealing with cache versus scratchpad memory. Concern was also expressed about communicating new architectural features as soon as possible. There is also a focus on algorithmic R&D needs, including the tools needed to implement solutions and ways to control energy usage, including finding hardware and allowing mathematics a way to cycle down and operate at the right speed. There seems to be a sense that we are losing control of the ecosystem. The workshop also discussed transition strategies and evolutionary versus revolutionary algorithmic solutions.

Panelists discussed game-changers including stochastic algorithms that could help deal with non-determinist behavior. They also discussed parallel in-time algorithms, and asked about external capabilities needed for extreme-scale solvers.

Need was expressed for a compiler tool so that an algorithm designer would not have to hand-code things and if there was a compiler tool that could make that work faster. Attendees were also asked to think about what mathematicians and computer science people should know beyond their usual knowledge bases to prepare for the era of extreme computing.

Discussion

Dr. Giles asked if there were any other revolutionary algorithms discussed beyond stochastic algorithms. Dr. Pao commented that revolutionary ideas are insufficiently imaginative. Stochastic algorithms are an idea that deals with resilience and that is something that still needs to be solved. The community is starting to think about this. The workshop committee is also analyzing other ideas. There are many interesting ideas including variable arithmetic.

Dr. Tang wondered about applications in the domain community. He is concerned about solutions for PF problems for machines already in existence in China, Japan and the one coming online at Oak Ridge. Dr. Tang asked if the workshop discussed a timeline or roadmap of things available today for the most advanced computing problems, particularly as the applications community relies on the solvers community. Dr. Pao noted that the SciDAC Institute's mission is to bring applications into the petascale era. SciDAC can officially use today's petascale machines. The workshop did not talk about this but it needs to advance. There is a lot of work involved in working with the people who make things work as shown by earlier data movement efforts. Dr. Hitchcock added that in the case of the ASCR facilities, they have carefully selected applications for two machines that capture the height of algorithms to ensure that there are guidelines for these machines.

Dr. Pao confirmed for Dr. Chapman that there was discussion on programming models and that this was addressed from a mathematician's view. People want to test new models, and wondered if they have an algorithm how they will know if what they are doing is right. Dr. Pao noted that earlier in the ASCAC meeting someone wondered if getting to 20 percent would ensure that what they are doing is right. She feels that this is being structure specific, looking at the data structure and the programming model. There was discussion about the programming model.

Dr. Chapman asked how a compiler might help. Dr. Pao noted that from a mathematician's view, there are many tools for performance metrics. There could be ways to design or benchmark an algorithm. The goal is to describe the requirements. Only one person in the workshop knew about compilers.

Dr. Pao hopes to have an FOA based on the workshop and requirements around May 2012. The funding level is not yet determined but will be around $4M.

Public comment

None

Presentation: A New Paradigm for Applied Mathematics

Dr. Sandy Landsberg, ASCR, gave a presentation on Applied Mathematics. This program does most of its work in partial differential equations (35 percent), optimization (15 percent) and UQ and stochastic systems (15 percent). The program involves supports approximately 110 projects at $40M per year. The program is thinking about increased fidelity, new multi-scale and multiphysics models, uncertainty quantification, the novel analysis of algorithms for streaming data, and mathematical challenges for dealing with new architectures.

Projects range in size from single and two-PI teams at funding of less than $250K per year to large teams or more than five researchers at more than $1M per year. This approach has been successful and in operation for about 50 years.

The new paradigm will support R&D for applied math models, methods and algorithms for understanding natural and engineered systems related to DOE’s mission. Currently this is being examined holistically then will look at challenges from their start and the problems to be solved.

While individual projects will be maintained, the new paradigm will be represented by the Mathematical Multifaceted Integrated Capability Centers (MMICC). This will enhance the impact of applied math on the DOE mission five and ten years out. The intent is to bring projects together on a large scale. The Centers will look at DOE-relevant systems for scientific discovery, design, optimization and risk assessment. There is also discussion about supporting high-risk and high-payoff projects. The structure will demonstrate clear relevance and impact to DOE mission.

A workshop in September 2011 demonstrated the need for new methods and solving complex systems holistically. An applied mathematics summit held in March 2012 discussed the themes for the MMICCs. Summit attendees discussed narrowing the themes to be able to focus on potential and impact, and to focus on areas ripe for investment. Topical areas are materials and chemistry for energy application, complex engineering systems such as the power grid with a focus on real-time and predictive capabilities, fluid-structures interaction with carefully chosen illustrative examples, and SC facilities and the math-related challenges found there such as how experiments are designed and integrated mathematically.

There were questions about how the Centers will be received in the math community and how the concept should be socialized. There is also a challenge in getting applied math researchers to work collaboratively. Early success is important and measures of success and next steps are being developed.

Summit attendees indicated that the MMICCs are a good idea, but long-term cross-cutting mathematics project are the mainstay of success for the DOE applied math program. The MMICC concept should let applicants define a grand challenge. For example, allowing the definition of one or more DOE-relevant science and engineering challenges that are abstracted into a set of math research challenges that must be addressed through a multi-faceted, integrated and iterative process. “Integrated” proposes looking at multiple facets together, while “iterative” means looking at the results repeatedly to find solutions that is optimized for the problem.

According to attendees, the Centers should be distinguished from areas in which there is already significant investment. These include fusion, nuclear energy, climate, computational fluid dynamics, and nuclear weapons and design code. This is so that mathematical aspects of these areas can be impacted.

MMICCs may consist of up to seven staff with postdoctoral researchers and will draw upon the individual projects. ASCR also believes that identifying effective center directors is critical and wants these individuals to devote a full amount of time to their respective facilities. Junior staff is important to support collaboration and a flexible structure is important. Some will have strong and loose collaborations. ASCR will socialize this with the applied math community and brief key groups to enable collaboration within this community.

In addition, the MMICCs will connect with DOE mission areas, and could fit with SciDAC partnership applications and exascale co-design centers. Where possible, these relationships may result in high-risk and high-payoff projects. This collaboration will be an enabler for the entire math program. No other agency has done something like this on such as scale.

The timing for this is the conclusion and continuation of other projects. Fifteen projects in multi-scale math that embraced math for complex systems will end in August 2012. And about 29 projects that are solving aspects of a much larger project have been going on for the seven years.

There will be a solicitation coming out in new few months and a pre-application process.

Discussion

Dr. Chan asked about coupling Centers with areas of collaboration and what the metric of success will be. Dr. Landsberg shared that the Applied Math Program is not looking for a robust system or robust computer suite but linkages between different modules. One can expect prototype codes and simulation results and integrated analyses to solve problems. In the long-term, the applications office will seek to adapt methods that have not previously been adapted.

Dr. Dongarra voiced concern that single PI awards would not fit in this model and that they have to be a part of the Centers. Dr. Landsberg clarified that small research teams would be maintained and there will still be room for traditional researchers. The Centers will take up less than 25 percent of the entire budget and are open to all. Teams will have mathematicians, computer scientists and domain scientists, and could be lead by universities, national laboratories, and industry.

Math innovation does not happen in a center, Dr. Berger commented, but by a mathematician at their desk. She is concerned that the individual PIs and young people who are not tied to senior people will be excluded. Dr. Landsberg explained that ASCR will grow the entire program while encouraging Centers to reach out to young people and make them aware of Center activities.

Dr. Tang commented that optimization is improved when algorithms move forward, but a common challenge is that advances in dealing with low memory core systems and local systems are needed. He asked if the better path is to work through a math center or do more with innovative new ideas and then include people. Dr. Landsberg conveyed that once the Centers are established, their focus areas will be known. If there is a gap, ASCR will step in. The Centers will have increased opportunity for younger scientists. If optimization method is not addressed in a cross-cutting project, she hopes it will be dealt with in the Centers. Care is needed to ensure that there is no duplication of efforts. Dr. Pao and Dr. Tang clarified that he is talking about introducing totally new algorithms. Dr. Landsberg responded with the hope that the Centers will achieve that and computing for new optimization in sum.

Dr. Chen wondered if picking new areas that have not been covered before will help cover new spaces of math motifs. Dr. Landsberg sees a combination of existing strong math capabilities and efforts to pull in and do data analysis and simulation. She believes that this will broaden the math portfolio but also leverage work that ASCR has supported over the past 10 years.

Dr. Landsberg explained that the Centers will have an impact for 10 years with support from ASCR for five years. The Centers could also have some core that is all co-located and could have broader geographic collaborations. Still, there is a lot to be said for co-location.

Dr. Giles asked where the ASCAC might be able to express more thoughts about the Centers. Dr. Hitchcock replied that individual comments could be sent by email to Dr. Landsberg. He stated that there might be a way to charge the ASCAC to look at the Centers' progress. A major challenge is knowing how to abstract the application domains that achieve mathematically applicable things such as fluid structures, or solving mathematical structures to achieve 787 wing design but not solving the wing design directly. He shared that there will be an annual review process and more frequent reviews to ensure that this is a path that ASCR wants to be on. This process could generate ASCAC feedback. Dr. Giles noted that the next COV in the annual rotation will be the mathematics program.

Public comment

None

Presentation: Scientific Collaboration for Extreme-Scale Science Workshop

Dr. Rich Carlson summarized the "Scientific Collaboration for Extreme-Scale Science Workshop" held on in December 6 – 7, 2011.[2] The workshop charge was to determine where to take large scale extreme science systems to make this useful for supporting research teams in the next decade. The workshop was Co-Chaired by Richard Mount of the SLAC National Accelerator Laboratory and David Skinner of LBNL.

Attendees identified critical science drivers, challenges, technical challenges, and determined the state-of-the-art to make sure that research is headed down a new pathway. They also looked at the exascale series of workshops to make machines useful and usable by a large portion of the user community.

Forty individuals on the organizing committee worked for three months to design the meeting. Fifty-four co-authors developed the workshop report. The report lists 12 different communities that were engaged in the workshop.

Major findings and recommendations from the workshop report are the following:

[2] https://indico.bnl.gov/getFile.py/access?resId=0&materialId=1&confId=403

- Unimpeded collaboration accelerates and empowers extreme-scale science
- Impediments to collaboration for extreme-scale science can be readily identified
- Removing the impediments and empowering collaboration requires advances in several areas

Workshop breakouts covered technologies for processes, data and teams. Attendees described how teams need to work together in terms of technologies for teams, data and processing.

One outcome was that even after decade of work there are a lot of primitive collaboration systems that hinder discovery. There are primary collaboration tools, but a need to move to digital collaboration systems to accelerate sharing systems to advance learning and make use of resources and devices that allow people to work together and get people to work together. A large scale program should be able to provide these types of systems to allow for large-scale collaboration to exist. A model for this is an application store for the iPhone.

Four major abstractions evolved from the workshop:

- Discovery – Things needs to be discoverable to advance the field and science
- Centrality – There is a need to centralize things to know if it will work and avoid making people install things
- Potable – Understand what works on a variety of machines and in different environments
- Connectivity – We have to know the origins of information and ensure that the metadata that goes with it is useful

The SC communities have a need for collaborative tools and services and a strong desire for these things to appear. The Next-Generation Networks for Science (NGNS) program in ASCR needs to continue to do this and work toward these solutions to develop a base research program and to make advances in the coming five to 10 years. Strategies are needed to transition NGNS research into science communities for longer-term sustainability. SciDAC could be used to transfer this into other domains.

Dr. Carlson shared that there is an existing solicitation to build this, and closed by thanking Drs. Chen, Tang and Williams for serving on the Workshop Committee.

Discussion

Dr. Hack asked about the comment that workshop outcomes can be built upon once the research component matures. Dr. Carlson replied that a long-term strategy is needed to keep these aims going. The workshop report can ensure that this point will be maintained.

Dr. Giles noted that in the context of the workshop, collaboration is between people. Dr. Carlson clarified that any scientific collaboration could be between people, machines, and agents, and that all fit in to this system. Dr. Giles shared that he was thinking that it is seemingly very hard to link users when building useful technologies. He wondered about using a layer of application developers who could contribute to this effort in its final stages. Dr. Carlson noted that attendees were specifically looking at this and how scientists of the future will work on these aspects.

Public comment

None

Overall public comment

None

The meeting was adjourned at 4:32 p.m. EST.

Wednesday, November 2, 2011

The Board Chair Dr. Giles started the session at 9:06 a.m. EST.

Presentation: Report on Inter-Agency Resiliency Workshop

Dr. Lucy Nowell of the ASCR summarized the "Inter-Agency Resiliency Workshop**"** held on February 21 – 24, 2012, in Catonsville, MD.[3] Participants were from the Federal government and included national laboratory representatives, as well as people from industry and academia.

The workshop addressed applications in terms of what has to be done to ensure an application will continue running through the extensive fault rate expected with an exascale platform. Researchers are trying to establish timelines for research. The most critical near-time research is the understanding of fault types and the frequency of occurrence. Currently this is unclear. Prediction in the long-term can mitigate faults before they occur, but prediction is very difficult. Attendees classified the types of R&D needed, supported by major investments and smaller research projects. A report will be issued in several weeks that will describe the workshop.

The group agreed that capabilities are needed to allow a particular application to be resilient through about 2015. They identified four areas of capability, specific tasks, and the kinds of research to be conducted. From a list of identified gaps, the group chose the top 10 priorities.

Attendees identified the target audiences for whom this work would matter most, and how to correctly do fault management. Work will be required for developers of applications to appropriately handle faults. The group tried to figure out who cares about resilience now. It seems that applications and checkpoint restart is being solved now. Developers do not care because they think it should be someone else's problem. People who develop runtime systems want the problem solvers to be vendors, yet vendors do not see this as their problem. No single group was identified as being responsible for the problem and there is currently no crisis.

A question posed was whether there is a problem that is really an issue for anyone and to which funding needs to be applied. The real question is that given the number of components in an exascale platform, with variable power, and the small scale of the hardware, how big is the problem going to be? The answer is unknown at present.

[3] https://stone.umd.edu.

Given these research issues, the next step in the process is the development of a roadmap. A one–day workshop will be held on June 6, 2012, in the Washington D.C. area to bring together laboratory representatives with key people from industry to sort out what part of the problems belong to the DOE. This was originally organized by someone from DOE and the issues that were raised in Catonsville were focused on the DOE, but the smaller workshop will really focus on the issues that should be of concern to DOE. Input will help determine if there is a need for a large workshop or to issue a solicitation, or it may determine that there really is not a problem. Dr. Nowell believes that a little research is needed to develop the scale of the investment needed to solve the problem, assuming that there really is a problem.

Discussion

The workshop included those who work in fault tolerance, represented by four attendees from industry. A larger-scale workshop would invite people from Google to share expectations regarding hardware and resiliency, and those from cyber security. She expressed that understanding is needed within the frame of exascale science, and determining what DOE can do itself without relying on others.

Dr. Giles echoed Dr. Nowell's earlier comments that researchers may assume accuracy with their results until they encounter roadblocks or problems. Dr. Nowell commented that the issue of silent errors was discussed, and that there is a need to rethink repeatable results and the characterization of accuracy in systems that are prone to silent faults. There are potential issues for hardware and she is concerned about the lack of a roadmap. Influencing hardware developers and vendors could provide answers sooner. In addition, there is a need to understand problems and help applications people learn what they can expect from new machines. Dr. Giles believes that a factor in developing understanding is the physics involved and not just the random faults in a machine where the sample is no longer known.

Public comment

None

Presentation: Program Response to the Networking Committee of Visitors

Dr. Thomas Ndousse-Fetter of ASCR shared the responses to COV recommendations to the Next-Generation Networks for Science Research Program (NGNS). The COV was held on October 10 – 11, 2011. Four COV members were from NSF, and most were program managers and were very familiar with the task. They have gone through prior COVs that informed their expectations. A caveat is that they were expected to compare DOE's process with that of NSF.

The COV was charged with reviewing NGNS management processes. It is focused on high performance networking tools and middleware to help the ASCR research communities utilize the capabilities of current and future computing infrastructure. The COV assessed DOE laboratory and university projects, and assessed the award process impact on the depth and breadth of the portfolio.

The COV noted that the NGNS is effectively managed and that the objectives of the High-Performance Network and High-Performance Middleware Program elements are aligned with DOE's aim to deliver needed capabilities. The COV found that NGNS uses effective mechanisms to monitor projects including reporting site visits and PI teleconference meetings.

From an international standpoint, NGNS has engaged top-level researchers and large-scale high-performance network infrastructure developers in first-class research and innovations. Persistent development has lead to world-class networking capabilities to enable the science critical to DOE missions and priorities.

The IT infrastructure supporting project archiving is not as good as that of NSF and others. SC is just beginning to put this together and SC-wide IT infrastructure similar to FastLane will address documentation problems identified by the COV.

Charge 1 for the COV was to examine the efficacy and quality of the peer-review process. It is recommended that NSGS seek active means to broaden participation in all phases from workshops to solicitations development and announcement. Program managers are charged with workshop development and the identification of attendees who are leaders in their fields. However, managers face constraints for a variety of reasons.

Specific to assessing the efficacy and quality of the processes used to solicit, review and document applications and proposal actions, the COV urged that the NSGS stay consistent in handling review criteria across solicitations and clarify the role of letters of intent (LOI) and enforce an LOI policy. The current accounting system may not support archiving all submission-related information, but a new system being developed by SC may take care of this.

The COV commented that solicitations should be more explicit in conveying the expectations for deployment on ESnet and other DOE networking infrastructure to ensure fair and appropriate proposal review.

The COV was concerned about having a central repository for reviewers, in their examination of proposal process fairness. Currently, a variety of mechanisms are used to define an individual and institution, their area(s) of specialty, and other things. The COV feels that there should be a panel summary reflecting discussions of each proposal. However, a summary is not required for the current SC peer-review guidelines. There are also some proposals that are well-reviewed but not funded, due to a variety of factors including the applicability of the proposed research to the SC mission. SC's policy is still evolving on this issue and at the time of the COV, SC was not required to compile data on rejected proposals. Now, SC is required to classify these as preferred proposals and describe why they were not funded.

In a review of solicitation development, the COV lauded NGNS's participation in the Early Career PI Program but was disappointed that no proposals were funded via this program. It was recommended that NGNS find ways to reach out and convey the objectives of NGNS to young investigators. This can occur through conferences and workshops attended by SC, and in inviting young researchers to join panel reviews to learn what DOE is doing. These activities might lead to proposal submission. NGNS is also encouraged to revisit the value of long-term and short-term research, as longer-term may give an opportunity to engage and attract young investigators.

The COV charge included assessing the efficacy and quality of processes used to monitor awards, projects and programs. NGNS uses a variety of techniques and the COV was able to see that through the documentation that they received.

The COV recommended that NGNS allow PIs to address reviewers' comments when awards are made and how the awards are managed. The NGNS office does take reviewers' comments into serious consideration, and reviewer comments are considered among the factors used to make decisions on awards and funding, and to manage projects.

NGNS was encouraged to formalize and document the negotiation of awards, particularly when budget reductions are discussed. In this instance, a project manager has to negotiate and re-scope a project to fit within established budget limits. This often requires a lot of communication and that should be archived to understand the decision process that occurs. The new archiving system will allow for this summarization versus storage within the email system.

Annual progress reports should be made available online, per the COV. This is not a requirement in current policy, however. Final reports are required to be published on the science.gov website, and the DOE does require an annual progress report.

Charge Two asked the COV to assess how the awards process has affected the breadth, depth, and national and international standing of the portfolio elements, particularly within the boundaries defined by DOE missions and available funding. The COV recommended that NGNS have clear strategies for future funding allocations between long-term fundamental research, near-term research and development, and testbed support. This can be difficult to do in instances when there is an award out there and allocation decisions have to be made. If demarcations or apportionments are made, the quality that will be received is not necessarily known. Typically, NGNS would implicitly state that in the announcement, but would not associate a specific budget limit within each category. A lack of good proposals could make this challenging.

NGNS was encouraged to engage the next generation of network researchers in R&D within the context of the DOE mission and Department priorities. NGNS excels in identifying research in fields such as computer science. DOE also leverages basic research in applied mathematics and does a good job reaching those young investigators. Involvement in society and association workshops and workshop sponsorship allows young researchers to learn about NGNS.

Lastly, NGNS collaboration with other federal funding agencies was recommended. NGNS was commended on its inter-agency collaboration and coordination, especial in organizing activities across the agencies.

Discussion

Dr. Williams asked about COV input on tools developed in the network, and the crossover into application areas and how that is supported over time. Dr. Ndousse-Fetter clarified that the COV only covered the processes used and not the program's scientific merit or the execution and output of research. However, the long-term sustainability of the tools and software that NGNS has developed have transitioned from R&D to application and this is an ongoing issue. NGNS does produce tools through research activities but lacks a mandate to sustain tools over a long period of time.

Some may want to reinvent NSF at the DOE, noted Dr. Blumenthal, yet it is recognized that the responses of panel summaries is not required by SC guidelines. He noted that a cross-cutting theme is to find ways to annotate elements of the processes to analyze them and track proposer quality, among other things. However, there seems to be a challenge simply in capturing data. Dr. Hitchcock responded that panel reviews are not required of DOE and panel summaries are not permitted in the DOE. He admires FastLane and hopes that DOE's IT structure will permit similar accomplishments. DOE has many tools and is making positive strides, but an SC approach will also have to deal with privacy issues. If information is kept in a system, then a Federal system of register is needed and that requires a lot of security and legal consent. For a single program, it makes no sense to do this. SC will take this one and with deal with legal counsel. Dr. Giles noted that several COVs raise the issue of needing documentation and that the response is often that this will be addressed by the next system. Dr. Hitchcock responded that the proposal receipt system is already up and the interface for the review process is coming.

In response to Dr. Huntoon's mention that the COV cited a lack of good proposals from young investigators, Dr. Ndousse-Fetter shared that NGNS will have to broaden outreach to this group and those who would submit proposals. The goal is to ensure that the program is advertised in the right places. However, system management must continue to aid solicitation quality. With each proposal, the submission is reviewed and steps are taken to improve the system. NGNS has a very good computational science graduate fellowship and can interview people who are entering it from other areas to inform the system. This is true for networking and computer science, and NGNS will have to keep improving.

Dr. Chan noted that the COV urged documenting the decision trail, including what is funded and why. He endorsed the recommendation and suggested that the archiving system be accessible by future COVs to look at that and improve the transparency of the whole process. Dr. Hitchcock replied that SC has already moved to keeping documentation ready for all decisions, and keeping it out of emails and in files that you can reliably find somewhere in a Microsoft file structure.

Public comment

None

Presentation: Update on Computer Science Committee of Visitors

Mrs. Christine Chalk described the COV for Computer Science Program. All programs do a COV on a five-year basis to look at processes and results. This is the second review for Computer Science. The COV will examine a random sample of proposals and reviews, efficacy and process quality, and active program and project monitoring. The COV is also asked to understand the impact of the Program on the DOE mission and objectives. In this case, the challenges are multi-core hybrid computing and peta-to-exascale scientific data management, and the Program's national and international standing in terms of other computer science research programs that are focused on the demands of HPC and the analysis of petascale datasets.

The COV is chaired by Tony Hey of Microsoft and Susan Graham of the ASCAC. Most U.S. citizens are barred from participating in this COV due to the status of some proposals.

The COV will occur on July 10 – 11, 2012, in Germantown, MD. There will be an update report delivered to the ASCAC in August 2012, with a final report due to ASCAC in November 2012.

Discussion

Dr. Giles encouraged ASCAC members to offer any comments to Tony Hey, Susan Graham, and others on topics that may emerge. Mrs. Chalk shared that one thing to comment on are aspects of the computing research portfolio from peta-to-exascale, and to take this opportunity to reflect on that in this context.

Dr. Chan wondered if there is a piece of the charge that addresses how the Program is bringing in new researchers. This is not explicit in the charge, said Mrs. Chalk, yet previous COVs have commented that this should be included. Some take this more earnestly, but DOE is not allowed to collect demographic information to allow that type of review. Dr. Chan urged distinguishing between demographics and the turnover of proposers, and Mrs. Chalk shared that every COV has commented on that with a desire to think about gathering proposals from new people.

Mrs. Chalk noted that the National User Facilities Organization User Science Exhibition would be held on March 28th in the Dirksen Senate Office Building and on March 29th in the Cannon Caucus Room in the Cannon House Office Building. This is a public event.

Public comment

None

Presentation: Final Report from Magellan and Update on Advanced Networking

Dr. Kathy Yelick of LBNL reported on Magellan and shared an update on the Advanced Networking Initiative (ANI).

HPC in science can be divided into the categories of scale, volume and data. An example of scale is climate change. A project called BISICLES uses adaptive mesh refinement to chart ocean flow into water. It looks at warming in the ocean and warming of the Antarctic ice sheet and correlations with sea level rise. In the last year, the effort was allocated 19M hours at NERSC.

Science through volume examines 10s of thousands of experiments and the ability to screen data and resulting problems. One example is protein folding and a database that stores the data and allows people to look through and identify properties that are applicable to their work. A materials project looks at manufacturing and cutting in half the time from the original discovery and to the manufacturing side. It helps narrow the range of materials to be discovered and learned about through experimentation leading to a smaller set of materials to be examined.

Science from data moves from simulation to image analysis, and capitalizes on this by modeling the appearance of supernovae and other analysis that automatically uses machine algorithms to look at data. Four of 10 science insights of the past decade are tied to computing work at LBNL and work has contributed to breakthroughs of the decade. Data rates from experimental devices will require exascale volume computing.

There are still huge data analysis challenges to be solved. The DOE has a unique space due to its computing facilities. Most have a huge increase in data sets and the need for ways to deal with data-intensive science, particularly with emergence of exascale computing.

Dr. Yelick described two ARRA projects – the ANI and Magellan.

ANI looked at the movement of large data sets and how to commercialize data distribution in 100 Gbps networking. This fit into the science in data category, and included the Prototype National Network and testbed for network R&D.

ANI relied on ESnet and traffic growth that is higher than that found in commercial networks. Forty DOE sites are linked to 140 other networks. The problem of sending many data sets is greater than sending one large data set. ESnet is the first continental scale network and will move to production this year and use dark fiber to deliver bandwidth.

ESnet's policy board had its first meeting in 2012 and concluded that outstanding people in operations must be preserved for this to be the most powerful networking tool in the world, and to leverage the dark fiber. The network research testbed is underway, dark fiber is being used, and additional research project support is available through the testbed.

ANI moved from a table-top testbed in 2009 with the purchase of dark fiber. In 2010, more fiber, optical gear and routers were purchased. A partnership with Internet2 was established in 2011, and by October 2012, there will be a complete network buildout and a transition to production before the end of the year.

One testbed project is Monitoring and Visualization of Energy in Networks (MAVEN). One problem with networking and energy is limited understanding of the use of energy in networks. A goal is to establish a baseline of energy use in networks and provide networking researchers with opportunities for improving energy efficiency. A network of centers could advance global optimization and find a way to make energy cheaper and more available. Data on this will become more available over time, and as ANI moves into production, there will be more data from the network and options like turning hardware off.

A second testbed project is end-to-end circuit service with OpenFlow. The premise is to tunnel and get past all of the protocols that are in the system overhead when sending data. A high performance network is possible through a system called OSCARS, but there is still a lot of work to configure networks for large data flows. Dynamic tunneling is on automatically and you can discover the end point to circuits. Using Remote Direct Memory Access (RDMA) produces 9.8Gbps out of 10Gbps from New York to Seattle, as shown at Supercomputing 2011. This can reduce CPU usage from 80 percent to four percent and still achieve network performance.

ANI's legacy is the existence of a unique 100G networking facility and other facilities will come into this as well. It will enable the first-of-its-kind big data science and having high bandwidth, among other things, and data transfer nodes will lead to big data science. Some in ESnet are already promoting these ideas and thinking about setting up a science DMZ.

Dark fiber is separate from the current connection network. It holds the potential for future ESnet upgrades and additional equipment can be purchased. In the short term, this is useful for networking research. Dark fiber is essential but setting up tests is difficult; industry and laboratory researchers could collaborate on this.

The other ARRA project, Magellan, looked at a DOE strategy for mid-range computing and a DOE cloud testbed for science. Part of this is setting-up a testbed hosted at ALCF at Argonne and NERSC. This speaks to science in volume.

The Gartner 2010 Emerging Technology Hype Cycle proposed that in 2009, cloud computing was moving to actual production. For some, cloud means virtualization and other things. The National Institute of Standards and Technology defines it broadly as consisting of resource pooling, broad network access, measured service, elasticity, and self-service.

Clouds make sense for science as resource pooling allows for economies of scale and running commercialized systems at higher rates of utilization of around 60 percent. Other sites have idle hardware and run at about 20 percent. HPC centers run at 90 percent utilization. Pay as you go (measured service) and HPC centers charge in hours, whereas commercial clouds charge in dollars. Researchers carefully control and manage hours like they do a budget, and might want guaranteed access to computing cycles. Elasticity impacts science and cloud computing as HPC centers allow job scale-up yet users wait in queues. At commercial centers, they do not queue but there is no guarantee of the quality of work. HPC centers address self-service through fixing some software to address version types of operating systems and compilers. This is especially true in large collaborations when using computing resources over a large community. This is seen in large hadron collider projects. Virtualization that comes out of DIY administration is useful here as you can build a software stack and make that available to a research group.

Magellan looked at questions about clouds and science and the usefulness and cost effectiveness. Magellan built a testbed at Argonne and one at NERSC using almost identical hardware but different numbers of nodes and other equipment aspects.

Grid computing is different. This deals with federating different resources together to give an impression that you have a distributed computing infrastructure. Cloud computing is about centralization and not about distribution. It looks at the movement of a workload from one center to another, and it gives you resilience and some form of surge protection.

When Magellen began, the science community was surveyed to gauge their interest in clouds. The leading response was a desire for access to additional resources and more cycles. The second was the ability to control software.

A demonstration of cloud technology for science was held in 2010. A testbed started in September 2009. It was deployed by February 2010 and had some early users in March. It was demonstrated for the Joint Genome Initiative in 2010 and most recently, shown in a STAR experiment that was running data analysis workload coming off of a detector.

The STAR project was a federated cloud project and specifically asked for ARRA support. It demonstrated the federated testbed at Argonne and NERSC together, and used realtime data from Brookhaven. It was the first time this amount was analyzed in real-time and started with 20 initial machines at NERSC and then more machines at ANL.

The performance of clouds for science was another part of Magellan. This was started by LBNL around the time that Magellan was created and used NERSC benchmarks and applications that covered algorithm and science space. The NERSC benchmarks do not use Monte Carlo or independent parallelisms as they are not criteria for cloud computing. It had a set of benchmarks, shrank down the job size, and used reduced data for cloud benchmarking. Small- and mid-sized problems were used and the benchmarks were used to understand the slowdown of the cloud relative to an HPC system. This resulted in the slowdown of many applications, especially PARATEC which had a 53 times slowdown. One that faired well was BLAST.

On the commercial side, Amazon, Yahoo, Google and others can set-up a high-speed network cloud and not over-subscribe their processors. Amazon is collaborating with the laboratory and seeks to make cloud computing better for science. The Berkeley lab team is now running in the Amazon cloud and it is not over-subscribed. Magellan is running without virtualization and within a cluster. It is much better in performance.

The first results and performance differences are being addressed. Infiniband was baseline and PingPong Latency was 40 times better. There are micro benchmark numbers, but greater concern about performance. Researchers looked specifically at PARATEC. Infiniband, network hardware and network protocol are all important for performance. When running the virtualization, the penalty is substantial. It was found that virtual overhead increased with the core count. The conclusion was that virtualization does hurt in running a parallel job but virtualization may cause load imbalance; there is a lot of communication happening.

Elasticity is an advantage gained through cloud computing. There is an assumption that clouds are cheap and infinitely available. Researchers looked at Hopper “Unleashed” which is Hopper in pre-production to understand what job size would look like if people had unlimited access to a cloud. Consequent analysis done in Franklin showed the on-demand response and the size that Franklin needs to be. The number of cores required to run a job immediately is 2,000. It is not practical to provide an elastic response time for these types of large jobs.

This work also looked at the cost of cloud computing for science. Cloud computing is about a business model and the ability to buy time in the cloud. In comparing HPC Centers with commercial options, it is recognized that the approach with science is that there is an unbounded market demand whereas commercial options charge for time. There are different models of scaling and workloads based on throughput. One issue in the commercial realm is getting data in and out of a cloud, whereas HPC centers can move large data sets.

Dr. Yelick shared that cost analyses favor centers. Assuming the commercial clouds run just as well as an HPC, they have a huge advantage. Purchasing the computer systems at NERSC at 1.38B hours equals the $180.9M Amazon cloud. NERSC's annual budget equal to commercial options is not equal to $200M and does include consulting, administrators and other resources. Commercial providers' response has been that they are not interested in this market.

There are other factors in the price differences between HPC and a public cloud. If a private data center, for instance, is 30 percent utilized, a commercial cloud is 60 percent utilized, and a HPC cloud is 90 percent utilized, then non-HPC options are more efficient as scientists are made to wait. The cost of people with HPC is another factor; the cheapest system in NERSC is Hopper with 150 cores. The most expensive are designated clusters for physics that have their own software stack and other components. The cost of power is another issue. Google and others make investments to get good prices on power, but all labs negotiate on good pricing for power. From the standpoint of energy efficiency and overhead, Center use is much better than the typical commercialized center – there is not much more to optimize as the system is running so efficiently. The cost of specialized hardware impacts an HPC center as there is an added cost because it has a larger workload. In considering the cost of commodity hardware, the public cloud buys 100 times more cores and they purchase huge quantities. The biggest factor in the comparison is profit and labs are not permitted to make a profit.

In pricing models for Amazon, Intel and AMD, it was found that the cost of a small instance at Amazon has dropped 18 percent over five years. The NERSC cores dropped by 10 times - from 20,000 to 200,000 cores from 2007 through 2011 - using a roughly flat budget. It demonstrates that Amazon, Google, others make a lot of money from their cloud offerings.

There are lessons that the HPC centers can draw upon from cloud environments. Centers are not good at providing a higher service level, users might need to be in control of access, and provisions may be needed for small groups that want to configure systems software. Other features associated with clouds include virtualization for the over-subscription of nodes, and the Map-Reduce programming model.

Among key findings, there are things that Centers can learn from clouds but Centers hold a cost advantage and are supportive of scientific workloads.

A final report on this work will be available on the ASCR website, and will include security models and benchmarking data, as well as cost and programming environment analyses.

Discussion

Dr. Sarkar asked what can be learned from the impact on application enablement, such as the exascale or writing software differently. Dr. Yelick shared that the STAR project included a lot of preliminary software work, but that once an investment is made in building a software stack it can be reused. Productivity use may be available for map reuse but not for productivity in general. LBNL is teaching an online software engineering course using a software course developed at Stanford. There are 50,000 students which caused difficulty when all were working on Google and downloading content on the same day.

Dr. Williams brought up ESnet and the goal to expand beyond the current four sites. Dr. Yelick confirmed the goal to expand but that this is budget dependent. Lawrence Livermore National Laboratory and Berkeley are on the short list, but there are other labs planned in the roll-out.

Dr. Williams asked about the interface with much slower networks at around 1 Gbps, and Dr. Yelick shared that ESnet has connections with the networking community. This was shown by the climate research done with colleagues in the U.K. and transferring data from Oak Ridge to NERSC. Originally, this was going to take two months but was sped up thanks to administrators on both ends tweaking their systems to move data. Still, this can be a physical limitation.

Dr. Giles thanked Dr. Yelick and her office for their work on Magellan and taking short-term ARRA funding and identifying and executing projects that will have long-term mission impacts.

Public comment

None

Board business

Dr. Giles thanked Dr. Hack, Dr. Manteuffel and Dr. Tang for their participation on the ASCAC Board.

Overall public comment

None

Adjourn

Dr. Giles adjourned the ASCAC meeting at 11:39 a.m. EST.