Minutes of the
Advanced Scientific Computing Advisory Committee Meeting
October 31, 2012
Teleconference

Because of the effects of Hurricane Sandy in the National Capitol Area and on travel in the eastern United States, this meeting, which was originally scheduled to be held in Washington, D.C., was abbreviated and conducted as a teleconference supplemented with the graphics of a web conference. A roll call established that a quorum of Committee members was in attendance via telephone.

ASCAC members present:

Marjory S. Blumenthal
Barbara M.P. Chapman
Vinton G. Cerf
Roscoe C. Giles (Chair)
Sharon C. Glotzer
Susan L. Graham
Anthony Hey
Gwendolyn L. Huntoon
Juan Meza
John Negele
Vivek Sarkar
Dean N. Williams

ASCAC members absent:

Marsha Berger
Vincent Chan
Jacqueline Chen
Jack J. Dongarra
Linda R. Petzold
Victoria White

About 55 others participated in the conference call.

The meeting was called to order by the chairman, **Roscoe Giles**, at 10:05 a.m. He announced that this was Daniel Hitchcock's last meeting as Associate Director (AD) because he was retiring. Giles offered his sincere appreciation for Hitchcock's service to the Office of Advanced Scientific Computing Research (ASCR) and to the scientific community.

William Brinkman was introduced to give the perspective of the Office of Science (SC). He also expressed thanks to Hitchcock for his years of service to DOE. Barbara Helland will be the temporary AD for ASCR. Expertise and leadership are needed to move forward with the Exascale computing initiative.

Planning for the Exascale computing initiative was started by Stephen Koonin; that plan was not received well in Washington because of the associated cost. The effort will stay at a low profile. An attempt is being made to put together a plan for the intermediate future, and it will be presented to ASCAC sometime in the future. Ways to cooperate with other parts of the government are being looked for.

The ensuing three months will be exciting in Washington. It is hoped that the lame-duck Congress will pass an FY13 budget. The planned possible sequestration of funds would be disastrous. It would cost DOE $400 million. DOE is working on the FY14 budget. It may be spring before the administration acts on it.

A lot of good things are happening. Titan is up and running and may beat out Sequoia as the biggest, fastest computer in the world. Both Argonne National Laboratory (ANL) and Oak Ridge

National Laboratory (ORNL) are upgrading their power infrastructure. Lawrence Livermore National Laboratory (LLNL) is encouraging the opening of its computer resources to industry.

The Higgs boson was discovered, and $\sin^2 \theta_{13}$ was measured. The Nobel Prize went to a researcher in protein structures who was funded by the Office of Biological and Environmental Research (BER). The Light Source 2 (LS2) is being finished at Brookhaven National Laboratory (BNL). The Linac Coherent Light Source (LCLS) is being used at Stanford, and the LCLS2 is being planned. The challenge for ASCR is figuring out a path to the exascale. The Long-Baseline Neutrino Experiment (LBNE) is controversial, but it is hoped that it can be fit into the budget. International Thermonuclear Experimental Reactor (ITER) is making good progress. All this news reflects a lot of good results in a large number of fields. The Office of Management and Budget (OMB) needs a long-range plan for facilities.

The prospect for future budgets is a constant level of funding. It is not clear when we will break out of the flat-budget scenario.

During the past 3.5 years, the energy programs have been advanced. Polls indicate that 70% of the population believes in climate change and are conscious of the issues. Hybrid cars are 3% of the U.S. market. New efforts to increase gas mileage are seeing fruition. Solar energy is coming on strong with 0.5-GW installations in California. Wind is moving upward, also; it is as economical as coal. Something needs to be done about carbon dioxide capture and sequestration. Carbon dioxide is now being used to enhance oil recovery; however, the costs of carbon dioxide are still too high. Small modular reactors are being promoted as a low-economic-risk means of expanding electricity generation in proportion to the expansion of demand.

Giles asked if there were any chance that information about the Exascale computing initiative might be forthcoming before the Supercomputing 2012 conference. The original $5 billion request was intended to ensure U.S. leadership in high performance computing. Brinkman replied that OMB has to approve the Exascale computing initiativereport, and that will not happen before Supercomputing 2012. It is still hoped that the United States will be the world leader in high performance computing. In addition, the Exascale computing initiativeand big data go hand-in-hand.

Cerf asked if there were any chance that other high performance computing agencies could get together to share resources. Brinkman said that SC is talking with the Advanced Research Projects Agency (ARPA) and the National Security Administration (NSA). Cerf noted that cloud computing will work well for certain algorithms and not others. He asked if DOE were checking to see if codes could be run on a cloud. Hitchcock responded that ASCR had explored the use of one of its codes and found that cloud computing was good for data analytics; however, the cloud's latency made the performance very bad for calculating partial differential equations.

Hey asked when JASON would report on their exascale study. Hitchcock replied that they will discuss this subject on the weekend following this meeting and produce a classified report and an open report.

Negele asked what was being done with the National Nuclear Security Administration (NNSA). Brinkman said that the discussions were tied up with the New Strategic Arms Reduction Treaty (New START), and budgets are being revisited.

Vivek Sarkar was asked to report on the activities of the ASCAC Data Subcommittee.

The Subcommittee has had input from several industry entities. The charge was to understand the similarities among and differences between these data challenges and the potential to leverage research investments to address issues spanning both the exascale and data-intensive science. The Subcommittee was to specifically address what investments are most

likely to positively affect both the exascale goals and the data-intensive science research programs, including data management at the next-generation facilities.

The Subcommittee interpreted the charge in terms of timeframe and decided that the focus should be on a 10-year timeframe. The scope clearly includes experimental facilities and scientific simulations that exemplify SC's unique role in data-intensive science. Upon completion of the study, some of the conclusions may be more broadly applicable (e.g., to NNSA and the Department of Defense). Other data-intensive applications (e.g., cyber defense) were ruled out of the scope of this study.

The desired outcomes are to make recommendations for identifying synergies between data-intensive and the Exascale computing initiative in the next 10 years within the context of furthering science, while identifying investments that are most likely to positively affect both data-intensive science goals and exascale computing goals. The Subcommittee will conduct a detailed study of exemplar scenarios from experimental facilities and scientific simulations to provide context for recommendations. There are several cost-cutting concerns: technology roadmaps, analytics and visualization, and software challenges.

The planned method was to build consensus on the scope through regular teleconferences between Subcommittee members and Office staff. The Subcommittee is dividing up responsibilities for different aspects of the study. There will be a detailed discussion of data-intensive scenarios in SC that will move on to a detailed discussion of SC investments in exascale computing and data-intensive computing. The exercise will end with the writing of draft and final reports.

Several recent ASCR workshops have focused on the exascale and data-intensive computing. There is an industry context of big data, also. The Subcommittee does not want to duplicate previous work. There are large volumes of observational data that necessitate data reduction at the data source; large volumes of simulation data that necessitate in situ and in-transit analytics; and large degrees of data sharing among nations that necessitate coordinated workflows.

Three examples of data intensive computing will be used to illustrate the challenge: high-energy physics [Large Hadron Collider (LHC experiments)], simulations (S3D combustion code), and climate [the Earth System Grid Federation (ESGF)].

The LHC detectors ATLAS [A Toroidal LHC Apparatus] and CMS [the Compact Muon Spectrometer] produce raw data at rates equivalent to 1 PB/s. These data are reduced in real time to more-manageable volumes (10 PB/year). Highly efficient networks and Grid technology enable the worldwide distribution of these data. Derived data sets are read-only after creation, thereby offering opportunities for new storage technologies in which reads perform better than writes.

On the issue of distributed vs. centralized computing, an experimental facility must be centralized, but a distributed approach to computing and storage can be adopted so computing resources around the world can be accessed and so international funding and intellectual capability can be leveraged. The costs of distribution include (1) added overhead for data movement and replication and (2) added complexity.

Simulations also produce massive amounts of data. The classic use of the petascale workflow model will not work at the exascale. It is expected that a simulation run will generate 1 PB of data with a dump every 30 min. Input/output (I/O) bandwidth constraints make it infeasible to save all the raw simulation data to persistent storage. Therefore, in situ and in-transit analyses are a necessity. A workflow must be designed to support the smart placement of analysis and

visualization, the tracking of large graphs, and the reduction of checkpointing size with in situ analytics.

The design considerations for data-intensive workflows and simulation include location of the computer resources; data access, placement, and persistence; and synchronization and scheduling.

Big data must be brought together with analytics. The significant challenges are the variety and complexity of data, the need to move large volumes of data quickly, and the volume of big-data analytics. At the same time, the information must be globally accessible, the infrastructure must be flexible, and the framework must be scalable.

The Earth System Grid Federation (ESGF) distributed data archival and retrieval system employs distributed and federated architecture, supports discipline-specific portals, supports both browser-based and direct client access, allows for single sign-on, and provides automated script in publication tools based on a graphic user interface (GUI).

The cross-cutting issues are:

- Solid-state memory will scale with Moore's Law and provide higher-performing and lower-energy storage than will disk storage
- Stacked memory will dramatically improve performance and energy use while supporting new algorithmic paradigms of moving computing to memory.
- Silicon photonics will significantly increase bandwidth and reduce the power needed to connect to optical networks.

Visualization must also allow getting to the information. Data and analytics must be collocated, and that influences the platform.

The last cross-cutting issues are the software challenges related to data-intensive science: The human cost will be a major bottleneck. Data reduction and analytics codes are often branch-intensive and do not make good use of graphics processing units (GPUs) or single instruction, multiple data (SIMD) vector parallelism. And the use of parallel I/O frameworks requires significant expertise.

The schedule is:

The Subcommittee's panel was formed in September 2012.
The kickoff teleconference was held September 20.
Additional teleconferences were held.
ASCAC is being updated today.
Information gathering will be conducted in November and December.
Report writing will start in January 2013.
The report is due March 2013.

Additional scientific domains and scenarios will be looked at; supercomputer-center costs and requirements for storage will be considered; and the current ASCR investments related to big data and the exascale will be reviewed.

Cerf had two concerns: algorithms and architectures. He noted that Ivan Sutherland had supported asynchronous design at the Turing Celebration seminars and asked if any work were going on in asynchronous computation. He also noted that multi-core designs have problems in getting data to the chip and getting data on and off the chip and asked if any research were going on in those areas. Sarkar replied that there *is* research going on in architecture design. Asynchronous computing is an important direction to go in. Algorithms must be designed to move data more effectively. Investment goals are not known to the Subcommittee at this time.

Hey asked what was meant by proxy applications. Sarkar replied that proxy applications reduce the workload and explore different mappings and applications.

Rick Stevens of Argonne National Laboratory asked where the input on biological sciences from the Office of basic Energy Sciences (BES) would come from. Sarkar responded that the Subcommittee had planned to contact Stevens and others in the next few months to get that information.

Giles stated that the Subcommittee should identify research recommendations. Sarkar answered that the point is well taken, and the Subcommittee would welcome additional input information.

Daniel Hitchcock was asked to report on the activities of ASCR.

When he walks out the door, the Office will be in good hands.

Synchronous algorithms are the kiss of death at the exascale. As one moves to a world where communication energy and communication time dominate performance, getting data where it needs to be just in time is a crucial issue. It is hoped that a solicitation for a data-intensive-science co-design center will be issued this year, budgets permitting. Applications that mimic the behavior of the whole system are being produced. They are small enough to be simulated to see how they perform on different chip architectures and also small enough to able to be run on designers' simulators to see which features have the most impact. Getting features on a chip is a relatively expensive proposition.

ASCR has done a lot of work with other SC offices to see what the computational needs will be. Two new staff members will help figure out how to deal with the light sources and other big-data producers and be a resource to the data community.

On the subject of budget uncertainties, there is nothing to add.

Giles and others expressed appreciation for Hitchcock's hard work over the years. Hitchcock noted that the community has gone from 56-kb networks to 100-Gb/wavelength networks and cutting circuits through the infrastructure in real time. Cerf noted that one thing that made the Internet flexible and scalable was an architecture that was carefully chosen not to do too much. One needs to look for loose-coupling computer architectures to provide such scalability in the future.

Negele asked if there were anything else to say about collaboration with NNSA. Hitchcock said that the Office has been working closely with its NNSA counterparts for 1.5 years. A strategic planning process is now in place. The Office has forged good relationships with laboratory managements at DOE and NNSA laboratories.

Giles asked if there were any comment on DOE's presence at Supercomputing 2012 the following week. Hitchcock said that ASCR is sending about a half-dozen people who will have a lot of work to do there. There will also be people from the national laboratories, although not as many as had been hoped.

Lucy Nowell was asked for a report on the program's response to the Computer Science Program (CS) Committee of Visitors (COV).

On the issue of improving the online information-management capabilities of the CS program, SC is in the process of establishing a new Portfolio Analysis and Management System (PAMS) to support program-management activities.

On the issue of establishing (1) a reviewer database that records areas of expertise, quality of past reviews, responsiveness, and conflicts of interest and (2) a principal investigator (PI) database that identifies previous successful and unsuccessful DOE proposals, these features will

be incorporated into PAMS. ASCR is also working to include more project information on its website.

On the issue of relaxing the present, very stringent approach to exclude reviewers with potential conflicts of interest to achieve a better balance between external reviewers and reviewers familiar with DOE's and ASCR's mission, ASCR program managers will continue to seek the most qualified reviewers available while complying with DOE guidance on conflict of interest. Addressing DOE and ASCR mission drivers is a primary responsibility of the program managers.

On the issue of providing a longer-term and more coherent schedule of solicitations, such a schedule is being developed as part of the planning for the Exascale computing initiative; however, budget uncertainties and continuing resolutions often interfere with plans. Having more program managers available helps.

On the issue of creating a new mechanism for funding the exploration of promising new ideas that might not conform to the planned solicitations, this action was seen as unneeded because unsolicited proposals and Early Career Research Program awards fulfill this need and are funded nearly every year (within budget constraints).

On the issue of encouraging the use of new technologies and new media, including social environments and hubs, to provide more efficient oversight, program managers will continue to use all the resources available to provide effective management and oversight of their portfolios.

On the issue of using better metrics for evaluating the need for workshops and other conferences, workshops are organized to facilitate a community dialogue with regard to potential technology developments during a period of rapid change. In addition, conferences foster collaboration among researchers within and across programs. ASCR aims to minimize the number of events and maximize their value.

On the issue of maintaining an appropriate balance between funding for the exascale research program and the base CS research program at the national laboratories, there are (1) an ongoing tension with respect to portfolio balance across several dimensions and (2) attempts to achieve an appropriate balance within and across programs, between national laboratories and universities, and between large and small universities. There is no set formula that provides funding to national laboratories rather than other applicants. The amount going to national laboratories has not changed in the past 3.5 years.

On the issue of devoting more prominence to research into energy-efficient computing, machine learning, and data analytics, it is noted that energy-efficient computing is a major thrust of the Exascale Computing Initiative. Multiple solicitations are planned in the near term to address challenges of data analytics. One problem is that some work cuts across other parts of ASCR.

On the issue of having review panels contain a mix of external and national-laboratory researchers, this is always a goal for panels within the constraints of conflict-of-interest requirements.

On the issue of having the Scientific Data Management and Analysis Program in the CS program work with the BES and BER experimental data communities, ASCR and BES held a joint workshop in October 2011. An additional joint workshop is in the planning stages and will include the experimental communities of BES, BER, and the Office of High Energy Physics (HEP).

On the issue of building expertise in machine learning and data mining in support of big data and data-intensive science, ASCR will continue to develop research programs and data analytics to meet the needs of SC. ASCR looks forward to the ASCAC Subcommittee's report on data-intensive science and the exascale to help define this work. However, some of this effort is outside the CS portfolio and is done elsewhere in ASCR.

On the issue of ASCR doing all that it can to ensure that it receive sufficient funding for exascale initiative, ASCR leaders will continue to make the case for the Exascale Computing Initiative and its role in assuring the international competitiveness of the United States.

On the issue of the ASCR program's maintaining its leadership role in high-end computing by continuing to engage with the international community, ASCR is participating in a variety of international activities related to supercomputing in general, in file systems and I/O, and in the Exascale Computing Initiative. These efforts are expected to continue within travel-budget limitations.

On the issue of ASCR's working with SC to do everything possible to secure adequate funding for the Exascale Computing Initiative and to protect U.S. leadership computing technology, ASCR continues to work with SC leadership to develop plans and justifications for the Exascale Computing Initiative to make sure that it is funded at a level adequate to protect U.S. leadership in supercomputing.

On the issue of ASCR's negotiating to be allowed to fill the remaining approved CS vacancies as quickly as possible, one position was advertised during the past summer, and the applicant-review process is under way. It is hoped that a second position will be advertised this year.

Hey thanked Nowell for the comprehensive response. He noted that, on the issue about machine learning, data mining, and analytics, those topics do not seem to be covered widely in the Applied Math Program. And on the topic of addressing the challenges of multicore hybrid computing and peta-to-exascale scientific data management, the COV was pleased that the Data Subcommittee was also looking at the relationship of observational data and also simulation data. The COV would like to see more support but recognizes that the Office is budgetarily constrained.

Giles asked if there has been the possibility of funding unsolicited proposals. Nowell replied, yes; a solicitation goes out each year with suggested topics for unsolicited proposals. The Office has funded unsolicited proposals every year but one (when there was a budget shortfall).
Giles asked if the workshops with BER and BES could be held given the travel restrictions. Nowell replied, yes and no. The restrictions effectively limit workshops to the Washington, D.C., area, which is very expensive for visitors.

David Brown of Lawrence Livermore National Laboratory asked about conflict-of-interest policies. Nowell answered that conflict of interest is defined as a direct financial relationship by a participant, spouse, or child with the proposers. The level of concern may depend on the size of the award. This is similar to the regulations used by NSF. A one-year window is looked at for conflicts of interest. Coauthorship contributes complexity to the process.

Hey asked about balance. Given the uncertainties involved with the exascale, national laboratory input is important. There seems to be some disinvestment at the national laboratories. Giles said that sequestration and the exascale plan will put added pressure on this issue.

Cerf said that, because of continuing resolutions and the possible sequestration, some committees have to provide prioritizations at different levels of expenditures because it is

sometimes it is not a strict hierarchy with regard to priority; depending on the budget level, some things take precedence because they are affordable. He asked if there were any desire for such input and if there was the will within the Committee to provide it. Hitchcock replied that in some fields there are (e.g., decadal) studies to gain a consensus of the community of where priorities lie. Some advisory committees have priority-setting subcommittees, but that course of action requires additional paperwork. Giles said that ASCAC tries to schedule meetings to have the greatest impact on budget planning. The embargo of budgetary data until it is submitted to OMB is consternating.

Daniel Hitchcock was asked to introduce the new charge to the Committee. The Applied Mathematics COV charge has been posted on the web. It is virtually identical to the Computer Science COV charge. This is now standard procedure. A Subcommittee chair is needed.

Giles asked if there were any Committee comments on the charge. There were not. Anyone wishing to serve should contact Giles in the next week or so. Hey volunteered that he did not want to be the chair of this COV. [Laughter.] Giles noted that the intermediate report would be made at the March ASCAC meeting, and the report would be due in August.

The floor was opened to public comment. There being none, the meeting was adjourned at 11:52 a.m.

Respectfully submitted,
Frederick M. O'Hara, Jr.
Recording Secretary
November 7, 2012