

Draft Minutes
Advanced Scientific Computing Advisory Committee Meeting
November 18, 2013
Crowne Plaza Hotel, Denver, Colorado

ASCAC members present:

Barbara M.P. Chapman	Anthony J. G. Hey
Jacqueline H. Chen	Gwendolyn L. Huntoon
Jack Dongarra	Vivek Sarkar
Roscoe C. Giles, Chair	

ASCAC members attending via telephone/Internet:

Marsha J. Berger	John W. Negele
Vincent Chan	Linda Petzold
Sharon C. Glotzer	Victoria A. White
Susan L. Graham	Dean N. Williams
Juan Meza	

ASCAC members absent:

Vinton G. Cerf

Also participating:

James Ang, Manager, Scalable Computer Architectures, Sandia National Laboratories
John S. Binkley, Associate Director, Office of Advanced Scientific Computing
Research, Office of Science, USDOE
A. S. Buddy Bland, Project Director, ORNL Leadership Computing Facility, Oak
Ridge National Laboratory
Christine Chalk, ASCAC Designated Federal Officer, Office of Advanced Scientific
Computing Research, Office of Science, USDOE (via telephone)
Phillip Colella, Computer Science Division, Lawrence Berkeley National Laboratory
Sudip Dosanjh, Division Director, National Energy Research Scientific Computing
Center
William Harrod, Director, Computational Science Research & Partnership Division,
Office of Advanced Scientific Computing Research, Office of Science, USDOE
Barbara Helland, Office of Advanced Scientific Computing Research, Office of
Science, USDOE
Robert Lucas, Director, Computational Sciences Division, University of Southern
California's Information Sciences Institute
Frederick O'Hara, ASCAC Recording Secretary
Douglas Ray, Associate Laboratory Director for the Fundamental & Computational
Sciences Directorate, Pacific Northwest National Laboratory
John M. Shalf, Chief Technology Officer, National Energy Research Scientific
Computing Center
Horst Simon, Deputy Director, Lawrence Berkeley National Laboratory
Marc Snir, Director, Mathematics and Computer Science Division, Argonne National
Laboratory

Kathy Yelick, Associate Laboratory Director for Computing Sciences, Lawrence Berkeley National Laboratory

About 50 others were in attendance at the Denver site; about 10 others were in attendance via telephone/Internet.

Chairman **Roscoe Giles** called the meeting to order at 1:02 p.m. MST and explained that this was a replacement for the meeting scheduled in October. He reviewed changes in the agenda and announced that public comment would be allowed after each agenda topic.

A roll call of members was held, and a quorum was affirmed.

Robert Lucas was asked to give an update on the activities of the Top-Ten Exascale Approaches Subcommittee. He listed the Subcommittee members and “additional contributors.”

The top-ten technical challenges as ascertained by the Subcommittee members are

- Energy-efficient circuit, power, and cooling technologies, which will have a significant effect on cost of ownership
- High-performance interconnect technologies to limit energy use and latency
- Advanced memory technologies to dramatically improve capacity and bandwidth; memory currently has a 3-D architecture, and the amount of memory per core or per flop is falling off from Moore’s Law; changes are coming and need to be integrated into the system
- Scalable system software that is power- and resilience-aware
- Data-management software that handles the volume, velocity, and diversity of data
- Programming environments to express massive parallelism, data locality, and resilience
- Reformulation of science problems and refactoring of solution algorithms for the exascale, maximizing locality and minimizing communication
- Ensuring correctness in face of faults, reproducibility, and algorithm verification as noise margins get smaller and more faults occur
- Mathematical optimization and uncertainty quantification for discovery, design, and decision
- Software engineering and supporting structures to enable scientific productivity and make software engineers more productive

Reformulating science problems and refactoring solution algorithms for the exascale is the biggest challenge to exascale computing.

Giles asked if the Subcommittee had a longer list that got truncated or condensed to the requisite ten challenges. Lucas replied that the Subcommittee had had lots more than ten. Twenty participants each entered their own top ten challenges, so some consolidation was conducted. At one point, the Subcommittee considered just three metaproblems: hardware concerns, systems operations and tool concerns, and mathematical issues (with software engineering as an outlier). The data issues had slipped by the Subcommittee, so that topic was added after the previous ASCAC meeting. The Subcommittee saw these three metaproblems as not being responsive to the charge, so it flattened out the long list to ten.

Dosanjh noted that the big challenge is integration of software and hardware. He asked if these issues at scale were seen as a problem. Lucas replied that there is a chapter in the report on that subject. That issue is addressed, but it is not one of the top-ten challenges. He showed the report outline. An executive summary will be published, also. The report will be expanded to cover topics brought up by ASCAC. The current draft is 87 pages; the penultimate revision must be completed by November 30. Thirty-four people are contributing to the report. Contributions will be cut off in the next few days, and then three to four editors will work on the final draft.

Jackie Chen asked where the science and mission drivers were. Lucas answered that this report is for the Office of Advanced Scientific Computing (ASCR) and will refer to the science and mission drivers but not focus on them.

Giles asked to review the charge; it was projected on the screen. The charge letter from Patricia Dehmer asked for input from the community. Specifically, it asked for ten technical approaches that will allow the exascale goals to be achieved (technical barriers and credible technical approaches to overcome those barriers). It also asked for a justification for and description of the expected impact on overall system performance for each topic.

John Negele noted that the long discussion of codesign early in the report does not respond to anything in the charge. It is distracting. He also questioned how the value of 20 MW was decided upon as a limit to power consumption. Lucas answered that the codesign chapter will be trimmed significantly and may be moved to the end of the report. Sarkar noted that there are other cross-cutting topics beside codesign. Dongarra said that the Subcommittee had discussed having a table of such cross-cutting topics that showed how each was related to each challenge. Giles stated that the top ten should be crisp and appear near the beginning of the report. The outline, as stated, covers all the concerns of the community. One could start by talking about the science produced and then go on to the top ten challenges, or one could go the other way around. The report has a broad readership. Hey said that the crosscutting themes should go at the end or as a subset of some of the top-ten topic discussions.

Sarkar said that this is an unusual charge. One should get to the point (the top ten) right away. Snir noted that there are areas that will be confronted by DOE, but others will be addressed by industry. He asked if this distinction were reflected in the report. Lucas said that the report is not being written in a vacuum, but it does not cover others' agendas.

Colella asked what things will drive ASCR's research agenda.

In response to the question about the 20-MW limit, Lucas replied that Bill Harrod had set the upper bound as a limit to what a facility could be expected to provide.

Simon noted that, about 4 years ago, there were three cross-cutting workshops. He asked what had been learned and what had changed. Lucas responded that the report does not discuss that explicitly. The recommendations from those workshops were adopted in this report. Nothing radically different is said here.

Ang stated that the Subcommittee believed that codesign underpinned *all* of the topics. There is support on the Subcommittee to move the discussion of codesign to the end of the report.

Giles said that the impact of codesign in these various technologies will be addressed. What ~~need~~ to be established is why DOE should work on this problem and how this effort

needs

contributes to DOE's mission. Codesign underpins that, and cooperation with industry serves the purposes of the Office and the Department. The report has to justify and describe why DOE should do it rather than having industry do it.

Jackie Chen asked whether codesign can be discussed both briefly at the beginning and in detail at the end. Huntoon commented that that would give the right level of detail. Sarkar noted that that introductory statement could also point out the relationship to DOE's mission.

Yelick asked if there is time to investigate these topics to fit into the exascale timeframe. In response, Lucas questioned what timeframe would be adopted. If the target were 9 to 10 years, the answer would be yes; if less, then no. The Subcommittee has talked about timelines, and it could adopt one, but the report does not have one (or a cost estimate) right now. Giles stated that his first inclination would be to say that ASCAC should steer clear of cost and timeline considerations. A timeline needs to reflect how these technologies fit into the path to the exascale. Lucas noted that there might be multiple potential outcomes (e.g., in the amount of power available). There is a lot of intellectual risk. Colella asked if that risk were assessed in the report. Lucas answered, no. A shot could be taken at it if ASCAC so desires. Giles pointed out that ASCAC was not asked to prioritize this list. One could assess progress that has been made since the early exascale reports. Sarkar said that the charge letter asked for credible impacts, and those impacts could touch on technical risk.

Jackie Chen said that a timeline is needed. These things have to occur in a logical order. A roadmap/timeline is needed for development.

Harrod stated that feature creep needs to stop for this report. Feature creep kills projects. Giles commented that a roadmap should not be prepared, but the need for a timeline should be noted. ASCAC is not in a position to assign dollars, either. Sarkar noted that the Subcommittee needs to be timely, also. Heavy editing will be needed to synthesize the large amount of contributions from the community.

Vincent Chan asked if each topic could be commented on in terms of effort versus impact. Giles said that that gets into feature creep. Lucas stated that the Subcommittee has already discussed these issues, and those discussions should not be reopened.

Shalf pointed out that the pervasive nature of the software environment is what makes the codesign effort so difficult. Giles said that that discussion is also broader than this charge.

William Harrod was asked to give an update on the ASCR exascale effort.

There has been a government shutdown, the budget has been squeezed, the Secretary of Energy has been briefed, and the Secretary called for a "real plan." The Office's staff has interacted with the Secretary of Energy Advisory Board (SEAB), which will meet on December 2–3. Giles noted that ASCAC will send a letter to that Advisory Board. Harrod continued that the Board will look at approaches, timelines, etc. and then make a recommendation to the Secretary. The Secretary has asked why it was believed that the exascale effort would be successful. The assumption by everyone involved is that it *will* be achieved, but a collection of approaches is needed to overcome all the barriers. The charge has already grown in scope. "What's next" might be a good appendix to the top-ten report. The top challenges need to be mapped onto a timeline. Earlier reports have shown that increased computing power is definitely needed. However, a 1000-fold increase in computing power will not be achieved by a business-as-usual approach.

Data plays a major role in computing. A strategy for moving forward includes (1) exascale R&D; (2) extreme-scale applications (developed before the computer shows up); and (3) facilities, including an “on-ramp” to exascale computing.

A standard definition of exascale computing is needed. It should include a specification of 10^{18} operations per second, a machine that is not a special-purpose computer, and a technology that sets a trajectory of progress. With a billion degrees of freedom and high resilience, productivity is important. This is not a “one-off” system. DOE’s investments should leverage the vendors’ investments and products. An exascale initiative would have an end date of 2020 at the earliest.

The challenges include parallelism/concurrency, reliability/versatility, energy efficiency, and memory/storage. System design issues include scalability, efficiency, time to solution, and reliability. Productivity issues include management of system complexity, portability of software, and generality. If one could do an exaflop in 500 cabinets, one could do a petaflop on a 19-inch rack! Codesign is essential.

A timeline has been developed for systems acquisition, application development, and R&D [codesign, FastForward, Design Forward (systems design and prototype build), and software technology]. DOE is going back to the vendors and other community members to see if this storyline makes sense.

One needs to have a goal perspective. Codesign centers have made nontrivial progress toward influencing hardware architectures.

FastForward is small-scale (\$62 million) jointly funded by the Office of Science (SC) and the National Nuclear Security Administration (NNSA) to initiate partnerships with industry players.

Design Forward is jointly funded by SC and NNSA with \$25 million in contracts out to AMD, Cray, IBM, Intel, and Nvidia.

A lot of investment has been made in software, including the 2012 X-stack. Modeling and simulation efforts are being funded. A start has been made on the resilient extreme-scale solvers. Four teams have been funded to look at these problems.

Sources were recently selected for Mathematical and Statistical Methodologies for DOE Data-Centric Science at Scale.

A new effort has been initiated in uncertainty quantification.

So, a start has been made on setting a new direction for future generations of computing.

Sarkar asserted that industry players will not do it on their own. He asked if there had been any discussion of synergies between (1) what is being built for commercial applications and (2) what will be needed for the exascale. Harrod replied, no.

Dongarra asked if this would be kicked off in 2015. Harrod replied that the real money would start in 2016, but it is not known what that amount will be. The Secretary will need to submit a budget for FY16 in February 2014.

Giles said that, the way the exascale is being approached, it seems that science and engineering goals are being put on a second-class status. The science needs to get done, and the exascale is needed to achieve that. The exascale initiative might build on what the commercial sector is doing, but it might not. The Department of Energy should be willing to invest in this initiative if it gets the science done. Harrod said that the argument always starts with the science. An exascale computer is not being built just to achieve the exascale.

Hey noted that hardware vendors must also be encouraged to develop the needed products. Chips must serve both the exascale and commodity applications. Harrod replied that a dialogue has been joined between DOE and the vendors as a result of the small investments that have already been made.

Snir said that the more unique the hardware is, the more unique the software will have to be. Harrod answered that there is a mass of problems in software.

The floor was opened to public comment.

Bland asked if any research would be funded in power supplies and cooling technology. Harrod replied that that area should not be forgotten. However, he did not know how interested vendors would be.

Jackie Chen asked whether the stack would be stable enough for application development by 2020. Harrod responded, no. Staff development and application development will be going on concurrently.

Steve Binkley was asked to present the view from Germantown.

Recent organizational updates in the Department of Energy include

- The establishment of the Office of the Under Secretary for Science and Energy, which will oversee seven offices, bringing in the applied-energy programs. The seven offices are Science, Fossil Energy, Energy Efficiency and Renewable Energy, Nuclear Energy, Electricity Development and Energy Reliability, Indian Energy Policy and Programs, and Technology Transfer)
- The reconstitution of the Secretary of Energy Advisory Board (SEAB), which had a fact-finding exercise recently
- The establishment of the Office of Energy Policy and Systems Analysis to do the *Quadrennial Energy Review*

Franklin Orr has been nominated as the Under Secretary for Science and Energy, and Marc Kastner has been nominated as the Director of SC. Also, Susan Graham has been nominated for membership in the President's Council of Advisors on Science and Technology (PCAST). Daniel Lehman will be retiring from his leadership position in Project Assessment by December 31.

There is a lot of latency in getting budget actions through Congress. The average of the House and Senate marks for ASCR could be close to the President's request. The Budget Control Act gives very strict guidelines for funding. ASCR's exascale proposals have garnered favorable opinions from Dehmer and Moniz. If there is a meaningful response from the Senate and the House, the FY14 budget would likely be about the same as that for FY12, which was a good year for ASCR. The Office currently has funds to carry it through mid-January of 2014.

In big data, ASCR has gotten the attention of the Office of Science and Technology Policy (OSTP) and the Office of Management and Budget (OMB). ASCR needs to set its plans for big data. Open data is the subject of policy statements from OMB and OSTP. The Office's staff is trying to figure out how to respond to these statements in terms of their application to proposals etc. ASCR's proposed data-management plans are being reviewed by OMB.

The applied-energy programs will need to be brought to the table for discussions about advanced computing. The Office of Nuclear Energy (NE) has a well-thought-out plan for advanced computing. Fossil Energy's applications and plans need to be developed for the FY15 budget.

Hey stated that big data goes along with machine learning (including data mining and other techniques). ASCR does not place any emphasis on machine learning. Binkley replied, so noted. It is a complicated area and requires a step up from the numerical analysis that ASCR has traditionally pursued. Harrod contended that machine learning *is* a major topic in ASCR.

Giles noted that the SC contribution to the facilities report mentioned a future Virtual Data Center and asked if there were still plans for such a center. Binkley replied that it is still on the Office's list, and how to do it will probably be figured out in the next 6 months.

Hey said that open access to the literature is also of interest to OSTP. Binkley said that that topic has been widely discussed across the government. The Office mentioned it in its reply to OSTP. The publishers are going to self-organize a repository that will make publications downloadable after 12 months. That seems to be okay to OSTP. During that 12-month embargo, DOE may make the final (accepted) manuscript available online.

Vincent Chan asked what the Department's position was on open access to computer codes. Helland replied that that topic will be addressed in ASCR's open-source policy.

Giles asked what role ASCR will play in the education of the next generation of mathematics and computer-science students. Binkley replied that OMB has been consolidating education programs (including fellowships) across the government. That has left mission agencies with a problem. The mission agencies are trying to regroup. The Secretary has told the offices to move forward with traineeship programs for mission-needed skills. Such programs will address some of the problems. Giles noted that there has been a useful means of meeting the needs for computer scientists for 20 years. **Computational** Helland interjected that the training programs would be modeled on the National Institutes of Health's programs; the money would go to universities. Giles asked what was happening to ASCR's Fellowship Program. Binkley said that carry-over money from FY13 will fund another cohort of fellows. Then we will see what OMB does for the next year.

Ray called attention to the effort in machine learning in the Office of High Energy Physics (HEP). ASCR's interest is focused on the need for data management by other offices. Binkley agreed that the bulk of data in SC comes from big experimental machines, where only a small portion of the data is collected. What is being observed at the Large Hadron Collider is a recognition that the techniques and facilities of ASCR may be more cost-effective to them than doing it on their own.

The floor was opened to public comment. There was none.

Giles initiated a discussion on the exascale effort and the top-ten report. He proposed that the Committee accept the draft top-ten report, including the list of ten challenges, as presented at this meeting. Sarkar seconded the motion

White asked whether there would be discussion about the exascale budget. Giles replied, yes, sometime in the future. Hey added that this Committee should not give the Subcommittee more requirements; the report needs to be finished by November 30.

Huntoon noted that this is a multistep process. Not everything can be put in this report. This report is a step of a process, not the whole process.

The question was called, and the motion passed unanimously.

Giles solicited Committee members' opinion on what efforts and activities ASCAC should pursue. Specifically, he asked members to comment on where they thought the exascale effort currently was.

Huntoon said that one has to put the proper and complete context around the top-ten report.

Sarkar noted that the exascale effort was behind schedule, which is disappointing because it is slowing the progress of science. The efforts of vendors are being leveraged. The same approach has to be taken to the exascale as in the case of other science facilities. It is easy to justify the infrastructure needed. Some specialized systems are needed for science missions, but one needs to be aware of hardware and software trends. Exascale delayed is exascale denied.

Hey commented that the exascale has an interesting agenda with a lot of technological challenges. The program makes it feasible to make commodity petaflop systems. A timeline and budget are needed.

Jackie Chen said that this program should not be delayed further. Codesign should be extended to additional algorithmic spaces in an integrated way.

Chapman noted that the excitement and anticipation of the exascale have been great. The community is looking to ASCR for leadership. More algorithms and applications need to be made exascale-enabled.

White said that good progress has been made in identifying problems but not in how to solve them. This is a big initiative; new funds are needed; and the size of the beast needs to be made known. Other projects (i.e., the International Linear Collider) have been more expensive than at first envisioned. One needs a budget and plan, not just talk about it.

Negele said that Sec. Moniz has configured DOE to place more emphasis on *energy*. ASCAC and ASCR need to emphasize the impacts of the exascale on each energy program. There is still more room for scientists to make key contributions to the effort [e.g., like Norman Christ did with quantum chromodynamics (QCD) computing, in which several capabilities were put on a single chip]. The scientists need to be engaged in pursuing a path to a machine.

Vincent Chan said that exascale computing needs to be moving to a new paradigm. Universities will not be teaching that paradigm anytime soon. Specialized training will be needed, and that training should be integrated into the top-ten topics. Balance will be needed, also.

Glotzer, as well, was surprised not to see education and training. The uncertainty about the fellowship program does not bode well for the next generation.

Williams believed that the list was good.

Petzold said that the need for science drives the need for the exascale; that idea needs to be made paramount.

Meza offered two issues that could be mentioned in the report: Algorithm developers need to be given access to exascale computers, and no mention was made of problems with mixed discrete and continuous variables.

Juan Meza was asked to report on the Applied Mathematics Committee of Visitors (COV).

The charge to the COV was

- To assess the efficacy and quality of the processes used to solicit, review, recommend, and document proposal actions and those used to monitor active projects and programs and
- To comment on (1) how the award process has affected the breadth and depth of portfolio elements and (2) the national and international standing of the program with regard to other applied-mathematics research programs.

The COV found that the Applied Mathematics Program is highly effective in its processes to solicit, review, recommend, and document proposal actions. The Applied Mathematics Program managers do an excellent job of monitoring all aspects of their portfolios. The overall breadth and depth of the Applied Mathematics portfolio is excellent. ASCR has a long tradition of supporting some of the best applied math research in the nation and also maintains an international leadership position in certain key mathematical areas.

The COV made a series of recommendations. In regard to the management of solicitations and awards, the COV recommended that

- Program managers should continue to look for ways to enhance the program's ability to attract new investigators.
- Program managers should be allowed to travel, as needed, to scientific meetings.
- The Computer Science Graduate Fellowship Program (CSGF~~F~~) should be expanded, and its funding should be doubled during the next 5 years; also, the program should remain within ASCR.
- Award rates for Applied Mathematics Program solicitations should be made publicly available so people would know what their chances are.

In regard to the monitoring of projects and programs, the COV recommended that

- There should be an annual meeting of the directors of the three Mathematical Multifaceted Integrated Capabilities Centers (MMICCs) to highlight technical achievements, share open problems to enable opportunistic collaborations, and share lessons learned about effective center management.
- The use of a standard reporting format should be instituted for the annual progress report to monitor the activities and achievements of the centers.

In regard to the breadth and depth of the portfolio elements, the COV recommended that

- A short-term business program should be developed at the MMICCs to enable promising researchers to develop collaborations with center members.
- A new interdisciplinary program of applied mathematics–statistics–computer science–facilities should be added that could drive the next generation of fundamental research applicable to the analysis of experimental/observational facilities' data.
- Outreach efforts should be continued to professional societies and research communities.

In regard to the national and international standings of the portfolio elements, the COV recommended that

- The review and annual reporting process for the MMICCs should include a listing of awards and accolades received by the project participants with brief summaries indicating the associated technical achievements; this report can be used to highlight the leadership role of the MMICCs.
- The Applied Mathematics Program should develop a set of key mathematical areas that will have the greatest impact on the DOE mission and in which they can either currently claim or plan to develop international leadership.

Giles commented that Meza had mentioned new key mathematical areas, and asked if those areas included machine learning. Mesa replied, yes; machine learning would be very useful.

Hey noted that the COV recommended additional funds for program-manager travel and asked if that helped the Office. Binkley replied, yes.

Vincent Chan asked if there were a chance for proposers to respond to reviewers' comments. He also stated that the Fellowship Program should be retained in ASCR because ASCR has the best knowledge of research fellows' placement. Meza replied that he did not believe that the COV would object to strengthening the recommendation about the ASCR fellowship program.

Giles asked if the COV were willing to accept that amendment. Meza responded, yes.

Giles called for the sense of the Committee on accepting the amendment of the COV report as made by Chan. The Committee was unanimous in accepting the amendment.

Giles moved to accept the amended COV report. Hey seconded the motion. A voice vote was held of those present in the room, and those members on the telephone were polled individually. The motion passed unanimously.

Giles opened the floor to new business. There was none. He opened the floor to public comment. There was none. He noted that the next regular meeting would likely be in March 2014. Christine Chalk stated that there would be a telephone vote on accepting the final Top Ten Report during December. Binkley noted that the SEAB will hold one or more deep-dive meetings in January to assess the future of the exascale program. The Top-Ten Report will be needed for those meetings.

There being no further business, Giles adjourned the meeting at 4:26 p.m. MST.

Respectfully submitted,
Frederick M. O'Hara, Jr.
Recording Secretary
Nov. 22, 2013