Draft Minutes
Advanced Scientific Computing Advisory Committee Meeting
Feb. 26-27, 2008, American Geophysical Union, Washington, D.C.

ASCAC members present:

F. Ronald Bailey
C. Gordon Bell
Marsha Berger
Jill P. Dahlburg, Chair
Roscoe C. Giles
James J. Hack
Thomas A. Manteuffel
Horst D. Simon
Ellen B. Stechel
Rick L. Stevens
Virginia Torczon
Robert G. Voigt, Co-Chair
Thomas Zacharia

ASCAC members absent:

David J. Galas

Also participating:

Melea F. Baker, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Christine A. Chalk, ASCAC Designated Federal Officer
Deborah Frincke, CyberSecurity Chief Scientist, Pacific Northwest National Laboratory
Michael Heroux, Numerical and Applied Mathematics Department, Sandia National Laboratories
Barbara J. Helland, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Daniel A. Hitchcock, Senior Technical Advisor, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Michael Holland, Office of Management and Budget
Gary M. Johnson, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Robert Lucas, Information Sciences Institute, University of Southern California
Barton Miller, Computer Sciences Department, University of Wisconsin
Frederick M. O'Hara, Jr., ASCAC Recording Secretary
Raymond Orbach, Under Secretary for Science, USDOE
Walter M. Polansky, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Michael R. Strayer, Associate Director, Office of Advanced Scientific Computing Research, Office of Science, USDOE
Katherine Yelick, Computer Science Division, University of California at Berkeley

About 40 others were in attendance.

Tuesday, February 26, 2008

Preliminaries: Chairperson Jill Dahlburg called the meeting to order at 8:58 a.m.

Raymond Orbach: Overview of the Office of Science

The FY08 Omnibus Bill was tough on the Office of Science (SC) but not so much on the Office of Advanced Scientific Computing Research (ASCR), whose FY08 appropriation was $11 million more than the President's Request. Congress was sending a message and giving a vote of confidence in ASCR's efforts. The Earth Simulator gave the United States a scare that it had lost its world leadership in an area in which it could not afford to be second best. Discussions have been opened about the social sciences' using DOE's computers, and DOE is working with 17 companies through the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program to open high-performance computing to industry. The leadership computing facilities (LCFs) are fully supported; the Oak Ridge LCF is reaching one petaflop, the Argonne LCF 500 teraflops, and the National Energy Research Scientific Computing Center (NERSC) at least 120 teraflops. The Energy Sciences Network (ESnet) is continuing to upgrade its high-speed optical networks; this critical capability will soon be operating at 40 Gbps, linking national laboratories, companies, universities, and states. It should soon be possible to integrate economics and social-behavior into the climate models, allowing policymakers to use the codes for simulations. The plan is to continue to make the LCFs available to the very best science through INCITE, nurture critical applications through Scientific Discovery Through Advanced Computing (SciDAC), provide direct support for leading-edge research groups, expand computational speed and productivity to the exascale, and work with key science applications to identify exascale opportunities. This program will require restructuring ASCR. With the FY08 Continuing Resolution, SC lost $0.8 billion from the President's Request. It is a warning not to take the budget for granted. The FY09 request is $800 million more than the FY08 appropriation and would restore funding for high-energy physics and the International Thermonuclear Experimental Reactor (ITER). Another year of budget cuts would be a disaster for SC. Science has Congress's support, but Congress needs to hear from scientists.

Discussion: Planning for the exascale, which has gotten here sooner than expected, indicates a downturn in funding in 2010 and then some upturns. Needed investments need to be identified. The LCFs and the Defense Advanced Research Projects Agency (DARPA) initiative have received good support; they will determine the architectures that will be used at the exascale. The United States will profit enormously from the international competition and collaboration being brought about by developments in high-energy physics. Blue Gene P will soon be surpassed by Jülich, and the INCITE program has two participants from abroad. The United States should take advantage of the initiatives of other countries. The codes of other DOE offices [e.g., climate, quantum chromodynamics (QCD) on a lattice, ITER, and high-temperature materials] depend on ASCR facilities. The ability to model across many scales is crucial to all the sciences.

Administration: A break was declared at 9:46 a.m. The meeting was called back into session at 10:05 a.m. It was pointed out that a report from Brown was now available, completing the work that Brown had reported on at the past two ASCAC meetings.

Michael Strayer: Overview of the Office of Advanced Scientific Computing Research

The ASCR funding profile is positive and increasing from the FY07 appropriation to the FY08 appropriation to the FY09 request both for Mathematical, Computational, and Computer Sciences Research and for High-Performance Computing and Network Facilities. A balance is needed between research and facilities. The ASCR appropriation for FY08 is $354.4 million, including $19.5 million to continue participation in the DARPA High-Productivity Computing Systems Partnership and an increase of $7.7 million for the Oak Ridge LCF. A new element is the establishment with the National Nuclear Security Administration (NNSA) of the Institute for Advanced Architectures and Algorithms with centers of excellence at Sandia National Laboratories and Oak Ridge National Laboratory for early investments in exascale software. The

FY09 budget request includes a new Joint Math-Computer Science Institute to develop algorithms for multicore computing facilities, support for leading-edge development teams, and a new petascale applied-math effort in the mathematics of large data sets. SciDAC will be expanded, and investments will be made in Next-Generation Networking for Science. ESnet will be upgraded with Internet2. A substantial sum will be invested in prototype architectures under partnerships with IBM and DARPA. Out of this will come computing with maybe tens of millions of cores. ASCR's research will focus on harnessing the potential of petascale systems and data sets, enabling new scientific applications requiring radically new algorithmic and computational approaches, building on lessons learned at the petascale, and developing new tools for handling data. Current investments in mathematics, libraries, and tools will be expanded, and new investments will be made in uncertainty quantification, approaches for managing large systems, operating systems, file systems, I/O issues, and R&D prototype testbeds. Computational partnerships will be increased in breadth through INCITE. Distributed networking will require a new focus on usability, inter-domain tools for managing federated optical networks, and innovative solutions to data sharing and workflow issues. Research-and-engineering-prototype investments will ensure that the development of high-performance computing resources continues and that the research challenges are understood. New data-facility initiatives will support the exponential growth of scientific data from observations and experiments. Ten new federal positions in ASCR have been approved. INCITE is making available 265 million processor hours for 55 scientific projects; these numbers will be increased in future years. Vern Paxson of Lawrence Berkeley National Laboratory (LBL) won the Grace Murray Hopper Award, and James Sethian of LBL was elected to the National Academy of Engineering.

Discussion: The review and study of algorithms are very complicated, especially at multiple scales. The fundamental equations need to be examined. One has to have an architecture in mind in which to solve them. The virtual institute will include postdocs, fellows, a strong university focus, and a large laboratory participation. The startup date will depend on the FY09 budget. The priorities have to be pursued: complex mathematics, multiscale, and algorithms require study and support. If funding does not appear in the FY09 budget, it will be written into the FY10 budget. About 27% of the INCITE teams and 30% of its cycles reflect SciDAC programs. Not all SciDAC programs are reflected in INCITE; there are other interests that are supported by the SciDAC program.

Ellen Stechel: Subcommittee on the Role and Efficiency of Networking and Networking Research

The network now underlies nearly every aspect of science. Simulation and modeling have become the third leg of scientific inquiry and constitute a major driver for cyber-resources. Scientific inquiry is moving toward more quantitative understanding of ever-more-complex systems, requiring data packaging, access, and mobility. The production of raw data and its archiving for reuse is increasing. A continuing trend is toward global collaboration, requiring data communication and data mobility. New cybersystems that combine and integrate existing facilities require a system-of-systems approach to scientific infrastructure. As a result, ESnet is critical to the conduct of science and is doing an exemplary job in architecting, deploying, and operating a high-performance network. Network capacity and service capabilities will likely increase by 3 or 4 orders of magnitude. DOE cannot depend solely on the commercial sector for its network needs. Promising network research concepts must be identified and moved through prototyping and experimental deployment to a production environment. The networks and networking services developed by a few leading-edge science communities must be leveraged and generalized to serve the broader DOE science community. The DOE science community should think broadly about how an unconstrained network resource might transform the conduct of science. Mechanisms should be established to manage the issues that arise from the system of systems. An external ASCR community should review the networking research program on a regular basis to maintain a 10-year research horizon and integration across SC. A deliberate strategy should be formulated to move concepts through testbed deployment to production. Next-generation networks will require end-to-end monitoring of the network, automated methods for monitoring and correcting the network, and higher-level services to manage resource utilization and workflow coordination. Data collection, archiving, curation, generation, pedigree assignment, and access should be researched in an integrated manner.

Discussion: Like ASCR, ESnet focuses on scientific impacts. The ASCR networks will build upon the sensor networks of the Office of Biological and Environmental Research (BER), which will be sources of large amounts of data. The concept of software as a service is being considered. The Subcommittee sees a continuing of the ESnet strategy to connect certain services and gatewaying to the rest of the universe. The network will become more involved and less invisible, it will require data standards, it will largely be a data handler, the networking services will need to help handle those data, and the network program will need to be closely integrated with those services. Data science also calls for and must be integrated with data management. No recommendations on funding were made by the Subcommittee; just strategy was recommended. There have not been many gains in networking in more than a decade. A number of research challenges need to be addressed, and not much research is being done by ASCR. The National Science Foundation (NSF) must be brought into this process.

Action: A vote on acceptance of the report resulted in 11 votes for and one abstention.

Ronald Bailey: Fusion Simulation Subcommittee

A panel has been selected that includes members from the fusion energy, mathematics, computer science, and facility and infrastructure communities. The panel is looking at what is critical and comparing those items with what ASCR is doing and can do. It will then consider how to mobilize different disciplines, laboratories, and institutions. It will be looking at facilities and infrastructure. It is holding weekly teleconferences and is collecting data. It will hold a workshop in April.

Discussion: The INCITE Committee of Visitors is empanelled. It has set a date in April to meet with DOE, the computing centers, and the user community. It will review the proposal-evaluation process.

Administration: The August ASCAC meeting will be held in Berkeley, California, on August 5-6 with optional activities for Committee members on August 4. The meeting was adjourned for lunch at 11:46 a.m. and reconvened at 1:00 p.m.

Barton Miller: Tree-Based Overlay Networks for Scalable Middleware and Systems

Today’s extremely large-scale systems are difficult to program in a scalable, effective manner. Tree-based overlay networks (TBONs) provide an immediate path to scalable tools and infrastructure and a research platform for new technologies. The front end of tools and middleware for scalable systems becomes a bottleneck both for control operations and data collection and processing. The problem can be decomposed using a tree-based overlay network (TBoN), with multicasting occurring down the tree and data reductions moving up the tree.

MRNet [Multicast/Reduction Network] is such a TBoN, allowing one to plug in filters (reduction operation) anywhere in the tree, and a computation gets done in log time. A number of applications now use TBONs for overcoming tool-startup latency or clock skewing by distributing performance data and control across the network and deriving global results by merging local data. With a system such as MRNet, simple and lightweight scalable tools can be more easily developed. Currently, a debugging tool (STAT) is being developed, in partnership with LLNL, to gather multiple stack traces over, currently, 100,000's of nodes and, soon, to millions. STAT merges these stack traces in an efficient and scalable way, providing debugging analyses and visualizations that allow a problem to be identified and fixed. STAT trace results can now be obtained from the largest of leadership class systems (BG/L) in less than a half second. MRNet also provides a basis turning almost any MRNet-based application or tool in one that is fault tolerant. The fault tolerance is based on the inherent redundancy present in such a tree-based computation, and on a new weak data consistency model.

Discussion: Comparing MRNet to Map/Reduce: Map/Reduce forms a simple but restrictive idiom for distributed computing. While MRNet offers a much more flexible substrate for such computations, handling finer-grain computations and high-throughput data processing, MRNet is not a programming idiom. Yahoo's new Pig Latin language, augmenting distributed computation with relational semantics, offers a better and more interesting alternative. Miller hopes to work with Yahoo to compile and map Pig Latin programs onto MRNet. Also in the MRNet new activities is an effort with TotalView to use MRNet to build a truly scalable version of their debugger.

Robert Voigt: Balance Subcommittee Report

ASCAC was charged with assessing balance within ASCR between facilities and research. The Subcommittee found that leadership in high-performance computing must be maintained; high-performance-computing hardware is becoming a worldwide commodity; access does not ensure scientific competitiveness; effective use must be made of high-performance computing to gain leadership; and increased computational complexity requires research in algorithms, tools, and software. Investments in research must be made to ensure that investments in LCFs realize their intended goals. Since FY05, investment in facilities has outstripped that in research. It is critical that ASCR maintain healthy LCFs; access to capacity computing is important. As the transition is made to the exascale, the LCFs may need to be upgraded and should be recompeted in 7 to 10 years. At that time, consideration should be given to migrating the current LCFs to capacity facilities. Architectural advances require increased investment in research into mathematics and computer science. ASCR should support core research that can advance scientific discovery, with some fraction of the research portfolio addressing new applications. In this effort, SciDAC has been very successful. It is recommended that a 50-50 balance be maintained between research and facilities. The Centers for Enabling Technology (CETs) should be evaluated on the basis of their productivity, and nonproductive centers should be pruned out. INCITE is very successful. A balance needs to be maintained in INCITE between scientific importance and code efficiency; currently, it favors LCF-ready applications. SciDAC-like groups should be formed for important applications, and the process for SciDAC access to LCFs should be streamlined. Investments in facilities must be made to stay in the international competition, but investments in research must be made to win that competition. [A break was declared at 2:24 p.m. The Committee was reconvened at 2:37 p.m.]

Discussion: The CETs need to be evaluated to see if they are producing the tools needed. The question was raised *why* a balance is needed among all these programs. The 50-50 split between research and facilities seemed to the Subcommittee to be a good number, but it is an arbitrary number. More "seed corn" might be needed. There is an opportunity to bring along programs in other offices, requiring more money in developing good computational algorithms for large-scale science. The emphasis should be on techniques that solve problems accurately and efficiently. Generally, the highest-risk projects require more development time to make efficient use of LCF. Every now and then, though, one should step back and see where the project is going. The intent of calling for milestones is not to discourage long-term funding for projects. With the devaluation of the dollar, a balance must be struck in money as well as in time and space, and all three need to be tracked. It is not obvious how to do research on the million-processor problem. A research portfolio tied to objectives was included to deal with that problem; the outputs of that portfolio can be tracked.

The terms "capacity" and "capability computing" are overused and should be abandoned; the terms "high-tech science" and "broader-tech science" should be used. NERSC is classified as a capacity facility but is running high-tech science. Facilities should be measured on the quality of the science performed, which is outside the control of the facilities. Costs have to be weighed against expected gain.

People who are running a production code do not want to take time out to rewrite that code for a new platform. It is important to engage the computer science and mathematics communities in the development of applications; then they "learn the applications language" and bring it back to the computer science and mathematics students and practitioners.

Balance cannot be static in a dynamic field. Evaluations of balance should be conducted over a period of time, and not with a snapshot. High-performance computing as a commodity is also a snapshot assumption. The data are noisy and should not be evaluated from a discrete data point; a discrete data point does not establish a trend.

This report implies that machines that are not top machines are unusable. Actually, a lot of science is done on such systems. The report should emphasize that, if there were no support for research, the future would not be good. One cannot give up older machines, but computer scientists have been porting their codes to the new machines because they are the machines of the future. That is why the report calls for more SciDAC-like projects to make use of the computing capabilities available. Some SciDAC-like activities are lumped into facility costs and do not show up in budgets (e.g., support funded out of operations to make science occur on Day 1).

A 50-50 balance is not the Golden Ratio. What is wanted is to build leadership-class capabilities, which was not the mission of ASCR 7 years ago. What should be considered is the demand for LCFs, how oversubscribed they are, their potential for attracting more research dollars, and what needs to be done to build out to the exascale. ASCR may need to fund architecture research when the DARPA project ends. A quantitative roadmap is needed. Staying on the Moore's Law path would be very expensive but must be done to attract partners and users.

Administration: In producing the balance report, the Subcommittee surveyed a lot of people. It also sent out six letters asking for 10-page documents. It got two or three such documents back. There was, therefore, not a lot of data in the report. There is now a draft report that needs to be changed. Each Committee member needs to identify what needs to be done to improve the report during the following two weeks. "Capacity" vs. "capability" needs to be standardized in usage. The Committee members were asked to forward suggested changes to the

Subcommittee chairs, and the Subcommittee and its leadership will consider the comments, decide what to do with the report, and report back at the August meeting.

Michael Heroux: Solver-Algorithm R&D for Scalable Applications and Architectures

Trilinos [“a string of pearls”] is a community project for algorithm development that targets all platforms, from a desktop to high-performance computers. It includes a core set of vector, graphics, and matrix classes that allow package developers to focus only on things that are unique to their packages. The Trilinos package includes discretizations, methods, core capabilities, and solvers. New packages will contain thread programming support, full multiprecision support, and mixed-precision algorithms. Trilinos needs to adapt industry techniques, but life-cycle models are generally developed from the point of view of business software and give little consideration to algorithmic development or parallel computers. Trilinos uses a promotional model that has three phases: research, production growth, and production maintenance. The research phase employs a project plan, configuration control, peer review, proof of correctness, and documentation. New node architectures (e.g., Clovertown, Barcelona, and Niagara2) offer opportunities (e.g., hybridization of multicores and multithreads) and challenges (e.g., bandwidth penalties and data-placement issues), but the Trilinos project has prepared for them from the beginning, with the majority of solver codes transparently porting to new nodes. Trilinos 9.0 will enable a variety of node-programming approaches for new algorithm development.

Discussion: The best algorithm work is being done in the applications. Some multicore techniques are quite old and are being reinvestigated.

James Hack: Report on Computational and Information Technology Rate Limiters on Climate-Change Science

A teleconference and a Subcommittee meeting were held in October 2007, and a preliminary draft report was presented to ASCAC and the Biological and Environmental Research Advisory Committee (BERAC) in November 2007. The bottlenecks to progress in climate-change science were seen to be computational solutions to requirements, software needs, algorithm needs, data management, networking, and collaboration. Model deficiencies need to be identified; most data come in small scales that are accumulated. Algorithms are needed to scale that data up. The next Intergovernmental Panel on Climate Change (IPCC) cycle will have significantly more complex models placing an enormous demand on operation counts. Examples of rate limiters are the initialization of the ocean, resolving the atmosphere, and introducing flexibility to the system. There will never be enough observational data to close the loop; modeling will be needed to do this within a well-managed, end-to-end enterprise. The current 1000-processor facilities are inadequate; new facilities are needed for data management, migration, and analysis. The allocation process may be suboptimal. The Subcommittee recommended:

- Continuing investment in facilities that are dedicated, configured, and managed to support integrated and multifaceted climate research and prediction
- Developing computational algorithms and scalable software to accelerate computational climate-change science
- Developing computational and theoretical foundations for new modes of climate simulation
- Pursuing a robust predictive capability of lower-probability/higher-risk impacts
- Developing a strong scientific understanding of leading-order carbon-cycle uncertainties

Discussion: Only minor clarifications have been made in the report since November. There have been discussions at the National Center for Atmospheric Research (NCAR) about incorporating economic and social behavior models into climate-change models. The chemistry issues are very slowly evolving, but they might not need to be coupled very closely; the same might be found for social-behavior models. If climate could be predicted on a regional scale, climate change could be related to crops. Solar variability is reflected in the upper-atmosphere calculations and chemistry and can be included in models. If more money were available, one should use it in a balanced approach; scalable algorithms, data management, etc. all have to be addressed. The time between the development of a successful application at an LCF and its use as a production code depends largely on where one truncates the resolution. One has to determine where the point of diminished returns is. One can parameterize and do one's calculations on small machines, but a lot of calculations are still needed. The turnaround rate is a couple of weeks for a simulation. As one builds higher-facility models, there is a reiteration of the model's design. It is difficult to determine where the optimum level is. The data-management needs of climate change are specialized. Remote-sensing data need to be introduced to a model during running, which is I/O intensive and sets the pace of those processes. All these data are available on the Earth Sciences Grid. The data sets used in the simulations are available for everyone to look at. Climate.com in the United Kingdom is using screensavers to do the climate simulations (a brute-force method) used to probe the strengths and weaknesses of the model. The models are available, also. Bundling INCITE runs may not be possible in scaling up some scientific questions. However, people are not getting enough computer time; the allocation process causes double jeopardy: getting the science approved *and* getting the cycles approved separately. This report will be a foundation stone on how science is performed in the United States.

Administration: The report was accepted unanimously. The floor was opened to public comment. The suggestion was made to Committee members to participate in Science Days on the Hill. Participants get a briefing from the American Association for the Advancement of Science (AAAS) and then go out and visit their congressmen and senators as a group. This is an opportunity for people from all disciplines to talk about science to their congressmen. There being no further public comment, the meeting was adjourned for the day at 5:11 p.m.

Wednesday, February 27, 2008

Preliminaries: The meeting was called into session at 8:00 a.m.

Katherine Yelick: Programming Models for the Petascale and Exascale

In *programming models*, the major question is what to virtualize (hide or expose) to get the most out of the machine. Partitioned global address space (PGAS) languages have static parallelism (they do not virtualize processes). They support distributed data structures, one-sided shared-memory communication, control over data layout, synchronization, and collective communication. They are not the final answer for the petascale because one-sided communication is faster than two-sided, global address space can be easier to program, multithreading is necessary, principled scheduling is necessary, and the combination of dynamic load balancing with locality control has new challenges. In *software,* automatic performance tuning is accomplished by running several versions of code to see which runs fastest. Programmers should write programs to generate code, not the code itself. Auto tuning finds a good performance solution by heuristics or exhaustive search. Experience with multicore machines is limited. Explicitly managed memory is easier to tune for. MPI [Message-Passing Interface] is not a viable model for the exascale; a new dynamic model is needed for software. In *algorithms*, memory, not flops, is important. The matrix should be read once, and then multiple steps should be carried out. Algorithmic design should address the bottlenecks of latency and bandwidth.

Discussion: Computations need to be programmed as a class, not individually; and there has to be tight performance feedback. High-Productivity Language Systems have a more dynamic execution code and challenging run-time ideas. They have unsolved problems. Virtualization is avoided at the petascale and exascale, although experiments with virtualization might be tried.

Rick Stevens: Report on Joint Subcommittee on Modeling and Simulation for GTL

The Joint Subcommittee is currently revising its preliminary findings and recommendations, preparing background material for the report from the panel presentations, producing a summary to provide context for the findings and recommendations, and forging linkages to two important National Research Council reports that intersect with DOE's modeling charge. The Subcommittee found that modeling and simulation play a critical role in integrating the understanding of biological mechanisms at multiple levels. Progress in integrative modeling during the past decade has been largely driven by a small number of groups with patchwork funding. ASCR's GTL and SciDAC activities are not currently supporting the development of integrated modeling and simulation. Integrative modeling and simulation efforts are highly dependent on the curation of genomics data and databases. There are not many integrated models of microbial systems. Obstacles to producing a predictive model for organismic engineering include the lack of integrated genomics databases and computational methods for curation, extension, and visualization of data; the lack of robust mathematical frameworks and supporting software for integrating models of metabolism with those of gene regulation; the lack of multiscale mathematics and associated software libraries and tools for integrating cellular processes of disparate scales; and the need to frame all computational biology in a computational and analytical theory that incorporates evolution as the basis for understanding and interpreting the results from comparative analysis. The Subcommittee made six recommendations:

- ASCR's Program Assessment Rating Tool goal for joint modeling and simulation with BER should be modified to read: (ASCR) By 2018, demonstrate significant advances in the capability to predict an organism's phenotype from its genome sequence through advances in genome sequence annotation, whole-genome-scale modeling and simulation, and integrated-model-driven experimentation.
- DOE should develop a research program aimed at achieving significant progress in predictive modeling and simulation in DOE-relevant biological systems.
- DOE should establish an annual conference highlighting the progress in predictive modeling in biological systems.
- The modeling and simulation research program should be supported by an explicit series of investments in the modeling technology, databases and algorithms, and infrastructure needed to address the computational challenges.
- DOE should support the curation and integration of genomics and related datasets to meet the needs of modeling and simulation in areas of energy and the environment that are not well supported by NSF and the National Institutes of Health.
- DOE should work with the community to identify novel scientific opportunities for connecting modeling and simulation at the pathway and organism level to modeling and simulation at other spatial and temporal scales.

Discussion: The Committee was polled to see if any modifications to the report were necessary. Comments included: There was a milestone that DOE could not meet because it was difficult to measure progress toward the original goal. There is no mention of a database effort. BER is supportive of the report and of tighter couplings with ASCR. The current BER program investments would not meet the PART goal because of the scale of support for ASCR's activities (which is off by a factor of 10) and because of the massive shift of BER's focus to cellulosic ethanol. Only a few million dollars are currently devoted to computational issues; the objective here is to change the face of biology. During the past decade, there was a fair amount of engagement between BER and ASCR; it is hoped that these programs will become even more engaged in the future. If the two offices were to adopt these PART goals, they would be on the way to achieving a specific goal. However, biological research is driven by bottom-up influences. The report does not seem to get ASCR sufficiently away from the ambitiousness of the PART goal and needs to be strengthened. The term "community" should be clarified and should include people who are crossing over discipline lines.

Action: It was decided by consensus to make changes to the report, send an interim report to BER before its April 3 teleconference, and approve the final report in the August ASCR meeting. A break was declared at 9:52 a.m., and the meeting was called back into session at 10:05 a.m.

Deborah Frincke: Transformation of DOE Cybersecurity

A grassroots DOE cybersecurity community has coalesced and holds two to three teleconferences a month to discuss key priorities, programmatic recommendations, and the need for cybersecurity R&D initiatives within DOE. It has found a need for investment in proactive, long-term, collaborative, testable, quantifiable, and scientific cybersecurity. Its objectives are to advance the state of the science in an Open Source and transparent environment by soliciting high-level advice and counsel from the classified communities. The community is focusing on identification of transformational, long term, and potentially higher risk research directions; identifying short-term payoffs; remaining aware of and responsive to shifts in the industry; and fostering collaboration with commercial, academic, and government interests. Staying ahead of the cyber threat, rather than remaining in a "catch and patch" mode of operation, would require advances in and applications of computer science, mathematics, and computer hardware. One possible game-changing approach to transformational cybersecurity might involve investment in predicting the actions of an adversary, or to deploy self-healing and resilient systems that could automatically block, thwart, or recover from damage. The community has issued a document, Transforming Cyber Security; is producing nine white papers; is bringing up a wiki https://wiki.cac.washington.edu/display/doe/Home ; is planning for an open-science R&D agenda and program; and is scheduling a series of Town Hall meetings ultimately intended to produce 10-year research roadmap.

Discussion: Sociologists, economists, and human-factor engineers need to contribute to these discussions.

Robert Lucas: The Exascale Computing Study

To determine what research needs to be done to enable computer vendors to develop petascale systems, four meetings were held, focusing on power, memory volume, programming, reliability, and packaging. Above 100 gigaflop/s, power increases with performance, with most power going to issuing instructions. This is a pervasive problem that requires a range of solutions: different processes and circuit technologies, optics, and lower-power memory.

Extended to the exascale, systems would be orders of magnitude away from the balance between memory size and performance seen today. The programming challenge would be how to handle 10 billion threads. Programming languages used at the exascale may need to be multicore, be multithreaded, have Single Instruction, Multiple Data (SIMD) extensions, and have explicit memory hierarchy. Memories and transmission lines are already fault-tolerance protected; fault-tolerant logic is a possibility. System packages must be small; distance equals latency, and power needs to be minimized. At the exascale, concurrency will increase 4 to 5 orders of magnitude, system balance will change dramatically, the number of successful codes will decline, and facilities will have to transform to adapt. Nonetheless, an exascale computing system within a decade is plausible. A number of significant problems will need to be overcome, so DOE should address them in partnership with DARPA and NNSA.

Discussion: The number of applications on exascale machines will shrink but not go to zero. If one trades off precision for accuracy, one can solve problems with fewer numbers. Repeatability is an issue, especially in parallel situations; it would be a great topic for the DOE mathematics program to investigate. If a set of science objectives were put forward, the vendors would be able to work backwards from the problems.

Michael Holland: Overview of ASCR from the Perspective of the Office of Management and Budget

Supercomputing and simulations do remarkably well with this administration. In 1999, the fight was about “big iron,” not efficiency. SciDAC was easier to sell than big iron. Eventually, big iron came in the form of the LCFs. ASCR has had a phenomenally good 7-year run. Currently, high-performance computing presents itself with a machine focus, and the push for the exascale sounds like a push just to be number one. High-energy physics took a huge whack in FY08 for doing just that. OMB buys the best science it can, not technology. Technology is bought because it buys science. If a new machine is needed, tell OMB the sweet spot in science that this machine will elucidate. A solid science case and a solid advisory-committee review are needed. SC now has three large facilities and needs a healthy debate on what the strategy should be to move forward. The people at OMB are proxies who need to be equipped to fight for ASCR. Currently, ASCR’s program is well balanced, reasonably productive, and well presented. The Joint Institute brings together applied math and computer sciences early in the process; ASCR should decide beforehand how to measure its success. In addition, funding has been requested for cybersecurity so ASCR can develop risk analyses for complex systems.

Discussion: Multiple facilities operate light sources because physicists go back and ask the communities what they want to be able to do and what they must be able to do. Then they design the machines to answer the most important science questions. Computer hardware developers need to adopt a similar strategy. The ASCR portfolio stops short of computational science, and the community therefore does not come from a technology standpoint. ASCR could be done away with, and the work could be farmed out to other offices. That would not be a good idea. ASCR is a service organization that helps other people answer their questions. That is an important role. The scientific-applications partnership could be expanded to bring the scientists in at the front end.

Administration: The floor was opened to public comment. There being none, the meeting was adjourned at 11:46 a.m.

Respectfully submitted,
Frederick M. O’Hara, Jr.
ASCR Recording Secretary
April 10, 2008