Minutes
Department of Energy and National Science Foundation
Nuclear Science Advisory Committee
Hilton Hotel, Gaithersburg, Md.
December 1, 2011

Members Participating:

Susan Seestrom, Chair
Peter Jacobs, Vice Chair
Robert Atcher
Jeffrey Binder
Jeffery Blackmon
Gail Dodge
Richard Furnstahl
Alexandra Gade
Carl Gagliardi
Susan Gardner
David Kaplan
Dmitri Kharzeev
Joshua Klein
Zheng-Tian Lu
Allison Lung
Curtis Meyer
Robert Tribble (via telephone)
Julia Velkovska

Members Absent:

Michael Bronikowski
Karlheinz Langanke

Others Participating:

Donald Geesaman
Timothy Hallman
Daniel Hitchcock
Bradley Keister
Robert McKeown
Steven Vigdor
Scott Wilburn

Presenters in Order of Appearance:

Timothy Hallman
Bradley Keister
Peter Jacobs
Daniel Hitchcock

About 35 others were in attendance during the course of the meeting.

Thursday, December 1

The meeting was called to order at 8:32 a.m. by the Chair, **Susan Seestrom**, who announced that the floor would be opened to public comment at the end of each discussion period in the Committee's agenda.

Timothy Hallman and Bradley Keister presented certificates to the outgoing members of the Committee; pointed out the strong, successful partnership between NSAC and the agencies; and congratulated the recipients of the certificates. A special thanks was expressed for the service and leadership of Susan Seestrom. Seestrom said that it had been a pleasure working with the NSF and DOE and with their personnel.

Timothy Hallman was asked to present the perspectives from the DOE Office of Nuclear Physics (NP). He pointed out that the Office supports discovery science, illuminates the properties of nuclear matter, and provides research tools. NP operates the Relativistic Heavy-Ion Collider (RHIC), Continuous Electron Beam Accelerator Facility (CEBAF), and Argonne Tandem Linac Accelerator System. The Holifield national user accelerator facility at Oak Ridge National Laboratory (ORNL) is being closed in FY12; a sizable user community needs to be transitioned.

At CEBAF, the 12 GeV CEBAF upgrade is proceeding on schedule. Operations were started on November 19, 2011 after a lengthy project-related shutdown.

The Facility for Rare Isotope Beams (FRIB) is planned to start civil construction in May 2012. This project is going exceptionally well. Michigan State University (MSU) is highly dedicated to this project, which will leverage DOE and NP technology developments at Argonne National Laboratory (ANL), Brookhaven National Laboratory (BNL), Thomas Jefferson National Accelerator Laboratory (JLab), Los Alamos National Laboratory (LANL), Lawrence Berkeley National Laboratory (LBNL), and ORNL. The recent Lehman review supported the preparations for a combined CD-2/3A strategy (the approval of a performance baseline combined with an approval of the start of limited construction) and the advancement of MSU funding in FY11 and FY12 to accelerate the conventional-facilities design and construction. A joint user group meeting was held in August 2011.

The Californium Rare Ion Breeder Upgrade (CARIBU) project has been completed, and all project milestones have been achieved. The first measurements have been made, providing new and improved measurements of isotope masses at the Canadian Penning Trap Mass Spectrometer and decay spectra at the Gammasphere. Current work is improving the yield, transmission, and beam purity. CARIBU beams will be available for the next Physics Advisory Committee (PAC) cycle, so proposals can be considered.

A number of universities and national laboratories have been integrated into the Isotope Program production strategy. Meetings with stakeholders were held throughout the summer.

The Electromagnetic Calorimeter (EMCal) project is designed to study jet quenching in the Large-Hadron Collider (LHC) heavy-ion collisions. The project is a collaboration among the United States, France, and Italy. Its completion marked a major milestone in the ALICE [A Large Ion Collider Experiment] at the LHC. The EMCal project was so successful that DCal [Digital Calorimetry Chip] scope enhancement was possible within the original total project cost ($15.5 million). This enhancement was a collaboration among the United States, China, France, Finland, Italy, and Japan.

GRETINA (the first stage of the Gamma-Ray Energy Tracking Array) was completed successfully at LBNL.

The Cryogenic Underground Observatory for Rare Events (CUORE) experiment is making good progress at the underground laboratory in Lignaro, Italy and the Majorana Demonstrator is proceeding at Sanford Laboratory in South Dakota.

The expected neutron electric dipole moment (nEDM) limit constrains extensions of the Standard Model and baryogenesis as the possible origin of matter/antimatter asymmetry. The nEDM project is a collaboration of 17 universities and 2 national laboratories with joint DOE and NSF funding.

The first heavy-ion (lead–lead) collisions at the LHC were accomplished at 2.76 TeV in November 2010. The observation first made at RHIC that particles containing charm quarks are just as suppressed as light quarks was surprising. These measurements have been extended in heavy ion collisions at the LHC. Additionally, at the LHC, quark-antiquark bound states (quarkonium) is being used as a standard candle to understand the effective temperature reached in the collisions.

In the Atom Trap Trace Analysis (ATTA) carried out at Argonne National Laboratory, a single atom can be trapped, which allows several important applications, such as gaining information on aquifer recharge and depletion (as has been done in Egypt and Yellowstone National Park).

JLab is pursuing an exciting search for a new light vector boson, the A′, in its investigation of possible new GeV-scale force carriers, which might constitute an important category of physics beyond the Standard Model.

New RHIC data show changes in the behavior of several observables at a collision energy of approximately 30 GeV per nucleon. These changes may signal the onset of the transition from hadronic matter to quark-gluon matter.

At Oak Ridge commissioning tests are underway for the NPDgamma experiment to research hadronic parity violation. Work is also ongoing to verify the claimed discovery of several new superheavy elements at a Laboratory in Russia.

Two new elements were to be named on the day of this meeting: element 114 (Flevorium) and element 116 (Livermorium), resulting from a US-Russia collaboration supported in part by the isotope production facilities at ORNL..

A new method has been developed for the precise determination of the isovector giant quadrupole resonances in nuclei at the High-Intensity Gamma Source (HIGS) at Duke University.

The Early Career selection process is going on, as is the Lindau Fellowship application process. A solicitation is out for a new cohort of Office of Science (SC) graduate fellowships that would begin in the fall of 2012. Applications are being accepted for internships, and a Visiting Faculty Program is being initiated.

NP's vision for U.S. nuclear science includes

- Transformative discovery science on the hot quantum chromodynamics (QCD) matter, quark structure, and limits of nuclear existence frontiers;
- Four leadership-class national user facilities for U.S. science;
- Targeted high-science-impact major items of equipment (MIEs); and
- Technological advances like the development of superconducting radiofrequency (SRF) cavities for high-power linacs, a restart of americium production for commerce, and important new nuclear data on decay heat in nuclear fuel.

The Office still does not have a budget for FY12, but the difference between the House and Senate marks is only $1.88 million. Assuming the FY12 presidential budget only requested the planned profile for the 12-GeV upgrade and FRIB with the rest of the program flat funded at the FY11 level, the NP increase would have been about $50 million relative to FY11. The House mark provides an overall increase of $11.886 million over FY11, and the Senate mark provides for an overall increase of $10 million with more than $33 million going to CEBAF and FRIB. Both the House and the Senate recognize the importance of these new research tools, but both the House and the Senate marks direct that the increases for the projects be provided in large measure from the existing NP base. These funding levels constitute a major challenge for NP; it is already requiring significant cuts and terminations to the base. Despite these cuts, the projects will be delayed and will undergo cost growth for either mark.

Looking beyond FY12, the Office got a passback from the Office of Management and Budget (OMB), which will be used to form the President's FY13 budget request. The OMB guidance is very challenging and indicates funding levels for which the vision expressed in the 2007 Long-Range Plan can not be carried out. A community exercise may therefore be needed early in FY2012 to solicit further guidance on the implementation of the 20067 Long Range Plan under the new guidance. An corresponding additional community exercise across the field may also be needed to gauge the strengths and unique capabilities of NP-supported research groups, taking into account revised priorities and plans within the field.

NP continues to support a high-impact world-class research effort with world-leading facilities and research tools. It is facing very challenging budgets. The Office will work with the community to mitigate impacts and to ensure continuation of the highest priority, highest impact nuclear science research.

Gardner noted that the Holifield users will become users elsewhere and asked what plan there was so they can continue their work. Hallman replied that the Office is working closely with the ORNL management to retain scientific and technical expertise and to continue research efforts. It is a work in progress.

Blackmon asked if these budget deliberations will delay the next long-range plan. Hallman answered that the office needs advice for the short term to deal with these budgetary challenges, and the formulation of that advice will probably delay the long-range plan effort. DOE needs to discuss this issue with NSF.

Kaplan asked if it were known what impact this OMB guidance would have. Hallman responded that it was still early in the budget formulation process and we would not know impacts until the 2013 President's request is released. Jacobs asked, given the House and Senate language, what the impacts would be on the FY12 budget numbers for FRIB and JLab. Hallman said that, for FRIB, the number would be $24 million rather than $30 million; the project team believes that that change can be mitigated. For JLab, the House number would be $40 million rather than $66 million for the 12-GeV upgrade; the Senate version was $55 million. The budget would be heavily impacted under the worst (House) scenario and would be considerably less so under the best scenario, where the impact could be minimized.

Furnstahl asked about the impact of the FY12 budget on high-performance computing. Hallman said that there is strong support for those efforts in NP.

Vigdor pointed out that the long-range plan takes about six months to develop. For guidance for the FY13 budget, one would need to start work on the long-range plan in February. Hallman said that the Office could not commission that plan until the President's budget request is released. It *could* do some preliminary organization with the hope of starting the planning effort in the spring.

Bradley Keister was asked to give an overview of NSF nuclear-physics activities. He emphasized that what NSF supports overlaps a lot with what DOE/NP supports. Laboratory Nuclear Astrophysics has been transferred to the Nuclear Physics. The National Superconducting Cyclotron Laboratory (NSCL) was just renewed for 5 years at a level that would permit some use of its new reaccelerator capability. Final approval is a National Science Board (NSB) action. The Board makes use of NSAC recommendations and priorities with respect to the science being performed at NSCL.

For FY11, funding for nuclear physics at NSF was $48.5 million. The NSF appropriation for Research and Related Activities (R&RA) was down 1%, nuclear physics and most other programs were down 3%,. the program continues to manage the impact of the FY09 American Recovery and Reinvestment Act (ARRA). Grants funded by ARRA will expire in FY12.

The FY12 NSF appropriation was part of a minibus bill (as opposed to an omnibus bill) that has been passed by Congress and signed by the President. The original request for R&RA was for $6.2 billion; the House version was flat from FY11; the Senate version was down from the FY11 level. The final R&RA appropriation was $5.7 billion, an increase of 2.4%. This is the earliest appropriation for NSF in 11 years.

The FY12 request for the Physics Division would have had a 3% increase in the President's request, but research would have increased 20%. However, these amounts were not appropriated and must now be modified.

The triennial competition for Physics Frontier Centers was completed in FY11. All four centers with expiring awards were successful in receiving renewals; one new center was added: the Institute for Quantum Information and Matter at the California Institute of Technology. Additional awards were made through the Domestic Nuclear Detection Office to perform nuclear-data measurements and detector development. PIs in the NP community received awards totaling $729,000 from the Major Research Instrumentation (MRI) Program.

In future funding, in addition to MRI awards, there is an ongoing Petascale Computing Resource Allocations (PRAC) to provide testbed access for petascale code development. The Blue Waters project has had a major transition to a new commercial partner, but it is now on track for completion.. The Domestic Nuclear Detection Office will also issue a solicitation with a deadline in spring 2012.

In personnel, Kyungseon Joo will be staying with the Nuclear Physics Program for a second year. He has an IPA appointment with the University of Connecticut.

McKeown asked if $700,000 for MRI was typical. Keister responded that $1.5 million is an average value, but that value fluctuates a lot from year to year.

Geesaman asked if there were any insight about when a decision would be made on the underground laboratory. Hallman said that there is a working assumption that early-science experiments would be done at Sanford Laboratory. There is money in the budget for dewatering and Sanford operations. There are regular meetings to monitor those activities. The Large Underground Xenon (LUX) detector and the Majorana Demonstrator will go underground next year. The Long-Baseline Neutrino Experiment (LBNE) is being assessed by the Office of High Energy Physics.

Peter Jacobs was asked to present the report of the Subcommittee on Fundamental Physics with Neutrons. Because of a conflict of interest, Seestrom recused herself and yielded the chair to the Vice Chair, Peter Jacobs.

The draft report was distributed in early October, and a revised report was finished November 11. The charge called for the Subcommittee to review and evaluate the current and proposed U.S. research program, scientific capabilities, and opportunities for fundamental nuclear physics with neutrons and to make recommendations of priorities consistent with the current (FY11) level of effort as well as projected resources. The scope of the charge included the full suite of fundamental neutron physics research opportunities in the United States and internationally, and their evaluation in the broader, world-wide context of fundamental symmetry measurements that test the Standard Model.

A Subcommittee was empanelled that reflected both theoretical and experimental physics. The scientific focus was on the search for an electric dipole of the neutron, neutron decay parameters (A, a, B, b, etc.), hadronic parity violation, and the neutron lifetime. The Subcommittee had a number of open meetings and telephone conferences, presented an interim report to NSAC in June, circulated a draft in October, and completed a final draft in November.

The principal *scientific* priorities found by the Subcommittee, ranked in descending order, are:

1. The search for a neutron electric dipole moment with the nEDM experiment.
2. Continuation of the Ultra Cold Neutrons: The "A" Correlation Experiment (UCNA) to obtain improved precision on λ, the ratio of the weak axial-vector to vector coupling constants of the neutron.
3. Completion of the NPD Gamma experiment to obtain a precision measurement of the weak isovector nucleon–nucleon–pion coupling constant.
4. Investment in the Nab apparatus [for the neutron a and b parameters] with the main goal to determine λ to unprecedented precision, using a complementary observable to that of UCNA.
5. Continuation of the National Institute of Standards and Technology (NIST) experiment to perform the most precise cold-beam-based measurement of the neutron lifetime.

The Subcommittee estimated that the five high-priority initiatives might be accommodated within a constant level of effort, exclusive of MIE construction funding, though moderate additional funding may be required. This ranking indicates a priority with which each effort should be supported, in the event of funding below the constant level of effort.

Gardner asked how these priorities were arrived at. Jacobs and Lu replied that it was based on an assessment of both physics significance and technical feasibility.

The Subcommittee found that the scientific motivation for EDM searches remains as compelling as ever. In particular, a measurement with sensitivity at the anticipated reach of the U.S. nEDM experiment ($\sim 4 \times 10^{28}$ e-cm) would have a profound impact on nuclear physics, particle physics, and cosmology, even in the event of a negative result. The U.S. nEDM project is the only technical concept among various worldwide efforts that is explicitly proposed with capabilities to reach this level of sensitivity.

The nEDM collaboration has already resolved many important technical challenges and developed a first-pass engineering design of the apparatus. However, significant further R&D is needed on several issues. The Subcommittee recommended that

1. The nEDM collaboration should immediately focus the bulk of its efforts on a well-structured and strategically targeted R&D plan to address the outstanding technical issues.
2. ORNL and LANL should jointly establish a standing Technical Review Committee (TRC) to review the R&D progress and to report periodically to the management of both institutions.
3. Long-lead-time procurements should be contingent upon resolution of the major outstanding technical issues in the measurement technique.
4. The agencies should provide continued support for 2 years.
5. In the event that major outstanding R&D issues remain unresolved after 2 years, consideration should be given to discontinuing the MIE Project and re-evaluating the U.S. strategy for achieving a precise neutron EDM measurement.
6. Strong support should be provided for the NPD Gamma experiment as the highest-priority measurement in hadronic parity violation, and every effort should be made to reach the design goal, an asymmetry determination of one part in 108.

7. Support should be continued for the UCNA experiment at LANL to improve the measurement precision of the A coefficient by exploring a cost-effective and expeditious path to the original design sensitivity of 0.2%. Parallel R&D should develop the experiment to measure the A coefficient with a sensitivity of 0.1% with the Nab spectrometer.
8. High priority should be given to acquiring new data with the cold-neutron-beam lifetime measurement at NIST following the facility's planned improvements.

The FY11 Presidential Request for neutron science totaled $9.33M. This amount is divided among: University Research ($1.3M), Laboratory Research ($3.8M), Fundamental Neutron Physics Beamline operations at the SNS ($330K), Capital Equipment excluding nEDM ($1.0M), and nEDM MIE ($2.9M). The constant-level-of-effort budget presented for FY12–16 corresponded to these amounts, adjusted for inflation. However, (1) Capital Equipment does not include funding for the nEDM project; (2) it is assumed that the Nab spectrometer magnet will be built with NSF instrumentation funds [this is in progress]; (3) FY10 Laboratory Research contains $650,000 for UCNA, which is not continued in successive years. Continuation of the UCNA experiment to achieve a 0.2% precision in the β-asymmetry parameter A, as recommended in this report, requires an additional $1.3 million over twoyears for LANL scientific staff and for operation of the UCN source; (4) further R&D and development of other LANL UCN experiments are assumed to be supported by LANL laboratory-directed research and development (LDRD) funds.

Scott Wilburn (LANL) said that, in FY13, LANL cannot cover that amount.

In terms of the workforce, all experiments provided full-time-equivalent (FTE) estimates for personnel. The "total workforce" defined a participant as devoting at least 0.2 FTE on one experiment. The field has 140 participants: 41 faculty, 27 research scientists, 25 postdocs, and 47 students (including the NIST effort). Thus, nEDM would have 31 FTEs: 11 at UCNA, 22 at NPD Gamma, 12 at Nab, and 4 at the NIST beam. The Subcommittee concluded that the field has adequate breadth and depth in its workforce.

The principal U.S. experimental initiatives provide excellent environments for technical innovations and for training of the next generation of scientists. However, the Subcommittee found that coordination of scientific effort and utilization of resources available in this area are not optimal at present. The Subcommittee's final recommendation was that consideration be given to establishing a standing committee to continually review and prioritize various initiatives in U.S. fundamental neutron science.

In its transmittal letter, the Subcommittee summarized its findings: The U.S. effort in fundamental neutron physics continues to be world-class, and there are compelling future opportunities in the United States in this area. A rank-ordered list of the five most important scientific priorities, as well as specific recommendations within each sub-area was developed. These five initiatives might be accommodated within a constant level of effort exclusive of MIE construction funding. Furthermore, the current workforce in neutron physics has sufficient scope and depth to carry them out. Most notably, the U.S. initiative to search for a nEDM at the Spallation Neutron Source (SNS) has unique capabilities to achieve a measurement sensitivity that will have profound impact on nuclear and particle physics as well as on cosmology. However, this promising approach still requires significant R&D, and the current nEDM effort should be focused on the most critical outstanding issues.

Geesaman asked if this report implies that these efforts can be completed under the constant-level budget. Jacobs replied that they *might* be. Geesaman noted that the cover letter implies that about $40 million in MIE would be added. Also, Nab is in the same boat. R&D might be covered under a constant level of effort, but equipment would not. Jacobs said that it was extremely difficult to make a concrete statement for the out-years. Geesaman suggested that the cover letter could state what could be done with existing funding. Jacobs agreed that this is a critical point to be made. The cover letter needs to distinguish between short term needs that can be accomplished with the existing level of funding and longer term needs that could require additional funding for construction.

A break was declared at 10:11 a.m. The meeting was called back into session at 10:38 a.m.

Jacobs offered a rewrite of the cover letter that excluded MIE and construction funding to address the issue raised by Geesaman.

Gardner pointed out that there were other issues of funding and that perhaps those issues should be reflected in the cover letter. Seestrom stated that whether an issue should be included or not is a matter of scale.

Lu stated that the volatility of budget issues is great and perhaps DOE and NSF should comment on the assumptions made. Hallman said that the Subcommittee received the charge and responded to it in the best way possible within the context of a particular funding scenario, which is the best that can be done. Neither the Subcommittee nor the agencies have a crystal ball and even the FY12 budget is unknown. The agencies will do their best to take the advice of the panel into account in managing the program, understanding that future funding constraints could impact what can ultimately be done. Blackmon asked how the priorities will be looked at by the agencies if there is a large budget cut. Specifically, will they cut from the bottom or distribute the cuts across the board. Jacobs said that the report is not prescriptive; it is a qualitative assessment and ranking. . It is up to the agencies how to implement these recommendations. The Subcommittee provides additional context to the agency staff. Keister pointed out that the Subcommittee cannot cover all scenarios, and no one has any idea what will happen to the budget in the future. Hallman said that what the agencies are seeking is advice on scientific priorities. The program offices will manage the programs as best as they can with that advice and within budgetary constraints; the Office may ask for advice from this Committee again as circumstances change.

Gardner asked whether, in the last sentence of the cover letter, it were prudent to rank UNCA so highly, given the lack of international competition. Jacobs replied that the Subcommittee got a lot of guidance, including the European perspective. Its Recommendation 9 calls for a standing subcommittee to deal with such developments (like the emergence of international competition). Gardner said that the sense of the argument is that cold neutron beams give different results from those of UCN bottle experiments; it is important that precision in cold-neutron-beam experiments be improved. Lu pointed out that UCNA uses a different technique. All of these are good experiments. It was difficult to come up with this ranking.

Binder stated that the idea of a standing committee is a good one and asked if it can be enacted. Hallman said that there are standing committees for some important topics. DOE and NSF would have to confer to see if this is one of those topics. Keister added that the agencies would have to consider *how* it would be done, also. Gardner said that the nuts and bolts aspects of setting up such a standing committee are important and asked how one composes such a committee with members that are not already heavily involved in the work that might be reviewed. Tribble asked if the standing committee would be put together by DOE or a national laboratory. Jacobs said that such a committee needs to be a collaboration among the three national laboratories together with the agencies, but it needs to be facility-driven.

Dodge asked if half the people were working on the high-priority experiments. Jacobs said that the Subcommittee did not calculate that out. Each experiment provided the Subcommittee with workforce levels. All of those workforce numbers were arrayed in a spreadsheet. It shows the solid core of workers in, say, nEDM.

Velkovska asked if these fractional people (FTEs) were working 100% in neutron science. Jacobs responded that these individuals were not necessarily working entirely in neutron science; the group is heterogeneous; some are 100% neutron physics, and some are multidisciplinary.

Jacobs suggested separate motions on the report and the cover letter. Hallman suggested polling the Committee's members individually to get their overall impressions. Each Committee member said that he or she would approve the report and the cover letter. Several members also congratulated the Subcommittee for the completeness and high quality of its work in preparing this report.

Jacobs noted that the report is from NSAC to the agencies and its language will reflect that relationship.

Gagliardi moved the approval of the report; Meyer seconded. The vote was unanimously in favor of the motion with Gardner and Seestrom recusing themselves from the vote.

Keister thanked the Subcommittee for its hard and excellent work. Hallman seconded that word of thanks.

Seestrom resumed the chair. The floor was opened to public comment. There being none, a break for lunch was declared at 11:30 a.m.

The meeting was called back into session at 1:11 p.m. **Daniel Hitchcock** was asked to present an overview of the activities of the Office of Scientific Computing Research (ASCR).

ASCR's job is to deliver world-leading computational and networking capabilities to extend the frontiers of science and technology. It provides computing facilities: high-end and leadership computers, research and evaluation prototypes, and the Energy Sciences Network (ESnet), which has 64 channels with 100 Gb per color. Two weeks before this meeting, a 10-Gb/sec network was compared to the new one. The difference was huge. The new network achieved 96 Gb/sec from application to application. It takes about a decade to get an idea developed into a usable system.

Scientific Discovery Through Advanced Computing (SciDAC) is now in its third generation. SciDAC 2 had 12 institutes and centers. No one knew which center to go to for consultation and advice. They are now being consolidated and currently number three:

- FASTMATH: Frameworks, Algorithms, and Scalable Technologies for Mathematics
- QUEST: on uncertainty
- SUPER: tools for diagnosing codes

An additional data-management center is to start up on February 2012.

The priorities for the science come from ASCR's partners [Fusion Energy Sciences, High Energy Physics, Nuclear Physics, Biological and Environmental Research, and Basic Energy Sciences (BES)]. Extensive discussions were held with these partners to combine the best applied mathematics, computer science, and networking with SC program expertise to enable strategic advances in program missions.

The future is in energy-efficient computing. Power consumption goes up as the square of the frequency, so frequency scaling has stopped and now many processors are being put on a single chip (up to 100 at this point). With current devices, 1 exaflop would require 200 MW. The target is to reduce that power usage by a factor of 10.

Dynamic random-access memory (DRAM) was designed in 1983 and is very inefficient; 60% of the power is used to keep the memory alive. As a result, energy usage is dominated by the capacitance. The hardware must be rethought to minimize the movement of data. One approach is to collocate different computations so one does not have to move the data. Another technique is to use concurrency. In addition, one must include hardware variability in uncertainty quantification (UQ); things like cosmic rays can introduce changes in stored data. One can also remap multiphysics to put as much work per location on the same die, include embedded UQ to increase concurrency, include data analysis for more concurrency, trigger output to move only important data off the machine, and reformulate to trade flops for memory use.

Codesign for the exascale was started this year to investigate the allocation of complexity, application designs, reformulation tradeoffs, UQ, new programming models and languages, massive multithreaded nodes, and ultra-lightweight operating systems. Three centers are looking at exascale issues: the Exascale Co-Design Center for Materials in Extreme Environments, the Center for Exascale Simulation of Advanced Reactors, and the Combustion Exascale Co-Design Center.

All hardware trends affect data-driven science (in many cases, in ways other than the computing intensity). Data from instruments still double every 18 to 24 months because the detectors are on the CMOS [complementary metal oxide semiconductor] feature size path. 100 gigabit per second per lambda networks are on the horizon with a lot of still-open questions. Disk read and write rates will fall further behind processors and memory. Significant hardware infrastructure is needed to deal with this development, which probably will not be replicated at the users' home institutions.

ASCR just finished a workshop on Data and Communications in BES to find out how to move the trigger closer to the camera, to handle 100-PB/year facilities, and to keep up with increases in luminosity without hugely increasing the power bills. BES's next big machine will produce data at a rate that is ten times that of the Linac Coherent Light Source (LCLS).

Klein stated that there are huge computer needs for a lot of detector simulation and event reconstruction, that those needs are not met by ASCR's program, and that there was a commonality throughout the physics community in the need for high-performance computing and the management of huge data sets. He asked what the nuclear physics community should do. Hitchcock responded that the architecture for event reconstruction has worked well. People who run experiments and ASCR are going to have to get together to make sure that advances go forward. During the past 10 years, most changes were in the interconnect. In the next 5 years, the changes will occur at the node. Therefore, there will be more opportunities for the two groups to work together. Jacobs pointed out that the commonality in need for high-performance computing goes beyond physics; for example, the Human Genome Project's Joint Genome Institute (JGI) produces huge amounts of data. Hitchcock noted that the JGI decided it was more cost-effective to install their machines at ASCR's National Energy Research Scientific Computing Center (NERSC) facility rather than at their own location. ASCR conducted the Magellan project to see how much could be saved across the SC. There is no good solution to the problem. It is on the margin of ASCR's mission. It is not known how it will play out.

Kaplan asked about the hierarchy of users and what the local user's machine would look like if the highest-performance machine were an exascale computer pushing huge amounts of data down to a laptop. Hitchcock replied that the local user will have 100 cores on a laptop. Rewriting the code for that laptop will make all previous conversions seem easy. There will be testbeds around to make that conversion faster and easier.

Furnstahl asked if there would be enough capacity to fulfill the needs of all the proposals that might come in. Hitchcock said that if the power cannot be brought down, there will not be enough capacity. If the power usage is brought down, there might be enough capacity. In the past, some universities have bought computers and only had half the power needed to operate them. Large computational systems use a massive amount of power. The more power they use, the more heat that has to be dissipated.

Seestrom asked about cooling strategies. Hitchcock said that cooling of the current systems can be improved by only about 15%. Bringing power use down is the only answer.

Hallman asked where else the exascale was being looked at. Hitchcock replied, Japan, Europe, and China.

Jacobs asked whether ASCR coordinated with the National Nuclear Security Administration (NNSA). Hitchcock replied, yes; the exascale work is a partnership with NNSA. The vendors have to spend about $2 billion to develop a new chip and they need a market to support those development costs. A partnership between DOE and other agencies is needed to produce such a market.

The floor was opened to public comments and for new business. There being none, the meeting was adjourned at 1:51 p.m.

These minutes of the Nuclear Science Advisory Committee meeting held at the Hilton Washington DC North, Gaithersburg, Maryland, on December 1, 2011, are certified to be an accurate representation of what occurred.

Susan J. Seestrom

Susan J. Seestrom
Chair, Nuclear Science Advisory Committee