ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE
to the
U.S. DEPARTMENT OF ENERGY

MEETING MINUTES

September 24-25, 2020

Teleconference

ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE

The U.S. Department of Energy (DOE) Advanced Scientific Computing Advisory Committee (ASCAC) convened a Teleconference on Thursday and Friday, September 24-25, 2020 via Zoom. The meeting was open to the public and conducted in accordance with the requirements of the Federal Advisory Committee Act. Information about ASCAC and this meeting can be found at http://science.osti.gov/ascr/ascac

ASCAC Members Present

Daniel Reed (Chairperson)
Keren Bergman
Martin Berzins
Vinton Cerf
Barbara Chapman
Jacqueline Chen
Silvia Crivelli
John Dolbow
Jack Dongarra
Thom Dunning
Tim Germann
Susan Gregurick
Anthony Hey
Gwendolyn Huntoon
Sandy Landsberg
Richard Lethin
David Levermore
John Negele
Vivek Sarkar
Krysta Svore

ASCAC Members Absent

Satoshi Matsouka

Also Participating

Barbara Helland, Associate Director, Advanced Scientific Computing Research (ASCR), DOE
Christine Chalk, ASCAC Designated Federal Officer, Program Manager, Oak Ridge Leadership Computing Facility (OLCF), ASCR, DOE
Chris Fall, Under Secretary for Science, DOE
Buddy Bland, Oak Ridge National Laboratory (ORNL)
Roscoe Giles, Boston University
Mike Heroux, Sandia National Laboratory (SNL)
Bruce Hendrickson, Lawrence Livermore National Laboratory (LLNL)
Harriet Kung, Deputy Director of Science Programs, Office of Science, DOE
Inder Monga, Lawrence Berkeley National Laboratory (LBNL)
Ceren Susut-Bennett, ASCR, DOE
Mark Taylor, Sandia National Laboratory (SNL)
Valerie Taylor, Argonne National Laboratory (ANL)
Rajeev Thakur, ANL
Jeffrey Vetter, ORNL

Attending

There were approximately 198 attendees in the teleconference.

OPENING REMARKS FROM THE COMMITTEE CHAIR, Dan Reed, ASCAC
Reed welcomed everyone and commented that despite the community's feelings of loss and isolation due to COVID-19, high performance computing (HPC) and the science collaborations are making important contributions to understanding the characteristics of the virus, its epidemiological spread, and the process of developing vaccines.

VIEW FROM WASHINGTON, Harriet Kung, Deputy Director for Science Programs

Kung discussed the DOE Office of Science (SC) reorganization, Quantum Information Science (QIS) research centers, the FY21 budget, COVID-19 response, and SC program activities.

The Office of Science has been working under a new organizational structure since April 2020. The impetus for the change was to better align SC with strategic goals. The reorganization affects the top levels of SC, primarily with the creation of the Principal Deputy Director, elimination of the Deputy Director for Resource Management, and retention of the Deputy Director for Science Programs and the Deputy Director for Field Operations.

Kung highlighted four new offices under the Principal Deputy Director: the Office of Accelerator Research & Development (R&D) and Production, the Office of Isotope R&D and Production, the Office of Strategic Planning and Interagency Coordination, and the Office of Diversity, Inclusion, and Research Integrity.

The Office of Accelerator R&D and Production will work closely with the science programs in SC-3 to strengthen the relationship and stewardship of accelerator R&D and production. The elevation of the Office of Isotope R&D and Production has raised the visibility of the critical, mission-central functions of the isotope program and gives it additional potential for growth and expansion. Both the Office of Strategic Planning and Interagency Coordination and the Office of Diversity, Inclusion, and Research Integrity's functions have been expanding their importance and visibility; having dedicated offices will serve the overall purpose and functions of strategic importance to the Office of Science.

In addition to the six program offices in the Deputy Director for Science Programs line, the Office of Grants and Contracts and the Office of Workforce Development for Teachers and Scientists are being co-located with SC-3 to enable the Office of Science to further strengthen the partnership between the science programs and these two offices. Two new offices that will be part of the SC-3 organization are the Office of Science Technical Information and the Office of Communications and Public Affairs.

Kung recognized Ceren Susut and Barb Helland's leadership in the QIS portfolio, a signature piece of SC investment. This is the first large-scale effort to cut across the technical breadth of SC. The extensive, community-wide input helped the development process, from scope to partnership model to management construct. The five QIS research centers have been awarded and collectively form a comprehensive investment by all six SC programs.

The FY20 enacted budget is $7.0B + $99.5M from the Coronavirus Aid, Relief, and Economic Security (CARES) Act for COVID-19 Science & Technology (S&T) response; the FY21 President's budget request is $5.838B. The House mark for SC overall is $7.05B. The emergency funding bill provides $6.25B to support user facilities and accelerate ongoing construction projects across the country. ASCR's budget has continued to reflect strong support at both the request and appropriation levels, having doubled over the past decade.

Due to COVID-19 SC has been exercising maximum flexibility to support people: permitting salaries and benefits, extending deadlines for proposals, granting no-cost extensions, and considering supplemental requests. Kung encouraged everyone to contact their research sponsor office and Office of Grants and Contracts for assistance with continued flexibilities.

SC has been given ~$100M of the CARES Act funding to address the COVID-19 response. Capabilities include light and neutron sources, nanoscience centers, computational resources, and scientists with deep expertise (testing, antiviral drug discovery, vaccine discovery, supply chain bottlenecks in advanced manufacturing, modeling, molecular and structural biology). DOE has stood up the National Virtual Biotechnology Laboratory (NVBL) to address the COVID-19 challenges. NVBL is a consortium of 17 DOE National Labs that coordinates in-coming requests to match gaps in technology or knowledge with the capabilities, facilities, and expertise at the labs. Part of the CARES Act funding has been applied to building up capabilities, including additional computing capabilities and other user facilities, as well as to stand up a number of compelling research projects. Initial activities focus on epidemiology, modeling, and logistics support; supply chain bottlenecks and prototypes; medical therapeutics; innovations on testing to address reagents; and the fate and transport of the virus in the environment (e.g., schools, restaurants, and other community gatherings).

Kung closed by noting program activities and reports across the other SC programs, and reminding everyone that the E.O. Lawrence nominations are due in October 2020.

Discussion

Cerf inquired about quantum algorithm development and ESnet6 funding. **Kung** stated that Helland and Susut will address these later.

REPORT FROM CROSS-CUTTING AI SUBCOMMITTEE, Tony Hey, ASCAC

The charge to the Artificial Intelligence (AI) for Science subcommittee was to deliver a report that assesses opportunities and challenges for AI for Science, identifies strategies ASCR can use, and creates a diverse committee across SC programs and agencies.

DOE's role as the lead agency for the AI for Science Initiative is due to its focus on data sets and the user facilities that create the data sets. DOE focuses on big computing, big data, and large scale science and has the capacity to assemble a science team with skills in mathematics and computer science together with software engineers and application scientists to create an end-to-end solution.

A successful AI for Science Initiative requires four components: application-specific solutions, research in AI algorithms and foundations, development of AI software infrastructure, and AI-specific computing architectures and hardware. Successful integration will require a full partnership between all of SC, engagement of national labs, involvement from universities and private industry, mechanisms for collaborative projects with other federal agencies, collaboration with expert international organizations, and an organized process for dissemination.

The subcommittee's key findings are that this is a once in a generation opportunity to take advantage of the growing convergence of AI, Data, and HPC; science can benefit from AI methods and tools; adoption of AI for Science will enable U.S. scientists to take advantage of new advances in user facilities; the generational shift encompasses the full spectrum of computing infrastructures; DOE labs are uniquely positioned to integrate AI/ML technologies; national impact will drive new industrial investments; a trained workforce is necessary; and partnerships will be critical.

The six recommendations include 1) the creation of a 10-year AI for Science Initiative, 2) a structure focused around four AI R&D themes (applications, algorithms, software, and hardware), 3) an Instrument-to-Edge-to-Compute Initiative, 4) training – focusing and retaining an AI/ML workforce, and 5) inter-agency and 6) international collaboration.

DISCUSSION AND VOTE ON AI SUBCOMMITTEE REPORT, ASCAC

Cerf commented on the differences between supervised and unsupervised learning and asked if unsupervised machine learning (ML) methods are capable of discovery. **Hey** said when the data comes off the machine it is not labelled; a graduate student could label a subset of the data, or one could use a Monte Carlo to generate artificial data to train on. On the other hand totally unsupervised learning is still uncharted and unknown – that is why there is a lot of interest in the research community for other ways of doing things. One of the challenges is to label data. There is no answer but the question is an important one. **Cerf** added that it is possible to do Monte Carlo designs of multi-layer neural networks and explore their capability.

Cerf reinforced the collaboration with the National Science Foundation (NSF) for two reasons: access to a large pool of talent, and topical resonance with the new NSF Director. **Hey** stated that he has spoken with the NSF Director who was immensely enthusiastic. **Cerf** introduced an example of agencies working together. DOE, NSF, the National Aeronautic and Space Agency (NASA), and the Defense Advanced Research Projects Agency (DARPA) all collaborated with Google on the internet. They formed FRICC, the Federal Research Information Coordinating Committee, who collaborated on funding. DOE and NASA were co-funding the predecessor to the Networking and Information Technology Research and Development (NITRD) Program which was doing cross-cutting coordination of research funding. Agency collaboration is possible, it has been demonstrated before and it can be done again.

Cerf expressed worry about Information to Edge (I2E) because of the potential to filter out potential new discoveries. **Hey** agreed that there are many opportunities in the control of complex apparatuses and facilities. I2E is an innovative way for ASCR to fund joint projects which accept inputs and expands out to new areas; it is an exciting project.

Negele noted that NSF announced a program to fund five institutes at $20M over 5 years. He asked if the subcommittee discussed ways to collaborate with them. **Hey** said that the subcommittee was aware of the centers and there has been discussion about DOE and NSF possibly combining in another center. Such a partnership would be welcomed by the AI community. **Reed** mentioned that the National Science Board is also interested in cross-agency collaboration.

Cerf called attention to a statement Dunning made in the chat on Zoom **–** the AI for Science project is not as defined as the exascale effort. AI is a much less understood territory and the roadmap, shown in the presentation, will be neither smooth nor easily laid out. **Hey** acknowledged and agreed with the statement, saying that part of the research program's role is to determine the services scientists will need; there will be an iterative loop in the R&D.

Dunning added that there will be a communication gap in the AI program. The Exascale Computing Project (ECP) is composed of individuals with a deep understanding of computing technologies and computation (computational scientists, computer scientists, and applied mathematicians). AI for Science will involve individuals who are not as conversant with AI and computing technologies. Dunning recommended funding initial multidisciplinary projects with the explicit goal of understanding the issues before initiating a large AI program. **Hey** explained that subcommittee members expressed interest in a SciDAC-like collaboration to keep in touch with ASCR's capabilities. He suspected there would be an evolution and a communication challenge. **Reed** added that in ECP the computing community was a substantial fraction of the leading authorities on HPC, whereas in AI there are many more players from various disciplines. **Doug Kothe** suggested that the incubation period proposed for the AI program (prior to the proposed formal project period) might be a good way to address Dunning's point as this is upstream R&D.

Richard Arthur (Chat) commented that perhaps a corresponding DOE-SC/NSF effort to AI for Science would be targeting theory development pursuits toward critical (AI-generated) black boxes that emerged out of the AI for Science process.

Chapman commended the subcommittee, expressing her appreciation for the broad scope that illuminated many aspects of AI and agreeing that collaboration and identification of overlaps will be important. Chapman drew attention to the Department of Defense's (DOD) AI research program that has little or no HPC; DOD is more interested in distributed systems. She commented that exploring overlaps in AI programs and research will be mutually beneficial.

ASCAC vote. Reed asked for any reservations to approve the report. Cerf moved to accept the report. Landsberg seconded the motion. Using a voice vote there was no opposition and the report was accepted unanimously.

VIEW FROM THE OFFICE OF SCIENCE, Chris Fall, Director of the Office of Science

Fall explained that the biggest issue on large project collaborations between NSF and DOE is the different stage-gate processes each agency uses. Securing a joint stage-gate process will be a significant accomplishment and would translate into more effective work between the two agencies.

Fall expressed his gratitude to ASCAC members for their participation, advice, and advocacy for ASCR. He said this is an amazing program doing amazing things for the country. For example the role of HPC, the machines, and the labs to address COVID-19; ECP's contribution to the country; the software stack – these are all profound accomplishments. It is breathtaking to see what SC is doing with huge science, analytics, predictions, changing the conditions and going back to do more science in real time; it is transformative. ASCR will play a role in the vision of data. The Associate Directors and Program Managers across SC worked out the idea of stewardship, management, and distribution of data. This will be put on the same footing with user facilities as a key pillar of SC. AI for Science is revolutionary; ASCR is going to edge into brain-inspired computing and other modalities. The point being DOE-SC is already thinking about what comes after exascale. For example, SC is working with the National Institutes of Health (NIH) on a potential joint project reminiscent of the Human Genome project and exascale will be big part of that. NIH wants to help the country but cannot do it without DOE's computing power.

Fall closed by updating ASCAC on the Intel situation and the Argonne exascale machine. He is confident the situation will be resolved to the benefit of the country and the program. The parties are still in conversation and he could provide no details.

Discussion

None.

Reed dismissed ASCAC for a break at 12:50 p.m. and reconvened the meeting at 1:10 p.m.

VIEW FROM GERMANTOWN, Barbara Helland, Associate Director of the Office of Science for Advanced Scientific Computing Research

Helland thanked ASCAC for their contributions and advice. Helland reviewed FY21 budget actions, ASCR personnel changes, ASCR years in review, and recognitions.

ASCR's FY20 enacted budget was $980M, the FY21 request was $988M, and the House Mark was $1.015B. Recommendations to DOE-SC include monies for AI/ML, Bioscience, Exascale Computing Initiative, and the QIS Research Centers. ASCR's portion of the QIS Centers includes testbeds, quantum applications, quantum internet, and quantum networking.

Helland shared updates on ASCR staffing, welcomed Hal Finkel as the new Lead for Compiler Technology and Programming Languages at the Argonne LCF (ALCF) and the two new American Association for the Advancement of Science (AAAS) Fellows, Sashwata Hier-Majumder and Jordan Thomas. Helland shared the departure of AAAS fellows James Ricci, Michael Nestor, Carolyn Vea, and Laura Biven. Rich Carlson, Robinson Pino, Bill Spotz, and Steven Lee will take over leadership in the Computer Science/ Data Management Analysis and Visualization portfolio.

Dr. Lucy Nowell, a recognized leader in visualization, passed away in June 2020. She was a computer scientist and program manager at government agencies including NSF and DOE. A virtual memorial for Dr. Nowell will be held at the IEEE Visualization 2020 conference in October.

Awards in FY20 included Early Career Research (7 awards), ASCR solicitations (5 awards), SciDAC Institute (2 awards), AI and Decision Support for Complex Systems (6 team awards), he Laboratory Base Math Program (9 awards), Laboratory Computer Science Machine Learning Research (4 renewals), and Data Management, Analysis and Visualization (8 awards). Helland also shared four accomplishments on the 1) COVID-19 rapid response, 2) 3D simulations for painting cars, 3) "Physics-Constrained" neural network, and 4) measuring quantum computers' capabilities.

The CARES Act provided funding to acquire additional computing facilities at ALCF, NERSC, Oak Ridge LCF (OLCF), LLNL, and Los Alamos National Laboratory (LANL), and to support a Tiger team to connect DOE's HPC with other agency efforts. The COVID-19 HPC Consortium was stood up on September 23, 2020; 90 of the submitted 171 proposals were approved. Six highlights were shared for NERSC, ALCF, and OLCF. ECP's early completion date is 2024 with CD-4 approval. ALCF's system may be delayed but OLCF's system, Frontier, is still on schedule to be delivered in CY21.

Helland mentioned four upcoming workshops, the Community of Interest (on Future Scientific Methodologies) in November 2020, Pioneering the Future of Federally Supported Data Repositories in January 2021, Data Reduction for Science Workshop in Winter 2020-21, and Randomized Scientific Computing – Algorithms for AI and Data Science at Scale in January-February 2021.

Two Sandia scientists were recognized – Tammy Kolda was elected as an Association for Computing Machinery (ACM) Fellow and Member of the National Academy of Engineering and Jacqueline Chen was chosen as a 2020 DOE-SC Distinguished Scientist Fellow. The NERSC Efficiency Optimization Team was also awarded the 2020 DOE Sustainability Award.

Discussion

Hey inquired about FY22 and FY23 budgets. **Helland** was restricted on providing details but indicated those budgets have been considered.

Cerf asked about the fate of new initiatives in the program in relation to the potential continuing resolutions. **Helland** explained that the new initiatives in FY21 are postponed until there is a budget. However, the activities that were started last year, with data and priority research, can move ahead. **Cerf** asked if there was a possibility to formulate the SC programs in such a way that includes forward-looking potential. **Helland** indicated that would require careful consideration with the science mission.

REPORT FROM EXASCALE TRANSITION SUBCOMMITTEE, Roscoe Giles, Boston University

Reed reminded ASCAC that the Exascale Transition Report had been approved at the April ASCAC meeting. Giles acknowledged the subcommittee members, ECP, ASCR research, the wider HPC community, ASCR and Advanced Simulation and Computing (ASC) leaders and staff, and SC program directors. Giles explained that the points raised by ASCAC in April have been incorporated into the text, both the text and discussions have been clarified, the executive summary, an acronym table, a summary of findings and recommendations, and references have been added. The findings and recommendation are unchanged except for editorial revisions and there has been improvement on the language.

The subcommittee was charged with examining ECP lessons learned, fundamental research investments, and new R&D priorities and to make recommendations on lessons learned, support software and hardware technologies, and inform ASCR's future investment strategy. There are findings and recommendations for each of the report themes: Advancing and Building on ECP (4 findings, 4 recommendations), Advancing ASCR Research (2 findings, 3 recommendations), Current and Future Workforce (3 findings, 5 recommendations), and National and International leadership (2 findings, 2 recommendations).

Discussion

Sarkar asked about investments in shared software stewardship. **Giles** commented that certain areas in the report discuss that feature, one is having predictable and consistent timetables in the appearance of opportunities thus allowing universities to plan ahead. There is also encouragement to use flexibilities with universities for doing rapid experimentation and prototyping that may be more difficult at the laboratories.

Reed expressed his worry about cultivating local and international talent and restricting visa intervals for international PhD students. The U.S. succeeds by being a vibrant, welcoming, and nurturing place for intellectual talent. Reed said it is important to encourage and support more members of underrepresented groups going into STEM.

WORKFORCE, Valerie Taylor, Argonne National Laboratory

Critical areas along the computing continuum include computer science, applied math, computational science, data science, computer engineering and more. Taylor's comments draw from three reports on the workforce and retention by the National Academies of Science, Engineering, and Medicine (doi:10.17226/24926; doi:10.17226/21739) and the ACM (retention-in-cs-undergrad-programs-in-the-us.pdf).

Taylor shared information on the challenges and opportunities to attract students from all levels as well as activities that are being pursued in STEM education at the labs and elsewhere. K-12 students need exposure to parallelism in novel ways. There is a need to add HPC to the 2-year curriculum. The challenge at the undergraduate level is determining how much of the curriculum focuses on the entire computing continuum. For graduate students summer internships and graduate courses in HPC are necessary. Finally, utilizing alternative pathways such as boot camps, curriculum changes, and focusing on the computing continuum are critical.

Partnerships offer many opportunities to bring the real world into the classroom and leverage project-based courses. These opportunities include guest lectures, course development, and certificates. Diversity can be increased through partnerships with minority serving institutions and by leveraging research opportunities and co-curricular activities.

It is important to consider all levels of the educational enterprise and to recognize there are multiple paths to degree completion. It is important to reward experience, tinkering, and curiosity. One recommendation is for ASCR to continue to support training and outreach efforts.

Discussion

Helland commented that the Office of Science can support undergraduates and graduate students, but only the labs can provide K-12 programs. DOE is looking for more opportunities with minority serving institutions and she welcomed ideas.

Berzins referred to an article by Moshe Vardi (doi:10.1145/3410470) about the challenge of getting domestic students interested in PhD pursuits and he asked if Taylor had any thoughts on how to change PhD programs to make them more attractive. **Taylor** suggested that giving students the opportunity to work on research and open-ended problems can be very exciting. Pointing out income acceleration and new opportunities through graduate degrees may entice some students. Seniors are focused on the job market which means it is important to communicate the advantages of a graduate degree early. **Reed** added that a common recommendation in National Academy studies about graduate school environments is shifting money away from Principal Investigator (PI) awards towards graduate fellowships.

Huntoon thanked DOE for the support they provide for the Women in IT Networking at SuperComputing (WINS) program. Many women from the WINS program began their studies at community colleges, usually in a non-technical program. There are also mid-career changes for women entering the information technology field. She asked if there are ways to capture that group of women. **Taylor** said that it is vital to expose associate degree students to STEM programs and HPC. Examples of some programs included Berkeley’s reentry program geared towards women and minorities to pursue a degree in computing or engineering, Texas A&M University’s Engineering Academies with two-year institutions for co-enrollment, and university articulation agreements with two-year institutions.

Sarkar said that there is a lack of qualified people to teach parallel programming, computing, and HPC, especially at two-year institutions. He suggested that lab scientists contribute by teaching this content and inquired about potential obstacles. **Taylor** explained that at 2- and 4-year institutions often the start-up phase (the computer science material to be taught and the platform to be used) is difficult. It is important to creatively disseminate work that has already been done for use in a course as a module.

Reed stated that entering and exiting the education system throughout life is increasingly the norm. Thinking about how to support that will be important. Reed referenced the National Science Board report on the Skilled Technical Workforce.

QUANTUM INFORMATION SCIENCE CENTERS, Ceren Susut-Bennett, ASCR

The National Quantum Information (NQI) Act (PL 115-368) was signed into law at the end of December 2018. NQI gave specific roles to three agencies (DOE, NSF, National Institute of Standards and Technology (NIST)). DOE's role, in part, was to establish at least two, but no more than five, QIS research centers. A Request for Information to gather community input was open from May to July 2019; 38 comments were received.

Between the FOA release in January 2020 and the mid-August applications deadline, responses were given, proposals were received, and selections were made on the five QIS Centers. The selections were based on five factors (merit review recommendations, pre-selection interview and clarifications, SC NQI Centers Working Group recommendations, and program policy). The five national QIS centers awarded in 2020 are Q-NEXT at ANL; C^2QA at Brookhaven; SQMS led by Fermilab; Quantum Systems Accelerator led by LBNL; and QSC at ORNL. Each QIS center targets every layer on the S&T Innovation Chain from fundamental science to applications. There are processes to enable technology transfer with partners to disseminate the research results. Each center is cross-cutting to ensure all the subtopics in the technical areas of interest are addressed.

The QIS Research Centers Portfolio has a diverse management structure and project management approaches, fully leverages SC programs across the lab complex, and reaches outside to industry, other agencies, and foreign institutions. The strategy is built on community engagement beginning in late 2014 with a roundtable discussion. The goal is to collaborate with industry to provide an innovation economy, to coordinate with other agencies, and to remain aware of the developments internationally.

The National QIS Research Centers are crucial to advancing the QIS goals. Several new elements were introduced: a whole of DOE-SC and whole of QIS approach to the FOA, a portfolio that addresses a DOE-SC-wide and a QIS-wide scope, a new structure for coordination within DOE-SC, a website about DOE-SC investments in QIS, and an emphasis on the ecosystem stewardship.

Discussion

Reed stated that there are lessons to be learned from ECP and AI for Science initiatives and there are things to be drawn from here that are applicable in other domains.

PUBLIC COMMENT

None.

Reed called the first day of the meeting to an end at 3:47 p.m. ET and reminded everyone about the continuation on Friday, September 25 via Zoom.

Friday, September 25, 2020

OPENING REMARKS FROM THE COMMITTEE CHAIR, Dan Reed, ASCAC

Reed called the ASCAC meeting to order at 11:00 a.m. ET, welcomed everyone to the meeting and introduced the first speaker.

INCORPORATING GPUs INTO EARTH SYSTEM SCIENCE, Mark Taylor, Sandia National Laboratories

Taylor discussed the Energy Exascale Earth System Model (E3SM) and the E3SM-Multiscale Modeling Framework (E3SM-MMF) approach (superparameterization).

Three overarching science drivers of E3SM are the water cycle, biogeochemistry, and cryosphere systems. E3SM focuses on new types of simulation campaigns to run on GPUs. GPUs will benefit the high workload simulations E3SM's "SCREAM" project, E3SM-MMF, and E3SM V4. There have been large investments in performance and portability strategy on Fortran + OpenMP and C++/Kokkos. Both of these require complete code refactoring or rewriting. This opportunity allows the Office of Biological and Environmental Research (BER) to replace legacy code, improve testing, and perform software engineering and verification. E3SM has upgraded their model to non-hydrostatic and rewritten it in C++. There has been steady progress in the community across machines and models and the first application is close to running on GPUs.

The E3SM-MMF approach addresses structural uncertainty in cloud processes by replacing traditional parameterizations with "superparametrization". GPUs will not make this model run faster but they will allow dramatically increased complexity while running at the same speed as the original model – an ideal approach for GPU acceleration. Cloud resolving models (CRM) improve observations much better than traditional resolution. The superparameterization approach runs well on GPUs, performance is dominated by a 2D CRM, and porting only the CRM is sufficient to make effective use of GPUs.

Taylor's concluding thoughts were that CPUs and GPUs are getting better, the E3SM V3's two approaches centered on cloud resolving simulations are made possible by GPU architectures, and E3SM V4 will efficiently run on both CPU and GPU architectures.

Discussion

There was a significant amount of discussion via the Chat feature in Zoom. **Cerf** (chat) asked if GPUs were low-precision, about sea/atmosphere interactions, and the meaning of verification and tracer. **Dongarra** clarified that GPUs can do 64, 32, and 16 bit flops. **Sarat Sreepathi** said the components (Ocean-atmosphere and other combinations) are propagated via a coupler component in a coupled configuration. A typical coupled model simulation includes atmosphere, ocean, sea-ice, land, land-ice, river, etc. to capture interactions and feedbacks between various components via the coupler. **Andy Saligner** explained that a tracer is a species that is transported in the air, like CO2, methane, and water vapor.

Verification means the model is coded correctly (**Andy Salinger**); it is a check on the implementation of the mathematical algorithms (**Carol Woodward)**. Validation, however, checks whether the code matches the physical system being targeted (**Carol Woodward**). **Larzalere** reminded everyone that the level of verification, validation, and uncertainty quantification needed depends on how the results are used – discovery science might not need much, but operating a nuclear reactor requires a lot more.

Berzins asked if Taylor knew, roughly, the performance characterized by n2 flops for a problem of size n. **Taylor** replied that at the moment he did not know, but explained that they are bringing in a CRM, a conventional low flops per memory access, and thus are not increasing that arithmetic intensity. Because of the workload, all the threads of the GPU can be used simultaneously with no synchronization. The 3X benefit is possible because of reliance on the bandwidth improvements of the GPU. **Berzins** commented that comparable performance may be realized with the Fugaku ARM.

Berzins asked about the cost of the energy efficiency related to the amount of rewriting time. While he agreed with E3SM's work with cloud resolving – using the power of the GPU to increase complexity nodes – the resources required for rewriting the codes make the GPU option impractical. **Taylor** agreed and said one concern is if the CPU and GPU trends will continue in parallel or if the GPUs will pull away. **Berzins** added that for many applications the memory bandwidth is key and that determines how the machines are actually designed.

Lethin asked what tools and approaches are being used for the correctness, verification, and validation in the software engineering. **Taylor** shared an example of the SCREAM project. Every subroutine is rewritten with the unit test. Every C++ routine, starting with the unit test, verifies with the Fortran code. SCREAM runs in special modes where the code is bit for bit and then adds property tests. Simple properties should be maintained by the subroutine and tested individually rather than via the traditional approach since these codes have a lot of legacy behind them. Then convergence tests are added. In special regimes convergence can be done, if designed carefully, off of single components. Finally, there is nightly automated testing. **Lethin** asked if there is any role for formally specifying performance conditions. **Taylor** said testing is expensive and takes a lot of time. However, several people have been pushing this hard and are slowly convincing the scientists this is important – the more testing the better.

Sarkar asked how much of the performance gap is a comparison between handwritten CUDA code and compiler generated GPU code from OpenMP rather than Fortran and C++. **Taylor** explained the code is a C++ array class called YAKL, which is similar to Kokkos. YAKL is easier to use with code coming from Fortran. **Sarkar** asked if the C++ code was handwritten or was a compiler tool used to generate the code. **Sarat Sreepathi** (chat) wrote that YAKL uses template metaprogramming with CUDA and other back-ends analogous to how Kokkos launches GPU kernels. It is a lightweight layer with Fortran-style multidimensional arrays support.

Lethin remarked that LBNL was architecting a special-purpose computer called Green Flash along the lines of the Anton machine for molecular dynamics but applied to climate. He asked if there is a role for specialization of hardware for E3SM. **Taylor** speculated that there is a role for special-purpose hardware, although commercial viability is questionable. The benefits of GPUs for some applications it is more dramatic than in climate. E3SM's role is to run cloud resolving effectively. How much can be done with custom hardware on E3SM is unknown. **Lethin** added that if the mission is critical enough then the commercial viability considerations might be set aside. **Reed** said there is a philosophical R&D question about if, as traditional semiconductor advances slow, the pendulum will swing back to ASIC for code acceleration.

James Ang (chat) wrote that the original plans for ECP Hardware Technology included "alternative paths" to explore purpose-designed hardware that was not based on existing hardware roadmaps. **Richard Lethin** (chat) responded that it would be good to see effort put into those alternate paths particularly for this application. **Dan Reed** (chat) added that the Microelectronics Basic Research Needs report considers alternative paths.

PREPARING FOR THE SUSTAINABLE DELIVERY OF THE DOE EXASCALE SOFTWARE STACK, Mike Heroux, Sandia National Laboratories; Rajeev Thakur, Argonne National Laboratory; Jeff Vetter, Oak Ridge National Laboratory

Heroux provided updates on progress in exascale sustainability and software technology (ST), Thakur discussed application development (AD) portfolio projects, and Vetter shared information on LLVM and hardware and integration (HI).

Kathryn Mohror from LLNL has joined the ST team replacing Rob Neely. One L4 level project, ExaWorks, was added. And Todd Munson has replaced Gary Smith as PI following his retirement.

ST tracking is focused on the third key performance parameters (KPP-3) – tracking capability integration into stakeholder environments. The dashboard to manage KPP-3 progress in Jira is available. Teams are encouraged to update the dashboard at least 2x per year as capabilities are integrated into stakeholder environments.

The Extreme-scale Scientific Software Stack (E4S) has the ability to create cross-institutional efforts for software development kits (SDK). E4S components can build from source or containers, is available in cloud form, and has GPU builds for AMD and NVIDIA with lots of access points. E4S community-driven quality commitments to users addresses some issues concerning dependence on 3rd party research software. This approach has migrated out to other SDK teams and a set of candidate policies on what it means to be part of the E4S stack have been created. Community policies and the community discussions surrounding those policies are important. Policies created for E4S and the math libraries are hard thought discussions about what it means to provide quality to the stakeholders.

The E4S documentation portal (Doc Portal) is new and will satisfy some user challenges. Doc Portal is an efficient and effective tool that allows developers to get noticed by new users and includes summary information and details about a product. The portal strategy is to have all documentation for the product reside in the product repository. Doc Portal relies upon the information that is in the product repositories and minimal information is requested from the development team concerning access.

The E4S Spack build cache and container build pipeline drastically reduces build time by using a cached build of a product. When Spack sees a repeated signature it simply grabs the pre-built binary data from the cache thus drastically reducing build time. Spack provides the technology and E4S is bringing the effort of keeping versioning, reducing so-called dependency help, and providing the latest capabilities for exascale and pre-exascale platforms.

E4S is not a closed system, not monolithic, not a commercial product, not a simple packaging of existing software, it is all of these things and more. E4S is now a portfolio and can speak with an aggregate voice, can expand to new domains, and can get better, faster, and cheaper by using the approaches being created.

Thakur discussed programming models and performance portability. The new exascale machines will have GPUs from multiple vendors meaning all software projects must support GPUs and support multiple vendor GPUs. In terms of programming between nodes and within a node the stress is on working well with the intra-node programming model, the hybrid programming aspect. MPI is expected to continue to solve this for most applications, but some Partitioned Global Address Space (PGAS) programming models face challenges in achieving performance and portability within a node.

The AD Milestone Report is available online and contains details about ECP applications and programming models used by ECP applications. Codes are using a variety of languages, but mainly a mix of C++ and Fortran. Kokkos and RAJA are C++ portability abstractions developed at LLNL and Sandia. ST has provided extra money to develop optimized back ends for Aurora and Frontier. Both efforts are organized as one project and they are collaborating on common features and common backend supports.

Kokkos has three components – core (C++ template-based library), tools (profiling, debugging, and tuning), and kernels (math libraries effort based on the Kokkos core). The team works closely with the International Standards Organization (ISO) on integrating Kokkos features into the C++ standard to enable the functionality Kokkos needs. RAJA also has three components – C++ kernel execution abstractions (RAJA), C++ array Aria abstractions (CHAI), and memory management (Umpire). RAJA has had the full support for HIP back-end public releases since January 2020 and is a key part of El Capitan Center of Excellence activities. The SYCL back-end is in progress, including SW4 and RAJA Performance Suite which are running on pre-Aurora systems now.

MPICH will be the primary MPI implementation on all three of the exascale platforms. Ongoing efforts include the MPI Forum and standardization, addressing ECP application-specific issues, work related to GPUs from multiple vendors, and a new library for efficient noncontiguous data communication. MPICH is taking advantage of the GPU RDMA. They have recently developed a new library called Yaksa for derived data types. MPICH has also made some improvements related to MPI + Threads, or hybrid programming.

Vetter discussed development tools and LLVM. LLVM is an infrastructure for creating compilers. The benefits of LLVM include its modular, well-defined IR, back-end infrastructure, state-of-the-art C++ frontend, CUDA support, scalable Link Time Optimization (LTO), and high-code quality. LLVM has a very permissive license that allows people to contribute and has been used in many languages including C++, Julia, Rust, Swift, and Tensorflow.

ECP is improving the LLVM compiler ecosystem by enhancing the implementation of OpenMP (SOLVVE), improving core optimization (PROTEAS-TUNE), developing open-source production Fortran frontend (FLANG), improving OpenMP profiling interface (HPCToolkit), optimizing template expansions (ATDM), and collaborating with many vendors. Vetter shared details of three examples – the Fortran effort (SOLLVE, PROTEAS-TUNE, +other contributors), the OpenACC in Clang (PROTEAS-TUNE), and the Speculative Loop Transformation Representation (SOLLVE).

The ECP LLVM integration and deployment team was asked to create a version of LLVM they could deploy on their own. The process is to mirror the entire LLVM nightly, then have branches (three currently) that ECP contributes to. There will also be periodic upstreaming and patching of LLVM. This integration and deployment plan allows any merge conflicts to be addressed on a daily basis and benefits from changes that are being made globally.

Discussion

Cerf (chat) asked if there is anything comparable to the Spack Forum for quantum. **Richard Arthur** referred to Quantum Algorithm Zoo (https://quantumalgorithmzoo.org/).

Levermore requested information on addressing access by malicious actors. **Heroux** explained that there is strong control on who is allowed to write to the code. A poll request (merge request) is submitted. That request is vetted by another developer and has to run through a battery of tests before it can be applied to the primary repository. More of the teams are adopting this best practice. **Vetter** added that LLVM has a vibrant developer community where all the changes are thoroughly vetted for coding standards, optimizations, and other features.

Lethin asked about access to the ECP software stack, and if the source code and Atlassian issue database are open to vendors and academics. **Heroux** explained that the source code is all open source. Anybody can access it in terms of reading it, but they cannot necessarily write to it. Teams are encouraged to have their primary issue database in their product suite, not in the ECP Atlassian platform. ECP tracks activity, but at a level of value proposition to ECP such as the regular portal, activities, and the KPPs. Access is generally given to people as business needs require – it is not generally open to everybody.

Berzins asked about the powerful levels of abstraction in C++ having led some developers to produce code that is almost compiler-proof, in terms of optimization, resulting in code that is easy to use but ultimately cryptic. **Heroux** said that the capability in KPP-3 protects against that, to some extent, because of the requirement that a capability be sustainably integrated in the stakeholder's environment. However, some products do require advanced C++ skills. **Vetter** added that is not just an issue in the C++ community. LLVM is working on more support for just-in-time compilation, autotuning, and link time optimization. **Berzins** followed up asking if there are any general principles – safeguards but not a set methodology. **Heroux** explained that many of the projects are discovery-based efforts that feed into a product-oriented pipeline. The existence of E4S and LLVM as ecosystems goes a long way to protecting against ultimately unusable approaches. By having safeguards in place and conducting regular checks on team progress, issues are quickly identified and corrected.

Reed posed a philosophical question – the great thing about ECP is it has developed a vibrant ecosystem and a large codebase. The bad part about ECP is that it has developed a large codebase. Along with considering maintenance, one must think about sustainability and balance in a long term support process that has bounded resources. **Heroux** responded that in a fundamental way the ECP L4 activities are not fundamentally different than what was going on before ECP. A large bulk of the funding ASCR is giving ECP is the same money they have been allocating for decades. ECP provided the value adds of the SDKs and E4S, which are allowing the development teams to feed into a well-structured and managed software ecosystem that will ultimately provide better quality for their software and get it out to users quickly and in a way that is cheaper than what has been done in the past. **Vetter** commented that in the Extreme Heterogeneity workshop the major concern was that now there are many branches (AMD GPU, Intel GPU, NVIDIA GPU, FPGAs) and this will quickly become a crisis. A model is necessary to support and keep up with different versions of the software and a gatekeeper is necessary to ensure consistency and community standards.

APPROACHING THE ESNET-6 ERA, Ben Brown, ASCR and Inder Monga, Lawrence Berkeley National Laboratory

ESnet is a user facility, a service provider to the service providers, a complex system tuned for science, and a world-class high performance network providing services and innovation to enable research. ESnet is, and has been, a wide-area network (WAN) connecting to the resources, the internet, and local area networks at the labs.

ESnet6 is a greenfield build of a new WAN, meaning for the first time ESnet will have complete command of the infrastructure. The themes of ESnet6 revolve around foundations (next generation infrastructure and services), innovation (testbeds and advanced networking R&D), and co-design (partnerships for new data solutions). ESnet's vision is that scientific progress will be completely unconstrained by geography. This goal leads to three strategic focus areas: 1) outstanding operations and planning, 2) information and tools for optimal network use, and 3) pioneering architectures, protocols, and applications.

ESnet6 architecture focuses on the ability to deliver terabit-scale performance with programmability and custom science services. ESnet is building the greenfield optical network including amplifiers, ~15K miles of fiber, 300+ sites, a packet core, a "low-touch path service edge", and a "high-touch path" to provide a variety of services. The CD-4 date is 2025 and early finish is 2023. The optical core build is almost complete. The first routers will be in the lab in a few weeks and the automation and installation process to deploy ESnet6 will occur in the next year. COVID has created a 3-month delay and some of the schedule contingency may be used to ensure mistakes are avoided. The optical substrate was built between April 2020 and August 2020. Customers were being transitioned from ESnet5 to ESnet6 as it was being built; 75% of all traffic is now on the new optical network and the remaining 25% will be transitioned by late November 2020. By 2022 there will be mix of 800Gbps to 1.2Tbps initial capacity deployed with the capability to add more bandwidth; most of the network will have base capacity of 400Gbps up from 100Gbps today.

Monga shared activities going on in ESnet including Packetscope, SENSE Automation, Intelligent edge – Data Transfer Nodes as Service, NetPredict, and BBRv2.

The "high touch path" is in the prototype phase. It is integrating edge computation, storage, and smart network interface controllers (SmartNICs). Personnel are working on Packetscope, similar to a telescope, which is a tool that uses field-programmable gate arrays (FPGA) and compute to take any network flow and get packet telemetry on a per packet basis. This will allow dynamic configuration and real-time visualization of the data.

The SENSE project, which is now part of ExaFEL, tackled the problem of moving data traffic where the end point does not have any controls and thus cannot make use of any automation at the instrument. This is now being used to automate and configure from the NIC on the server to the network to the NIC on the end point server.

Intelligent edge provided a new opportunity to offer Data Nodes as a Service by containerizing data transfer nodes and offering those containers as a service. This will give people the choice of a large number of data transfer protocols that they can deploy. Since these will be deployed on ESnet6 it will be possible to go from one edge to the other edge of the network without requiring the end sites to make changes or the end site administrators to learn new protocols or new ways of doing things.

NetPredict is a deep learning model to predict available network traffic paths based on past traffic patterns. The goal is to make NetPredict a more intelligent traffic engineering controller to traffic engineer big data science data flows, or elephant flows, more appropriately.

ESnet has started doing evaluations of BBRv2 based on the anticipation that in the future there will be small buffer/ high buffer delay product environments that require new protocols. ESnet has started an internal project, using 40G data transfer nodes, to build a testing harness for BBRv2 use for high-throughput data transfer.

Three examples of co-design that ESnet is involved in are the Gamma-Ray Energy Tracking Array (GRETA), Advanced Light Source (ALS) and National Center for Electron Microscopy (NCEM), and Caching Compact Muon Solenoid (CMS) and A Toroidal LHC ApparatuS (ATLAS) projects. The ESnet prototyping and testbed group and network engineers worked on the Forward Buffer and the fiber plan for the GRETA data pipeline. ESnet is working closely with ALS and MCEM to put the edge node (FPGA) architecture into use for scientific processing. Finally, ESnet is working on a caching platform with High Energy Physics' CMS and ATLAS and will build the analysis to evaluate effectiveness.

ESnet is supporting DOE and the national strategy on quantum and 5G by enabling the testbed in Long Island, working with Fermilab to use ESnet capabilities in Chicago, being part of the quantum blueprint effort, and participating in the 5G Enabled Energy Innovation workshop. The ESnet architecture is being leveraged to design NSF's mid-scale infrastructure project called FABRIC. Finally, ESnet wants to partner with NSF to provide an at-scale network for network researchers at every U.S. university to ensure that network research stays healthy.

Discussion

Cerf asked if Monga's term "hollow core" is synonymous with hollow fiber. **Monga** explained the hollow core is segmented tunnels where the tunnels get the packet information.

Cerf inquired if the optical design implemented for ESnet6 is capable of carrying both conventional and quantum photon communications. **Monga** said the current one cannot carry quantum photon information. The classical network must be working in concert with the quantum network to make the quantum network happen.

Cerf asked if there was a noticeable change in the interfaces during the transition from ESnet5 to ESnet6. **Monga** said the same router is being kept so there is no noticeable difference.

Cerf asked if the BBRv2 team has looked at the QUIC protocol from Google. **Monga** responded that the QUIC protocol targets the browsers and end clients while BBRv2 is targeting high-throughput activities.

Reed dismissed ASCAC for a break at 1:38 p.m. and reconvened the meeting at 2:00 p.m.

DISTINGUISHED SCIENTIST TALK, Jackie Chen, Sandia National Laboratories and ASCAC

Chen thanked ACSR for the honor which comes with a three-year sponsorship. During her Fellowship, Chen will construct a software framework for In Situ Reduced-Order surrogate modeling for Direct Numerical Simulation (DNS) of Turbulent Combustion at the exascale.

In the last 10 years she and a team have developed a combustion, high-fidelity suite of codes known as Pele. This is a suite of reacting flow partial differential equation solvers. The challenge problem, when Frontier is available, is to address key questions around the distribution of reactivity in preparation of a mixture, multiple injections, durations, fuels needed to shape the right reactivity gradients in time and space, and rate-controlling reactions.

A second code, S3D, solves compressible reactive Navier-Stokes energy and species continuity equations and uses higher order finite difference method for spatial differentiation and for advancement in time. S3D can treat detailed reaction kinetic and molecular transport models. The code has been refactored a number of times from MPI codes to OpenMP, to OpenACC, and more recently using Legion.

Because of scalability and portability of S3D the team is now able to simulate full-up laboratory-scale flames with detailed chemistry in full 3D turbulence. This is allows the team to sit down with experimentalists an design experiments, make direct comparisons, and validate diagnostic techniques.

To move to exascale and eventually to yotta scale there is still a problem that ‘real’ fuels, those used in aviation or in diesel trucks, are described by chemical models with very high dimensionality. The goal is to find a low-dimensional manifold in the species-temperature space that serves as a surrogate for the full system dynamics of the reacting flow. There has been a push in the last 10 years to look at Empirically-derived Low Dimensional Manifolds (ELDM) that are constructed from data.

Chen wants to develop a framework to perform in situ reduced order surrogate models for DNS at the exascale for combustion. Chen’s computational framework will run at scale on Summit and on future exascale machines, will keep the DNS data generator to create the raw data and identify the quantities of interest features, will build and detect anomalous behavior, and will provide reduced- or course-grained data through filtering and averaging processes. In addition to the traditional DNS solver the framework will include in situ ML algorithms, include static and time-dependent bases for reduced order models that can be applied to highly transient systems, and use convolutional neural networks.

Elements of the project are 1) to look at reduced order modeling to reduce the high dimensional composition space needed to describe multi-component gasoline and diesel surrogates to enable DNS of turbulence chemistry interactions for the PACE Engine Consortium; and 2) to replicate the high dimensional computation space with low dimensions using principle component (PC) transport.

To date the PC transport has been trained with static principal component analysis (PCA) from existing observational data. The next step is to evolve the PCs dynamically during the course of a full simulation. Reduction and modeling of the transport terms has to be done on-the-fly during the DNS thus creating a complex workflow. Metrics must be developed to determine when the principle components are dynamically updated. Ideas on how to do this include augmenting the second moment and higher moments, considering the 4^{th} order Kurtosis Moment, defining anomaly detection with PCA transport to update, and doing co-kurtosis.

Another method to use and evaluate in this framework is real-time reduced-order modeling on a time-dependent basis. This has two modalities, both observationally driven and model driven. The idea is to develop a method that will adaptively generate a time-dependent basis that will capture the strongly transient phenomenon along a given trajectory in the system. This will be put together using the testbed S3D programming in Legion which is a dynamic task-based system. Legion is a data-centric parallel programming style developed to work well on heterogeneous distributed machines with accelerators and other kinds of compute resources. The method automates aspects of getting high-performance (data-level parallelisms but also task parallelism) and it automates the detail of how tasks are scheduled and how data is moved around on the machine. Legion has been refactored using the compile code Regent. Regent is a task-based programming model that is built on Legion runtime, ports to exascale machines, and automatically compiles for different GPUs by adding only one command.

The last part of the infrastructure is the deep learning framework, Flex Flow, which is also built on top of Legion. It distinguishes itself from the commercial deep learning frameworks in that it has automatic search methods based on Markov Monte Carlo search algorithm to find high performance data partitioning. Flex Flow dramatically improves locality and scalability and reduces the large scale training that is anticipated for machines like Summit from days to hours.

This fellowship is an opportunity to promote scientific and academic excellence in ASCR and Basic Energy Sciences (BES) research through collaborations between labs and universities. We will practice co-design, train students and post-docs, develop an open source scalable and portable in-situ DNS/ML framework, generate high-fidelity data, and engage the LCFs.

Discussion

Dunning congratulated Chen on the award and all the work she has done in her career. He said the progress that has been made over time shows what can be done when DOE makes long-term investments in a particular area. The collaboration between BES and ASCR on these types of problems illustrates what a “whole of Office of Science” activity can accomplish.

Levermore asked how ML helps budget and navigate uncertainties that are derived empirically, by observations, by theory, or when navigating model space. **Chen** said while that is a research topic unto itself, her team can address sensitivity – how sensitive is the prediction of emissions (e.g. soot) dependent upon getting certain elementary reaction rate properties correct. Propagating the forward sensitivity analysis using these time-dependent reduced-order models will provide that kind of information. Those sensitivities can then be communicated to the chemists. **Levermore** relayed that in the 1980s LLNL used a very low-dimensional, known to be bad, model and built it up with sensitivity analysis – growing the model rather than reducing it.

Reed congratulated Chen for her phenomenal progress.

REPORT FROM THE SUBCOMMITTEE ON 40th ANNIVERSARY ACCOMPLISHMENTS, Bruce Hendrickson, Lawrence Livermore National Laboratory

Hendrickson shared that the two ASCR@40 documents (report and highlights) are both complete and available at https://computing.llnl.gov/ascr-at-40.

Chris Fall and Harriet Kung expressed interest in a set of materials for outreach and advocacy on behalf of ASCR, based on the BES one-pagers, to communicate ASCR's impact to stakeholders. These have a common structure and layout and are being created by the Krell Institute. The 11 topics are computational science, mathematical foundations, uncertainty quantification, networking, collaboration tools, big data and visualization, parallel processing, architectures, facilities, workforce/ Computational Science Graduate Fellowship (CSGF), and open source scientific software.

Discussion

Reed thanked Hendrickson and all those who worked through the three documents.

A FEW THOUGHTS ON HPC, Buddy Bland, Oak Ridge National Laboratory

Bob Meisner introduced Buddy Bland by sharing some background stories. Bland acquainted the U.S. Air Force with supercomputing as they set up the first Cray computer at Kirkland Air Force Base. Following the U.S. Air Force, Bland went to ORNL to stand up their first Cray computer, the XMP. Twenty years later Bland was helping Oak Ridge become a leader throughout the world. As Program Director for the OLCF, Bland provided three number one computers for the nation. More importantly he delivered the tools the nation needs for its most demanding problems. When the nation called, those machines that Bland brought to full operating capability were the things that enabled scientists to provide the answers like no others in the world. From megaflops to exaflops Bland has dedicated his life to serving his country and we want to recognize that.

Bland thanked Chalk and Helland for inviting him to speak and Meisner for the introduction. Bland began by sharing that he has worked on systems from minicomputers to vector machines to massively parallel machines to hybrid CPU/GPU systems. Each architectural change reflected the end of one type of technology and the birth of another technology to address performance improvements. For example, the vector machines that used emitter-coupled logic eventually shifted to complementary metal-oxide-semiconductors (CMOS), and now at the end of Moore's Law and Dennard scaling CMOS's transition is being determined. But the hardware is only a piece of the performance story. Software improvements account for a large part of the performance improvements of applications. Better software has come in the form of mathematical methods, improved solver precision, improved accuracy and performance, better development tools, and new versions of applications that take advantage of the latest architecture.

Much of ASCR's investments in these areas have come through the SciDAC program and ECP which is supporting ~50 applications, 70 software tools including libraries, compilers, debuggers, profilers, etc. The complexity of the systems has also grown tremendously. The Cray-1 had approximately 200K gates and used 70 sqft of floor space, a single IBM Power 9 processor that runs Summit has 8B transistors and each NVIDIA voltage CPU has 21B transistors, between 2-4 transistors per gate and Summit uses ~8,000 sqft of floor space.

The software complexity is just as daunting. Programming a hybrid CPU/GPU system is more difficult than using traditional HPC CPU-only machines. The Cray-1 had a single CPU with several functional units that provided small levels of parallelism. Some of today's systems have 10s to 1000s of nodes and billions of threads of execution. Managing multiple levels of parallelism and different instruction sets between CPUs and GPUs in a system takes a lot of skill and experience along with the right tools. On top of that there are fewer vendors, for example the PGI compiler software that is on virtually all of the computers was recently purchased by NVIDIA, as were ARM and Mellanox. NVIDIA aspires to be, and is becoming, a world leader in HPC but there are fewer and fewer opportunities for others to use the same technology such as the compilers and debuggers.

The facilities are also amazing. When Bland was in graduate school in Computer Science he never thought about the space, power, and cooling needed for computers. ORNL now has an order of magnitude more space available for computing – moving from 7,000 sqft in a room above the library on a floor that could hold 60lbs per sqft to now 70,000 sqft with raised floors that can hold 600lbs per sqft (Summit sits on bare concrete and can hold even more weight). The power and cooling capacity has increased from 2MW to 100MW, and in some cases there has been a need to build tools. For example, the ORNL facilities team built a dolly bolted on the side of a disk cabinet to roll it down the hall to maintain its integrity. The computing people, along with the facilities people, make these systems work.

LCF requirements are outpacing the commercial marketplace. More cycles are needed than can be purchased on the commercial market. The number of leadership systems is small. Because there is a relatively small number of vendors trying to design, build, and sell the leadership systems and R&D funding through government channels will be required.

Traditional HPC applications require very high memory bandwidth and interconnect bandwidth in addition to high performance in the floating point operations per second. Not all commercial applications require that very high bandwidth; if bandwidth is the most expensive thing to deliver, and if most systems sold do not require it, vendors are not going to produce those systems without somebody paying for the engineering work or for the R&D. While it is possible to build commercial cloud systems that can meet these leadership requirements, the cost of the bandwidth is prohibitive unless those cycles are being continuously sold and the system is highly utilized – DOE runs its systems between 90-100% utilization most of the time. For leadership systems the cost of using the commercial cloud is at least 2-3 times more than having the computers in-house.

The transition to accelerator-based nodes using CPUs and GPUs has provided a big increase in performance and efficiency. Quantum computing has promise to provide similar gains, but it has a lot of complexity as well. AI may provide answers to some of the problems without having to compute the answer but AI is not appropriate for all types of problems. The cost of the current leadership systems is approaching $1B per system; with a nominal 5 years lifetime that is extremely expensive – this is a very capital intensive business.

An all-of-the-above strategy to build the next generation systems is required. The strategy requires R&D into new devices that are more efficient, leveraging commercial technology, performing calculations in the most efficient way, aggressively turning off circuits that are inactive, and continuing to improve algorithms to be more efficient.

A computer that does these things might be an exascale ecosystem or it might have more heterogeneity to improve efficiency. Such a computer will utilize off-the-shelf technology wherever possible to reduce costs, it will consist of distributed powerful computers and storage, each component will be tailored to the task to be done at a specific location, and it will only move data and the calculations to the exascale computer when needed to solve a problem. There may even be federated instruments at the edge.

Predictions in the technology space are dangerous. Bland advised taking what he said with a grain of salt. Bland closed by stating he could not think of a better place or a better group of people to have spent the last 40 years with and he will miss everyone.

Discussion

Reed thanked Bland for his contributions and remarked that not many people have seen that many orders of magnitude changes in technology in their professional career.

Monga questioned the comment that the supercomputing bandwidth needs to be an order of magnitude more than the commercial clouds or data centers. He suggested that it is the commercial sector that is driving bandwidth rather than supercomputers. **Bland** clarified that he was talking about the bandwidth needs of the applications, memory bandwidth as well as the interconnect bandwidth in the parallel systems. **Monga** said he was referring to network bandwidth interconnecting nodes within the data center. **Cerf** confirmed that the data centers are huge consumers of the capacity. **Reed** added that there was a scale at Microsoft he had never seen before.

Bland commented that supercomputers run a different programming model requiring a group of processors working together on a single application. **Cerf** countered that commercial data centers must take in tons of data, index it, and provide answers in .018 seconds. Most commercial data centers do not to that for certain kinds of computation like those discussed today. However, there will be a shift in that direction as the commercial demands increase and as government requires the use of data centers for more conventional computation.

PUBLIC COMMENT AFTER WHICH ASCAC WILL ADJOURN FOR THE DAY

None.

Respectfully submitted,
Tiffani R. Conner, PhD, PMP, AHIP October 29, 2020