Meeting Minutes
ADVANCED SCIENTIFIC COMPUTING ADVISORY COMMITTEE (ASCAC)
December 12, 2018
Teleconference Meeting

ASCAC Members Present

Keren Bergman
Martin Berzins
Barbara Chapman
Jacqueline Chen
Silvia Crivelli
John Dolbow
Jack Dongarra
Thom Dunning
Tim Germann
Tony Hey
Gwendolyn Huntoon
Richard Lethin
John Negele
Daniel Reed (Chairperson)
Vivek Sarkar
Krysta Svore

ASCAC Members Absent

Vinton Cerf
Susan Gregurick
David Levermore
Satoshi Matsuoka
Linda Petzold
Dean Williams

Also Participating

Katerina Antypas, National Energy Research Scientific Computing Center, (NERSC) Lawrence Berkeley National Laboratory (LBNL)
Christine Chalk, ASCAC Designated Federal Officer (DFO), Advanced Scientific Computing Research (ASCR), Office of Science (SC), Department of Energy (DOE)
Sudip Dosanjh, NERSC, LBNL
Barbara Helland, Associate Director, ASCR, SC, DOE
Daniel Jacobson, Oak Ridge National Laboratory (ORNL)
Thorsten Kurth, NERSC, LBNL
Maureen Leavitt, Oak Ridge Institute for Science and Energy (ORISE)
Paul Messina, Argonne National Laboratory (ANL)
Andrew Siegel, Exascale Computing Project (ECP), ANL

Attending

Laura Biven, SC, DOE
Buddy Bland, ORNL
Ron Brightwell, Sandia National Laboratory (SNL)
Ben Brown, ASCR, SC, DOE
David Brown, LBNL
Rich Carlson, ASCR, SC, DOE
Scott Collis, ANL
Erik Draeger, Exascale Computing Project (ECP), Lawrence Livermore National Laboratory (LLNL)
Mark Guiton, Cray, Inc.
Bruce Hendrickson, LLNL
Tony Hey, Science and Technology Facilities Council, United Kingdom
Paul Hovland, ANL
Doug Kothe, ECP, ORNL
Sriram Krishnamoorthy, PNNL
Alex Larzelere, Council on Competitiveness (COC)
Carolyn Lauzon, ASCR, SC, DOE
Steven Lee, ASCR, SC, DOE
Michael Martin, LBNL
Sonia McCarthy, SC, DOE

Despina Michael-Grigoriou, Microsoft Computer Games and Emerging Technologies Research Lab (GET Lab)
Christopher Miller, ASCR, SC, DOE
Kathryn Mohror, LLNL
Jeff Nichols, ORNL
Lucy Nowell, ASCR, SC, DOE
Claire Oto, DOE
Mike Parks, SNL
Tom Peterka, ANL
Mike Procario, DOE
Michelle Schwalbe, National Academies of Sciences, Engineering, and Medicine (NASEM)
Jim Sexton, LLNL
James Stewart, SNL
Sharlene Weatherwax, SC, DOE
Mike Weaver
Carol Woodward, LLNL
S.J. Ben Yoo, University of California, Davis

Wednesday, December 12, 2018

OPENING REMARKS FROM THE COMMITTEE CHAIR

Christine Chalk, ASCAC DFO, confirmed a quorum present at 11 a.m. Eastern Time (ET). She stated a quorum was unattainable for the second day of the meeting; therefore, it was cancelled. **Thom Dunning,** ASCAC, chaired the meeting until **Dan Reed,** ASCAC Committee Chair, arrived. Dunning called the meeting to order and introduced the first speaker.

ATTACKING THE OPIOID EPIDEMIC: DETERMINING THE EPISTATIC AND PLEIOTROPIC GENETIC ARCHITECTURES FOR CHRONIC PAIN AND OPIOID ADDICTION

Daniel Jacobson, Chief Scientist for Computational Systems Biology, ORNL, presented a computational approach to examining traits and relationships in complex biological systems. While conventional genome-wide association studies (GWAS) are applied for single traits, this new approach converts what might be thousands of single trait results into networks with nodes for genes and phenotypes to identify relationships. The algorithms can conduct this activity over the entire system, and rank the order of genes that support the opening question. Results are effectively integrating numerous lines of evidence that can be used to test theories or hypotheses without the same concerns associated with conventional GWAS. For example, Type I and II errors can be significant in conventional GWAS, but mitigated with this new approach. More work is needed to address epistasis.

Jacobson presented two applications: genetic relationships with three-dimensional protein structures in cells for biofuel plants, and analyzing environmental variables within 50 years of climate data over space and time with plant characteristics. Examining epistasis moves from additive to multiplicative genetics and will require even more sophisticated analyses and computational capacity. Tools such as explainable artificial intelligence and single nucleotide polymorphisms (SNP) correlation networks are under development to produce integrated scalable drosophila phenotyping.

The algorithm CoMet was used to process the Analytics for Medical Precision to Improve Outcomes (MVP-CHAMPION) database from the Veterans Administration in a project examining the genetic basis for pain and opioid addiction. This work was recognized recently with a Gordon Bell prize. These projects were unthinkable before Summit. Projects in discussion

now would take 2B Titan central processing unit (CPU) hours, or 340,000 years on a laptop, but only 17 days on Summit.

Discussion:

None.

UPDATE ON EXASCALE APPLICATIONS

Andrew Siegel, Director of Application Development, ANL, provided an update on the 33 subprojects comprising the ECP Application Development focus area. These projects are pushing science forward in their domain areas. Simulation success must be clear, but there are no comparable baseline metrics from other projects. Siegel demonstrated the complexity and specialized architecture with project summaries from the SC and the National Nuclear Security Administration (NNSA) and described six common challenges these applications face to prepare for exascale.

Siegel summarized the co-design project at the Center for Efficient Exascale Discretizations (CEED), led by Tzanio Kolev at LLNL. CEED is intended to affect as many application projects as possible. CEED goals include developing algorithms and software to enable more efficient high performance computing (HPC) simulations on a variety of hardware. CEED is working better than imagined, with major impacts on the community. Domain scientists are the primary audience for this work and these experts are hard to find. Siegel underscored the need for early adoption of intermediate (100 petaflop) systems to track progress and identify issues while also balancing domain expertise and performance engineering. Siegel closed by stating facilities' engagement programs are critical to achieve their goal.

Discussion

Berzins expressed his appreciation for the outstanding work presented and asked about the thinking beyond current generations of machines to "future-proof" software designs. **Siegel** stated this is a very important issue. Many of the projects have a life outside of ECP and want to continue beyond ECP. While there is a lot of interest in programming with one degree of separation from hardware, Figures of Merit (FOM) are still high and must be met; cases where performance is in a tiny fraction of the code can achieve this balance. Siegel cited the successes of Kokkos, RAJA and Legion in addressing this concern and acknowledged this will be something to watch.

Dunning stated commitment to support new applications is important and he cautioned against integrating software if it is not supported in the future. **Siegel** acknowledged the point and explained there is a need to discuss potential adoption of long-term support. **Dunning** expressed his hope that it will be a significant topic of discussion for the team planning the activities beyond ECP.

Alex Larzelere and **Sonia McCarthy** expressed their appreciation, stating it was a great presentation.

EXASCALE DEEP LEARNING FOR CLIMATE ANALYTICS

Thorsten Kurth, Application Performance Specialist, NERSC, LBNL, acknowledged the importance of the collaboration between DOE labs, software and hardware vendors, and shared that the project recently earned a Gordon Bell prize. Kurth presented the computational development aspects of the effort. Exascale is needed to achieve one square kilometer resolution, and this scale presents three challenges: data management issues to feed the data fast

enough to keep the graphics processing units (GPUs) busy, multi-node synchronization across 27,360 GPUs, and convergence and accuracy in large effective batch sizes.

Using the Tensorflow deep learning framework and Horovod distributed-training enabling framework, provided the balance of scale, flexibility, and time. Two test systems, Piz Daint in Switzerland and Summit in the USA, demonstrated portability to other computing systems. One key development is application of gradient lagging to increase the number of GPUs. The system achieved a sustained rate of 999 petaflops/second with a peak of 1.13 exaflops/second using 4,560 nodes, 27,360 GPUs. Kurth concluded that deep learning and HPC converged to achieve exascale performance.

Discussion:

Hey stated the computations on 27000 GPUs are impressive but was interested in the scientific results of this work. **Kurth** explained this presentation demonstrated "proof of concept" and that further work is underway to get the climate results. He deferred comments on the results until the effort is complete.

Svore thanked Kurth for a great presentation. **Sharlene Weatherwax** thanked Dan Jacobson and Thorsten Kurth for demonstrating the exceptional value of the partnership between DOE's Biological and Environmental Research (BER) and ASCR programs.

The meeting paused for a break at 12:50 pm and resumed at 1:02 pm. **Chalk** confirmed a quorum was in place.

REPORT FROM THE SUBCOMMITTEE ON 40th ANNIVERSARY ACCOMPLISHMENTS

Messina presented a high-level summary of the subcommittee's documented ASCR impacts over the past 40 years. He recommended reading the files provided which represent drafts at various stages of preparation. ASCR-funded research impacted almost all of the supercomputing ecosystem through collaboration among national labs, academia and industry. The subcommittee's report will include such things as the important role of software policy and type of license and the value of block/base funding over long periods to achieve breakthroughs. ASCR challenges going forward will also be summarized; for example, the need to emphasize high quality sustainable software to drive basic scientific discoveries for the next generation.

Messina summarized highlights of the draft report, including computational science, applied math, facilities and the impact on industry and the workforce. Several case studies of research efforts that developed into applications are also in progress. Finally, Messina acknowledged the tremendous efforts of the subcommittee and specifically Eli Dart and Inder Monga.

Discussion

Dunning thanked **Messina** for the presentation and commented on its broad look at many accomplishments of ASCR. **Dunning** added that people in the field understand all of this but the larger population does not.

Sarkar thanked **Messina** for an excellent presentation and appreciated seeing all of ASCR's accomplishments. **Sarkar** posited the accomplishments described over 40 years are a result of venture investments made in early years. He noted early career staff in computer science and applied math perceive a reduction in this type of research investment funding (not initially tied to a high quality deliverable). **Sarkar** stated the community would appreciate a reflection of

this issue in the report. **Messina** said part of the subcommittee's charge is to give advice on future funding strategies and the issue will be included in the report.

Dan Reed joined the meeting, thanked **Dunning** for standing in for him, and introduced the next speaker.

VIEW FROM GERMANTOWN – UPDATE ON ASCR

Barb Helland, Associate Director of ASCR, reminded ASCAC that ASCR has had an approved budget of $935.5M since October with specific funds for each of the three leadership computing facilities (ORNL, NERSC, LBNL), the Computational Sciences Graduate Fellowship program, ESnet and Scientific Discovery through Advanced Computing (SciDAC). Congress specified $13 million of the SciDAC funds will be used for the algorithm and method development for the Artificial Intelligence (AI) and Big Data focused on the needs of SC user facilities. Congress further specified this is the only funding within SC for this initiative and a DOE must submit a briefing before that portion of SciDAC funds is released for use. DOE will submit this briefing soon.

While the overall research funding is higher than last year, funding for Applied Math and Computer Science projects is reduced. ASCR implemented some adjustments to ease the impact on the research budget. Argonne Leadership Computing Facility (ALCF), NERSC, and Oak Ridge Leadership Computing Facility (OLCF)) each contributed $1 million to the SBIR program to reduce the amount coming from research. Basic Energy Sciences (BES) is interested in partnering with ASCR, and there is additional funding for quantum algorithms. Releasing the $13M for SciDAC will help as well. Exascale Computing was funded as requested.

Helland highlighted ASCR's significant achievements, including two 2018 R&D 100 awards (Darshan at ANL and Multinode Evolutionary Neural Networks for Deep Learning (MENNDL) and SuperComputing 18's (SC18) two Gordon Bell awards and the continued ranking of Summit as the world's most capable computing system.

Helland listed reviews planned in 2019 and provided perspective questions based on the ECP Annual review conducted in October 2018. She stressed the importance of coordination between ECP and the facilities standing up the system to inform schedules and milestones. Three workshops are planned in late January and early February: In Situ Data Management, Applied Math Principal Investigators Meeting, and the cross SC Quantum Information Science Principal Investigators Meeting. Of special note, ideas for the proposed DOE Quantum Information Centers submitted by PIs and selected by Associate Lab Directors will be presented during a lightning round at this workshop

Helland referred to a recent presentation of NASEM study results on sexual harassment of women, climate, culture and consequences. She will include this presentation at the April meeting. ECP and the facilities standing up the system to inform schedules and milestones. Three workshops are planned in late January and early February on In Situ Data Management, Applied Math Principal Investigators Meeting, and Quantum Information Science. Of special note, ideas for the proposed DOE Quantum Information Centers submitted by PIs and selected by Associate Lab Directors will be presented during a lightning round at the Quantum workshop.

Helland announced that Lucy Nowell will retire at the end of 2018. Helland acknowledged Nowell's substantial leadership over many years and indicated plans are underway for a recognition of Nowell in January.

Discussion

None.

UPDATE ON NEW SUBCOMMITTEE

Dan Reed, Chair, ASCAC, summarized the charge issued by **Helland** in September to assemble a subcommittee to identify key elements of ECP that need to be transition into ASCR programs or other new SC initiatives after the end of the project. Previous ASCAC Chair, **Roscoe Giles,** will chair this subcommittee and is currently working to populate it. Reed directed anyone with ideas to speak with Professor Giles.

Discussion

None.

NERSC-9 UPGRADE

Sudip Dosanjh, NERSC Division Director, LBNL and **Katie Antypas**, NERSC-9 Project Director, presented a summary of the status of NERSC-9 and how this system will serve the broad community of NERSC users. Planning for NERSC-9 began in 2015 and the system will be delivered in 2020. Lessons learned and successes of NERSC-8, as well as ALCF and OLCF facilities, were given serious consideration. NERSC-9 is the first system to consider data management equal to simulation. The system is named "Perlmutter", after the 2011 Nobel Prize winner who pioneered combining largescale simulations with experimental data. The system will be three to four times faster than Cori, with both CPU only nodes and GPU plus CPU nodes. Estimated capacity is 8 billion hours for the CPU nodes and at least 16 million hours for the combined GPU/CPU nodes. A NERSC Exascale Science Application Program (NESAP) will be formed for Perlmutter. The call for proposals will open December 18, 2018.

Valuable information to guide NERSC-9 development came from a survey of users' software package usage and preferences. Results indicated that Python is ubiquitous, Jupyter is very popular, and MongoDB is preferred for many-use databases

The current timeline indicates the GPU Rack for Cori will be available to NESAP users at the end of 2018. The contract to produce Perlmutter is in place, with system delivery in the fall of 2020. NESAP users should have access in early 2021 and system acceptance in the fall of 2021.

Discussion

Lethin asked how the NVLink and Cray GPUs will bridge into the general fabric of the machine and what software abstractions will be available. **Antypas** explained they are connected with NVLink 3 and the GPU to CPU link is PCI Gen 4. The interactions with slingshot and the network interface card (NIC) will be with PCI Gen 4. To achieve a balanced network for the CPU-only node, there is one NIC per node whereas there are four NICs per node for GPU. This design offers substantial bandwidth capabilities outbound and inbound.

Martin Berzins expressed his appreciation for the presentation and asked about considerations for users who lack the compute intensity for effective use of the GPU, even if their codes are mapped onto the GPU. He suggested this is a complex issue to consider for future machines, and relates to the previous theme that basic research is needed to solve this and other issues. **Dosanjh** agreed stating this issue was part of the reason for a CPU partition on NERSC-9. If users want to

make progress and need more computing, NERSC is willing to work with them, but they also need to work with the broader community.

REPORT ON BASIC RESEARCH NEEDS (BRN) FOR MICROELECTRONICS WORKSHOP

Reed and **Helland** provided a summary of the BRN on Microelectronics held with BES, ASCR, and High Energy Physics (HEP) in October 2018. The objective was to identify transformative science challenges to move technology forward 10-30 years in the future. To maintain our competitive position, DOE must look beyond exascale to a co-design innovation ecosystem that includes materials, chemistry, devices, systems, architectures and algorithms. Researched must be closely integrated, with emphasis on energy-relevant applications. The scope of the BRN included examination of Complementary metal–oxide–semiconductor (CMOS) and beyond CMOS technologies, excluding quantum information science.

Four panels were convened to determine priority research directions (PRDs). The panels were big data collection, analytics, and processing for SC facilities; Co-design for HPC beyond exascale; Power control, conversion and detection; and Crosscutting themes. The five PRD are: Flip the current paradigm: Define innovative materials, devices, and architecture requirements driven by applications, algorithms and software; Revolutionize memory and data storage; Reimagine information flow unconstrained by interconnects; Redefine computing by leveraging unexploited physical phenomena; and, Reinvent the electricity grid through new materials, devices, and architectures. The report will provide details for the PRD as well as details of each panel discussion. Publication target is February 2019.

Helland thanked Dan Reed, Cherry Murray, and Supratik Guha for co-chairing the workshop, specifically recognizing their effort and time to identify the scope of the committee. She also thanked Gil Herrera (SNL), Andy Schwartz (BES), Robinson Pino (ASCR), and Eric Colby (HEP) for their efforts.

Discussion

Lethin asked about the action plan and the timeframe. **Helland** explained the report is critical to make the case for a budget. Given the 2020 budget will go to Congress in February, this report will likely influence the 2021 budget and funding opportunities will follow. **Reed** added this is the fundamental partnership in action. The committee generates ideas for the future and agencies create programs and approaches to make them reality.

PUBLIC COMMENT

None.

Reed adjourned ASCAC at 3:03 p.m. eastern time.

Respectfully submitted,
Maureen Leavitt
ORISE
January 11, 2019