Summary of ASCR PI Meeting

Scott Klasky, David Rogers and John Wu

Executive Summary

The 2017 ASCR Computer Science Principal Investigators Meeting: Resilience, SSIO, Design Space and SDMAV was held on March 14-16, 2017 in Bethesda MD. A diverse group of 94 ASCR researchers and program managers, with expertise across a wide range of scientific and technical domains, met to discuss their current work, generate ideas for future collaboration, and identify common themes across the portfolio. Through publications and software artifacts, they highlighted their work advancing the goals of the Department of Energy ASCR office.

The purpose of the meeting was to review the current state of research through panel sessions, breakouts, and poster sessions. A primary focus was to provide a forum for the Program Manager to assess the current state of research through technical talks and poster sessions. Participants were also asked to:

1. Understand the evolving hardware architecture for next-generation supercomputers and future exascale computing systems and evaluate the impact the changing architecture will have on their research;
2. Understand the goals and the current state of the Exascale Computing Project (ECP); and
3. Discuss the requirements to support scientific discovery at extreme scale, and to integrate various research activities and projects so they provide the innovative data management, data analytics and visualization technology solutions needed.

Modern scientific discovery is encountering a number of constraints, as the scale of data from large scientific experiments (computational, experimental, and observational) continue to increase faster than the increase of computing power. The topics discussed in this meeting focused on storage, I/O, data management, data analytics, data visualization, in situ data processing, power/energy, reliability, and programming models.

There were several broad themes and common opportunities for research collected from the sessions:

- It is critical for scientists to understand and take advantage of the tradeoffs across the high performance computing (HPC) solution space (power/energy and performance, resiliency vs. performance, storage hierarchy and data artifacts) through measuring, modeling, and creating tools for managing these tradeoffs. Scientists must be able to make decisions about tradeoffs so they can create workflows that address their specific problems.
- Co-designed solutions across technical domains (algorithms, I/O, in situ analysis, workflows, etc.) continue to offer the best approach to having the maximum impact on science. Tools that simplify workflows, facilitate problem definition and create efficient solutions on a range of platforms are critical to impact next-generation scientific discoveries.
- Usability and composability of workflows, applications and analysis tools is critical to enable the community to move forward towards science goals. We need to continue to create solutions that make things easier for scientists, allowing them to focus on the science, increase the value of the data artifacts created, and promote effective decision-making about the tradeoffs available at extreme scale. We must also create effective methods of evaluating and testing these tradeoffs.
- We need novel methods for interactively exploring and querying the complex data artifacts created during extreme scale runs. This includes in situ methods for automatic and human-assisted selection of data of interest, domain-specific methods of appropriately reducing data, flexible methods of creating valuable data extracts, and interactive tools for exploring these artifacts.

Detailed Information from Panels

The meeting was organized into nine panels of related work. Researchers presented summaries of their current work, and the entire group of attendees participated in a discussion following each panel. Breakout sessions were held at the end of each day to encourage further interplay of ideas. Detailed notes for each session are summarized in the following sections.

Panels on **Reducing the Time for Data Movement to and from Storage** and on **Exascale Storage Challenges** focused on I/O and storage related research. This is a critical area due to the ever-increasing mismatch between I/O and compute speeds, and it requires scientists to make intelligent tradeoffs between what is computed, what is stored, and how data is moved in order facilitate access (in both short and long term). This encompasses both experimental and simulation data. The area is made more challenging by rapidly changing and increasingly complex memory hierarchies. The panel presentations and discussions emphasized common challenges, but also dealt with areas requiring different approaches. One consistent theme that was highlighted was the impact this area has on the productivity of scientists, programmers and analysts. Main discussion points included the following:

- Compression and data reduction approaches and data refactoring are critical, due to the constraints of the I/O and compute mismatch. With larger simulations, more data must be saved, but storage will continue to be a limited resource on future systems.
- In-situ methods are commonplace, but diverse and flexible solutions are critical in order to have impact across science domains.
- A diverse set of approaches is needed, including both automated and human assisted methods, feature-based, statistical, and highly-compressed, lossy data reduction approaches, and streaming solutions. These will support a wide range of science use cases.
- Progressive exploration and data refactoring approaches provide an opportunity for different modes of exploration of large data, and can make more of a simulation's results available to scientists. Since information foraging is difficult at large scale, these techniques can provide methods that encourage scientists to touch and explore a high percentage of the data, despite the fact that there will be more of it than ever before.
- Novel I/O approaches are needed for the future, though the existing POSIX-based approaches will still be widely used in the short and mid-term. Legacy codes must access existing files, and realistically it will be years before POSIX is replaced. New approaches include disaggregation, which allows use of different memory hierarchies, but there are a variety of approaches to this. New concerns including managing power/energy and effectively using new non-volatile technologies will likely require an enhanced interface to resolve data placement tradeoffs. Possible approaches include object stores, name-value systems, and versioned storage.
- Naming was a point of discussion in several contexts. Further research is needed to address the solvability of naming, approaches to granularity and versioning, and how naming can be done in context.
- There is a growing need for research and tools for capture, curation and utilization of metadata and provenance. This will have an impact across science domains.
- Data will be consumed in a variety of ways that may not map well to the way that the data was stored. Therefore, methods and tools for data sampling, transformation, and reorganization are important. These must be flexible, to adapt to a range of use cases.
- It is recognized that emerging technologies such as FPGAs, cloud computing and can have an impact in this domain, and this must be reflected in ongoing research.
- The range of expectations across different science domains leads to very diverse requirements in future I/O and storage.

Panels on **Understanding and Managing Resilience**, on **Understanding, Optimizing and Programming for Power,** and on **Modeling and Understanding** brought together projects focused on constraints such as power and energy, programming for resilience, and understanding the tradeoffs in these and other spaces through modeling. The combination of approaches being explored promotes decision-making by domain scientists and computer scientists while developing scientific codes, and when running specific problems on specific hardware. The panels discussed the concerns about data read/write optimization, system performance, and resilience. Main discussion points included the following:

- Prediction and modeling are critical to help scientists make decisions in software design, algorithms, workflows and analysis. The complex tradeoff space requires both fine-grained and system-wide data from future systems about reliability, power and energy. Application teams and scientists want to know what areas of this complex system they can control. This requires a wealth of data from hardware, which we can only get through cooperation by vendors.
- Science on next generation HPC systems presents many clear resilience challenges, and solutions discussed have a wide range of complexity and flexibility.
- There is a spectrum of failure modes, with optimal performance at one end, and full reliability at the other. Scientists require solutions along this spectrum, which necessitates R&D into solutions in the tradeoff space on both current and future hardware systems. Management of resilience, power and energy is not a one-size-fits all activity, and flexible, well-modeled options must be available so individual codes/projects can make intelligent choices.
- It is important to show the tradeoff between power/energy and both performance and quality of service for low-level hardware as well as high-level capabilities. In addition, users must understand tradeoffs in accuracy, performance and power/energy.
- Innovations in hardware (custom silicon, float-based computation, FPGAs, etc.) can have a significant impact on this area of research. It is critical to have holistic, flexible (with respect to data types and time resolution) and real time system monitoring in order to influencing future hardware and software decisions.
- The growing hardware complexity is seen as an impediment to domain scientists. Several attendees espoused the philosophy "hide as much as possible, expose what's asked for" to manage architectural diversity, complexity, and performance portability.

Panels on **Data Analysis**, on **In situ Visualization and Infrastructure** and on **Scientific Workflows** and on **Exploratory Visualization Tools and Techniques** focused on data approaches and infrastructure for getting the best science possible out of existing and future facilities. The tools and techniques are critical to advancing scientific discovery, due to future system constraints, and the difficulty of foraging and exploring increasingly large and complex collections of data. Main discussion points included the following:

- Novel solutions for interactive exploration of data artifacts, with human-in-the-loop feedback, are critical to future data analytics. Reducing data through in-situ techniques, feature detection, compression and sampling will not by itself make data easier to explore and understand.
- In situ approaches to handling extreme scale data are central to providing options to scientists in storing the best data for their specific science needs, and in optimizing system performance. Future research should include optimizing algorithms for performance on next generation architectures. There are commonalities, but the best architecture choice for a particular analysis will be domain and application dependent.

- If individual science teams develop solutions in this space – especially at the workflow level - it may be difficult for those solutions to impact other domains. Developing shared expertise and capabilities, in contrast, will have a high impact on many scientific applications.
- Expressing complex operations through abstractions such as domain-specific languages and workflows can promote efficient use of capabilities across science problems.
- Usability, portability and flexibility of proposed solutions are critical for adoption, and continued research is critical to address developing science requirements. Leadership in these areas can be a distinguishing capability, relative to other countries.
- In general, we need data analysis and visualization algorithms that operate without user intervention, or with limited intervention. New algorithms can include machine learning, automatic feature detection, and domain-specific approaches, among others. A generalized approach to this will enable us to combat the increasing complexity in hardware, so we can continue to have impact across science domains.
- We need to enhance our methods of browsing and exploring large collections of data artifacts created through in situ methods. Future data pipelines will likely include automated and event-driven components, and it may not be easy for scientists to understand and assess what data is available after a simulation or experiment. Automated identification of "important" or "surprising" results will be crucial. Additionally, we need to support interactive post-processing of extracts and exploratory tools for next-generation extracts, including approaches that optimize perception and cognition on reduced data.
- It is important to develop a catalogue of workflows currently in use that includes details of the software architecture to enable us to identify potential common practices and tools. Today, many workflows are developed to support specific applications, or by an individual science team. These tend to be very domain specific. As a community we lack system software mechanisms to ease development of new workflows. Common building blocks are possible (in particular for managing data between components), and would be extremely valuable to the larger science community.