Research Opportunities in Operating Systems for Scientific Edge Computing

Contributors: Pete Beckman (ANL), Christian Engelmann (ORNL), Shantenu Jha (BNL), Jack Lange (Univ of Pittsburgh)

DOE Point of Contact: Hal Finkel <hal.finkel@science.doe.gov>

May 18, 2021

As scientific experiments generate ever-increasing amounts of data, and grow in operational complexity, modern experimental science demands unprecedented computational capabilities *at the edge* -- physically proximate to each experiment. While some requirements on these computational capabilities are shared with high-performance-computing (HPC) systems, scientific edge computing has a number of unique challenges. In the following, we survey current trends in system software and edge systems for scientific computing, associated research challenges and open questions, infrastructure requirements for operating-systems research, communities who should be involved in that research, and the anticipated benefits of success.

Current Trends

Edge computing is set to become an important feature in future HPC environments as scientific workloads begin to incorporate live data collection and analysis. Examples of this trend include the need to stream data directly from large scale scientific instruments and experimental installations to high end computing environments, large scale data collection and analysis needed by distributed sensor platforms, and real time computational steering of “smart” infrastructure including experimental and manufacturing processes. This trend will serve to dramatically expand the definition of traditional HPC resources to include much more diverse computational environments executing on a wide range of hardware and software platforms. Integrating this new infrastructure into existing HPC architectures and environments will pose a significant challenge across many levels, both technical and organizational.

The current state of practice in experimental science heavily involves human-in-the-loop activity for controlling instruments and computing resources, analyzing data, steering ongoing experiments and planning future studies. In several cases, the data stream is too complex for online control by humans, necessitating multiple runs with different parameter settings and neglecting online control. This results in reduced precision of experiments, particularly in materials synthesis applications.

Current human-in-the-loop solutions for steering instrument experiments with computational analysis have been developed on an ad-hoc basis for specific instruments and are limited in scope. Data formats, communication protocols and human-machine interfaces are usually incompatible with other instruments, analysis approaches and the envisioned autonomous online real-time control. For example, replacing instrument data analysis with AI or using AI in decision-making processes for steering or planning experiments may require wrangling of instrument data, interaction between instruments and edge computing resources, modification of vendor software to make instruments steerable, common data formats, and human-machine interaction for planning, observation and control. Also, computational simulations evaluating or planning experiments may have experiment data as input and experiment control data as output, requiring compatible interfaces between simulations and experiments.

Research Challenges

The integration of edge and HPC infrastructures will necessitate a re-examination of existing trust and security models deployed by modern HPC centers. Modern environments tend to rely predominantly on user based access controls enforced at the edge of HPC environments. These mechanisms rely heavily on out-of-band vetting of individual users and strong authentication mechanisms to ensure that only trusted entities are able to gain access to HPC systems. As the edge of HPC environments expands to include remote hardware installations and large numbers of cheap and widely distributed sensor devices, there will be a significant challenge to update existing security models to keep pace. Developing new security models suitable for this new system model poses a significant research challenge that will require a significant rethinking of existing trust models, authentication mechanisms, access control mechanisms, and security policies.

The DOE’s recent Artificial Intelligence (AI) for Science Report [1] outlines the need for smart systems, instruments and facilities to enable science breakthroughs with autonomous experiments, “self-driving” laboratories, smart manufacturing, and AI-driven design, discovery and evaluation. The DOE’s recent Computational Facilities Research Workshop report [2] identified smart systems and facilities as a broad challenge area with enabling automation and eliminating human-in-the-loop requirements as a cross-cutting theme.

One of the major bottlenecks for science is the limited speed at which experiments can receive feedback from computation and theory. Accelerating this cycle requires connected instrumentation with local/instantaneous computation using edge computing resources where feasible and with remote/urgent computation using leadership computing resources when necessary. Creating this cross-linking infrastructure enables computational analysis of experiment data for steering ongoing experiments and planning future studies. For example, experiment data analysis may discover anomalies in a sample that need to be further investigated using the same or another instrument. Similarly, a biological sample may have only a limited lifetime, requiring computational analysis of experiment data to be performed promptly to allow further study. Also, experiments in manufacturing may use data analysis to discover and treat manufacturing defects, requiring a real-time feedback control loop.

Infrastructure Requirements

Updating HPC security models to support edge based resources will require a number of infrastructural mechanisms. Incorporating large numbers of edge resources will dramatically increase the scale of the environment and will require scalable access control mechanisms that are able to support much more heterogeneous computing resources and more complex organizational structures. To effectively adapt to these environments will require access control mechanisms that support a high degree of semantic expressiveness in order to effectively encode increasingly complex and abstract policies. In addition access control policies will need to expand to include heterogeneous and dynamic collections of edge based hardware resources which will likely exhibit large amounts of churn and varying degrees of availability. This will require the expansion of security models to enable them to dynamically establish trust with potentially ephemeral edge resources in order to ensure end-to-end data security and access control. Trusted edge resources will need to support the ability to provide trusted execution environments with hardware level primitives as well as attestable and authenticatable software environments. Building on these features, edge enabled environments will require new protocols and secure/trusted Operating System/Runtime environments to enable the environment to incorporate the edge computing resources into the broader HPC ecosystem.

Smart instruments, “self-driving” laboratories and smart manufacturing employ machine-in-the-loop intelligence for autonomous decision-making. Human-in-the-loop needs are reduced by an automated online control that is capable of collecting experiment data, analyzing it, and taking appropriate operational actions in real time to steer ongoing or plan the next experiment. It may be assisted by a “black box” AI trained offline with archived data and/or with synthetic data created by a digital twin. It may also rely on causal models, reinforcement learning or advanced statistical methods. Human interaction for experiment planning, observation and steering is still part of the automated online control.

A software framework for connecting instruments with edge and center computing resources is needed that is capable of collecting, transferring, storing, processing, curating and archiving data in common formats. It also must be able to communicate with instruments and computing and data resources for orchestration and control, and with humans for critical decisions and feedback. Programming interfaces are needed that leverage community and commercial/custom software for instruments, automation, workflows and data transfers. This infrastructure must be able to orchestrate resources, control resources and transfer data across multiple administrative domains. It serves autonomous robot-controlled laboratories utilizing edge computing resources as well as DOE’s leadership instruments utilizing leadership computing facilities.

In order to support existing and future use-cases and requirements, OS (middleware) for scientific edge-computing must include agile mechanisms to provision distributed resources with collective properties, as opposed to individual properties, and possibly inconsistent capabilities and availability. Top down and centralized resource federation is unlikely to scale, be resilient, or even responsive to dynamic changes. Integrated application-systems decision making -- which combine elements of top-down with decentralization, promise scale, resilience and responsiveness. In order to achieve the necessary control and information flow for integrated application-systems resource and workload management, the infrastructure must also present unified abstractions and interfaces for resource and workload management to distributed applications. For example, user-facing abstractions and interfaces must include explicit performance and quality of service measures; system-facing abstractions and interfaces must include information about resource availability, and control for configuration and distributed resources selection.

Benefits of Success

Expanding HPC platforms to seamlessly integrate edge resources will allow future HPC systems to dramatically expand their utility by supporting emerging classes of workloads based on large scale and distributed online data collection. As experimental data collection needs increase, IOT sensor platforms proliferate, and the need for HPC computing resources extends more and more into edge devices, the ability to leverage large scale supercomputing resources in service of these workloads will provide significant new capabilities that would otherwise remain out of reach.

Autonomous experiments, “self-driving” laboratories, smart manufacturing, and AI-driven design, discovery and evaluation using a combination of edge and center computing and data resources enable faster science breakthroughs with autonomous steering of ongoing and planning of the next experiments. Computational analysis of experiment data in a computing environment that includes dedicated edge resources close to the instrument and DOE’s leadership center resources accelerate the feedback from computation and theory and significantly advance the speed at which science breakthroughs are achieved. For example, autonomous robot-controlled laboratories can perform experiments 24/7 using Bayesian design of experiments. In another example, autonomous experiments at DOE’s leadership experimental facilities can make use of DOE’s supercomputers, using edge computing capability to automatically orchestrate data transfer and supercomputer allocation reservation and job launch.

Contributing Research Communities

The work in building an infrastructure ecosystem that connects instruments with edge and center data and computing resources involves the edge computing community, the distributed systems community, the HPC global OS community, HPC operations personnel at computing facilities, the AI/ML and data analytics community, the decision sciences community, and, most importantly, the instrument science community.

Open Questions

Since most efforts in connecting instruments with edge and center data and computing resources involve problem-specific solutions and some available community software, the biggest open questions revolve around the general concepts, distributed hardware and software architectures, communication protocols and interfaces, resource orchestration and control, and data transfer, management and provenance. Additional open questions exist in leveraging AI/ML, data analytics and advanced statistics for computational steering and planning of experiments, including using existing software solutions. Other interesting open questions include the efficient and effective integration of human-machine interfaces and the role of virtual instruments as digital twins.

References

[1] AI for Science Report. March 2020. URL https://www.anl.gov/ai-for-science-report

[2] DOE National Laboratories' Computational Facilities – Research Workshop Report. ANL/MCS-TM-388. February 2020. URL https://publications.anl.gov/anlpubs/2020/02/158604.pdf