Reference ID: 11226652534_Bougie

Submission Date and Time: 12/16/2019 2:14:18 PM

This contribution was submitted to the National Science Foundation in response to a Request for Information, https://www.nsf.gov/pubs/2020/nsf20015/nsf20015.jsp. Consideration of this contribution in NSF's planning process and any NSF-provided public accessibility of this document does not constitute approval of the content by NSF or the US Government. The opinions and views expressed herein are those of the author(s) and do not necessarily reflect those of the NSF or the US Government. The content of this submission is protected by the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode).

Consent Statement: “I hereby agree to give the National Science Foundation (NSF) the right to use this information for the purposes stated above and to display it on a publicly available website, consistent with the Creative Commons Attribution-NonCommercial-NoDerivatives 4.0 International License (https://creativecommons.org/licenses/by-nc-nd/4.0/legalcode)."

Consent answer: I consent to NSF's use and display of the submitted information.

Author Names & Affiliations

Submitting author: Devin Bougie - Cornell High Energy Synchrotron Source

Additional authors: Werner Sun, CHESS; Elke Arenholz, CHESS; Joel Brock, CHESS

Contact Email Address (for NSF use only): (hidden)

Research domain(s), discipline(s)/sub-discipline(s)

Photon Science

Title of Response

Data and Workflow Management for Photon Science

Abstract

Over the past five years, the volume of data produced at CHESS has undergone a phase transition. Previously, CHESS users only collected a few gigabytes per week across all beamlines. Today, is it not unusual for a single user group to collect 10-20 terabytes of data in two days. In parallel, the algorithms and tools used to analyze and reduce this data has become so computationally intensive that users may

not have the necessary resources at their home institutions and rely on CHESS resources for the analysis. Although we have already taken steps to standardize storage and computing solutions across the facility, the future will require a more comprehensive approach to managing workflows and data.

Question 1 (maximum 400 words) – Data-Intensive Research Question(s) and Challenge(s). Describe current or emerging data-intensive/data-driven S&E research challenge(s), providing context in terms of recent research activities and standing questions in the field. NSF is particularly interested in cross-disciplinary challenges that will drive requirements for cross-disciplinary and disciplinary-agnostic data-related CI.

In the next five years, two related challenges in data handling can be foreseen. First, the volume of data produced by x-ray detectors will continue to grow to the point where the data will need to be reduced before it can be written to the central data store. Second, analysis techniques will become so computationally intensive that the data will need to be analyzed on multiple compute nodes in parallel, which requires repackaging and transporting the data for efficient analysis and aggregating the results from all the nodes. Both of these challenges introduce the requirement to track and manage data through its entire life cycle (provenance) including distributed storage tiers for raw, reduced, and analyzed data, along with associated metadata at each step.

Question 2 (maximum 600 words) – Data-Oriented CI Needed to Address the Research Question(s) and Challenge(s). Considering the end-to-end scientific data-to-discovery (workflow) challenges, describe any limitations or absence of existing data-related CI capabilities and services, and/or specific technical and capacity advancements needed in data-related and other CI (e.g., advanced computing, data services, software infrastructure, applications, networking, cybersecurity) that must be addressed to accomplish the research question(s) and challenge(s) identified in Question 1. If possible, please also consider the required end-to-end structural, functional and performance characteristics for such CI services and capabilities. For instance, how can they respond to high levels of data heterogeneity, data integration and interoperability? To what degree can/should they be cross-disciplinary and domain-agnostic? What is required to promote ease of data discovery, publishing and access and delivery?

As a first step towards managing the exponential growth in data volume at CHESS, we must develop real-time processing techniques to reduce the amount of data written to disk, using a combination of high-performance computing hardware and specialized data acquisition software. In addition to improving the quality and quantity of useful data produced, this will enable "real-time" decision making – guiding the experiments and data collection while it is taking place. As data volume grows, so does computational complexity. As a result, high-performance compute farms must be equipped with analysis frameworks that present a common interface for all users to make efficient use of the computational facilities. This framework would make use of advanced database, data compression, and cacheing techniques. It would include developing database infrastructure and applications to streamline discoverability and retrieval of user data, with the long-term goal of an optimized workflow that seamlessly connects users with their data. Improving and streamlining access to parallel computation resources will also enable the rapid prototyping and refinement of new analysis methods that would benefit the photon science community at large. Encompassing the above growth in data volume and computational complexity is a need for a comprehensive data management solution for our truly heterogenous environment with a wide variety of experiments, workflows, software frameworks and tools, complexity, users (with individual data privacy and control requirements), metadata, and data types, complexity, and formats. This includes the full data provenance, from collecting and storing raw data and associated metadata, processing and reducing that data on the fly, learning from and making real-time decisions on experiments, post-processing that data for deeper analysis and more complex rendering, and identifying and storing "relevant" data and results for long-term reference and analysis. In some cases, this includes significant simulation work that is required to inform the experiments and provide the basis for preliminary analysis of the live data. Flexible API-driven database applications to manage this wide variety of data and improve the accessibility and discoverability of the data are required, including a common workflow for collecting, storing, and analyzing this data. From the practical CI hardware perspective, the absence of standardized storage technology beyond 10Gb is one limitation to handling the exponential increase in data rates. We expect in the next three years or so the industry as a whole will converge on 50 or 100GbE. However, the industry is currently in a state of transition with networking switches and devices moving from 40Gb to 50 and 100Gb, but the available storage options are still limited to 10 and 40Gb.

Question 3 (maximum 300 words) – Other considerations. Please discuss any other relevant aspects, such as organization, processes, learning and workforce development, access and sustainability, that need to be addressed; or any other issues more generally that NSF should consider.

At CHESS, the vast heterogeneity of experiments, hardware and software tools, and analysis frameworks introduces significant challenges in developing a common infrastructure for managing the experimental lifecycle. The data management workflow will need to accommodate a wide variety of experiments with disparate and at times conflicting requirements. Historically, the software for controlling, collecting, and analyzing data from CHESS beamlines has been developed separately over decades by individual domain scientists at each beamline. As such, there is no common set of tools, workflows, or techniques deployed even within our single facility. In addition, basic software development best practices such as version control, branching, merging, and tagging has not been universally employed. This will be required to ensure the portability and reproducibility of our simulations, analysis, and experiments. Developing the tools to successfully process and reduce raw data in real-time, and analyzing the important data to inform the live experiment, both represent possible areas for collaboration with computer scientists. There is the potential for great gains by evaluating the existing software and tools used, and then optimizing the existing code to most effectively take advantage of the computational resources available. Real-time feedback from preliminary analysis of data will help researchers make the most efficient use of their limited beamtime at CHESS, there-by accelerating the rate of scientific discovery. The wide range of user expertise and comfort with computational frameworks and data management will require systems and resources that are intuitive, easy to use, and easily discoverable for both the domain scientists collecting the data and the end users analyzing the data.

-- End Submission –