SUMMARY OF COMMUNITY INPUTS FROM SC19

A Report from the LCCF Birds of a Feather Session Held 21 November 2019, Denver CO

Omar Ghattas, Dan Stanzione,
Rick Stevens, John West

TACC
TEXAS ADVANCED COMPUTING CENTER
LEADERSHIP-CLASS COMPUTING FACILITY

This material is based upon work supported by the National Science Foundation under Award No. 1925096. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

Summary of Community Inputs from SC19

NSF OAC-1925096
Texas Advanced Computing Center

Title	SC19 Requirements BOF Summary
Document ID	LCCF-REQ.001
Owner	John West Dan Stanzione

Change History

Date	Author	Description of Change
Dec 2019	John West	Initial document

Table of Contents

1 Introduction

This report presents a summary of inputs received from science, computational, and technology members of the NSF and broader HPC community at a Birds of a Feather (BoF) session held during the SC19 conference held 17-22 November, 2019 in Denver, CO. The event was hosted by Omar Ghattas (UT Austin), Dan Stanzione (Texas Advanced Computing Center, TACC), Rick Stevens (U Chicago), and John West (TACC) on Thursday, 21 November 2019 from 12:15pm - 1:15pm. The abstract for the session is below:

> *The NSF in the US has begun a planning process for a National-Scale HPC facility providing roughly an 0.5 – 1.0 exaFLOPS in computing capability. This BoF is one of a series of sessions that will guide the design of this facility. The Texas Advanced Computing Center is seeking input from both the science and technology communities regarding the key capabilities this facility must provide to effectively support large scale open science. The session will begin with a brief discussion of the roadmap for the proposed facility, followed by audience discussion and Q&A to stimulate a broad-based discussion of requirements.*

Manish Parashar, director of the NSF Office of Advanced Cyberinfrastructure, provided context for the session in terms of the NSF's larger long-term vision for cyberinfrastructure (CI). His talk touched on the growing role of streaming data and AI/ML in the science workflow, both of which are especially important as larger scale instruments come on line (the SKA, and so on). With respect to the distributed research enterprise, what is the "right" configuration for CI like the LCCF: is there a need for physically distributed components, for example data distribution and/or ingest centers located regionally or near specific instruments or groups of researchers? Parashar also touched on the issue of time scale: the design for the LCCF will be the centerpiece of NSF CI for 5-15 years, so the design needs to look to the long term and be flexible enough to adapt as unanticipated requirements emerge.

Then the session continued with framing comments by Stanzione; these are reflected in the session questions included in the next section.

Stevens commented on the reality that experimentalists are becoming a larger share of the traditional scientific computing user base, a trend that is reinforced each time a new class of instrument (telescopes, colliders, and so on) comes on line. This shift is something that could potentially change both the service and architecture emphases of CI providers and HPC centers. The use of code surrogates to do simplified science that provides fast estimates that are "good enough" is growing in some communities, and early proofs of functionality exist for applications that achieve speedups by alternating between ML kernels for speed and a full physics solution for accuracy and compliance with physical laws governing the simulation. If it continues this trend argues for systems that integrate robust ML performance alongside, rather than instead of, traditional floating-point operations needed for solution of systems of equations. Related to this issue is the observation that there are over 100 startups making AI accelerators, and the three major chip vendors – AMD, NVIDIA, and Intel – have all announced future products with integrated GPU/CPU capabilities on a single die, sharing memory. The shift from accelerator connected by PCI-e to a partnership of equal peers (in terms of system resource access, especially memory) represents an important shift in computing and will change the balance that software designers must achieve between moving work to a PCI-e connected accelerator or doing it in the CPU. Finally, as computing uptake continues in branches of science that have not been traditionally large users of compute CI – for example, linguists and sociology – these communities may stress elements of CI design that have not been emphasized, requiring adaptation in the types of systems and services we field.

The NSF design goal for the proposed Leadership Class Computing Facility (LCCF) is a 10x the capability of Frontera. Ghattas focused on the ways in which scientists may choose to take advantage of this – and many future – increases in the size of CI to improve various aspects of their research. For example, today's advanced climate models are already running at 5km resolution, and in most disciplines, there is a resolution threshold below with our understanding of the physics and governing equations begins to break down. However, assuming a factor of two refinement in each direction is possible before we reach that stage in most disciplines, 16x more computing capability is needed (assuming an explicit solution in three spatial dimensions and one time dimension). Many domains such as weather or ice sheet modeling do not model all of the physical processes that drive the effects we observe in the real world in order to make computations tractable with current resources. Incorporating those effects into fully coupled simulations will likely require a dramatic increase in capabilities (up to three orders of magnitude).

2 Motivating Questions

During the session introduction Stanzione provided several questions as a framework for the discussion.

2.1 What are the science challenges the community will face in 2025-2030?

- Current workloads at ever larger scales?
- How much will the workload shift towards AI/Data, and what impact will that have on the facility?
- How much will be driven by instruments and the data they produce – how much is truly "HPC" versus throughput at scale?
- What are the new science codes workflows we should target?

2.2 Put another way, how would you define 10x?

- 10x the problem size on current codes? On a new set of codes?
- 1/10th the execution time on current problems?
- 10x the total system throughput?
- 10x the users?

2.3 What are the architectures that support these challenges?

- How much heterogeneity/acceleration are you willing to tolerate?
- How much will the workload shift towards AI/Data, and what impact will that have on the facility?
- A single computing capability, or multiple specialized or domain-specific ones?
- What about requirements for data and filesystems?

2.4 What are the services the community will need in 2025-2030?

- Put another way, what is the scope of the facility?

- Computing – obviously we will provide computing, for Simulation/AI/Analytics, but ---
- Traditional Batch, interactive, event driven, streaming, persistent service support?
- Data?
- Storage – long term? Short term? Metadata and provenance? Support for reproducibility?
- Software? (Can we get to 10x without improvements?)
- Application and Algorithm support?
- Performance Profiling, Programming model and tools, workflow support, Persistent, composable computation and data services, Support for community software, Continuous integration, container repositories?

2.5 What should be central vs. distributed?

The word "facility" in LCCF implies a single entity, but a distributed facility is possible. What would that look like?

- Multiple compute centers?
- Regional Visualization Centers?
- Regional User Support Centers?
- Regional Visualization/Data Access?

3 Summary of Responses

Philip Maechling, Associate Director for Information Technology at the Southern California Earthquake Center, noted that there is a missing link between the process of discovery and being able to broadly apply the discovery as usable information about a phenomenon. Models need to be validated so that they are relevant in practice, but the characteristics of model validation workloads do not fit well into current NSF models. There is a potential need to directly interface large-scale CI and observation instruments and evaluate forecasts on a daily basis: that kind of routine processing is not practical today.

David Martin, Industry Partnerships and Outreach Manager at the Argonne Leadership Computing Facility, asked about the implications of federation and sharing both processing and data between data centers. Stanzione observed that the Frontera project at TACC is already building connections between that system and public clouds for hybrid workflows. Stevens noted that quantum computing may require partnering. Robotic experimental labs are yet another example where large-scale observational facilities can benefit from direct connections to CI, where the observational and computing systems work together on the front end to collect and transform data into a higher value product before a human even enters the processing loop. He suggested we may need a much more holistic thought process about the science and operational pieces, possibly where batch work may have to become pre-emptible in order to satisfy observational processes with real time constraints.

Philip Maechling remarked that with respect to the incorporation of ML into the more traditional approaches to scientific computing often requires training today. A catalog of data available to the public for training related to specific problems would be a valuable asset to the community. Stevens noted that there is precedent for this in the machine learning community itself, and pointed to the ImageNet example. New databases of the kind Maechling is talking about could leverage what already exists, and use crowd-sourcing, such as hackathons, to clean up the data. This raises questions about how to ensure that the needs of the communities that need those data can be supported over the long run. Is there a role here for automated processes or unsupervised learning to manage and curate datasets?

Claire Porter, UMN Polar Geospatial Center, talked about her community's need for stable, long-term data storage. In her work the assemble unique datasets during an observation period and then may need "10 years to get through it and publish". She emphasized the disruption that occurs when a storage CI provider or project is terminated and they scramble to move and rehost the data, then reestablish their workflows. She reflected on the need for a predictable storage service with predictable capacity that support publishing of data sets to be shared with other users. Some of these data sets can be quite large, and she also mentioned the requirement for high data throughput in order to reduce processing to reasonable time frames. In response Stevens commented that it would be valuable to have a "catalogue" of high throughput, high impact applications along with cases studies that to make the requirements more concrete.

Outreach and Training. There was a general discussion about outreach and training. An attendee from GA Tech noted that they have about 3000 users, with a lot of high throughput computing. Their experience with users (which is also consistent with experiences in the Department of Defense and other NSF users at TACC) is that users don't generally want to invest in making their codes perform better or use resources more efficiently as long as they are able to meet their own project deadlines – in fact, they often actively resist such help. Stevens suggested that one thing they are trying with their user community is to automate the processes of code improvement. Some efficiencies can be gained without invasive access to the code. West observed that for community codes with a large user base there may be an opportunity to work with the science team behind the code and have performance improvements incorporated into the main code itself, with benefits automatically accruing to code users. Unfortunately, many users, especially users of large amounts of cycles, do their own builds of community codes. While these builds often include custom or experimental science that is key to the user's research, the codes themselves are often not compiled effectively and the experimental code is often very inefficient. So opportunities to actualize this potential benefit aren't as abundant as would appear at first. Stevens also suggested that an approach that may achieve this same thing is to incorporate machine learning into codes such that they can adapt themselves to the performance environment in which they are running each time they run, providing an opportunity for improved performance even when any particular group of users isn't using the "canonical" (and, thus, the optimized) version of the code.