

Supporting Big Data Analytics at the NASA Advanced Supercomputing (NAS) Division

Piyush Mehrotra, L. Harper Pryor

{piyush.mehrotra, [laura.h.pryor}@nasa.gov](mailto:laura.h.pryor@nasa.gov)

NASA Advanced Supercomputing (NAS) Division, NASA Ames, Moffett Field CA 94035

NASA is the curator and host for large quantities of observation and simulation data. By some estimates, NASA's approximately 100 current active missions, observing Earth, the solar system and deep space, produce more than 50 PBs of data annually. Similarly, large-scale simulations run on its supercomputers are starting to produce terabytes and sometimes petabytes of data from each run. As NASA's flagship site for computational science and engineering at scale, NAS supports a user base that is at the forefront of data intensive and data driven science. Our users' codes use and generate very large datasets and analyzing these datasets to extract knowledge is a fundamental part of their workflows. Recently we at NAS reached out to many of our users to better understand the challenges they face in executing these workflows. We have categorized these challenges according to the major components of the workflows – discovery of data and tools, access to and movement of data, storage and management of data, algorithms/tools for performing the analysis/analytics and finally dissemination of the results.

For NASA data, discovery is challenging because in addition to sheer volume, we have a great variety of data, it resides in many archive centers, and the data is scientifically complex. Also because it is hosted in physically distributed storage centers, our users often have to move and manage very large amounts of data and hence need tools to meet this challenge. And then of course we have the challenge for the future of providing them with easy access to exascale processing platforms that meet both their computational and analytic needs.

Guided by the results of interviews with our users as well as many studies that have addressed the convergence of big data and big compute, we have augmented the NAS architecture roadmap with several initiatives aimed at meeting these challenges and satisfying the big data analytic needs of our users. In this paper we describe two of these that also happen to address key needs that emerged from last year's BDEC Workshop: 1) higher level support for scientific workflows to make the challenges of working with big data and big compute more transparent to the user, and 2) tighter integration of compute engines with analytic engines. We will also describe one of the research challenges that we think has the potential to provide more efficient analytics on some big data of particular interest to NASA.

The first initiative is part of a project called the NASA Earth Exchange, or NEX [NEX]. NEX is a collaborative platform that improves the availability of Earth science data, models, analysis tools and scientific results through a centralized environment that fosters knowledge sharing, collaboration, innovation and direct access to compute resources. In simpler terms, NEX is intended to bring together big data and big compute for the Earth science domain. Accordingly, NEX needs to help users develop, capture, manage, and re-use data driven scientific workflows that can harness all of these resources. To accomplish this, we have implemented VisTrails on Pleiades, NASA's premier supercomputer, and built tools to capture workflow information into MongoDB, which is then sent to VisTrails [Zhang 2013]. We are also working with ParaView both as a candidate alternate workflow tool and as a companion tool to VisTrails. Note that these tools will also address provenance, another need discussed in the BDEC 2013 Workshop report. Use cases that we want to support range from allowing a user to piece together datasets and tools into a workflow that addresses a scientific hypothesis to managing the wide-area

workflows associated with deriving new data products from satellite data residing at other datacenters, both within NASA and at other agencies including USGS, NOAA, and DOE. Our vision is to provide an environment capable of capturing the workflow so that it can be shared with colleagues who can then repeat the experiment and/or tweak the input data/algorithms to generate new knowledge.

The second initiative addresses increasing integration of analytic capability – in this case for visualization – with compute capability. For many years, NAS has pursued tighter integration of large-scale visualization using our 16 x 8 panel hyperwall with codes executing on our multi-petaflop Pleiades supercomputer [hyperwall]. Where in the past, visualization was a post-processing activity that could only begin after a compute run was complete, we developed the capability to integrate the visualization engine driving the hyperwall (a 128 node cluster) into the same InfiniBand fabric as Pleiades, giving it access to storage resources in our Lustre filesystem. This allows us to set up data streams from computation nodes to the visualization nodes via the InfiniBand I/O fabric while the code is executing. The intermediate data can be rendered and the images examined concurrently with execution (for computational steering) or stored for later analysis. This provides scientists temporal fidelity at much lower storage cost. We are continuing to enhance both the hardware and software environment to enable scientists to easily and efficiently use these capabilities for visualization and analytics.

The two initiatives we just described deploy reasonably mature technology to address the big data challenges our users face. For the longer term, one of the research challenges we face in NASA is how to take advantage of what we know about our big data – in particular our Earth science data – to make analytics on these data more efficient. We don't want to touch all of the data if we don't have to. We know a lot about the structure of the data that might be used to steer the computation toward the subsets of the data that are applicable to the query – and not use the subsets we know are not relevant. This is the good news side of what we know about our data – the structure could be used to make processing more efficient. The bad news side is that there is a lot of complexity hiding behind the data and this complexity is critical to using it properly. For example, we talk about remote sensing of atmosphere and land temperatures from space, but a satellite does not really measure temperature. It measures radiance. To derive a temperature requires a lot of knowledge about the sensor – things like the transfer function of the detectors. As another example, for a satellite that is nominally in a sun synchronous orbit, what if the orbit has drifted? So there is a tremendous amount of information in the metadata and an aspect of the discovery challenge is: How do we make the knowledge in the metadata more easily accessible to the user? What are the approaches (representation, tools, algorithms) necessary to reason about the metadata.

Today NAS operates in the realm of petaflops and petabytes, but we know that the not too distant future is exaflops and exabytes. In that future, we can foresee many research and technical challenges as we attempt to support the entire workflow of NASA scientists including: discovery, access, transportation, management, and dissemination of big data along with tools for analysis and analytics of such data. We must never lose sight of the fact that our product is the scientific and engineering knowledge that we extract from big data.

1. [hyperwall] http://www.nas.nasa.gov/hecc/resources/viz_systems.html
2. [NEX] <http://nex.nasa.gov>
3. [Zhang 2013] Zhang, Votava, Lee, Chu, Li, Liu, Liu, Xin, Nemani, "Bridging VisTrails Scientific Workflow Management System to High Performance Computing," 2013 IEEE Ninth World Congress on Services (SERVICES), pp. 29-36, 2013.