Results from Large Scale STAR Raw Data Reconstruction on NERSC HPC

By R. Jefferson Porter (LBNL), Jerome Lauret (BNL), Mustafa Mustafa (LBNL), Jan Balewski (LBNL)

Introduction

The STAR experiment at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory (BNL) is a leader in the study of strongly interacting QCD matter, the Quark-Gluon Plasma (QGP), that is generated in energetic heavy ion collisions. STAR consists of a large, complex set of detector systems that measure the thousands of particles produced in each collision event. Detailed analyses of billions of such collisions have enabled STAR scientists to make fundamental discoveries and measure the properties of the QGP.

STAR is operated with continuous (24x7) running for about 6 months each year, during which it records large volumes of data from RHIC collisions of ions or protons. In recent years, STAR datasets have reached billions of events, with data volumes at the multi-Petabyte scale. Raw data signals collected by the detector electronics are processed using sophisticated pattern recognition algorithms to generate the higher-level datasets that are used for physics analysis. This initial processing stage, known as raw data reconstruction, requires many months of processing on dedicated STAR computing resources at the RHIC and ATLAS Computing Facility (RACF) at BNL for each year of data-taking.

In an environment where new data is collected every year, even normal delays to prepare calibrations or tune software algorithms can cause processing backlogs. In addition, introduction of new state-of-the-art subsystems bring new challenges to understand and calibrate these devices, resulting in even longer delays. STAR scientists are highly motivated to reduce or eliminate such backlogs and have pioneered the use of non-dedicated computing resources external to RACF, such as Cloud resources and High-Performance Computing (HPC) systems at NERSC. Elastic use of such facilities by High Energy and Nuclear Physics experiments has until now been limited largely to Monte Carlo (MC) simulations, whose processing tasks are relatively easy to organize, require no input data, and produce modest amounts of data. In contrast, raw data reconstruction is a complicated process with rigorous quality assurance metrics, which consumes and generates vast amounts of data and is currently the most CPU-consuming activity carried out by STAR computing operations.

A recent processing backlog threatened to delay the analysis of a novel STAR data set acquired in 2014 and 2016. At that time, the NERSC Cori Phase I system had just been commissioned, which included a number of features designed to support large scale data processing. STAR physicists, with support of NERSC staff, developed a processing framework capable of running STAR reconstruction on Cori. Based on this development, STAR applied for and was awarded a 2017 over-target allocation at NERSC from the NP program office. The goal was to use that allocation to process parts of the 2016 data and evaluate the use of NERSC HPC for complex, data-driven workflows, with large serial tasks typically run as High Throughput Computing (HTC) workloads. The allocation was successfully applied to reconstruction of the entire 2016 STAR raw data set for d+Au collisions at 200 GeV. This report presents the results and conclusions from that pioneering effort.

Physics Motivation

A major focus of the RHIC physics program is to quantify properties of the strongly-coupled quark gluon plasma (sQGP) formed in RHIC collisions. One promising avenue for studying those properties is the measurement of heavy quark distributions and correlations. The STAR collaboration constructed and operated a state-of-the-art, silicon MAPS pixel sensor detector, the Heavy Flavor Tracker (HFT), to make precise, differential measurements of heavy flavor production at RHIC. The HFT was installed and commissioned in 2014, with large minimum bias data sets from top energy Au+Au collisions collected in 2014 and 2016 and a reference data set from d+Au collisions taken in 2016. The parameters of these data sets are listed in Table 1.

Table 1. STAR datasets recorded during the 2014 and 2016 running periods

Data Sample	Minimum Bias collision events	Approx. raw data volume (PB)	Approx. output data (PB)	Est. Processing time (hours)
2014 Au+Au 200 GeV	1.0×10^9	1.5	0.75	25×10^6
2016 Au+Au 200 GeV	2.0×10^9	3.0	1.5	50×10^6
2016 d+Au 200 GeV	0.4×10^9	0.25	0.2	5×10^6
Total	3.4×10^9	4.75	2.45	80×10^6

First results on azimuthal anisotropy in D^0 production extracted from the 2014 data were presented at the Quark Matter 2017 conference, followed quickly by a publication in Physics Review Letters[1]. As illustrated in Figure 2 below, the measurement showed the first direct evidence that heavy, rare charm quarks participate strongly in the collective motion of quark-gluon-plasma in semi-central Au+Au collisions. These precision measurements enable detailed comparison to theoretical calculations, providing new measurements of the QGP and new insights into its structure and dynamics.

In order to fully exploit the Au+Au collision data, scientists need to compare such measurements with those obtained from d+Au collisions, in which a large-volume QGP is not expected to be generated. While the 2016 Au+Au collision data was being reconstructed, the reference d+Au data set became ready for processing and was selected as the best candidate for reconstruction on Cori. This data set was chosen because of its timely interest to the scientific community, and because the primary STAR computing resources were fully occupied processing other data. In addition, the dataset is large enough to fully test the capability of running the data-intensive HTC workflow at scale, yet modest enough to allow the full dataset to be processed within the over-target allocation.

---

[1] L. Adamczyk *et al.* (STAR Collaboration) Phys. Rev. Lett. 118**, 212301**

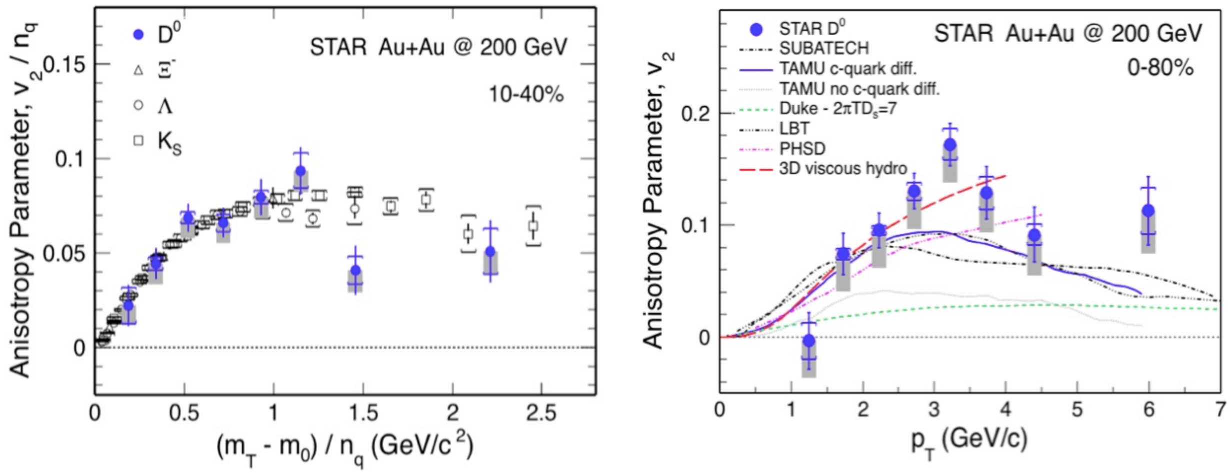

Figure 1. ***Left:*** *v_2 normalized by the number-of-constituent-quarks (n_q) as a function of transverse kinetic energy (normalized by n_q) for D^0 mesons and light hadrons in 10-40% central Au+Au collisions at 200 GeV.*
Right: *v_2 as a function of p_T for D^0 in 0-80% Au+Au collisions at 200 GeV compared to model calculations.*

Reconstruction Framework at NERSC

The reconstruction framework for running on NERSC HPC systems was first developed in 2016 by STAR physicists at LBNL and BNL, with assistance from NERSC and ESnet engineers. The primary features of this development have been reported elsewhere[2], along with initial performance results that illustrated highly efficient data processing on a small but sustained set of Cori CPU resources.

Several features of the system are critical for achieving reproducible and scalable processing of STAR datasets on Cori. One such feature is Shifter, the NERSC Linux container solution in which application software can be packaged with its full software stack; Shifter enabled the reconstruction to be run within its fully vetted Scientific Linux 6.4 environment, avoiding the need to port the extensive code base directly onto the OS of Cori, SuSE Linux Enterprise Server 12. In addition, the capability to efficiently instantiate a snapshot of the STAR MySQL calibrations database on each node removed potential bottlenecks due to network latency and database server contention, which would otherwise depend on the number of concurrent running jobs[3]. That capability was facilitated by Shifter, by which a per-node XFS filesystem can be instantiated at runtime to contain the STAR MySQL snapshot. The workload was segmented into data chunks, sized to be processed in several hours using all cores on a single node. This mapping of data to job-size allowed processing to scale out to the number of nodes

---

[2] "STAR Data Reconstruction at NERSC/Cori, an adaptable Docker container approach for HPC" Mustafa Mustafa, Jan Balewski, Jérôme Lauret, Jefferson Porter, Shane Canon, Lisa Gerhardt, Levente Hajdu, Mark Lukascsyk accepted for publication CHEP 2016.

[3] Initial tests showed that such bottlenecks became critical at job rates well below our target: 1000s of concurrent jobs.

available. The framework of service modules used for specific tasks (job submission, job validation, data QA and management) built around a central production database, was capable of dynamically scaling out the processing task without direct intervention, limited only by the rate at which data was transferred from BNL to NERSC and the amount of real-time CPU capacity available to the framework. The central production database, run on a local MongoDB server, was used as the source for a web-based monitoring system developed to check on the workflow performance and operations.

Operations

With the d+Au data set chosen, the system was configured for operation. A Shifter image with the STAR software version vetted for this production was deployed and an up-to-date copy of the calibration database was created. A number of data files were processed with the system and the results were subjected to QA tests, including comparisons with the same data processed at BNL. These tests were repeated several times during which the system remained stable and the results reproducible.

The network transfer bandwidth between BNL and NERSC, initially optimized in 2016, was re-evaluated and showed new asymmetries in bandwidth capacities between the sites. These were rectified with upgrades to STAR grid infrastructure at BNL and network tuning by project engineers[4]. The final data transfer services were implemented and deployed, and based on a 'pull' model, in which raw data was pulled from BNL to NERSC by services run at NERSC with the output data pulled back to BNL by BNL services. As such, each sited managed its own storage systems. The data transfers relied on GridFTP clients to move data between the GridFTP servers running at BNL and NERSC.

During the ramp-up of the 2017 processing, the network bandwidth and transfer framework capacity was capable of transferring several many 10s TB/day, easily meeting our processing goals[5]; however, infrastructure at BNL (staging disks, HPSS tape access) was sometimes oversubscribed with ongoing STAR experiment operations. To protect against such resource contention interrupting the transfer pipeline and producing a loss of available compute cycles, the project received a large dedicated disk allocation on the NERSC Cori scratch filesystem for the duration of the campaign. Sufficient raw data was kept on site such that even an interruption lasting days would not reduce the CPU utilization rate.

The processing framework ran with little intervention during the actual production. The number of jobs dynamically filled the amount of capacity available, and the service modules were simply stopped and restarted for routine machine maintenances. The only direct intervention during the campaign was to redirect jobs onto special node reservations when they became available, as noted in the next section.

[4] https://drupal.star.bnl.gov/STAR/star-newsletter/february-2017#com

[5] Neither the wide area network bandwidth capacity of ESnet nor the data transfer infrastructure at NERSC were limiting factors. In fact, only 1 of the 12 data transfer nodes (DTNs) available at NERSC were used during the project.

Processing Results

The production campaign officially began in late June. As noted earlier, the system is designed to scale out dynamically, limited only by availability of input data and CPU capacity. Since the large disk buffer at NERSC provided sufficient input raw data to the framework, the processing rate was set by the fair-share algorithm in the Slurm batch system on Cori. Production ran steadily for well over a month, maintaining 200 jobs queued for processing and sustained periods of running on 50 nodes (1,600 cores), the maximum allowed by the Slurm queue structure for jobs limited to 2 or fewer nodes. At that rate, no aspects of the framework were being stressed and the CPU utilization efficiency remained well above 98%. Therefore, the project requested and received a dedicated two-day reservation of 100 nodes (3200 cores). During that reservation period, the node availability was only about 83% due mainly to a rare Luster filesystem event; however, the CPU utilization efficiency of the available nodes by the STAR workflow system remained high, at 98.5%.

The success observed during the 100-node reservation indicated that the system should continue to perform well at larger scales and over longer sustained periods. To test this capability, a week-long 200-node (6,400 cores) reservation was requested. This represented nearly 10% of the Cori Phase I system and therefor generated some concern, prior to approval of the reservation, about whether the framework could scale to make full use of those resources over an entire week. The project's monitoring tools were able to show that the framework had been extremely fault tolerant and efficient throughout the 100-node reservation. The 200-node reservation was approved and, during the week-long reservation, the overall CPU utilization efficiency was again high at +99.3% for the 1.1M cpu-core-hours consumed: a testimony to the stability of the production pipeline and Cori system.

The project was interested in evaluating utilization of a burst capacity significantly larger than the dedicated facility at RACF. After the 200-node reservation, the amount of d+Au data that remained to be processed could be completed by a much larger reservation over a short period. The project obtained a two-day 800-node (25,600-cores) reservation, more than double the full capacity of the STAR reconstruction facility at RACF. The stability of the system during the reservation is illustrated in Figure 2. As a test of the burst capability, processing was started cold, with no jobs running. It took 12 minutes to reach full scale (inset in Figure 2), and continued uninterrupted for the two-day period. The overall utilization efficiency was +97.6% of the 1.23 M cpu-core-hours consumed. A total of 15.4k raw data files were processed on first attempt with 2 failed jobs, yielding a better than 99.9% success rate. A total of 106M d+Au collisions were processed during the 48 hours. In retrospect, the entire production could have been run in a little more than a week time with such a large reservation.

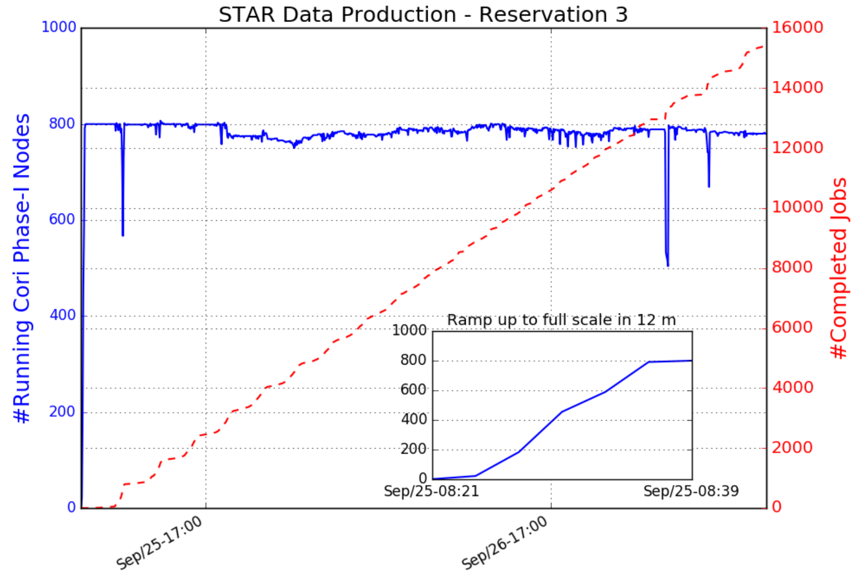

Figure 2. STAR processing during the 800-node, 25,600-core reservation. The blue curve shows the number of concurrent jobs (LH-axis, Cori nodes in use) over the 2-day reservation and during the 12-minute initial startup time (inset). The red curve shows the steady increase in number of completed jobs (RH-axis).

Conclusions and Outlook

STAR physicists, with support from engineers at LBNL and BNL, have built and deployed a novel data processing framework for running STAR raw data reconstruction on NERSC HPC systems. That system was optimized and used to reconstruct the entire 2016 STAR dataset for d+Au collisions at 200 GeV on NERSC Cori. The data, which were urgently needed by the STAR collaboration, are now being analyzed, and physics publications are being written based on them.

Within the context of filling STAR data production needs, the project successfully tested the use of NERSC HPC for complex data-driven workflows by carrying out a massive serial HTC workload, running on over 25,000 CPU cores at better than 97% end-to-end efficiency. In doing so, we evaluated and deployed a number of solutions to data processing challenges on HPC systems. Such solutions included a restructured data-to-job mapping for flexible scaling and support of high-concurrency access to calibration data. Furthermore, we found that the data transfer infrastructure had sufficient capacity to meet throughput requirements of several 10s of TB per day[6]. The framework remains in place and will continue to be used by STAR. It likewise serves as a prototype and reference to the NP experiment community for their future use of HPC systems.

[6] Use of disk buffers for staging input data insured high CPU utilization during interruptions in the data transfer pipeline.