

Data Reconstruction for the sPHENIX experiment

Ross Corliss^{a,*} *for the sPHENIX Collaboration*

^a*Center for Frontiers in Nuclear Science, Stony Brook University,
100 Nicolls Rd, Stony Brook, NY 11790, USA*

E-mail: ross.corliss@stonybrook.edu

The sPHENIX detector is a next-generation experiment under construction at the Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory. Starting in 2023 it will collect high statistics data sets from ultra relativistic Au+Au, p+p and p+Au collisions. The readout is a combination of triggered readout for calorimeters and streaming readout for the silicon pixel/strip detectors and the time projection chamber (TPC). sPHENIX does not employ higher level triggers – only a small subset of events is built online for monitoring purposes – which makes it unique among NP/HEP experiments. Events are assembled from multiple input streams as part of a multi-pass reconstruction which includes calibration and space charge distortion corrections for the TPC data. This reconstruction will run near realtime within a fixed latency of when the data were taken. To meet its physics requirements, sPHENIX has developed state-of-the-art reconstruction software based on the "A Common Tracking Software" (ACTS) package which was adapted to reconstruct the TPC data. The raw data will be processed at the Tier 0 for the RHIC experiments - the Scientific Data Computing Center (SDCC) at BNL. The Production and Distributed Analysis (PanDA) system was chosen as workload management system to handle the complexities of our workflow.

Here we describe the details of the data processing for the sPHENIX experiment.

*41st International Conference on High Energy physics - ICHEP2022
6-13 July, 2022
Bologna, Italy*

*Speaker




Figure 1: The sPHENIX Detector, with major detector components identified. Beginning from the collision point, particles pass through the tracking detectors (MVTX, INTT, and TPC), then electromagnetic, and finally, hadronic calorimeters. Between the inner and outer hadronic calorimeters is the coil and cryostat of the 1.4 T superconducting magnet. The TPOT detector is directly outside the TPC outer radius, covering $\sim 90^\circ$ in ϕ . Forward components, MBDs and EPD, are not shown.

1. Introduction

Over the past two decades, experiments at the Relativistic Heavy Ion Collider (RHIC) and the Large Hadron Collider (LHC) have revealed the existence of, and probed the nature of, Quark-Gluon Plasma (QGP). Originally proposed in 2010, the sPHENIX experiment is a substantial upgrade to PHENIX that will substantially extend the precision with which the properties of the QGP can be measured at RHIC. In addition to the performance of the new detector components, this extended reach depends on the sPHENIX data acquisition system itself, which is designed to take data at 15 kHz in order to record extremely large samples of minimum-bias data. Detector construction is nearly complete; the experiment will begin data-taking in early 2023.

sPHENIX's physics program [1] has four main areas. Jet structure measurements (e.g. jet-photon momentum balance) probe varied momentum and angular scales, Upsilon spectroscopy allows variation of the size of the probe, and open heavy flavor measurements probe on different momentum and parton mass scales. In parallel, cold QCD measurements will study cold nuclear effects including spin and transverse momentum. The kinematics of these measurements overlap with those at the LHC, providing the ability to measure the same observables in QGPs with different initial conditions.

2. The sPHENIX Detector

sPHENIX (Fig. 1) is an azimuthally symmetric suite of tracking detectors and calorimeters based around and within the 1.4 T superconducting solenoid used by the BaBar experiment. [2] The innermost detector is the MAPS¹ Vertex Detector (MVTX), made up of three cylindrical layers of silicon pixel staves between 2.3 and 3.9 cm from the collision point. Outside this is the Intermediate Tracker (INTT) from 6.0 to 12.0 cm, consisting of two layers of silicon strip detector staves. At

¹Monolithic Active Pixel Sensor

larger radius, the Time Project Chamber (TPC) covering radius from 20 to 78 cm. The TPOT is a MicroMegas-based tracking layer primarily intended to assist in TPC distortion corrections and is the only ϕ -asymmetric component.

Outside the tracking detectors there is a tungsten and scintillating fiber sampling Electromagnetic Calorimeter (EMCal) followed by an Inner Hadronic Calorimeter (iHCal) composed of aluminum and scintillating tiles. On the outside of the magnet cryostat, the Outer Hadronic Calorimeter uses steel and scintillator, and serves as the solenoid's flux return.

In the forward regions, the Minimum Bias Detector (MBD) reuses PHENIX's Beam-Beam Counters, 3 cm quartz radiators with mesh dynode photomultipliers, and the Event Plane Detector (EPD) consists of two wheels of scintillating tiles.

3. Data Acquisition

RHIC is a mature, well-understood machine, and delivered luminosity is expected to be high on day one. Depending on the species, peak event rates are expected to be as high as 50 kHz. The beams collide with a small crossing angle to increase the fraction of collisions that occur within the acceptance of the MVTX and INTT.

The sPHENIX Data Acquisition system (DAQ) is designed to record data at a maximum rate of 15 kHz independent of event size. Maximizing the overall physics output requires considering the time-averaged performance of the full computing chain, from collision through to final processing of data: Each year, production of physics data from raw event data must complete before the next run period begins. The ultimate data volume in each year depends on experiment and RHIC uptime, and are conservatively estimated to be ~ 70 , ~ 80 , and ~ 180 PB in the first (13 weeks Au+Au), second (21 weeks p+p,p+A), and third (24.5 weeks Au+Au) year of sPHENIX running, respectively.

The sPHENIX DAQ is a hybrid system, in which all calorimeters, as well as the MBD and EPD, are triggered, while all tracking detectors are capable of free streaming. In both cases, data from every beam crossing is digitized and buffered on the Front End Modules/Electronics (FEM, FEE). In the calorimeters, the FEMs read out on the arrival of a trigger signal, while the FEEs for the tracking detectors can stream automatically.

To avoid sampling bias and access untriggerable observables, the main trigger for sPHENIX is provided by the MBD, imposing only a vertex position cut via tube timing. Other calorimeter-based jet patch or high- p_T triggers are also being studied.

3.1 Triggered Readout

The readout chain for the triggered detectors inherits from the PHENIX readout. FEMs buffer data from multiple bunch crossings, until triggered externally, at which point data is sent to Data Collection Modules (DCM2s) [3] re-used from PHENIX, which in turn transmit to sub-event buffers (SEBs). Each SEB handles only the subset of the total event, corresponding to the DCM2s that feed into it, hence the full event not available on any node.

3.2 Streaming Readout

Data flow for the tracking detectors is conceptually similar: they push data out in streaming mode to ATLAS-designed Front-End Link eXchange (FELIX) [4] cards, which are interfaced




Figure 2: Schematic drawing of data flow for triggered (top) and streaming (bottom) detectors. A combined total of 61 SEBs and EBDCs provide packets through the network switch to buffers boxes, which have enough disk to collect several stores' worth of data. The buffer boxes have twin disks, one filling from the detector while the other transmits to the tape and processing facilities.

directly to Event Buffer and Data Compressors (EBDCs), off-the-shelf computers which serve the equivalent role of SEBs, after which handling is identical to the triggered systems.

Because of the $\sim 13 \mu\text{s}$ drift time it takes for an event to be fully read out from the TPC, tracks from earlier and later events are present in the drift volume as well as the in-time event. In particular, with the bunch crossings occurring every 110ns, tracks from the following 118 additional crossings will be partially recorded in the original time window. sPHENIX can opportunistically extend the streaming window in order to fully capture these: Each additional 110 ns completes the data from an additional bunch crossing. By tuning the size of the extended window to match an acceptable data rate, substantially more minbias events can be captured. At a benchmark of $\sim 60 \text{ Gbps}$, this hybrid mode would allow the tracking data from $\sim 10\%$ of all min bias collisions (100-1000 times more than triggered) to be recorded in $p + p$ and $p + \text{Au}$ running, which is crucial to the open heavy-flavor and cold QCD programs.

3.3 Writing Data

Both triggered and streaming data packets feed into buffer boxes, which serve two roles: They level the incoming data rates and allow us to send the average, rather than the peak rates to HPSS. They also offer about 5 days worth of data storage in case of a problem with the HPSS system, allowing the experiment to continue to take data in that situation. During data taking, the various SEBs and EBDCs collectively write about 60 individual data files, containing data from that particular portion of the detector, that get load-balanced across the buffer boxes.

Data is written to one of two file systems on each buffer box while the other transfers data to tape storage. When the first one hits a high-water mark, the second (which has completed its transfer) is cleared, and their roles are switched. This ensures that no disk is competing between read and write, minimizing disk head movement and dramatically increasing throughput.

4. Event Reconstruction

In order to analyze collisions, the data from each detector corresponding to a particular event must be assembled. In PHENIX, this was done as part of the data acquisition with an Event

Builder, before writing to tape. Event-building in sPHENIX is offline instead, integrated with event reconstruction — reassembled raw events are not written to tape, but assembled only as-needed during production. Shifting the event-building step offline reduces the complexity of the online system, which no longer needs to synchronize and assemble numerous packets in realtime. It also greatly enhances the reliability of the system overall: If errors occur the event building can simply be re-done, while with an online event builder the data were lost. An online system must also handle the peak data rate, while the offline system is buffered by disk, so that the event building only has to handle the average rate. This substantially reduces the computing resources needed to avoid data loss.

Due to the asynchronous nature of an online Event Builder, the events in a given data file were not strictly time-ordered in PHENIX, and consecutive events generally ended up in different data files. Conversely, the events in individual files in sPHENIX are strictly consecutive. This enables easy access to the history of a given detector, such as the deposited charge in the TPC that is used for various corrections. Calibrations will likewise apply to contiguous blocks as the file is read. In addition, since the location of detector information is deterministic, reconstruction steps only need to load the data they will use. The price of this is that reprocessing data from tape is significantly more involved, since many files have to be staged together in order to reconstruct any event.

4.1 Calorimeter Calibration

Calorimeters are calibrated through a combination of LED monitoring and beam-data driven methods. Absolute calibrations in the EMCal are performed with an iterative fitting of π^0 and η mass spectra, and in the HCals through monte-carlo based jet and high-pT particle spectra. Drift of relative gains is monitored through tower energy spectra and LEDs systems. Dedicated LED and pedestal scans are performed when beam is not present.

4.2 Tracking and TPC Corrections

Track reconstruction uses the A Common Tracking Software (ACTS) framework, adapted for the sPHENIX detectors [5], in particular the TPC. The most challenging calibration effect in tracking is the handling of distortions in that detector. These distortions factor into three components: Static distortions, caused by $\vec{E} \times \vec{B}$ effects, are not expected to vary during running, and will be mapped by a directed laser system in advance. Average distortions are caused by the average charge in the chamber, primarily due to slow variation in beam conditions, and are expected to be stable over ~ 20 minutes. Fluctuations are the event-by-event variation of space charge due to the ion-backflow from particular events, monitored by flashing a UV laser on a test pattern on the TPC's central membrane.

Track reconstruction takes place in two passes in order to accommodate the correction schemes for these distortions: While static distortions can be corrected in advance, and fluctuations can be handled in a time series, in order to derive the corrections for the average distortions, a large sample of tracks is needed in each voxel of the TPC. These statistics require tens of minutes of data, resulting in a common correction map that must then be applied to every event from that set.

4.3 Production

Production itself is coordinated with the sPHENIX Handy Remote Execution Koordinator (SHREK), which simplifies the task of defining jobs, staging data sets, and documenting production campaigns with a single source description for the complete production. Through SHREK, jobs are mapped onto the PanDA workflow management system, which submits jobs on local and remote compute resources, while Rucio handles the cataloguing of input data, intermediate job products and final outputs. Version control of production jobs, including all scripts and macros, is handled by SHREK through an interface with Github.

4.4 Online

Where possible calibration and monitoring are done without collating full events. Calorimeter calibrations, for instance, can be probed with local sub-events, readily available from separate data files without synchronization. In some cases, however, data across multiple files or detectors is necessary, for data integrity checks or track reconstruction. It is expected that a small fraction of the data will be assembled to enable online monitoring of TPC performance, and other global issues.

5. Outlook

The sPHENIX experiment will be the first new general-purpose detector at RHIC since STAR and PHENIX began taking data in 2000, and will continue to probe the nature of the Quark Gluon Plasma, as well as other topics in hot- and cold-QCD through an anticipated three-year program. In addition to standard jet and high-pT triggers, sPHENIX's tracking detectors have streaming readout, with the potential to record large, truly unbiased datasets and provide access to otherwise untriggerable events. In order to maximize the physics impact between commissioning and the beginning of construction of the future EIC, sPHENIX's data production must run with an on-average fixed latency with respect to data taking, hence efficient use of computing resources is essential. The construction of sPHENIX is nearing completion, with first data scheduled for early 2023.

References

- [1] A. Adare et al. An Upgrade Proposal from the PHENIX Collaboration. 2015. arxiv:1501.06197.
- [2] sPHENIX Collaboration. sPHENIX Technical Design Report, 2019. <https://indico.bnl.gov/event/7081>.
- [3] S.S Adler et al. PHENIX on-line systems. *NIM A*, 499(2):560–592, 2003. The Relativistic Heavy Ion Collider Project: RHIC and its Detectors.
- [4] J. Anderson et al. FELIX: a PCIe based high-throughput approach for interfacing front-end and trigger electronics in the ATLAS Upgrade framework. *Journal of Instrumentation*, 11(12):C12023, dec 2016.
- [5] Joseph D. Osborn et al. Implementation of ACTS into sPHENIX track reconstruction. *Computing and Software for Big Science*, 5(1), oct 2021.