The Readout of the sPHENIX Tracking System

Martin L. Purschke, for the sPHENIX Collaboration

***Abstract*—The sPHENIX Collaboration at RHIC is upgrading the PHENIX detector in a way that will enable a comprehensive measurement of jets in relativistic heavy ion collisions. The upgrade will give the experiment full azimuthal coverage within a pseudorapidity range of $-1.1 < \eta < 1.1$.**

We have presented the status of the sPHENIX project at the last conference. Since then, we have selected the technologies for the tracking system, which will consist of a silicon detector based on MAPS (Monolithic Active Pixel Sensors), followed by an Intermediate Tracker (INTT), and then by a TPC. We will present the envisioned design of the tracker readout, and explain the challenges with the high data rates generated by this readout method, which could reach as much as 80 GBit/s. We will show the design of our data acquisition that can cope with those data rates, and present the status of the ongoing R&D.

The sPHENIX Apparatus

sPHENIX [1] is a proposal for a new experiment at RHIC capable of measuring jets, jet correlations and upsilons to determine the temperature dependence of transport coefficients of the quark-gluon plasma. A reference design of the new apparatus is shown in Fig. 1.

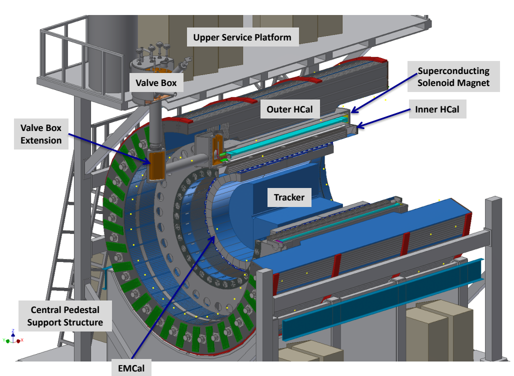

Fig. 1. View of the sPHENIX detector with its component subdetectors.

Figure 2 summarizes the expected increase in statistical reach of hard probes measurements in A+A collisions, showing the most up to date R_{AA} measurements of hard probes in central Au+Au events by the PHENIX Collaboration plotted against statistical projections for sPHENIX channels measured after the first two years of data-taking. While the existing measurements have greatly expanded our knowledge of the QGP created at RHIC, the overall kinematic reach is constrained to less than 20 GeV even for the highest statistics measurements.

sPHENIX will greatly expand the previous kinematic range studied at RHIC energies (in the case of inclusive jets, the data could extend to 80 GeV/c, four times the range of the current

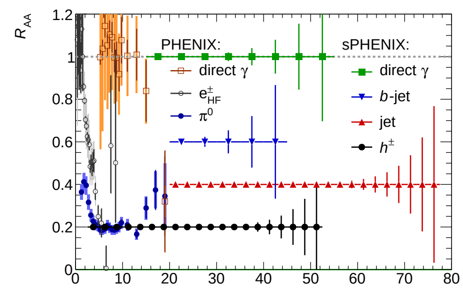

Fig. 2. Statistical projections for the R_{AA} of various hard probes vs p_{T} in 0–20% Au+Au events with the sPHENIX detector after two years of data-taking, compared with a selection of current hard probes data from PHENIX.

PHENIX π^0 measurements) and will allow access to new measurements entirely (such as fully reconstructed b-tagged jets).

The upgraded detector will be based on the former BaBar magnet and will include a high resolution and low mass tracking system for reconstruction of the Upsilon states, a new electromagnetic calorimeter, and, for the first time at a RHIC experiment, a hadronic calorimeter.

A. The sPHENIX Tracking System

The technologies for the tracking system have recently been selected. The tracking will consist of a silicon detector based on MAPS (Monolithic Active Pixel Sensors), followed by an Intermediate Tracker (INTT), and then by a TPC (Fig. 3).

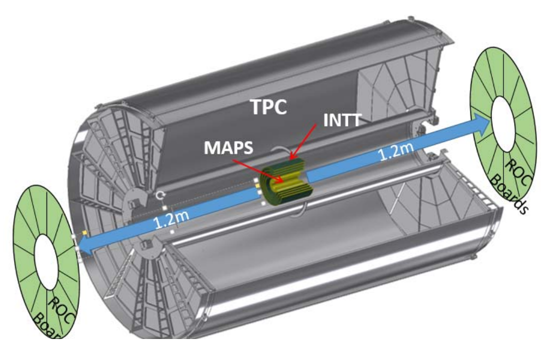

Fig. 3. The sPHENIX tracking system consisting of the MAPS detector, the Intermediate Tracker (INTT), and the TPC.

With the exception of the INTT, which is based on the former PHENIX Forward Vertex Detector and also re-uses the readout technology, both the MAPS detector and the TPC introduce new readout hardware and strategies.

In order to achieve high event rates, the TPC needs to be read out in continuous, or streaming, mode, without the use of a gating grid. The resulting continuously sampled waveforms of the TPC sensors must then be processed and correlated with the actually triggered events of the full detector (fig. 4).

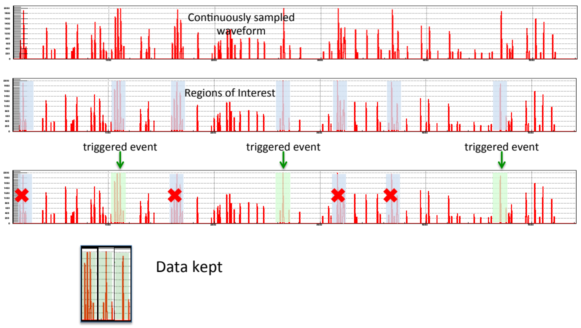

Fig. 4. A conceptual overview of the TPC "streaming" readout. The front-end electronics continuously samples the waveforms. A processing system selects "regions of interest", indicated in this example as amplitudes above a threshold. Further processing selects those regions that correlate with triggered events.

Because the TPC is using the ALICE SAMPA ASIC for its readout, one of the possible readout cards is the ALICE "Common Readout Unit" (CRU). We are also looking into a card in use in the ATLAS experiment, called "FELIX". Either card, which we call *Data Aggregation Module* (DAM), is PCIe-based and has a large number (up to 48) of fiber inputs to connect to the front-end cards. The boards have powerful FPGAs for waveform processing. The processed data from the DAMs are read out by the *Event Buffering and Data Compressor* machines (EBDC). Those systems hold the data from the respective subset of connected readout channels, which still need to be combined into full events that contains all the data from one collision.

This task is performed by the Event Builder, which consists entirely of commodity PCs running Linux. The SEBs that hold the *sub-events* from the calorimeter and other non-streaming detectors, and the EBDCs are connected to a high-end network switch, which is central to the Event Builder. It must be able to sustain the aggregated bandwidth in a non-blocking fashion. Also connected to the switch are a number of *Assembly and Trigger Processors* (ATPs). Through the network switch each ATP receives the sub-events from a given collision and combines them into a fully assembled event.

While the number of SEBs and EBDCs is determined by the topology of the front-end and fixed for a given configuration, the number of ATPs is not fixed and can be adjusted to match the load. We use the ATPs to perform a late-stage, distributed compression of the data before they are sent to the so-called *Buffer Boxes* that receive the data from the ATPs and provide local storage capacity. The compression has traditionally yielded savings of about 45% – 100GB of data shrink, on average, to 55GB.

The Buffer Boxes are designed to limit the number of concurrently written output files to a reasonably small value by receiving the compressed buffers from the ATPs and writing them to disk. They provide about 80 hours of local storage capacity, which will help us to ride out short-term outages of the tape storage system without the need to stop taking data. In addition, the local buffering levels the changing data rates of the experiment and allows us to transfer the *average*, rather than the peak, data rate of the experiment to the long-term HPSS tape storage system located at the RHIC computing facility.

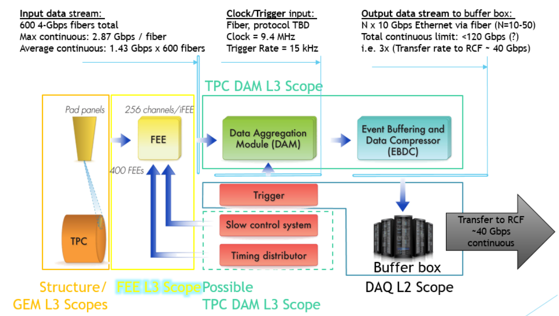

Fig. 5. An overview of the TPC electronics chain. FEE cards housing SAMPA chips are located on board of the detector. Zero suppressed, untriggered data flows to Data Aggregation Modules (DAMs). From there, the TPC data joins the main stream flow of the sPHENIX DAQ.

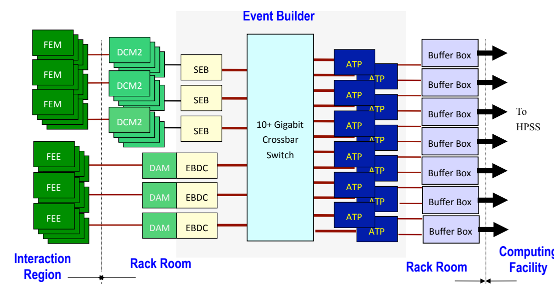

Fig. 6. Overview of the event builder design. The data are digitized in the Front-End Modules and zero-suppressed and packaged in the Data Collection Modules. The data from a given collision are initially distributed over many SEBs and EBDCs. The data from one collision are collected in the ATP's, which sees the full complement of data of that collision for the first time. The ATP compresses the data before transmitting them to the *Buffer Boxes*, from where the data are transferred to a long-term storage system.

Summary

Focusing on the tracking system, we will describe the goals and design of the sPHENIX experiment. We will present the envisioned design of the streaming readout, and explain the challenges with the high data rates generated by this readout method, which could reach as much as 80 GBit/s. We will show the design of our data acquisition to cope with those data rates, and present the status of the ongoing R&D. By the time of the conference, we will likely have a number of prototype setups, and some actual performance data.

REFERENCES

[1] A. Adare et al. An Upgrade Proposal from the PHENIX Collaboration. 2014.