

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS

GEREMEAS, SARDINIA, ITALY

1–6 OCTOBER 2023

A demonstrator for a real-time AI-FPGA-based triggering system for sPHENIX at RHIC

J. Kvapil^{a,*}, G. Borca-Tasciuc^b, H. Bossi^c, K. Chen^d, Y. Chen^d, Y. Corrales Morales^c, H. Da Costa^a, C. Da Silva^a, C. Dean^c, J. Durham^a, S. Fu^e, C. Hao^f, P. Harris^c, O. Hen^c, H. Jheng^c, Y. Lee^c, P. Li^f, X. Li^a, Y. Lin^a, M. X. Liu^a, A. Olvera^e, M. L. Purschke^g, M. Rigatti^h, G. Roland^c, J. Schambachⁱ, Z. Shi^a, N. Tran^h, N. Wuerfel^j, B. Xu^f, D. Yu^k and H. Zhang^f

^aLos Alamos National Laboratory, Bikini Atoll Rd, Los Alamos, NM 87545, U.S.A.

^bRensselaer Polytechnic Institute, 110 8th St, Troy, NY 12180, U.S.A.

^cMassachusetts Institute of Technology, 77 Massachusetts Ave, Cambridge, MA 02139, U.S.A.

^dCentral China Normal University, No. 152, Luoyu Rd, Wuhan 430079, China

^eUniversity of North Texas, 1155 Union Cir, Denton, TX 76205, U.S.A.

^fGeorgia Institute of Technology, 225 North Ave, Atlanta, GA 30332, U.S.A.

^gBrookhaven National Laboratory, PO Box 5000, Upton, NY 11973, U.S.A.

^hFermilab, PO Box 500, Batavia, IL 60510, U.S.A.

ⁱOak Ridge National Laboratory, PO Box 2008, Oak Ridge, TN 37831, U.S.A.

^jUniversity of Michigan, 500 S State St, Ann Arbor, MI 481091, U.S.A.

^kNew Jersey Institute of Technology, 323 Dr Martin Luther King Jr Blvd, Newark, NJ 07102, U.S.A.

E-mail: jakub.kvapil@lanl.gov

ABSTRACT: The RHIC interaction rate at sPHENIX will reach around 3 MHz in pp collisions and requires the detector readout to reject events by a factor of over 200 to fit the DAQ bandwidth of 15 kHz. Some critical measurements, such as heavy flavor production in pp collisions, often require the analysis of particles produced at low momentum. This prohibits adopting the traditional approach, where data rates are reduced through triggering on rare high momentum probes. We explore a new approach based on real-time AI technology, adopt an FPGA-based implementation using a custom designed FELIX-712 board with the Xilinx Kintex Ultrascale FPGA, and deploy the system in the detector readout electronics loop for real-time trigger decision.

KEYWORDS: Detector control systems (detector and experiment monitoring and slow-control systems, architecture, hardware, algorithms, databases); Trigger algorithms; Trigger concepts and systems (hardware and software)

ARXIV EPRINT: [2106.08353](https://arxiv.org/abs/2106.08353)

*Corresponding author.

Contents

1	Motivation	1
2	sPHENIX detector	1
2.1	The DAQ-AI data flow	2
2.2	Latency breakdown	3
3	Model description	4
4	Generation of the GNN IP core	4
5	Summary and outlook	5

1 Motivation

Realizing the science potential of modern nuclear physics (NP) experiments at colliders relies on the collection and processing of very large data that often exceeds the current DAQ bandwidth and budgetary limitation. We propose to develop real-time AI technologies, implemented in the detector readout electronics loop, that address these challenges for the next generation of NP experiments at RHIC and Electron-Ion Collider (EIC). First, we will deploy a demonstrator that is being developed under the current Fast-ML project [1] for the pp running in the sPHENIX experiment in 2024, and then generalize our approach for applications in experiments at the EIC, using future generations of EIC detector technologies.

The triggered readout rate of sPHENIX is limited to 15 kHz due to the design of the calorimeter readout system, which places limitations on the overall data volume and the expected collision rate. In pp collisions, RHIC delivers collision rates of 3 MHz, limiting sPHENIX to collect less than 1% of heavy-flavour (HF) events of the total pp (and p+Au) rate when using triggered readout. The extended Streaming readout (SRO) of the tracking detectors can further improve the statistics up to 10% of the total luminosity. The goal of this project is to sample the remaining luminosity further enhancing the collected data samples. The aim is to deploy a future system on the EIC to identify the (non)interesting Deep-Inelastic-Scattering processed in the e+p/A collisions.

2 sPHENIX detector

The sPHENIX detector [2] is located at the RHIC accelerator complex in Brookhaven National Laboratory (BNL), with data taking in 2023–2025. The sPHENIX detector schematic and MVTX and INTT detectors are shown in figure 1. The tracking detectors (MVTX, INTT, TPC, TPOT) are capable of SRO. Even-though they are able to record all data, the data volume of the TPC detector exceeds the capability of the computing center and therefore a down-selection of what to save must be done. The aim of this project is to reconstruct tracklets from the silicon detectors in order to search for signatures of HF decays based on unique topology, and provide an additional trigger to sPHENIX.

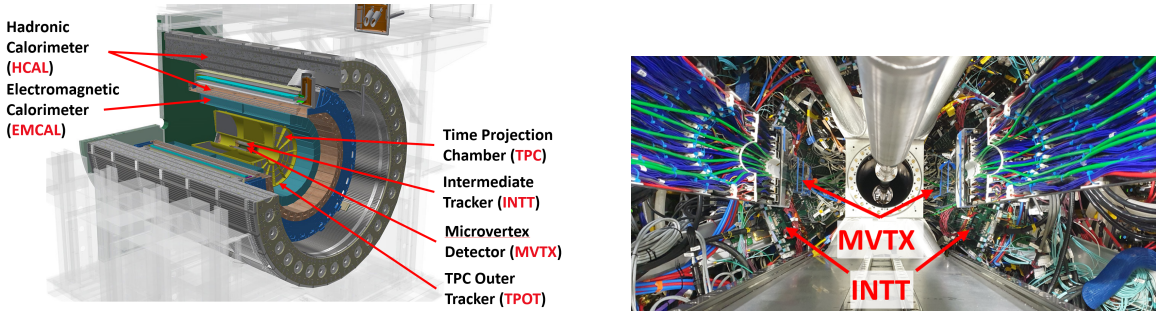


Figure 1. Left: the sPHENIX detector with highlighted central barrel detectors. Right: installation of MVTX and INTT detectors.

Silicon detectors. The three innermost tracking layers (the MVTX detector) are based on Monolithic Active Pixel Sensors (MAPS), the ALPIDE, originally developed for the ALICE experiment. The ALPIDE offers very fine pitch of $27 \mu\text{m} \times 29 \mu\text{m}$ with a collision event time resolution of $5 \mu\text{s}$. The MVTX contain a total of 270M channels. The next two layers (the INTT detector) contain silicon strip sensors (manufactured by Hamamatsu) with a pitch of $78 \mu\text{m} \times 16$ (or 20) mm with 360k channels in total. Both detectors are shown in figure 1 (right).

sPHENIX readout, trigger and timing distribution. The schematic of the sPHENIX readout chain is shown in figure 2. The tracking detectors' Front-End Electronics (FEE) sends data to the Event Buffer and Data Compressor (EBDC) through the FELIX (FELIX-712) [3] interface card. The calorimetry detectors' Front-End Modules (FEM) send data to the SubEvent Buffer (SEB) through the Data Collection Module (DCM2). The trigger and timing information is generated by the Global-Level-1 trigger system (GL1). This information, in addition to detector-specific timing parameters and other book-keeping data, is adapted for each detector by Granule-Timing Modules (GTMs) that also transmit a copy of the accelerator clock to the front-ends, and the trigger information. The GL1 can currently generate up to 64 different triggers derived from 4 fiber inputs that carry the collision information from the front-ends, in addition to 4 hardware signal inputs that can be used for various purposes. The most used trigger inputs are from the Hadronic Calorimeters for the cosmic and high energy jet triggers (normally not pre-scaled) and Minimum-Bias Detector for the beam collision trigger (heavily pre-scaled). The goal is to provide an additional high efficiency trigger input for HF events in pp collisions.

2.1 The DAQ-AI data flow

Data from the MVTX and INTT are transmitted from the FELIX-DAQ board to the AI-Engine which is also based on FELIX board (FELIX-AI). It was decided that the AI-Engine will be hosted on a widely HEP community supported FELIX board for the following 3 reasons. First, FELIX offers 48 high-speed optical links to receive data; second, to reuse its Wupper module for the PCIe communication and associated software tools; and third, to use the sPHENIX infrastructure for tracking detectors that was built around the FELIX cards. The FELIX-AI houses the raw data decoder, event builder, clusterizer, and GNN models to provide fast tracking and the trigger decision. The trigger decision is sent via a LEMO cable to the GTM. Since the decision is based on an event topology, a reference point (beam spot), which changes in time, must be precisely known and monitored. A GPU based feed-back system will be in place to process the data from the buffer boxes, reconstruct the beam spot position and update the position in the FELIX-AI. The schematic of the data flow is shown in figure 3.

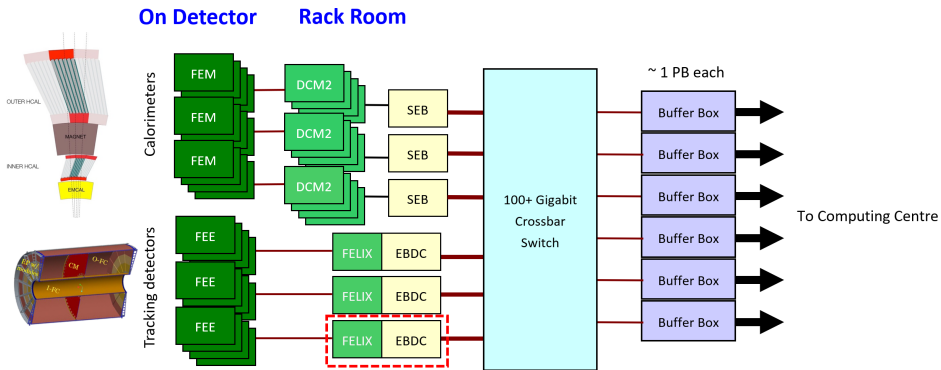


Figure 2. A schematic of sPHENIX readout chain. The calorimetry data are sent from Front-End Modules (FEM) to the SubEvent Buffer (SEB) through the Data Collection Module (DCM2), and the tracking data from Front-End Electronics (FEE) are sent to the Event Buffer and Data Compressor (EBDC) through a FELIX interface card.

There are 144 optical links running at 3.2 Gbps per fiber for the MVTX alone, thus two engines will be used, one for each MVTX/INTT hemisphere. This will allow to have 24 links for MVTX and 24 links for INTT. The FELIX optical links have been tested up to 14 Gbps with BER $< 10^{-16}$ with an external loop-back measurement. INTT offers excellent time resolution to tag 100 ns RHIC bunch-crossing time to assign a unique timing to each collision event.

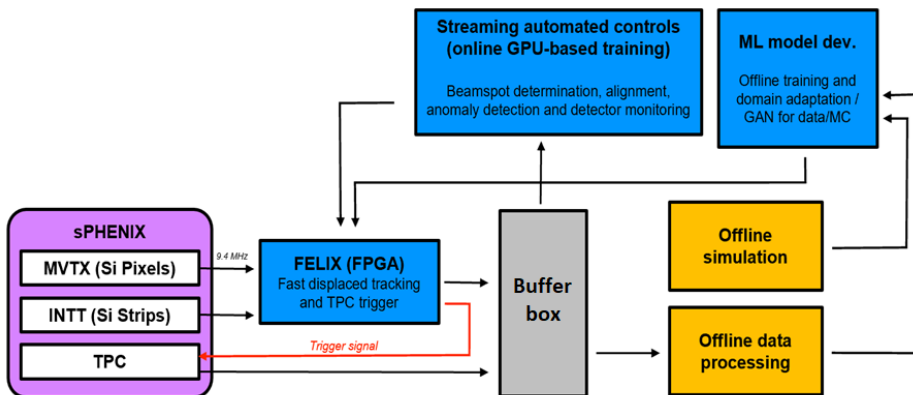


Figure 3. The schematic of the DAQ-AI data flow.

2.2 Latency breakdown

The TPC buffers can hold up to about 30 μs of data and that is the maximum hard limitation. Our stretched goal is to deliver the trigger decision to the global system with in 10 μs which could allow us to capture data from other faster detectors. The MVTX contributes up to 5 μs latency, the cables between the interaction region and counting house contribute around 0.3 μs (81 m fibres). Forwarding data to the FELIX-AI and decoding them takes up to 0.6 μs , which results from the maximum depth of FIFO holding the decoded hits (128). This leaves around 4 μs for the AI engine to perform tracking and trigger decision. The real latency of the decoding and data transmission depends on the occupancy, which is estimated to be around 50 physics hits per chip per event. The FELIX-AI needs to ensure the latency is fixed, either by delaying the trigger decision, or vetoing the processing by decoding the bunch-crossing numbers from MVTX and INTT.

3 Model description

The model is based on Graph Neural Network (GNN) using PyTorch and PyTorch geometric. The aim is to reconstruct the decay topology of the tracklets and search for secondary (displaced) vertices, which is one of the prominent features of HF decays. The displaced vertex will be $\sim O(100 \mu\text{m})$ away from the primary vertex which has resolution $\sim O(10 \mu\text{m})$. We have developed a novel method for treating events as graphs consisting of tracks as nodes and interconnection between tracks when they belong to the same particle decay rather than hit graphs [4]. A 15% improvement is observed by estimating the transverse momentum, p_T , based on the tracks. There are three stages in event processing: hit clustering, track reconstruction, and trigger decision. We use the GNN models to reconstruct tracks, remove outliers (TrackGNN), regress the track momentum and displaced vertices on the corresponding tracks, and to estimate the probability that the event is a trigger. The flow chart is shown in figure 4. Table 1 summarizes the efficiency and background rejection rates at two levels of signal/noise ratios, that is, 1% and 0.1%. The signal/noise ratio is computed using heavy flavor simulations where the D^0 decay products fall within the sPHENIX detector volume even if the tracks do not have sufficient momentum to reach all the tracking detectors. The test ratios can be compared to signal/noise ratio at sPHENIX under these circumstances which is expected to be approximately 0.5 %. Note that the efficiency decreases with the signal/noise ratios. The apparent difference in efficiency is due to statistics of samples, however, the purity is decreased by a factor of 10 as expected. This is already comparable with sPHENIX partial streaming of 10% MB events with $6.2 \pm 0.9 \%$ purity for $D^0 p_T \geq 2 \text{ GeV}$ [5].

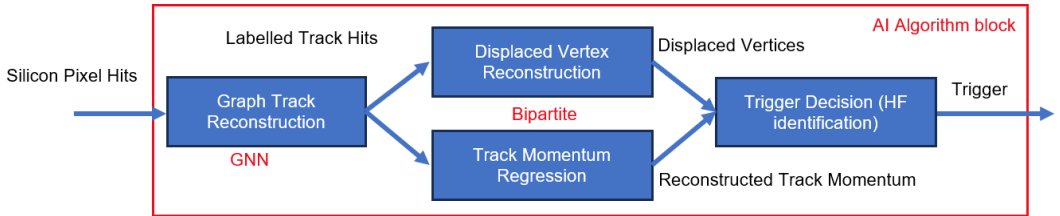


Figure 4. Flowchart of the GNN network.

Table 1. Efficiency and Background Rejection with 1% and 0.1% signal/noise Ratios.

1% signal/background ratio			0.1% signal/background ratio		
Bg. rejection	Efficiency	Purity	Bg. rejection	Efficiency	Purity
90%	75.6%	7.1%	90%	76%	0.75%
99%	19.4%	16.4%	99%	23.2%	2.3%

4 Generation of the GNN IP core

We propose to use High-Level Synthesis (HLS) tools to generate synthesizable Register Transfer Level (RTL) code, e.g., Verilog, to implement the software sPHENIX model onto FPGA. The steps are as follows. First, we manually translate the sPHENIX model into synthesizable C code and feed it into the HLS tool, Vitis HLS [6]. Then, we perform hardware optimizations of the model in HLS following the FlowGNN architecture [7], which is the state-of-the-art GNN architecture on FPGA.

The target is to have a model implemented on FPGA that can process 100–200 nodes (hits) and 200–500 edges (hit connections) within $\sim O(10 \mu\text{s})$.

FlowGNN architecture with on-FPGA implementation. The current TrackGNN model in sPHENIX we are using has one GNN layer, which includes 4 multi-layer perceptron (MLP) layers for both node and edge embedding with a dimension of 8. The proposed architecture follows the message-passing framework in FlowGNN [7]: the node embeddings are processed first, followed by an adapter to orchestrate the node information to the correct edge processing units for edge embedding computation and message aggregation. We also use quantization to reduce the data precision and to reduce the memory and computation requirements. We use `ap_fixed<18, 6>` for node embeddings, edge embeddings, and model weights and bias, and use `ap_fixed<21, 9>` for messages and input node features. The initial study FPGA board is the Alveo U280, approximately twice as big as FELIX-712 used for the final deployment. The resource utilization is as follows: 194 K (14.9%) LUT, 214 K (8.2%) FF, 406 (20.2%) BRAM, and 488 (5.4%) DSP. The processing latency is measured on-board in an end-to-end fashion, including graph and weight loading, model computation, and results readback. The latency for an average-sized input (92 nodes and 142 edges) is $8.82 \mu\text{s}$ at a 285 MHz clock. The time-size scatter plot for nodes and edges are shown in figure 5, where the x-axis is the time spent processing one graph, and the y-axis is the graph size in terms of the number of nodes and edges. Compared to CPU calculations, the FPGA was 99.86% accurate.

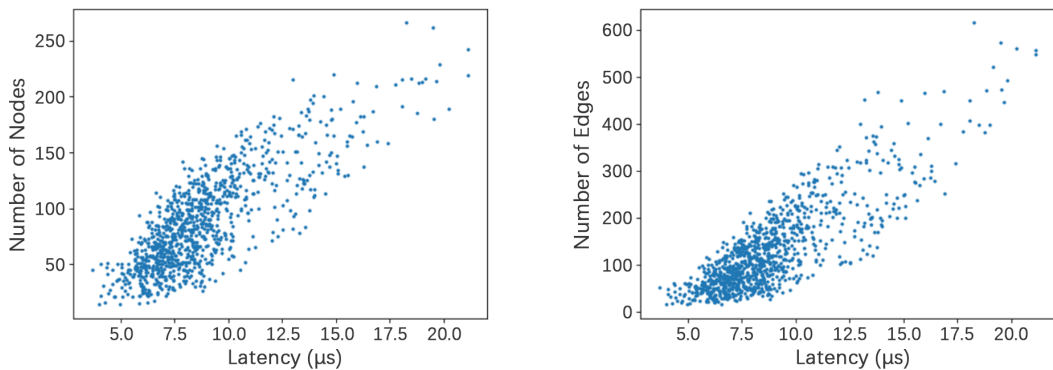


Figure 5. The time-size scatter plot of the TrackGNN model measured at Alveo U280. Left: latency values with numbers of nodes. Right: latency values with numbers of edges.

5 Summary and outlook

This paper describes how the streaming data of selected tracking detectors can be used to form an additional trigger input. Several modules of the FELIX-AI required for the trigger decision have been tested independently on different target devices. The final development is focused towards combining them into a single design and deploy on FELIX-712 boards. A second iteration of the TrackGNN model has achieved a latency of $8.82 \mu\text{s}$ at a 285 MHz clock, and this is sufficient for the target maximum latency limitation of $30 \mu\text{s}$. However it is approximately double from our stretched goal of around $4 \mu\text{s}$ in order to deliver the trigger decision within $10 \mu\text{s}$. A further improvement is expected by providing more lightened version of the model and by targeted optimisation of the design. Furthermore, a parallel effort to translate the model using the hsl4ml framework [8], which is a generalized package to translate neural networks into an IP core just started.

Acknowledgments

This project is funded by the United States Department of Energy, funding calls FOA-0002490, FOA-0002875 and Los Alamos National Laboratory LDRD program.

References

- [1] NP 2022 Accelerator R&D | U.S. DOE Office of Science(SC). <https://science.osti.gov/np/Research/DOE-NP-Accelerator-R-D-PI-Meeting/DOE-NP-Accelerator-R-D-PI-Meeting-Presentations-2022> [accessed 19-12-2023].
- [2] sPHENIX collaboration *sPHENIX Beam Use Proposal*, [sPH-TRG-2020-001](#) (2020).
- [3] K. Chen et al., *A Generic High Bandwidth Data Acquisition Card for Physics Experiments*, *IEEE Trans. Instrum. Measur.* **69** (2019) 4569.
- [4] T. Xuan et al., *Trigger Detection for the sPHENIX Experiment via Bipartite Graph Networks with Set Transformer*, in *Machine Learning And Knowledge Discovery In Databases*, Springer (2023) p. 51–67 [DOI:10.1007/978-3-031-26409-2_4].
- [5] C. Dean, *$D^0 \rightarrow K^- \pi^+$ modelling and selection in min-bias Au+Au simulations at sPHENIX*, [sPH-HF-2023-001](#) (2023) [accessed 20-12-2023].
- [6] Vitis High-Level Synthesis, Xilinx, <https://www.xilinx.com/products/design-tools/vivado/high-level-design.html> [accessed 19-12-2023].
- [7] R. Sarkar, S. Abi-Karam, Y. He, L. Sathidevi and C. Hao, *FlowGNN: A Dataflow Architecture for Real-Time Workload-Agnostic Graph Neural Network Inference*, in the proceedings of the *2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, Montreal, QC, Canada (2023), p. 1099–1112 [DOI:10.1109/HPCA56546.2023.10071015].
- [8] F. Fahim et al., *hls4ml: An Open-Source Codesign Workflow to Empower Scientific Low-Power Machine Learning Devices*, in the proceedings of the *tinyML Research Symposium 2021*, Online Conference, U.S.A. (2021) [arXiv:2103.05579].