## Intelligent Experiment Through Real-Time AI

Fast Data Processing and Autonomous Detector Control for sPHENIX and Future EIC Detectors

> Ming Liu Los Alamos National Laboratory for the Fast-ML Team

DOE Fast-ML Presentations November 30, 2022

# Today's Presentations

#### **1. Overview**, 10'

- Ming Liu (LANL)

#### **2.** Physics simulation and AI-ML algorithms, 10'

- Dantong Yu (NJIT)/Cameron Dean(MIT)/Zhaozhong Shi(LANL) /Tingting Xuan(SBU/NJIT)/Hang Qi(MIT)/Hao-Ren Jeng(MIT)/Beilei Jiang(NTU)

#### 3. HLS4ML and firmware implementation, 5'

- Micol Rigatti (FNAL)/Nhan Tran(FNAL)/Phil Harris(MIT)

#### **4. Demonstrator implementation**, 5'

- Jakub Kvapil (LANL)/Yasser Corrales(MIT)/Noah Wuerfel(LANL)/Jo Schambach(ORNL)/Kai Chen(CCNU)/Lang Lei(CCNU)/Beilei Jiang(NTU)













# Overview

- Ming Liu

Fast-ML Status and Plan @DOE Presentations

## **Project Goals and Deliverables**

## Selective streaming real-time AI and autonomous detector control:

Deliver a demonstrator for p+p and p+Au running for sPHENIX -> generalizable for applications in experiments at the EIC



## Leaderships and Technical Roles

Los Alamos National Laboratory, Ming Xiong Liu, (Lead Principal Investigator) Fermi National Laboratory, Nhan Tran, (Co-PI) Massachusetts Institute of Technology, Gunther M Roland (Co-PI) New Jersey Institute of Technology, Dantong Yu (Co-PI)

Leadership structure of the team The project team will be led by Lead Principal Investigator, Dr. Ming Xiong Liu of LANL, who is accountable to the DOE program leadership for the project's overall success. The team shares the responsibility and accountability for success. Within that structure, lead roles are assigned to co-Principal Investigators (co-PIs), also referred to as key personnel. Dr. Liu will be the lead for hardware design. Dr. Gunther Roland will be the physics lead in sPHENIX and EIC. Dr. Tran will be the lead for Co-Design of AI software and Hardware. Dr. Yu will be lead for Deep Neural Networks Software Design.

#### New teams joined later in early 2022:

- ORNL, sPHENIX/EIC readout integration, Dr. Jo Schambach (sPHENIX MVTX and ePIC readout lead )
- FELIX-AI-Trigger hardware integration, Dr. Kai Chen (FELIX developer at BNL for ATLAS, also sPHENIX)
- NTU, Data acceleration, Prof. Song Fu (SC, data acceleration)

## Technical Approaches and Highlights - I

Objective 1 – Design, build, simulate, and benchmark a prototype streaming readout system with AI-based fast online data processing and autonomous detector control system that meets the physics and engineering requirements. To support this objective, we first aim to generate a large volume of simulation data for heavy flavor decay events. We plan to design a prototype in the simulated and the real sPHENIX experimental environment and later apply the technology in the high luminosity EIC experiments at RHIC. Our objective is to create a working prototype that serves as a baseline and template for future upgrades. With this prototypical working solution, we target to improve the heavy flavor samples from the current 0.05% yield to more than 10+%. (Task 1)





8b/10b MVTX/INTT data (KC705) to FPGA/AI Engine (VC709)

Jakub

## Technical Approaches and Highlights - II

• Objective 2 – Design advanced deep neural networks commensurate with sPHENIX/EIC streaming data requirements. We aim to design deep neural networks with the following goals: (a) network size: neuron weights that fit in the FPGA block RAMs (BRAMs) of the FELIX cards in sPHENIX/EIC experiments, (b) handling the extremely low signal-to-noise ratio of hit images due to the sparse readout of the high-resolution MVTX and INTT detectors, (c) performance improvements: 10% improvement over state-of-the-art triggering algorithms, and (d) minimal performance gap between simulated data and real experiment readouts, and outstanding generalization capability. (Task 2)



## Technical Approaches and Highlights - III

Objective 3 – Deploy advanced deep neural networks within the FELIX system that are capable
of real-time reconstruction of heavy flavor events at high throughput. With the development of
advanced deep neural networks, a parallel strategy is needed to ensure that these networks can be
designed to operate at low latency and high throughput on the FELIX FPGA cards. This challenge
involves detailed AI/hardware co-design to ensure that the desired algorithms can be fit within existing
resources, and can achieve full throughput. (Task 3)



#### **Primary focus:** achieving low latency, real-time processing of data, and deployment of algorithms with high efficiency



Micol/Phil

## sPHENIX Readout and AI-ML HF Trigger Integration



## From sPHENIX to ePIC: Streaming + AI/ML DAQ



Fast-ML Status and Plan @DOE Presentations

# Summary and Outlook

- Carried out full sPHENIX physics and detector simulations of heavy quark and QCD backgrounds
- Developed preliminary AI-algorithms for sPHENIX HF triggers
- Completed MVTX SRO
- INTT SRO demonstrated
- Implemented a toy AI-algorithm in HLS4ML in FELIX
- Work in progress to implement full sPHENIX trigger in a simplified hardware
- On track to complete a demonstrator in 2023

#### Future plan/proposal:

- Extend project by 2 years, 2024 ~ 2025
- Implement the demonstrator for sPHENIX p+p run in 2024
- Develop ePIC TDR of SRO with AI/ML for EIC CD2(2024) and CD3(2025) based on our work

	From DOE proposa						
TASK I	TASK II	TASK III	TASK IV				
Year 11. Software development and system design. We will first perform detailed sPHENIX physics and detector simulations to design a real-time fast data processing and autonomous detector control and calibration system. In the meantime, we will survey currently available AI models and design a system for offline training and domain adaptation for data and MC. The physics and detector simulation results and the performance of hardware are used to tune the AI algorithms.2. Hardware development and system integration. We will take advantage of the streaming readout capability of the sPHENIX tracking system to implement continuous readout of two fast silicon tracking subsystems, MVTX and INTT. A FPGA based fast tier-1 AI system will be developed to identify heavy flavor (HF) events in $p+p$ collisions,							
By Q2 ► Generate open heavy flavor and QCD back- ground events for simula- tions (LANL, MIT)	► HF trigger algorithm development for FPGA (FNAL, LANL, NJIT)	<ul> <li>MVTX streaming read- out (LANL)</li> <li>INTT streaming read- out (MIT)</li> </ul>	<ul> <li>Beamspot interaction and readout simulation (FNAL)</li> <li>Displaced tracks and anomaly simulation (FNAL, MIT)</li> </ul>				
<ul> <li>By Q3</li> <li>▶ Develop fast tracking algorithms using MVTX and INTT hit information (LANL, MIT, NJIT)</li> </ul>	► Design real-time GPU training machine (MIT, NJIT)	▶ hls4ml implemen- tation and algorithm development (FNAL, MIT)	<ul> <li>Preliminary design of streaming and automated controls of online GPU- based training system (MIT, NUT)</li> </ul>				
By Q4 ► Complete a prelimi- nary design of HF trig- ger AI offline (FNAL, LANL, NJIT)	▶ ML, Graph NN train- ing, by NJIT and MIT	► FPGA implementa- tion of HF trigger with MVTX and INTT (All)	► Simulation and training (MIT, NJIT)				
	Yea						
We will focus on the system	n integration and continue to	b improve and benchmark th	e performance of software,				
By Q5 & Q6 ► Interface between AI system and MVTX detector Data Input by (FNAL, LANL) ► Interface between AI system and TPC Readout Control (FNAL, LANL)	► Design new GNNs (Encoder, Attention, Particle-Net) algorithms with hls4ml (MIT, NJIT)	<ul> <li>hls4ml customization for FELIX board (FNAL, LANL)</li> <li>FPGA, GPU system in- tegration and evaluation (MIT, NJIT)</li> </ul>	► GPU deployment for autoencoder and training (FNAL, MIT)				
By Q7 ► Continue to improve algorithms for HF tag- ging (LANL)	► Improve algorithm with hls4ml on FPGA (FNAL, NJIT)	<ul> <li>Multi-FELIX Board Integration (LANL)</li> <li>Validation and test with FELIX boards (All)</li> </ul>	► ML model and domain adaption update (MIT, NJIT)				
By Q8 11/11/22 ► Benchmark system performance with sPHENIX or test beam data (All)							

**Table 2: Tasks and Milestones** 



# Physics simulation and AI-ML algorithms

- Dantong Yu

Fast-ML Status and Plan @DOE Presentations

# Al Trigger Pipeline

1. Fetch events from event buffer to Processing

2. Data Pre-processing Clustering

- 3. Tracking + Outlier hits Removal
- 4. Triggering
- 5. Triggers on TPC







# Simulations

- Most simulations have focused on sPHENIX
  - PIC detector simulations for EIC are under rapid changes
- We simulated 2 different physics events (all 200 GeV pp collisions)
  - > Minimum bias with all heavy flavor decays rejected for background
  - >  $D^0 \rightarrow K^- \pi^+$  in minimum bias for signal
- Simulations have been improved throughout the year
  - Added full service material for MVTX
  - Added realistic hit duplication in MVTX from pixel pulse length

## Dataset

- > 8 million Events 50% signal/noise.
- > 8 million Events 0.1% signal/noise.

# Algorithm Flow

- Strategy is to construct a large scale AI algorithm with all elements of trigger
  - Algorithm is factorizable into core physics components

  - Emulates the normal reconstruction workflow Use of core components allows for intermediate physics validation



# Tracking as Graph Neural Networks



- Graph neural networks are a popular way of posing tracking problems in physics
  - Sparse hits in space do not fit traditional ML architectures
     Growing sub-field of geometric deep learning
- Data is structured as a graph of connected hits
  - Connect plausibly related hits using geometric constraints and learn best associations and parameters of connected hits (tracks)

## Three Types of Potential Interactions in Model Event Decay



# Bipartite Graph Networks with Set Transformer (BGN-ST) Model Architectures



## Set Encoder with Bipartite Aggregator (SEBA) Blocks



## Experiments: Physics Driven BGN-ST

Table 2: Comparison to Baseline Models with Estimated Radius.									
	with	LS-radius		without radius					
Model	#Parameters	Accuracy	AUC		#Parameters	Accuracy	AUC		
Set Transformer	80,002	86.17%	91.75%		79,810	72.04%	78.92%		
GarNet	$284,\!210$	86.22%	91.81%		284,066	72.59%	79.61%		
PN+SAGPool	780,934	86.25%	92.91%		780,678	69.22%	77.18%		
BGN-ST	$363,\!426$	$\mathbf{87.56\%}$	93.22%		$363,\!170$	74.13%	81.81%		

1% signal/ba	ckground rat	io	0.1% signal/background ratio			
Background Rejection	Efficiency	Purity	<b>Background Rejection</b>	Efficiency	Purity	
90%	72.5%	7.25%	90%	78%	0.78%	
95%	48.9%	9.78%	95%	50%	1.0%	
99%	15.0%	15.0%	99%	17%	1.7%	
99.33%	10.5%	15.74%	99.33%	11.0%	1.65%	

## Status and Outlook

## • Demonstrated BGN-ST outperforms selected state-of-the-art methods <sup>[1,2]</sup>

- Improves the task accuracy and AUC score by about 15%
- Architecture benefits from pairwise interactions between tracks and allows a two-way scattering and gathering for effective information exchange and adaptive graph pooling
- Adopts the physics-aware concept and introduces explicit physics properties such as transverse momentum
- Next step: From optimal performance of state-of-the-art methods, develop firmware implementations
  - Understand what is feasible and condense the model into latency and resource constraints

[1] Xuan, Y. Zhu, G. Borca-Tasciuc, M. X. Liu, Y. Sun, C. Dean, Y. C. Morales, Z. Shi, D. Yu, End-To-End Pipeline for Trigger Detection on Hit and Track Graphs in, Accepted by Thirty-Fifth Conference on Innovative Applications of Artificial Intelligence, IAAI.

[2] Xuan, G. Borca-Tasciuc, Y. Zhu, Y. Sun, C. Dean, Z. Shi, D. Yu, Trigger Detection for the sPHENIX Experiment via Bipartite Graph Networks with Set Transformer in European Conference on Machine Learning and Principles and Practice of Knowledge Discovery in Databases (ECML-PKDD 22).





# hls4ml translation and firmware implementation

- Micol Rigatti and Phil Harris

Fast-ML Status and Plan @DOE Presentations

## Hardware and Firmware Implementations

Selective data streaming in sPHENIX will use the FELIX board. FELIX is a 16-lane Gen-3 PCIe card with 48 transmitters and receiver optical links. The on-board FPGA is a Kintex Ultrascale XCKU115FLVF1924-2E



## Hardware Configuration



A second FELIX is connected to the first FELIX through the optical transceivers, as a **dedicated FPGA hardware for <u>smart control and real-time</u> <u>decision-making</u> for TPC readout in the selective data streaming architecture.** 

The **AI Engine** will search for displaced tracks to identify tracks from **heavy quark decays** that are pointing away from the nominal beam center. Such a signal will initiate the readout of the TCP detector.

The firmware of the AI Engine will implement the **Neural Network** deployed using *his4mi*, the PCIe interface, and the optical link interface.

## Schedule – completed







## On the Host Computer, the **FELIX driver** and the **software application** is been installed.



The **data movement** between the BNL711 and the Host Computer is been tested. The data are emulated inside the BNL711.



The hsl4ml workflow is been tested with the Kintex Ultrascale XCKU115 as a target generating the IP of a **simple NN** 



The IP of the simple NN is in the integration phase with the **Wupper** module (**PCIe engine**)



## Schedule – next steps and challenges



## **Test** the functionality of the NN and the data routing mechanism



The FELIX card is a **data router**. It needs an application to be instructed on the data movement.

The **application** relies on an **OS** and a **driver** interfacing the FELIX card.

- Inserting the NN in the FELIX Firmware will change how the driver interfaces with the card
- The testing part right now relies on the interplay of Hardware, Firmware, and Software



Interfacing the board to test the NN is the other real challenge

## Fit the NN in the FELIX Firmware complying with timing constraints.

The workflow of hls4ml generates an IP considering available the resources of the target FPGA. We need to route the NN considering the already routed FELIX Firmware.



2

Meeting the timing in this condition is the real challenge

**NEXT STEP: the analysis of the FELIX Firmware in terms of resources** to drive the hls4ml implementation of the NN



**NEXT STEP**: the integration of the IP of the simple NN is in with the Wupper module (**PCIe engine**)

# GNN Implementation in hls4ml

### First version of hls4ml-based Graph encoder available



Critical optimization needed in construction of Graph-based mapping

## GNN Planned Upgrade in hls4ml

Design	$(n_{\rm nodes}, n_{\rm edges})$	RF	Precision	Latency [cycles]	II [cycles]	DSP [%]	LUT [%]	FF [%]	BRAM [%]
Throughput-opt.	(28, 56)	1	ap_fixed<14,7>	59	1	99.9	66.0	11.7	0.7
Throughput-opt.	(28, 56)	8	ap_fixed<14,7>	75	8	21.9	23.8	4.7	0.7
Resource-opt.	(28, 56)	1	ap_fixed<14,7>	79	28	56.6	17.6	3.9	13.1
Resource-opt.	(448, 896)	1	ap_fixed<14,7>	470	174	56.6	25.0	7.4	16.5
Resource-opt.	(448, 896)	8	ap_fixed<14,7>	1590	520	5.6	25.0	7.4	16.3

@200 MHz, 1590 Cycles  $\rightarrow$  7.5 $\mu$ s

### Active work is underway to improve the GNN implementation

- Base implementation has been updated twice in 2022
  - > A third update will start in mid December, focusing on 3 examples
    - Example 1: Tri-muon reconstruction with the LHC (muon endcaps)
    - Example 2: Heavy flavor tracking at sPHENIX
    - Example 3: Silicon strip tracking at LHC
  - Current Graph encoder to be optimized (adjacency matrix computation)
    - Aiming to centrally rework this with core hls4ml developers
    - Will be a central project across several domains



# Demonstrator Implementation

- Jakub Kvapil

Fast-ML Status and Plan @DOE Presentations

## Demonstrator Development @ LANL, ORNL and CCNU

## Why a demonstrator?

- sPHENIX DAQ still underdevelopment
   ➢ Parallelize development and be ready for final deployment
- 2. Not enough FELIX boards fabricated yet
  - The AI core logic will be implemented on VC709 which is the FELIX protoboard (share similar resources and is supported)



## **Demonstrator Implementation**



Both KC705 and VC709 will be replaced by FLX712 at deployment stage

## sPHENIX RAW Data Simulation

- In order to test the full loop feedback AI system detector hits must be known.
- Code developed to transform MC JSON hit patterns into MVTX and INTT raw data streams





LANL

## KC705 Raw Data Transmission

- Simulated data must be streamed to the AI VC709/FLX712 FPGA
- KC705 is used to read the MC data via PCI XDMA and transmit them using g-Links
- One link for MVTX and one for INTT



CCNU

## FLX712 Raw MVTX Data Transmission to AI-FPGA



ORNL

# **Timeline and Outlook**



# Backup slides

## Event Timing: Reject out of time MVTX hits with INTT



11/30/22

# Feed MVTX/INTT MC Hits to AI Engine



# sPHENIX and EIC Schedules

#### sPHENIX Run Plan: 2023-2025

• pp and pAu run in 2024

Year	Species	$\sqrt{s_{NN}}$	Cryo	Physics	Rec. Lum.	Samp. Lum.
		[GeV]	Weeks	Weeks	z  <10 cm	z  <10 cm
2023	Au+Au	200	24 (28)	9 (13)	3.7 (5.7) nb <sup>-1</sup>	4.5 (6.9) nb <sup>-1</sup>
2024	$p^{\uparrow}p^{\uparrow}$	200	24 (28)	12 (16)	0.3 (0.4) pb <sup>-1</sup> [5 kHz]	45 (62) pb <sup>-1</sup>
					4.5 (6.2) pb <sup>-1</sup> [10%-str]	
2024	$p^{\uparrow}$ +Au	200	-	5	0.003 pb <sup>-1</sup> [5 kHz]	$0.11 \ {\rm pb^{-1}}$
					0.01 pb <sup>-1</sup> [10%-str]	
2025	Au+Au	200	24 (28)	20.5 (24.5)	13 (15) nb <sup>-1</sup>	21 (25) nb <sup>-1</sup>

#### EIC Project Plan (as of 11/05/2022)

- ePIC TDR for CD2/CD3A, 2024
- Final design/construction, 2025

## EIC Reference Schedule - V3



## EIC: ePIC DAM Candidates: ATLAS "FELIX"



Current ATLAS Phase 1/ sPHENIX FELIX BNL-712v2

(Hao Xu, BNL, DAQ WG meeting 7/2021)



Assembled 24ch FLX-181 with 25 Gbps FireFly FMC+



#### Current Design of FELIX FLX-182



Fast-ML Status and Plan @DOE Presentations

## **FELIX Status**

## FLX-182 Status

- Design passed FELIX review, will be sent out for fabrication in this week
- First assembled board is expected to be delivered in early September 2022
- 7 boards will be produced if there's no big design issues, by December 2022
- Small production for more boards is possible once FPGA is available

## **Plan for 48-ch FELIX**

- FPGA: Versal Premium, e.g. VP1552
- Transceivers: Up to 100+ GTYP/GTM
- PCle Gen 5 up to 16 lanes
- If FPGA is available as planned, design will start in Q1 of 2023, first board is expected to be available in Q3 2023.

## **Architecture and Interfaces**

- PCIe Gen 4 x 16 lanes
- Transceiver
  - Transceiver Type: Samtec FireFly transceiver
  - Transceiver Speed: up to 10 Gb/s ("CERN-B") or 25 Gb/s
- Number of Optical Connectors per Card
  - At least 24 bi-directional connections to front-end electronics
  - A separate bi-directional connection to the TTC/BUSY system
- Configuration
  - Boot from JTAG/QSPI/SD card
  - Remote FPGA configuration from Multiple Flash Partitions
- DDR4/Flash Memory/SD card
- I2C
- External Electrical Interface
- Voltage Protection
- Temperature Protection