Summary of TDAQ Sessions
CPAD Instrumentation Frontier Workshop 2021

Georgia Karagiorgi, Columbia University
and
Isobel Ojalvo, Princeton University

March 22, 2021

Exciting developments in Trigger and DAQ Instrumentation

Day 1

12:00	Mu2e TDAQ and slow control systems	Antonio Gioiosa	Stony Brook, NY	12:00 - 12:15
	Development of the Mu2e electromagnetic calorimeter front-end and readout electronics	Franco Spinella et al.	Stony Brook, NY	12:20 - 12:35
	Does anybody really know the time it is?	Roger Rusack	Stony Brook, NY	12:40 - 12:55
13:00	FELIX: the Detector Interface for the ATLAS Experiment at CERN	Alexander Paramonov	Stony Brook, NY	13:00 - 13:15
	The Expandable Modular ATCA hardware design for high-energy physics experiment applications	Alexander Madorsky	Stony Brook, NY	13:20 - 13:35
	Break		Stony Brook, NY	13:40 - 14:00
14:00	Liquid Argon Time Projection Chamber Trigger Development with MicroBooNE and SBND	Daisy Kalra et al.	Stony Brook, NY	14:00 - 14:15
	Trigger and Data Acquisition for the Mu2e-II experiment	Richard Bonventre et al.	Stony Brook, NY	14:20 - 14:35
	Streaming data acquisition system for CLAS12 Forward Tagger	Mariangela Bondi'	Stony Brook, NY	14:40 - 14:55
15:00	Designing a 30 MHz GPU trigger, the LHCb experience	Vladimir Gligorov	Stony Brook, NY	15:00 - 15:15
	Detector design for a Muon Collider experiment	Sergo Jindariani	Stony Brook, NY	15:20 - 15:35

Day 2

12:00	Real-time analysis in Run 3 with the LHCb experiment	Mika Vesterinen	Stony Brook, NY	12:00 - 12:15
	Towards an Interpretable Data-driven Trigger System for High-Throughput Physics Facilities	David Miller	Stony Brook, NY	12:20 - 12:35
	Global Trigger for the ATLAS Phase-II upgrade	Jochen Heinrich	Stony Brook, NY	12:40 - 12:55
13:00	The Challenges of Machine Learning at the Edge	Audrey Corbeil Therrien	Stony Brook, NY	13:00 - 13:15
	Coprocessors as aservice for accelerated inference of DL algorithms	Jeffrey Krupa	Stony Brook, NY	13:20 - 13:35
	Break		Stony Brook, NY	13:40 - 14:00
14:00	Triggering on Long-Lived Particles decaying to Hadronic Showers in CMS Muon System	Ka hei martin Kwok	Stony Brook, NY	14:00 - 14:15
	Real-time Artificial Intelligence for Accelerator Control: A Study at the Fermilab Booster	Christian Herwig	Stony Brook, NY	14:20 - 14:35
	hls4ml: enabling real-time deep learning in HEP trigger and DAQ systems	Jennifer Ngadiuba et al.	Stony Brook, NY	14:40 - 14:55

A total of 18 contributions spanning Mu2e, ATLAS, MicroBooNE, SBND, LHCb, CMS, related experiment-specific upgrades/developments, and experiment-agnostic developments and applications

Common TDAQ Themes in 2021

- Next-generation TDAQ Systems developed as scalable systems with high level of parallelization, for high-throughput facilities, with an eye toward “self-driving” systems (see, e.g., Pie in the Sky talk by D. Miller)

- New TDAQ designs and performance are benefiting from **AI developments**, **hardware processor advancements**, and **design tools developments**

- High degree of cross-experiment collaboration and co-development

Talk by Antonio Gioiosa

Highlights: **Mu2e TDAQ and slow control systems**

Mu2e requirement is to process 200k events/s
Readout scheme designed with capability of 40 GB/s data readout → 280 MB/s to disk

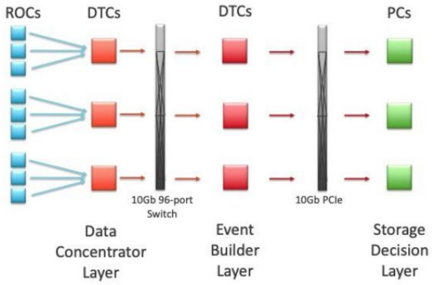

- Mu2e Experiment is under construction at Fermilab and will be ready for data taking next years
- Mu2e TDAQ and slow control are in large part developed according to the requirements (200K events/s for data taking) and hardware tests are going on
- Slow control integration in the online DAQ system, *otsdaq*, provides an advanced slow controls monitoring, an interface to send *otsdaq* front-end DAQ hardware, data processing and DQM slow controls informations to EPICS, and a real configuration and Integration with the *otsdaq* State Machine

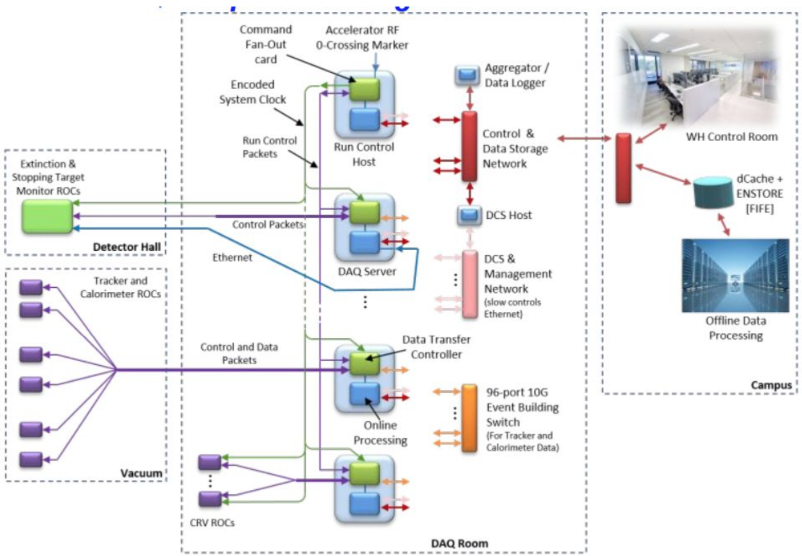

Online DAQ is off-the-shelf DAQ (otsdaq) based on artdaq

Provides a library of supported front-end boards and firmware modules, and integrated Run Control GUI and readout software

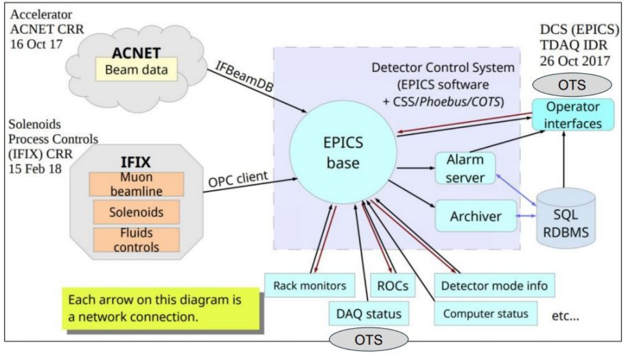

Talk by Franco Spinella

Highlights: **Development of the Mu2E electromagnetic calorimeter front-end and readout electronics**

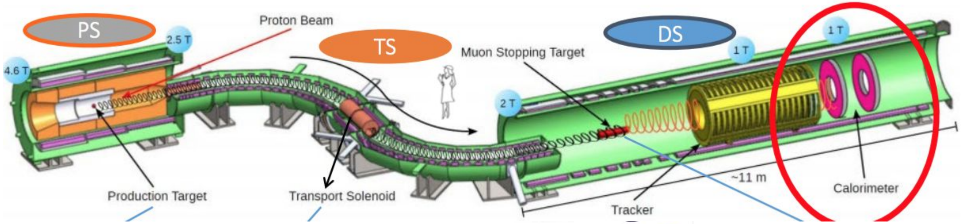

DIRAC: custom waveform digitizer

Intense campaign of tests to determine components

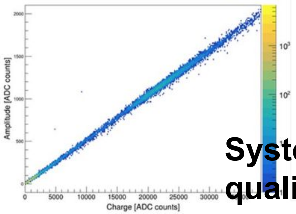

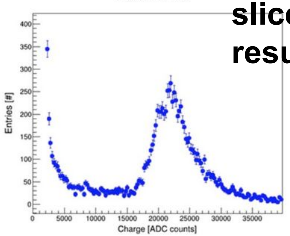

Fig.B: Energy distribution deposited by cosmic rays (ADC counts)

System design qualified with slice test results

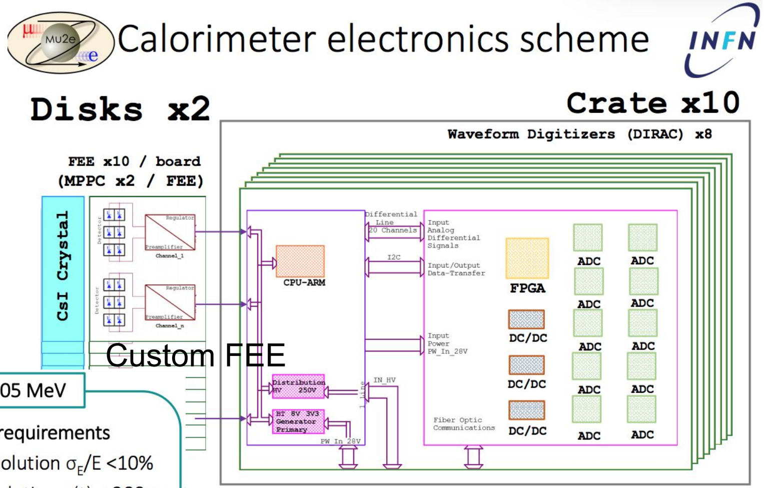

@ 105 MeV

Calorimeter requirements

- energy resolution $\sigma_E/E < 10\%$
- timing resolution $\sigma(t) < 200$ ps
- position resolution < 1 cm
- Work in vacuum @ 10^{-4} Torr
- 1 T Magnetic Field

- Sampling frequency of *200 MHz*
- ADC with *12 bits resolution*

CsI crystals + SiPMs
High radiation environment

Talk by Roger Rusack

Highlights: **"Does Anybody Really Know What Time it is?"**

State-of-the-Art Today

LpGBT used to distrubute a high precision clock derived by clock-recovery from the 2.56 Gbs downlink control signal.

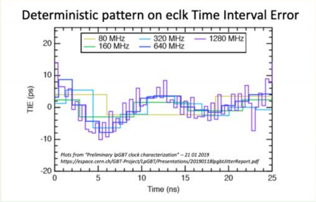

From Talk by T. Kugathasan December 2020.

LpGBT-v0

- Random jitter 2.2 ps
- Deterministic jitter peak-to-peak 25 ps.

Source identified and LpGBT-v1 expected to reduce deterministic jitter.

As we push the limits of time measurements to ~1 ps, we need to have reference clocks that are stable at < 1ps.

How Do We Correct For Drifts?

To solve the problem of how to align clocks that drift we have made a multi-cell planar wave guide in TSMC 65 nm process.

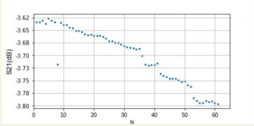

Digitally Controlled Phase Shifter — DCPS

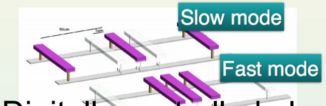

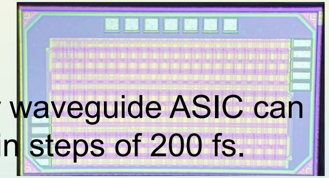

Digitally controlled planar waveguide ASIC can delay digital clock signal in steps of 200 fs.

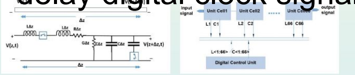

Measured delay step is 200 fs

'Pure' Clock Distribution System

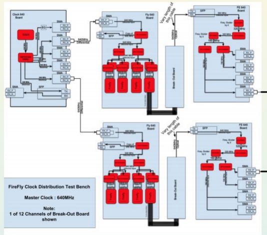

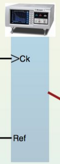

Microsemi 5275A

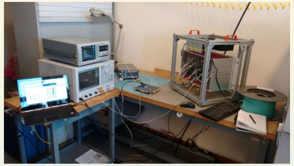

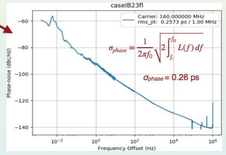

RF Quality Fanout from ON Semiconductors.
NB7VQ1005MMNG — max data rate 10 Gbps.

Phase Noise Plot 0.1 Hz - 1 MHz

Use DDMTD to measure time drifts averaged over many cycles.

Basic method that goes back to FM is radio is to heterodyne the signal.

Digital Dual Mean Time Difference (DDMTD) circuit*

Offset clock with $f_{off} = f_{ref}(1 - 1/N)$

f_{ref}

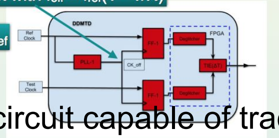

DDMTD Schematic

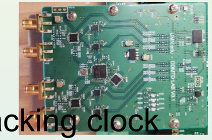

First version of DDMTD

Low-cost circuit capable of tracking clock drifts at sub-ps level

Pure clock distribution system demonstrates sub-ps jitter levels, exceeding current state-of-the-art with clock recovery.
→ Can deliver stable clock with low jitter and low wander.

CPAD Trigger/DAQ Mar 22, 2021

Talk by Alexander Paramonov

Highlights: **FELIX: The new detector readout system for the ATLAS experiment**

FELIX
Front End LInk eXchange

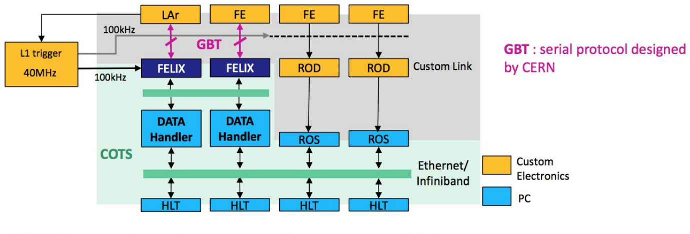

- FELIX is a router between custom serial links and a commodity network
  - Takes advantage of the latest technology to simplify the ATLAS readout
- In LHC Run-3 (2021-2023) FELIX will be used for selected detectors and trigger systems.
  - FELIX firmware and the software are mature
  - Most of the boards have been produced
- Ongoing efforts:
  - Integration with the ATLAS on-detector (front-end) systems
  - Development of FELIX board and firmware for LHC Run-4.

FELIX hardware platform

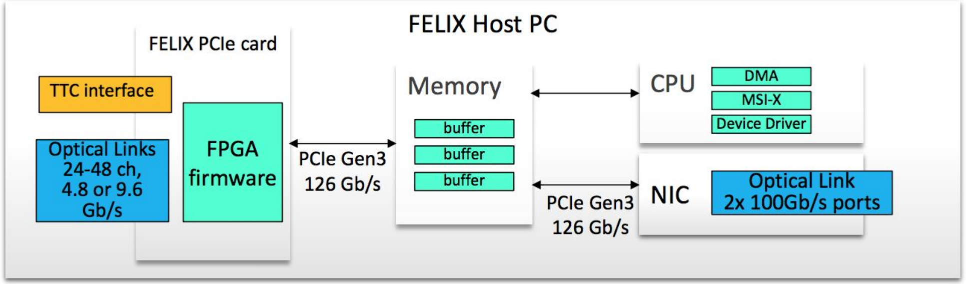

PCIe card with FPGA chip + Host PC + NIC

Talk by Alexander Madorksy

Highlights: **X2O: A Modular ATCA Design with Cost Optimization Options**

A modular system in ATCA standard, designed with cost optimization in mind, and customizable for user's needs
Approach: eliminate base board, use large heat sink as mechanical platform,
Several modules: FPGA modules attached to heat sink, optical module near standard front panel, power module near backplane
Cost range: $11-13k, for 10G vs 25G QSFP devices

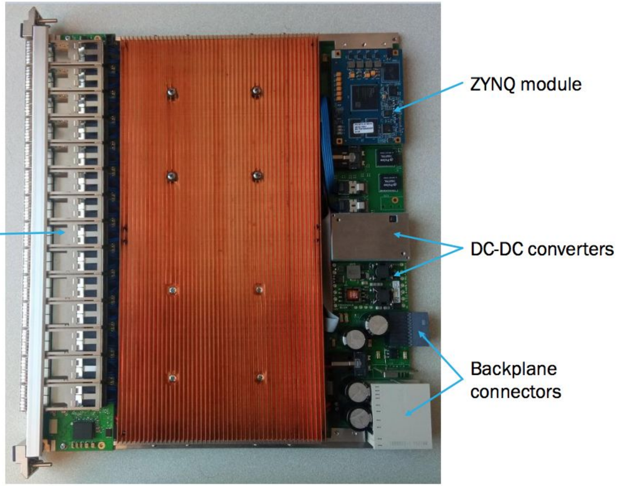

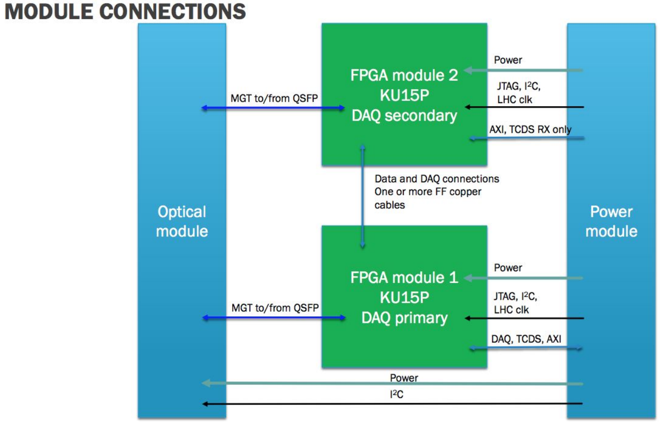

X2O KU15P MODULE

Prototypes of all modules produced and tested (power, FPGA, QSFP)
Completely assembled board tested (including thermal tests at full power), and three units in operation (UF, CERN, UCLA), plus more modules in assembly.

Talk by Daisy Kalra

Highlights: **Liquid Argon TPC Trigger Development with MicroBooNE & SBND**

Planned in situ demonstrations of TPC triggers (including ML based) with currently & soon to be operating MicroBooNE & SBND LArTPCs.

Triggers to be demonstrated first online in CPU/GPU, and then in FPGA over the next 1-2 years.

hls 4 ml

Sample	Train Size	Test Size	Accuracy (%)			Inference Time (ms)
			ϵ_{NB}	ϵ_{LE}	ϵ_{HE}	
NB	12,023	4,027	99.53	0.47	0.12	
LE	12,050	3,970	4.01	94.48	1.51	1.6±0.1
HE	10,137	3,417	3.63	6.15	90.22	

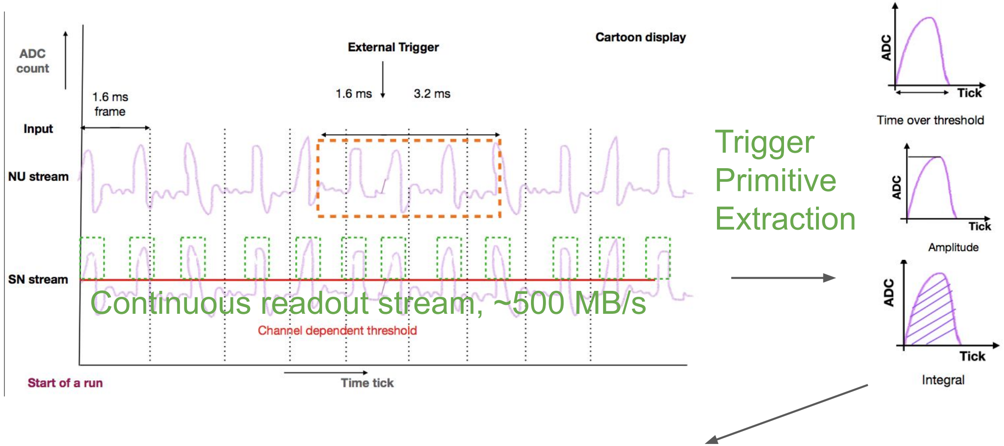

Higher Level Trigger Decision (clustering, or image classification)

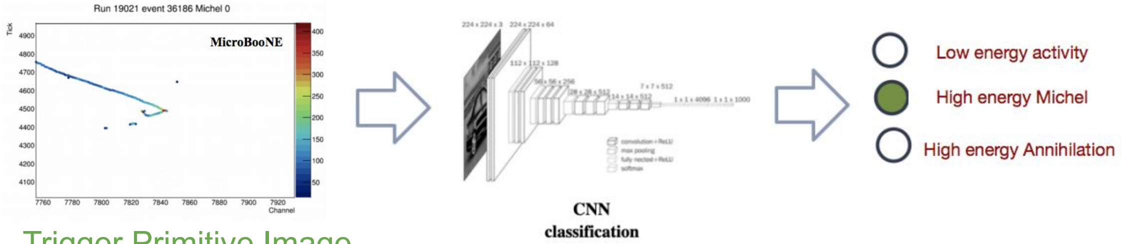

Trigger Primitive Image

CPAD Trigger/DAQ Mar 22, 2021

Talk by Richard Bonventre

Highlights: **Trigger and DAQ for the Mu2e-II experiment**

Physics motivation:
Improve sensitivity to charged-lepton flavor violating (CLFV) neutrino-less conversion of a nuclear bond muon into an electron by an order of magnitude over Mu2e.

- Also 10x radiation levels and 20x data rate!

Derived from A. de Gouvea, P. Vogl, Prog. Part. Nucl. Phys. 71 (2013) 75

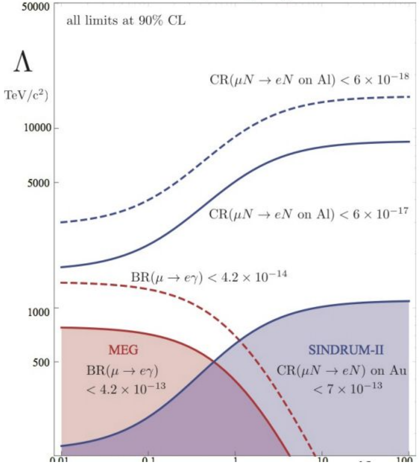

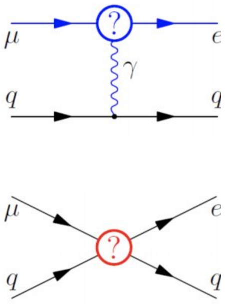

Mu2e-1 Trigger DAQ consists of multiple layers starting with Readout Controllers

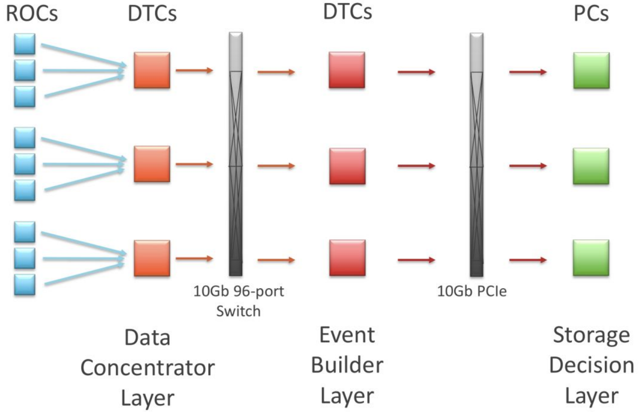

Mu2e-II Trigger DAQ Options: Trigger-less TDAQ, GPU co-processors, FPGA pre-filtering, FPGA pre-processing and trigger primitives

- Trigger-less TDAQ - requires 10x more hardware
- HLT parallelization non-trivial!
- L1 Hardware trigger implemented on FPGAs utilize HLS
- 2-level TDAQ system running on FPGAs

CPAD Trigger/DAQ Mar 22, 2021

Talk by Mariangela Bondi

Highlights: **Streaming DAQ system for CLAS12 Forward Tagger**

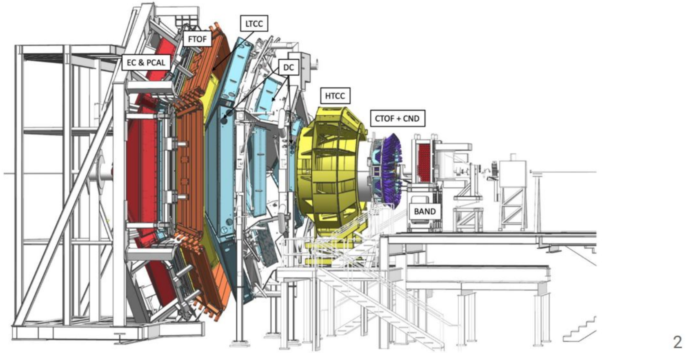

Installed at JLAB Hall-B

- Investigation of the structure of the proton and neutron in ground and excited states

Recent progress towards testing streaming read out of data with full TriDAS chain to study performance

- Must reduce 50 MB/s to ~4 MB/s
- Perform reconstruction in real time
- Full chain tested successfully
- New AI algorithm improves mass resolution and real time clustering

Prototype is being used as the basis for developing a large system for the entire CLAS12 detector

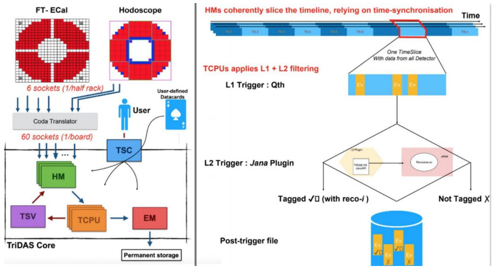

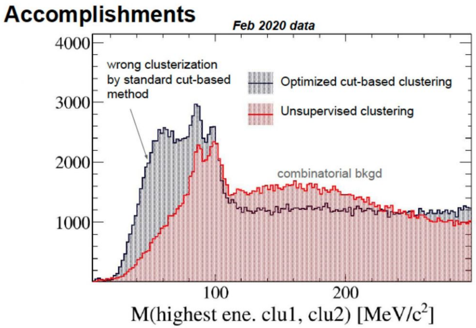

Talk by Vladimir Gligorov

Highlights: **Designing a 30 MHz GPU trigger, the LHCb Experience**

0.5 - 1% momentum resolution

15 + 29/p_T[GeV] micron impact parameter resolution

(1+10/√E[GeV])% ECAL resolution

Efficient and high-purity identification of all five stable charged particle types (pion, kaon, proton, electron, muon) over the momentum range 2-100 GeV

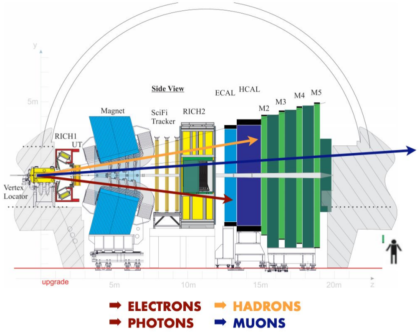

A general-purpose forward spectrometer at the LHC, optimized for heavy-flavour physics

Typical triggers select signal needles in Standard Model haystacks
LHCb needs to sort and compresses haystacks of needles — Real Time Analysis!

LHCb endeavored to move from a triggered to streaming reconstruction system after Run 2

- Required track reconstruction at 30 MHz!
- Initial endeavor to put single algorithms on GPUs, eventually found porting "everything" to GPUs was more cost-effective (though more technically difficult)

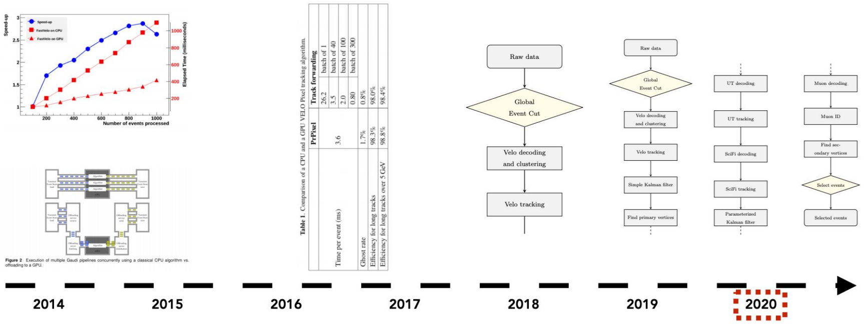

- Integration tests show this is ready for data taking
- Only O(200) GPUs required!

Talk by Sergo Jindariani

Highlights: **Detector for a Muon Collider (TDAQ)**

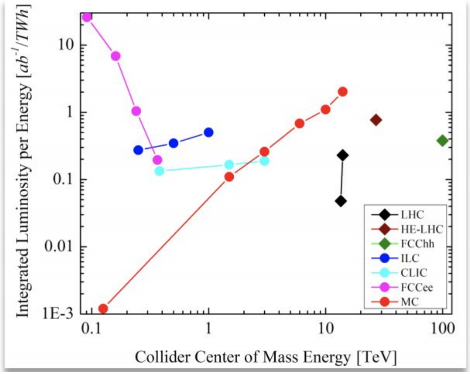

Physics potential for a Muon Collider is formidable, but so is the Beam Induced Background (BIB) Due to the Muon's short lifetime

Data rates and timeline should allow for streaming readout, some R&D needed for FE chips

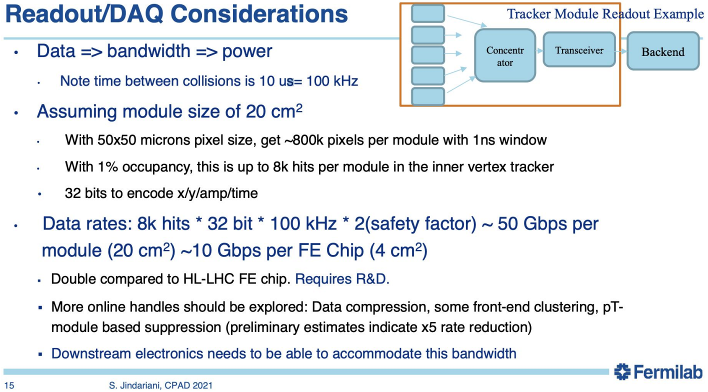

Readout/DAQ Considerations

- Data => bandwidth => power
  - Note time between collisions is 10 us= 100 kHz
- Assuming module size of 20 cm^2
  - With 50x50 microns pixel size, get ~800k pixels per module with 1ns window
  - With 1% occupancy, this is up to 8k hits per module in the inner vertex tracker
  - 32 bits to encode x/y/amp/time
- Data rates: 8k hits * 32 bit * 100 kHz * 2(safety factor) ~ 50 Gbps per module (20 cm^2) ~10 Gbps per FE Chip (4 cm^2)
  - Double compared to HL-LHC FE chip. Requires R&D.
  - More online handles should be explored: Data compression, some front-end clustering, pT-module based suppression (preliminary estimates indicate x5 rate reduction)
  - Downstream electronics needs to be able to accommodate this bandwidth

Fermilab

15 S. Jindariani, CPAD 2021

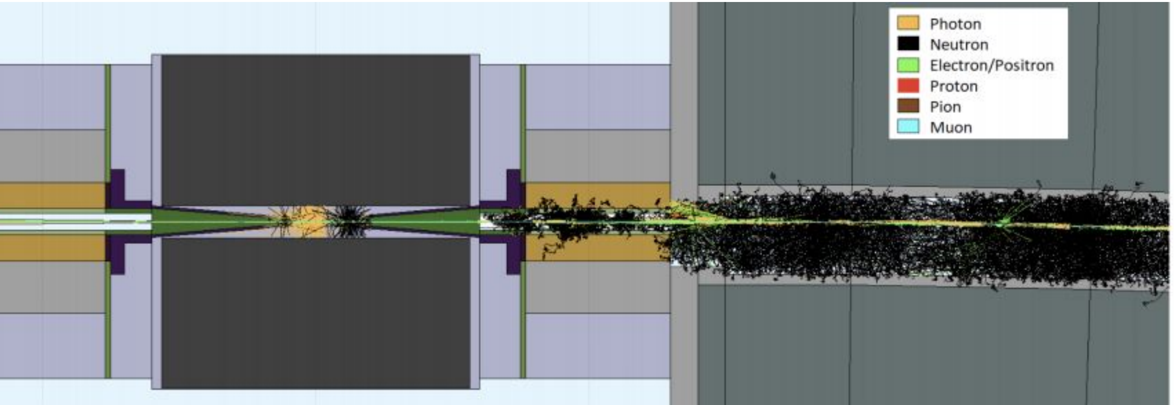

- Data rate for tracker alone are 30 Tbps with 1ns readout window
- Calorimeters will contribute ~the same
- 100-200 readout boards at 10-20 Gbps

CPAD Trigger/DAQ Mar 22, 2021

Talk by Mika Vesterinen

Highlights: **Real time analysis in Run 3 with the LHCb Experiment**

JINST 14 (2019) P04013

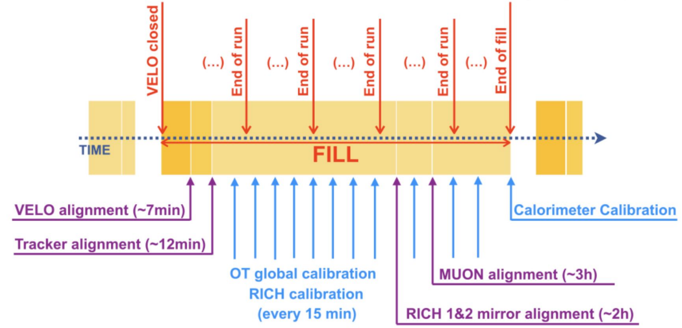

((~7min),(~12min),(~3h),(~2h)) - time needed for both data accumulation and running the task

LHCb streaming DAQ system requires continuous calibration online in order to reconstruct offline-quality objects in real-time

- Developing a global calibration system shiwhc will run every 15 minutes during the fill

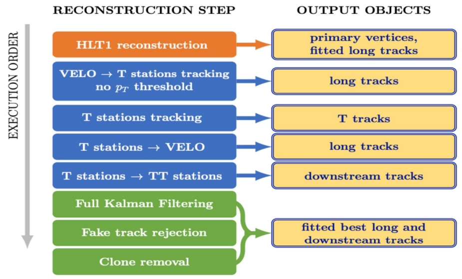

Plus full upfront RICH and CALO reconstruction

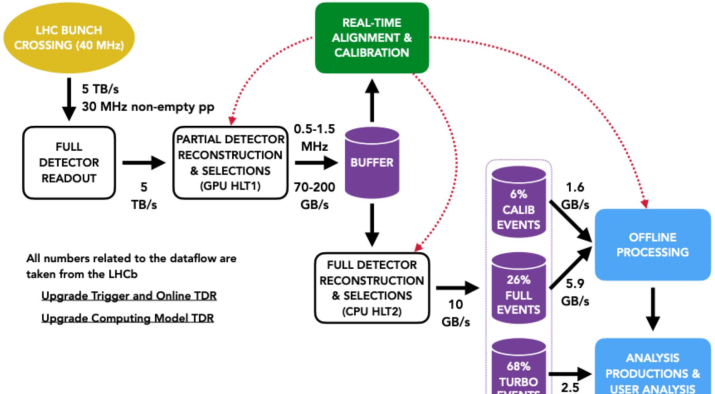

See Vava's talk on the GPU HLT1

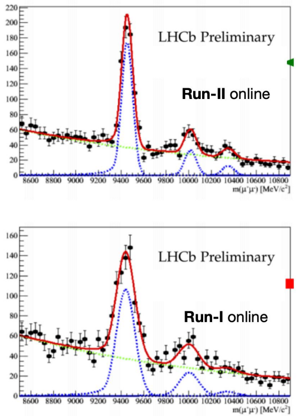

LHCb-FIGURE-2020-007

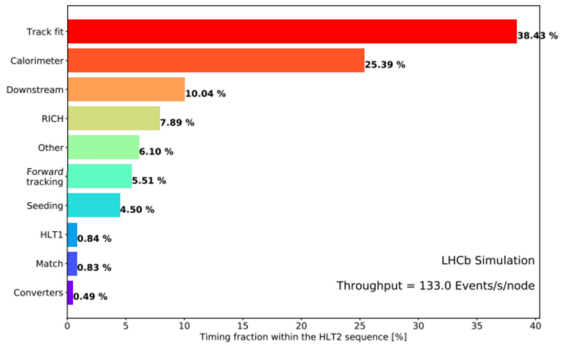

Talk by David Miller

Highlights: **Towards an Interpretable Data-driven Trigger system for High- Throughput Physics Facilities (Pie-in-the-Sky talk!)**

Envisioning a Self-Driving trigger system

What has been learned such that an update is merited?

- Interpret the output of the algorithm
- "Why" was the event triggered?
- What trigger algorithm was "most important" to the trigger decision?

What are the impacts of those updates?

- Given a definition of the resource cost of a set of triggers, how can we optimize the algorithm execution and usage to minimize that resource usage?
- Cost might include bandwidth considerations, CPU time, data preparation, etc

Future work: stream-based active learning

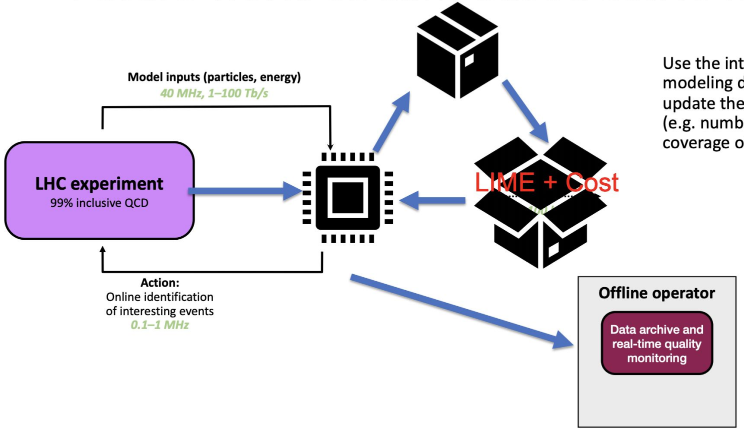

Use the interpretable and cost-effective modeling described in the previous slides to update the trigger selections and algorithms (e.g. number of jets) required to maintain coverage of key physics processes.

Goal to minimize the "cost" (Rate!)
… While maximizing the performance!
- Study performed using CMS open-data
- Future work to implement stream-based active learning

CPAD Trigger/DAQ Mar 22, 2021

Talk by Jochen Heinrich

Highlights: **Global Trigger for the ATLAS Phase-II Upgrade**

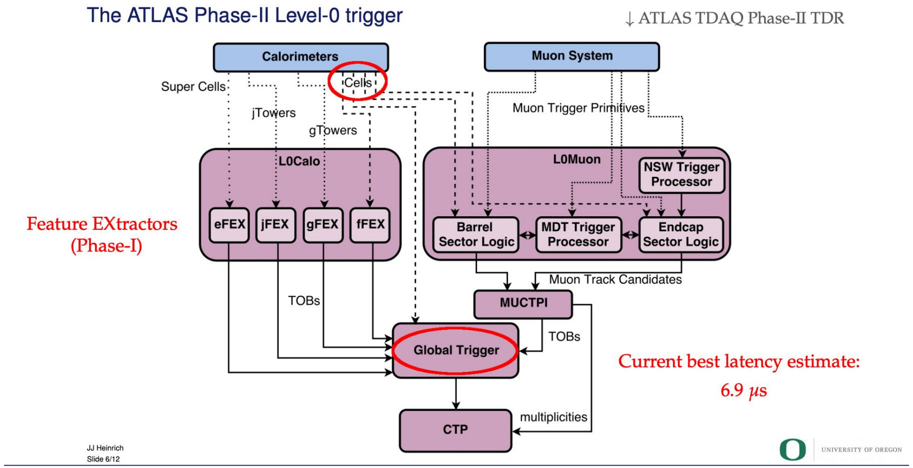

Preparation for HL-LHC
GT trigger for Phase II is primarily implemented in FW (FPGAs)
The Single-Layer system will allow rate to increase from 100 kHz to 1MHz with 10 microseconds latency

Example: Breaking down 3D clustering algorithms to several 2D clustering algorithms

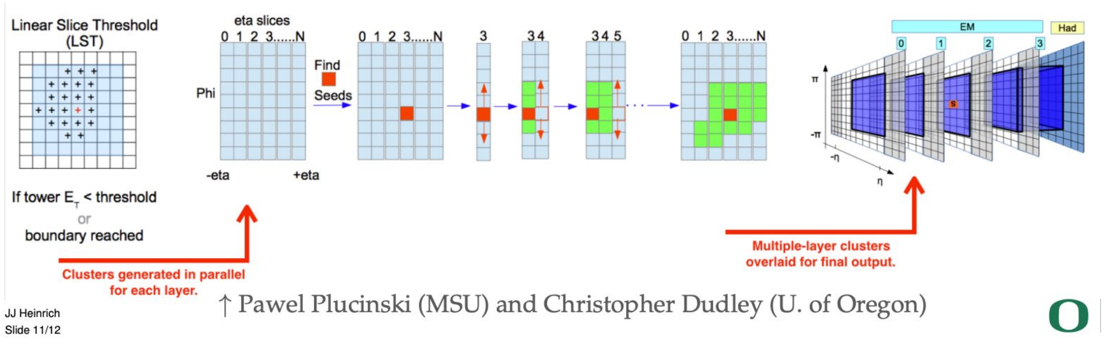

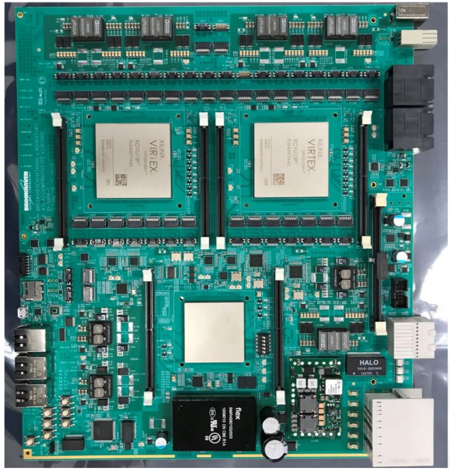

CPAD Trigger/DAQ Mar 22, 2021

Talk by Audrey Therrien

Highlights: **The Challenges of Machine Learning at the Edge**

Why do we need Real-Time ML ?

New generation sensors and detectors have :

- More pixels;
- Faster sampling rates;
- Better sensitivity and dynamic range;
- Larger surface area;
- Too much data!

Traditional programming

Input + Program = Output

Machine Learning

Input + Output = Program

Benefits of ML

- Recognizing patterns
- Recognizing anomalies
- Non linear regression → reconstruction
- Faster more flexible programming
- Lower computational burden
- Fast inference
  Low latency decision

CPAD Trigger/DAQ Mar 22, 2021

Talk by Jeffrey Krupa

Highlights: **SONIC: coprocessors as a service for accelerated inference of DL algorithms**

SONIC: Services for Optimized Network Inference on Coprocessors Framework for integrating GPUs and FPGAs as a service (aaS) into physics workflows

Case studies of integrating GPUs/FPGAs aaS into:
- LHC experiments: GPU, FPGA
- neutrino experiments: ProtoDUNE
- Gravitational waves: LIGO denoising

As-a-service paradigm introduces coprocessors to HEP with minimal changes to pre-existing computing workflows

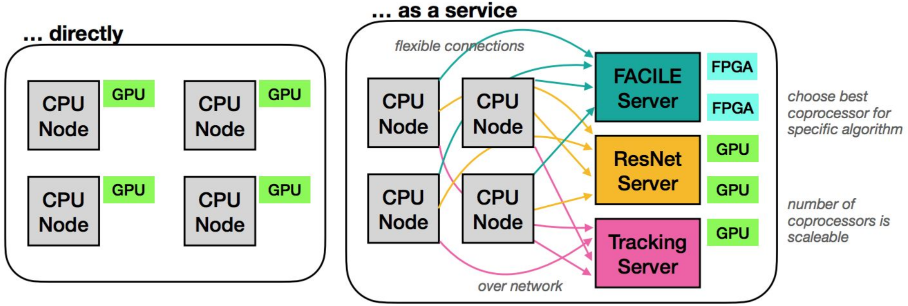

Demonstration of scaled CMS HLT speed-up with hadron calorimeter reconstruction (low- latency, high batch) performed on GPUs and FPGAs

2007.10359

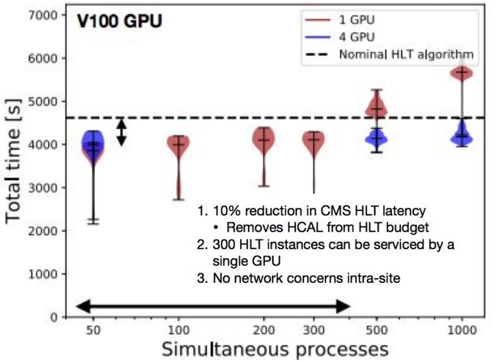

2010.08556

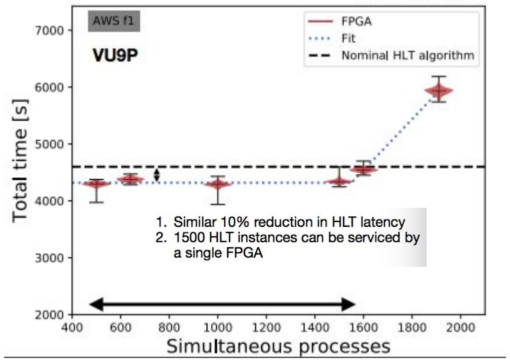

Mar 22, 2021 CPAD Trigger/DAQ

Talk by Martin Kwok

Highlights: **Triggering on Long-Lived Particles decaying to Hadronic Showers in CMS Muon System**

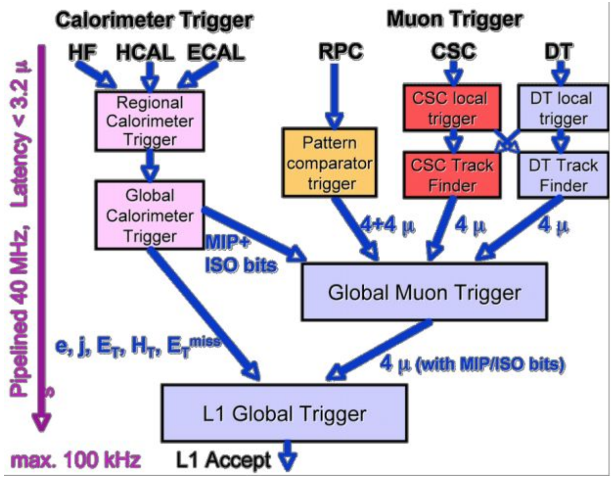

Development of new trigger path targeting long-lived particles in CMS, using the muon system as a sampling calorimeter.
Dedicated trigger for showers in CSC could extend CMS's reach for LLPs to >1m lifetime.
Can improve acceptance by 20-30x; on-track to be brought online in CMS Run 3.

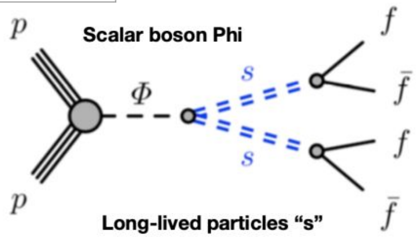

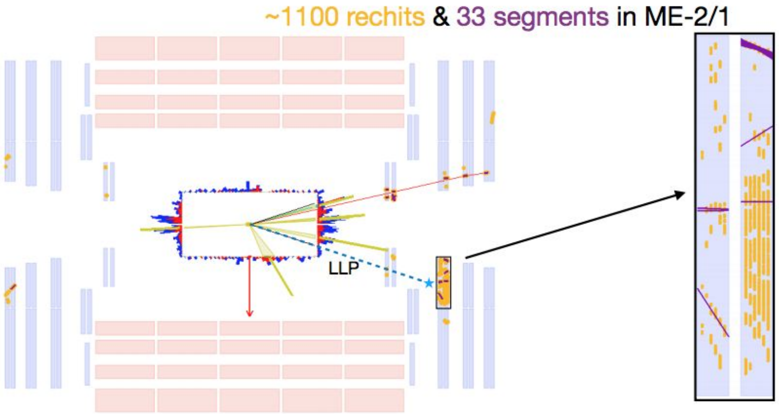

CPAD Trigger/DAQ   Mar 22, 2021

Talk by Jennifer Ngadiuba

Highlights: **hls4ml enabling real-time deep learning in particle physics**

- **hls4ml** is a library for automatic translation of deep learning models to FPGA firmware for inference with ultra low latency
- **First target applications:** hardware trigger of LHC experiments and detector front-end electronics

Applications for the hardware trigger at LHC experiments, e.g.
CMS Phase-2 L1 trigger upgrade:

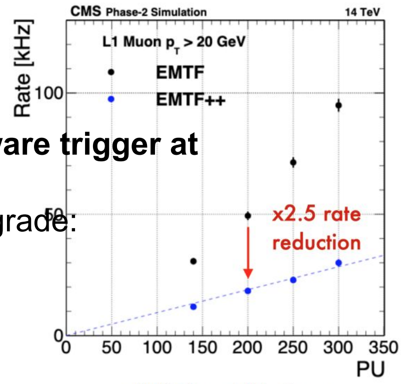

CMS Phase-2 L1 trigger upgrade TDR

Recent developments/library expansions:
Quantization-aware training and pruning [arXiv:2006.10159]
Convolutional neural networks [arXiv:2101.05108]
Custom architectures as graph neural networks [arXiv:2008.03601,arXiv:2012.01563]
+ more...

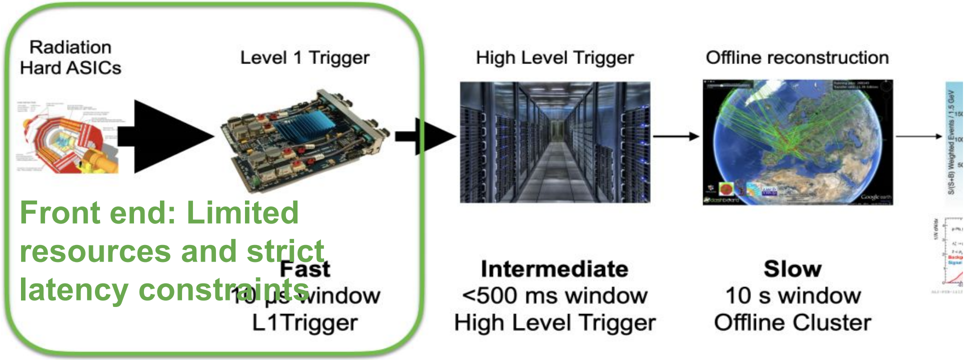

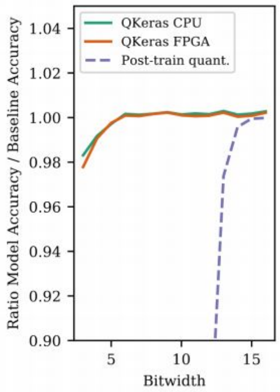

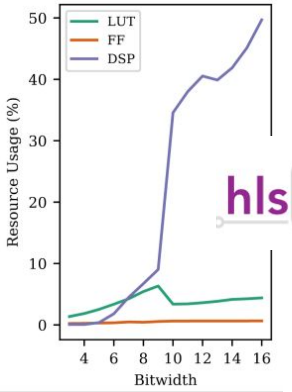

Talk by Christian Herwig

Highlights: **Real-time AI for Accelerator Control**

A study at the Fermilab Booster:
Megawatt proton beam with high proton per pulse density and minimal beam losses is required to meet DUNE requirements → use of ML regulator to enhance beam control [arXiv:2011.07371]

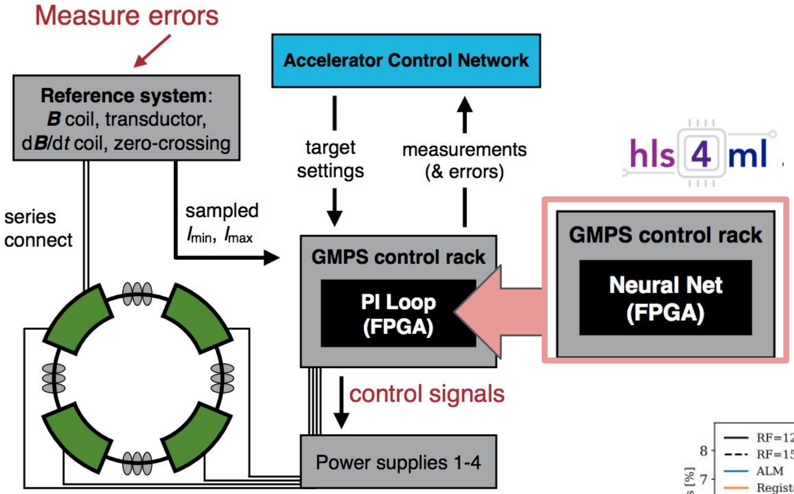

Can naturally incorporate many inputs. Offers potential for "live" adjustments to the algorithm parameters while in operation.

Inputs: current values for the five important signals

Parameter	Details [Units]
B:IMINER	Setting-error discrepancy at injection [A]
B:LINFRQ	60 Hz line frequency deviation [mHz]
B:VIMIN	Compensated minimum GMPS current [A]
I:IB	MI lower bend current [A]
I:MDAT40	MDAT measured MI current [A]

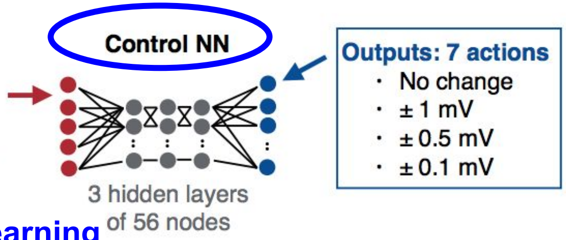

+ reinforcement learning

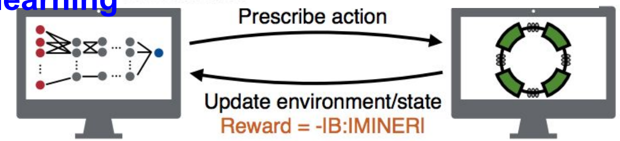

+ design footprint minimization

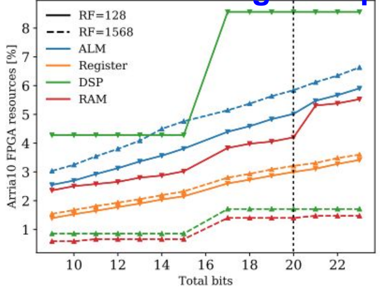

Comfortably fit within 6% of the target Arria10 FPGA's resources.

Can trade serial / parallel designs to trade resources for latency.

reuse factor	DSP	BRAM	MLAB	ALM	Register	Latency
128	130	114	229	21.4 k	51.2 k	2.8 μs
224	74	100	1420	40.2 k	78.3 k	4.1 μs
1568	26	38	357	24.9 k	54.9 k	17.2 μs
Available	1518	2713	···	427 k	1.7 M	···

Appears to perform better on historical booster data

Deployment planned for this spring for in situ demonstration

Final Thoughts

Mastering extreme environments and data rates in HEP experiments:
Future HEP detectors will involve extreme environments and exponential increases in data rates to explore elusive phenomena. ... To do so requires the intimate integration of intelligent computing with sensor technology.

[BRN Study Group, Identified Grand Challenges]

Many ongoing efforts reflect challenge and need for integrating intelligent computing into TDAQ systems.

BRN TDAQ PRD's:
21: Achieve on-detector, real-time, continuous data processing and transmission to reach the exascale
22: Develop technologies for autonomous detector systems
23: Develop timing distribution with picosecond synchronization