Challenges and R&D for DAQ in Particle Physics Experiment

Kai Chen

With input from many colleagues

Brookhaven National Laboratory

December 10, 2019

BROOKHAVEN NATIONAL LABORATORY

Typical Data Acquisition System

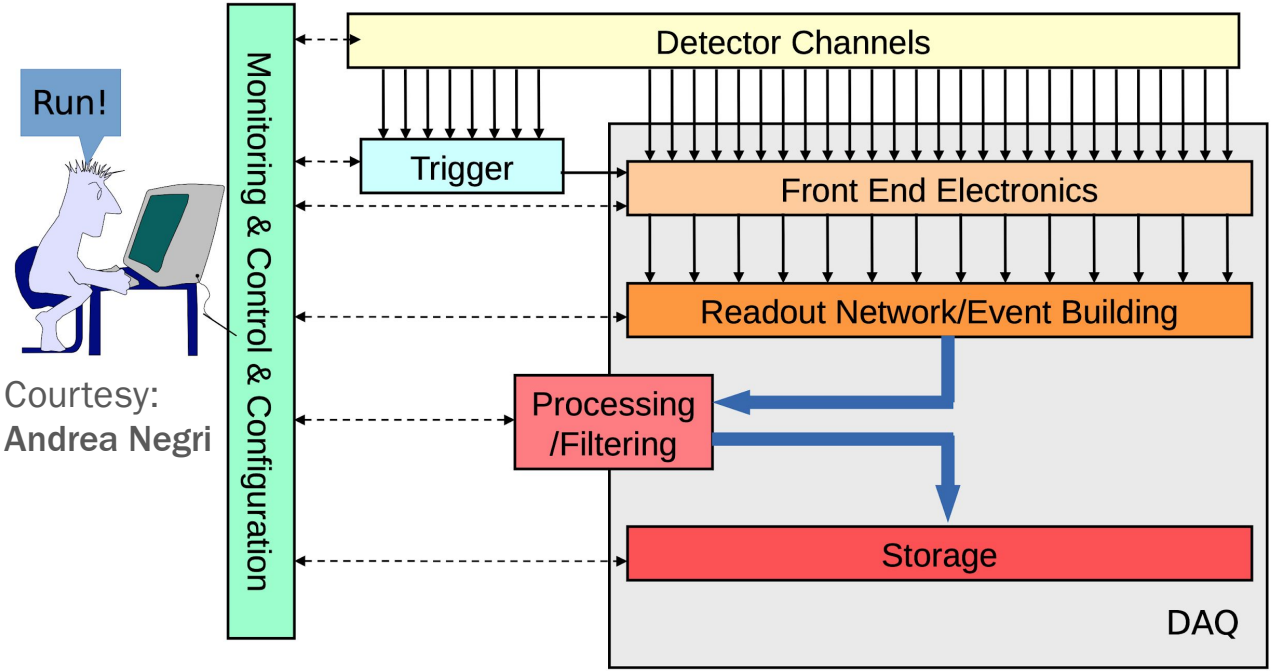

- Triggered readout
  - ATLAS
  - CMS
  - ALICE
- Streaming readout
  - LHCb (Run-3)
  - EIC (in R&D)
- Hybrid readout
  - sPHENIX
  - ProtoDUNE-SP
  - SBND
  - DUNE

Intensity Frontier

Energy Frontier

How Much Data is Generated?

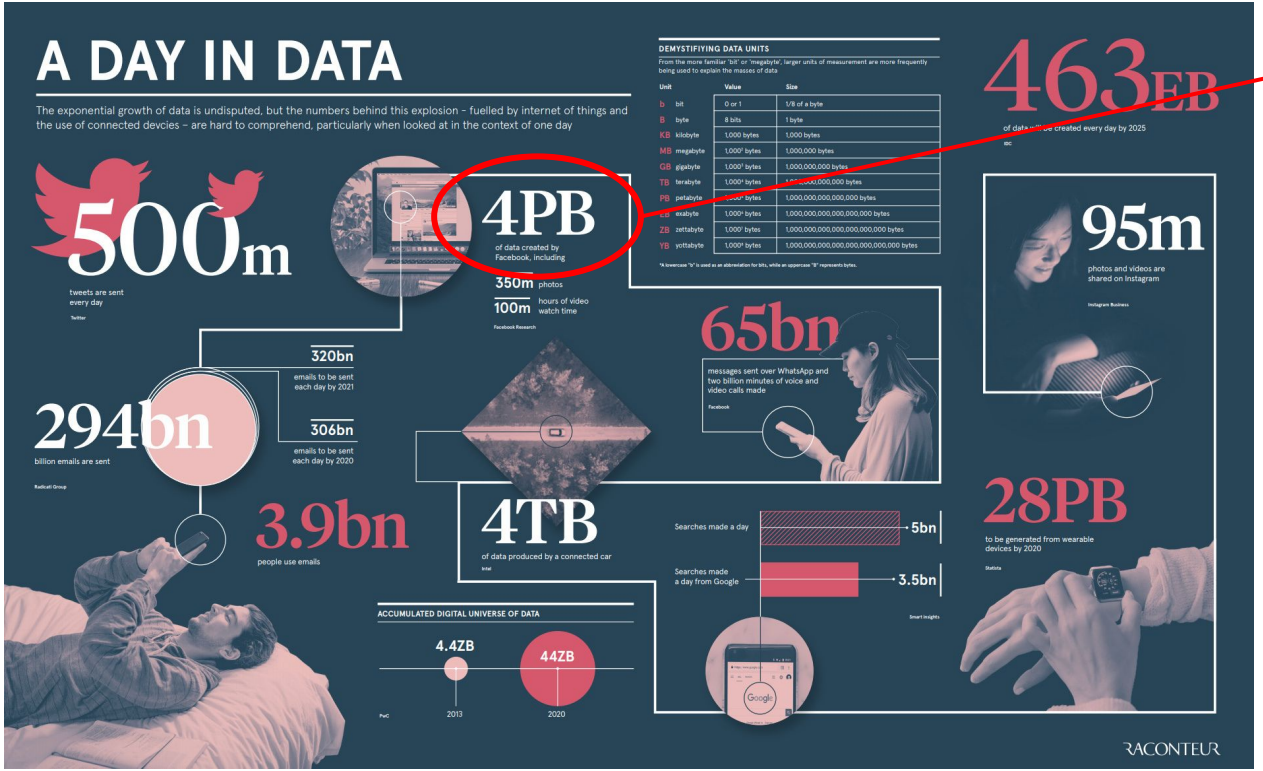

4PB/day for Facebook

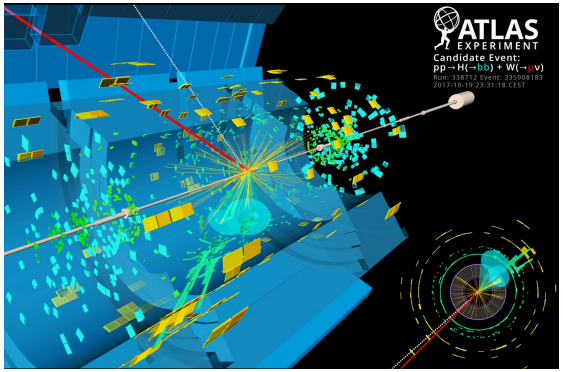

ATLAS raw data: ***~1PB/s***
after zero-suppression:
~0.5Pb/s

FCC-hh: ~10Pb/s

Image: **Raconteur**

Example 1: ATLAS DAQ in Run-2 (2015-2018)

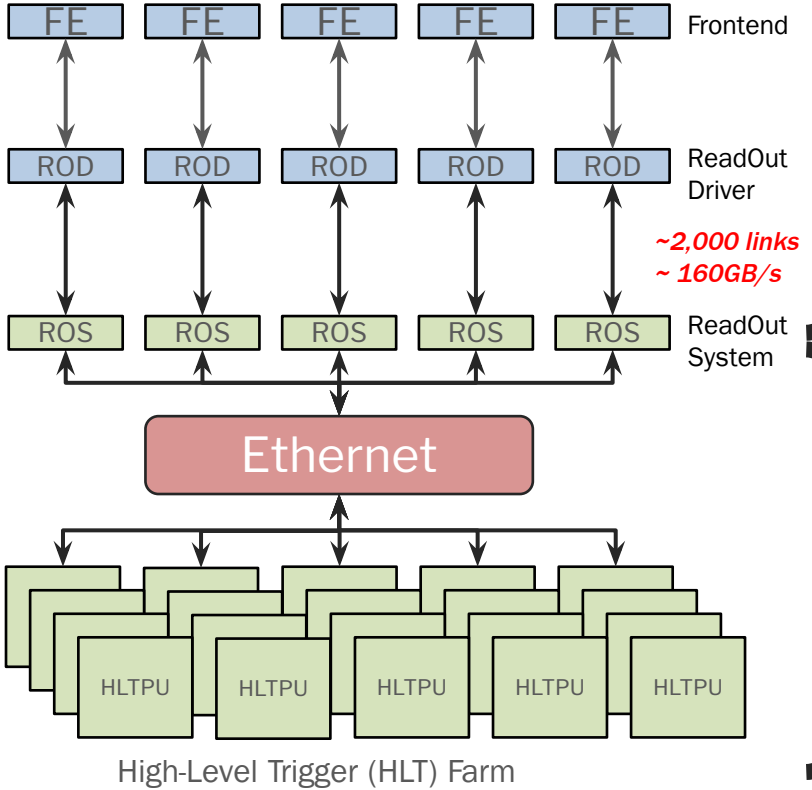

Custom electronic components

- ROD:
  - *Detector-specific custom hardware (mainly VMEbus)*
  - *Perform initial data processing and formatting*
- ROS:
  - *First common stage of DAQ system*
  - *Data buffered in custom PCIe I/O card (RobinNP)*
  - *Buffers and serves data fragments for HLT.*
- HLT:
  - *Uses full event tracking information*
  - *Performs more complete analysis of event*

PCs (COTs)

* COTs: Commercial off-the-shelf

ATLAS DAQ for Run-3 (2021-2023)

Moving common hardware nearer to detector. Exploit commodity electronics where possible.

- ***FELIX***
  - PCIe cards hosted in a server
  - Connect directly to detector front-end electronics or trigger hardware
  - Receive and route data from detector directly to clients over high bandwidth network
  - Distribute clock, L1 trigger and control signals to front-ends
  - Able to interface ASIC/FPGA with GBT protocol, or FPGA with high bandwidth 'FULL mode' protocol
- ***SW ROD***
  - Software process running on servers connected to FELIX via high bandwidth network
  - Common platform for data aggregation and processing – enabling detectors to insert their own processing software into data path
    - Previously performed in ROD hardware
  - Buffer data and serves it upon request to HLT
    - Interface indistinguishable from legacy readout (ROS
- ***Control and monitoring applications also now distributed among servers connected to data network***

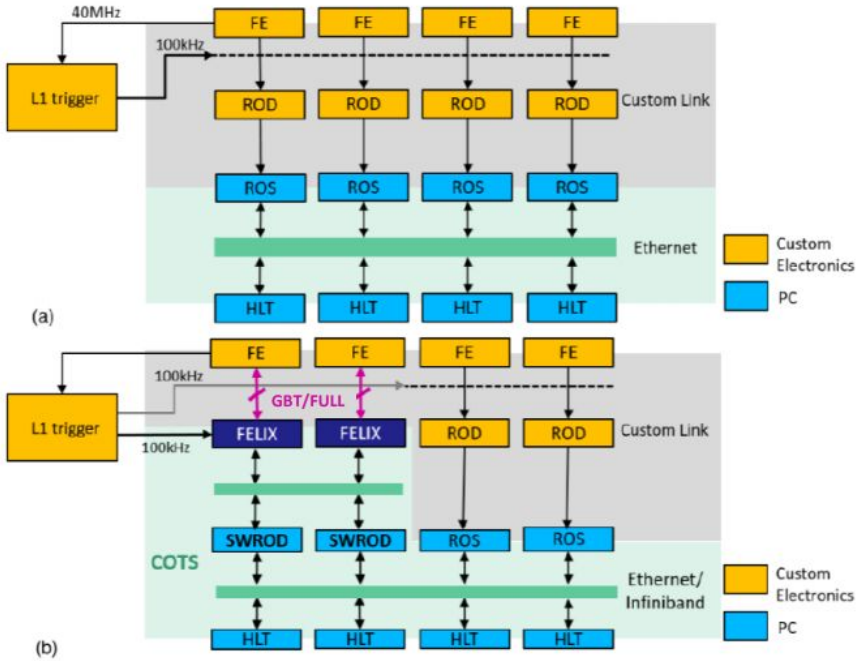

ATLAS DAQ in Run-4

Raw data per event:
~1.6 MB => ~5MB

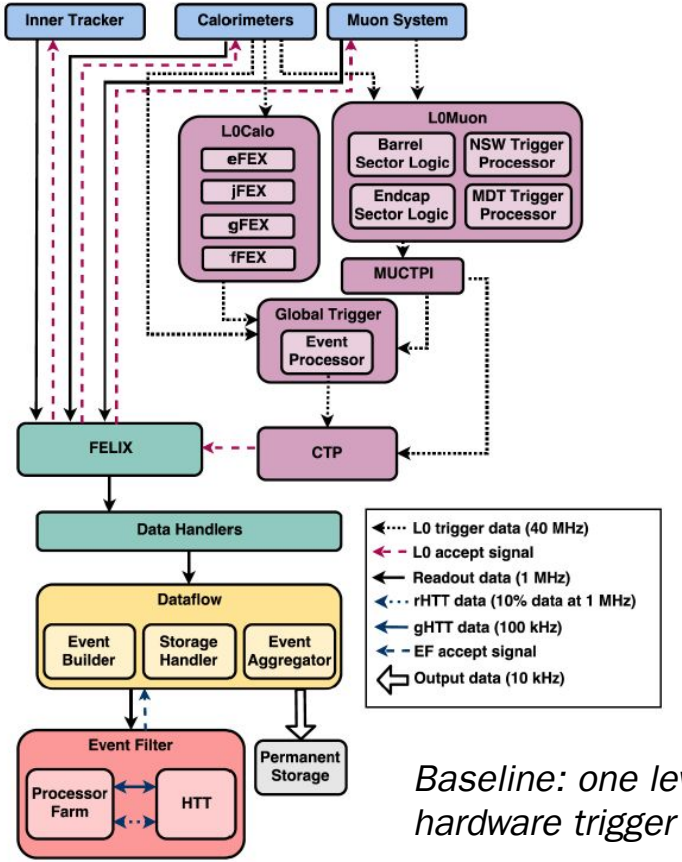

Baseline: one level hardware trigger

- Raw data from detector channels: ~Pb/s
- Billions of channels: 5,000,000,000 pixel channels
- FELIX for data readout with hardware trigger
  - ~20,000 fiber optical links
  - ~**10 Gbps radiational-hard** links with front-end
  - ***~42 Tb/s***
- ~480 Gb/s for storage.

Dataflow:
- ***Event Builder*** builds event records and manages the storage volume of the Storage Handler system
- ***Storage Handler*** buffers event data before and during processing by the Event Filter
- ***Event Aggregator*** collects, formats and transfers the output to CERN permanent storage

Example 2: ProtoDUNE-SP

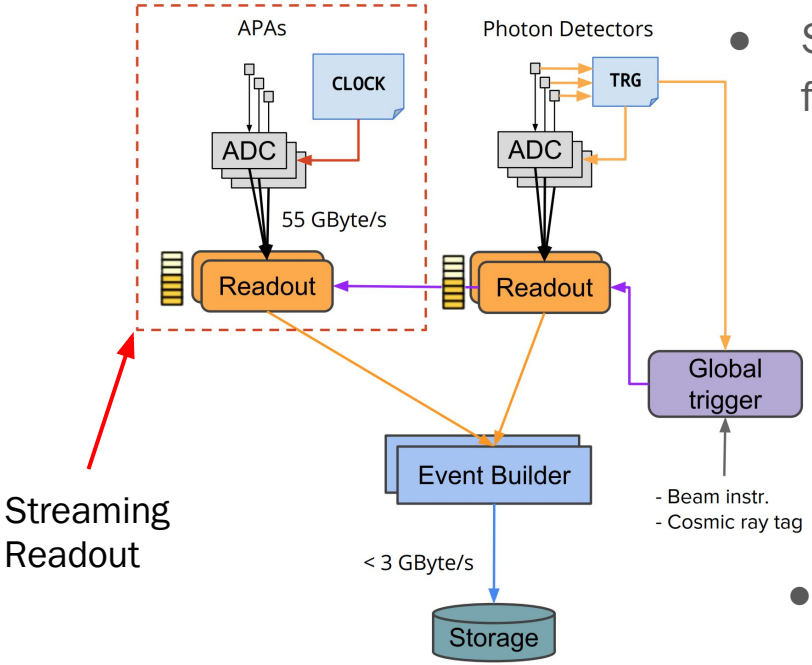

- Streaming readout for APAs
  - 2MHz sampling
  - **6** APAs
  - 15360 channels/APA
  - **~440Gb/s**

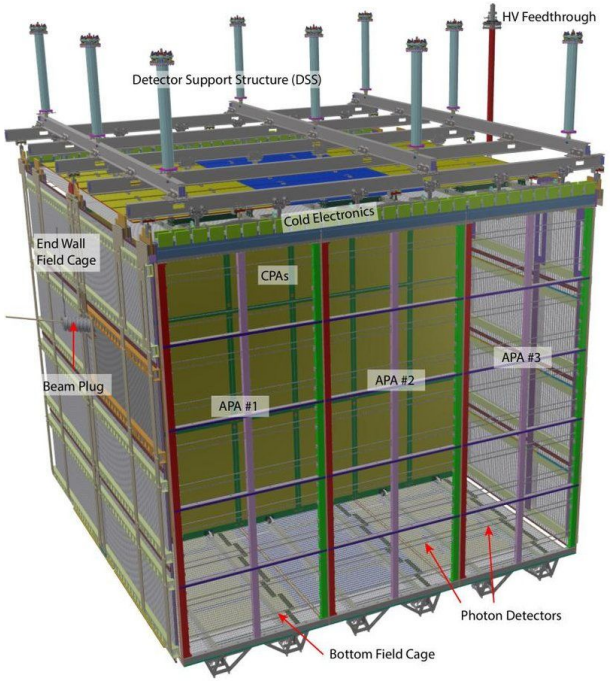

- For DUNE:
  - **150** APAs for one 10kTon module
  - **~12Tb/s**

Example 3: sPHENIX

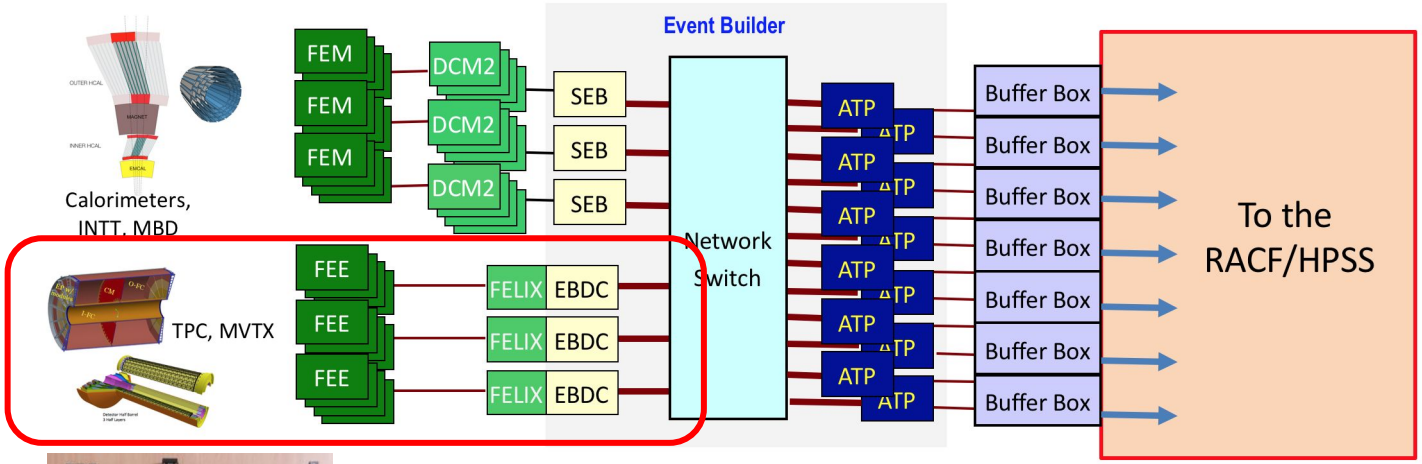

Hybrid with triggered calorimeters

Similar architecture applies to proposed EIC experiments, Ref: Streaming Workshop: https://indico.bnl.gov/event/6383/

MVTX RU, 200M ch
ALPIDE

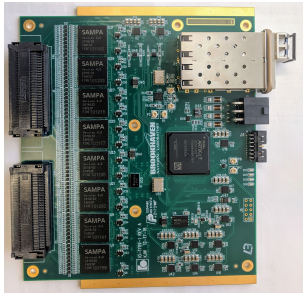

TPC FEE (trigger-less), 160k ch
SAMPA with **zero suppression**
~5.8 Tb/s

FELIX FLX-712, 40 links are used per card

Example 4: PUMA Experiment in Cosmology Frontier

Packed **U**ltra-wideband **M**apping **A**rray

See Paul's talk on Tuesday

- A next-generation cosmic survey using **intensity mapping** of the 21-cm emission from neutral hydrogen
- Interferometric array of 32,000 six-meter dishes closely packed
- Dual-polarization feeds, compact on-antenna electronics
- Redshift range 0.3 < z < 6 corresponding to 1100 < ν < 200 MHz
- Primary science goals:
  - Probing physics of dark energy in the pre-acceleration era
  - Searching for signatures of inflation
  - Probing the transient radio sky (fast radio bursts and pulsars)

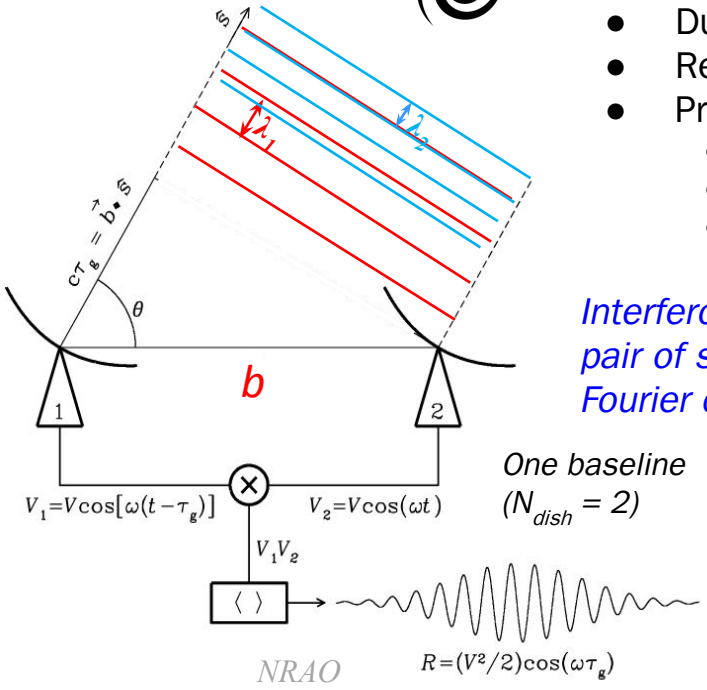

One baseline (N_{dish} = 2)

Interferometer measures the sky image directly in the Fourier space. Every pair of stations provides a baseline, measure a 'Visibility', which is a Fourier component of the image.

Challenges for DAQ

Jitter, phase calibration and stability of the clock distribution are critical.

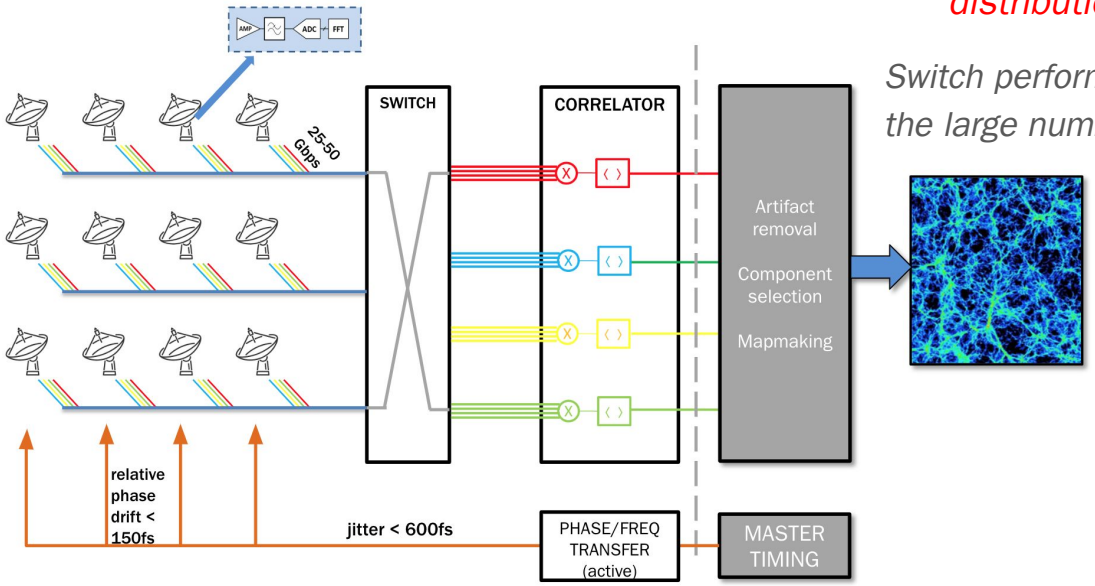

Switch performs frequency de-multiplexing data stream from the large number of antennas.

32,000 dishes (1500m diameter)

- **1.5 Pb/s** to the SWITCH
- Need **100PFlops** for correlator:
  - **~700kW** power consumption
- The digital and analog functions on the antennas need another **~1MW**
- **40 Gb/s** to tape

Joint R&D is ongoing at BNL for the readout

LDRD: *Experimental Cosmology with 21cm Hydrogen Intensity Mapping*

LDRD: *High-Throughput Advanced Data Acquisition for eRHIC, Particle Physics and Cosmology Experiments*

BROOKHAVEN NATIONAL LABORATORY

R&D for Detectors

- **Radiation hard for energy frontier experiments.**
- **Higher spatial granularity;**
  - LAr based **EM calorimeter** for FCC-hh:
    - Increased granularity in noble liquid calorimeters with fine segmented readout electrodes: $\Delta\eta x\Delta\varphi \approx 0.01x0.01$ (4x better for the 2nd layer), 8 longitudinal layers.
    - Increasing signal density of feedthroughs to ~ 50/cm^2 which is a factor ~5-10 more than in ATLAS

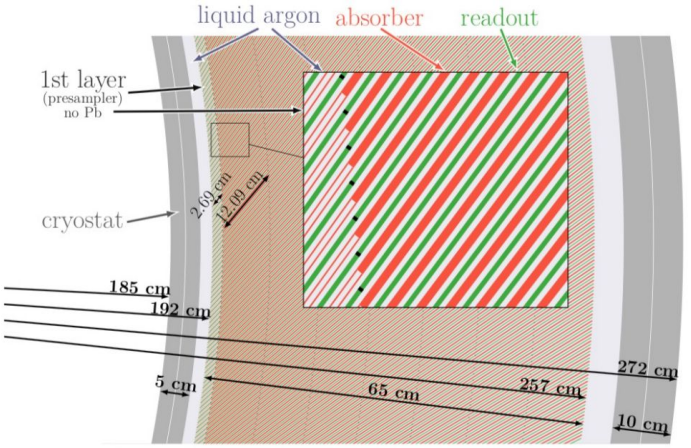

Courtesy: **Martin Aleksa**

- **Faster pixel detector: HV/HR-CMOS to realize large depleted area & high charge collection efficiency.**
- **Better energy, timing resolution, examples:**
  - HGTD: new pixelated silicon detector in the end-cap for ATLAS, to provide timing information (~30ps) for 4-D reconstruction and pile-up contamination reduction (factor of 6).
- **Low power consumption and low noise in Cryogenic environment**
  - Examples: **LArTPC**
    - Wire based APA (Anode Plane Assembly) => PCB based APA
    - Pixelated Anode with charge readout: LArPix, QPix

BROOKHAVEN NATIONAL LABORATORY

R&D for LArTPC Readout

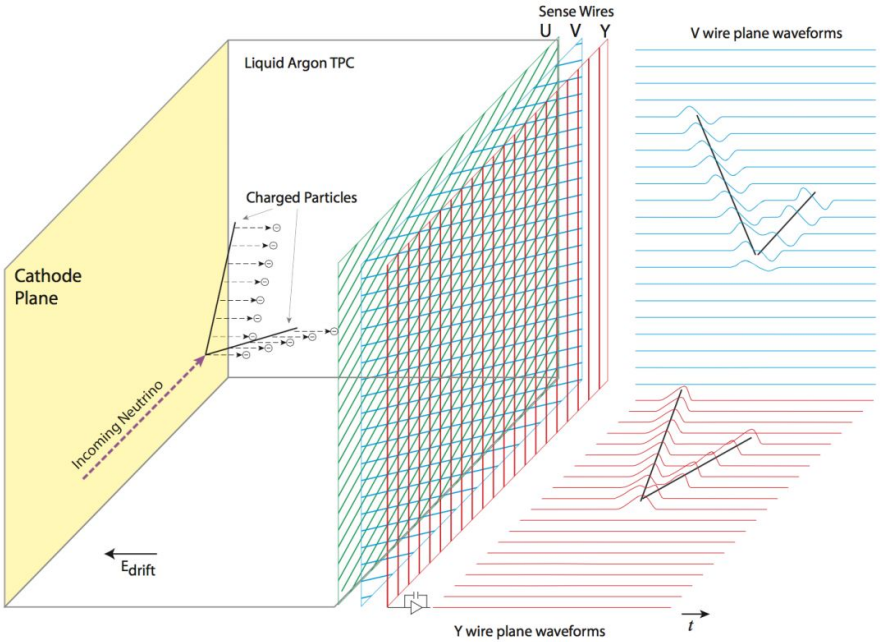

The MicroBooNE detector schematics

- U, V: induction planes
- Y: collection plane
- 3 mm pitch in all plane for wires

- **MicroBooNE**: CMOS Analog Front-End ASIC in LAr (PA+Sh+Drv); cold cable transfer analog signal to warm electronics for digitization;
- **SBND/ProtoDUNE**: digitization is in front-end cold electronics

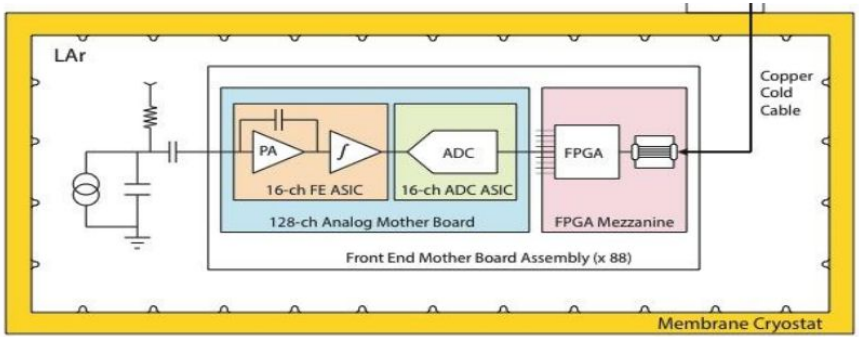

Advantages of cold electronics:

- Better SNR:
  - *Closer to wire electrodes;*
  - *Lower noise in LN_2*
- Reduce the number of cryostat penetrations

R&D for DUNE Detector

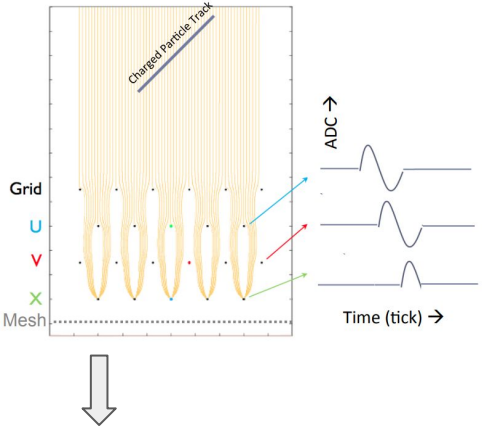

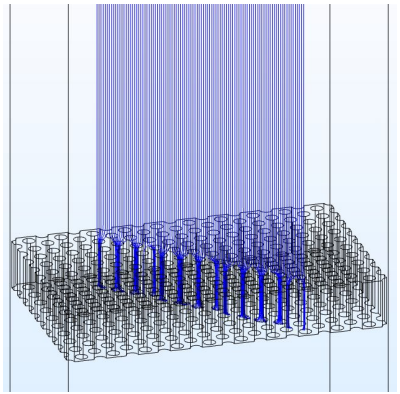

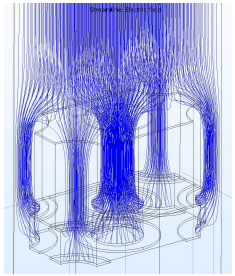

Electron Paths through the PCB Holes

Fragile wires are replaced by **robust copper strips**:

- Robust and easy to maintain the wire pitch and uniformity
- Easy for mass production, scale-up and modulation
- Strips in the front (screen plane) and intermediate layers (induction plane) sense induction signal
- Strips on last plane (collection plane) collect the ionization electrons.
- 3mm pitch of the readout strips

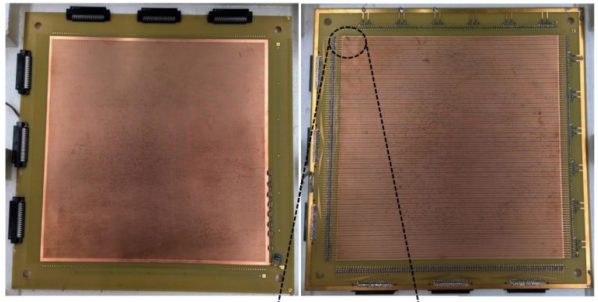

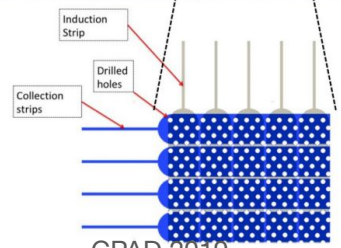

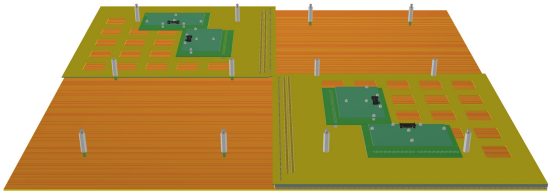

Integrating the FE electronics (FEMB of ProtoDUNE) on the PCB

The noise of LArTPC on the sensitive wire/pad are capacitive noise, the THGEM structure is equivalent to a parallel plate capacitor that would increase the noise.

More details: Bo Yu's slides

Data Transmission and Compression in Front-End

Data transmission: higher bandwidth, radiation hard, lower mass, lower power consumption

Electrical links between front-end ASIC and high-speed transmitter

- For RD53, ATLAS: up to 6 m @ 1.28 Gbps;

High-speed fiber optical links:

- R&D towards 28G/56G

Wireless transmission:

- R&D by groups like WADAPT for tracking detector: 60G band and 240G carrier have been demonstrated.
- Data rate 1/10 carrier frequency (OOK, BPSK)

Front-end ASIC/electronics to reduce the transmitted data volume:

- Self-triggering for analog circuitry
- Data compression in digital domain: ALPIDE, SAMPA
- On-detector intelligence

High Speed Links Development @ CERN

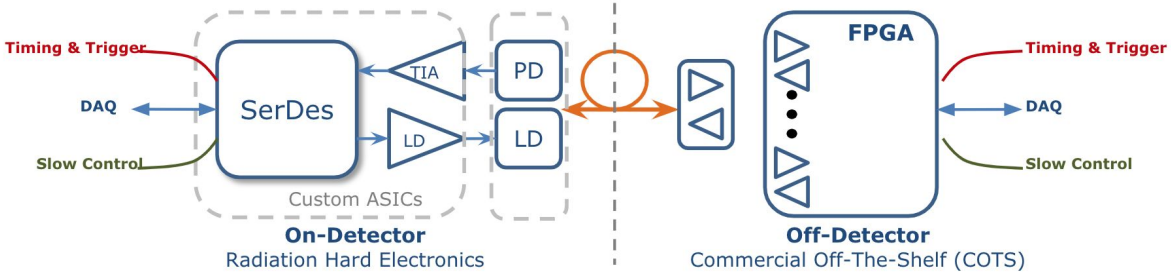

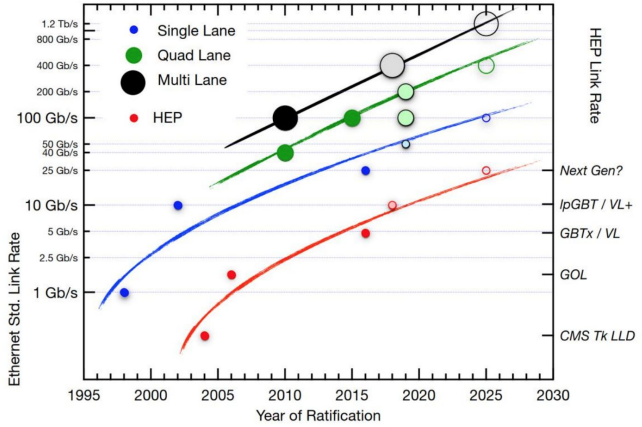

GOL: Gigabit Optical Link
GBT: GigaBit Transceiver
LpGBT: Low-Power GBT

	CMOS	Power Consumption	Link Speed	TID
GOL for LHC (e.g. ALICE)	250 nm	400 mW/chip	Uplink: 1.6Gb/s	~10Mrad
GBT for LHC Run-3	130 nm (1.5V)	980 mW/chip	Bidirectional: 4.8Gb/s	~100Mrad
LpGBT for HL-LHC Run-4	65 nm (1.2V)	500 mW (5.12 Gbps) 750 mW (10.24 Gbps)	Uplink: 5.12/10.24 Gb/s Downlink: 2.56 Gb/s	~200Mrad

High Speed Links Development @ CERN

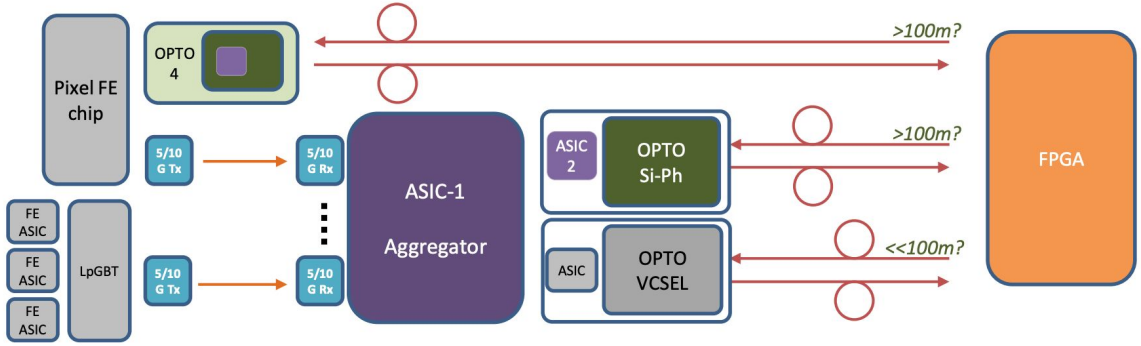

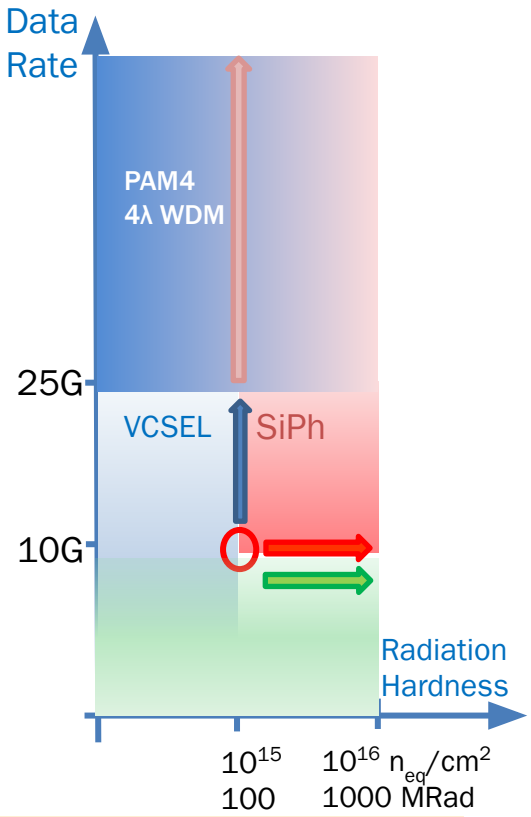

Activity	Task	Description
ASICs	ASIC-1	Very high data rate aggregator/transmitter
	ASIC-2	Optoelectronics drivers
	ASIC-3	Low-mass electrical cable transmission (active cable)
FPGA	FPGA-1	FPGA-based system testing and emulation
OPTO	OPTO-1 & 2	Silicon Photonics System & Chip Design
		Silicon Photonics Radiation Hardness
	OPTO-3	Next-generation VCSEL-based optical link
	OPTO-4	Silicon Photonics packaging

"Down scaled" Activity

"Dropped" Activities

28Gbps NRZ / 56Gbps PAM4 Transmitter with **28nm** CMOS

Si-Photonics: integration of optoelectronic devices in a "Photonic Si chip", by using WDM: 40Gbps NRZ is possible. **Mach-Zehnder Modulator** is also insensitive to NIEL.

Sources from **CERN EP Department**

High Bandwidth COTS Solutions for Back-End

- Xilinx: **112Gb/s PAM4** SERDES will be supported by Versal devices; Intel has also demonstrated 112Gb/s PAM-4 Transceiver I/O
- 200G/400G will be available for single lane with **coherent optical transmission**
- **Terabit Ethernet**: 800 Gbit/s and 1.6 Tbit/s may become IEEE standard in 2025
- **PCIe Gen 6** with PAM4: 128GB/s per 16 lanes (specification to land in 2021); Extended PCIe like CCIX >200GB/s is possible.

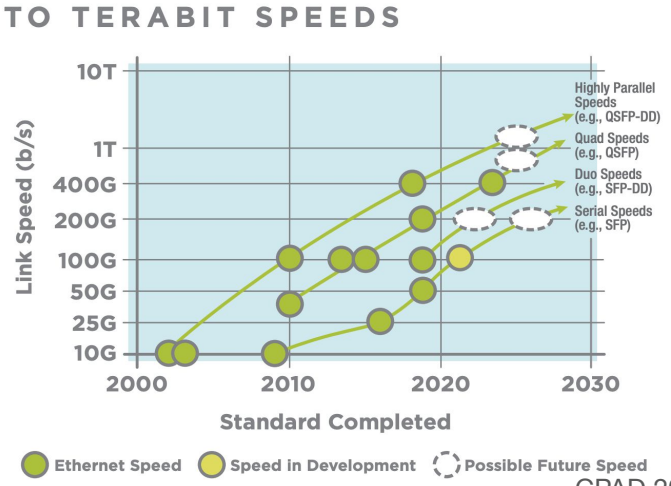

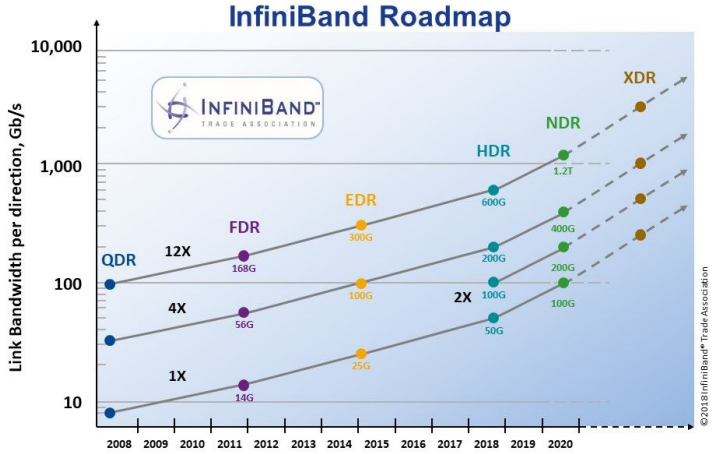

Summary

- The R&D for future detector will improve spatial granularity, energy and timing resolution. This will remarkably increase the raw data volume to be handled.
- Fast development of industry solutions provide inputs for R&D of the DAQ system.
- R&D of the DAQ system should be integrated with Front-End readout electronics, to meet the overall bandwidth requirement, power and space limit, and allow global optimization.
- For experiments with huge data like ATLAS, FCC-hh (~10Pb/s, needs a few 100,000 links). To make the streaming readout (triggerless) be possible, R&D should also be carried out in the detector side.
  - *Wireless transmission*
  - *Radiation hard high-speed serializer and optoelectronics*
  - *Data compression*
  - *On-detector intelligence*
  - *Self-triggering*
- ***Most R&D directions will need collaboration internationally and coordination globally to leverage the current knowledge base, development experience and available expertise.***

Thanks for your attention!

Intent-Based Network Functions

Image: **Cisco**

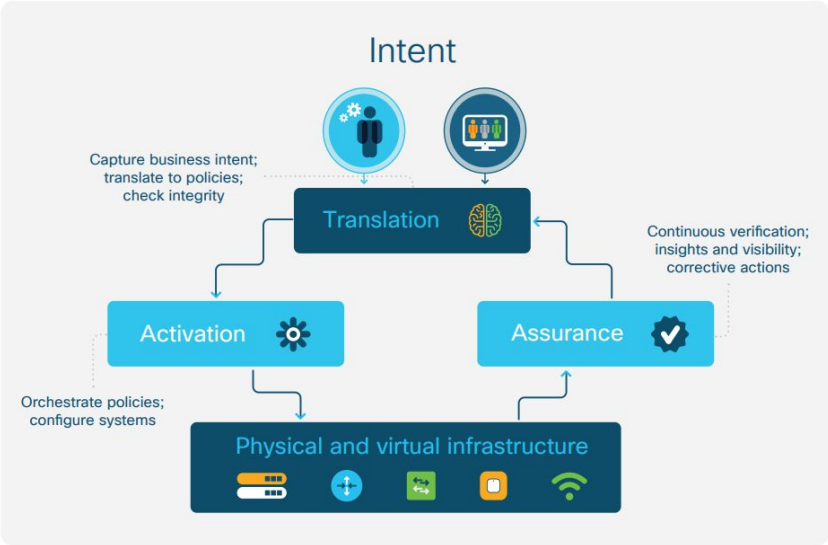

Intent-Based Network: use AI (Artificial Intelligence), ML (Machine Learning), MR (Machine Reasoning) to **automate administrative tasks across a network**.

- *Fast detection and response to faults*
- *Improve data collection efficiency*

Translation: enables network operators to express intent in a declarative and flexible manner, expressing what the expected networking behavior is that will best support the business objectives, rather than how the network elements should be configured to achieve that outcome.

Activation: installs these interpreted policies from captured intent into the physical and virtual network infrastructure using network-wide automation.

Assurance: maintains a continuous validation-and-verification loop, to continuously check that the expressed intent is honored by the network at any point in time.

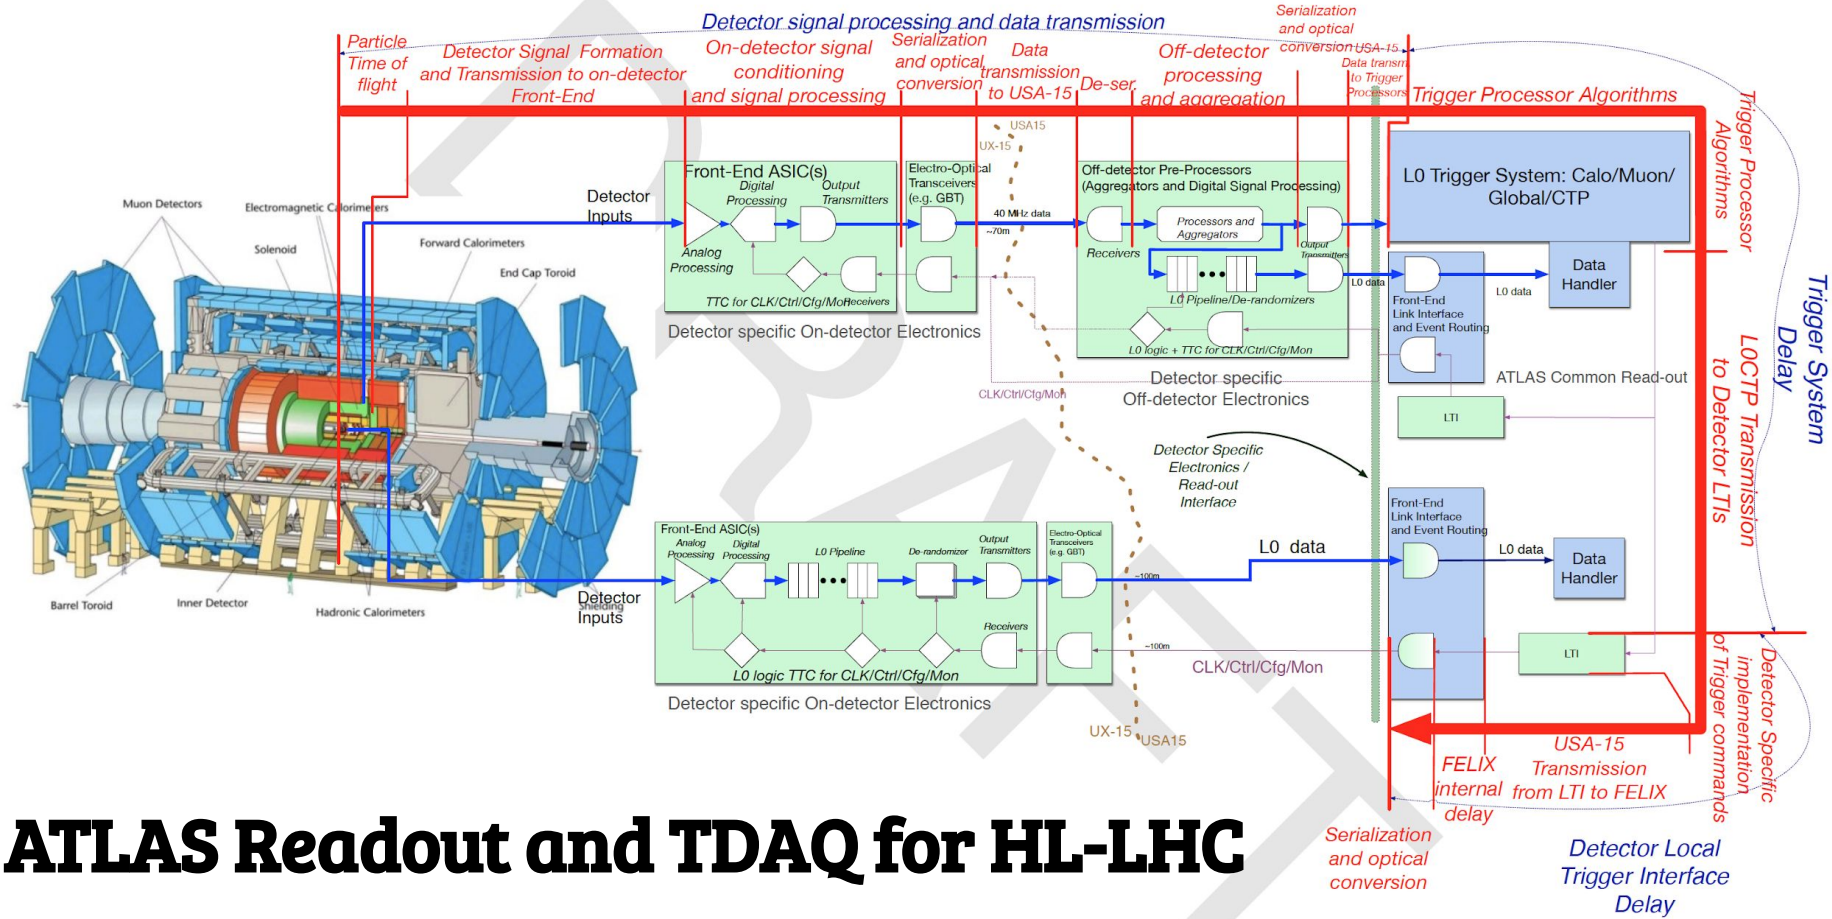

ATLAS Readout and TDAQ for HL-LHC

Network Implementation & Dataflow for ATLAS

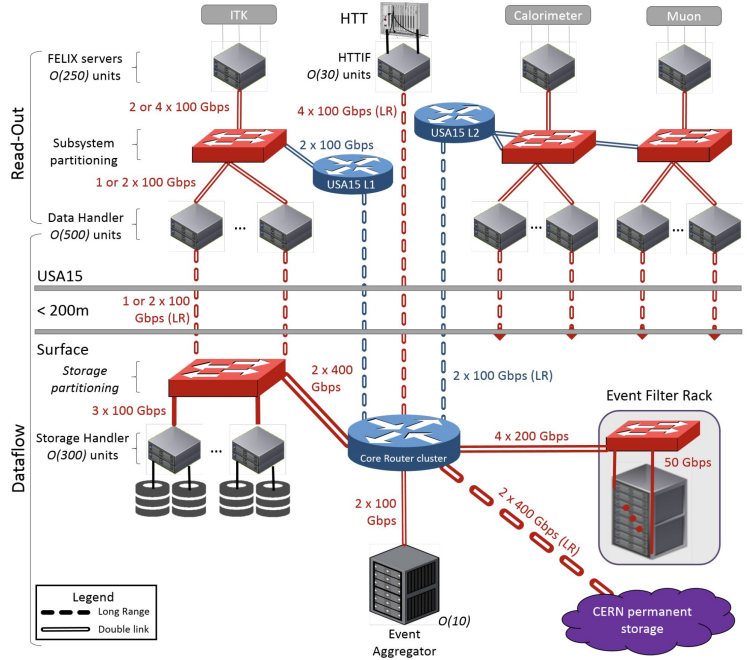

Possible network implementation using Ethernet technology

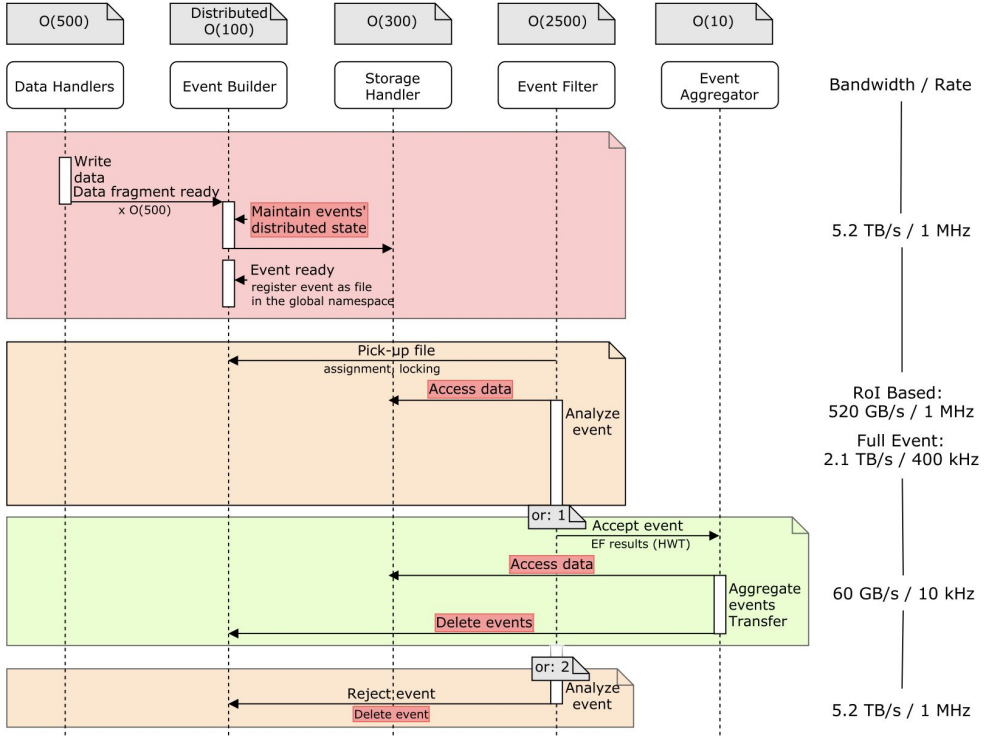

Logical communications between different components of the Dataflow system

R&D: 3D LArTPCs

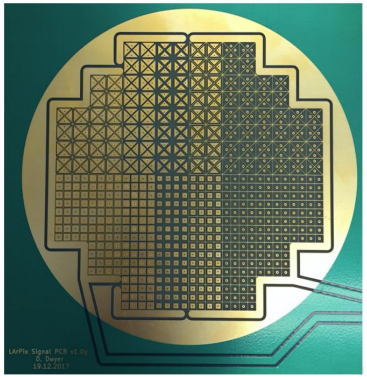

LArPix-v1: 10 cm diameter, 3mm pitch, 832 pixels

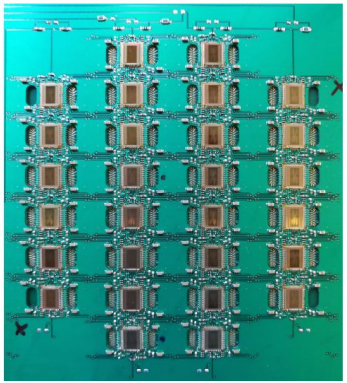

photograph of the back side of the readout assembly

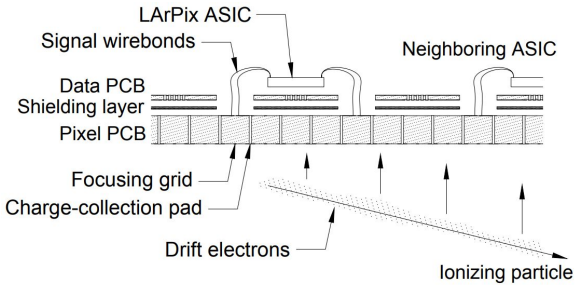

- Pixelated TPCs have been used by ALICE, sPHENIX: gaseous TPC
- For LArTPC, challenges are:
  - Cryogenic operation - cold ASIC for digitization & readout
  - Thermal constraints - low power consumption
- 32-ch LArPix-v1: 180 nm CMOS;
  - Low power: 62 uW/ch
  - Low noise: ~1.1mV (~275 e) in LN^2 bath
  - Self-triggering: avoid digitization and readout of mostly quiescent data

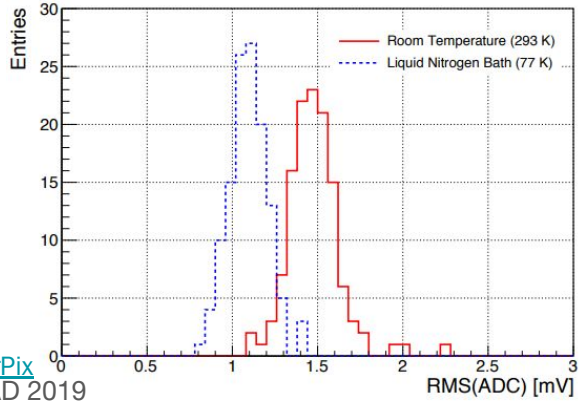

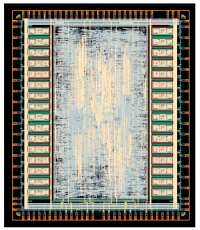

[Peter Madigan's talk on Sunday: Pixelated LArTPC R&D: LArPix]

BROOKHAVEN NATIONAL LABORATORY

LHCb for Run-3

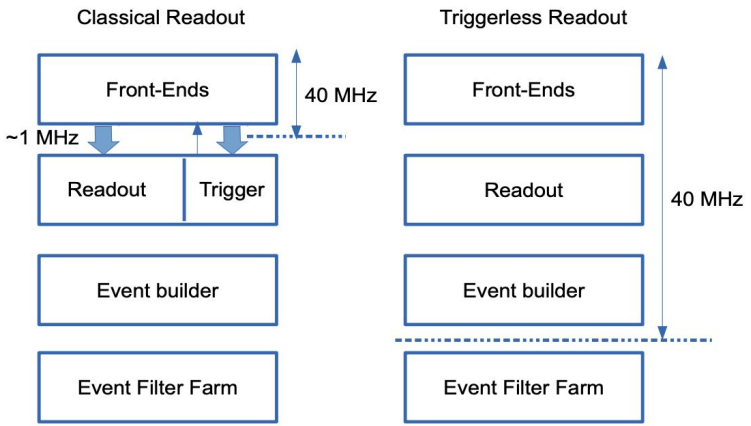

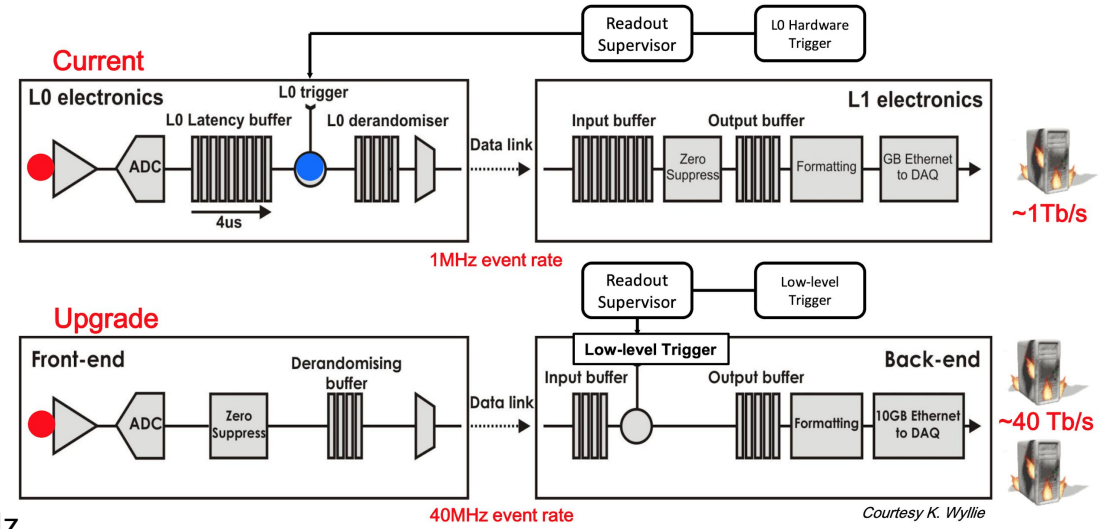

Triggerless readout, which can read out 40 MHz

- ~15000 optical links
- ~ 500 readout boards
- ~24 links in average on each board
- ~100 kbytes per event
- ~4 TB/s aggregate bandwidth

Data is compressed (zero-suppression) *in front-end ASIC (reduce # of 4.8 Gb/s links to about 1/6)*

HGTD for ATLAS

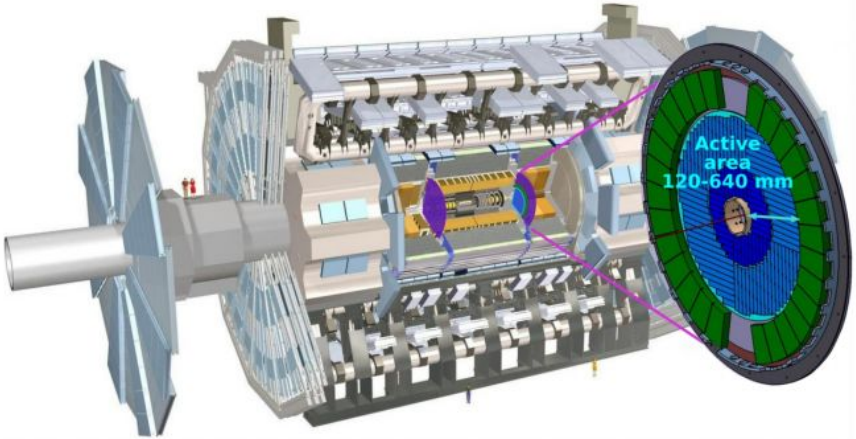

- Challenges for readout:
  - ~1200 up links for main readout
  - Clock dispersion less than 10 ps across a wide range of frequencies and over the detector acceptance.

HGTD to replace existing MBTS.

4-D reconstruction: reducing the pile-up contamination in tracks and vertexes:

- reducing the pile-up by factor of 6 if nominal Gaussian beam profile with 45mm or 175 ps spread in z-direction.

Pseudorapidity coverage	$2.4 < \|\eta\| < 4.0$
Thickness in z	75 mm (+50 mm moderator)
Position of active layers in z	3435 mm $< z <$ 3485 mm
Radial extension:	
Total	110 mm $< R <$ 1000 mm
Active area	120 mm $< R <$ 640 mm
Time resolution per track	30 ps
Number of hits per track:	
$2.4 < \|\eta\| < 3.1$	2
$3.1 < \|\eta\| < 4.0$	3
Pixel size	1.3×1.3 mm^2
Number of channels	3.54M
Active area	6.3 m^2