The DAQ systems of the DUNE Prototypes at CERN

Karol Hennessy on behalf of the DUNE collaboration

July 11, 2018

University of Liverpool

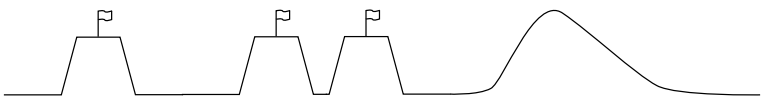

The Deep Underground Neutrino Experiment

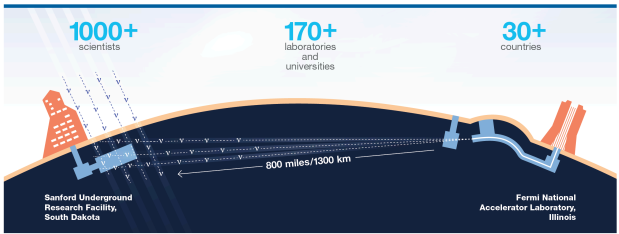

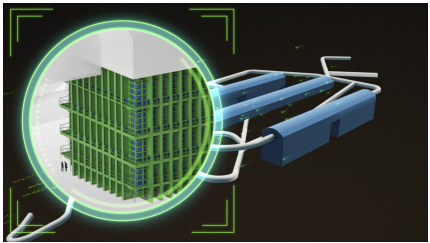

- Intense beam of ν_μ (or $\bar{\nu_\mu}$) fired 1300 km at a large detector
- Studying CP violation in the lepton sector, proton decay, supernovae
- Beam from Fermilab
  - Muon neutrinos from 1.2-2.3 MW proton beam
- ...to Far Detector
  - 4 caverns housing 10+ kt liquid argon (LAr) detectors

DUNE Far Detector

- Four 10-kt (fiducial) **liquid argon time projection chamber (LAr-TPC)** super-modules
- gives excellent 3D imaging and energy measurement capability
- **Single and dual phase detectors**
- Integrated photon readout

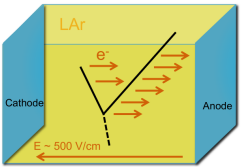

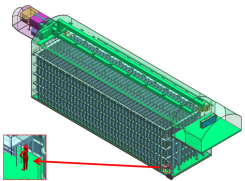

- 150 Anode Plane Assemblies (**APAs**) [2.3×6 m]
- 384,000 readout wires

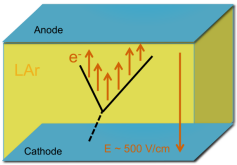

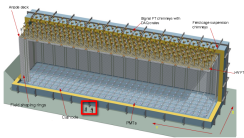

- signal amplification in gas phase
- 80 Charge Readout Planes (**CRPs**) [$3 \times 3\,\mathrm{m}^2$]

ProtoDUNEs

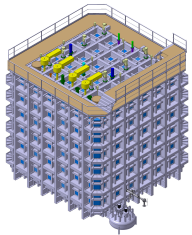

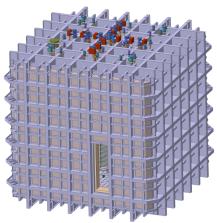

Single Phase

- Ionization signals (collection + induction) read out in liquid volume
- As used in ICARUS, ArgoNEUT/LArIAT, MicroBooNE
- Long-term operation/stability demonstrated by ICARUS T600

Dual Phase

- Ionization signals *amplified* and detected in gaseous argon above the liquid surface
- Being pioneered by the WA105 collaboration
- If demonstrated, potential advantages over single-phase approach

ProtoDUNEs

- Both at surface
  - lots of **cosmics** - dominant signal source (kHz)
- Goals
  - Demonstrate the viability of components and solutions for Far Detector modules
  - Learn procedures for construction, assembly, commissioning
- They look similar but...
  - Separate beam lines
  - **Very different electronics and readout strategies**

ProtoDUNEs

EHN1 - August 2016

EHN1 - June 2018

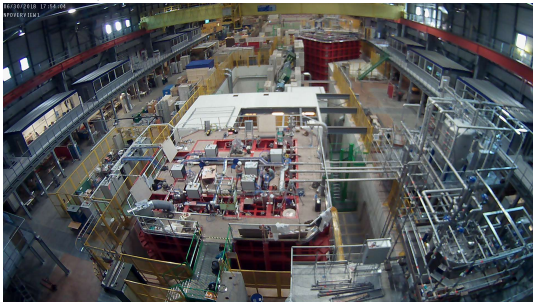

- Very rapid timescale!

ProtoDUNEs

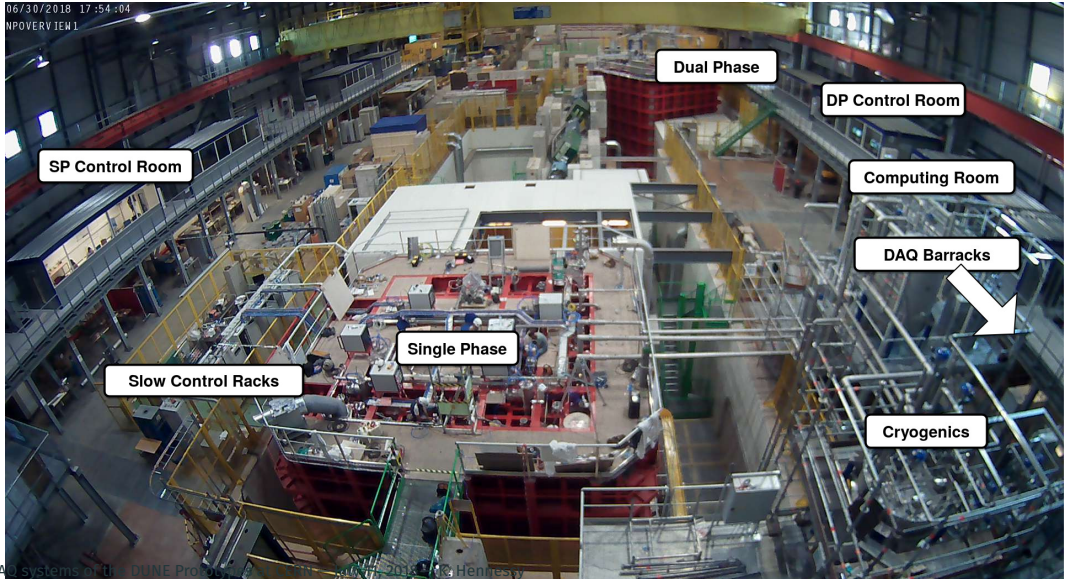
06/30/2018 17:54:04
NPOVERVIEW1
Dual Phase
DP Control Room
SP Control Room
Computing Room
DAQ Barracks
Single Phase
Slow Control Racks
Cryogenics

LAr-TPC Event Reconstruction - from a DAQ perspective

- Very different from collider physics
- Ionized charge drift is *slow* - 2.5 ms (7.5 ms dual-phase)
- Crudely speaking...
  - **LHC detectors take "photos"** to capture and event every 25 ns
  - **LAr-TPC detectors take "video"** - many snapshots to capture the 2.5 ms drift
- Consequence - event sizes are much larger, but there are fewer of them
- Direct effect on dataflow management, storage, etc.

		L0/L1 accept rate
ATLAS/CMS	1-2 MB	100 kHz
LHCb	∼50 kB	1 MHz
ProtoDUNE-SP	∼250 MB	25 - 100Hz
ProtoDUNE-DP	∼420* MB	25 - 100 Hz

*uncompressed

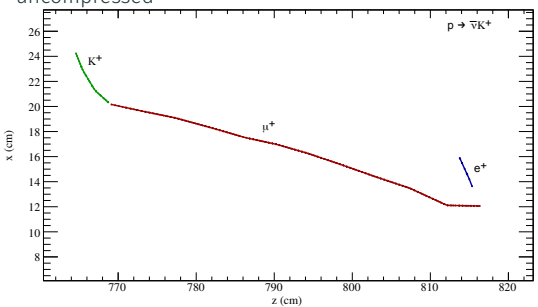

Simulated nucleon decay event $p \to \bar{\nu} K^+$, with a subsequent leptonic decay of the K^+. This event has been fully reconstructed, and the different tracks found by the automated reconstruction are shown by different colors.

Single Phase DAQ

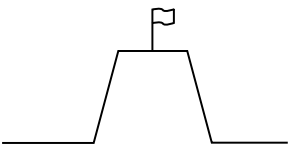

Single-Phase DAQ

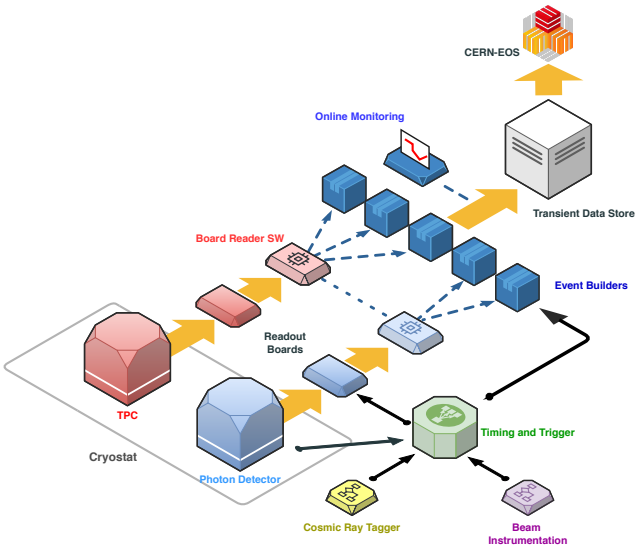

- DAQ is **optically isolated** from cryostat to minimise noise to front-end
- Data is **triggered and compressed** online in hardware and software prior to event building
- **Events sent offline at up to 20 Gb/s**

ProtoDUNE-SP DAQ Interfaces

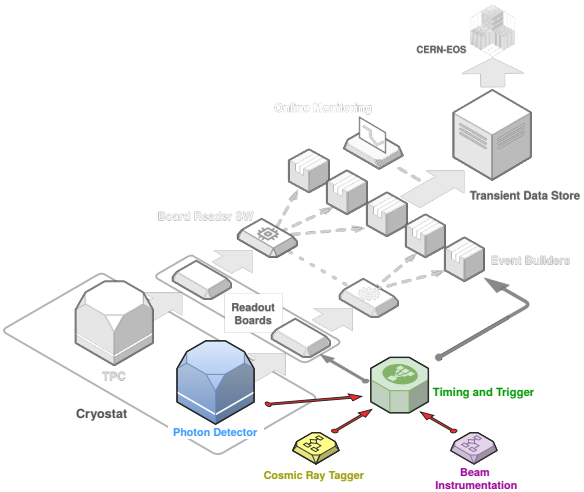

- Trigger Inputs

ProtoDUNE-SP DAQ Interfaces

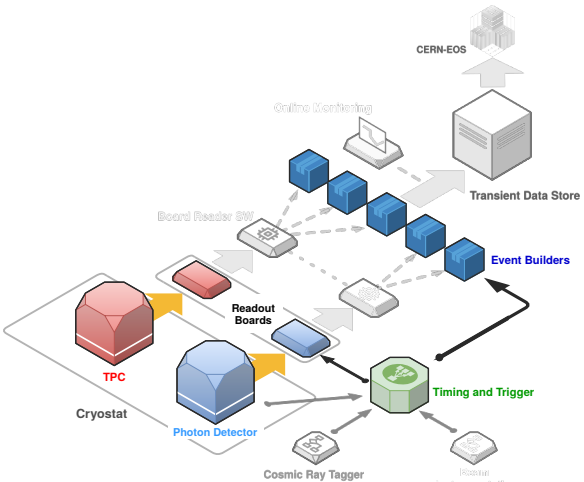

- Timing and Trigger to Front-Ends and Event Builders

ProtoDUNE-SP DAQ Interfaces

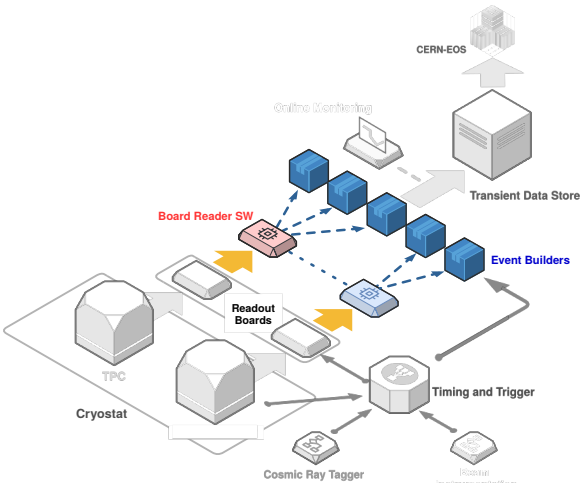

- Data fragments to Event Builders

ProtoDUNE-SP DAQ Interfaces

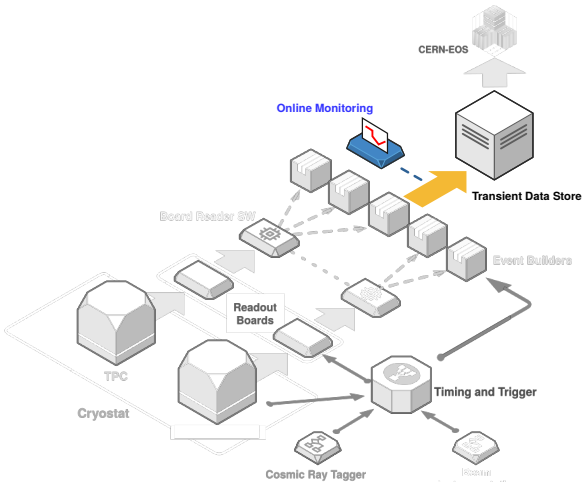

- Data to Temporary Storage

ProtoDUNE-SP DAQ Interfaces

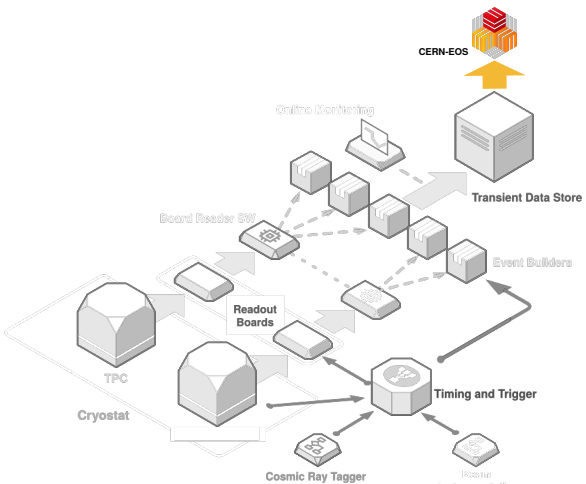

- Data to Offline

Front-end Electronics

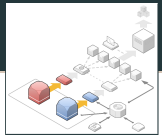

TPC Warm Interface Boards (WIBs)

- **Interface from cold electronics to DAQ** with shielding and local real-time diagnostics
- Source: ProtoDUNE Front-end Motherboards (FEMBs)
- Multiplex data from 4 FEMBs
- Output: Optical links to DAQ, towards readout systems and slow control

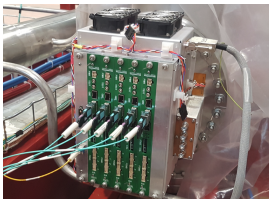

Photon Readout (SSP)

- The Silicon Photomultiplier Signal Processor (SSP) prototype module
- High-speed waveform digitizer
- Current sensitive, differential input amplifiers
- Good noise performance over long cables
- 12 channels per module

Timing system

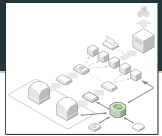

- Provides a **50 MHz clock** to all endpoints
- Multiplexed 8b10b encoded data stream
- Endpoint CDR circuit recovers clock and data
- Interface for CERN SPS spill signals (start of spill, etc.)
- Dedicated trigger interface - timing system data stream provides trigger distribution
- Interface to backpressure signals - trigger inhibit
- GPS driven, **64-bit timestamps**
  - provided to event builders for fragment matching
  - provide unambiguous event time in data irrespective of file, run, etc.
- **Partitionable** system

Trigger system

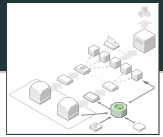

- Input from **Beam instrumentation, Photon Detector system, Cosmic Ray Tagger**
- Configurable at run time
  - multiple trigger levels, prescales, veto
- **Tightly coupled to timing system**
- Dedicated data stream for calibration

Central Trigger Board

- Based on Xilinx Zynq 7020 - MicroZed System-on-Chip
- 100 inputs - optical, NIM, TTL, LVDS...

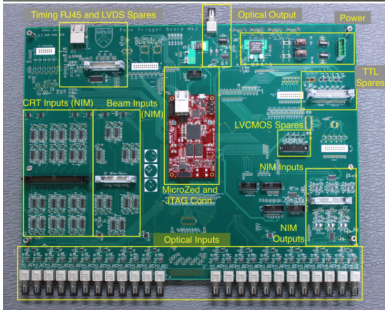

TPC Readout - RCE

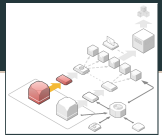

- Reconfigurable Cluster Element
  - **ATCA-based readout solution**
  - Used in several experiments - LSST, Heavy Photon Search, ATLAS Muon...
  - Custom Rear Transition Module (RTM) for experiment interfaces
    - 4× QSFP+ input transceivers for ProtoDUNE-SP
  - **High Level Synthesis** for C++ based algorithms on FPGA fabric

ProtoDUNE-SP

- TPC input data rate - 64 Gb/s
- Triggers and compresses data

The RCE Platform

- High density/high performance
- Nine clustered processing elements on a board
- Zynq 7045 SoC
- Dual core ARM A9
- 1 GB DDR3 RAM
- 10 Gb/s onboard switch

Poster on RCE from Ka Vang Tsang

TPC Readout - FELIX

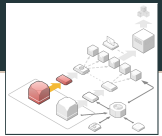

- **F**ront**E**nd **LI**nk e**X**change
- Designed for ATLAS LAr Calorimeter Phase-I Upgrade
- **PCIe based readout solution**
- ProtoDUNE-SP
  - readout 1 APA (one-sixth of TPC)
  - Software trigger selection
  - **Software compression** (can be accelerated with Intel QuickAssist (QAT) technology)

Detailed talk on FELIX from Enrico Gamberini - Thurs @ 11am

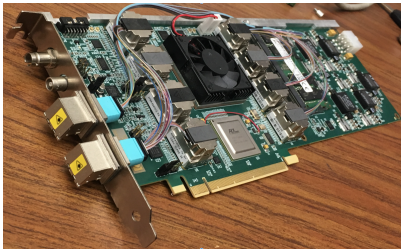

The FELIX Platform

- Card for ProtoDUNE-SP: FELIX BNL-711
- Xilinx Kintex Ultrascale
- 48 duplex optical links @ 14 Gb/s
- PCIe Gen3 x16 lanes ($\approx$ 100 Gb/s)
- Onboard DDR4 up to 16 GB
- GBT for front-end communication

Dataflow software - artDAQ

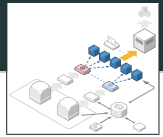

- artDAQ - data acquisition software toolkit
- **configuration delivery** and **data readout** of front-end hardware
- event building, data logging
- infrastructure for filtering, compression and online analysis
- infrastructure for real-time data quality monitoring

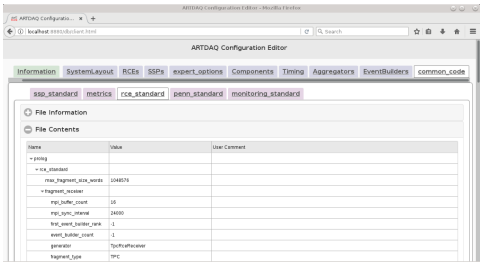

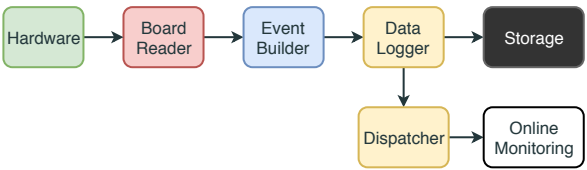

- configuration database for storing/retrieving HW&SW config
- Used on many experiments - DarkSide-50, LArIAT, mu2e, SBND, ICARUS

Run Control

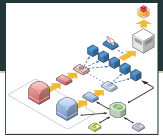

- Based on CERN's **JCOP** (Joint Controls Projects) extension to Siemens WinCC-OA framework
  - Used by all LHC experiments
- Interfaces with artDAQ
- Finite State Machine
- Partitionable system allows for parallel operation separate parts of the system

- **Primary front-end to data taking**
  - interacts with Run database backend
  - catalogs run information, and submits to logbook
  - essential for error information and diagnosis

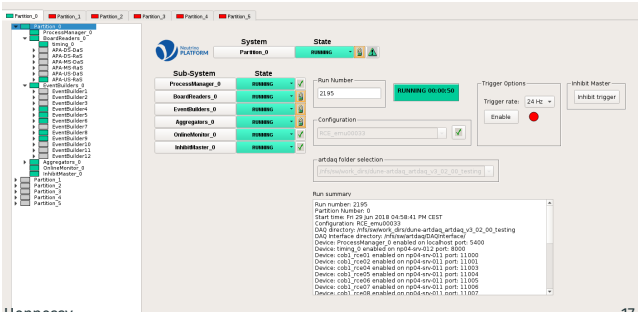

Operational Monitoring

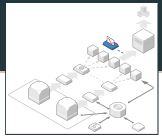

- **Metric reporting** from artDAQ components
- Display and archiving of trends for different quantities
- **Archived** to CERN central Oracle databases
- Shows:
  - data sizes, rates
  - fragment and event sizes and rates
  - individual or aggregated metrics

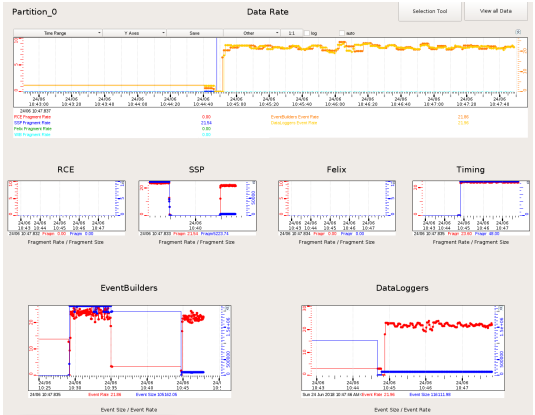

Online Monitoring

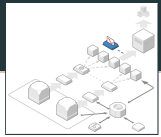

- Monitor detector performance during data taking
- Provide feedback for calibration
- Used to qualify the APAs in “cold box*” tests during installation
- Based on *art* and *LArSoft* physics software for LAr-TPC experiments
- DataLogger sends events to one or more dispatcher processes
- Dispatchers route events to the online monitoring processes

**test of electronics and front-end performance in liquid nitrogen bath*

- RAW decoders unpack the raw data to perform low-level analysis
- ArtAnalyzers perform high-level analysis of the unpacked events
- Histograms are saved and propagated to the **web display** - Monet

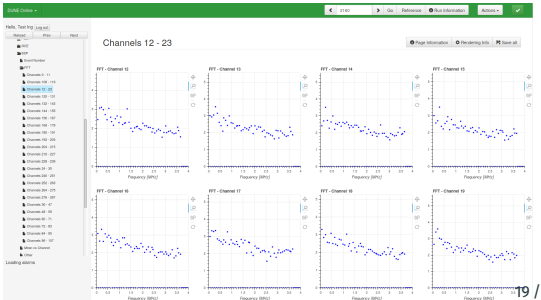

Dual Phase DAQ

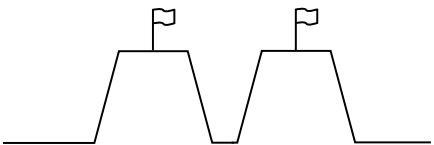

DAQ back-end network structure

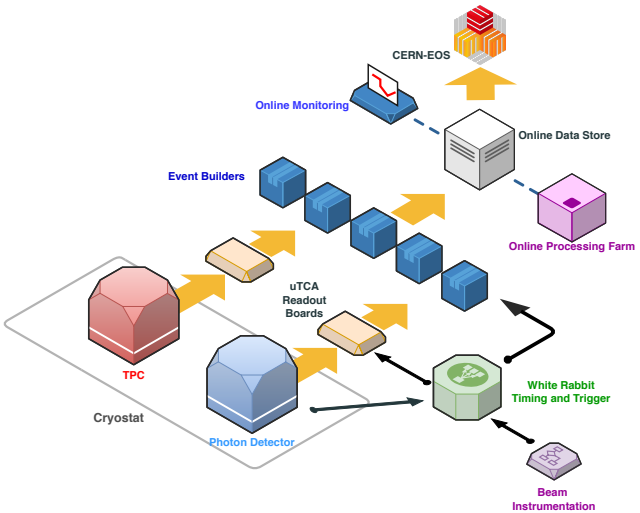

- Very large data volume continuously written to disk
- Lossless compression close to front-end
- Large Processing farm for online reconstruction and data quality
- Events sent offline at up to 20 Gb/s

Electronics

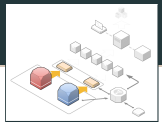

- Cryogenic ASIC amplifiers are **externally accessible**
- Digital electronics accessible at top of cryostat
- Architecture based on μ**TCA** standard
  - 12 crates, 10 AMC cards, 64 ch/card
  - 1 crate for light readout
  - 64 ch AMC digitisation cards

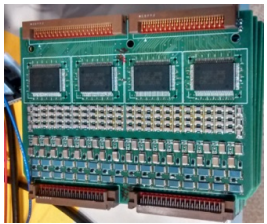

Front-end amplifier card

μTCA crate

Electronics on cryostat

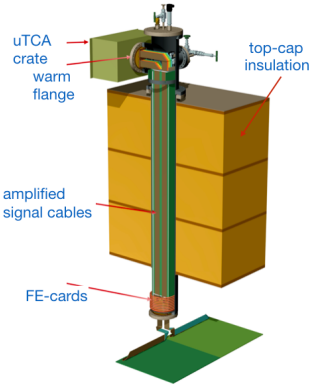

AMC readout cards

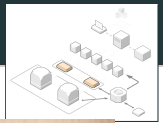

- Charge Readout Plane (CRP)
  - Intel Cyclone V GX FPGA with NIOS processor
  - 2.5 MHz, 12 bit, 10 GbE connectivity
  - Lossless Huffman-like **compression** (factor 10)
  - μTCA backplane provides dedicated WR clock and trigger transmission
- Light Readout (LRO)
  - Basic architecture derived from charge readout AMC cards, 16 LRO channels per card
  - External trigger ± 4 ms around beam spill
  - Internal Light ReadOut trigger from CATIROC ASIC for out of spill data
    - will acquire 1 drift window for these LRO triggers
  - 14 bit digitiser @ 40 MHz reads PhotoMultiplier tubes

Timing & Trigger

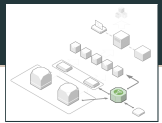

- **White Rabbit** distribution system
  - Sub-1 ns accuracy
- external triggers from beam instrumentation, cosmic ray counters, or light readout system
- μTCA DAQ architecture integrated with White Rabbit network
- White Rabbit slaves in μTCA crates

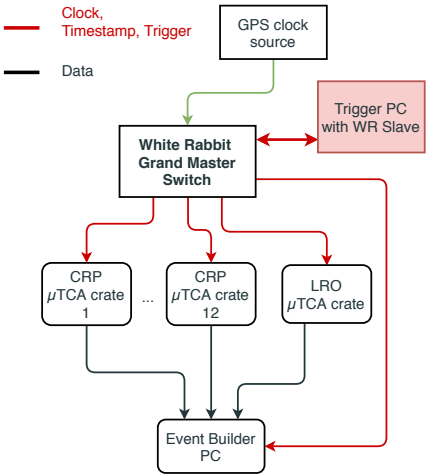

Event Builders

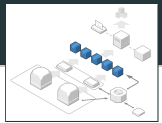

- Two Level-1 event building PCs
  - DELL R730, 256 GB RAM
  - 2 Intel X710 Quad Port 10 Gb/s
  - 1 Mellanox Connect-X3 Dual Port 40 Gb/s
  - Collates data for a drift (each PC does half-detector)
- Four Level-2 event building PCs
  - DELL R630, 128 GB RAM
  - 1 Mellanox Connect-X3 Dual Port 40 Gb/s
  - Assembles whole event from half-event fragments; prepares multi-event files for writing offline
- Interconnectivity via Brocade ITX7750 26 port switch at 40 Gb/s

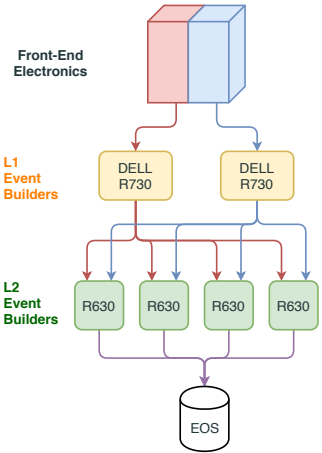

Storage/Processing

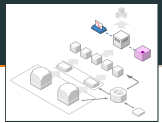

- Online processing and storage facility
- **1 PB** storage
- **384 cores** for processing
- internal bandwidth 20 GB/s
- key element of online analysis
  - identification of cosmics, purity, gain, event filtering

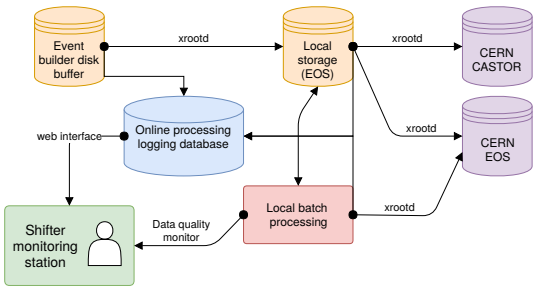

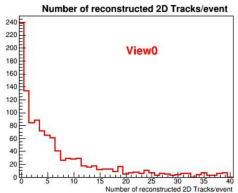

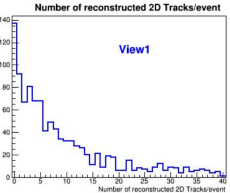

$3 \times 1 \times 1$ dual-phase prototype

- "proto"-prototype for DUNE-DP
- Exchange of electronics tested
- **Demonstrating first results** of charge amplification using cosmic rav muons

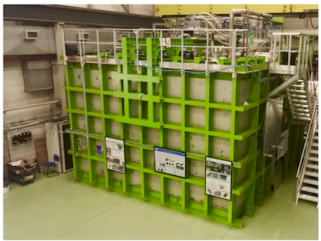

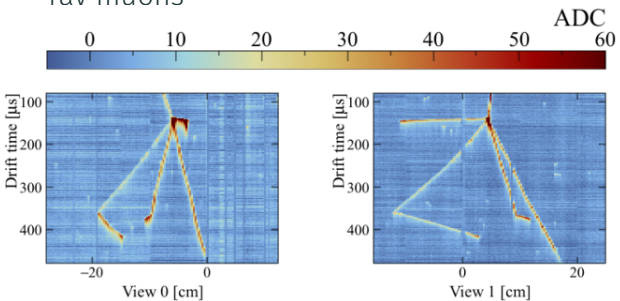

Cosmic ray (raw) events recorded in the $3 \times 1 \times 1\text{m}^3$.

Data Challenge

- "DC2" involving SP, DP, and offline computing from CERN to Fermilab
- Single Phase
  - Data copy without disturbing DAQ operation
  - Sustained full 20 Gbit/s bandwidth EHN1 to EOS
- Dual Phase
  - Steady 20 Gbit/s EHN1 to EOS for 24 hrs
  - Very few errors, and causes fixed shortly thereafter

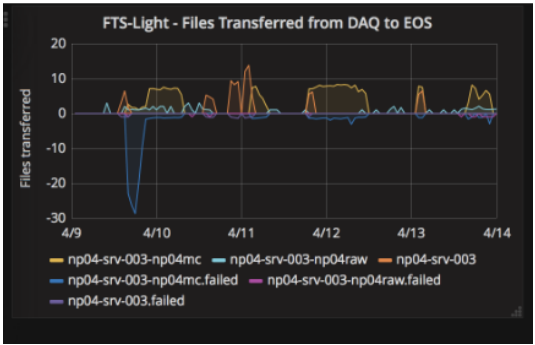

Final Remarks

- ProtoDUNE single and dual-phase are essential milestones on the roadmap to DUNE
- Qualify electronics and DAQ solutions and provide early physics performance feedback
- Largest LAr-TPC and test-beam experiments today
- Several DAQ solutions under study
- Looking forward to providing first results by the end of 2018

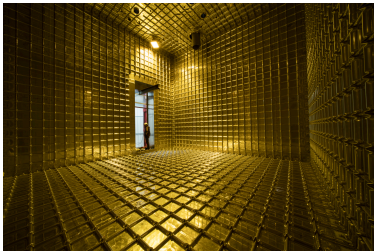

Inside the cryostat (before module installation)

Thank You