The NOvA Data Acquisition System

A highly distributed, synchronized, continuous readout system for a long baseline neutrino experiment

Andrew Norman, Fermilab

for the NOvA Collaboration

NOvA Physics

- The recent measurements of a large θ_{13} have greatly changed the landscape of the measurements NOvA will make.
- With $\sin^2 2\theta_{13} = 0.095$ NOvA will not only make an independent measurement of θ_{13} through ν_e appearance but will measure both ν & anti-ν modes:

$$P(\nu_\mu \to \nu_e) \quad \text{and} \quad P(\bar{\nu}_\mu \to \bar{\nu}_e)$$

- The measurements together provide significant information on the neutrino mass hierarchy and can establish/constrain CP violating in the neutrino sector through the phase δ_{CP}.

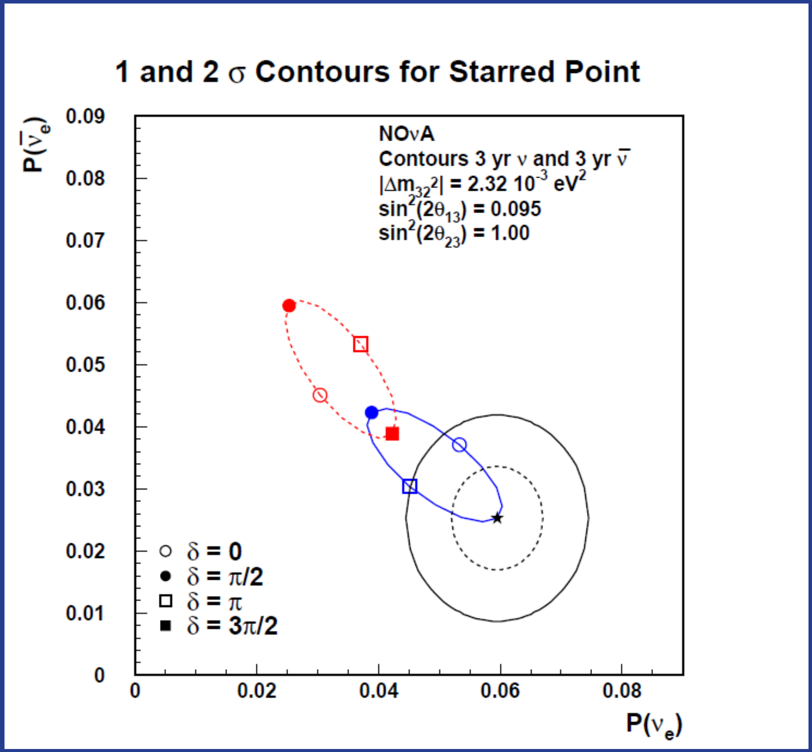

NOvA sensitivities to neutrino mass hierarchy and CP phase. Plot shows the parametric contours in δ for normal (blue) and inverted (red) mass hierarchies. NOvA 1 and 2σ contours for measurements corresponding to δ=3π/2 in the normal hierarchy.

NOvA

- NOvA is a program to investigate the properties of neutrinos
- It includes:
  - Doubling of the Fermilab NuMI beam power to 700 kW
  - An 15kTon totally active *surface* detector, 14 mrad off axis at 810km (first oscillation max for 2GeV)
  - A 220 Ton totally active near detector
- It has been optimized as a segmented low Z calorimeter/range stack to:
  - Reconstruct EM showers
  - Measure muon track
  - Detect nuclear recoils and interaction vertices
- It presents a challenge for modern computing and data acquisition

The 3 NOvA Detectors

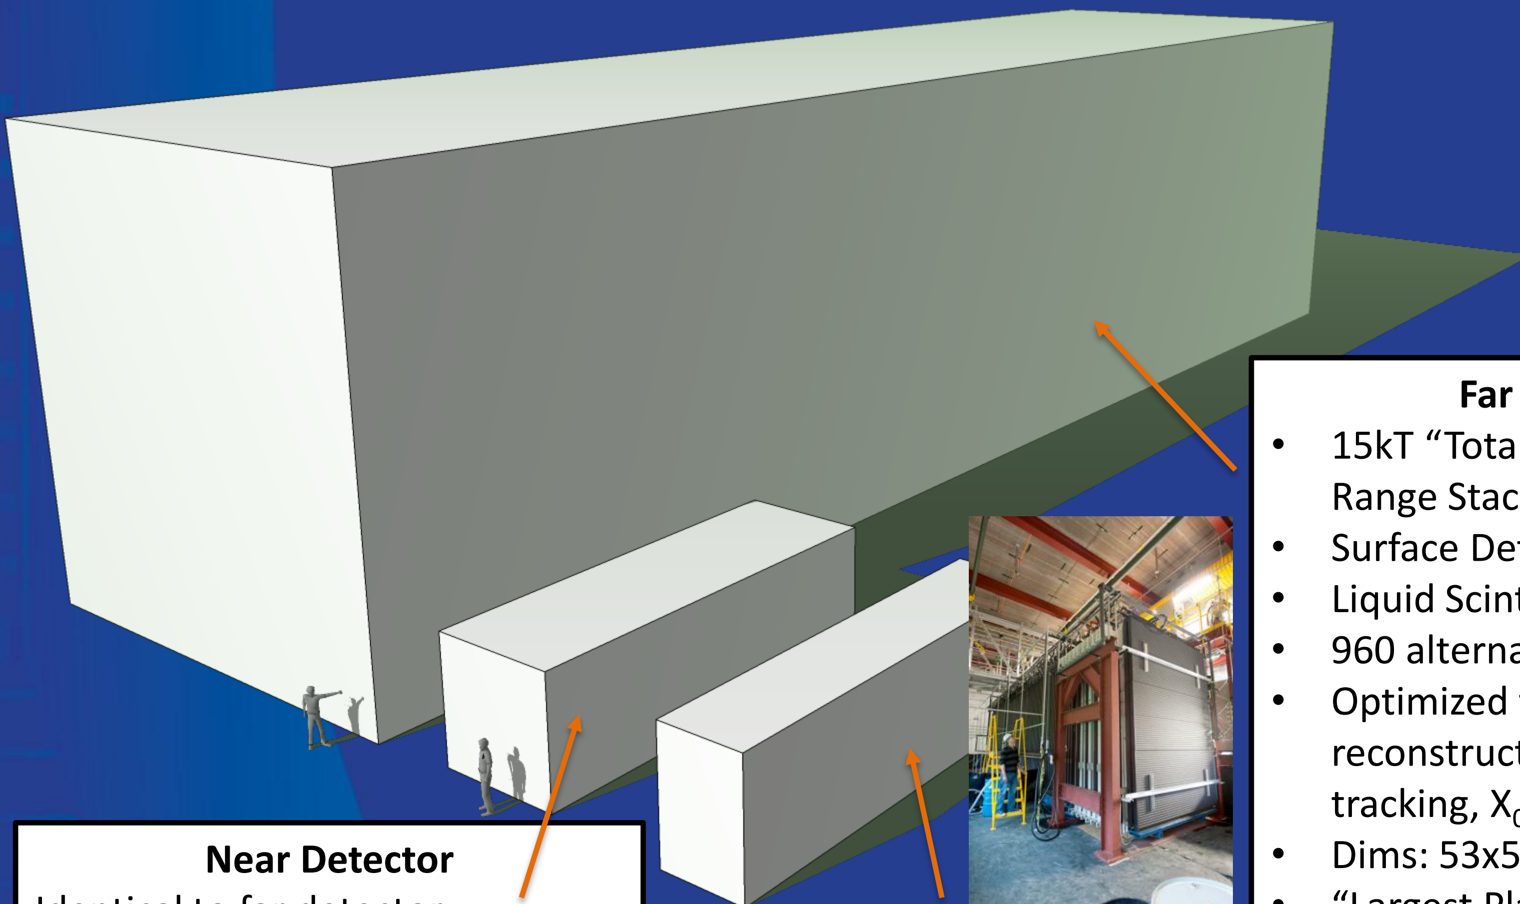

Far Detector

- 15kT "Totally Active", Low Z, Range Stack/Calorimeter
- Surface Detector
- Liquid Scintillator filled PVC
- 960 alternating X-Y planes
- Optimized for EM shower reconstruction & muon tracking, X_0 ≈40cm, R_m≈11cm
- Dims: 53x53x180ft
- "Largest Plastic Structure built by man"
- Began construction May 2012
- First operation est. Sep. 2012 (cosmics)

Near Detector

Identical to far detector
1:4 scale size
Underground Detector
Optimized for NuMI cavern rates -
- 4x sampling rate electronics

Near Det. Protype

In operation 2010-Present on surface at FNAL in NuMI and Booster beam line

The 3 NOvA Detectors

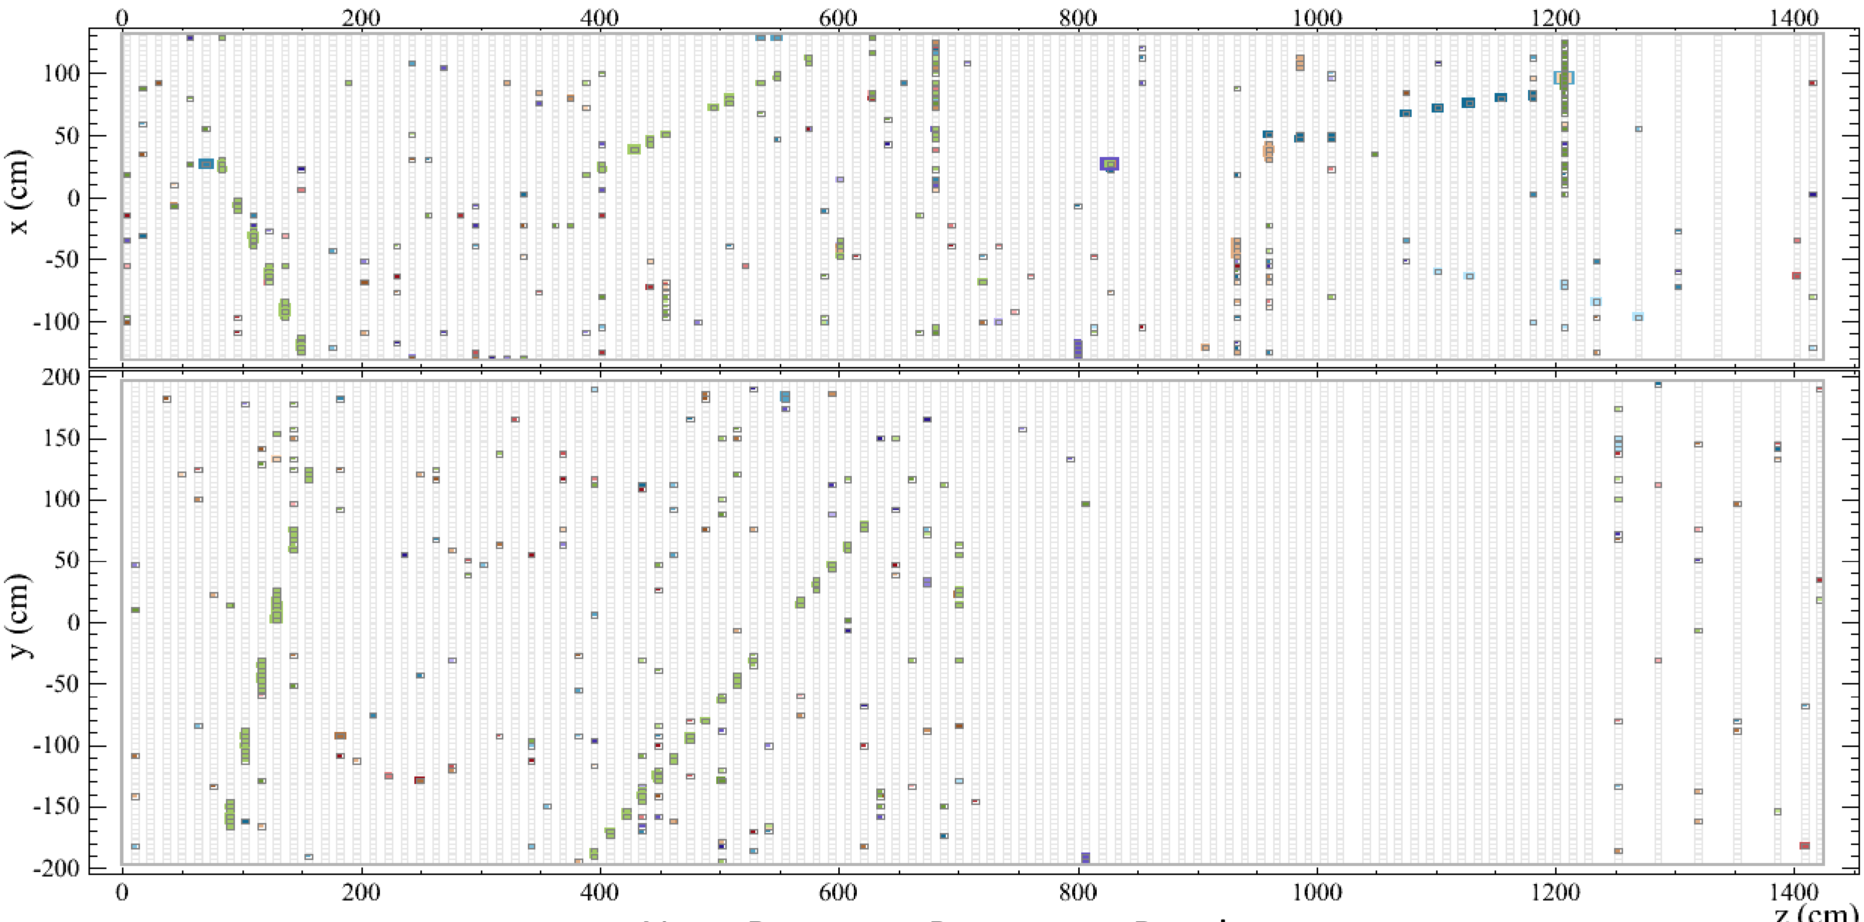

Near Detector Prototype Readout

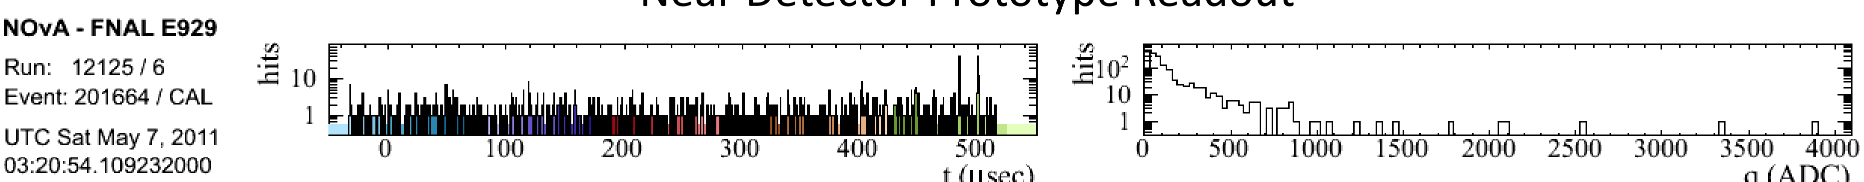

Making the Core Measurements

- To make the core physics measurements NOvA needs to be able to do the following:
  1. *Send* a narrow (10μs) pulsed beam of mainly ν_μ to a near detector (small) and far detector (massive) where the beam will interact.
  2. *Correlate* interactions in each detector with the beam
  3. *Compare* the rates of interaction to determine appearance/disappearance rates and energy spectra

The Correlation Problem

- There is no way to provide hard timing signals to the far detector 810km away.
- There is no way to "predict" the beam spill
    - NuMI spills are slotted into each accelerator super cycle but:
        - Timeline structure changes based on which experiments are running
        - Absolute [wall] time alignment of the super cycles will vary based on the timeline.
        - Unable to verify a spill until it has actually happened (i.e. kicker fired AND protons were recorded on target)
- The far detector is a surface detector
    - "Activity" base triggering won't work
    - There is always significant activity in the detector

The Correlation Solution

- Design a system that doesn't need a traditional trigger and can tolerate extremely long latencies in readout
- NOvA solution is continuous readout with absolute synchronization:
  1. Instrument every channel with TDC
  2. Synchronize every channel of the detector to every other channel (*global sync*)
  3. Readout the detectors continuously
  4. Time stamp every hit and store every hit in deep buffers
  5. Time stamp the beam spill at FNAL and send the information to the far detector (max latencies ≈ 20s baseline)†
  6. Match the hit data with the beam spills
  7. Record the data

†Baseline is 20s, actual buffer sizes allow for 60+ min

Readout Overview

- NOvA is a free running readout
  - The electronics are always live, always digitizing.
    - Every channel of the detector is being continuously digitized.
  - The whole system is completely deadtime-less.
  - The entire raw detector data stream (up to 4.3GB/s) is actively buffered in a large computing farm

- Triggering is asynchronous & decoupled from the readout
  - Trigger information packets are sent from FNAL over the internet to the far detector.
    - Average one way flight time is 11ms.
    - The system is designed to tolerate at least 20s of latency.
    - Real buffering can be much deeper (60+ minutes)
  - Buffered data that has a time overlap with a trigger window is saved to permanent storage

The Readout Cell

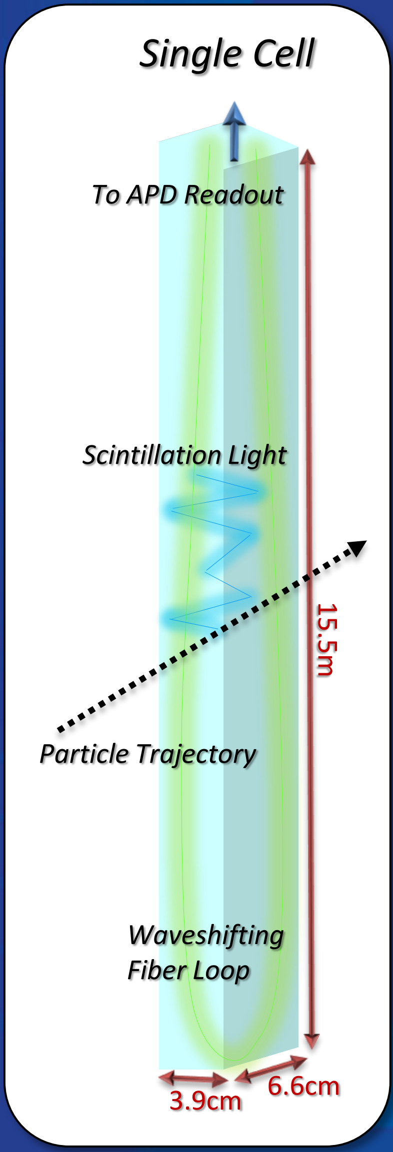

NOvA Detection Cell

- The base detector unit 3.9x6.6cm cell 15.7m long, filled with a mineral oil based liquid scintillator.
- Passage of MIP through the cell (longitudinal to beam) results in dE/dx ≈12.9 MeV across the cell.
  - Roughly 10% of energy loss is in the PVC wall
  - Yields 10-12MeV of deposition in the scint.
- The measured light output is 30-38 p.e. from the far end of the cell into the APD readout
- Light yield gives a minimum Sig/Noise 10:1 (far end)

- Signal is amplified and shaped
- Channel is continuously sampled at 2MHz to obtain wave form
- Zero suppression is performed at (0.5-0.66) MIP (15-20 p.e.) via a dual correlated sampling (dcs) algorithm
- Results in a single cell, lower energy detection threshold of ≈6-8MeV (far end)

- Fits to the shape of the waveform recover the timing and pulse height information

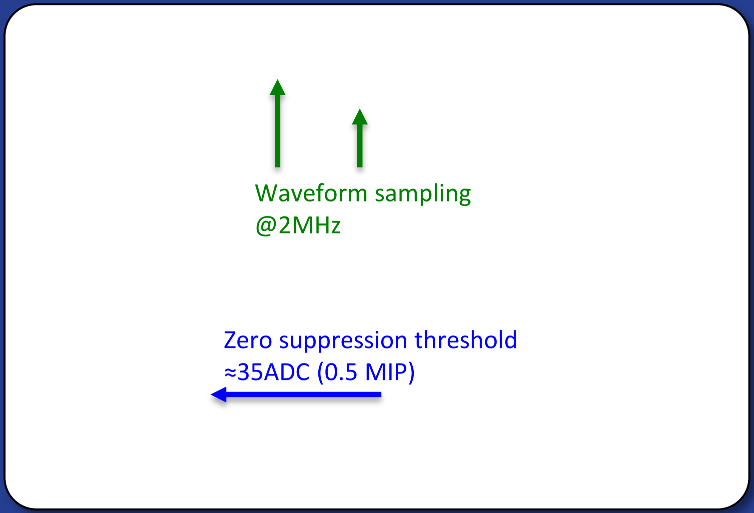

Data Concentrator Module

- The DCM is designed to aggregate the continuous data stream 64 front end boards (2048 detector cells) into partially time sorted data block
- All the data in one DCM is one geographic quadrant on the detector
- The DCM combines a large FPGA with a single board computer to build 5ms
  - Event building performed in both the FPGA and the CPU
- Run custom embedded linux with support for NOvA firmware
- Demonstrated to handle 24MB/s data rate on near detector

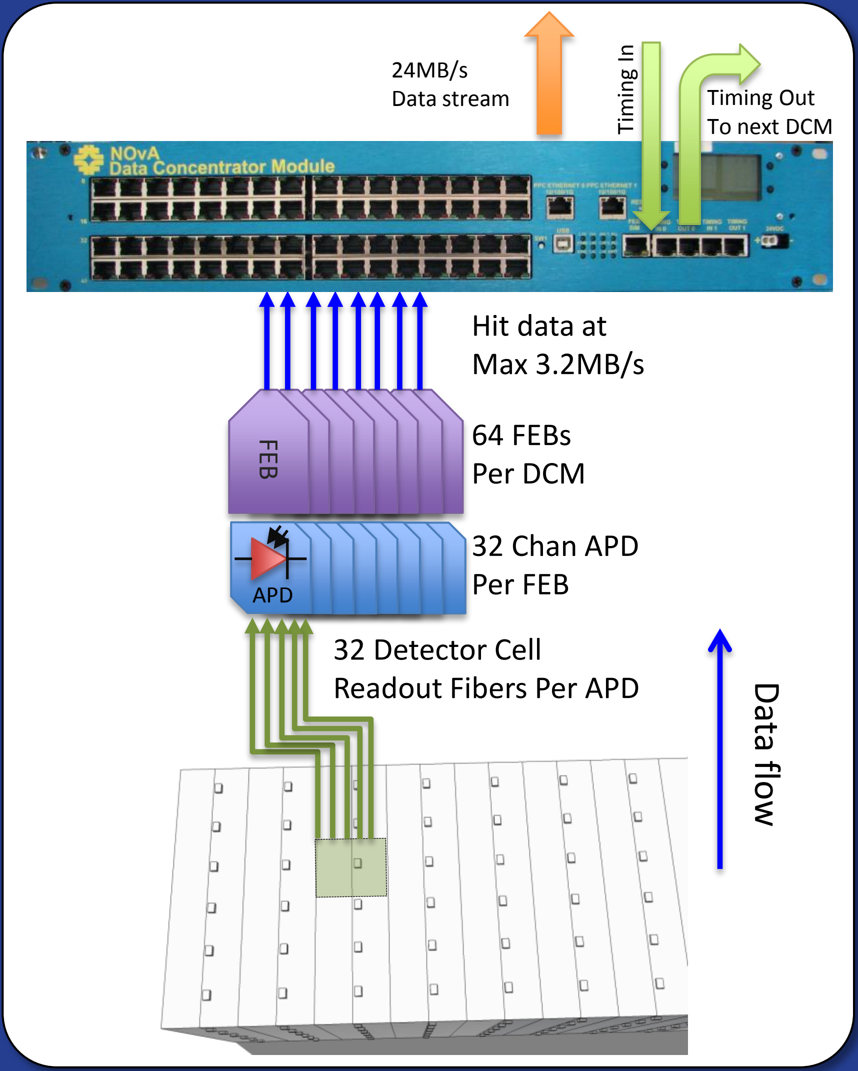

Instrumented Detector

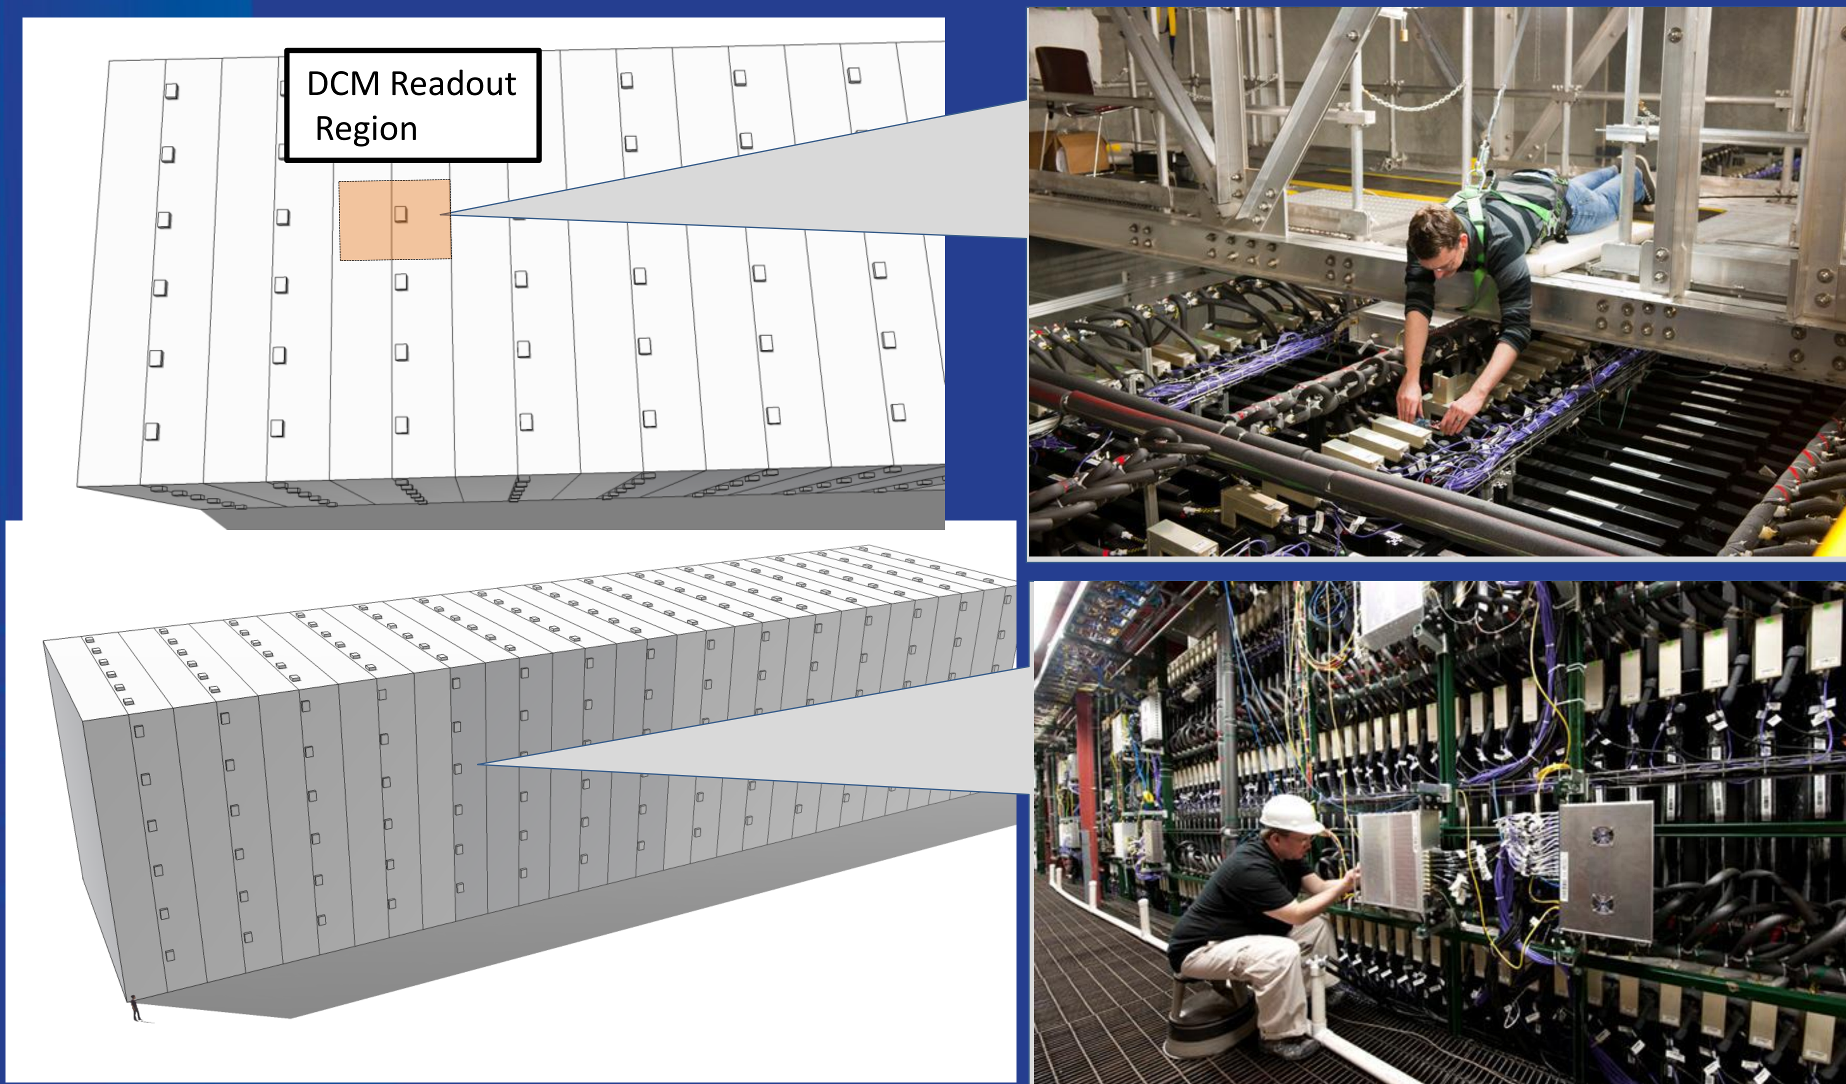

Timing System

- NOvA designed a custom timing system implemented on ARM/PowerPC + FPGAs which provides:
  - A 56 bit timestamp counter on every device
    - Timestamp counter can operate at 64MHz
  - A self calibrating synchronization system that can set/align the timestamps
    - We use GPS to link sites with universal “wall time”
    - Align “NOvA Time” to GPS with epoch offset to establish higher resolution timing
  - System uses a hierarchical topology to distribute:
    - Stable system clock (16/64MHz)
    - Timing Command Data
    - Synchronization pulses
    - Loopback calibration

Timing System Master/Slave

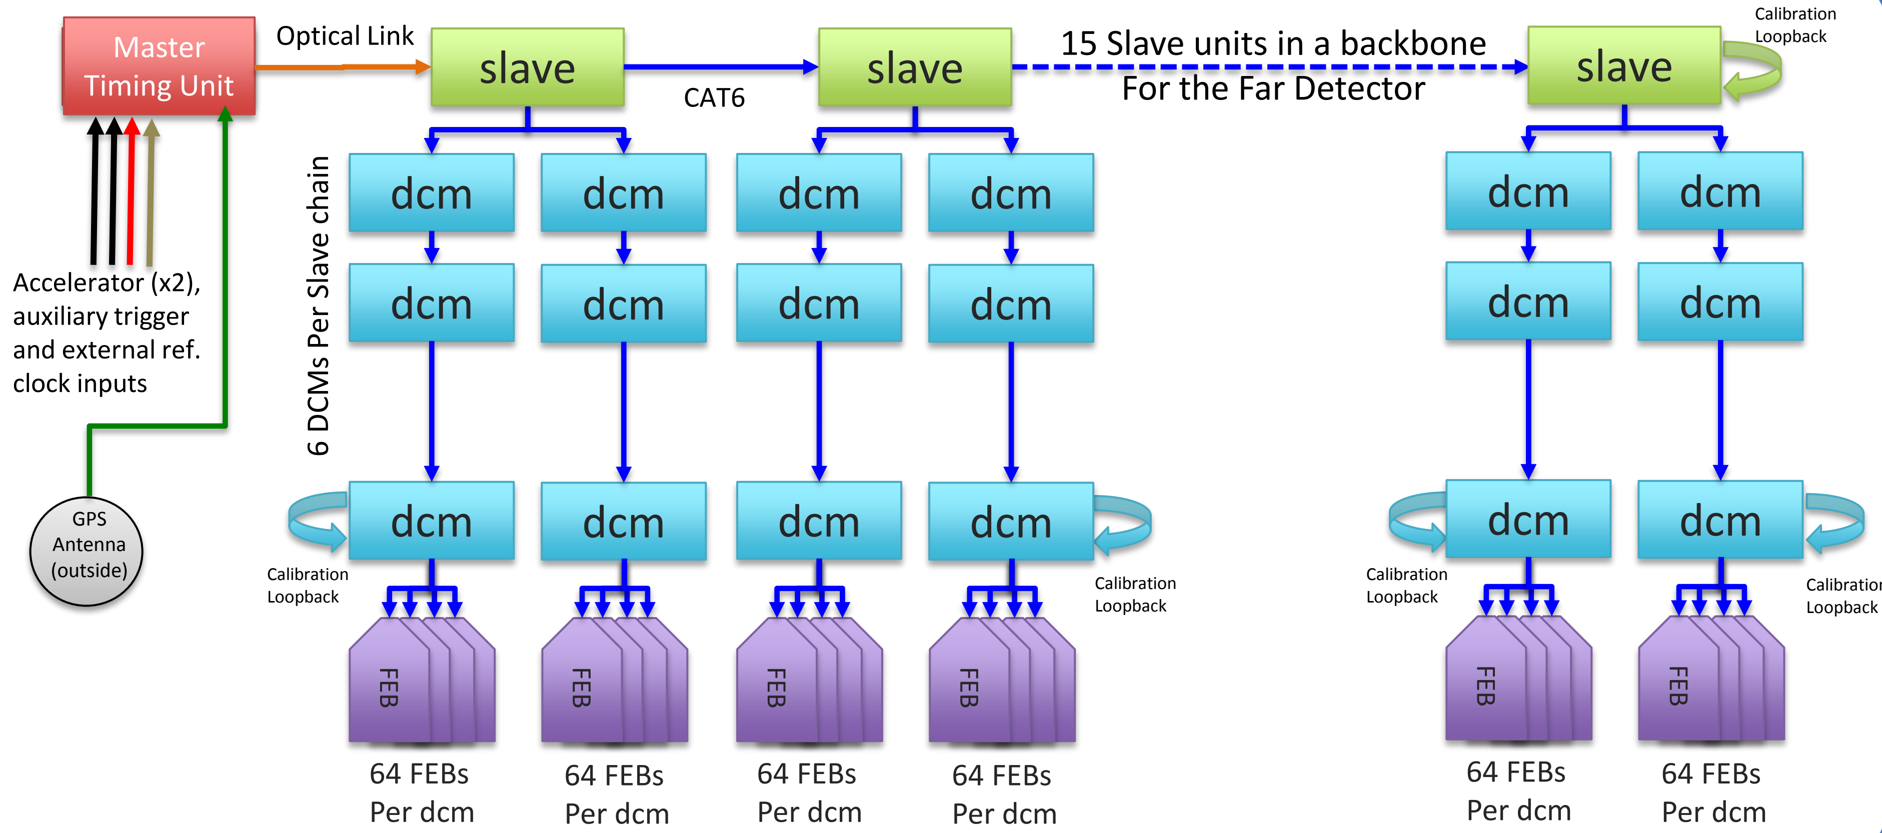

Timing System Master/Slave

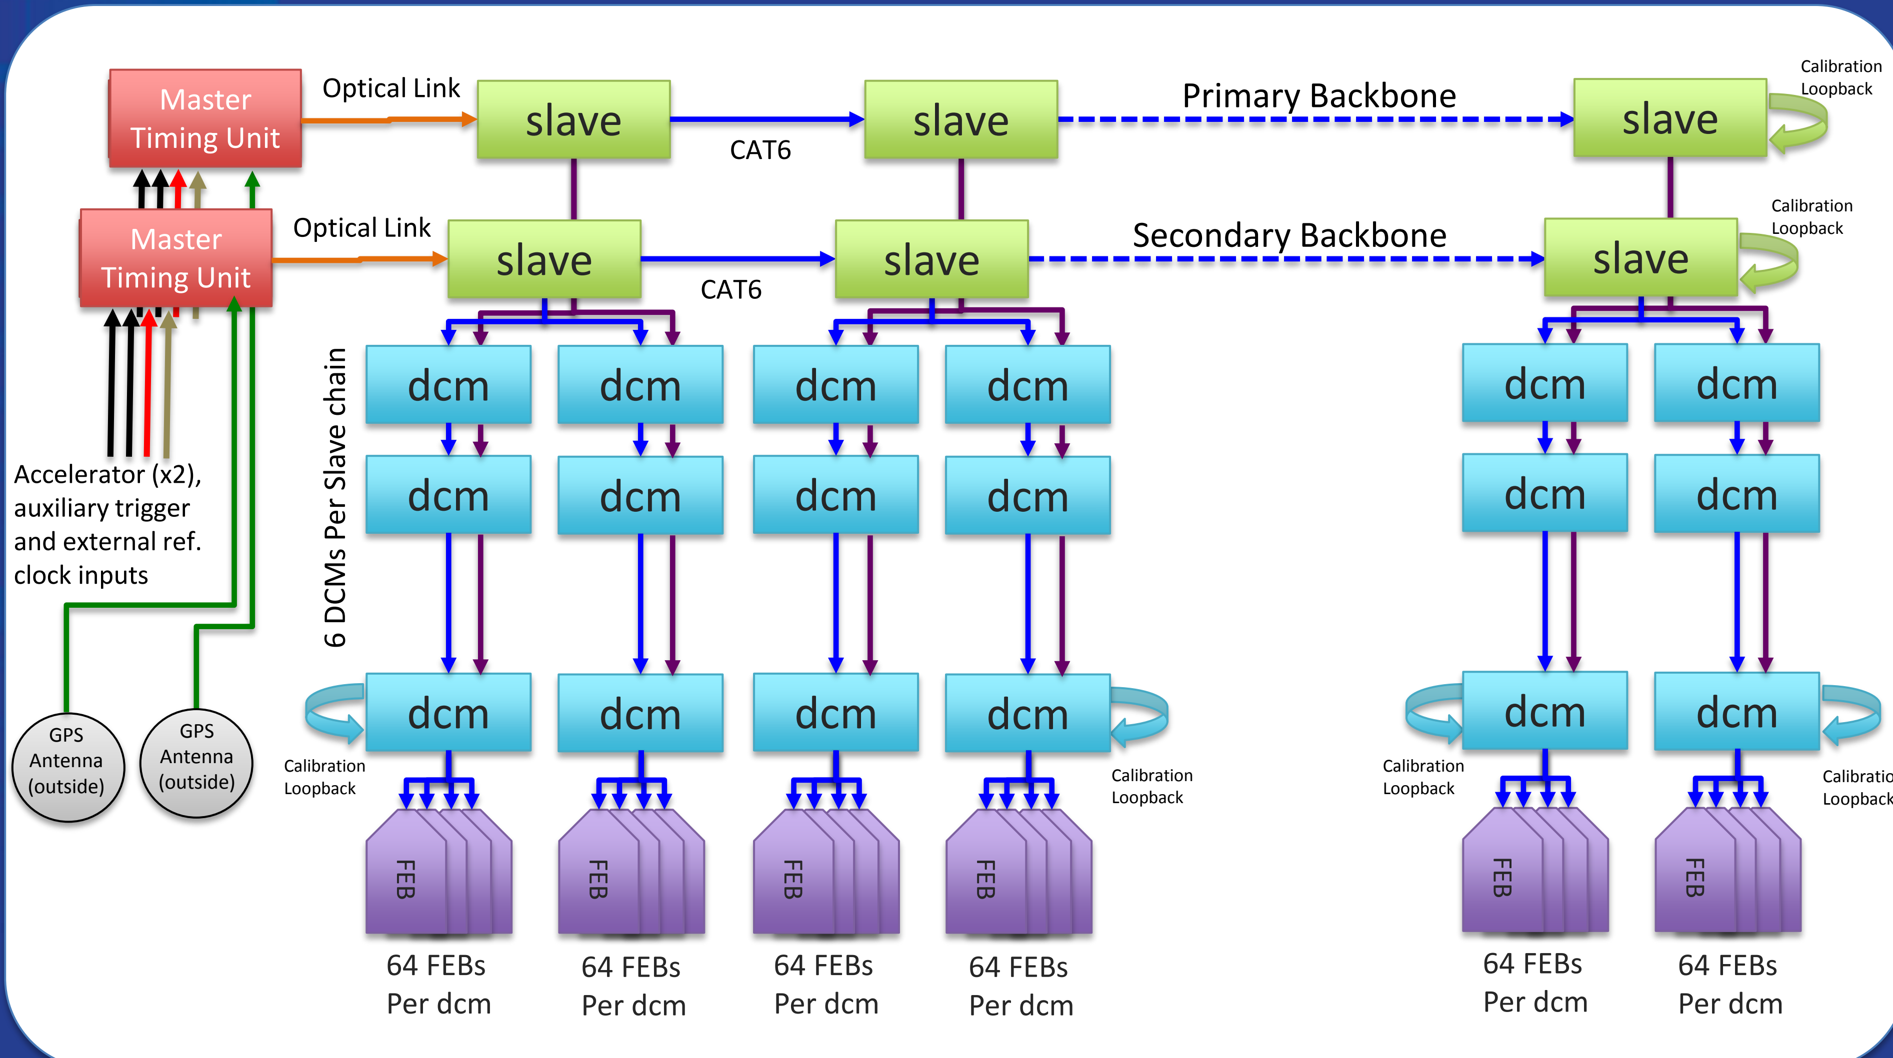

Timing System Master/Slave

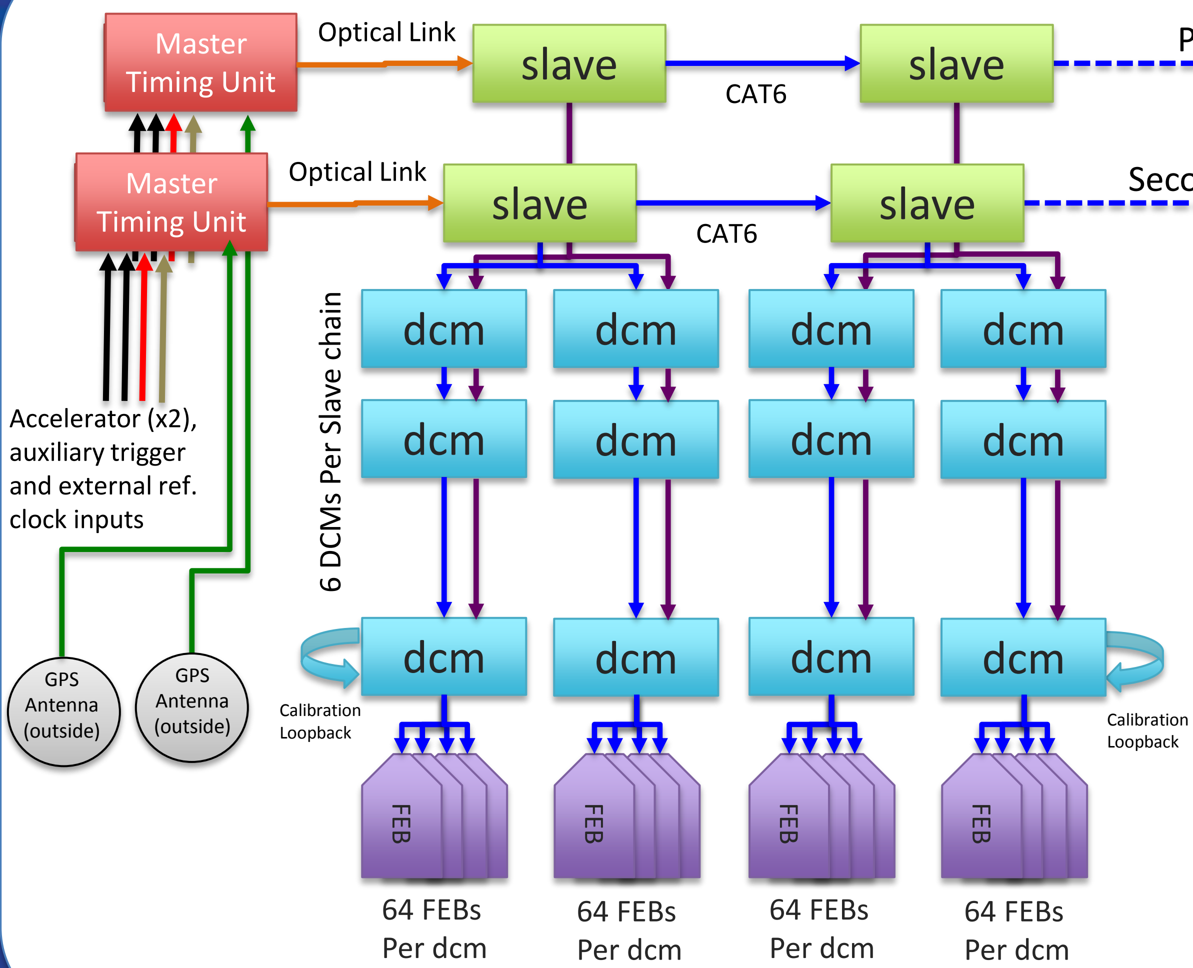

NOvA Far Detector Timingsystem & teststand in production for monitoring and diagnostics of Ash River GPS

Synchronized DAQ

When you are done you have:

Device	Daughter	Number
Master	1 per chain	3 per site
Slave	15 per chain	30 far/4 Near
DCM	6 per chain	30 chains, 180 units 14 near†
FEB	64 per DCM	11520 far *576* near
Readout Cells	32 per FEB	368,6400 far *18,432* near

Sync'd and counting in lockstep.

Synchronization

- How do you actually sync 386,000+ channels of readout?

You take a cue from Bell Telephone:

"At the tone the time will be…."

Sync System

Step 1

Use GPS to determine the 64bit NOvA time at the next 1s boundary (64MHz precision)

Send this time over the serial command channel.

The timing circuit loads the bytes into the "time stamp preset" registers.

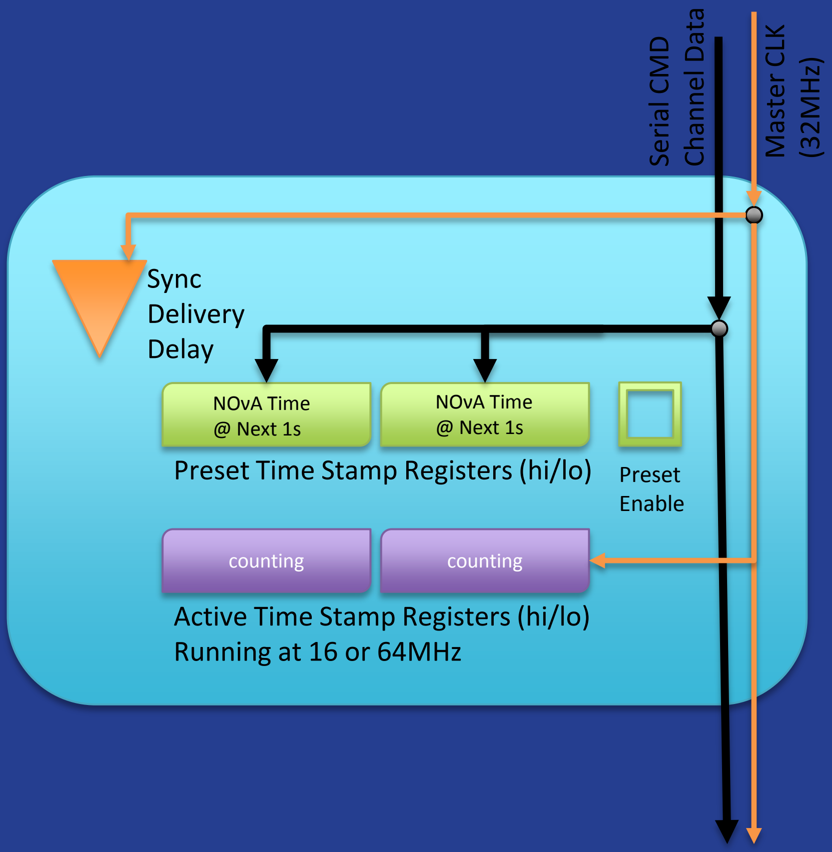

Sync System

Step 2

Use the serial command channel to set the “Preset Enabled Arm” register.

This Arms the system to react when it receive the next sync pulse.

Having a specific Arm functions allows the sync line to be used for differ purposes like:

- Self Delay Calibration
- Timing Fiducially Diagnostics
- External Sync. Verification

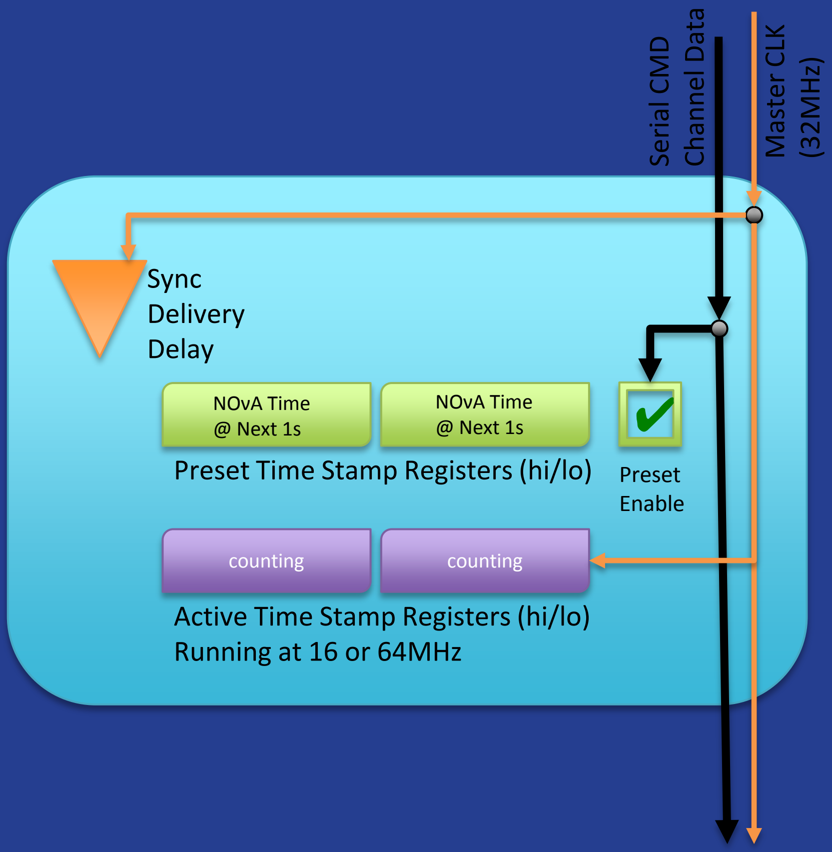

Sync System

Step 3

Send a pulse on the dedicated SYNC line at a predetermined time BEFORE the next 1s boundary.

The pulse is buffered and retransmitted at each unit.

Each unit is calibrated to know its offset from its upstream and downstream partners in the chain.

The pulse is not actually delivered to the circuit until a calibrated number of master clock cycles have elapsed on this unit.

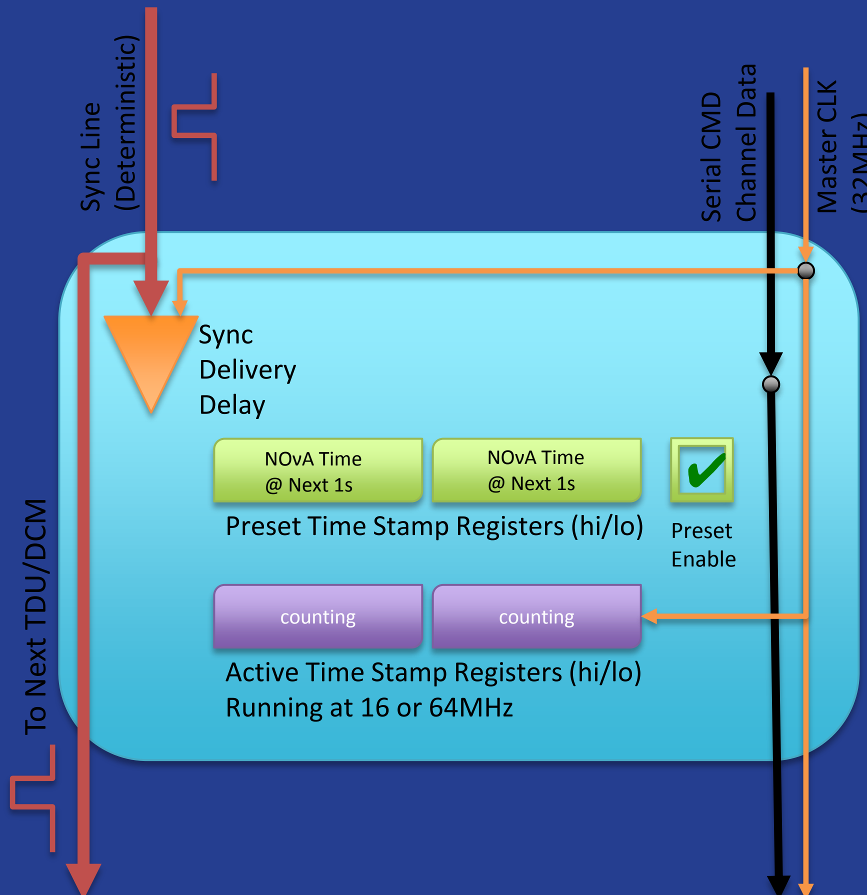

Sync System

Step 4

Upon reception of the delayed sync, the values of the preset register are latched into the active time stamp registers.

The ARM register is cleared and the system continues counting on the next master clock cycle with the synchronized time.

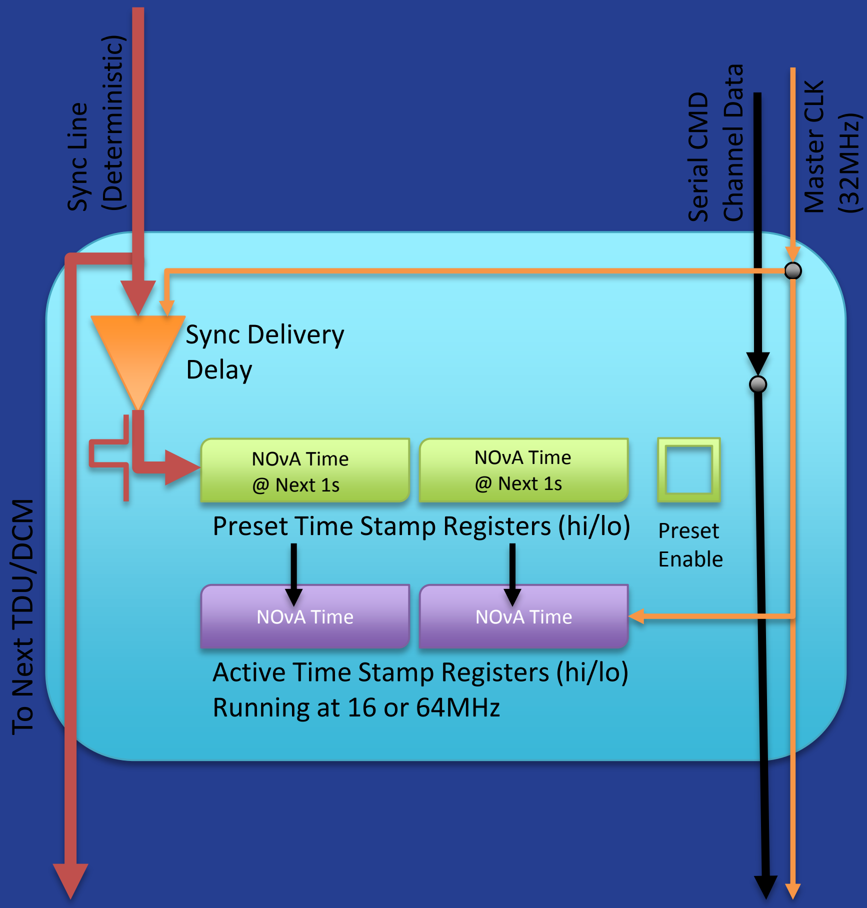

Fully Synchronized Detector

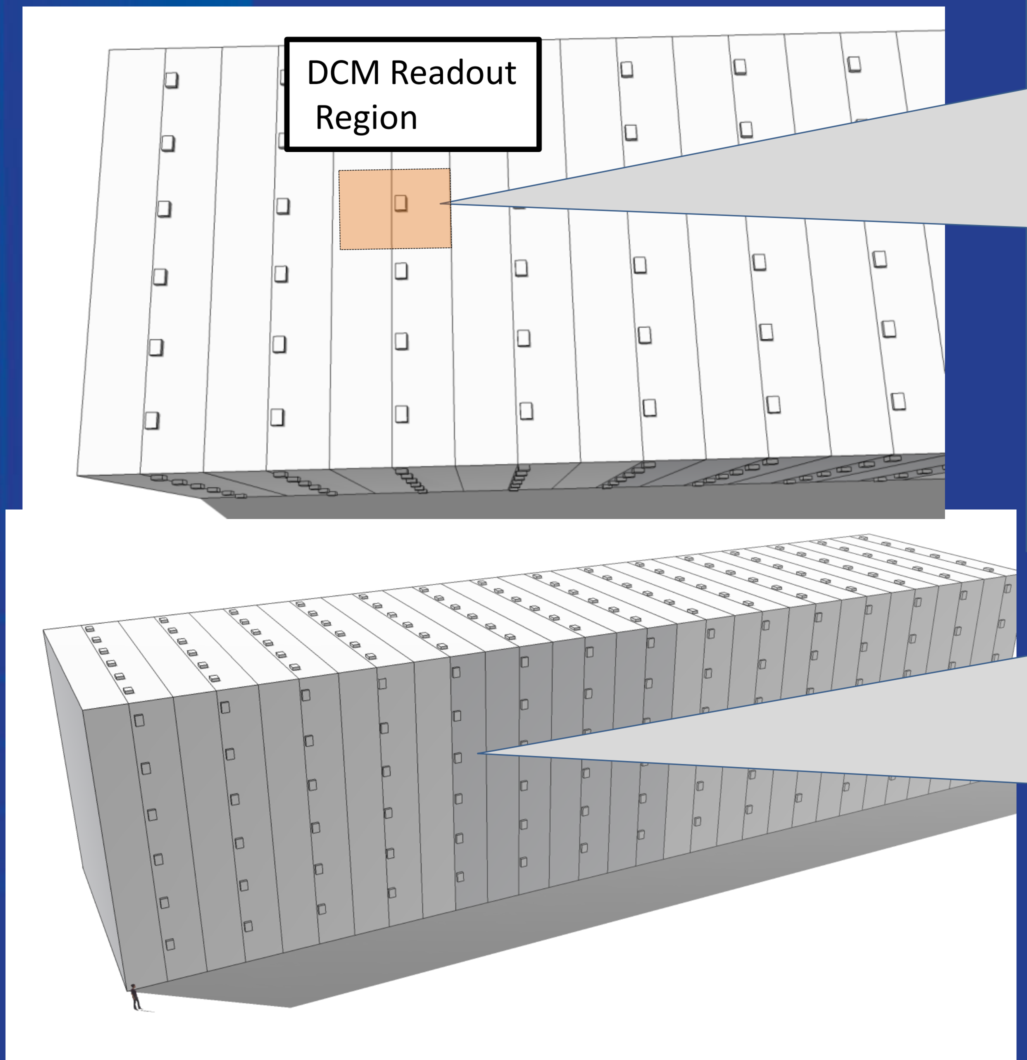

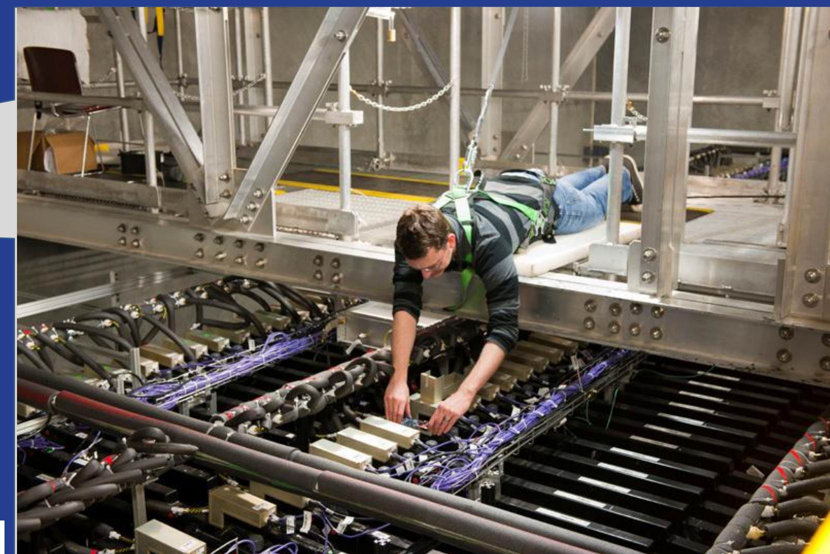

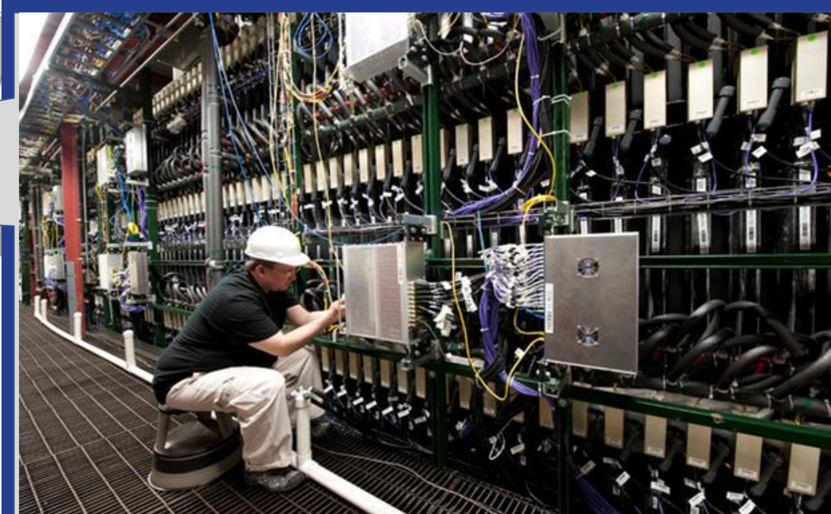

Fully Synchronized Detector

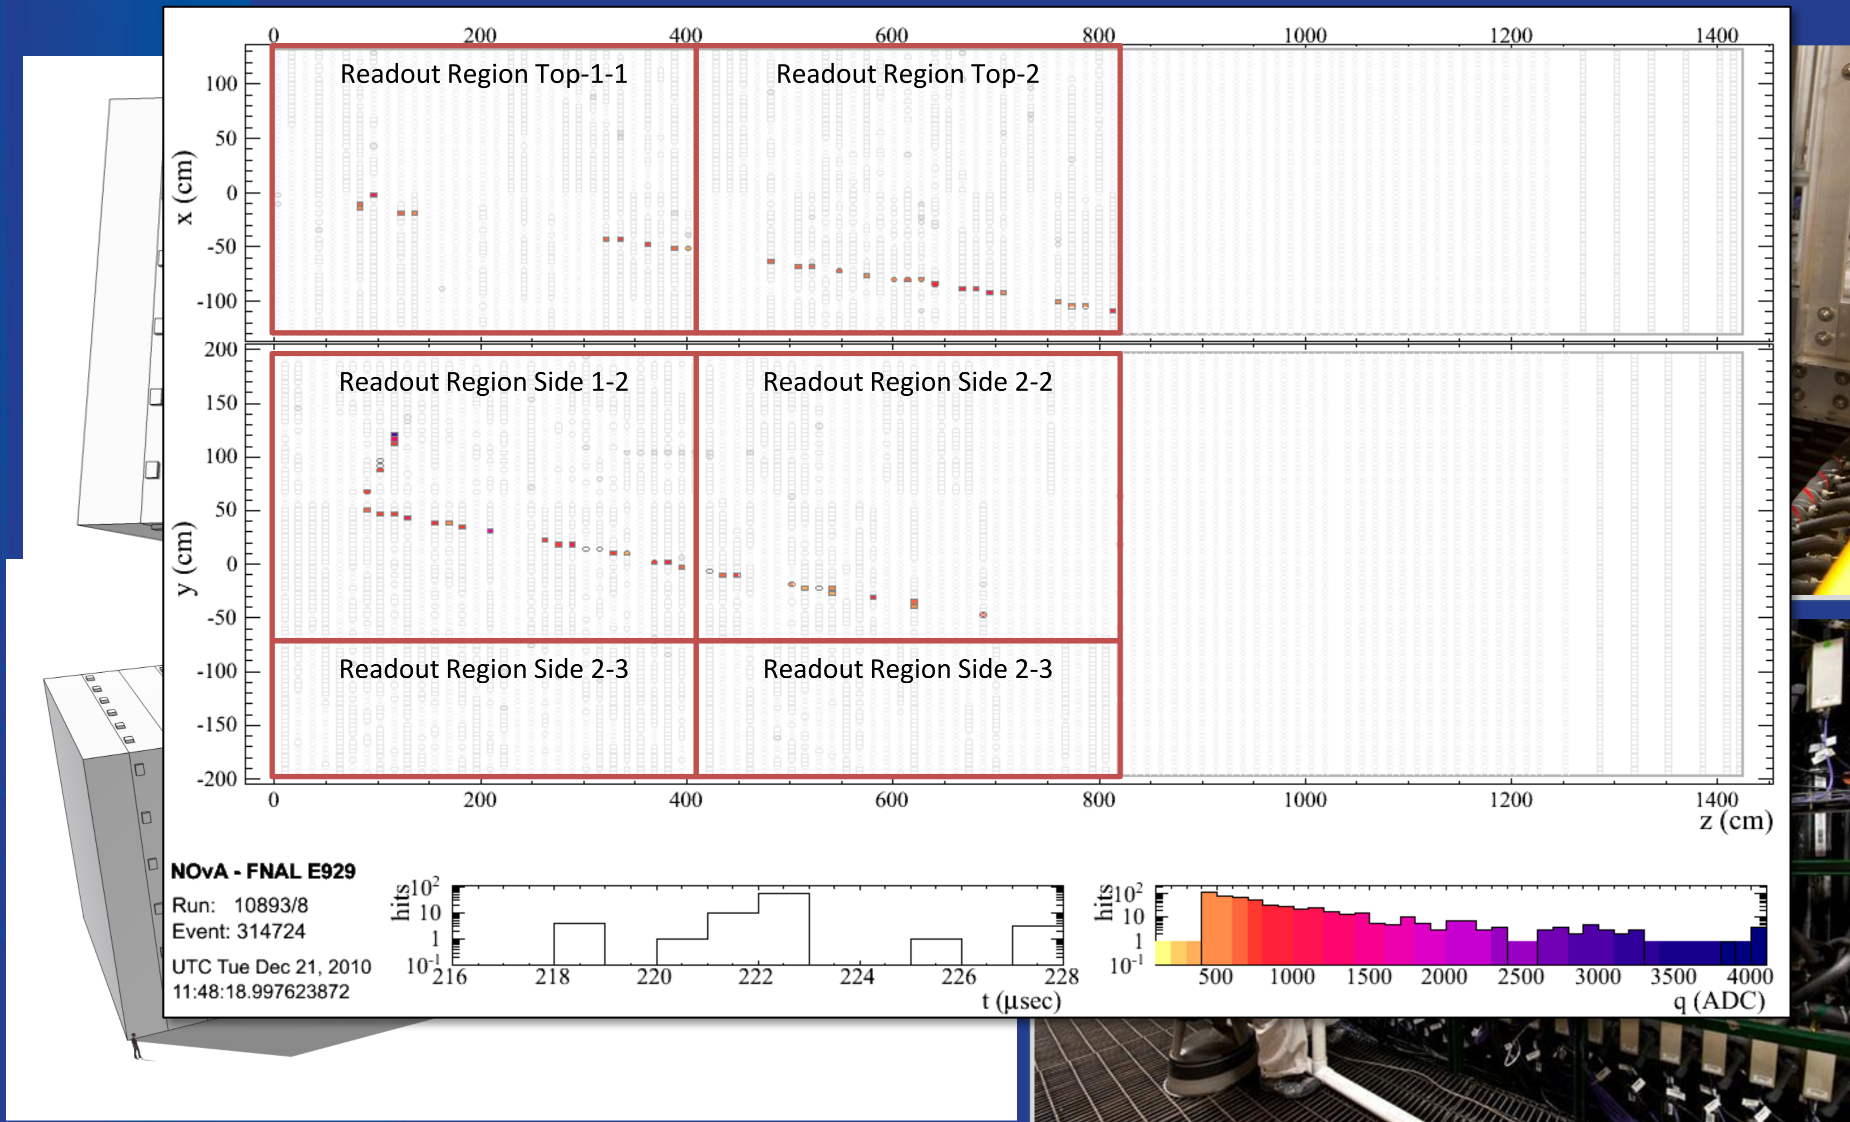

Readout Buffering

- The data from DCMs are sent to a farm of "buffer nodes"
- Every 5ms the data stream from ALL DCMs on the detector is transmitted into a single buffer node.
- This gives a complete 5ms snapshot of the detector in that buffer node
- The destination buffer node is then rotated in a round-robin pattern through the entire cluster (1s return to start with nominal 200 buffer nodes)
- Projected data rates allows us to buffer up to 69min of raw, min bias readout system wide *(baseline buffering 20s)*

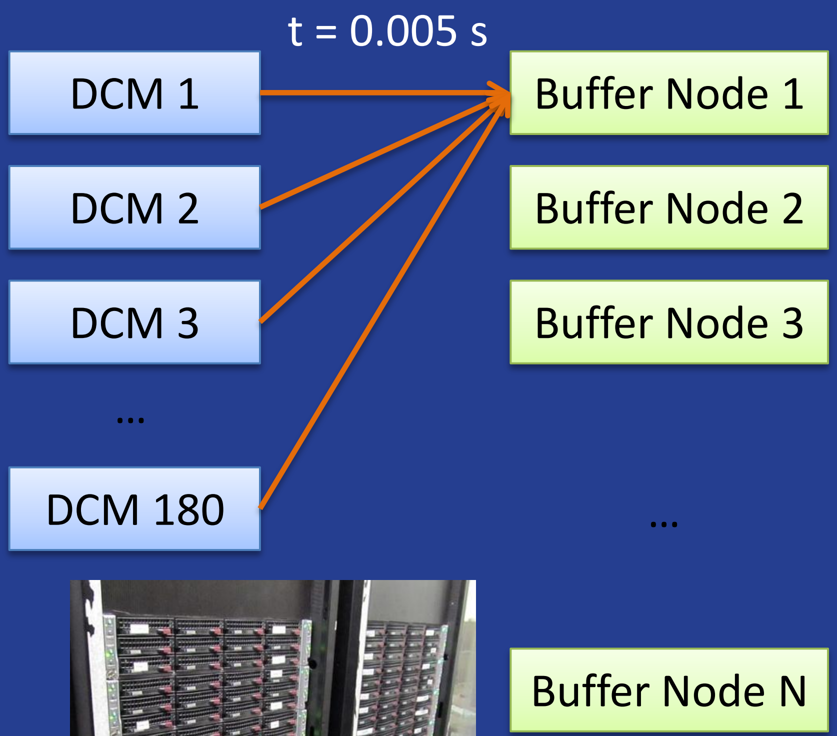

Buffer computing at far detector

Readout Buffering

- The data from DCMs are sent to a farm of “buffer nodes”
- Every 5ms the data stream from ALL DCMs on the detector is transmitted into a single buffer node.
- This gives a complete 5ms snapshot of the detector in that buffer node
- The destination buffer node is then rotated in a round-robin pattern through the entire cluster (1s return to start with nominal 200 buffer nodes)
- Projected data rates allows us to buffer up to 69min of raw, min bias readout system wide *(baseline buffering 20s)*

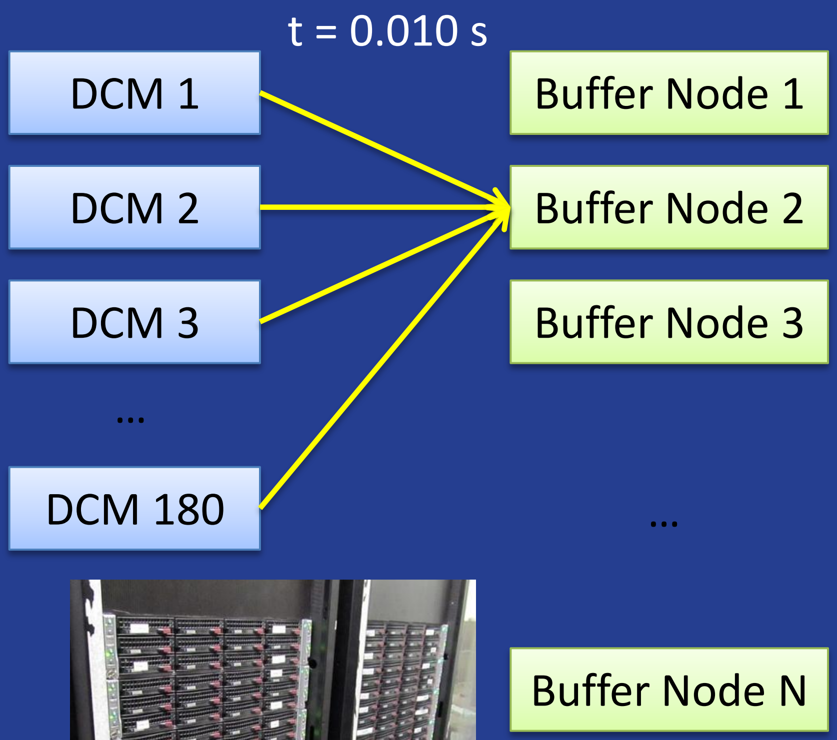

Buffer computing at far detector

Readout Buffering

- The data from DCMs are sent to a farm of “buffer nodes”
- Every 5ms the data stream from ALL DCMs on the detector is transmitted into a single buffer node.
- This gives a complete 5ms snapshot of the detector in that buffer node
- The destination buffer node is then rotated in a round-robin pattern through the entire cluster (1s return to start with nominal 200 buffer nodes)
- Projected data rates allows us to buffer up to 69min of raw, min bias readout system wide *(baseline buffering 20s)*

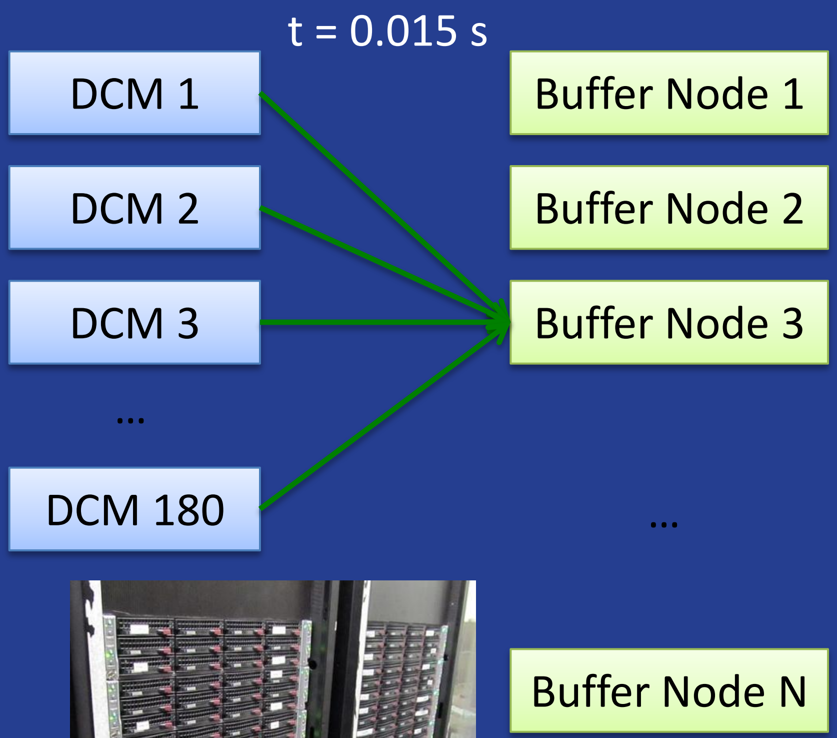

Buffer computing at far detector

Readout Buffering

- The data from DCMs are sent to a farm of “buffer nodes”
- Every 5ms the data stream from ALL DCMs on the detector is transmitted into a single buffer node.
- This gives a complete 5ms snapshot of the detector in that buffer node
- The destination buffer node is then rotated in a round-robin pattern through the entire cluster (1s return to start with nominal 200 buffer nodes)
- Projected data rates allows us to buffer up to 69min of raw, min bias readout system wide *(baseline buffering 20s)*

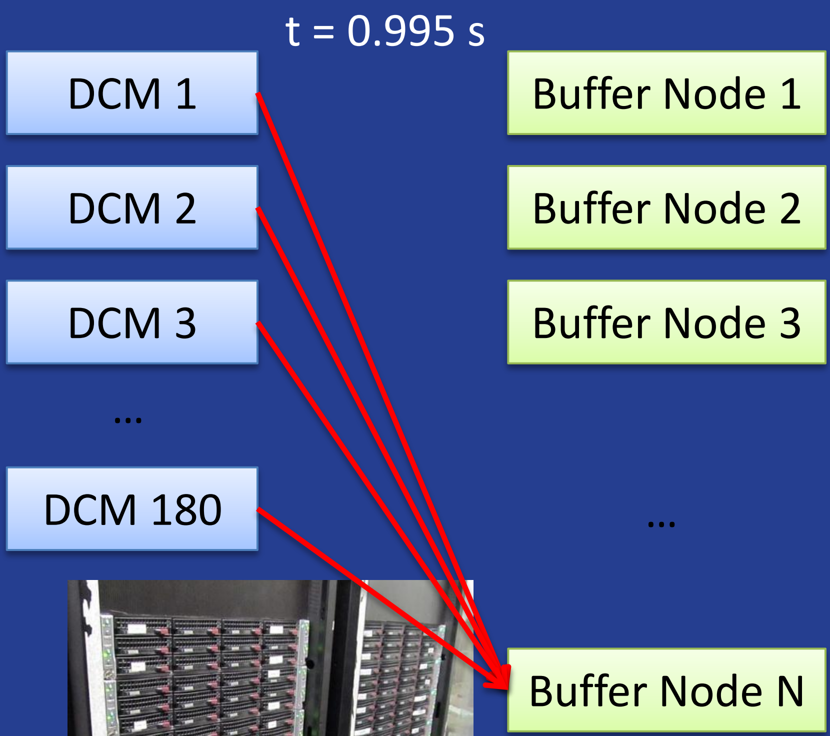

Buffer computing at far detector

Beam Triggering

- For the core NOvA beam related physics topics $\{\theta_{13}, \theta_{23}, \delta_{CP}, \Delta m^2\}$ the only "triggers" you need are:
  - NuMI Beam spill
  - Calibration pulser
- *These are completely open zero bias selections.*

- NuMI Beam Spill Trigger
  - Centered on the GPS time of the spill at FNAL
  - Extract a window around that time to get side bands
  - For intial running we use a wide window (500μs) to locate and characterize the beam
  - For production running we narrow the window to 30μs to retain sideband information, and take additional calibration data in the off spill intervals

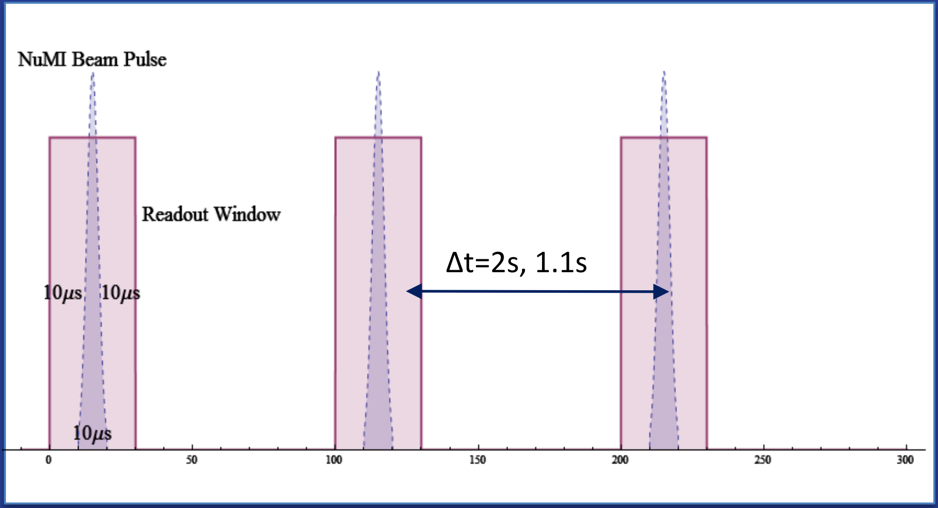

Nova beam trigger structure, 30μs trigger window corresponding to beam spill + sidebands

Beam Triggering

- For the core NOvA beam related physics topics $\{\theta_{13}, \theta_{23}, \delta_{CP}, \Delta m^2\}$ the only "triggers" you need are:
  - NuMI Beam spill
  - Calibration pulser
- *These are completely open zero bias selections.*

- NuMI Beam Spill Trigger
  - Centered on the GPS time of the spill at FNAL
  - Extract a window around that time to get side bands
  - For intial running we use a wide window (500μs) to locate and characterize the beam
  - For production running we narrow the window to 30μs to retain sideband information, and take additional calibration data in the off spill intervals

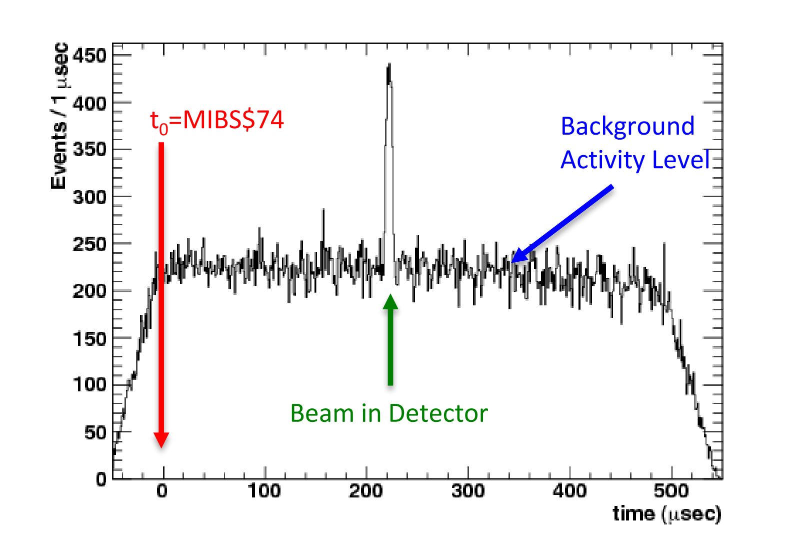

NuMI Beam Trigger window from 2011/2012 near detector running. Beam appears in the detector 220μs after the MIBs $74 signal to the extraction kicker

Data Driven Trigger

- It’s important to note that ALL the data (≈4.3GB/s) is buffered in our buffer farm
- But only 30μs of each 2s is written out as physics beam triggers

- Being able to examine the full data stream in real time while it is in the buffers gives us the opportunity to generate “data driven triggers”
- Because the detector is 100% live these triggers have tremendous potential for accessing Non-accelerator physics, calibration:
  - Improve detector calibrations using “horizontal” muons
  - Verify detector/beam timing with contained neutrino events
  - Search for exotic phenomena in non-beam data (magnetic monopole, WIMP annihilation signatures)
  - Search for directional cosmic neutrino sources
  - Detect nearby supernovae

DDT Design Overview

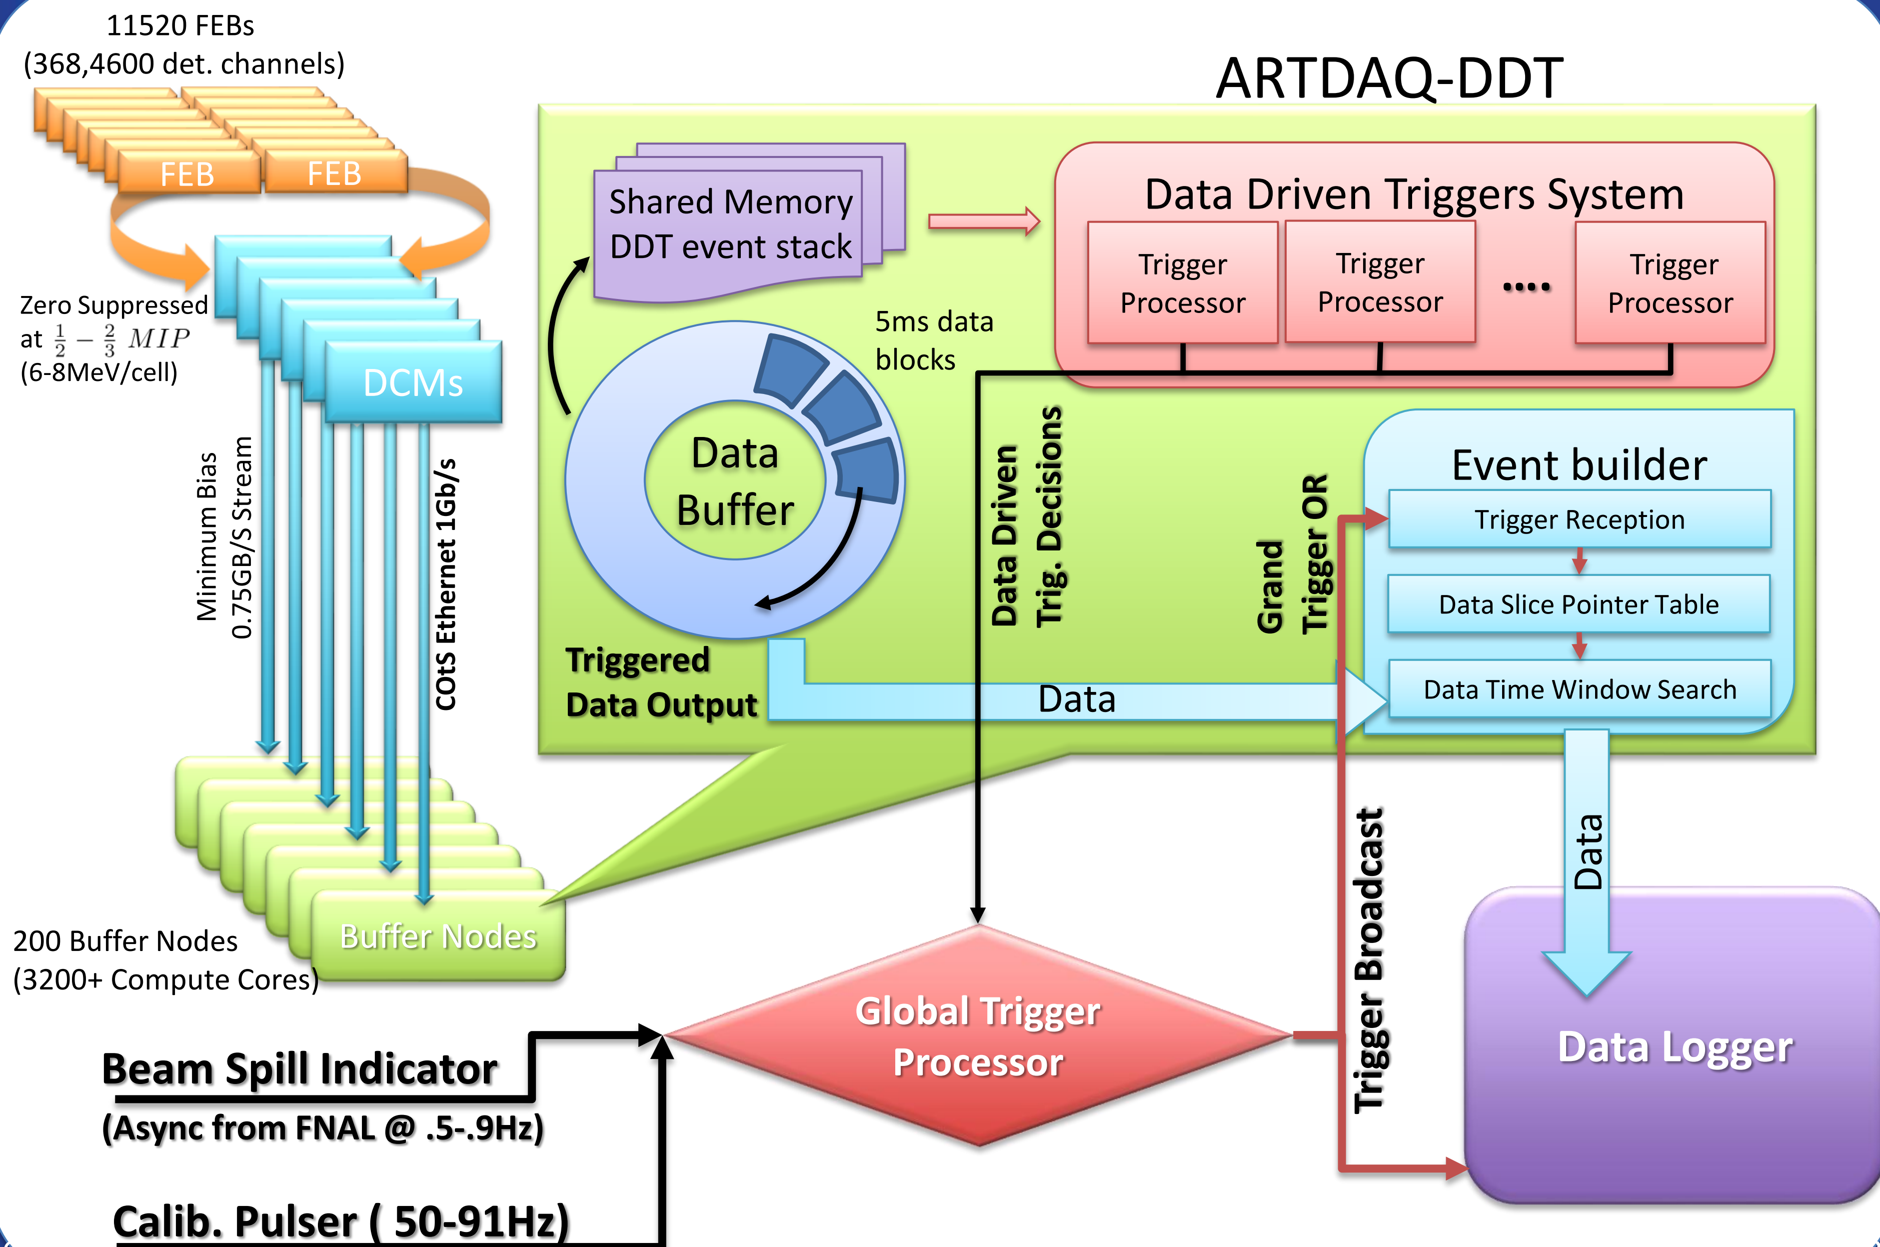

ARTDAQ-DDT

- Data Driven Triggering (DDT) is implemented in NOvA using ARTDAQ
  - Event processing framework (see talk C.Green)
  - Tailored to read from shared memory event buffers which are populated directly by the NOvA event builder
  - Partitioned off from primary DAQ chain for system stability
  - Permits the shared use of algorithms between online/offline:
    - transparent simulation of triggers, efficiencies, etc…
    - Leverages work on offline reconstruction & analysis modules
    - Simplifies development, testing, validation cycles

Example: Hough DDT

- Real world track finding in X/Y calorimeter
  - Linear Hough transform
  - 2D view version is a example of N^2 complexity algorithm
    - Can be implemented in serial, partitioned or massive parallel processing (i.e. GPU exploration)
    - Allows for tuning of both the algorithm and framework
  - For NOVA-DDT we implemented a global hit transformation on 5ms time windows
    - 500× nominal beam spill window
    - Allows for simultaneous searches for both fast and slow (β=1-10^{-4}) track topologies
    - Hough method is expandable into dim > 2 using timing & pulse height information

Example: Hough DDT

Transform applied Globally to X/Y Views

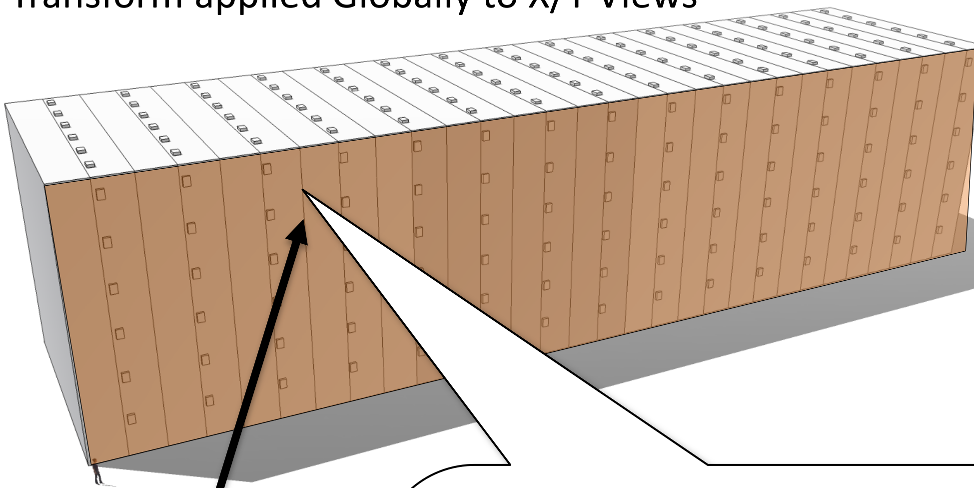

Local track features detected

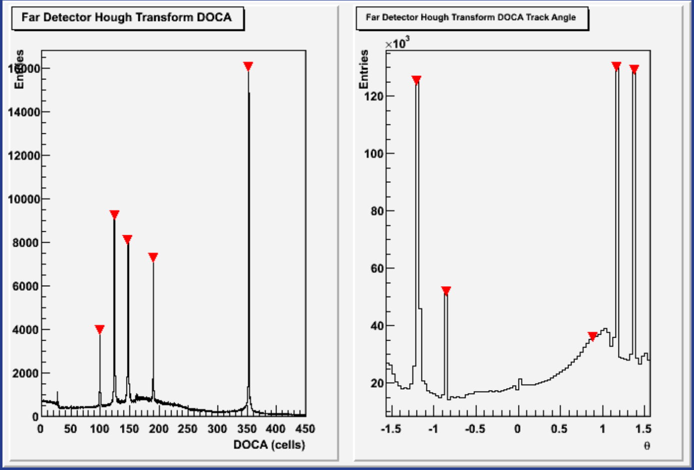

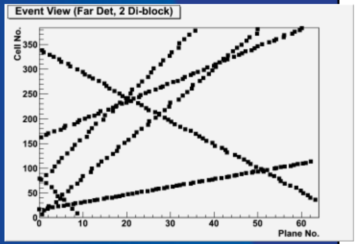

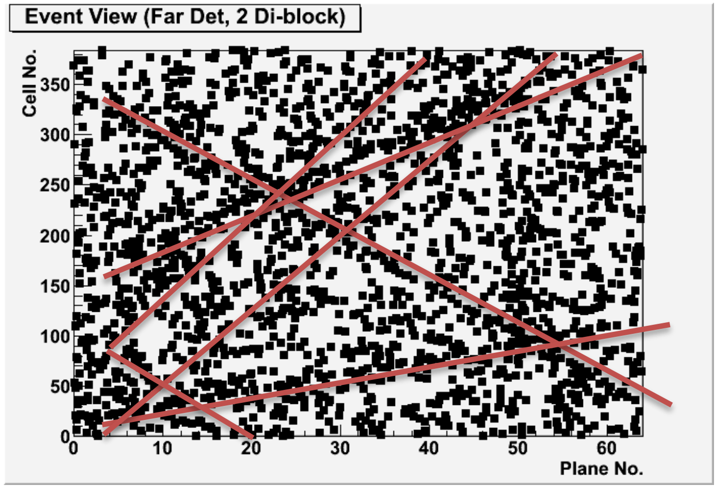

Extraction of tracks from long Δt (5ms) slices

Trigger decisions are issued based on accumulations/peaks in the Hough coordinates

ARTDAQ-Nova-DDT Performance

- The Hough algorithm was implemented in the ARTDAQ framework to provide a data driven trigger (DDT) test.
- The DDT system was run on the near detector prototype data stream
  - Surface detector with zero suppression threshold set to match far detector data rates
  - Total readout ≈1/30 of far detector readout
  - 5ms (500x beam window) slices of data were processed as a single block (no partitioning)
  - Goal was to achieve 60ms decision
- Real performance without parallelization, partitioning or algorithm specific optimization:
  - 98ms to decision
  - 2% of processing time spent in framework overhead
  - CPU utilization was $1/16^{th}$ of available

Positions us well for larger scale tests and algorithm development

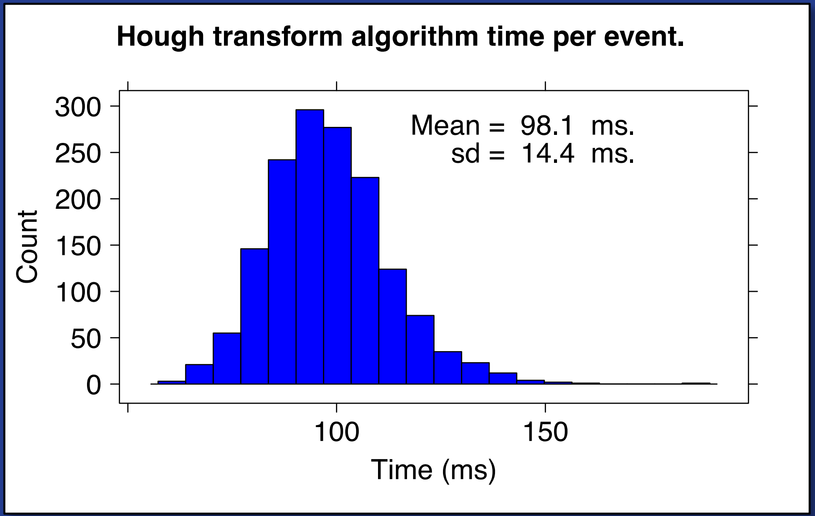

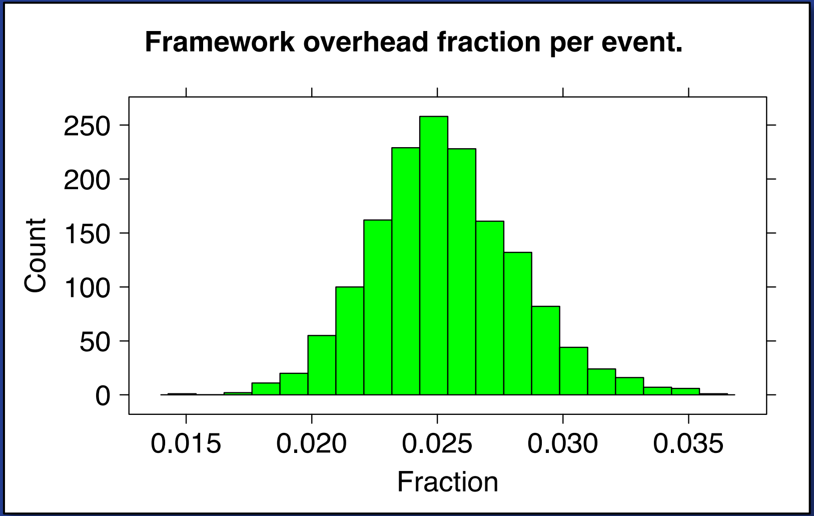

Things I didn't cover…

- Predicting & Detecting correlated faults across systems (see next talk by Qiming Lu)
- Database systems for accelerator/beam information (see poster by I.Mandrichenko)
- Data handling and cataloging (see poster by A.Lyon)
- Message passing & State machine architectures for controlling 400+ independent event builders (ask me!)

The Future

- We have built the foundation for a unique free-running DAQ system
- We have solved the problem of correlation and synchronization
- We understand how to perform continuous readout and buffering
- We are exploring the power and opportunities of Data Driven Triggering

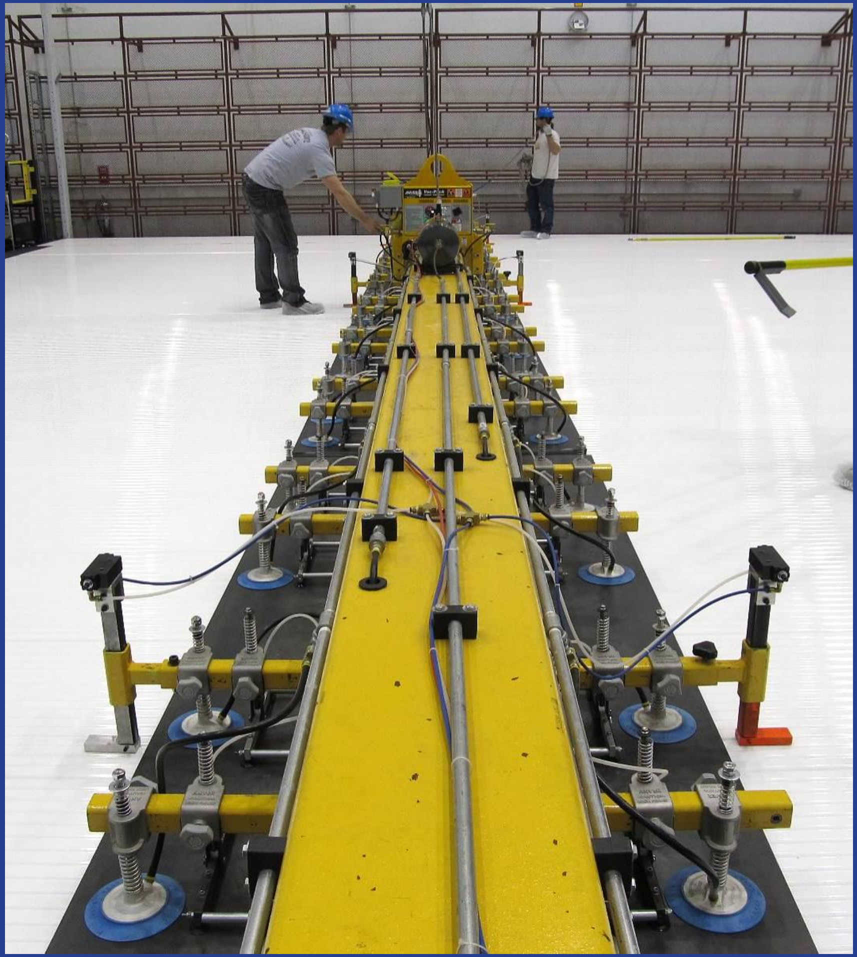

5/15/2012: Detector plane 1 dry stack