



# **NOvA Computing**

FIFE WG Meeting

**Gavin S. Davies**

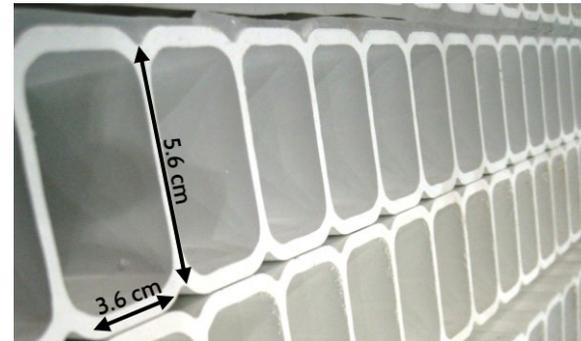
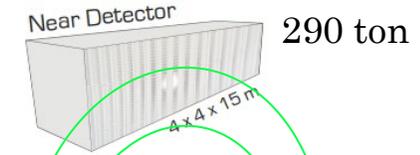
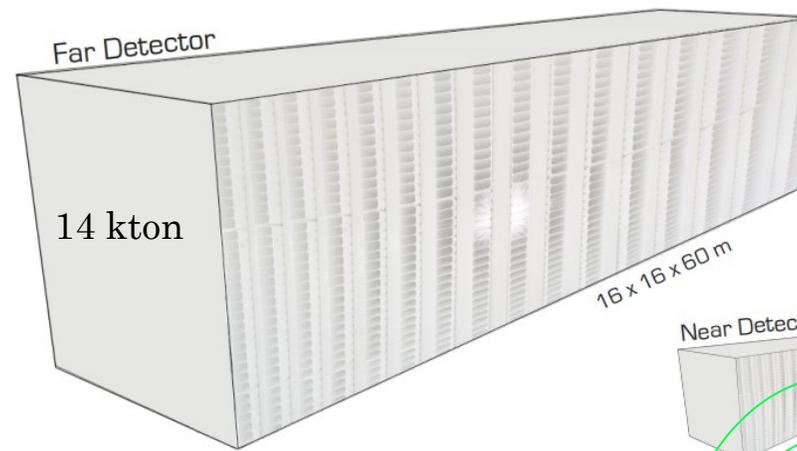
University of Mississippi

**On behalf of the NOvA Collaboration**

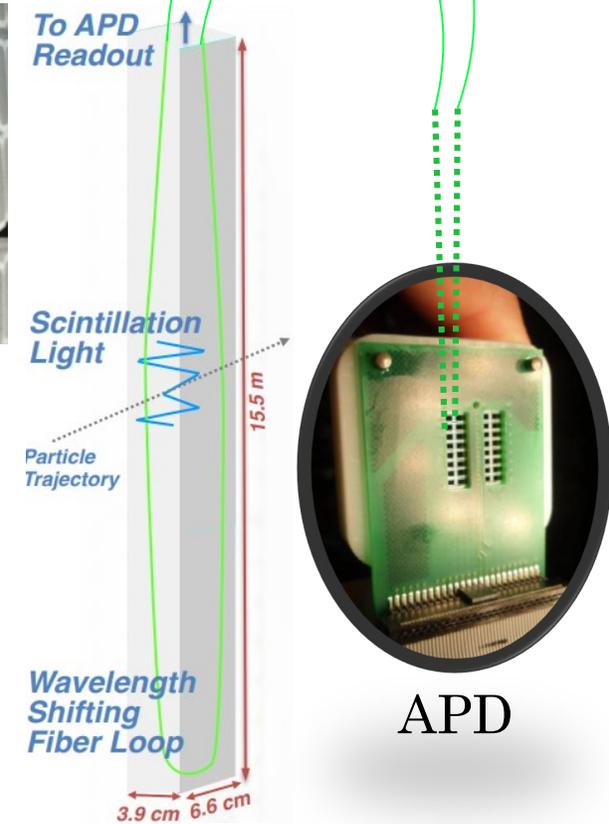
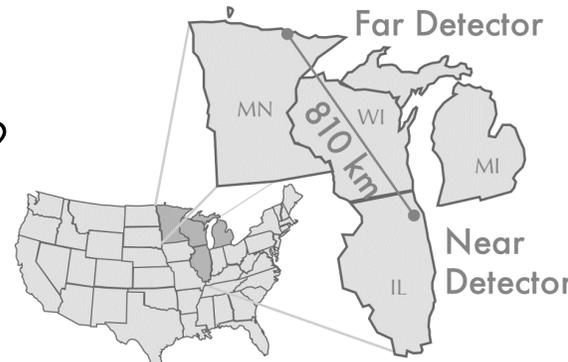
**August 4<sup>th</sup>**

# NOvA

- ❑ Long-baseline neutrino oscillation experiment
  - ❑ High power/high purity neutrino and antineutrino beams from Fermilab's NuMI facility
  - ❑ At 14 mrad off-axis, energy peaked at 2 GeV
  - ❑ 2 functionally identical detectors separated by 810 km
    - ❑ ND on-site at Fermilab
    - ❑ FD in Ash River, Minnesota
- ❑ NOvA addresses many compelling questions surrounding the nature of neutrino mass
  - ❑ What is the Neutrino Mass Hierarchy?
  - ❑ Is there CP symmetry violation in neutrinos?
  - ❑ Is there more to it than 3 x 3 PMNS



PVC extrusions (15% TiO<sub>2</sub>) filled with liquid scintillator (mineral oil with 5% pseudocumene)





# NOvA @ Neutrino22 and ICHEP22

[https://neutrino2022.org/program/detail\\_program?day=3](https://neutrino2022.org/program/detail_program?day=3)

J.Hartnell (Sussex): Neutrino 2022  
“New Results from the NOvA Experiment”



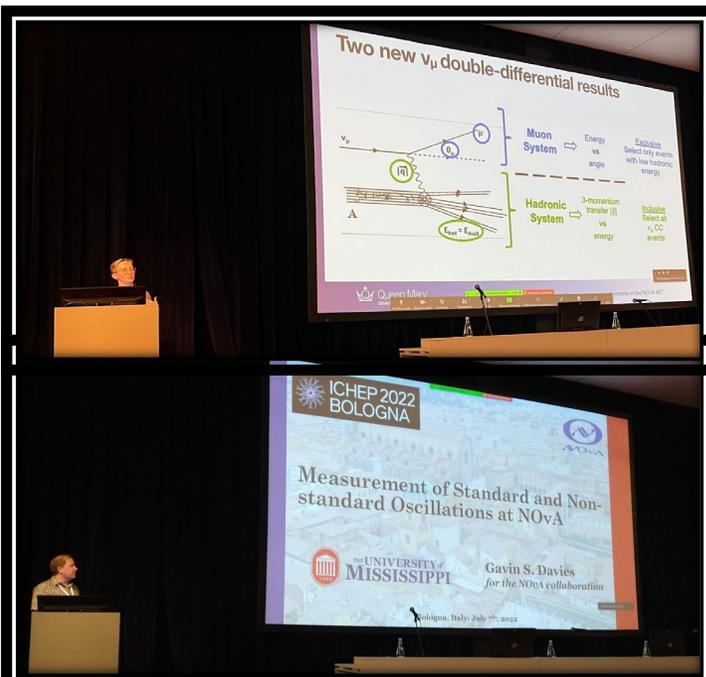
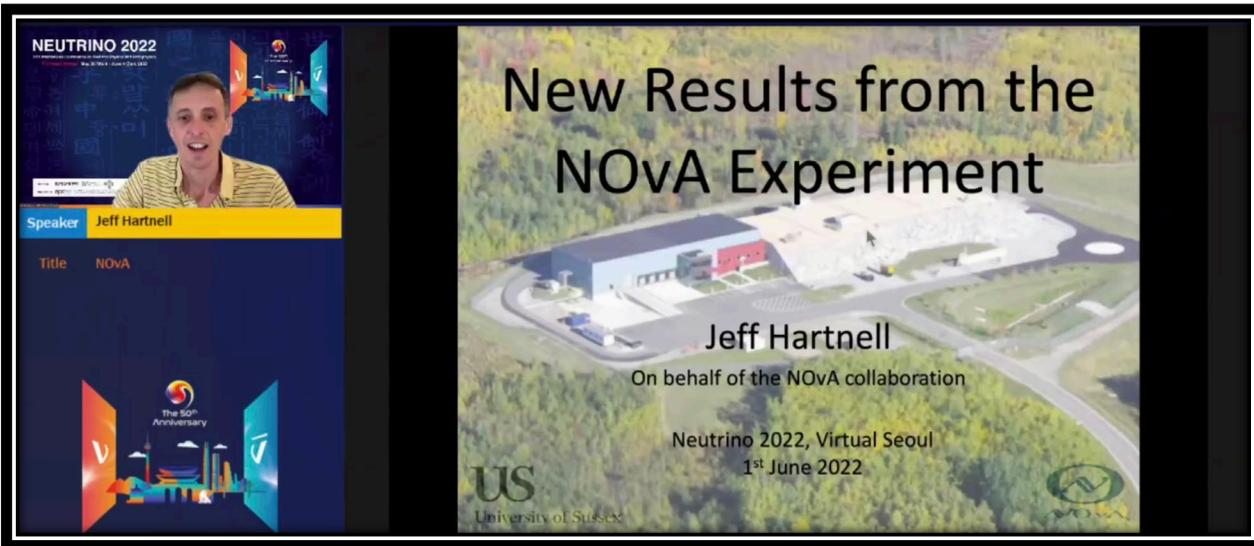
Significant NOvA data-theory discrepancies in scattering measurements

PMNS oscillation model holding up to increased scrutiny

- No sign of sterile neutrino
- NSI do not improve description of data
- Good agreement with other PMNS measurements (T2K; reactors)

Plethora of new NOvA results around the corner!

- About 50% of expected data collection still to come



L. Cremonesi (QMUL): ICHEP 2022  
“Neutrino scattering measurements at NOvA”

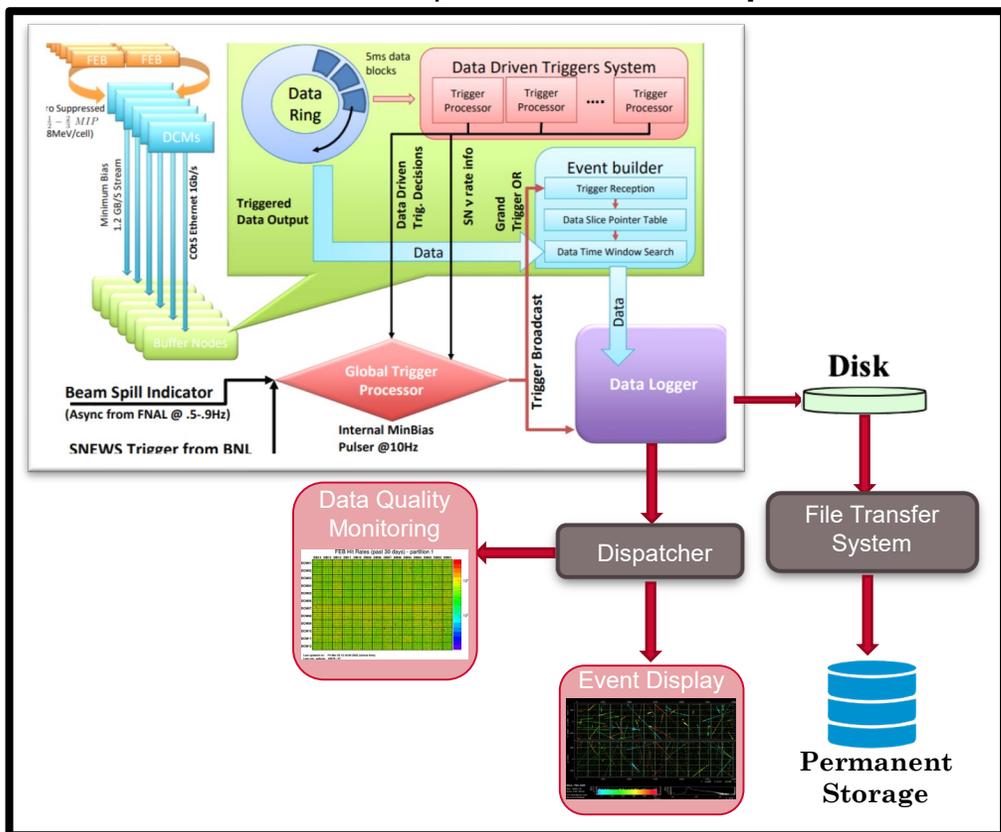
<https://agenda.infn.it/event/28874/contributions/169658/>

G. Davies (U. Mississippi): ICHEP 2022  
“Measurement of Standard and Non-standard Oscillations at NOvA”

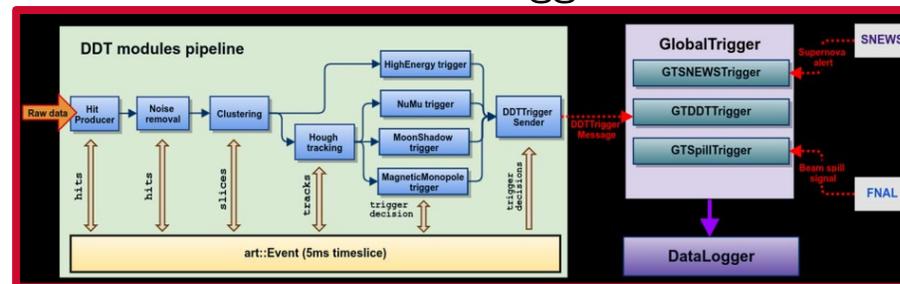
<https://agenda.infn.it/event/28874/contributions/169624/>

# NOvA framework

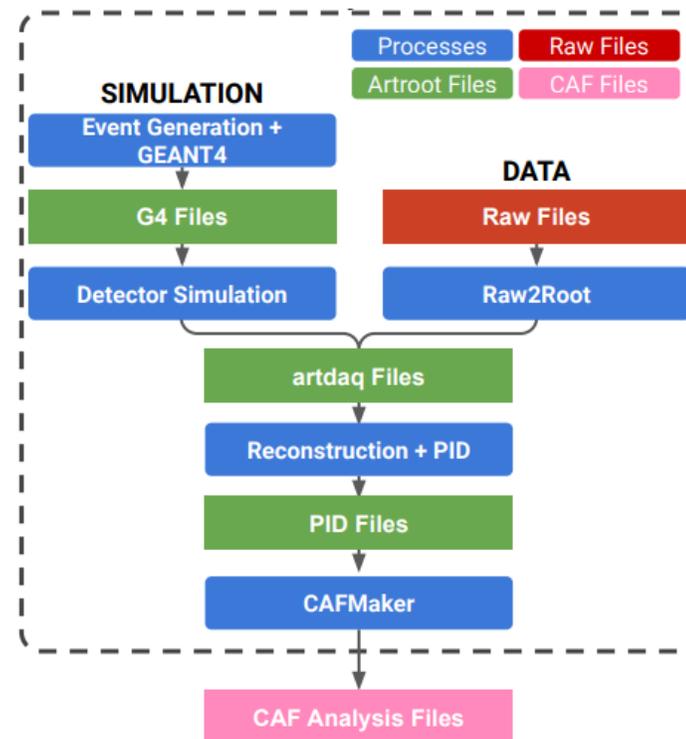
## NOvA Data Acquisition (novadaq)



## NOvA Data Driven Triggers (novaddt)



## NOvA Offline Software (novasoft)



NOvA ensemble is three main systems:

NOvA Data Acquisition (**novadaq**)

NOvA Data Driven Triggers (**novaddt**)

NOvA Offline Software (**novasoft**)

} All built around **art** framework

# NOvA Offline Software

novasoft **art** instance used for simulation, calibration, reconstruction, particle identification

Hosted on **github** repository

- ❑ main always builds. Individuals work on branches
- ❑ Codeowners philosophy assigns individuals to various parts of the codebase
  - ❑ If a pull request (PR) touches ones package a review from individual(s) is requested in order to approve the PR

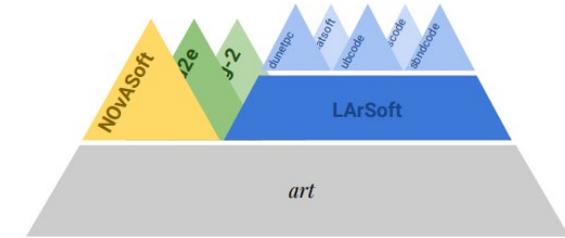
novasoft builds handled by Software Release Tools (SRT) [1], unix based software management system (NOvA also supports a Cmake/mrb build system)

- ❑ Uses a “test release” philosophy; just have the one or two pkgs in working area
- ❑ GNUmakefiles handle building against full software stack and shared libraries
  - ❑ Have to be careful about circular dependencies
- ❑ Partial, sparse clone for github cloning of repository in tandem with Git LFS

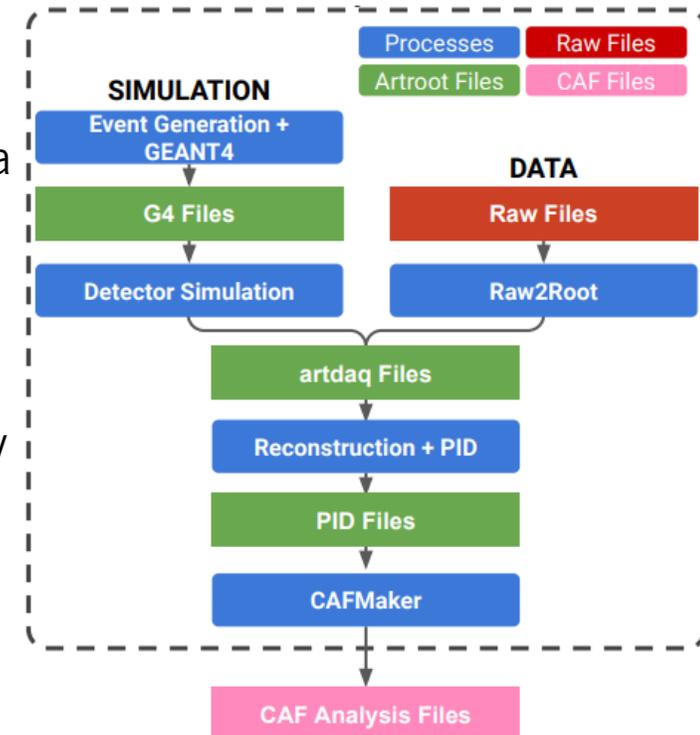
All releases of novasoft hosted as an OSG CernVM File System (cvmfs) repository

External product (ROOT, art etc) access via Fermilab-developed Unix Product Support (UPS) and Unix Product Distribution (UPD)

- ❑ Supports multiple versions of a product and build-types per version



## NOvA Offline Software (novasoft)



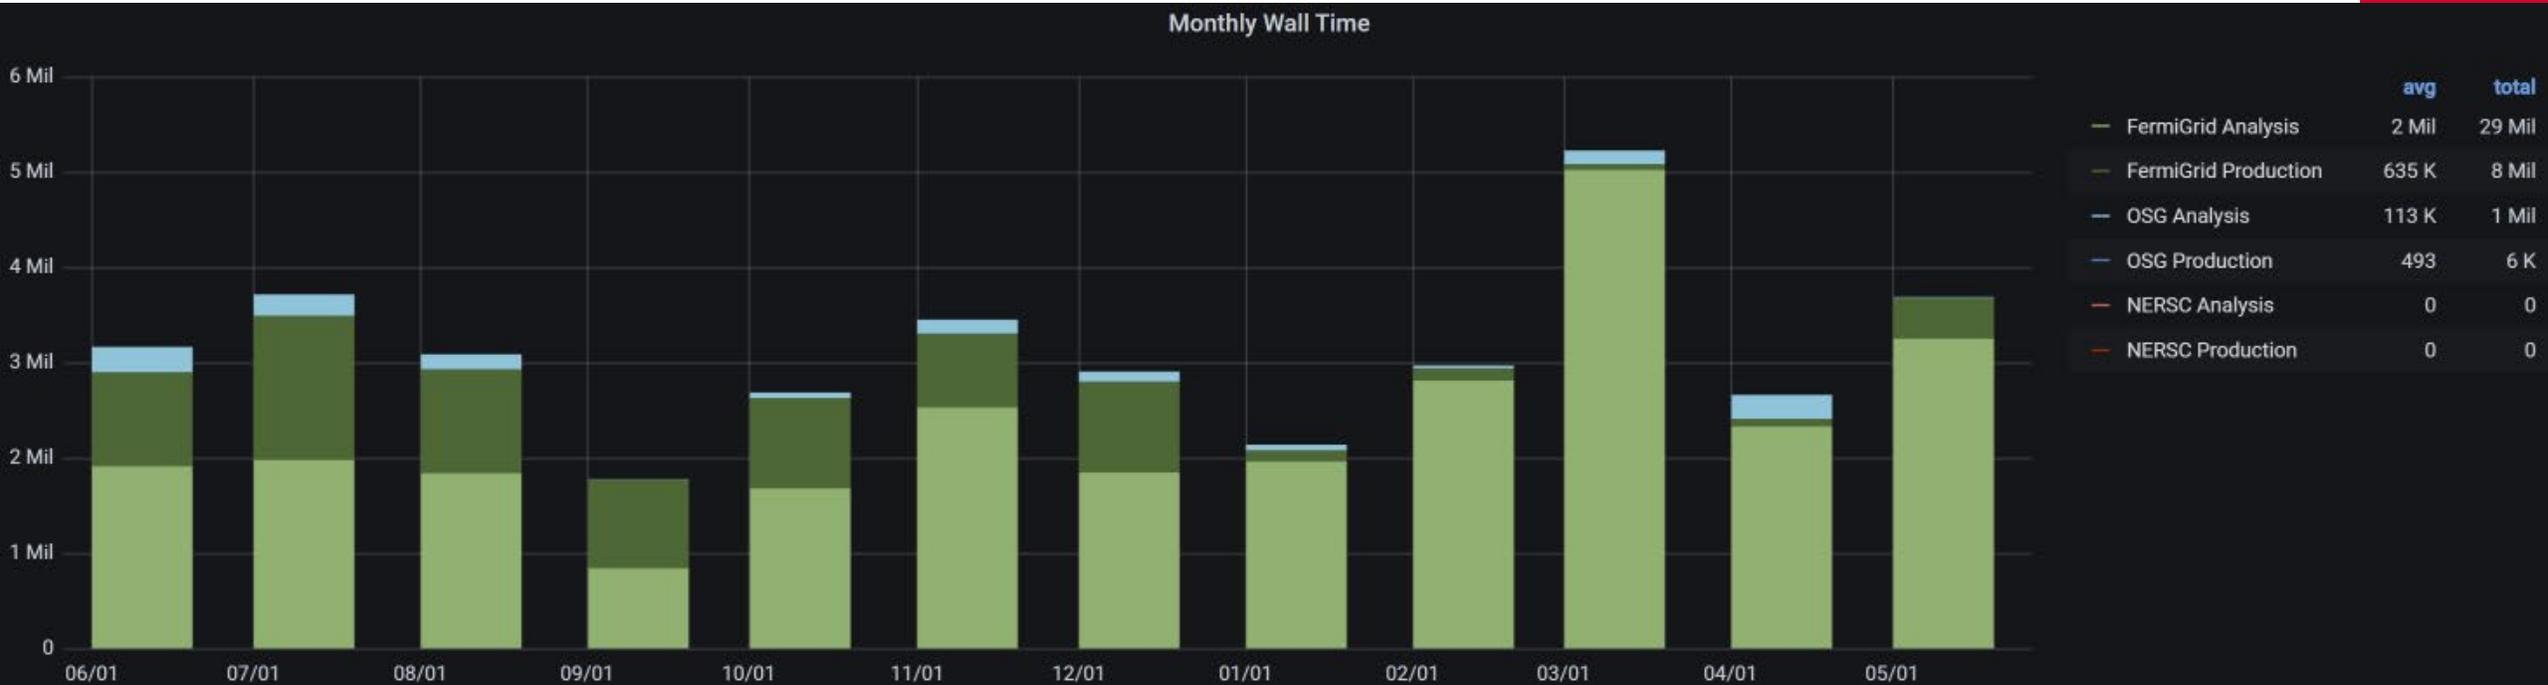
[1] J. Amundson, CHEP 2000: <https://inspirehep.net/literature/538395>

# CPU: Experiment Usage past year

Security changes severely hampered offsite submission ability; ~ 1 year to recover.

NERSC: ~ 90 Million hours; not shown on FIFEmon

@ALCF used ~ 160 GPU node hours (tried using OSG, not very successful: availability/quality too low)



Able to run offsite again. Last lingering issue was our CAFAna analysis jobs

Thanks, Marc Mengel : Helped discover that we need to clear CPATH and LIBRARY\_PATH env. vars.

(bleeding into container environment).

# NOvA Builds, CI, FTS

novasoft release builds handled by Jenkins  
Built with parallelism  
Nightly builds of "main"



✓	⚙️	novasoft-git_SRT_slf7_nightly_build	1 day 0 hr - #292	2 mo 12 days - #221	1 hr 56 min
✓	⚙️	novasoft-git_SRT_slf7_nightly_build_output	22 hr - #231	N/A	10 min
✓	⚙️	novasoft-git_SRT_slf7_release_build	4 days 13 hr - #52	1 mo 17 days - #45	1 hr 45 min
✓	⚙️	novasoft-git_SRT_slf7_release_build_output	4 days 11 hr - #22	N/A	8 min 15 sec

Run a series of continuous integration tests for Jenkins builds

Utilize Fermilab's Continuous Integration Project system which provides integration builds of software code and reports on current and past builds, with detailed information on pre- and post-install tests of the software

Build ?	Start Time ?	Build Type ?	checkout ?	build ?	ci_tests ?
<a href="#">nova_ci/13296</a> (NOvA)	2022-05-23 17:15:06	slf7 maxopt	✓	✓	✗

**In progress:** Github web-hooks allow to additionally trigger builds with PRs (with assistance from Vito)  
Additionally, updating CI jobs

**Need:** FTS build with python3

Example of failing CI

# NOvA Production

NOvA follows a traditional HEP computing strategy where event-driven analyses access centrally-stored data and Monte Carlo simulated events by distributing data to parallel jobs running on grid computing clusters.

Relies heavily on ROOT and GEANT4.

**novaproduct** github repository for custom python-based production tools for interfacing with Fermilab's FIFE-managed services for:

- job submission (**jobsub**)
- monitoring (**fifemon**)
- data access/cataloging (**SAM**)
- data transfer (**ifdhc**)

Provides a simple user interface for command-line or configuration file submission

Users/Production specify dataset, number of jobs, container, onsite vs. offsite etc.

Allows for manager control of defaults (for example, specifying recommended offsite locations)

Require production jobs to be exactly reproducible, exact same random number seed and same input files.

Achieve by storing one fhicl file for each job submitted (for simulation)

# Cosmic Rejection

Fully processing the currently recorded cosmics would take  $> 2$  years of continuous running

Considering that only 1 in  $10^7$  cosmics will be background candidates to the oscillation analysis

Therefore, we would like to run a pre-reconstruction filter (neural network-based) to remove the most obvious cosmics

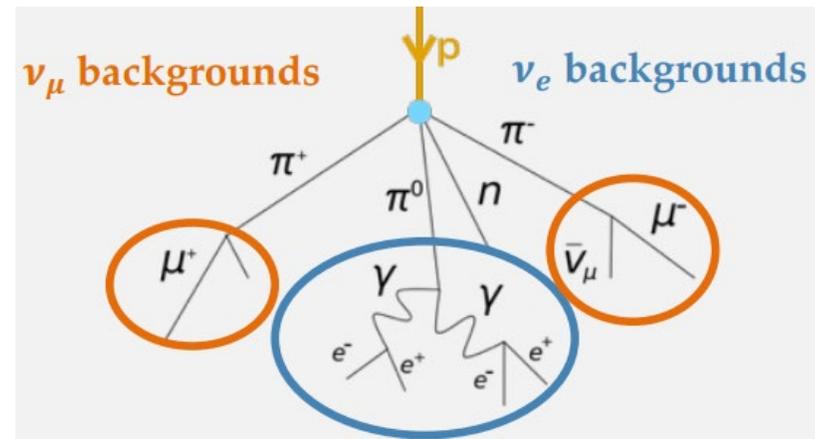
- ❑ This network should have no software or reconstruction dependencies, so that in the future we are free to change, for example, our calibration tunes or clustering algorithms

Network imagined as a “one-pass” network

Leverage NOvA’s excellent timing resolution

- ❑ Split the  $550 \mu\text{s}$  readout window into  $18 \mu\text{s}$  “time slices”.
- ❑ Only process time slices where there is a neutrino candidate

The network uses a ResNet-18 backbone, with a Siamese structure, taking two pixel maps as input (one in the XZ view, and another in the YZ view)

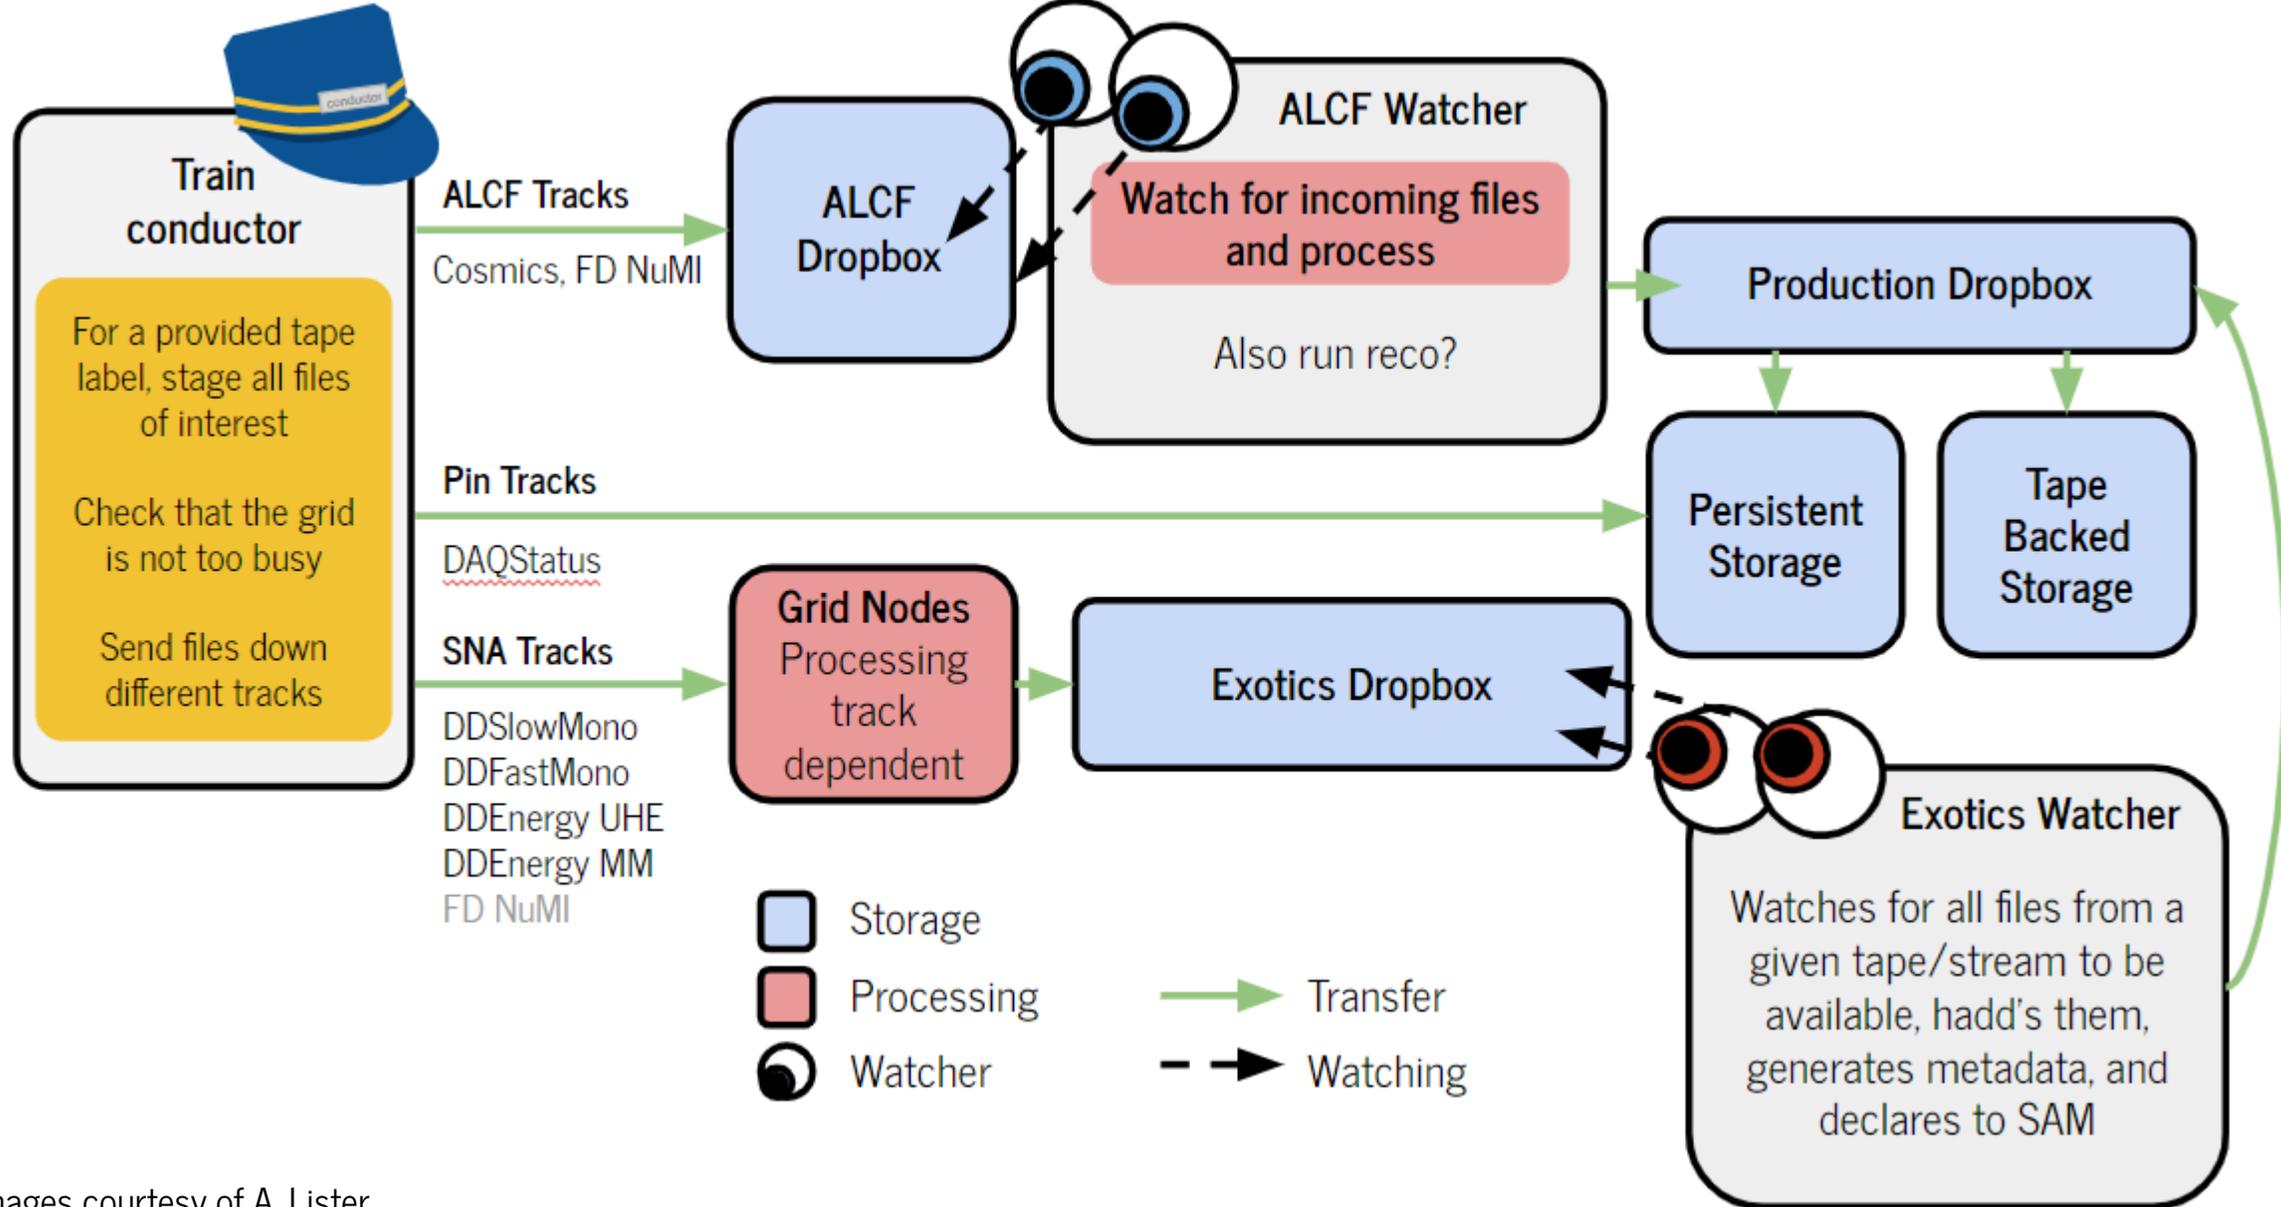
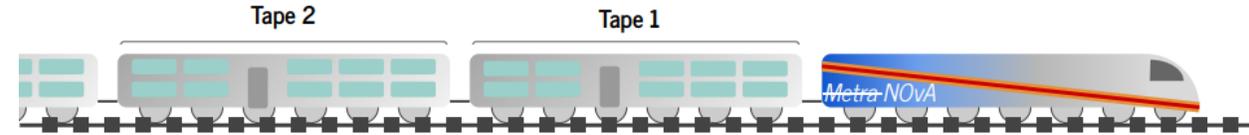


The Neural Network outperforms the traditional cosmic rejection algorithm in all samples

Data Sample	Traditional Cosmic Rejection (%)	Cosmic Rejection Neural Network (%)
$\nu_e$	93.21	99.71
$\bar{\nu}_e$	92.81	99.82
$\nu_\mu$	93.22	99.20
$\bar{\nu}_\mu$	92.82	99.20
$\nu$ NC	93.24	97.08
$\bar{\nu}$ NC	92.79	96.82
Cosmic $\nu$	7.80	5.00

The challenge is processing the large volume of data on tape – but doing once and only once  
 $O(1 \text{ million})$  files per data stream (7 streams to process)  
Go tape-by-tape, run network at ALCF

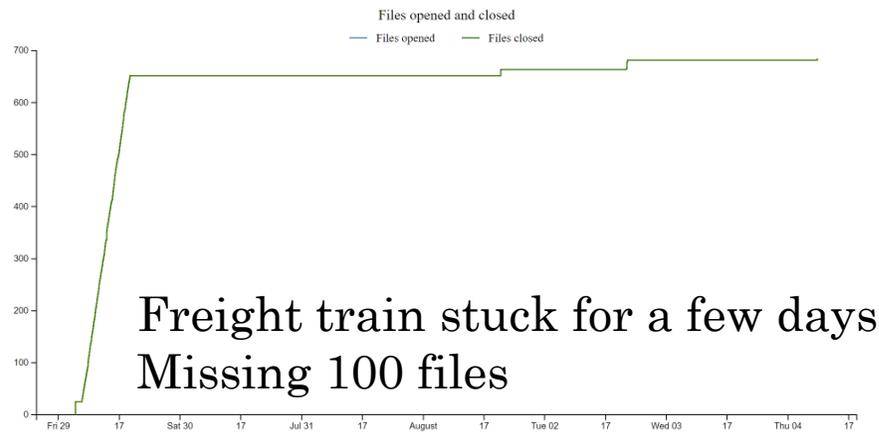
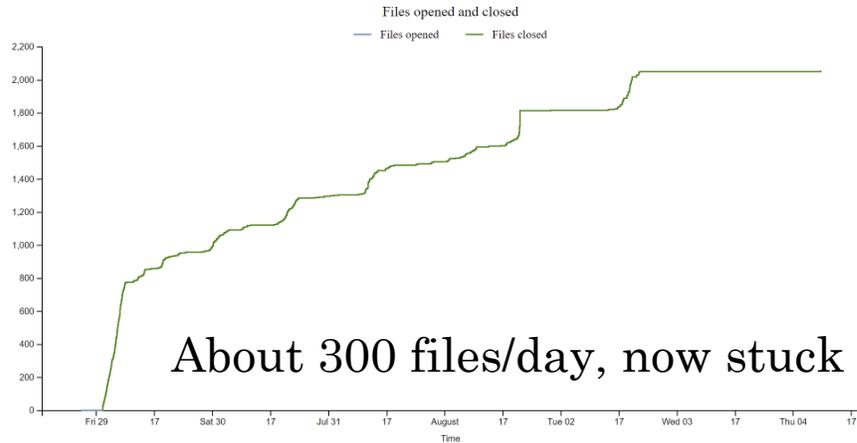
# NOvA Freight Train Model



Images courtesy of A. Lister

# Prestaging

Varying performance prestaging from tape



Generally about 200-300 files per day  
Open communication with tape/storage folks

Mount one tape, don't dismount until all files from tape are staged  
In practice, unclear if it is worth prestaging one tape at a time

Additionally, reported a broken SAM command:

```
SAM = samweb_client.SAMWebClient("nova")  
SAM.listFilesAndLocations(defname=defi)
```

will give wrong results if the definition is built with *with stride* or *with limit*

# CPU Adaptations Going Forward

Freight train scheme in place for filtering raw data; access once and never again.

Unique data transfer challenge between FNAL and ALCF

HEPCloud jobs at NERSC interest for HPC GPUs notably with Perlmutter coming online and successes already shown, including at ALCF.

But SCD support needed: Examples, documentation, likely to personnel time to aid integration with existing workflow (limited by our own FTE availability too)

With large-scale GPU availability, opens up additional avenues for R&D reconstruction work planned.

# Thank you!



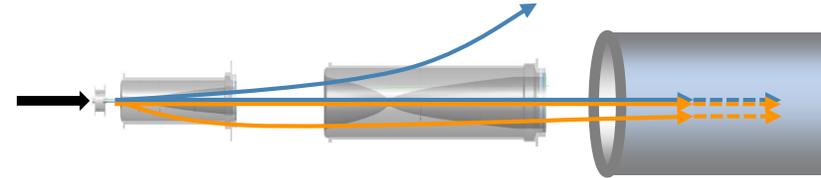
<http://novaexperiment.fnal.gov>



# NOvA Simulation

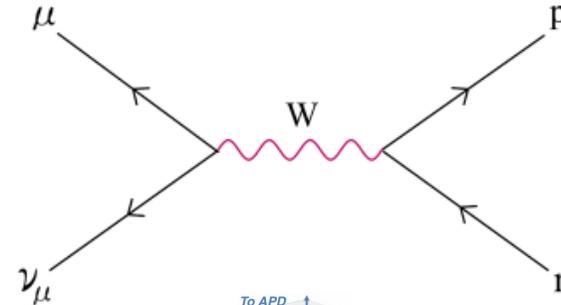
## Neutrino Flux

G4NuMI: GEANT4-based simulation  
Reweighted to incorporate external  
hadronic production data



## Neutrino Interactions

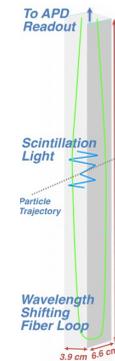
Simulated with GENIE 3.0.6  
Tuned to external data and NOvA ND Data

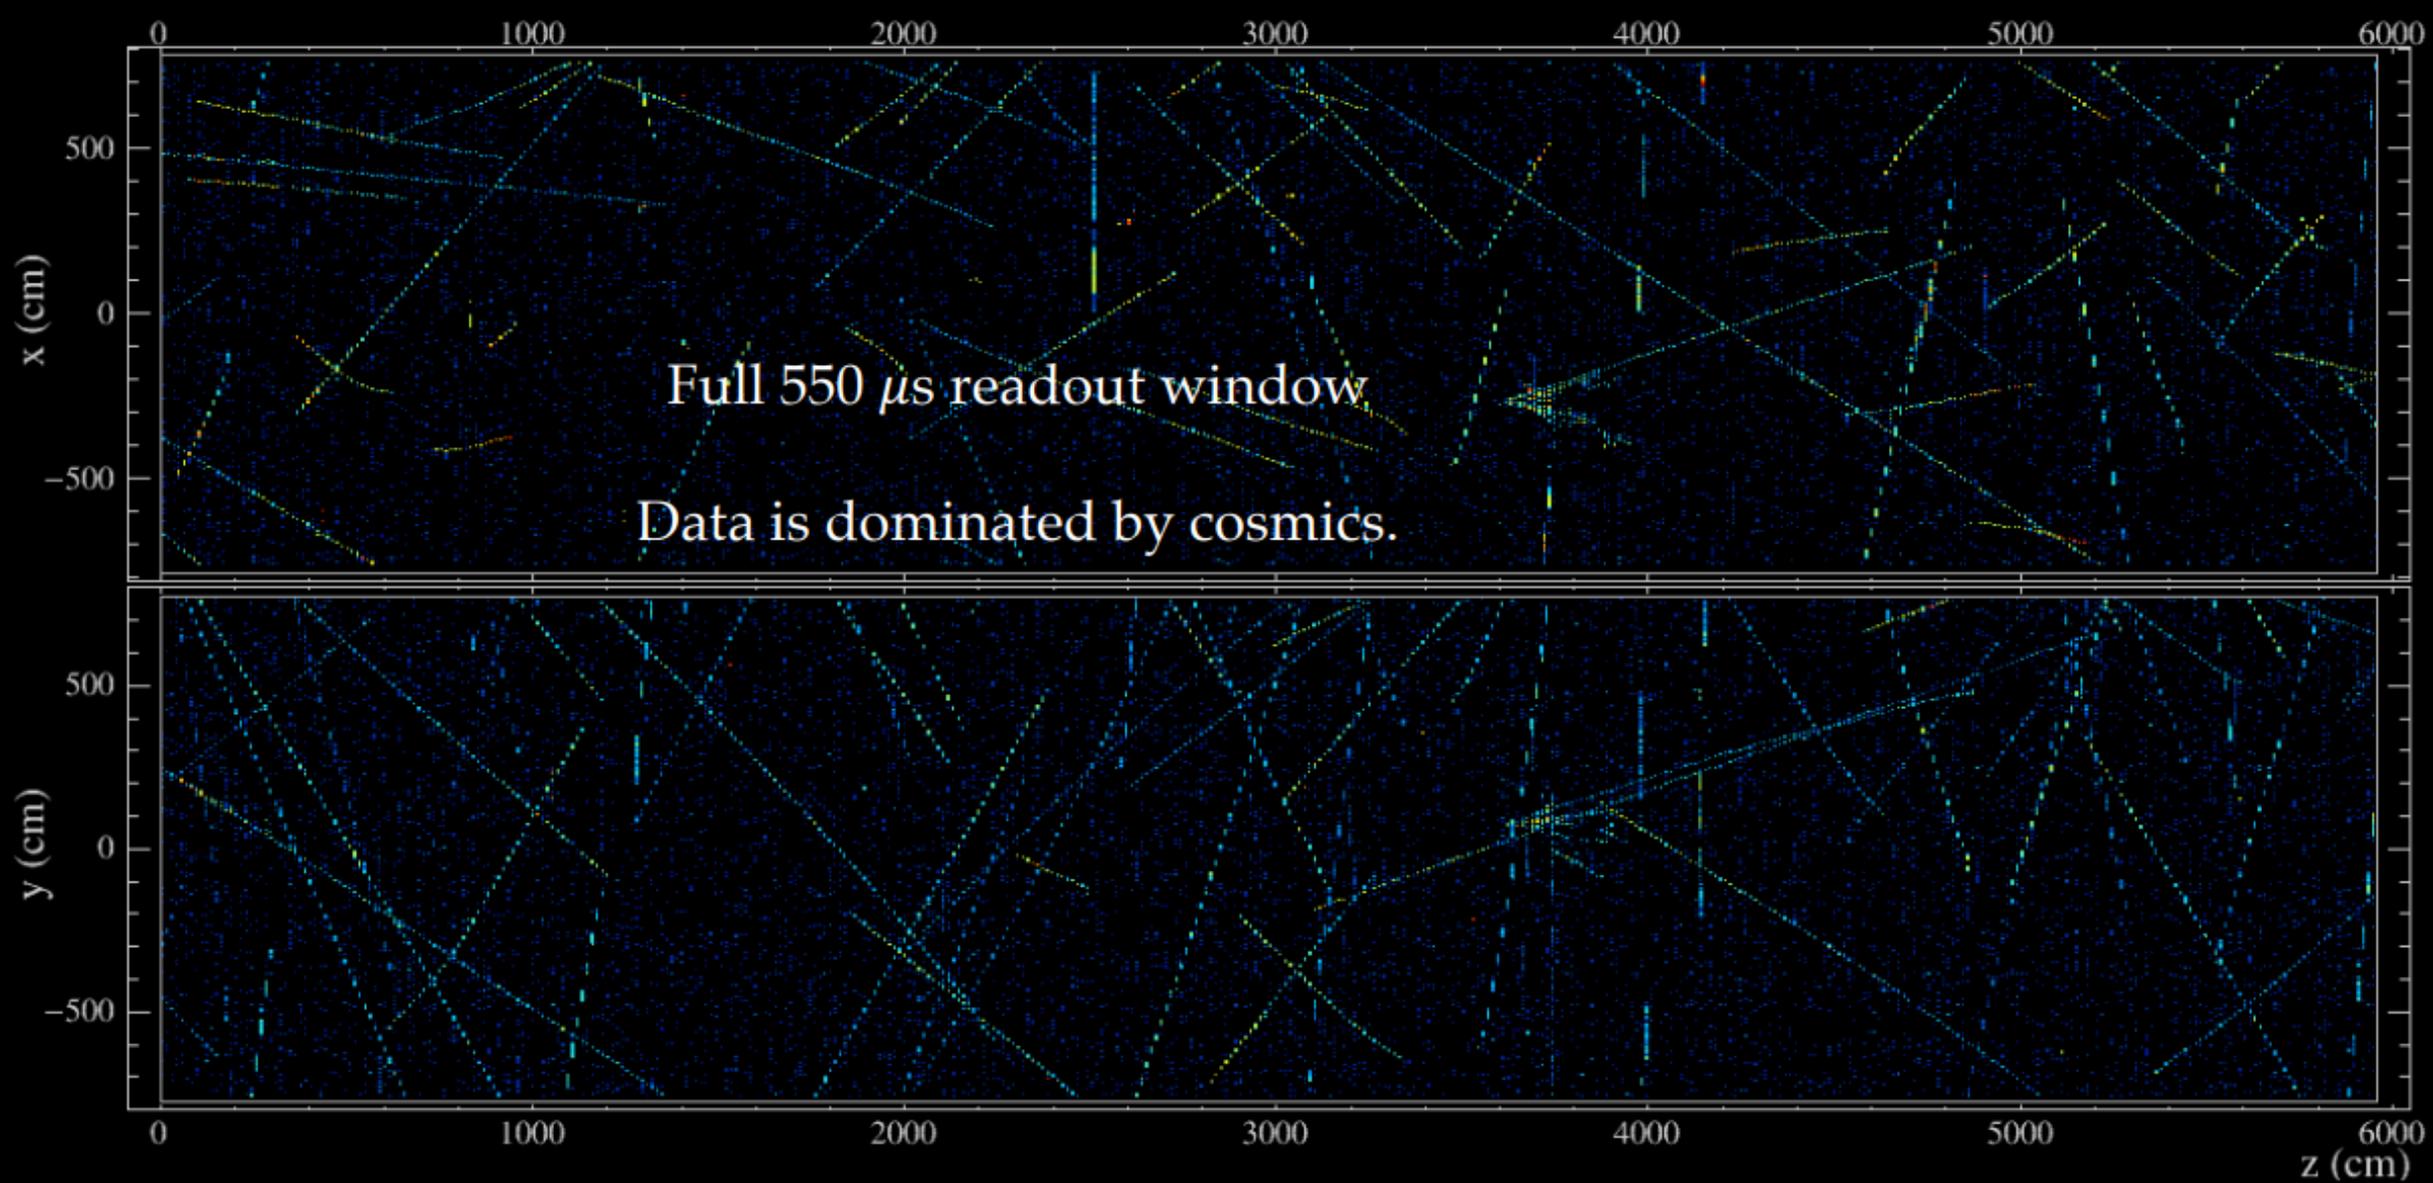


## Detector Response

GEANT4 simulation of the detector's response  
to charged particles and light propagation

In addition, we need to overlay real background data  
with “signal” information; cosmic rays in FD, rock events in ND  
Or in some cases overlay single-particle events onto nominal Monte Carlo  
Leverage art product mixing





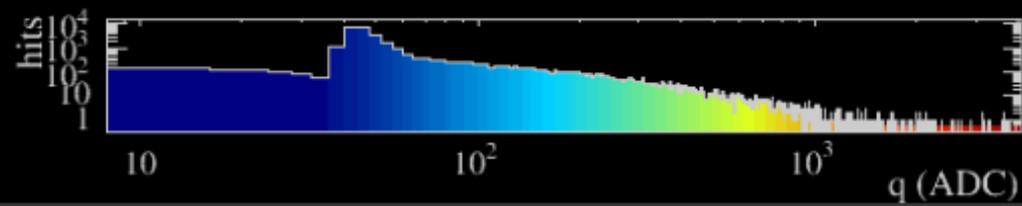
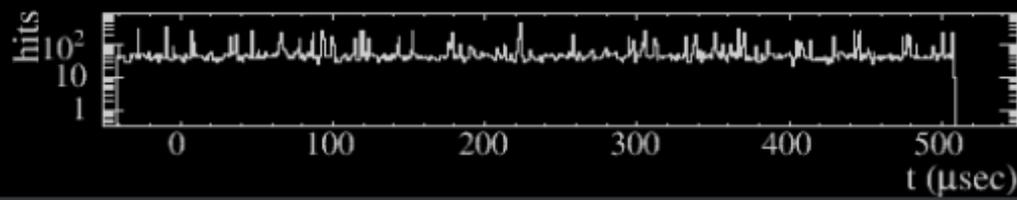
**NOvA - FNAL E929**

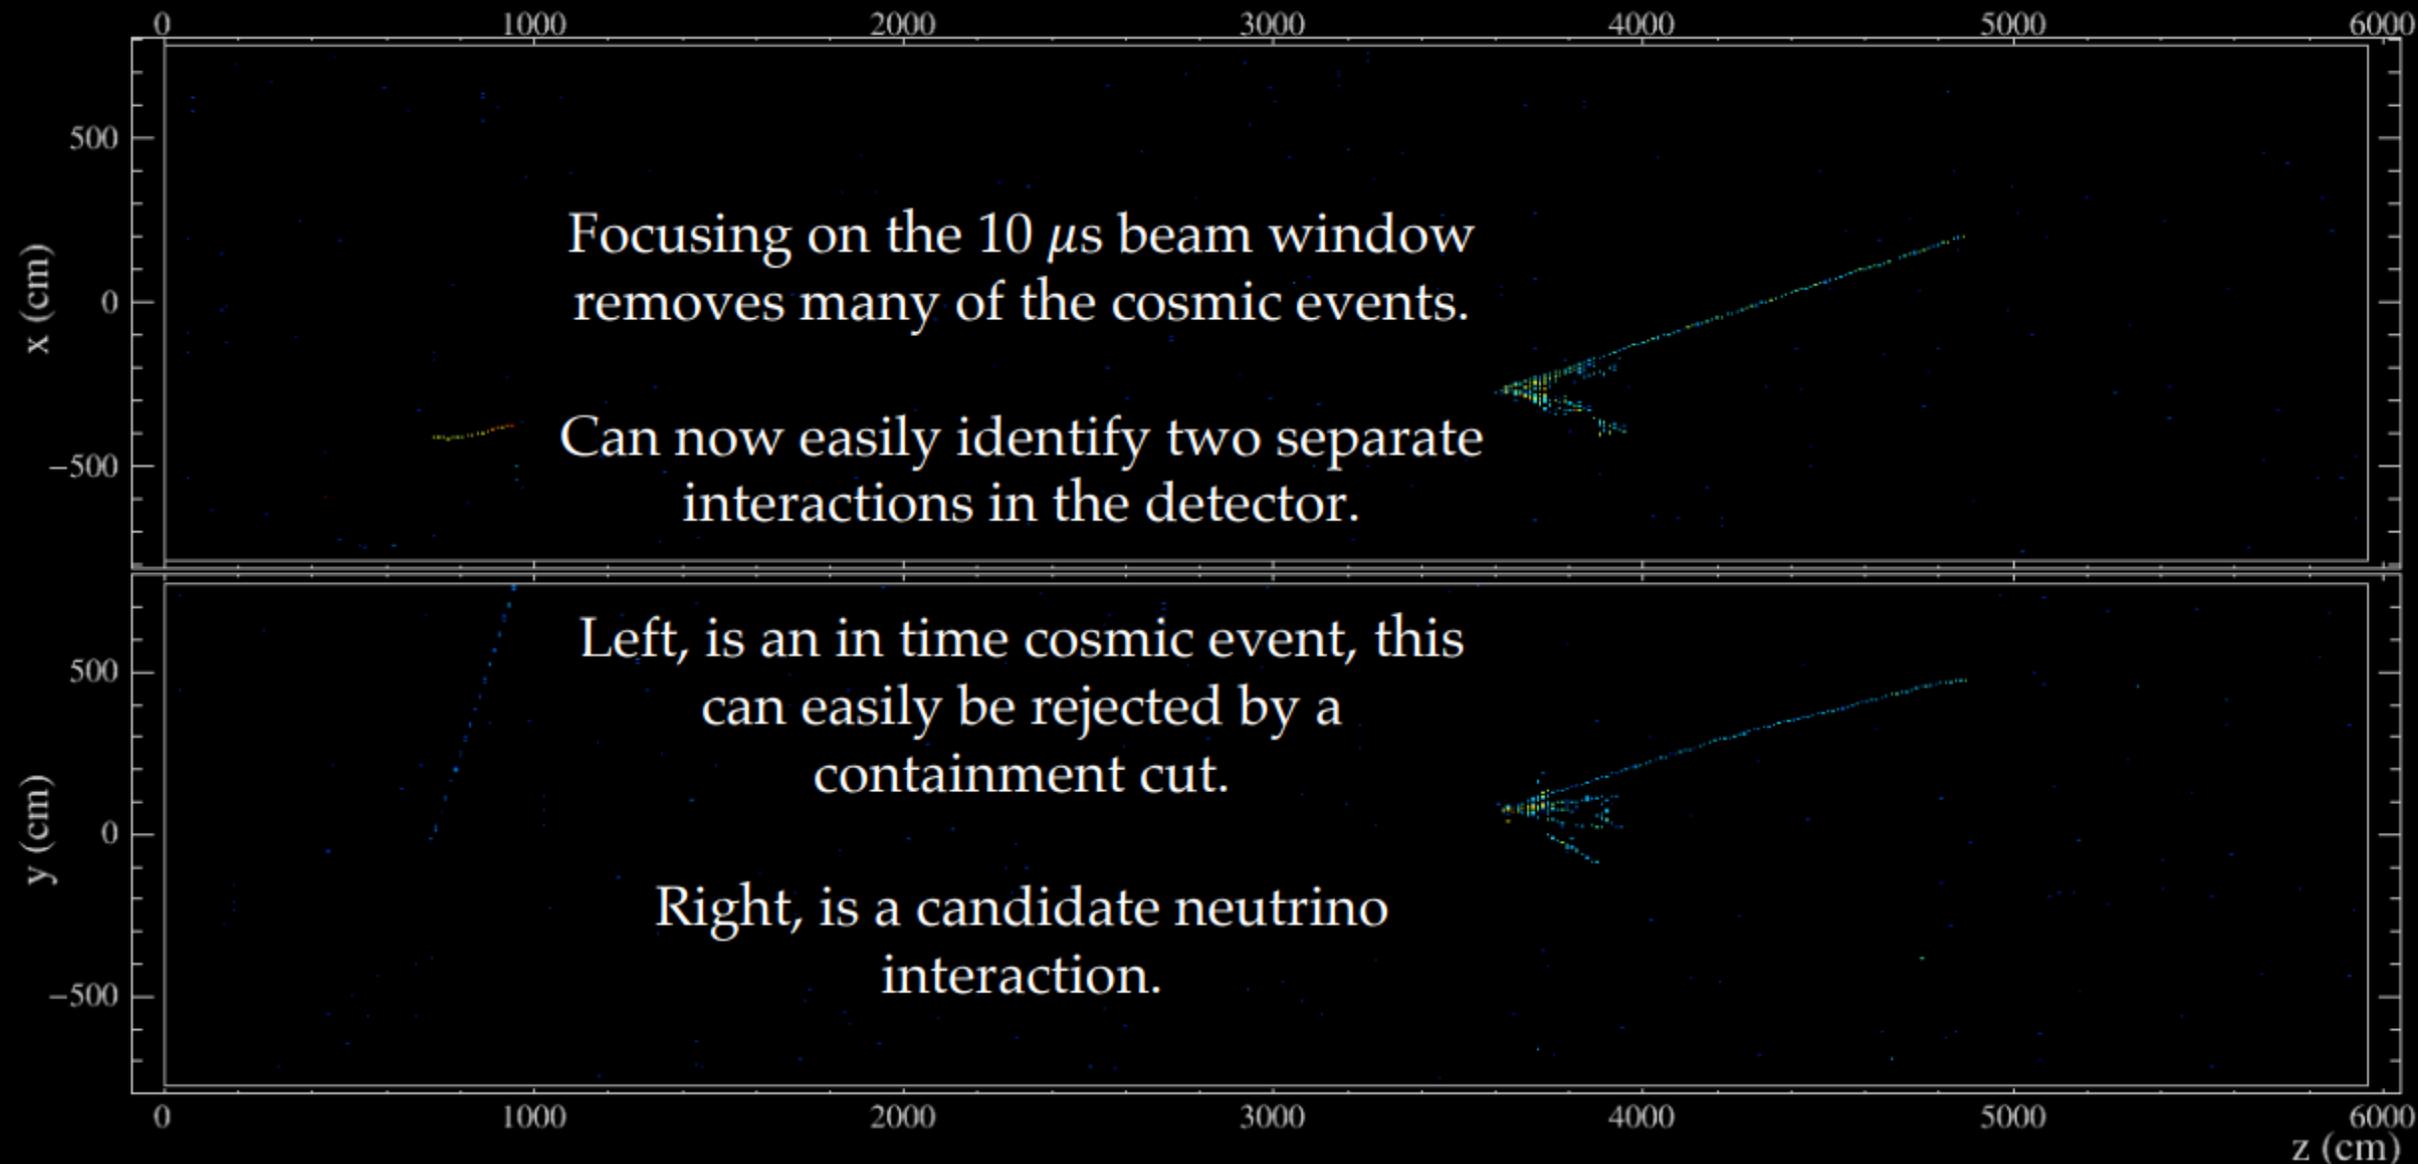
Run: 18620 / 13

Event: 178402 / --

UTC Fri Jan 9, 2015

00:13:53.087341608





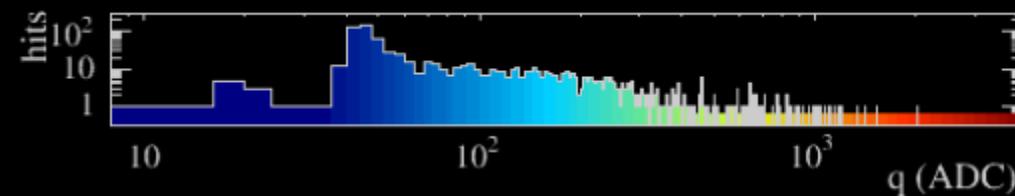
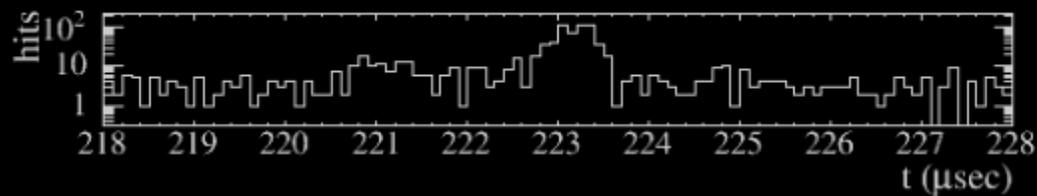
**NOvA - FNAL E929**

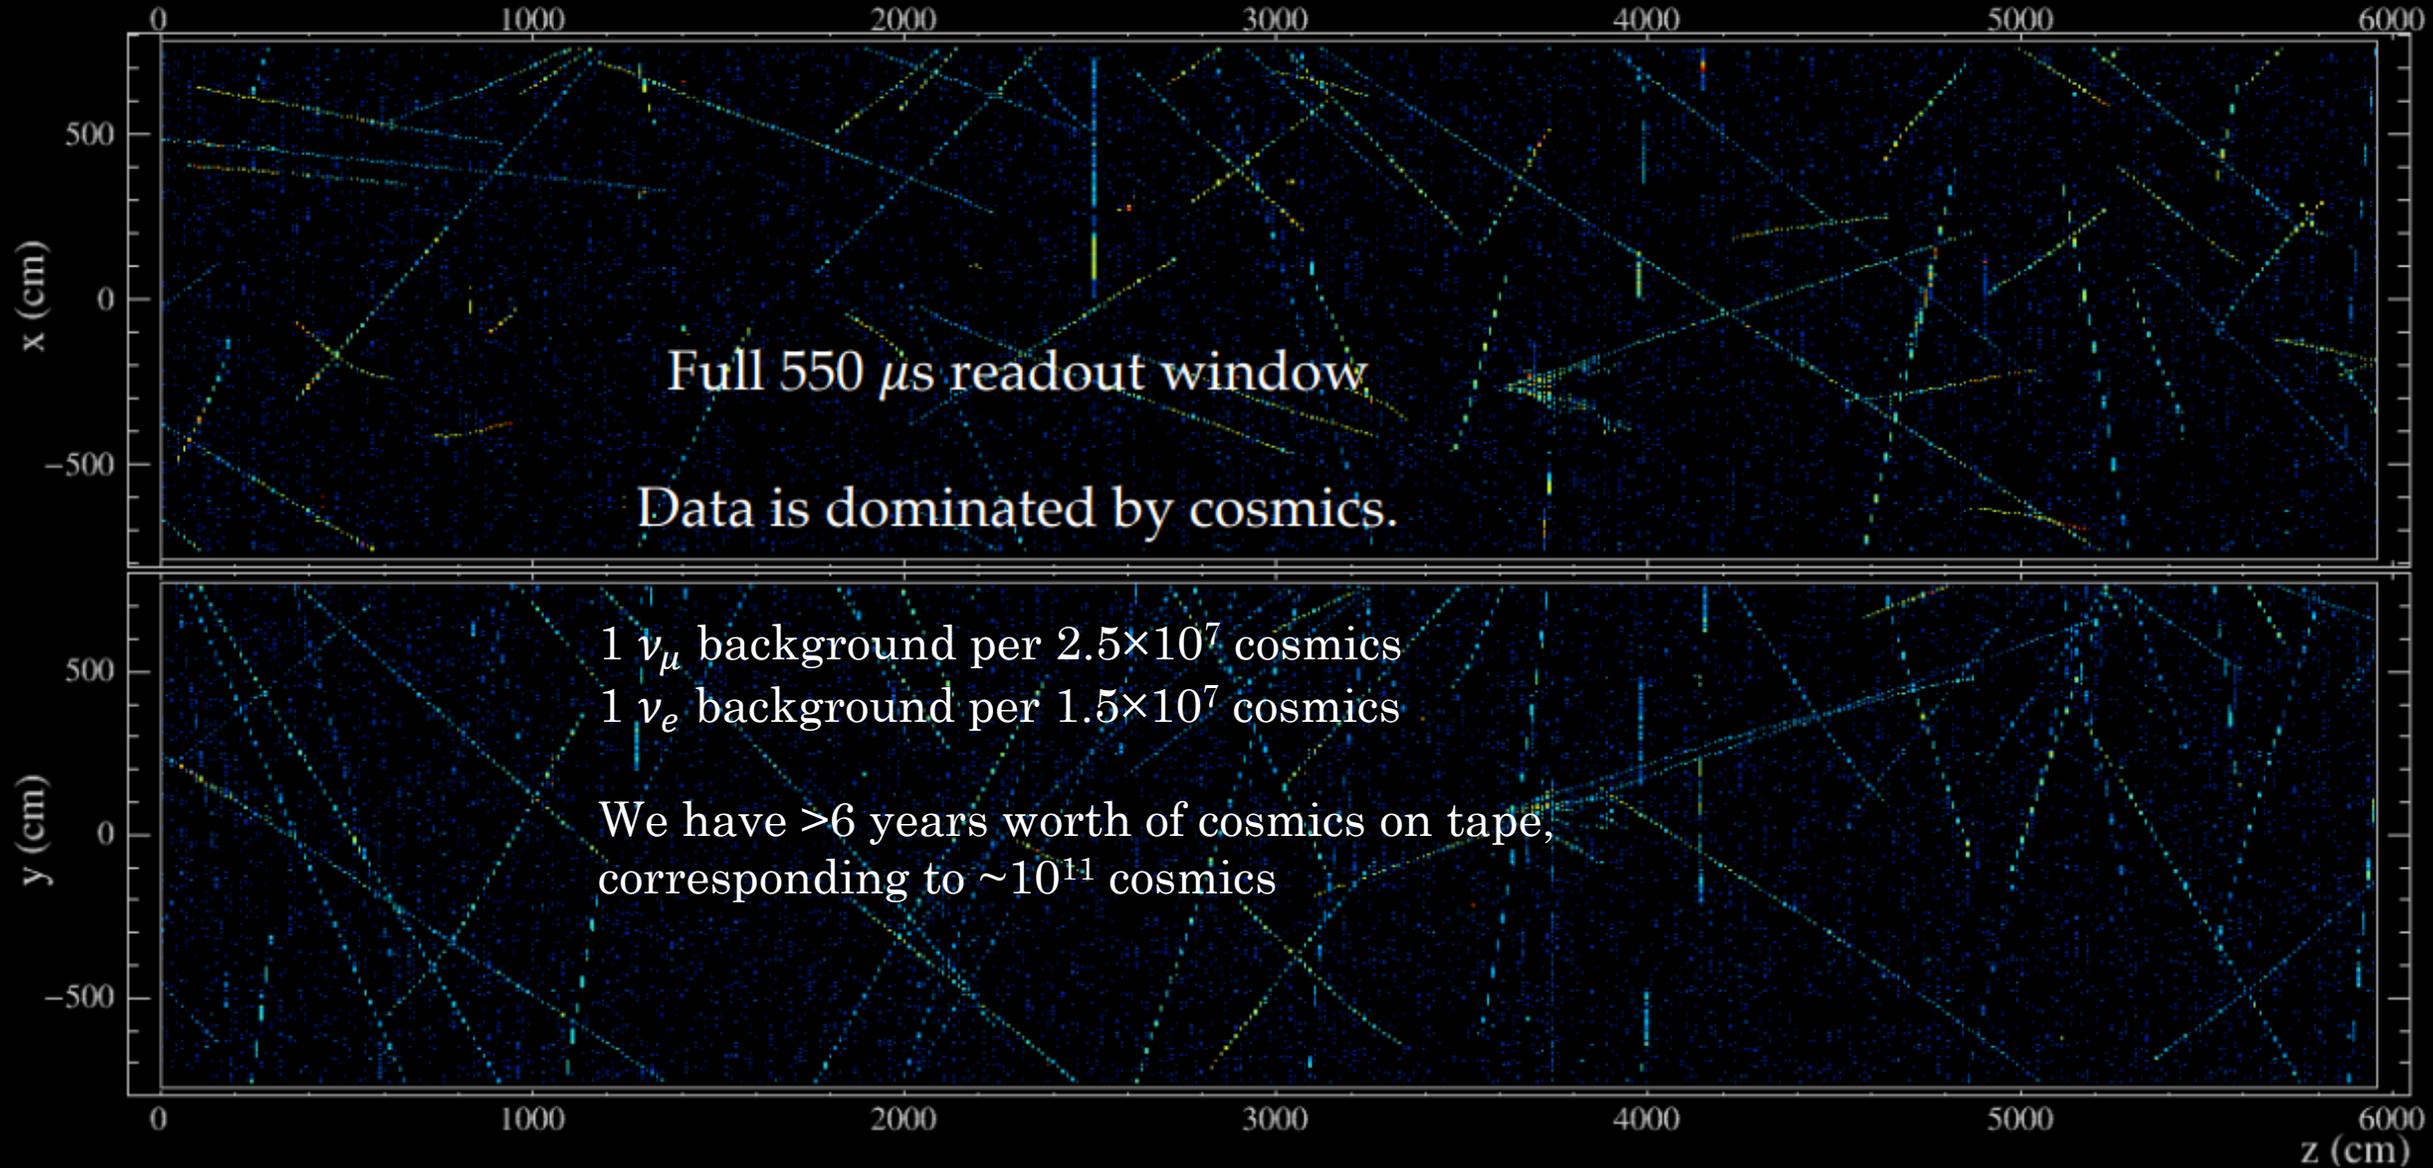
Run: 18620 / 13

Event: 178402 / --

UTC Fri Jan 9, 2015

00:13:53.087341608

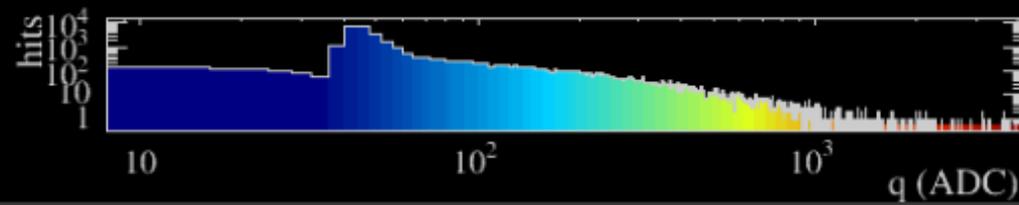
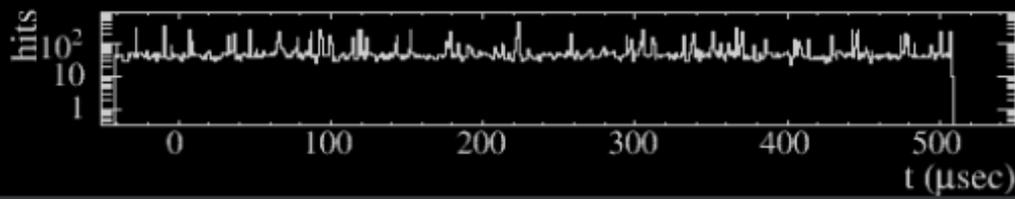




**NOvA - FNAL E929**

Run: 18620 / 13  
Event: 178402 / --

UTC Fri Jan 9, 2015  
00:13:53.087341608



# NOvA Analysis

Outputs of each production stage are stored in art-ROOT format

Stored objects are complex and difficult to read as regular TTree's

*art* analyzers are used for simple analysis.

The standard NOvA tool for analysis (originally developed for oscillation analyses) is **CAFAna** [2]

On NOvA, final stage of production is to produce CAFs (Common Analysis Files)

- ❑ Standard analysis-level file format; conventional ntuples, ROOT TTree
- ❑ Each record represents a single neutrino candidate (or "slice")
- ❑ Additional experimentation with HDF5 representation of CAFs
- ❑ Written in C++, also a python interface enabled by pyROOT

Basic CAFAna histogram object is a *Spectrum*; a series of bin contents, internally manipulated by Eigen [3]

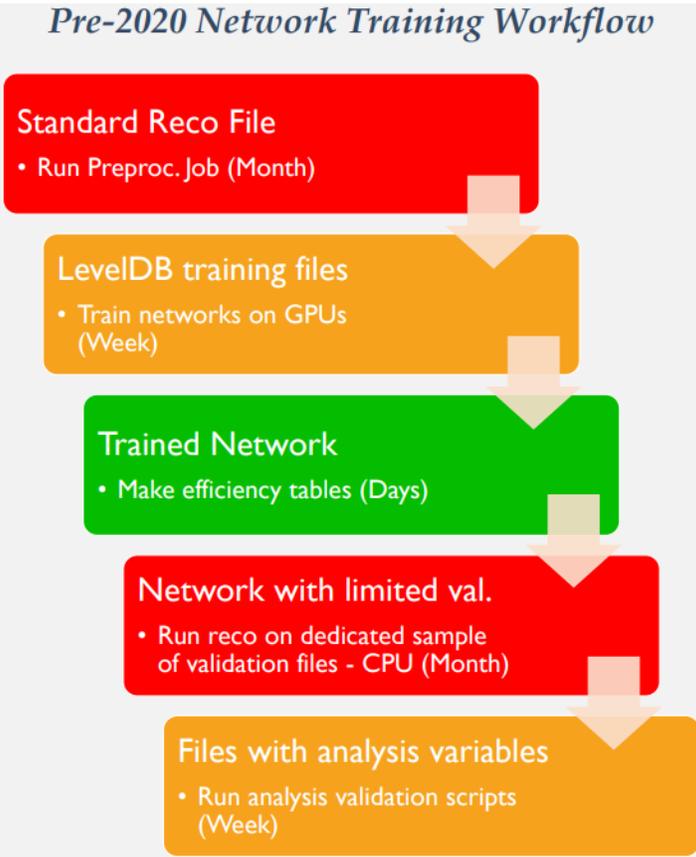
Multi-dimensional spectra supported

Fully integrated with Fermilab **jobsub** infrastructure and **SAM** data management system

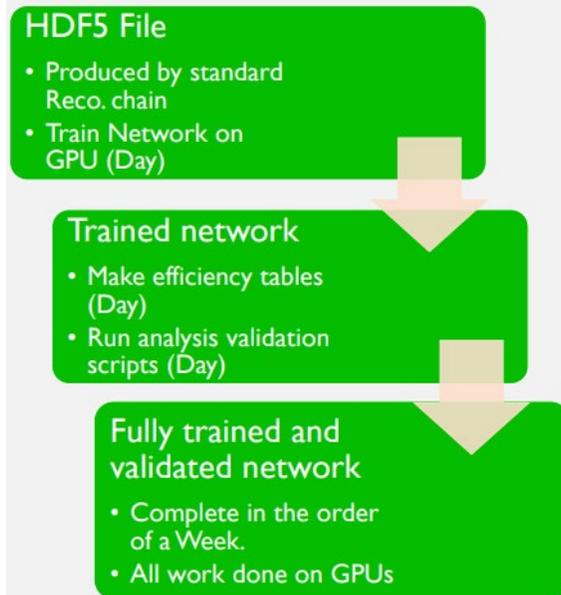
[2] The CAFAna framework for neutrino analysis, [arXiv:2203.13768](https://arxiv.org/abs/2203.13768) [hep-ex]

[3] G. Gael, J. Benoit et al., Eigen v3, <http://eigen.tuxfamily.org> (2010)

In order to pursue our physics goals NOvA has to be able to classify neutrino interactions with high efficiency and purity. Machine Learning is one of the tools that we use to do this. NOvA was the first HEP experiment to use a CNN in a physics measurement to classify candidate neutrino interactions



*2020+ Network Training Workflow*



Transitioned from training networks using LevelDBs, Root and Caffe to using HDF5, Pandas and Tensorflow

❑ Found that all aspects of training were significantly faster, and that network inference was roughly 7 times faster

A process which previously took months, can now be done on the order of a week.

❑ This also allows us to train many network variants to explore how systematic uncertainties (eg detector calibration) affect network performance.

A Convolutional Neural Network Neutrino Event Classifier: JINST **11** P09001

Context-Enriched Identification of Particles with a Convolutional Network for Neutrino Events Phys. Rev. D **100**, 073005

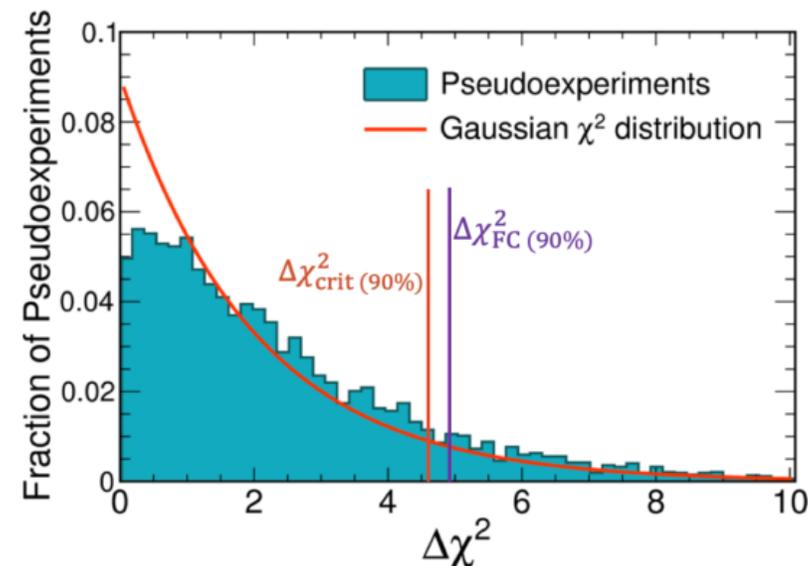
# NOvA HPC

The frequentist Unified Approach [4] or Feldman-Cousins approach is commonly used in particle physics (and by NOvA) to construct confidence intervals when Wilks' theorem [2] conditions are not satisfied

- ❑ Low statistics
- ❑ Physical boundaries
- ❑ Binary choices in hypothesis tests

Principle: Pseudoexperiments are generated and fitted to build empirical test statistic distributions which may deviate from standard  $\chi^2$  distributions.

Ultimately, requires the generation and fitting of **millions of pseudo-experiments**: extremely **computationally intensive** approach.



Constructing confidence intervals with traditional computing resources like the Fermigrid/OSG would take >6 months

Generation and fitting of millions of independent pseudo-experiments is a **highly parallelizable** problem

NOvA, in collaboration with SciDAC-4, developed a High Performance Computing framework capable of leveraging the power of supercomputers like NERSC

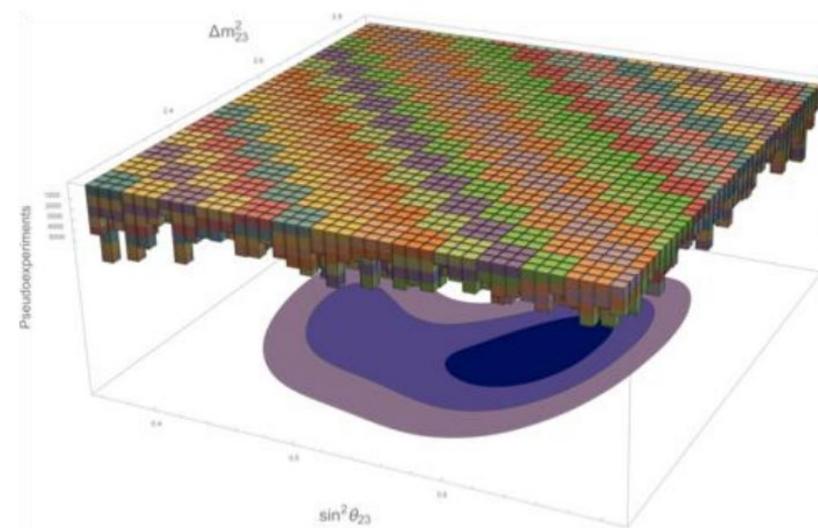
[4] G. Feldman, R. Cousins. doi:10.1103/PhysRevD.57.3873

[5] S. Wilks, doi:10.1214/aoms/1177732360

# NOvA HPC

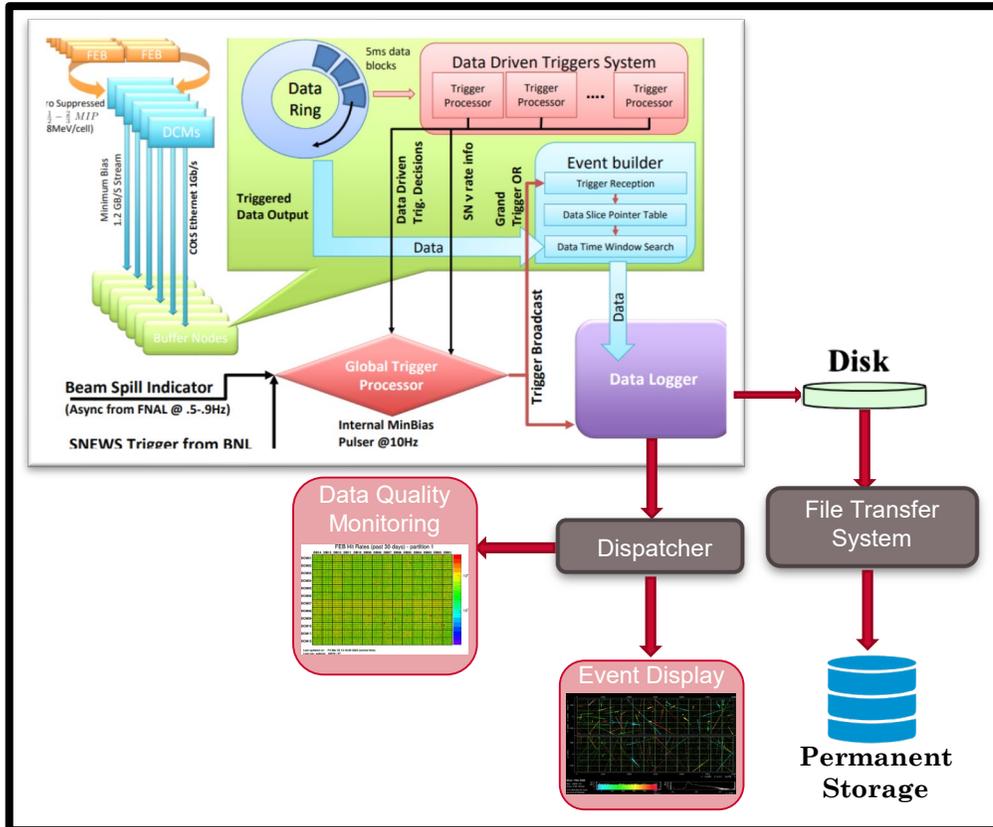
Feldman-Cousins framework with HPC platforms:

- ❑ Docker image with our OS, software stack, scientific libraries, physics inputs etc
  - ❑ easily converted to shifter at NERSC or singularity at ALCF
- ❑ DIY [6] implementation
  - ❑ Large Message Passing Interface (MPI) program ( $10^5$  s of MPI ranks)
  - ❑ easy decomposition of the parameter space across the computing nodes (allocate computing resources where necessary)
  - ❑ Enables dynamic rank communication
  - ❑ Convenient job configuration
- ❑ Replaced ROOT with Eigen library for linear algebra: 30% speed-up
- ❑ Successful FC campaigns at NERSC since 2018
  - ❑  $>10^8$  NERSC hours so far, and counting
  - ❑ results obtained in a few days, compared to several months prior
- ❑ HPC tools enable previously computationally prohibitive analysis techniques to be employed, which will help NOvA continue to produce exciting physics results



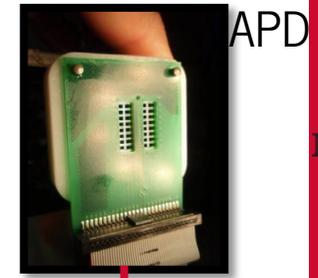
# NOvA DAQ

## NOvA Data Acquisition (novadaq)



Continuous readout rather than “traditional” trigger systems  
 Can then inspect the data after the fact - Ensures we capture everything; then over to DDT or GT

32-pixel Avalanche Photodiodes (APDs) are read out by Front End Boards (FEBs)



APD

Spatial slice of detector data from 64 FEBs collected by PPC+FPGA-based Data Concentrator Modules (DCMs)



FEB

DCMs time order the data in 5μs lists of raw samples for 5ms at a time (stage 1 event building)



DCM

180 different DCMs send their slice of data to a different round-robin “Buffer Node” each 5ms

- ❑ Buffer Node then has 5ms of complete data as an entry in a circular buffer
- ❑ Data rate is ~1.2GB/s

Event Dispatcher allows access to events in shared memory to run online monitoring and event display

See live events for yourself: <https://nusoft.fnal.gov/nova/public/>

# NOvA DDT

## Global Trigger (GT)

Generates a minbias pulser: 10 Hz, 500  $\mu$ s window

Electronics monitoring data requests, 1 Hz

GT manages externally generated triggers:

- Beam: NuMI spill, BnB spill, 1 Hz accelerator heartbeat
- SNEWS (Supernova Early Warning System)
  - heartbeat/min, once-daily full SN readout test
- LIGO
  - heartbeat/min, weekly GW trigger if LIGO is running
- Data Driven Triggers (DDT)
  - artdaq jobs: copy of raw data via shared memory segment
  - ~15 s to make decisions on a 5 ms block

Papers –

- Supernova:
  - Triggering: Phys. Rev. D **104**, 063024
  - Detection: JCAP 10 (2020) 014
- Multi-muons from Cosmic Rays
  - FD: Phys. Rev. D **104**, 012014
  - ND: Phys. Rev. D **99**, 122004
- Magnetic Monopoles: Phys. Rev. D **103**, 012007
- Multimessenger Coincidences w/ GWs from LIGO:
  - Phys. Rev. D **101**, 112006
  - Phys. Rev. D **104**, 063024

## NOvA Data Driven Triggers (novaddt)

