

DUNE OFFLINE COMPUTING



Oregon State

Heidi Schellman, Oregon State University
For the Computing Consortium




U.S. DEPARTMENT OF
ENERGY

Office of Science

Outline

- DUNE detector and physics
- What are our unique challenges
- How are we different and similar from the LHC experiments?

- **More info at: DUNE Offline Computing Conceptual Design Report**
 - <https://arxiv.org/abs/2210.15665>



Superposition


Particles can exist as a superposition of states:

Examples:

Electrons in benzene

Polarized light

Your cat



<http://hyperphysics.phy-astr.gsu.edu/hbase/phyopt/imgpho/polcls.png>

A major goal for DUNE is the study of neutrino interactions



ν_e



ν_μ



ν_τ

Flavor Basis
(Interactions)

$$\begin{bmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{bmatrix} \begin{bmatrix} c_{13} & 0 & s_{13}e^{-i\delta_{CP}} \\ 0 & 1 & 0 \\ -s_{13}e^{i\delta_{CP}} & 0 & c_{13} \end{bmatrix} \begin{bmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{bmatrix}$$

$$= \begin{bmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{CP}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{CP}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{CP}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{CP}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{CP}} & c_{23}c_{13} \end{bmatrix}$$



ν_1



ν_2



ν_3


Mass Basis
(Motion)

2 different views of the same neutrinos

Simple example: two flavor mixing

- Assume that the weak interaction states ν_e and ν_μ are mixtures of the mass states ν_1 and ν_2

$$\begin{bmatrix} \nu_e \\ \nu_\mu \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \nu_1 \\ \nu_2 \end{bmatrix}$$



$$\begin{bmatrix} \nu'_e \\ \nu'_\mu \end{bmatrix} = \begin{bmatrix} \cos \theta & \sin \theta \\ -\sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} e^{i(E_1 t - p_1 L)} & 0 \\ 0 & e^{i(E_2 t - p_2 L)} \end{bmatrix} \begin{bmatrix} \cos \theta & -\sin \theta \\ \sin \theta & \cos \theta \end{bmatrix} \begin{bmatrix} \nu_e \\ \nu_\mu \end{bmatrix}$$

Get a beat frequency that depends on $(m_1^2 - m_2^2) L/2E$


The quantum wavelength of a 2 GeV **muon neutrino** is $\sim 10^{-16}$ m

But it is actually a superposition of the 3 mass types of neutrinos which have slightly different wavelengths – the beat wavelength between the types is about 2000 km.

Bottom line – propagation can change a **muon type neutrino**




into an **electron type neutrino**



We know neutrinos mix, are anti-neutrinos exactly the same?


- If the mixing matrix is complex, NO!
- An experiment to see if matter and anti-matter are different would
 - Make lots of muon neutrinos at 2-5 GeV energy
 - Build a detector ~1000 km away
 - Measure the neutrino conversion rate
 - Measure the anti-neutrino conversion rate




If neutrinos and anti-neutrinos act differently, we have a clue about the early universe when weak interactions helped form matter and anti-matter.

But there is this pesky earth in between...

CP violation is not the only source of asymmetry



Mass ordering unknown



We know the mass differences well but...

The mass ordering is unknown!!

Electron neutrinos interact differently with matter due to the W diagram, but the other neutrinos only see the Z


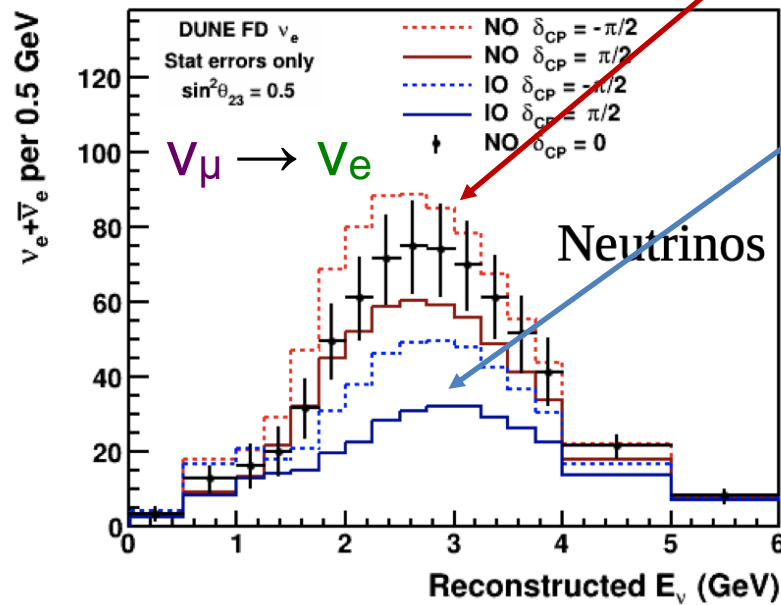
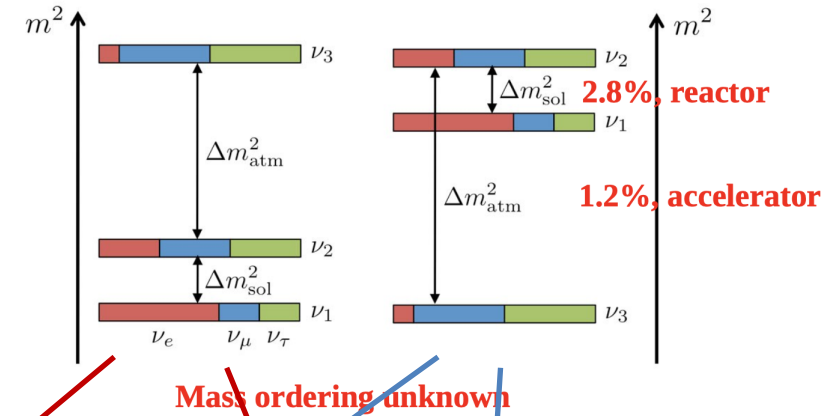
The sign of the matter effect flips if the mass ordering flips

You can measure this if you go through enough matter.

What we expect to see at DUNE

- Here mass ordering shows up as the difference between **red (normal ordering)** and **blue (inverse ordering)** lines
- **CP violation** shows up as the difference between **solid** and **dashed** lines
- The information is in the energy dependence and the ν -anti- ν differences

These are ~10% of events




measuring neutrino oscillations is not easy

- Need a long distance ($L \sim 1200$ km)
 - Need a REALLY strong neutrino beam
 - Need a REALLY big neutrino detector
 - Need to tell electron from muon neutrinos
 - Need a LOT of people (and \$)
- Two current efforts
 - NOvA (Illinois to Minnesota)
 - T2K (across Japan)
 - Several future efforts
 - JUNO (reactor neutrinos)
 - HyperK (across Japan but bigger)
 - **DUNE (Illinois to South Dakota)**
 - ESSvSB (Sweden)

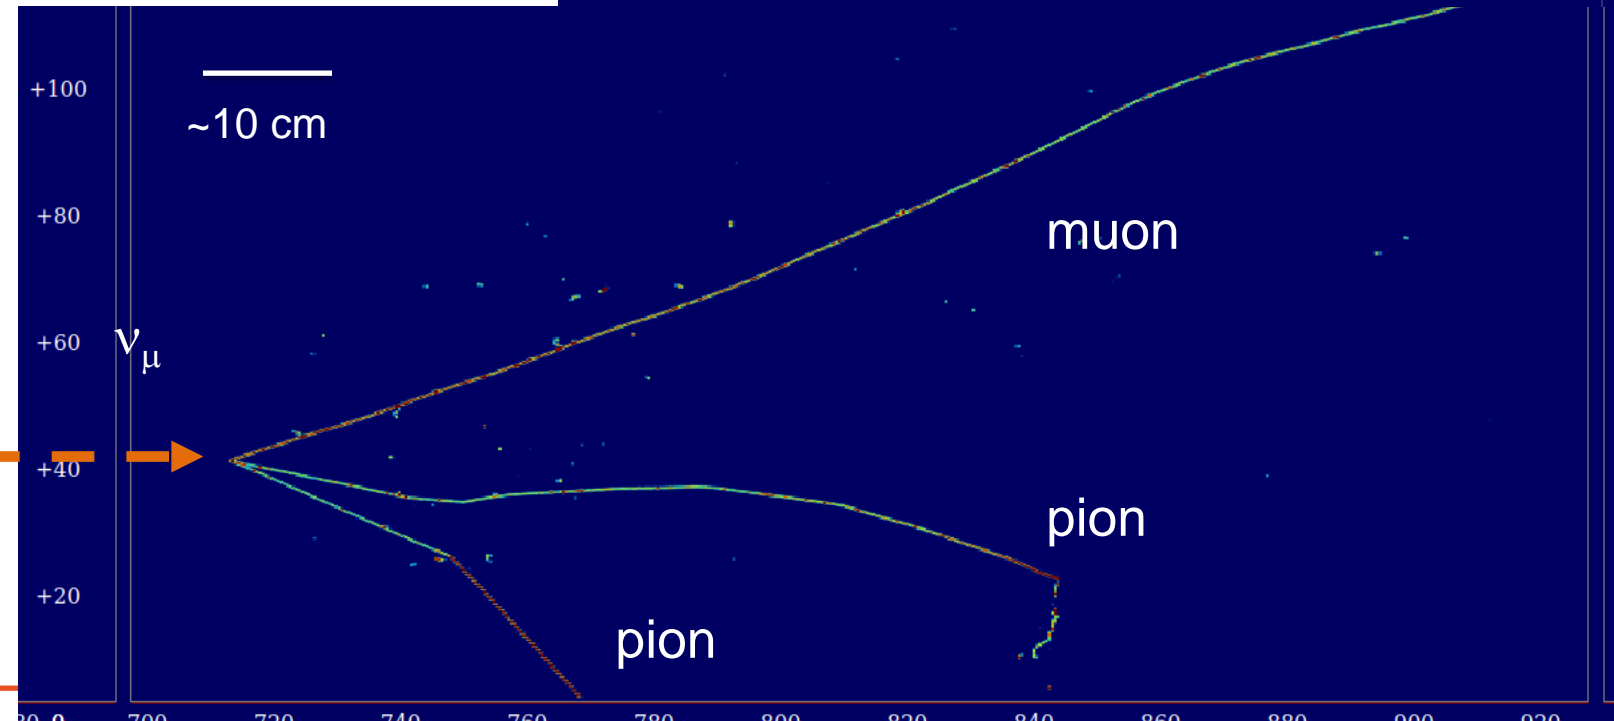
How do you tell a ν_μ from a ν_e



Liquid Argon detectors can distinguish muons from electrons



Problem is you need to instrument $\sim 50,000\text{ m}^3$ with cm granularity and no dead material




Put a huge LAr detector “DUNE” in the Homestake Gold Mine


FERMILAB, IL

Make a very powerful neutrino beam

HOMESTAKE, SD


Run for 10 yrs.






Experimental facilities at 1478 m
(about 1 mile down)
Two vertical access shafts

Long-Baseline Neutrino Facility South Dakota Site



Four cryostats filled with liquid argon

Each of the four cryostats will contain 17,000 tons of liquid argon at 89 K (-184°C or -299°F)





External Dimensions: 19 m x 18 m x 66 m

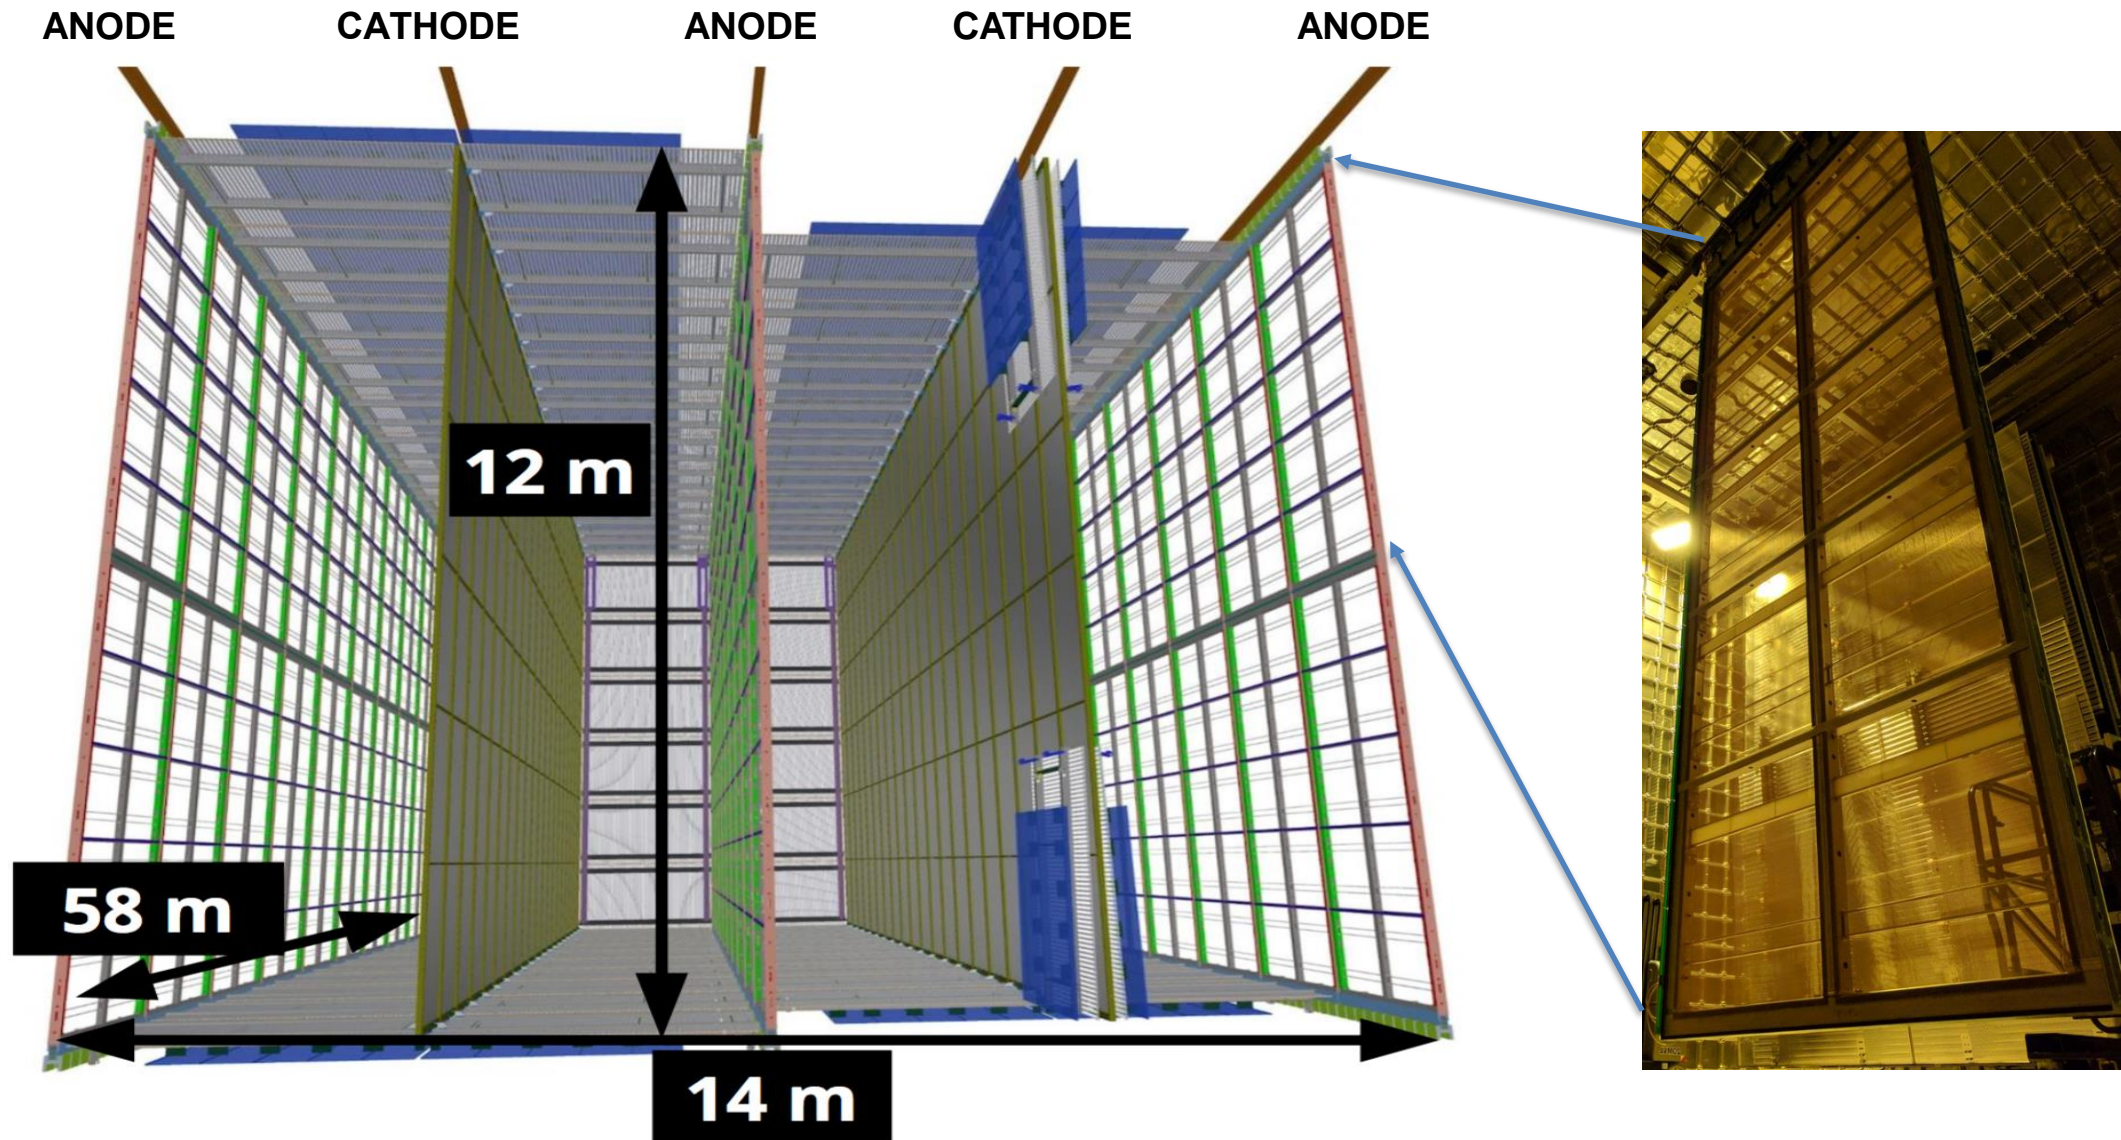
Volume of 4 modules is ~ Weniger Hall



Technical solution: Liquid Argon Time Projection Chamber (LArTPC)



Horizontal Drift Detector (Module 1)



Expected data rates/properties per Far detector module

Process	Rate/module	size/instance	size/module/year
Beam readout	41/day	3.8 GB	30 TB/year
Cosmic rays	4,500/day	3.8 GB	6.2 PB/year
Supernova trigger	1/month	140 TB	1.7 PB/year
Solar neutrinos	10,000/year	≤ 3.8 GB	35 TB/year
Calibrations	2/year	750 TB	1.5 PB/year
Total			9.4 PB/year

- Set a cap at 30 PB/year – allows 2 modules + commissioning/test
- Data taking rate is ~ 10Gbs DC, with 100Gbs bursts for SNB and calibration

Physics → Computing Challenges

Fine segmentation needed for electron-photon discrimination:

Sub-cm-level segmentation over very large volumes drives the number of channels and data volumes.

4 x 400K channels x 6000 14-bit samples for far detectors

12M channels for the near detector

Low-energy (< 10 MeV) thresholds for astrophysical neutrinos:

The need to optimize the low-energy threshold drives our need to carefully record waveforms with minimal processing and thus drastically increases the raw data volume.

Precise energy calibrations:

The large volumes, complex E field configurations, liquid motion, and potential variations in electron lifetime and drift velocity make it necessary to have large calibration data samples that span the full FD detector volumes. Large cosmic ray and artificial calibration samples will dominate the total data volumes from the FD.

Supernovae:

A supernova neutrino burst candidate will generate 320 TB of (uncompressed) data across the first two modules, resulting in thousands of data files produced over a 100s period. Supernova physics drives the need for fast data transmission from the FD to computing facilities. The drastically different time scale of SNB physics also places requirements on the software framework.

Near Detector integration:

Integration of disparate detector technologies into a coherent whole. **Very high occupancy !!!**

Analysis and parameter extraction:

Neutrino interaction samples are generally simpler than event records at colliders. However, final parameter extraction using large numbers of nuisance parameters (and or ML) is still a computationally intense problem and will likely require efficient utilization of HPC resources.


Time scale

- The first far detector module is scheduled for late 2028
 - Starts commissioning the DAQ system several years before
 - Do calibrations and astrophysics
- The near detector and beam arrive in 2031
 - Do oscillation physics
 - Do precision neutrino interaction physics
- Take data for a long time!

- But first, we have prototypes!



- ProtoDUNE first test
- We built 6 Anode Plane Assembly's (APA)
- Put them in the cryostat
- fill with ~ 800 tons of purified Argon



Non-trivial geometry/calibration


Good news:

Volume filled with uniform material
simulation codes really like this


Bad news:

Field non-uniformities
Liquid flow
The liquid charges up
Impurities

1/25 prototype still generates 180 MB of
data/readout



Fields lines



ProtoDUNE-SP Data

took data for 6 weeks in 2018 @ CERN



protoDUNE raw events are each about 75 MB (compressed),
at 10-25Hz – 2 GB/sec peak

We read out 8M interactions

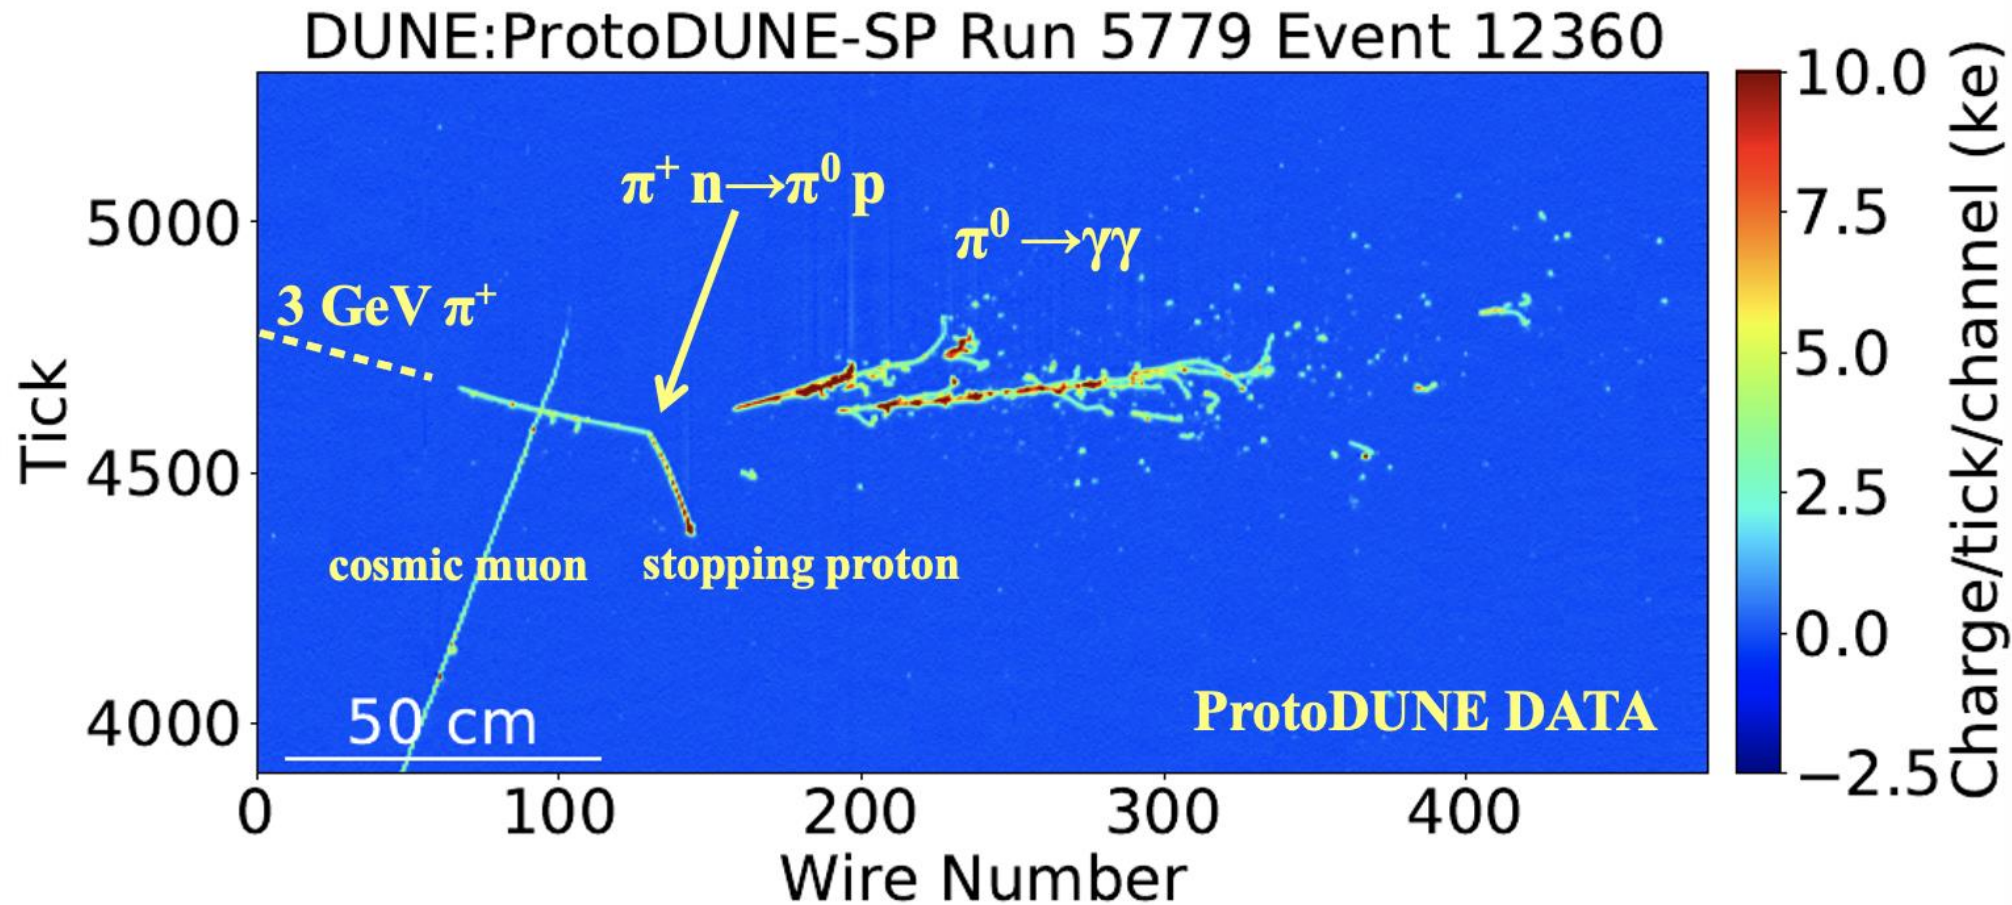
0.4 TB of data,

1.6 PB of simulation

And so far, we've reprocessed 4-5 times







A ProtoDUNE Data Event



Crucial validation of technology

Signal processing for 1 wire plane

- 3 ms of data
 - 6,000 12-bit samples on ~2560 wires
- 10 MB/plane/view
- Remove known bad channels
- Remove coherent noise
- Apply 2-D deconvolution
- Store only regions of interest
 - x 10 data reduction
 - Takes ~30 sec/plane
- ProtoDUNE had 6 wire planes x 3 views
- Full DUNE module has
 - 150 planes → ~2-8 GB/readout



Memory management for Far Detectors



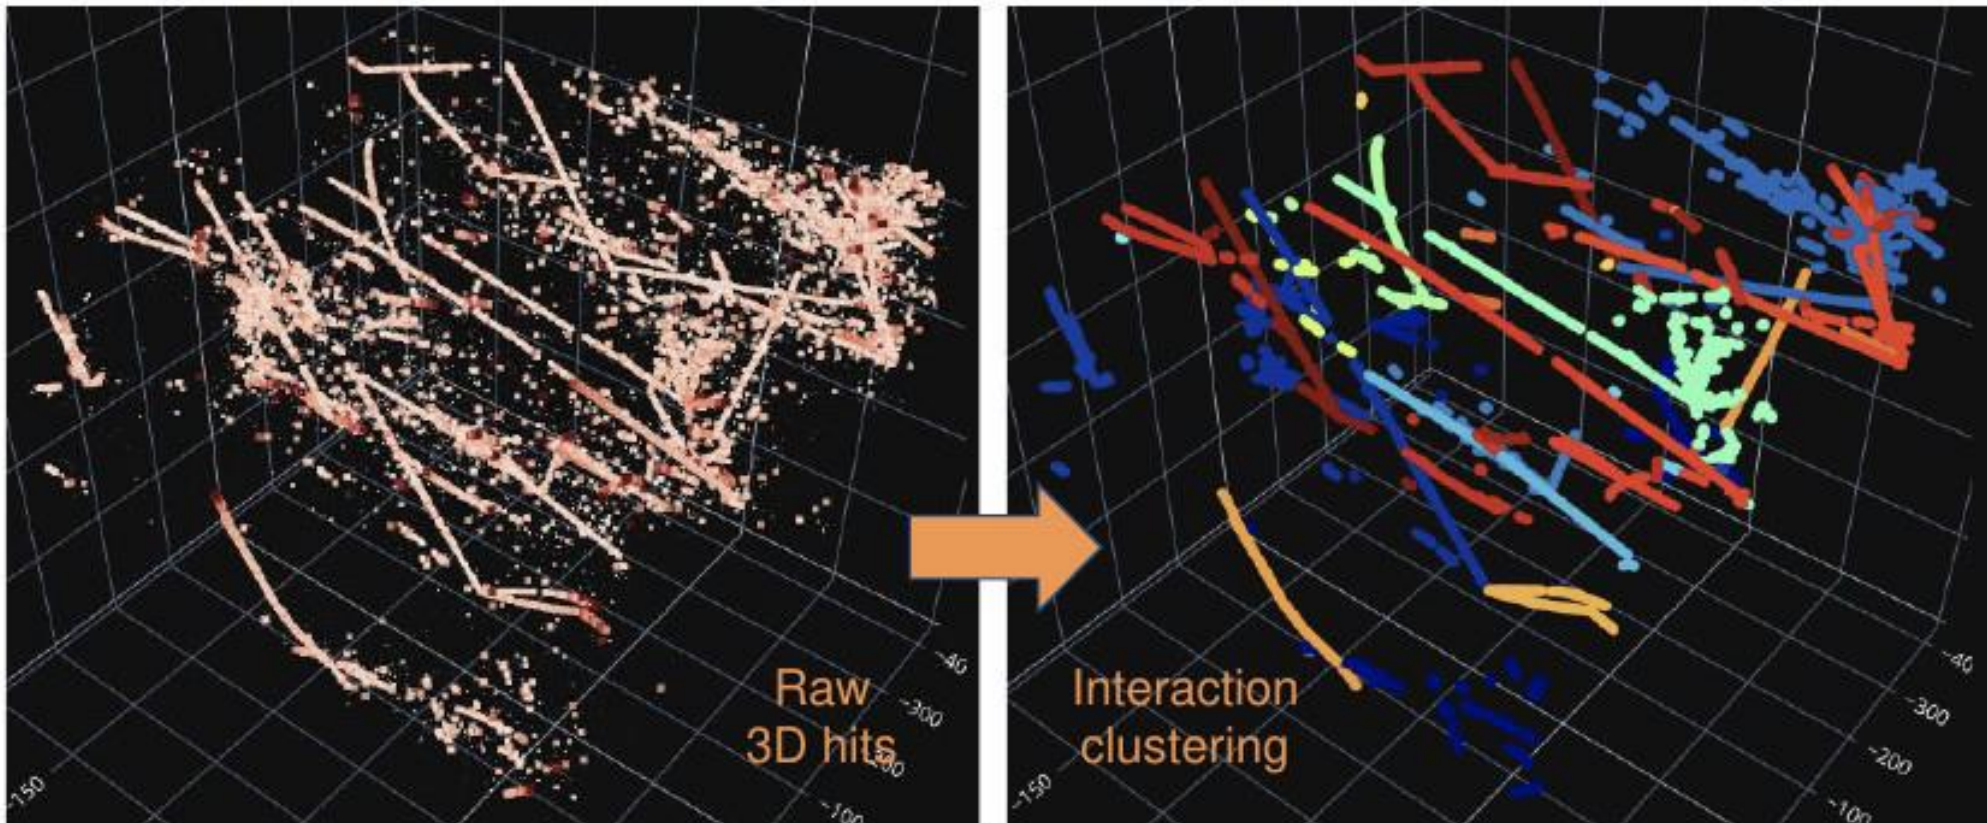
Current framework (Art) optimized for LHC (CMSSW) and small IF experiments.






DUNE data comes in larger chunks
Need to redesign the plane.

Near Detectors

- **10-20 neutrinos** + stuff from interactions in the upstream rock every **10 μs beam spill at 1 Hz**
- Drift time in the argon is several 100 μs , use scintillation light to sort out interactions
- 12M channels with high signal/noise X 15M spills/year



Near detector prototypes




- 2x2 prototype of pixel detector – running at Fermilab in high intensity neutrino beam this summer
- Simulated on GPUs at NERSC/Perlmutter see: *JINST* 18 (2023) 04, P04034

There's more! Supernovae

- SN1987a in the Large Magellanic Cloud. This was 168,000 light years away! But they had just turned on some underground detectors.





The Kamiokande and IMB experiments saw it in neutrinos!





DUNE will be able to see a lot more

- A supernova in our galaxy will deposit 100's to 1000's of neutrinos in our detector over 100 seconds



DUNE should be able to see the next one!

- DUNE should be sensitive to nearby (Milky Way and friends) supernovae. Real ones are expected every 30-200 years
- Supernova readout = 100 sec, one false alarm/month
- 100 sec readout from 4 modules = ~ 600 TB ... takes 10 hrs to read at 100 Gb/s
- Need to process the data quickly!!!



We can use track directions to find the supernova location

DUNE FD-Data for Supernova

If the 180 MB of **protoDUNE** data from one readout was a **few boxes**, a **far detector module for 3 ms** is a **semi-trailer** and a **supernova for 100s** is a **container ship**



Networking for DUNE

- Write 30 PB/year SURF → FNAL
 - Possibly 360 TB as fast as possible for **Supernova!**
 - **Withstand 7-day outage**
 - Implies > 1 PB buffer at SURF and at least one 100 Gb/s link
 - Secondary link for controls/backup
- Host lab must also distribute data for processing to global sites
- ESNET!!!



Compute Challenges Summary

- **Large memory footprints**
Data comes in **30-100 MB size chunks** in aggregations of up to **320 TB** → framework improvements before FD operations
- **High occupancy**
in near detector (20 neutrinos + stuff from upstream interactions)
- **Storing and processing data on heterogeneous international resources**
Current solution is to use mainly single-threaded code and **OSG/WLCG standard configurations** where possible.
Data challenges in progress to test our new computing models.
- **Machine learning**
How do we support **GPU parallelism** and **ML development and deployment** across the wide range of platforms available to us?
There is ongoing work on **HPC** integration.
- **Efficient and sustainable use of resources**
Design for efficient use of global resources in workflow design but also document, train and do continuous integration to make certain that codes are robust and efficient.
- **Keeping it all going**
Experiments run for decades; people and operating systems turn over much more quickly. Considerable effort will always be needed to maintain and deploy
 - Security and authentication
 - DUNE specific databases
 - Code integration and management
 - All the underlying services and code we build on

Bottom line for computation

- DUNE data comes in large but pretty uniform chunks:
 - Each TPC view need to be processed as a whole – 10MB and there are lots of them
 - photon detection wave forms are smaller but non-trivial
 - Near detector has very high occupancy → many pixels
- Simulation is difficult – more like IceCube than ATLAS
 - low energy thresholds and high granularity over large spatial volumes
 - the detector is largely transparent so photon ray tracing over 10s of meters is also a significant challenge.
- Bottom line
 - reconstruction CPU/byte is similar to an LHC experiment at 1-10 sec/MB
 - but the data comes in much large chunks
 - Uniform detector makes very parallel GPU use feasible (and demonstrated)

Backup: DUNE Computing Resource Model

- Less “tiered” than current WLCG model
- Sites provide “services”
 - Large sites provide **disk stores** in ~ 1PB chunks + CPU
 - Smaller sites **contribute CPU**
 - Specialized sites: HPC/tape
- plan to use **common tools** (WLCG/OSG/Rucio...) for most services
- participation in the **HSF** process important to provide and integrate new solutions
- We’re about 10% of an LHC expt...

