Fermilab | U.S. DEPARTMENT OF ENERGY | Office of Science

Computing for the DUNE Long Baseline Neutrino Oscillation Experiment

Michael Kirby
for the DUNE Collaboration

ICHEP 2020

July 29, 2020

DUNE DEEP UNDERGROUND NEUTRINO EXPERIMENT

Outline Computing for DUNE

- Introduction to the DUNE Experiment and physics goals
- DUNE Near & Far Detectors designs
- ProtoDUNE detectors at CERN
- Unique DUNE computing challenges
- DUNE Computing Model
- Performance on the ProtoDUNE data
- Future prospects

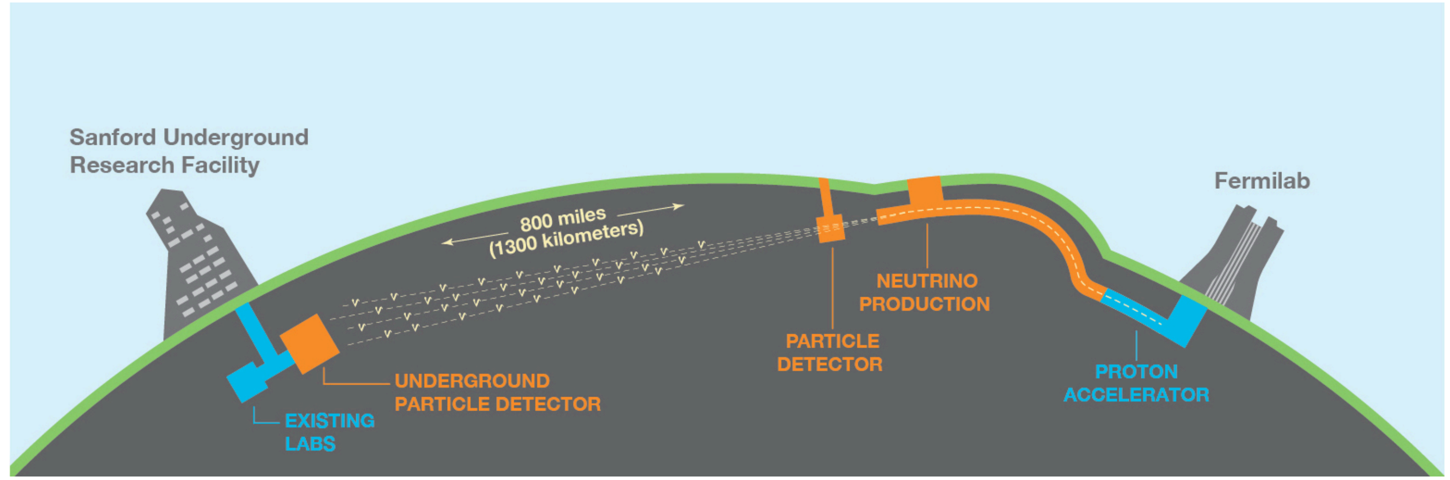

Fermilab

DUNE Experiment Overview

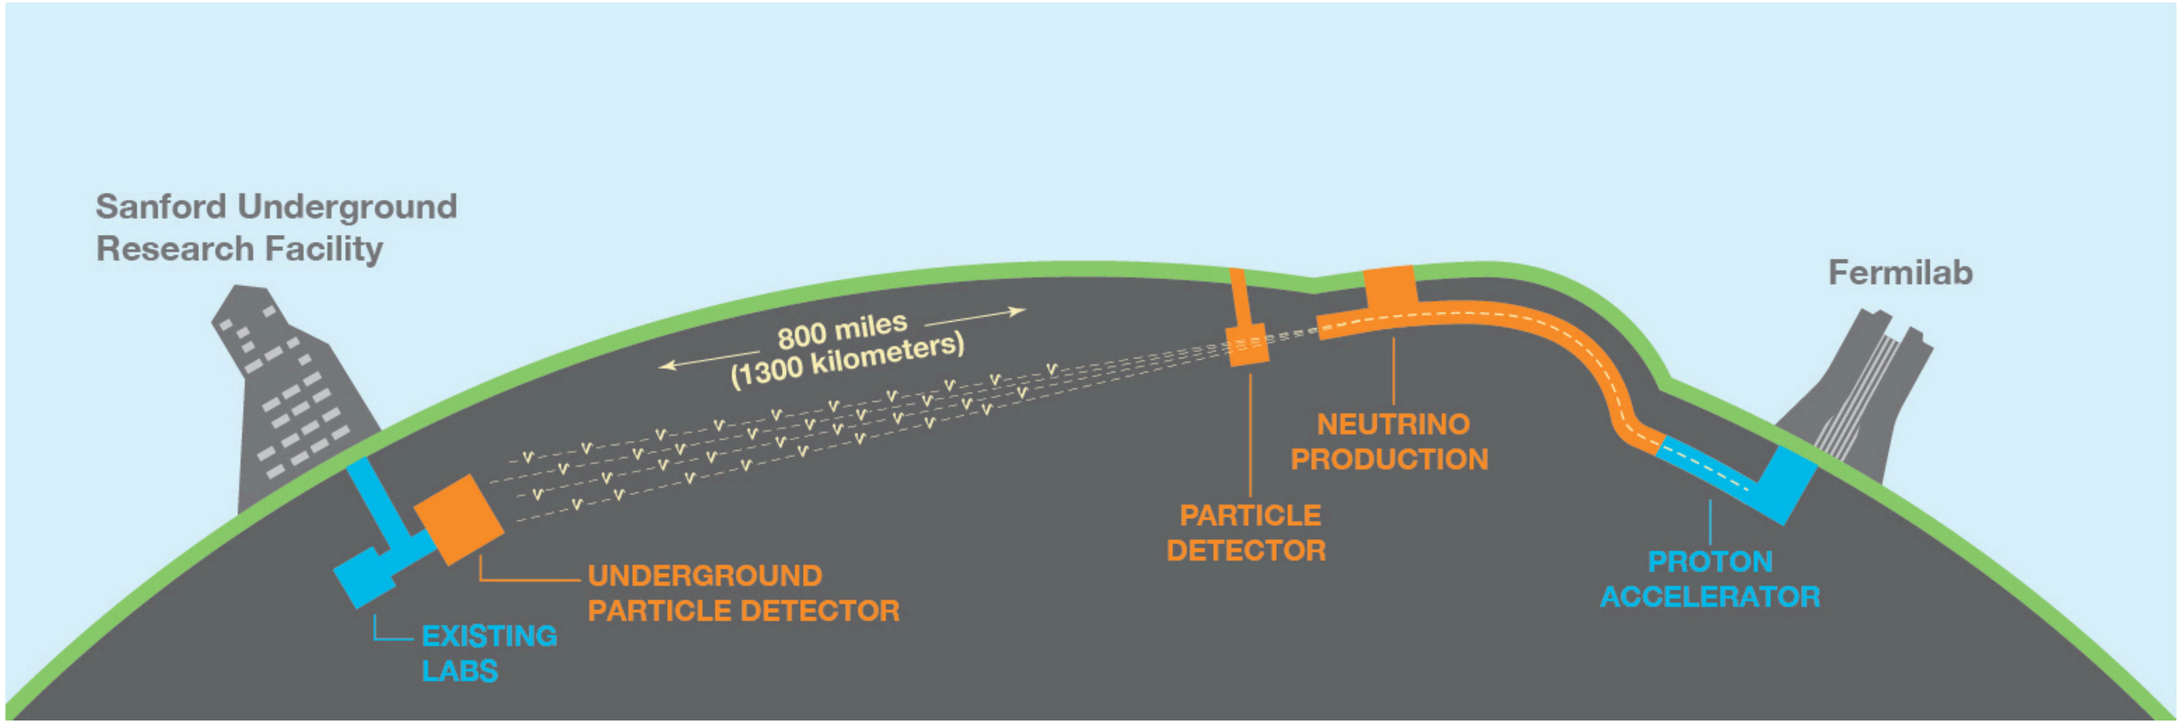

- neutrino experiment measuring neutrino oscillation parameters (mass ordering, matter vs antimatter asymmetry, unitarity), proton decay, supernova neutrinos, and more.
- Far Detector consists of 4 LAr TPC modules at 4850 ft underground in Lead, SD (SURF)
- Near Detector (proposed design) at Fermilab near the neutrino production
- baseline of 1300 km and neutrino beam optimized for oscillation measurement sensitivity
- Two prototypes at CERN - (ProtoDUNE Single Phase - ProtoDUNE Dual Phase)

DUNE Far Detector at Sanford Underground Research Facility

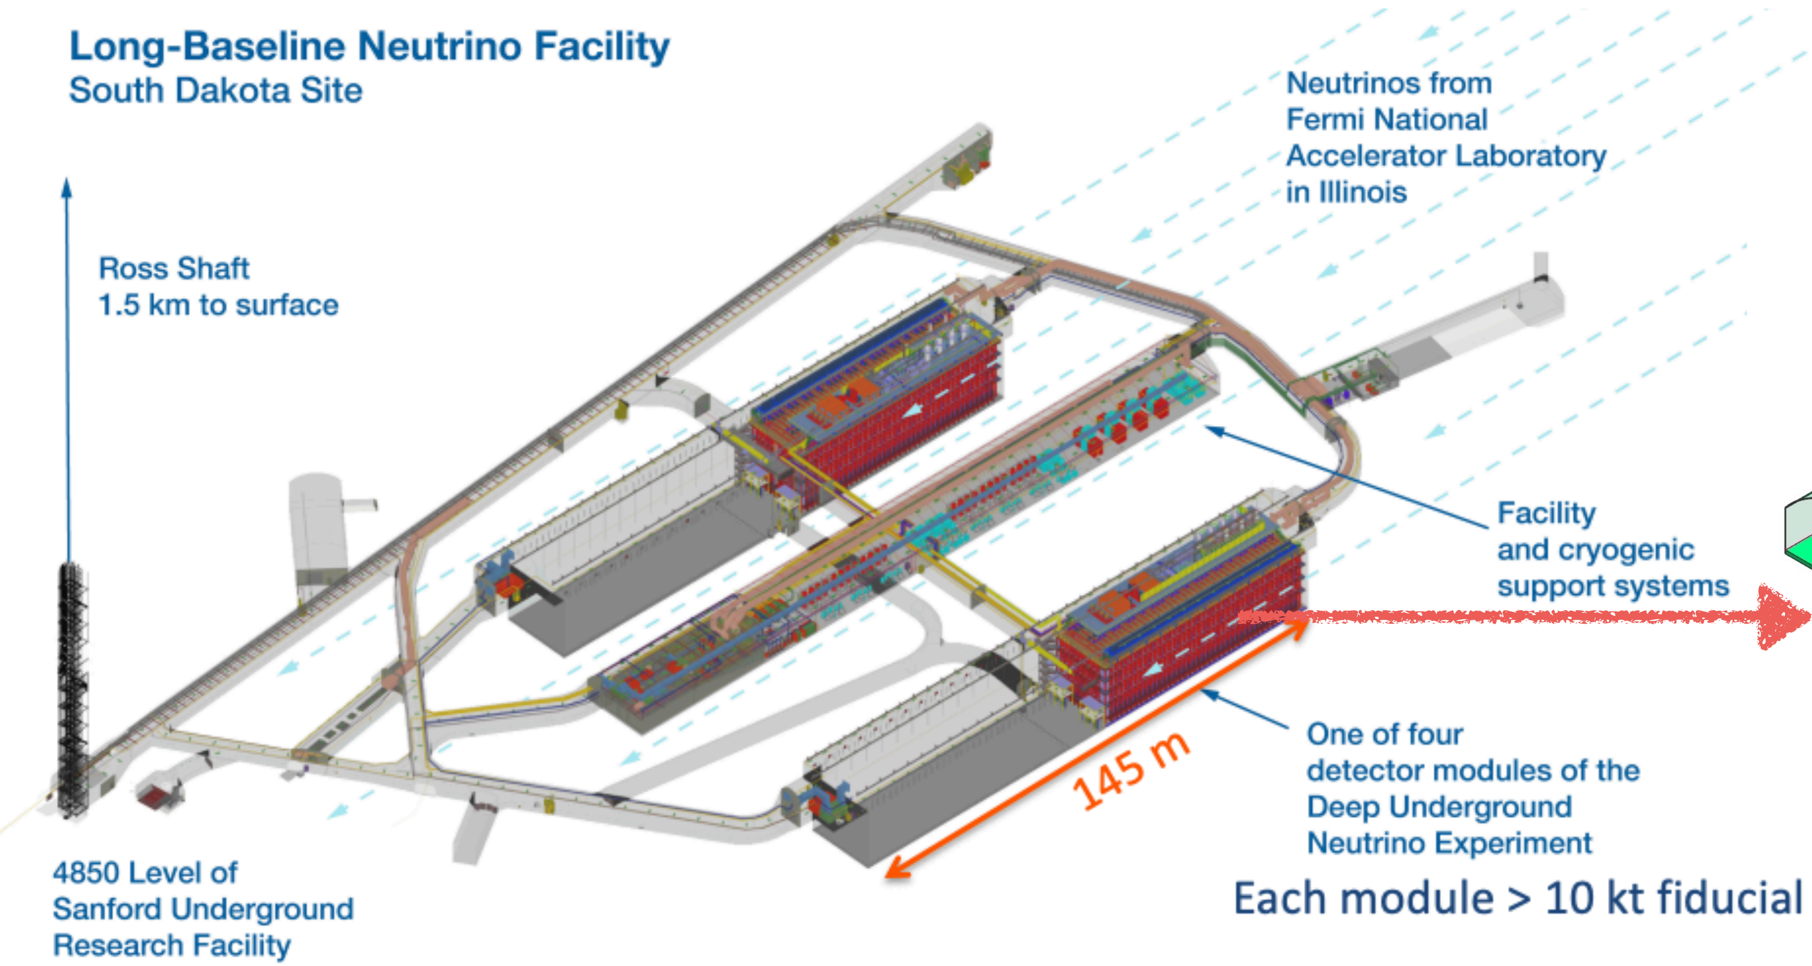

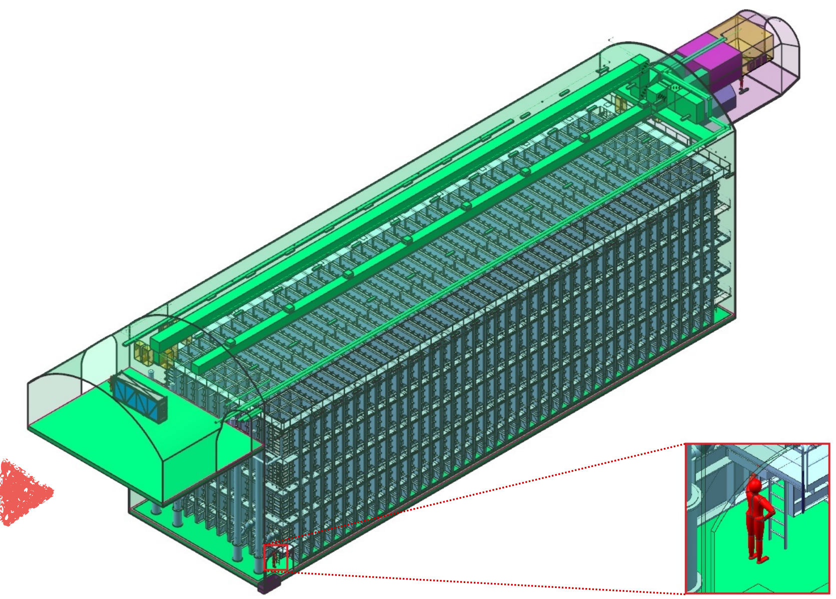

Each module > 10 kt fiducial

Slide: Ed Blucher

Single Phase Liquid Argon Time Projection Chamber

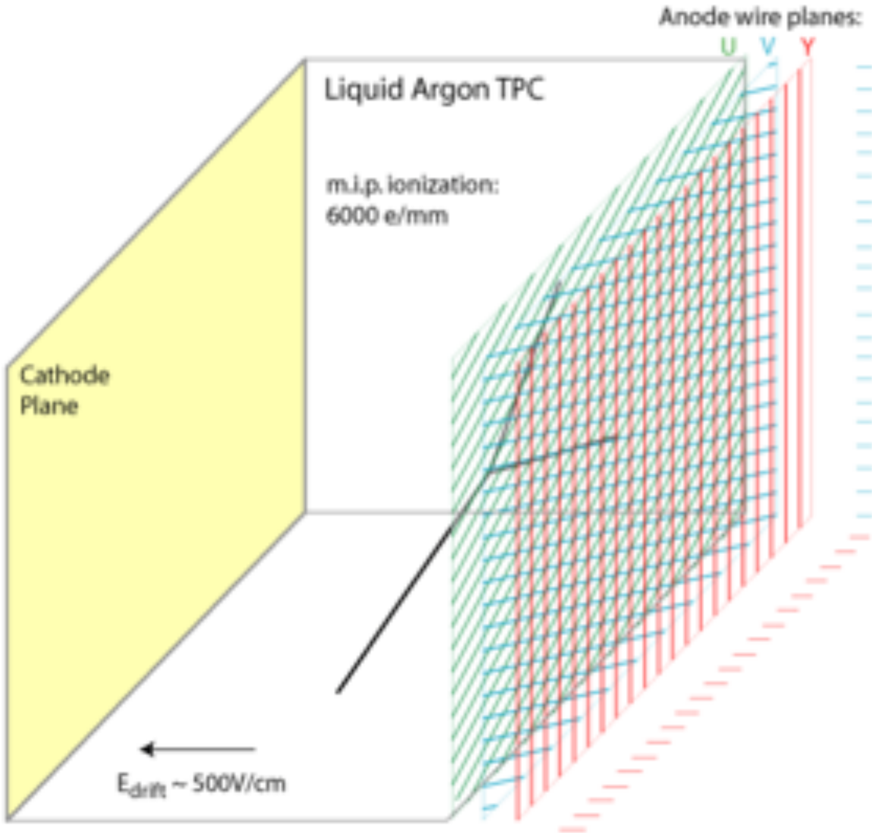

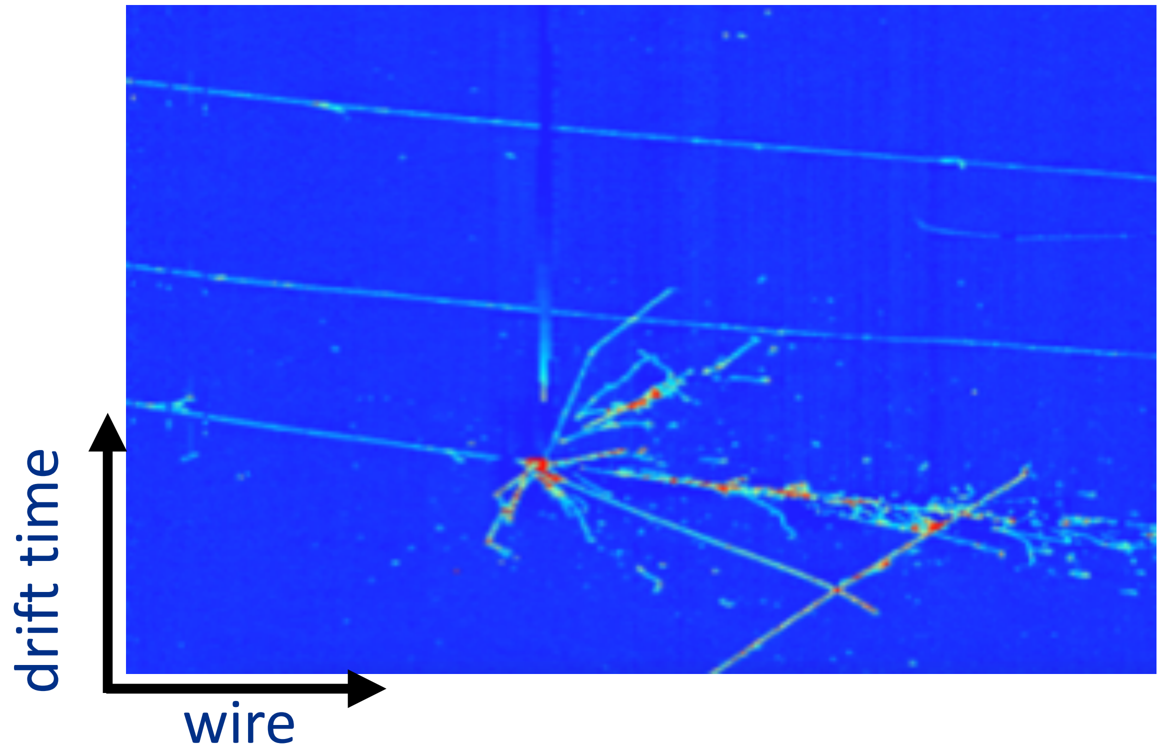

- High spatial and calorimetric resolution - necessary for particle ID, ν-energy measurement
- drift time set by electric field, sampling rate and ADC precision set by physics goals and cost
- two induction planes and one collection plane for the DUNE single-phase Far Detector modules
- one of the future far detector modules will be dual-phase LAr TPC

DUNE Far Detector Design

- First Far Detector build will be a single phase LAr TPC
  - 17 kT of liquid Argon
  - 150 Anode Plane Assemblies tiled on center and walls
  - 180 kV electric field across each drift volume -> 5.4 ms drift time
- Beam timing triggered readout for oscillations physics analysis
  - normal neutrino-beam trigger record is 5.4 ms
  - 12-bit ADC sampled every 0.5 μs
  - 2560 channels per APA
  - 6 GB uncompressed or 2-3 GB compressed
  - 5000 events per day -> 5-10 PB/year/module
- time-extended readout window of far detector module varies greatly
  - continuous readout (SuperNova), calibrations, etc
  - DAQ designed with greater bandwidth, but reduced with trigger, zero suppression, and compressed data format
- goal of 30 PB/year from Far Detector dominated by calibration data
- reconstruction of signals and hits spatially independent within an Anode-Plane Assembly, but 2D deconvolution and FFT require time stitching
- processing of a single trigger record can generate multiple "events" - consider these events to be causally separable regions of interest
- creation of analysis events to minimize data volume and facilitate additional processing

Far Detector SP Module

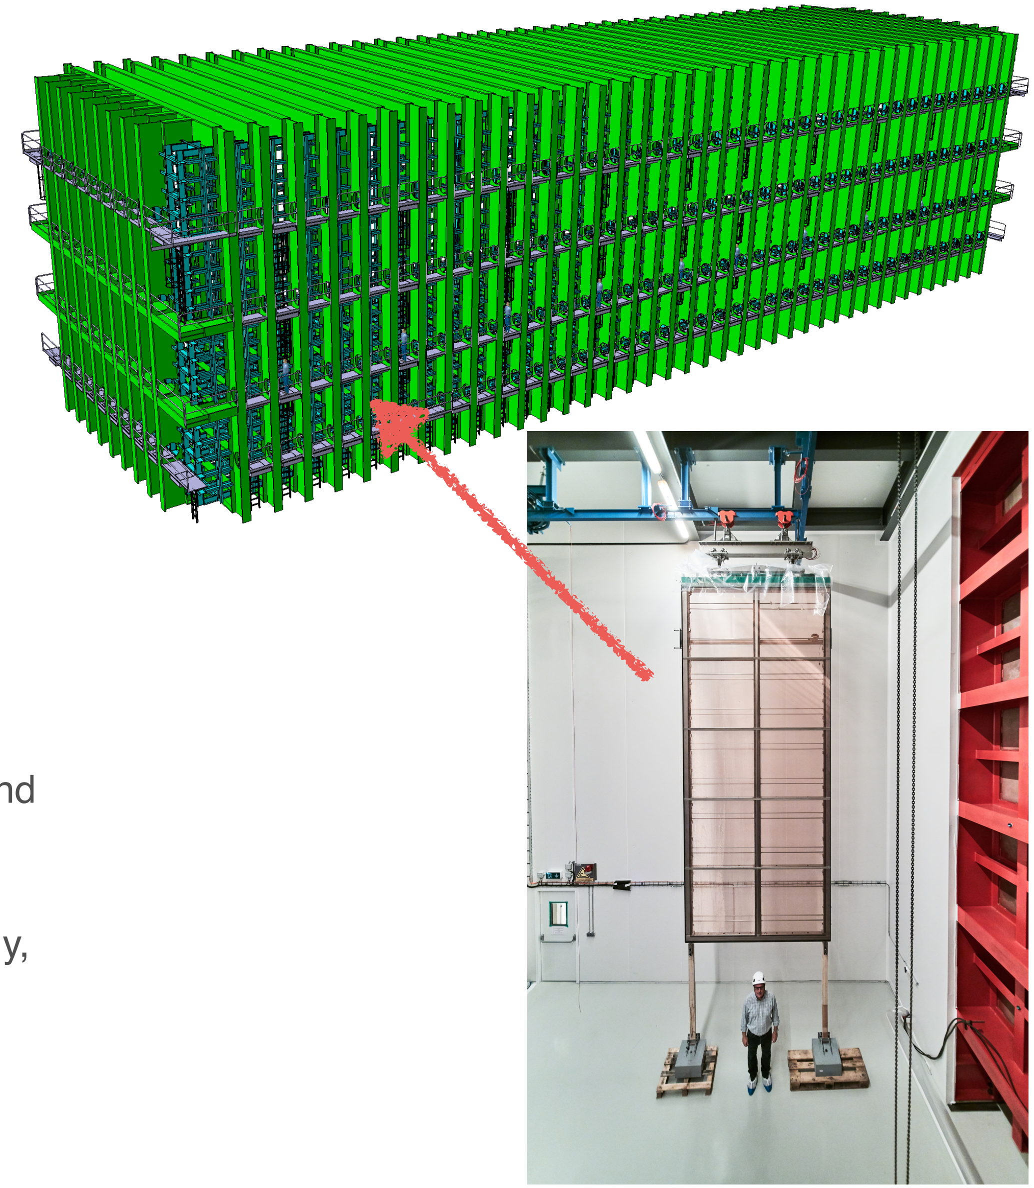

Single APA

Fermilab

DUNE Far Detector Design

- First Far Detector build will be a single phase LAr TPC
  - 17 kT of liquid Argon
  - 150 Anode Plane Assemblies tiled on center and walls
  - 180 kV electric field across each drift volume -> 5.4 ms drift time
- Beam timing triggered readout for oscillations physics analysis
  - normal neutrino-beam trigger record is 5.4 ms
  - 12-bit ADC sampled every 0.5 μs
  - 2560 channels per APA
  - 6 GB uncompressed or 2-3 GB compressed
  - 5000 events per day -> 5-10 PB/year/module
- time-extended readout window of far detector module varies greatly
  - continuous readout (SuperNova), calibrations, etc
  - DAQ designed with greater bandwidth, but reduced with trigger, zero suppression, and compressed data format
- goal of 30 PB/year from Far Detector dominated by calibration data
- reconstruction of signals and hits spatially independent within an Anode-Plane Assembly, but 2D deconvolution and FFT require time stitching
- processing of a single trigger record can generate multiple “events” - consider these events to be causally separable regions of interest
- creation of analysis events to minimize data volume and facilitate additional processing

Far Detector SP Module

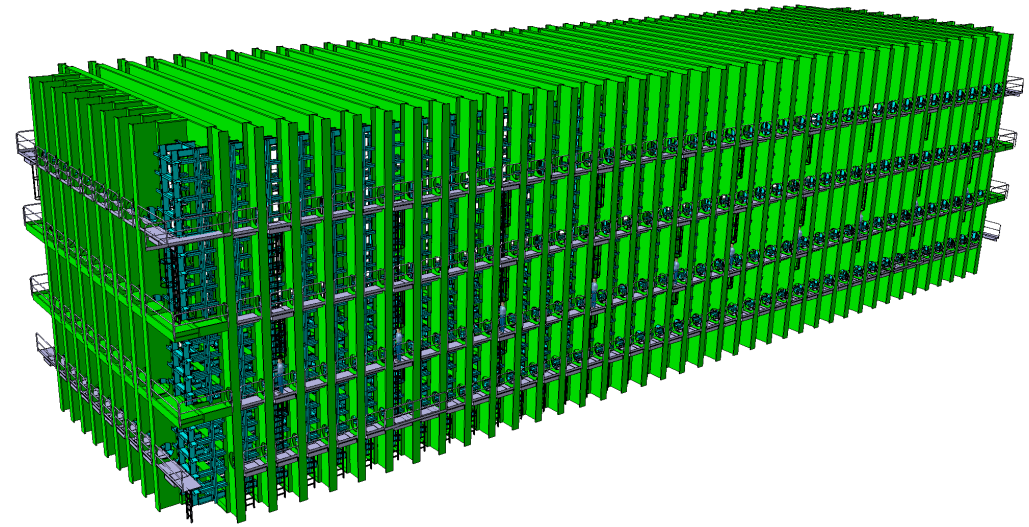

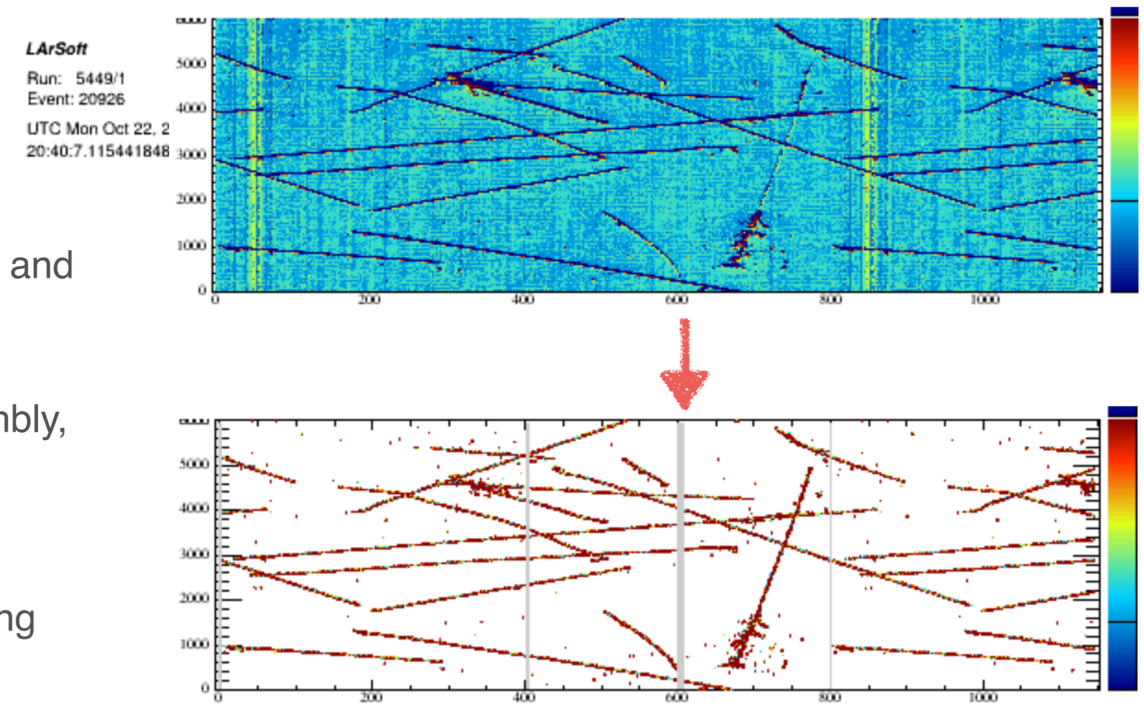

DUNE Far Detector Supernovae Detection

- FD designed to be sensitive to nearby (Milky Way neighborhood) supernovae events
- estimated to occur every 30-200 years, but a false alarm tuned to once a month
- 100 sec readout implies
  - 1 channel = 300 MB uncompressed
  - 1 APA = 768 GB uncompressed
  - 1 module = 115 TB uncompressed
  - 4 SP modules = 180 TB compressed
  - takes 4-5 hrs to transfer at 100 Gb/s
- Need to detect and process the data quickly for multi-messenger observations can follow

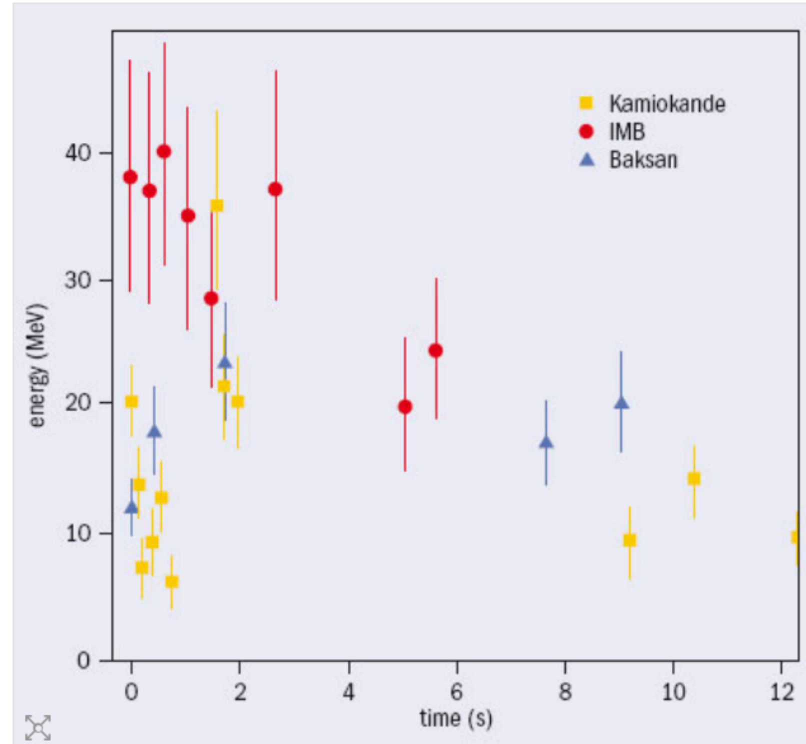

SN1987A neutrino events observed by Kamiokande, IMB and Baksan showed that the neutrino burst lasted about 13 s.

Supernova 1987A

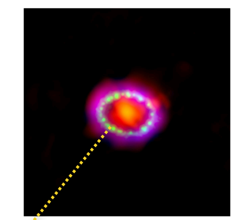

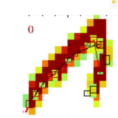

30 MeV

ν_e CC

CERN Courier, Jan 2007

Fermilab

DUNE Far Detector Supernovae Detection

- FD designed to be sensitive to nearby (Milky Way neighborhood) supernovae events
- estimated to occur every 30-200 years, but a false alarm tuned to once a month
- 100 sec readout implies
  - 1 channel = 300 MB uncompressed
  - 1 APA = 768 GB uncompressed
  - 1 module = 115 TB uncompressed
  - 4 SP modules = 180 TB compressed
  - takes 4-5 hrs to transfer at 100 Gb/s
- Need to detect and process the data quickly for multi-messenger observations can follow

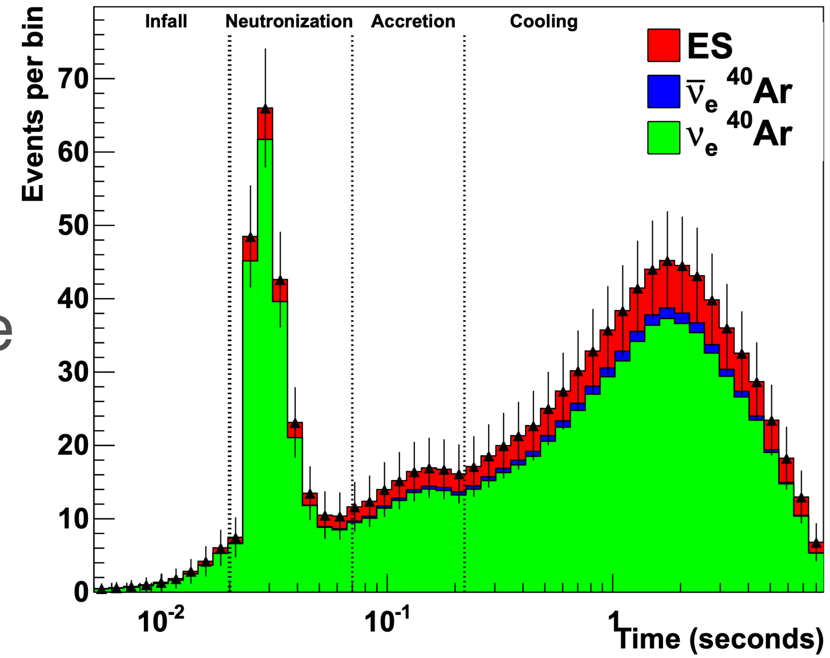

Supernova 1987A

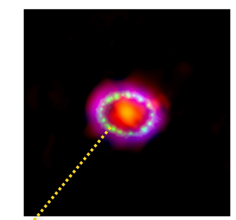

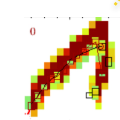

30 MeV

ν_e CC

CERN Courier, Jan 2007

ProtoDUNE Detectors at CERN

- Construct small single-phase and dual-phase LAr TPC (~770 tonnes LAr)
- Single-phase constructed with 6 APAs (instead of 150)
- Dual-phase constructed with 4 Charge Readout Planes

Figure by Thomas Kutter

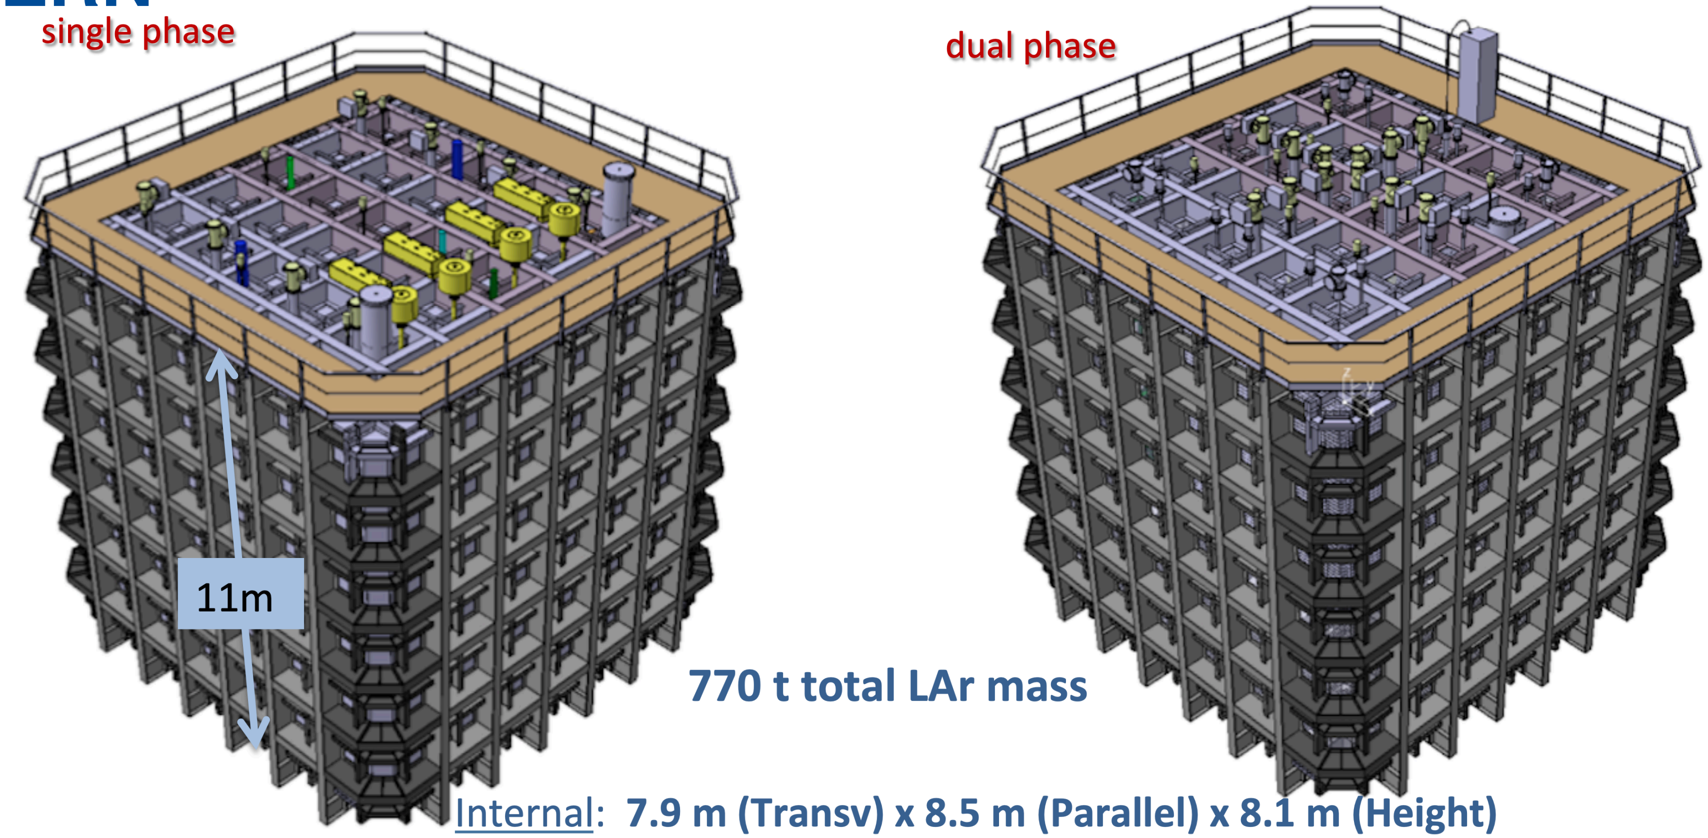

770 t total LAr mass

Internal: 7.9 m (Transv) x 8.5 m (Parallel) x 8.1 m (Height)
External: 10.8m (Transv) x 11.4 m (Parallel) x 11.0 m (Height).

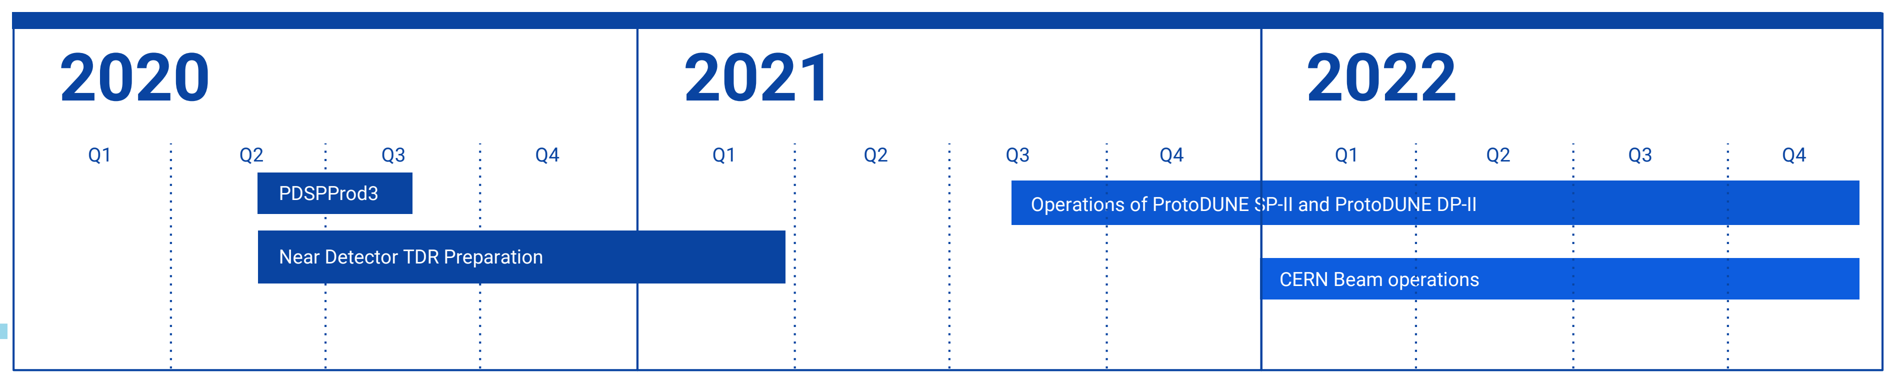

(Proposed) Near Detector Design

Slide: Ed Blucher

- critical part of the oscillation analysis
- precise cross section measurements to improve modeling
- anticipate high occupancy with proximity to beam source and lower overburden
- readout rate ~1 Hz
- at least 1 PB/year
  - but likely much larger
- greater simulation complexity and statistics compared to modeling of Far Detector

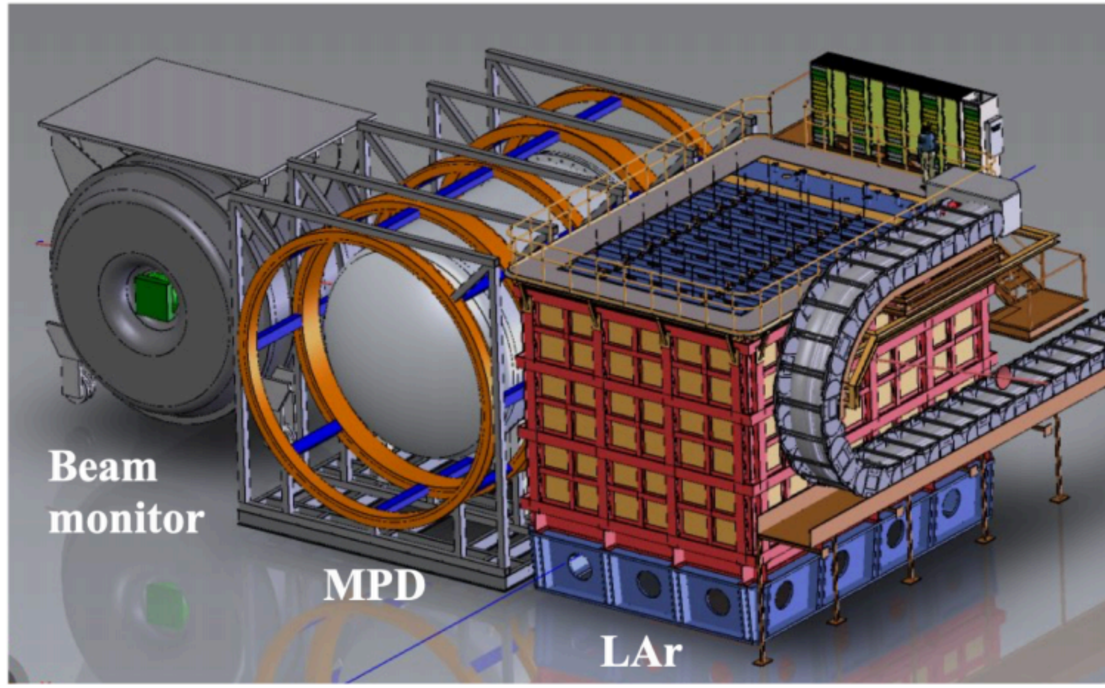

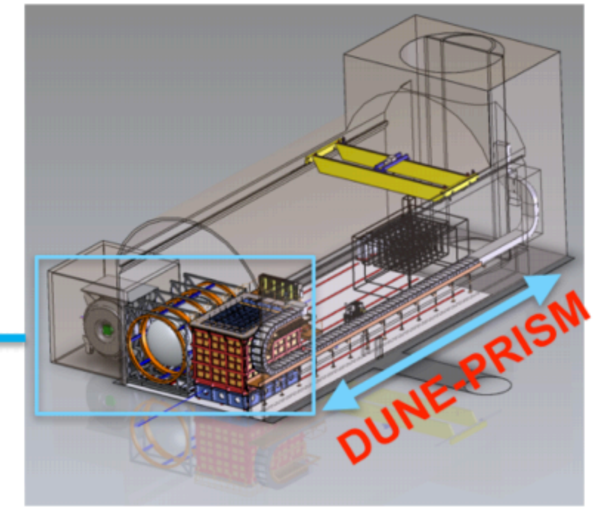

- LAr: Highly segmented LAr TPC (ArgonCube)
- MPD (Multi-purpose detector): High Pressure Gas Argon TPC, Calorimeter, and muon system magnetized by superconducting coils
- Beam monitor: High density plastic scintillator detector with tracking chambers and calorimetry in KLOE magnet
- DUNE-PRISM: Movement of LAr+MPD transverse to the beam, sampling different E_ν

DUNE Computing Challenges Summary

- Need to transfer and store 30 PB/year from Far Detector at SURF
- Handle burst transfer and processing of time-extended trigger records (Supernovae)
  - 180 TB compressed data from 100 s of streaming data - 4 hours to FNAL @ 100 Gb/s
  - 30,000 CPUs to analyze data within 4 hours and provide 5° pointing
- Creation of ROI after signal processing and noise removal
  - reduce data volume from 6 GB/evt
- Calibration data on similar scale as Supernovae (370 TB/year/10kt for laser, 290 TB/year/10kt for PNS)
- Near Detector data volume potentially comparable to Far Detector (trigger rate, occupancy)
- Detector Simulation challenging for 1 MeV scale at FD and occupancy at ND

DUNE Computing Model

DEEP UNDERGROUND NEUTRINO EXPERIMENT
COMPUTING CONSORTIUM

- Working from HSF DOMA model for sites
  - Archive Center - tape/staging
  - Data Center - disk + CPU
  - Compute Center - CPU + cache
  - Analysis Facility - cpu + cache
  - HPC - (HPC, IaaS)
- Goal is to have resource split between FNAL and other institutions - 25%/75%
- FNAL store one copy of all data on tape, replica copy distributed around the world
- additional resources potentially available from commercial cloud, institutional clusters, and HPC

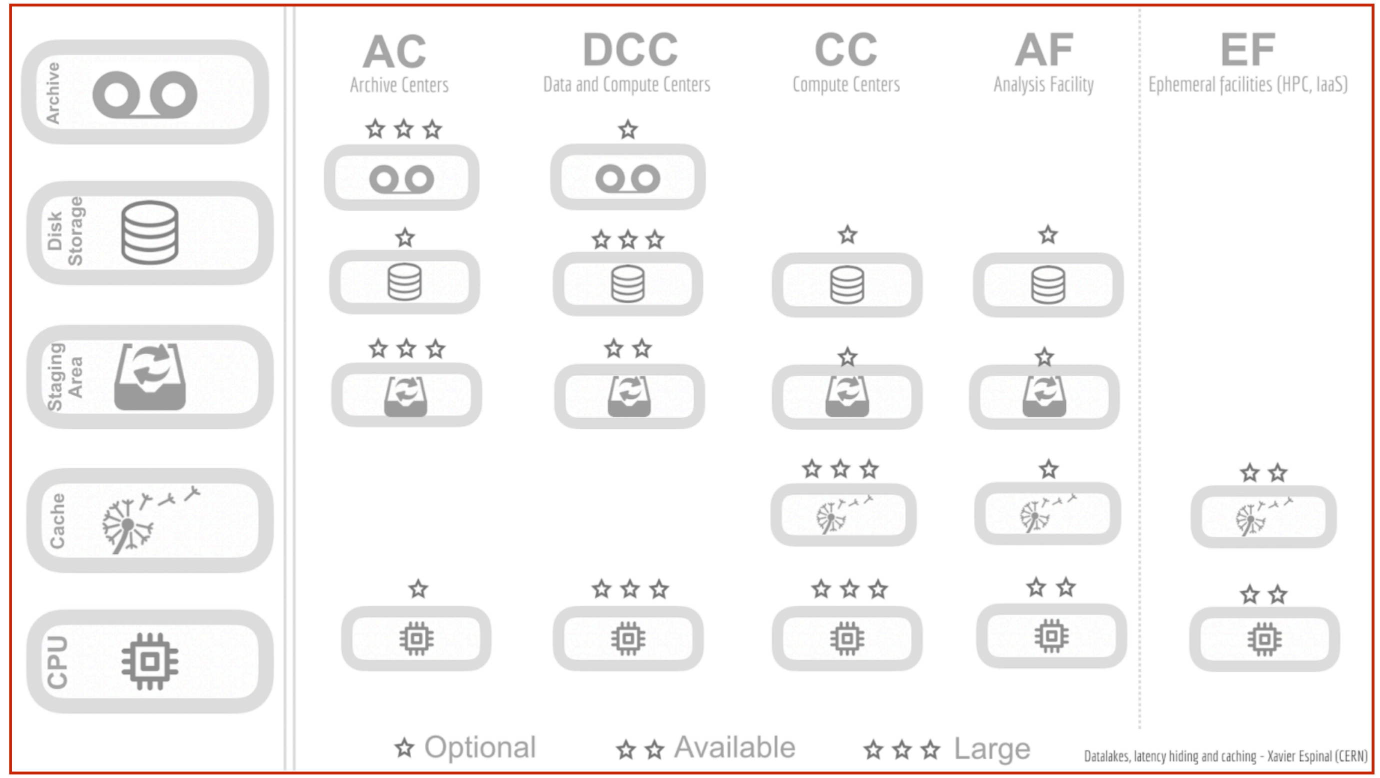

Data Access in DOMA, HSF/OSG/WLCG Joint Workshop J-LAB Newport News, VA 19-23 March 2019

Implementation of Computing Model with Data from ProtoDUNE - SP

- ProtoDUNE Single Phase beam Oct - Nov 2018
- Time-of-Flight and Cherenkov tagged events
  - 300k π - Momentum 1, 2, 3, 6-7 GeV/c
  - additional e, p, K events
  - 8M total beam events - 600 TB raw data
- Additional > 50M cosmic events
  - 2 PB raw data
  - varying the purity, HV, Xenon doping
- utilize this large sample of data to test the Computing Model for DUNE
- See next talk by Steve Timm - “DUNE Data Management Experience with Rucio”

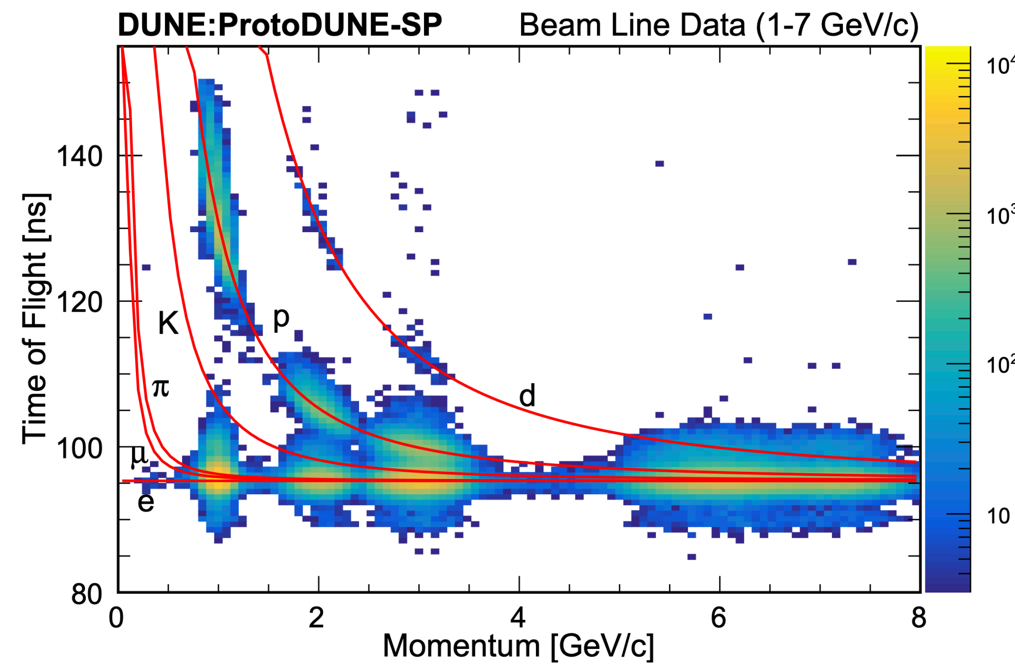

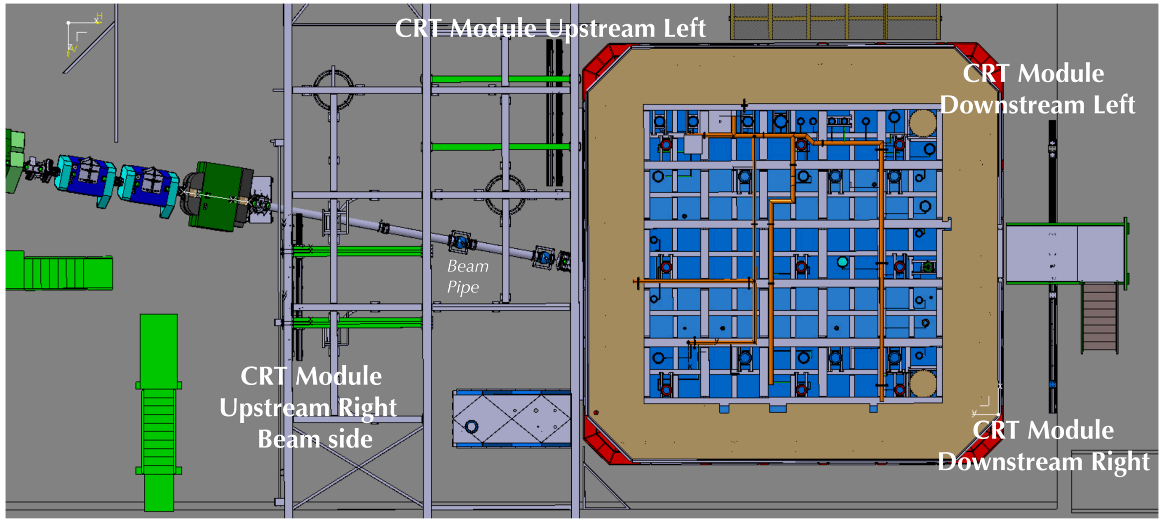

ProtoDUNE Reconstruction Processing

- Based on art framework and LArSoft physics software with ProtoDUNE specific modules
  – WireCell, Pandora libraries utilized extensively
- Processing a 100 event 8 GB file takes ~500 sec/event (80 sec/APA)
  – Signal processing is < 2 GB memory
  – Pattern recognition is < 4 GB memory
  – ProtoDUNE SP event is 75 MB/evt
- reduction to 20 MB/evt - 2 GB files
- with 25 Hz trigger rate is equivalent to DUNE Far Detector beam trigger stream
- “First results on ProtoDUNE-SP liquid argon time projection chamber performance from a beam test at the CERN Neutrino Platform” https://arxiv.org/abs/2007.06722

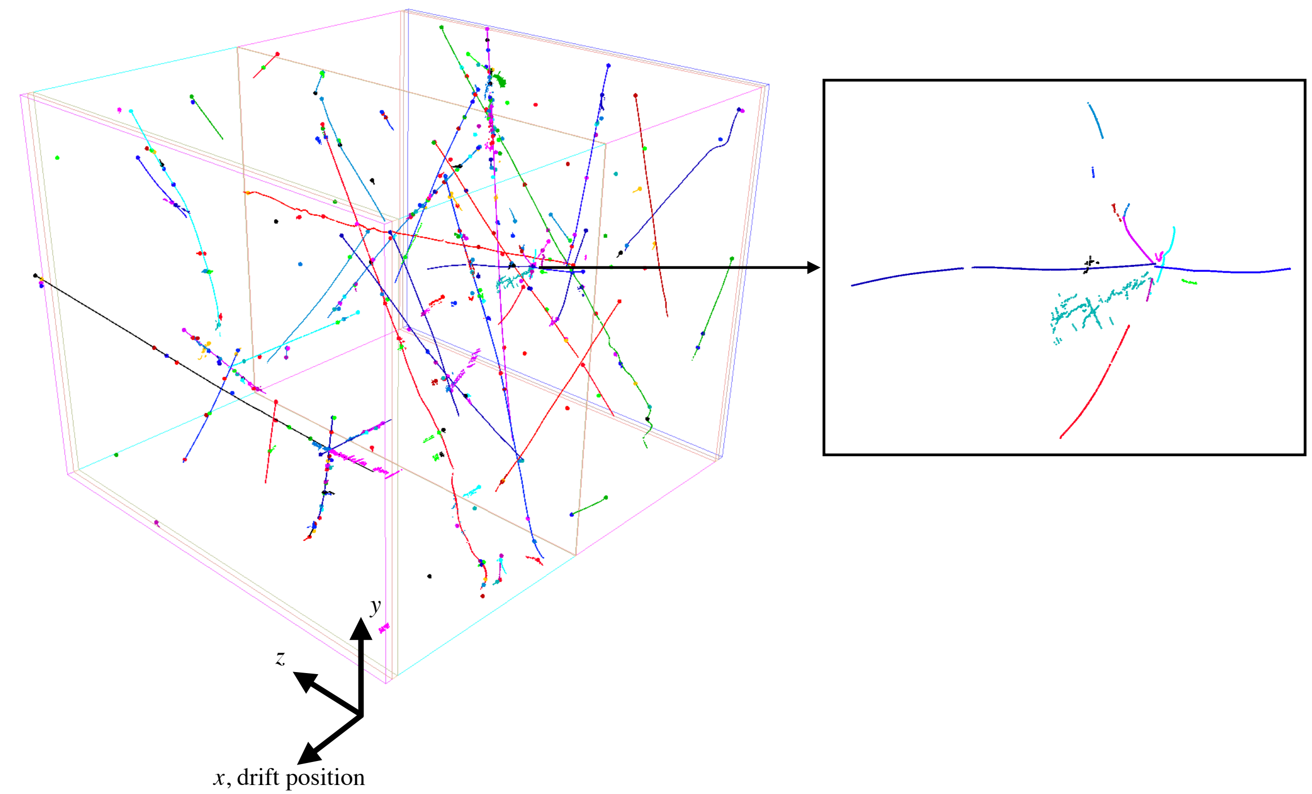

Fermilab

ProtoDUNE Production Processing

- DUNE Site working group doing excellent job incorporating new Compute Elements (CE) and Storage Elements (SE)
  - Continue to add resources from sites around the world - 36 sites
  - Addition of Storage Elements continues - 13
- Soon undertake ProtoDUNE Single Phase Production version 3 (PD-SPProd3)
- Data processing on distributed computing
  - (FNAL ~50% similar to previous usage)
- Utilizing NERSC SuperComputer Cori allocation for simulation generation (10000 simultaneous jobs running ~40% of total DUNE CPU hours)
- Anticipate using 80 - 100 M CPU hours/year in during ProtoDUNE II operations

CPU Hour fraction Feb 1 - May 20, 2020

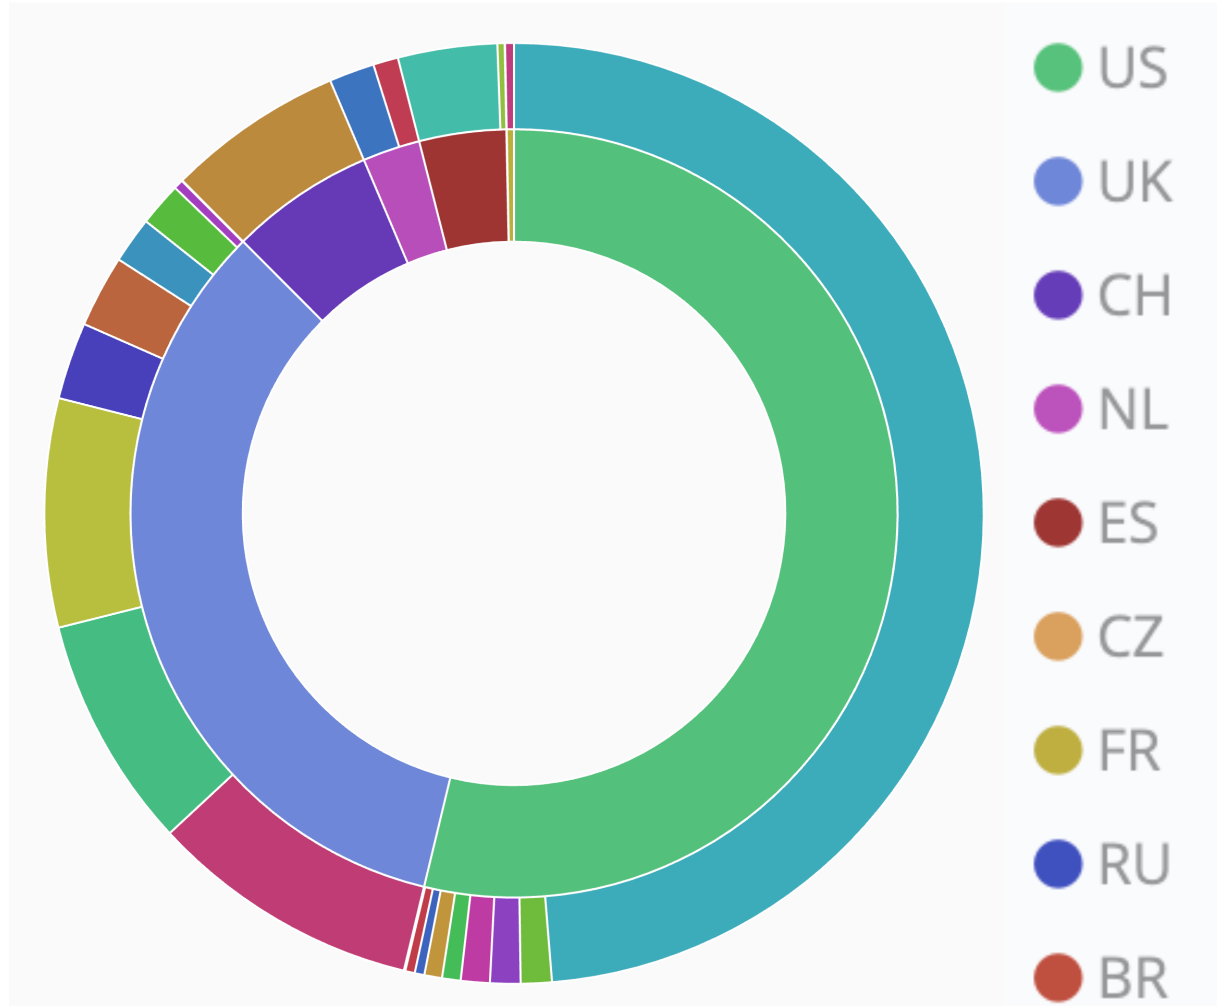

Inner Circle is country

Thanks to Ken Herner for plots

ProtoDUNE Production Processing

- DUNE Site working group doing excellent job incorporating new Compute Elements (CE) and Storage Elements (SE)
  - Continue to add resources from sites around the world
  - Addition of Storage Elements continues
- Soon undertake ProtoDUNE Single Phase Production version 3 (PD-SPProd3)
- Data processing on distributed computing
  - (FNAL ~50% similar to previous usage)
- Utilizing NERSC SuperComputer Cori allocation for simulation generation (10000 simultaneous jobs running ~40% of total DUNE CPU hours)
- Anticipate using 80 - 100 M CPU hours/year in during ProtoDUNE II operations

CPU Hour fraction Feb 1 - May 20, 2020

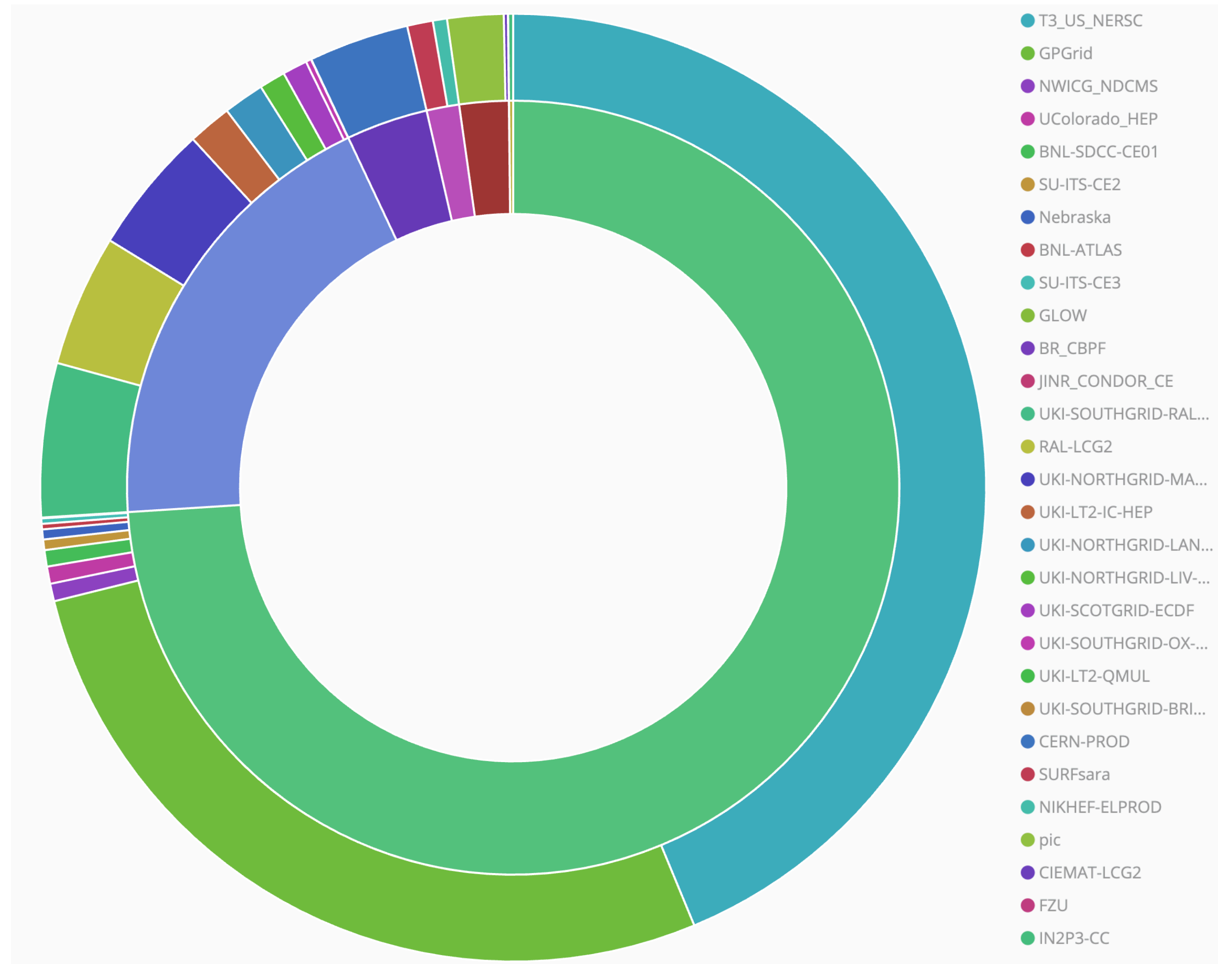

Adding in HPC @ NERSC Cori

Thanks to Ken Herner for plots

Tape and disk storage 2018-2022

Total DUNE Storage

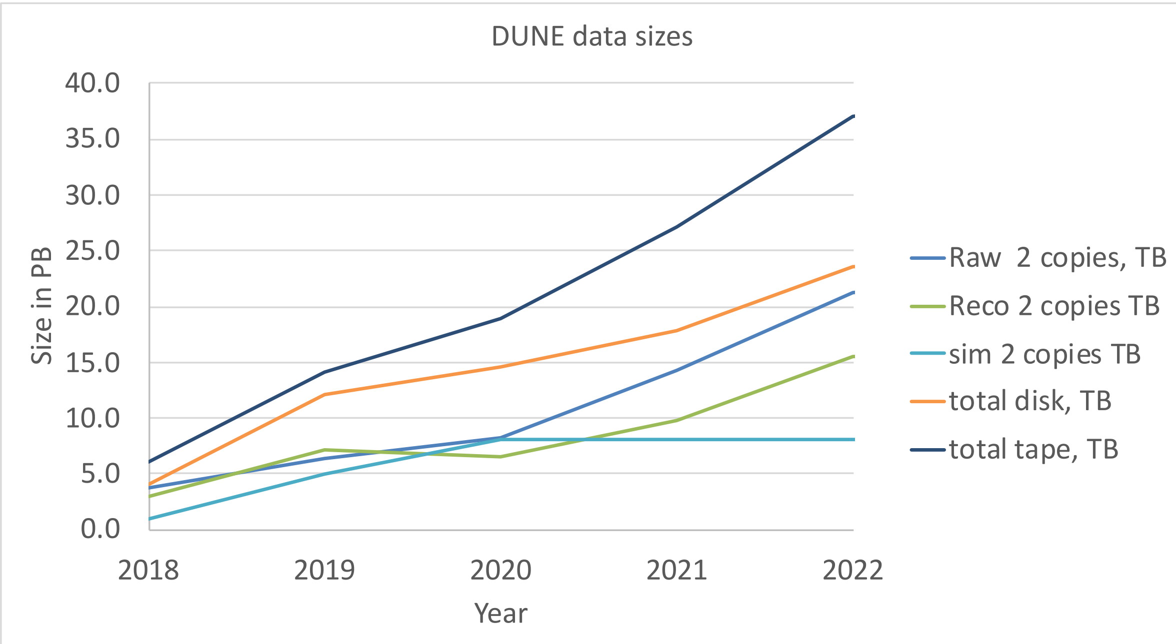

FNAL DUNE Storage

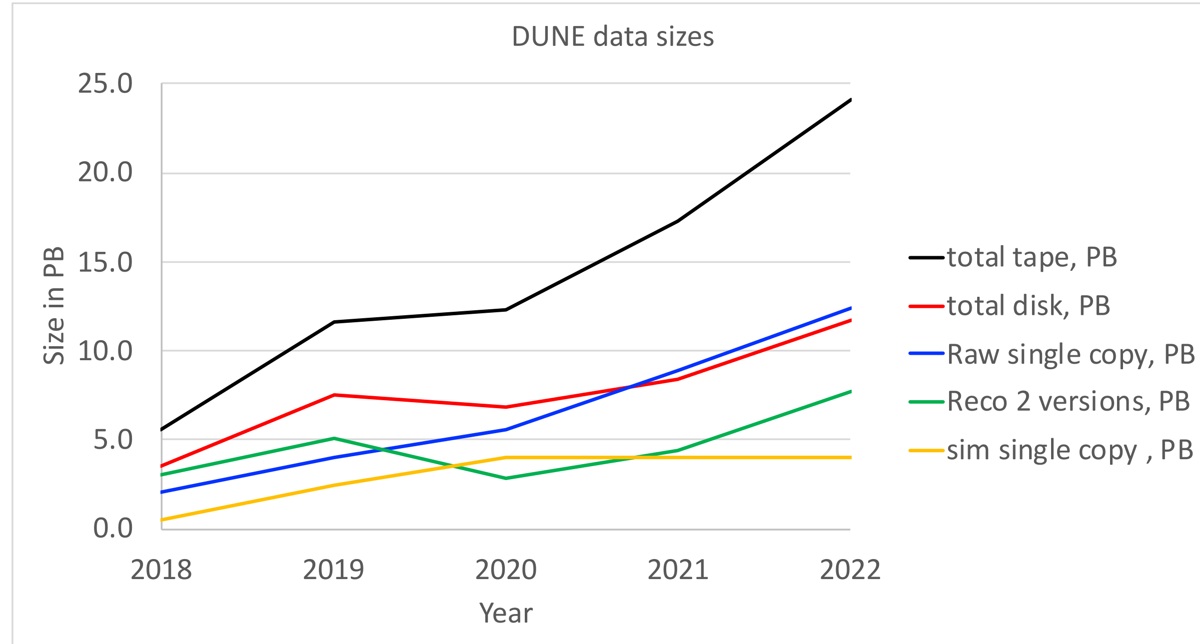

- Computing Model for DUNE Storage
  - 2 archival copies of raw, derived, and simulated data - 1 copy at FNAL, second copy distributed institutions
  - production processing of SP and DP data and matching simulation twice per year
  - 2 or 3 copies of active derived and simulated datasets on disk - dataset stays active for 1 year

Computing Infrastructure Work

- Working to incorporate additional tools for tracking, accounting, and real-time monitoring status of the collective DUNE infrastructure
- Transition to a Global Pool for universal submission from around the world - HTCondor
- working with WLCG and OSG to evaluate, adapt, and incorporate pre-existing tools
- Computing Resource Information Catalogue - CRIC
- WLCG Experiment Test Framework - ETF
- Networking - LHCONE, perfSONAR (**perf**ormance **S**ervice-**O**riented **N**etwork monitoring **AR**chitecture)

Open Science Grid

WLCG
Worldwide LHC Computing Grid

CRIC
Computing Resource Information Catalog

Fermilab

Future Prospecting

Homestake Gold Mine - Lead Historic Preservation Org

- Continue to expand the global nature of DUNE Computing - adopt new site structure, contribution agreements, and support models
- Event Data Model poses potential challenge - 6 GB/evt -> analysis format
  - Machine Learning techniques being developed based on close to raw data
  - capitalize on APA/wire plane/chan parallelism
  - promising new results in Inference as a Service (Cloud GPU) suggest excellent potential
- Continue to collaborate with OSG/HSF/WLCG to test, integrate, and develop features for infrastructure/middleware software
- Data Management and workflow management will be important focus of Computing Model in the near future

Fermilab

backup

Computing Model Policies

- Tape Storage
  - two copies of all raw data for security
  - FNAL provides storage for an archival copy of all raw data for DUNE (ND, FD, protoDUNE)
  - Rucio Storage Elements (RSEs) around world provide storage for 2nd copy
  - FNAL provides storage of derived datasets with lifetime of 2 years
  - FNAL provides storage for single copy of simulated data
  - RSEs around the world provide storage for second copy of simulated data
- Disk Storage
  - two or three copies of every active derived dataset on disk at any time
  - two derived datasets will be active at any one time
  - latest two active derived dataset staged to disk at FNAL
  - two or three copies of every active simulated dataset on disk at any time
  - two simulated datasets will be active at any one time (matching active derived dataset)
- From these policies can development estimates for resource needs

DUNE Computing Outline

DEEP UNDERGROUND NEUTRINO EXPERIMENT
COMPUTING CONSORTIUM

- Computing Consortium Activity
- Grid Operations and Production Campaigns
  - ProtoDUNE-SP Reprocessing
  - ProtoDUNE-DP Processing
- Data Management
- Fermilab Computing Resource Scrutiny Group
- Computing Infrastructure Work
  - CRIC
  - ETF
  - Networking
- Future plans for Computing CDR and TDR
- Database progress and updates
- Analysis Framework Task Force update

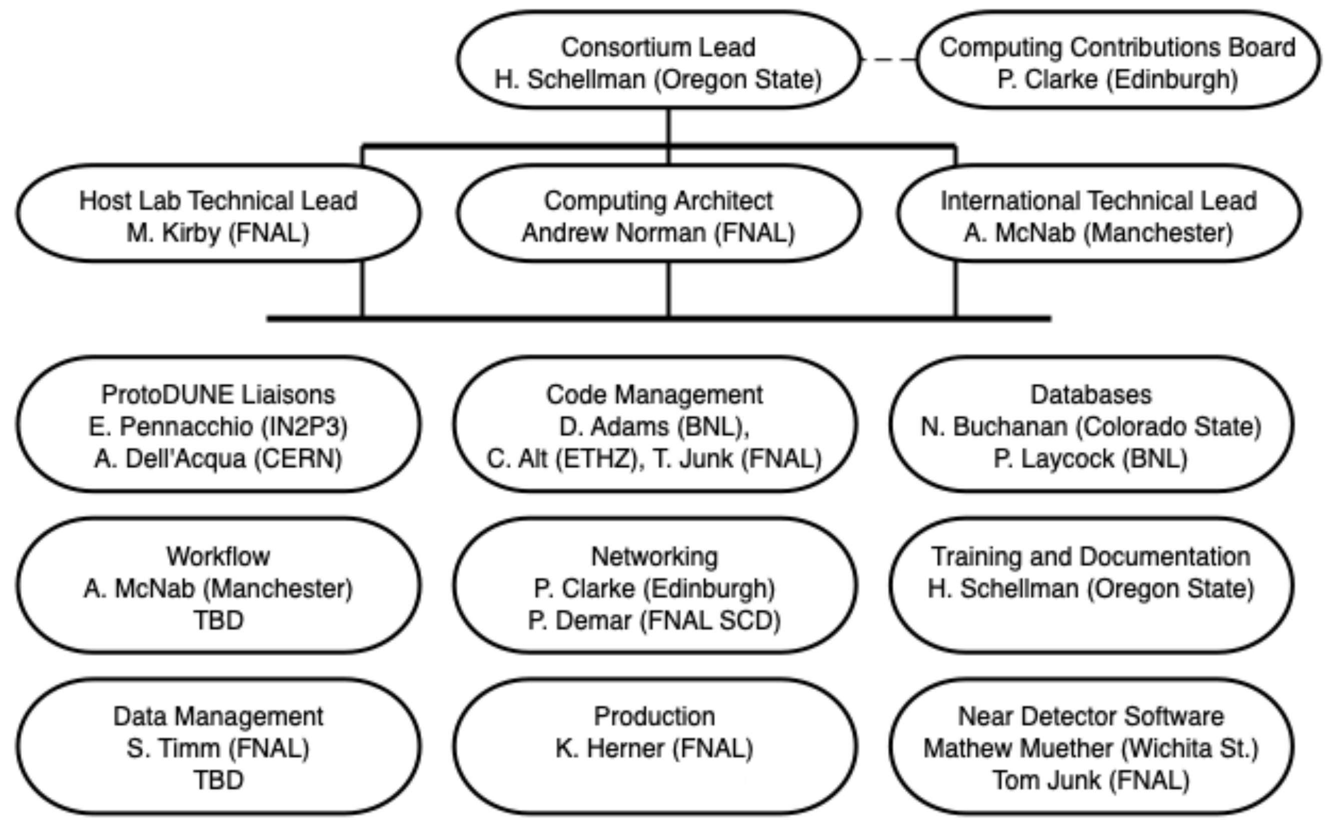

DUNE Experiment Physics Goals

The quantum wavelength of a 2 GeV muon neutrino is ~ 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 

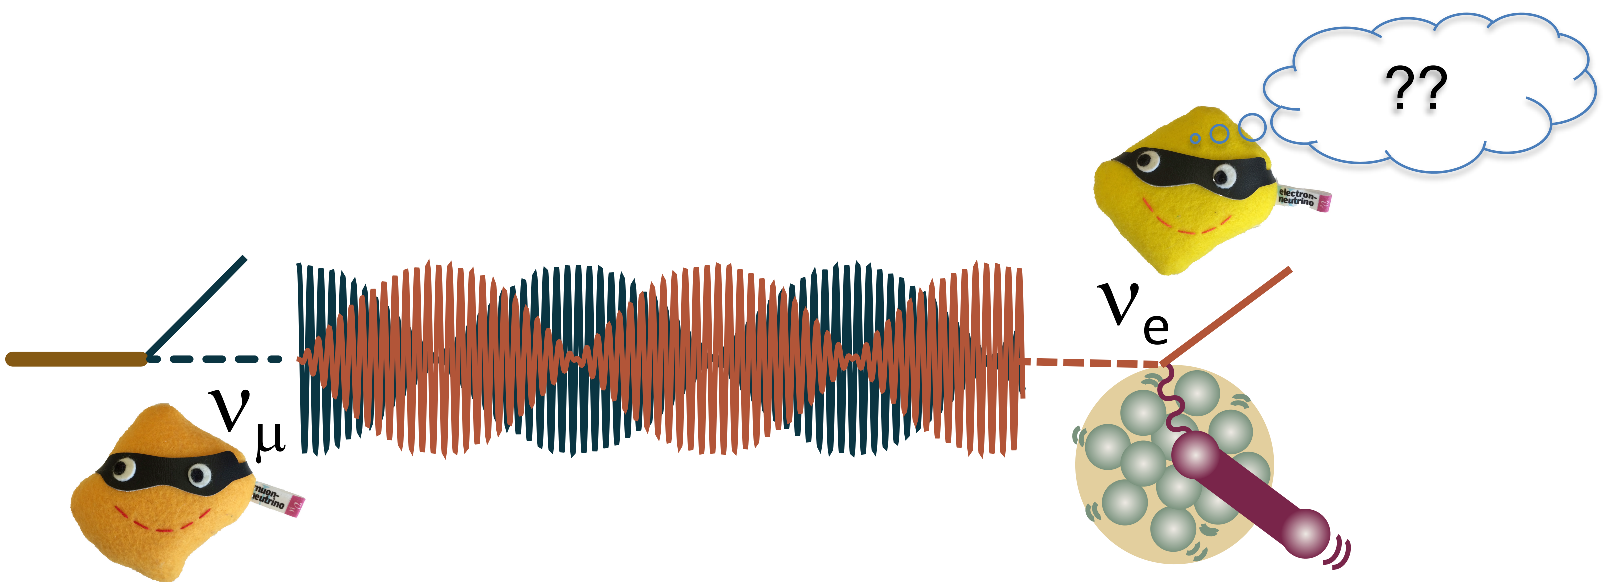

Fermilab
Slide courtesy of Heidi Schellman