DUNE Software and Computing Incomplete Overview for NPPS

Brett Viren

Physics Department

BROOKHAVEN NATIONAL LABORATORY

NPPS – 2019-06-05

Outline

Experiment

Online Computing (Far Detector DAQ)

Offline Computing

DUNE Experiment and Physics

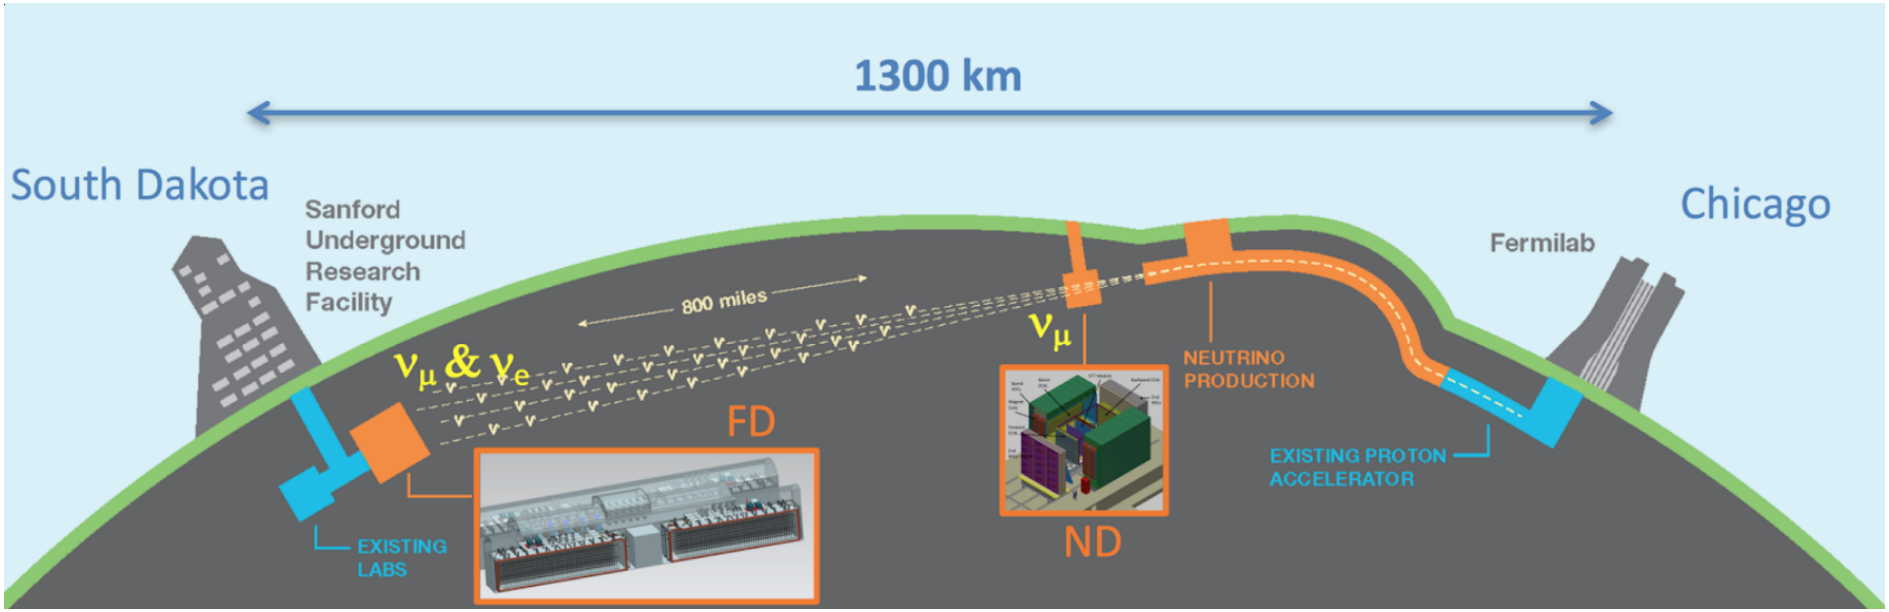

- Long-baseline neutrino beam: discover ν CP-violation, ν mass hierarchy, precision ν oscillation parameter measurements.
- Nucleon decay: targeting SUSY-favored modes ($p \to K^+\bar{\nu}$)
- Supernova Neutrino Burst (SNB): sensitive to the galaxy, sensitive to ν_e (complementary to water-Chernkov).

DUNE Far Detector Operation

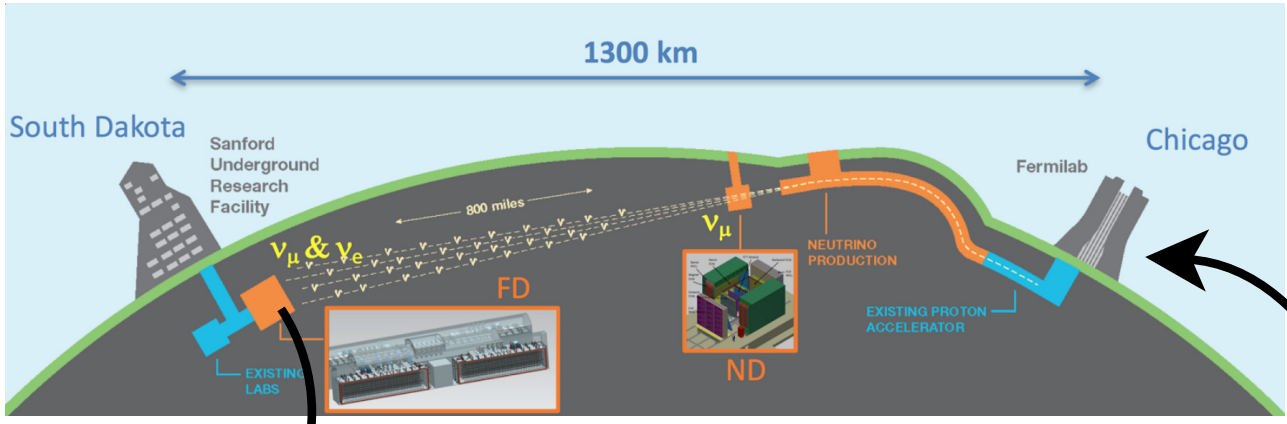

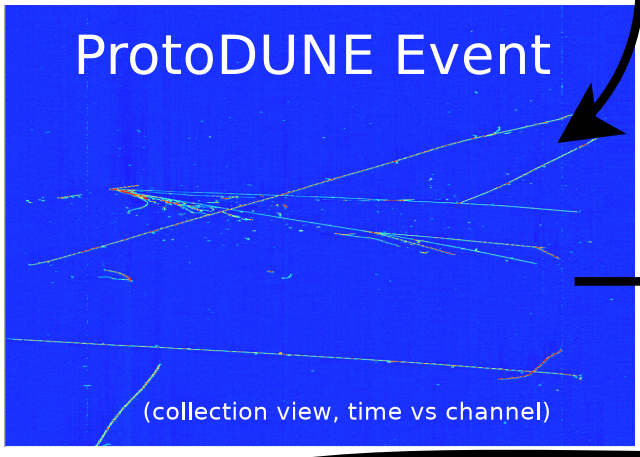

4 far detector modules,
10 kt each

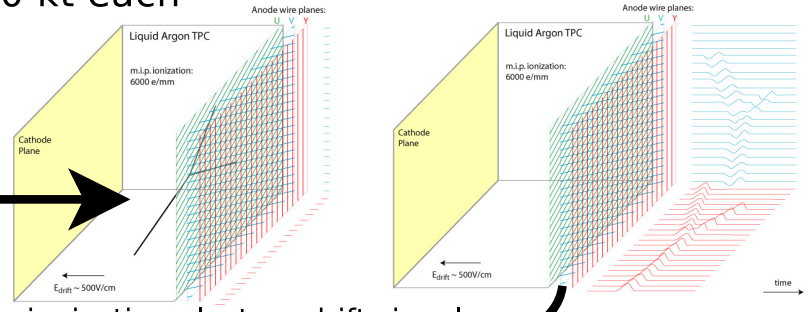

ionization electron drift signals

30 PB/year
raw data transfer
to FNAL and others

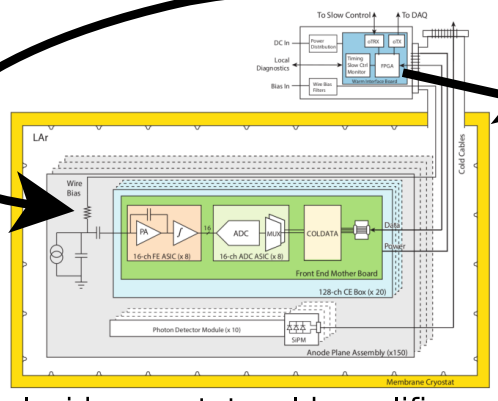

Inside cryostat: cold amplifier,
shaping, ADC, FPGA, copper out

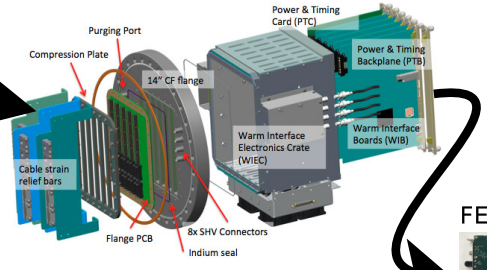

Warm digital electronics,
copper to optical,
data stream multiplexing

FELIX

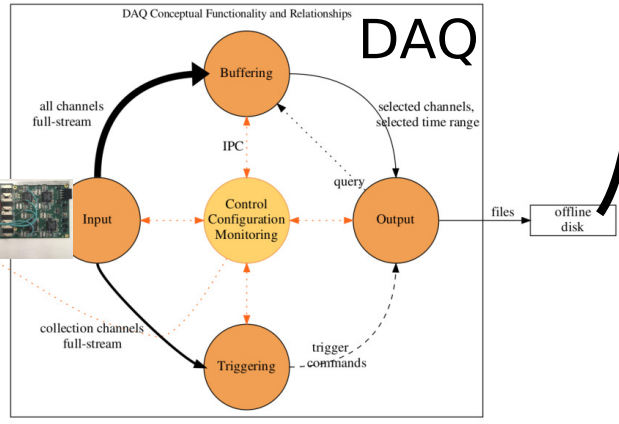

10 kT TPC module:
384,000 channels x 2 MHz x 12 bit = 9.2 Tbps

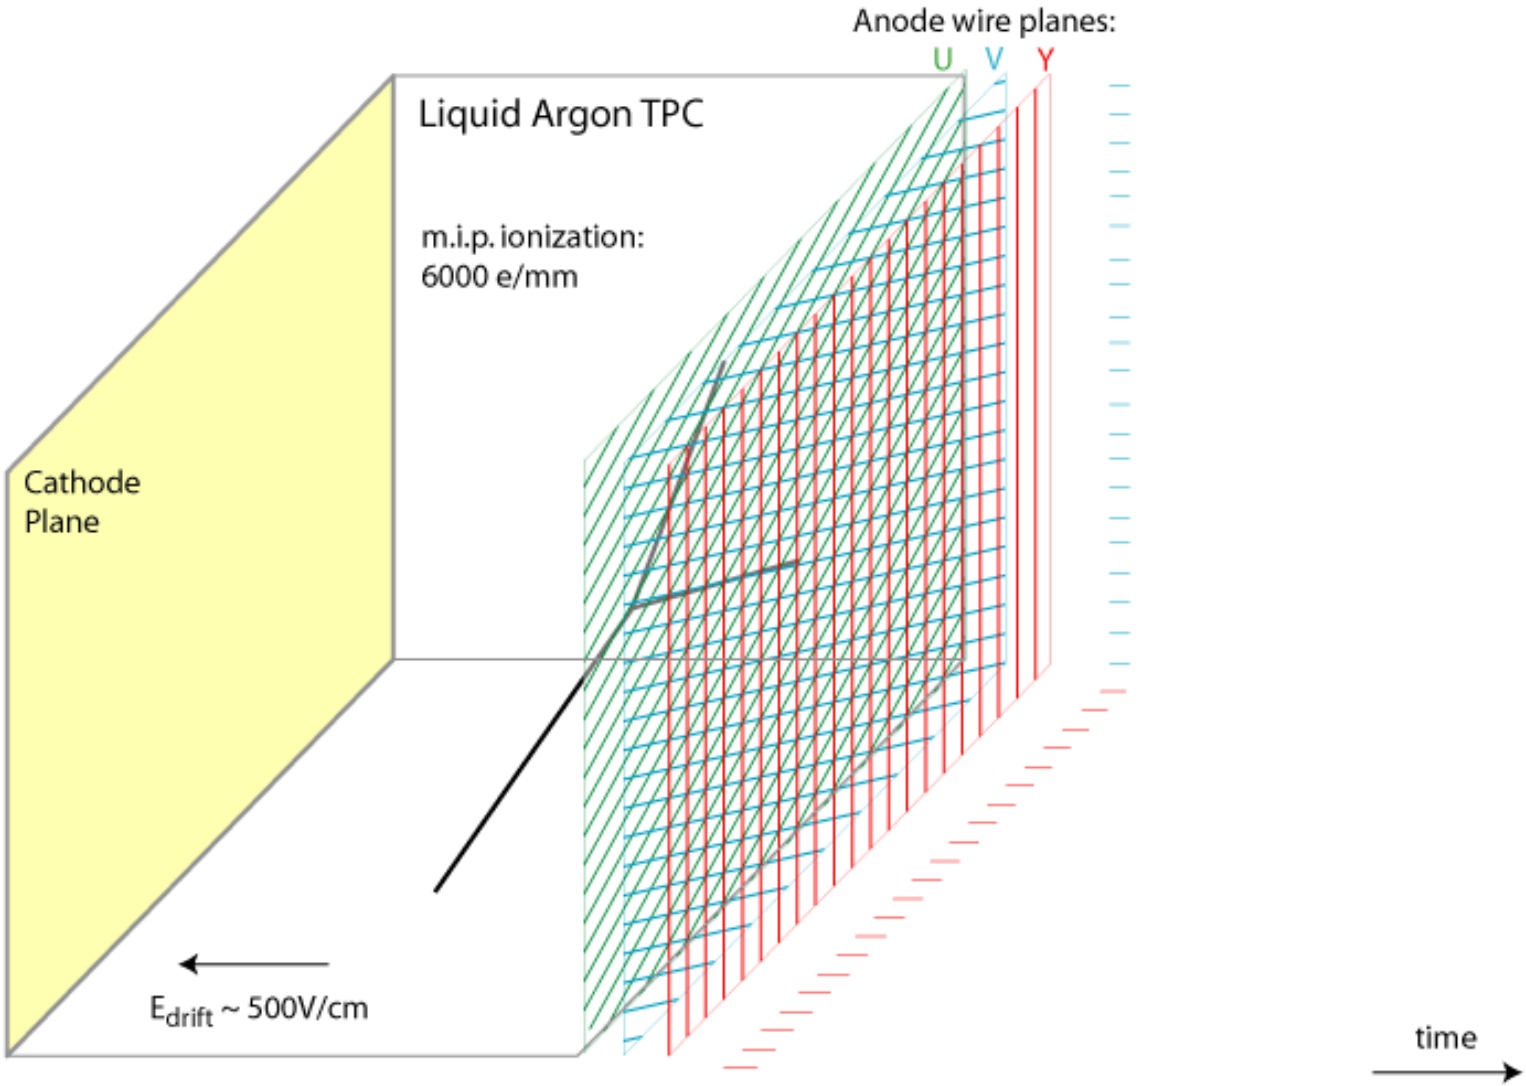
Anode wire planes:
U
V
Y
Liquid Argon TPC
m.i.p. ionization:
6000 e/mm
Cathode
Plane
E_drift ~ 500V/cm
time

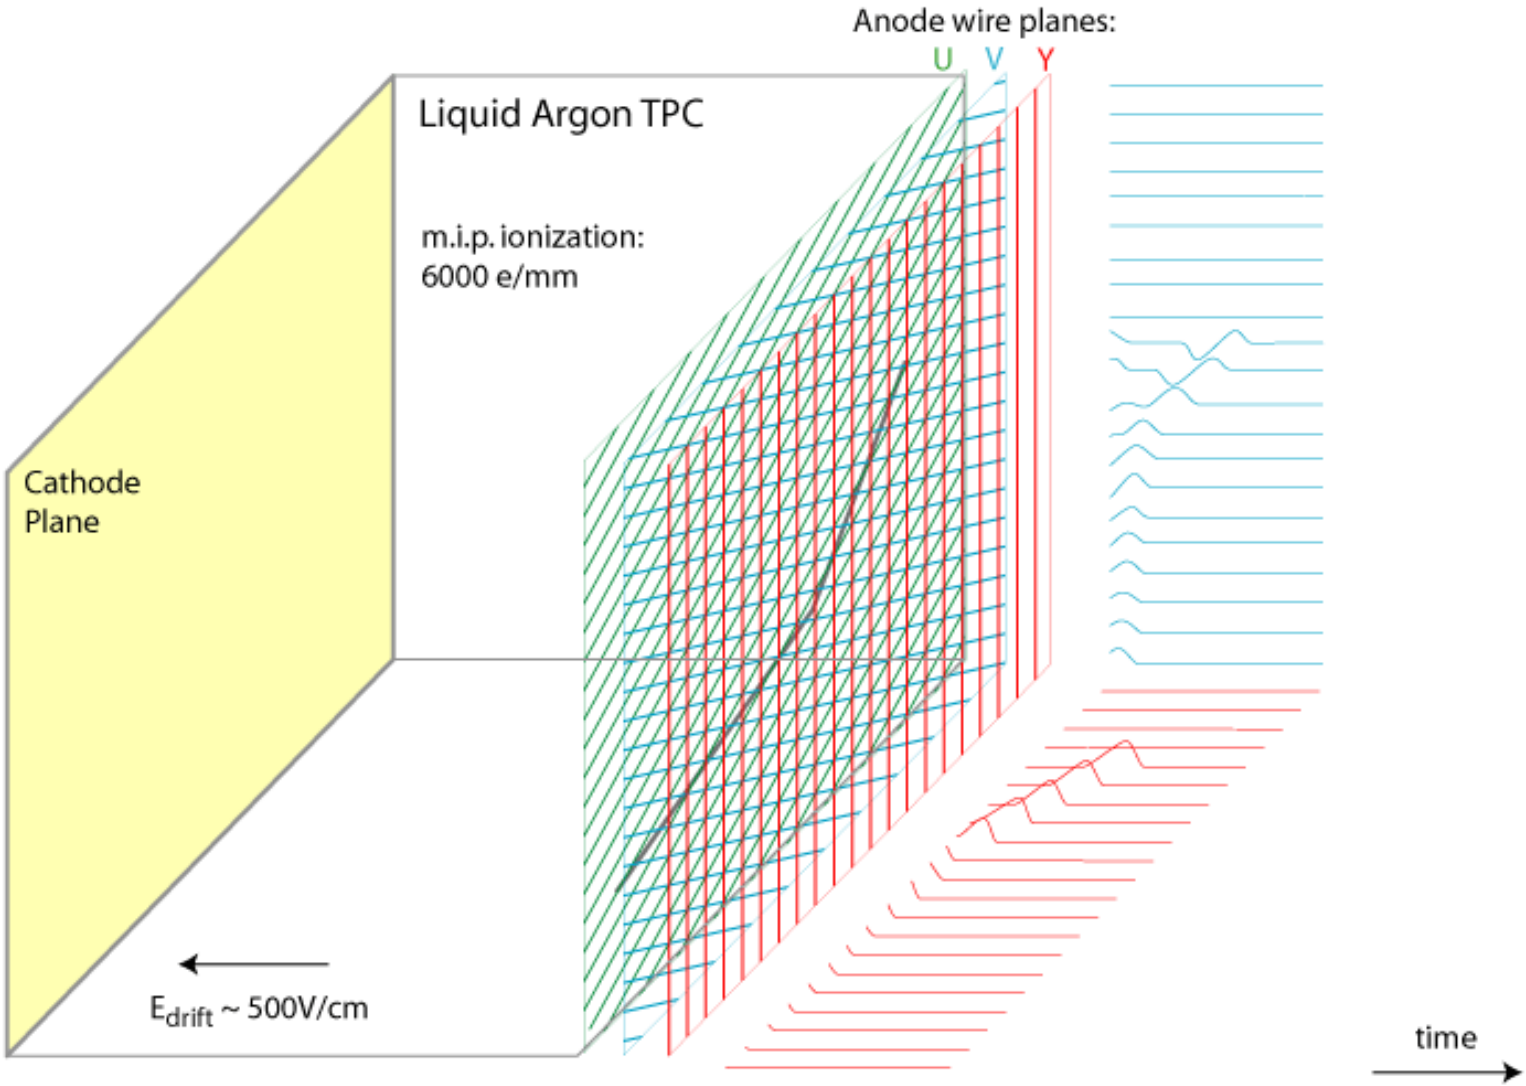
Anode wire planes:
U
V
Y
Liquid Argon TPC
m.i.p. ionization:
6000 e/mm
Cathode
Plane
E_drift ~ 500V/cm
time

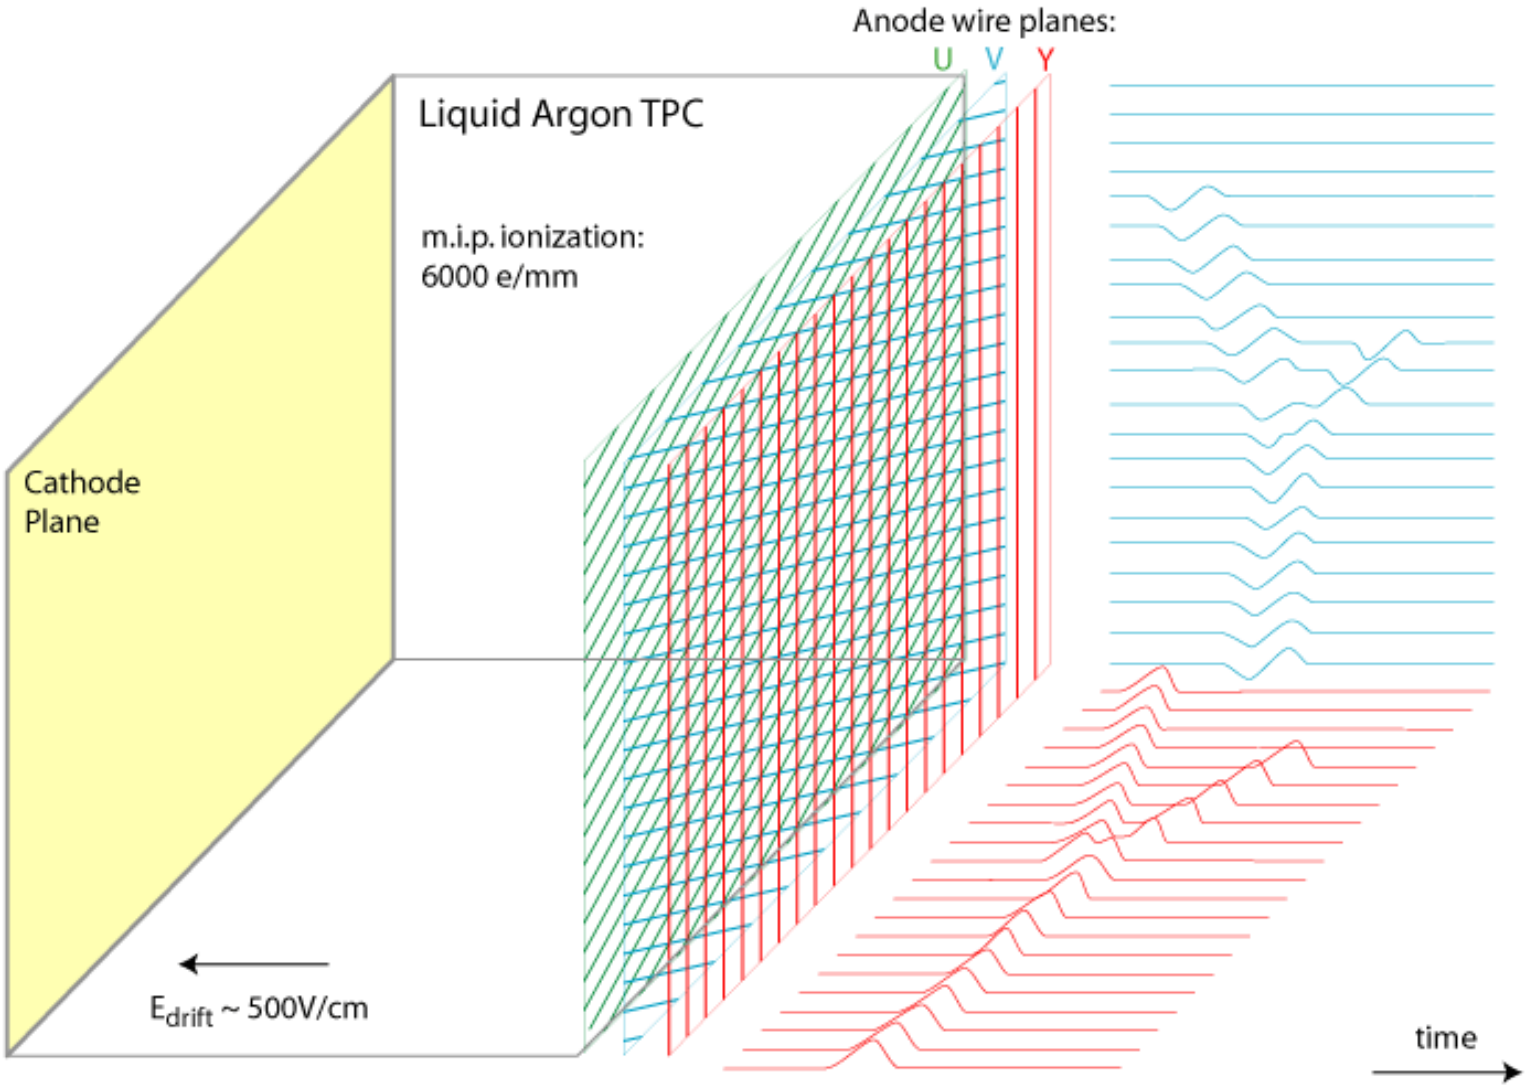
Anode wire planes:
U
V
Y
Liquid Argon TPC
m.i.p. ionization:
6000 e/mm
Cathode
Plane
E_drift ~ 500V/cm
time

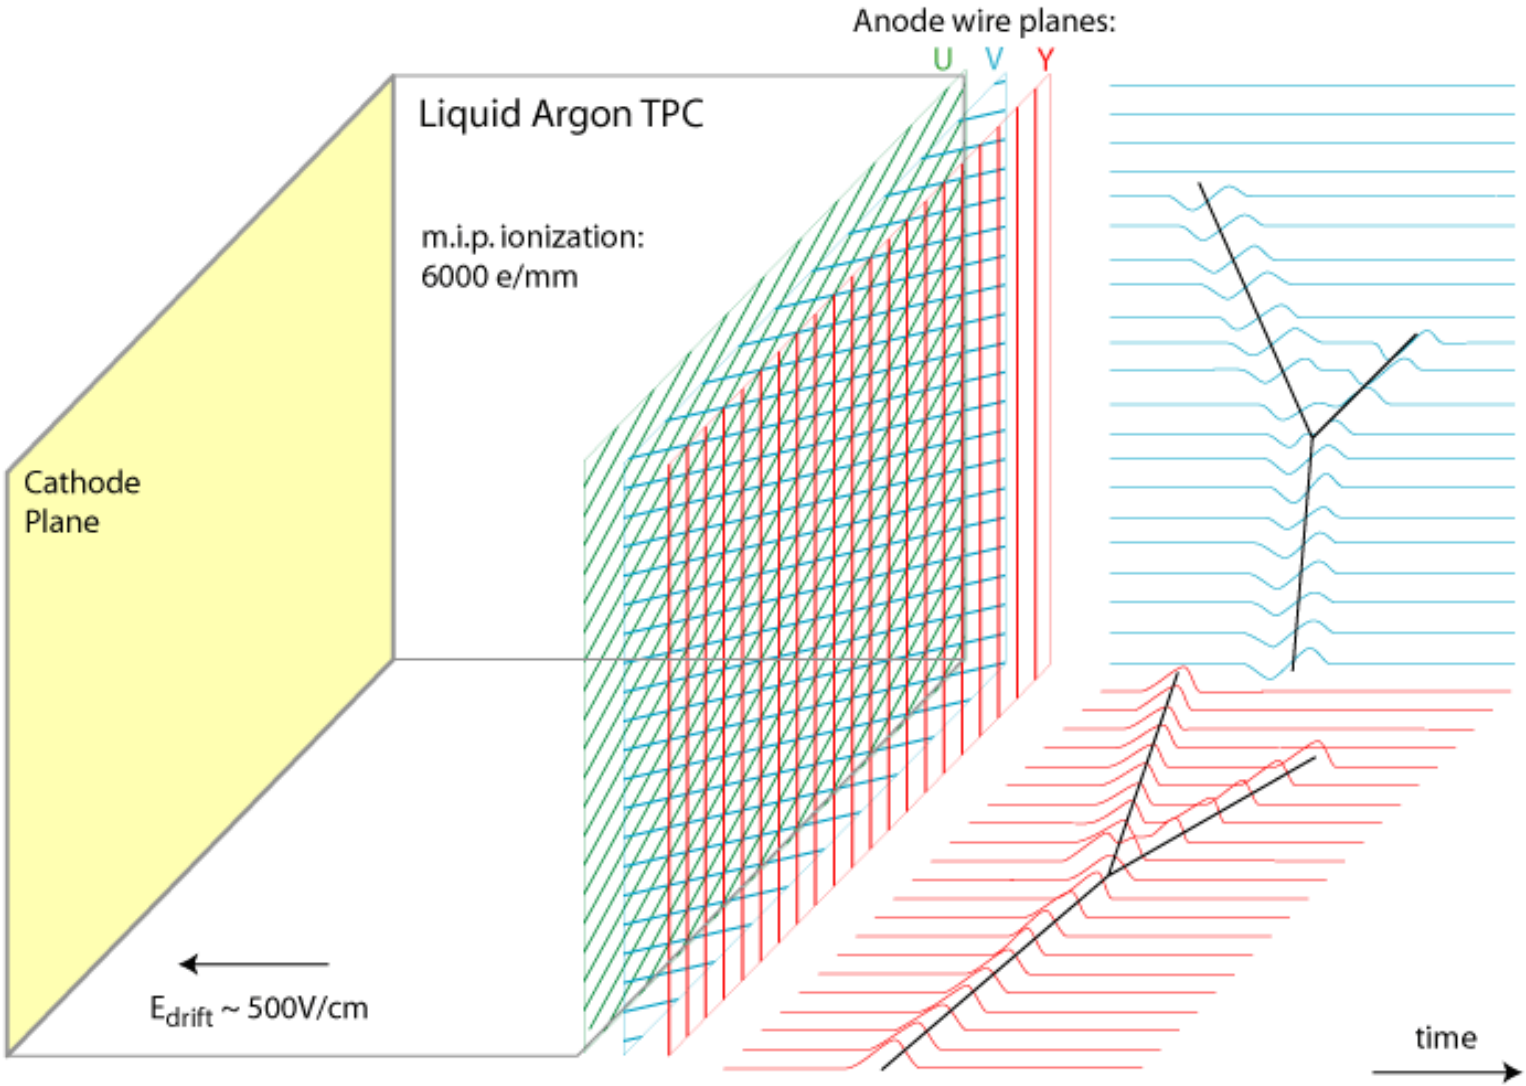
Anode wire planes:
U
V
Y
Liquid Argon TPC
m.i.p. ionization:
6000 e/mm
Cathode
Plane
E_drift ~ 500V/cm
time

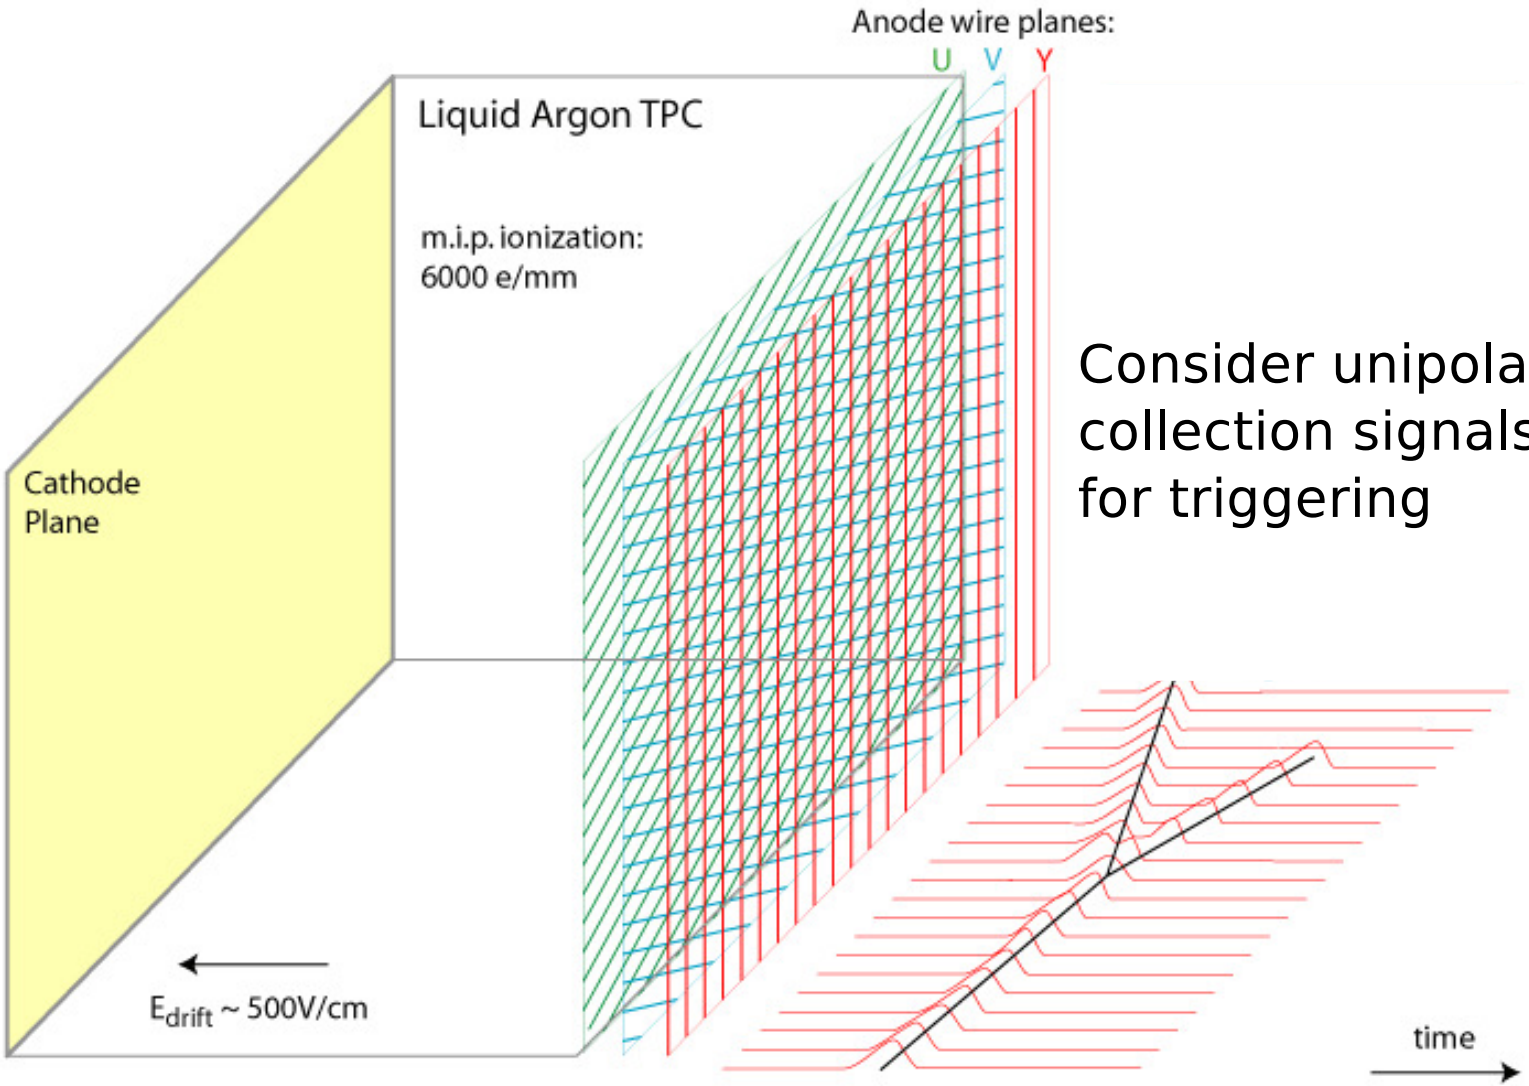

Consider unipolar collection signals for triggering

DUNE Far Detector DAQ Overview

- Common DAQ for 4 loosely-coupled 10kton modules.
  - LAr 3-plane wire readout, LAr/GAr 2-plane strip readout, future LAr pixel?
- $\mathcal{O}(10^6)$ channels, 2 MHz, 12bit waveforms, 3 – 4 TB/s into DAQ.
- Output to tape: 30 PB/year $\approx$ 3000$\times$ reduction.
- Physics drivers of DAQ design:
  - **supernova neutrino burst**: 10s of pre-trigger buffer, 100s full readout
  - **natural ^{39}Ar decay**: 10 MHz, 0.5 MeV endpoint energy, reject
- Self-triggering on ionization activity (largely software-based)
- Fermilab's **artDAQ** used now in protoDUNE, expected to provide basis for DAQ back-end.
  - considering to aggregate triggered data "event" via distributed file system directory (eg, **glusterfs**), or as entry in key-value store (eg, **DAQ-DB**), following ATLAS R&D and technology studies.

FELIX as DUNE FD DAQ Input Interface

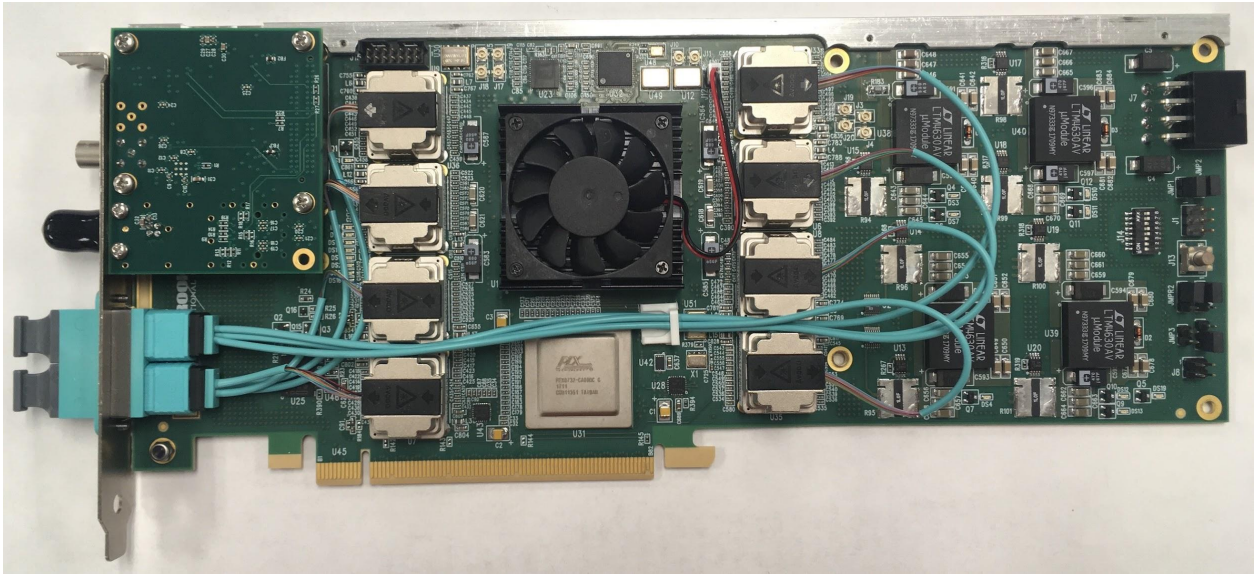

FELIX collaboration: BNL, ANL, Bologna, CERN, FNAL, Irvine, Nikhef, UCL and Weizmann
BNL: hardware design and co-development of firmware.

- Thin custom hardware and FPGA between detector electronics and DAQ's commodity computing.
- R&D shared with ATLAS, protoDUNE, sPHENIX, ~~Belle2~~, others.
- Powerful FPGA, 48 optical I/O (460 Gbps), support for daughterboard.
- Commodity host PC interface: PCIe gen3 x16, (16 GB/s).
  - goal: 75 front-end PCs, each with 2 FELIX PCIe gen3 boards (per 10kt)
  - stretch goals: reduce PC/board count by 2 – 4 with PCIe gen4

Data Transfer and File Catalog

- For prototype detectors (ProtoDUNE) transfer raw data from CERN to FNAL via **Fermi-FTS** using **SAM** for data catalog.
- For DUNE, expecting to transition to **Rucio** for online→offline and production data management.
- Replacement for SAM under consideration.

Major Offline Processing Stages

Signal processing noise filters and detector response deconvolution. Heavy use of FFTs. Output signal-regions-of-interest $\approx 100\times$ data reduction.

3D Imaging reconstruct ionization activity patterns. Fast, compressed sensing techniques.

Conventional reconstruction clustering, track/show modeling, pattern recognition.

Machine learning dense and sparse CNN, Graph NN, GANS.

Wire-Cell Toolkit

- Provides leading LAr TPC signal+noise simulation, noise filtering, signal processing, 3D imaging algorithms.
  - pattern recognition, charge/light matching in prototype
- Toolkit supports *data flow programming* paradigm.
  - dynamic plugin system, comprehensive configuration via Jsonnet
- Abstract DFP graph execution engine, multiple implementations
  - Default is low-memory, single-threaded
  - Experimental multi-thread based on Intel TBB
  - Future, multi-node engine possible
- Developed and maintained by BNL.
  - Initially for MicroBooNE, now for ProtoDUNE, ICARUS, DUNE....
- Runs stand-alone CLI or embedded in Fermilab's *art*/LArSoft framework

WCT Job Graph for ProtoDUNE 3D Imaging

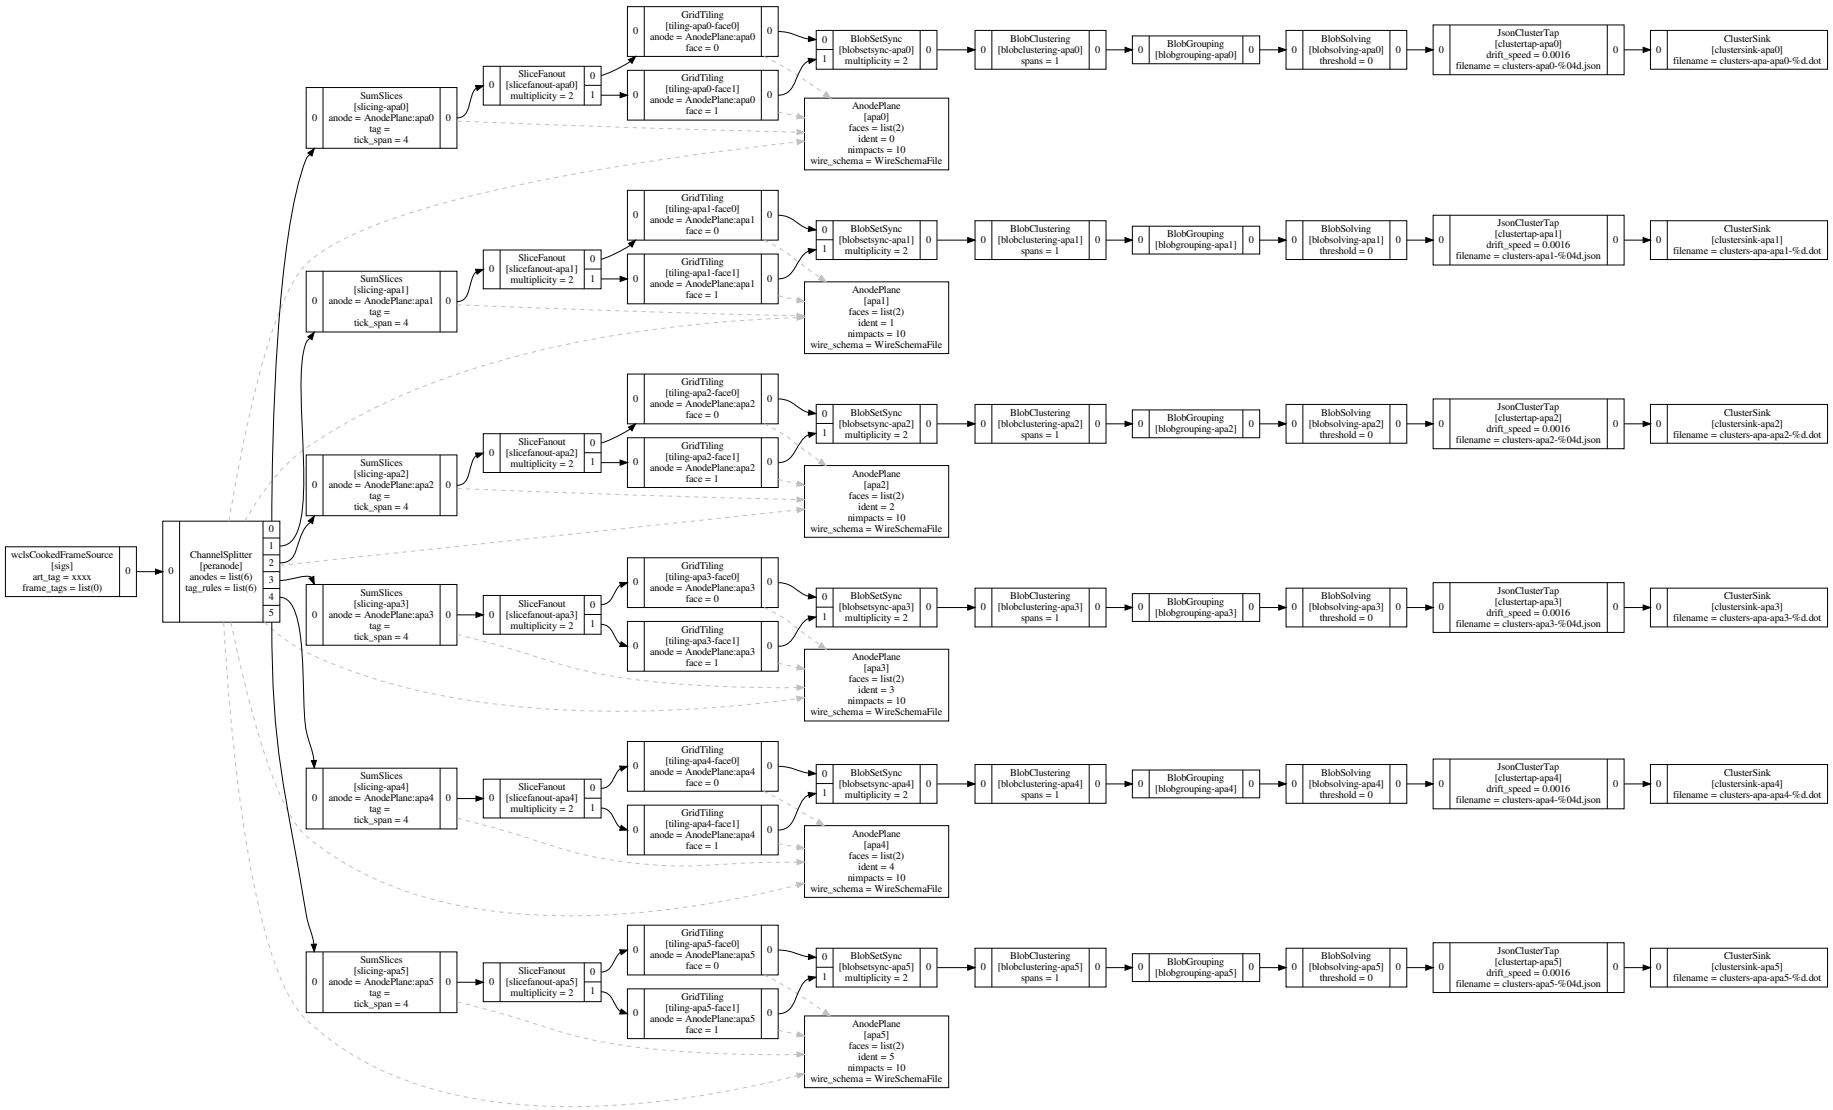

"Bee" BNL's Web-based Visualization

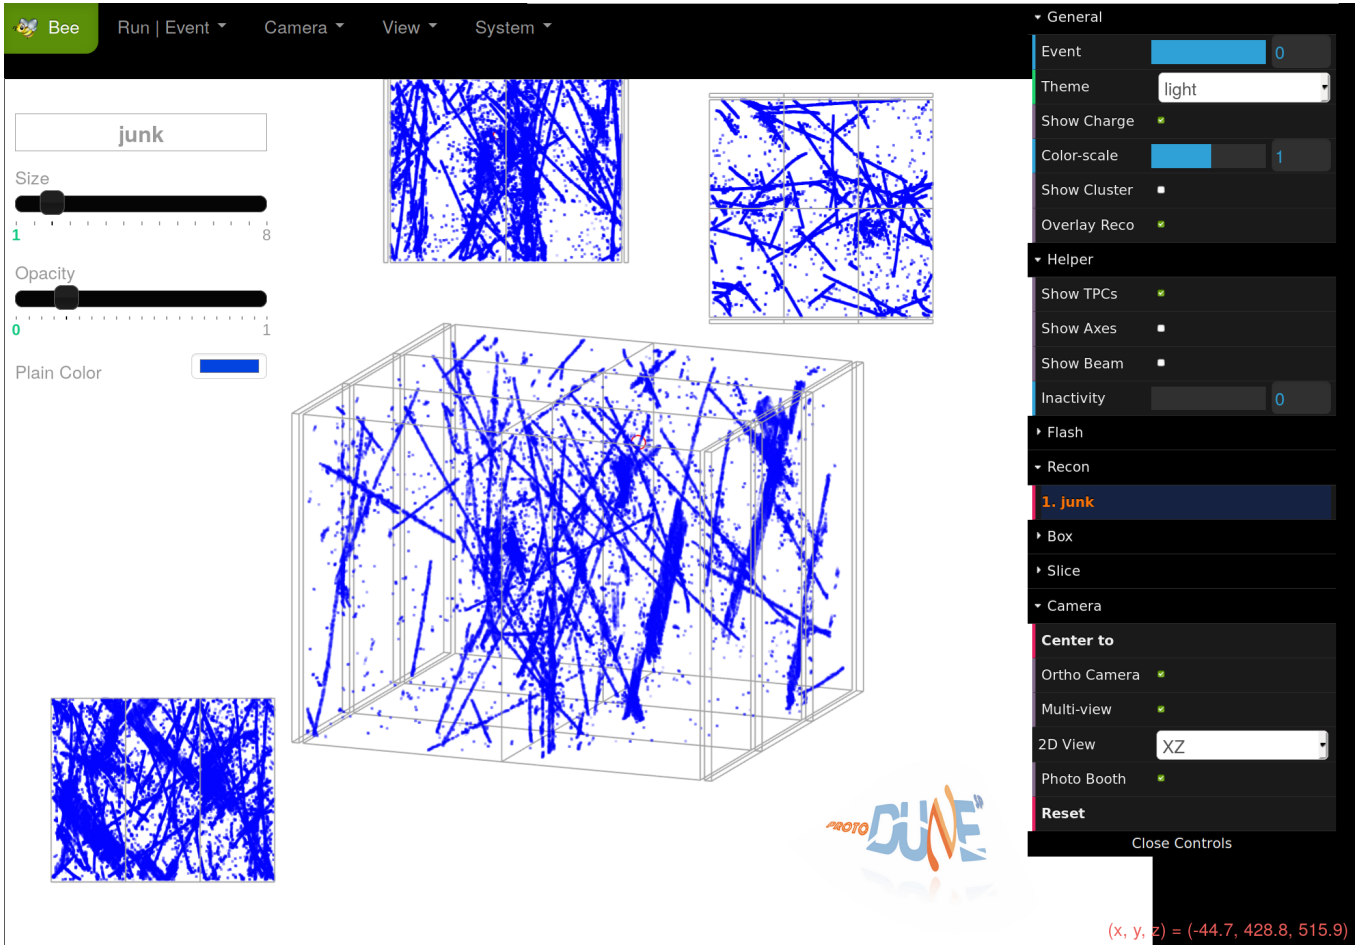

https://www.phy.bnl.gov/wire-cell/bee/set/bccdc2d7-16e1-4363-8034-032d9fe2de50/event/0/

Developed/maintained by Chao Zhang. WebGL, GPU accelerated. PC/phone browsers. JSON data, user uploads.

art/LArSoft

art Fermilab's event processing framework forked from CMSSW and used by several Intensity Frontier experiments. Fermilab is now considering to "merge" *art* back into CMSSW (or something).

LArSoft "application" layer for general LArTPC detectors, shared by multiple experiments, provides data model, *art* "modules" (equiv to Gaudi "algorithms") and services. Maintained by Fermilab, contributions from many. Includes a Wire-Cell Toolkit integration layer.

dunetpc Further "application" layer on top of LArSoft with DUNE-specific modules/services.

- Provides Wire-Cell its interface to Geant4 and raw detector data.
  - *art*/LS support effectively required for running code in production DUNE jobs.
- LArSoft includes "competing" drift simulation + signal processing based on simple, 1D field response functions.
  - somewhat faster but less correct compared to WCT's 2D versions.

art/LArSoft/dunetpc/WCT package deps

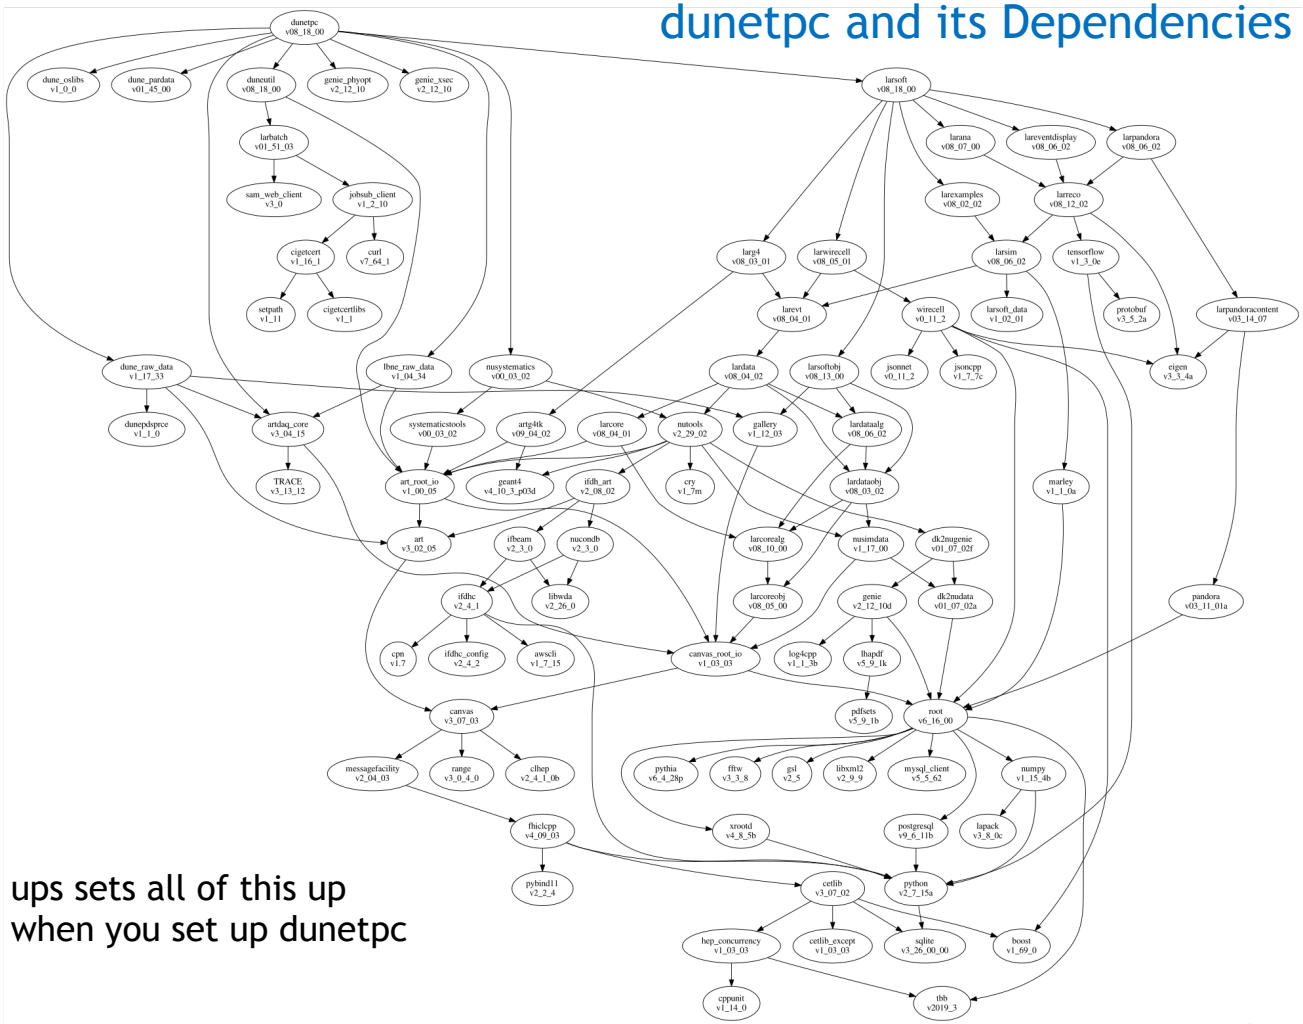

Mike Kirby, DUNE Collab Meeting May 2019.

Coordination with Worldwide LHC Computing Grid

- WLCG undergoing a transition in the organizational structure towards a Scientific Computing Infrastructure Steering Committee
- Approved middleware and interfaces (CREAM, ARC, HTCondorCE, dCache,EOS, SRM, etc)
- Tickets (GGUS), registry (GOCDB), and accounting (APEL) from EGI

- DUNE is now an **observing member** of the WLCG Management Board - announced at the HSF/OSG/WLCG Workshop at JLab this spring
- will participate in Global Deployment Board (GBD) for infrastructure planning
- DUNE has already started to participate at an operational level through GGUS tickets - report and aid in identification of site specific issues (UK and French sites currently)
- Hoping to have a GDB meeting in the US later this year with focus on DUNE-relevant areas
- **Technical contributions** made with development and evaluation of services (not focused on getting everything into production operations)
  - Edinburgh + FNAL: development work on RUCIO for Data Management
  - RAL-PPD + Manchester: exploring DIRAC capabilities for Workload Management (FNAL kx509 certs and file catalog filed w/ metadata from SAM)
  - RAL-PPD: Running small MC simulation on WLCG and RAL SE integration

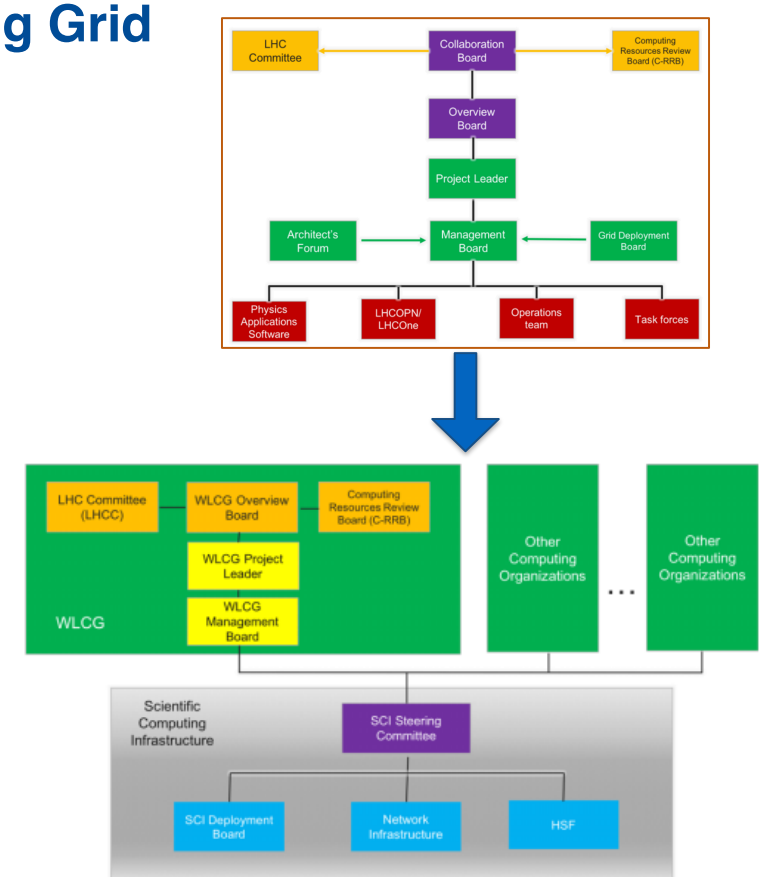

Fermilab DUNE

Mike Kirby, DUNE Collab Meeting May 2019.

Initial DUNE FD CPU Estimate

- Current protoDUNE-SP data production
  - 42 M triggered events. 1.8 PB raw data sample.
  - Processed 8M "good" events in 2.5 M core-hours (300 core-years)

  $\Rightarrow$ 1500 core-years for full sample
- DUNE 10kt module estimate $\sim$8000 core-year/year-of-data.

As is, it is modest but lacks some things:

- Three other 10kt modules + near detector.
- Full reconstruction software still in development, new algorithms may require substantially more CPU.
- Resources to support machine learning not yet fully understood
  - Requires high-quality, somewhat slower "2D" Wire-Cell simulation and signal processing.
  - Training set size $\mathcal{N}_{MC} = X \cdot \mathcal{N}_{data}$ where X is 10-100?

DUNE / Wire-Cell HPC/GPU "strategies"

- DUNE CPU requirements are maybe modest (not ATLAS), but can DUNE afford to **not** try to exploit HPC resources?
- BNL CSI + Physics EDG with CCE funds somehow under LArSoft umbrella (?) are evaluating Wire-Cell Toolkit for HPC/GPU porting
  - Accelerate the many FFTs for signal simulation and signal processing.
  - A lot of "heuristic" code remains not GPU-friendly.
  - Will it be a good match to HPC? Will GPUs be under-utilized?
- Working w/ machine learning collaborators to use Wire-Cell
  - WCT's quality 2D sim/sigproc a must for reasonable 2D CNN.
  - WCT sigproc/ROIs and 3D imaging natural to feeds "sparse array CNN" and Graph NN techniques, friendly on limited GPU RAM
    - Dense 2D CNN typically must strongly down sample their images
  - Can we make a highly-GPU accelerated chain: WCT sim, sigproc, 3D img + DL training to effectively utilize GPU?
  - Considering adding HDF5 format support to WCT.
- New EDG post-doc on this starting next month.

DUNE T.B.D.

Some major items with computing implications are still unknown:

near detector measure neutrino cross sections, project neutrino flux to FD. Detector design not yet finalized. Data rates, CPU requirements unknown.

far detector full makeup of all 4 FD modules not yet known. For sure 1 SP, maybe 2 SP, maybe 1 DP. Maybe 4th with pixel readout, potentially producing far more data than the other technologies.

BNL/DUNE Computing

- The DUNE Computing "Consortium" is solidifying now.
- Fermilab mgt, "wants to break down the computing fortress model"
- BNL can offer expertise in many needed areas: Rucio, SAM file catalog replacement, conditions DB, workflow mgt, HPC/GPU (CSI).
- Above my pay grade: how to develop trust and a collaborative plan, agreement, something between the two labs?
- I had many good discussions during DUNE collab meeting with FNAL computing mgt, LArSoft mgt and other fellow software/computer types.
- Pervasive, positive expressions for increasing collaboration with BNL on DUNE software and computing.

DUNE Computing Meetings

Upcoming workshops:

Data Model 14-16 Aug 2019 at BNL. data structure, DAQ/Offline handoff, production processing requirements.

Computing Model 9-11 Sep 2019 at FNAL (likely). computing infrastructure, contributions from other lab/uni.

Weekly Monday meetings (via Zoom):

- 10:00 "global computing sites"
  `https://indico.fnal.gov/category/827/`
- 10:30 "core computing"
  `https://indico.fnal.gov/category/496/`