Proposal to add DUNE to the OSG Council

Ken Herner for the DUNE Collaboration

CHEP 2019

13 Dec 2019

DUNE Introduction

DUNE is an international large-scale neutrino experiment hosted by Fermilab

Three major discovery areas

Origin of Matter

DUNE scientists will look at the differences in behavior between neutrinos and antineutrinos, aiming to find out whether neutrinos are the reason the universe is made of matter.

Unification of forces

DUNE's search for the signal of proton decay—a signal so rare it has never been seen—will move scientists closer to realizing Einstein's dream of a unified theory of matter and energy.

Black hole formation

DUNE will look for the gigantic streams of neutrinos emitted by exploding stars to watch the formation of neutron stars and black holes in real time, and learn more about these mysterious objects in space.

https://news.fnal.gov/wp-content/uploads/dune-fact-sheet.pdf

DUNE and ProtoDUNE

- DUNE
  - Future long-baseline neutrino experiment; near (FNAL) and far (SURF) detectors
  - Far det: 4 liquid argon TPCs
- ProtoDUNE
  - Two LAr TPC detectors, 1/20 size of regular DUNE far detectors
  - Single-phase operational in 2018
  - Dual-phase operational in 2019
  - Beam tests in 2018; another post-LS2

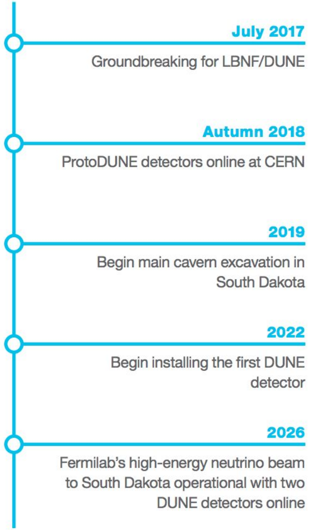

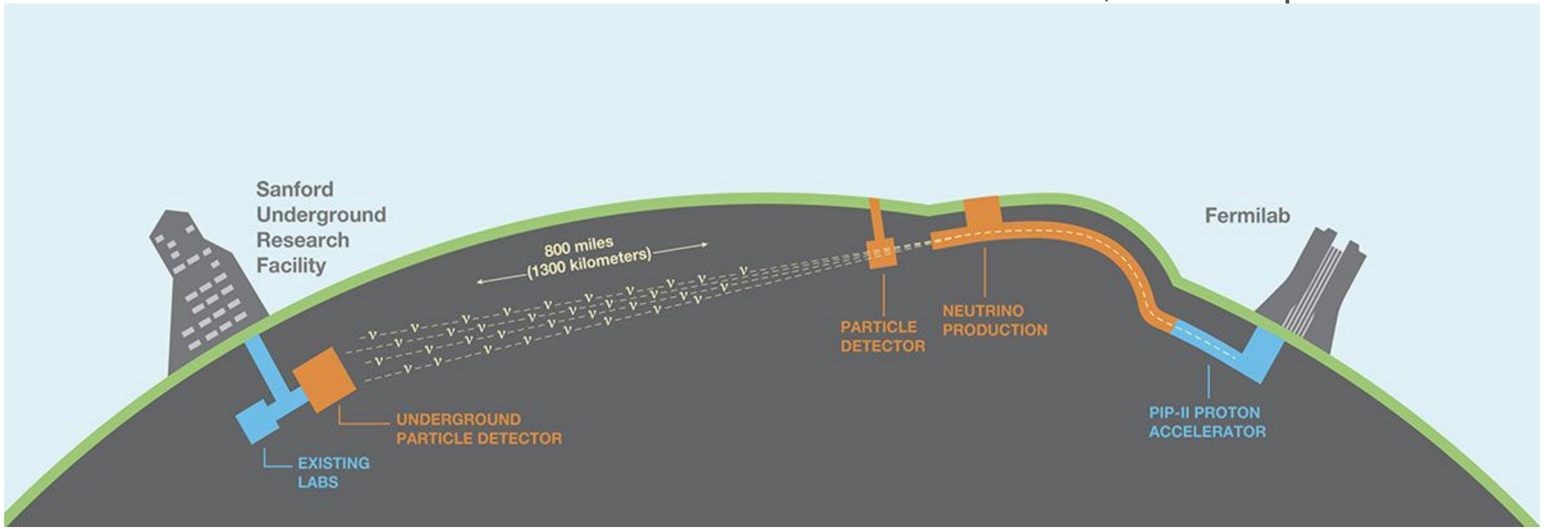

Fermilab

Far Detector

40-kt (fiducial) liquid argon
time projection chambers
- Installed as four 10-kt modules

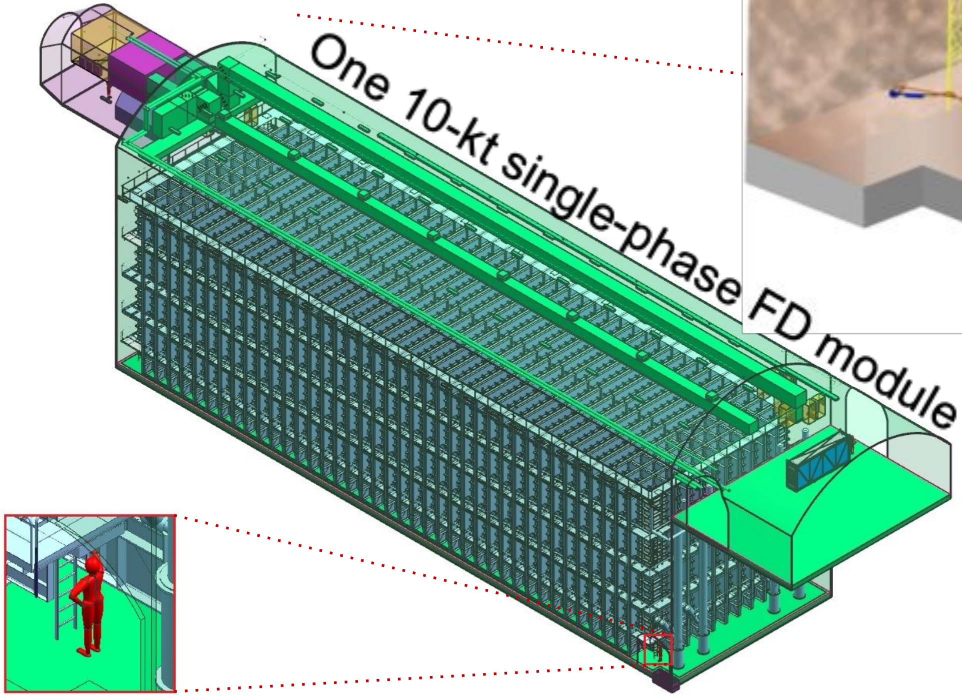

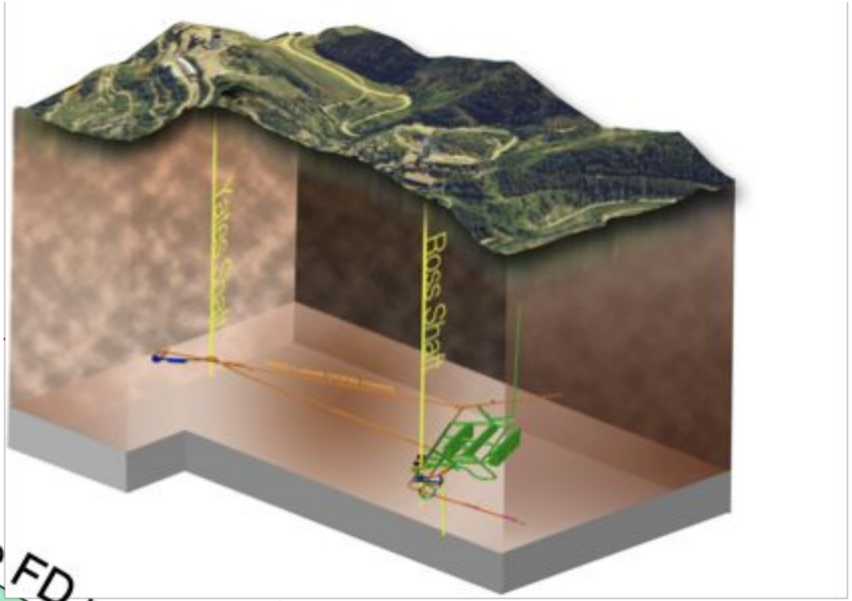

Sanford Underground Research Facility (SURF)

Ryan Patterson

- **4850' level at SURF**
- First module will be a **single phase LAr TPC**

ProtoDUNE inside EHN1 at CERN

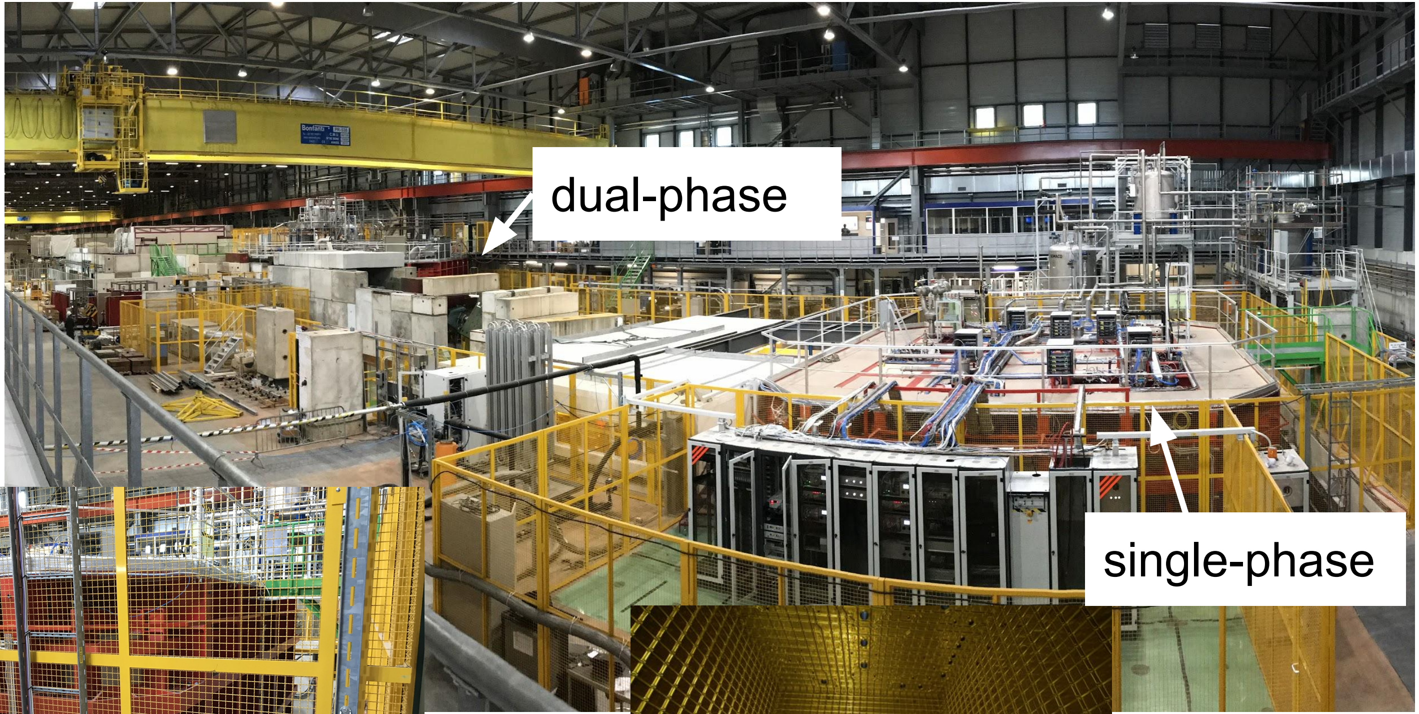

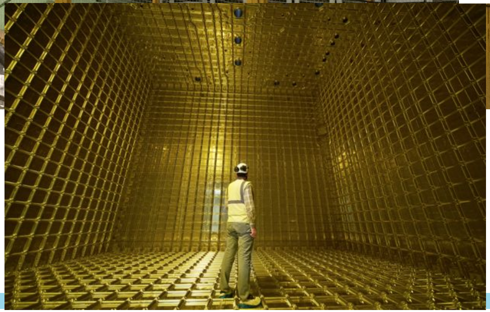

Far Detector Data Volumes

- The first far detector module will consist of 150 **Anode Plane Assemblies (APAs)** which have 3 planes of wires with 0.5 cm spacing. Total of **2,560** wires per APA
- Each wire is read out by 12-bit ADC's every 0.5 microsecond for 3-6 msec. Total of **6-12k** samples/wire/readout.
- Around 40 MB/readout/APA uncompressed with overheads **~6 GB/module/readout**
- 15-20 MB compressed/APA **~2-3 GB/module/readout**
- Read it out ~5,000 times/day for cosmic rays/calibration **~3-4PB/year/module (compressed)**

(x 4 modules x stuff happens x decade) = ….

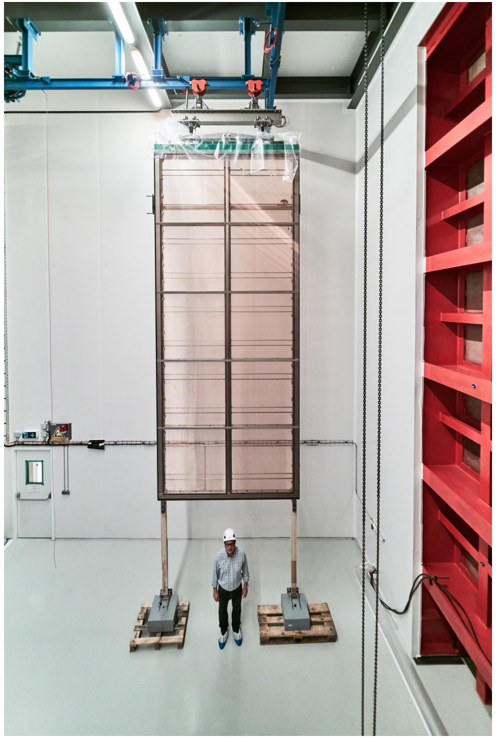

1 APA – 2,560 channels
150 of these per FD module

Fermilab

Far Detector Data Volumes

- The first far detector module will consist of 150 **Anode Plane Assemblies (APAs)** which have 3 planes of wires with 0.5 cm spacing. Total of **2,560** wires per APA
- Each wire is read out by 12-bit ADC's every 0.5 microsecond for 3-6 msec. Total of **6-12k** samples/wire/readout.
- Around 40 MB/readout/APA uncompressed with overheads **~6 GB/module/readout**
- 15-20 MB compressed/APA **~2-3 GB/module/readout**
- Read it out ~5,000 times/day for cosmic rays/calibration **~3-4PB/year/module (compressed)**

  (x 4 modules x stuff happens x decade) = ….

And there's a near detector too!

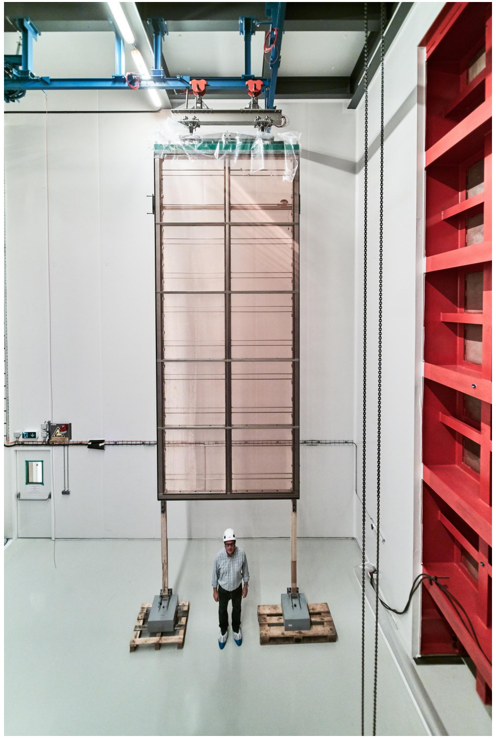

1 APA – 2,560 channels
150 of these per FD module

More fun with supernovae

- DUNE should be sensitive to nearby (Milky Way and friends) supernovae. Real ones are every 30-200 years but we expect 1 false alarm/month

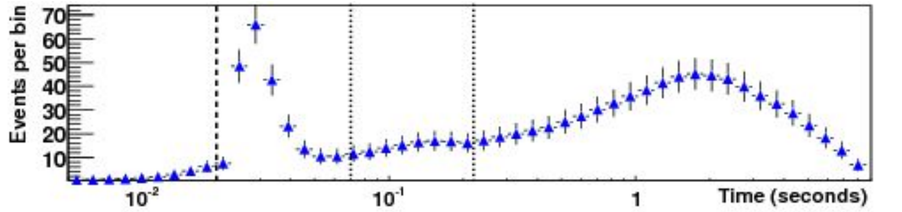

- Supernova readout = 100 sec, one trigger/month
- 100 sec readout implies
  - 1 channel = 300 MB uncompressed
  - 1 APA = 768 GB uncompressed
  - 1 module = 115 TB uncompressed
  - 4 SP modules = **460 TB** … takes 10 hrs to read at 100 Gb/s
  - Dual Phase technology has higher S/N □ smaller per module
- Some calibration runs will be similar in scope….

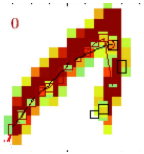

30 MeV ν_eCC

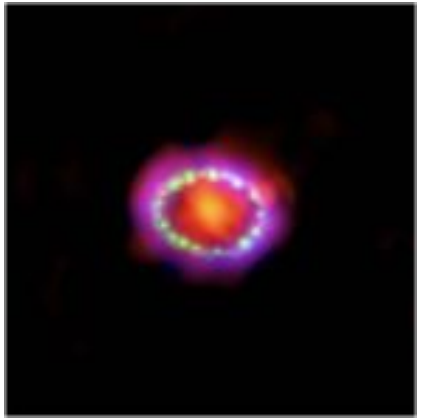

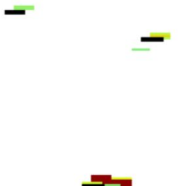

10 MeV NC

$\nu + A \rightarrow \nu + A^*$

CPU Needs

- ProtoDUNE data (with beam) more complex than future far detector data
  - Reconstruction currently typically requires 2.5 - 3.5 GB RAM; some steps can use multiple cores
- ~30 PB/yr of far detector data expected to require O(100 M) CPU hours/yr for reconstruction
  - Roughly 12k cores DC
- Reprocessing passes will be at least this much
- Simulation will be on this scale as well
- Near detector CPU requires still being formulated, but could be greater than far detector
- ...And then there's analysis. So far seeing about 50-50 analysis-production, but experience tells us that won't last
- 2021-24 will be busy with simulation, SW R&D, ProtoDUNE Run 2 processing
- All in all, **expect to be at LHC scales** (maybe not quite HL-LHC scales)

The Collaboration

Now over 1,200 collaborators in over 30 countries

Roughly the size of LHCb, ⅓ of ATLAS or CMS

Continuing to grow!

Members have significant experience with OSG from prior experiments

Fermilab

The DUNE Computing Consortium

Many of these institutions are already involved in OSG and/or WLCG

DUNE now has observer status on the WLCG management board and the GDB

Institution	Country
CBFP	Brazil
Unicamp	Brazil
York Univ.	Canada
CERN	CERN
FZU	Czech Republic
CCIN2P3	France
Indian groups	India
KISTI	Korea
Nikhef	Netherlands
Bern	Switzerland
CIEMAT	Spain
Edinburgh	UK
GridPP	UK
Manchester	UK
Queen Mary Univ.	UK
RAL/STFC	UK

Institution	Country
Argonne	USA
Berkeley	USA
BNL	USA
Colorado State	USA
CU Boulder	USA
Fermilab	USA
Florida	USA
LBNL	USA
Minnesota	USA
Northern Illinois Univ.	USA
Notre Dame	USA
Oregon State University	USA
SLAC	USA
Texas, Austin	USA

DUNE's Current Relationship with OSG

Current setup: Job submission

- Resource/slot provisioning is with GlideinWMS, widley used in OSG (setup shared with other FNAL IF and muon expts.)
- DUNE software built for both SL6/7
- Copyback is generally to FNAL dCache, other sites demonstrated
- Exploring creation of a global gWMS pool similar to CMS; would allow for additional submitter resources to come online
- OSG prescription for setting up new sites works extremely well for DUNE
- DUNE regularly reports in OSG Production meetings; KH is an AC

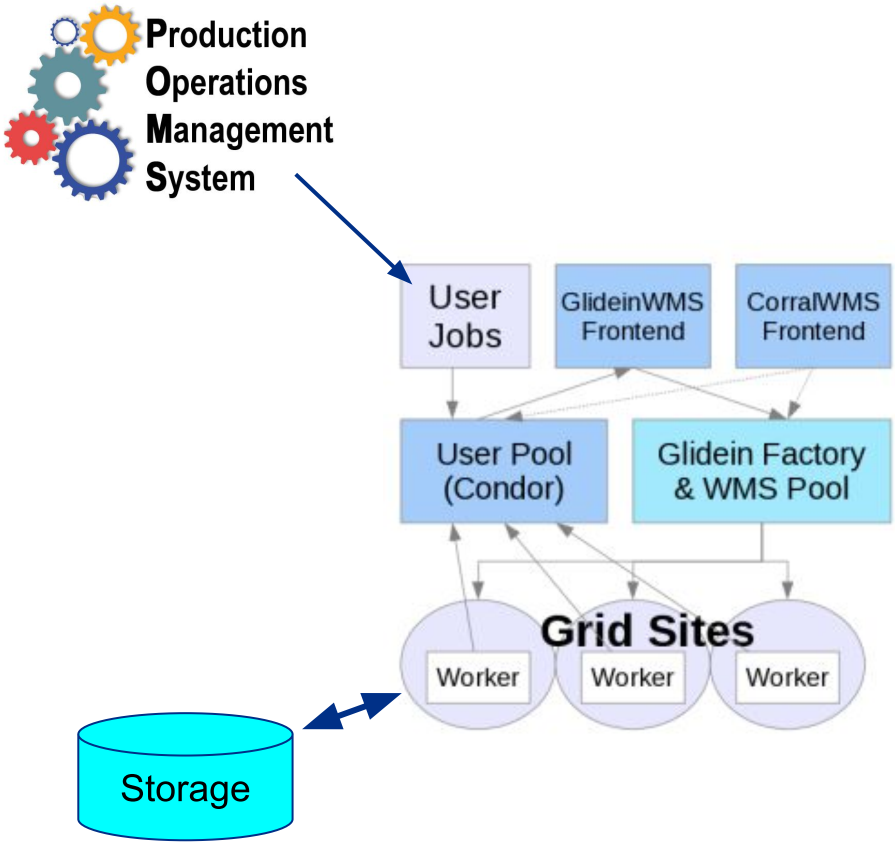

International Contributions

DUNE already getting significant contributions from international partners

In 2019 so far, **49% of production wall hours are from outside USA**

Actively working to add more sites and countries-- making this easy is critical

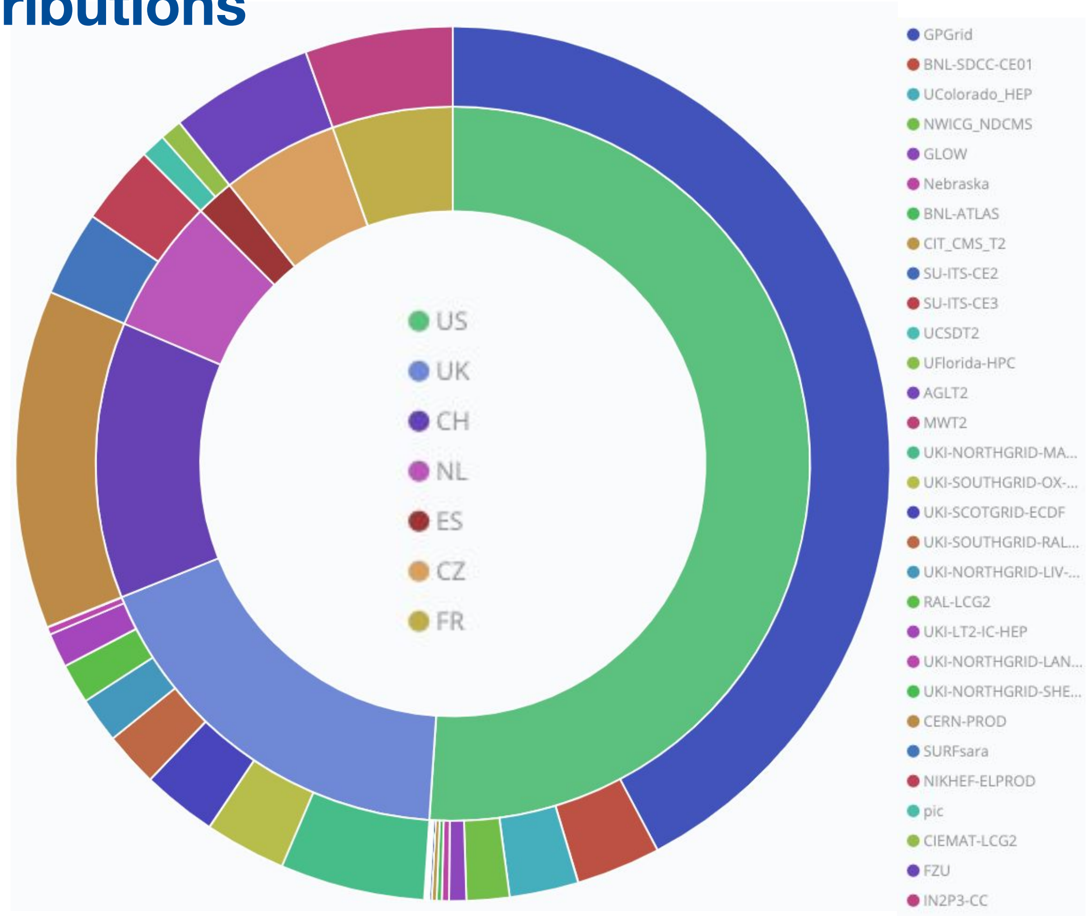

Fermilab

Current Setup: Data movement

- DUNE using the FNAL SAM system for file catalog and delivery
- Data replication being handled by Rucio instance
- Most input streamed with xrootd; output usually returned via gridftp (can easily use other protocols as needed)
- **Auxiliary file input (needed for MC generation) now handled via StashCache**; used heavily in Spring 2019 (1.75 PB transferred)

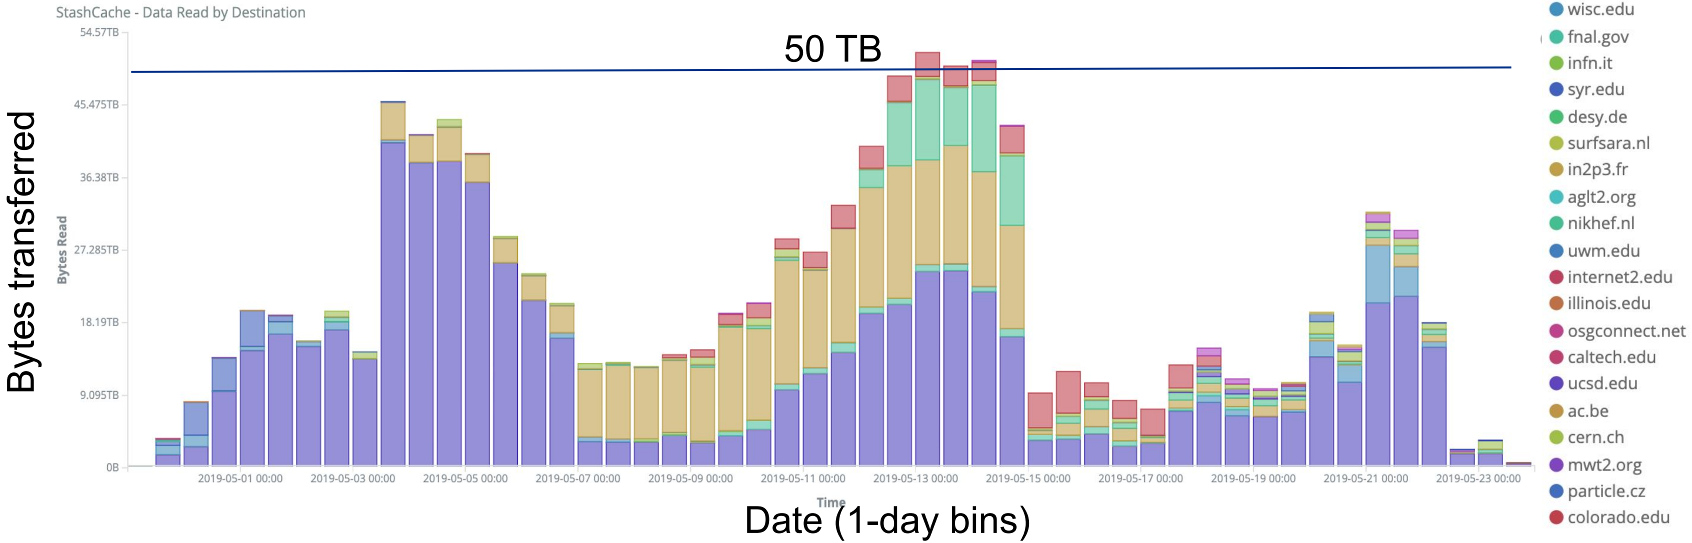

Fermilab

Setting the current scale

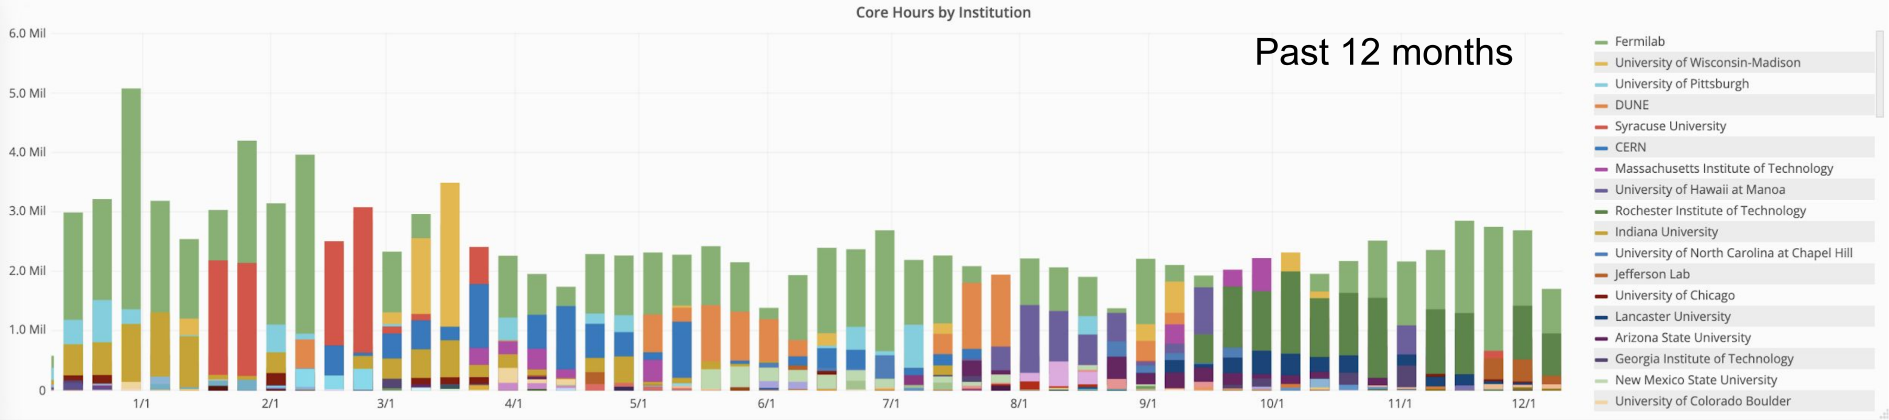

Projects

PI Name	Organization	Field of Science	Project Name	Core Hours ▾
Francis Halzen	University of Wisconsin-Madison	Astrophysics	IceCube	50.81 Mil
Thomas Robert Junk	DUNE	High Energy Physics	dune	37.37 Mil
Achille Petrilli	CERN	Particle Physics	cms.org.cern	27.74 Mil
Joe Boyd	Fermilab	High Energy Physics	nova	26.02 Mil
Joe Boyd	Fermilab	High Energy Physics	gm2	21.53 Mil

DUNE is about 75% of IceCube right now, and increasing!

How DUNE's joining the council benefits everyone

- DUNE will be largest neutrino (also largest non-LHC HEP?) experiment; represents large fraction of the US community
- **DUNE wants to utilize common solutions** wherever possible and partner with OSG, HSF, etc. on development
- DUNE will attract newer community members who may not have been involved in other large-scale HEP experiments in the past
  - DUNE's council membership will help keep these community members aware of trends in distributed computing and can help steer development in mutually beneficial ways

Summary

- DUNE will be the world's largest neutrino experiment
  - Already has world's largest LArTPC
- DUNE is successfully building on proven technologies (in many cases pioneered by OSG effort); interested in continuing to do that
  - Some new technologies and method will be required of course; shared development is ideal
- As largest neutrino experiment, DUNE will attract new community members. As they support DUNE, a strong relationship with OSG provides additional resources to everyone and sends a message that each values the other

DEEP UNDERGROUND NEUTRINO EXPERIMENT

BACKUP

- Processing chain exists and works for protoDUNE-SP
  - Data stored on **tape** at FNAL and CERN, staged to dCache in 100 event 8GB files
  - Use **xrootd** to stream data to jobs
  - Processing a 100 event 8 GB file takes ~**500 sec/event** (80 sec/APA)
    - Signal processing is < 2 GB of memory
    - Pattern recognition is 2-3 GB
  - Copy 2 GB output back as a single transfer.
  - TensorFlow pattern recognition likes to grab extra CPU's (fun discussion)
- Note: ProtoDUNE-SP data **rates** at 25 Hz are equivalent to the 30 PB/year expected for the full DUNE detector. (Just for 6 weeks instead of 10 years)
- ProtoDUNE-DP
  - Data transfer and storage chain operational since August – up to 2GB/s transfer to FNAL/IN2P3
  - Reconstruction about to start

CPU needs

RECONSTRUCTION

- ProtoDUNE events are more complex than our long term data.
  - ~**500** sec to reconstruct 75 MB compressed – 7 sec/MB
  - For FD, signal processing will dominate at about 3 sec/MB
  - < 30 PB/year of FD data translates to **~100 M CPU-hr/year**
  - That's ~ **12K cores** to keep up with data. But no downtimes to catch up.
- Near detector is unknown but likely smaller.

ANALYSIS (Here be Dragons)

- NOvA/DUNE experience is that data analysis/parameter estimation can be very large
  - ~ 50 MHrs at NERSC for NOvA fits

LAr TPC Data Processing

- hit finding and deconvolution
  - **x5 (ProtoDUNE) -100 (Far Detector)** data reduction
  - Takes 30 sec/APA
  - Do it 1-2 times over expt. lifetime
- Pattern recognition (Tensorflow, Pandora, WireCell)
  - Some data expansion
  - Takes ~30-50 sec/APA now
  - Do it ? times over expt.
- Analysis sample creation and use
  - multiple^2 iterations
  - Chaos (users) and/or order (HPC)