DUNE computing

Andrew McNab
University of Manchester

for the DUNE Computing Consortium

Overview

- This talk focuses on DUNE computing, rather than software

- DUNE recap: protoDUNE, Near and Far Detectors
- Timescales
- Scale of the problem
- Current computing status
- Technologies

DUNE recap

- *"Make neutrinos at FNAL then detect some of them in South Dakota (and maybe supernovae and proton decays)"*
- Really 3 detectors: protoDUNE; Near and Far Detectors

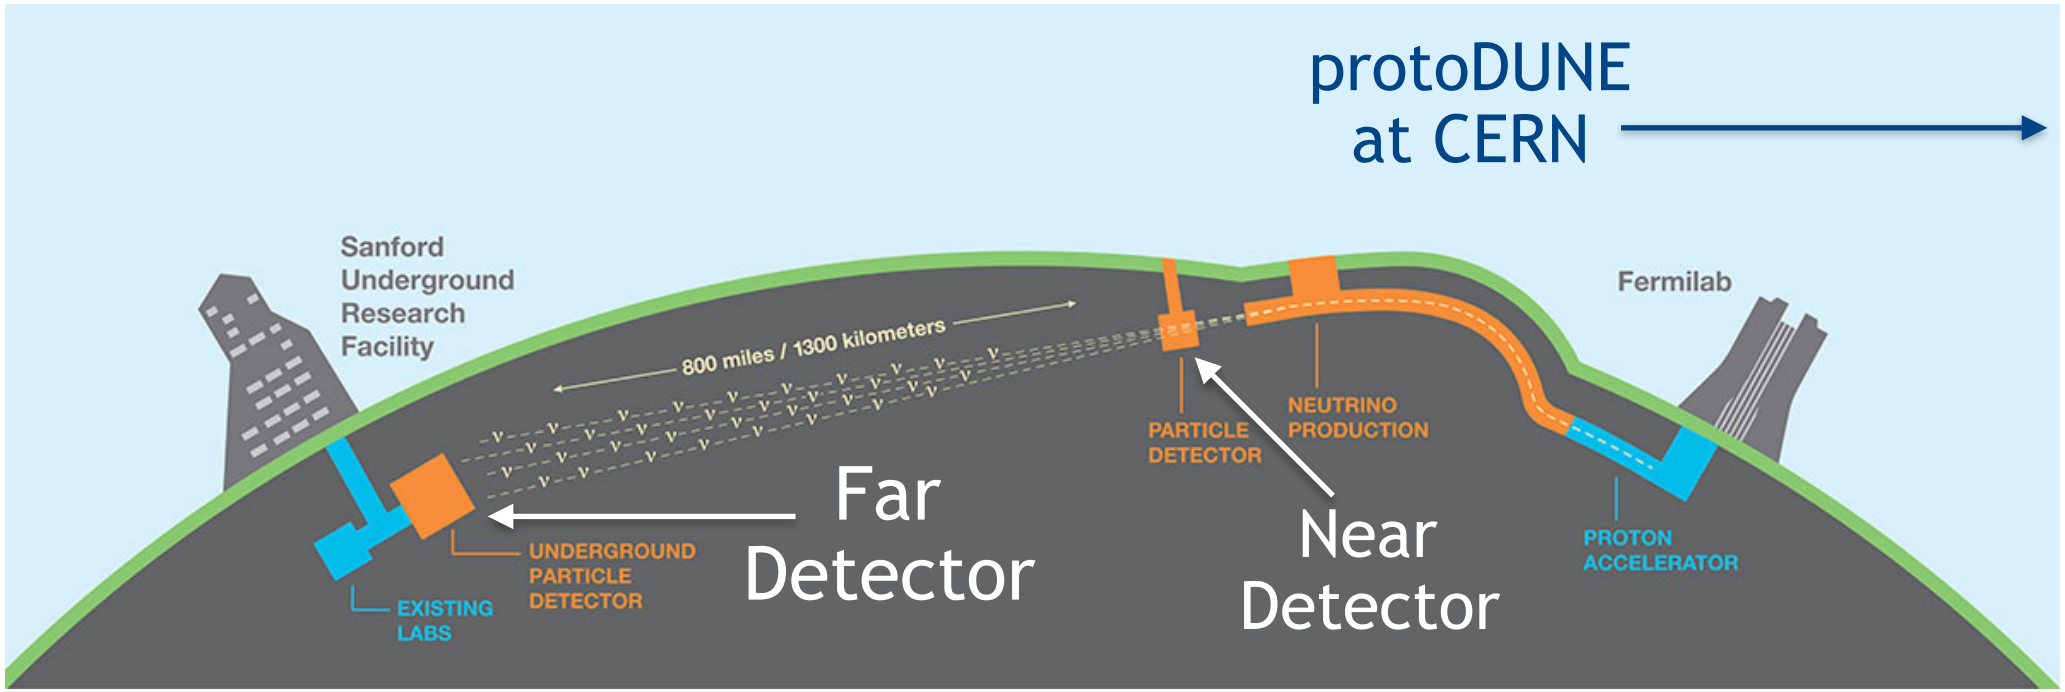

protoDUNE and Far Detector

- The eventual Far Detector will comprise four modules, each ~17000 tons of liquid argon
  - Time Projection chambers, either Single (SP) or Dual Phase (DP), plus photon detection for timing

- protoDUNE SP and DP at CERN are being used to test full size components
- But much smaller than an FD module: only 800 tons of LAr
  - SP is still the largest monolithic LAr TPC

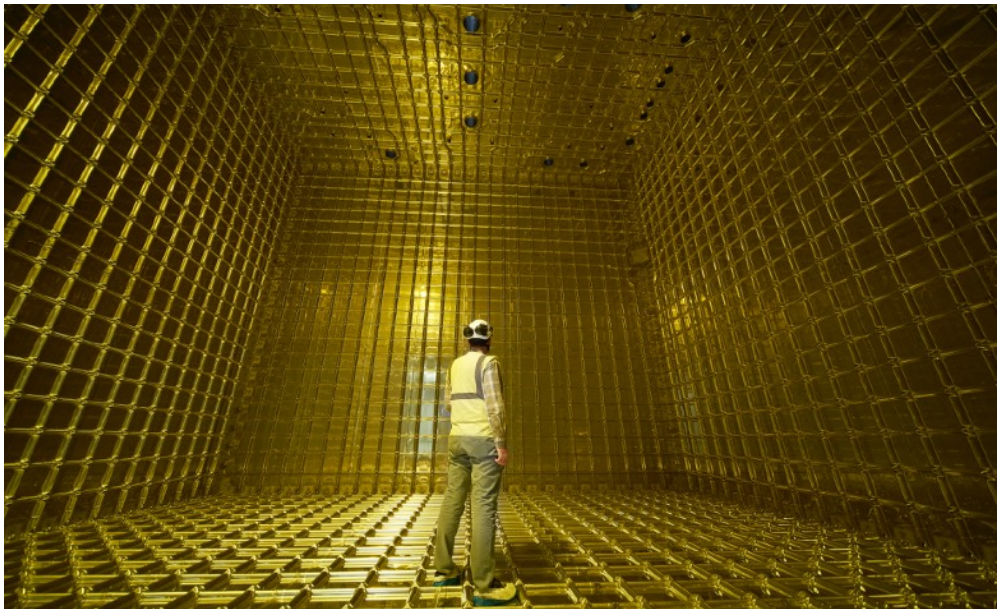

Timescales

CERN	protoDUNE-SP data taking with beam (and then cosmics)	Autumn 2018
CERN	protoDUNE-DP installed and start cosmics	Summer 2019
CERN	protoDUNE-II-SP and DP installed. Cosmics, and then beam	2021/22
S.Dakota	DUNE FD #1 installed. Cosmics	2025
S.Dakota	DUNE FD #2 installed. Cosmics	2026
S.Dakota	DUNE FD #1 and #2 data taking with FNAL beam	2026

protoDUNE events

- The TPC is continually sensitive, with triggering deciding what snapshots to record
- Quite like photographing a diffusion cloud chamber!
- The picture shows a protoDUNE event, with a beam neutrino causing a cascade but lots of cosmic background tracks
- More like 3D images than hits in multiple sub detectors

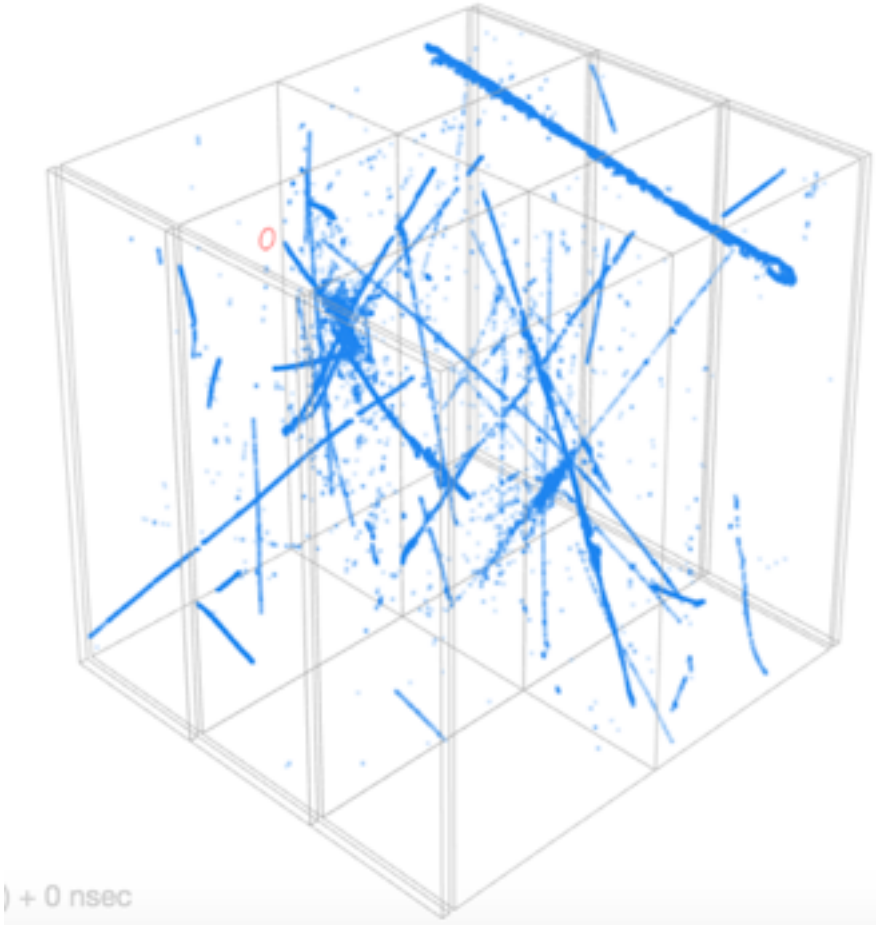

Event sizes

- protoDUNE raw events are each about 180 MB, at 10-25Hz
  - Compare ~2 MB for ATLAS/CMS p-p
  - And ~8 MB for ALICE Pb-Pb
  - During data taking had to reduce time window to cope
- The eventual Far Detector is 100x bigger than protoDUNE
  - We do just read out where the event has happened
  - But we also need to be able to record supernovae, with long time windows and event sizes up to TB
- So we are similar to the astronomy projects that need some access to slots significantly larger than 4GB/processor

Processing the data

- protoDUNE took about 2PB of raw data in 6 weeks of data taking with beam, leading to 1PB reconstructed data
- Aiming for 30PB/year raw data from eventual Far Detector
  - Near Detector data rate is likely to be larger
- FNAL will act as a “Tier-0” source of raw data, with data from FD handled internally and then exported from FNAL
  - But need a distributed approach after that to provide sufficient resources to process and analyse the data
- Experience of the LHC experiments and WLCG tells us that this is already manageable
  - Not trivial though of course

Current computing status

Computing resources

- Storage
  - dCache/pnfs at FNAL, EOS/Castor at CERN, Echo at RAL
  - Recruited several UK Tier2 sites last year (DPM, dCache)
  - More storage being on-boarded: CC-IN2P3, PIC/CIEMAT
- Jobs
  - Grid submission via Glidein WMS
    - So HTCondor-CE, ARC, CREAM are all ok
    - Onboarding process much faster than when we started (eg NIKHEF and CIEMAT last week)
  - Able to use capacity provisioned by HEPCloud
  - Already a small amount of HPC resources being used

protoDUNE worldwide data processing

- Location of grid jobs November 1-24, during protoDUNE data taking.
- A total of ~250,000 reconstruction and simulation jobs were run.
- Up to 17,000 jobs at once
  - ~10 hrs/job (up to 24hrs)
- 60% were external to the dedicated resources at FNAL
- **So protoDUNE already used WLCG sites and implicitly WLCG infrastructure**

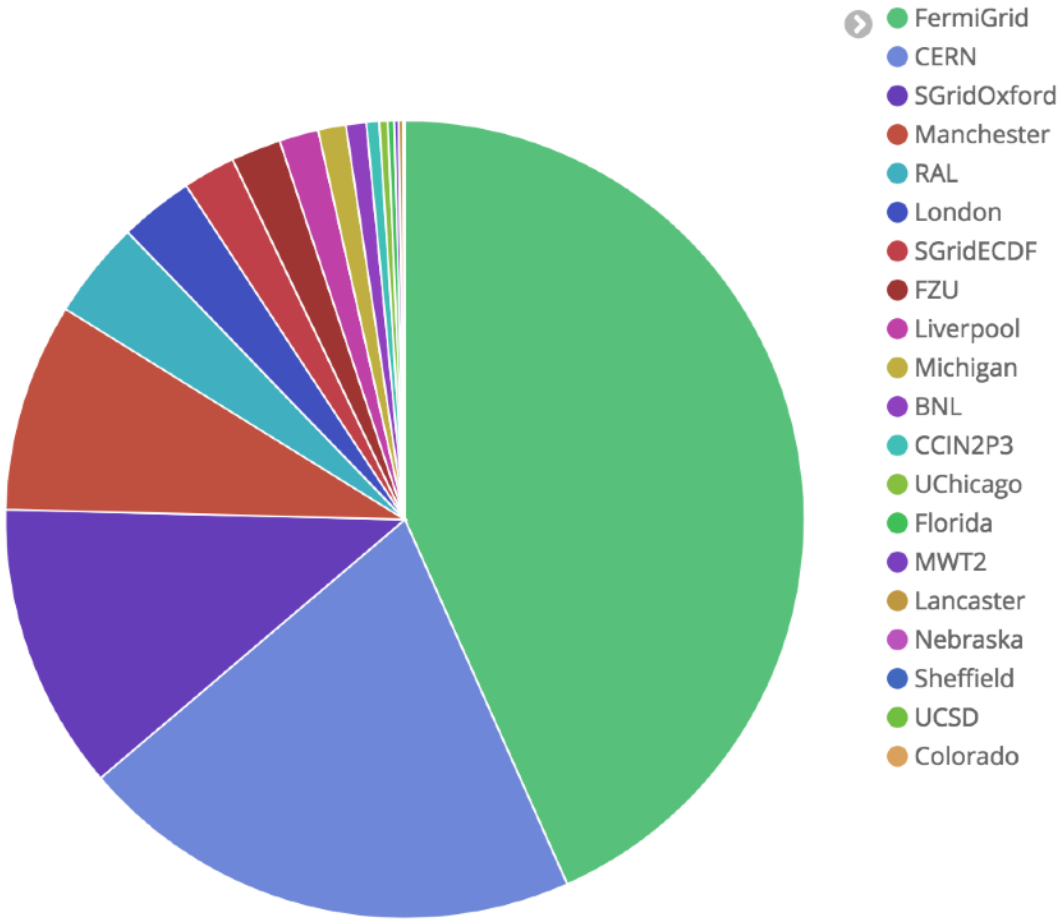

2019 jobs to date, excluding FNAL

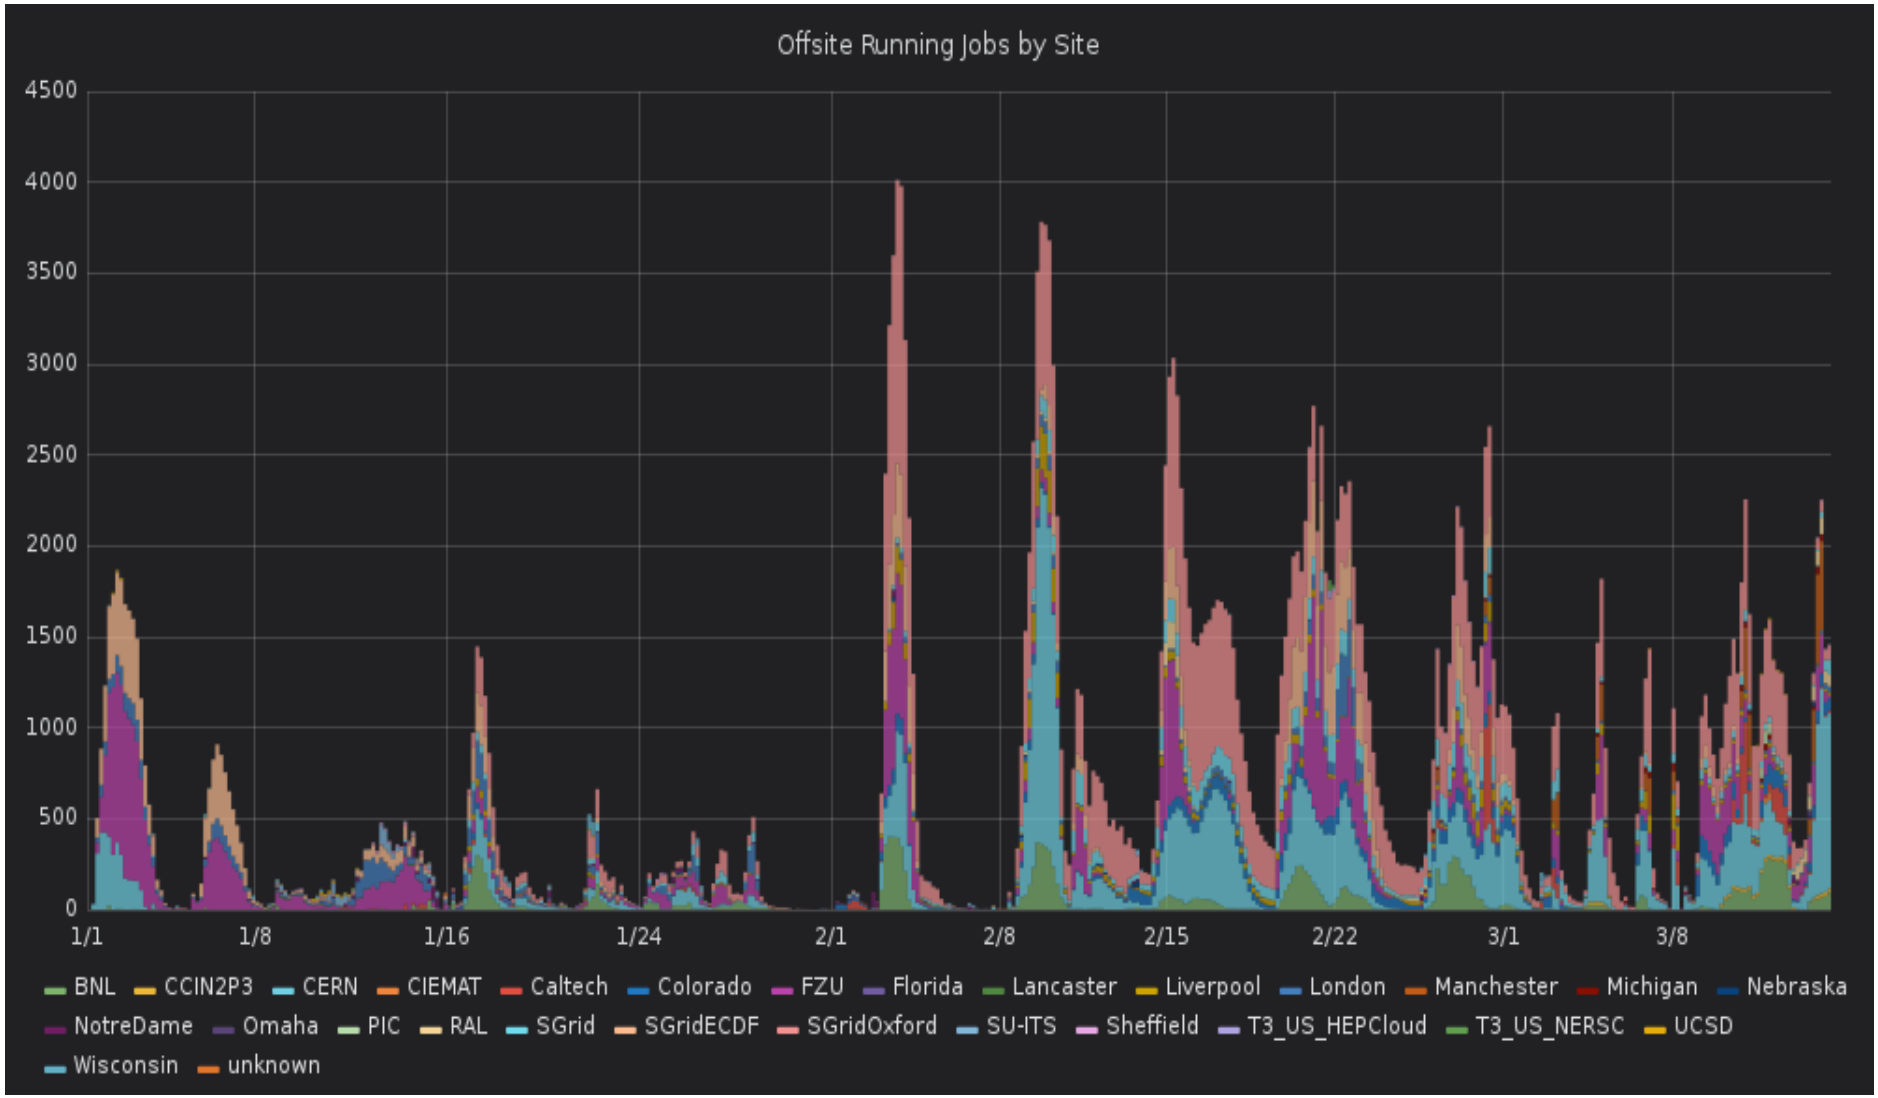

Technologies in use

- Transitioning from “everything at FNAL/CERN” to a grid model
- Inherited long-standing FNAL general platforms (eg SAM and GlideinWMS pool) and FNAL neutrino platforms (eg FIFE)
  - SAM has file metadata and replica locations
  - Most jobs are submitted to Glidein pool and sent out to sites
    - Lots of direct batch submission at FNAL/CERN still by users
- Starting to use RUCIO to place data at sites
  - Developing SAM/RUCIO interface

Technologies going forward

- Already decided to base data management around RUCIO
- Need somewhere to put the file-decscription metadata that's currently in SAM
  - RUCIO may be able to handle this in future, but leaning towards a standalone database
- Need a workload management system that is aware of RUCIO's knowledge of replica locations
  - Starting to evaluate DIRAC, with DIRAC/RUCIO interface in mind - but other solutions are in the mix
- Need a more automated production management system built on top of the WMS

Summary

- DUNE has a series of milestones before the start of data taking with beam in 2026
  - Thanks to protoDUNE SP and DP we have real data to work with as we develop the computing model and systems
- Process of moving to a distributed, grid model is already well underway
  - Many Tier-1/Tier-2 sites already recruited and more welcome
- Decided to use RUCIO for data management
- Evaluation of workload management and higher level production management options has started

Backup slides

ProtoDUNE @CERN (from Heidi Schellman)

Two walls of the cryostat are covered with 3 planes of wires spaced 0.5 cm apart. Total of **15,360** wires

The electrons take ~ 3msec to drift across and you need to detect and time them for the full time

Each wire is read out by 12-bit ADC's every 0.5 microsecond for 3-5 msec. Total of around **6,000** samples/wire/readout.

Around **230 MB/readout** → **80-100 MB compressed**

ProtoDUNE was read out at **10-25 Hz for a 6 week test run**

2.5 GB/sec --> < 1 GB/sec after compression

One issue – this is a **1%** prototype of the real 4-module beast

The big one won't read out as often….