

What DUNE is doing that GridPP sites should know

Andrew McNab
University of Manchester

GridPP52,
Ambleside, 2024




Overview

- DUNE Overview
- Recent/future activities
 - DC24, DDC25, protoDUNE
- DUNE software and computing context
 - RS&DC: Pandora, Rucio/justIN
 - Monday ops meetings
- GPUs in DUNE
 - justIN GPU support
 - Pandora
 - ND

DUNE / LBNF overview

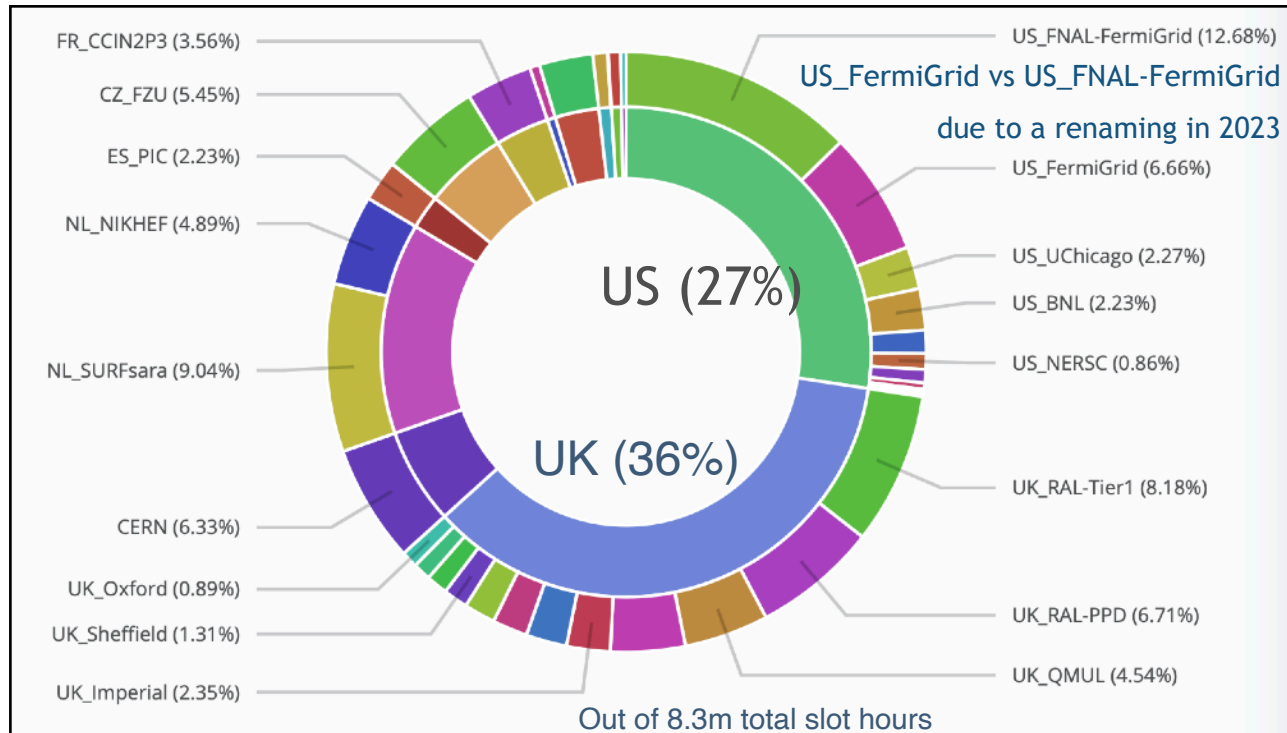
- “Make neutrinos at FNAL then detect some of them in South Dakota (and maybe supernovae and proton decays)”
 - Construction underway. Start taking data later this decade
 - protoDUNE experiments running at CERN **now** and taking data to mid September
- DUNE is part of the **Long** Baseline Neutrino Facility (LBNF)



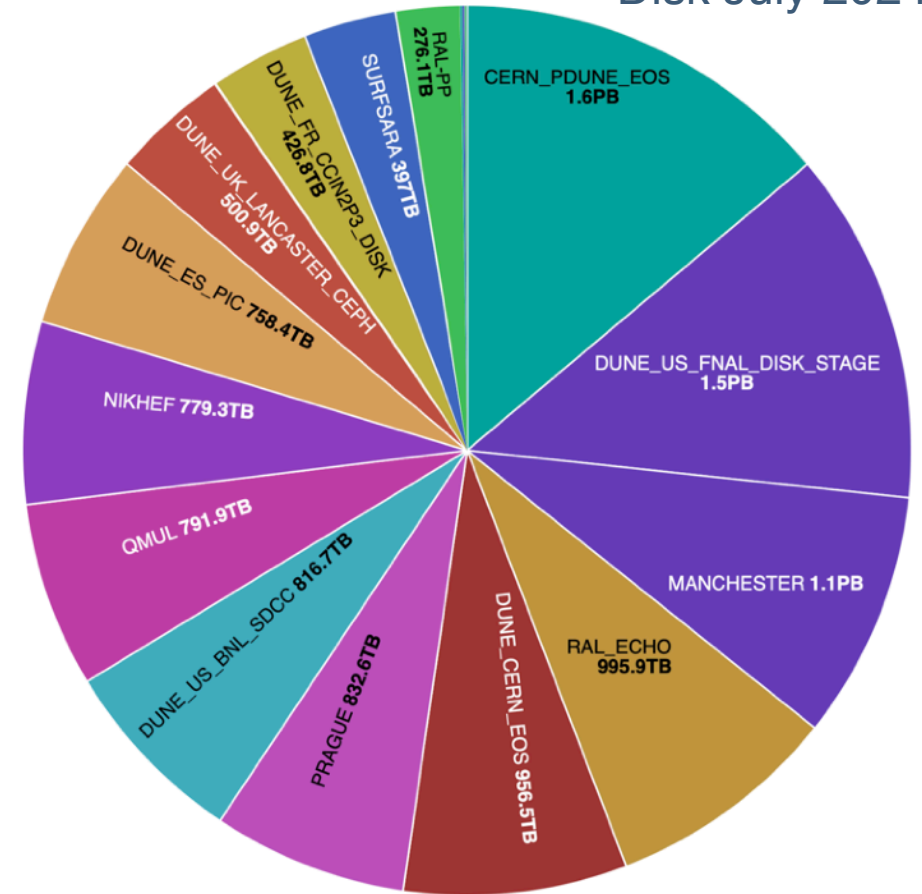
The UK is the main non-DOE contributor

This is very visible and very much appreciated
We make sure to credit GridPP and IRIS

CPU used by DUNE production, 2023

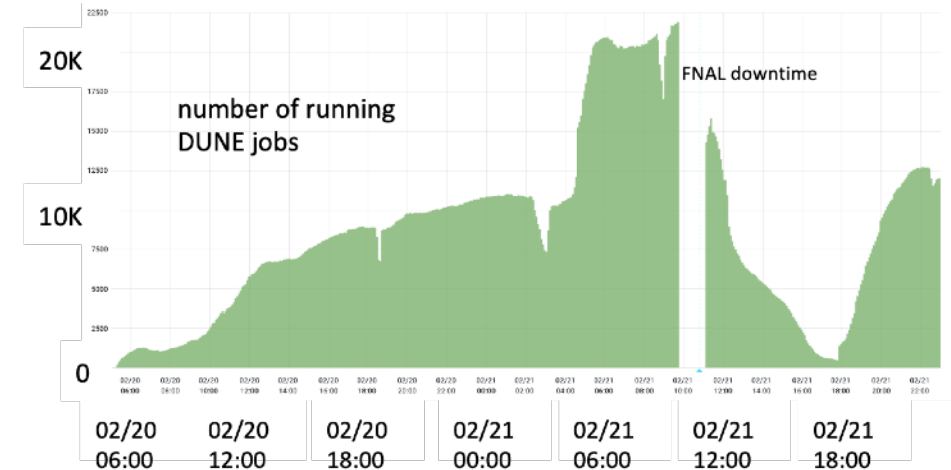



Disk July 2024




DUNE in WLCG DC24

- DUNE has run several private data challenges and participated in DC24 this spring
- Aim was to test BNL (as SURF) to FNAL tape; keep up processing; and processing supernova candidates
 - All at 25% of final far detector goals, which will be comparable to LHCb *now*, so very achievable
 - In practice, only the keep up processing was pushed to the target scale, including some last minute reconfigs of our services
- This was an important preparation for protoDUNE data taking this summer
 - We did another validation closer to the start of data taking, as the cryostat filled
- We plan to do DUNE DCs in the spring each year when there isn't a WLCG DC



protoDUNE at CERN in 2024

- Horizontal Drift cryostat refilled in April
 - Cosmics then SPS beam from June to Sept 18th
 - (Liquid Ar to be transferred to Vertical next month, and beam for it next year)
- Data transferred from DAQ into Rucio-managed storage by DUNE ingestion daemons at CERN
 - Data primarily sent to Fermilab
 - Keep up processing also done at other sites
 - All managed with new justIN workflow system
 - Outputs to “nearby” storage then transferred to Fermilab with Rucio rules



DUNE UK RS&DC project


- “Reconstruction Software and Distributed Computing” - funded 1st April 2024 to 30 Sept 2028
- Continues WP1 of the DUNE UK construction project
 - That’s the **Reconstruction Software** side, including the Pandora framework
 - and the **Distributed Computing** part that concerns us
- Distributed Computing is divided into
 - Data Management effort is based at Edinburgh
 - Significant DM effort at Fermilab/Brookhaven too - mostly Rucio
 - Workflow Management effort is at Manchester and RAL-PPD
 - Workflow uses systems and expertise at Fermilab, but we provide ~all the dev effort - new justIN workflow system plus HTCondor

Workflow Management



- Centered on the justIN workflow management system developed by Manchester and RAL-PPD
- This replaces some of the book-keeping for workflows previously provided by Fermilab product SAM from the Tevatron era
 - and ties together MetaCat (replacing SAM's metadata system to find filenames)
 - and Rucio (replacing SAM's replica management system to find copies of files)
- Underneath the system still uses the global DUNE HTCondor Pool and GlideInWMS
 - But instead of using jobsub + Fermilab schedds, we have dedicated HTCondor schedds for justIN on the RAL OpenStack Cloud
 - Adds just-in-time matching of jobs to HTCondor slots based on nearby unprocessed files for that workflow

- justIN was first used by DUNE globally for the data challenge at the end of 2022
- justIN was used for DUNE's part of the spring WLCG Data Challenge 2024
- After that justIN became the basis for all DUNE central production campaigns
 - More recently, justIN is being used for keep-up processing of protoDUNE data

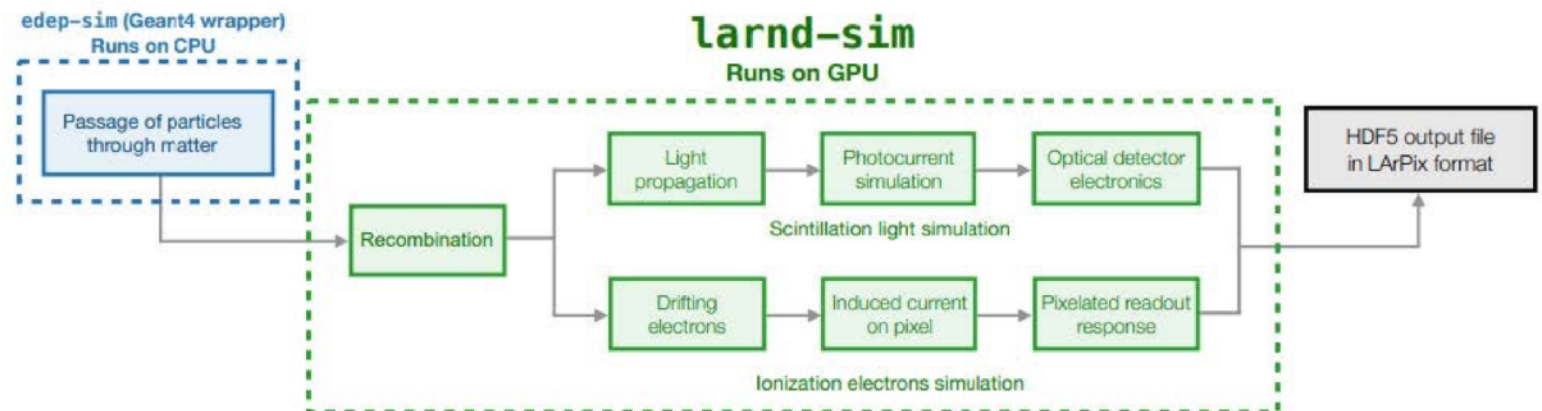


DUNE computing operations

- We have weekly operations meetings on Mondays - 2:30pm UK time
 - <https://indico.fnal.gov/category/1207/>
 - Sites are very welcome to attend! Problems, successes, updates, curiosity :)
- There's also a Fermilab mailing list
 - dune-computing-sites on the Fermilab listserv - ask me if you want to be added
- We try to use GGUS tickets
 - But sometimes contact starts as email and continues like that
- For adding new storages, Rucio is under our direct control
- For adding new compute endpoints (CE/queue), we need to contact the OSG pilot factory admins - we do this on your behalf

DUNE and GPUs

- Pandora developers are making heavy use of GPUs for training machine learning modules
 - Need interactive access and want to move to grid too
- The Near Detector software is highly dependent on GPUs
 - GPUs used both in NDLAr simulation and ML based reconstruction
 - Currently run on Perlmutter at NERSC - lots of GPUs
- We don't yet have an agreed figure for how many GPUs we are going to try to get globally ...



GPU support in justIN

- justIN's wrapper jobs gather and record info about job slots
 - Including GPU presence, memory, model, driver version etc
 - Used to direct jobs for a workflow that needs GPUs to suitable queues at sites
 - Currently this is GPU yes/no, but will allow memory etc requirements in future
- Typically, a site gives the pilot a number of GPUs based on the queue/classads whatever
 - GlideInWMS/HTCondor inside the pilot then matches GPUs based on ClassAds that we generate for the justIN jobs
- Work underway to add NERSC Perlmutter to DUNE HTCondor pool (via HEPCloud), and then have justIN workflows running at NERSC
 - Then transfer ND workflows to justIN, and so allow them to use non-NERSC GPUs
- For Pandora developers needing GPUs, we plan to provide a pathway from interactive development, to running justIN jobs in its interactive test mode, to running them on the grid

Summary

- Funding via IRIS to the GridPP sites is making a major and very visible contribution to DUNE computing resources
- STFC funded RS&DC project has started and is key to DUNE reco, data management and workflow management
 - Sites interact with old friends (Rucio, glideInWMS/HTCondor) and new (justIN) being developed by the project
- DUNE already has a major requirement for bulk GPU, which is currently provided by NERSC
 - We are integrating GPU support into the workflow system at this early stage and plan to use GPUs at sites that make them available