DUNE Computing

Raja Nandakumar

For the DUNE collaboration

DUNE

- Experiment hosted by US DOE
  - Fermilab : Neutrino source, Near Detector
  - SURF : Far Detector
    - ~1500m underground (minimise cosmics)
- Aim : To study neutrino physics
  - CPV in neutrino flavour mixing
    - Measure leptonic CPV to better than 3σ Over 75% of the range of δ_{CP}
  - Take data from accelerator and supernovae

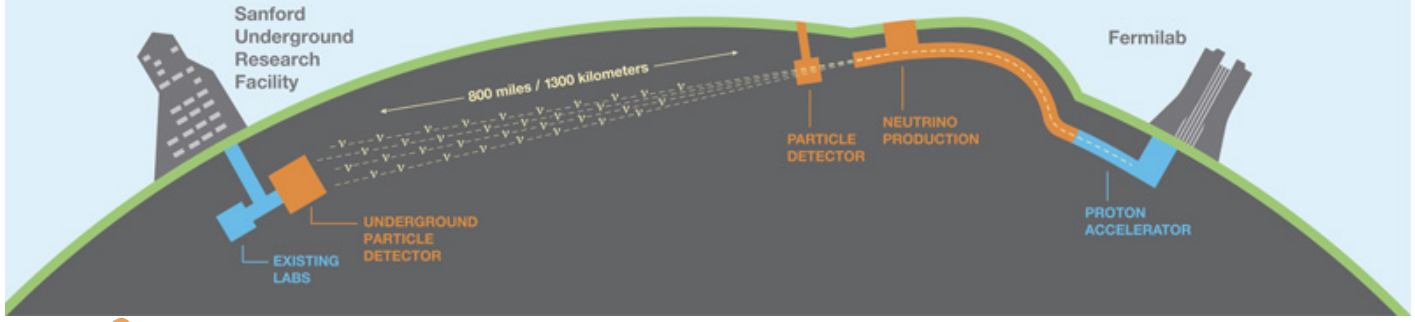

DEEP UNDERGROUND NEUTRINO EXPERIMENT

Numbers

- Detector : 40 KT of liquid Argon in active volume
  - -185°C
- Fermilab :
  - Wide band of protons from LBNF beamline
    - Proton Energy of 60 GeV / 120 GeV
    - ν_μ Energy ~ 0 – 6 GeV
- Supernova (SNB):
  - All neutrino flavours
  - Timescale ~ 10 seconds
  - Energy ~ 10 – 30 MeV

- Data : ~ 30 PB / year (raw)
  - Dominated by Calibration + Beam + Cosmic ray data
  - Each trigger can be several GB
    - Several TB for a supernova candidate
  - Large buffer at SURF for storing possible supernova burst

Computing overview

- Tier-0 : Fermilab
  - Detector site : SURF
  - 100 GB/s connection to FNAL
  - Tape storage at Fermilab, CERN, Tier-1 sites
  - Disk storage at many more sites (e.g. WLCG Tier-2s)
  - Pledged and opportunistic cpu
- Software
  - Simulation and analysis : larsoft, art
    - Replacement for art framework being developed
  - (Re)Use common tools developed for HEP Experiments
    - HEP Software Foundation Collaboration
    - WLCG
    - OSG, ...

Analysis software

- larsoft
  - Liquid Argon Simulation, Reconstruction and Analysis software
  - Originally developed by ArgoNeuT, MicroBoone
    - Use Liquid Argon Time-projection Chamber
  - Now used by multiple experiments
    - SBND, protoDUNE, …
  - Fully supported by DUNE
    - Extend to support multi-threading, …
  - Extended to include CNNs
    - Will need GPU based hardware in future
  - Demonstrated to run on 260+K core super computer
    - 6.5M cosmics reconstructed in 1 hour

Offline computing

- Computing resources primarily set up from Fermilab
  - Tier-0
- Distributed computing paradigm
  - Resources all around the world
    - Pledged, opportunistic, depending on need
    - Ability to analyse remotely
      - Even if there is no local storage
    - Learn from LHC experiments
      - Mainly CMS, Atlas
  - Flexible
    - Wide array of technologies deliver resources
    - Robust
      - Fault tolerant – resources may go off / on depending on circumstances
      - Add new resources as they come available

- dCache, CEPH, EOS, CTA, …
- ipv4, ipv6
- HPC, WLCG, OSG, …

Offline computing (pre-2020)

- Tools originally developed to support analysis at Fermilab
  - Production jobs
    - Poms for production submission
  - User jobs
    - Jobsub at Fermilab
  - Monitoring
    - Fife-mon
  - Experiment data management
    - Sam (Serial Access to Metadata)

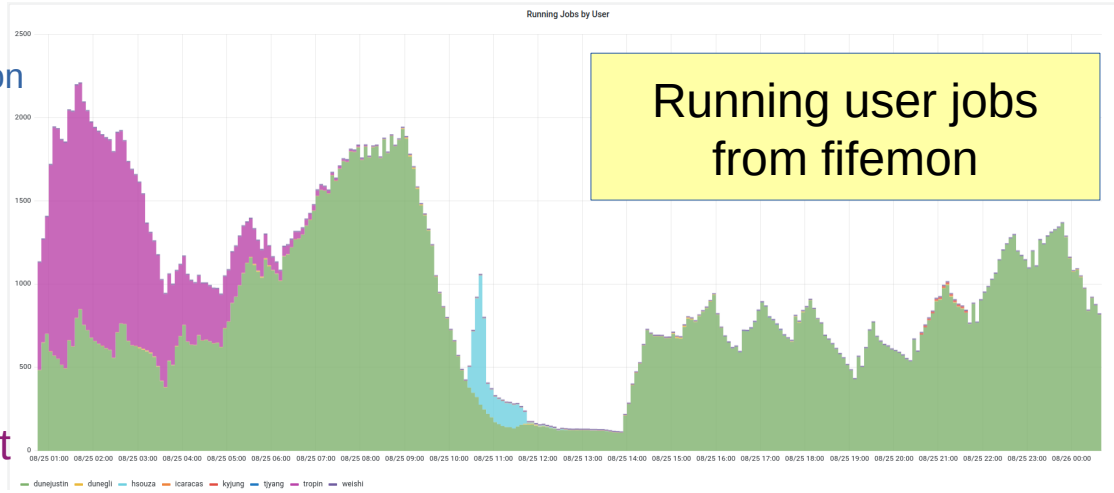

- Other tools
  - Grid job submission : glidein WMS (from CMS)
  - Fermilab storage technology : dCache
  - Software distribution : CVMFS

Computing considerations

- Support distributed resources
  - Slide 6 for example
- High rate of input data
  - Similar to LHC Run-2
- Use well tested technologies
  - Compare technologies when needed
- Best practices in computing techniques
  - Make robust systems, minimise failures

- Job submission : glidein WMS
  - Keep, as DUNE users have extensive experience with it
  - Also has ongoing support from CMS
  - Operationally handled by UCSD computing team

Data management

- Originally samweb
  - Not easily adaptable for DUNE
  - Available technologies supporting expected data sizes and rates
    - Rucio (ATLAS,...)
    - DFC (LHCb,... : Dirac File Catalog)
- Use Rucio
  - With metacat as metadata catalog
  - Also supports the concept of data sets
  - Comes as a stand-alone package
  - “Easily” integrated with other systems

Workflow Management

- Jobs submitted as “pilots”
  - glideins / placeholder jobs which pick up actual jobs based on available resources
- Requirement : (Very) Late binding of job to pilot
  - Allow dynamic prioritisation of work
  - Minimal wastage of cpu resources
  - Well tested principle in LHC experiments
- New development from DUNE
  - justIN – Just-IN-Time workflow scheduler

justIN

- Experiment agnostic
- Developed in the UK
  - Manchester and RAL
  - Dockerised – fast set up and (almost) anywhere
    - e.g. as a gateway to HPCs
  - Extensive documentation and monitoring available
- Pilot jobs talk to justIN servers
  - Pick up the best job for the pilot
    - Based on location, resources, …
- DUNE has transitioned to using justIN for production management
  - Last jobs with poms finished in February 2024
  - Now also accessed by users

Example workflow

Description	PDHD Keepup (2024-08-26T0000 -- 2024-08-26T1200)
State	running
Message	Moved to running state by justIN
MQL	files where created_timestamp >= 2024-08-26T0000 and created_timestamp < 2024-08-26T1200 and core.run_type=hd-protodune and core.file_type=detector and core.data_tier=raw and core.data_stream in (physics, cosmics)
Scope	hd-protodune-det-reco
HTCondor Group	group_dune.prod_mcsim
Finding	every 1 hour(s) to 2024-08-28 23:59:59
Next find	2024-08-26 07:37:48
Submitted by	
Created	2024-08-26 06:05:21
Submitted	2024-08-26 06:05:21
Started	2024-08-26 06:25:36
Finished	
Events for this workflow	

File states per stage

Stage ID	Files	Finding	Unallocated	Allocated	Outputting	Processed	Not found	Failed
1	552	0	271	280	0	0	1	0

Job states per stage

Stage	Total	Submitted	Started	Processing	Outputting	Finished	Notused	Aborted	Stalled	Jobscript error	Outputting failed	None processed
1	384	81	10	280	0	13	0	0	0	0	0	0

justIN time: 2024-08-26 06:57:17 UTC justIN version: 01.01.06

AWT Monitoring from justIN

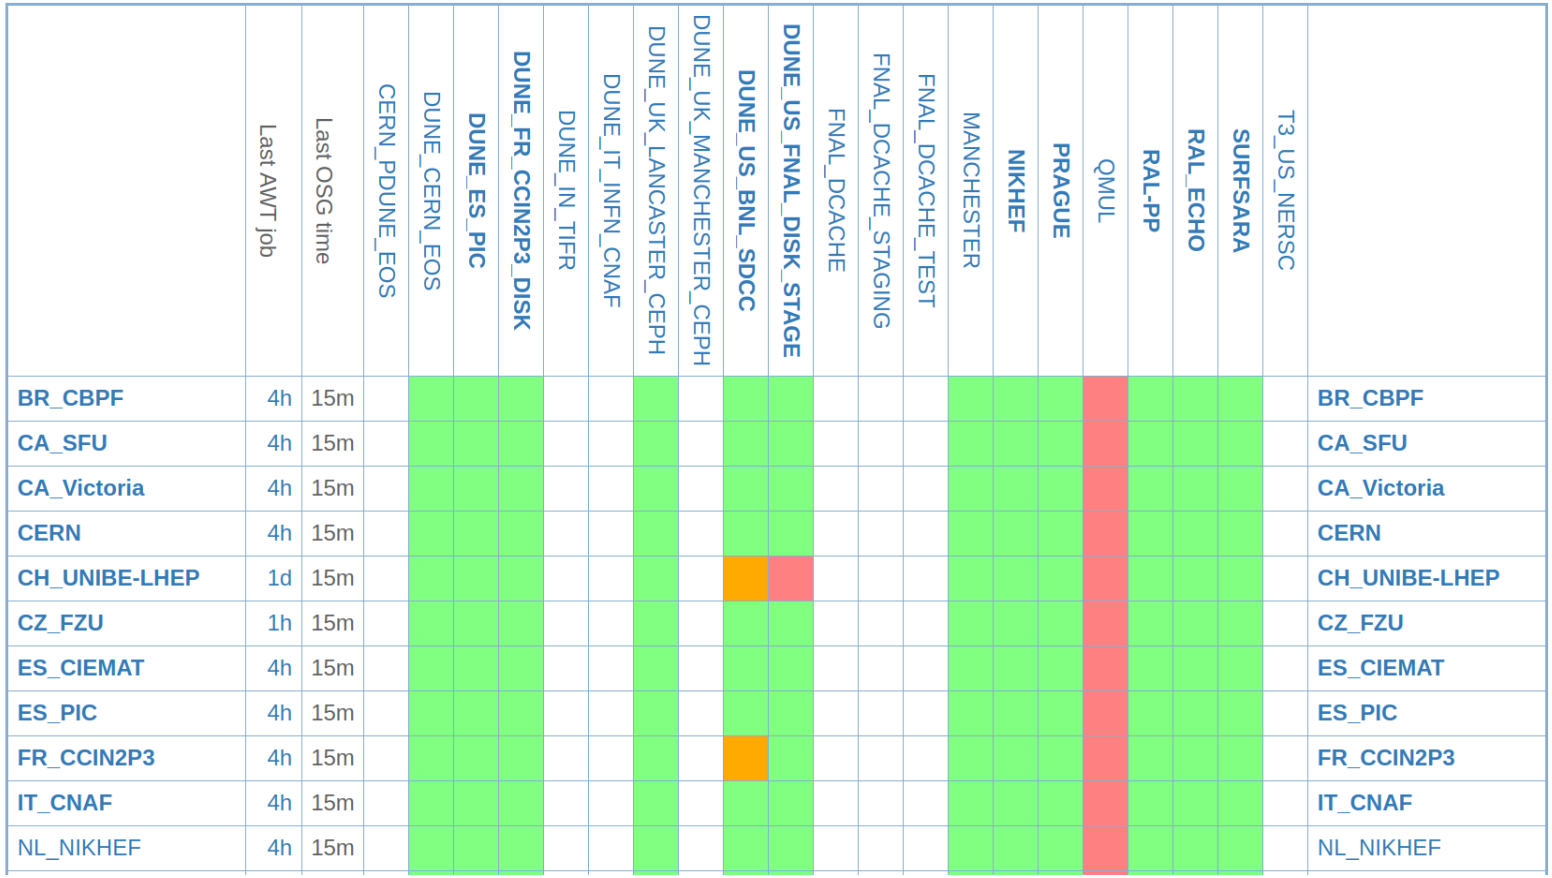

DEEP UNDERGROUND NEUTRINO EXPERIMENT

Summary

- DUNE computing has evolved a lot
  - Almost at optimal configuration (for now)
  - Demonstrated ability to work with the resources pledged
- Many updates being implemented
  - Features requested for justIN, rucio
- Some updates coming up
  - Full ipv6 compatibility
    - Mandated by US DOE
- On the way to getting ready for taking data in 2029+
  - We have already successfully processed protoDUNE data