DUNE Computing Requests for 2023

Computing Contributions Board Meeting

Heidi Schellman and Michael Kirby

Physics → Computing Challenges

Fine segmentation needed for electron-photon discrimination:

Sub-cm-level segmentation over very large volumes drives the number of channels and data volumes.

Low-energy thresholds for astrophysical neutrinos:

The need to optimize the low-energy threshold drives our need to carefully record waveforms with minimal processing and thus drastically increases the raw data volume.

Precise energy calibrations:

The large volumes, complex E field configurations, liquid motion, and potential variations in electron lifetime and drift velocity make it necessary to have large calibration data samples that span the full FD detector volumes. Large cosmic ray and artificial calibration samples will dominate the total data volumes from the FD.

Supernovae:

A supernova neutrino burst candidate will generate 320 TB of (uncompressed) data across the first two modules, resulting in thousands of data files produced over a 100 s period. Supernova physics drives the need for fast data transmission from the FD to computing facilities The drastically different time scale of SNB physics also places requirements on the software framework.

Near Detector integration:

Integration of disparate detector technologies into a coherent whole. New pixel LArTPC geometry.

Analysis and parameter extraction:

Neutrino interaction samples are generally simpler than event records at colliders. However, final parameter extraction using large numbers of nuisance parameters (and or ML) is still a computationally intense problem and will likely require efficient utilization of HPC resources.

DUNE DEEP UNDERGROUND NEUTRINO EXPERIMENT

Compute Challenges

The CDR has lots of details on existing or proposed solutions to our major technical challenges

- **Large memory footprints**
  Data comes in 30-100 MB size chunks in aggregations of up to 320 TB → framework improvements before FD operations. We can handle protoDUNE data but more work is needed for FD.
- **Storing and processing data on heterogeneous international resources**
  Current solution is to use mainly single-threaded code and OSG/WLCG standard configurations where possible.
  There is ongoing work on HPC integration.
  Data challenges in progress to test our new computing models.
- **Machine learning**
  How do we support ML development and deployment across the wide range of platforms available to us?
- **Efficient and sustainable use of resources**
  Design for efficient use of global resources in workflow design but also document, train and do continuous integration to make certain that codes are robust and efficient.
- **Keeping it all going**
  Experiments run for decades; people and operating systems turn over much more quickly. Considerable effort will always be needed to maintain and deploy
  - Security and authentication
  - DUNE specific databases
  - Code integration and management
  - All the underlying services and code we build on

DUNE Computing Resource Model

- less “tiered” than current WLCG model —> flatter model proposed by HSF DOMA working group
- collaborating institutions (or groups of institutions) provide significant disk resources (~1PB chunks)
- plan to use common tools for most services
- participation in the HSF process important to provide and integrate new solutions

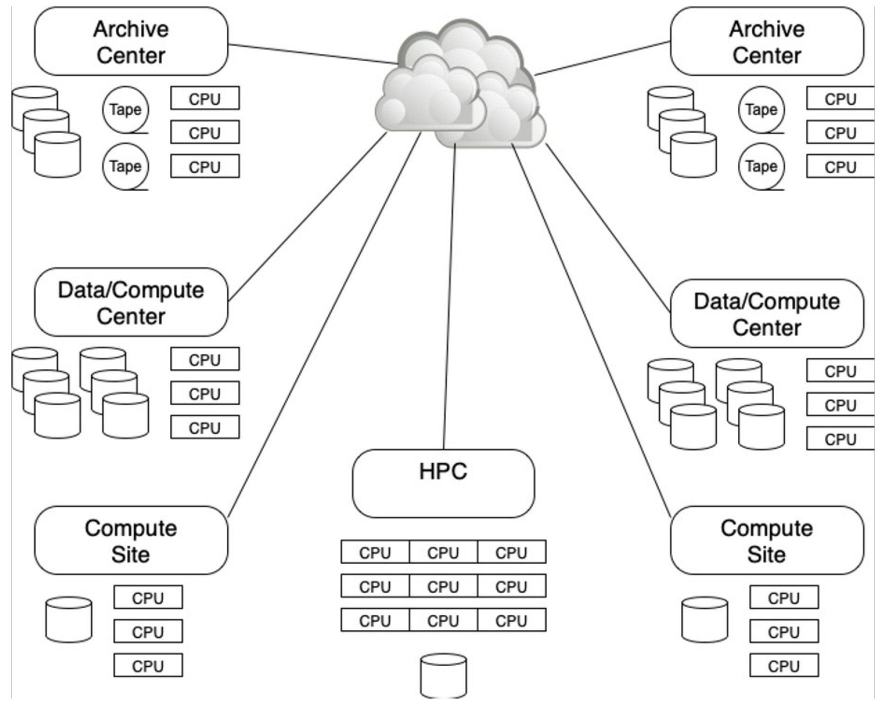

Basics of the resource model

- Keep raw data on disk for 1 year, on tape to end of expt.
  - For protoDUNE – 1 copy each at CERN/FNAL
  - For DUNE 1 copy at FNAL, 1 copy at other institutions
- Reconstruct full sample every year (protoDUNE for 4 years, DUNE to end of expt.)
- Do new simulation campaigns each year
- Keep simulation and reconstructed data on disk for 2 years (always have 2 versions)
  - One copy in Americas, one in Europe where possible (model assumes 1.5 copies)
  - No need to stage from tape until it ages out
- One copy of reconstruction/simulation -> tape as it can be redone if necessary.

- **CPU estimates** are based on measurements from ProtoDUNE data and existing simulations and **for the FD/ND have large uncertainies**.

DEEP UNDERGROUND NEUTRINO EXPERIMENT

Projected disk needs from CDR

These have been revised since the CDR but general estimates still hold.

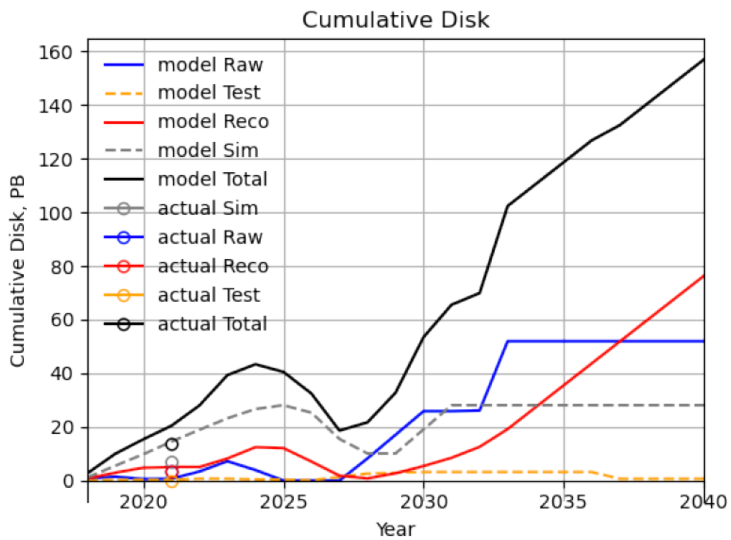

We are ~10% of CMS

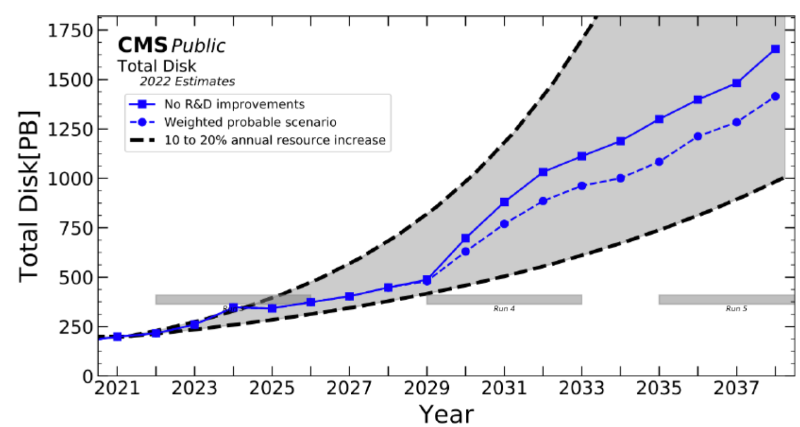

Figure 6.4: Estimated size of various disk samples in PB for DUNE and CMS at the HL-LHC for comparison. This estimate includes retention policies and multiple copies. The points show actual use in 2021 which was lower than planned due to delays in distributing second copies of samples to remote sites.

DEEP UNDERGROUND NEUTRINO EXPERIMENT

	2021	2021	2022	2022	2022	2023	2023
	Pledge (PB)	Disk Actual	Pledge (PB)	Disk Alloc (PB)	Disk Used	Standard Request	Modified Request
BR	0.00						
CA				0.05	0.05		
CH	0.20		0.20				
CZ	0.30		1.00	1.13	0.51		
ES	0.50		0.72	0.72	0.01		
FR	0.50		0.50	0.50	0.13		
IN	0.75		0.75	0.10	0.00		
IT							
NL	1.90		1.90	1.90	0.42		
RU			0.50	0.50	0.50		
UK	4.00		4.00	3.83	3.12		
US BNL	0.50		0.50	1.00	0.50		
US - other							
National	8.65	0.00	10.07	9.73	5.24	15.40	12.94
CERN	2.20		3.00	4.00	2.50	2.60	4.00
FNAL	2.20		7.60	8.86	8.85	7.80	8.86
Total	13.05	0.00	20.67	22.59	16.59	25.80	25.80

Table 4: Summary of disk pledges, allocations and usage for 2021-2022 with model request for 2023. This is based on the 2022 CCB tables which are available in indico [2, 3]. These numbers are derived from the rucio reports in Table 3 and may not be complete.

Updates from 2022-12-07

Country	RSE	Total Allocation(PB)	Used(PB)	Free(PB)	Free(%)
CERN	DUNE_CERN_EOS	1.230	0.720	0.464	41.5%
CERN	CERN_PDUNE_EOS	3.010	1.500	1.300	48.9%
CZ	PRAGUE	1.130	0.517	0.554	54.1%
ES	DUNE_ES_PIC	0.720	0.492	0.207	31.6%
FR	DUNE_FR_CCIN2P3_XROOTD	0.013	0.000	0.012	100.0%
NL	SURFSARA	**0.025**	0.007	0.017	72.7%
NL	NIKHEF	0.991	0.422	0.517	57.4%
UK	RAL_ECHO	1.000	0.862	0.126	13.9%
UK	RAL-PP	0.100	0.100	0.000	0.0%
UK	QMUL	1.100	0.680	0.382	38.2%
UK	MANCHESTER	1.080	1.000	0.042	4.3%
UK	LIVERPOOL	0.001	0.001	0.000	22.9%
UK	LANCASTER	0.550	0.515	0.032	6.3%
US	DUNE_US_BNL_SDCC	0.600	0.506	0.086	15.7%
CERN	Total	4.240	2.220	1.764	41.6%
CZ	Total	1.130	0.517	0.554	49.0%
ES	Total	0.720	0.492	0.207	28.8%
FR	Total	0.013	0.000	0.012	92.3%
NL	Total	1.016	0.429	0.534	52.6%
UK	Total	3.831	3.158	0.582	15.2%
US	BNL	0.600	0.506	0.086	14.3%
Total	**National**	**7.310**	**5.102**	**1.975**	**27.0%**

Projected CPU needs from CDR (not memory weighted)

These have been revised since the CDR but general estimates still hold.

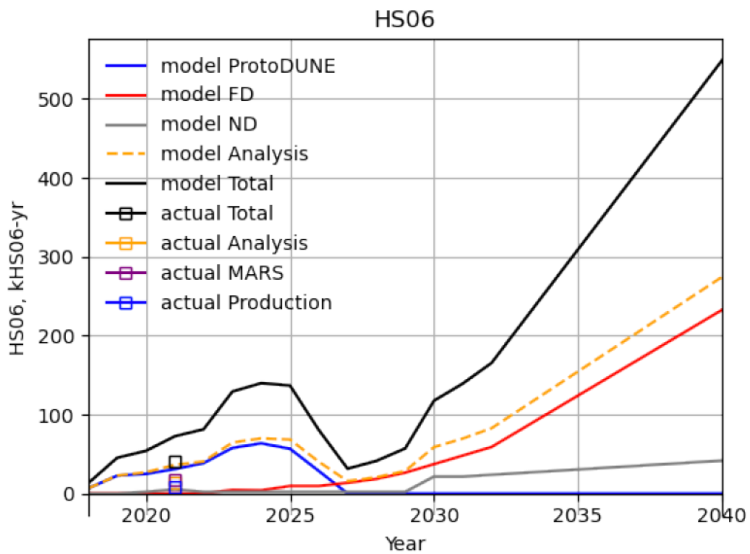

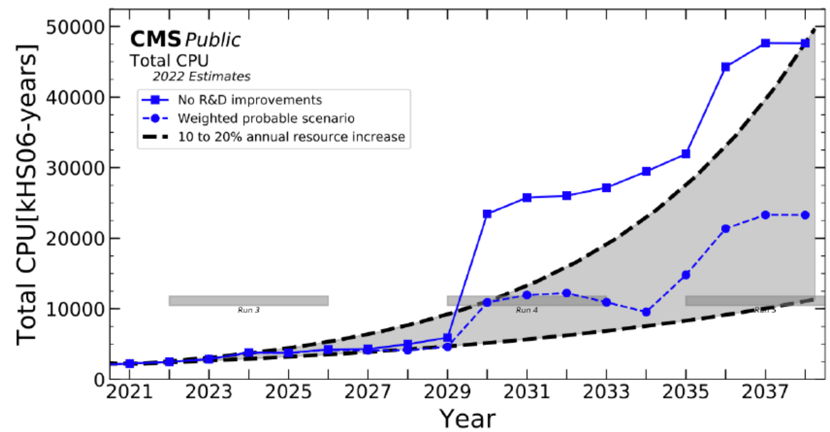

Figure 6.8: Estimated CPU needs for DUNE compared to CMS. The units are kHS06-years (2020 vintage CPU's are ∼11 HS06 [80] per core) assuming 70% efficiency. CPU utilization in 2021 was lower than the model due to the absence of a yearly reconstruction pass.

- 5000 to 15,000 simultaneous processes

DUNE DEEP UNDERGROUND NEUTRINO EXPERIMENT

Important point on CPU usage in DUNE jobs

- DUNE differs from other HEP experiments in **frequently requiring more memory/core** than is available at particular sites.
- For example an 8-core pilot with 16 GB of available memory may only accommodate four reconstruction processes. As a result, we make our requests in terms of **memory-weighted-core wall time (MWC)** with the base memory being 2000 MB.
- This maps reasonably well to the memory weighted slot-time returned by HTCondor and the slot-time reported in EGI statistics. Sites that offer more (or less) than 2000 MB/core can scale their contributions up by the local memory/core.
- The wall-time estimates in the model are created by estimating the number of simulated and raw events taken and then scaling by the measured CPU time on a gpvm corrected to wall-time by the estimated efficiency (default 70%) and for a memory utilization factor that takes into account the differing memory needs for different applications. Here we assume that **analysis takes 3000 MB**, **reconstruction takes 4000MB**, and **simulation takes 6000MB**.
- **Contributions** are then **requested in units of MWC-time** which is wall-time×2000 MB units.

DEEP UNDERGROUND NEUTRINO EXPERIMENT

Example

- An example of a national pledge in MWC might be
  - 1000 cores with 2GB available/core for a year
  - 500 cores with 4 GB available/core for a year
  - 100 cores with 8 GB available/core for a year.
- The MWC pledge would then be

$$1000 \times 2GB/2GB +$$
$$500 \times 4GB/2GB +$$
$$100 \times 8GB/2GB$$
$$= 2400 \text{ MWC-year units}$$

CPU	2021 pledge	2021 use	2022 pledge	2022 MWC used	2023 request (MWC)
BR	100	19	200	16	
CA		6		40	
CH	200	1		0	
CZ	1560	128	2400	78	
ES	500	47	512	122	
FR	310	34	250	48	
IN	450	79	450	126	
IT	-	-		-	
NL	696	210	788	340	
RU		8	1000	39	
UK	1000	947	1000	912	
US OSG	1250	593	1000	354	
National	6066	2072	7600	2075	7585
CERN	3310	306	950	298	3792
US Fermilab	3310	3059	1945	2966	3792
Total	12686	5437	10495	5339	15169
Model		7912		9855	15169

Table 5: Summary of DUNE wall-time pledges and contributions for 2021 and 2022. The 2022 actual numbers are memory-weighted. Individual nations are listed and then merged (with US OSG) into a National section.

DEEP UNDERGROUND NEUTRINO EXPERIMENT

EGI accounting

Country	Site	Total Slot-Years	Percent
BR	CBPF	**13**	**0.82%**
CA	CA-SFU-T2	**0**	**0.00%**
CH	UNIBE-LHEP	**1**	**0.06%**
CZ	praguelcg2	**120**	**7.52%**
ES	Total	**198**	**12.42%**
FR	IN2P3-CC	**60**	**3.76%**
NL	Total	**184**	**11.53%**
UK	Total	**1388**	**87.03%**
TOTAL	Total	1,595	

EGI accounting for Nov21-Oct22

Differs from the FNAL memory-weighted #'s

NL, CA EGI < FNAL
CH, UK, ES, CZ, FR EGI > FNAL

IN/US not included

https://accounting.egi.eu/wlcg/countries/elap_processors-year/COUNTRY/Year/2021/11/2022/10/custom-dune/localinfrajobs/

Revised estimates for CPU with memory weighting

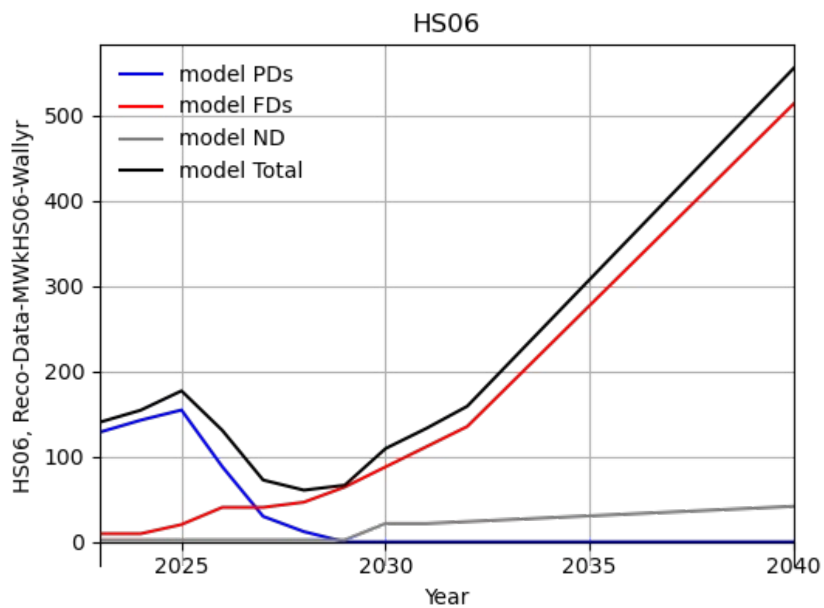

Here the main change is the memory-weighting relative to 2000 MB reference machines.

DEEP UNDERGROUND NEUTRINO EXPERIMENT

Summary of request for 2023

		Disk (PB)	Modified Disk (PB)	Tape(PB)	CPU (MWC-years)
Model		25.80	25.80	45.5	15,169
Request					
	FNAL	7.80	8.86	36.2	3,792
	CERN	2.60	4.00	9.2	3,792
	National	15.40	12.94	0.1	7,585
	Total	25.80	25.80	45.5	15,169

Table 1: Proposed pledges for 2023. Disk pledges are based on existing CERN and FNAL contributions with National contributions making up the rest of the model request. Tape pledges reflect the dominant use of CERN and FNAL for archival storage of data. CPU pledges are in units of memory-weighted-core-years and assume Fermilab and CERN each pledge 25%.

- Disk request includes existing FNAL and CERN contributions
- Tape request reduced to 100 TB from National sites for testing, will increase in later years.
- CPU request is now memory-weighted, assumed data taking in 2023.