

DUNE Timeline and needs

Paul Laycock, with thanks to Mike Kirby

- DUNE aims to release a Computing CDR this year, understanding what DUNE's computing needs are is a work in progress. We still have a little time, DUNE will start in 2027 +
- The ProtoDUNE (I and II) experiments at CERN can be used to test large LArTPC detectors although they're not at the same scale as full DUNE - ProtoDUNE II starts next year
- Full DUNE will record up to 30 PB/year of raw data, first stage processing sees a data reduction of ~100 and the CPU guesstimate is at the 10k core scale
- Challenges (my view) are mainly related to hybrid computing (some processing steps look very GPU-friendly), the 6GB nominal event size and not to mention the 100TB supernova events - challenging for a software framework, data management and WMS/WFM
- FNAL is host lab and has projects like HEPCloud, plus significant investment in in-house Rucio expertise - I asked Mike Kirby (Computing Consortium Technical lead) for his take (next slides)

The Future of DUNE Computing

From the brain of Kirby

- Computing infrastructure
 - emphasis on a flatter computing model to rely upon the much greater distributed infrastructure around the world - we are not building a new open computing grid as done by WLCG/OSG
 - delivery of CPU/disk/tape will naturally be associated with each other but not strongly coupled to each other - (very) late binding of jobs is part of the model
 - late binding allows for smaller reliance on HA resources and associated cost in effort
- Workload management
 - the low-level tools for queuing and resource provisioning are the purview computer scientists. In the next 5 years, DUNE hasn't yet found a need or expertise to transition away from current solutions (GlideInWMS right now)
 - custom configurations and feature requests will be part of our future
 - bursting to HPC for raw data processing may become more important but is not unique to DUNE
- Workflow management
 - DUNE will have some novel tasks and workflows
 - development of those tools will be important to our successful and efficient use of resources
 - we have expertise and effort to accomplish these tasks within house and to focus on DUNE's requirements
- Data Management
 - because of unique data volume ratios (raw data:production:ntuple:etc) and flatter computing model (CPU location vs data location) and reliance on networking, data management will be very important
 - flatter computing model also lends itself to data lakes and new SE models
 - strong effort in contributing to Rucio (Quality of Service), Analysis Metadata Catalog, Data Delivery, and monitoring has been one focus of DUNE Computing
 - management of data replication policies, monitoring data popularity, and providing centralized management to raw data (6+ GB/evt) will be important
- Software Management
 - Modernization of software management will be focused on github/spack/containers
 - Containers are still dependent upon operating systems and support therein - follow lead of FNAL and CERN
 - portability of code will be important, but planning on using portability tools to accomplish this
- Databases
 - I defer to the esteemed Database Lead, Paul Laycock, Ph.D.
- Analysis
 - here be dragons - there is a very strong risk that we will be in a troubling situation of having many, many analysis formats (FD osc, FD non-osc, ND-3DST, ND-GAr, ND-Prism, a second FD osc, Supernova)
 - effort is needed to make sure this doesn't happen

Research and Development

From the brain of Kirby

- AI/ML - there's so much potential here at every level - trigger, raw data, analysis, data management that it's hard to focus on one place with a specific goal. A well-defined problem likely has a straightforward non-ML solution that is 15% less performant with a known path. What do we pursue?
- Analysis tools/data reduction research - columnar analysis, large-scale data reduction research using data lakes with access to large $O(100PB)$ datasets could have huge impact
- HPC already shown to have significant impact for fitting and marginalization of nuisance parameters, but how do we get that back to the individual analysis level instead of just the collaboration global fit
- ...

Summary

Back to Paul...

- Based on his slides, and the fact that FNAL is making significant investment in their own tools (e.g. HEPCloud) I asked Kirby if this was a good summary:
 - WMS/WFM are covered at FNAL, a good goal could be to show complementarity and share information on e.g. how best to exploit HPCs
- This was met with a thumbs up from Kirby!