



DUNE Software & Computing Challenges

Paul Laycock for the DUNE collaboration

BROOKHAVEN
NATIONAL LABORATORY



Neutrino physics




Figures taken from: <http://www.hyper-k.org/en/neutrino.html>

- **The flavour states** of neutrinos are mixtures of **the mass states**
- As neutrinos happily travel along at less than the speed of light, the relative contributions of **the mass states** changes (they have different wavelengths), and thus **the flavour states** that we see via their interactions with matter also change

DUNE concept

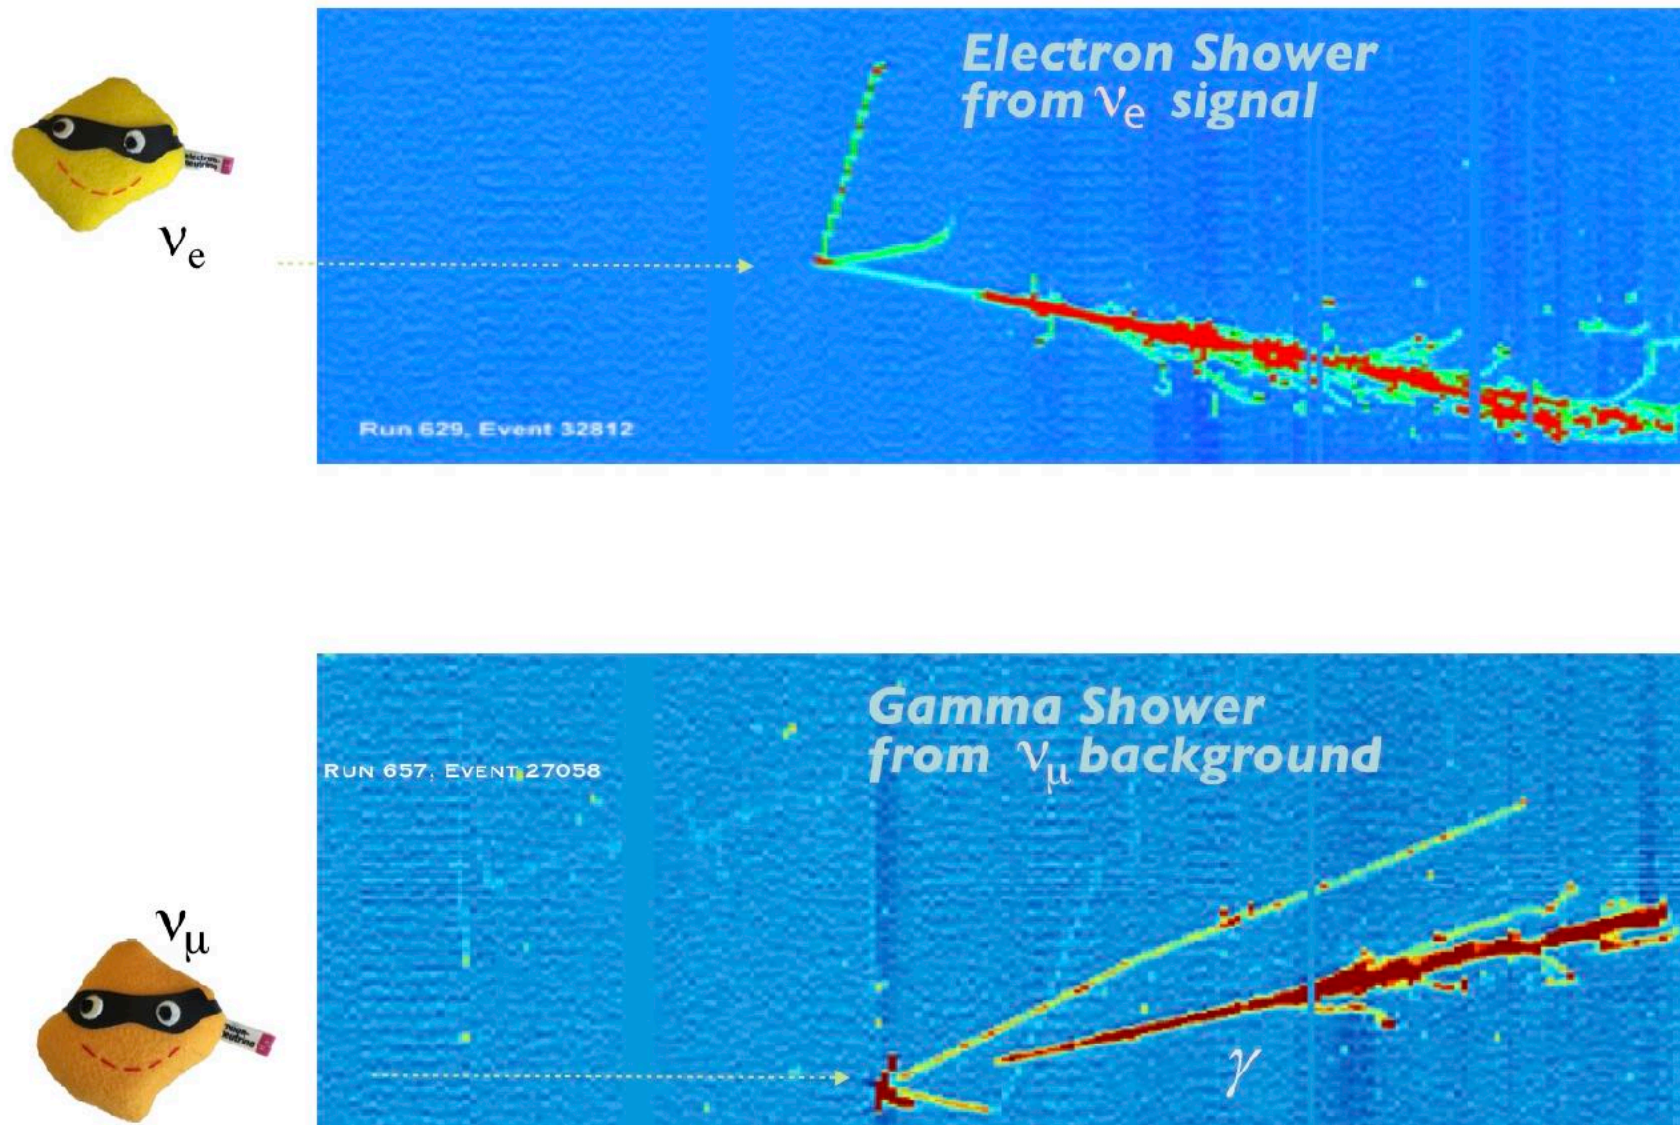
<https://www.dunescience.org/>



- **DUNE** will produce a beam of largely **muon neutrinos** at **Fermilab**
- 1300 kilometres away, some of the **muon neutrinos** should become **electron neutrinos**
- Comparing the conversion rates of *neutrinos with anti-neutrinos* could show evidence of *CP violation*
 - **DUNE** will measure this at the few % level, unprecedented for a neutrino experiment
- The enormous **DUNE** Far Detector is also sensitive to *proton decay* and *supernovae*

DUNE physics

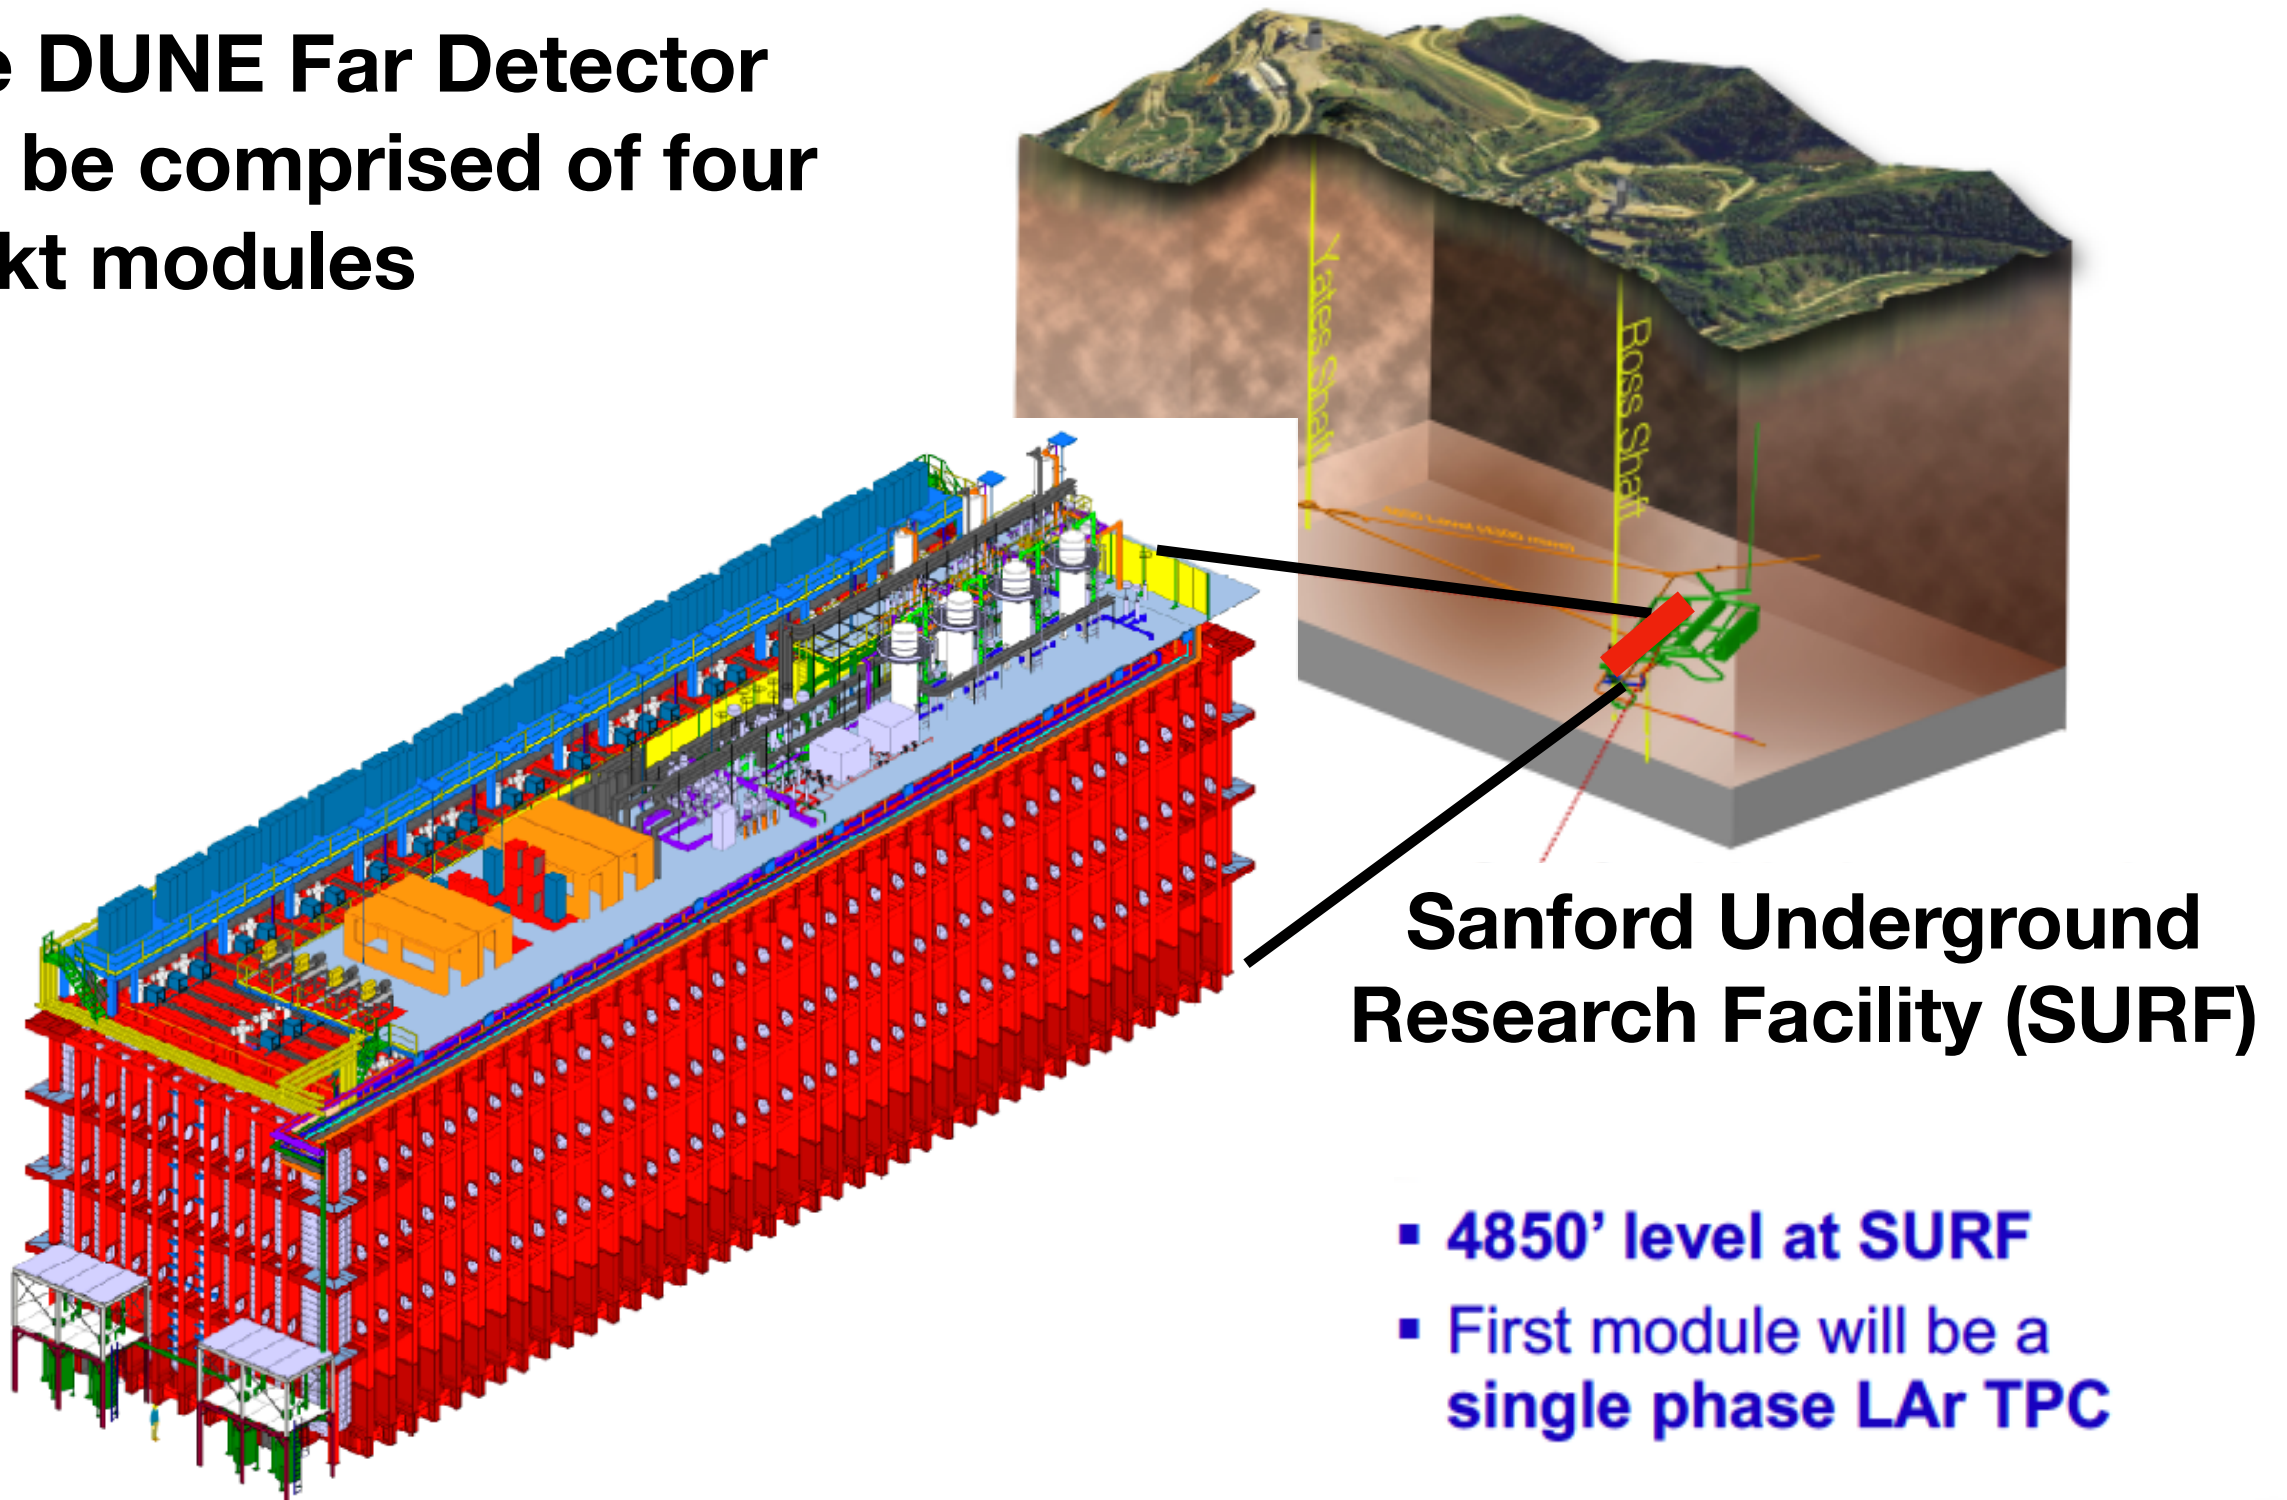
Image credit: ArgoNeuT



- The experiment relies on distinguishing **electron neutrino** appearance events (*top*) from the **muon neutrino** interactions
- LAr TPCs are the detector technology of choice, see Haiwang's talk from earlier today

DUNE far detector




The DUNE Far Detector will be comprised of four 17-kt modules



DUNE far detector

1 APA


150 APAs = 1 module



- A single module will comprise **150 Anode Plane Assemblies (APAs)**
- Each **APA** has 3 planes of wires and a total of 2560 wires
- Every wire is readout with 12-bit ADCs every $0.5 \mu\text{s}$ for $\sim 6\text{ms}$
- Around **40 MB** uncompressed readout per **APA**

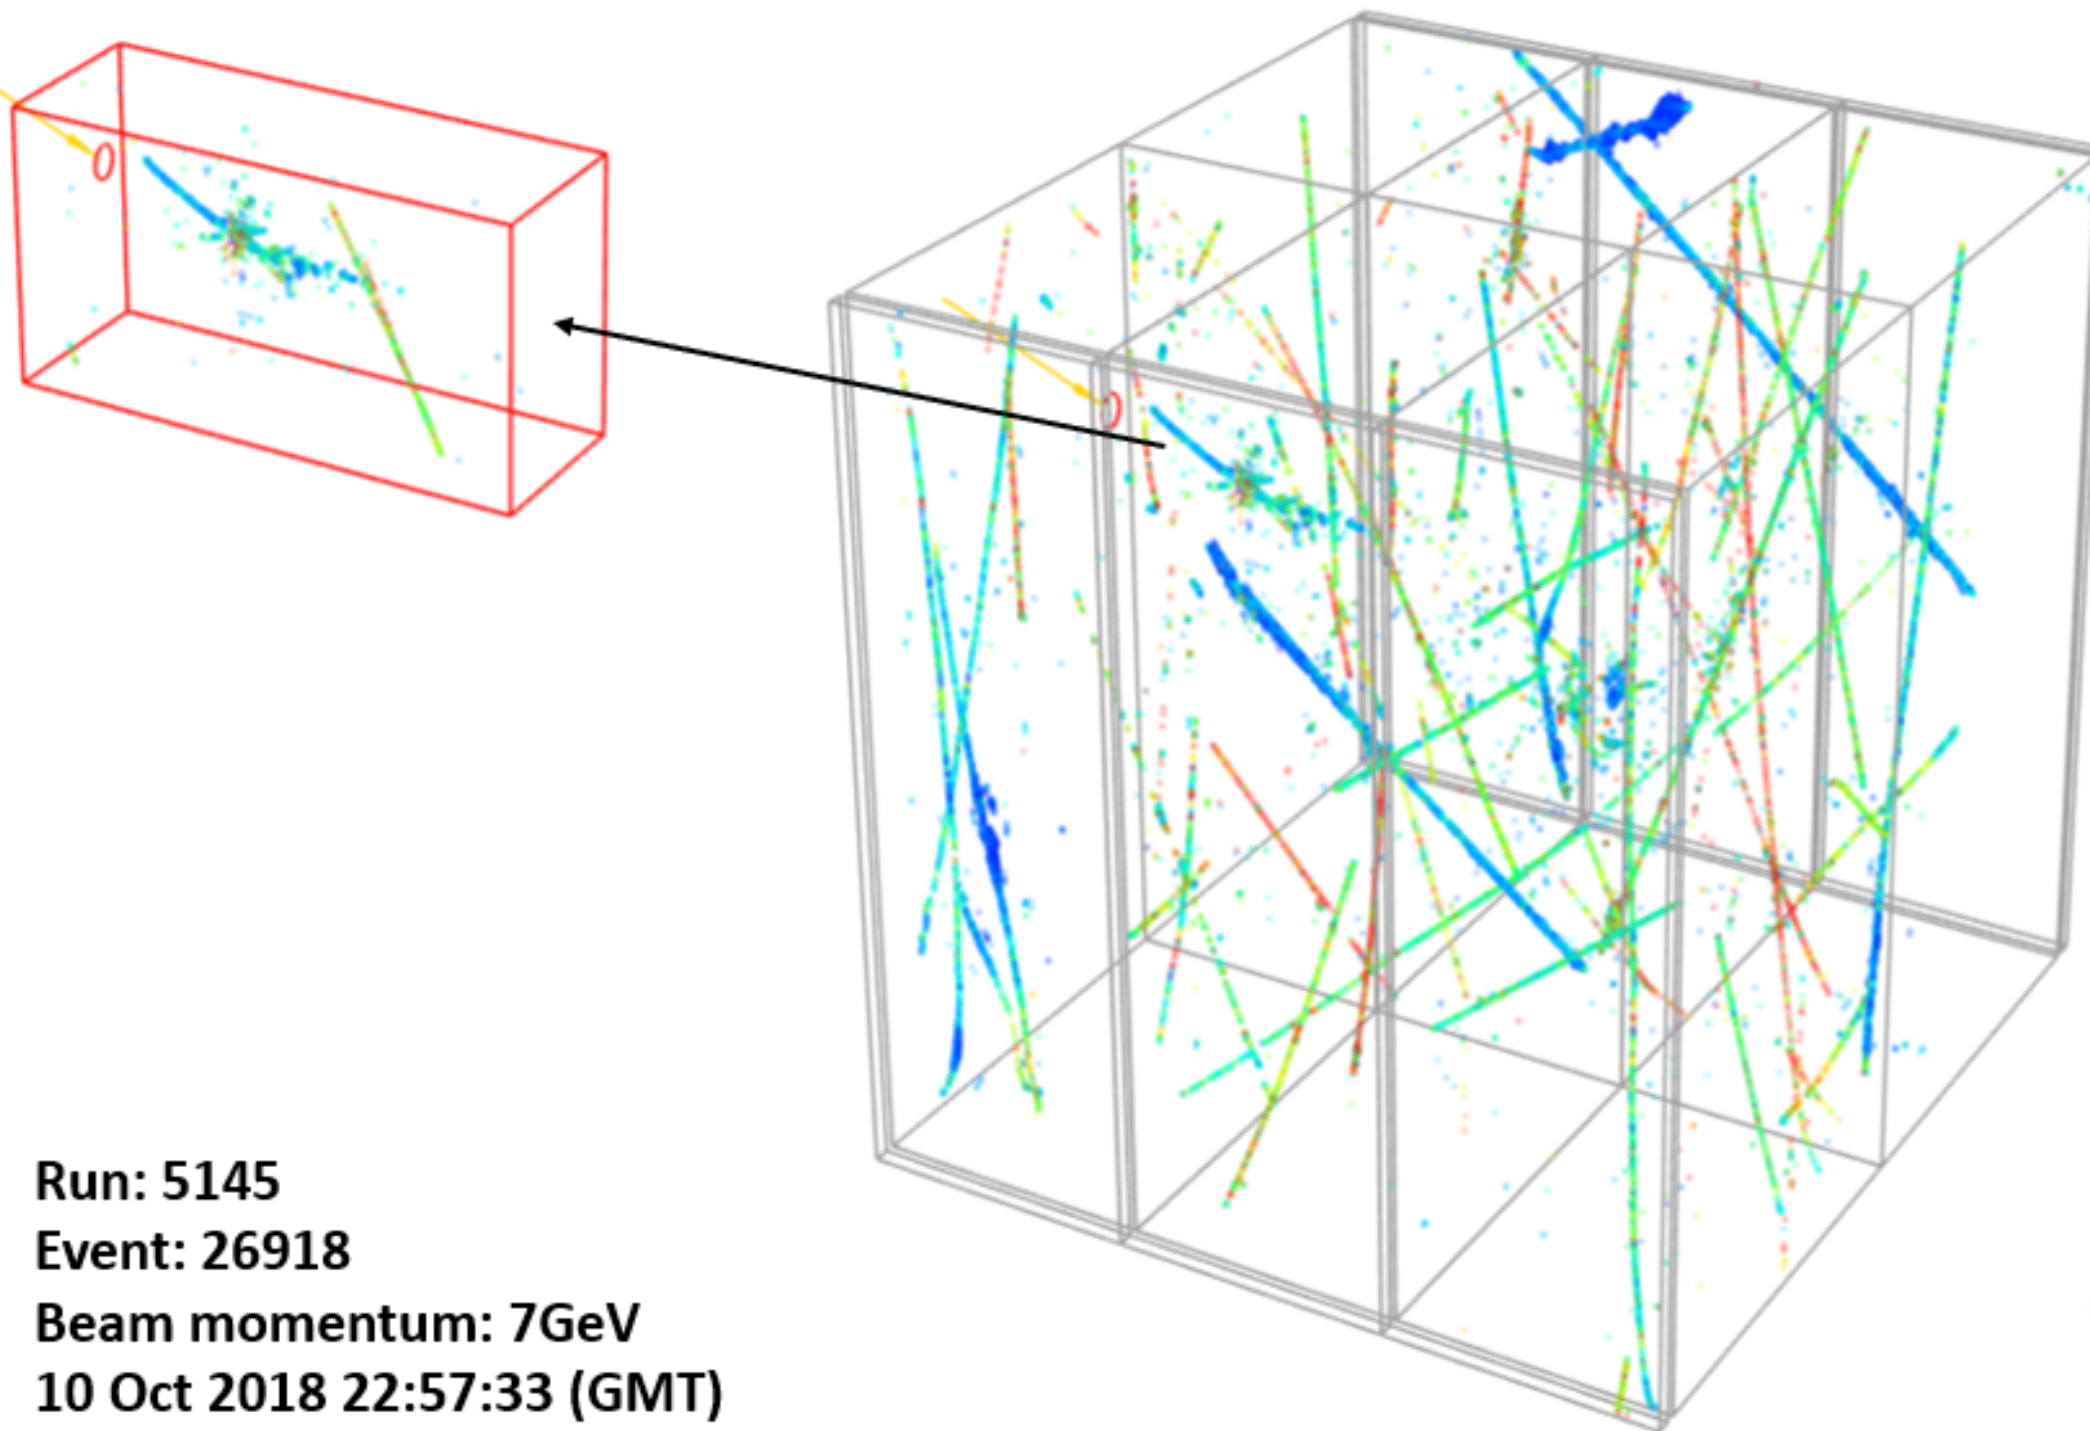
A LAr TPC in action

Animation by Bo Yu (BNL)



- Ionisation from a neutrino interaction being drifted to three anode plane wires
- Every wire is readout with 12-bit ADCs every $0.5 \mu\text{s}$ for $\sim 6\text{ms}$


ProtoDUNE data



Run: 5145
Event: 26918
Beam momentum: 7GeV
10 Oct 2018 22:57:33 (GMT)



DUNE data



DUNE data and “events”

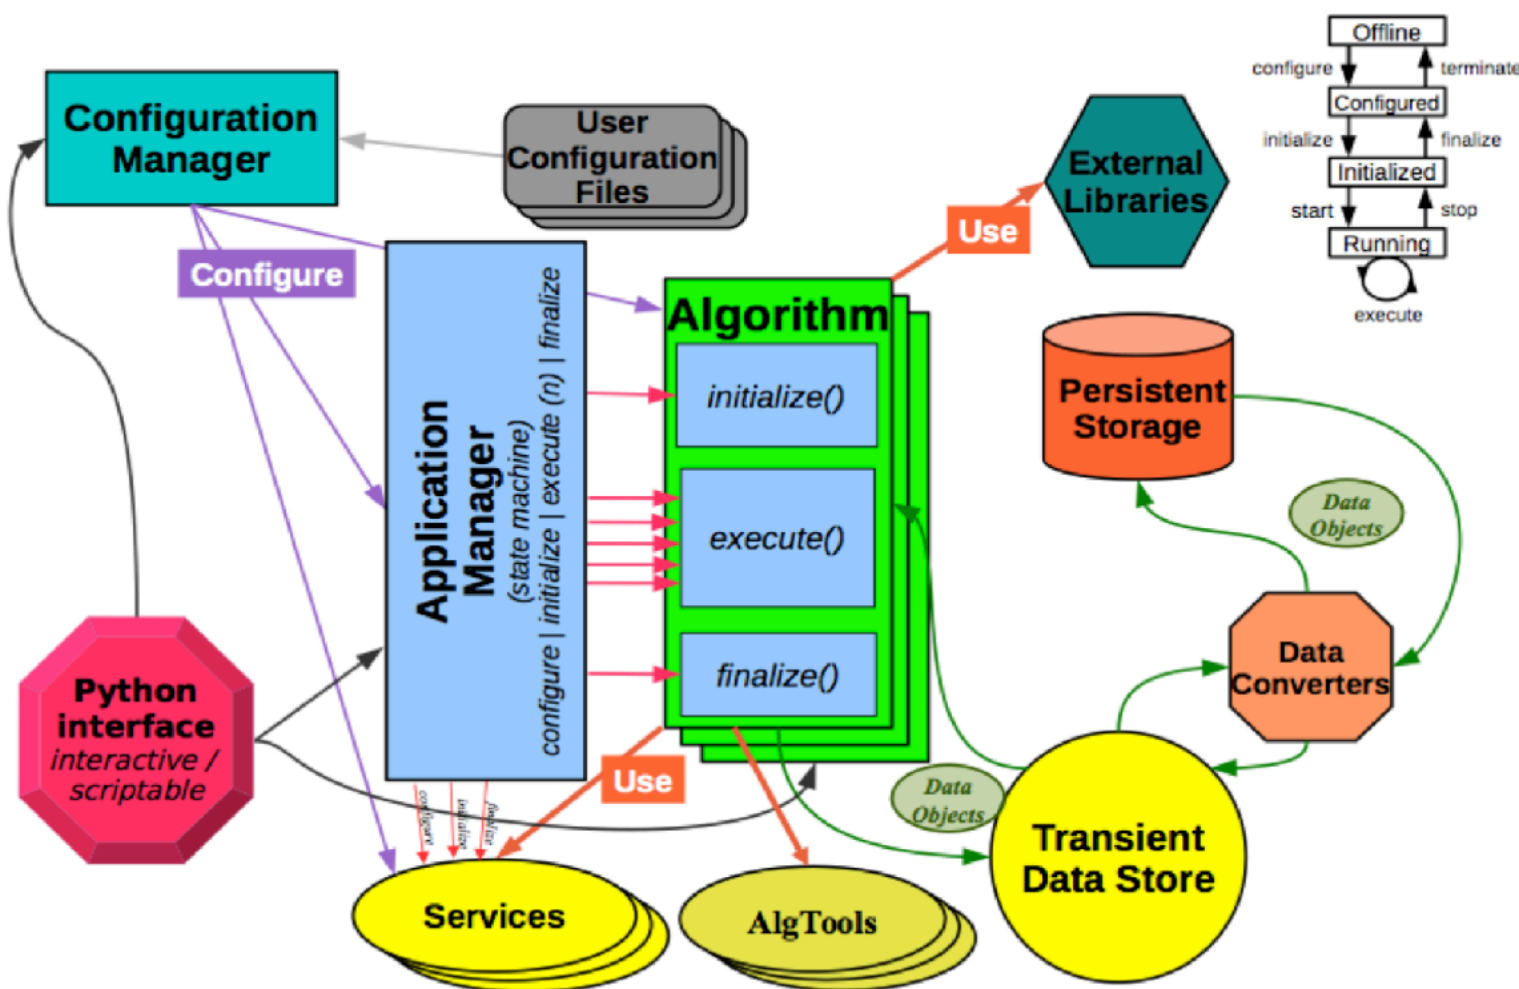
Testing!

Process	Rate/module	size/instance	size/module/year
Beam event	41/day	6 GB	47 TB/year
Cosmic rays	4,500/day	6 GB	9.7 PB/year
Supernova trigger	1/month	115 TB	1.4 PB/year
Calibrations	2/year	750 TB	1.5 PB/year
Total			12.9 PB/year

*per module
numbers*

- | | <i>size in MB</i> | <i>% of 2GB</i> | |
|---|-------------------|------------------------|--|
| • RAW event on ATLAS | 3 | 0.15% | |
| • RAW event of one APA | 40 | 2% | |
| • RAW event on ProtoDUNE | 200 | 10% | <i>approximate
uncompressed size</i> |
| • RAW event on DUNE | 6,000 | 300% | |
| • RAW event DUNE Supernova | 460,000,000 | 230,000,000% | |
| • Software frameworks generally read the RAW event into memory, copy the data to intermediate data products and have an output buffer for writing | | | |
| • <i>ProtoDUNE is already stretching the limit of this paradigm</i> | | | |
| • Moving on to the full DUNE Far Detector, one module of 150 APAs is 6GB | | | |

DUNE software framework



Sketch of the Athena/Gaudi framework

- The software framework that performs reconstruction and simulation of **DUNE** data, and analysis of those data products, needs particular attention to memory
- **DUNE** currently uses **Art/LarSoft** on (Proto)**DUNE** data, but there are clearly **BIG** challenges ahead of us and we need to understand exactly what those are
- In time-honoured tradition, **DUNE** made a Taskforce to investigate

DUNE software framework

- Software framework elements to consider:
 - **I/O**: transient / persistent data models and memory management
 - **State Machine**: configure, initialize, execute, finalize
 - **Configuration**: user config vs pre-defined
 - **Conditions**: interface, IOVs, DB backends including local
 - **Algorithms / Tools / Services**: state and concurrency
 - **External software packages** (*particularly AI/ML*)
 - **Heterogeneous computing**: *accelerators of all kinds*
- Investigate the physics use cases
 - *Are we including physics analysis as well?*

DUNE processing complexity

RAW Transforms and Hit Finding


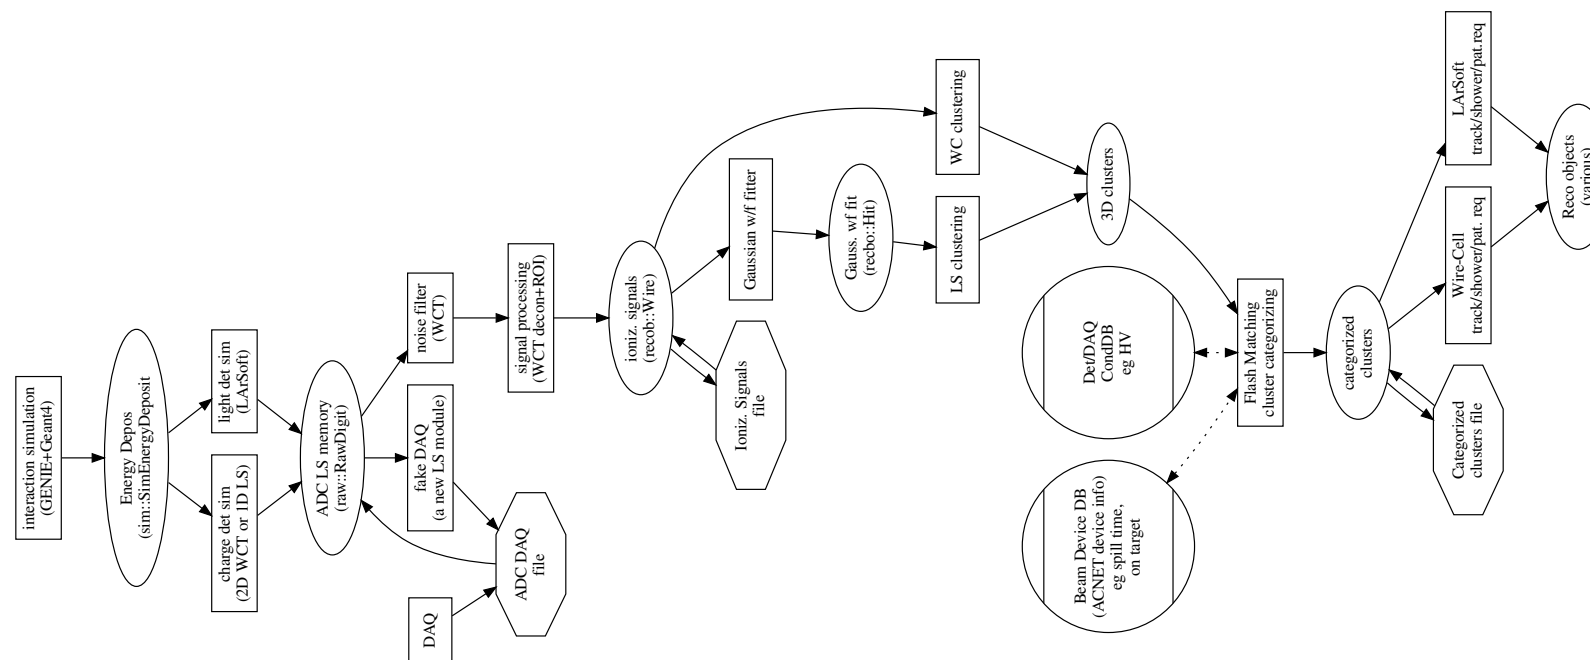
Traditional Reconstruction

Advanced Reconstruction

Physics Object Ident/Classification

Interaction Selection

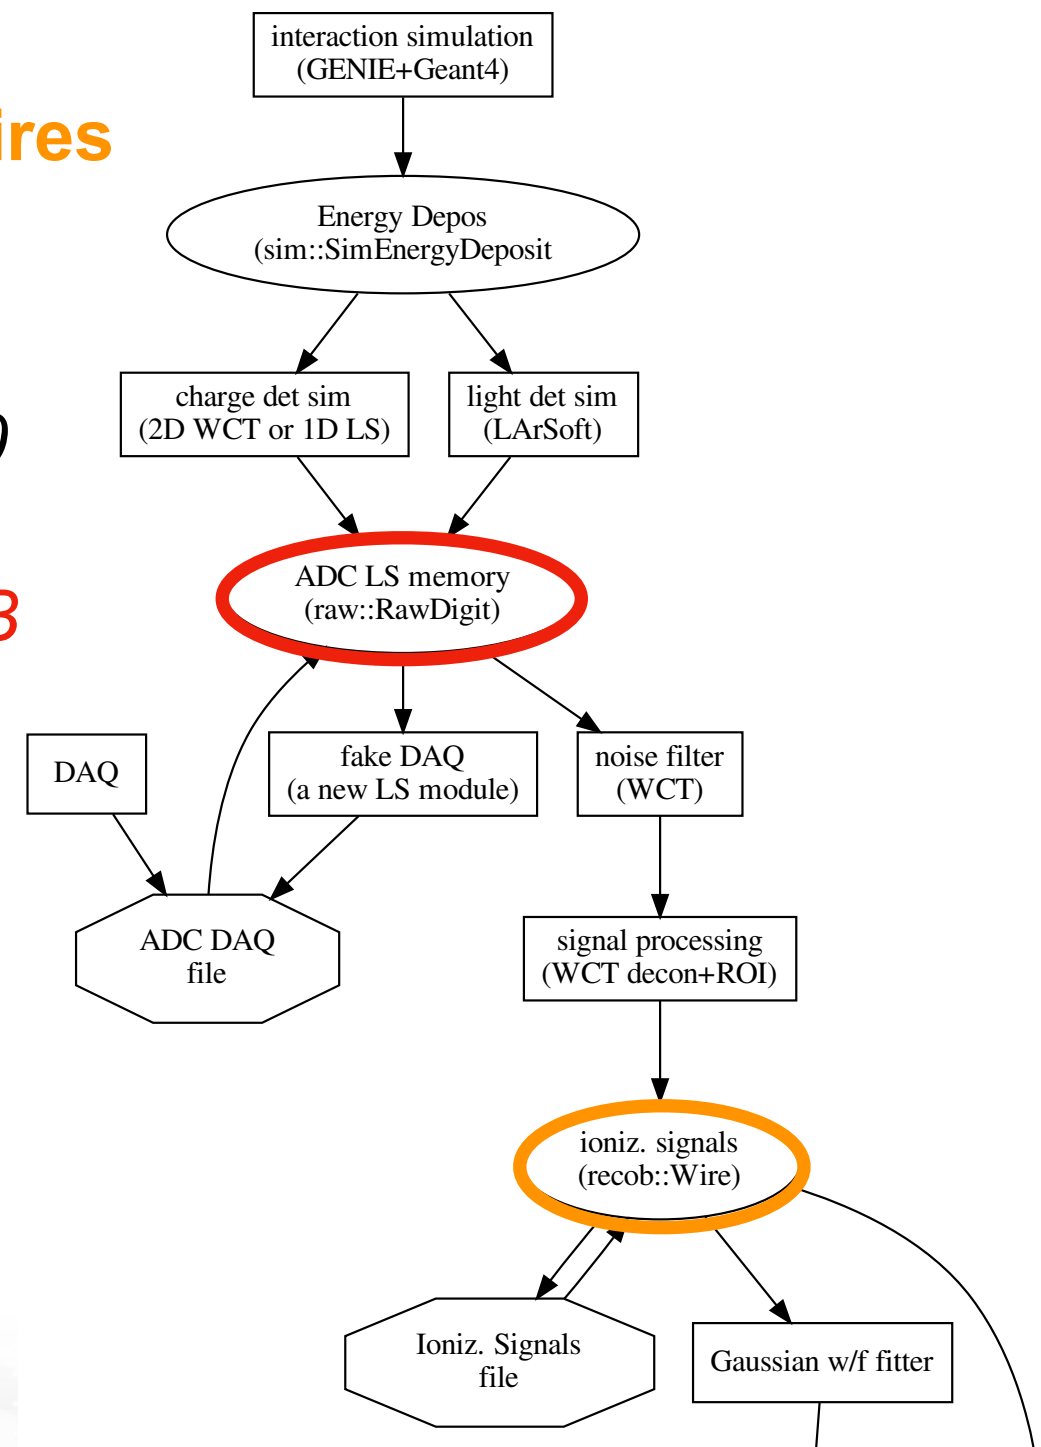
Ensemble Analysis



- **First stage processing steps are the most difficult**, processing the huge **Raw** digits to produce **Wires**
 - Homogeneous data should be well-suited to modern architectures, but there is a LOT of data!
- **Unique challenges for DUNE:- 6GB nominal readout, 1/2 petabyte supernovae**
 - Delivering this data efficiently to heterogeneous computing hardware that hasn't been built yet will be challenging ! Learn lessons from other projects
 - **Significant data reduction of two orders of magnitude in this first step**


DUNE raw data processing

- THE challenge is getting from **Raw digits** to **Wires**
- Raw digits uncompressed for **DUNE** is **~6GB**
 - 2 MHz readout, 6ms window = 12k samples
 - 12k samples * 12 bits * 2560 channels * 150 APAs = **~ 6GB**
 - *The framework will see something like 30GB*
 - inflated to 32 bits and two copies
- **ProtoDUNE** works because it's 6/150th
- For **DUNE** we will need *sub-event processing*
- ...and then there are the *supernovae*...




DUNE supernovae

See talk by K. Scholberg




- **DUNE** will be sensitive to real supernovae events roughly once every **30 years**, but we expect a *test or false alarm once a month*
- **A supernova trigger will read out the detector for 100 seconds**
- That makes for almost half a Petabyte of data in one “event”
- Ideally we would process this quickly and give input to the SuperNova Early Warning System (SNEWS)
- *But it's half a Petabyte !*

Calibration system



- Bring in 266 nm light via optic fibers to the APA and illuminate photocathode (Al) tabs on the CPA.
- Precise estimation of electron drift time
- Diagnosis: check if detector is awake
- Measure drift velocity and field distortions
- Vertex reconstruction on CPA

PE targets on the CPA



- Far detector calibrations will be even bigger than a supernova, about **3PBs** in total (for all four modules) and would also ideally be processed quickly to optimise detector performance

Software framework requirements

- The DUNE Taskforce collected all of the physics-driven requirements and delivered a report (right):
- Converged on **43 requirements** in broad categories:
 - *Memory management*
 - *Data and I/O layer*
 - *Configuration requirements*
 - *Concurrency and Multithreading*
 - *Reproducibility and provenance*
 - *Random numbers, machine learning and conditions*
 - *Physics analysis*

DUNE Software Framework Requirements Taskforce Report Executive Summary

This taskforce report was commissioned by the DUNE computing consortium. The scope of the report was to provide an enumeration of the needs of the DUNE experiment, as driven by its physics mission, in regards to a software framework for data processing and analysis.

The taskforce was composed of representatives from the different physics missions of DUNE, scientists with extensive experience with large scale data processing and analysis from outside of DUNE¹ and technical experts in software framework design, including the current conveners of the HSF² frameworks working group.

Task force members

Co-chairs - Andrew Norman (FNAL) and Paul Laycock (BNL)

DUNE members - David Adams (BNL), Adam Aurisano (U. Cincinnati), Chris Backhouse (UCL), Mary Bishai (BNL), Claire David (York), Tom Junk (FNAL), Tom LeCompte (ANL), Chris Marshall (LBL), Brett Viren (BNL)

Advisors - Brian Bockelman (Madison), Chris Jones (FNAL), Kyle Knoepfel (FNAL), Liz Sexton-Kennedy (FNAL), Vakho Tsulaia (LBL), Peter Van Gemmeren (ANL)

General framework requirements

For ease of reference, this executive summary lists the enumerated framework requirements defined by the Software Framework Requirements Task Force, more info on the task force can be found here:

https://wiki.dunescience.org/wiki/Software_Framework_Requirements_Task_Force

Brave readers are encouraged to read the full document to understand the context and nuances of each of the requirements, the wording here is the same as the full text. While there may be overlap, the complete set of requirements as derived from various considerations is presented and no attempt at reducing this list is made here, rather that is left to framework designers when drawing up specifications. Considerations on "Utilities" and "Desired Features" are also presented at the end of the document to capture useful discussions and provide additional context.

We list in the following the formal requirements determined by the taskforce.

¹ Included scientists from ATLAS, CMS, Belle II, NOvA, MicroBooNE, CDF and D0

² High Energy Physics Software Foundation

DUNE using HSF expertise

- The DUNE Taskforce collected all of the physics-driven requirements and delivered a report (right):
- **Q. Can we leverage the expertise and experience from the wider community?**
- **A. We asked for help from HSF!**
 - A panel of framework experts from different experiments, assembled by the **HSF Framework group convenors**, will provide guidance to help DUNE work out the next steps
 - Workshop later this month to give us feedback - **thank you HSF!**
- *Expect this to be useful for the broader LArTPC community (SBND, Icarus, MicroBooNE...)*

DUNE Software Framework Requirements Taskforce Report

Executive Summary

This taskforce report was commissioned by the DUNE computing consortium. The scope of the report was to provide an enumeration of the needs of the DUNE experiment, as driven by its physics mission, in regards to a software framework for data processing and analysis.

The taskforce was composed of representatives from the different physics missions of DUNE, scientists with extensive experience with large scale data processing and analysis from outside of DUNE¹ and technical experts in software framework design, including the current conveners of the HSF² frameworks working group.

Task force members

Co-chairs - Andrew Norman (FNAL) and Paul Laycock (BNL)

DUNE members - David Adams (BNL), Adam Aurisano (U. Cincinnati), Chris Backhouse (UCL), Mary Bishai (BNL), Claire David (York), Tom Junk (FNAL), Tom LeCompte (ANL), Chris Marshall (LBL), Brett Viren (BNL)

Advisors - Brian Bockelman (Madison), Chris Jones (FNAL), Kyle Knoepfel (FNAL), Liz Sexton-Kennedy (FNAL), Vakho Tsulaia (LBL), Peter Van Gemmeren (ANL)

General framework requirements

For ease of reference, this executive summary lists the enumerated framework requirements defined by the Software Framework Requirements Task Force, more info on the task force can be found here:

https://wiki.dunescience.org/wiki/Software_Framework_Requirements_Task_Force

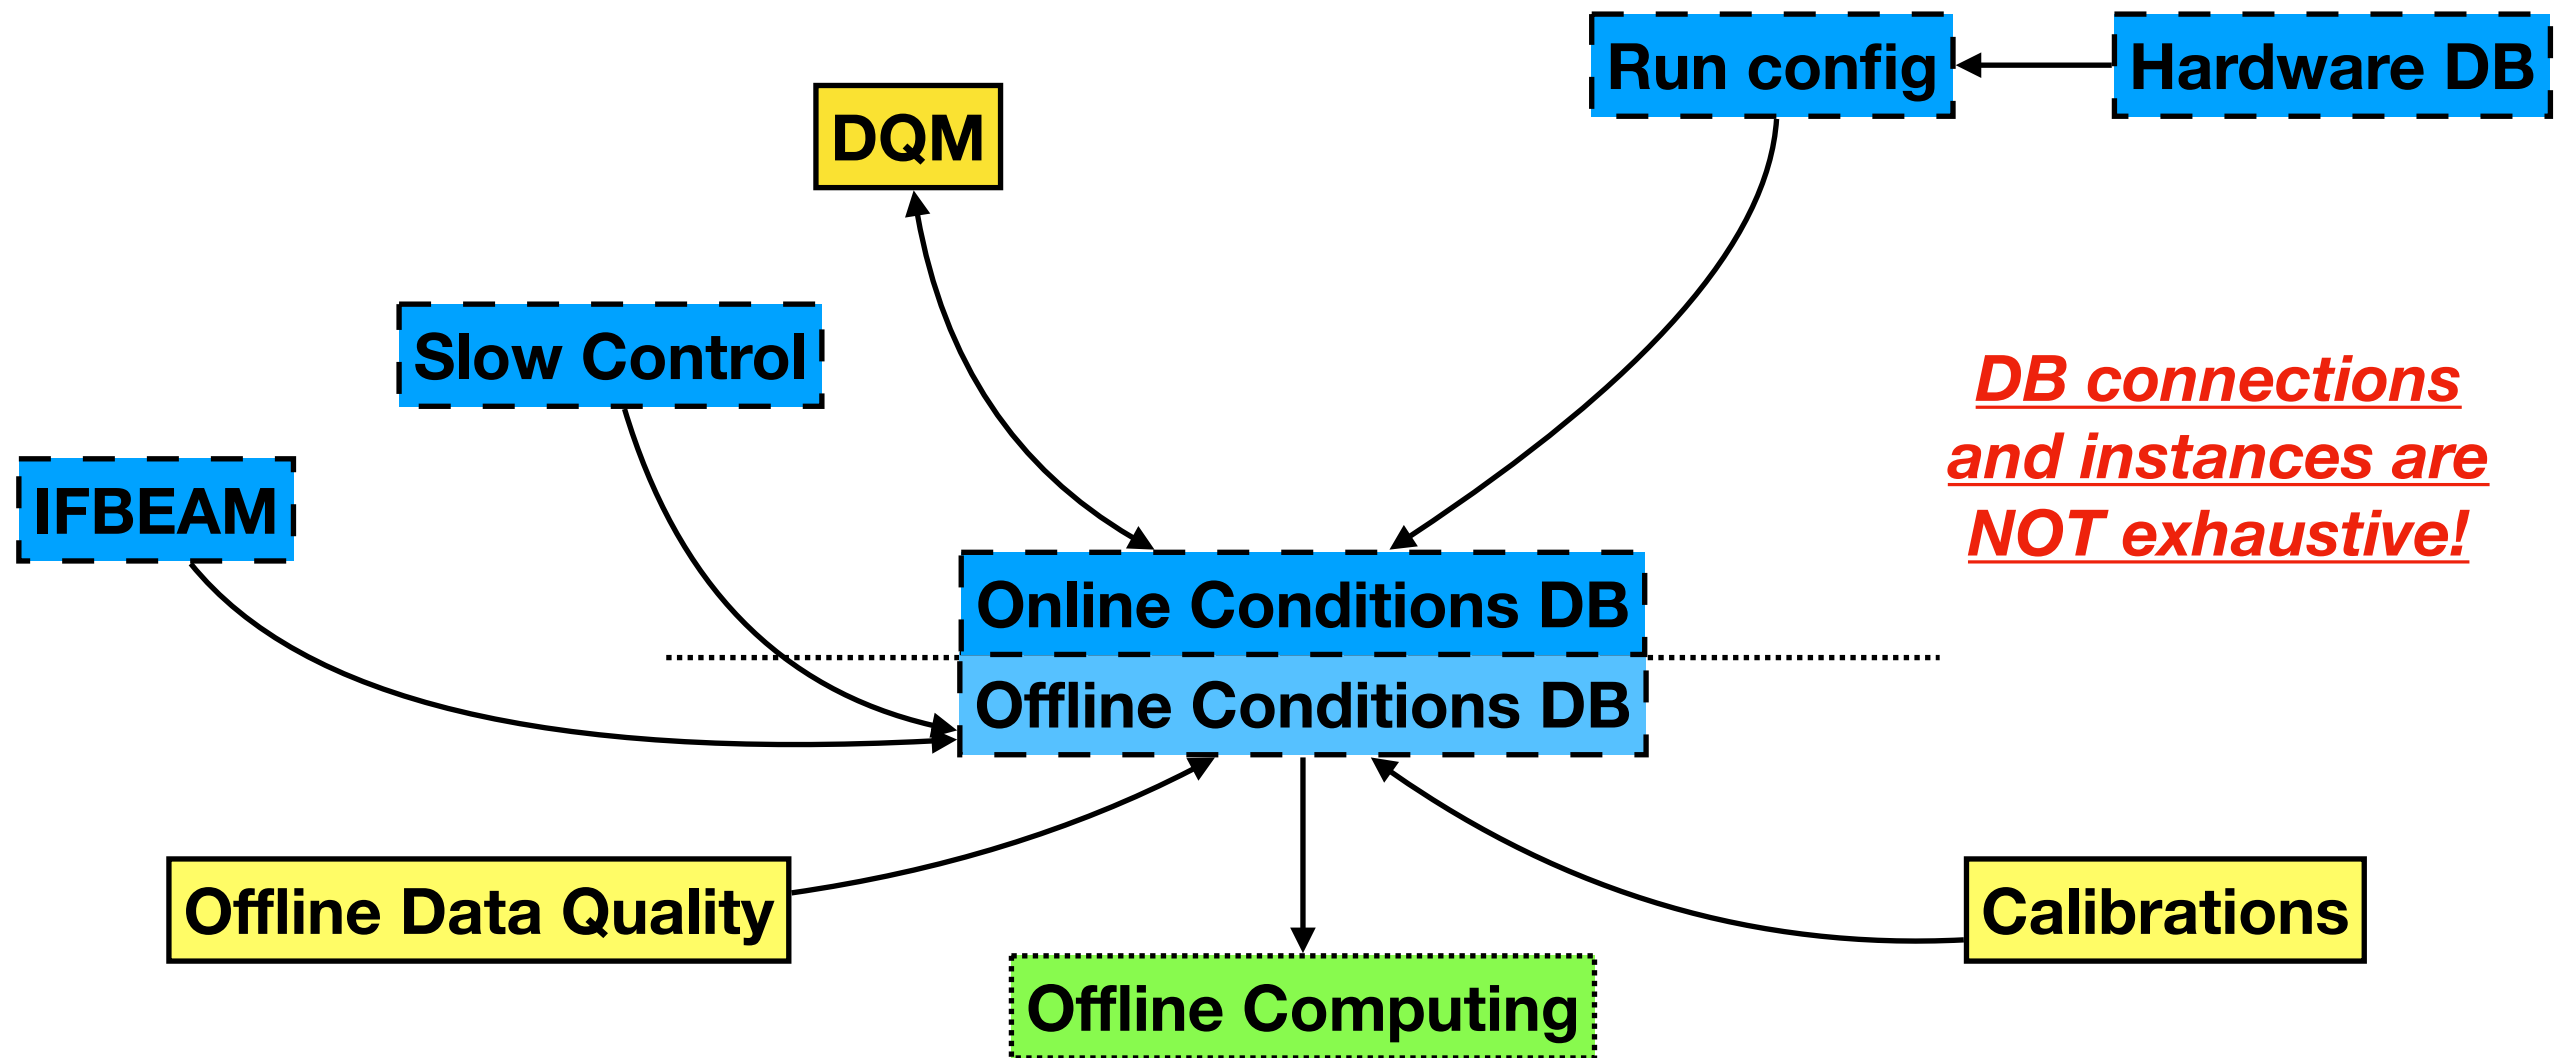
Brave readers are encouraged to read the full document to understand the context and nuances of each of the requirements, the wording here is the same as the full text. While there may be overlap, the complete set of requirements as derived from various considerations is presented and no attempt at reducing this list is made here, rather that is left to framework designers when drawing up specifications. Considerations on "Utilities" and "Desired Features" are also presented at the end of the document to capture useful discussions and provide additional context.

We list in the following the formal requirements determined by the taskforce.

¹ Included scientists from ATLAS, CMS, Belle II, NOvA, MicroBooNE, CDF and D0

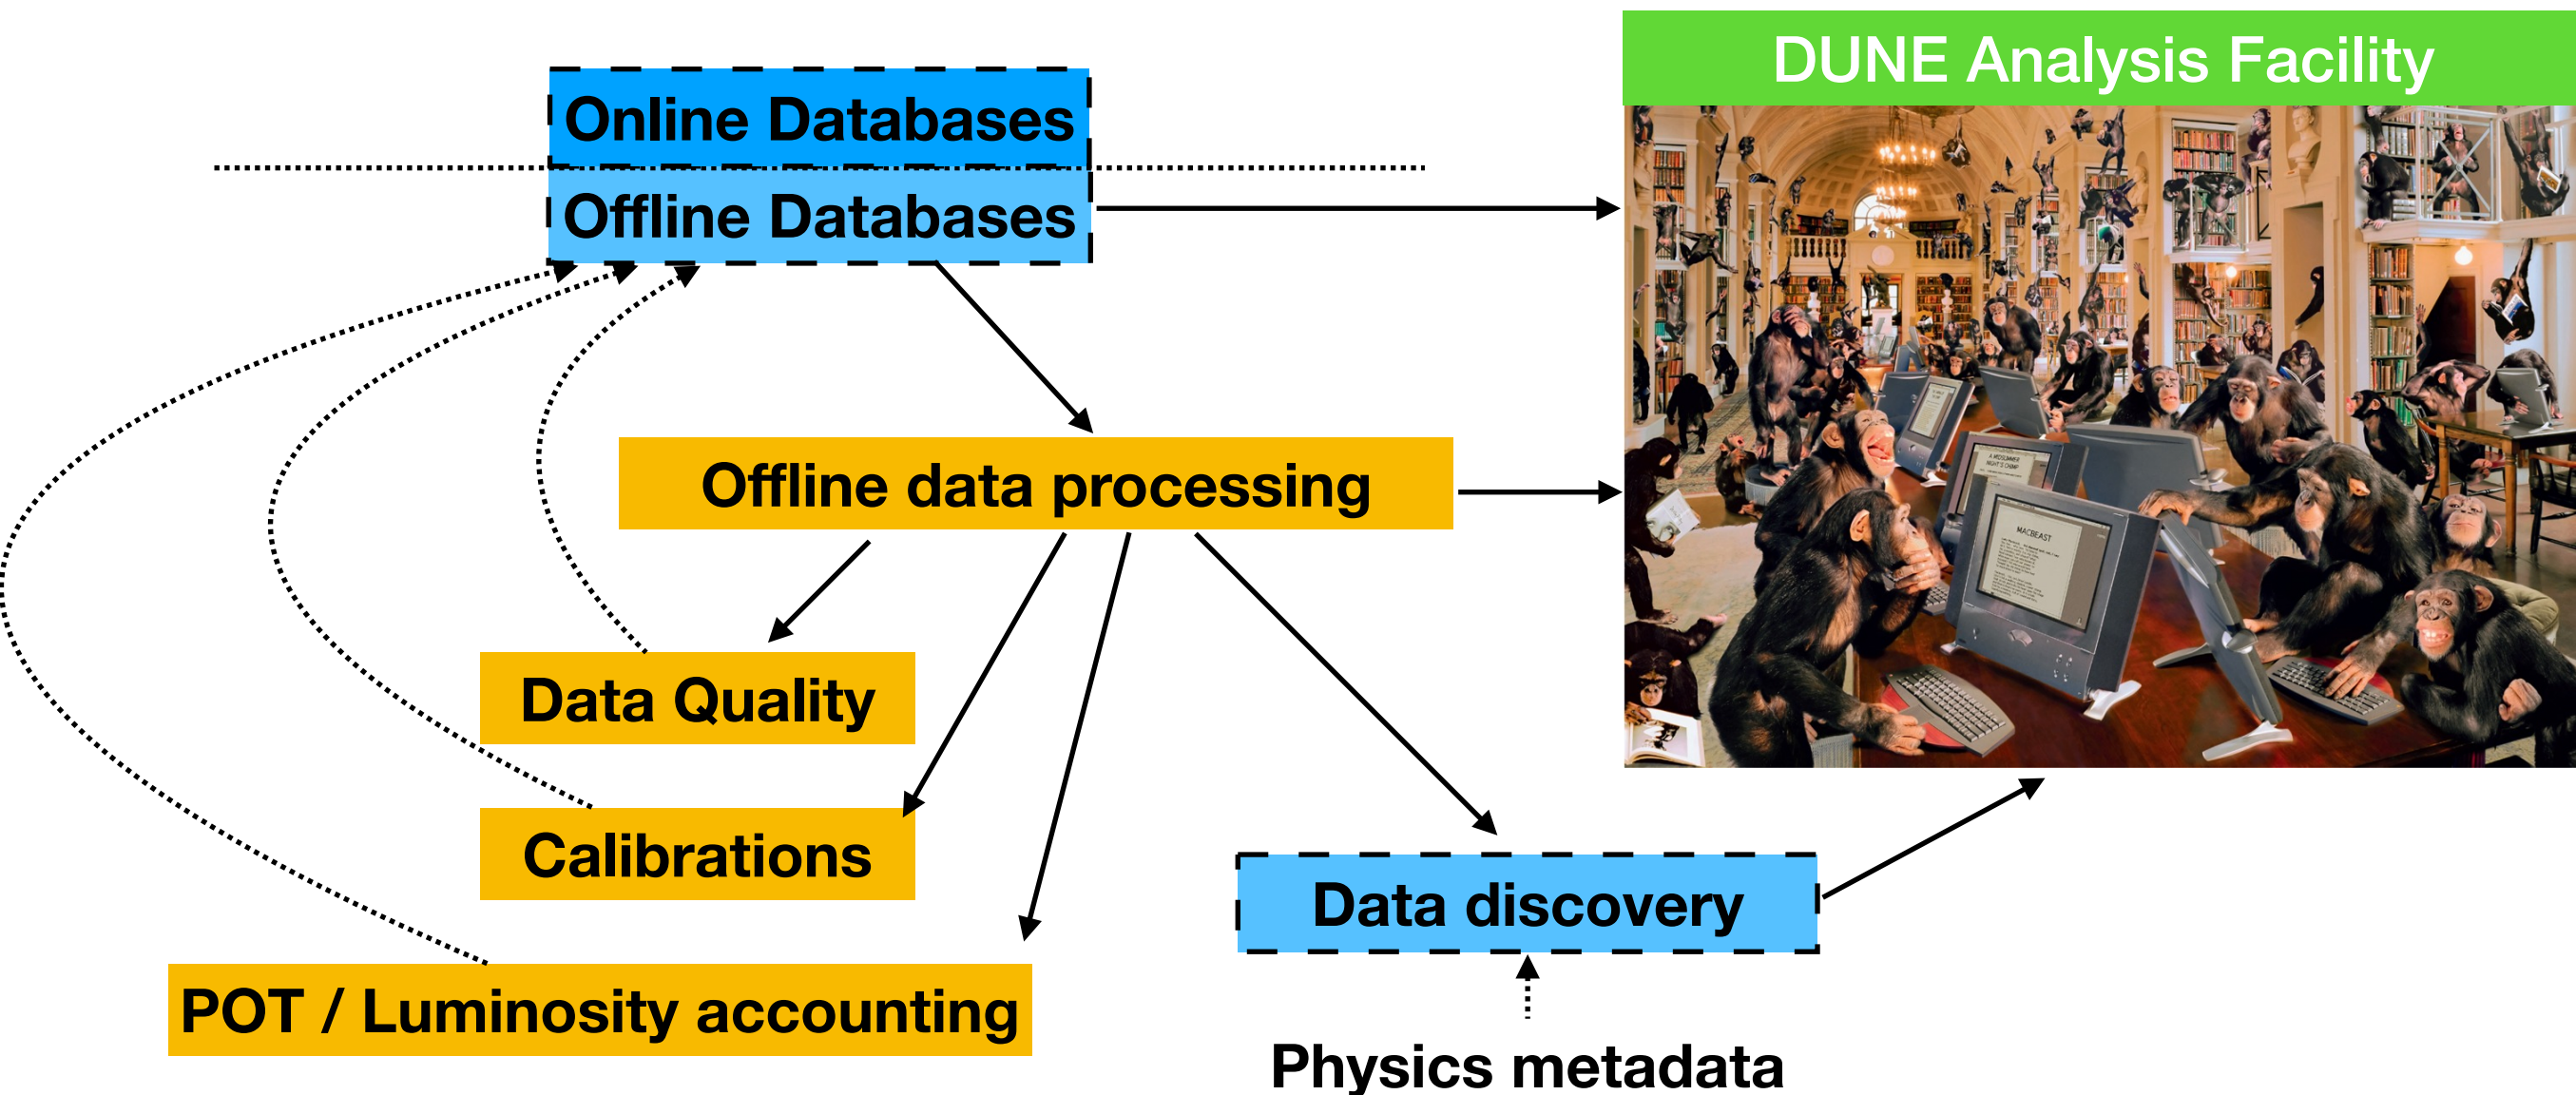
² High Energy Physics Software Foundation

Databases



- **Online metadata producers write directly to online databases**
 - Either dedicated databases where appropriate, or the Conditions DB
 - Some online metadata is needed offline, filtered to the Conditions DB
- **Offline metadata producers write to the offline Conditions DB**
 - Some offline metadata is needed online, via the Conditions DB

DUNE data processing workflows



- Preferable to have one coherent ***interface*** to metadata for offline
 - The shape of DUNE data and calibrations may pose some interesting technical challenges
 - Dynamic data discovery (give me all of the data that fulfils these conditions) is an expectation of neutrino physics, data management will use **Rucio**, see talk on MetaCat by Igor

DUNE computing model




Image from the WLCG-DOMA group

- Still in the process of understanding DUNE needs with ProtoDUNE data
- Leverage existing solutions as much as possible, e.g. **Rucio** is already adopted
 - Working **with** the broader community, e.g. DUNE contributed **S3** and **Swift** object store support to **Rucio**
 - This pays off as the community grows and we are all fixing common software
- Engaging with **HSF** and **WLCG** for experience and expertise as much as solutions

Summary

17-21 May 2021
Europe/Paris timezone

Search...

Overview

- Timetable
- Contribution List
- My Conference
- My Contributions
- Registration
- Participant List
- Guide for Participants
- Videoconference Rooms
- Mattermost Channels
- Instructions for Presenters
- Review and Proceedings
- Jobs@vCHEP**

25th International Conference on Computing in High-Energy and Nuclear Physics

vCHEP2021

Welcome! The CHEP conference series addresses the computing, networking and software issues for the world's leading data-intensive science experiments that currently analyse hundreds of petabytes of data using worldwide computing resources.

vCHEP 2021 will be held *as a virtual event* between **Monday-Friday 17th-21st May 2021**.

Starts 17 May 2021, 14:50
Ends 21 May 2021, 18:10
Europe/Paris

[vCHEP Mattermost Signup](#)

- The size and shape of DUNE data poses some very interesting challenges
- We have lots of interesting work to do, and there are opportunities to join us
 - Visit the [Jobs@vCHEP](#) page to find out more!

Backup

DUNE CPU needs

RECONSTRUCTION

- ProtoDUNE events are more complex than our long term data.
 - ~500 sec to reconstruct 75 MB compressed – 7 sec/MB
 - For FD, signal processing will dominate at about 3 sec/MB
 - < 30 PB/year of FD data translates to ~100 M CPU-hr/year
 - That's ~ 12K cores to keep up with data. But no downtimes to catch up.
- Near detector is unknown but likely smaller.

ANALYSIS (Here be Dragons)

- NOvA/DUNE experience is that data analysis/parameter estimation can be very large
 - ~ 50 MHrs at NERSC for NOvA fits
- Not incredibly demanding CPU needs but no downtime and the calibration and supernovae use cases are particularly demanding, up to ~100k cores needed

DUNE Detector

Quantity	Value	Explanation
Far Detector Beam: Single APA readout Single APA readout APAs per module Fullmodule readout Beam rep. rate Signal processing CPU time/APA Signal processing CPU time/input MB Memory footprint/APA	41.5 MB 16.6 MB 150 6.22 GB 0.83 Hz 40 sec 2.5 sec/MB 0.5-1 GB	Uncompressed 5.4 ms ×2.5 compression Uncompressed 5.4 ms Untriggered from MC/ProtoDUNE compressed input ProtoDUNE experience
Supernova: Single channel readout Four module readout Trigger rate	300 MB 460 TB 1 per month	Uncompressed 100 s Uncompressed 100 s (assumption)