Data Analysis at the Intensity Frontier

Philippe Canal, FNAL

HEP analysis eco-system workshop

22 May 2017

Thanks and Acknowledgements

Kim Siang Shaw, University of Washington, g-2

Rob Kutsche, FNAL, Mu2e

Jiyeon Han, FNAL, MINERvA

Adam Lyon, FNAL, g-2

Alexander Himmel, FNAL, Nova, DUNE

Kenneth Richard Herner, FNAL, FIFE

Michael Kirby, FNAL, MicroBooNE, NOvA, LArIAT

Oliver Gutsche, FNAL

The presentation is based on input they passed along, all over-simplification, over-generalization, omissions and mistakes are of course mine.

Neutrino Experiments

Detect Neutrinos and measure their masses

Important Standard Model measurement, candidates for dark matter

Produce neutrino beam and direct it to far detector, compare near and far detector measurements

- **Current experiment: NOvA**
  - –500 miles to Ash River, Minnesota
- **Future experiment: DUNE/LBNF**
  - –800 miles to Sanford Underground Research Facility, in Lead, South Dakota

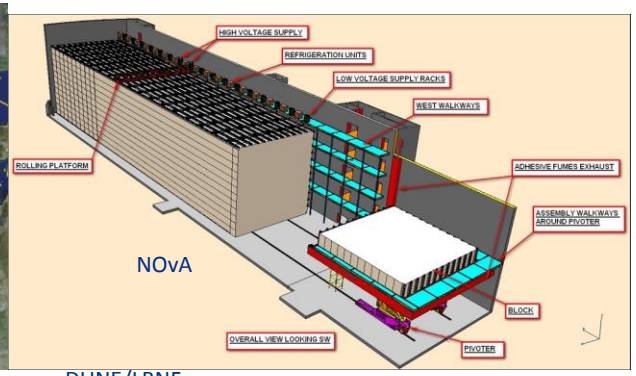

NOvA

DUNE/LBNF

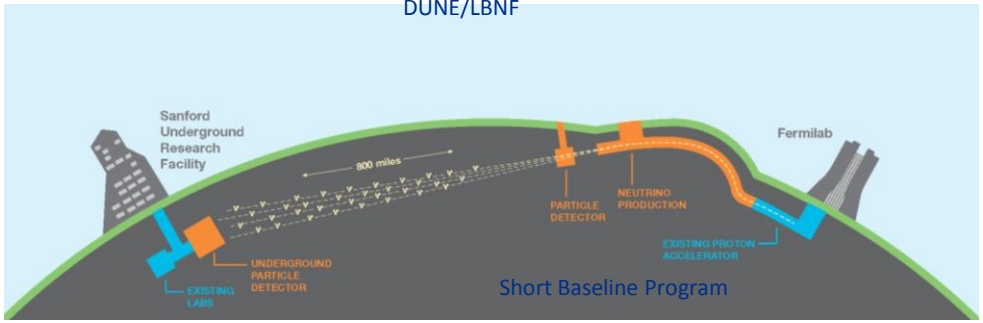

Short Baseline Program

MicroBooNE Scientific Goals

- MicroBooNE is a liquid argon time projection chamber (LArTPC) experiment with multiple physics and R&D goals.
  - Resolve the low-energy excess of electron-like neutrino interactions observed by MiniBooNE in the booster neutrino beam (BNB).
  - Improve measurements of n-Ar cross sections in the O(1 GeV) energy range.
  - MicroBooNE is part of a LArTPC R&D program aimed at developing LArTPC technology (both hardware and software) for use in increasingly higher mass neutrino detectors.
  - Supernova neutrino detection, in event of a nearby supernova.

MicroBooNE
Reconstructed Data Volume & CPU Requirements

	FY16	FY17	FY18	FY19
Data events / year	1.29E+07	2.22E+07	2.22E+07	2.22E+07
Reconstructed Data Event Size (MB)	120	120	120	120
Initial Reconstruction Data Size (TB)	1548	2664	2664	2664
Reprocessed Events	5.00E+05	2.50E+06	5.73E+07	5.00E+06
Reprocessed Data Size (TB)	60	300	6876	600
Total Data Size (TB)	1608	2964	9540	3264
Reconstruction CPU / event (s/ev)	800	800	800	800
Initial reconstruction (cpu hours)	2.9E+06	4.9E+06	4.9E+06	4.9E+06
Reprocessing (cpu hours)	1.1E+05	5.6E+05	1.3E+07	1.1E+06
Total reconstruction (cpu hours)	3.0E+06	5.5E+06	1.8E+07	6.0E+06
Total reconstruction (cpu years)	340	627	2017	690

Common Threads

- Art Framework
  - simulation and running reconstruction on simulation and data.
  - One exception MINVERvA uses Gaudi
- ROOT files
  - Will **require** ROOT for **life** of experiment
- Fermilab FIFE grid tools
- SAM for dataset and metadata handling
- Use of UPS and likely move to SPACK
- Smaller experiment but combined approach scale of LHC experiments

g-2 data tiers

Tier definition: A proposal

- Proxies for Tier[0,1,2]:

  - **raw**: recorded binary data

    RawIslands (t, V)

  - **reconstruction**: reconstructed data at detector component level

    CrystalHits (t, E, χ^2)

  - **physics**: reconstructed physics objects readily for data analysis

    CaloClusters (t_{clus}, E_{clus}, w_{clus})

8

Data File Workflow

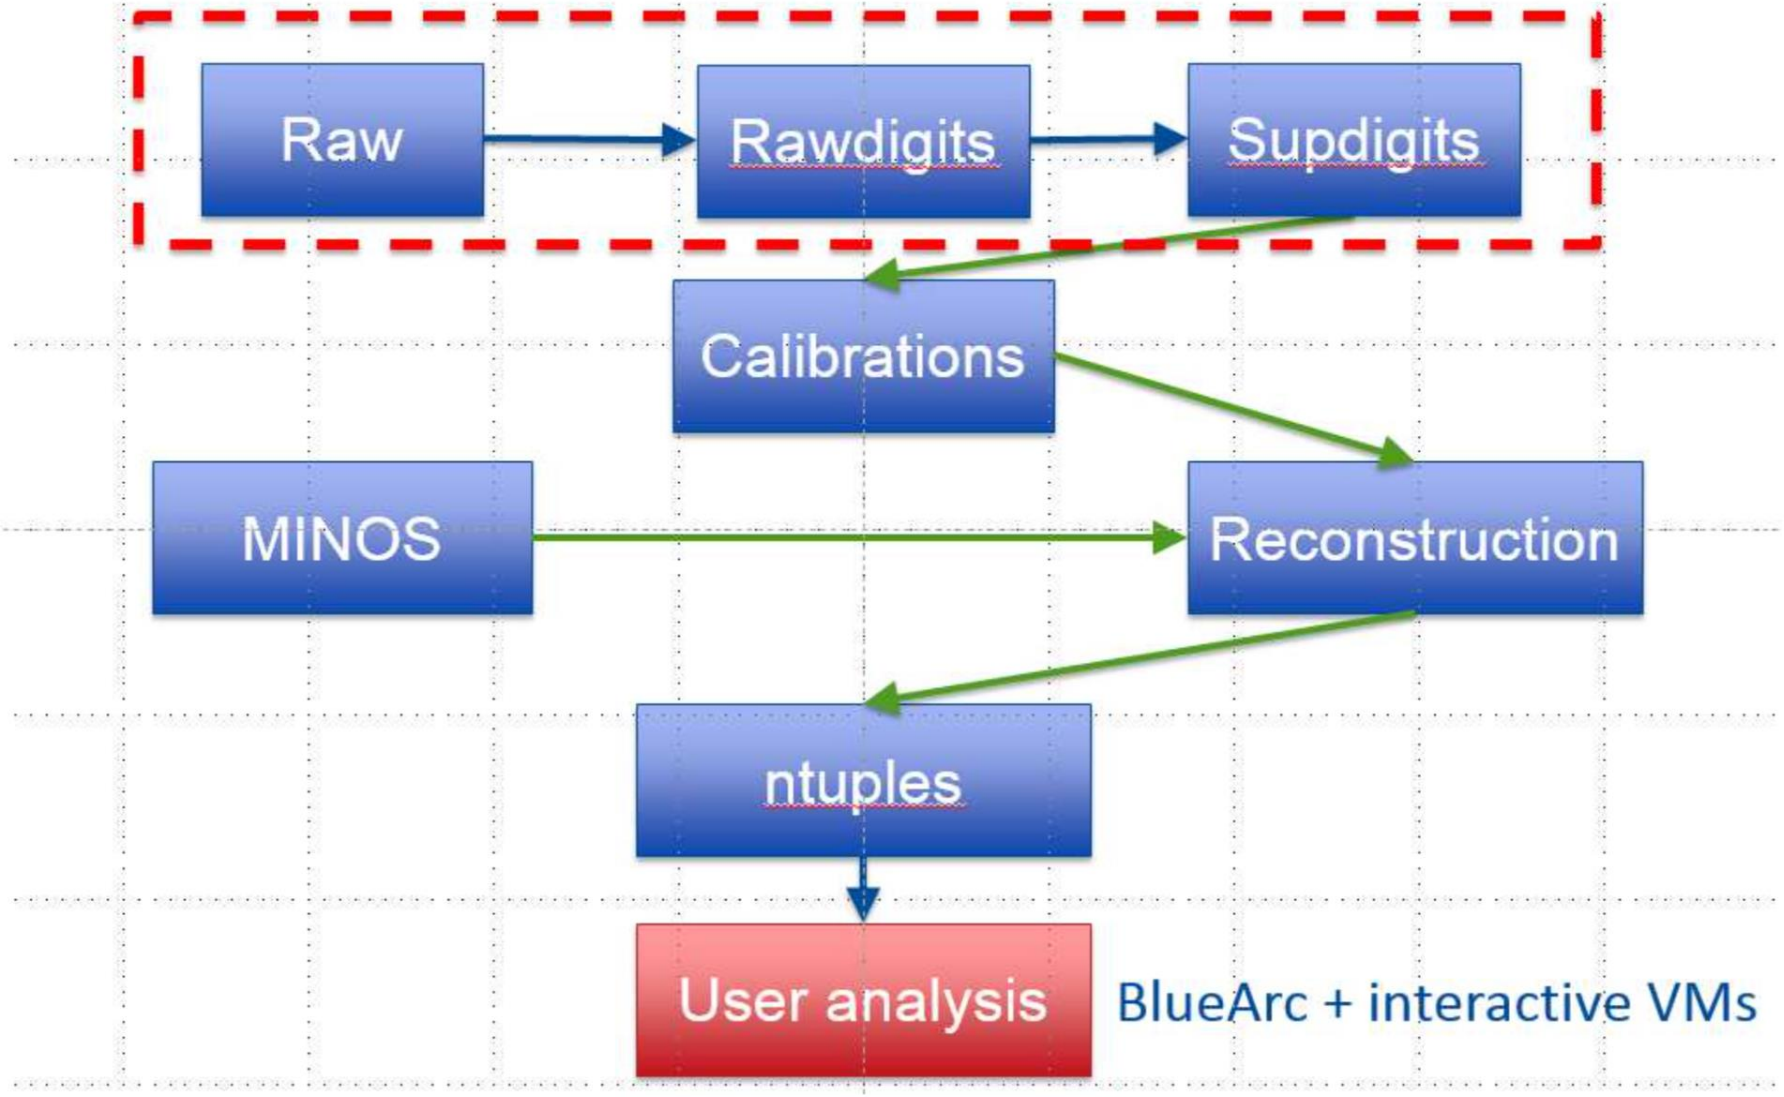

Common problems, common solutions

- Bulk of Intensity Frontier experiments on average are 1-2 orders of magnitude smaller than LHC experiments
  - Often lack sufficient expertise or time to tackle all problems, e.g. software frameworks or job submission tools
  - Also much more common to be on multiple experiments in the neutrino world
- By bringing experiments under a common umbrella, can leverage each other's expertise and lessons learned
  - Greatly simplifies life for those on multiple experiments
- Common modular software framework is also available (ART, based on CMSSW) for most experiments

Except for Analysis …

- g-2
  - “Gallery” from Root, Python (with PyRoot) and R.
  - Bare ROOT file produced by private Art analyzer.
  - Convert data to HDF5 and analyze with Python or R (soon)
- NOvA
  - binary -> artroot -> reconstruction -> CAF -> dCAF
- LArIAT
  - binary -> artroot -> reconstruction ; root ntuple.
  - numpy/scipy format file for ML feedback into reconstruction files
- MicroBooNE
  - binary -> artroot -> reconstruction -> anatree
  - HDF5 format file for ML based on images the LArTPC response

Except for Analysis …

- Mu2e
  - Currently each group its own format/workflow
    - Art-module producing histogram
    - TTree with object + interactive ROOT
    - Flat ntuples + interactive ROOT
  - But the 3 groups all use common classes/functions to make cuts and to compute physics quantities
  - Plan for a committee to define an Analysis Model; including, but is not limited to, the software tools and data formats
- NOvA
  - “ntuple” format with summary information for analysis.
  - Some significant added functionality for NOvA, DUNE analysis
- MINERvA
  - All of data/analysis workflow is managed by Gaudi tool

GIRD usage for analysis?

- MINERvA
  - Uses FNAL Grid for production, extending to OSG
  - Long term plan to use OSG for analysis users.
- g-2
  - We mainly use the Grid for MC production and reconstruction. Not so much for analysis yet. We use the Fermilab FIFE grid tools. We use SAM for dataset and metadata handling.
- NOvA / LArIAT / MicroBooNE
  - Usage of grid for analysis step: experiments don't have their own grids, just FIFE GPGrid and OSG
  - xrootd hasn't yet fully propagated to all the experiments, but is starting to take hold. Not a game changer for grid use.
  - SAM4Users has made the SAM file catalog usable by individual users and has been really successful in NOvA and MicroBooNE,

Machine Learning

- NOvA: Caffe
  - Open-source machine learning software from the Berkeley Vision Lab.
- LArIAT
  - numpy/scipy format file for ML feedback into reconstruction files
- MicroBooNE
  - HDF5 format file for ML based on images the LArTPC response

Wishes

- Along with multithreading and vectorization, we would like to see tighter integration with python (g-2)
- Worth some thought about how our data formats can connect with other open-source tools (NOvA)
  - We have one solution, which is to convert, but the community as a whole has been exploring a variety of solutions.
- access to GPUs is the biggest limiter right now it seems. But I think that should be solved in the next 1-2 years. Access to data is going to be the biggest problem. Large datasets (MB/event is huge) and organization of the datasets for skimming/slimming etc. (NOvA / LArIAT / MicroBooNE)
- **Not** mentioned: concern about performance
  - *“because the data sets that reach analyzers will likely be small, even by today’s standards“*

Special Mention

Whatever the new tools turnout to be, the transition plan should include ***efforts*** to mentor senior people in the tools so that can continue to be ***effective*** mentors of analysis teams that use these tools.
During the transition from FORTRAN to C++ the community did a bad job of this and the community lost a lot of physics expertise.

BACKUP SLIDES

MicroBooNE Scientific Goals

- MicroBooNE is a liquid argon time projection chamber (LArTPC) experiment with multiple physics and R&D goals.
  - Resolve the low-energy excess of electron-like neutrino interactions observed by MiniBooNE in the booster neutrino beam (BNB).
  - Improve measurements of ν-Ar cross sections in the O(1 GeV) energy range.
  - MicroBooNE is part of a LArTPC R&D program aimed at developing LArTPC technology (both hardware and software) for use in increasingly higher mass neutrino detectors.
  - Supernova neutrino detection, in event of a nearby supernova.

Minos Scientific Goals for FY17 and FY18

- MINOS+ used the medium energy neutrino beam developed for the NOvA experiment through June 2016
  - Three-flavor neutrino measurements
  - Sterile neutrino searches in multiple channels
  - Non-standard interaction studies
  - Large extra dimensions investigation
- Conferences and publications
  - Moriond is next conference on agenda
  - Various PRD and PRL publications currently in the works
    - Standard oscillation, sterile models, electron neutrino appearance, etc...

Minos Production

- Timely data reconstruction processed on Fermigrid
  - Run on daily basis – covering previous day's data
  - ~500-800 CPU-hours nightly [under MINERvA]
- Raw data occupies 15.0 TB in Enstore and Dcache
  - 1.2 TB from last year
  - Keep on Dcache disk to reduce tape risk
  - Write vault copy to separate tape robot
- Reconstructed data uses 98.3 TB + 23 TB on dcache and bluarc, respectively
  - Full pass over all data takes 20 TB MINOS, 15 TB MINOS+
  - Plan to transition to dcache completely
    - Data transition complete, MC transition pending

Muon g-2's Scientific Goals

- Installation Phase
- Last vacuum chamber in place
- Detector installation begins
  - –Successful calorimeter readout and laser calibration test
- We hope for a commissioning run this year
- First data taking now in FY18 due to accelerator shutdown schedule

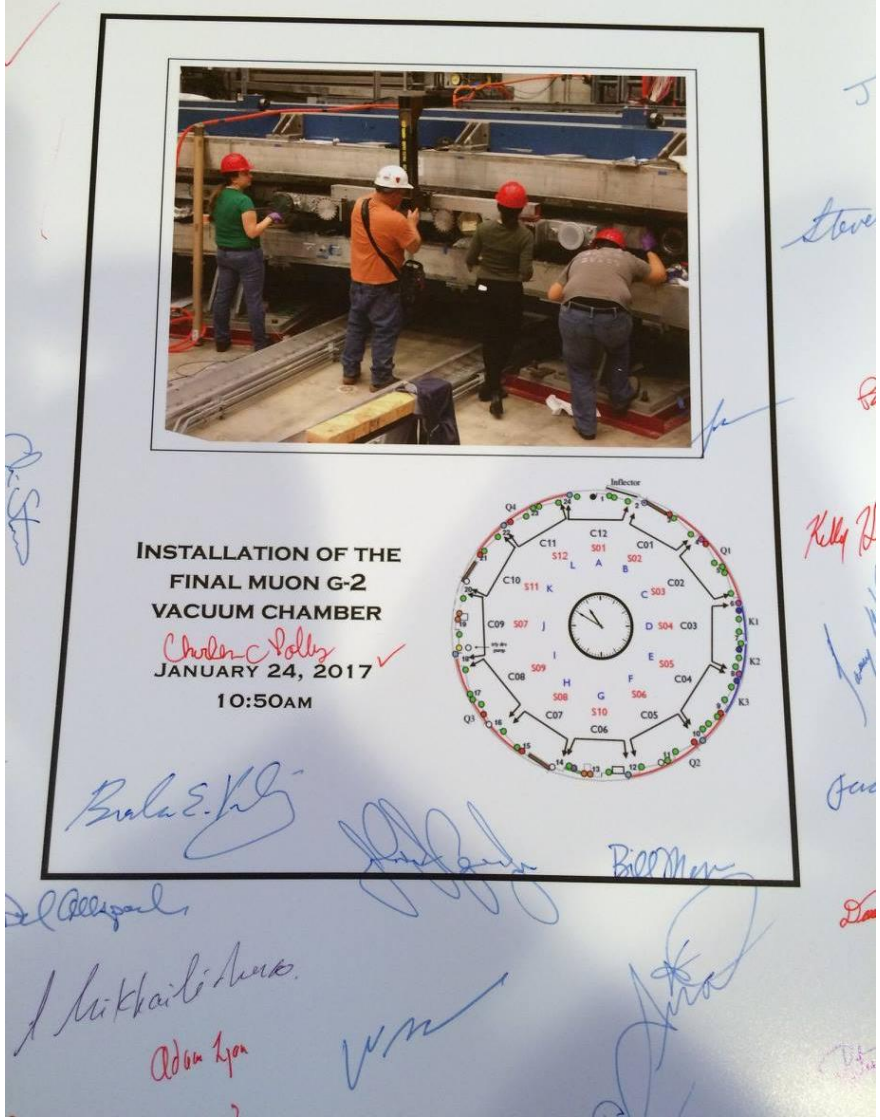

g-2 Tape needs

- We have month-to-month estimates of data taking rates
- Review recommends 2 copies of RAW data
- RAW data is MIDAS DAQ output
  - RAW data should be closest to the DAQ
- Reconstructed data (*art* format) is just a little smaller than RAW (want all waveforms in reconstructed data)
- Add another 0.5 * Reconstructed for simulation and analysis

- FY17: 1.8 PB new tapes needed (mostly for simulation with small data size/not writing out all waveforms)
- FY18: 7.2 PB new tapes needed
- FY19: 8.2 PB new tapes needed

NOvA's Scientific Goals

- Last year NOvA presented results at the NEUTRINO2016 which were among the highlights of the conference.
  - First 2σ evidence of non-maximal θ_{23}
  - Exclusion of a region of θ_{23}-δ_{cp} space at > 3σ
- This would not have been possible without significant computing resources and expert help from SCD!
- We are planning significant updates to these results over the next two years:
  - In late summer 2017 we plan to present an update of the θ_{23} analysis with neutrino data.
  - At NEUTRINO2018 (June) we plan to present an update of the θ_{23}-δ_{cp} analysis with neutrino and antineutrino data.
  - As in 2016, 2018 is a deadline we cannot miss – so we are planning towards it now.

NOvA Large disk requirements

- We accumulate a large amount of data because our far detector is on the surface (~8 billion cosmic rays/day).
  - This comes to ~2 PB/year, will continue through FY19
- Processing multiplies this since ART copies data products to downstream files.
  - Production campaigns get larger with time.
  - We expect 2017 and 2018 to be 1.5 PB each, and 2019 to be 2 PB.
  - Since half of the 2018 campaign will be during 2017, we get the numbers below:

Requirements:

- 8.6 PB on tape now, additional tape needed in…

2017	2018	2019
4.25 PB	2.75 PB	4.0 PB

Scientific goals

LArIAT (~ 80 collaborators): LArTPC with the following physics and R&D goals

SCIENCE OUTLOOK:

ELECTRON VS PHOTON SHOWER DISCRIMINATION

Experimental confirmation for the separation efficiencies (MC determined) - key feature of LArTPC technology

MUON SIGN DETERMINATION (W/OUT MAGNETIC FIELD)

Explore a LArTPC feature never systematically considered (decay vs capture in LAr)

STUDY OF NUCLEAR EFFECTS

Pion Absorption, π^0 from $\pi^\pm$ Charge Exchange, Elastic Cross-Section

Kaon interaction channels

Antiproton annihilation (relevant for n-nbar oscillations)

DEVELOPMENT OF A NEW CONCEPT FOR LAR SCINTILLATION LIGHT COLLECTION

Relate energy deposited to **charge** and **light** for an improved calorimetric energy resolution

LArIAT **Tape storage summary estimates**

Data type	2015 Size [TB]	2016 Size [TB]	2017 Size [TB]	2018 Size [TB]	Total Size in TB
Raw data	2.3	14	14		30.3
Sliced data	1	6	6		13
Reconstruction (guess adds 10%)	0.1	0.6	0.6		1.3
Beam simulation files	0.005 (6 conditions/1000 spills)	-	-		0.005
Detector Simulation + Beam Simulation (10xreal data?)			70	20?	90
					~ 150

MINERvA Scientific Goals

MINERvA is a neutrino cross section experiment measuring neutrino-nucleus interactions over a wide variety of nuclei. The experiment has produced 16+ cross section publications using the low energy tune of the NuMI beam and is looking forward to publishing results using the very large NOvA era NuMI tune. Continue to take neutrino and anti-neutrino data.

This requires keepup processing and monitoring of incoming data from both the MINERvA and MINOS near detector.

Using the recent production looking at 10E20 POT of neutrino data ME analyses are laying the foundation for the next generation of publications.

DUNE Scientific Goals for FY17 and FY18

Two main thrusts in DUNE for 17/18:

- Design work/Sim toward TDR
  - Need to demonstrate sensitivities **(full sim)**
- ProtoDUNE Experiments
  - Single Phase
  - Dual Phase projects

- ProtoDUNE experiments have well defined physics measurements (charge particle interactions on LAr)
- Critical Technical Measurements are technology demonstrations
  - Need to show noise, readout, performance characters at DUNE detector scale.

Time critical: Scheduled beam window ends with start of CERN long shutdown.

- Additional Goals:
  - Improve LArTPC reconstruction techniques
  - Near detector design studies (many options LArTPC, GArTPC, Spectrometer…)

- No specific conference deadlines

Current Software Frameworks in NOvA

ART framework for producing simulation and running reconstruction on simulation and data.

- *Shared framework with other intensity frontier experiments.*
- *Uses root file format.*
- **Maintains "persistency" so processing history is stored for all the data products in the file.**

NOvA-specific "ntuple" format with summary information for analysis.

- *Based on ROOT file format.*
- *Some significant additional functionality for handling NOvA (and now DUNE) analysis.*

Machine learning software framework: Caffe

- *Open-source machine learning software from the Berkeley Vision Lab.*

Future Needs (Nova)

NOvA will continue to need ROOT for the for the lifetime of the experiment.

- *This is the common element in all of our storage.*
- Generally, need support of old version for an extended period of time.
  - We have only recently begun the Root5 → Root6 transition.

Worth some thought about how our data formats can connect with other open-source tools.

- *We have one solution, which is to convert, but the community as a whole has been exploring a variety of solutions.*

Mu2e

Mu2e is still 3 years away from early commissioning data and another year or so until physics data.

We have 3 small groups of people doing sensitivity and background studies on reconstructed simulated events; each group has rolled their own. Due our early stage and limited manpower we have not yet created a standard way to do analysis. We are in the process of staffing a committee and charging it to define an Analysis Model; the charge will include, but is not limited to, the software tools and data formats.

Mu2e

There are two groups each of which writes their own standardized TTree/TNtuples and does its final analysis in interactive root. The third group has it's own art modules that fill the final histograms directly from the Mu2e art data products.

The 3 groups all use common classes/functions to make cuts and to compute physics quantities.

Among the TTree/TNtuple groups there are some TTrees that have objects or collections of objects on the branches but there are also some true TNtuples and some TTrees that are basically just flat ntuples on each branch.

Mu2e

My opinion is that the analysis phase of Mu2e will not be a technical challenge because the data sets that reach analyzers will likely be small, even by today's standards.

There is one other topic I would like to mention. Whatever the new tools turnout to be, the transition plan should include efforts mentor senior people in the tools so that can continue to be effective mentors of analysis teams that use these tools. During the transition from FORTRAN to C++ the community did a bad job of this and the community lost a lot of physics expertise.

MINERvA Experiment
Data process/Analysis Workflow

Jiyeon Han (Univ. of Pittsburgh)
for Analysis eco-system workshop

MINERvA General Workflow

- *MINERvA uses Gaudi framework for the experiment*
  - *Gaudi framework was originally developed by LHCb*
  - *All of data/analysis workflow is managed by Gaudi tool*

- Data / MC workflow
  - Data and MC uses slightly different workflow :
    - MC requires simulation stage, overlay stage
    - Data requires calibration stage
  - Next two slides show the workflow of MC vs. Data

- *All of production and analysis ntuplizing procedure are done using FNAL grid(GPGrid) system*
  - *Plan to use OSG (open science grid) for production process this year*
  - *OSG usage will be extended to analysis users, but not near term*

MINERvA General Workflow

- Data file format :
  - Most of workflow use ROOT file format from calibration stage
  - Some of part (ex. MINOS stage and part of calibration) uses database (MySQL, MariaDB)
  - Ascii file format also is also used to cover parameter inputs

- Code manage/development is controlled by “cvs”
  - Plan to implement MINERvA software into “cvmfs” for OSG project

MC File Workflow

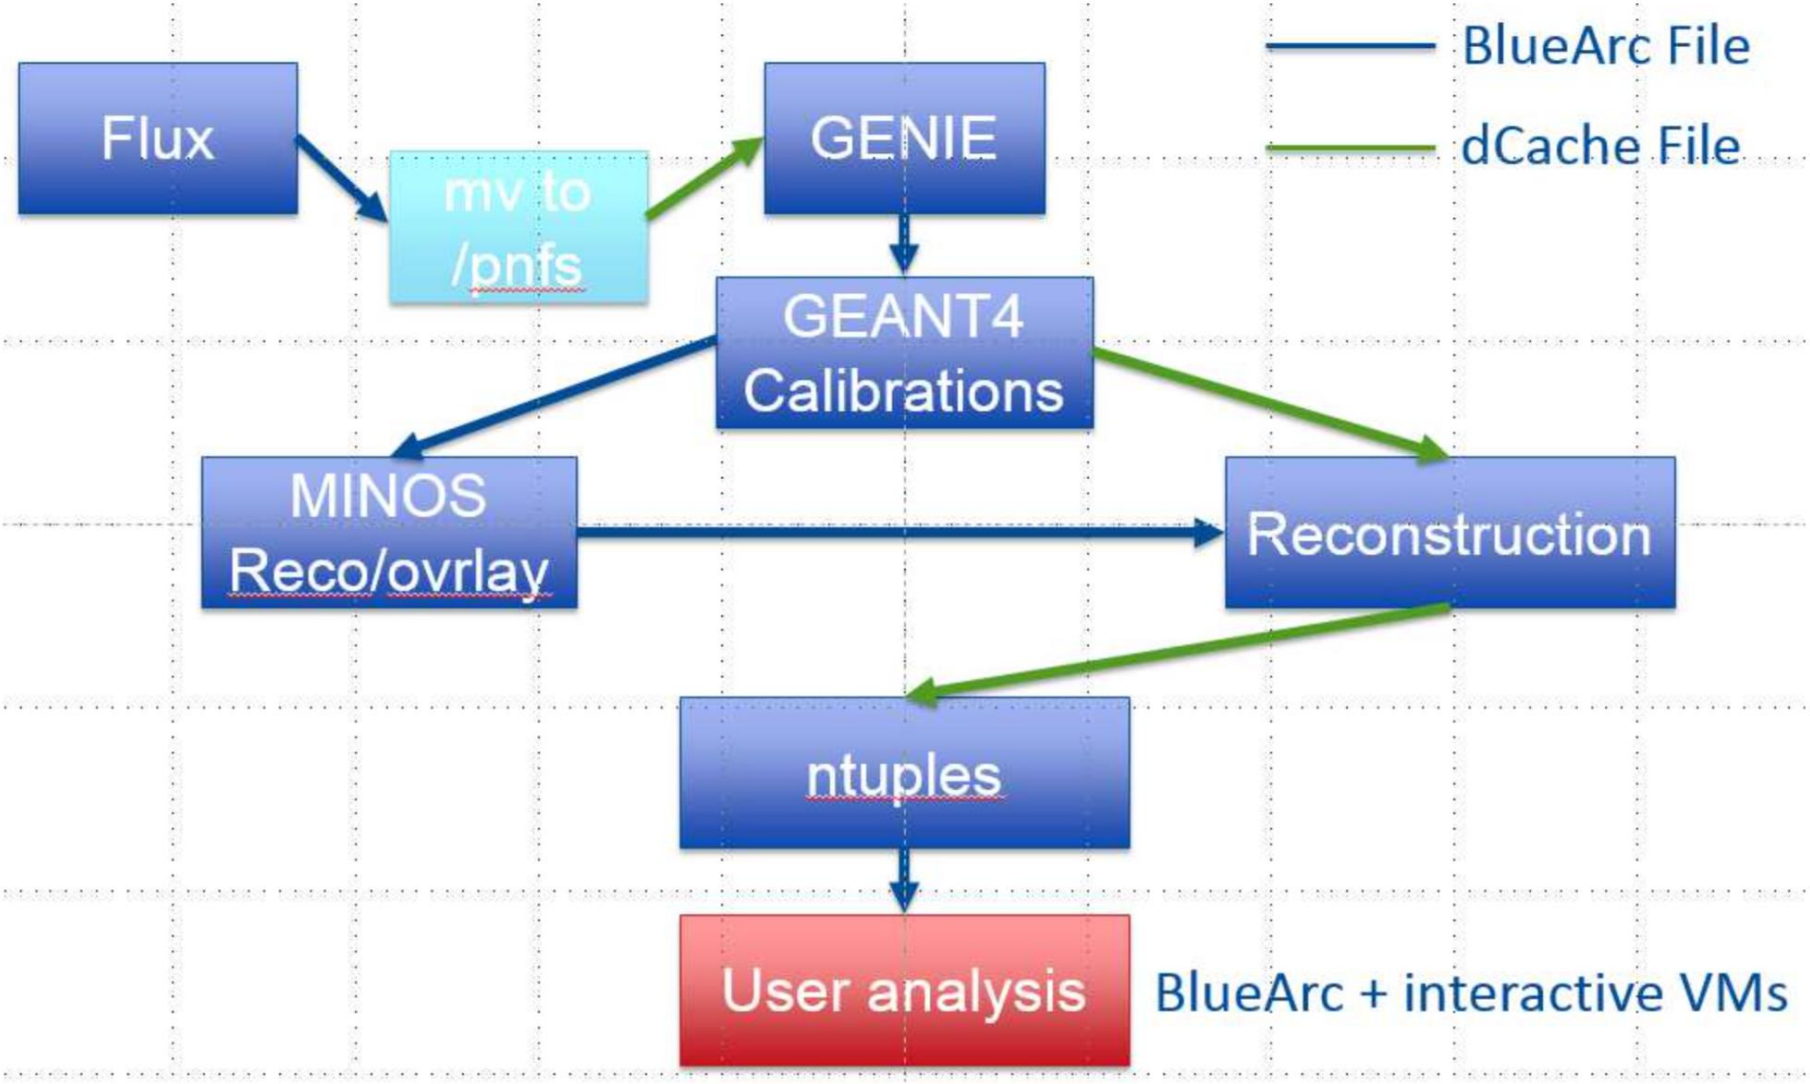

MINVERvA are reducing BlueArc usage since last year, so the output of workflow (ntuple) is copied into dCache area now

g-2 MIDAS is our DAQ system

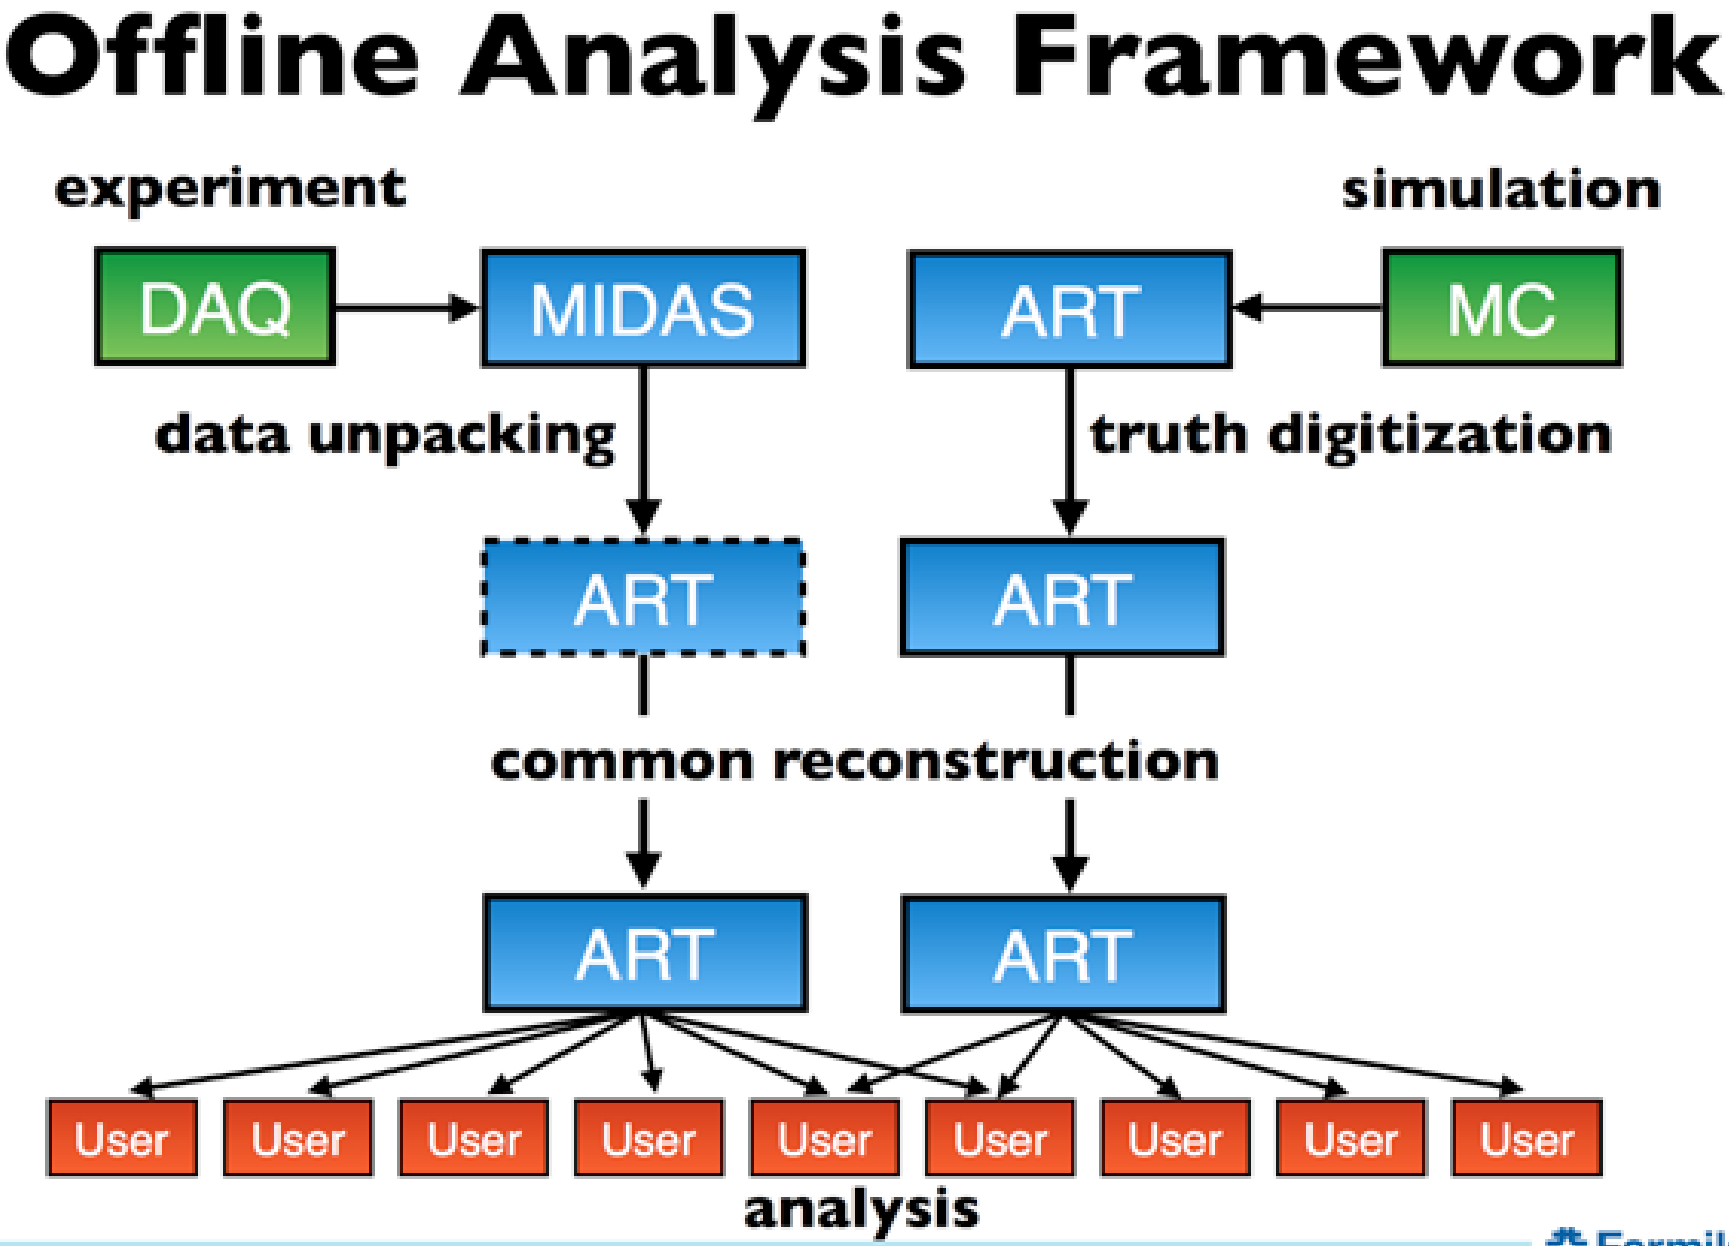

g-2 Calibrations

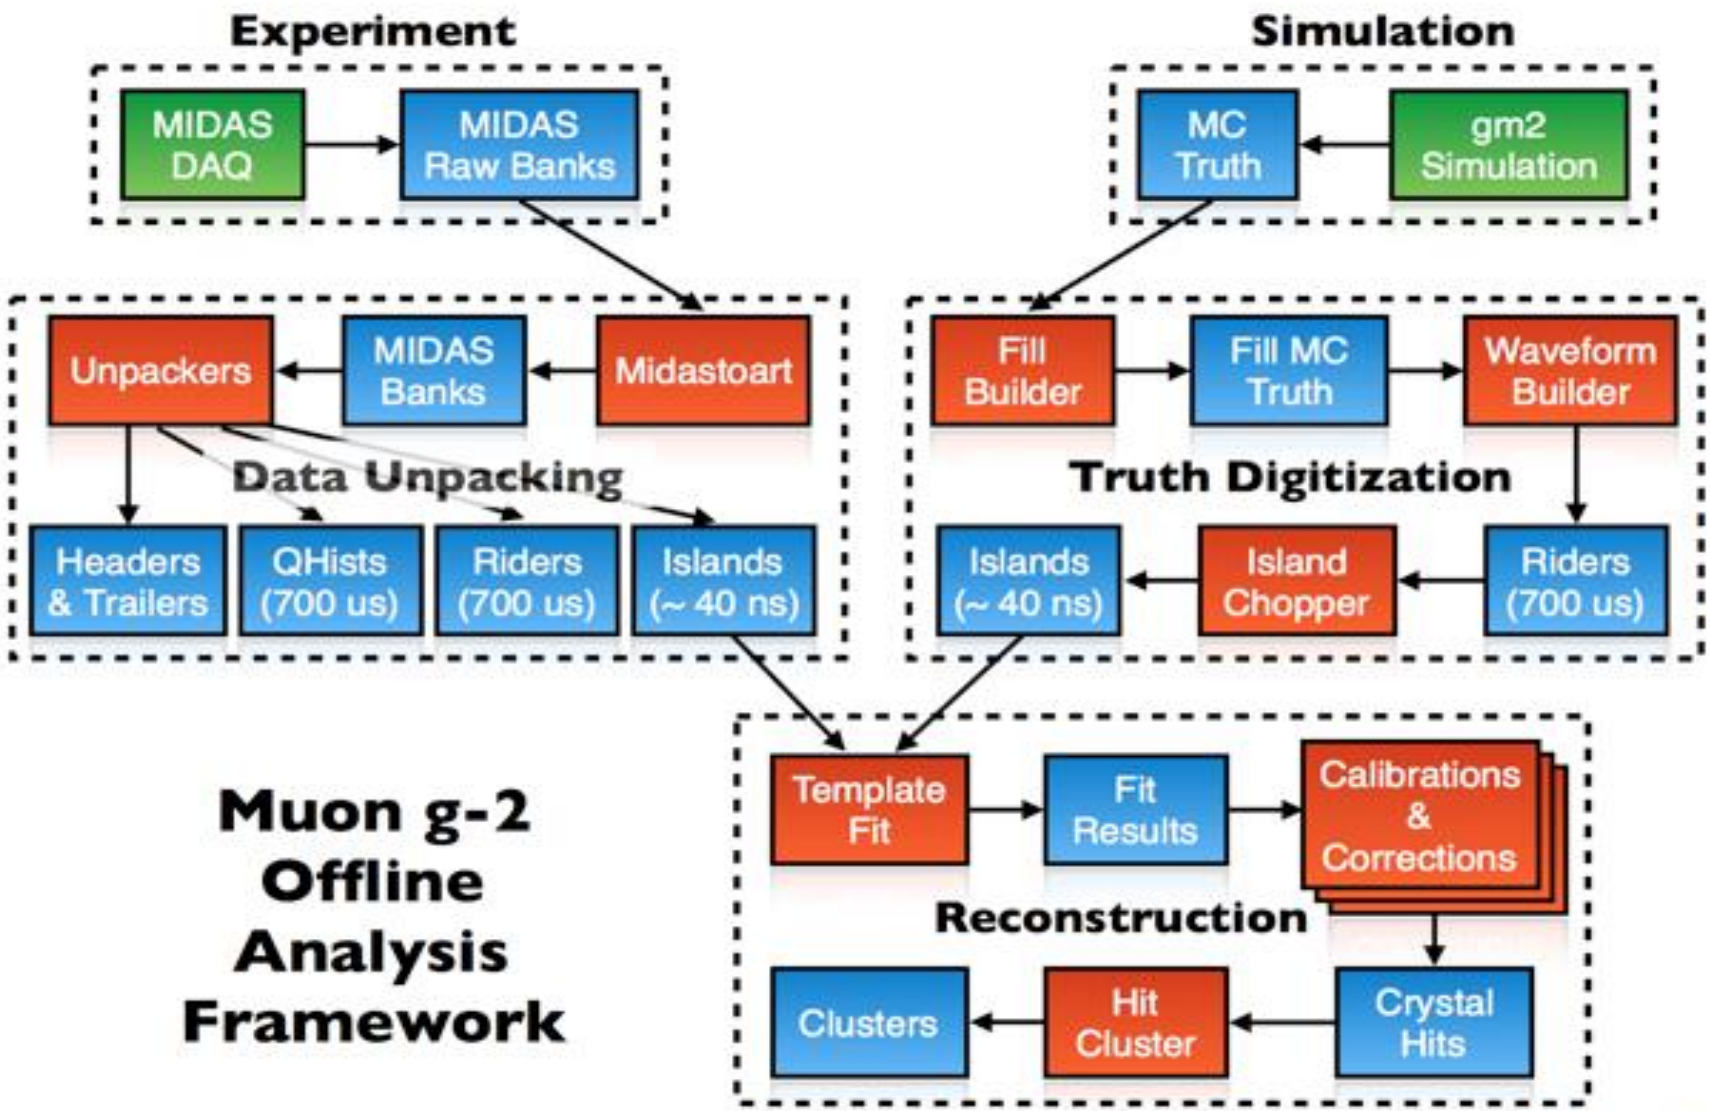

Software development workflow

We have several git repositories on the Fermilab Redmine system (plans to move to Github, but no time). We distribute libraries, executables and externals via CVMFS. We use Fermilab "ups" as our package manager, but may switch to Spack at some point.

Art is our framework.

g-2 Validation

- We run some unit and integration tests. We also run a physics verification package (produces many diagnostic plots).

g-2 Event data access

We have several analysis level workflows:

- *Art itself links against Root, so all of Root is in art. Can do analysis directly in art; though few do this.*
- *Art analyzer modules write out Root histograms and/or Trees. Then bare Root for analysis*
- *"Gallery" is a library that can loop over events in an art data file (art data files are Root files, but with a specific structure not easy to read with bare root). Gallery can be used within Root, Python (with PyRoot) and R.*
- *Convert data to HDF5 and analyze with Python or R (we aren't doing this yet)*
- Except for art in batch jobs, these tools do not scale to large data

g-2 Conditions

- We'll use a Postgres database. The art framework will have routines to read from the database. This is a work in progress.

g-2 Central common steps vs individual analysis

- *Analysis seems to be at the individual level at the moment. It may change as we ramp up and get more people*

g-2 Using the Grid

- *We mainly use the Grid for MC production and reconstruction. Not so much for analysis yet. We use the Fermilab FIFE grid tools. We use SAM for dataset and metadata handling.*

g-2 Improvements

- *Along with multithreading and vectorization, we would like to see tighter integration with python. The questions on the slides don't really apply to g-2.*

NOvA data flow

binary -> artroot -> reconstruction -> CAF -> dCAF

- *binary - is the raw data from the detector*
- *artroot - root representation of the data within the binary file*
- *reconstrution - output of the standard reco algorithms keeping everything*
- *CAF - root ntuples that contains objects in trees but without all the overhead of the dependency and associations within the reco output - these are the files that are used by most analysis groups that form the starting point of analysis datasets*
- *dCAF - derived CAF files that are skimmed and slimmed version of the CAF files within an analysis group for a specific channel or study.*

LArIAT data flow

binary -> artroot -> reconstruction

- *Those files have similar description as for NOvA, but there isn't a CAF file which is just a root ntuple.*
- *They do also use a numpy/scipy format file to do Machine Learning but it is then put back into a module that runs over the reconstruction files which are artroot*

MicroBooNE data flow

binary -> artroot -> reconstruction -> anatree

- *anatree - this is a flat root ntuple derived from the reconstruction artroot files. There is slimming and skimming done at this stage*
- *MicroBooNE also using the HDF5 format to do Machine Learning and Deep Learning image techniques based upon the images of the LArTPC response, these are translations of artroot raw data files.*

NOvA / LArIAT / MicroBooNE

- condition flow: not jet calibrations, but something like cluster calibrations is more appropriate
- Software development workflow: there are definitely private frameworks, but they are on the individual and group levels. Mostly for doing DL/ML algorithms
- Validation: is very experiment dependent but none are as robust as the LHC experiments right now
- Event data access: I don't know of anyone using non-HEP infrastructure, just software frameworks.
- Conditions and non event data: almost everyone in IF at Fermilab are using the SCD conditions database

NOvA / LArIAT / MicroBooNE

- *Usage of grid for analysis step: experiments don't have their own grids, just FIFE GPGrid and OSG, xrootd hasn't yet fully propagated to all the experiments, but is starting to take hold. But not really changed the way that people think about utilizing the grid, SAM4Users has made the SAM file catalog usable by individual users and has been really successful in NOvA and MicroBooNE*,
- *Improvements / Near Term Future: access to GPUs is the biggest limiter right now it seems. But I think that should be solved in the next 1-2 years. Access to data is going to be the biggest problem. Large datasets (MB/event is huge) and organization of the datasets for skimming/slimming etc.*

FIFE: Data and Job volumes

- Nearly 5 PB new data catalogued over past 6 months
- Average throughput of 1.8 PB/wk through FNAL dCache
- Typically 15K simultaneous jobs running; peak over 30K
- **Combined numbers approaching scale of LHC experiments**

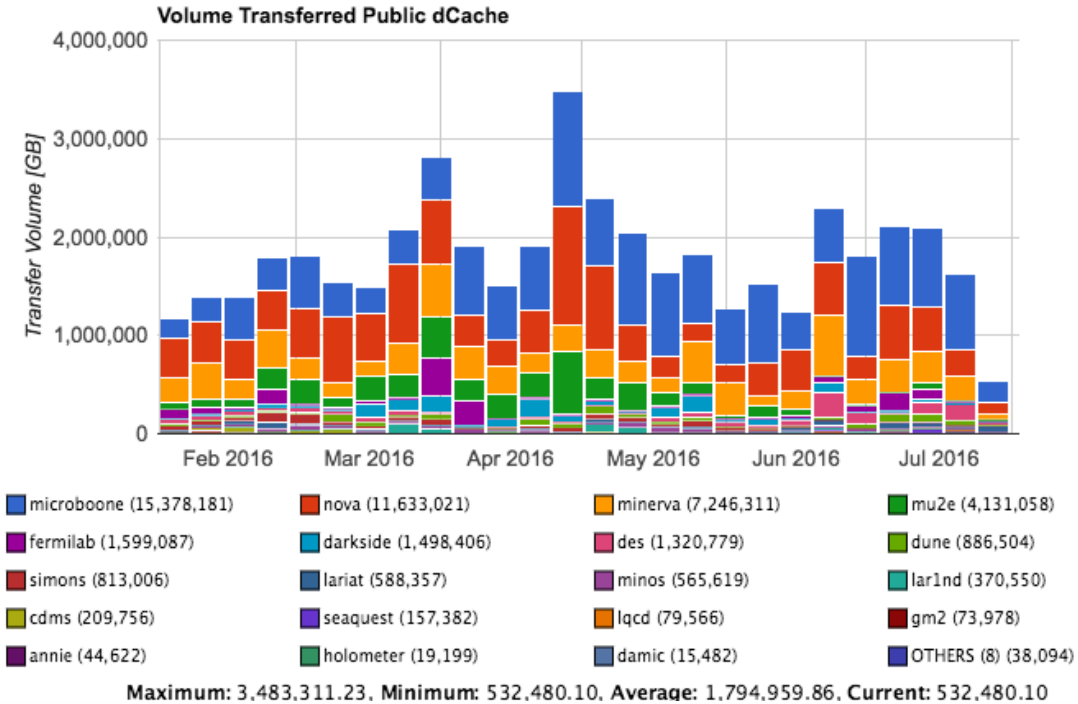

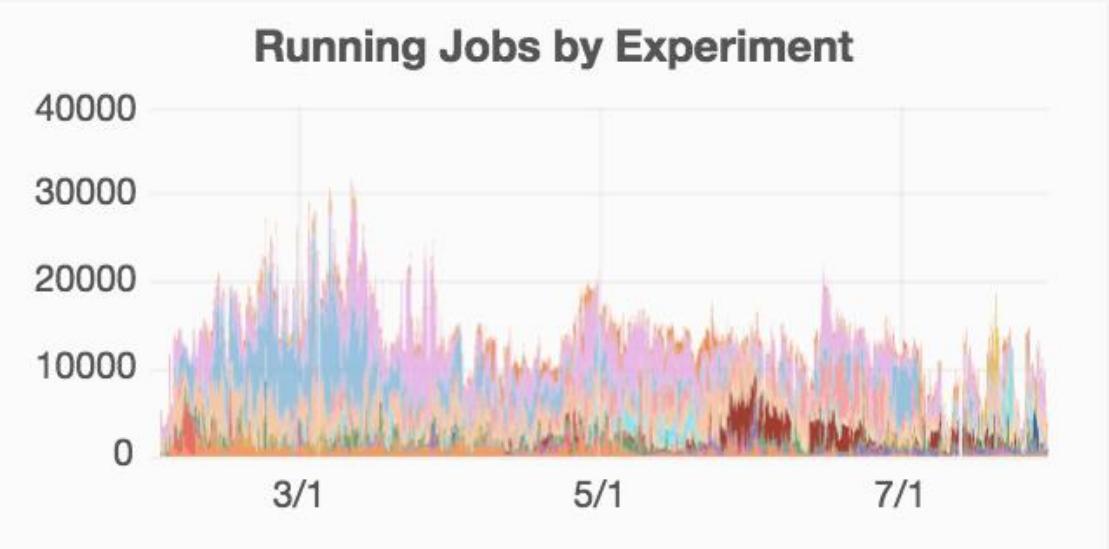

FIFE: A Wide Variety of Stakeholders

- At least one experiment in energy, intensity, and cosmic frontiers, studying all physics drivers from the P5 report, uses some or all of the FIFE tools
- Experiments range from those built in 1980s to fresh proposals

LArIAT