

DUNE FIFE Report

Tom Junk
FIFE Workshop
June 1, 2015

Outline

- Organization
- Workflows
- Schedules and plans
- Batch Systems
- Data Handling
- DUNE-ification
- Things that work well
- Things that don't work so well

The LBNE Physics Tools and Software and Computing Org Chart as of May 2014

A lot has changed

1) Turnover

2) New Collaboration

→ New working groups have yet to be set up

→ We're working with the old ones, but with some new names in a few boxes.

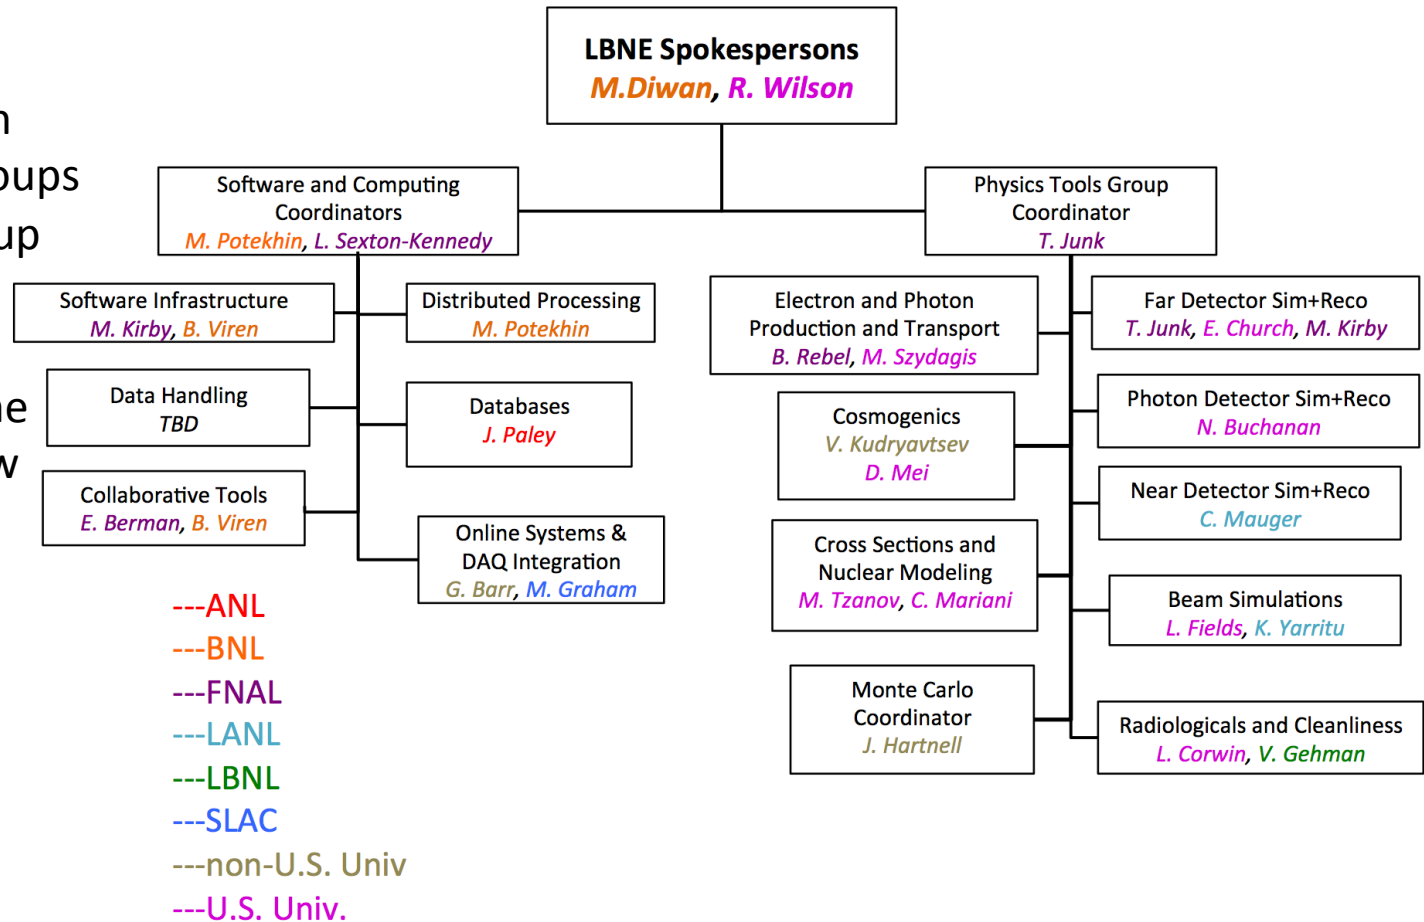


Chart from
L. Sexton-Kennedy,
May 2014 DOE Review

The LBNE Physics Tools and Software and Computing Org Chart as of May 2014

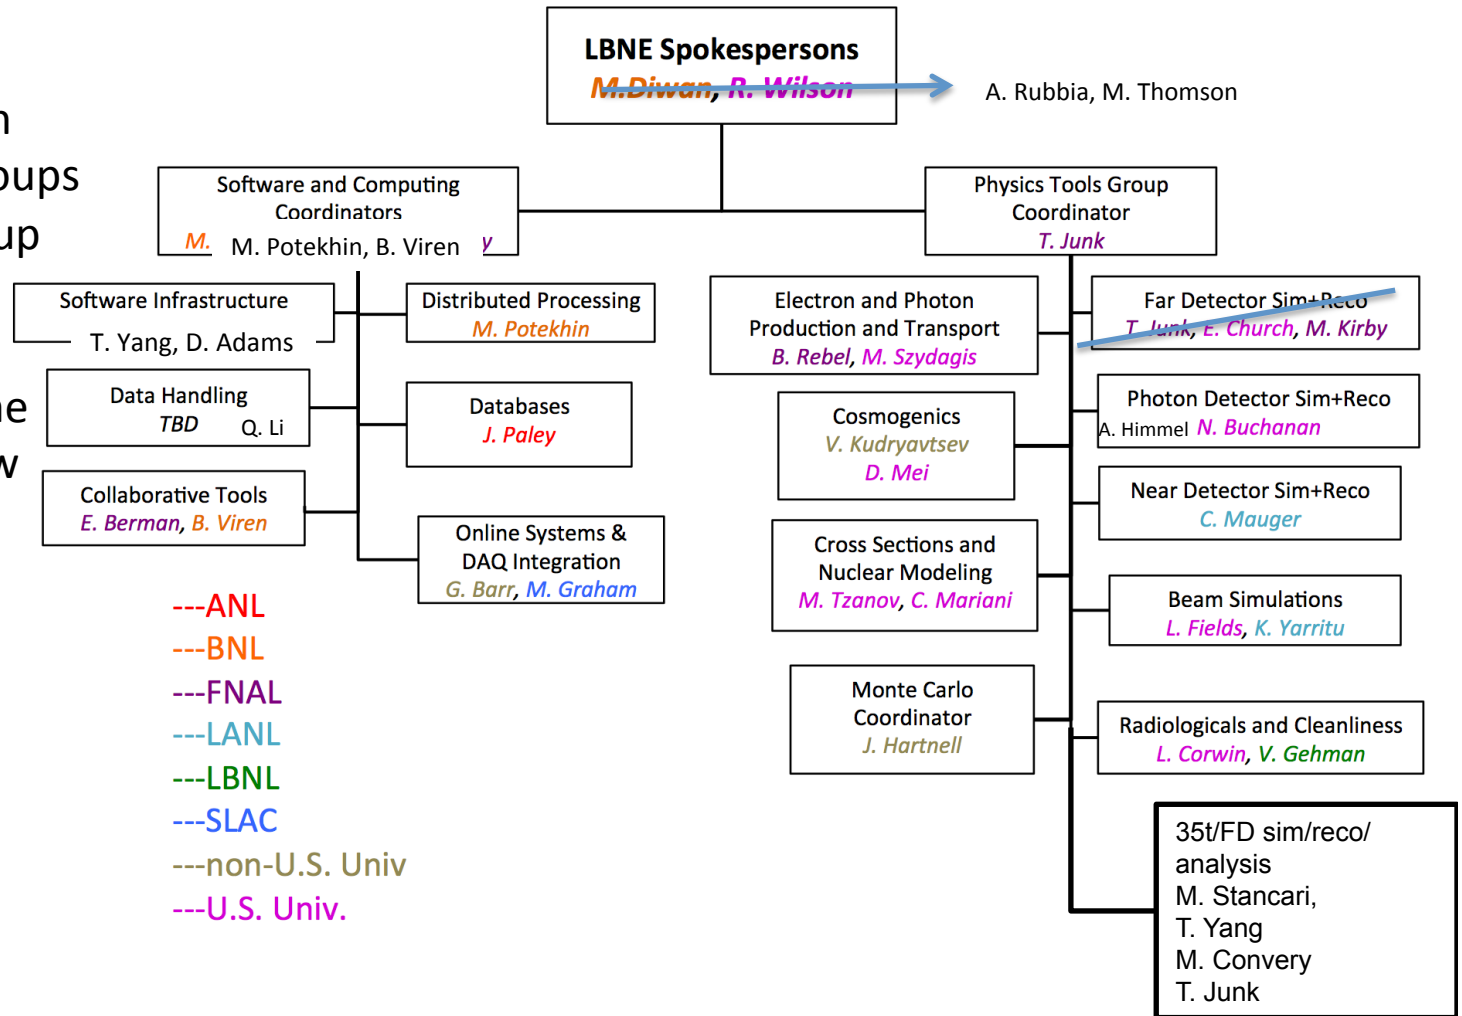
A lot has changed

1) Turnover

2) New Collaboration

→ New working groups have yet to be set up

→ We're working with the old ones, but with some new names in a few boxes.



Some but not all updates applied.

DUNE Organization

- Two new spokespersons:
 - André Rubbia (ETH Zürich)
 - Mark Thomson (Cambridge)
- Executive Board
- Institutional Board
- Technical Coordinator: Eric James
- Resource Coordinator: Chang Kee Jung

Need: Technical Board to review design proposals

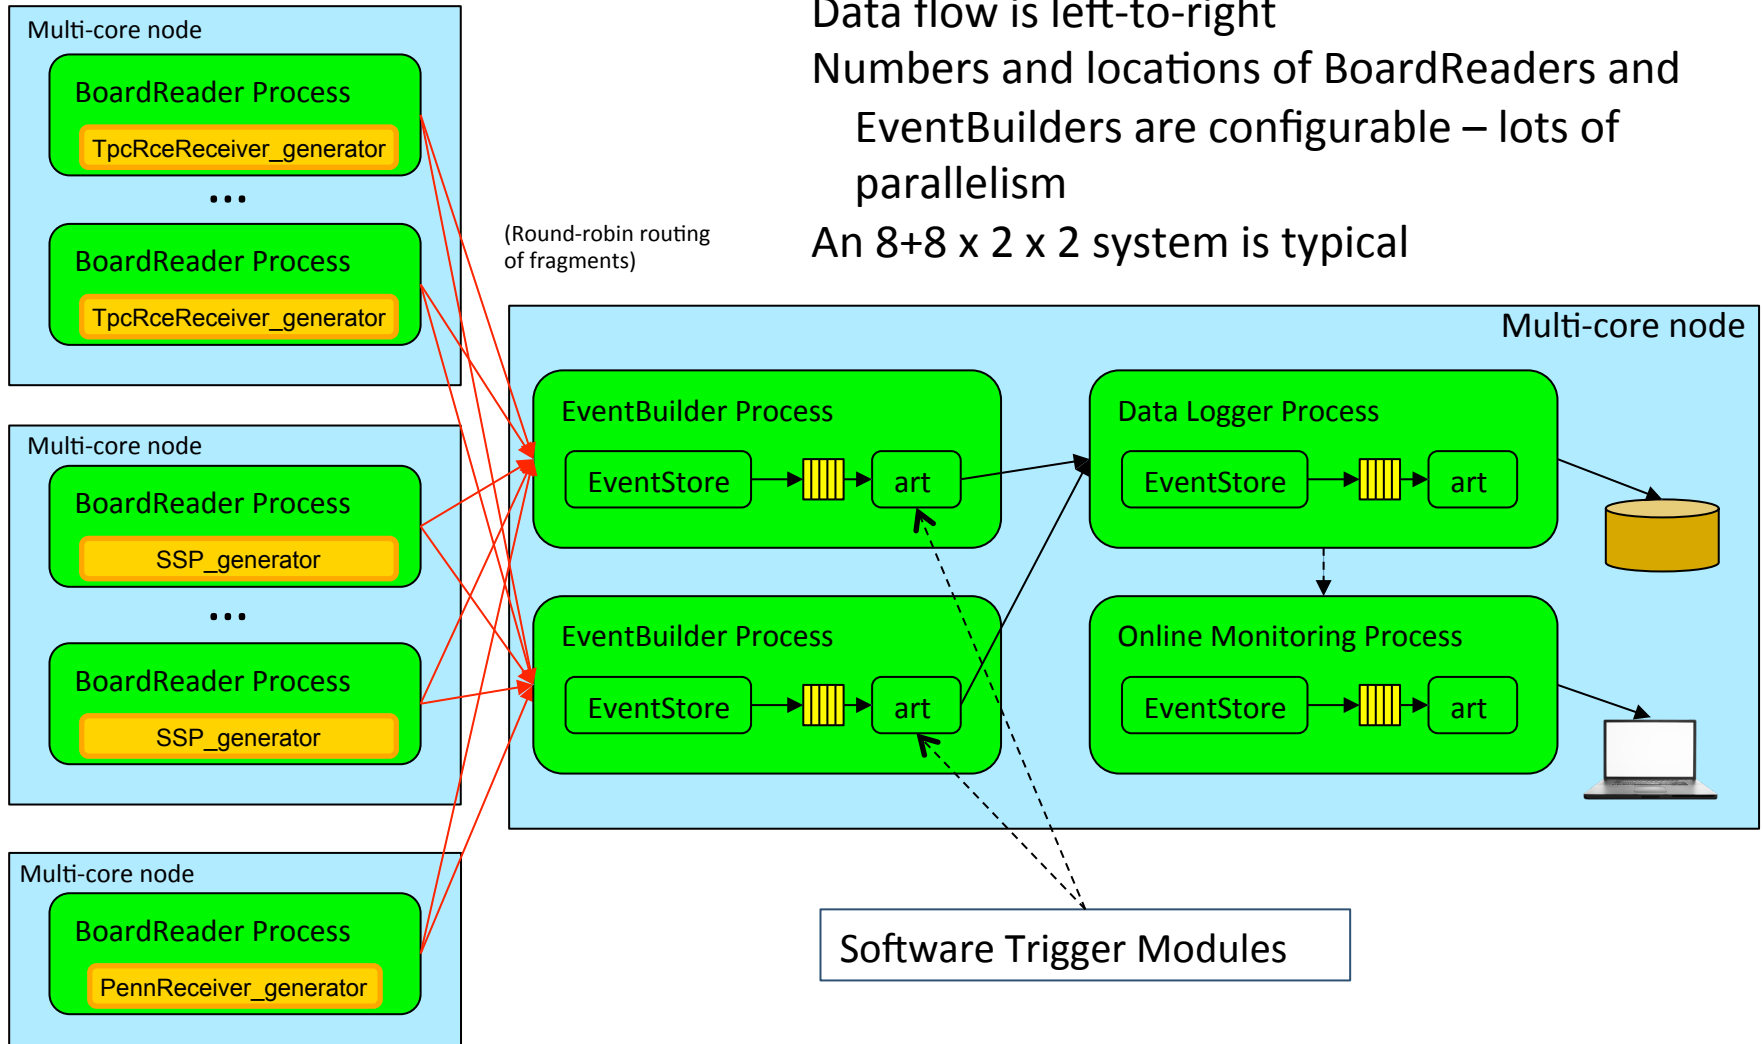
DUNE CD-1 Refresh Task Forces

- **Far Detector:**
Jim Stewart, Marco Zito
- **Near Detector and ND Physics:**
Raj Gandhi, Roberto Petti, Sanjib Mishra
- **Long-Baseline Oscillation Physics:**
Lisa Whitehead, Thomas Patzak
- **Nucleon Decay:**
Jon Urheim, Josh Klein
- **Atmospheric Neutrinos:**
Alessandra Tonazzo, Hugh Gallagher
- **SNB Physics:**
Kate Scholberg, Ines Gil-Botella
- **Systematics (LBL):**
Dan Cherdak, Elisabeth Worcester
- **Beam Requirements (for LBL physics)**
Alberto Marchionni, Laura Fields
- **DAQ/Online**
Giles Barr, Jacques Marteau
- **Reconstruction:**
Andy Blake, Tom Junk
- **Computing Infrastructure**
Brett Viren, Maxim Potekhin

35-ton Software and Workflows

- Online: *artdaq* is used as the framework for the board readers, event builders, aggregators, and online monitoring.
- Online monitoring incorporated into *artdaq* processes. Provides DQM modules with events without slowing down the DAQ.
- Convenient to use *artdaq* online and *art* offline as the online monitoring code more easily ports from one environment to the other.
- Data handling to use FTS online. Dropboxes already set up. Needs work to set up and check
 - Metadata extraction, file transfer to dropboxes
 - validation that transfer was successful
 - file cleanup

35-ton *artdaq* Architecture

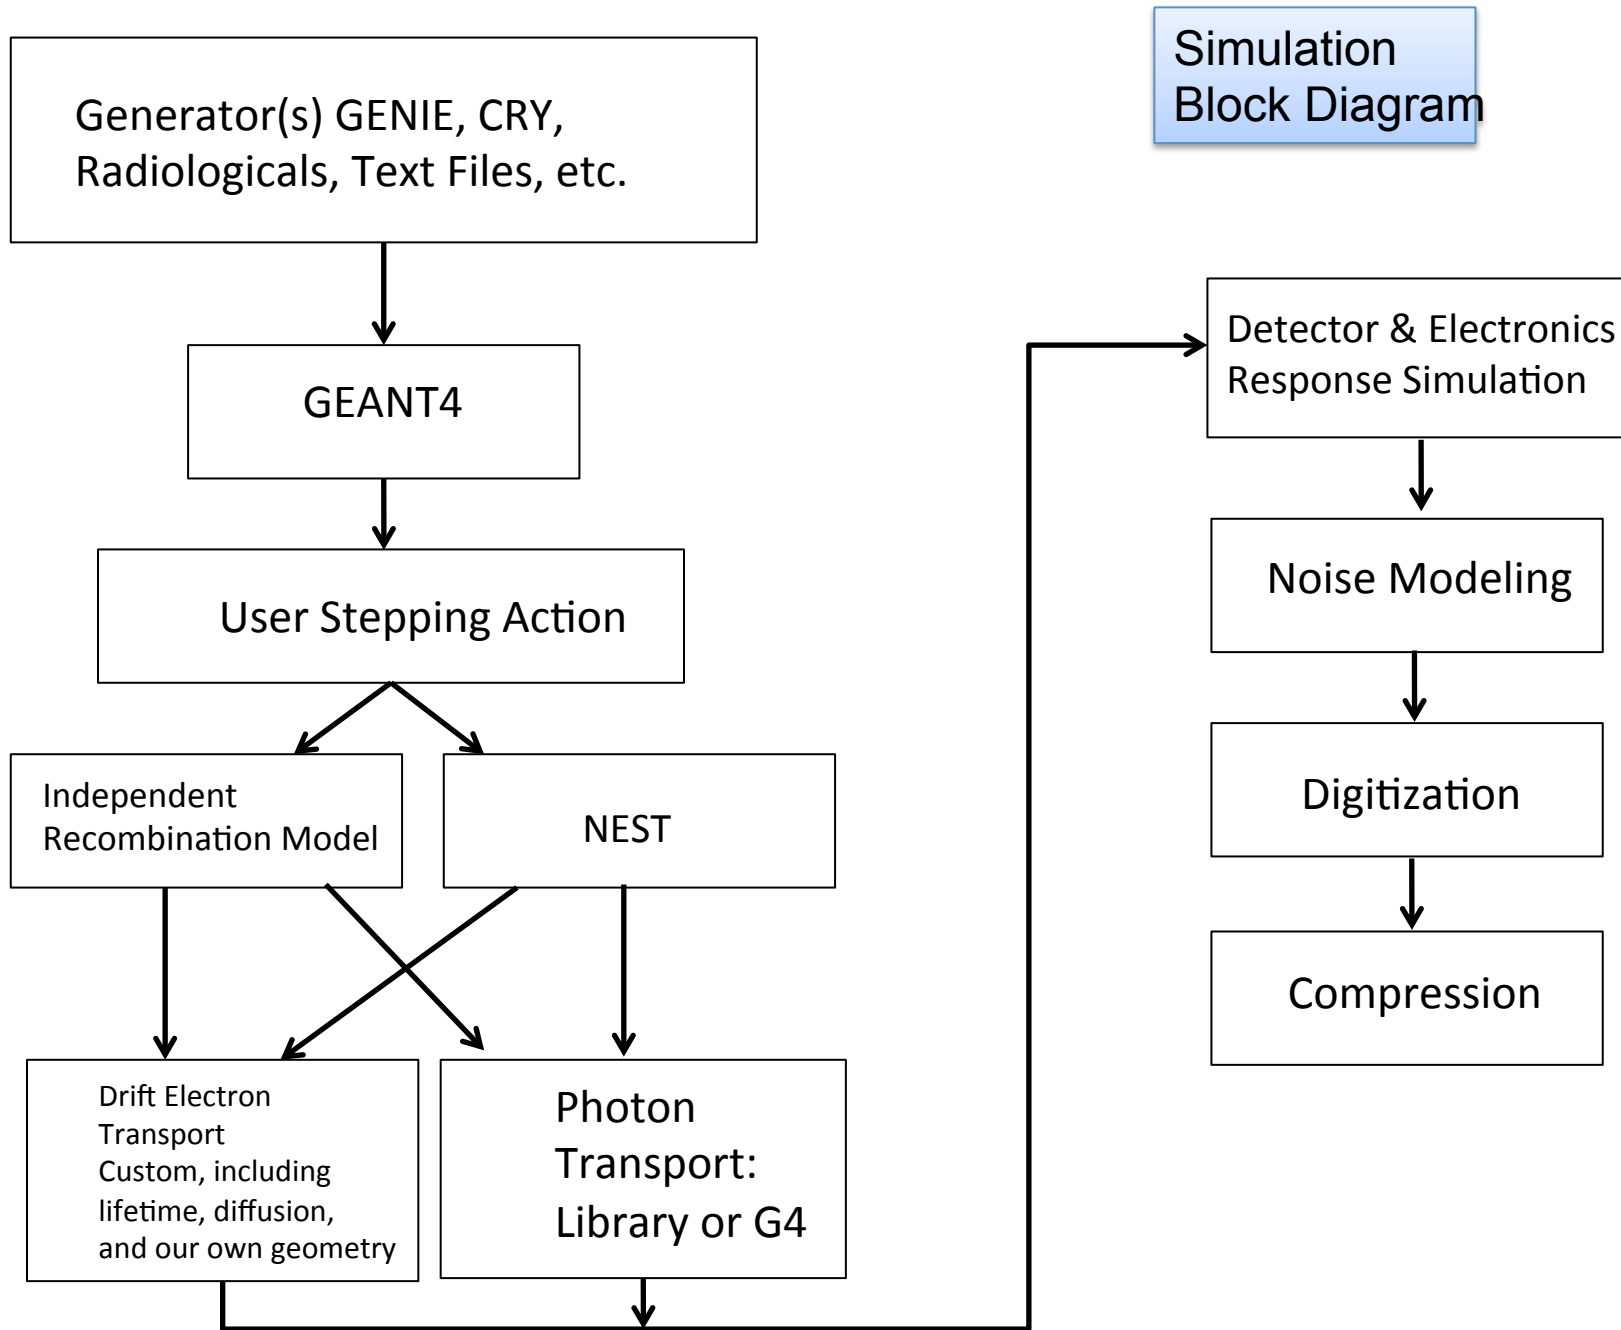


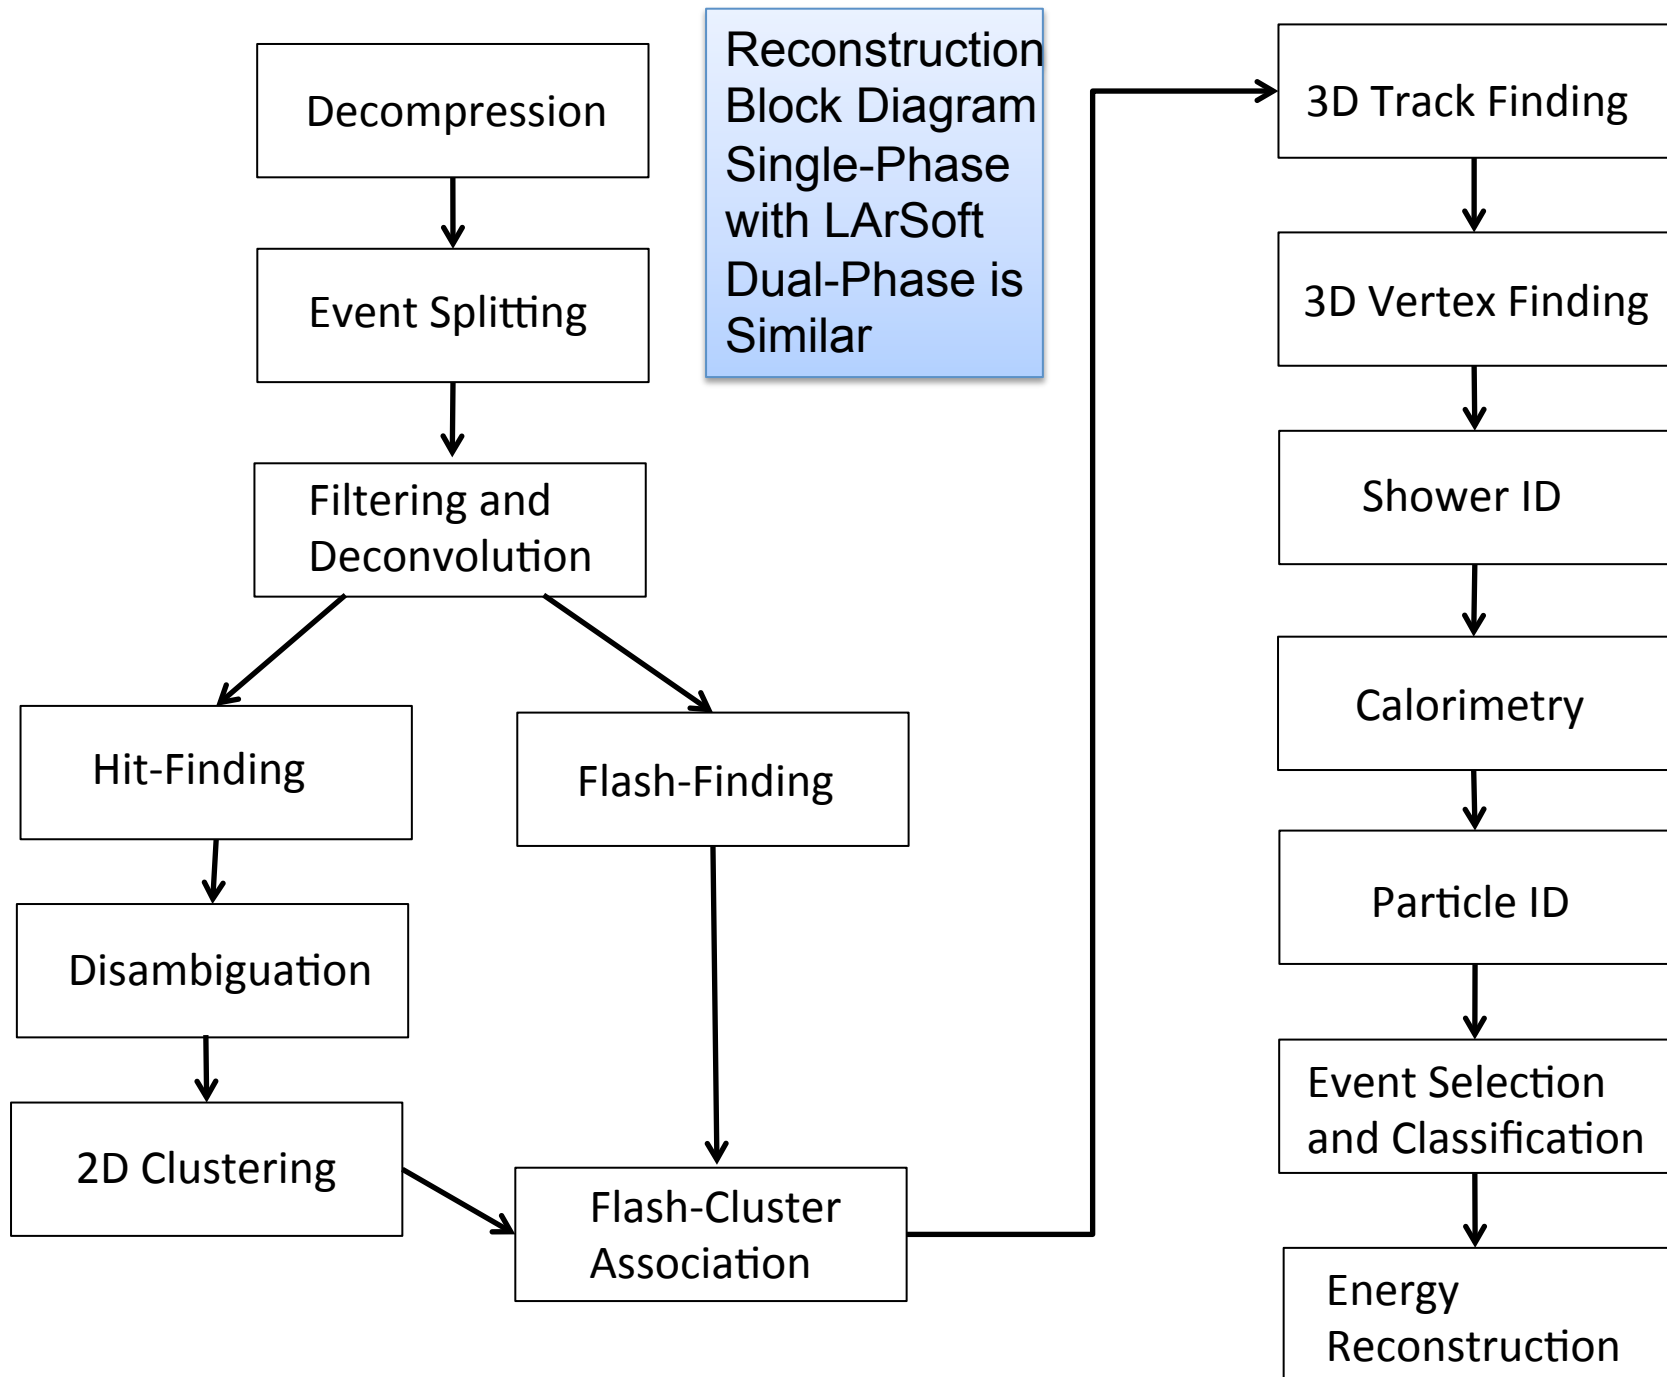
35-ton Offline Software and Workflows

- We use the *art*-based LArSoft suite of liquid-argon simulation and reconstruction tools
- Sharing of code with MicroBooNE and ArgoNeuT has sped the development immensely
- LBNE, and now DUNE, has a large base of collaborators, but most are working at small fractions of their time.
This is changing! As data-taking approaches, we are getting more FTE's engaged in software development.
- Much of the effort is in adapting code developed elsewhere for 35-ton and general DUNE use
PANDORA, ICARUS software, LArSoft software optimized for single-TPC detectors with non-wrapped wires. Adapting `calwire_ROI` code to the version we had split off for 35t caldata

35-ton Offline Software Workflows

- Other effort directed at 35-ton (or DUNE) specific issues: wire ambiguity resolution, 35-ton event display
- Generally useful tools as well
 - stopped muon analysis
 - π^0 selection (online trigger, offline analysis)
 - clustering tools
 - calorimetry
 - photon detector reco tools
- Most effort is development with little production needed.
- Two MC Challenges already this year (10K events each)
project.py – adapted from MicroBooNE is very convenient for this purpose.





CERN Single-Phase Prototype

- Just getting started, though the schedule is tight. Supposed to take data in 2018 for a few months.
- Proposal ready. Studies using LArSoft already performed, but no dedicated software yet.
- Need repository, geometry, detector-specific modules
 - geometry choice changed at the April Collaboration meeting.
 - also allowing flexibility in case it needs to be reconfigured for the maximum drift length.
- Would like data to be stored at CERN but also at collaborating institutions. Fermilab is also an obvious place to store these data. Of order 1 petabyte of data are anticipated.

Far Detector Design

- CD-1R FD design review May 19, 20
- Director's review – June 2-4
- Recommendations from the design review panel
 - Use GEANT-based simulations and automated reconstruction to inform design choices, such as wire pitch, angle, maximum drift length, wrapping, APA placement, and zero-suppression choices.
 - Goal is to tie the design choices to physics sensitivity
 - Currently done with Fast MC – also a FIFE consumer

Beam Simulations

- Workflow – self-contained GEANT4-based program runs and makes root trees. Uses: batch system, dCache
- Many repetitions for different designs of targets and horns.
- Interface to FastMC – changing flux requires reruns of FastMC to compute sensitivity
- Laura Fields has parameterized the FastMC response to make it even faster so that a genetic algorithm can sample thousands of flux distributions and pick the best one.
- Output storage space has limited the group's processing – ramping up usage of dCache.
- To do soon – explore running jobs offsite

MARS Experience

- Simulation of beamline, target chase, shielding, decay pipe, absorber.
- Computes radiation dose, temperatures, and beam flux.
- Has its own representation of the geometry
- Large amounts of CPU (4M hours/year), small amounts of storage – BlueArc is fine for storing results (8 TBytes requested). Probably can use dCache as well

DUNE-ification

Many (all!) of DUNE's FIFE resources are labeled "lbne" – need new names at the very least, and new instances of some of the resources. The ball is rolling to get this to happen.

- Collaboration affiliation on the ID/account request page
- Collaboration affiliation on the Service Desk web interface
- gpvm interactive node names
- the Virtual Organization (LBNE's was its own VO not part of the Fermilab VO, and DUNE should be the same)
- BlueArc mount point names /lbne/data, /lbne/data2, /lbne/app, /grid/fermiapp/lbne, etc.
- FIFEMON monitoring group
- Code repositories and Redmine pages
- DocDB (we already have a new DUNE DocDB but aren't using it: docs.dunescience.org)
- Collaboration web page: old: lbne.fnal.gov, new: dunescience.org

What Works Well

- Large grid CPU resource. Users are very happy to find they can submit thousands of jobs and they run.
- Large dCache disk resource – haven't run into problems with the space
- BlueArc is convenient for interactive work – developing and local storage for running interactive root jobs to make plots
- No one has complained about unmounting /lbne/data and /lbne/data2 on the grid worker nodes.
- Remote builds and code distribution now works
 - scisoft.fnal.gov
 - CVMFS
 - worch

What Works Well

- Much faster MRB builds! Thanks so much!
- Help from art team in writing an event splitter/stitcher input source.
- I tried some time ago to use XRootD to access files in dCache (test was onsite). Some DUNE users are interested in using XRootD to distribute data.
- BlueArc Quota overage e-mails now go to the users directly.
- SCD has proactively helped the beam simulations team with storage and file transfer.

What Doesn't Work So Well

- FIFEMON
 - Missing data on plots
 - Slow response
 - History graphs connect lines between unequally spaced data samples that can misleadingly be interpreted as batch usage when in fact there was none
- BlueArc Quotas and Shared accounts (lbneana)
If multiple users transfer files back to BlueArc from batch jobs with lbneana ownership (not using gridftp), then lbneana's quota can fill up and multiple users lose their job output.
Solution: use gridftp, dCache. Not defaults.

What Doesn't Work So Well

- no quotas on /lbne/app. Smaller disk than /lbne/data. Some users have used /lbne/app for data files. Takes some human effort to track these users down and get them to move or delete their files. Often big usage on /lbne/app is unintentional.
- Offsite users sometimes have problems with authentication
 - Kerberos and firewalls (only one user so far and that was some time ago)
 - KCA certificates. get-cert runs on some platforms and users struggle with others – kx509 is sometimes difficult to use on non-FNAL supported platforms.Can be a barrier to the annual VO Acceptable Use Policy signing.

What Doesn't Work So Well

- gridftp timeouts can cause job output loss. Need gridftp option to ifdh cp to retain file ownership by the user. But a timeout results in failure with this method.
- The ifdh cp -d option gotcha
- Submitting thousands of jobs one by one with jobsub_submit can be slow (possibly user just needs to use -N)
- Deleting thousands of jobs is also tedious.
- Multiple users have requested an option to jobsub_rm that deletes all jobs owned by a specific user.
- One user asked for a test queue that starts short jobs right away for testing purposes. Can have strict limits – short time, maximum jobs running per user = 1, for example.

What Doesn't Work So Well

- Some power users have mentioned that I/O speed using ifdh cp to BlueArc results in low job efficiency and long wait times.
- A user has asked for files that fail to transfer back from a batch job due to exceeding quotas (or other reason) to go to a temporary safe spot so users have some time to recover the output instead of losing it.
- Some users found the documentation rather jargony and made assumptions of prior knowledge. Sometimes users find it hard to locate the the right documentation
- Users plead for backwards compatibility.
- New jobsub_client commands require many non-default options to be specified in order to work

Extras

artdaq features

