Long-term vision for LArSoft

ICARUS Perspective

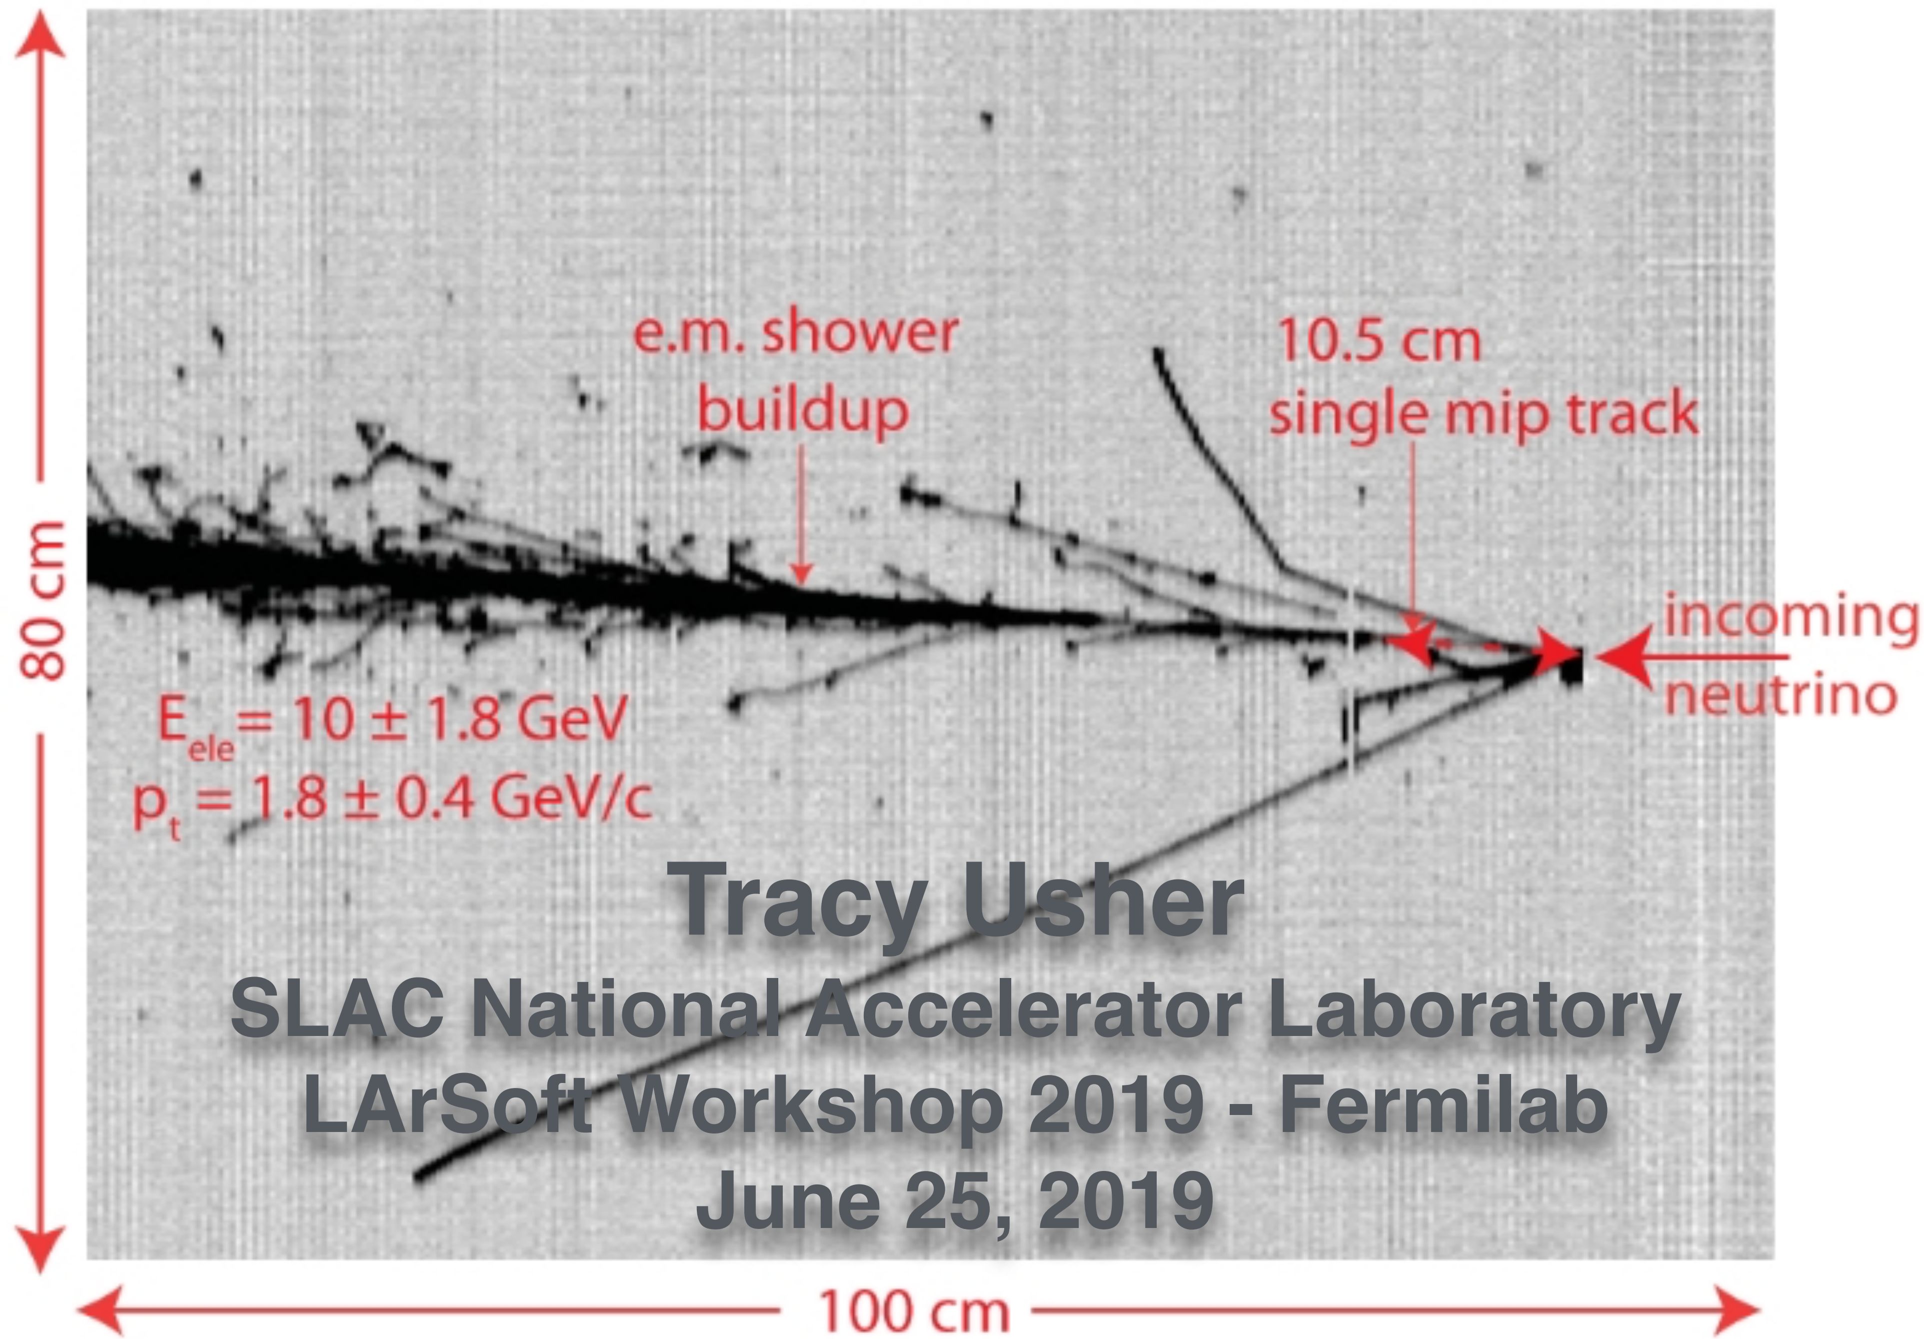

Tracy Usher

SLAC National Accelerator Laboratory

LArSoft Workshop 2019 - Fermilab

June 25, 2019

The General Plan

- Brief Introduction to ICARUS from the Software point of view
- Basic software components for simulation and reconstruction
- Current workflow
- Opportunities for improved computing techniques

- Aim to point out some of the areas where ICARUS is already stressing the “standard” implementation of LArSoft based simulation and reconstruction
    - Will primarily focus on the TPC simulation/reconstruction

ICARUS

- ICARUS is an acronym for:
  "Imaging Cosmic And Rare Underground Signals"
  - Originally conceived to do as the name suggested…
  - Three year physics run performed at LNGS starting in 2010
    - Neutrino beam from CERN's Neutrinos to Gran Sasso (CNGS)
  - The result of some 20+ years of development of Liquid Argon TPCs as high resolution particle imaging detectors
    - From ideas first presented by Carlo Rubbia in 1977*
- ICARUS is the first large scale Liquid Argon TPC
  - The precursor to the entire program based here at Fermilab
- ICARUS has moved to Fermilab to be part of the Short Baseline Program
  - Completing the installation now, expect commissioning in the Fall

ICARUS Snapshot

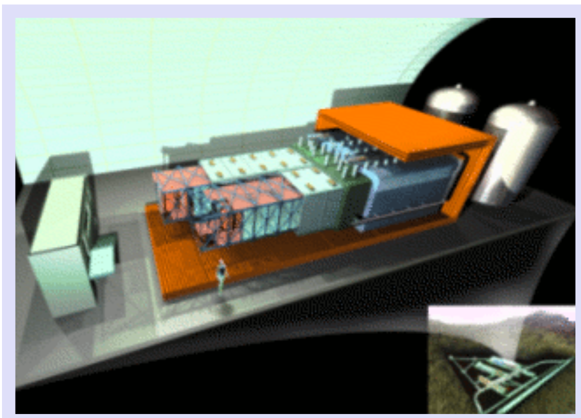

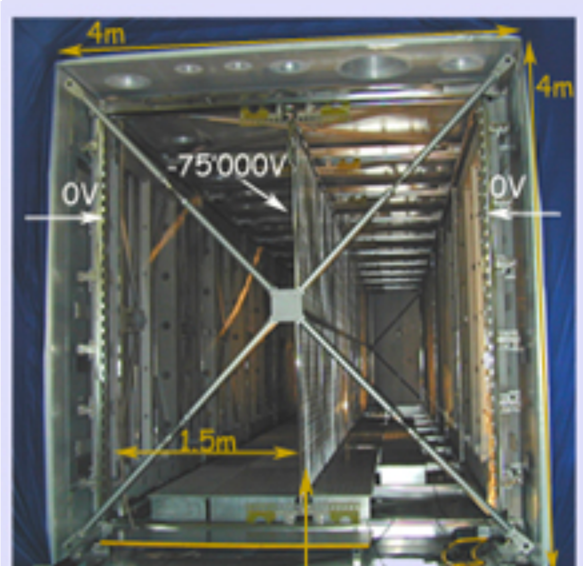

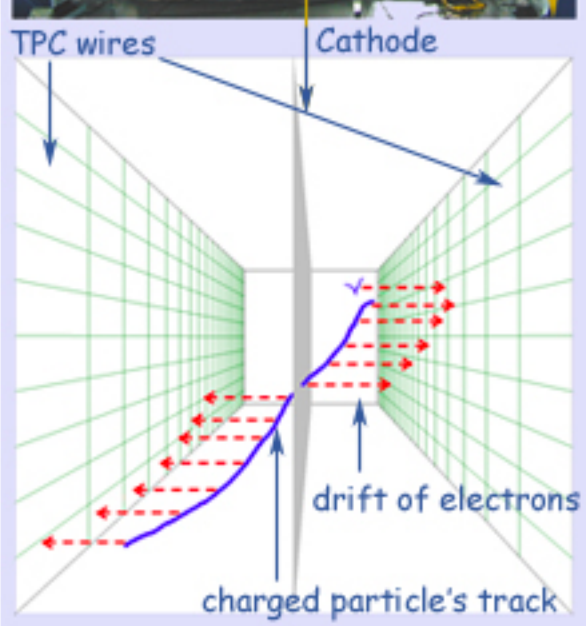

- ICARUS (T600) consists of two modules each placed in its own cryostat (2 x T300)
  - Each module contains two TPCs separated by a common cathode
  - Each TPC has three readout planes but in a different order from “normal”:
    - First induction - horizontal (Y) wires split in the middle
    - Middle induction - U wires (+60^{o})
    - Collection - V wires (-60^{o})
    - Active area is ~18 m x ~3.1 m
    - >13,000 wires/TPC - ~53,000 total
  - Each TPC also has 90 PMTs, 360 total
    - 500 MHz clock rate, ~ 2 ms readout window
- Also CRT system

SBN Detector Comparison

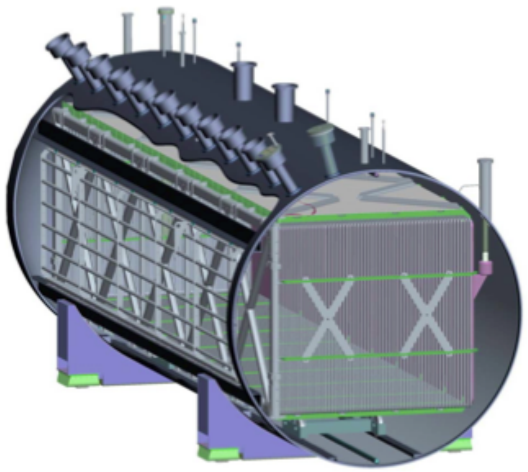

MicroBooNE

- 170 tons of LAr with ~87 ton active volume located 470m from the BNB source
- Single TPC - 8,256 wires in 3 sense planes oriented +60, -60 and 0 degrees with 3 mm pitch
- 36 phototubes located behind the anode plane
- CRT, cold electronics, 2.5 m drift, 273 V/cm

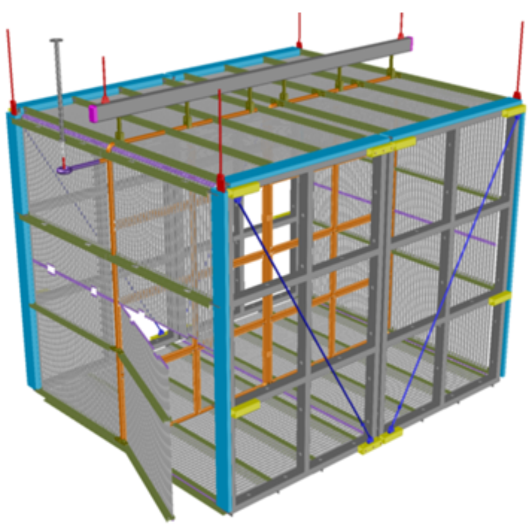

SBND

- 220 tons of LAr with ~112 ton active volume at 112m from the BNB source
- 2 TPC's in 1 Cryostat with 11,264 total sense wires
- Each TPC has 3 sense planes oriented at +60, -60, and 0 degrees with 3 mm pitch
- PMT's, ARAPUCA, TPB coated reflector foils and Light Guide Bars
- CRT, cold electronics, 2 m drift, 500 V/cm

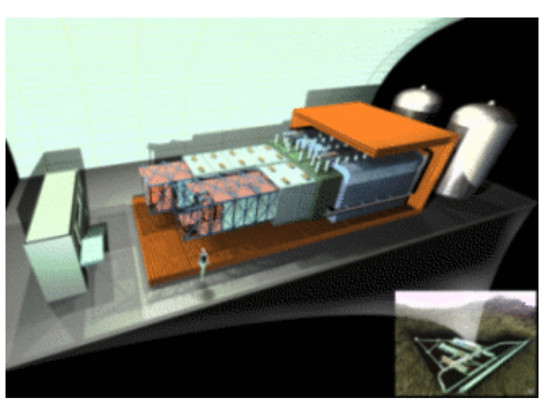

ICARUS

- 760 tons of LAr with ~470 ton active volume located 600 m from the BNB source
- 4 TPC's in 2 Cryostats with 53,248 total sense wires
- Each TPC has 3 sense planes oriented at +90, +60 and -60 degrees with 3 mm pitch
- 90 phototubes in each TPC located behind the anode plane
- CRT, warm electronics, 1.5 m drift, 500 V/cm

ICARUS requires a significant increase in computing resources

Simulation/Reconstruction

- MicroBooNE has pioneered automated event reconstruction for large scale Liquid Argon TPCs
  - Obviously, first step for ICARUS is to transfer that knowledge!
    - In other words, steal now and ask for forgiveness later!
- Follow the basic steps in LArSoft (as outlined by Erica)
  - Event generation <— standard code from LArSoft
  - Geant4 simulation of events <— standard code from LArSoft (…)
  - Detector simulation <— ICARUS specific implementations
  - Event reconstruction <— Now the fun really starts…
- As they say, the devil is in the details…
  - Particularly given the increased number of channels for the TPC

ICARUS Geometry

- Everyone says “my code will just work”
- To date everyone has been wrong
- Why?
    - ICARUS has horizontal wires, not vertical
        - It was originally optimized for detector Cosmic Rays
    - The first induction plane are the the horizontal wires.
        - The collection plane are “V” oriented wires.
    - The horizontal wires are split
        - To date we are treating the horizontal wires as continuous
        - Our best man is on this, hopefully we have a working geometry with split wires in the next few weeks
    - There are many pieces of code which treat “U”, “V” and “W” interchangeably with planes 0, 1, 2.
        - I think this is STILL true

TPC Detector Simulation

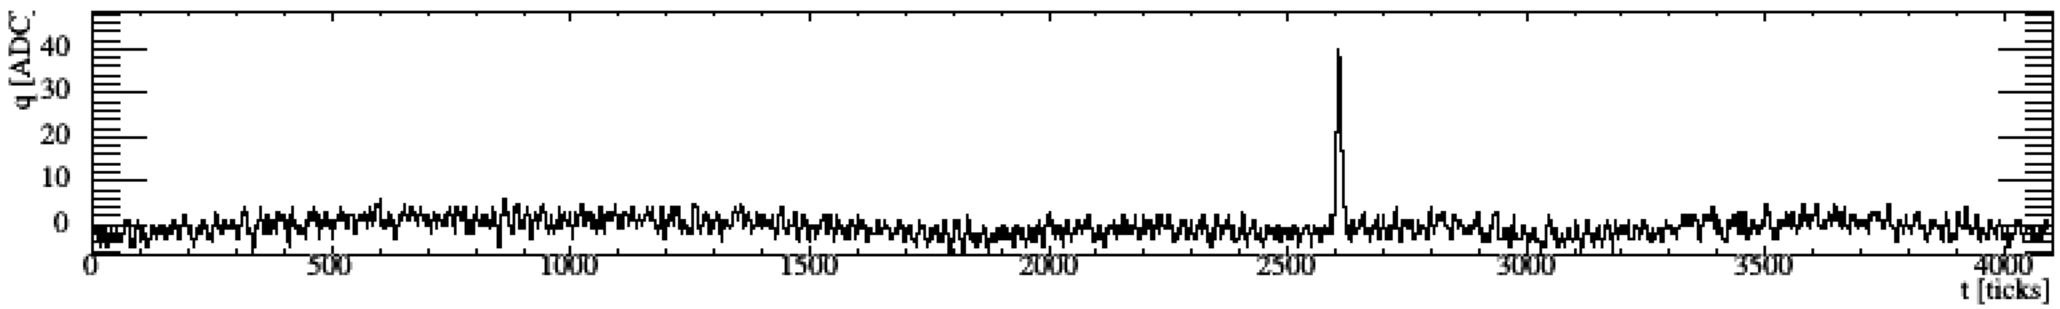

- Goal is to output waveform per wire
  - Currently use a 1D "TPC response" to simulate signal on wire
    - Input are the "SimChannels" from the Geant4 stage (drifted electrons)
    - 1D convolution of the TPC field response coupled with electronics response
  - Add noise to the channels using model based on previous experience
    - Incoherent component
      - Sampled from observed frequency response from previous running plus from testing of new electronics at CERN test facility, including significant low frequency oscillation
    - Coherent component
      - Common across 32 channels sharing the same motherboard (similar to what was seen in MicroBooNE
- Note: Currently only output channels on shared motherboards where at least one channel has seen a signal
  - Full detector simulation at this time is very time consuming

Waveform Comparison

Data from LNGS Running

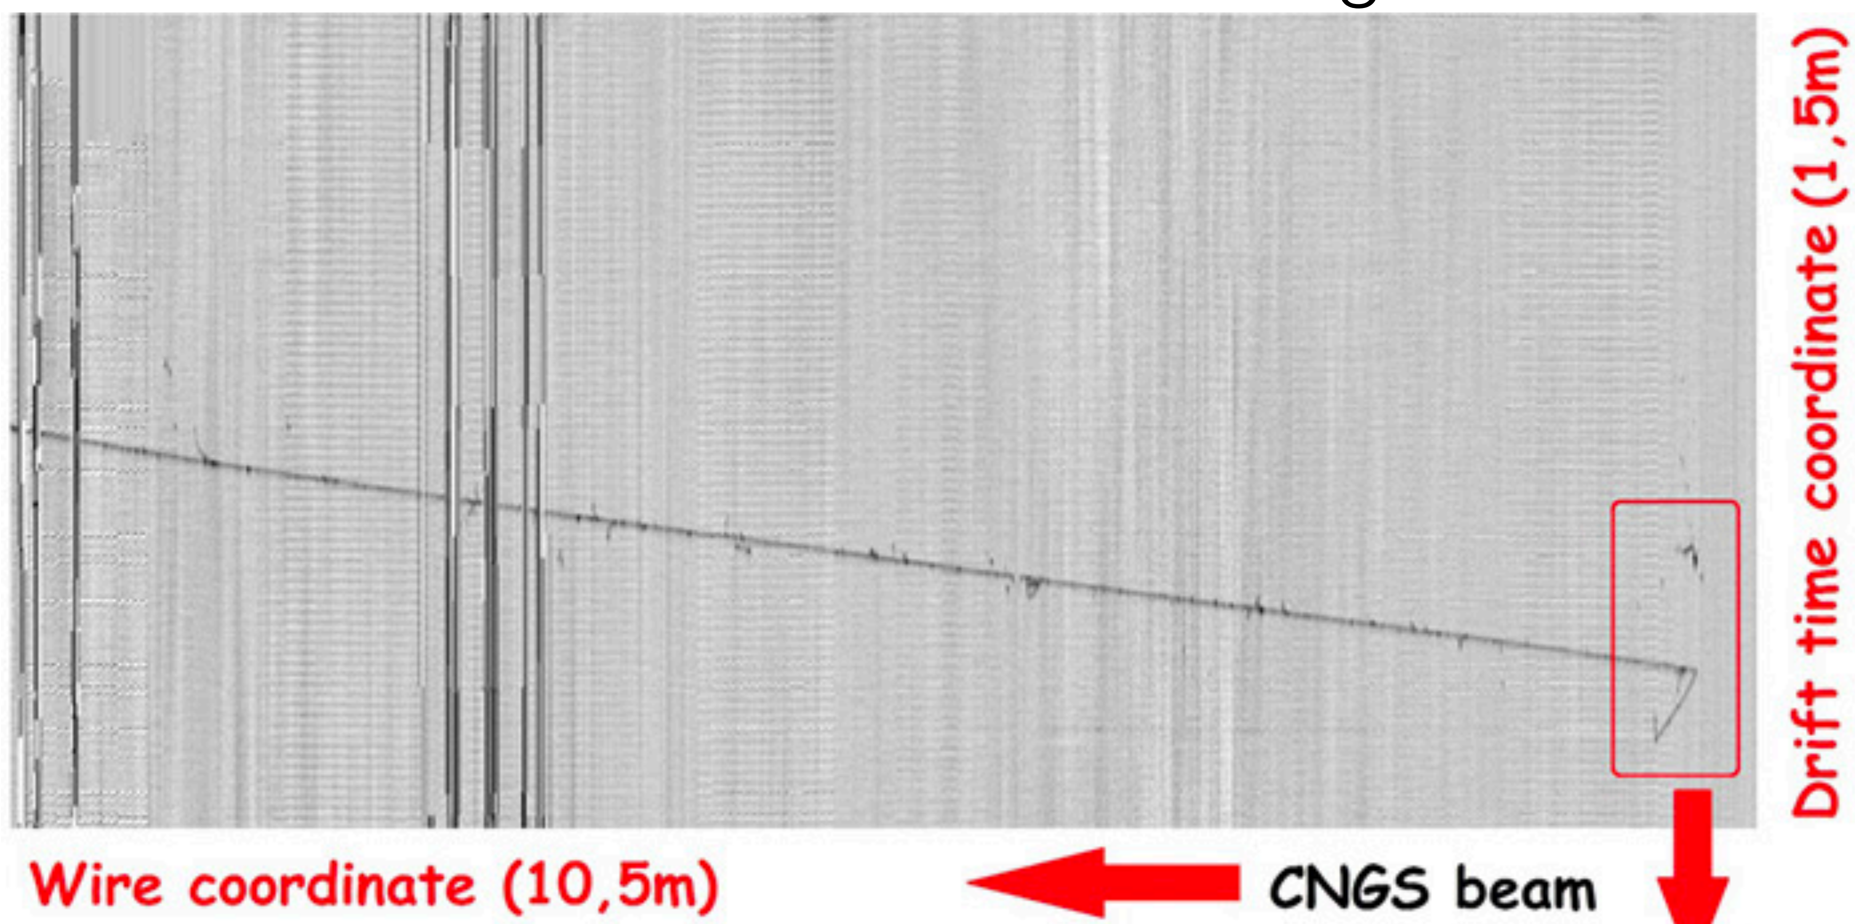

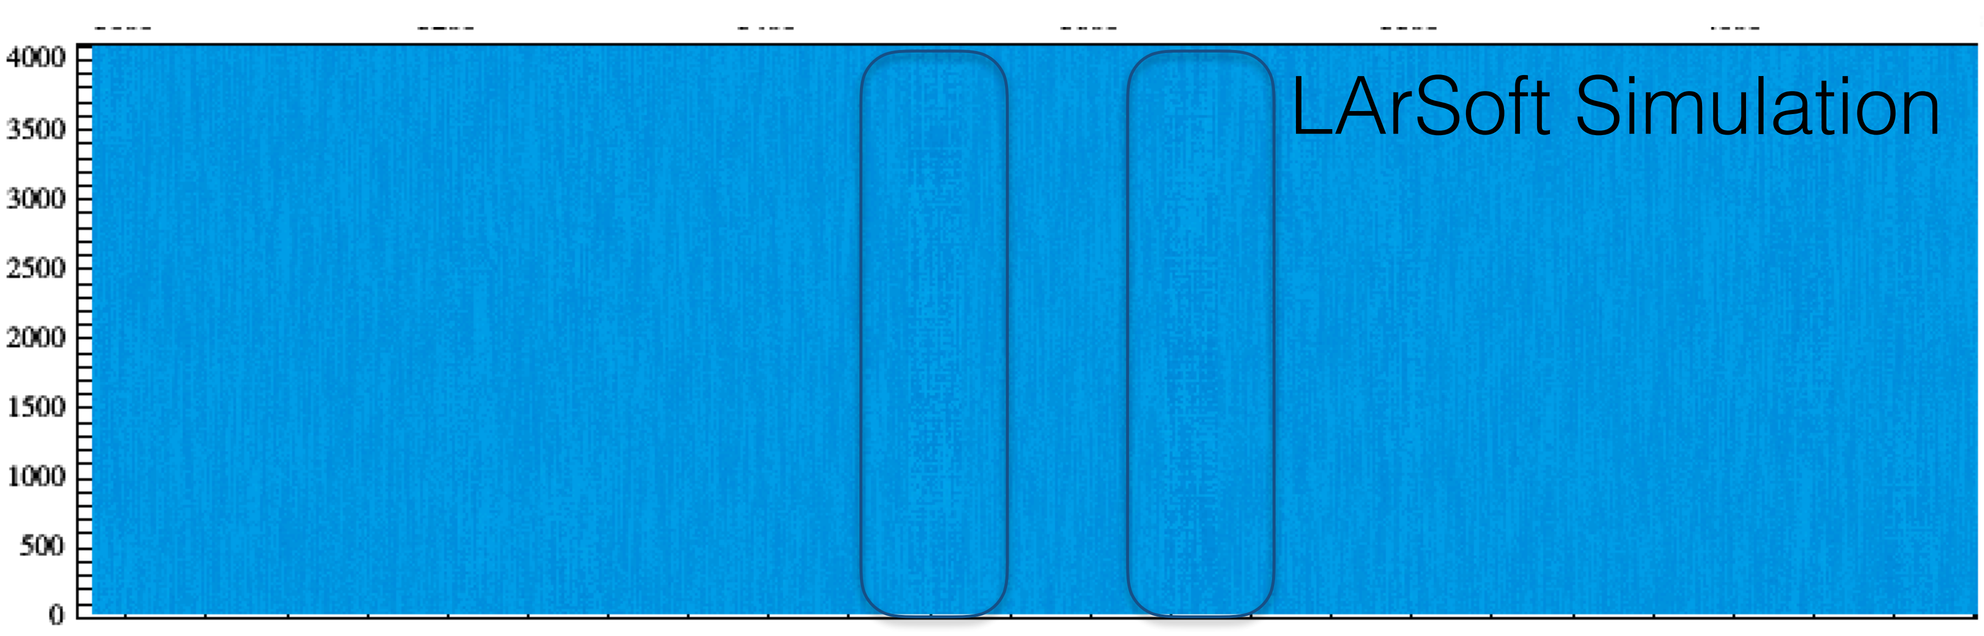

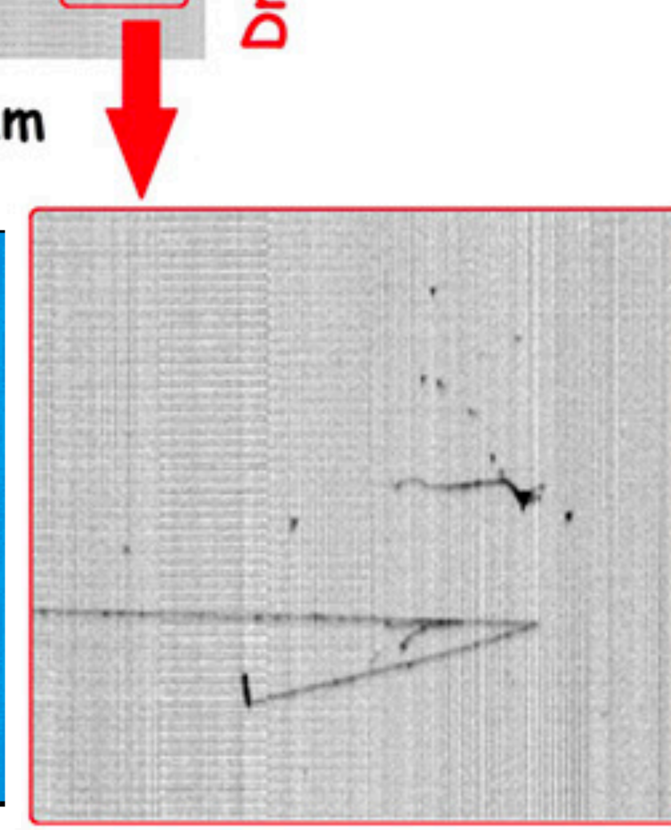

PMT Detector Simulation

- ICARUS specific PMT simulation starts with the Geant4 stage
  - Photon lookup library used to give number of photons seen by each of the phototubes due to the ionization in the liquid argon
    - Too time consuming to try to track each photon
    - Library pre-built from dedicated jobs tracking photons from individual voxels in the liquid argon
  - Currently have photon library generated for a single cryostat
    - Extended through symmetry considerations to cover both cryostats
  - Problem at the moment: reading into LArSoft it expands out to be several GB of storage…
    - And is currently not accounting for all aspects of light simulation that we ultimately will want
  - Investigations are under way to address this
    - e.g. Gianluca Petrillo is providing code to make further use of symmetries in how the library is stored in LArSoft hoping to significantly reduce its footprint

Reco Step 1: Signal Processing

- ICARUS has two paths for signal processing
- LArSoft/MicroBooNE approach
  - Noise Filter to remove artifacts in the raw TPC data
  - Deconvolution of the raw signals
    - Transform bipolar induction plane signals into gaussian shaped unipolar signals and, ideally, make response on all three planes uniform
    - Currently:1D deconvolution (convolved field and electronics responses)
    - Returns ROIs around potential signal regions
  - Hit finding using the "gaushit" finder.
- ICARUS approach
  - Run "raw hit finding" on the noise filtered waveforms directly
  - Historical approach, will not discuss but note it is cpu intensive

Reco Step 1: Signal Processing

- Noise filter:
  - Primarily involves two types of operations
    - Fast Fourier Transforms to remove noise at particular frequencies
    - Operations across parallel sets of channels
      - Removing coherent noise involves forming an average waveform from 32 adjacent channels and then subtracting this from each channel
  - There are also multiple passes through channels to categorize each
    - Determine pedestal then subtract, get full and truncated rms, etc.
- Deconvolution:
  - Fast Fourier transform on each channel
- Hit Finding
  - Fit multiple gaussian shapes to candidate peaks
    - Currently done by root Tminuet… looking forward to coming change!

Signal Processing - Future

- Currently what I've discussed for noise filtering and deconvolution is running in LArSoft
- The WireCell folks from Brookhaven joined ICARUS this year and are now actively working on updating their noise filtering and 2D deconvolution code to work with ICARUS
  - This will also include the WireCell 2D drift simulation as well
- The time scale for this work is tentative but they are aiming to integrate into the ICARUS workflow on the order of October
  - This is the target time for the beginning of commissioning
- Currently they are also addressing cpu and memory issues encountered with integration to ProtoDUNE so idea is that these techniques will transport to ICARUS too
- Still… we have been asked to keep the current signal processing chain active to make sure we have something available for first data…

Reco Step 2: “Cleaning up”

- ICARUS utilizes warm electronics located outside of the cryostats
  - Target Signal-to-Noise ratio is ~12 for the collection plane, slightly worse on the induction planes (compare ~40 for MicroBooNE)
- This is particularly problematic on the middle induction plane where the bi-polar signal struggles to stand out from the noise
- For example:

Each of the black points represent a “gaushit” hit

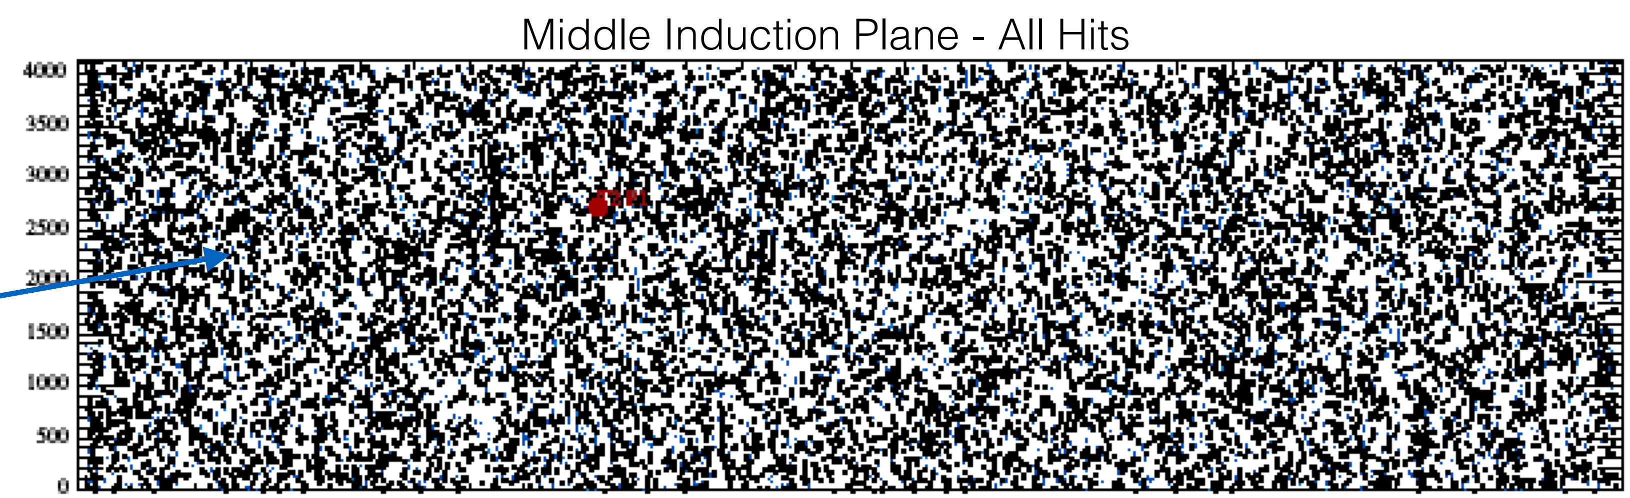

- Current solution is to weed out the spurious hits by creating 3D space points and then outputting a new hit collection comprising the 3D space points.
- For example:

Side Note: 3D space points are output and converted, via LArCV to an output stream to be used by the Deep Learning folks

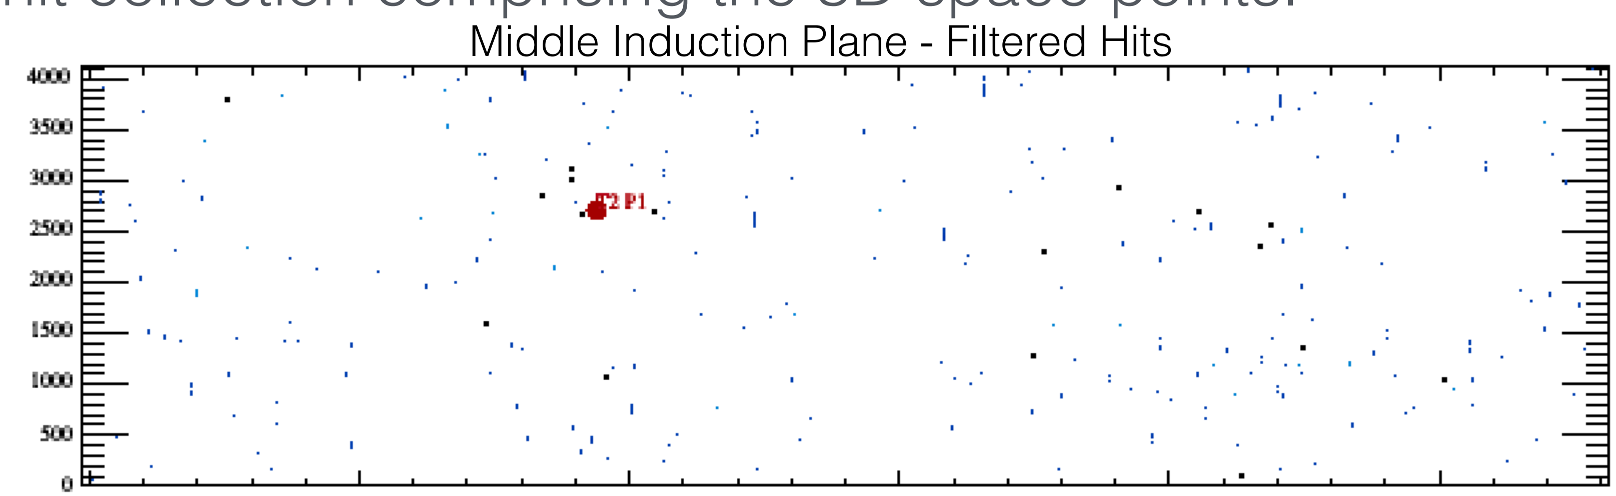

Reconstruction - all the rest

- Once we have hits we follow the “normal” processing chain
  - Currently, SBN experiments using pandora for the common track/ shower pattern recognition
  - Followed by tracking fitting
  - Eventually to include calorimetry, particle ID, etc., etc.
- Also have available PMT optical hit finding and flash reconstruction
  - Current code stolen from MicroBooNE (ok, stolen by original author)
  - Even though fine time resolution, waveforms are zero suppressed so PMT reconstruction not resource intensive
- Also have CRT reconstruction but not well developed at this point

Workflow - In the Beginning

- Generally followed lead of MicroBooNE with separate steps for
  - Event generation
  - Event simulation (Geant4 step)
  - Detector simulation
  - Reconstruction
- With this workflow quickly ran into problems with simulations of Cosmic Ray events
  - Even suppressing channels with no signal, a typical CR event will have output on >40k channels
  - Trying to handle TPC data from all four TPCs in one data object puts pressure on both memory and cpu
  - (it didn’t help to have a couple of memory leaks…)

Workflow - New Strategy

- New strategy to break into logical pieces:
  - Perform signal processing steps on a per TPC basis
    - Noise Filtering, deconvolution and hit finding
    - Results in four sets of output objects at each stage
  - Space Point creation on Cryostat basis
    - Results in new “filtered” hit collections on a cryostat basis
  - Pandora pattern recognition on a cryostat basis
- Requires detector simulation to also stage its output on a TPC basis - four sets of RawDigits
  - Note that this is what the daq people are telling us they plan to do for data taking
- Primary goal is to try to reduce the overall memory footprint for any single execution of one of the algorithms
  - Making it more of a ProtoDUNE size problem at each step…

Current Bottlenecks

- Possible areas where I can see multithreading helping
  - The noise filtering stage is the current loss leader. Currently we are only running channel categorization and FFT to remove low frequency noise component - not yet activating coherent noise removal. This task takes 90-120 seconds/event
  - The “gaushit” finder is next tall pole primarily because of the situation with the middle induction layer and the need to do many multiple gaussian fits on candidate waveforms
    - Improving signal finding would help significantly
    - New fitting algorithm should really really help quite a bit
  - The deconvolution step is the next on the list
    - The new format for the FFT algorithm could help significantly
- These three are already broken by TPC, these could be candidates for testing multi-threading in art? Pointers on how to test welcome!

What Would be Next?

- The SpacePoint building is done with Cluster3D which is currently configured to build SpacePoints (and output a new hit collection) on a Cryostat basis.
  - Building SpacePoints requires looping over hits from each of three planes, currently this is the most time consuming part of the algorithm (as it also does clustering and basic pathfinding as well)
  - Trying to optimize this process could be useful but probably needs consultation with experts
- Pandora is outside the scope of the folks on ICARUS
  - I know the Pandorans are thinking about HPC'ifying their code
- As well, when we switch to WireCell signal processing it will be an in-house BNL project (I assume).
- At this time remaining downstream code is not causing memory or cpu time issue… (stay tuned!)

The Last Slide

- SLAC has a very active Deep Learning group now
    - Currently focusing on developing 3D pattern recognition using a combination of various networks and key algorithms
- Utilizes simulation/data from LArSoft
    - Package “LArCV” contains code to translate LArSoft data structures to their LArCV equivalents
- Network training is performed on GPUs
    - SLAC has a setup of Tesla GPUs sufficient for basic level training
- Are there future plans for LArSoft to try to adapt to GPUs?
    - Ok, Adam’s talk…