# Lessons Learned from the DUNE 35-ton Prototype

M.Convery, T. Junk. M. Stancari

June 2, 2016

## 1 Introduction

Experiments studying neutrino properties with LArTPCs placed in accelerator neutrino beams have taken data successively. However, it is not straightforward to scale all aspects of their design to the multi-kTon size required by DUNE. Instead, the DUNE single-phase LArTPC design employs many novel ideas that are not yet tested in a fully integrated system. The 35-ton prototype detector was built to evaluate the performance of many of these ideas in an integrated system, including a membrane cryostat, a printed circuit board (PCB) based field cage design, waveguides coupled to SiPMs for light collection and detection, cold TPC readout electronics inside the cryostat using pre-amplification and digitization ASICs, wrapped wire planes and a triggerless (continuous readout) data acquisition system.

This document briefly describes the 35-ton components, the history of commissioning, and running to date from the perspective of the collaboration scientists actively involved in the program. The details of the project management aspect of the 35-ton are not covered here. The bulk of the document recounts the lessons learned, big and small, from the 35-ton experience, focusing on those relevant to protoDUNE and the DUNE far detector. The lessons are grouped by category. Each lesson has a brief description of the incident, observation or choice that occured during the 35-ton experience, followed by the consequence, impact or result. The individual lesson is closed with any specific advice for the future.

There are common themes that permeate this collection of lessons. These themes are the real value of the collected lessons as they point to flaws in organization, commitment and communication.

- The group of experts that design and construct each component must be actively involved in the subsequent integration and commissioning of the system to ensure a quality result. This must include physicists who are interested in analyzing the data to be produced. Continuous physical presence and focused attention of at least one expert from each sub-system during the entire period of installation, commissioning and initial data taking is essentialto address integration issues that inevitably arise however much preparation is done. This person must have access to appropriate resources form the responsible institution in the case of crisis. A working apparatus is the result only when the members of each group consider it their mission to optimize the performance of the integrated system, not to solely deliver their component.
- The 35-ton effort suffered from the lack of personpower and the attrition of members in critical roles. Long schedule delays were responsible for some of the loss of staff and expertise. Personnel who were available for a limited amount of time were unable to contribute to key schedule items that were delayed. Across the board, the 35-ton prototype was given lowest priority when conflicting with commitments to other experiments. Turnover situations should be managed carefully to ensure complete information

transfer. Additional risk is present when a physicist, technician or engineer undertakes a task for which he or she is poorly qualified. These situations need to be recognized early, communicated to management, given increased oversight, and allocated additional contingency.

- There needs to be a grounding and shielding plan to which all parties agree. This plan must encompass all disciplines (mechanical, electrical, cryogenic, cryo controls). The team must be committed to adhering to this plan at every stage, and acknowledge that adherence often comes with inconveniences and delays. G&S experts need to be available and vigilant during installation and commissioning.
- The 35-ton required large investments of both effort and resources from each sub-system group over a 2 year period. Despite its prototype nature, an integrated system test of this scale should protect this investment and minimize the risk of individual component failure that could jeopardize the execution of the integrated system tests. Each component should be extensively tested on the bench and deemed functional before being integrated into the system. In particular, bench tests of the TPC cold electronics must be done with a full size (or close to full size) APA attached and done inside a faraday shielded enclosure.

It is important to remember that the 35-ton is a prototype that planned for a short, two month run with cosmic rays as its only source of data. One should not expect levels of stability or performance that would be necessary for a long-term, beam-based experiment. In any case, the personpower available was not sufficient to achieve these levels. The 35-ton prototype succeeded in its mission to evaluate the new design elements. However, in some cases a minimum level of subsystem performance was not achieved.

# 2 Description of the 35-ton detector and the phase 2 run

Figure 1 shows the geometry of the 35-ton TPC. As described above, the 35-ton TPC was constructed using similar techniques as are planned for the Far Detector. The anode plane was not monolithic as in ICARUS and MicroBooNE. Rather, it consisted of four modular Anode Plane Arrays (APAs), each of which had 512 wires. Each APA was "double-sided" in the sense that it could detect drifting electrons impinging on either side. One APA is thus associated with two separate drift volumes. In LArSoft language, each of these volumes is called a "TPC". In the 35-ton, the length of the drift volume on one side was close to the length expected in the DUNE Far Detector. The other used the remaining space available in the cryostat and was much shorter, as a result. Each side of the APA had its own collection plane wires. The induction wires were wrapped around the APA and thus were shared by the two sides.

The four APAs were built to dimensions that allowed them to be installed through the very narrow man-hole of the 35-ton. The two outer APAs cover the full height of the active volume, but the middle APA position was covered by two, shorter APAs. This allowed study of horizontal gaps between APAs as well as vertical, both of which will be present in the FD.

Embedded between the wire planes of each APA is a number of waveguide Photon Detectors (PDs) with their attached Silicon PhotoMultipliers (SiPM). All of the APAs, except for APA1 also contained a grounded grid plane between its collection plane and its PDs. This provided shielding between the PDs and the wire planes.

The TPC signals were amplified, digitized and transmitted by Front End Mother Boards (FEMBs) mounted on the APAs. Each FEMB handled 128 channels, so that four were required for each APA and 16 for the whole TPC. The digital signals were transmitted at 1Gb/second through cables that exited the TPC through a "flange board" mounted on the top of the cryostat. This board also provided connections for power

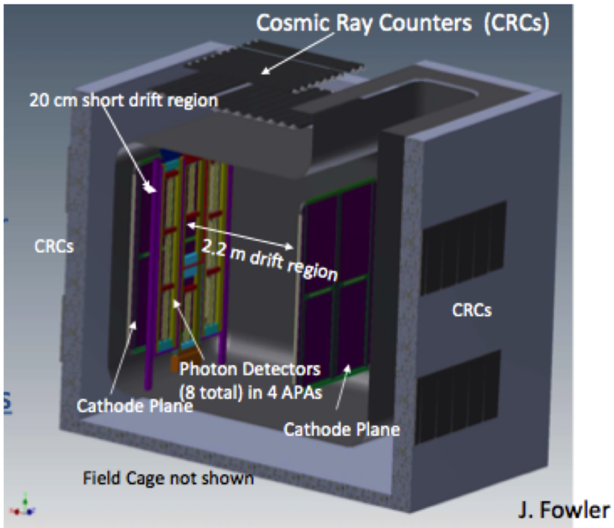


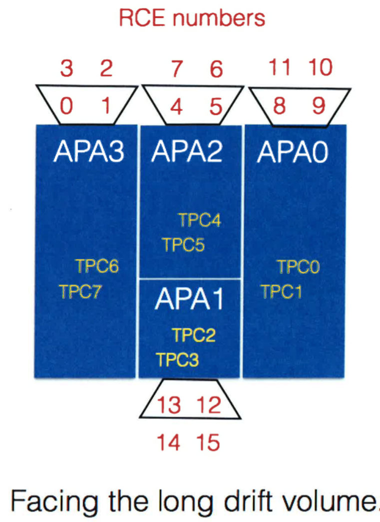


Figure 1: Cutaway drawing of the 35-ton prototype (left) and map showing APA offline numbering (right).

cables and the PD analog signal cables. Once the digital TPC signals exited the cryostat, they were converted to optical signals by the Optical Converter Board. From there, they were readout by Reconfigurable Cluster Elements (RCEs), which are FPGA-based data processing nodes. There was one RCE for each FEMB and the two components used identical numbering. Eight RCEs are contained in one Cluster On Board (COB), of which there were two. The analog PD signals were digitized by SiPM Signal Processor (SSP).

The first run of the 35-ton prototype tested the membrane cryostat technology and verified that the required electron lifetime of 3 ms was achievable inside a membrane cryostat. This early run included a minimal amount of diagnostic equipment in the cryostat, just purity monitors and temperature sensors. Detector installation for the second run took place from September - December 2015. The cooldown and filling of the cryostat took place in December 2015 and January 2016. The argon filtration/recirculation system was turned on February 11. Detector commissioning began soon thereafter, and continued until the start of the data run on March 7. The data run ended prematurely on March 19. A pipe connected to a vibrating pump fatigued and broke, resulting in unrecoverable contamination of the liquid argon. A series of dedicated studies was performed after the contamination, including observing the stability of the HV distributions system (field cage, cathode, and feedthrough) at different cathode voltage settings, various DAQ studies and measurements to identify the sources of noise in the TPC and photon detector systems. The system is currently warming up and further noise studies will be done. Once the detector components are removed from the cryostat, the readout electronics will be further investigated on a test bench.

## 3 Detector construction

The APA planes were constructed well in advance of installation and contained PD modules of early design. This was the first time that modular APAs with a design similar to that of the DUNE FD have been constructed. Their construction was successful and identified several possibilities for improvement that will be employed for ProtoDUNE and beyond. Similarly, this is the first time that PD modules of DUNE design have deployed in a large TPC. Although they were of an early design and thus have lower light yield than current models, their construction and installation yielded insights that will be valuable for future TPCs.

### 3.1 APA construction

**Contact:** Dan Wahl, Bob Paulos

**Item** Broken induction plane wires during APA cold tests, see DUNE docDB 913 *APA winding.*

**Result** A redesign of the wire interface boards and the glue strip that avoids anchoring the wire to two different boards was already implemented.

**Contact:** Michelle Stancari, Tom Junk, Mike Wallbank

**Item** Identifying the broken wires in offline channel space continues to be an issue. The location of the broken wires was recorded but the coordinate system is ambiguous - preserved under a 180 degree rotation around the long axis of the APA. The definition of the coordinate system used was not well documented, and the documentation was incorrect. The abundance of dead channels (13%) complicates verification of the position of the broken wires. In addition, the orientation of the APAs was not recorded when they were installed, and muon tracks were required to resolve this.

**Result** Six iterations of the online-offline channel map were necessary before it converged. Using tracks tagged by the counters was necessary to define and verify the channel map. The initial incorrect versions of the channel map delayed other offline tasks including offline data quality monitoring, and consumed personpower. Some of the broken wires and their compromised neighbors have yet to be identified in the offline channel numbering.

**Advice** In advance of APA assembly and testing, the convention for identifying individual wires and recording the installed orientation should be agreed upon and documented. A mark on one corner of the APA may be necessary to break the rotation ambiguity. A plan to verify the channel map with muon tracks crossing every wire in a particular plane should be developed in advance, and may require the placement of small external counters and a special DAQ trigger.

**Item** The photon detectors were installed inside the APAs roughly 24 months before the APAs were installed in the cryostat because they had to be inserted into the frame before the wires were wrapped around the frame.

**Result** The photon detectors used in the 35-ton were two generations old by the time data were taken. Protecting the TPB coating from visible light makes handling the APAs more complicated and there is greater risk that the light output degrades. The APA frame design was modified to allow insertion of the photon detectors later in the APA assembly process.

**Advice** Both to protect the photon detectors and to ensure that the best quality detectors are installed, do not anticipate the insertion of photon detectors in the current assembly plan. The insertion should happen as late in the process as possible.

## 4 Detector Installation and Integration

The 35-ton offered an excellent opportunity to practice installation of a TPC into a membrane cryostat - an operation that had never been performed before. The installation proceeded according to schedule and also identified a number of potential problem areas.

## 4.1 Installation

**Contact:** Jack Fowler, Lee Scott

### 4.1.1 Trial installation

**Item** A trial installation was performed in the Fall of 2014 at PSL with PVC piping delineating the cryostat walls. Final components were used.

**Result** Detector installation in cryostat was extremely smooth. It was well planned and executed. The trial assembly specifically allowed the group to

- check that all parts actually fit together
- check that all of the appropriate fasteners and tools were available. Some special or modified tools were needed for certain instances
- verify that access was possible to install the parts
- learn what types of access equipment was necessary
- rework some parts where unforeseen interferences occurred
- make an accurate assessment of number of people required to perform the tasks
- do a thorough hazard analysis of the tasks
- confirm the sequence in which the parts should/could be assembled

**Advice** A trial installation is absolutely necessary. We should have added a placeholder for the electronics and faraday shield. These were not available at the time of the trial assembly. Not having these objects in place allowed for much easier access than turned out to be the case during the actual installation where these objects were difficult to work around. A cabling mock up prior to this would have been helpful. An *as built* survey of the cryostat membrane may be necessary in some congested areas to eliminate interferences during installation.

### 4.1.2 Forgotten Photon Detector Cable

**Contact** Alan Hahn, Alex Himmel

**Item** A bundle of extension cables connects the cables from the detectors (TPC and PDS) to the flange board at the cryostat boundary. When the APAs and the photon detectors were cabled, one set of photon detector cables was left disconnected at the detector-extension cable junction. This one set of detectors has a different cable routing - underneath the APAs instead of above the APAs - making them vulnerable to installation mistakes. The cabling was done by technicians due to the confined space training requirement. The checkout procedure followed by the PDS group did not uncover this problem and they gave the go ahead to proceed with installation. The procedure consisted of measuring the noise level on each channel with the DAQ readout and was performed by PDS group members inexperienced with the hardware instead of experts that have worked extensively with the various test stands. Alan Hahn performed a diode test (putting a test voltage across the cable measuring the resistance for both polarities of the voltage) on all channels of the PDS system a couple weeks after the checkout. This test definitively identified an open connection, however the cable junction was no longer accessible.

**Result** One sixth of the coverage (xx% of the channels) for light detection did not work. One photon detector design was not characterized.

**Advice** The detector checkout procedures should be written down and reviewed by the integration team in advance. The results of the checkout should be presented to the integration team after execution but before proceeding with installation. Abnormal or exceptional situations that have the potential to cause mistakes should be identified in advance and triple checked during installation.

#### 4.1.3 High Resolution Photographs

**Contact** Tom Junk

**Item** The location of the photon detectors was poorly recorded and the TPC channel identification system used by PSL to denote the location of broken wires was ambiguous.

**Result** We were able to have confidence in the detector details used in offline reconstruction because of the excellent quality of the photographs of the detector. One can see missing wires, non-uniformities of the field cage, the amount of overhang between the field cage and the APA's, the deflector, and the lack of mesh on one of the APA's.

**Advice** High-resolution photographs of the detector as assembled are essential. There should be line items in the schedule for taking such photographs. They should be stored in a public place where they can be viewed easily.

#### 4.1.4 Capacitive Wire Pulsing (The Tickler)

**Contact** Mark Convery, Michelle Stancari

**Item** After detector installation, but before real data is taken, the only diagnostics available for testing the TPC are measurements of noise and of pulses injected directly into the elecronics. Very useful additional information may be obtained by performing a capacitive pulsing test, in which signals are induced on the wires themselves by pulsing a copper plate placed in proximity to the wire planes, as shown in Figure 2 (left). In the 35-ton test, we called this plate "The Tickler". The charge, $q$,induced on a wire is given by $q = C\Delta V$, where $C$ is the capacitance between the plate and the wire, and $\Delta V$ is the height of the driving pulse. The typical capacitance between the plate and a wire in the outermost plane is roughly 0.1pF, implying that a $\Delta V$ of about 1 V should induce a 2000 ADC count signal, given a gain of about 10 mV/fC and 0.5 mV/ADC. A typical induced signal is shown in Figure 2 (right). Induced signals are largest on the outermost induction wire plane. Due to the shielding effect, the signals are smaller on the second induction plane and smaller still on the closer collection plane. Signals on the "far" collection plane are likely not visible due to shielding and distance. A detailed calculation could likely yield precise pulse height predictions for a given plate position, although this was not done for the 35-ton test.

**Result** These tests yielded the following information that was not available from noise measurements alone:

- Confirmation of the connection between each wire and the FEMB.
- A test of the map between electronics channel and physical wire position.
- An early data-set containing "real" pulses that is useful for testing the whole analysis chain including timing and hit finding.

**Advice** If possible, these tests would provide useful information during the commissioning of future detectors such as ProtoDUNE-SP and the FD.

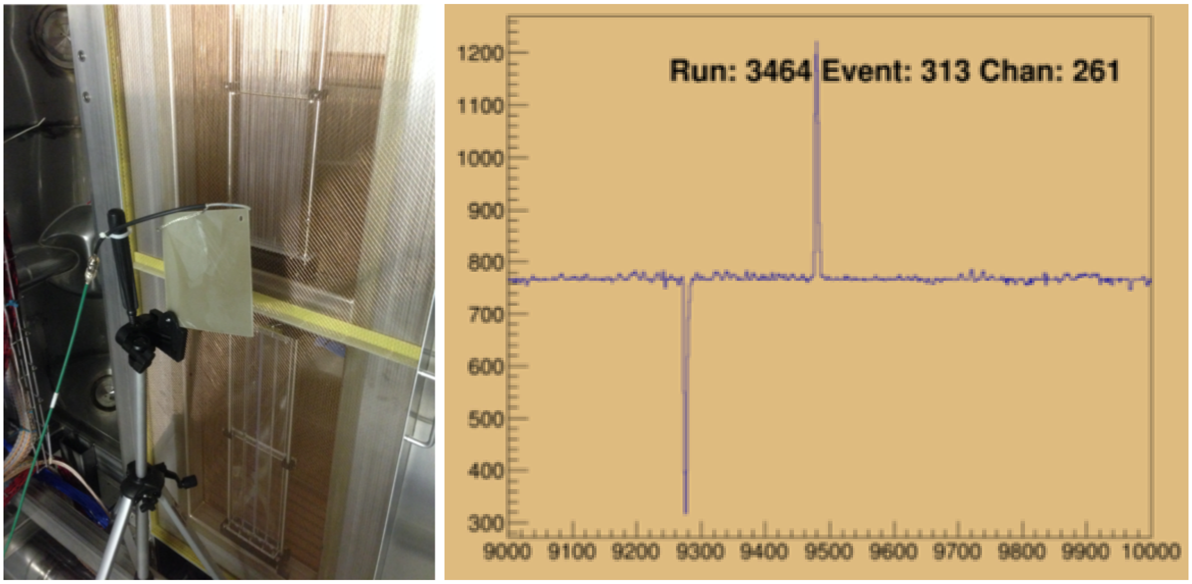


Figure 2: Copper plate placed in proximity to the wire planes, which was used to induce signals on the wires (left). Waveform showing a typical induced signal (right).

## 4.2 Installation and Integration Management

**Item** We lost Joe Howell as installation and integration manager in the Spring of 2016. He was not replaced.

**Result** Many of Joe's responsibilities were already completed, such as the installation of the external counters. The rest were taken over by Jack Fowler (and Russ Rucinski? and Alan Hahn?). However, Jack, Russ and Alan were already heavily involved in the project, and their available time saturated. Many items did not get enough attention during installation due to the lack of personpower.

**Advice** When possible, plan the turnover of critical people in advance. Add additional dedicated personpower during the transition. Any known holes in staffing should be watched carefully by management.

## 4.3 Grounding and Shielding

### 4.3.1 Noise sweep

**Contact** Linda Bagby

**Item** A noise sweep was performed inside the 35-ton cryostat before APAs were installed. Several floating (not grounded) pipes that introduced noise inside the cryostat were identified. The pipes were subsequently grounded.

**Result** The noise introduced into cryostat by these pipes was eliminated before detector installation.

**Advice** Perform a noise sweep inside the cryostat frequently as cryogenic piping and other external equipment is installed to identify noise sources early.

### 4.3.2 Integration

**Contact** Jack Fowler, Terri Shaw

**Item** Photon Detector (SiPM) mechanical holders contain metal which was not connected to the APA frame. Grounding rules would have required a solid ground connections between these elements.

**Result** Unknown

**Advice** The interface of all detector elements to detector ground must be documented and reviewed.

## 4.4 Feedthrough/Cables/Connections

### 4.4.1 Flange Board

**Contact** Terri Shaw

**Item** The Flange Board schematics were thoroughly checked and were OK. However, the Flange layout had to go through three iterations due to errors such as part cells being built with incorrect shape and pinout. In the end, parts still needed to have pins physically trimmed back in order to mount them on the board without creating shorts and soldermask had to be physically removed from pads it was covering. The mechanical technician at Penn who designed the layout and stuffed the board did not have the appropriate experience or supervision to perform this job well. Alan Hahn discovered we had significant leakage on the wire bias voltages which go through the Flange Board. This has yet to be understood. It was discovered late and required a late re-routing of the wire bias cables. The re-routing of the cables did not conform to the G&S guidelines, and in fact introduced a small amount of noise inside the cryostat.

**Result** The layout iterations and cleanup (pin trimming and solder mask) delayed the vertical slice test of the DAQ 4 months and consumed person resources at Fermilab. This would have delayed the entire installation schedule had the cold electronics been delivered on time. It did delay the vertical slice test significantly.

**Advice** QA/QC needs to be done by the group responsible for the component and before integration of the component. When unexpected problems require solutions outside the plan, the solution should be reviewed by the G&S committee. Ideally the group responsible for the component should propose and implement the solution. (The group that designed the flange board was aware of the impedance specifications for the feedthrough connections by the G&S committee).

**Contact** Terri Shaw

**Item** Several connectors chosen for use at the Flange board interface are friction lock. While this may save space, it is a poor choice for cable retention. High speed GORE cables, while possessing a "lock", are easy to wiggle and pop free.

**Advice** Choice of cable connectors should be reviewed.

**Contact** Terri Shaw

**Item** Many of the cables we require are custom and will have a long manufacturing leadtime. Getting the high speed cables specified and produced took ∼1 year from initial vendor contact.

**Result** No adverse affect as other schedule delays were more severe.

**Advice** While not all cables will take this long, we should not underestimate the time required for this task.

**Contact** Terri Shaw

**Item** With the larger wire size required for the power feeds, the crimped pins in the box connectors tend to pop out if handled roughly. This happens even though everything is within spec of the connector and pins.

**Result** Time lost during cable fabrication and installation.

**Advice** Alternatives to the use of crimped pins in box connectors should be considered.

**Contact** Jack Fowler

**Item** The connectors for the LV and control cables in the 35t used very small loose hardware for attachment.

**Result** NEAR MISS: During installation, there was a high risk of small hardware being dropped inside the field cage or the APA.

**Advice** No loose hardware should be used. If any hardware is needed it should be captive to the connector.

**Contact** Jack Fowler

**Item** For some of the 35t control cables, the connectors were inverted on the ribbon cables.

**Result** These cables were forced to immediately make 180 degree turn in a very congested area.

**Advice** Attention to the connector orientation should be specified on the cable drawings.

**Contact** Jack Fowler

**Item** Kapton tape was specified to bundle harnesses from Gore. It was later determined that Kapton tape was problematic - many people had experiences where the adhesive was dissolved by LAr.

**Result** We removed all of the Kapton tape and replaced it with Tefzel cable ties, at the cost of time and sticky fingers.

**Advice** A better method to bundle cables would be to identify a LAr compatible wrap or jacket to organize and protect these bundles. Testing of all materials should be performed before any purchases/drawings are made.

#### 4.4.2 TPC readout cabling

**Contact** Jack Fowler

**Item** The APAs had cable connection points on both (front and back) faces of the upper frame of the APA. The upper frame consists of an electronics plate or fin as it is sometimes referred to. The cold electronics (CE) boards were mounted to this fin by standard standoffs. The entire fin area with CE boards was surrounded by a SS housing referred to as the faraday cage (FC). Inside the FC, there were interconnects directly between the boards through board mounted connectors. LV power and control signals were brought to the board via an internal harness inside the FC. These internal harnesses had connectors mounted through the FC for the cold cables to attach. There was not enough space for these internal harnesses and they needed to be attached to the FC panels before assembling them to the APA.

**Result** This made the assembly very difficult.

**Advice** For the far detector, the FC panels are longer (2.3 m) and the cable bundles will have 8 to 10X more connectors. Trial runs of the cable routing and assembly should be done before the design is final. In both the FD and ProtoDune, 2/3s of the APAs will only have access from one face. The other face will be ¡ 15 cm from the cryostat wall and likely 40 to 50 cm from the cryostat ceiling. There will be no direct access to the rear face of the FC to make cable connections and this will be a very congested area. All cable connections will need to be on one face of the FC.

#### 4.4.3 APA Numbering

**Contact** Mark Convery

**Item** At least three different numbering schemes were employed to identify individual APAs:

- Offline numbering as described above and used throughout this document
- "Online" numbering, where the APA number follows the RCE numbering, so that RCEs 0-3 belong to APA0, etc.
- Construction-based numbering: "3-board", "4-board", etc.

**Result** Confusion due to people using different numbering schemes

**Advice** A single numbering convention should be adopted early on for ProtoDUNE. This is likely to be the offline convention, since it will defined first. This convention should be well documented, publicized and enforced. Construction "serial numbers" can be used, but should be abandoned soon after installation.

# 5 Cryogenics and Purity

The Phase II 35-ton cyrogenic system reused much of the system built for Phase I, which in-turn reused much of the system built for the LAPD test. As a result, the commissioning was relatively straightforward and the system provided good purification capability very quickly. On the other hand, the LAPD system was likely not built with such a long life in mind. So, the system was not as easy to use or robust as will be needed for a beam experiment.

## 5.1 Placeholder

### 5.1.1 Installation

**Contact** Alan Hahn, Russ Rucinski

**Item** We had problems establishing the 35-ton internal LAr pumping. The pumps had been rebuilt after phase 1, but had never been "run in". Finally putting 3 phase 208 AC on pump motor broke it free, afterward the normal AC drive worked fine.

**Result** It took 7-10 days to resolve this situation, caused a delay in commissioning. No impact to the overall schedule as many items also had delays.

**Advice** Nothing goes into the cryostat without a functionality test, even if it worked fine last time.

### 5.1.2 Operations

**Contact** Alan Hahn

**Item** It was painful to use the cryo interface, it was extremely slow.

**Result** Anyone not familiar with the system (i.e. not a cryo expert) was not able to look at the data from the system. If shifters and detector experts have easy access to this information, they will look for trends and correlate those trends to what they see.

**Advice** Need to develop a better interface into Cryo controls and alarms

### 5.1.3 Pipe failure resulting in severe argon contamination

See Fermilab internal root cause analysis.

# 6 TPC readout Electronics

The 35-ton electronics represented the next step in Cold Electronics development with the integration of a cold ADC chip and digital data transmission together with the previously developed cold analog ASIC. Although the development took longer than expected, and a number of unexpected problems were observed, the 35-ton experience will be invaluable for future projects, such as SBND and ProtoDUNE, that intend to use very similar electronics.

## 6.1 Schedule delays

**Item** The delivery of the FEMBs to Fermilab was several months late, and this delay was announced two weeks at a time.

**Result** Other groups were affected by this and perhaps could have adjusted the allocation of their resources better if a more realistic picture of the situation was available.

**Advice** Better communication. Management should be vigilant about how schedule delays in one area can affect others, and proactivly minimize losses when possible.

## 6.2 Checkout procedures

**Item** The checkout procedure to verify that the FEMBs had not been damaged during transport or installation was not provided until requested by the technical coordinator, and after the installation of the boards on the APAs. The CE group did not come to Fermilab to perform the installation or the testing.

**Result** Non-experts did the installation and testing. Serious noise issues were observed but not understood. The latter would not have been different if the experts were present, but the testing may have finished quicker. Local personpower was significantly over subscribed and unable to do the job efficiently.

## 6.3 Miscellaneous

### 6.3.1 Mounting hardware

**Contact** Terri Shaw

**Item** The FE board stack requires several screws and spacers to hold the board stack together and to mount it to the APA. This hardware has not been well documented and was supplied late.

**Result** There was no time to do cold cycle tests.

**Advice** The ability of the hardware to go through a cold cycle and remain tight should be considered; do we need to specify belleville washers?

**Item** Significant board warpage was seen in several boards mounted on the APAs.

**Result** Unknown, but a possible explanation for the problems described in Section 15.2.3

**Advice** We need to make sure the proper mounting hardware and PCB layout and materials allow for construction of a planar board. Cold cycles cannot be allowed to induce significant warpage.

**Item** At least two cases of “dirty” pins on the cold FE board stack were found during checkout at DAB.

**Result** This required a lot of time to track down.

**Advice** Connector pins should be cleaned.

**Item** The ground connection between the FEMBs and the APA must be low impedance. In the 35-ton this relies on the mounting hardware.

**Result** Unknown

**Advice** Characterize the impedance of the 35-ton connections, design a better connection if necessary. Measuring the impedance of this connection must be a part of the QA/QC procedure for each individual APA.

#### 6.3.2 Documentation

**Contact** Terri Shaw

**Item** At one point we had out of date schematics and connector pinout information for the FPGA board in the FE board stack.

**Result** This led to power/ground pins being reversed when the first card was powered. It also required re-working of several cable harnesses which had been built to the wrong specifications.

**Advice** Up-to-date documentation must be available.

#### 6.3.3 Test bench noise measurements

**Contact** Alan Hahn, Terri Shaw, Linda Bagby

**Item** The FEMBs and wire planes were not tested together in a realistic test. a coupling of the power bus into the front end chips was missed during bench testing. Excess amounts of noise were observed in the FEMBs several times during the 18 months before installation. Each time, the source of the noise was only briefly investigated due to lack of appropriate tools and personpower.

**Result** A coupling of the power bus into the front end chips was missed during bench testing. This was a large contribution to the noise level, but not all of it. The excessive noise level in the 35-ton rendered many of the planned investigations and measurements impossible.

**Advice** A combined electrical and mechanical design team is needed. A test facility with adequate shielding so that a subset of the detector can be operated at room temp outside the cryostat is fundamental. This system will have long wires and very low noise electronics so a **large shielded enclosure** is required for this test to be meaningful. Bench tests need to show the expected level of warm noise before anything is installed in the cryostat.

#### 6.3.4 Diagnosing noise problems

**Contact** Linda Bagby, Terri Shaw, Marvin Johnson

**Item** All channel waveforms are digitized before they leave the cryostat. This prevents the use of commercial instruments such as spectrum analyzers to study detector issues. FFT analysis was performed on the digital signals written to disk by the DAQ, but this process required 10-15 minutes and significantly slowed down the noise investigations. Because the grid plane is floating except for the connection to the voltage supply outside the cryostat, it makes a decent antenna to study the noise inside the cryostat.

**Result** Identifying the sources of noise in the cryostat was severely hindered by the lack of an analog signal. The experts could not use the tools they are accustomed to, in particular spectrum analyzers. The long turn around time for a measured frequency spectrum limited the extent of the testing.

**Advice** The analog signal from one channel on every board must exit the cryostat. The best system is to have the output of input amplifier on the front end ASIC routed (with an analog buffer/driver) to an output pad on the chip and then brought out of the cryostat, but this removes the channel from the main data stream.

**Contact** Michelle Stancari, Linda Bagby

**Item** The measured noise levels increased during installation, however a single cause could not be identified. The noise levels were monitored only sporadically and the settings of the preamp ASIC, which confused the comparison of different noise measurements, were not readily available.

**Result** While the bulk of the unexpected noise likely has a root cause in a sub-system or integration flaw and not an inadvertent ground loop created during installation, small amounts of noise have been attributed to the building overhead lights and the TPC bias HV feedthrough. A lot of time was spent recreating the installation process (connecting and turning on one component at a time) because noise measurements were not systematically done, recorded and analyzed after each step of installation.

**Advice** Continuously monitor the in-cryostat electronics during installation, cool-down and commissioning. For example, running a pulsing system at the end of every working shift. Identify a person or group of people responsible for examining the data.

# 7 Trigger

The trigger system of the 35-ton, which was also responsible for the readout of the external scintillator counters, enabled the DAQ to write to disk the most interesting of the cosmic rays that passed through the 35-ton. The development of a multi-purpose module that is interfaced with ArtDAQ may have applications in future experiments.

**Contact** Josh Klein, Nuno Barros, Michelle Stancari, Rick Van Berg

**Item** The Penn Trigger Board was delivered and commissioned very late, though it was not ultimately responsible for any delays. It was tested and commissioned at the last minute even though a test stand was available for doing so more than six months in advance. Due to changes in personnel, agreements on the logic details and channel mapping between Penn and Fermilab were lost. This postdoc who took on responsibility for the project needed some time to come up to speed and needed better support to complete the job in a timely fashion. A consequence of the late integration was that a DAQ expert was no longer available to assist with the integration.

**Result** The cabling of the counters had to be redone to conform to the channel map assumed by the board layout. Board testing and integration happened simultaneously with DAQ commissioning and detector monitoring during the cooldown, and these items competed for scarce resources.

**Advice** The staffing plan needs to anticipate possible changes in personnel and include adequate time contingency for the transfer of information and coming up to speed.

# 8 Photon Detectors

The 35-ton represented the first opportunity to install and operate PDs of the design planned for the DUNE FD in a real TPC. The lessons learned about the operation of the PDs themselves as well as the SSPs used to read them out will be invaluable for future experiments, such as SBND and ProtoDUNE.

## 8.1 Commissioning

**Contact** Alex Himmel

See DUNE docdb 913 *Photon Detector Operations*

**Item** There was no plan developed ahead of time for commissioning the photon detectors due to lack of personpower.

**Result** For most of the run, self triggered PD data was limited to special PD only runs, otherwise data rate was too high for the DAQ. Initial self triggered photon detector data used only 3 SSPs. That was all the team could get working in time. Gains and thresholds were not set in a consistent way and changed during data taking. This information is not recorded anywhere easily accessible.

**Advice** Have PDS monitoring software ready in advance. Define a procedure for setting gains and thresholds in advance. Block out time in the schedule for calibration studies after cooldown. Run-by-run gain and threshold settings need to get put in the database in an automatic way.

# 9 DAQ

The 35-ton DAQ was the first opportunity to test the idea of a "continuous readout" of LArTPC data, which is necessary to achieve the non-beam physics goals of the DUNE FD. In this system, the front end electronics transmits all of its data to the DAQ, without compression of any kind. The DAQ can then decide which data to keep to stay within its available bandwidth. In principle, this could mean saving data for all times, but with a high level of compression. In practice, for the 35-ton, we used an external trigger to select the interesting times.

The system performed well in operation and wrote more than 30 TBytes of data to disk. The development and operation of this system will be extremely valuable for future LArTPC experiments. Much of the hardware, firmware and software will be directly reused in ProtoDUNE-SP.

## 9.1 DAQ Hardware

### 9.1.1 Lack of adequate DAQ room/environment

**Contact** Giles Barr

**Item** The site selected for the 35t phase 2 test was unsuitable for delicate electronics because of the combined presence of dust and humidity. No provision was made for an electronics hut. The consequence was that NEMA-12 rack solution was used, which is a non-standard configuration for which little experience existed among the people.

**Result** See section 9.1.2.

**Advice** Adequate environment should be planned for the external readout electronics and computers.

### 9.1.2 Rack protection

**Contact** Alan Hahn, Linda Bagby, Mark Convery

**Item** The 35-ton DAQ computers and RCEs were housed in air-conditioned NEMA 12 racks to protect against dust and humidity. The AC unit that cooled the rack could be disabled. Three different times the AC unit was disabled with the rack doors closed. The first two incidents were caught quickly. The third was caught only several hours after the rack doors were closed. At roughly the same time, the network switch reported a temperature ¿ 60C and the deputy run coordinator (Mark Convery) was not able to operate the DAQ (the RCEs did not communicate) remotely. He decided to check the rack at PC4 and found doors closed and the AC off. Opening the doors temporarily mitigated the situation. The RCEs sustained a small amount of mechanical damage to their front panel latches that required a couple of days to mitigate.

**Result** Downtime while RCEs were repaired. NEAR MISS on loss of critical components (RCEs and COBs) for which we had no spares on site.

**Advice** Keep spares on site. The cost of spares needs to be included in the overall project cost. All expensive or irreplaceable equipment should be housed in racks with temperature interlocks. The need for drip sensors should be evaluated.

#### 9.1.3 COB failure due to ECR modification failure

**Contact** Alan Hahn, Matt Graham, Mark Convery, Terri Shaw

**Item** The Fermilab ECR (Electronics Review needed for Operational Readiness Clearance) led to a requirement for an additional fuse to be added to the COB board. This fuse eventually made contact with a ground rail and overheated, causing damage to the COB. Since it was a non-standard modification, SLAC had to modify an existing COB and ship it to FNAL.

**Result** Getting a spare COB modified and shipped from SLAC during the Christmas holidays took 10 calendar days.

**Advice** Keep spares on site. Include cost of spares in the project cost. Non-standard hardware modifications require special attention to ensure reliability.

#### 9.1.4 Optical Link Stability

**Contact** Mark Convery, Matt Graham

**Item**The optical links between RCEs and FEMBs were initially unstable after installation at PC4. This was tracked down to two causes: the power supply of the Optical Converter Board did not have *Remote Sense* and the Nova timing unit was not providing a clean clock to the RCEs due to excessive cable length. After these were fixed, the only link problems observed were due to power cycling or other interventions.

**Result** Initial commissioning at PC4 was slower than necessary. No schedule delay. It took much longer than it should have to identify the too-long cable. There were too few people doing too many things, our single contact with the Nova timing experts had limited availability.

**Advice** Ensure that all of the *simple* aspects of detector installation are done properly before attempting more complex operations. The ability to perform hardware resets of the links may have helped with debugging.

#### 9.1.5 DAQ development

**Item** The photon detector readout electronics, called SSPs, were delivered on time (2014) and integrated smoothly into the DAQ.

**Contact** Giles Barr

**Item** Only 16 cold electronics boards were delivered, the exact number needed to instrument the detector. It is necessary to have a spare to maintain a vertical slice during the installation period. This way, DAQ development can continue during the installation and configuration of the DAQ computers at PC4 (2 weeks). Any problems, such a link errors that are found during installation can be investigated in a lab environment, where test equipment, such as scope probes, could be applied to the things on the cold side.

**Advice** Maintain the vertical slice test stand during DAQ installation and commissioning.

**Contact** Giles Barr

**Item** The online monitoring could have been available earlier and been more extensive. This would have helped early DAQ development. The addition of gateway2 and gateway3 (extra computers for monitoring) happened very late, and caused confusion for the network topology. We note that the DAQ is funded by the project, but the monitoring is not because it is categorized under software and computing. The scientific effort from the collaboration for non-project items was very scarce and over-subscribed.

**Advice** Include the data quality monitoring in the DAQ planning from the beginning. When allocating manpower for software and computing tasks, give DQM priority.

## 9.2 DAQ operations

### 9.2.1 Need for Local DAQ Experts

**Contact** Mark Convery, Giles Barr

**Item** The 35-ton DAQ system was initially difficult for non-experts to operate. Because the DAQ was relied upon for detector checkout prior to installation, proper support needed to be supplied. Unfortunately, the schedule for CE delivery slipped quite a bit. Local experts were available during the originally scheduled commissioning period of Sept-Dec 2014. When the CE arrived at FNAL (June 2015), the experts who had planned to provide local DAQ support during detector checkout were no longer available. The DAQ commissioning actually occured from Sept 2015- February 2016. The funding for experts, both salary and travel, had already been spent. Suitable replacements, local or remote, were seldom available.

**Result** Some of the APA checkout at DAB took longer than necessary. The DAQ commissioning at PC4, in particular the network reconfiguration, took longer than necessary.

**Advice** All subsystems must respond to schedule delays of other subsystems. Responding to these delays requires careful management of resources, so that they are not all spent before they are needed. If more funding is required to make such responses, it must be found.

**Contact** Alan Hahn

**Item** The DAQ servers were not setup as Fermi Standard. The experts with sysadmin rights were not local.

**Result** We were halted several times because the experts were not available to troubleshoot problems or change configurations after hardware changes. A couple instances we had to wait for a few days.

**Advice** Local experts are necessary. Critical roles like sysadmin need multiple individuals that together provide 24/7 availability. Fermilab SCD can offer more support if/when the system conforms to Fermilab standards and security policies.

**Item** The photon detector readout electronics (SSPs) did have good techinical support throughout the run from Argonne. When problems were found, they were investigated and mitigated quicky. Both instances involved a firmware modification. The support was reactive and not proactive. The Argonne group did not participate in integration discussions before or during commissioning unless it was specifically requested.

**Result** The product delivered by Argonne was high quality and well supported. There was very little involvement of Argonne scientists in the 35ton running, and the SSPs could have been more extensively tested. It was smoothly integrated into the DAQ because the UK personnel were available at the scheduled integration time.

### 9.2.2 Photon Detector Data Rate

**Contact** Alex Himmel

**Item** The data rate from the photon detector system is not trivial when reading out the waveforms at 128 MHz in self triggered mode, even for a zero-suppressed TPC with no noise at all. Nominally, the 35-ton's 63 active PDS channels require 70 MB/s for cosmic rays, 20 MB/s for argon-39 and possibly much more for radon. This amount of data exceeds the TPC data stream and does not fit in the DAQ pipeline. This was realized a couple of months before the start of the run, but only after probing questions to the PDS group from the operations team. The PDS group had a default readout mode from their test stands which became the initial 35-ton readout plans via inertia. The PDS group had not estimated data rates for the entire system, and was unaware of the DAQ throughput limitations.

**Result** Once a data rate estimate was requested from the PDS and a PDS data budget specified by the DAQ group, the issue was well defined. The PDS group developed different running modes, including a header-only mode in which hit information was preserved but the raw waveforms were discarded and a weeder-mode where the event builder discarded waveforms for which there was no TPC data or for which only one set of photon detectors had a signal above threshold. With these options, we were able to keep the data rate from the PDS from causing DAQ bottlenecks. The weeder-mode development required the time of a DAQ expert.

**Advice** Every group needs to be committed to optimizing the performance of the integrated system. Identification of issues like this one in advance was facilitated by two things in the 35-ton experience: joint meetings discussing the expected performance of the integrated system and individuals tasked with evaluating the integrated system for potential performance issues.

### 9.2.3 Data-taking Stability

**Contact** Mark Convery

**Item** The DAQ software was still undergoing development during the Cosmic Ray run. This led to some instability that required shifter intervention on a time scale of 30 to 60 minutes. This level of instability is not unexpected for a Cosmic Ray prototype, but would be unacceptable for a beam-based experiment

**Result** The data-taking efficiency was somewhat lower than ideal

**Advice** For ProtoDUNE, the DAQ software needs to be frozen well before the beginning of the beam run. Automatic procedures should be put into place to correct common crashes.

### 9.2.4 Configuration Management

**Contact** Mark Convery

**item** The configuration of the DAQ and trigger includes the following items:

- The hardware components to use in a particular run
- Some aspects of component configuration
- Trigger conditions and prescales

These were each handled in different ways, some of which required expertise that cannot be expected of a shifter.

**Result** Due to shifter errors, the DAQ was sometimes not in the desired configuration when data was taken.

**Advice** A unified Configuration Management scheme is needed that is operable by a shifter and is guaranteed to configure the DAQ properly.

### 9.2.5 Backpressure

**Contact** Giles Barr

**Item** Backpressure was omitted from the 35-ton DAQ software.

**Result** The DAQ had no way to relieve backpressure and would eventually crash if sufficient resources, such as disk-writing bandwidth, were not available. The mitigation was to lower the trigger rate manually, but this is not a good solution going forward.

**Advice** Implement a backpressure relief system for the ProtoDUNE DAQ. (This is already underway as a high priority)

## 9.3 Data Structure

### 9.3.1 Event Structure

**Contact** Michelle Stancari, Tom Junk, Matt Graham, Mark Convery

**item** The data format from the DAQ was inconvenient for nearline monitoring and downstream analysis. The event start and stop times were not synchronized with the triggers. Instead, the events started at uniformly-spaced times, with the intention of taking continuous data. Because of the DAQ bottleneck, we resorted to triggering on coincidences in the external counters, which enriched a sample of muons traversing the detector in interesting positions and directions. The data from events which were scheduled to be read out but which did not contain triggers was simply discarded, although the headers were retained, and the "event counts" of the output files thus are mostly empty headers without data payloads. The alignment of the data payloads from triggered interactions and the output events was uncorrelated. The ADC samples then in general needed to be collected together from neighboring DAQ events to make one consistent offline event. We had originally planned on dividing large DAQ events into smaller pieces centered on the interaction of interest, but the same input reformatter was able to stitch together neighboring events as well.

**Result** The main consequence is that no online filtering of events was possible, because the monitoring could only look at one DAQ event at a time, and the relevant drift window straddled two DAQ events most of the time. Online monitoring could not look at interaction-level quantities, like hit occupancy The data were sliced and stitched before being fed to the nearline monitoring, resulting in a 10-15 minute delay before being displayed. A significant portion of the DAQ file content ($\sim$20%) was empty headers because we triggered only at 1 Hz. The event reformatting (*slicing*) step was computationally cumbersome, due to the formatting of the data by channel instead of by time tick (see next lesson). The slowness of the reformatting of events caused us to make a second, reformatted copy of the data for subsequent analysis.

**Advice** Any DAQ system configured for continuous readout will have this issue. The pain can be minimized with an optimal DAQ event length. One may consider doing the reformatting in the event builder before writing the data to disk, although this will require ample memory. Alternately, the original events can be written to disk, but the reformatted events are passed to the online monitoring.

Having multiple parallel DAQ streams, each ending in separate file writes, has been proposed to alleviate the observed DAQ bottleneck. This will make reformatting in the event builder difficult if not impossible. Furthermore, the job of reformatting events after the file write step takes on an added dimension because the other half of the drift window for an interaction will not only be in another DAQ event, but likely also in different file. Protodune has chosen to employ triggered readout only, and therefore does not need to address this issue.

### 9.3.2 Data Format

**Contact** Tom Junk

**item** The data format was arranged by channel and not by time tick. A nanoslice consists of a group of ADC samples for 16 channels for a single time tick. The offline raw data format sorts the data the other way – all time samples for a single channel are laid out in a row in contiguous memory. While it isn't difficult to reformat the data, it takes time to compute the addresses to look up the ADC values in the online data format and copy the resulting values, and the fewer times this happens the better. More importantly, however, the online data format is inconsistent with zero-suppression in the input data, as some channels may be suppressed for a time tick that others are not. A new format more accommodating of zero-suppressed data – blocks with beginning and ending times for individual channels – was proposed but did not get commissioned for the 35-ton run. As the signal-to-noise ratio was low, it was not clear whether zero suppression would have worked, and was de-prioritized. We will need to write non-zero-suppressed data as well as some kind of compressed version for ProtoDUNE.

**Advice** Investigate the difference in computing time for both format styles. Consider changing the format of the DAQ files.

#### 9.3.3 Channel Map

**Contact** Tom Junk

**item** The channel map was applied at the slicing stage. This is natural, as the data format after the slicing stage matches that in Monte Carlo, and we didn't want software routines that knew whether they were looking at data or at MC processing the data.

**Result** Every time errors in the channel map were discovered, we had to re-make the sliced samples.

**Advice** Consider moving the online-offline channel map to a producer module in larsoft, especially if it becomes run-dependent. It may be more painful to move it than to reprocess data a few extra times.

#### 9.3.4 Emulators

**Item** Software emulators of different hardware allowed parallel development to occur

**Result** When hardware arrived later than expected, such as the PTB, it could still be integrated very quickly.

**Advice** Implement software emulators for hardware components when necessary.

## 10 Monitoring

Data Quality Monitoring is essential. Significant effort went into developing the DQM software and interfaces before the run. We learned enormous amounts about the prototype just from the online event displays and the histograms, and learned it quickly. As early as the DAQ vertical slice test, the online DQM plots were used for diagnostics and feedback. For future experiments, the online monitoring should be developed as soon as a DAQ vertical test stand is available.

#### 10.0.1 Online Event Display

**Contact** Michelle Stancari, Tom Junk

**Item** Despite the availability of a nearline event display with more information (both collection and induction plane wires and t0 corrected waveforms), shifters and experts alike almost exclusively used the online event display. This display was accessed through a web page rather than a larsoft job running in ROC west and could therefore be watched easily from remote locations.

**Advice** Online monitoring which is easily accessed will be used much more. A web interface to display the plots removes many barriers.

#### 10.0.2 Need for pedestal subtraction

**Contact** Michelle Stancari

**Item** The mean pedestal value, or baseline, varied channel-to-channel by as much as 200 ADC counts. With signals on the order of 50 counts, this means that the event displays and the nearline monitoring needed to do pedestal subtraction channel by channel in order to make meaningful, useful plots. The code to insert and retrieve run dependent pedestal values from the database was not ready immediately.

**Result** We were blind online and nearline for the first few days. Effort was spent on a temporary solution offline until the database interface was ready.

**Advice** If the large variation in baseline values persists in the next versions of the readout electronics, make sure that all the tools needed to determine pedestals on the fly and access the values online is in place before the run begins.

### 10.1 Nearline Monitoring

## 11 Commissioning and Operations

We estimate the critical mass of experts for 35-ton commissioning and operations requried to run smoothly to be a factor of 3 times more than the personpower that was available, and this factor does not include the sub-system experts who should also have been present.

**Item** There was not a single person in charge of *control room* commissioning, which includes DAQ, pedestals, detector monitoring, electronics, and trigger. Instead we had a rotating run coordinator who served for 1 week at a time. This was purely a personpower issue, and was recognized as not optimal.

**Result** Confusion over goals of a particular day. Commissioning efforts were not always organized or efficiently scheduled. Documentation was sparse and incomplete.

**Item** It was decided not to run owl shifts because of the staffing issues, especially because of the drain on local experts from being on call during the night.

**Result** This was absolutely the correct decision, even if it lengthened the time to take a given amount of data. Shifters were often not needed during working hours because the experts were debugging the system.

**Item** When the start of the data run was delayed (multiple times) the decision was made not to cancel or reschedule shifts for collaborators, but instead to use these shifters to help commission the DAQ.

**Result** Gave us a little more personpower when personpower was critical at the expense of a large author list.

## 12 External Scintillation Counters

**Contact** Michelle Stancari

**Item** The counter system had limited scope in the original 35-ton plans - provide the t0 truth for studying the time resolution of the photon detectors and triggers for commissioning/special studies. The counters were invaluable during commissioning in ways not originally envisioned. They allowed us to trigger and filter the data to identify muons tracks close to the wire planes. With poor purity, we still had quick confirmation of a working detector after turning on the recirculation system. Isolating a sample of muons traveling almost horizontally was trivial offline, and these set of rare muon tracks are very useful for commissioning. They were used to verify the channel map, to measure the purity in the TPC without track reconstruction (24 hour turnaround), to measure the drift velocity (confirm the drift field value) and to check the trigger/DAQ timing with muons tracks crossing the wire planes. Our offline reconstruction code needed to be tuned to deal with the high noise levels, and was not available until after the run. The position information from the counter hits is currently needed for track reconstruction.

**Advice** Muons traveling more horizontal than vertical (from one side of the cryostat to the opposite side) are rare but essential for commissioning the detector and characterizing its performance. A counter system segmented along the sides of the cryostat, especially segmented along the drift direction, with the individual counter hit times recorded in the data stream facilitates both the acquisition and reconstruction of a useful cosmic ray data sample.

# 13 Software and Computing

**Contact** Tom Junk

**Item** It takes many different computing skills, authorities, and access to resources to make an experiment or event a prototype of the scale of the 35-ton prototype to work. Not only did the DAQ groups need to supply equipment and expertise in order to manage the computers, they had to be integrated into both the environments at DAB and PC-4. Rack space, power, cooling, and connections had to be established, not only for the computers providing DAQ functions but also those needed to interface with the Fermilab network and switches. This work was underestimated/missed in the planning of the DAQ group. Alan and Tom were left to find personnel outside the DAQ group to identify and locate or procure missing hardware, install it, commission it, test it, and maintain it. Software components also needed to be installed, integrated, and monitored.

**Result** Schedule delays, no DAQ system for noise monitoring at the beginning of installation

**Advice** Most specialized skills are needed for only short amounts of time. One-time (or at least infrequent) setup of networks and software require the right people with the skills, equipment, and authority to do the job, but many setup jobs only take a few hours or days to complete. One therefore needs a smaller number of people dedicated to the experiment who can "do anything" or at least identify what the problems are, and a much larger resource pool to draw from for short-term projects that fit their expertise.

**Advice** Fermilab has the right staff with the right skills and equipment to do the job. We were pleased with the efforts on our behalf extended by the Scientific Computing Division (SCD) to help us set up our DAQ and online computers and associated software. Fermilab's infrastructure and tools – artdaq, FTS, dCache, SAM, enstore, databases, electronic logbook, art, LArSoft, build tools, and the redmine repository all made the job much easier. We did have to manage our own custom pieces and plug-in components contributing to each of these. Learning about how to use an existing tool is much easier than re-inventing the tool. Fortunately the tools provided by Fermilab are highly aligned with our goals.

**Advice** Sometimes redesigning and re-implementing software to make a 35-ton-specific tool is faster and more convenient than re-using a much larger tool from elsewhere. For example, re-using an LHC data transfer system instead of writing our own would have slowed us down.

**Advice** While many resources are present at Fermilab, requests are not always acted upon immediately. Time must be allocated for requests to be completed, and sometimes service providers must be reminded of the urgency of getting some tasks done. Especially for larger tasks, meetings may need to be set up in order to get something done, even if it is just installing software.

**Item** Spare computers left over from Tevatron Run II let us meet our goals, but they had reliability problems. Specifically, approximately half of the disks on lbne35t-gateway02 needed replacement in the Summer of 2015 before installation of the computer at DAB. The newer computers supplied by the UK institutions did not exhibit hardware failures. We found we can live with computers with reduced reliability for the 35-ton prototype run as long as we have ready spares that can be turned on rapidly. We knew that lbne35t-gateway02 was not as reliable as we would have liked, and identified a similarly capable gateway machine, lbne35t-gateway04 as its spare. When lbne35t-gateway02's RAID controller failed, we were in the process of getting lbne35t-gateway04 configured, but it was not complete. The repair of lbne35t-gateway02 was quicker than finishing the configuration of lbne35t-gateway04. Had we been in a beam experiment, we would have lost two days' worth of data, and more if we had to wait for lbne35t-gateway04 to be configured.

**Advice** Extended warranty coverage on computing hardware is useful. Service technicians from the vendor come out more readily – less paperwork – when the computer is under warranty. Hot spares should be in place if downtime is costly.

**Advice** Geographically dispersed groups can work together, but members need to meet and communicate frequently. Misunderstandings arise when communication is inadequate.

**Advice** Time-zone differences, and holiday schedule mismatches, can introduce delays. These did not cause trouble for the 35-ton prototype computing, but one must watch out for ProtoDUNE which may have a tighter schedule.

**Item** Having two system administration domains (some nodes were administered by the UK groups and others by Fermilab) introduced complexity and delays. The lbnedaq*.fnal.gov machines were not Kerberized and thus could not be visible on the Fermilab network. This initially required two transfers of the data – first to a gateway machine and then to dCache. A technical solution in which the non-Kerberized machines would be protected by an Access Control List in the switch allowed a more direct route while maintaining security, but it was introduced after a time and required meetings and arranging for the right people to help out and do the work. There was some wariness of installing applications on the daq nodes that might slow down their performance. At the very least, file copy services needed to be run. The UK groups did not allow installation of TotalView on the daq nodes.

**Result** Fermilab personnel had to be pulled in to reconfigure the DAQ machines in a way that conformed to FNAL security policies and facilitated faster file transfer. Asking them to support a system that they did not design was tricky.

**Advice** It is best to have a spectrum of options – Plans B, C, D ... in case Plan A either fails or does not perform adequately. An example is that transferring files directly from lbnedaq6 to dCache instead of being first copied to the gateway computer(s) constituted a backup plan in case the disk writes were slow on the gateway computers. We would have lost the ability to run near-line monitoring programs on those files and would still need a subset of those files. Another fallback option in case there was a bottleneck on one of the gateway computers (likely due to slow disk writes) was to run file transfers in parallel using both gateway computers. Had we needed to compress files locally, this may have been more attractive. The networking and computing setup should be reviewed

**item** Many operational issues get ironed out after the first encounter with them, including ones that are foreseen but not necessarily prioritized. Power cuts exercise many rare or at least infrequent cases. An example of this was our FTS instance not restarting after the power cut.

**Item** The scratch pools in dCache were frequently very busy. Writing files to them could be slow at times, and there was a large dispersion in the times required to write to the scratch areas. Using other areas in dCache improved the consistency of transfer speed.

**item** dCache reliability was not great. There was at least one planned outage per month, and additionally more unplanned operational glitches.

**item** The checksum check after the first file copy from lbnedaq6 to lbne35t-gateway02 was sped up by reading the copy in the disk cache memory. This checks the network integrity but not the disk write and readback integrity.

**Advice** It is useful to check every step of file copying with a checksum. The most common failure producing corrupt data files was an outage that started during the transfer of a file, resulting in a partial file on the receiving end. SAM caught these quickly in that the size mismatched. We checked the file sizes and checksums when transferring each file from lbnedaq6 to lbne35t-gateway02 and re-initiated the transfer if a mismatch was detected. But we initially lacked a machine that performed checksums after the copy to dCache over NFS. When we started using FTS, there was a checksum provided by SAM.

**Item** File writes are much slower ($\times 4$) than file reads.

**Advice** In a system bottlenecked by the outgoing transfer network speed, it may make more sense to simply take more data than to try to recover a file corrupted by a failed transfer. The 35-ton prototype was bottlenecked, but not by file transfer, but instead upstream with the artdaq step. We never were in this situation.

**Item** The artdaq crash rate was quite high – we lost a significant fraction of data ($\sim$10%) due to unclosed ROOT files. Too-small files induce another penalty – closing a file and opening a new one required starting a new run, which required re-initializing the RCE's, which introduced deadtime.

**Advice** File size is important. If files are too large, transfers time out, and more data are lost in case of a failure. ROOT files that are not properly closed due to the application writing them crashing were not recoverable. We should look into ways of mitigating this. The art group has this on its list of items to work on and has asked for our timescale for ProtoDUNE in order to schedule this work.

**Item** Our data-acquisition system crashed with trigger rates above 1 Hz. This was due to inadequate buffers available for the fragment generator. John Freeman suspects the bottleneck was the 1 GBit bandwidth of the local network inside the air-conditioned rack. Even though the outgoing bandwidth was 1 GBit, having a higher bandwidth for communication between DAQ computers would have allowed us to take data at a faster rate than we did. Upgrading the internal network speed in the local network would have required new NIC cards in each computer as well as a new switch.

**Item** The DAQ system was capable of writing files much larger – we got one in excess of 700 GBytes. It could not be transferred using FTS due to timeouts. It also does not fit in anyone's quota on BlueArc, and would not fit on a batch worker.

**Advice** We should put in place a policy that files larger than a preset size will not be handled automatically but will require human intervention to split up or discard. 2-5 GBytes is a good size for a raw data file.

**item** The compression performance for 35-ton data using the gzip libraries in ROOT, and also the gzip utility itself, amounted to approximately a factor of two reduction in data size, almost independent of the requested compression level. The CPU required to compress the data was significant. About 3.5 minutes of CPU are required to compress a 10 GB file with compression level 1, and compression level 6 takes about 17.5 minutes. CPU's are not getting faster very quickly, but more cores are being packed into servers. We probably had enough CPU and disk I/O bandwidth between lbne35t-gateway02 and lbne35t-gateway04 to run compression on the output files. As it is, the bottleneck was not the 1 GBit network out of PC-4, but rather was upstream, so there was no pressing need to compress the data. Compressing the data in the artdaq logger process would have worsened the backpressure issue and reduced the trigger rate. Storing uncompressed raw data on tape is not the most efficient use of the tape, but the next step, *sliced* data, was compressed offline where the CPU and I/O resources are much greater. Only about 70 TB of raw data were written (including DAQ tests before the actual run), and so the tape needs were not magnified much by failure to compress the data.

**Advice** The ProtoDUNE datasets are anticipated to be a factor of ten to thirty larger, and it may be a better use of resources to compress the data on the online cluster.

**item** The channel map was difficult to get right and to verify. The map in LBNE DocDB 10145 from the APA installation took time to digest and also contained some errors, such as exchanging every other collection wire. There is a high degree of symmetry in the system – the APA's can be installed frontwards or backwards equally well, and there were 16! ways of assigning RCE's to the APA's, so the initial channel map files still had another step to do before they could be used offline. Since the APA's have asymmetries in their broken wires, the as-installed positions and labels need to be made clear. Validation with data showed that we had several errors in the channel map, and some that were very difficult to disentangle with the data. The wrapping of the wires makes checking the induction-plane channel map more difficult.

**item** Keeping expertise in data acquisition development is also important. Lessons learned about the artdaq internals are bound to the relevant experts.

**Item** Camera data needs to be put in a format that can be analyzed and correlated with other data.

**Item** Data in iFIX (slow controls and monitoring) also needs to be exported to be useful offline.

**item** It is important to keep good track of run configuration parameters for each run.

**item** Timestamp formats varied from subsystem to subsystem. This was addressed with software, but did take time to understand and verify.

**item** File cataloging is important, as well as distributing instructions to collaborators on access to the files. We are still learning about how to arrange the good-run list as different data collection periods are good for different analyses. The individual analyses themselves may rely on small amounts of data targeted at low-noise running with specific electron lifetime, hardware performance, and high-voltage constraints.

**item** We used very little of the LBNE metadata specification. The run number, the start and stop times, the file size, the number of events, the subdetectors contributing (which RCE's, which SSP's, and whether the Penn trigger board contributed) were logged as metadata, to be used in accessing files later. This was expected, as many metadata fields previously envisioned were relevant only for Monte Carlo or for other DUNE detectors.

# 14 Reconstruction Software and Data Analysis

The 35-ton represented the first opportunity for the DUNE collaboration to analyze real data. As such, it has inspired a significant amount of work on DUNE-specific issues, such as wrapped-wire disambiguation and stuck code mitigation.

Much of the software developed will be used by future experiments, such as ProtoDUNE-SP.

### 14.1 Stuck Code Mitigation

**Contact** Tom Junk

**item** The stuck codes could be mitigated by interpolating ADC samples at nearby times. If there were too many stuck codes in a row, then this interpolation becomes less accurate. Here mitigation means that hit finding was possible. The correlated noise could be mitigated by subtracting off the median of the difference between neighboring channels' ADC values at a particular ADC tick and the corresponding pedestal. Not all of the noise was correlated, and much of the noise was left. Some of the noise may still have been correlated and simply out of phase however. The correlated noise subtracter has the effect of also subtracting off signals that arrive on neighboring wires at the same time, that is, from tracks that travel perpendicular to the electric field (parallel to the APA planes).

**Advice** Dedicated studies to understand the impact of stuck codes and correlated noise subtraction on energy resolution need to be performed.

### 14.2 Ambiguity Resoluation

**Contact** Tom Junk

**Item** The 44.3° and 45.7° angles of the $U$ and $V$ planes relied on hits being detected in all three views in order to break the wire-wrapping ambiguity. With just two views, the ambiguity is unbreakable, unless a cluster reaches the end of the APA or a wrapping boundary. The issue is that with the low signal-to-noise ratio observed in the 35-ton prototype, it is difficult to see tracks in all three views simultaneously. If a track travels perpendicular to the wires of a plane and the electric field, then a minimum amount of charge is available for detection on the wires of that plane. If a track travels nearly parallel to the wires in a plane and perpendicular to the electric field, it leaves a much larger signal in a short amount of time on fewer wires, which is more likely to stand out above the noise. The issue is that tracks cannot be oriented so that they leave large signals on each wire in each of the planes – at least one of the planes is left with a small signal on each wire. This makes breaking ambiguities difficult or impossible, and the physics measurements on such events will have to be made with the data collected from two views and the ambiguities broken with the information from the external counters.

**Item** Dead channels also contribute to ambiguity-resolution difficulties. If an ASIC is dead (due to too many power cycles for example), many contiguous channels are lost, and so entire clusters of hits become visible in only two views, even under the best signal-to-noise ratio conditions.

### 14.3 Recognition of Contributions

**Contact** Tom Junk

**item** Keeping analyzers engaged after the run is a difficult but important task. It is easier to get collaborators to commit to a week of shifts than a more lengthy analysis task. Writing reconstruction software (or any software) must be rewarded. Publications, authorship, conference talks, leadership roles, and career advancement are the standard rewards.

## 15 Appendix: Electronics Problems without Current Solutions

During the operation of the 35-ton prototype, a number for problems, in addition to the unexpectedly high noise, were observed with the Cold Electronics that have not, as of June 2, 2016, been understood or corrected.

As such, they cannot be classified as "lessons learned". Rather, we list them here in hopes that their causes and fixes will be found.

## 15.1 High Noise State

Soon after the LAr fill was completed on February 2, the TPC electronics was observed to enter a "High Noise State" (HNS). In this state, the typical ADC RMS of every active TPC channel was more than 100, as compared to the "normal" state where it was close to 10. During the operation of the 35-ton, the following observations were made about this state:

- The electronics would sometimes enter this state seemingly spontaneously. It is possible that an external trigger of some kind initiated the HNS, but these were not reliably identified.
- The HNS would typically persist for several hours, but would usually eventually stop by itself and the electronics return to the "normal" noise state
- The HNS could usually be stopped by toggling the LV power to one or more APAs
- During data-taking, we typically waited for HNS to "fix itself" due to fears of damaging the CE by too much power toggling.
- The "three board" APA (offline APA1, RCEs 12-15) seemed to be related to the HNS. When its LV was turned off, the HNS occurred much less frequently. For this reason, much of our data was taken with APA1 turned off.

## 15.2 Non-noise Electronics Problems

### 15.2.1 Stuck Code

The "stuck code" problem and its software mitigation has been well documented elsewhere. We include it here for completeness.

### 15.2.2 Initial Dead Channels

A number ( 2%) of isolated dead channels were present from the beginning. These were very likely observed during pre-installation testing at BNL.

### 15.2.3 Additional Dead Channels After Cooldown

During the initial warm testing of the 16 FEMBs installed on the 35-ton APAs, all FEMBs were functional and only isolated dead channels were observed. During the initial cooldown of the 35-ton, channels 64 to 127 of FEMB 3 were observed to "die". Their pedestal ADC mean values and RMS values dropped to zero. This would seem to indicate a problem with the digital part of the FEMB. Subsequently, over the course of several weeks, 5 more groups of 16 channels each were observed to exhibit similar behavior. However, in this case, the pedestal ADC mean remained at normal values, and the RMS value dropped very low but not zero. This would seem to indicate a problem with the analog part of the FEMB, perhaps related to the FE analog ASIC. These two problems led to 208 dead channels. In total, there were 275 of the 2048 channels that were dead for the whole run.

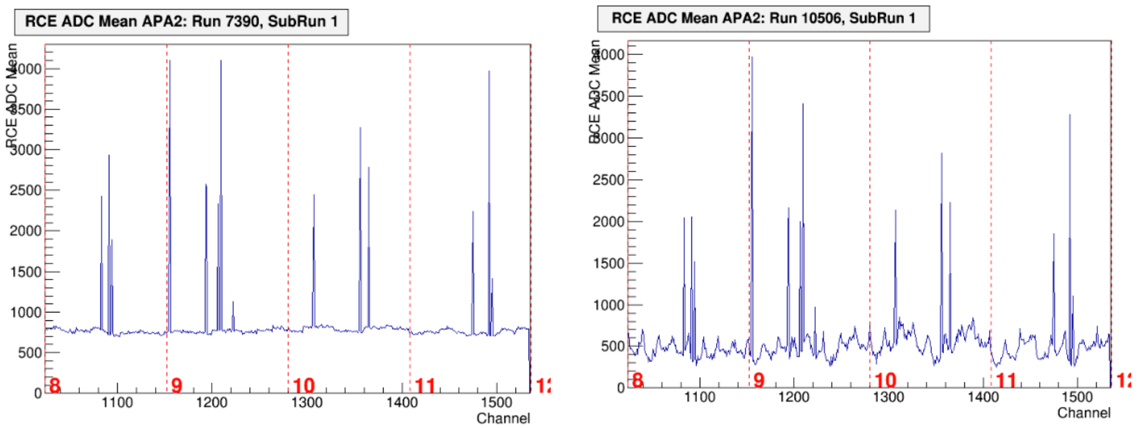


Figure 3: ADC mean values when warm (left) and cold (right).

#### 15.2.4 Dead FEMBs

One of the stranger observed behaviors of the FEMB occurred on February 22 when the LV power supplies for FEMBs 9 and 10 became current limited - indicating a likely short circuit. Several checks were done to confirm that this short was inside the cryostat and not related to the power supply or external cabling. After February 22, we turned off these LV supplies and ran without them.

However, after the March 4 site-wide power outage, the LV power supplies were found to be no long current limited, and the FEMBs became functional again. How the power outage could cause this is a mystery. Inspection of the boards after warm-up may yield some answers.

#### 15.2.5 Pedestal Variation when Cold

As shown in Figure 3 (left-hand side), the warm pedestals were relatively uniform except for spikes due to dead channels. In contrast, as shown in the right-hand side of the figure, the cold ADC means displayed a wide variation. The "wavy" pattern follows the boundaries of the FE ASICs and may indicate a problem with power distribution, or other channel-dependent effects.

The variation of the pedestal or baseline channel to channel was much larger than naively expected by the online/offline crew. A histogram of these values at room temperature for functioning channels is shown in the left plot of figure 4 and for liquid argon temperature on the right plot.

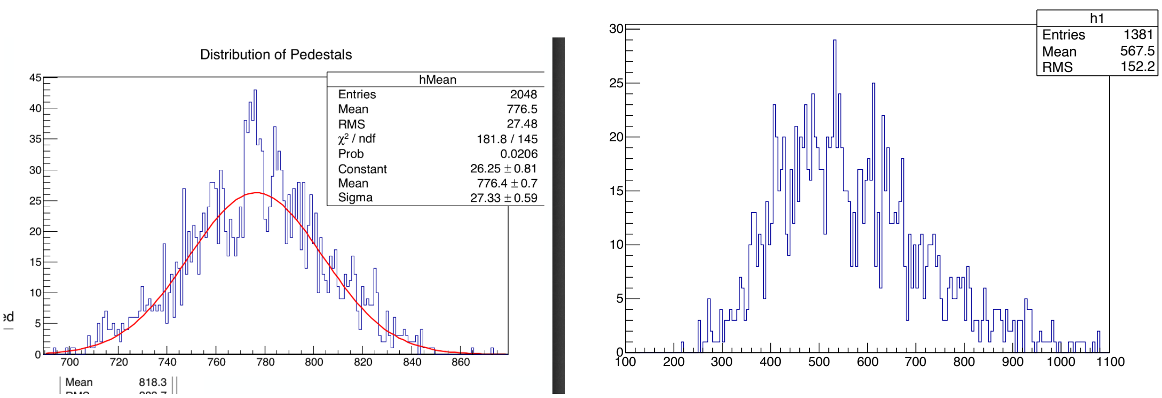


Figure 4: The distribution of pedestal means, or baseline, for functional channels at room temperature (left) and liquid argon temperature (right).