DUNE SINGLE PHASE DAQ DESIGN CURRENT STATUS: 30 OCT 2016

This is a summary of the current status of the DAQ design for ProtoDUNE single phase, and also a guideline to the documentation we have created ahead of our design review on November 3-4. Time is tight in the ProtoDUNE schedule and the design is still evolving in a few areas, so we are approaching the review in a way that is common with other reviews by focussing one talk on each of the main sub-system areas by a member of that implementation team. The advantage is that this gives a good view of each corner of the design, and how that part of the team sees it interacting with the whole. The disadvantage is that we are addressing aspects of the charge questions spread among the talks. To facilitate understanding the charge, we list the questions, along with a brief description of how we think we are meeting them and pointers to where in the documentation that is shown here.

INTRODUCTION

The DAQ is needed for NP04 ProtoDUNE detector operation at the CERN Neutrino Platform in 2018. It is critical for us to be on time, and to minimise the times when the DAQ development is taking centre stage, for the most part, we must work in parallel with the provision of the other bits of ProtoDUNE. It is not critical for the DAQ to be the same as the final DAQ for DUNE in 2021/2, indeed, it is better to delay as long as possible in some respects to take account of higher performance solutions available with more modern commercial products. Also, we have ways of testing the final DAQ solution in the lab with test data sources and so are not reliant on ProtoDUNE for that (although parts we do demonstrate is an advantage).

Part of the design comes from the 35t LBNE prototype DAQ, but a lot of it is new and provides: steps to making a more robust system, provision for taking beam data, adapting to improvements in the upstream TPC electronics, providing hooks to allow new ideas to be tried or a combination of these. Some of the legacy system from the 35t is helpful in the short-term as we can use the existing system immediately to test improvements step-by-step rather than waiting for complete new systems to be ready. This is important because of the short timescale.

There are places where ProtoDUNE has special needs that are different to DUNE due to the high arrival rate of particles in the test beam in comparison to the underground detectors where there is one cosmic ray muon per minute. In these cases, we have mainly added features to the design following the classical fixed-target experiment paradigm. Examples are beam trigger addition, changing the records in the output files that are normally referred to as events from ‘time blocks’ to ‘event triggers’, adding a 100% robust backpressure system.

REVIEW CHARGE QUESTIONS (AND STATUS OF ANSWERS)

QUESTION 1

Does the DAQ system design meet the science and engineering requirements for NP04? Does the design provide sufficient flexibility for alternates? Are the science and engineering requirements/justifications sufficiently complete and clear?

The science requirements are laid out in M. Thomson's introductory talk for the review. The table in his talk (which is from the TDR) gives the main data rate, readout window and compression factor expected and these are the key requirements for the DAQ. The definition of which triggers are good, and the requirements on the rejection rate (selection performance) of the experiment are not in the scope of the DAQ group (the beam instrumentation group are directing the implementation of that logic based on the studies from the beam-physics working group). The introduction of the FELIX system (described in the introduction and the FELIX talks in the review) is an example where flexibility for alternates is provided. By modifying firmware in the RCEs and/or SSPs and timing, we can provide hooks for other aspects that are important for the final design. The justification is a balance between doing something that is on-time, not expensive and fulfils the data collection requirements of the main goals of ProtoDUNE (that are outlined in Mark Thomson's talk) and doing something that showcases the potential DAQ ideas for DUNE (that can otherwise be demonstrated in the lab later).

QUESTION 2

Are DAQ system risks captured and is there a plan for managing and mitigating these risks?

The risks are identified in the review documentation that captures risks that are perceived from this point onwards. The risks are also called out in the individual talks. They are mainly worries of failure to capture an interface completely, requiring rework, delays in external equipment arriving (people standing around and needing to return to CERN), and resources (travel money, people being moved between projects etc.). We believe they are mainly under control, that we have some plan-Bs and that there is still flexibility in running more activities in parallel to mitigate some delays. The risks are not yet managed other than by the DAQ leadership providing a response when a delay or other problem occurs.

There are several risks that were identified early on, where mitigation solutions have been implemented to our satisfaction already. One example was to spread the knowledge and responsibility in the leadership among many people so that the absence of a single person would not cause scheduling delays. Another was to ensure that there is a constant level of ground-personnel at CERN to avoid delays from unanticipated things becoming long.

A register of the risks have been described in the document attached to Giovanna Lehmann Miotto's talk.

QUESTION 3

Does the design lead to a reasonable production schedule, including QA, installation and commissioning? Does the DAQ schedule allow sufficient time for testing of other components?

The lead-in between the plans of the DAQ group and the needs of the other parts of the experiment is addressed mainly in Karol Hennessey's talk. The status and adaptation of this schedule is tracked (docdb-1548), a link to it is provided in the review documents. It is divided into sections, to focus on the external milestones from the DAQ group within ProtoDUNE NP04. Within these sections, we track expected hardware delivery dates, test outcomes and other major planned activities. The overall schedule milestones are largely based on the ends of each of these sections of activities.

QUESTION 4

Does the documentation of the DAQ system technical design provide sufficiently comprehensive analysis and justification for the design adopted?

An introduction to the DAQ design is in the TDR. The justifications for the main design choices are in the introduction to this document, in Giovanna Lehman Miotto's talk and documentation and in other talks.

QUESTION 5

Is the DAQ system scope well defined and complete? Are all interfaces to other systems: Cold Electronics, Computing, EOS, beam instrumentation and Photon Detector systems documented, clearly identified and complete? Is the cabling and power well defined and understood? If any parts of the DAQ design impact the grounding and shielding are they understood and adequate?

System scope was defined by Eric James in the 'internal proposals' July 2016 to be "Includes Optical fibers from Warm Interface Boards (WIBs) to RCE boards; RCE boards; COBs; aTCA crates; cables from RCE boards to event builder (EVB) computing infrastructure; EVB computing infrastructure; timing distribution board; cables from timing distribution board to WIBs, RCEs, SSPs and EVB; RCE board firmware; EVB software, Run Control (RC) software, and Data Quality Management (DQM) software. Also trigger logic and the means of distributing trigger signals to RCEs, SSPs and other readout boards (but not the trigger counters themselves, or the readout electronics for using their hits offline)".

The Ethernet switch that aggregates the SSPs is in the Photon Detector scope so compatibility with the SSP Ethernet interface can be tested and so that grounding/noise is considered along with the rest of the photon detector installation.

It is unclear at present how much of the scope for the calibration lies with the DAQ group, however this is not crucial yet. Calibration for the SSPs follows the system in use in the 35t and pulser calibration for the TPCs is handled by the electronics group. There is no laser system in ProtoDUNE. The DAQ group will implement the control/status of the calibration hardware and the incorporation of the calibration run modes within the run control software.

Interface documentation (docdb entries attached to Giovanna Lehmann Miotto's talk):

1. To TPC Cold Electronics: (docdb-415, docdb-1701, docdb-1394)

2. To Computing, EOS: the interface is a files written to disk in an agreed directory structure. Additional metadata to be passed to the offline computing will be determined in due course. Data will be transferred to EOS using the tools described in docdb-1212.
3. To Beam instrumentation: Paola Sala's talk documentation
4. Cosmic Ray Tagger: Camillo Mariani's talk
5. To Photon Detectors: Martin Haigh's talk documentation , docdb-1571, docdb-910
6. Cabling, power, grounding: though not fully defined, the rack space and power budget in the counting room dedicated to DAQ has been agreed upon (see Geoff Savage's talk); any grounding issue is eliminated be ensuring that any connection between the DAQ and the detector is optical;

QUESTION 6

Is the software architecture suitable, including Event Builder, Run Control, Online Monitoring, Timing, Triggering and Databases? Are there sufficient resources for the required software effort?

Kurt Biery and Wes Ketchum's talks. As highlighted in Karol Hennessey's talk, more effort is needed for the implementation of data monitoring: the infrastructure to provide events to monitoring processes is in place. We think that most of the architecture is tested from the 35t, and that the shortcomings that were identified have been addressed (Kurt's talk). The run control is a standard CERN package and is different from the 35t. The 35t run control is being maintained as long as needed to allow hardware testing and integration with artDAQ to proceed in parallel with the interfacing of the CERN run control with artDAQ. We do not yet have sufficient resources for the online monitoring, in particular the part needed for the early APA tests.

QUESTION 7

Are the DAQ specifications of commercial units and design drawings/part-lists of custom hardware sufficiently complete to demonstrate that the design can be constructed, installed and operated safely and efficiently?

As indicated in each individual component talk, the elements composing custom hardware are well defined. The server characteristics for data flow, control and monitoring are not yet fully specified. We have extensive experience of the commercial computers similar to the ones proposed. Procurement of the computers will be at the very start of the next UK financial year (April 2017), which is why we have not finalised the parts list yet. Regarding network equipment, we will remain aligned with the devices in use at CERN for both central IT and most of the experiments, since this provides on site support (spares as well as guidance for installation and administration).

QUESTION 8

Are operation conditions listed, understood and comprehensive? Are interfaces to calibration systems and plans well understood? Are proposed triggering schemes sufficiently well understood? Has appropriate consideration been made for collection of both zero suppressed and non-zero suppressed data?

As outlined in several of the talks, there is one main operating mode, in which data are collected at a high rate within one 4.8s SPS spill and transferred out also in the interspill gap. For the TPC data, the buffering is provided in the RCEs, FELIX host and board readers. The RCE applies trigger cuts and compression as the data is written to the buffer (real time). The FELIX system writes the data in full real time to the host memory and sends it to the boardreader computers to apply the trigger selection and compression. Both systems have sufficiently large buffers to be able to accumulate data in excess of the instantaneous output capacity and profit from the inter-spill time to stream out data. Likewise for the SSPs, data accumulation and buffering is performed in the board readers.

See comment about calibration systems in answer to question 5 above.

See comment about trigger design in answer to question 1 above.

We have decided in discussion with management that all the data will be collected non-zero suppressed. The only data reduction to be applied is to (a) use the trigger to select full-drift times where a desirable beam particle is present and (b) to apply lossless compression. Functionality will be implemented to remove the bottom ADC bits in a lossy fashion for high-noise channels (experience learned from 35t).

To collect long periods of consecutive non-zero suppressed data, a chain of events can be defined. Up to about 500ms of data can be collected in this way by essentially coalescing all the triggers from one 4.8s spill together.

Another running mode can be associated to cosmic data taking (when no beam is available). In this case an appropriate trigger rate will be set, allowing for continuous data taking without using a spill/inter-spill structure.

Data taking with beam and cosmic data taking may be superimposed, by allowing a low rate of cosmic triggers to fire in the inter-spill time.

QUESTION 9

Are the DAQ system analyses sufficiently comprehensive for safe handling, installation and operation at the CERN Neutrino Platform? Is the installation plan sufficiently well developed?

See Geoff Savage's talk. It should be noted that all DAQ equipment will be housed in a counting room, off-detector and that thus handling is simpler than for other ProtoDUNE systems.

QUESTION 10

Have applicable lessons-learned from previous LArTPC devices been documented and implemented into the QA plan? Are the DAQ quality control test plans and inspection regimes sufficiently comprehensive to assure efficient commissioning and adequate operational performance of the NP04 experiment?

Previous LArTPC devices include argoNEUT, MicroBooNE, ICARUS, the 35t-LBNE prototype and the ongoing 3x1x1 dual phase prototype. We have many of the implementers of the MicroBooNE DAQ on ProtoDUNE (Georgia Karagiorgi, Wes Ketchum, Eric Church). There was a review of lessons learned from the 35t closeout document at (docdb-1315 linked to Karol Hennessy's presentation). The ways in which the lessons-learned have factored into our plans are in Karol Hennessey's talk.   We still need more online monitoring effort especially for the early part of the testing phase.