Close-Out

Gustaaf Brooijmans, Leslie Camilleri, Philippe Farthouat, Sigve Haug, Niko Neufeld, Andrew Norman, Wainer Vandelli, Jos Vermeulen

ProtoDUNE DAQ Design Review

November 3-4 2016

Charge

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?
2. Are DAQ system risks captured and is there a plan for managing and mitigating these risks?
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?
4. Does the documentation of the DAQ system technical design provide sufficiently comprehensive analysis and justification for the design adopted?
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?
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?
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?
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?
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?
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?

Reviewers

- Gustaaf Brooijmans (Columbia)
- Leslie Camilleri (Columbia)
- Philippe Farthouat (CERN)
- Sigve Haug (BERN)
- Niko Neufeld (CERN)
- Andrew Norman (Fermilab)
- Wainer Vandelli (CERN)
- Jos Vermeulen (NIKHEF)

Architecture, System Testing and Exploitation

- Question: 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?
- Yes. The system presented will meet the needs of the experiment to accumulate beam data at a rate of 25 Hz under the given assumptions of noise rates and data compression ratios.
- Yes. The design presented is modular and allows for alternative systems or technologies to be applied for the basic readout of the TPC. There are other portions of the design which were not presented with alternative designs but were shown to be in a very advanced state and consistent with a low technical risk.
- Yes. The science and engineering goals are clear and well defined. However the timelines for completing the science goals within the aggressive timeline that has been presented and in the context of the hard cut off of the LS2 date, present a significant risk if the commissioning work leading up to the proposed 120 day run period were to be delayed or were to run into unforeseen difficulties (e.g. noise rates) which had to be addressed to meet the operational assumptions that were presented.

Architecture, System Testing and Exploitation

Findings:

- The DAQ system has been designed, and the network and other communications bandwidths sized to support a baseline acquisition rate of 25 Hz under the assumption of a factor 4 reduction in data size from the use of loss-less compression.
- The actual compression ratio that is achievable is dependent on the noise levels (coherent or incoherent) that are present in the final detector setup.
- Compression factors better than 2 have been demonstrated on the 35t prototype which had unforeseen sources of noise.
- The network bandwidth required to operate the system is estimated at 14 Gb/s with an available bandwidth of 20 Gb/s to the ENH1 site.
- There are plans for a full vertical slice/integration test bed.

Architecture, System Testing and Exploitation

Comments:

- Because the compression ratios are highly dependent on the structure of noise in the system, it will be advantageous to have measurements of any noise that is either intrinsically present in the production APAs or in their specific installation in the ENH1 facility. Early characterization of the noise sources will allow for better determination of the final compression factors that are achievable and set the overall trigger rates that the experiment can be run at.
- The current network bandwidth utilization is at approximately 70% of the available. This may not leave sufficient head room for additional network traffic that protoDUNE will need during commissioning for other tasks (monitoring, prompt processing, recovery from outages, etc…) It may be possible to mitigate this limitation by increasing the bandwidth to EHN1.
- The vertical slice has several critical path items and potentially long lead time procurements. These may affect the overall schedule for the availability of first vertical slice.
- The design allows for some alternates. Up to the COB continuous read out is in principle possible. The design includes use of FELIX for one APA, however, ACTA COB read out for all 6 APAs is available. Downstream there is redundancy depending on network (10 Gb/s) and server capacity.

Architecture, System Testing and Exploitation

Recommendations:

- The experiment should investigate the possibility of accelerating the procurement/shipping of the COTS computing (or obtaining analogous systems) in order for the core components of the DAQ readout systems to be ready ahead of the integration and testing of components for the full vertical slice. This may help in mitigating time pressures that would other arise in the context of any delays in the availability of the ENH1 space or other critical path items that may otherwise delay the vertical slice.

Architecture, System Testing and Exploitation

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

- Yes. Risks for each of the subsystems were presented and were reasonable. Many of these risks are collected and managed centrally in a risk registry. The protoDUNE project should re-examine the cataloged risks and add to them any risks that were identified in this review. In particular some of the risks have that have been identified should be mitigated and managed and have their status tracked by the project. It should be made clear through the organizational structure of the DUNE and protoDUNE organizations who holds the different risks and how they are mapped onto the current working groups of experiments.

Architecture, System Testing and Exploitation

FINDINGS:

- Risks for each subsystem were presented with the exception of the cosmic ray tagger which presented that they had identified no risks.

COMMENTS:

- The presented risks seemed reasonable and inline with the presented design for the DAQ system.
- The lack of an enumerated risks in the integration of the CRT with the rest of the protoDUNE DAQ and experimental setup is probably due to the extremely recent nature of the decision to include the CRT in the system.
- A major risk is the EHN1 schedule. A decision on preinstall in a different location is to be taken in January. The team should not hesitate to do this if longer delays in EHN1 are expected by then.

RECOMMENDATIONS:

- Review the risk registry and evaluate the status and validity of currently held risks as they pertain to the protoDUNE experiment. Add to the formal risk registry major risks that were identified during this review and determined to be appropriate for inclusion in the DUNE/ protoDUNE registry.

- When a formal design for the integration of the CRT with the rest of the protoDUNE experiment and DAQ is determined, the risks should be evaluated and included in the formal risk registry if appropriate.

Fermilab DUNE

Architecture, System Testing and Exploitation

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?

- Yes. The design that was presented is very modular and will lend itself to a well defined production, QA and installation plan. This plan was presented in detail for in 2017. The timelines for 2018 leading up to actual detector commissioning are not as developed.
- Yes. Most of the subsystems appear to have adequate time in the base schedule for testing. It may be beneficial to accelerate these schedules to mitigate unforeseen delays due to the hard deadline imposed by the long shutdown.

Architecture, System Testing and Exploitation

FINDINGS:

- There is a detailed WBS that outlines the schedule for component delivery, acceptance and integration.
- The subsystems are modular and in many cases provided for discrete testing (i.e. without dependent systems)

COMMENTS:

- The timelines are tight and there is a hard deadline with the long shutdown.
- The experiment may want to consider revisiting the WBS to see if there are areas where it can be accelerated to provide more schedule contingency. This is especially true for the activities in 2018 which were not as detailed in the presented plans.

RECOMMENDATIONS:

- None

Architecture, System Testing and Exploitation

- Question 5: Is the DAQ system scope well-defined and complete?
- Findings:

The engineering requirements are partly clear. About 470 Gb/s in total from 6 APAs and some data from photon detection system and beam signals. Up to the ACTA crates it is clear. The requirements for the online cluster are not clear. Disk buffer size (about 300 TB) and upload link to T0 are clear. Justifications are clear up to the ACTA crates.

- Comments:

The online cluster size and performance requirements are crucial for defining cooling and power and rack capacity. It is not completely evident that three plus one spare rack will be enough. Numbers like about 15 servers for board readers and event builders are there. From overview diagram it is not clear what is needed.

The 20 Gb/s uplink to T0 is close to writing speed (10-24 Gb/s). Recovery of a filled online buffer may be tight.

- Recommendations:
  - Update main reference drawing of system with consistent bandwidth links, CPU, RAM and writing and network numbers, server numbers. Define the downstream end of DAQ system, e.g. stops with the writing and FTS and no re-read and write of files on online buffer. (see also data flow recommendation) High priority, within a month.
  - Check the possibility for mitigation periods with higher band width to T0.

Architecture, System Testing and Exploitation

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?

- Findings :

The scope is well defined. There is some uncertainty around what should happen on the online buffer.

- Comments:

Additional read and write on the online buffer would have a large impact on online buffer hardware and rack space.

- Recommendations:

  - Decide asap if reread and write of data should happen on the online buffer.

Timing, Trigger and Throttling

- Findings:
  - The synchronisation of the different elements of the detectors is based on the distribution a phase controlled clock to all parts and of a 64-bit time stamp generated by all parts and propagated through the DAQ
  - The necessary precision on the phase clock at the end point does not need to be better than 1 ns
  - Synchronous commands can be sent to all or some group of destinations. This can be used for resync if needed
  - The stability of the phases of the clocks at all destinations can be monitored thanks to a loop back capability
- Comments:
  - The system assumes that there is a fixed latency between the front-end and the place where the complete time stamp is inserted in the data block. That must be a design constraint for the WIB (for instance) which is using the serialisers of an FPGA. That constraint could be removed if a full time stamp is formed by the WIB before transmission to the next stage
- Recommendations:
  - The documentation is to be updated and should include all requirements to the timing system (e.g. precision) as well as requirements to other parts involved (e.g. fixed latency data transmission)

TPC & SSP Readout

- Findings:
  - 5 of the 6 Anode Plane Arrays will be readout by the RCE system, the 6th by a FELIX system. FELIX is seen as R&D project
  - RCE hardware is available for reading out the 6th array in case there are issues with FELIX preventing reliable acquisition of data of the 6th array
  - The RCE hardware is identical to the hardware used for the 35t prototype, but firmware and software is or has to be adapted to the protoDUNE SP environment. Apart from firmware for data compression, firmware is maintained by RCE development group

TPC & SSP Readout

- Findings cont'd:
  - Backup hardware for protoDUNE SP RCE system will be available
  - ATCA control is the same as for the 35t prototype
  - Firmware for testing WIB interfacing available / tested
  - RCE system without compression should be available in time for first Anode Plane Array test. No role for FELIX in this test
  - RCE system needs data compression in firmware to transfer the data into its memory, but not much. Using additional DMA controllers will make it possible to stream uncompressed data to the memory, but the amount of memory available and the network output bandwidth may limit the trigger rate

TPC & SSP Readout

- Findings cont'd:
  - Board Reader processes for the RCE and for the FELIX system were not described, whereas the focus for the SSP readout was the Board Reader process
  - The FELIX system is built from commercially available PCIe cards, already purchased, and server PCs. The BNL-711 card, baseline for the ATLAS Phase-1 FELIX system, may be an alternative, if available in time, then one card in one PC is needed
  - FELIX firmware is the same as ATLAS full-mode FELIX firmware; there is interest from PNNL in adding compression
  - FELIX: baseline one VC709 and one HTG710 in two PCs. Both cards in one PC would be possible, but 2 PCs preferred in view of compression in software (after trigger)
  - The WIB firmware used for the RCEs needs to be modified for FELIX. The full-mode protocol will be used. Testing not yet planned
  - SSP Readout will be tested in Oxford vertical slice
  - In general: no principle problems

TPC & SSP Readout

- Comments:
  - "Customer" relation with SLAC RCE development group, not a protoDUNE collaborator, potentially problematic, but good support and interest has been reported.
  - From RCE presentation: "waiting for Bristol-provided firmware blocks before we work on DTM firmware"? Answer to question on firmware development was that all development is at SLAC by RCE developers and by protoDUNE collaborators. In the document Oxford is, but Bristol is not listed in the table with team members.
  - RCEs / FELIX: it is advisable to send a small fraction of the data uncompressed
  - FELIX: for software data compression the use of a GPU or Xeon Phi could be of interest.
  - FELIX: provided that a sufficiently fast network is used it will be possible to output all data received untriggered and uncompressed.
  - The nature of the data handled by the SSP Readout Board Reader, although not a subject for the review, needs attention.

TPC & SSP Readout

- Recommendations:

  - RCEs: evaluate software and firmware status, determine what needs further work, what needs to be tested and plan accordingly. Also support for running the system (tests, data taking) will need to be discussed. Appoint a coordinator who calls and chairs meetings, defines agendas (Indico), is responsible for planning and reports to or is part of a higher layer of management. Use svn or git repo accessible for all developers for source code, VHDL/Verilog and for documentation. Complement with e.g. Twiki, Sharepoint, JIRA

  - FELIX: same as for RCEs

Trigger Inputs

- Charge

Beam Instrumentation

6. *Is the software architecture suitable, including Event Builder, Run Control, Online Monitoring, Timing, Triggering and Databases?*

The BI event building will be during tier-0.

The pLAPPD readout integration must be studied.

Cosmic Ray tagger

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?*

   This an existing system already used in the Double Chooz experiment.

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?*

The CRT Event building is currently assumed to happen off-line. This should be reviewed.

Institutes and individuals who will work on the CRT DAQ and monitoring need to be identified.

Trigger Inputs

- Findings:

- *Beam Instrumentation*

It consists of scintillating fibres for tracking and momentum measurement, Cerenkov counters and two possible TOF systems. Trigger signals within a few hundred nanoseconds and data fragments will be provided.

It is provided and maintained by the CERN Beam group (except for one TOF system, pLAPPD, is provided by FNAL) and is used for incident particle identification and momentum measurement, as well as for triggering PD.

The trigger signals are sent to the CTB as well as 11 bits carrying trigger planes TOF and Cerenkov information. The actual data is sent to BI storage for retrieval and merging with PD data off-line during the first data treatment pass (tier-0) after conversion to artDAQ dormat. Warning and End of extraction will also be provided.

Synchronization will be provided by time stamping data using White Rabbit (WR) with a precision of 700ps. WR distributes a 1pps signal and a 10 MHZ clock to BI and PD. PD will derive its 50MHZ clock from it.

BI will be ready summer 2017 and pLAPPD is being developed.

Cosmic Ray tagger

It consists of unused Double Chooz Outer Veto (OV) planes each plane made up of 64 scintillating strips readout by WLS fibres and one multichannel pmt. Up to 40 planes are being considered for deployment in front and behind PD for beam halo muon identification. An additional plane may be deployed on top of PD. The read out will be identical to Double Chooz providing bits for each hit counter and, optionally, pulse height information. The read out is asynchronous to PD data taking and will be merged off-line using time stamps. A trigger signal consisting of a coincidence between planes of orthogonal orientation or of an OR of all planes will be generated in an OV trigger board and provided to CTB.

All components have been tested but will be visually inspected. A test stand will be provided at Virginia Tech.

Trigger Inputs

- Comments:

Beam Instrumentation:

The Beam group has a long standing expertise in providing beam signals and triggers and this should allow a smooth use of the BI information.

The Fermilab microchannel plate based pLAPPD integration in the PD DAQ and monitoring must be studied and implemented.

Monitoring the beam profile and quality will not be part of the PD on-line monitoring but can be made through a separate beam monitor terminal provided by BI.

Cosmic Ray tagger:

This is a new addition to the PD hardware. The counter, readout and trigger generation has been shown to function as expected in Double Chooz.

Whereas all components have originally been tested, it would be advisable to recheck them if this has not been done recently.

Trigger Inputs

- Recommendations:

Beam Instrumentation:

Cosmic Ray Tagger

The CRT is a large system (up to 2560 channels): Identify the institutes and individual people involved in the CRT DAQ and monitoring as well as their FTE commitment.

As soon as the CRT final configuration is determined, sort out grounding schemes, cable lengths and layout, positioning of the trigger box, LV and HV supplies and readout PC's. Include the time stamp in this.

Run Control and Monitoring

- Charge questions:
  - 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?
    - Run Control/Operational Monitoring/Configuration architecture appear suitable for protoDUNE
    - There seems to be limited resources in the (JCOP-based) software development

Run Control and Monitoring

- Findings:
  - The team is planning to use a combination of artDAQ and JCOP to provide a complete set of functions for Run Control, Configuration and Operational Monitoring
  - The artDAQ framework will be used for the initial on-site testing and commissioning while the JCOP-based functionalities are developed
- Comments:
  - JCOP is a powerful toolkit providing a large set of base tools. It is used for the Run Control of other large experiments (LHCb)
  - artDAQ is a well established data-acquisition framework which embeds Run Control/Configuration/ Operation monitoring functionalities
  - The team provided a initial scheme for interfacing a JCOP-based Run Control with other existing components
  - Plans for software logging and debugging and operational monitoring seems to be sufficient for the protoDUNE needs
  - Achieving a complete set of functions (i.e. from run control to operator alerting) in a JCOP environment will require significant development, integration and testing efforts
  - Separation of responsibilities with respect to the Run Control between artDAQ and Run Control team may have to be clarified. This will help evaluating the level of available manpower
- Recommendations:

Dataflow

- Findings:
  The dataflow will be based on the established artDAQ framework developed and maintained by Fermilab and already used for the 35t. At the hardware level it is based on standard x86 servers and Ethernet (Gigabit and 10-Gigabit). At the core of the network a single Brocade ICX7750 device is going to be used. Event-building is performed at about 24 Gbit/s, assuming a data-compression achieved of about a factor 4. An online storage system will be capable to receive these data and have a capacity for up to three days. artDAQ includes a mechanism for online monitoring, where (sampled) data can be sent to dedicated servers without the possibility of back-pressure on the actual data-flow.

- Comments:

  - A single switch for all data-traffic between the DAQ and monitoring components will be used. This switch will have to cope with a lot of sources with quite variable traffic patterns, including occasional spikes in required bandwidth.

  - A back-pressure free monitoring data feed has been presented, but a framework for the actual monitoring processes and resource estimations for the same seem to be missing. For instance the data will need to be decompressed, the events are large, and the rate is low (125 events / spill) so that it is not clear how much down-sampling is desirable. In the monitoring there seems also to be a lack of person-power.

- Recommendations:

  - Make a clear inventory of all components connecting to the switch, with worst-case assumptions, before procurement. As much as possible it should be attempted to test the switch before by generating traffic similar to what is expected in the actual operation, in particular including peak-loads.

  - Identify more effort for the technical side of the monitoring as early as possible. One of the first tasks of such a person could be to estimate the computing and network resource needs of the monitoring as a function of sampling rate and complexity of the desired processing.

Management

- Findings:
  - The ProtoDUNE DAQ team was put together over the summer
  - A proto-orgchart was presented
  - No set of milestones was presented
  - The risk registry provided is uneven in its description of risks
  - The schedule provided is quite detailed for the next few months

Management

- Comments:
  - Given the very small time since inception, the team has made remarkable progress
  - While ProtoDUNE is a test bed, its importance to the DUNE project makes early success almost into a requirement
  - The DAQ management team would benefit from clear definition of responsibilities through an org chart
  - A top-down resource estimate for the different DAQ aspects would allow direction of new effort to tasks that need it most
  - Milestones (spaced every few months) would allow objective progress assessment and taking appropriate actions when needed
  - The risk registry should be fleshed out to include impact and mitigation strategies for all identified risks
  - It is important to identify any items that might require full procurement procedures and their impact on the schedule

Management

- Recommendations
  - Complete the DAQ org chart before the end of the year
  - Establish key milestones (spaced every few months) for each DAQ subarea by the end of the year