DEEP UNDERGROUND
NEUTRINO EXPERIMENT

Report of the Design Review of the

ProtoDUNE DAQ

3–4 November 2016

CERN

1) Introduction

The ProtoDUNE DAQ system was reviewed on November 3 and 4 at CERN. The review committee consisted of Gustaaf Brooijmans (Columbia, chair), Leslie Camilleri (Columbia), Philippe Farthouat (CERN), Sigve Haug (BERN), Niko Neufeld (CERN), Andrew Norman (CERN), Wainer Vandelli (CERN) and Jos Vermeulen (NIKHEF). Extensive documentation is linked from the indico page[1] in the form of documents and talks. The ProtoDUNE DAQ system is designed to provide reliable data acquisition as soon as possible, and not as a prototype DAQ platform for DUNE. It therefore extensively builds on existing tools, on both the hardware and software fronts. The committee very much supports this approach given the very short time available to provide a working system. Nevertheless, significant customization work is required.

This report first addresses findings, comments and recommendations specific to individual DAQ components, and then answers the charge questions.

2) Architecture, System Testing and Exploitation

1. 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 through the use of loss-less compression. During the beam spills (2 x 4.8s in the 48s SPS supercycle), ProtoDUNE expects to take data at 25 Hz with event size reduced to 60 MB, i.e. at 1.5 GBps rate. The average rate from beam events is thus 2.4 Gbps over the whole supercycle. Including cosmics, the experiment expects to write an average of 2.7 Gbps during the 2018 run. At 25 Hz, with 100% data taking and accelerator efficiency, the experiment would collect the 6 million desired physics events in 41 days.

   The actual compression ratio that is achievable depends on the (coherent and incoherent) noise level in the ProtoDUNE detector. Compression factors better than 2 have been demonstrated on the 35t prototype which had unforeseen sources of noise.

   The available bandwidth from ENH1 to Tier-0 is 20 Gbps.

   There are plans for a full vertical slice/integration test bed including a cold box, which will also serve for acceptance tests as detector elements arrive at CERN.

2. 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 design allows for some alternates. Up to the COB continuous read out is in principle possible. In the design ATCA COB (RCE) are used for the readout of five APA's and FELIX for the sixth one. However, ATCA COB read out for all 6 APAs is available. Downstream, there is redundancy depending on network (10 Gb/s) and server capacity.

The online cluster size and performance requirements are crucial for defining cooling, power and rack capacity. It is not completely evident that three plus one spare racks will be enough, in particular depending on the resources needed for online monitoring, given the large single event size.

The DAQ system scope is generally well-defined. There is some uncertainty around what will happen on the online cluster.

3. 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 otherwise 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.

3) Risks

1. Findings

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

The risk registry provided was uneven in its description of risks

2. Comments

The presented risks seemed reasonable and inline with the presented design for the DAQ system.

The lack of 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 pre-installation 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.

The risk registry should be fleshed out to include impact and mitigation strategies for all identified risks

3. 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.

4) Timing, Trigger and Throttling

1. Findings

The synchronisation of the different elements of the detector is based on the distribution of 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.

2. 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 before transmission to the next stage.

3. Recommendations

Update the documentation and include all requirements to the timing system (e.g. precision) as well as requirements to other parts involved (e.g. fixed latency data transmission)

5) TPC & SSP Readout

1. 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 an 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 the RCE development group.

Spare replacement hardware for the protoDUNE SP RCE system will be available.

ATCA control is the same as for the 35t prototype.

The RCE firmware for testing the WIB interfacing is available and tested.

The RCE system without compression should be available in time for the first Anode Plane Array test. There is no role for FELIX in this test.

The 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.

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 baseline is one VC709 and one HTG710 in two PCs. Both cards in one PC would be possible, but 2 PCs are preferred in view of compression in software (after trigger). 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 sufficient.

The FELIX firmware is the same as the ATLAS full-mode FELIX firmware; there is interest from PNNL in adding compression firmware.

The WIB firmware used for the RCEs needs to be modified for FELIX. The fullmode protocol will be used. Testing for this is not yet planned.

SSP Readout will be tested in the Oxford vertical slice.

2. Comments

“Customer” relations with the SLAC RCE development group, not a protoDUNE collaborator, are potentially problematic, but good support and interest has been reported.

From the RCE presentation: “waiting for Bristol-provided firmware blocks before we work on DTM firmware.” However, the answer to a question on a firmware development was that all development is at SLAC by RCE developers and by protoDUNE collaborators. In the documentation 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.

It would be worth investigating if the FELIX firmware can be adapted instead of the WIB firmware.

3. Recommendations

For the RCEs as well as for FELIX: 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 for the RCEs and a coordinator for FELIX, who call and chair meetings, define agendas (Indico), are responsible for planning and report to or are part of a higher layer of management. Use an svn or git repo accessible for all developers for source code, VHDL/Verilog and for documentation and complement this with the use of e.g. Twiki, Sharepoint, JIRA.

6) Beam Instrumentation

1. Findings

This 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, 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 are 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 format. 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 700 ps. WR distributes a 1pps signal and a 10 MHZ clock to BI and PD. PD will derive its 50 MHZ clock from it.

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

2. Comments

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.

3. Recommendations

None

7) Cosmic Ray Tagger

1. Findings

This 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 set of planes 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 offline 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.

2. Comments

   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.

3. Recommendations

   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 and readout PC’s. Include the time stamp in this.

8) Run Control and Monitoring

1. 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.

2. 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 and Operational monitoring functionalities.

   The team provided an 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.

3. Recommendations

   None

9) Data Flow

1. 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.

2. 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 (25 Hz during the 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.

3. Recommendations

   Make a clear inventory of all components connecting to the switch, with worst-case assumptions, before procurement. It should be attempted to test the switch before operation 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.

10) Management

1. Findings

The ProtoDUNE DAQ team was put together over the summer.

A proto-orgchart was presented.

No set of milestones was presented.

The schedule provided is quite detailed for the next few months.

2. 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.

It is important to identify any items that might require full procurement procedures and their impact on the schedule.

3. 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.

11) Answers to Charge questions

a. 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. 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. 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.

b. 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.

c. 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 2017. The timelines for 2018 leading up to actual detector commissioning are not as developed. Most of the subsystems appear to have adequate time in the base schedule for testing. It may be beneficial to accelerate these schedules where possible to mitigate unforeseen delays due to the hard deadline imposed by the long shutdown.

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

   Yes. The design is generally well-documented. One weakness is the resource requirement for monitoring, since events are very large. Given its impact on the online cluster size, this should be determined soon.

e. 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?

   Yes, with the same caveat as above regarding monitoring. Readout of the pLAPPD needs to be added. No grounding issues should arise if all connections to the detector are over optical links as planned. The remaining copper connection between a switch and the SSP front-end is to be carefully handled.

f. 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 Run Control/Operational Monitoring/Configuration architecture appear suitable for protoDUNE, but there seems to be limited resources in the (JCOP-based) software development.

g. 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?

Yes, but the data flow to monitoring should be added.

h. 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?

   Mostly, yes. The compression firmware under development for the RCEs should provide for the ability to take some data uncompressed.

i. 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?

   Mostly, yes. The installation plan is very fresh but the responsible person will move to CERN full-time soon and should be able to address this.

j. 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?

   Yes. The 35t lessons are documented and implemented in the plans. There are test plans from all subsystems.

A. Review Charge

DUNE Design Review:
ProtoDUNE Single Phase DAQ System
3–4 November 2016
Charge

The Committee is requested to review the DUNE Single Phase DAQ system technical design and determine if it is at a state commensurate with that needed for NP04 ProtoDUNE detector operation at the CERN Neutrino Platform in 2018.
In particular, the review team is asked to address the following questions:

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?

The committee should present its findings, comments, and recommendations in a closeout meeting with DUNE management on November 4. The committee should provide a final written report by November 18.

B. Review Committee

Name	**Phone**	**email**
Gustaaf Brooijmans (Columbia) Chair		gusbroo@phys.columbia.edu
Leslie Camilleri (Columbia)		camil@nevis.columbia.edu
Philippe Farthouat (CERN)		philippe.farthouat@cern.ch
Sigve Haug (Bern)		sigve.haug@lhep.unibe.ch
Niko Neufeld (CERN)		Niko.Neufeld@cern.ch
Andrew Norman (FNAL)		anorman@fnal.gov
Wainer Vandelli (CERN)		Wainer.Vandelli@cern.ch
Jos Vermeulen (NIKHEF)		i73@nikhef.nl

C. Review Agenda

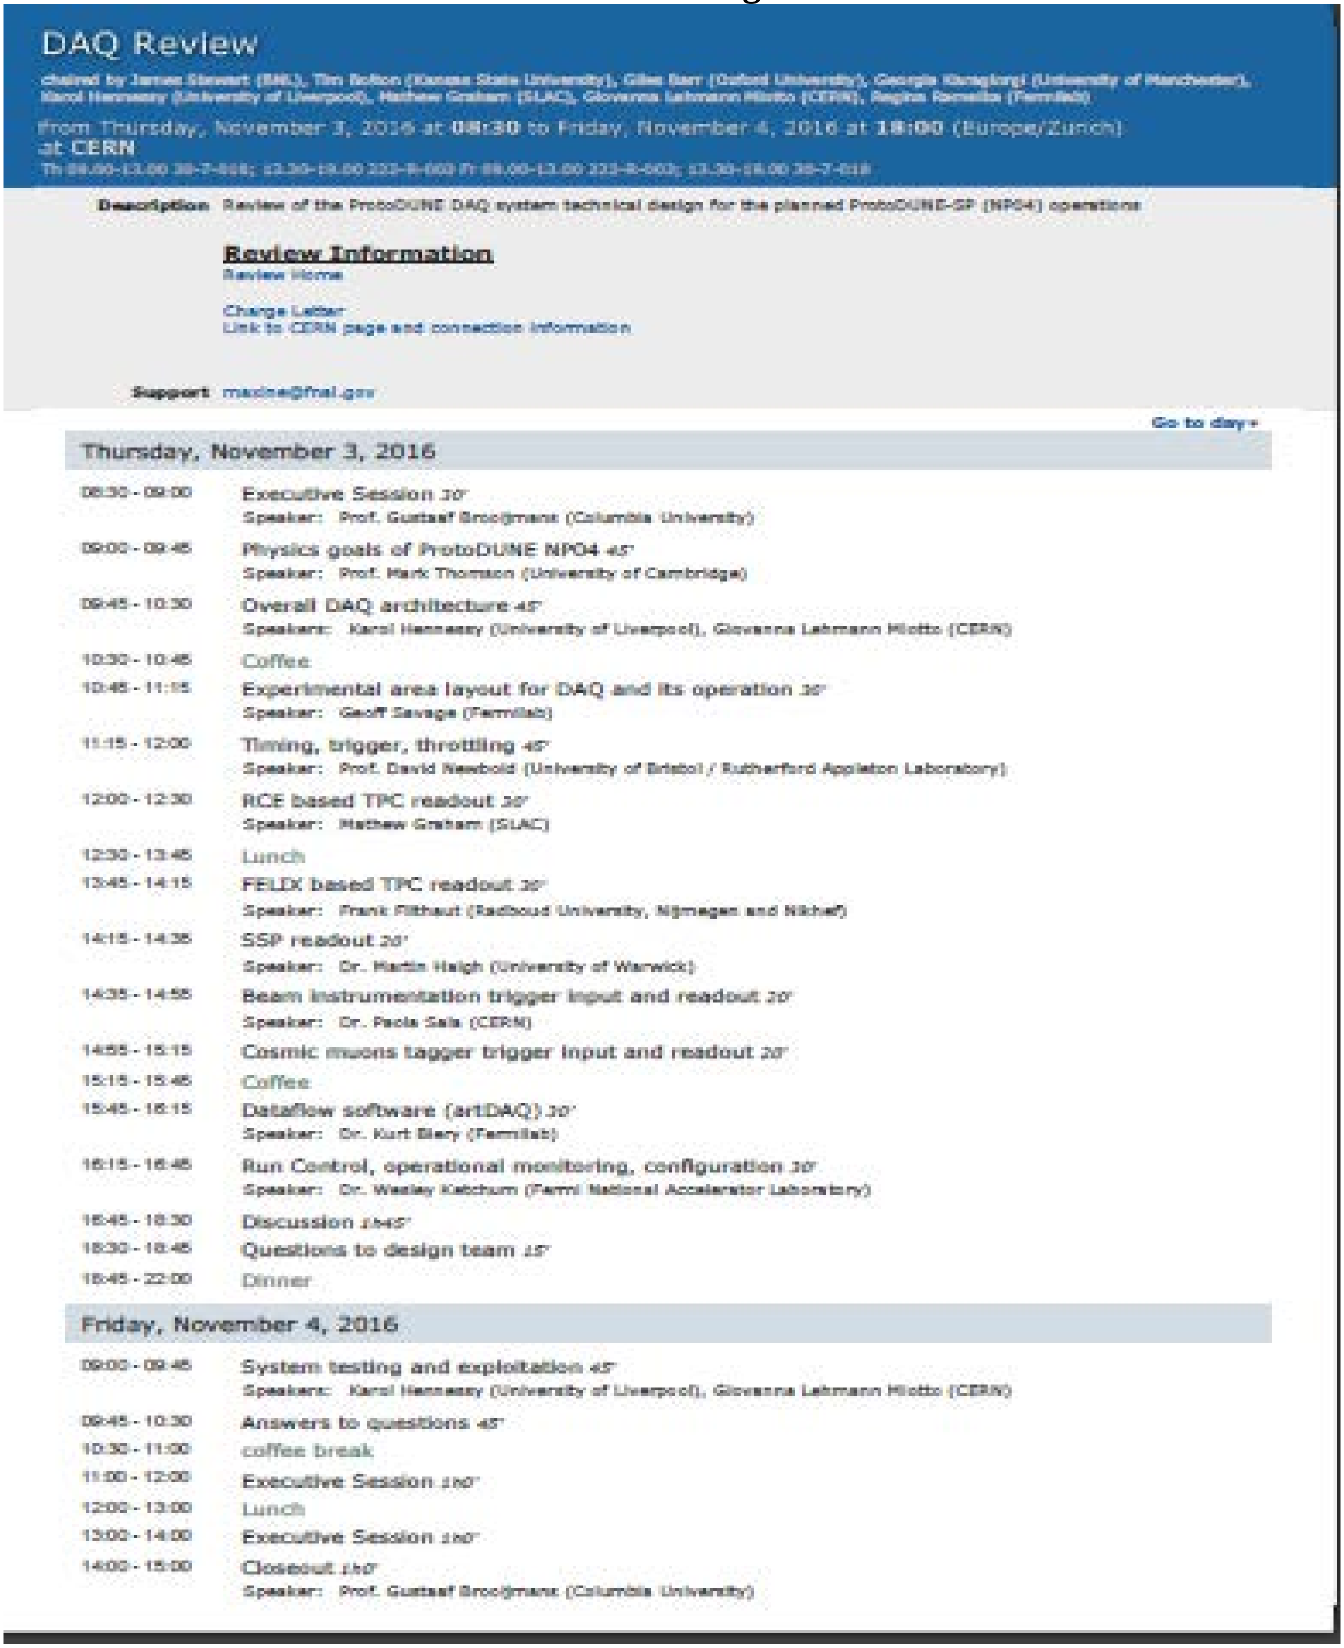

DAQ Review

chaired by James Stewart (BNL), Tim Bolton (Kansas State University), Giles Barr (Oxford University), Georgia Karagiorgi (University of Manchester), Karol Hennessy (University of Liverpool), Mathew Graham (SLAC), Giovanna Lehmann Miotto (CERN), Regina Rameika (Fermilab)

from Thursday, November 3, 2016 at **08:30** to Friday, November 4, 2016 at **18:00** (Europe/Zurich)
at **CERN**
Th 09.00-13.00 30-7-018; 13.30-19.00 222-R-003 Fr 09.00-13.00 222-R-003; 13.30-19.00 30-7-018

Description Review of the ProtoDUNE DAQ system technical design for the planned ProtoDUNE-SP (NP04) operations

Review Information
Review Home

Charge Letter
Link to CERN page and connection information

Support maxine@fnal.gov

Thursday, November 3, 2016

08:30 - 09:00	Executive Session *30'* Speaker: Prof. Gustaaf Brooijmans (Columbia University)
09:00 - 09:45	Physics goals of ProtoDUNE NP04 *45'* Speaker: Prof. Mark Thomson (University of Cambridge)
09:45 - 10:30	Overall DAQ architecture *45'* Speakers: Karol Hennessy (University of Liverpool), Giovanna Lehmann Miotto (CERN)
10:30 - 10:45	Coffee
10:45 - 11:15	Experimental area layout for DAQ and its operation *30'* Speaker: Geoff Savage (Fermilab)
11:15 - 12:00	Timing, trigger, throttling *45'* Speaker: Prof. David Newbold (University of Bristol / Rutherford Appleton Laboratory)
12:00 - 12:30	RCE based TPC readout *30'* Speaker: Mathew Graham (SLAC)
12:30 - 13:45	Lunch
13:45 - 14:15	FELIX based TPC readout *30'* Speaker: Frank Filthaut (Radboud University, Nijmegen and Nikhef)
14:15 - 14:35	SSP readout *20'* Speaker: Dr. Martin Haigh (University of Warwick)
14:35 - 14:55	Beam instrumentation trigger input and readout *20'* Speaker: Dr. Paola Sala (CERN)
14:55 - 15:15	Cosmic muons tagger trigger input and readout *20'*
15:15 - 15:45	Coffee
15:45 - 16:15	Dataflow software (artDAQ) *30'* Speaker: Dr. Kurt Biery (Fermilab)
16:15 - 16:45	Run Control, operational monitoring, configuration *30'* Speaker: Dr. Wesley Ketchum (Fermi National Accelerator Laboratory)
16:45 - 18:30	Discussion *1h45'*
18:30 - 18:45	Questions to design team *15'*
18:45 - 22:00	Dinner

Friday, November 4, 2016

09:00 - 09:45	System testing and exploitation *45'* Speakers: Karol Hennessy (University of Liverpool), Giovanna Lehmann Miotto (CERN)
09:45 - 10:30	Answers to questions *45'*
10:30 - 11:00	coffee break
11:00 - 12:00	Executive Session *1h0'*
12:00 - 13:00	Lunch
13:00 - 14:00	Executive Session *1h0'*
14:00 - 15:00	Closeout *1h0'* Speaker: Prof. Gustaaf Brooijmans (Columbia University)

---

[1] https://indico.fnal.gov/conferenceOtherViews.py?view=standard&confId=12392