External Review of the NOνA Front-end Electronics Review Report

19 February 2010

Summary

This is the final report from the NOνA electronics external review conducted by John Anderson (ANL), Steve Chappa (FNAL), Maurice Garcia-Sciveres (LBNL), Jim Patrick (FNAL), and Claudio Rivetta (SLAC). This was a "virtual review" in that there were no formal presentations to the committee. Instead, documentation was made available in the NOνA docdb and a list of subsystem experts provided. We talked directly and via e-mail with some of the engineers involved. The charge was given in the overview document for the review (NOνA-docdb-4342) and is given below in Appendix I. The scope was taken to include the APD, FEB, DCM, Timing, and Power Distribution systems and overall system considerations among these. We focused on technical issues and did not examine in detail the cost and schedule except what was included in the technical documentation. Thus we were not able to judge the effect of costs on design choices.

The overall system design seems sound, and the collaboration has made good progress on the individual parts. Prototypes for most components have been built and tested, and quantities for the IPND test beginning this summer are in the process of procurement. The IPND will be the first time many components are tested together and it is extremely important for validating all the designs. For the IPND test to be successful there needs to be strong on-site coordination, and a plan of tests to be performed and measurements made before signing off on full production for the far detector. Another topic that raised many questions was the detector electrical distribution and grounding. We suggest a focused review be done on that topic. Below we offer more detailed comments on each of the subsystems followed by global system issues.

APDs

- The test time of 333 days is estimated based on single pass testing. A 2nd pass contingency should probably be assumed so that any issues uncovered during integration may be addressed with re-testing. This would require a doubling of the planned test speed. This requires adding a second mainframe meter, the cost seems sufficiently modest that this should be included in the testing plan.

- Only functional testing at nominal operating conditions is proposed, which provides no information on operating margin. A more parametric test program could be valuable (for example measure the change in performance with temperature). This would require making measurements at 2 or more temperatures. While the nominal operating temperature is -15C, in practice the value could be changed in response to problems or new constraints. While prototypes have been tested at different temperatures, we recommend the same be done to at least a subset of the production modules.

- A validation of the test results in real operating conditions would demonstrate sensitivity of the tests. That is, take APDs that passed the tests, build them into prototypes, measure the performance in the prototypes, and correlate these results with the test data. While this has been done for a somewhat similar system in CMS, this should still be done here as well.

- Gains within an APD array will not be equalized, but a maximum variation is specified. The acceptance cut on gain uniformity should take into account any time dependence observed in CMS.

- In the APD boards, the HV distribution/signal path comprises a large loop area. This loop is just at the input of the FEE amplifier (the most sensitive area of the detector). It increases the susceptibility of the FEE electronics to spurious magnetic fields. In the APD board re-design (up-date of version 4.0FEB) try to improve this issue in the layout. Similarly, if it is possible, allocate more wires for –2.5V in the flat cable connecting the APD to the FEB (-2.5V is the signal current return).

- There is currently no plan to include Photodiodes and LED flashers into the final modules for in-situ for calibration. Gain evolution of time will be tracked using cosmic rays and by taking advantage of the narrow band nature of the beam. We suggest a small scale flasher system could be implemented in just a few modules in order to provide a separate calibration with different systematics that would cover the full dynamic range. Perhaps this could be done on an initial "phase 0" deployment of far detector modules.

Front-End Board (FEB)

- We and the collaboration note that the ADC chip seems to be working without problems in CMS where the requirements are somewhat more demanding.

- For the readout chips a detailed testing plan needs to be made and reviewed. There should be some contingency in case the yield is less than expected

- The repeatability of the TEC circuit performance (minimizing the ripple currents) is a potential issue. While the design uses magnetic components that should give consistent performance this should be verified with the devices built for the IPND. Reliability and failure modes are also important to understand. There seems to be a good plan to address these and other potential issues with the TEC.

- The flow-rate of the water cooling system at TEC level is 2mL/sec. That implies a flow-rate for each chiller of 24 GPM which seems high. We recommend evaluating lower flow rates at the IPND test. Lower flow rates could simplify the chillers and reduce the cooling power.

- Given the large number of FEBs in the system, reliability should be a consideration in the engineering design. How important this is depends on the impact of bad FEBs on the data, and the time required to identify and replace them. If they have not already done so the collaboration should provide input on the reliability requirements. Also the mechanical layout should make it straightforward to replace bad FEBs with a high probability of success.

- Various prototyping activities are under way for the near term (Near detector, vertical slice test at Caltech, CDF Hall engineering prototype). It was not clear what aspects of the FEB and APD will be tested in each of these prototyping efforts. It may be useful to spell out, for each prototype, what results should be obtained related to the APD and FEB. Increased coordination of electrical system integration for these various efforts would probably help.

- It is planned to store in a database measurements made during production testing. The collaboration should do studies correlating these measurements with measurements made in the IPND and installed detector.

Data Concentrator Module (DCM)

- The links to the 64 FEBs are DC coupled which saves a lot of space that would be devoted to capacitors. Given the modest cable lengths and appropriate grounding this seems reasonable but needs to be verified in a realistic operational environment. In case there is still uncertainty when production quantities have to be ordered, a possibility would be to put pads on the boards. Though this would require allocating this extra space on the boards.

- Having been put on hold in 2008, the design is now some years old. An assessment of projected future parts availability should be done before finalizing the production version.

- With the large number of cable connections to the FEBs, attention needs to be given to appropriate strain relief, robustness of the cable connectors, and ease of reliable DCM replacement. We understand this is being addressed in a mockup at Lab 6.

- When estimating the power draw, reasonable assumptions need to be made for CPU utilization and number of switching FPGA elements. Also the startup currents due to FPGA initialization can be large relative to normal currents.

- While the hardware development is done at Fermilab, there currently appears to be no local expertise for the low level software (boot loader, operating system, and device drivers). This is now handled remotely by the University of Virginia. Until such time as the boards including the firmware and FEB communication are debugged, this is likely to be inefficient. We recommend trying to obtain on-site support from the computing division at least for the initial IPND phase.

- It appears that only very limited testing with FEBs has been done. There is substantial communication involved, not just data but timing, control, and status information. A high priority should be placed on debugging this communication with multiple FEBs.

- We recommend that remote reset lines be available so that boards may be rebooted without requiring a power cycle. Remote access to the console serial output would also be useful in diagnosing situations where a processor doesn't boot.

- There is a substantial amount of software required for these boards. There is the low level software mentioned above, FPGA programming of the DCM and FEBs, configuration of FEBs,

DAQ and DCS monitoring, diagnostics, as well as the normal data acquisition. The software has to be robust against synchronization or other errors from FEBs. There seems to be quite a lot to get done by the IPND test this summer, especially given that minimal attempts have been made to communicate with FEBs. A prioritized list of the essential software elements to operate the system for the IPND test should be made and specific people assigned to them.

- Transmission of data to the buffer farm at the far detector goes through a 3 layer network of 48 port switches and relies on traffic shaping by the DCMs to minimize contention in the switches. Depending on the details of the switches, this scheme might be sensitive to CPU load or other transient issues on the DCMs, or systematic data imbalances among DCMs. Ideally some modest test of this scheme should be done. An alternative, though likely more costly, would be to use enterprise class switches such as have been used in recent years by D0, CDF, and CMS. A system based on one would likely alleviate the need for traffic shaping and reduce the buffering required on the DCM.

DC Power Distribution System

In the course of reviewing the power distribution system, a number of issues were found in the documentation and corrected or clarified by the responsible parties. We note that complete and accurate documentation will be required in order to receive Operational Readiness Clearance for the detector and recommend the documentation be reviewed again prior to submission for that. Schematics should include the final connection of cable shielding and also wire characteristics. And also they should include the final filter or decoupling capacitors.

High Voltage Distribution

- The CAEN HV channels individually float. The HV power distribution must be safely connected to ground to avoid developing a high voltage potential with respect to ground on metallic parts that should be a potential zero with respect to GND. The 450VRET line has to be connected to GND via a 1K-10K resistor at the PDB (no jumper, solid connection). Similarly any cable shield or metallic cover from the power supply to the patch panel has to be referenced to ground with a permanent connection (resistors, no jumpers). The external shield of the HV cable (Triaxial Belden 9222) should be connected for safety reasons to ground at the HV power supply end. The connection of the shield to ground at the PDB chassis is better from an electrostatic shielding point of view but can be unsafe. Because the HV power supply is located in a different area than the PDB, disconnection of either the HV cable from the PDB chassis or the PDB to GND can induce a potential hazard if the HV power supply is turned ON.

- When fabricating the HV patch panel, refer to the Fermilab ES&H Manual, Section 5045, page 5045-1, for requirements when mounting the coaxial connectors: Use of D-holes for mounting, connector shell solidly grounded to panel, etc. These requirements are the same when using triaxial connectors. This means that the outer shield of the triax cable (guard shield) must remain grounded at all times. (Assuming that the inner shield is used for the channel's HV return connection.) As described above, if the outer shield is connected to ground only at the PDB end, then the cable will lose its ground when disconnected from the PDB chassis. Thus, if it is desired

for operational purposes, to ground the guard shield at only one place, this place must be at the source (the patch panel and then to the CAEN chassis through the rack metal and safety ground connection). Ideally, the outer shield should be grounded at both ends.

- Second, a point on the cable/wiring connections between the AP1520P output connector and the patch panel triaxial connectors: the insulation on the individual wires needs to be rated at 600 V, not the usual 300 V used for normal hook-up wire. Also, if the terminations on the coaxial connectors are solder-tail connections, then these connections are considered exposed and need to be covered. This covering can either be local to the patch panel or by enclosing the rack where the panel is mounted. Either way, these connections cannot be open to touch by personnel.

- The coaxial connectors should have three shell bayonets (as opposed to the usual two bayonets used on standard signal BNC connectors). This is to prevent any compatibility with the BNC connectors used for low-voltage signals. If two bayonet shells are used, the patch panel must be clearly labeled for exclusive use with the HV.

Low Voltage Distribution

- The cable is used for PDB-FEB connection is 6 conductor non-paired 18AWG. We would suggest using a twisted pair or shielded twisted pair cable.

- Fuses at the output of the power supplies should have a maximum current limit lower than the current capacity of the wire. Those currents are above the nominal current specified in the documents. If the current limits in the power supplies are set at a level similar or just above the fuse rating, the fuse is not going to blow-up in case of load failures. Fuses are included in case of failure of the electronic current limit of the power supply or other power supply failure where the current limits do not operate.

Power Distribution Box

- We note there is no remote voltage or current monitoring of loads driven by the PDB. Though the FEBs and DCMs can provide some information, this requires them to be mostly operational. Most often if there is a power problem they won't be able to do this. The plan is to rely on the online data monitoring system to identify parts of the detector with missing or bad data. As there are a variety of possible causes of missing data, we note this scheme will make it somewhat more difficult to localize failures.

- The grounding of the PDB for high frequency signals has to be evaluated in the context of the overall grounding scheme of the detector. The 0AWG wire is considered ground just for safety purposes.

Safety and Interlocks Issues

Since current protection is a safety issue and also there is a large overlap between electrical safety and the DC Power Distribution's proper operation, most of the safety issues regarding current protection, proper conductor sizing, etc were addressed in the Preliminary Design Review.

- The primary safety issue, as concerns the electronics, is the use of the 450 VDC from the Caen power supply mainframes. Most of the hazards identified and outlined in the Preliminary Design Review document (NOνA docdb-3689) have been addressed: The individual output channels can limit its output current. There is a 1 mA current limiter circuit for each output on the distribution PC-board. Also, I see no need for a protection interlock within the power distribution box.

- The second safety issue is with the rack infrastructure. There was no description of the AC power distribution for an equipment rack. Since multiple Weiner power supplies will be housed within one rack, there should be some thought as to how to safely distribute the 208 VAC power within that rack. Even if the building's wiring is not yet known, an electrical drawing for the AC wiring and grounding within a rack needs to be done and reviewed.

- The current plan is to not implement a full rack protection interlock system as was done for MINOS (and also at least for CDF and D0 in past experiments at the lab) due to the high false alarm rate at MINOS. We did not have enough detailed information on what will be in each rack to judge the risk this presents. Since these racks are housed in an enclosed building mostly below ground level we would suggest, at a minimum, to have smoke detection interlocks for each rack that houses high power electronics. In the event of a smoke detection, the AC power can be dropped to reduce the probability of activating the building's fire suppression system. The only protection against an overload condition that can lead to smoke or fire will be proper fusing at each stage of DC power usage. In this case, the potential problem will be to insert more than one fuse in a single power circuit or improperly sizing the fuses for the handling of potentially large in-rush currents.

  For the AC power, protection of the (208 VAC or 240 VAC) wiring in each rack will then be an over-current device (three-pole breaker or fuse). The location for this breaker (or fuse but a breaker is recommended) can be at a service subpanel only if the rack is fed by an individual circuit. If a subpanel circuit feeds more than one rack, each rack in turn must have an over-current breaker, size to the rack's load, as part of the rack's AC distribution. If both 208VAC and 240VAC supply voltages (as indicated in the power distribution document) are to be used, then a breaker will be required for each voltage.

  Some type of chassis, therefore, will be required to house these AC breakers. The addition of a reliable shut-off relay and smoke detector should not require that much more additional effort.

Safety Review Issues

In April of 2009, Steve Chappa of the committee conducted a preliminary (advisory) design review of the DC power distribution system. The results of this review were documented in NOνA

docdb-3689-v1. From the recommendations made in this document and from subsequent discussions, several changes to the PC-boards used in the power distribution box (PDBs) were made. These changes are reflected in the current document (NOνA docdb-4356) that describes the power distribution system.

There are five areas of concern that exist:

1. For PC-board layout with respect to the PDB's modules, there should be sufficient clearance between the card guides and any copper lands or traces.

2. We note the original card-edge connectors for the PDB modules have been replaced with a pinned connector design. We assume the chosen connector and layout will support the required loads with some margin, has pins that are guarded and that are not easily bent on insertion.

3. There should be no "HOT swapping" of these modules. There is just too much power at the backplane connectors. Any incorrect insertion of the modules into the backplane connector will cause a short as there is no fuse protection. We understand the collaboration agrees with this.

4. There exists a potential for various equipment to have voltage reference potentials that may exceed the CM rejection of the serial data transmission. I would recommend that a systematic grounding diagram be drawn up and reviewed so that potential reference grounding problems, if any, can be identified. This diagram should also include safety grounding connections so that any AC powered equipment will not be allowed to "float" and thus provide a safety hazard.

5. Protection for the backplane from physical damage or contact is a concern, given the upstream fusing of the power conductors supplying the DC power to the PDB. It was not originally obvious to us from the documentation the chassis is to be enclosed. We have been assured that it will be so this should not be an issue.

In the document on page 15, section 7.1, it describes the routing of the high-current conductors from the Weiner supplies to the distribution boxes within 1.25-inch conduit. Running conductors in conduit, as opposed to running conductors in open wire baskets, reduces the ability of any heat generated by the current within the conductors to dissipate. Following the NEC fill requirements as is planned should alleviate this concern.

Timing Distribution System

The source material used for the review includes the Timing System Distribution document (NOνA docdb-4354) and two discussions with Vince Pavlicek and Rick Kwarciany. The system described is based on that developed by Bill Haynes for the Accelerator's BPM upgrades. Thus the system does have an operational history. In addition, the readout scheme for the detector does not require timing accuracies at the nano-second resolution. However, being a serially distributed timing system, the following concerns are listed:

- The serial distribution of the timing signals means that a failure in any one of the slave Timing Distribution Units (TDUs) will knock out the operation of all "downstream" modules. Similarly for the DCMs daisy chained off of each TDU. There should be diagnostic procedures to identify the location of a failure in the chain. Some kind of diagnostic signal, from an LED to a scope output may be useful. Also, the use of relays on slave modules that would bypass a module without power would allow the system to continue to function normally.

- The system design includes some provisions for redundant timing links but there did not seem to be a plan for how or even whether to use this capability. We note this would most likely not help in the case of a dead module in the chain.

- The jitter of the 10 MHz clock derived from the GPS unit can vary widely among vendors. We suggest carefully evaluating candidate units. This appears to have been done.

- With resources scarce for supporting the embedded PPC processors in the DCMs it would seem like using another flavor of processor (ARM) in the timing distribution boards would exacerbate this situation. However it appears there is a history of support for this processor in the computing division so perhaps it is not an issue.

- We recommend that remote resets be provided to bring the boards to a known state in case of loss of sync. Power cycling as a routine reset mechanism is not acceptable.

- So far, there has been no system testing of a loop with several modules to test various aspects of the relative timing compensation, adjustment, signal (serdes) generation, etc. This prototyping testing is to be done during the above-ground test detector scheduled for June '10. This does not leave a lot of time for revisions and or corrections.

- The last concern involves the cables and connectors chosen (CAT6 cables using RJ45 connectors). These connectors are not very robust and are susceptible to physical damage. Also, most cables terminated with these connectors use a foil shield that, if not mindful of the bending radius, can become damaged. This can produce intermittent noise and reliability issues. If this cable/connector arrangement is to be used, testing of the cable/connector's integrity needs to be done before and during installation. The same comment applies to the cables between the DCMs and FEBs.

Global System Issues

- The far detector electronics must operate at Ash River for many years presumably with very minimal on-site staff. Components need to be designed for reliability, and also in a way that allows for straightforward pin-pointing of failures, and ease of reliable replacement. In addition to the normal readout software, some diagnostic software will likely be needed. Some people we talked to recognized these issues, but they are not mentioned much in the documentation. We suggest that attention be given to remote reliability and maintenance issues during the IPND testing and the development of prototypes for the far detector.

- There does not appear to have been very extensive integration testing of different components prior to ordering quantities required for the IPND. Some APD/FEB work has been done and will be done at the CalTech vertical slice facility. But there appears to have been minimal testing of FEB/DCM integration, and integration of anything with the timing and power distribution system. It may be more efficient to do initial integration tests on test stands before installation on the IPND. And we hope the overall schedule allows taking into account experience gained from the IPND before production quantities are ordered for the full detector.

- We recommend developing a list of goals and milestones for the IPND tests. There are many tests to be done and measurements to be made that will provide essential input before signing off on production of the full far detector electronics. These need to be done in a timely manner in order to meet the overall schedule.

- We did not examine in detail the overall electrical distribution system and grounding. We strongly recommend a separate focused review on this issue.

- Off-site institutions providing hardware for the IPND need to have a substantial Fermilab presence through the remainder of this year. This will be especially important as many components have not yet been tested together. There should be an on-site coordinator for the IPND electronics/DAQ effort.

- While we suspect it is already the plan, we recommend the far detector deployment be staged. That is individual blocks be completed serially and go into operation while installation and checkout of new blocks continues. That way integration issues can be identified early and applied to later blocks.

Appendix I – Charge

Review Purpose and Scope

"This external review of the electronics and data acquisition system for NOνA is intended to constitute a final system review of the chosen technologies for the NOνA experiment. As the designs represent significant amounts of engineering and testing effort, the review should be limited in scope to serve only as a review of the feasibility and any flaws of the system as designed. Examination of design, operational, and safety issues can be recommended by the reviewers."