May 12, 2017

DUNE Cold Electronics Task Force Report

Per Hansson Adrian[1], Hucheng Chen[2], David Christian[3], Grzegorz Deptuch[3], Angelo Dragone[1], Daniel Dwyer[4], Vernon Emerson[2], Carl Grace[4]

[1] SLAC, [2] Brookhaven, [3] Fermilab, [4] LBNL

The charge given to our group (in italics) and our responses follow.

- *Review and concisely restate the requirements for DUNE far detector single phase LAr TPC cold electronics and how these requirements are related to the various goals of the experiment, including measurements of the neutrino beam, measurements of possible nucleon decay, and astrophysical measurements.*

In our review of the requirements for the DUNE far detector single phase LAr TPC cold electronics, we have focused on the requirements for the front-end amplifier/shaper and the ADC. These requirements are intimately coupled to the design of the TPC and to assumptions about LAr purity as well as to features of the physics measurements DUNE hopes to make. We have failed in the effort to concisely restate these requirements, but generally speaking we confirm the requirements that have been developed in the past by the collaboration. A large amount of work has been done with the goal of deriving electronics requirements from the physics goals by members of the DUNE collaboration as well as by members of other collaborations, including SBND, ArgoNeut, and MicroBooNE. We agree that a single document with a concise statement of requirements and how they flow from physics needs would benefit the DUNE collaboration and be especially useful for future reviews. We suggest that this topic be revisited once the DUNE simulation and reconstruction techniques are more mature.

The most important requirement for the cold electronics is low noise. Low noise is required to reliably identify hits from isolated minimum ionizing particles passing close to the cathode plane, but even lower noise is crucial to good two-track separation, hit matching (correlating signals in the different wire planes), and vertex resolution. Very low noise is necessary to enable the most ambitious astrophysical measurements. Making a precise quantitative requirement on the noise level, and importantly how sensitive the physics goals are to it, will require more study and may depend on how the raw signal is deconvolved. **The noise should be much less than 1000 e- equivalent**.

The required timing resolution is driven by the need for good two-track separation and vertex resolution. Both vertex resolution and two-track separation are limited by the wire spacing and gaps between planes in the anode plane assemblies. **Given the design of the DUNE TPC, the fastest usable shaping time is ~1 microsecond. Assuming approximately Gaussian shaping, this implies an ADC sampling rate of ~2 MHz.**

The dynamic range requirement is also derived from the need for good vertex resolution. A significant fraction of interactions of the neutrino beam in the liquid argon will produce many more than one highly ionizing proton. The dynamic range should probably be specified by a requirement that the fraction of events in which signals induced by tracks at the vertex exceed the electronics dynamic range be less than a given percentage. The dynamic range and noise

requirements together determine the number of bits that an ADC must have. The current dynamic range requirement of $\sim 10^6$ e- is obtained by considering the energy deposited by five slow protons within 5mm of the neutrino interaction vertex. **Our analysis confirms the desirability of using a 12-bit ADC.** Although LSB-level resolution over the full range is not necessary, a linear ADC provides the most straightforward way to digitize both bipolar and unipolar signals. We note that some aspects of the non-accelerator program (specifically nucleon decay and the study of low energy neutrinos from e.g. supernova events) do not require either microsecond time resolution or large dynamic range and would be better served with longer shaping time if this resulted in lower noise.

- *Review the status of the elements of the cold electronics that have been developed or are currently in development and determine which of these elements:*
    - *Meet DUNE requirements in their current form,*
    - *Are likely to meet DUNE requirements after minor redesign,*
    - *Are unlikely to meet DUNE requirements without significant redesign.*

We limit our response to the status of the cold ASICs.

We believe that the front-end ASIC might meet DUNE requirements in its current form and is likely to meet requirements after minor redesign. The noise performance achieved in small-scale test setups and in MicroBooNE is very good. It is important that all sources of noise in the 35 Ton test be understood, but it is unlikely that this understanding will necessitate a fundamental redesign of the front-end ASIC. A minor redesign could address problematic features of the current ASIC, including the non-uniform 200 mV baseline.

We believe the current ADC ASIC is unlikely to meet DUNE requirements without significant redesign. The "domino" ADC relies on excellent transistor matching, and transistor matching is worse at LAr temperature than at room temperature.

There has not yet been a prototype of the COLDATA communications ASIC, but the current approach using 65 nm CMOS is very likely to succeed. A phase lock loop capable of operating at room temperature, at LAr temperature, and at intermediate temperatures has been fabricated and proven to work well. The characteristics of the 65 nm transistors at LAr temperature have been measured, and "cold corner" models have been developed for use by the CAD tools. A digital library of standard cells using transistors with 90 nm channel length (for extended lifetime at LAr temperature) has been developed for use with the Cadence synthesis, timing analysis, and automatic place and route tools.

- *Propose a plan to complete the development of DUNE cold electronics. This plan should balance ultimate performance, reliability, power consumption, cost, schedule, and risk. Summarize the alternatives that have been considered in the development of the proposed plan. If consensus cannot be reached on the best plan, present pros and cons of the different approaches preferred by committee members.*

We have considered a range of options for the DUNE cold electronics, including the use of commercial electronics and the option of locating only the analog front-end inside the cryostat. The SLAC group has performed a preliminary "system trade study" as input to our discussions (see Appendix). The options included in this study were:

  - Option 1: Retain the currently planned architecture and develop a new ADC.
  - Option 2: Develop a fully integrated system-on-a-chip including all functionalities in a single chip.
  - Option 3: Drive analog signals out of the cryostat and use commercial ADCs located outside the cryostat.
  - Option 4: Use a commercial ADC located in the liquid argon (and perhaps use an FPGA in the liquid argon).

Our consensus opinion is that the risk of failure of commercial electronics in liquid argon (option 4) due to accelerated aging is too high to consider the last option except as a last resort. We also believe that the current plan of locating the ADC close to the front-end amplifier/shaper and multiplexing the resulting digital data is superior to a solution in which analog signals are driven out of the cryostat (option 3).

Given our conclusion that the current front-end ASIC is likely to meet DUNE specifications and that the "domino" ADC is unlikely to meet DUNE requirements, even with significant redesign, we believe that a new ADC must be developed. Given the need for this ADC to digitize bipolar signals from the induction planes and (mostly) unipolar signals from the collection planes, the ADC should be a simple linear ADC. Three ADC architectures, all of which rely primarily on the matching of passive components rather than active components, appear to be viable. The three choices are a sigma-delta ADC, a successive approximation register ADC, and a pipelined ADC.

If the ADC is implemented in 180 nm CMOS it could use the same voltage rails as the Front-End ASIC. On the other hand, if it were implemented in a more advanced technology, the ADC might achieve better performance and/or might dissipate significantly less power.

We recommend that the development of a new ADC be begun. The first step in the development of the new ADC should be the choice of

architecture and the target technology node (these choices are closely coupled). We expect that a first prototype ADC could be designed in approximately one year.

The SLAC group feels strongly that given the tight schedule and the short falls encountered during the development of the current solution a risk mitigation alternative path is required. For risk mitigation purposes, they recommend that both option 1 and option 2 be pursued in parallel until enough results are available to select the better option. Their preliminary trade study indicates that option 2 (development of a single system-on-a-chip) may have many advantages from a system integration point of view and they also recommend also that further work on the trade study be done to quantify and apply weighting and prioritization to the trades.

Appendix: System Trade Study

The path forward must carefully consider not only that the final system will meet the ultimate performance requirements discussed above but also be a viable option in terms of overall risk to schedule and cost.

Summary of available options and path forward

The review has identified four options; three were eventually selected for a so-called system trade study. The goal of the trade study is to provide a qualitative and fair comparison between identified system solutions for a set of metrics chosen as the most important for the project. It is important to stress that the focus for the system study is meant to be wider than focusing on the ASIC itself and rather equally consider advantages/disadvantages of all aspects of the system. While a quantitative study of the trades for each system can be made, this would require more time and effort.

Four solutions were identified as viable but only three were included in the trade study. The fourth option using commercial ADCs was determined to be inherently less attractive and is not included in this study. A summary of the solutions is given below.

Option 1	• Multi-ASIC solution (currently planned architecture):
	  • Use existing Front-end ASIC developed by BNL
	  • Re-design of the ADC ASIC
	  • Design of a digital back-end ASIC multiplexer and serializer ASIC
	  • This includes variation of this solution:
	  - New digital back-end ASIC where the ADC, multiplexer and serializer are integrated into one ASIC

Option 2	• Fully integrated system-on-a-chip solution: • Fully integrated system on chip where the Front-End amplifier, ADC and digital back-end with multiplexer and serializer are all in one single ASIC

Option 3	• Analog readout from chamber using warm ADC and digital backend:
	  • Use existing Front-End ASIC developed by BNL as starting point
	  • May require adding serial decoder/decoder interfaces and further support to meet the requirements from driving analog signals to the bulk-head feed-throughs to the warm ADC electronics up to 20 meters away.
	  • Could possibly be included on the existing ASIC or on a separate cold support ASIC.

Option 4	• Use of Cold commercial ADCs *Not included in the following trade study.*

A summary of the system trade study is presented below. The trades were selected based on their impact on the system and are, by choice, high-level entities. Further studies would be able to drill down into the trades to provide a more quantitative comparison.

Trade	Option 1: Multi-ASIC solution	Option 2: Fully integrated system-on-a-chip solution	Option 3: Analog Cold solution
Performance risk of ASICs (e.g. digital/analog xtalk within ASIC)	***Lower,*** low-noise analog blocks and ADC and high-speed digital backend blocks are on separate ASICs.	***Higher***, all blocks are on same ASIC. But: can be measured early when prototype ASIC is available, then risk can be retired.	*Similar to option 1*
# of ASICs to be designed	***Higher***, amplifier ASIC exists, so need two (ADC and digital backend, one if ADC is integrated with backend).	***Lower***, need only one type of ASIC.	***Lower,*** need one (use existing amplifier ASIC but may need re-design and/or new drive/control ASIC and bulkhead)
System performance risk (e.g. digital/analog xtalk within board/between boards, pick-up)	***Higher,*** more complex boards, more active components on board (18), more traces.	***Lower,*** one active component.	***Higher,*** many long high-performance analog lines.
Power Dissipation	***Higher,*** need inter ASIC drive, plus can't share resources between circuits.	***Lower***	***Higher,*** high power analog drivers.
Reliability (likely scales with number of components on board, interconnections, number of solder connections, vias, traces)	***Lower/worse,*** 3 ASIC types, 18 ASICs on each board (8 Front-end ASICs, 8 ADC ASICs, 2 cold digital ASIC)	***Much Higher/better,*** single ASIC/board, minimum of solder connections, traces, vias, I/O, etc.	**Lower/worse,** many traces will be a challenge for the feedthrough design.
ASIC production & test cost (scales with number and types of ASICs, also scales with feature size technology)	***Higher,*** more ASICs and thus more test setups, possibly higher cost with smaller feature size technology (65nm) used for ADC.	***Lower,*** single ASIC type, moderate 130nm feature size (one type of ASIC)	***Medium***
System complexity, size of PCB, production & test cost (PCB fab., assembly and testing)	***Much Higher,*** more complex PCB, more vias/traces, loading, debugging, test)	***Lower,*** fewer active components on board, less pins to solder.	Medium
System integration, performance risks	Somewhat higher(?), more exposed "antennas", traces?	**Somewhat lower?**	**Very high,** long analog transmission lines, only find out real problems

			very late with full system.
Overall Cost (scales with number and types of ASICs, complexity of PCBs and cabling/interconnections)	***Higher,*** more ASICs, more complex boards, more labor and M&S	***Lower***	Higher, many feedthrough's needing high fidelity, high performance cabling.
Schedule	More manpower may translate to longer duration (depends on available manpower)	Completely depends on results from integrated ASIC. Might need more iterations, may need fewer.	About same as option 1, still needs new ASIC.
Risk	Higher long term risk	Higher short term risk	***Very high long term risk***, system behavior far from characterized before full system test (incl. ground-loops, pick-up, etc.)

While a qualitative study cannot give the full answer, it does provide useful insights. While option 3 do pertain some advantages from the cold electronics perspective, mainly related to being simple, it has a few serious system-level disadvantages that makes it least favorable. These drawbacks are related to the high long-term risk associated with a system based on driving analog lines over large distances. Essentially there are aspects of the system that cannot be tested until the complete system is assembled in place which is a large disadvantage from a risk perspective.

From the above a single integrated ASIC solution (option 2) carry many long-term benefits. It does however present a much higher short-term risk than the current option (option 1) which needs to be considered. We think that this added short-term risk can be mitigated by moving towards a prototype integrated ASIC in one year time-scale. If a working prototype is demonstrated the largest risk will be retired and the advantages of this approach can be exploited. There is a finite risk that the noise performance required cannot be met in a fully integrated solution in time for DUNE.

While a fully integrated ASIC has been demonstrated for other projects (e.g. low-noise pixel cameras for LCLS) it is still not a proven solution in the same sense as option 1. This option, being the current solution selected for DUNE has of course the added benefit of a long history of development that carry weight going forward. Nevertheless, this solution does have significant amount of development and testing ahead of it with challenges and associated risks to be retired. Some of the risks are considered more long-term than a fully integrated ASIC solution due to the added complexity of the system.

Considering the system performance risks and the time left until planned CD-2, it is suggested that more than one path forward is pursued until working ASICs are available and some risks can be retired.