Cold Electronics Task Force Report

David Christian

May 16, 2017

DUNE Plan for Cold Electronics

- 128 channel Front-End Mother Boards including 3 types of ASICs:
  - 16 channel amplifier/shaper (8 chips)
  - 16 channel ADC (8 chips)
  - 64 channel communications ASIC (2 chips)

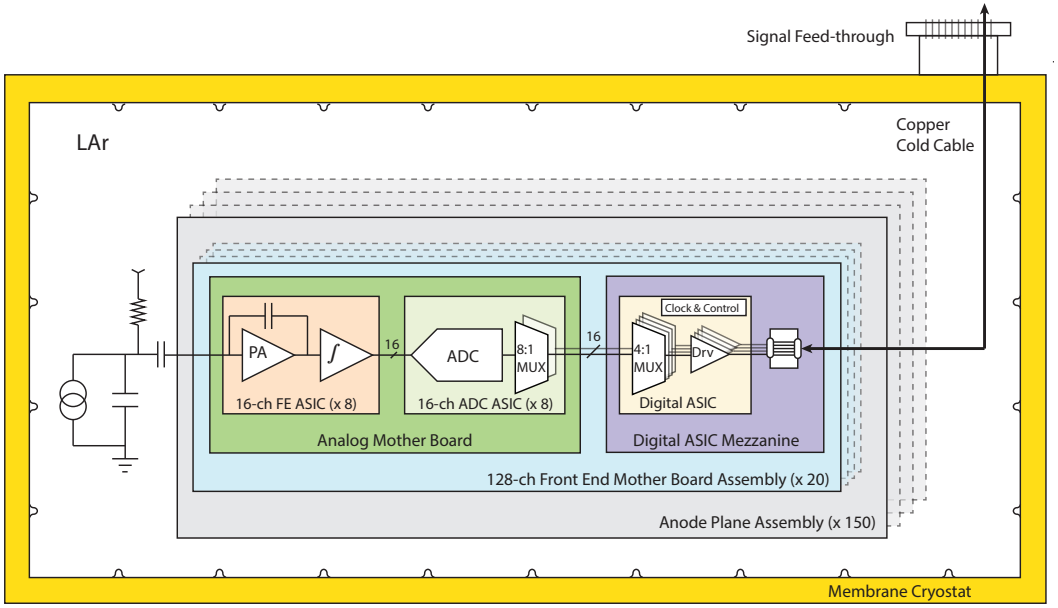

October Review

- Prototype ADC had problems including non linearity and “stuck codes.”
- Review committee expressed doubt that a 12-bit “domino” ADC could work in LAr, and also doubt that a 12-bit ADC is required.
- In February, Eric James responded by establishing a task force to review DUNE’s plan for custom integrated circuits and recommend a path forward.

Task Force Membership

- Group of physicists and engineers from the four US/HEP labs with ASIC design groups:
  - David Christian, FNAL (Chair)
  - Gregory Deptuch, FNAL
  - Hucheng Chen, BNL
  - Emerson Vernon, BNL
  - Per Hansson Adrian, SLAC
  - Angelo Dragone, SLAC
  - Daniel Dwyer, LBNL
  - Carl Grace, LBNL

Task Force Meetings

- Members worked independently and ten teleconference meetings were held, from Feb. 22 – May 10.
- Meeting topics included:
  - Review of DUNE goals & associated physics signals
  - Signal formation in a LAr TPC
  - Lifetime of MOS transistors at LAr temperature
  - Status of FE ASIC
  - Discussion of dynamic range and sampling frequency requirements
  - Review of ADC architectures
  - Discussion of options/response to the charge

Charge #1

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.

- A large amount of work on requirements has been done by members of the DUNE collaboration as well as by others including SBND, ArgoNeut, and MicroBooNE.
- We focused on the requirements of the front-end amplifier/shaper and ADC.
- We took as constraints the existing TPC design (including drift field, maximum drift length, Anode plane assembly details such as wire pitch and gaps between wires) and estimates of drifting electron lifetime.
- Even with this narrow focus and constraints that limit the range of tradeoffs that can be made, we failed in the attempt to concisely state requirements, but we confirm the correctness of the requirements that have been developed in the past.

Noise Requirement

- Low noise is of paramount importance.
  - Low noise is required to reliably identify hits from isolated MIPs passing close to the cathode.
  - Even lower noise is crucial to good two-track separation, hit matching (correlating signals from the different views), and vertex resolution.
  - Very low noise is necessary to enable the most ambitious astrophysical measurements.
  - ***The noise should be much less than 1000 e- equivalent***.

Timing Resolution Requirement

- Derived from the need for good two-track separation and vertex resolution.
- ***Given ~5 mm anode wire spacing and ~5 mm gaps between anode wire planes, the fastest usable shaping time (assuming Gaussian shaping) is ~1 microsecond.***
- ***This implies an ADC sampling rate of ~2 MHz.***

Dynamic Range Requirement

- Derived from the need for good vertex resolution.
  - A large fraction of events include >1 highly ionizing proton.
  - Specification should be extracted from simulation by requiring that less than (10%) of events include signals from tracks near the vertex that exceed the dynamic range of the electronics.
  - The current dynamic range requirement of ~10^6 e- comes from the energy deposited by 5 slow protons within 5 mm of the primary vertex.

ADC Requirements

- The least significant bit (LSB) bin width follows from the requirement that the ADC not add significantly to the electronics noise.
- 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.
- ***Our analysis confirms the desirability of using a 12-bit 2 MSPS ADC***.

Charge #2

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 limited our response to the status of the cold ASICs.

FE Status

- We believe that the FE 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 this will necessitate a fundamental redesign of the FE ASIC.

ADC Status

- We believe that 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 much worse at LAr temperature that at room temperature.

Communications Chip Status

- There has not yet been a prototype of the COLDATA communications ASIC, but the current approach using 65nm CMOS is very likely to succeed.
- A phase lock loop has been made and proven to work well.
- The characteristics of the 65nm 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 90nm channel length (for extended cold lifetime) has been developed for use with the Cadence synthesis, timing analysis, and automatic place and route tools.

Charge #3

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.

Options Considered

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

Consensus Opinion

- The risk of failure of commercial electronics in LAr (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).

Consensus Opinion (continued)

- We believe a new ADC must be developed.
- Given the need 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, appear to be viable:
  - Sigma-delta
  - Successive Approximation Register (SAR)
  - Pipelined

Consensus Opinion (continued)

- We recommend that the development of a new ADC be begun. The first step should be the choice of architecture and the target technology node (these choices are closely coupled). We expect that a first prototype could be designed in approximately one year.

SLAC Trade Study of Options 1-3

Trades	Option 1: Current 3 ASIC solution	Option 2: Single ASIC solution	Option 3: Analog Cold solution
Performance risk of ASICs (e.g. digital/analog xtalk within ASIC)	**Lower risk,** low-noise analog blocks and ADC and high-speed digital backend blokcs are on separate ASICs.	**Higher risk**, 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 needs new drive/control ASIC)
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: 1 ASIC/board, minimum of solder connections, traces, vias, IO, etc.	Lower/worse (lot of feedthru's)
ASIC production & test cost	Higher (more ASICs, scales with types of ASICs, test setups). Also scales with feature size of technology used	Lower (one type of ASIC)	Medium
System complexity, size of PCB, production & test cost (PCB fab, assembly, test)	Much Higher (more complex PCB, fab cost higher, more vias/traces, loading, debugging, test)	Lower (less 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, only find out issues very late
Overall Cost	Higher, needs more ASICs, more complex boards, more labor and M&S	Lower	Higher (many feedthru's, cables)
Schedule	More FTEs so longer duration? (Depends on FTE avail)	Completely depends on results from integrated ASIC. Might need more iterations maybe fewer?	About same as option 1, still needs new ASIC.
Risk	Higher long term risk	Higher short term risk	Very high long term risk (incl ground-loops, pick-up, etc)

Second Opinion

- 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 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 may have many advantages from a system integration point of view and they also recommend that further work on the trade study be done to quantify and apply weighting and prioritization to the trades.