DUNE Requirements and short/long-term plans

for SBN-DUNE synergy meeting

Friday November 20th 2015

Presented by Giles Barr

Introduction

- We have collected the information for four configurations of DUNE
  1. ProtoDUNE – single phase prototype
  2. WA105 – dual phase prototype
  3. One 10kt module of the final DUNE, single-phase detector
  4. One 10kt module of the final DUNE dual-phase detector

- The ordering of the questions here is the same as in the template sent round last week (hopefully). The questions are in this colour.

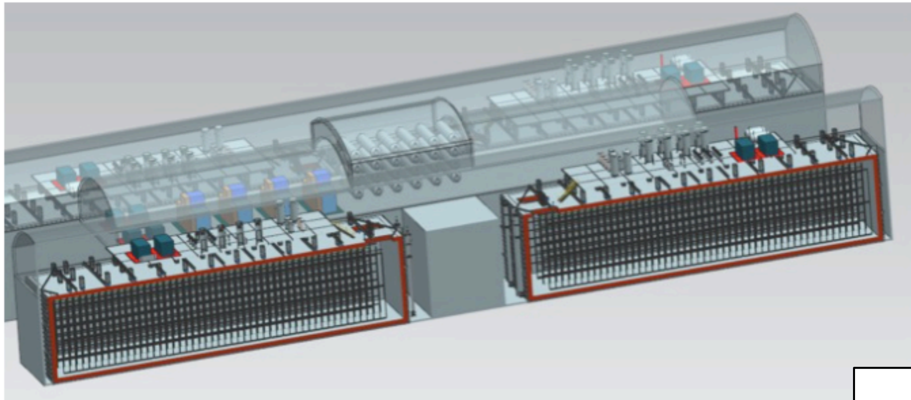

4 identical cryostats W, H, D, 15.1m x 14.0m x 62m

The four 10-kt modules can use the single-phase or dual-phase options

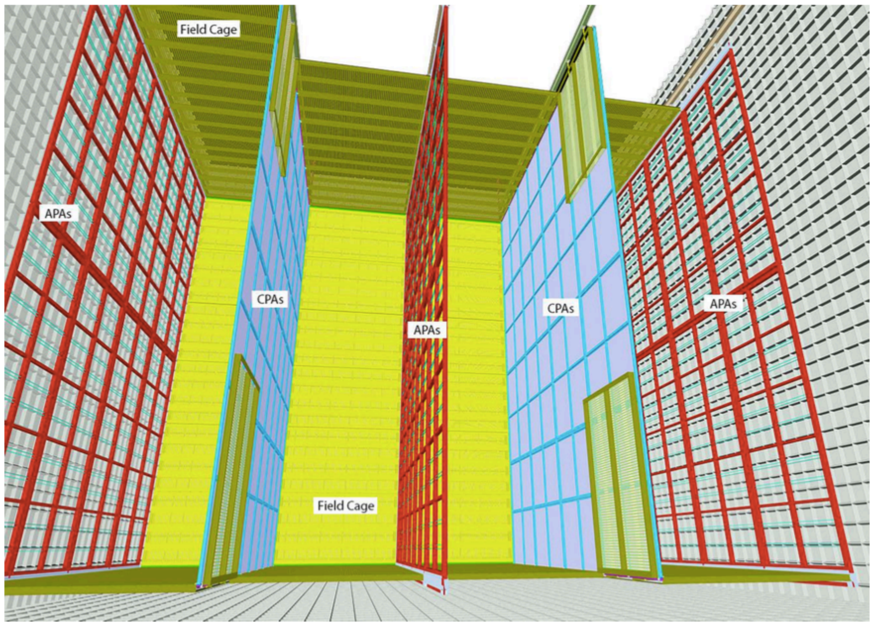

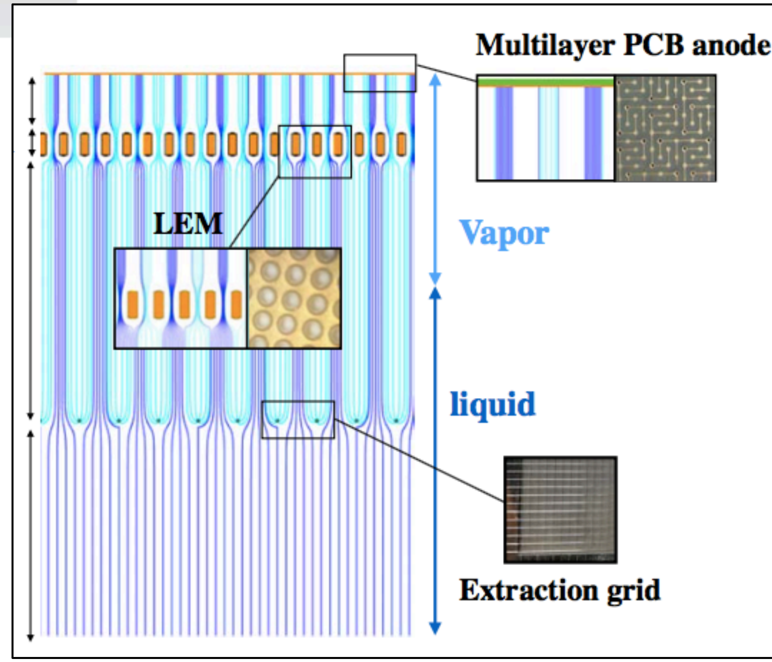

Single-phase (liquid only)

CERN-platform -prototype is ProtoDUNE

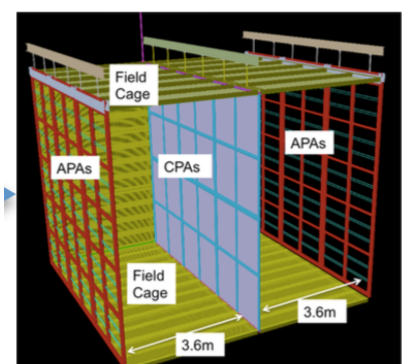

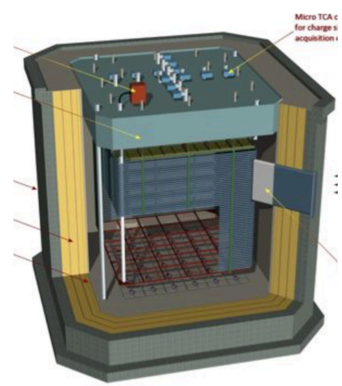

Dual-phase (liquid & gas)

CERN-platform -prototype =WA105

Timescales

- Give an indication of timescales
  - For protoDUNE, aim is to take beam data by June 2018, i.e. stable DAQ by March 2018. We need to urgently solidify our plans now, both technically and for identifying funding sources (together).
  - For DUNE, install one 10kT sub-detector per year. The electronics and DAQ need not be identical for each and a scheme for running different systems is sketched out in CD1-R. 1^{st} 10kT being installed in 2021.
  - Latest the technology decisions can be taken for 10kt #1 is likely around 2018 (intention is to be as late as possible to benefit from newer technology, and for funding options to be clear)
  - For other 10kt modules, wait at least until the single-phase/dual-phase decision is made, and possibly later.

There is plenty to do right away: (1) ProtoDUNE, (2) Carefully understand and plan the algorithms and data handling on whichever platform is used.

Handling of links from Front-End electronics (1)

	Single Phase		Dual Phase	
	ProtoDUNE	One 10kt module	WA105	One 10kt module
#sub detectors	1	1 to 4 modules	1	1 to 4 modules
Data pretrigger per detector GB/s	46	1,152	38	768
Data posttrigger per detector GB/h	O(2000)	O(500)	O(1600)	O(400)
Supernova event size (no compress)		11 TB		8 TB
#channels per APA	2560			
#APA in detector	6	150		
#channels in detector	15360	384,000	7680	153,600
Digitization frequency, #bits, data size	2 MHz, 12bit 12 bit		2.5 MHz 12 bit 16bit	
Drift distance (time)	3.6 m horizontal (2.4ms)		6m vertically (4ms)	
Beam trigger rate	100 Hz	1 Hz	100 Hz	1 Hz
Beam window	4.5 s	10 us	4.5 s	10 us
Time between spills	?	1 s	?	1 s
Other important triggers		continuous		continuous

Raw rate for 10s, no compression

To estimate these numbers: No compression (yet)
Underground: 1 trig/s, 2 drifts, 3% ROI
At EHN1: 50 trigs/s in beam, 1 drift, 30% ROI

- How many sub-detectors will be deployed?
- What's the data throughput for each sub-detector? And overall?
- Explain these requirements in terms of number of channels, amount of data per channel, digitization frequency, trigger window duration (e.g. how long is a drift time) and trigger rate.

Handling of links from Front-End electronics (2)

- How many links, and of which type, you will need to handle from Front-End electronics? What's the maximum level of compression that can be realized in order to reduce the total number of links?
  - Single phase links with ADCs in cold, push data for 128 channels on one link (so a bit under 1Gbit/s). A synchronization command is required at the 2MHz digitization rate.
  - Links further downstream have not been determined yet, could be commodity Ethernet
  - Compression studies will be done with 35t data
- What synchronization is needed among the different elements of the FE and possibly with other external inputs?
  - Rigorous synchronization at (a) 64MHz (=PD digitization rate), (b) 1-PPS, (c) Run starts/ends is required. GPS tag needed at far detector to match with spills at near detector
  - Careful monitoring that synchronization is good is required, tagged with each event. (Calibration pulser events, read counter from last second at each 1-PPS, test pulse input).

Integration of the trigger system

- Do you need any backpressure of the DAQ towards the trigger distribution?
  - For protoDUNE – yes, applied at the trigger distribution
  - For DUNE underground, data are taken continuously (to be sensitive to NDK, atmospheric, etc.) so a 'diagnostic' backpressure system is needed on the continuous streams before the trigger is defined, **however in normal operation, it should NEVER need apply dead time**. Additionally, a backpressure system at the point the data are triggered is needed
- What hardware will be used to implement the trigger logic?
  - protoDUNE – Trigger counters, beam Cerenkov, beam gate
  - DUNE – Computers, asynchronously with data collection
- Will the same trigger signal be distributed to all FE units of the different sub-detectors or is a segmentation of data acquisition foreseen (e.g. Region of Interest)?
  - protoDUNE – Unknown, needs study soon
  - DUNE – ROIs will be needed, we realized this while writing CD1-R.

Event building

- Should the data packets coming from all sub-detectors be merged into one single file?
  - Ideally yes. For merging the four 10kT sub detectors in final DUNE, this could be done 'offline' about an hour after data taking, to maintain the independence of the four DAQ systems.
- What additional information will be associated with each "event"? Will this be stored in an event header or is a different structure foreseen?
  - Slow control data will be stored in a separate database. However, the facility to insert 'out-of-band' data into the event stream would be useful.
- Is there any specific requirement on the data format?
  - Offline requires art format files, however, DAQ could write something else and then it can be changed (e.g. swizzled). Currently we plan to use artDAQ and write art files.
- Is any compression needed? If so, of which order and at which point of the event building?
  - For both protoDUNE and final DUNE, this still needs careful study. (Likely answer is YES!, input will come from 35t)

Data flow

- Do you foresee any zero suppression? If so, at which point in the chain?
  - protoDUNE: Unknown, physics impact of doing it needs urgent study, we do not yet have a fully worked out scheme for dealing with all the data protoDUNE will generate. (Note, since the final detector can take data with no zero suppression, we must have the capability of collecting data like this in protoDUNE)
  - Full DUNE: No, not before event building, except perhaps for events for very low energy-neutrino studies. Hit finding is used in some schemes for forming trigger primitives, but full data is kept for positive trigger decisions.
  - Full DUNE: Potential to have 'pre-offline' zero suppression, to reduce the data load on the offline. However, the DAQ should not limit the capability to collect the full data.
- Do you need to handle multiple streams of data packets for each "event"?
  - Yes

Monitoring tools

- What are the mandatory checks of integrity of the data? Is the information contained in the header enough or will the opening and reading of the files be needed?
  - Online monitoring will look at all data to make hit maps etc.
- Is there any specific need on the type of interface of the monitoring tools? Are they supposed to be interactive, e.g. with the possibility to remotely apply changes to some plotting parameters, or passive – i.e. user looks at plot only?
  - Currently, user has lots of plots to look at, but passively. This allows archived plots to be useful for comparison.
- What features are foreseen for keeping track of data flow?
  - Some rudimentary ones are available for the 35t, but we are looking for better ones.
- Which nodes of the processing chain must be remotely accessible? With what level of protection?
  - Unclear as yet. We need the capability to run remote shifts, but maybe we can push the needed data to the remote shifter for increased security?

Block diagrams + level of 'fixedness' of design

Provide simple block diagrams of the data flow, from Front-End electronics to post-processing, indicate for each piece if:

- it already exists, under development or still to be designed;
- it is strictly defined or open to slight changes or can be completely substituted with a different solution.

- Block diagrams are on the next page

For DUNE, of course, **we are open to new ideas** as we are still searching for new funding opportunities.

For protoDUNE, **we need to hurry up and decide**, even though we are still looking for funding there as well.

An important difference between DUNE and all the other detectors being considered today is that DUNE is very big (1.5M channels in 40kt). So it is **important to consider the strategy to divide the work up** as well as the technical choice.

Block diagram

- A generic pictures of the DUNE DAQ that is emerging below, where the 'Trig hit finder' module generates trigger primitives, and the 'Full-data buffer' retains non-zero suppressed events until the trigger decision has been made.

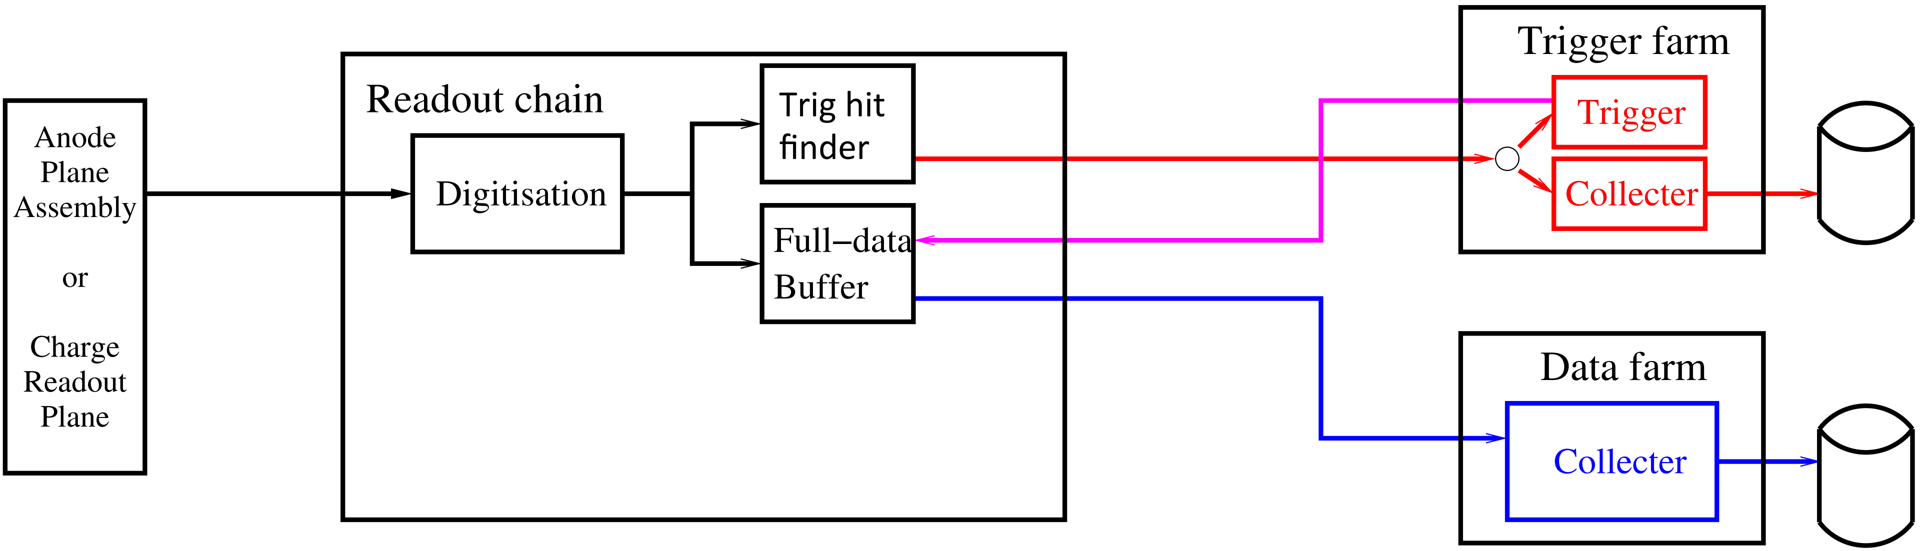

- In this general concept, the location of the 'Trig hit finder' and 'full data buffer can change to optimize the design – in FPGAs/SOCs, or commodity computers. Of course, where they both happen in one place, these functions need not be very distinct. – see diagrams on next slide.

Block diagram (2)

THF = Trigger hit finder
Buffer = Full data buffer

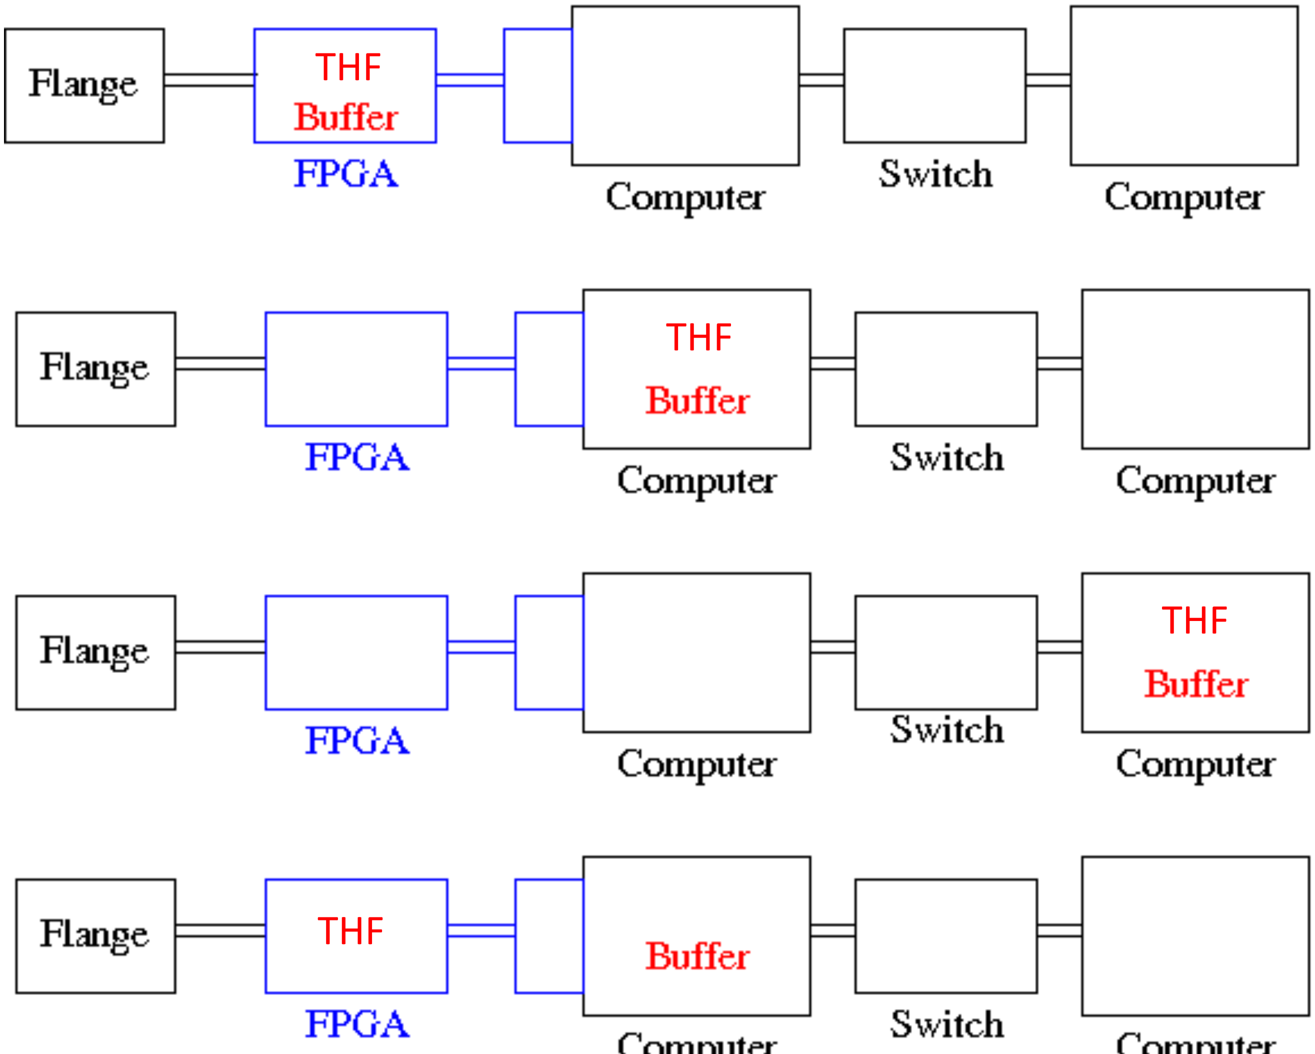

Synchronization – discussion starter questions

Several layers:

(1) Sync data bit streams to start of the digitization tick
(2) Ensure all FADCs 'flash' simultaneously
(3) Count clock ticks since previous fiducial
(4) What is optimum fiducial interval
   – (e.g. 1 frame, 1 pulse per second, start of run?)
(5) Start run and end run concurrently everywhere

Cross checks:

- If we distribute triggers by timestamp, then it is helpful to have a crosscheck – do we need a separate test-pulse distribution network? Can the calibration system do it for us?
- If we distribute triggers by 'NIM pulse', i.e. a rising edge occurs a fixed delay after the time of interest, we get a cross check

In each case, how do we record the info to do the cross check?

Potential synergy between SBN and DUNE – can we collectively arrive at the optimum timing distribution solution?

Short term plans for DUNE

- Run the 35t prototype and analyze the data
- Define ProtoDUNE plans
- Have set up a task force to make simulation framework for data reduction studies – aim to simulate (a) physics performances of whatever we finally propose and (b) data rate requirements, derandomization issues etc.
- Refine plans for cavern requirements, investigate power consumption for ALL viable alternatives. Solution for option to transfer raw data to surface needed.
- Set up further design task forces:
  - Software framework
  - Hardware solutions
  - Synchronization
  - Run Control
  - Slow control

Plan is to see what input we get from this meeting before starting them.

Some longer term plans for DUNE

- Incorporate many things from prototypes and SNB detector experiences
  - Lessons learned
  - Good design tricks
  - Learning how to do QA and QC – simple diagnostic tools, modular diagnostic tools (i.e. the most technical expert person not required to do it)
  - Working with grounding rules from electronics groups
- Maintain portfolio of solutions to providing DAQ for the four caverns so that funding can be applied for as opportunities arise.
  - In most cases, it is important to be successful for the team applying for funding to own some of the intellectual input to the design.
  - Although the DAQ in the four caverns need not be the same, we should deliberately steer towards whichever is the better solution, rather than maintaining several alternatives.
- Installing and debugging is a huge task, so we need a better plan soon,
  - including identifying teams, and to know what the procurement options and constraints are.