DUNE FD DAQ

Georgia Karagiorgi
Columbia University

DUNE Monthly Collaboration Call
March 8, 2019

DUNE FD DAQ

- The DUNE FD DAQ system design has progressed rapidly in recent months
- This talk summarizes:
  - **DAQ Baseline Design**
    - DUNE FD SP DAQ baseline design has passed internal CDR review in December 2018
  - **DAQ Data Selection System**
  - **Design Justification** (ProtoDUNE and simulations)

DUNE FD DAQ Design

CUC: DAQ Front-End lives here

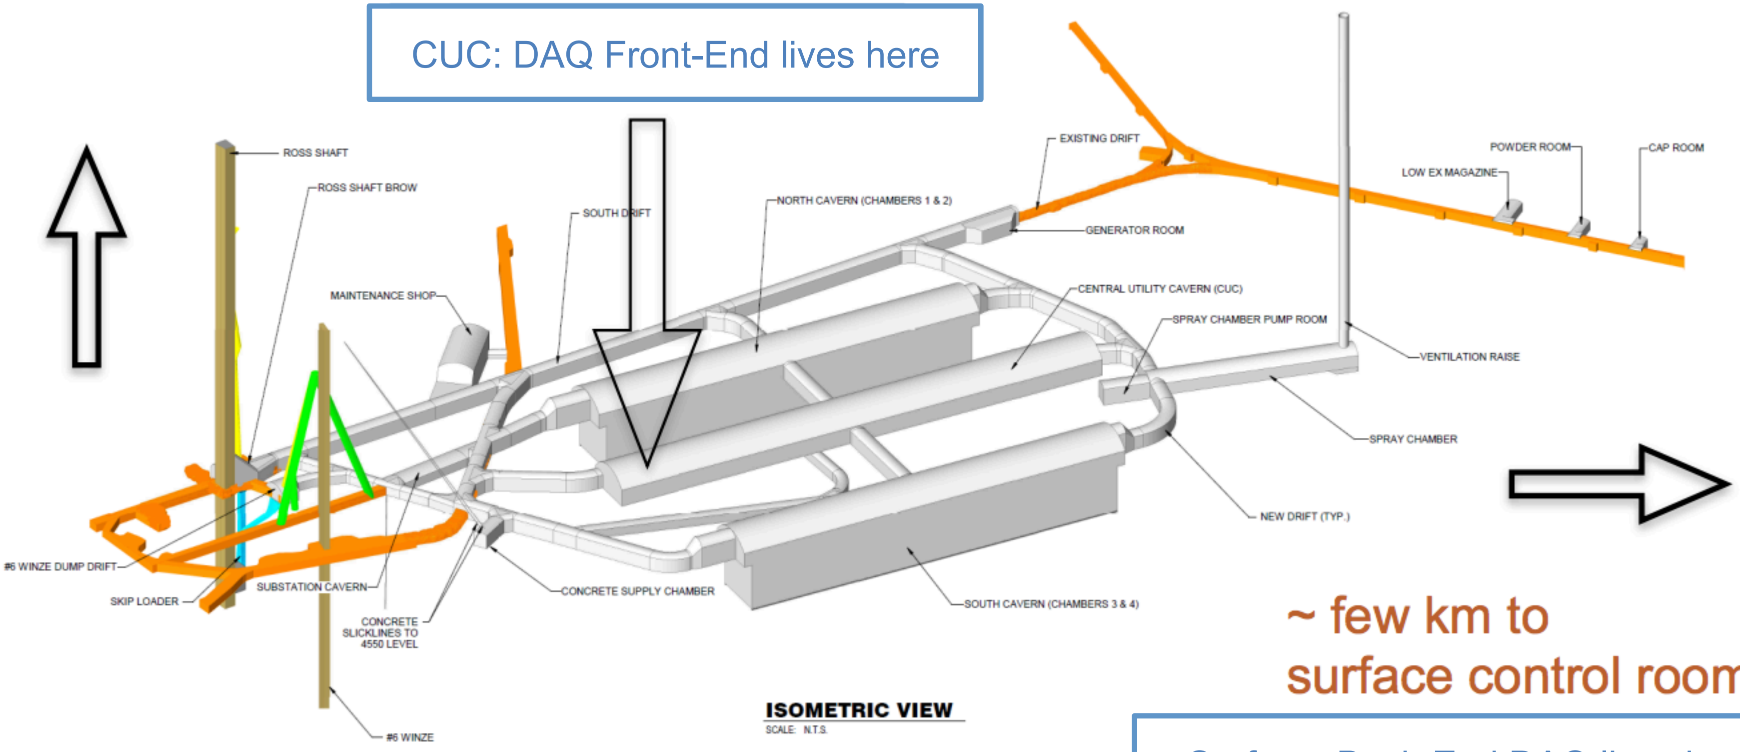

~ few km to
surface control room

Surface: Back-End DAQ lives here

- DAQ is split between 4850ft level and surface: Bulk of processing/buffering underground, minimizing data traffic to surface.
- Strong power, cooling, space constraints in CUC (600kW for all 4 modules, ~50 racks)

DUNE FD DAQ Design (SP+DP)

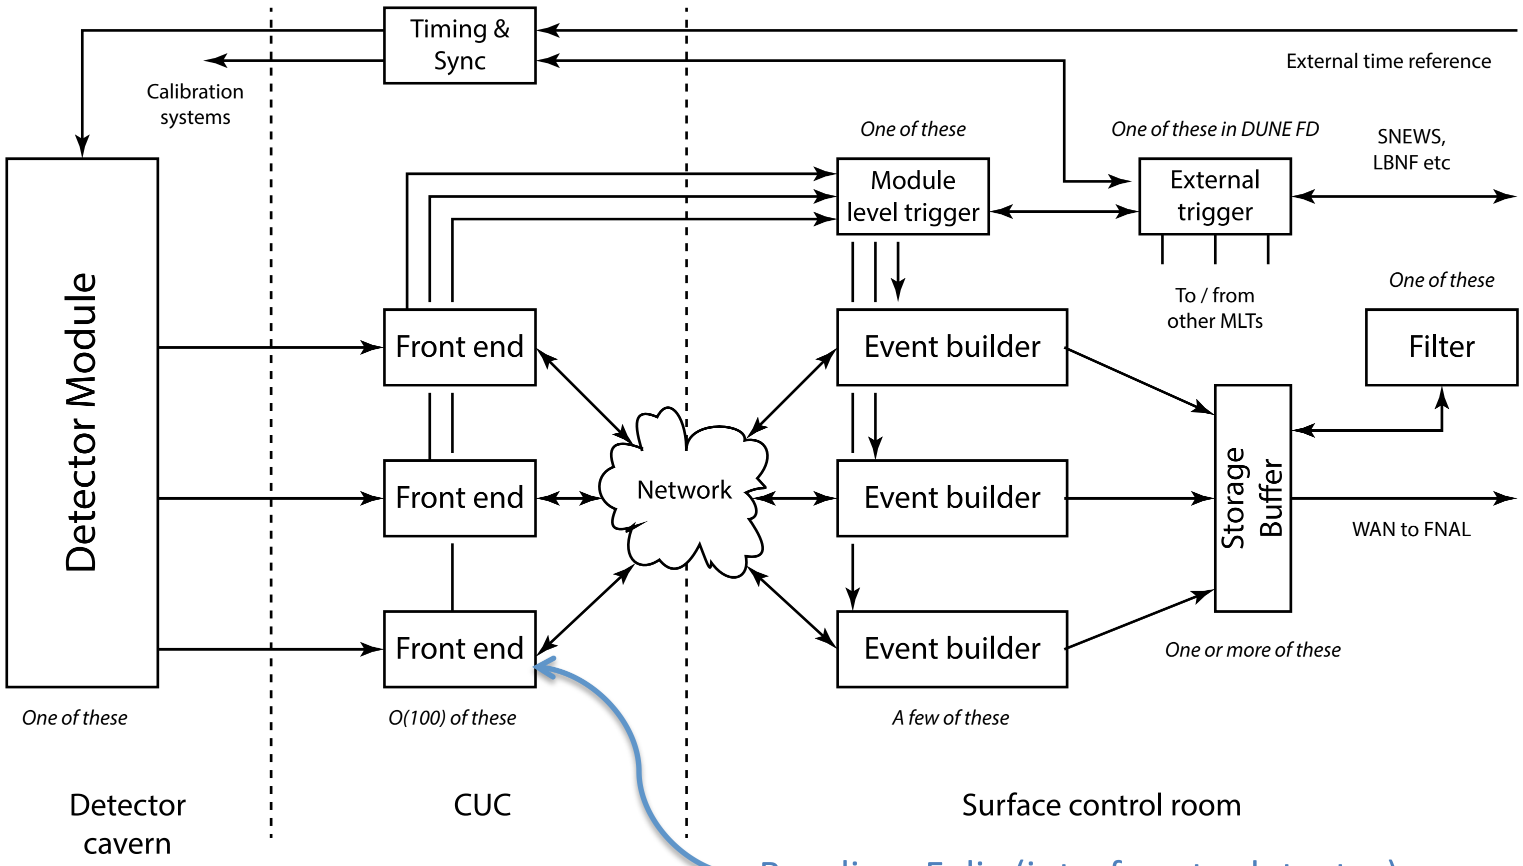

Baseline: Felix (interface to detector)
+ FPGA co-processor + CPU

Not shown: Control Paths

DUNE FD DAQ Design (SP+DP)

- **Scope:**
  - Begins at fibers from the detector (electrically decoupled from cryostats)
  - Ends at WAN network interface (fibers to FNAL via ESnet)
  - Provides common computing and network services for other systems
  - All slow control and safety functions are outside DAQ
- **Functionality:**
  - Provides basic timing and synchronization for sub-detectors
  - Receives, synchronizes, compresses, buffers streaming data from sub-detectors
  - Extracts trigger primitives from data, summarizing local activity in detector; makes local, module, and cross-module trigger decisions
  - Builds “event records” from selected space-time volumes, buffers, and relays those to permanent storage
  - Carries out local data reduction and filtering as required

DUNE FD DAQ Design (SP+DP)

- **Design principles:**
  - A single, scalable system design for all detector modules
  - Ability to self-trigger with zero dead-time and with high efficiency for targeted physics signals and to record triggered data covering the full detector
  - Evolutionary design: begin with very conservative design for first module
  - Preserve possibility of additional capacity as required

- **Key challenges:**
  - Supernova Burst (SNB) requirements: large buffering upstream, low fake trigger rates
  - Long (“permanent”) commissioning state: partition-able system
  - Difficult-to-access location: reliability and remote operation
  - Power, cooling and space considerations

DUNE FD DAQ: SP-focused in next slides

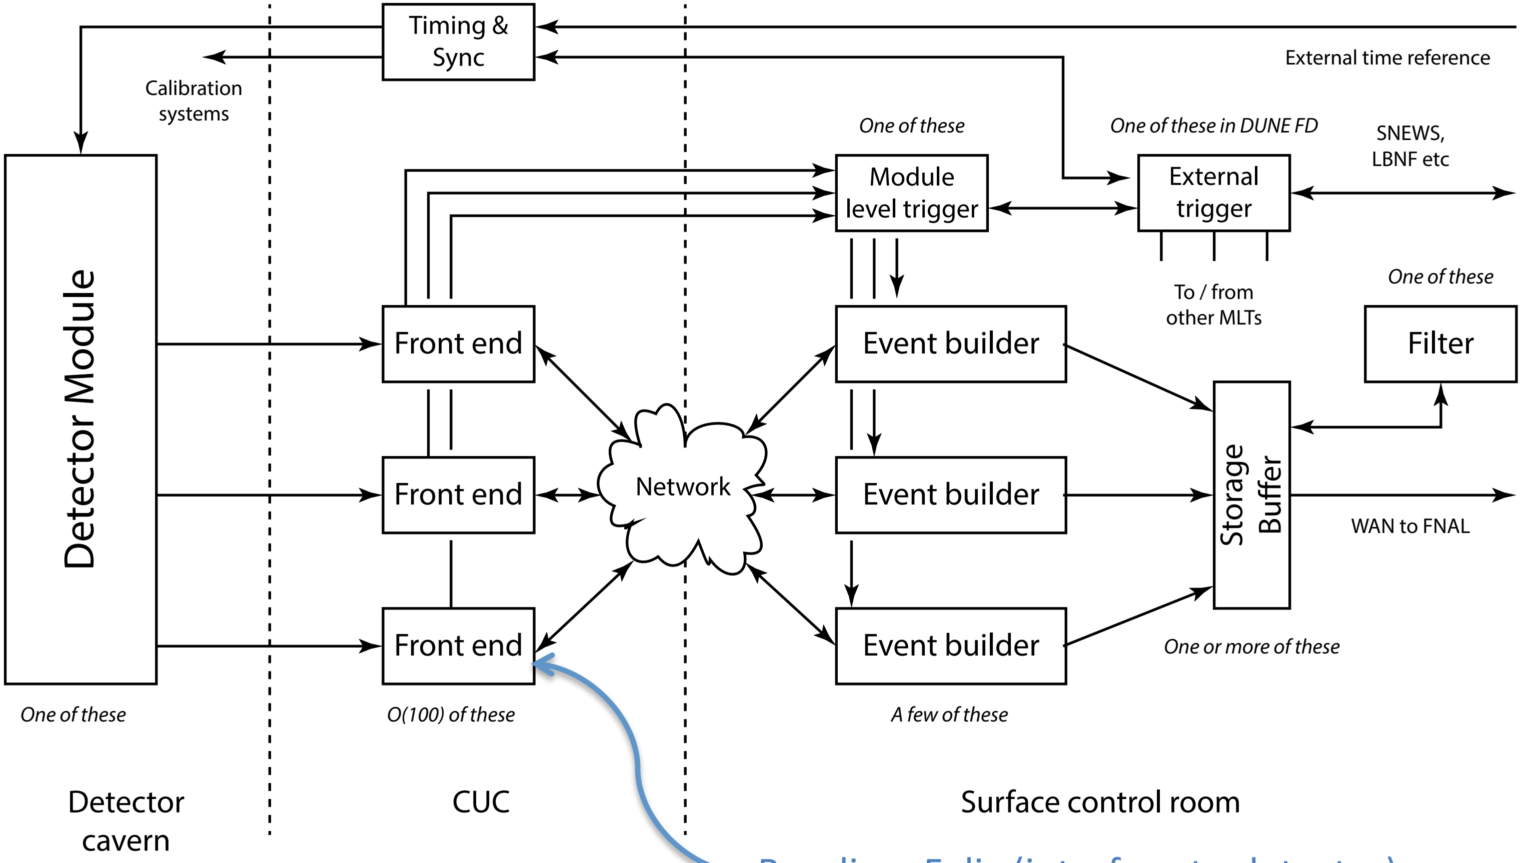

Baseline: Felix (interface to detector)
+ FPGA co-processor + CPU

Not shown: Control Paths

DUNE FD SP DAQ Baseline Design

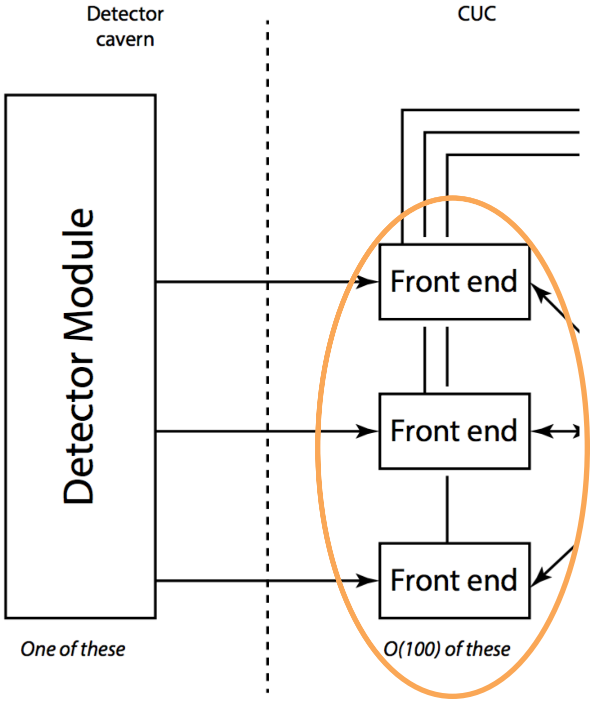

For 1 SP Module:

150 APAs → 150 Detector Links
75 Felix Boards
150 Co-processor (FPGA)
75 Felix Host servers (CPU)
75 Data Selection servers
+networking
+CCM

PDS System → Assumed at a level of 10% of TPC
6-8 Felix Boards
6-8 Felix Host servers
6-8 Data Selection servers
(possibly → on surface)
+networking
+CCM

DUNE FD SP DAQ Baseline Design

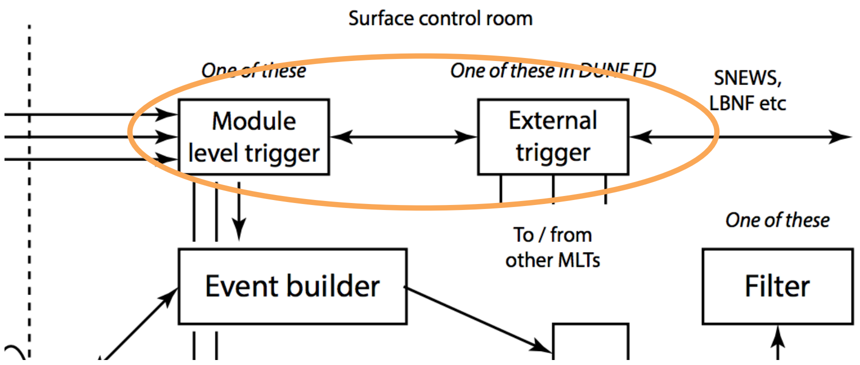

For 1 SP Module:

1 (+1) Data Selection (MLT) server
+networking
+CCM

Common to all four modules:
1 (+1) Data Selection (ETM) server
Timing System Interface
+ networking
+CCM

Also: Trigger Primitive Local Storage System
(~1PB)

DUNE FD SP DAQ Baseline Design

For 1 SP Module:

30 Event Builder server
+networking
+CCM

5 Data Selection (HLT) servers
+networking
+CCM

Storage buffer (~1PB)

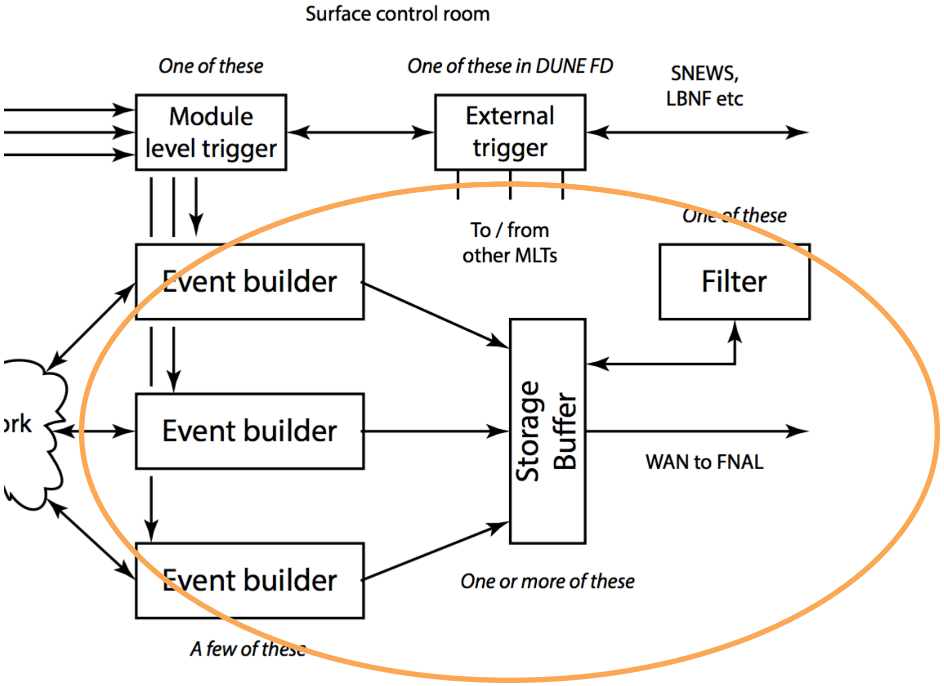

DUNE FD SP DAQ Baseline Design

Timing System:

ProtoDUNE design

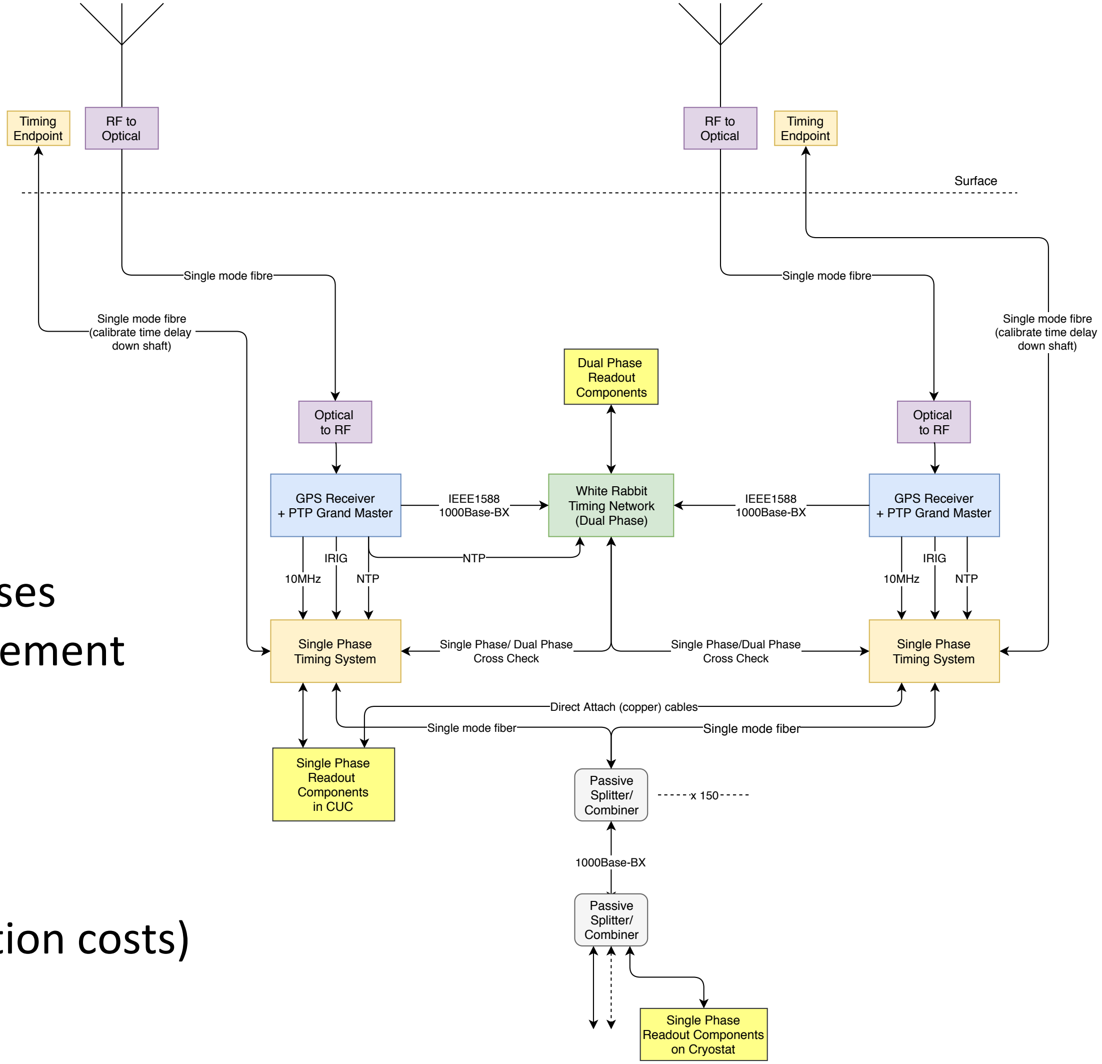

Control, Configuration and Management (CCM) for Back-End DAQ:

4 Servers for Config,
Run Control, Databases

4 Servers for Data Management
and Transfer

Infrastructure:

Facilities costs (16 racks,
professional installation costs)

IT infrastructure

Substantiating Design: 1. ProtoDUNE-SP

Demonstrated Hardware Components:

- **Front-end Readout:**
  2 out of 6 APAs are now read out by FELIX-based Front-end system
  (~DUNE FD system without Co-processor(s) and only 1 APA/FELIX board)

  → interface to front-end electronics, scalability, data flow
  → host server requirements and specifications
  → platform for further system development (co-processor, trigger primitives)

- **Back-end DAQ:** Event Builder farm, CCM machines, disk buffers

  → system partitioning, scalability
  → platform for further system development (CCM, IPC, Data Flow Orchestrator)

- **Data Selection/Timing:**
  (External Triggering)

  → Timing Distribution System

Substantiating Design: 1. ProtoDUNE-SP

- **Key challenge for DUNE FD**:
  (what is **new with respect to ProtoDUNE**, MicroBooNE, SBND, ICARUS…):

  Continuous self-triggering of detector

  - Self-triggering demonstration is being planned for ProtoDUNE
    (target: 2019, and ~2021).
  - Other detectors (e.g. ICARUS, MicroBooNE) have demonstrated self-triggering in coincidence with external gates (data and trigger rates are capped).
    In DUNE-FD: Such ``throttling" limits physics sensitivity to off-beam physics.

→ Need intelligent trigger (Data Selection) design and validation on reliable simulation, plus accurate noise and background predictions.

(This has been the Data Selection WG focus thus far).

Data Selection sub-system

Baseline Data Selection Strategy and Hierarchy: TPC-based

Lowest Level

Trigger Primitives:	“hit on a wire/channel”
Trigger Candidates:	“clusters of hits”
Trigger Command: (prompts “event record”)	“localized (HE) interaction” “extended (SNB) interactions”
Filter:	down-selection of event records

Highest Level

Data Selection sub-system

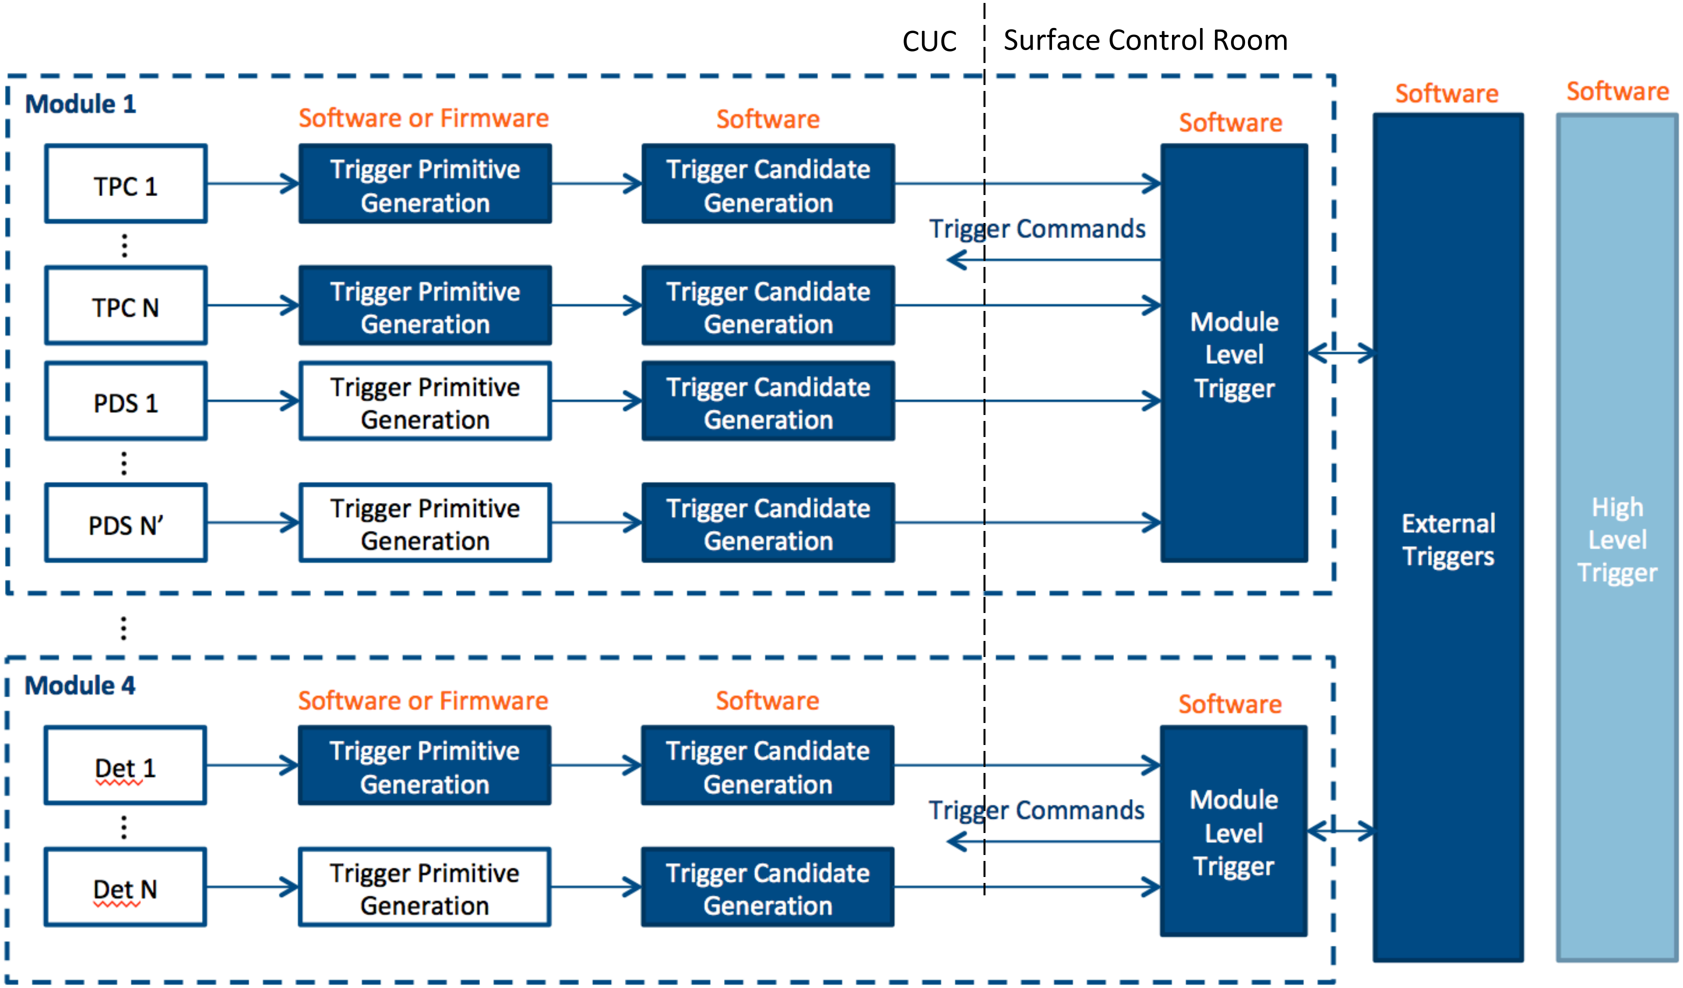

Block diagram of DUNE DAQ Data Selection sub-system, illustrating hierarchical structure of sub-system design.

Substantiating Design: 2. Simulations

Expected Trigger Rates:

Requirement: 100 MeV deposited energy threshold for **“localized (HE) interaction” trigger with >99% eff.** (→ 5.4ms of losslessly compressed data from entire module to event builder; rate up to 1 Hz)

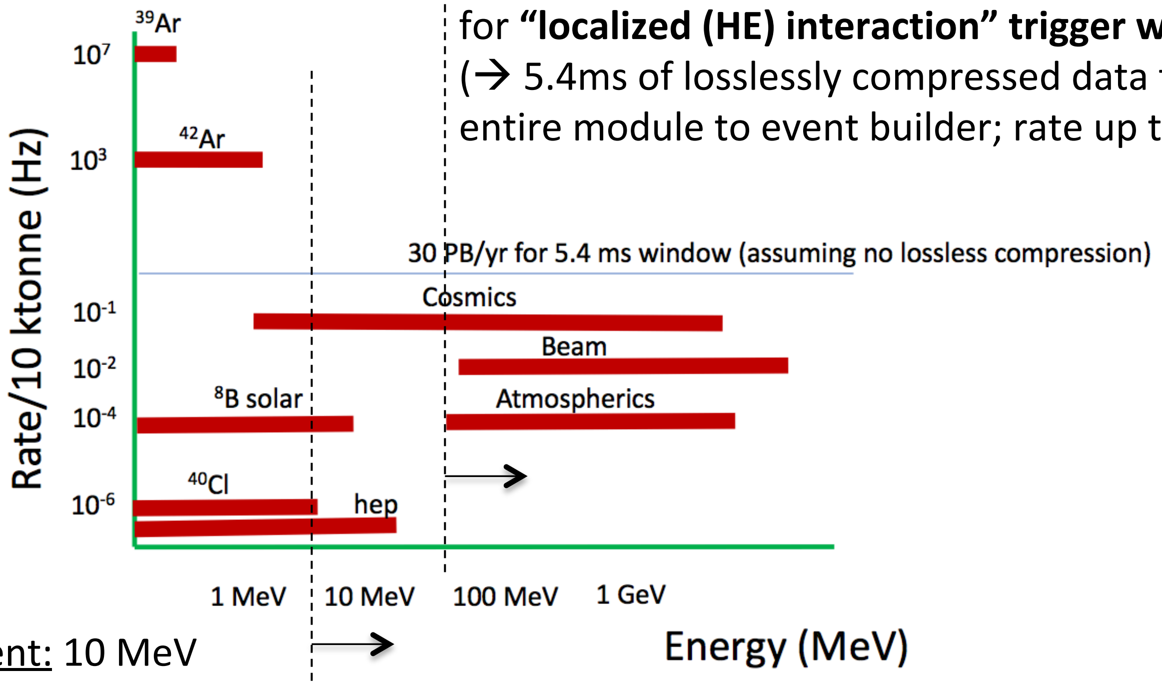

Requirement: 10 MeV → deposited energy threshold for localized LE interactions with sufficiently high eff. as input to **“extended (SNB) interactions” trigger = SNB trigger with >90% galactic coverage** (→ 100s of losslessly compressed data from entire module to event builder; rate ~1/month)

Substantiating Design: 2. Simulations

TPC Trigger Primitive Simulations (SP)

Significant effort dedicated on trigger primitive simulations (see backup slides):

- Charge collection **efficiency and fake rates** due to noise and radiologicals have been studied as a function of hit threshold with MonteCarlo
- CPU **trigger primitive generation speed on CPU** (4 cores) was demonstrated to keep up with expected raw data rates, offline
- Trigger primitive **rates measured at ProtoDUNE-SP** in situ. Effort on understanding and removing contribution from cosmics/cosmogenics and (known) noisy channels is ongoing.
- Full-stream, single-APA, **online trigger primitive generation on CPU** (10 cores) demonstrated at ProtoDUNE SP

Substantiating Design: 2. Simulations

TPC Trigger Candidate Simulations (SP)

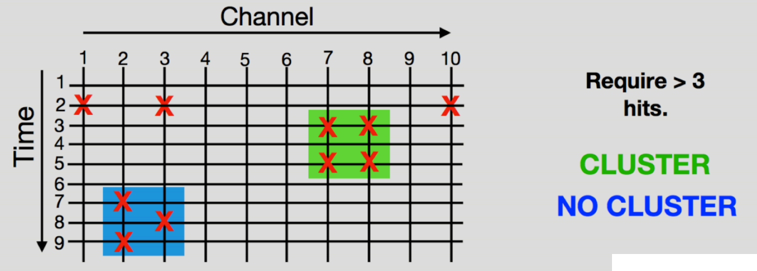

Trigger Candidates formed
by clustering Trigger Primitives
in Channel (collection) vs. Time

Trigger Candidates
with sufficient energy
can serve as Trigger Commands
for **"localized (HE) interaction"**
events: beam, atmospheric
neutrinos, nucleon decay, neutron-antineutron oscillation, cosmics

Simulation demonstrates
>99% trigger efficiency for >100 MeV
visible energy.

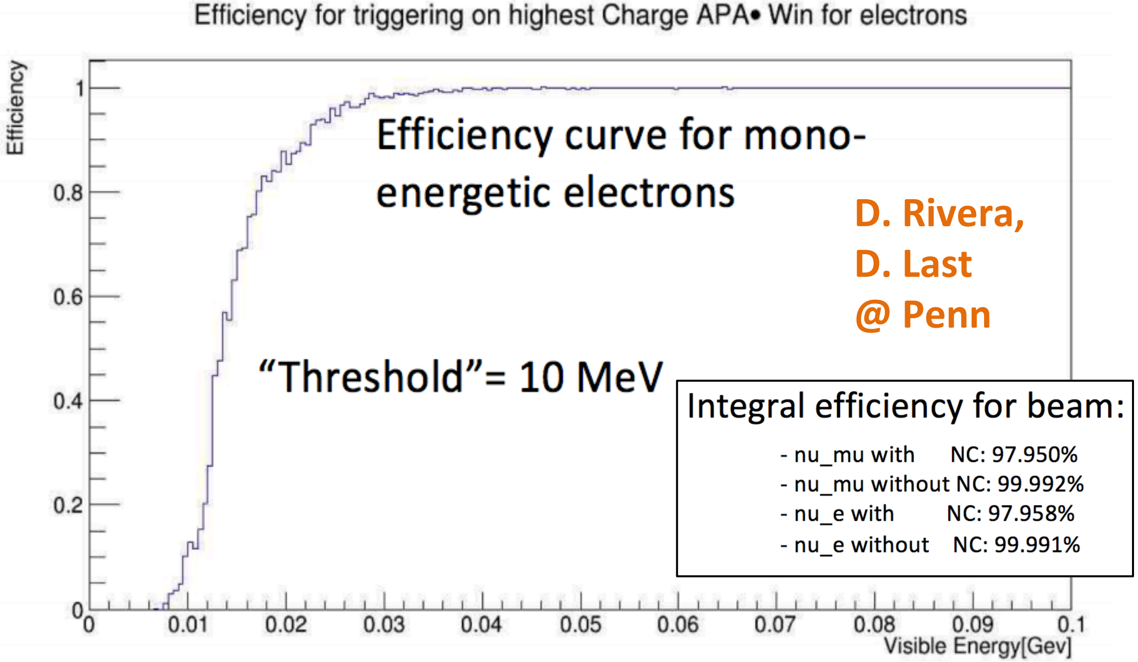

Efficiency curve for mono-energetic electrons

D. Rivera,
D. Last
@ Penn

"Threshold"= 10 MeV

Integral efficiency for beam:

- nu_mu with NC: 97.950%
- nu_mu without NC: 99.992%
- nu_e with NC: 97.958%
- nu_e without NC: 99.991%

Substantiating Design: 2. Simulations

TPC Trigger Candidate Simulations (SP)

Low-energy trigger candidates serve as input to SNB trigger:

Simulation demonstrates ~20-30% trigger candidate efficiency for individual SN neutrino interactions.

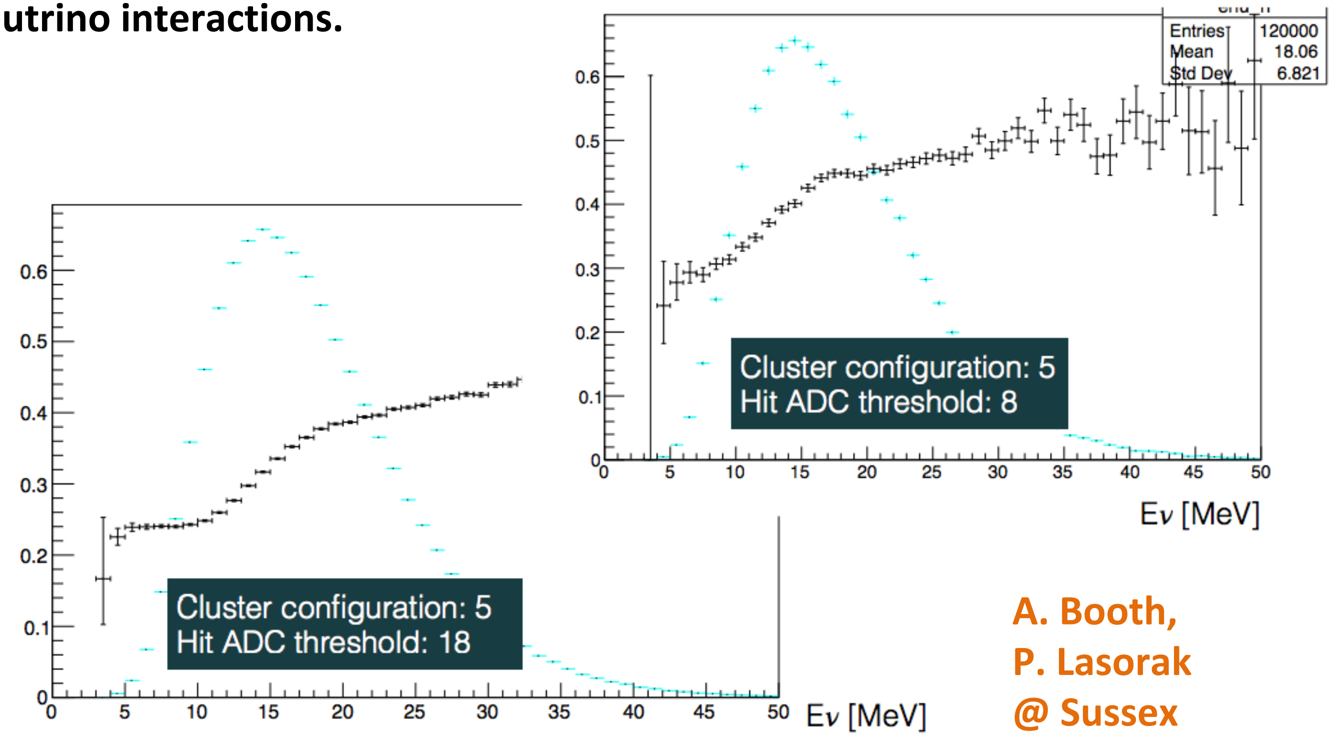

A. Booth,
P. Lasorak
@ Sussex

Substantiating Design: 2. Simulations

TPC SNB Trigger Simulation (SP)

SNB trigger: Multiplicity-based: low-energy trigger candidates over (up to) 10 seconds. (An energy-weighted multiplicity count scheme could be applied to increase efficiency/minimize background.)

A. Booth,
P. Lasorak
@ Sussex

Dominant Radiological Backgrounds

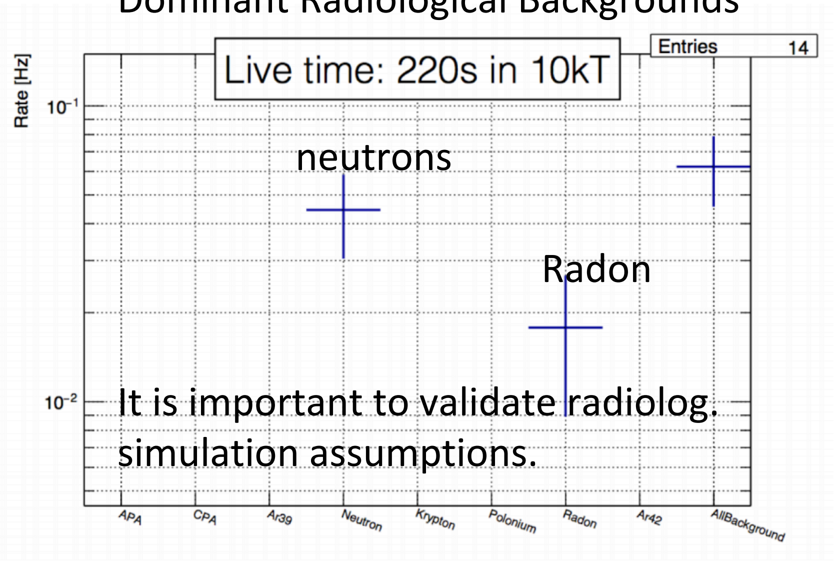

It is important to validate radiolog. simulation assumptions.

SNB trigger studies demonstrate high (>90%) galactic coverage while keeping fake SNB trigger rate to 1/month.

Substantiating Design: 2. Simulations

Trigger Rates and ProtoDUNE-SP Event Builder experience → Extrapolation to DUNE FD SP Back-End DAQ

- **Simulation studies support physics and rates requirements**:
  - \>99% trigger efficiency for "localized (HE) interactions," and
  - \>90% galactic coverage for SNB's

  can be maintained while keeping:
  - "localized (HE) interaction" trigger rate to ~ 0.1 Hz (<1 Hz)
  - "extended (SN) interactions" (SNB) trigger rate to ~ 1/month
- Additionally, **ProtoDUNE-SP Back-End DAQ** has demonstrated our ability to keep up with 1/25th the size of a single DUNE FD SP module, for trigger rates up to 40 Hz and 3 ms readout window (x5 design goal)
- Allows **confident scaling of ProtoDUNE-SP Back-End DAQ** system to DUNE FD SP scale.

Substantiating Design: 2. Simulations

TPC HLT (Filter) SP Simulations:
E.g. CNN-based APA-frame image classification/filtering in GPUs

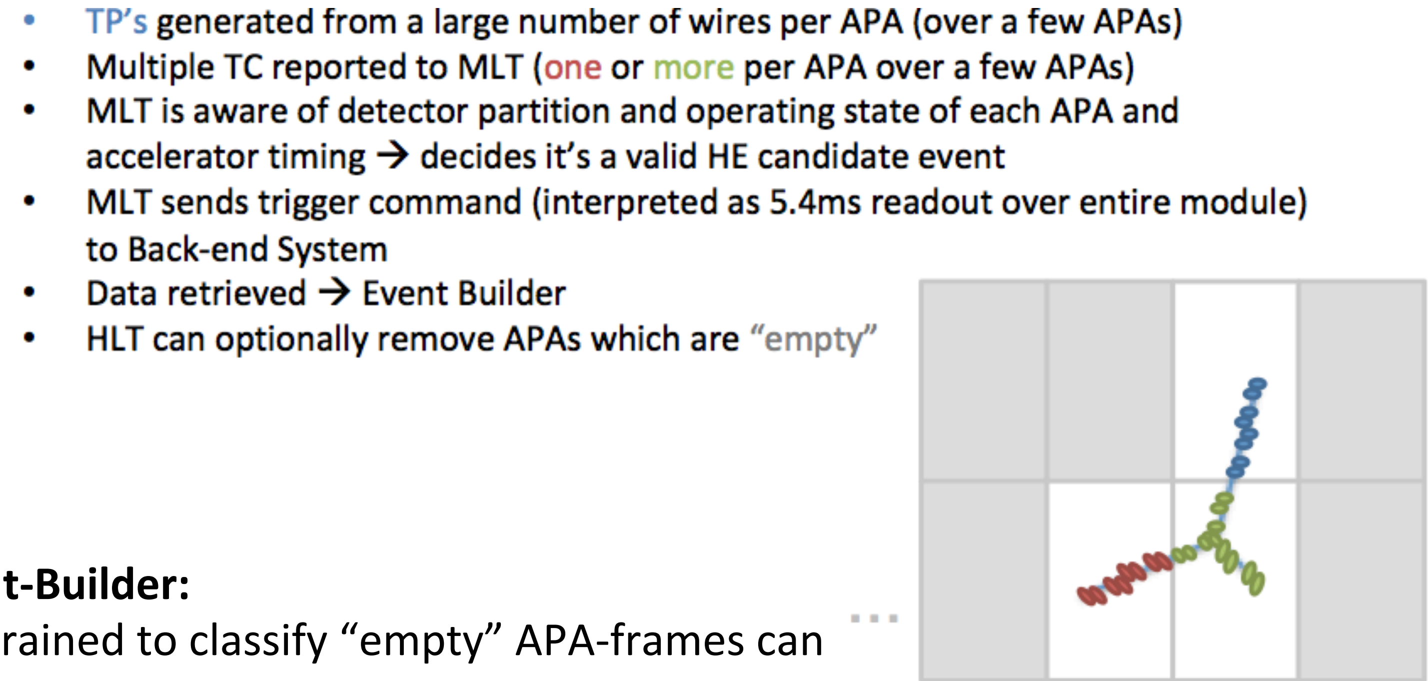

- TP's generated from a large number of wires per APA (over a few APAs)
- Multiple TC reported to MLT (one or more per APA over a few APAs)
- MLT is aware of detector partition and operating state of each APA and accelerator timing → decides it's a valid HE candidate event
- MLT sends trigger command (interpreted as 5.4ms readout over entire module) to Back-end System
- Data retrieved → Event Builder
- HLT can optionally remove APAs which are "empty"

Post-Event-Builder:
E.g. CNN trained to classify "empty" APA-frames can provide further radiological+noise reduction
(e.g. empty frames can be dropped post event builder).

Substantiating Design: 2. Simulations

TPC HLT (Filter) SP Simulations:
E.g. CNN-based APA-frame image classification/filtering in GPUs

- TP's generated from a small number of wires per APA (over multiple APAs)
- Multiple TC reported to MLT (~one per APA over multiple APAs and over an extended period of time (longer than 1 drift))
- MLT is aware of detector partition and operating state of each APA and accelerator timing → decides there are valid LE candidate events → decides this is a valid SNB event
- MLT sends trigger command (interpreted as 100s readout over entire module) to Back-end System
- Data retrieved → Event Builder
- HLT can optionally remove APAs which are "empty"

Post-Event-Builder:
E.g. CNN trained to classify "empty" APA-frames can provide further radiological+noise reduction
(e.g. empty frames can be dropped post event builder).

Substantiating Design: 2. Simulations

TPC HLT (Filter) SP Simulations:
E.g. CNN-based APA-frame image classification/filtering in GPUs

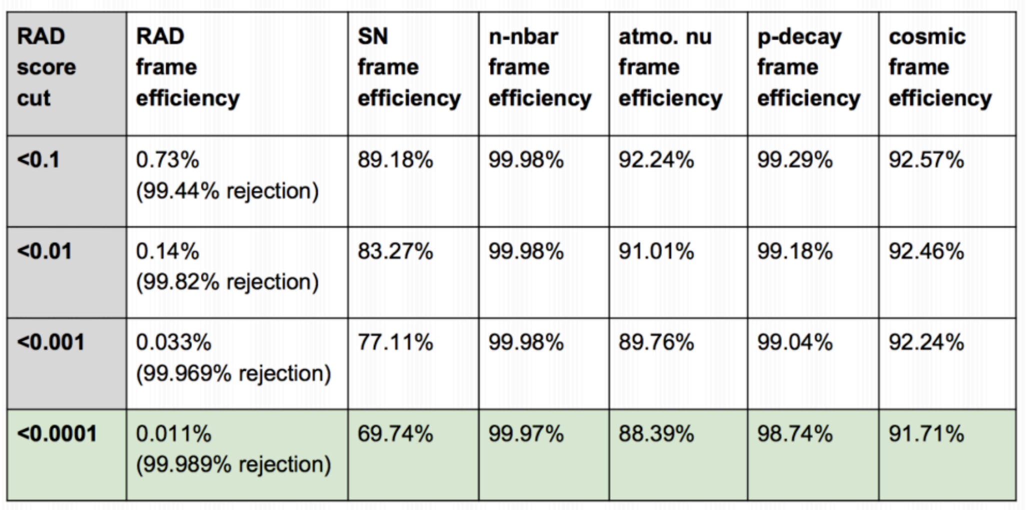

RAD score cut	**RAD frame efficiency**	**SN frame efficiency**	**n-nbar frame efficiency**	**atmo. nu frame efficiency**	**p-decay frame efficiency**	**cosmic frame efficiency**
<0.1	0.73% (99.44% rejection)	89.18%	99.98%	92.24%	99.29%	92.57%
<0.01	0.14% (99.82% rejection)	83.27%	99.98%	91.01%	99.18%	92.46%
<0.001	0.033% (99.969% rejection)	77.11%	99.98%	89.76%	99.04%	92.24%
<0.0001	0.011% (99.989% rejection)	69.74%	99.97%	88.39%	98.74%	91.71%

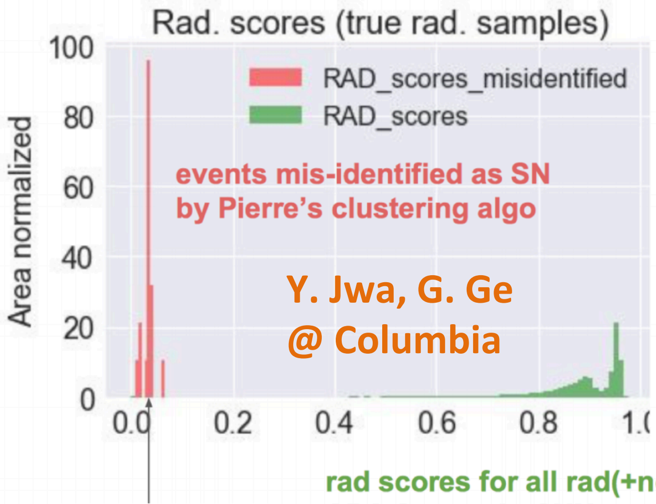

events mis-identified as SN by Pierre's clustering algo

Y. Jwa, G. Ge @ Columbia

rad scores for all rad(+noise) APA-frames in CNN study

This scheme has been shown to correctly tag mis-identified events for low-energy clustering trigger approach.
CNN vgg16b network performance: **20-30ms classification time per APA-frame** on single GPU card.

Closing remarks

DUNE FD DAQ Conceptual Design (Baseline) is complete.

Substantial efforts in recent months have gone into substantiating design and costing, through **simulations and ProtoDUNE-SP** experience.

What has been presented is being incorporated in the **TDR**.

Coming up:

- **ProtoDUNE-DP**
- **Parallel (simulation) efforts ongoing on PDS-based data selection**, for SNB trigger (see talks at January collaboration meeting)
    - SP: P. Lasorak @ Sussex, D. Pershey @ Duke
    - DP: J. Soto, A. Gallego Ros @ CIEMAT

…and groups already thinking about future evolutions of the DAQ design!

Backup Slides

Co-processor design and FELIX integration plans

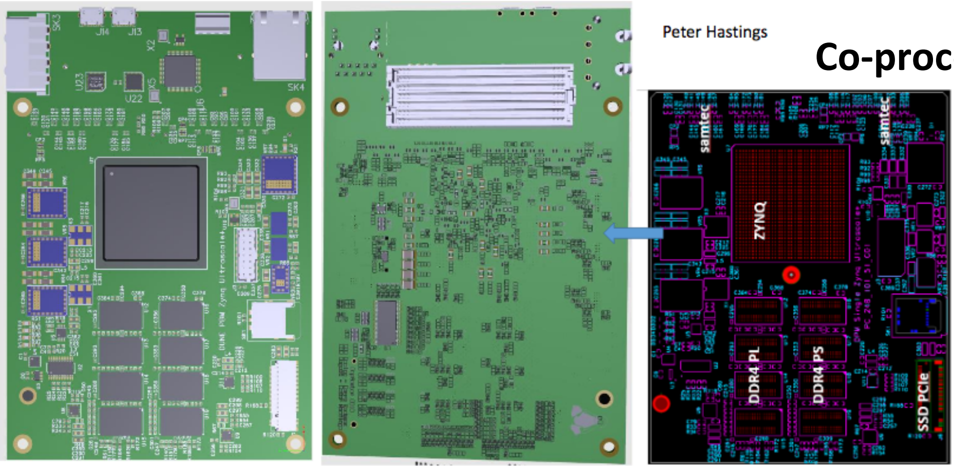

Peter Hastings

Co-processor layout

FELIX layout

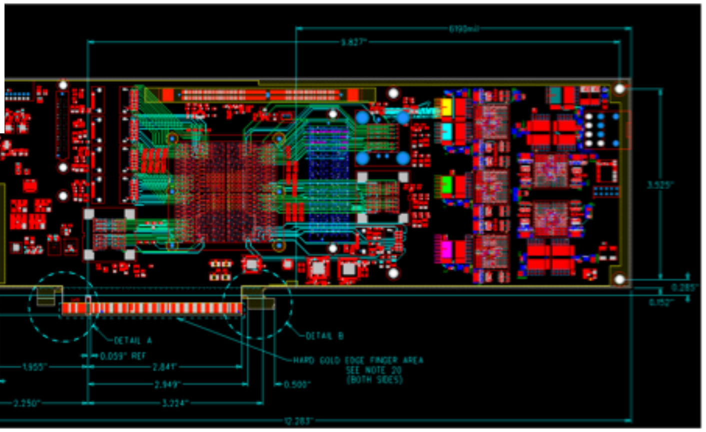

Co-processor firmware structure

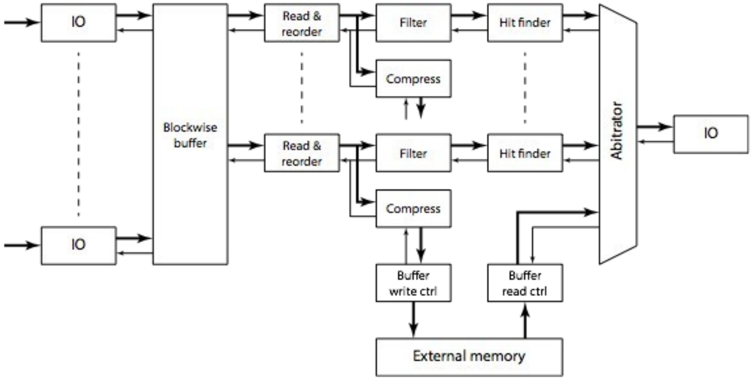

- Substantial progress in next-generation FELIX + prototype co-processor

Co-processor design and FELIX integration plans

- Constraints:
  - Large data volume from Single Phase TPC
  - Relatively poor signal/noise ratio
  - Tight power limits for DAQ
  - Data from all readout (DP , SP / TPC , PDS) received into Felix PCIe board
- Baseline solution: Perform some functions on FPGA co-processor to SP TPC
  - reduce bandwidth to server PCs
  - Reduces total system power (still under investigation)
  - Approach successful at LHC experiments
    - Re-using firmware development infrastructure from CMS
- Functions on FPGA:
  - Hit finding
  - Lossless data compression
  - Data buffering for SNB triggering (O(10-seconds))
  - Storage of losslessly compressed data for SNB triggers (O(100-seconds))

Co-processor design and FELIX integration plans

- DUNE variant of Felix
  - Connection to co-processor
    - Uses VITA standard FMC+ connector.
  - Uses same firmware / software as Atlas Felix
    - …. which is large majority of development effort
- For more details see [talk at DUNE Collaboration meeting, September 2018]

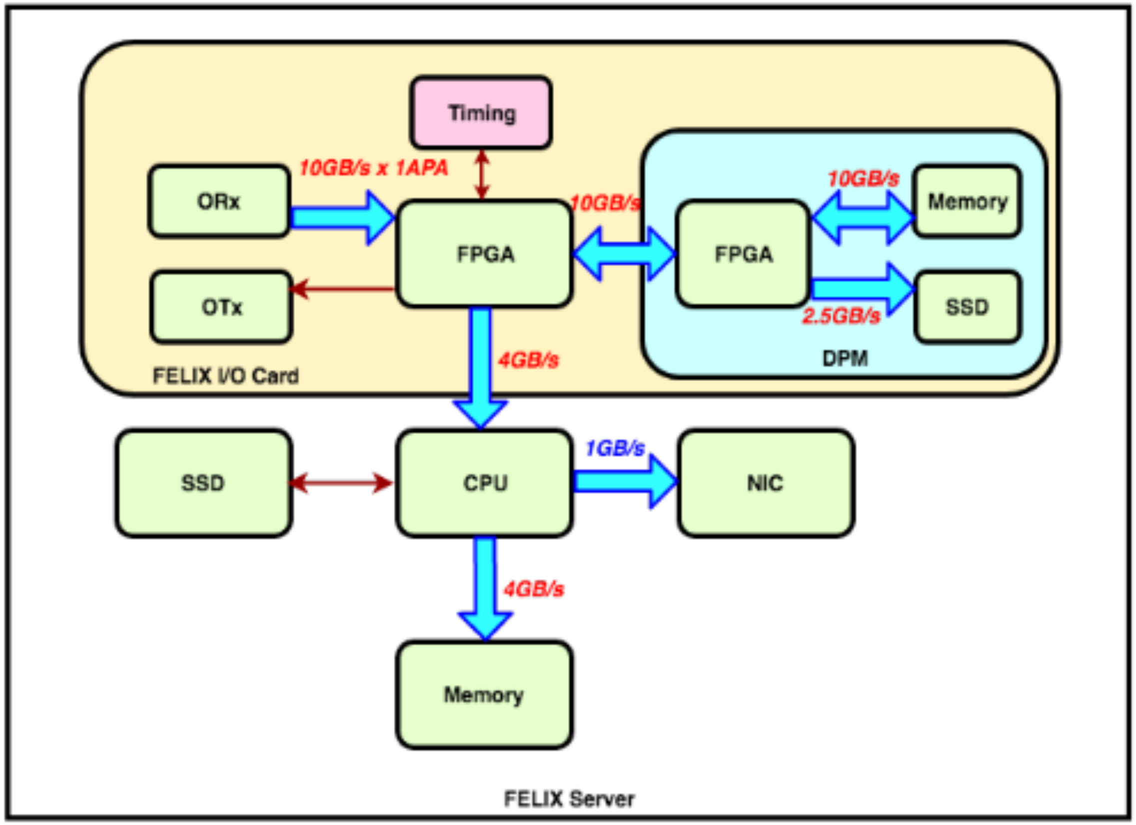

Co-processor design and FELIX integration plans

- Initial studies have shown that hit finding in FPGA straightforward
  - current algorithms consider each wire separately
  - Low FPGA resource usage
- "Low cost" data compression in FPGA straightforward
  - Encode differences between samples with fixed table.
  - Achieves almost same performance as "best possible" compression.
  - Very low FPGA resource usage
  - "Best possible" compression, e.g. Huffman compression with dynamic tables, tried at ProtoDUNE but uses significant FPGA resources.
- Working towards "full chain" tests in Lab then ProtoDUNE

Co-processor Firmware

- Time-multiplexing framework
  - Logic runs at high multiple of TDC sample rate. Reduces resource usage. Used extensively at LHC
  - Data divided up into packets and sent along communication channels between processing blocks
  - Each processing block cycles through data from 64 wires
  - For more details see [talk at DUNE Collaboration meeting, Jan 19]()

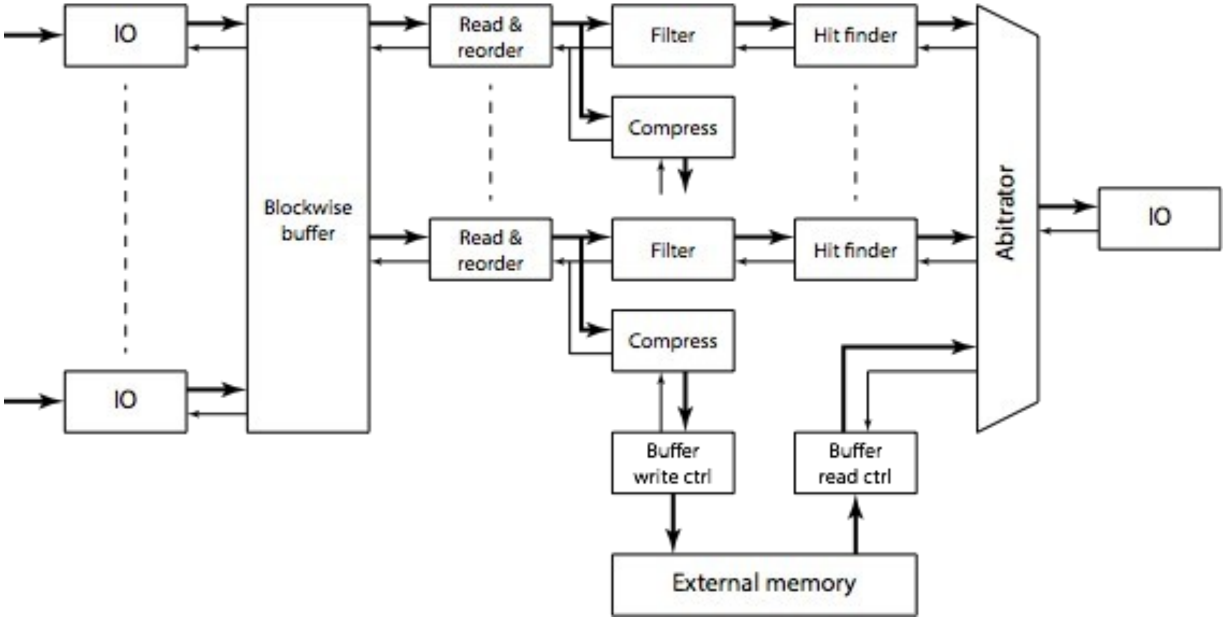

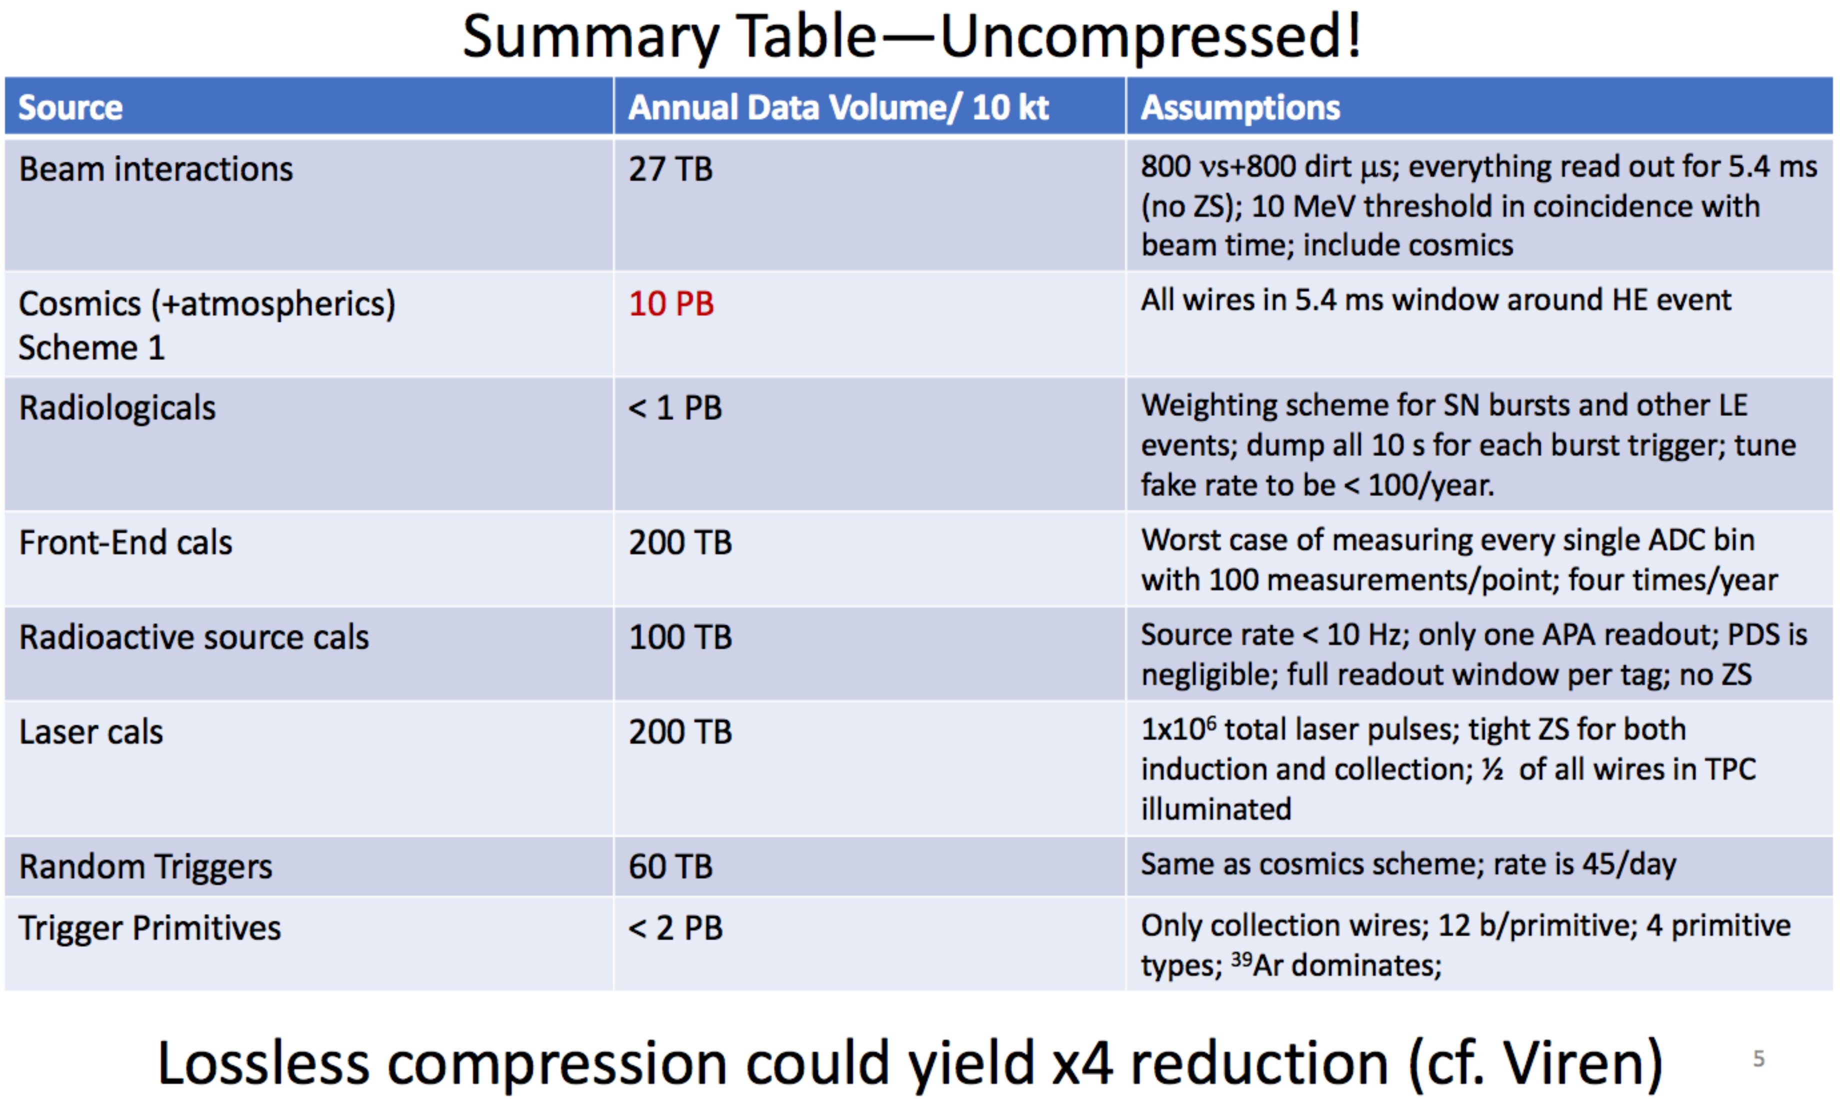

Summary Table—Uncompressed!

Source	Annual Data Volume/ 10 kt	Assumptions
Beam interactions	27 TB	800 νs+800 dirt μs; everything read out for 5.4 ms (no ZS); 10 MeV threshold in coincidence with beam time; include cosmics
Cosmics (+atmospherics) Scheme 1	10 PB	All wires in 5.4 ms window around HE event
Radiologicals	< 1 PB	Weighting scheme for SN bursts and other LE events; dump all 10 s for each burst trigger; tune fake rate to be < 100/year.
Front-End cals	200 TB	Worst case of measuring every single ADC bin with 100 measurements/point; four times/year
Radioactive source cals	100 TB	Source rate < 10 Hz; only one APA readout; PDS is negligible; full readout window per tag; no ZS
Laser cals	200 TB	1x10^6 total laser pulses; tight ZS for both induction and collection; ½ of all wires in TPC illuminated
Random Triggers	60 TB	Same as cosmics scheme; rate is 45/day
Trigger Primitives	< 2 PB	Only collection wires; 12 b/primitive; 4 primitive types; ^{39}Ar dominates;

Lossless compression could yield x4 reduction (cf. Viren)

Uncompressed Event Rates

Localized HE triggers (1 SP module):

5.4ms x 150 x 2560 x 2 MHz x 12 bit = 6.22 GB

Extended SNB triggers (1 SP module):

100s x 150 x 2560 x 2 MHz x 12 bit = 115 TB

Trigger Primitive Definitions

- Currently in TDR: Conservative TP Model:

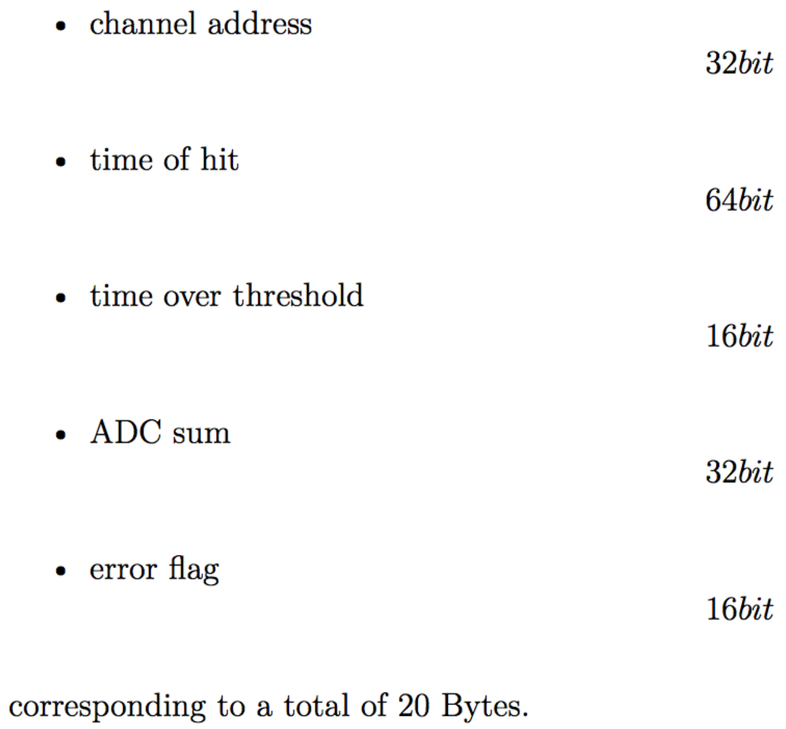

• channel address	$32bit$
• time of hit	$64bit$
• time over threshold	$16bit$
• ADC sum	$32bit$
• error flag	$16bit$

corresponding to a total of 20 Bytes.

- Expected to be dominated by Ar39 radiological backgrounds (~100Hz per channel for 0 threshold)

Substantiating Design: 2. Simulations

TPC Trigger Primitive Simulations

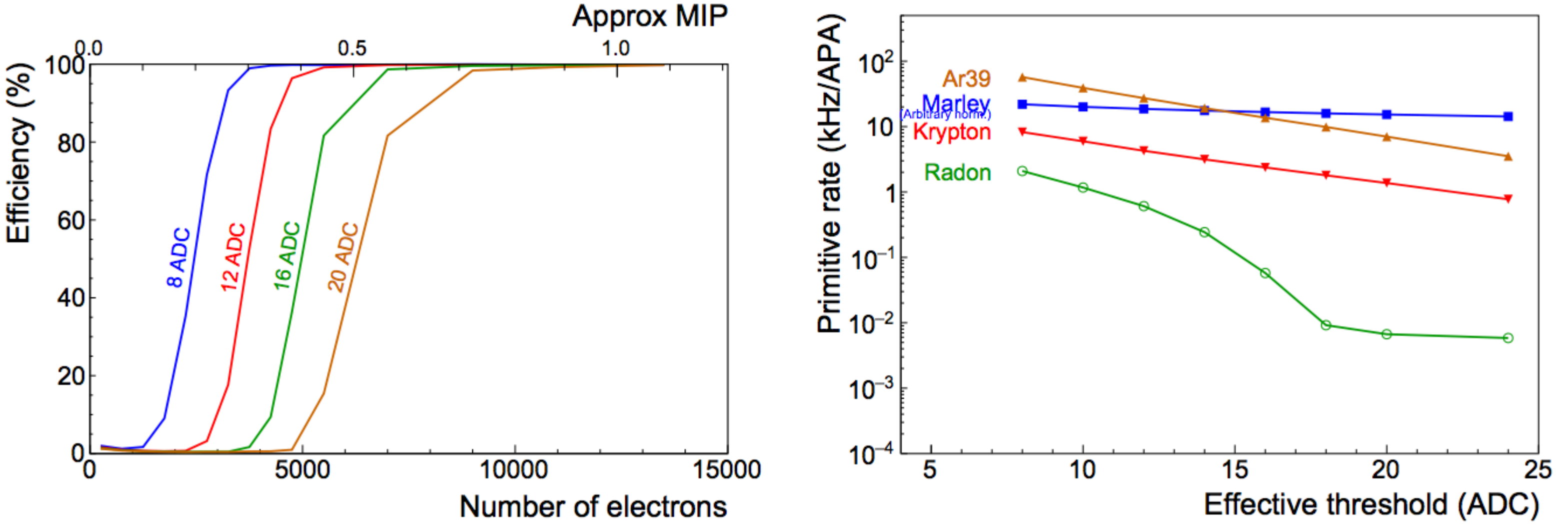

Figure 5: Efficiency and hit rate, default noise. The rate of hits due to noise (ie, not matched with any true energy deposition) is too low to show up on this plot.

Ongoing effort on trigger primitive generation:
Efficiency and TP rates due to noise and radiologicals.

Substantiating Design: 2. Simulations

TPC Trigger Primitive Simulations

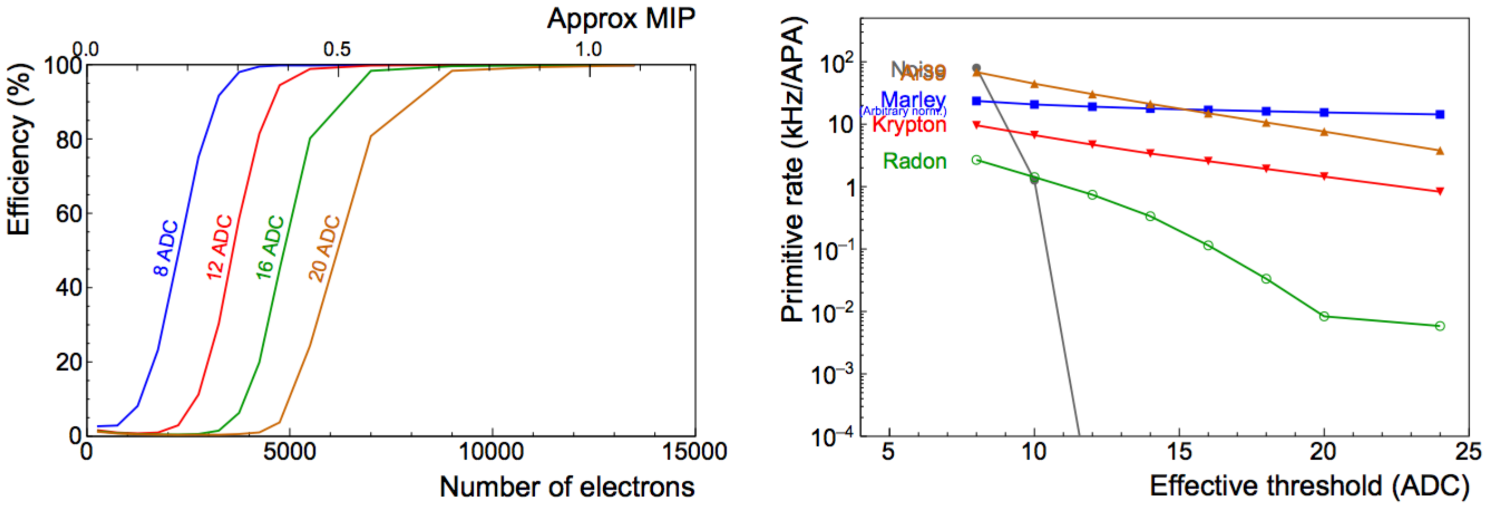

Figure 6: Efficiency and hit rate, with default noise increased by 50%.

Effect of noise increase.

Substantiating Design: 2. Simulations

TPC Trigger Primitive Simulations

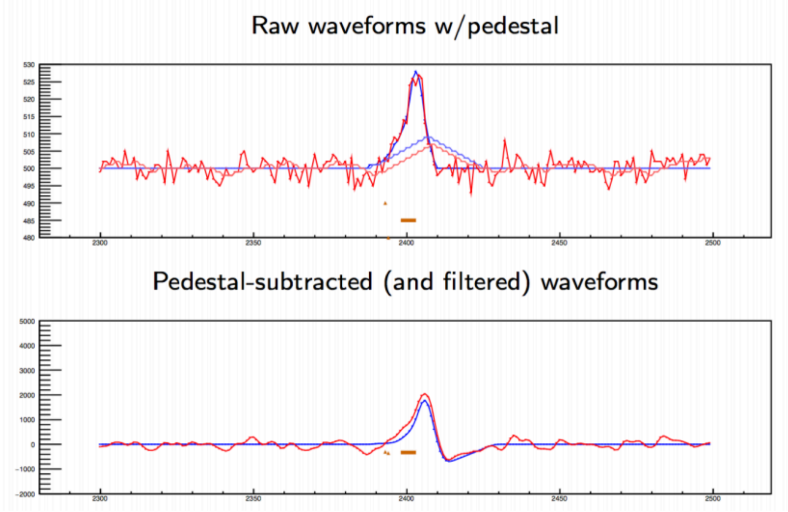

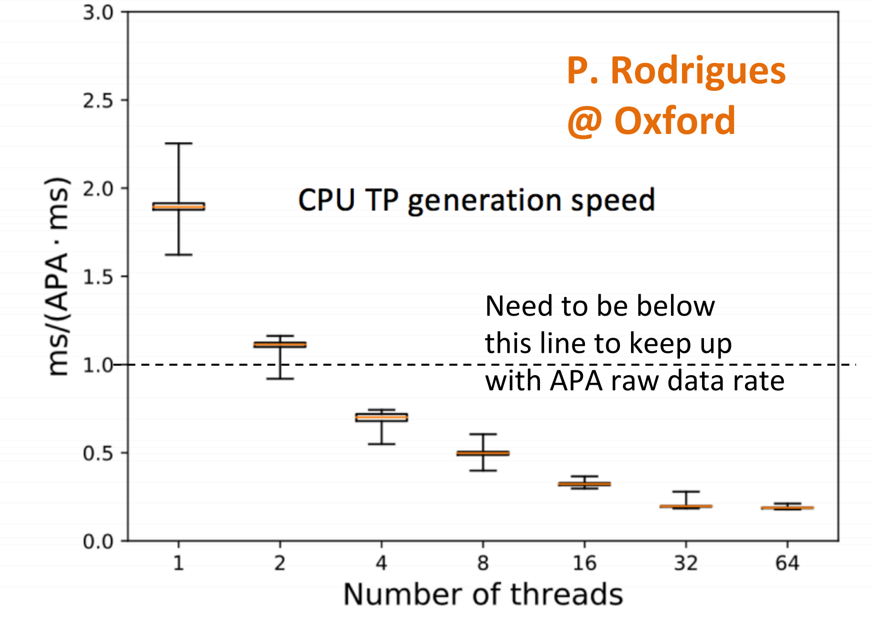

P. Rodrigues @ Oxford

CPU TP generation speed

Need to be below this line to keep up with APA raw data rate

Figure 7: Time for Trigger Primitive generation in milliseconds of processing over APA milliseconds, as a function of the number of threads. The test was done on an existing Xeon Gold 6140 system. The boxes indicate the median and interquartile range of times for various iterations of the test, and the whiskers indicate the minimum and maximum after 1000 repeats of thet test. By 4 threads (cores), the algorithm can keep up with detector data rate.

Ongoing effort on filtering (simulation based) and trigger primitive generation speed.

ProtoDUNE-SP Trigger Primitives

"Pure" 39Ar rate should be ~33kHz/ APA (half of DUNE, since only one side of APA is active).

Currently working on understanding additionally expected contribution from cosmics and (known) noisy channels in ProtoDUNE-SP.

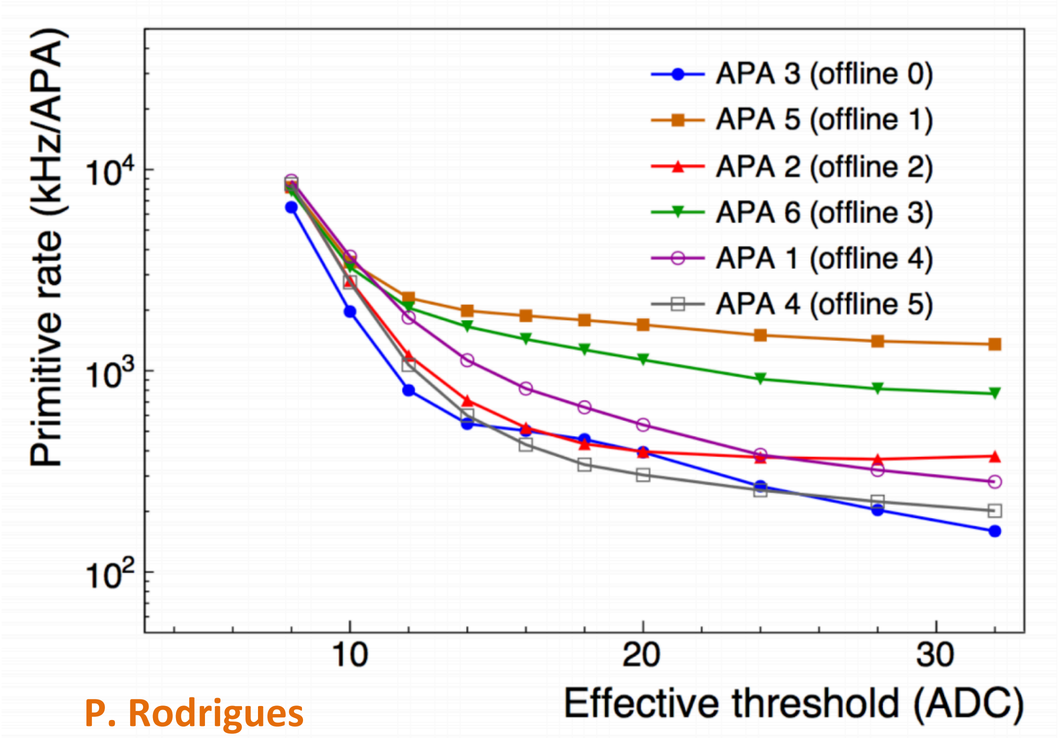

P. Rodrigues @ Oxford

Also, full stream, single APA trigger primitive generation on CPU (10 cores) demonstrated at ProtoDUNE-SP!

ProtoDUNE-SP Noise Studies

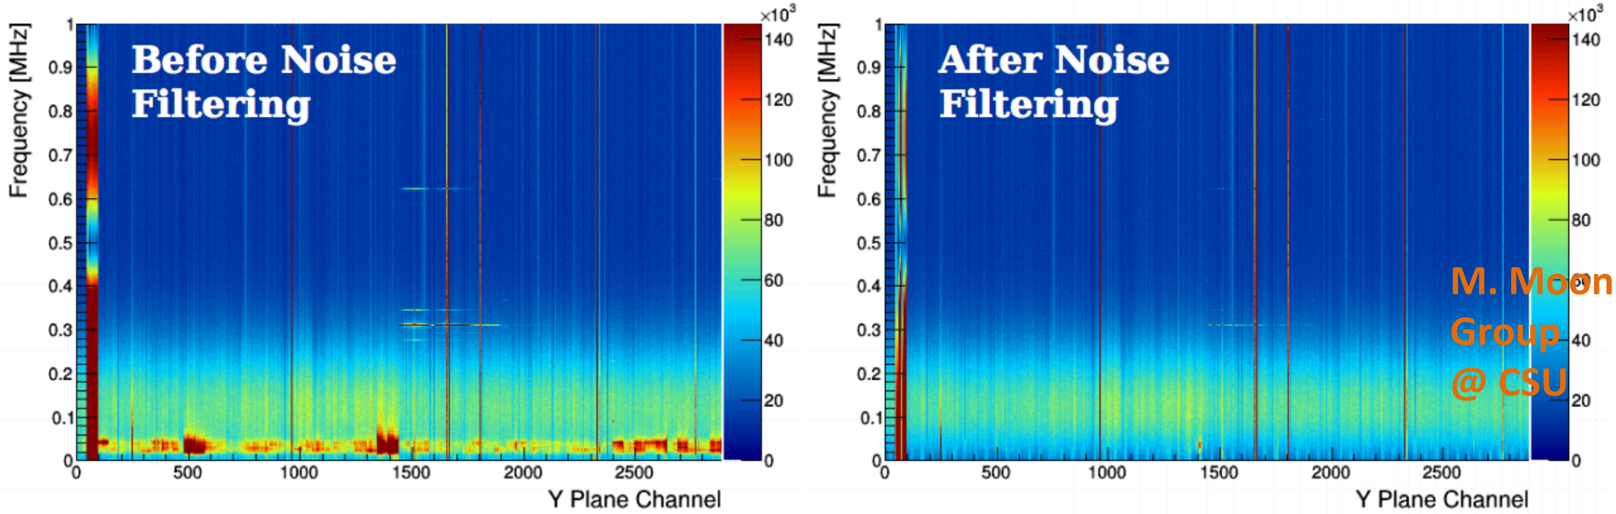

M. Mooney Group @ CSU

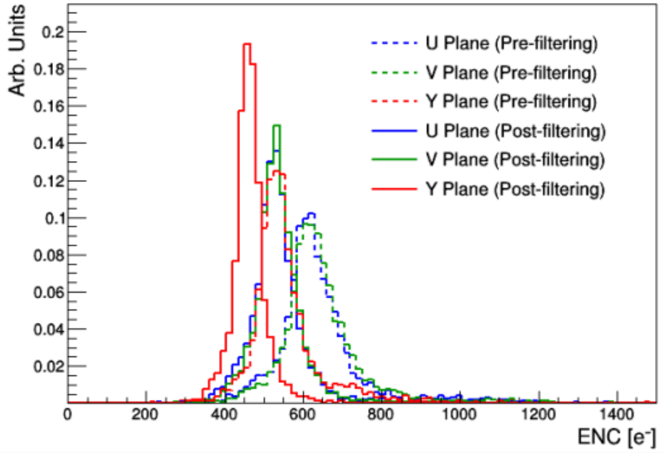

- Coherent noise @ 10-40 kHz
- Remove by subtracting median ADC value across every 16 channels, independently for each tick
- ENC: ~100 e^- effect

ProtoDUNE-SP Ar39 Rates

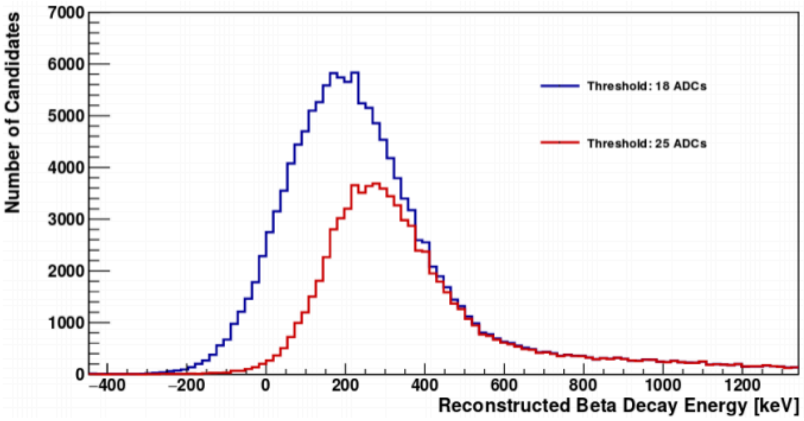

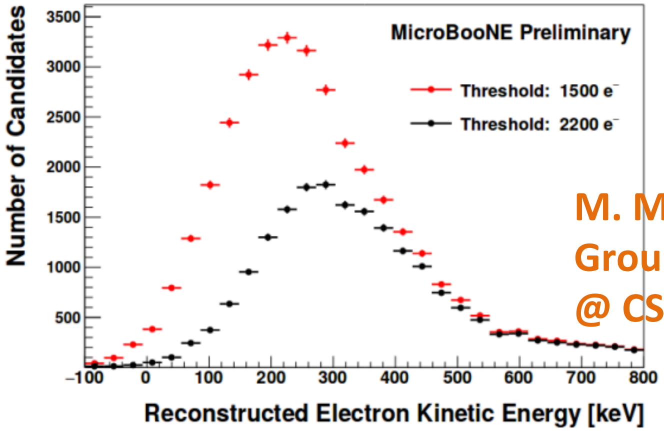

M. Mooney Group @ CSU

- After the three fixes, spectrum looking much closer to that observed at MicroBooNE
  - Still broader due to longer wires, thus higher noise
  - Lower energy reach due to less recombination at 500 V/cm (MicroBooNE runs at 273 V/cm)
- Next: fit to MC templates and extract **rate**

DUNE FD SP DAQ Schedule & Milestones

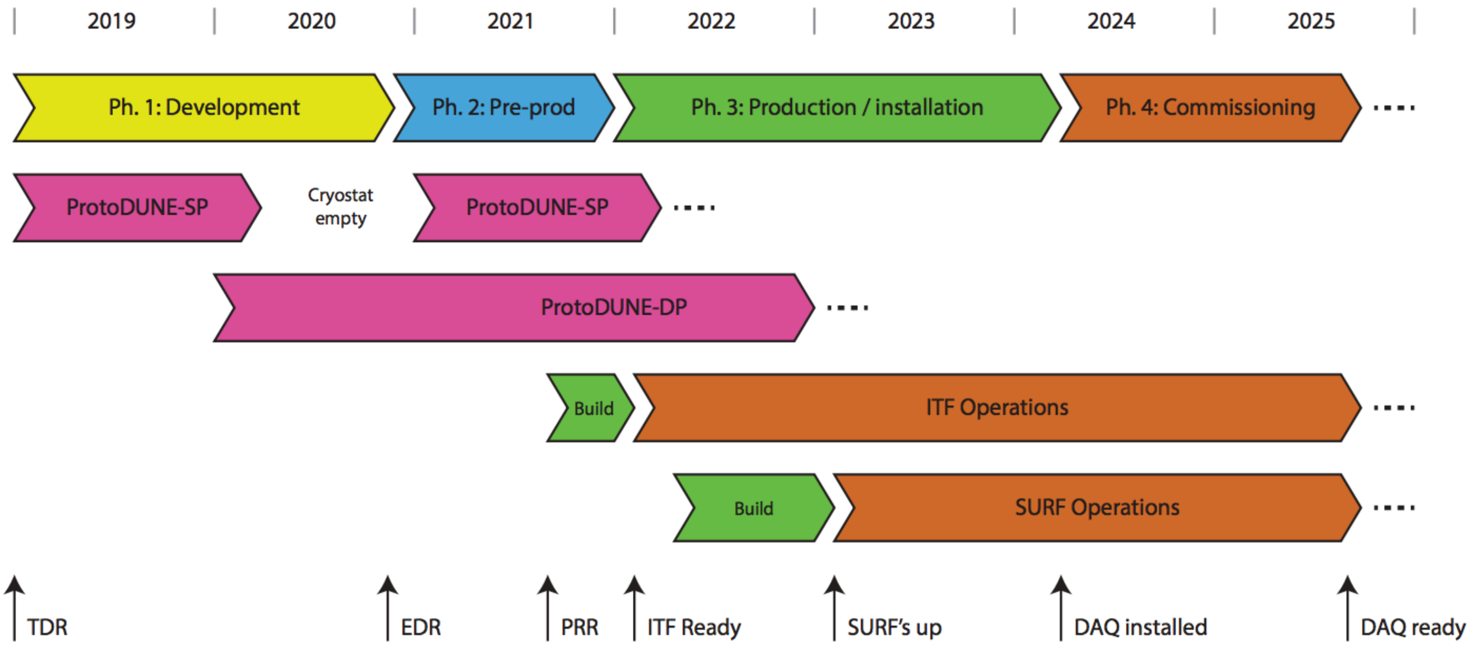

Current status:

- Done: Establish baseline design
- In Progress: Demonstrate baseline design in protodune
- Planned: Complete engineering design

Consortium Membership

Country	Institute	Primary interests	Links
CERN	CERN	FELIX software	
France	Lyon	DP interface	DPE
Japan	Iwate	DP interface	DPE
Japan	KEK	DP interface	DPE
Japan	NIT Kure	DP interface	DPE
Netherlands	NIKHEF	FELIX firmware	
UK	Birmingham	Co-processor	
UK	Bristol	Timing system; firmware	
UK	Edinburgh	Computing interface	Computing
UK	Imperial	CCM; firmware	
UK	Liverpool	PD-SP coordination	APA
UK	Oxford	Co-processor; data selection	
UK	RAL	CCM; firmware; infrastructure / installation	TC
UK	Sussex	Data selection	
UK	UCL	Firmware	
UK	Warwick	PD-SP PDS	PDS

Country	Institute	Primary interests	Links
USA	BNL	FELIX hardware	CE
USA	Columbia	Data selection	
USA	Duke	Data selection	
USA	FNAL	Back-end	Many
USA	Iowa	Data selection	
USA	Minnesota	Infrastructure / instalation	
USA	Notre Dame	Data selection	Calibration
USA	Penn	Data selection; timing system	
USA	PNNL	TBD	
USA	SDSMT	Data selection	
USA	SLAC	CE interface	CE
USA	UCD	DQM	
USA	Colorado	Data Selection	
USA	UCI	TBD	
Portugal	1 institution	Calibration interface	Calibration
Italy	1 institution	Data selection	
Canada	3 institutions	FELIX	
Czech republic	1 institution	CCM	
France	1 institution	Firmware	DPE

Consortium Organization

- Consortium Leader: *Dave Newbold (STFC)*
- Technical Leader: *Georgia Karagiorgi (Columbia)*

- TDR Editors: *Brett Viren (BNL) & Georgia Karagiorgi (Columbia)*

Working groups:

- **WG1.** Architecture
  - *Lead: Giles Barr (Oxford) & Giovanna Lehman-Miotto (CERN)*
- **WG2.** Hardware
  - *Lead: David Cussans (Bristol) & Matthew Graham (SLAC)*
- **WG3.** Data Selection
  - *Lead: Josh Klein (U. Penn.)*
- **WG4.** Back-end
  - *Lead: Kurt Biery (FNAL)*
- **WG5.** Integration/Infrastructure
  - *Lead: Alec Habig (UMN – Duluth)*