The APx Board for the CMS Phase 2 L1 Calorimeter trigger: Testing and Performance

Pallabi Das[1], Sridhara Dasu[3], Piyush Kumar[2], Bhawna Ghomber[2], Tom Gorski[3], Stephanie Kwan[1], Andrew Loeliger[1], Luis Alberto Perez Moreno[1], Alexander Savin[3], Victor Shang[3], Varun Sharma[3], Aleš Svetek[3], Jesra Tikalsky[3], **Isobel Ojalvo[1]**

1 Princeton University
2 University of Hyderabad, India
3 University of Wisconsin - Madison

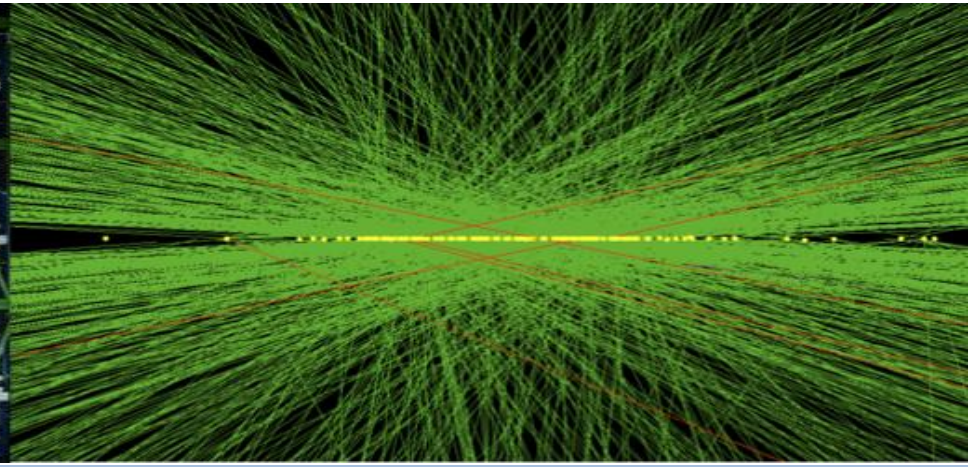

HL-LHC Upgrades

- L1 trigger upgrade:
  - More granular input information: 25 times increment for calorimeter trigger.
- First time inclusion of
  - tracker information
  - Correlator trigger
- Advanced and more complex algorithms.
- Usage of:
  - Large FPGAs: Virtex UltraScale+.
  - High-speed optical links: 25 Gbps.
- Latency: 12.5 µs
- Rate: 750 kHz
- Replacement of electronics infrastructure
  - from µTCA to ATCA standard
    - Rack
    - Crate
    - Board

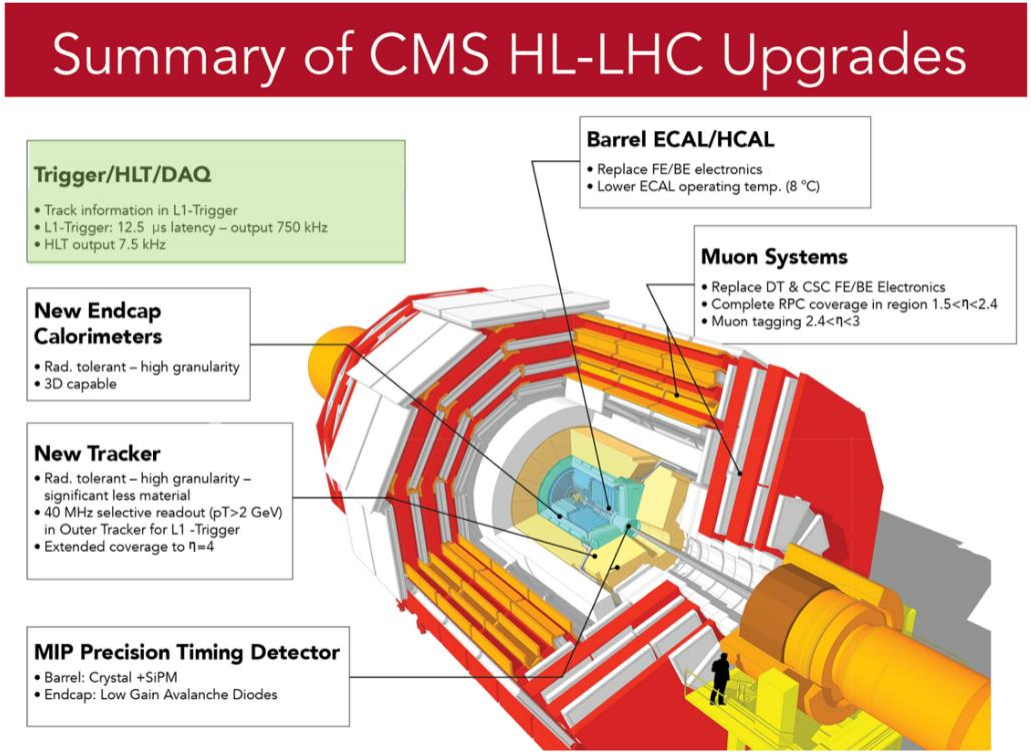

Trigger Systems

Data is selected for offline analysis with the 2-tiered **trigger system**

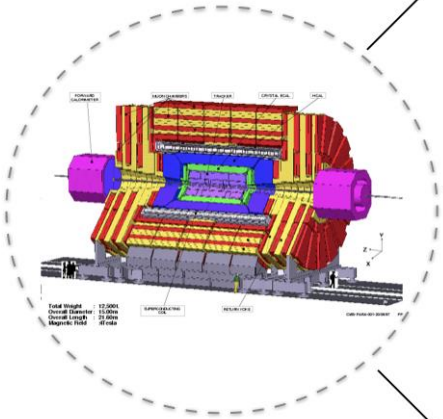

40MHz Collision Rate
~8 MB/event

Level 1 Trigger

Detector data stored temporarily in buffers for 12.5 uS

L1A 750kHz

High Level Trigger

Full Detector Information Available
-> trade off: Fast Algos

Final Decision ~7.5 kHz

Full Detector Readout for RECO and Offline Analysis

Phase 2 CaloTrigger Block Diagram

- The trigger system as a whole will have inputs from Tracker, Calorimeter, Muon Systems
  - Calorimeter trigger has inputs from the Barrel Calorimeter, Forward Hadron and HGCAL
- Latency:
  - Entire System: 12.5µs
    - Calorimeter trigger allotted 4 µs

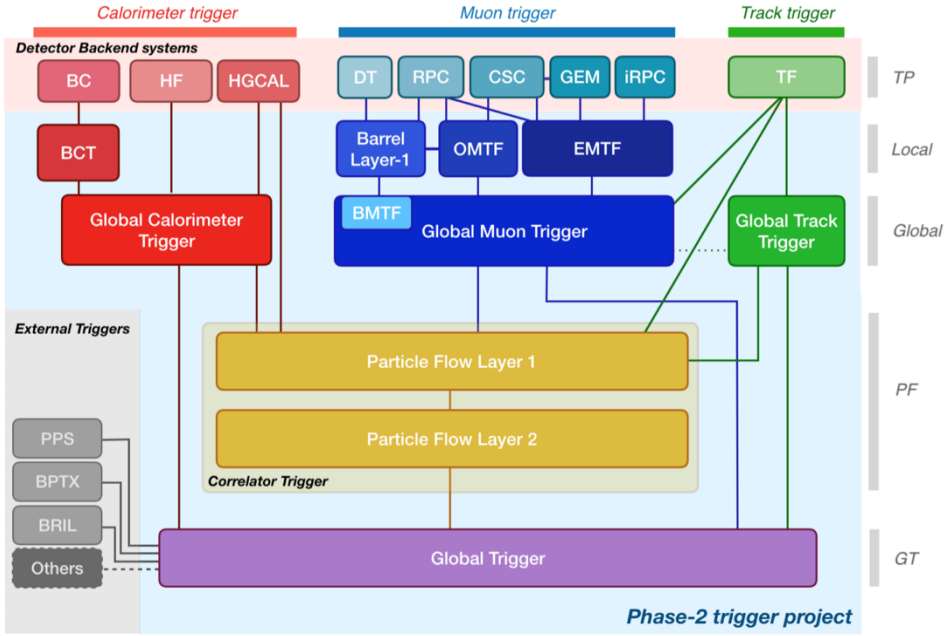

Phase-2 CaloTrigger Cluster Finding

- The figure shows each ECAL crystal as a pink square.
- The core cluster sums the energy in a 3x5 window of crystals (light blue) centered on the seed crystal (red).
- Shower shape energy deposit sums are calculated around the seed crystal in 2x5 and 5x5 windows (dashed blue and purple lines).
- Bremsstrahlung energy is detected in adjacent 3x5 windows (yellow spread in ϕ).
- Isolation is computed in a 7x7 tower window around the tower containing the seed (too large to display).

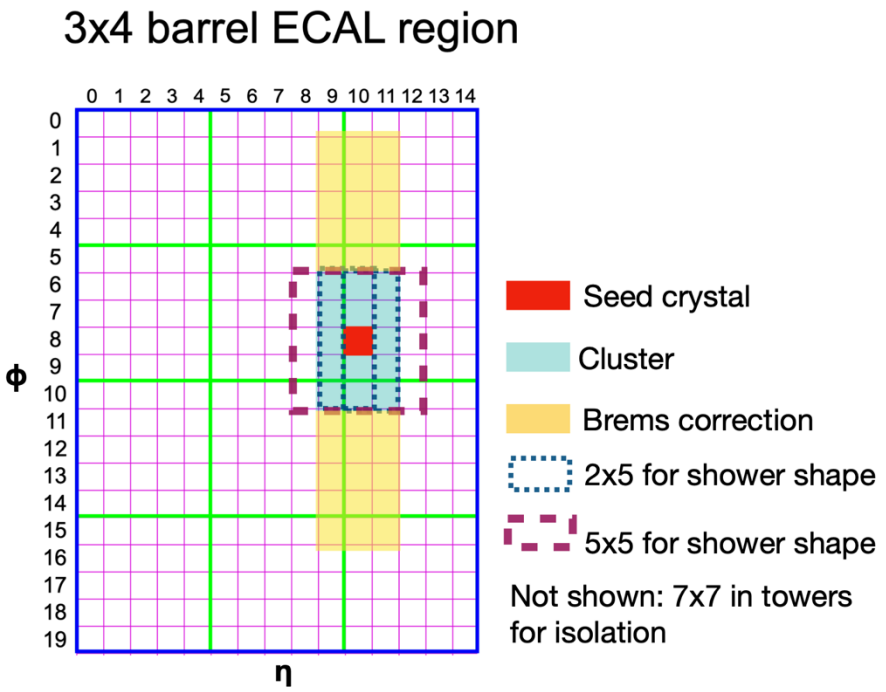

Phase-2 CaloTrigger Cluster Finding

- The figure shows each ECAL crystal as a pink square.
- The core cluster sums the energy in a 3x5 window of crystals (light blue) centered on the seed crystal (red).
- Shower shape energy deposit sums are calculated around the seed crystal in 2x5 and 5x5 windows (dashed blue and purple lines).
- Bremsstrahlung energy is detected in adjacent 3x5 windows (yellow spread in ϕ).
- Isolation is computed in a 7x7 tower window around the tower containing the seed (too large to display).

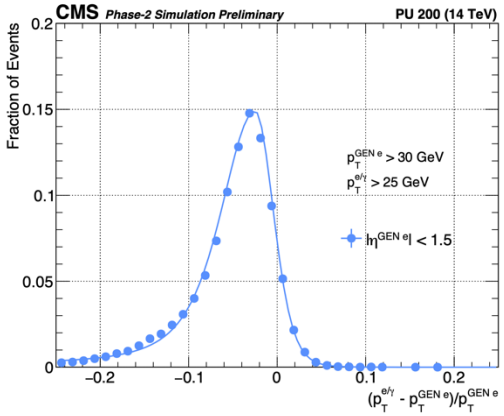

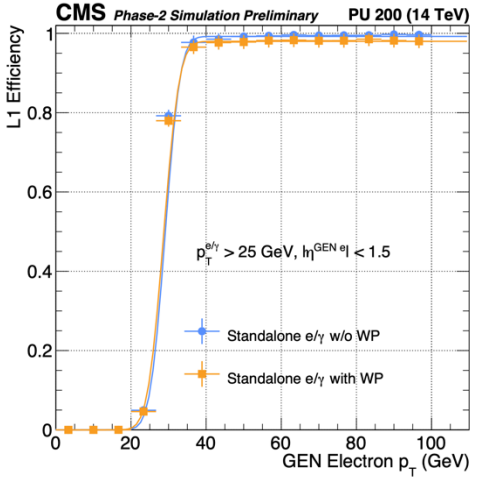

Calo Trigger Hardware: APxF Card

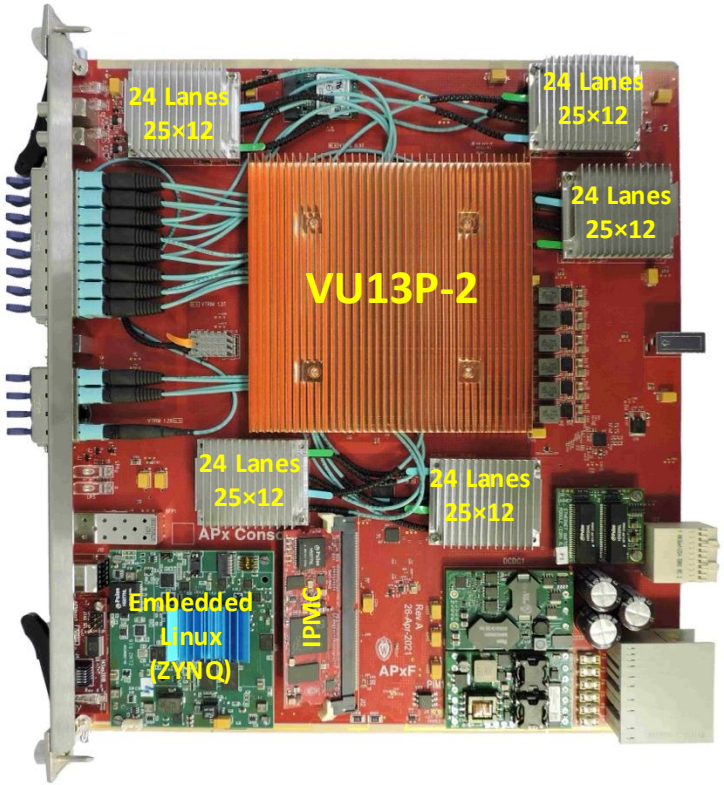

- VU13P-2 Speed FPGA
- 120 Tx/Rx Lanes of 25×12 Fireflys for trigger primitive I/O
- 4 Lanes of 28×4 Firefly for readout
- ZYNQ-7000 based IPMC
- ZYNQ MPSoC-based Embedded Linux for online control

APxF Rev K

- ELM2: Custom ZYNQ MPSoC Linux endpoint for APx ATCA Blades
  - Used on all APx Cards until APxF Rev B
- Kria K26: Generic Xilinx ZYNQ MPSoC
- Migrate ELM2 to K26
  - K26 costs less, but lacks GbE PHY, SD card and refclk synths and needs 5V supply
  - Smaller footprint than ELM opens main board space to implement needed circuitry
- Will use K26 modules in Calo Trigger production
- First units currently undergoing assembly

APx Firmware Shell

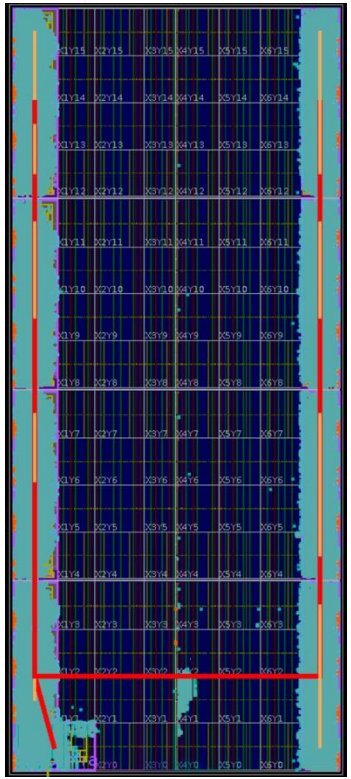

Empty Shell

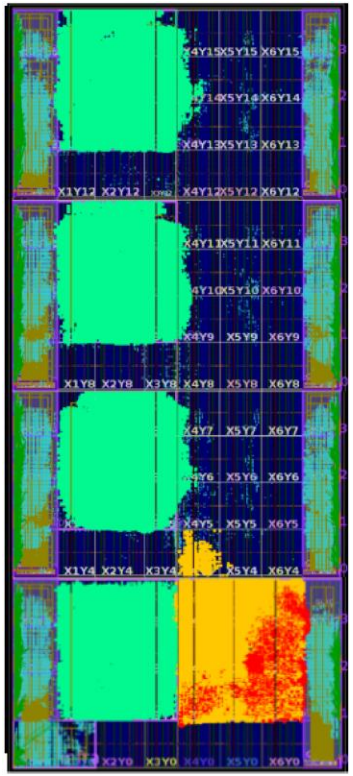

Shell + RCT algo

- Generic FPGA project
  - Built-in MGT Protocol Blocks (CSP), LHC Timing (TCDS2)
  - Slow control via AXI-mapped register file to Linux endpoint
- Shell constrained to the periphery, leaving center of FPGA for algorithm use
  - 64-bit word AXI stream interface between links and algorithm
- Each Tx/Rx channel has a 1024-word playback/capture buffer
  - Can read/write the AXI stream at the link/algo boundary
  - Advanced modes for playback/capture, including dynamic switching on playback between buffer/live data, and captures triggered by link errors.

CMS Standard Protocol (CSP)

- Trigger Data in CMS Level 1 Trigger transported using CSP
- CSP Essentials:
  - Based on 64b66b style encoding with control and data words
  - Operating modes defined for 16G and 25G line rates, both synch and async to LHC clock
  - Data words carry trigger primitive information for the Level 1 Trigger
  - Control words carry overheads:
    - Packet CRCs, Transmitter ID words, BC0, filler characters needed to sustain line rate
    - No overheads in data packets (100% physics payload)
  - Protocol reliability mechanisms to resist and recover from synchronization loss from transmission errors

APx 25G CSP Link Latency

APxF Frontpanel
NIM 50 Ω I/O
(3.3V LVCMOS)

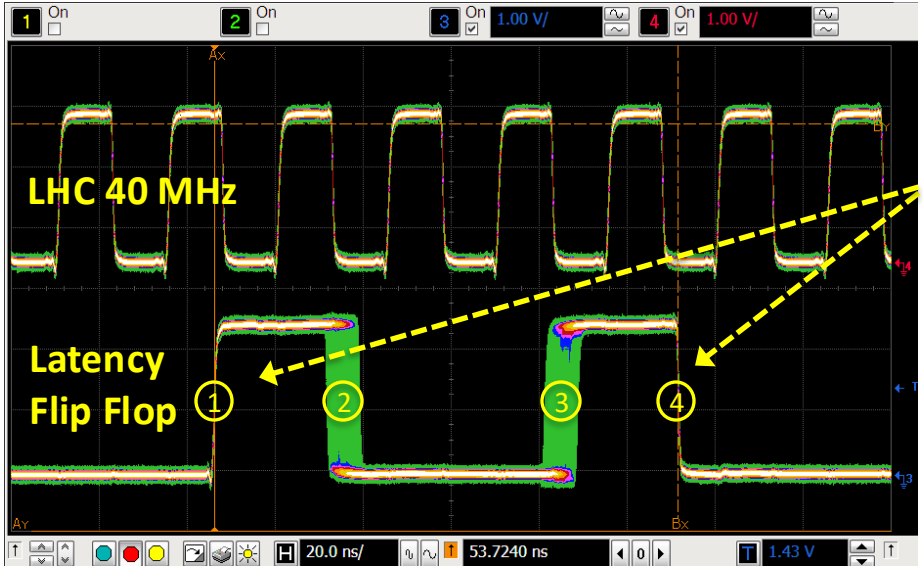

Latency Flip Flop Events for BC0-tagged data word:

1. Word at CSP Tx input port (rising edge, @algo clock)
2. Word at Tx MGT input port (falling edge, @link clock)
3. Word at Rx MGT output port (rising edge, @link clock
4. Word at Rx algo input port (falling edge, @algo clock)

- Oscilloscope measurement of APx CSP minimum link latency on APxF card
- Asynchronous Flip Flop tracks BC0-tagged word from Tx algo to Rx algo
- Measured latency with 1m of external fiber is **~4.3 Bx** (plus extra Bx per 5m of fiber)

Algorithm De-skewing With 25G Links

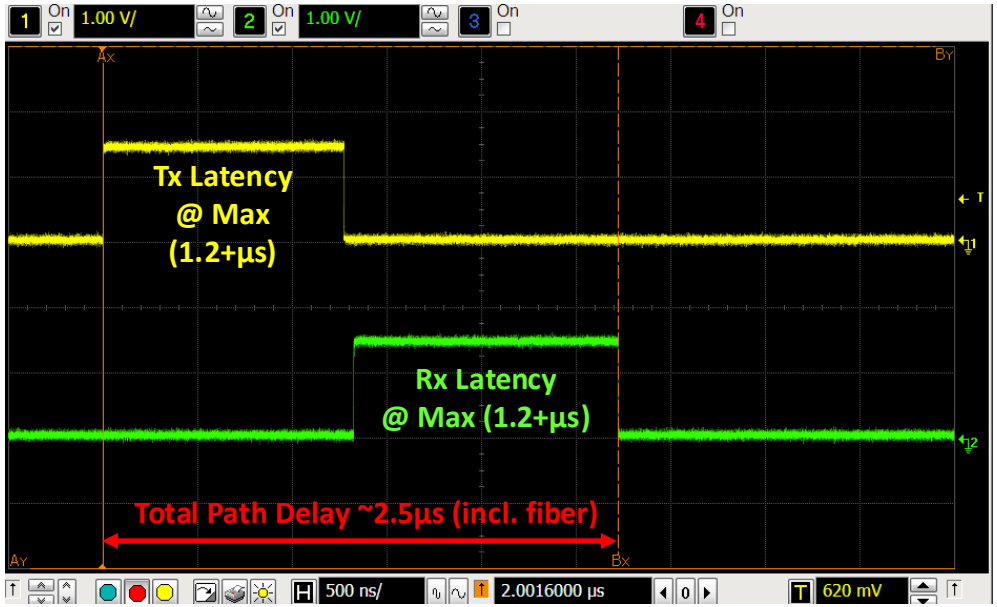

- Endcap data arrives at Global Calo Trigger much later than Barrel data
- APx CSP logic uses 512-word BRAM FIFOs to cross algo/link clock domains
- These FIFOs can also provide extra latency on data algo paths
- Tx side: bulk delay, build time options of 0, ~0.5, 1.0, 1.25 µs at 25G
- Rx side: algo-aligned delay, up to ~1.25 µs, controlled at run time at 25G
- Up to **~2.5 µs** of de-skew per Tx/Rx pair, fully channel-specific at 25G
- Eliminates need for internal algo-based de-skewing (e.g. UltraRAM)

ATCA Crate Infrastructure for Production

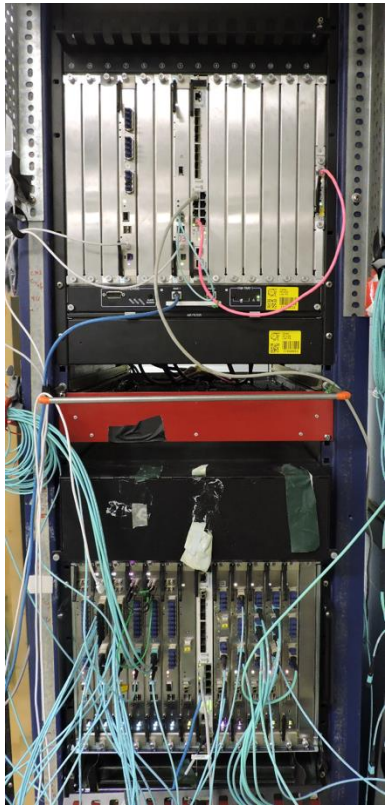

- Two ATCA crates in Wisconsin lab
- Common ersatz LHC timebase established between them in hub slot 1
  - Optical connection between two APx Test Hub (ATH) Cards, with possible future upgrade to CMS DTH cards
  - LHC 40 and BC0 to 24 blade slots
- Sufficient slots to allow newly-assembled cards to undergo extensive performance qualification of optical links (~2E15 bits/link/day @25G)
- Lab capacity for additional crates as needed

25×12 Firefly Evaluation

- The story of Samtec 25×12 parts is well-documented elsewhere
- Our approach has been *to get as many 25×12 links in operation as possible running actual protocols on actual boards and to count errors*
- Past Test Highlights:
  - Dec 2023: Success running 192 lanes of 2023-vintage Tx devices with 2022 Rx devices for >31 hours without any CRC errors
  - Jan 2024: During rigorous testing a VCSEL fails and Samtec subsequently acts to change VCSEL devices
  - May 2024: 24 lanes of new VCSEL devices are received from Samtec and rigorously tested to 3E16 bits/lane with no CRC errors observed
  - June 2024: 1E17 bits/channel reached on new VCSEL eval devices with no CRC errors
- From evaluation to date the 25×12 devices appear to have adequate performance (BER <1E-15) with availability the primary issue

All 28x4 Firefly Approach – APx4

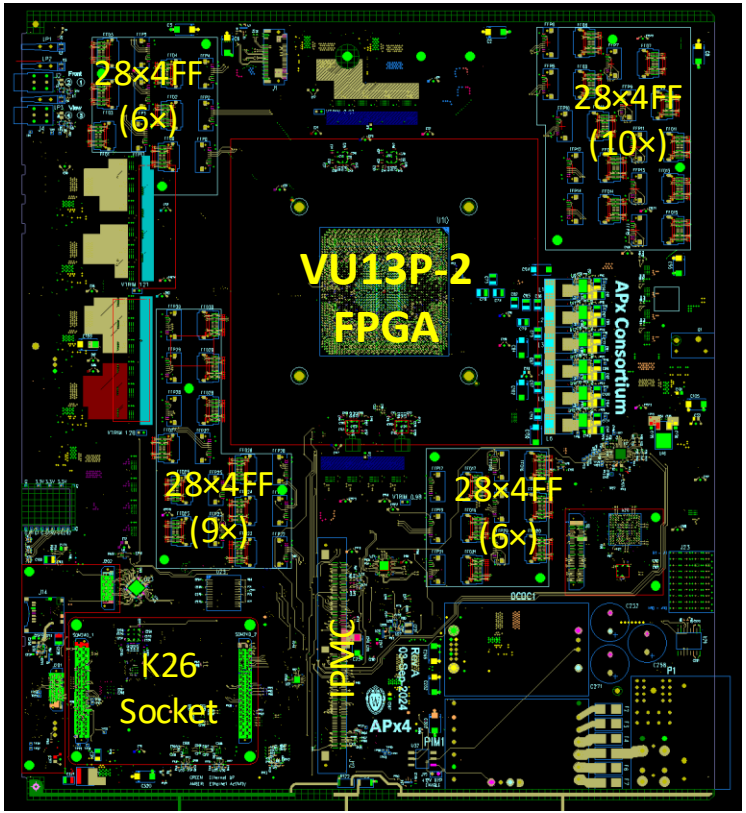

- Same VU13P-2 FPGA as APxF
- Replace 25×12 Firefly devices with 28x4 devices—total of 31 in all
- Same Linux (K26) and IPMC circuitry
- Redesigned frontpanel, Firefly power distribution and Firefly heat sinks
- PCB layout complete
  - Waiting on evaluation of K26 circuit in APxF Rev K before building boards

Firmware Validation

- Calorimeter Trigger consists of 24 regional cards and 10 global cards
- The physical footprint of the subsystem—especially the optical fiber plant—is too large to fully construct outside of the experimental site
- Through the use of virtualization of cards and link connections, full scale testing of all bitfiles and all link paths is possible

FPGA Environment for Algorithm Slice Tests

- FEAST – Environment for algorithm development and validation
- Supports mixed physical and virtual link test systems including link alignment
  - Virtual links implemented using APx shell playback/capture buffer technology
- Bitfiles run on target FPGAs with target MGT channel allocations
- Construct full-scale, multi-layer, virtualized installations and test on actual FPGAs anywhere using FEAST engine

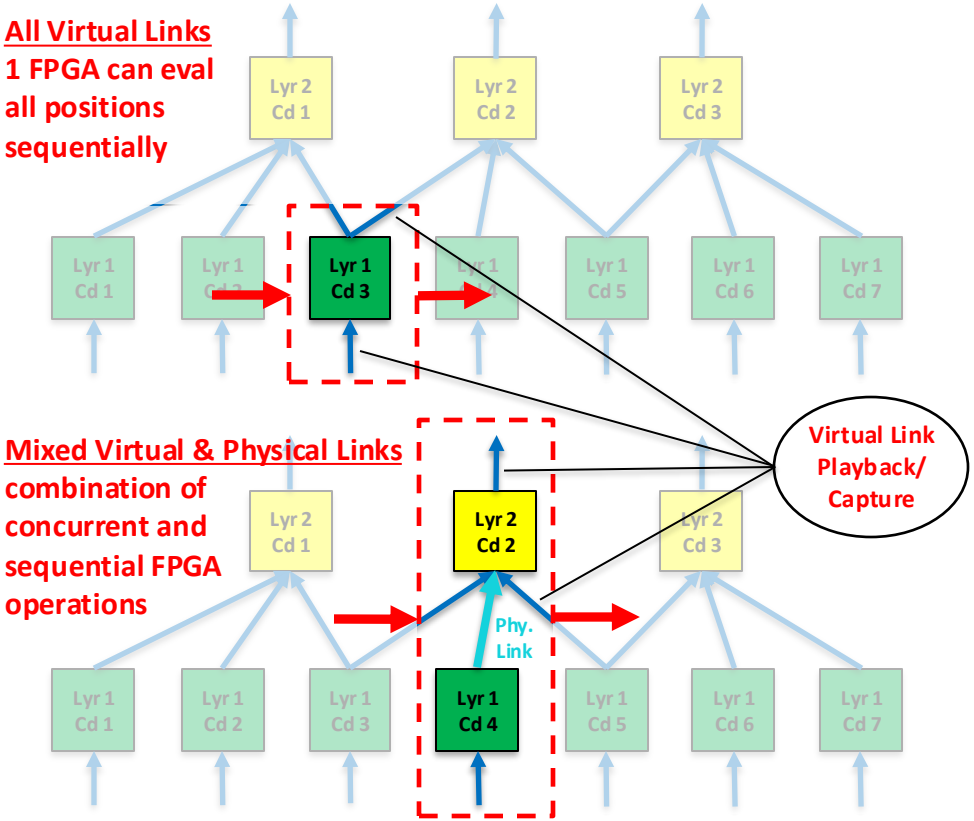

FEAST Example Job Diagram

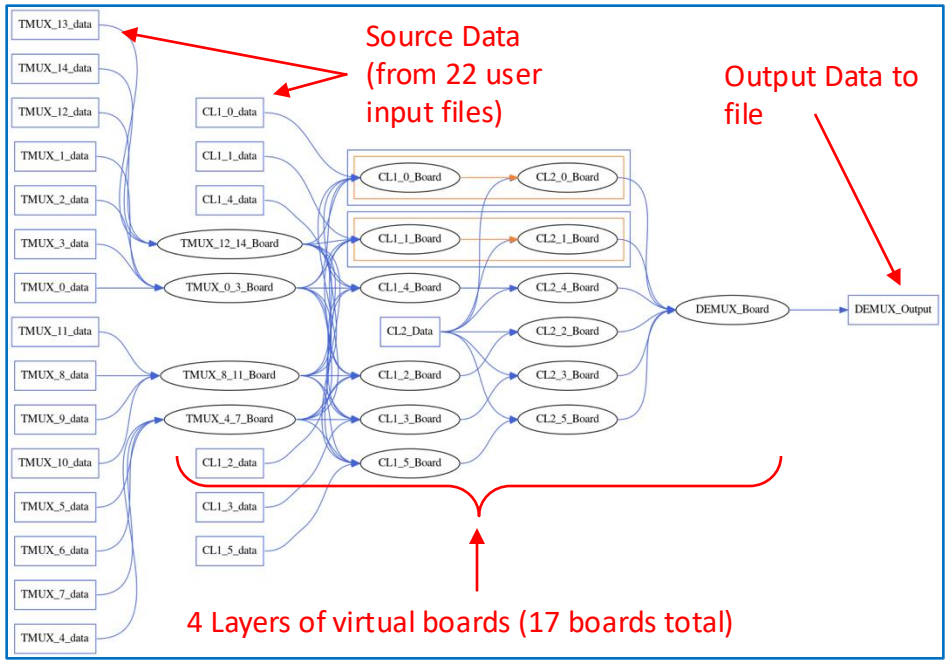

(Diagram automatically generated by FEAST sequencer)

- Auto-generated image from analysis of the example job
- Rectangles represent input/output data files
- Ovals represent virtual boards
- Labels are user-supplied
- Lines represent links (data paths between boards/files)
  - **Blue** → for virtual links
  - **Orange** → for physical links
- Visual confirmation of system configuration

Calorimeter Trigger Prototype

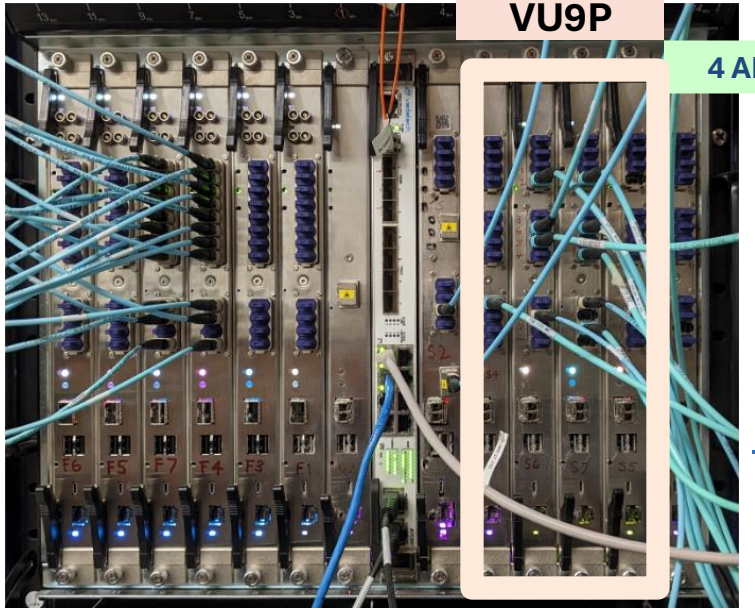

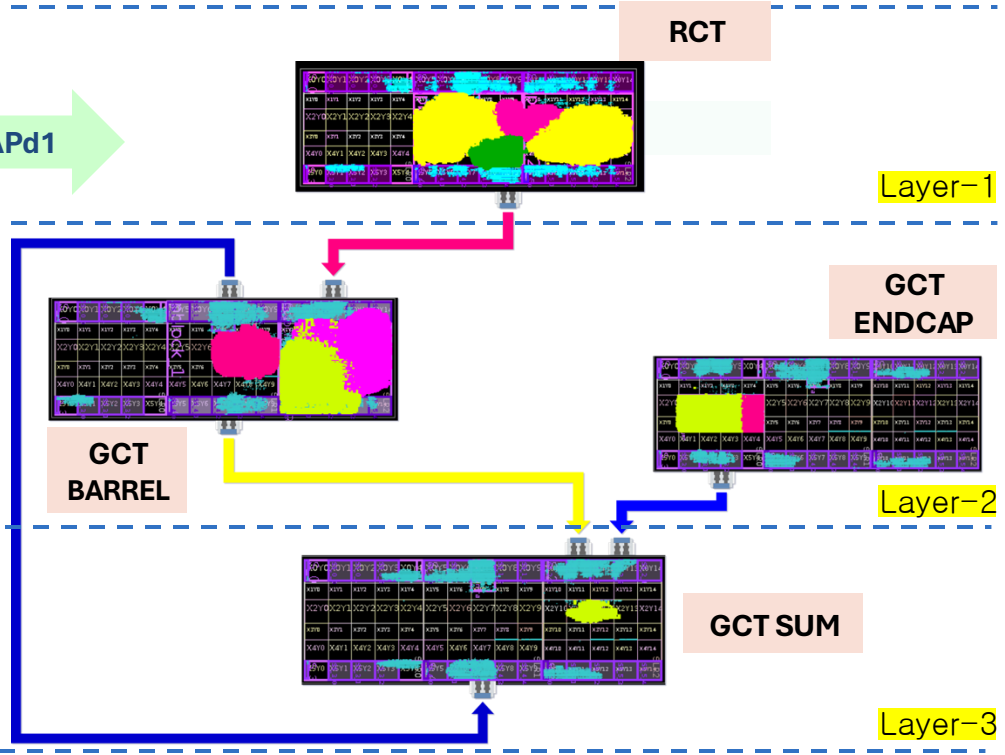

Global Calorimeter Trigger: Firmware

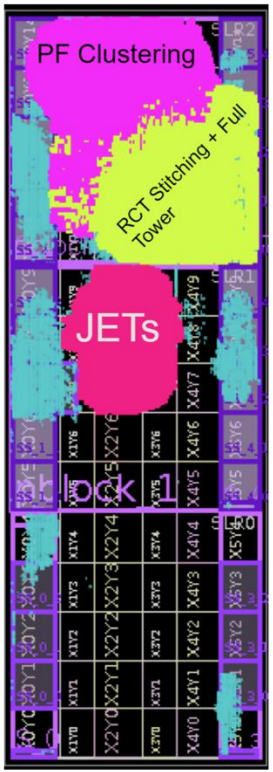

- VU9P has multiple SLRs
  - Algorithms spread across multiple SLRs
    - Crossing SLRs is not a trivial exercise
  - SLR2:
    - Stitch e/gamma clusters across RCT card boundaries, create GCT towers
    - PF Clustering: 3x3 clustering of GCT towers, sent to correlator trigger for further processing
  - SLR1:
    - Create GCT jets in barrel

Summary

- APxF card successfully produced
  - Now being used for core FW development
  - APxF Rev K undergoing assembly
- Experience gained during design, installation, commissioning and operation of the current Level-1 Calorimeter Trigger System
  - Firmware Shell separates physics algorithms from core firmware
  - Algorithms built with HLS
  - Firmware Validation using “FPGA Environment for Algorithm Slice Tests”
    - Construct full-scale, multi-layer, virtualized installations and test on actual FPGAs anywhere using “FEAST” engine

Global Calorimeter Trigger: Firmware

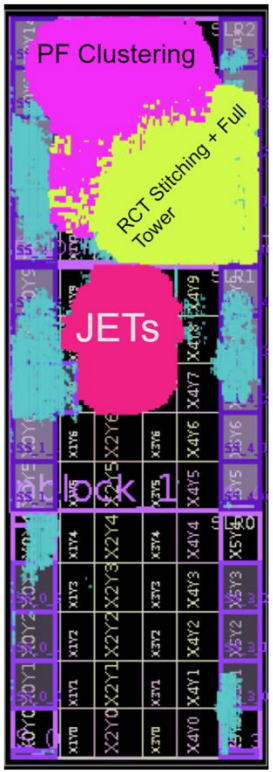

- VU9P has multiple SLRs
  - Algorithms spread across multiple SLRs
  - SLR2:
    - Stitch e/gamma clusters across RCT card boundaries, create GCT towers
    - PF Clustering: 3x3 clustering of GCT towers, sent to correlator trigger for further processing
  - SLR1:
    - Create GCT jets in barrel
  - Output:
    - 48 links to Correlator Trigger
    - 6 links to GCT sum board
  - Latency: 692 ns out of 3 µs budget
  - Resource Utilization
    - Look-up Tables (LUTs): 22%
    - Flip-flops (FFs): 18%
    - Digital Signal Processor (DSP): 0%