




CMS Barrel Calorimeter Read-out and Trigger Primitive Generation

ACES 2020 - May 27, 2020

Stephen Goadhouse, University of Virginia
Nikitas Loukas, University of Notre Dame



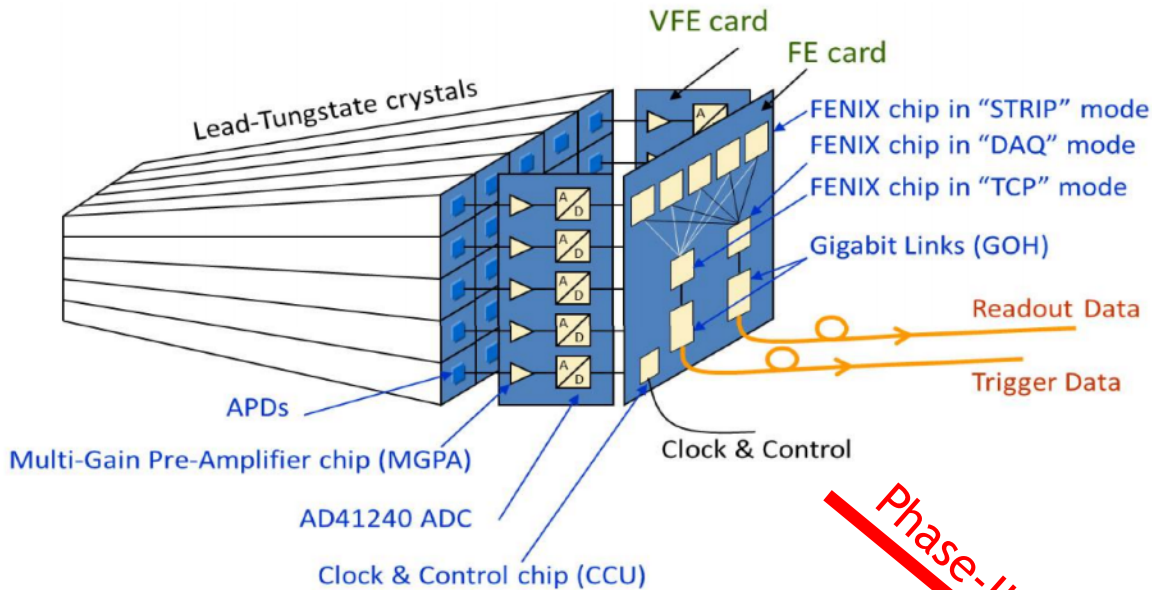


The APx Consortium

- Pooling of efforts in ATCA Processor hardware, firmware and software development
- Multiple ATCA processors and mezzanine board types
- Modular design philosophy, emphasis on platform solutions with flexibility and expandability
- Reusable circuit, firmware and software elements
- APx application areas in CMS Phase 2 Upgrade: Barrel Calorimeter, Muons, Trigger




ECAL Barrel after phase-II upgrade – Front-End



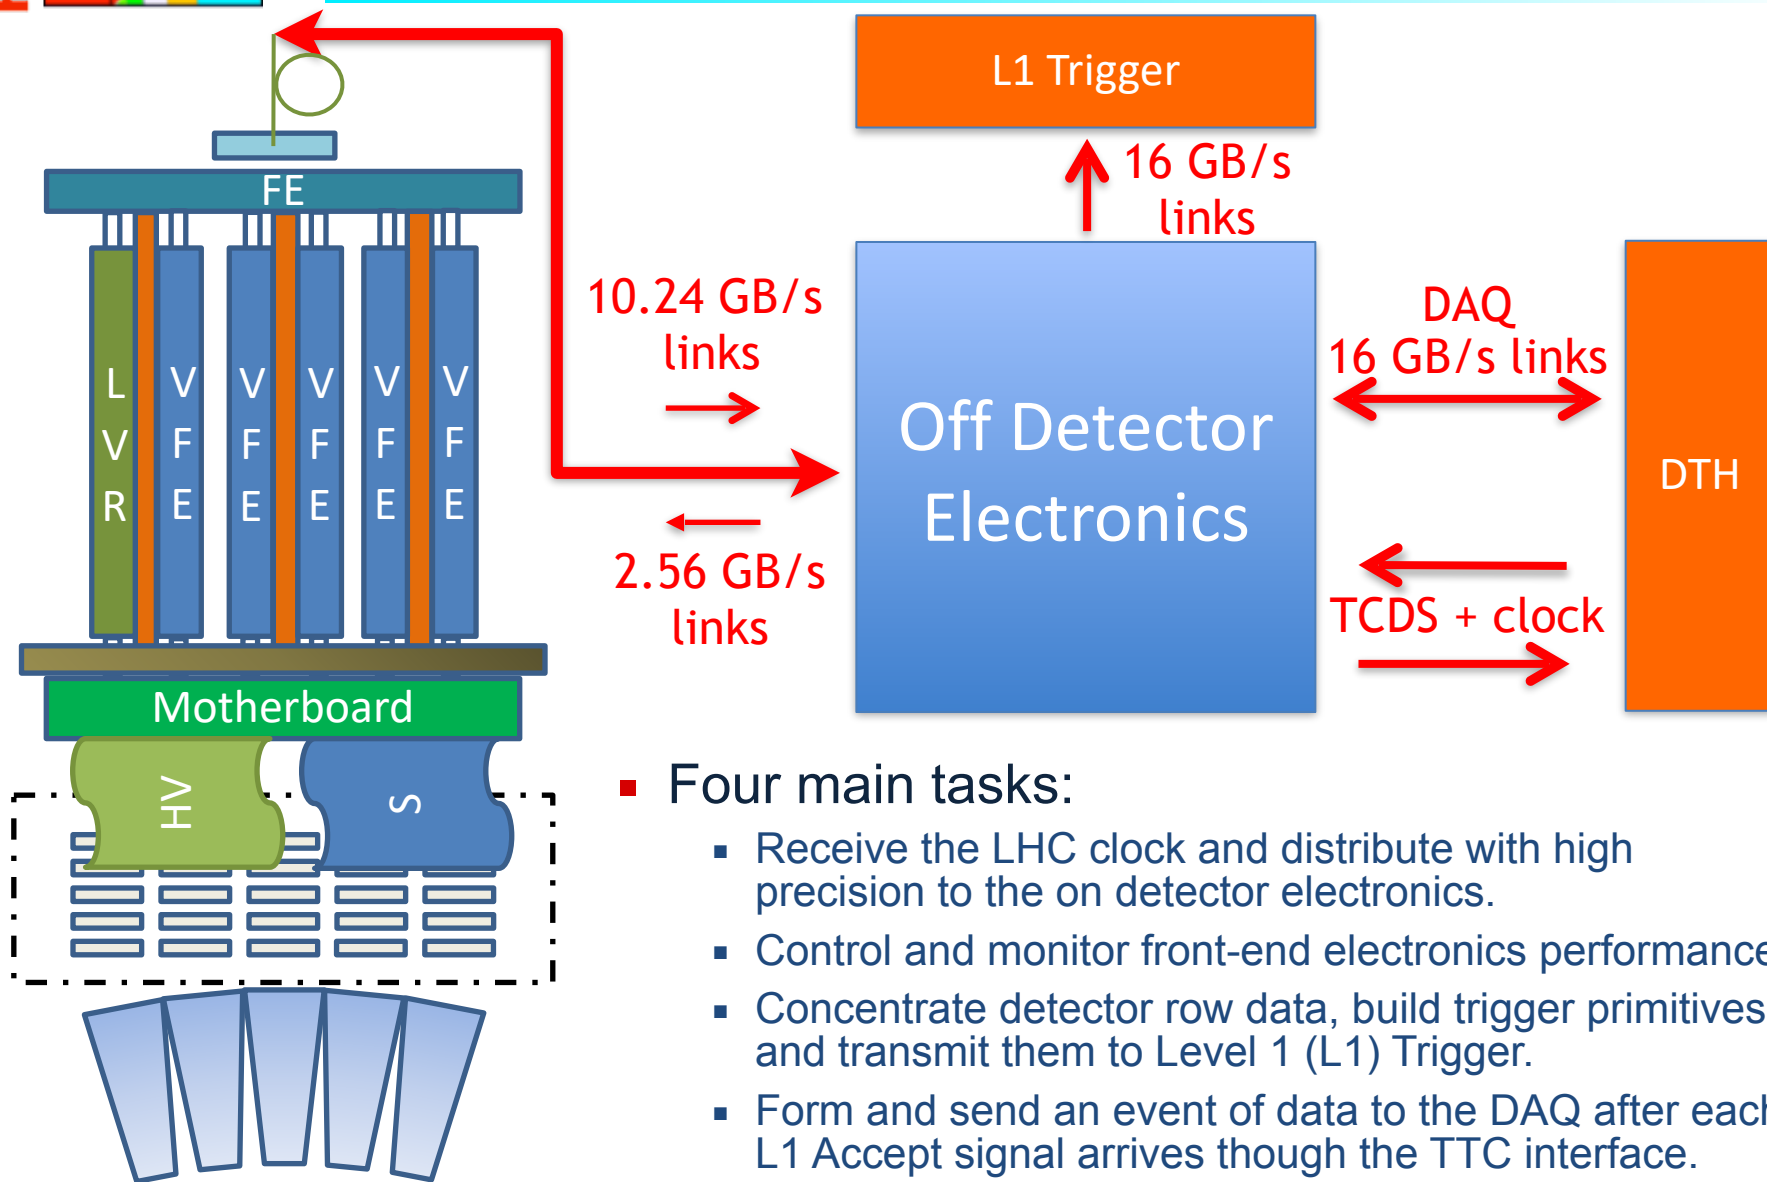
- With legacy Electromagnetic Calorimeter (ECAL), the Trigger Primitives (TPs) are built on detector, or Front-End (FE), by the FENIX chips.
- The FENIX chips combine data from 5 crystals into a strip, analyzes the digitized signal and performed Trigger and Readout functions.
- 0.8 Gb/s links are sending all data to the VME Back-End.
- The data sampling rate is 40 MHz.

Phase-II



- After phase-II, ECAL will calculate the TPs off detector in the Back-End (BE).
- FE will have IpGBT chips that concentrate row ADC data and send them to BE via Versatile+ links.
- The “master” IpGBT in FE will distribute the high precision clock and will control the VFEs via I2C buses.
- The sampling rate after the upgrade will be 160 MHz which is required for spike rejection.
- Link utilization per tower: Upstream 4 links at 10.24Gb/s with FEC5 and downstream 1 at 2.56Gb/s.

ECAL Barrel after phase-II upgrade – Back-End

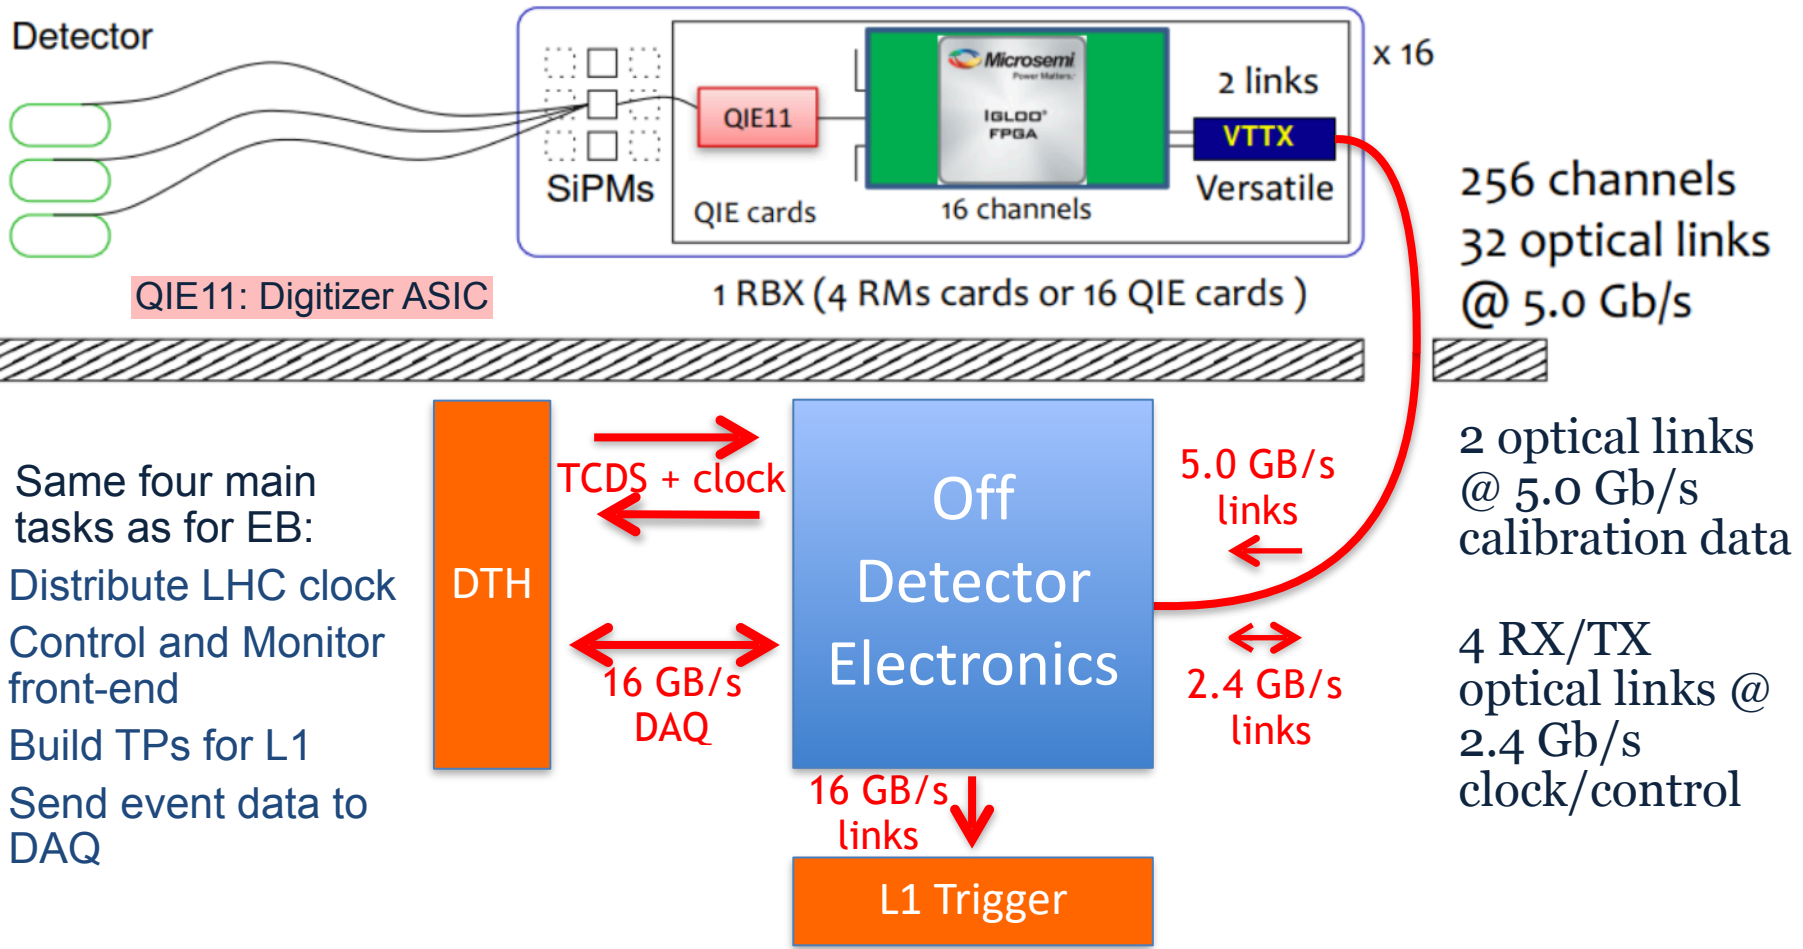


Four main tasks:

- Receive the LHC clock and distribute with high precision to the on detector electronics.
- Control and monitor front-end electronics performance
- Concentrate detector row data, build trigger primitives and transmit them to Level 1 (L1) Trigger.
- Form and send an event of data to the DAQ after each L1 Accept signal arrives through the TTC interface.

HCAL Barrel after phase-II upgrade

- Installed Barrel Hadron Calorimeter (HB) Phase-1 upgrade of FE will remain through Phase-2



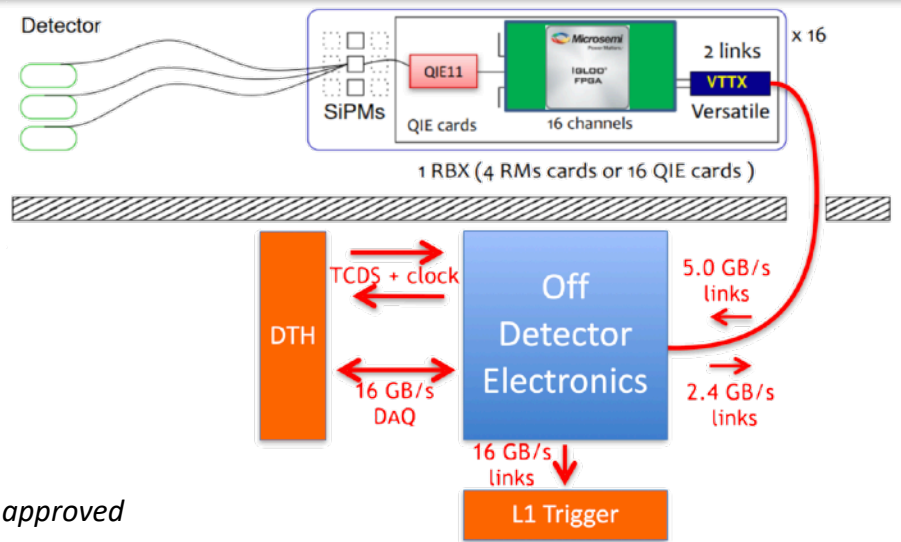

- HCAL Barrel will move to a common back-end platform with ECAL

Barrel Calorimeter Processor (BCP)

- The off detector electronics of both ECAL Barrel (EB) and HCAL Barrel (HB), will be based on the Barrel Calorimeter Processor (BCP) ATCA board.
- Baseline BCP
 - 128 Multi Gigabit Transceivers (MGT) total
 - Up to 64 TX and 96 RX MGTs for FE
 - Up to 8 TX/RX MGTs for DAQ
 - Up to 32 TX MGT for Trigger
 - Up to 16 TX/RX MGTs for BCP-to-BCP
 - Up to 8 TX/RX dedicated to BCP operations
- Both EB and HB will use the same protocols for Level-1 Trigger, TCDS and the cDAQ interfaces (though DTH)
- Both EB and HB will need DSP processing for Trigger Primitive Generation
- Additionally, EB needs processing for
 - Data decompression
 - Anomalous signal veto (DSP)

| Total BCPs in System | # BCPs | # ATCA crates |
|--------------------------------|--------|---------------|
| ECAL Barrel (EB) | 108 | 12 |
| HCAL Barrel (HB) | 18 | 2 |
| HCAL Forward (HF) [†] | 9 | 1 |
| HCAL Outer (HO) [†] | 9 | 1 |

[†] - if approved



FPGA Resource Requirements



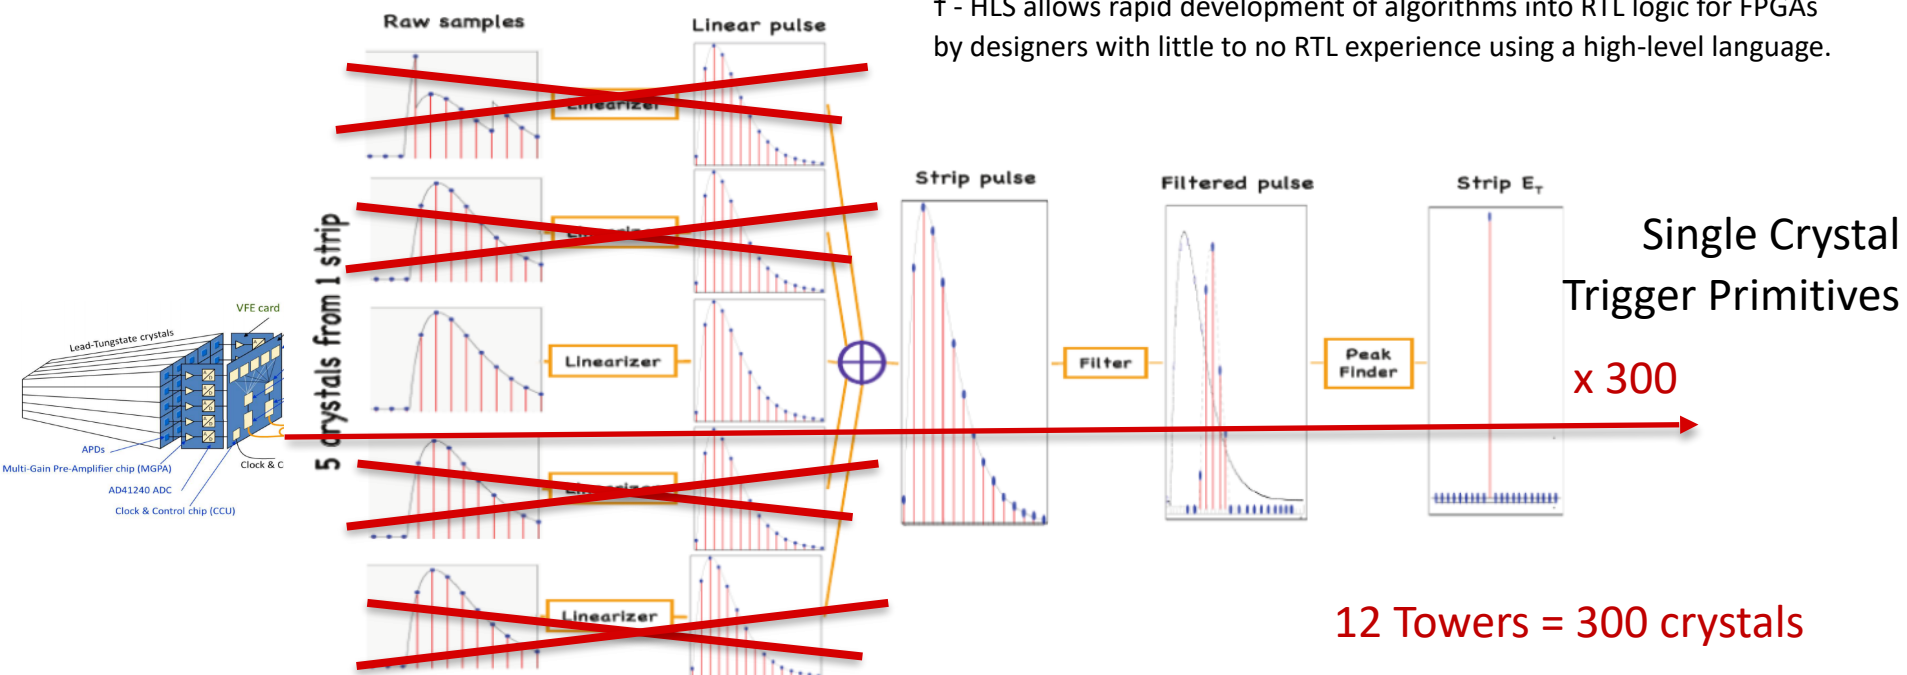
ECAL Barrel Trigger Primitive Generation

- Each crystal provides data for a Trigger Primitive
- Basic needs
 - 1) Reconstruct signal amplitude and BX identification
 - 2) Anomalous signal, or “spike”, identification
 - 3) Time measurement
- Algorithms necessary to extract these information will run on the BCP
- Use High Level Synthesis (HLS) to implement version 0 and measure preliminary values of latency and resources needed
- In the following slides we report about 1) and 2)

Amplitude reconstruction using the weights method

- In the legacy system, the amplitude is reconstructed with the ‘weight method’ (FIR) implemented in the FENIX chip.
- First step was to modify the whole algorithm so that it works per channel instead of 5x5 crystal matrices
- Logic was implemented using both VHDL and High Level Synthesis[†] (HLS) and successfully compared with results of simulation in CMSSW. Only a starting point since algorithm is not optimized for the Phase 2 TIA signal shape. It allows to evaluate the resources needed.

[†] - HLS allows rapid development of algorithms into RTL logic for FPGAs by designers with little to no RTL experience using a high-level language.



Resources utilization with the weights method

Preliminary



| Resources for 300 channels | KU115 (BCP baseline) | | KU15P | |
|----------------------------|----------------------|------------|----------|------------|
| | VHDL | HLS | VHDL | HLS |
| LUT | ~ 8 % | ~ 17 % | ~ 11 % | ~ 22 % |
| FF | ~ 3 % | ~ 11 % | ~ 4 % | ~ 14 % |
| DSP | ~ 27 % | ~ 27 % | ~ 76 % | ~ 76 % |
| Latency (w/ 160MHz clock) | 8 (2 BX) | 1 (1/4 BX) | 8 (2 BX) | 1 (1/4 BX) |

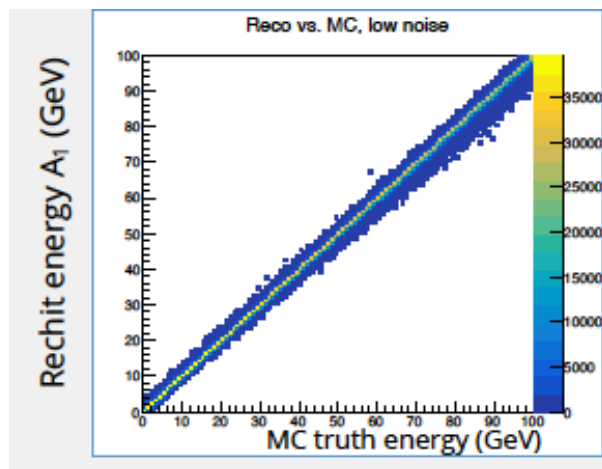

N.Loukas
L.Lutton
N.Marinelli
(NotreDame)

- Single crystal trigger primitive generation – weights method.
 - VHDL synthesis - based on FENIX chip
 - HLS synthesis - based on FENIX chip
- High-Level Synthesis (HLS) optimizes FPGA resources or latency
 - It depends how you write HLS code
- KU115 looks more appropriate choice for BCP than KU15P



Amplitude reconstruction with multi-fit (1/2)

- Linearized Multi-fit is a possible TPG algorithm based on extracting the signal amplitude while mitigating the out of time pile-up (OOT PU). Data are fit with the sum of a signal template plus N components for PU.
- Method developed and largely used offline for Run2
- We started investigating the possibility of applying it online in a very simplified/affordable form
- Only signal and Bx0 +1 and -1 templates used, 12 samples at 160 MHz. Signal shape is the TIA. HLS implementation based only on additions and multiplications.
- Tested with Toy Monte Carlo (MC) in input. First results encouraging. HLS reproduces well the input signal. Resources (next slide) need to be optimized.



L.Lutton
N.Marinelli
(NotreDame)
J.Hakala
(Virginia)



Amplitude reconstruction with multi-fit (2/2)

Preliminary



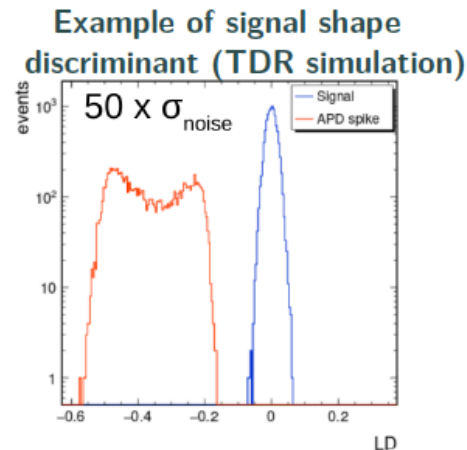
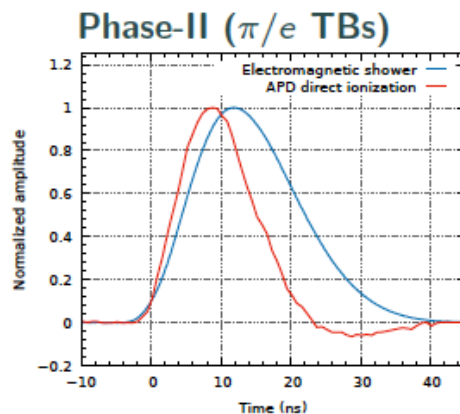
| Resources for 300 channels | KU115 (BCP baseline) | KU15P |
|----------------------------|----------------------|-----------|
| LUT | ~ 12 % | ~ 15 % |
| FF | ~ 9 % | ~ 11 % |
| DSP | ~ 43 % | ~ 120 % |
| Latency (w/ 160MHz clock) | 3 (3/4 BX) | 3 (3/4BX) |

L.Lutton
N.Marinelli
(NotreDame)
J.Hakala
(Virginia)

- Single crystal primitive with linearized multi-fit – VERY preliminary version
- The number of DSPs in an FPGA is an important parameter
 - Considering that today the cost of an FPGA with a larger number of DSPs is 3 times higher than our current choice, KU115 is still a good choice.
- 43% DSP can be reduced by increasing slightly the latency which currently is low (<1BX)

Spike Killing (1/2)

- Spike killing will be based on the signal shape analysis: the TIA pulse will allow online discrimination between spiked and scintillation pulses (<https://cds.cern.ch/record/2283187/files/CMS-TDR-015.pdf>)



- Algorithm based on linear discriminant has been tested offline and proved successful during test beam 2018
- The algorithm was translated into HLS to provide first evaluation of time and resources

Spike Killing (2/2)



| Resources for 300 channels | LD2 algorithm <u>KU115</u> | LD3 algorithm <u>KU115</u> |
|-------------------------------|-------------------------------|-------------------------------|
| LUT | ~ 9 % | ~ 17 % |
| FF | ~ 8 % | ~ 19 % |
| DSP | ~ 21 % | ~ 70 % |
| Latency (w/ 160MHz clock) | 3 (3/4 BX) | 5 (1+ 1/4 BX) |

- Two different algorithms were tried: L2 (use 2nd order polynomial) and L3 (3rd order polynomial) which clearly lead to different resource consumption

Algorithm summary

- If we add resources from different algorithms showed in the previous slides we would sum:
 - LUT ~ 1/3 of KU115
 - FF ~ 1/3 of KU115
 - DSPs ~ 2/3 (with LD2) of KU115
 - ~ 2 BX latency for single crystal TPG + spike killing which should both be wisely enlarged to save resources

- Are we done? Not at all !
 - A lot more to be implemented and tested
 - But we have some estimates on our needs for DSPs which are needed by algorithms and are not needed for other parts of the firmware like buffering, interfacing, controlling, etc ...


- These preliminary DSP requirements support FPGA pick:
 - XCKU115 Xilinx Kintex UltraScale FPGA





Demonstrator Electronics BCP V1

BCP V1 - The Off-Detector Demonstrator Electronics



- ATCA form-factor
- Single FPGA: XCKU115-2E
- Half capability of a baseline BCP
- 64 total bi-directional Multi-Gigabit Transceivers (MGTs)
- Supports 12 ECAL towers (or 1 HCAL wedge)
 - 12 TX / 48 RX IpGBT @ 10.24Gbps
 - 15 L1 Trigger @ 16Gbps
 - 4 DAQ @ 16Gbps
 - 2 TCDS / LDAQ @ 10.3Gbps
 - 1 ELM Control @ 10.3Gbps

BCP V1 - It's Alive!

- Initial board power-up went smoothly with no issues
- IPMC, ESM and ELM worked “out of the box” with full Ethernet connectivity
- After development on IPMC, could manually check individual DC/DC circuits
- Initial results are good with stable and low noise power
- Stepped through bring-up of ELM, clock ICs (Si5345), FPGA & FireFly
- All successful

Board Designers:



Stephen Goadhouse



Nikitas Loukas

Samtec FireFly
16 Gbps FOTRX
60 channels

Xilinx KU115
UltraScale FPGA



Embedded Linux Mezzanine (ELM)
Zynq-based board controller from UW

Ethernet Switch Module (ESM)
from UW

Mini ATCA BP
from Alex Madorsky



Zynq-IPMC management controller from UW

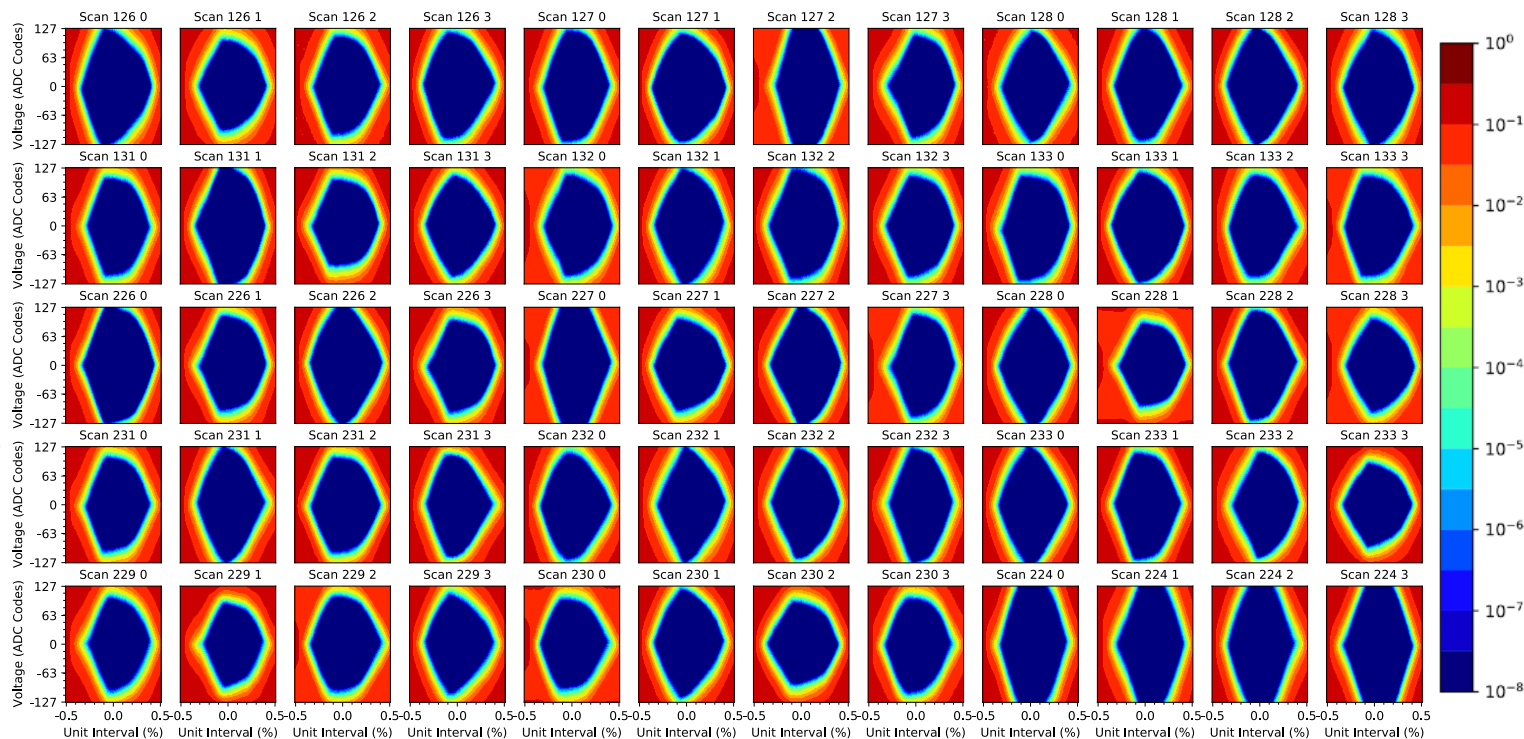


BCP V1 - So do the links work?

Yes!

- Programmed FPGA with iBERT core and connected FireFly together with optical loopbacks to test links and capture eye diagrams
- setup: PRBS-31, no signal equalization (incl. FireFly), dwell BER: 1×10^{-8}

- Most FireFly are “16Gbps”, running at their max of 15.7Gbps with 64ps UI
- Scans 224 0-3 use 25G FireFly running at UltraScale max of 16.3Gbps
- PRBS-31 ran 6 hours: 0 errors BER: 2.7×10^{-15} (per channel)



Graph from modified version of [iBERTpy](#) by [mvsoliveira](#)

BCP V1 - What we know works

DC/DC Circuits
all across board

Board Temp.
Voltage &
Current
Sensors


FPGA SFP →

IPMC
interface →

Ethernet →

Host USB →

Embedded
USB JTAG
Programmer
And Serial
Consoles →



Overall, all issues
have been minor
with work-arounds
that allow us to
progress

← Clock Generators

← Ethernet
Switch Module



← ATCA Power
-48V to +12V




Design Process Highlights

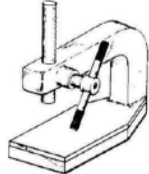
BCP V0.5

- One of the most expensive mistakes is incorrect connector placement
- So, after primary components placed in BCP V1 design, a branched version was quickly created and produced
 - Used outer layers + 2 int. GND/PWR layers - FR4 - very cheap
 - Minimal routing to power and connect IPMC and ELM for early development (UART, Ethernet, some sensor I2C + JTAG)
- Press-fit install practice
- Verified connector placement & clearance
- Verified ATCA fit
- Verified front panel
- Verified ATCA -48V to +12V power circuits
- Had customized IPMC design for BCP V1 before we had the V1 boards
- Discovered issue with ESD strip - fixed in V1
- Changed front panel vendor



BCP V0.5 - Press-fit Components



- Press-fit connector compliant pins forced into holes w/o solder
- Requires tight hole tolerance
- Use Arbor press for even, vertical force 
- Most components require 100 to 225 Kg of total force
- Zone 2 conn. recommends 7-9 Kg per pin x 160 pins = 1.4 metric tons of force!
- Designed custom PCB support and top tooling to work with UW's base
- Focuses force where needed and protects connectors from damage
- Top tools have exact same height off of PCB so can set a safety stop on press
- BCP custom tooling designs [here](#)




D. Parenti & J. Mitchell
circuitsassembly.com




WISCONSIN
UNIVERSITY OF WISCONSIN-MADISON
Support: T.Gorski

BCP V1 - PCB Stack-Up

- High-speed digital signals benefit from low loss engineered laminate materials
 - Isola Tachyon 100G material
 - Very Low loss dielectric & ultra smooth copper
 - Top-of-the-line material
 - Considered saving a few \$ with a mid-range material
 - Ultimately decided that we would take every advantage we could to maximize success
- Blind via span layers 1-9 and 10-18
 - Routed signals so that no back-drilling was needed
- All vias filled with non-conductive epoxy and over-plated with copper
 - Provides more reliable pads, vias and surface mount connections
- All GND and most Power layers are 1 oz copper for higher ampacity
- Have many power segments so decided to bump stack-up from 16 to 18 layers for more power layers
 - Does add a slightly longer blind via stub
 - Reduced layout time



BCP V1 - Other Features




14-pin
Tag-
Connect
instead
of
0.1”
header
33%
smaller

- Tied a few unused MGTs and I/O to SFPs
- SMA and LEMO debug clock and I/O ports
- Three Si5345s for very flexible clock tree
 - Every MGT has LHC synchronous and asynchronous reference clocks
 - Spare outputs tied to spare inputs
 - Will be used to optimize clock tree and experiment with jitter measurements
- Manual control over power enables for testing DC/DCs
- Legacy TCDS Receiver circuit from AMC13


FTDI FT4232H USB Quad UART/MPSSE

- Single USB Device Port
- IPMC Serial Console
- ELM Serial Console
- IPMC JTAG Programming / Debug (Xilinx Virtual Cable)
 - Complements XVC access of ELM and FPGA over Ethernet
 - Networked computer or ELM on other BCP can host server
- Access to I2C and SPI busses for direct control
 - May use for semi-automated production test



BCP V1 - Summary


- A lot has been done, even with COVID-19 quarantine
- A lot left to do
- Core functions work and much more
- We have a very capable development platform





Backup

ECAL Barrel architecture (1/2)



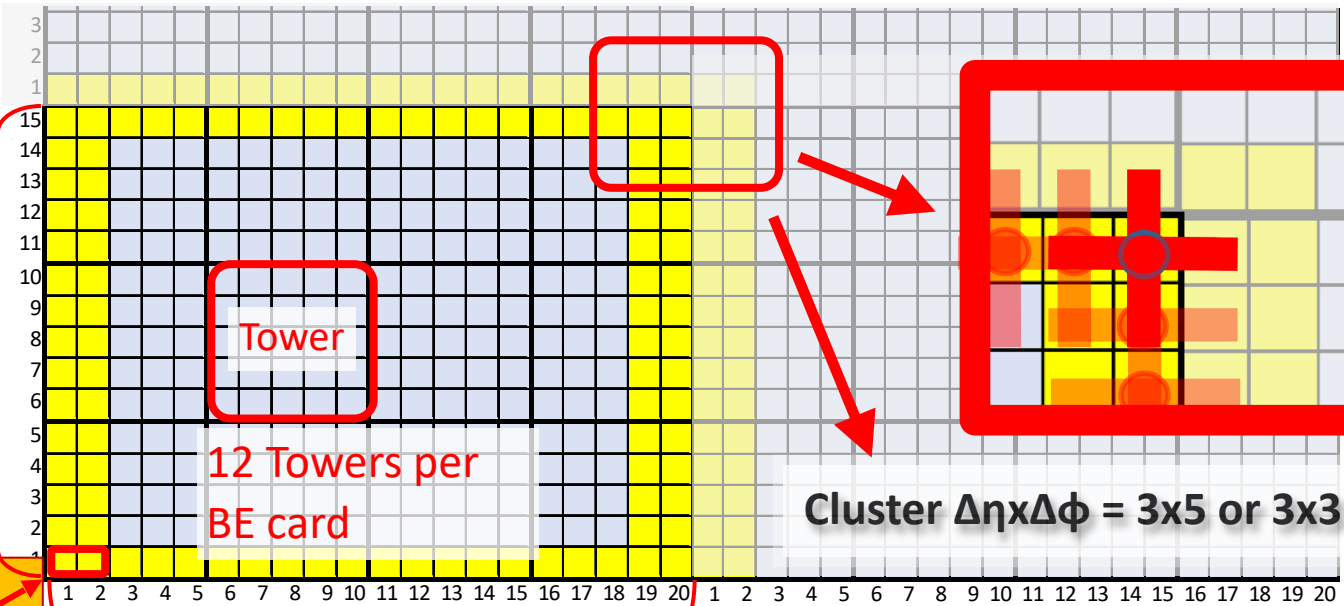
- FE ↔ BE links: 12240
- BE to trigger links: 3060
- Each baseline BCP has 2 FPGAs and supports 24 ECAL towers, 3 BCPs / SM ϕ
- Total number of BCPs: 3 x 36 = 108

ECAL Barrel architecture (2/2)

Right or left neighbors:
 30 Crystals x 16 bits
 x 40MHz = 19.2Gb/s
 => 2 links @16Gb/s

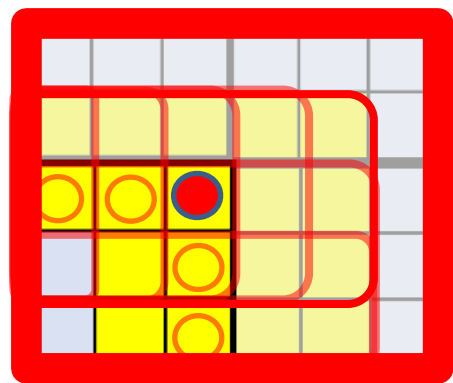
Corner neighbors:
 2 Crystals x 16 bits x
 40MHz = 1.28Gb/s =>
 2 links @~1Gb/s

6 Transceivers in total
 4 Right/left
 4 Top/Bottom
 8 SERDES
 1 or 2 on each corner




Top or bottom: 20
 Crystals x 16bits x
 40MHz = 12.8Gb/s =>
 1 or 2 links @16Gb/s

Cluster $\Delta\eta \Delta\phi = 3 \times 5$ or 3×3 xtls



→ Developed a solution for optical sharing crystal of data between FPGA regions, based on technology used for Phase 1 L1 Calo Trigger upgrade.

HCAL Barrel architecture



- Each baseline BCP has 2 FPGAs and supports 2 HCAL wedges or 1/2 BCP / RBX


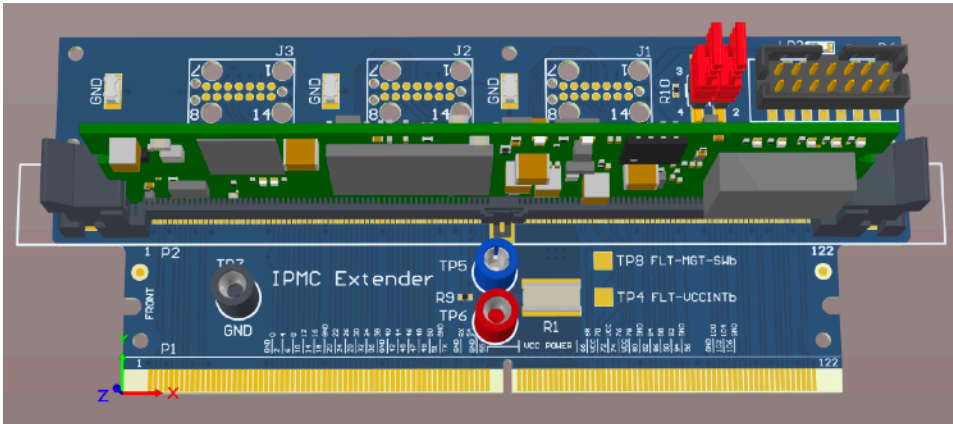
- Total number of BCPs: $0.5 \times 36 = 18$



Other HCAL Subsystems

- HCAL Forward (HF): fully supported by a single ATCA crate of BCPs
 - 72 Clock/Control/Status Links: 4.8 Gbps or 2.4 Gbps
 - GBT-encoded
 - 864 Data links: 5.0 Gbps links
 - HF data links are very similar to HB data links.
 - 16 calibration links: 5.0 Gbps links
 - HF data links are very similar to HB data links
- HCAL Outer (HO): fully supported by a single ATCA crate of BCPs
 - 840 Data Links: 1.6 Gbps 8b10b-encoded
 - 36 Fast Control Links: 80 Mbps TTC-encoded

IPMC Extender



- Need to test FPGA voltage sequencing
 - Bad sequencing could harm the expensive FPGA
- Power Enable signals are not fully broken out in BCP V0.5 nor BCP V1
- Probe wires soldered to BCP V1 not possible since these signals are under DIMM socket
- Solution: a simple IPMC extender board
- Can be used stand-alone or in BCP
- Breaks-out needed PWR-EN / PGOOD pins
 - Jumper to Eval Boards for the DC/DC circuits to test with BCP V0.5
 - Sequence can be verified in real-time with oscilloscope
- Also add manual 12V disconnect so can bring up initial board with only 3.3V for IPMC without a bug in IPMC s/w also bringing up 12V
- Added a sense resistor to measure IPMC current draw to better sense requirements of limited Management 3.3V rail
- <https://twiki.cern.ch/twiki/bin/view/CMS/lpmcExtender>