# Firmware Development for the Phase-2 Upgrade of the Calorimeter L1 Trigger for the CMS Detector at the LHC

Piyush Kumar & Bhawna Gomber in Collaboration with
University of Hyderabad, Princeton University, and University of Wisconsin-Madison


Compact Muon Solenoid

# High-Luminosity LHC (HL-LHC)

- Luminosity: indicate the performance of an accelerator
  - Proportional to: number of collisions that occur in a given amount of time
  - higher the luminosity: the more data the experiments can gather
- Aim: to deliver a much larger dataset for physics to the LHC experiments
- Pile-up: Number of simultaneous proton-proton interactions (~200)
  - With high pile-up, need more advanced selection algorithms at L1 trigger
- This increased datasets will help in the high precision measurements of:
  - Standard model (SM)
  - new territories beyond the SM (BSM)

| | Instantaneous Luminosity | Pile-up (average) | Integrated luminosity |
|---|---|---|---|
| Run-2 | $2.1 \times 10^{34}$ cm$^{-2}$s$^{-1}$ | 55 | 160 fb$^{-1}$ (4 years) |
| HL-LHC (baseline) | $5 \times 10^{34}$ cm$^{-2}$s$^{-1}$ | 140 | 3000 fb$^{-1}$ (10 years) |
| HL-LHC (ultimate) | $7.5 \times 10^{34}$ cm$^{-2}$s$^{-1}$ | 200 | 4000 fb$^{-1}$ (10 years) |

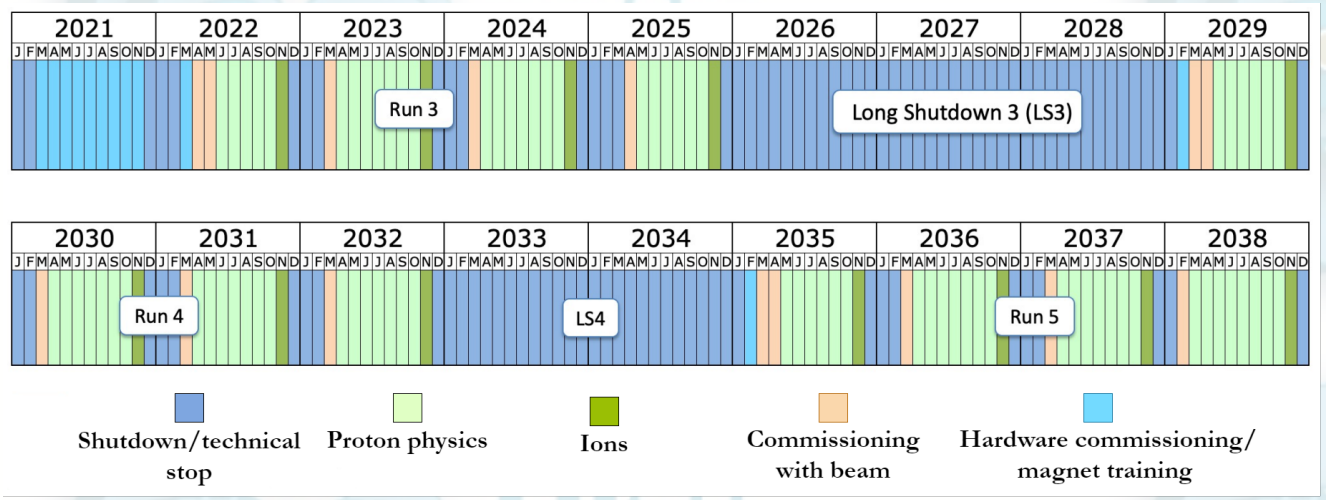


Fig: HL-LHC timeline

# CMS HL-LHC upgrade

- The CMS detector planned upgrade for the HL-LHC era:
  - New pixel and strip tracking detector
  - New high-granularity calorimeter (HGCAL) of the endcap
  - New frontend/backend electronics for the:
    - Barrel calorimeter
      - Electromagnetic calorimeter (ECAL)
      - Hadronic calorimeter (HCAL)
    - Muon system
      - Drift tube (DT)
      - Cathode strip chambers (CSC)
  - 40 MHz Scouting system
    - can be used to scrutinize the collision events and identify potential signatures unreachable through standard trigger selection processes
  - L1 trigger:
    - Inclusion of the tracker information
    - Extensive usage of:
      - large FPGA (Virtex UltraScale+/Kintex UltraScale)
      - high-speed optical links (28 Gbps)

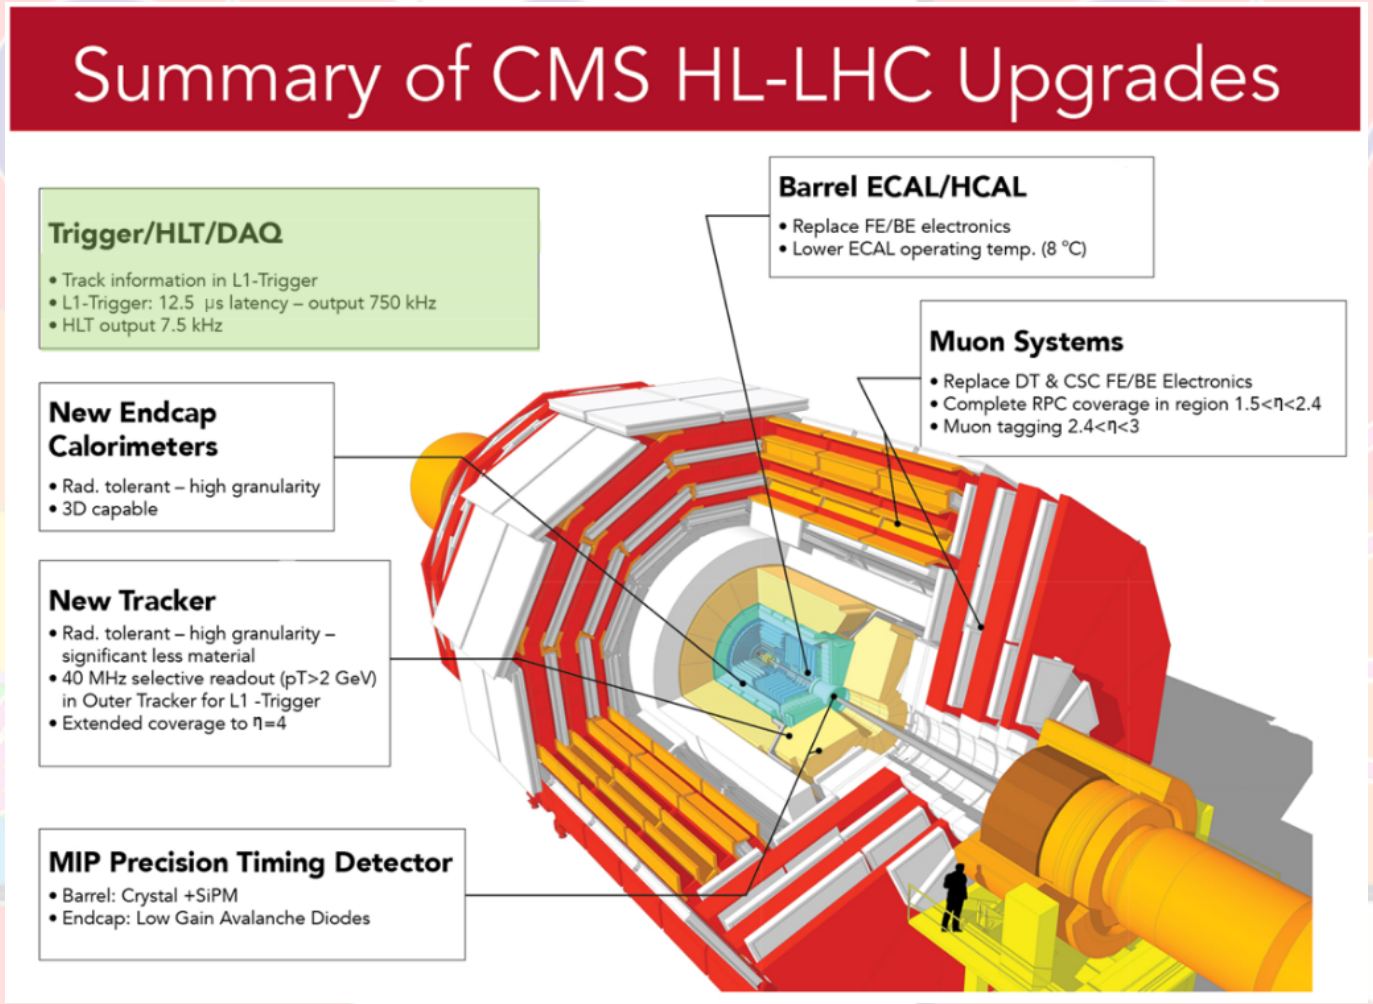


Fig: CMS detector HL-LHC upgrade

# L1 trigger principle

- At design parameters the LHC produces:
  - ~$10^9$ events/second in CMS detectors.
  - each event is ~ 1 MB.

- $10^9$ events/s x 1 Mbyte/events = $10^{15}$ bytes/s = 1 PB/s (1 Petabyte/second)

- Problem:
  - It is impossible to store and process this large amount of data
- Solution:
  - a drastic rate reduction has to be achieved
    - Level-1: 40 MHz to 750 kHz
    - High level trigger (HLT): 750 kHz to 7.5 kHz

- A trigger is designed to reject the uninteresting events and keep the interesting ones for physics.

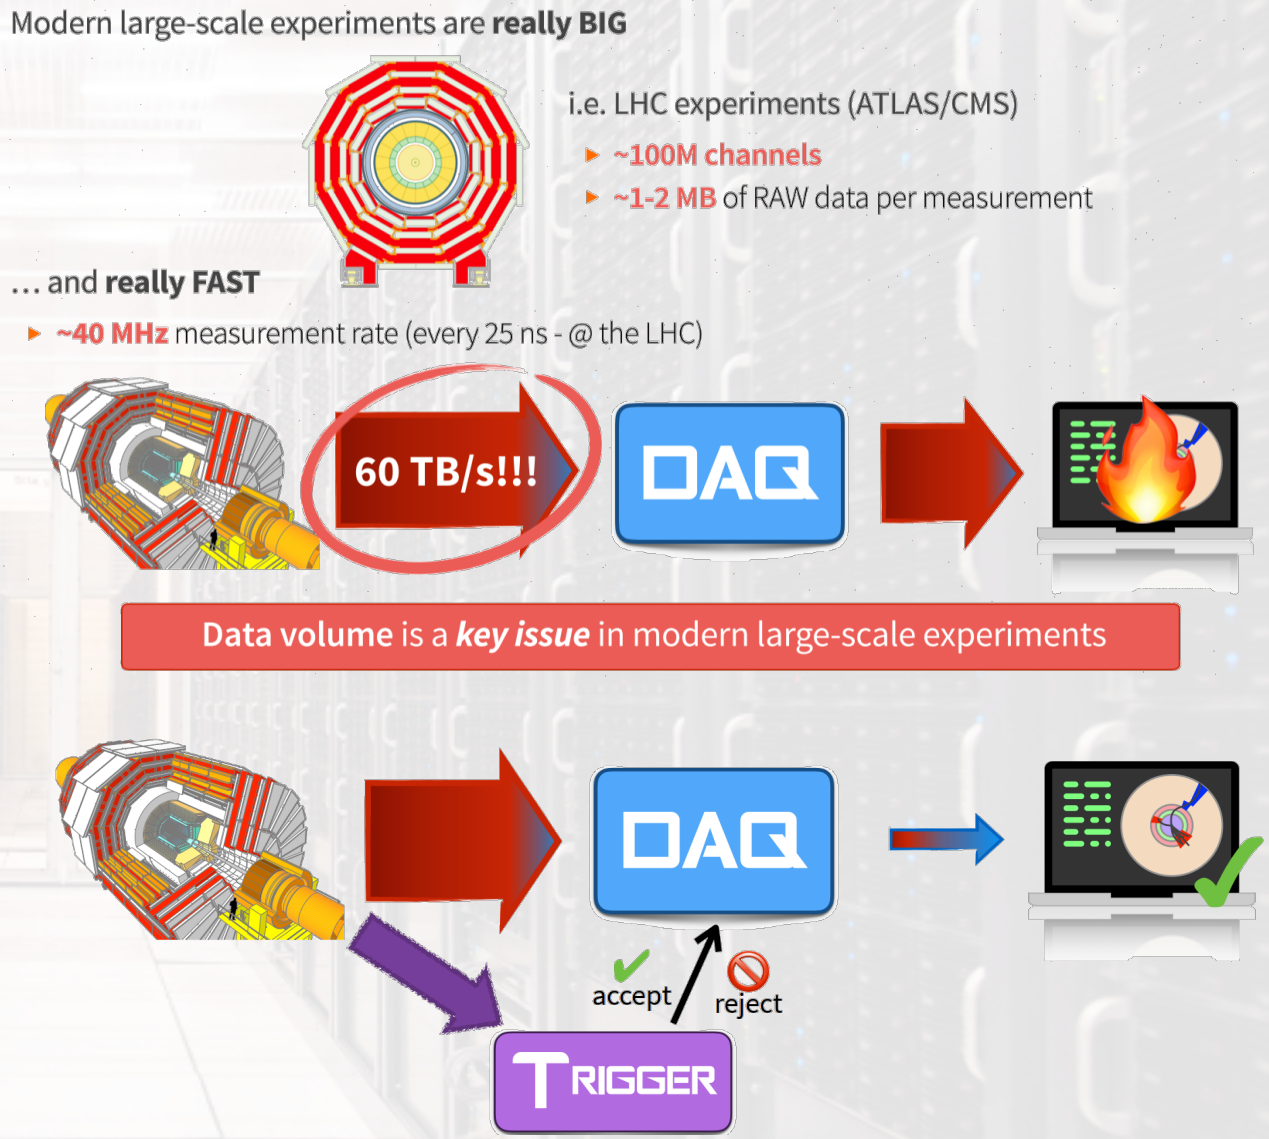


Fig: Trigger system

# CMS L1 trigger principle – a hardware perspective

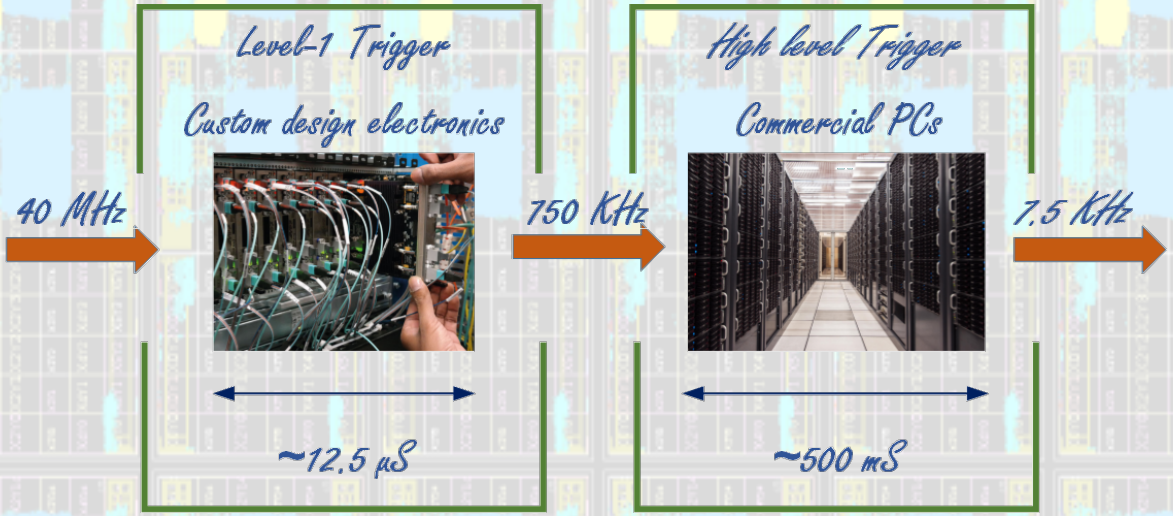


CMS two-level triggering

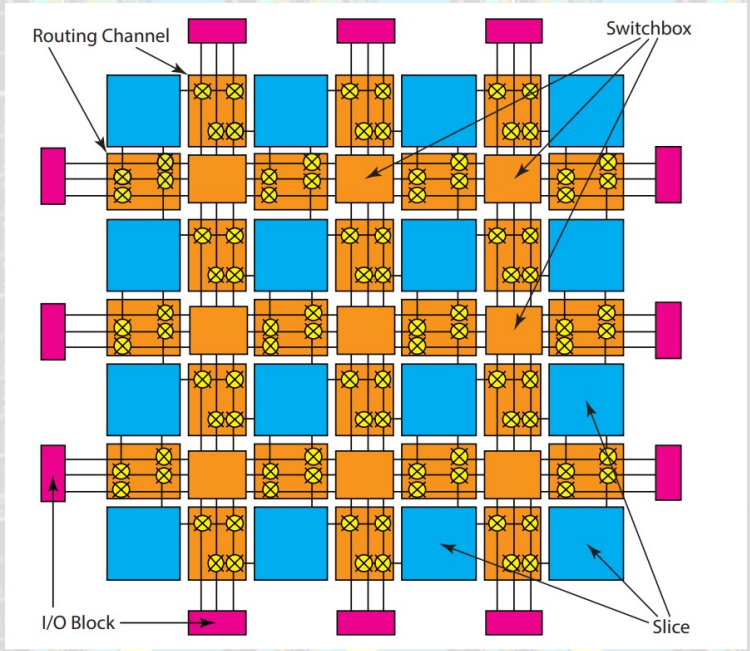


FPGA architecture


Virtex UltraScale+ Device

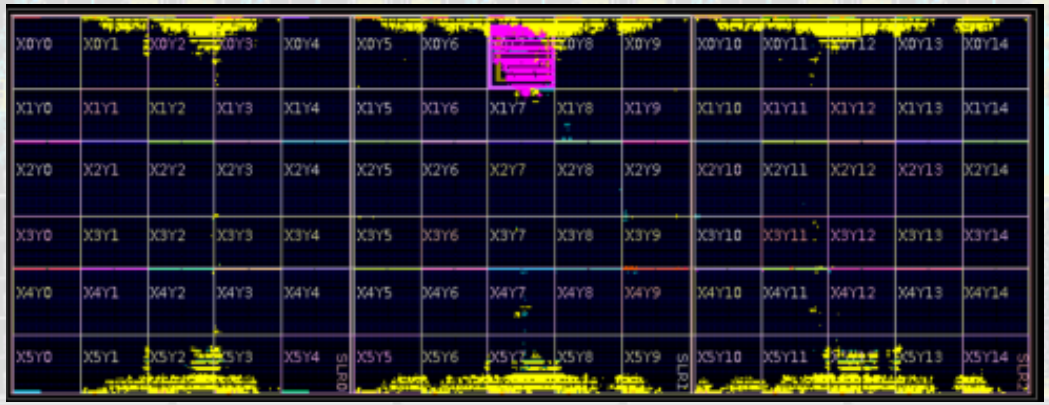


FPGA logic firing

# Calorimeter L1 trigger architecture (Barrel/Endcap/HF)

- Calorimeter trigger is processed in the following steps:
  - Barrel: Regional calorimeter trigger (RCT) and Global calorimeter trigger (GCT)
  - HF and HGCAL: GCT

- The barrel calorimeter covers the geometry of $|\eta| = 1.48$
  - Processed in two steps
    - RCT: currently 36 XCVU9P (3 SLR#) FPGAs
    - GCT: 3 XCVU9P FPGAs

- RCT geometry for the FPGA processing:
  - $17\eta \times 4\phi$ of the barrel (for 36 FPGAs)
  - $17\eta \times 6\phi$ of the barrel (for 24 FPGAs)

- GCT geometry for the FPGA processing:
  - 12 unique RCT ($17\eta \times 4\phi$ ) + 4 neighbours
  - 8 unique RCT ($17\eta \times 6\phi$ ) + 4 neighbours

- Input rate: 40 MHz
- Increased output rate: 100 kHz => 750 kHz
- Increased latency: 3.8 µS => 12.5 µS

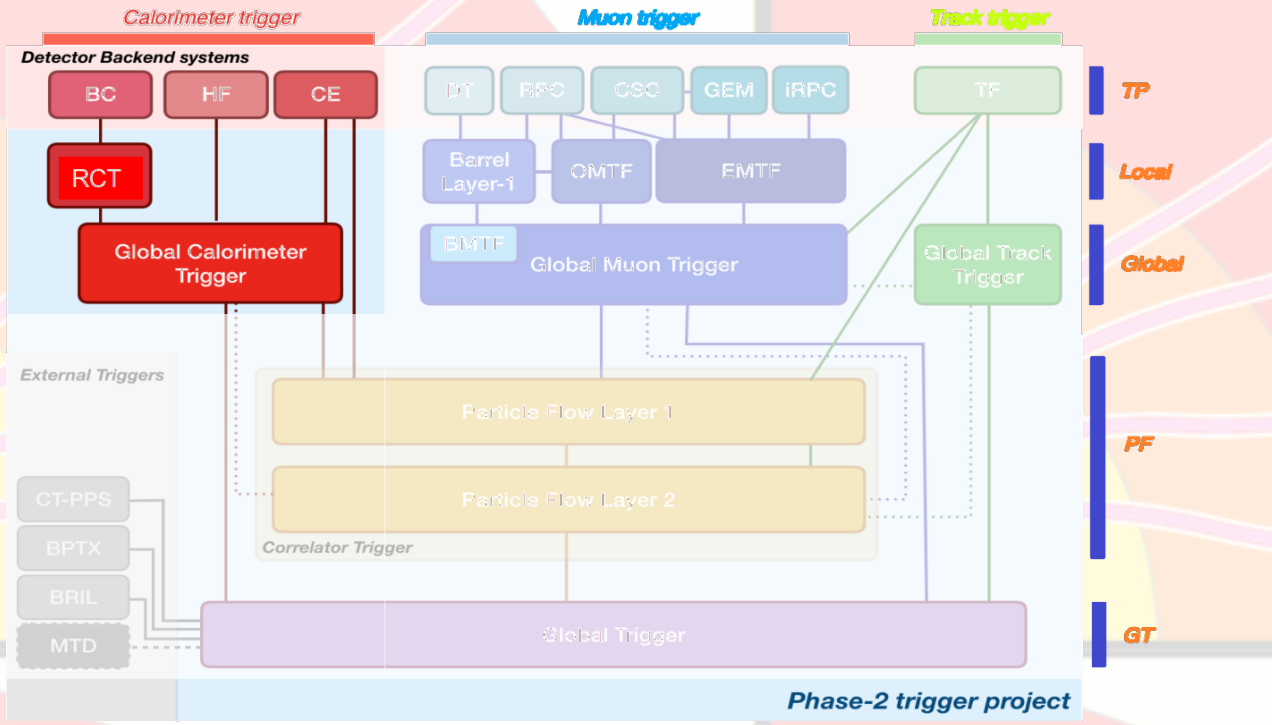


Fig: Calorimeter trigger architecture

#: Super logic region

# Barrel calorimeter segmentation

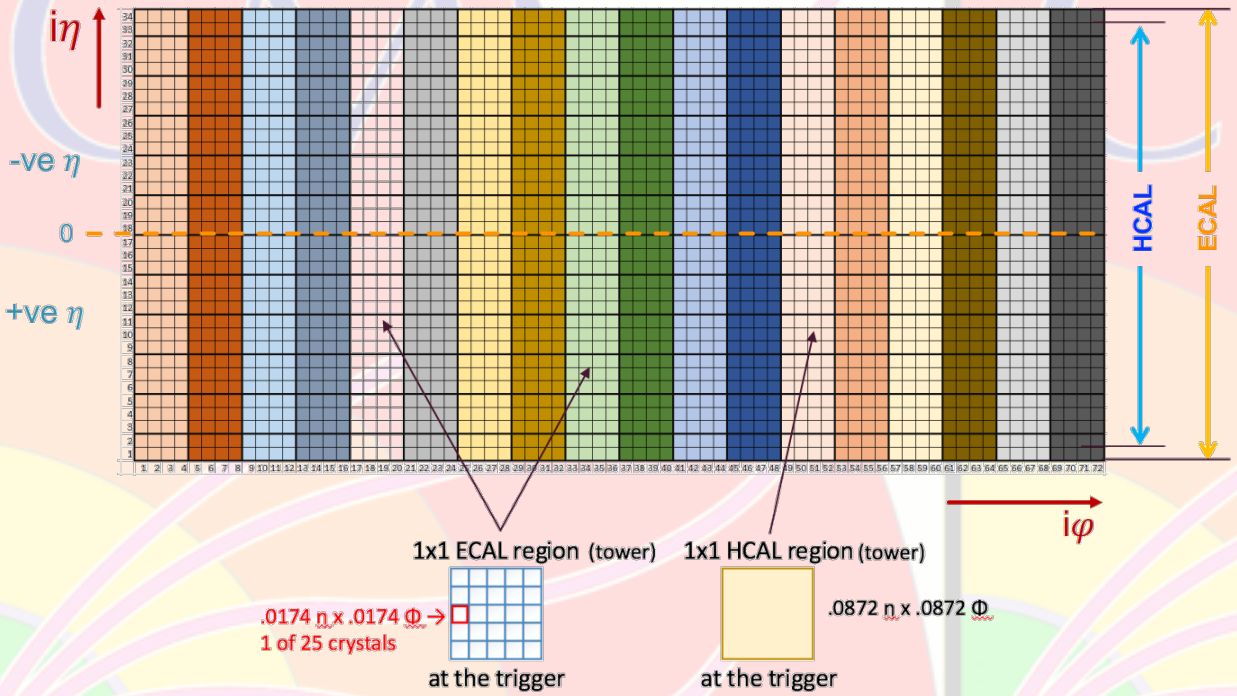


Barrel calorimeter segmentation

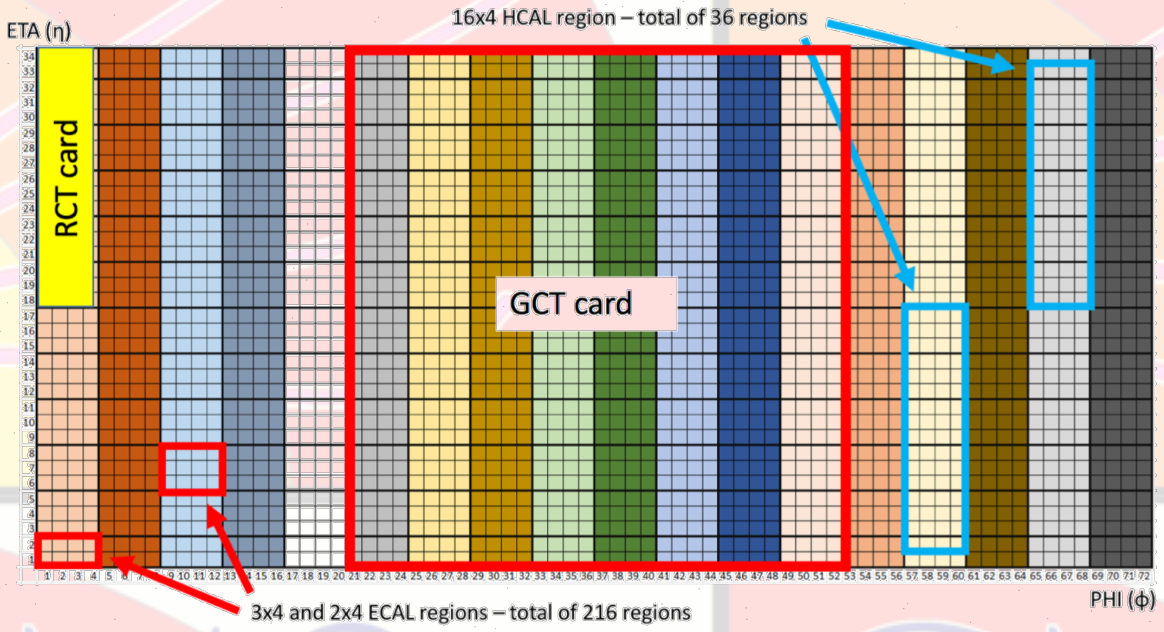


RCT and GCT processing region(s)

# Trigger Firmware Development

- The trigger algorithms are implemented by using a high level synthesis tool
  - Rapid prototyping
  - Codes are written in C++
  - HLS synthesizes the code to generate the HDL (hardware description language) and
  - Provide an early estimate of latency and resource utilization
  - Increased ease of collaboration and code sharing for algorithm design

**Project Flow:**

- Algorithm development/MC simulation
- HLS algorithm implementation and simulation
- HLS algorithm integration into firmware shell
- Standalone testing on hardware
  - ($R_x$ buffer -> Algorithm -> $T_x$ buffer)
- Testing and final validation of results through the entire chain
  - MC == HLS Sim == HDL Sim == HW Results

*Aim is to write HLS algorithms in a framework agnostic way*

Performance Estimates

Timing (ns)

Summary

| Clock | Target | Estimated | Uncertainty |
|---|---|---|---|
| ap_clk | 4.17 | 2.917 | 1.25 |

Latency (clock cycles)

Summary

| Latency | | Interval | | |
|---|---|---|---|---|
| min | max | min | max | Type |
| 158 | 158 | 6 | 6 | function |

Summary

| Name | BRAM_18K | DSP48E | FF | LUT | URAM |
|---|---|---|---|---|---|
| DSP | - | - | - | - | - |
| Expression | - | - | 0 | 520 | - |
| FIFO | - | - | - | - | - |
| Instance | 0 | 0 | 280881 | 397283 | - |
| Memory | - | - | - | - | - |
| Multiplexer | - | - | - | 12656 | - |
| Register | 0 | - | 7808 | 512 | - |
| **Total** | 0 | 0 | 288689 | 410971 | 0 |
| Available | 4320 | 6840 | 2364480 | 1182240 | 960 |
| Available SLR | 1440 | 2280 | 788160 | 394080 | 320 |
| Utilization (%) | 0 | 0 | 12 | 34 | 0 |
| Utilization SLR (%) | 0 | 0 | 36 | 104 | 0 |

Fig: Vivado-HLS performance estimates of trigger algorithm

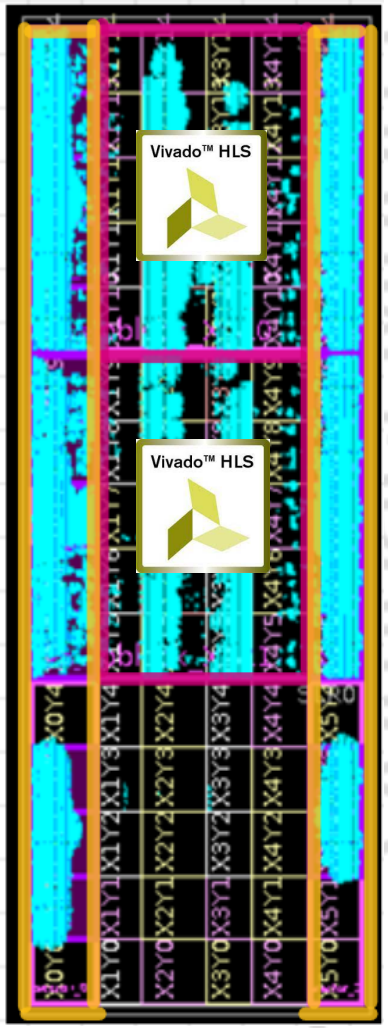


Fig: Trigger algo device implementation

# RCT algorithm

- Input:
  - ECAL: 17x4x5x5 crystals
  - HCAL: 16x4 towers
- Output: calorimeter clusters and towers information
  - Cluster energy
  - Tower energy
  - Peak (or seed) crystal energy and positional information
- Input-Output link BW: 25 Gbps
- Implemented only in 2 SLR (SLR2 and SLR1)
- Bitstream is successfully tested using the 113 bunch-crossing of LHC data (standalone mode).
  - Generated by MC simulation
- Latency: 269.466 ns (1 μs budget) @360 MHz

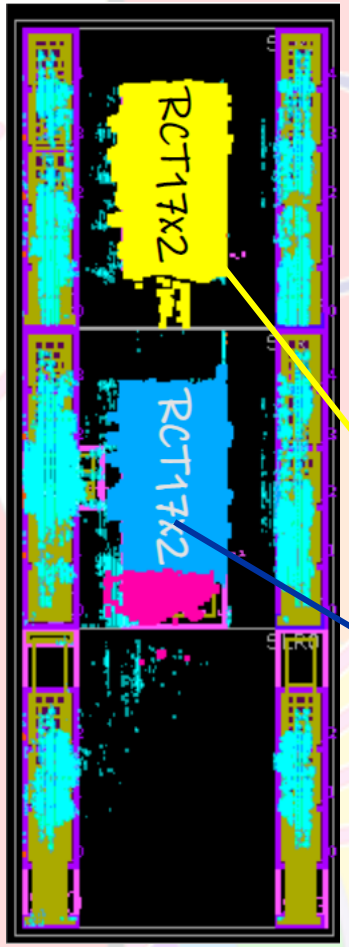


RCT17x2

RCT17x2

Two RCT17x2 combination

APx-FS-R2 (Firmware shell)

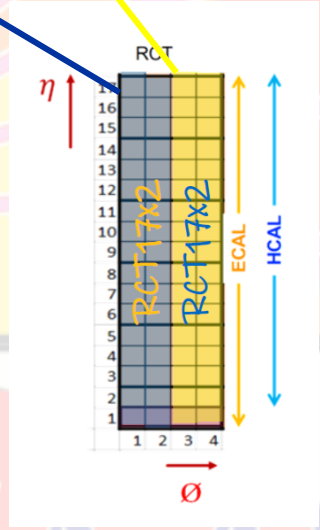


Fig: RCT algo Floorplan

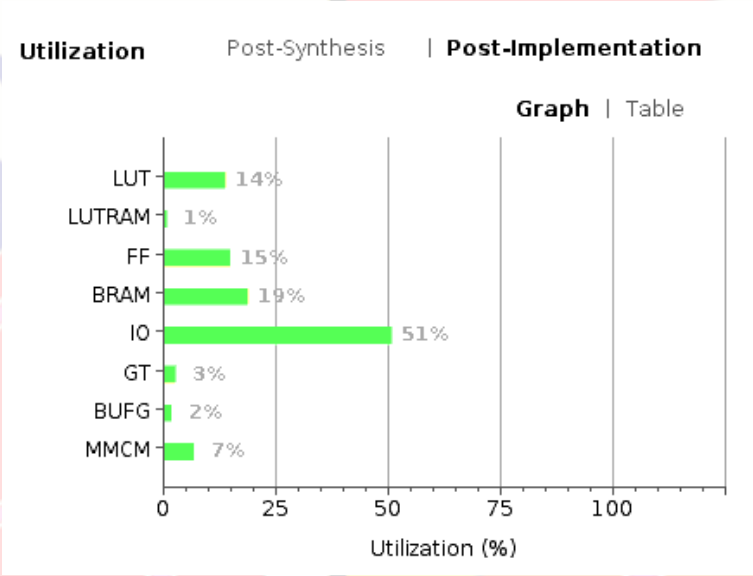


Fig: Utilization

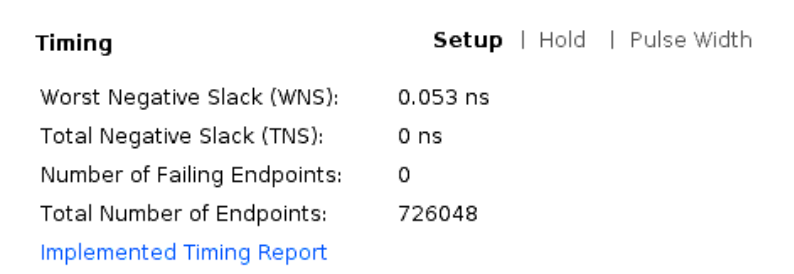


Fig: timing

# GCT algorithm

- GCT stitches makes clusters forwards to Correlator and Global trigger (GT)
- Sorts EGamma and hadronic clusters
- Firmware is developed for the MET algorithm
  - Processes 12 unique RCT cards
- MET utilizes the CORDIC algorithm
  - Calculate the cos and sin of LHC azimuthal angle ($\phi$)
- Calculates $E_x$, $E_y$ for missing transverse momentum (MET)
- Bitstream is tested successfully using output of 12 RCT cards (standalone mode)
- Latency: 163.902 ns

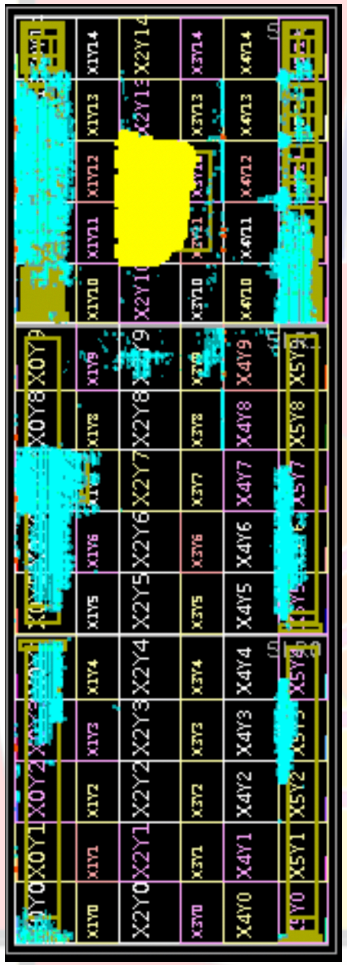

MET

APx-FS-R2 (Firmware shell)

Fig: MET floorplan

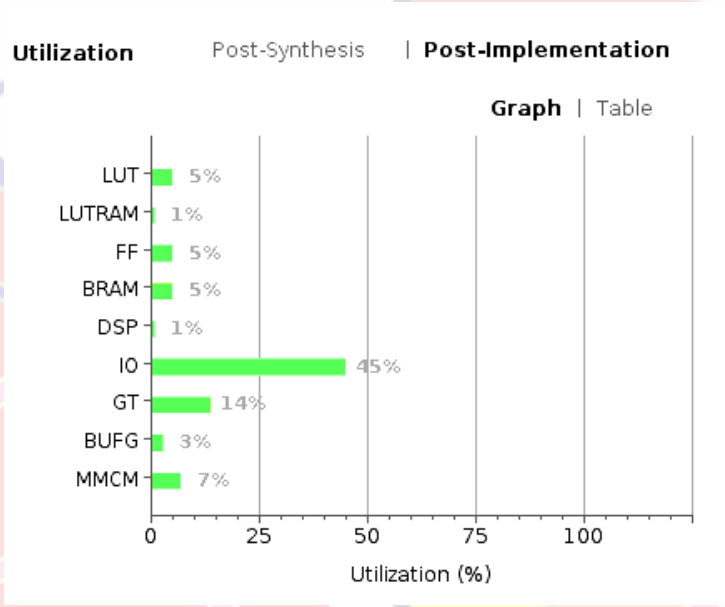


Fig: Utilization

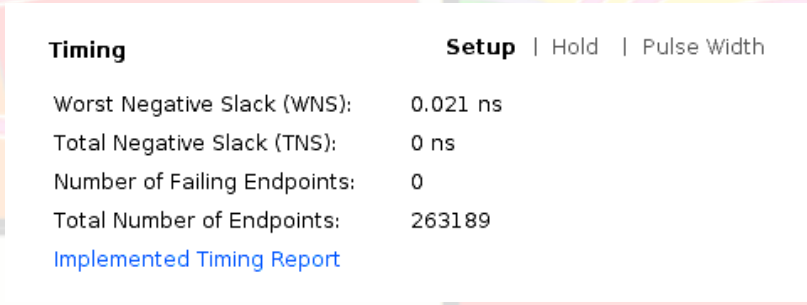


Fig: timing

# RCT and GCT multi-board test

- Setup at University of Wisconsin Madison lab
- 4 APd1 board (first Advanced Processor demonstrator) employed
  - 3 for RCT
  - 1 for GCT

Project Flow (Multi-board):

- Algorithm development/MC simulation
- HLS algorithm implementation and simulation
- HLS algorithm integration into firmware shell
- Standalone testing on hardware
  - ($R_x$ buffer -> Algorithm -> $T_x$ buffer)
- Multi-card testing and final validation of results through the entire chain
  - MC == HLS Sim == HDL Sim == HW Results

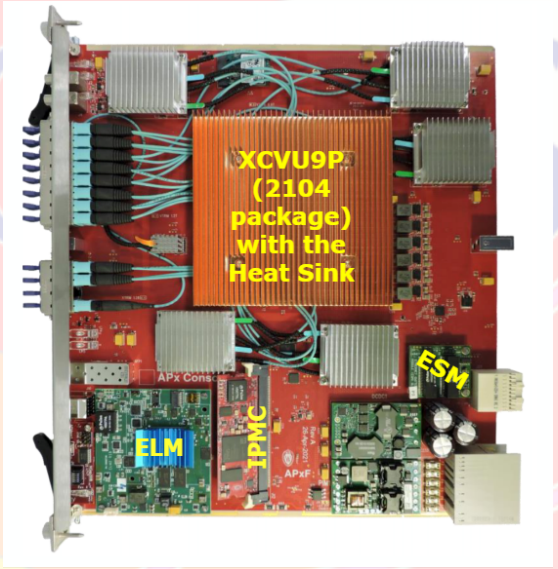


Fig: APd1 prototyped board

Fig: APd1 board at UW lab

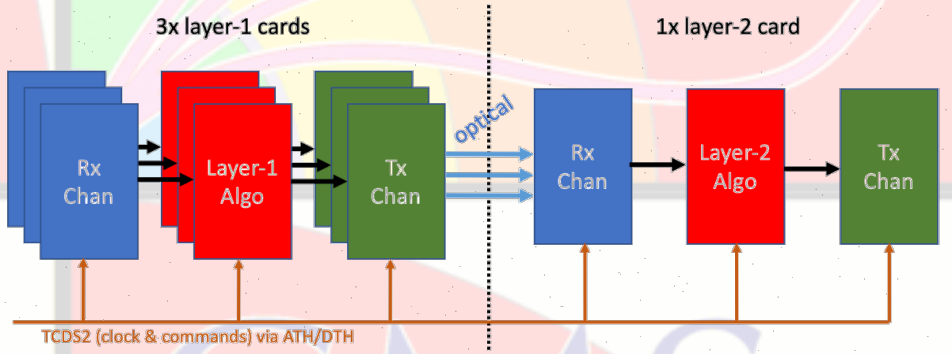


Fig: Multi-board test chain

# RCT and GCT multi-board test

Test Vectors from Phase-2 MC samples: ggH(bb), QCD, VBFHToTauTau

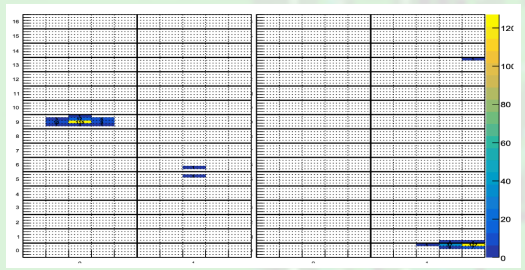

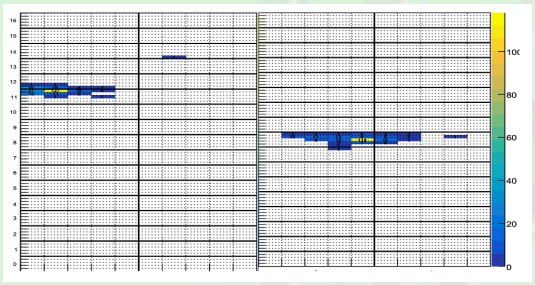

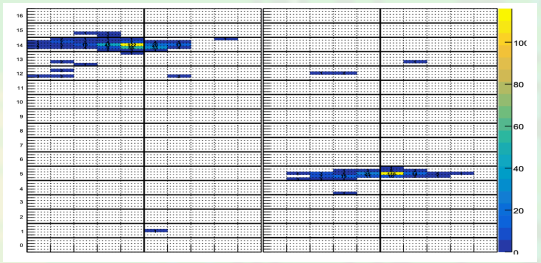

RCT-1 Links:68

RCT-2 Links:68

RCT-3 Links:68

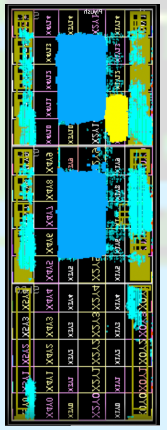

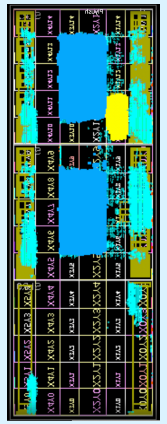

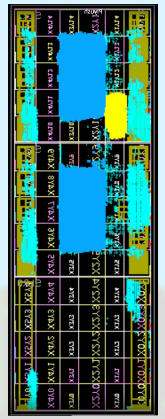

4 output links from each RCT board

Via optical fibres

GCT

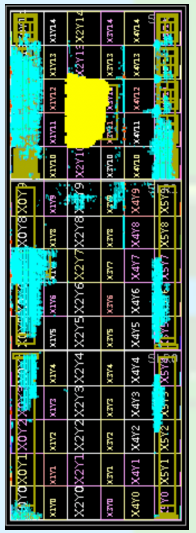

GCT receives 12 input links (3 RCT cards)

Data captured

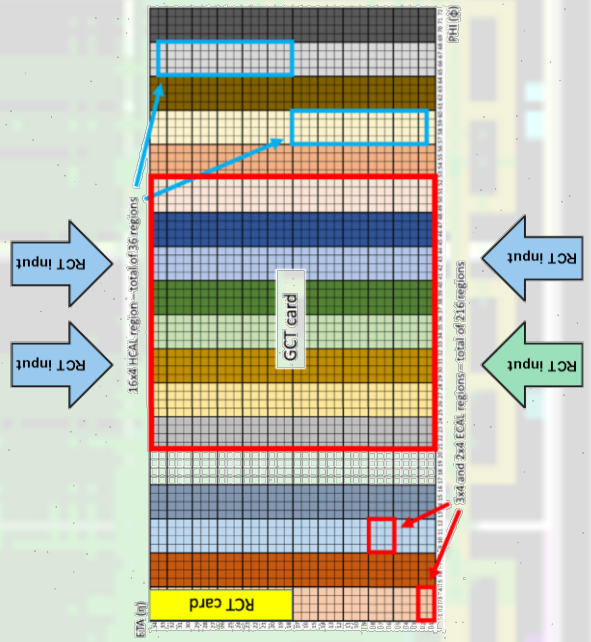

# RCT and GCT multi-board test

Test Vectors from Phase-2 MC samples: ggH(bb), QCD, VBFHToTauTau

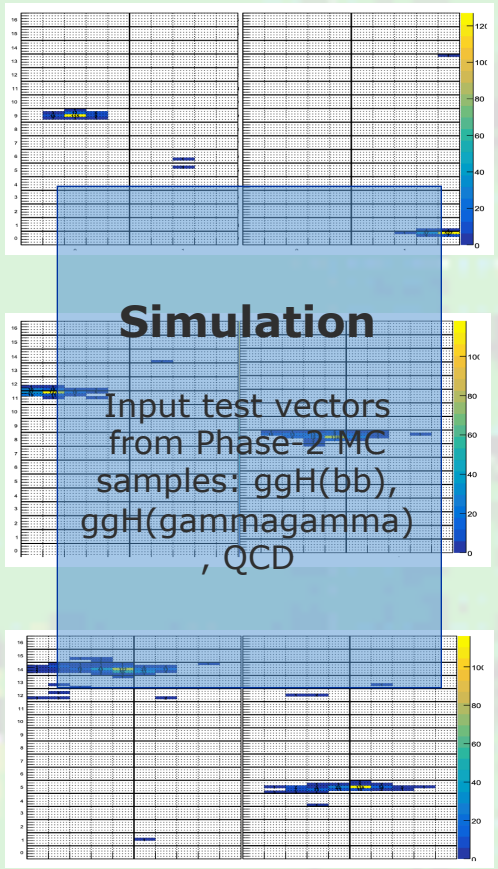

**Simulation**

Input test vectors from Phase-2 MC samples: ggH(bb), ggH(gammagamma), QCD

RCT-1 Links:68

RCT-2 Links:68

RCT-3 Links:68

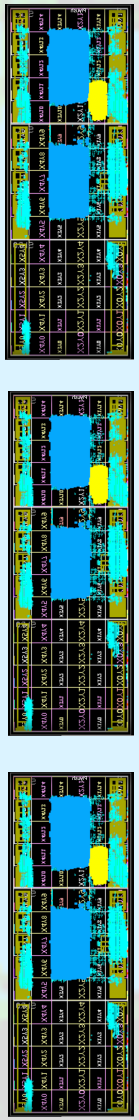

**RCT Output** (each card)

**MC outputs == HLS SIM == HW Buffer**

GCT

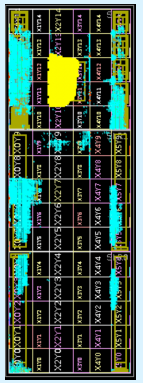

Data captured

## GCT Output: Simulation Agrees with Hardware 100%

| | Link_out_0_bitfile | link_out_0_ref | comparison | link_out_1_bitfile | link_out_1_ref | comparison |
|---|---|---|---|---|---|---|
| 2 | 000000030000001B9 | 000000030000001B9 | Match | 0000000200000086 | 0000000200000086 | Match |
| 3 | 0102000400060000 | 0102000400060000 | Match | 0194006006000700000 | 0194006006000700000 | Match |
| 4 | 0000005600000001 | 0000005600000001 | Match | 0000011000010003 | 0000011000010003 | Match |
| 5 | 000000000000003D | 000000000000003D | Match | 00000003000001D9 | 00000003000001D9 | Match |
| 6 | FF23FFFDFFFD0000 | FF23FFFDFFFD0000 | Match | 01A9000700090000 | 01A9000700090000 | Match |
| 7 | 0000FF3FFFFFFFFE | 0000FF3FFFFFFFFE | Match | 000000D200010003 | 000000D200010003 | Match |
| 8 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 9 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 10 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 11 | 000000030000013A | 000000030000013A | Match | 0000000200000081 | 0000000200000081 | Match |
| 12 | 0114000700000001 | 0114000700000001 | Match | 01AF000900000001 | 01AF000900000001 | Match |
| 13 | 0000005600000001 | 0000005600000001 | Match | 0000011000010003 | 0000011000010003 | Match |
| 14 | 000000000000002C | 000000000000002C | Match | 0000000300000150 | 0000000300000150 | Match |
| 15 | FF14FFFC00000000 | FF14FFFC00000000 | Match | 01C5000B00000001 | 01C5000B00000001 | Match |
| 16 | 0000FF3FFFFFFFFE | 0000FF3FFFFFFFFE | Match | 000000D200010003 | 000000D200010003 | Match |
| 17 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 18 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 19 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 20 | 011A000500010000 | 011A000500010000 | Match | 00D8000300000000 | 00D8000300000000 | Match |
| 21 | 0102000400060000 | 0102000400060000 | Match | 0194006006000700000 | 0194006006000700000 | Match |
| 22 | 00000000006B0000 | 00000000006B0000 | Match | 0000000001020000 | 0000000001020000 | Match |
| 23 | FFD2000000000000 | FFD2000000000000 | Match | 0160005000100000 | 0160005000100000 | Match |
| 24 | FF23FFFDFFFD0000 | FF23FFFDFFFD0000 | Match | 01A9000700090000 | 01A9000700090000 | Match |
| 25 | 00000000FF560000 | 00000000FF560000 | Match | 0000000000DD0000 | 0000000000DD0000 | Match |
| 26 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 27 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 28 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 29 | 011A000500010000 | 011A000500010000 | Match | 00D8000300000000 | 00D8000300000000 | Match |
| 30 | 0114000700000001 | 0114000700000001 | Match | 01AF000900000001 | 01AF000900000001 | Match |
| 31 | 00000000006B0000 | 00000000006B0000 | Match | 0000000001020000 | 0000000001020000 | Match |
| 32 | FFD2000000000000 | FFD2000000000000 | Match | 0160005000100000 | 0160005000100000 | Match |
| 33 | FF14FFFC00000000 | FF14FFFC00000000 | Match | 01C5000B00000001 | 01C5000B00000001 | Match |
| 34 | 00000000FF560000 | 00000000FF560000 | Match | 0000000000DD0000 | 0000000000DD0000 | Match |
| 35 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 36 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 37 | 0000000000000000 | 0000000000000000 | Match | 0000000000000000 | 0000000000000000 | Match |
| 38 | 0000000101AC0005 | 0000000101AC0005 | Match | 0000000100DD0002 | 0000000100DD0002 | Match |
| 39 | 0000000301701F0 | 0000000301701F0 | Match | 0000000401801C7 | 0000000401801C7 | Match |
| 40 | 0000000000000094 | 0000000000000094 | Match | 000000010000011D | 000000010000011D | Match |

Tested for 113BXes

# Summary and Future Plan

- RCT algorithm firmware is developed and tested successfully on the APd1 board with 113 BX TVs using three different physics TVs.
  - Latency: 269.466 ns (within the required budget)
- MET algorithm is implemented and successfully tested on the APd1 board using RCT generated output.
  - Latency: 163.902 ns
- Multi-board test using 4 APd1 board at UW lab is carried successfully with the following specification
  - RCT: 3 APd1 board
  - GCT: 1 APd1 board
  - Both the algorithms are running @360 MHz clock
- Complete chain of *MC == HLS Sim == HDL Sim == HW Results* from RCT (3 APd1) board to GCT (1 APd1) board is verified successfully.
- Several algorithms at GCT front are in development phase:
  - Taus
  - Jets
  - Sums

धन्यवाद

Thank You

# Acknowledgement


- Piyush Kumar and Bhawna Gomber acknowledges the support from IOE, University of Hyderabad through Grant Number UOH-IOE-RC2-21-006

CASEST

CENTRE FOR ADVANCED STUDIES IN
ELECTRONICS SCIENCE & TECHNOLOGY


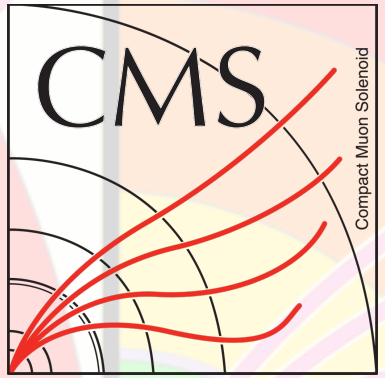

# BACKUP...

## Xilinx Stacked Silicon Interconnect (SSI) Technology

- The SSI technology integrate multiple Super Logic Region (SLR) components placed on a passive Silicon Interposer (fig 3).

- Each SLR contains the active circuitry common to most Xilinx FPGA (Field programmable gate array) devices. This circuitry includes large numbers of:
  - 6-input LUTs (Look-up tables)
  - Registers
  - I/O components
  - Gigabit Transceivers (GT)
  - Block memory
  - DSP blocks
  - Other blocks

- The device we are using for our synthesis and implementation is based on Xilinx SSI technology and support three SLRs.
  - Xilinx Virtex UltraScale+ xcvu9p flgc2104-1-e FPGA

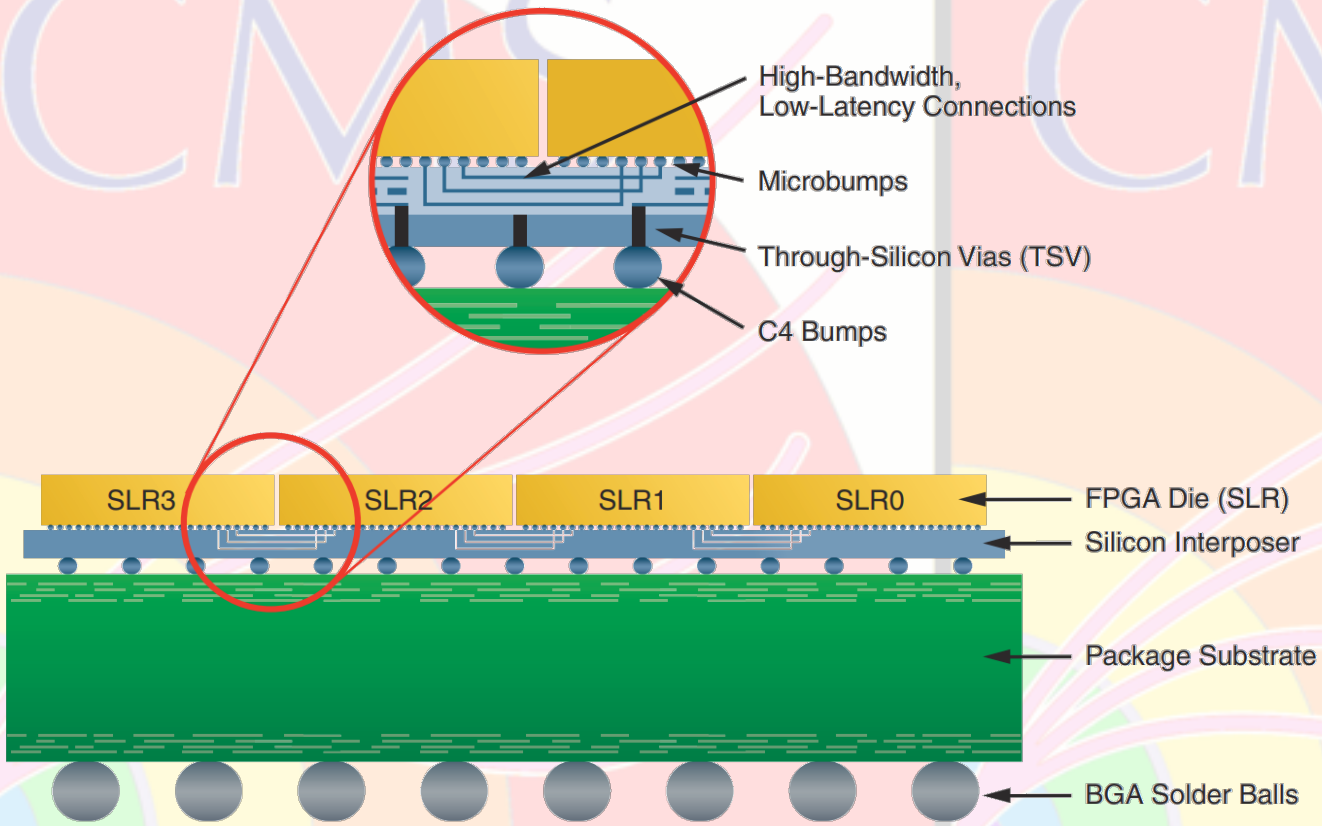


Fig 3: Xilinx FPGA Enabled by SSI Technology*

*: *UG872 Large FPGA Methodology Guide*

# Barrel Calorimeter Segmentation (New)

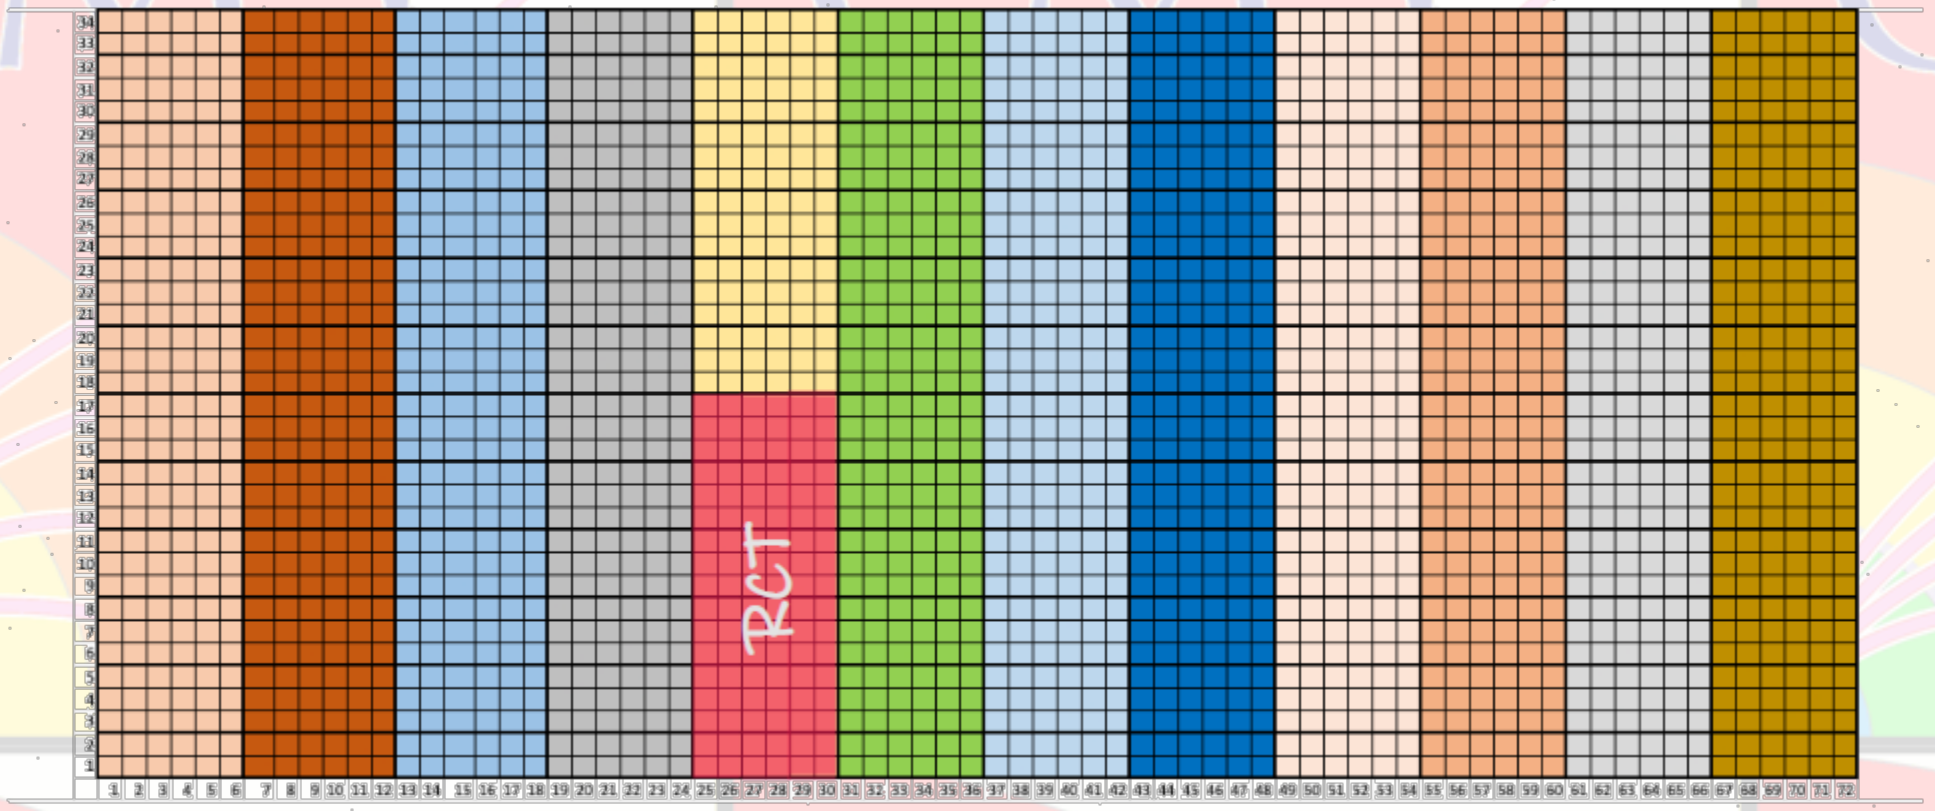


Fig 2: Barrel calorimeter segmentation (new)

# Physics bandwidth vs Algo clock @ 16G

| | |
|---|---|
| LHC BC Clock [MHz] | 40.08 |
| Word Bit Size | 66 |
| **Line Rate [Gbps]** | **16.00000** |
| Max Theoretical Words/Bx | 6.04851 |

| | TM1 | | | TM6 | | | TM18 | | |
|---|---|---|---|---|---|---|---|---|---|
| **Bx Frame Length (TM interval)** | **1** | **1** | **1** | **6** | **6** | **6** | **18** | **18** | **18** |
| **Words/Frame** | **4** | **5** | **6** | **24** | **30** | **36** | **72** | **90** | **108** |
| Equiv. Words/Bx | 4.00 | 5.00 | 6.00 | 4.00 | 5.00 | 6.00 | 4.00 | 5.00 | 6.00 |
| Equiv. Bits/Bx | 256 | 320 | 384 | 256 | 320 | 384 | 256 | 320 | 384 |
| **Data Rate [Gbps]** | 10.58 | 13.23 | 15.87 | 10.58 | 13.23 | 15.87 | 10.58 | 13.23 | 15.87 |
| **Filler Rate [Gbps]** | 5.42 | 2.77 | 0.13 | 5.42 | 2.77 | 0.13 | 5.42 | 2.77 | 0.13 |
| Average Filler Words/Bx | 2.05 | 1.05 | 0.05 | 2.05 | 1.05 | 0.05 | 2.05 | 1.05 | 0.05 |
| Average Filler Words/Orbit | 7300.89 | 3736.89 | 172.89 | 7300.89 | 3736.89 | 172.89 | 7300.89 | 3736.89 | 172.89 |
| Average Filler Words/Frame | 2.05 | 1.05 | 0.05 | 12.29 | 6.29 | 0.29 | 36.87 | 18.87 | 0.87 |
| **Payload Bits/Frame** | **256** | **320** | **384** | **1536** | **1920** | **2304** | **4608** | **5760** | **6912** |
| **Algo Clock @ 64b i/f [MHz]** | **160.32** | **200.4** | **240.48** | **160.32** | **200.4** | **240.48** | **160.32** | **200.4** | **240.48** |

# Physics bandwidth vs Algo clock @ 25G

| | |
|---|---|
| LHC BC Clock [MHz] | 40.08 |
| Word Bit Size | 66 |
| **Line Rate [Gbps]** | **25.78125** |
| Max Theoretical Words/Bx | 9.74613 |

| | TM1 | | | TM6 | | | TM18 | | |
|---|---|---|---|---|---|---|---|---|---|
| **Bx Frame Length (TM interval)** | **1** | **1** | **1** | **6** | **6** | **6** | **18** | **18** | **18** |
| **Words/Frame** | **7** | **8** | **9** | **42** | **48** | **54** | **126** | **144** | **162** |
| Equiv. Words/Bx | 7.00 | 8.00 | 9.00 | 7.00 | 8.00 | 9.00 | 7.00 | 8.00 | 9.00 |
| Equiv. Bits/Bx | 448 | 512 | 576 | 448 | 512 | 576 | 448 | 512 | 576 |
| **Data Rate [Gbps]** | 18.52 | 21.16 | 23.81 | 18.52 | 21.16 | 23.81 | 18.52 | 21.16 | 23.81 |
| **Filler Rate [Gbps]** | 7.26 | 4.62 | 1.97 | 7.26 | 4.62 | 1.97 | 7.26 | 4.62 | 1.97 |
| Average Filler Words/Bx | 2.75 | 1.75 | 0.75 | 2.75 | 1.75 | 0.75 | 2.75 | 1.75 | 0.75 |
| Average Filler Words/Orbit | 9787.22 | 6223.22 | 2659.22 | 9787.22 | 6223.22 | 2659.22 | 9787.22 | 6223.22 | 2659.22 |
| Average Filler Words/Frame | 2.75 | 1.75 | 0.75 | 16.48 | 10.48 | 4.48 | 49.43 | 31.43 | 13.43 |
| **Payload Bits/Frame** | **448** | **512** | **576** | **2688** | **3072** | **3456** | **8064** | **9216** | **10368** |
| **Algo Clock @ 64b i/f [MHz]** | **280.56** | **320.64** | **360.72** | **280.56** | **320.64** | **360.72** | **280.56** | **320.64** | **360.72** |

# Project hierarchy and floor planning

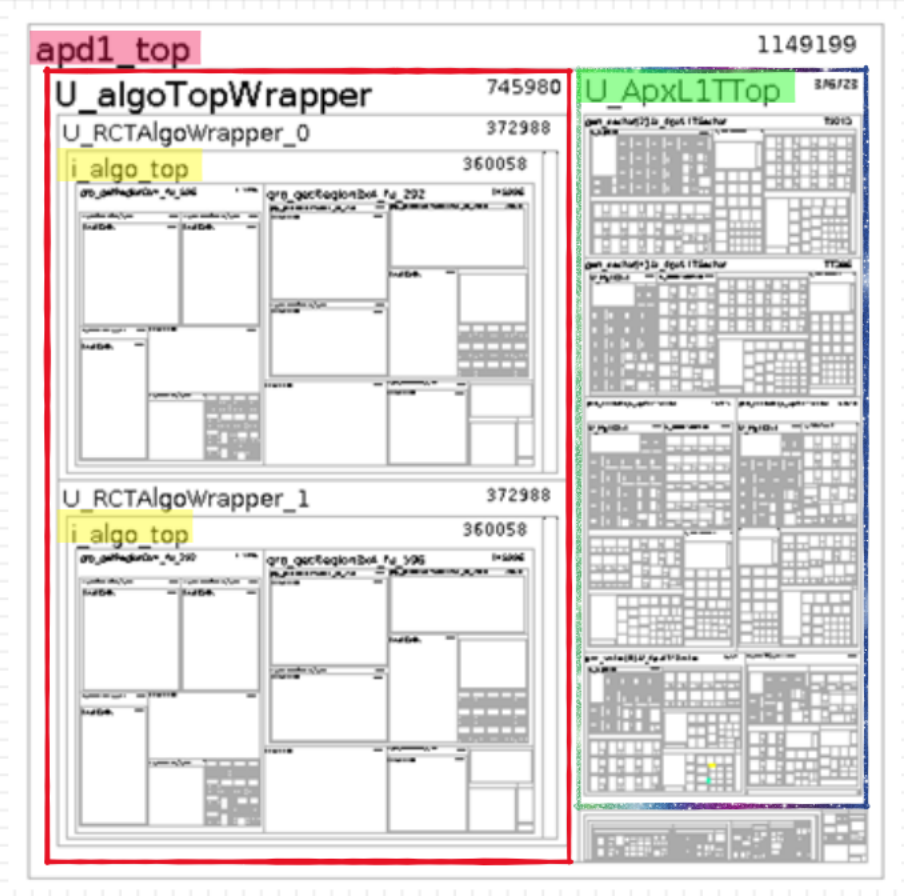


Fig 22: Project hierarchy in Vivado

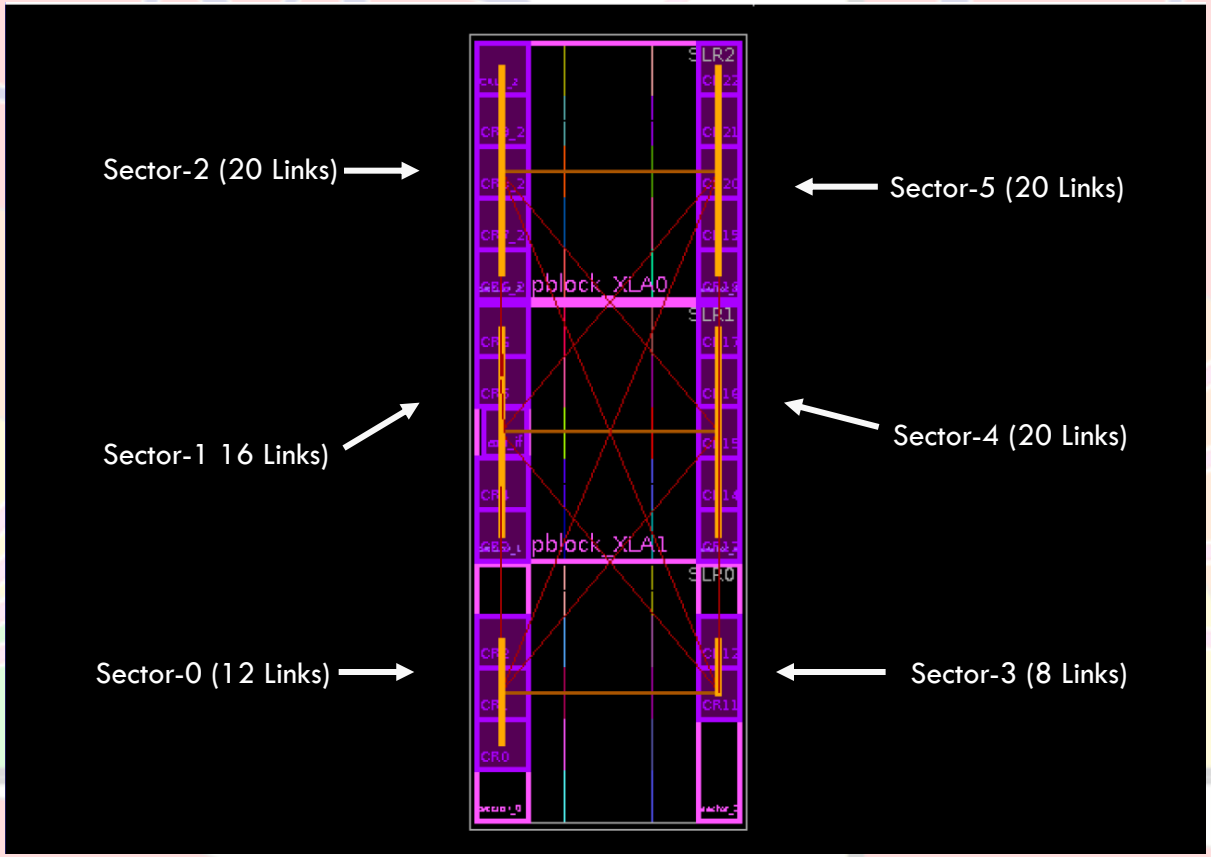


Fig 23: Project floor planning

# APx Firmware shell

Iridis – 64b66b-based optimized signaling method and firmware cores for CMS Trigger applications