

The CMS Level-1 Calorimeter Trigger for the HL-LHC

Piyush Kumar & Bhawna Gomber (on behalf of the CMS collaboration)

CASEST, School of Physics, University of Hyderabad, Hyderabad, Telangana, India

CALOR 2022 - 19th International Conference on Calorimetry in Particle Physics University of Sussex, UK

High-Luminosity LHC (HL-LHC)



- Aim: to deliver a much larger dataset for physics to the LHC experiments

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




Fig 1: HL-LHC timeline

Summary of CMS HL-LHC Upgrades





Fig 2: HL-LHC CMS detector upgrade highlights

L1 trigger principle

- At design parameters the LHC produces:
 - $\sim 10^9$ collisions/second in CMS detectors.
 - each event is ~ 1 MB.
- 10^9 collisions/s \times 1 Mbyte/collision = 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.
- A **trigger** is designed to reject the uninteresting events and keep the interesting ones.

Modern large-scale experiments are **really BIG**



i.e. LHC experiments (ATLAS/CMS)

- ▶ **~ 100 M channels**
- ▶ **~ 1 -2 MB** of RAW data per measurement

... and **really FAST**

- ▶ **~ 40 MHz** measurement rate (every 25 ns - @ the LHC)



Data volume is a **key issue** in modern large-scale experiments

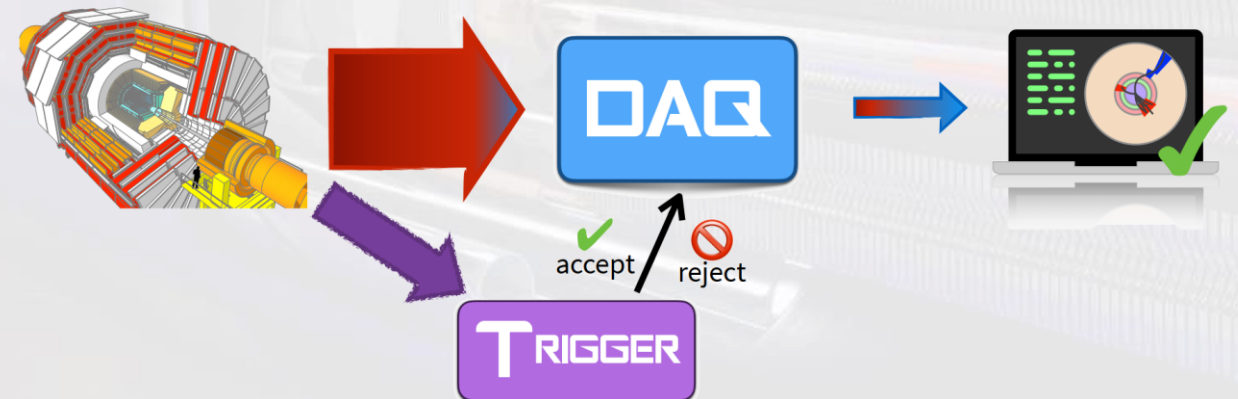


Fig 3: Trigger system

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
- RCT geometry for the FPGA processing:
 - $17\eta \times 4\phi$ of the barrel (36 FPGA)
 - $17\eta \times 6\phi$ of the barrel (24 FPGA)
- GCT geometry for the FPGA (3 XCVU13P) processing:
 - 12 unique RCT ($17\eta \times 4\phi$) + 4 neighbours
 - 8 unique RCT ($17\eta \times 6\phi$) + 4 neighbours

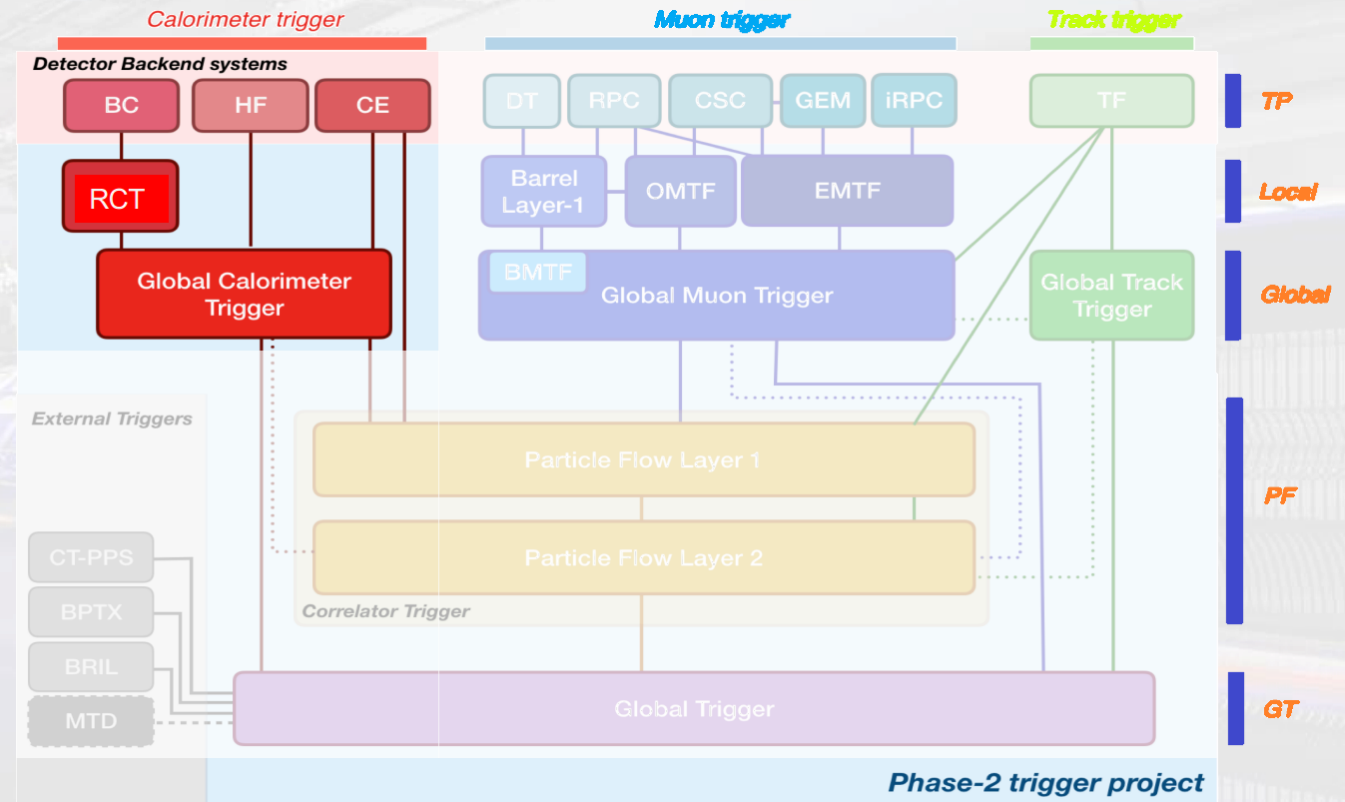


Fig 4: Calorimeter trigger architecture

Barrel Calorimeter Segmentation (TDR)




Fig 5: Barrel calorimeter segmentation

Calorimeter Trigger Development

- The trigger algorithms are implemented by using Xilinx Vivado-HLS (high level synthesis) tool
 - Rapid prototyping
 - Codes are written in C++
 - HLS synthesizes the code to generate the RTL and
 - Provide an **early estimate of latency and resource utilization**
 - Increased **ease of collaboration** and code sharing for algorithm design
- Downstream:
 - Integration of the algo with the firmware shell (**orange box**) that provides
 - MGT link instantiation
 - TCDS connectivity
 - DAQ support,
 - and an AXI interface to the controlling system
 - Uses HDL wrapper for integration (**magenta box**)

```

== Performance Estimates
=====
+ Timing (ns):
  * Summary:
  +-----+-----+-----+-----+
  | Clock | Target | Estimated | Uncertainty |
  +-----+-----+-----+-----+
  | ap_clk | 4.17 | 2.917 | 1.25 |
  +-----+-----+-----+-----+

+ Latency (clock cycles):
  * Summary:
  +-----+-----+-----+-----+
  | Latency | Interval | Pipeline |
  | min | max | min | max | Type |
  +-----+-----+-----+-----+
  | 32 | 32 | 6 | 6 | function |
  +-----+-----+-----+-----+
    
```

```

=====
== Utilization Estimates
=====
* Summary:
+-----+-----+-----+-----+-----+
| Name | BRAM_18K | DSP48E | FF | LUT | URAM |
+-----+-----+-----+-----+-----+
| DSP | - | - | - | - | - |
| Expression | - | - | 0 | 4 | - |
| FIFO | - | - | - | - | - |
| Instance | - | - | 49827 | 78752 | - |
| Memory | - | - | - | - | - |
| Multiplexer | - | - | - | 56 | - |
| Register | 0 | - | 3360 | 32 | - |
+-----+-----+-----+-----+-----+
| Total | 0 | 0 | 53187 | 78844 | 0 |
+-----+-----+-----+-----+-----+
| Available SLR | 1440 | 2280 | 788160 | 394080 | 320 |
+-----+-----+-----+-----+-----+
| Utilization SLR (%) | 0 | 0 | 6 | 20 | 0 |
+-----+-----+-----+-----+-----+
| Available | 4320 | 6840 | 2364480 | 1182240 | 960 |
+-----+-----+-----+-----+-----+
| Utilization (%) | 0 | 0 | 2 | 6 | 0 |
+-----+-----+-----+-----+-----+
    
```

Fig 6: Vivado-HLS performance estimates of trigger algorithm




Fig 7: Trigger algo device implementation

RCT Algorithm

Regional Calorimeter Trigger (RCT) creates e/γ clusters and towers and sends them to Global Calorimeter Trigger (GCT)

- Algorithm is divided in three part
 - RCT8x4:**
 - Implemented in SLR1
 - Processes the $8\eta \times 4\phi$ RCT regions
 - only ECAL.
 - RCT9x4**
 - implemented in SLR2
 - processes the $9\eta \times 4\phi$ RCT regions
 - ECAL
 - $16\eta \times 4\phi$ HCAL data.
 - RCTSUM**
 - implemented in SLR2
 - combines both the algorithm and sends the output to the GCT.

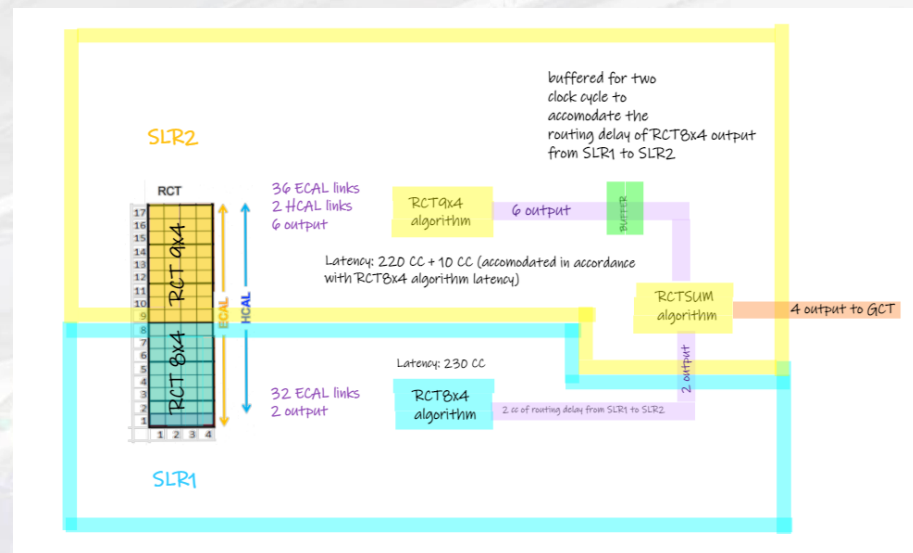


Fig 8: RCT algorithm organisation and dataflow




Fig 9: e/gamma cluster making in RCT algorithm

Timing (ns)				
Summary				
Clock	Target	Estimated	Uncertainty	
ap_clk	4.17	3.491	1.25	
Latency (clock cycles)				
Summary				
Latency	Interval			
min	max	min	max	Type
230	230	6	6	function

Summary					
Name	BRAM_18K	DSP48E	FF	LUT	URAM
DSP	-	-	-	-	-
Expression	-	-	0	24202	-
FIFO	-	-	-	-	-
Instance	8	0	303544	464948	-
Memory	-	-	-	-	-
Multiplexer	-	-	-	16292	-
Register	30	-	23821	1813	-
Total	38	0	327365	507255	0
Available	4320	6840	2364480	1182240	960
Available SLR	1440	2280	788160	394080	320
Utilization (%)	~0	0	13	42	0
Utilization SLR (%)	2	0	41	128	0

Fig 10: RCT algorithm HLS results

RCT slice test

Implemented in 2 SLRs of XCVU9P FPGA

- The bitstream is generated and successfully tested on the APd1 board.
- The implementation is scalable for the region of $17\eta \times 6\phi$ (can use 3 SLRs).



Fig 11: Prototyped APd1 board

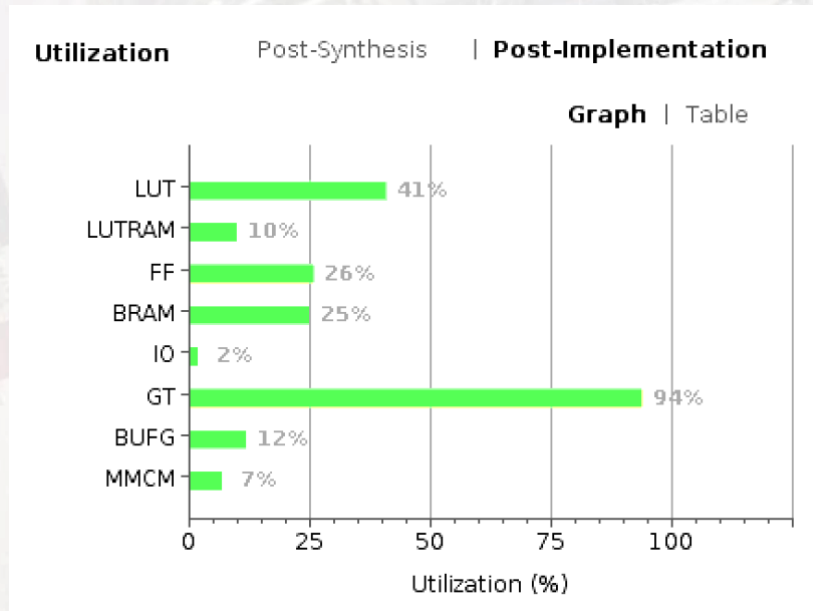


Fig 12: Timing and utilization summary

Timing	Setup	Hold	Pulse Width
Worst Negative Slack (WNS):	0.004 ns		
Total Negative Slack (TNS):	0 ns		
Number of Failing Endpoints:	0		
Total Number of Endpoints:	1425432		

[Implemented Timing Report](#)




Fig 13: RCT Device Implementation

GCT Slice test (RCT to GCT)

- Tested on a single card:
 - Replicate the 4 RCT output links x5 (20 input) ~ GCT processing 5 RCT cards
- This GCT algorithm (stitches the E/gamma between the RCT cards and produces the final towers) is synthesized in HLS.
- Implemented in
 - XCVU9P FPGA
 - Clock: 240 MHz
 - Link bandwidth: 16 gbps

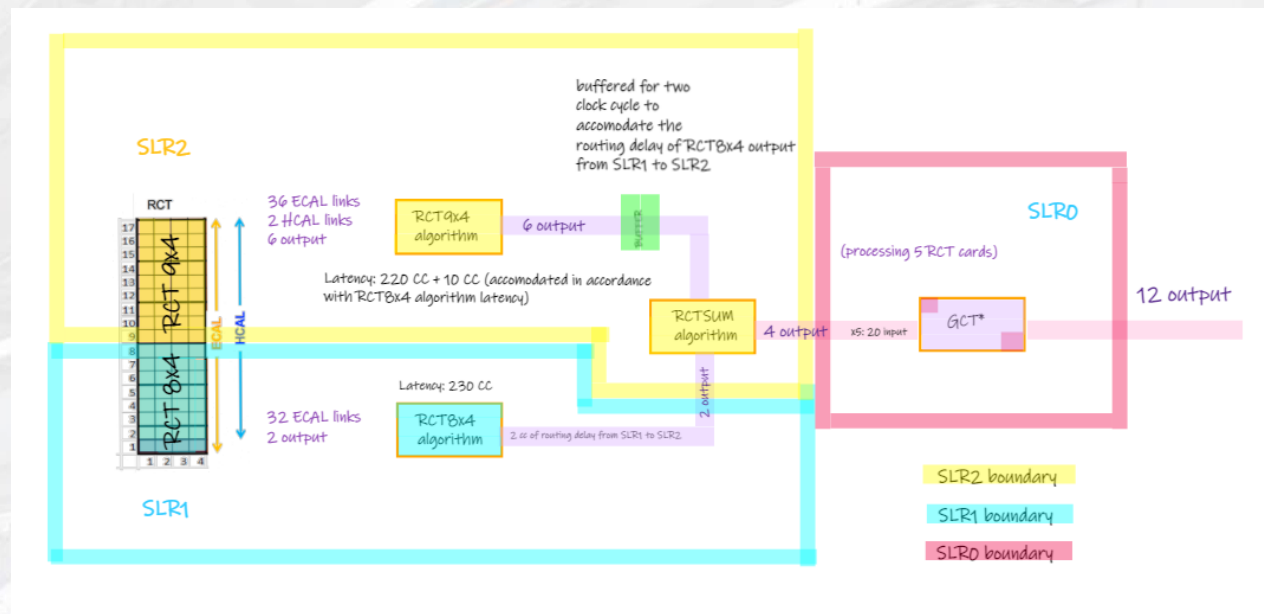


Fig 14: RCTTDR and GCT algorithm implementation in three SLR

Performance Estimates

Timing (ns)

Summary

Clock	Target	Estimated	Uncertainty
ap_clk	4.17	2.909	1.25

Latency (clock cycles)

Summary

Latency		Interval		Type
min	max	min	max	
120	120	6	6	function

Utilization Estimates

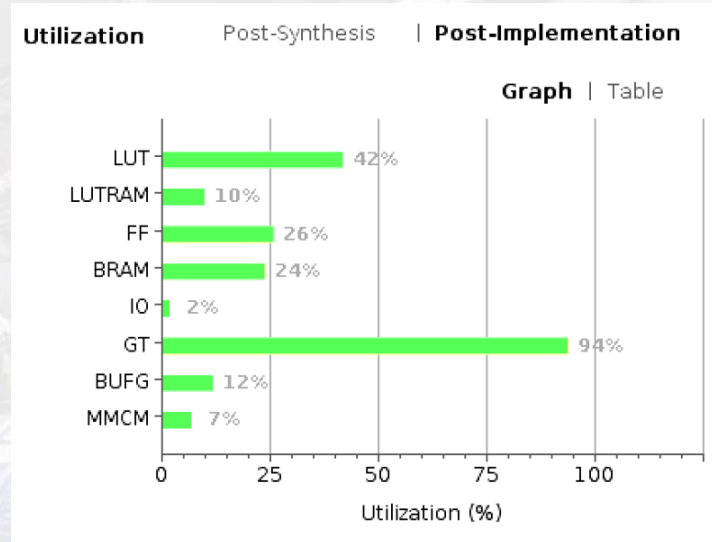
Summary

Name	BRAM_18K	DSP48E	FF	LUT	URAM
DSP	-	-	-	-	-
Expression	-	-	0	1444	-
FIFO	-	-	-	-	-
Instance	-	-	27703	146555	-
Memory	-	-	-	-	-
Multiplexer	-	-	-	56	-
Register	0	-	82036	36864	-
Total	0	0	109739	184919	0
Available	4320	6840	2364480	1182240	960
Available SLR	1440	2280	788160	394080	320
Utilization (%)	0	0	4	15	0
Utilization SLR (%)	0	0	13	46	0

Fig 15: GCT algorithm HLS results

GCT Slice test

- The bitstream is generated and the project passes the timing constraints.
- Following are the algorithms device placement:
 - RCT8x4: SLR1
 - RCT9x4: SLR2
 - RCTSUM: SLR2
 - GCT: SLR0
- Post implementation device utilization is within the boundary.
- The bitstream is successfully tested on the APd1 board.



Timing	Setup	Hold	Pulse Width
Worst Negative Slack (WNS):	0.019 ns		
Total Negative Slack (TNS):	0 ns		
Number of Failing Endpoints:	0		
Total Number of Endpoints:	1434128		

[Implemented Timing Report](#)

Fig 16: Utilization and timing summary (setup)




Fig 17: GCT device implementation

$$F_{\max} = 1/(4.167-0.019) \sim 241 \text{ MHz}$$

Summary

- RCT algorithm is developed in Vivado-HLS and implemented and tested on APd1 board (based on Xilinx XCVU9P FPGA)
 - Clock speed: 240 MHz
 - Link bandwidth: 16 Gbps
 - Project passes the timing constraints
 - Tested successfully
- GCT algorithm is developed in Vivado-HLS and implemented together with the RCT algorithm and in tested on APd1 board (based on Xilinx XCVU9P FPGA)
 - Clock speed: 240 MHz
 - Link bandwidth: 16 Gbps
 - Project passes the timing constraints
 - Tested successfully
- HGAL data de-multiplexer work is going on
- GCT to Correlator Time-multiplexer work is going on
- Several GCT algorithms are in pipeline

Acknowledgement

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


धन्यवाद

Thank You



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)

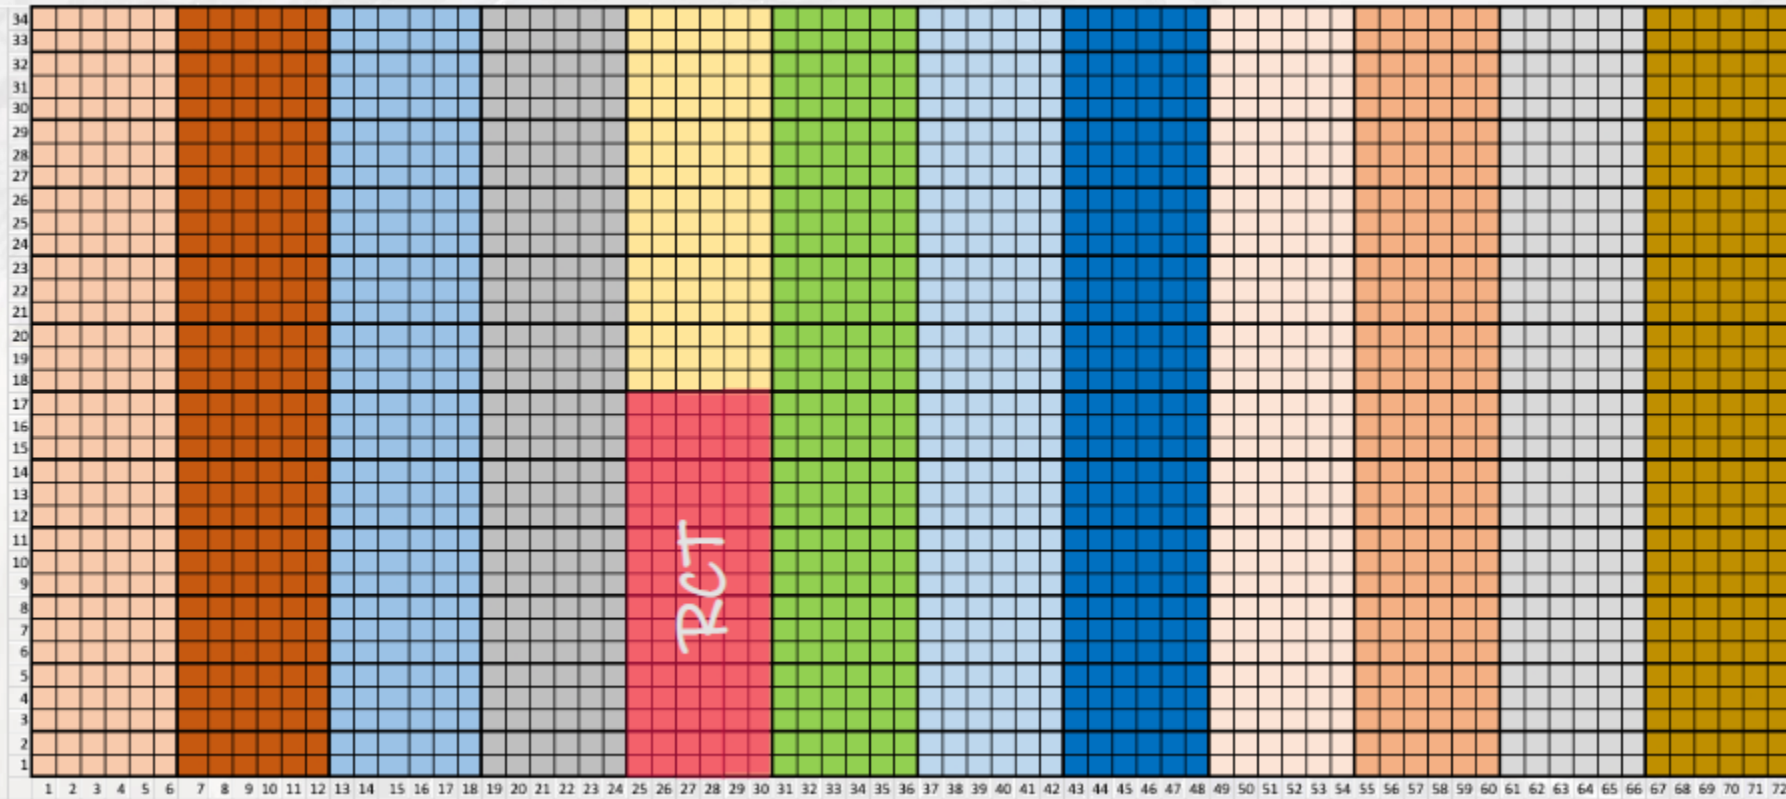


Fig 2: Barrel calorimeter segmentation (new)



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

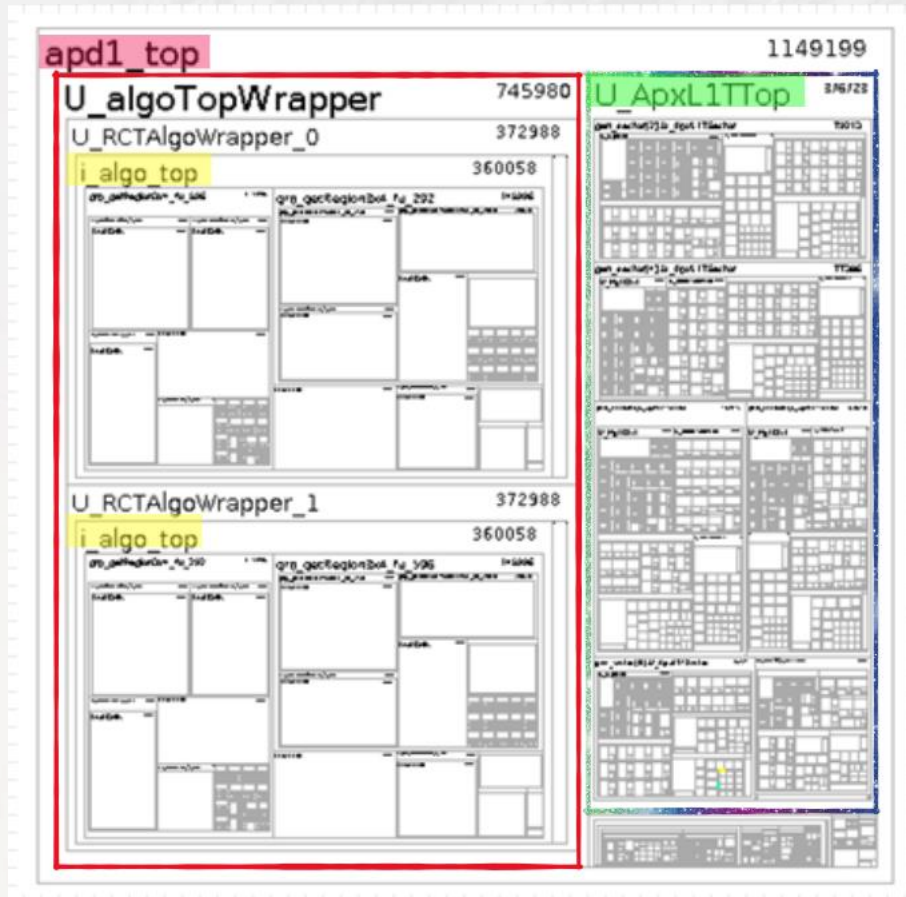


Fig 22: Project hierarchy in Vivado






Fig 23: Project floor planning

HGCAL to GCT slice test

HGCAL towers configuration: $72\phi : 20\eta$



20 → 13 (combine towers at high η)



HGCAL-

Barrel

HGCAL+


HGCAL sends information at TMI18 and 25 GBps:
 $18 * 9 \text{ clocks/bx} = 162 \text{ 64b words}$



30 words are reserved for towers 16b each tower:
 $30 * 4 = 120 \text{ towers, } 131 \text{ for clusters}$

12 fibers: the whole HGCAL (+ or -)
 $12 * 120 = 1440 \text{ towers } (72\phi : 20\eta)$

Towers at high η need to be combined:
 we make 13 towers in η



HGCAL to GCT slice test

- We prepared a GCT code that receives 12 fibers with 162 words
- Selects 30 words with towers
- Unpack towers and combine them if needed
- Array 72x13 is stored
- Next step will be to run jets/taus/sums on this information
- Can use the code for tests when hardware configuration will allow
- Can make some tests with existing configuration emulating HGCAL output to GCT

```

+ Timing (ns):
  * Summary:
  +-----+-----+-----+-----+
  | Clock | Target| Estimated| Uncertainty|
  +-----+-----+-----+-----+
  | ap_clk |  2.78|   2.490|       0.83|
  +-----+-----+-----+-----+
    
```

```

+ Latency (clock cycles):
  * Summary:
  +-----+-----+-----+-----+
  | Latency | Interval | Pipeline |
  | min | max | min | max | Type |
  +-----+-----+-----+-----+
  |  12|  12|   1|   1| function |
  +-----+-----+-----+-----+
    
```

```

* Summary:
+-----+-----+-----+-----+-----+
| Name | BRAM_18K| DSP48E| FF | LUT | URAM|
+-----+-----+-----+-----+-----+
| DSP | - | - | - | - | - |
| Expression | - | - | 0 | 265 | - |
| FIFO | - | - | - | - | - |
| Instance | - | - | 6634 | 10288 | - |
| Memory | - | - | - | - | - |
| Multiplexer | - | - | - | 762 | - |
| Register | 0 | - | 2624 | 32 | - |
+-----+-----+-----+-----+-----+
| Total | 0 | 0 | 9258 | 11347 | 0 |
+-----+-----+-----+-----+-----+
| Available SLR | 1440 | 2280 | 788160 | 394080 | 320 |
+-----+-----+-----+-----+-----+
| Utilization SLR (%) | 0 | 0 | 1 | 2 | 0 |
+-----+-----+-----+-----+-----+
| Available | 4320 | 6840 | 2364480 | 1182240 | 960 |
+-----+-----+-----+-----+-----+
| Utilization (%) | 0 | 0 | ~0 | ~0 | 0 |
+-----+-----+-----+-----+-----+
    
```

Fig 27: GCT HGCAL receiver synthesis results