B02 -Trigger algorithm development, performance, and demonstration

Nhan Tran

HL LHC CMS CD-1 Review

October 22-24, 2019

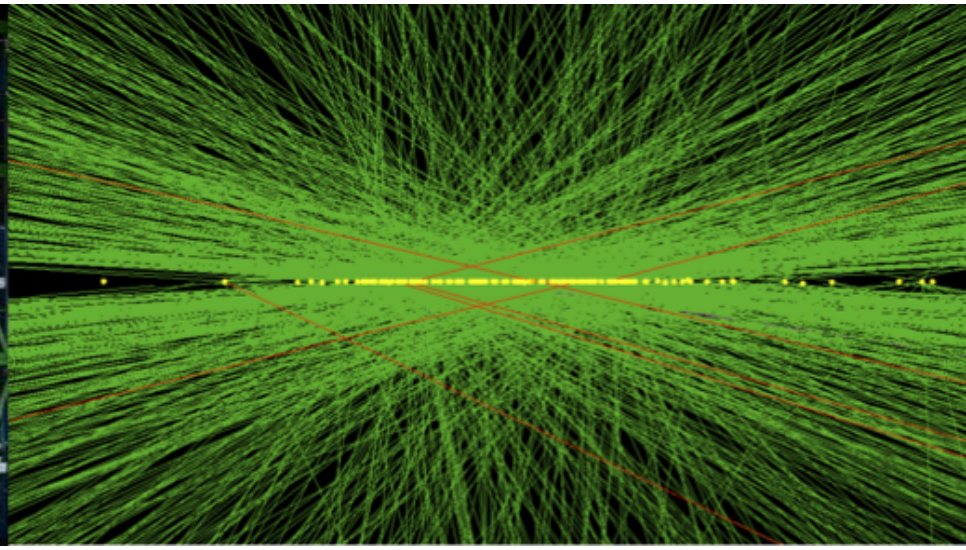

Outline

Technical design

- Algorithms overview and physics performance
- Algorithm design
- Firmware demonstration

Management aspects

- Risk, quality assurance
- Milestones and progress
- *Cost and schedule (see talk later by Jeff)*

Brief Biological Sketch

Nhan Tran

Wilson Fellow (Fermilab)

- L3 Manager: Correlator trigger
- Development of Particle Flow and PUPPI in L1 trigger
- Lead on hls4ml: high level synthesis for machine learning

Postdoc (Fermilab)

- Track trigger ASIC development and testing for Vertically Integrated Pattern Recognition Associative Memory (VIPRAM)
- Development of PUPPI algorithm

TECHNICAL DESIGN

Trigger Scope Overview

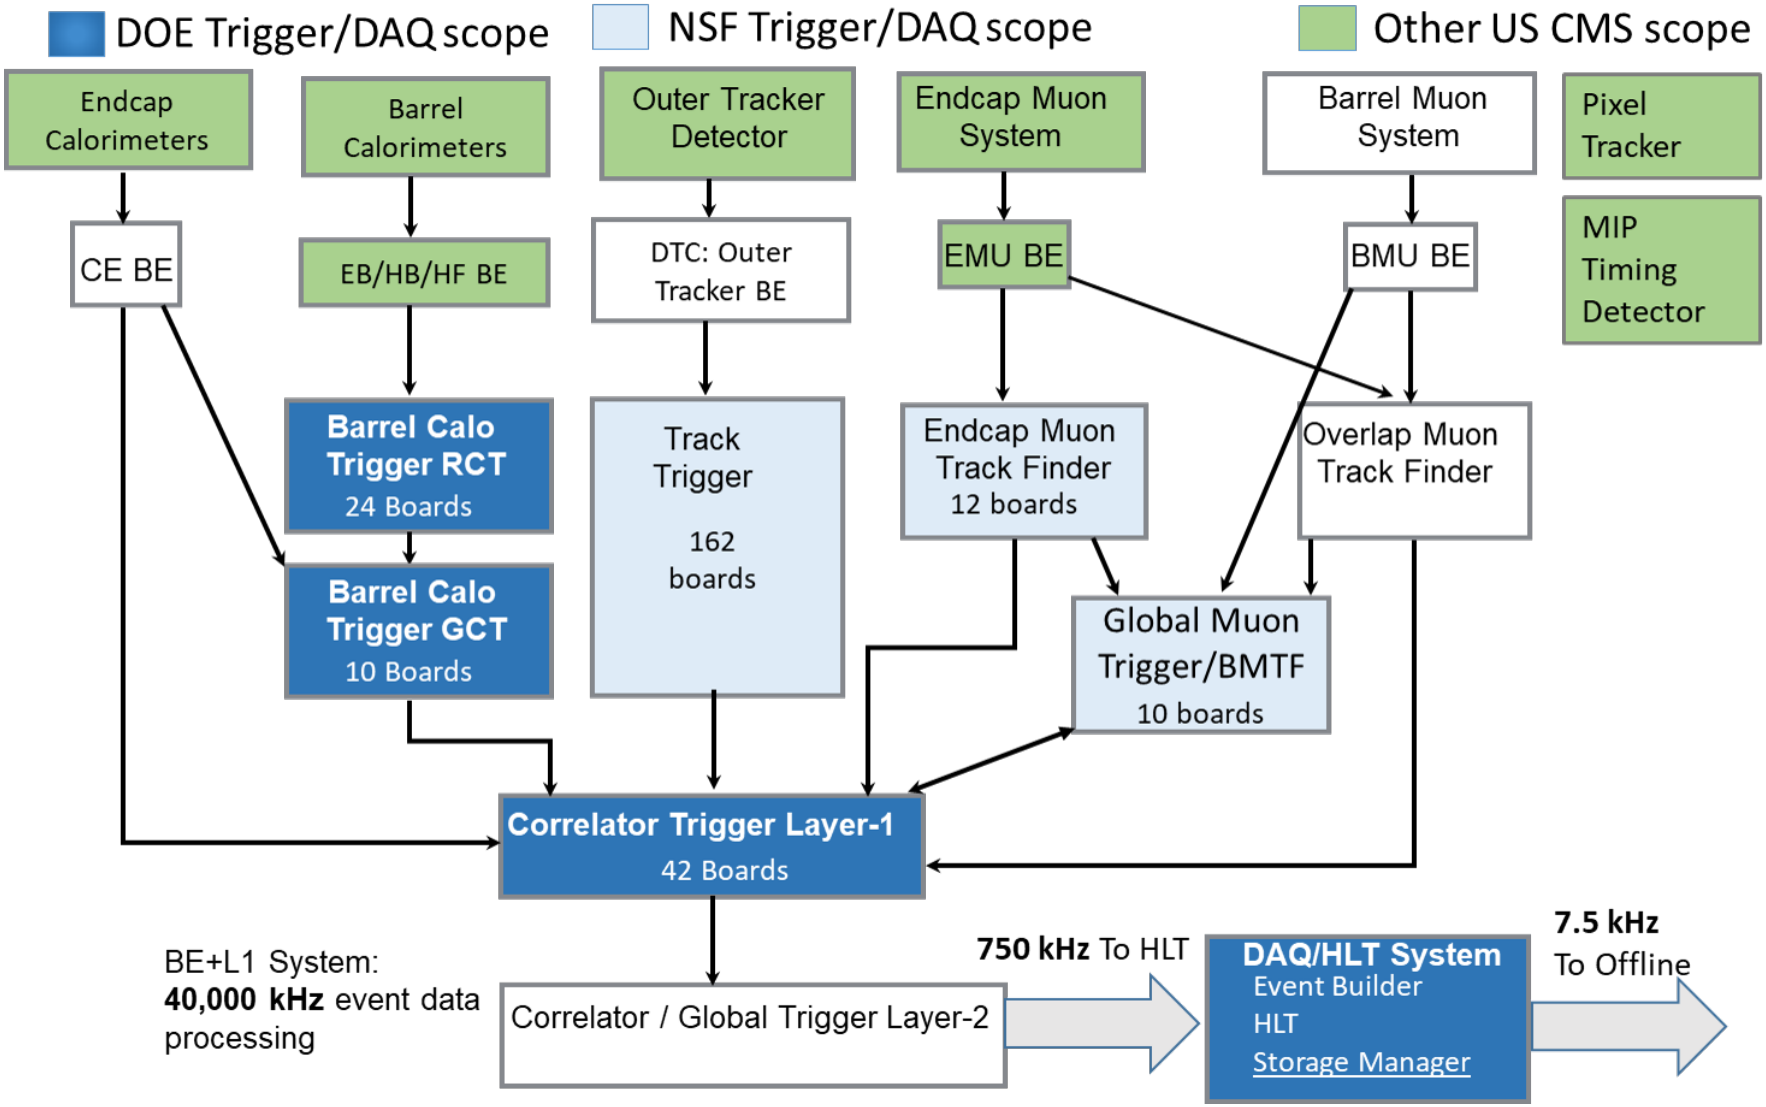

Design Considerations

Maintain performant trigger under high luminosity conditions

- Upgrade L1 trigger accept rate: 750 kHz
- Upgrade L1 trigger total latency: 12.5 μs

Detector/Trigger Upgrades

- Tracking trigger for tracks with pT > 2 GeV
- New high granularity endcap calorimeter
- Full crystal readout of barrel ECal
- New muon detectors for improved high η coverage and higher granularity readout

DOE trigger scope

402.6.3: Calorimeter (regional barrel and forward calorimeter, global)

402.6.5: Correlator trigger (combining muon, calorimeter, tracker inputs)

General Algorithm Strategy

Deliver a suite of algorithms which cover both **robustness** and **optimized physics performance**

Single system triggers*

Robust, simpler algorithms

Global Calorimeter Trigger objects

Track-only Trigger objects

** muon system only triggers in NSF scope*

Multi-system optimized reconstruction

More complex, performant algorithms

Track + muon + calorimeter correlated (particle flow and PUPPI) trigger objects

Trigger overview

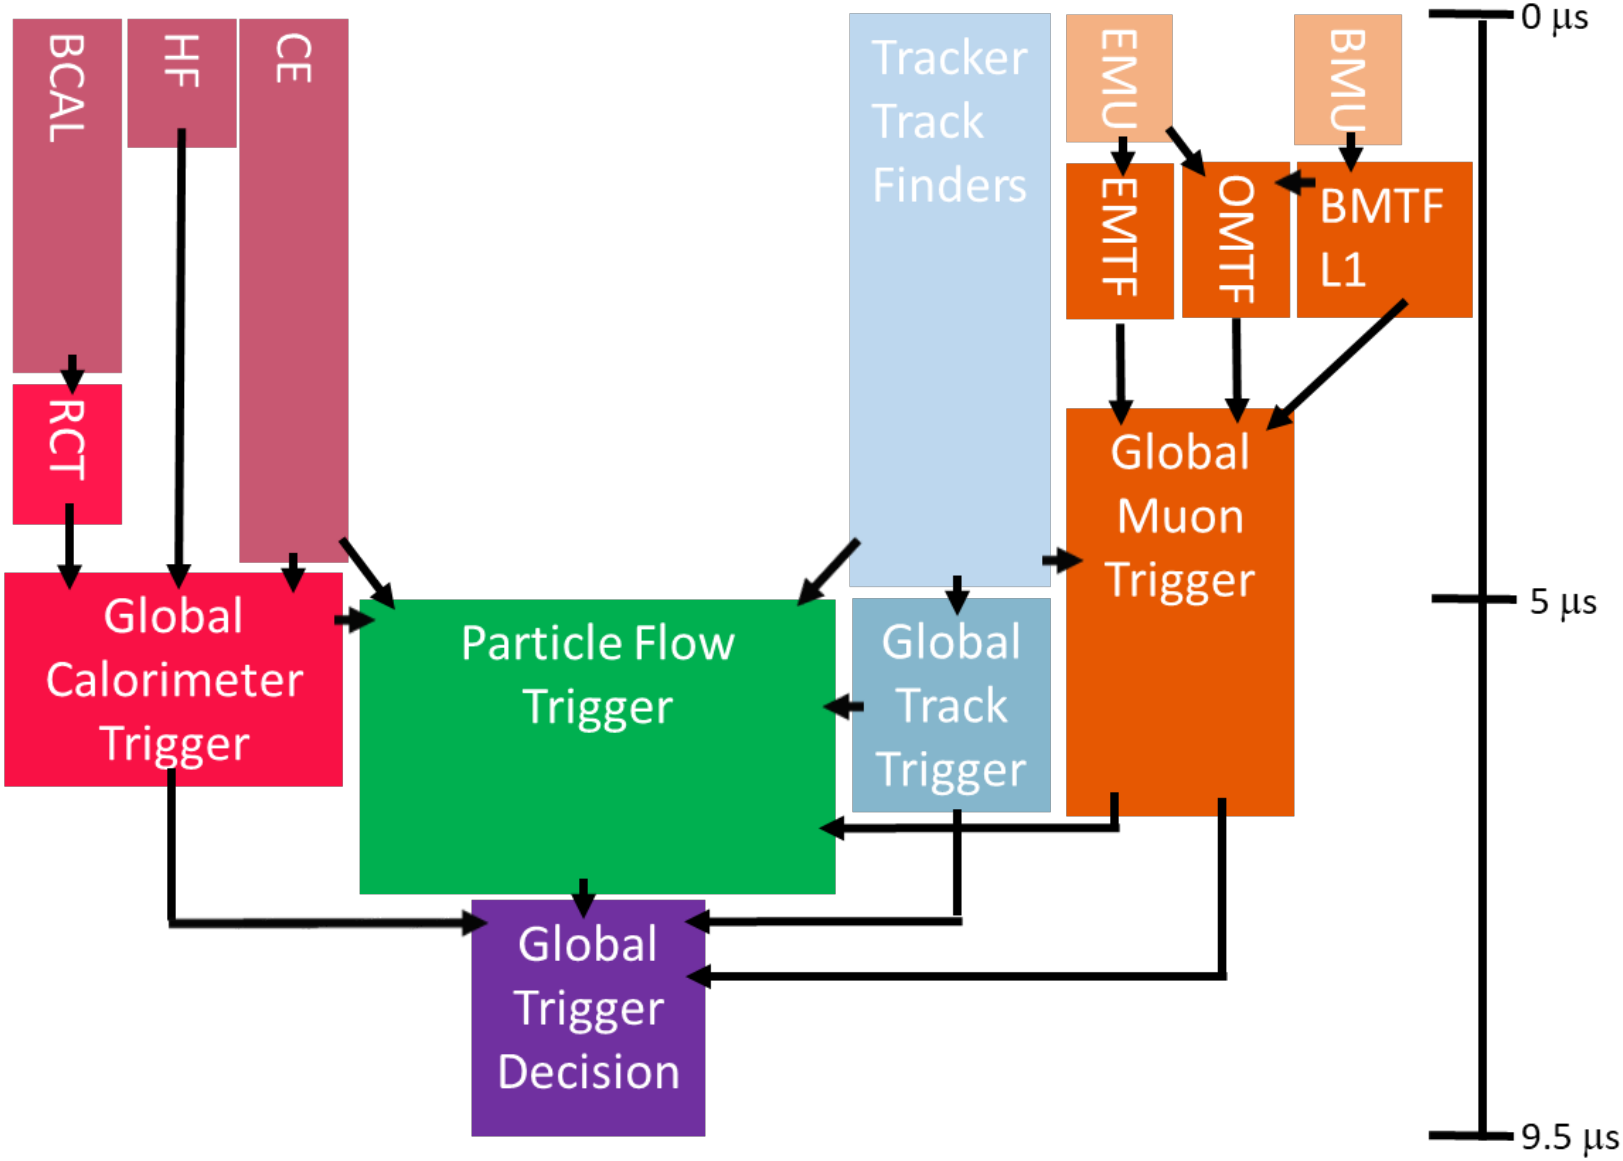

Trigger overview

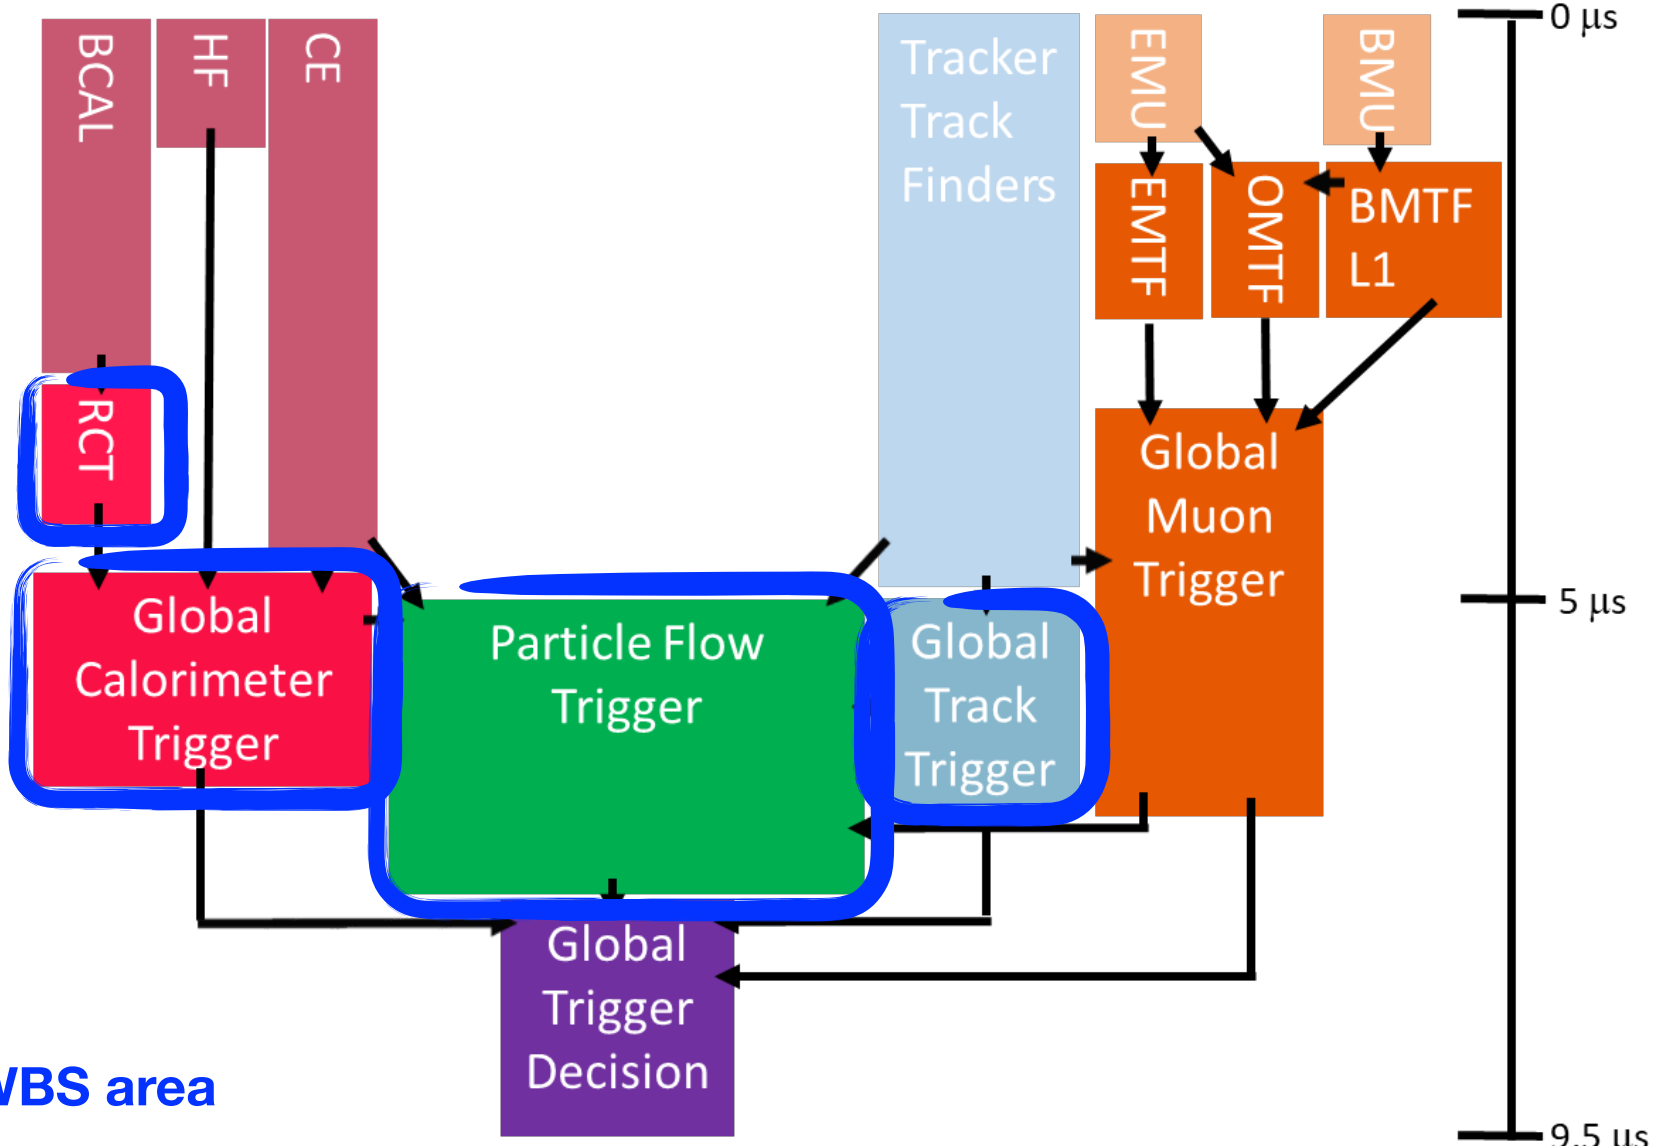

This WBS area

Deliverables

Devise algorithms and study physics performance for the correlator (Particle Flow) and calorimeter triggers

Implement and integrate algorithms into firmware

Within latency, bandwidth, and resource requirements

From *robust*, single-system algorithms to *optimized* multi-system algorithms

Calorimeter and Track-only objects

Particle Flow with PUPPI

Functional algorithm diagram

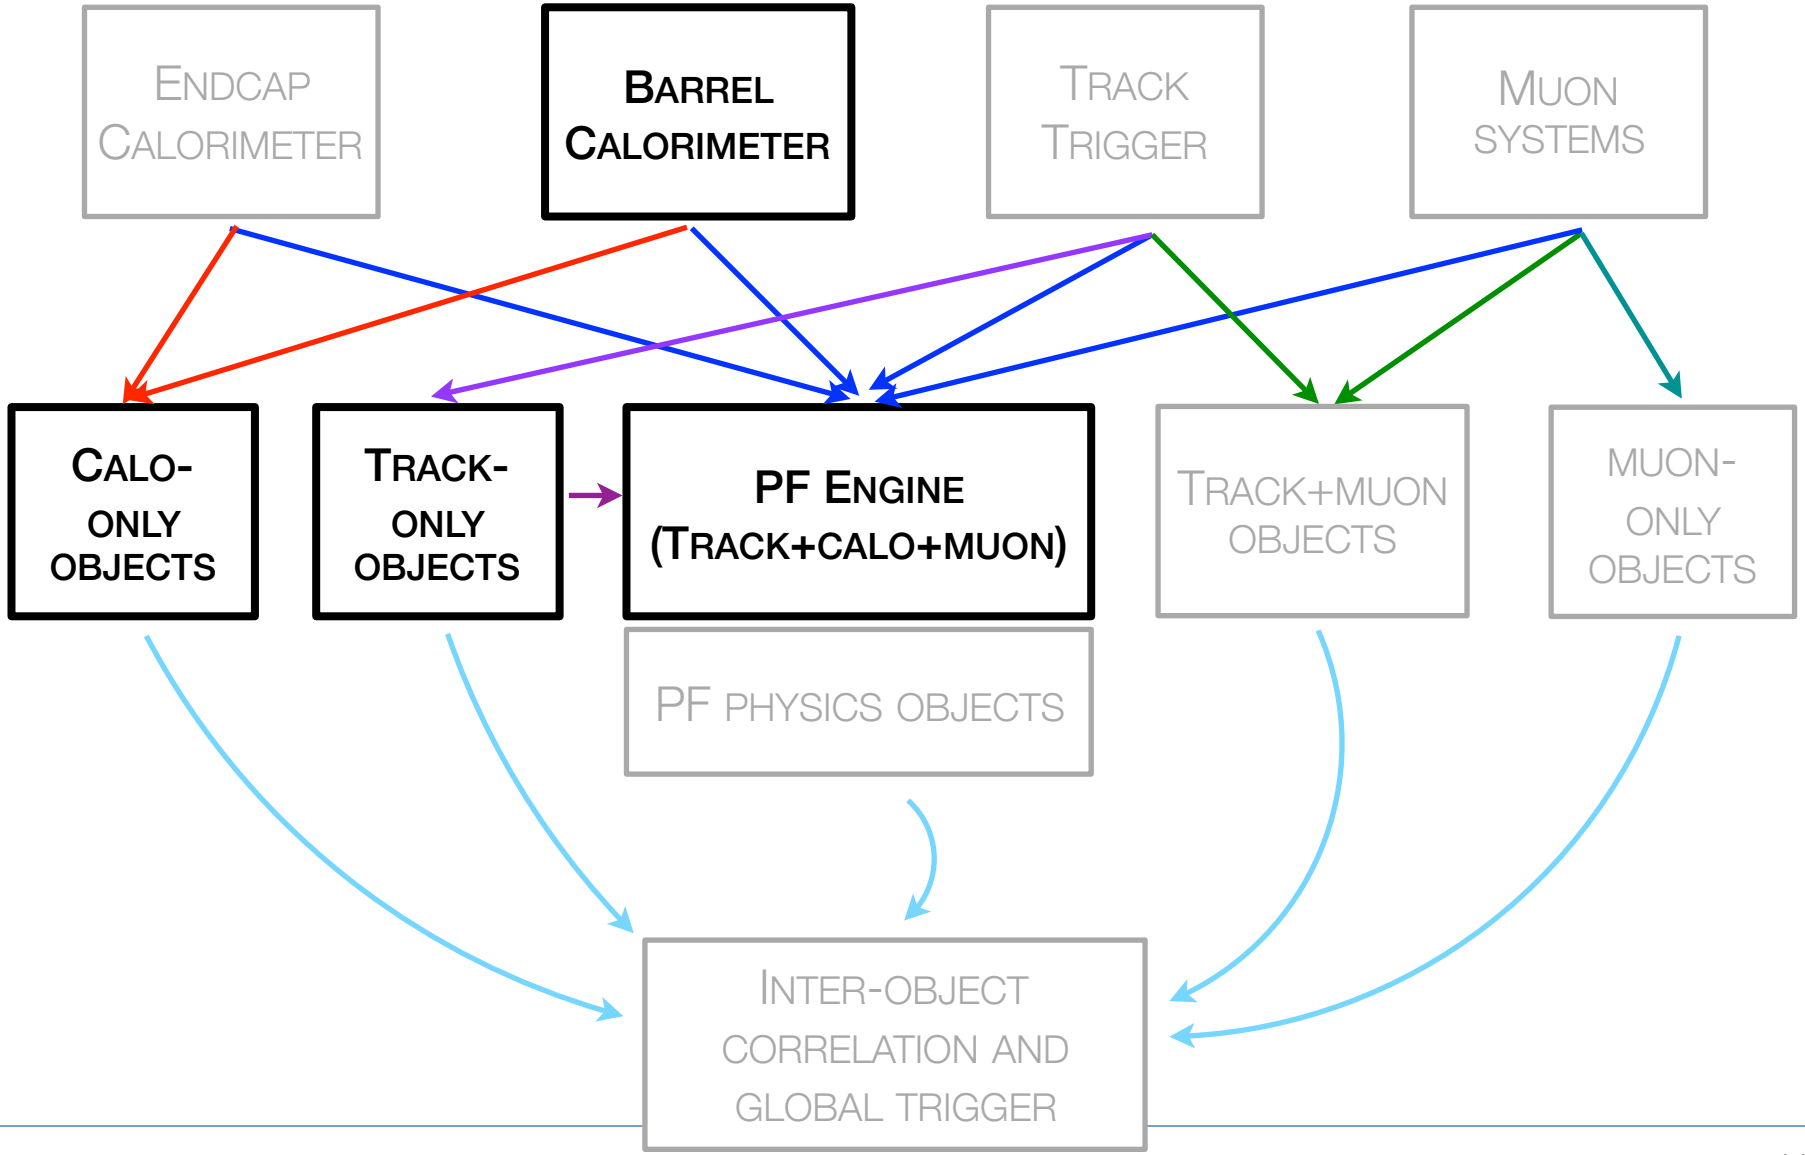

Functional algorithm diagram

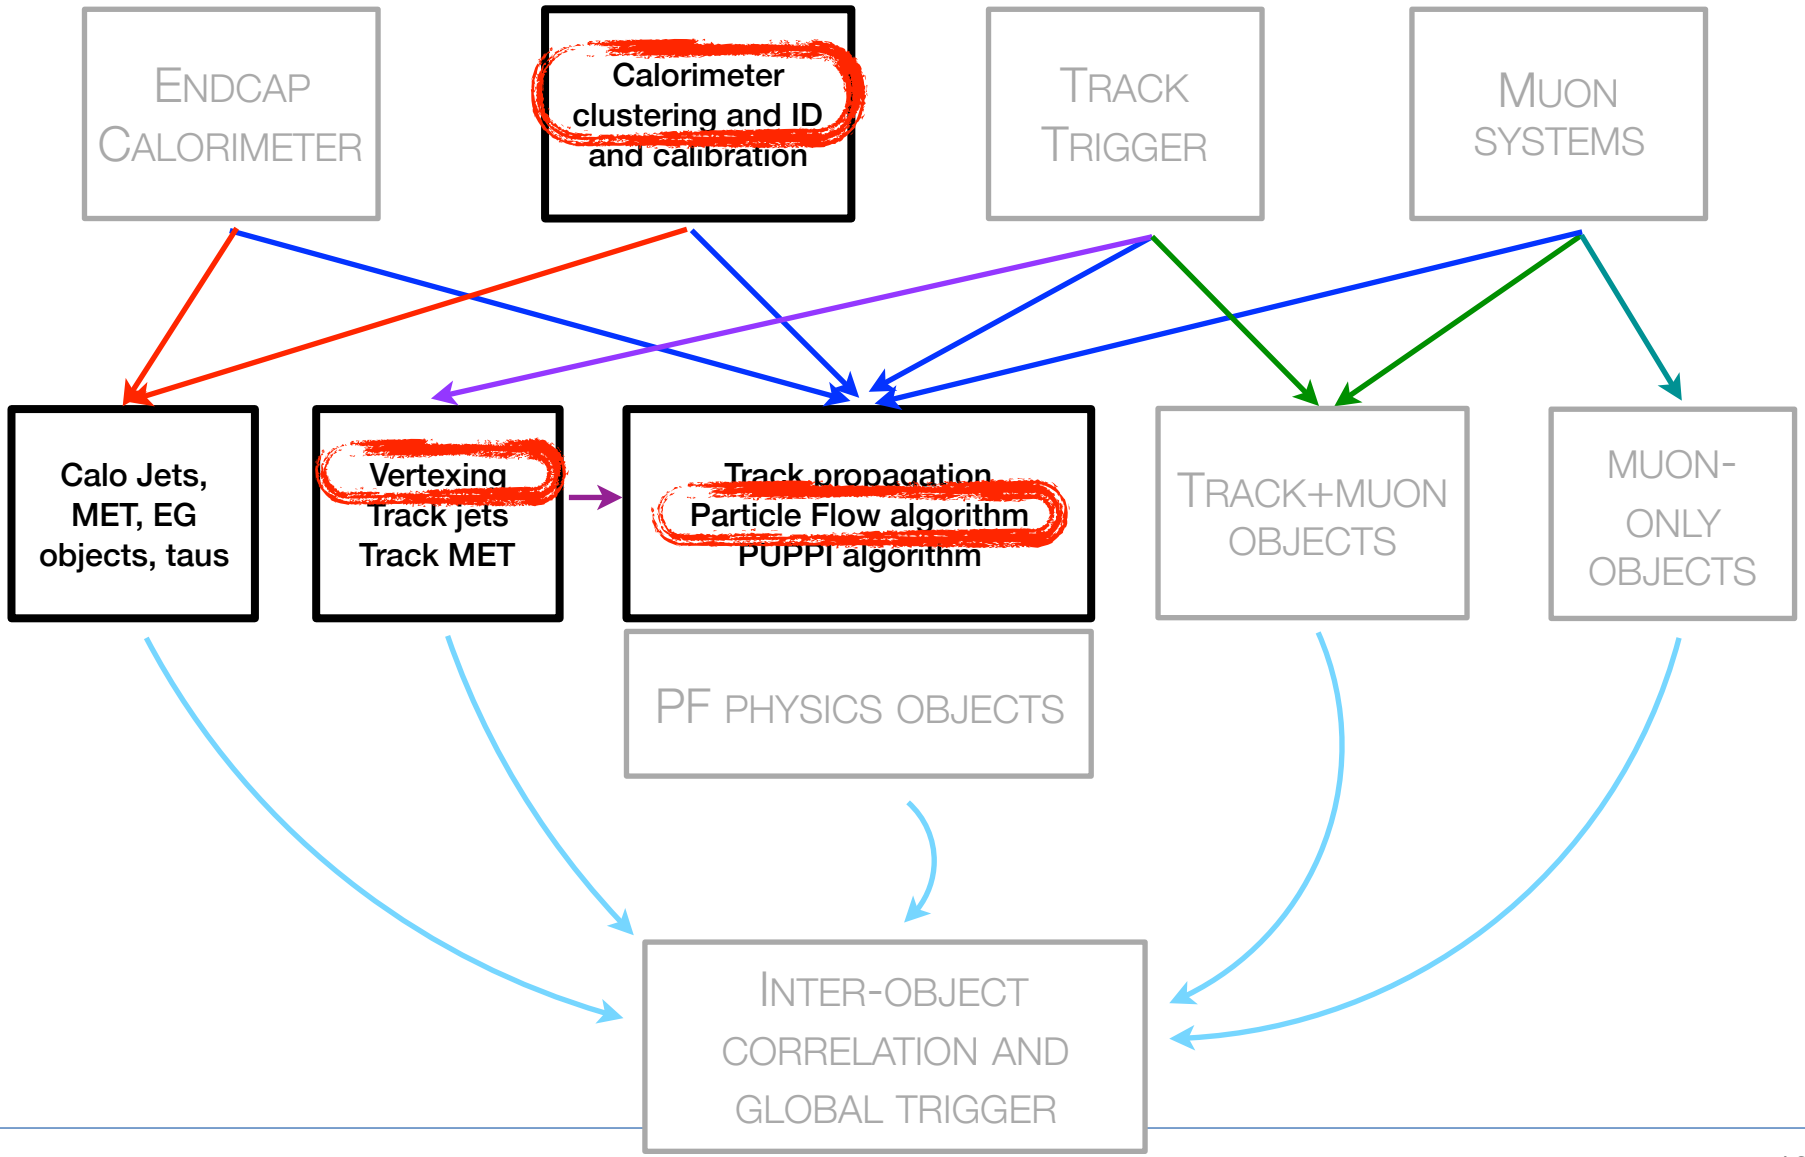

Algorithm development

Steps of algorithm development path:

Understand physics case

Define algorithm and interfaces (inputs & downstream)

Develop firmware, estimate latency and resources

Integrate within full trigger architecture

Demonstrate full implementation feasibility

n.b. difficult to present all steps for all algorithms

1) **Calorimeter clustering** example

2) **Vertexing** example

2) **Particle flow** example larger focus, most complex in terms of algorithm size and input interfaces

A note on algorithm development

FPGA development of algorithms in languages like VHDL or Verilog (RTL) have long development cycles and require a lot of engineering support

New tools: **HLS, high level synthesis**

C-level programming with specialized preprocessor directives which synthesizes optimized firmware

Particle flow example:

Core first version of firmware developed in 2-3 months using HLS, only physicists

Engineering firmware support still required (of course!) — our experience: system interfaces, infrastructure, and signal routing, etc.

Calorimeter clustering

Barrel **Regional Calorimeter Trigger** takes as input single ECal crystals and HCal towers to build clusters

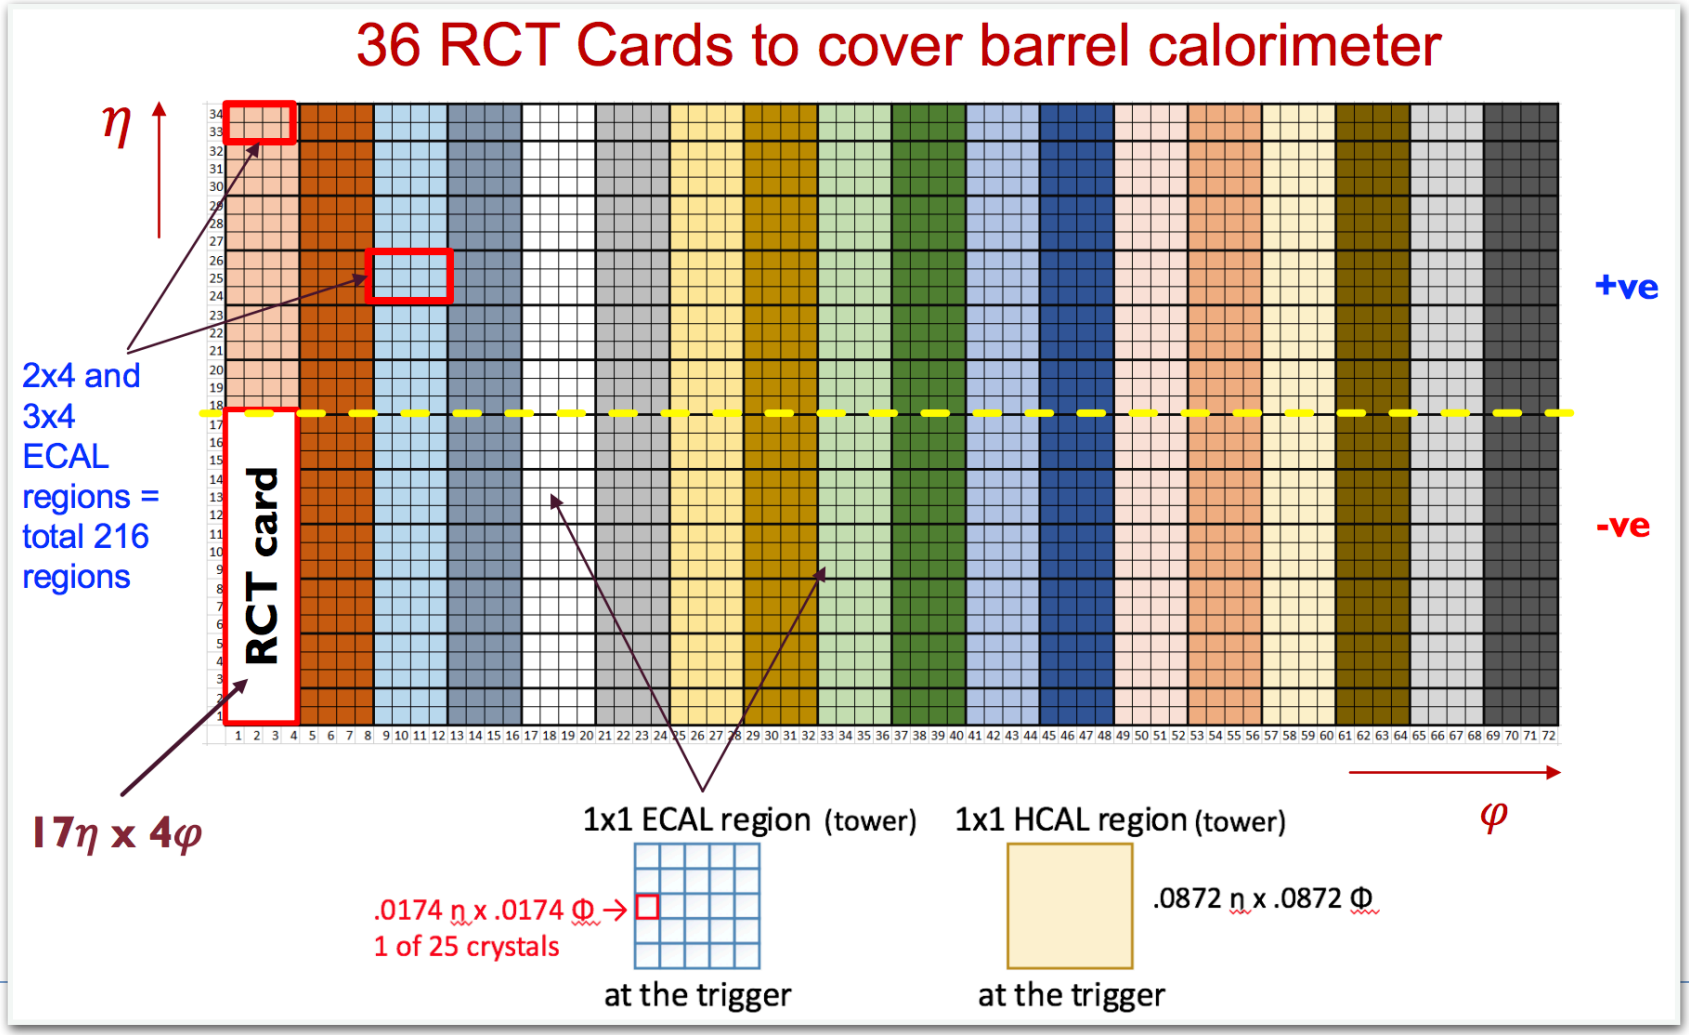

Calorimeter clustering

- **Input:** 17x4x5x5 ECAL crystals and 16x4 HCAL towers.
- 1-RCT card covers 17η x 4φ towers
- Divide card in regions of 3η x 4φ towers to make clusters.

- Building clusters in 3η x 4φ region:
  - Search for seed crystal > 1 GeV
  - Make 3x5 clusters at crystal level
  - Select maximum of 5 highest ET cluster in 3η x 4φ region

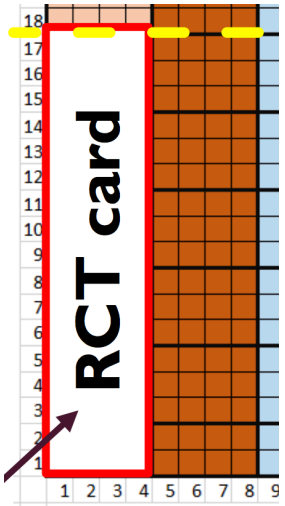

- Move to next 3x4 towers and then do the merging around the neighbors if cluster is at the boundary of the tower

- For 1-RCT card, there are 5 – (3η x 4φ) regions = 30 clusters
- Sort and send a maximum of 12 highest ET clusters

- To these 12 highest ET clusters, if there is a HCAL tower behind the ECAL tower, HCAL ET is also added to the cluster.

- **Output:** 12 Clusters (ECAL + HCAL)

Calorimeter clustering — resources

First Barrel RCT algorithm has reasonable resource usage and latency of 40 clock cycles @ 240 MHz for large resource FPGA (VU9P)

Two 17x2 blocks split across SLRs for better floor planning

Candidate for smaller form factor FPGA

Xilinx VU9P, 240 MHz

```
================================================================
== Utilization Estimates                          360 MHz
================================================================
* Summary:
+-----------------+---------+-------+---------+---------+-----+
|       Name      | BRAM_18K| DSP48E|    FF   |   LUT   | URAM|
+-----------------+---------+-------+---------+---------+-----+
|DSP              |        -|      -|        -|        -|    -|
|Expression       |        -|      -|        0|      504|    -|
|FIFO             |        -|      -|        -|        -|    -|
|Instance         |        0|      -|    46529|    35250|    -|
|Memory           |        -|      -|        -|        -|    -|
|Multiplexer      |        -|      -|        -|     8106|    -|
|Register         |        0|      -|    32007|     7680|    -|
+-----------------+---------+-------+---------+---------+-----+
|Total            |        0|      0|    78536|    51540|    0|
+-----------------+---------+-------+---------+---------+-----+
|Available        |     4320|   6840|  2364480|  1182240|  960|
+-----------------+---------+-------+---------+---------+-----+
|Utilization (%)  |        0|      0|        3|        4|    0|
+-----------------+---------+-------+---------+---------+-----+
```

Similar resource utilization for clock @ 240 MHz

Vertexing algorithms

Vertexing can be done in parallel to particle flow but is needed for pileup mitigation techniques

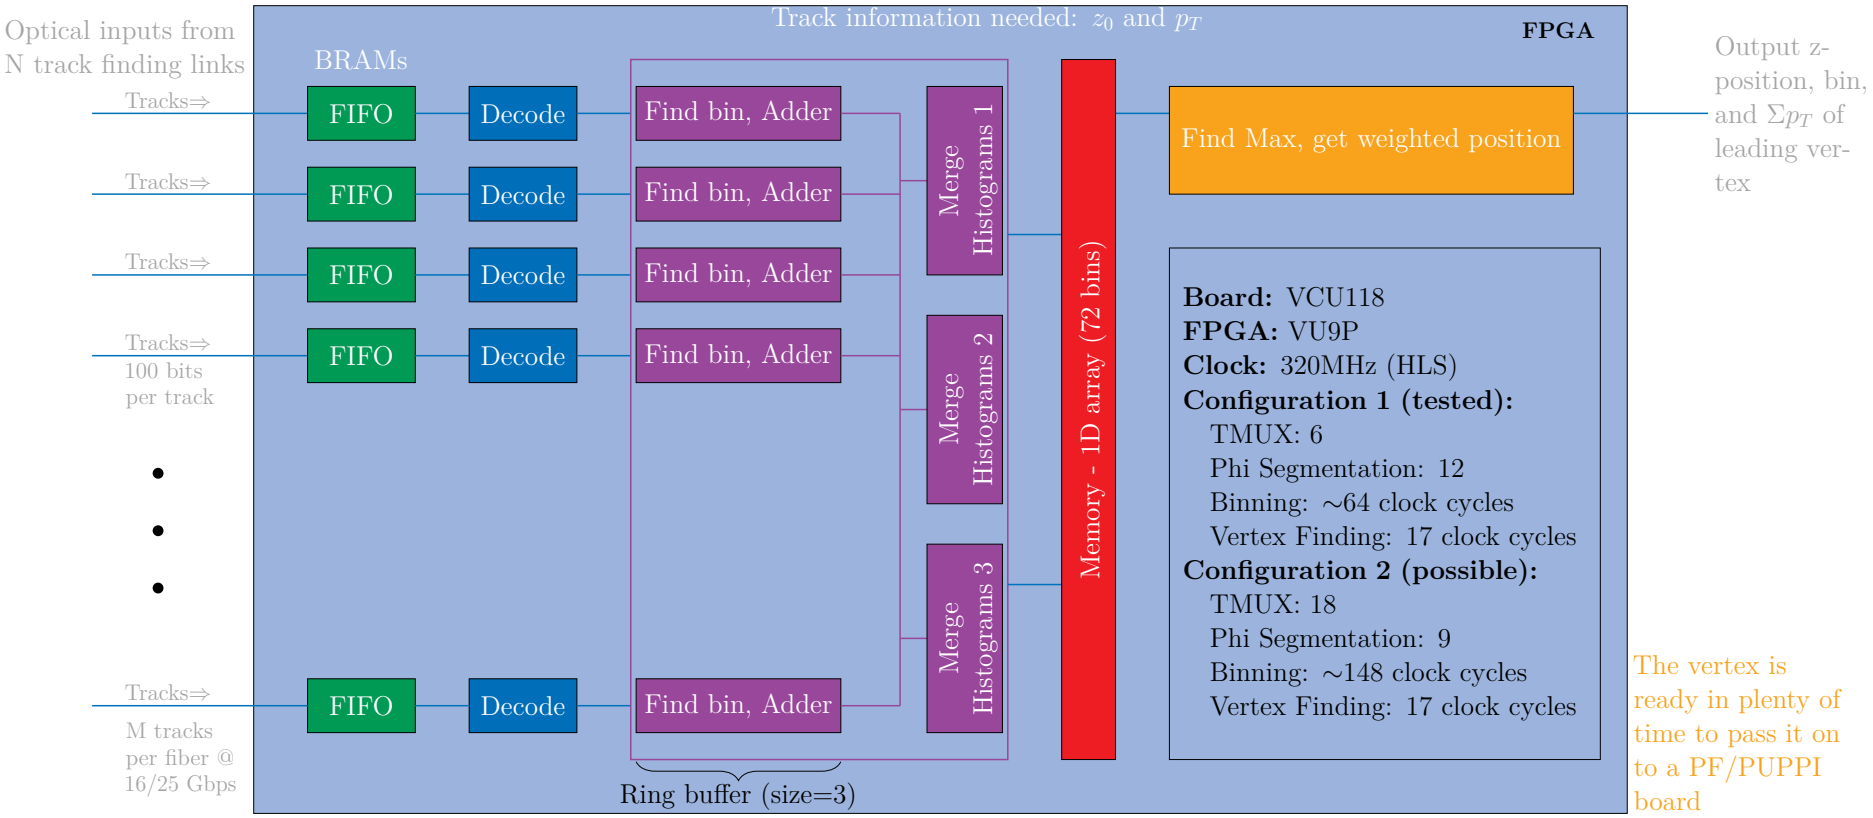

First “fast histogramming” algorithms implemented as a baseline

Vertexing algorithms

Algorithm implemented in HLS with VHDL wrapper to support feeding in tracks and reading out results

VU9P	FF	LUT	BRAM	DSP
Vertexing	117735	154979	1	1
Available	2364480	1182240	2160	6840
% used	5	13	0	0

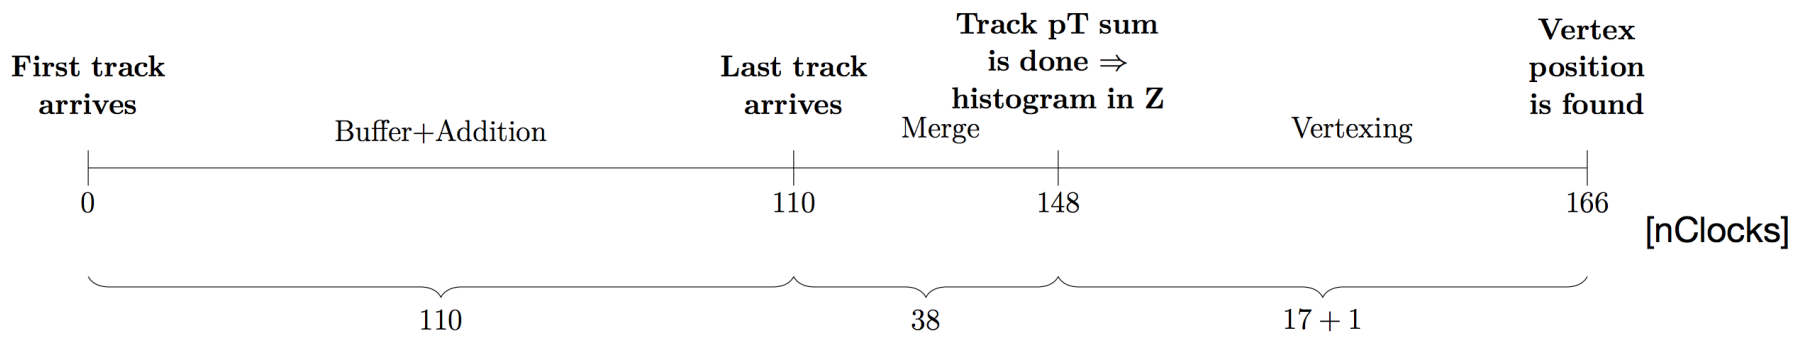

Algorithm finished in time to pass to PF+PUPPI

Particle Flow Engine

Use inspiration from offline reconstruction for best performance

Particle **F**low:

efficient combination of complementary detector subsystems

particle interpretation of the event, improves any single system energy/ spatial **resolution**

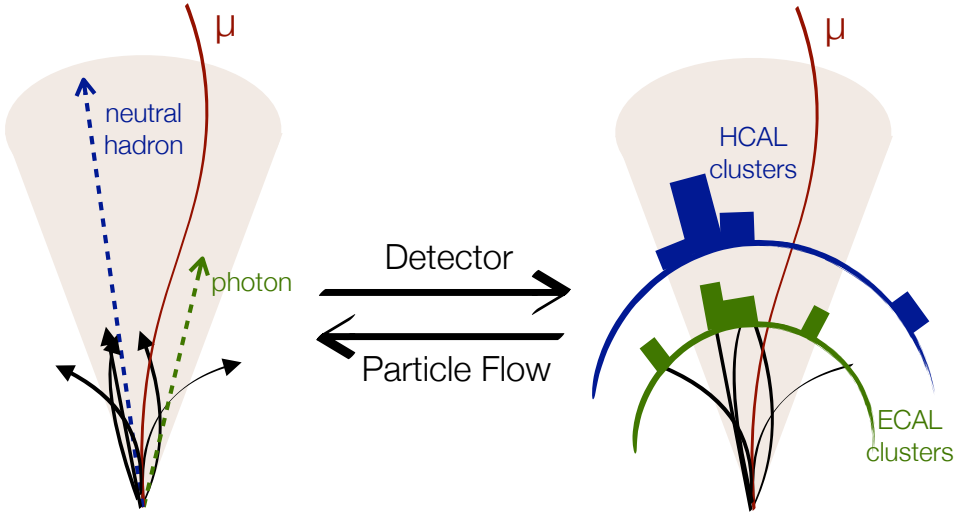

Detector	p_T-resolution	η/Φ-segmentation
Tracker	0.6% (0.2 GeV) – 5% (500 GeV)	0.002 x 0.003 (first pixel layer)
ECAL	1% (20 GeV) – 0.4% (500 GeV)	0.017 x 0.017 (barrel)
HCAL	30% (30 GeV) – 5% (500 GeV)	0.087 x 0.087 (barrel)

PF, offline experience

Large gains from PF on jet and MET resolutions

arXiv:1706.04965 [PF paper]

Particle flow impact

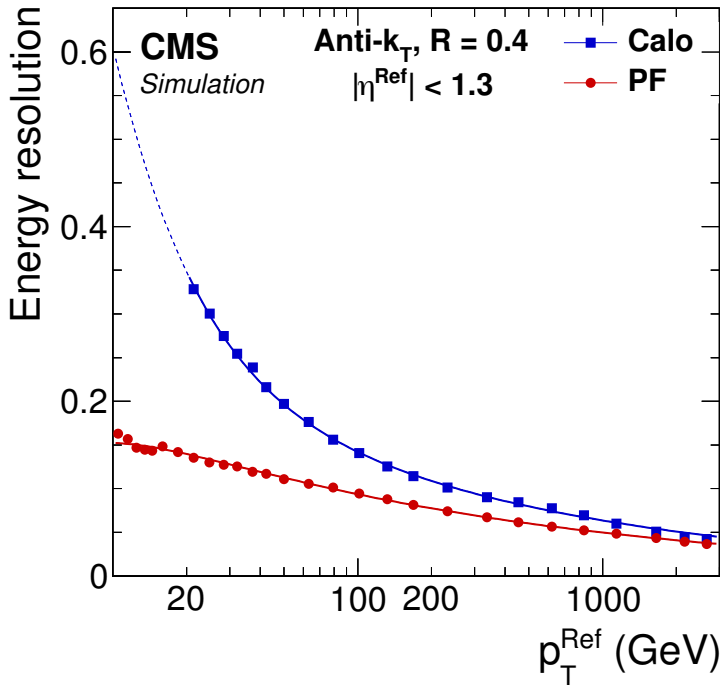

improved jet pT resolution

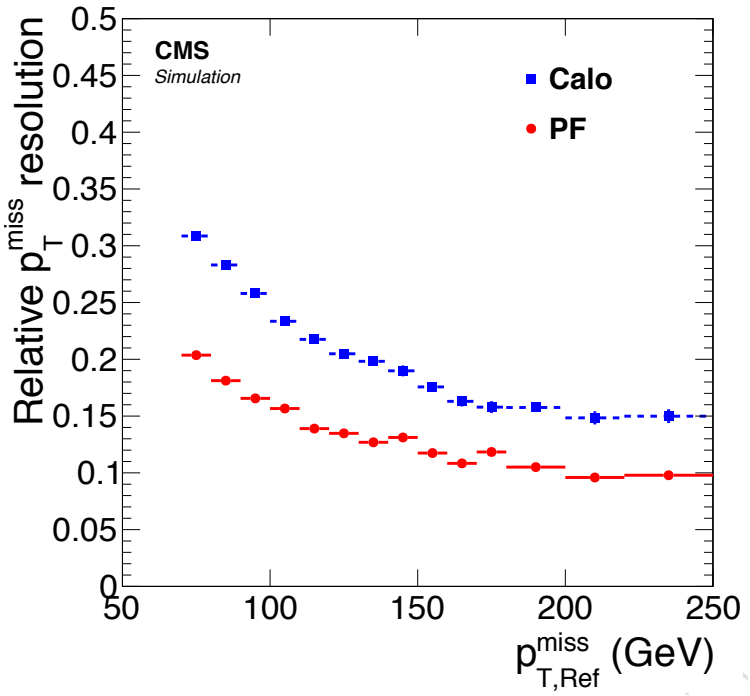

improved missing pT resolution

PF+PUPPI schematic

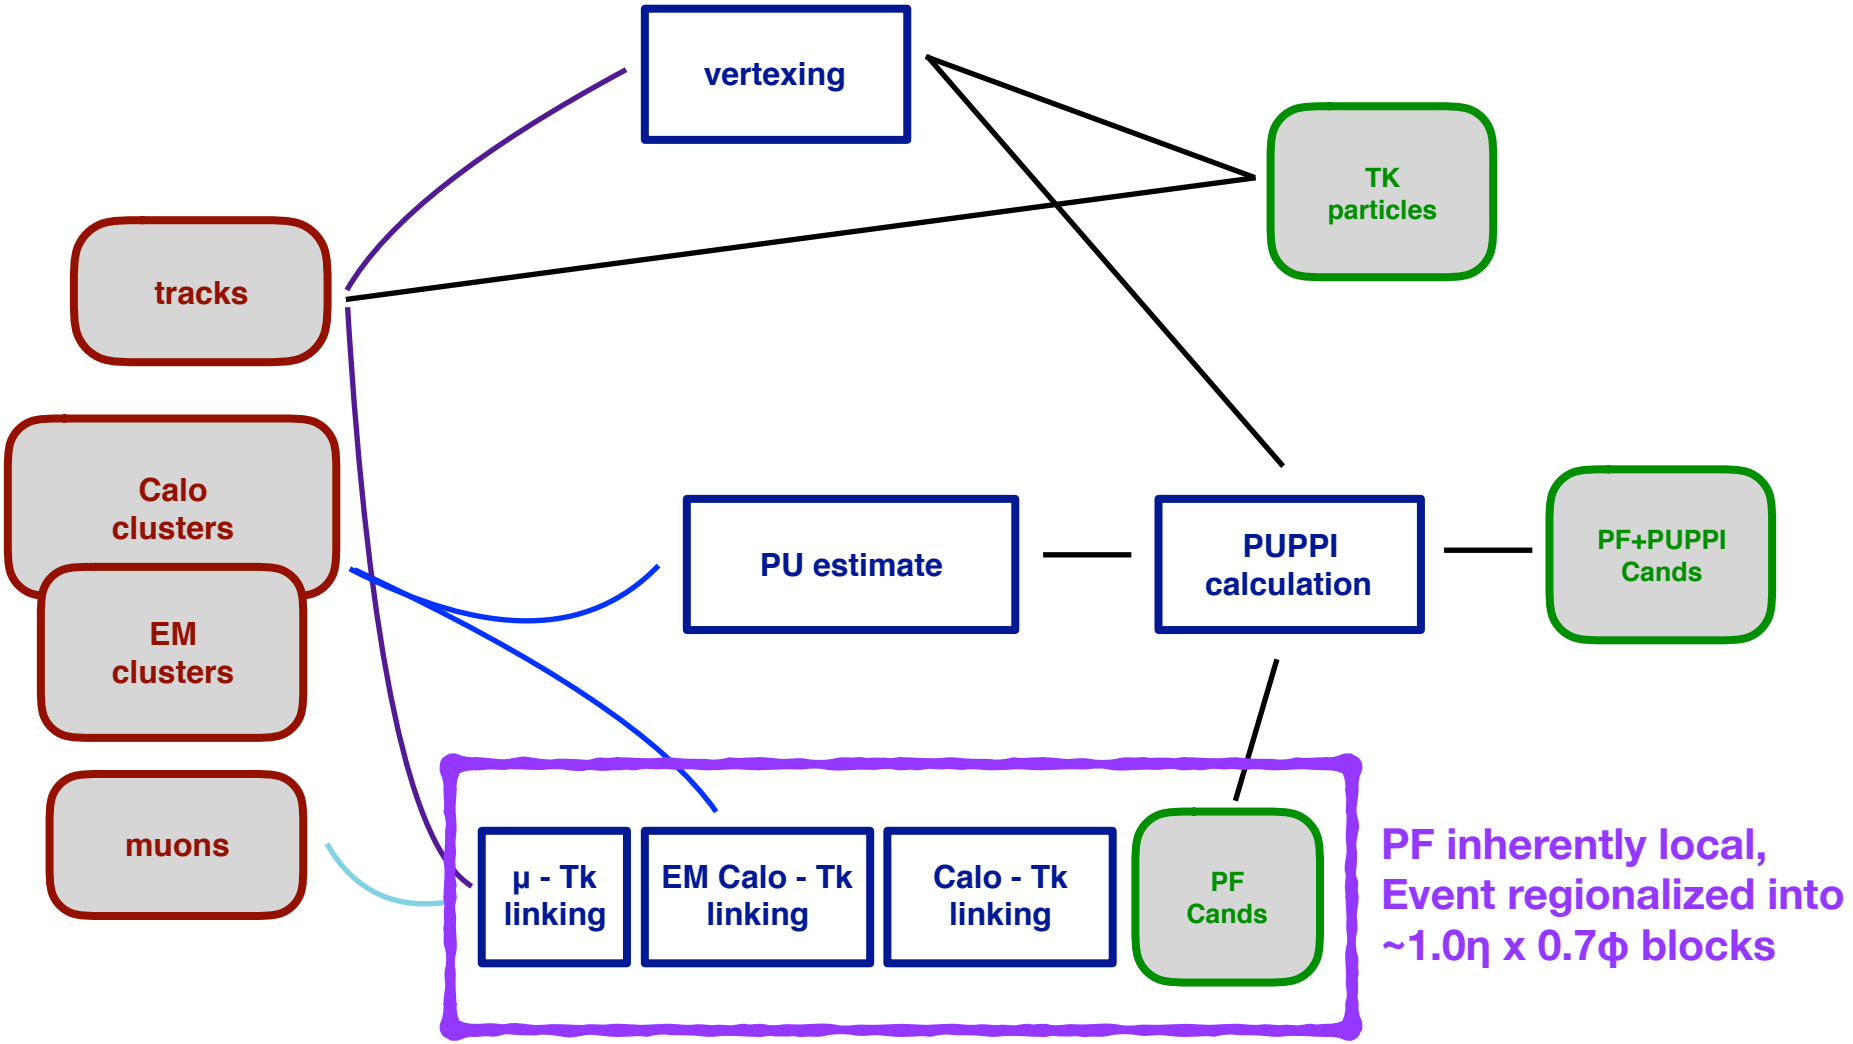

Particle flow regions

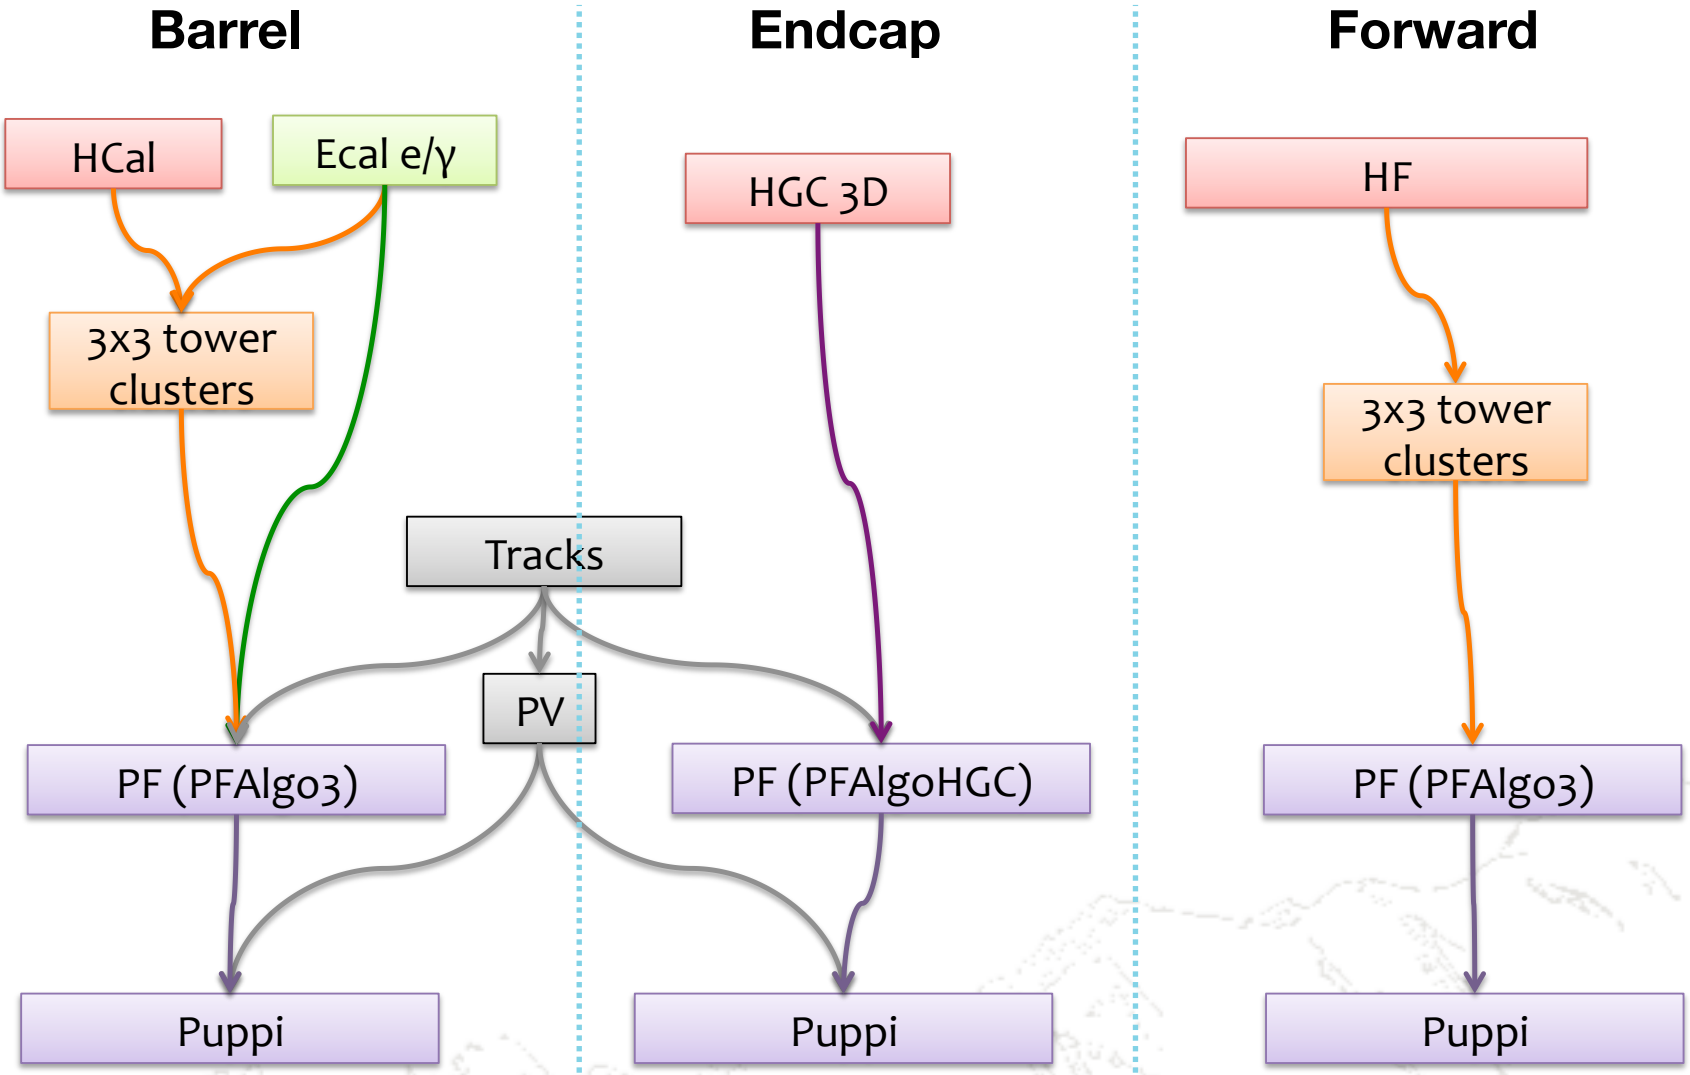

Jet and MET performance

PF+PUPPI algorithms bring significant improvement for hadronic trigger objects

- Continual improvements to algorithms
- Jet algorithms still offline style, work in progress

multijet trigger

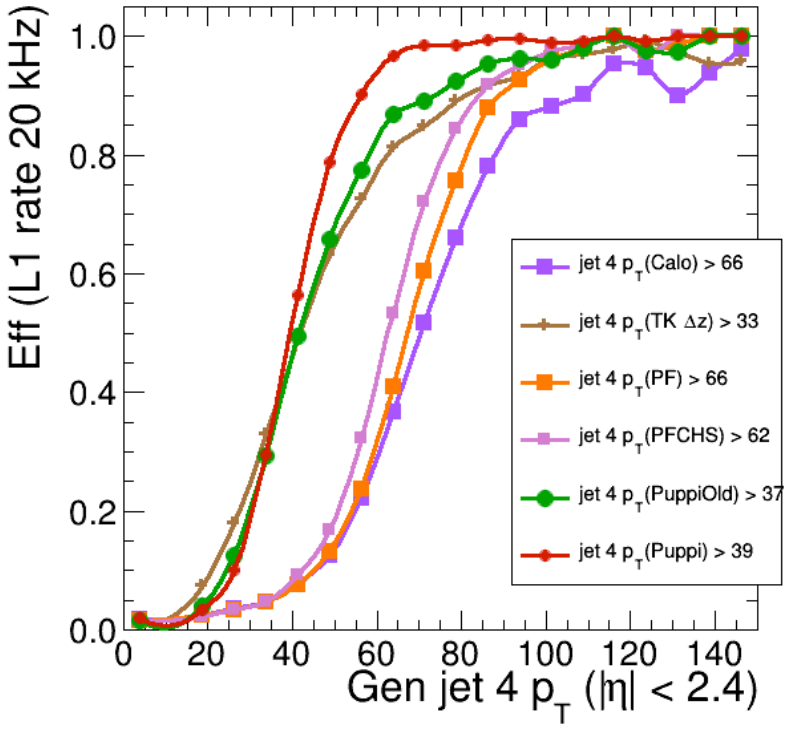

HT trigger

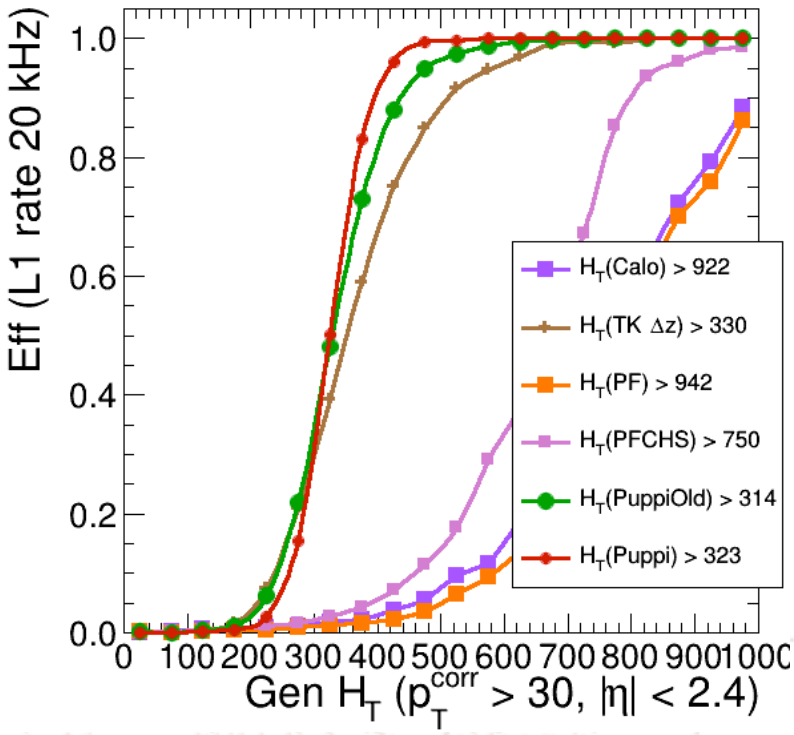

Input definitions

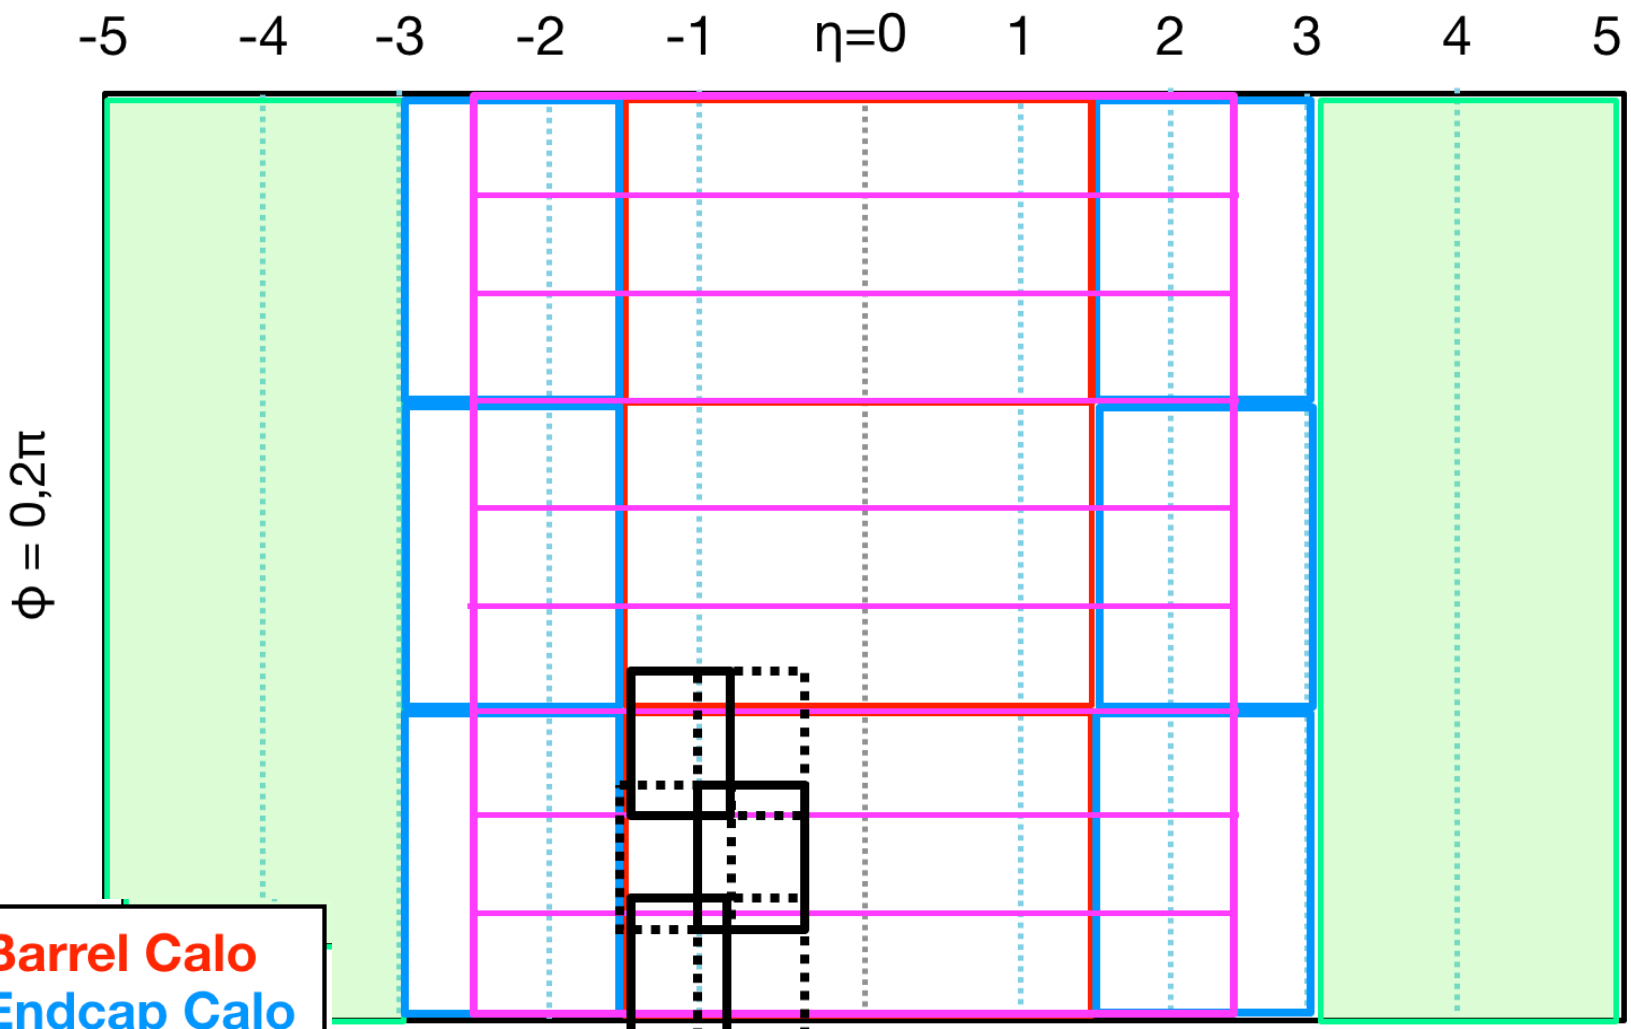

Resources

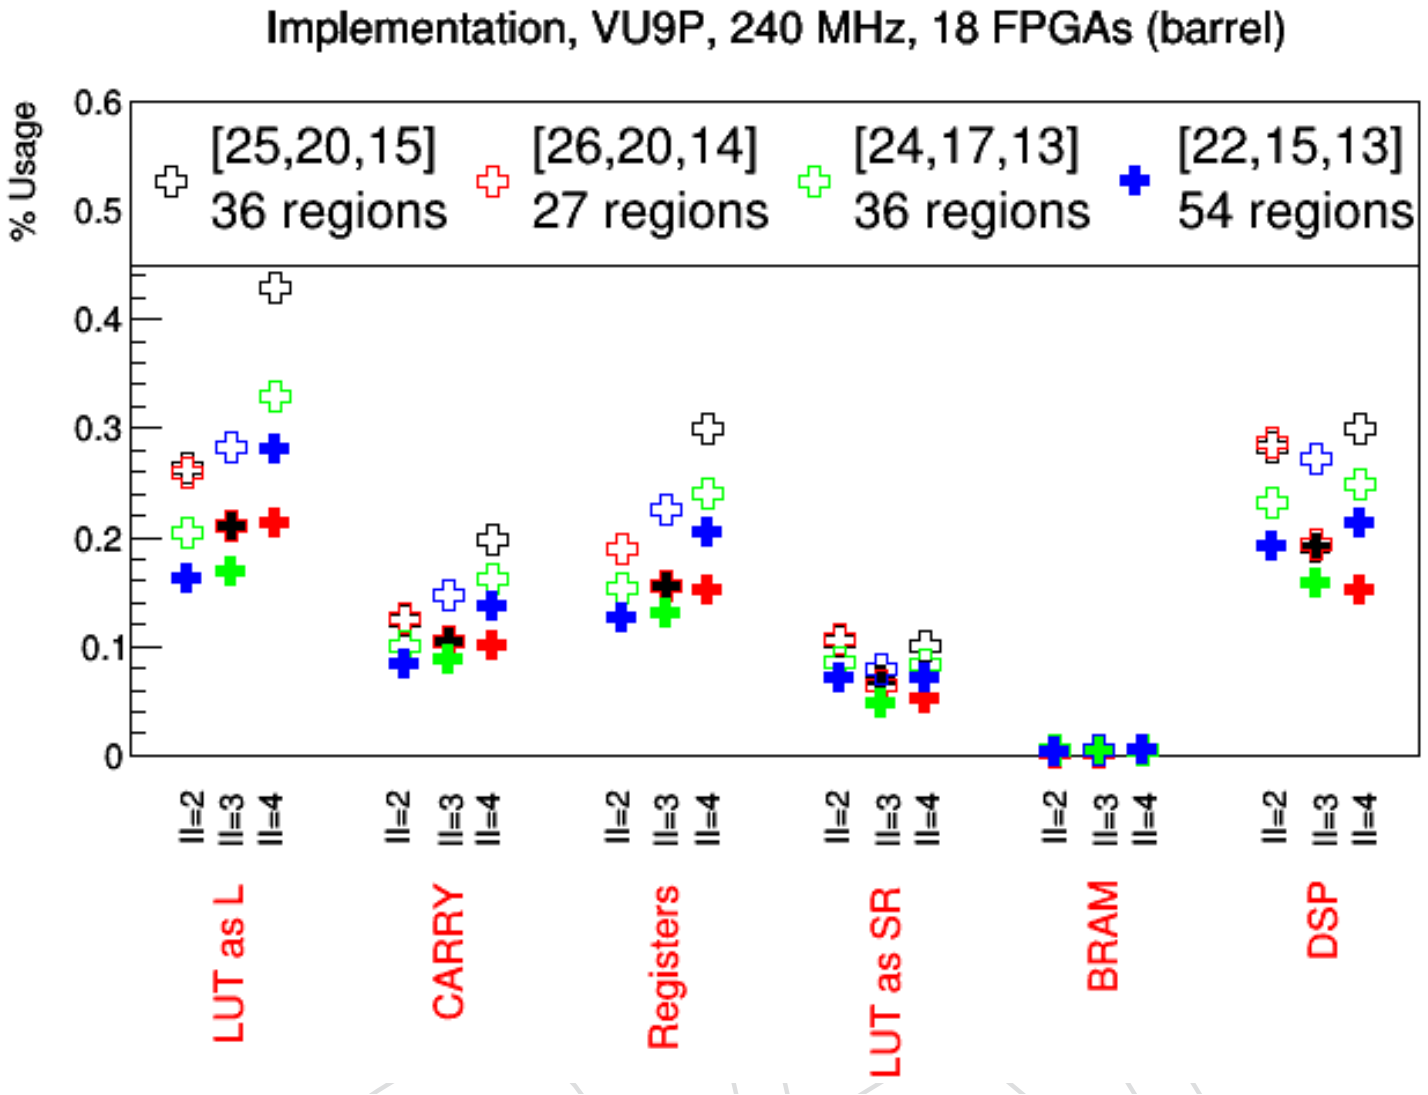
Implementation, VU9P, 240 MHz, 18 FPGAs (barrel)
% Usage
0.6
0.5
0.4
0.3
0.2
0.1
0
[25,20,15]
36 regions
[26,20,14]
27 regions
[24,17,13]
36 regions
[22,15,13]
54 regions
II=2
II=3
II=4
LUT as L
CARRY
Registers
LUT as SR
BRAM
DSP

Architecture (Exploded View)

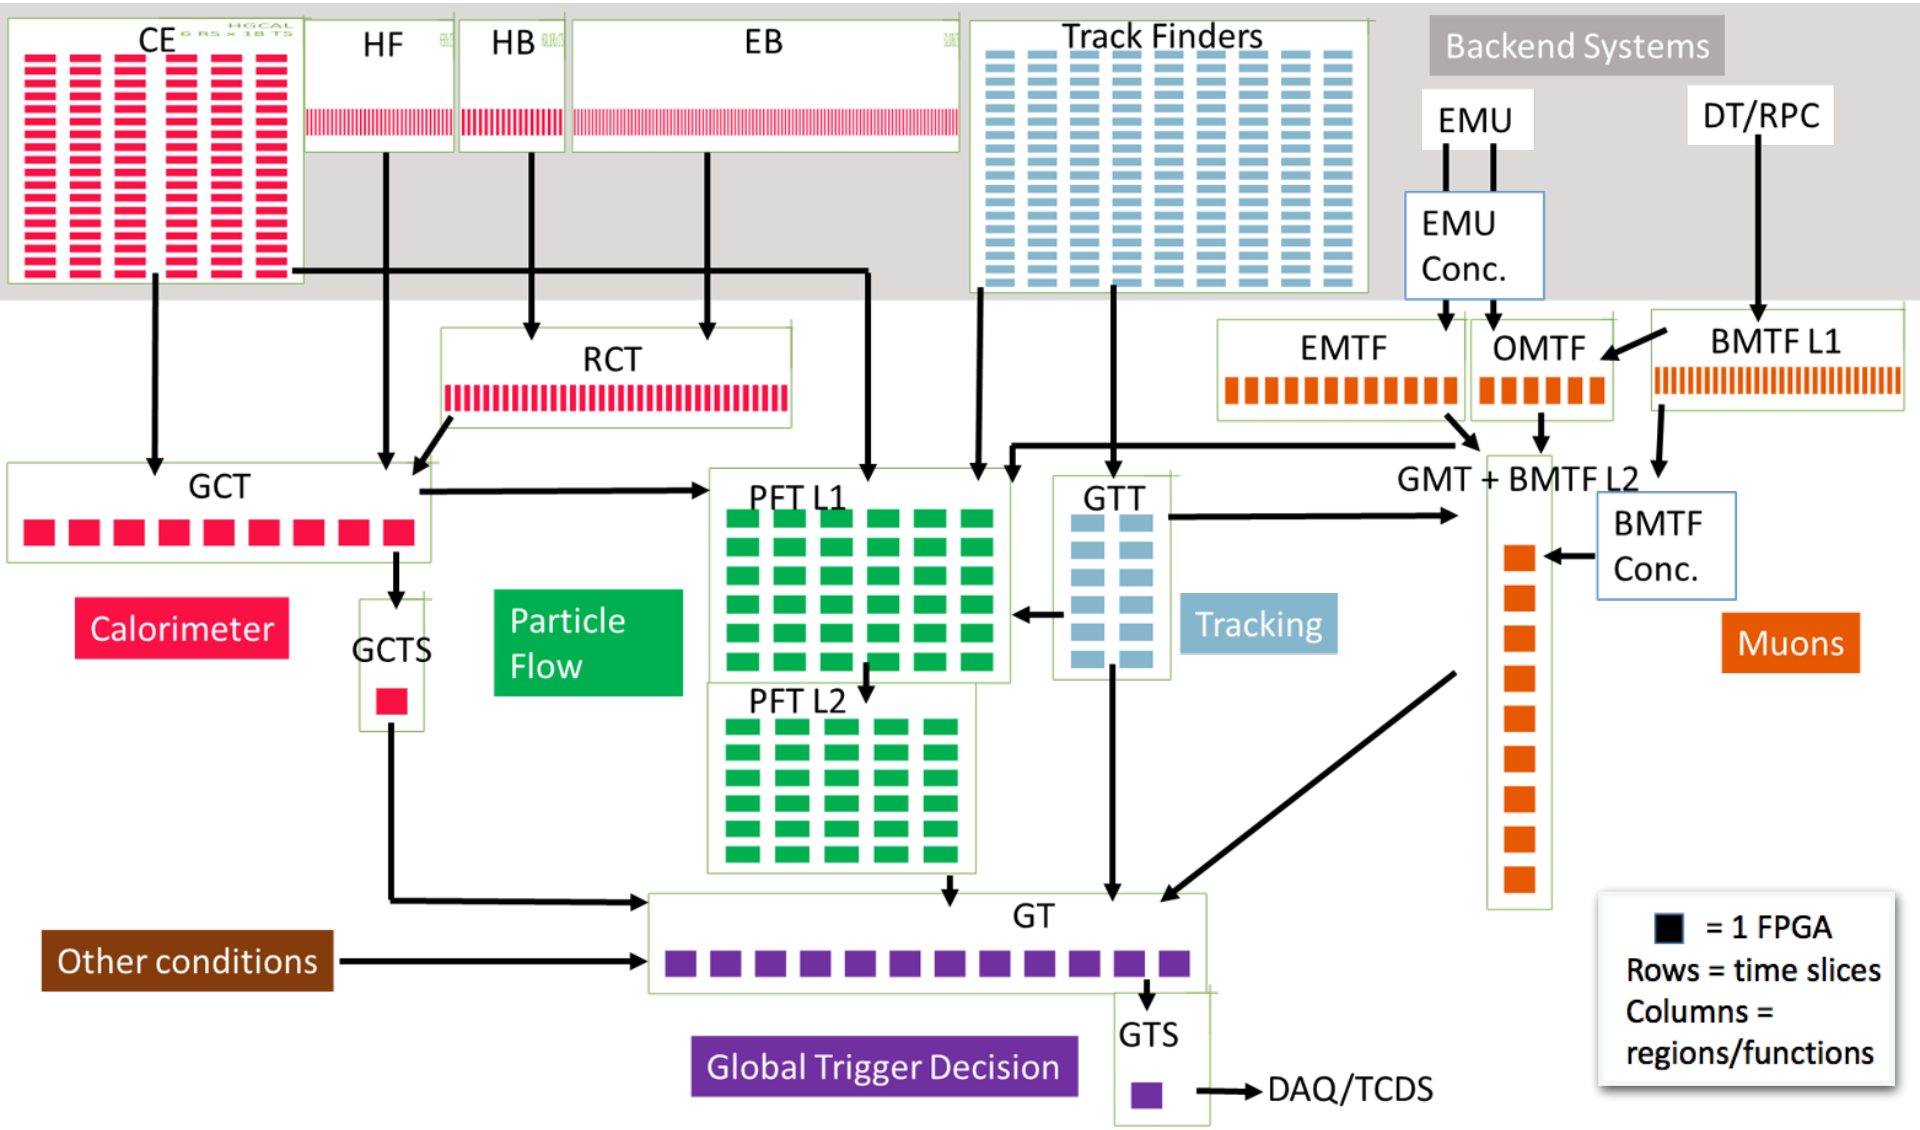

Latency

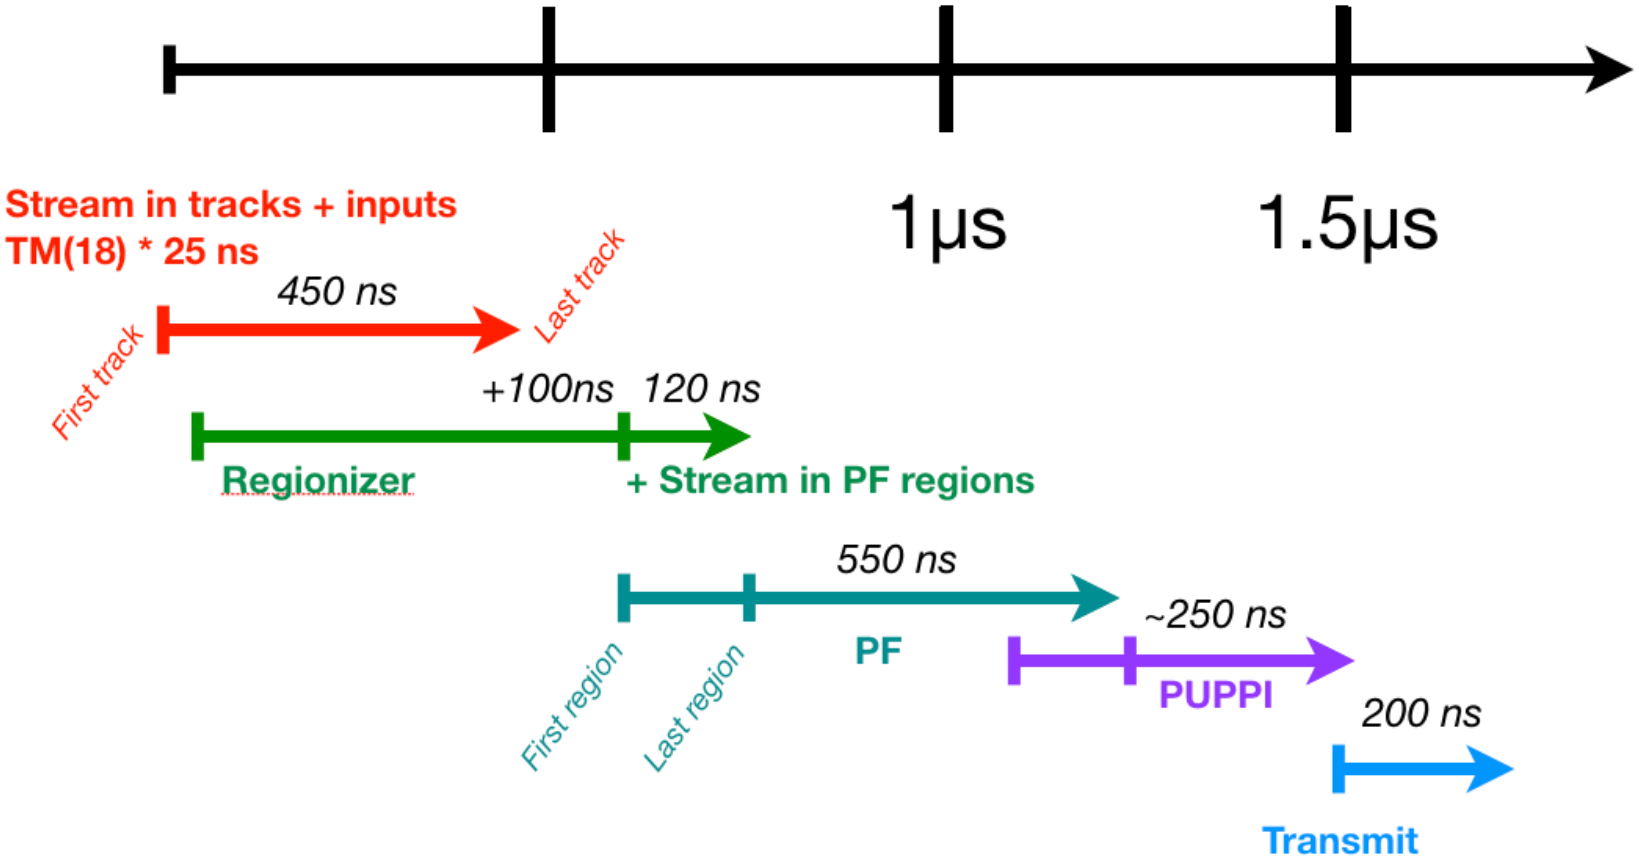
1µs
1.5µs
Stream in tracks + inputs
TM(18) * 25 ns
450 ns
First track
Last track
+100ns
120 ns
Regionizer
+ Stream in PF regions
550 ns
First region
Last region
PF
~250 ns
PUPPI
200 ns
Transmit

Demonstration Firmware

(More in Sridhara's talk)

- Gen-0 (in operation at UW and CERN)
  - CTP7-based (3 cards, 96 total active links @ 10.0G available)
  - HLS interface grafted onto Phase 1 8b10b link infrastructure
  - 64-bit link interfaces
- Gen-1 (in operation at UW)
  - APd1 single card setup
  - Iridis-style 64b66b transport
  - Early APx shell functionality
  - Asynchronous processing clock
- Gen-2 (underway)
  - Multicard test setup
  - Iridis-style signaling and APx shell environment
  - Common emulated TCDS timebase in ATCA/MicroTCA crates using CTP7

Versioned in Github/Gitlab

Demonstration

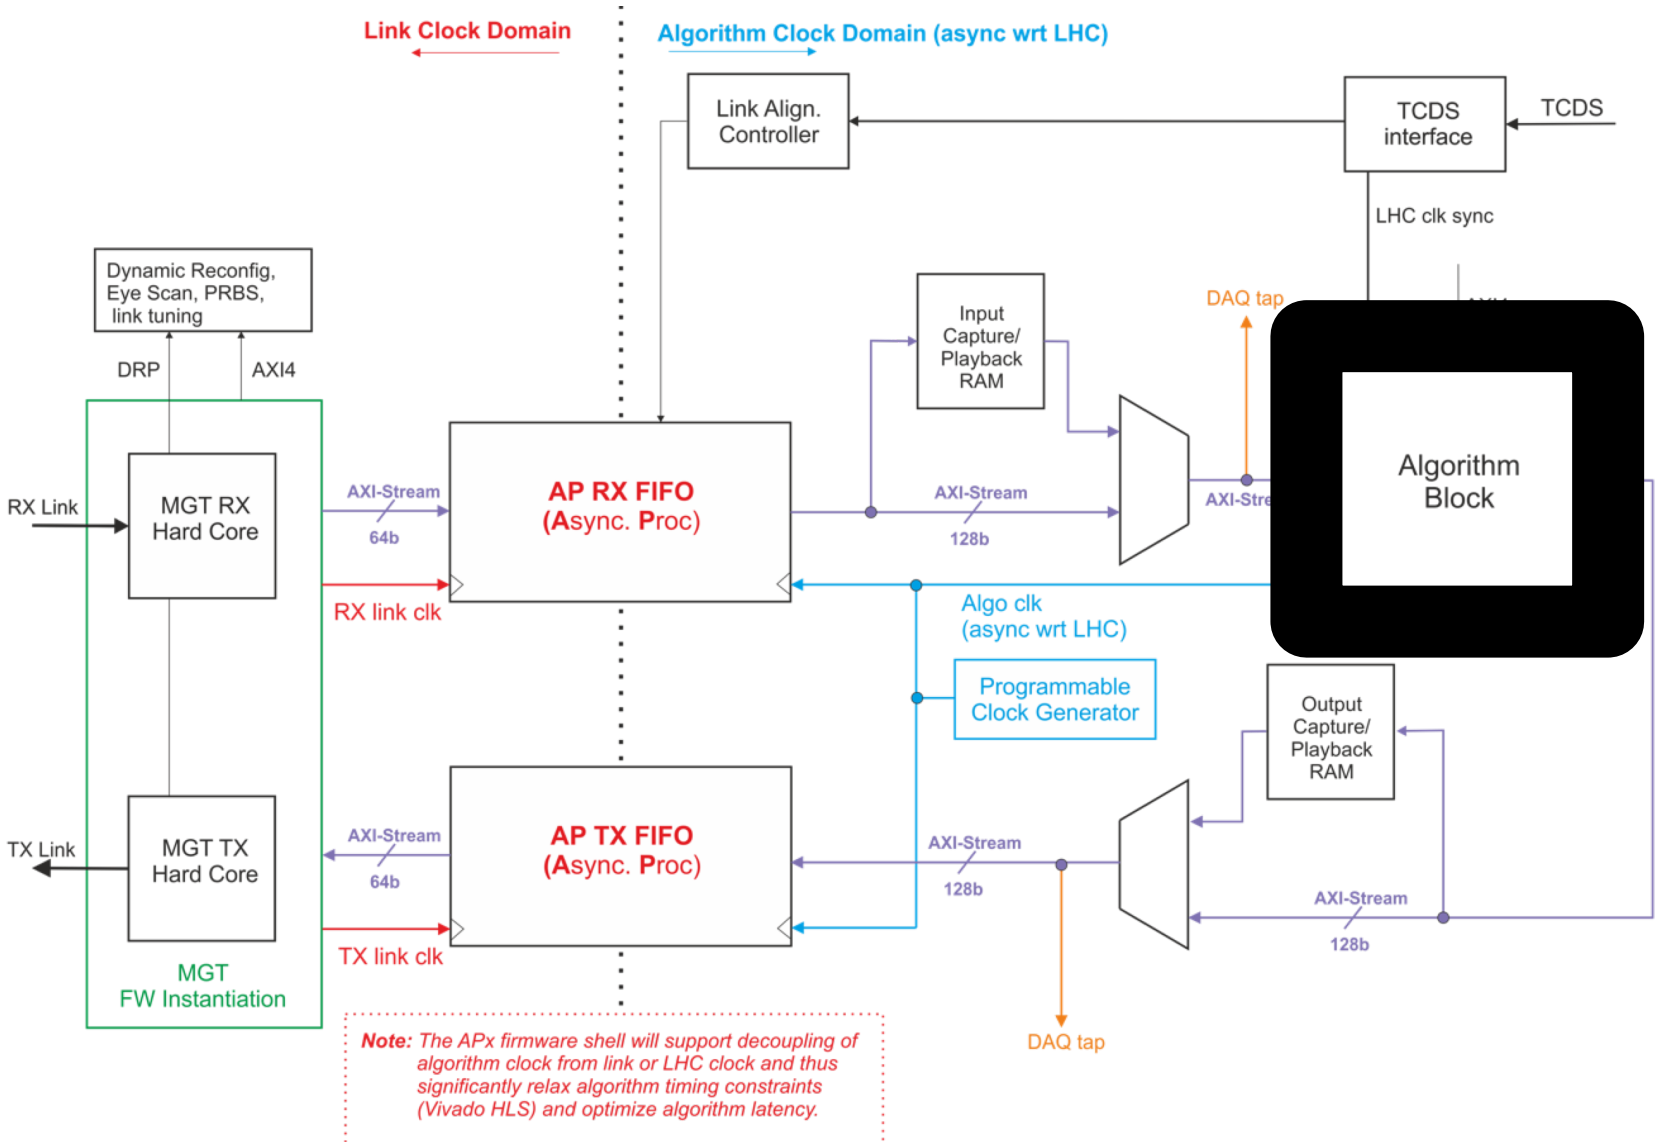

Demonstration

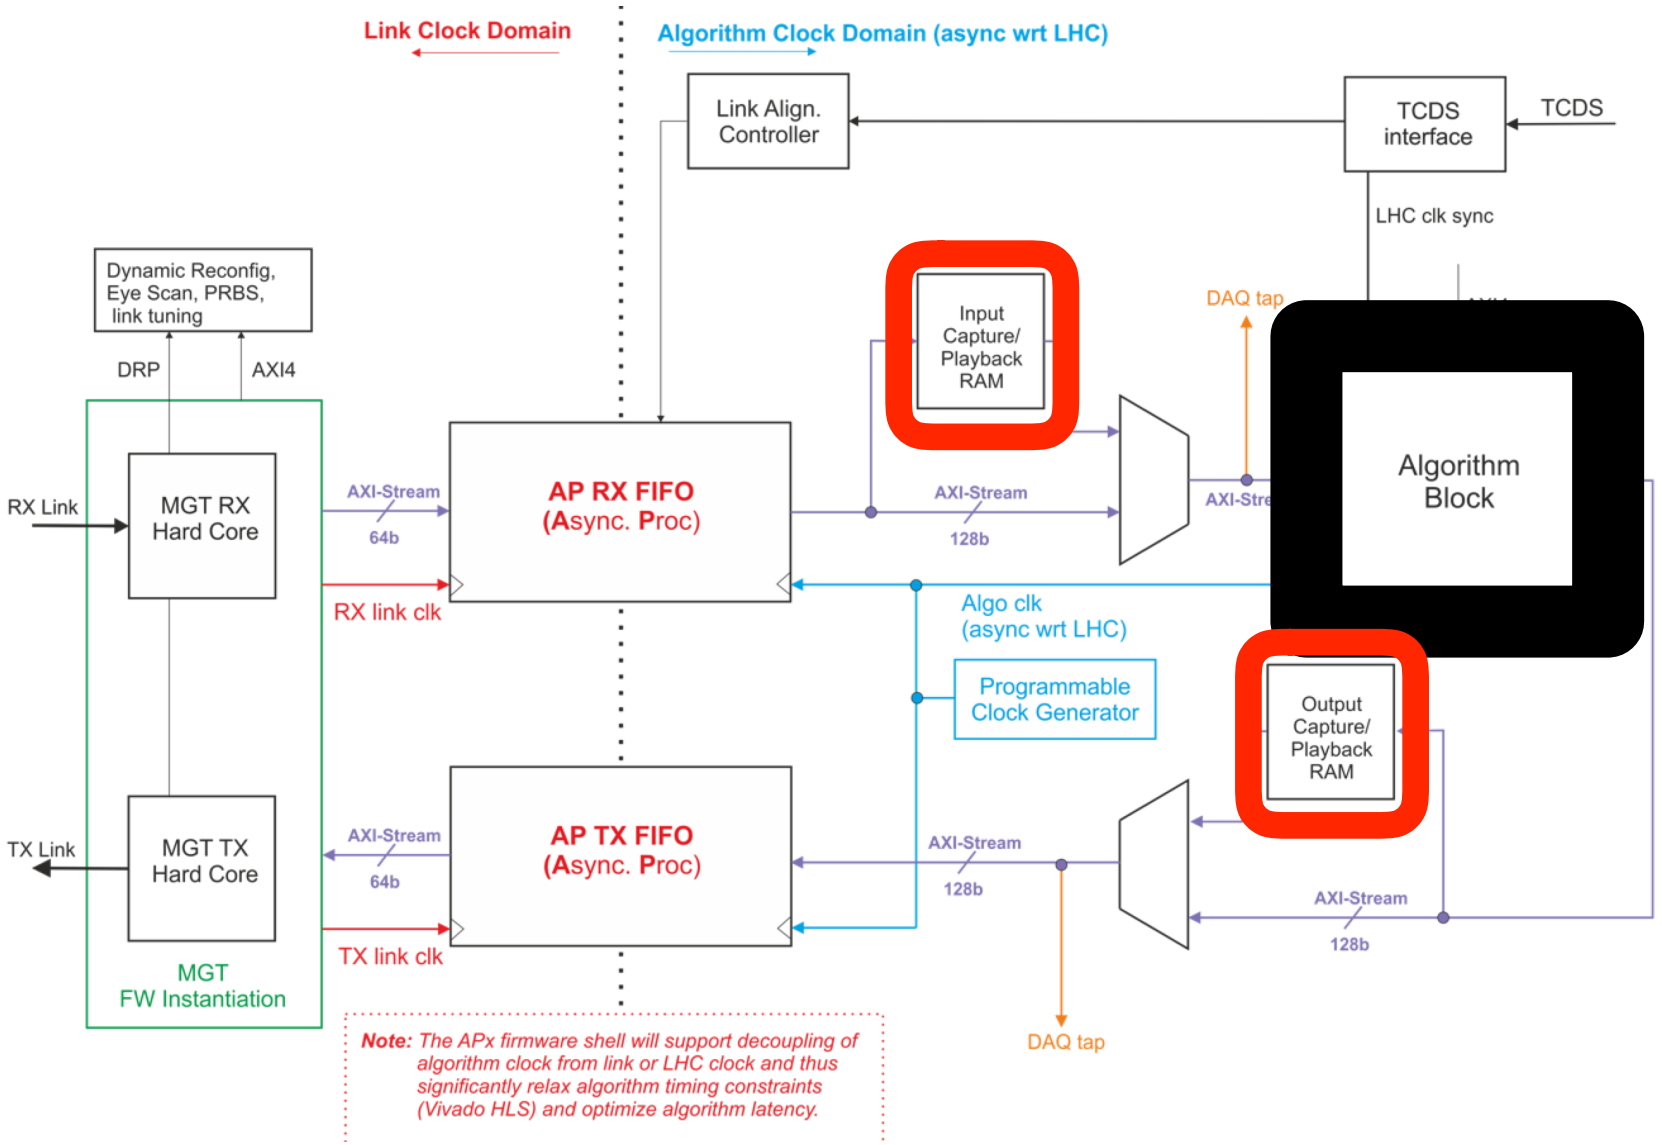

Bitwise simulation

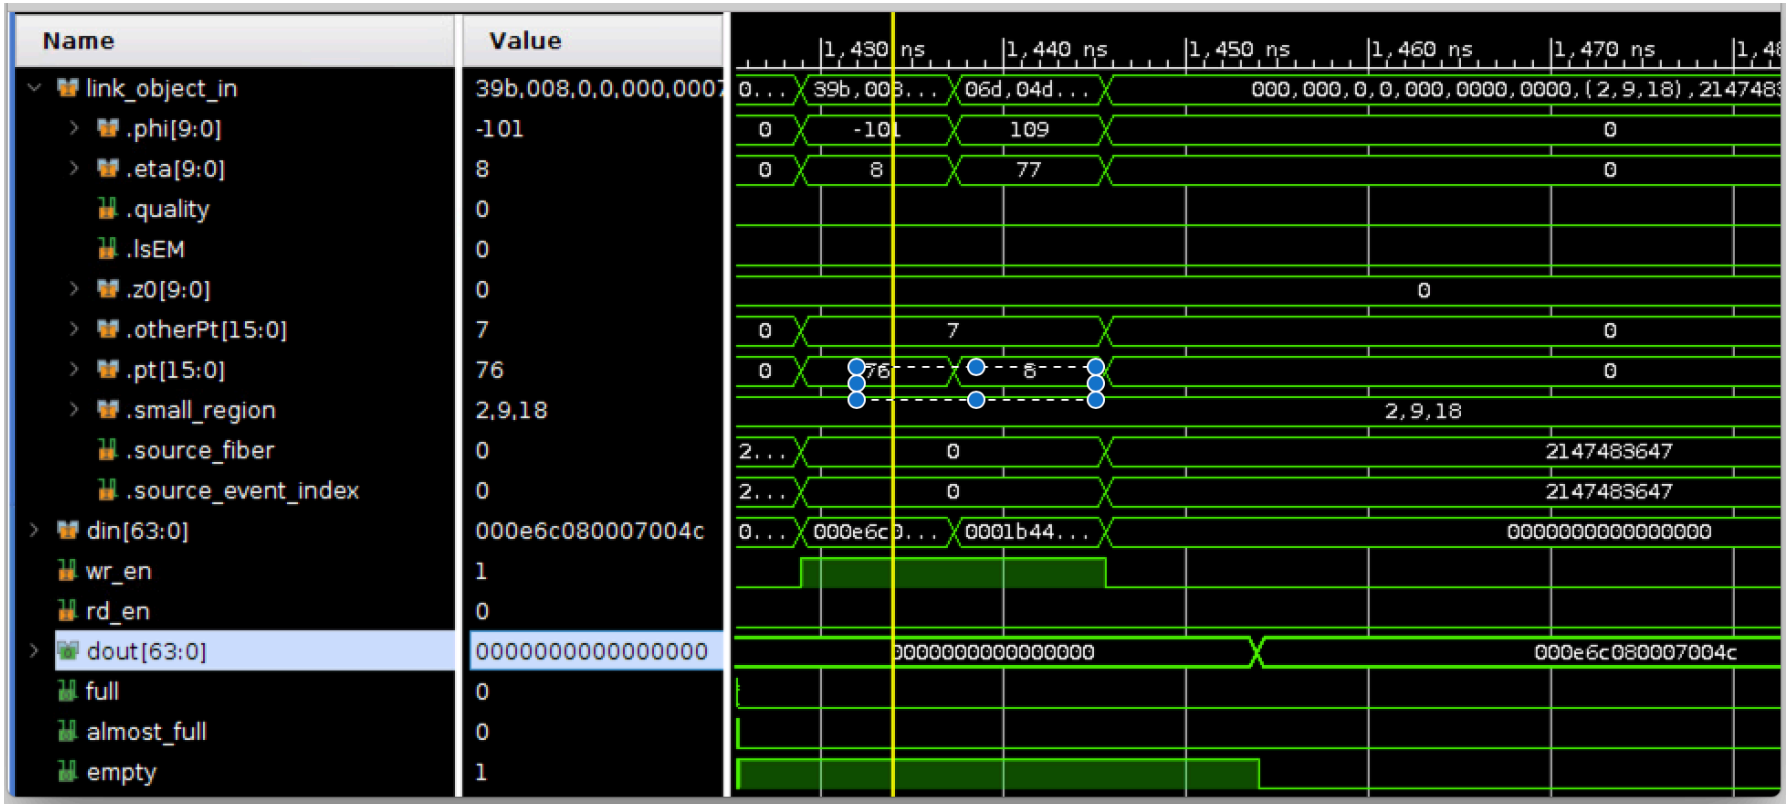

Demonstration

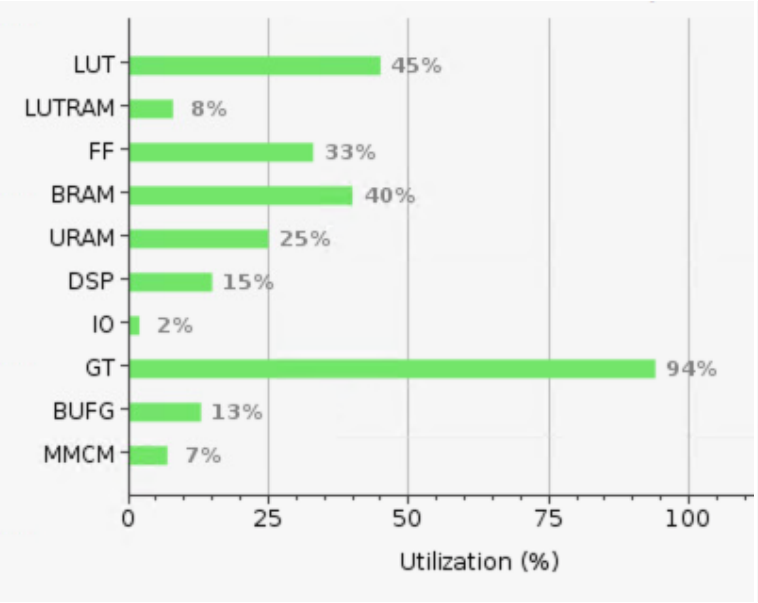

Resource	Utilization	Available	Utilization %
LUT	528286	1182240	44.69
LUTRAM	47220	591840	7.98
FF	785310	2364480	33.21
BRAM	871	2160	40.32
URAM	243	960	25.31
DSP	1021	6840	14.93
IO	9	416	2.16
GT	98	104	94.23
BUFG	227	1800	12.61
MMCM	2	30	6.67

Utilization within project requirements

Less than 50% (70%) for algorithm (total) firmware

Recent milestone: all elements integrated and meeting timing with full place & route

Demonstration

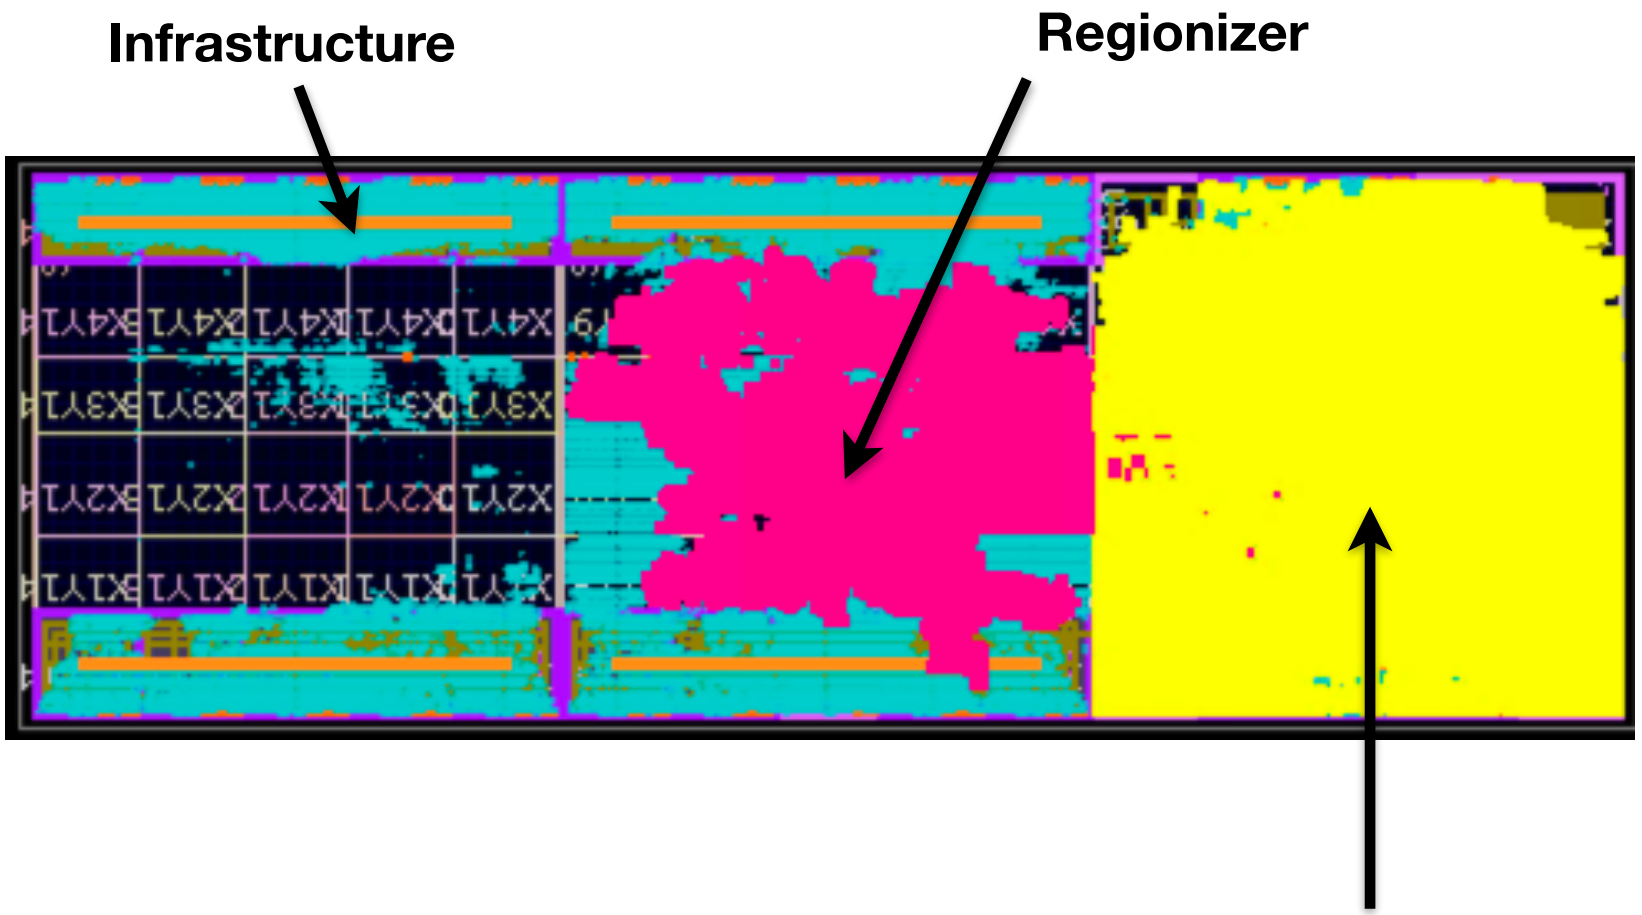

Summary of algorithm status

	baseline algo	firmware
Clustering		
ID		
Calibration		
Track prop		
PF block		
Vertexing		
PUPPI		
trk jet		
τ's		
calo e/γ		

Suite of algorithms to meet physics needs (menu) demonstrated

Firmware for most resource intensive algorithms within system requirements to meet mission need

R&D status

Achieved

- Full simulation framework for studying algorithm physics performance
- Algorithm development with High Level Synthesis (HLS) tools and bitwise validation of firmware/software
- Integration of HLS IP blocks into firmware infrastructure in multiple modes (direct to infra, through intermediate VHDL)

Needed before production

- Final specification of algorithm interfaces for board-to-board communication (with USCMS and iCMS technologies)

MANAGEMENT

Institutions and contributed labor

Contributing institutions

Clustering and ID: UW

Calibration: MIT, Fermilab, UIC

Track propagation: TAMU

Muon-track correlation: UCLA, UF, TAMU, Fermilab

Vertexing and track-based objects: CU Boulder, Rutgers

Particle Flow and PUPPI: MIT, Fermilab, UIC

Calo-based objects: UW

Risk register

Threat	RT-402-6-02-D	TD - Board or parts vendor non-performance (DOE)
Threat	RT-402-6-03-D	TD - I/O performance does not meet requirements (DOE)
Threat	RT-402-6-04-D	TD - Additional board redesign is required (DOE)
Threat	RT-402-6-05-D	TD - Additional firmware development is required (DOE)
Threat	RT-402-6-06-D	TD - Baseline FPGA does not satisfy requirements (DOE)
Threat	RT-402-6-90-D	TD - Key Trigger or DAQ personnel need to be replaced (DOE)
Threat	RT-402-6-91-D	TD - Shortfall in Trigger or DAQ scientific labor (DOE)

RT-402-6-05: Risk Rank = Medium, firmware does not meet technical specifications
Probability: 20%, Impact: $10-60k
Mitigation: allocate more time to on-project engineering, or new hire

RT-402-6-90: Risk Rank = Low, due to lack of base funding
Probability: 30%, Impact: $0-292k
Mitigation: replace with costed labor, other USCMS or iCMS institutions

QA/QC

Quality assurance plan (cms-doc-13093,cms-doc-13318)

Flows down from the engineering requirements

- Acceptance plans based on QC activities
- For algorithms, basic requirements on resource and latency met with each firmware release
- Important: SW and FW stored and maintained in repositories

TD-QC-004	Design Verification and Measurement/ Testing	The total (algorithm) firmware will be designed to execute within the required latency budget and not occupy more than 75% (50%) of chip resources involves analysis of logic utilization estimates during firmware synthesis.	Acceptance of each firmware release will involve validation activities that verify that the firmware fits within the chip resource limits and executes within the required latency.

ES&H

- As with entire project, we follow the Integrated Safety Management Plan (cms-doc-13395) and have documented our hazards in the preliminary Hazard Awareness Report (cms-doc-13394)

Summary

Algorithm performance and firmware have progressed since 2018 CD1

Algorithms for: barrel calorimeter trigger, global calorimeter trigger, correlator (including vertexing, track-based objects)

Full demonstration system for algorithm firmware progressing

First demonstration performed

In sync with iCMS milestones for TDR in 2019

ADDITIONAL MATERIAL