Highlights of TWEPP 2023

Tom Williams (RAL)

RAL seminar, 23/11/2023

TWEPP 2023

TWEPP = Topical Workshop on **Electronics** for Particle Physics

The workshop covers all aspects of electronic systems, components and instrumentation for particle and astro-particle physics such as: electronics for particle detection, triggering, data-acquisition systems, accelerator and beam instrumentation.
Operational experience in electronic systems and R&D in electronics for LHC, High Luminosity LHC, FAIR, neutrino facilities and other present or future accelerator projects are the major focus of the workshop.

TWEPP 2023

TWEPP = Topical Workshop on **Electronics** for Particle Physics

The workshop covers all aspects of electronic systems, components and instrumentation for particle and astro-particle physics such as: electronics for particle detection, triggering, data-acquisition systems, accelerator and beam instrumentation.
Operational experience in electronic systems and R&D in electronics for LHC, High Luminosity LHC, FAIR, neutrino facilities and other present or future accelerator projects are the major focus of the workshop.

TWEPP 2023

TWEPP = Topical Workshop on Electronics for Particle Physics

The purpose of the workshop is:
- Present original concepts and results of research and development for electronics relevant to particle physics experiments as well as accelerator and beam instrumentation at future facilities;
- Review the status of electronics for running experiments and accelerators;
- Identify and encourage common efforts for the development of electronics;
- Promote information exchange and collaboration in the relevant engineering and physics communities.

TWEPP 2023

TWEPP = Topical Workshop on Electronics for Particle Physics

The purpose of the workshop is:
- Present original concepts and results of research and development for electronics relevant to particle physics experiments as well as accelerator and beam instrumentation at future facilities;
- Review the status of electronics for running experiments and accelerators;
- Identify and encourage common efforts for the development of electronics;
- Promote information exchange and collaboration in the relevant engineering and physics communities.

TWEPP 2023

A few plenary sessions: invited talks

Two parallel sessions, talks grouped into tracks:

- ASICs
- Optoelectronics & links
- Power, grounding & shielding
- Production, testing & reliability
- Programmable logic, design & verification tools and methods
- Radiation tolerant components & systems
- Module, PCB & component design
- System design, description & operation
- Trigger and timing distribution

Two poster sessions

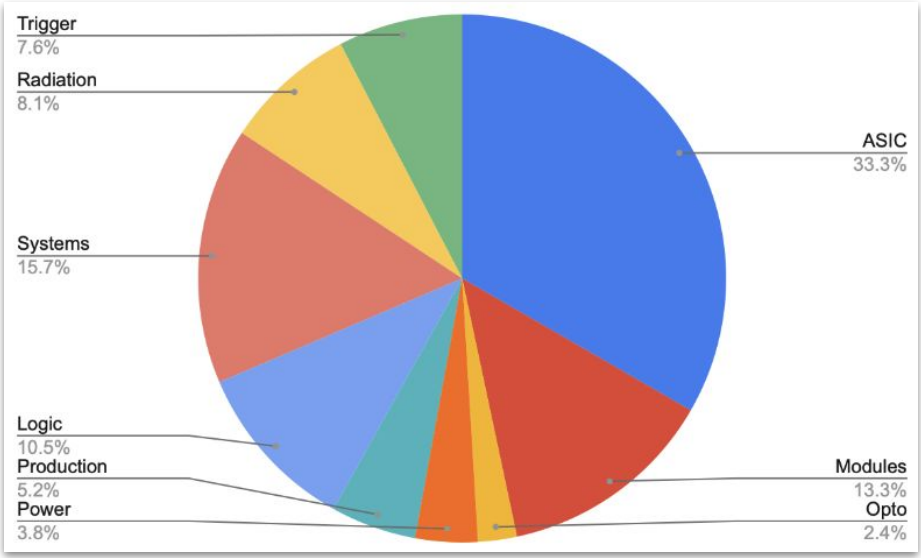

Abstract submissions by topic

TWEPP 2023

A few plenary sessions: invited talks

Two parallel sessions, talks grouped into tracks:

- ASICs
- Optoelectronics & links
- Power, grounding & shielding
- Production, testing & reliability
- Programmable logic, design & verification tools and methods
- Radiation tolerant components & systems
- Module, PCB & component design
- System design, description & operation
- Trigger and timing distribution

Two poster sessions

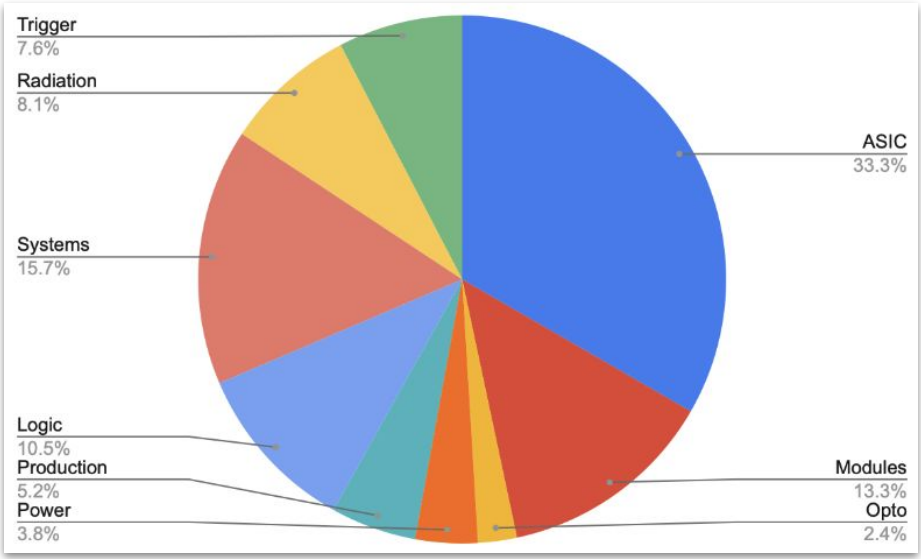

Abstract submissions by topic

Topics I'll focus on today …

Location, Location, Location

Calaserena Village, near Cagliari (Sardinia)

- Resort in Geremeas
  - Very small village, mainly holiday homes & resorts
  - 45 minute drive from Cagliari
  - Surprisingly cheap
    - Out of season — temperature 20 - 25C *(from my perspective, ideal conference weather!)*

- In previous years, hosted in a city or at least a large town. Potential downside of resort?
  - Stuck there all week
    - What if the food is bad?
    - *Even worse: What if the bar closes early …*

- On the other hand: Private beach!

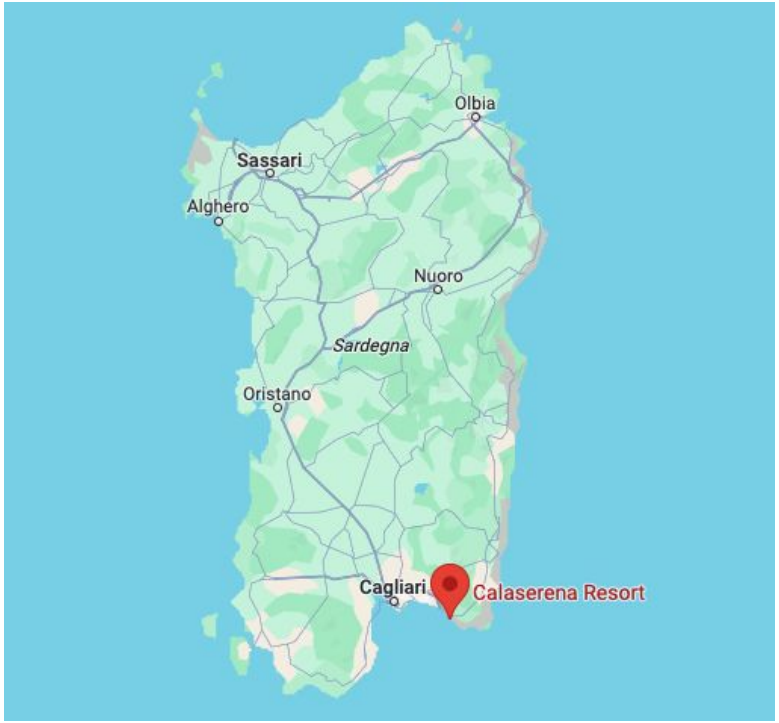

Location, Location, Location

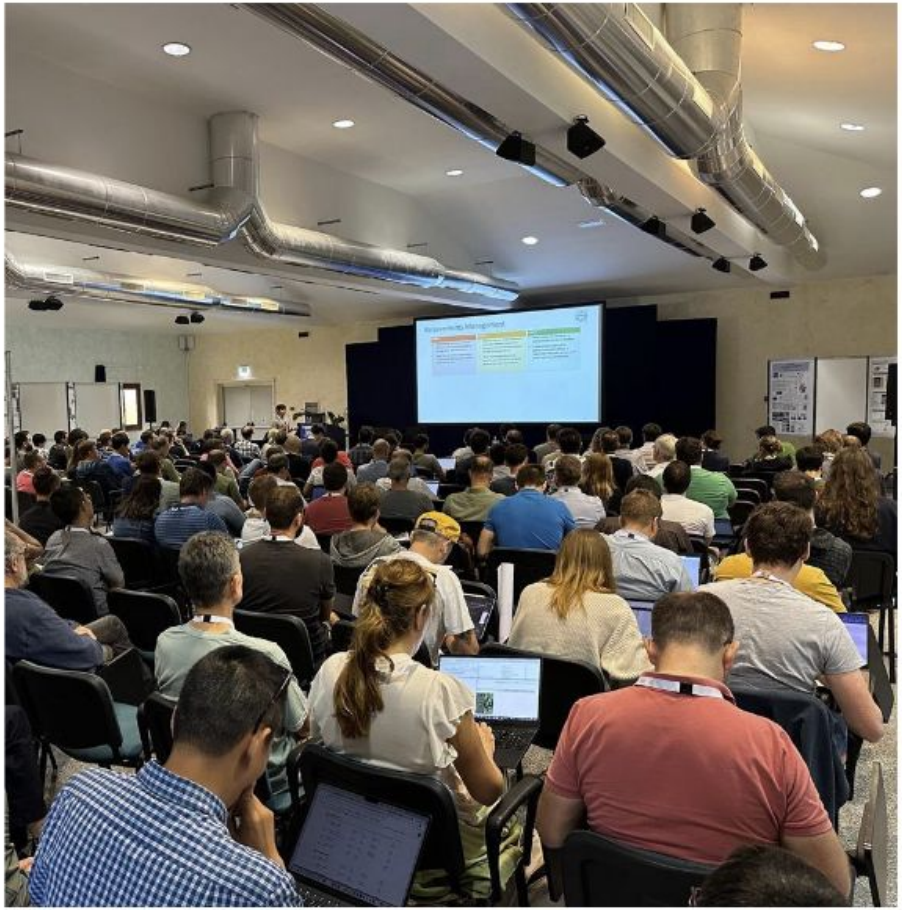


TWEPP 2023

Food & drinks

Photos stolen from closeout slides

Familiar faces

Quite a few people from RAL and ex-PPD'ers

- Dave Newbold
- Luigi Calligaris, Luis Ardila
- Weiming Qian
- Marcus French, Mark Willoughby
- ASIC design team

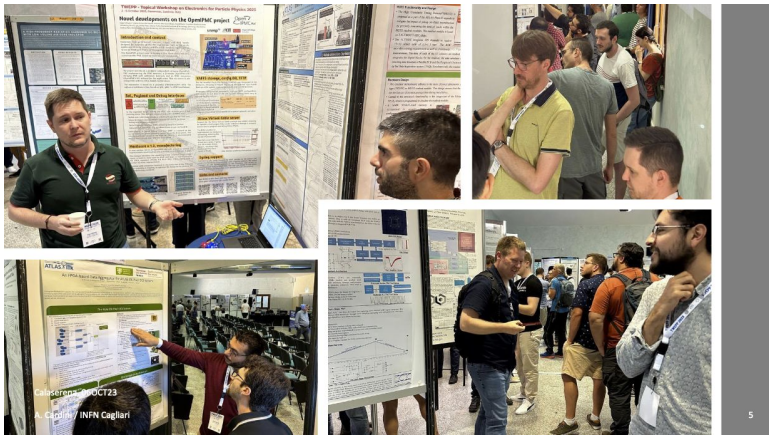

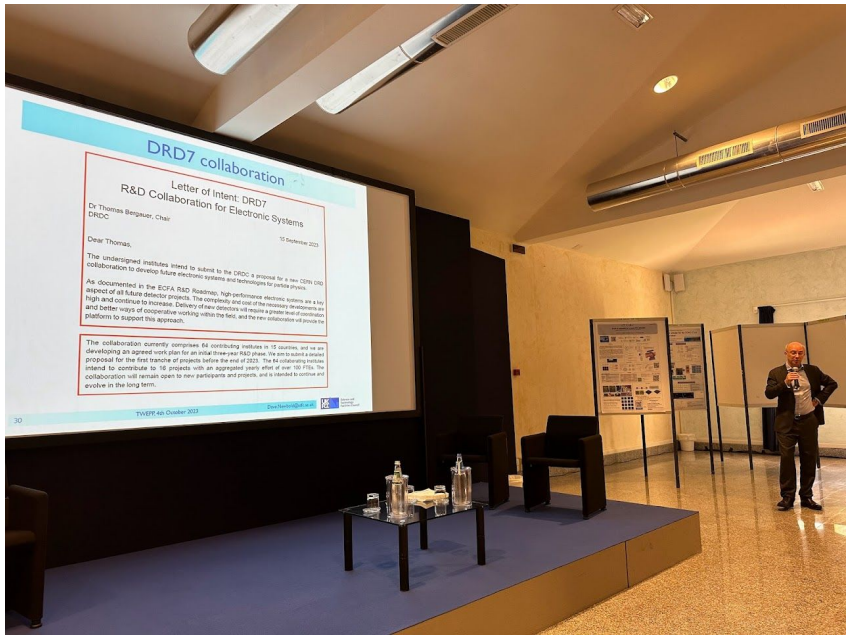

Onto the ~~physics~~ electronics …

Many, many talks & posters ⇒ this summary necessarily focuses on what I found interesting (and could understand)

Mostly go chronologically:

- Current experiments
- HL-LHC upgrades
- R&D for future experiments
  - DRD / DRD-ish

Current experiments

CMS Level-1 Trigger: Anomaly detection

CMS Level-1 Trigger selects 100kHz of bunch crossings, making decision within 3.8 µs

- FPGA-based boards connected by optical links
- Calorimeter & muon subsystems: Perform reconstruction of relevant objects
- Global trigger: Applies menu of trigger paths = single-/multi-object kinematic & quality cuts

Recent development: Detect ‘anomalous events’, by implementing autoencoder neural network in Global Trigger

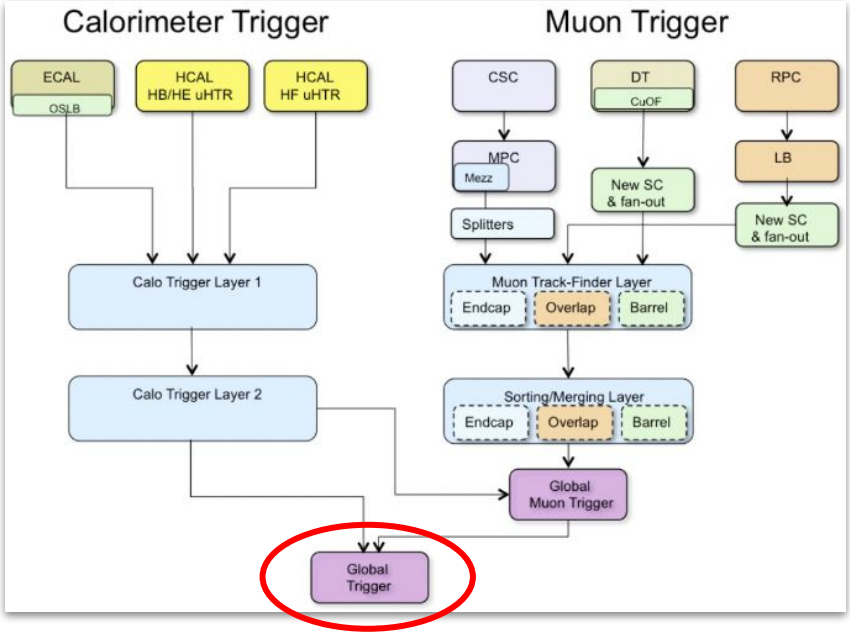

CMS Level-1 Trigger: Anomaly detection

Why Anomaly Detection?

Problem:

Traditional trigger strategies rely on a priori knowledge of signal or generic kinematic selections.

What if we miss new physics because we don't have the right trigger?

Solution:

Triggering on "anomalousness" offers an answer that is both

1. Signal agnostic - Applicable to signatures that we have not had the foresight or person-power to target specifically
2. Highly sensitive – Can boost signal efficiency to signatures limited by L1 trigger bandwidth

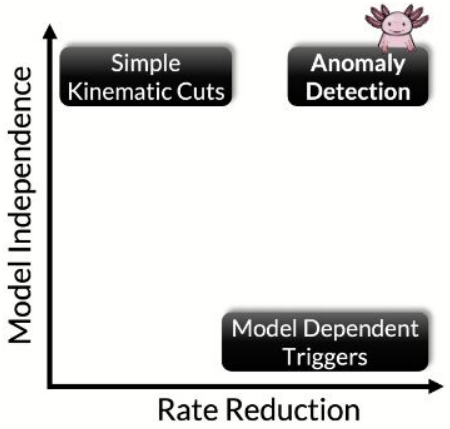

CMS Level-1 Trigger: Anomaly detection

What is *AXOL1TL*?

Anomaly eXtraction Online Level-1 Trigger aLgorithm

- Variational autoencoder (VAE) trained on real unbiased data to detect outliers
- Information bottleneck created by small-dimensional latent space enforces efficient encoding ⇒ learning
- Calculated from standard Global Trigger (μGT) quantities
  - (pT, η, ϕ) hardware integer inputs from: 1 p_T^{miss}, 4 e/γ, 4 μ, and 10 jets

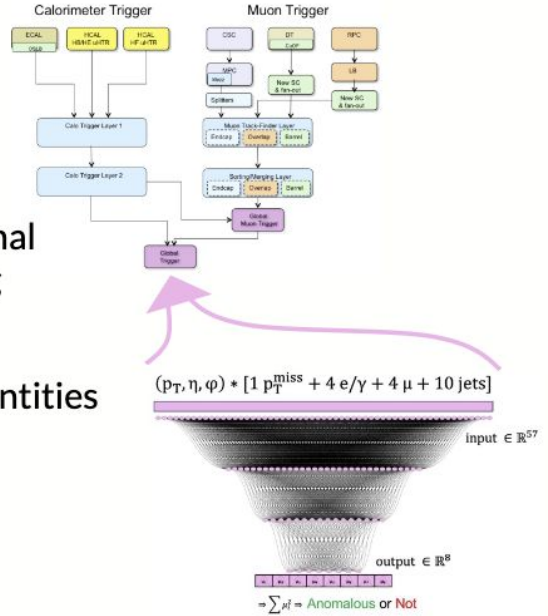

CMS Level-1 Trigger: Anomaly detection

Model Design

Level-1 Trigger constraints informed design

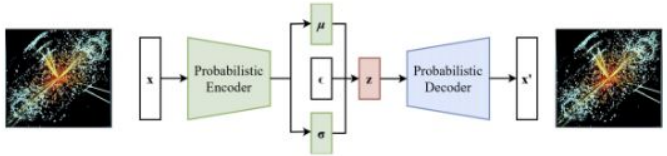

$$\text{Loss} = (1-\beta)\|x-\hat{x}\|^2 + \beta\frac{1}{2}(\mu^2+\sigma^2-1-\log\sigma^2)$$

Reconstruction term — Full regularization term

CMS Level-1 Trigger: Anomaly detection

Model Design

Level-1 Trigger constraints informed design

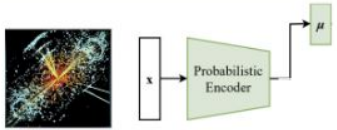

$$\text{Loss} = \cancel{(1-\beta)\|x-\hat{x}\|^2} + \beta\frac{1}{2}(\mu^2 + \cancel{\sigma^2 - 1 - \log\sigma^2})$$

Reconstruction term — Full regularization term

- Remove decoder network
  - Significant latency & resource savings, minimal performance degradation
- Remove latent σ term from loss calculation
  - Saves even more on timing, negligible performance degradation

CMS Level-1 Trigger: Anomaly detection

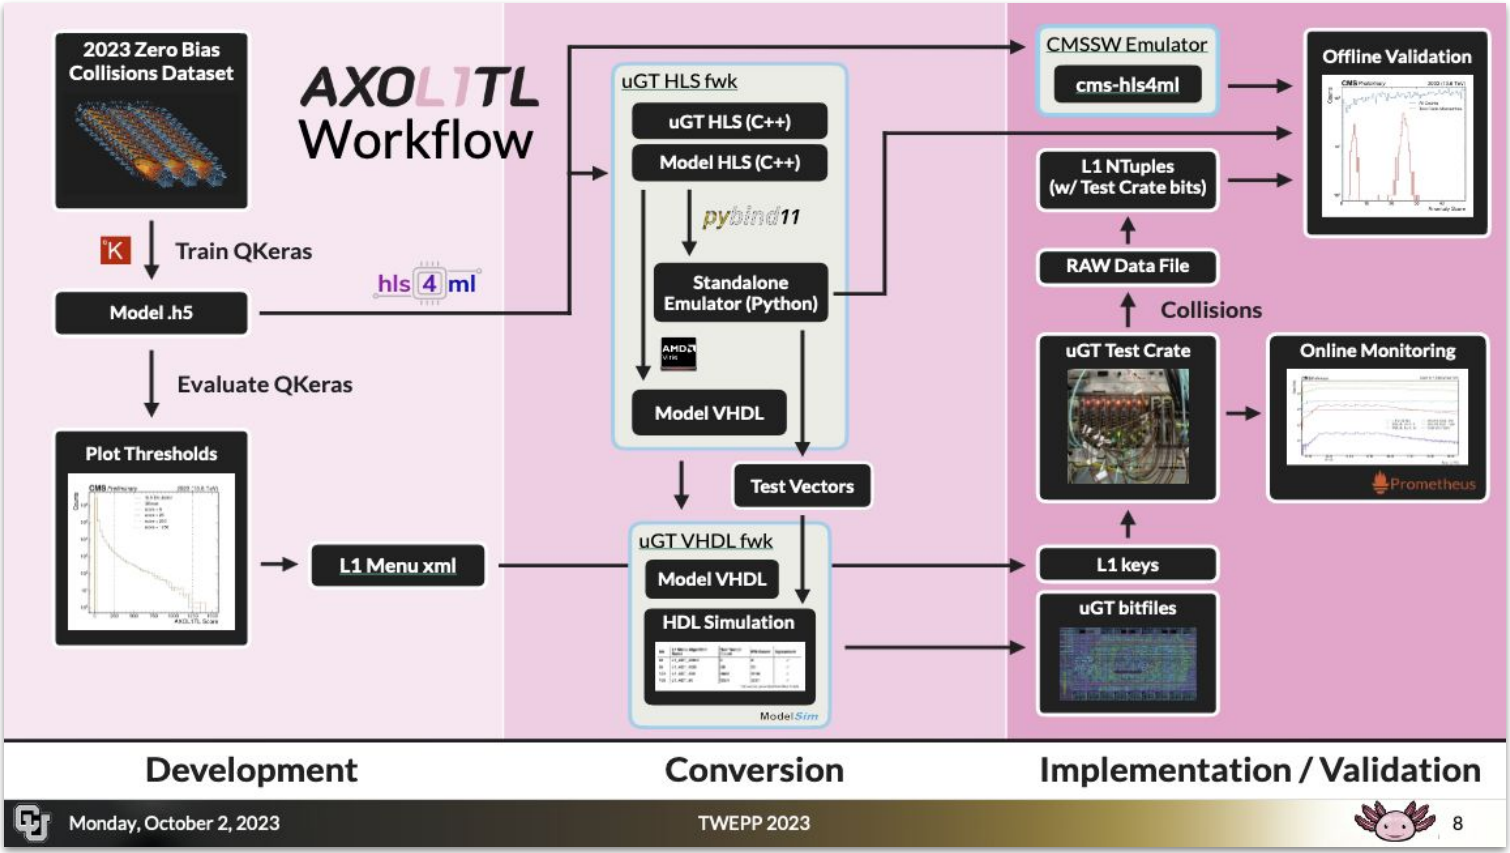

CMS Level-1 Trigger: Anomaly detection

Test Crate Validation

- For certain runs, Test Crate decisions are recorded in 2023 data files
  - Use these bits to validate emulation and show rate agreement
- Minimal (~1%) mismatches between trigger hardware and emulation
  - Mismatches clustered near decision boundaries, most likely due to rounding issue

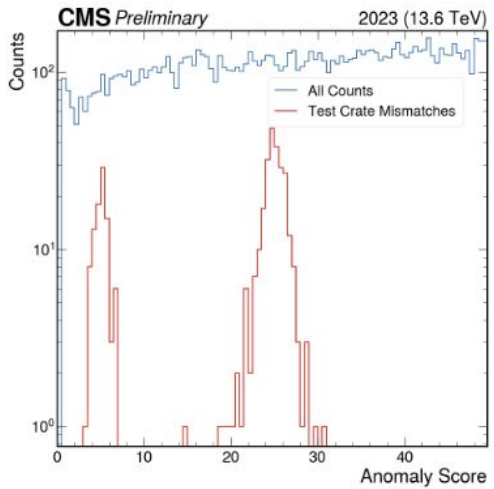

L1 Menu Algorithm Name	Test Crate Count	Standalone Emulator Count	Mismatches
L1_ADT_20000	1	1	0
L1_ADT_4000	742	741	19
L1_ADT_400	21236	21229	253
L1_ADT_80	25468	25481	93

GNNs in an FPGA @ sPHENIX

sPHENIX experiment

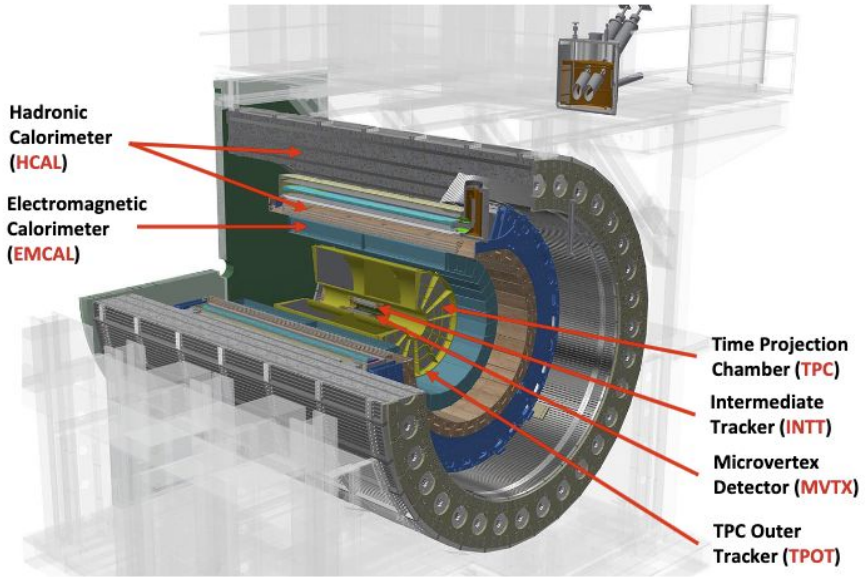

- Located at RHIC accelerator at BNL (USA)
- ~56 MHz accelerator clock with ~9.3 MHz BC
- Running period 2023-2025
- ~4m long, ~5m high, 1000 tons
- Tracking detectors (MVTX, INTT, TPC, TPOT) and calorimeters (EMCAL, HCAL)
- 1.4 T Magnetic Field, $|\eta| \leqslant 1.1$
- Tracking detectors capable of streaming readout, but unable to save all TPC data.
- 15 kHz designed Trigger Rate

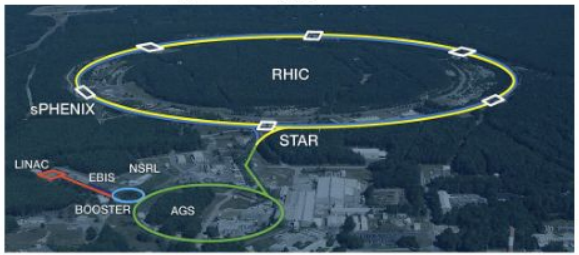

GNNs in an FPGA @ sPHENIX

Motivation – Heavy Flavour

- Integrate the AI-based heavy flavour trigger system demonstrator into the sPHENIX experiment for p+p run in 2024 to R&D its feasibility, requirements, and constrains
  - Heavy-flavour (HF) events are very rare ~1% of Minimum Bias (MB) events at RHIC energy
  - RHIC collision rate is around 2-3 MHz, sPHENIX readout 15 kHz (DAQ - 300 Gb/s)
    - Trackers are Streaming Readout (SRO) capable, but can't save all TPC data
  - 10% trigger-enhanced SRO increases HF MB rate ~ 300 kHz
  - **ML HW tagging aims to sample remaining 90% of the luminosity using the tracklet reconstruction from the silicon trackers**

- The aim is to deploy future system on Electron-Ion Collider (EIC)
  - AI-based electron tagging with streaming readout to identify the (non)interesting Deep-Inelastic-Scattering (DIS) processes in the e+p/A collisions.
    - based on the measured scattering electron energy and direction

GNNs in an FPGA @ sPHENIX

The ML algorithm – TrackGNN

ECML PKDD 2022, Sub 1256

- Based on Graph Neural Network (GNN)
  - Detector and physics knowledge improves prediction
  - Based on PyTorch and PyTorch Geometric
- Initial training on simulated data from MVTX and INTT
  - On GPU - NVIDIA Titan RTX, A500, and A6000
- Topological selection of HF signals on FPGA
  - Tracking and clustering must be done on FPGA
- Beam-spot and anomaly detection on GPU based feed-back system

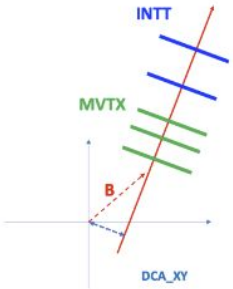

- We propose a novel method to treat the events as track graphs instead of hit graphs. This method is driven by the physics (transverse momentum)
  - Estimate momentum based on silicon hits -> 15% improvement on trigger decision

GNN in an FPGA @ sPHENIX

Generation of the GNN IP core – two parallel efforts

1. Team lead by the Georgia Institute of Technology (GIT)
   - Direct translation of the sPHENIX TrackGNN model to IP using HLS
   - Model
     - 5 layers, each layer: 64 dim 4 layers for node and 64 dim 4 layers for edge embedding
   - Goal: 100-200 nodes, 200-500 edges
   - Implementation
     - 100 nodes, 140 edges
     - Measured Start-to-end latency
       - 150 us @ 130 MHz, 130 us @ 180 MHz
     - Still needs 10-20x speedup!

Target: 5µs

Utilization (Alveo U280)	
LUT	308K (23.7%)
FF	378K (14.5%)
BRAM	1025 (50.8%)
DSP	1426 (15.8%)

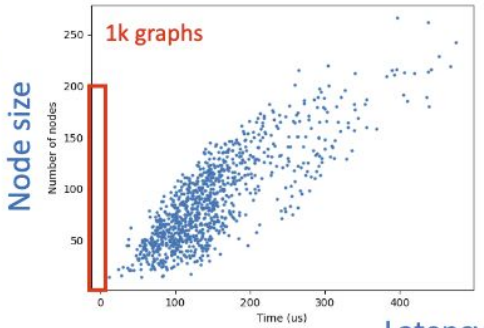

Latency

   - Fast-paced development 380 us (25th August) -> 150 us (4th September) @ 130 MHz
     - Attempts to increase clock to 300 MHz failed on timing constrains
     - Detailed latency breakdown and parallelism exploration ongoing
     - Might require model changes

Close discussion between model developers and FPGA engineers

GNN in an FPGA @ sPHENIX

Generation of the GNN IP core – two parallel efforts

arXiv:2112.02048
arXiv:2103.05579

2. Team lead by the Massachusetts Institute of Technology (MIT) and Fermilab (FNAL)
   - Based on High Level Synthesis for Machine Learning (hls4ml), a generalized python framework for machine learning inference in FPGAs

- Third main upgrade underway, focusing on 3 examples
  - Example 1: Tri-muon reconstruction with the LHC (muon endcaps)
  - Example 2: Heavy flavor tracking at sPHENIX
  - Example 3: Silicon strip tracking at LHC

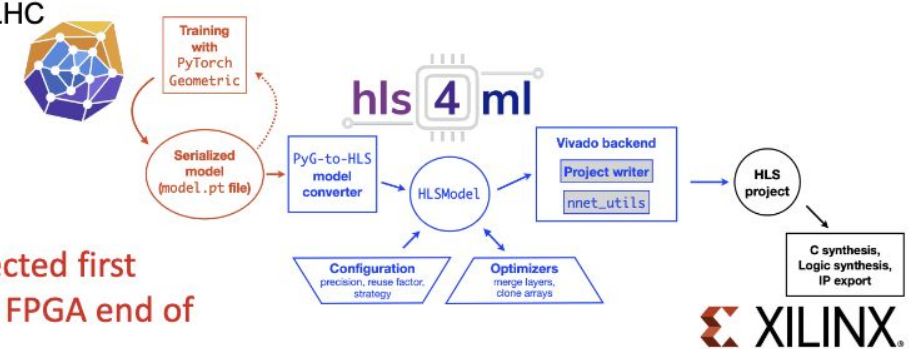

Initial translation just started, expected first version of the TrackGNN model on FPGA end of October 2023

HL-LHC upgrades

Serenity

A common solution

- Single backend board design used in tracker, HGCal, level-1 trigger, MTD, BRIL. **> 700pcs**
- VU13P FPGA with 124x bidirectional optical links (25Gb/s for backend)

Beyond the hardware …

- Boards provided with common infrastructure FW & SW — everything apart from the system-specific algorithm firmware
  - Focus of effort at RAL

Collaboration of 8 institutes (UK, Germany, France, Italy, India, China) formalised: Serenity consortium

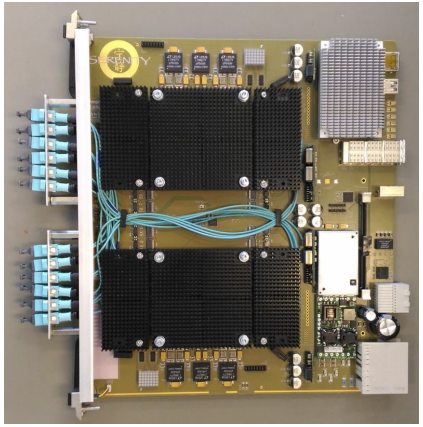

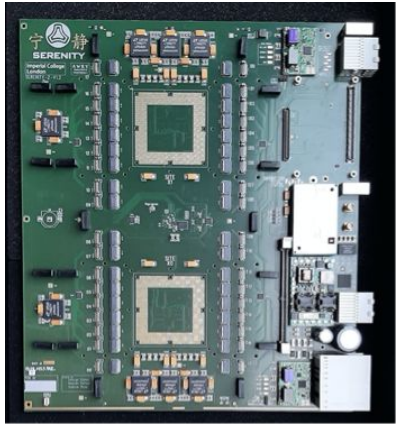

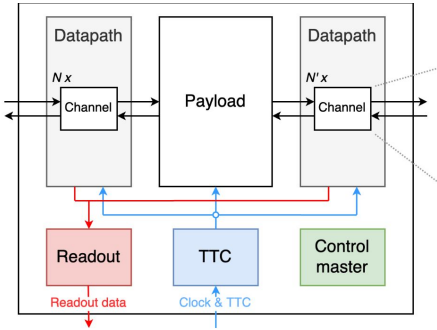

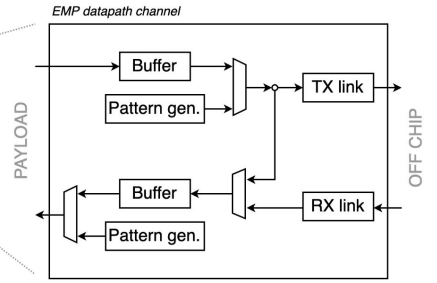

SERENITY

Lessons learned while developing the Serenity-S1 ATCA card

Torben Mehner on behalf of the Serenity consortium
and the CMS Tracker group

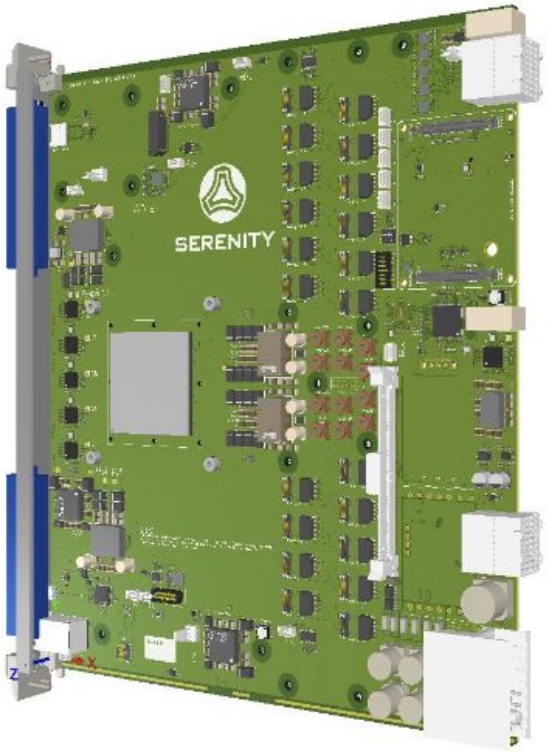

5th October 2023

Serenity

Board Overview

SERENITY

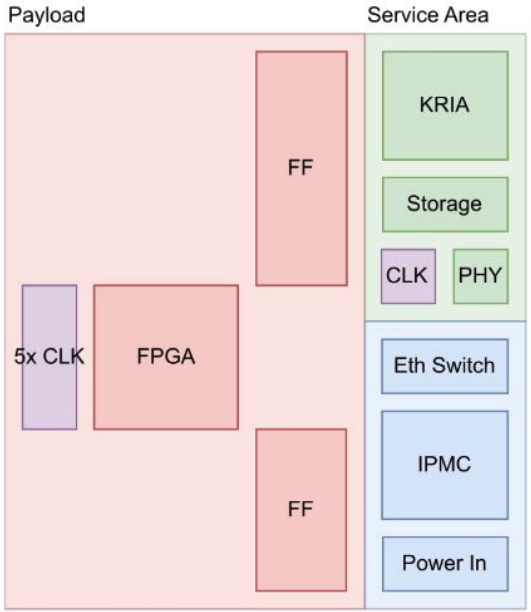

- Board Infrastructure
  - Xilinx KRIA SoM
  - Clock, power, PHY
  - SD, SSD
- ATCA Infrastructure
  - Backplane connectors
  - IPMC (OpenIPMC DIMM module)
  - Power input
  - Ethernet switch
- Payload
  - FireFly optical transceivers
  - VU13P FPGA
  - Clocks

ZynqMP board management: https://iopscience.iop.org/article/10.1088/1748-0221/17/03/C03009
OpenIPMC: https://ieeexplore.ieee.org/document/9465210

Serenity

Component Shortage Mitigation – PLL

SERENITY

- Zero-delay jitter cleaner phase-locked loop
- Skyworks Si5395A not available
- Evaluated ZL30274 (dual PLL) - P. Hazell, S. Baron
- Accumulated jitter <1ps (1 kHz - 10 MHz)
    - Virtex Ultrascale+ requirements met with LDO power supply

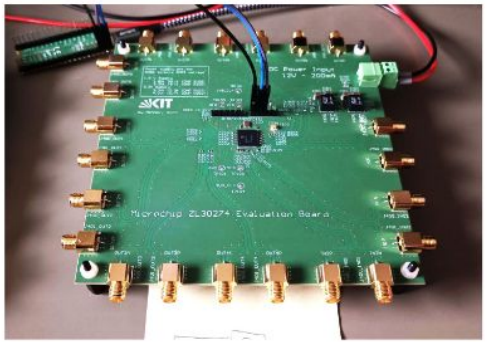

Phase noise measured vs. max phase noise requirements [dBc/Hz]

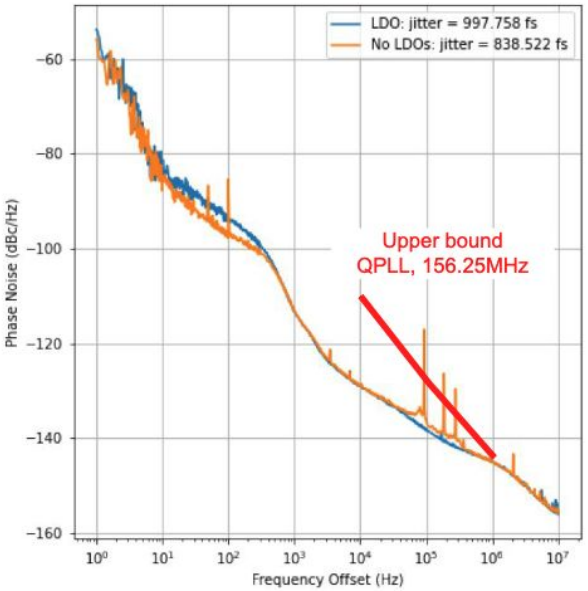

Serenity

Component Shortage Mitigation – ATCA connectors

- TE has stopped producing Zone 1 connectors
- EPT will stop production in May 2024
- Zone 1 connectors are still produced by
  - Positronic VPB series
  - Conec ATC22* series
- Adapt footprint to be multi-vendor compliant
  - Mechanical alignment pins
- CERN has bought all connectors for Serenity production (Zone 1 and Zone 2) + 15%

Samtec FireFlys

Samtec FireFlys

Used extensively in HEP backend electronics

- Custom module for frontend links (i.e. lpGBT)
- 28Gbps parts used for sending data between backend boards
  - 4-channel bidirectional modules
    - Worked fine for several years
  - 12-channel unidirectional modules
    - ✓ Same footprint as frontend parts
    - ➢ Issues encountered with beta TX parts (Jul 22)
      - Some channels died during qualification
      - Traced to VCSEL *(another company)*
    - **Samtec rep. gave comprehensive summary of source of problems at TWEPP**
      - VCSEL fixed; Samtec improved their QA.
    - Testing new parts on Serenitys right now!

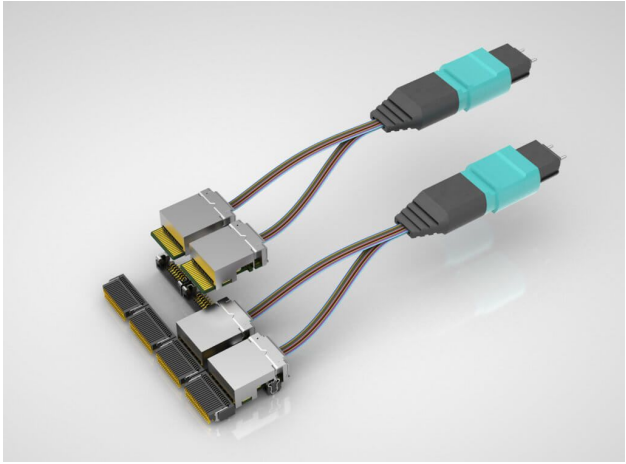

Science and Technology Facilities Council

ATLAS Global Common Module (GCM)

Brookhaven National Laboratory · ATLAS EXPERIMENT · Science and Technology Facilities Council

GCM Technology Choices

- Platform
  - ATCA front board
    - 1 MUX node + 1 GEP node per board
- FPGA
  - Versal Premium adaptive SoC VP1802 for MUX node
  - Versal Premium adaptive SoC VP1802 for GEP/gCTPi node
  - VSVA5601 package
    - 75x75, 0.92mm pitch
    - Minimizing crosstalk
- Optics
  - 20 Firefly 28G parallel optical engines
    - 12 ch/optical engine
- PCB
  - 26 layers
  - Via-in-pad
  - Backdrill

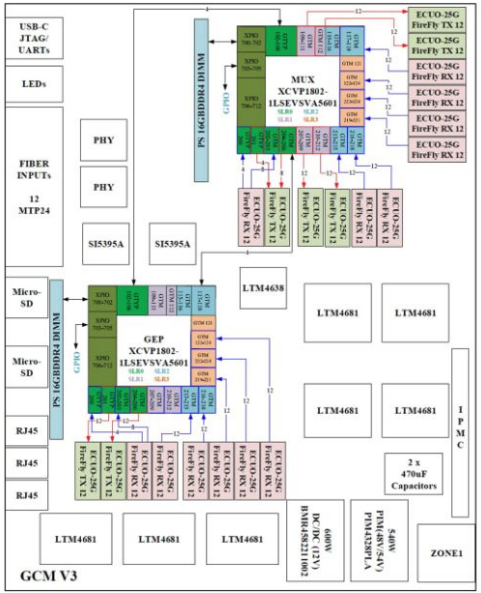

Science and Technology Facilities Council

ATLAS Global Common Module (GCM)

Brookhaven National Laboratory | ATLAS EXPERIMENT | Science and Technology Facilities Council

Power Design Challenges

- Each adaptive SoC VP1802
  - Minimum of 9 power rails
- GCM with two adaptive SoC VP1802
  - ~20 power rails in total

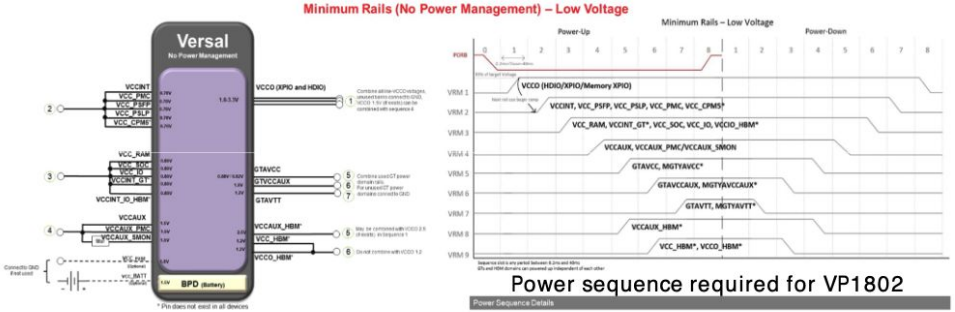

Science and Technology Facilities Council

ATLAS Global Common Module (GCM)

Brookhaven National Laboratory | ATLAS EXPERIMENT | Science and Technology Facilities Council

GCM Thermal Simulation

- Heatsink design and simulation are outsourced
  - Firefly optical engine, vertical block

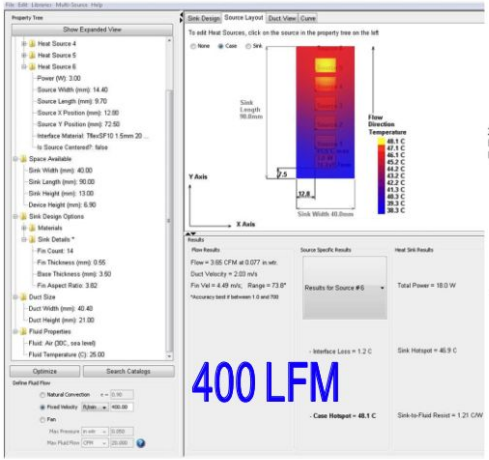

Provided by Alpha Novatech

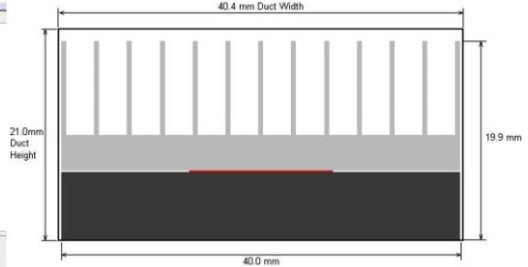

- 6 FFs in vertical row, It is OK to meet 50 ℃ target with 40mm x 90mm x 13mm heatsink.

CMS Level-1 Trigger

- Decides which events should be read out
- Further selection downstream by HLT

Challenges for phase-2

- Luminosity: 1.5 → 7.5 x 10^{34} $cm^{-2}s^{-1}$
- Pile up: ~ 60 → 200
- Design goal: Retain the same thresholds
  - Extend acceptance for some final states (e.g. LLP)

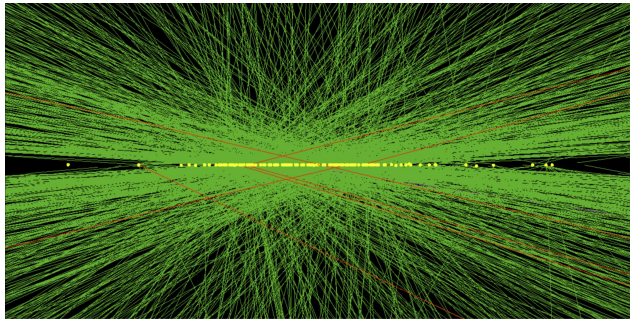

L1 Trigger seeds	Offline Threshold(s) at 90% or 95% (50%) [GeV]	Rate $\langle PU \rangle = 200$ [kHz]	Additional Requirement(s) [cm, GeV]	Objects plateau efficiency [%]
Single/Double/Triple Lepton (electron, muon) seeds				
Single TkMuon	22	12	$\|\eta\| < 2.4$	95
Double TkMuon	15,7	1	$\|\eta\| < 2.4$, $\Delta z < 1$	95
Triple TkMuon	5,3,3	16	$\|\eta\| < 2.4$, $\Delta z < 1$	95
Single TkElectron	36	24	$\|\eta\| < 2.4$	93
Single TkIsoElectron	28	28	$\|\eta\| < 2.4$	93
TkIsoElectron-StaEG	22, 12	36	$\|\eta\| < 2.4$	93, 99
Double TkElectron	25, 12	4	$\|\eta\| < 2.4$	93
Single StaEG	51	25	$\|\eta\| < 2.4$	99
Double StaEG	37,24	5	$\|\eta\| < 2.4$	99
Photon seeds				
Single TkIsoPhoton	36	43	$\|\eta\| < 2.4$	97
Double TkIsoPhoton	22, 12	50	$\|\eta\| < 2.4$	97
Taus seeds				
Single CaloTau	150(119)	21	$\|\eta\| < 2.1$	99
Double CaloTau	90,90(69,69)	25	$\|\eta\| < 2.1$, $\Delta R > 0.5$	99
Double PuppiTau	52,52(36,36)	7	$\|\eta\| < 2.1$, $\Delta R > 0.5$	90
Hadronic seeds (jets,H_T)				
Single PuppiJet	180	70	$\|\eta\| < 2.4$	100
Double PuppiJet	112,112	71	$\|\eta\| < 2.4$, $\Delta\eta < 1.6$	100
PuppiH_T	450(377)	11	jets: $\|\eta\| < 2.4$, $p_T > 30$	100
QuadPuppiJets-PuppiH_T	70,55,40,40,400(328)	9	jets: $\|\eta\| < 2.4$, $p_T > 30$	100,100
E_T^{miss} seeds				
PuppiE_T^{miss}	200(128)	18		100
Cross Lepton seeds				
TkMuon-TkIsoElectron	7,20	1	$\|\eta\| < 2.4$, $\Delta z < 1$	95, 93
TkMuon-TkElectron	7,23	3	$\|\eta\| < 2.4$, $\Delta z < 1$	95, 93
TkElectron-TkMuon	10,20	1	$\|\eta\| < 2.4$, $\Delta z < 1$	93, 95
TkMuon-DoubleTkElectron	6,17,17	0.1	$\|\eta\| < 2.4$, $\Delta z < 1$	95, 93
DoubleTkMuon-TkElectron	5,5,9	4	$\|\eta\| < 2.4$, $\Delta z < 1$	95, 93
PuppiTau-TkMuon	36(27),18	2	$\|\eta\| < 2.1$, $\Delta z < 1$	90, 95
TkIsoElectron-PuppiTau	22,39(29)	13	$\|\eta\| < 2.1$, $\Delta z < 1$ $\Delta R > 0.3$	93, 90

CMS Level-1 Trigger (2)

Phase-2: Meeting the challenge

- Specification: A bit more breathing room …
  - Latency: 3.8 → 12.5 microseconds
  - Max. accept rate: 100 → 750 kHz
- Inputs
  - Finer granularity
  - Particle trajectories from tracker *(NEW!)*
    - ✓ Identify primary vertex
    - ✓ Particle-flow-style reconstruction
- Re-build system using latest technology
  - UltraScale+ FPGAs
    - ✓ More flip flops , LUTs, RAM, I/O …
  - Optical links: 10 → 25Gbps

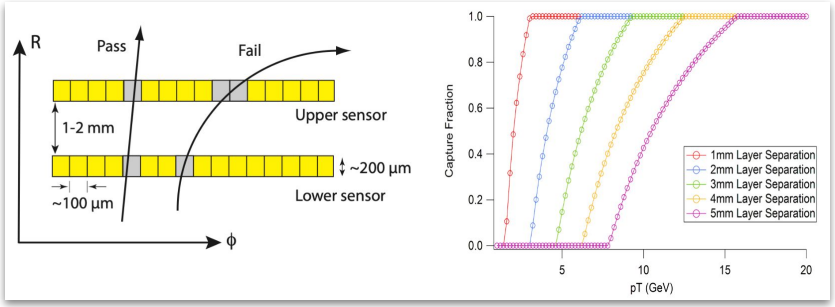

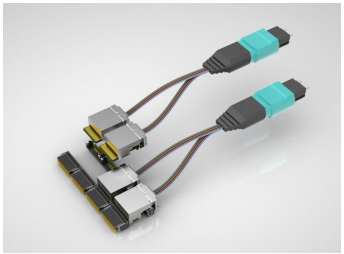

CMS Level-1 Trigger: Phase-2 architecture

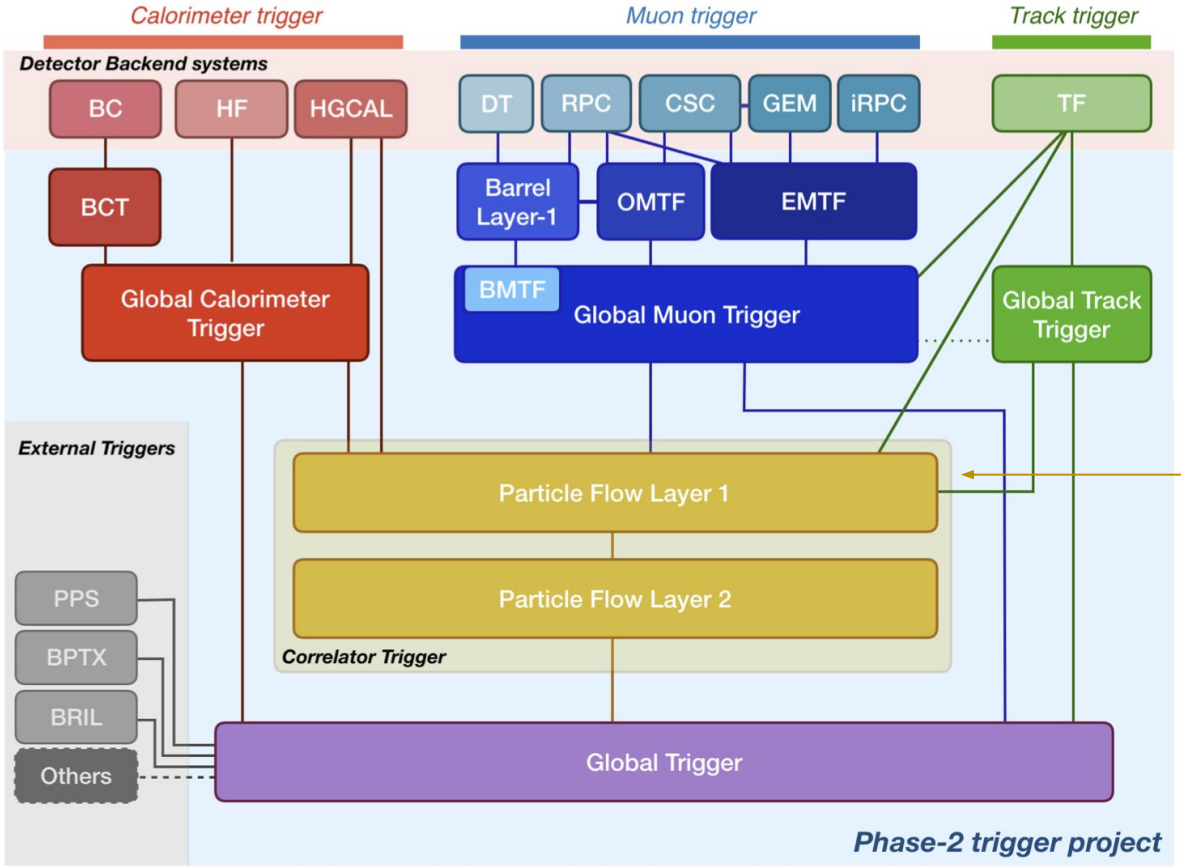

Independent subsystems for calorimeter, muon & tracker inputs

Correlator
Combines information from all detectors using particle-flow-based reco./ID algorithms → electrons, photons, tau, jets & sums

CMS Level-1 Trigger: Integration tests

Science and Technology Facilities Council

The subject of my poster

- Final system: O(200) boards of 4 designs, implementing 20 algos, connected by several thousand links
- Factorise testing via well-defined I/O interfaces
  - Fully validate algorithms first in single-board tests

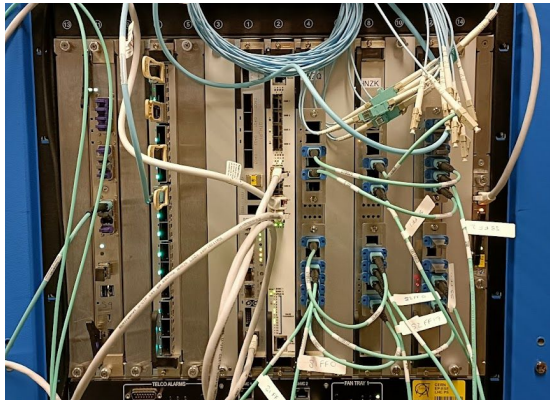

- ✓ Extensive tests of link protocol *(error injection)*
- ✓ Verified most algorithms in single-board tests
- ✓ Several multi-board slice tests performed
- ✓ Latency currently 8.6 µs, less than 9.5 µs target
- ✓ Online SW: Already very mature
  - Reliably controlling several boards — incl. for several of the slice tests

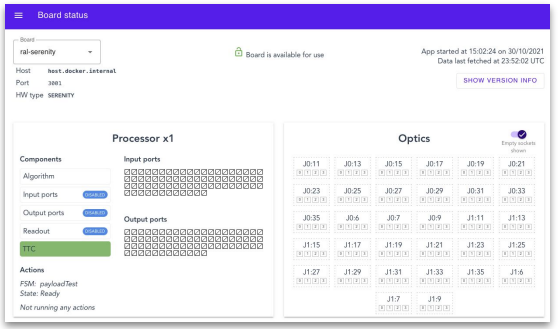


CMS
FPGA Design with High Level Synthesis
Methodology, gains, and pitfalls
Michalis Bachtis
University of California, Los Angeles
On behalf of the CMS Collaboration
TWEPP 2023

Lookup tables

```
const ap_uint<72>  lookup[512] = {
        0,1,2,3,4,5,6,7,8,9,10,11,12,13,14,15,16,17,18,19,20,
        21,22,23,24,25,26,27,28,29,30,31,32,33,34,35,36,37,38,
        39,40,41,42,43,44,45,46,47,48,49,50,51,52,53,54,55,56,
        57,58,59,60,61,62,63,64,65,66,67,68,69,70,71,72,73,74,
        75,76,77,78,79,80,81,82,83,84,85,86,87,88,89,90,91,92,
        93,94,95,96,97,98,99,100,101,102,103,104,105,106,107,
        108,109,110,111,112,113,114,115,116,117,118,119,120,
        121,122,123,124,125,126,127,128,129,130,131,132,133,
        134,135,136,137,138,139,140,141,142,143,144,145,146,
        147,148,149,150,151,152,153,154,155,156,157,158,159,
        160,161,162,163,164,165,166,167,168,169,170,171,172,
        173,174,175,176,177,178,179,180,181,182,183,184,185,
        186,187,188,189,190,191,192,193,194,195,196,197,198,
        199,200,201,202,203,204,205,206,207,208,209,210,211,
        212,213,214,215,216,217,218,219,220,221,222,223,224,
        225,226,227,228,229,230,231,232,233,234,235,236,237,
        238,239,240,241,242,243,244,245,246,247,248,249,250,
        251,252,253,254,255,256,257,258,259,260,261,262,263,
        264,265,266,267,268,269,270,271,272,273,274,275,276,
        277,278,279,280,281,282,283,284,285,286,287,288,289,
        290,291,292,293,294,295,296,297,298,299,300,301,302,
        303,304,305,306,307,308,309,310,311,312,313,314,315,
        316,317,318,319,320,321,322,323,324,325,326,327,328,
        329,330,331,332,333,334,335,336,337,338,339,340,341,
        342,343,344,345,346,347,348,349,350,351,352,353,354,
        355,356,357,358,359,360,361,362,363,364,365,366,367,
        368,369,370,371,372,373,374,375,376,377,378,379,380,
        381,382,383,384,385,386,387,388,389,390,391,392,393,
        394,395,396,397,398,399,400,401,402,403,404,405,406,
        407,408,409,410,411,412,413,414,415,416,417,418,419,
        420,421,422,423,424,425,426,427,428,429,430,431,432,
        433,434,435,436,437,438,439,440,441,442,443,444,445,
        446,447,448,449,450,451,452,453,454,455,456,457,458,
        459,460,461,462,463,464,465,466,467,468,469,470,471,
        472,473,474,475,476,477,478,479,480,481,482,483,484,
        485,486,487,488,489,490,491,492,493,494,495,496,497,
        498,499,500,501,502,503,504,505,506,507,508,509,510,511};

ap_uint<73> lookupAndAdd(const ap_uint<32>& a,const ap_uint<9>& addr) {
    return a+lookup[addr];
}
```

- An adder that reads a LUT and and adds a constant
- A LUT of 512x72 bits instantiated
  - In Ultrascale architecture = 1 BRAM
- At 100 MHz
  - Latency of 1 cycle [to read the ROM]
- At 500 MHz
  - Latency of 2 cycles [automatic register in the output of ROM]
- When increasing the width of more than 72 or if we add more than 52 entries
  - Automatically instantiates more BRAMs

Playing with the pipeline, reusing logic

- In Time Multiplexed designs: implement one algorithm core and feed several chunks of data. As an example let's assume:
  - 6 sets of three 16 bit numbers (a,b,c) are arriving in the system
  - For each set we need to calculate a+b*c [with a DSP]
- We have a fully pipelined option and an option to reuse logic

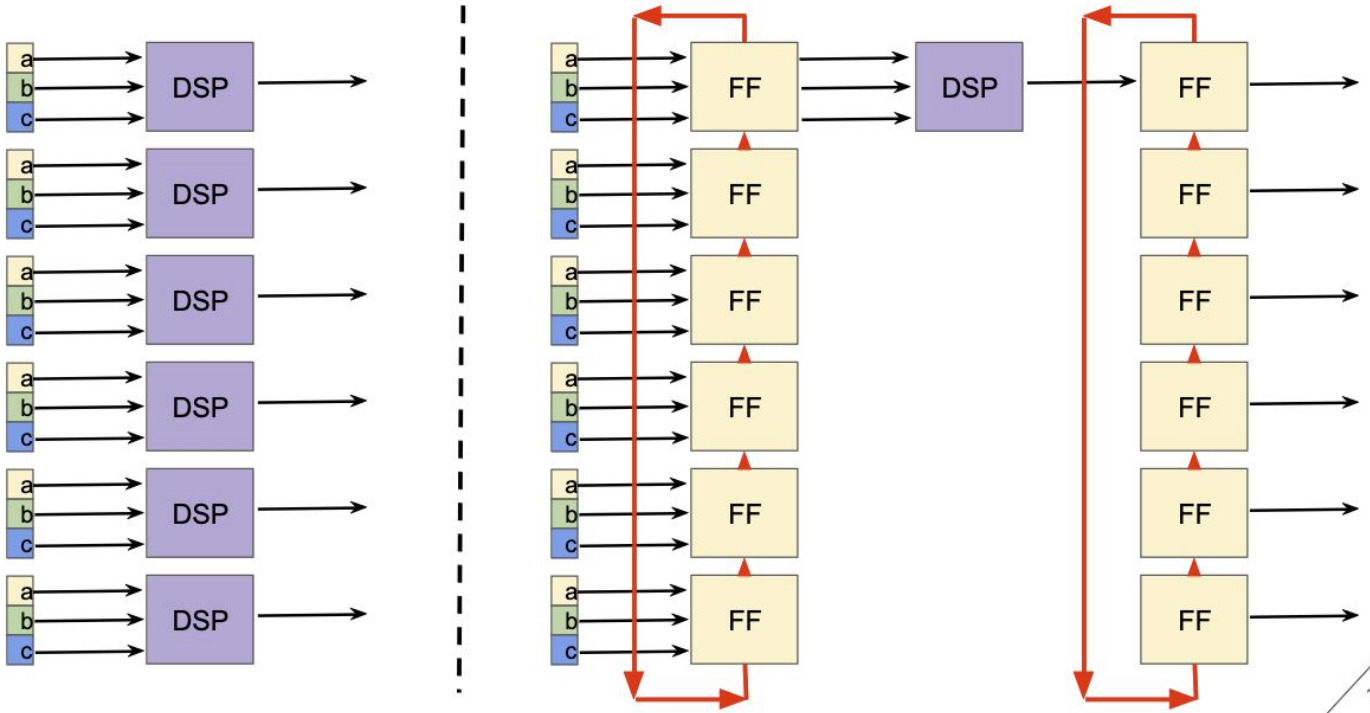

Pitfall. What did the compiler really build?

- We did not specify anything in the code to force a shift register…
  - We could have “helped” the compiler by mimicking the array manipulations in C
- In fact the design was implemented with muxes

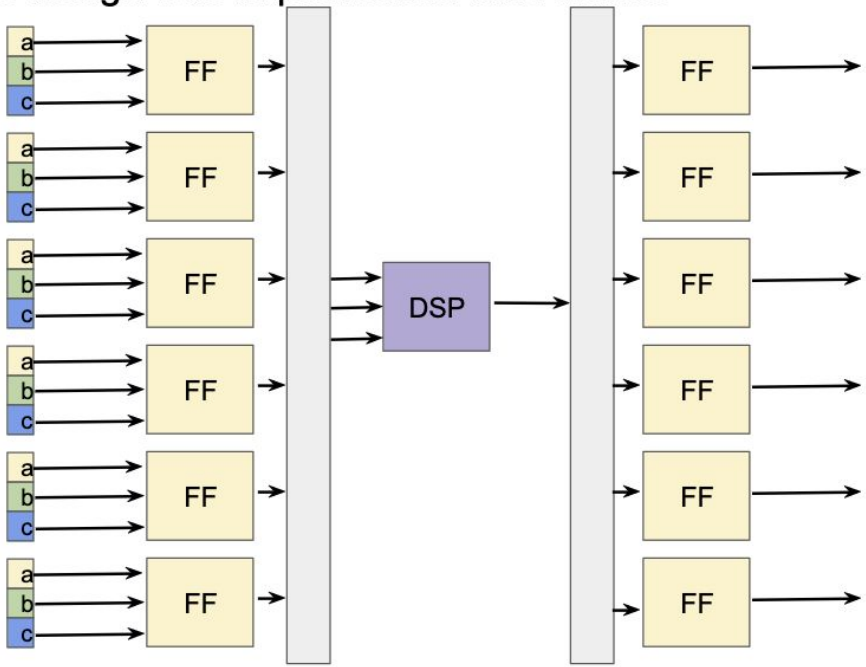

- Does it matter?
  - Sometimes it does because in large designs the routing congestion could get the implementation to fail..

HLS algorithm + HDL data management+glue

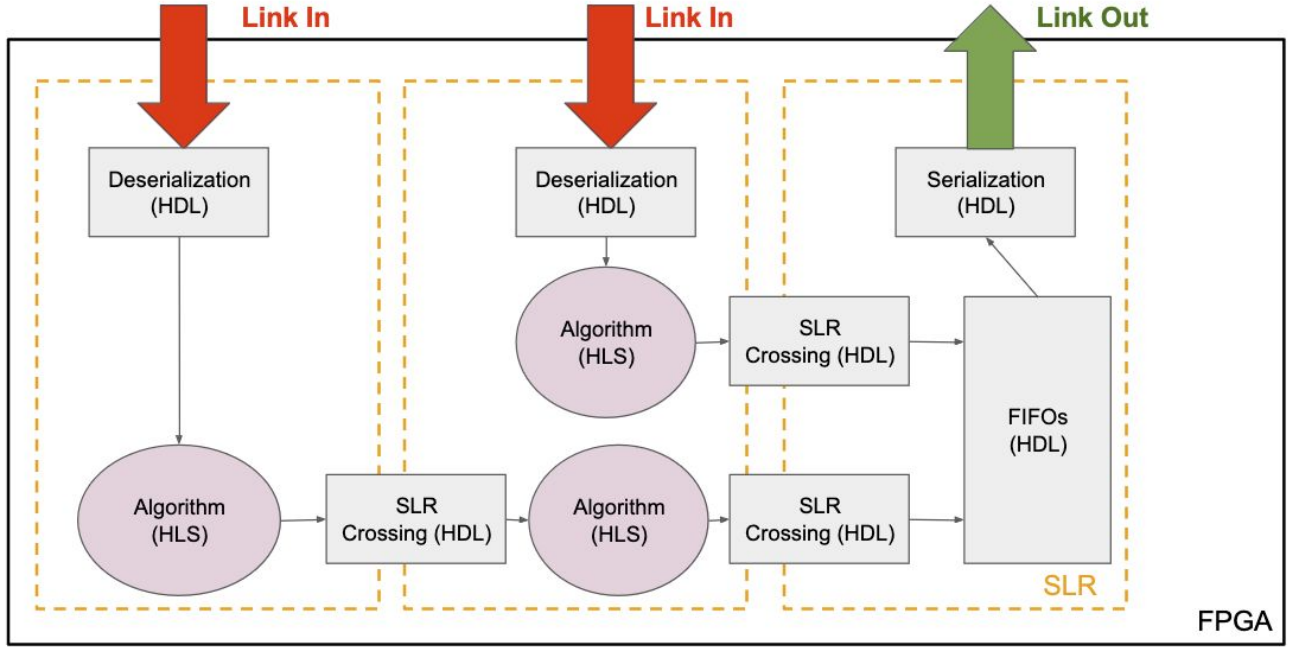

- Optimal results obtained by both HLS cores and HDL
- Algorithms → HLS
- Data management, SLR crossing etc → HDL

CMS 40MHz scouting

The CMS Level-1 trigger upgrades for CMS Phase-2

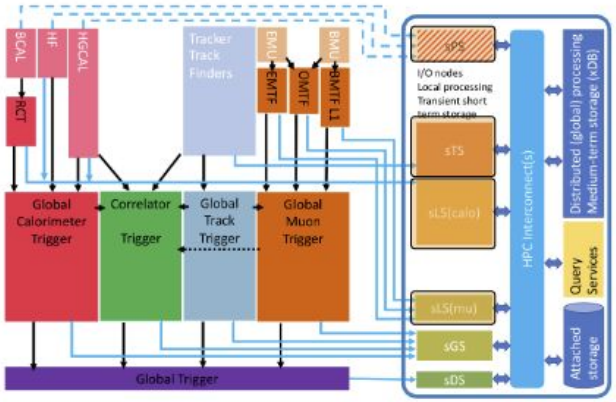

Increased budget latency and rate:

- 3.8 μs → 12.5 μs
- 750 kHz of L1 output

Advanced object reconstruction on FPGA:

- **Global Calorimeter Trigger (GCT)** and **Global Muon Trigger (GMT)** (higher granularity)
- **Global Track Trigger (GTT)** (tracker tracks, vertex finding)
- **Correlator Trigger (CL2)** (Particle Flow)
- **Global Trigger (GT)** (with more complex algos)
- Resolution similar to offline level

Level-1 trigger Data Scouting (L1DS) at 40 MHz LHC bunch crossing rate

- Collect and store the reconstructed particle primitives of the L1 processing chain at the full bunch crossing rate
- Enable study of exotic signatures that cannot be fit into the trigger budget
- **Global Trigger decisions ⇒ sDS**
- **Correlator, Global Track, Global Calorimeter and Global Muon Trigger ⇒ sGS**
- **Can later be extended to include other systems in later stages ⇒ sLS**

CMS 40MHz scouting

The physics potential of a Level-1 trigger Data Scouting system

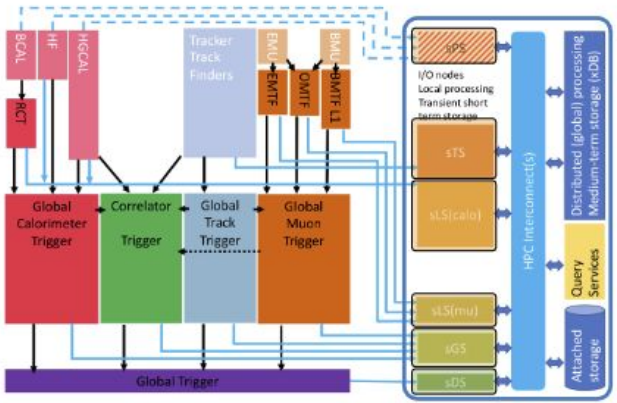

Phase-2 L1DS main physics plans

- Use when possible to study with L1 resolution
- High combinatorics: $W \to 3\pi$, $D_s\gamma$, $H \to \rho\gamma, \phi\gamma$, …
- High rate: multiple soft (b-)jets, displaced (soft) leptons, …
- Heavy Stable Charged Particles (HSCPs) over multiple BXs

Monitoring at the bunch crossing rate

- L1 trigger pre-/post-firing without special configurations
- Per-bunch luminosity measurements

Level-1 trigger Data Scouting (L1DS) at 40 MHz LHC bunch crossing rate

- Collect and store the reconstructed particle primitives of the L1 processing chain at the full bunch crossing rate
- Enable study of exotic signatures that cannot be fit into the trigger budget
- Global Trigger decisions ⇒ sDS
- Correlator, Global Track, Global Calorimeter and Global Muon Trigger ⇒ sGS
- Can later be extended to include other systems in later stages ⇒ sLS

CMS 40MHz scouting

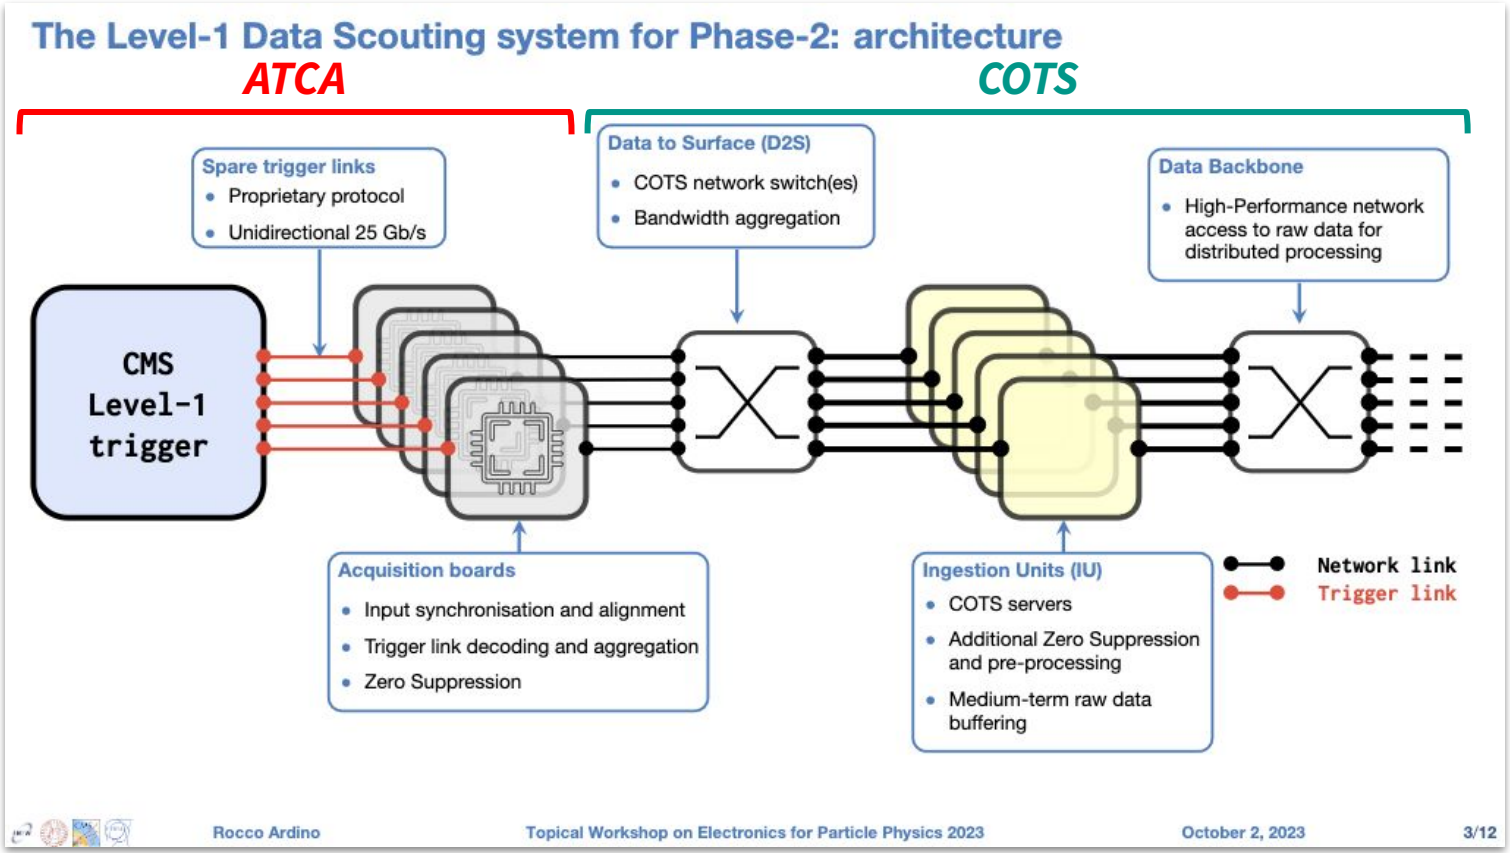

CMS 40MHz scouting

Run-3 demonstrator of the L1DS

For LHC Run-3, L1DS demonstrator to readout multiple sources of the CMS L1 trigger:

- Very heterogenous system
- 3 boards (KCU1500, SB-852, VCU128), different output technologies (DMA, TCP/IP)

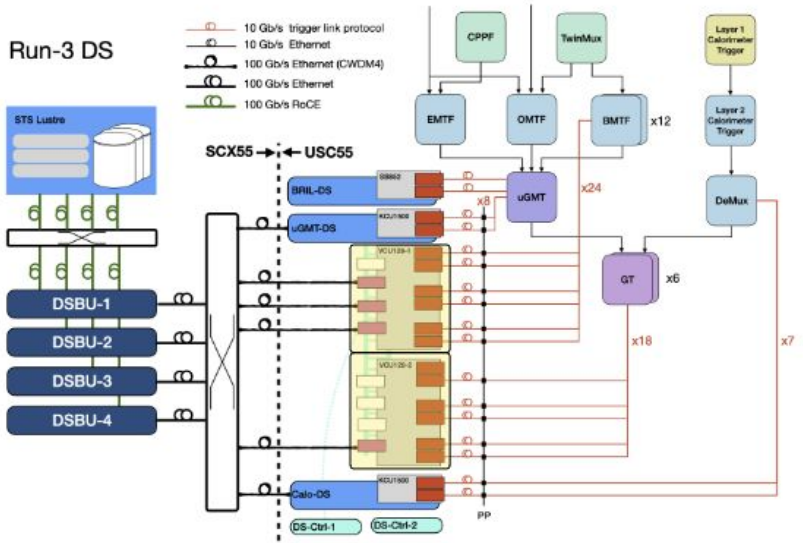

Global Muon Trigger (μGMT) ⇒ GMT muons
- 8 × 10 Gb/s links from 1 trg processor → 1st KCU1500 and SB-852
- **DMA output** through PCIe gen3 to host server

Calorimeter Trigger (DeMux) ⇒ jets, e/γ, τ_h, MET
- 7 × 10 Gb/s links from 1 trg processor → 2nd KCU1500
- **DMA output** through PCIe gen3 to host server

Barrel Muon Track Finder (BMTF) ⇒ barrel stubs
- 24 × 10 Gb/s links from 12 trg processors → 1st VCU128
- **TCP/IP output** to surface (3 × 100 Gb/s links)

Global Trigger output (μGT) ⇒ decision bits
- 18 × 10 Gb/s links from 6 trg processors → 2nd VCU128
- **TCP/IP output** to surface (1 × 100 Gb/s link)

CMS 40MHz scouting

Xilinx VCU128 boards setup in Run-3 L1DS demonstrator

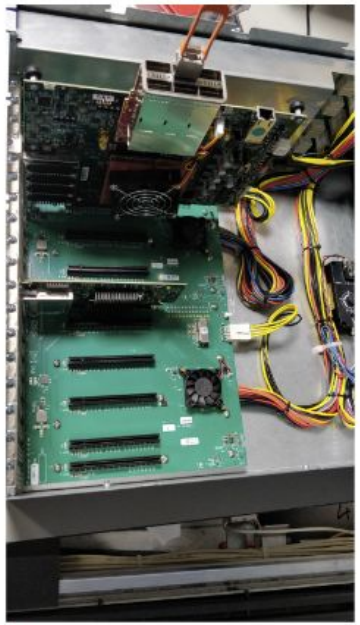

(a) Test setup in CMS DAQ laboratory

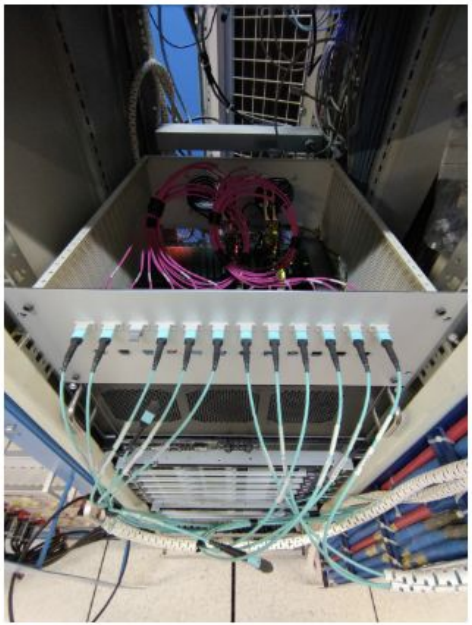

(b) Point 5 service cavern production system

Setup for VCU128 scouting boards:

- *One Stop Systems* **PCIe bus** to accomodate multiple PCIe boards
- 2 × 5 PCIe (3.0 × 16) Slot Expansion
- Control server connected to PCIe bridge for control and monitor of boards

Production system in P5 service cavern:

- 1st VCU128: connected to 12 × BMTF processors
- 2nd VCU128: connected to 6 × GT processors
- Both boards on same PCIe tree
- **Output links from CMS service cavern → surface (~100m)**

CMS 40MHz scouting

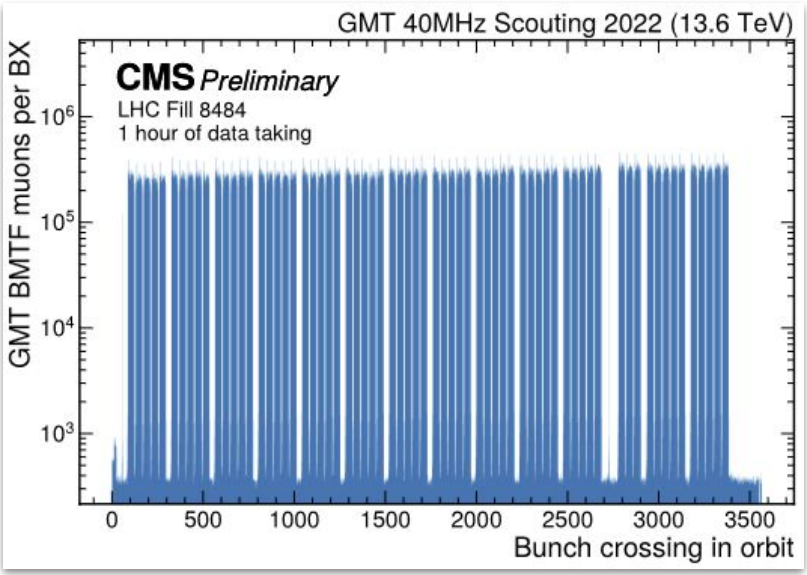

GMT muon occupancy per bunch crossing for barrel muons, matching filling scheme.

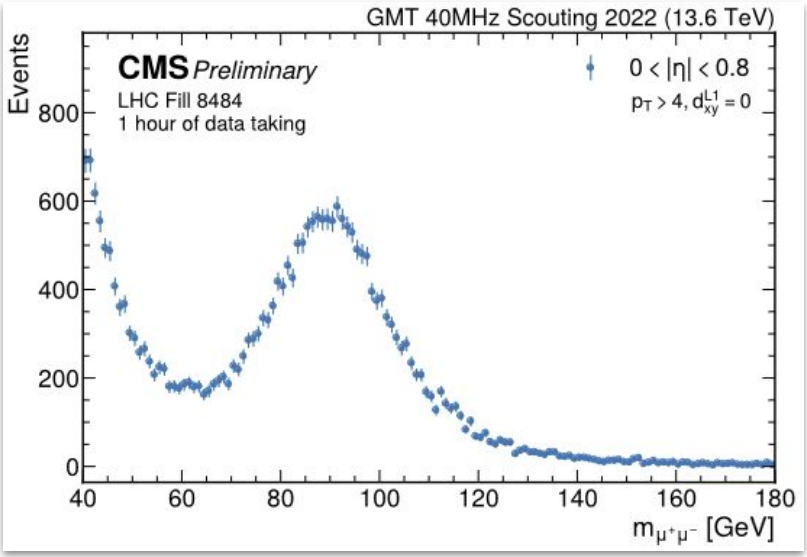

Di-muon invariant mass distribution, after recalibration of L1T muon pT estimate

Longer-term R&D

Electronics for Colliders: What's Next?

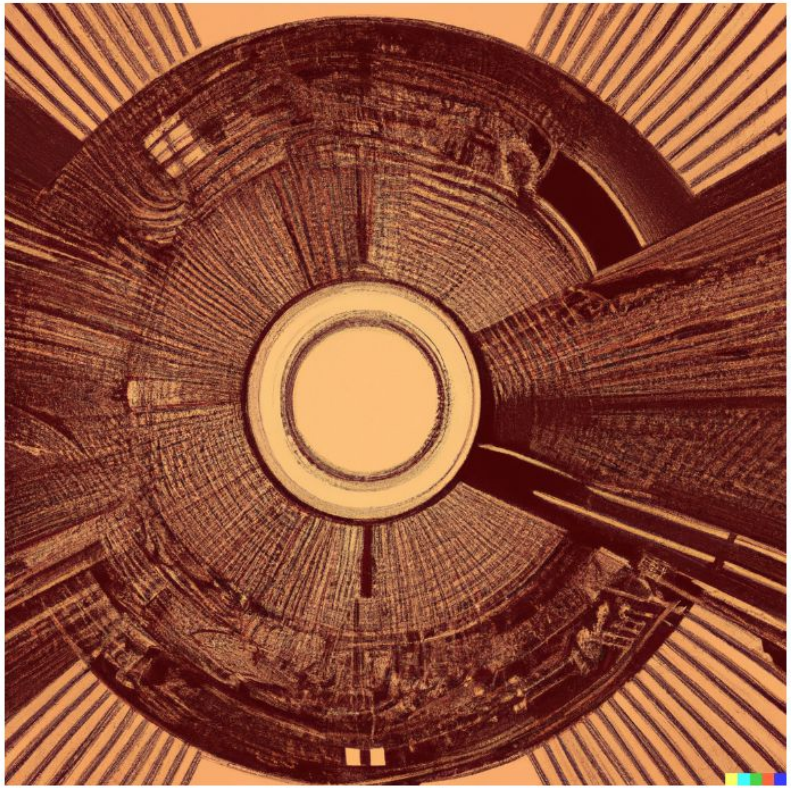

DALLE-2: "Electronics for a future particle collider"

Dave Newbold, STFC

Science and Technology Facilities Council

Types of colliders

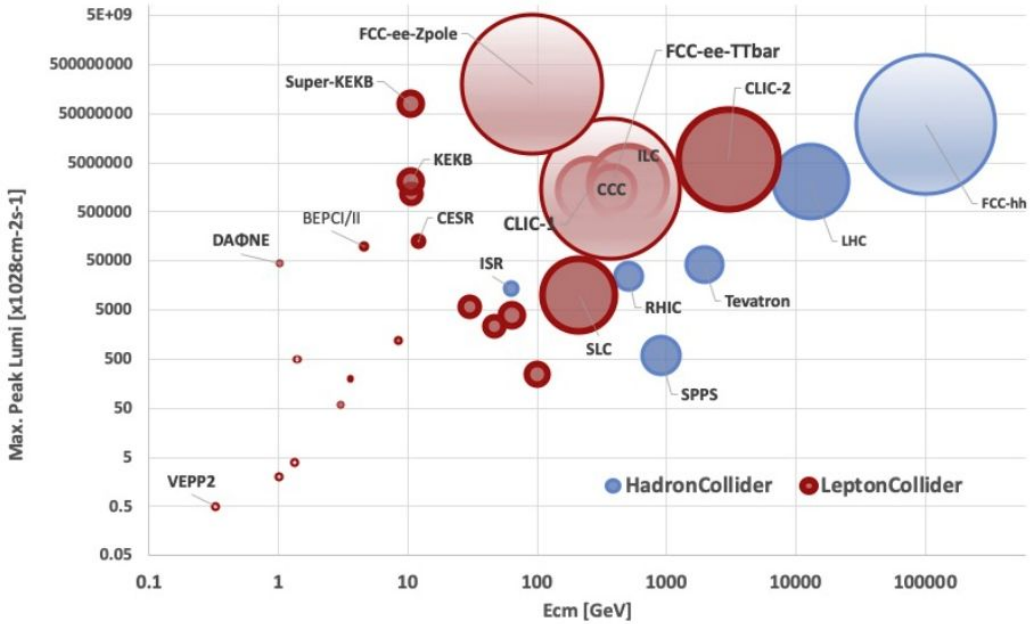

http://arxiv.org/abs/arXiv:2209.04009

- There are new games in town since 2020 (actually not so new)
  - Muon collider (Higgs factory to multi-TeV)
  - Electron-proton collider in the LHC tunnel
  - Hybrid (plasma / RF) asymmetric electron-proton collider

Themes in detector requirements

- Improved granularity
  - More channels, more data, more multiplexing
  - Reconstruction / data reduction requires distributed data
- Improved precision
  - More data (but improved data reduction possible?)
- 4D and 5D techniques (space + time + energy deposition)
  - More data, need for high-accuracy timing distribution
- Low power -> reduced cooling -> less material
  - Requirement for advanced technology nodes, better integration
- Advanced data handling
  - Sophisticated data reduction at front end to cope with backgrounds
- Design and simulation of full systems
  - Co-optimisation of sensors, electronics, and algorithms

Analysis of R&D Needs

		DRDT	< 2030	2030-2035	2035-2040	2040-2045	> 2045
Data density	High data rate ASICs and systems	7.1		*			
	New link technologies (fibre, wireless, wireline)	7.1					
	Power and readout efficiency	7.1		*			
Intelligence on the detector	Front-end programmability, modularity and configurability	7.2					
	Intelligent power management	7.2		*			
	Advanced data reduction techniques (ML/AI)	7.2					
4D-techniques	High-performance sampling (TDCs, ADCs)	7.3					
	High precision timing distribution	7.3					
	Novel on-chip architectures	7.3					
Extreme environments and longevity	Radiation hardness	7.4					
	Cryogenic temperatures	7.4					
	Reliability, fault tolerance, detector control	7.4					
	Cooling	7.4		*			
Emerging technologies	Novel microelectronic technologies, devices, materials	7.5					
	Silicon photonics	7.5					
	3D-integration and high-density interconnects	7.5		*			
	Keeping pace with, adapting and interfacing to COTS	7.5					

SPS fixed target · Belle II · ALICE LS3 · PIP-II/LBNF/DUNE · ALICE 3 · LHCb (≥ LS4) · ATLAS/CMS (≥ LS4) · EIC · LHeC · ILC (Tracking) · ILC (Calorimetry) · FCC-ee (initial detectors) · CLIC (Tracking) · CLIC (Calorimetry) · FCC-hh (initial detectors) · FCC-eh (initial detectors) · Muon collider

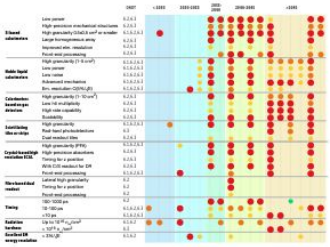

TF6: Calo

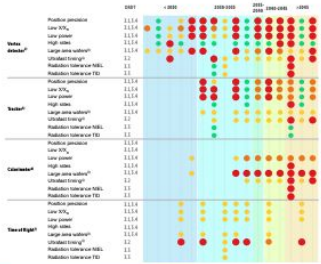

TF3: Solid state

- Must happen or main physics goals cannot be met
- Important to meet several physics goals
- Desirable to enhance physics reach
- R&D needs being met

ASIC design verification for HEP

ASICs for HEP research

Past (LHC era)	Present (HL-LHC era)	Future
• Predominantly analog chips • Relatively simpler digital circuits • Mainly hand-crafted designs • E.g., APV25, PACE, K-chip, FE-I, PSI46	• Mixed signal designs with complex digital circuits • Digital-on-top design methodology • Radiation, power and area were the main drivers • E.g., lpGBT, RD53, MPA, SSA, ABCstar, CIC etc.	• Complex digital circuits to handle data reduction • On chip data processing (AI/ML/RISC-V) • E.g., PicoPix

Key Observations

- **Digital complexity is increasing** in ASICs for HEP experiments
- Technology scaling
  - 250nm -> 130nm -> 65nm -> 28nm -> ??
  - There is a general industry trend that as technology nodes becomes smaller, there is an increasing shift toward digital ASICs
- Design methodologies
  - Full custom designs -> Analog-on-top -> Digital-on-top -> ??
  - It is efficient to implement complex features in digital domain (thanks to EDA tools)
- Design teams
  - Specialized functions
  - Large collaborative efforts (E.g., RD53)

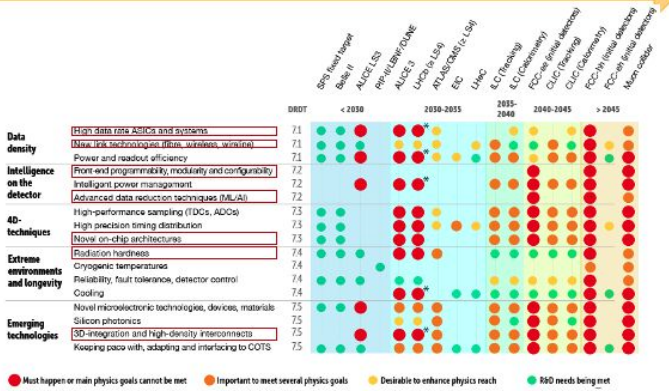

ASIC design verification for HEP

Verification – a bottleneck

- ASIC verification is more complex and time consuming than design
- Verification is resource intensive
  - New skill in HEP community

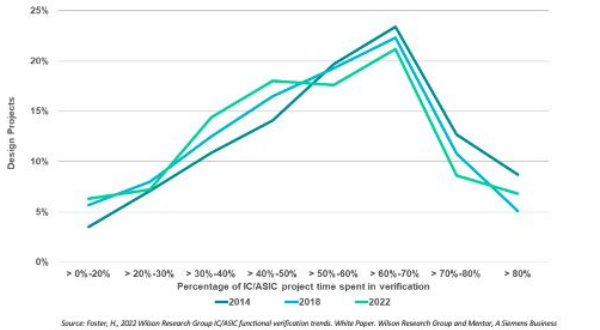

Source: Foster, H., 2022 Wilson Research Group IC/ASIC functional verification trends. White Paper. Wilson Research Group and Mentor, A Siemens Business

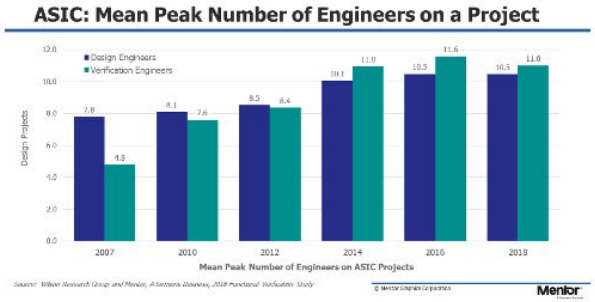

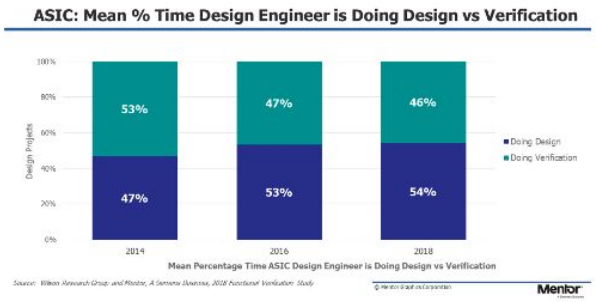

7

ASIC design verification for HEP

Unleashing verification quality and efficiency

- Verification activities in HEP ASIC community
  - Several ASIC projects already have dedicated verification experts to handle verification
  - Future ASICs will require more verification resources
  - We need to lay a strong foundation to ensure that verification activities are efficient and effective
- This talk focusses on processes and practices which improves quality of verification and makes it efficient
  - Improves overall quality of ASIC design project

- Maturity model
  - A set of structured levels that describe how well the behaviors, practices and processes of an organization can reliably and sustainably produce required outcomes
  - Can be use used as
    - **a framework to discuss quality and efficiency**
    - benchmark for comparison
    - an aid for introspection
    - a tool for appraisal

Capability Maturity Model (CMM)

8

ASIC design verification for HEP

HEP Verification Maturity Model

	Level 0	Level 1	Level 2
Project Management	Processes are unpredictable, poorly controlled and often reactive. Lack of process standardization leads to inconsistency in project execution and outcomes. There are no well-defined roles and responsibilities. The role of project-manager itself is organic (one of the team members rise to the occasion)	Project management processes are established. Projects are planned and tracked and controlled. Success still depends on individual project managers abilities to follow established processes	Standardized project management processes are defined and documented across the organization. These processes are understood and followed by all the participating members. Project manager uses quantitative metrics to monitor projects
Reviews	Reviews are conducted on demand. There is no defined process or criteria for the review. They are usually viewed by the design team as bureaucratic process with no real outcomes (checking off a tick-box to meet submission guidelines)	Standard review processes are defined and communicated. Roles and responsibilities of review participants are defined. Review outcomes, action items, and resolutions are tracked. Reviews are typically conducted as condensed meetings (usually during major milestones) where reviewers only provide feedback at a very high level	Reviews are conducted at various stages of the project. Review stages include architectural reviews, design/verification strategy review, project progress reviews and sign-off reviews. Reviewers are expected to dive into greater project details and provide elaborate feedback
Requirements Management	Requirements are communicated between individuals in meetings or through other informal means. Some form of scattered documents (usually presentations) that contains requirements exist	Requirements are written down and exchanged between stakeholders through a central document usually shared via network drives. These requirement documents typically stay out of sync as projects advance with new requirements	Requirements are managed in a structured database using an industry standard requirements database. A clear process is defined for adding/removing/modifying a requirement and its impact on project planning and execution
Design Documentation	No structured documentation. Designers, design users and system developers often rely on reverse engineering the source code to understand design features. Some documentation is produced on demand – typically for reviews, conference presentations etc.	Design documentation exists. However, there is no clear distinction between various kinds of design documents (architecture specifications, micro-architecture documents and user manuals). There is no clear understanding of roles and responsibilities for document creation, review, approval, and maintenance	Develop and maintain standard document templates, style guides and terminology guidelines to ensure uniformity, consistency, unambiguity, and consistency. There is clear distinction between several design documents (architecture specifications, micro architecture documents and user manual. Processes are defined for roles and responsibilities for document creation, reviews approval and maintenance
Verification Documentation	None	No elaborate verification documentation. Test lists are written down and maintained	Verification documents created using standardized templates are maintained. These documents typically include verification requirements, verification strategy, verification plan (test plan, coverage plan, sign-off checklist). Verification plan is maintained in a structured database using an industry standard verification planner that integrates well into regression testing tools and facilitates metric driven verification
Source code management	Little to no version control, leading to potential conflicts and loss of code changes	Code changes are tracked through version control. However, lack of standards and workflows makes source code inconsistent	Source code management processes and workflows (proper branching and merging strategies) are established and followed. Coding standards and conventions are documented and enforced. Code reviews and continuous integration with automated testing are integrated into source code management process
Reproducibility of results	None	Results from some of the design activities are reproducible by the design team. There are still several manual steps involved which makes reproducibility hard	All the results from architectural exploration to sign-off are fully reproducible by any user at any point in the future. All the flows are fully automated and make no user-specific assumptions
Compute resources and Tools	Unplanned: Use what is available. Usually, project execution suffers towards the submission when the demand on Compute resources and Tool licences increases dramatically	The design is aware of compute and tool licensing requirements and plans the projects such that peak utilization can be absorbed. Several compute resources are typically under utilized	The design team can predict compute and tool usage based on historical data. Resource utilization is optimized
Verification Methodology	No methodology exists. The design team creates an ad hoc way of writing testbenches and testcases. The verification content becomes unmaintainable over time as project advances	The design team uses some structured way of writing testbenches and tests. The methodology is repeatable across projects that the design team works on. However, since the methodology is unique to the design team, it makes reuse of content across design teams harder. Source code becomes unmaintainable when the key members of the design team leave the project	Verification methodology is unified, standardized and reusable. Clear coding conventions for code style and idioms are established. Emphasis on maintainability of code. The methodology also considers reusability of the environment for several design activities such as performance analysis, architectural exploration etc. Moving engineers across projects can be achieved efficiently
Stimulus	Bit banging, directed tests usually written by the designer to check immediate correctness of the feature. Limited coverage of design space by the stimuli	Constrained random test vectors are used. Tests are run with several random seeds to maximize coverage and reduce corner cases	Scenario driven that integrates data from physics simulations along with constrained random vectors. There is close interaction between verification engineers, DAQ developers and detector physicist
Correctness	Manual checks by viewing waveforms and reading logs	Automatic checking with reference model and scoreboard	Scoreboards considers Key Performance Indicators. Scoreboards are precise for some features and approximate for other features
Verification Metrics	Checklist of testcases and features exercised. Tracked manually through verification engineer	Code coverage is collected and analysed. Some functional coverage is coded and collected towards the end of the project. Functional coverage is not exhaustive	Functional coverage is planned early in the project. Progress of verification is tracked using functional and code coverage metrics collected in regular weekly regressions. Design engineer implements inline micro-architectural coverage in RTL. The designer uses this as a metric to measure the quality of stimuli written by verification engineer (virtuous cycle of quality checks)
Debug	Manual. Typically involves viewing waveforms and following signals to understand the failure scenario and then root cause the failure	Some debug infrastructure (e.g., data logging, configuration logging etc.) is built into the verification environment which facilitates faster understanding of test scenario that caused the failure	Automatic. The regression testing environment can parse the logs of failing tests, pinpoint the causes of failures and notifies responsible engineers to further root cause or fix the identified issue
GL Simulations	Unplanned/ad-hoc	Late start with limited stimuli	Early planning, which involves decomposition of netlist simulation into several flavours which optimizes run-time. With early planning one can provide feedback into design partitioning that can simplify setup of GL simulation hierarchies
SEE Simulations	None. Design team relies on engineer's judgement and beam testing to gain confidence on SEE robustness of the design	Some SEE simulations are performed towards the end of the project. SEE injection is limited to a sub-set of nodes which designer thinks are critical. Checking is limited - typically restricted to making sure nothing catastrophic occurs. SEE simulations are usually an afterthought. Testbench and testcases go through significant updates to support SEE simulations	SEE simulations are not considered as an add-on to functional verification but are treated to be a part of it. The environment, tests, and checkers are all built to tolerate errors that may be observed during SEE campaigns. SEU/SET injection must be integrated into scenario-based test cases

ASIC design verification for HEP

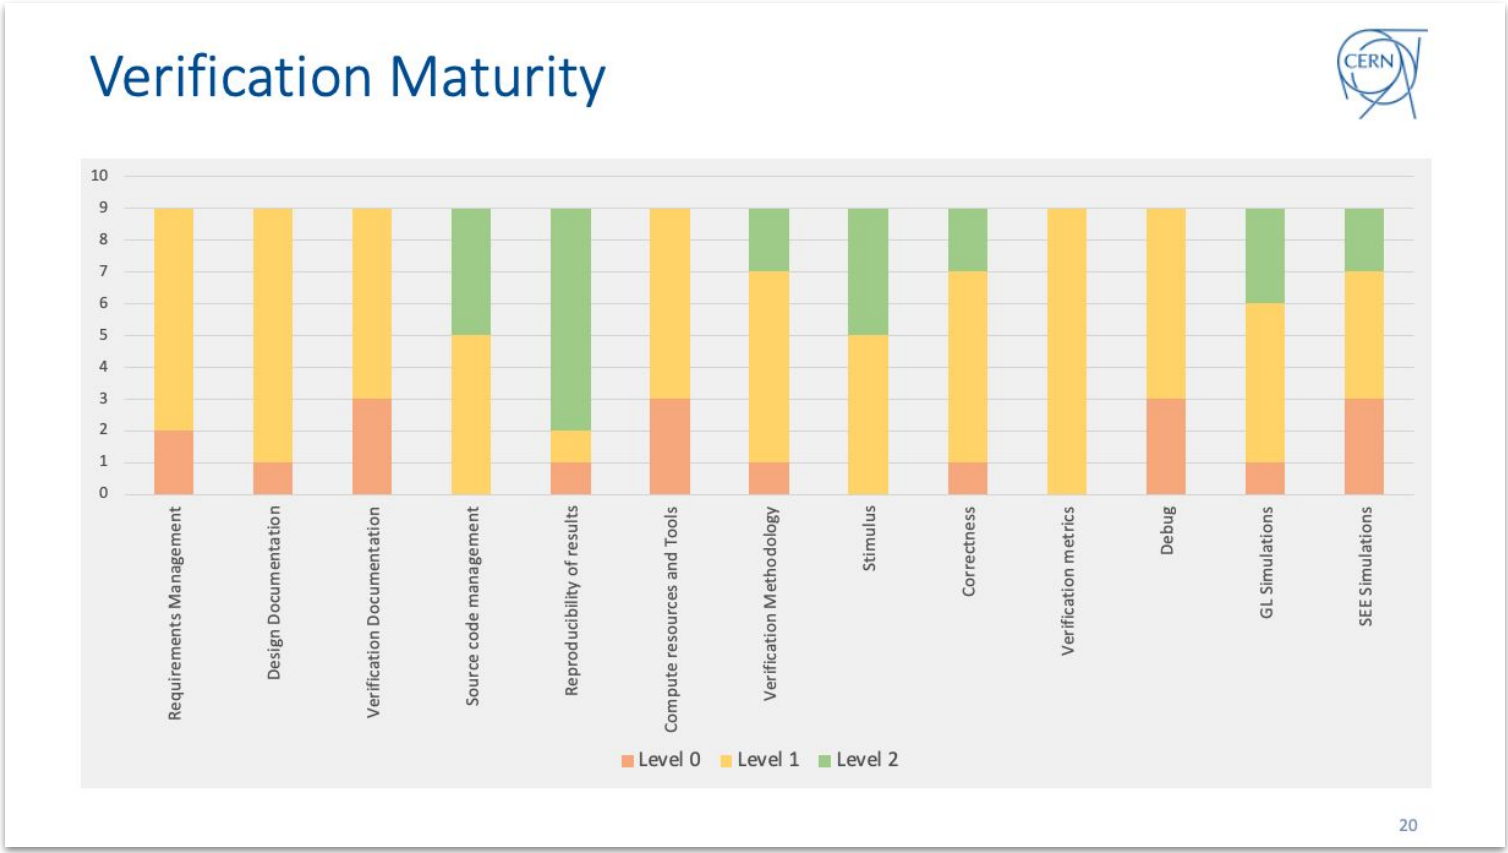

ASIC design verification for HEP

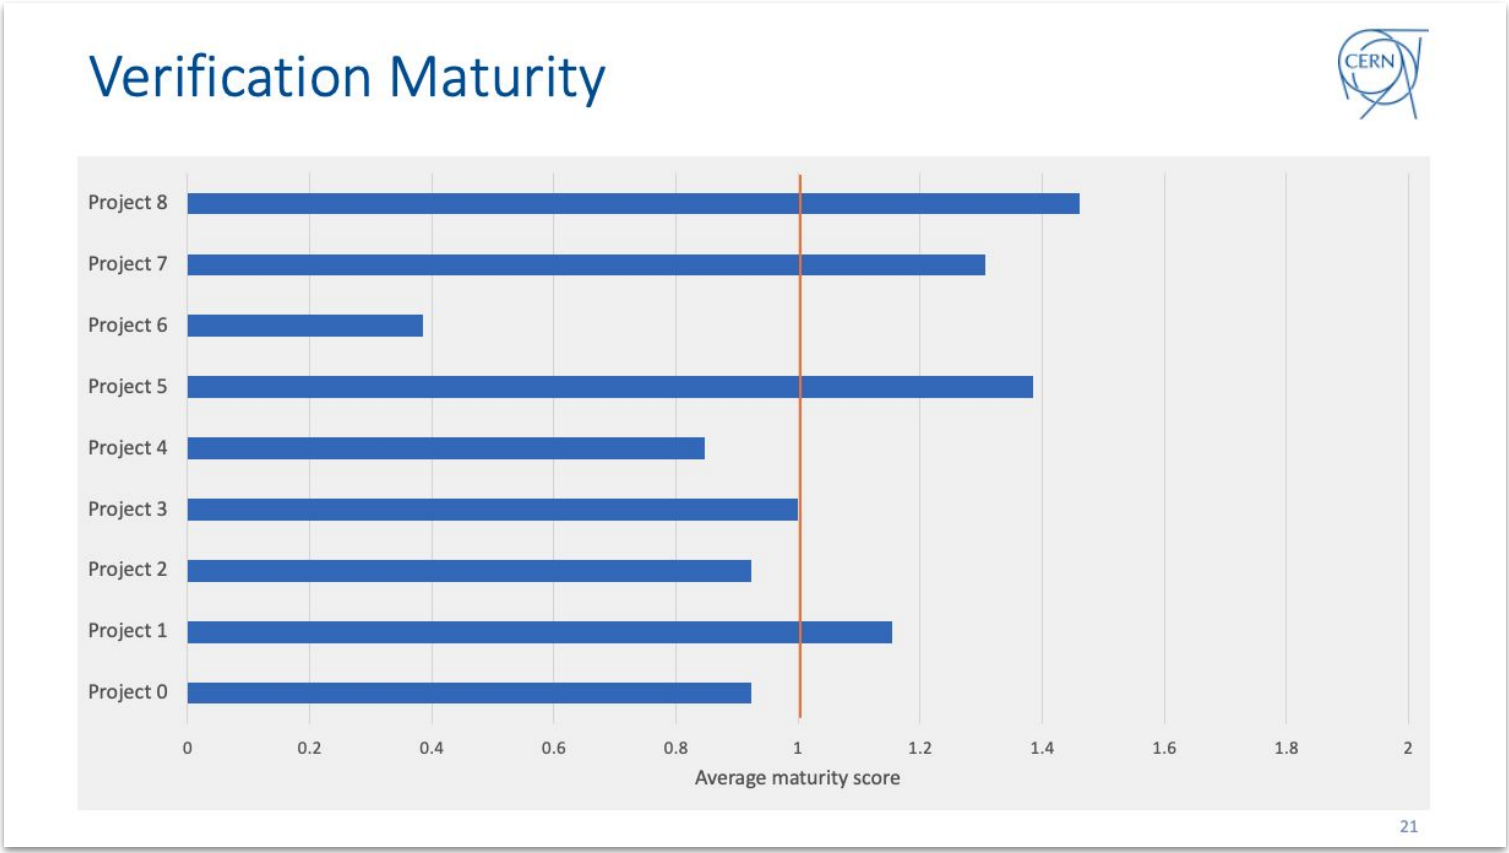

28nm IP development programme

RAL TD contribution to CERN common IP library

Amplifier IP

A range of amplifiers have been developed to validate the analogue technology limits of 28nm CMOS.

- Rail-to-Rail class AB Precision Amplifier
  - <800µW power consumption
  - 70x42um^2
  - Input range of 100-700mV
  - Input offset <3mV
  - High drive strength for high loads on-chip/off-chip
- PMOS input class AB High speed Amplifier
  - <700µW power consumption
  - 63x53um^2
  - Input range of 100-410mV
  - Bandwidth 100MHz at 1pF load
  - Input offset <5mV
  - High speed for driving fast signals or references.
- NMOS input class AB High speed Amplifier
  - <700µW power consumption
  - 60x40um^2
  - Input range of 390-700mV
  - Bandwidth 100MHz at 1pF load
  - Input offset <5mV
  - High speed for driving fast signals or references.

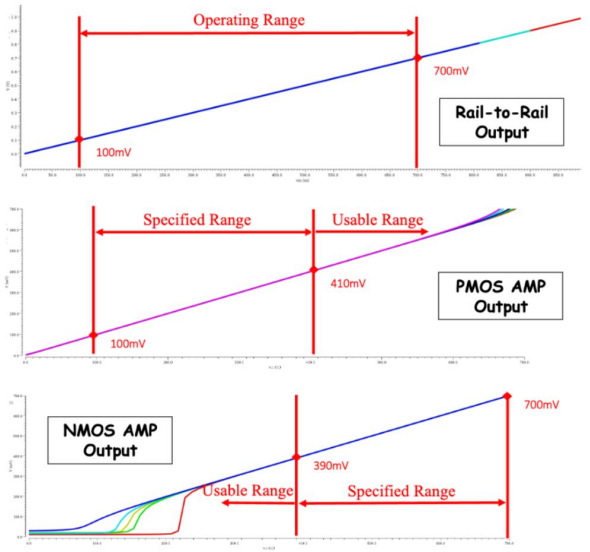

+ bandgaps, reference drivers, DACs, 1Gbps serializer, …

Front-end RDMA over converged Ethernet

Introduction on RDMA and RoCE

In a DAQ system a large fraction of CPU resources is engaged in networking rather than in data processing; common network stacks that take care of network traffic usually manipulate data through several copies.

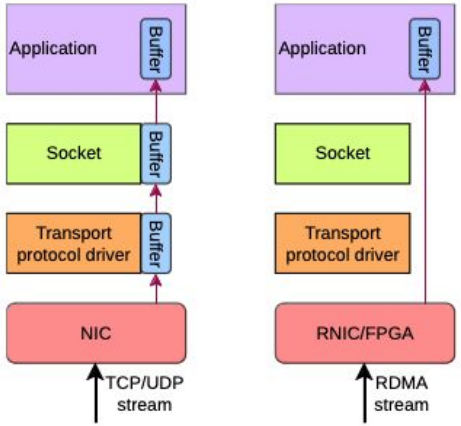

Remote Direct Memory Access (RDMA), as the name suggests, allows read and write operations directly in the target machine(s). This implies no OS involvement allowing high-throughput and low-latency applications.

This requires RDMA enabled NICs on both ends (RNIC) that perform the DMA, reducing the CPU load.

Front-end RDMA over converged Ethernet

What is FERoCE?

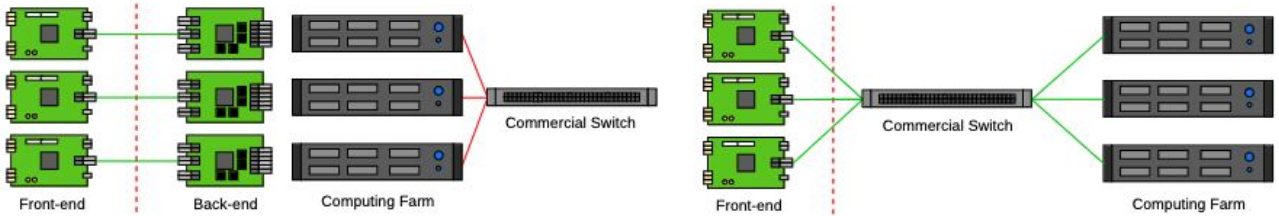

Back-end boards required to get the data, and send it to the computing farms. This requires multiple custom cards and custom boards

Front-end boards send data already packaged within an ethernet frame allowing switching and routing. Choosing the proper protocol allows the use of COTS switches

ETH RDMA network stack library has been chosen for the first prototype. Some of its characteristics:

- Entirely written in HLS (Vivado 2019.1)
- It targets Xilinx FPGA with PCIe connection
- 10/100 Gb/s speeds
- It supports UDP, TCP and RDMA

Systems @ ETH zürich

Front-end RDMA over converged Ethernet

Real-time Firmware Simulation

Start form ETH network stack entirely developed in HLS. Functionalities and features must be understood: real-time firmware simulation with real network traffic.

- Works on Linux machines: Tun/Tap devices
- It makes use of DPI-C interface of SystemVerilog: C code in our testbench!
- Tap device exchanges raw ethernet frames between simulation and Linux network stack
- We can capture such frames and study them

Simulation with Synopsys VCS. XGMII interface directly form Xilinx MAC

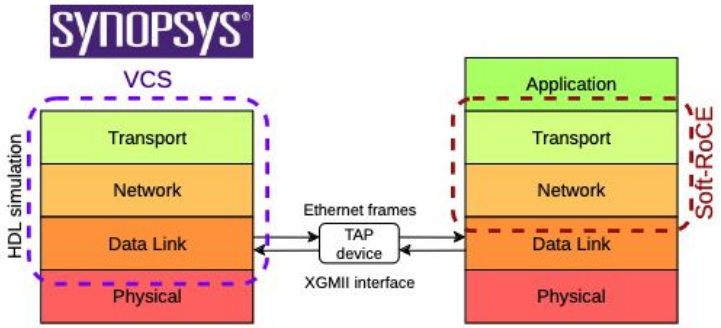

Soft-RoCE used to capture and store in memory data sent. Enable fast verification of the stack without going through sythesis/implementation every time.

Once the stack has been verified, firmware can be eventually built (Resources? Performances? Is timing closure reached?)

Front-end RDMA over converged Ethernet

Some results -Wireshark-

Used Wireshark to capture Ethernet frames coming out of the simulation.

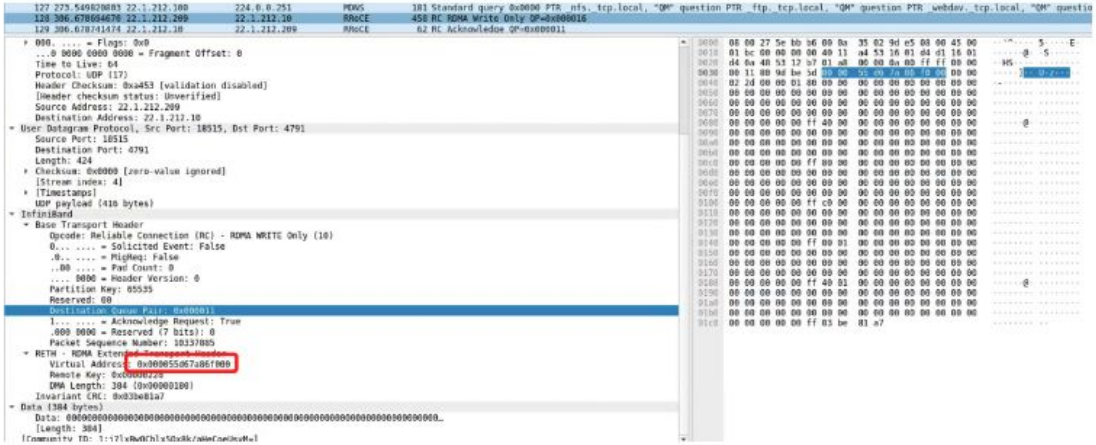

In this frame we can check:

- Queue Pair number
- RDMA OP Code
- IP addresses
- Memory addresses

RISC-V: Fault tolerance via triplication

Fachhochschule Dortmund
University of Applied Sciences and Arts

Introduction

- Many custom ASICs have a similar structure:

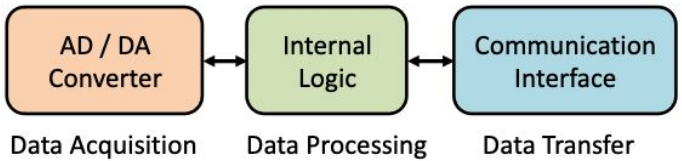

- Design and verification of a custom ASIC is complex and time-consuming
- Reuse of generic blocks possible (ADC, voltage regulators, etc.)
- Adaptation of internal logic difficult, custom to original application

- Internal data processing logic replaced by with RISC-V processing system
  - Adaptation to new application / Bugfixes via firmware updates
- Hybrid detector with RISC-V-based microprocessor SoC

RISC-V®

RISC-V: Fault tolerance via triplication

Fachhochschule Dortmund
University of Applied Sciences and Arts

STRV-R1 – Architecture

- RV32-IMC Core
  - 3 stage pipeline
  - Multiplication extension
  - 50 MHz @ 1.2V
  - Fully triplicated core
- SRAM shared between instruction & data
  - Flexible memory layout
  - IMEM & DMEM data bus can access whole SRAM address range
  - RISC-V pipeline stalls during load & store instructions to SRAM
  - load & store to peripherals simultaneously possible
- JTAG Interface
  - JTAG TAP & debug module
  - Non-volatile debug ROM with debug ISR

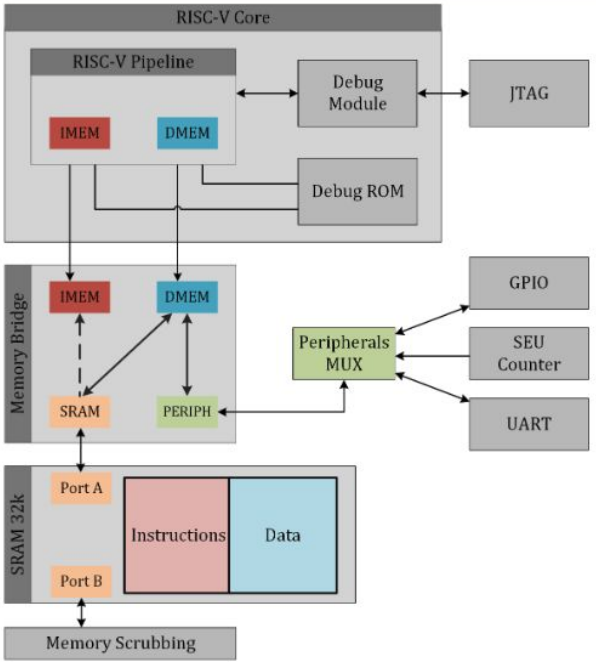

RISC-V: Fault tolerance via triplication

Fachhochschule Dortmund
University of Applied Sciences and Arts

STRV-R1 – Heavy-Ion Irradiation SEFI

- Despite the SEE mitigation techniques SEFIs o
  - SEFIs observed during heavy-ion Irradiation
  - Average improvement over SEU cross-section
    - At low LETs (<16 MeV.cm²/mg): 2800x
    - At high LETs (>32 MeV.cm²/mg): 7700x
- Estimated SEFI rate in HL-HLC environment
  - SEE particle flux 1 × 10^9 p/cm²/s
  - 2.2 Chip level SEFI per hour

Cross-section	L_0 $[\frac{MeV cm^2}{mg}]$	$\sigma_{HI\infty}$ $[cm^2]$	$L_{0.25}$ $[\frac{MeV cm^2}{mg}]$	$\sigma_{p\infty}$ $[cm^2]$
SEU	< 1.0	4.27×10^{-2}	18.87	2.66×10^{-9}
SEFI	< 5.7	2.95×10^{-6}	10.32	6.15×10^{-13}
Timing	< 3.3	2.86×10^{-5}	19.95	1.58×10^{-12}

Cross-section	$\sigma_{p\infty}$ $[cm^2]$	[a] Φ $[\frac{Hz}{cm^2}]$	Event rate $[Hz]$	Events / h $[\frac{1}{h}]$
SEU	2.66×10^{-9}	1×10^9	2.66	8.14×10^3
SEFI	6.15×10^{-13}	1×10^9	6.15×10^{-4}	2.21

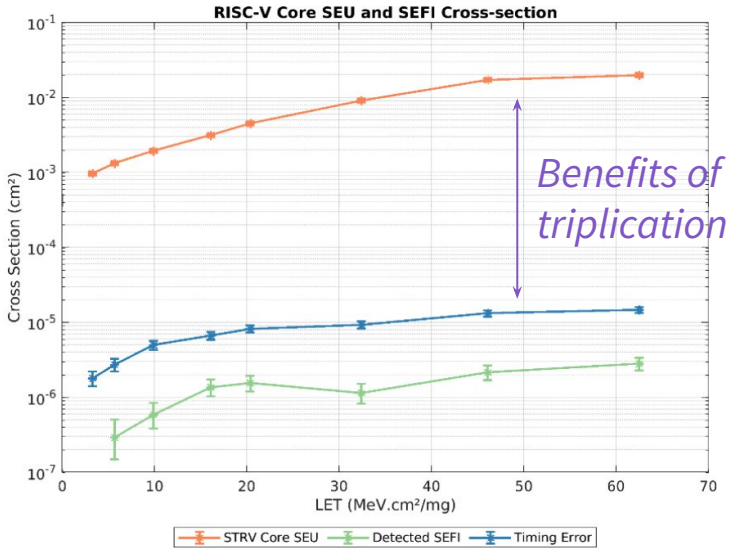

RISC-V: SEU detection

Motivation

Security Features

- Ibex can implement a set of extra features to support **security-critical** applications
- Main strategy: Ibex core can detect external attacks due to corrupted states
- Alerts provided by dedicated signals

Research Question:

Can these built-in security features be used to detect SEUs within the Ibex core?

- ✓ Alert outputs
- ✓ Dual core lockstep
- ✓ Bus integrity checking
- ✓ Register file write enable glitch detection and ECC
- ✓ Hardened PC
- ✓ Shadow CSRs

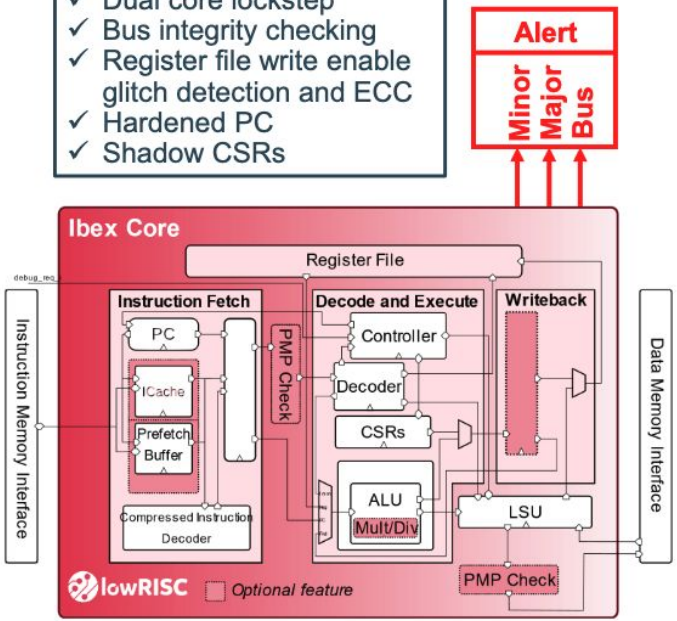

https://ibex-core.readthedocs.io/en/latest/03_reference/security.html

RISC-V: SEU detection

Research Methodology

Testbench architecture

- CoCoTB testbench
  - Ibex RTL code
  - Python models for SoC
    - Data/Instruction memory
    - Stdio
    - …
  - Random SEU injection

  (Pre-pass with Genus to extract flip-flop list)
- Application code compiled and loaded in I-memory
- Xcelium RTL simulator

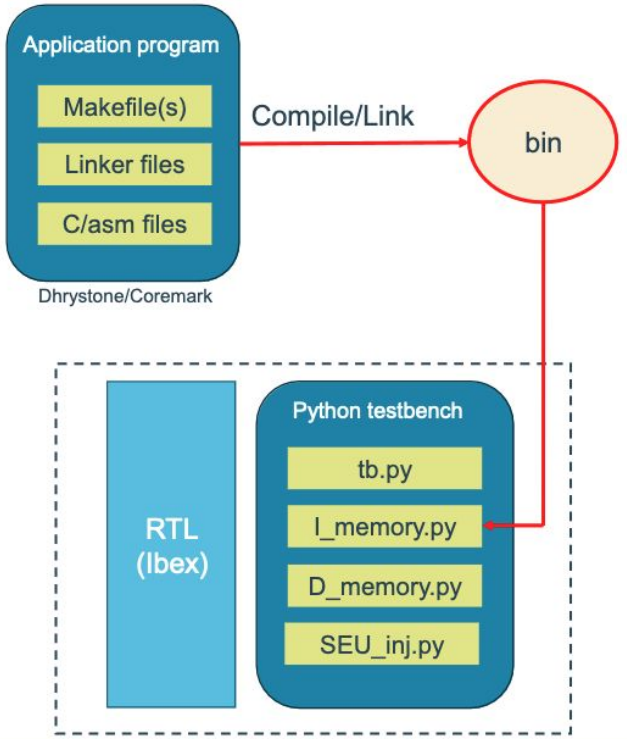

RISC-V: SEU detection

Fault Injection Simulation Results

Results by symptom

Target	Total flips	Alert major internal	Alert major bus	Alert minor	No error	Undetected flips
data_req_o	299600	7230	2723	0	292370	0
data_we_o	299600	7314	2743	0	292286	0
data_be_o	299600	8024	2784	0	291576	0
data_addr_o	299600	12692	2784	0	286908	0
data_wdata_o	299600	13460	2784	0	278201	7939
data_wdata_intg_o	299600	13457	2784	0	278204	7939
instr_req_o	299600	7279	2753	0	292321	0
instr_addr_o	299600	7523	2855	0	292077	0

↓

TB found CRC error but alert was low

Any questions?

CMS Level-1 Trigger: Anomaly detection

Model Performance

- AXOL1TL is trained with unbiased data collected by the CMS Experiment during 2023 with √s=13.6 TeV
  - 10.5 million events used – 50% for training, 50% for setting thresholds
- Dotted lines represent the score thresholds implemented in the Global Trigger Test Crate
- Significant performance improvement on various SM and BSM signals by adding AXOL1TL to the 2023 trigger menu
  - Signal samples are Monte-Carlo generated
  - Table shows performance improvement for a Higgs decaying to 2 (pseudo-) scalars to bottom quarks

h->a(bb)a(bb)

AXOL1TL Rate	1 kHz	5 kHz	10 kHz
Signal Efficiency Gain	46%	100%	133%

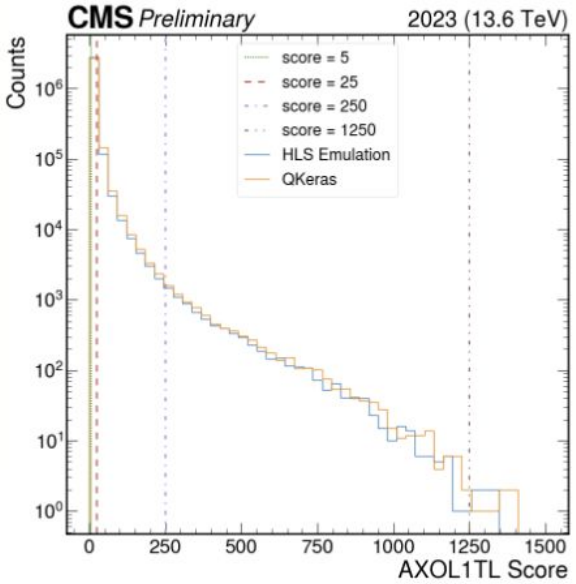

CMS Level-1 Trigger: Anomaly detection

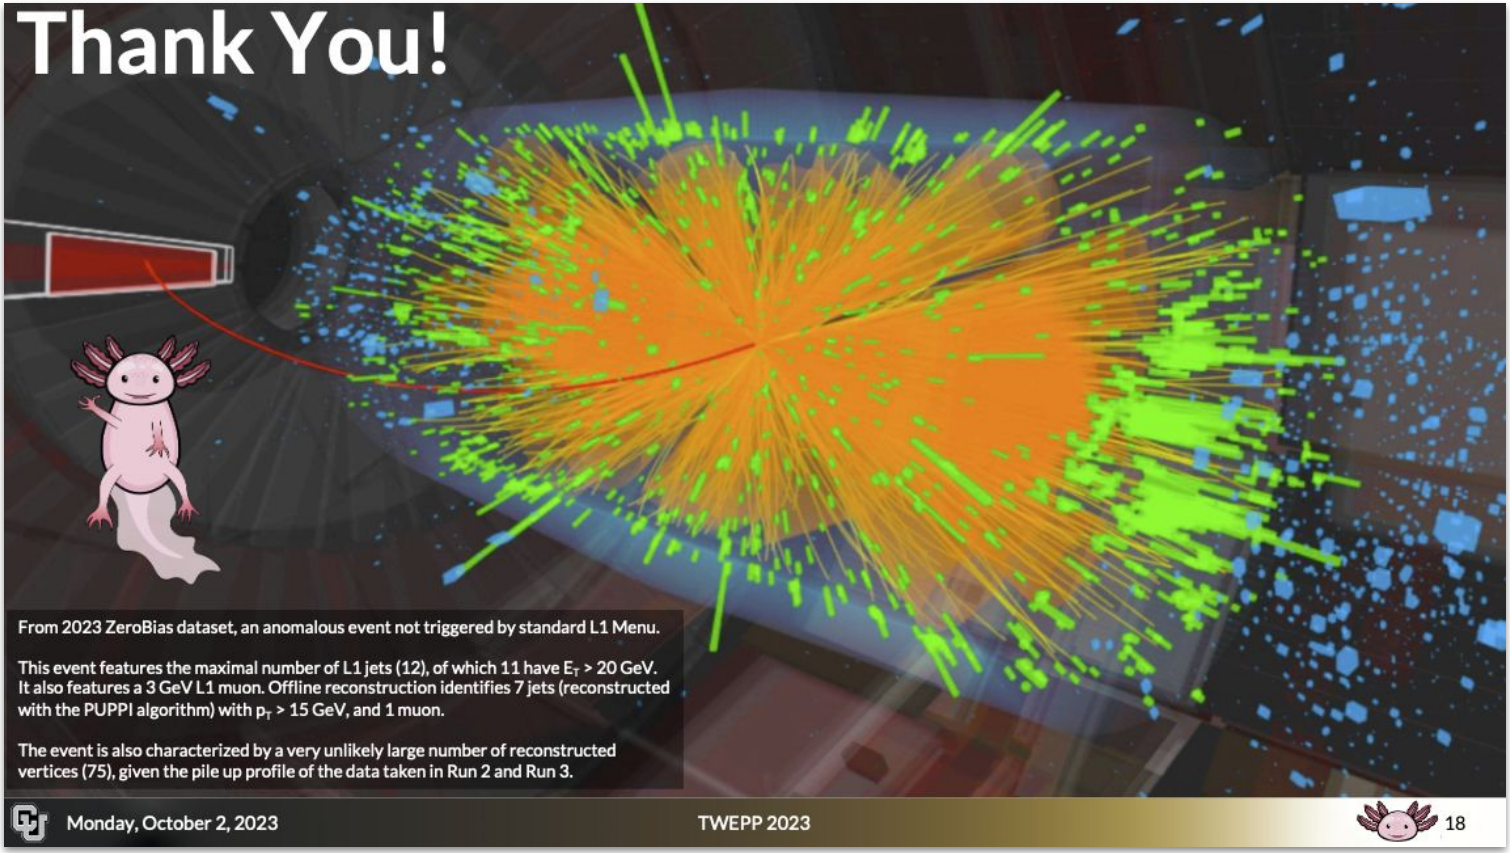

GNN in an FPGA @ sPHENIX

The DAQ–AI Data Flow

- Motivation to use FELIX board:
  - To reuse the PCIe implementation (16-lane Gen-3) and software tools provided by the FELIX developers
  - on-board FPGA is a Kintex Ultrascale XCKU115FLVF1924-2E
- The decision signal of heavy flavor event from the AI-Engine will be sent out via the LEMO connectors to the sPHENIX GTM/GL1 system to initiate the TPC readout in the triggered mode
- GPU based feed-back system for the beamspot monitoring

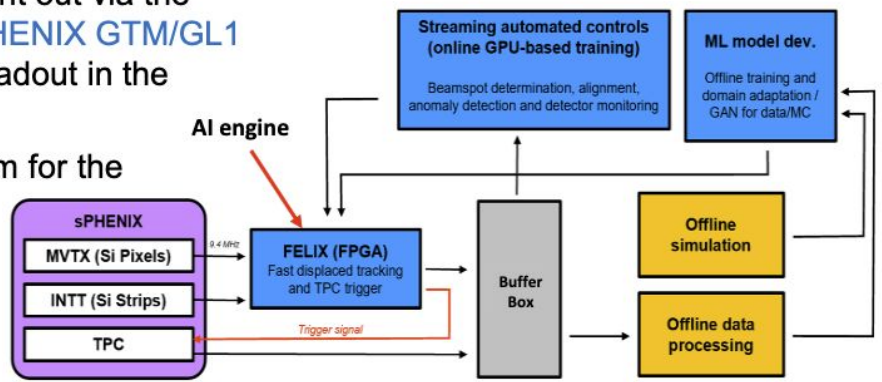

ATLAS Global Common Module (GCM)

Science and Technology Facilities Council

Brookhaven National Laboratory | ATLAS EXPERIMENT | Science and Technology Facilities Council

Global Trigger Architecture

- Data for a single Bunch Crossing are time-multiplexed onto a single node

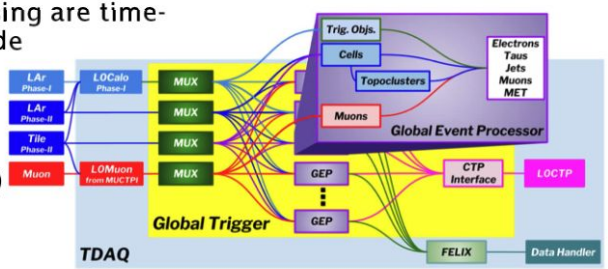

- Three functional layers
  - Multiplexing Processor (MUX)
    - ~56 nodes
  - Global Event Processor (GEP)
    - ~49 nodes
  - Global Central Trigger Processor Interface (gCTPi)
    - 1 node

- One hardware platform for all three functional layers
  - Global Common Module (GCM)
    - Simplify firmware/software development
      - Sharing the common infrastructure
    - Simplify the long-term maintenance
      - Reducing the number of spare modules needed

Track reconstruction in FPGAs @ LHCb

LHCb in Run 5&6 ?

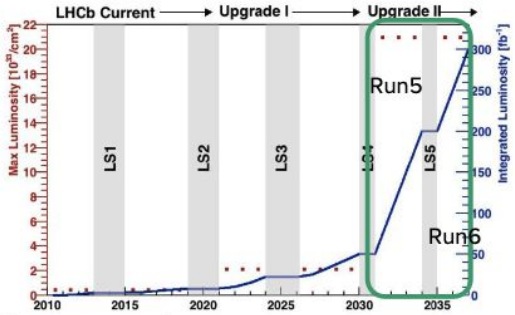

- Target: ~300 fb^{-1}
- Pile-up: ~40
- To keep the same performance in more difficult conditions, timing will be required in some sub-detectors
- **200 Tb/second data produced**
- **More processing has to be performed earlier in the DAQ Chain to reduce data offline**
- **Moving to a "heterogeneous-computing" paradigm**

LHCb-TDR-023

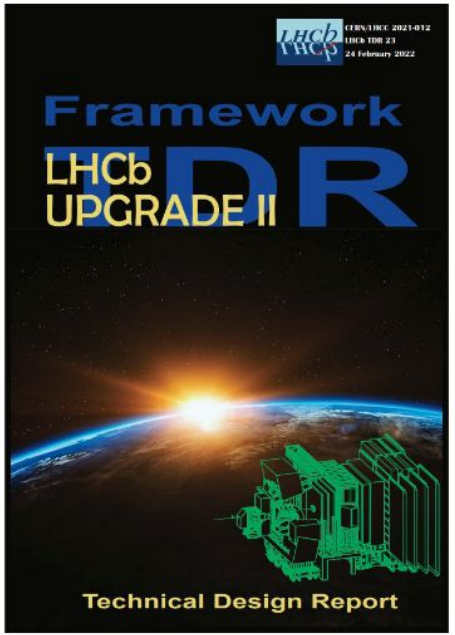

Track reconstruction in FPGAs @ LHCb

Real time tracking with FPGAs

- Modern FPGAs can perform parallel data processing with high throughputs, low latencies and better energy efficiency than CPUs and GPUs (for certain tasks)
- **This talk:** demonstrator system for real-time tracking on FPGAs with the **"artificial retina"** architecture to reconstruct tracks in the Vertex Locator

PCIe 16x board, 1 Intel Stratix 10 FPGA, 16 optical links

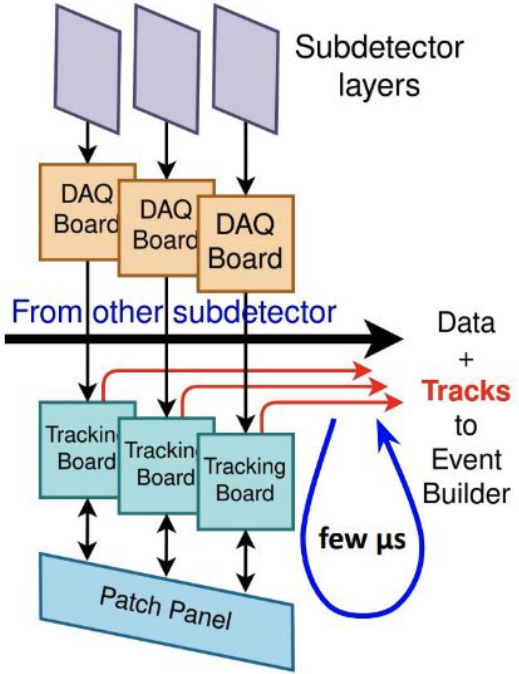

Track reconstruction in FPGAs @ LHCb

The “artificial retina” architecture [NIMA 453 (2000) 425-429]

Track parameter space divided into cells (pattern tracks)

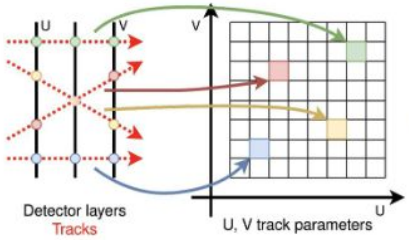

Each cell computes a weighted sum of hits near the reference track

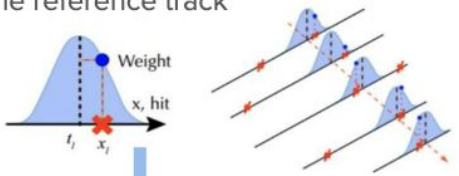

- Reconstructed tracks correspond to local maxima in the matrix of cells response
- Final track parameters from interpolating responses of nearby cells.

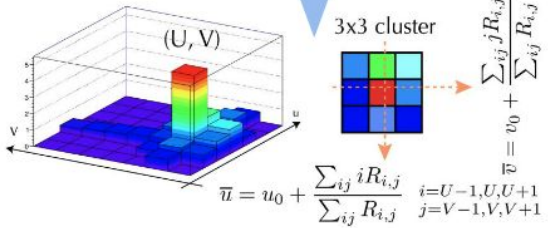

- **Cells work in parallel**: high-throughput and low-latencies
- **FPGA size limitations overcome by spreading cells over several chips** (without increasing latency).

Track reconstruction in FPGAs @ LHCb

The RETINA demonstrator at the testbed

- Simulated data used for high rate tests
- **Now live data from the LHCb monitoring farm**
- Demonstrates that a RETINA based tracking on FPGAs possible in HEP experiments:
- Current setup:
  - Reconstructs tracks of a VELO quarter
  - Spread over multiple PCIe-hosted FPGA cards. **8 cards** are sufficient
  - Scalable to cover the whole detector with additional FPGA cards.

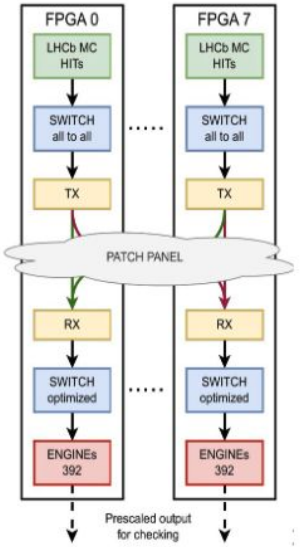

Track reconstruction in FPGAs @ LHCb

Distribution network

- As the RETINA algorithm is spread among several boards, a distribution is needed to exchange hits among boards:
  - 8 nodes full-mesh network
  - 28 full-duplex links at 25.8 Gbps
  - Total bandwidth 1.41 Tbps

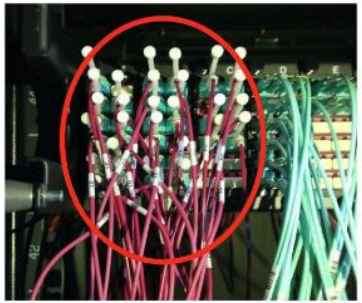

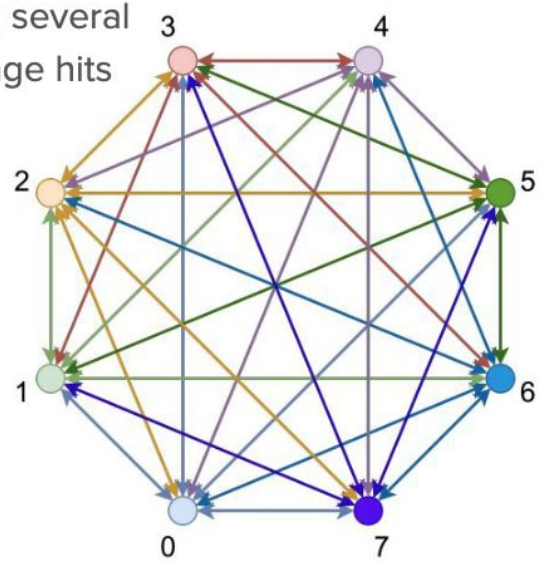

SystemC framework for architectural modelling

ELECTRONIC SYSTEM DESIGN FLOW IN HEP

Limitation of the currently used design flow:

- Based only on a low-abstraction level description of the system (hardware level).
- Architecture exploration is time and resources-heavy
- in multi-chip modules/detector, single chip are optimized separately

"Standard" design flow:

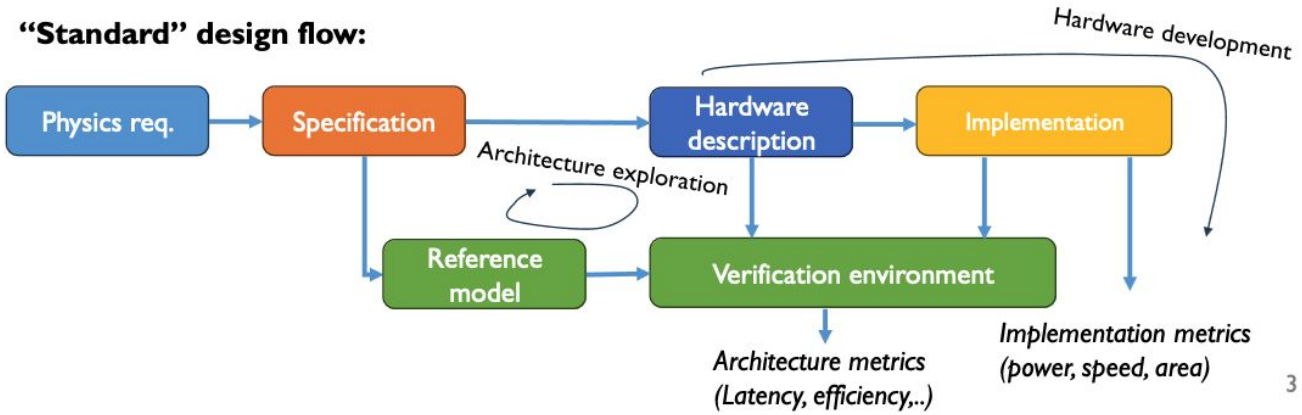

3

SystemC framework for architectural modelling

ELECTRONIC SYSTEM LEVEL APPROACH

Develop a **high abstraction level** description of the system, from front-end to back-end, for:

- architecture exploration
- new feature development
- reference model development

Requires a self-contained environment for Virtual prototyping

Proposed design flow:

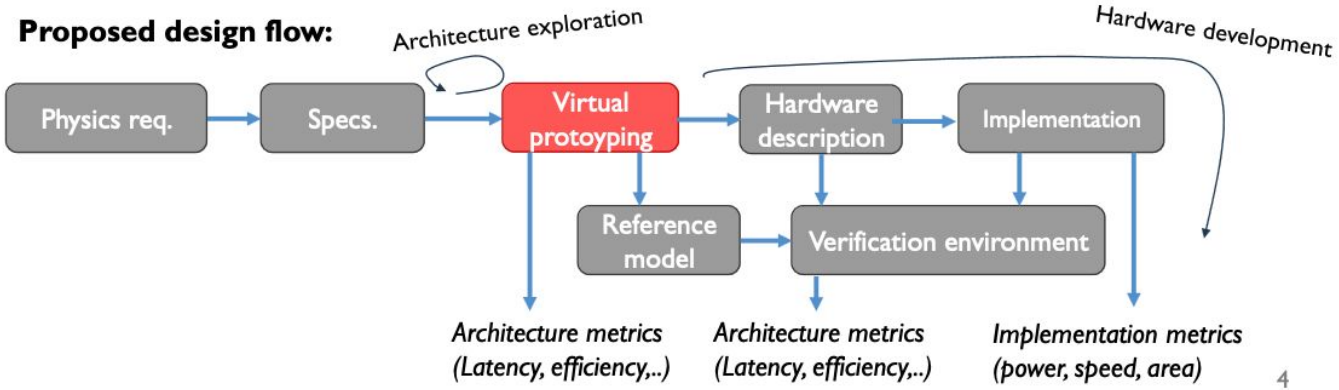

4

SystemC framework for architectural modelling

PixESL: an Electronic System Level prototyping framework

Open source:

- The model is based on C++ and SYSTEMC
- Performance analysis are based on Python

User-friendly:

- User and developer roles are separated
- The framework supports architectural and network configurability (structure, memory, arbitration, interconnections)

Reusable:

- Generalized layers and standardized packet transport (TLM)
- A library of layer types, functional components, and packet transport types
- Common integrated metrics analyzer

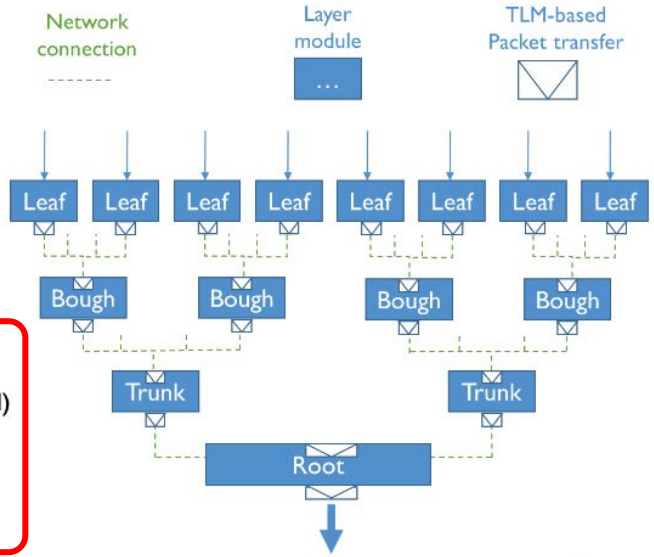

5

SystemC framework for architectural modelling

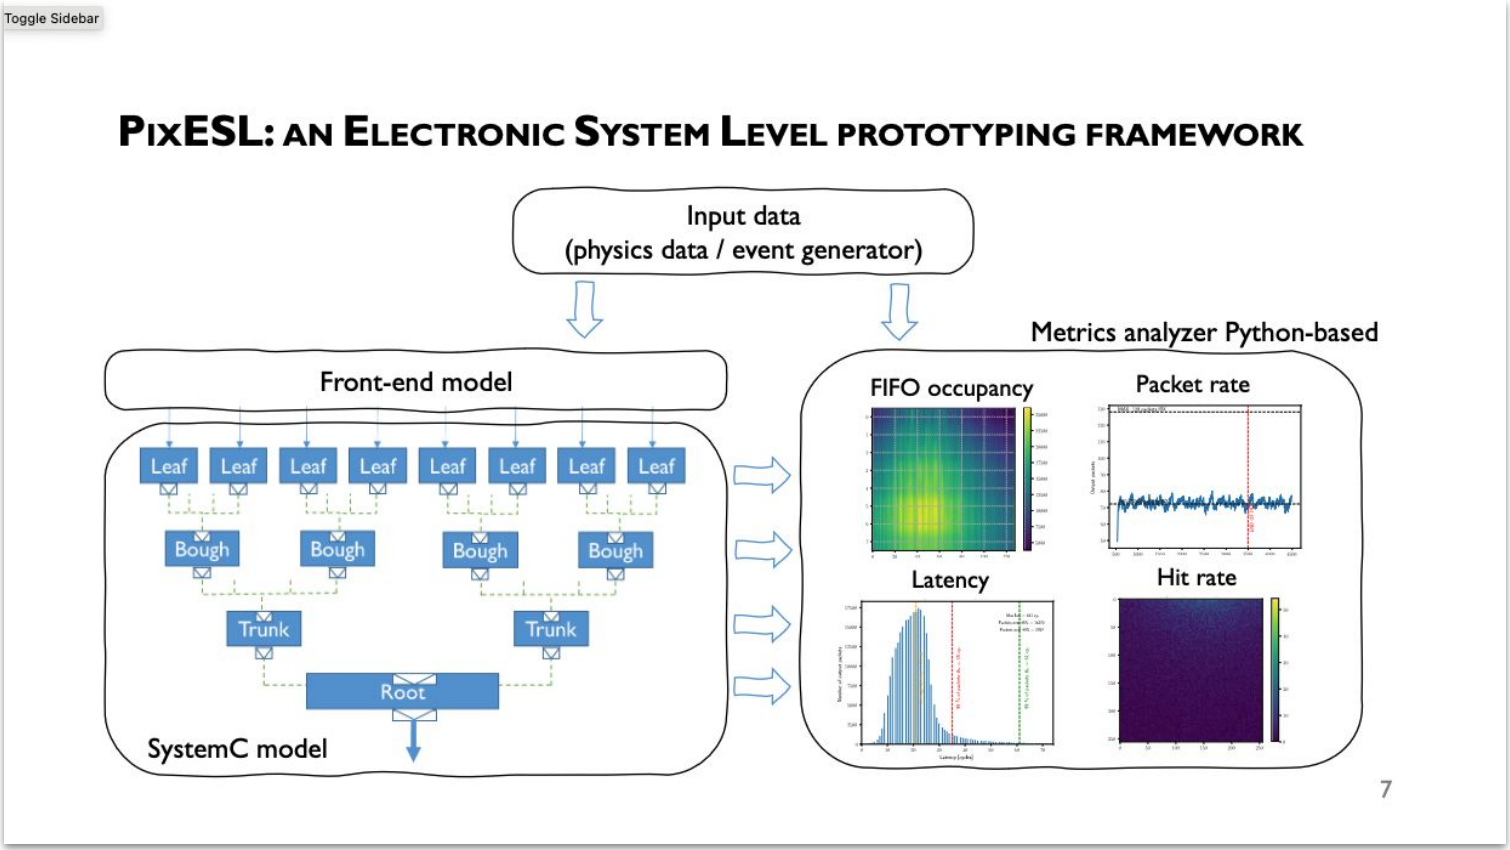

SystemC framework for architectural modelling

LHCb Velo upgrade II architecture exploration

The upgrade aims at a 4D pixel detector.

https://cds.cern.ch/record/2844669/

Main readout challenge:

extreme occupancy (x2 Velopix)

Flow:

Model Velopix (VELO upgrade I ROC)

Simulate higher occupancy events

Find bottlenecks

Optimize architecture

Repeat!

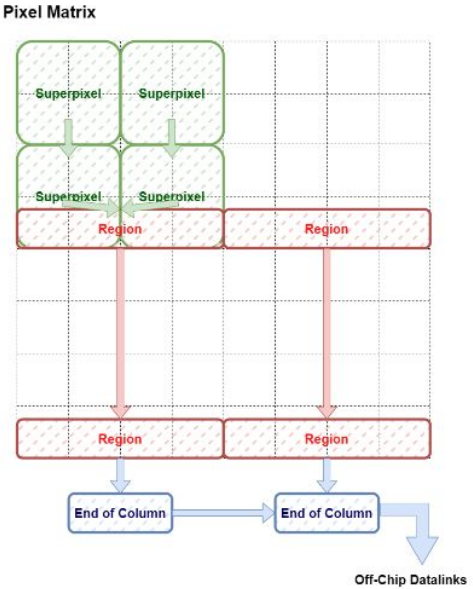

10

SystemC framework for architectural modelling

ON-CHIP PACKET SORTING DESIGN SPACE EXPLORATION

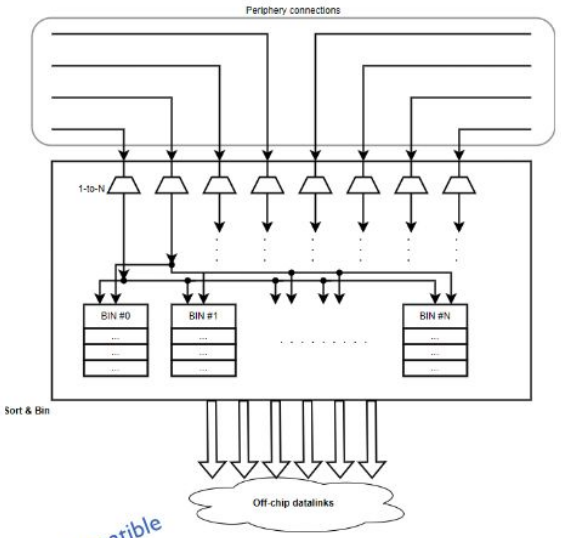

HLS compatible

SystemC description of the module

The number of bin depends on the latency of the readout efficiency and on the target grouping efficiency

Number of bins

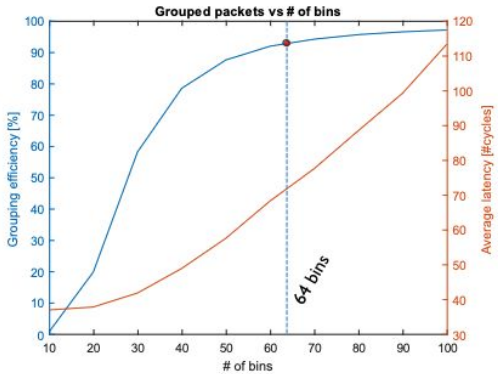

Target: >90% grouping

16