



UNIVERSITÀ
DEGLI STUDI
DI PADOVA



CMS Level-1 trigger Data Scouting firmware prototyping for LHC Run-3 and CMS Phase-2

Topical Workshop on Electronics for Particle Physics 2023

R. Arduino^{1,2,3} for the **CMS collaboration**

¹Università degli Studi di Padova

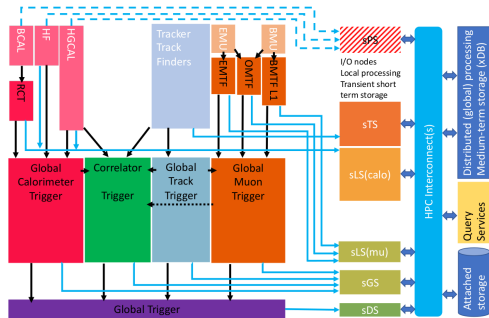
²INFN Sezione di Padova

³CERN, Geneva, Switzerland

October 2, 2023

Introduction

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



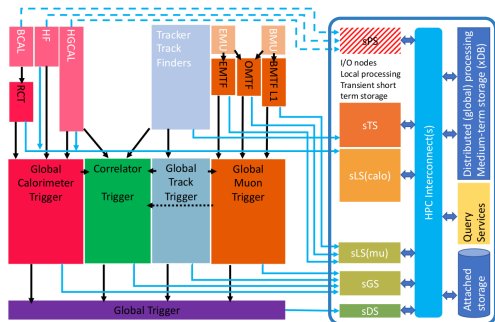
Increased budget latency and rate:

- $3.8 \mu\text{s} \rightarrow 12.5 \mu\text{s}$
- 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**

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



Increased budget latency and rate:

- $3.8 \mu\text{s} \rightarrow 12.5 \mu\text{s}$
- 750 kHz of L1 output

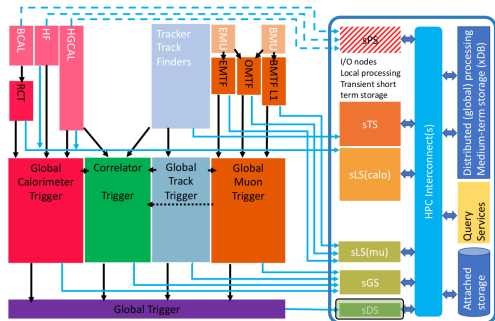
Advanced object reconstruction on FPGA:

- **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

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



Increased budget latency and rate:

- $3.8 \mu\text{s} \rightarrow 12.5 \mu\text{s}$
- 750 kHz of L1 output

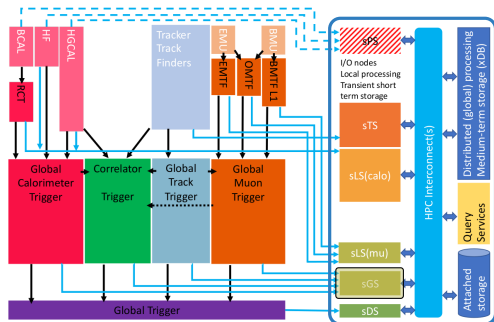
Advanced object reconstruction on FPGA:

- **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** \Rightarrow **sDS**

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



Increased budget latency and rate:

- $3.8 \mu\text{s} \rightarrow 12.5 \mu\text{s}$
- 750 kHz of L1 output

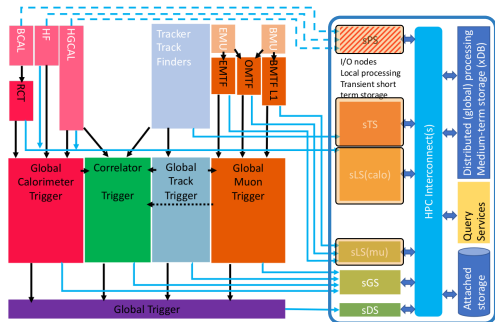
Advanced object reconstruction on FPGA:

- **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** \Rightarrow **sDS**
- **Correlator, Global Track, Global Calorimeter and Global Muon Trigger** \Rightarrow **sGS**

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



Increased budget latency and rate:

- $3.8 \mu\text{s} \rightarrow 12.5 \mu\text{s}$
- 750 kHz of L1 output

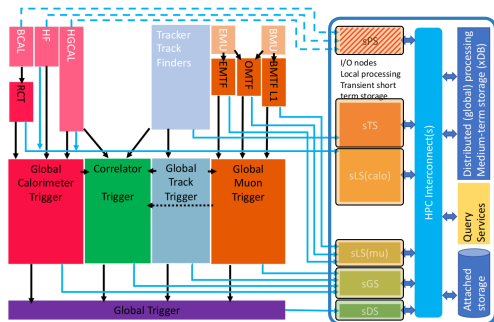
Advanced object reconstruction on FPGA:

- **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** \Rightarrow sDS
- **Correlator, Global Track, Global Calorimeter and Global Muon Trigger** \Rightarrow sGS
- **Can later be extended to include other systems in later stages** \Rightarrow sLS

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: $W \rightarrow 3\pi, D_s\gamma, H \rightarrow p\gamma, \phi\gamma, \dots$
- 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** \Rightarrow sDS
- **Correlator, Global Track, Global Calorimeter and Global Muon Trigger** \Rightarrow sGS
- **Can later be extended to include other systems in later stages** \Rightarrow sLS

The 40 MHz scouting system for Phase-2

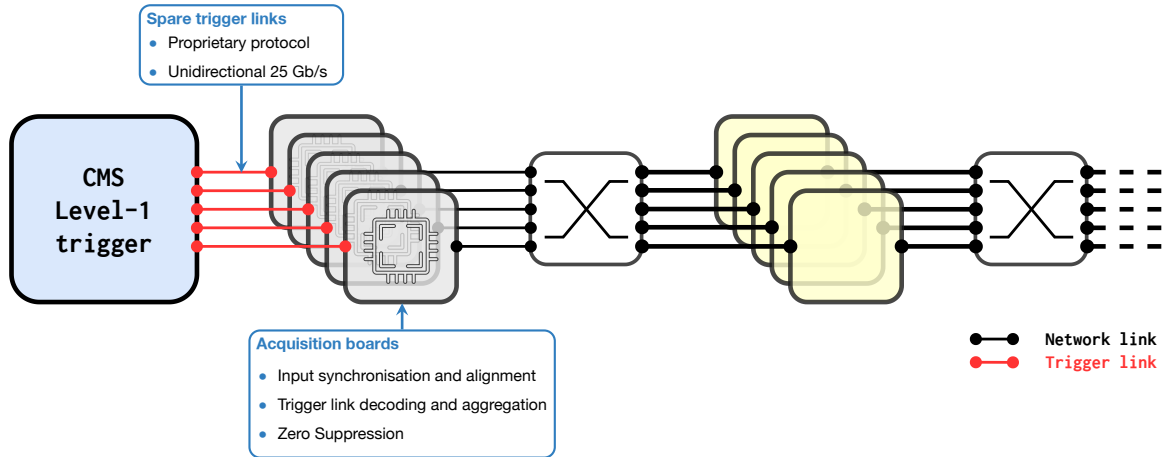
The Level-1 Data Scouting system for Phase-2: architecture



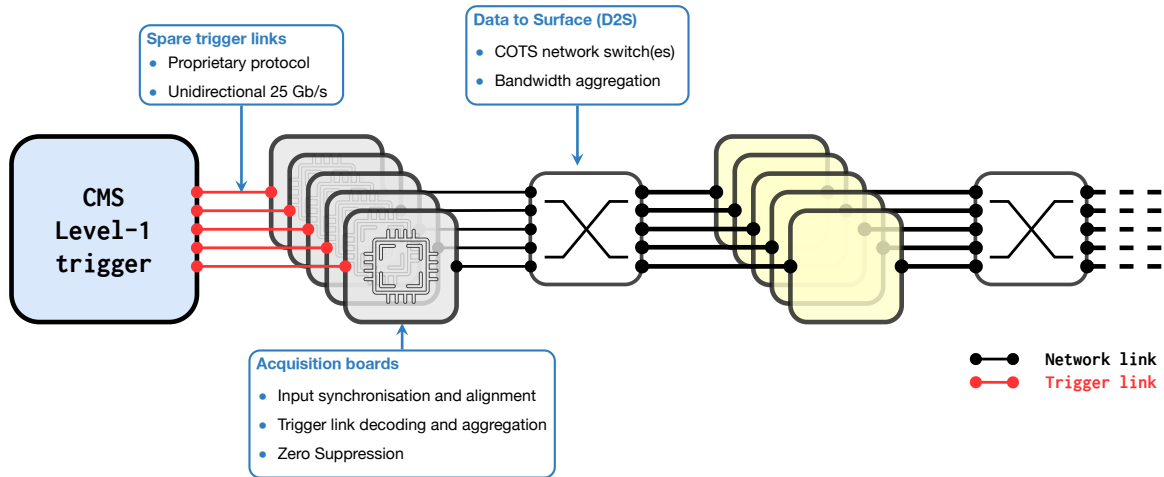
The Level-1 Data Scouting system for Phase-2: architecture



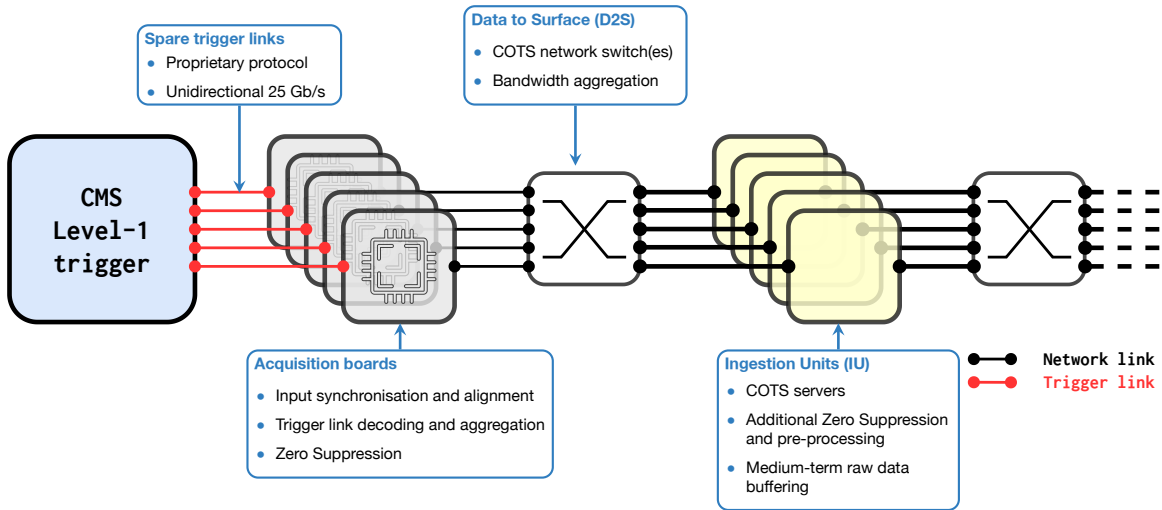
The Level-1 Data Scouting system for Phase-2: architecture



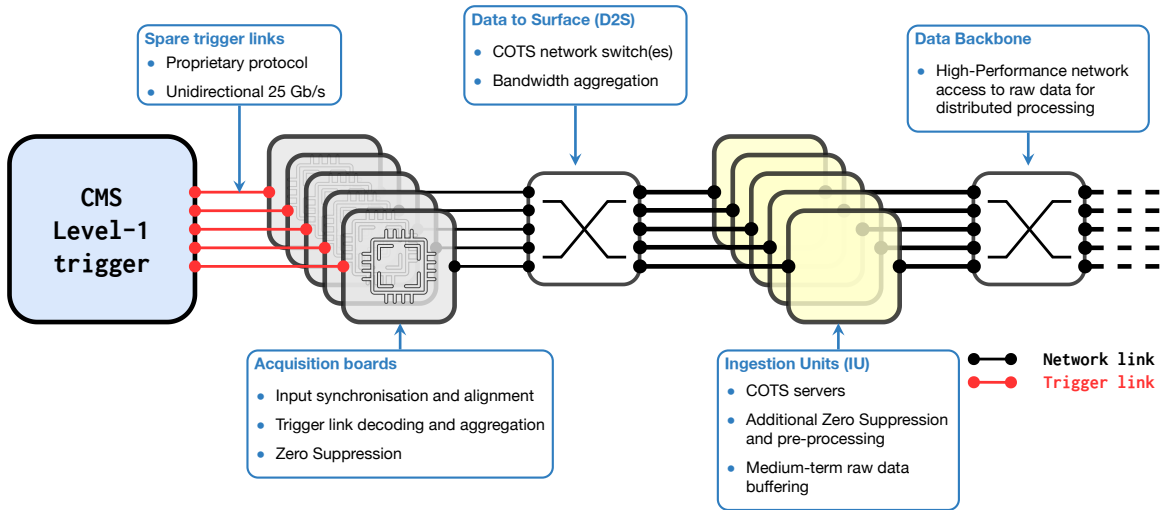
The Level-1 Data Scouting system for Phase-2: architecture



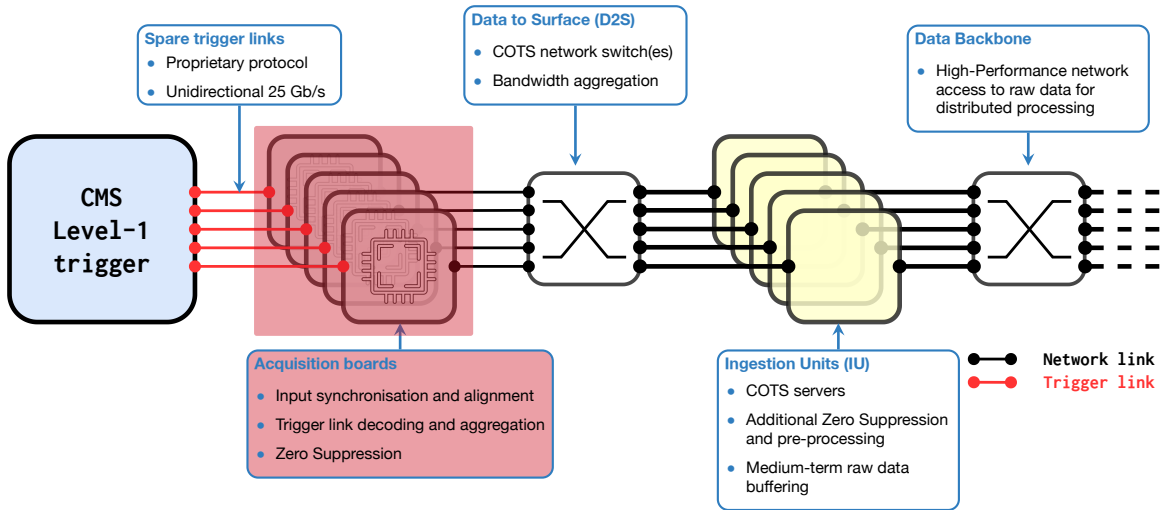
The Level-1 Data Scouting system for Phase-2: architecture



The Level-1 Data Scouting system for Phase-2: architecture




The Level-1 Data Scouting system for Phase-2: architecture



DAQ-800 readout board for cDAQ and L1DS

Scouting with DAQ-800 hardware platform:

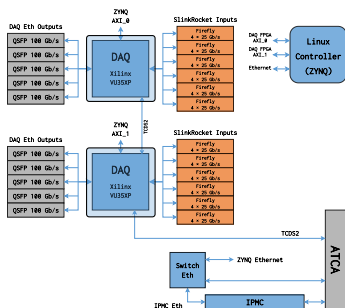

- CMS Phase-2 central DAQ readout board with ATCA form factor
- Production starting in 2024, 5 prototypes will be available in 2024 Q1



DAQ-800 readout board for cDAQ and L1DS

Scouting with DAQ-800 hardware platform:

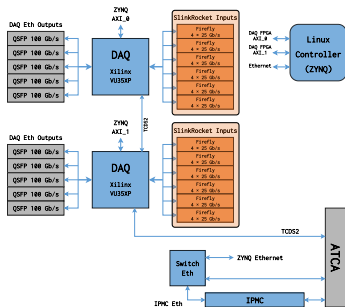

- CMS Phase-2 central DAQ readout board with ATCA form factor
- Production starting in 2024, 5 prototypes will be available in 2024 Q1
- **2 × Xilinx Virtex Ultrascale+ VU35P FPGAs with 8 GB of HBM memory** for output buffering
- **6 × 4 FireFly inputs / FPGA** ⇒ **48 × 25 Gb/s total input** (total: 1.2 Tb/s)
- **5 × QSFP outputs / FPGA** ⇒ **10 × 100 Gb/s total output** (nominal total: 800 Gb/s)



DAQ-800 readout board for cDAQ and L1DS

Scouting with DAQ-800 hardware platform:

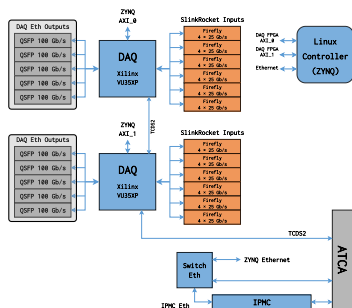

- CMS Phase-2 central DAQ readout board with ATCA form factor
- Production starting in 2024, 5 prototypes will be available in 2024 Q1
- **2 × Xilinx Virtex Ultrascale+ VU35P FPGAs with 8 GB of HBM memory** for output buffering
- **6 × 4 FireFly inputs / FPGA** ⇒ **48 × 25 Gb/s total input** (total: 1.2 Tb/s)
- **5 × QSFP outputs / FPGA** ⇒ **10 × 100 Gb/s total output** (nominal total: 800 Gb/s)



DAQ-800 readout board for cDAQ and L1DS

Scouting with DAQ-800 hardware platform:

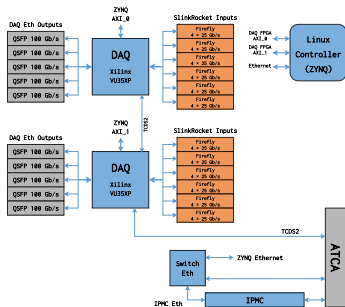

- CMS Phase-2 central DAQ readout board with ATCA form factor
- Production starting in 2024, 5 prototypes will be available in 2024 Q1
- **2 × Xilinx Virtex Ultrascale+ VU35P FPGAs with 8 GB of HBM memory** for output buffering
- **6 × 4 FireFly inputs / FPGA** \Rightarrow **48 × 25 Gb/s total input** (total: 1.2 Tb/s)
- **5 × QSFP outputs / FPGA** \Rightarrow **10 × 100 Gb/s total output** (nominal total: 800 Gb/s)



DAQ-800 readout board for cDAQ and L1DS

Scouting with DAQ-800 hardware platform:


- CMS Phase-2 central DAQ readout board with ATCA form factor
- Production starting in 2024, 5 prototypes will be available in 2024 Q1
- **2 × Xilinx Virtex Ultrascale+ VU35P FPGAs with 8 GB of HBM memory** for output buffering
- **6 × 4 FireFly inputs / FPGA** ⇒ **48 × 25 Gb/s total input** (total: 1.2 Tb/s)
- **5 × QSFP outputs / FPGA** ⇒ **10 × 100 Gb/s total output** (nominal total: 800 Gb/s)
- **Baseline scouting configuration** (assuming a minimal zero suppression of 30%) **fits in 7 × DAQ-800**



Xilinx VCU128 development board

Development of needed firmware on **Xilinx VCU128** development kit:

- 1 × Xilinx Virtex Ultrascale+ VU37P FPGA, similar to a DAQ-800 but with half the connectivity
- 6 × QSFP from mezzanine (HT-Global) ⇒ Up to 24 input links
- 4 × QSFP from board ⇒ Up to 4 × 100 Gb/s total output
- 2 × 4 GB HBM memory stacks, 16 × 256 MB slots / stack
- TCP/IP sender core near-identical to DAQ-800, receiver and pre-processing modules scouting-specific
- **Development and validation in scouting Run-3 demonstrator**

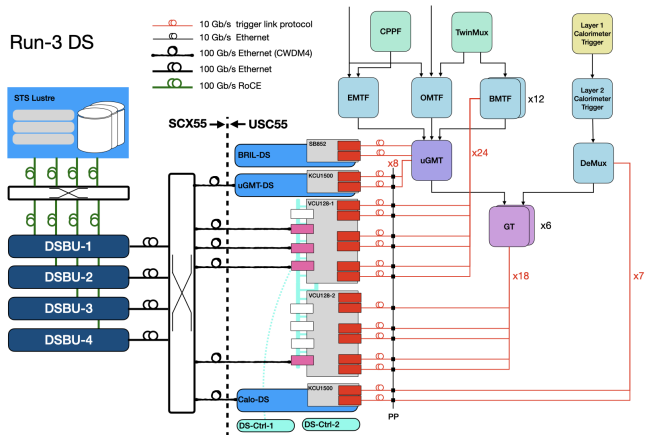


**The L1DS
Run-3 demonstrator**

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)

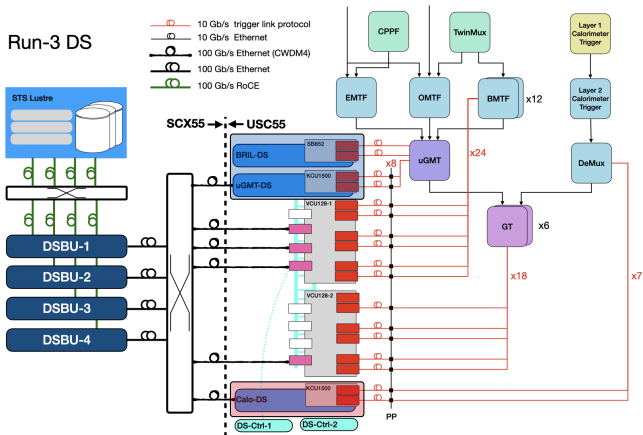


Run-3 demonstrator of the L1DS

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)

Run-3 DS



Global Muon Trigger (μ GMT) \Rightarrow GMT muons

- 8 \times 10 Gb/s links from 1 trg processor \rightarrow 1st KCU1500 and SB-852
- **DMA output** through PCIe gen3 to host server

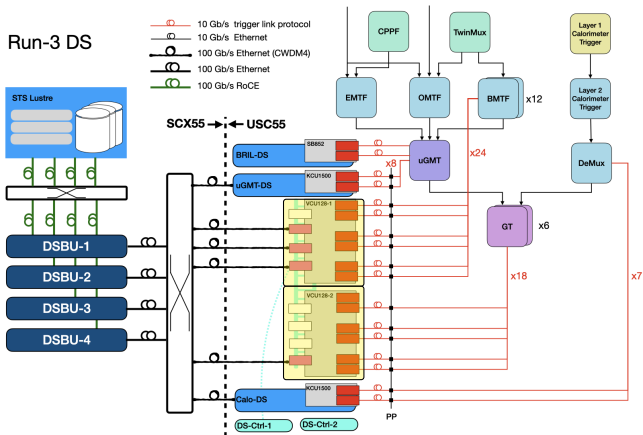
Calorimeter Trigger (DeMux) \Rightarrow jets, e/γ , τ_h , MET

- 7 \times 10 Gb/s links from 1 trg processor \rightarrow 2nd KCU1500
- **DMA output** through PCIe gen3 to host server

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 (μ GMT) \Rightarrow GMT muons

- 8 \times 10 Gb/s links from 1 trg processor \rightarrow 1st KCU1500 and SB-852
- **DMA output** through PCIe gen3 to host server

Calorimeter Trigger (DeMux) \Rightarrow jets, e/γ , τ_H , MET

- 7 \times 10 Gb/s links from 1 trg processor \rightarrow 2nd KCU1500
- **DMA output** through PCIe gen3 to host server


Barrel Muon Track Finder (BMTF) \Rightarrow barrel stubs

- 24 \times 10 Gb/s links from 12 trg processors \rightarrow 1st VCU128
- **TCP/IP output** to surface (3 \times 100 Gb/s links)

Global Trigger output (μ GT) \Rightarrow decision bits

- 18 \times 10 Gb/s links from 6 trg processors \rightarrow 2nd VCU128
- **TCP/IP output** to surface (1 \times 100 Gb/s link)

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 accommodate multiple PCIe boards
- 2×5 PCIe (3.0 \times 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 $12 \times$ BMTF processors
- 2nd VCU128: connected to $6 \times$ GT processors
- Both boards on same PCIe tree
- **Output links from CMS service cavern \rightarrow surface (~ 100 m)**


Firmware structure: optical link input data receiver and decoder



What are the trigger boards sending?

- Phase-2: proprietary protocol, 25 Gb/s links, 65b/67b
- Run-3: 10 Gb/s links, 8b/10b encoding
- **BX data is a fixed record of 6 × 32b frames / link**
- Data frames sent at 250 MHz

Firmware structure: optical link input data receiver and decoder



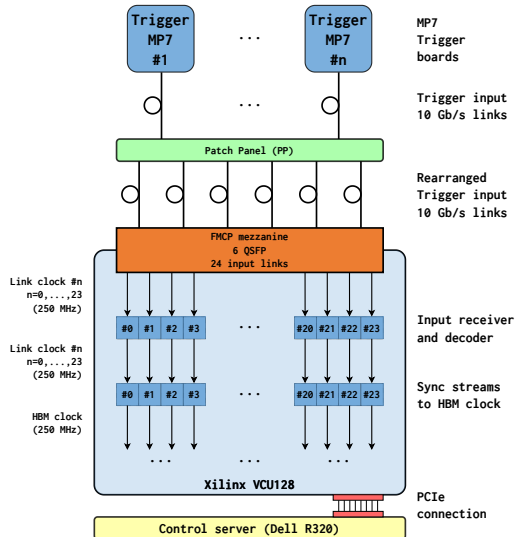
What are the trigger boards sending?

- Phase-2: proprietary protocol, 25 Gb/s links, 65b/67b
- Run-3: 10 Gb/s links, 8b/10b encoding
- **BX data is a fixed record of 6 × 32b frames / link**
- Data frames sent at 250 MHz

Input receiver and decoder:

- Receive up to 24 input links at 10 Gb/s
- Data recovered clock for each input link
- **Input decoding (data/key/invalid words)**

Firmware structure: optical link input data receiver and decoder



What are the trigger boards sending?

- Phase-2: proprietary protocol, 25 Gb/s links, 65b/67b
- Run-3: 10 Gb/s links, 8b/10b encoding
- **BX data is a fixed record of 6 × 32b frames / link**
- Data frames sent at 250 MHz


Input receiver and decoder:

- Receive up to 24 input links at 10 Gb/s
- Data recovered clock for each input link
- **Input decoding (data/key/invalid words)**


Input streams sync:

- **Sync** stream #n from data recovered clock...
- **...to HBM clock** (250 MHz)

Firmware structure: links alignment, data reduction and processing




Firmware structure: links alignment, data reduction and processing




BX-wise input links aligner:

- Align using BX start (green) and known dimension (6 frames)




Firmware structure: links alignment, data reduction and processing

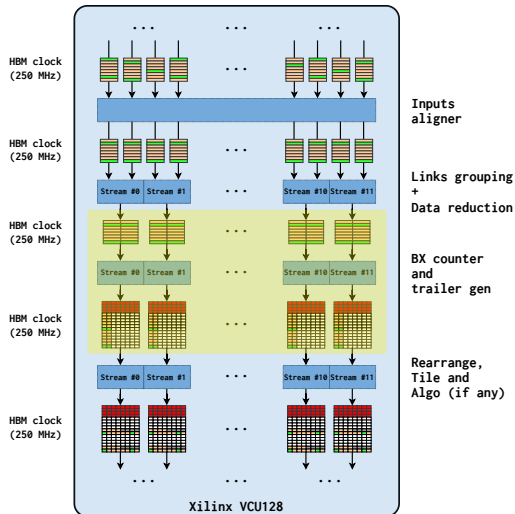


Group links and Data Reduction:

- **Mark BX as "suppressed" based on condition** (e.g. no valid muon stubs, no firing GT algo, ...)




Firmware structure: links alignment, data reduction and processing




BX counter and trailer generator:

- Pad to 256b frames (HBM alignment)
- Assign BX number, given known BX start (green)
- Add trailer at end of orbit with suppressed BX info




Firmware structure: links alignment, data reduction and processing




Rearrange and algo (if any):

- “Transpose” BX grouped data to 256b frames
- **Apply algo (e.g. neural network \Rightarrow [hls4ml](#))** and store in spare space (grey) or send to another stream



Firmware structure: links alignment, data reduction and processing



BX-wise input links aligner:

- **Align using BX start** (green) **and known dimension** (6 frames)

Group links and Data Reduction:

- **Mark BX as “suppressed” based on condition** (e.g. no valid muon stubs, no firing GT algo, ...)

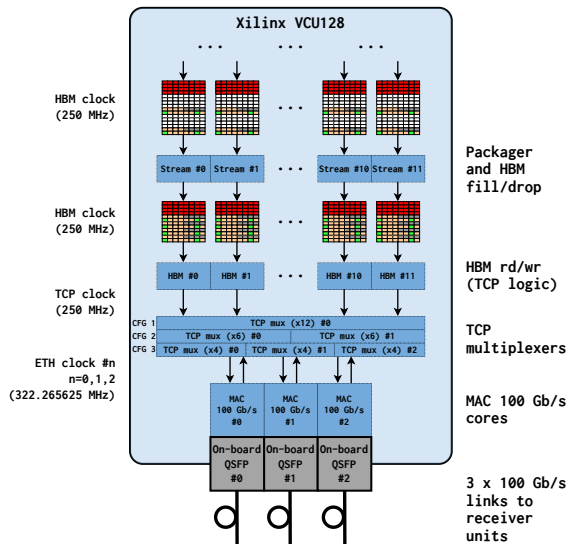
BX counter and trailer generator:

- **Pad to 256b frames (HBM alignment)**
- **Assign BX number, given known BX start** (green)
- **Add trailer at end of orbit with suppressed BX info**


Reshape and algo (if any):

- “Transpose” BX grouped data to 256b frames
- **Apply algo (e.g. neural network \Rightarrow [hls4ml](#))** and store in spare space (grey) or send to another stream

Firmware structure: packager, HBM logic, TCP/IP core




Firmware structure: packager, HBM logic, TCP/IP core

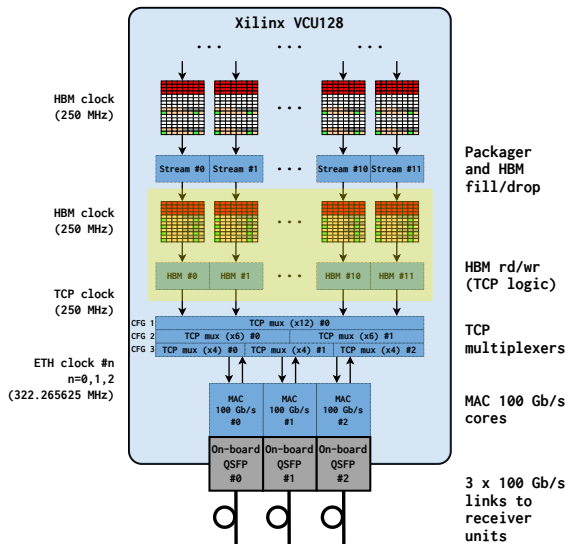


Packager:

- Encapsulation of N orbits in a “Scouting Block”
- Header and trailer added to block
- Fill HBM and handle backpressure




Firmware structure: packager, HBM logic, TCP/IP core

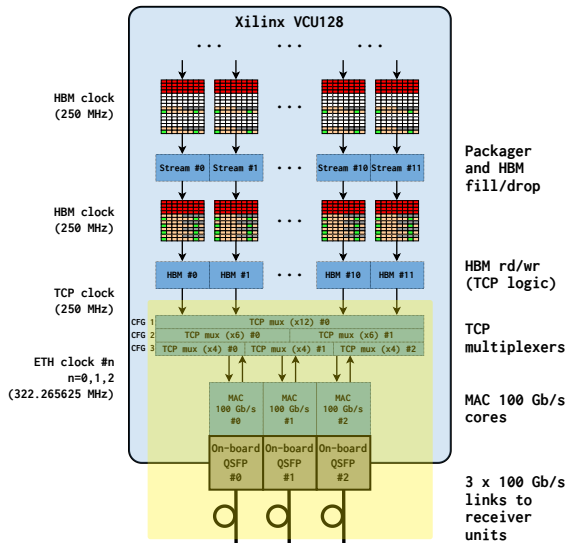


HBM write/read:

- HBM slots are 256b aligned and work at 250 MHz
- **Reserve space for header, fill HBM with payload**, after that write header in reserved space
- HBM reading part connected to TCP logic

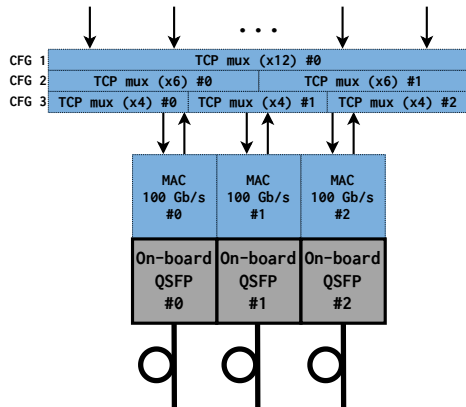


Firmware structure: packager, HBM logic, TCP/IP core

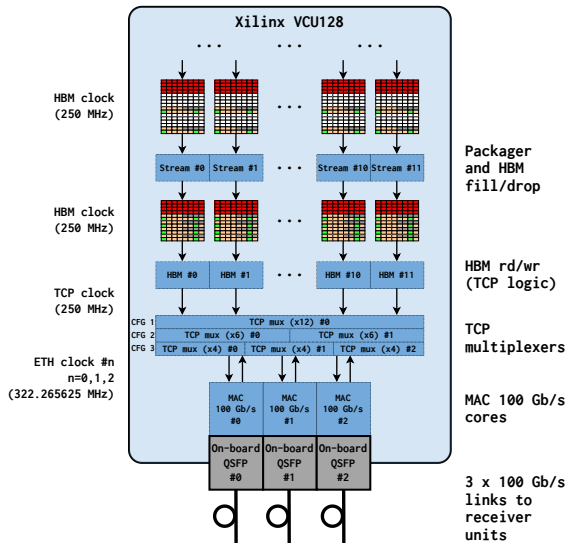


Output streams:

- TCP streams are multiplexed
- **4, 6 or 12 streams per optical 100 GbE output**



Firmware structure: packager, HBM logic, TCP/IP core



Packager:

- Encapsulation of N orbits in a “Scouting Block”
- Header and trailer added to block
- Fill HBM and handle backpressure


HBM write/read:

- HBM slots are 256b aligned and work at 250 MHz
- Reserve space for header, fill HBM with payload, after that write header in reserved space
- HBM reading part connected to TCP logic

Output streams:

- TCP streams are multiplexed
- 4, 6 or 12 streams per optical 100 GbE output

Resource usage and remarks




Legend:

- Trigger input links, alignment
- Processing (Zero Suppression, merging, algo, packager)
- HBM write/read, TCP logic
- 100 GbE output
- Monitoring (PCIe, I2C) and debug (ILA, VIO)

Resource	Utilization	Available	VU37P	VU35P extrapolation	
			Utilization %	Available	Utilization %
LUT	199478	1303680	15.30	871680	22.88
FF	328338	2607360	12.59	1743360	18.83
BRAM	490	2016	24.31	1344	36.45
URAM	48	960	5.00	640	7.50
DSP	6	9024	0.07	5952	0.10

Resource usage and remarks



Legend:

- Trigger input links, alignment
- Processing (Zero Suppression, merging, algo, packager)
- HBM write/read, TCP logic
- 100 GbE output
- Monitoring (PCIe, I2C) and debug (ILA, VIO)



Resource	Utilization	Available	VU37P	VU35P extrapolation	
			Utilization %	Available	Utilization %
LUT	199478	1303680	15.30	871680	22.88
FF	328338	2607360	12.59	1743360	18.83
BRAM	490	2016	24.31	1344	36.45
URAM	48	960	5.00	640	7.50
DSP	6	9024	0.07	5952	0.10

Remarks:

- **Floorplan for 24 × input links, 12 × HBM slots, 3 × 100 GbE cores**
- No machine learning algo included ($O(10^2 - 10^3)$ DSPs)
- TCP Streams to Dell R7515 machines (AMD EPYC 7502P 32-Core)
- TBB pipelines for DAQ software (receive, unpack and write raw data)

Conclusions



Conclusions



Level-1 Data Scouting is technically feasible:

- System planned for CMS Phase-2 upgrade (2029)
- **DAQ-800 platform as trigger readout board**
- **Goal: real-time analysis and tiny events stored for later analysis**

Conclusions




Level-1 Data Scouting is technically feasible:

- System planned for CMS Phase-2 upgrade (2029)
- **DAQ-800 platform as trigger readout board**
- **Goal: real-time analysis and tiny events stored for later analysis**





Run-3 demonstrator already collecting data:

- Started with Global Muon and Calo Trigger scouting
- Extending with BMTF and GT scouting
- **Xilinx VCU128 board for Phase-2 fw development**
- Public results: [CMS-DP-2022-066](#), [CMS-DP-2023-025](#)
- (A lot) More to come!

Run-3 DS



Conclusions



Level-1 Data Scouting is technically feasible:

- System planned for CMS Phase-2 upgrade (2029)
- **DAQ-800 platform as trigger readout board**
- **Goal: real-time analysis and tiny events stored for later analysis**

Run-3 demonstrator already collecting data:

- Started with Global Muon and Calo Trigger scouting
- Extending with BMTF and GT scouting
- **Xilinx VCU128 board for Phase-2 fw development**
- Public results: [CMS-DP-2022-066](#), [CMS-DP-2023-025](#)
- (A lot) More to come!

Additional R&D for Phase-2:

- Test new technologies, e.g. “DAQ-800” board with Versal HBM VH1782
- New ideas for output link protocols, e.g. RoCEv2 (see [Gabriele’s talk](#))

**Backup
Slides**

A new Compact Muon Solenoid experiment for High-Luminosity LHC

- Instantaneous luminosity of up to $7.5 \cdot 10^{34} \text{ cm}^{-2}\text{s}^{-1}$, average number of pp collisions per bunch crossing up to 200
- **Significant detector and hardware upgrade for Compact Muon Solenoid (CMS) experiment**

Level-1 trigger

- Tracks in L1 trigger at 40 MHz
- Particle Flow selection
- 750 kHz L1 output

Tracker


- Increased granularity
- Track-trigger at 40 MHz
- Extended coverage to $|\eta| \lesssim 3.8$

High-Granularity Calorimeter Endcap

- 3D showers and precise timing
- Si, Scint+SiPM in Pb/Cu-W/SS
- 30 ps TOF resolution

Beam Radiation Instrumentation and Luminosity

- Beam abort and timing
- Beam-induced background
- Bunch-by-bunch luminosity: 1% offline, 2% online
- Neutron and mixed-field radiation monitors



DAQ and High-Level Trigger

- Full optical readout
- Heterogenous architecture
- 60 TB/s event network
- 7.5 kHz HLT output

Muon system

- DT and CSC FE/BE new read-out
- New GEM/RPC $1.6 < |\eta| < 2.4$
- GEM coverage up to $|\eta| = 2.9$

Barrel EM calorimeter

- ECAL crystal granularity readout at 40 MHz
- 30 ps e/γ TOF resolution

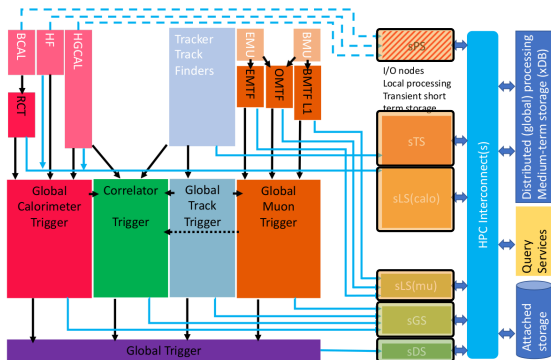
MIP timing detector

- Barrel layer: crystal + SiPM
- Endcap layer: low gain avalanche diodes
- 30 ps TOF resolution

The Level-1 Data Scouting system for Phase-2: architecture

Trigger data captured from spare outputs of L1T boards:

- Same 25 Gbps serial optical links and protocol used for the Level-1 interconnects
- **Global Trigger (GT) decisions** \Rightarrow **sDS**
- **Correlator (CL2), Global Track (GTT), Global Calorimeter (GCT) and Global Muon Trigger (GMT)** \Rightarrow **sGS**
- **Can later be extended to include other systems in later stages** \Rightarrow **sLS**



Scouting system	Source	Objects	Possible extensions
sDS	GT	Final decision word	
sGS	GTT	Vertices, track combinations with thresholds	Zero-suppressed list of input tracks
	GCT	Standalone calorimeter objects	
	GMT	Standalone muon objects	
	CL2	Correlator objects (jets, electrons, etc.)	Zero-suppressed PUPPI candidates
sLS	CL1	PUPPI candidates	Zero-suppressed PF candidates $ \eta \leq 3$ instead

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

Level-1 trigger Data Scouting 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

Phase-2 L1 Data Scouting first physics thoughts/plans

- Enable study of exotic signatures that cannot be fit into the trigger budget
- Use when possible to study with resolution available at Level-1 trigger
- High combinatorics for the L1 budget: $W \rightarrow 3\pi$, $W \rightarrow D_s\gamma$, $H \rightarrow \rho\gamma$, $H \rightarrow \phi\gamma$, ...
- High rate for the L1 budget: multiple soft (b-)jets, displaced (soft) leptons, ...
- Slow or long-lived objects: Heavy Stable Charged Particles (HSCPs) across multiple BXs
- Scouting on the whole set of L1 tracks: Soft Unclustered Energy Patterns (SUEPs)

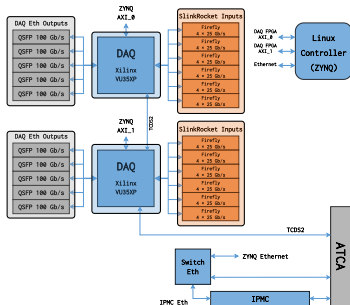

Monitoring at the bunch crossing rate

- Unaffected by issues in readout system, e.g., if detector is blocked due to excessively high rate from a trigger object
- Study Level-1 trigger pre- and post-firing without special trigger configurations
- Per-bunch luminosity measurements

DAQ-800 readout board for cDAQ and L1DS

Scouting with DAQ-800 hardware platform:

- CMS Phase-2 DAQ readout board with ATCA form factor
- Production starting in 2024, 5 prototypes will be available in 2024 Q1
- **2 × Xilinx Virtex Ultrascale+ VU35P FPGAs with 8 GB of HBM memory** for output buffering
- **6 × 4 FireFly inputs / FPGA ⇒ 48 × 25 Gb/s total input** (total: 1.2 Tb/s)
- **5 × QSFP outputs / FPGA ⇒ 10 × 100 Gb/s total output** (nominal total: 800 Gb/s)
- **Baseline scouting configuration** (assuming a minimal zero suppression of 30%) **fits in 7 × DAQ-800**



Scouting system	Source	Links (baseline)	Links (upstream ZS)	fraction of DAQ-800 board inputs
sDS	GT	12	12	0.25
sGS	GTT	24	24 + 48 (Tracks ZS)	0.5(+1)
	GCT	6	6	0.125
	GMT	18	18	0.375
	CL2	30	30 + 24 (PUPPI ZS)	0.625 (+0.5)
sLS	CL1	216 (PUPPI)	84* (PF Inj3 ZS)	4.5 (1.75)
Total		306	246	6.375 (5.125)

Firmware structure: machine learning inference

Deploy ML applications directly on FPGA in data scouting pipeline:

- [hls4ml](#) for generation of neural networks in Hardware Description Language
- Recalibration of Level-1 muon primitives: [CMS-DP-2022-066](#)
- Real-Fake Level-1 muon pair classification: [CMS-DP-2022-066](#)
- Muon barrel stub primitives fit (work in progress!)

