Trigger and DAQ: Challenges and Opportunities

2026 Update: European Strategy for Particle Physics

Jun 23–27 2025, Venice Lido

Thea Klæboe Årrestad (ETH Zürich)
Dorothea vom Bruch (CPPM Marseille)

Timelines and synergies

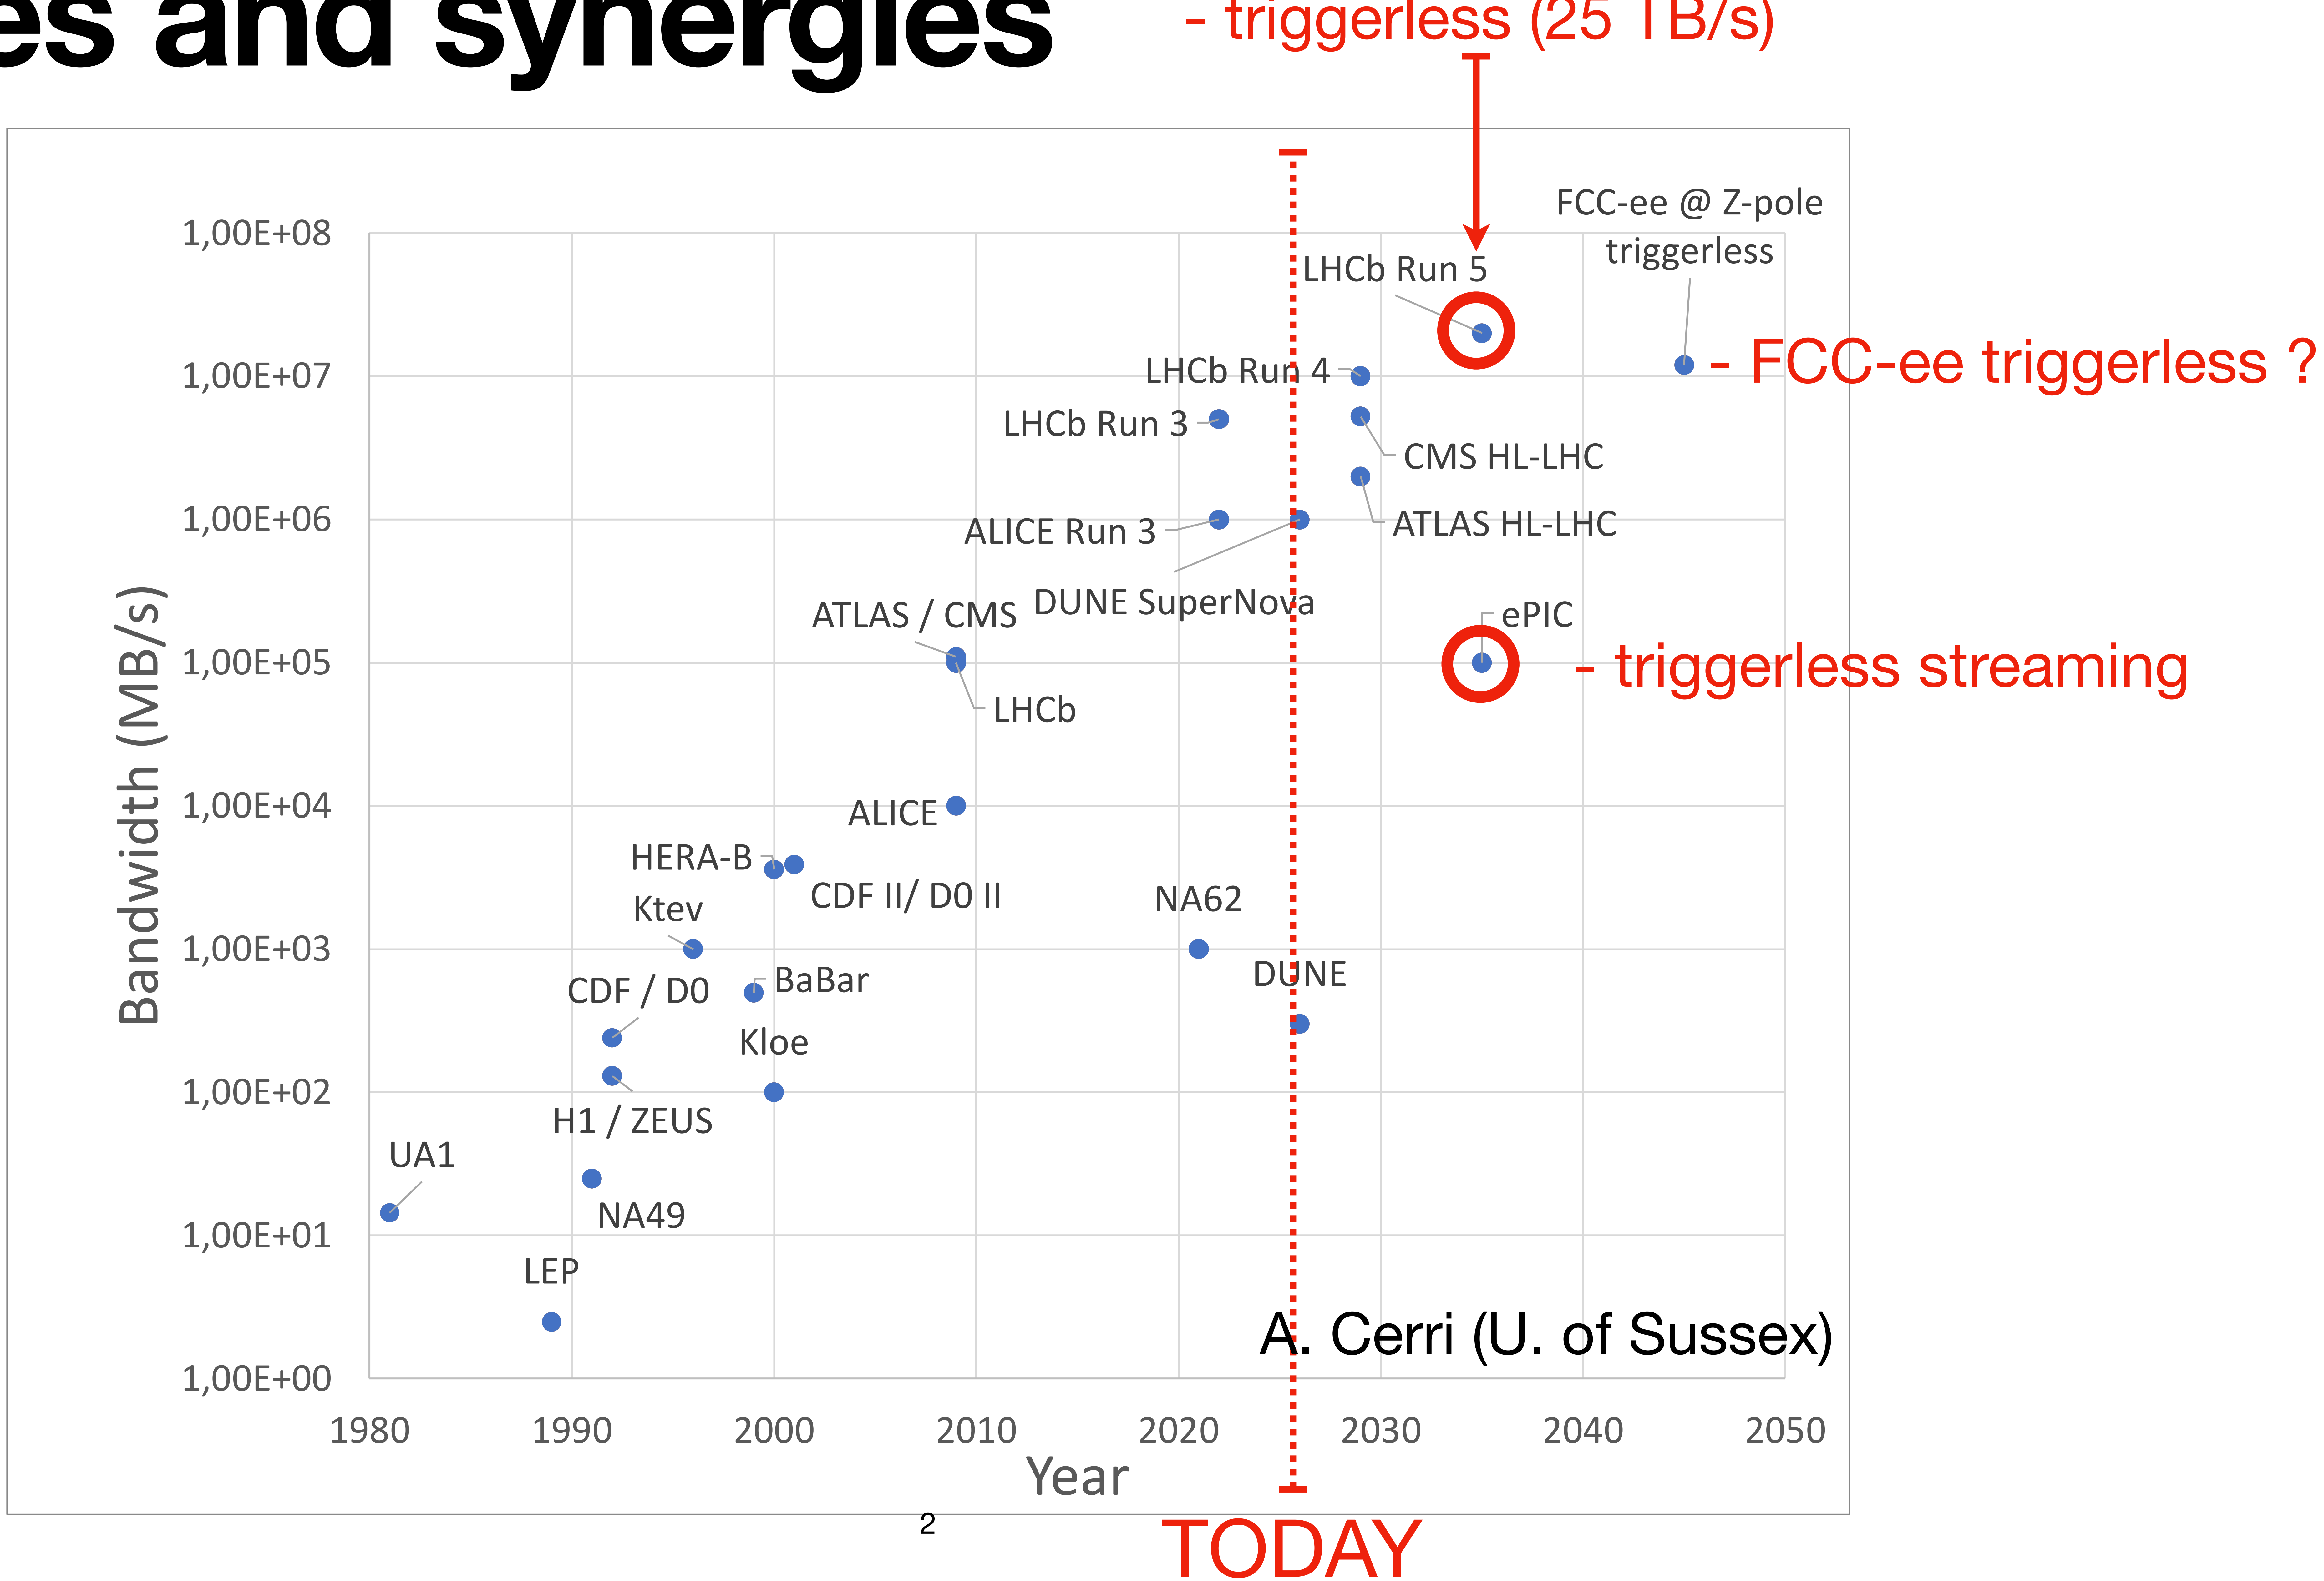

#241 #95

FCC-ee TDAQ challenges

2.3 × 10^{36} cm^{-2} s^{-1} at Z-pole!

- FCC-ee at Z-pole:
  - highest instantaneous luminosities ever achieved
  - each sub-detector **must be capable of read-out at 50 MHz**
  - physics rate ~100 kHz, **large background rate in innermost layers**

- Aim is “trigger-less” design: software -based online selection

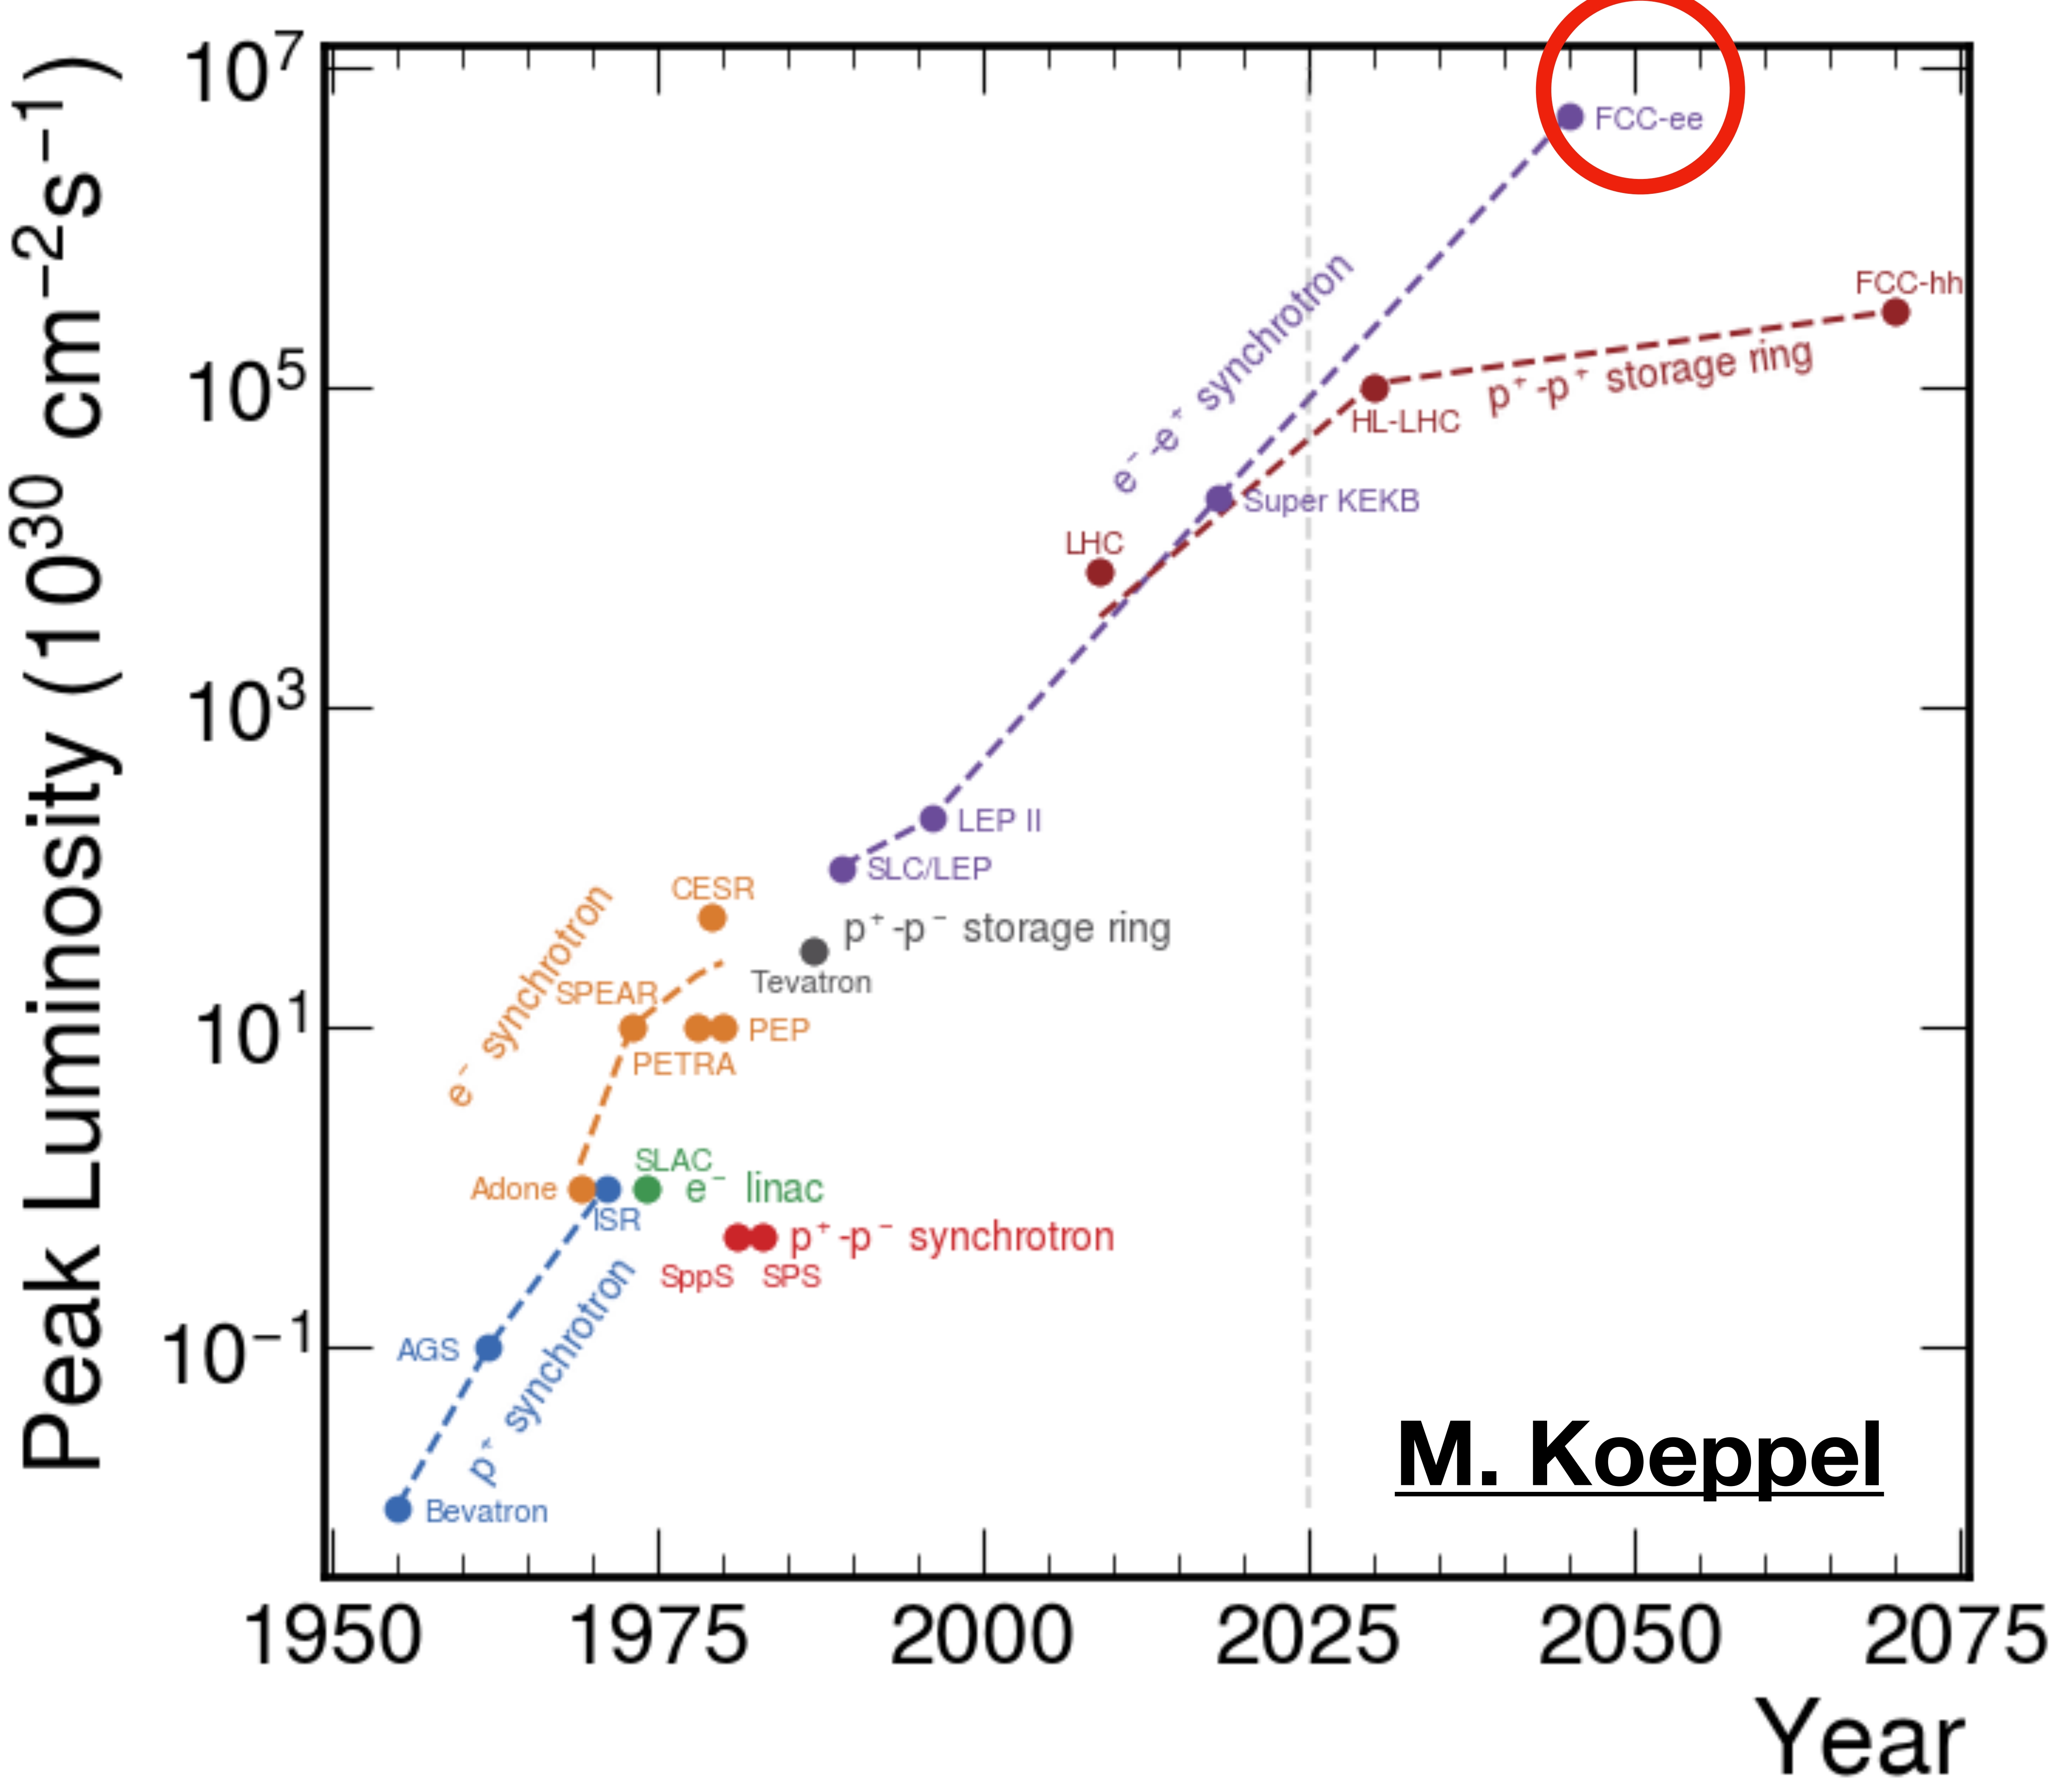

FCC-ee TDAQ challenges

- TDAQ requirements **highly dependent** on detector technology. Potential challenges:
  - TPC **cannot read out every 20 ns**, hardware-based first filtering might be necessary
  - Noise from **increasingly granular**
  - **Limited by single read-out ASIC data transmission capacity**

Rate estimates, FCC-ee Z-pole (F. Bedeschi)

Subdetector	Untriggered
Idea vertex	~1 TB/s
Idea DCH	~500 GB/s
Idea DR Calorimeter	~10 TB/s
Idea Luminometer	~20 GB/s
Idea Muon System	~400 MB/s
Total	**11.7 TB/s**

1) Extreme hit rates in inner layers

2) Fine granularity, but noisy

Read-out of vertex detectors

- Large background from incoherent pairs creation (mostly forward, but substantial fraction within vertex detector acceptance)
- For vertex detector → 200 MHz/cm² for innermost layer
  - Untriggered: **24.4 Gbit/s** / 2-chip module
  - **2.2 Tb/s for Layer 1** (x10 less in Layer 2)
- Current MAPS technology **cannot match this rate**, all FCC-ee detector concepts foresee MAPS
  - Max readout speed achievable on chip →**3.2** - 6.4 Gbit/s (current ARCADIA)
- Trigger-less readout seems extremely challenging (also: can we even handle this output?)
- **Need on-detector compression!**

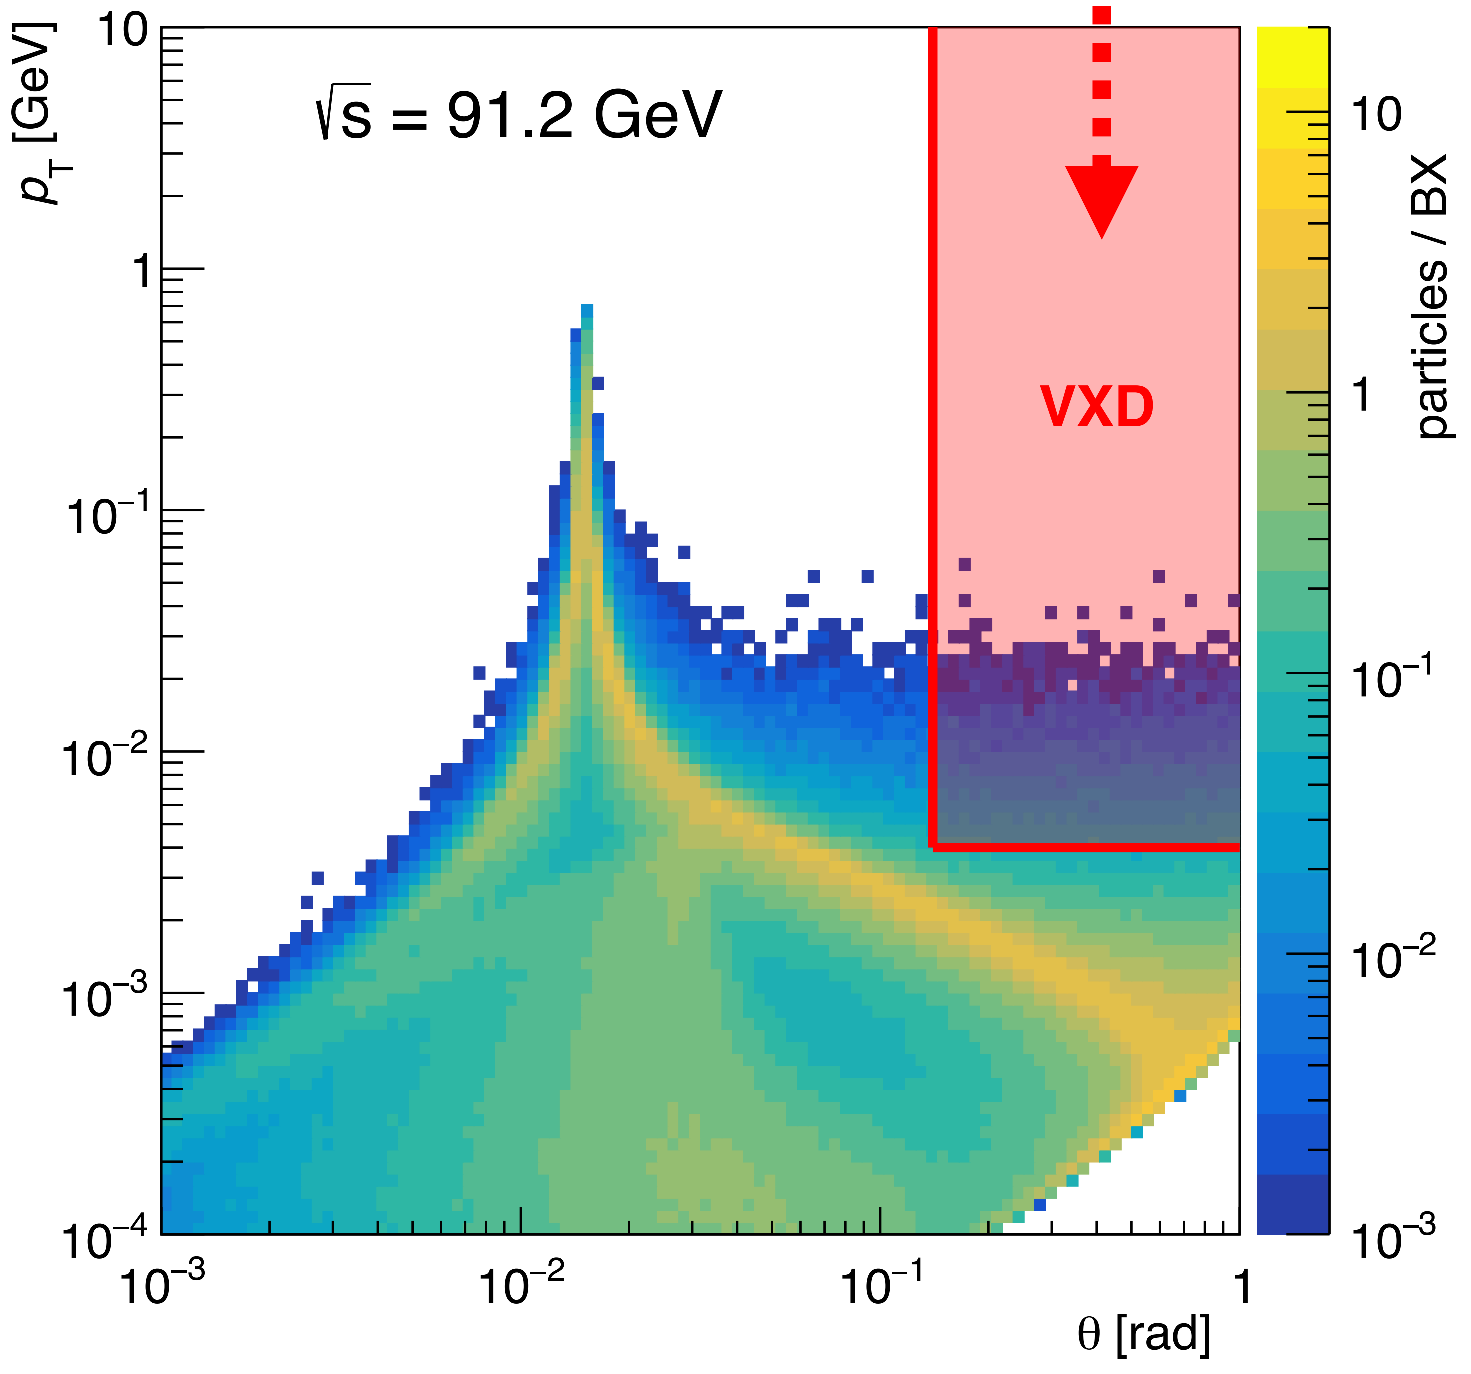

See Fabrizio's and Daniela's talk

FCC-hh TDAQ challenges

- Increased granularity and acceptance:
  - 130→1000 pile-up, x6 increase in peak luminosity
  - 250 TB/s for calorimetry + muon systems (x10 HL-LHC ATLAS/CMS).
  - O(1) PB/s of tracker data.
- Full read-out of calo/muon systems ~possible
- Tracker readout impossible at full rate (even with readout technology available, radiation hardness and infrastructure constraints)
- To keep trigger rates at acceptable levels, offline algorithms must be migrated to trigger!
  - Track readout+hardware trigger for outer layers (>20 cm)

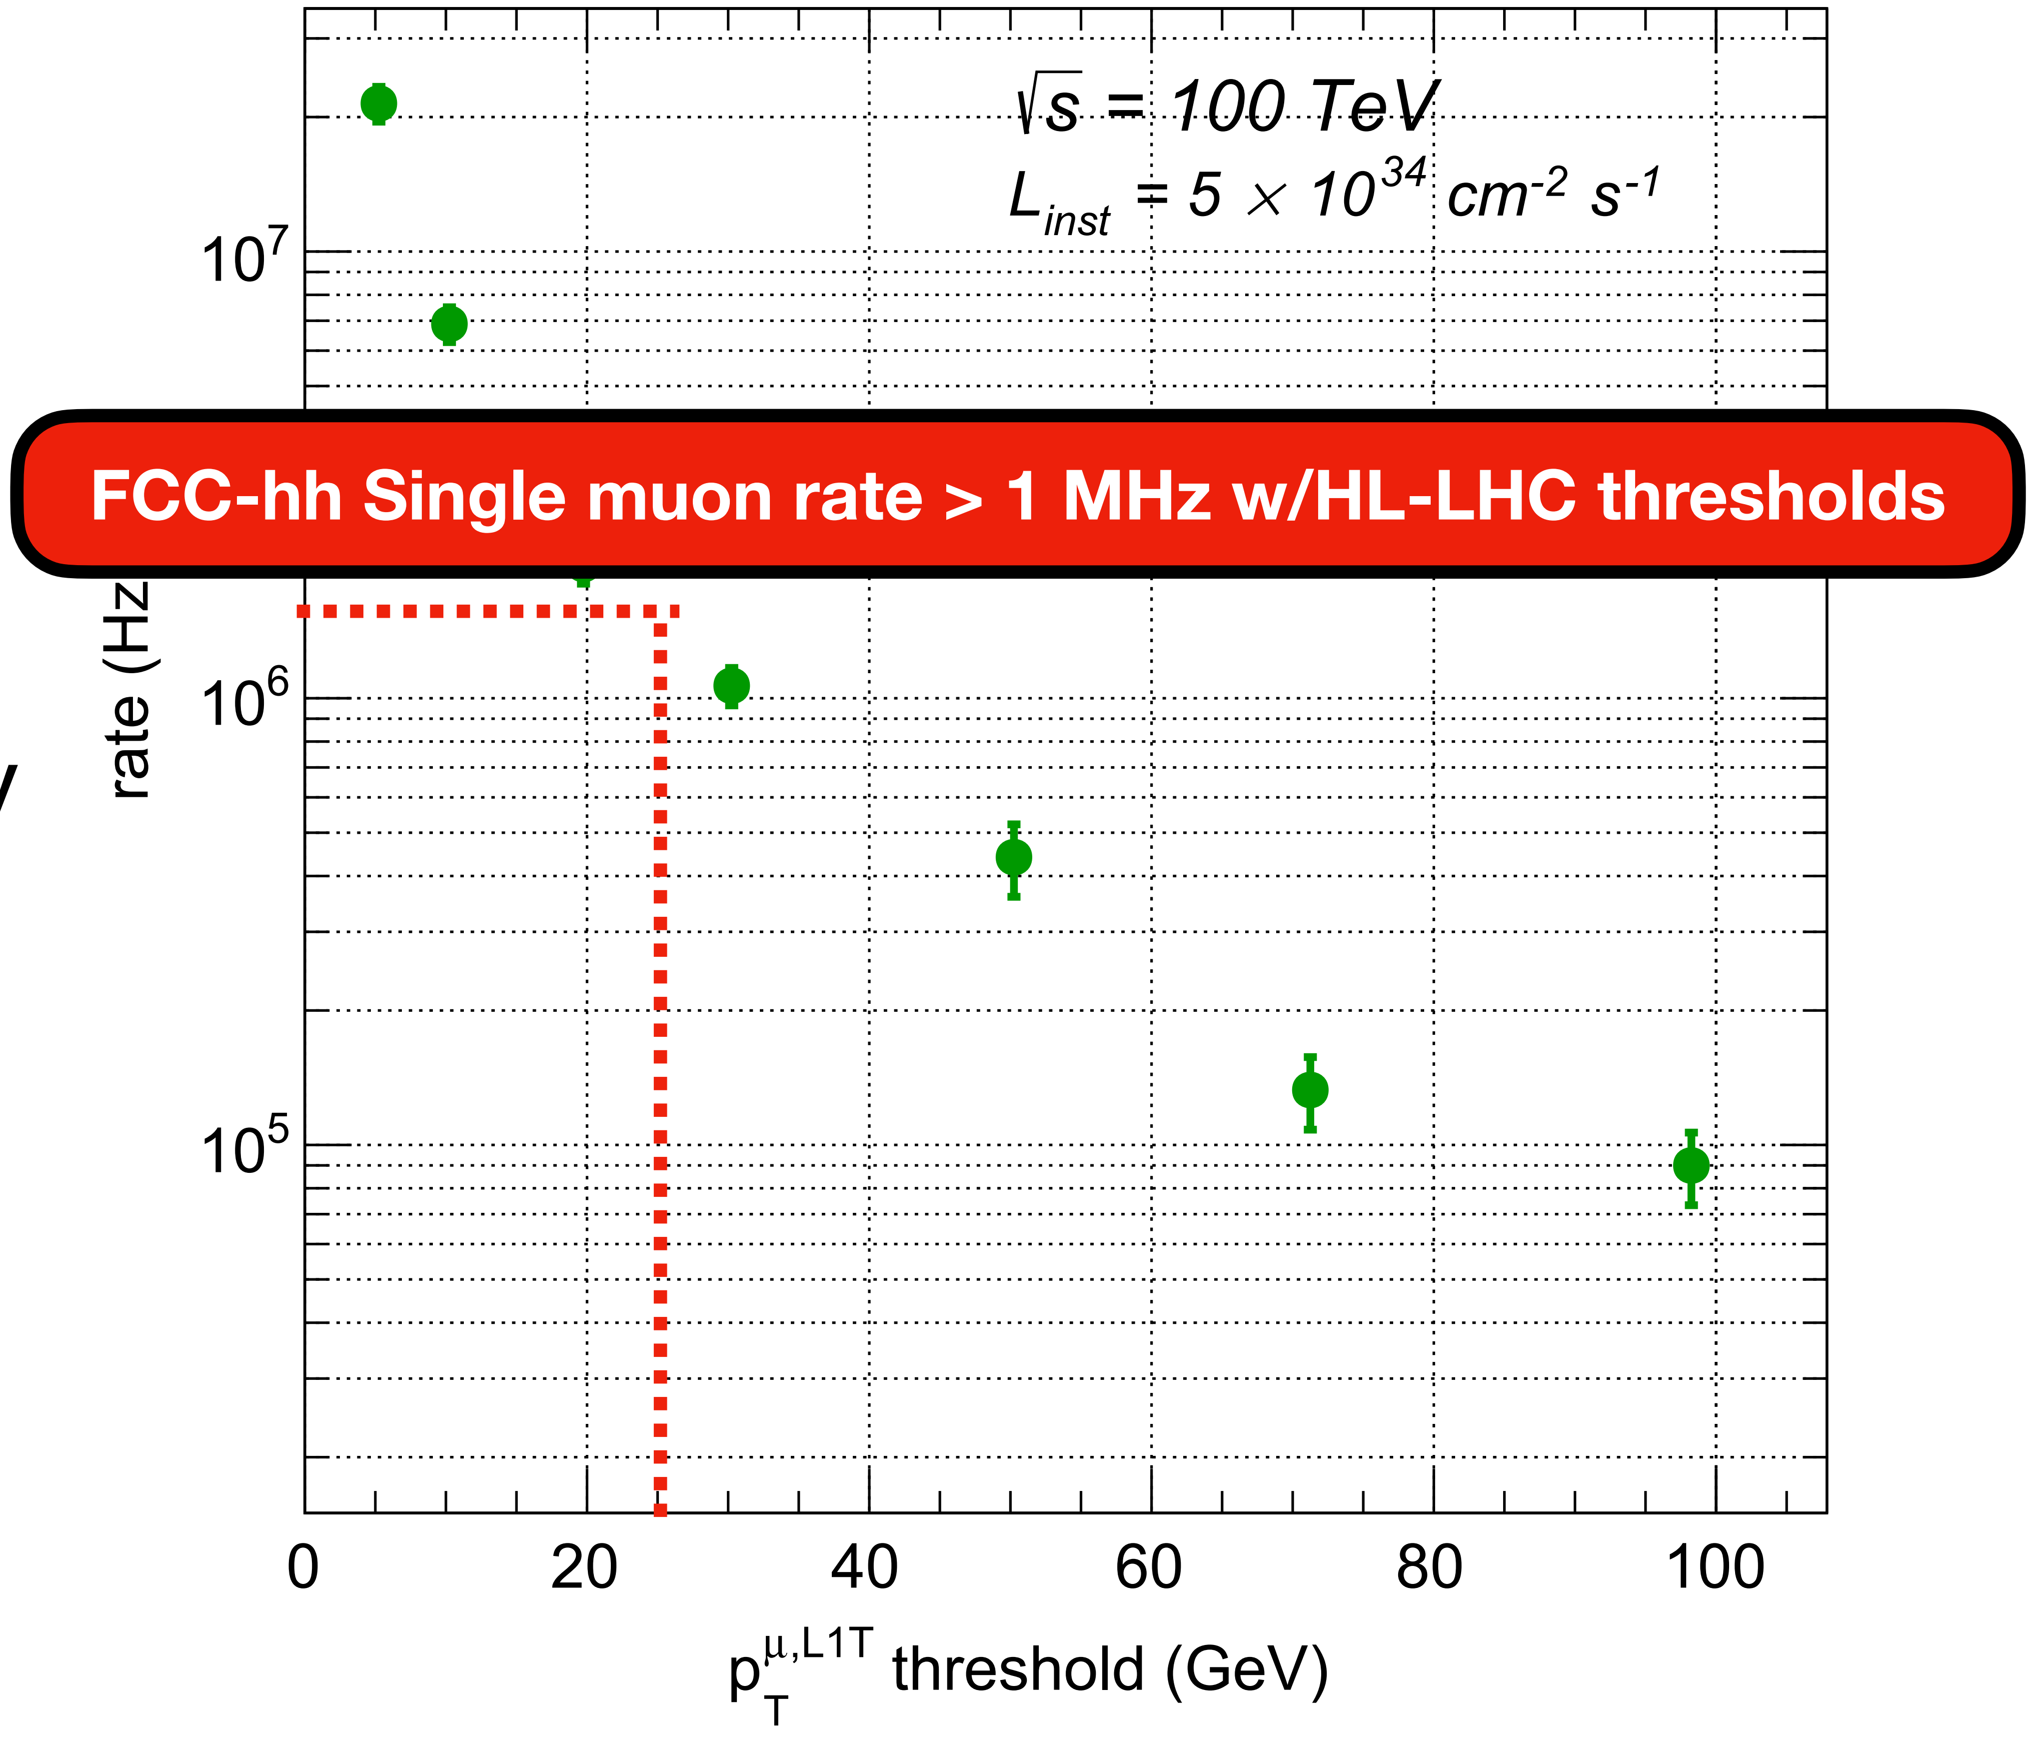

Trigger Type	Threshold at 14 TeV (GeV)	Threshold at 100 TeV (GeV)
Single Muon	25	78
Photon/ Electron	30	150
Jet	120	300

Abada, A., Abbrescia, M., AbdusSalam, S.S. et al. FCC-hh: The Hadron Collider

How do we address these?

Opportunities for novel designs

- FCC-ee in streaming trigger less mode faces significant challenges
  - ~ 160 Pb/year for FCC-ee detectors
  - Must **reduce data directly at source, aggregate and stream to offline, while staying within tight material and power budgets**
  - Can learn from HL-LHC LHCb & ALICE, EIC ePIC how to do **real time physics extraction with offline-like resolution**, requiring real-time calibration and alignment
- FCC-hh can read out calorimeter and muon systems at full rate with tomorrows technology
  - Tracker (probably) not! Will need **compression/reconstruction at-source and/or track trigger for a PU=1,000 scenario**
- **Must invest in the current and upcoming experiments to develop workflows that harness emerging hardware technologies!**

Data readout bottleneck

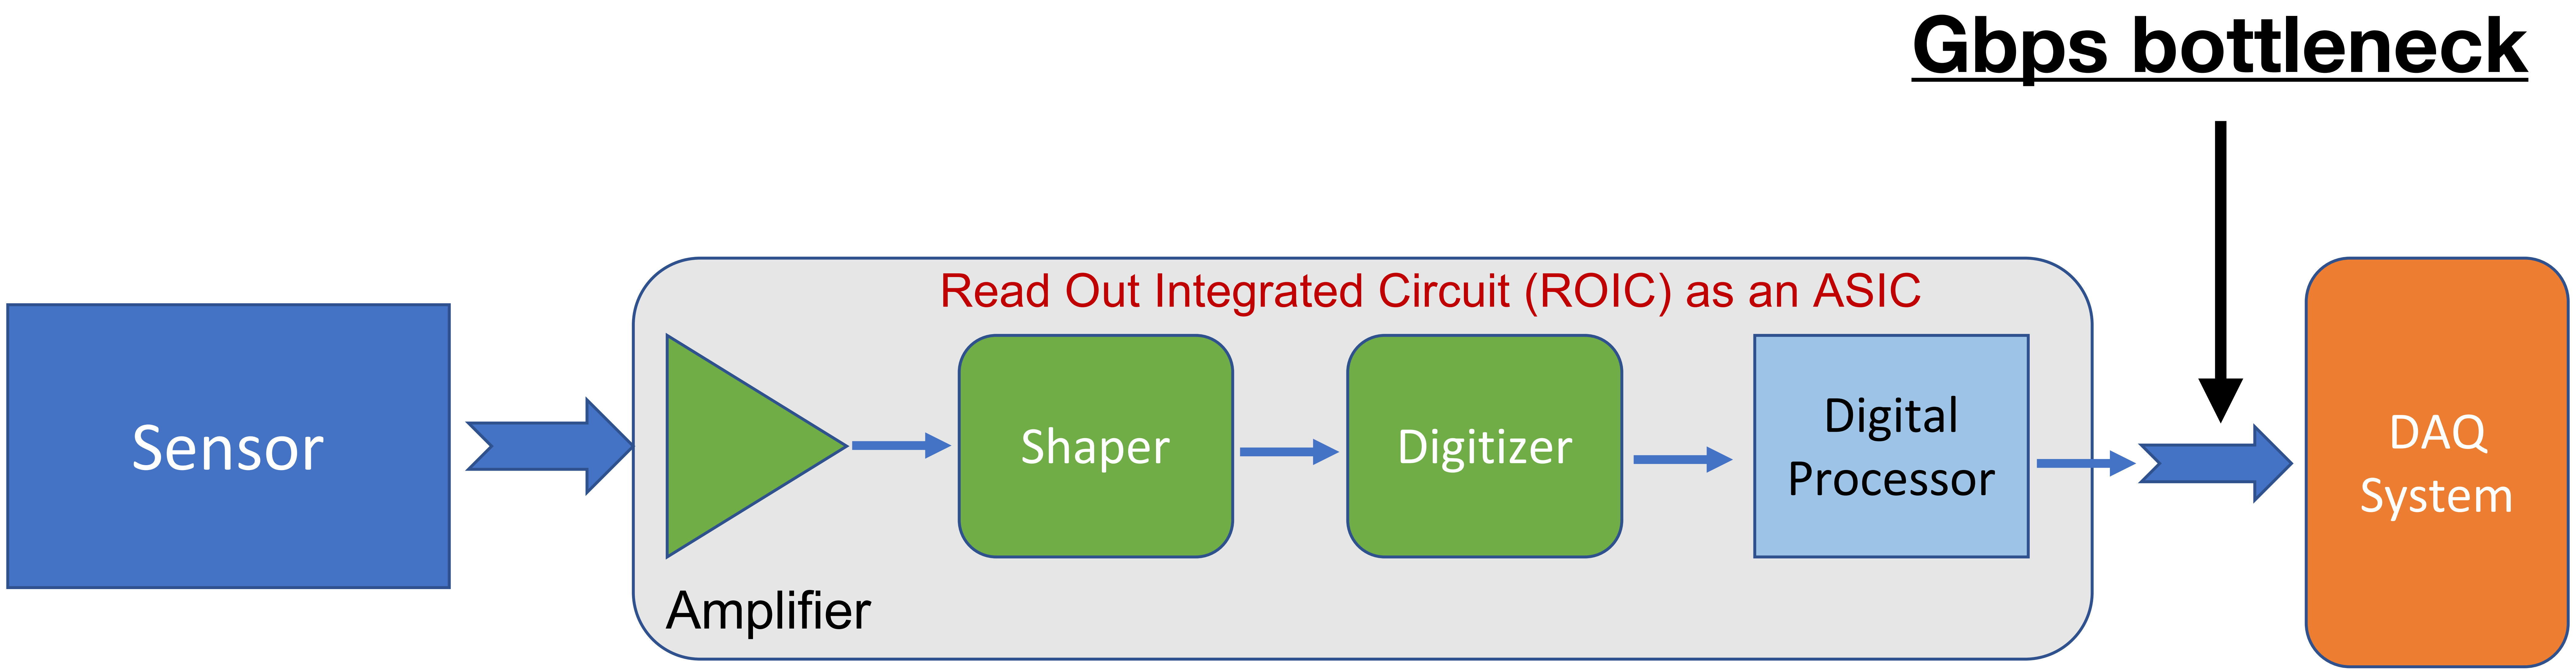

Example vertex detector inner layer from before (#241 FCC-ee):

32 bits pixel data in inner layer @ 200 MHz/cm²

➡ **24.4 Gbit/s**

Two options:

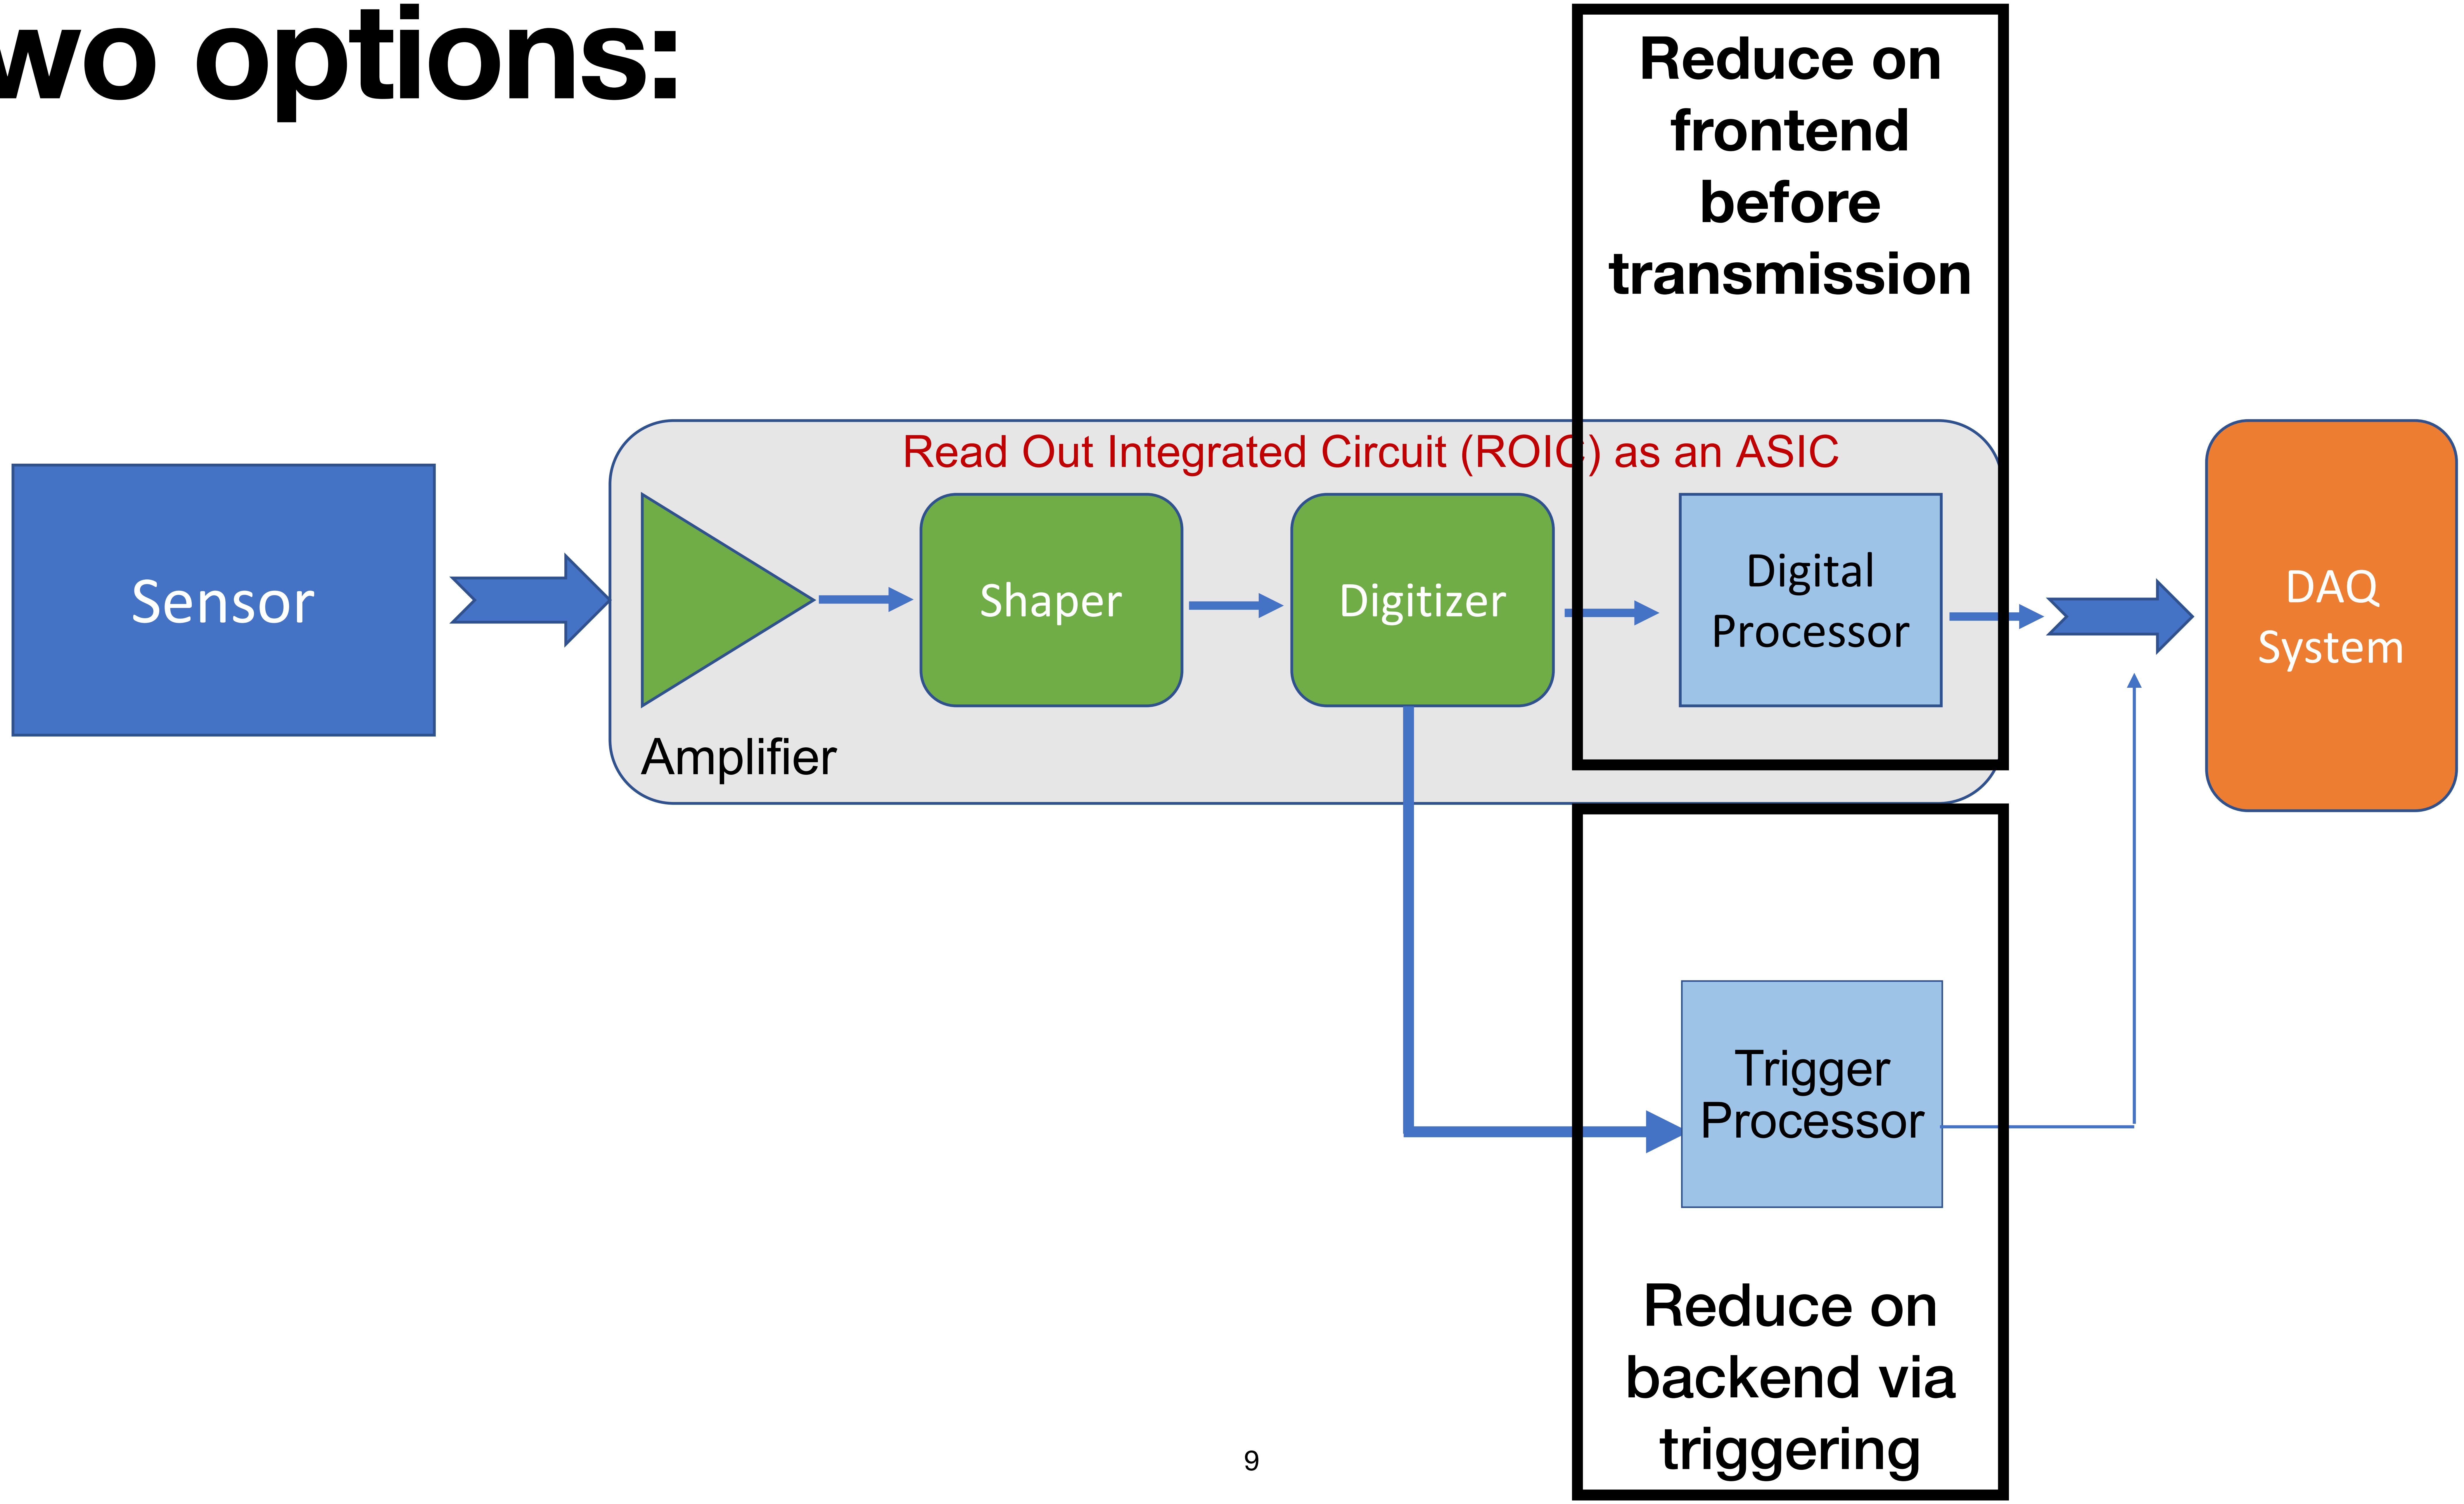
Reduce on frontend before transmission
Read Out Integrated Circuit (ROIC) as an ASIC
Sensor
Amplifier
Shaper
Digitizer
Digital Processor
DAQ System
Trigger Processor
Reduce on backend via triggering

Two options:

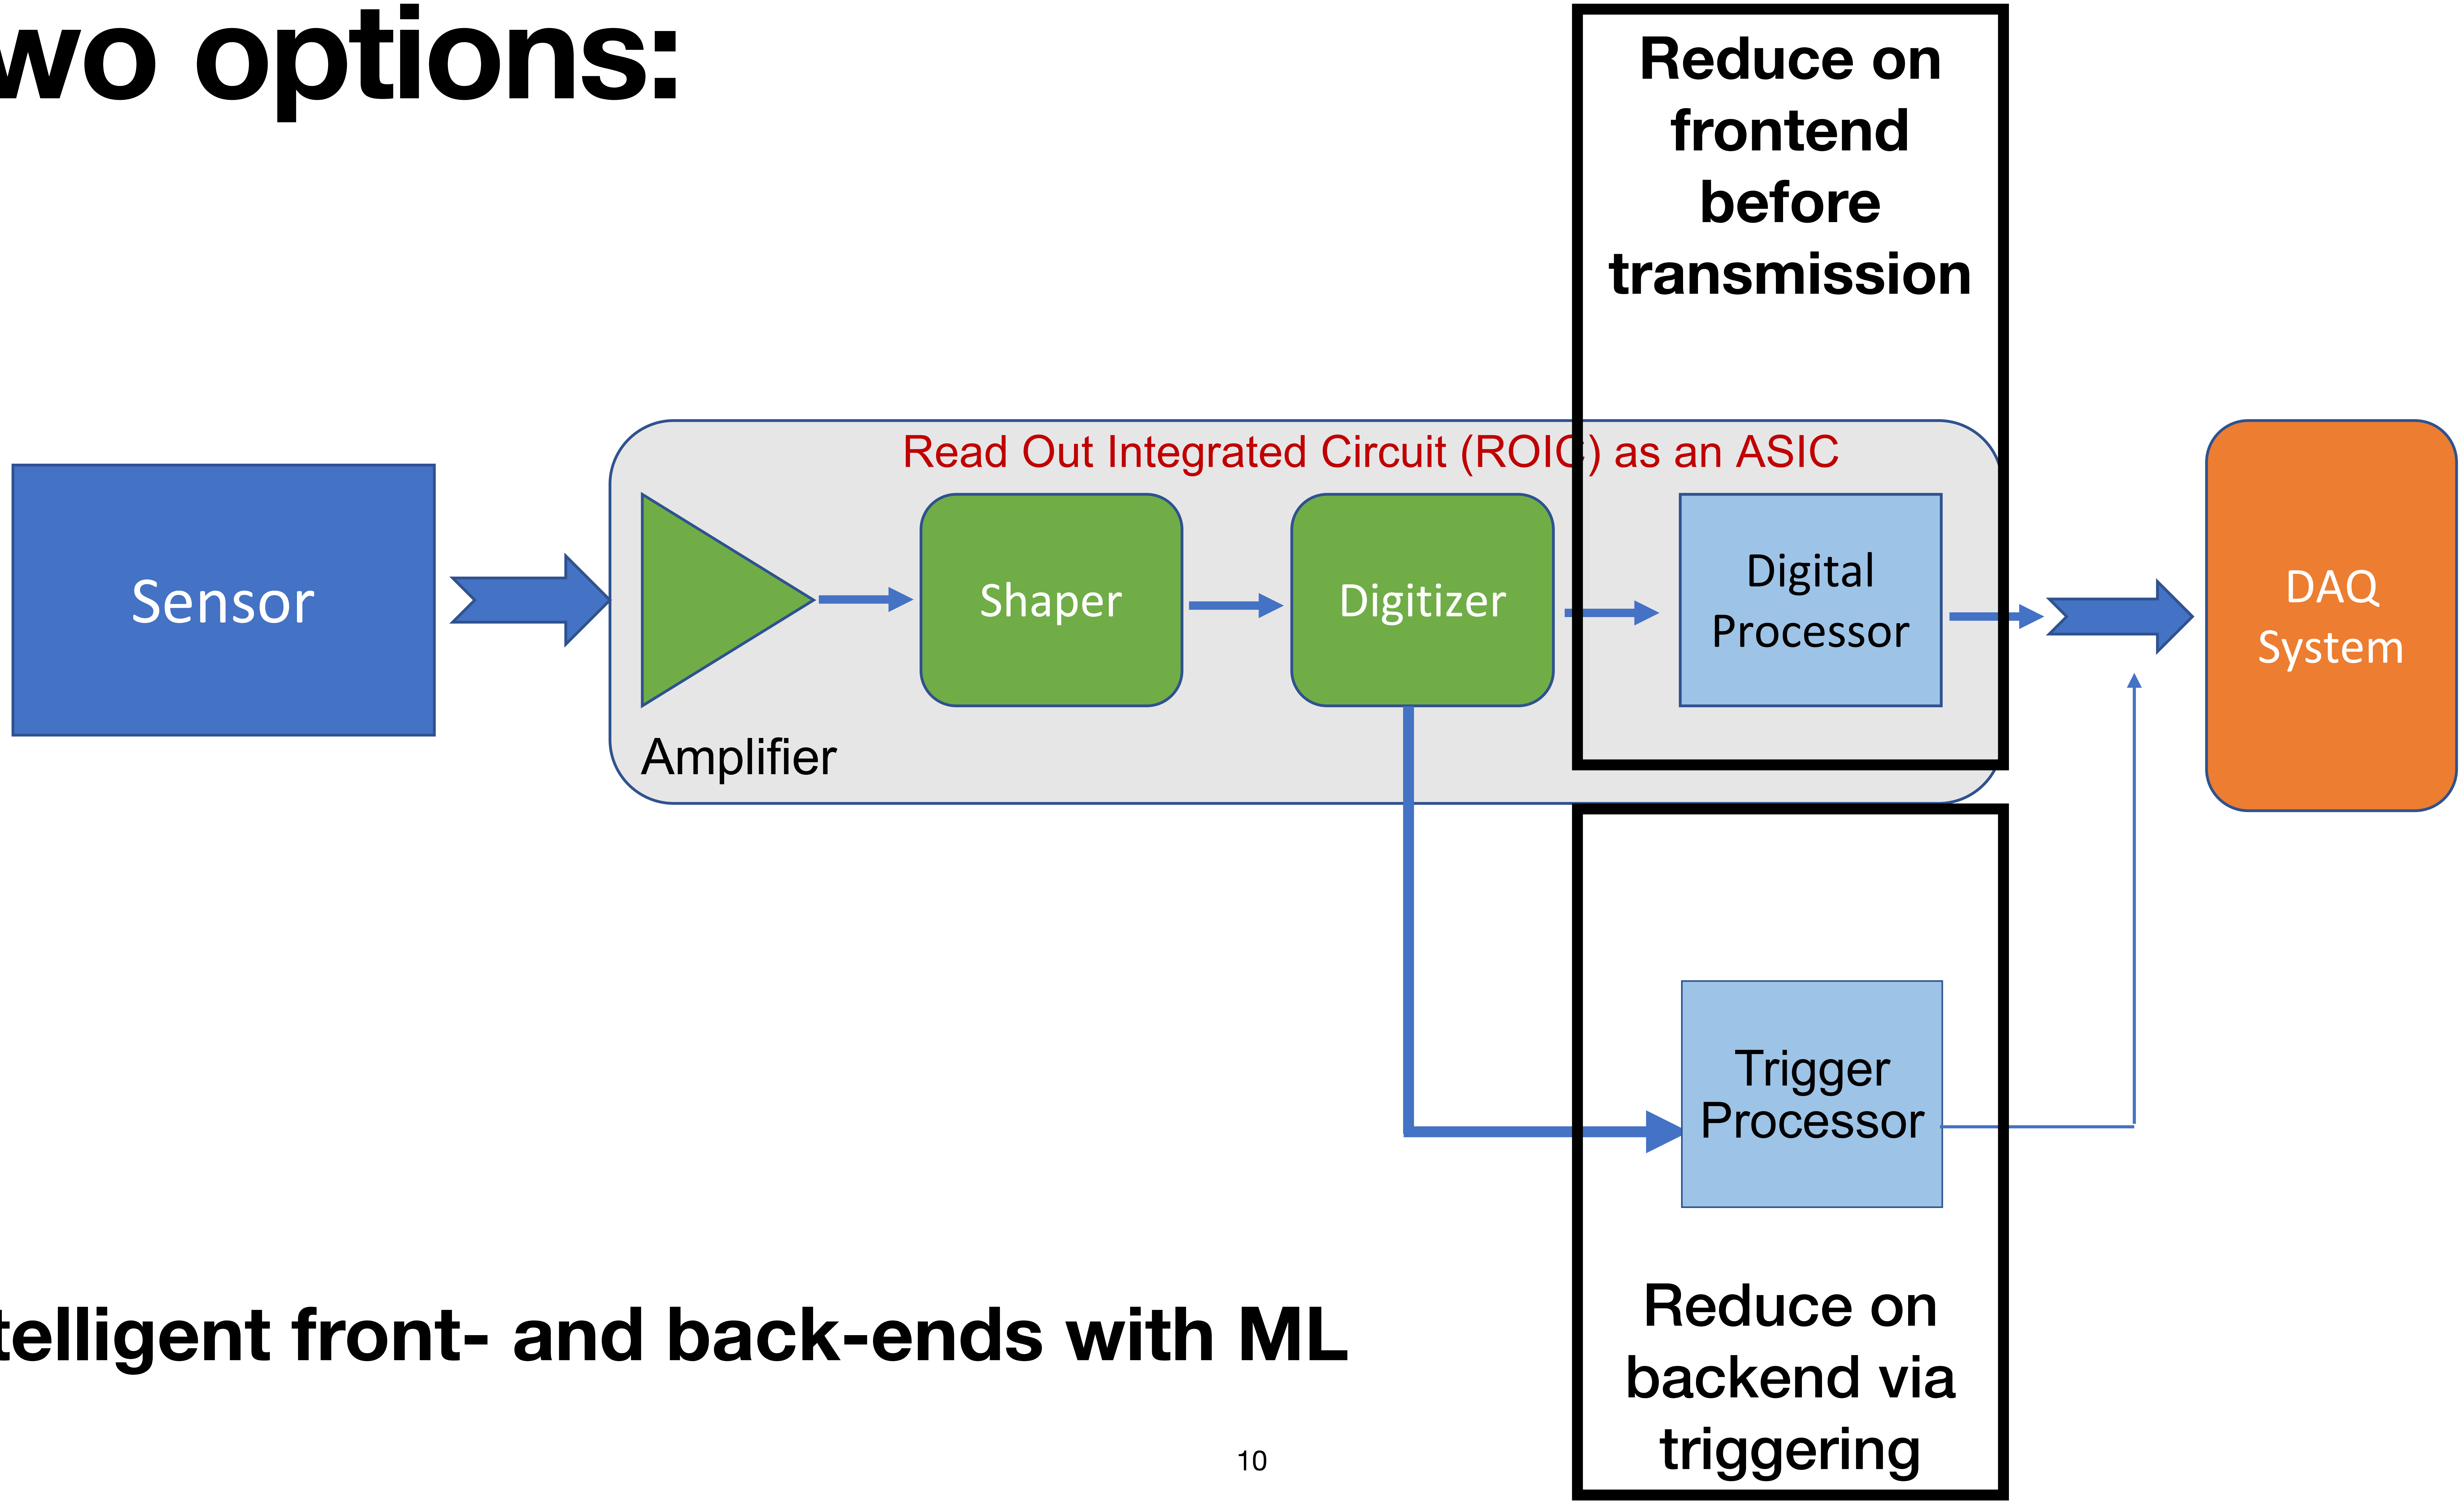

→**Intelligent front- and back-ends with ML**

Fast Machine Learning for HEP

ML on specialised hardware (ASICs, FPGAs) could help us address our data processing challenges

HEP tasks not well-represented by industry-driven tools

Have developed our own tools!

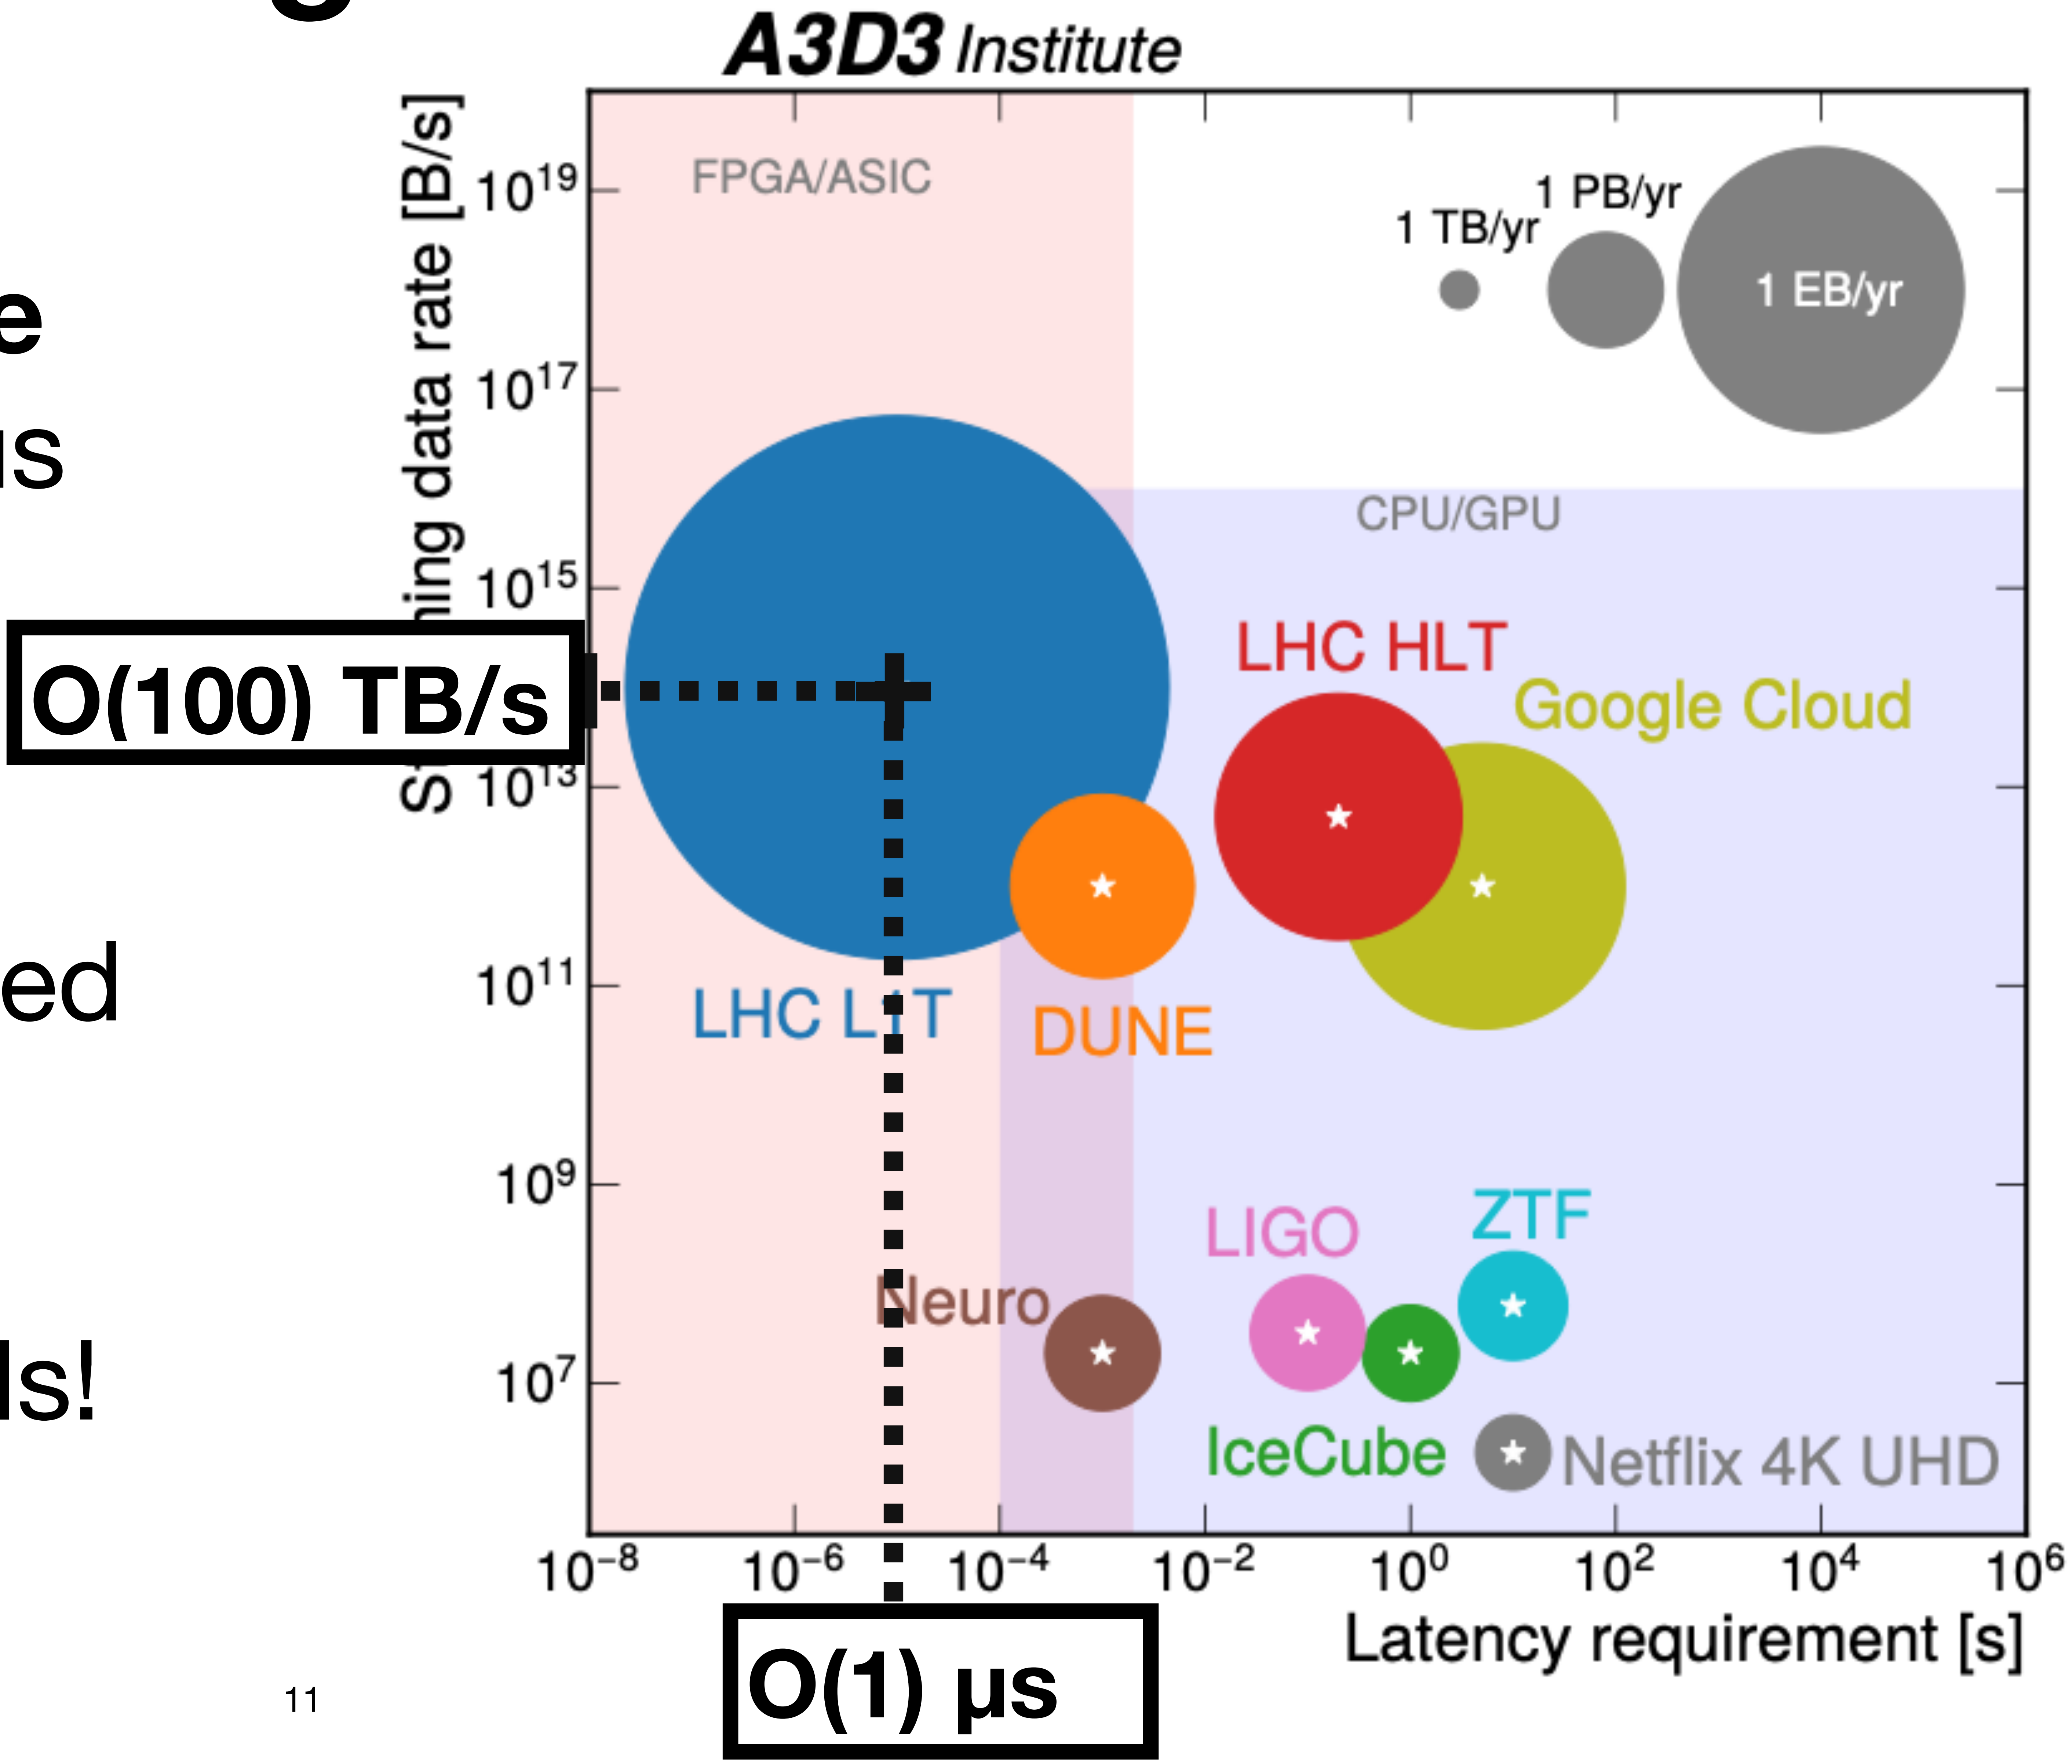

ML inference at low latency

HEP Tools and Communities

Model (quantized/pruned)

HEP quantization libraries:

QONNX

HGQ

QKERAS

HEP hardware ML libraries:

hls4ml

Conifer

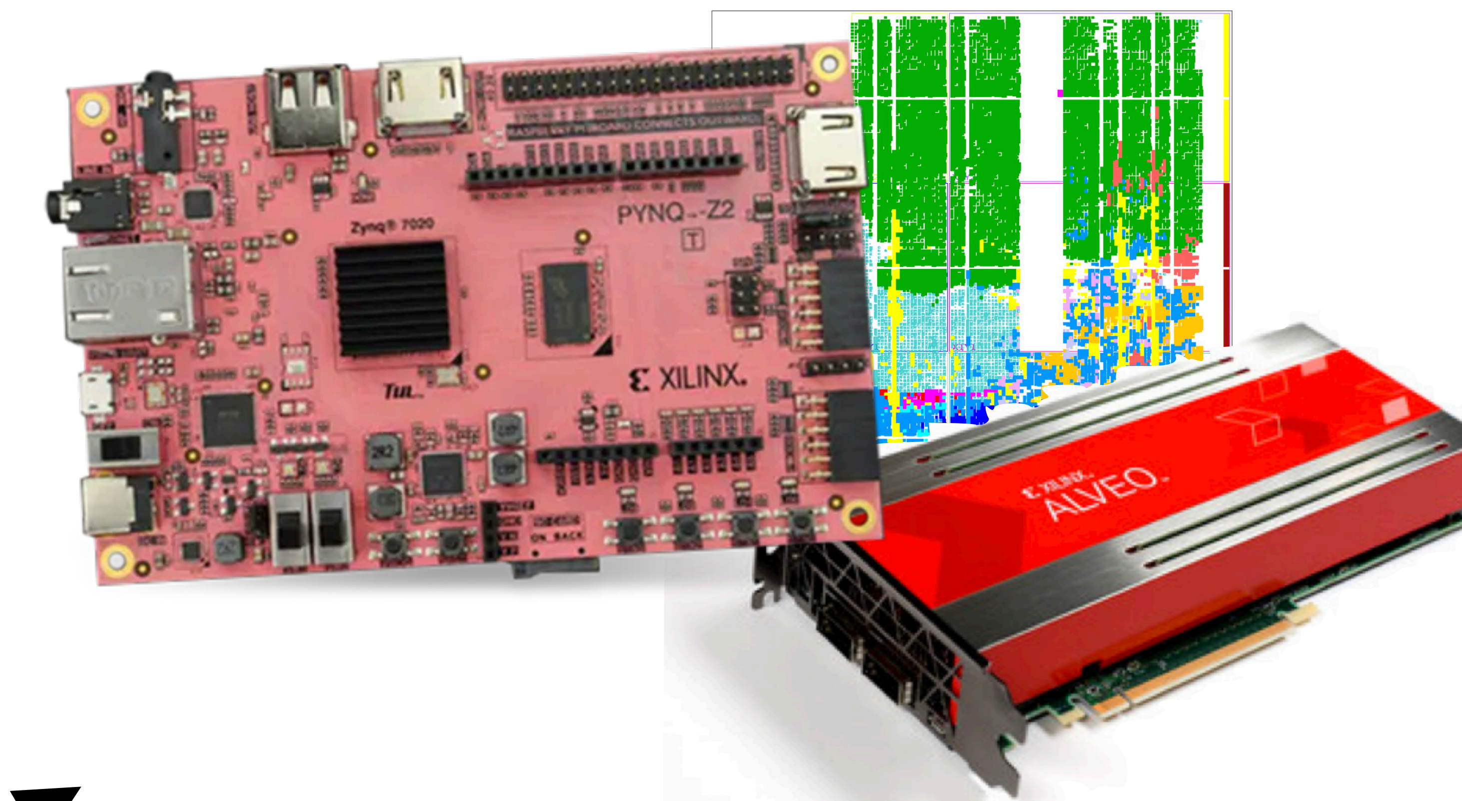

Co-processing kernel (Xilinx accelerators/SoCs)

FPGA custom designs (eg trigger algorithms)

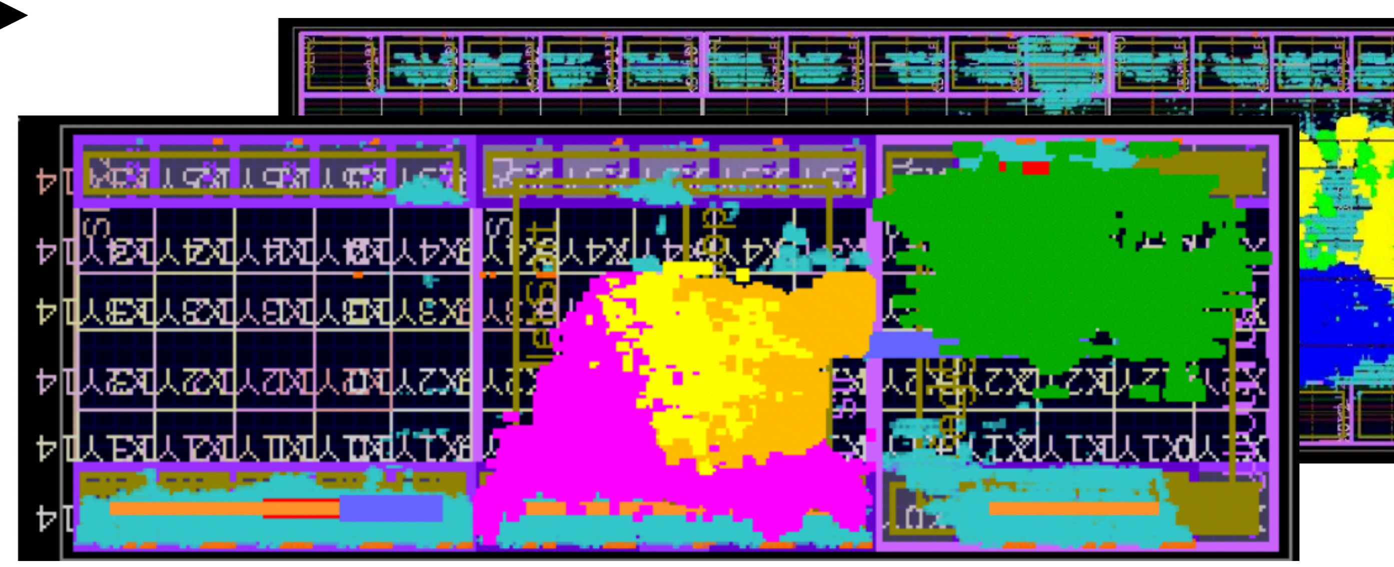

ASICs

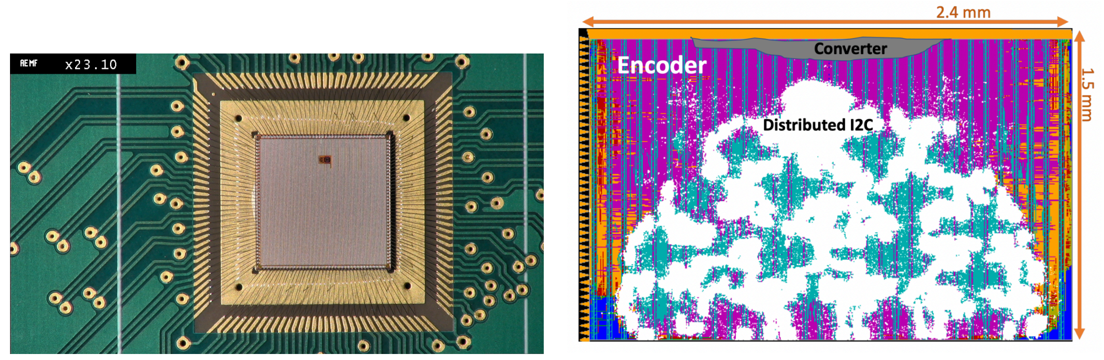

2024: Neural hardware triggers making decisions in LHC experiments!

CMS Experiment:
- Anomaly detection in 50 ns
- 300 events/second

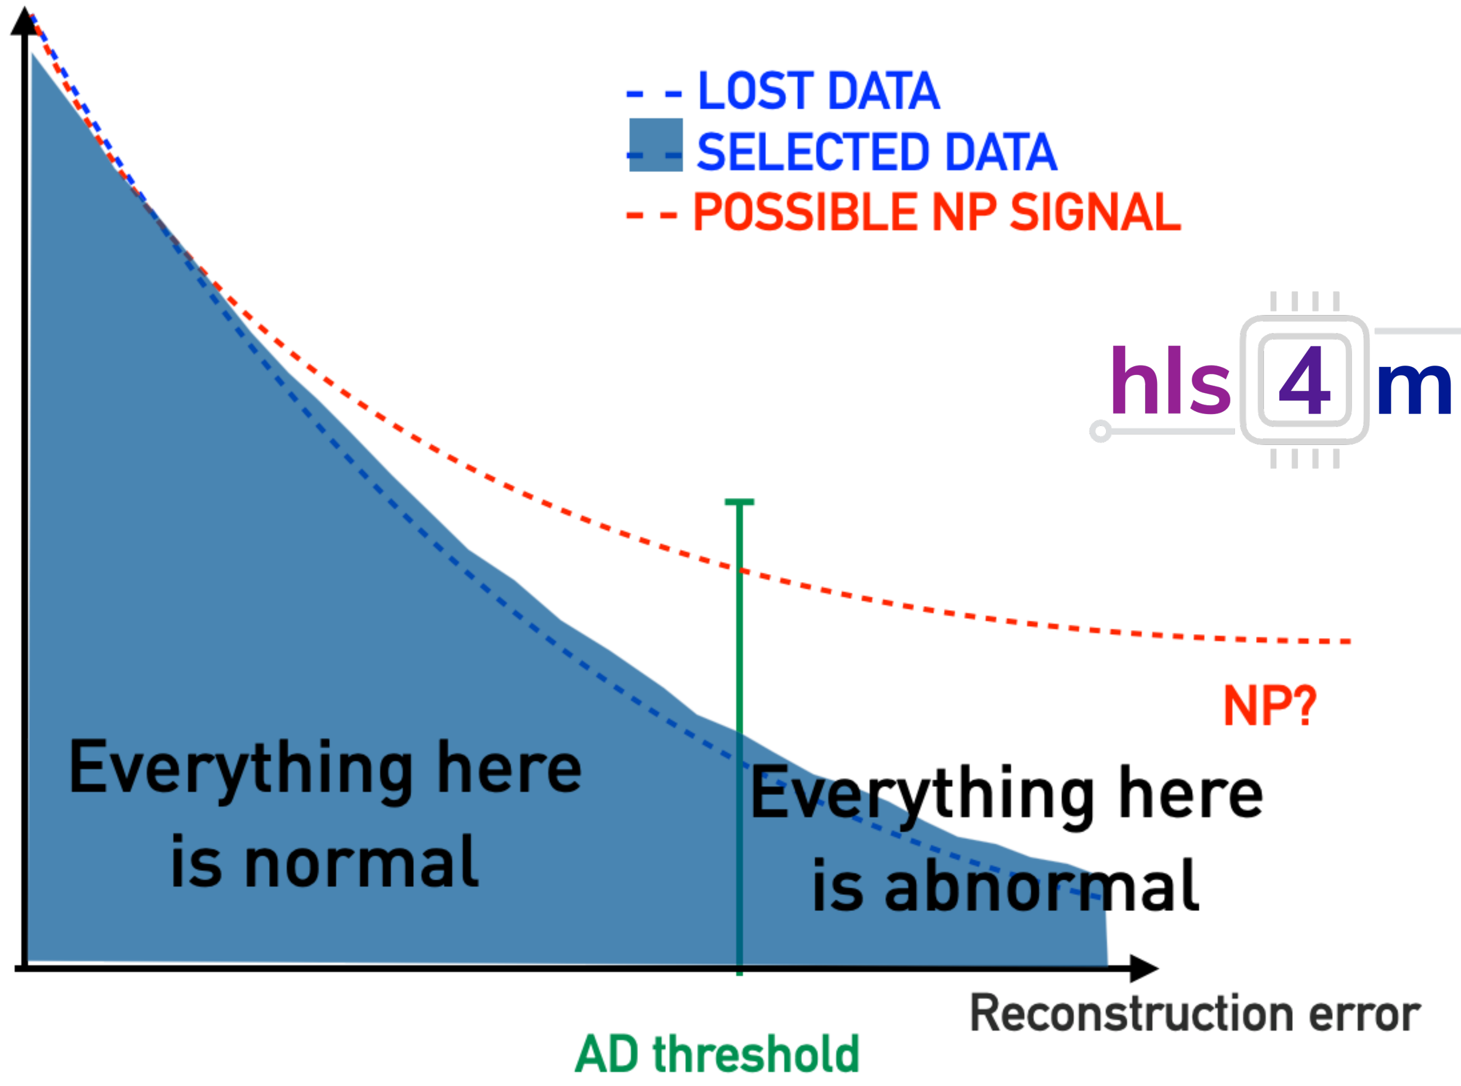

CMS DP2023_079

ATLAS Experiment:
- BDT selecting candidate τ lepton events in <100 ns

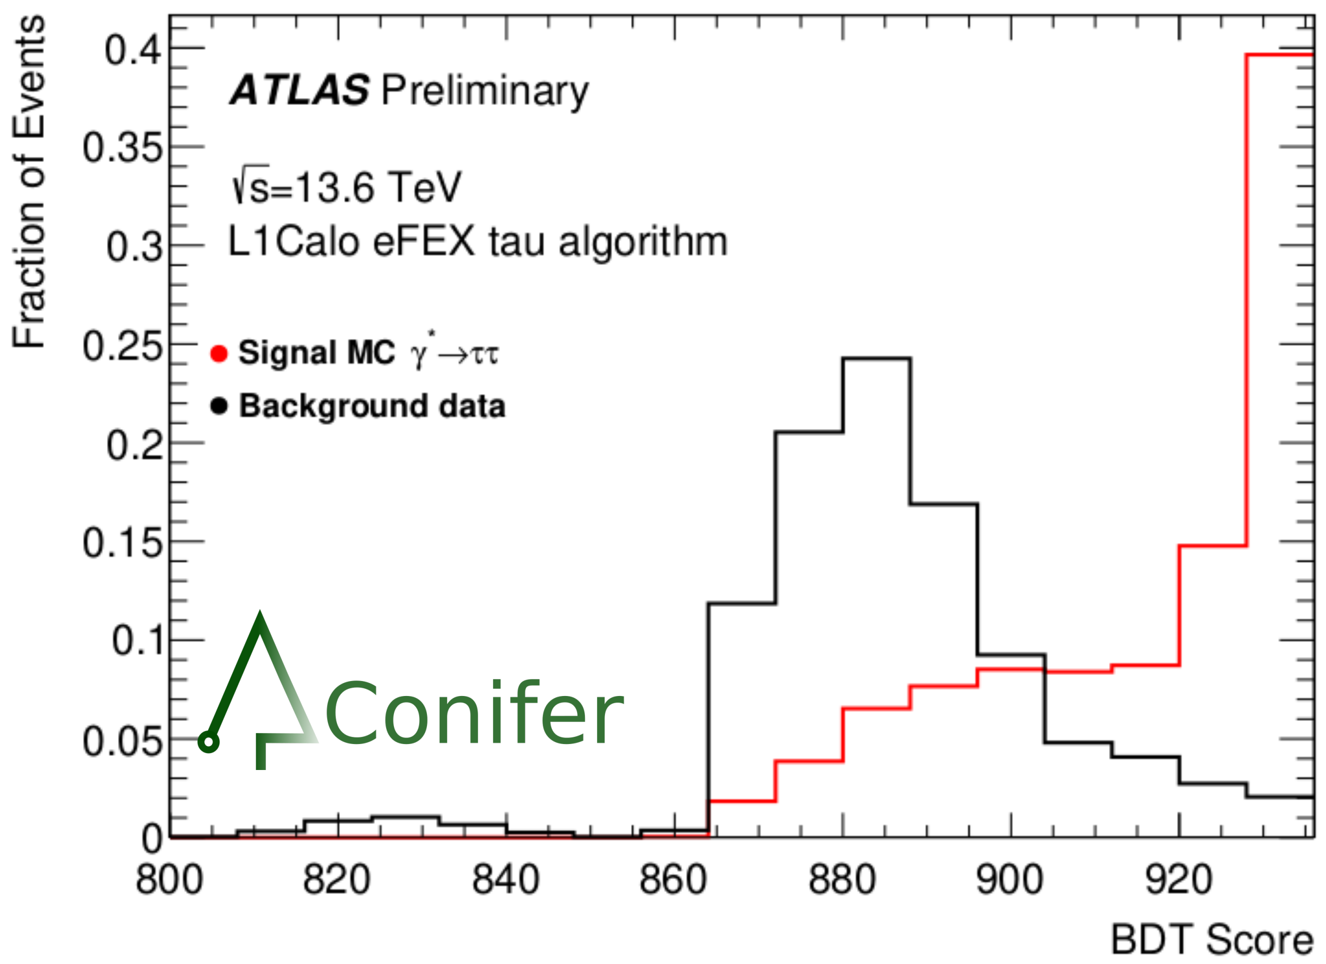

L1CaloTriggerPublicResults

2024: Neural hardware triggers making decisions in LHC experiments!

CMS Experiment:

- Anomaly detection in 50 ns
- 300 events/second

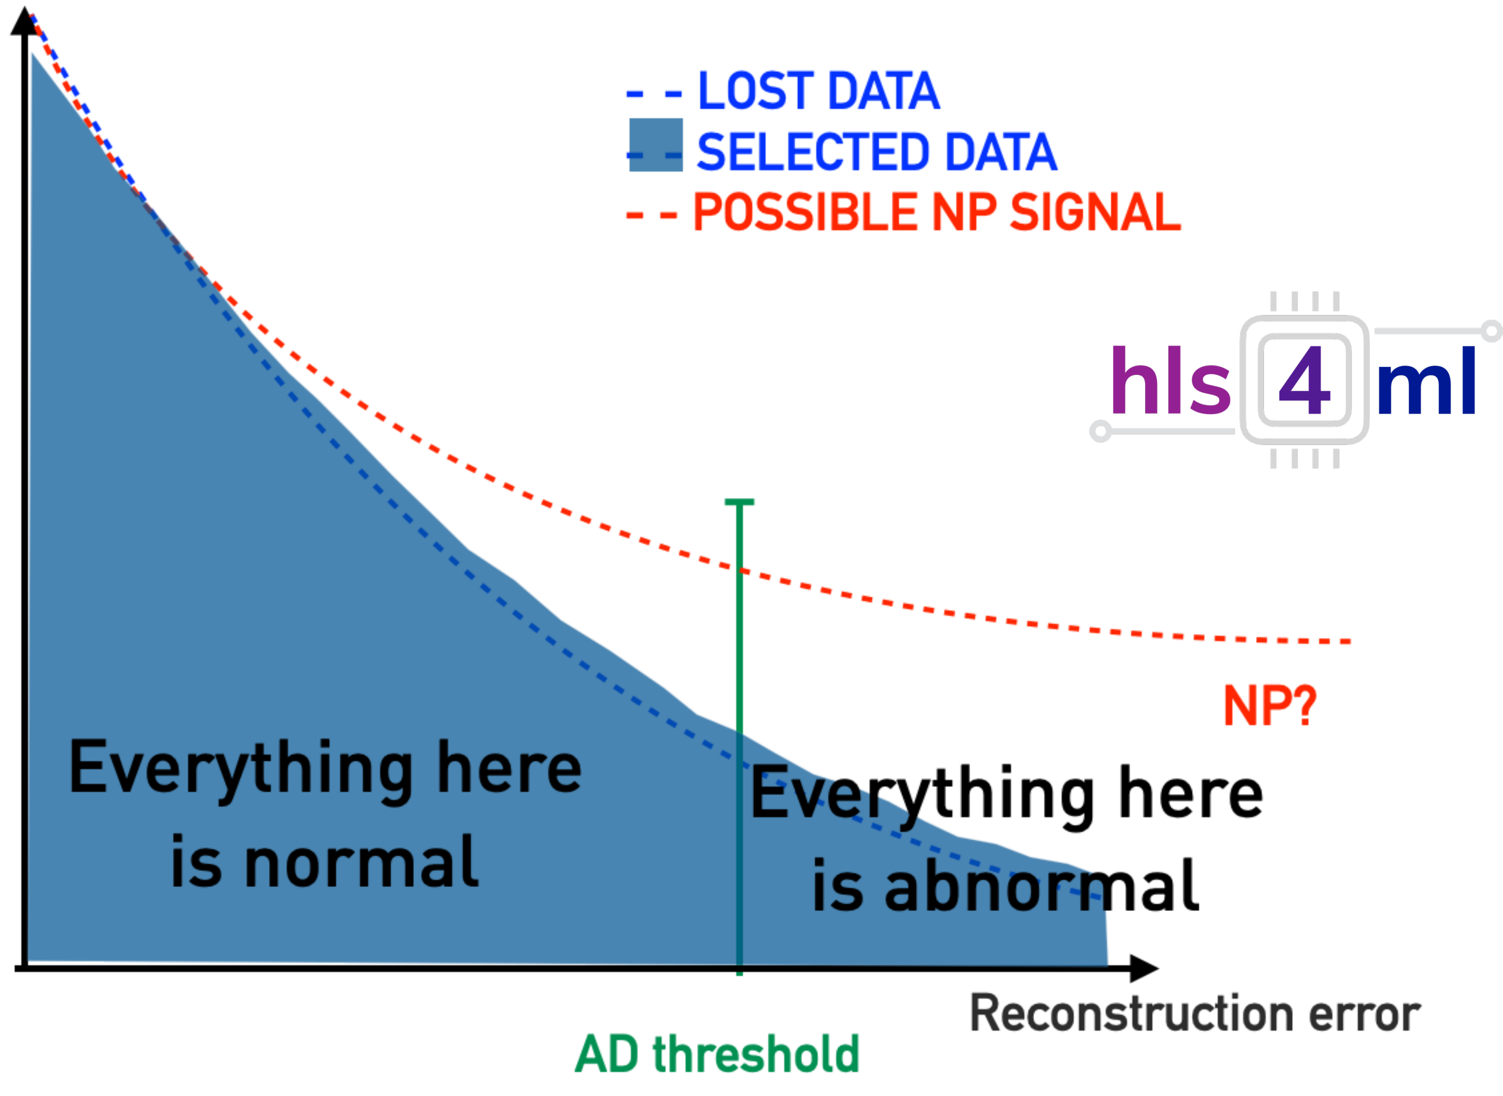

CMS DP2023_079

ATLAS Experiment:

- BDT selecting candidate τ lepton events in <100 ns

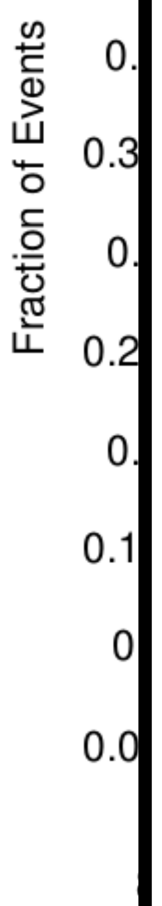

2031: ~20 billion inferences/s during HL-LHC

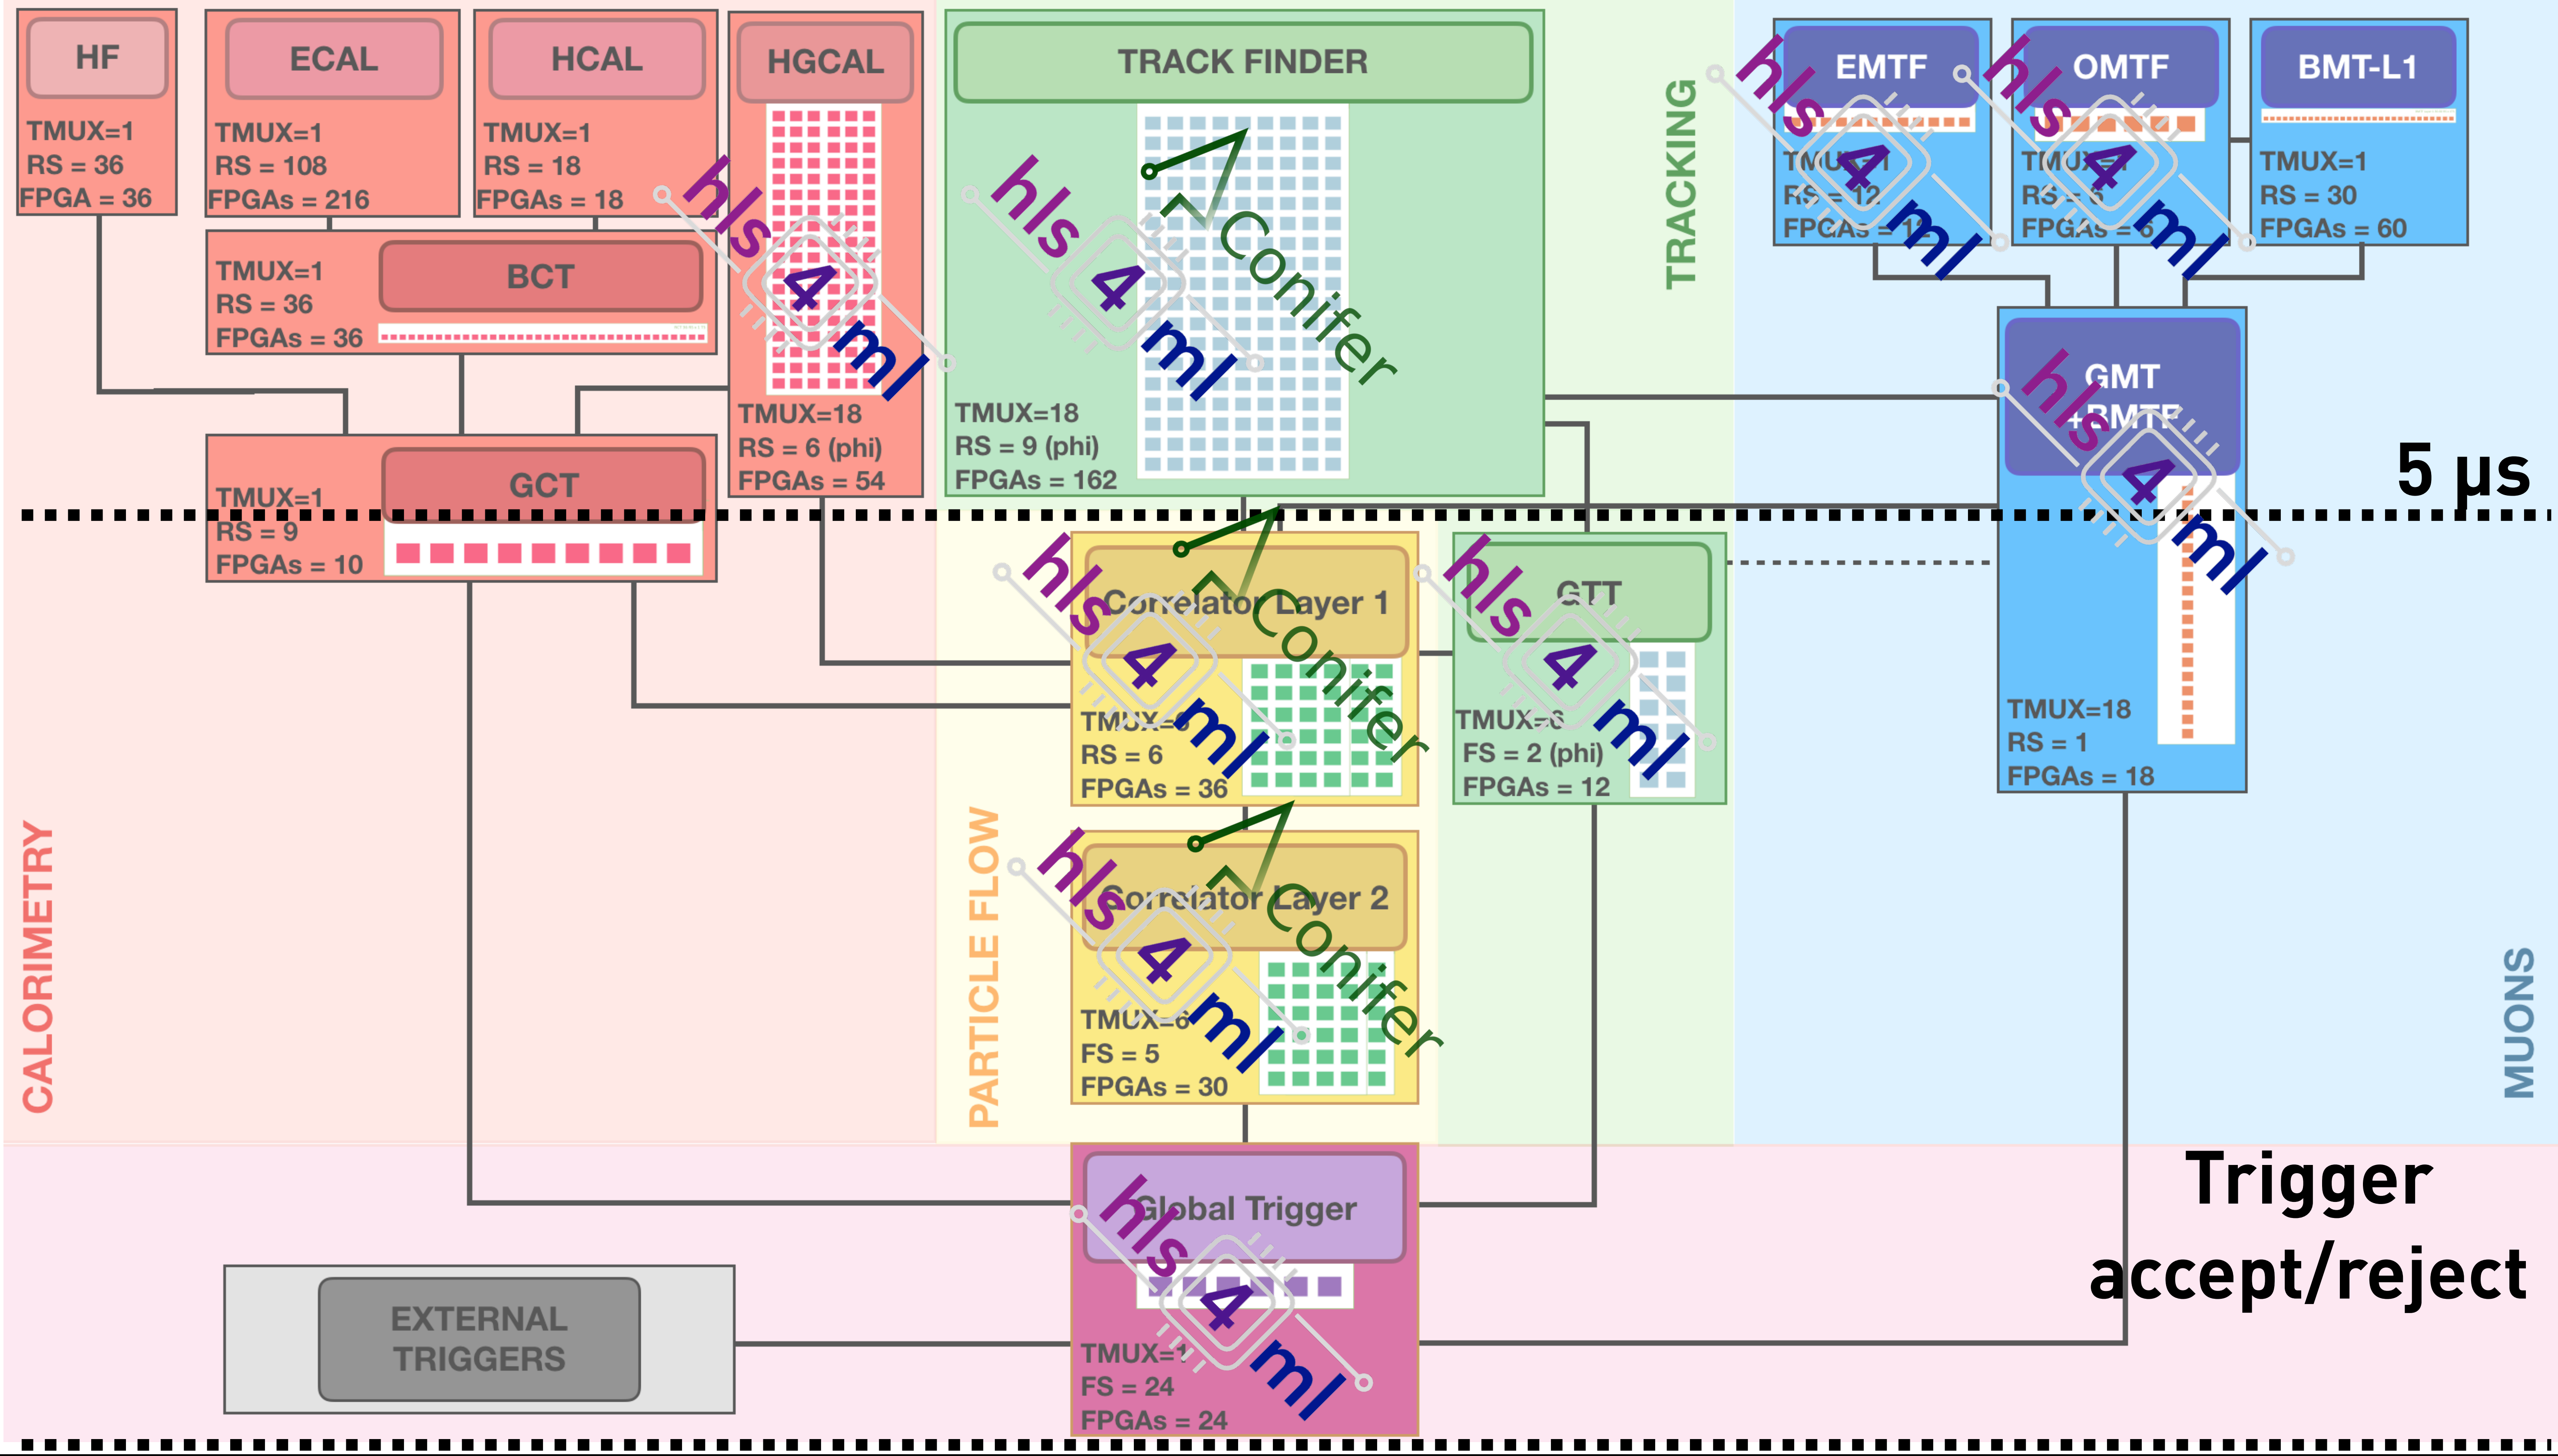

Frontend reduction

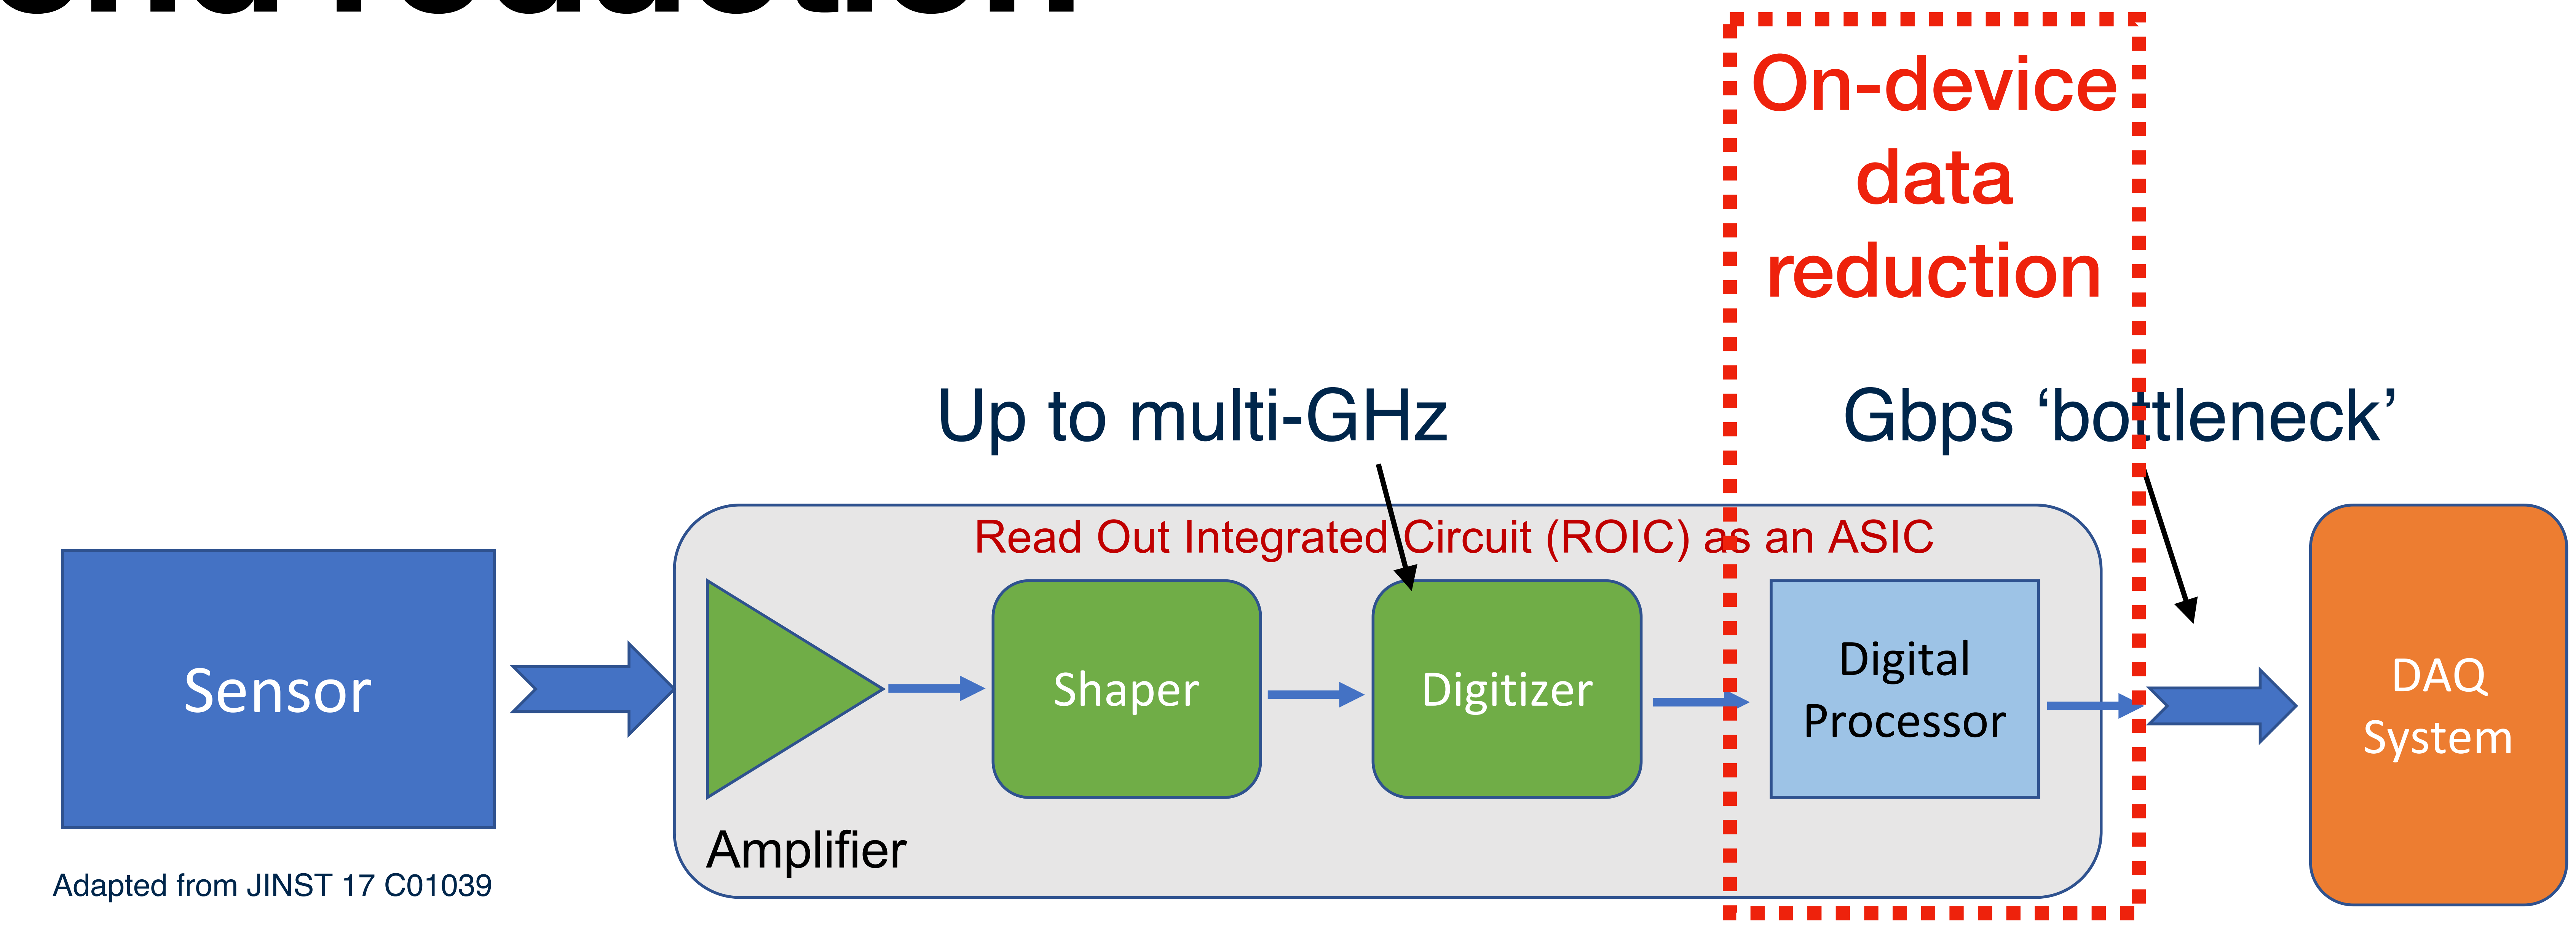

- highly granular detectors, background → high data transfer rate
- limited by **available bandwidth** of electrical and optical links
- 2 solutions: **frontend reduction** or **higher capacity** transfer

QKeras + hls4ml

Smart Pixels

https://fastmachinelearning.org/smart-pixels/

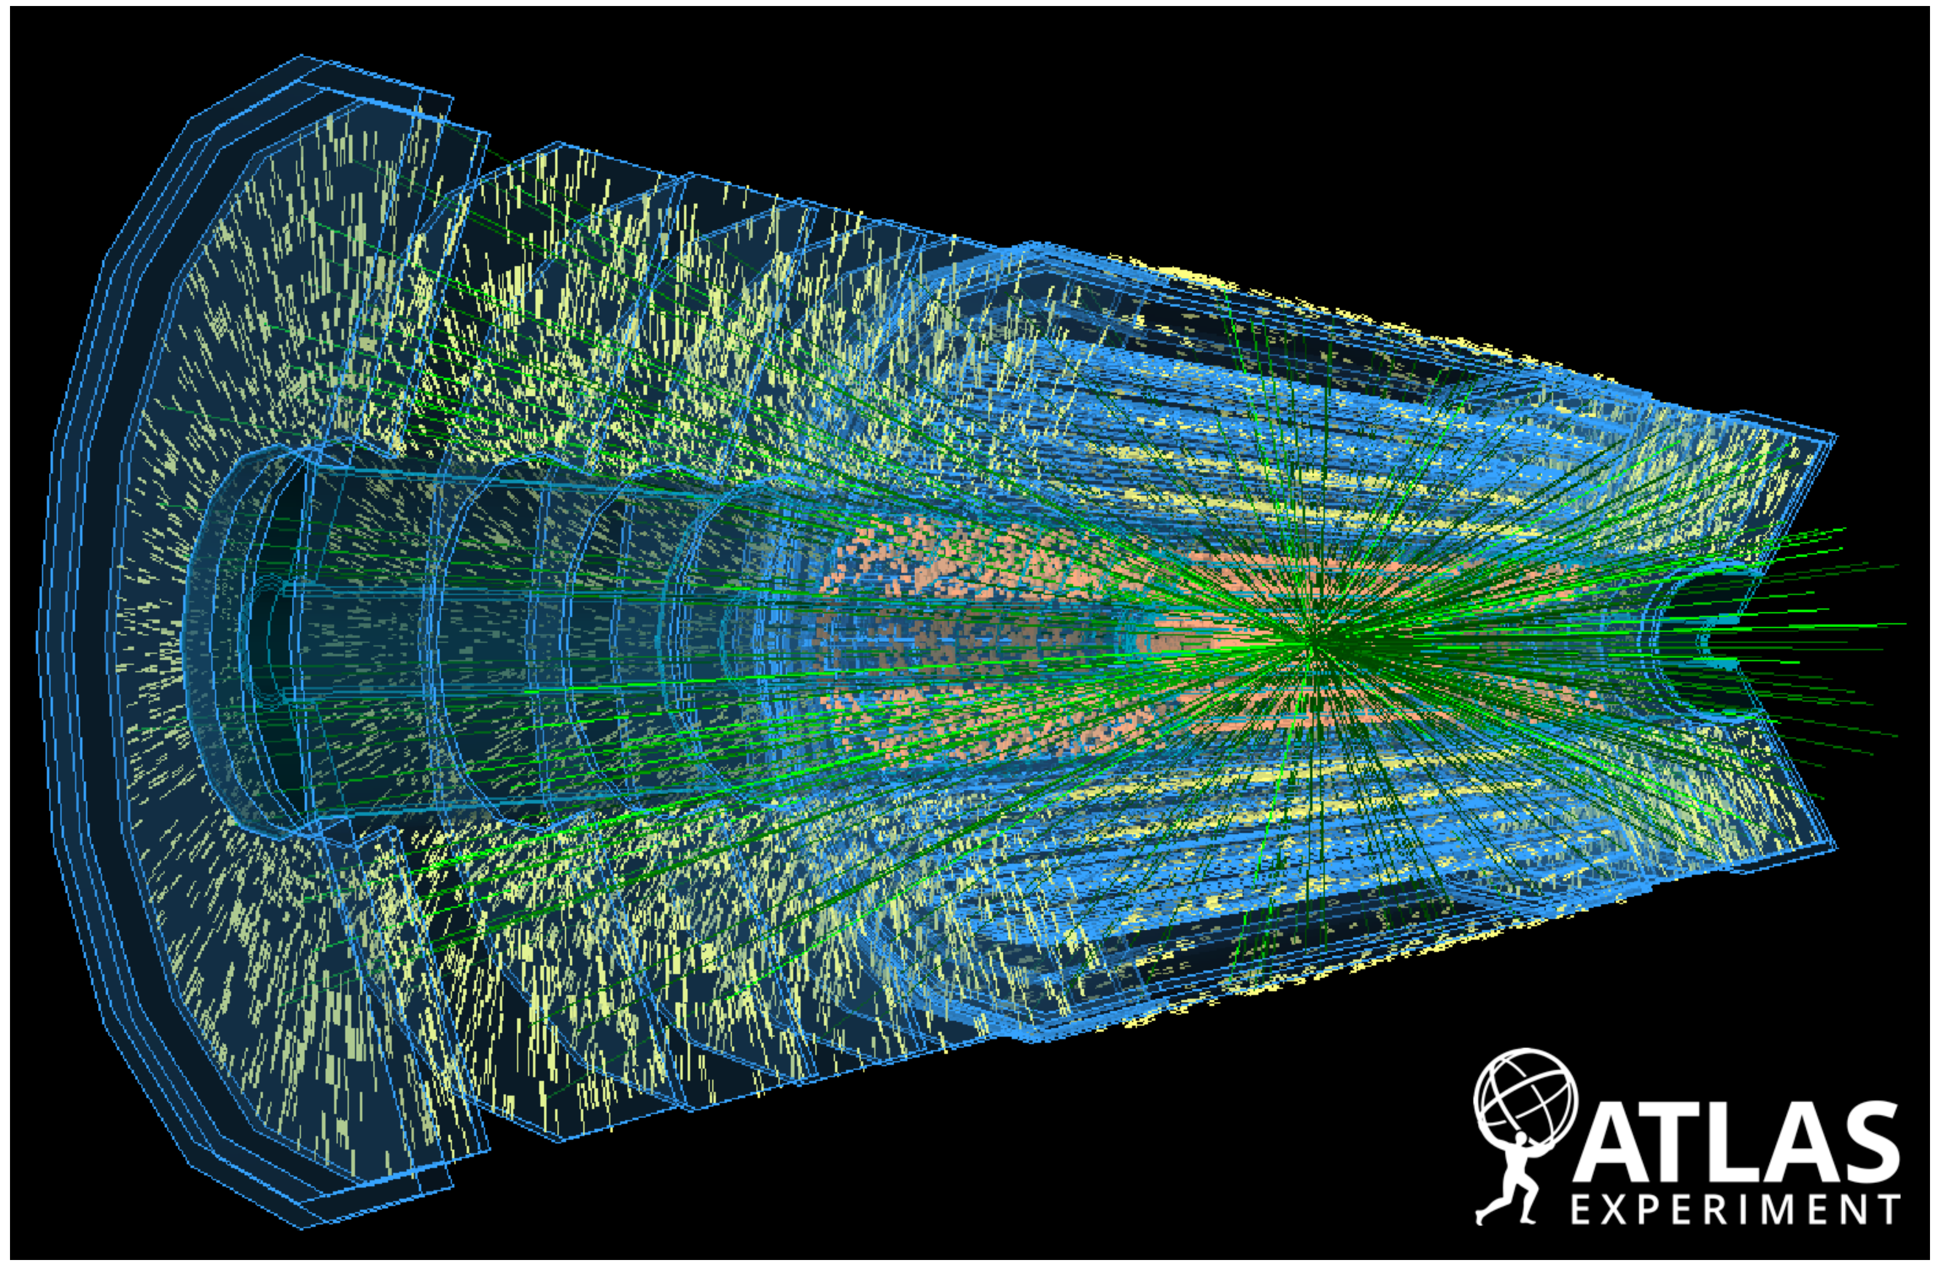

- Reduce silicon data via **in-pixel intelligence**
  - **frontend filtering**: discard low-p_T track data (< 2 GeV)
  - **feature extraction**: Extract particle position and angle in pixel front-end ASICs from charge in single pixel layer
- Bandwidth savings of 57-75%!

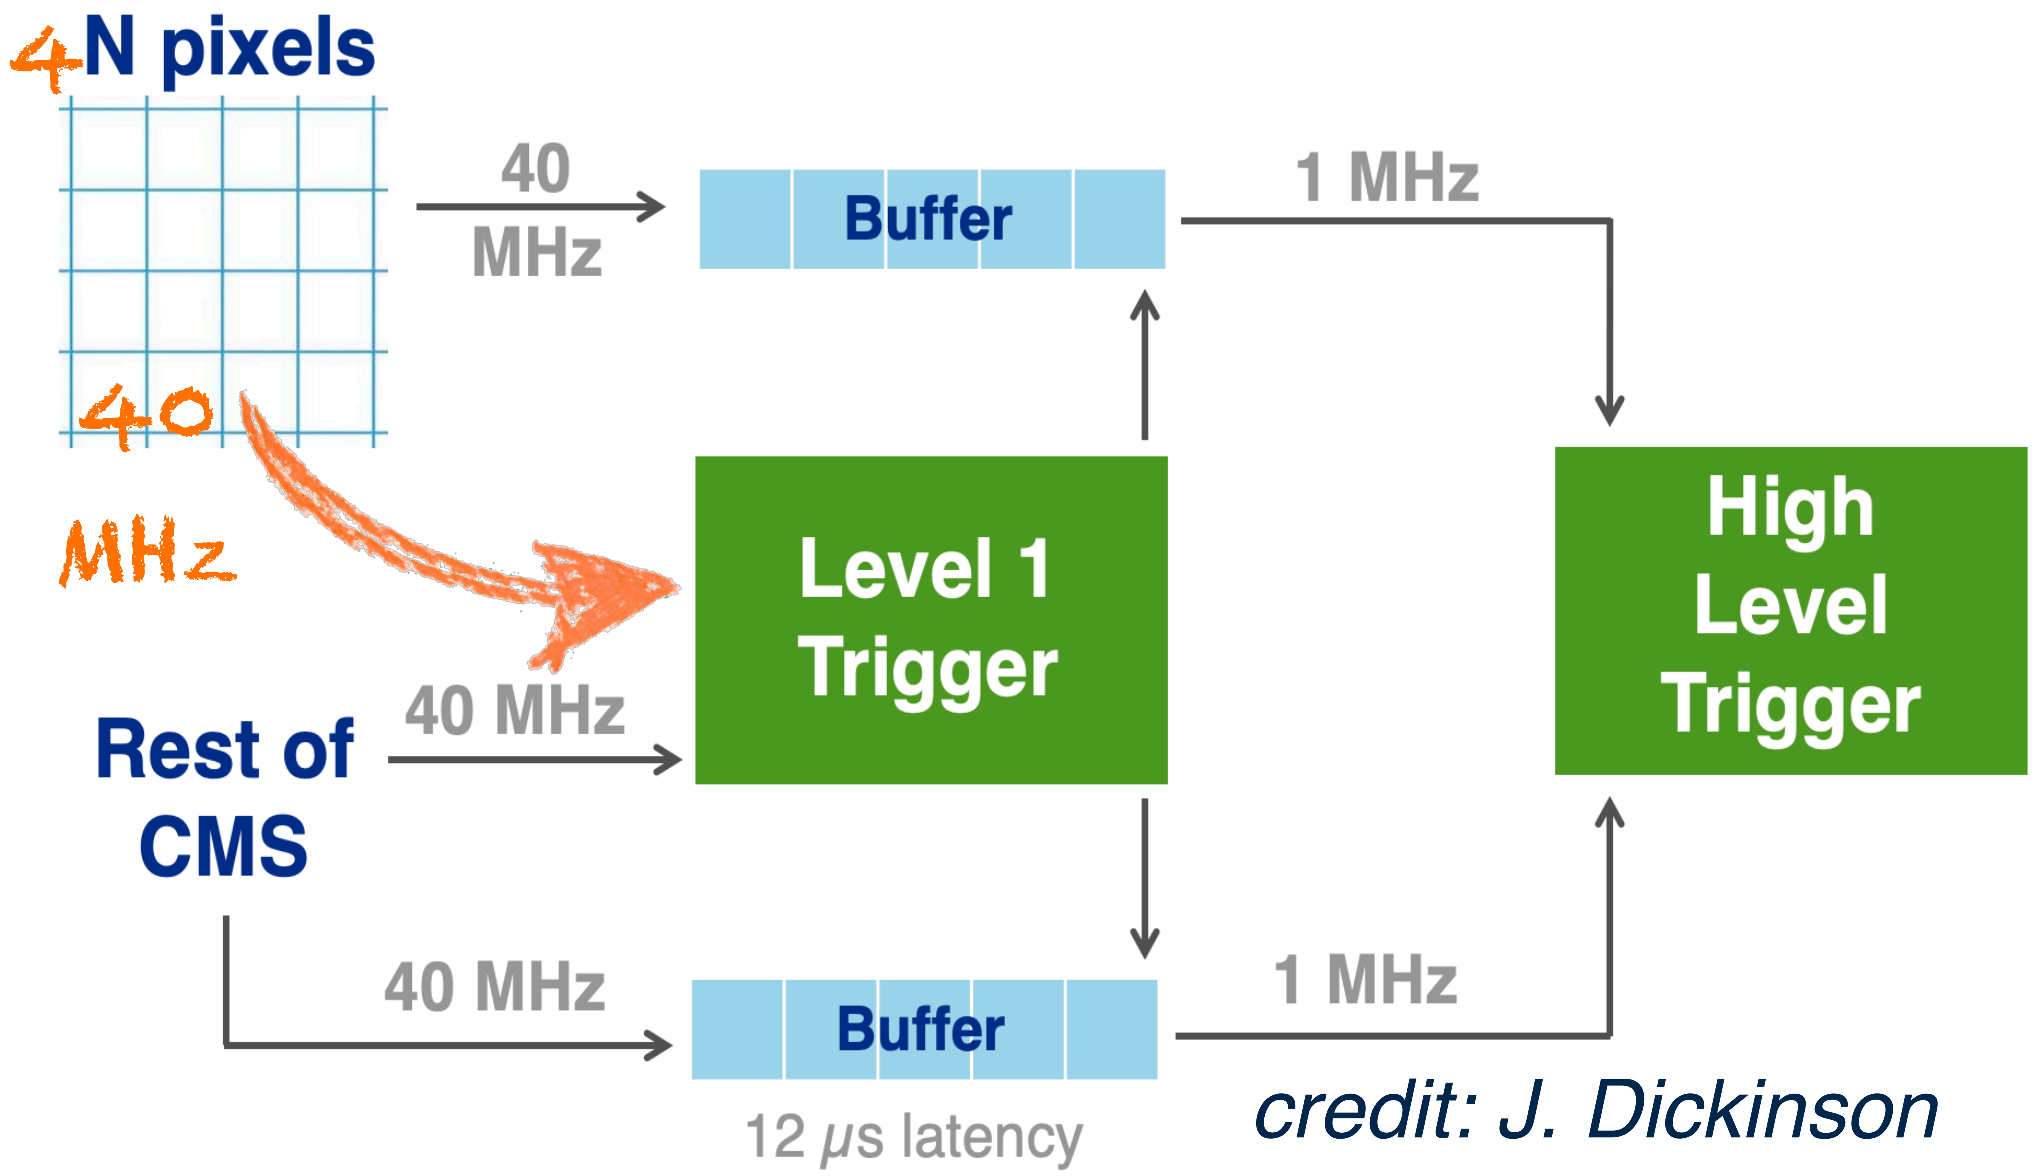

credit: J. Dickinson

Reconfigurable logic in ASIC design

The Embedded FPGA framework

- Pathway to implementing ML "at source”
  - Fully reconfigurable logic on detector frontend
- Open source (FABulous, OpenFPGA)
  - potential to apply to variety of subsystems/ fields (SuperKEKB, FCC-ee, DUNE, free-electron lasers)

28nm CMOS ASIC (1x1mm)

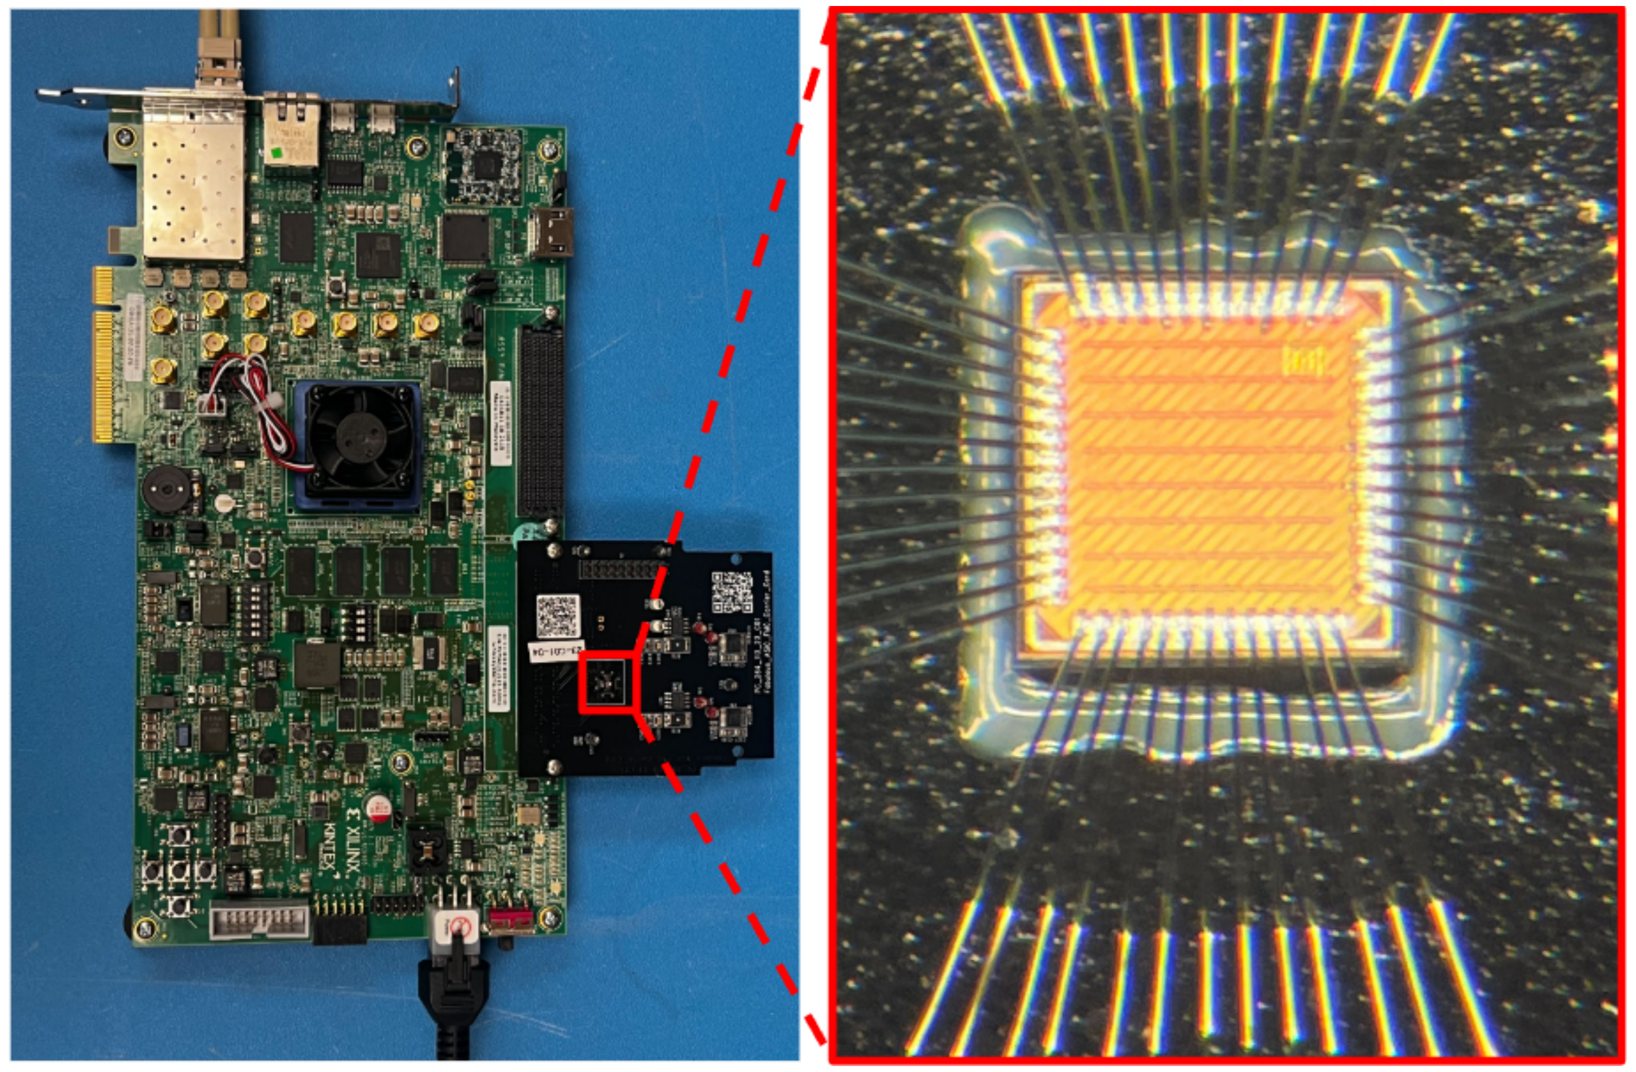

Gonski et. Al

ML compression: Calorimeter data concentrator

10,000 ECONs with ML inside going live in HL-LHC

Di Guglielmo et Al, IEEE TNS 68.8 (2021)

Data transfer bottlenecks

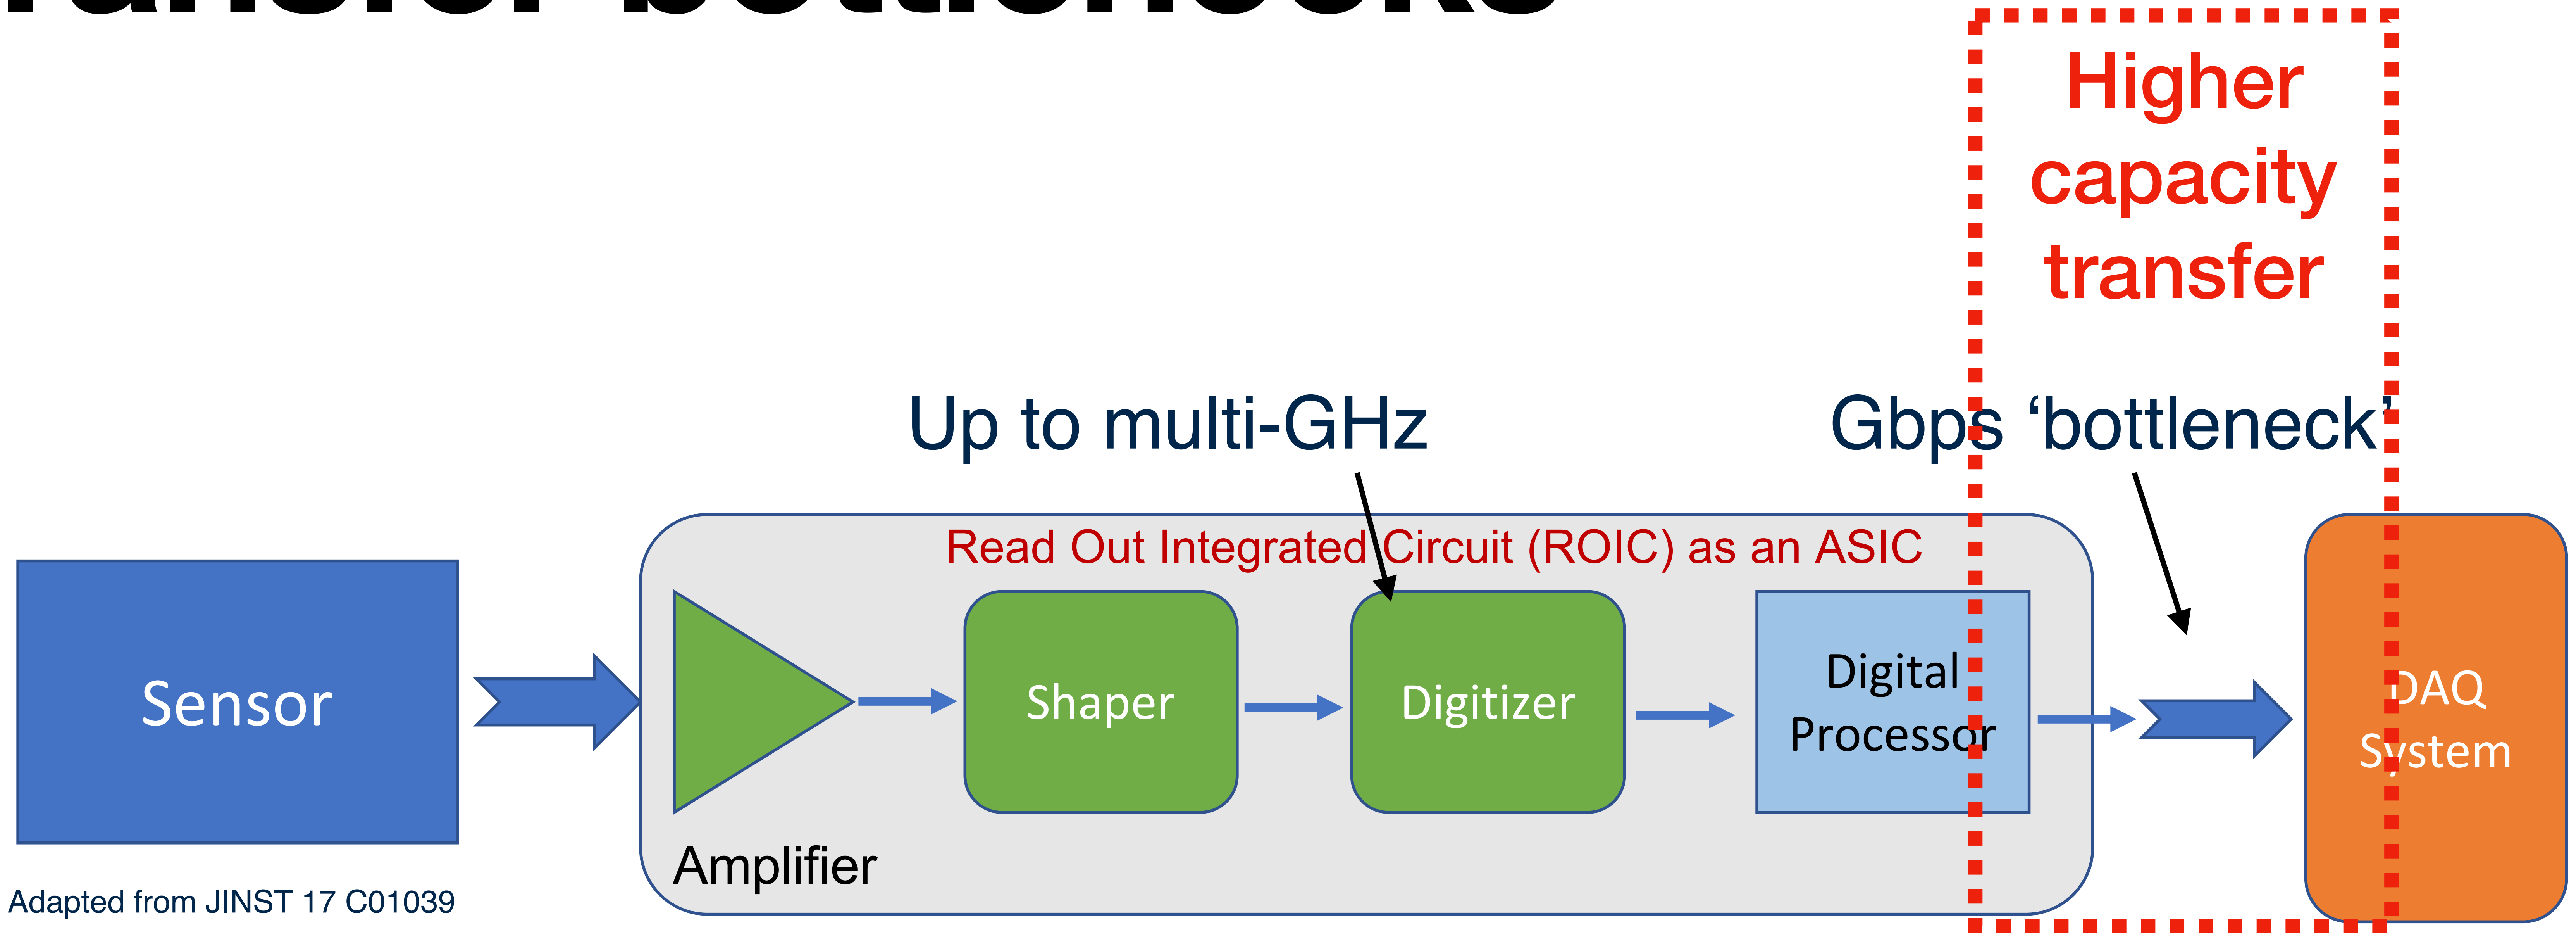

Adapted from JINST 17 C01039

- 2 solutions: **frontend reduction** or **higher capacity** transfer

ESPPU26 #95 #247 #233 #93 #211 #272

Wireless data transmission

Reducing material budget for detector readout

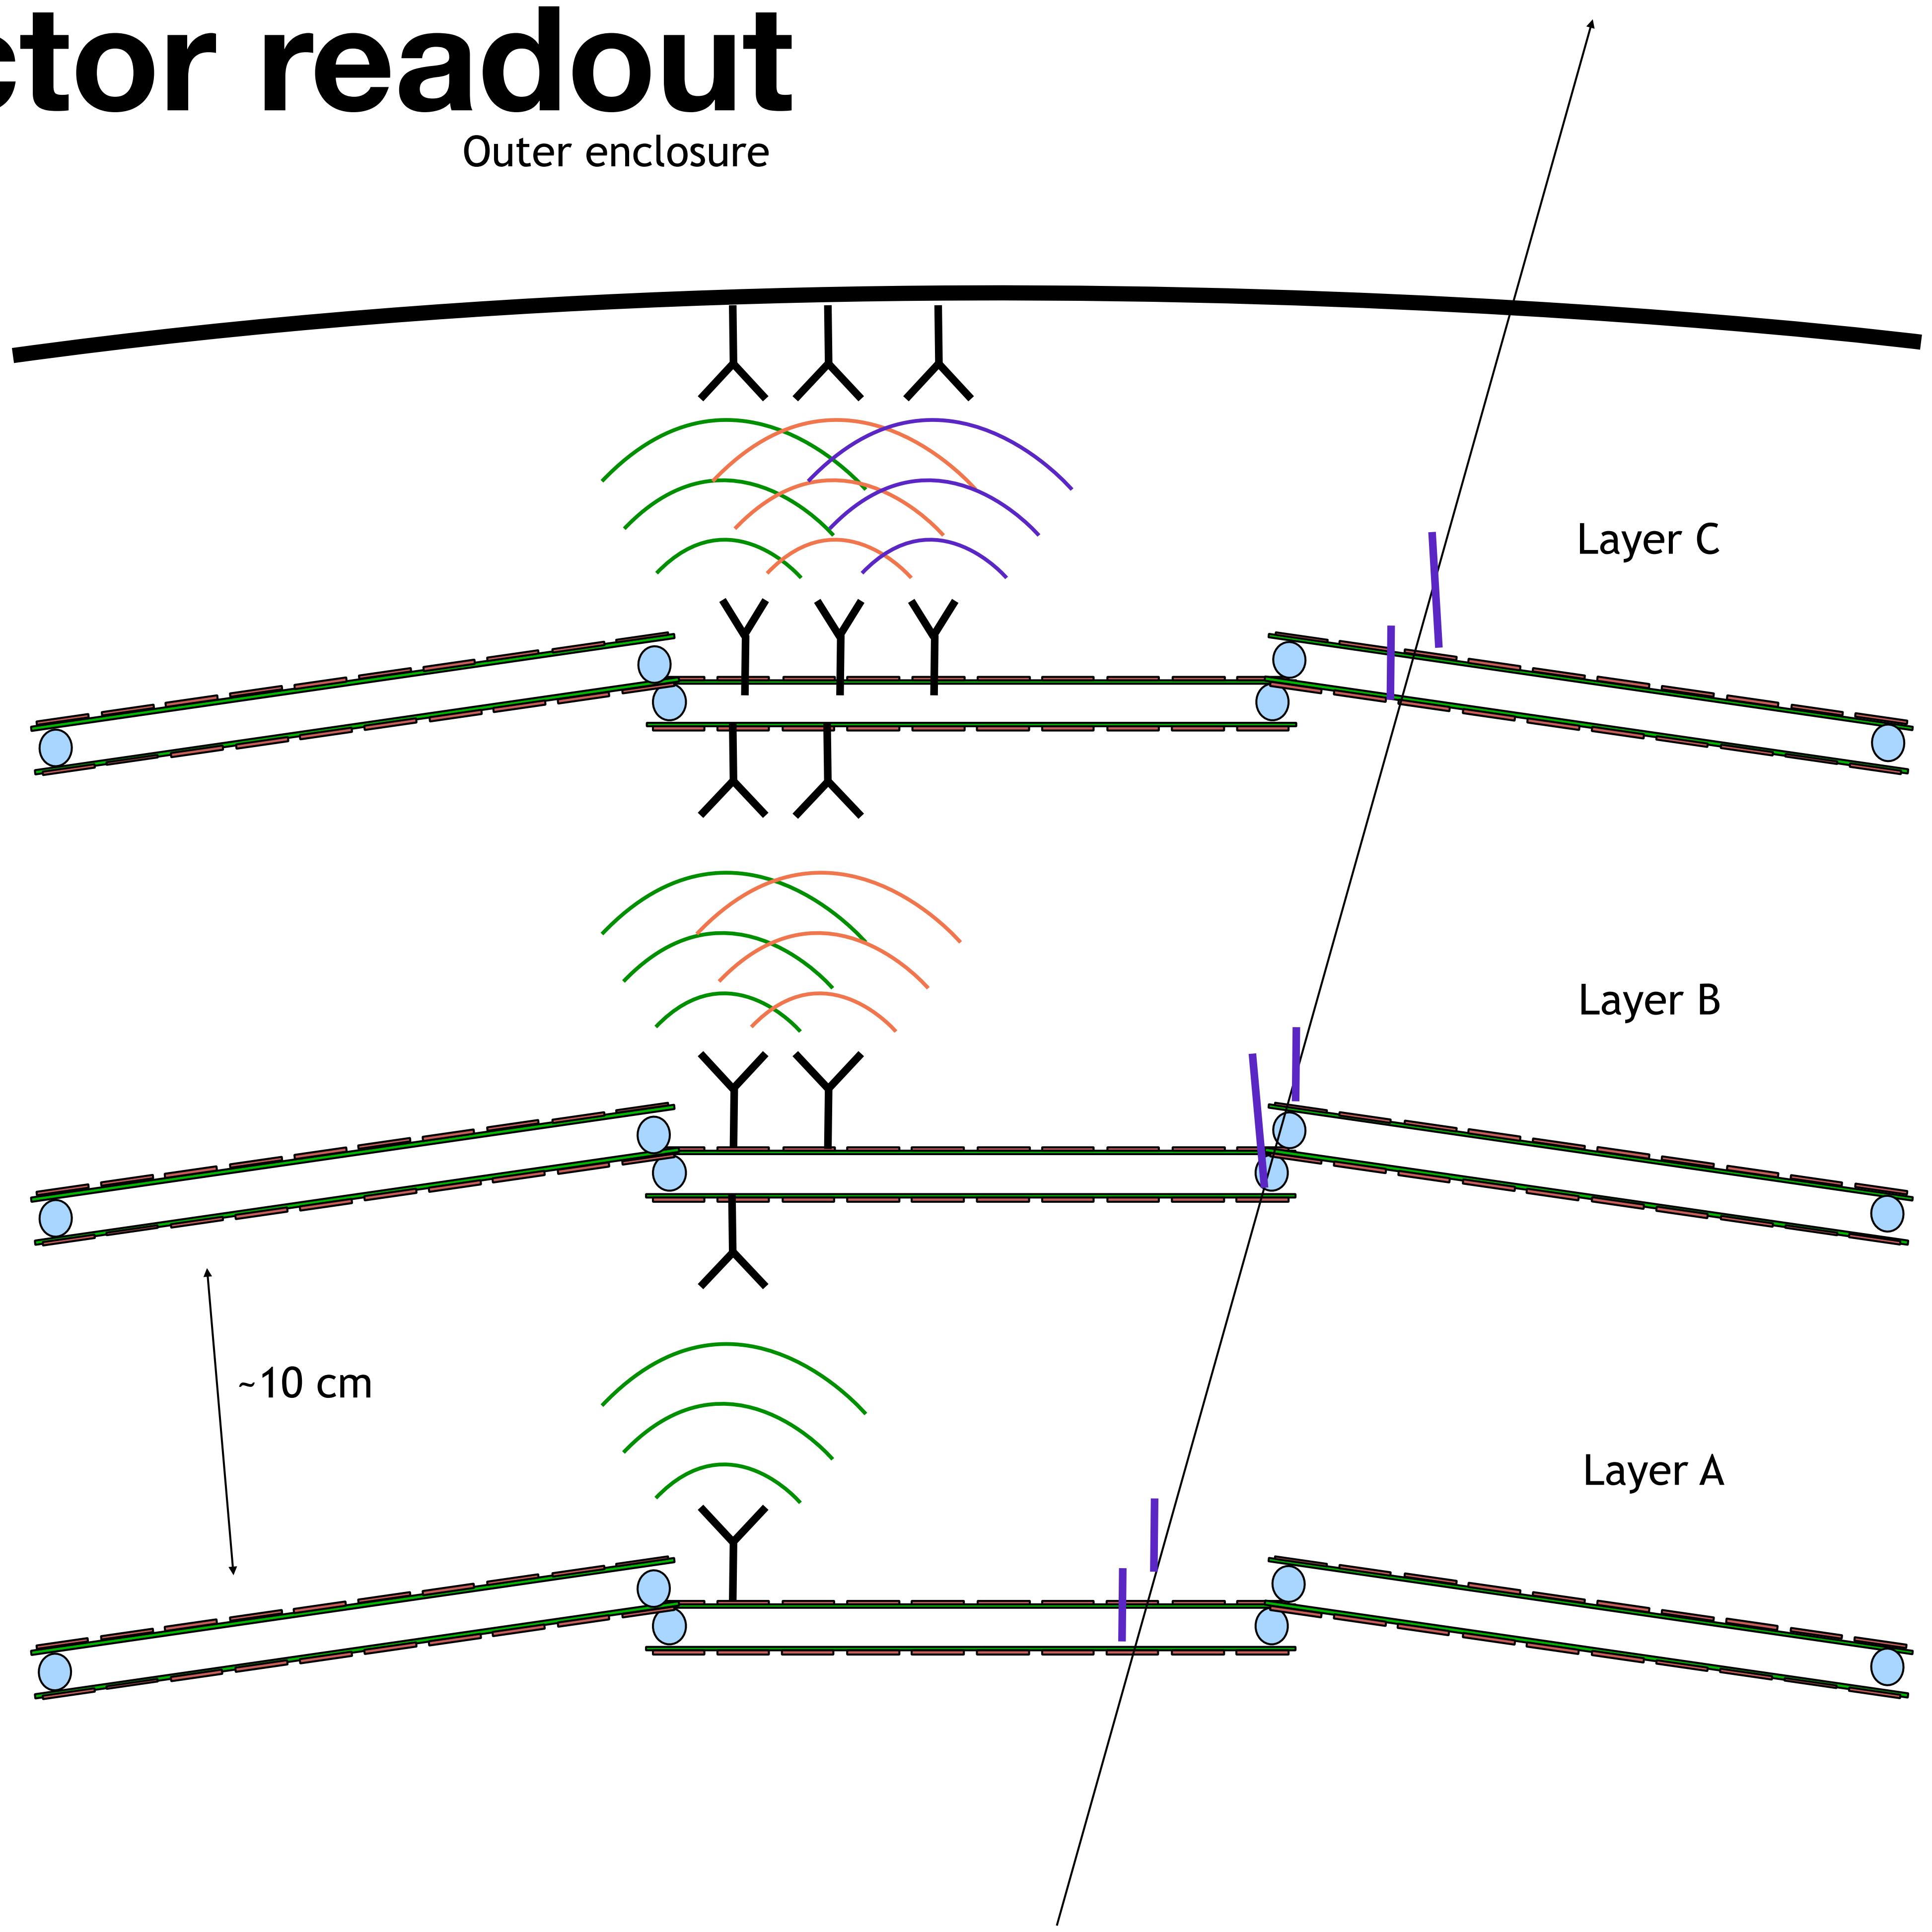

- Increase Gbps w/o increasing material
- Send single signal to several receivers, saves cabling
- Cost reduction, simplified installation/ repair, reduction in dead material
  - Especially important for future tracking detectors

Few Gbps possible with 802.11ac/ad WiFi!

Wireless Allowing Data And Power Transmission (DRD-7.1c)

Intelligent back-ends

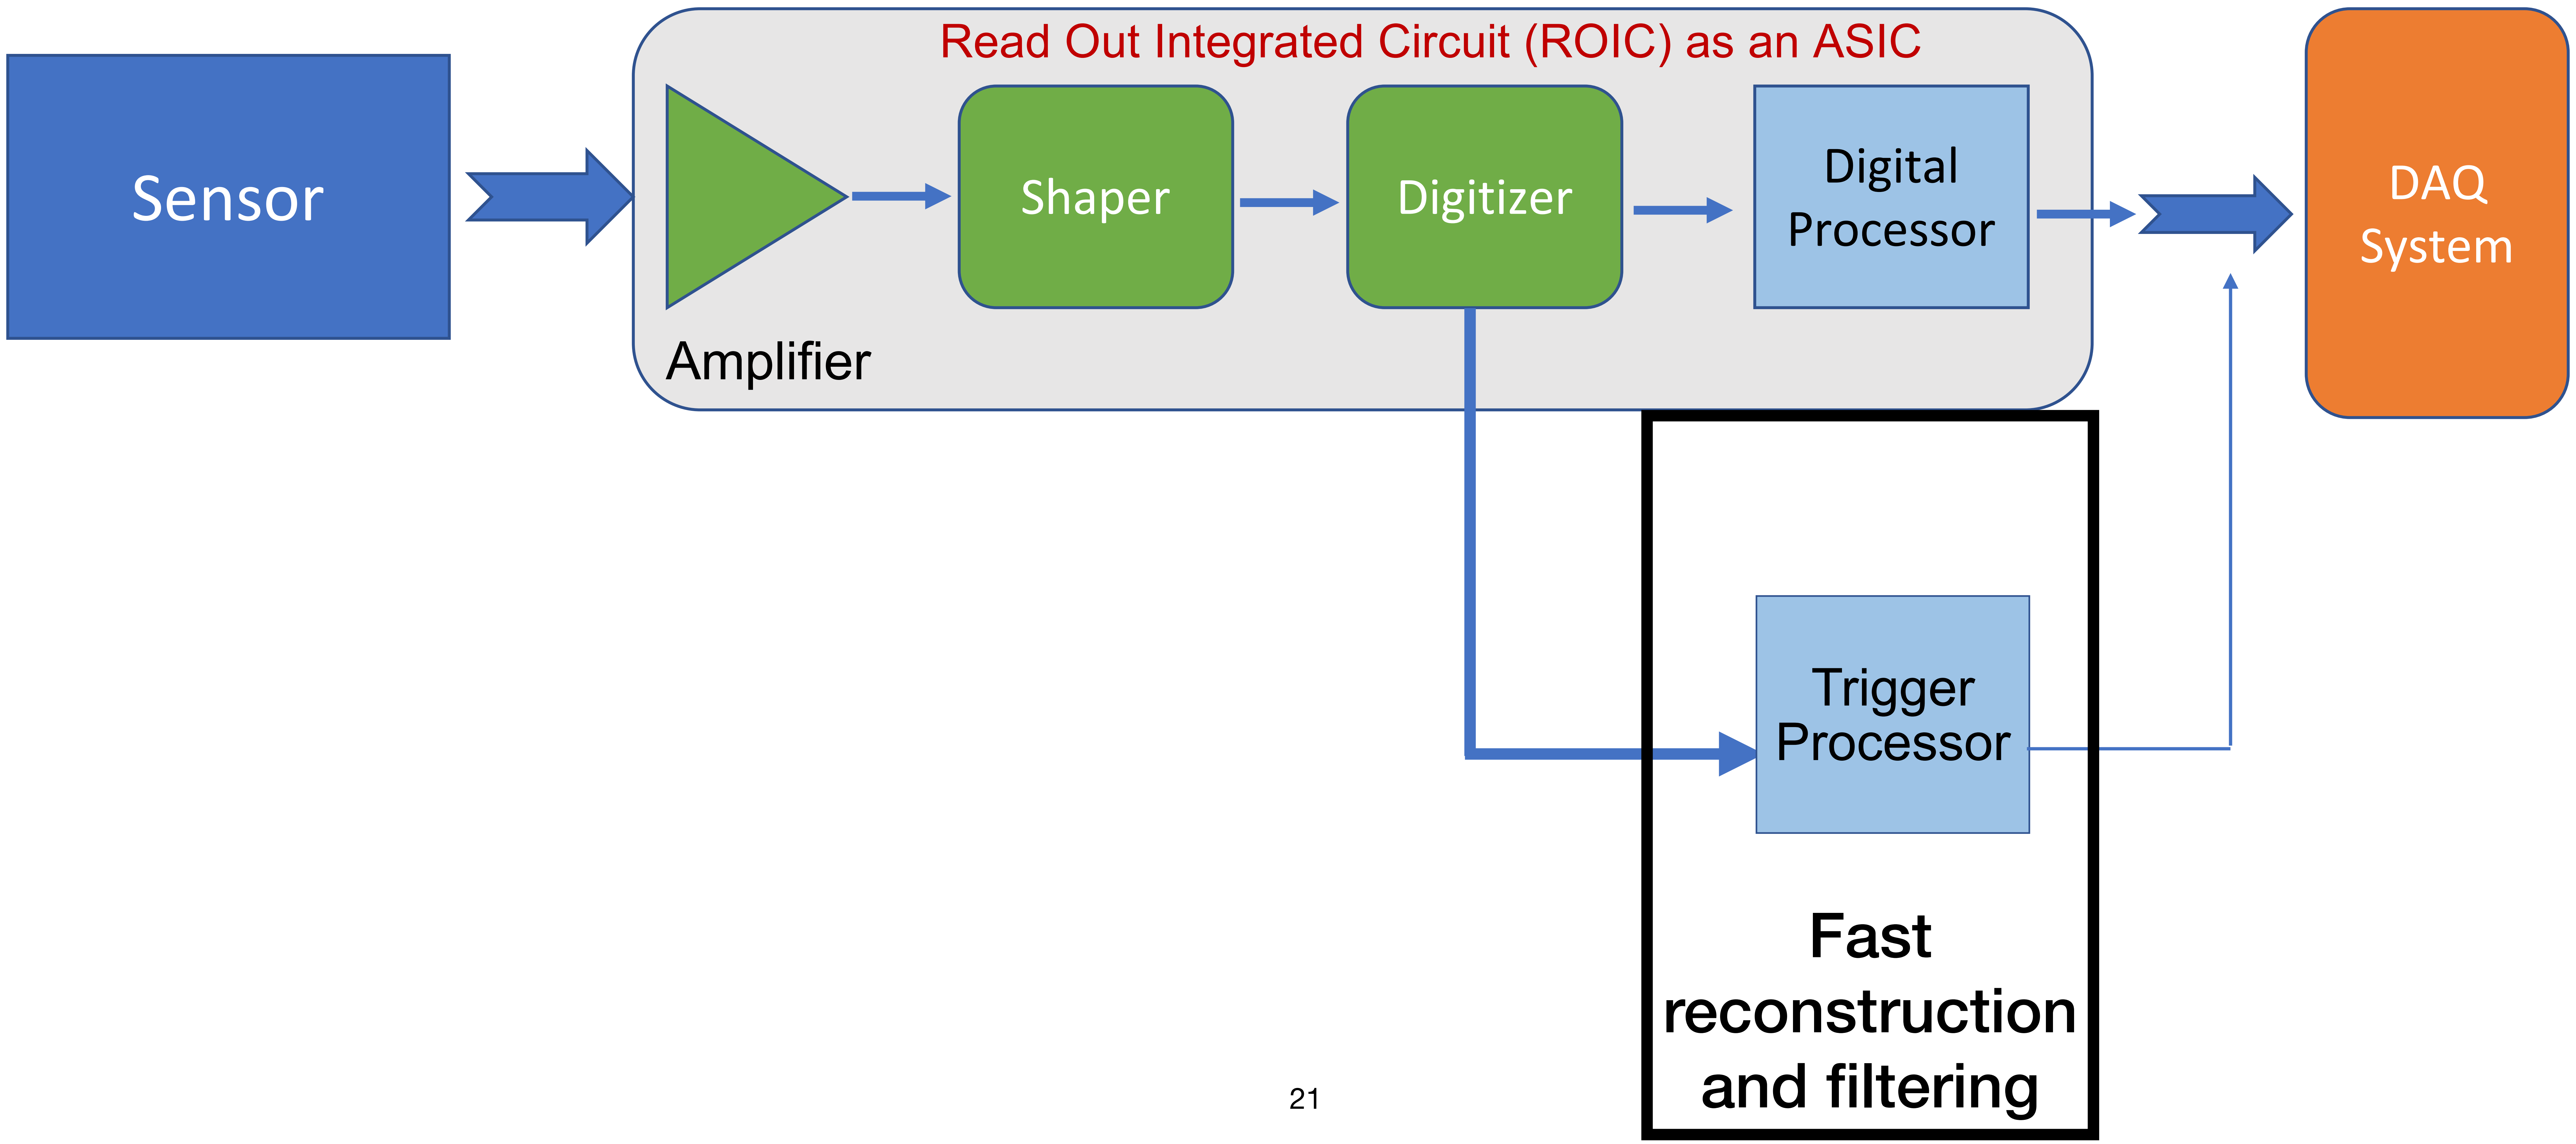

Real-time tracking

- HL-LHC CMS will read out tracks with p_T>2 GeV and run L1 track reconstruction
  - reduce rate by x10 by filtering low momentum hits
  - Offline like reconstruction: track seeding, building and Kalman Filter to identify final track candidates and determine track parameters
  - Track quality with BDT implemented with Conifer

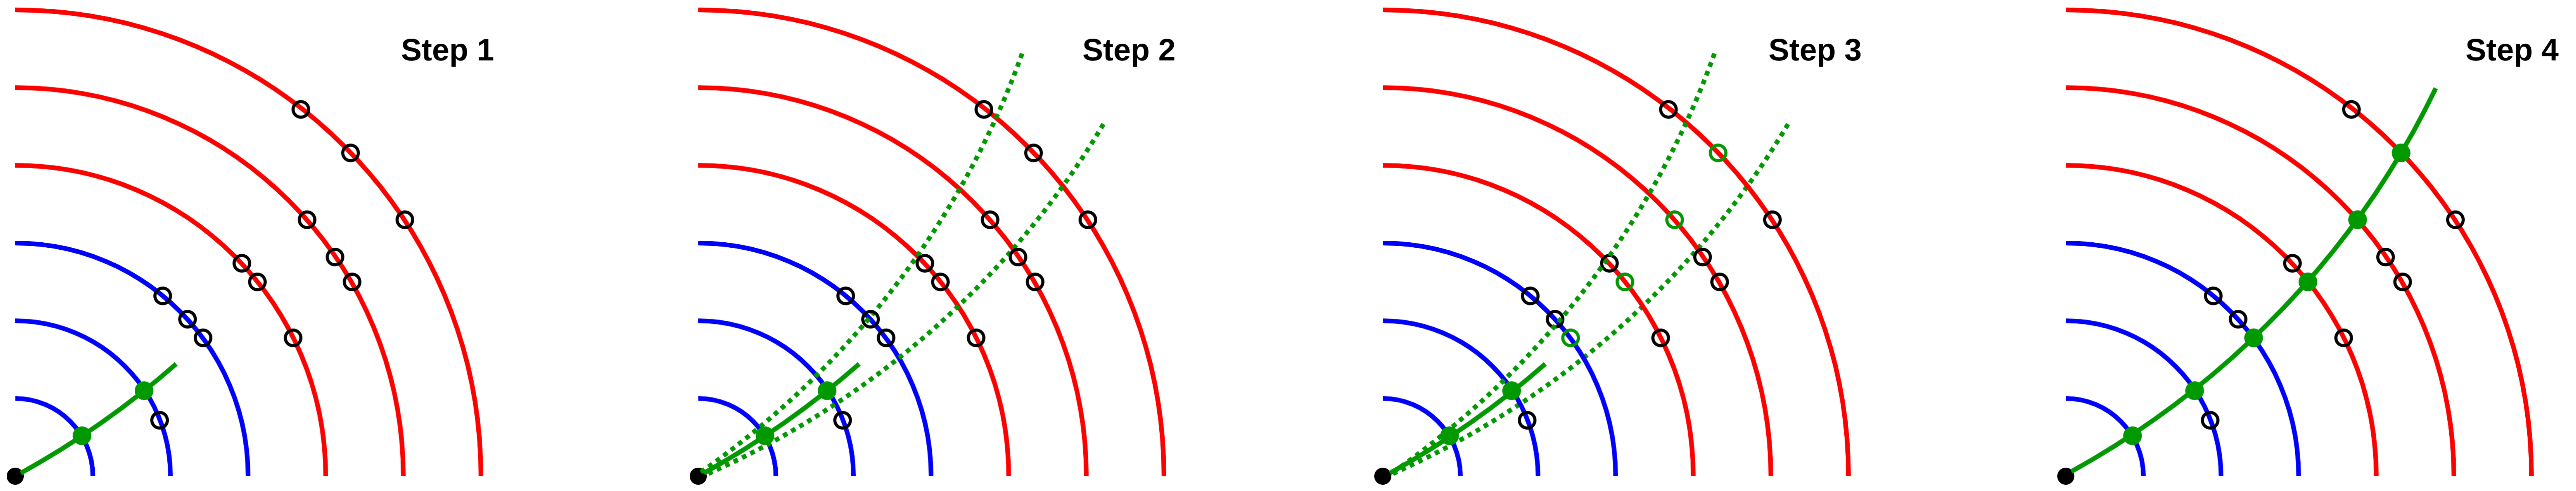

  - ….. **in 3.4 μs!**

Hardware jet building and flavour tagging

HL-LHC CMS will run with

FPGA track finding, particle flow, jet clustering and flavour tagging in < 12 μs

arxiv:411.13596 #205

Real-time track triggers
BELLE-II

- Belle-2 uses neural track hardware trigger for reconstruction of vertex, and azimuthal and polar angles of single particle track
- ML for GNN-based offline tracking
  - Work ongoing on bringing this to the hardware trigger

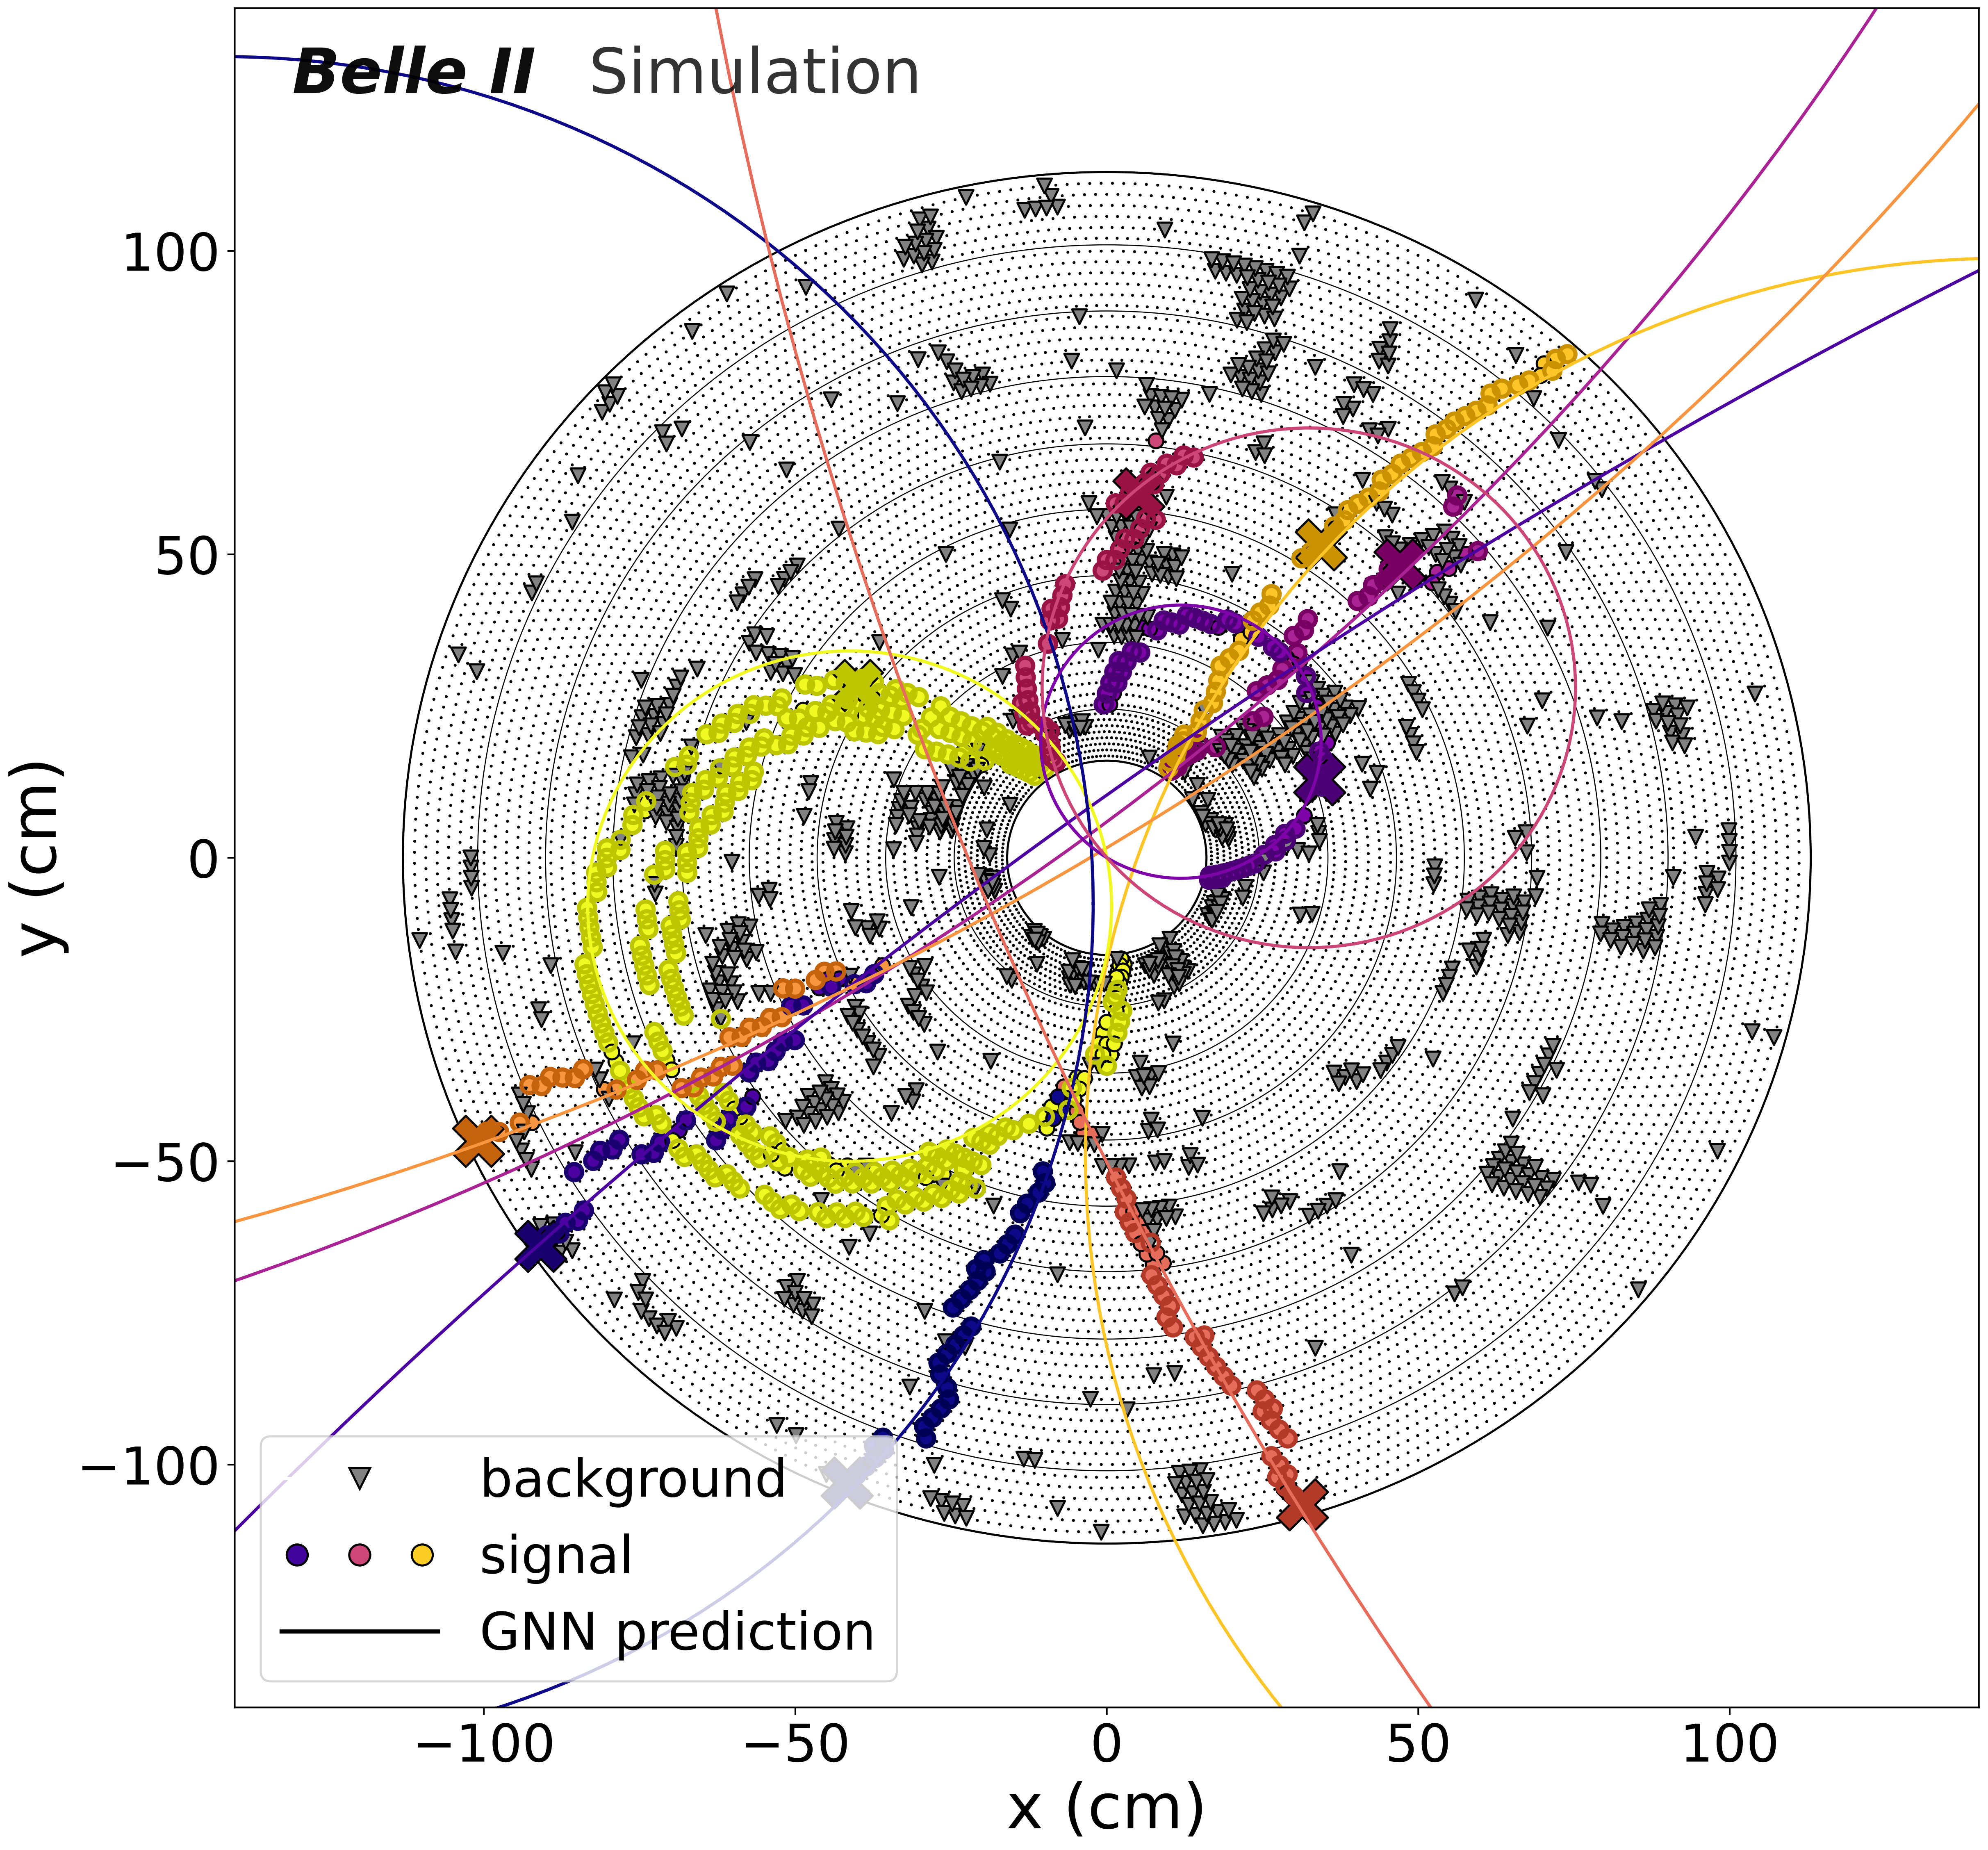

Trigger-less design: LHCb Run-3

- Run 3 event selection entirely in software, heterogeneous platform
  - FPGA-based DAQ cards (low-level bit manipulations)
  - high-speed dedicated network cards (memory ops for transfer)
  - GPUs for data processing (large-scale parallel problems)
- Lessons learned in Run3:
  - Expertise in Low-level network simulations
  - Architecture that supports fast adaptation to emerging tech
  - **Full read-out, track reconstruction and lepton PID at 30 MHz**
- Run-4 R&D:
  - Early track reconstruction on FPGA boards (RETINA)
  - Investigating emerging processors as GPU alternative (ALLEN)

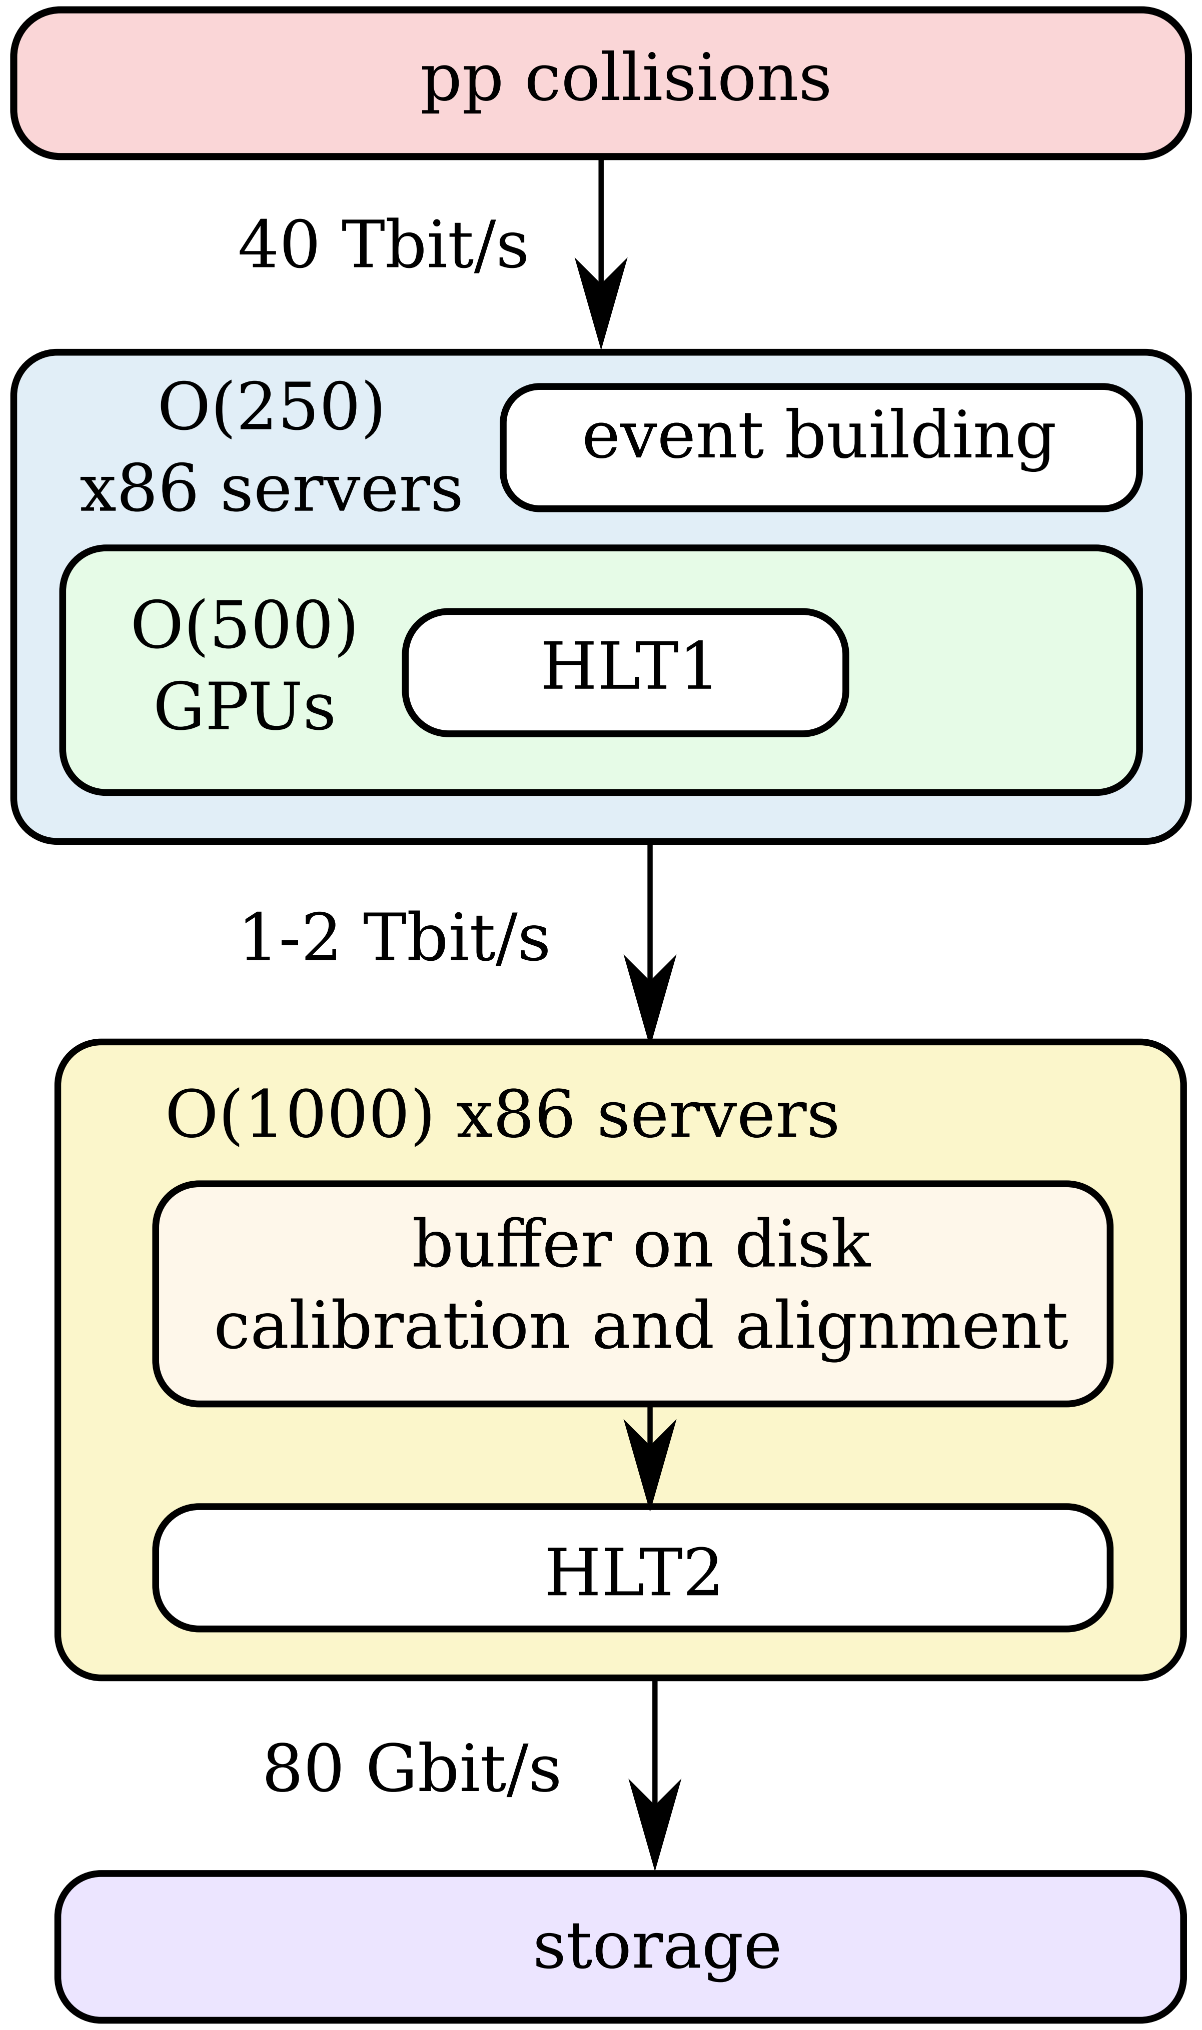

Aaij, R., Albrecht, J., Belous, M. et al. Allen: A High-Level Trigger on GPUs for LHCb.

ESPPU26 #127 #148 #64

Heterogeneous TDAQ architecture

- ALICE, CMS & LHCb deploy **heterogeneous software triggers with CPU and GPU architectures since LHC-Run 3**
  - independently developed software frameworks, integrated with experiment software
  - Collaboration **crucial to share common tools** and experience with emerging processors

LHCb

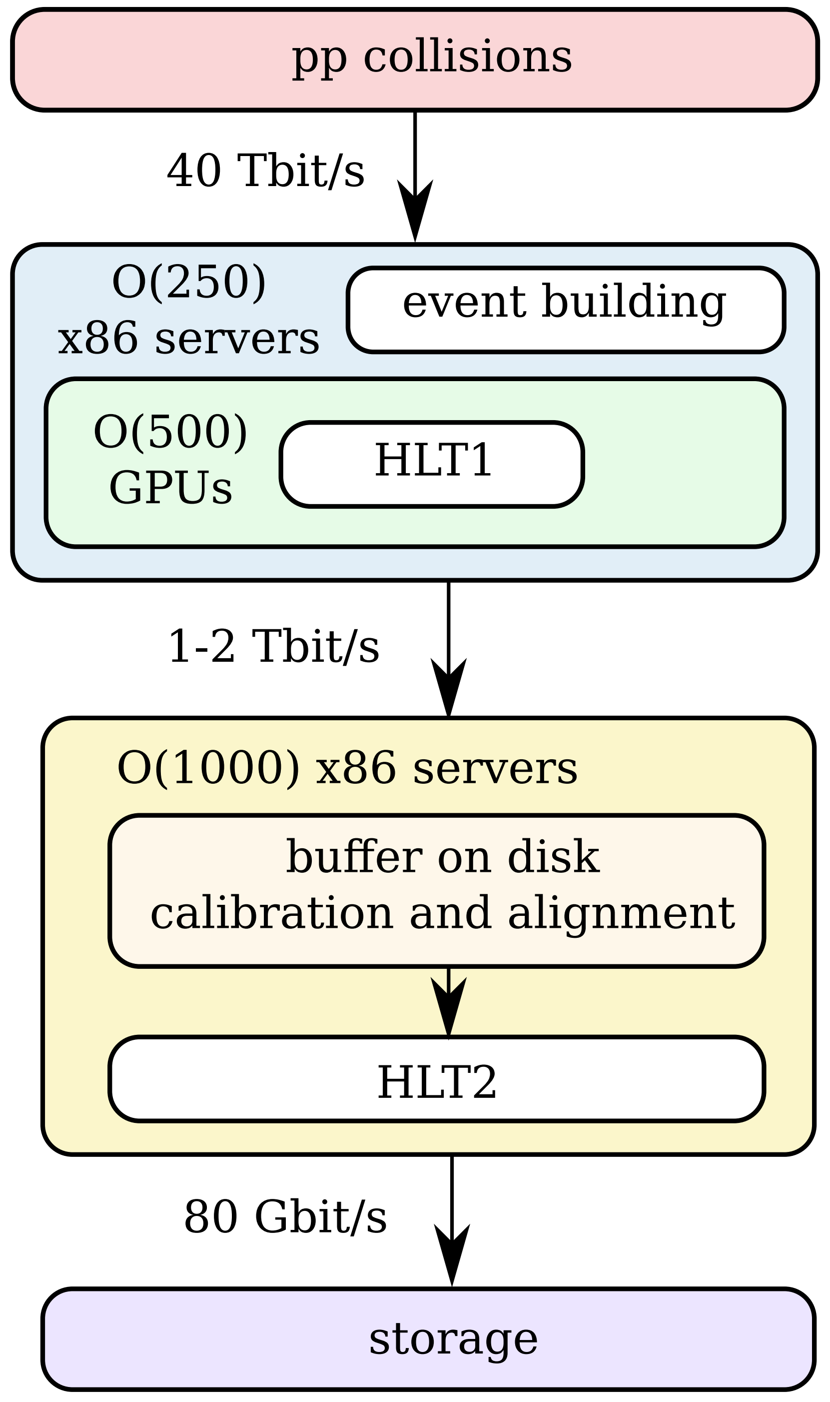

CMS

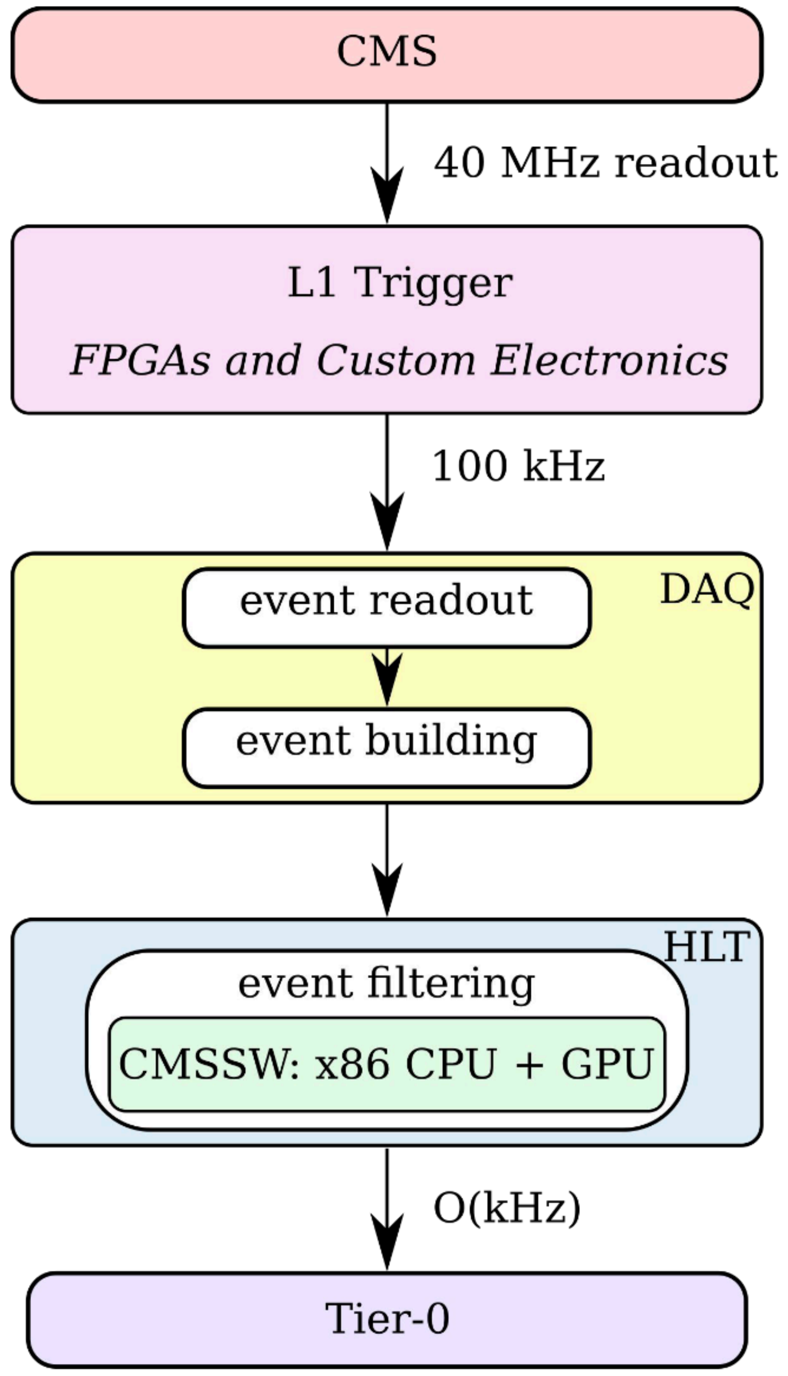

ALICE

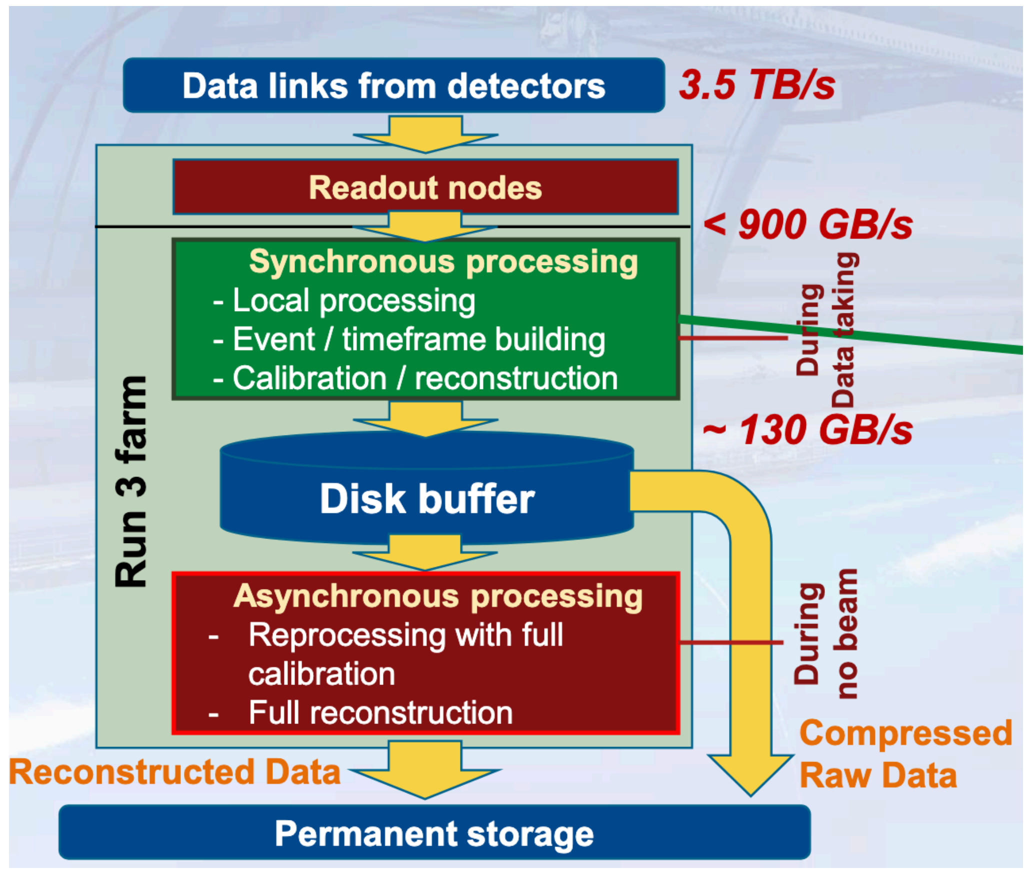

Real-time alignment & calibration

Reaching the FCC-ee precision needs

- High quality alignment & calibration crucial to minimise systematic uncertainties at future colliders and experiments
- Will need to cope with
  - high channel counts
  - timing-based reconstruction for 4D alignment
  - test case in ePIC and LHCb Run5
- Need fast calibration (limited buffering)
  - heterogeneous software frameworks developed for HLTs and advanced distributed computing techniques

LHCb's real-time alignment & calibration system deployed in 2018

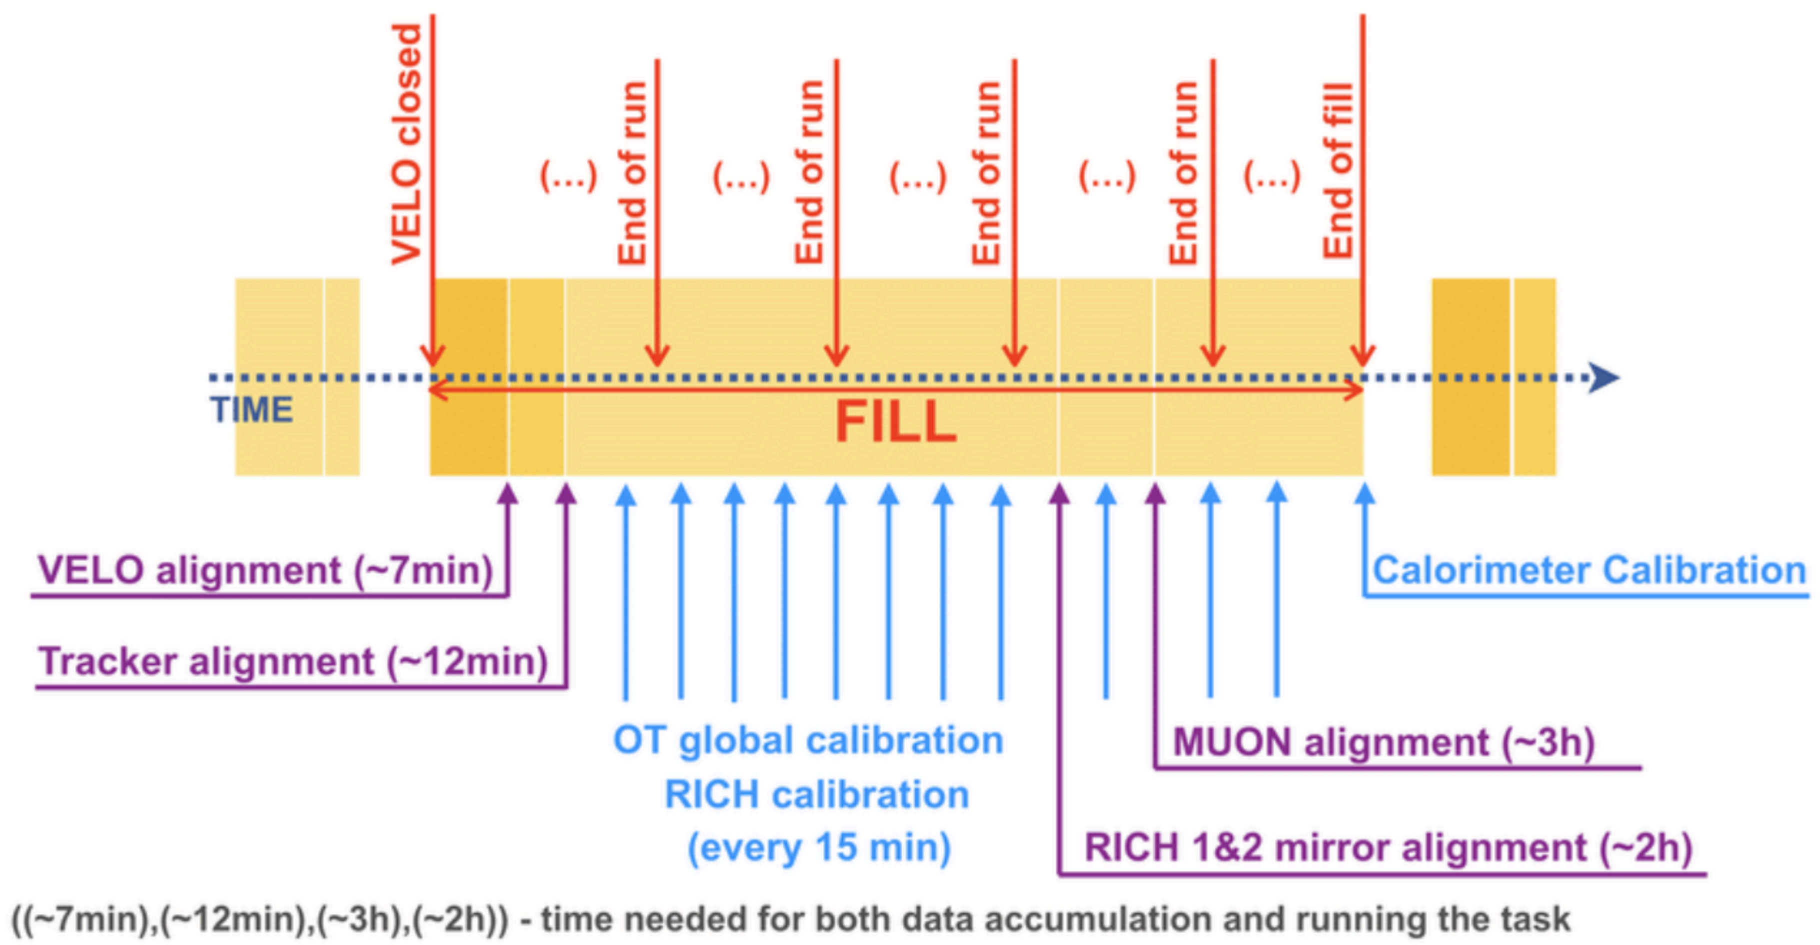

https://iopscience.iop.org/article/10.1088/1748-0221/14/04/P04013

CMS automated alignment calibration deployed since 2022

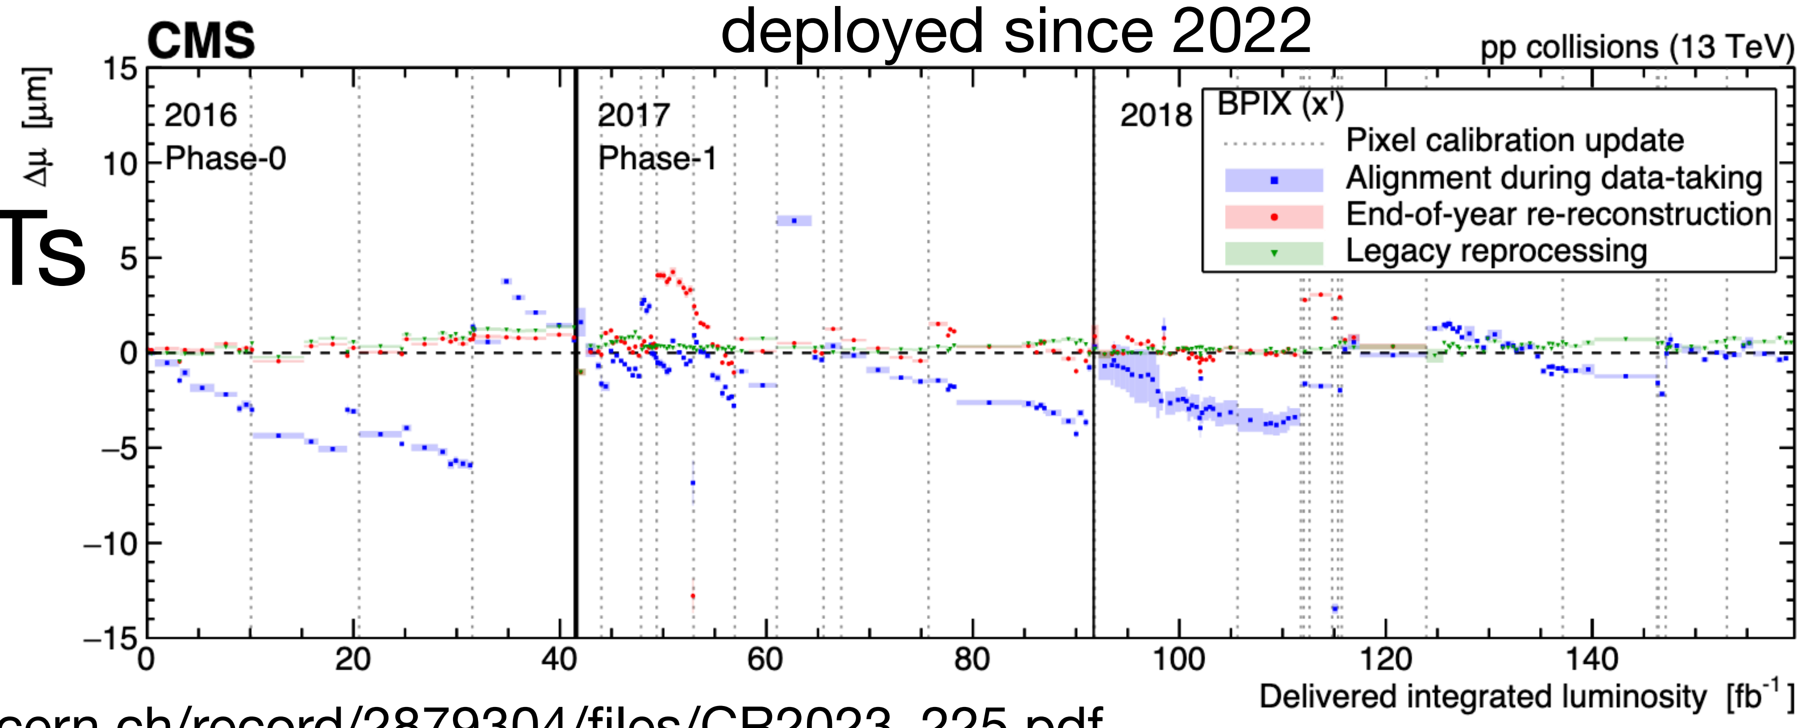

https://cds.cern.ch/record/2879304/files/CR2023_225.pdf

DRD7: R&D Collaboration on Electronics and On-Detector Processing

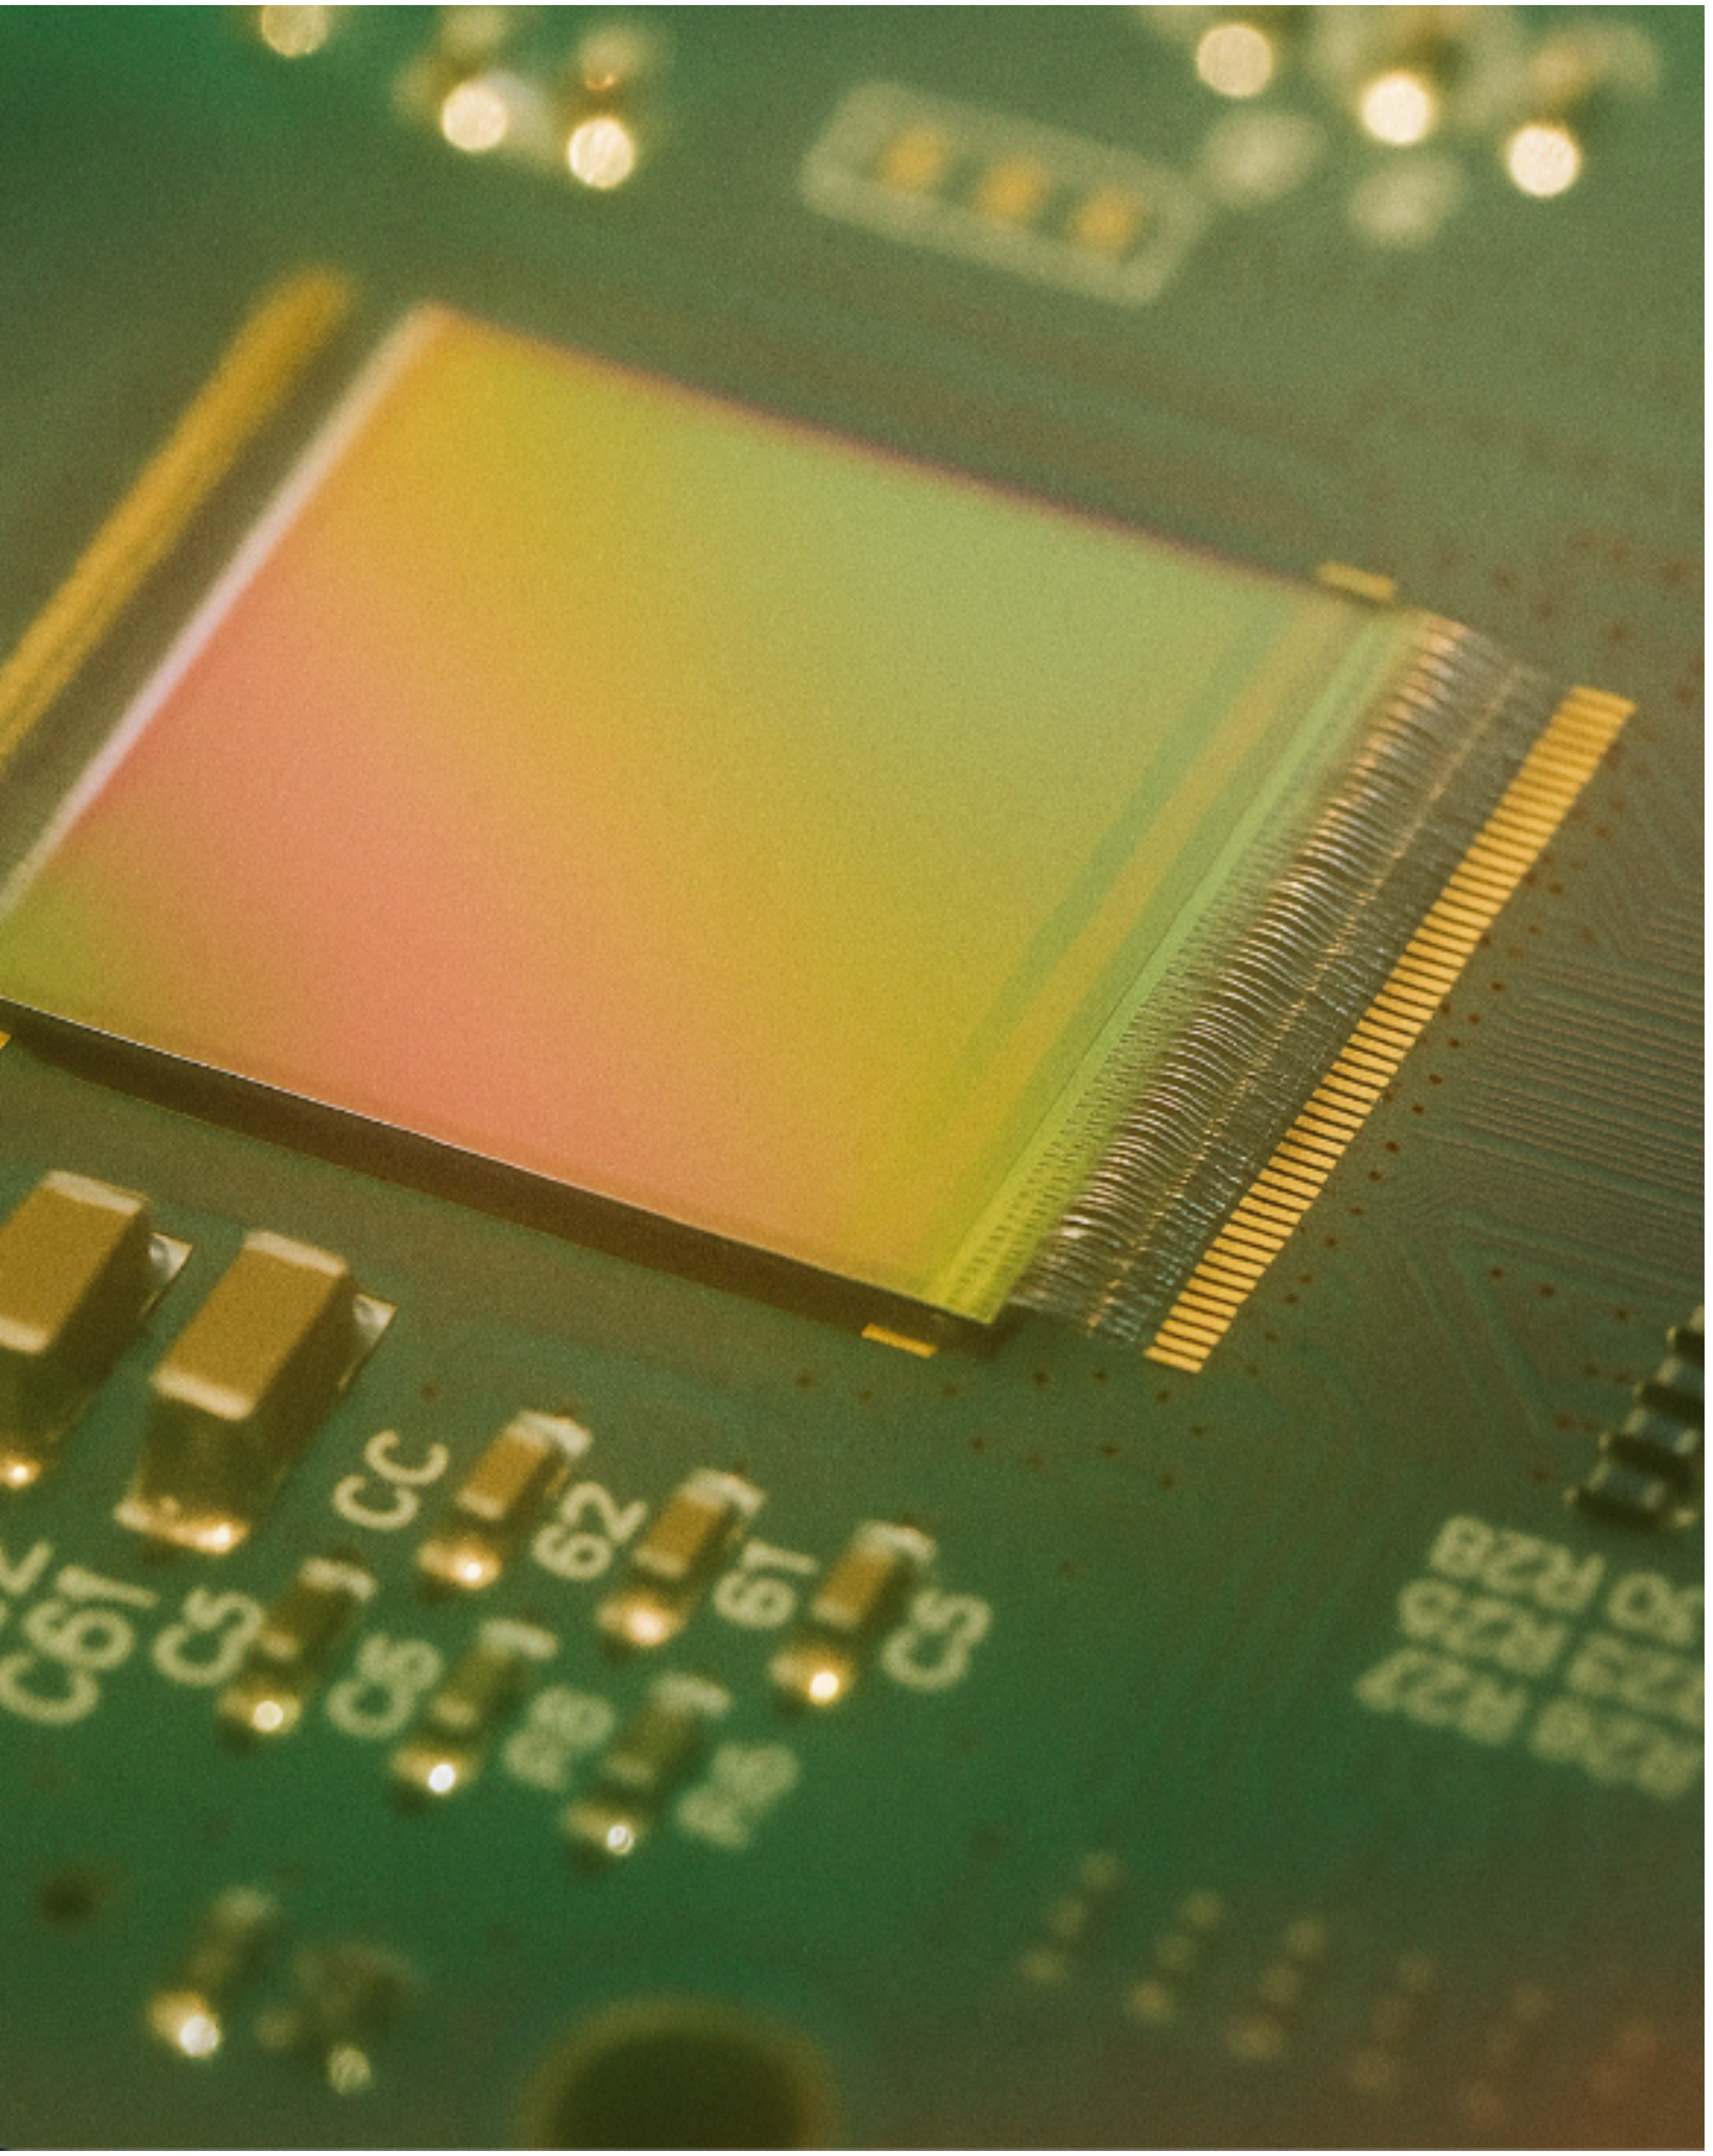

WP7.1: Data density and power efficiency

WP7.2: Intelligence on the detector

WP7.3: 4D and 5D techniques

WP7.4: Extreme environments

WP7.5: Backend systems and COTS components

WP7.6: Complex imaging ASICs and technologies

WG7.7: Tools and Technologies

Extend to include:

- **Real-time inference on specialized hardware.**
- **Heterogeneous real-time software tools**

in close collaboration with the Fast Machine Learning Foundation , CERN NGT and EuCAIF AI-RDs

AI R&D collaborations

EuCAIF proposal for scalable, robust AI through cross-domain collaboration

- **AI for Data Processing:** Front-end electronics, trigger
- **AI for Detector and Accelerator Control:** Accelerator performance, calibration, system monitoring.
- **AI for Detector Optimization:** Differentiable programming, reinforcement learning to maximize detector performance
- **AI for Event Reconstruction:** Tracking, calorimetry, end-to-end foundation models.

Conclusion

- **Exciting challenges ahead to make sure TDAQ is not the precision bottleneck**
    - read-out of increasingly granular detectors within tight material, power budgets
    - high fidelity on-detector compression for high-precision triggers or to reduce steaming data load for processing and storage
- AI/ML tools are essential to meeting these challenges
    - Must develop and maintain versatile heterogeneous frameworks and platforms
- Cross-experiment organisational support for frontend/backend ML developments represented in DRD collaborations/AI-RDs