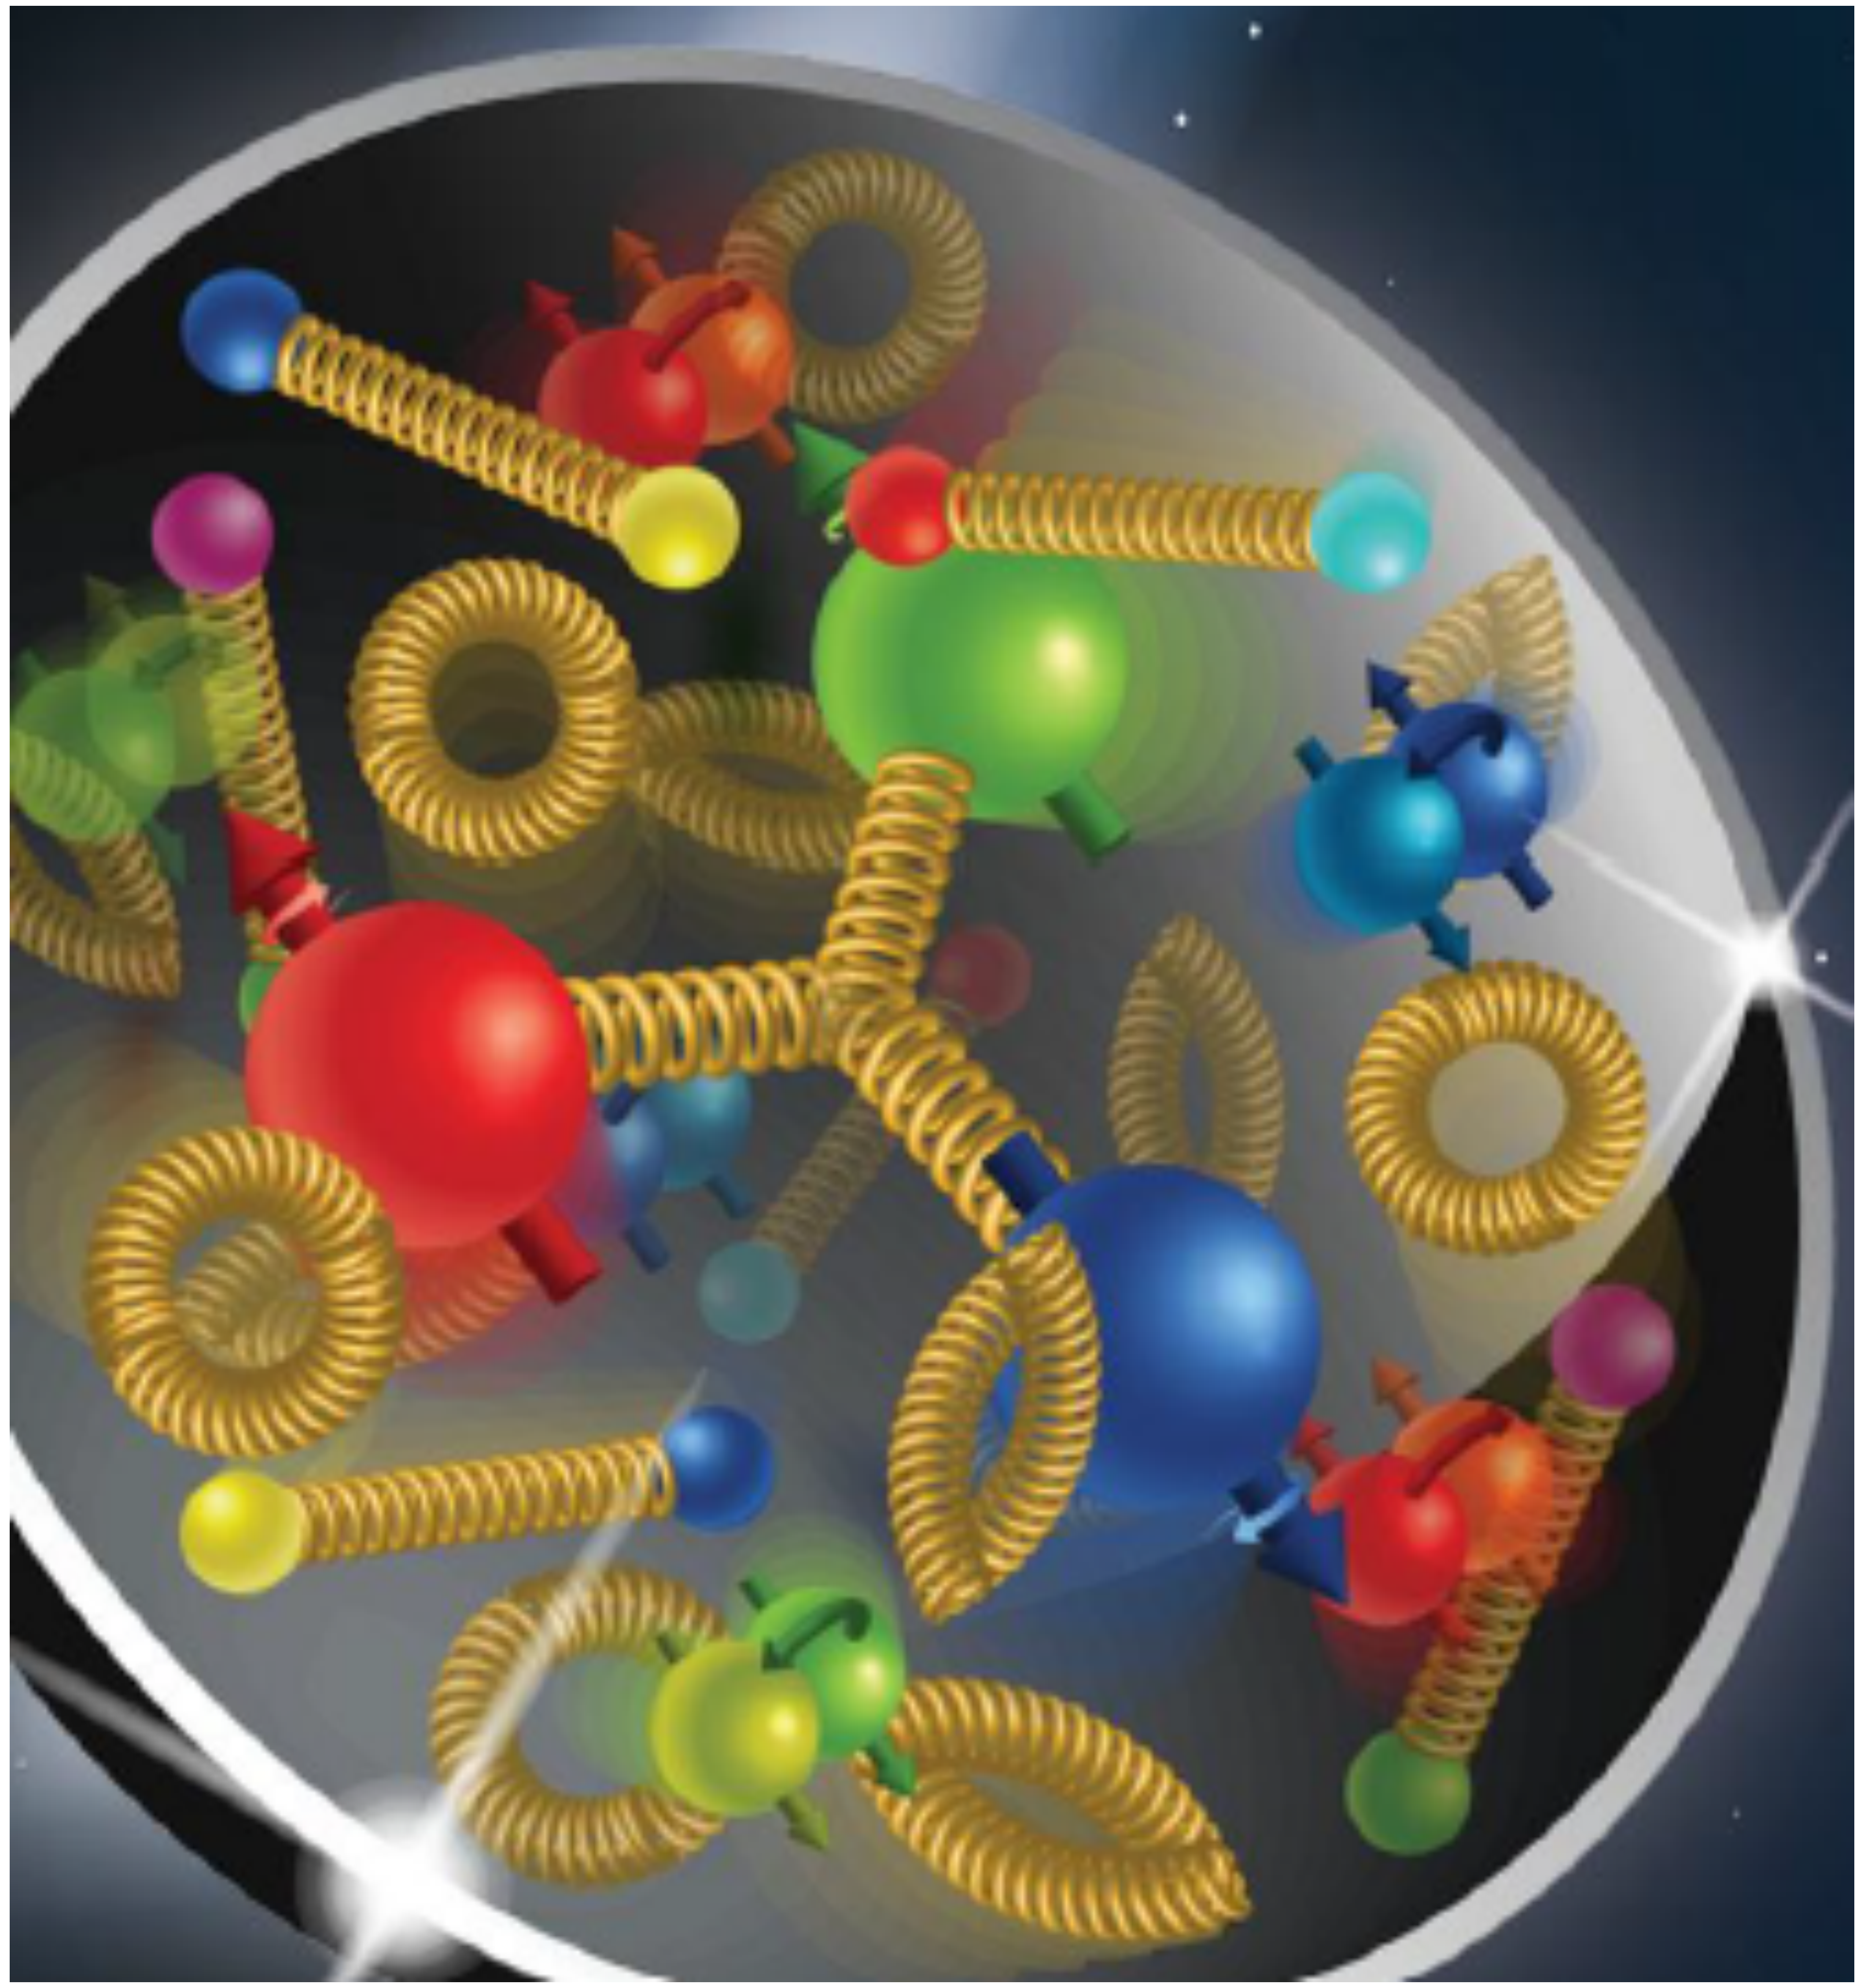

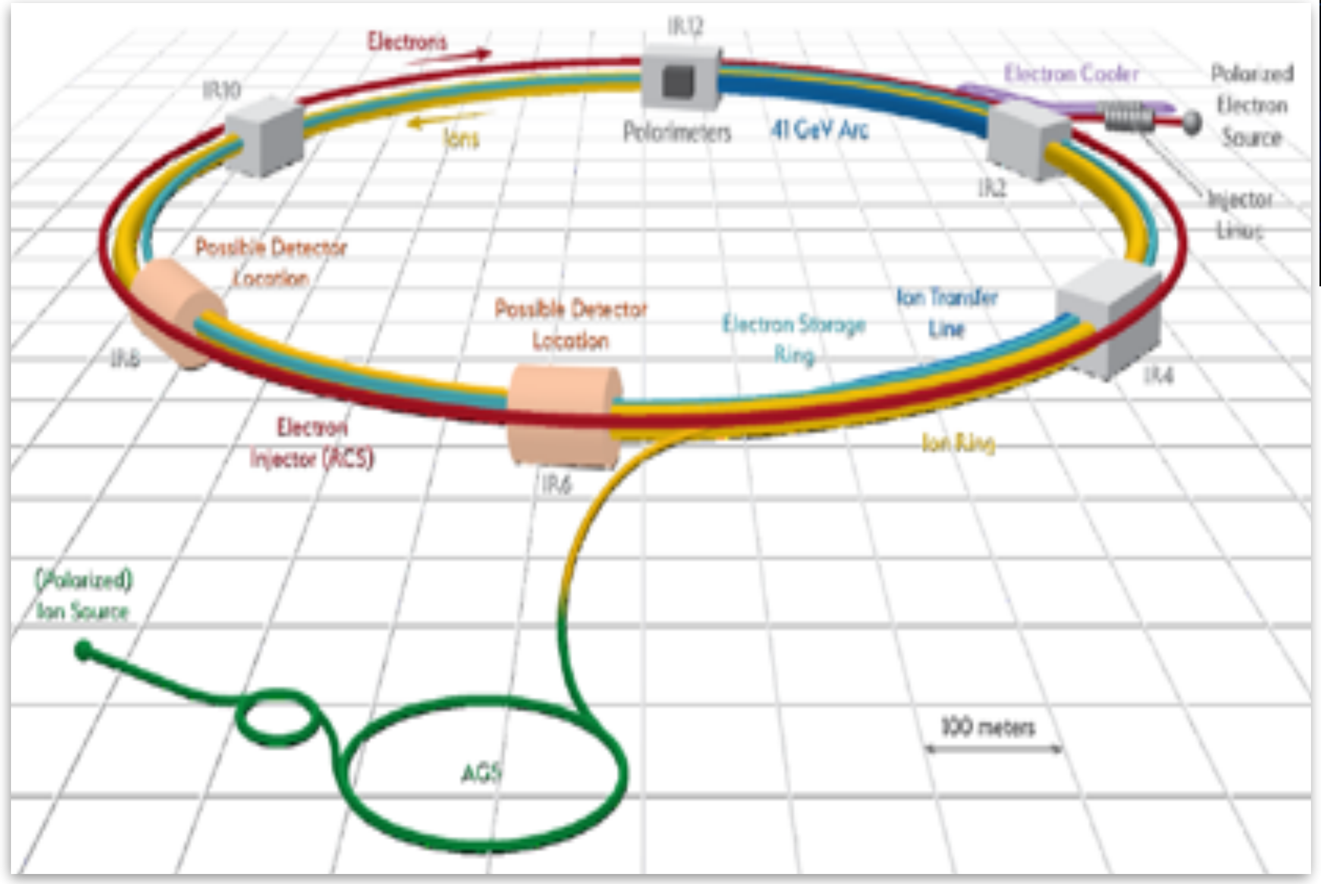

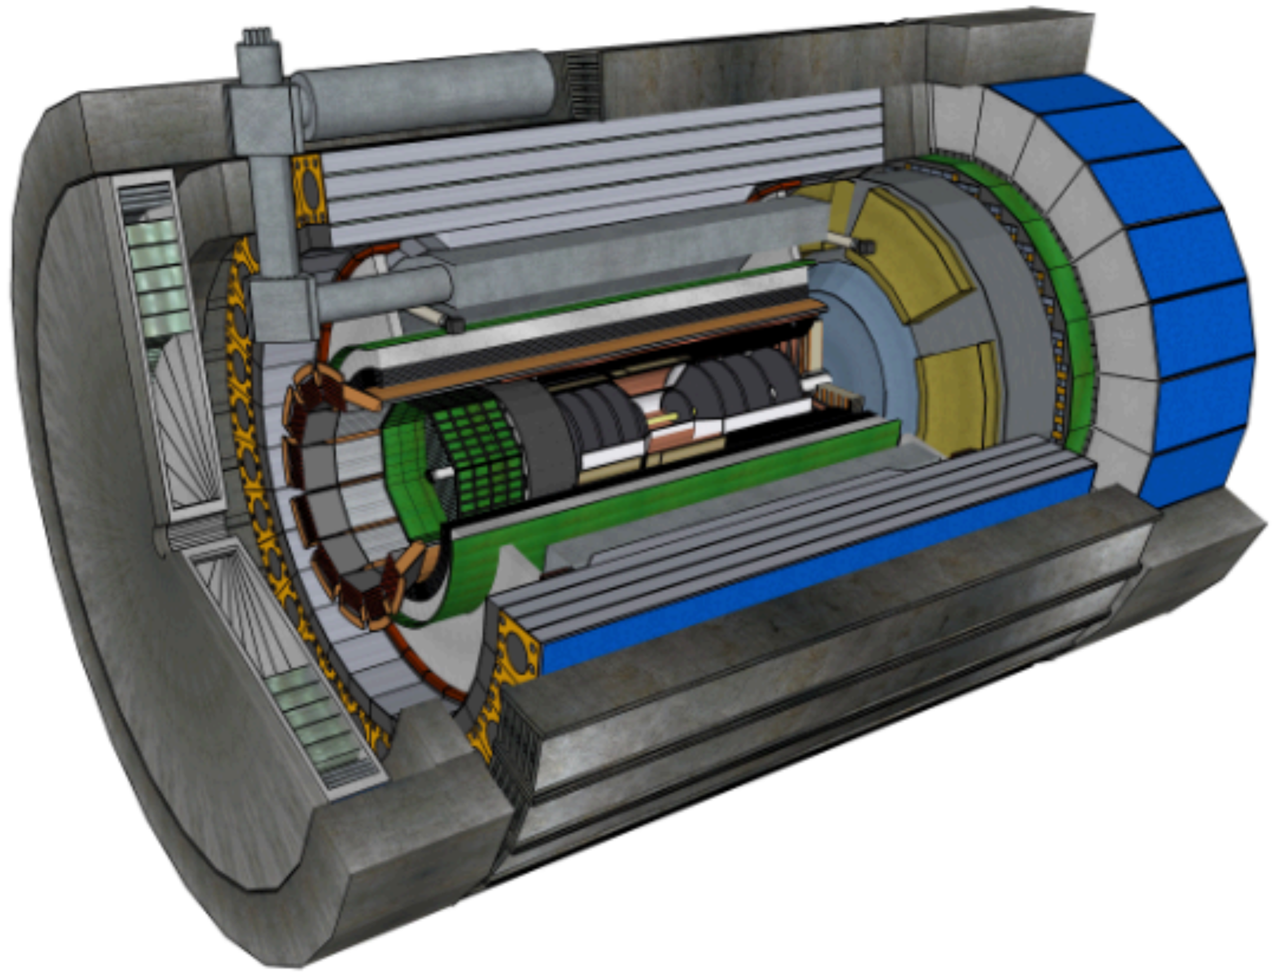

SRO WG Calibrations & Alignment: Summary of XIII Streaming readout workshop (SRO-XIII)

M.Battaglieri (INFN), M.Diefenthaler (JLab), T.Gunji (TokyoU), J.Landgraf (BNL), T.Wenaus (BNL)

ePIC Brookhaven National Laboratory Jefferson Lab e-JLab12 INFN

ePIC Streaming readout

Streaming readout Workshop SRO- XIII

https://agenda.infn.it/event/47630/overview

December 9-11, 2025 Catania (Italy)

- 70 registered participants(25-30 in person/day)
- 33 talks
- 2.5 days of presentations and discussions
- Significant pattendance from CERN, Japan, US

Topics

- **Streaming Readout Paradigms**
  Evolution from triggered to fully streaming DAQ architectures and free-running detectors.
- **Real-Time Processing & Orchestration**
  Online reconstruction, scheduling, data management, and orchestration frameworks for streaming data.
- **Heterogeneous Computing & Accelerators**
  Use of GPUs, FPGAs, and mixed architectures for low-latency data processing.
- **Machine Learning in Online Systems**
  Ultra-low-latency ML inference for triggering, tracking, calibration, and real-time decision making.
- **Streaming-Optimized Hardware & Frontends**
  ASICs, digitizers, high-speed links, and detector electronics designed for continuous data flow.
- **Calibration, Infrastructure & Cross-Experiment Experience**
  Autonomous calibration, infrastructure requirements, and lessons learned across major experiments (ePIC, LHC, neutrino, astroparticle, and gravitational-wave detectors).

Real time analysis on heterogeneous architectures with Allen Core
LHCb and ePIC use cases

Gonzalo Díaz López – LPNHE

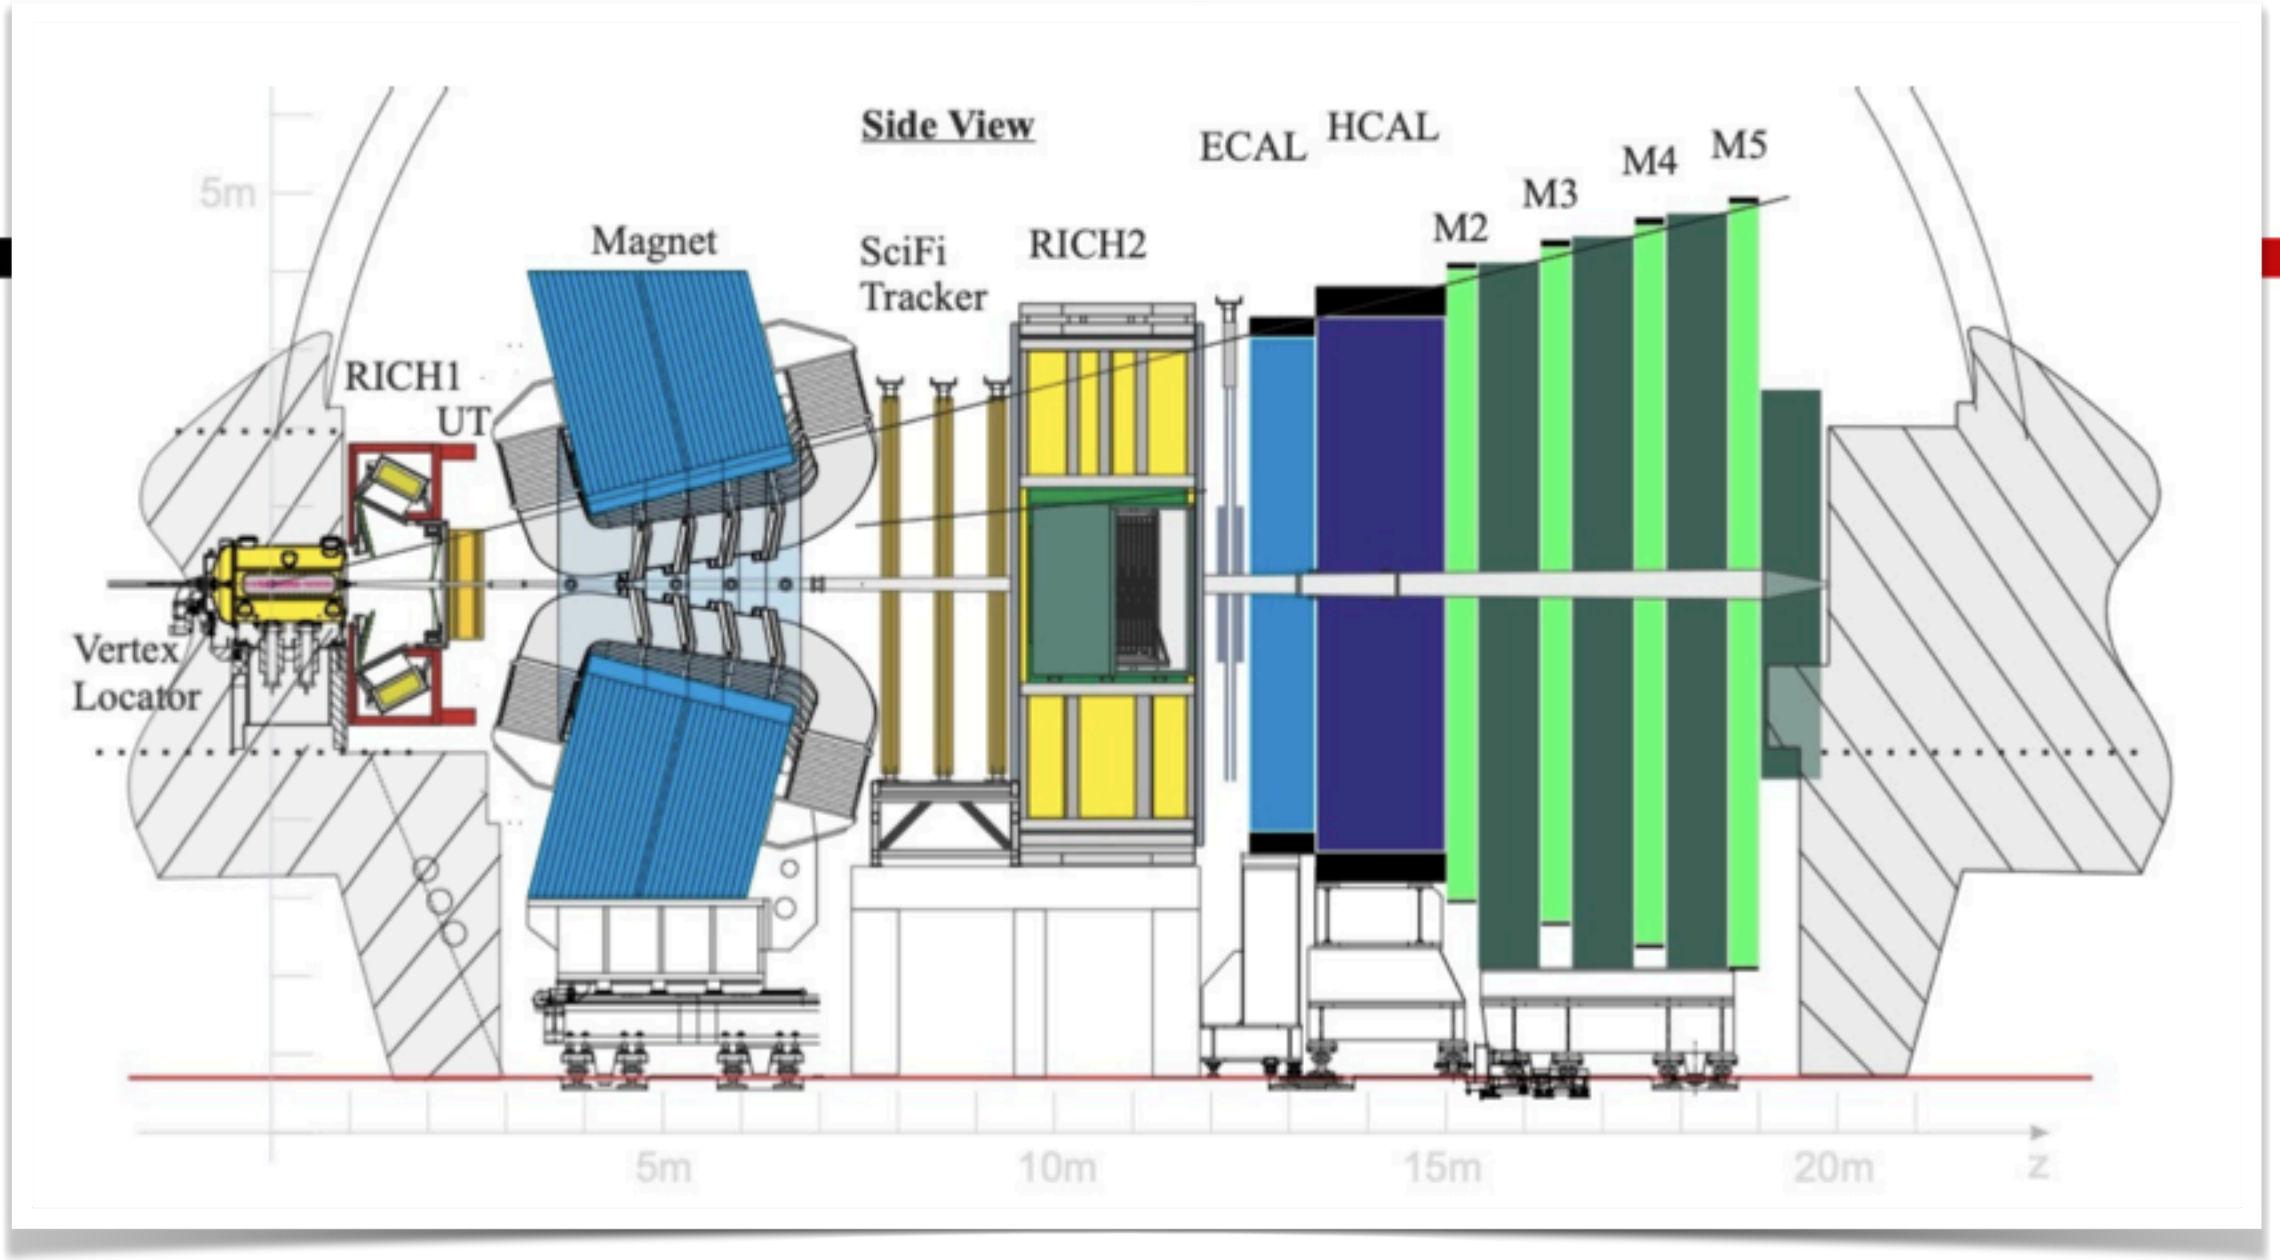

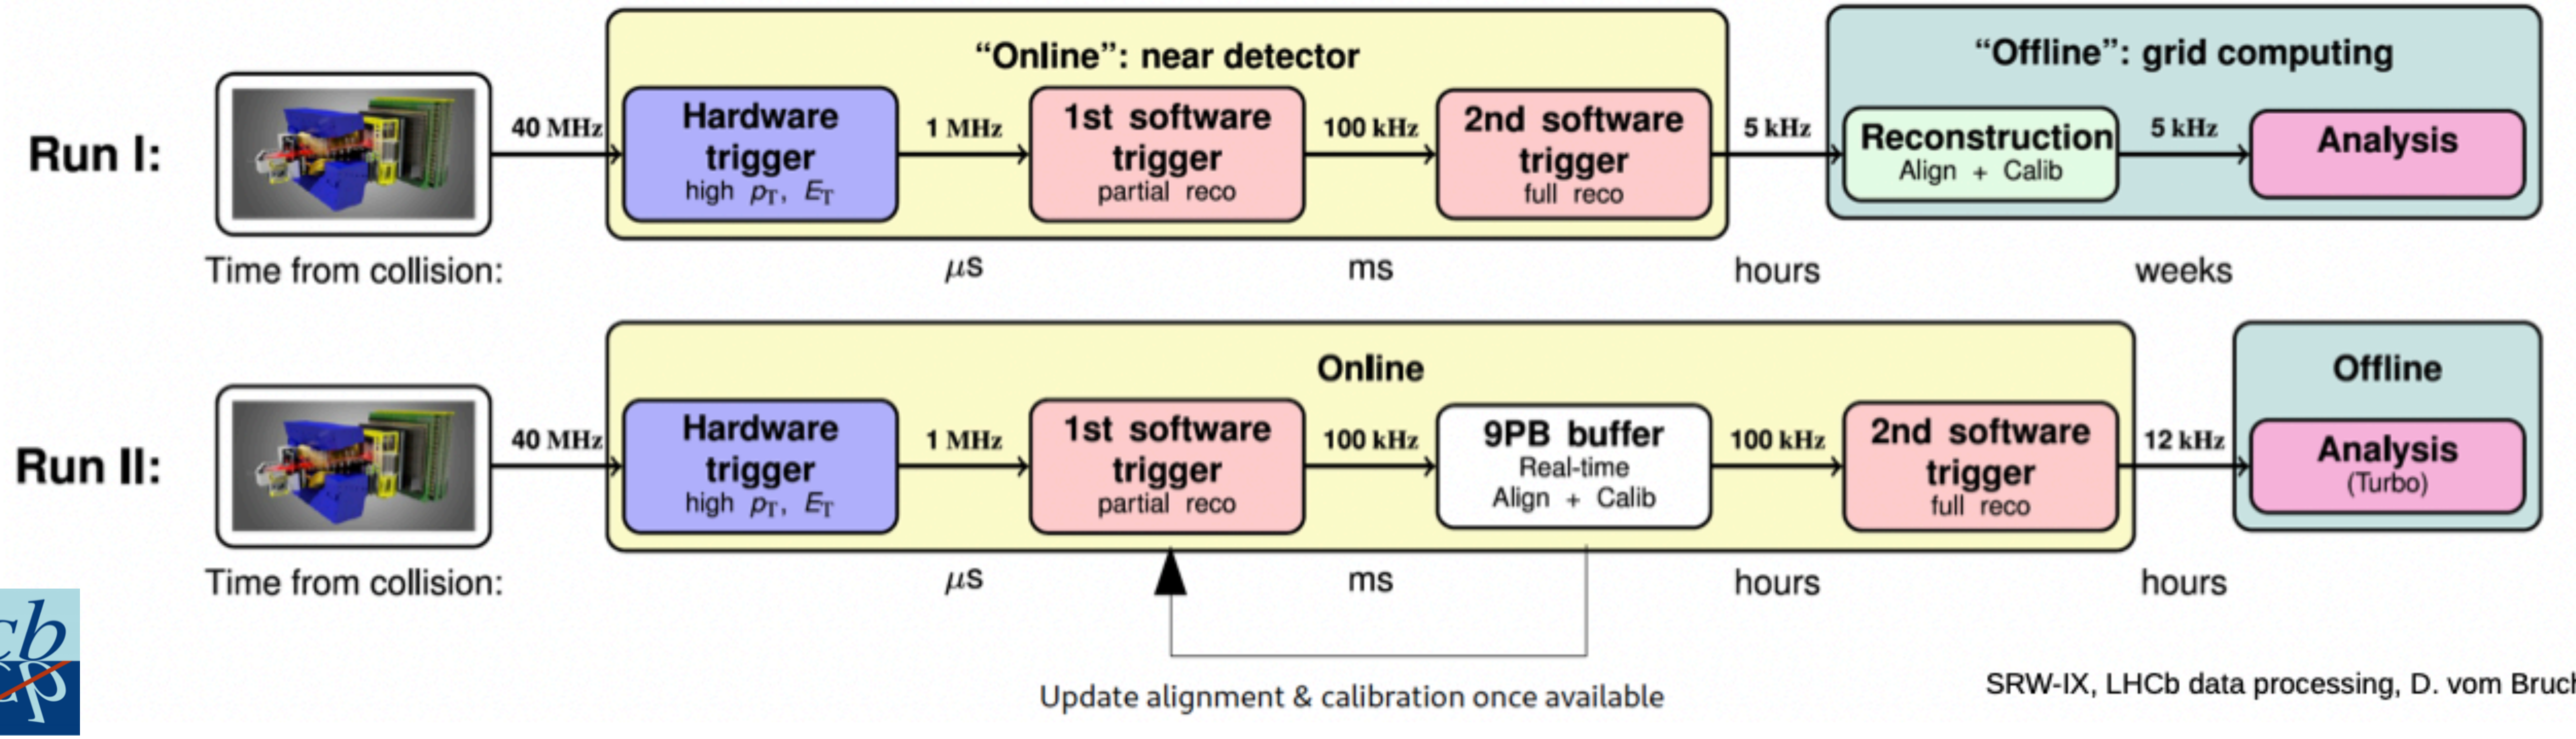

SRW-IX, LHCb data processing, D. vom Bruch

Run III:

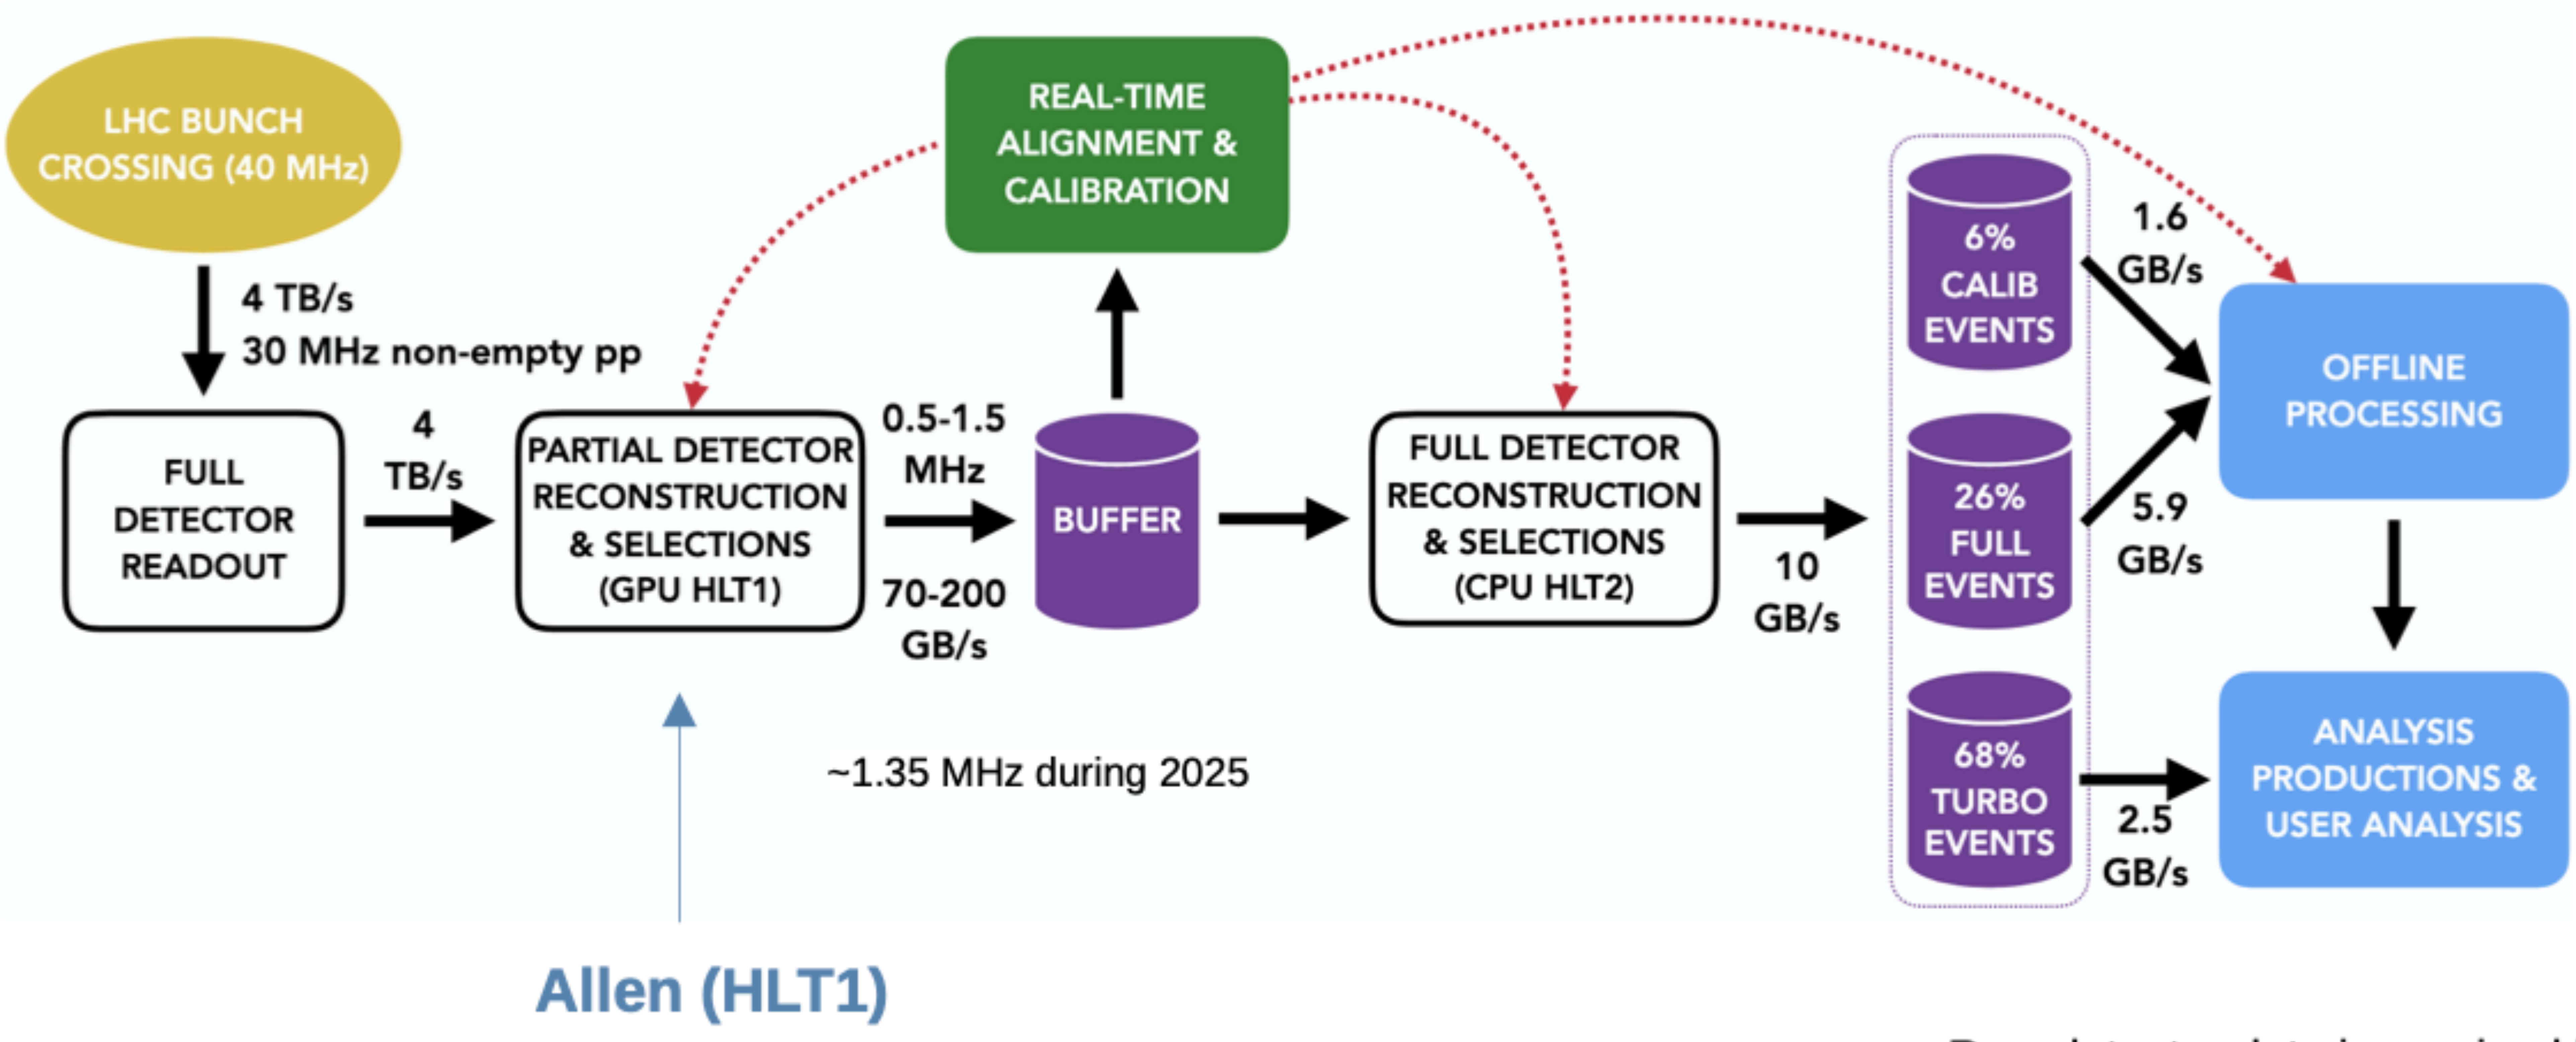

- Raw detector data is received by FPGA cards
- 173 event builder servers:
  - aggregate data from sub detectors
  - 3 GPUs/server run HLT1 (~ 500 Nvidia A5000) – extended from TDR 2 GPU/server during data taking period
- HLT1 output temporarilly stored on buffer (~40 PB) while alignment & calibration is processed
- HLT2 is processed on computing farm with 250k CPU cores

M.Battaglieri - INFN/
T.Gunji TokyoU

Allen Core – a general processing framework for heterogeneous architectures

- **Decouple** framework from application
- **Refactoring** and clean up
- **Simplify** user experience
- Improve adaptability to different **devices**
- Provide up-to-date **documentation** and **maintainability**
- **Interoperability**
  - Generalise services (provide interfaces for external projects)
  - APIs to core functionalities (e.g. Streams)
  - Modular
  - Smooth integration with LHCb's workflow
- Work **in progress**, to be fully developed during LS3 (2026-2030)
- First version to be released during **Q1 of 2026** – ODD demonstrator
- Interest from other experiments (**ePIC**)

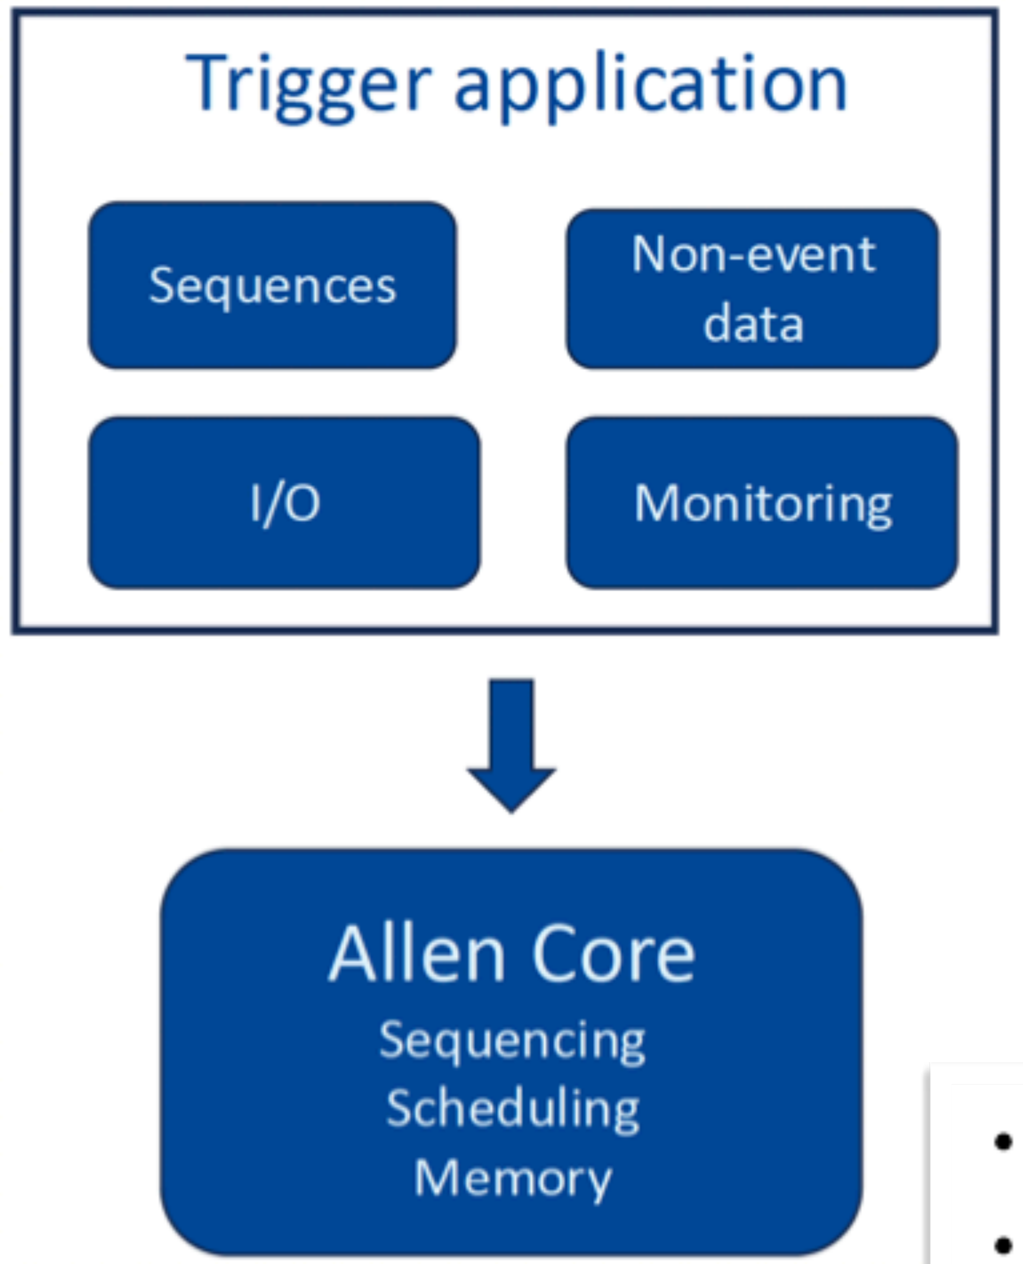

- Near future: similar conditions for **Run 4**
- Major challenge for **Run 5 (Upgrade II)**:
  - ~x10 luminosity increase
  - pileup from ~5 to ~35, O(10ps) timing needed
  - HLT1 input bandwidth from 4 to 25 TB/s
- Processing the full reconstruction on GPUs seems nowadays the only viable option
- Best suited architecture to be determined
- Keep two level trigger (HLT1 and HLT2) to be determined
  - Need to run HLT2 on GPUs

Allen4EIC project

- **Interest** in using Allen Core as online reconstruction and calibration tool for ePIC (Echelons 1-2)
- **A first application** as a demonstrator for the calibration of the ePIC backward electromagnetic calorimeter (under design and construction at IJCLab)
- **Deploy a pipeline** on a French computing cluster, emulating a future EIC Echelon 2
- Allen Core use case **from trigger filter to real-time** reconstruction tool
  - External application: ePIC's data model, geometry and algorithms

ALLEN 4EIC A HETEROGENEOUS, OPEN-ACCESS FRAMEWORK FOR REAL-TIME HIGH-THROUGHPUT DATA PROCESSING

HORIZON-INFRA-2025

- LHCb's paradigm shift from a hardware to a software trigger was very successful, in part due to the increasing feasibility of using GPUs for high-end, high-throughput applications over the past decade
- Allen (HLT1) trigger application targets only LHCb use case right now – Allen Core framework project
- Working on an Allen Core application demonstrator using Open Data Detector to be released in Q1 2026
- Allen Core relevant in LHCb's future context – running HLT2 on GPUs

Data and memory management, scheduling, and experience about using GPUs in ALICE online and offline

David Rohr for the ALICE Collaboration, CERN
Streaming Readout Workshop XIII
9.12.2025
drohr@cern.ch

Targeting to record large minimum bias sample.
- Access low S/B "untriggerable" signals
- All collisions stored → no trigger
- Continuous readout → data in drift detectors overlap
- Recording time frames of 2.8 ms of continuous data, instead of events
- 100x more collisions → need online compression to reduce data volume
- Using common Online Offline (O^2) framework for reconstruction, simulation and analysis

→ GPUs speed up online (and offline) processing

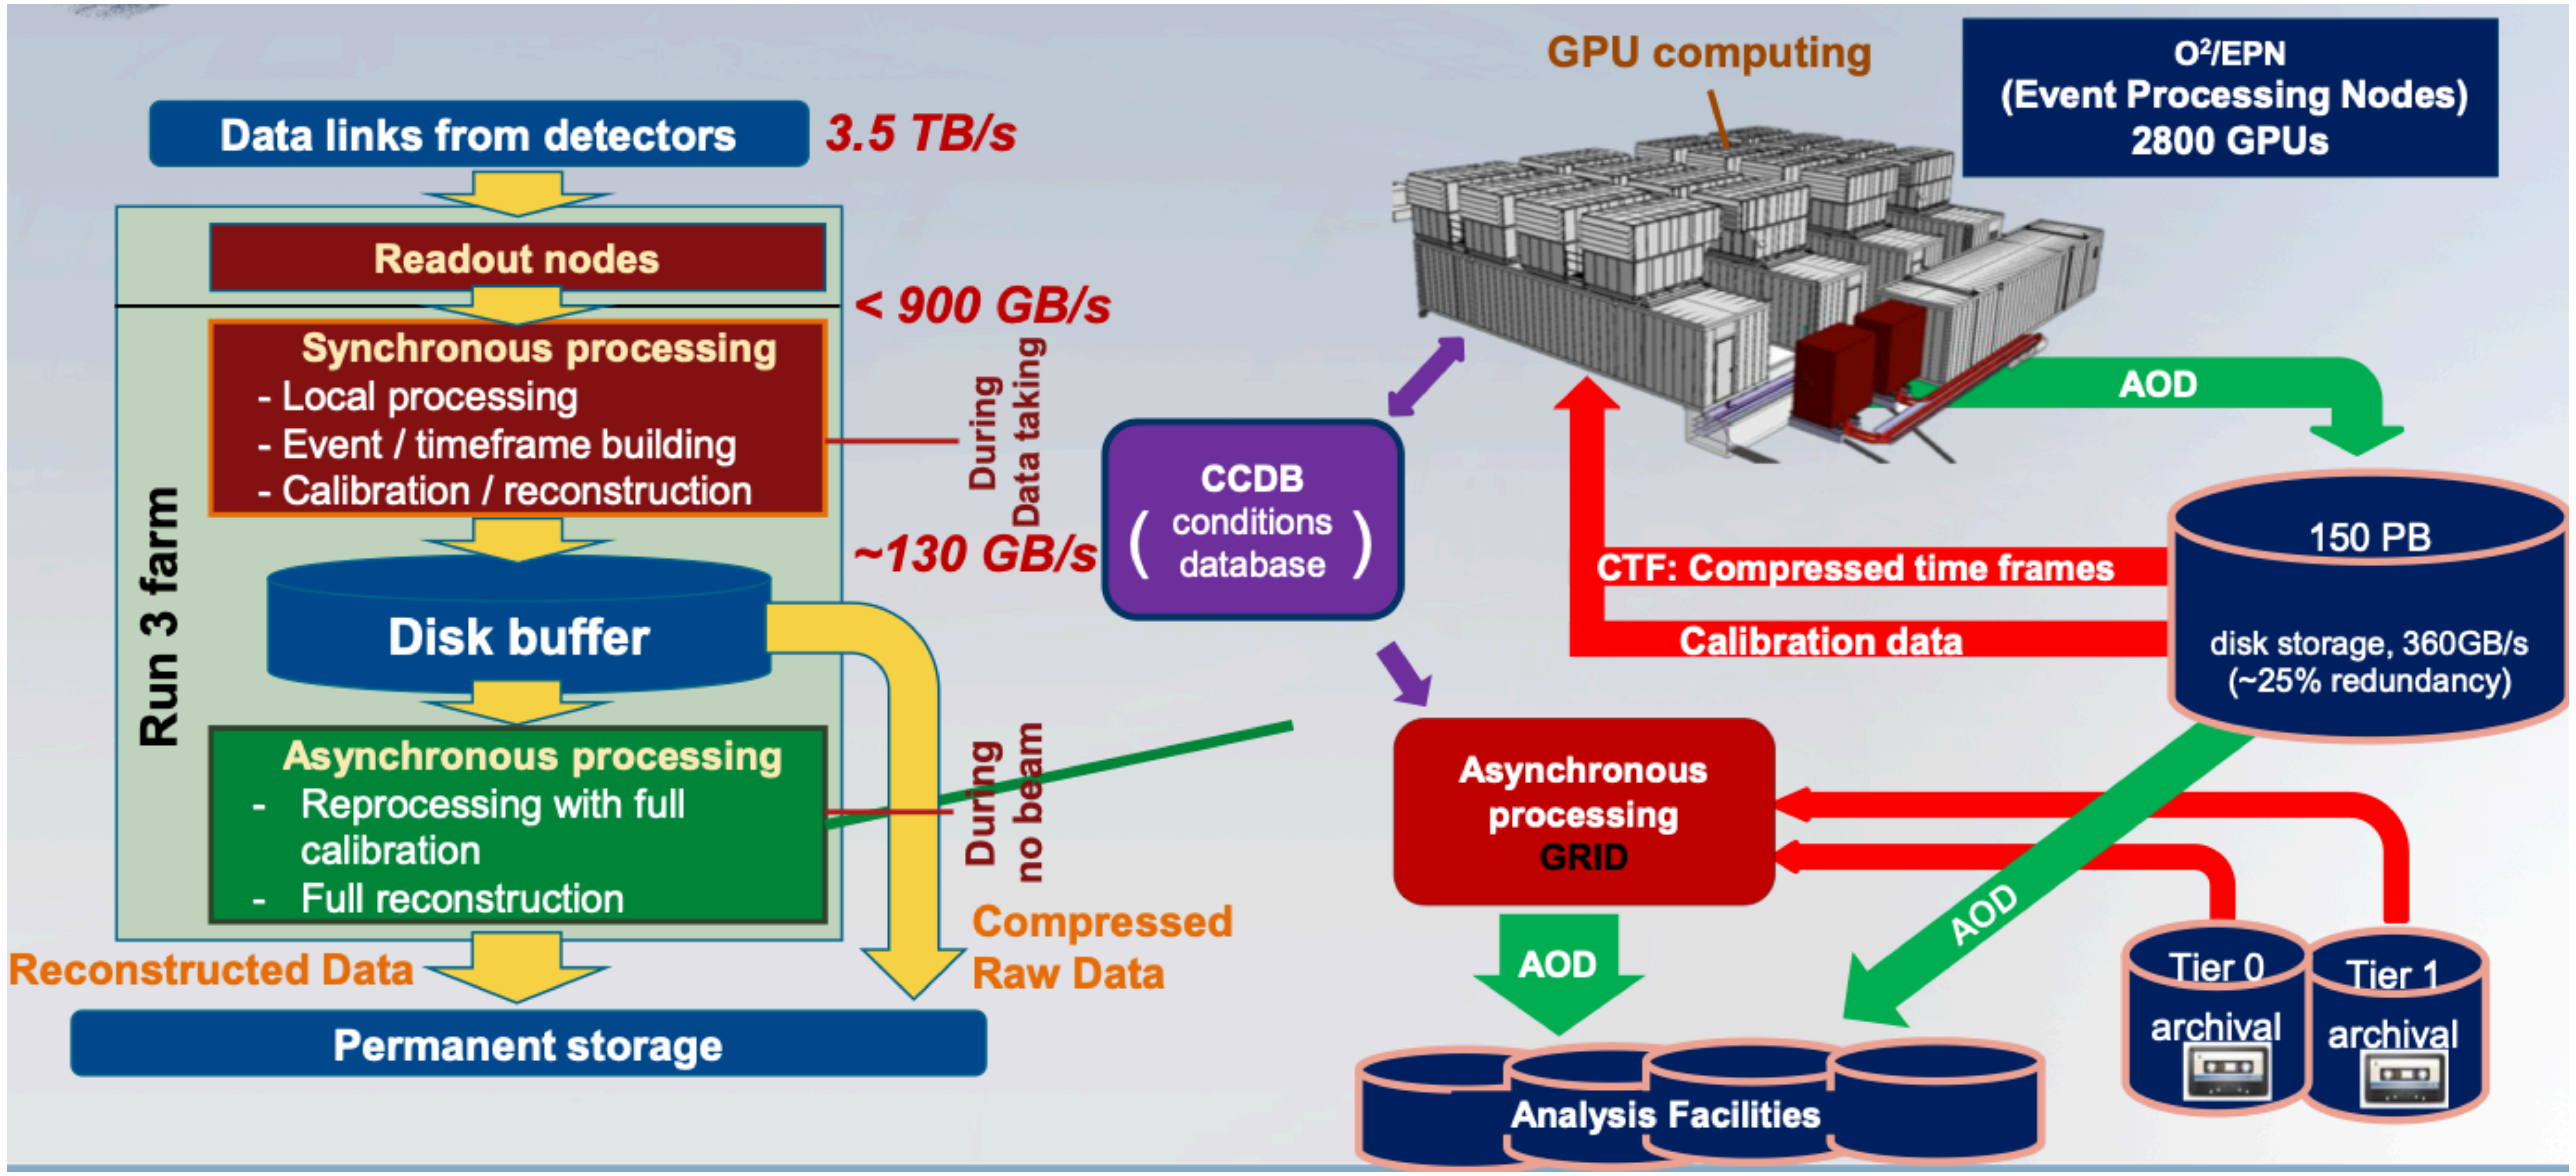

- **Online processing:**
  - Extract information for **detector calibration**:
    - Previously performed in 2 offline passes over the data after the data taking
    - Run 3 avoids / reduces extra passes over the data but extracts all information in the sync. processing
    - An intermediate step between sync. and async. processing produces the final calibration objects
    - The most complicated calibration is the correction for the TPC space charge distortions
  - **Data compression**:
    - TPC is the largest contributor of raw data, and we employ sophisticated algorithms like storing space point coordinates as residuals to tracks to reduce the entropy and remove hits not attached to physics tracks
    - We use ANS entropy encoding for all detectors
  - **Event reconstruction** (tracking, etc.):
    - Required for calibration, compression, and online quality control
    - Need **full TPC tracking** for data compression
    - Need tracking in all detectors for ~1% of the tracks for calibration
    - → **TPC tracking dominant part**, **rest** almost negligible (**< 5%**)

Needs tracking of 1% of tracks

Needs 100% TPC tracking

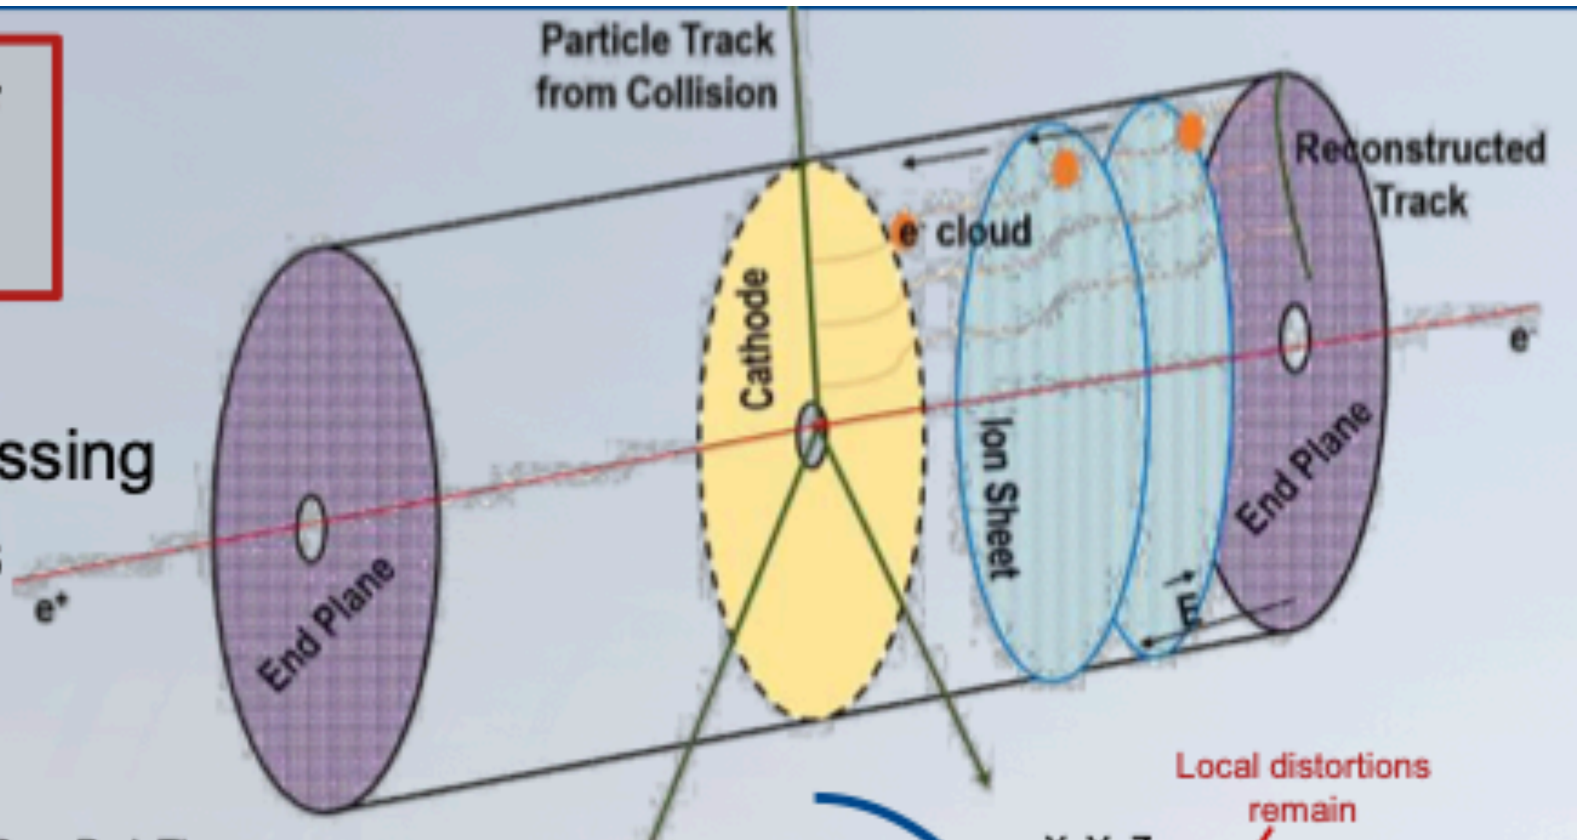

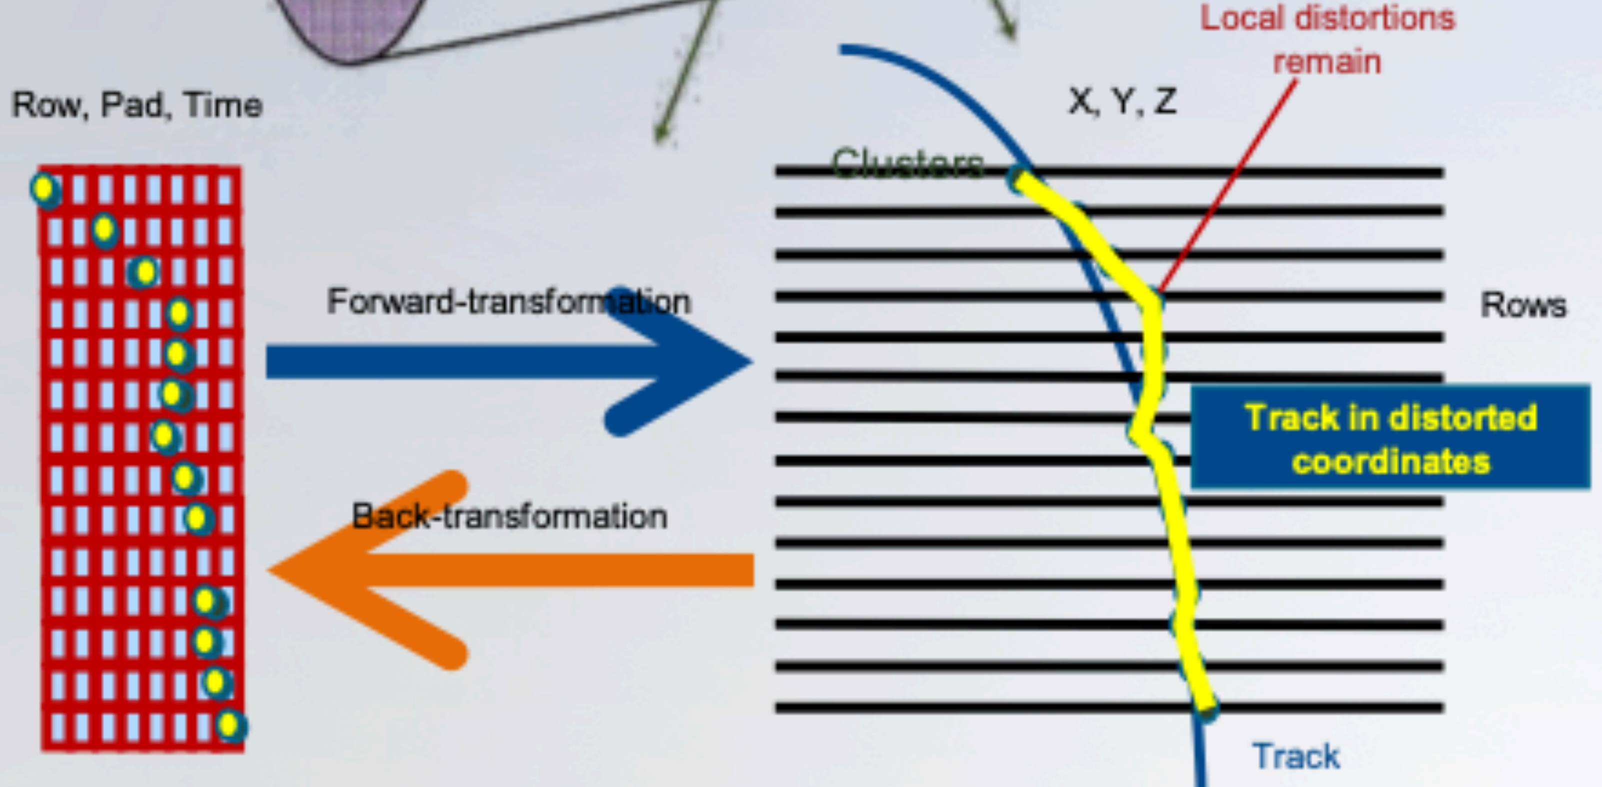

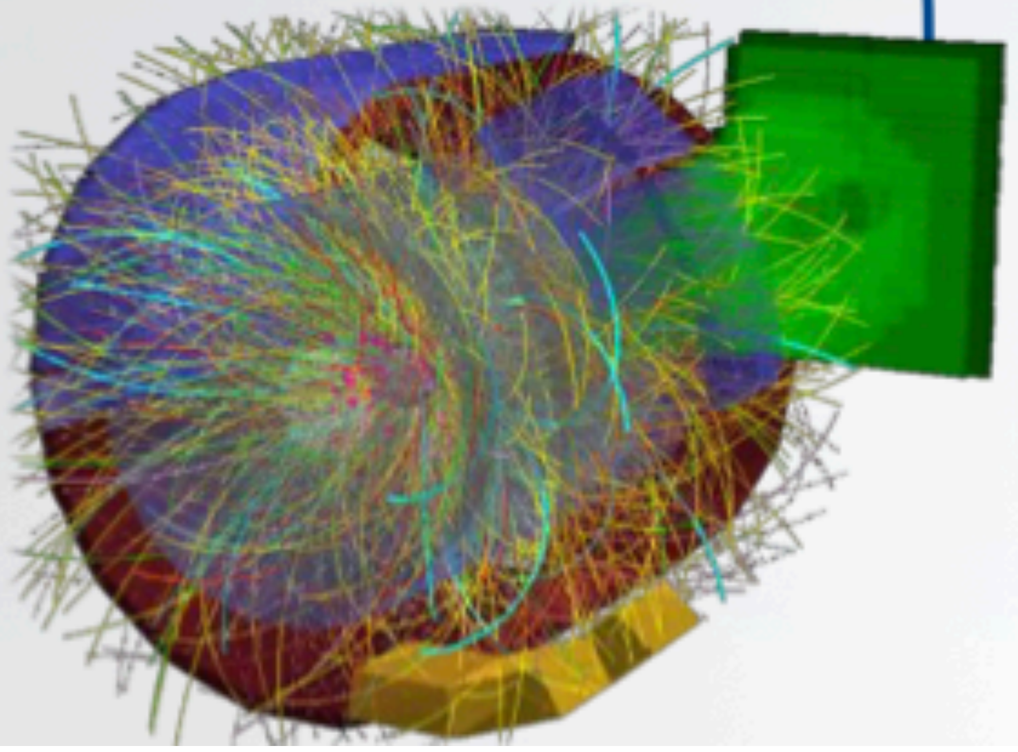

- **Offline processing:**
  - **Full reconstruction**, **full calibration**, **all detectors**
  - TPC part faster than in synchronous processing (less hits, no clustering, no compression)
  - → **Different relative importance of GPU / CPU** algorithms compared to synchronous processing

Needs 100% of everything

M.Battaglieri - INFN/
T.Gunji TokyoU

Introduction to
The KM3NeT data acquisition system

Tommaso Chiarusi
Emidio Giorgio
Francesco Benfenati Gualandi

INFN

Sezione di Bologna
Laboratori Nazionali del Sud

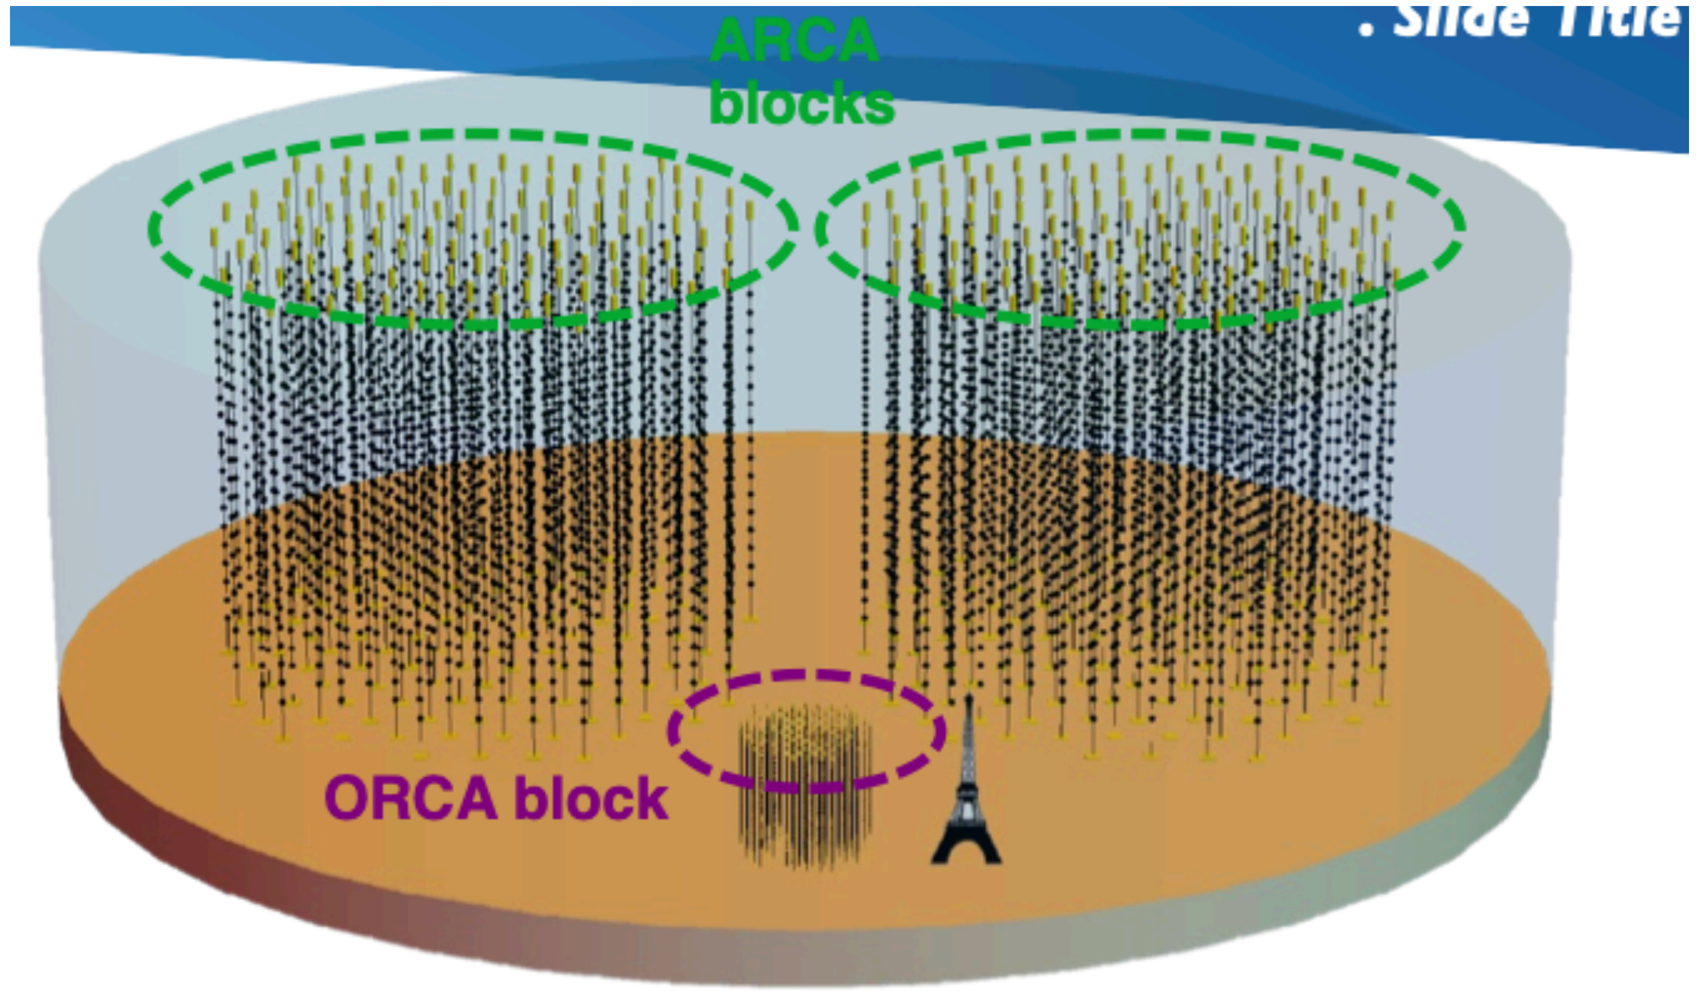

	ARCA	ORCA
Location	Italy (Sicily)	France (Toulon)
Anchor depth	3450 m	2450 m
Distance from shore	100 km	40 km
DUs	115x2 blocks	115
DU horizontal spacing	90 m	20 m
DOM vertical spacing	36 m	9 m
DOMs/DU	18	18
PMTs/DOM	31	31
Instrumented water mass	1 Gton	7 Mton
DUs deployed so far	48-53	33

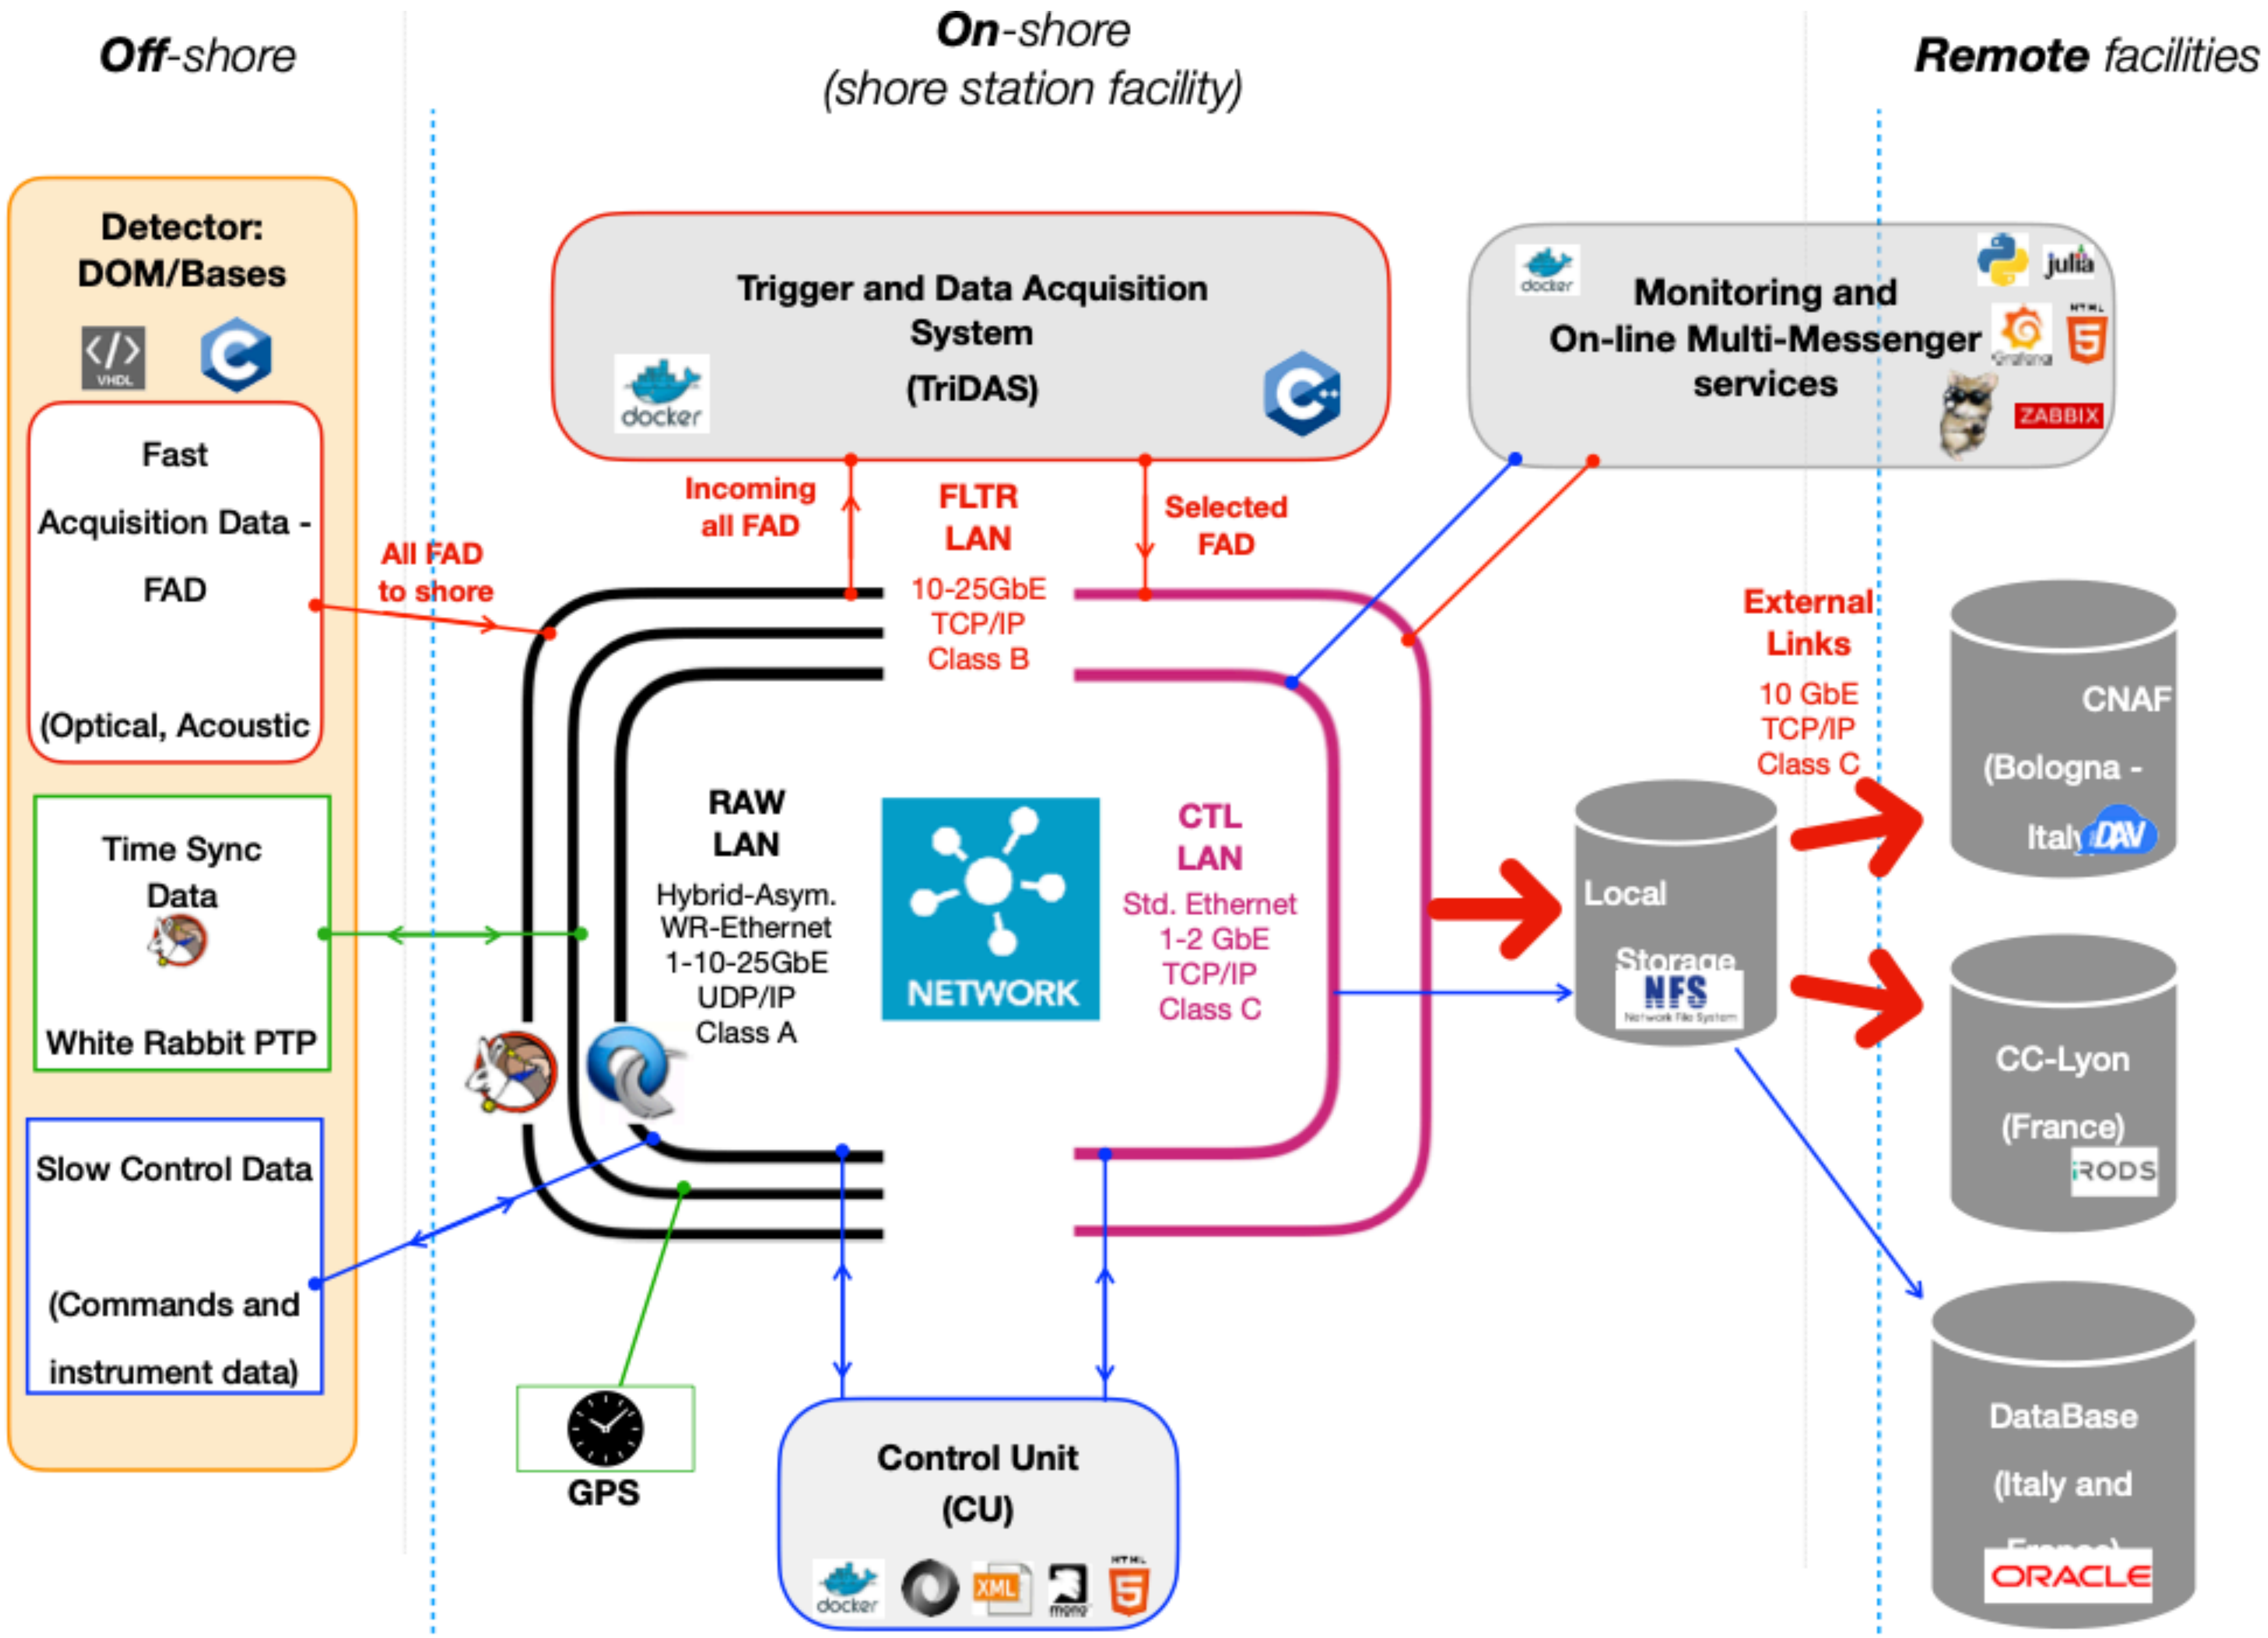

Basic triggers

L0: all hits over threshold (i.e. all hits sent by the CLBs)
L1: pairs of hits of the same DOM within 25(10)ns.
L2: further constraints applied to L1 hits (e.g. space angles btw PMT axes)

Trigger settings passed to the Data Filters via the run setups by the Control Unit

Higher-trigger level

- **3D-Trigger** - general concept:

1. A minimum n. of ***consecutive*** L2 s ≥ N_{th} within a ΔT (at least $n_{DOM} \geq 2$ or 5)
2. *3D*-causality filter : $|t_i - t_j| \leq |\vec{x}_i - \vec{x}_j|\frac{n}{c} + T_{MaxExtra}$
3. The trigger is set if the n. of satisfying hits is ≥ N'_{th}

- **3D-Muon/Shower**
  Assumes an extended track-like / short pulse shape for the event topology
- **MX-Shower**
  Cluster one L2 with causality-combined L0s.
- **Supernova (SN)**
  Combines L1 with additional constraints (e.g. multiplicity of L0 hits)

Trigger algorithms are developed within a large C++ software framework, Jpp.
The same codes are used for the on-line DAQ as well as off-line analysis.

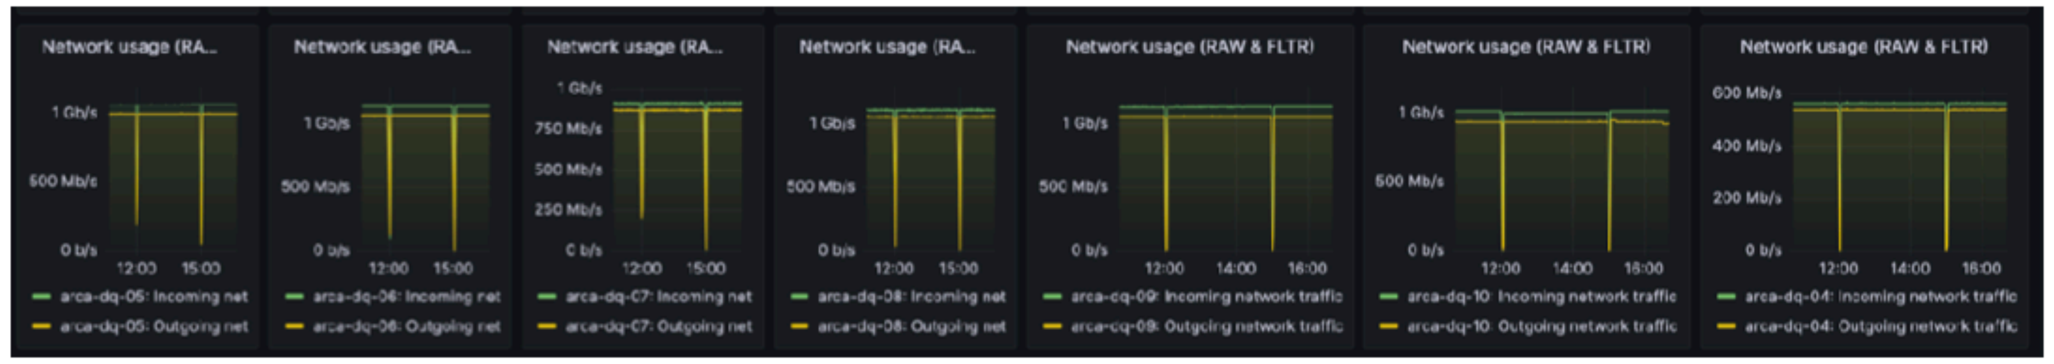

DataQueue level: receive and route to Data Filters (O+A)

to Data Writer (i.e. data filtered by ~3 order of magnitude)

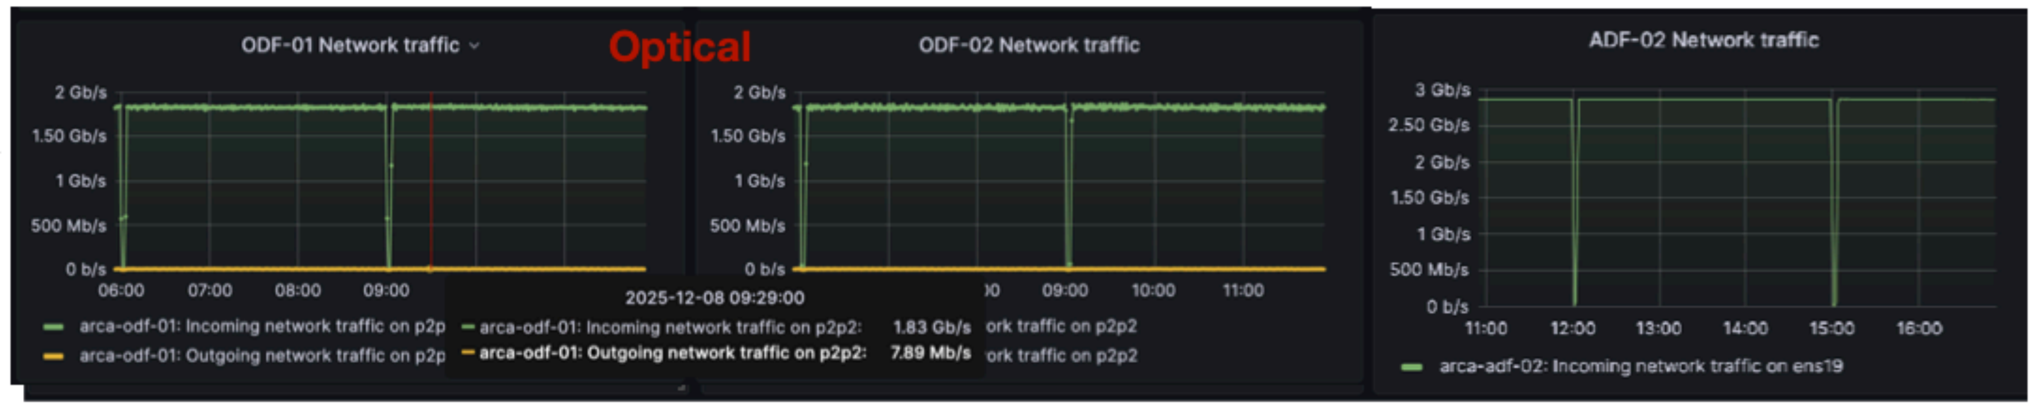

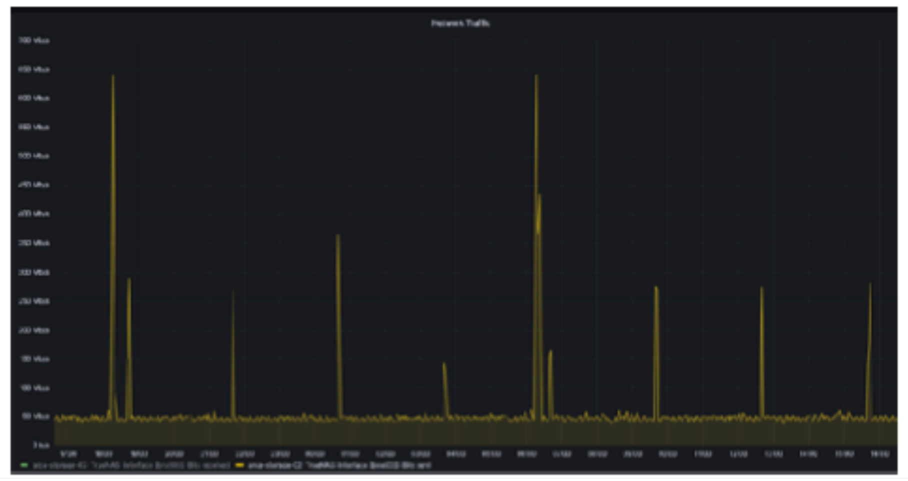

Periodic data transfer to permanent storage to CC-Lyon and @CNAF (Tier 1)
- /3h DAQ
- /12h Reco Online

M.Battaglieri - INFN/
T.Gunji TokyoU

→ **The strawman view: an optical cavity which has to stay locked**

Introduction
DAQ architecture
Calibration and control
Online data processing

S.Viret (IP2I Lyon)
On behalf of the Virgo collaboration
(s.viret@ip2i.in2p3.fr)

PART 1
⇒ Hard-real time part (***down to 10kHz***): front-end, control loops

⇒ Interfaced with slow control (*lock acquisition, monitoring and operation*)

⇒ Full info stored permanently in raw data stream (***~60Mb/s***)
PART 2

⇒ h(t) and quality flag ***(<0.1Mb.s)*** are then distributed by the ITFs to computing centers for low latency analysis.
PART 3

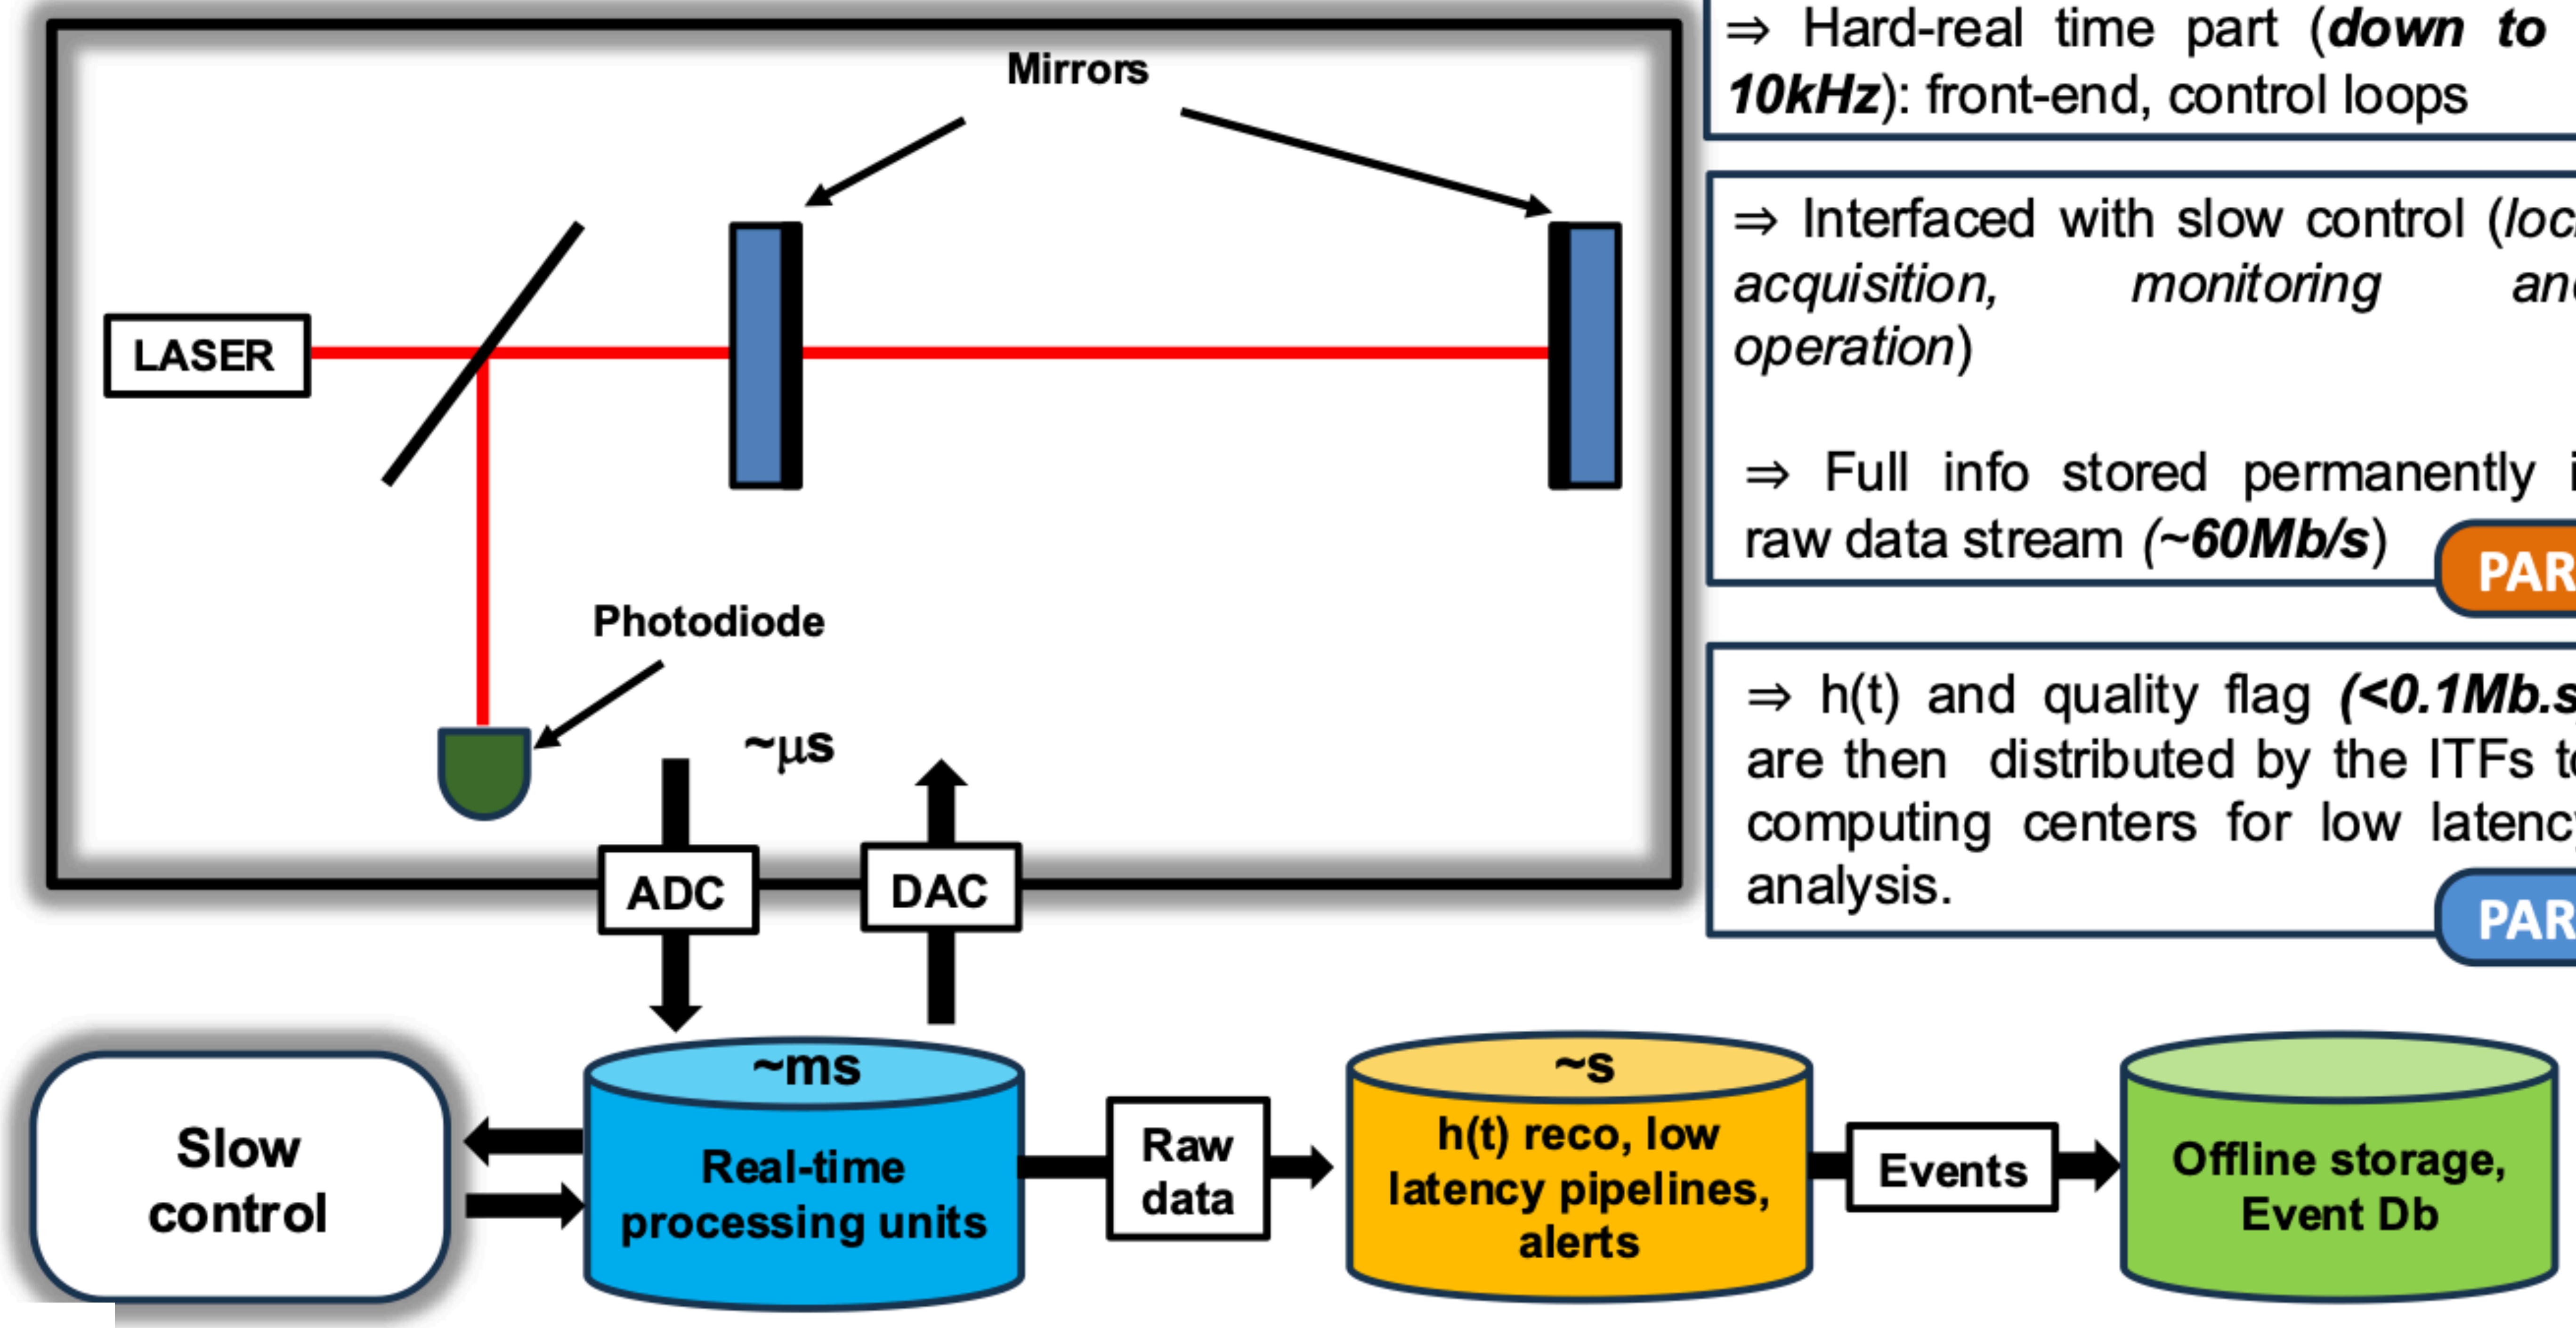

⇒ **Virgo is triggerless by design**. The deformation **ΔL/L=h(t)** (*aka strain*) is continuously recorded as soon as the ITF is locked.

⇒ One 4kHz signal per ITF… Looks simple on paper, but:

- Complex inter-calibration of the system is constantly needed
- Strain should be shared online with other ITFs
- Low latency data processing is required to alert other facilities

- **DSP boards** (***INFN***) for real-time control of the main mirrors position.

→ **Online calibration, real time control**

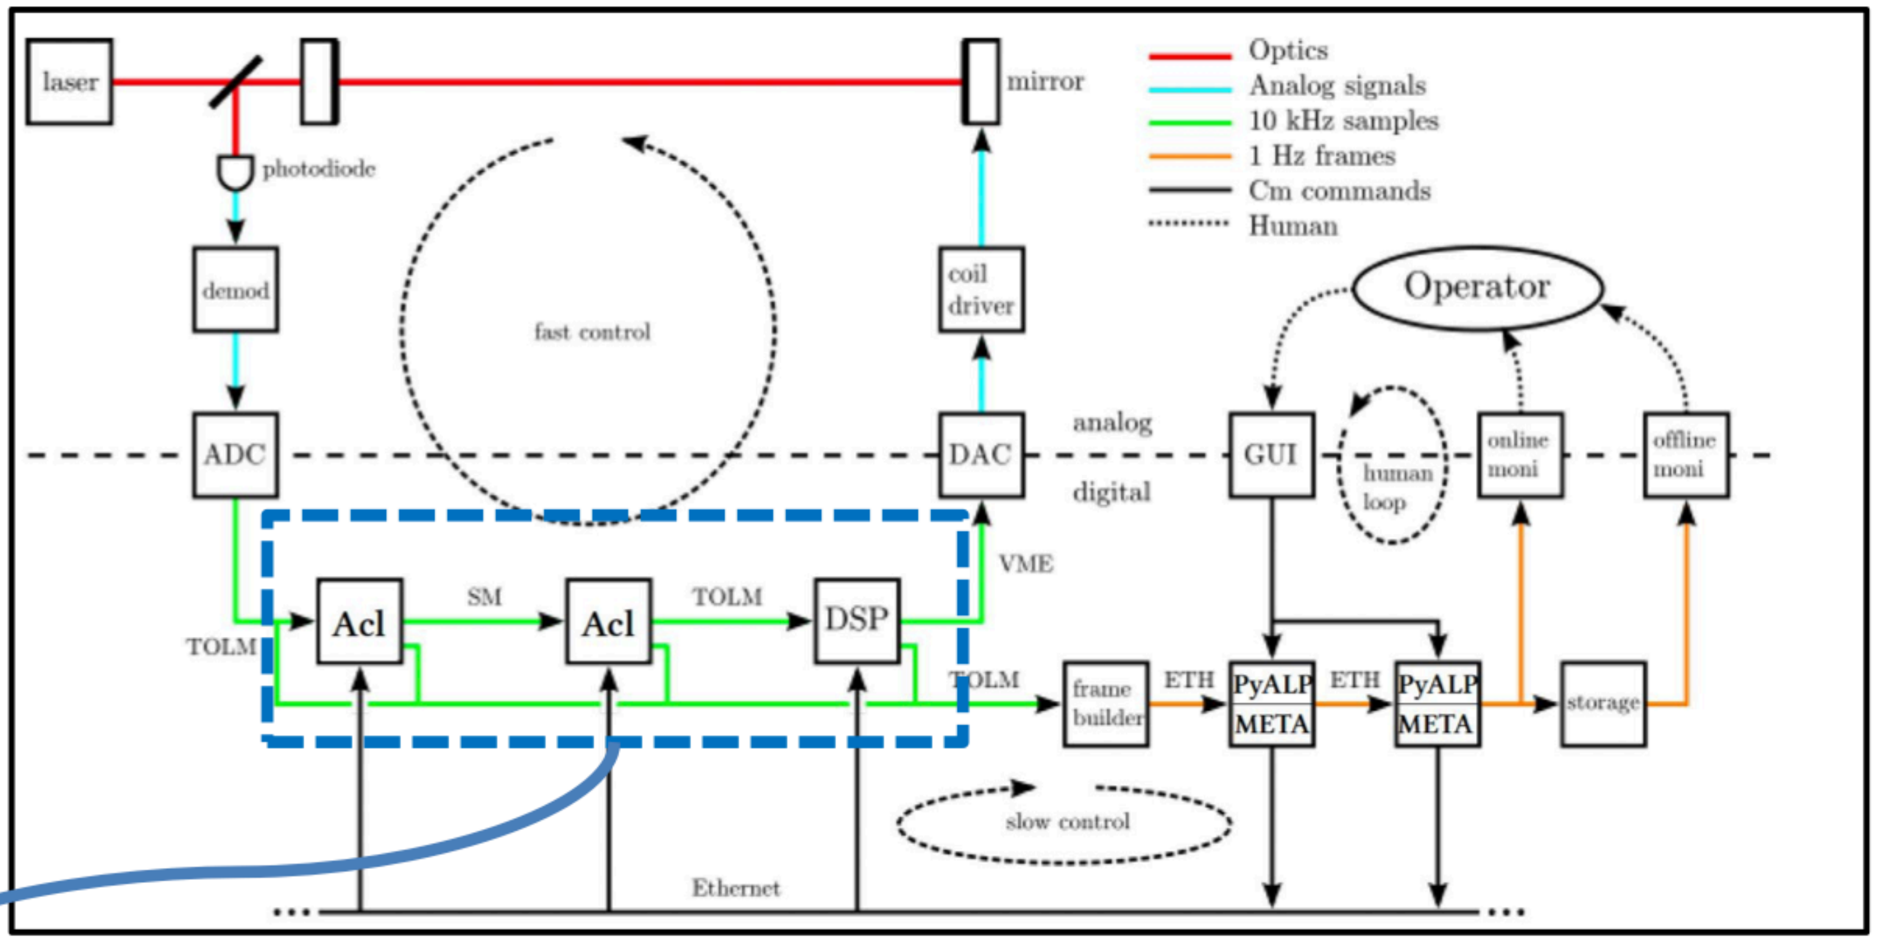

⇒ Real time control and calibration loops are handled either on **RTPCs** or on the **DSP slices** via dedicated custom low-level softwares: **ACL** for RTPCs (*C code*) and **DSPcode** for DSPs (*Assembly code*)

→ **The online analysis flow: low latency steps**

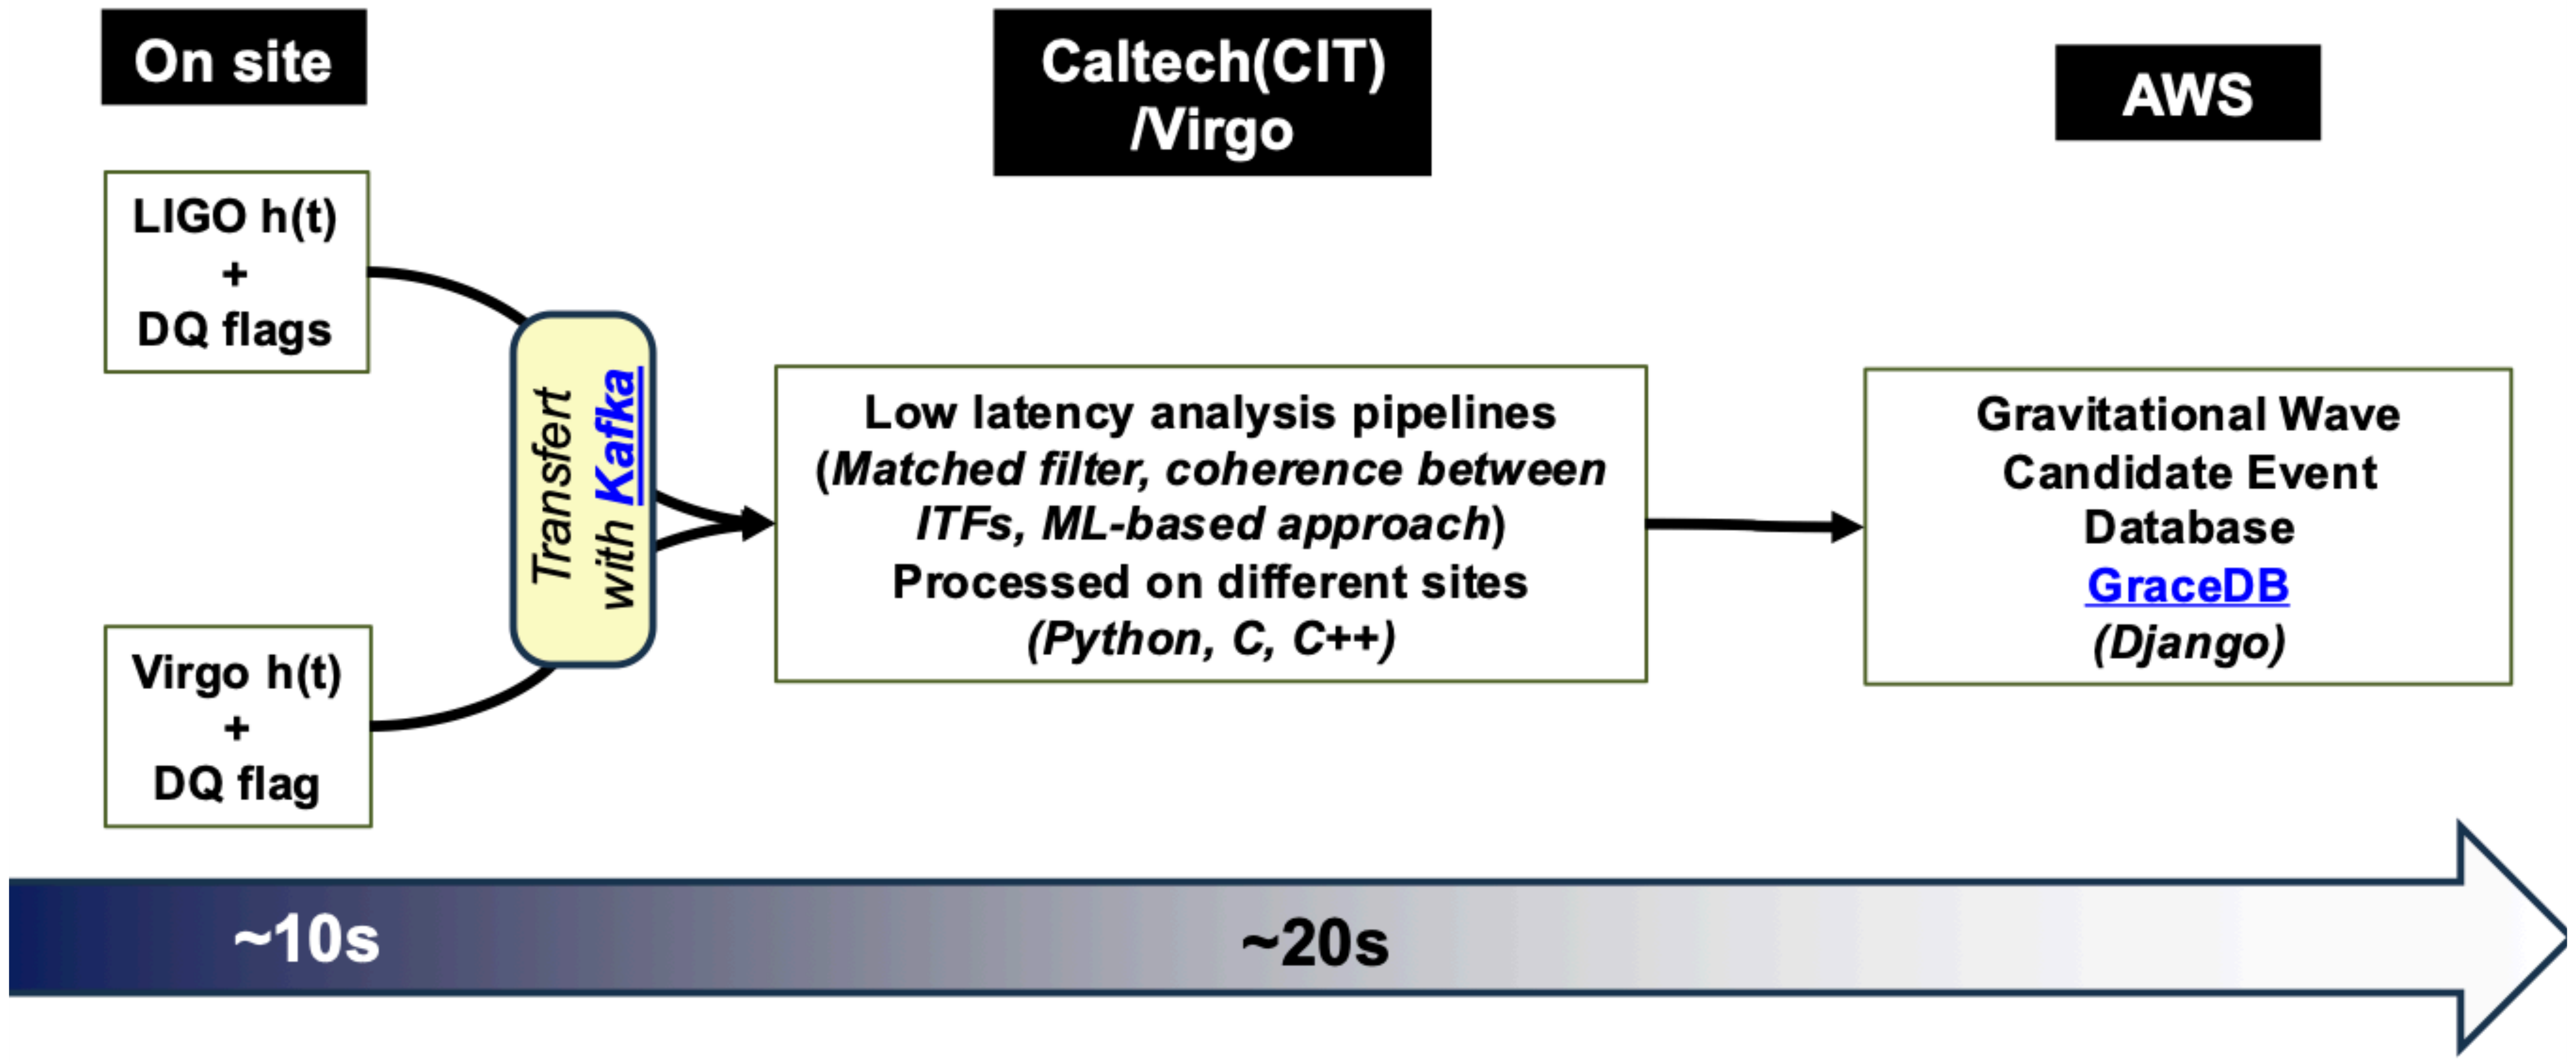

⇒ Less than 30s between an interesting event and the publication of the candidate on the DB. Low latency pipelines are highly parallelized in order to meet those requirements. CPU-based for the moment.

⇒ Reducing this time and possibly transfering more than h(t) will be a challenge for the next generation of ITFs, which will start to experience GW pileup.

Lesson learned on calibration&alignement from existing experiments using SRO

- Each experiment is unique but all rely on C&A as part of the SRO pipeline
- C&As are essential for a rapid turn around from data to physics
- C&As is a synergic effort between DAQ/electronics (online) and reconstruction (offline)
- Experiments forced to develop real-time calibration in connection with a SRO model did succeed
- Some experiments have a similar multi-Echelons structure that should be studied to get inspired
- Hardware's performance has an impact on the procedures (GPU? FPGA? CPU?)
- The framework shall integrate C&As into both online/offline rec with different level of accuracy/complication
- C&A frameworks and procedures were developed over time, building on failures and success (experience!)
- ePIC is trying to develop a SRO framework which should work from day-0 -> a lot of work should be done in the preparatory phase (detector development, bench tests, experimental campaigns, …)

Developing a calibration & alignment framework for ePIC is feasible but to avoid a long time for converging requires work and coordination among experts in different areas: detector, DAQ/electronics, software developers (rec and sim), physics analysis

2026 priorities: - coordination; - bottom-up use cases;