Functions of a Dynamic Archive
a HEP perspective

PUNCH4NFDI TA5 – WP3 Workshop

DESY-Zeuthen 18/01/2024

A. Redelbach

FIAS Frankfurt Institute for Advanced Studies

Overview of project

In high-energy (particle/heavy ion) physics there seems to be no perfect analogue to a dynamic archive

However, several of its functionalities are also used

Highlighting some existing concepts and projects

Many technical details can be added (if there is more time available)

Connection to document WP3-1

Quoting from concept document

In relation to dynamical filters, workflow execution timescales can broadly be divided into three categories:

• **Immediate**: Workflows running on archive data but continuously and at a speed sufficient to influence real-time sensor operations.

• **Emulation**: Improved versions of the real-time workflows are typically developed based on archived data. The relevant time-scales are thus the estimates for execution times when carried out at a real-time setup.

• **Offline**: Many archive queries are never meant to be used for real-time operations, but rather to perform the final data analysis. Such pure offline studies are instead limited by the total load which the archive can be put on.

Online triggering in HEP

Traditional hierarchical scheme

In (traditional) hierarchical triggering schemes a fixed set of trigger conditions exists.
If the data stream supersedes a predefined (and possibly prescaled) signal threshold, a corresponding logic is 'fired'.

That scheme includes hardware- or software-based triggering and logically combined single trigger selections.

The decision process and workflow is rather static than dynamic, i.e. the set of decisions remains typically stable during a ongoing period of data taking, a so-called run.

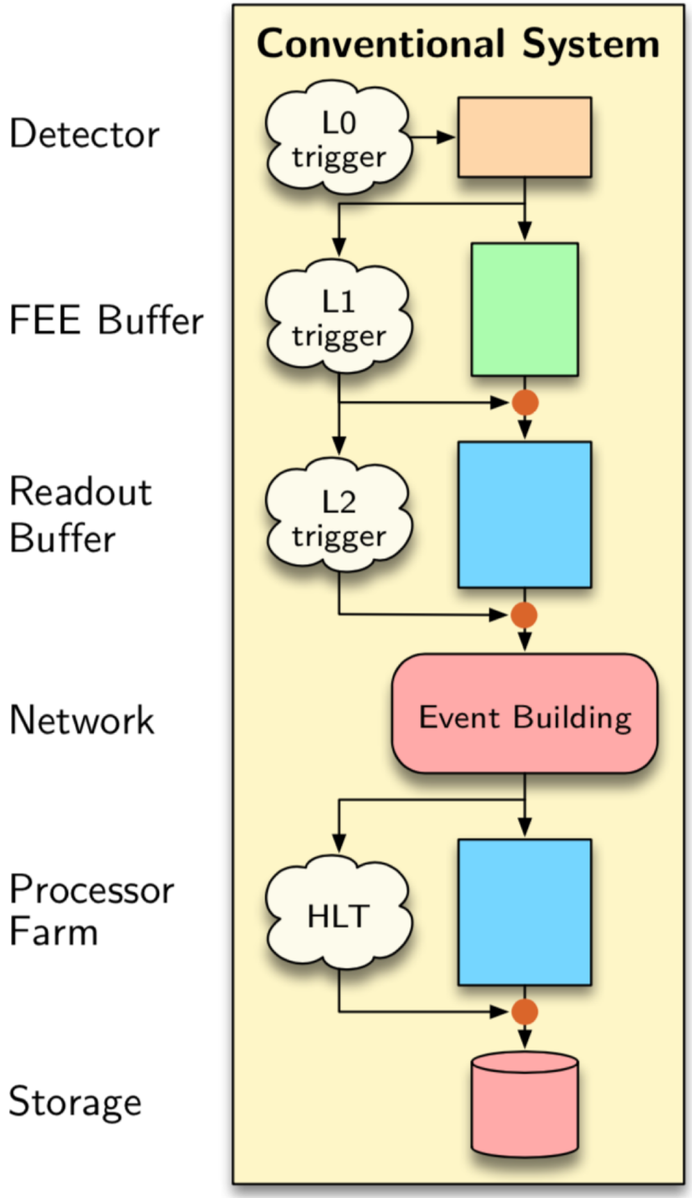

Online triggering in HEP

Towards software-based scheme including calibration feedback loop

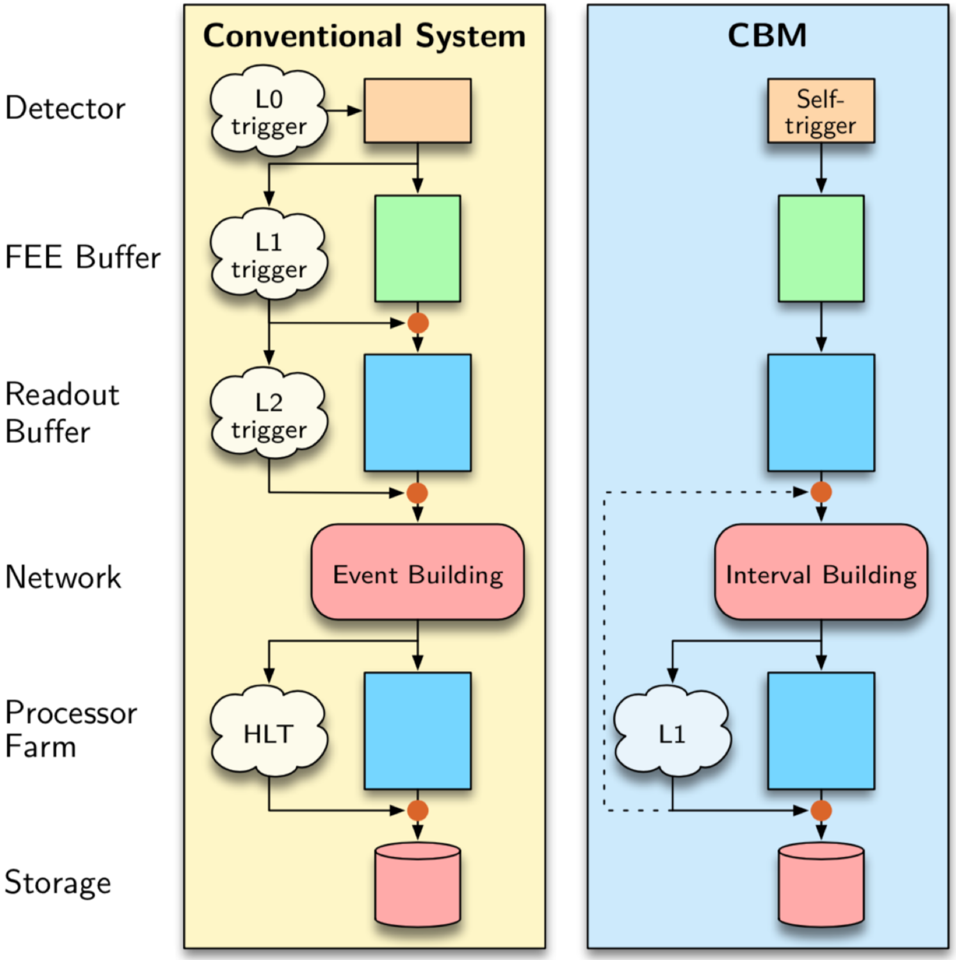

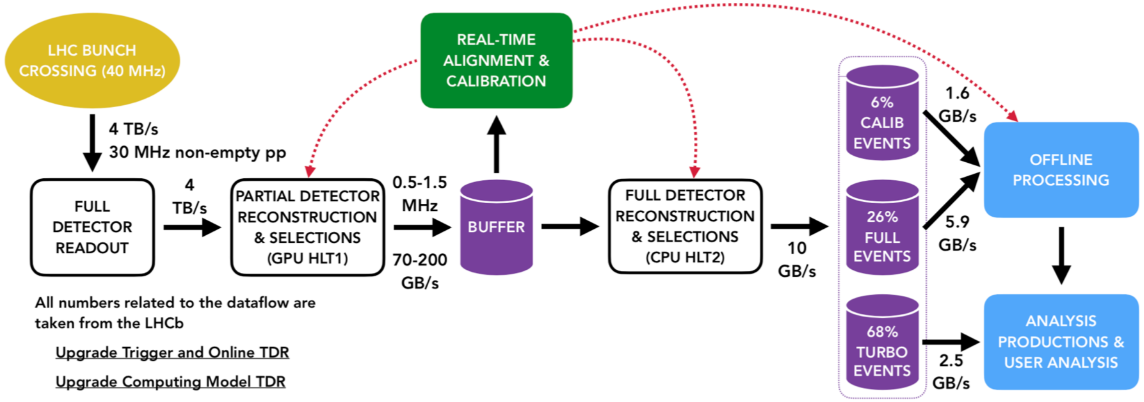

Figure 3: Current data processing pipeline of the LHCb experiment for proton-proton collisions [6, 7]. Arrows indicate data flow, which are annotated with event and data rates.

Testing online/offline workflows: mCBM data challenges

Emulation of scaled workflows

Execute data chain using synthetic data (pre-recorded mCBM data)
Data merged and split again into 80 raw data streams

Monitoring of performance data using InfluxDB/Grafana - subsequent performance analysis

Container operation tricky but successful - binaries built into Docker/Apptainer container

4 nodes sending recorded raw data (timeslice format)

Parallel operation of 10+ processing nodes
Per node: Parallelization using OpenMP -16 cores per process

Increasing replay speed

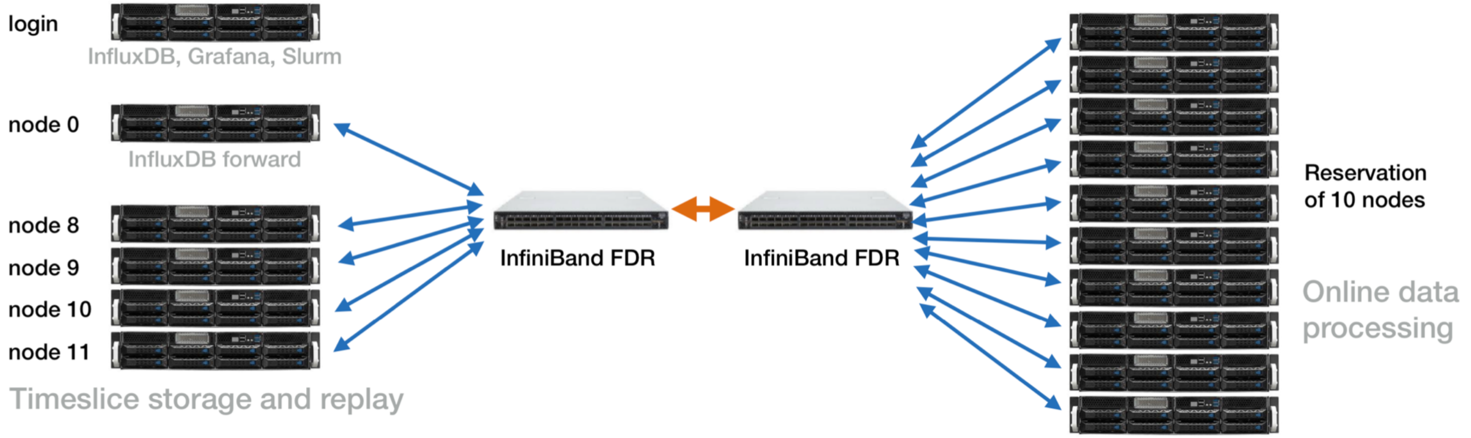

Validation of online/offline filtering

Evaluation of trigger efficiencies (reconstruction of Z → $\mu\,\mu$ peak)

Idea: Efficiency measurements with the **tag and probe method** utilize any known correlation between two separately reconstructed objects, e.g. the tracks from the decay of a known particle.

With stringent selection criteria on the tag object’s side and the correlation to the probe, one can ensure the identity of the probe object without direct, hard cuts on this object.

Well known decay of the Z boson to pair of muons is often used.

1. A muon candidate is selected, using measurements and tight selection criteria „tag muon“.
2. A probe muon is selected (with looser criteria).
3. **Tag and probe muons** are consecutively **combined** and further pair criteria are applied: Opposite charge, origin from the same vertex, their **invariant mass** at the nominal Z boson mass.

→ Example of a well-established HEP method for validation in online and offline workflows

Validation of online/offline filtering

Evaluation of trigger efficiencies (Z → mu mu peak, tag and probe method)

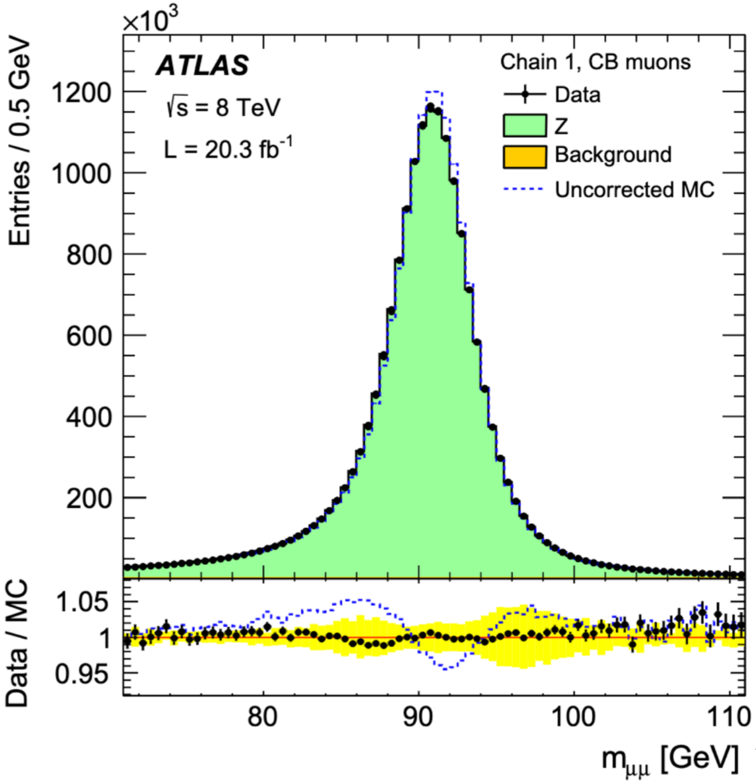

Eur. Phys. J. C (2014) 74:3130

Resolution/width and position of Z mass peak is a measure of the quality of tracking

Comparison with reference distributions

Used in offline quality control and also online quality control

Works also for relatively low numbers of (muon) tracks

Invariant mass(x-axis):
Estimate of mass of decaying particle

Message-based distributed data taking: FairMQ

Data-flow based model: Message Queues based multi-processing
Works locally and across most networks (e.g. Ethernet, InfiniBand, Shared Memory Transport)
The framework does not impose any format on messages.

Actor model:

- Standalone processes ("devices") perform a task (e.g. track finding) and communicate with each other via messages (mediated by a queue).
- No locking, each process runs with full speed
- Easier to scale horizontally to meet computing and throughput demands (start/add new instances)

Building block (FairMQ Devices)

- Device takes/passes ownership of data
- Framework user sees only the callback to his algorithm
- Different channels can use different transport engines

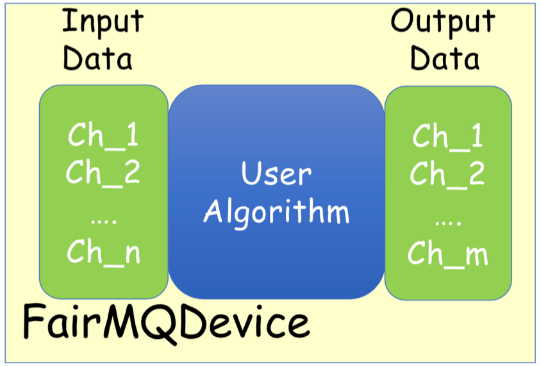

Tested for complex real-time topologies:
Integrated in the online farm directly connected to the ALICE detector

More details: M. Al-Turany [ALFA - A framework for building distributed applications]

Online calibration: ALICE CCDB

Calibration and Conditions Database at a glance

- Central store of calibration and condition data
- Metadata stored separately from the serialized calibration data
- Data distribution using a set of reliable Grid storage elements
- **Millisecond resolution for object Interval of Validity**
- **Expected O(100Hz) updates and O(20kHz) queries**
- HTTP(s) for metadata queries
- HTTP(s) and/or XrootD for data access
- REST API for querying it

Offload to ALICE computing grid:
- Blobs are uploaded to several Grid SEs
- Local disk used as buffer and cache only
- Metadata queries performed locally
- Clients are redirected to read from the Grid SEs
- Bandwidth scales with the number of replicas
- Location-aware sorting of WAN addresses

→ Database replication with native PostgreSQL replication

Repository(version for local development):
https://gitlab.cern.ch/grigoras/ccdb-local

More details:
Slides of C. Grigoras at TA5 meeting

References:
https://journals.agh.edu.pl/csci/article/view/4831

CHEP contribution

Online calibration: ALICE CCDB

Continuous synchronization of “online” and “offline” databases

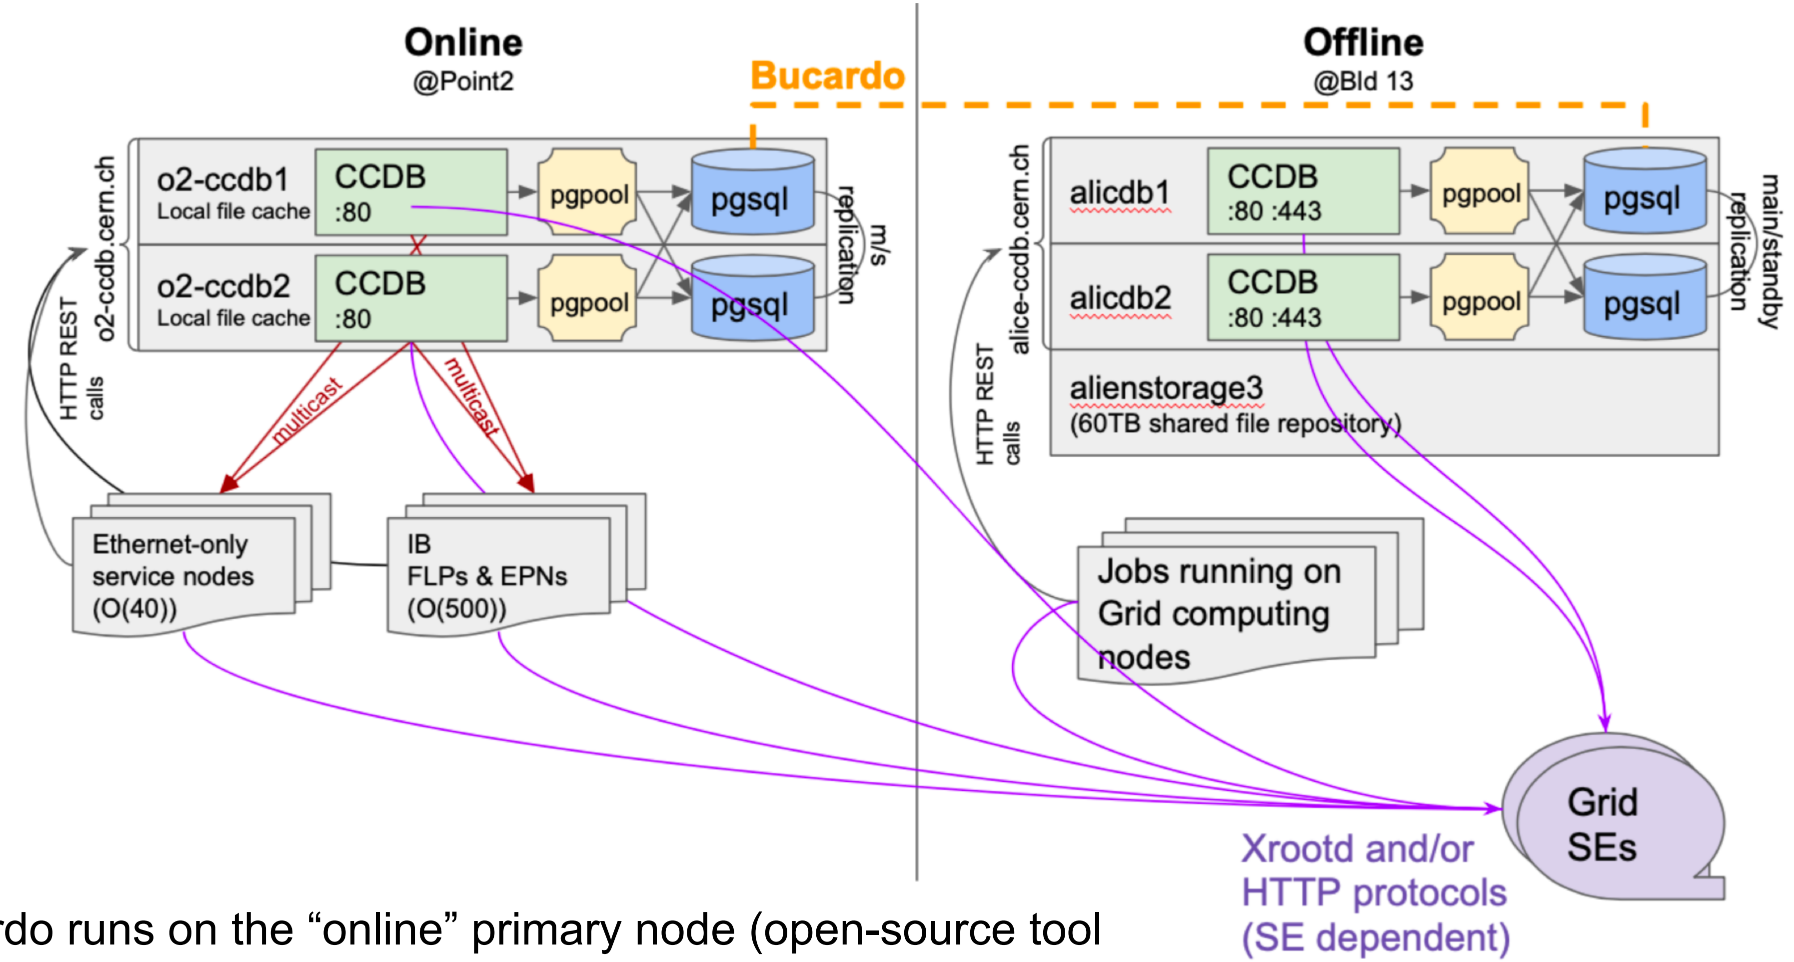

Bucardo runs on the “online” primary node (open-source tool supporting multi-primary replication for PostgreSQL)

PostgreSQL databases store metadata about condition files.

Anomaly detection: Monitoring and new physics (WP5)

Anomaly detection is a typical area of application of neural networks:
detect events which are different from regular event patterns

Anomalies may appear due to
- signals not foreseen in usual physics models → **new physics**
- detector behavior → **predictive maintenance**

Irregular detector data:
- dead or noisy channel, including correlated noise
- deviation from regular detector performance, e.g. excess or drift in alignment, calibration, ...

TA5 – WP5

WP5-1 (Q3/2024): Development of ML prototypes for anomaly detection and predictive maintenance
WP5-2 (Q3/2024): Interference recognition and mitigation schemes for transient discovery leading to a robust triggering system
WP5-3 (Q3/2026): Expansion of the concept to a generalized toolkit for predictive maintenance and anomaly detection
WP5-4 (Q3/ 2026): Evaluation of the ML approaches by analyzing false-alarm rates and online feedback

https://indico.desy.de/event/39436/contributions/144012/attachments/82644/108976/ML_FPGA_ParticlePhysics_2023_AStraessner.pdf

Towards anomaly-based triggering for new physics

Very active research field in HEP community

Ephemeral Learning -- Augmenting Triggers with Online-Trained Normalizing Flows
https://arxiv.org/abs/2202.09375

Method compresses an entire data set by learning a generative model for events in terms of network parameters

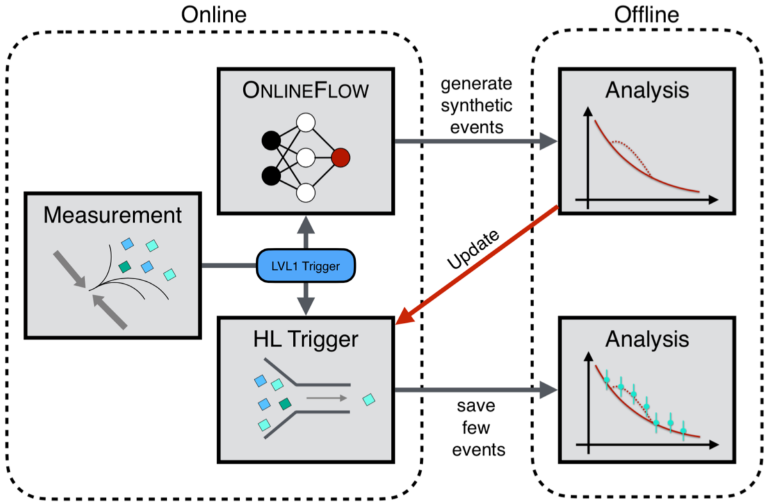

„If necessary, we adjust the trigger to take new data accordingly (online) and analyze that data (offline).“

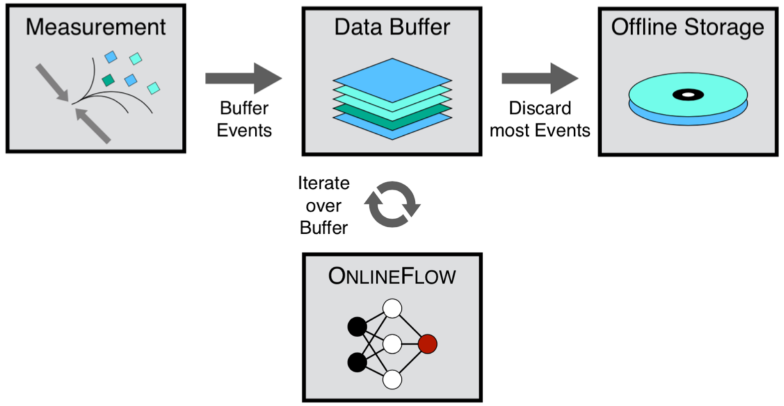

„ONLINEFLOW (bottom center) is trained online on the data in the buffer and iterated until the buffer is replaced.“

Further reading:

LHC Olympics 2020: A Community Challenge for Anomaly Detection in High Energy Physics
https://arxiv.org/abs/2101.08320

A Living Review of Machine Learning for Particle Physics
https://iml-wg.github.io/HEPML-LivingReview/

Also activities in PUNCH-TA3: Anna Hallin - Anomaly detection with ML in HEP

Continual learning approach

- **Concept:** Train a model with a continuous stream of data
- Learns from a sequence of partial experiences rather than all the data at once

→ Adapts (by desing) to a changing data stream
→ No need to store (all) previous training data
→ For supervised learning need a ~continual stream of labelled data which might not be accessible
→ Fast ML settings this make deployment on FPGA difficult

Encouraging results from C. Brown et al.: Studies for ML in CMS Phase-2 L1T demonstrating that Continuous Learning has a larger impact on the robustness of the ML model

C. Brown for CMS at CHEP 2023

Avalanche framework
https://arxiv.org/abs/2104.00405

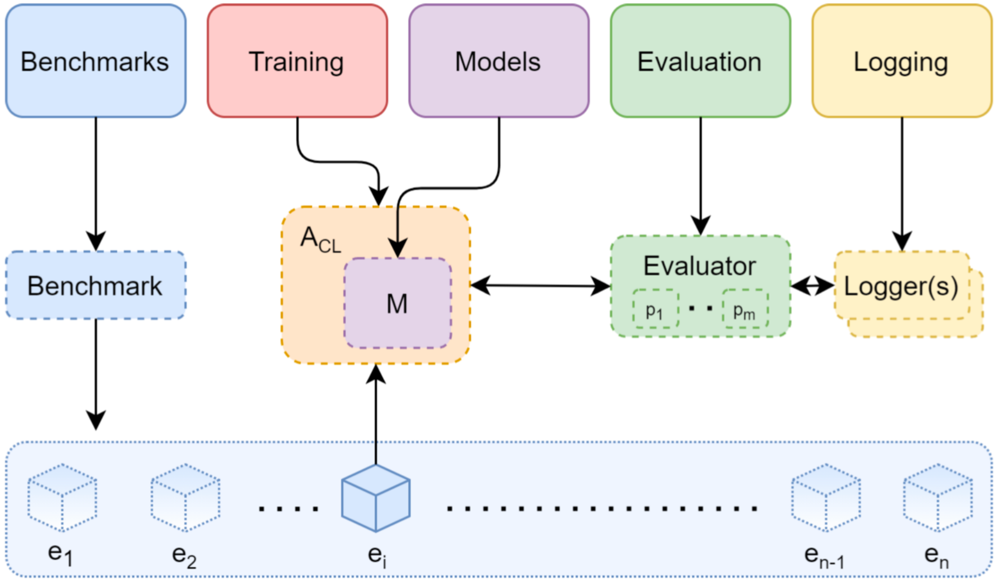

A *Benchmark* generates a stream of experiences e_i which are sequentially accessible by the continual learning algorithm A_{CL} with its internal model M.

Possible next steps

Highlighting some interesting projects and frameworks in HEP

Invitation of experts for more detailed discussions
→ Discussions at the end of the workshop

Input for brainstorming session

Forming of teams working on sub-topics tomorrow