The O^2 project

A new computing system for the ALICE Experiment

Vasco Barroso
CERN

Outline

- Introduction
  - ALICE
  - O^2 Project
- O^2 Farm infrastructure
- O^2 Quality Control
- O^2 Physics data reconstruction
- Conclusion

Introduction

ALICE: A Large Ion Collider Experiment

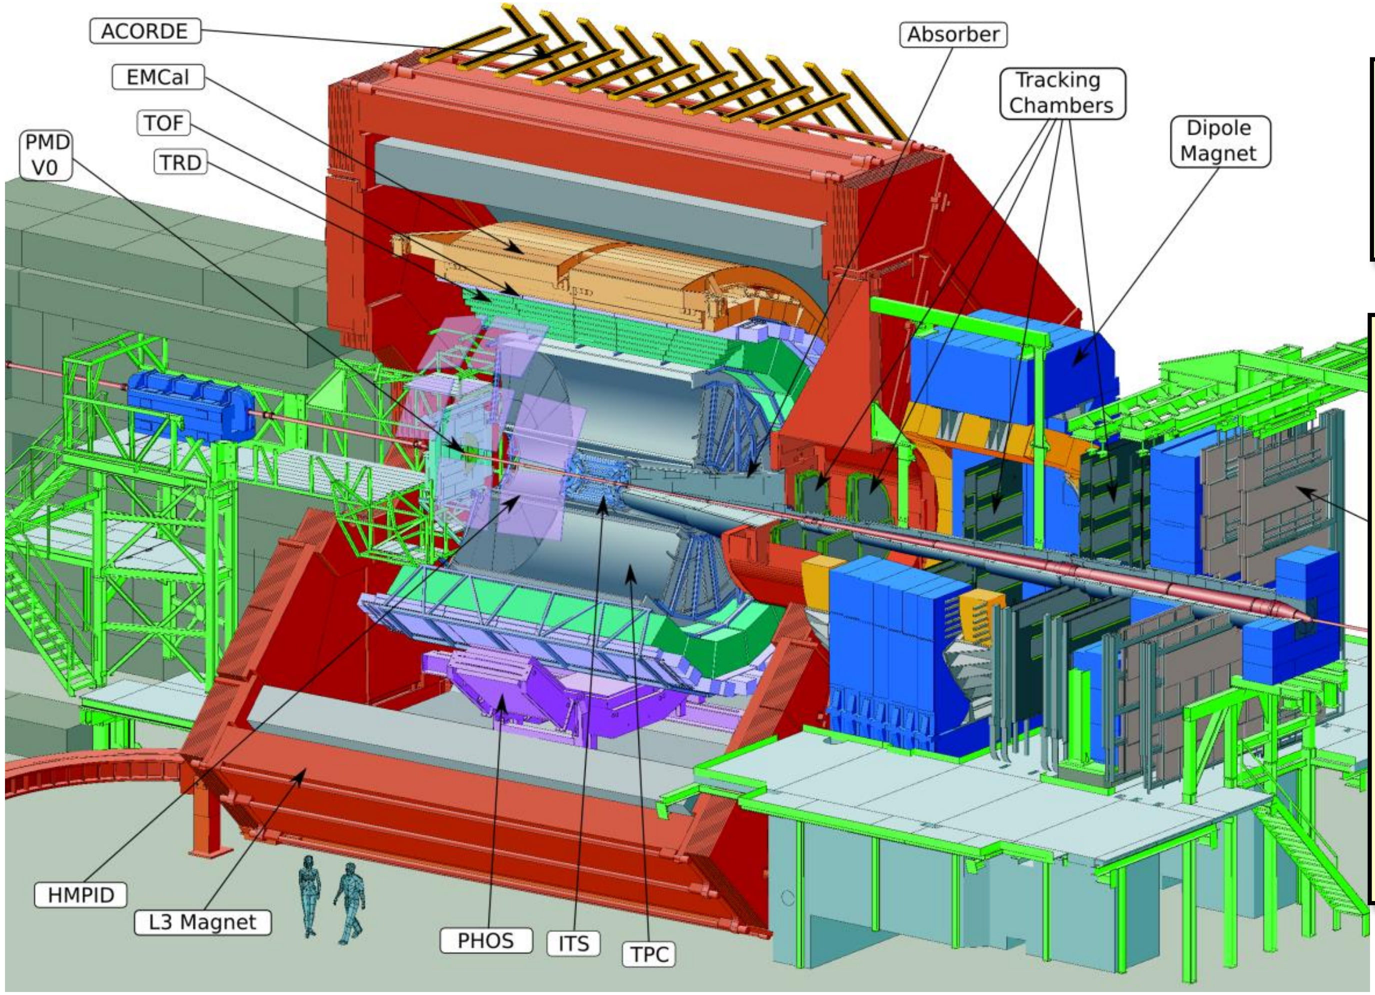

Detector:
18 technologies
Size: 16 x 26 meters
Weight: 10,000 tons

Collaboration (Jan '18):
1800 Members
176 Institutes
41 countries

A brief history of ALICE

1990-1996: Design
1995: Technical Proposal
1992-2006: R&D
2004: Computing Technical Design Report
2000-2009: Construction, Installation, Commissioning
2010-2013: Run 1 (Operation)
2013-2014: LS1 (Long Shutdown 1)
2015-2018: Run 2 (Operation)
2019-2020: LS2: major ALICE Upgrade
2021-2023: Run 3 (Operation)
2024-2026: LS3
2026-**2029**: Run 4 (Operation)
Project lifetime of 40 years

Introduction

The O² project in a nutshell

Requirements

1. LHC min bias Pb-Pb at 50 kHz
   ~100 x more data than in 2010
2. Rare physics processes with very small signal over background ratio
3. Triggering techniques very inefficient if not impossible
4. 50 kHz > TPC inherent rate
   Support for continuous read-out

New computing system

➔ Read-out the data of all interactions

➔ Compress these data intelligently by online reconstruction

➔ One common online-offline computing system: O²

Unmodified raw data of all interactions shipped from detector to online farm in triggerless continuous mode

HI run 3.4 TByte/s

Baseline correction and zero suppression
Data volume reduction by cluster finder.
No event discarded.
Average compression factor 6.6

500 GByte/s

Data volume reduction by online tracking.
Only reconstructed data to data storage.
Average compression factor 5.5

100 GByte/s

Data Storage
1 year of compressed data

20 GByte/s

Tier 0, Tiers 1 and Analysis Facilities

Asynchronous (hours) event reconstruction with final calibration

Detector Electronics
9000 GBTs links

270 First-Level Processors
Hw acc: FPGAs

Switching Network

1500 Event Processing nodes
Hw acc: GPUs

Switching Network

Write 120 GB/s
Read 75 GB/s
Capacity: 60 PB

Introduction

The O^2 project hardware facility

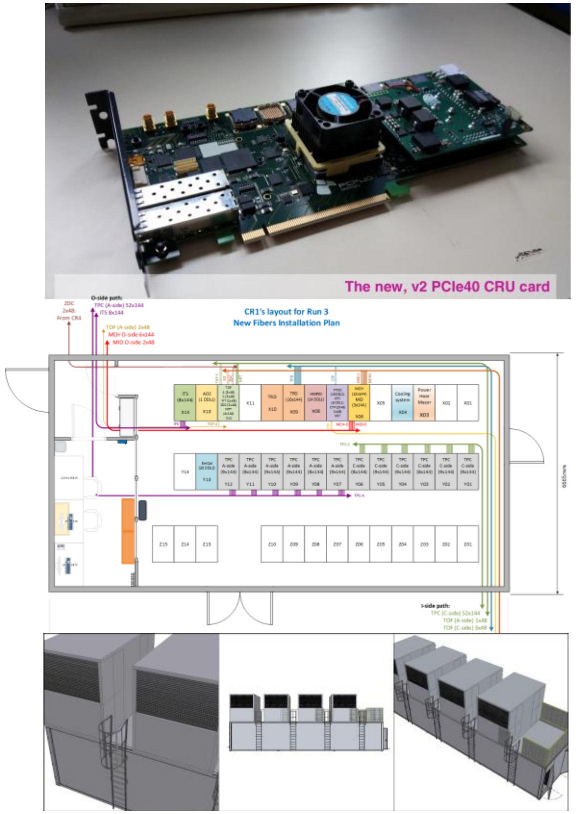

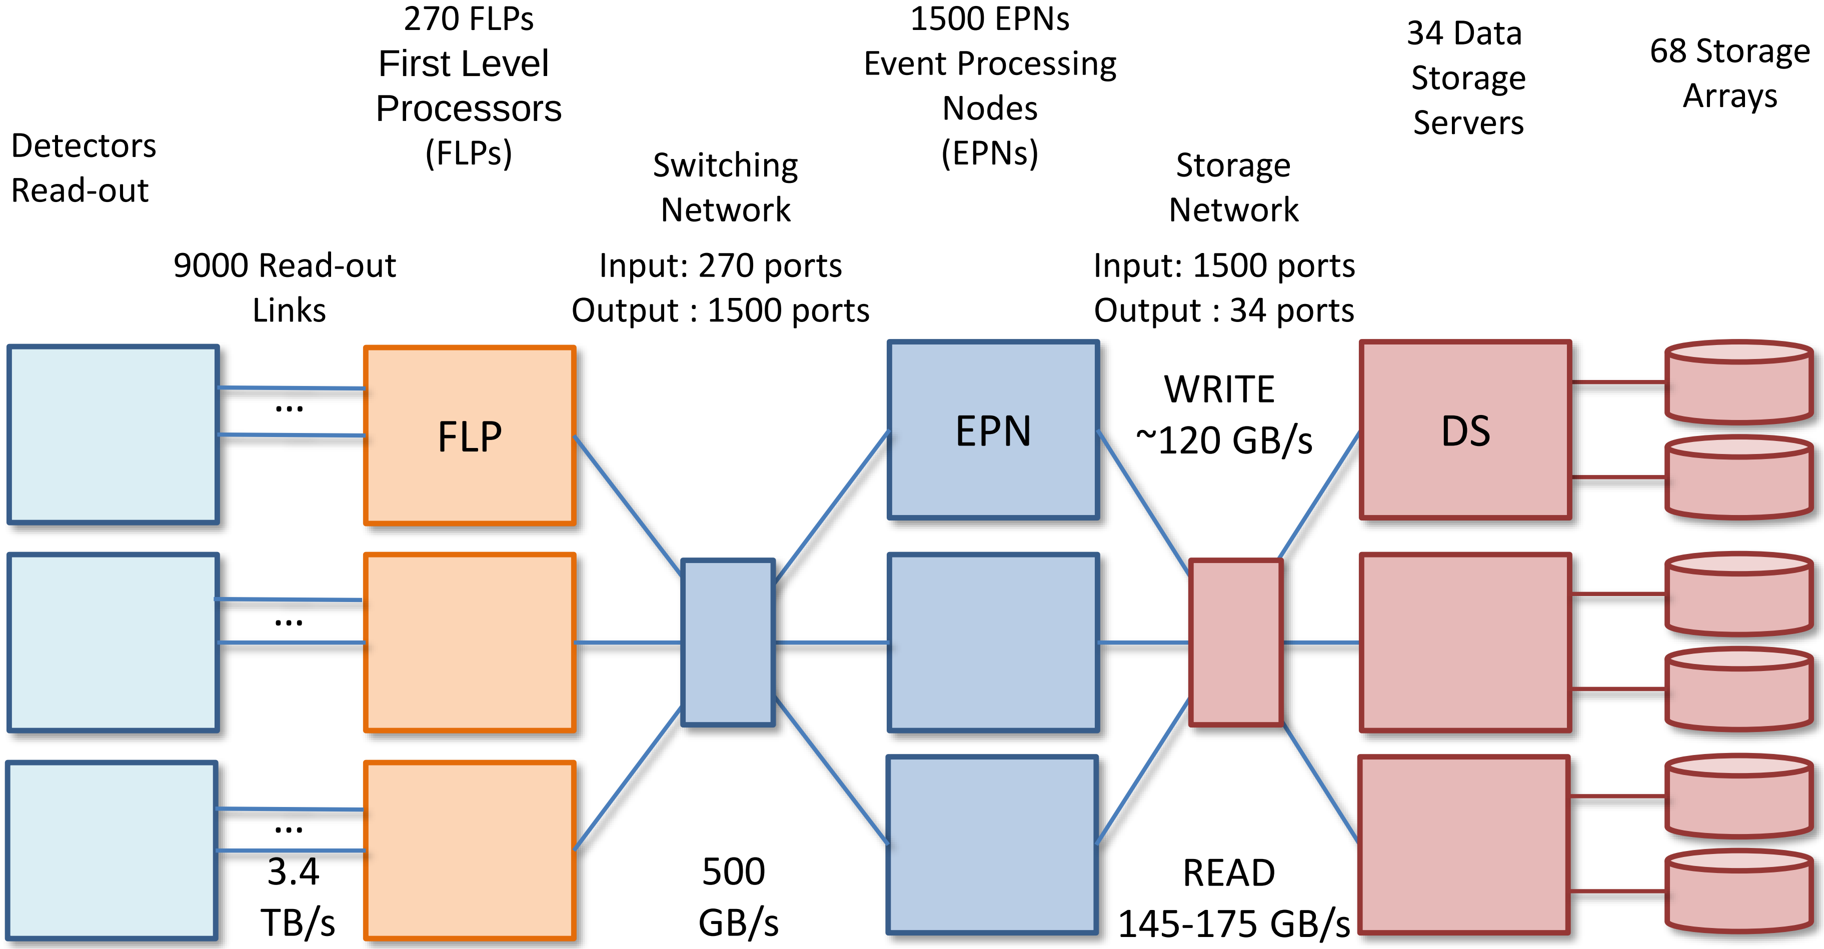

Introduction

The O^2 project timeline

2012	2015	2018	2019	2020	2021
• ALICE Upgrade Letter of Intent	• O^2 Technical Design Report • Software development starts	• First tests with detectors • O^2 dataflow integration	• Detector full-scale tests • O^2 global integration	• O^2 and detector commissioning • Deployment in experimental area	• LHC restart

Introduction

The O² project software architecture and project organization

- Message-based multi-processing
  - Microservices
  - 100k processes
- 17 thematic working groups
  - ALFA : ALICE and FAIR
- Embrace widely-used tools
  - Technology refresh

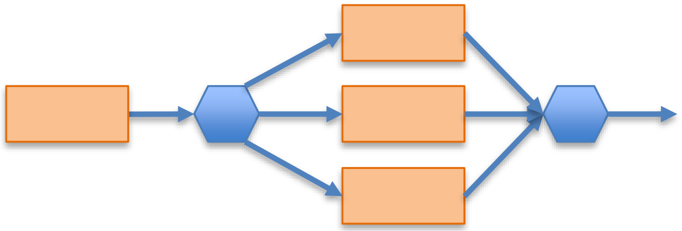

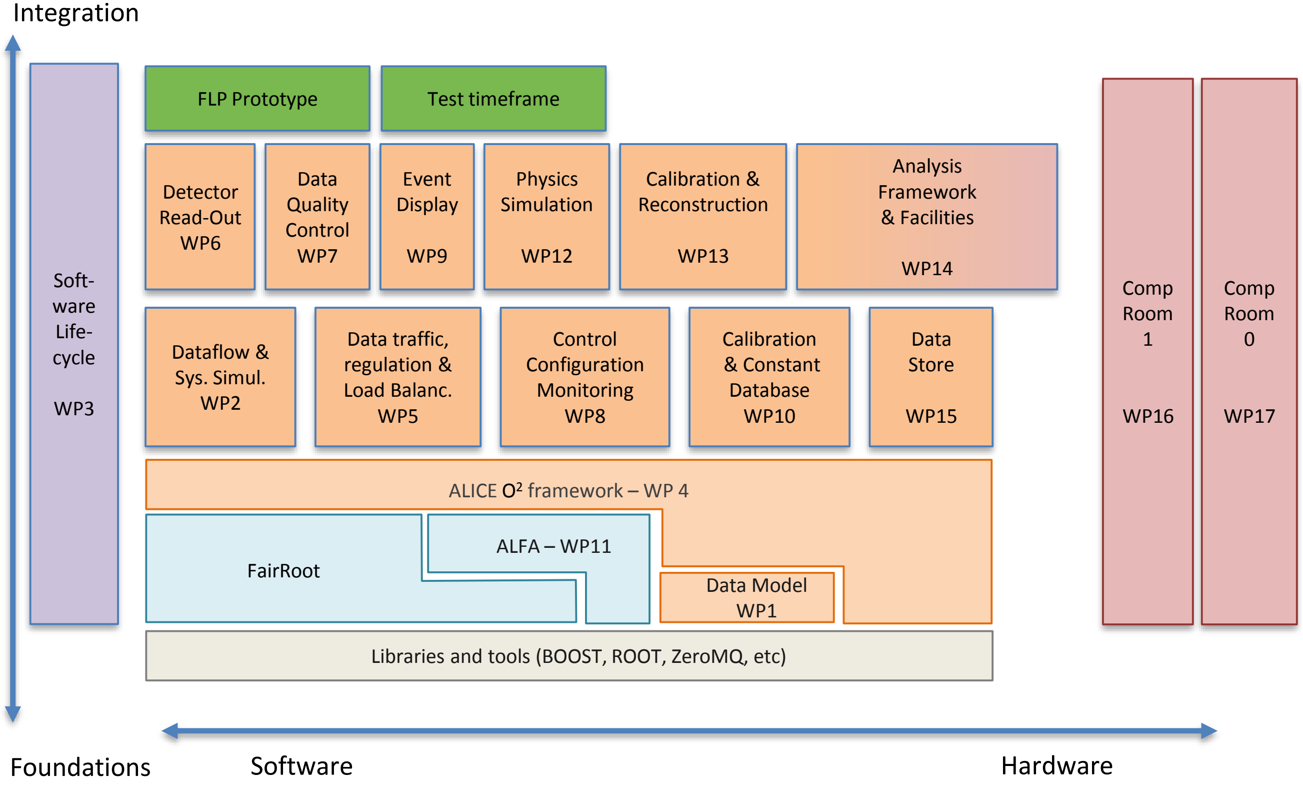

O^2 farm infrastructure

Monitoring

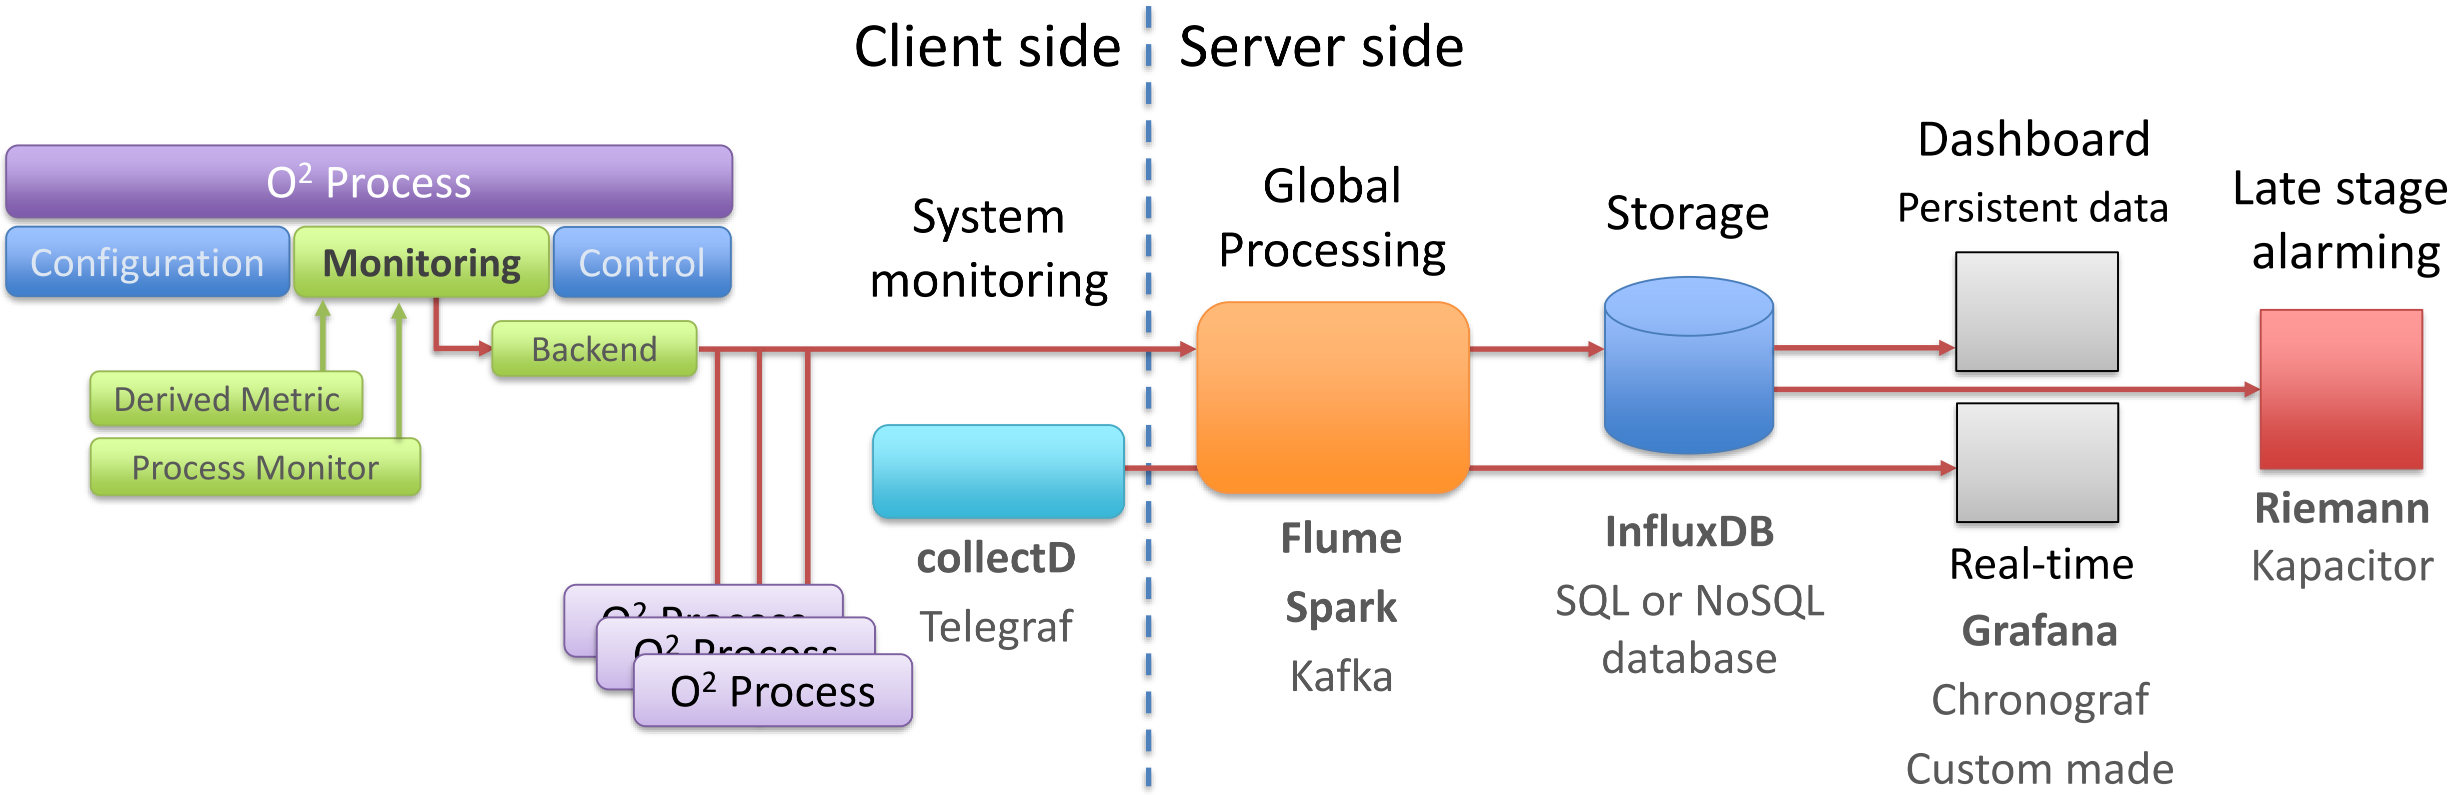

O^2 farm infrastructure

Monitoring

1. collectd
   - Collects system performance metrics
2. Flume
   - Moves metrics from multiple sources into centralized store
   - Basic processing
3. Apache Spark
   - In memory data processing
4. InfluxDB
   - Time series database
   - Continuous queries
   - Retention policy
5. Grafana
   - Time series visualization tool
6. RIEMANN
   - Alarming

O^2 farm infrastructure

Control and Configuration

- Resource Management
- Scheduling
- Deployment
- Process Configuration
- Service Discovery

- Integration of ALICE subsystems
- Automation of operations

O^2 farm infrastructure

Automation

- In the pipeline
  - React to LHC operations
  - Horizontal scaling with luminosity
  - Unresponsive node replacement

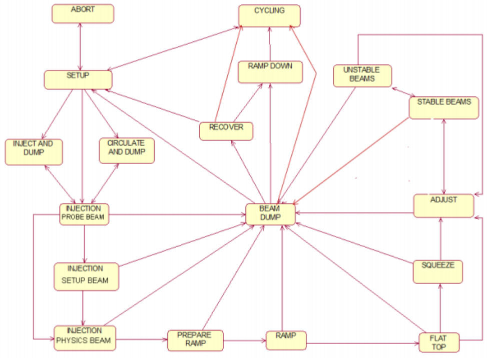

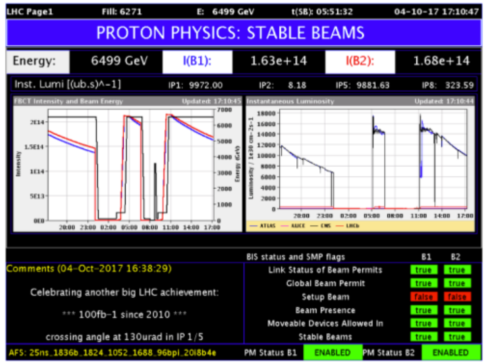

- Being explored
  - Predictive maintenance
    - Log analytics
    - Deep learning

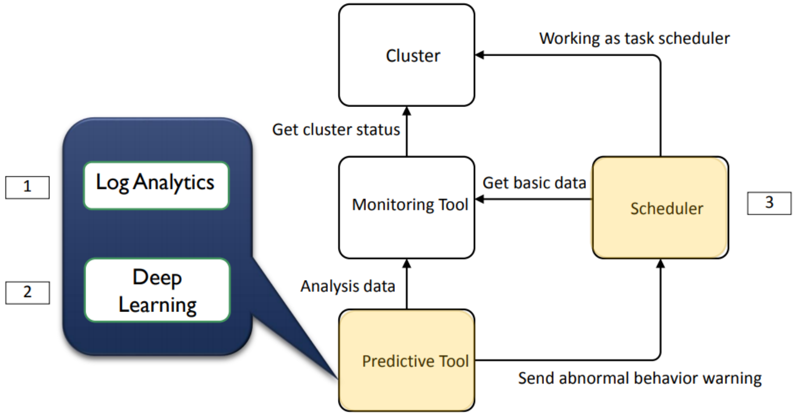

KMUTT, Thailand

Tiranee Achalakul
Phond Phunchongharn

O^2 Quality Control

Feedback on the quality of the physics data

- Quality Control
  - Online (quasi real-time, synchronous with data taking)
    - Make sure to record high quality data
    - Identify and solve problem(s) early
  - Offline (asynchronous)
    - Make sure to analyse high quality data
- 25000 plots generated per minute
- How to assess their quality ?
  - Humans -> ~100 plots
  - Expert written algorithms -> ~5000
  - Machine Learning techniques

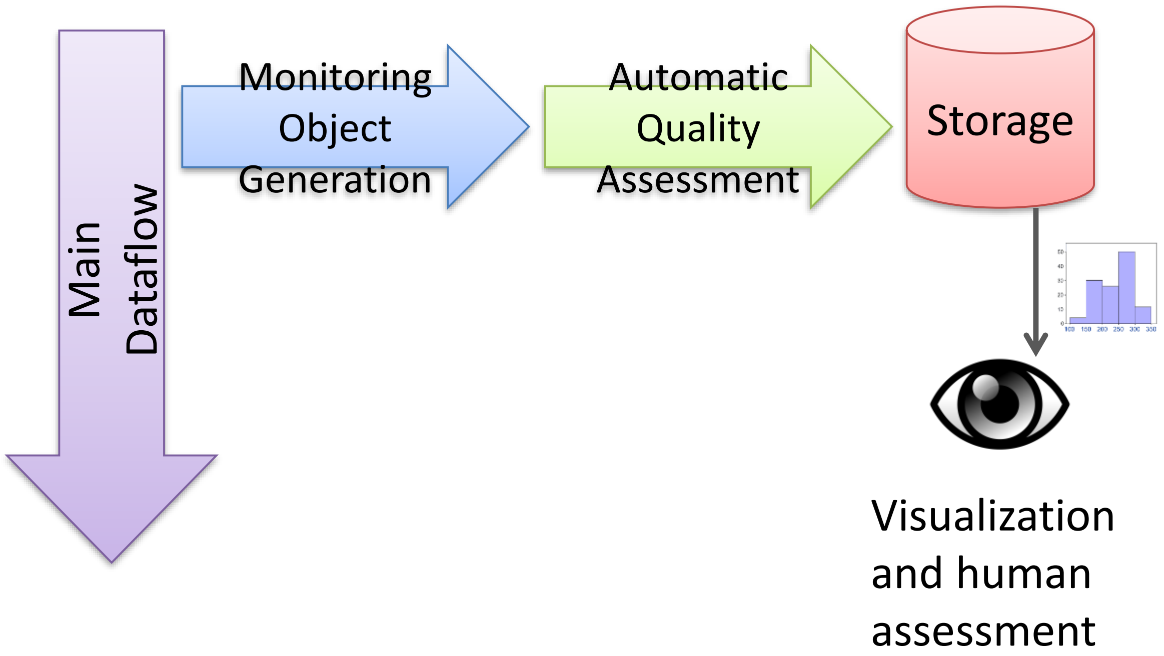

O^2 Quality Control

Assigning quality labels using Machine Learning algorithms

- Supervised learning
- Automatic assignment of quality labels
- Self-improvement, adapt to new conditions
- Several methods tested on existing data
  - Most promising is the *Random Forest* combined with *adaptive boosting*

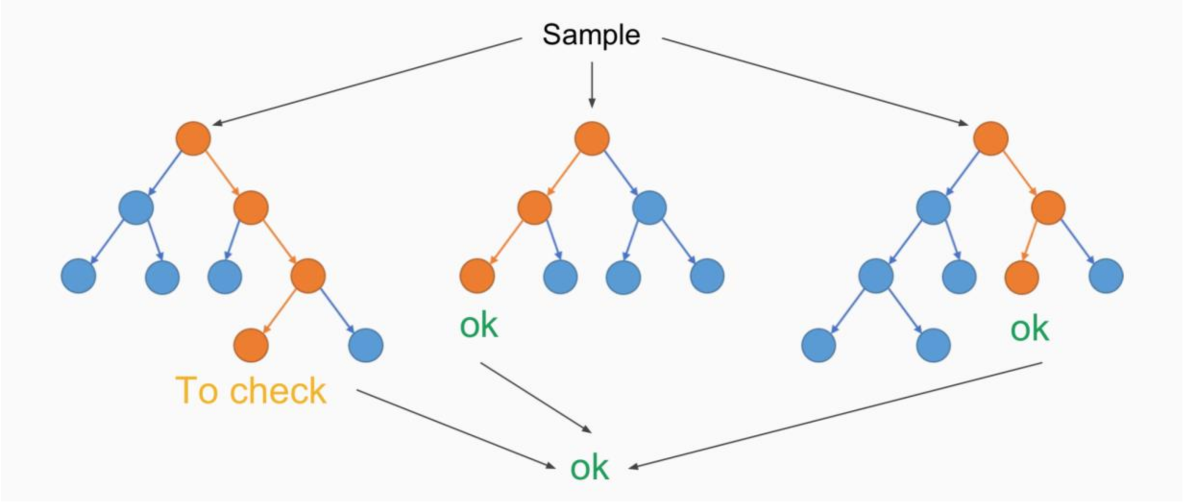

- Results with old data sets
  - Predicts the quality label in 75% of the cases with over 95% accuracy

O^2 Quality Control

Anomaly detection using Generative Adversarial Networks (GANs)

- Unsupervised algorithm to detect anomalies in the QC histograms
- GANs trained on healthy/good data (A)
- Good distributions of energy generated in real time (B)
- Anomalies defined as a difference between GAN-generated and real data points (C)

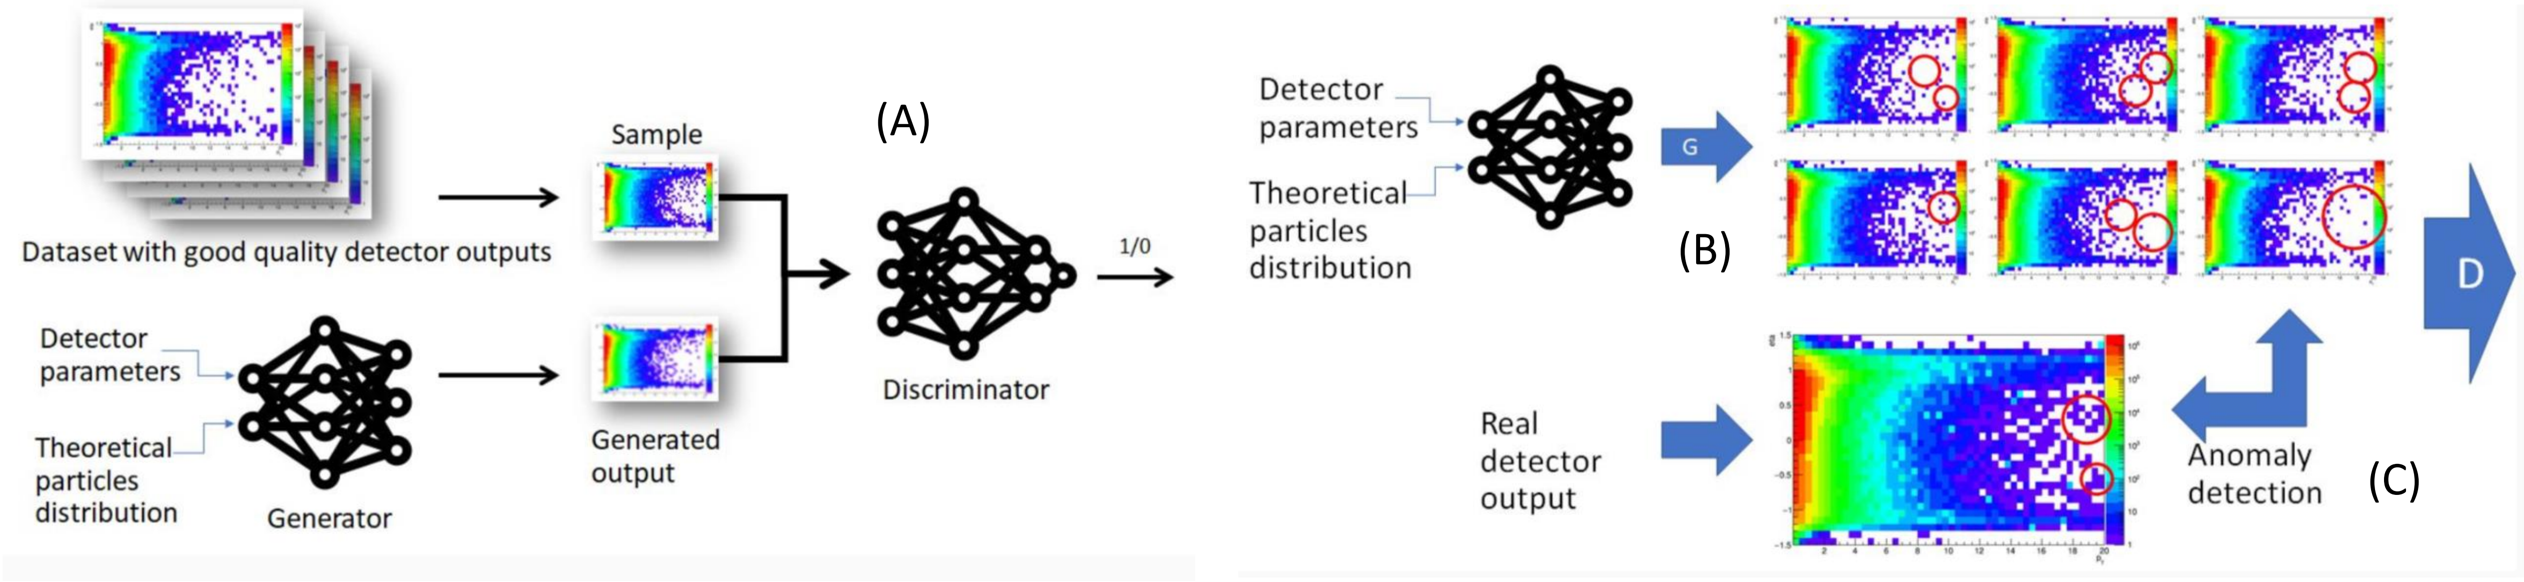

O^2 Physics Data Reconstruction

Clustering and Tracking with GPUs and machine learning

- The normal reconstruction chain is:
  - Clustering, tracking, vertexing, some higher level reconstruction, and then analysis (by and large).

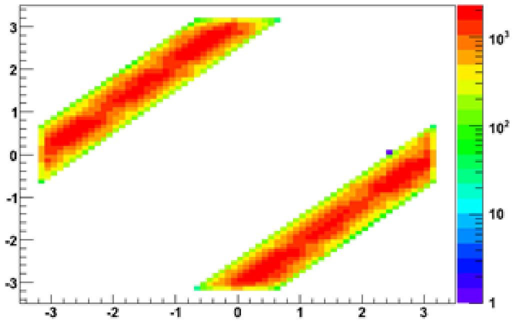

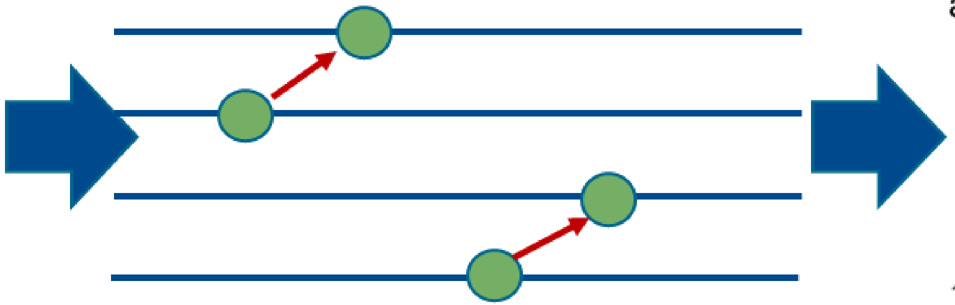

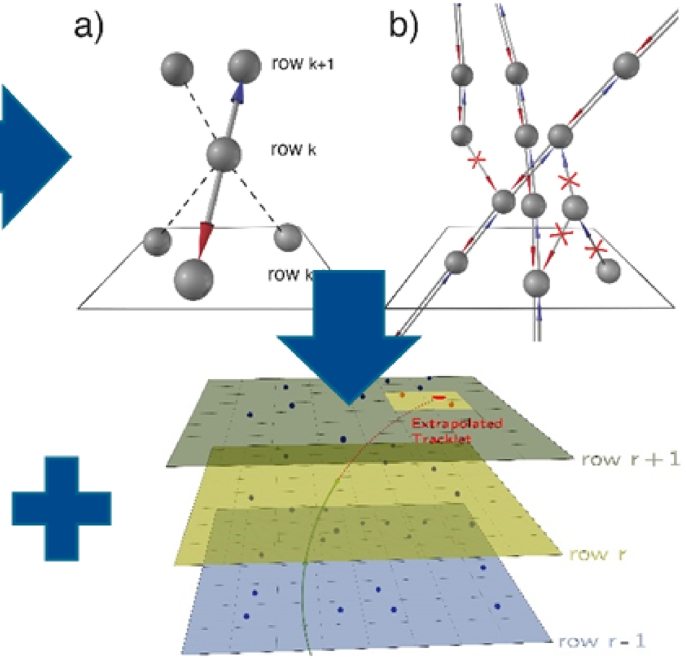

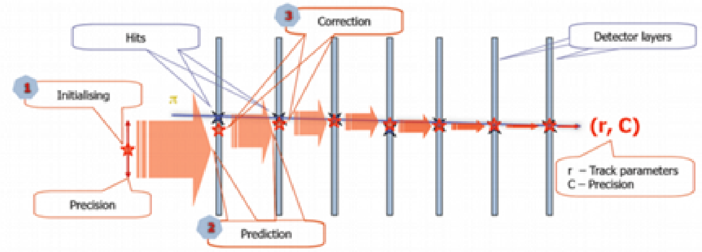

CERN R&D proposal

David Rohr (FIAS)
Sandro Wenzel (CERN)
Giulio Eulisse (CERN)

O^2 Physics Data Reconstruction

Clustering and Tracking with GPUs and machine learning

- The normal reconstruction chain is:
  - Clustering, tracking, vertexing, some higher level reconstruction, and then analysis (by and large).

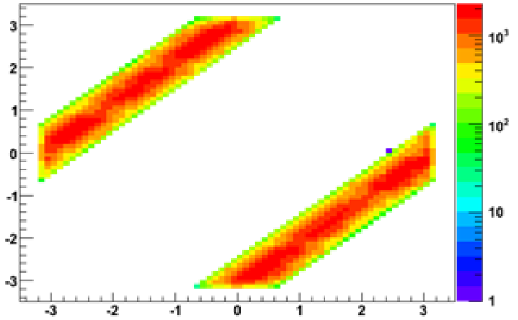

Why don't we use the charge directly?

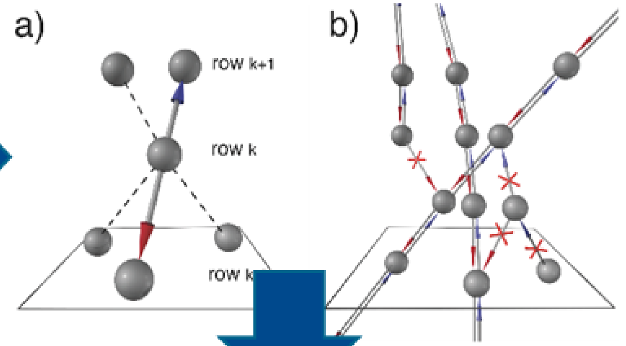

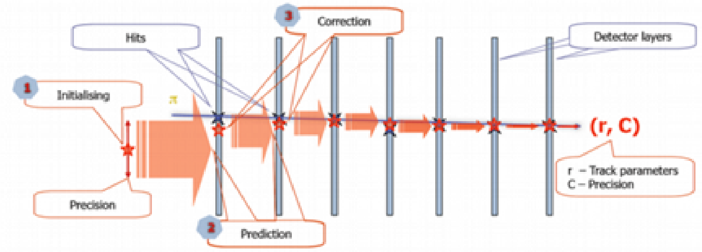

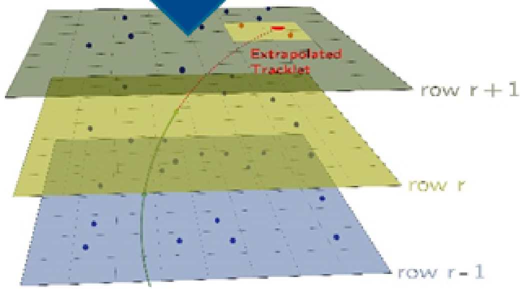

CERN R&D proposal

David Rohr (FIAS)
Sandro Wenzel (CERN)
Giulio Eulisse (CERN)

O^2 Physics Data Reconstruction

Clustering and Tracking with GPUs and machine learning

- From charge distribution to tracklets with **Machine Learning**
  - **“visual” task**, so **neural networks** might be good at it (they can identify cats on pictures, can’t they).

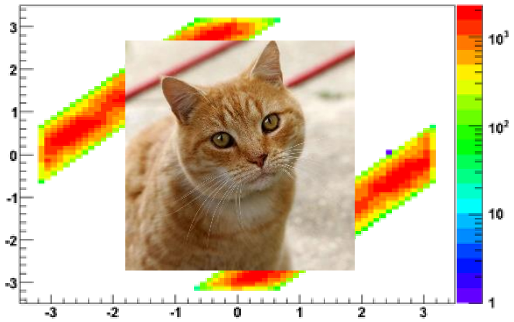

Why don’t we use the charge directly?

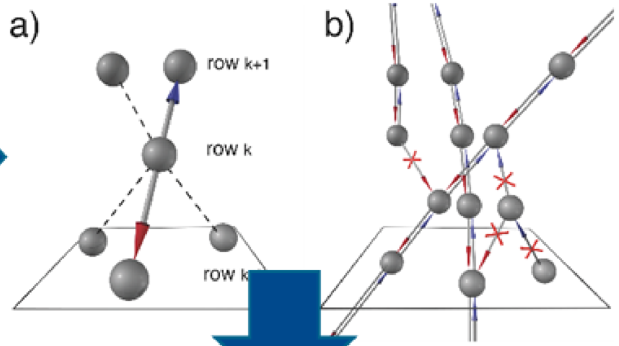

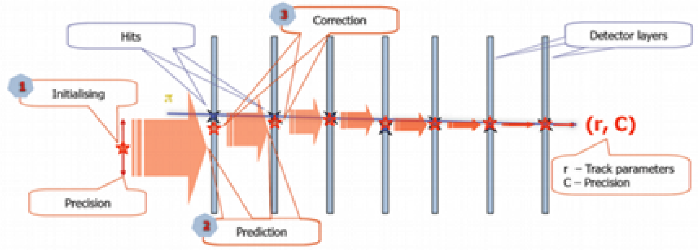

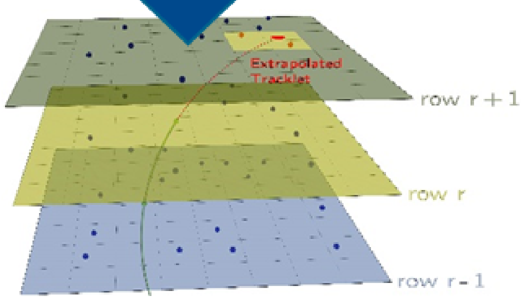

CERN R&D proposal

David Rohr (FIAS)
Sandro Wenzel (CERN)
Giulio Eulisse (CERN)

Conclusions

- O^2 is a new and challenging project
- Working hard on multiple fronts
- Several ideas on automation
  - Some in the plan, others over the horizon