Chapter 4
Computing

Working Group

Workshop da Rede Nacional de Física de Altas Energias (RENAFAE) de 2022 – 25–28/abr/2022

Section 4.1

Introduction

Introduction

RENAFAE

Brief introduction to the computing models used in HEP

- ❏ HPC, GPU, advanced analysis techniques (ML, etc.)
    - highlighting the **synergy among the experiments**
    - opportunities the development of such technology may bring to the country ⇒ applications

- ❏ Brief summary of current usage of computing arsenal
    - offline, online, software, analysis techniques, etc.

- ❏ Prospects of **common** future needs and **joint initiatives** of the groups supported by RENAFAE

For reflection…

Some topics:

- NOW: very important opportunity to work together towards a common goal that benefits all!
- Word of caution ⟶ historical perspective
  - Several previous initiatives to build a common project were not successful…
  - Lack of historical/cultural know-how to effectively work for a common goal
  - Learning process is fundamental ⟶ is in progress!
- High Energy Physics Community ⟶ seems to have reached maturity
  - Creation and consolidation of RENAFAE
  - May have reached “critical mass” of scientists
  - International experience may have helped
  - Showing signs of willingness to gather efforts towards a common goal ⟶ National Project
- Active community of young scientists formed during these years ⟶ excellent sign!
- Good luck to all of us!!

4	Computing		
4.1	Introduction	**Sandra/Gilvan**	Brief introduction to the computing models used in HEP (HPC, GPU, Advanced analysis techniques, like ML), highlighting the synergy among the experiments and the opportunities the development of this kind of technology can bring to the country in terms of applications.
4.2	High Performance Computing	Rogério Iope (UNESP)/ Renato Santana (CBPF)	Brief introduction to HPC, discussing its many components and needs, like network infrastructure and high demanding data processing
4.2.1	GPU Programing	Murilo (UFRJ)	Brief introduction to GPU, the main usage in HEP, possible applications and advantages in this technology
4.2.2	Network Infrastructure	Marcio Costa (UNESP)/Jadir Marra (UNESP)	Description of Brazilian network infrastructure used by HEP, current projects and demands for the future

4.3	Advanced Data Processing	Thiago Tomei (UNESP)/ Gustavo Gil (UFRGS)	Brief introduction to the demands in terms of analysis techniques in HEP (online systems, detector simulations and data analysis), connecting to modern computing algorithms (as Artificial Intelligence)
4.3.1	Online Systems	Sandra Amato (UFRJ)/ João Victor (UFRJ)	Describe the concept of online systems in HEP, current and future projects in Brazil
4.3.2	Detector Simulation	Laura Paulucci (UFABC)/Antonio Vilela (UERJ)/ Tiago Silva (USP)	Brief introduction to detector simulation, basic techniques, current and future projects
4.3.3	Data Analysis Techniques	Marisilvia (USP)/ Cesar Bernardes (UFRGS)	Brief introduction to data analysis and techniques
4.3.4	Detector Operation Optimization	Bernardo Peralta (UERJ)	
4.3.5	Software Engineering	Rodrigo Torres (UFRJ)/ Fernando Guimarães (UFRJ)	
4.3.6	Data Quality and Monitoring	Edmar Egídio (UFBA)	

Section 4.2

High Performance Computing

GPU Programming

Network Infrastructure

High Performance Computing / Offline Data Processing

Introduction

LHC experiments annually process more than an exabyte of data using an average of 500,000 distributed CPU cores, to enable hundreds of new scientific results from the collider.

However, the resources available to the experiments have been insufficient to meet data processing, simulation and analysis needs over the past five years as the volume of data from the LHC has grown.

The problem will be even more severe for the next LHC phases. High Luminosity LHC will be a multiexabyte challenge where the envisaged Storage and Compute needs are a **factor 10 to 100** above the expected technology evolution.

Our HEP community needs to start discussions on how to evolve current computing and data organization models in order to introduce changes in the way we use and manage the IT infrastructure, focused on optimizations to bring performance and efficiency, not forgetting simplification of operations (= infrastructure automation).

The HL-LHC Run will begin operations in 2027/28, with expected data volumes to increase by at least an order of magnitude as compared with present systems.

The LHC experiments' needs in terms of data processing can be partially (opportunistically) covered by commercial cloud providers or general-purpose HPC systems. For storage we should consider federating resources - the **data lake** concept.

Data Lake: the goal is to consolidate geographically distributed data storage systems connected by fast network with low latency. Benefits: flexibility, scalability, lower (and shared) costs.

Joint initiatives for the RENAFAE Project

Some topics for discussion

- **Processing:** containers, clouds, kubernetes clusters, HPC systems (e.g. Santos Dumont)
  - ex.: OSG => PRP (U.S.), EOSC (Europe)
- **Storage:** federated / regional data storage - 'data caching for Science'
  - a 'scientific data lake', to be shared by the HEP/RENAFAE community
  - ex.: OSiRIS (Open Storage Research Infrastructure), IRIS-HEP DOMA
- **Networking:** map common bottlenecks and work together with RENs (RNP, REDNESP); deploy tools for shared monitoring (e.g. PerfSONAR - MaDDash); spread the concept of 'Science DMZ' among the host institutions
  - REDNESP (upcoming network upgrade)
  - RNP - PADEX (https://www.rnp.br/servicos/experimentos-avancados/eciencia/padex)
- **GPUs:** consumer GPUs are not suitable for large-scale AI/ML/DL projects and demands
  - sharing one or two high-end GPU servers among the local HEP/RENAFAE community
- . . .

Network infrastructure in São Paulo (REDNESP)

GPU systems for large-scale projects and datacenters

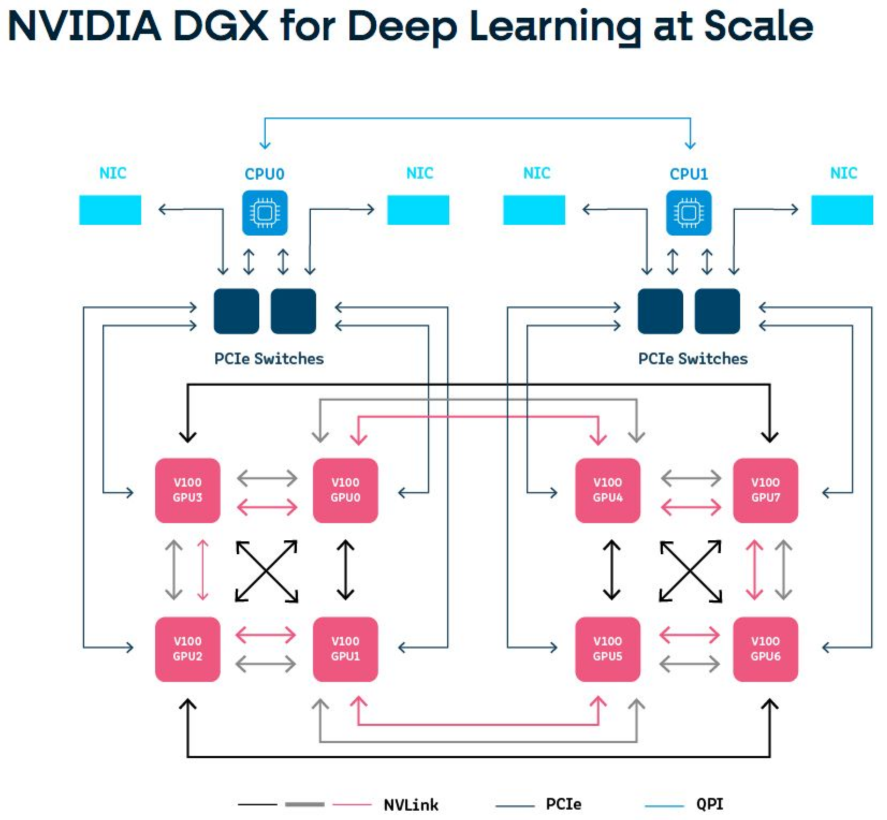

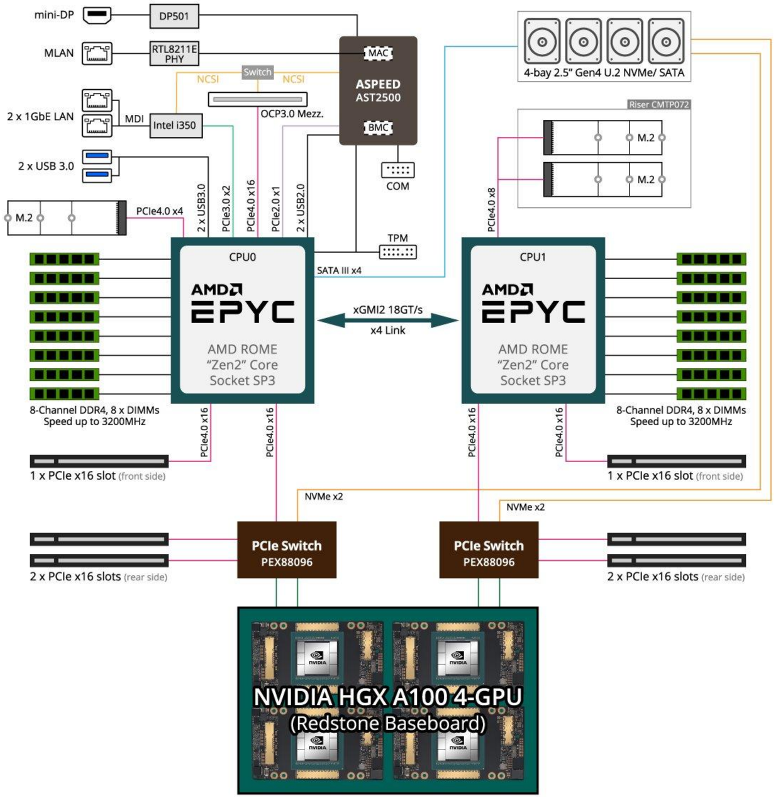

https://www.supermicro.com/en/products/rackmount-workstations

OSiRIS (Open Storage Research Infrastructure)

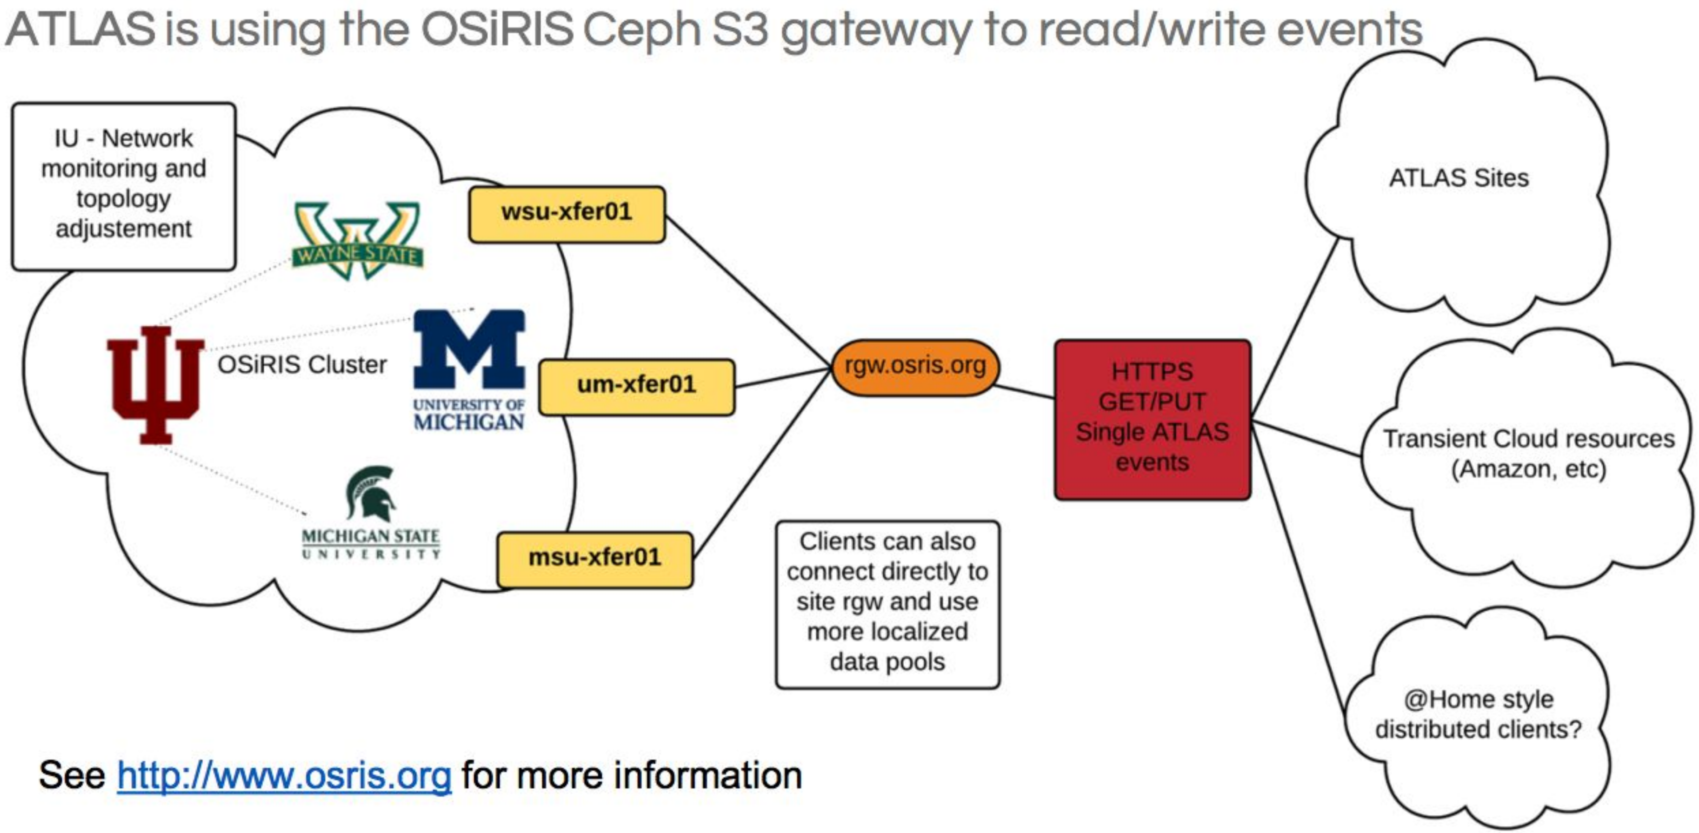

See http://www.osris.org for more information

GPU Programming

High Energy Physics has a large investment in software in critical parts of the physics production pipeline: real-time and offline analysis software

Algorithms need a huge amount of arithmetic on independent data, considering bandwidth as priority over latency.

GPUs have very suitable application in real-time selection and most of experiments are already using it or planning to use it. Offline applications allow complex analysis to be developed

Contributions: Infrastructure tests, software development and cards acquisition.

Table 2. Overview of GPU usage for real-time analysis in various HEP experiments.

Experiment	Main task processed on GPU	Event / data rate	Number of GPUs	Types of GPUs tested	Date for employment	References
NA62	RICH ring pattern reconstruction	10 MHz / 2.5 Gbit/s	1	Nvidia K20c, P100	Tested in 2017 & 2018, planned for 2021	[7, 8]
Mu3e	Track- & vertex reconstruction in the pixel tracker, data selection	20 MHz / 32 Gbit/s	$\mathcal{O}$ (10)	Nvidia GTX980, GTX1080, RTX1080Ti	2021	[9]
CMS	Decoding of raw data, clustering, pattern recognition in the pixel detector	100 kHz / -		Nvidia RTX2080, K20	Planned for 2021	[10, 11]
ALICE	Track reconstruction in the TPC	< 500 Hz Pb-Pb or < 2 kHz p-p / < 100 Gbit/s	64	Nvidia GTX480	2010–2013	[12]
ALICE	Track reconstruction in the TPC	< 1 kHz Pb-Pb or < 2 kHz p-p / < 384 Gbit/s	180	AMD S9000	2015–2018	[12]
ALICE	Track reconstruction in three sub-detectors	50 kHz Pb-Pb or <5 MHz p-p / 30 Tbit/s	$\mathcal{O}$ (2000)		2021	[13, 14]
LHCb	Decoding of raw data, clustering, track reconstruction in three sub-detectors, vertex reconstruction, muon identification, inclusive selections	30 MHz / 40 Tbit/s	$\mathcal{O}$ (500)	Nvidia RTX2080Ti, RTX6000, V100	2022	[15]

D. vom Bruch 2020 JINST 15 C06010

Section 4.3

Advanced Data Processing

Online Systems

Detector Simulation

Data Analysis

Detector Operation Optimization

Data Quality and Monitoring

Cronograma LHC

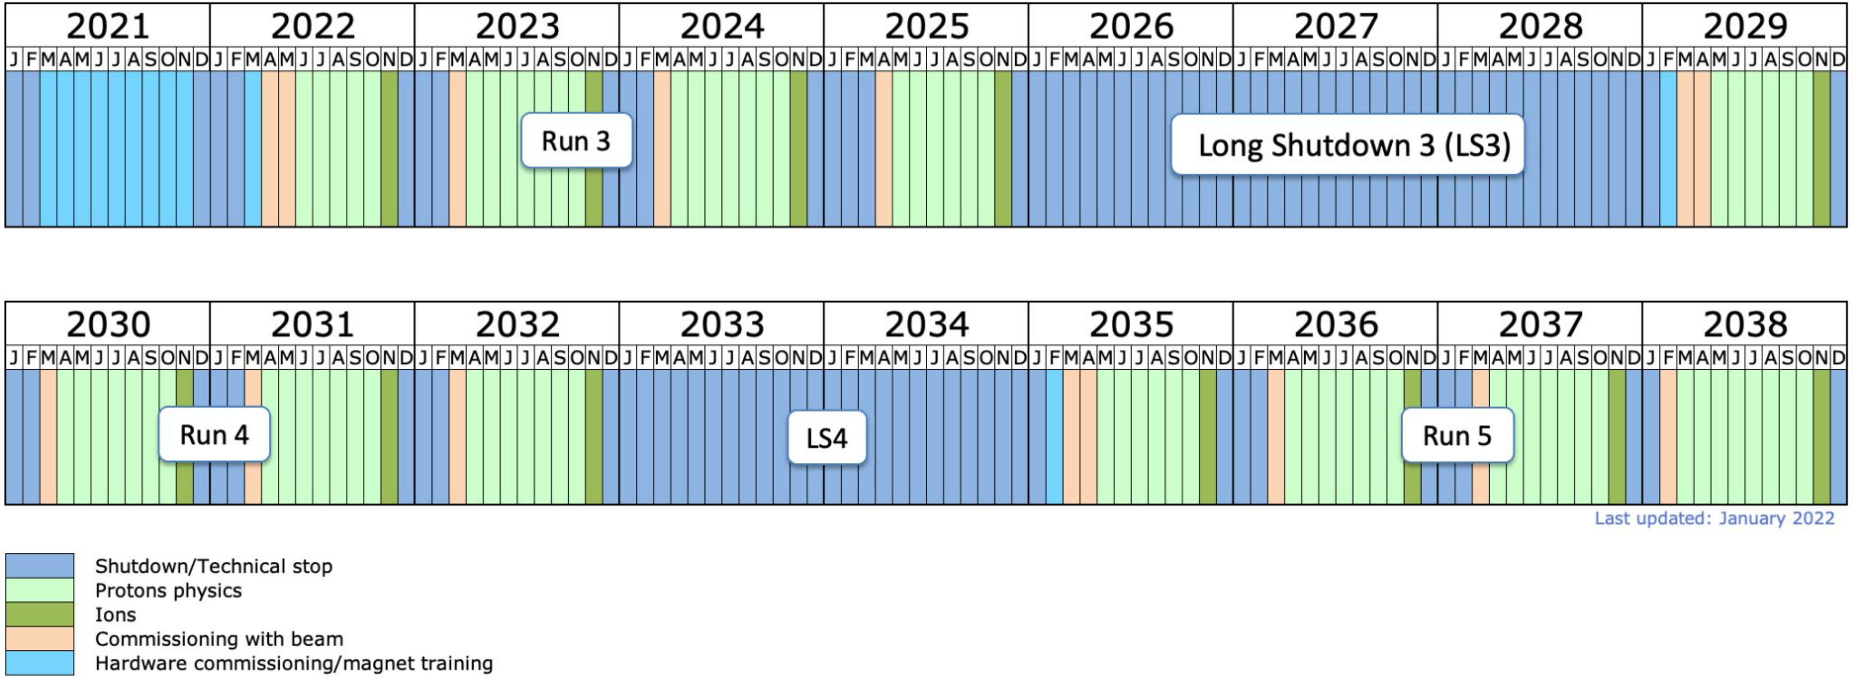

ATLAS and CMS HL-LHC Computing Projections

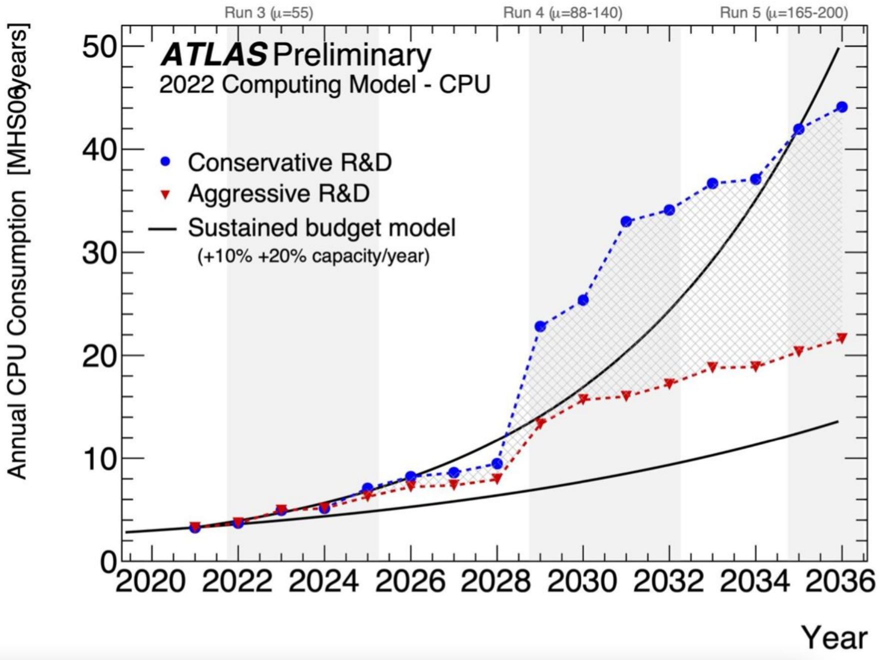

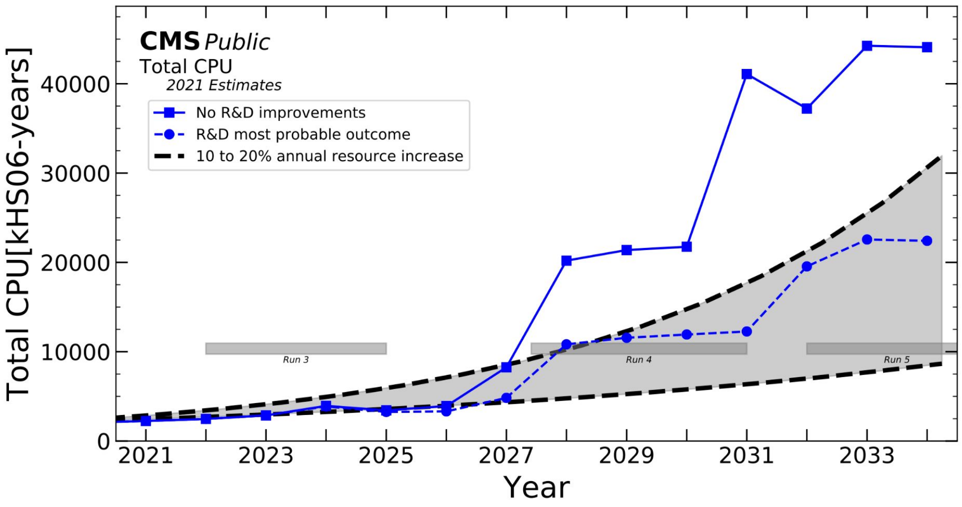

Online Systems

4.3.1

Applications for High-level trigger

- Solutions with HPC
  - GPU
  - FPGA

Data acquisition

Online Systems – ATLAS

Online systems are often used in HEP experiments to classify events of interest and reduce background incidence during collisions:

- Sequential strategies combine the use of hardware and software to preserve the physics of interest and increase background noise rejection;
- Usually composed of feature extraction algorithms based on calorimetry and hypothesis testing using shower discriminant variables;
- In ATLAS, the frequency of pp collisions at 40 MHz and the total (without trigger) data output rate would be ~70 TB/s;
  - Not everything can be stored: It must be efficient in the physics of interest and still reject most non-relevant events.

Online trigger systems are developed dedicated to particles of interest:

- In the case of electron detection in ATLAS, sequential strategies use calorimetry information to classify electrons, reducing the incidence of background noise at each step.

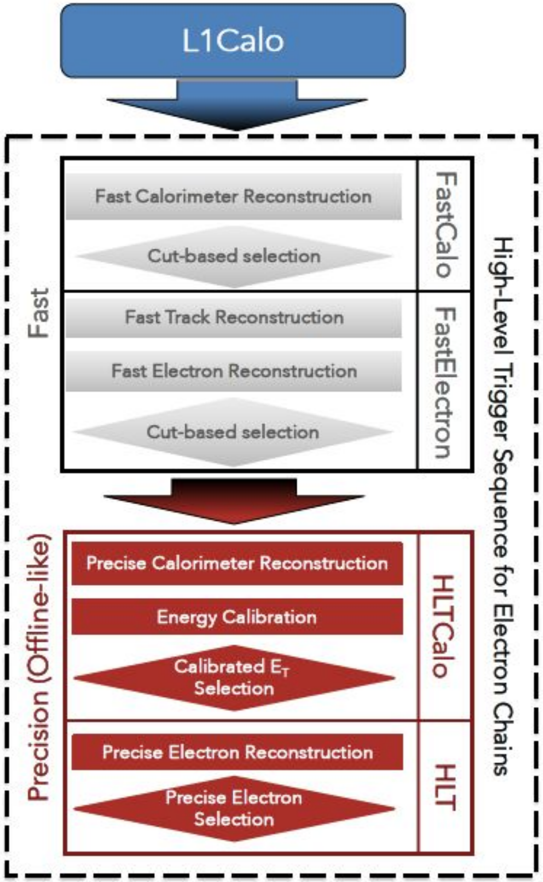

Online Systems – CMS

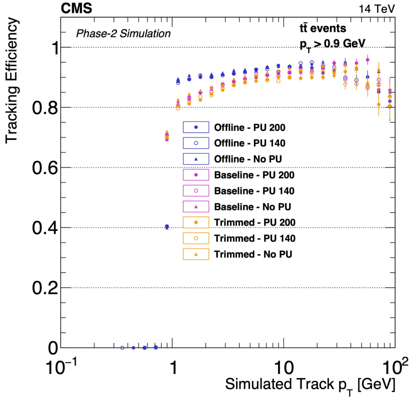

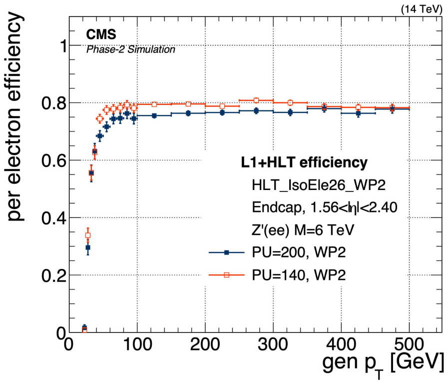

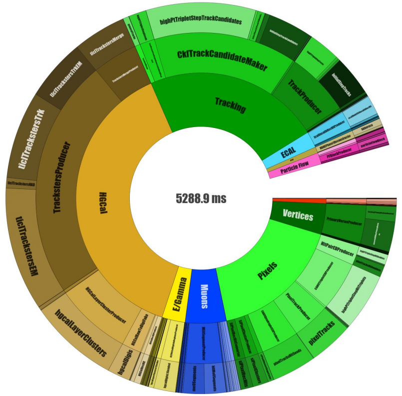

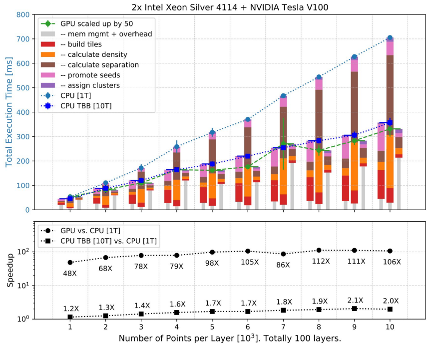

HL-LHC conditions

- 75 Hz/nb, <PU> = 200

DAQ requirements

- Input: 750 kHz, 50 Tb/s
- Output: 7.5 kHz, 50 GB/s

HLT requirements

- Same performance as Phase-1
- Budget: 16 MCHF

TDR results

- High-efficiency physics objects
- Simplified menu (50% target rate)
- Roadmap for timing reduction
- Heterogeneous architectures (GPU)

https://cds.cern.ch/record/2759072/

Online Systems – LHCb

Goal: accumulate large data samples.

Challenge: filter uninteresting events while keeping high efficiency for signal events.

Events built at high rate with a online (real-time) analysis similar or the same as performed offline. Alignment and calibration performed online before complete event reconstruction.

Due to storage and bandwidth capacity limitations, different solutions can be implemented such as reducing event size and using high-performance computing techniques (FPGA, GPUs, …).

Contributions: Development to software infrastructure, validation of software projects, operation and maintenance of the system.

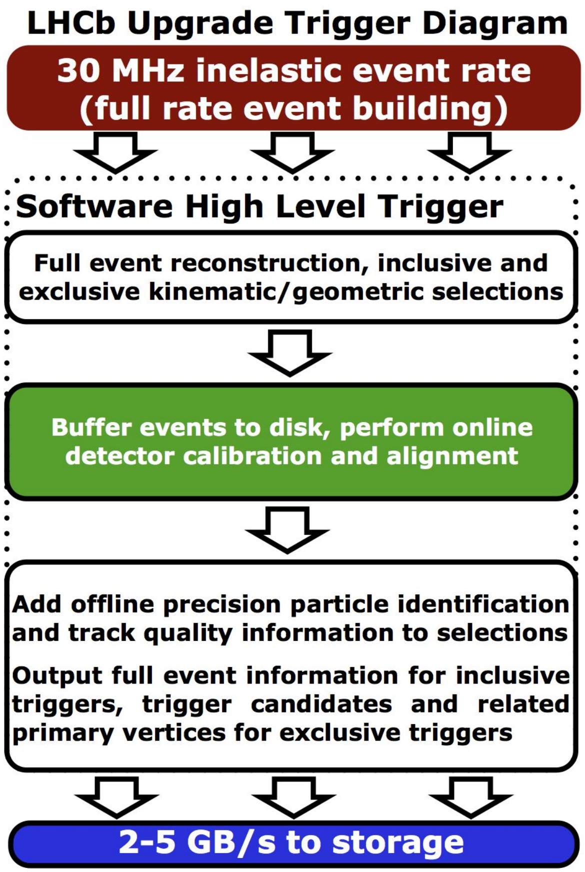

Detector Simulation

Section 4.3.2

Integration of detector simulation & simulation frameworks

Detector characterization and performance

Design decisions

Reconstruction tools

Fast and Full simulation

- Established frameworks
- Generic simulation tools

Machine Learning techniques

- Neural networks (GANs, DNNs,...)

Simulation challenges

Section 4.3.2: Detector simulation

Integration of detector simulation & simulation frameworks

- Geant V: https://arxiv.org/abs/2005.00949 (Geant4 is “only game in town”)

Modelling of new technologies & detector optimization

Example: Simulation of Micro-Patterned Gaseous Detectors (MPGDs)

Challenges in processing of simulation (measurables in terms of number of processed events, dataset size, etc.)

Fast and *Full* simulation (Collaborations)

- AtlFast: https://arxiv.org/abs/2109.02551
- CMS FastSim: https://arxiv.org/abs/1701.03850
- LHCb: ReDecay: http://arxiv.org/abs/1810.10362
- Future colliders (FCC, ILC)
- DUNE
- …

Fast simulation (Generic simulation)

- Delphes: https://arxiv.org/abs/1307.6346 (“only game in town” for large detectors?)
- GGS: https://arxiv.org/abs/2104.10395 : (optimised Geant4 for small- and medium-sized detectors)

Section 4.3.2

Machine Learning for Detector Simulation

CaloGAN: fast simulation of electromagnetic showers in calorimeters
https://arxiv.org/abs/1712.10321

GAN for fast simulation of the ATLAS calorimeter
https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-SOFT-PUB-2020-006/

GAN for Fast simulation of muons in SHiP experiment:
https://arxiv.org/abs/1909.04451

Bounded Information Bottleneck Autoencoder for High Granularity Calorimeters:
https://arxiv.org/abs/2005.05334

MPGDs – Simulation of Degradation and Aging

Goals:

- Provide interpretation for experiments
- Constraint models
- Evaluation of material migration within the detector
- Increase explainability of aging process

Requirements:

- Better databases of Penning factor
- Data on resonant charge exchange
- Data of molecular breakup
- Integration with gas flow simulation
- Transport of complex charged molecules

Implement integration between microscopic and macroscopic simulations!

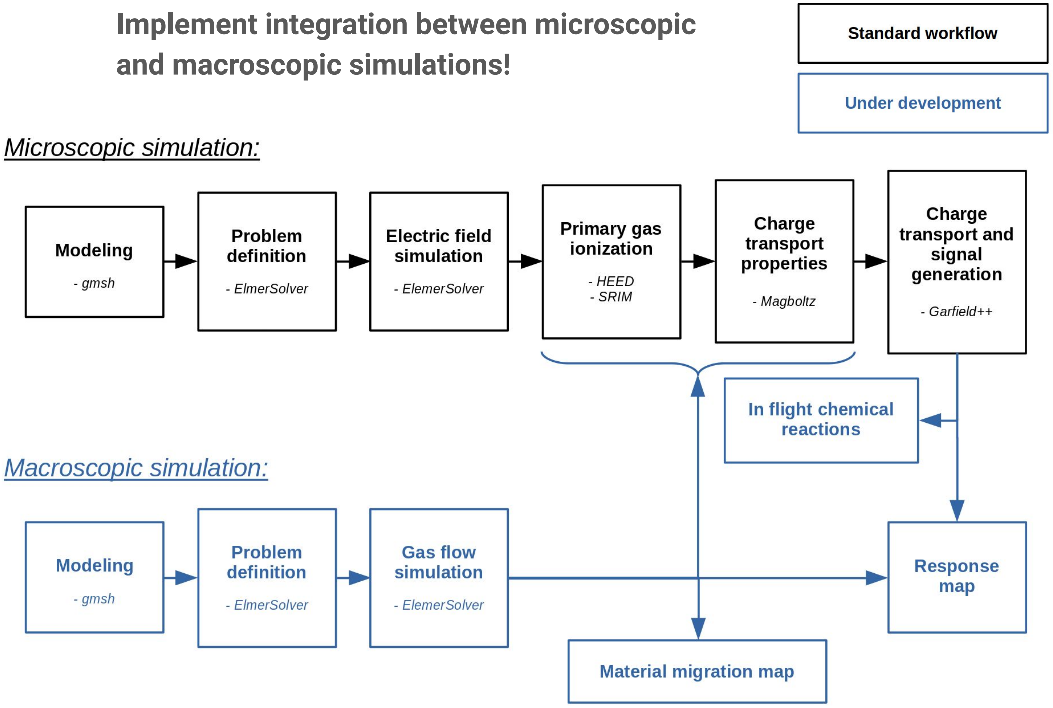

MPGDs Simulation — Microsized Optimizations

Goals:

- Explore alternative geometries for new types of MPGDs
- Mitigate charging up and ion backflow effects
- Integrate with 3D printing formats for prototyping
- Simulate artifacts of 3D printing process

Requirements:

- Enable flexible definition of geometries in the simulation software
- Protocols for mesh size tests
- Data visualization and interpretation

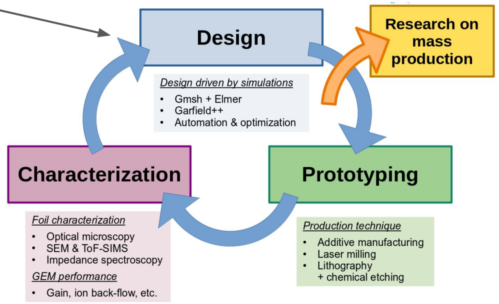

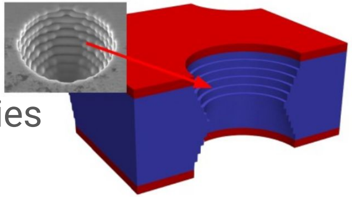

Preliminary simulation results

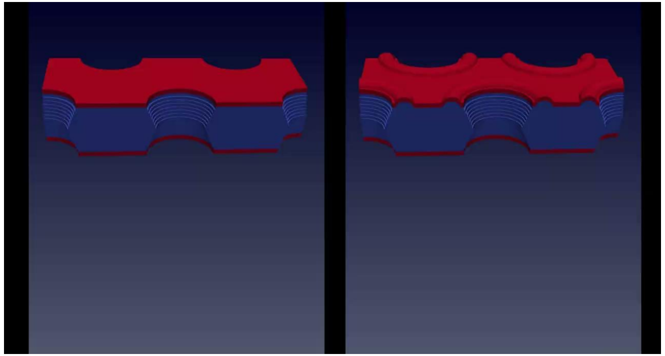

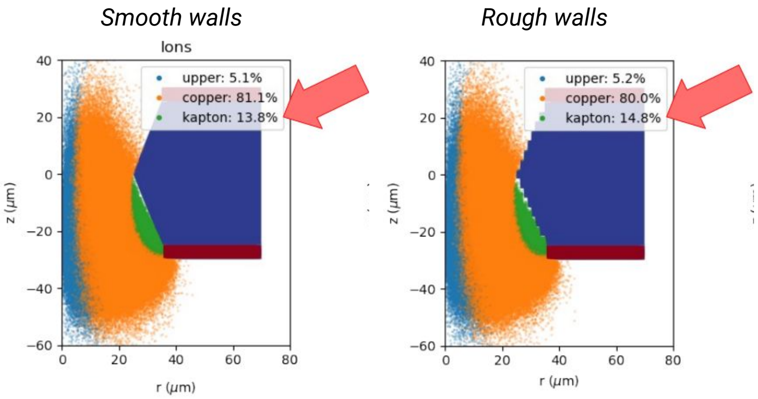

Expected increment of ~7% in charging up effect due to roughness of internal walls

Offline Event Reconstruction and Data Analysis

Section 4.3.3

- Event generation
  - Generating four-momenta
  - Monte carlo integration
- Offline event reconstruction
  - Tracking
  - Calorimeters
  - Hadronic jets
  - Particle flow & single particle reconstruction
- Data access and manipulation
  - Columnar data structures
  - Data analysis in the cloud
  - Histograms, fits, statistics, interpretations
  - Corrections, syst. uncertainties, unfolding
- Relationship with other areas (outside HEP)

Offline Event Reconstruction (1)

Tracking

Exa.TrkX: https://arxiv.org/abs/2103.06995 – Geometric Deep Learning

- Metric Learning
- Graph Neural Networks (GNNs)

PataTrack: https://www.frontiersin.org/articles/10.3389/fdata.2020.601728/full

- Initially used in CMS pixel tracking
- Fully focused in heterogeneous reconstruction

Acts: Acts Common Tracking Software:

- Initially based on ATLAS tracking
- Has realistic GNN-based algorithm: https://arxiv.org/abs/2103.00916
- Links:
  - https://acts.readthedocs.io
  - https://arxiv.org/abs/1910.03128
  - https://arxiv.org/abs/2106.13593

Offline Event Reconstruction (2)

Calorimeters

Deep Learning + Cellular Automaton for LHCb Electromagnetic Calorimeter

- https://doi.org/10.3390/app112311467
- Deep CNNs can learn rules of CAs (current LHCb algorithm)
- Reconstruct clusters in nearly constant time

TICL: The Iterative CLustering for CMS HGCal

- https://doi.org/10.1051/epjconf/202125103013
- 2D layer clusters from density-based algorithm
- 3D clustering with physics-focused iterations and masking

GNNs + Graph Condensation for CMS HGCal

- https://arxiv.org/abs/2204.01681
- End-to-end ML: clustering classification and energy/position regression
- GravNet (https://arxiv.org/abs/1902.07987) blocks and object condensation loss function

Data Access and Manipulation (1)

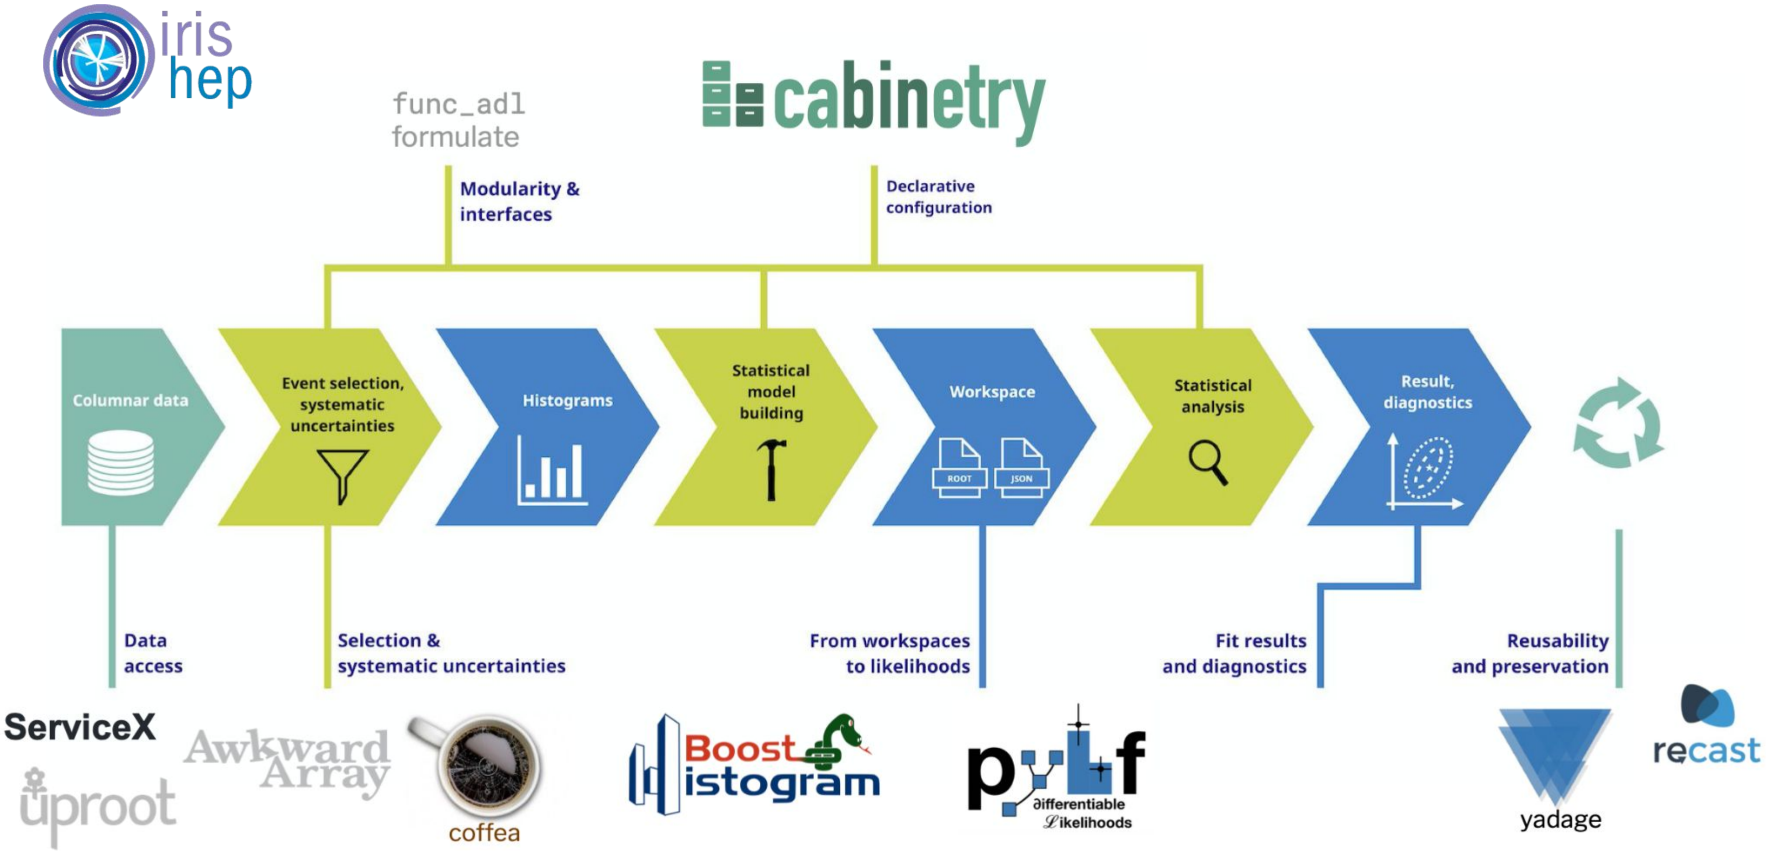

Data Access and Manipulation (2)

RENAFAE

Columnar data structure

- ROOT: RDataFrame: https://root.cern/manual/data_frame/
  - TTree → RNTuple: https://root.cern/doc/master/md_tree_ntuple_v7_doc_README.html
- Awkward Array: https://awkward-array.org/quickstart.html
  - TTree via Uproot: https://uproot.readthedocs.io/en/latest/index.html

Data analysis in the cloud

- All based on an a Jupyter notebook interface: https://jupyter.org/
- SWAN (Service for Web based ANalysis): https://swan.web.cern.ch/swan/
  - File storage/sharing: CERNbox
  - Computing: CERN Spark clusters
- Coffea-casa: https://coffea.casa/
  - Computing: Nebraska T2 via Kubernetes + HTCondor + Dask
- “INFN effort”: https://doi.org/10.1051/epjconf/202125102045
  - Computing: INFN via Kubernetes + HTCondor + Spark

Data Access and Manipulation (3)

RENAFAE

Statistics and (Re-)Interpretation

Pyhf: pure-Python implementation of HistFactory models

- https://doi.org/10.21105/joss.02823
- https://pyhf.readthedocs.io/
- Builds on top of well-loved RooFit / RooStats
- Together with published full-likelihoods, allows REPRODUCIBILITY

Reproducible research data analysis platform (REANA)

- https://cds.cern.ch/record/2652340
- https://reanahub.io/

Or, maybe the likelihood is *actually* intractable 🙂

- Madminer: https://arxiv.org/abs/1907.10621

Detector Operation Optimization

Section 4.3.4

- Energy loss in matter
  - Charged particles
  - Tracking
- Calorimeters
  - Electromagnetic
  - Hadronic shower
- Particle identification
- Operation
  - Electronics
  - Calorimeter energy reconstruction
  - Alignment and calibration systems
  - Current algorithms and perspectives

Detector Operation Optimization

Motivation:

- Detectors are crucial for getting to Frontier Physics

Goals:

- Detector design is based on a deep understanding of particle/ matter interactions
- Extremely complex experiments are built based on important principles such as:
  - Interactions of radiation with matter
  - Sensors and read-out principles
  - Computational constraints

Optimizations:

- Deal with unprecedented conditions (pile-up)
- Use of computational intelligence strategies for both online and offline operations. Examples:
  - Electron trigger in ATLAS
  - Calorimeter Energy estimation in ATLAS TileCal

Electron interactions with several systems

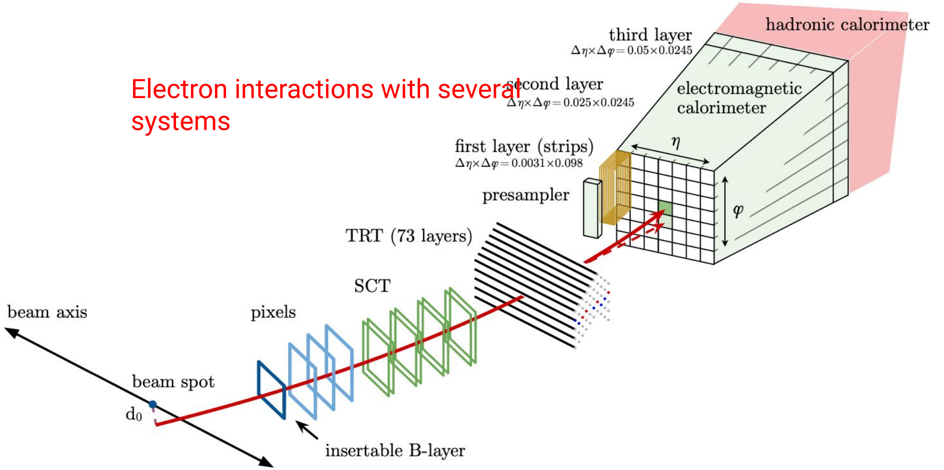

Energy estimation for calorimeter cell based on linear estimation (optimal filter) assisted by a neural network

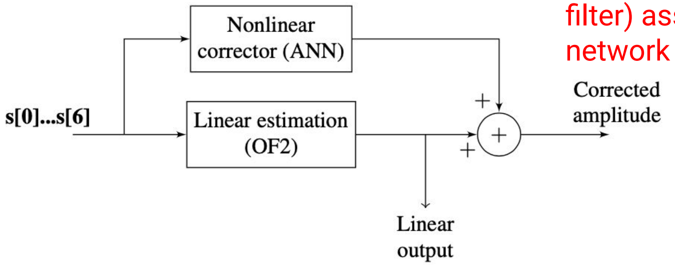

Detector and Computing Operation Optimization

HL Trigger com ML em FPGA?

HL Trigger em GPU?

Data scouting with ML?

Operational Intelligence (OpInt) project: https://doi.org/10.3389/fdata.2021.753409

- Increase of automation in computing operations, reduce human interventions.
  - Machine learning
  - Data mining
  - Log analysis
  - Anomaly detection

Data Quality and Monitoring

Section 4.3.6

- Anomaly detection
- Data certification
    - Deep autoencoders
- Trigger rate inspectors
- Machine Learning proposals
    - AutoDQM
- Applications
    - Collider Machine
    - Other experiments

Data Quality and Monitoring (1)

Data Quality and Monitoring workflows are often used in HEP, primarily for:

- Allows to check for possible problems during data taking (electronic issues, readout, calibration or computation);
- Allows to observe the distributions of shower shapes at each step of Trigger;
- Prevents data collected with any problem from being used in physics analysis;

Main contributions of the Brazilian Cluster to the ATLAS experiment in DQ & Monitoring:

- Online monitoring software for reconstruction algorithms;
- Development of webdisplays for comparison between distributions;
- Development of the framework for offline monitoring of egamma triggers

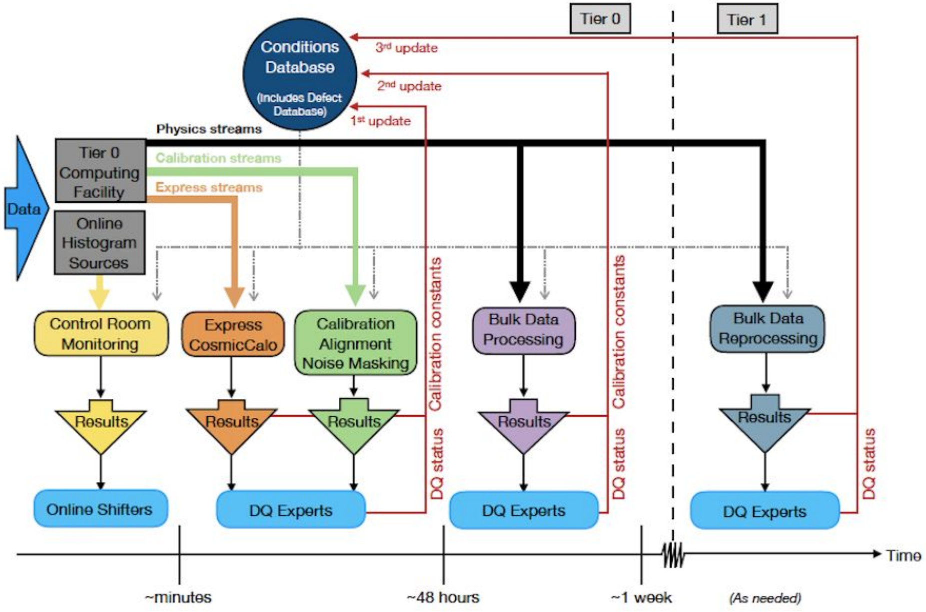

Schematic diagram illustrating the nominal Run2 operations workflow for the data quality assessment of ATLAS data [1].

[1] ATLAS COLLABORATION - ATLAS data quality operations and performance for 2015–2018 data-taking. Journal of Instrumentation, 15, 2020

Data Quality and Monitoring (2)

AutoDQM

- Tool to improve current DQM within CMS, based on semi-automation of the monitoring process
- Development is ongoing for Run 3 of the LHC, including addition of other CMS subsystems and additional functionality
- Modularity of the tool allows for easy integration of subsystems and further comparators
- Potential for implementation of more powerful ML techniques to compare reference and data runs, and across multiple runs and histograms
    - Clustering algorithms (DBSCAN, k-means)
    - Autoencoders
    - Time correction in autoencoders and PCA
    - Combining multiple runs for larger luminosity datasets
- Links
    - https://github.com/AutoDQM/AutoDQM
    - https://github.com/jkguiang/AutoDQM/wiki

Discussion