CERN

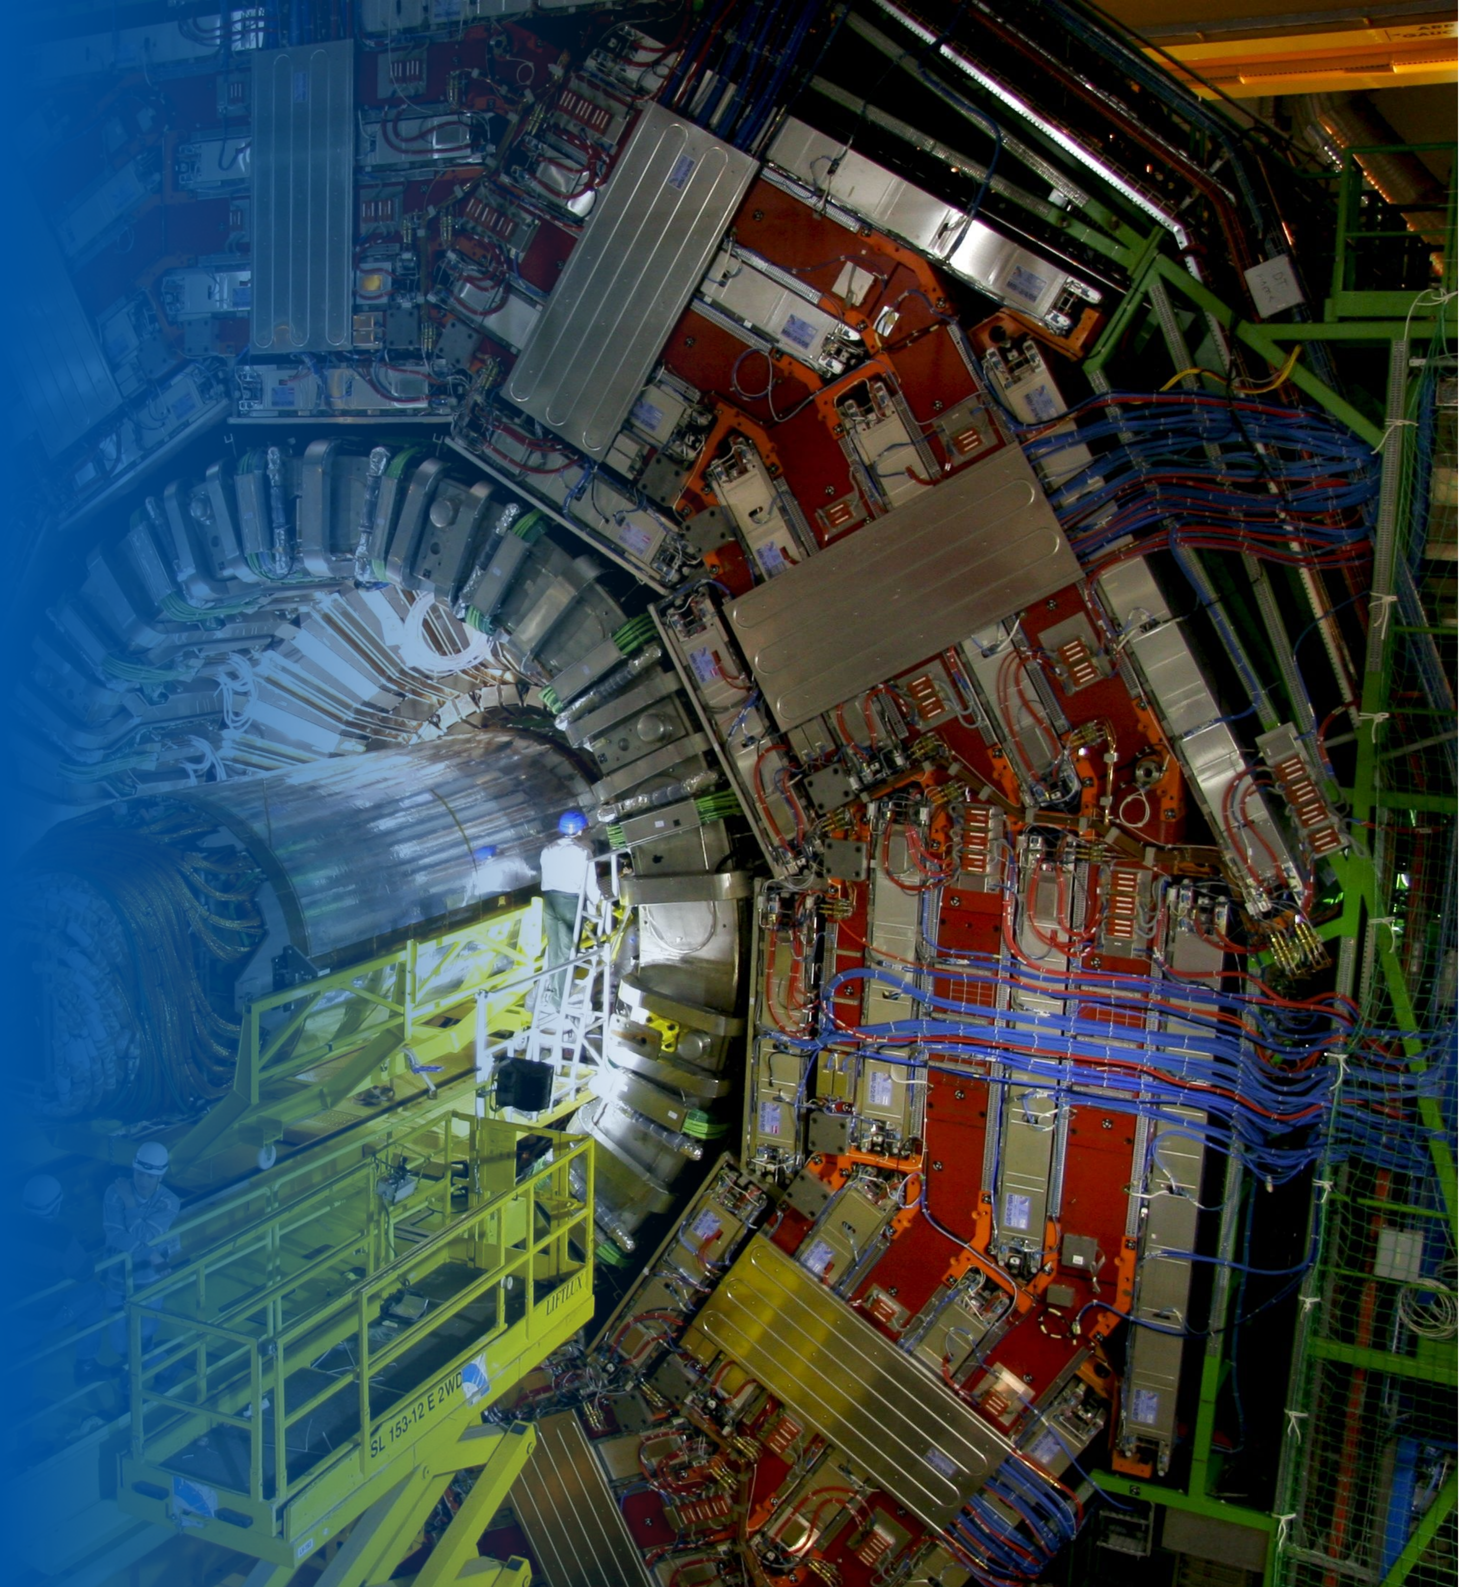

The Trigger of the CMS Experiment

Andrea Bocci, CMS Trigger Coordinator

3rd KSETA Plenary Workshop 2016
22-24 February 2016

Outline

- Introduction
  - the Compact Muon Solenoid experiment
  - why we need a trigger system
- the Level 1 Trigger
  - limitations and challenges
  - architecture and performance
- the Data Acquisition system
  - the Read Out and Builder Units
  - the Filter Farm
  - the Storage Manager and Transfer System
- the High Level Trigger
  - limitations and challenges
  - architecture and performance
- Conclusions

Introduction

Compact Muon Solenoid

- is a collaboration of over 3000 people from 180 institutes and 40 countries
- is one of the four large experiments at the Large Hadron Collider at CERN
- has collected over 30 fb^{-1} of collisions data, and published over 460 papers

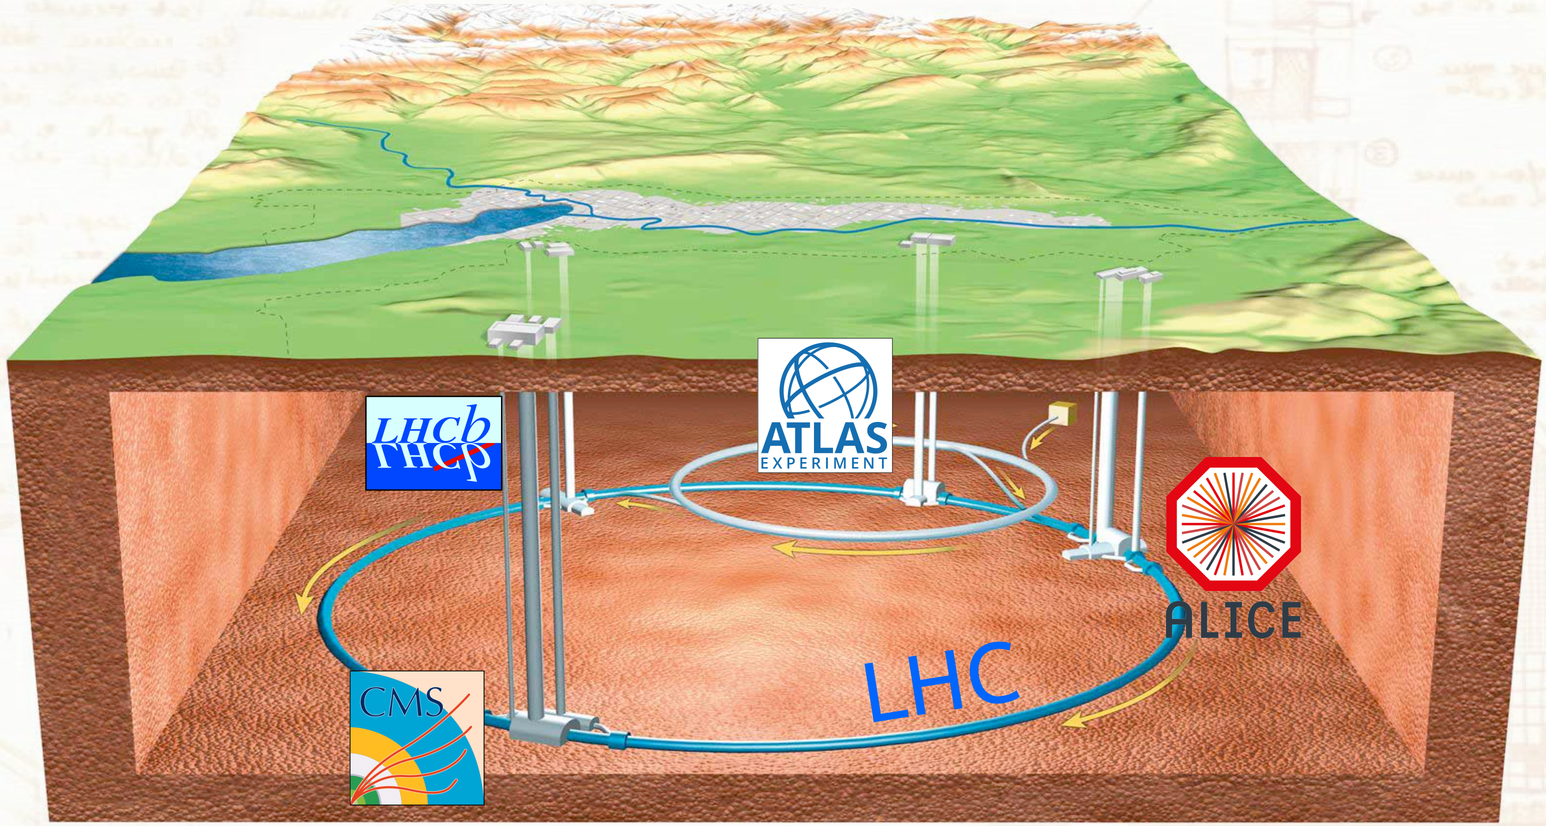

the CMS detector

- Silicon Tracker
  - 65M pixels and 10M strips
- Electromagnetic Calorimeter
  - 75k lead tungstate crystals
- Hadronic Calorimeter
  - sampling calorimeter
- Muon Detectors
  - Drift Tubes (barrel)
  - Cathode Strip Chambers (endcap)
  - Resistive Plate Chambers (both)
- Event rate and size
  - full read out of the detectors limited to 100 kHz
  - “raw” event size: 500 kB

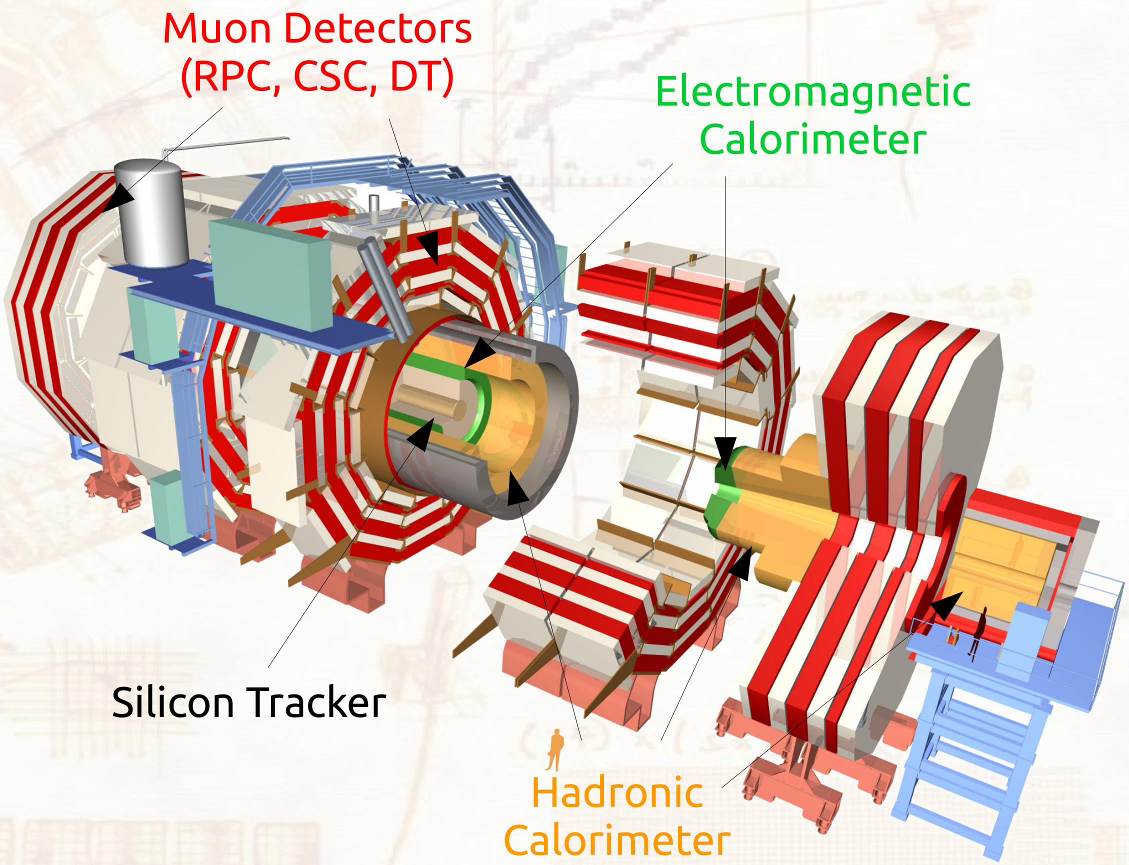

LHC operations in 2016

- two proton beams, colliding at an energy of **13 TeV** ⟶ **σ_{pp} = 80 mb**
  - 25 ns bunch spacing
  - ~2700 proton bunches, colliding every 11.2 kHz
  - collisions at **30 MHz**

- expected peak luminosity
  - **1.0×10^{34} $cm^{-2}s^{-1}$** = 10 $nb^{-1}s^{-1}$

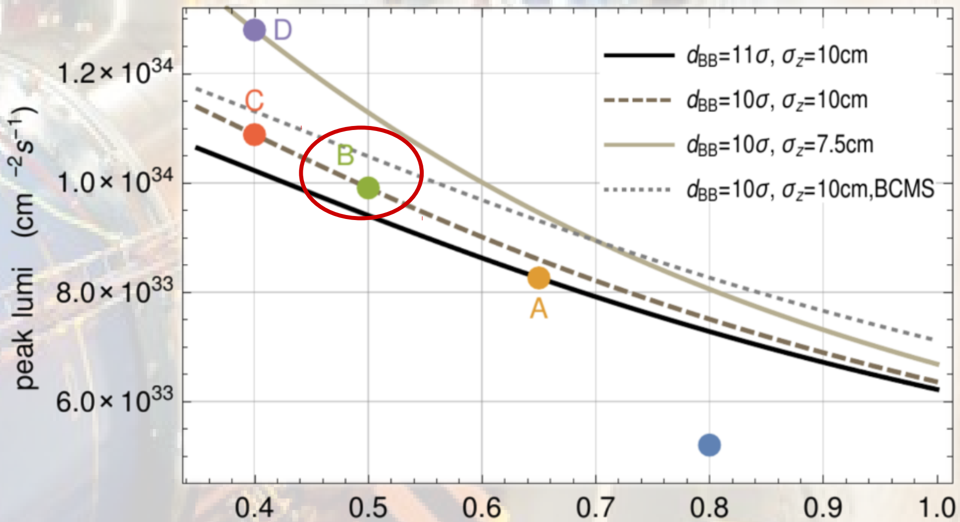

- pileup !
  - 10 $nb^{-1}s^{-1}$ × 80 mb = **800 MHz**
  - average pileup at peak luminosity **~ 26**

Why do we "trigger" ?

- last year, the LHC produced at CMS

  - ~ 3 million Z → μμ events

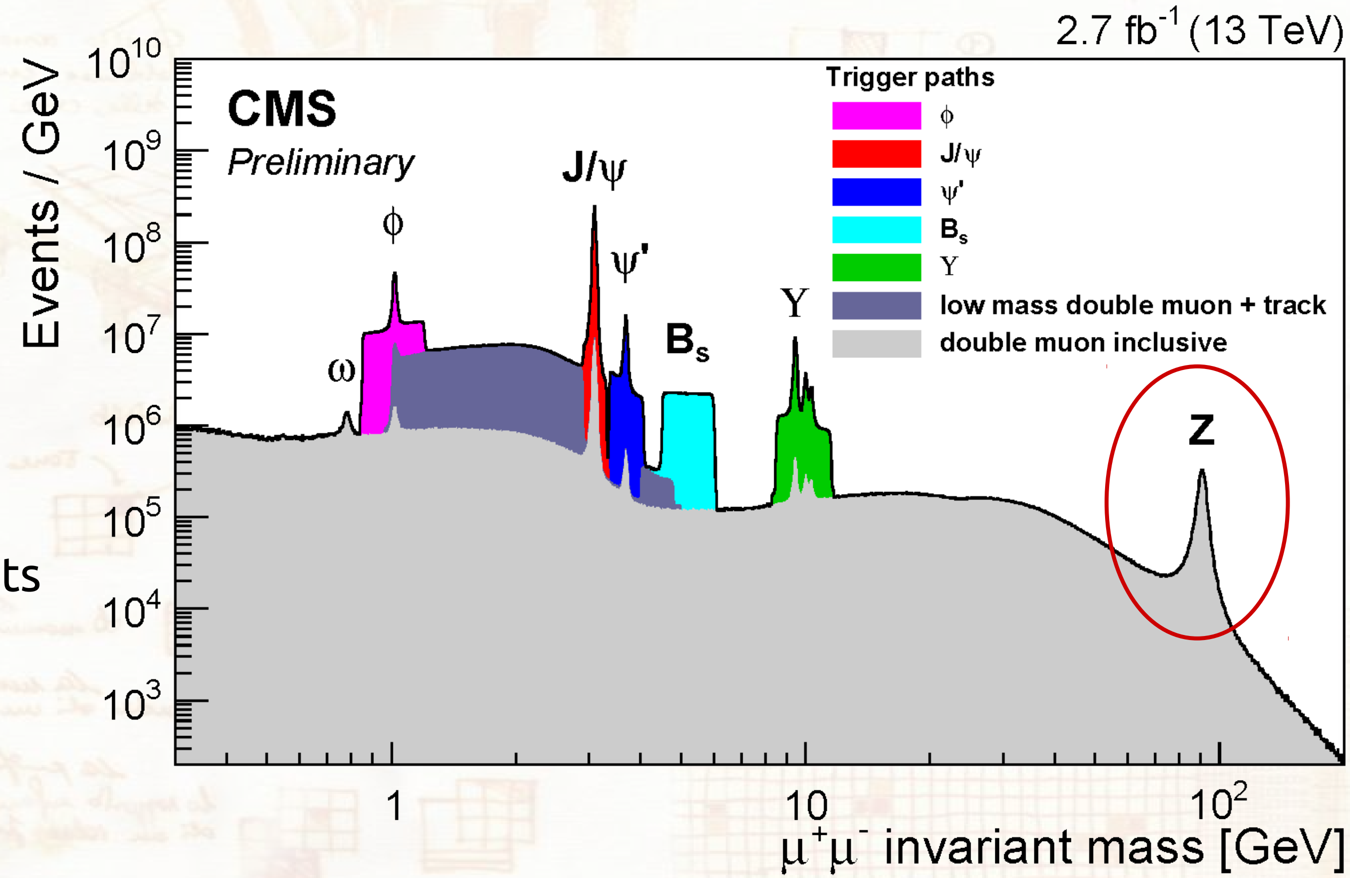

Why do we "trigger" ?

- last year, the LHC produced at CMS
  - ~ 2 × 10^{14} pp collisions
  - ~ 3 million Z → μμ events

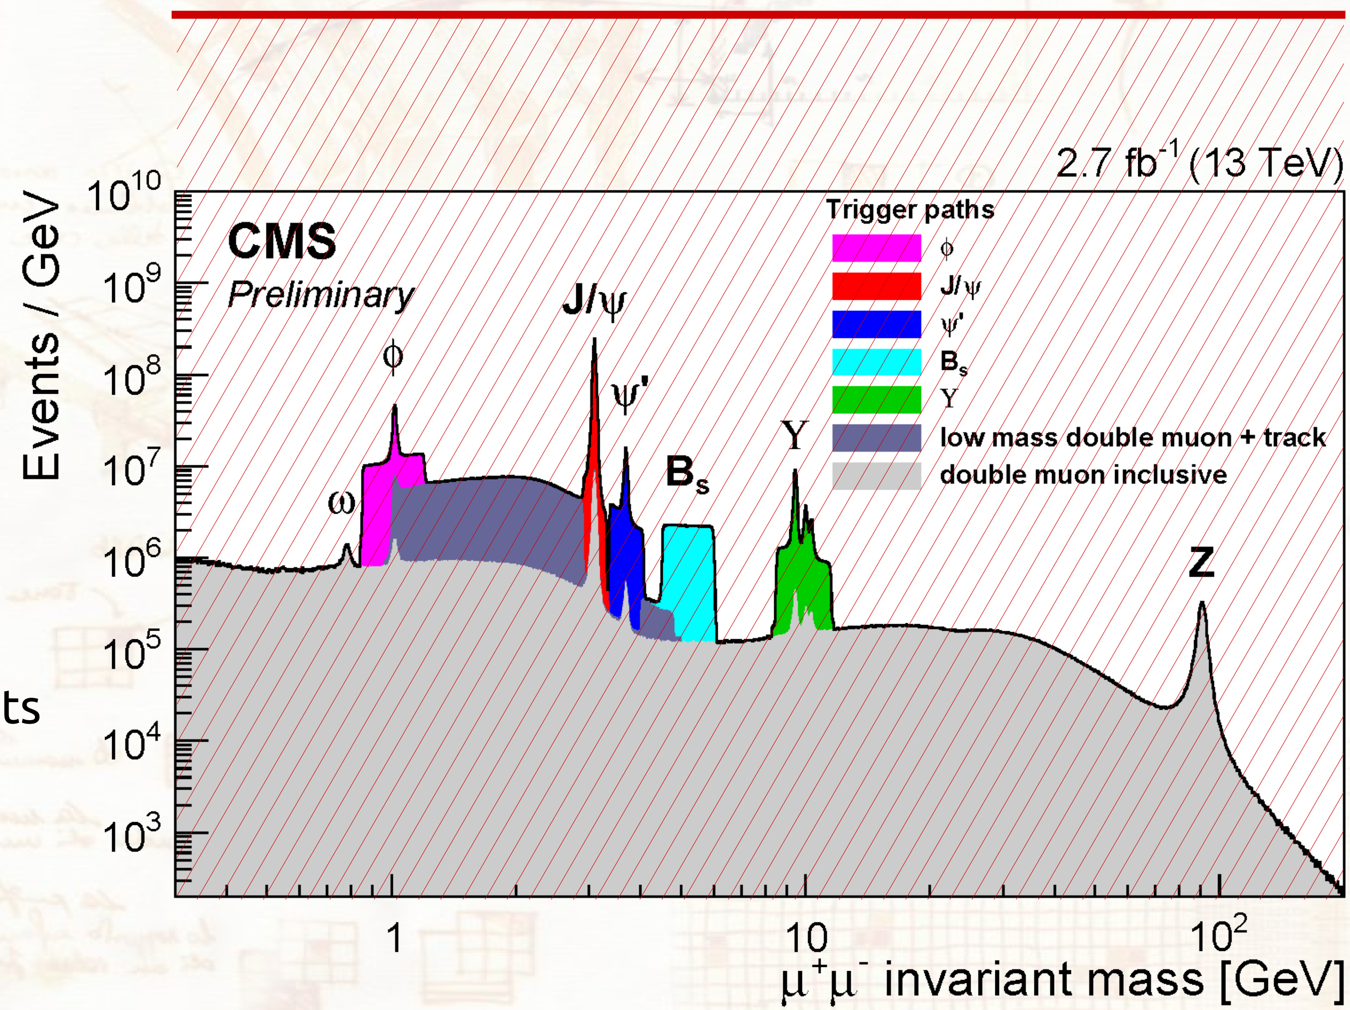

To reduce the event rate …

… from the **interaction rate** of the machine, to a rate that allows

- the full detector read out
- **storage on disk**, offline reconstruction and analysis

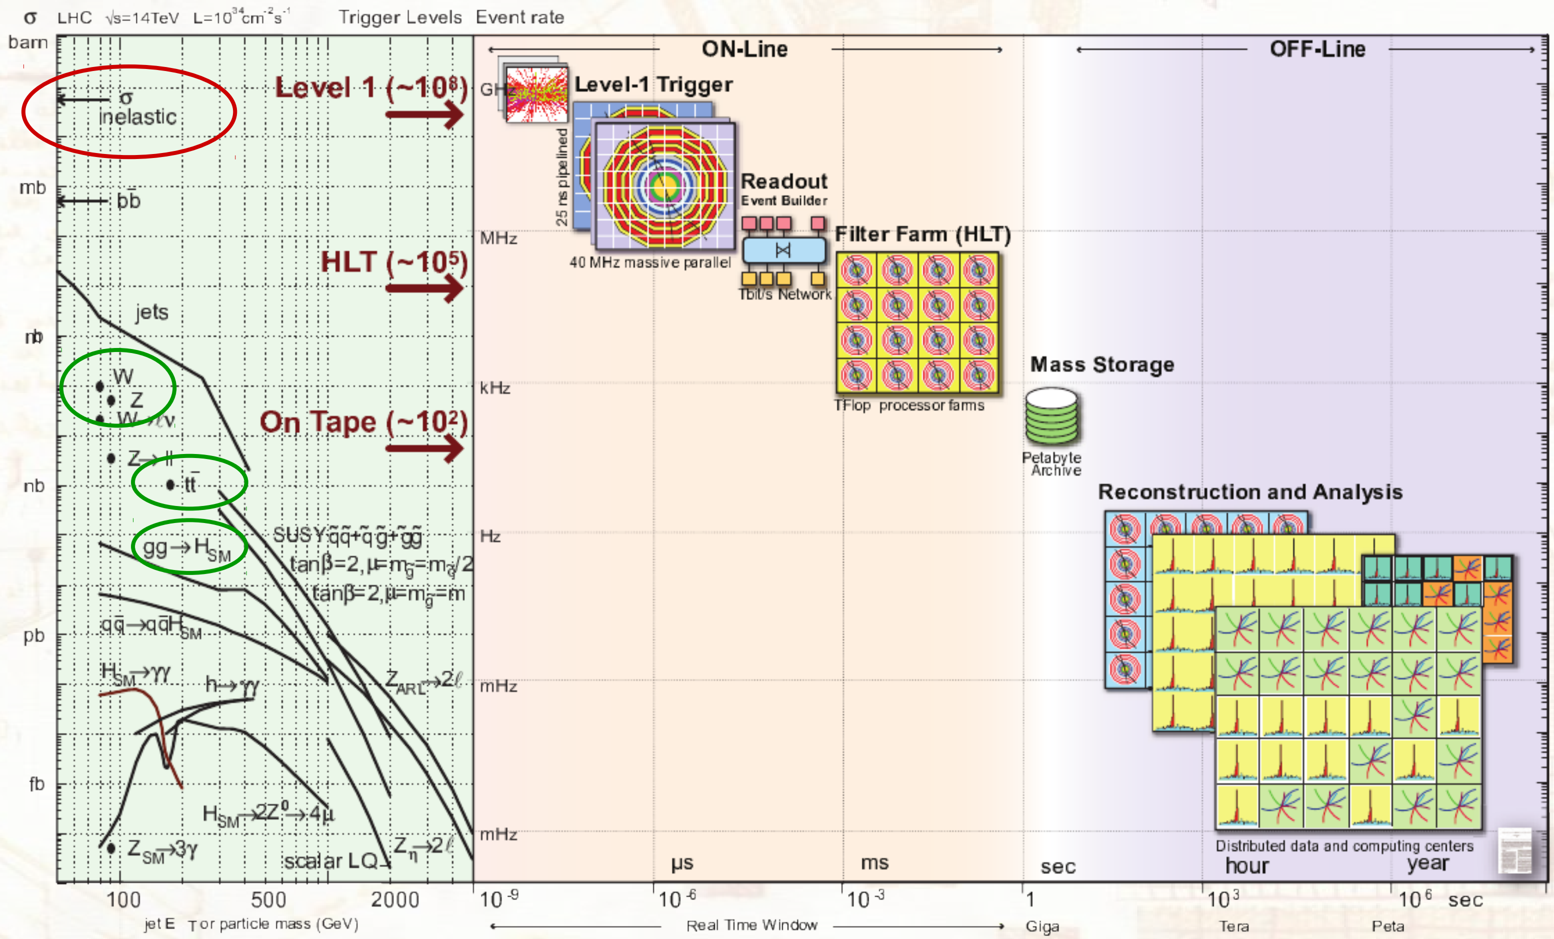

- while maintaining as much as possible the **physics reach** of the experiment.

Challenges of a trigger at LHC

- **low latency and processing time**
  - analyse as quickly as possible all the collision events
  - massively parallel

- **high purity**
  - discard **the uninteresting events** as soon as possible
  - layered approach

- **high efficiency**
  - keep as many as possible of the **interesting physics events**
  - high level algorithms and calibrations close to those used in the offline analyses

Two – layer approach

- different architectures
  - no “one size fits all” solution
  - exploit all the information available at each step

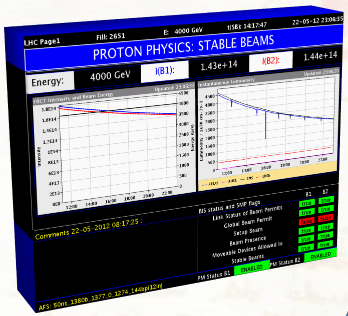

LHC

40 MHz

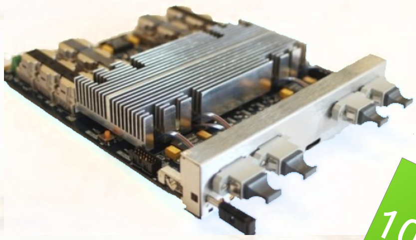

Level 1 Trigger

100 kHz

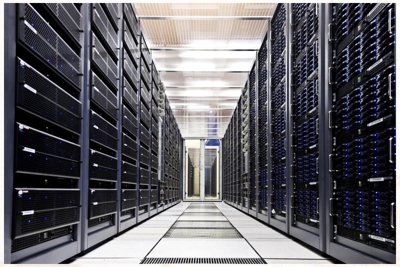

High Level Trigger

~1 kHz

Offline reconstruction and analyses

- successive steps
  - reduce data rate
  - increase granularity and complexity

Level 1 Trigger

Level 1 Trigger

- fast readout of the detector, with a coarse granularity

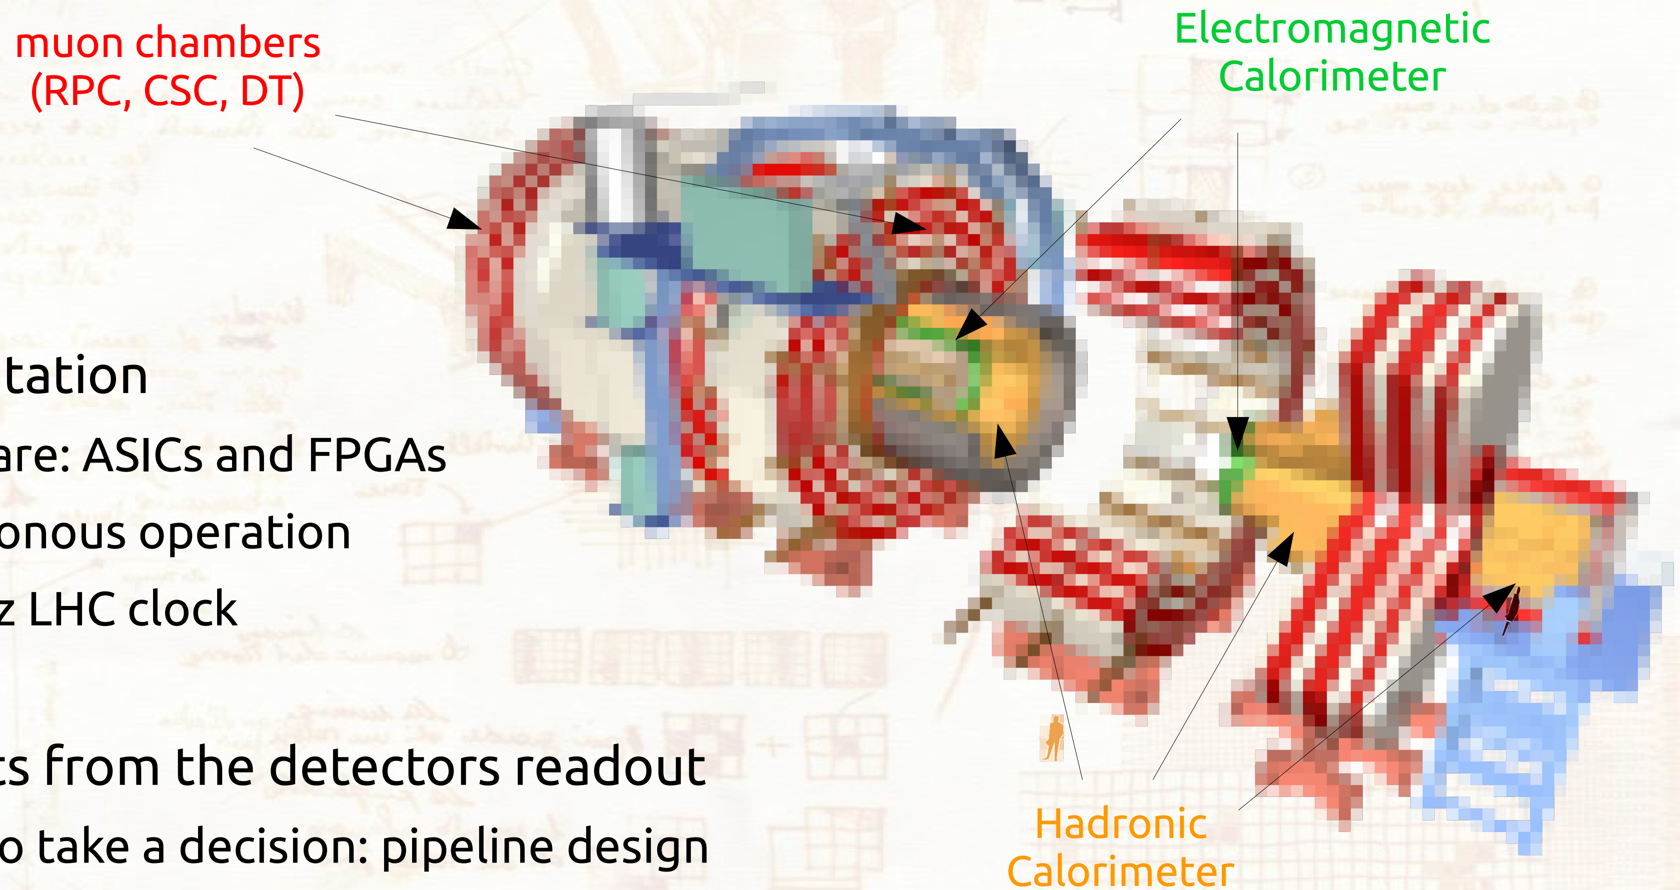

- implementation
  - hardware: ASICs and FPGAs
  - synchronous operation
  - 40 MHz LHC clock
- constraints from the detectors readout
  - ~4 μs to take a decision: pipeline design
  - readout: 100 kHz maximum output rate

2009 – 2012

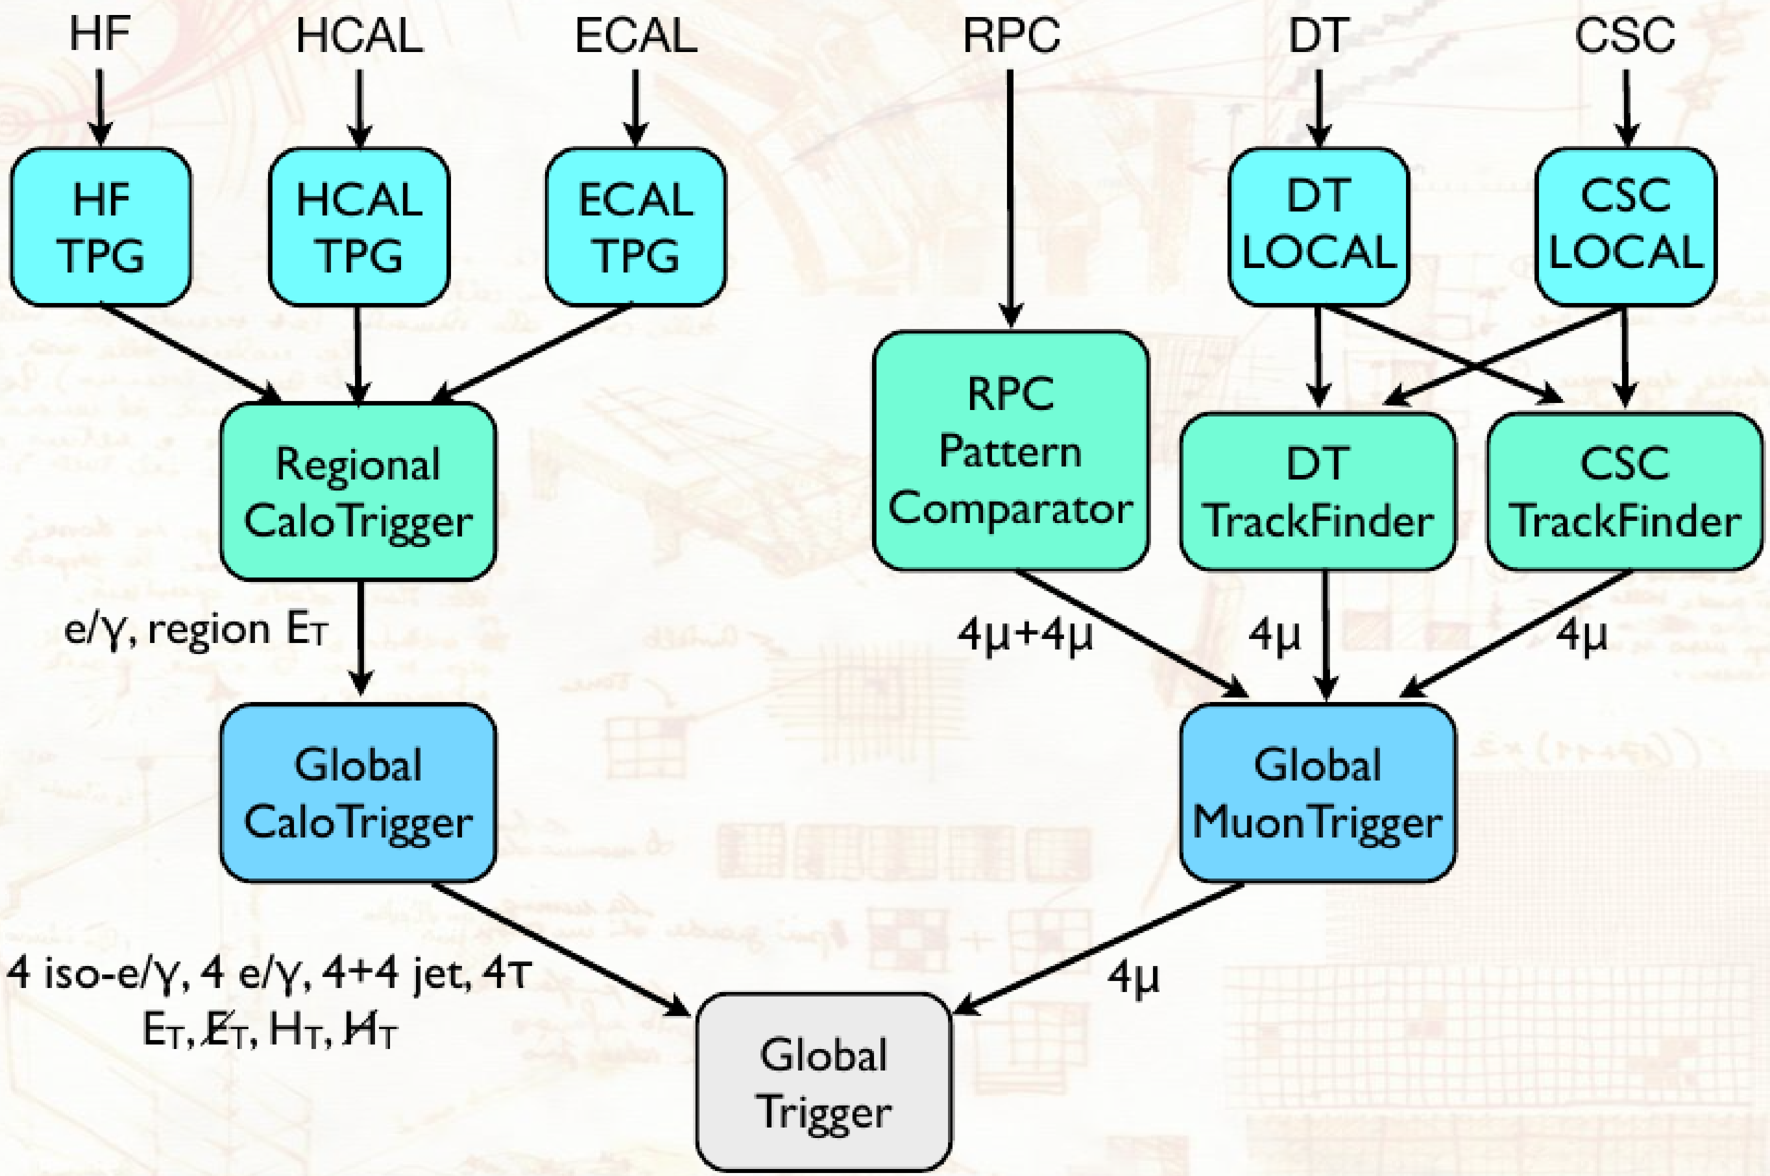

HF
HCAL
ECAL
RPC
DT
CSC
HF TPG
HCAL TPG
ECAL TPG
DT LOCAL
CSC LOCAL
Regional CaloTrigger
RPC Pattern Comparator
DT TrackFinder
CSC TrackFinder
e/γ, region ET
4μ+4μ
4μ
4μ
Global CaloTrigger
Global MuonTrigger
4 iso-e/γ, 4 e/γ, 4+4 jet, 4τ
ET, ET, HT, HT
4μ
Global Trigger

2015

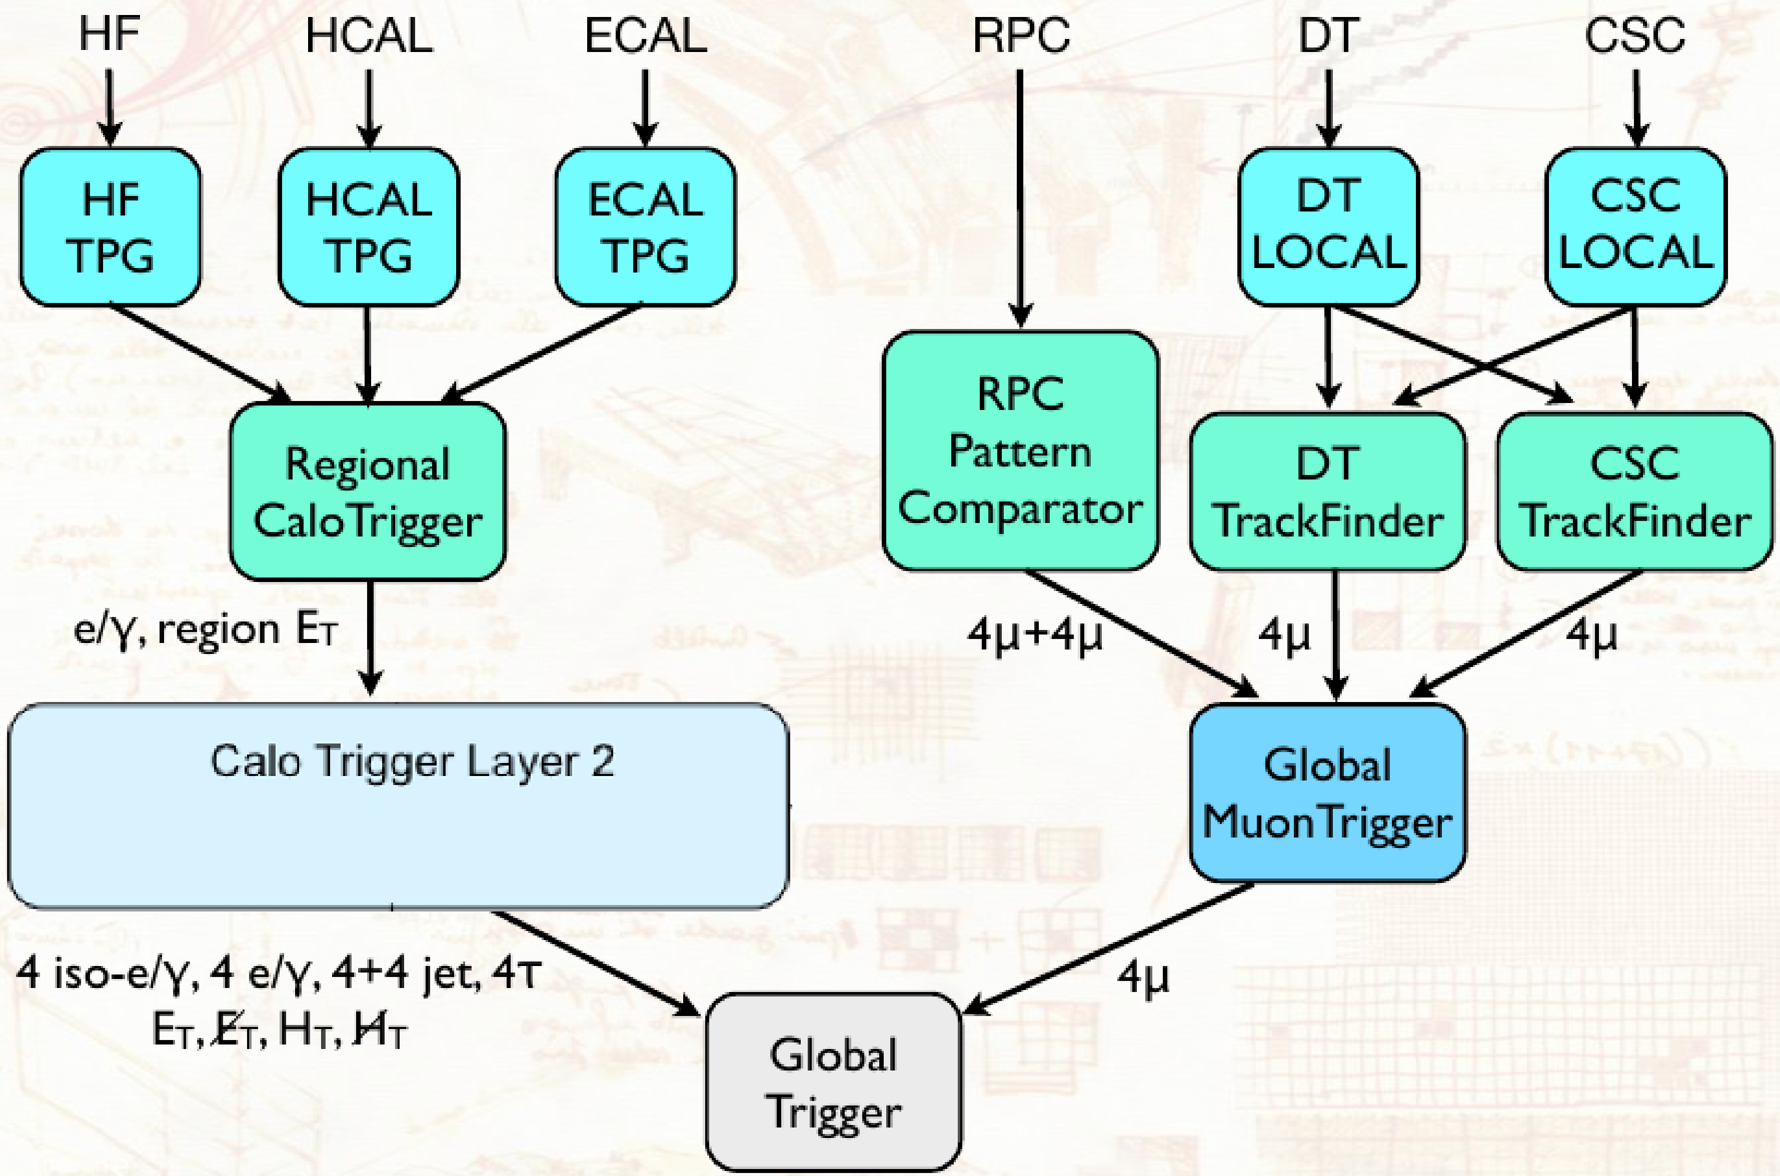
HF
HCAL
ECAL
RPC
DT
CSC
HF TPG
HCAL TPG
ECAL TPG
DT LOCAL
CSC LOCAL
Regional CaloTrigger
RPC Pattern Comparator
DT TrackFinder
CSC TrackFinder
e/γ, region ET
4μ+4μ
4μ
4μ
Calo Trigger Layer 2
Global MuonTrigger
4 iso-e/γ, 4 e/γ, 4+4 jet, 4τ
ET, ET, HT, HT
4μ
Global Trigger

2016 – ...

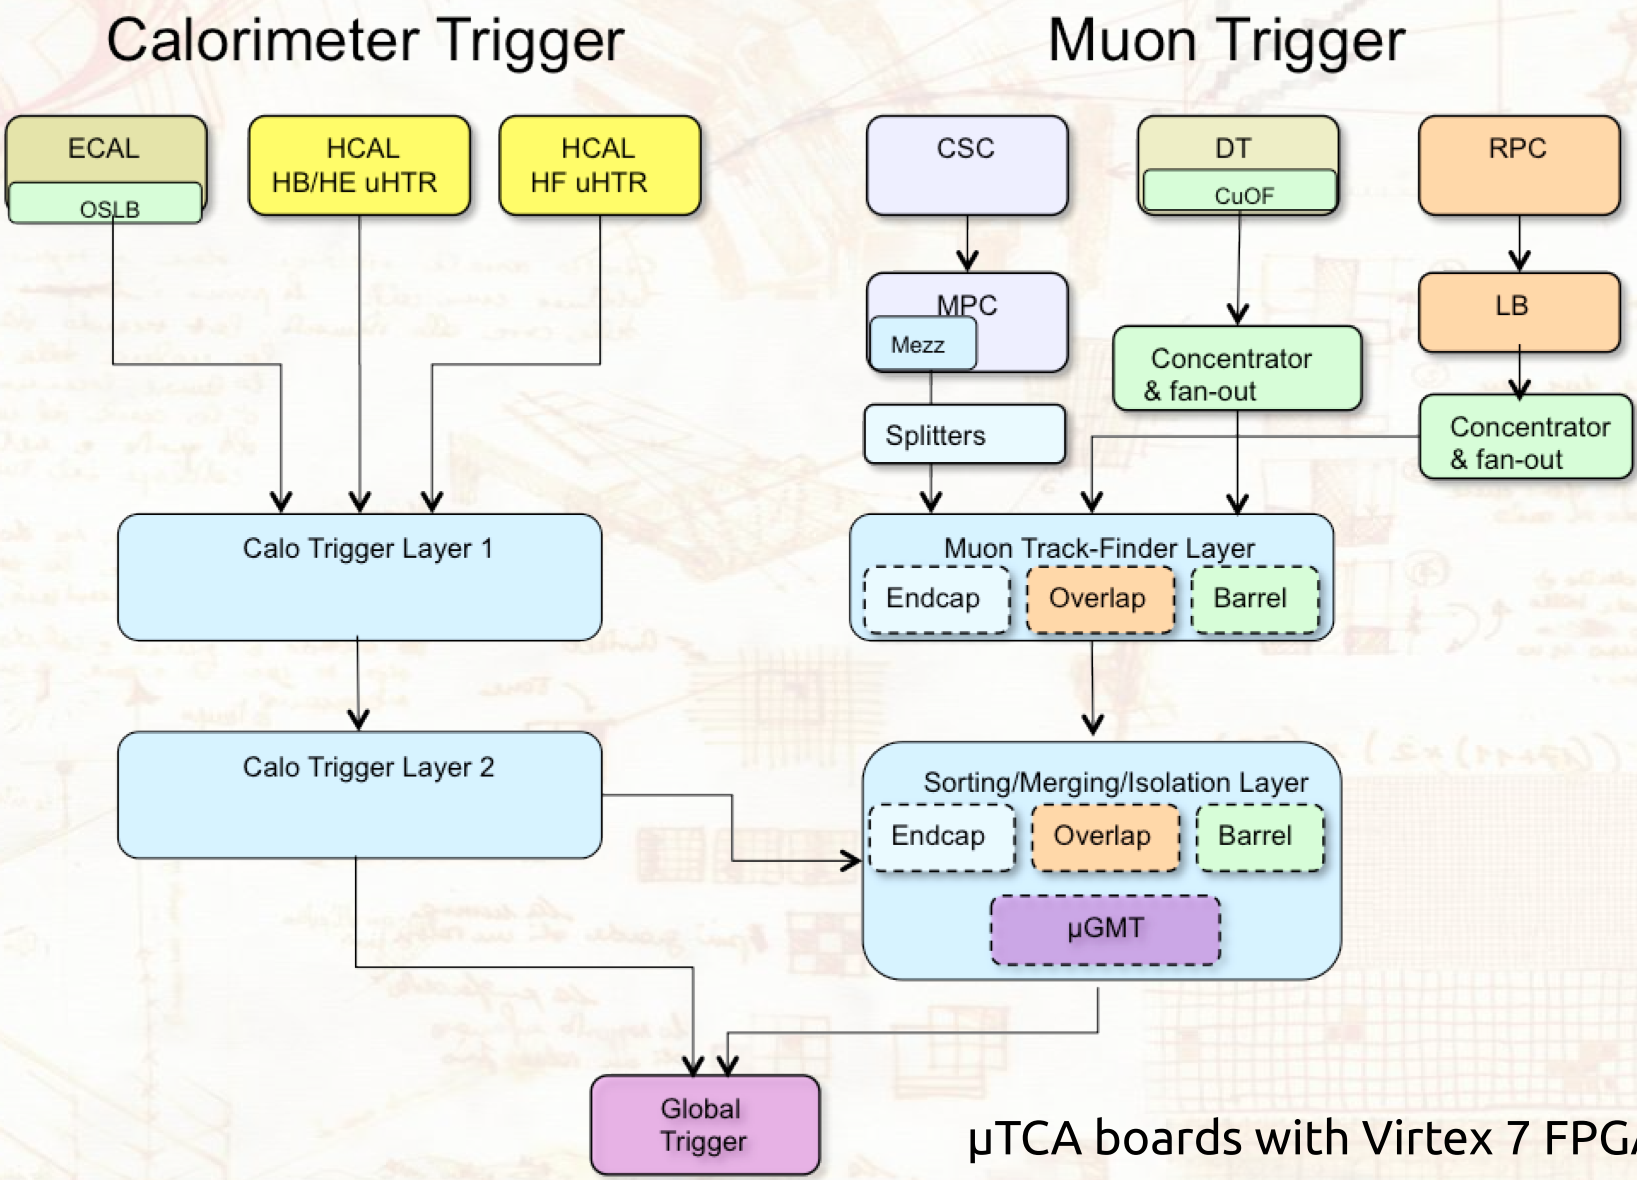

μTCA boards with Virtex 7 FPGAs

L1 Muon Trigger

- **Barrel Muon Track Finder**
  - track segments are identified in the DT
  - eventually, add RPC information
- **Overlap Muon Track Finder**
  - combine in formation from DT, RPC, CSC
- **Endcap Muon Track Finder**
  - track segments are identified in the CSC
  - eventually, add RPC information
- **all Track Finders**
  - assign η, φ, p_T and quality to each candidate
- **μGMT – μTCA Global Muon Trigger**
  - use this information to remove duplicates and combine the candidates from the Barrel, Overlap and Endcap track finders
  - each candidate is assigned η, φ, p_T and quality
  - select **8 leading muon candidates**
    - high quality candidates used for single muon triggers
    - different quality criteria for di-muon and cross-triggers

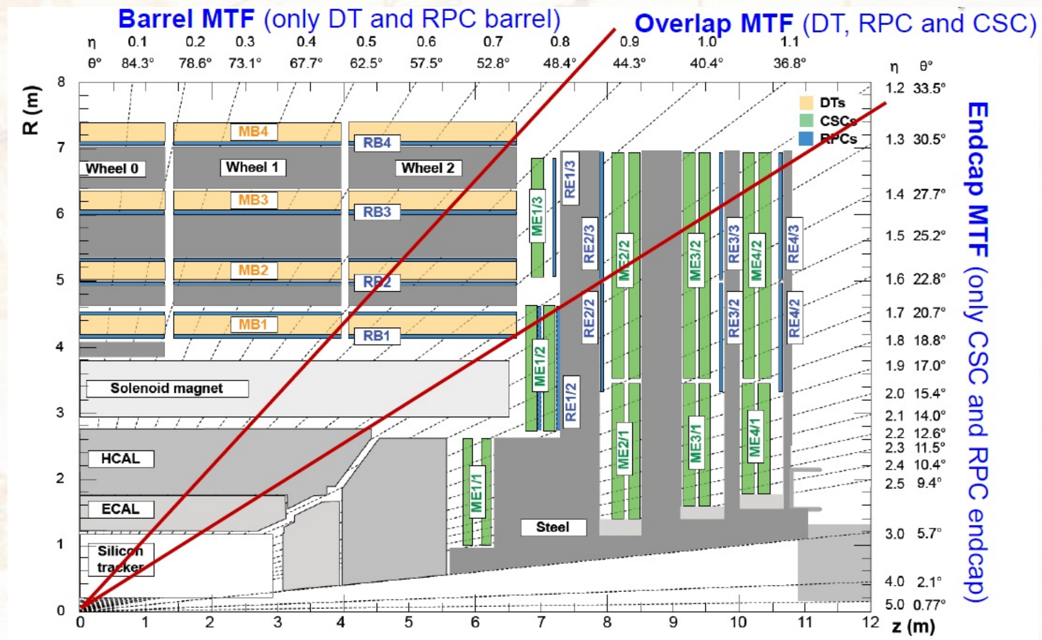

L1 Muon Trigger

rate reduction by a **factor 2 ~ 3**, with a **similar** efficiency

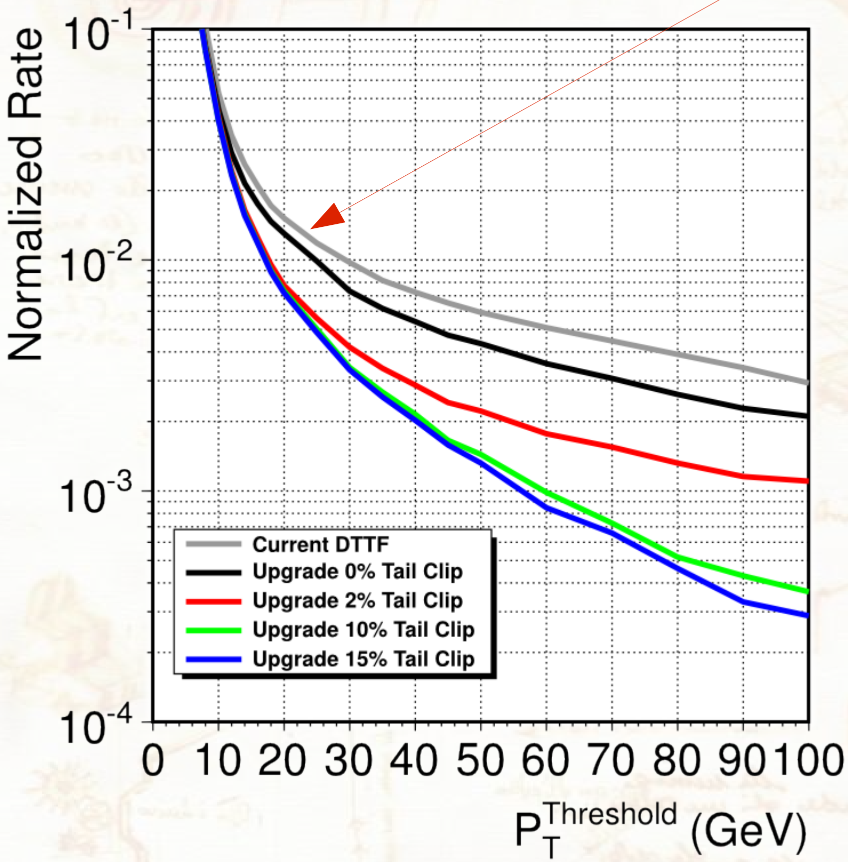

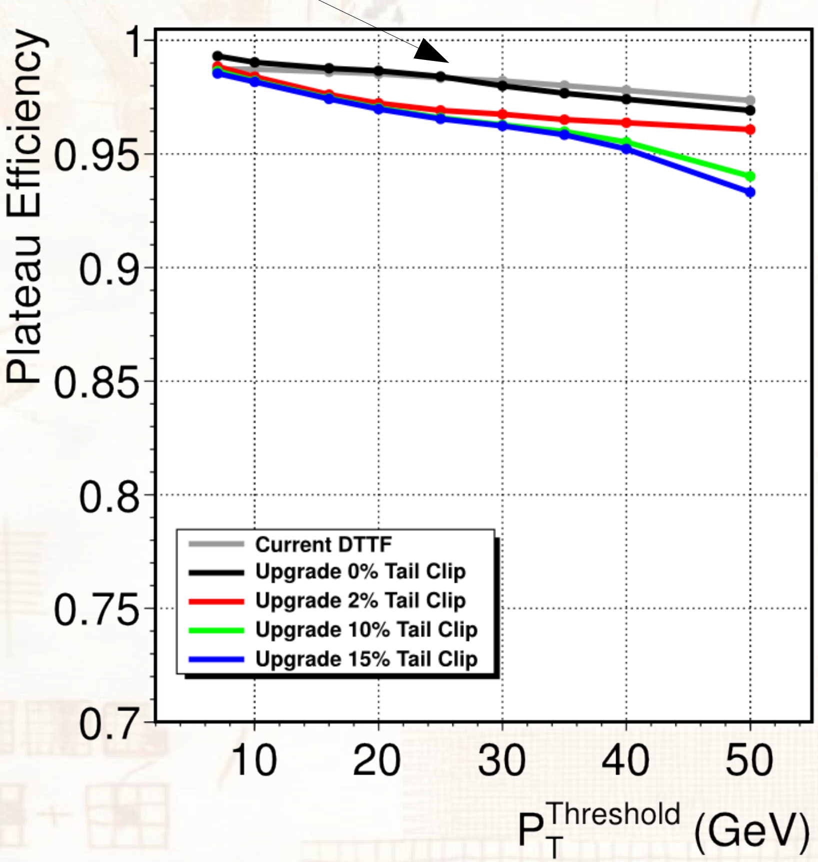

new muon pT assignment (bigger LUTs, post-processing)

L1 Calo Trigger

- **Layer 1**
  - combines inputs from ECAL (5x5 crystals) and HCAL into "trigger towers"
  - applies position- and energy-dependent calibrations

- **Layer 2**
  - pattern recognition: finds Jets, E/Gamma and Tau candidates
  - applies pileup subtraction and computes isolation
  - computes global quantities: ET / MET, HT / MHT

L1 Jets

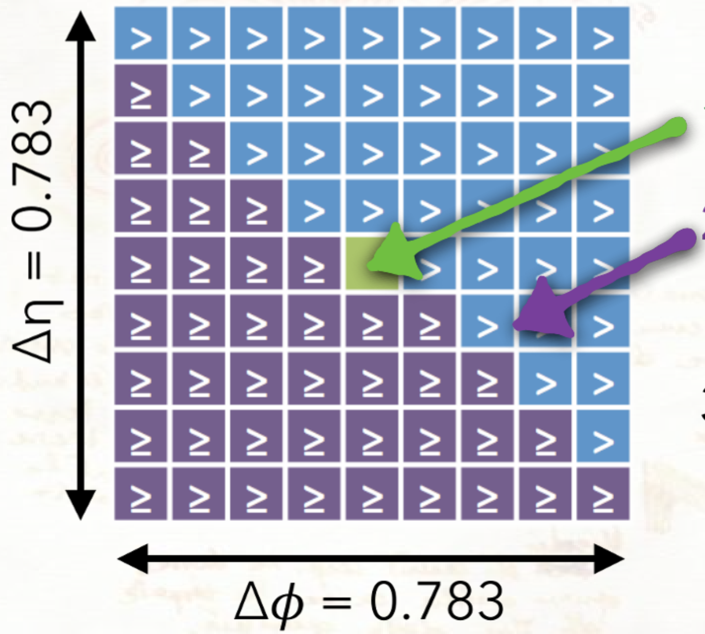

1. Look for local maximum above jet seed threshold
2. Apply mask to surrounding tower (antisymmetric to avoid double-counting of overlapping jets)
3. Jet $E_T = \sum$ TT in 9×9, jet centre = highest energy TT

Local correction to jet using energy in TT strips surrounding jet area:

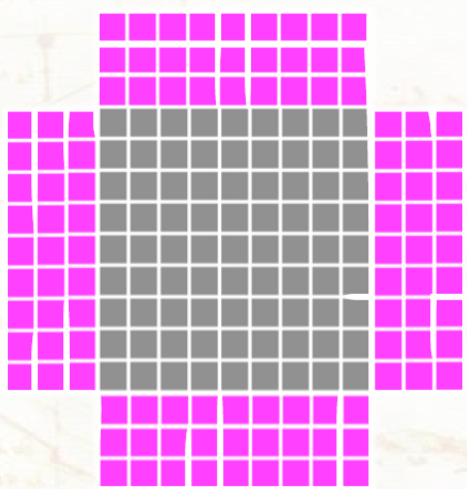

Subtract total energy in **lowest** 3 sides - avoids overcorrection from overlapping jets

Use 3 strips to mitigate fluctuations, but keeping correction local and robust

L1 E/Gamma

- discrimination between jets and e/gamma candidates
  - H/E (ratio HCAL / ECAL energy)
  - cluster shape

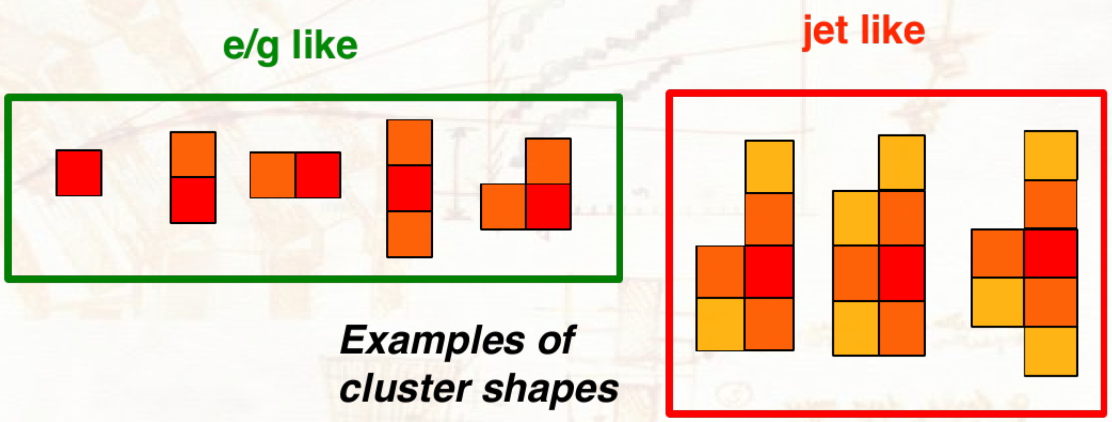

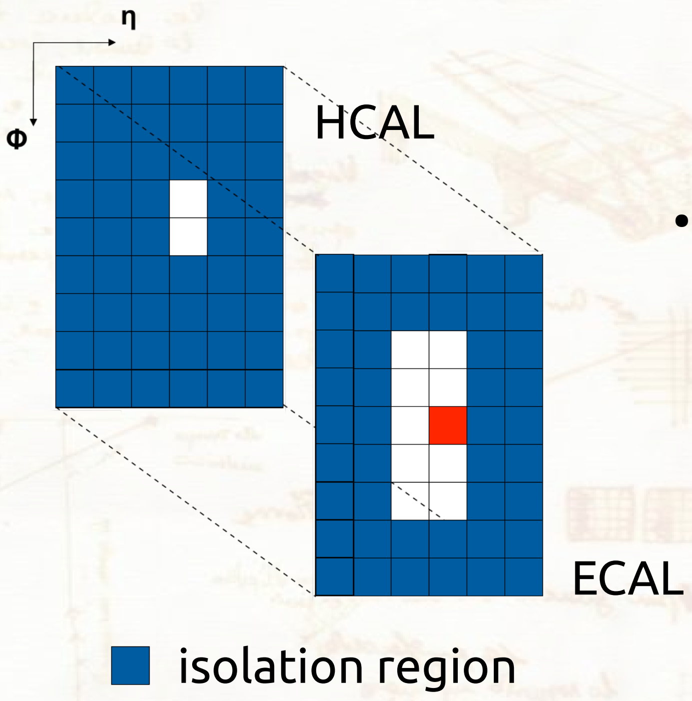

- isolation
  - energy in a 6×9 region
  - $E+H_{6x9} - E_{2x5} - H_{1x2}$ < isolation cut
  - function of position and pileup

- estimate pileup from the number of trigger towers above threshold

L1 Taus

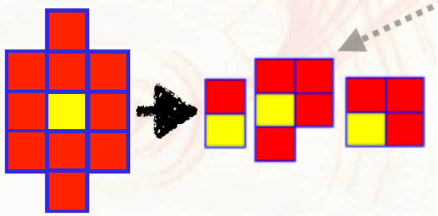

Clustering

- TT are regrouped into clusters
- Basic object to form L1 tau candidates

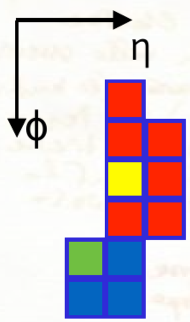

Merging

- Search for neighbors in a defined path (tau decay products can be spread out)
- ~**15%** merged, **85%** non merged

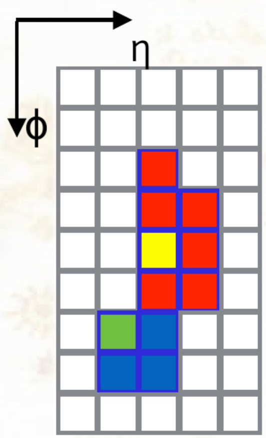

Isolation

- Compute isolation as E(6x9) - E (tau)
- **PU subtraction at hardware level**

Calibration

- Improves energy resolution

Shape veto

- Cluster with certain shapes are rejected
- Additional background rejection

L1 Global Trigger

- receives candidates from the Muon and Calo triggers
  - 8 muon candidates
  - 12 e/gamma candidates (isolated and non-isolated)
  - 12 jets
  - 8 tau candidates (isolated and non-isolated)
  - energy sums (ET/MET, HT/MHT, …)
  - external conditions (e.g. beam coincidence or anticoincidence, …)

- combines the candidates in up to 512 algorithms
  - single and multiple-object triggers (single muon, di-jets, …)
  - cross-triggers (muon + EG, …)
  - η, φ, p_T cuts on each candidate
  - $\Delta\eta$, $\Delta\varphi$, ΔR topological correlations among candidates (muon within a jet, …)
  - invariant mass cuts among different candidates

L1 Global Trigger

- receives candidates from the Muon and Calo triggers
  - 8 muon candidates
  - 12 e/gamma candidates (isolated and non-isolated)
  - 12 jets
  - 8 tau candidates (isolated and non-isolated)

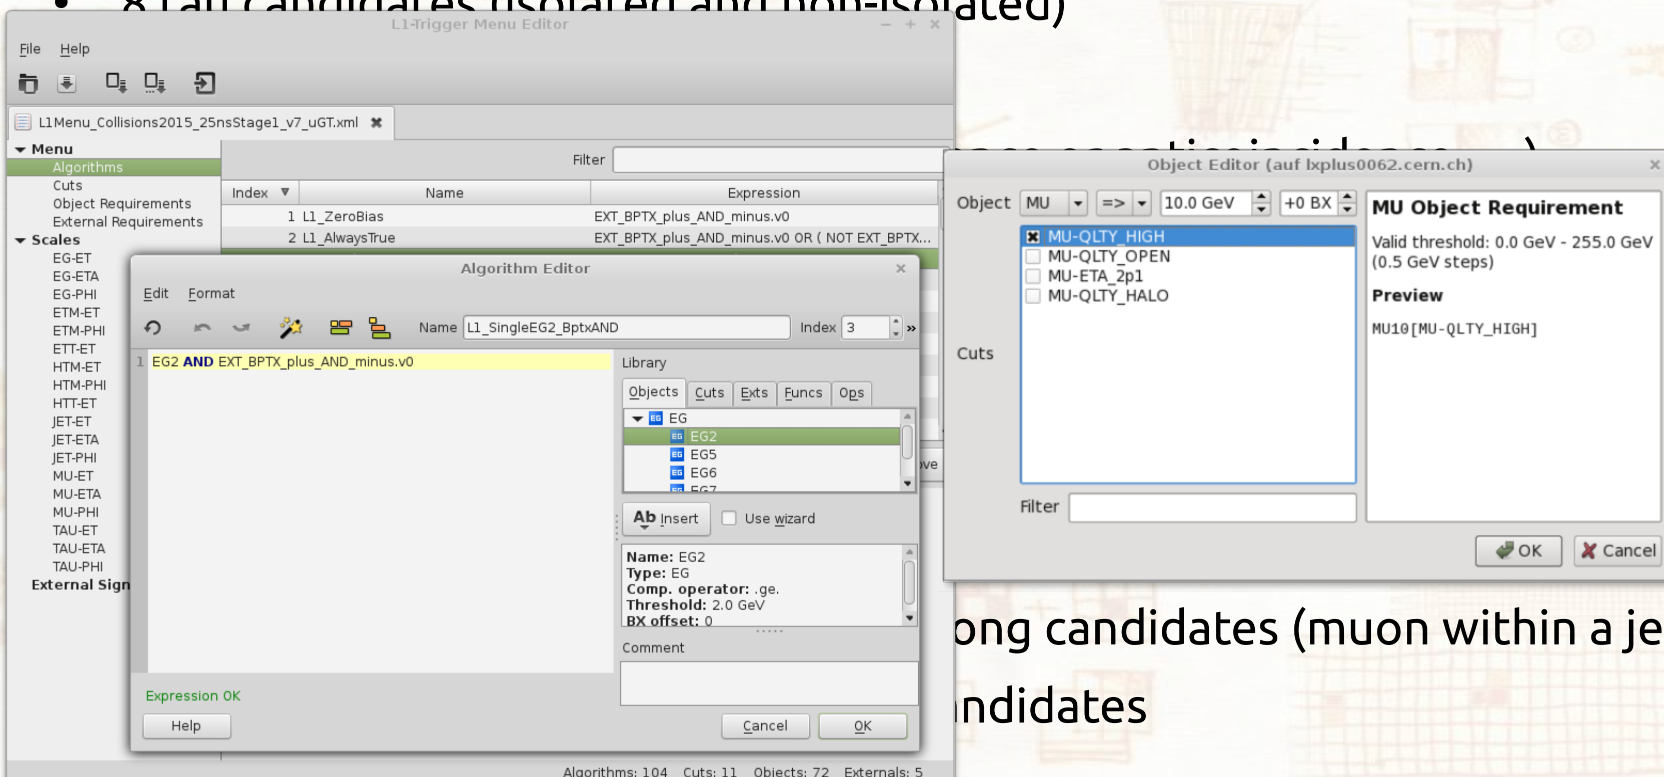

- …
  - …ong candidates (muon within a jet, …)
  - …ndidates

- python/Qt4 Trigger Menu Editor to implement the various algorithms

(some) L1 Algorithms used in 2015

- L1_SingleMu16
- L1_DoubleMu_10_3p5
- L1_DoubleMu0er16_WdEta18
- L1_DoubleMu_10_0_WdEta18
- L1_TripleMu0, L1_TripleMu_5_5_3
- L1_QuadMu0

- L1_SingleJet128
- L1_DoubleJetC84
- L1_TripleJet_84_68_48_VBF
- L1_QuadJetC40

- L1_ZeroBias

- L1_SingleMuOpen_NotBptxOR
- L1_SingleJetC32_NotBptxOR

- L1_SingleEG20
- L1_SingleIsoEG18er
- L1_DoubleEG_15_10
- L1_TripleEG_14_10_8

- L1_DoubleIsoTau28er
- L1_DoubleTau40er
- L1_Mu16er_TauJet20er
- L1_IsoEG20er_TauJet20er_NotWdEta0

- L1_Mu4_EG18, L1_Mu5_EG15, L1_Mu12_EG10
- L1_Mu5_DoubleEG5, L1_DoubleMu6_EG6

- L1_HTT100
- L1_ETM50
- L1_Mu0er_ETM40, L1_Mu10er_ETM30

L1 rates

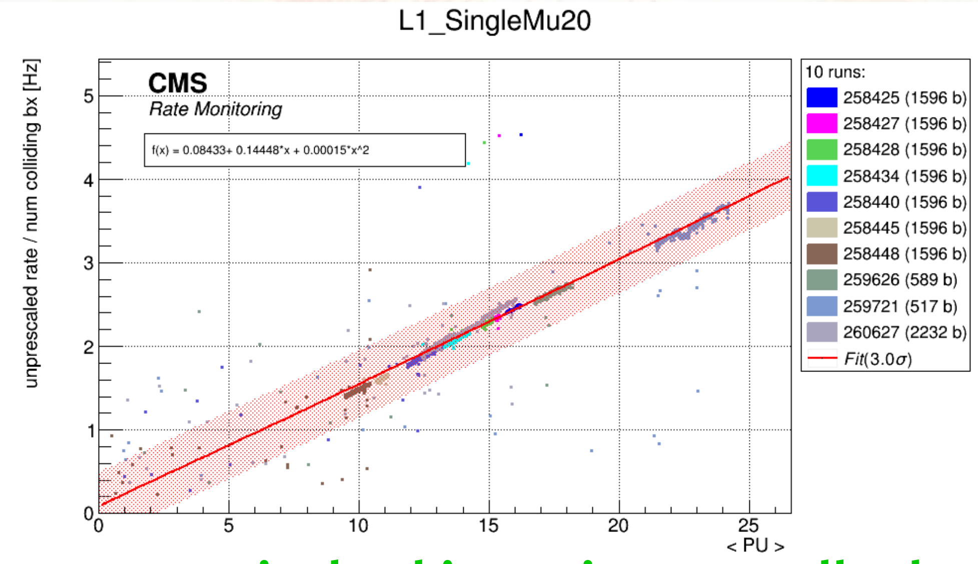

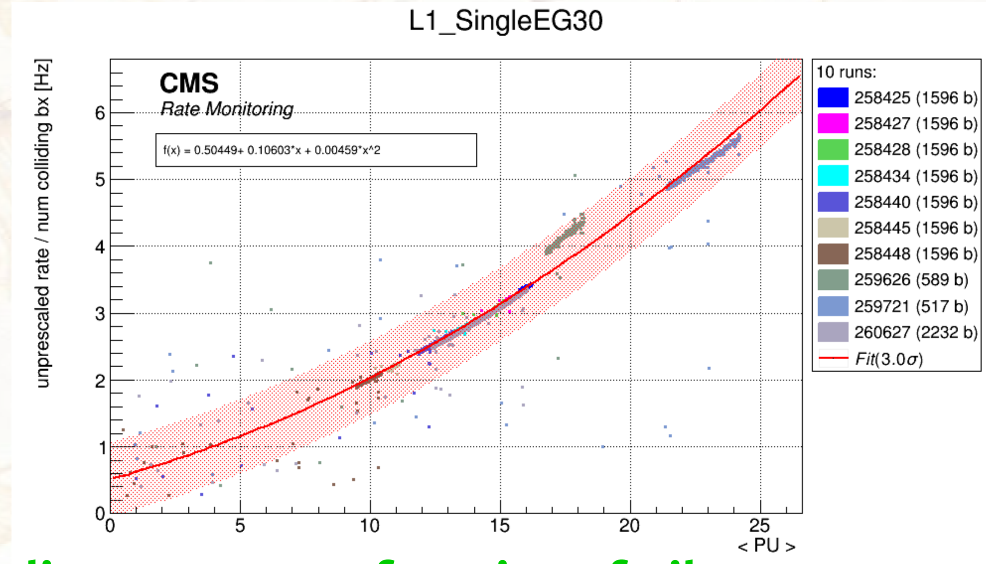

single object trigger usually show a linear rate as a function of pileup

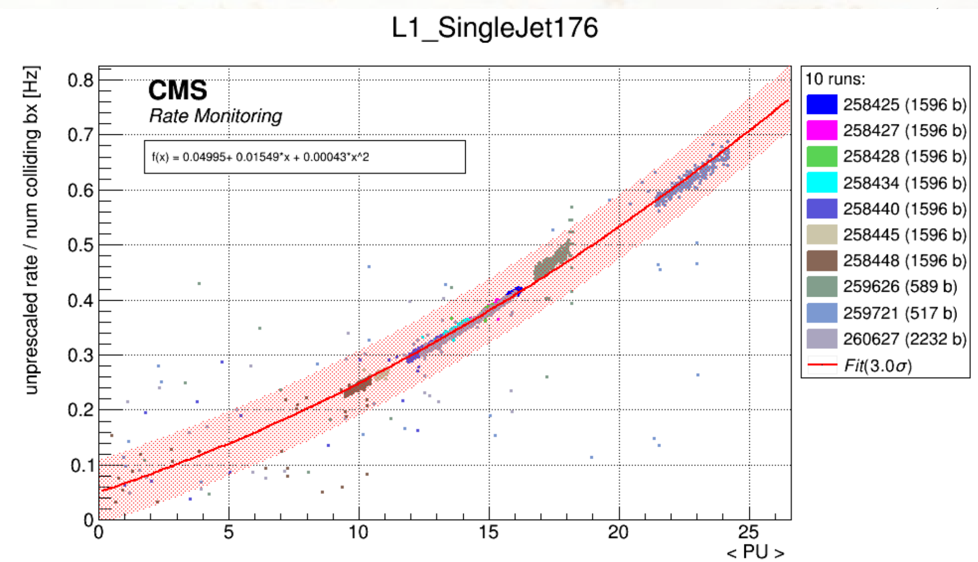

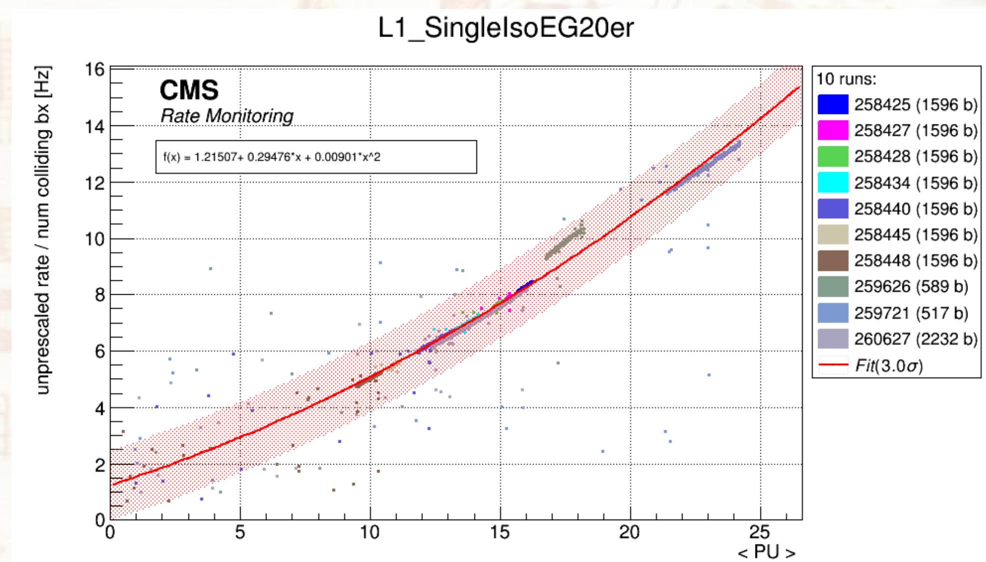

L1 rates

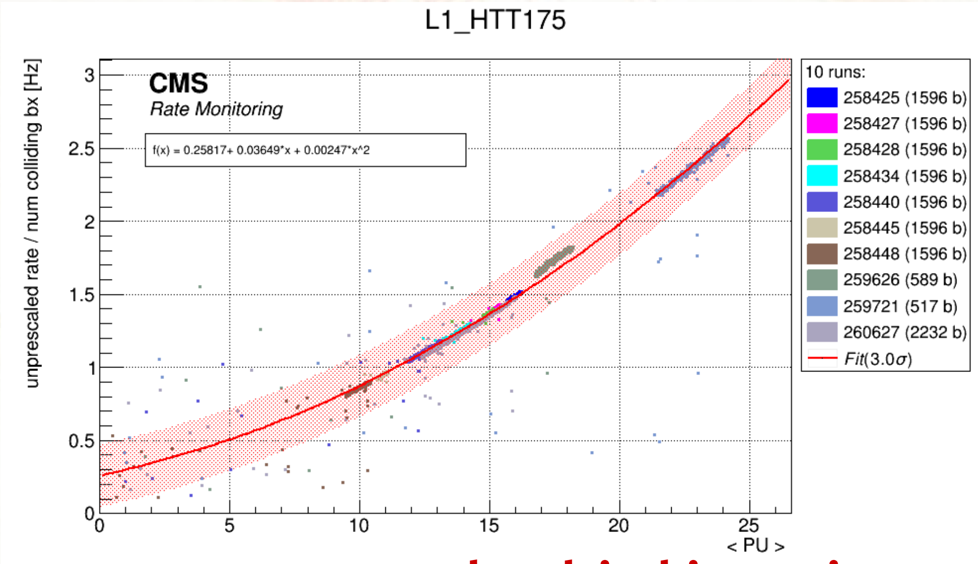

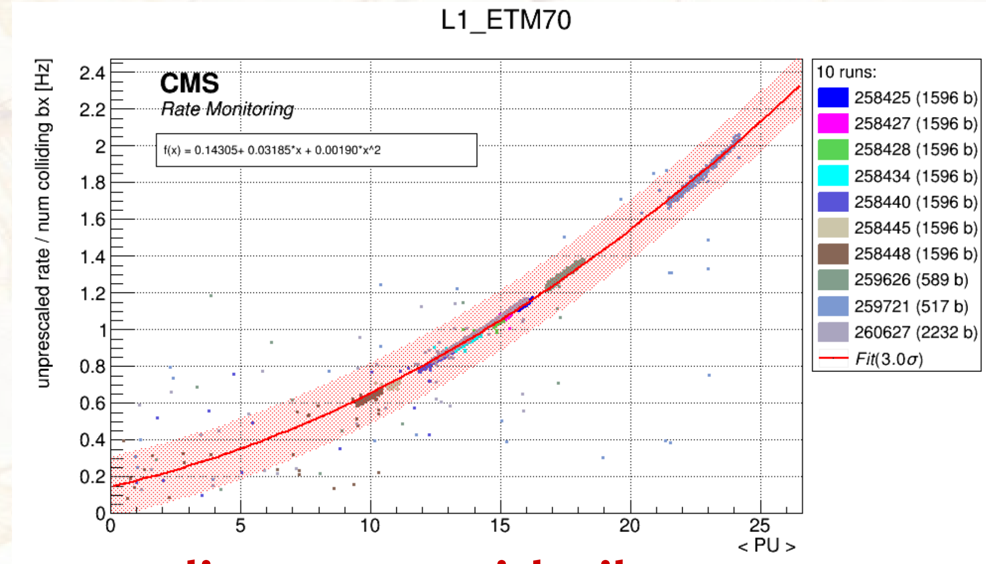

sums and multi-object triggers show non-linear rates with pileup

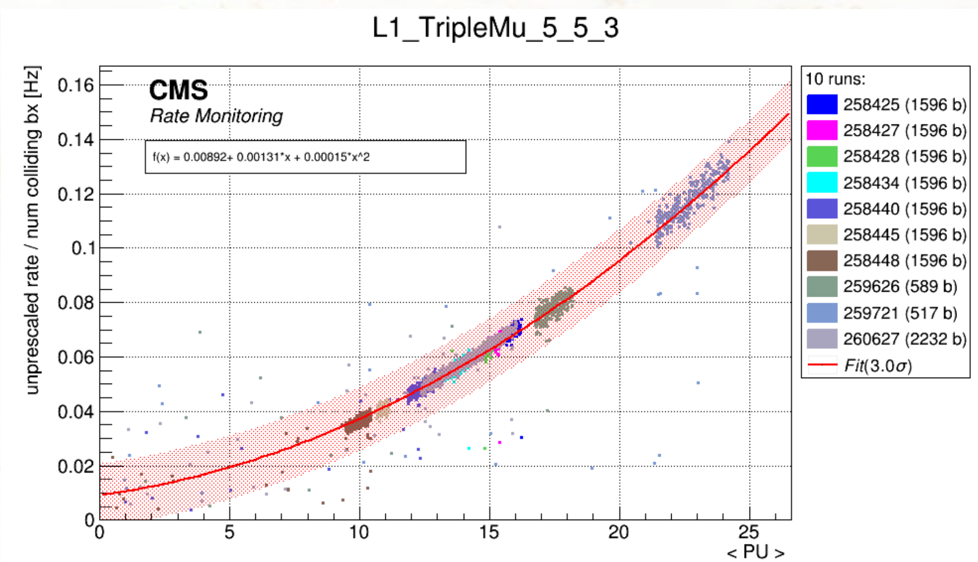

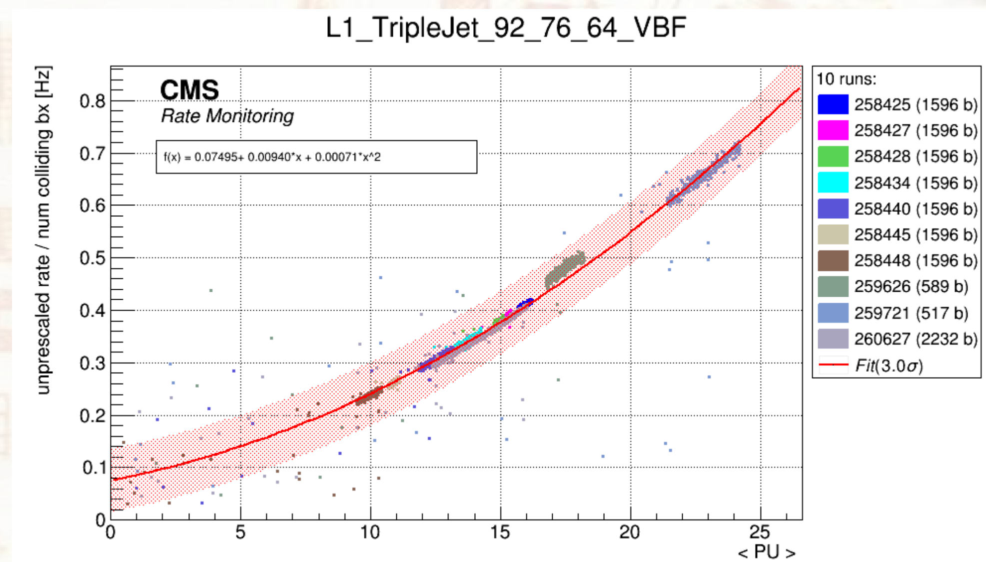

Data Acquisition

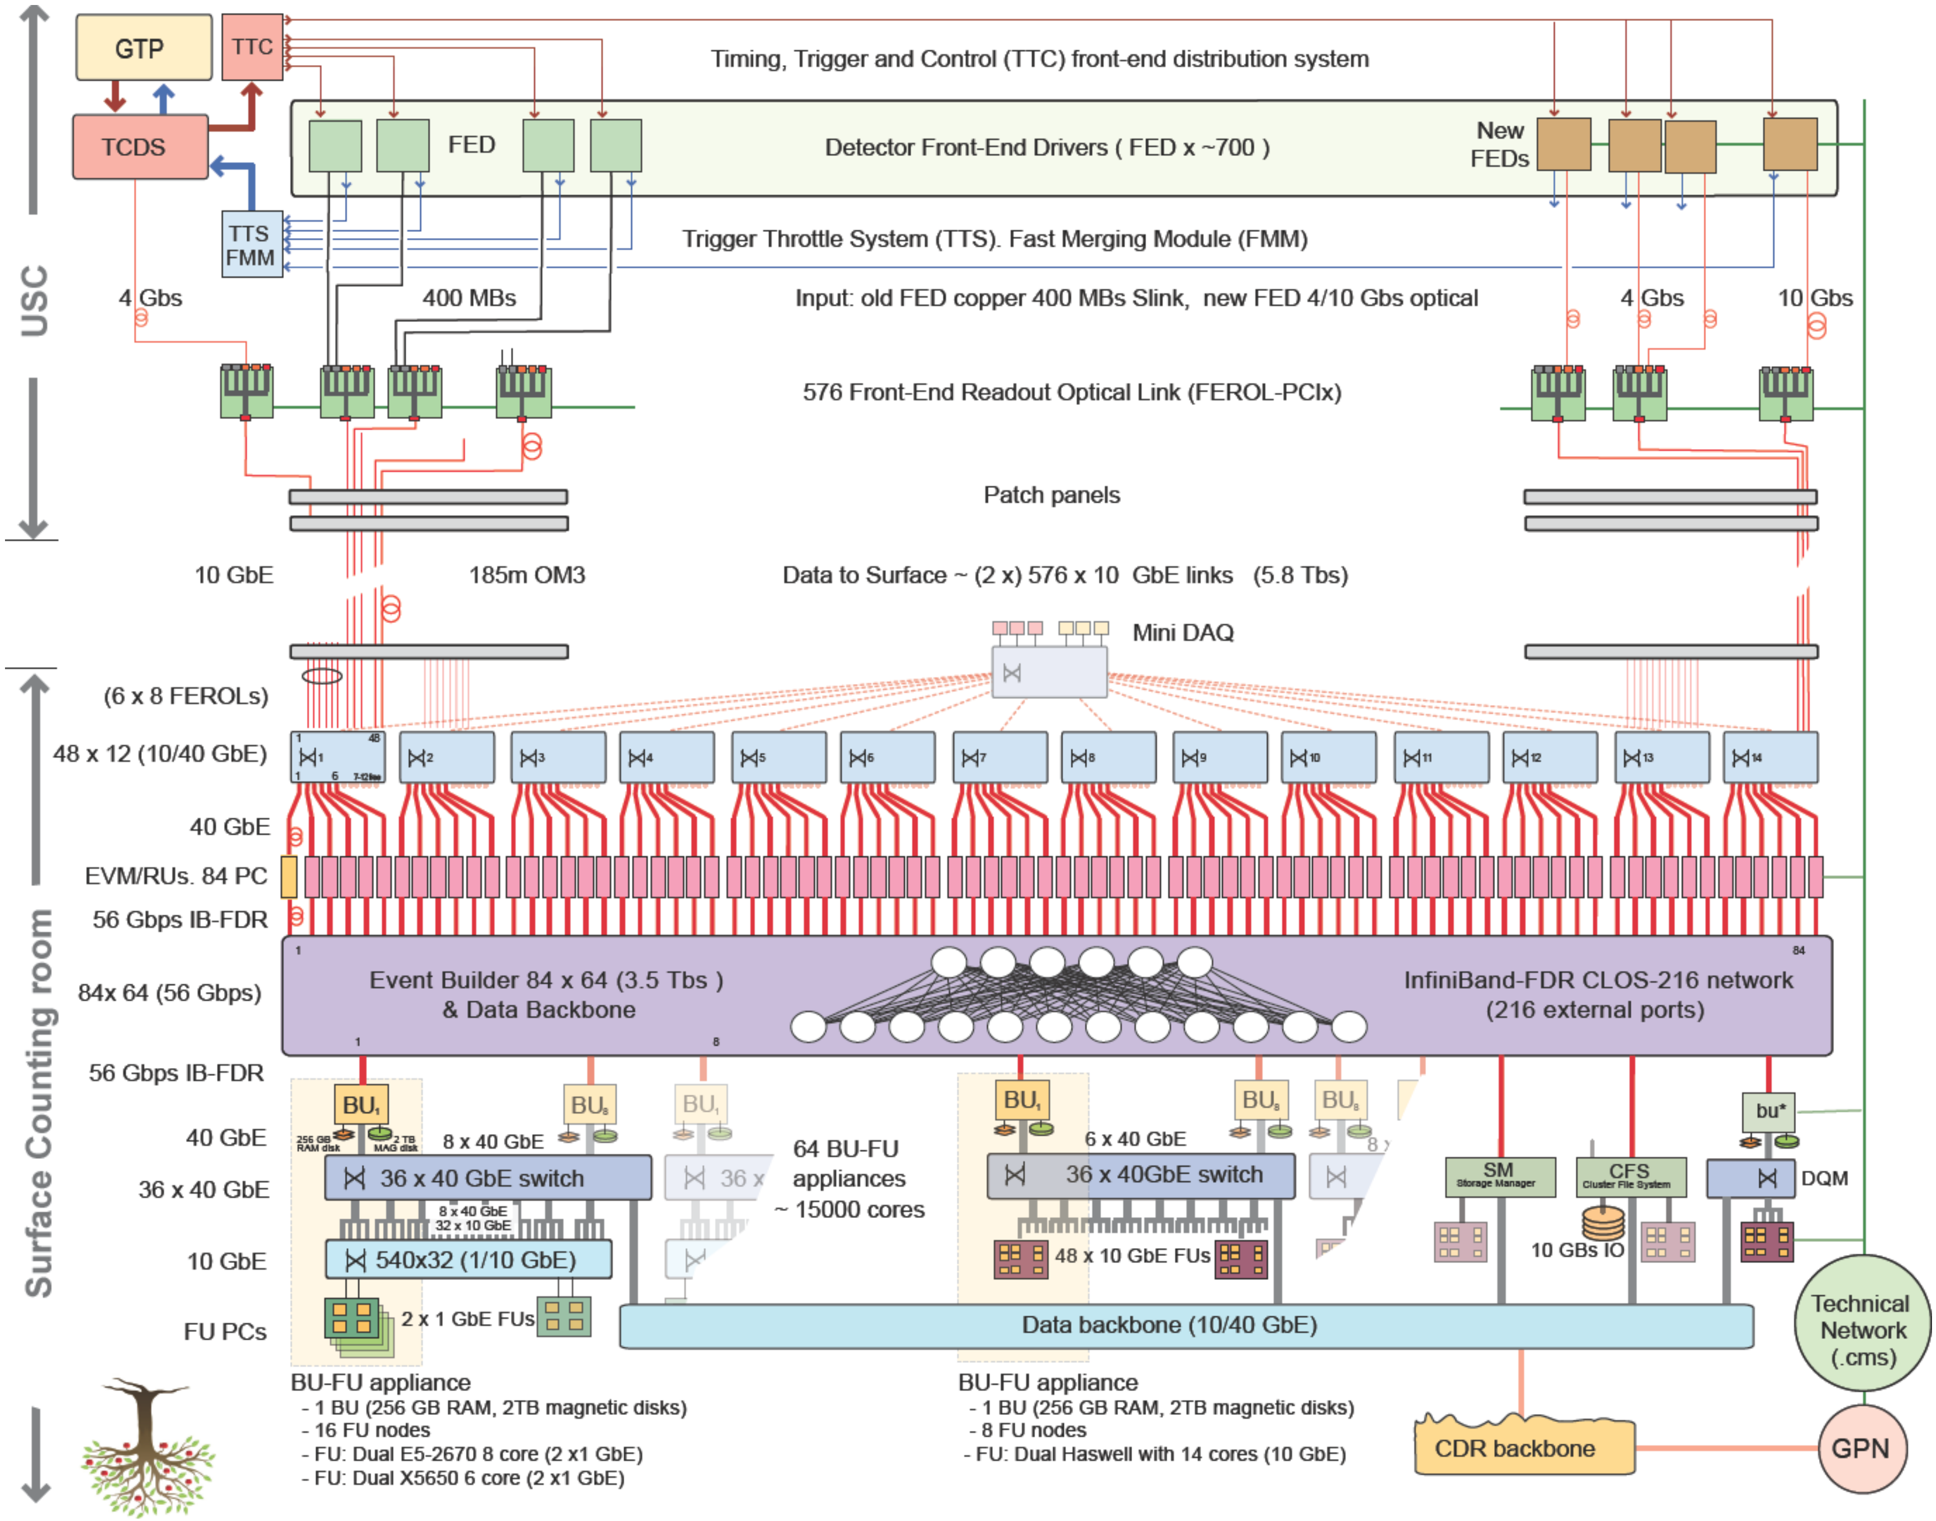

GTP
TTC
TCDS
Timing, Trigger and Control (TTC) front-end distribution system
FED
Detector Front-End Drivers (FED x ~700)
New FEDs
TTS FMM
Trigger Throttle System (TTS). Fast Merging Module (FMM)
USC
4 Gbs
400 MBs
Input: old FED copper 400 MBs Slink, new FED 4/10 Gbs optical
4 Gbs
10 Gbs
576 Front-End Readout Optical Link (FEROL-PCIx)
Patch panels
10 GbE
185m OM3
Data to Surface ~ (2 x) 576 x 10 GbE links (5.8 Tbs)
Mini DAQ
(6 x 8 FEROLs)
48 x 12 (10/40 GbE)
40 GbE
EVM/RUs. 84 PC
56 Gbps IB-FDR
84x 64 (56 Gbps)
Event Builder 84 x 64 (3.5 Tbs) & Data Backbone
InfiniBand-FDR CLOS-216 network (216 external ports)
56 Gbps IB-FDR
Surface Counting room
40 GbE
36 x 40 GbE
10 GbE
FU PCs
8 x 40 GbE
36 x 40 GbE switch
540x32 (1/10 GbE)
2 x 1 GbE FUs
64 BU-FU appliances
~ 15000 cores
6 x 40 GbE
36 x 40GbE switch
48 x 10 GbE FUs
SM
CFS
10 GBs IO
DQM
Data backbone (10/40 GbE)
Technical Network (.cms)
BU-FU appliance
- 1 BU (256 GB RAM, 2TB magnetic disks)
- 16 FU nodes
- FU: Dual E5-2670 8 core (2 x1 GbE)
- FU: Dual X5650 6 core (2 x1 GbE)
BU-FU appliance
- 1 BU (256 GB RAM, 2TB magnetic disks)
- 8 FU nodes
- FU: Dual Haswell with 14 cores (10 GbE)
CDR backbone
GPN

the Filter Farm

- each HLT node in the Filter Farm runs an “HLTd” service
  - wait for the start of a new run
  - spawn and monitor the health of the HLT jobs

- in 2015, the HLT farm was composed by three type of computers
  - **288** dual Westmere Xeon processor, 2x 6 cores (from 2011)
  - **256** dual Sandy Bridge Xeon processor, 2x 8 cores (from 2012)
  - **360** dual Haswell Xeon processor, 2x 12 cores (from 2015)

- **16192** cores
  - taking into account the different performance of the machines, this corresponds to **~14400** "Haswell-equivalent" cores
  - **~144 ms/ev.** running at an L1 rate of **100 kHz**
  - **~125 ms/ev.** scaled to the performance of the machine used for the benchmarks

- in 2016
  - replace the Westmere machines with brand new Broadwell machines
  - **324** dual Broadwell Xeon processor, 2x 14 cores
  - 21800 cores, **~22000** “Haswell-equivalent” cores
  - increase the farm processing power by 40%~50%

Storage Manager and Transfer System

- Storage Manager
  - collect the output "streams" from the filter farm
    - one file per stream per "luminosity section" per HLT job
  - low-latency, hierarchical merge of data into smaller number of files
    - micro-merger
      merge files on each HLT node
    - mini-merger
      merge files at the "builder unit" level
    - macro-merger
      merge files for the whole farm
    - additional metadata files insure the integrity of the data through the chain
  - based on a distributed Lustre FS
    - copy-less merge (aggregate fragments without reading/writing them back)

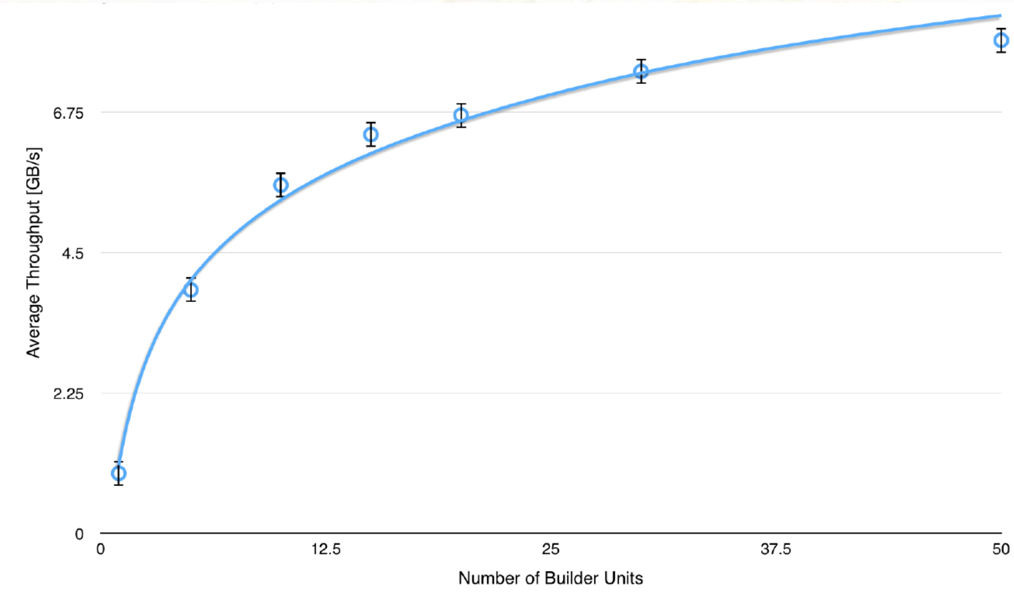

LFS bandwidth benchmarking

- Transfer System
  - transfer the fully merged files to the Tier-0 for offline processing and reconstruction

Storage Manager and Transfer System

- Storage Manager
  - collect the output "streams" from the filter farm
    - one file per stream per "luminosity section" per HLT job
  - low-latency, hierarchical merge of data into smaller number of files
    - micro-merger
      merge files on each HLT node
    - mini-merger
      merge files at the "builder unit" level
    - macro-merger
      merge files for the whole farm
    - additional metadata files insure the integrity of the data through the chain
  - based on a distributed Lustre FS
    - copy-less merge (aggregate fragments without reading/writing them back)

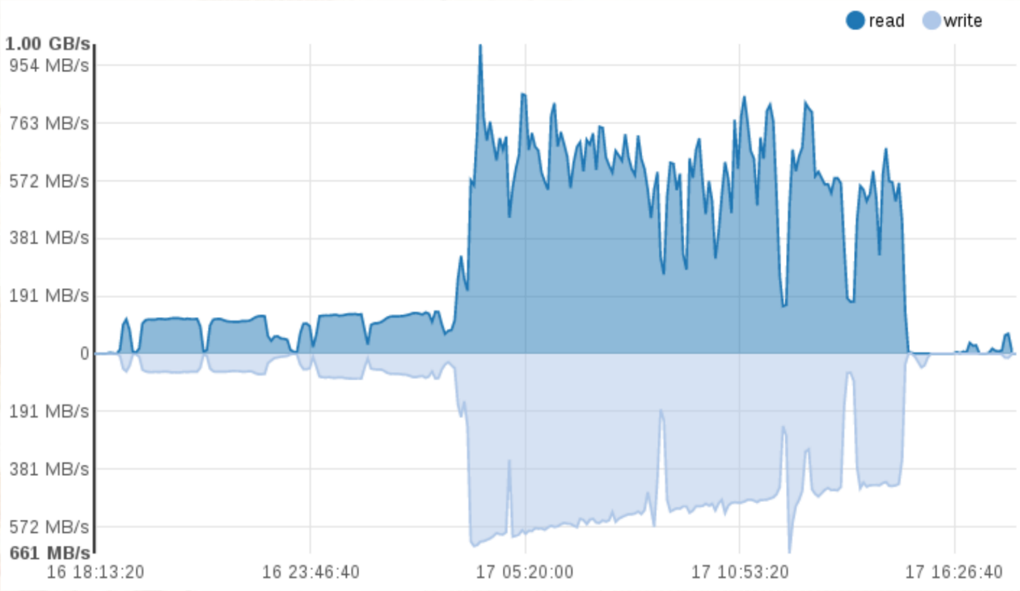

- Transfer System
  - transfer the fully merged files to the Tier-0 for offline processing and reconstruction

High Level Trigger

High Level Trigger

- full detector readout, full granularity

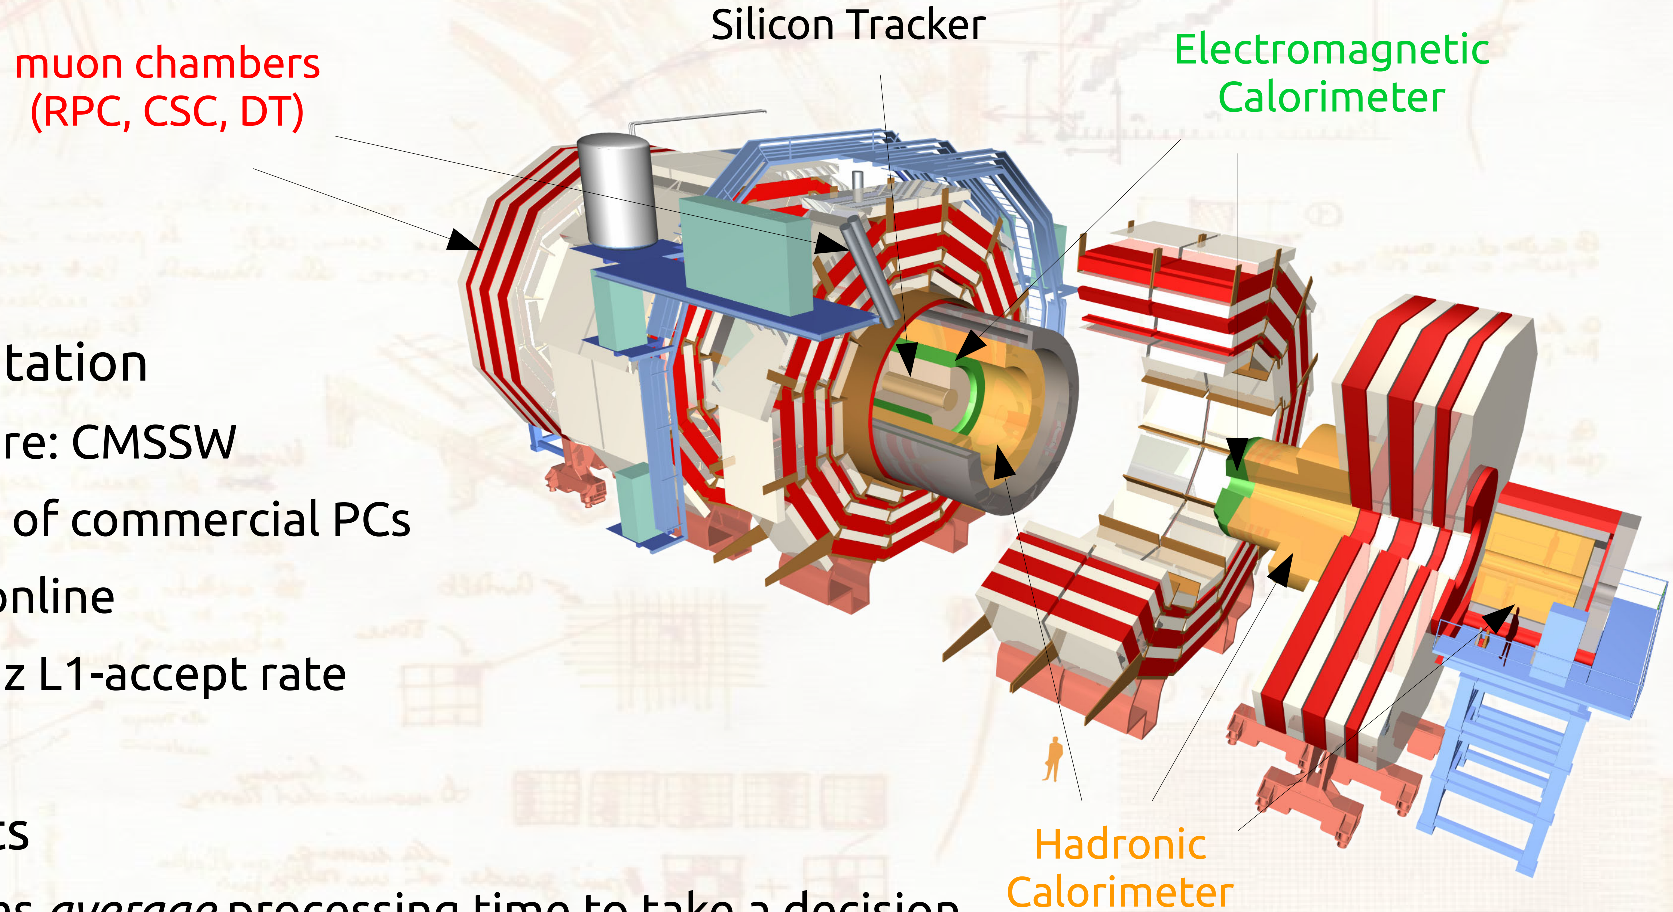

- implementation
  - software: CMSSW
  - cluster of commercial PCs
  - quasi-online
  - 100 kHz L1-accept rate
- constraints
  - ~200 ms *average* processing time to take a decision
  - 1 kHz *average* output rate (limited by offline resources)

CMSSW

- CMSSW software
  - written in C++, configured in Python
    - uses ROOT files (in special "EDM" format) for persistency
  - organised in modules, corresponding to C++ classes
    - data processing: Producers, Filters, Analyzers
    - organised in Sequences and Paths
    - Event Setup: access to conditions and calibrations, based on different IOVs
    - Services: tools, e.g. performance monitoring, logging
  - multi-threading
    - multiple events processed at the same time

- HLT uses same software as the Offline simulation, reconstruction and analysis
  - code and configuration optimised for speed
  - similar or same algorithms used for HLT and Reconstruction
  - calibrations close to those used in the Prompt Reconstruction

trigger "path"

CMSSW and the HLT

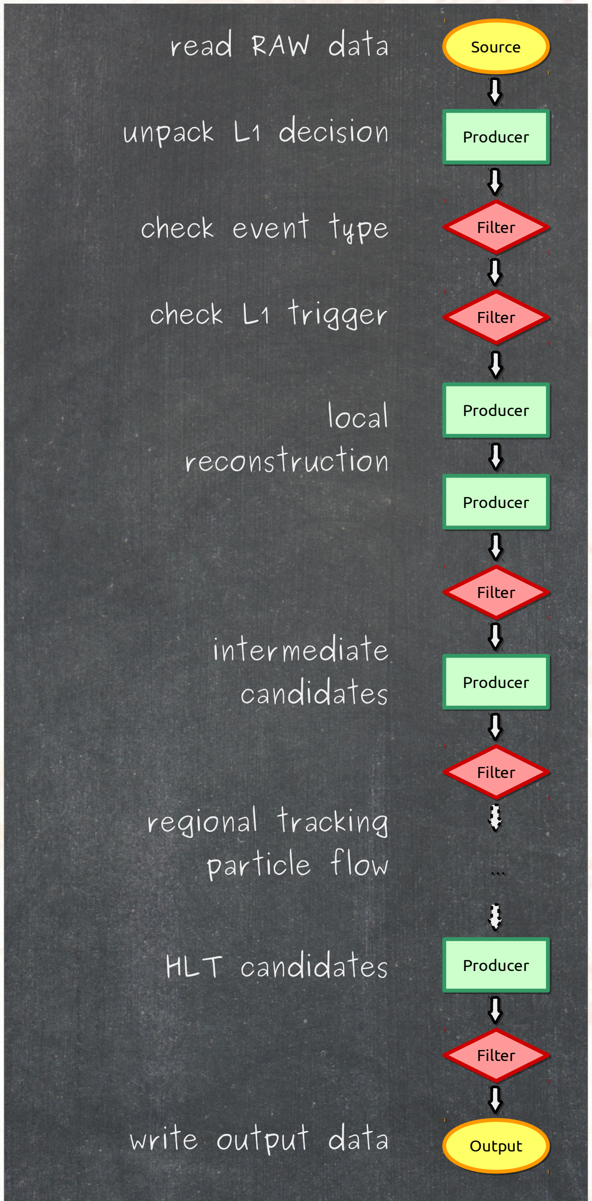

trigger “path”

- Trigger Paths
  - succession of Producers, Analyzers, and Filters
  - Filters can stop the execution along a Path
    - other Paths are not affected
    - logically parallel execution
  - stop each Path as early as possible
    - reject events not meeting the selection criteria
- what methods can we use to speed up the HLT ?
  - regional reconstruction
    - around L1 candidates
  - reject often, reject early
    - intermediate reconstruction steps
    - reject events as soon as possible
  - modularity and reuse of the reconstructed quantities
  - good enough reconstruction
    - trade large speed gains for small accuracy drops

HLT menu

single source:
RAW data
events selected by the L1 Trigger

trigger paths
run independently
of each other

common modules and sequences
are **shared**
across different paths

selected events are output to
different **data streams**
with different rates,
content and size

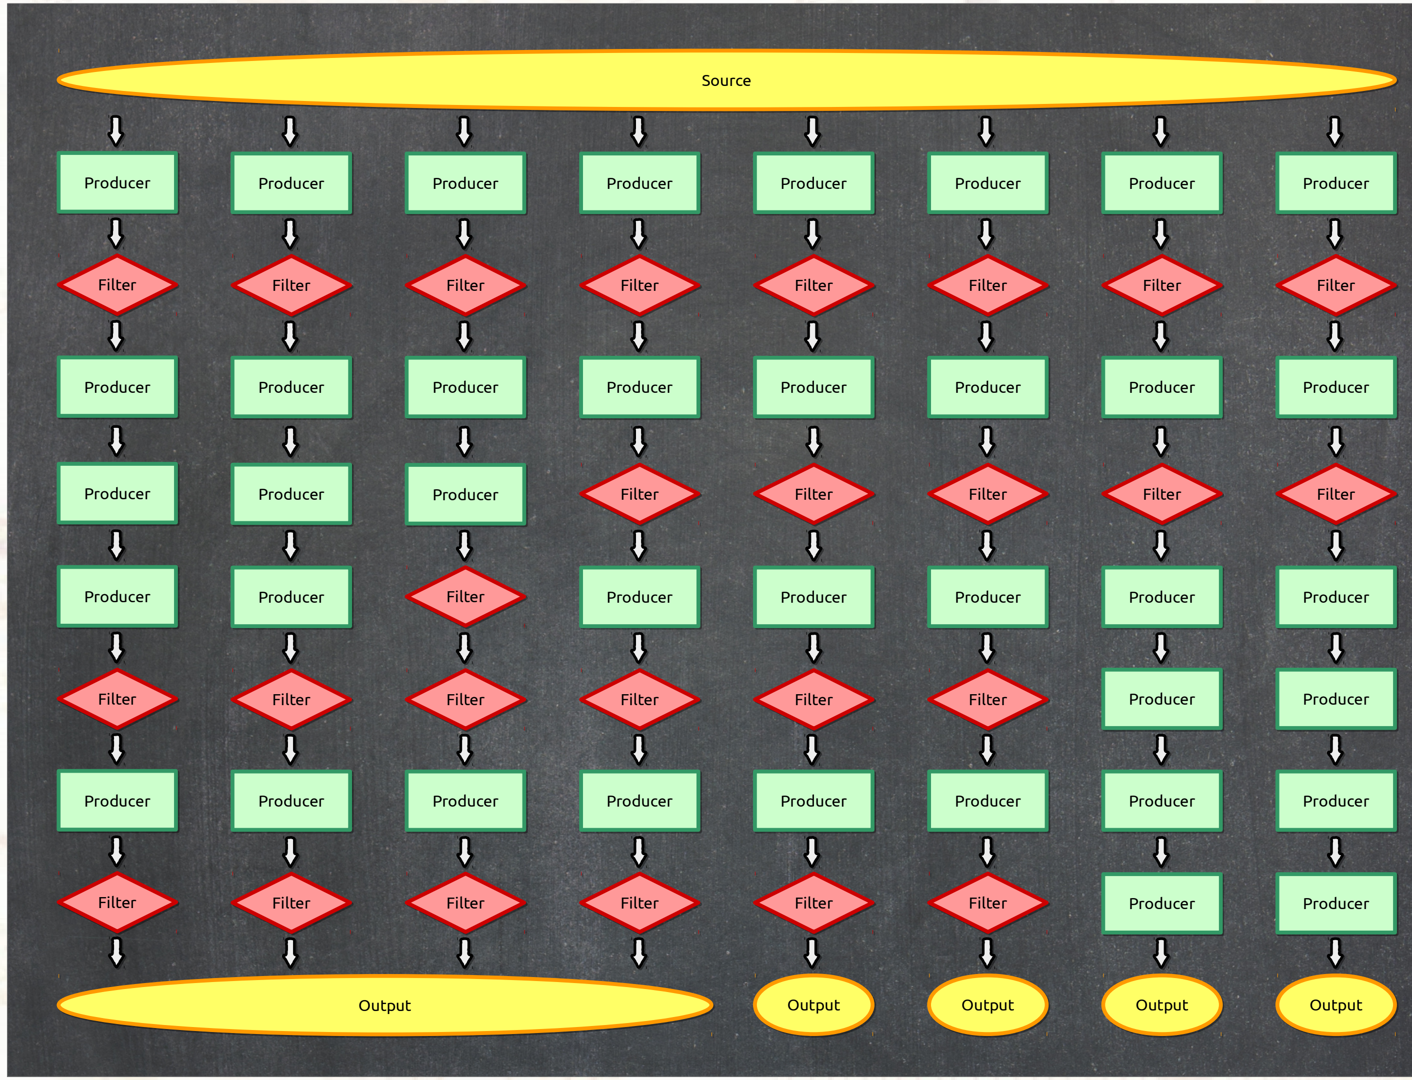

High Level Trigger objects

what can we reconstruct at HLT ?

- muons
  - “L2” stand alone muons
  - “L3” global and “tracker” muons
- photons
  - based on ECAL superclusters
  - calorimeter-based id
- electrons
  - based on ECAL superclusters, pixel tracks, and GSF tracking
  - calorimeter and track-based id
- general
  - particle-flow based isolation
  - pileup correction for isolation and jet energy
- taus
  - particle flow reconstruction
- jets, MET, HT
  - calorimteric jets and MET
  - particle flow-based jets and MET
  - pileup correction to Jets and HT energy
- b-tagging
  - secondary vertex reconstruction
  - soft-lepton based b-tagging
- but also
  - razor, α_T, dE/dx, …
  - jet substructure, …
  - displaced decays for non-prompt searches

the CMS HLT "menu" in 2015 - muons

Single muon (global or tracker-based)

- isolated muon, pT > 18 GeV
- isolated muon, η < 2.1, pT > 17 GeV + tau cross trigger
- isolated muon, η < 2.1, pT > 16 GeV + MET cross trigger

- non-isolated muon, η < 2.1, pT > 45 GeV
- non-isolated muon, pT > 50 GeV

Double muons (global or tracker-based)

- loosely isolated muons, pT > 17, 8 GeV
- non-isolated muons, pT > 17, 8 GeV, same sign
- non-isolated muons, pT > 27, 8 GeV

Triple muons

- triple muons, pT > 12, 10, 5 GeV

A whole family of muon triggers for "quarkonia" and B physics

Muon + E/Gamma cross triggers

- loosely isolated muon, pT > 8 GeV, loosely isolated electron, pT > 17 GeV
- loosely isolated muon, pT > 17 GeV, loosely isolated electron, pT > 12 GeV

- double muon, pT > 9 GeV, single electron, pT > 9 GeV
- single muon pT > 8 GeV, double electron, pT > 12 GeV

muon performance at HLT

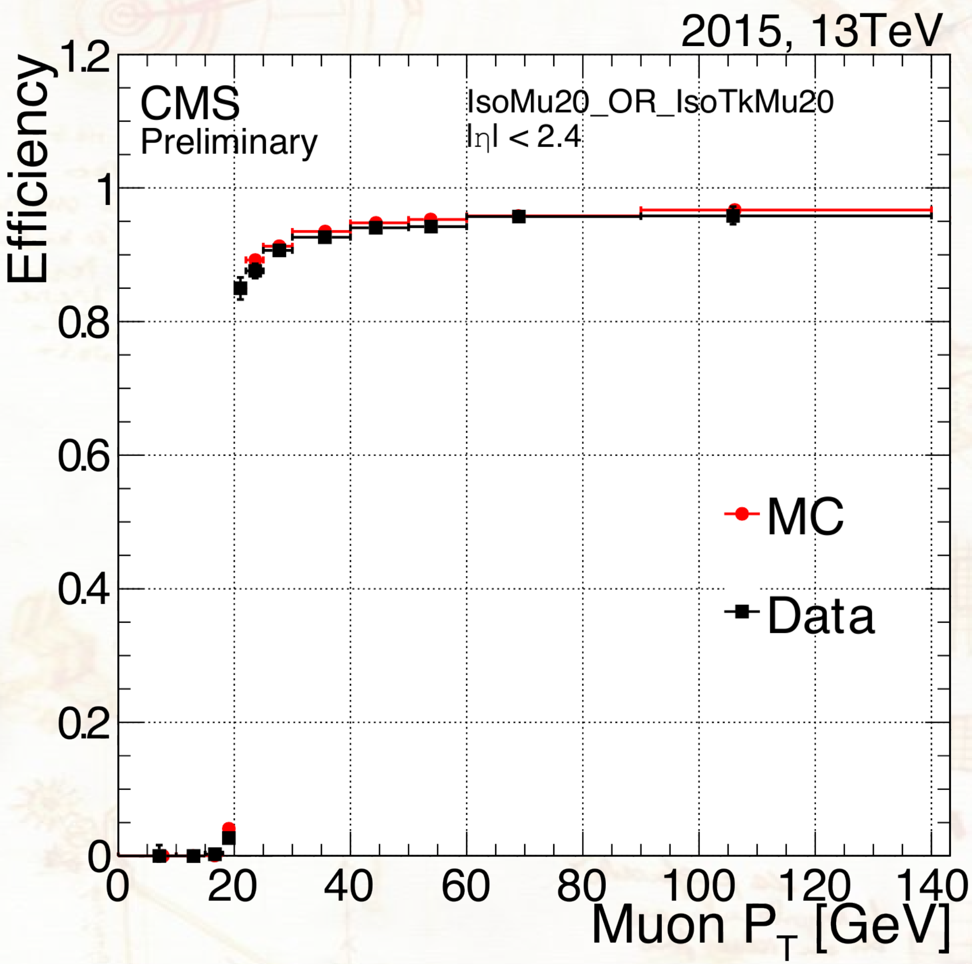

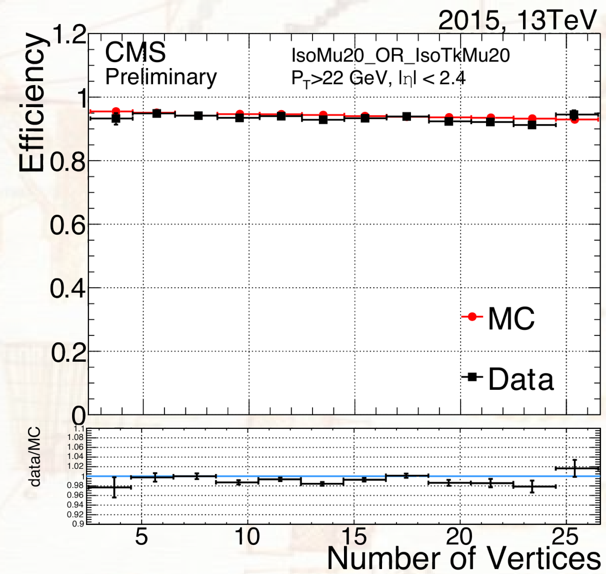

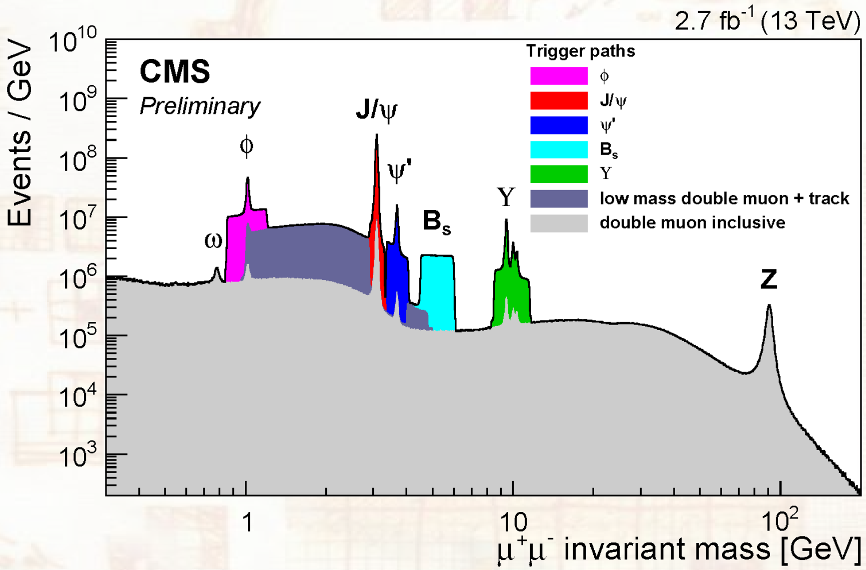

the CMS HLT “menu” in 2015 – e/ɣ

Single electron

- loosely isolated electron, η < 2.1, pT > 22 GeV
- loosely isolated electron, pT > 23 GeV (disabld at 7e33, backed up by cross-triggers)
- loosely isolated electron, pT > 27 GeV
- non-isolated electron, pT > 105 GeV

Double electrons

- isolated double electrons, pT > 17, 12 GeV
- loosely isolated double ele., η < 2.1, pT > 24, 22 GeV
- non-isolated double electrons, pT > 33 GeV

Triple electrons

- triple electron, pT > 16, 12, 8 GeV

Single photon

- isolated photon, pT > 36 GeV, MET or VBF cross triggers
- non-isolated photon, H/E selection, pT > 165 GeV
- non-isolated photon, no H/E selection, pT > 250 GeV

Double photons

- isolated double photons, pT > 30, 18 GeV, electron veto, mass cuts
- isolated double photons, pT > 36, 22 GeV
- non-isolated double photons, pT > 60 GeV

electron performance at HLT

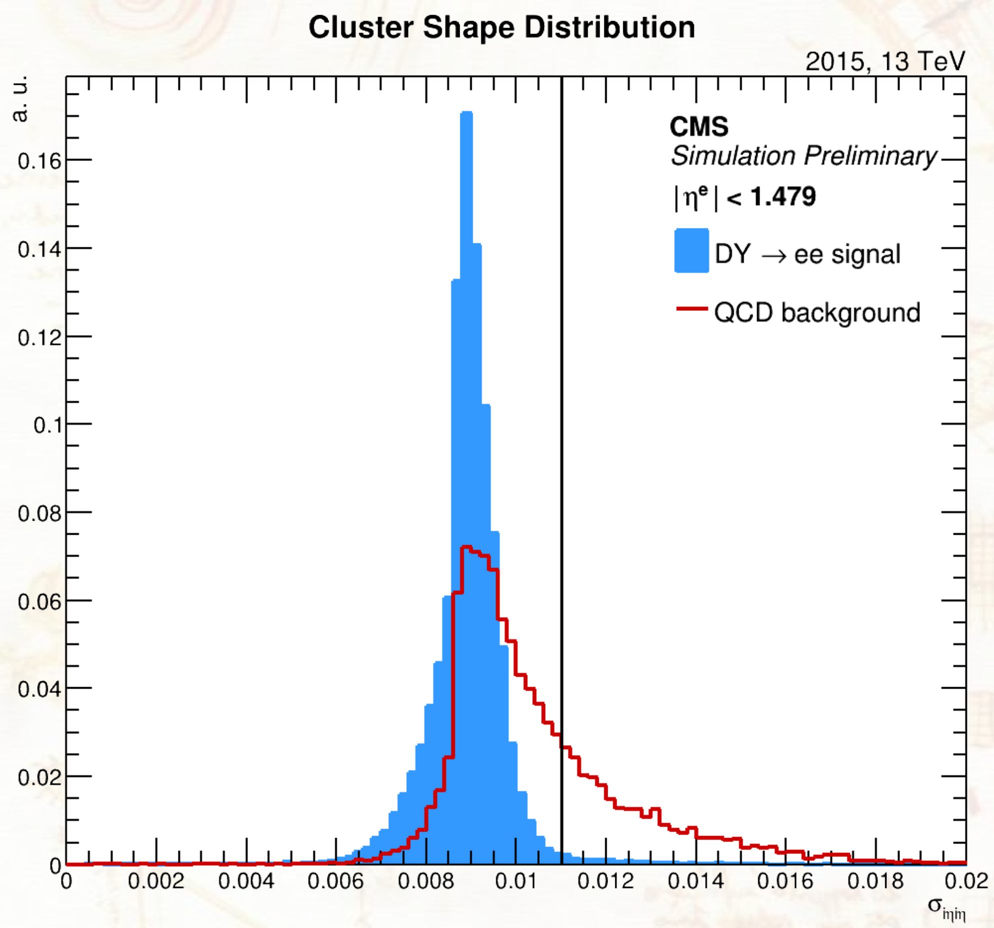

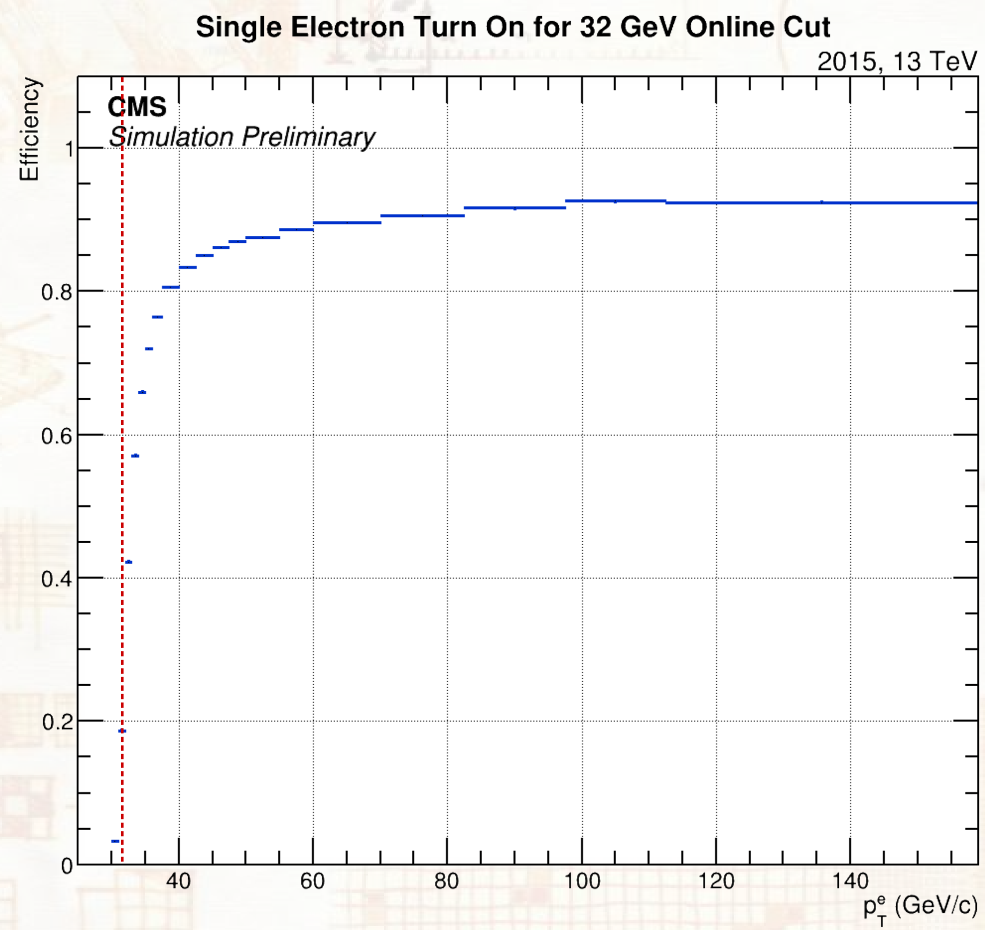

the CMS HLT “menu” in 2015 - calo

Jets

- single particle flow jet, pT > 450 GeV
- single calorimetric jet, pT > 500 GeV
- di-jets, pT >

A familiy of triggers for displaced jet topologies

B-tagged jets

- double pf. jets, pT > 160 GeV (pT > 100 GeV central), double b-tagging
- quadruple pf. or calo. jets, single, double or triple b-tagging

Taus

- double isolated tau, η < 2.1, pT > 35 GeV
- isolated tau, η < 2.1, pT > 50 GeV + MET cross trigger

HT, MHT and MET

- particle flow HT > 800 GeV
- calorimetric HT > 2000 GeV
- calorimetric HT > 450 GeV (for parking)

- particle flow MET > 170 GeV
- calorimetric MET > 200 GeV
- particle flow MET > 90 GeV, MHT > 90 GeV
  - and higher threshold backups, as these triggrs have a very large rate at high pilup

- HT plus multijets
- HT plus high-mass wide jets
- A family of α_T triggers
- A family of Razor triggers
- MET cross-triggers with muons or monojets

HLT output and streams

- the HLT is responsible for splitting the data in different streams
  - different purposes
  - different event content
  - different rates

- physics, calibrations, monitoring, etc.

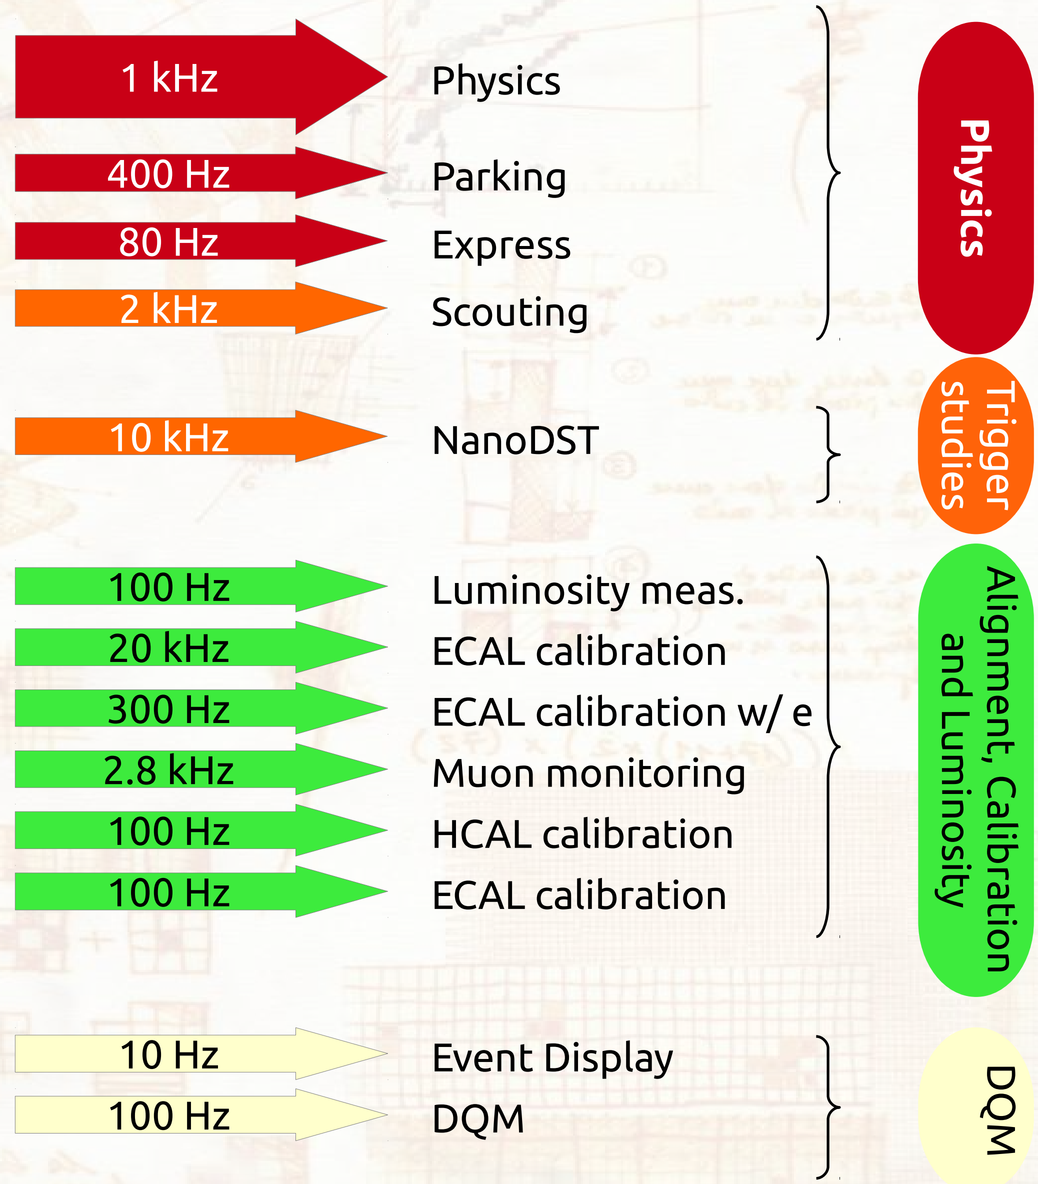

HLT output and streams

- Physics streams
  - full event content
  - prompt reconstruction
  - events for physics analyses

- Express stream
  - subset of the Physics streams
  - reconstructed immediately
  - used for monitoring and calibration of the detectors
- Parking
  - full events content
  - not (prompt) reconstructed
- Scouting
  - reduced event content
  - events for physics analyses

HLT output and streams

- **NanoDST stream**
  - saves trigger information for 10% of all L1-accepted events
  - used for trigger studies
- **AlCa streams** collect events for dedicated calibration workflows
  - only a fraction of the detector is read: small event size, high rate
- **DQM streams**
  - online monitor of the detector conditions, online reconstruction performance, etc.

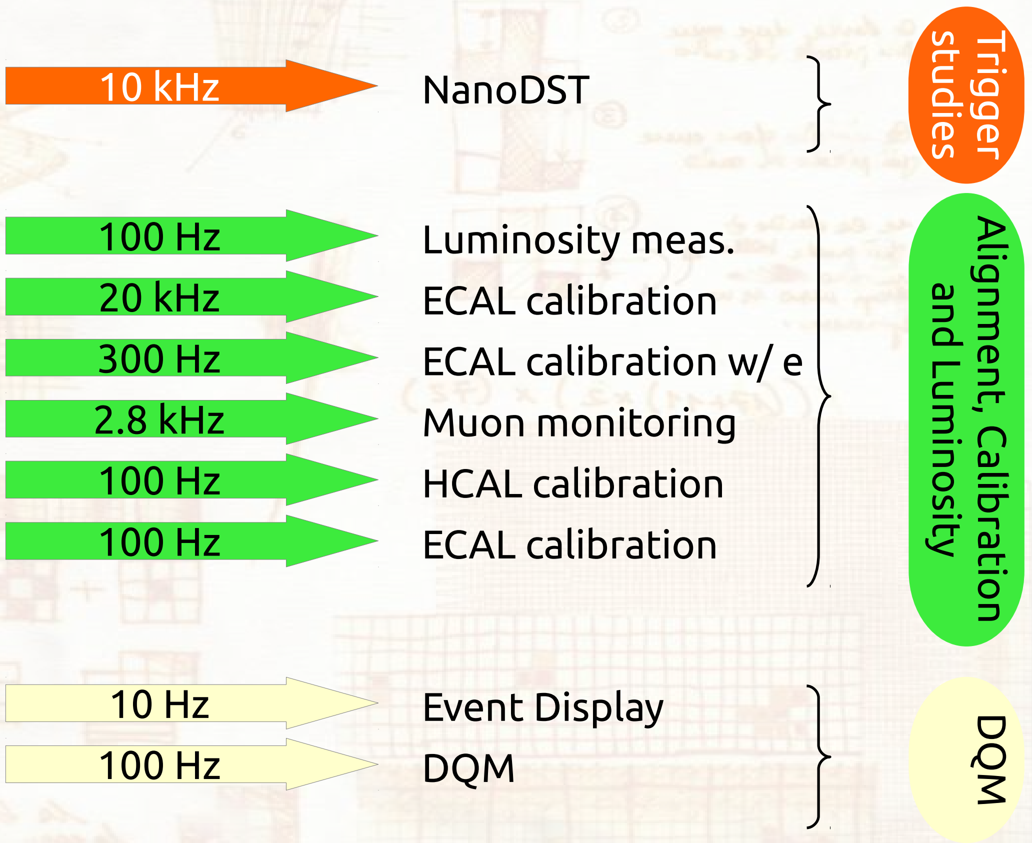

Data Scouting at CMS

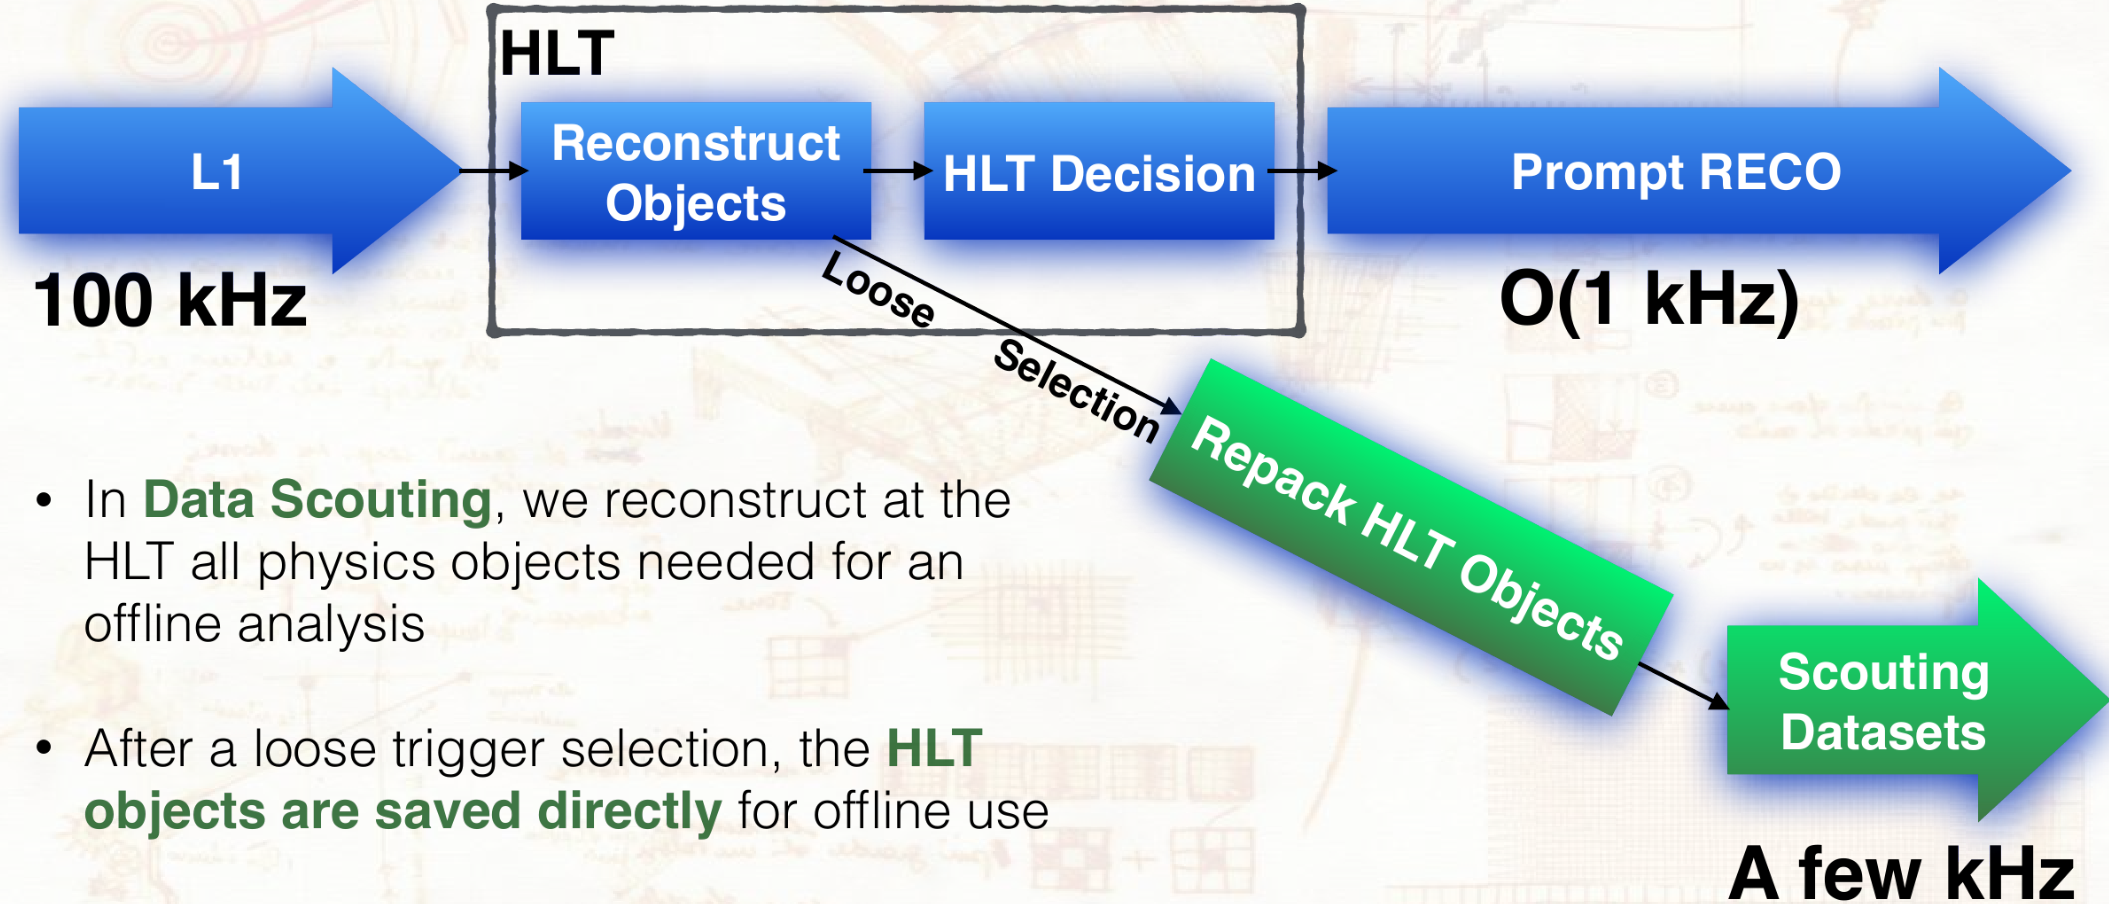

- In **Data Scouting**, we reconstruct at the HLT all physics objects needed for an offline analysis
- After a loose trigger selection, the **HLT objects are saved directly** for offline use
- The event is **not** sent to Prompt RECO, and no RAW data is saved

Data Scouting at CMS

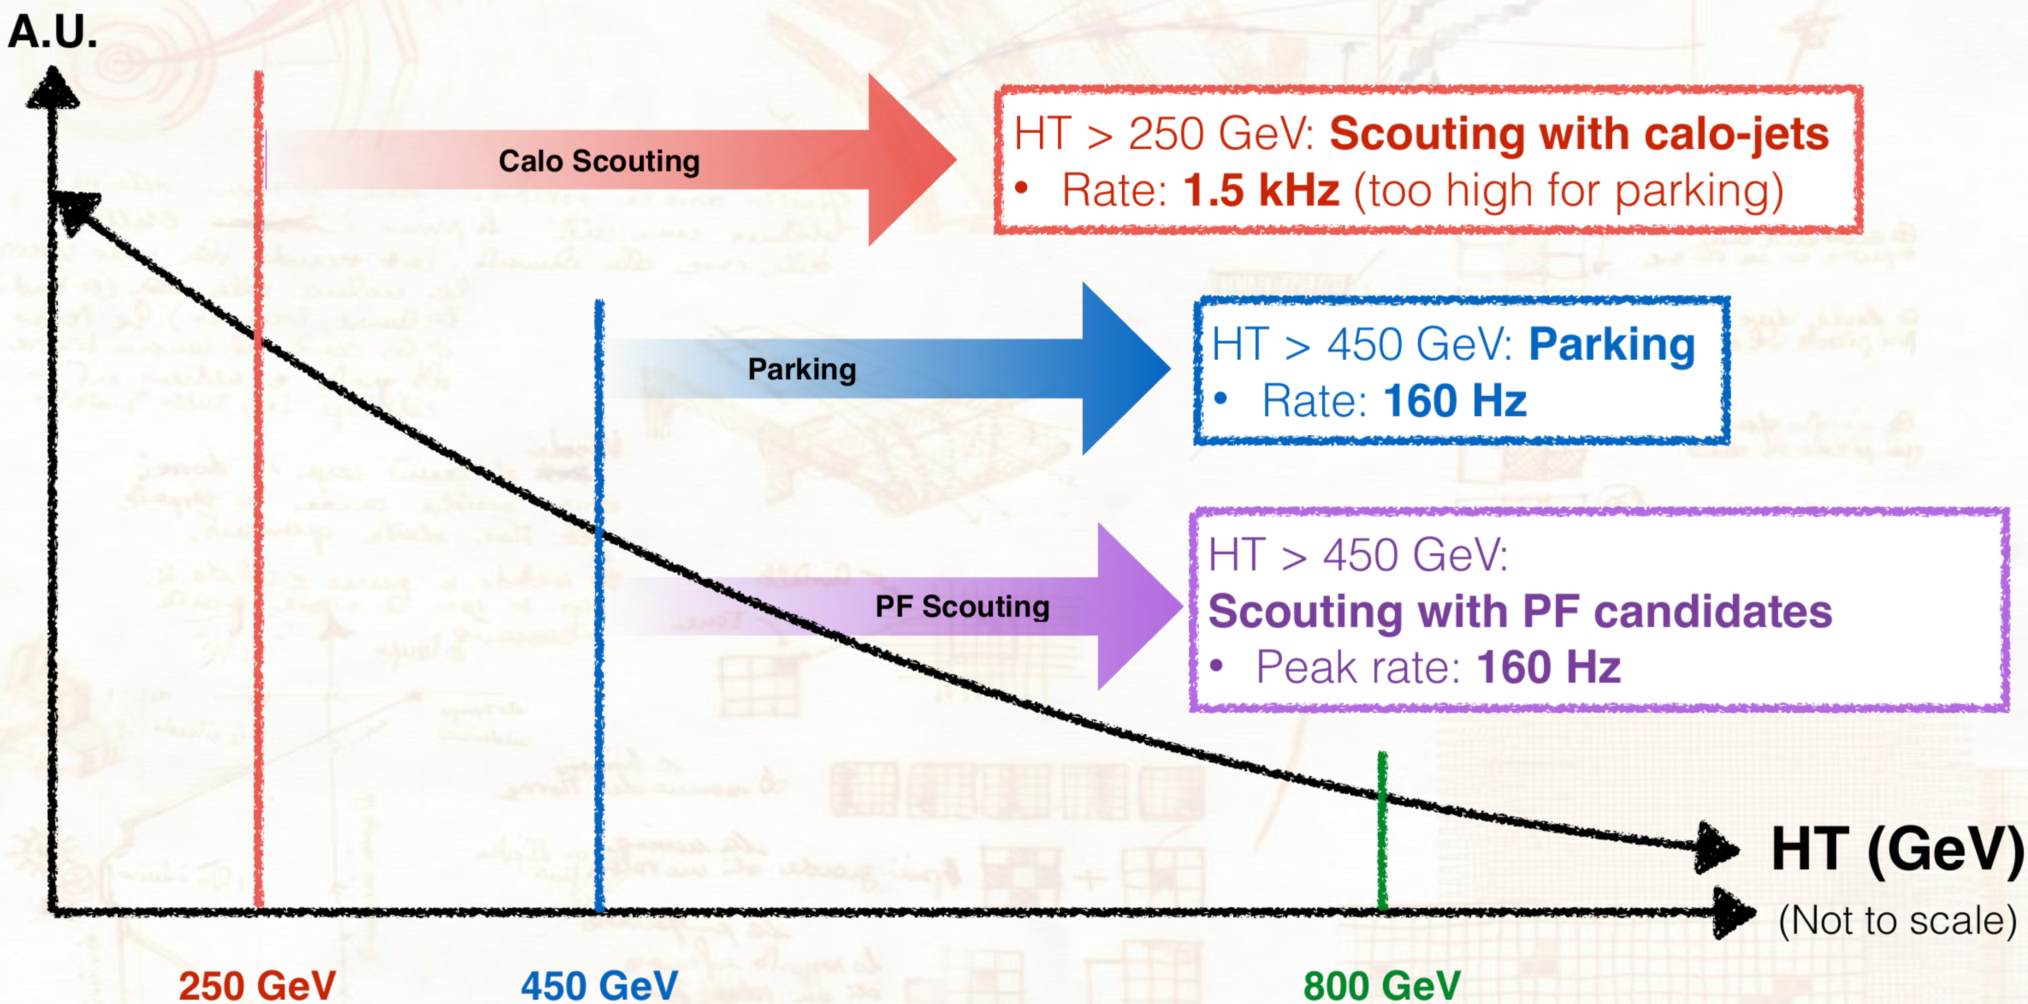

HLT processing time

look at L1 information

fast accept/reject

reconstruction of leptons and calorimetric object

e.g. muons, photons

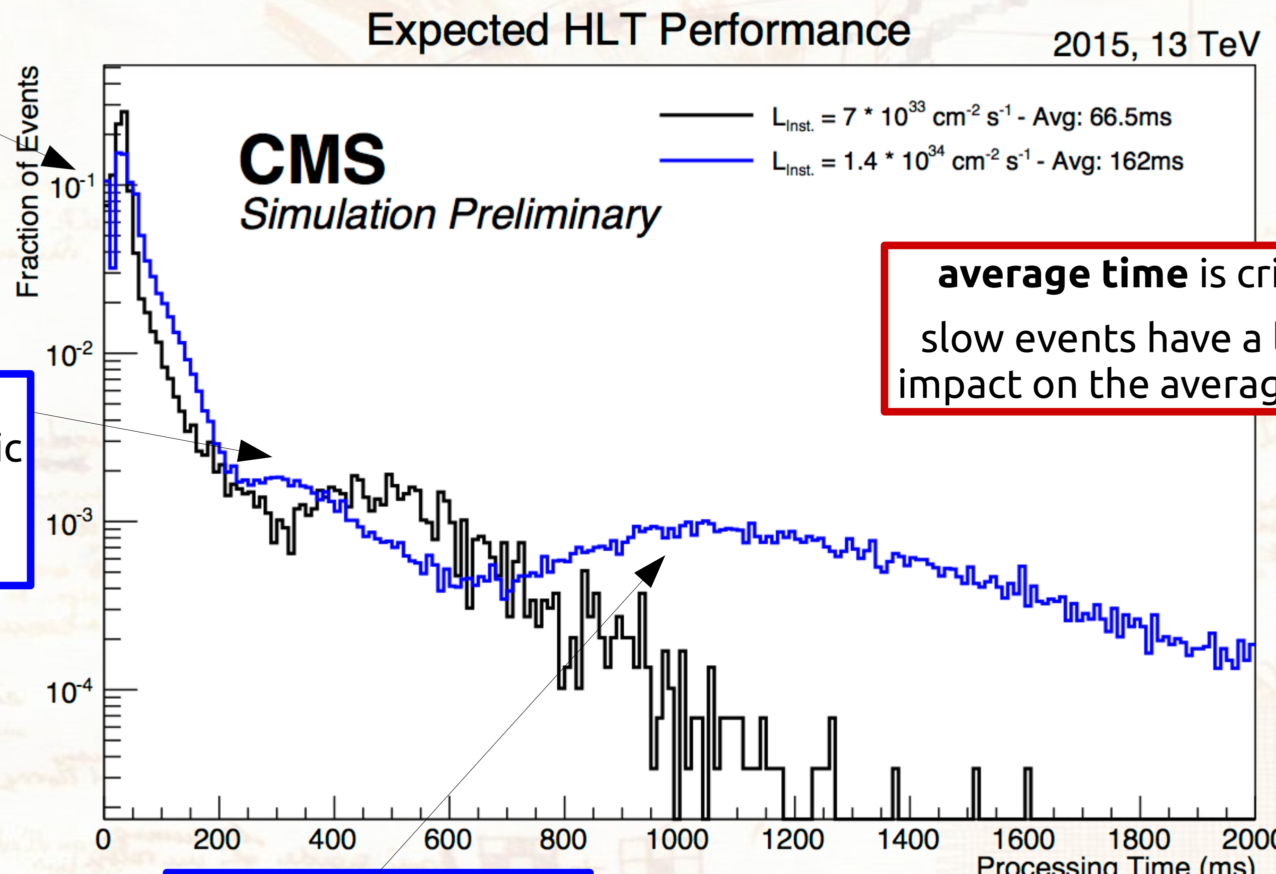

average time is critical

slow events have a large impact on the average time

full track reconstruction and particle flow

e.g. jets, tau

HLT processing time

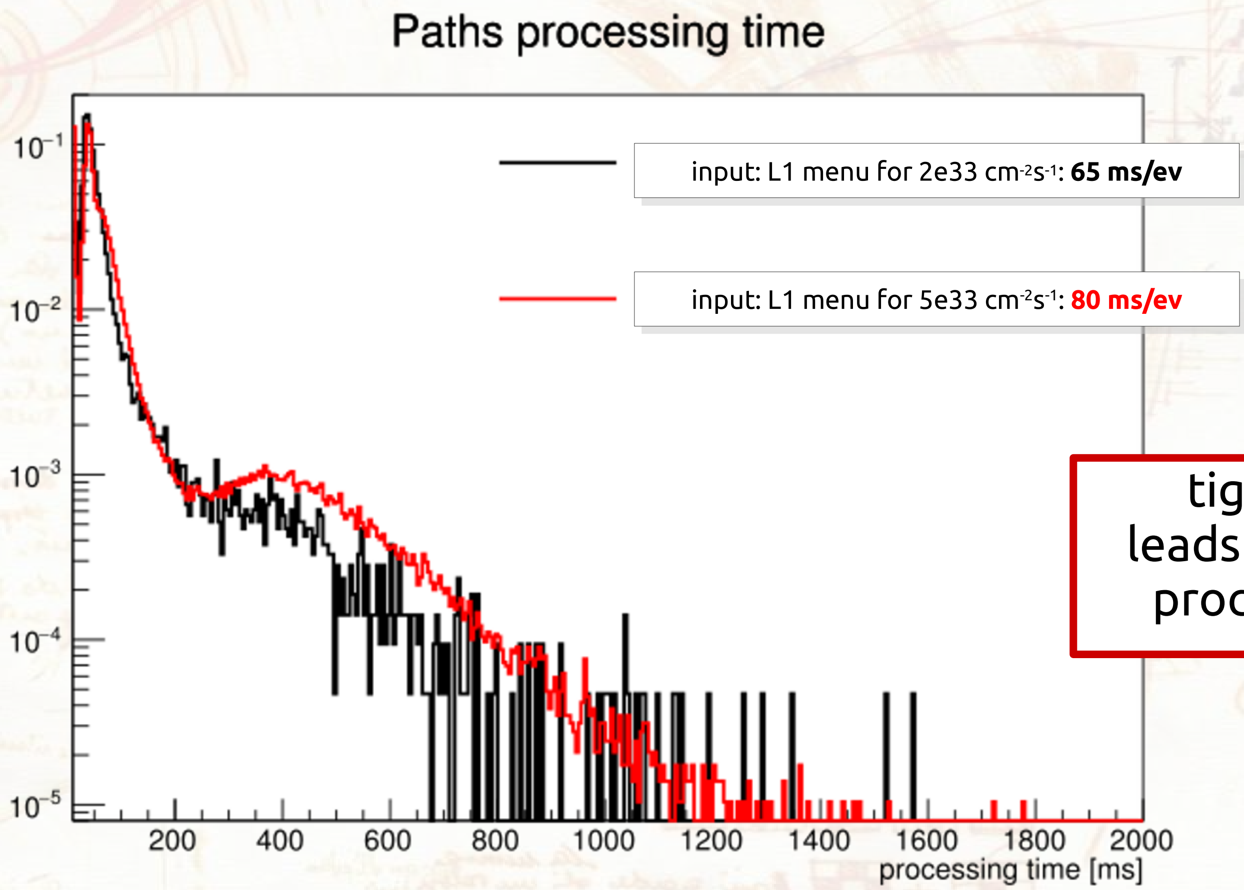

tighter L1 selection leads to a higher average processing time at HLT

HLT and multithreading

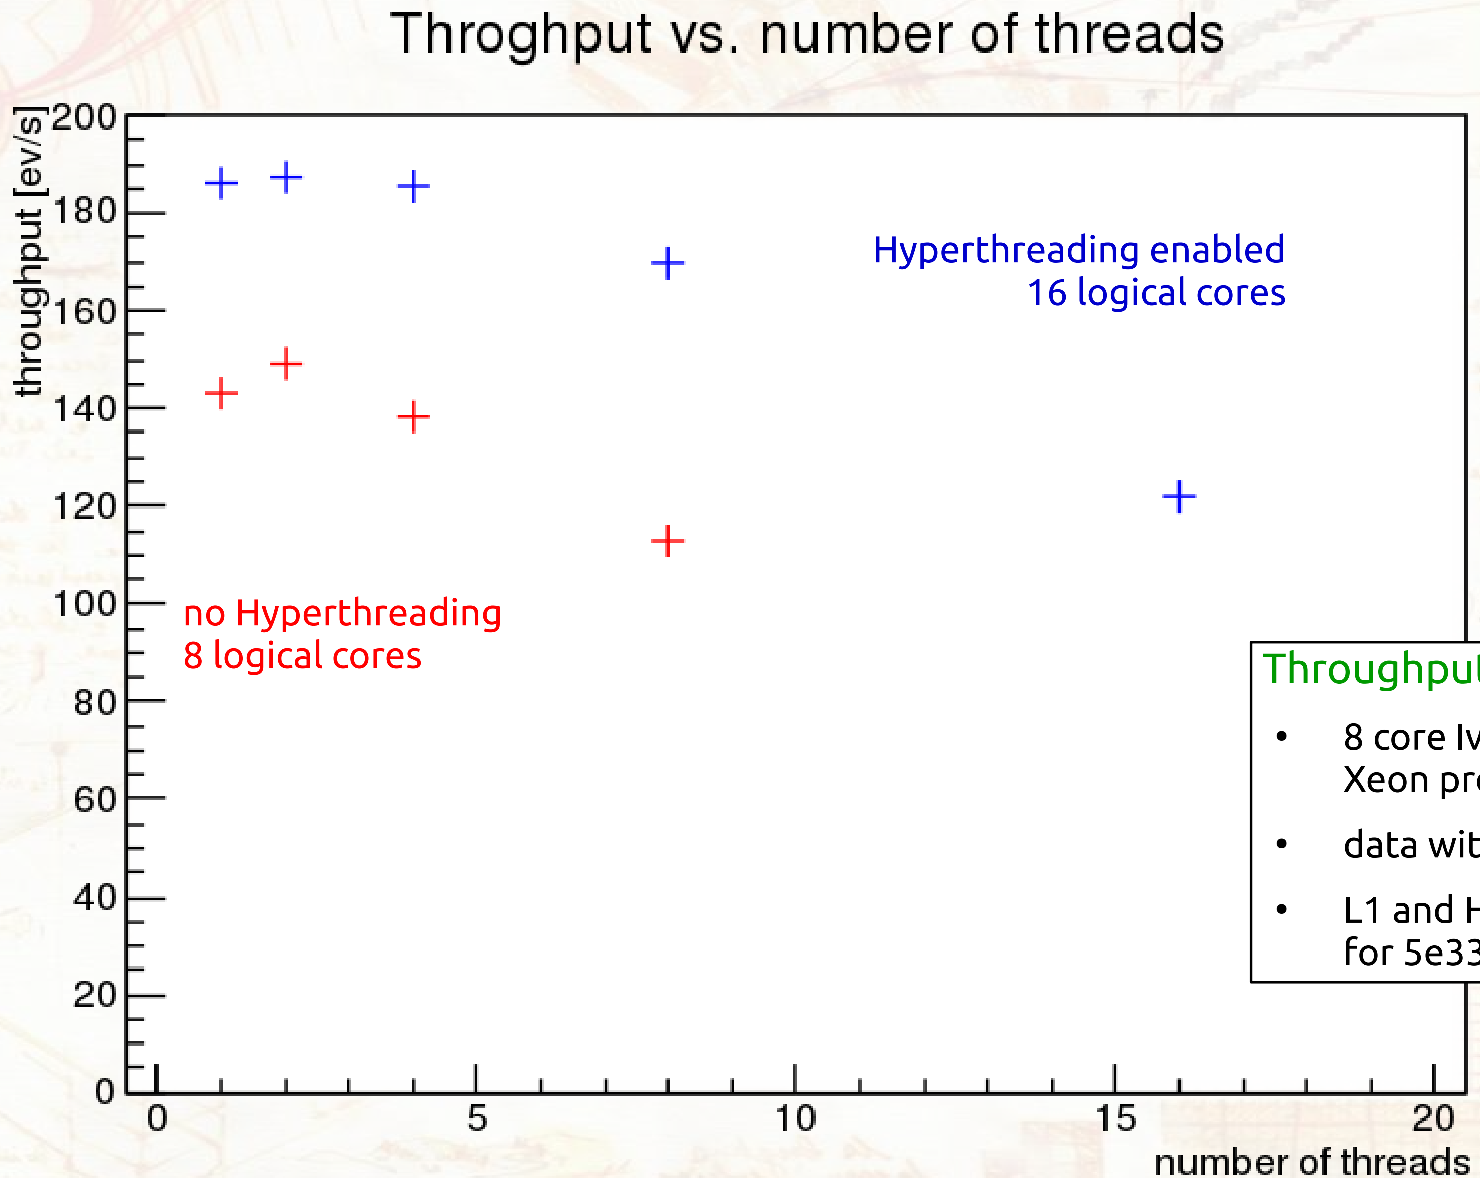

Throughput (ev./s)

- 8 core Ivy Bridge Xeon processor
- data with pileup ~17
- L1 and HLT menus for 5e33 $cm^{-2}s^{-1}$

HLT and multithreading

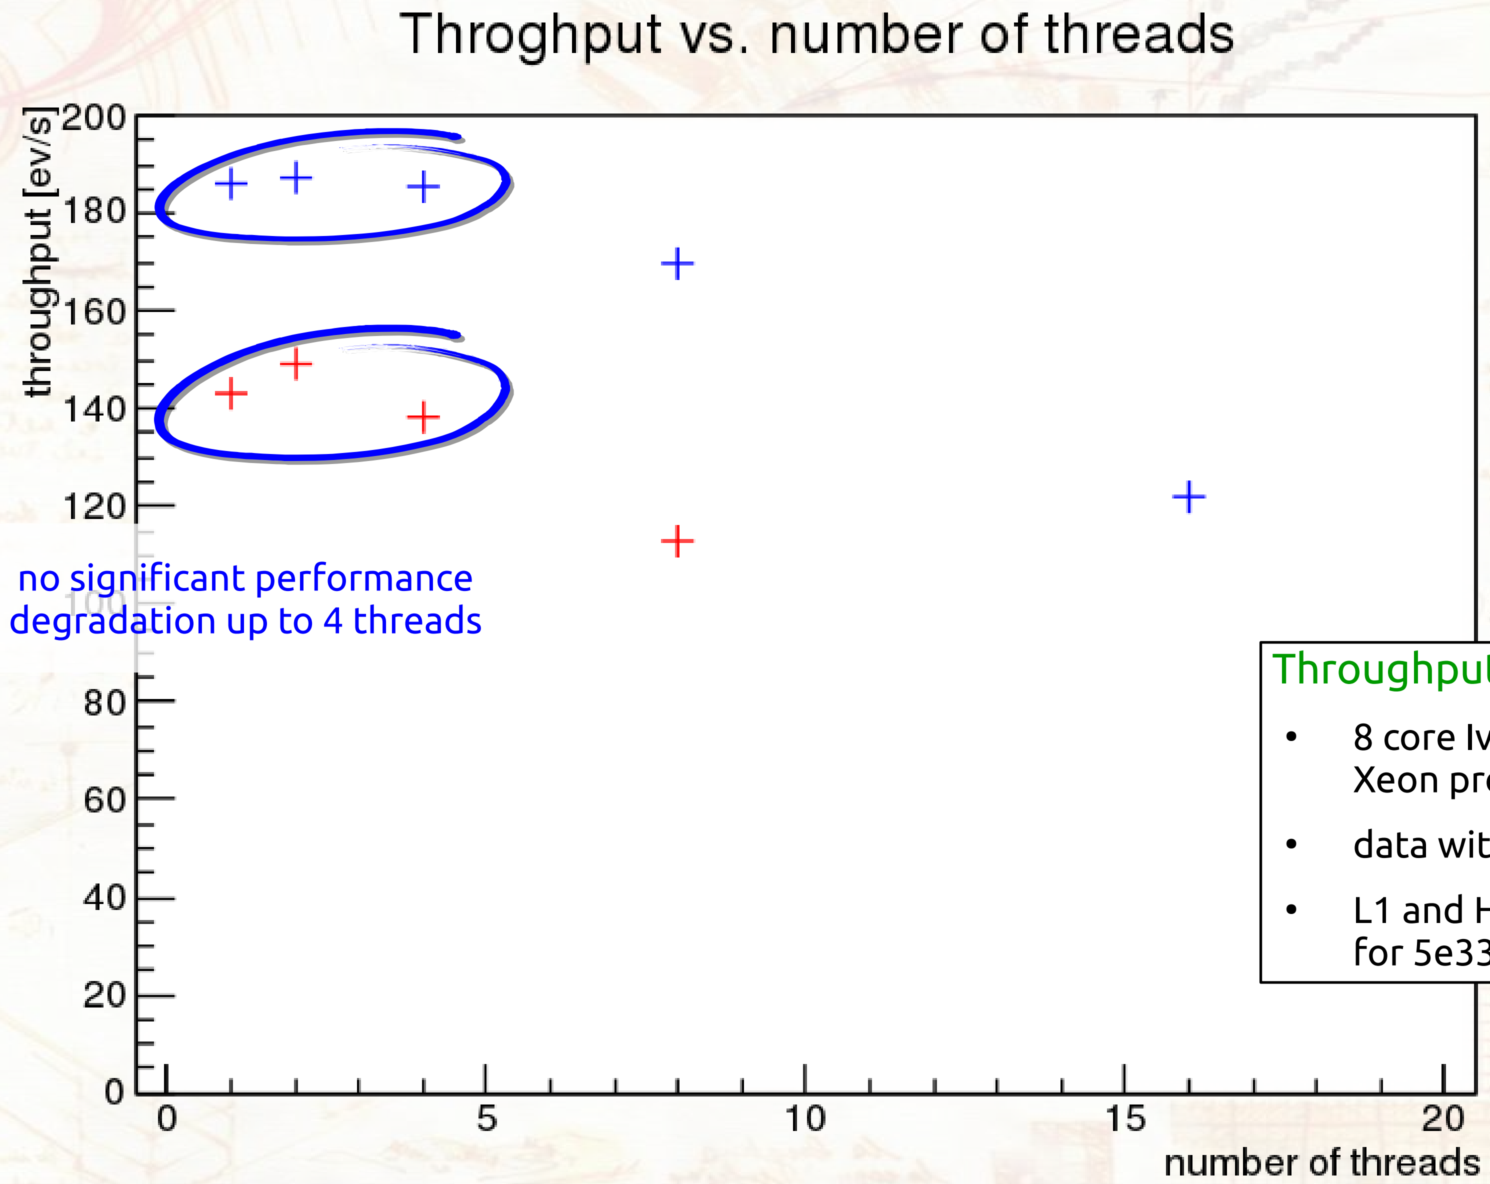

Throughput (ev./s)

- 8 core Ivy Bridge Xeon processor
- data with pileup ~17
- L1 and HLT menus for 5e33 $cm^{-2}s^{-1}$

HLT and multithreading

configurations that use too much memory

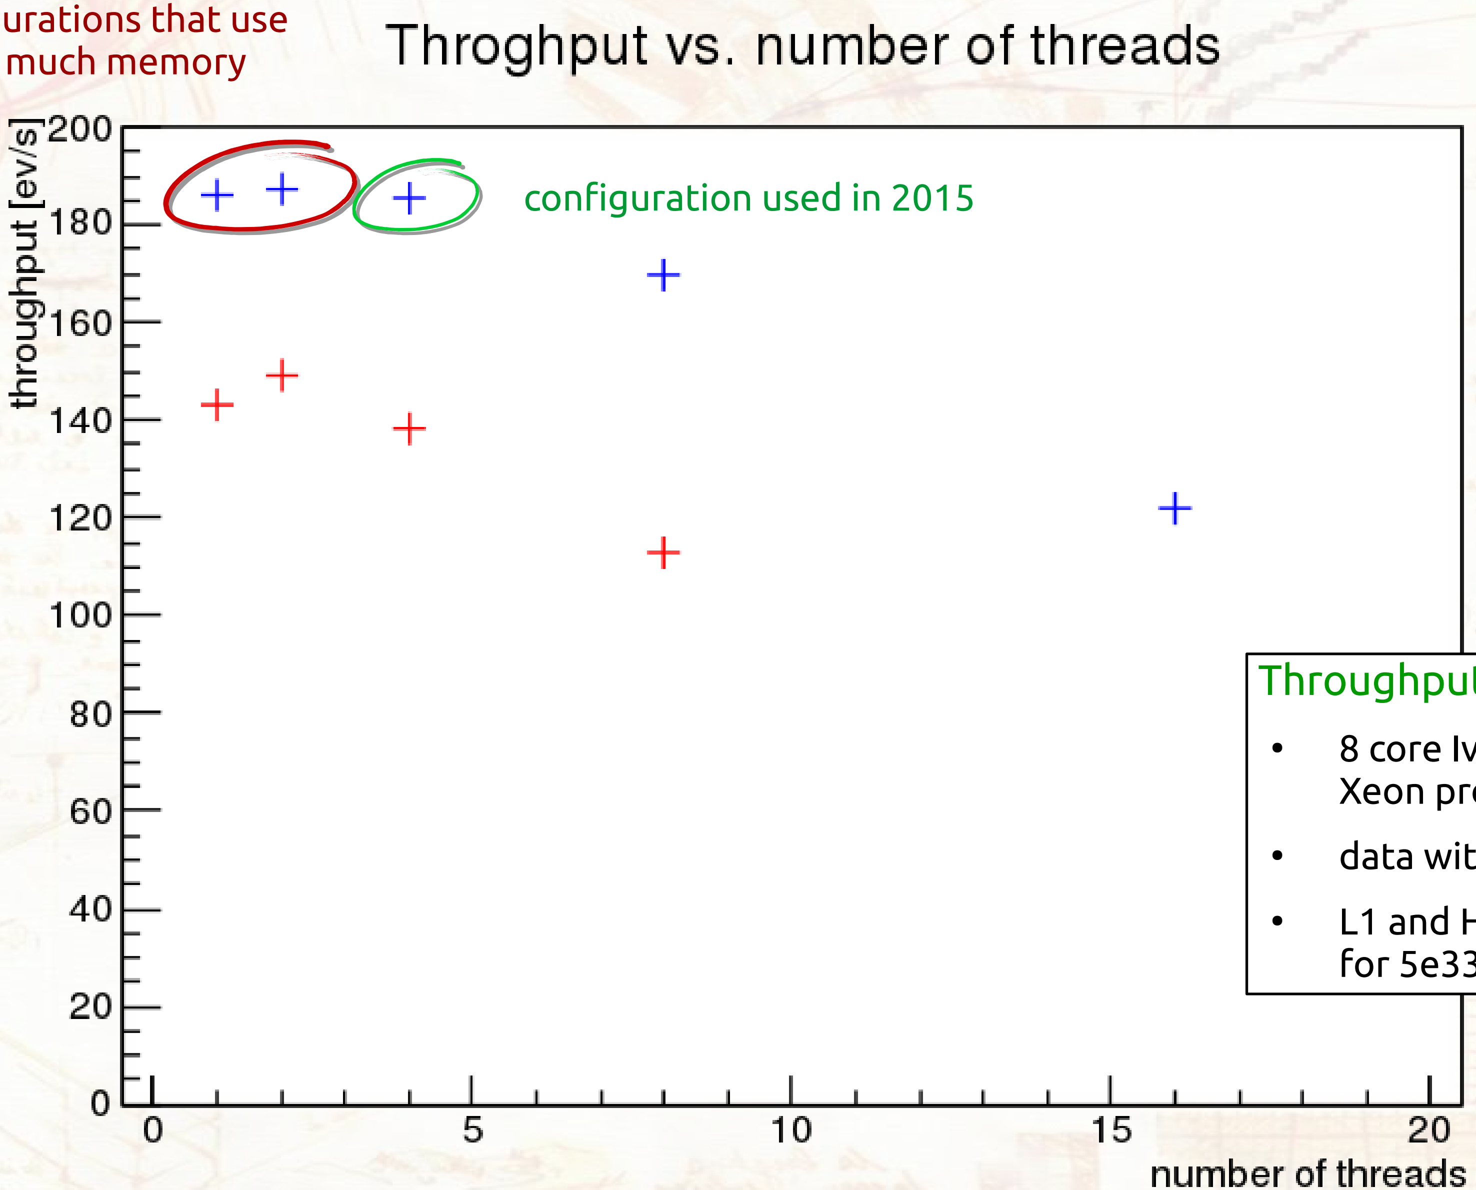

Throughput (ev./s)

- 8 core Ivy Bridge Xeon processor
- data with pileup ~17
- L1 and HLT menus for 5e33 $cm^{-2}s^{-1}$

Filter Farm – extrapolation to 2016

Level 1 Trigger rate – extrapolation to 2016

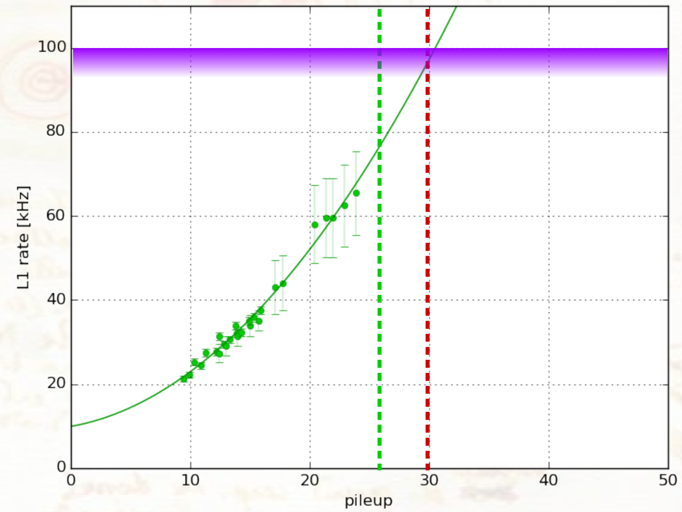

HLT CPU usage – extrapolation to 2016

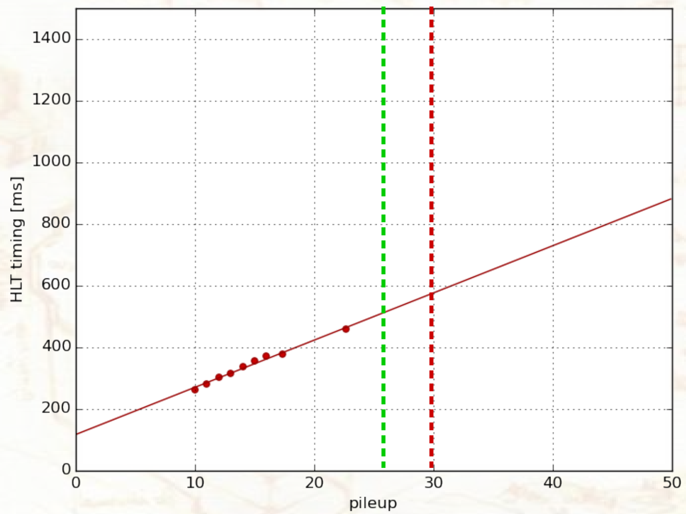

Filter Farm occupancy – extrapolation to 2016

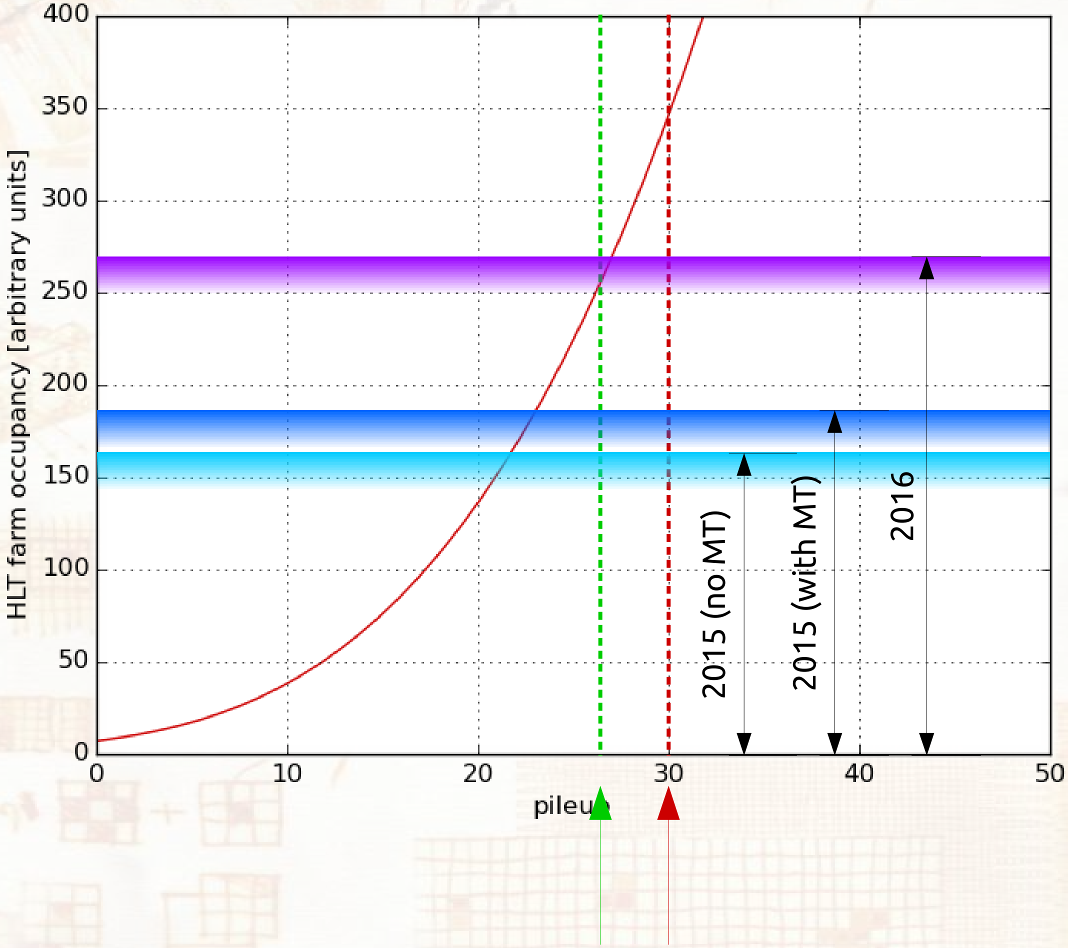

LHC scenarios for 2016

Questions ?