Summer Students Lectures 2015

Triggers for LHC Physics

Andrea Bocci

Introduction

- at a luminosity of **1^{34} cm^{-2}s^{-1}**, or **10 Hz/nb**, the LHC will produce **~0.8 billion** inelastic proton-proton collisions **per second**

- to save all of these collision events, the larger experiments would need to read, process, transfer, and store, tens of TB per second, or **hundreds of PB per hour**

- but do we even need such large amount of data ?

- the more **interesting physics** processes are **much, much rarer** than the inelastic proton-proton scattering !

Introduction

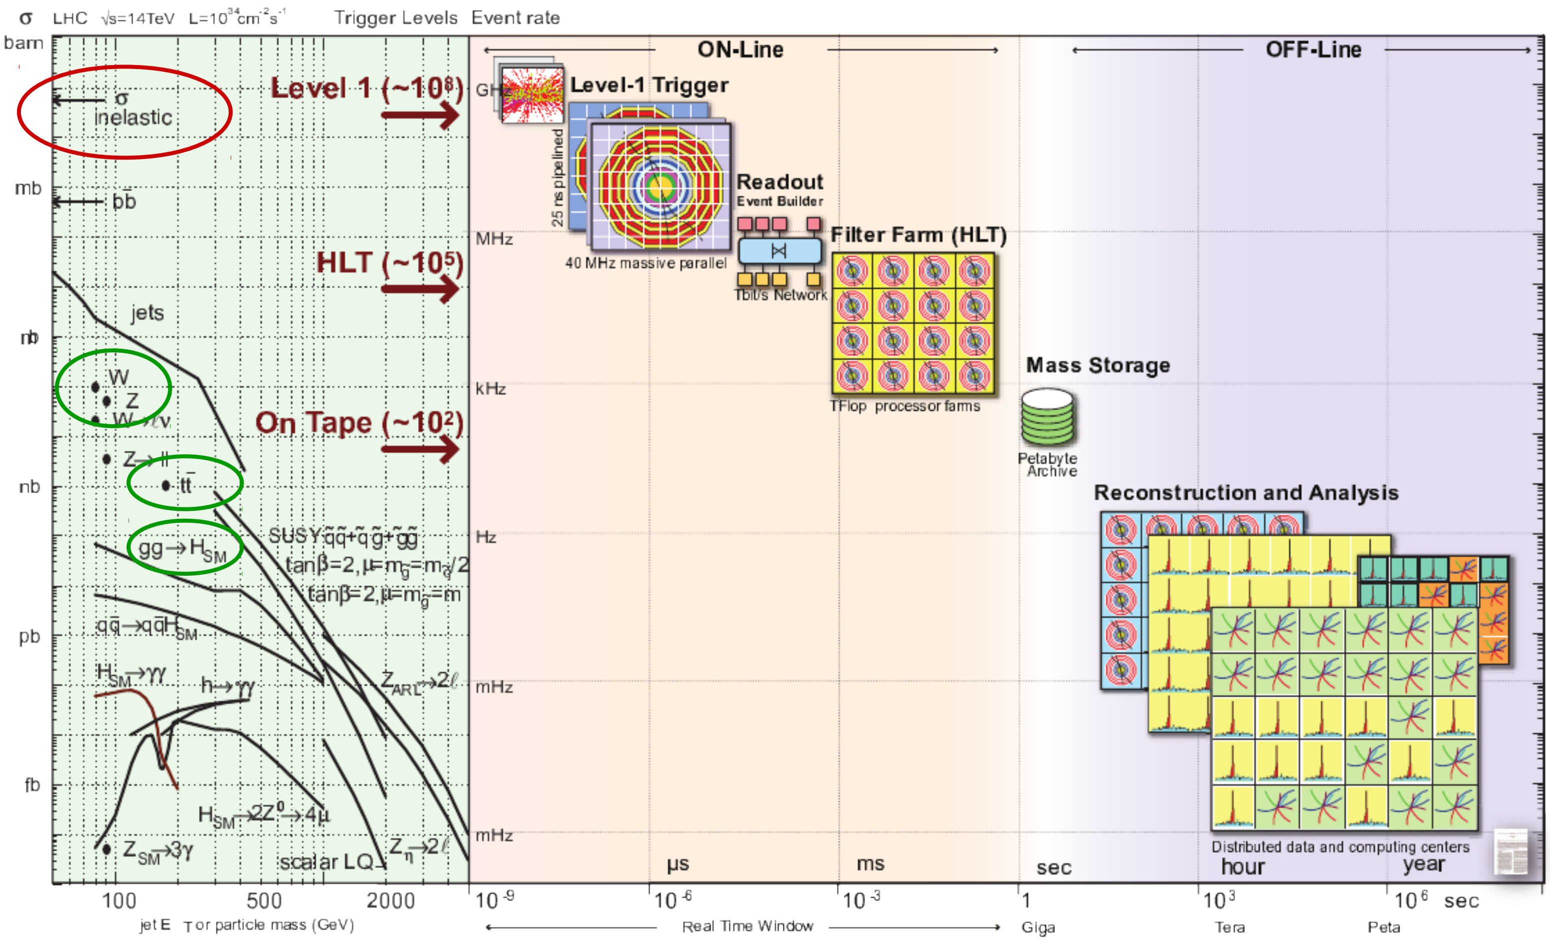

Introduction

- at a luminosity of **1^{34} cm^{-2}s^{-1}**, or **10 Hz/nb**, the LHC will produce **~0.8 billion** inelastic proton-proton collisions **per second**

- at the same luminosity, we expect the production of around **1000 W and Z per second**, **1 $t\bar{t}$ pair per second**, few **Higgs bosons per minute** …

Introduction

- at a luminosity of **1^{34} $cm^{-2}s^{-1}$**, or **10 Hz/nb**, the LHC will produce **~0.8 billion** inelastic proton-proton collisions **per second**
- at the same luminosity, we expect the production of around **1000 W and Z per second**, **1 $t\bar{t}$ pair per second**, few **Higgs bosons per minute** …
- finding one Z boson is like finding a **single person** in **Stockholm** !
- finding a $t\bar{t}$ decay is like finding a **single person** in **all of Europe** !

EVERY SECOND !

Introduction

- at a luminosity of **1^{34} cm^{-2}s^{-1}**, or **10 Hz/nb**, the LHC will produce **~0.8 billion** inelastic proton-proton collisions **per second**

- at the same luminosity, we expect the production of around **1000 W and Z per second**, **1 $t\bar{t}$ pair per second**, few **Higgs bosons per minute** …

- finding one Z boson is like finding a **single person** in **Stockholm** !

- finding a $t\bar{t}$ decay is like finding a **single person** in **all of Europe** !

- finding a Higgs boson is like finding a **single person** on the whole **Earth** !

Can you find me ?

Can you find me ?

What does it mean …

- … to “trigger” ?

- Analyse **as quickly as possible** all the collision events

  → **low latency and processing time**

- **Discard** events as soon as possible, as quickly as possible, the events deemed **not interesting**

  → **high purity**

- While keeping as many as possible of the **interesting physics events**

  → **high efficiency**

Trigger Signatures

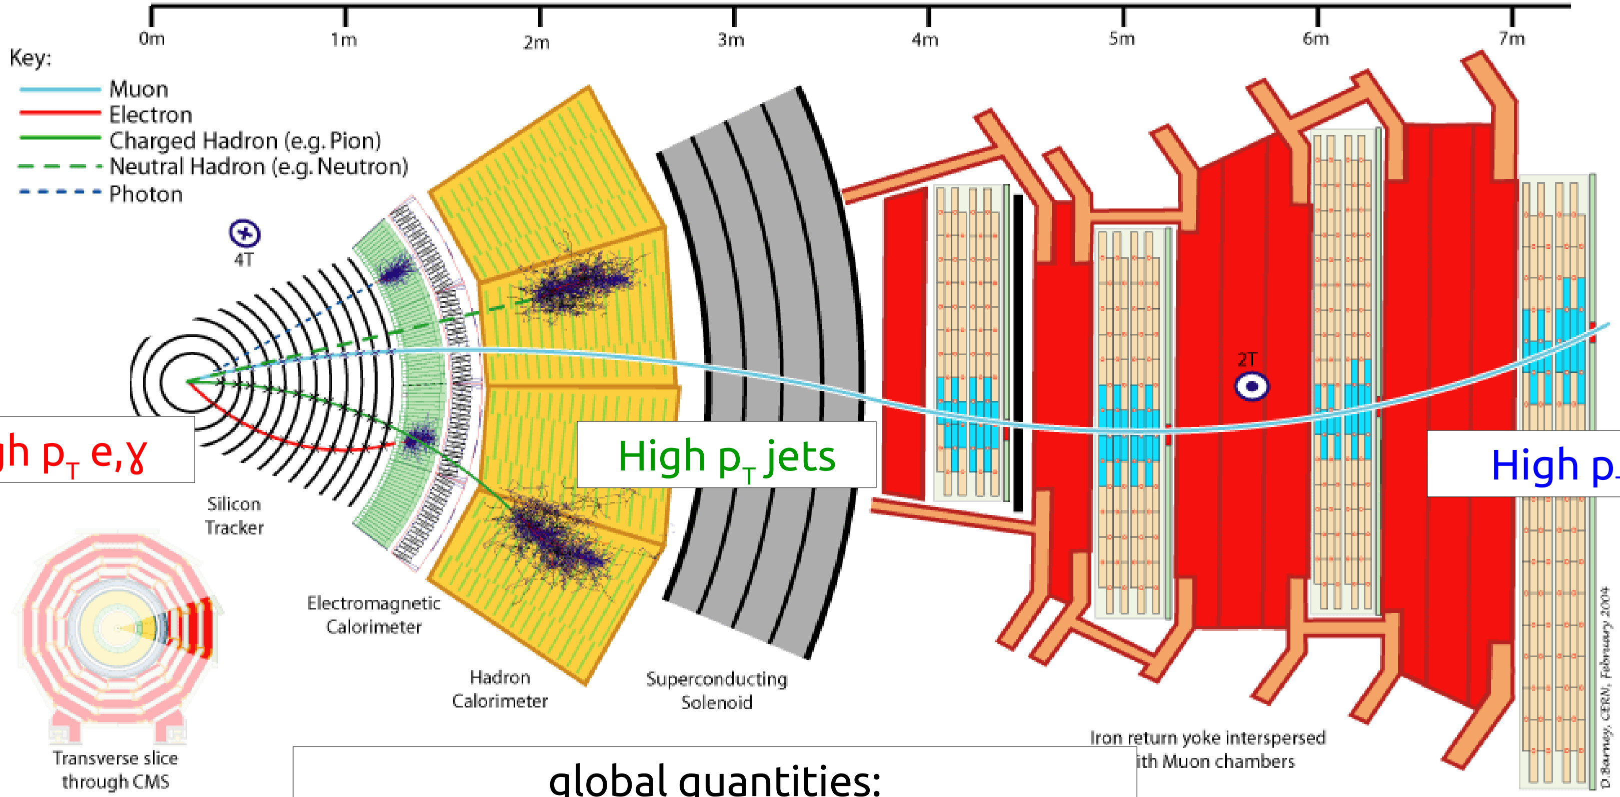

High p_T e,γ

High p_T jets

High p_T μ

global quantities:
total and missing transverse energy

A layered approach

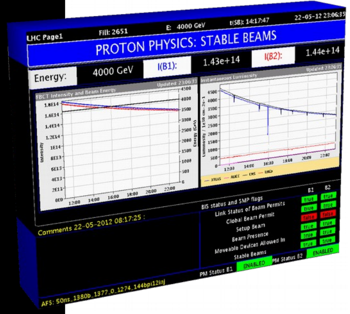

LHC

~ 40 MHz

Level 1 Trigger

100 kHz

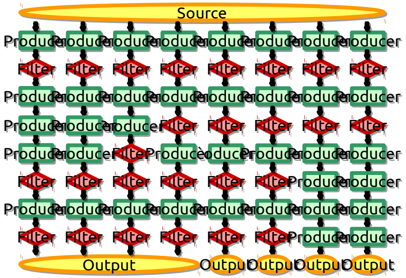

High Level Trigger

~1 kHz

Offline

- no “one size fits all” solution
- take advantage of the information available at each step

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

Level 1 Trigger

Level 1 Trigger: CMS

- fast readout of the detector, with a coarse granularity

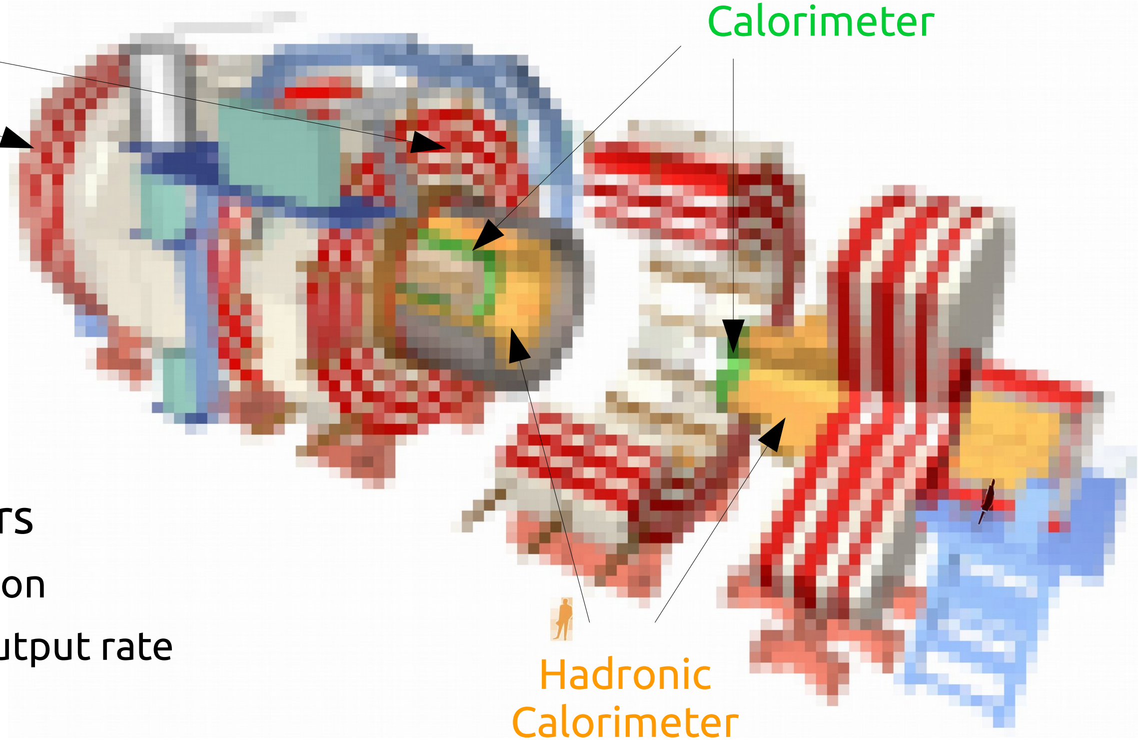

- implementation
  - hardware: ASICs and FPGAs
  - synchronous operation
  - 40 MHz LHC clock

- constraints from the detectors
  - pipeline: ~4 μs to take a decision
  - readout: 100 kHz maximum output rate

L1 Calorimeter Trigger

- Signatures for several physics objects
  - electrons, photons (EM only)
  - Jets, τ leptons (EM + hadronic)

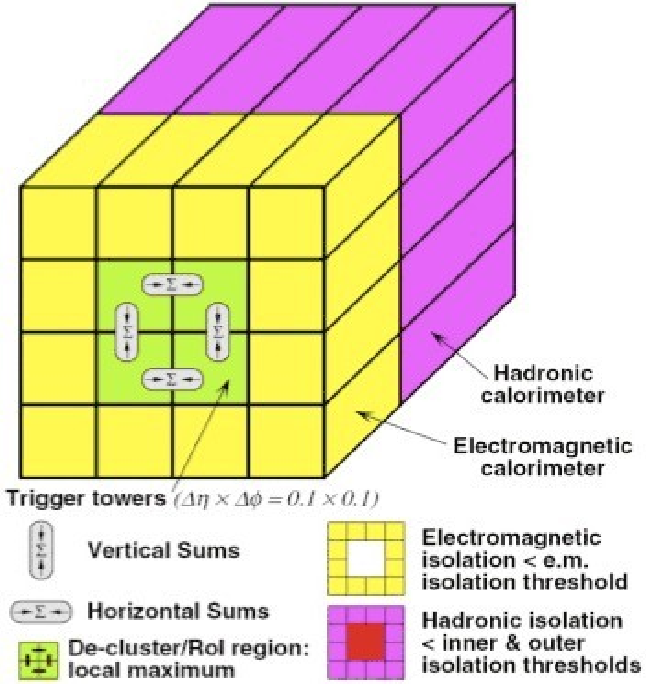

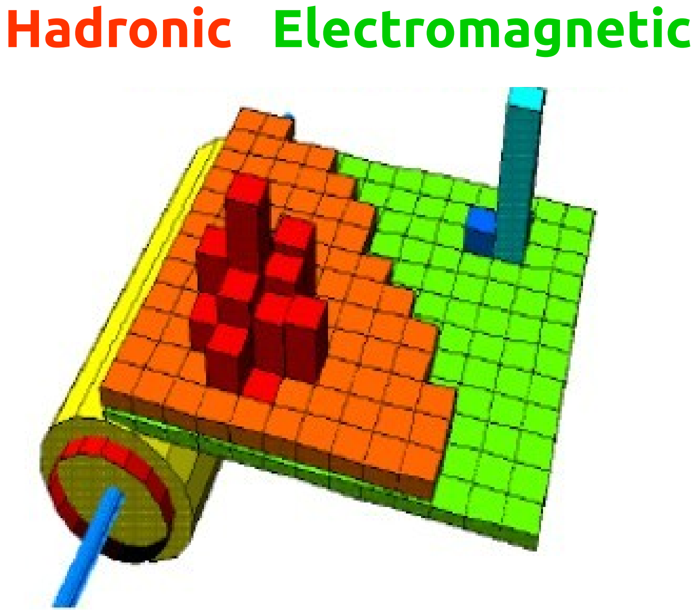

ATLAS e/ɣ trigger

- sum energy from calorimeter cells into towers
- search in 4x4 tower overlapping, sliding window
- clusters from the local maximum within a window

L1 Muon Trigger

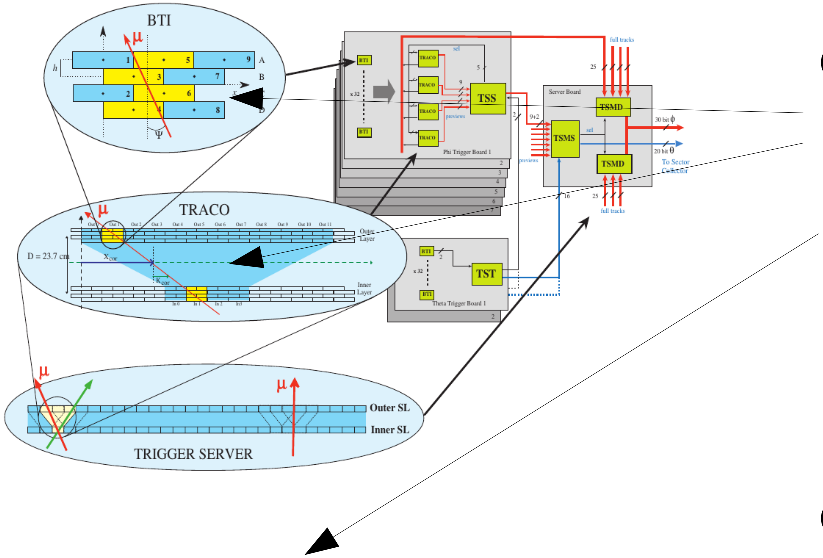

CMS L1 Muon Trigger

- inside each detector, hits and track segments are built into muon candidates
- measure p_T from muon track bending in the magnetic field
- for each candidate, assigne η, φ, p_T and quality
- for each detector, select 4 best candidates

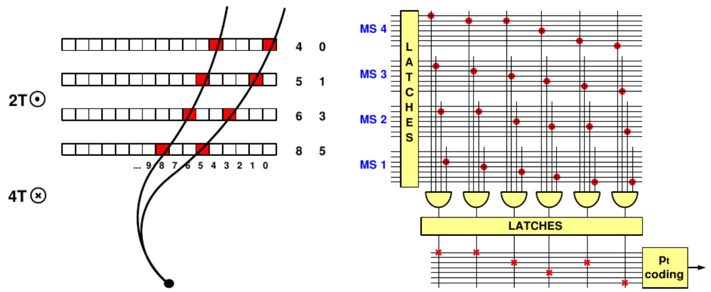

CMS Global Muon Trigger

- combine the candidates from the DT, CSC and RPC detectors
- merge and removes duplicates
- select 4 leading muon candidates

Bringing it all together

- **combine** the information from the L1 trigger candidates
  - ATLAS Topological and Central Trigger Processor
  - CMS L1 Global Trigger

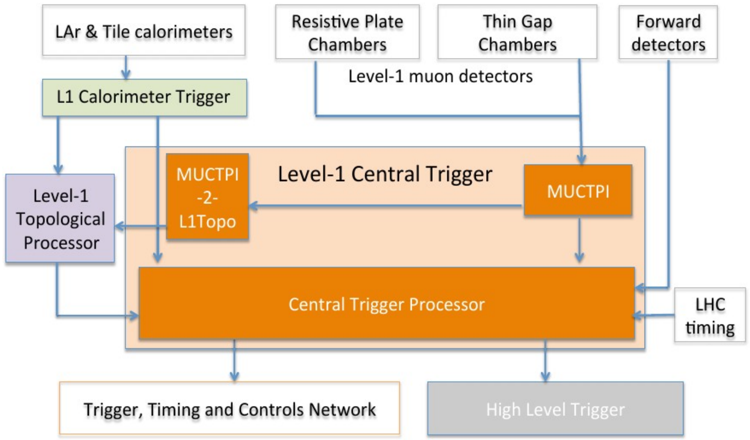

- if the event passes any of the L1 trigger algorithms …
  - one or more candidates above thresholds
  - correlation between candidates (Δη, Δφ, …)
  - etc.

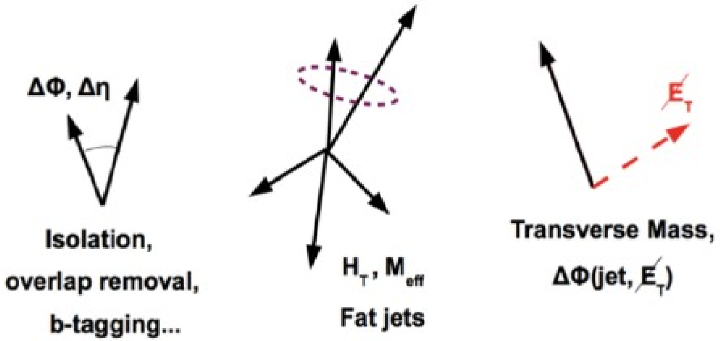

- … **accept the event**
  - perform the full readout of the detector
  - send the event to the High Level Trigger

But, wait …

- … **where does the limitation to the Level 1 Trigger rate come from ?**

- the information in input to the L1 trigger is available for every collision event
  - by itself, the L1 trigger could even accept every event

- accepting an event means triggering the full readout of the detector
  - this is limited by the **detector electronics**, front-end boards, and DAQ system
  - for CMS and ATLAS, the limit is **100 kHz**
  - for LHCb, the limit is **~ 1 MHz**

- trying to read more events causes *back pressure* and *dead time*

Dead time

- *dead time* is the fraction of time that a detector is unable to process its input
  - because it is already processing the previous events
  - because of problems in the readout, etc.

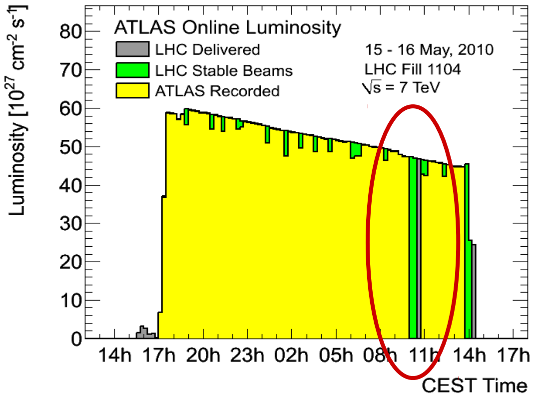

Example of *dead time* due to **detector readout** issues

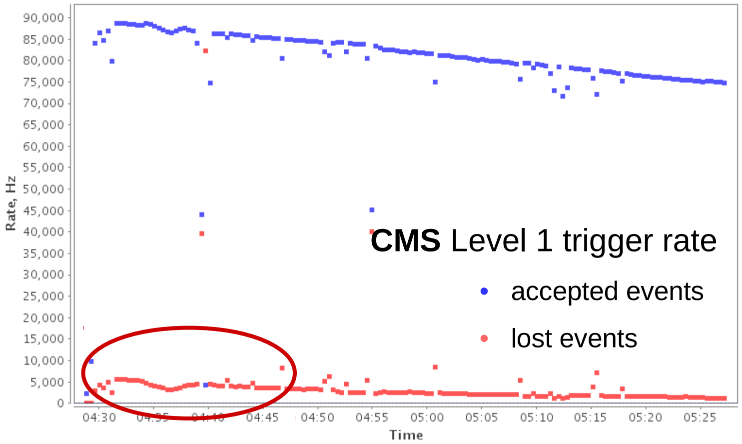

Example of *dead time* due to high trigger rate

- dead time is *unbiased*, that is, it does not depend on the events being lost
  - must be taken into account when considering the integrated luminosity !

High Level Trigger

High Level Trigger

- the events accepted by the Level 1 Trigger are
  - read out from the front-end electronics
  - assembled in by the DAQ
  - reconstructed, analysed and filtered by the High Level Trigger

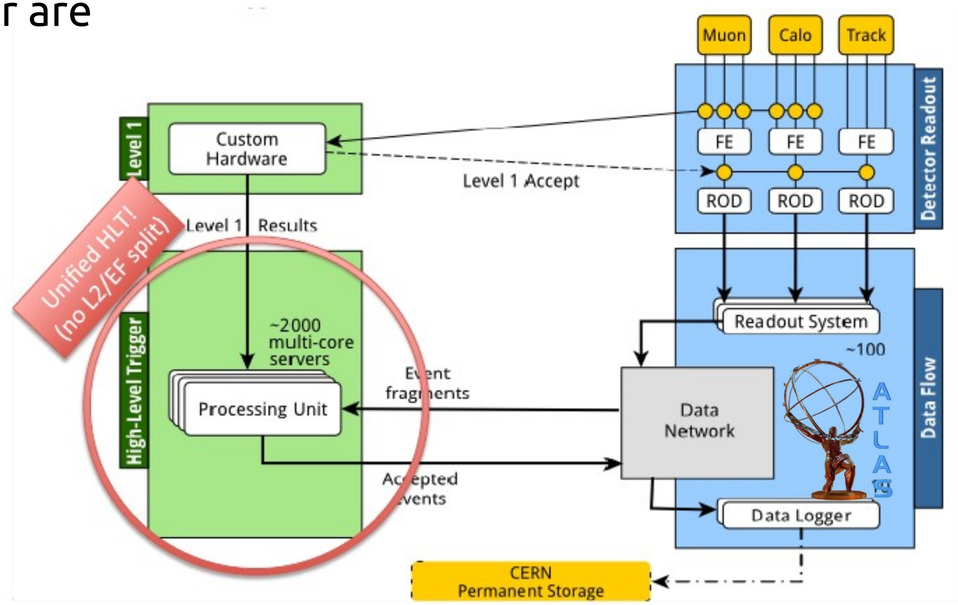

- the High Level Trigger is software-based
  - C++, ROOT, experiment's framework, ...
  - runs on a cluster of commercial computers
  - process events in parallel
  - runs at the L1-accept rate (100 kHz ~ 1 MHz)
  - reduces the event rate to 1 kHz ~ 10 kHz

- the performance of the HLT is as close as possible to the offline reconstruction
  - similar algorithms and calibrations, optimised for speed
  - selection criteria looser than the final analyses

High Level Trigger: CMS

- full readout of the detector at 100 kHz, with full granularity

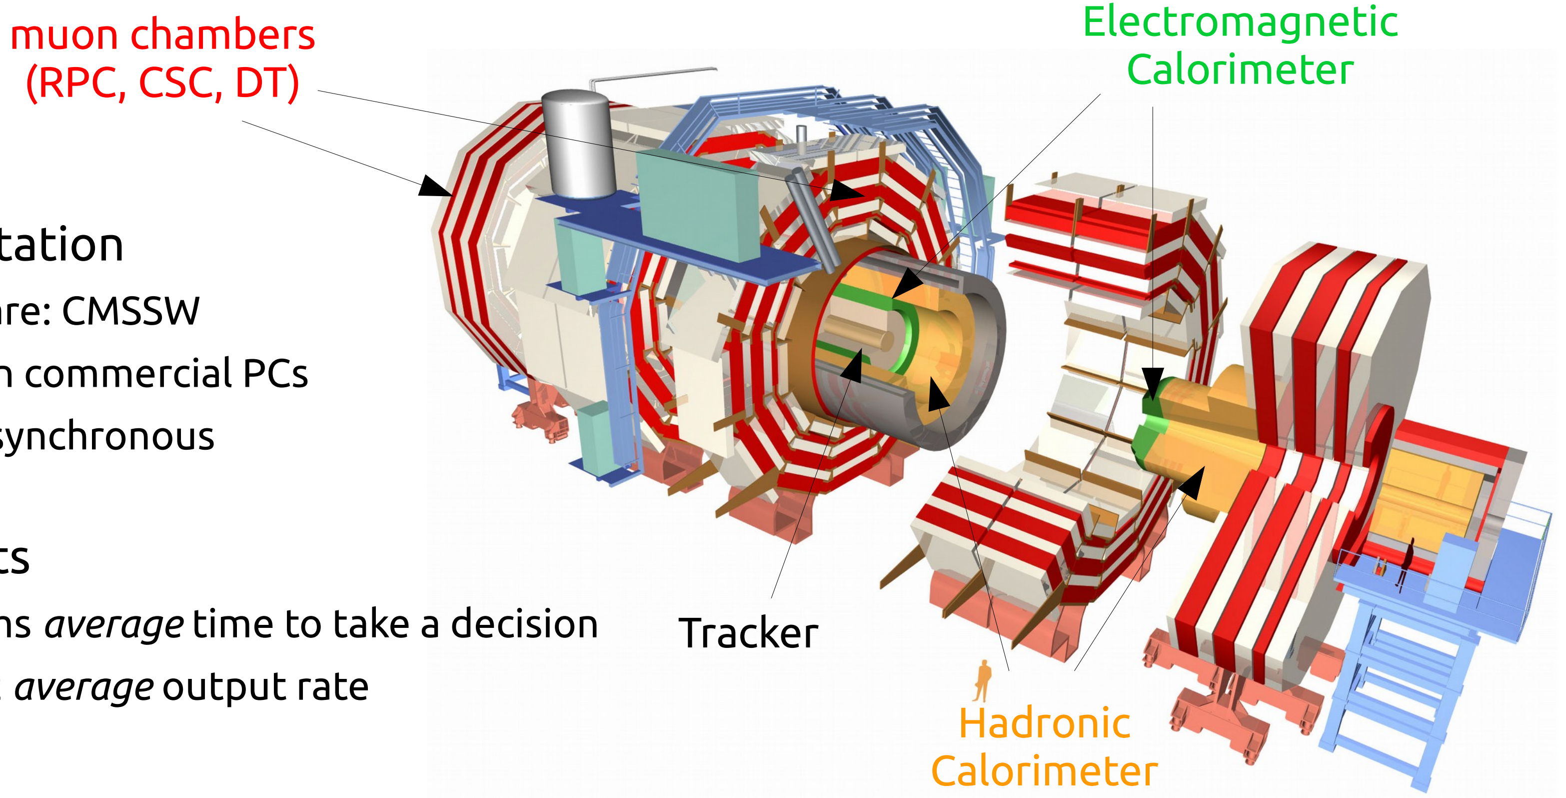

- mplementation
  - software: CMSSW
  - runs on commercial PCs
  - quasi-synchronous

- constraints
  - ~160 ms *average* time to take a decision
  - ~1 kHz *average* output rate

High Level Trigger: CMS

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 and rejection
- b-tagging
  - secondary vertex reconstruction
  - soft-lepton based b-tagging
- but also
  - razor, α_T, dE/dx, …
  - jet substructure, …

HLT constraints – processing power

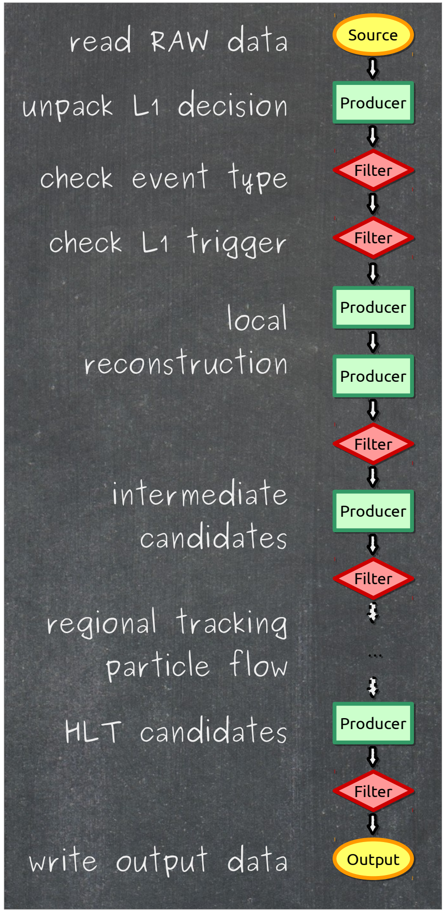

trigger “path”

- the amount of time that the HLT can use to take a decision is limited by the available processing power, for example
  - 20'000 CPU cores
  - 100 kHz input rate

  gives the HLT an *average* of 200 ms per event

- 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 path structure

improve resolution
increase reconstruction time

the simplest HLT paths:
pass-through of L1 selection
no additional selection
no bias with respectto L1

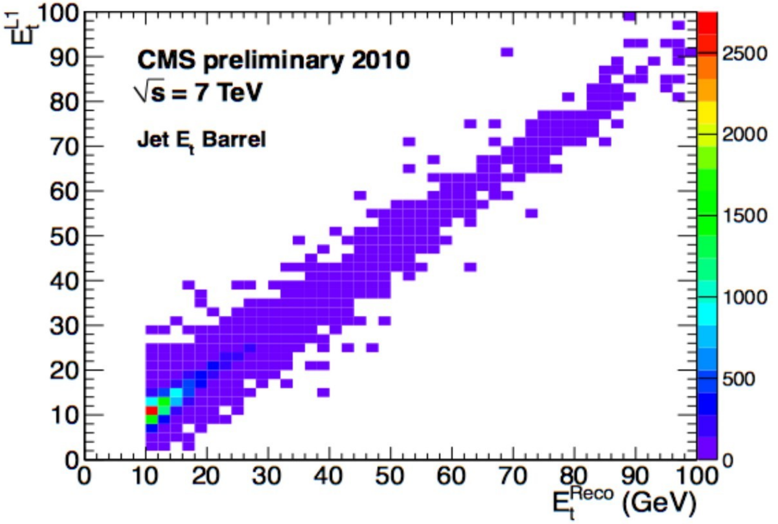

HLT path structure

improve resolution
increase reconstruction time

the simplest HLT paths:
pass-through of L1 selection
no additional selection
no bias with respectto L1

next step:
confirm L1 object with full granularity
fast reconstruction technique
improved resolution

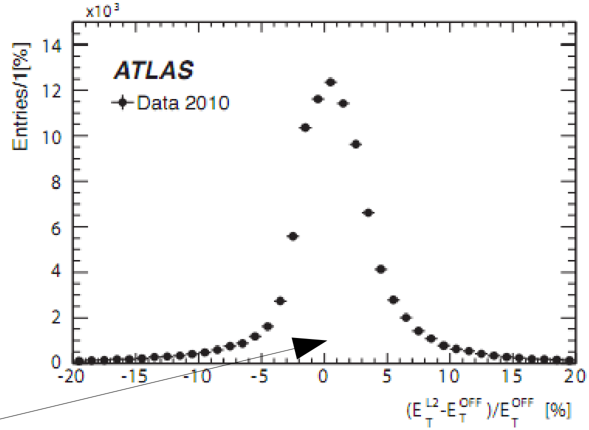

HLT path structure

improve resolution
increase reconstruction time

the simplest HLT paths:
pass-through of L1 selection
no additional selection
no bias with respectto L1

next step:
confirm L1 object with full granularity
fast reconstruction technique
improved resolution

continue adding complexity
improve quality of trigger object,
approaching offline resolution

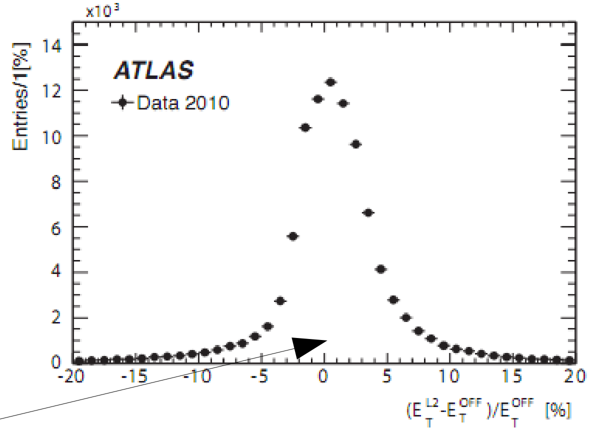

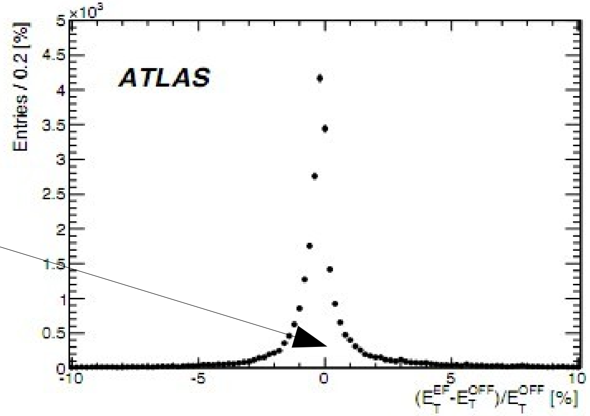

HLT Timing

look at L1 information

fast accept/reject

reconstruction of leptons and calorimetric object

e.g. muons, photons

full track reconstruction and particle flow

e.g. jets, tau

average time is critical

slow events have a large impact on the average time

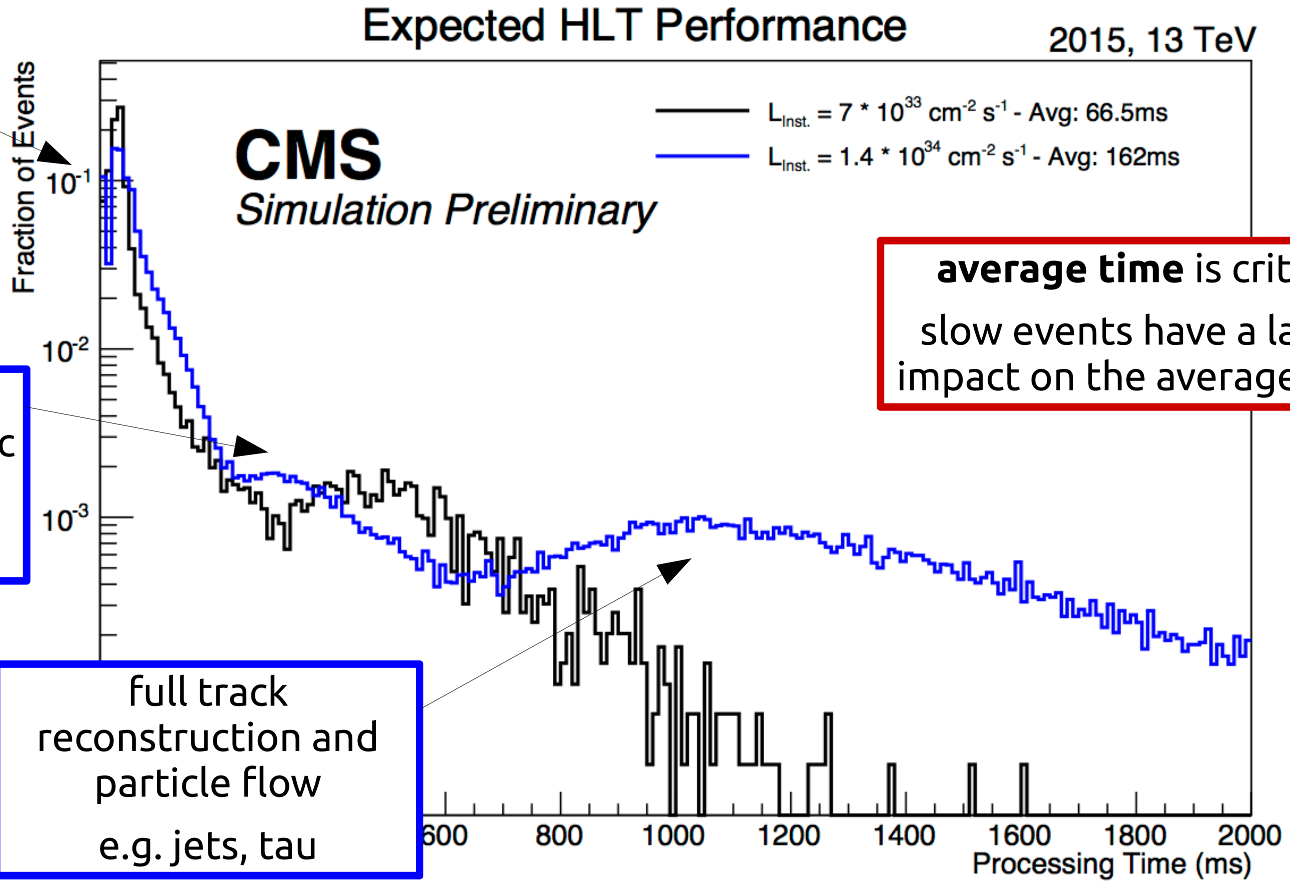

HLT menu

single source:
RAW data
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

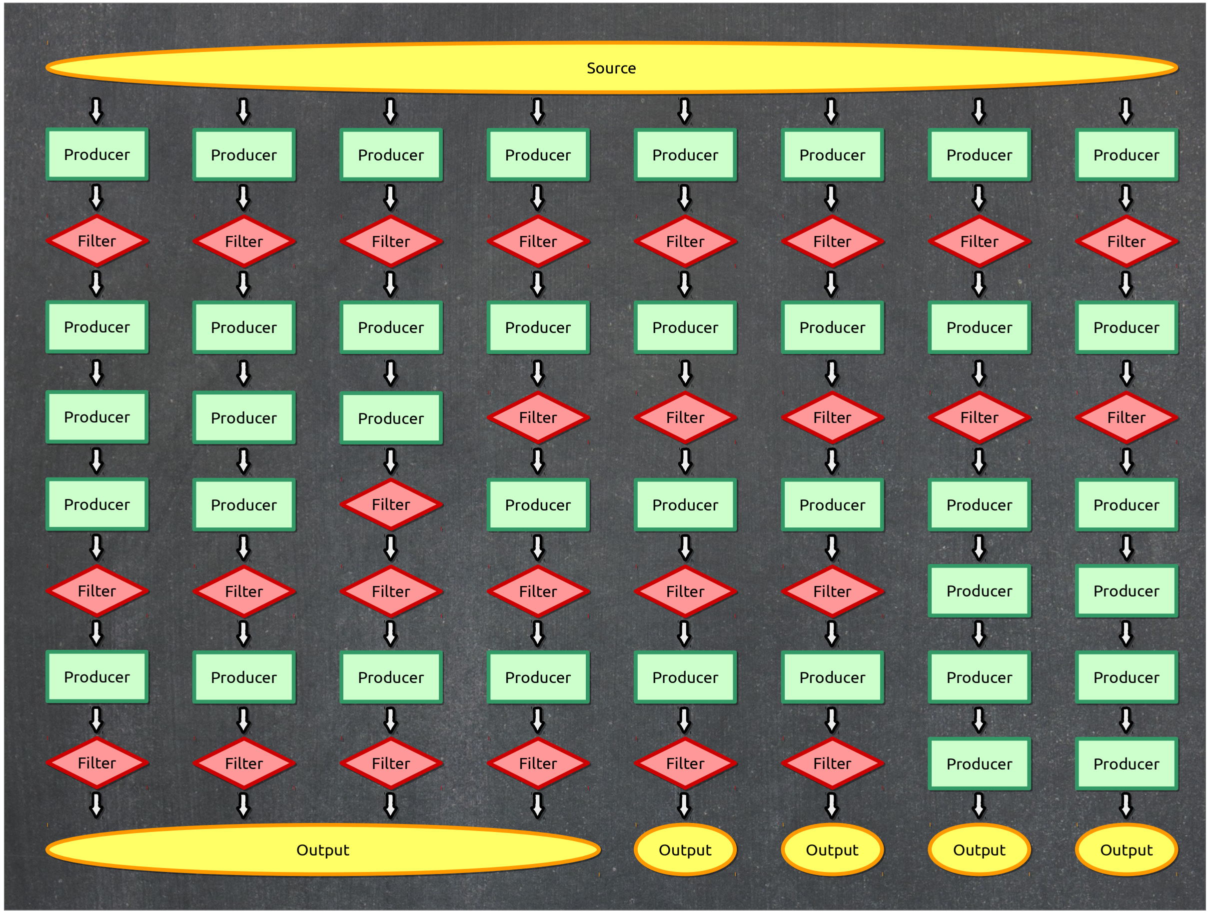

example: HLT electrons

- start from L1 e/g seeds with suficient energy
- reconstruct the cluster in the EM calorimenter
  - is the **cluster energy** high enough ?
  - does the **cluster shape** look like an electron or photon ?
  - **reject hadrons** (compare EM and H calorimenter energy)
  - is the candidated **isolated** in the calorimeters ?

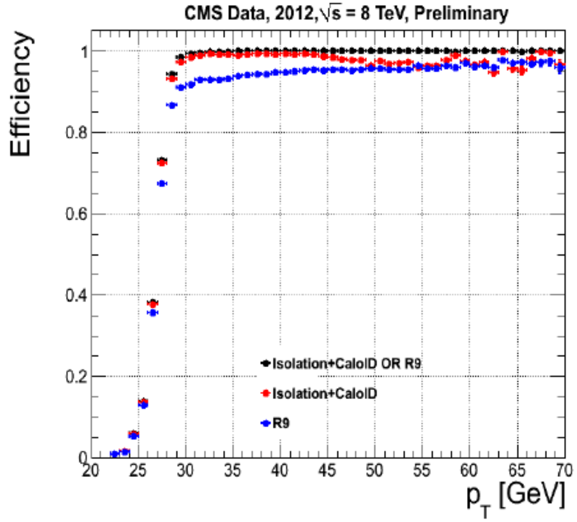

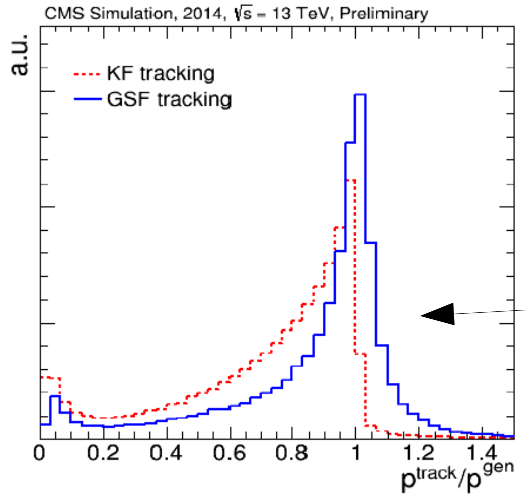

- look for **electrons**
  - recontruct the tracks in the pixel detector
  - is there a pixel track **pointing to the cluster** ?
  - **dedicated tracking** with the full tracker
  - is the track **compatible with the cluster** ?

example: HLT muons

- step 1: fast
  - read inner tracker and muon detector in a **region of interest** around the L1 muon candidates
  - assign muon pT using fast **look up tables**, based on the muon and tracker hits
  - is the pT high enough ?
- step 2: accurate
  - extrapolate to the collision point and reconstruct the muon track
  - is the muon isolated ?
  - compute pT using the tracking information
  - is the pT high enough ?

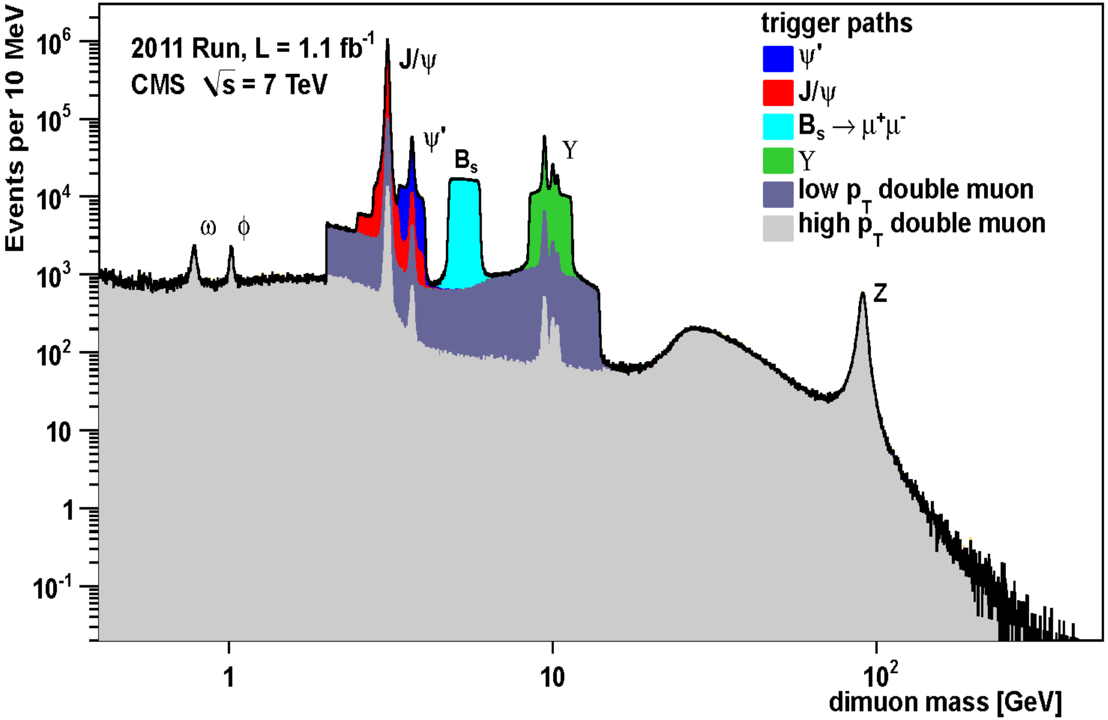

HLT and DAQ comparison

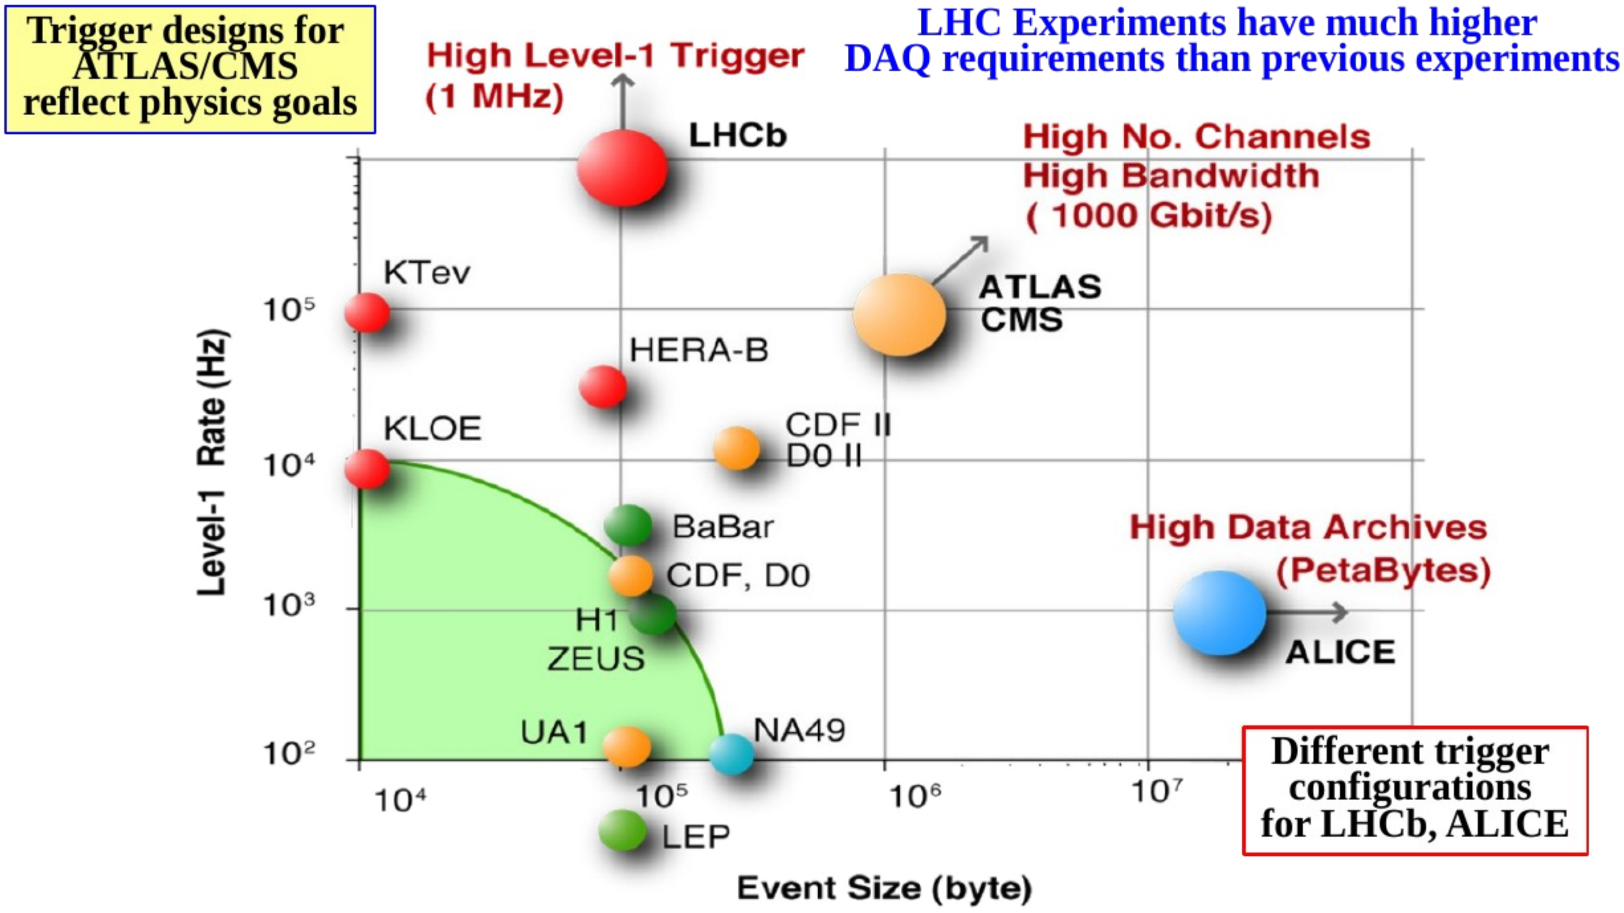

ALICE Trigger

Unique ALICE constraints

- Low rate of Pb-Pb collisions
- Very large events
- Slow tracking detector (TPC)

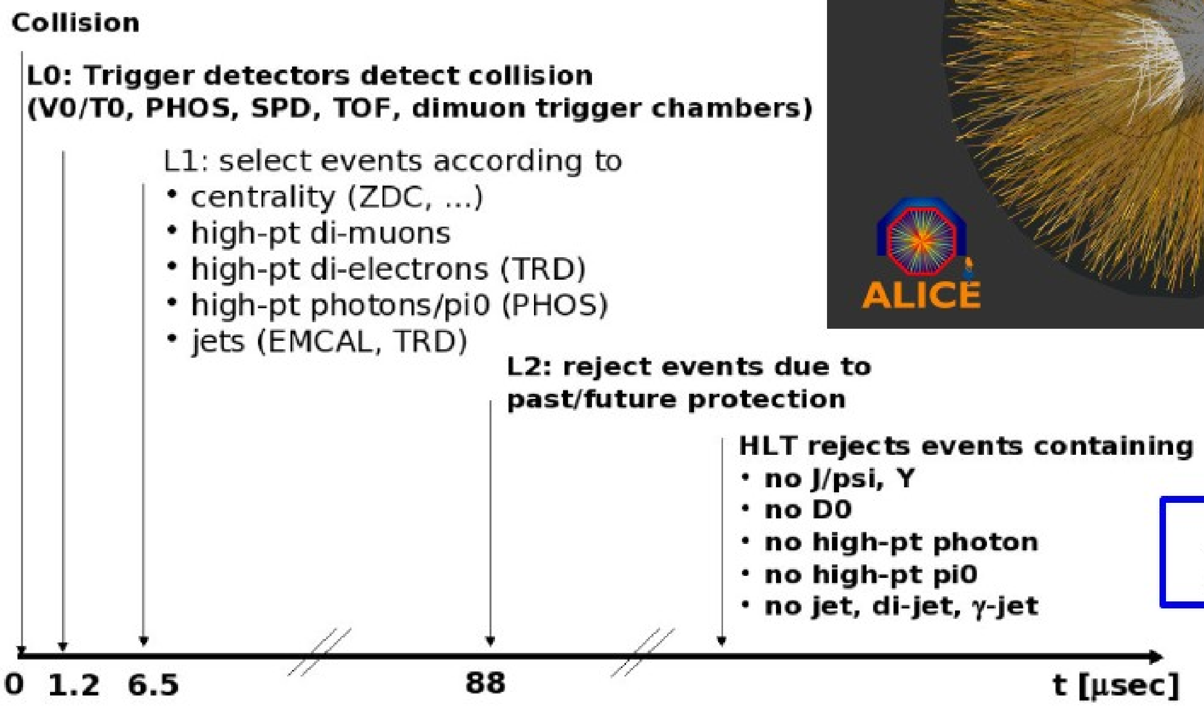

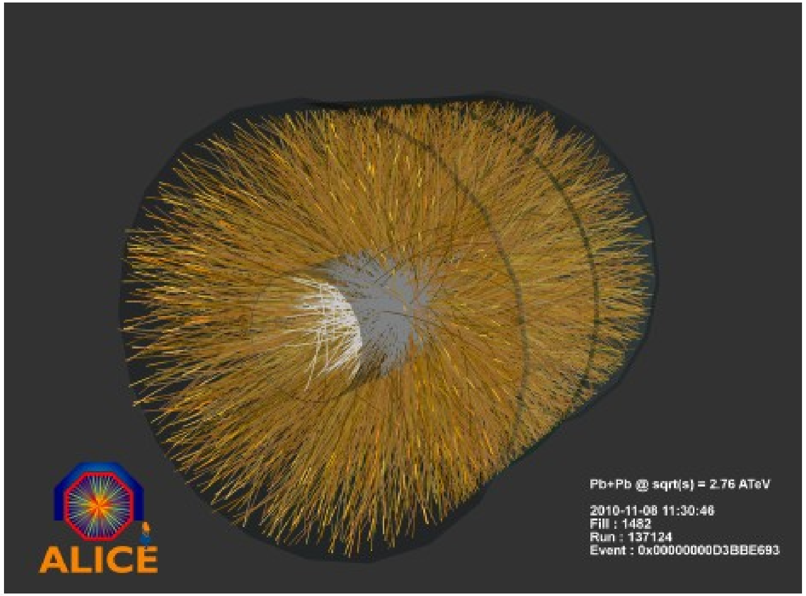

Three levels of hardware triggers

ALICE High Level Trigger

- HLT processing time dominated by track reconstruction
  - massively parallel problem
  - resort to a parallel solution and hardware acceleration:
- tracking on GPUs !

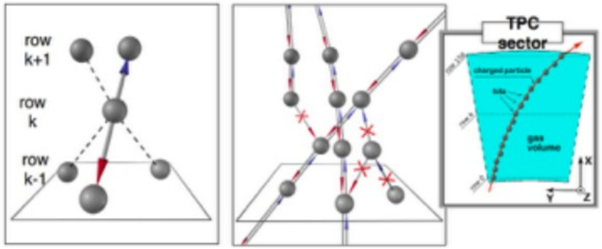

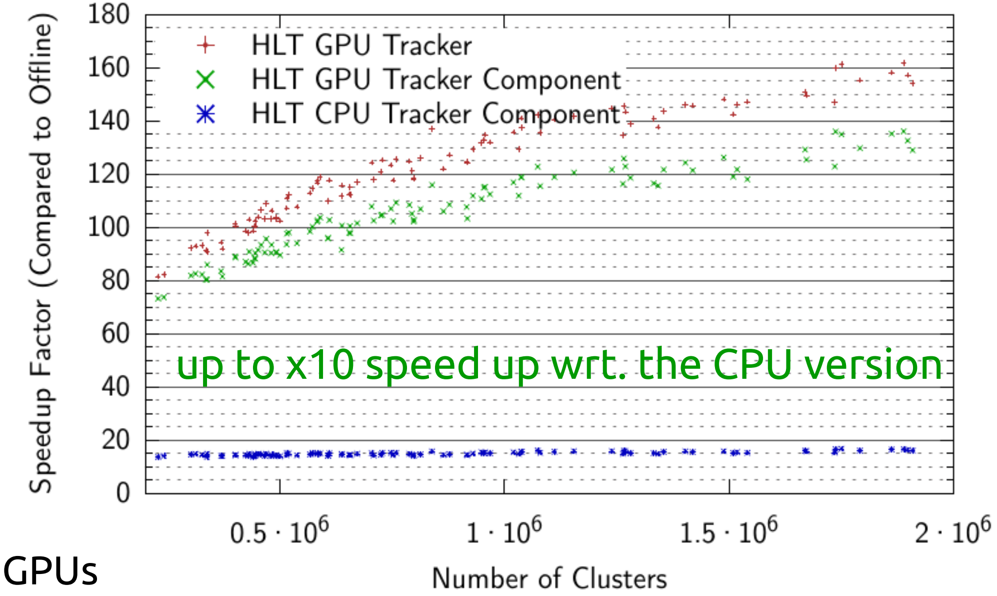

- cellular automaton track finder:
  - Run I farm: CUDA on nVidia GPUs
  - Run II farm: OpenCL on AMD FirePro GPUs
  - CPU version (x86 + OpenMP option)

LHCb Trigger

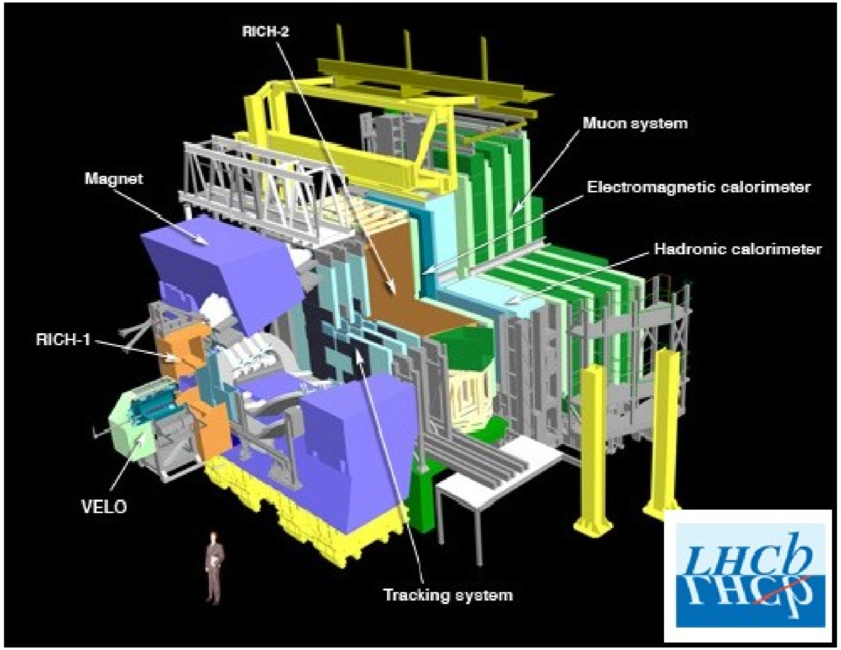

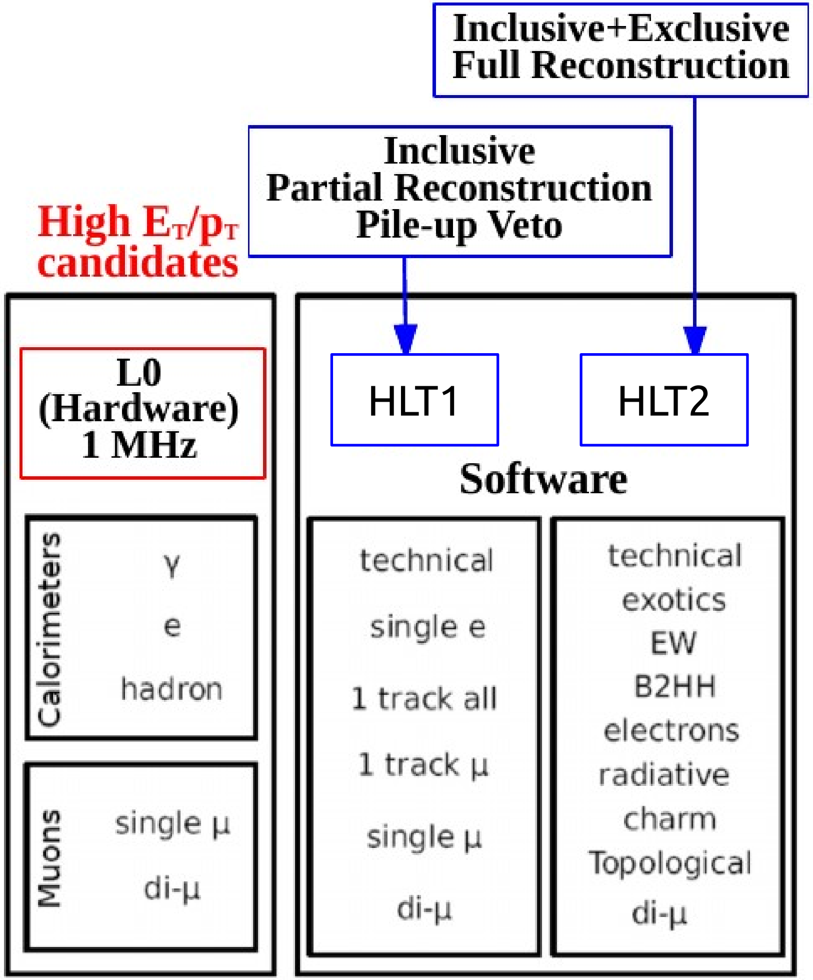

LHCb High Level Trigger

- LHC delivers *stable beams* during ~30% of the running period
  - HLT farm is idle ~70% of the time
  - is it possible to take advantage of these idle resources ?
- defer part of the HLT

- HLT1
  - first stage preselect events
  - runs quasi-online, at L1 accept rate
  - stores accepted events on a local disk
- HLT2
  - second stage performs the final event filtering
  - runs later, after HLT1 has terminated
- bonus points
  - HLT2 can use offline-level aligment

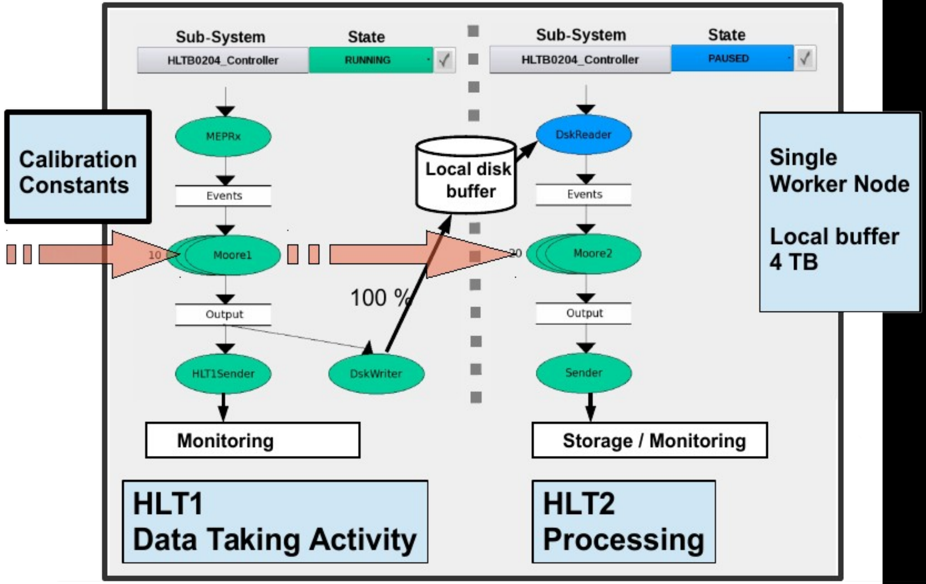

Triggers and Analyses

Triggers and Analyses

- As far as the data is concerned, the trigger is the first step towards publication
- But the order is a bit backward for physicists
- Why?

Triggers and Analyses

- Physicists start with an analysis idea
  - Determine what you want to look for (i.e. where you want to go)
  - Then figure out how to select the data
- There is little point in trying to do an analysis if every “interesting” event fails the trigger
- Want to build a trigger that has loose requirements that you tighten up offline
- Design a trigger to meet analysis goals, but...

Triggers and Analyses

- There are hundreds to thousands of physicists on an LHC collaboration
  - All are competing for the same resources
  - Only O(100) Hz of collision data available
    - At L = 10^{34}, this is roughly the rate of $W \to \ell\nu$ production!
- How do you make sure your (very important) data is kept for later analysis?
  - Need to meet physics needs with limited bandwidth
- Cutting at the trigger level throws away data forever
  - Potential bias to events that you analyze
  - Loss of interesting data

"The Trigger does not determine which Physics Model is right, only which Physics Model is left"

Trigger Efficiency

- In order to determine a cross section, you need to know your selection efficiency
  - Detector acceptance
  - Reconstruction efficiency
  - Trigger efficiency
- Your trigger is used to collect your data
  - You cannot blindly use your data to study efficiency
- Need an unbiased measurement of trigger efficiency
  - Random sample of pp collisions
  - Events collected by an orthogonal trigger
  - Use events collected by a looser (prescaled) trigger
  - Tag-and-Probe sample

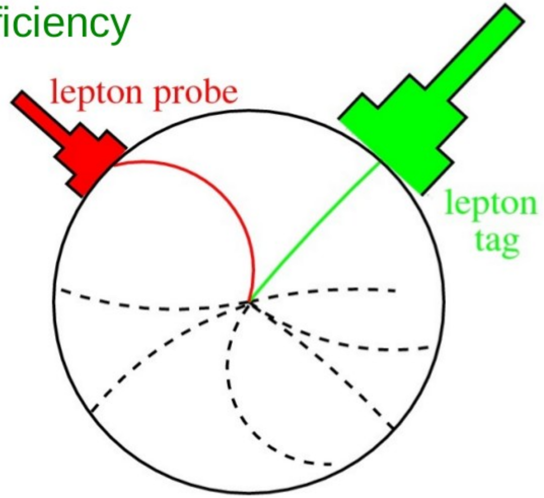

Trigger Efficiency

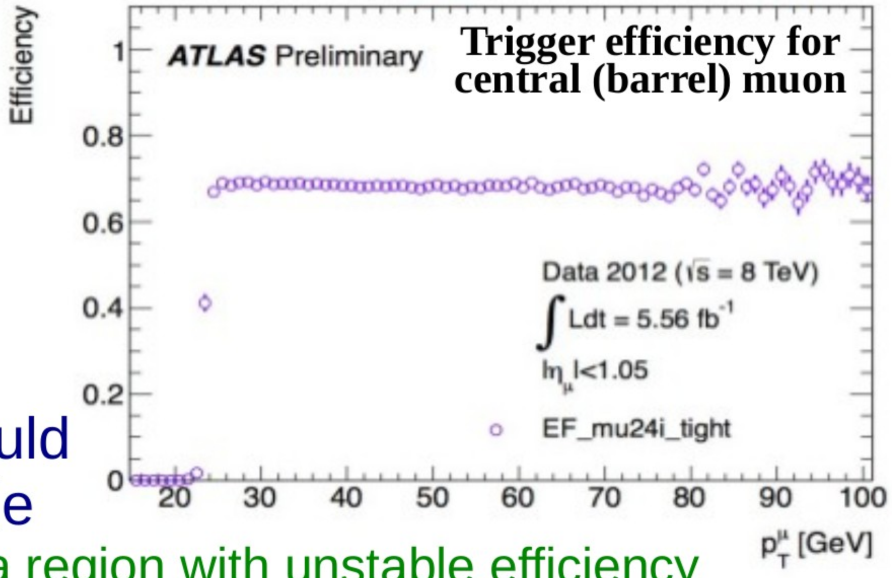

- Trigger efficiency is usually measured as a function of p_T and/or detector position
- We often speak of a trigger "turn-on" curve
- The turn-on curve should be as sharp as possible
  - Prevents working in a region with unstable efficiency
- Even when flat, the efficiency may not be 100%
  - Important to consider in the analysis

Turn-on curve

CMS Electron Trigger
Early in 8 TeV run

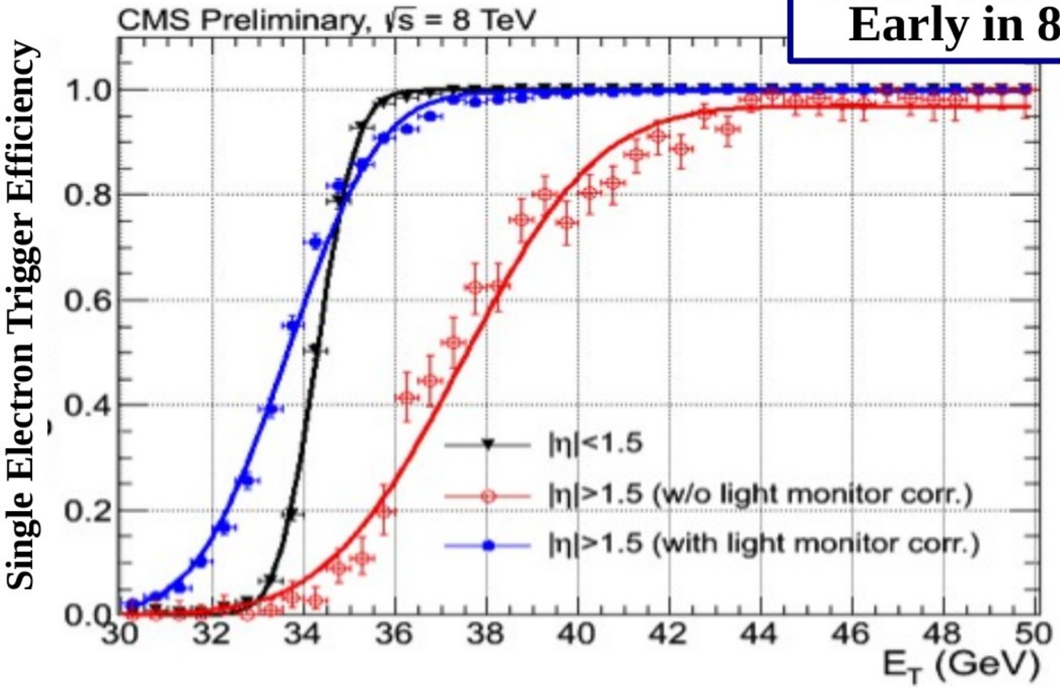

Adjust trigger conditions to account for a changing detector
Increased luminosity, increased light loss in CMS EM calorimeter

Building a trigger

- Imagine you need events with a Z boson
  - Standard Model, Higgs→ZZ, useful for Z' searches, ...
- How do you collect these events online?

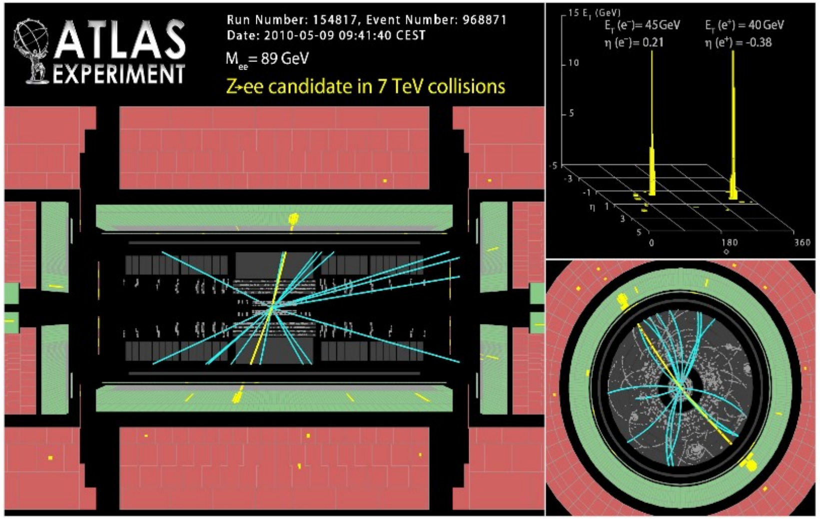

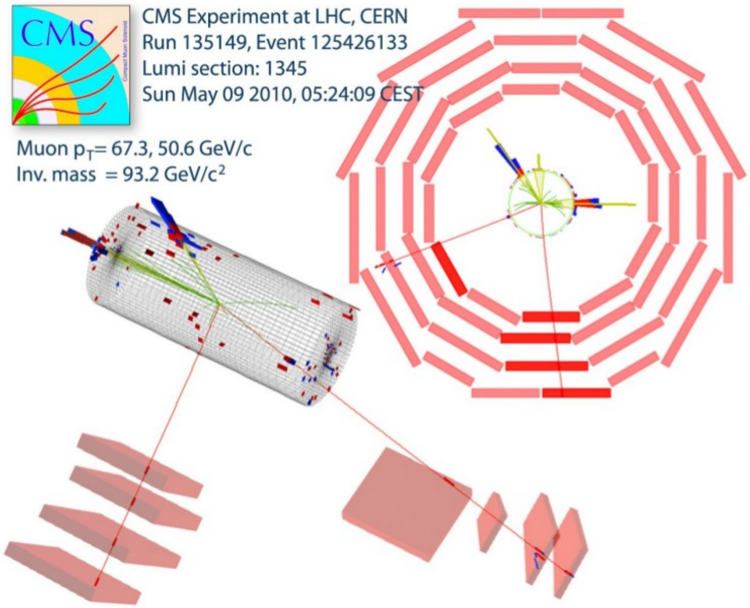

Building a trigger

- Isolated high p_T leptons are rarely produced in a typical pp collision
  - Every Z decay has two of them!
  - So, construct a trigger that requires high p_T leptons
- General strategy for building a trigger
  - The simpler, the better
  - Be as inclusive as possible
  - Robust design
  - Redundancy

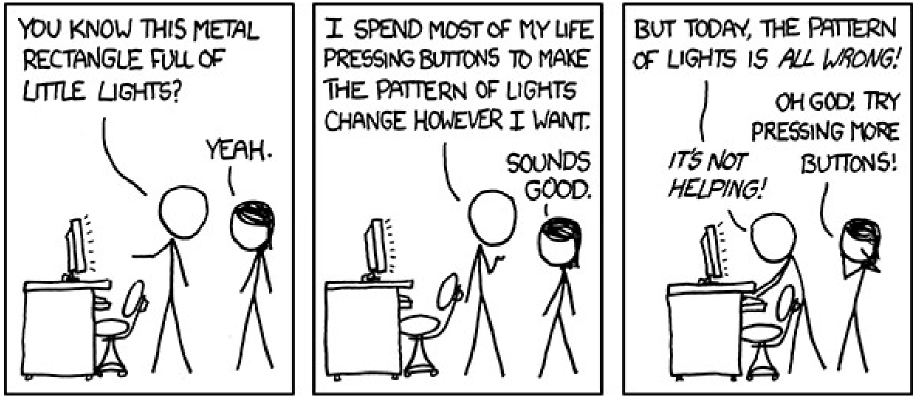

Building a trigger

- Simple triggers are
  - Easier to commission
  - Easier to debug
  - Easier to understand
- If possible, create a new (tighter) trigger from an older (more inclusive) trigger
  - At high rate, or limited bandwidth, more inclusive triggers tend to be prescaled

Trigger Strategy

- Simple
- Inclusive
- Robust design
- Redundancy

Evolution of a trigger

- Initially, we started with a single lepton trigger
  - Efficiency for Z events was very high
  - Take our (hypothetical) single muon trigger as an example
    - Let's say we estimated the muon efficiency to be 90% using tag and probe techniques
    - Our trigger efficiency for $Z \rightarrow \mu\mu$ should be...

Evolution of a trigger

- Initially, we started with a single lepton trigger
  - Efficiency for Z events was very high
  - Take our (hypothetical) single muon trigger as an example
    - Let's say we estimated the muon efficiency to be 90% using tag and probe techniques
    - Our trigger efficiency for $Z \rightarrow \mu\mu$ should be...**99%**

81%	**9%+9%=18%**	**1%**
Probability that both muons triggered the event	Probability that only one muon triggered the event	Probability that neither muon triggered the event

Evolution of a trigger

- By using minimal (simple) trigger strategies, we have nearly 100% efficiency in our selection
- By making our trigger more complicated by adding a second muon (or electron), our efficiency drops
  - Must account for such effects in the analysis

81%	**9%+9%=18%**	**1%**
Probability that both muons triggered the event	Probability that only one muon triggered the event	Probability that neither muon triggered the event

Trigger Strategy

- So, we wish to collect events with Z decays online
  - What should we do?
- Easiest solution: Use single lepton triggers
  - Two leptons (electrons or muons) from the Z as either could trigger the event
  - If you choose a double lepton trigger, you are insisting online that both leptons pass trigger requirements
    - Best to wait until you must do this
    - Determined by LHC conditions, physics goals

Trigger Strategy
- Simple
- Inclusive
- Robust design
- Redundancy

What is done online cannot be undone...

Trigger Strategy

- What happens if your trigger has a large rate?
  - Remember, we can only save O(100) events/second
- Possible solution: Get Help!
- Hopefully many physics analyses (besides yours) could use the same trigger
  - Likely we are not the only group looking for lepton triggers
    - Standard Model: **Z**, W, top
    - SUSY
    - Exotic signatures
    - ...
- A trigger is easier to keep if most of the collaboration is using it

Trigger Strategy
- Simple
- Inclusive
- Robust design
- Redundancy

Trigger Strategy

Don't design your trigger expecting this...

- Your trigger is going online, so it should run on every kind of event
- Prepare for "real life", which includes pathological events
- Minimize (to ZERO) the number of crashes due to trigger design

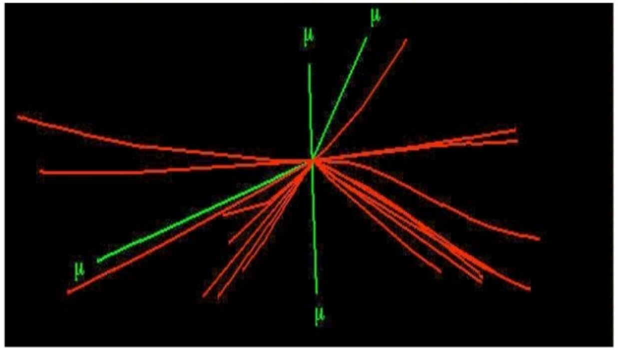

...when life might look like this

$H \to ZZ \to 4\mu$
(and 25 pileup events),
with and without
$p_T > 25$ GeV track requirement

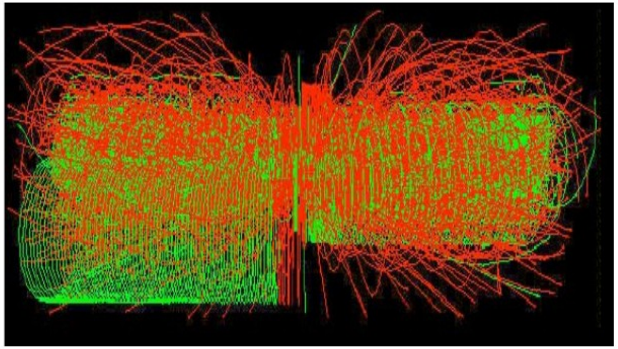

Trigger Strategy

- Simple
- Inclusive
- Robust design
- Redundancy

Trigger Strategy

- It is very useful if your analysis can be selected using more than one trigger
  - Will help understand any potential trigger bias
  - If one trigger has problems (detector or LHC conditions leading to higher rate), you can still get your data
- Try to introduce tighter triggers online before they are necessary
  - Allows triggers to collect data before they are strictly necessary
  - Provides consistency for physics analysis, opportunity to study new trigger on existing data

Trigger Strategy
- Simple
- Inclusive
- Robust design
- Redundancy

Trigger Strategy

- Trigger strategy with a concrete example
- Collecting Z events using single electron, single muon triggers
  - High p_T, isolated leptons are rare in pp collisions
    - Much of the physics (and hence the detectors) designed around this fact
  - Lots of consumers in the community, so we can use a "common" trigger
  - (Let's assume that the trigger has been robustly tested and is working without problems online)
- We have back-up (redundant) triggers in place and ready for higher luminosity
  - Single electron/muon triggers with tighter requirements
  - Double electron, double muon triggers also ready

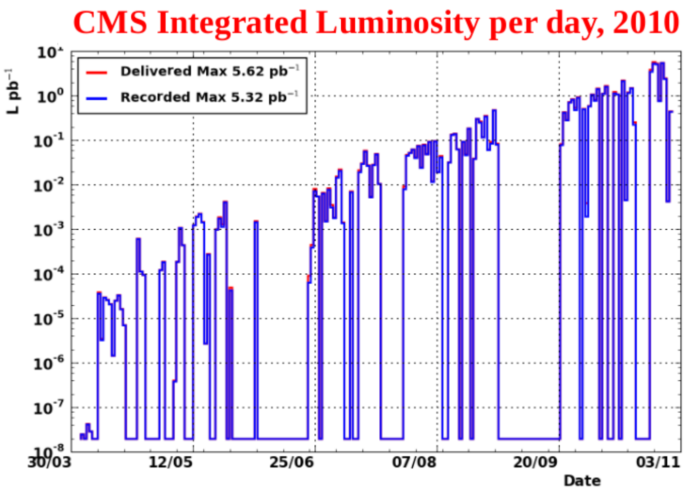

Trigger Strategy

Once you have the data, analysis awaits!

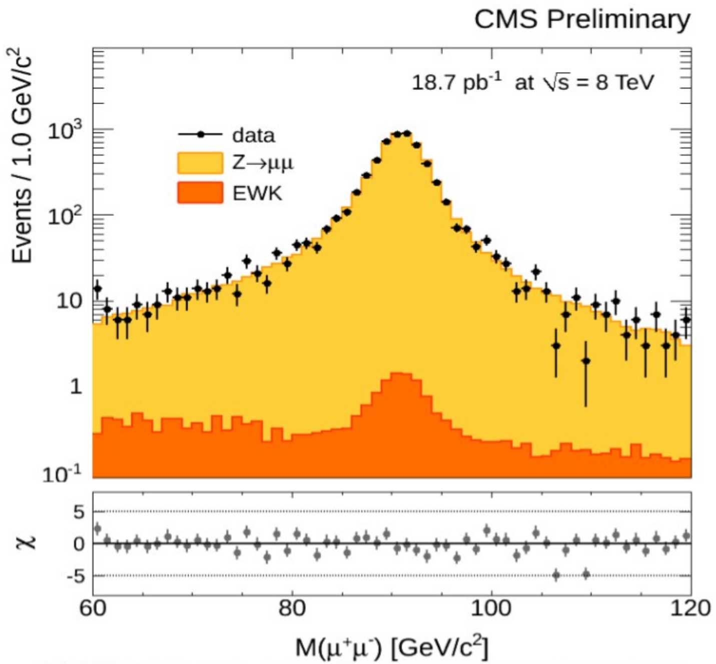

W/Z cross section measurement at 8 TeV (CMS, SMP-12-011)

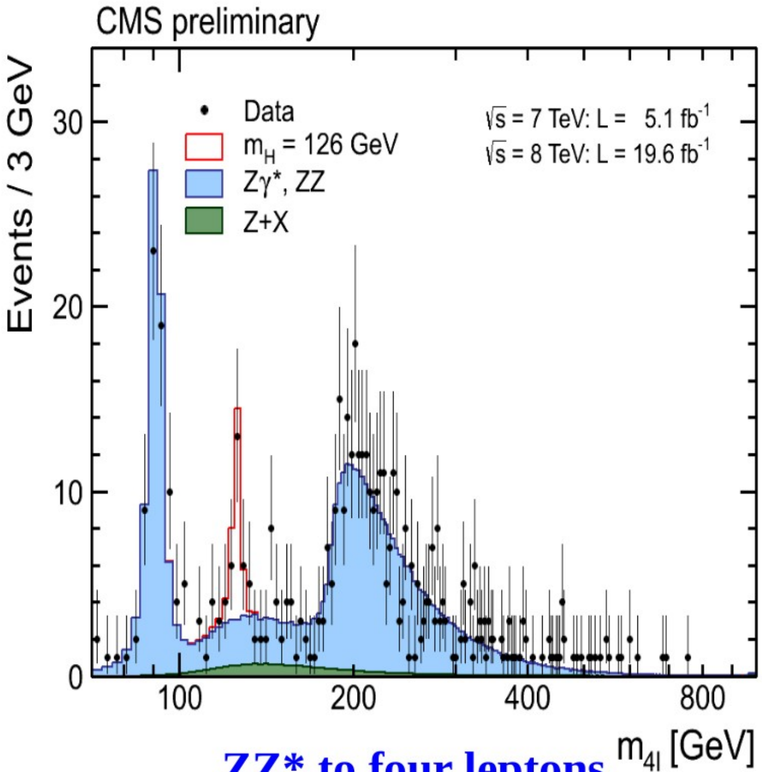

ZZ* to four leptons 2011/12 data (CMS, HIG-13-002)

Trigger Strategy

- You should always look ahead, even when working with the data you have
  - Always more to explore, additional properties to investigate
- The LHC is constantly improving
  - Higher instantaneous luminosity, so rate of W, Z, H, ... production constantly increasing
- Very likely that our first trigger idea is now obsolete
  - Improvements in software will increase efficiency
  - Additional filters in trigger path increase purity
    - But these filters reduce efficiency
  - Is it time to move to double electron/muon triggers?

Most Important: How do our trigger choices impact the analysis, and how do we adapt?

Conclusions

Conclusions

- the role of the trigger is to **maximise the physics** reach of an experiment
  - within the constraints of the detector
  - data acquisition, online and offline storage and processing
- reducing the event rate from the LHC collision rate
  - to what the detector can actually read out
  - to what can be written to disk and analysed
- **all the events that are not selected by the trigger are lost, forever !**
- choose a trigger strategy for your analyses as **efficient** and **robust** as possible
  - **simple** and **inclusive** triggers
  - **redundat** possibilities