Data acquisition and Trigger (with emphasis on LHC)

Monika Wielers (RAL)

- Introduction
  - Data handling requirements for LHC
- Design issues: Architectures
  - Front-end, event selection levels
- Trigger
- Future evolutions
- Conclusion

DAQ challenges at LHC

- Challenge 1
  - Physics – Rejection power
  - Requirements for TDAQ driven by rejection power required for the search of rare events
- Challenge 2
  - Accelerator – Bunch crossing frequency
  - Highest luminosity needed for the production of rare events in wide mass range
- Challenge 3
  - Detector – Size and data volume
    - Unprecedented data volumes from huge and complex detectors

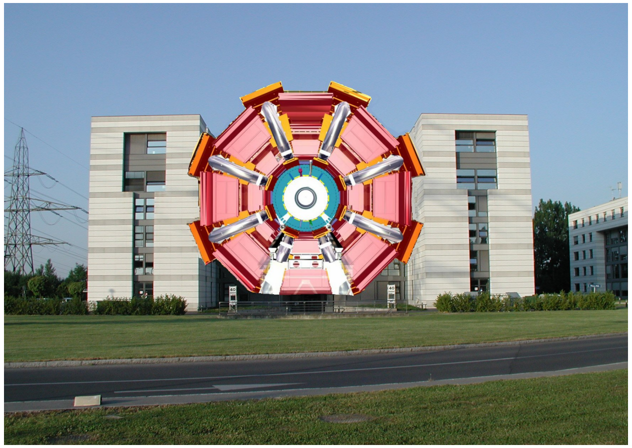

Challenge 1: Physics

- Cross sections for most processes at the LHC span ~10 orders of magnitude
- LHC is a factory for almost everything: t, b, W, Z…
- But: some signatures have small branching ratios (e.g. $H \to \gamma\gamma$, BR ~10^{-3})

Process	Production Rate $10^{34}\ cm^{-2}s^{-1}$
inelastic	~1 GHz
bbbar	5 MHz
$W \to l\nu$	150 Hz
$Z \to l\nu$	15 Hz
ttbar	10 Hz
Z’	0.5 Hz
H(125) SM	0.4 Hz

- $L=10^{34}\ cm^{-2}s^{-1}$: Collision rate: ~10^9 Hz. event selection: ~$1/10^{13}$ or 10^{-4}Hz !

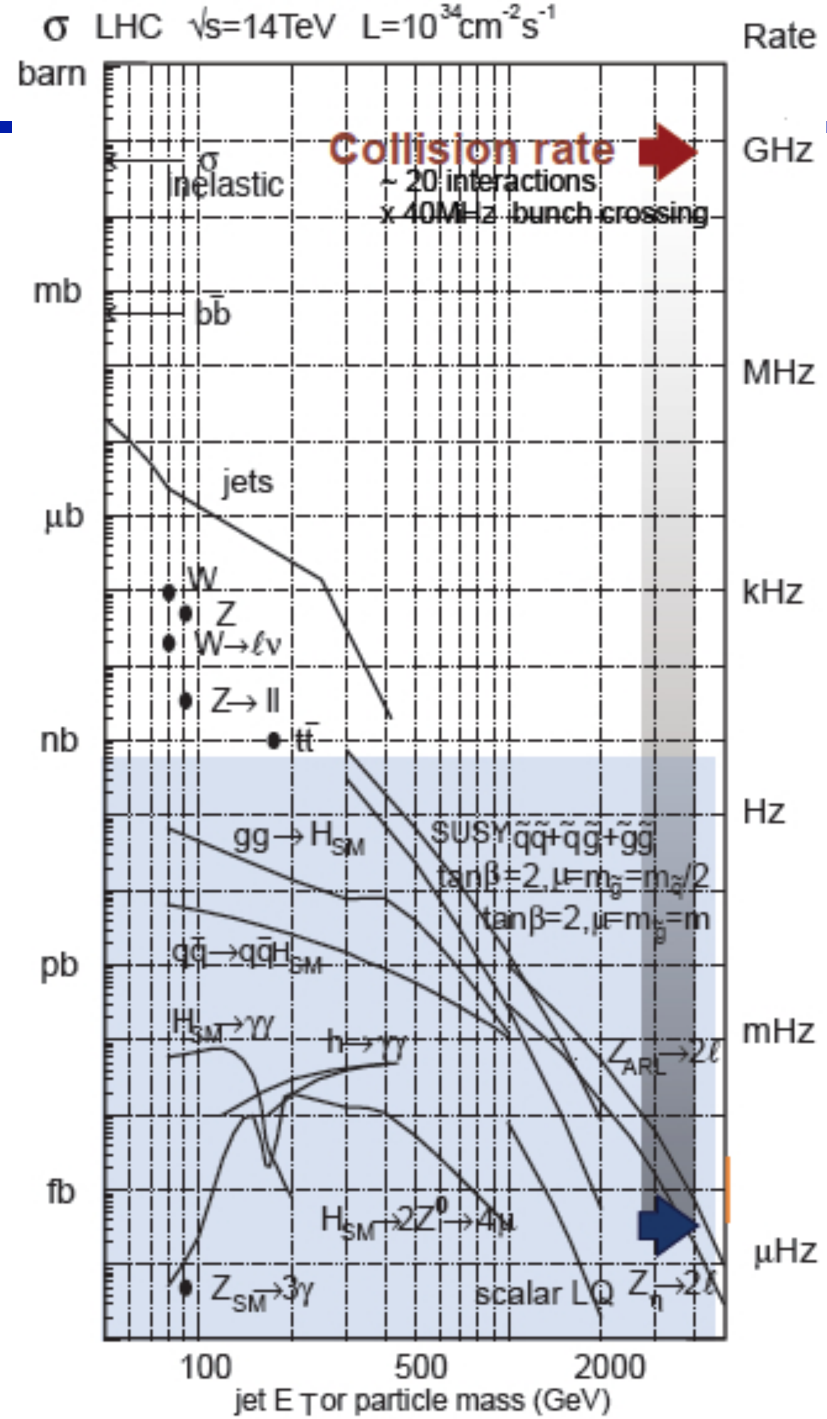

Challenge 1: Physics

- Requirements for TDAQ driven by the search for rare events within the overwhelming amount of "uninteresting" collisions
- Main physics aim
  - Measure Higgs properties
  - Searches for new particles beyond the Standard Model
    - Susy, extra-dimensions, new gauge bosons, black holes etc.
  - Plus many interesting Standard Model studies to be done
- All of this must fit in ~1 kHz of data written out to storage
- Not trivial, $W \rightarrow l\nu$: 150 Hz @ 10^{34} $cm^{-2}s^{-1}$
  - "Good" physics can become your enemy!

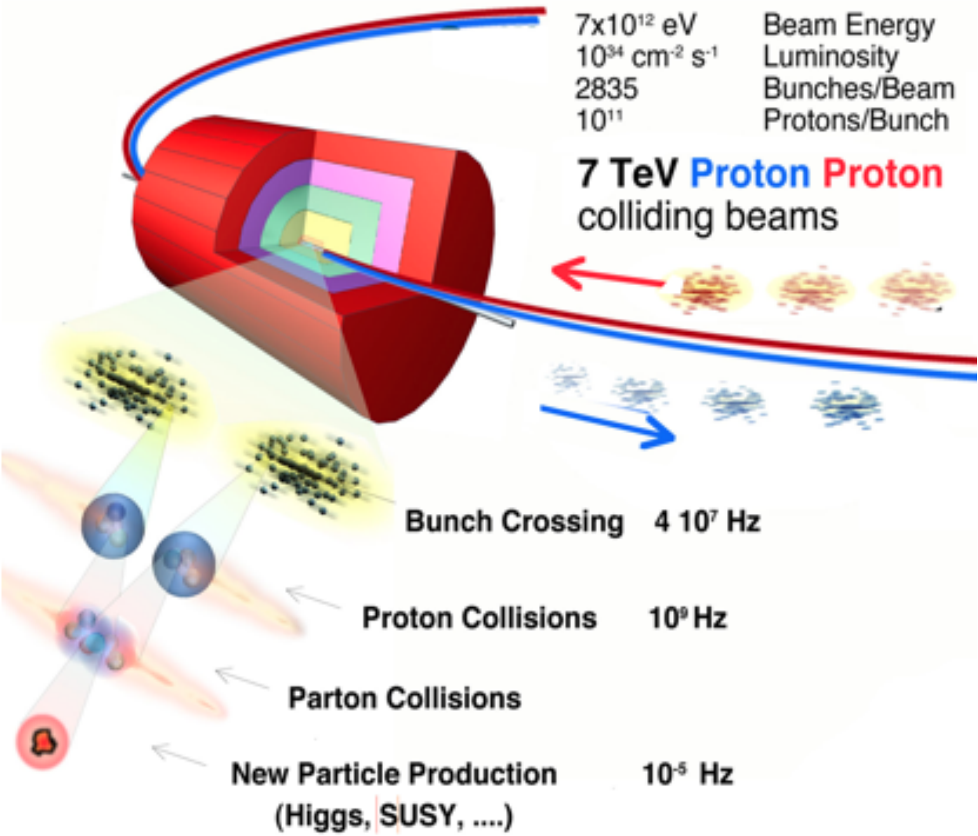

Challenge 2: Accelerator

- Unlike e^+e^- colliders, proton colliders are more ‘messy’ due to proton remnants
- Multiple collisions per bunch crossing
  - Currently ~20-30 overlapping p-p interactions on top of each collision (pile-up) ➔ >1000 particles seen in the detector!

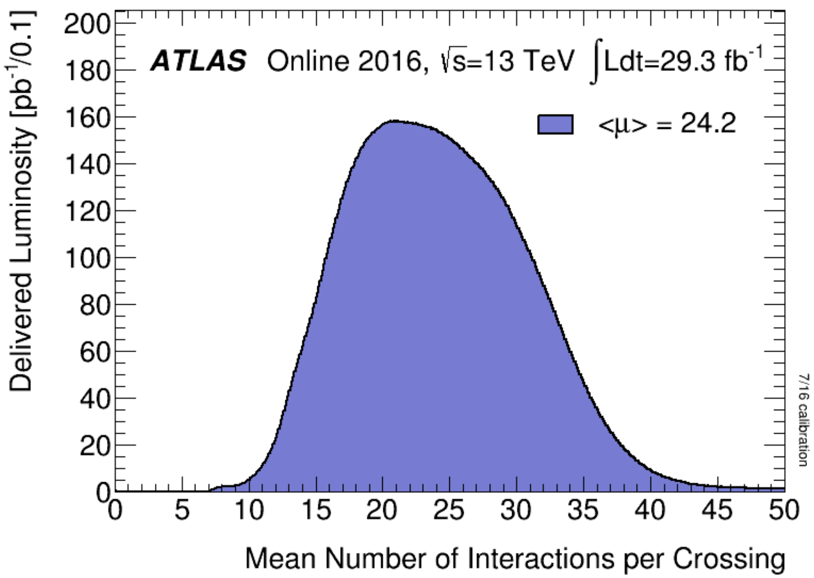

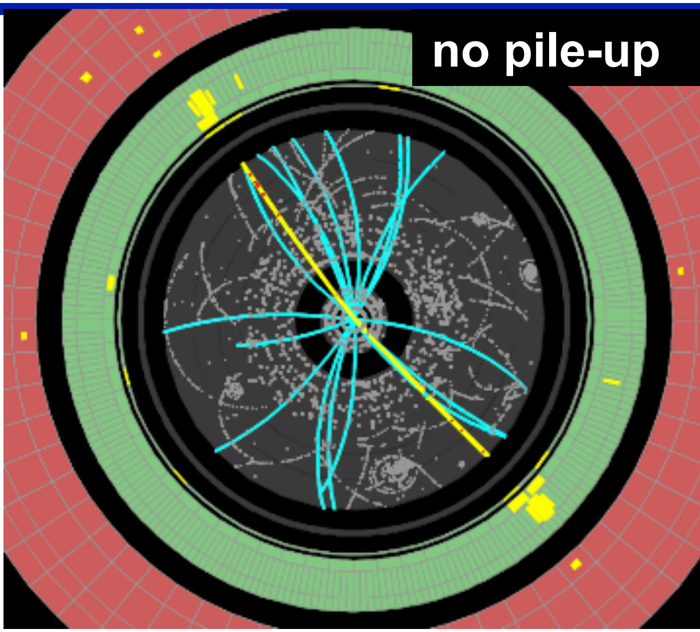

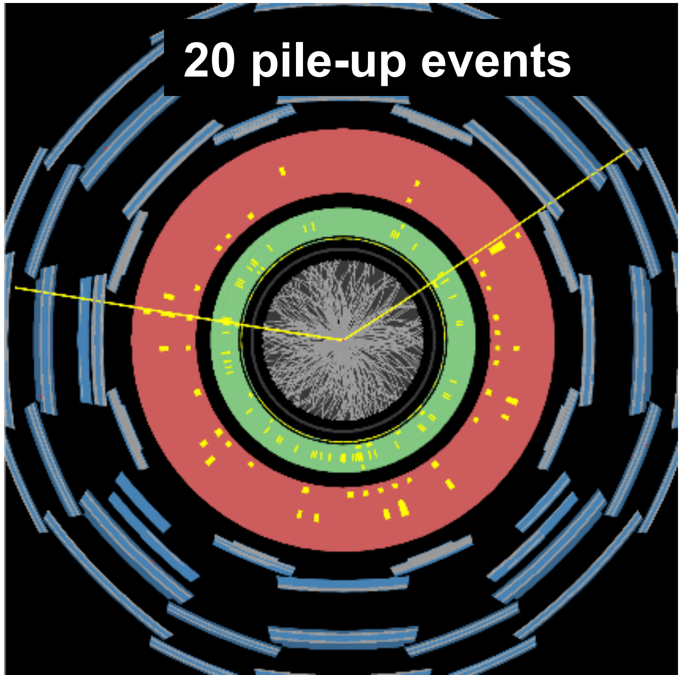

Challenge 3: Detector

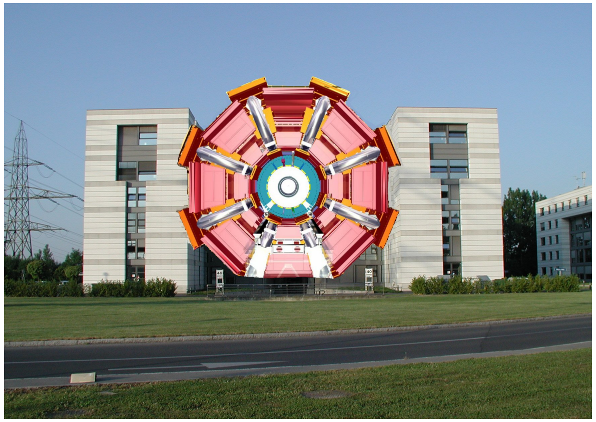

- Besides being huge: number of channels are $O(10^6\text{-}10^8)$ at LHC, event sizes ~1 MB for pp collisions, 50 MB for pb-pb collisions in Alice
  - Need huge number of connections
- Some detectors need > 25ns to readout their channels and integrate more than one bunch crossing's worth of information (e.g. ATLAS LArg readout takes ~400ns)
- It's On-Line (cannot go back and recover events)
  - Need to monitor selection - need very good control over all conditions

Let’s build a Trigger and DAQ for this

- What do we need?

Let’s build a Trigger and DAQ for this

- What do we need?
  - Electronic readout of the sensors of the detectors (“front-end electronics”)
  - A system to collect the selected data (“DAQ”)

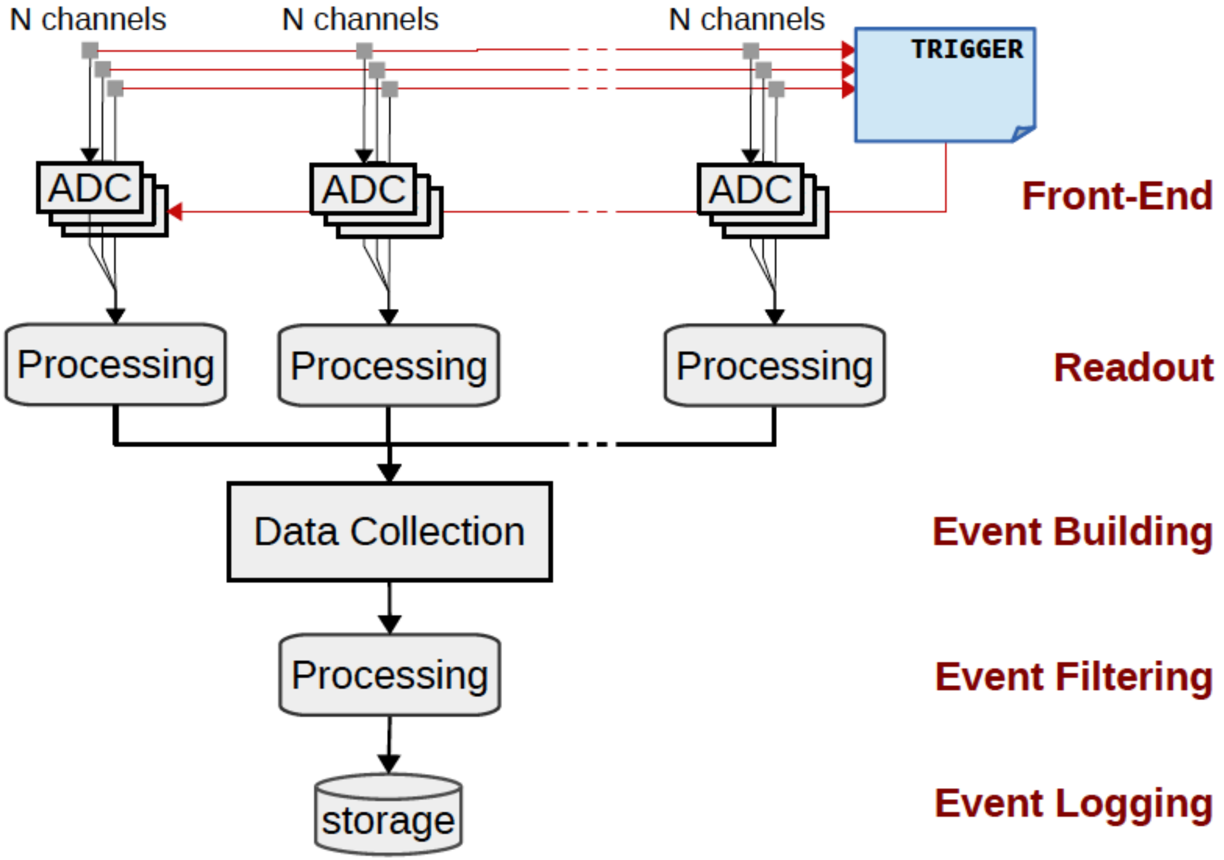

Let's build a Trigger and DAQ for this

- What else do we need?
  - A system to keep all those things in sync ("clock")
  - Data belonging to the same bunch crossing must be processed together
  - Particle time of flight, cable delays, electronic delays all

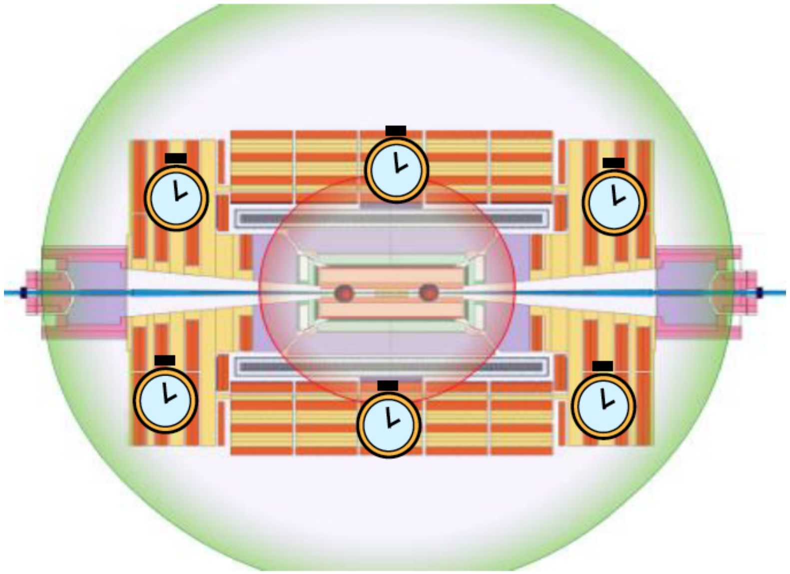

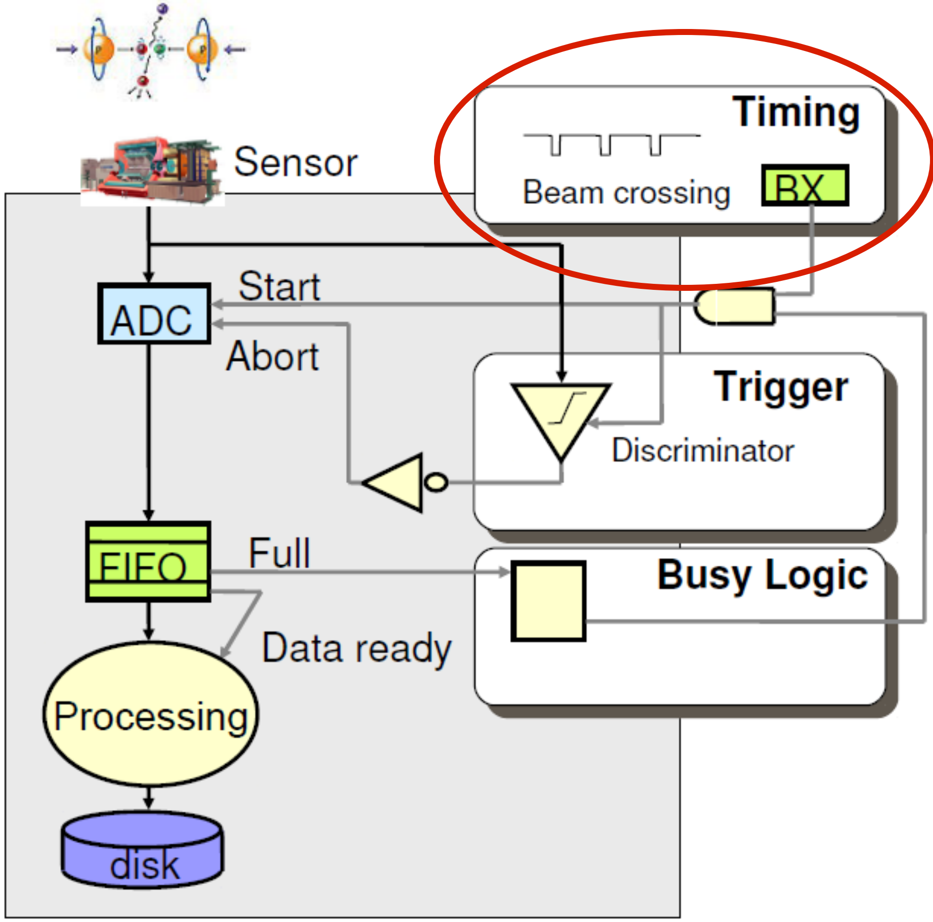

Let’s build a Trigger and DAQ for this

- What do we need?
  - Electronic readout of the sensors of the detectors (“front-end electronics”)
  - A system to collect the selected data (“DAQ”)
  - A system to keep all those things in sync (“clock”)
  - A trigger – multi-level due to complexity

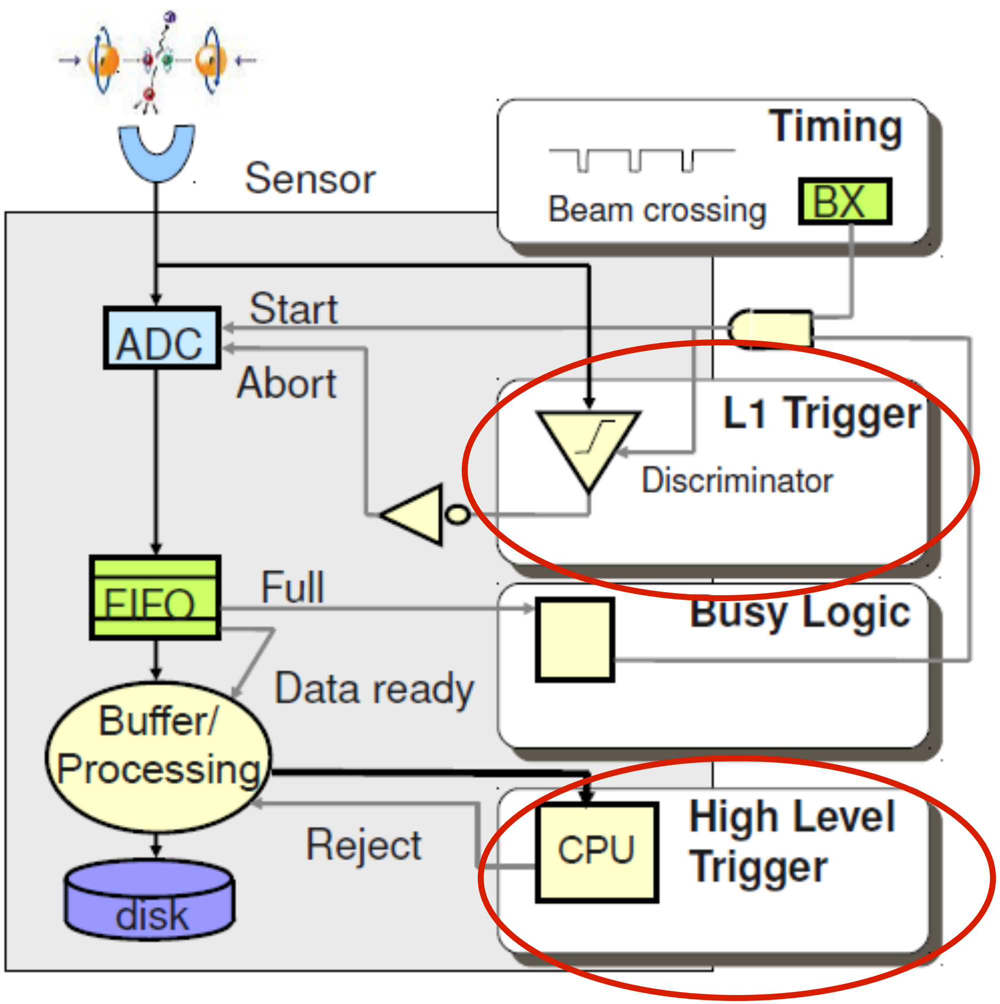

Let’s build a Trigger and DAQ for this

- What do we need?
  - Electronic readout of the sensors of the detectors (“front-end electronics”)
  - A system to collect the selected data (“DAQ”)
  - A system to keep all those things in sync (“clock”)
  - A trigger – multi-level due to complexity
  - A Control System to configure, control and monitor the entire DAQ

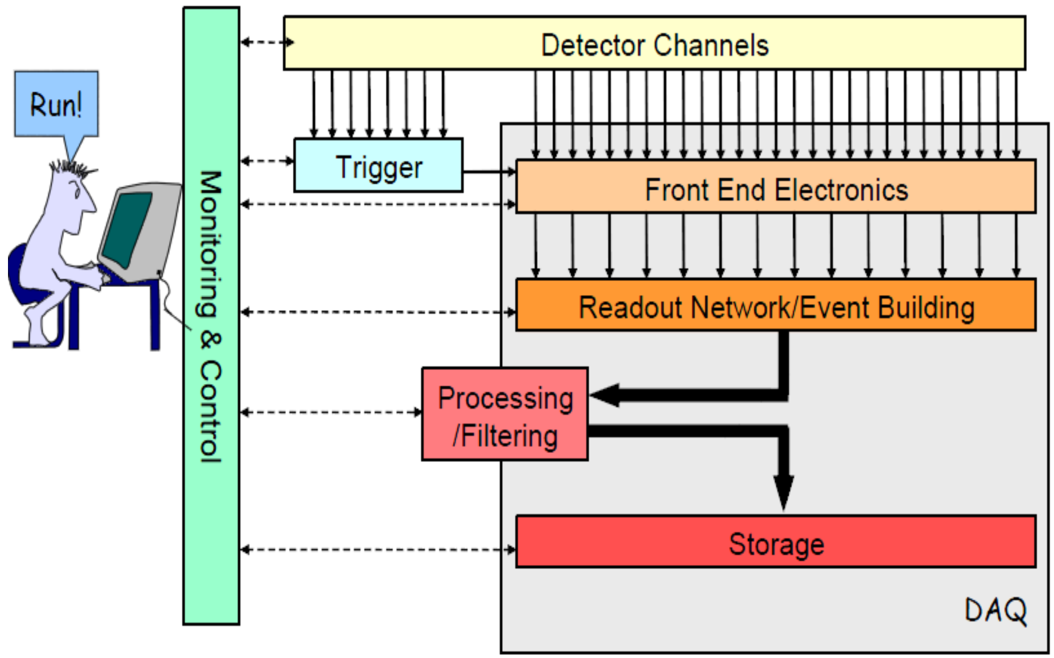

Let’s look more at the trigger part…

Multi-level trigger system

- Sometime impossible to take a proper decision in a single place
  - too long decision time
  - too far
  - too many inputs
- Distribute the decision burden in a hierarchical structure
  - Usually $\tau_{N+1} >> \tau_N$, $f_{N+1} << f_N$
- At the DAQ level, proper buffering must be provided for every trigger level
  - absorb latency
  - De-randomize

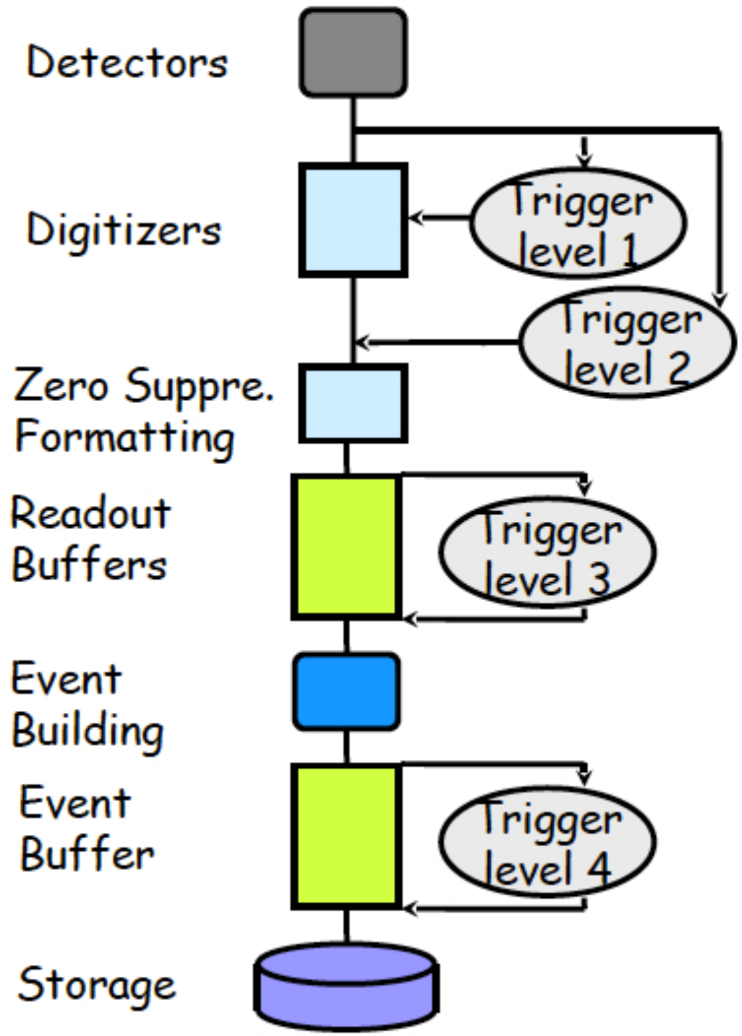

LHC DAQ phase-space

- When LHC experiments were designed back in the 90’
  - Raw data storage capped at ~ 1 PB / year per experiment

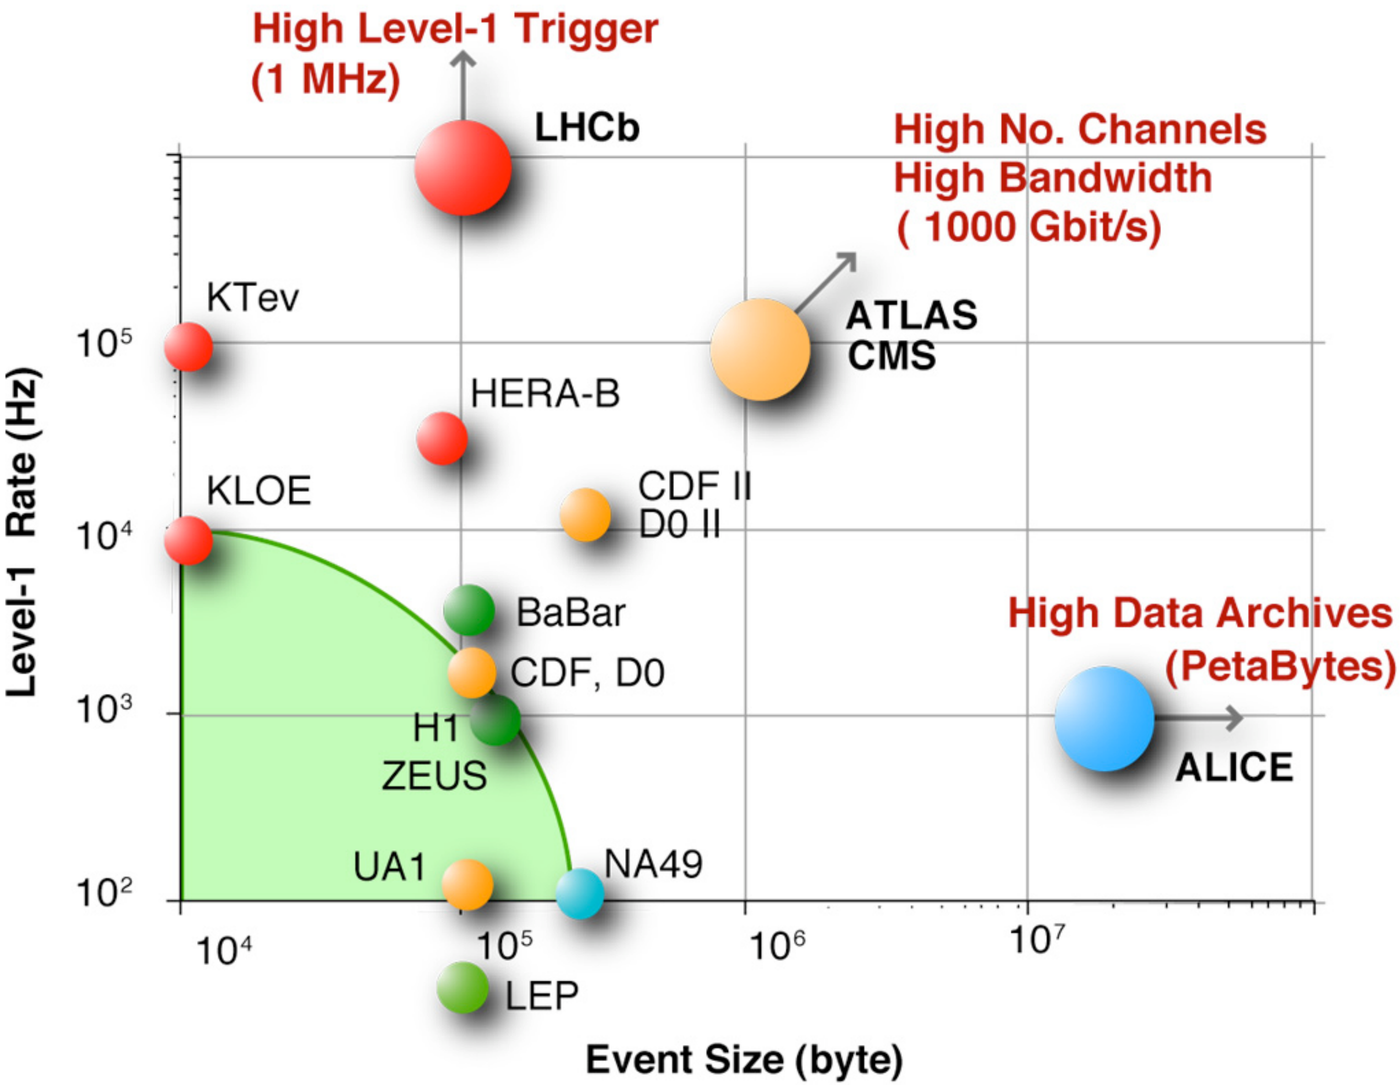

Hardware Trigger (L0, L1)

- Custom electronics designed to make very fast decisions
  - Application-Specified Integrated Circuits (ASICs)
  - Field Programmable Gate Arrays (FPGAs)
    - Possible to change algorithms after installation
- Must cope with input rate of 40 MHz
  - Reduce rate from 40 MHz to ~100 kHz
  - Otherwise cannot process all events
  - Event buffering is expensive, too
- Use pipeline for holding data during L1 processing
  - Digital/analog custom front-end pipelines
  - Parallel processing of different inputs as much as possible

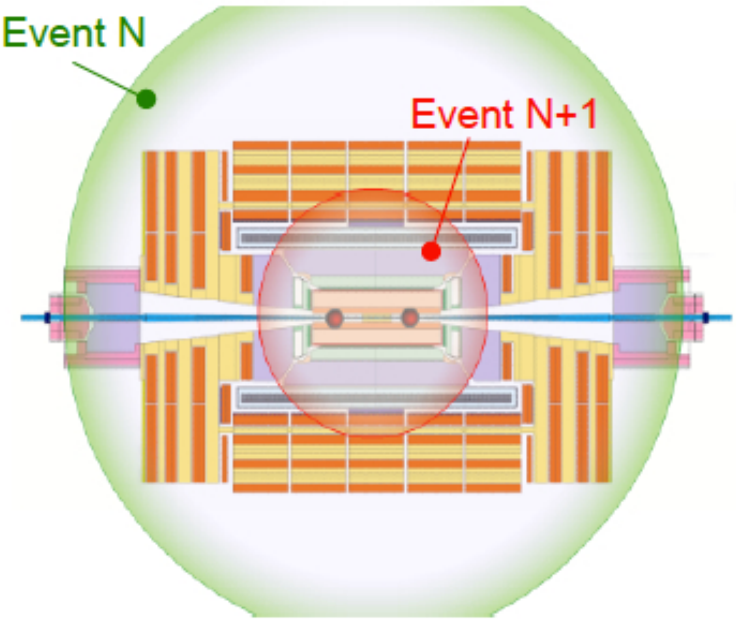

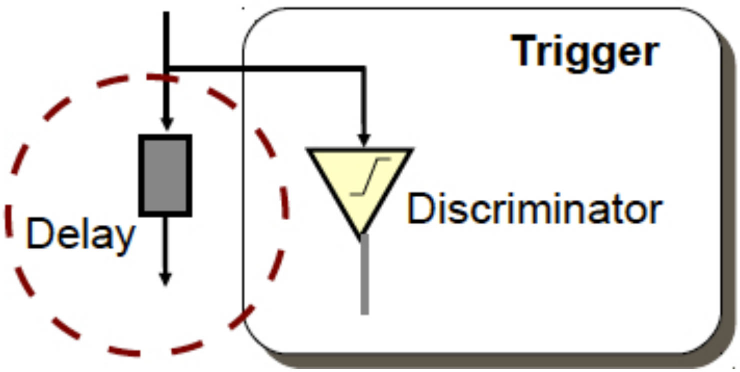

Trigger Latency

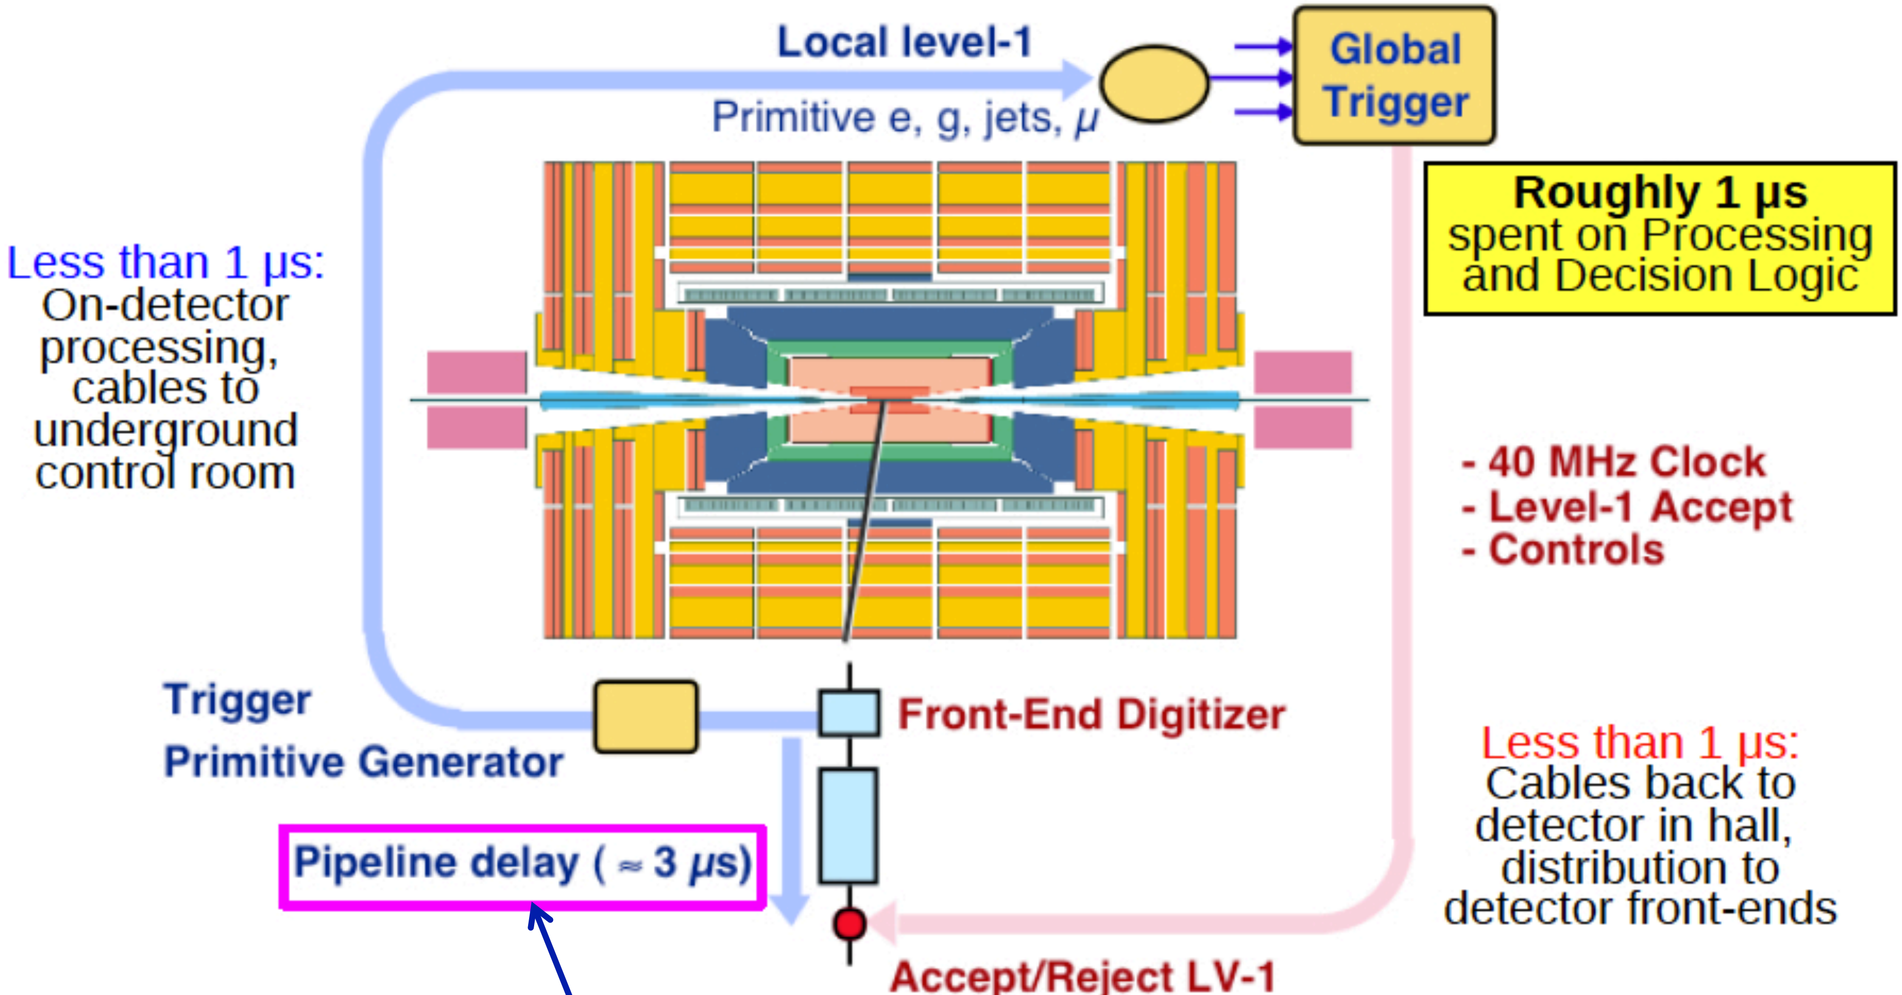

This time determines the depth of the pipeline

L1 Trigger in ATLAS

- Calorimeter and muons only
- Simple algorithms on reduced data granularity
- Selection based on particle type, multiplicities and thresholds
- Reject the bulk of uninteresting collisions

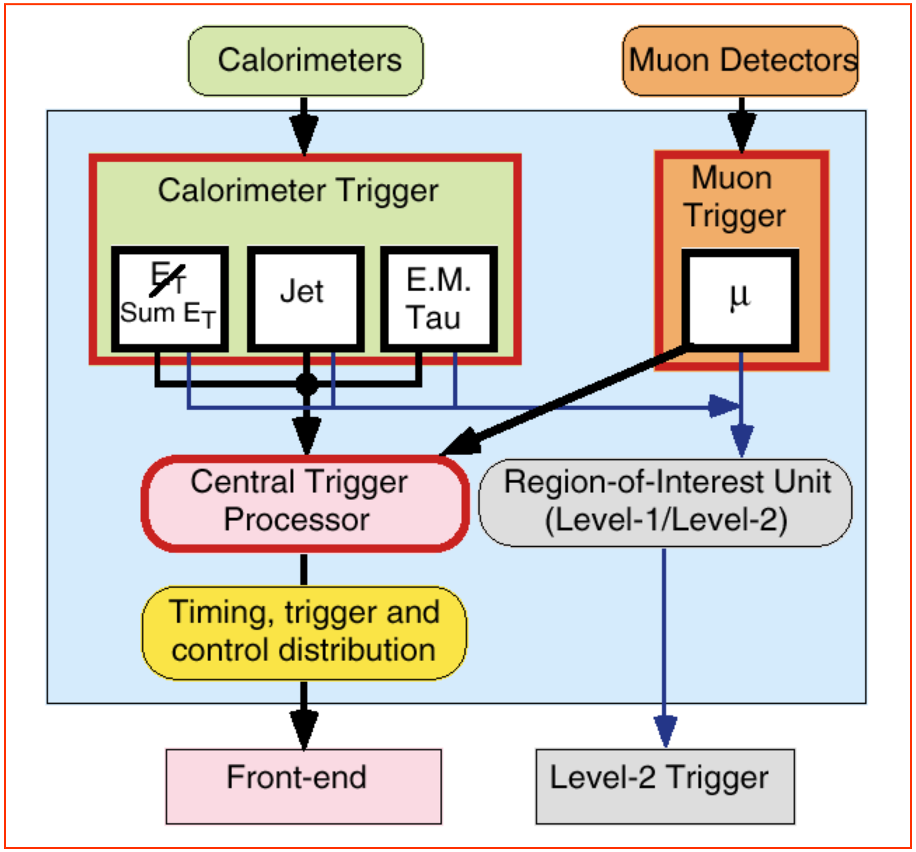

ATLAS L1 calorimeter trigger

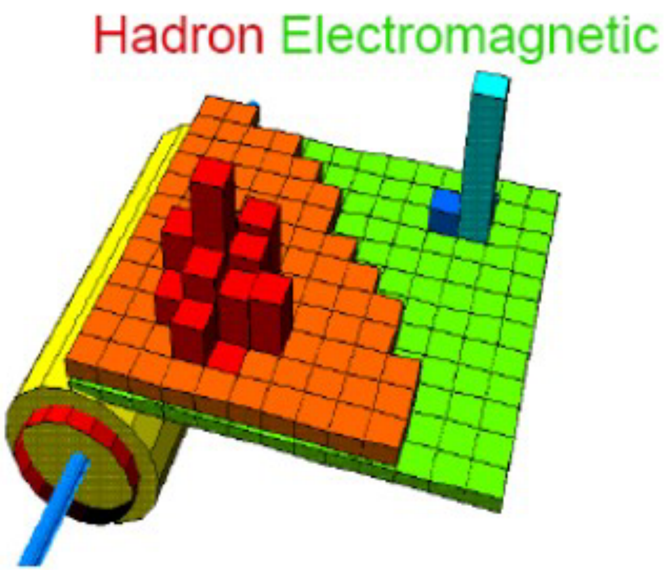

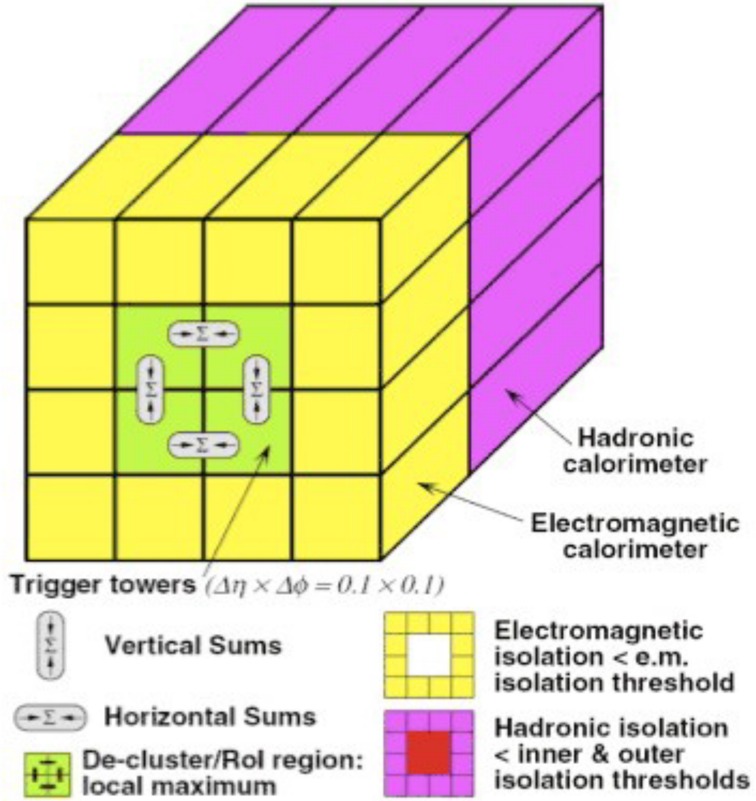

- Example: ATLAS e/γ trigger
  - Sum energy in calorimeter cells into EM and hadronic towers
  - Loop over grid and search in 4x4 towers for a local maximum 1x2 (2x1): cluster
- Can do something similar for other particles: jets, tau or sum the energy of all towers: missing E_T

CMS L1 muon trigger

Curved p_T-dependent muon path requires fast pattern recognition

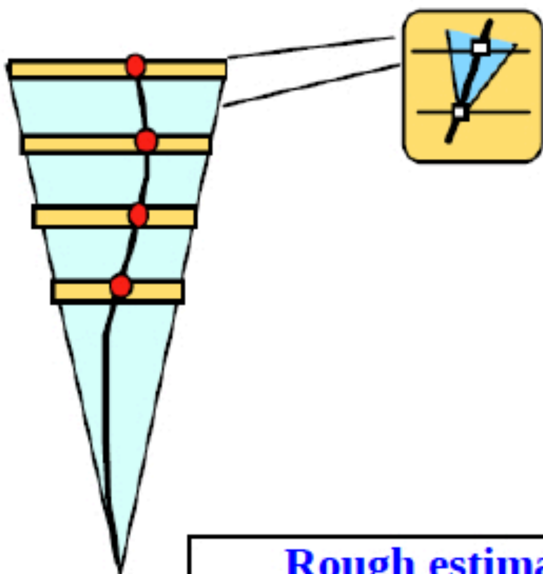

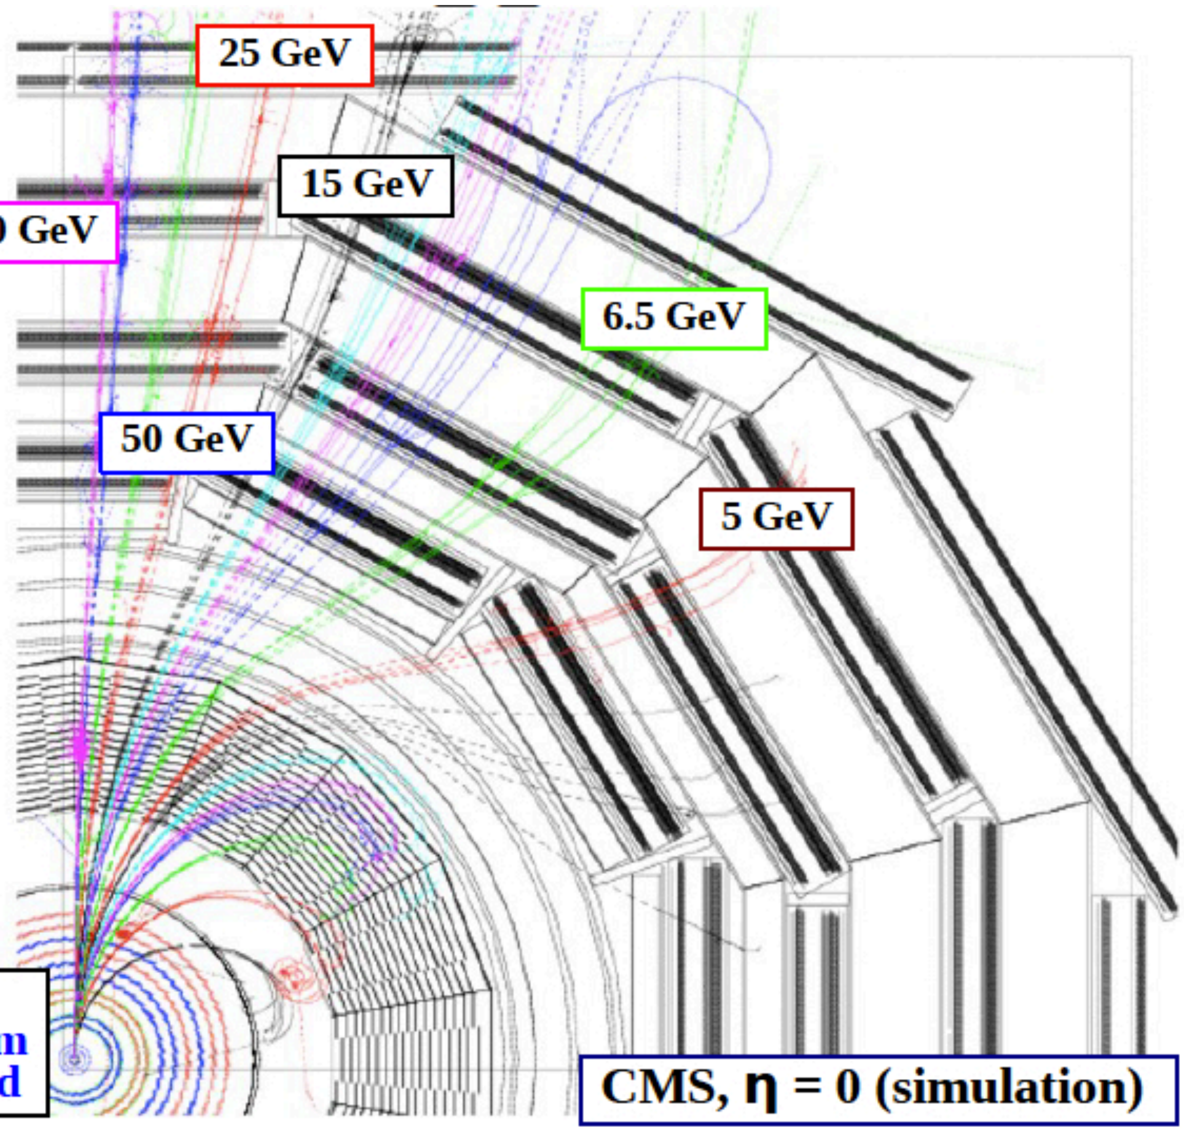

Rough estimate of muon p_T determined from bending in magnetic field

Central/Global Trigger

- Now we have the information on the particle candidates found by L1 in the detector
  - We know type, location and E_T/p_T threshold passed
  - Can also look at topological information
    - E.g. lepton opposite ETmiss, invariant mass of 2 leptons…
- Need to decide if this event is of any interest to us
  - This needs to be made quickly

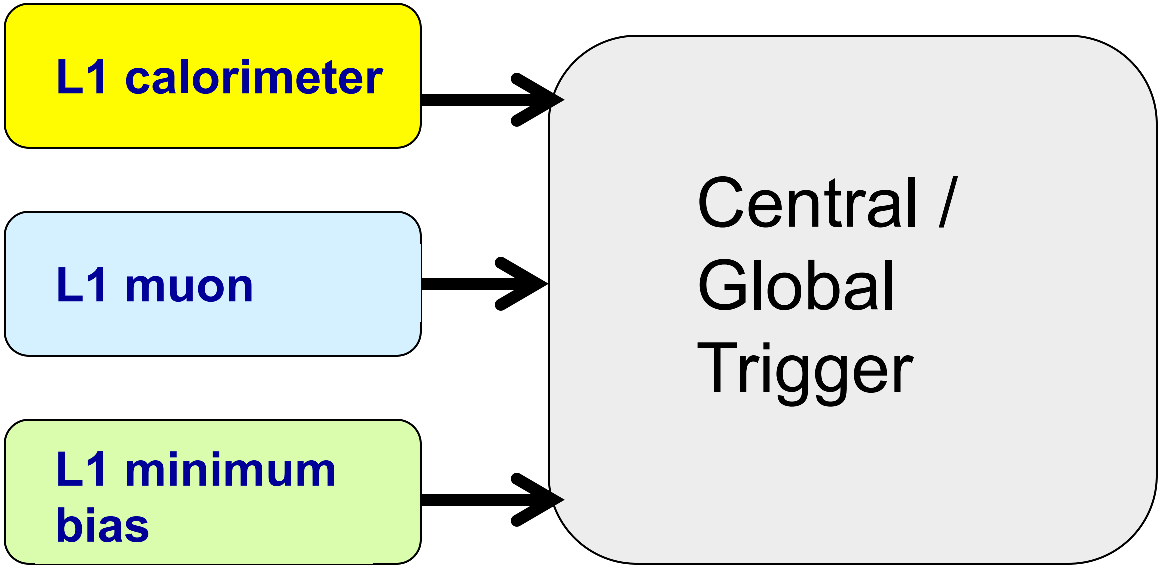

Software Trigger: Higher Level Trigger (HLT)

- L1 selected a large rate (up to 100 kHz) of events that "might be interesting"
  - These events are not kept yet (rate too high for storage), but sent to the HLT for additional filtering
- Use network-based High Level Trigger computer farm(s)
  - commercially available HW organized in a farm

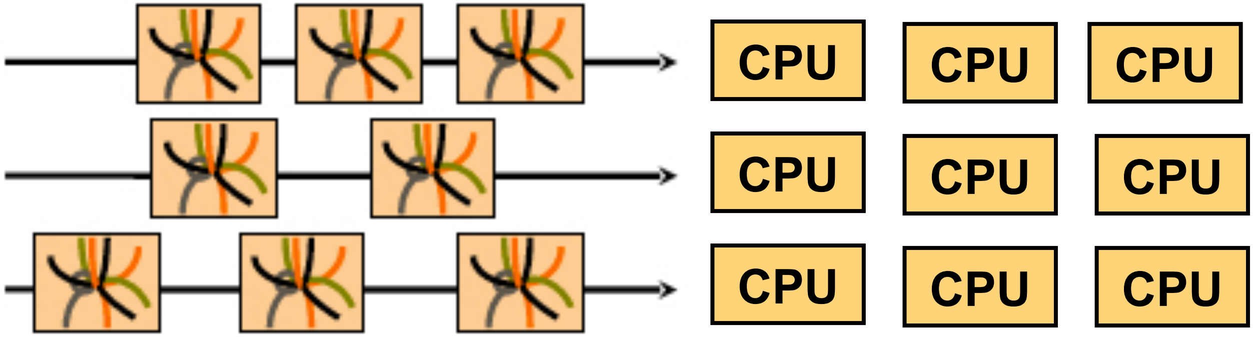

HLT Example: Muon

- Muons in CMS:
  - Reconstruct and fit tracks using only muon system
  - Continue if sufficient p_T
  - Combine tracker hits with muon system to improve p_T measurement
  - Keep the event if p_T is large enough
- Muons in ATLAS:
  - At Level 2, using detector information from the region around the L1 muon candidate, assign muon p_T based on fast look up tables
  - Extrapolate to the collision point and find the associated track
  - Is the muon isolated in the tracker, calorimeters?
  - Refine selection at L3 using offline-based reconstruction, recompute p_T
- More on HLT in next lecture

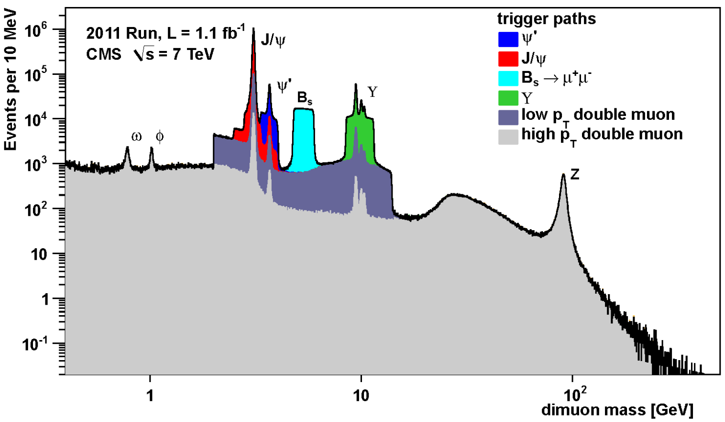

Higher Level Trigger

- Massive commercial computer farm
- Each CPU can process individual event or run multi-threaded
- Resources are still limited
  - Offline: Full reconstruction takes seconds (minutes)
  - Online latency: ms - s (input rate dependent)
- Need to reduce rate to O(1 kHz)
  - Note, output rate mainly driven by offline resources (CPU / disk space)

The ATLAS Trigger/DAQ System

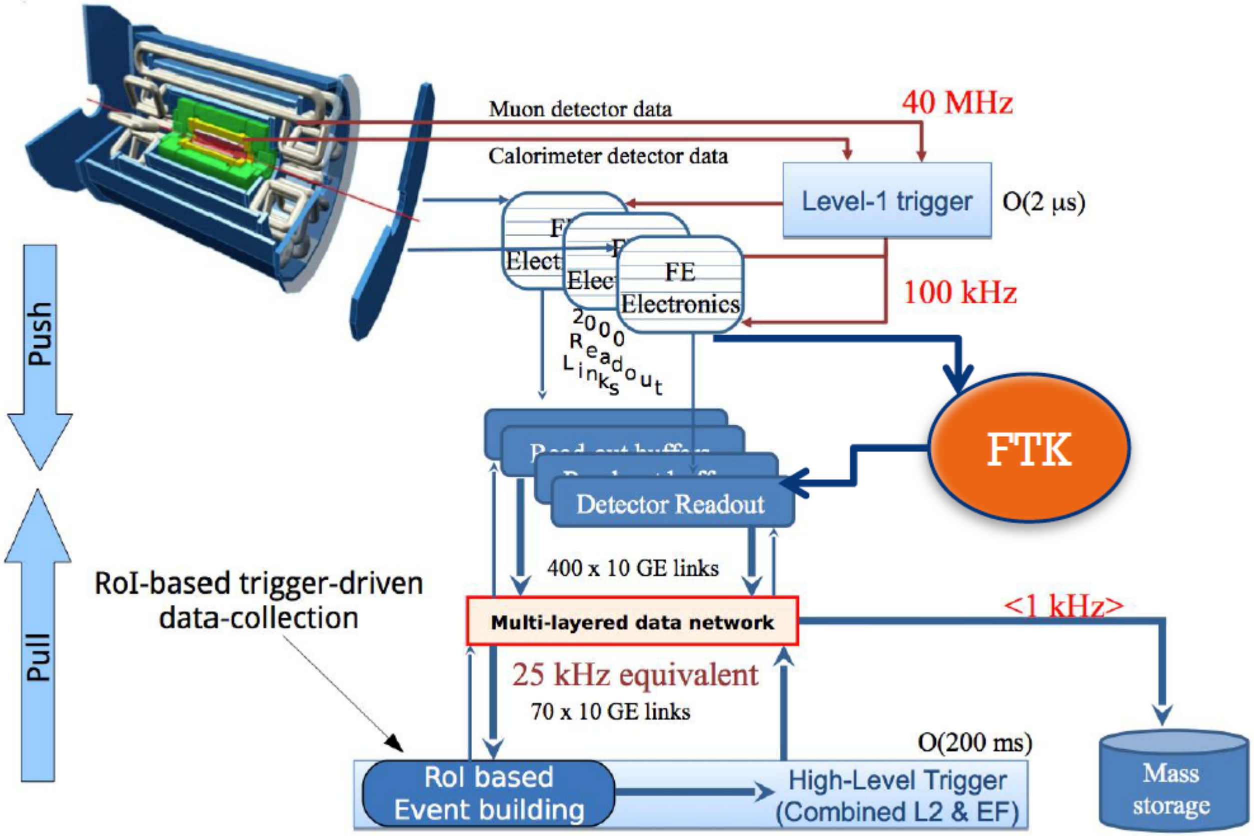

DAQ/HLT

- Overall Trigger & DAQ architecture: 3 trigger levels
- Level-1:
  - 2.5 μs latency
  - 100 kHz

- HLT: run L2 and EF in one farm
- Average output rate: ~1 kHz (physics), ~2 kHz (calib/monitoring)
- Processing time: 0.2s on average
- Average event size 1.5 - 2 MB

The ATLAS Special Features

- On-demand event building seeded by Region of Interests
  - No need to analyse the whole event in HLT, just look at regions flagged at L1 (e.g. regions with e/γ, μ, τ, jet candidates
  - On average look only at ~5% of the data

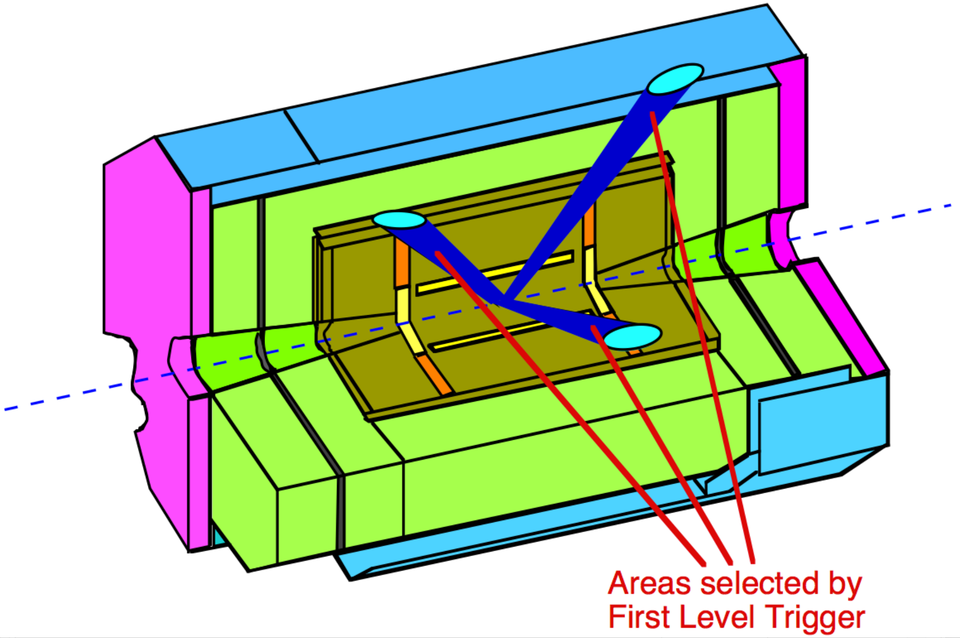

- L2 and EF run on same CPU within one farm (new in 2015)
  - Provides efficient coupling between subsequent selection steps, reducing duplication of CPU usage and network transfer
  - Allows flexible combination of fast and detailed processing

The CMS Trigger/DAQ System

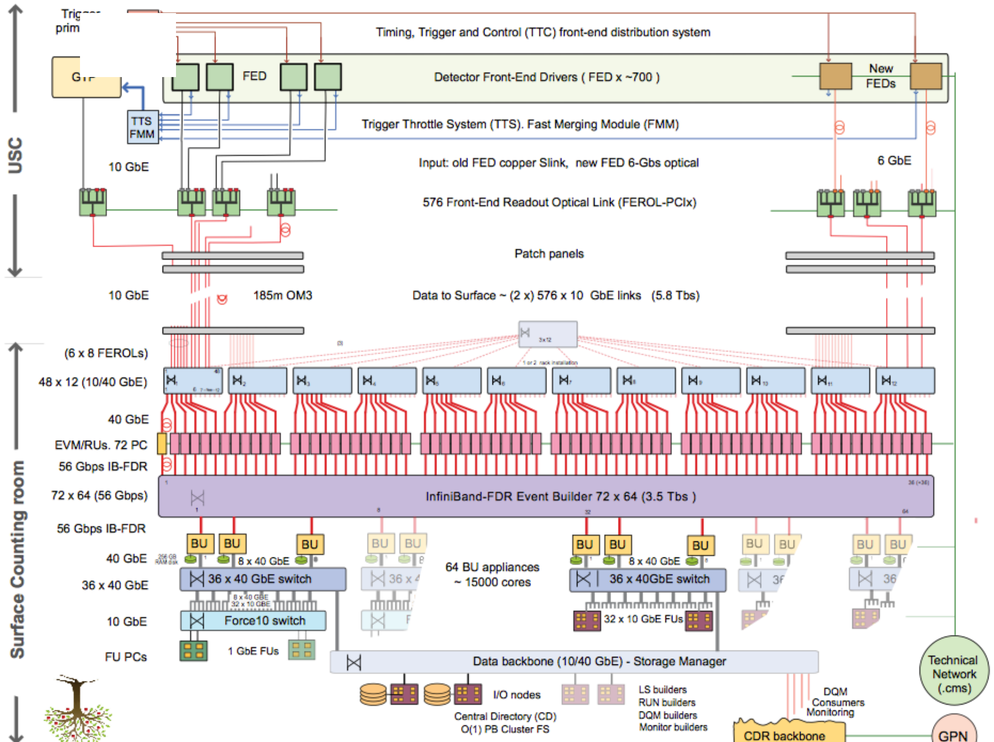

- Overall Trigger & DAQ architecture: 2 trigger levels
- DAQ & HLT decoupled via inter-mediate shared temp. storage
- Level-1:
  - 3.2 µs latency
  - 100 kHz output
- DAQ/HLT
  - Event building at full L1 rate
  - Average output rate: ~1 kHz
  - Average event size 1.5 Mb
  - Max. average CPU time: ~160 ms/event

The CMS Special Features

- 2 stage event building!
- 1st stage:
  - Combine fragments into super-fragment in RU (Readout Unit) builder
  - Event building in builder units which then write events to transient files on RAM disk
- 2nd stage:
  - Serve complete events to trigger farm.
- DAQ and HLT decoupled via intermediate shared temporary storage (new in 2015)

Detector front-end

Front-End Readout Optical Link

Data Concentrator switches

Readout Units

Event Builder switch

Builder Units

Filter Units (HLT)

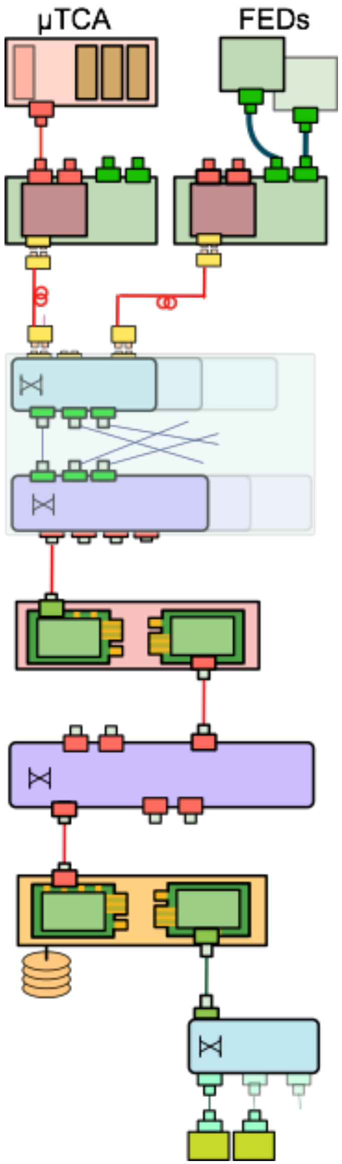

The LHCb Trigger/DAQ System

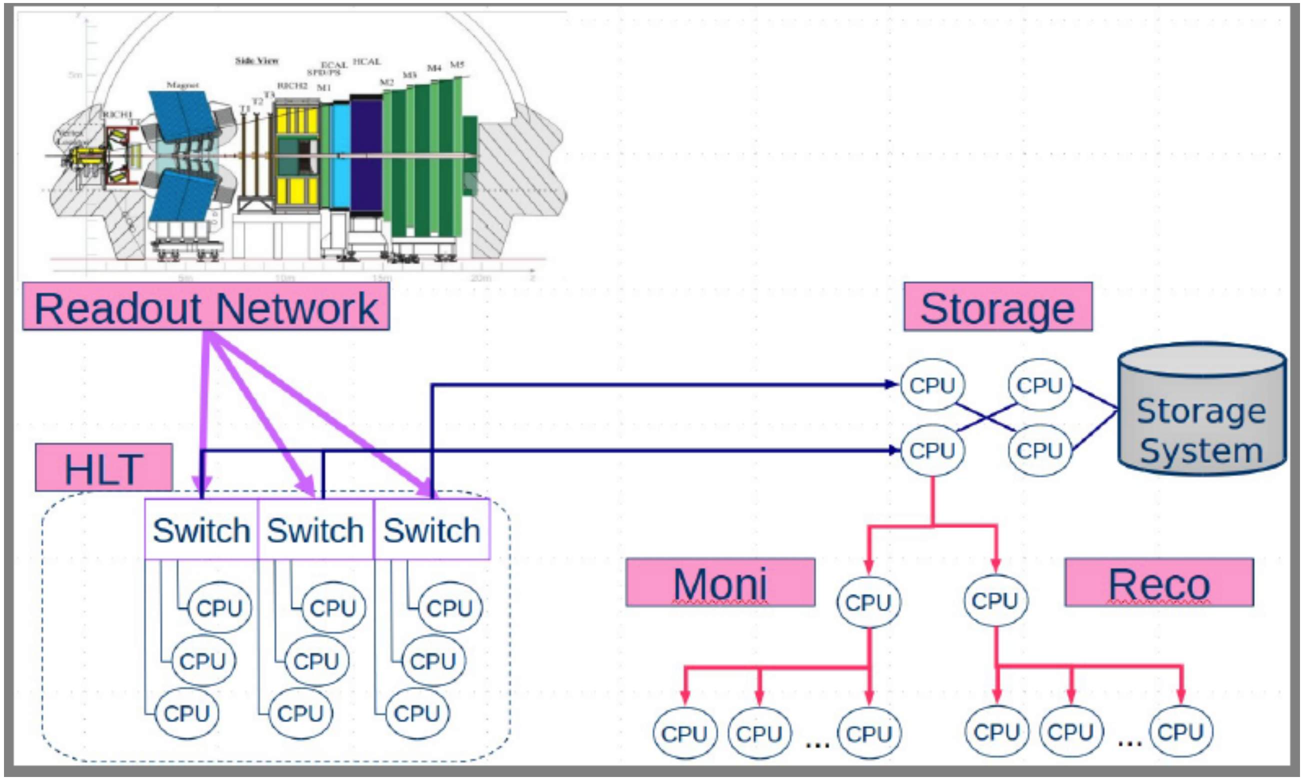

- Overall Trigger & DAQ architecture: 3 trigger levels
- Level-0:
  - 4 μs latency
  - 1 MHz output
- DAQ/HLT
  - L1: spot displaced high p_T tracks, output 100-200 kHz
  - L2: full event reconstruction
  - ~34 (650) ms @ L1 (L2)
  - Average output rate: 12.5 kHz,
  - Average event size 50 kB

The LHCb Special Features

- HLT decoupled from data flow via local temporary storage!
  - Using periods without beam boost CPU usage by 200 %
- Full offline-quality reconstruction available online
  - Alignments done at beg of fill, calib done per run
- Turbo Stream + Tesla Application:
  - Store full information of trigger candidates, remove most of detector raw data
  - Save more than 90% space
  - Ideal for very high signal yield [millions]
  - Very quick turn around [24 h]

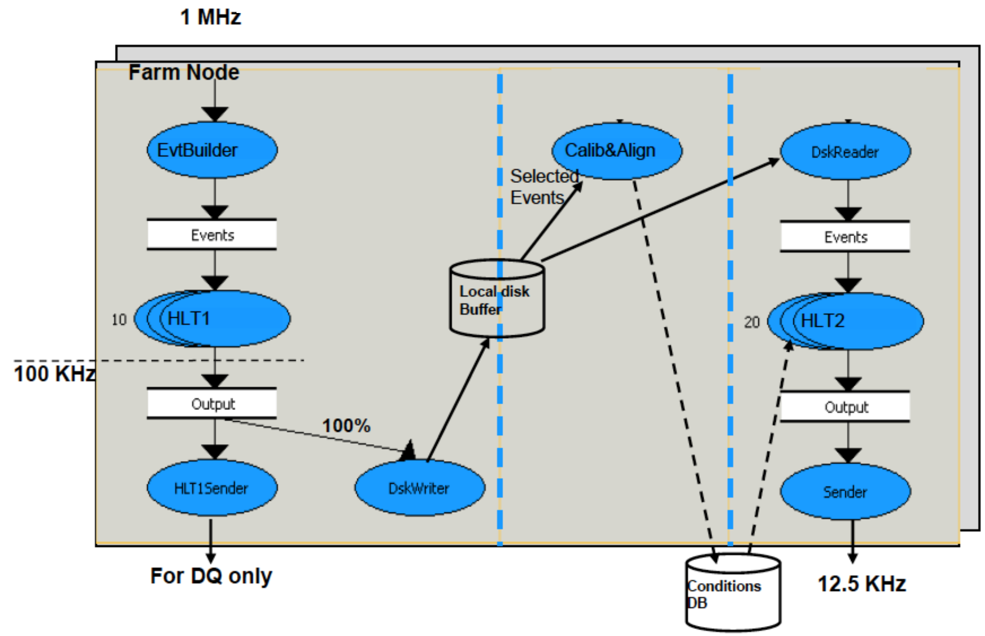

The ALICE Trigger/DAQ System

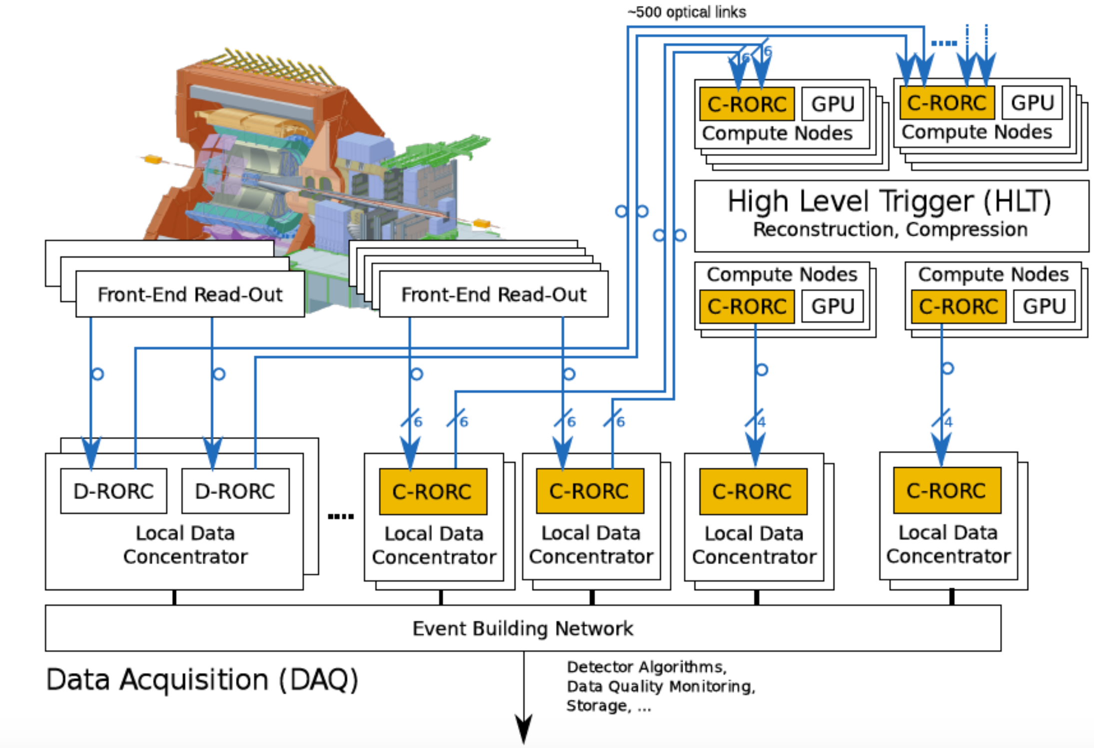

- Alice has different constraints
  - Low rate: max 8 kHz pb+pb
  - Very large events: > 40MB
  - Slow detector (TPC ~ 100 μs)
- Overall Trigger & DAQ architecture: 4 trigger levels
- 3 hardware-based trigger, 1 software-based:
  - L0 – L2: 1.2, 6.5, 100 μs latency
  - L3: further rejection and data compression

The Alice Special Features

- Deal with huge events
  - 3 hardware level triggers
  - Heavy utilisation of hardware acceleration: FPGA + GPU
  - Use of data compression in trigger

Towards the Future

- Experiments upgrade every time the conditions provided by the accelerator change
  - Preparations start well in advance
  - The 4 LHC TDAQ systems are already planning major upgrades
    - ALICE & LCHb will upgrade for Run 3
    - CMS and ATLAS will mainly upgrade for Run 4
- Guiding Principles
  - Physics goals
  - Accelerator conditions
  - Technology reach
  - Cost
- Rapidly evolving area

Towards the Future

- Alice
  - Support for continuous read-out (TPC), as well as triggered read-out
    - Read out the data of all interactions at a maximum rate of 50kHz (upon min bias trigger)
  - One common online – offline computing system: O2
- LHCb
  - (Triggerless) Read-out @ 40 MHz + full software trigger
  - Data centre at the surface
- CMS
  - Hardware-based track trigger
- ATLAS
  - Hardware based track trigger after very first trigger level

Summary

- Challenge to design efficient trigger/DAQ for LHC
  - Very large collision rates (up to 40 MHz)
  - Very large data volumes (tens of MB per collision)
  - Very large rejection factors needed (>10^5)
- Showed data acquisition used in LHC experiments
- Introduction to basic functionality of trigger
- We'll look in detail at the trigger aspects in the next lecture
  - That one will be less technical and more physics-oriented!

Backup

Trigger/DAQ parameters

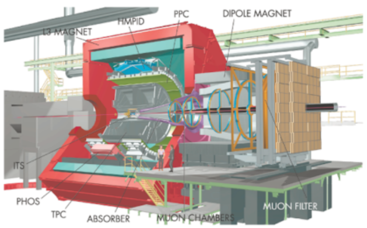

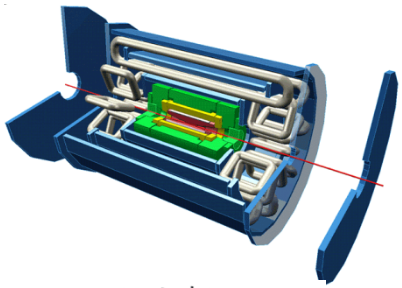

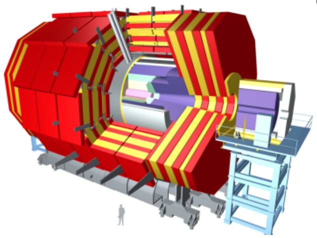

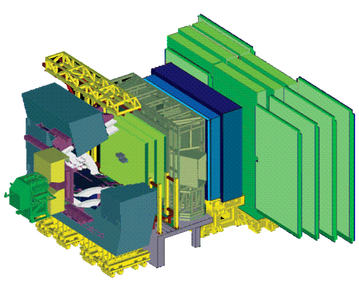

No.Levels Trigger	**Level-0,1,2** Rate (Hz)	**Event** Size (Byte)	**Readout** Bandw.(GB/s)	**HLT Out** MB/s (Event/s)
4	Pb-Pb **500** p-p **10^3**	**$5\text{x}10^7$** **$2\text{x}10^6$**	**25**	**1250** (10^2) **200** (10^2)
3	LV-1 **10^5** LV-2 **$3\text{x}10^3$**	**$1.5\text{x}10^6$**	**4.5**	**300** ($2\text{x}10^2$)
2	LV-1 **10^5**	**10^6**	**100**	**~1000** (10^2)
2	LV-0 **10^6**	**$3.5\text{x}10^4$**	**35**	**70** ($2\text{x}10^3$)

TDAQ comparison

	ATLAS	CMS	LHCb	ALICE
"L1" Latency [μs]	2.5	3.2	4	1.2/6/88
Max "L1" output rate [kHz]	75	100	1000	~2
Frontend readout bandwidth [GBytes/s]	120	100	40	25
Max HLT avg. latency [ms] (upgrade with luminosity)	L2: 40 EF: 1000	50 (in 2010)	20	
Event building bandwidth [GBytes/s]	4	100	40	25
Trigger output rate [Hz]	~200	~300	~2000	~50
Output bandwidth [MBytes/s]	300	300	100	1200
Event size [MBytes]	1.5	1	0.035	Up to 20

Data handling requirements