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
- Upgrades
- 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

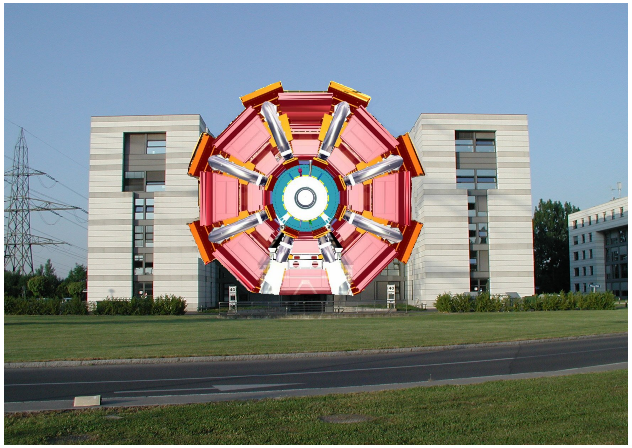

Challenge 1: Physics

- 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. $H\rightarrow\gamma\gamma$, BR $\sim 10^{-3}$)

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

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

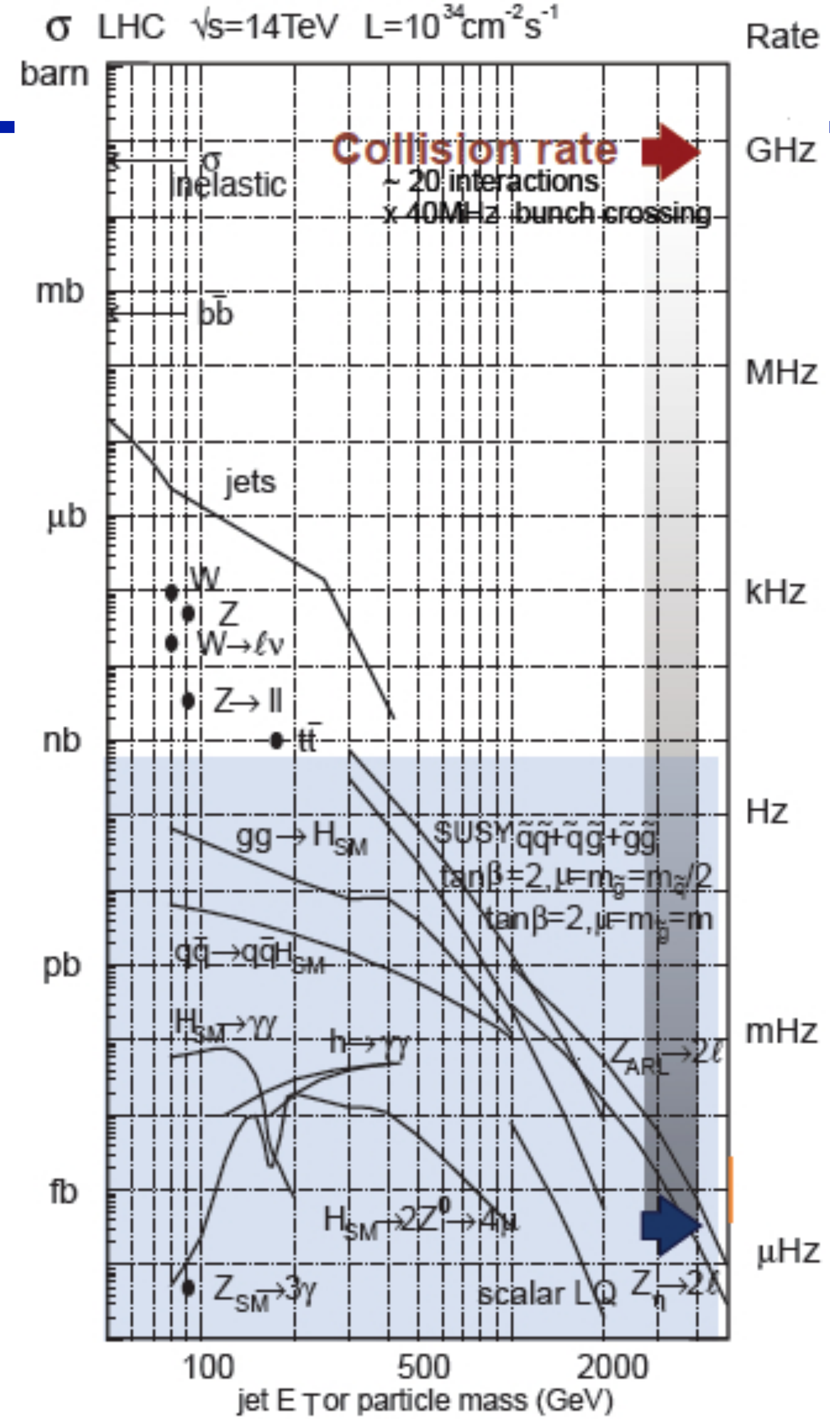

Challenge 1: Physics

- 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 ~300-1000 Hz of data written out to storage
- Not as trivial, W→lν: 150 Hz
  - “Good” physics can become your enemy!

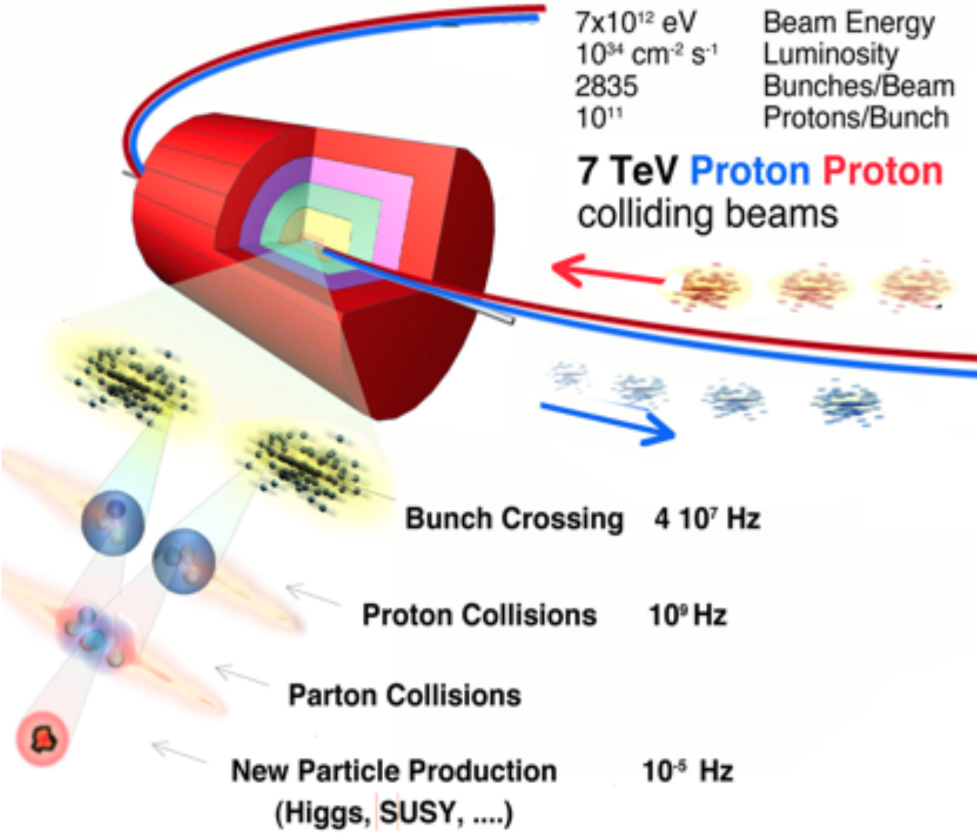

Challenge 2: Accelerator

- Unlike e^+e^- colliders, proton colliders are more ‘messy’ due to proton remnants
- In 2012 LHC already produced up to 30 overlapping p-p interactions on top of each collision (pile-up) ➔ >1000 particles seen in the detector!

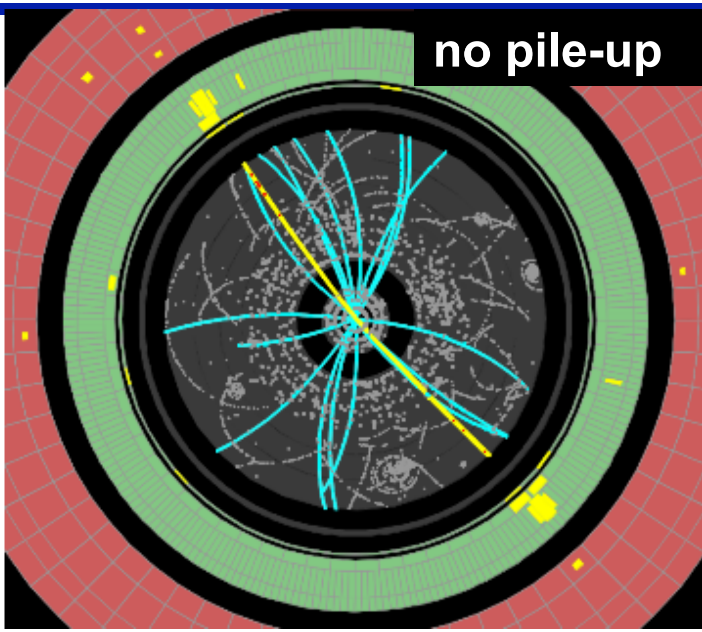

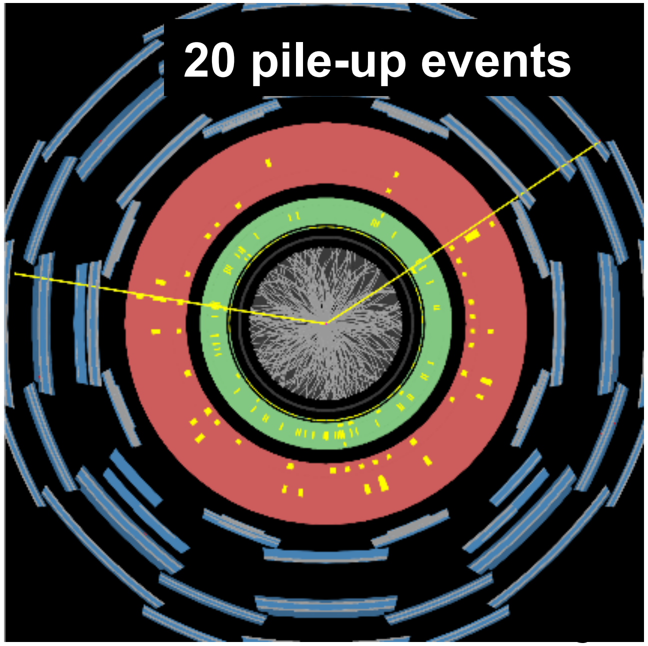

Challenge 3: Detector

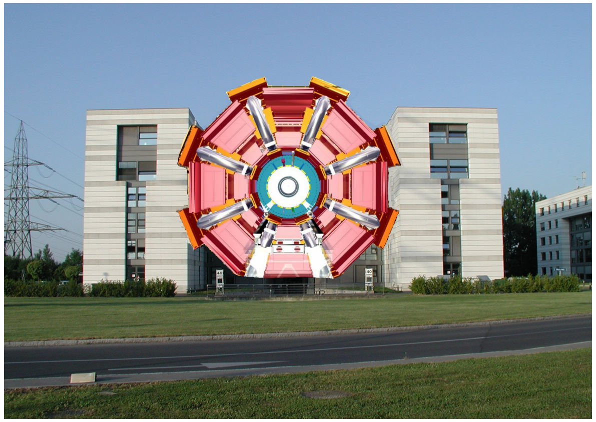

- Besides being huge: number of channels are O(10^6-10^8) at LHC, event sizes ~1.5 MB for pp collisions, 50 MB for pb-pb collisions in Alice
  - Need huge number of connections
- Some detectors need > 25ns to readout their channels and integrate more than one bunch crossing's worth of information (e.g. 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")

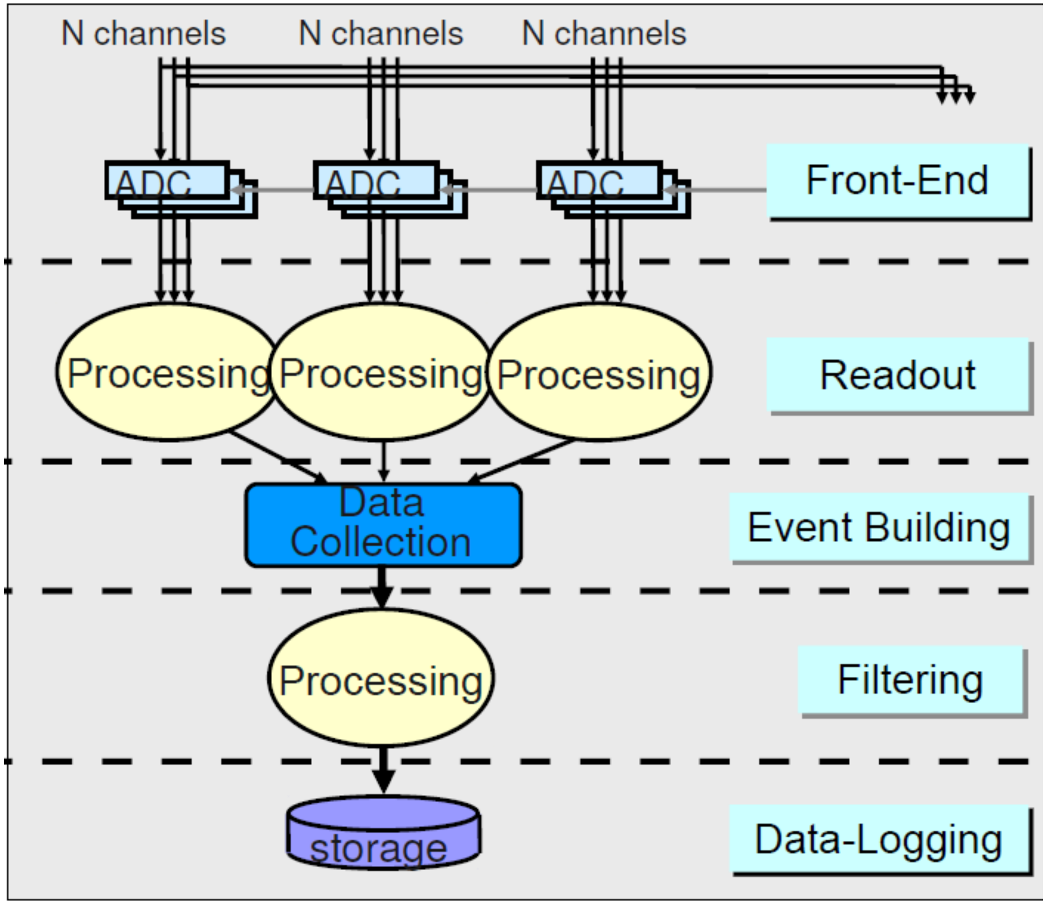

Let's build a Trigger and DAQ for this

- What 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")

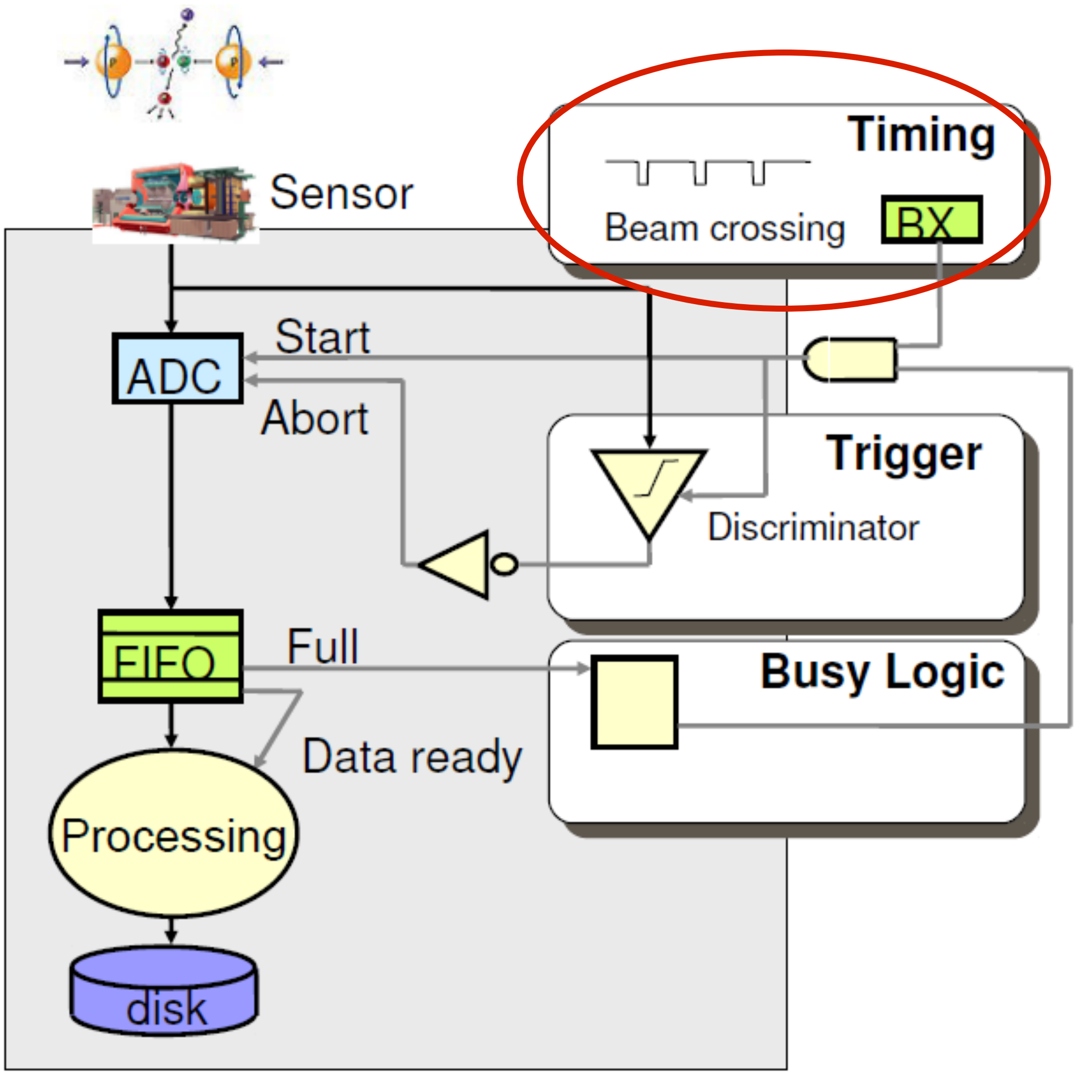

Let's build a Trigger and DAQ for this

- 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

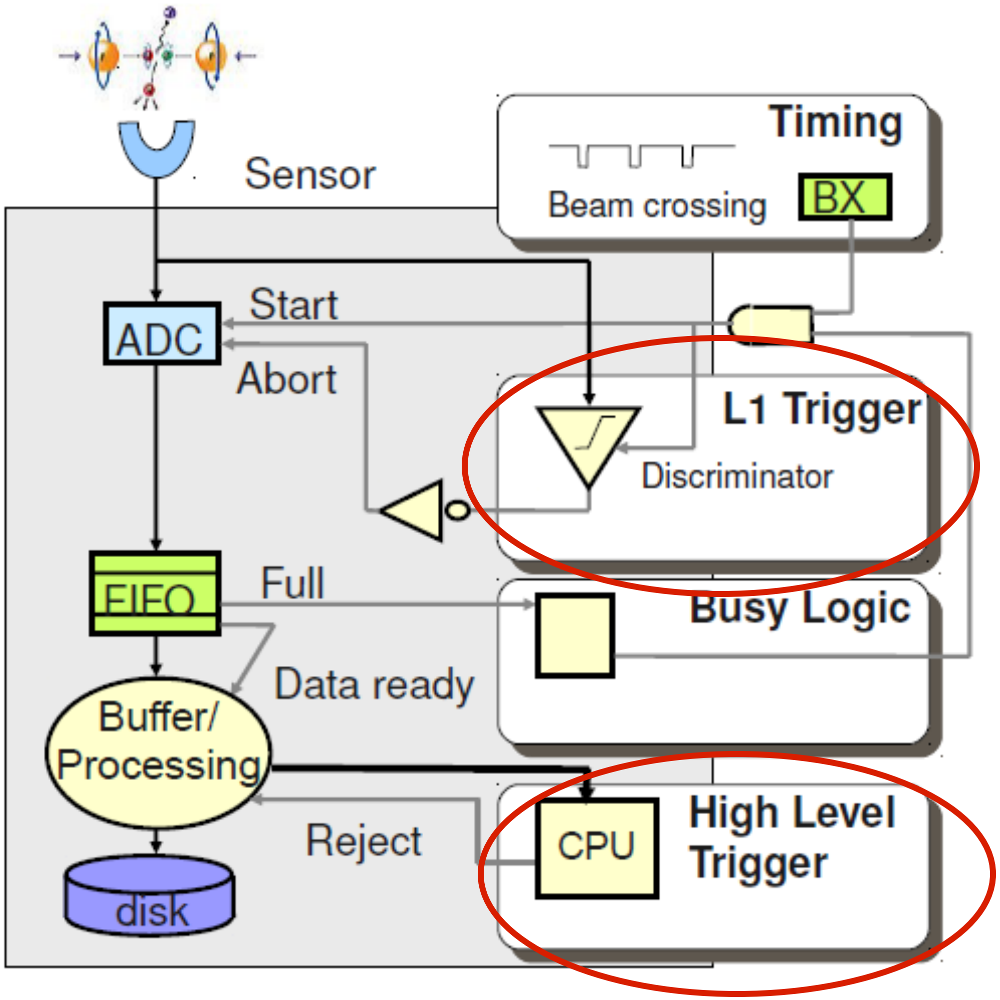

Let's build a Trigger and DAQ for this

- 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

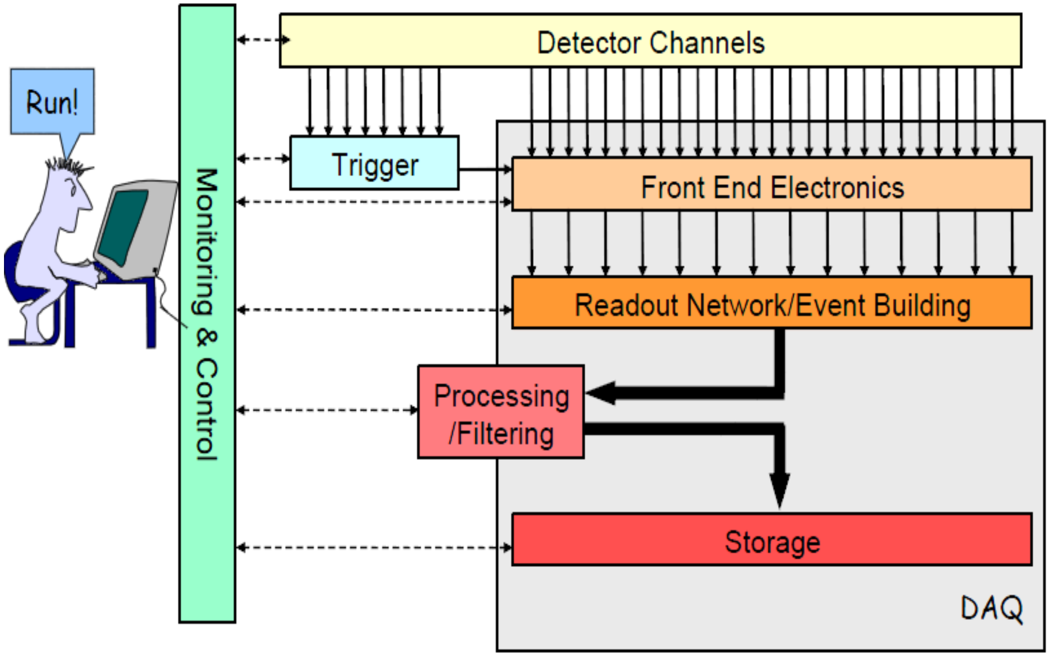

Let's look more at the trigger part…

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

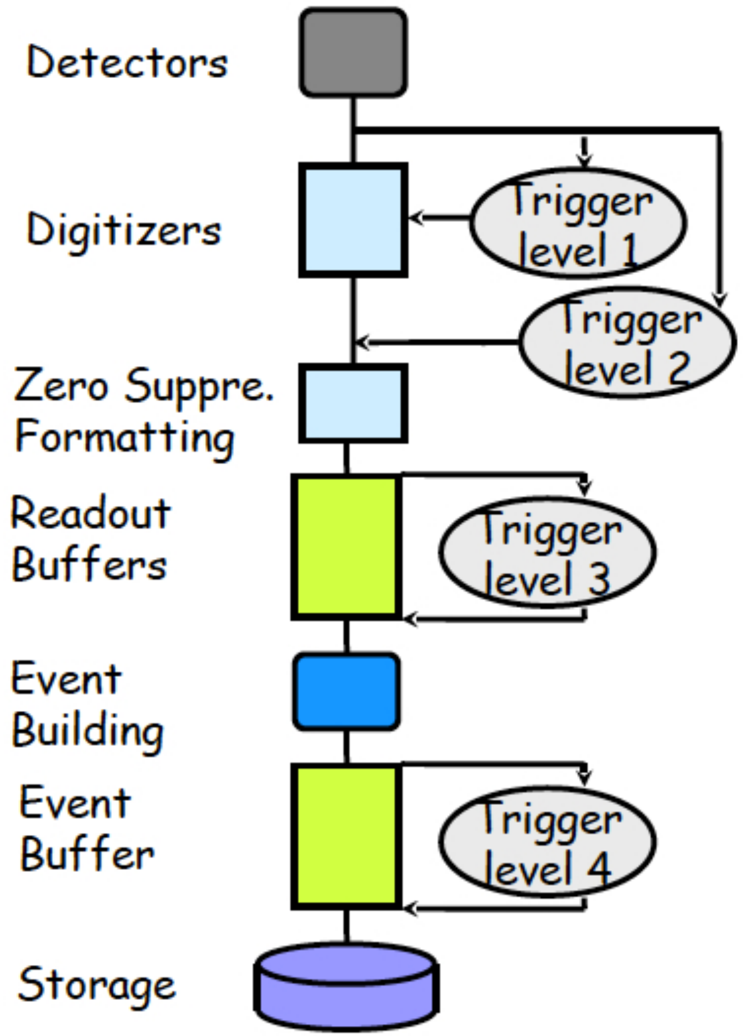

LHC DAQ phase-space

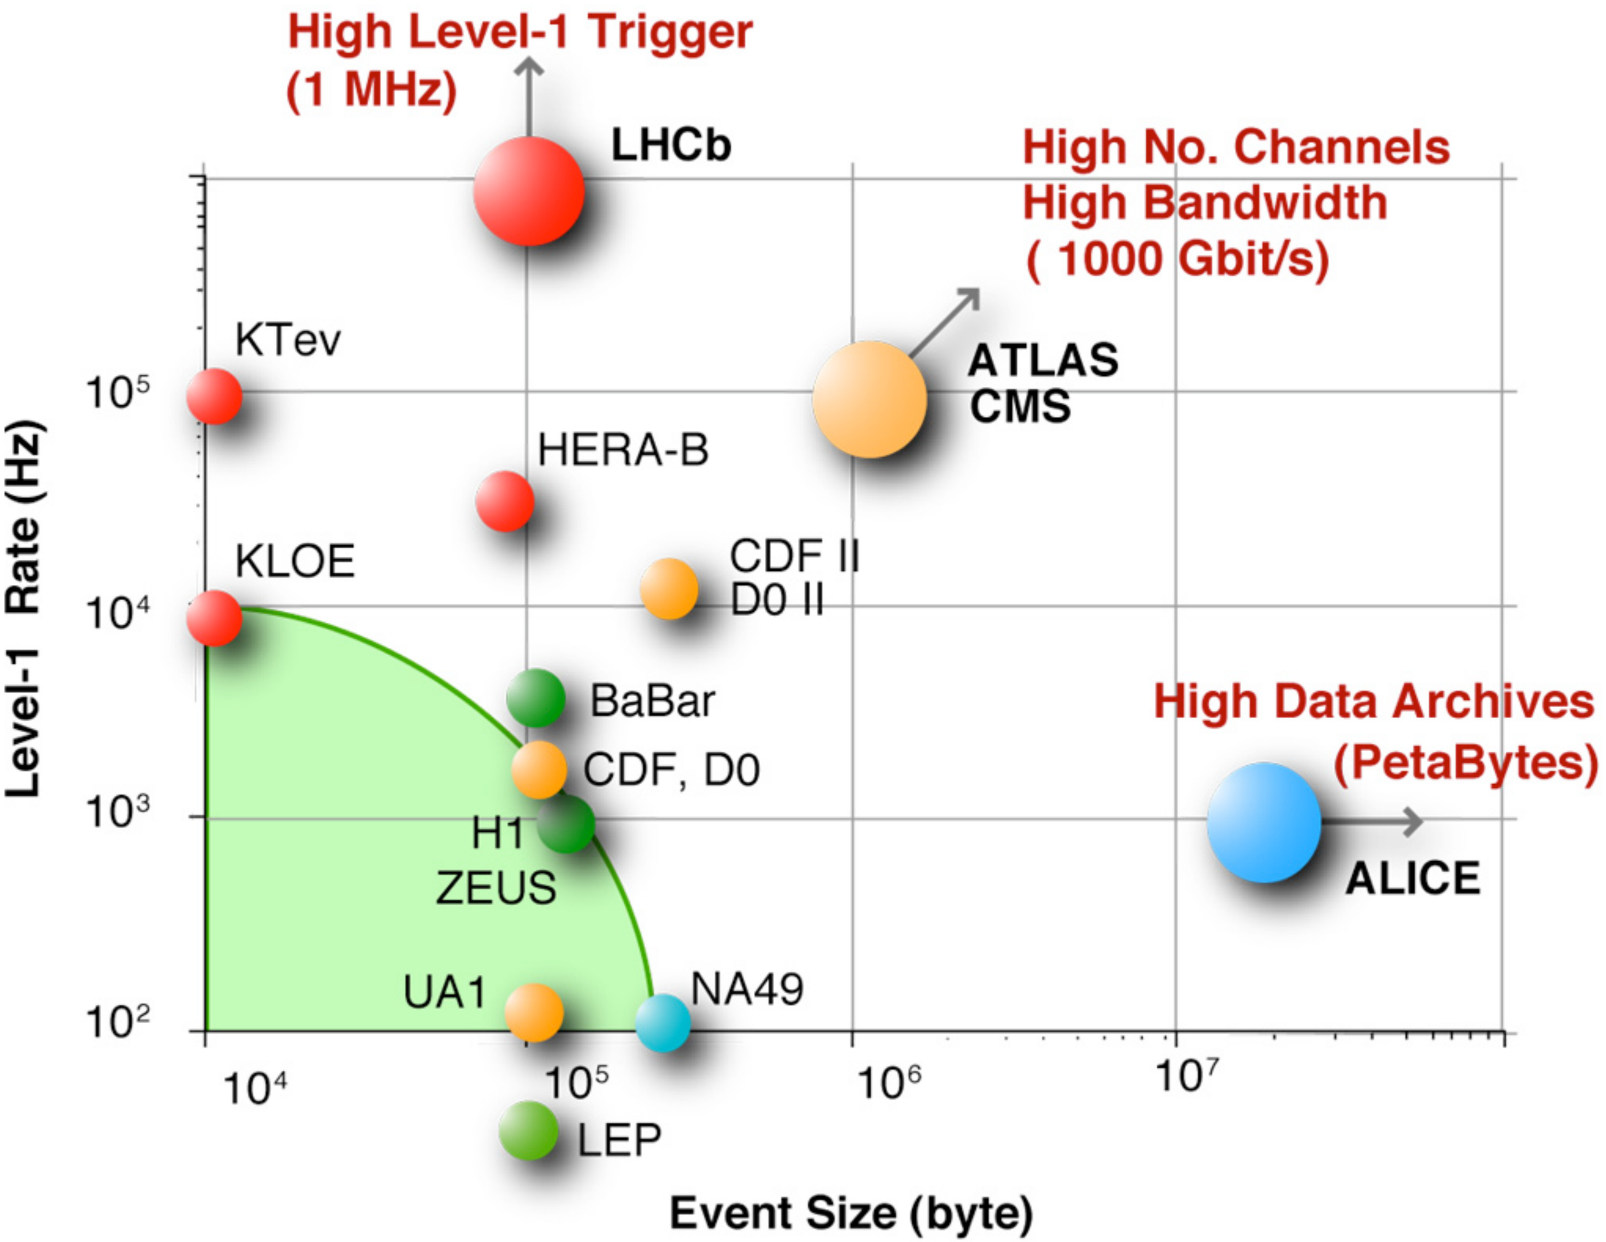

[

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

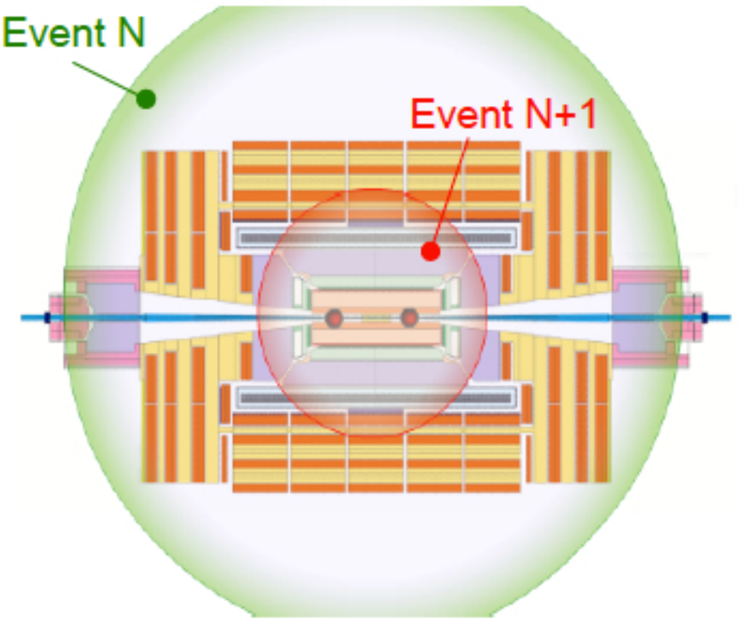

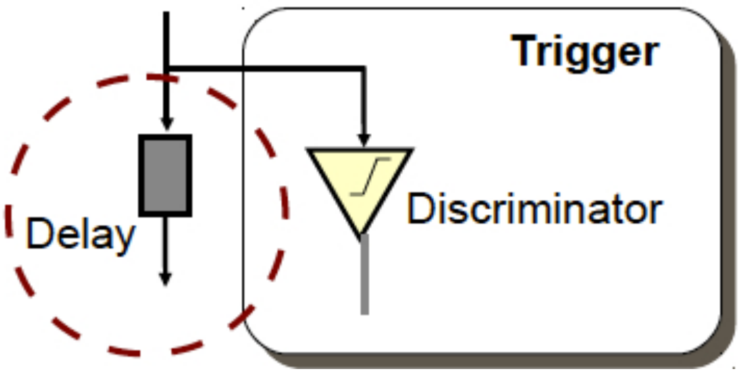

Trigger Latency

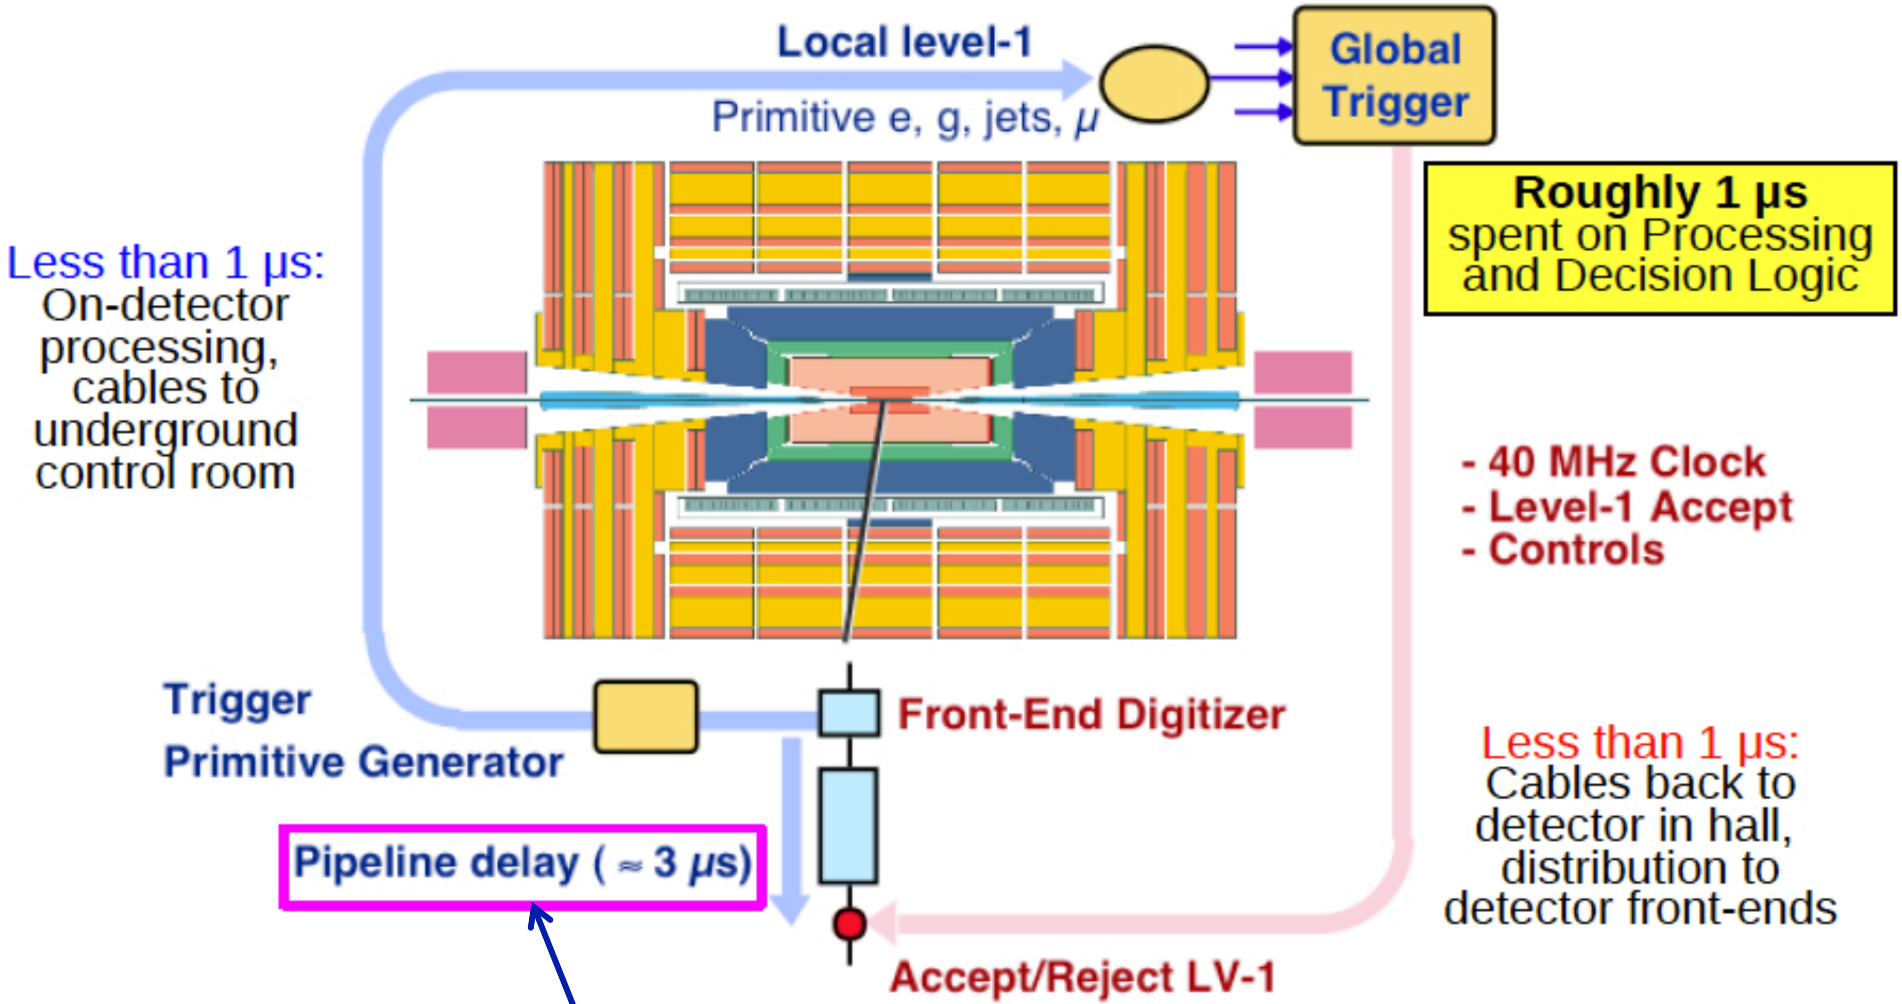

This time determines the depth of the pipeline

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

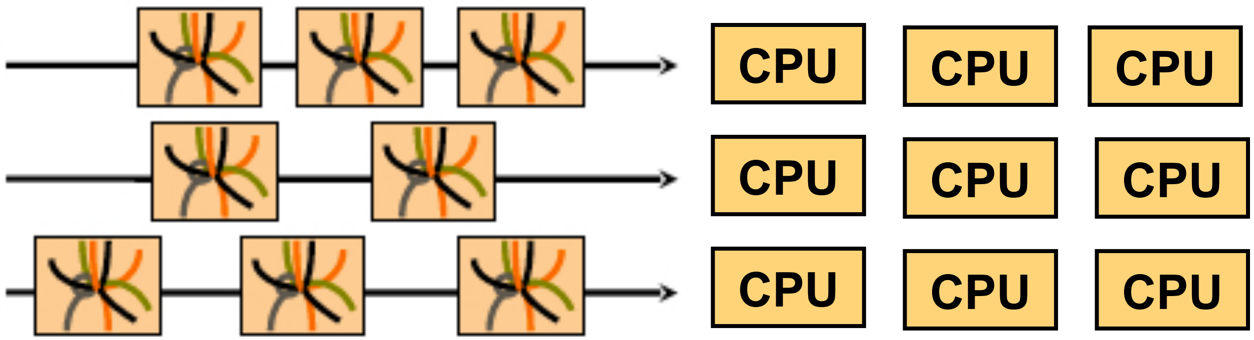

Higher Level Trigger

- Massive commercial computer farm
  - ATLAS: L2 and L3 handled by separate computing farms in 2012
    - Roughly 17k CPUs that can be freely assigned to either
  - CMS: Single computing farm (roughly 13k CPUs in 2012)
- Parallel processing, each CPU processes individual event
- Resources are still limited
  - Offline: Full reconstruction takes seconds (minutes)
  - Online latency: ms - s (input rate dependent)
- Need to reduce rate to O(few 100 Hz)
  - Note, output rate mainly driven by offline

The ATLAS Trigger/DAQ System

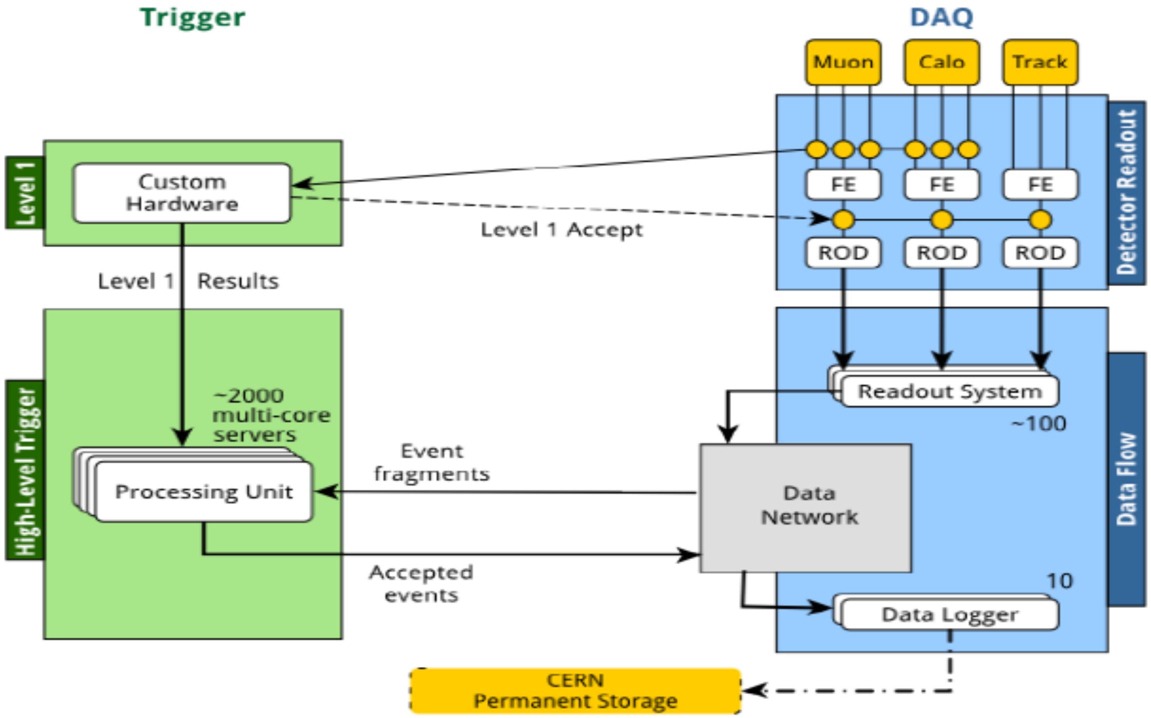

DAQ/HLT

- Overall Trigger & DAQ architecture: 3 trigger levels
- Level-1:
  - 2.5 μs latency
  - 75 kHz output in 2012, 100 kHz in 2015

- Analyse regions around particles identified at L1 or whole event
- Average output rate in 2012: 400 Hz prompt, 200 Hz “parked”, ~1kHz in 2015
- Processing time: few seconds
- Average event size 1.5 MB in 2012, ~2 MB in 2015

The CMS Trigger/DAQ System

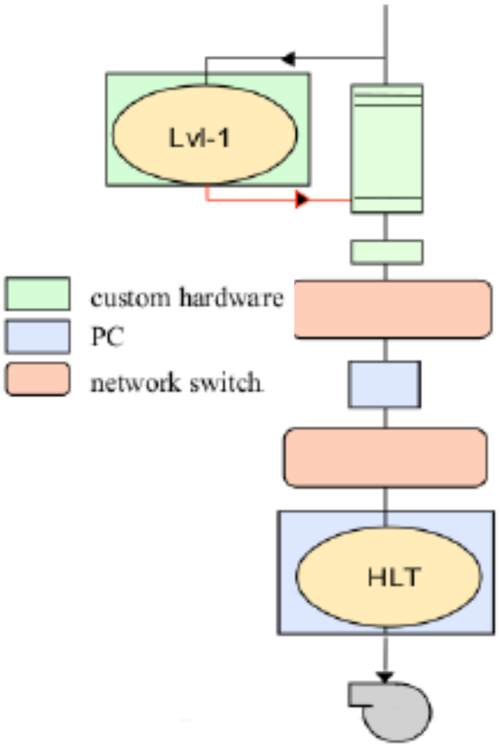

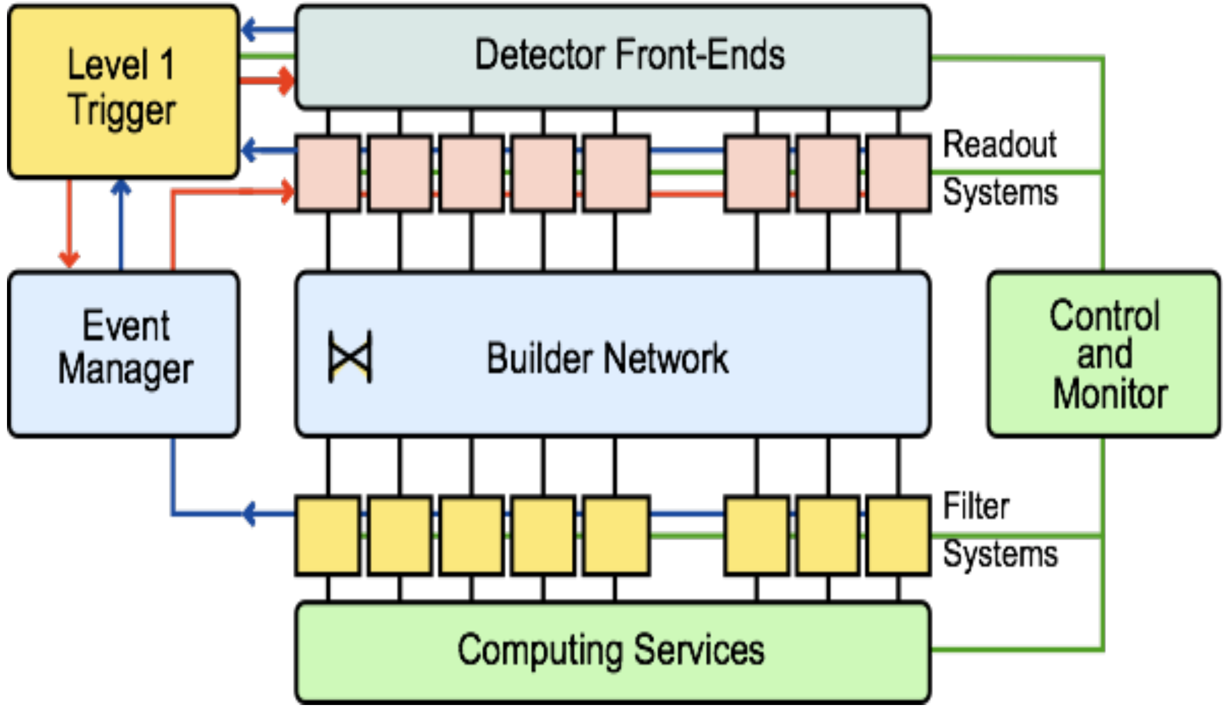

- Overall Trigger & DAQ architechture: 2 trigger levels
- Level-1:
  - 3.2 μs latency
  - 100 kHz output

- DAQ/HLT
  - Event building at full L1 rate
  - Average output rate in 2012: 350 Hz prompt, 300Hz “parked”, ~1 kHz in 2015
  - Average event size 1 MB in 2012, 2 Mb in 2015
  - Average CPU time few 100 ms

The LHCb Trigger/DAQ System

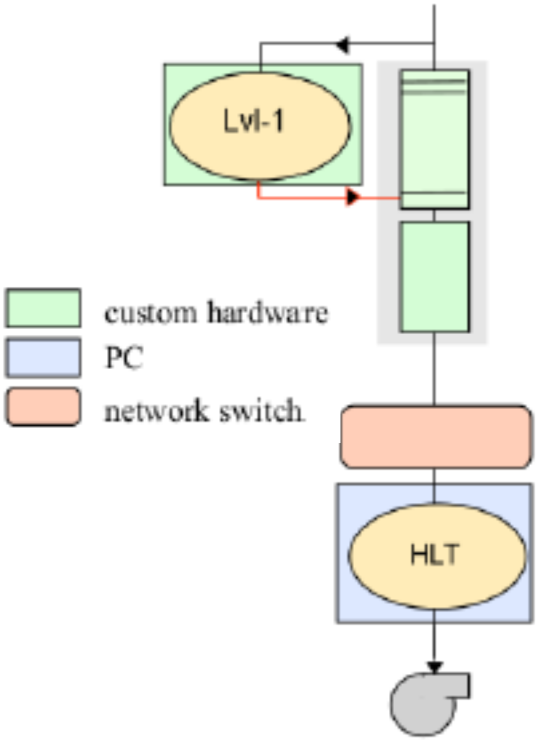

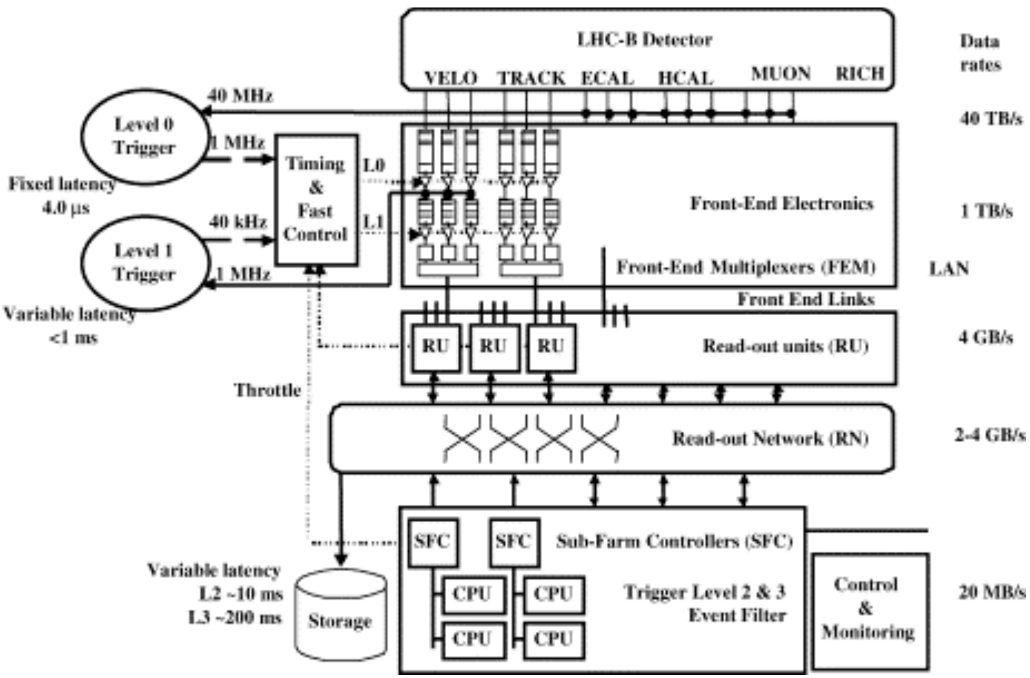

- Overall Trigger & DAQ architechture: 3 trigger levels
- Level-0:
  - 4 µs latency
  - 1 MHz output

- DAQ/HLT
  - L1: look displaced high p_T tracks, output 70 kHz
  - L2: full event reconstruction
  - Average output rate in 2012: 5 kHz, 2015: 12.5 kHz
  - Average event size 35 kB in 2012, 60 kB in 2015

The ALICE Trigger/DAQ System

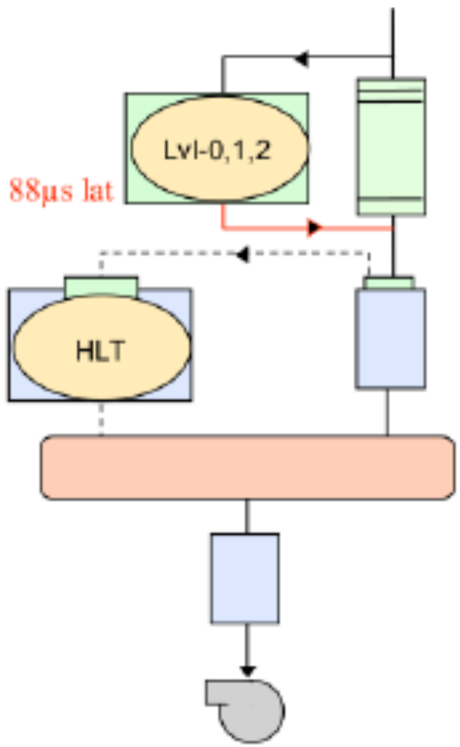

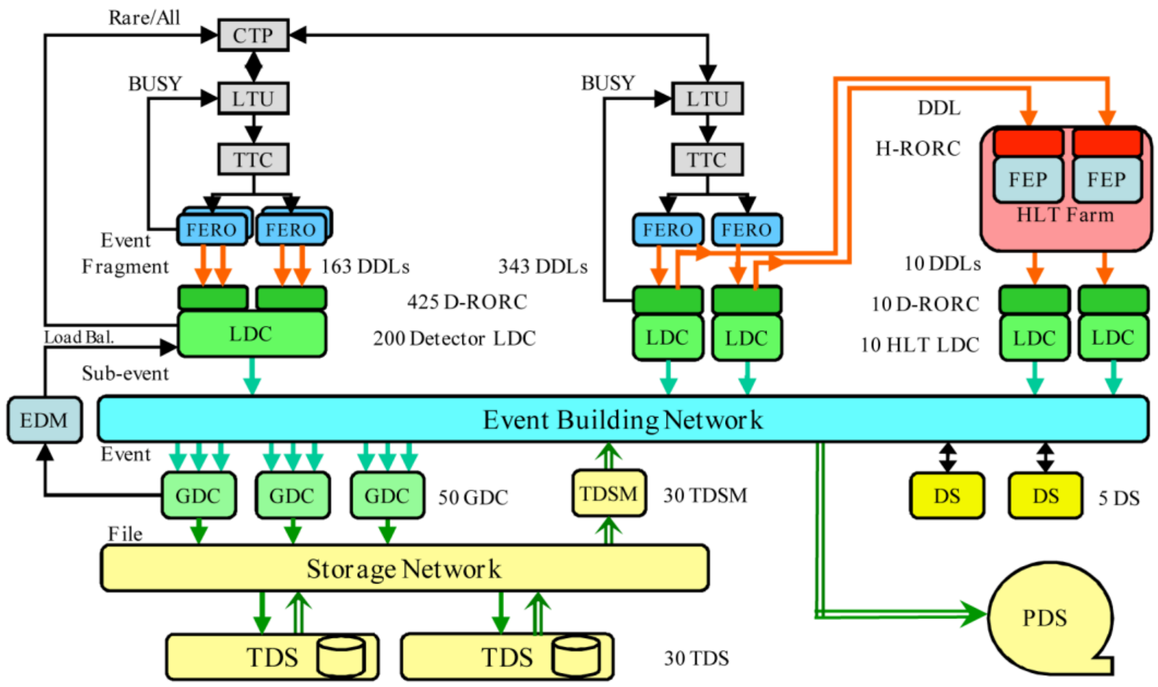

- Alice has different constraints
  - Low rate: max 8 kHz pb+pb
  - Very large events: > 40MB
  - Slow detector (TPC ~ 100 μs)
- Overall Trigger & DAQ architecture: 4 trigger levels
- 3 hardware-based trigger, 1 software-based:
  - L0 – L2: 1.2, 6.5, 88 μs latency
  - L3: further rejection and data compression

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

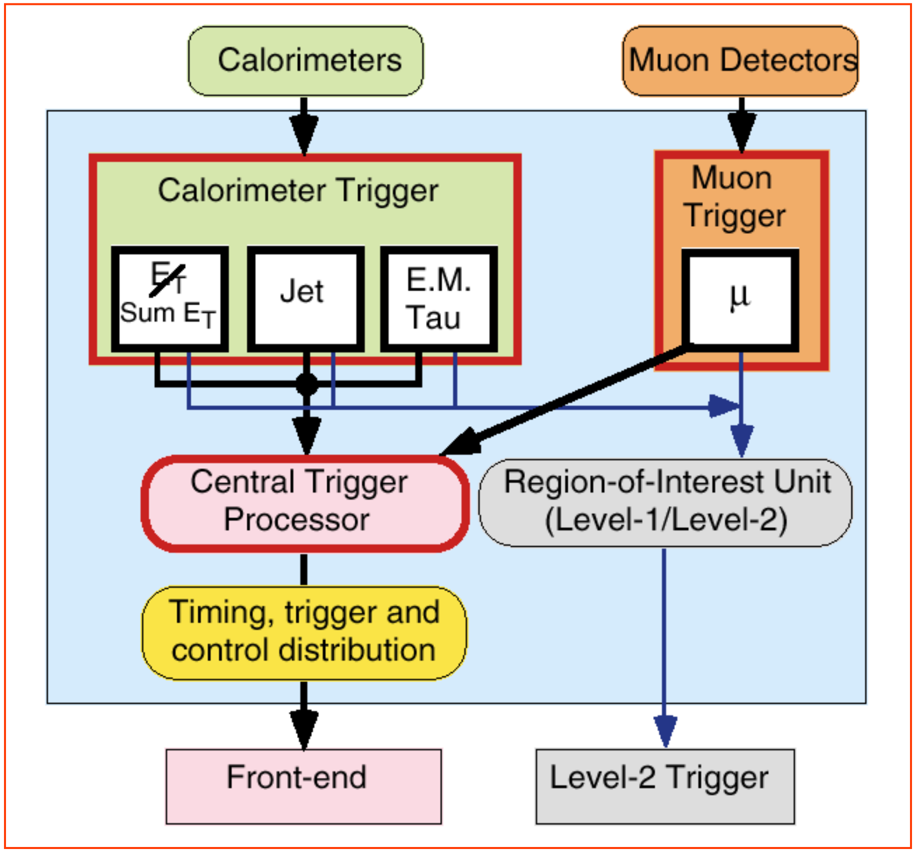

ATLAS L1 calorimeter trigger

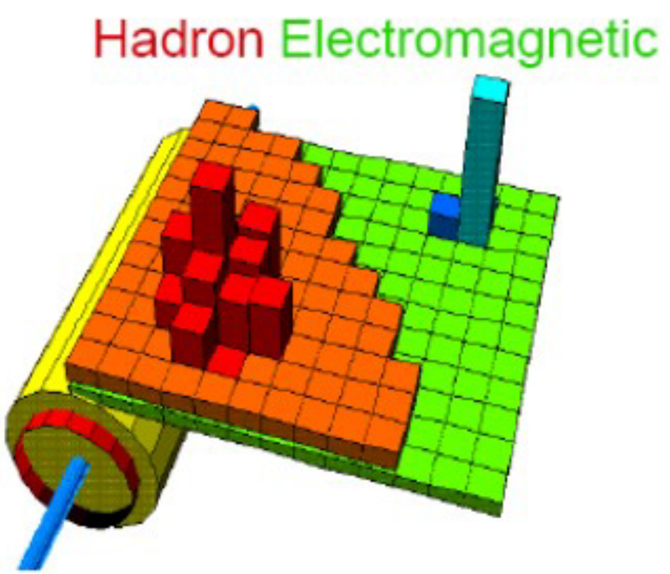

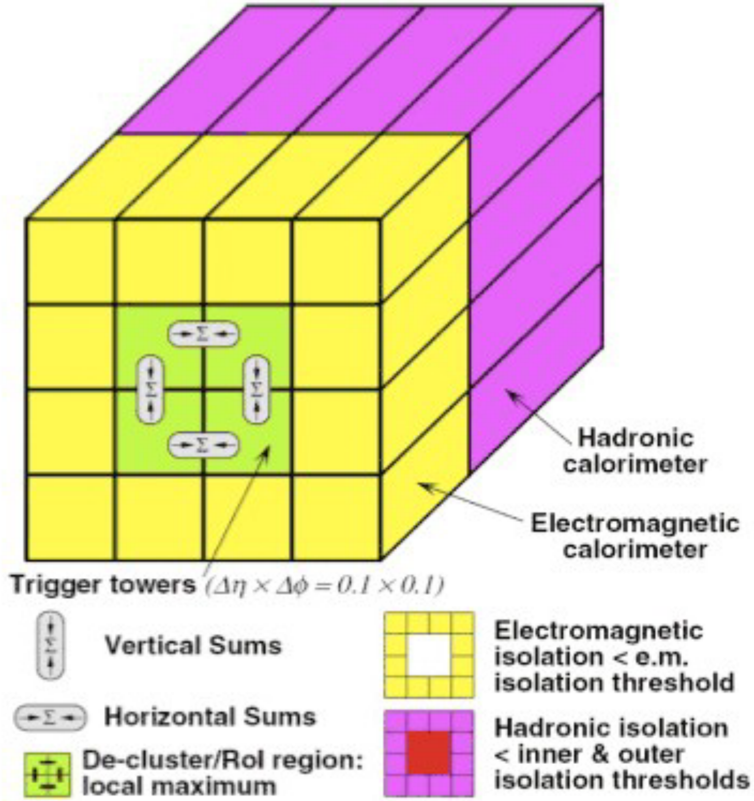

- Example: ATLAS e/γ trigger
  - Sum energy in calorimeter cells into EM and hadronic towers
  - Loop over grid and search in 4x4 towers for a local maximum 1x2 (2x1): cluster
- Can do something similar for other particles: jets, tau or sum the energy of all towers: missing E_T

CMS L1 muon trigger

Curved p_T-dependent muon path requires fast pattern recognition

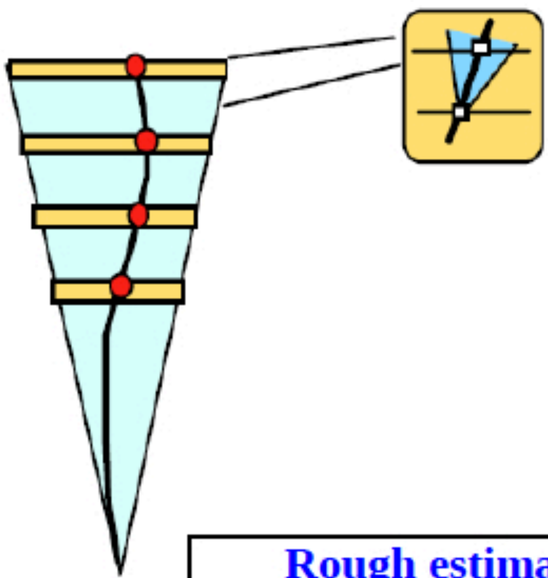

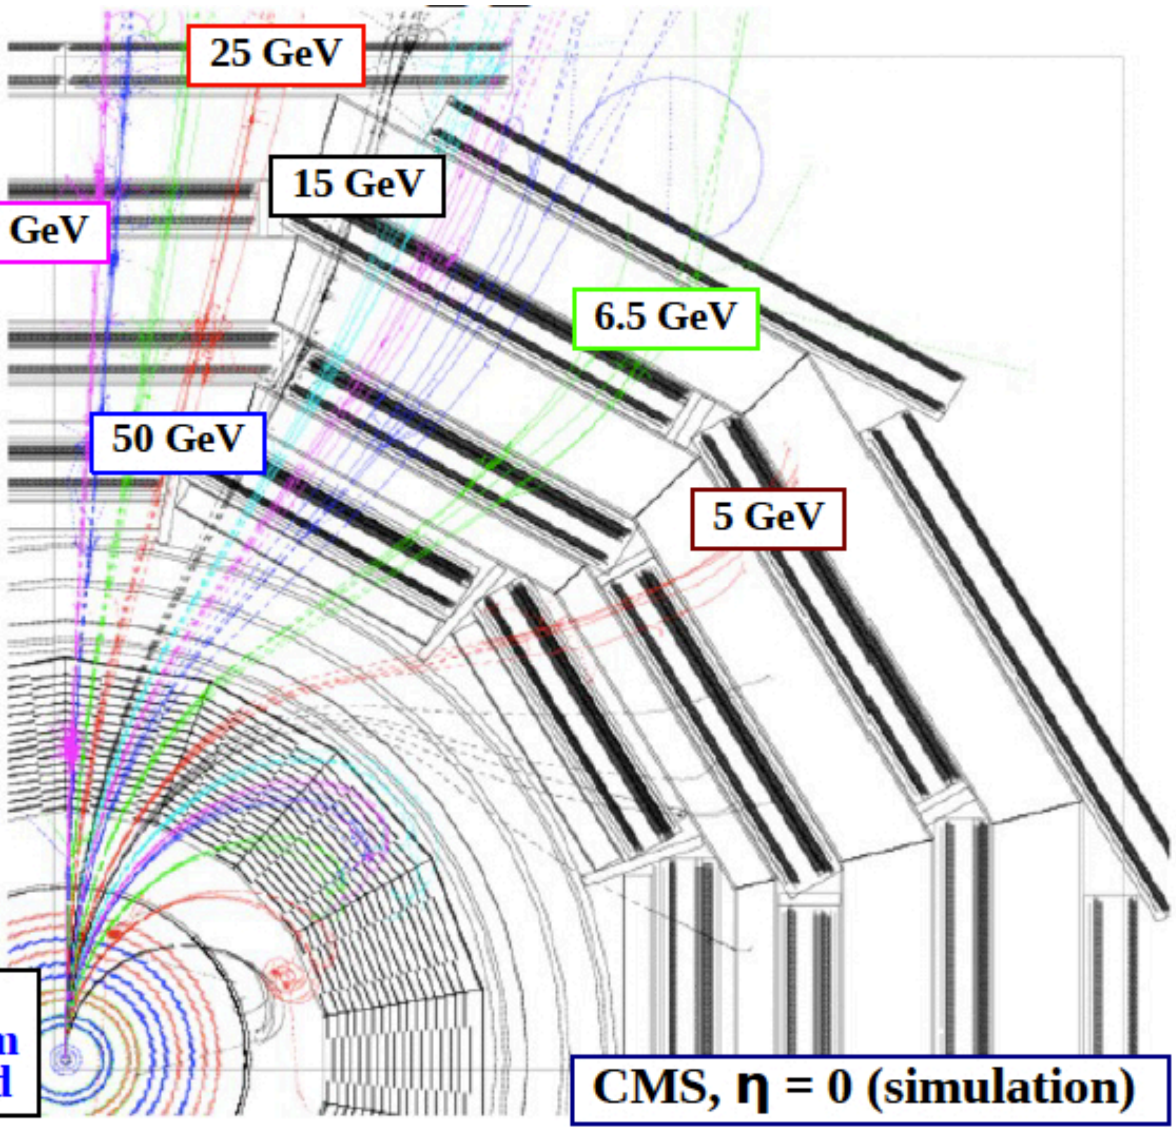

Rough estimate of muon p_T determined from bending in magnetic field

CMS, $\eta = 0$ (simulation)

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

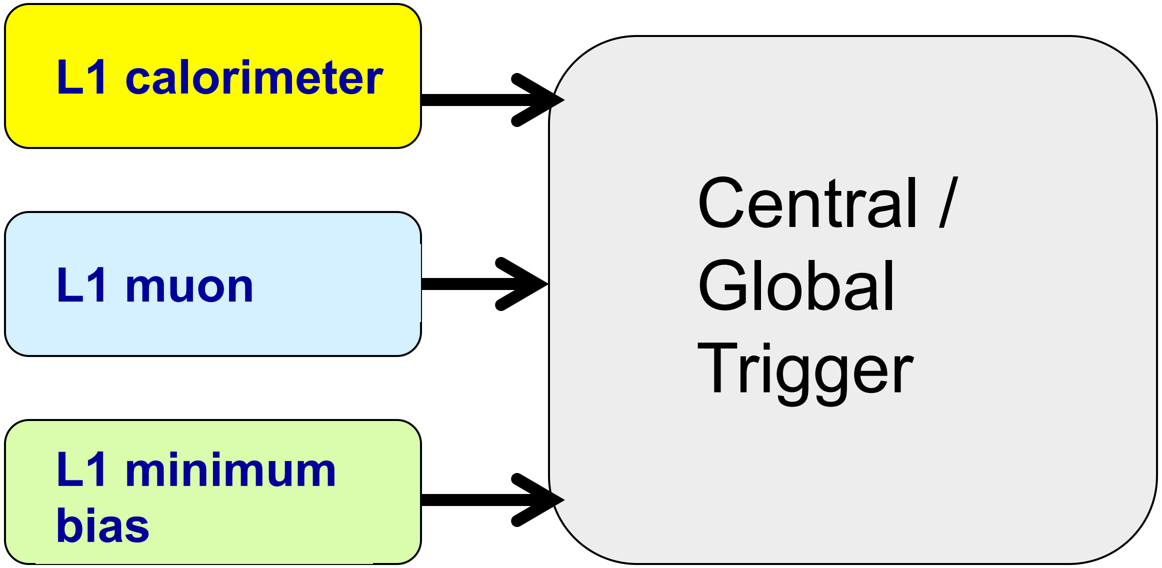

HLT Example: Muon

- 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

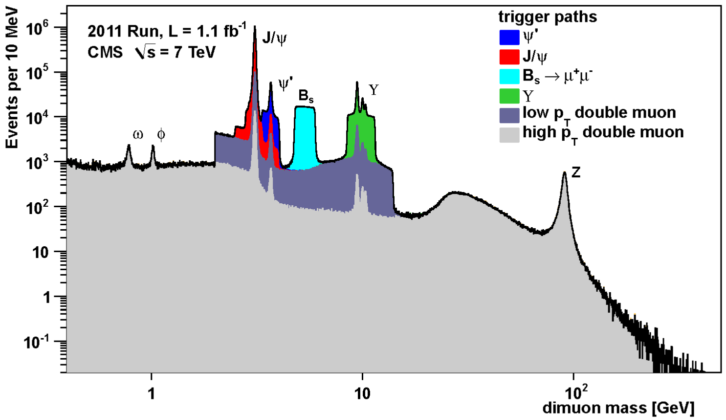

Upgrades

Long Shutdown

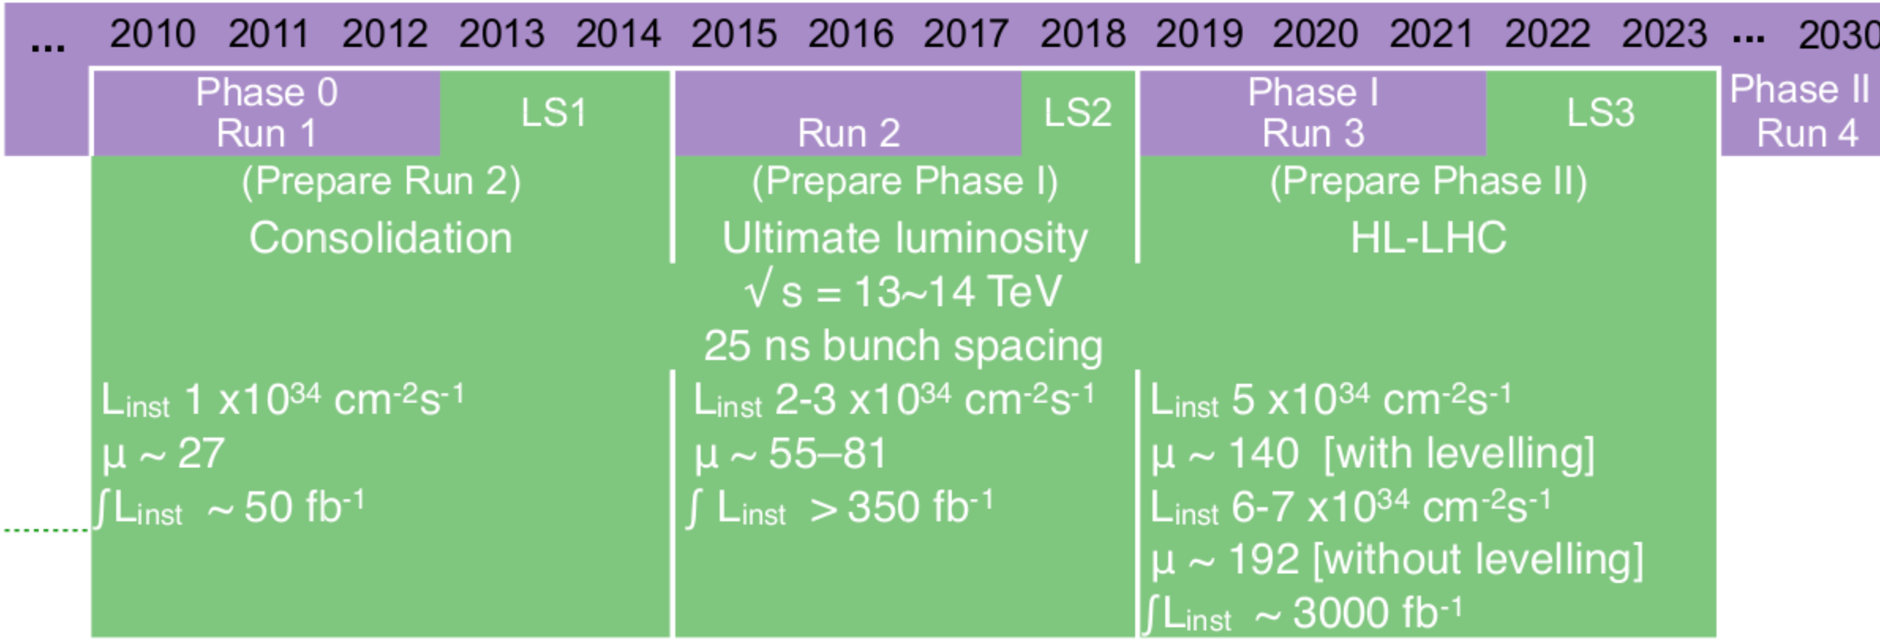

- LHC data acquisition system backbones installed >5 years ago
  - Very stable running in last 3 years, better than we were hoping for
- Current shutdown is occasion to
  - Upgrade core systems and review architectures
  - Introduce new technologies, retire obsolete ones
  - Follow changes on the detector side
- Prepare for challenges of Run2 (and Run3)

Long Shutdown

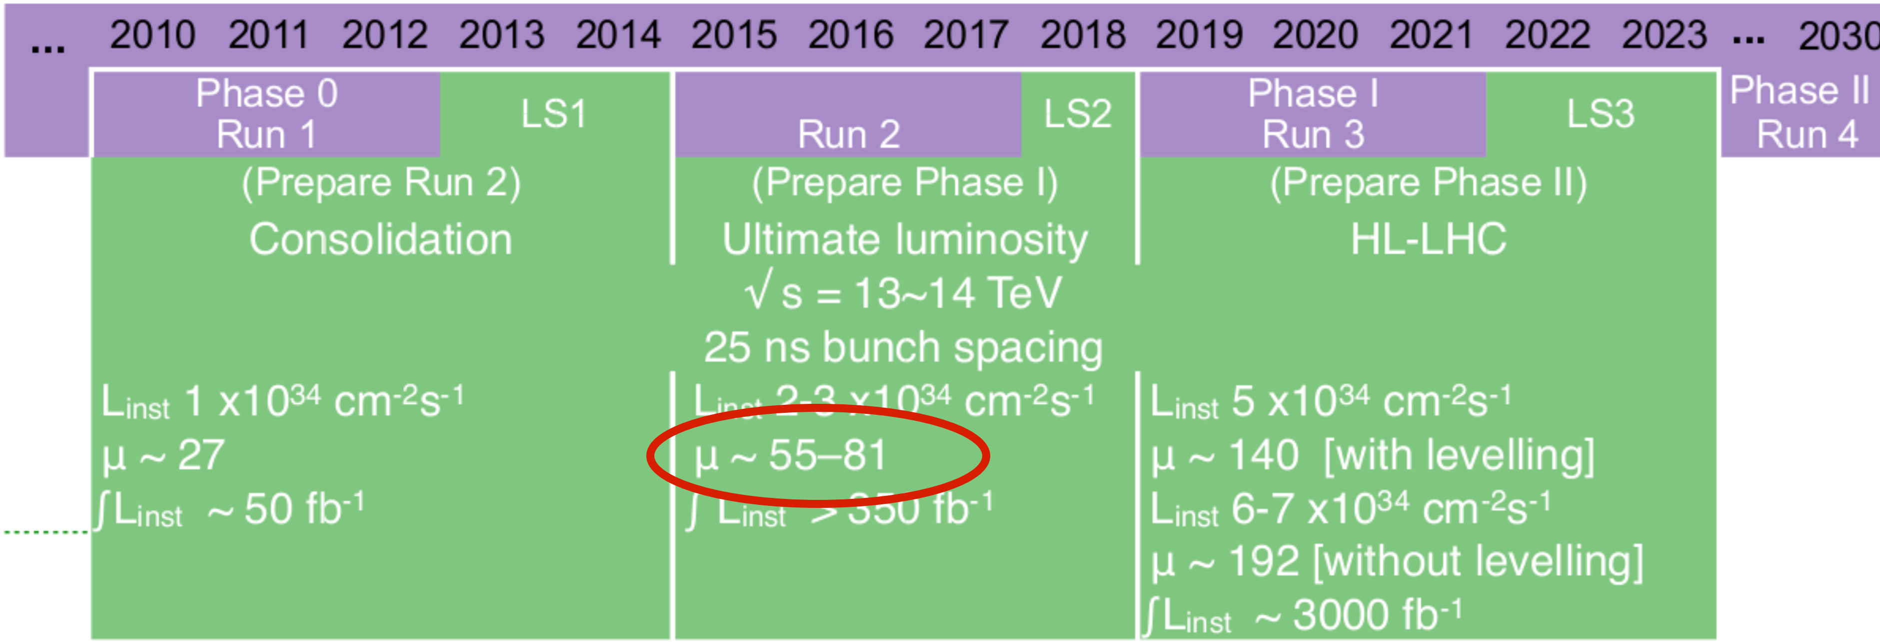

- LHC data acquisition system backbones installed >5 years ago
  - Very stable running in last 3 years, better than we were hoping for
- Current shutdown is occasion to
  - upgrade core systems and review architectures
  - introduce new technologies, retire obsolete ones
  - follow changes on the detector side
- prepare for challenges of Run2 (and Run3)

Pileup Issues

- CMS Simulation: 300 GeV H→ZZ→eeμμ at various luminosities

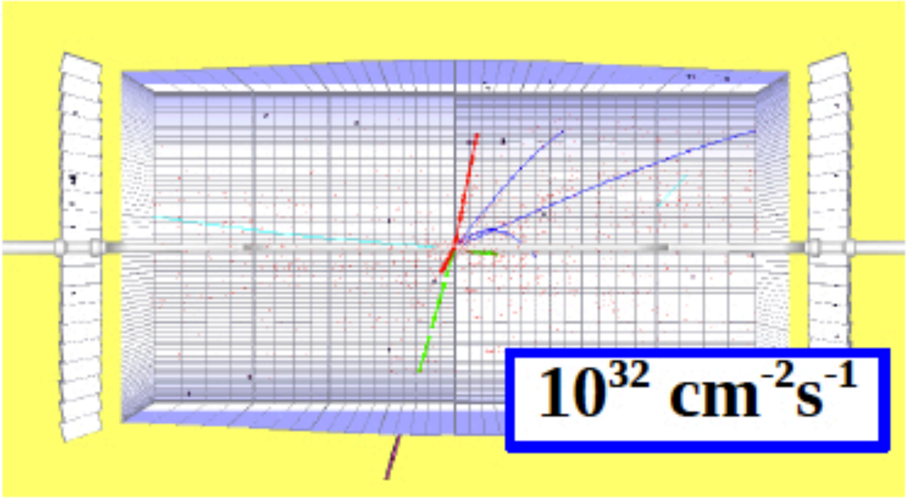

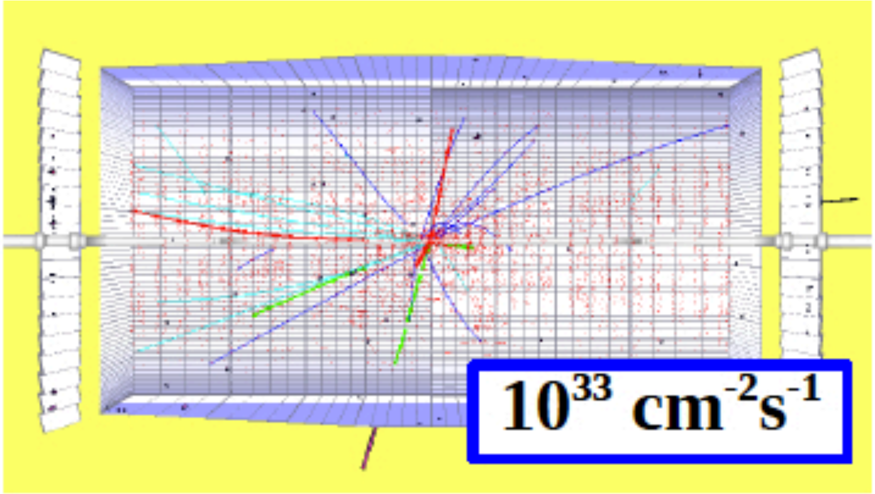

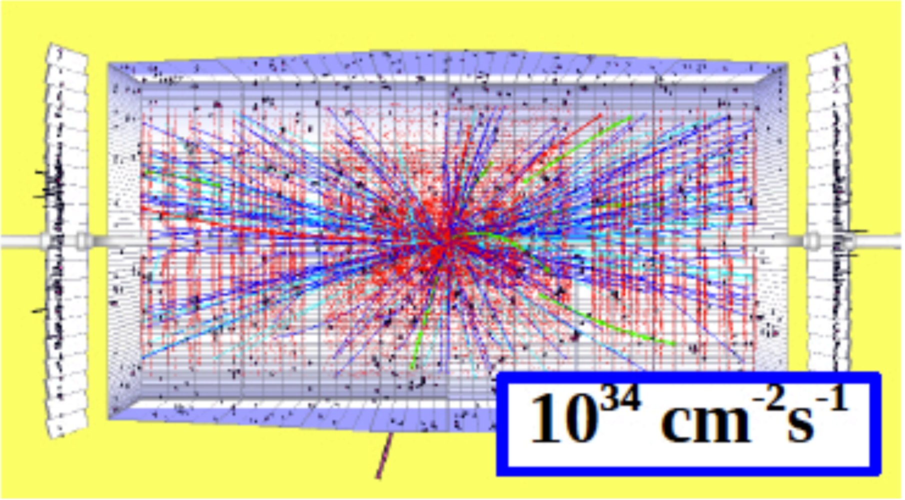

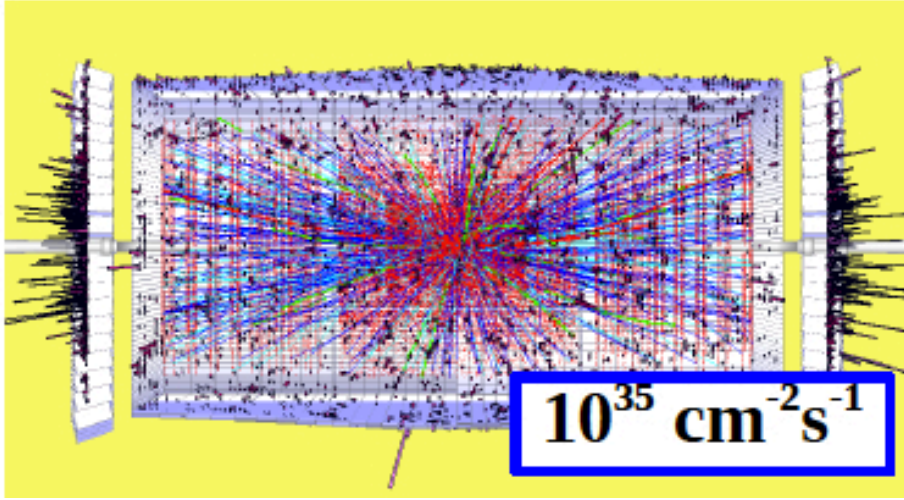

Run 2 challenges

- Increased pileup
  - more complex events → increased computing needs, affects trigger efficiency and rejection power
  - larger data size → bandwidth and storage

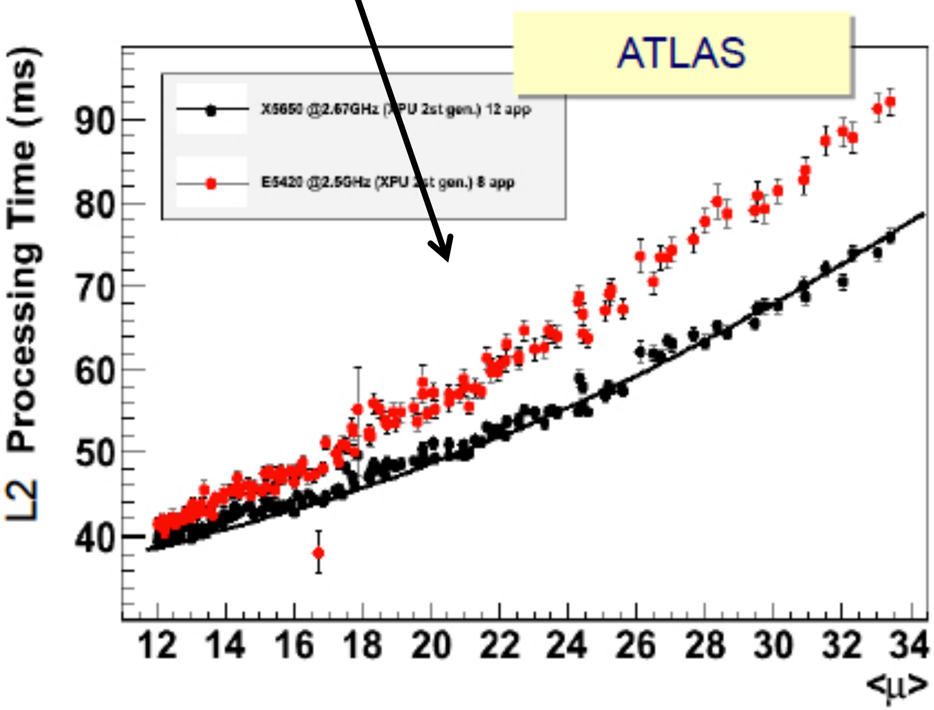

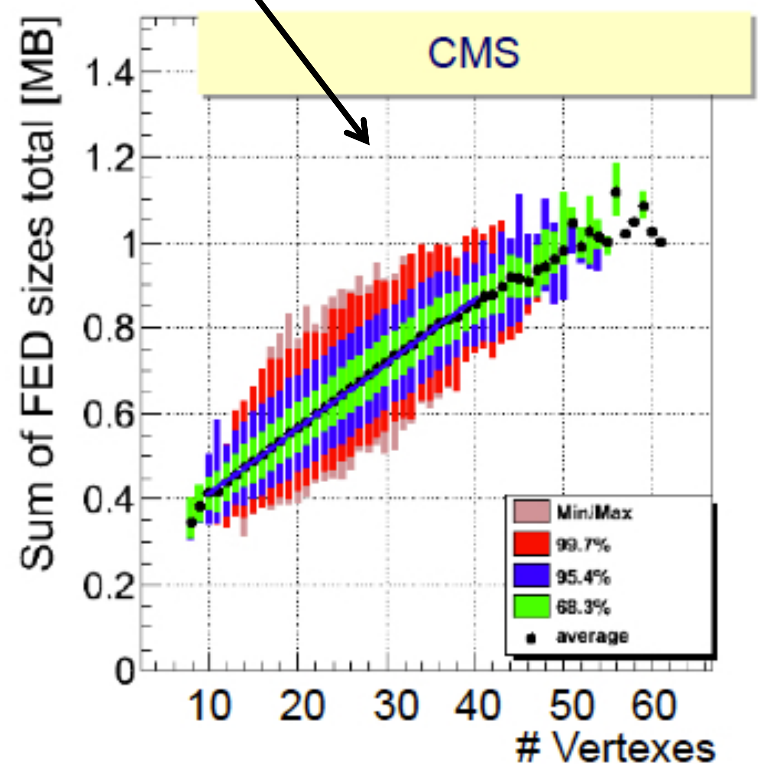

Upgrades for Run 2

ATLAS

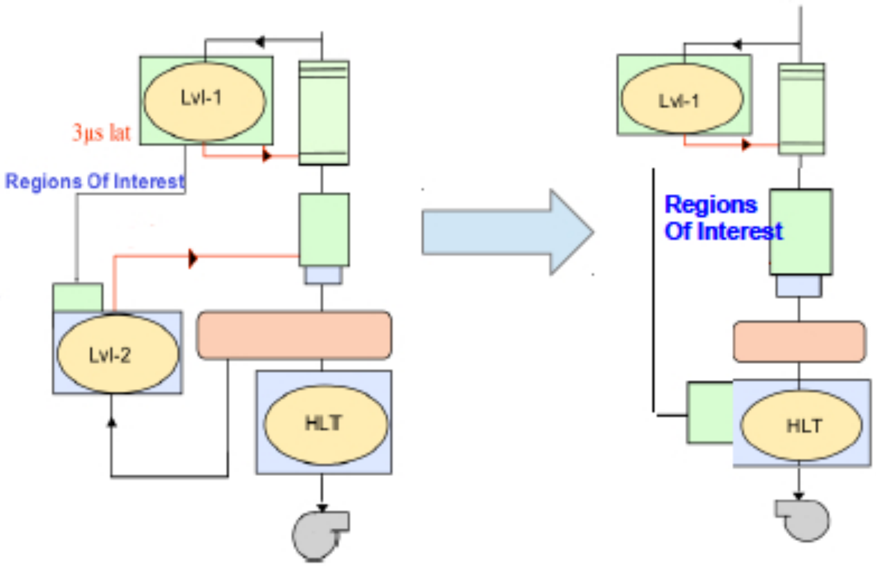

CMS

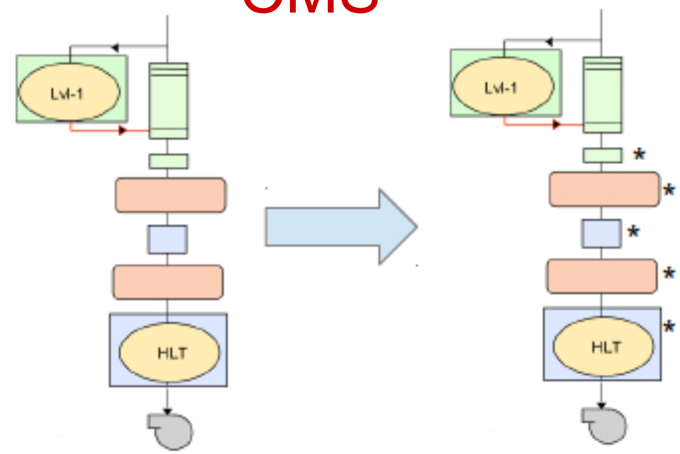

- Merge L2 and L3 into a single HLT farm
  - preserve Region of Interest, but diluted the farm separation and fragmentation
  - increased flexibly, computing power efficiency

- No architectural changes, but all network technologies replaced
  - Myrinet → Ethernet
  - Ethernet → Infiniband
- Filebased event distribution in the farm
  - achieve full decoupling between DAQ and HLT

LS2 and beyond

Alice

- 500 Hz → 50 kHz of PbPb interactions
  - Importance: physics with low S/B
  - Implies need to store 500 PB/ month (1 HI period)
- Data volume reduction
  - Online full reconstruction
  - discard raw data
- Combined DAQ/HLT/offline farm
  - COTS, FPGA and GPGPU

LHCb

- 1 MHz → 40 MHz readout and event building → trigger-less
  - trigger support for staged computing power deployment (only "soft" L0)
  - Need full event reconstruction with track finding and fitting plus particle identification to extract "interesting" event
- On detector zero suppression → radhard FPGA
- 4 TB/s event building
  - 100 kB/event

- On the long term, all experiments looking forward to significant increase in L1 trigger rate and bandwidth. Alice and LHCb will pioneer this path during LS2

Summary

- Challenge to design efficient trigger/DAQ for LHC
  - Very large collision rates (up to 40 MHz)
  - Very large data volumes (tens of MBytes 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)
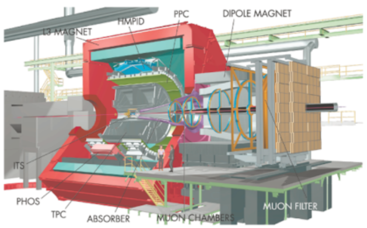	**4**	Pb-Pb **500** p-p **10^3**	**$5x10^7$** **$2x10^6$**	**25**	**1250** (10^2) **200** (10^2)
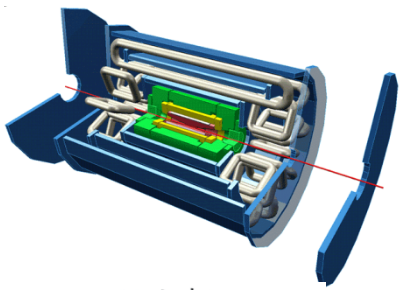	**3**	LV-1 **10^5** LV-2 **$3x10^3$**	**$1.5x10^6$**	**4.5**	**300** ($2x10^2$)
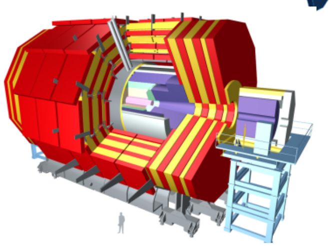	**2**	LV-1 **10^5**	**10^6**	**100**	**~1000** (10^2)
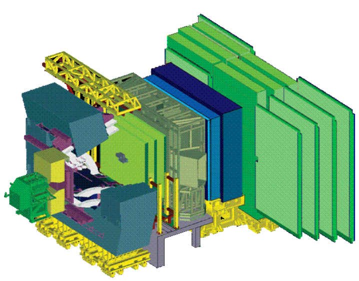	**2**	LV-0 **10^6**	**$3.5x10^4$**	**35**	**70** ($2x10^3$)

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