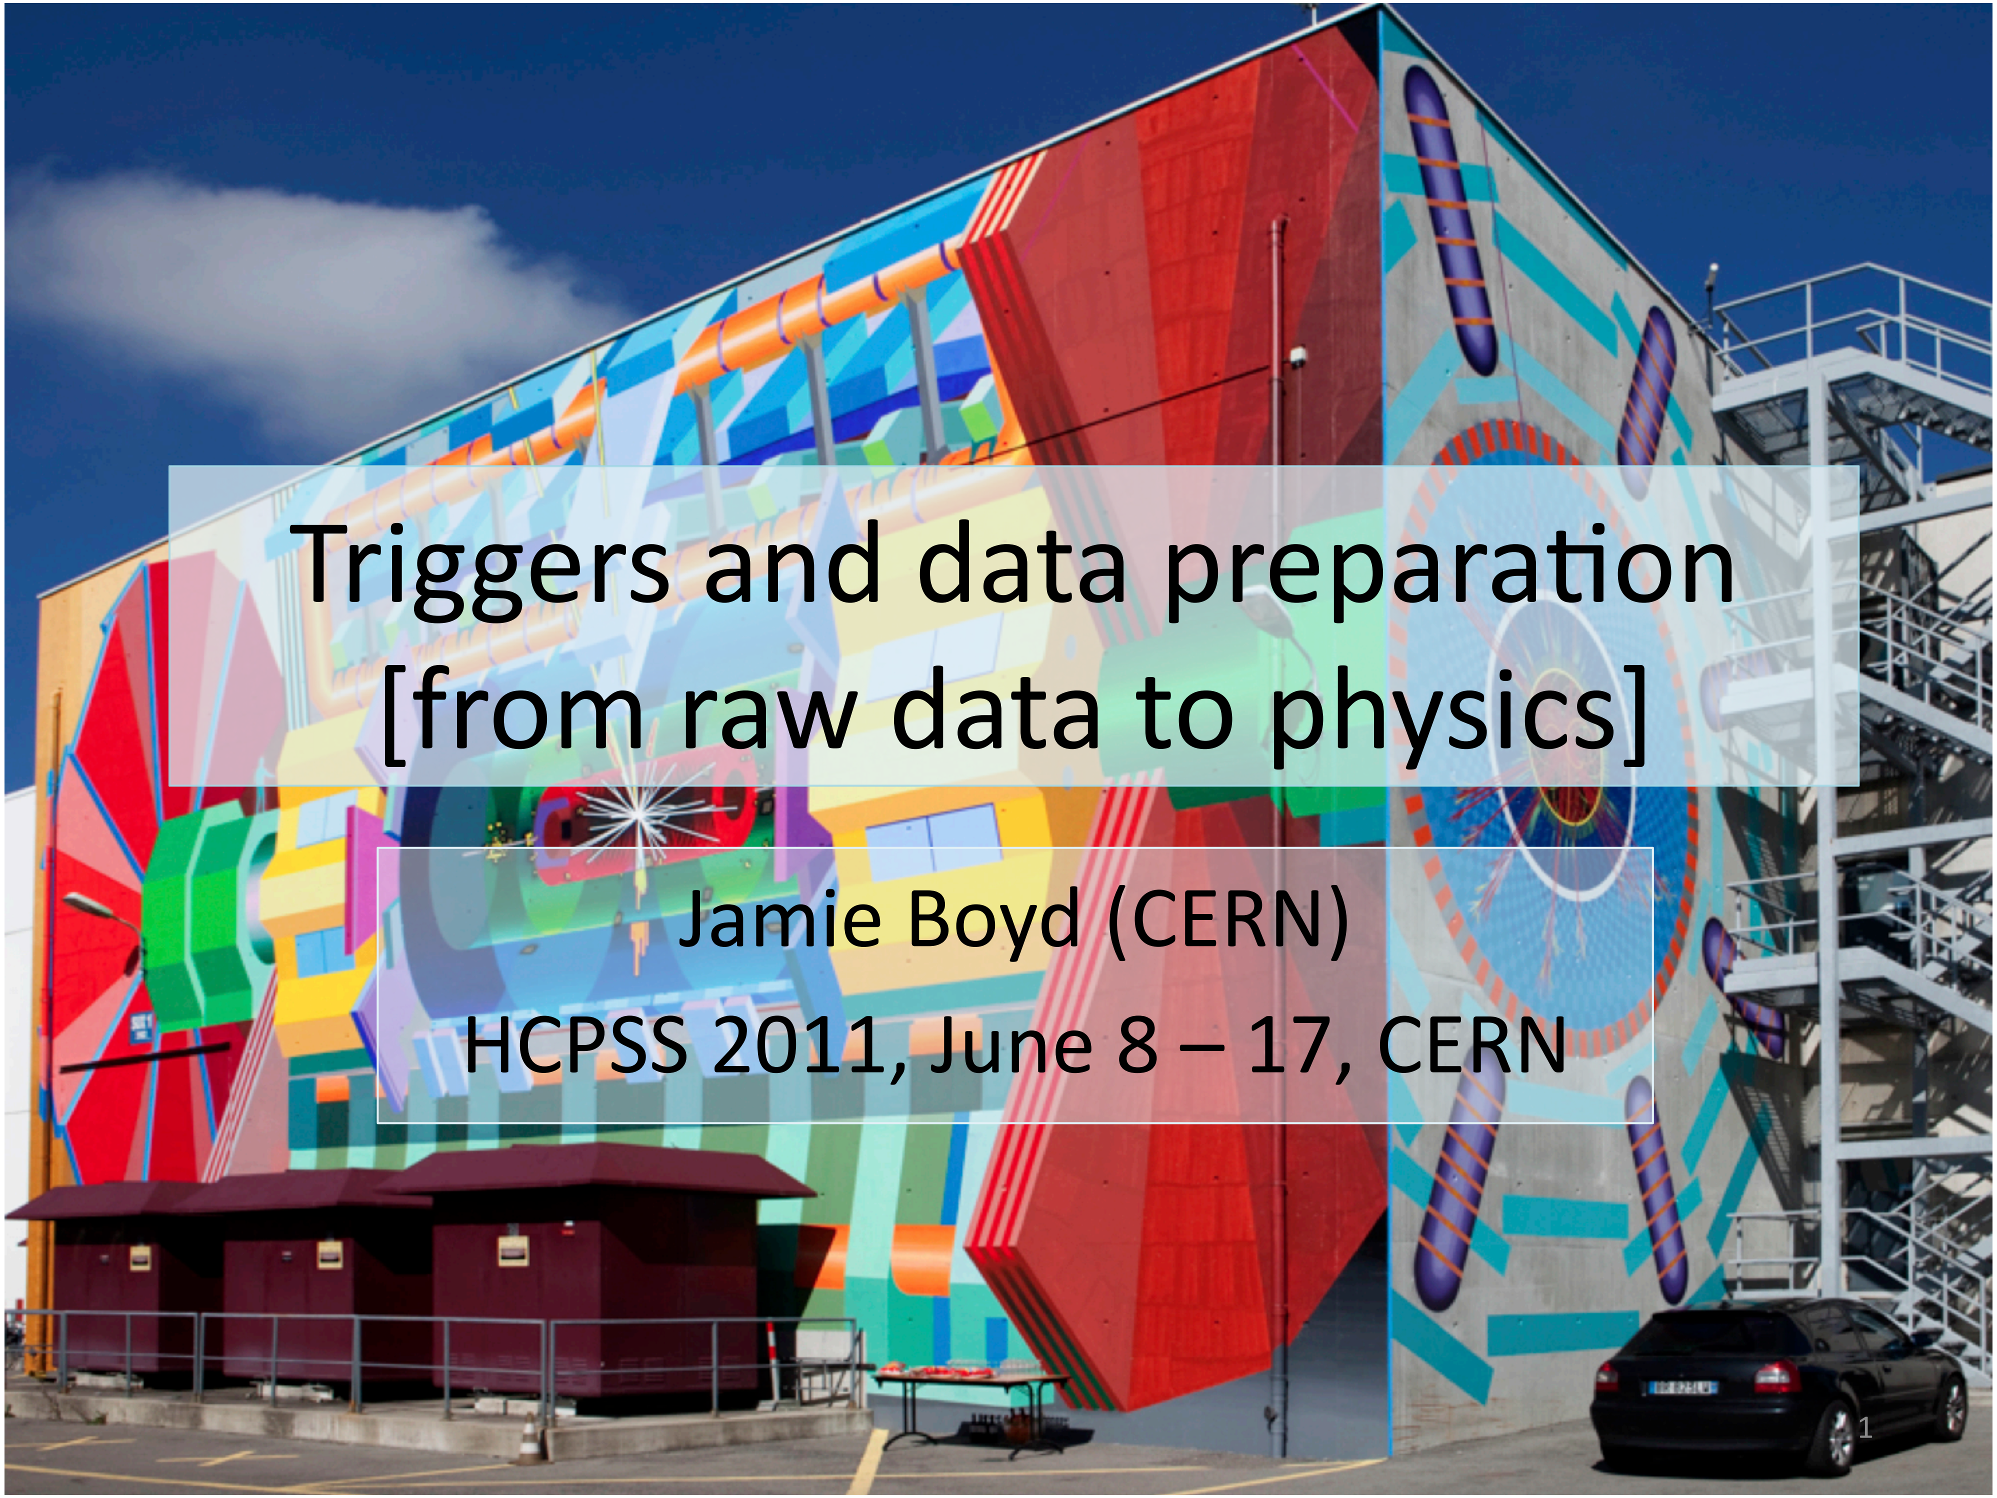

Triggers and data preparation [from raw data to physics]

Jamie Boyd (CERN)

HCPSS 2011, June 8 – 17, CERN

Lecture 1: The Trigger

Triggers and dat eparation
[from raw d to physics]

oyd (CERN)

HC 11, June 8 – 17, CERN

PreAmble

- 1 hour is much too little time to discuss Triggering at the LHC
- Since the LHC has started running at high luminosity only recently I really want to touch on the latest challenges faced by the LHC experiments
    - To give an idea about what currently are the main issues and the near future challenges
- There are a list of references in the backup which have more details on many of the things touched on in this lecture
- Since I work on ATLAS the lecture has a slight bias towards this experiment BUT I do want to show some information from other experiments (mostly CMS)

Outline

- Introduction
  - Why we need a trigger ?
  - Basic Trigger/DAQ concepts
- Triggers at the LHC experiments
  - Design principles
  - Examples
  - Comparisons between experiments
- Trigger operations, menus & analysis
  - Trigger efficiency measurement
  - Trigger menus
  - Operational issues

Introduction…

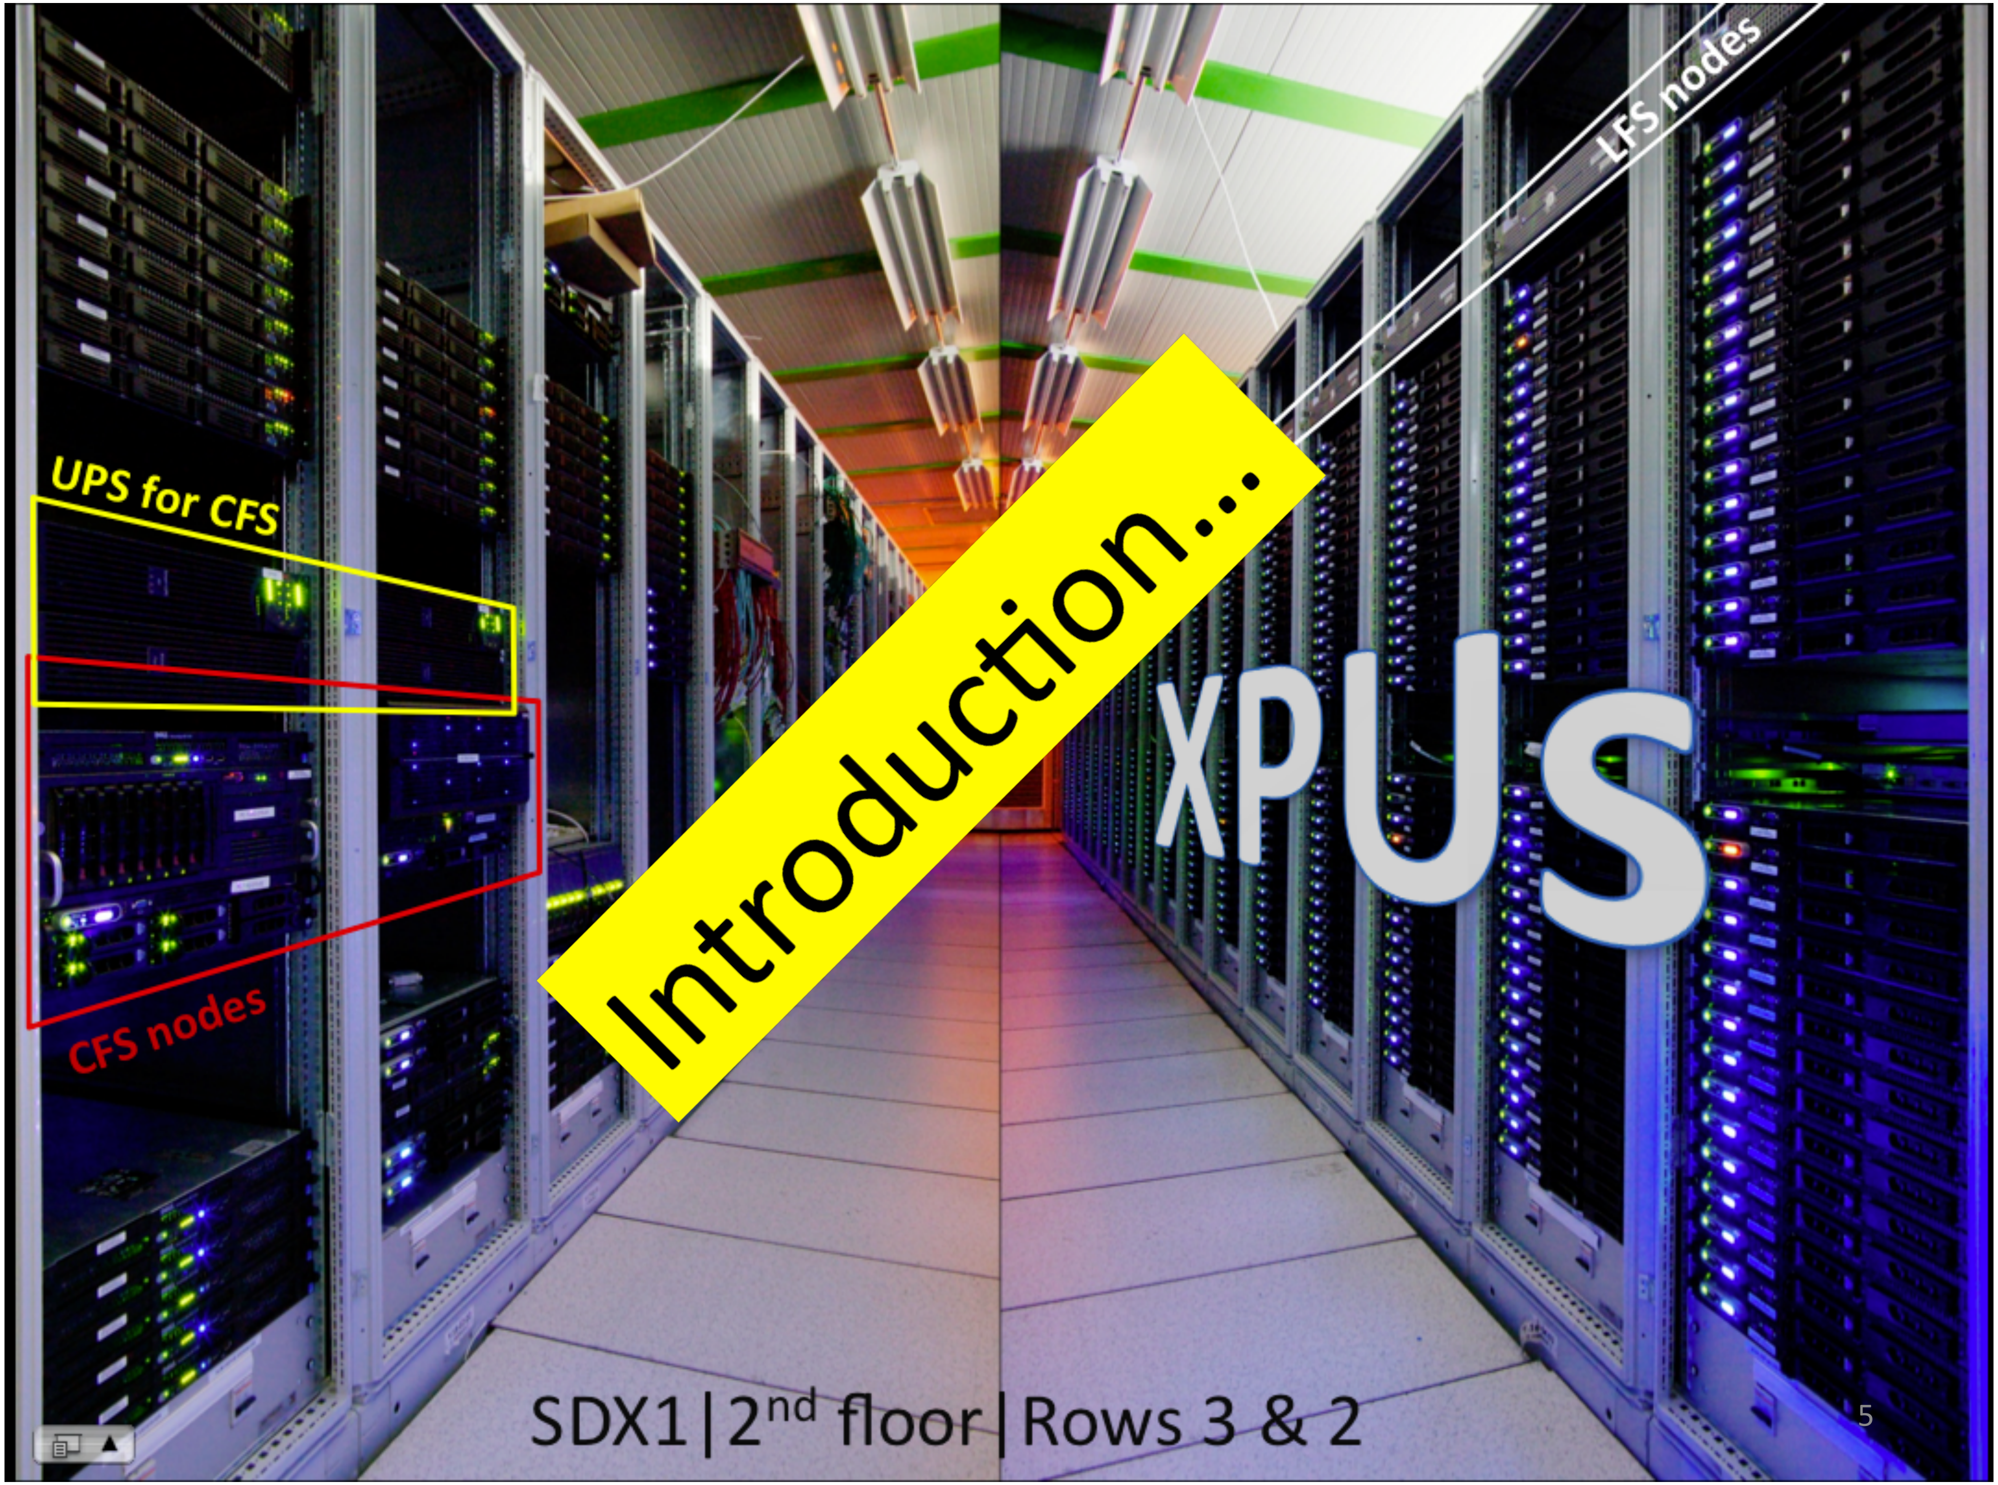

SDX1 | 2nd floor | Rows 3 & 2

LHC Status

LHC status	Startup (Nov 2009)	End 2010	Current (June 2011)	Design (start 2014?)
√s (TeV)	0.9	7	7	14
Lumi ($cm^{-2}s^{-2}$)	10^{27}	10^{32}	10^{33}	10^{34}
Lumi (pb^{-1}/15hr fill)	0.000004	0.4	40	400
N-Bunches	1	350	1000	2800
Pileup (μ)	<<1	3	7	25
Bunch Spacing (ns)	No bunch trains	150	50	25

The LHC parameters have been evolving since start up and these directly affect trigger and DAQ performance. This means any number or plot only makes sense for a given luminosity scenario. In this talk I have plots/numbers taken from various stages in this procedure, as well as some for the 'design' luminosity – Remember these are not always comparable.

Why we need a trigger?

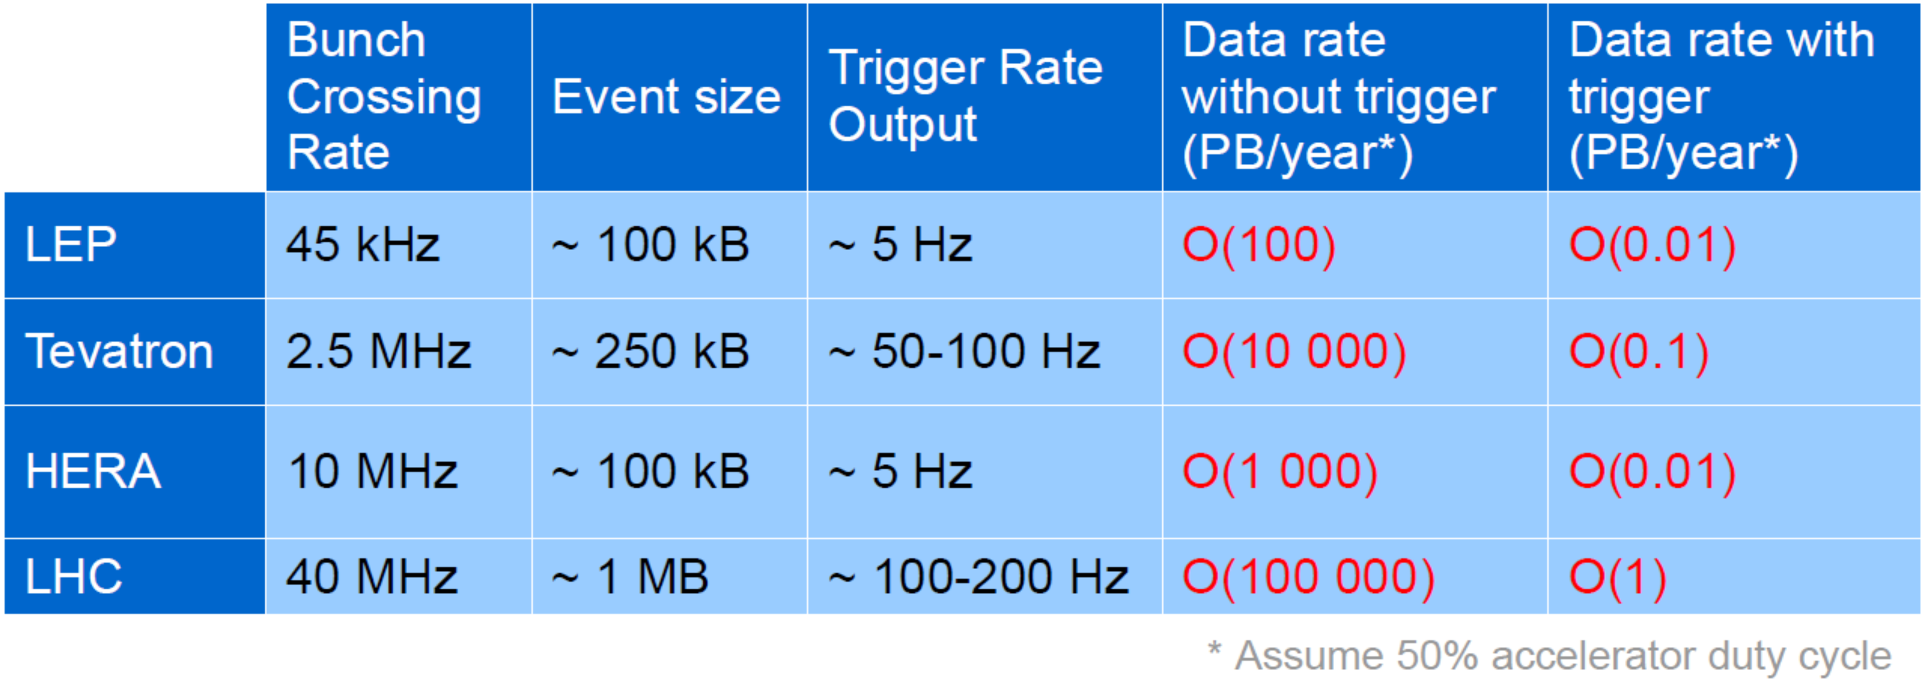

	Bunch Crossing Rate	Event size	Trigger Rate Output	Data rate without trigger (PB/year*)	Data rate with trigger (PB/year*)
LEP	45 kHz	~ 100 kB	~ 5 Hz	O(100)	O(0.01)
Tevatron	2.5 MHz	~ 250 kB	~ 50-100 Hz	O(10 000)	O(0.1)
HERA	10 MHz	~ 100 kB	~ 5 Hz	O(1 000)	O(0.01)
LHC	40 MHz	~ 1 MB	~ 100-200 Hz	O(100 000)	O(1)

* Assume 50% accelerator duty cycle

Without a trigger the output data size is impossibly large:
Too big to store
Too big to reconstruct, and analyze

Rejection factor of trigger typically 10^4-10^5

p p(bar) Cross Sections

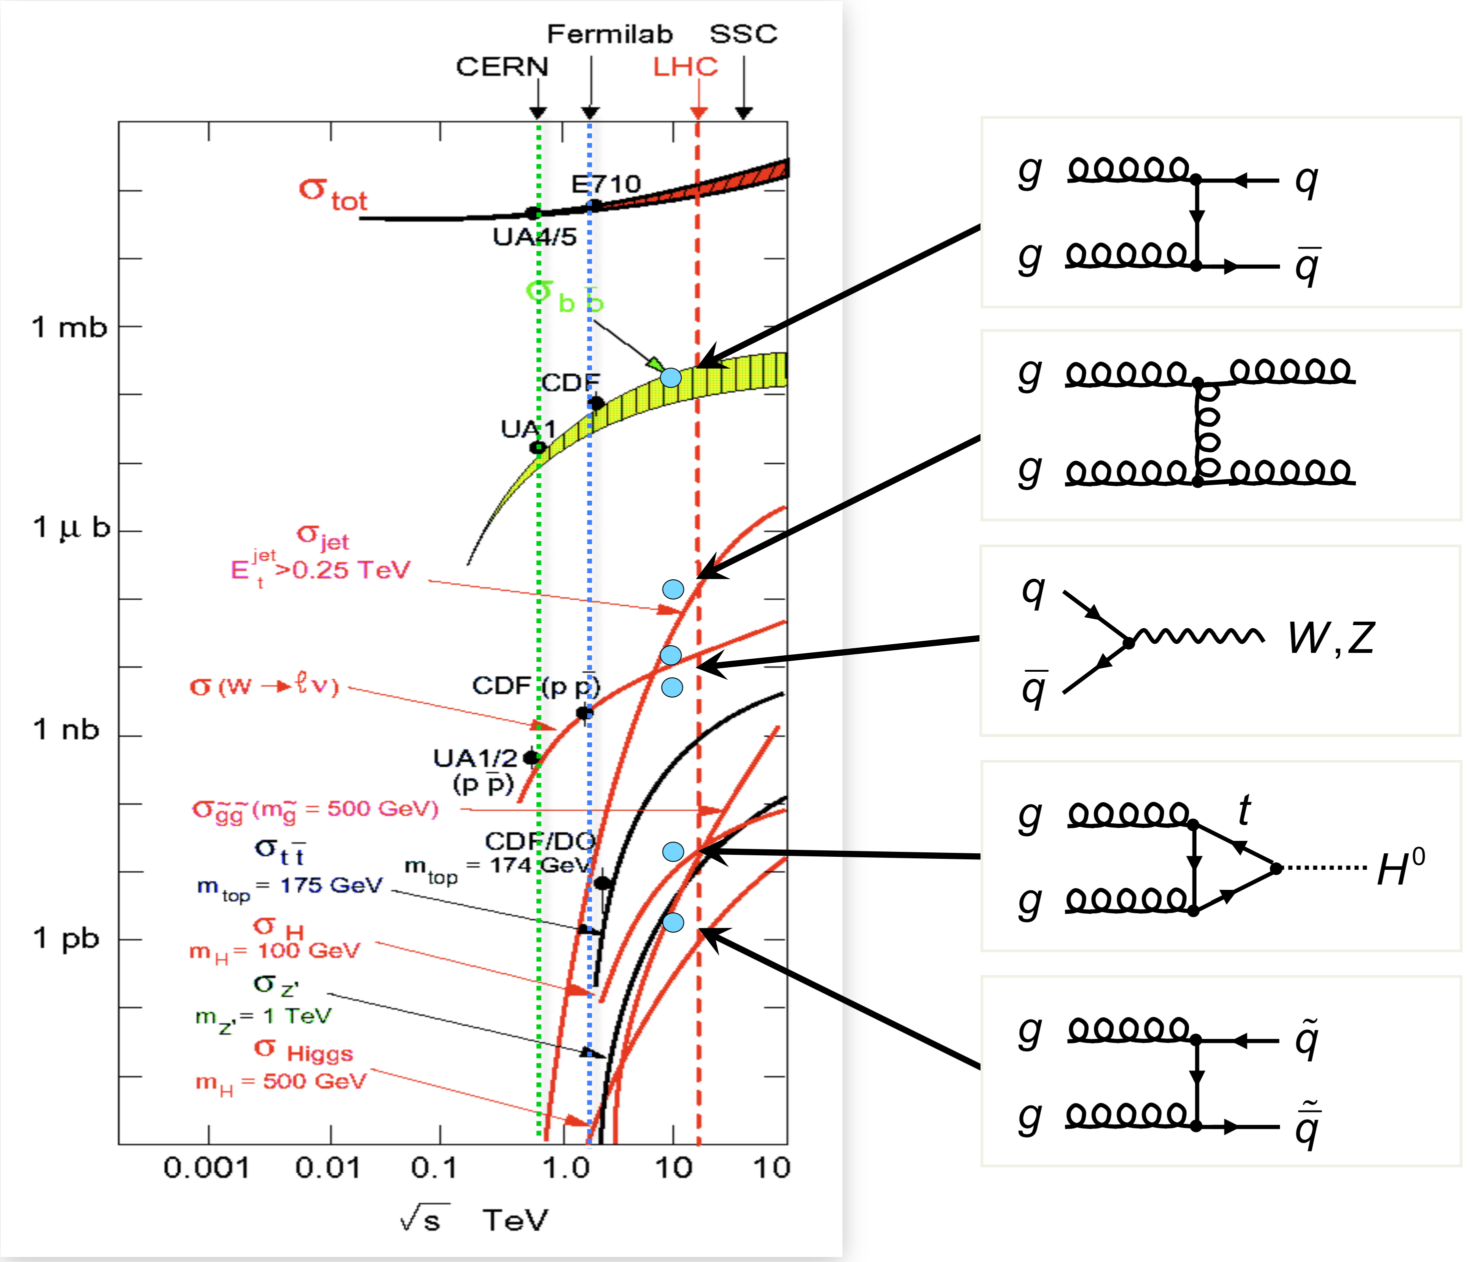

Quark-flavour production

High-p_T QCD jets

W, *Z* production

gluon-to-Higgs fusion

squarks, gluinos
(m ~ 1 TeV)

p p(bar) Cross Sections

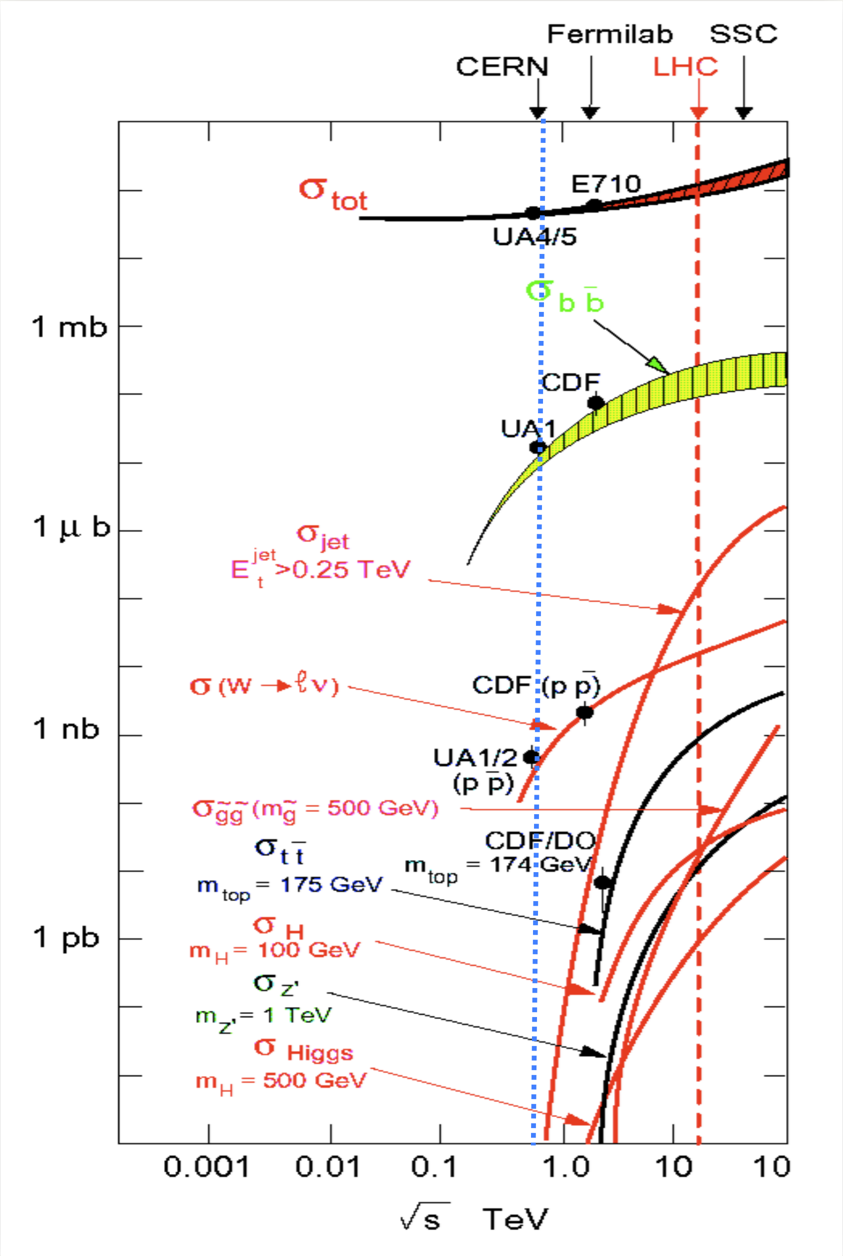

Process	Cross section (nb) at 14 TeV CM energy	Production rates (Hz) at $L=10^{34}$ cm^{-2}s^{-2}
Inelastic	10^8	10^9
$b\bar{b}$	5×10^5	5×10^6
$W \to \ell\nu$	15	150
$Z \to \ell\ell$	2	20
$t\bar{t}$	1	10
Z' (1 TeV)	0.05	0.5
$\tilde{g}\tilde{g}$ (1 TeV)	0.05	0.5
H (120 GeV)	0.04	0.4
H (180 GeV)	0.02	0.2

Recall: $1\ \text{pb}^{-1} \triangleq 10^{36}\ \text{cm}^{-2}$

$\Rightarrow 15\ \text{nb} \times 10^{34}\ \text{cm}^{-2}\text{s}^{-1} =$

$15\ \text{nb} \times 10^{-2}\ \text{pb}^{-1}\text{s}^{-1} = 150\ \text{Hz}$

p p(bar) Cross Sections

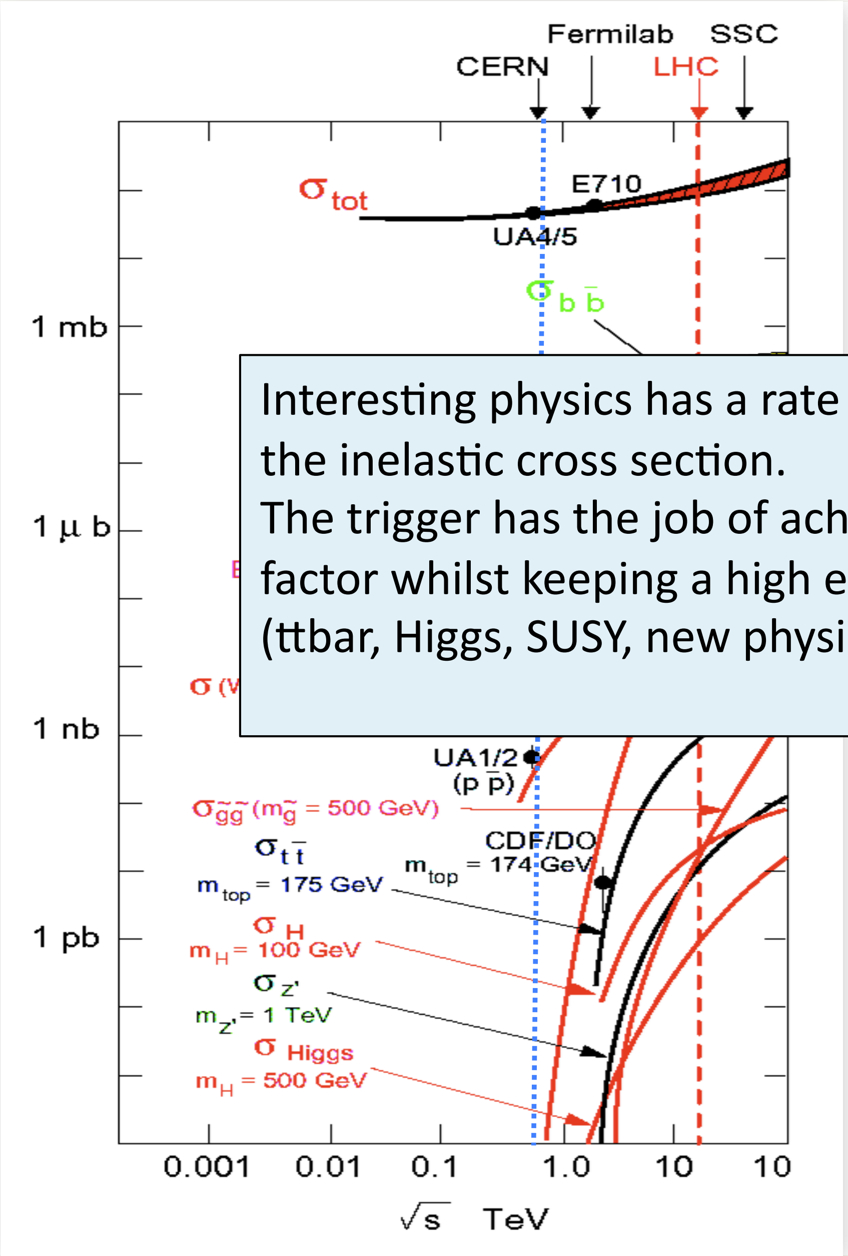

Interesting physics has a rate many of orders of magnitude less than the inelastic cross section.
The trigger has the job of achieving the required ~10^4-10^5 rejection factor whilst keeping a high efficiency for interesting physics processes (ttbar, Higgs, SUSY, new physics we have not thought of yet!)

Process	Cross section (nb) at 14 TeV CM energy	Production rates (Hz) at $L=10^{34}$ cm^{-2}s^{-2}
Inelastic	10^8	10^9
		10^6
		50
		20
		10
		.5
$\tilde{g}\tilde{g}$ (1 TeV)	0.05	0.5
H (120 GeV)	0.04	0.4
H (180 GeV)	0.02	0.2

Recall: 1 pb$^{-1} \triangleq 10^{36}$ cm^{-2}

$\Rightarrow$ 15 nb $\times$ 10^{34} cm^{-2}s^{-1} =

15 nb $\times$ 10^{-2} pb^{-1}s^{-1} = 150 Hz

Trigger must be…

- Fast (low latency)
  - Low dead-time
- Have a Large rejection factor (order of 10^4-10^5)
- Highly efficiency for interesting events
  - Events which are rejected are lost forever
  - Must be able to measure the efficiency
  - Bias must be understood and manageable
    - At hadron collider by definition the trigger will cut into the physics
- Affordable
- Flexible
  - Cannot foresee all possible interesting signals
- Redundant
  - Need to be able to cope with something not working properly

Physics signatures

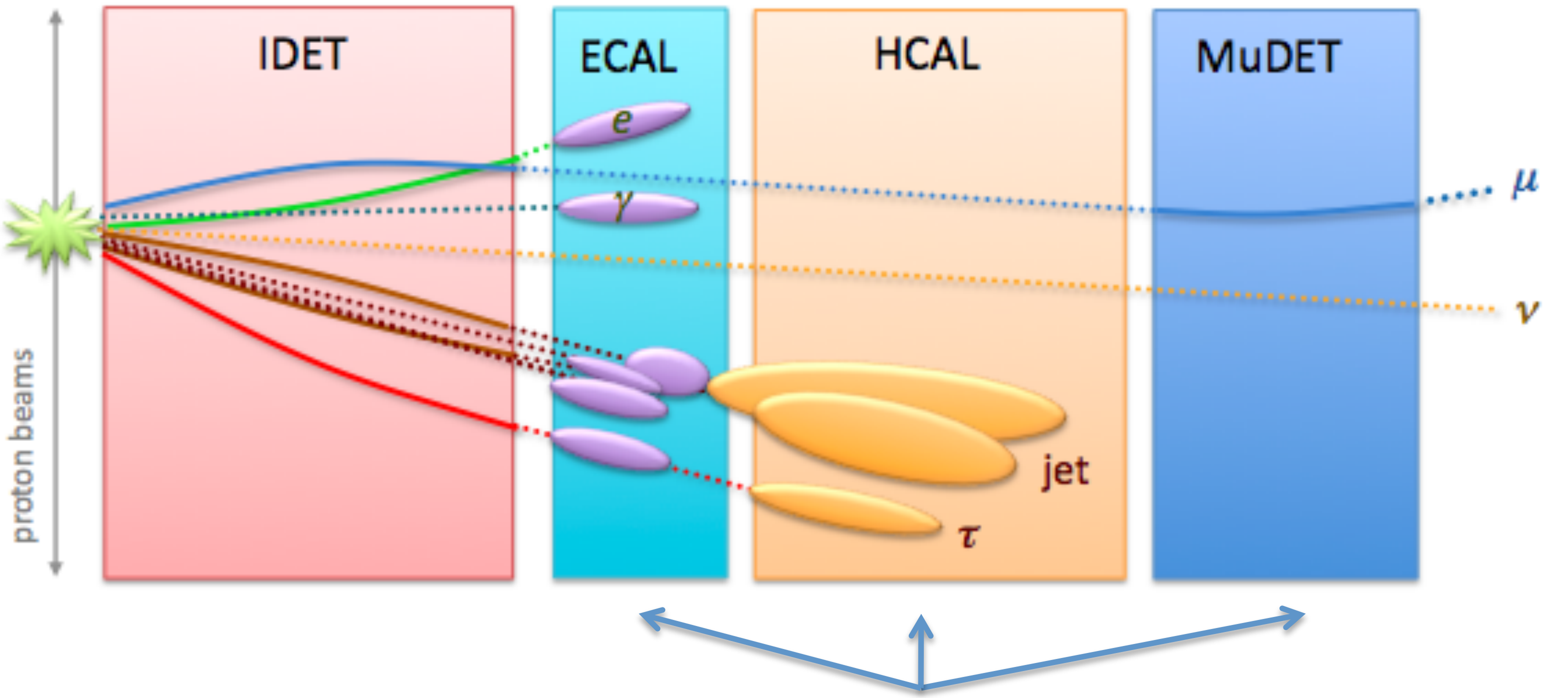

All systems used in HLT

Used in L1 trigger

Trigger on:

- High P_T isolated leptons (Electrons and Muons)
- High P_T isolated photons
- High P_T Jets
- High Missing transverse Energy or Total transverse energy
- Multi object triggers (combinations of above – eg. Jet + Lepton)

Trigger Basics

- **Not possible** to achieve the required background rejection whilst keeping a high signal efficiency in 25ns
- Use **multi-level trigger**
  - Level 1 (L1)
    - Runs in hardware
    - Very fast
  - Higher Level Trigger (HLT)
    - Runs in software
    - More complex reconstruction
- Buffer detector data at various stages whilst the trigger is making its decision

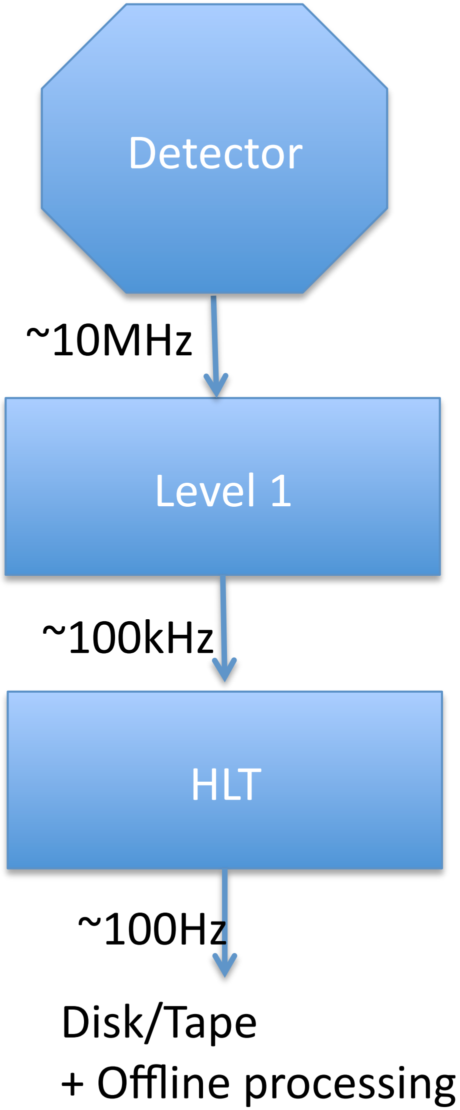

Dead-time

- "Dead-time" is the fraction of time an experiment **cannot record data**
- Sources of dead-time:
  - Read-out and trigger dead-time
  - Operational dead-time
    - Examples: start/stop data-taking periods, high voltage trip, power problem…
- Note that trigger dead-time logic is required to prevent triggering on another event before the detector is fully read-out!
  - In practice have lots of buffering and pipelined electronics to minimize dead-time
- Note trigger dead-time veto's events so causes loss of efficiency but **NO BIAS**
  - Eg. If you have a very busy event this may cause dead-time but that will veto proceeding events. Must not veto the busy event. Otherwise we could lose all blackhole events

LHC experiment Trigger systems

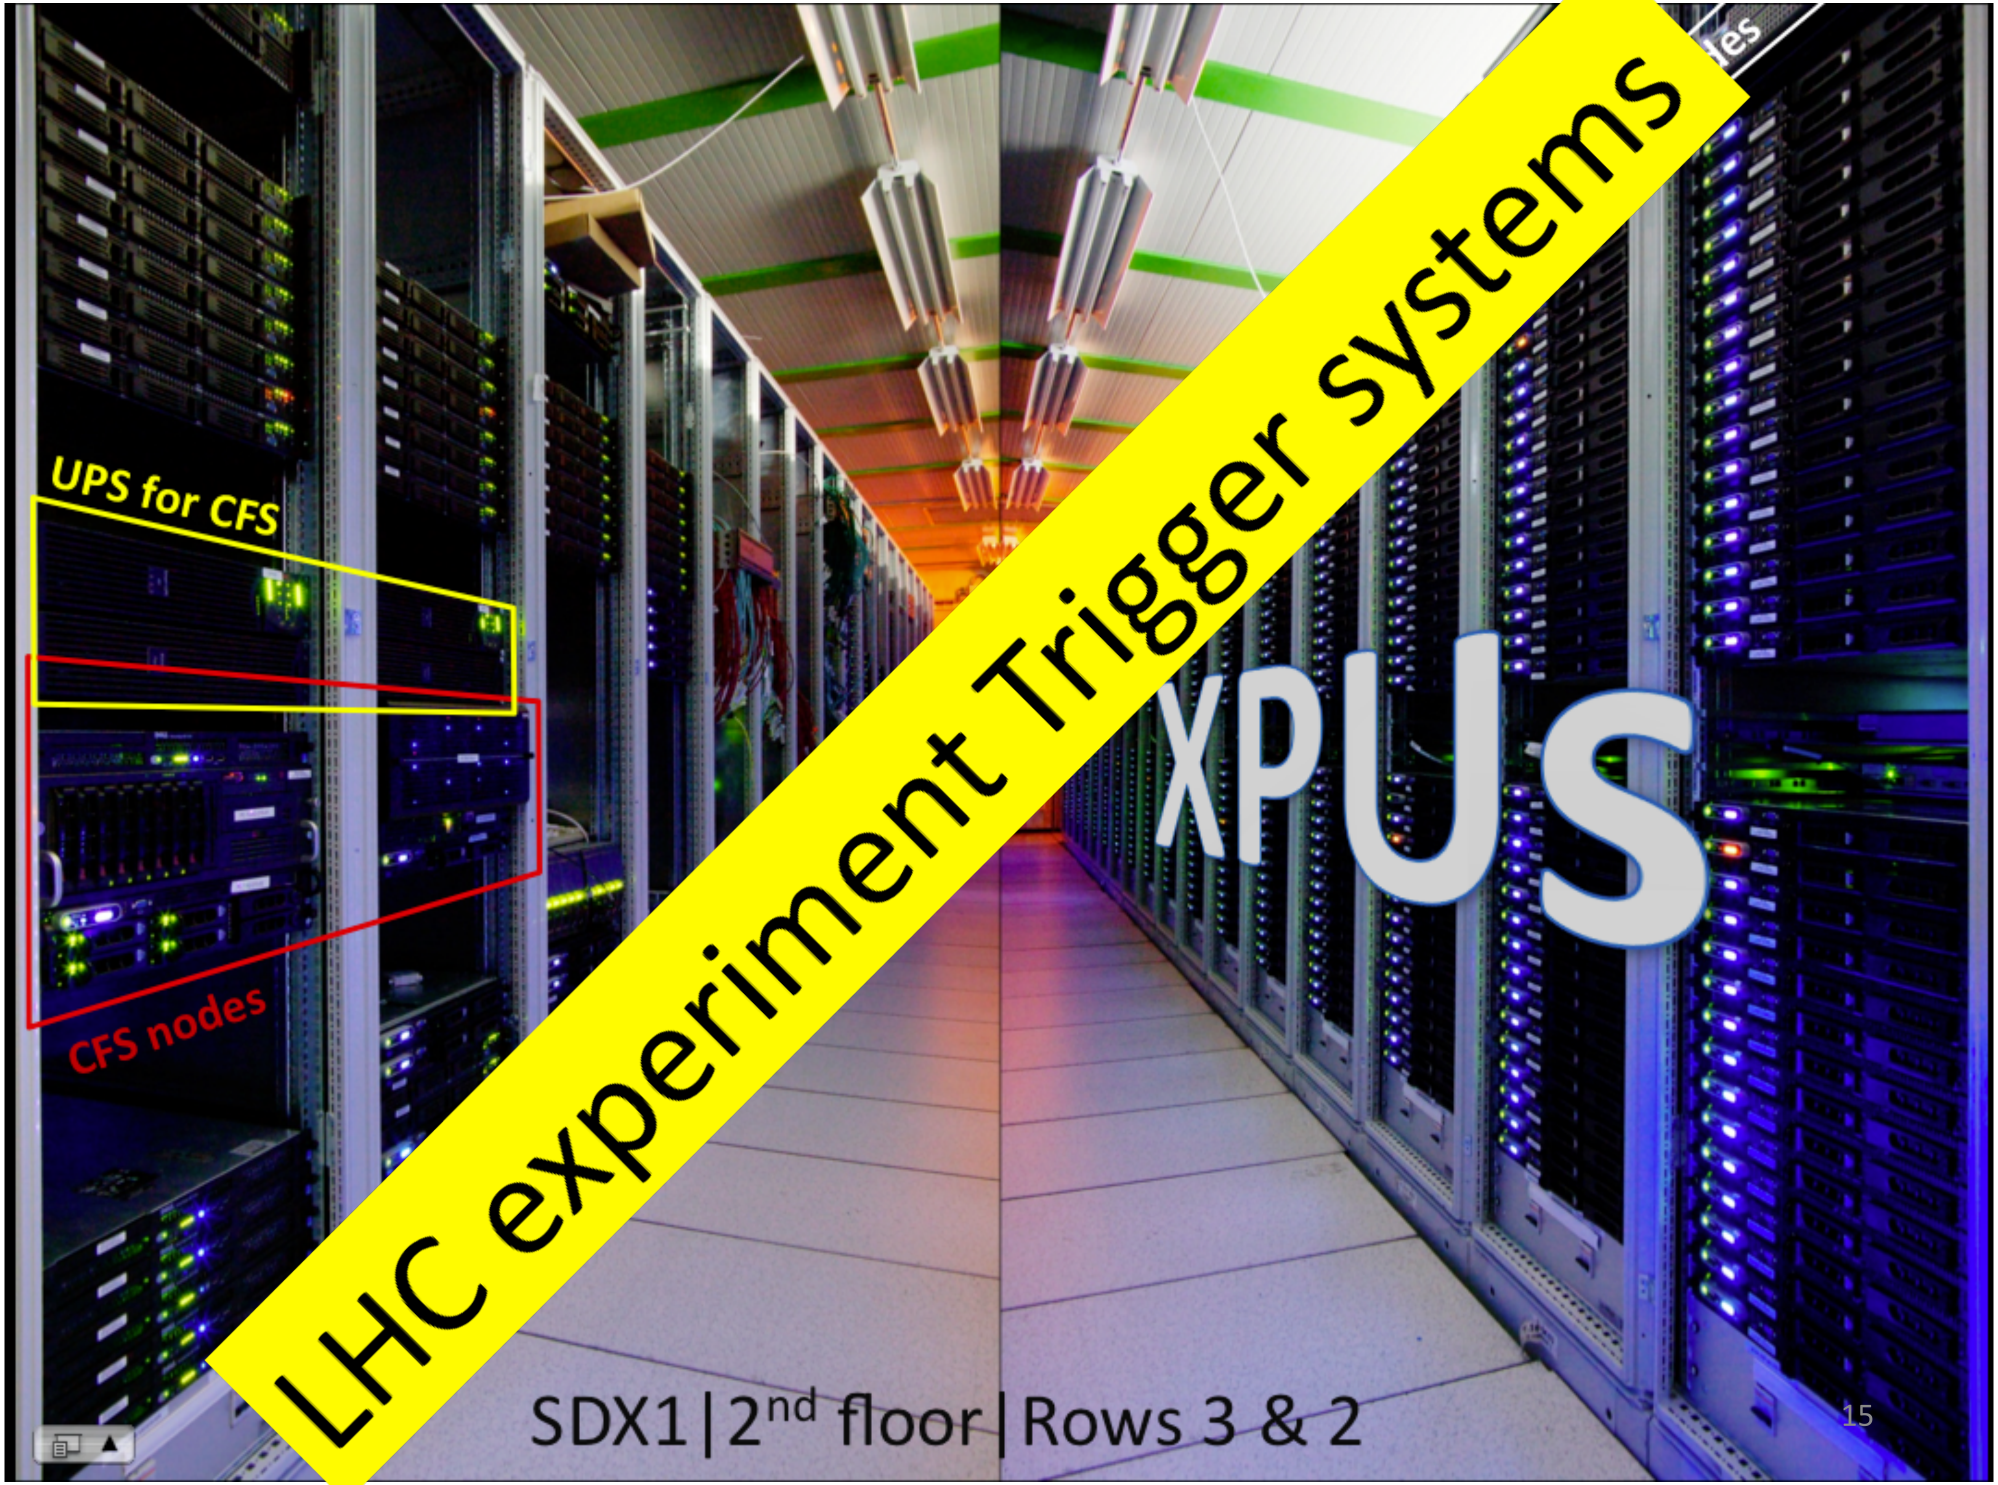

SDX1 | 2nd floor | Rows 3 & 2

Physics Triggers @ LHC

- At ATLAS/CMS the trigger system is quite similar
- L1 trigger (design: 40MHz -> ~100kHz)
  - Hardware (custom made electronics)
  - Input from Muon Detectors and Calorimeters (El, Ph, Jets, Tau and Missing Energy)
    - No track trigger at L1
  - Uses coarser granularity than full detector allows
  - Fixed latency ~3µs
- High Level Trigger (HLT) (design: ~100kHz -> ~300Hz)
  - Software trigger (running on commercial PC farm)
  - Includes tracking detectors
    - Allows b-jet triggers – from displaced vertex inside jet
  - Full detector granularity gives better rejection than L1
  - Many sub steps
  - Variable latency average ~100ms

Level 1 trigger design

- To achieve required high rates
  - Must use custom electronics
    - Application-specific integrated circuits (ASICs)
    - Field programmable Arrays (FPGA)
  - System must be pipelined (processing many events at the same time)
    - ATLAS 2.5μs latency is 100 bunch crossings
    - If not pipelined this would latency would be dead-time
  - Parallel processing of many different parts of the event at the same time
  - These kind of systems have a fixed latency – it always takes the same amount of time to come to a decision independently of the input data
    - Quite different from usual software

CMS Calo trigger algorithm card

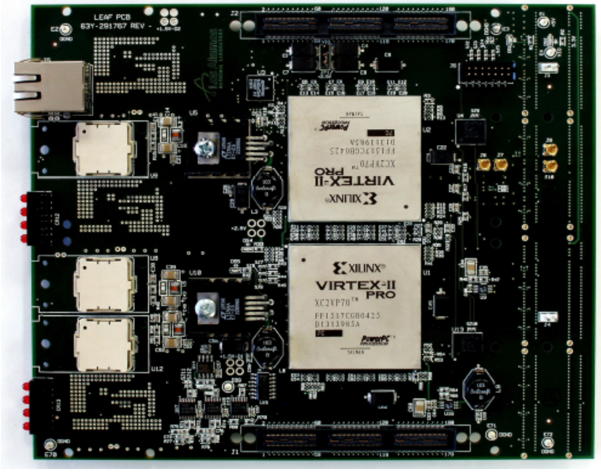

ASIC chip for ATLAS muon trigger

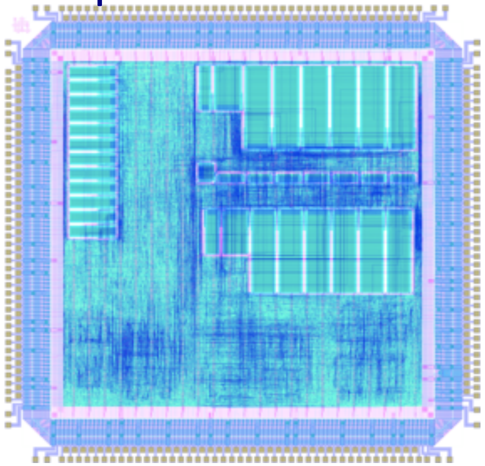

Aside – Trigger Latency

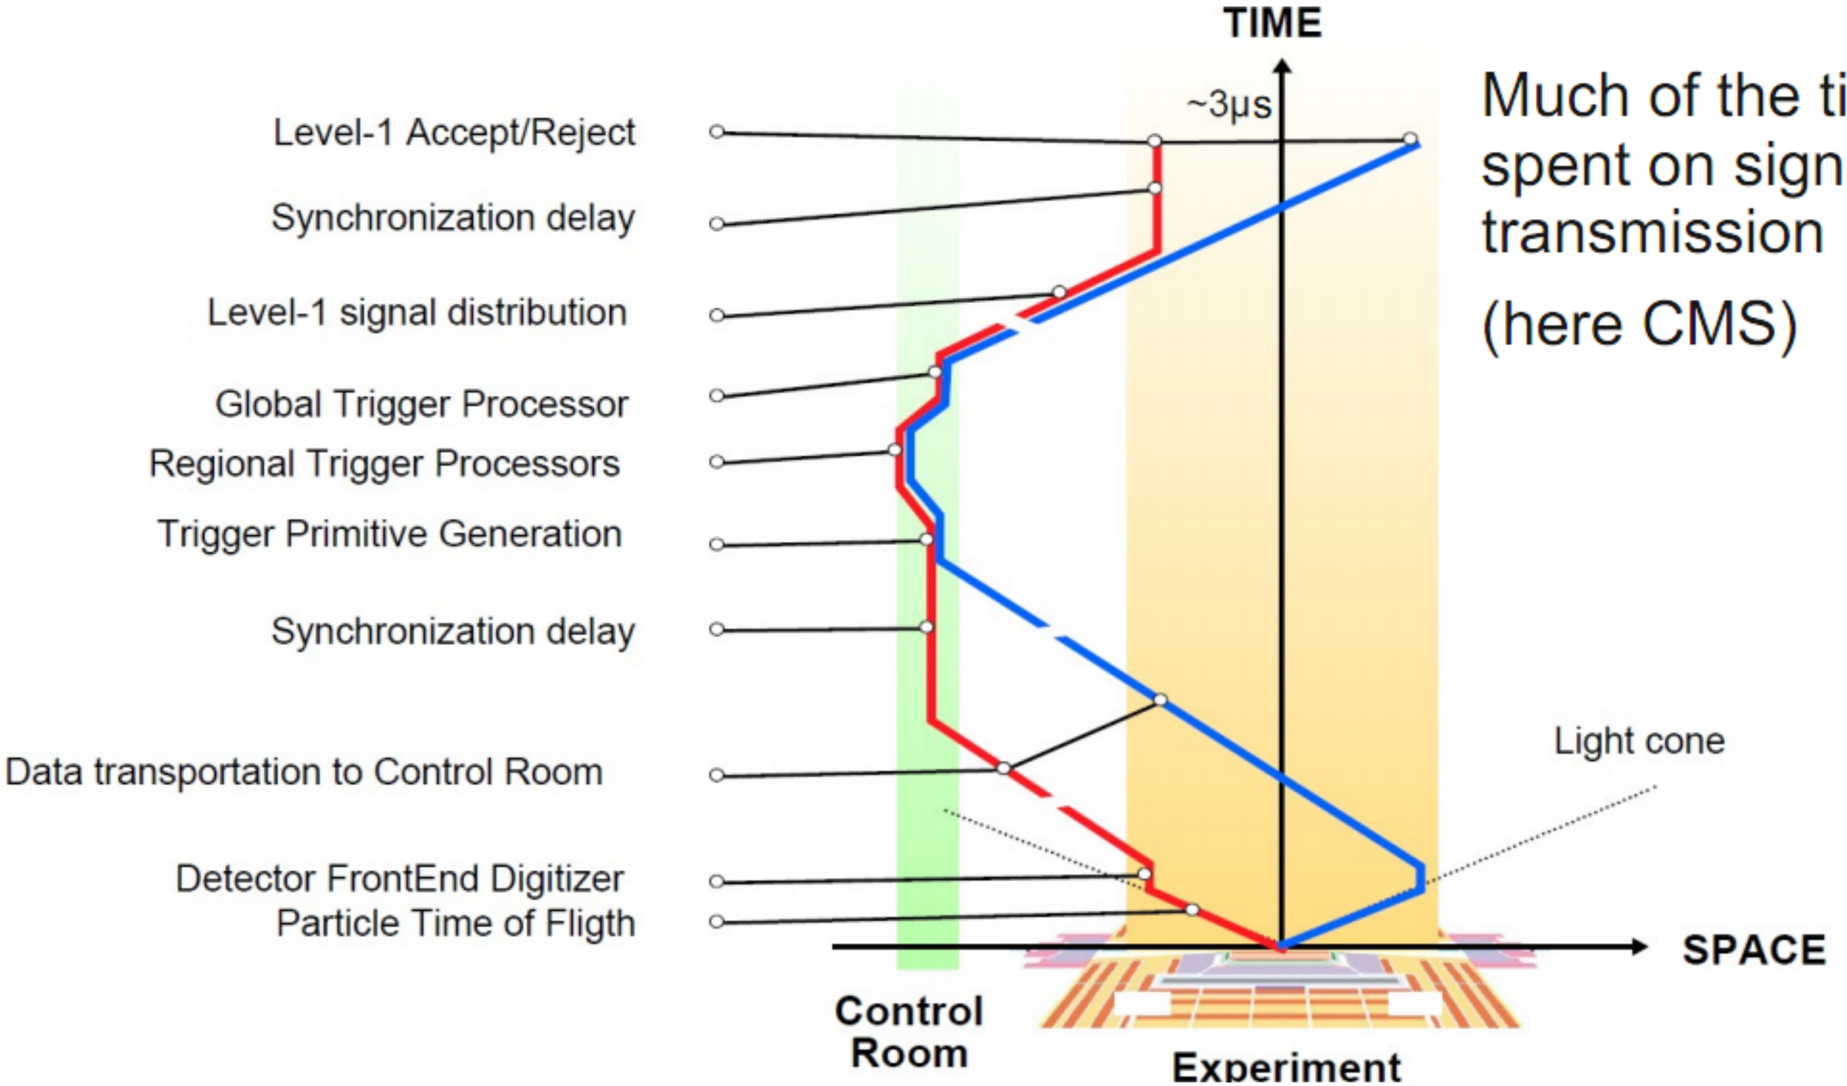

Much of the time spent on signal transmission

(here CMS)

L1 Muon Trigger (CMS)

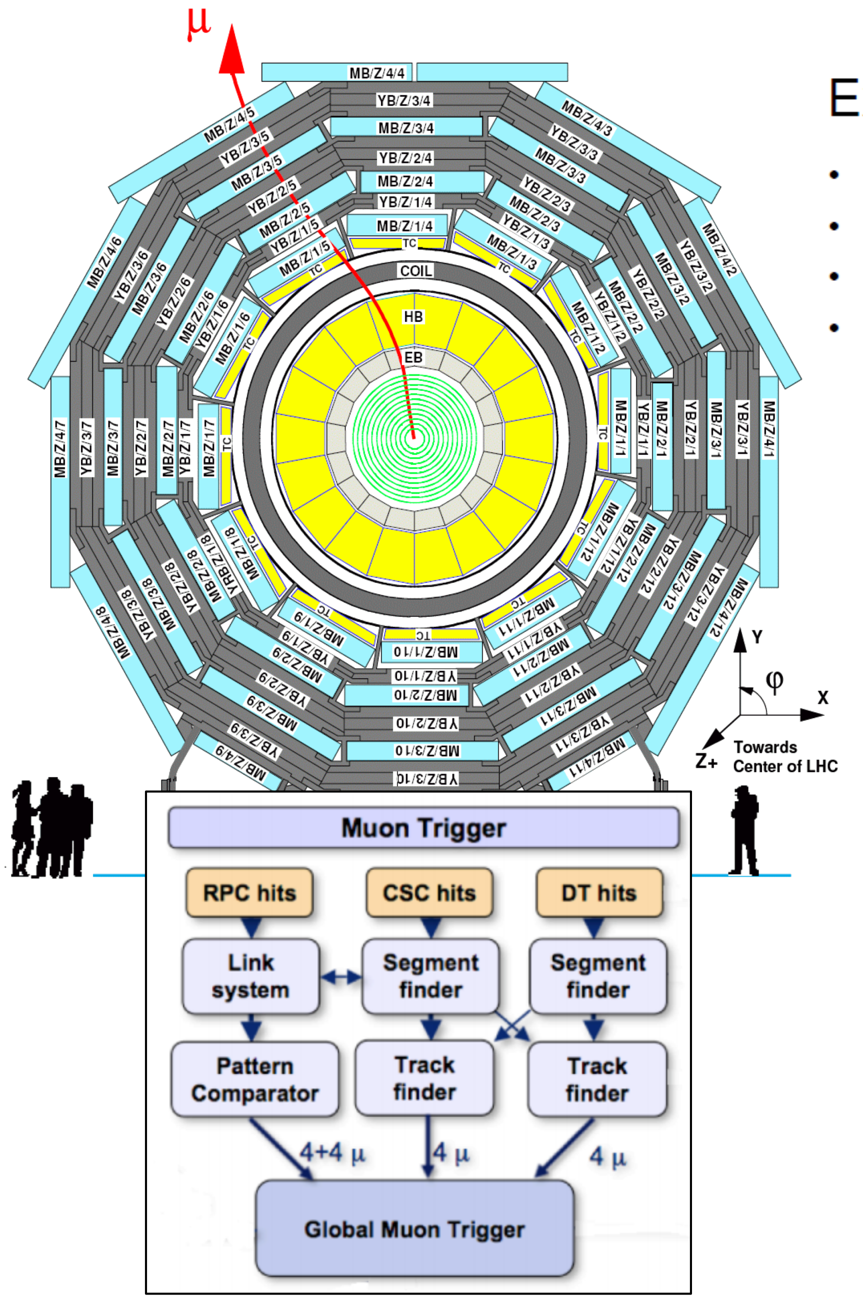

Example: trigger with drift tubes in barrel:

- Reconstruct local segments on chambers using ASICs
- Segment position, p_T and quality sent to Track Finder
- TF combines segments to form μ track using FPGAs (LUT)
- Typical p_T resolution 20%

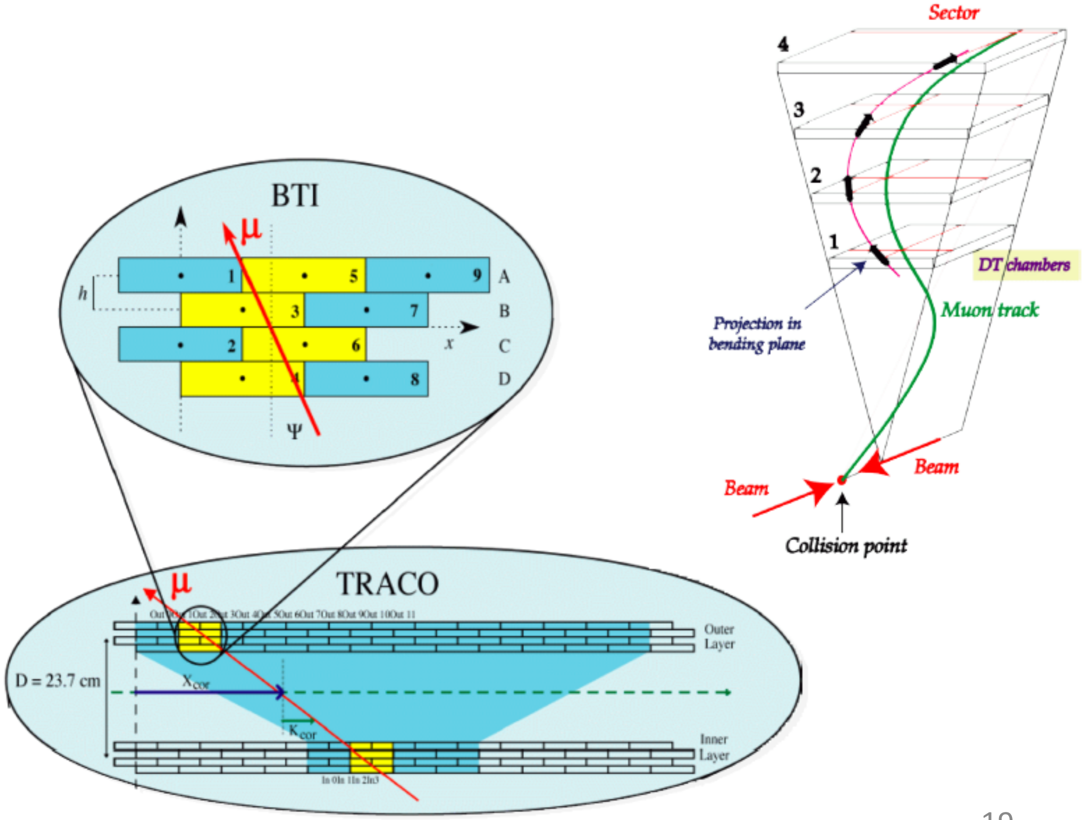

L1 Calorimeter Trigger (ATLAS)

L1 Calorimeter trigger for:

-Electrons & photons
    - Narrow (isolated) cluster in EM calorimeter only

- Jets
    - Large energy deposit in both EM and hadronic calorimeters

-Taus
    - Like a narrow Jet

-Sum Et and Missing Et
    -Summing over all towers in the calorimeter

Example – Electron/Photon trigger in ATLAS:

Cells summed to towers of 0.1x0.1 in ηxϕ
(EM and hadron towers separate)

Search in 4x4 overlapping, sliding window

Cluster local maximum in window

Possible to require isolation
(Max energy in EM ring around or
in hadronic towers behind)

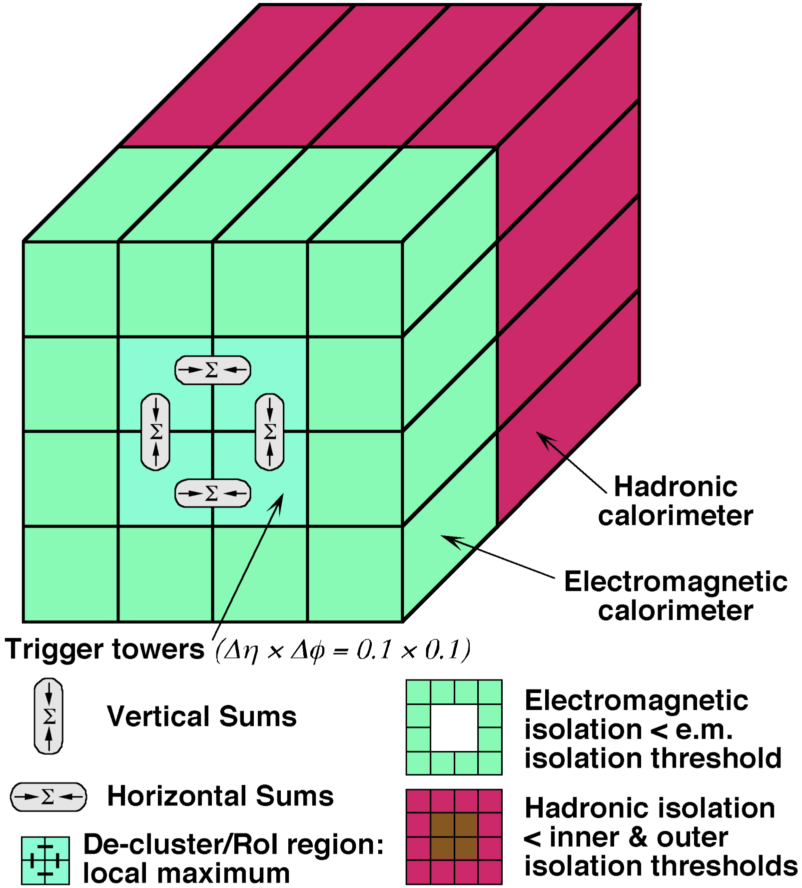

ATLAS L1 Calorimeter Trigger

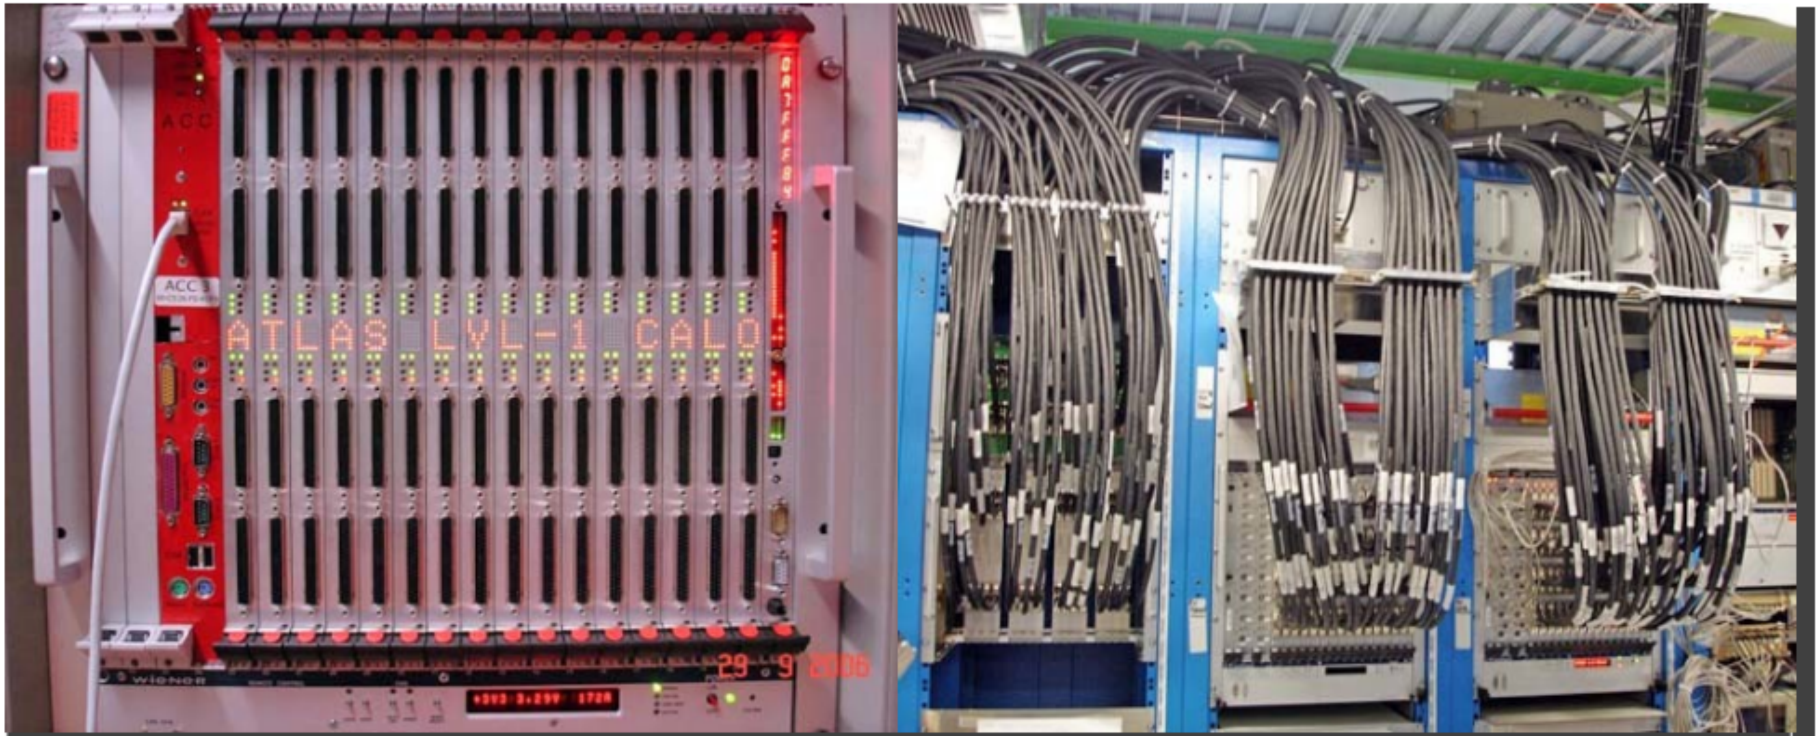

Level-1 Calorimeter Pre-processor crate

Analogue trigger cables received in electronics cavern

In total 27 VME crates in the full L1 calorimeter trigger

L1 Central Trigger

- The outputs of the L1 Muon and L1 Calorimeter triggers as well as various other L1 inputs are fed into the L1 Central (or Global) Trigger
- This has the job of
  - Time synchronization of inputs
  - Combination of inputs
  - Applying prescales
  - Applying BUSY veto (dead-time)
  - Monitoring L1 objects
  - Issuing the L1 accept
- The CTP **issues a L1 decision** (accept/reject the event) **every 25ns**

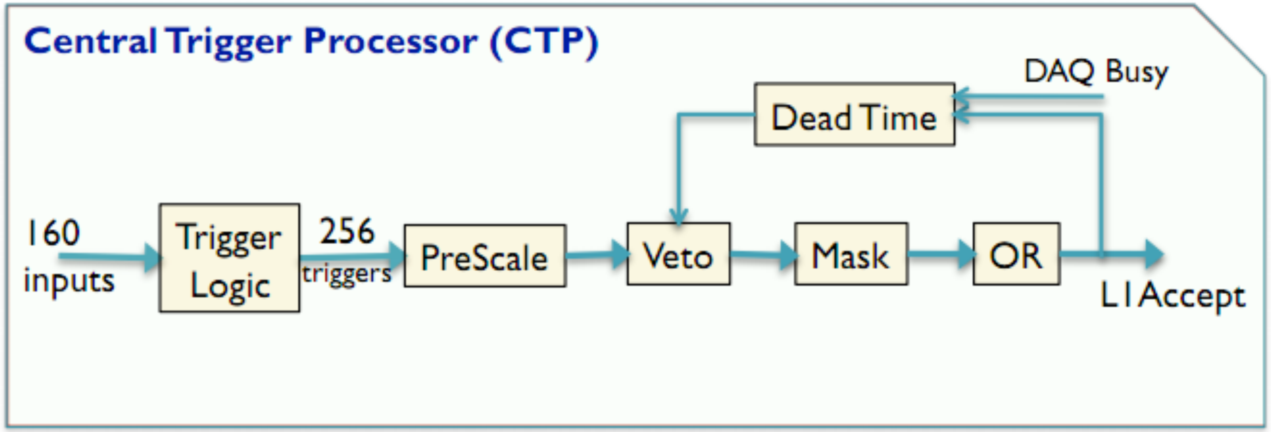

Small aside:

Dead-time in ATLAS

ATLAS has two types of trigger dead-time:

-'Simple' dead-time: after a trigger veto 5 BCs (125ns) currently leads to ~1-2% dead-time (at 50kHz L1 rate)

-'Complex' dead-time: No more then 8 triggers in 80μs to avoids bursts of triggers – currently gives ~0 dead-time

Ratio of yellow to gray is total dead-time. Average dead-time in this fill is ~2%. Spikes from short detector readout problems.

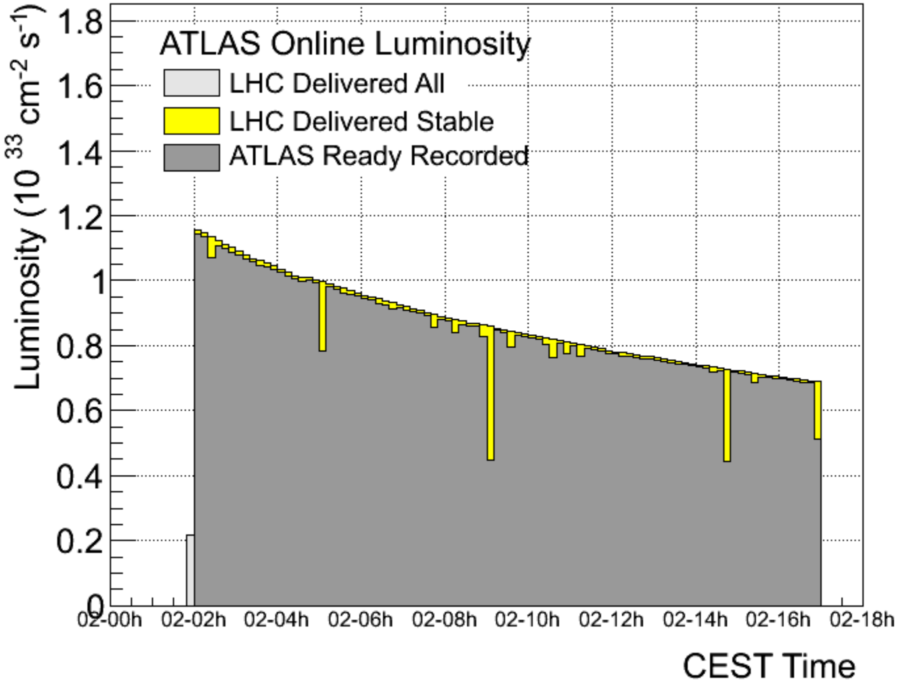

Luminosity measurement based on rate of hits in special luminosity detectors which are not in the normal DAQ – need to be corrected for dead-time.

Why no track trigger at L1?

- LHC is designed to have 25 overlapping pp interactions per bunch crossing
  - Currently average is ~6 interactions per crossing
- This gives a huge number of low P_T tracks from the **pileup interactions**
- L1 track trigger very challenging
  - Transmitting tracking detector data (~100M channels) out of the detector at 40MHz requires a huge amount of power and bandwidth
  - Perform the track finding in the required latency (~3μs)
- For current ATLAS/CMS physics program track trigger at L1 not needed
- For sLHC upgrade L1 tracking triggers being considered
  - Can be used for track isolation or matching track to El or Mu
    - Will especially be needed by CMS with worse standalone muon resolution
  - Very much work in progress

Z->$\mu^+\mu^-$ event in recent ATLAS data with high pileup (11 reconstructed vertices)

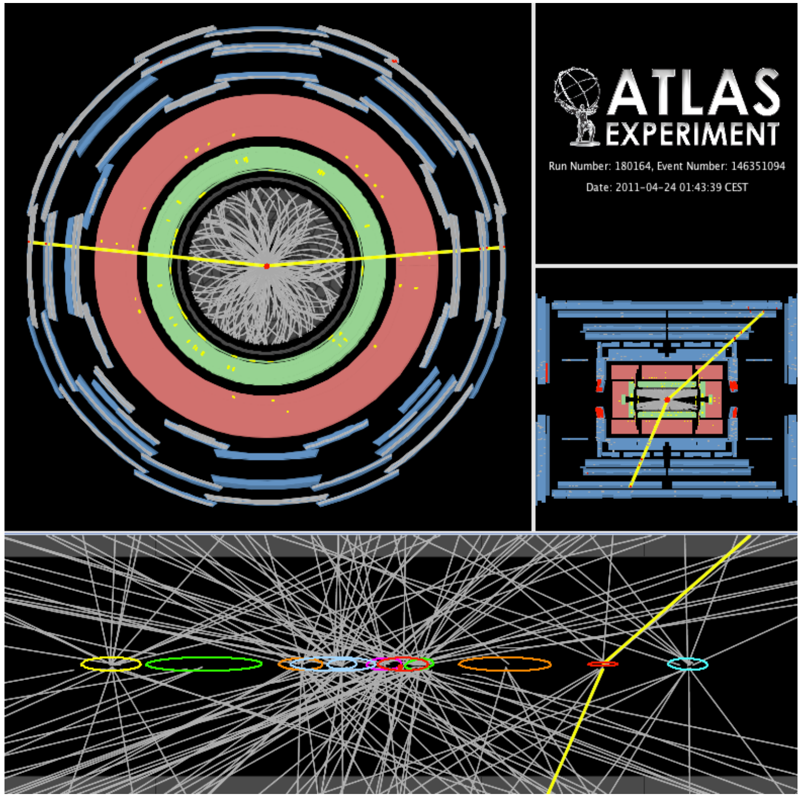

HLT Design principles

- HLT runs in software
  - Relatively slow so run on many PCs at same time
  - Very flexible
  - Software is optimized for overall speed
- Run in parallel many ‘Trigger lines’ each corresponding to a sequence of algorithms aiming at a physics signature
  - Often have several hundred Trigger lines

Example of an electron Trigger Line

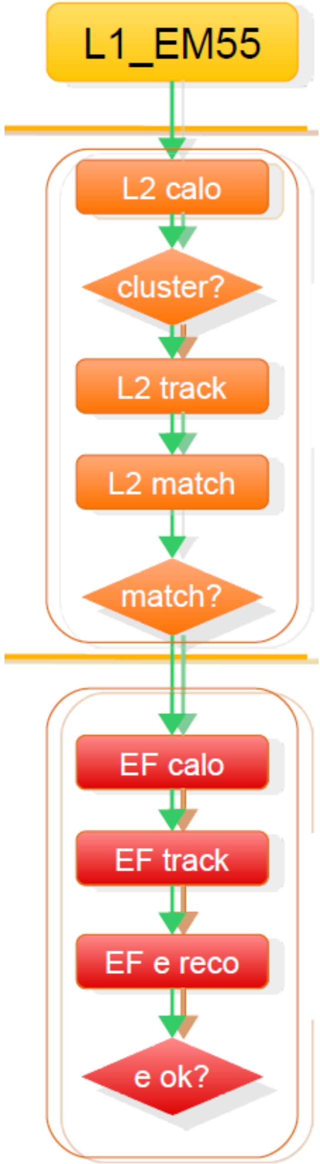

HLT Design principles

- Many tricks to speed it up!
- Regional reconstruction
- Sequential selection
  - For first algorithms only use information from a subset of detectors (eg. Many events rejected before accessing ID information)
- Early rejection
  - All algorithms executed stepwise and in parallel
  - As soon as an algorithm fails stop that trigger line
  - As soon as all trigger lines in an event are stopped reject the event
  - If a trigger line is accepted we still complete the processing of other lines as this information is useful for studying multi-object triggers and evaluating trigger efficiencies (see later)

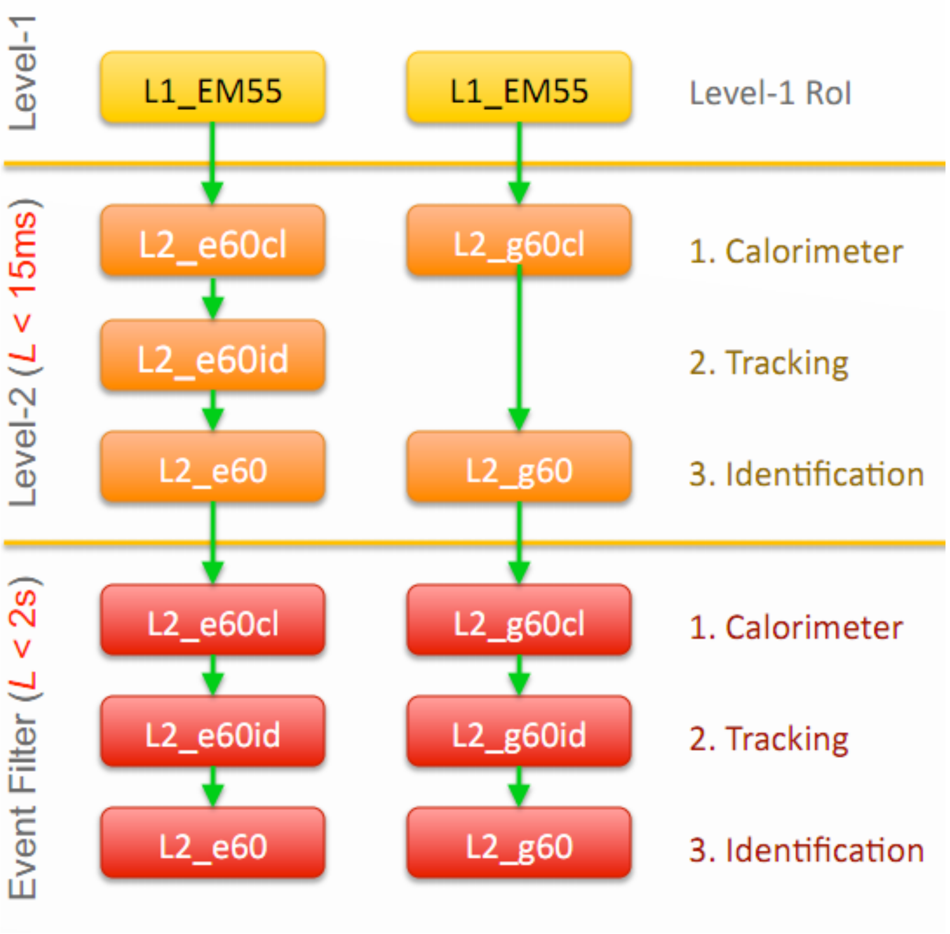

HLT Design Principles

Increasing complexity and CPU time
Decreasing number of events to run over

- First HLT algorithms often confirm L1 object with higher granularity detector information
- Then can run fast version of object reconstruction with moderate resolution
- Finally run offline like reconstruction with (near) optimal resolution

LHCb muon trigger momentum resolution

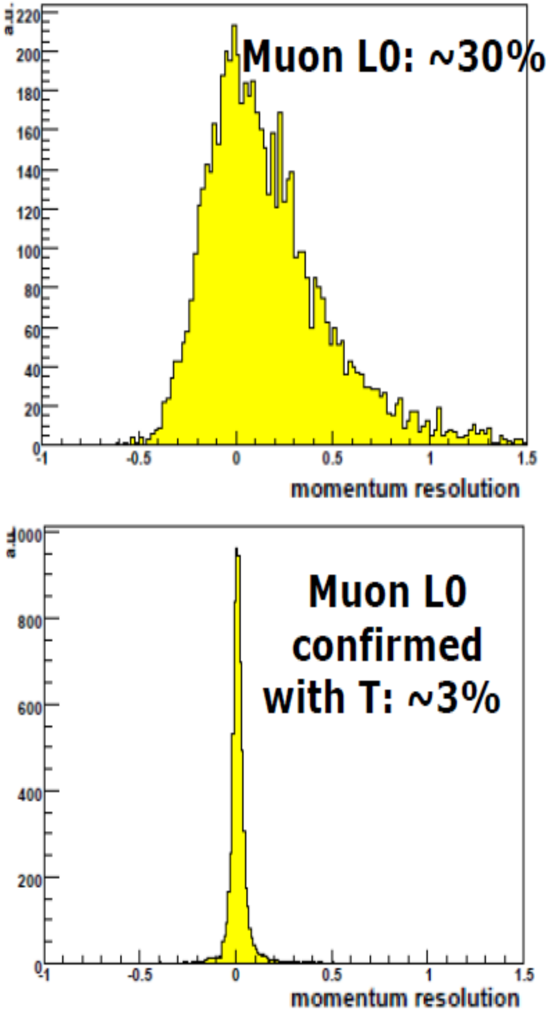

HLT Design Principles

Increasing complexity and CPU time
Decreasing number of events to run over

- First HLT algorithms often confirm L1 object with higher granularity detector information
- Then can run fast version of object reconstruction with moderate resolution
- Finally run offline like reconstruction with (near) optimal resolution

E_T resolution with respect to offline for EM clusters

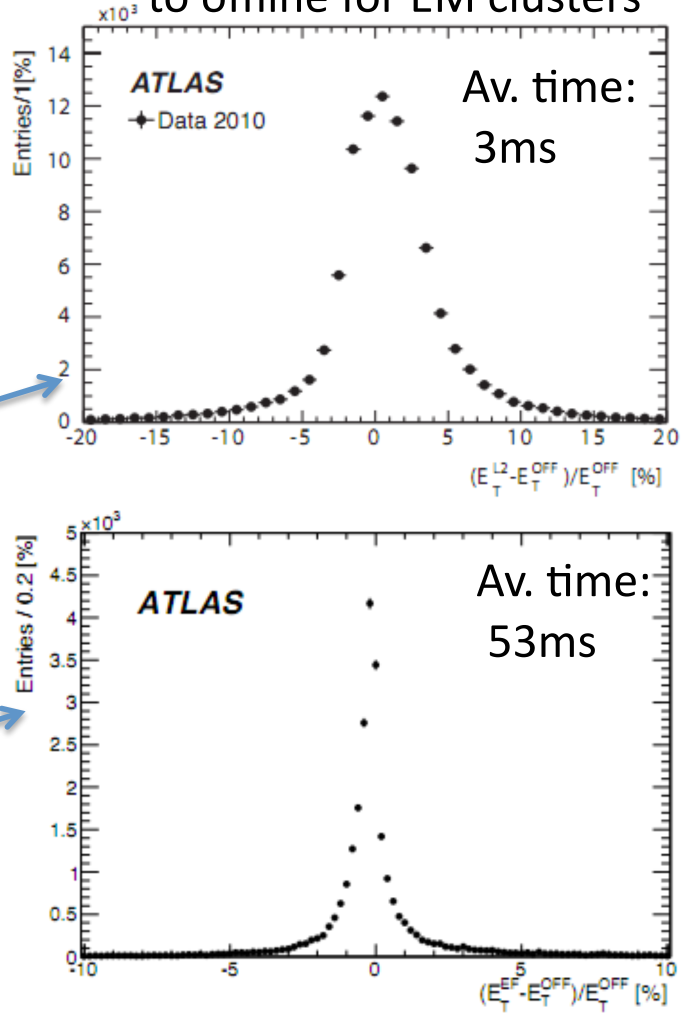

Av. time: 3ms

Av. time: 53ms

HLT Design Principles

Increasing complexity and CPU time
Decreasing number of events to run over

- First HLT algorithms often confirm L1 object with higher granularity detector information
- Then can run fast version of object reconstruction with moderate resolution
- Finally run offline like reconstruction with (near) optimal resolution

Total timing for 2 levels in ATLAS HLT

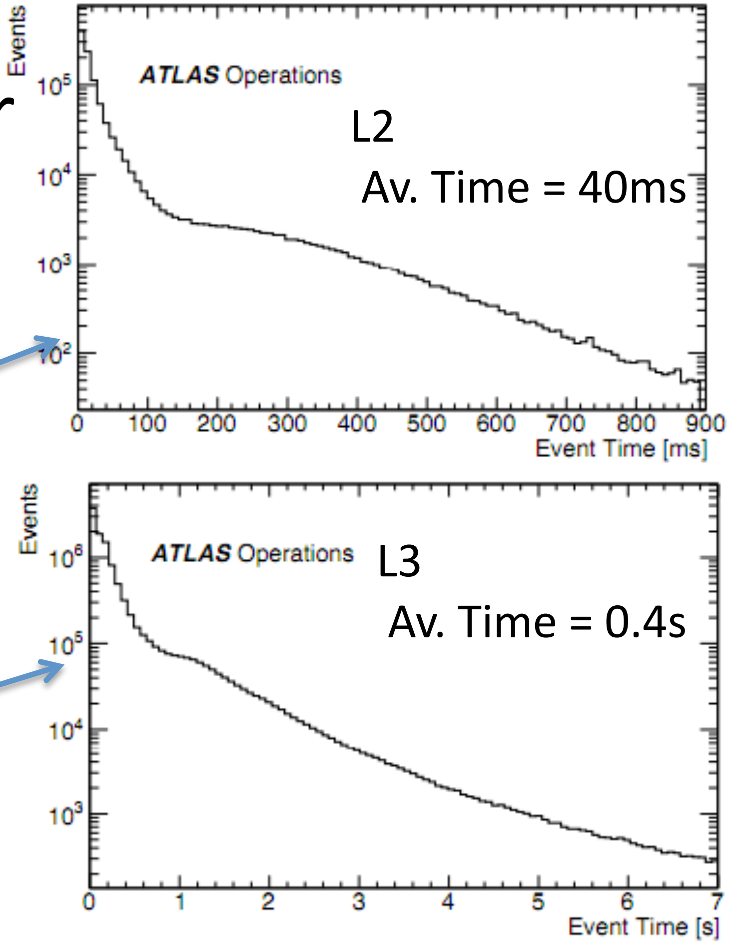

Electrons & Photons in the HLT

CMS E/gamma HLT workflow

L2 Step

- Spatial matching of ECAL clusters with e/γ candidates at L1
- Superclusters are formed
- ET cut applied
- Calorimetric (ECAL+HCAL) isolation

L3 Photons

- Tight track isolation

L3 Electrons

- Electron track reconstruction
  - Spatial matching of ECAL cluster and pixel track
- Loose track isolation in a "hollow" cone

ATLAS has a similar workflow. Amongst other thing cut on the shower shape variables:

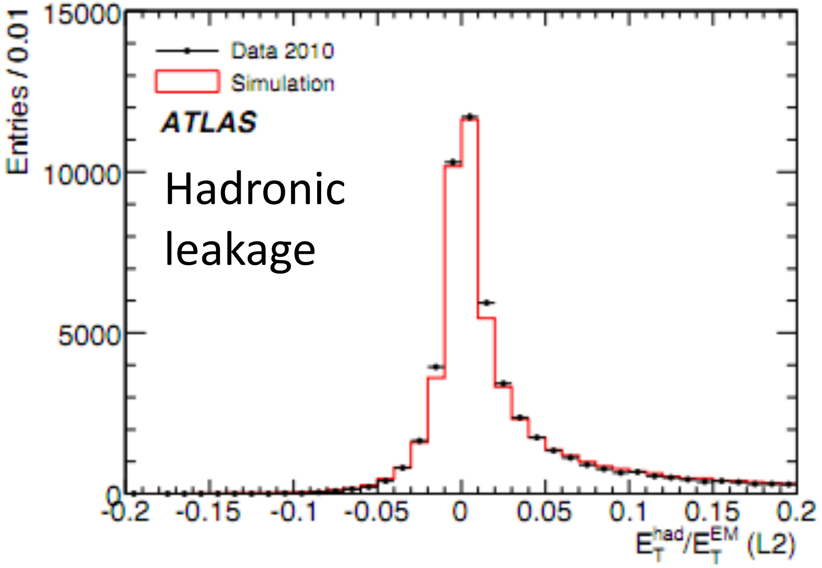

Hadronic leakage

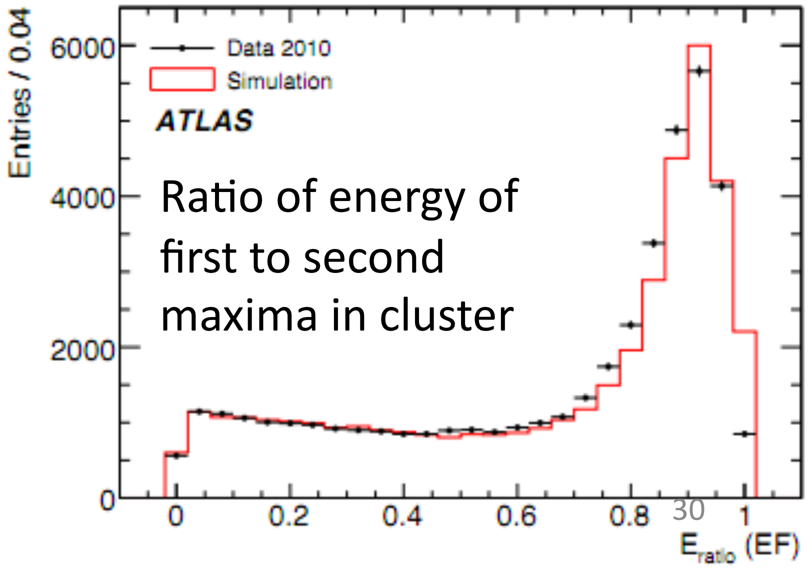

Ratio of energy of first to second maxima in cluster

Muons at HLT

CMS Muon HLT workflow

L2 Step

- Re-fit hits in the muon chambers with the full granularity
- Reco in L1 region of interest

L3 Step

- Combine tracker hits with L2 objects:
  - **Include tracker hits**
  - **Matching tracker and muon informations**
  - **Better transverse momentum measurement**

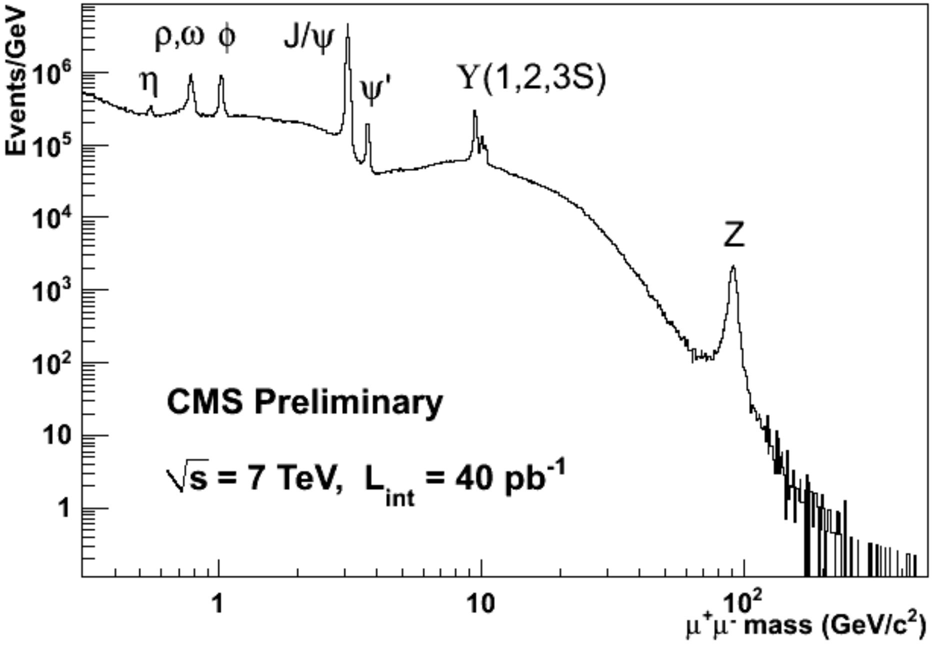

Muons at HLT

ATLAS Muon HLT workflow

High Level Trigger

- Level 2 uses hits in Muon Spectrometer (MS) in an RoI defined by L1.
  - Select MDT hits and fit based on MDT drift time and position information and assign pT using fast look up tables.
- Additional refinements:
  - Tracks extrapolated to IP correcting for material and alignment and associate track from Inner Detector in RoI
    - Referred to as "Combined"
  - Isolation in Calorimeter & Tracking
- Event Filter:
  - Same as L2, but fits performed on hits in MDT and Tracking instead of using LUT.

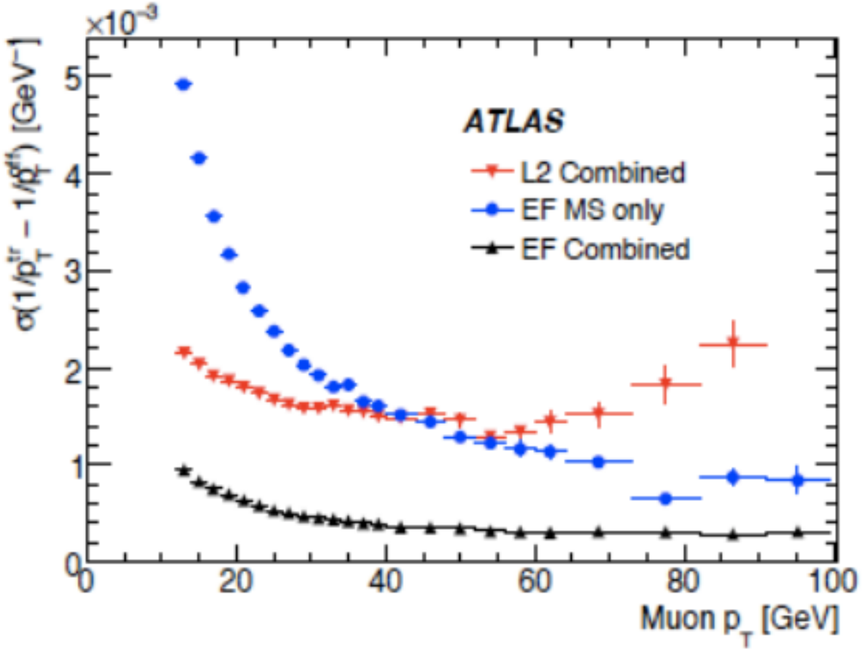

Comparing LHC expt. trigger systems

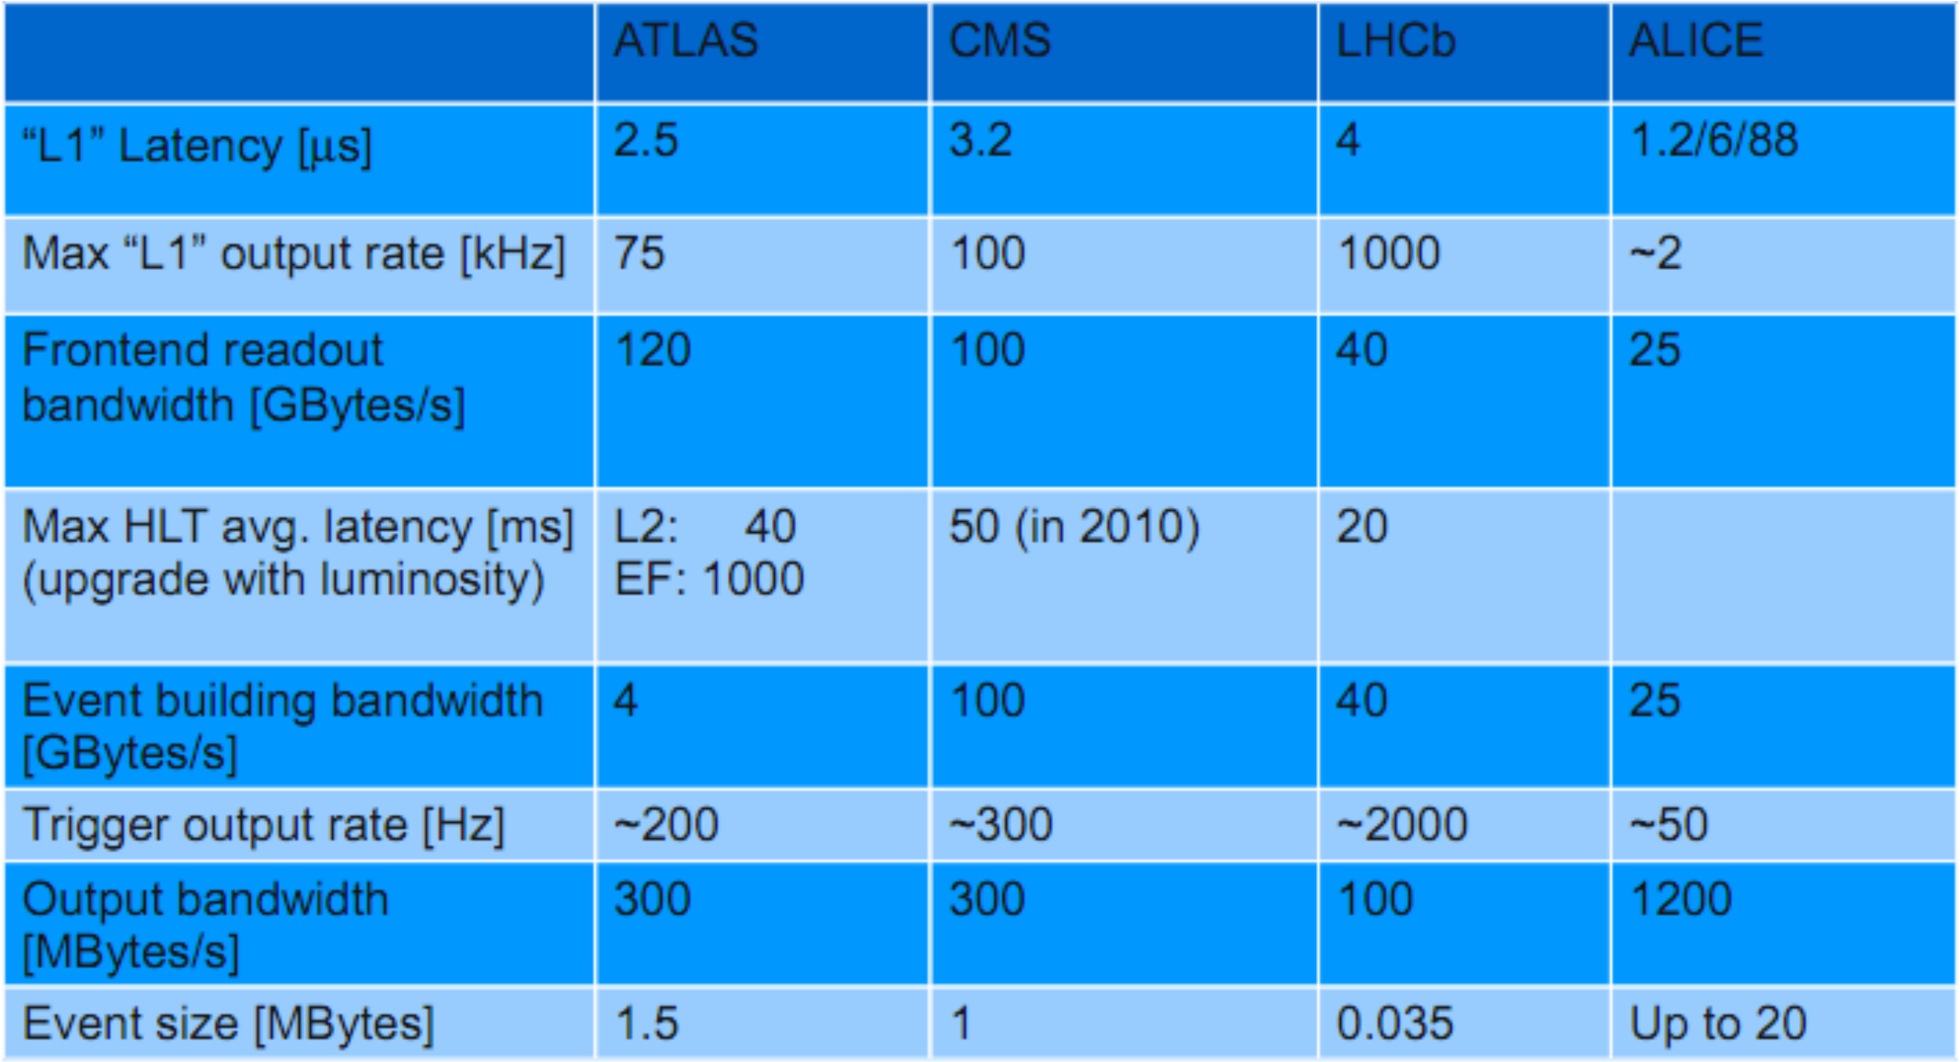

	ATLAS	CMS	LHCb	ALICE
"L1" Latency [μs]	2.5	3.2	4	1.2/6/88
Max "L1" output rate [kHz]	75	100	1000	~2
Frontend readout bandwidth [GBytes/s]	120	100	40	25
Max HLT avg. latency [ms] (upgrade with luminosity)	L2: 40 EF: 1000	50 (in 2010)	20	
Event building bandwidth [GBytes/s]	4	100	40	25
Trigger output rate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

Taken from Brian Petersen's academic training lecture.
Some numbers not directly comparable – but to give a feeling.

Comparing LHC expt. trigger systems

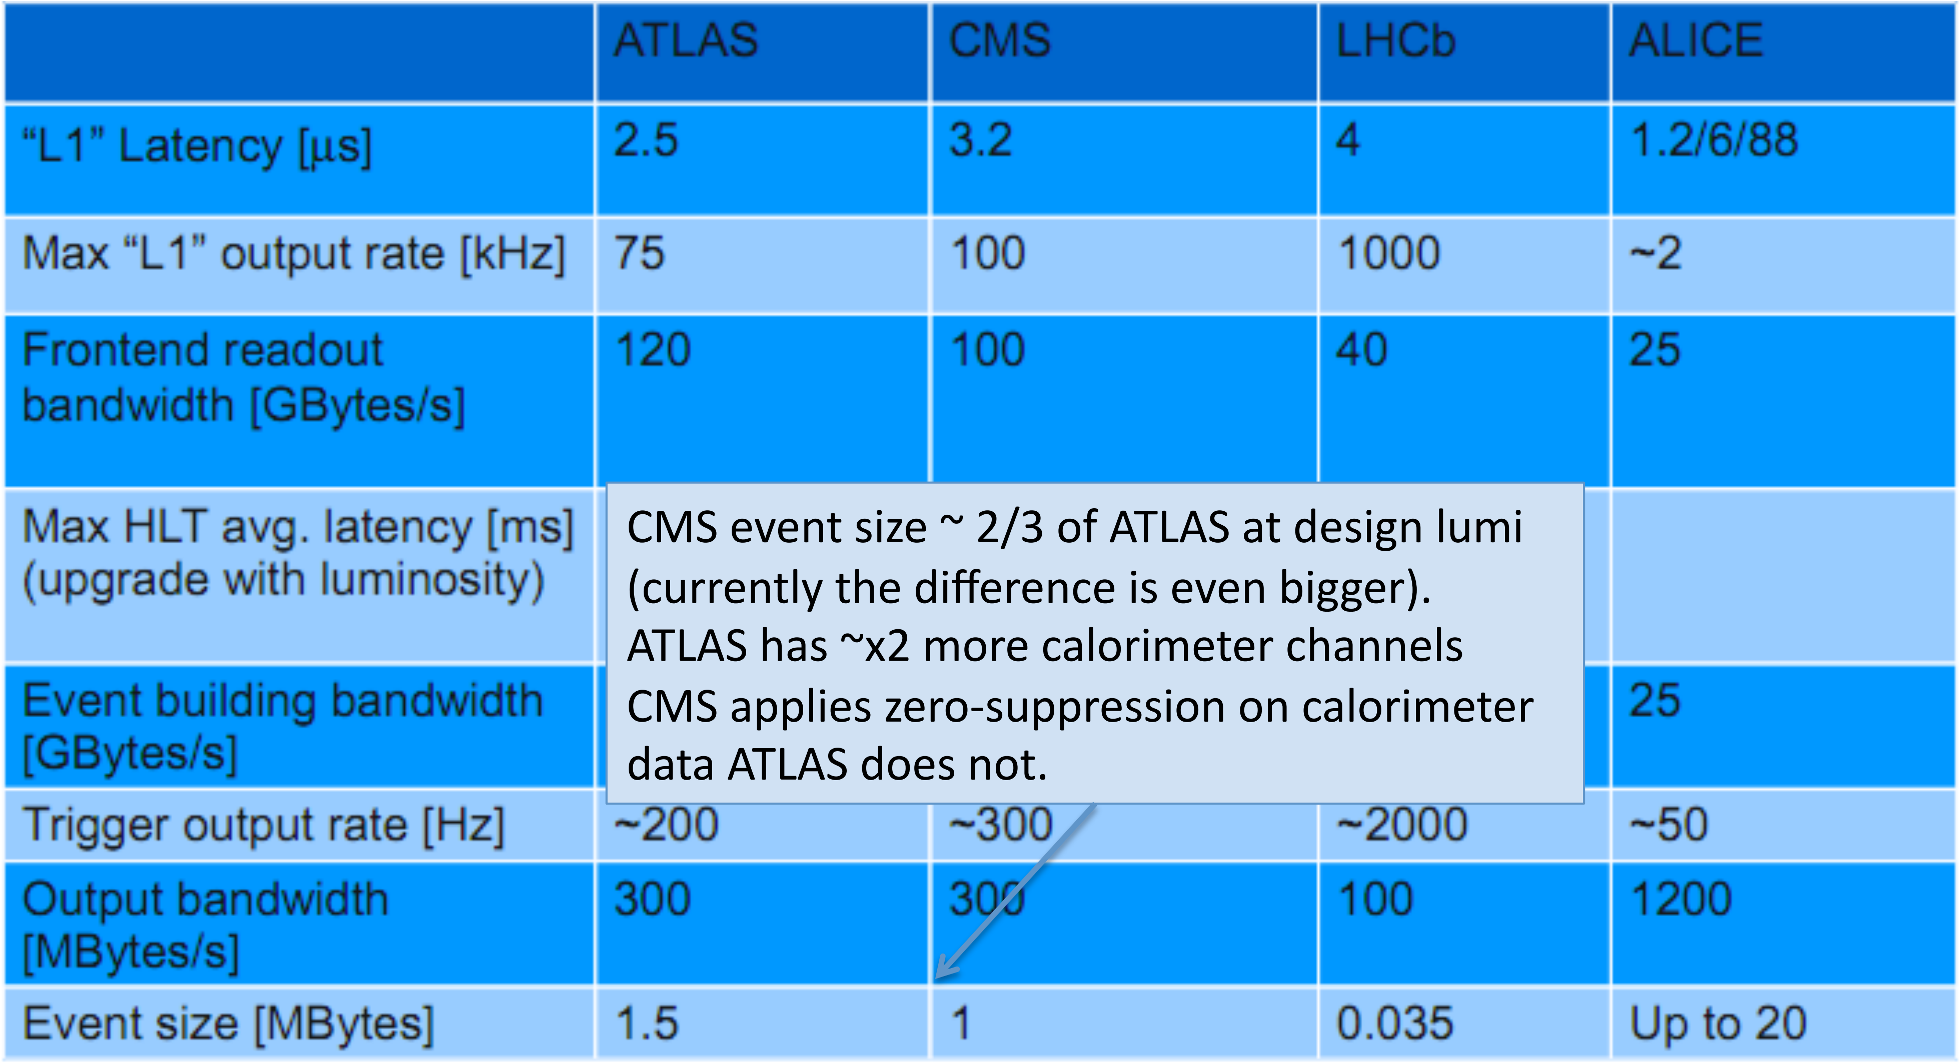

	ATLAS	CMS	LHCb	ALICE
“L1” Latency [μs]	2.5	3.2	4	1.2/6/88
Max “L1” output rate [kHz]	75	100	1000	~2
Frontend readout bandwidth [GBytes/s]	120	100	40	25
Max HLT avg. latency [ms] (upgrade with luminosity)				
Event building bandwidth [GBytes/s]				25
Trigger output rate [Hz]	~200	~300	~2000	~50
Output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

CMS event size ~ 2/3 of ATLAS at design lumi (currently the difference is even bigger).
ATLAS has ~x2 more calorimeter channels
CMS applies zero-suppression on calorimeter data ATLAS does not.

Taken from Brian Petersen’s academic training lecture.
Some numbers not directly comparable – but to give a feeling.

DAQ comparisons

LHC experiments pushing the DAQ limits by an order of magnitude compared to previous experiments. LHCb and ALICE have quite different requirements.

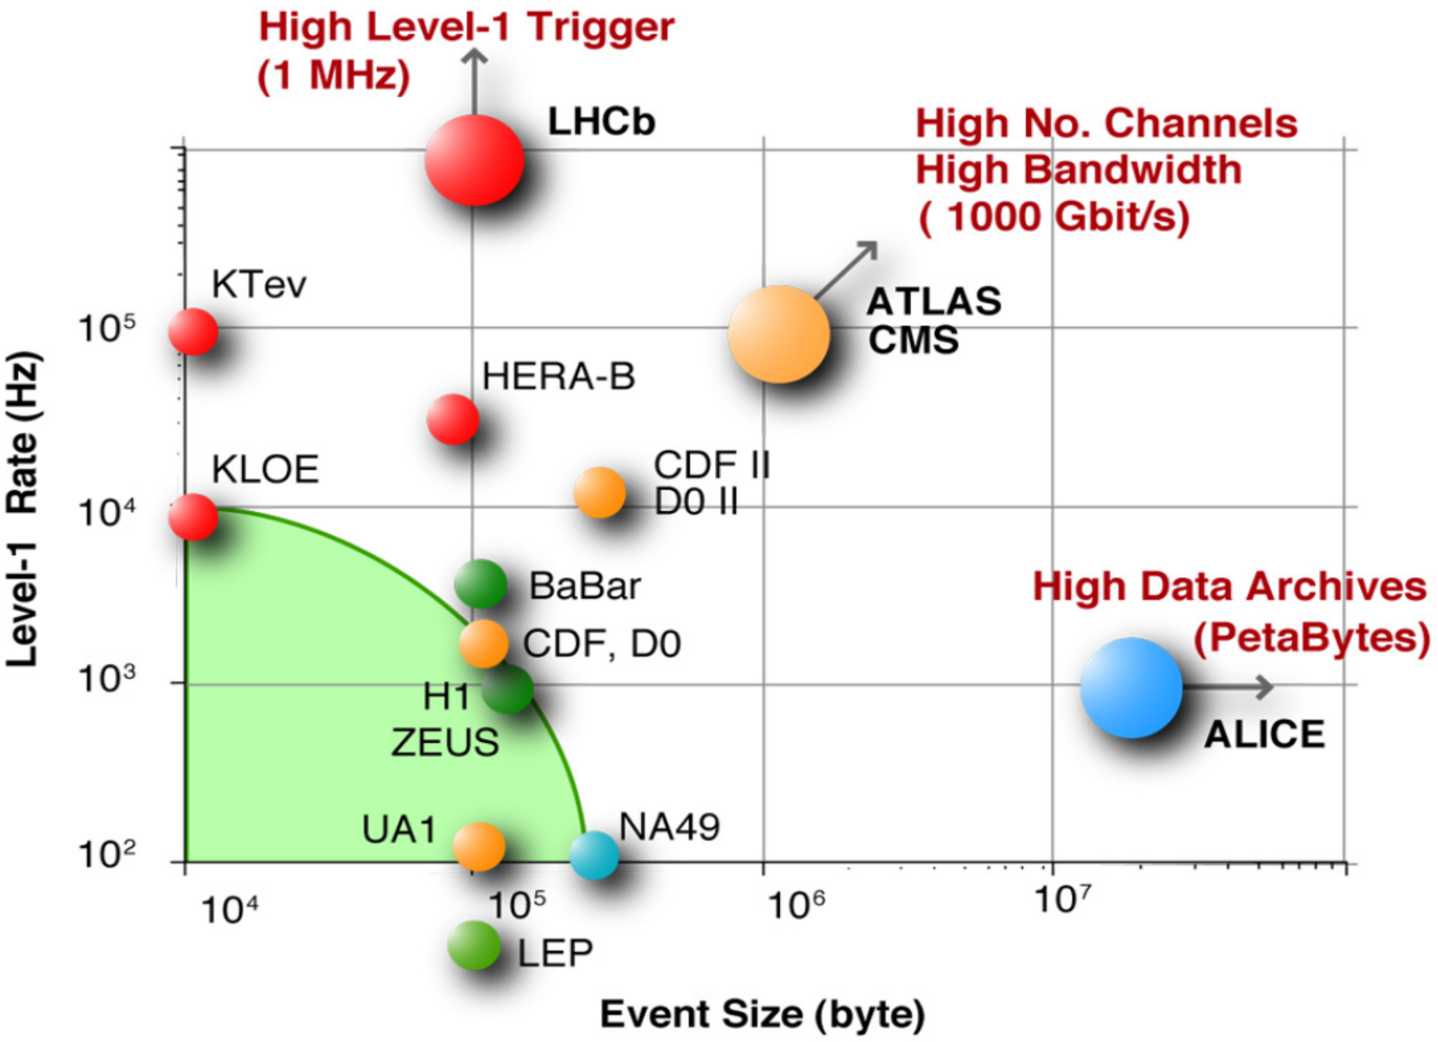

Trigger operations, menus & analysis

SDX1 | 2nd floor | Rows 3 & 2

Trigger Efficiency determination

- To measure a cross-section need to know the efficiency of the trigger to select your events
- Usually parameterize object trigger efficiency as function of offline reconstructed P_T , η
  - "Turn-on curve"

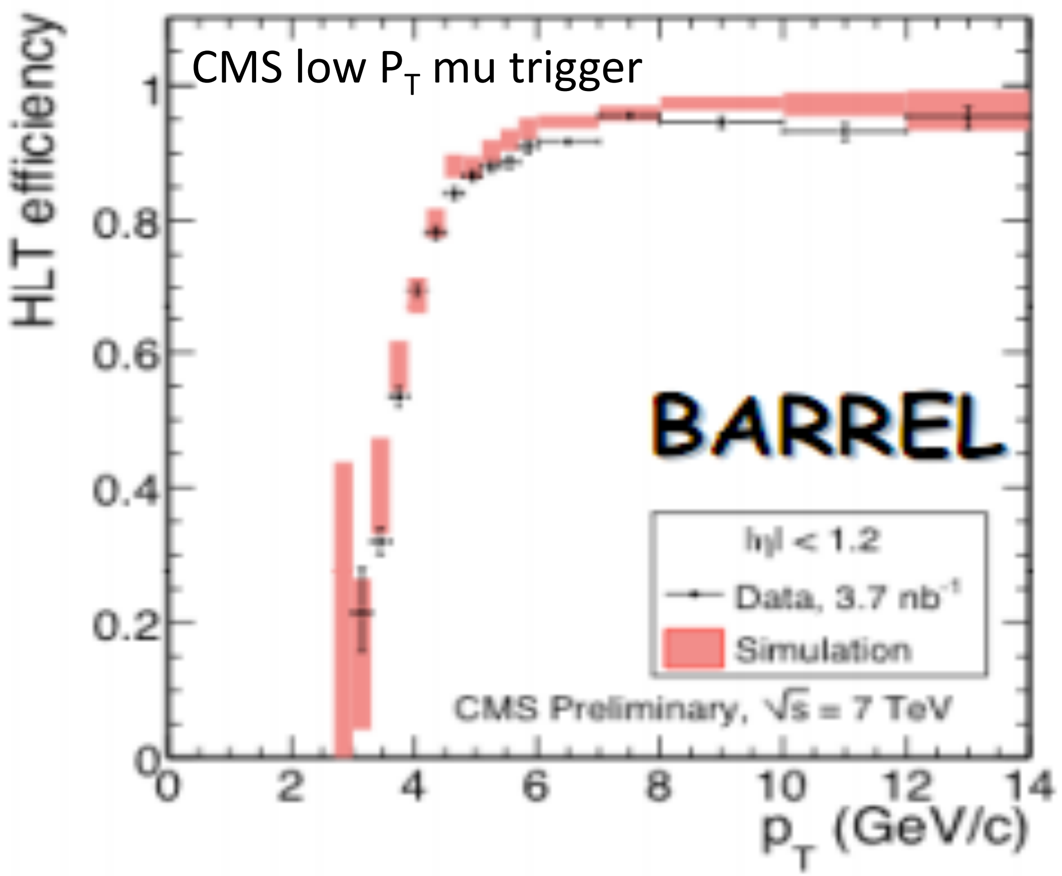

Want the turn-on curve to be as sharp as possible.
This allows you to use more of the triggered events in physics analysis without worrying about a varying trigger efficiency.
It is not sharp because you are cutting at different trigger levels on quantities with different resolutions, different cut values etc…
Different objects have different turn-on curves – Jets/MET usually worse.

Trigger Efficiency determination

- Ways to measure the efficiency
  - Tag & Probe (using eg. Z, W, J/ψ)
  - Using orthogonal or passthrough triggers
  - Boot strapping from (prescaled) lower threshold triggers (only gives relative efficiency)
  - Simulation – better to use data-driven method if possible

Data-driven efficiency measurement correctly takes into account time varying detector/ trigger problems if on same data sample as analysis

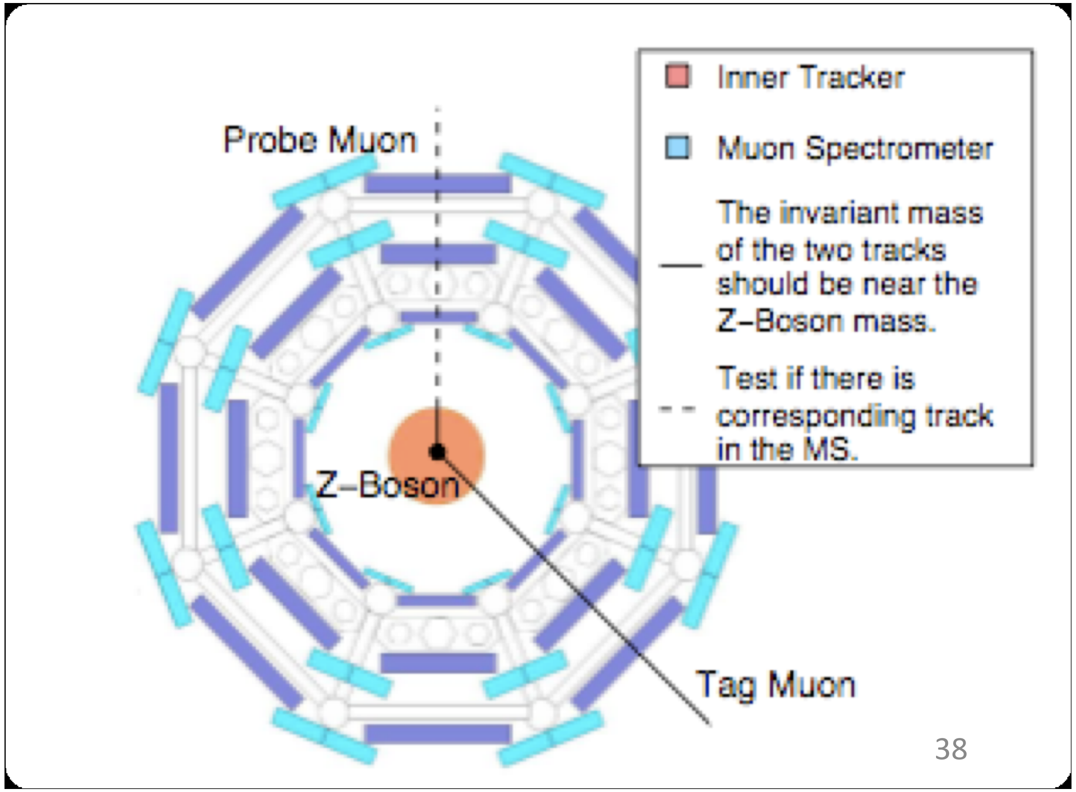

More turn-on curves

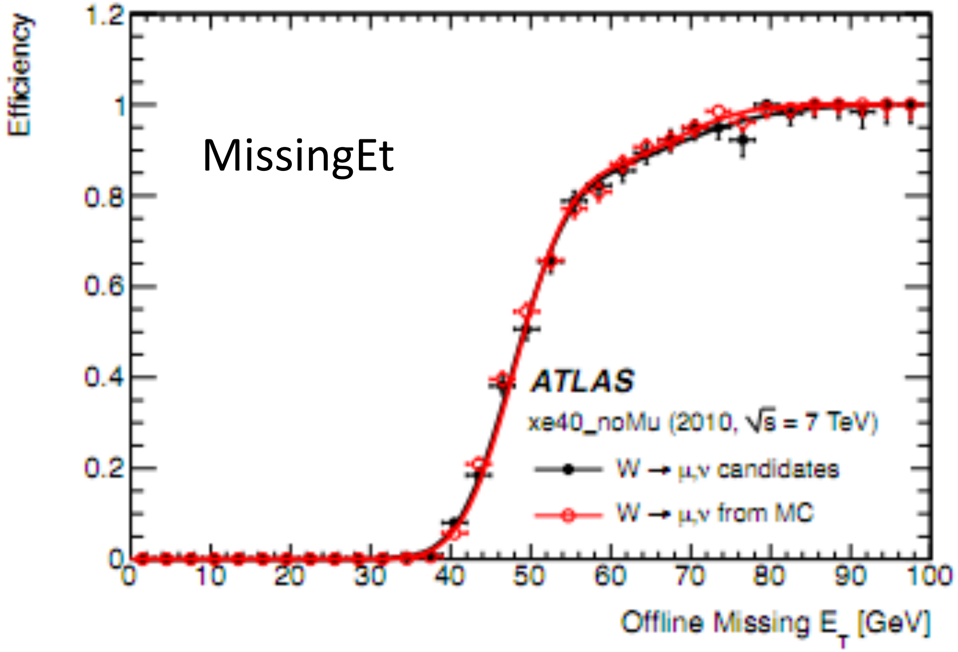
MissingEt

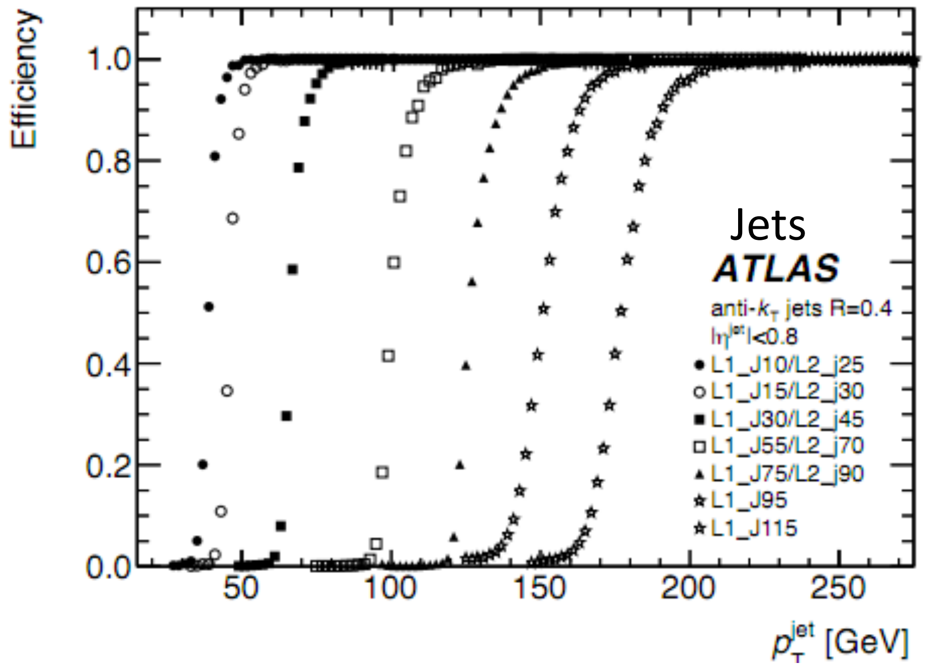
Jets

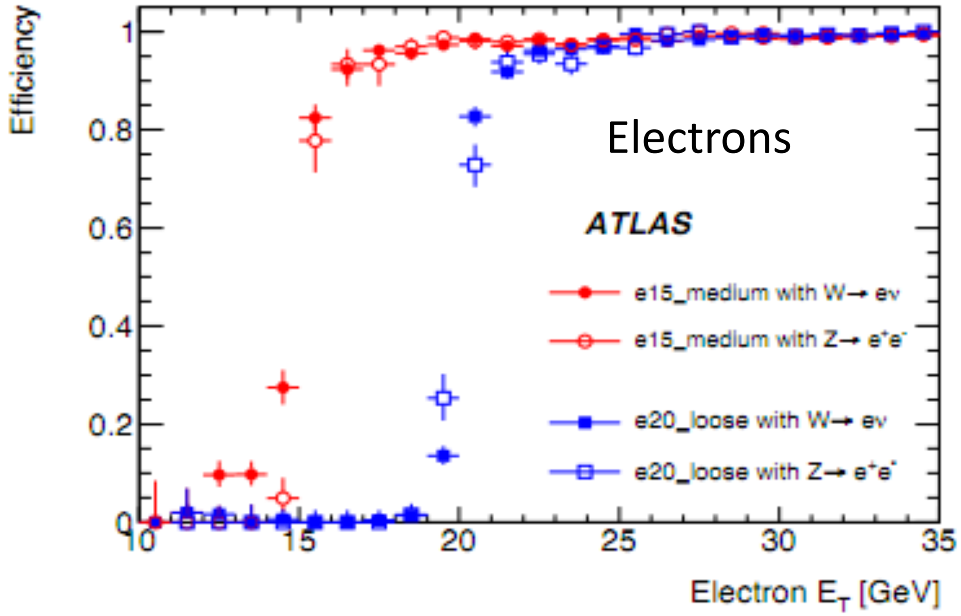
Electrons

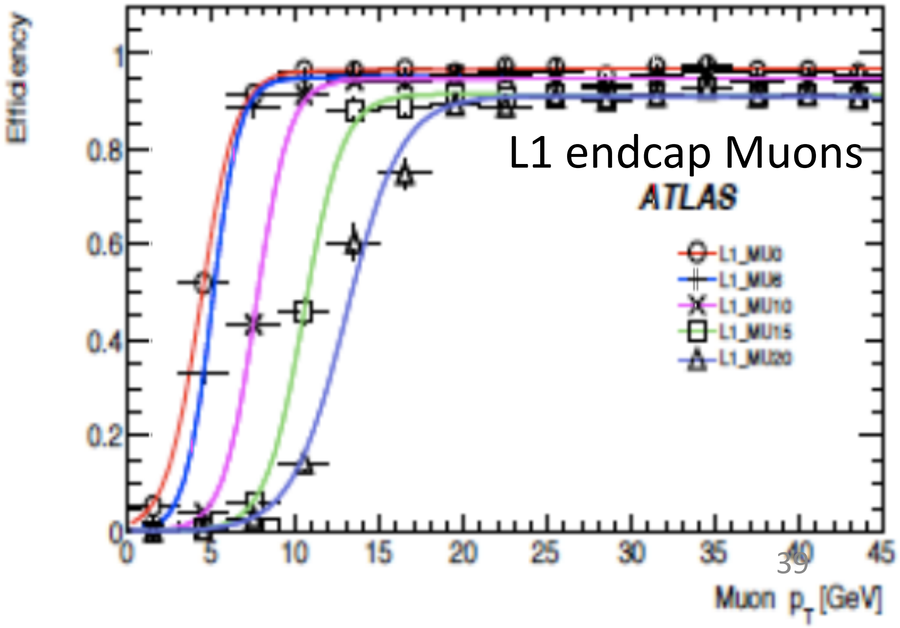
L1 endcap Muons

Prescales & Passthroughs

Not all triggers need to be recorded at full rate so apply **prescale**
- Prescale of 10 means we accept 1 event in 10 of the triggered events
- Often want to just sample low E_T events (for calibrations/efficiency measurements)
- Some triggers might just be too high rate
  - B-physics triggers often need to be prescaled (need to take into account prescale in luminosity determination)

(ATLAS HLT also allows fractional prescales eg. pass 2 events in 3)

For efficiency measurements and trigger monitoring run with a low rate of **passthrough** triggers
- Run full trigger on these events but **record them independently of the trigger decision**
  - Random triggers passthrough L1
  - Standard L1 triggers with passthrough at the HLT

Many different passthrough triggers running in ATLAS trigger configuration with low rate (<1Hz each)

Trigger information in the Offline & Simulation

- In ATLAS store the following trigger information in the event data
  - Trigger decision: 256 L1 bits + a few hundred HLT decision bits
    - Before and after prescale and busy veto
  - Trigger features (eg. L1+HLT trigger object information)
    - Can take up a lot of space in the outputs of reconstruction
- Trigger objects needed for:
  - Monitoring the trigger is working as expected
  - In analysis often want to match offline analysis object to the triggered object in order to apply trigger efficiency correction
  - Useful for data driven trigger efficiency studies
  - Needed for designing new triggers
- Very useful to be able to **re-run the HLT offline** with updated configuration
- Also need trigger in simulation
  - Difficult to simulate all trigger configurations used – can make analysis difficult
  - Store event and trigger data independent of simulated trigger decision
  - **Accurate L1 simulation crucial** as this cannot be run offline

Trigger Menu's

- **Dividing up** the available **trigger bandwidth** between the different possible triggers is an extremely difficult task
- The trigger menu (or table) says which triggers are taken with which (if any) prescales
- The trigger menu needs to **respect the rate restrictions** of the system
  - Not just the final output rate, but internal limits of the architecture (L1, EventBuilding etc..)
- The menu doesn't just contain physics triggers - divided into
  - Primary triggers for physics analysis – normally unprescaled
  - Supporting triggers for efficiency measurements and performance studies – run with fixed bandwidth
  - Calibration triggers for calibration and alignment

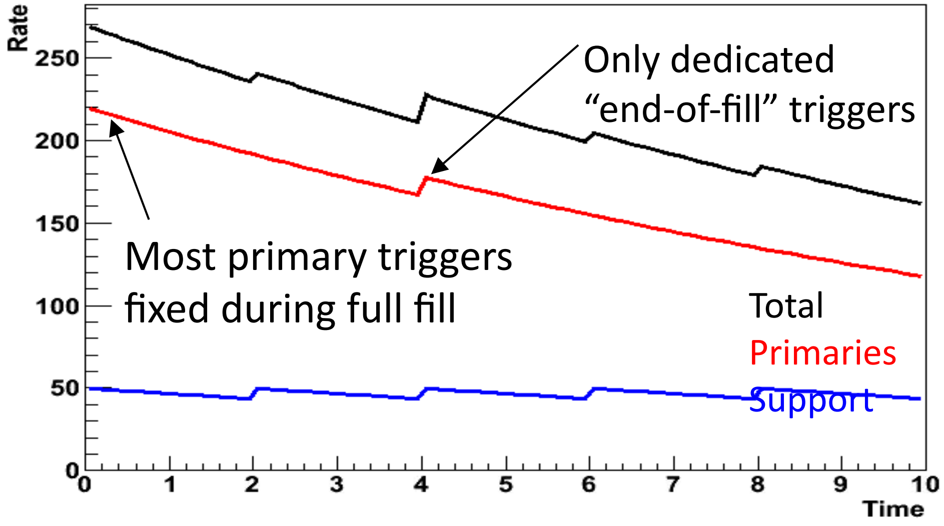

Only dedicated "end-of-fill" triggers

Most primary triggers fixed during full fill

Trigger Menu’s

- **Dividing up** the available **trigger bandwidth** between the different possible triggers is an extremely difficult task
- The trigger menu (or table) says which triggers are taken with which (if any) prescales
- The trigger menu needs to **respect the rate restrictions** of the system
  - Not just the final output rate, but internal limits of the architecture (L1, EventBuilding etc..)
- The menu doesn’t just contain physics triggers - divided into
  - Primary triggers for physics analysis – normally unprescaled
  - Supporting triggers for efficiency measurements and performance studies – run with fixed bandwidth
  - Calibration triggers for calibration and alignment

ATLAS 10^{33} primary triggers break down like:

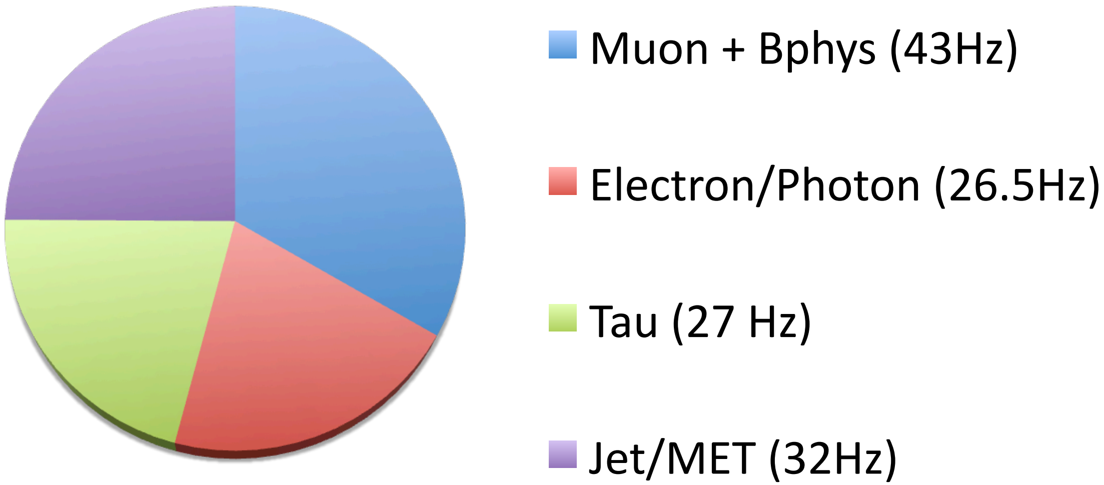

Example ATLAS trigger menu for L=10^{33}

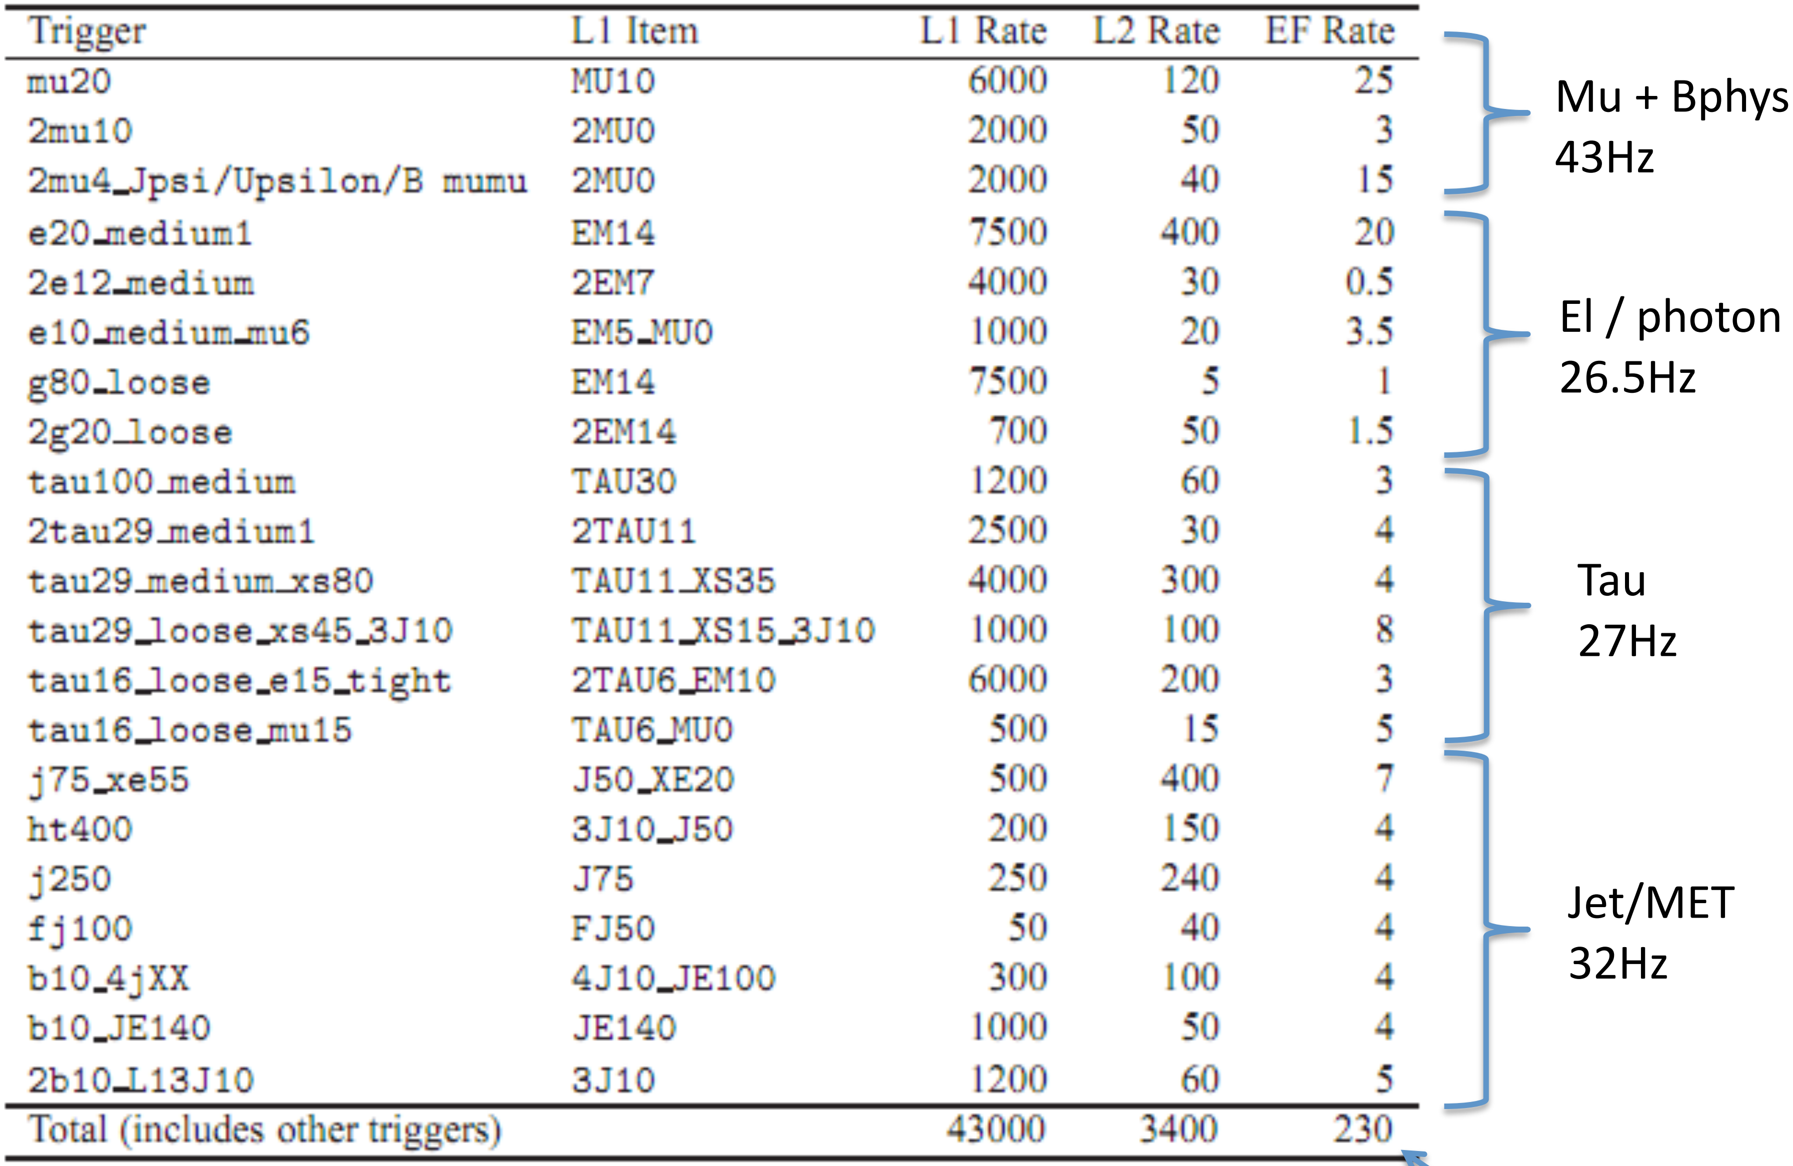

Trigger	L1 Item	L1 Rate	L2 Rate	EF Rate
mu20	MU10	6000	120	25
2mu10	2MU0	2000	50	3
2mu4_Jpsi/Upsilon/B mumu	2MU0	2000	40	15
e20_medium1	EM14	7500	400	20
2e12_medium	2EM7	4000	30	0.5
e10_medium_mu6	EM5_MU0	1000	20	3.5
g80_loose	EM14	7500	5	1
2g20_loose	2EM14	700	50	1.5
tau100_medium	TAU30	1200	60	3
2tau29_medium1	2TAU11	2500	30	4
tau29_medium_xs80	TAU11_XS35	4000	300	4
tau29_loose_xs45_3J10	TAU11_XS15_3J10	1000	100	8
tau16_loose_e15_tight	2TAU6_EM10	6000	200	3
tau16_loose_mu15	TAU6_MU0	500	15	5
j75_xe55	J50_XE20	500	400	7
ht400	3J10_J50	200	150	4
j250	J75	250	240	4
fj100	FJ50	50	40	4
b10_4jXX	4J10_JE100	300	100	4
b10_JE140	JE140	1000	50	4
2b10_L13J10	3J10	1200	60	5
Total (includes other triggers)		43000	3400	230

Mu + Bphys
43Hz

El / photon
26.5Hz

Tau
27Hz

Jet/MET
32Hz

Totals need to take into account overlap between different triggers

Menu from ~April 2011.
Not all triggers shown.

Example ATLAS trigger menu for L=10^{33}

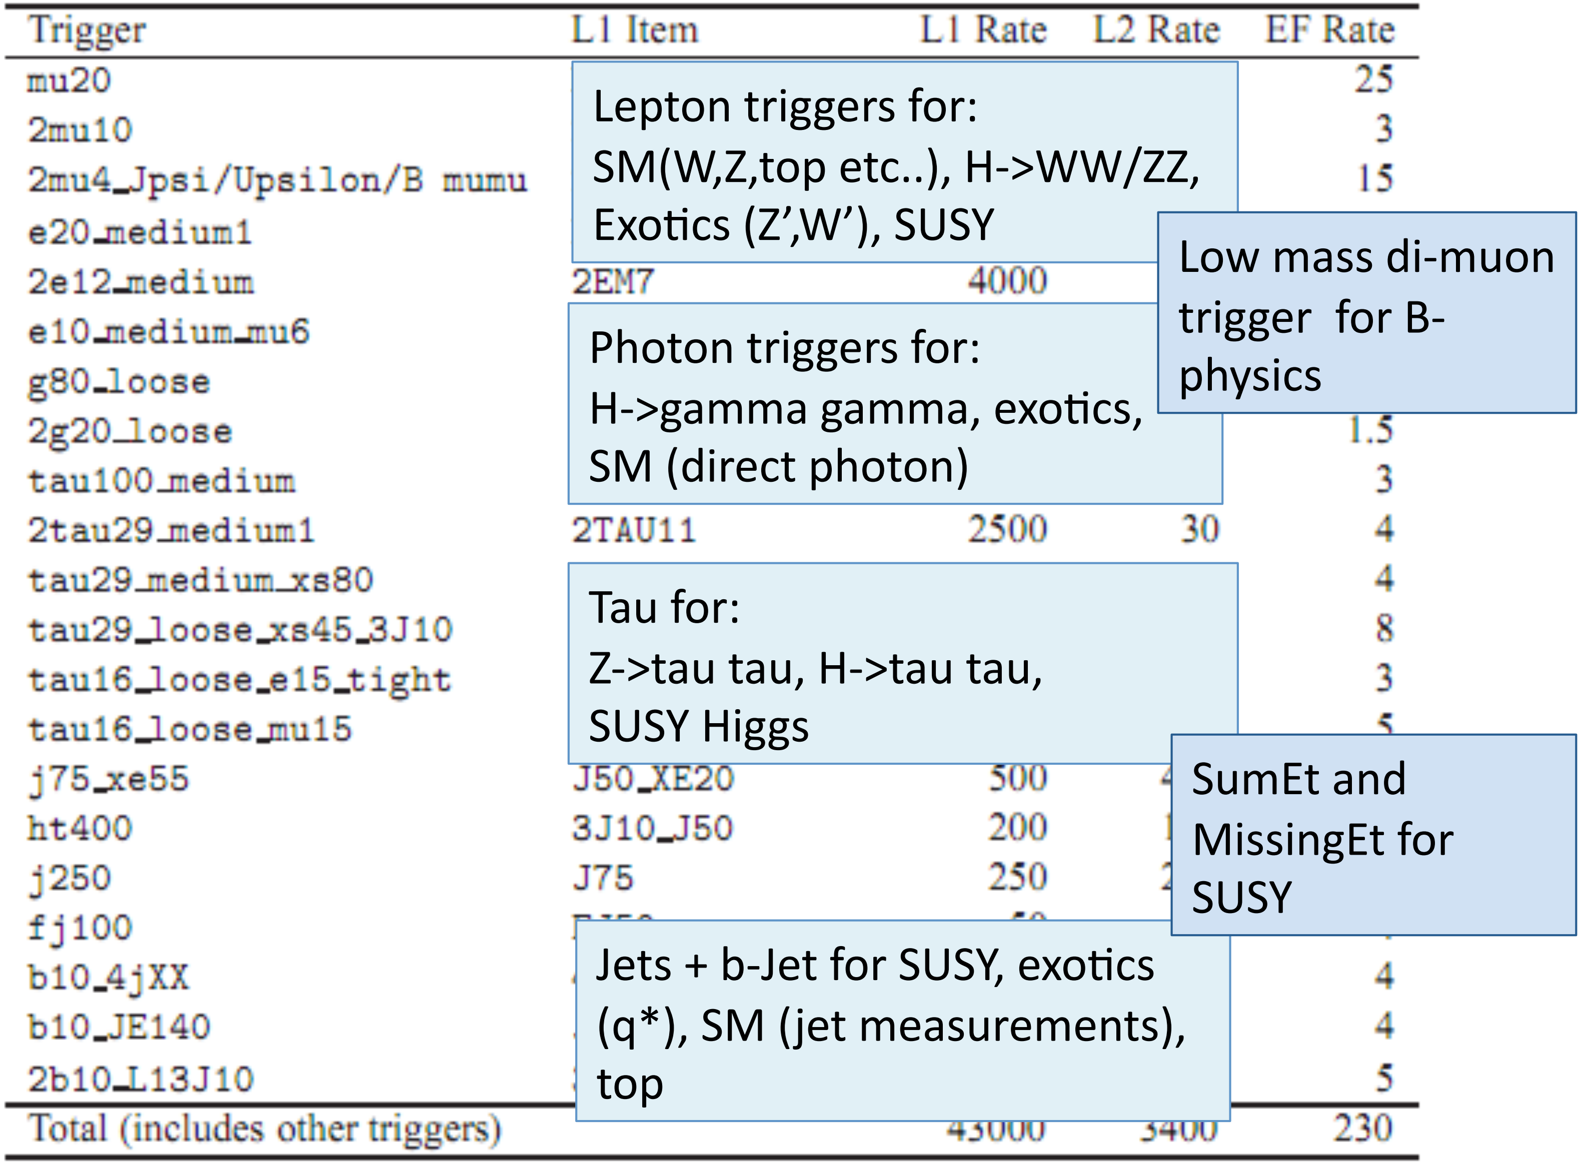

Trigger	L1 Item	L1 Rate	L2 Rate	EF Rate
mu20				25
2mu10				3
2mu4_Jpsi/Upsilon/B mumu				15
e20_medium1				
2e12_medium	2EM7	4000		
e10_medium_mu6				
g80_loose				
2g20_loose				1.5
tau100_medium				3
2tau29_medium1	2TAU11	2500	30	4
tau29_medium_xs80				4
tau29_loose_xs45_3J10				8
tau16_loose_e15_tight				3
tau16_loose_mu15				5
j75_xe55	J50_XE20	500		
ht400	3J10_J50	200		
j250	J75	250		
fj100				
b10_4jXX				4
b10_JE140				4
2b10_L13J10				5
Total (includes other triggers)		43000	3400	230

Lepton triggers for:
SM(W,Z,top etc..), H->WW/ZZ,
Exotics (Z',W'), SUSY

Low mass di-muon trigger for B-physics

Photon triggers for:
H->gamma gamma, exotics,
SM (direct photon)

Tau for:
Z->tau tau, H->tau tau,
SUSY Higgs

SumEt and MissingEt for SUSY

Jets + b-Jet for SUSY, exotics (q*), SM (jet measurements), top

Example Physics usecases....

Example ATLAS trigger menu for $L=10^{33}$

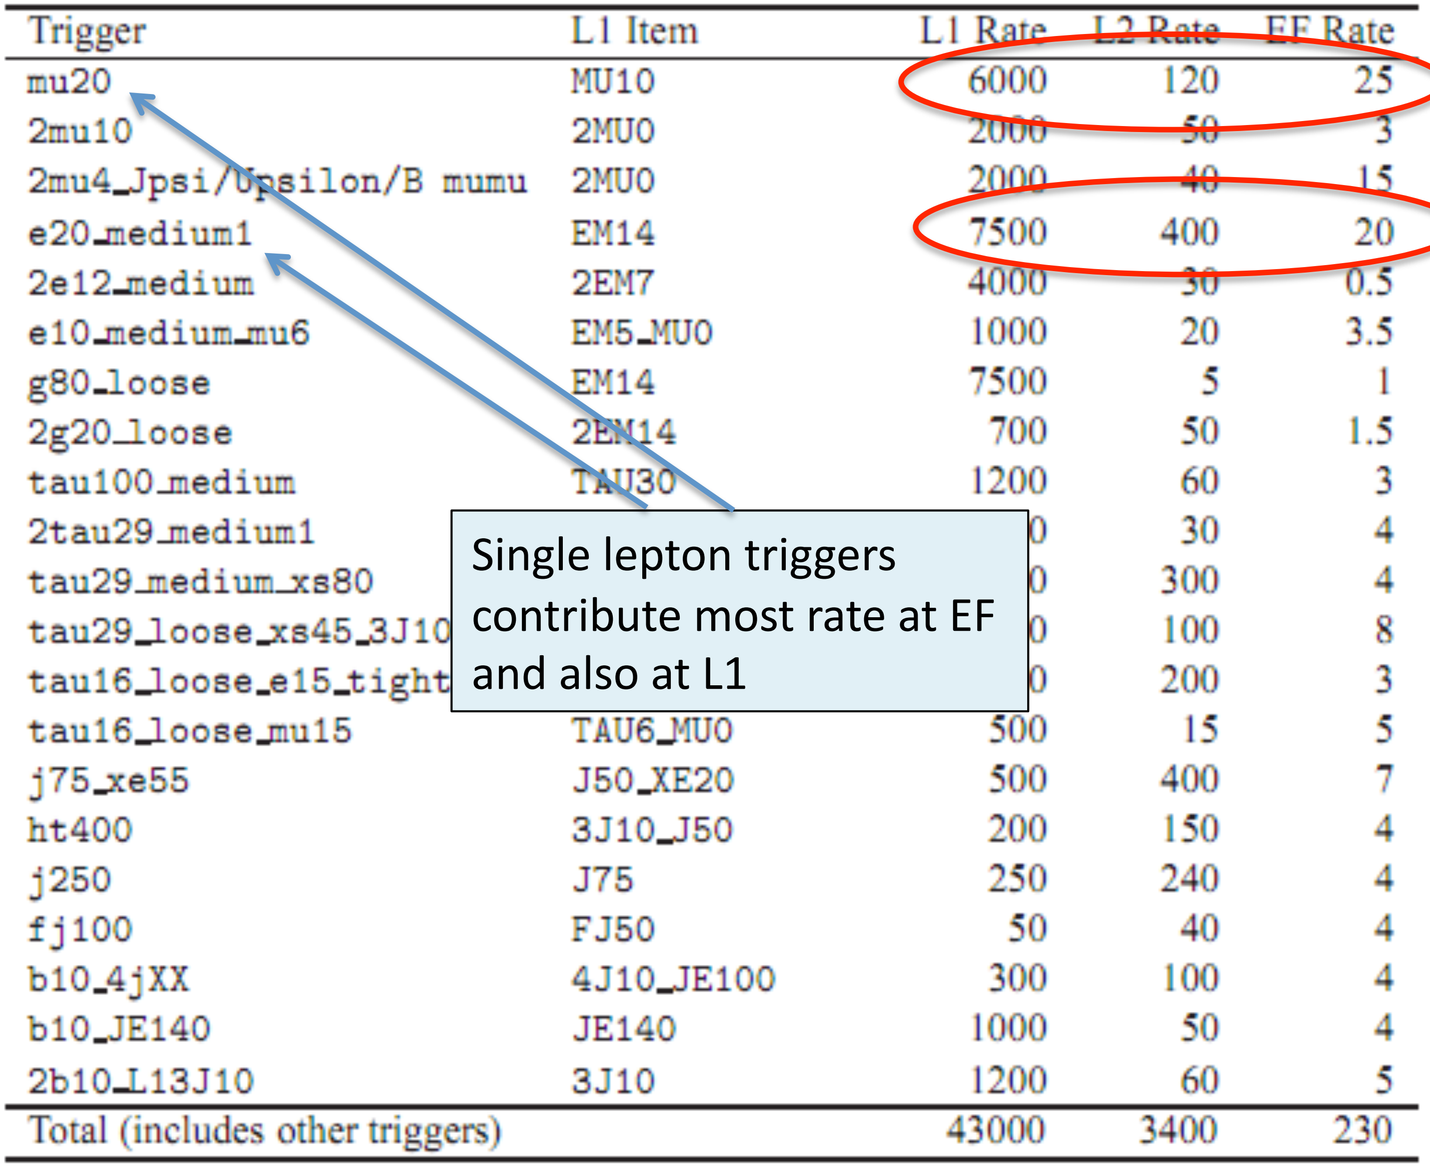

Trigger	L1 Item	L1 Rate	L2 Rate	EF Rate
mu20	MU10	6000	120	25
2mu10	2MU0	2000	50	3
2mu4_Jpsi/Upsilon/B mumu	2MU0	2000	40	1.5
e20_medium1	EM14	7500	400	20
2e12_medium	2EM7	4000	30	0.5
e10_medium_mu6	EM5_MU0	1000	20	3.5
g80_loose	EM14	7500	5	1
2g20_loose	2EM14	700	50	1.5
tau100_medium	TAU30	1200	60	3
2tau29_medium1	[illegible]	[illegible]	30	4
tau29_medium_xs80	[illegible]	[illegible]	300	4
tau29_loose_xs45_3J10	[illegible]	[illegible]	100	8
tau16_loose_e15_tight	[illegible]	[illegible]	200	3
tau16_loose_mu15	TAU6_MU0	500	15	5
j75_xe55	J50_XE20	500	400	7
ht400	3J10_J50	200	150	4
j250	J75	250	240	4
fj100	FJ50	50	40	4
b10_4jXX	4J10_JE100	300	100	4
b10_JE140	JE140	1000	50	4
2b10_L13J10	3J10	1200	60	5
Total (includes other triggers)		43000	3400	230

Single lepton triggers contribute most rate at EF and also at L1

Mu+Bphys
43Hz

El / photon
26.5Hz

Tau
27Hz

Jet/MET
32Hz

Other triggers

- As well as the usual physics object triggers mentioned previously we also have triggers for monitoring, calibration and various exotic physics signals
- For monitoring beam backgrounds we trigger on events which have **unpaired** (non-colliding) bunches passing through the detector
- For noise monitoring we trigger on **Empty** bunches (25ns clock ticks where there are no bunches in the detector)
- Empty bunch triggers are also used for physics analyses looking for new particles which can stop in the detector and decay at a later time (stopped gluino's, R-hadrons etc…)
- Calibration triggers use partial event building to readout only a certain detector region for calibration purposes
  - eg. CMS Ecal calibration using high rate triggers reading out only clusters from π^0 reconstructed in the HLT
  - ATLAS also uses similar techniques

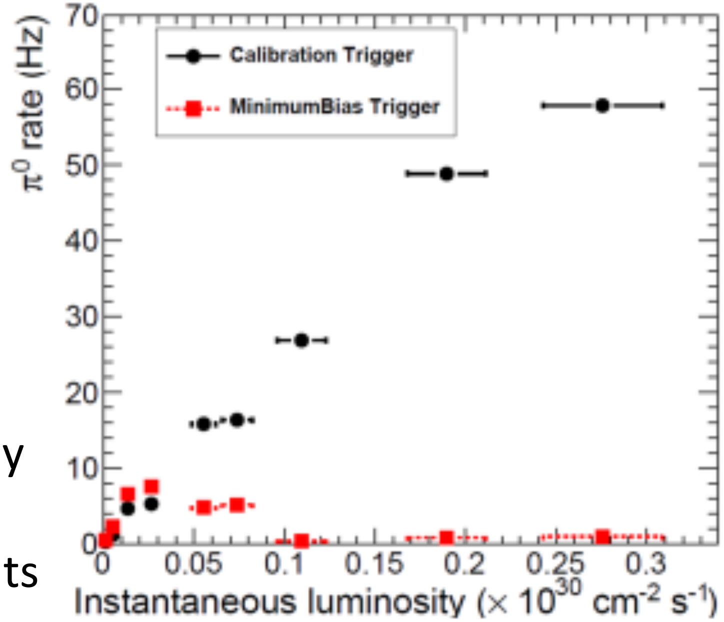

Also have:

- Zero-bias triggers (random filled bunch trigger)
  - Useful for studying pileup effects and for MC overlay
- Forward detector triggers
  - Useful for forward physics, luminosity measurements and beam background monitoring

Tightening the trigger menu

- As the luminosity increases we need to ‘tighten’ the menu to keep the output rates within the limits
- Many different options
  - Raise P_T/E_T thresholds
  - Tighten selection (eg. cut tighter on shower shape variables)
  - Apply isolation (at L1 or HLT or both)
  - Sharpening the turn-on curve with respect to the offline allows better use of the triggers (eg. Using more sophisticated algorithms and calibrations in the trigger)
  - Multi-object triggers designed for specific analyses can also help
- All of the above can effect physics analysis
  - Often different physics groups would prefer different solutions
    - Exotics would prefer higher thresholds, Standard Model prefer tighter selection
- Increased pileup can effect certain triggers more than simple luminosity scaling – mostly Missing E_T and Total E_T triggers
- **Flexible trigger system essential to be able to cope with increasing luminosities**

Tightening the trigger menu

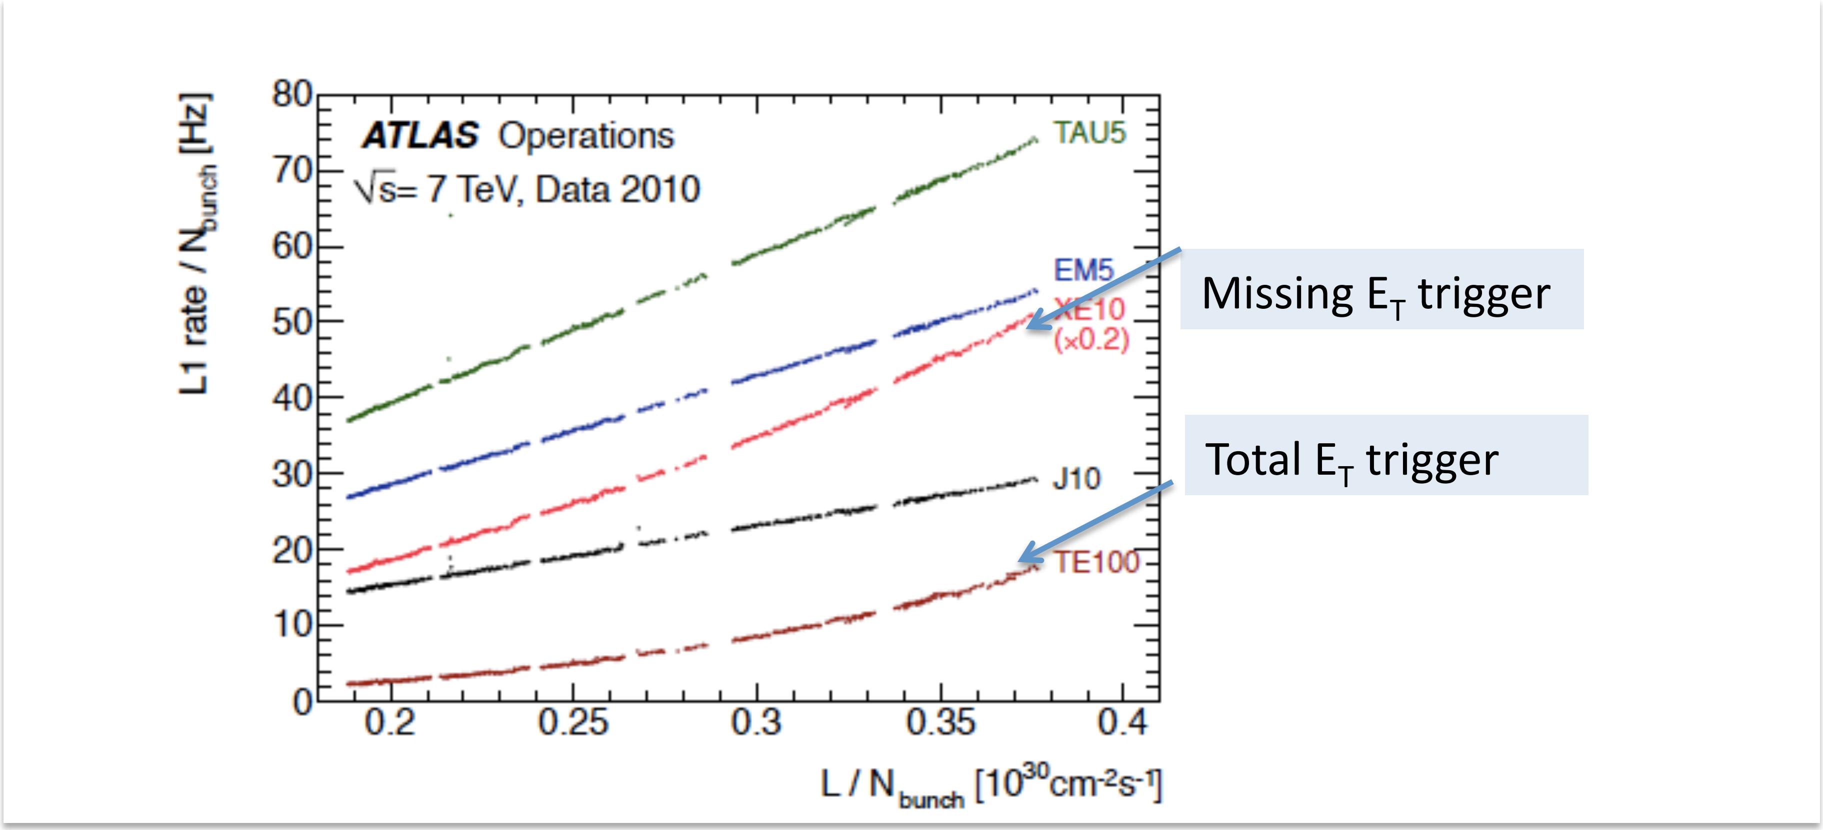

Missing E_T trigger

Total E_T trigger

- Increased pileup can effect certain triggers more than simple luminosity scaling – mostly Missing E_T and Total E_T triggers
- **Flexible trigger system essential to be able to cope with increasing luminosities**

Predicting rates for higher lumi

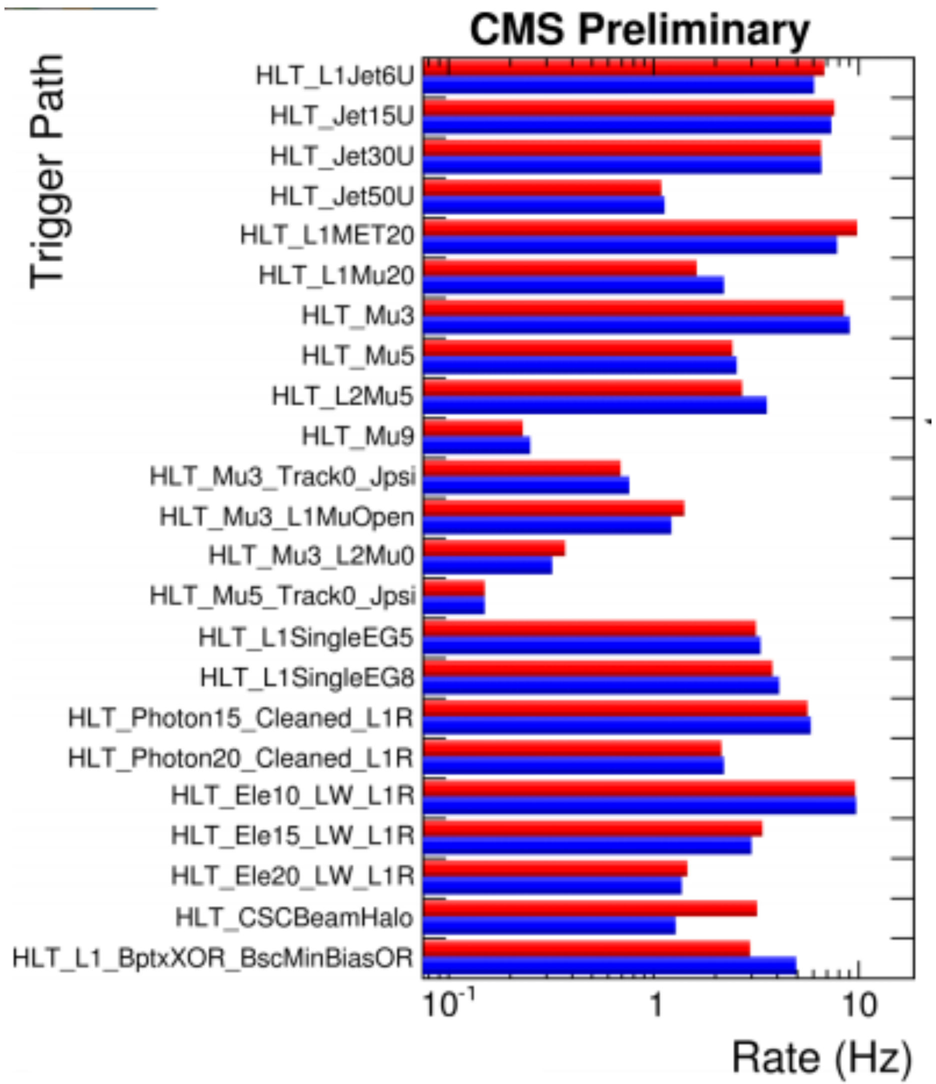

In order to develop trigger menu's for the future we need to predict the rates at higher luminosities.
Examples from ATLAS and CMS.
Blue is the predicted rate (obtained from extrapolating observed rates at lower lumi).
Red is the observed rate.

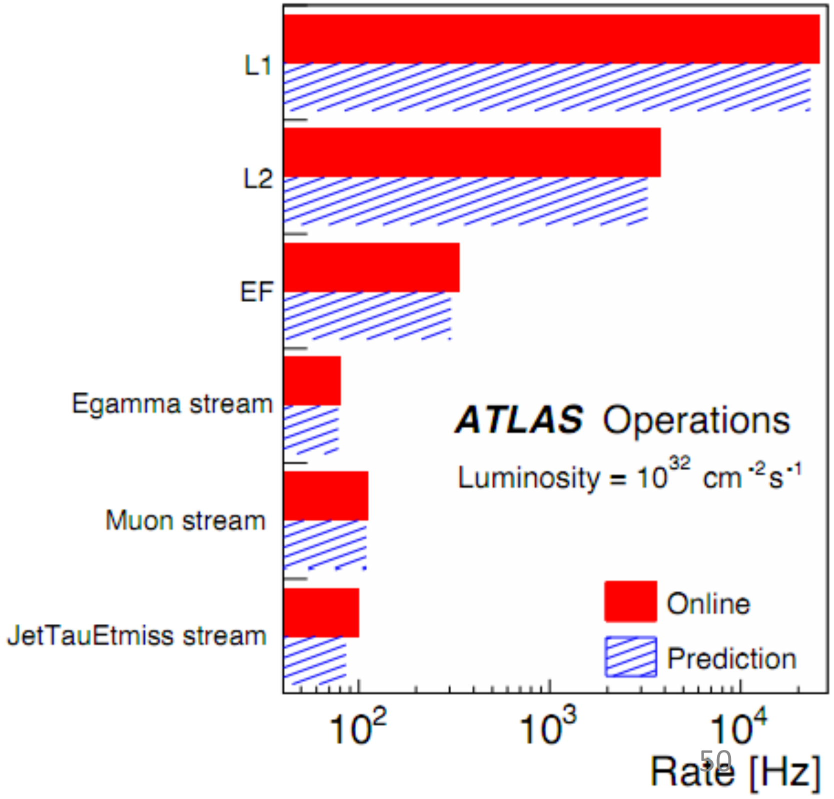

Extrapolation based on large minimum bias samples of data, and assumes the rates scale linearly with luminosity. Corrections for changes in pileup needed for some triggers.

Operational issues – HLT crashes

- What to do if an HLT algorithm crashes or timeouts?
- You could think that a rare crash in the HLT can just be treated as an additional dead-time (ie. It just reduces the luminosity of the dataset by a small amount)
- BUT have to very careful that such crashes are **not biasing the physics**
  - Ie. If all blackhole events (which are very busy) cause the trigger to timeout even if this is a tiny fraction of the events processed it would mean we will not discover blackholes
- In ATLAS if an HLT algorithm crashes or timeouts we send the event to a debug stream
  - We then run the trigger on these events offline to see if they would have passed the trigger (for timeouts)
  - If so we inject them back into the physics streams so they are used by analysis:

  ATLAS Search analyses routinely run over the debug stream events
  - In practice this is very rare with a few timeouts a day
- **HLT algorithm need to be extremely robust**
  - Many of these algorithms ran over several hundred billion of events in 2010

Operational issues – Problems during a run

- Many problems can occur during a run which can effect the trigger
  - Noisy regions in the detector causing very high trigger rate
- Can lead to high dead-time in the system which means effectively losing physics events
- Sometimes can mask effected detector region at input to trigger
- Or can change trigger prescales during the run so can disable effected triggers

Spikes in missing E_T trigger rate from hot tower in the calorimeter 'fixed' with a prescale change

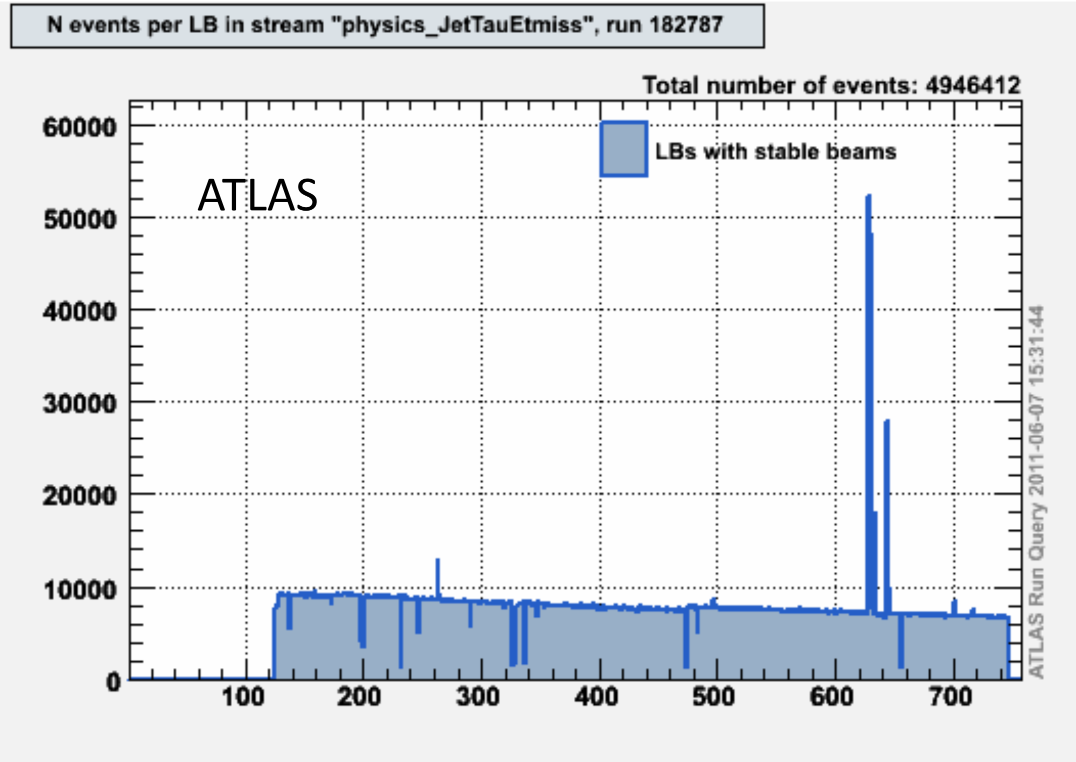

Operational Issues – changing the Trigger

- There can be many good reasons to change things in the trigger
  - Change menu to be able to run at higher lumi
  - Change calibration constants, algorithm logic in order to have better rejection – sharper turn-on curves
  - Change algorithms to bugs
- However any change **needs extremely careful validation**
  - Can be very work intensive
- Also from the physics analysis point of view want a **stable trigger**
  - Can not remake all MC simulation samples with updated trigger configuration
  - Need to make sure any changes do not kill an important analysis
- **Need to balance trigger improvements versus stability** (disruption from change)
  - Best to keep all changes until a well defined time if possible

Operational Issues – Trigger monitoring

- The overall trigger rates and dead-time is one of the most important things that is monitored during data-taking
  - A sudden increase in rate or dead-time indicates a detector problem
- More sophisticated monitoring of the trigger algorithms themselves
  - Resource usage, algorithm errors, trigger object η-ϕ maps, trigger object efficiency and resolutions
- **The trigger is a very complicated system and extensive monitoring is essential for smooth and efficient operation**

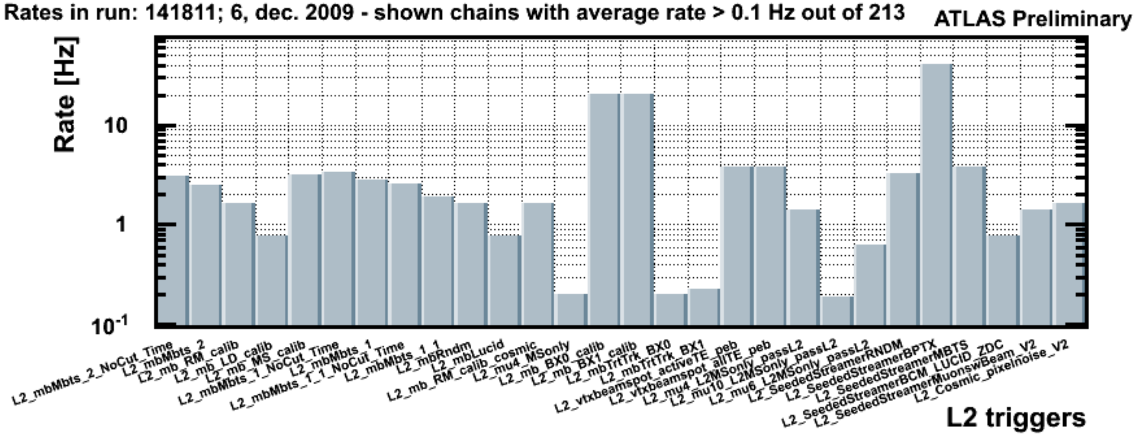

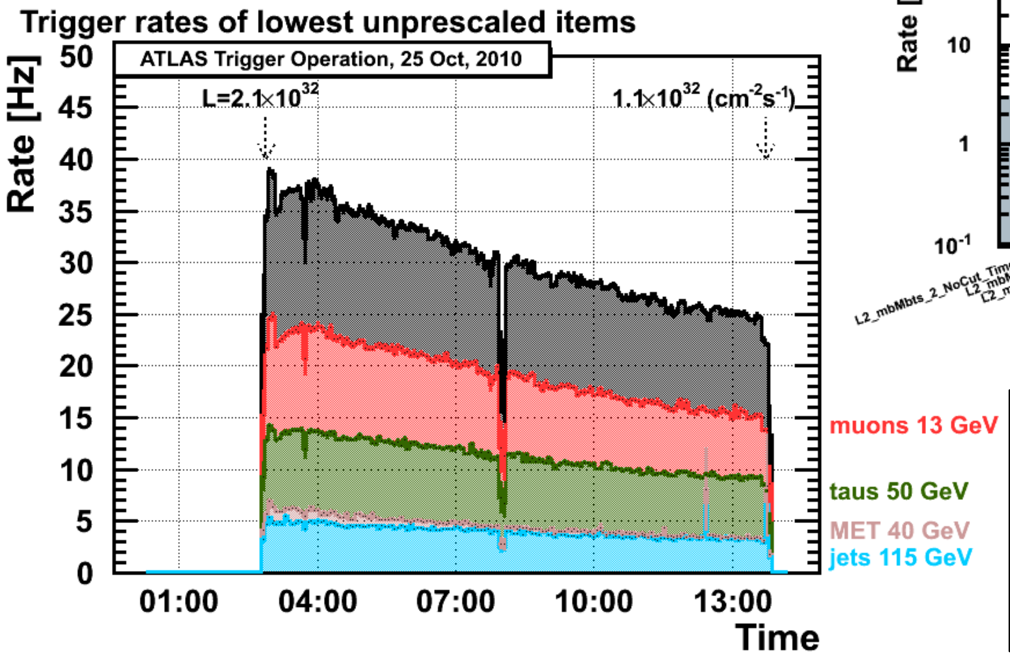

If the trigger is not working correctly this directly effects all data-taking.

Any problem needs to be fixed or bypassed in the minimum time (day or night)

Summary

- Designing and building LHC experiment trigger system is **very challenging**
  - Higher collision rate, Event size and much more complex detectors compared to previous experiments
- Achieved with **complex** and **flexible** systems using
  - Calorimeter and Muon hardware trigger at L1
  - All detectors at full granularity at HLT
- All LHC experiment Trigger systems are **running extremely well** so far
- Current challenges
  - **Increasing luminosity** and high pileup
  - **Smooth** and **efficient operation**
  - Upgrade for sLHC (lumi of ~10^{35}, pileup of 500!)

References

- Much of this lecture taken from the following lectures/talks/papers between them these have much more details so read on if you are interested!
- Previous Schools:
  - https://indico.fnal.gov/getFile.py/access?contribId=44&sessionId=36&resId=0&materialId=slides&confId=3532
  - https://indico.fnal.gov/getFile.py/access?contribId=49&sessionId=42&resId=0&materialId=slides&confId=3532
  - http://hoecker.home.cern.ch/hoecker/talks2/hoecker_hcpss09-1.pdf
  - http://hoecker.home.cern.ch/hoecker/talks2/hoecker_hcpss09-2.pdf
- CERN Academic Training:
  - http://cdsweb.cern.ch/record/1350789/
- Very cutting edge ATLAS trigger conference talk at TIPP 2011 conference:
  - http://indico.cern.ch/contributionDisplay.py?sessionId=19&contribId=21&confId=102998

Many thanks to….

Brian Petersen, Wainer Vandelli, Andreas Hoecker, Brigitte Vachon, Alex Tapper, Wesley Smith, Jim Brooke, David Berge, Srini Rajagopalan, Thorsten Wengler, David Salek…

“Region of Interest” in ATLAS

- In ATLAS the HLT is divided into 2 levels – L2 and EventFilter (EF)
- L2 only has access to detector data in a “Region of Interest” around the L1 trigger objects
  - L2 only requests a few % of the data
- L2 algorithms run with full detector granularity on the data in this region (and include tracking)
- The L2 output rate is the Event-Building rate (design ~2kHz – current 4kHz)
  - The RoI means the Event Building rate is much reduced
- EventFilter runs more detailed HLT algorithms with access to the full event (output rate ~300Hz)

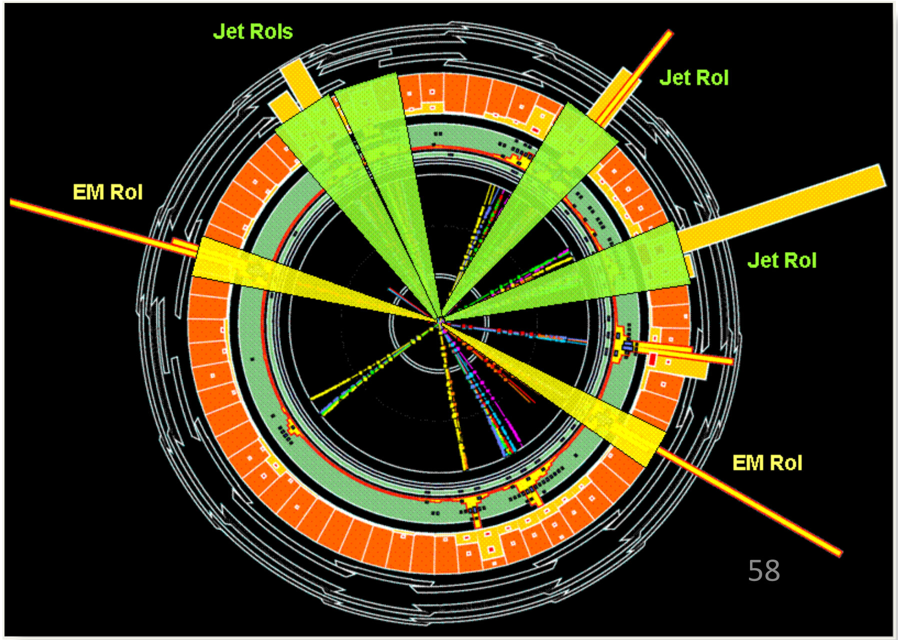

ATLAS and CMS HLT Concepts

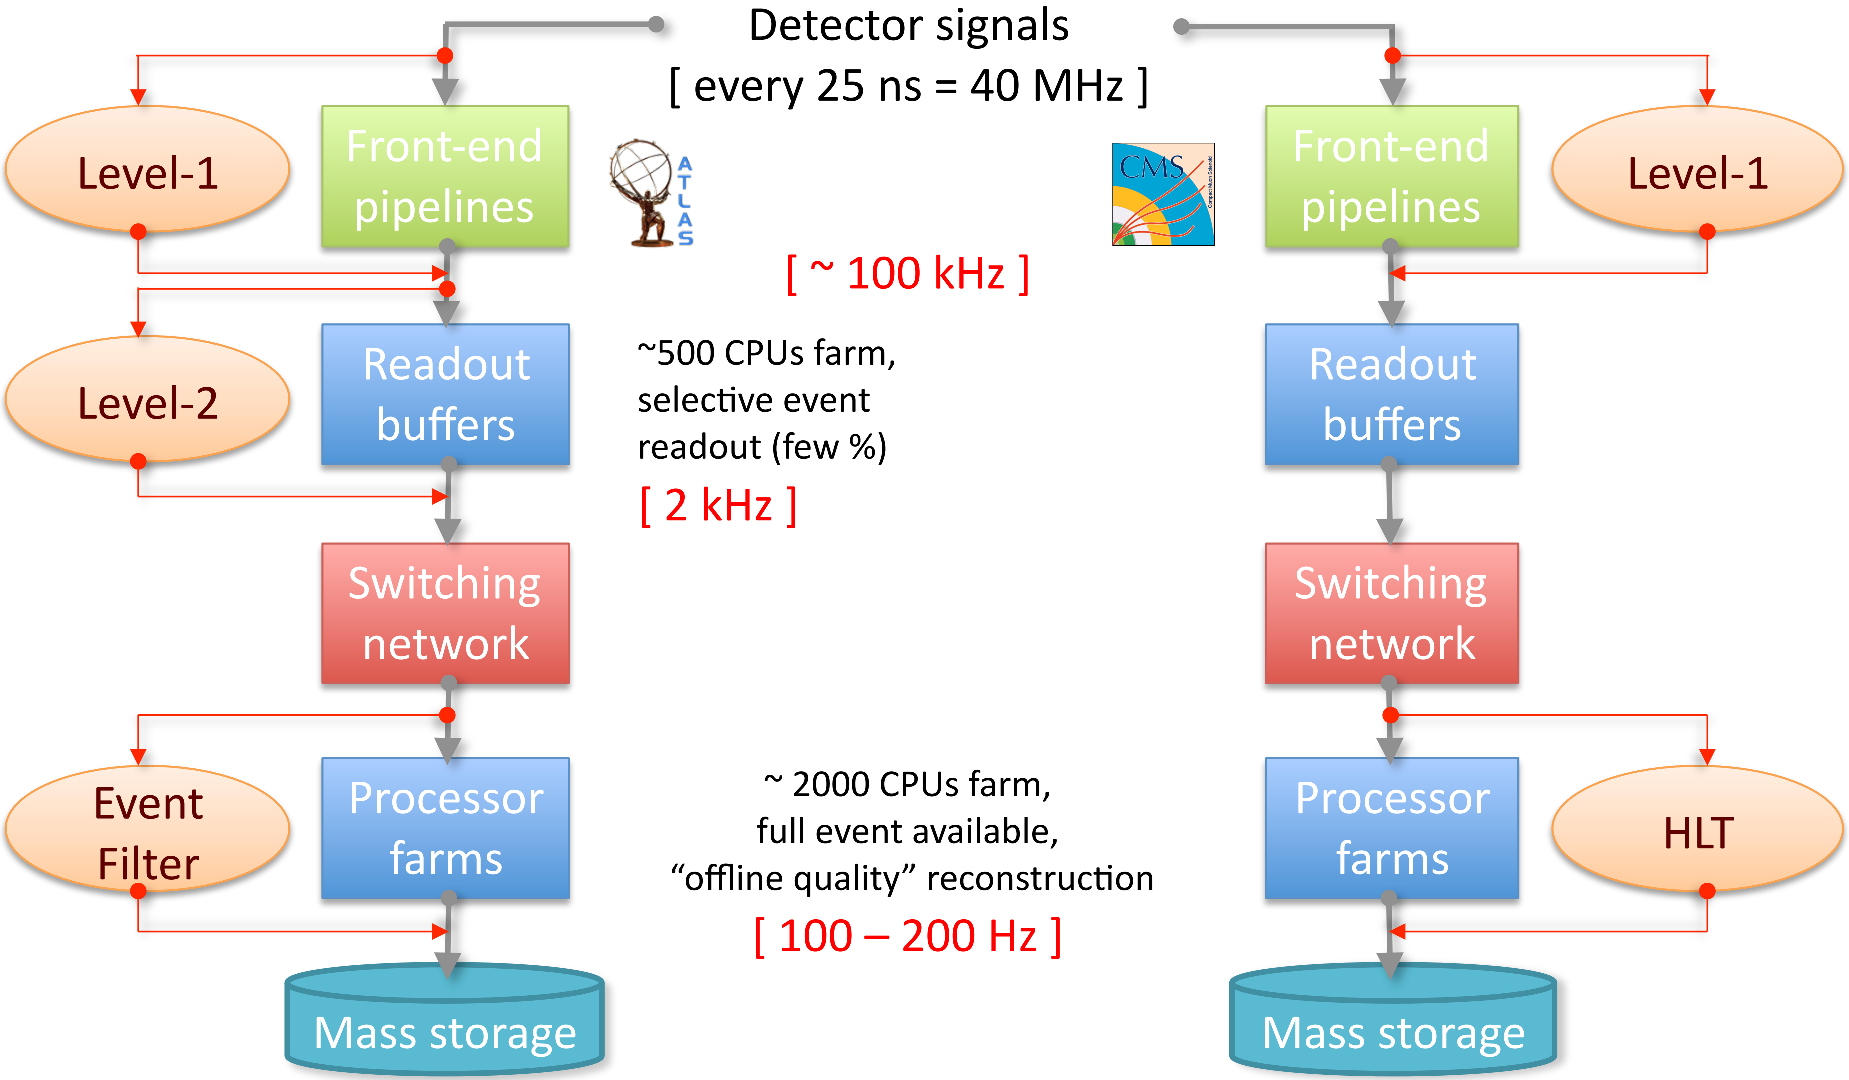

Detector signals

[every 25 ns = 40 MHz]

Level-1

Front-end pipelines

[~ 100 kHz]

Level-2

Readout buffers

~500 CPUs farm,
selective event
readout (few %)

[2 kHz]

Switching network

Event Filter

Processor farms

~ 2000 CPUs farm,
full event available,
“offline quality” reconstruction

[100 – 200 Hz]

Mass storage

Front-end pipelines

Level-1

Readout buffers

Switching network

Processor farms

HLT

Mass storage

Numbers are approximate design values

Trigger & Data Flow

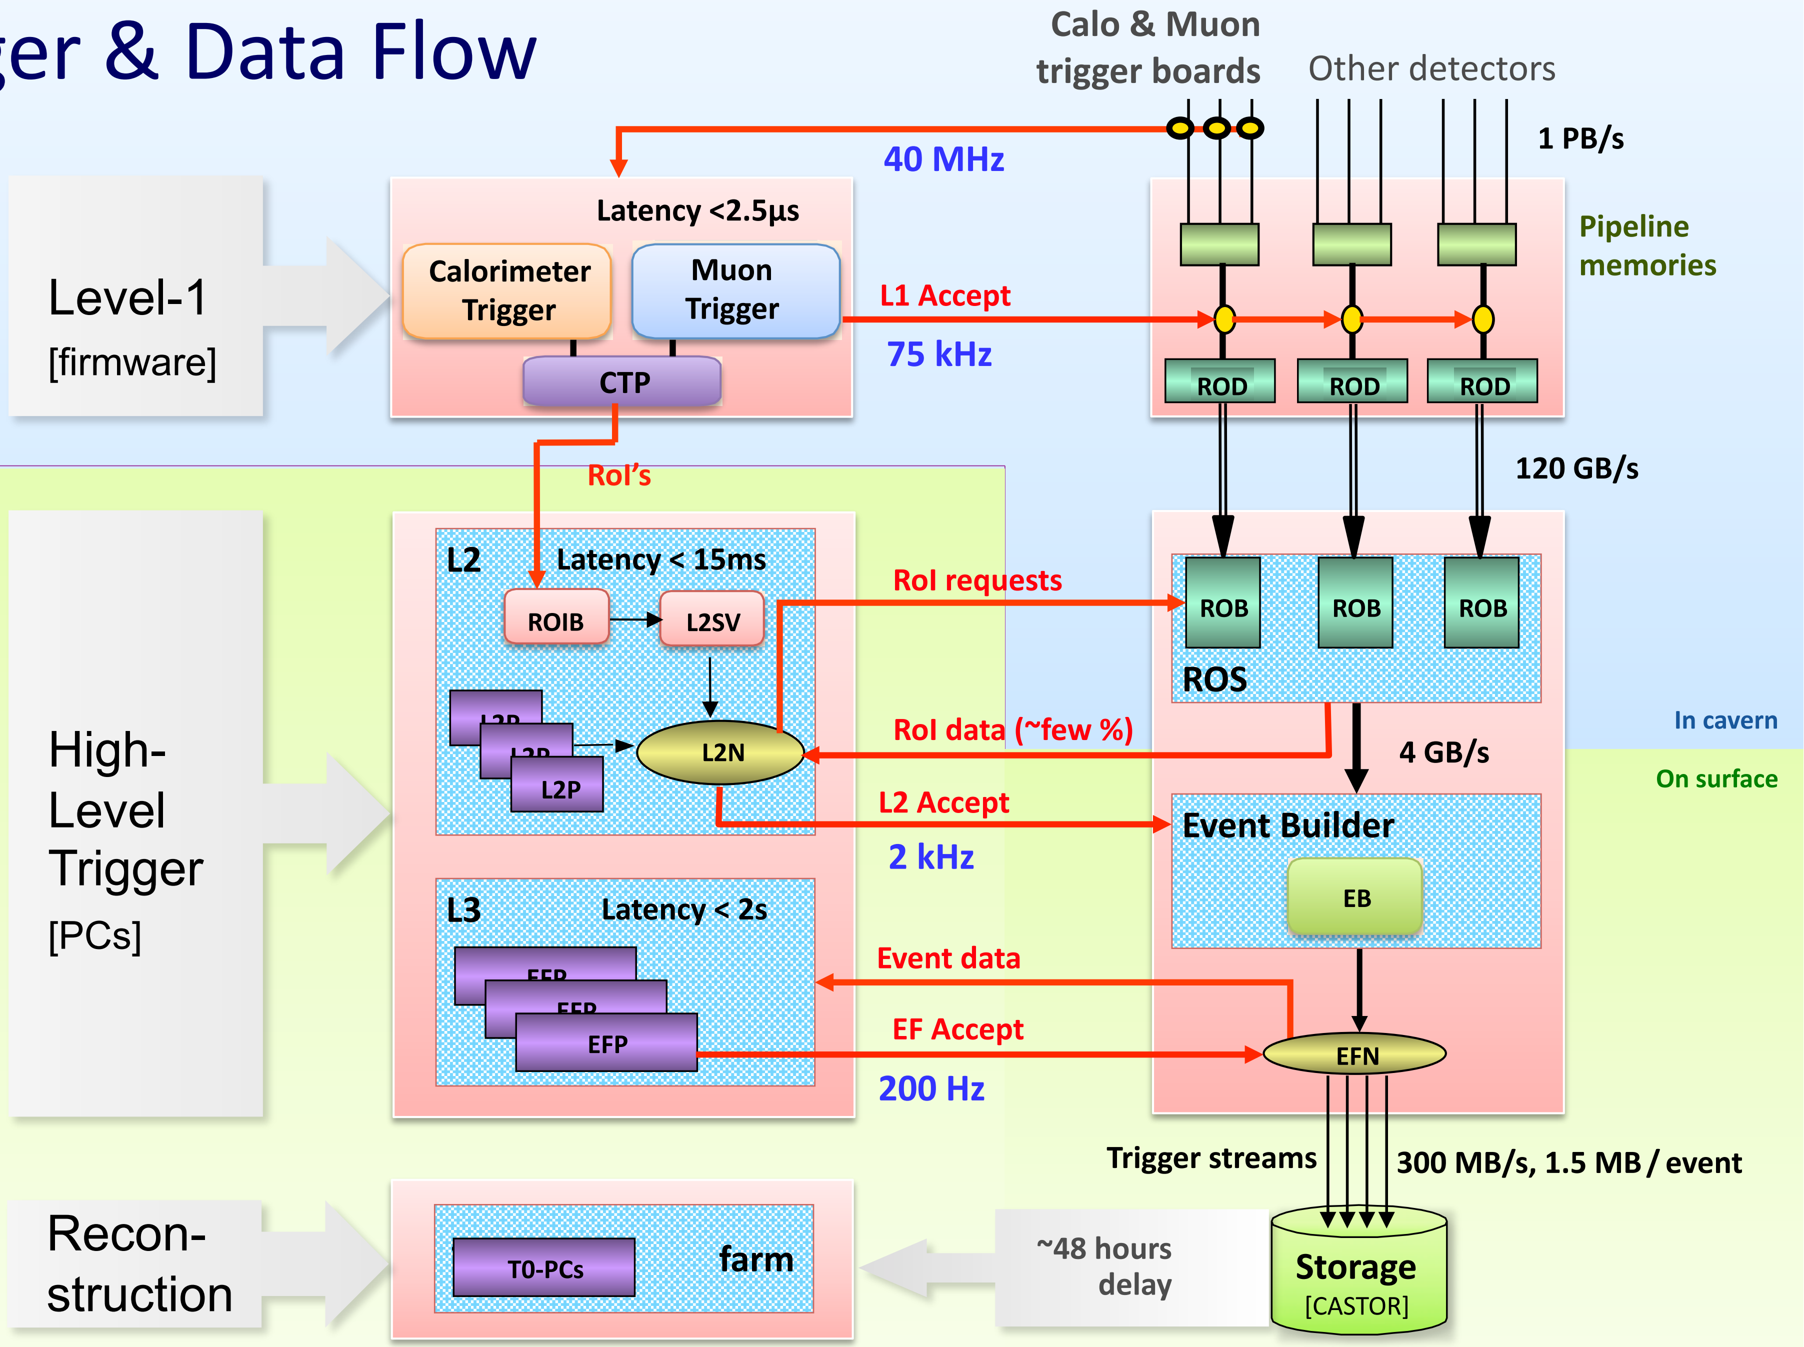

Calo & Muon trigger boards

Other detectors

1 PB/s

40 MHz

Level-1
[firmware]

Latency <2.5µs

Calorimeter Trigger

Muon Trigger

CTP

L1 Accept

75 kHz

Pipeline memories

ROD

ROD

ROD

120 GB/s

In cavern

On surface

RoI's

High-Level Trigger
[PCs]

L2

Latency < 15ms

ROIB

L2SV

L2P

L2P

L2P

L2N

RoI requests

RoI data (~few %)

L2 Accept

2 kHz

ROB

ROB

ROB

ROS

4 GB/s

In cavern

On surface

Event Builder

EB

L3

Latency < 2s

EFP

EFP

EFP

Event data

EF Accept

200 Hz

EFN

Trigger streams

300 MB/s, 1.5 MB / event

Recon-struction

T0-PCs

farm

~48 hours delay

Storage
[CASTOR]

Numbers are design values

ATLAS TDAQ rates 2010 → 2011

ATLAS Data Trigger Info

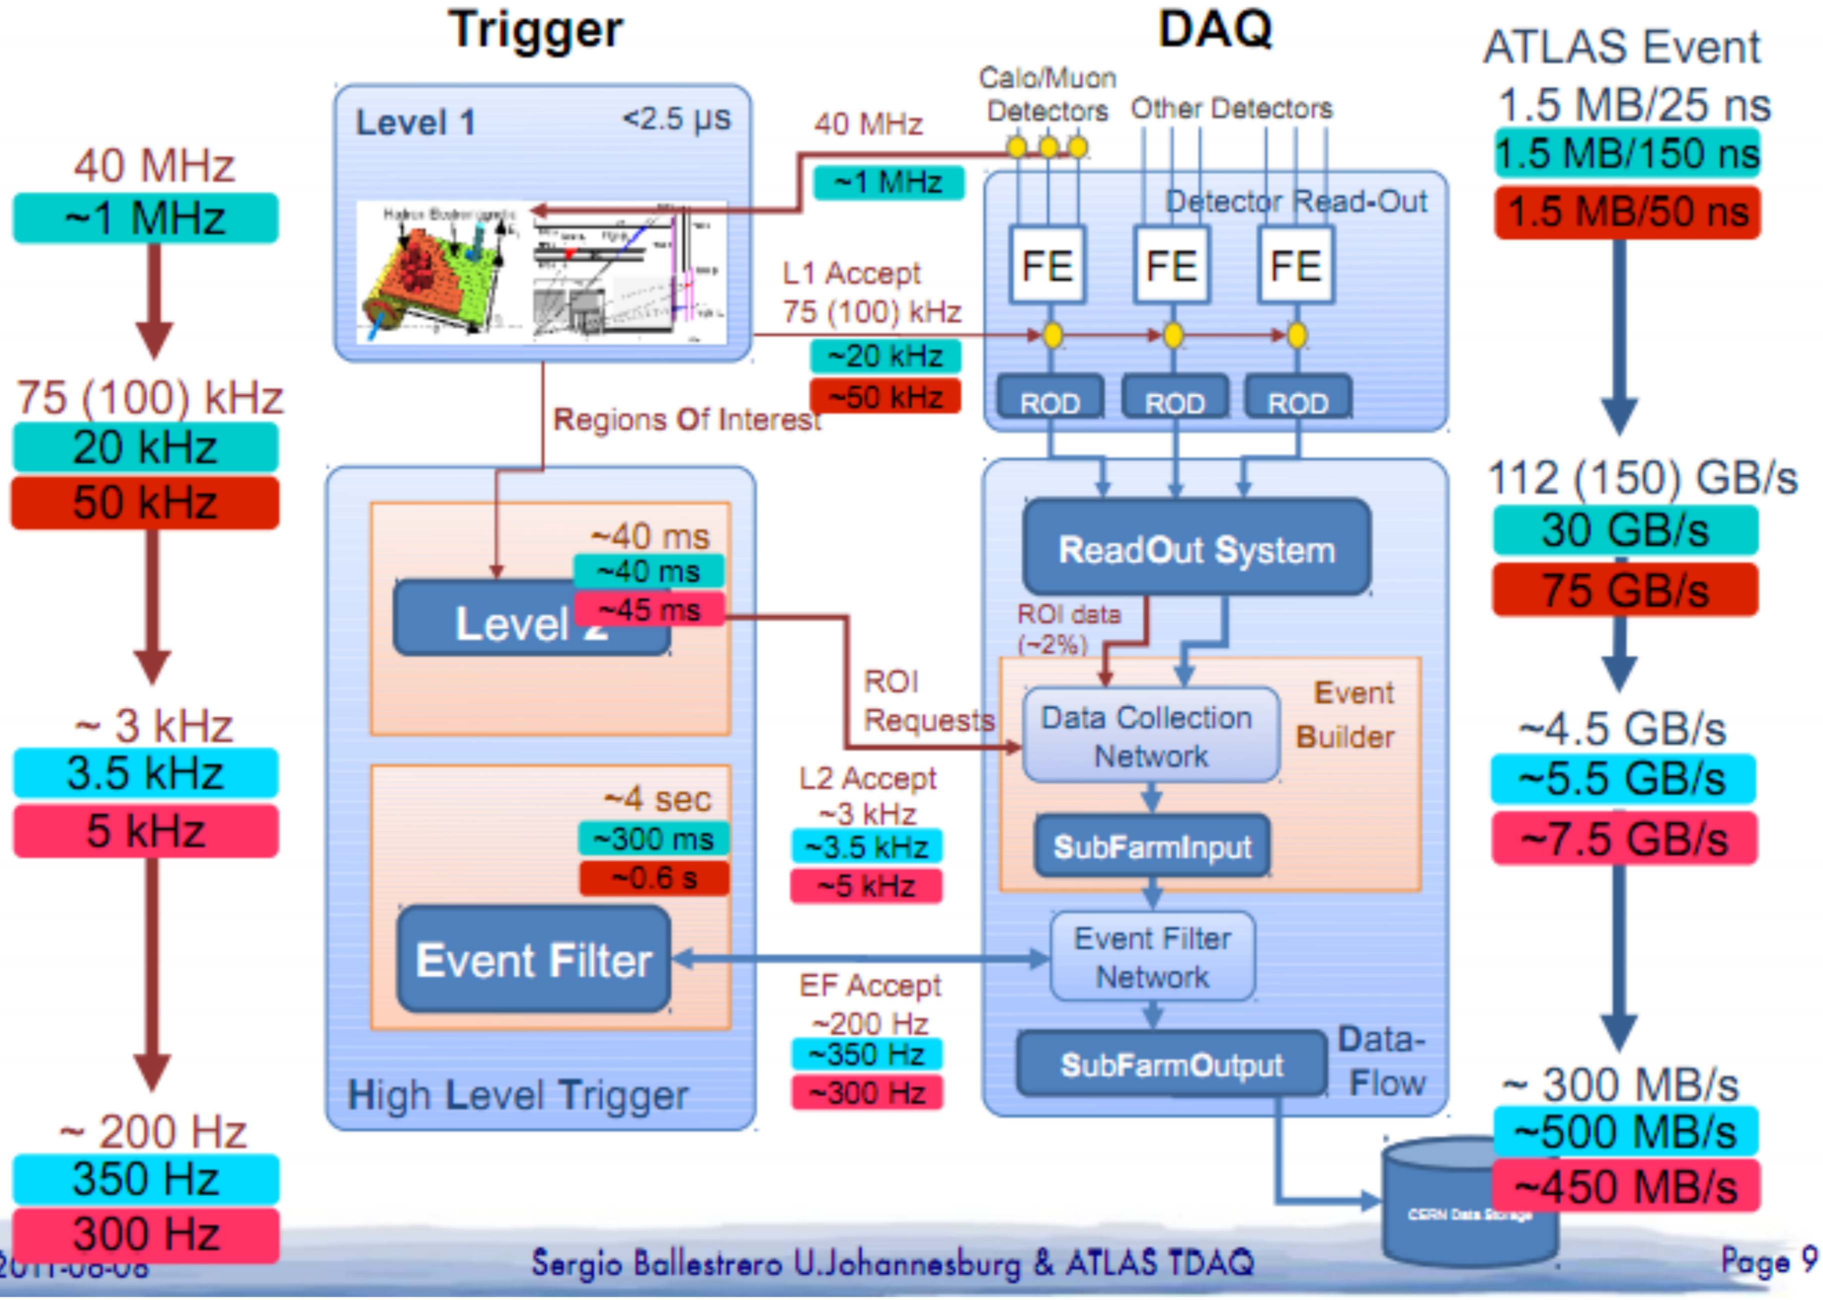

40 MHz
~1 MHz
75 (100) kHz
20 kHz
50 kHz
~ 3 kHz
3.5 kHz
5 kHz
~ 200 Hz
350 Hz
300 Hz

Trigger

Level 1 <2.5 μs

Regions Of Interest

~40 ms
~40 ms
~45 ms
Level 2

~4 sec
~300 ms
~0.6 s
Event Filter

High Level Trigger

40 MHz
~1 MHz

L1 Accept
75 (100) kHz
~20 kHz
~50 kHz

ROI
Requests

L2 Accept
~3 kHz
~3.5 kHz
~5 kHz

EF Accept
~200 Hz
~350 Hz
~300 Hz

DAQ

Calo/Muon
Detectors

Other Detectors

Detector Read-Out

FE
FE
FE

ROD
ROD
ROD

ReadOut System

ROI data
(~2%)

Data Collection
Network

Event
Builder

SubFarmInput

Event Filter
Network

SubFarmOutput

Data-
Flow

CERN Data Storage

ATLAS Event
1.5 MB/25 ns
1.5 MB/150 ns
1.5 MB/50 ns

112 (150) GB/s
30 GB/s
75 GB/s

~4.5 GB/s
~5.5 GB/s
~7.5 GB/s

~ 300 MB/s
~500 MB/s
~450 MB/s

LHC bunch structure

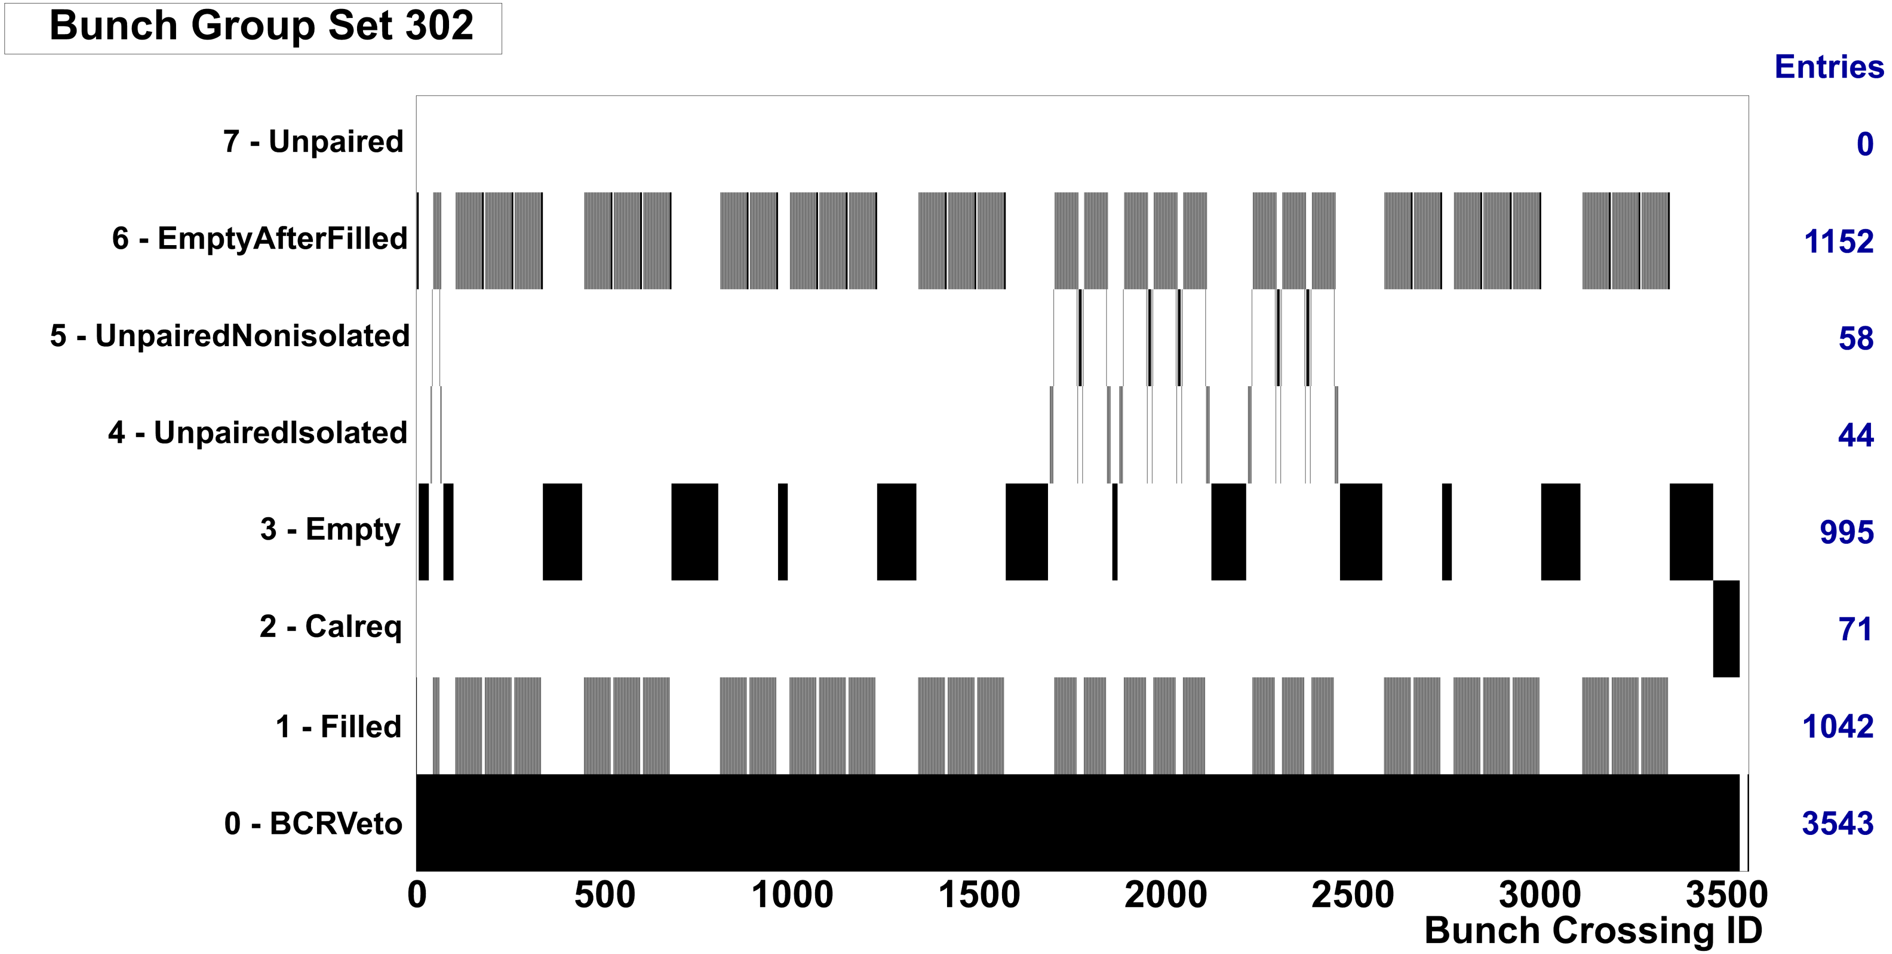

The LHC beam is structured in ~3000 25ns buckets (bunch crossing).
Each bunch-crossing can be classified as Filled (for colliding bunches), Empty (for no bunches) and Unpaired (for a bunch in 1 beam but not the other).
Above is a graphic of this for a recent fill.

Operational strategy

- Typical trigger menu provides several threshold points and designed to work over order of magnitude in Luminosity
- Keep primary triggers as stable as possible, not just over the run but over long data periods
  - Highly desirable for analyses
- As Luminosity falls over fill:
  - Enable additional triggers
    - Requiring high L1 rate.
- Keep supporting and other monitoring triggers at fixed output rate.
- Rates adjusted during the fill with “prescale sets”, deployed during a run without the need for a stop/re-start.

Simulated rate evolution in an LHC fill

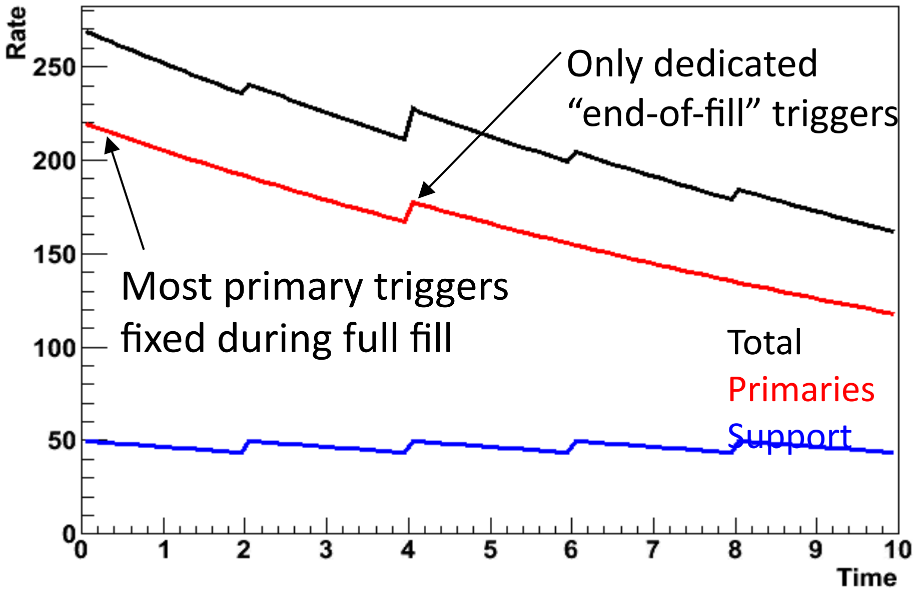

ATLAS L1 Latency

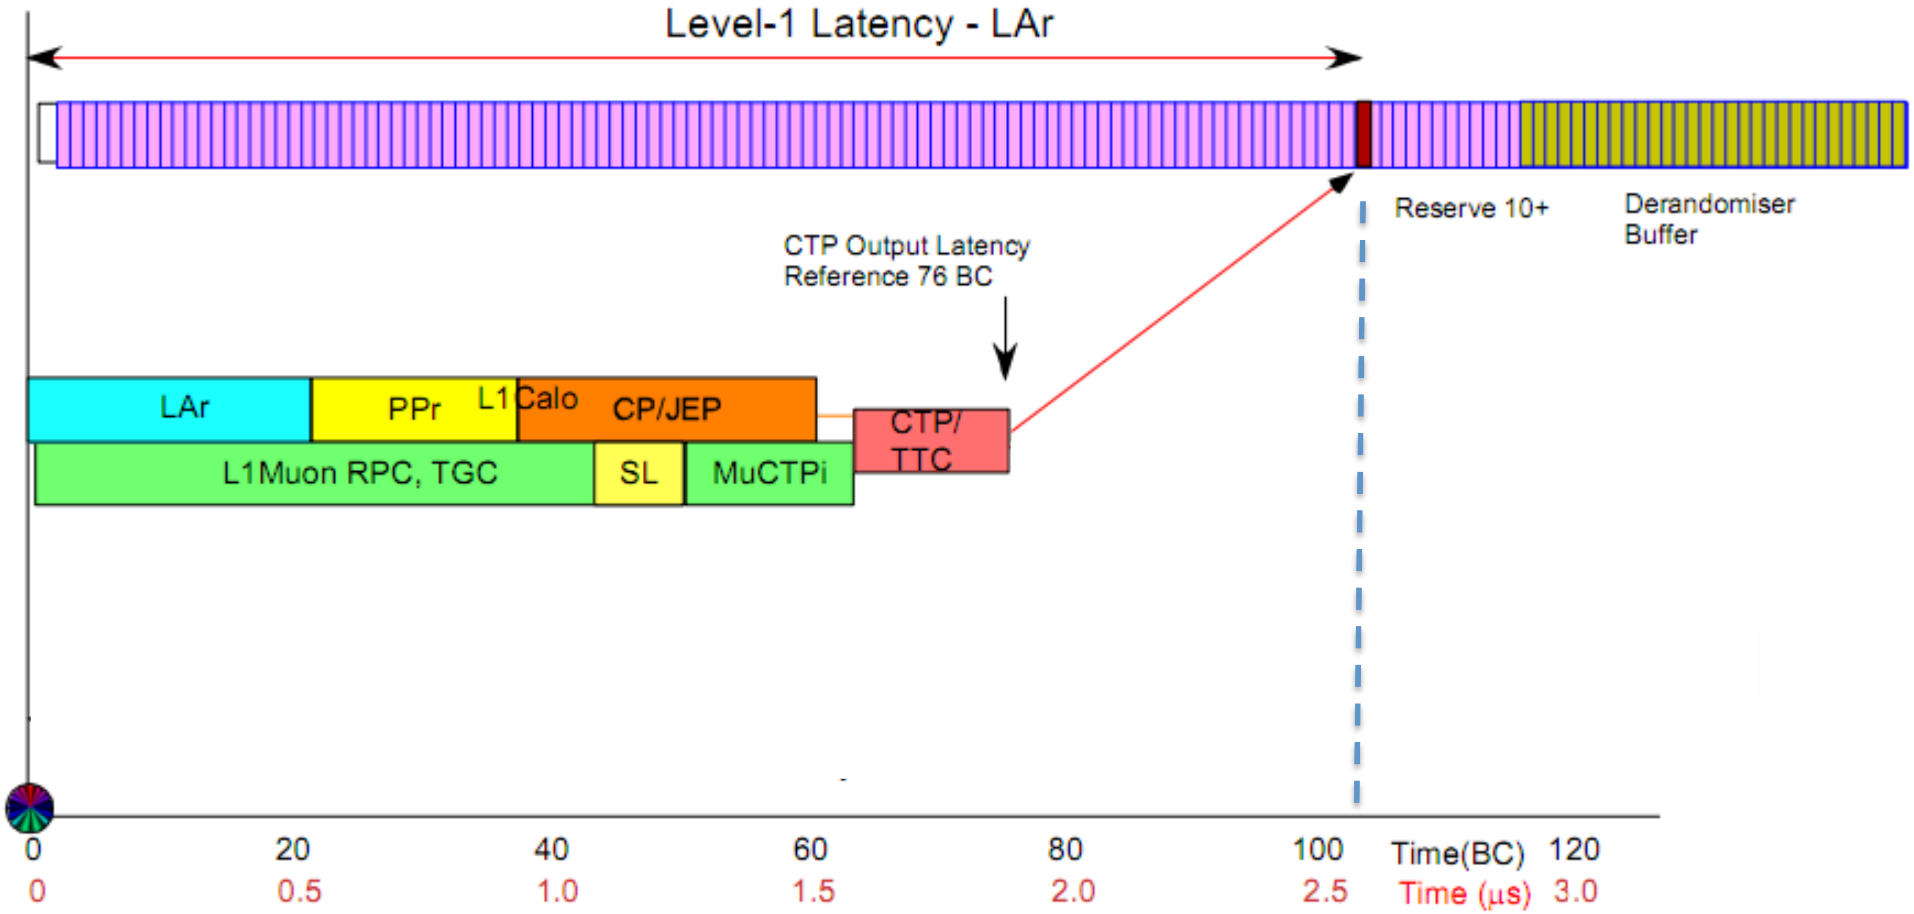

More on Trigger dead-time

In theory if the trigger rate is too high then the dead-time will limit the rate. However do not want the dead-time to control the rates as this will veto interesting physics events at the expense of less-interesting higher rate events.

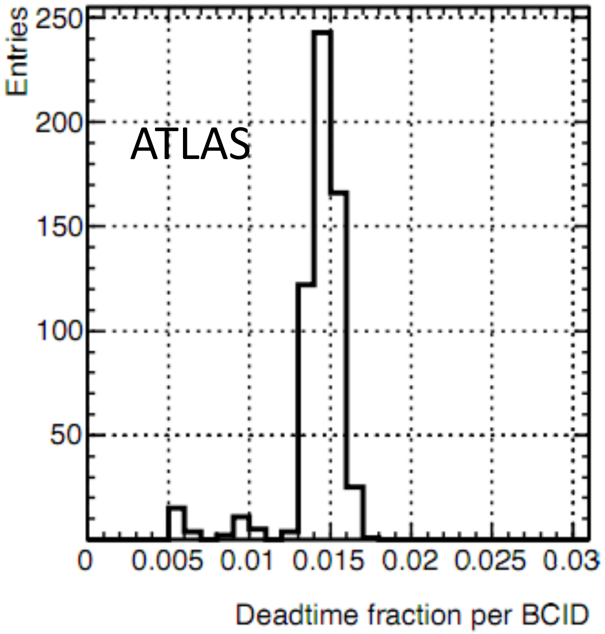

Dead-time primarily related to the L1 trigger rate.
However because the DAQ system is pipelined – everything in the online path can cause dead-time. HLT algorithms taking too long will mean the HLT fills up and the next L1 triggered events can not be processed in time – this feeds back to the L1 trigger being inhibited.
Even writing the data to tape can cause dead-time if it is too slow.

ATLAS L1 Calorimeter Trigger Logic

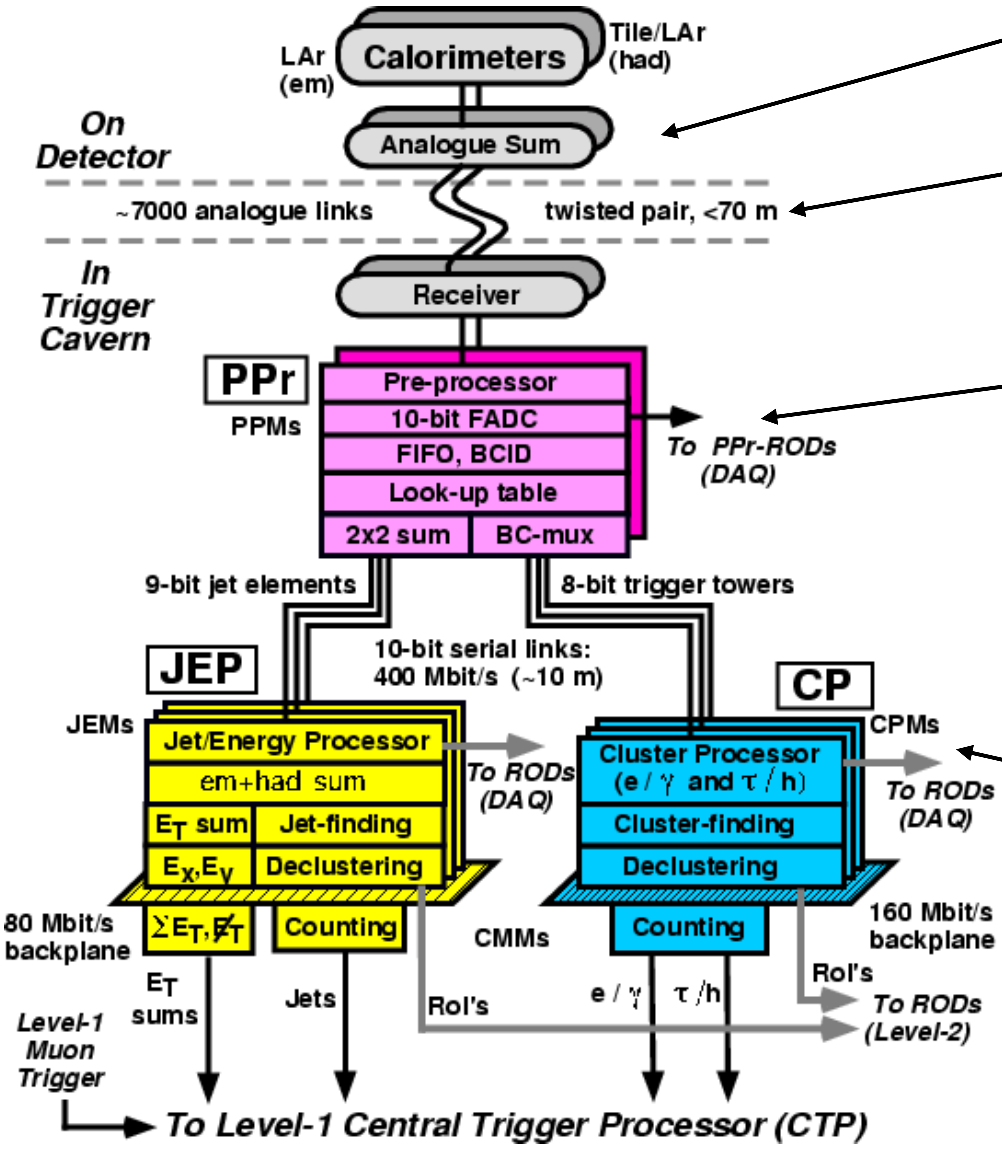

Calorimeter cells in analog sum on detector

Transmitted to underground counting room

Signal timing adjusted, signal digitized and converted to tower ET (*ADC* and *ASICs*)

Cluster identification and isolation requirement (in *FPGA*)

A simple Triggered Experiment

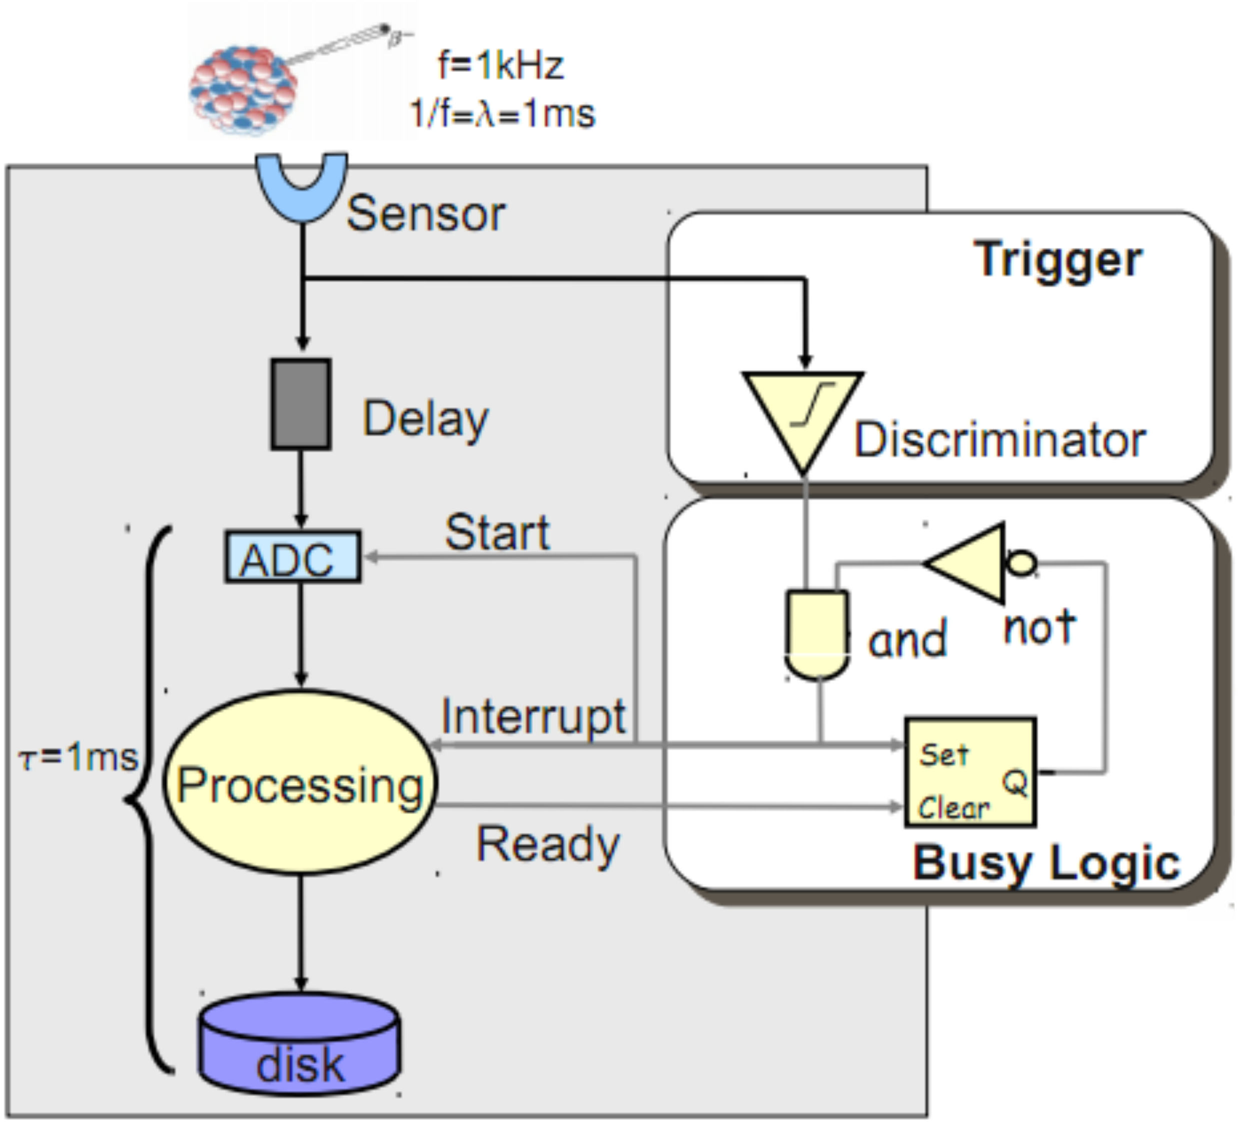

Eg. Measuring Beta-decay. Stochastic process so need a trigger.

Need to delay the signal to be in time with the trigger.

Have BUSY logic to prevent a new trigger while current event is being processed/recorded.

BUSY causes deadtime if input rate is too fast compared to latency of trigger, processing and readout system.

Reducing the deadtime with buffering

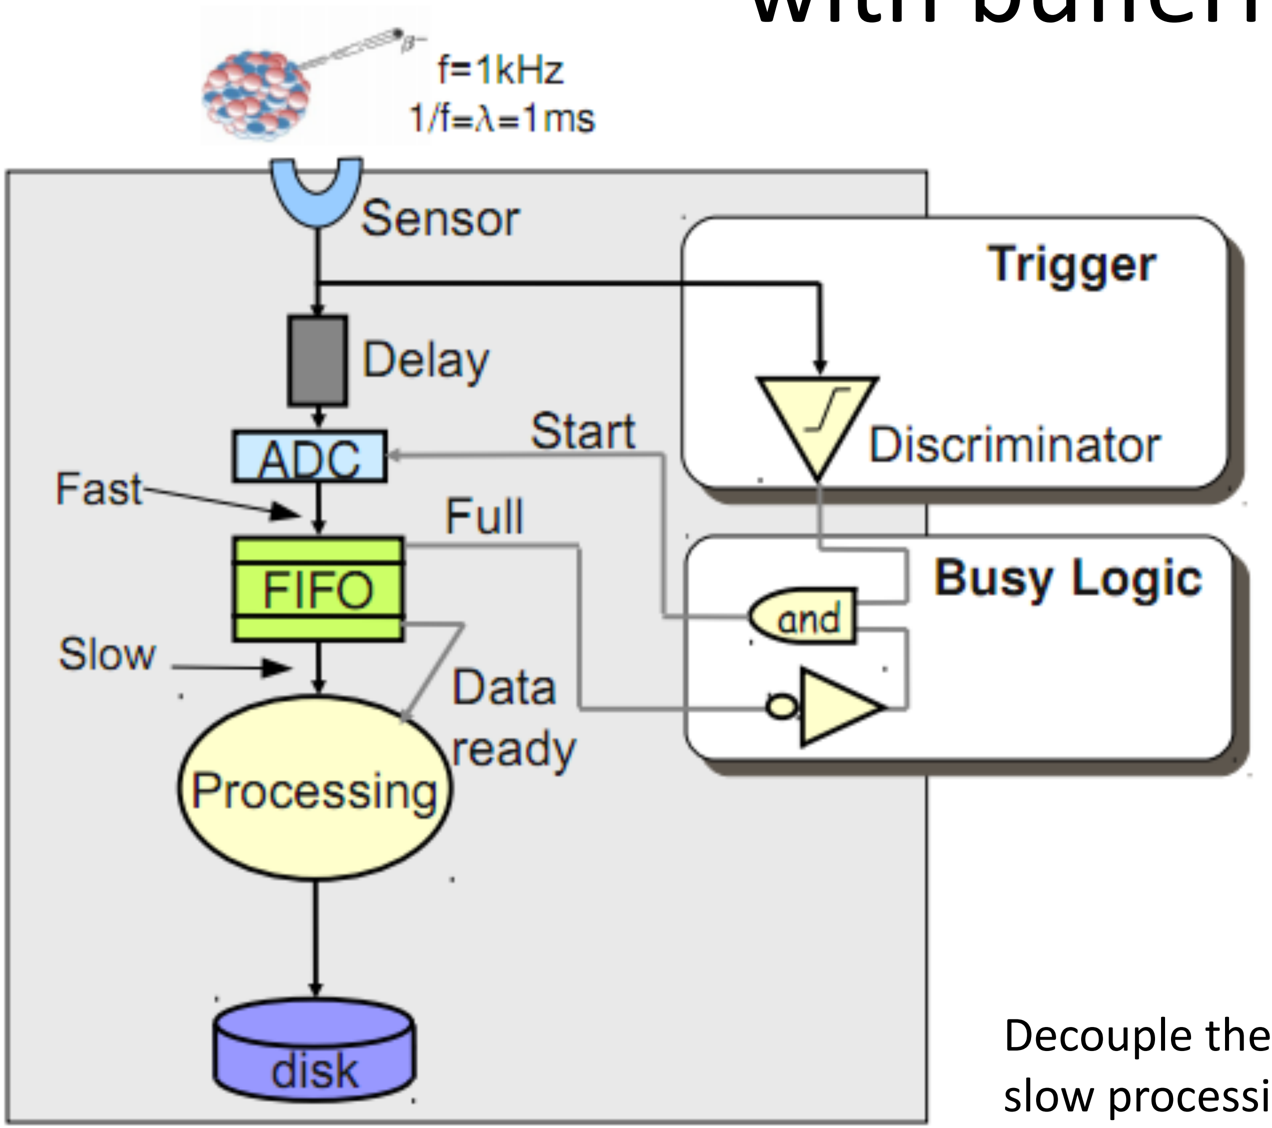

Smooth out fluctuations (derandomize) by introducing an fast, intermediate buffer

Organized as a queue First-In, First-Out (FIFO)

Decouples the fast front-end (ADC) from slow readout

Decouple the fast ADC from the slow processing and readout.

Increasing rejection

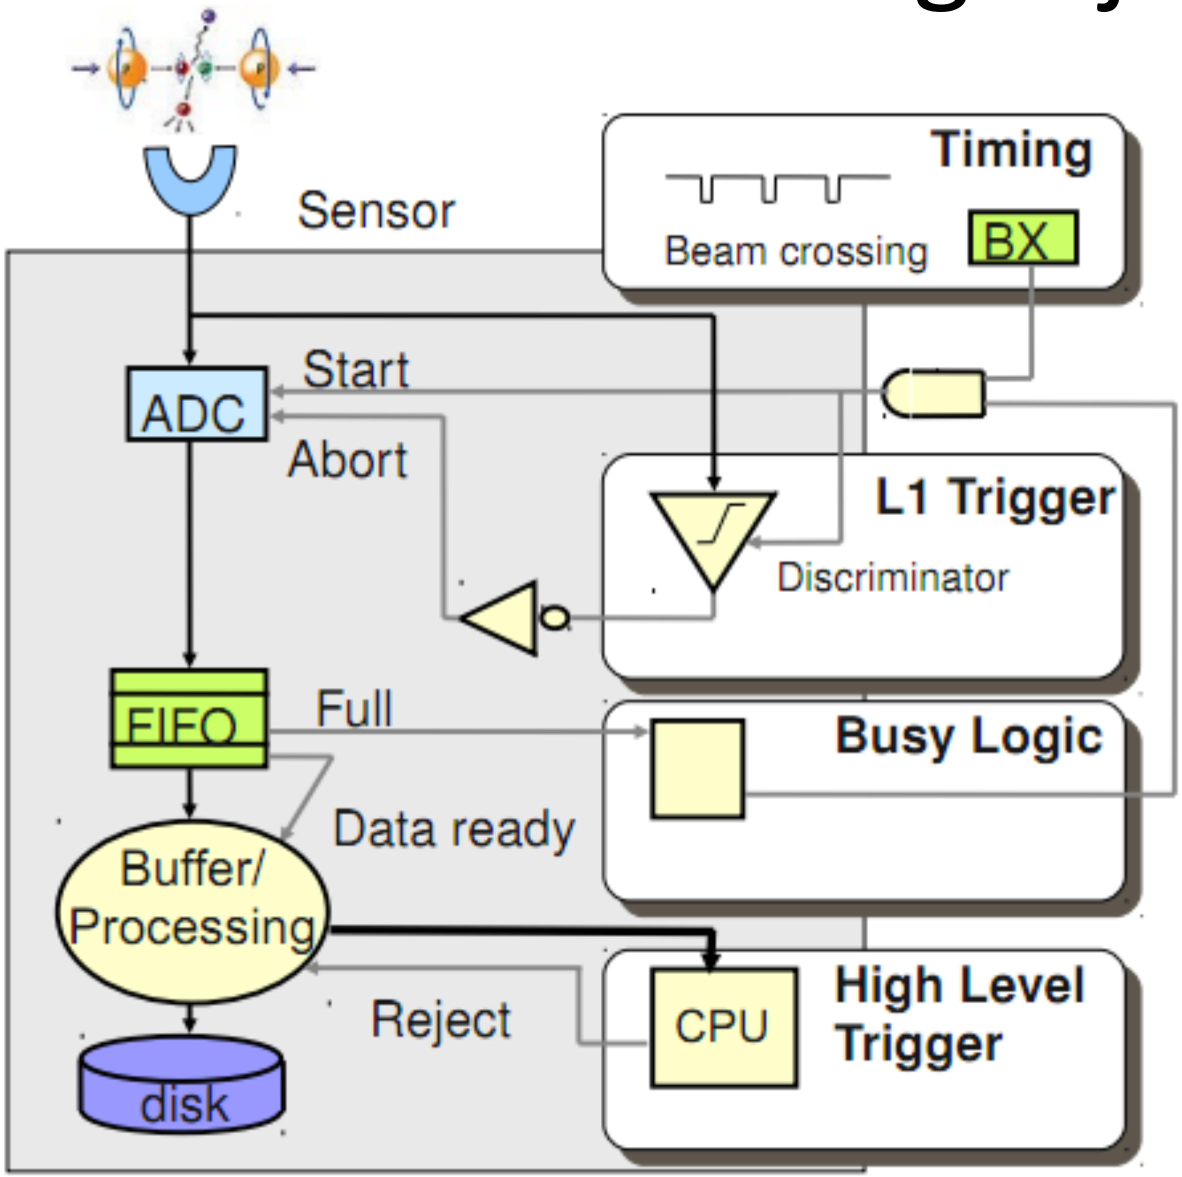

For optimal data reduction can add trigger level between readout and storage (High-level trigger)

Has accessed to some/all processed data

High Level Trigger (HLT) has increased rejection power but longer trigger latency.

Scaling up

Need to impose structure with well-defined interfaces between components