

From collisions to analysis

Lecture 1

Mika Vesterinen
University of Warwick

5-6 September 2019




Lecture 1 outline:

- Introduction
- Trigger and DAQ basics
- Low-level trigger
- High-level trigger

Lecture 2 outline:

- Real-time analysis
- Interaction with analysis
- Upgrades



Acknowledgements to G. Raven and V. Gligorov for their slides from previous CERN/FNAL school lectures on these topics

Collisions to analysis



data reduction

Simulation

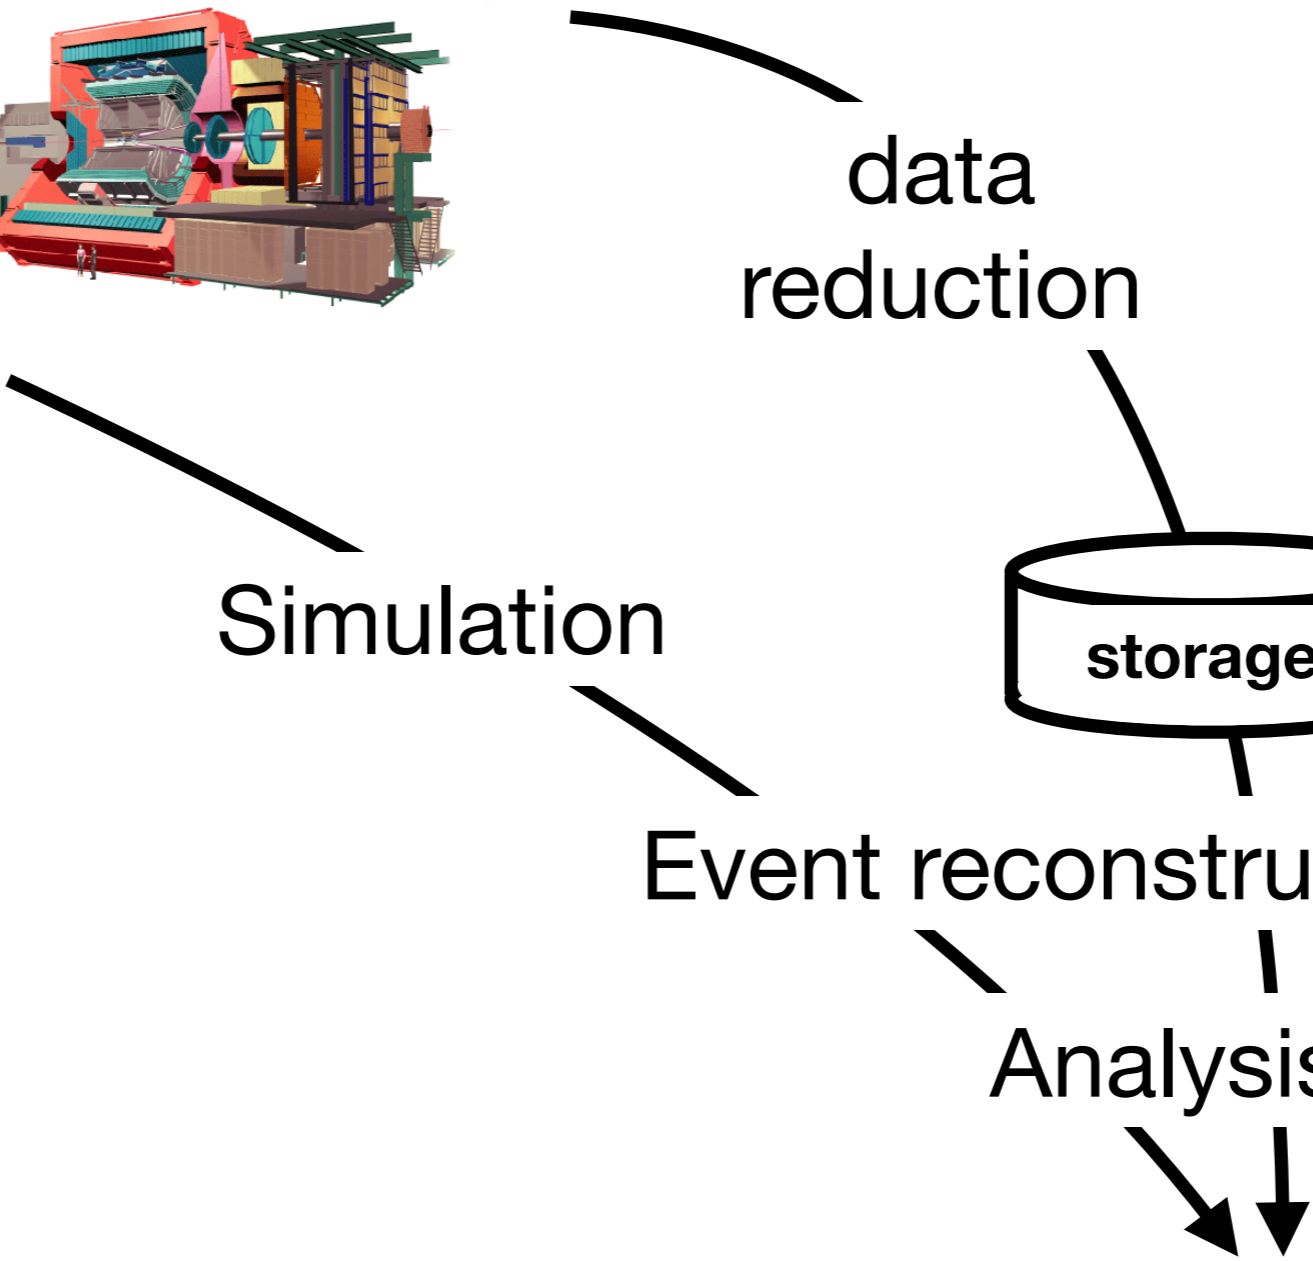


~few PB/year/
experiment
currently

Event reconstruction

Analysis

Physics results



Collisions to analysis



E.g. CMS has $\sim 100\text{M}$ channels, giving a typical event size of $\sim 1\text{ MB}$, and 40 TB/s @ 40 MHz .

Simulation

data
reduction

LHC year
 $\sim 5 \times 10^6\text{ s}$

$2 \times 10^5\text{ PB/yr}$

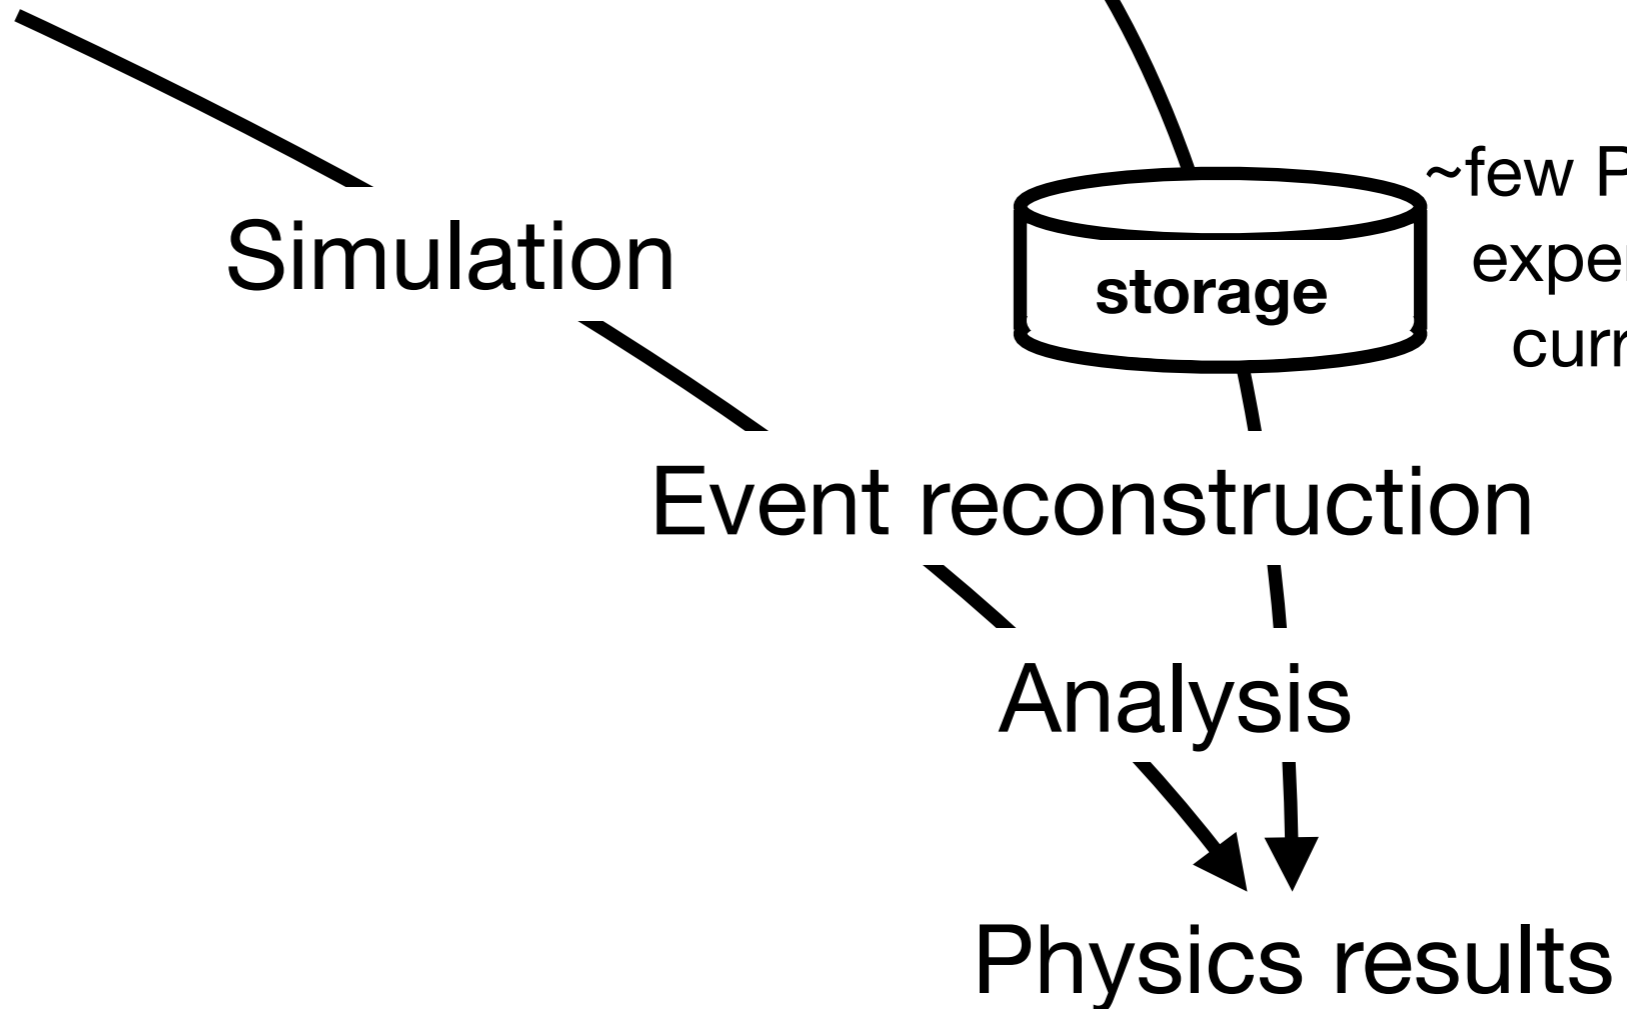
storage

$\sim \text{few PB/year/}$
experiment
currently

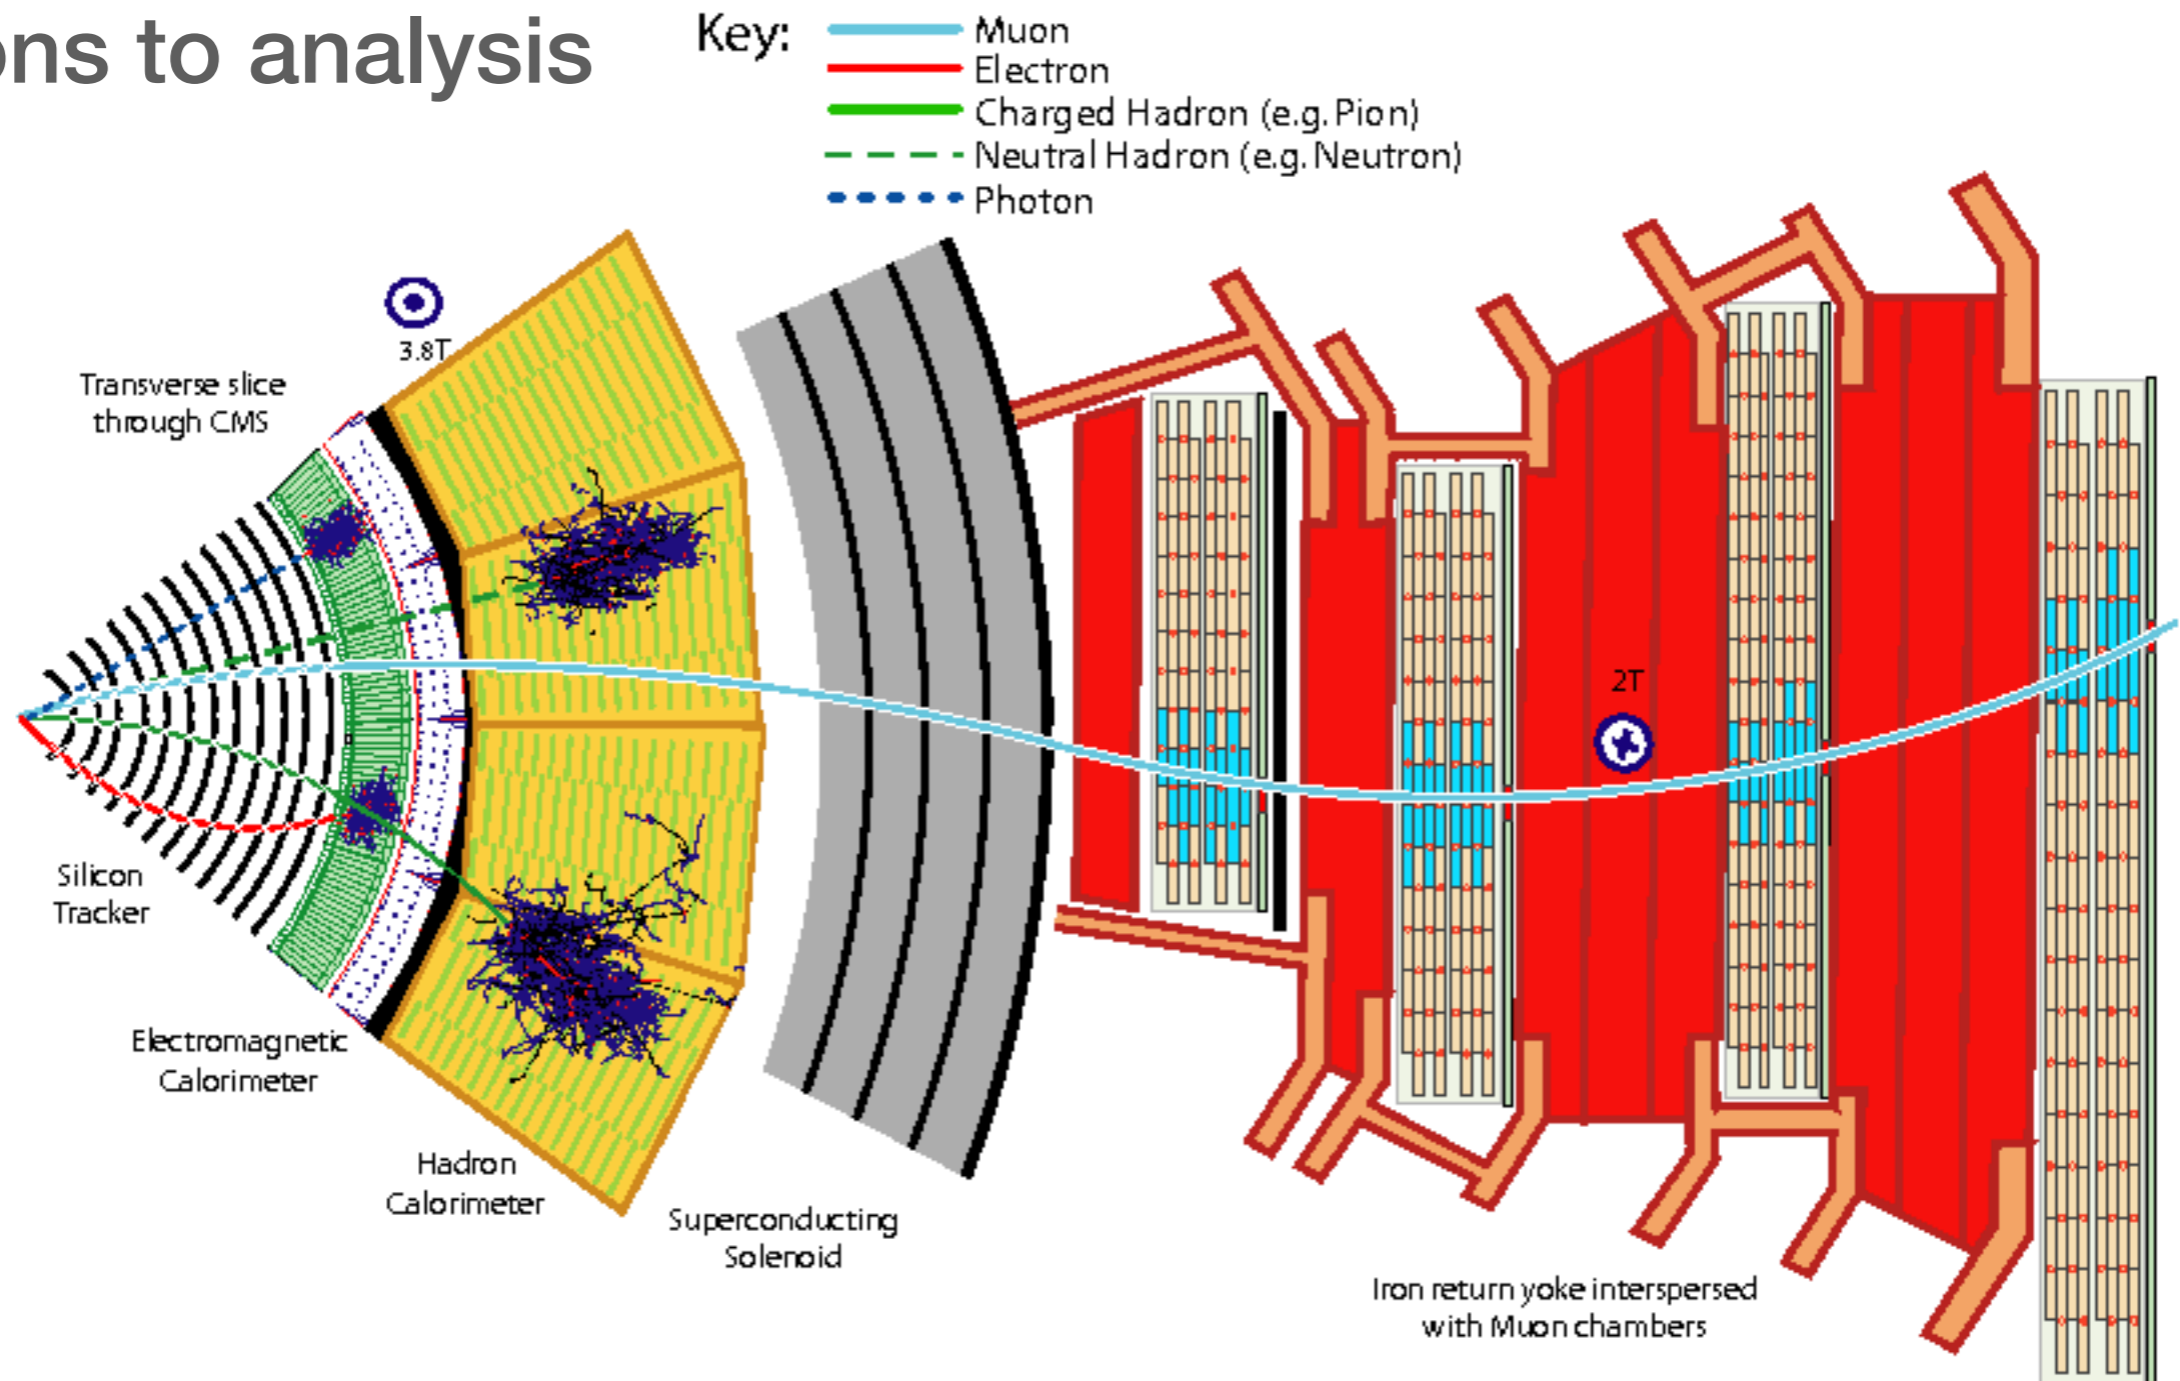
Event reconstruction

Analysis

Physics results



Collisions to analysis





Event reconstruction converts raw data hits to:

- Tracks, ECAL/HCAL clusters.
- e , γ , μ , τ
- Composite objects: missing E_T , H_T , vertices

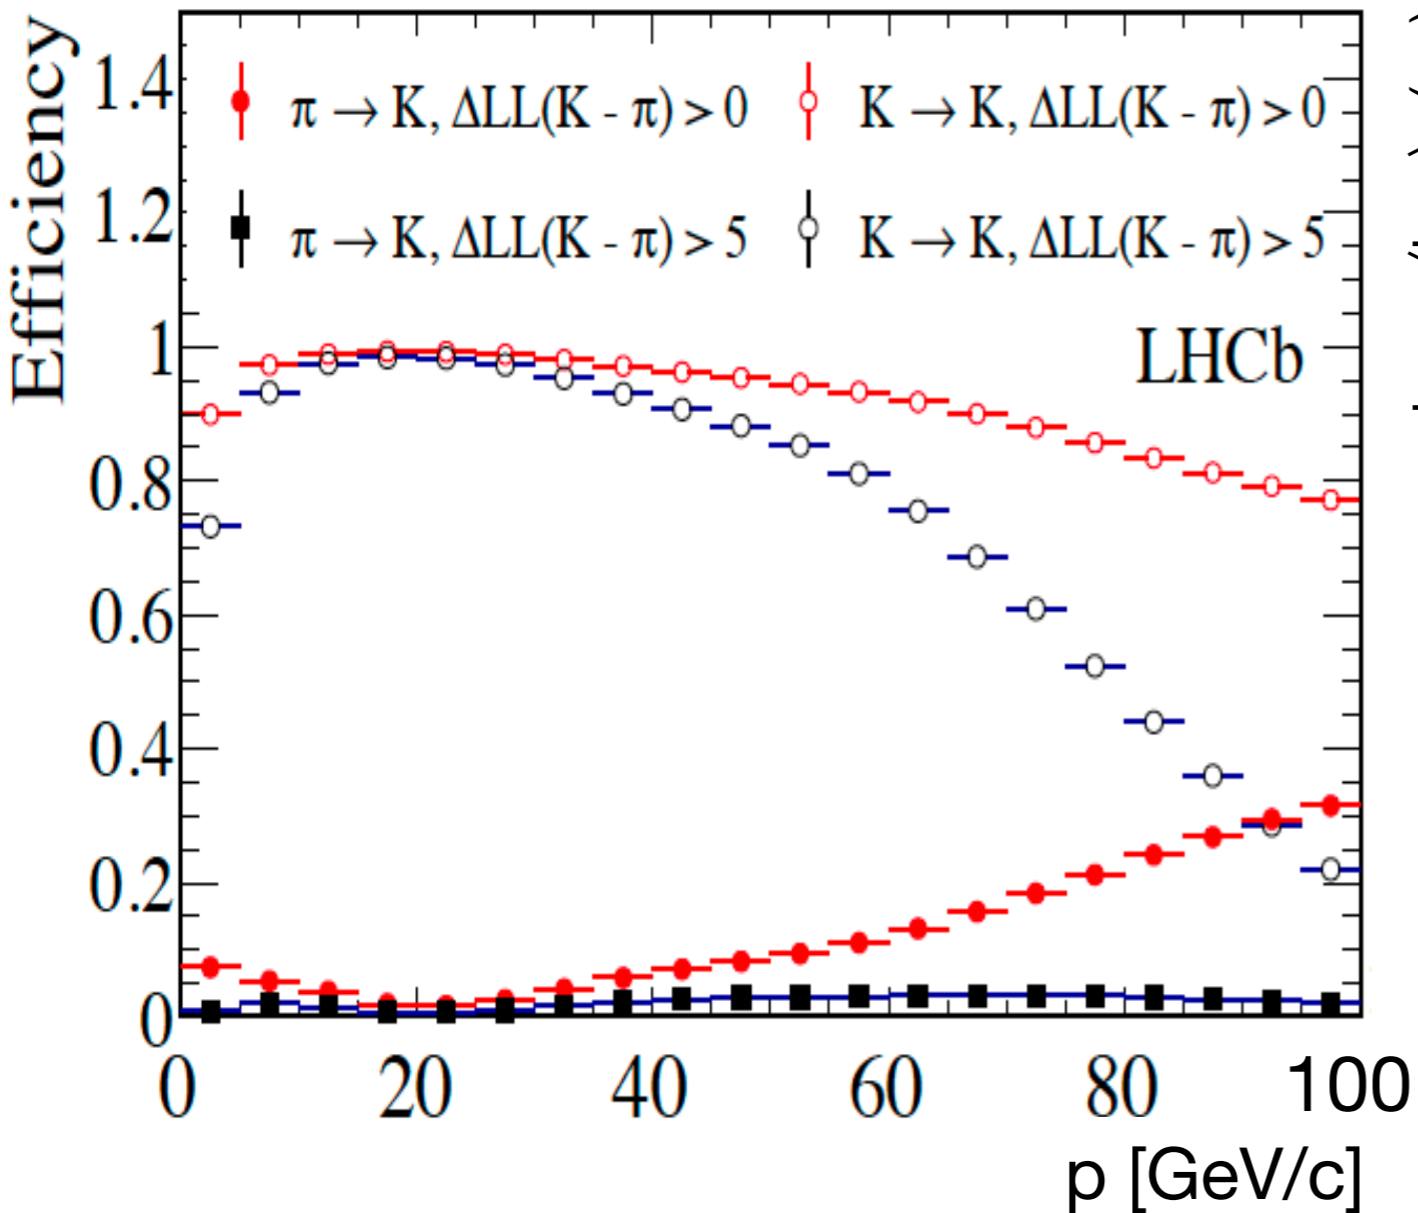
Charged particle identification (π, K, p)

JINST 14 (2019) P04013




Charged particle identification (π, K, p)

JINST 14 (2019) P04013

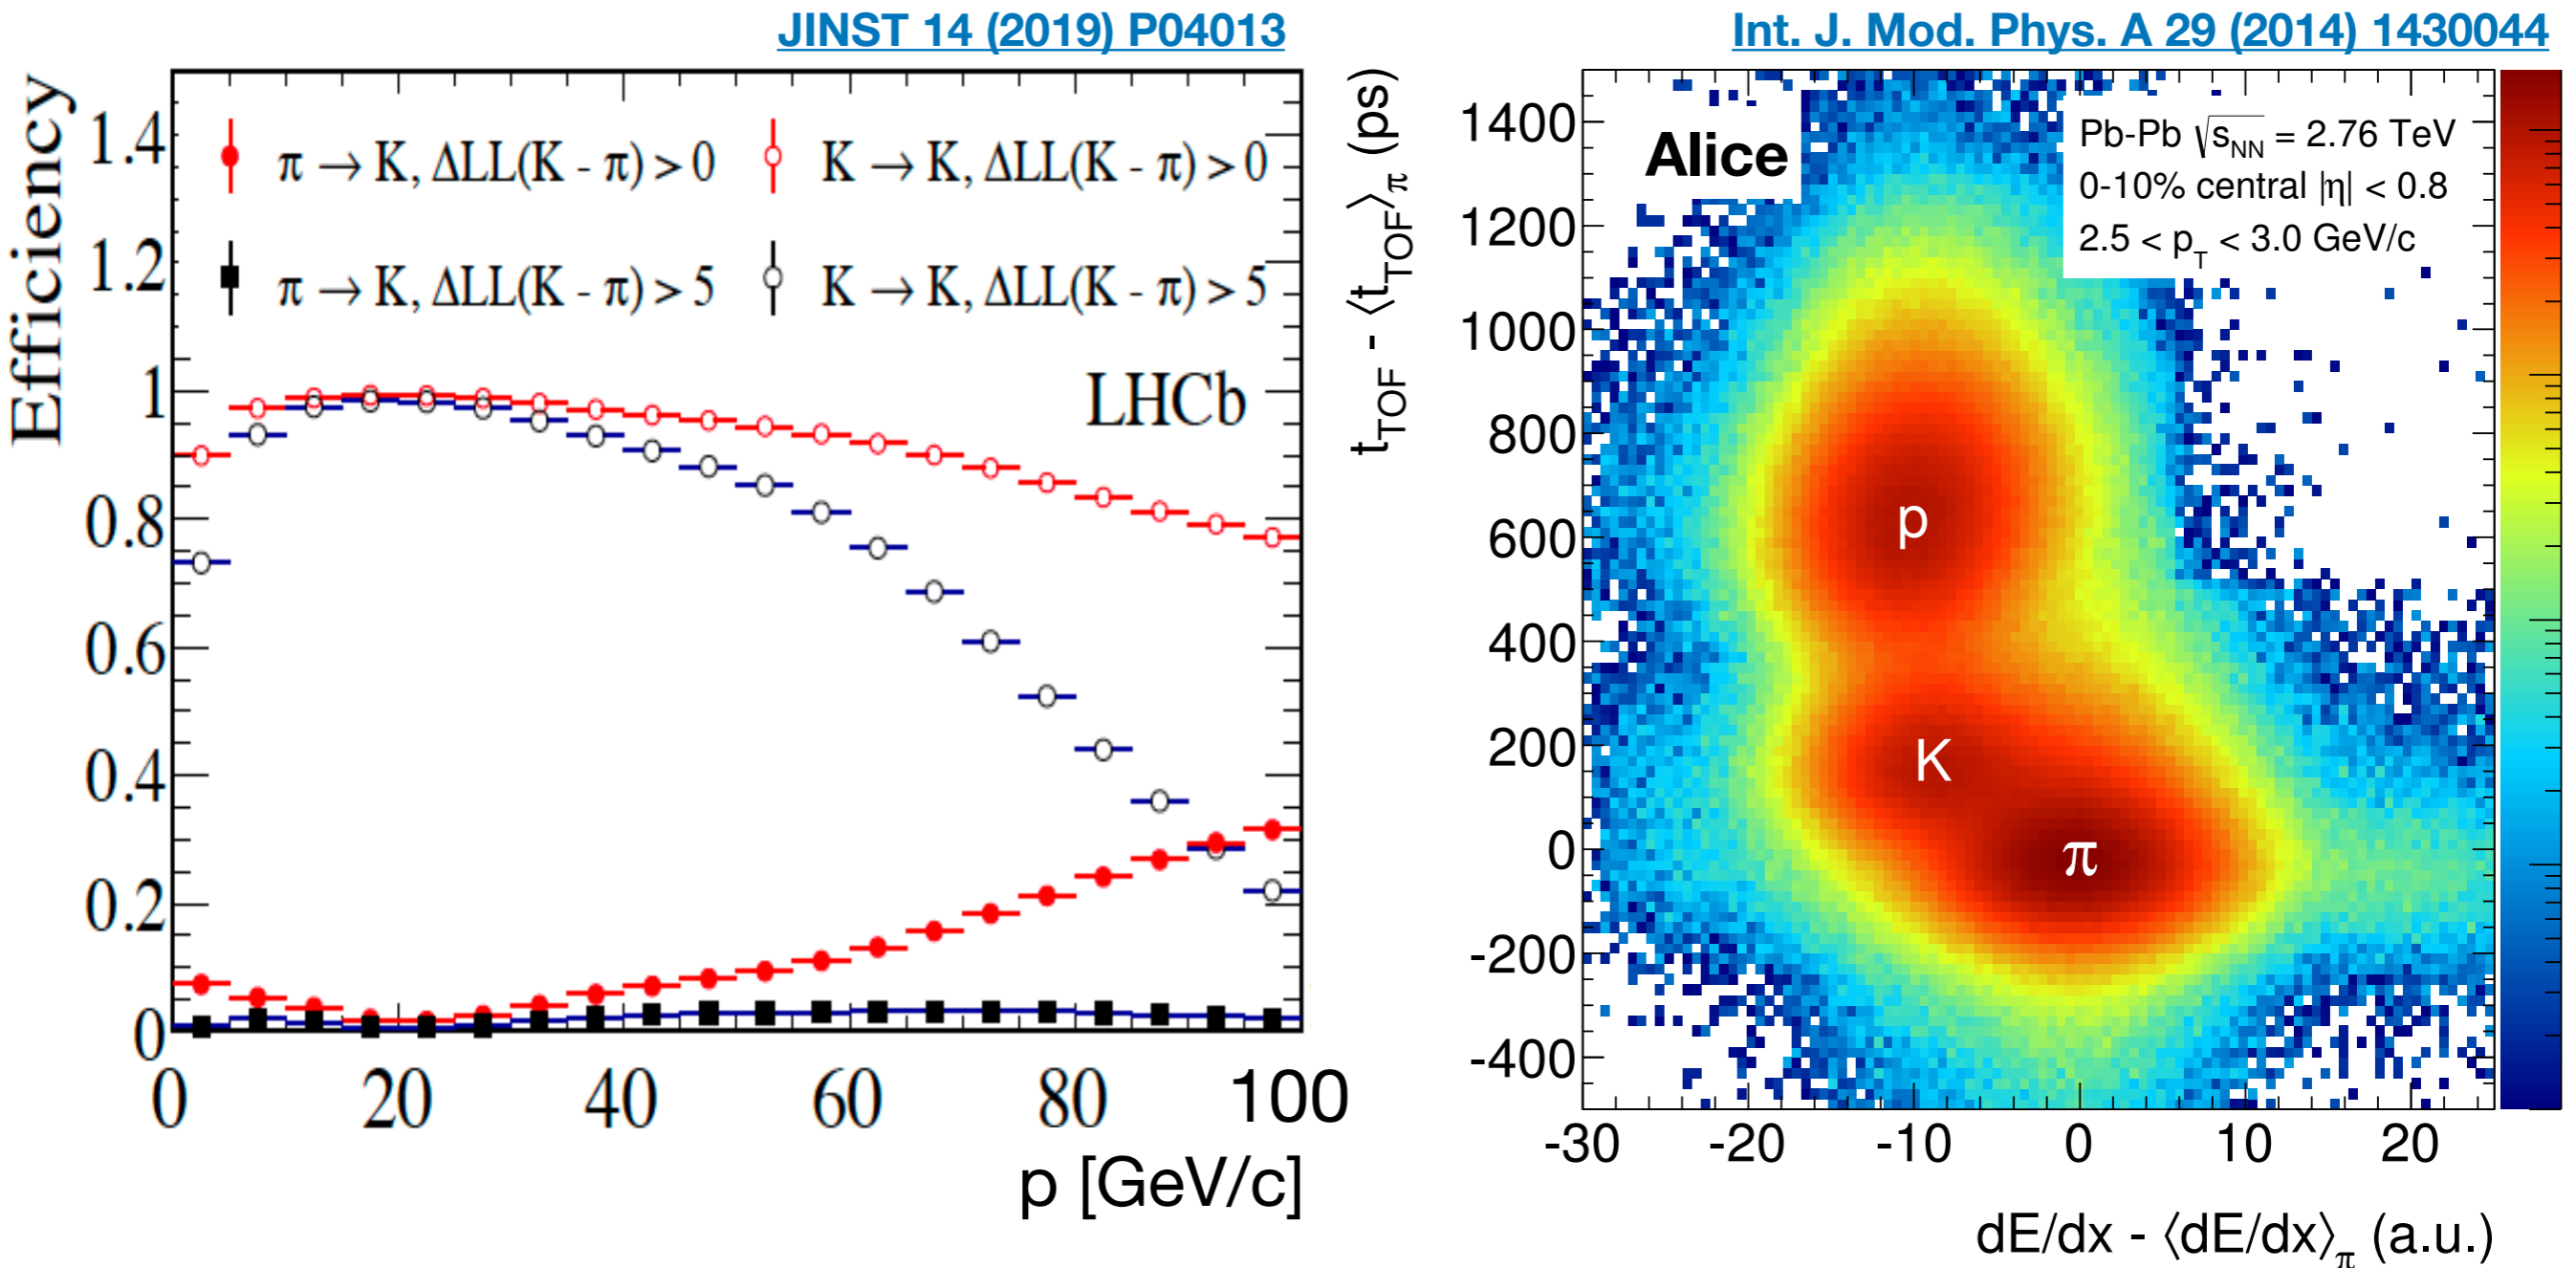


Int. J. Mod. Phys. A 29 (2014) 1430044



Requires precise alignment and calibration of the detector.

Charged particle identification (π, K, p)

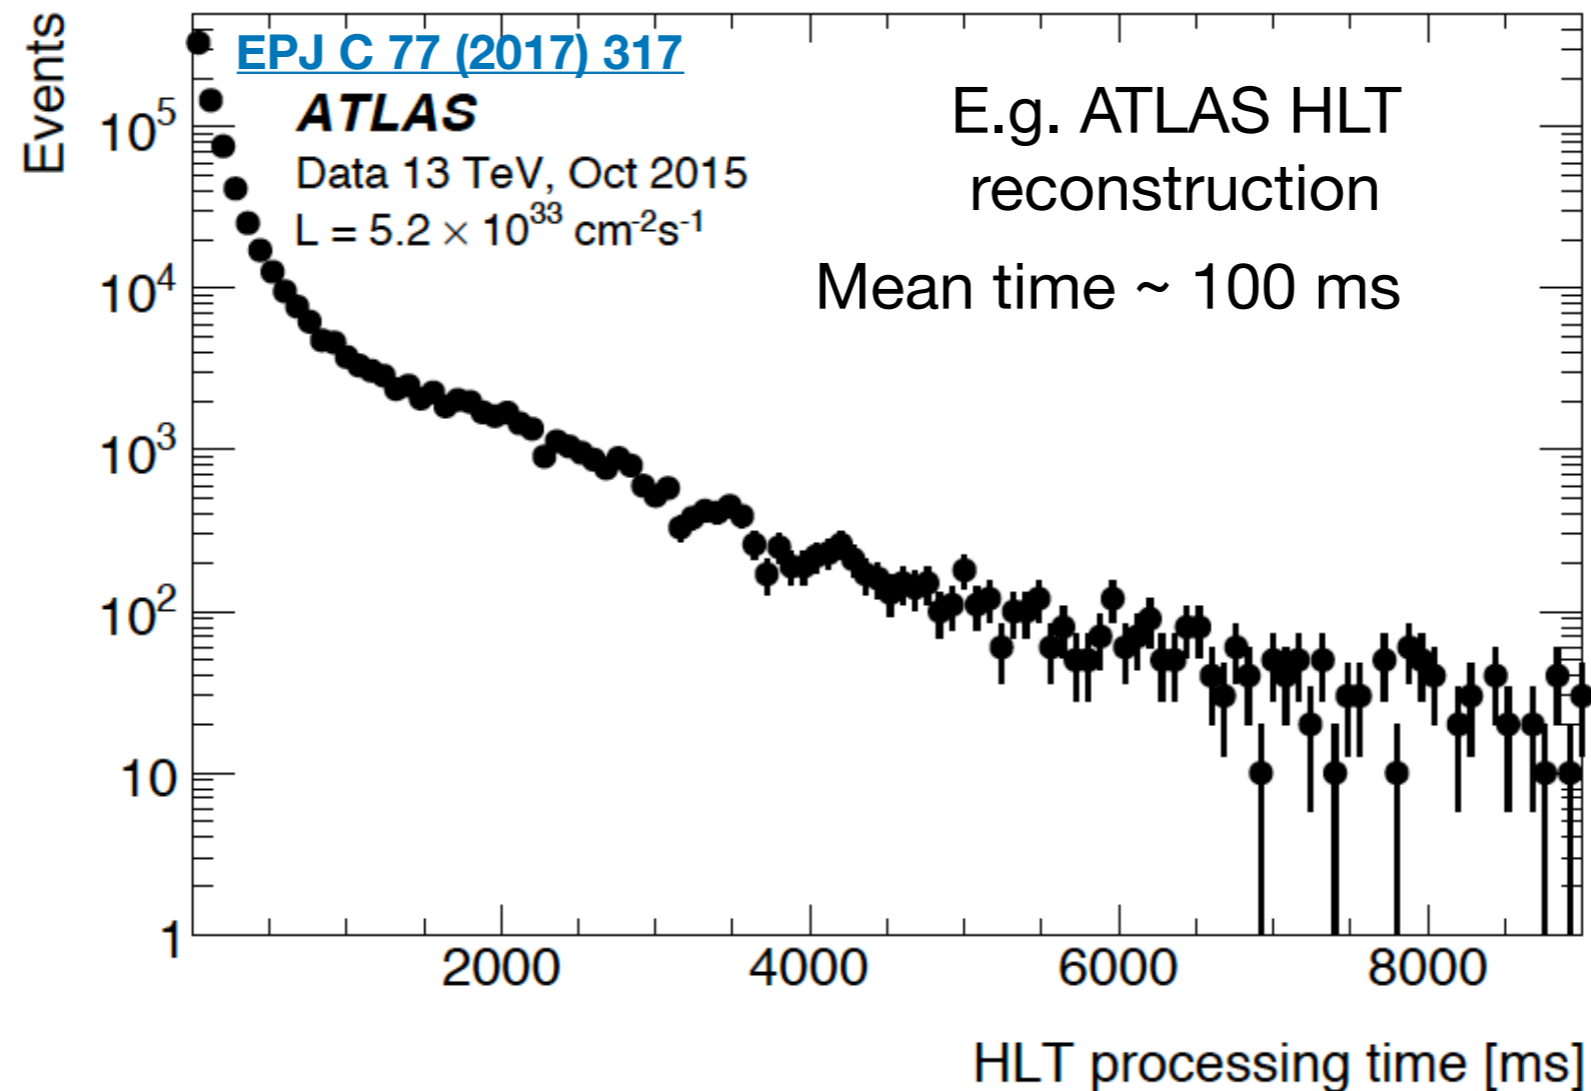


Requires precise alignment and calibration of the detector.

Full event reconstruction involved complicated algorithms, which are typically best suited to the high flexibility of CPUs.

Cost of full event reconstruction

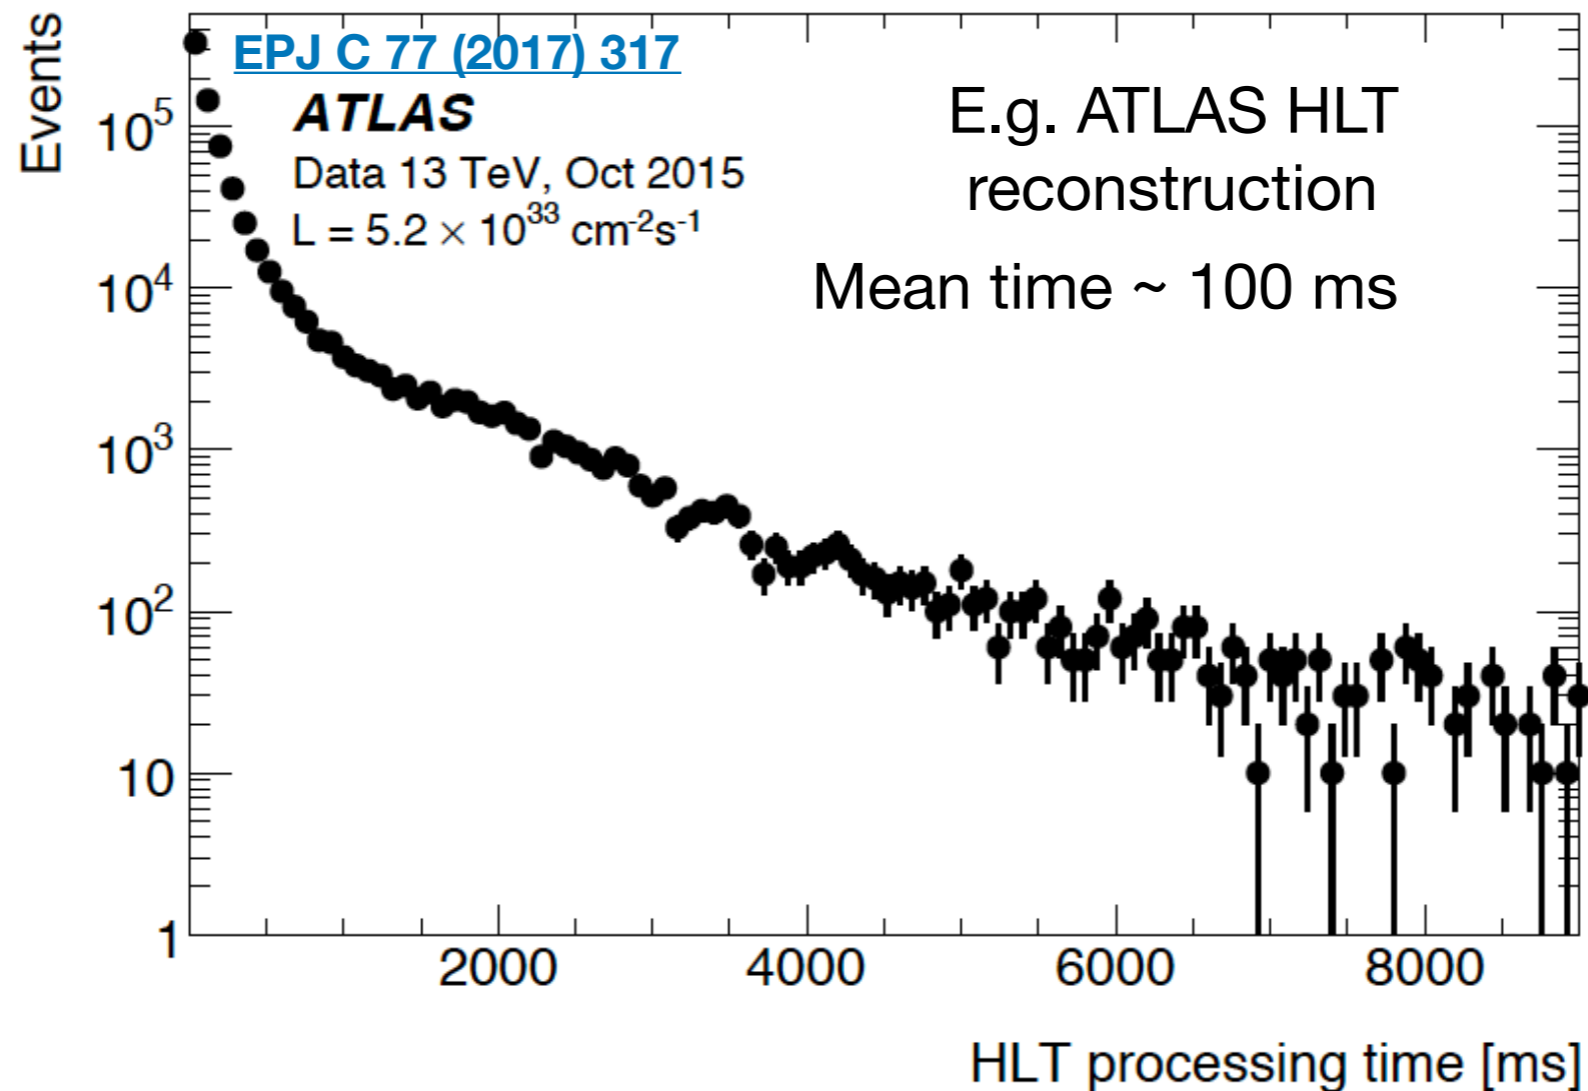
A ballpark figure for LHC experiments is 1 second / CPU process, but compromises can be made in exchange for speed, e.g.:



Could we run this on 30 MHz of bunch crossings?

Cost of full event reconstruction

A ballpark figure for LHC experiments is 1 second / CPU process, but compromises can be made in exchange for speed, e.g.:




Could we run this on 30 MHz of bunch crossings?

Requires 3 million CPU cores - not affordable!

Cost of full event reconstruction

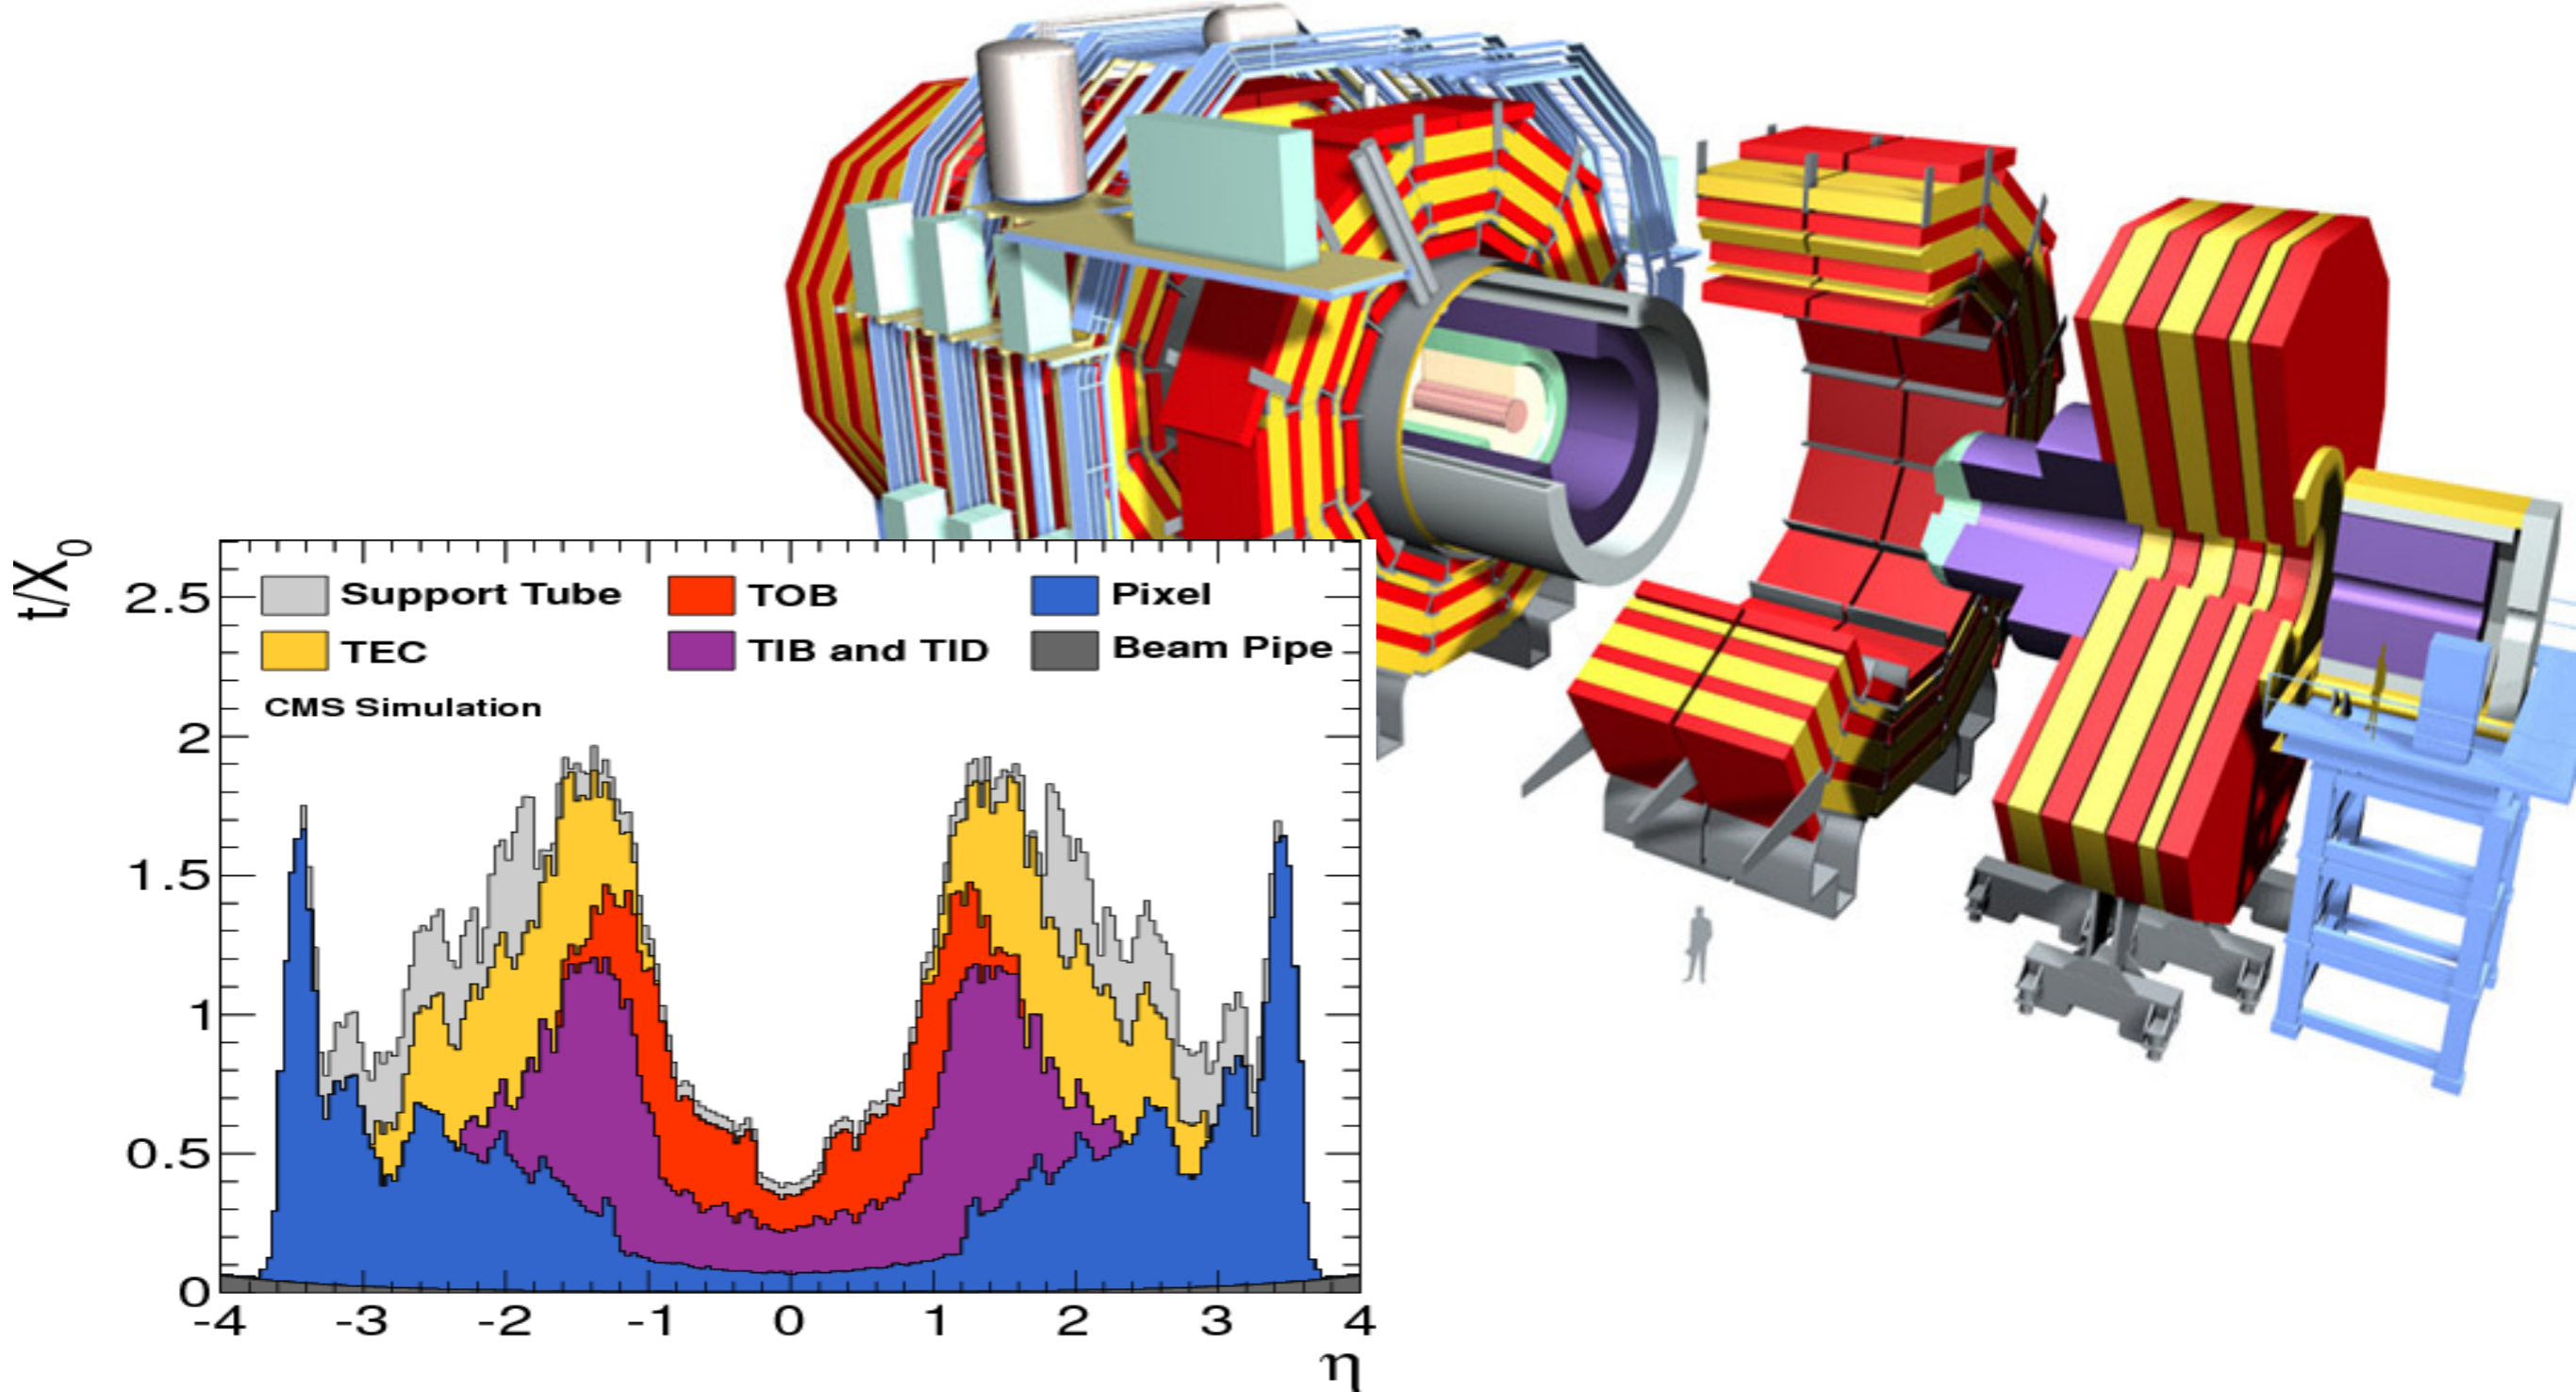
A ballpark figure for LHC experiments is 1 second / CPU process, but compromises can be made in exchange for speed, e.g.:



Could we run this on 30 MHz of bunch crossings?

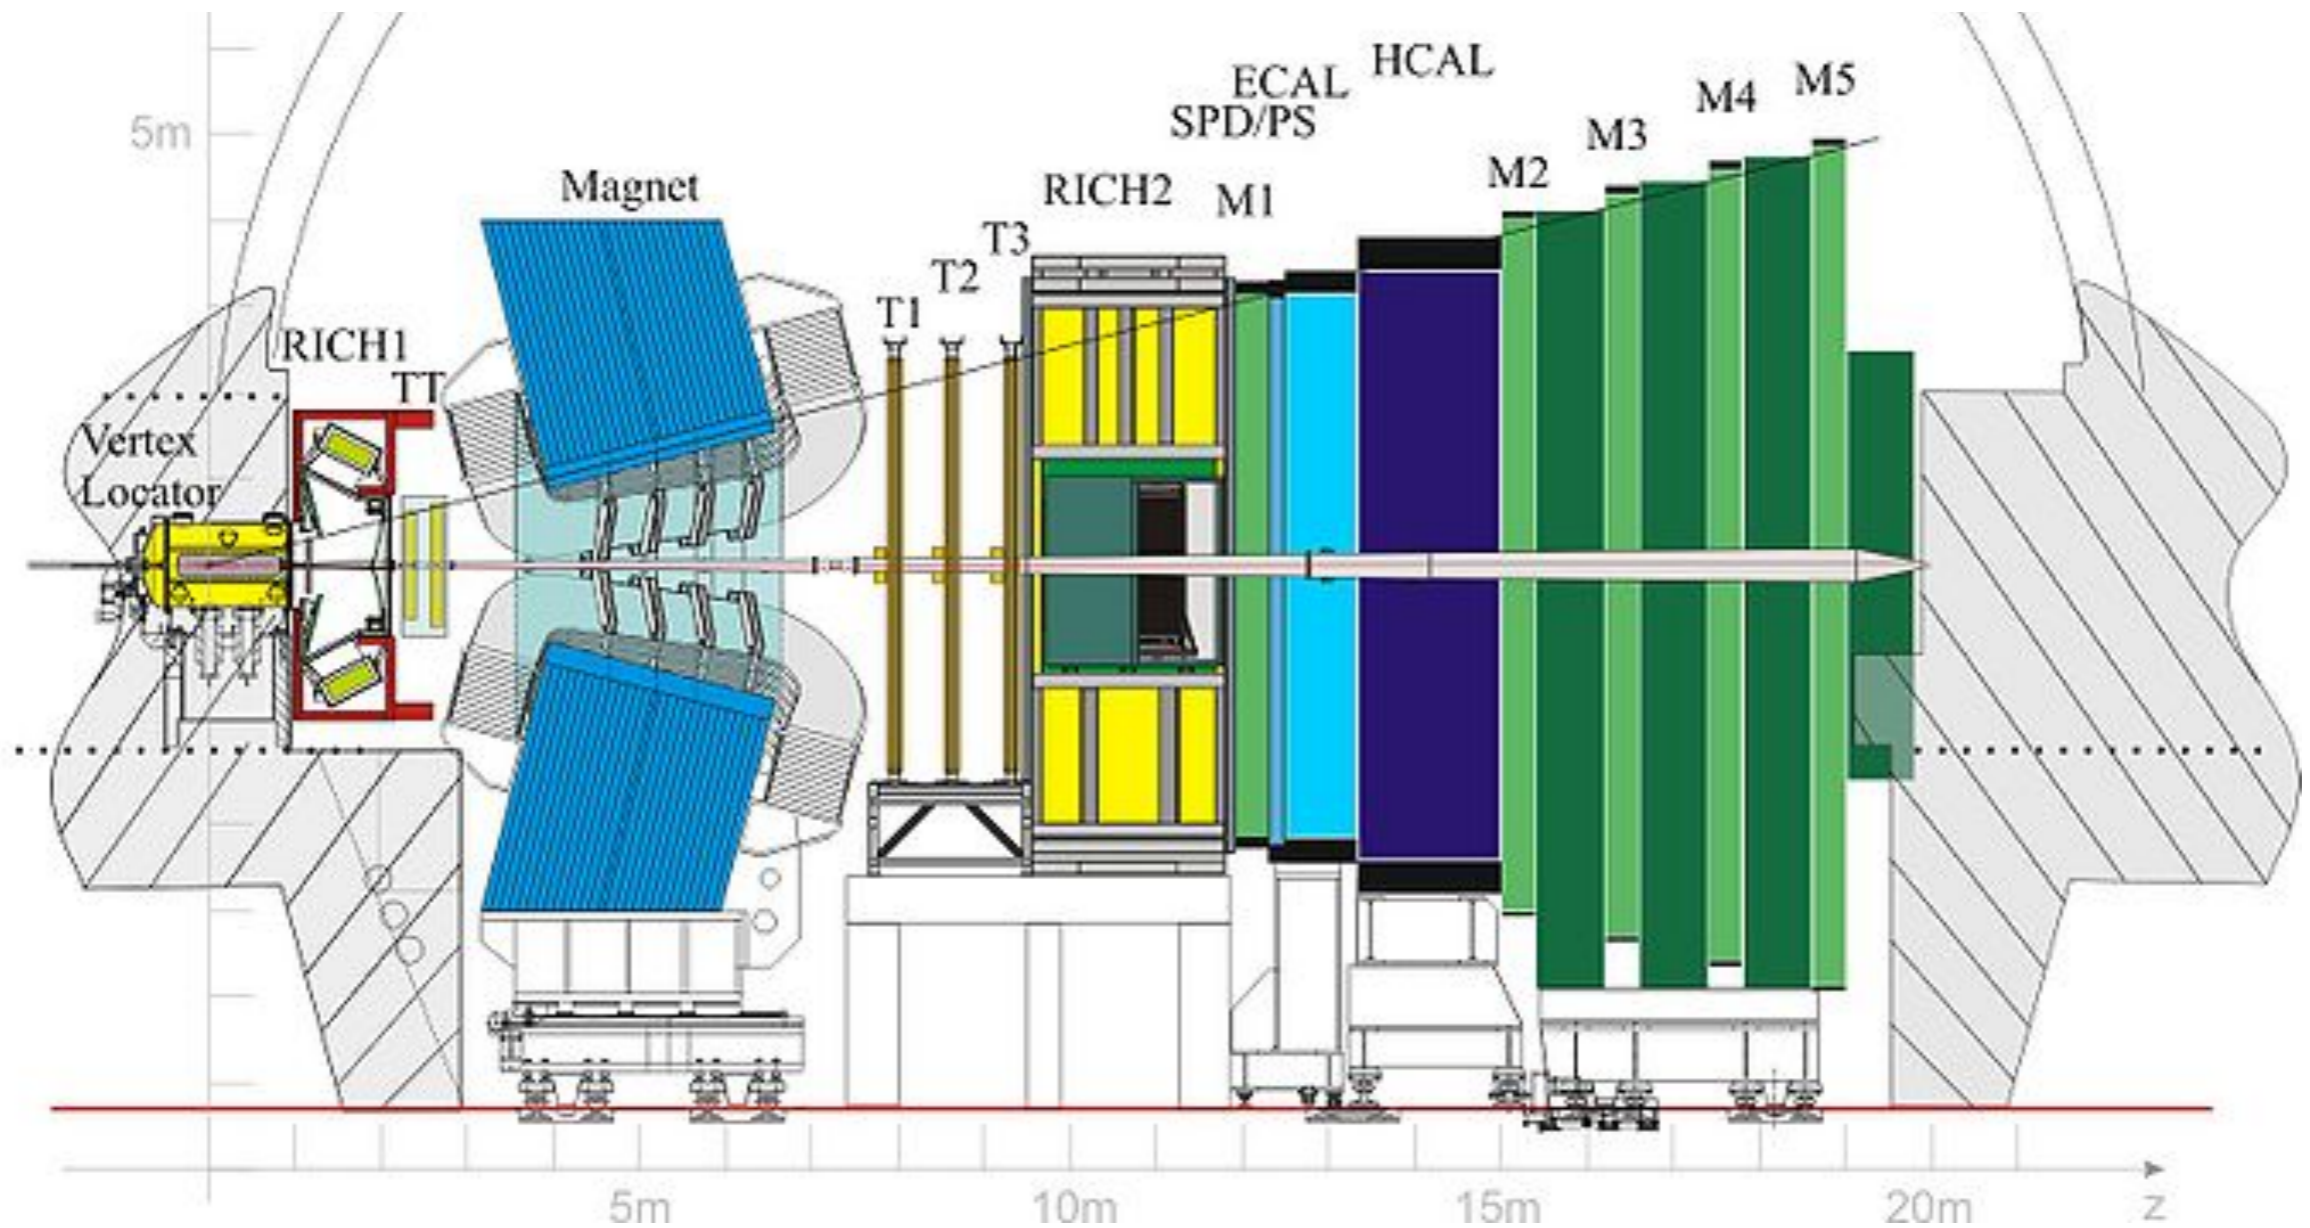
Requires 3 million CPU cores - not affordable!

Note: not all CPU processes have the same speed, so better to talk in terms of *throughput* of a fully loaded multicore PC.



Even if we could afford to process/store all of the raw data, we couldn't power/cool the necessary electronics without spoiling the material budget.


Not such a constraint for LHCb



Its smaller events, and unique geometry, allow the LHCb upgrade to have trigger-less readout at a luminosity of $2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. [LHCb-TDR-016](#)

Or for the relatively lower event rates in PbPb

The Alice upgrade will have continuous readout at 50 kHz in PbPb collisions.




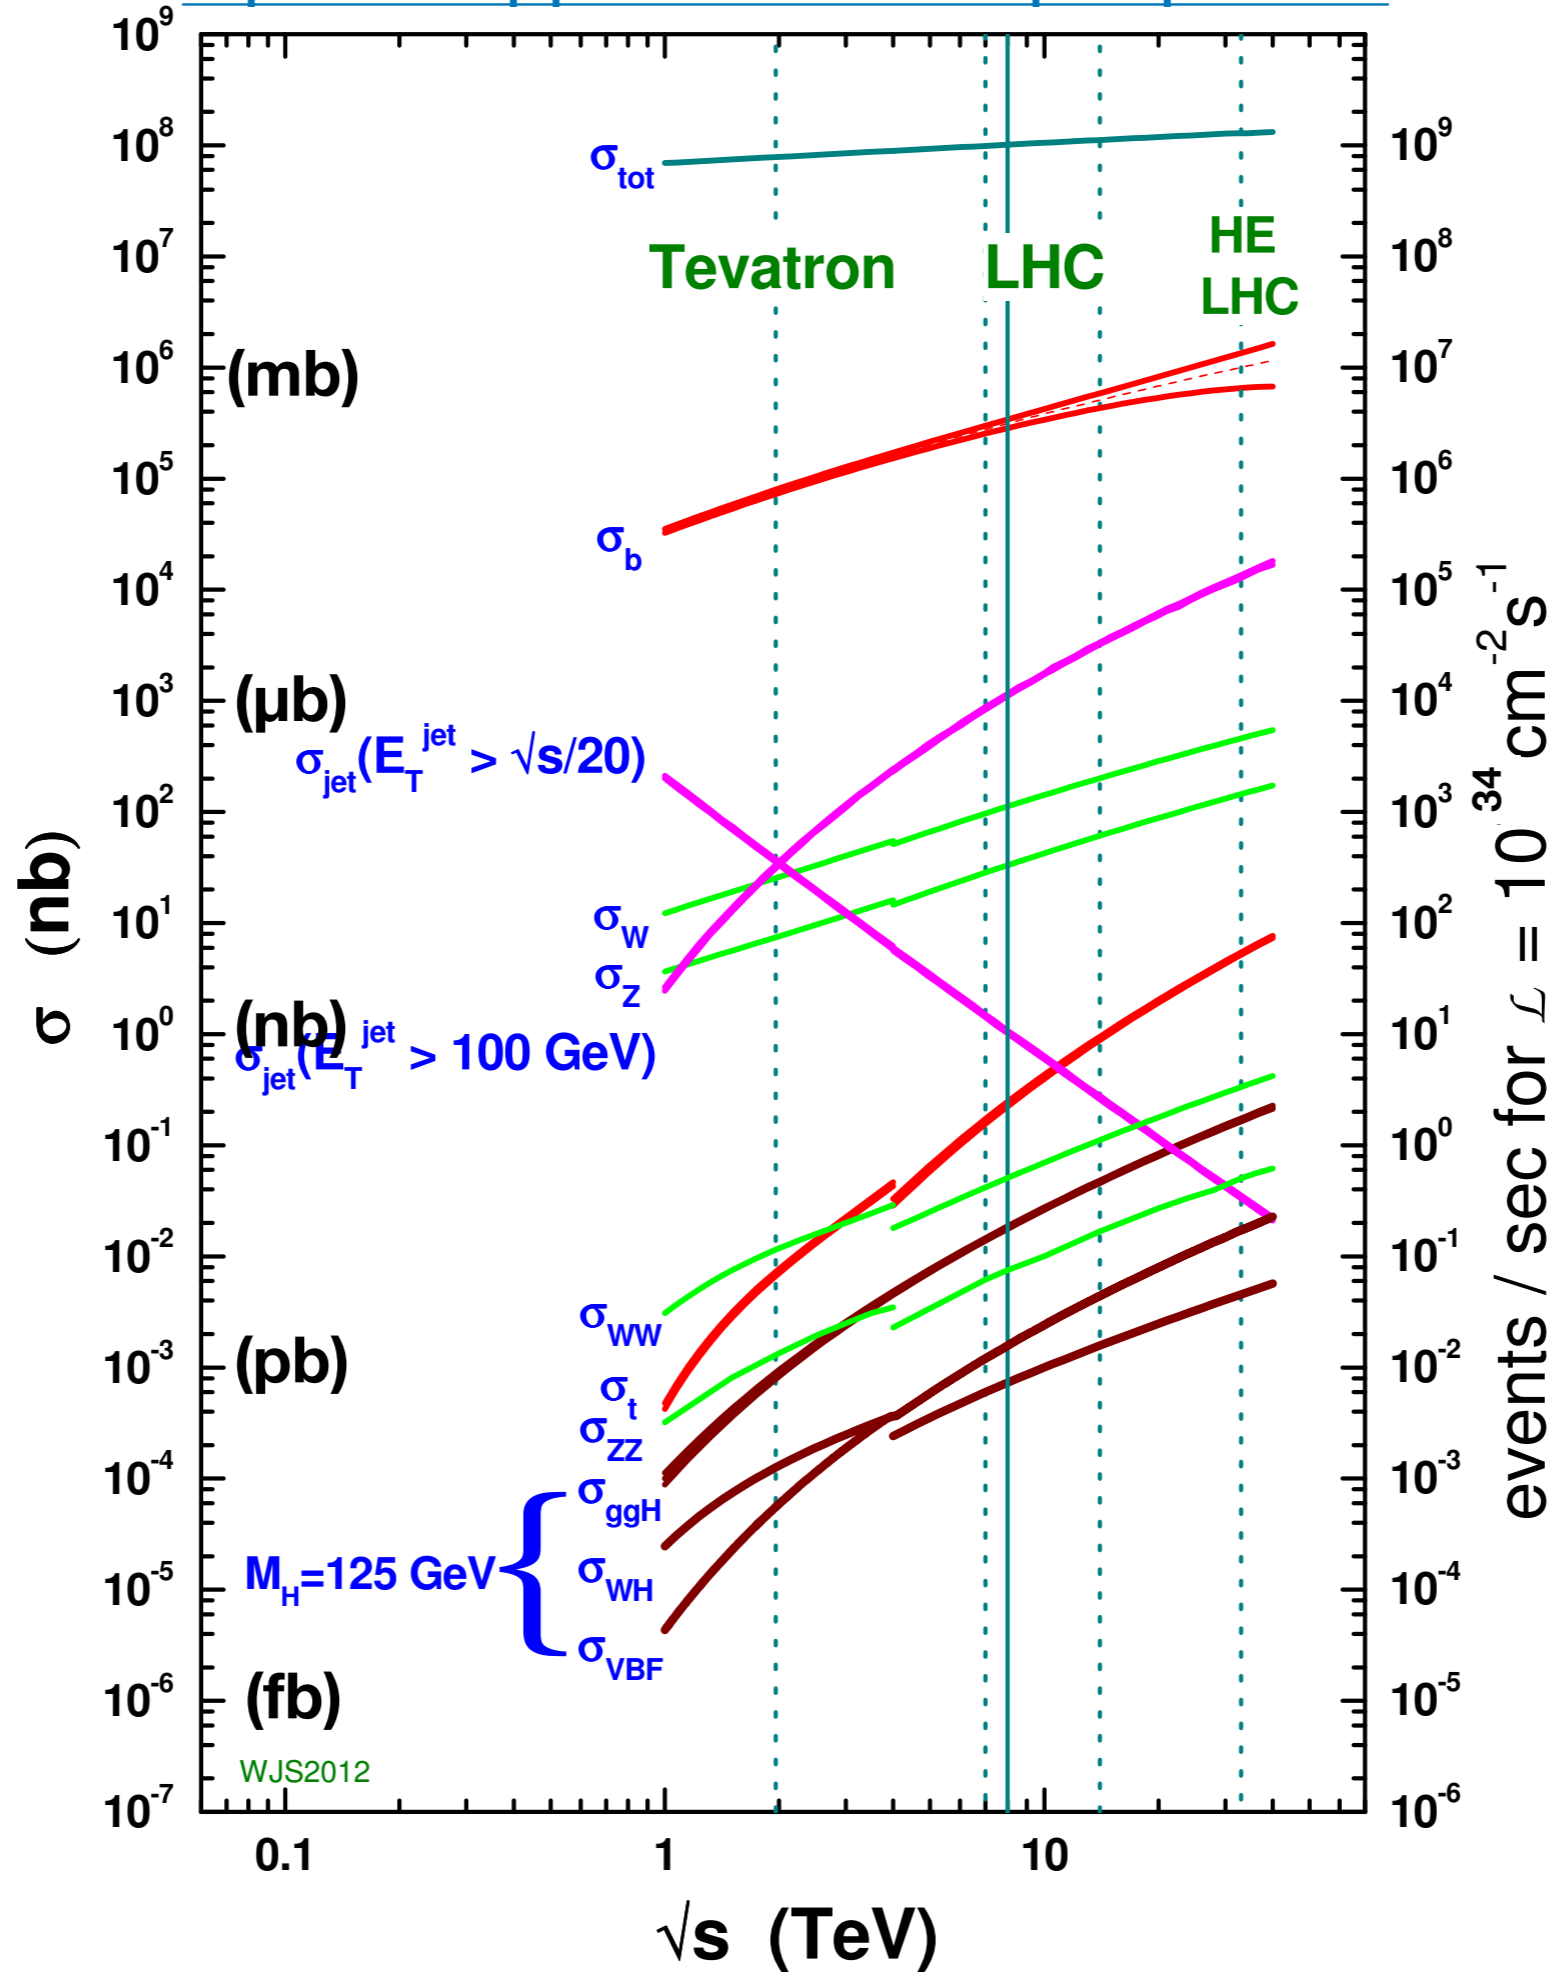
Detector	Input to Online System (GByte/s)	Peak Output to Local Data Storage (GByte/s)
TPC	1000	50.0
TRD	81.5	10.0
ITS	40	10.0
Others	25	12.5
Total	1146.5	82.5

Collisions to analysis



E.g. CMS has $\sim 100\text{M}$ channels, giving a typical event size of $\sim 1\text{ MB}$, and 40 TB/s @ 40 MHz .





1 GHz

1 kHz

events / sec for $\mathcal{L} = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$

What do the GHz of background interactions look like?

$$\frac{1}{N_{\text{ev}}} \frac{dN_{\text{charged}}}{d\eta} \sim 5$$

~25/interaction in $|\eta| < 2.5$


What do the GHz of background interactions look like?

$$\frac{1}{N_{\text{ev}}} \frac{dN_{\text{charged}}}{d\eta} \sim 5$$

~25/interaction in $|\eta| < 2.5$

The tracks are *soft*:

- 97% have $p_T < 2$ GeV
- 99% have $p_T < 3$ GeV
- etc..



What do the GHz of background interactions look like?


$$\frac{1}{N_{\text{ev}}} \frac{dN_{\text{charged}}}{d\eta} \sim 5$$

~25/interaction in $|\eta| < 2.5$


The tracks are *soft*:

- 97% have $p_T < 2$ GeV
- 99% have $p_T < 3$ GeV
- etc..

@ pileup ~ 30, there are still
~10 charged particles
with $p_T > 3$ GeV.



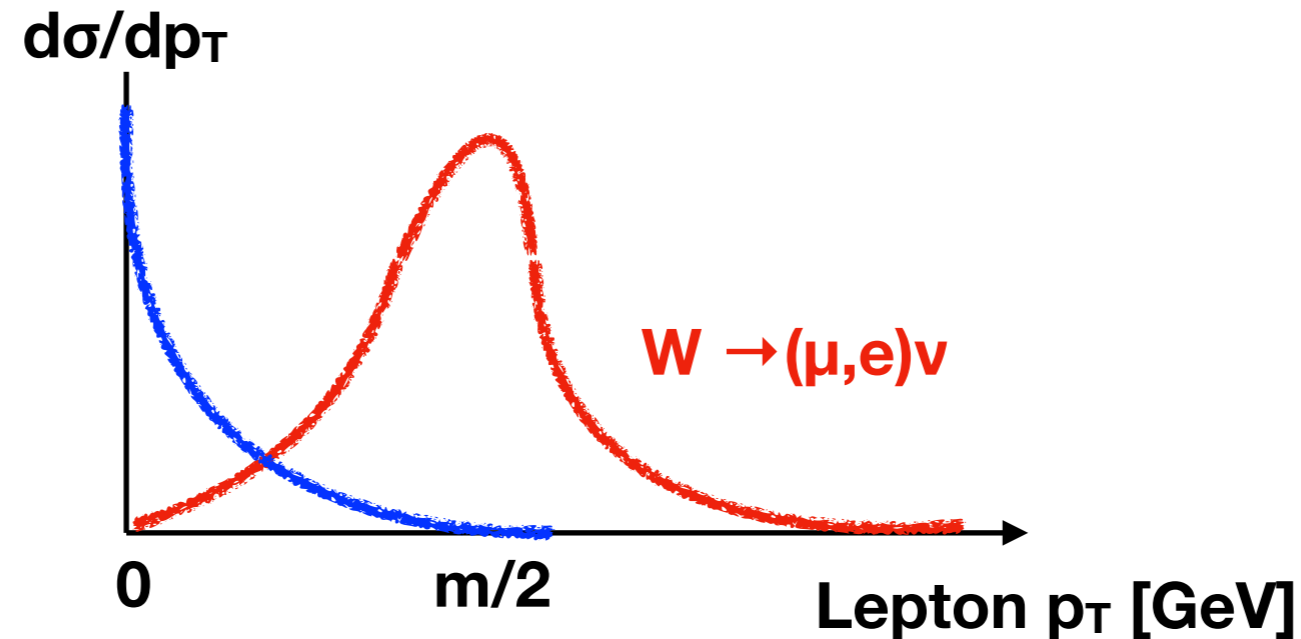
Mass and p_T



Final state particles get softer with higher multiplicity decays and/or more complicated cascades.

Signals without any leptons will always be difficult.

Mass and p_T



Final state particles get softer with higher multiplicity decays and/or more complicated cascades.

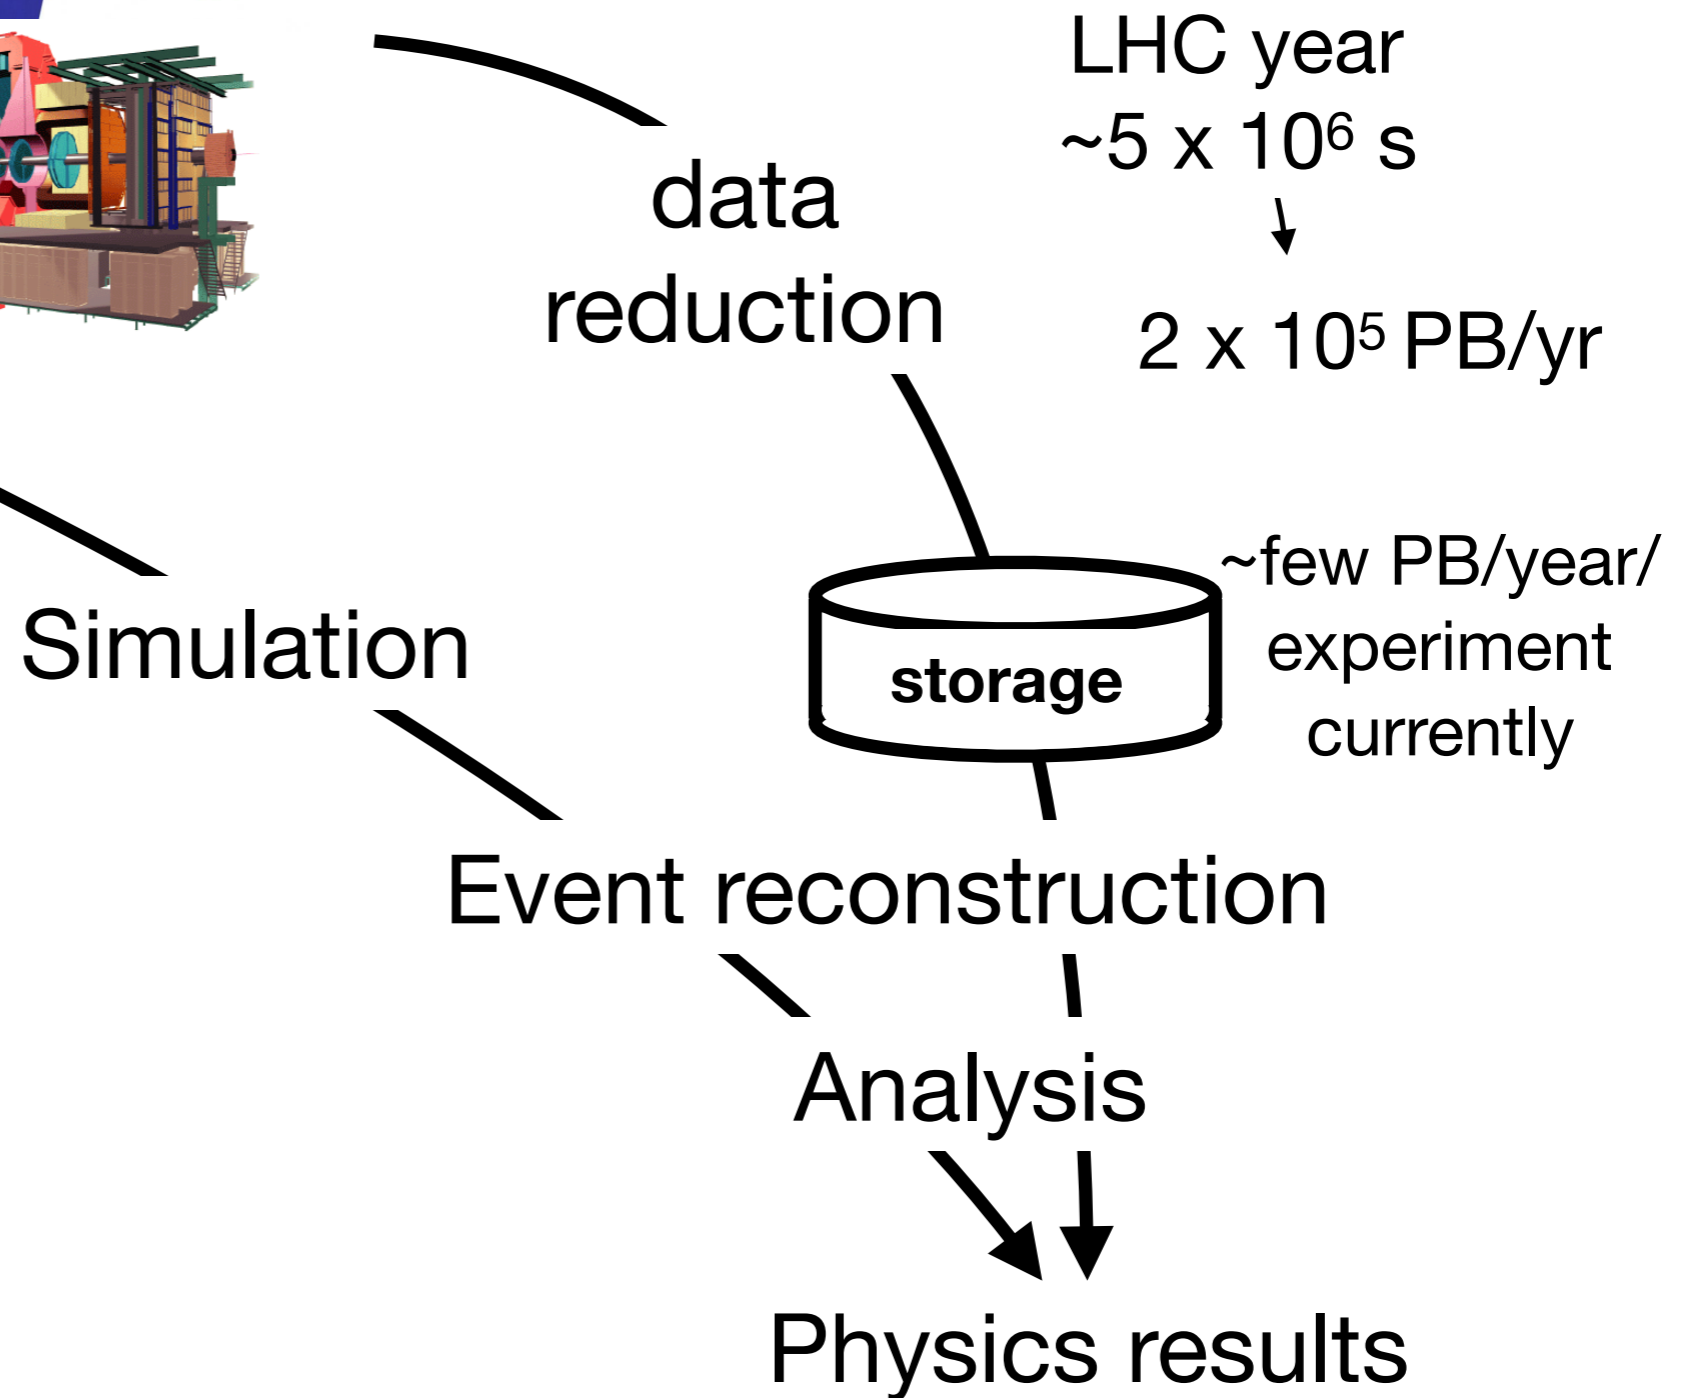
Signals without any leptons will always be difficult.

At an absolute minimum we must have single lepton triggers with p_T thresholds below ~ 25 GeV, without *prescales*.

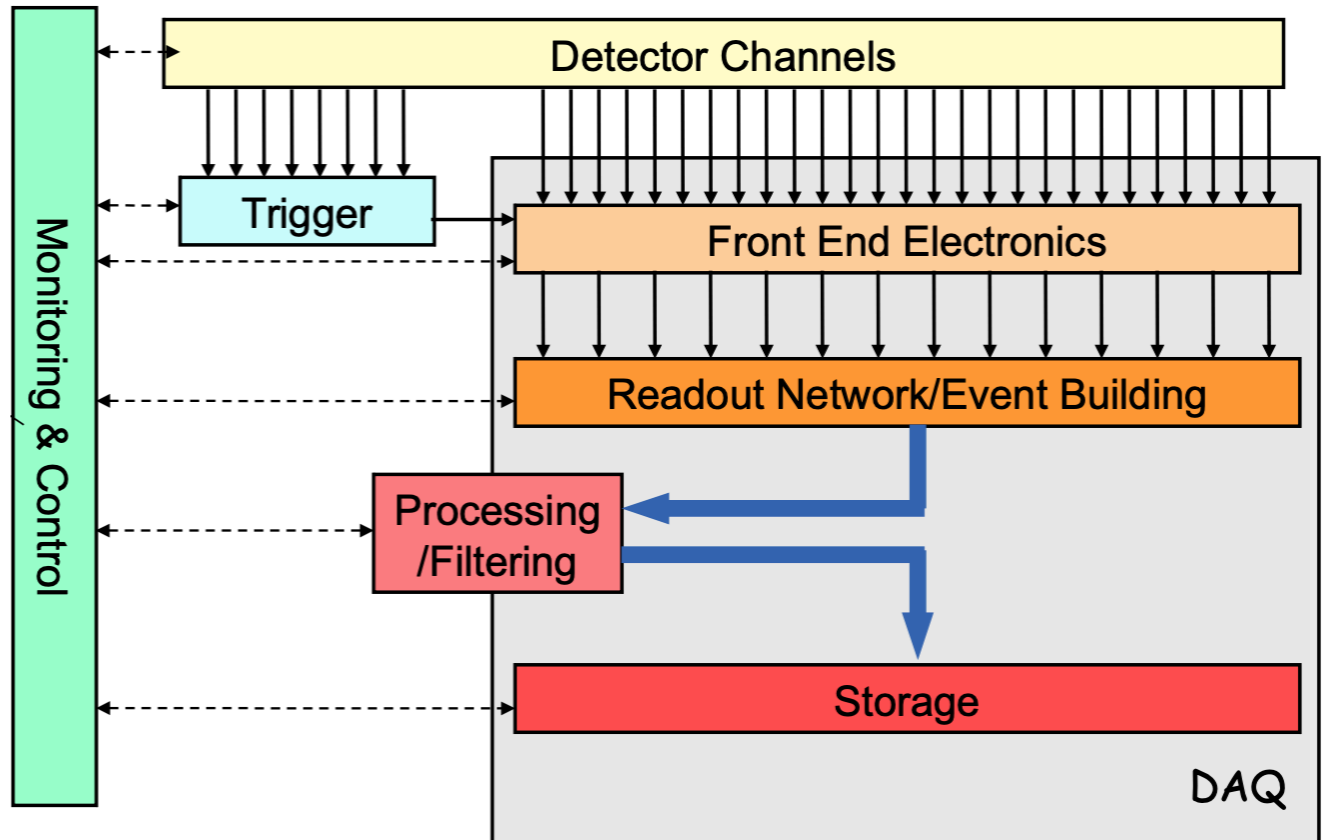
Collisions to analysis



E.g. CMS has $\sim 100\text{M}$ channels, giving a typical event size of $\sim 1\text{ MB}$, and 40 TB/s @ 40 MHz .



Collisions to analysis

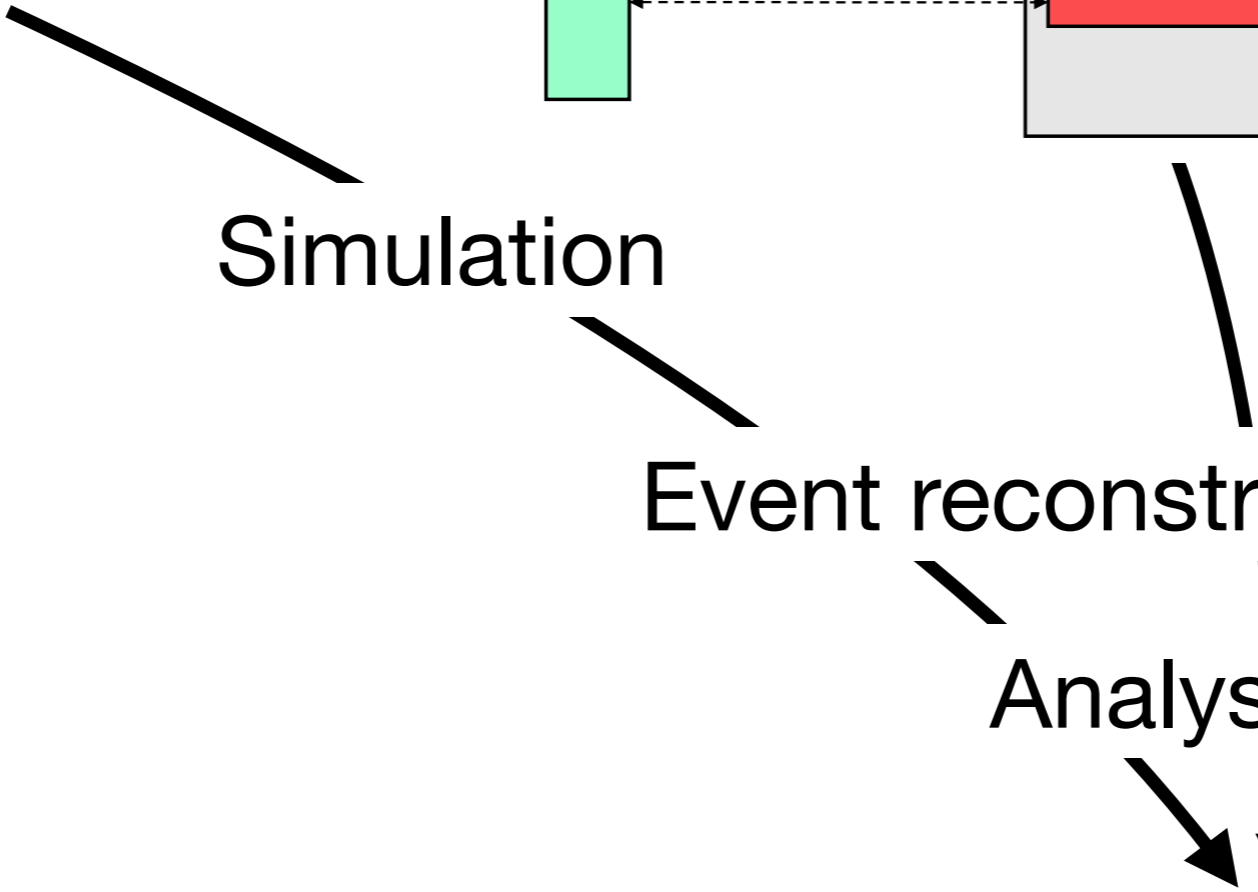


Simulation

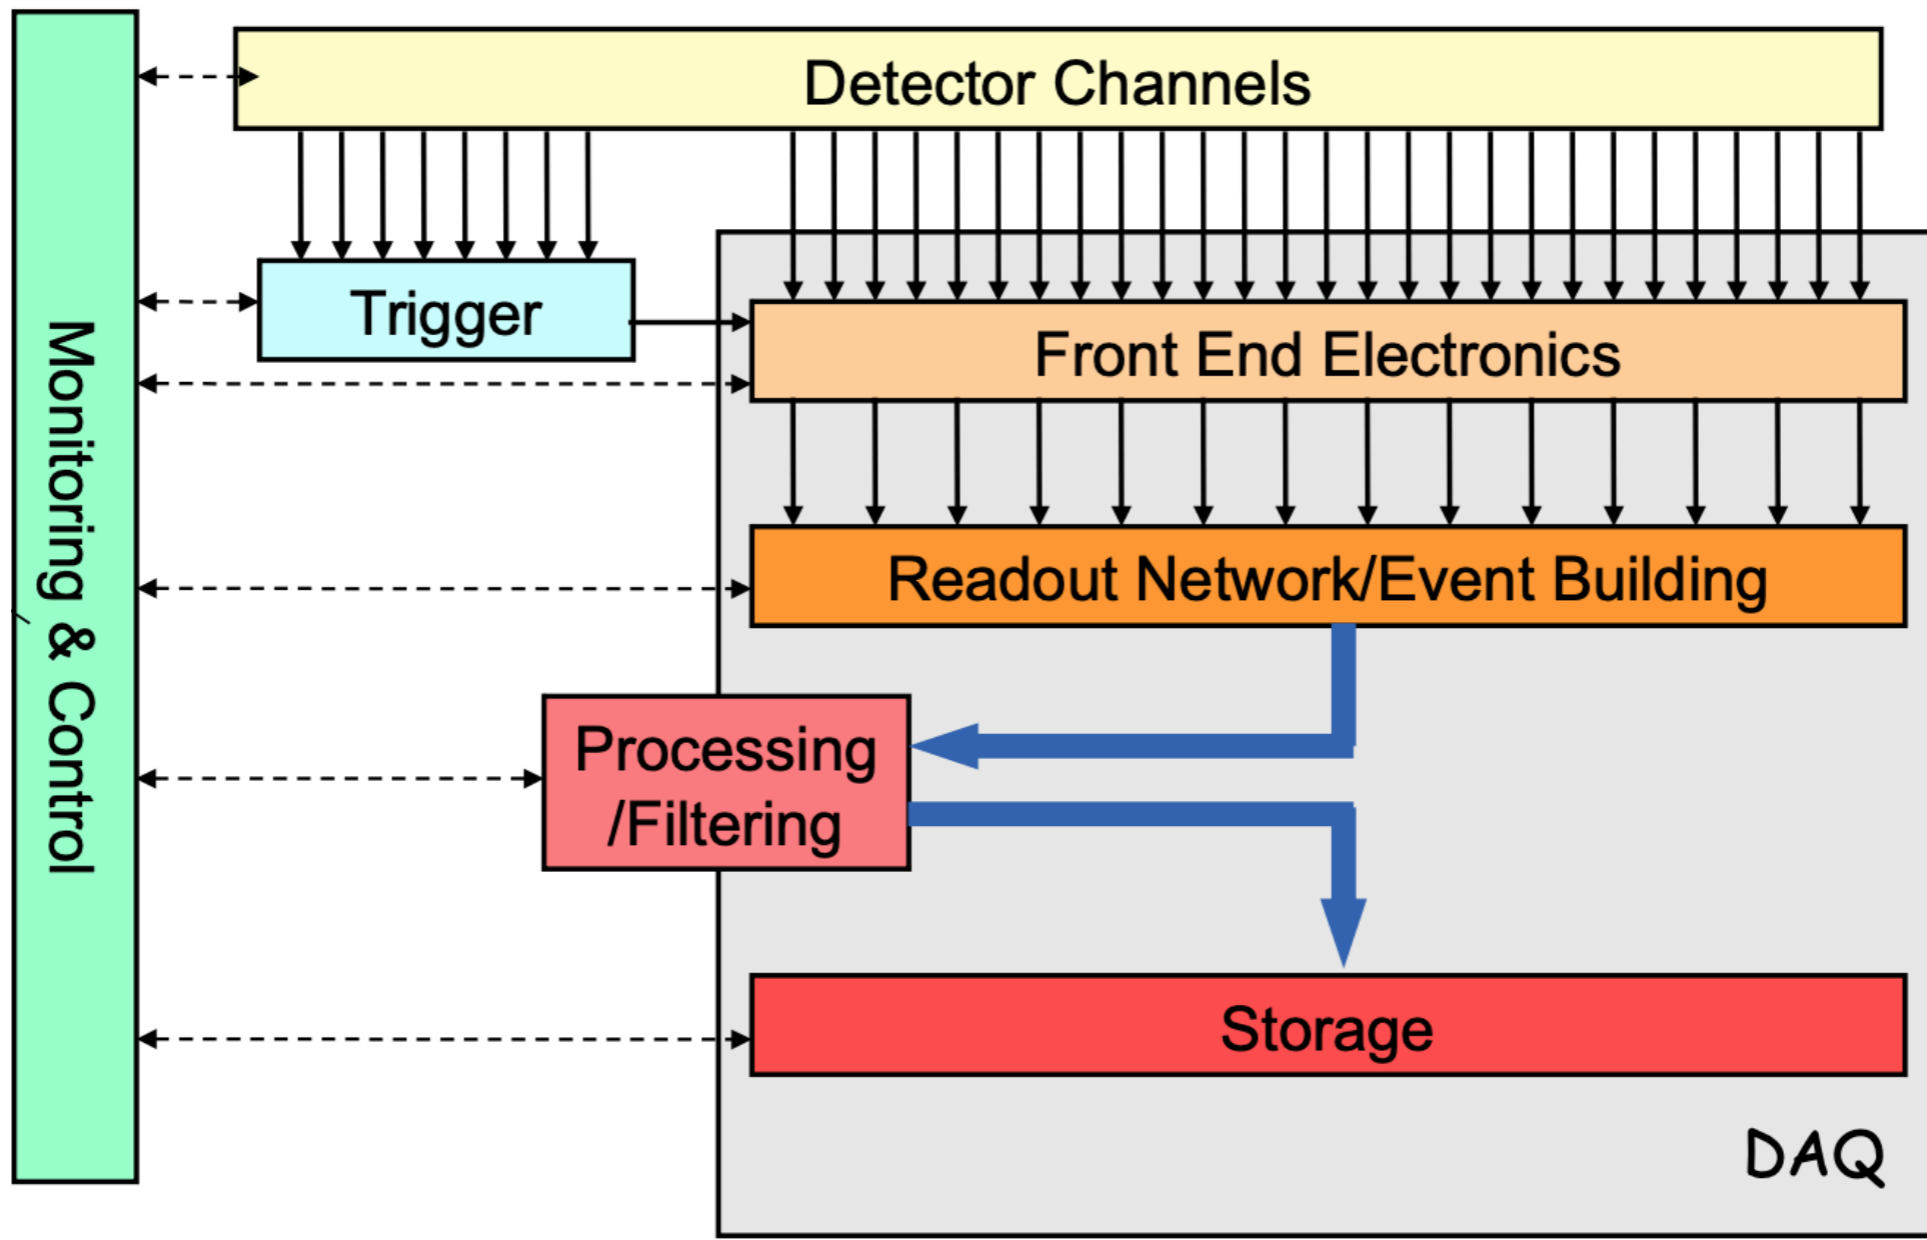
Event reconstruction

Analysis

Physics results




Trigger and DAQ



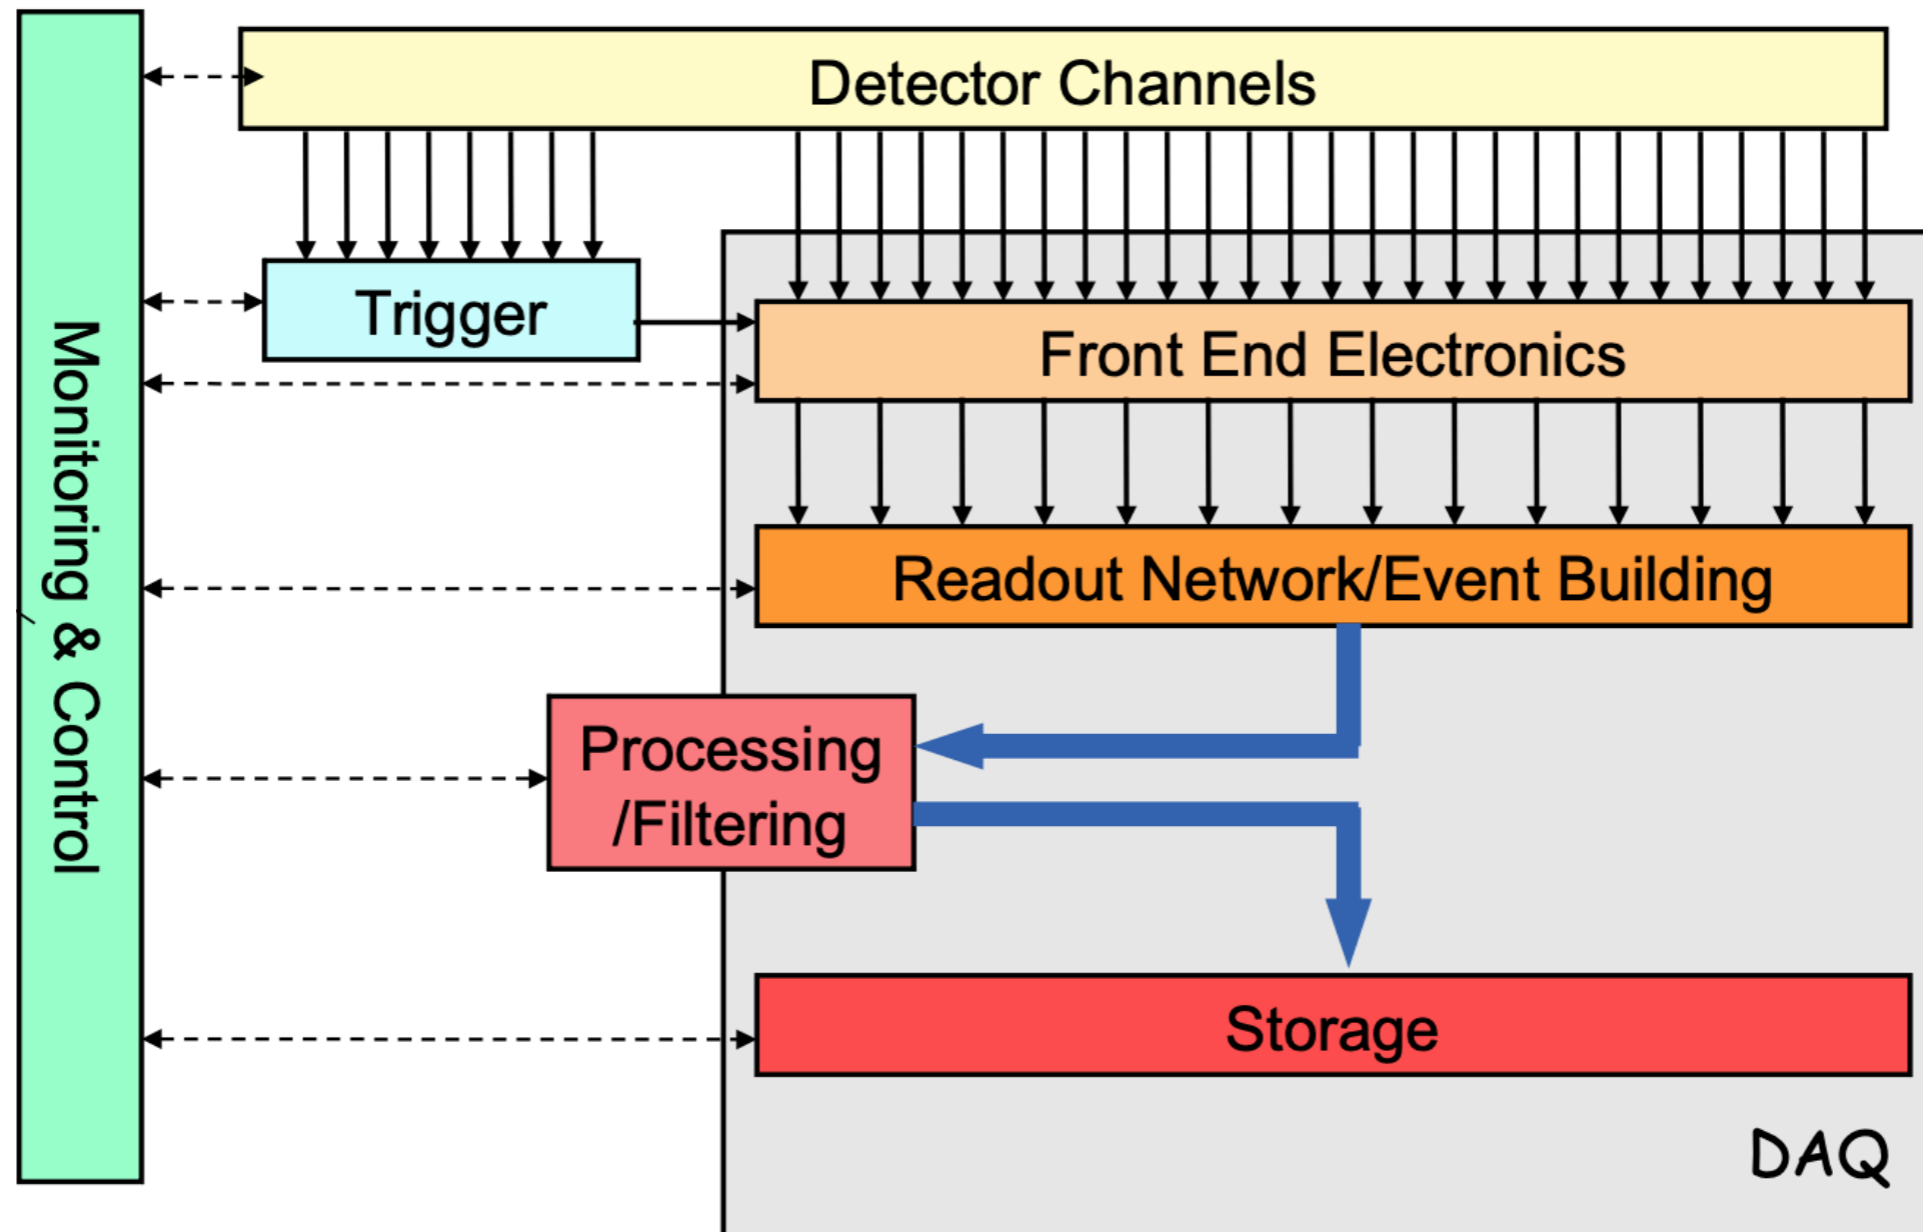
The 2015 ISOTDAQ [slides](#) by A. Negri provide a nice explanation of trigger and DAQ concepts.

Trigger and DAQ



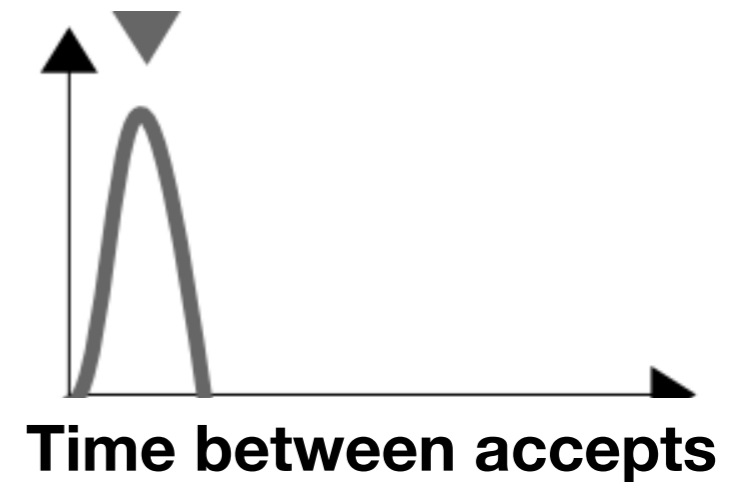

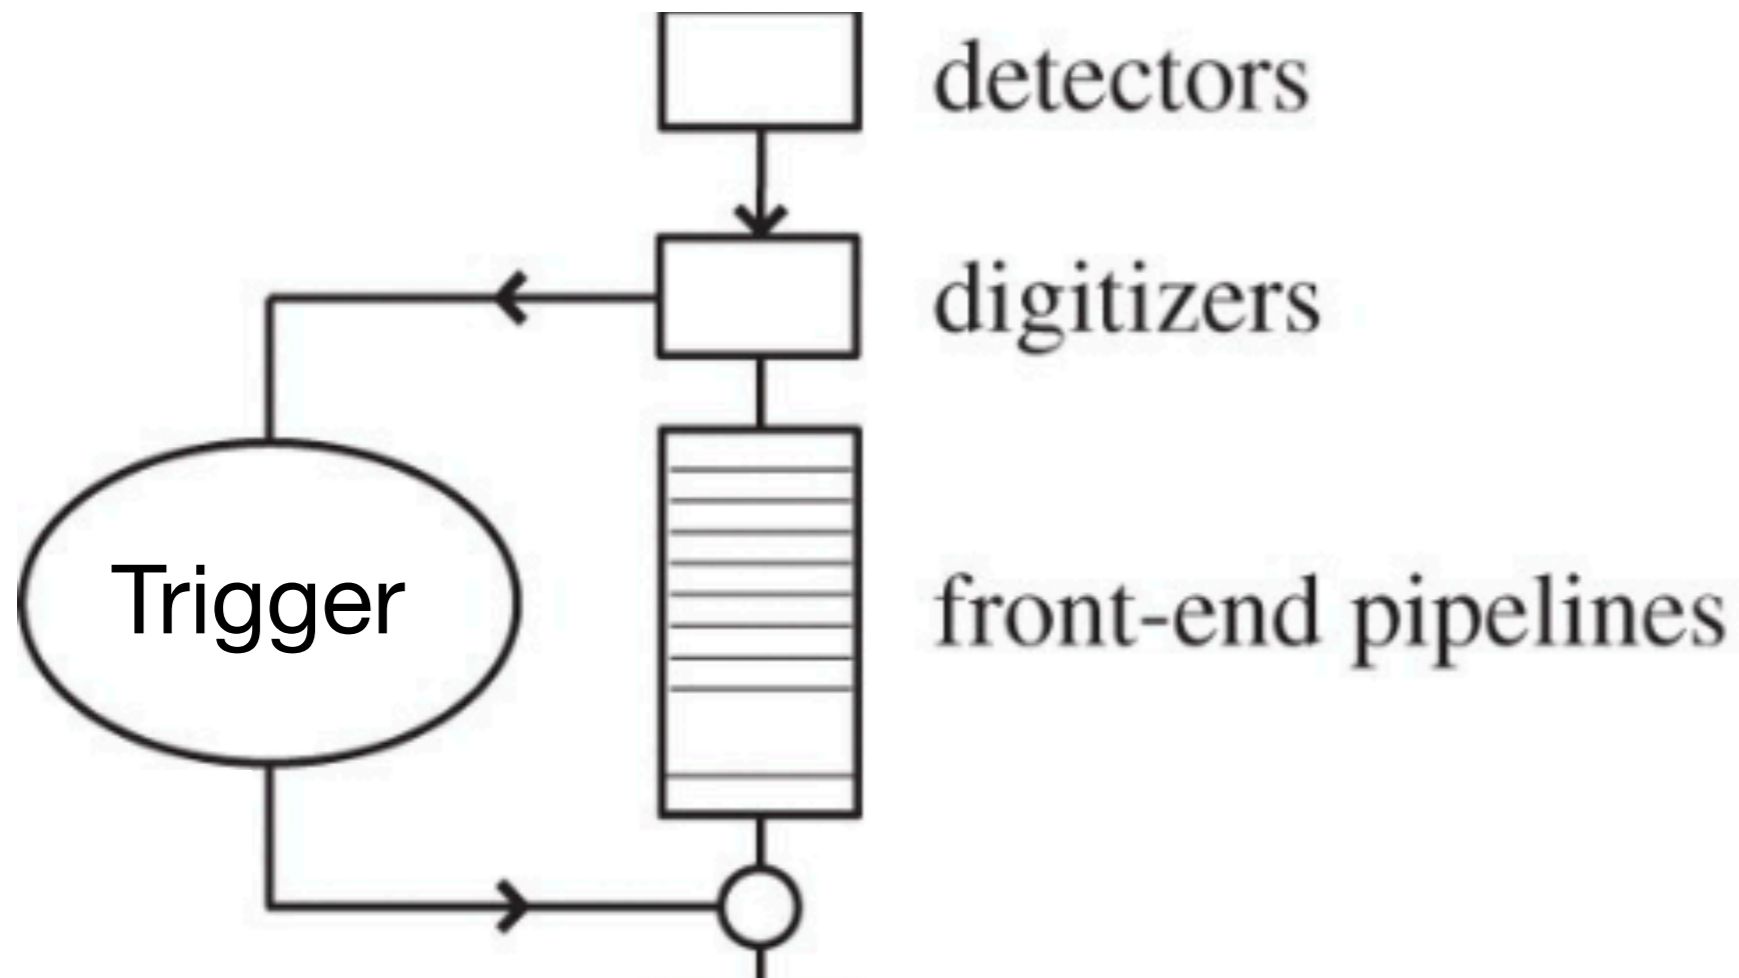
- The trigger takes part of the data and decides, with accept rates of ~ 100 kHz for ATLAS/CMS, whether to send the event for further processing.

Trigger and DAQ



- The trigger takes part of the data and decides, with accept rates of ~ 100 kHz for ATLAS/CMS, whether to send the event for further processing.
- The decision must be made and returned in a fixed *latency* of a few μs , while the data are buffered in front-end *pipelines*.

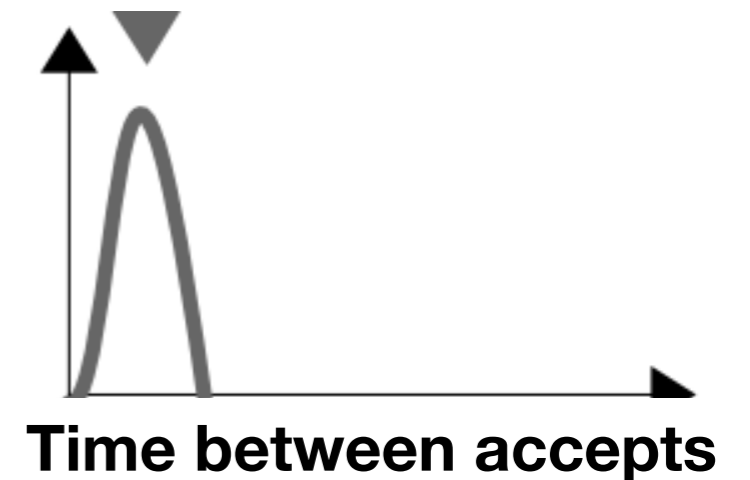
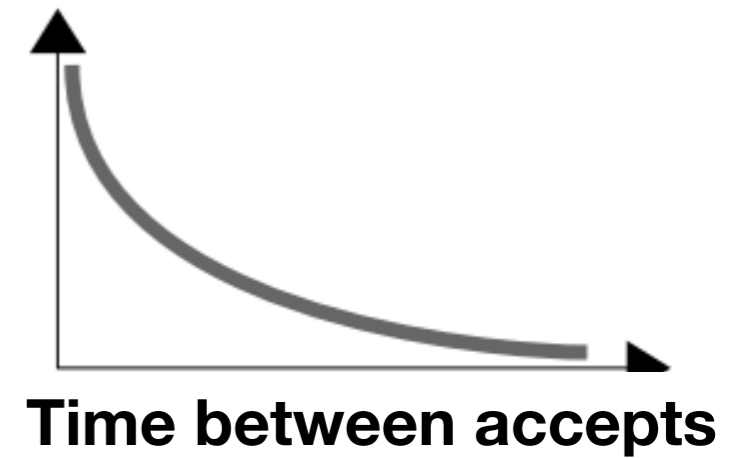

Deadtime and de-randomisation



The *pipelines* buffer the data over the trigger latency.

They also de-randomise the data, reducing *deadtime* due to subsequent processing stages being *busy*.

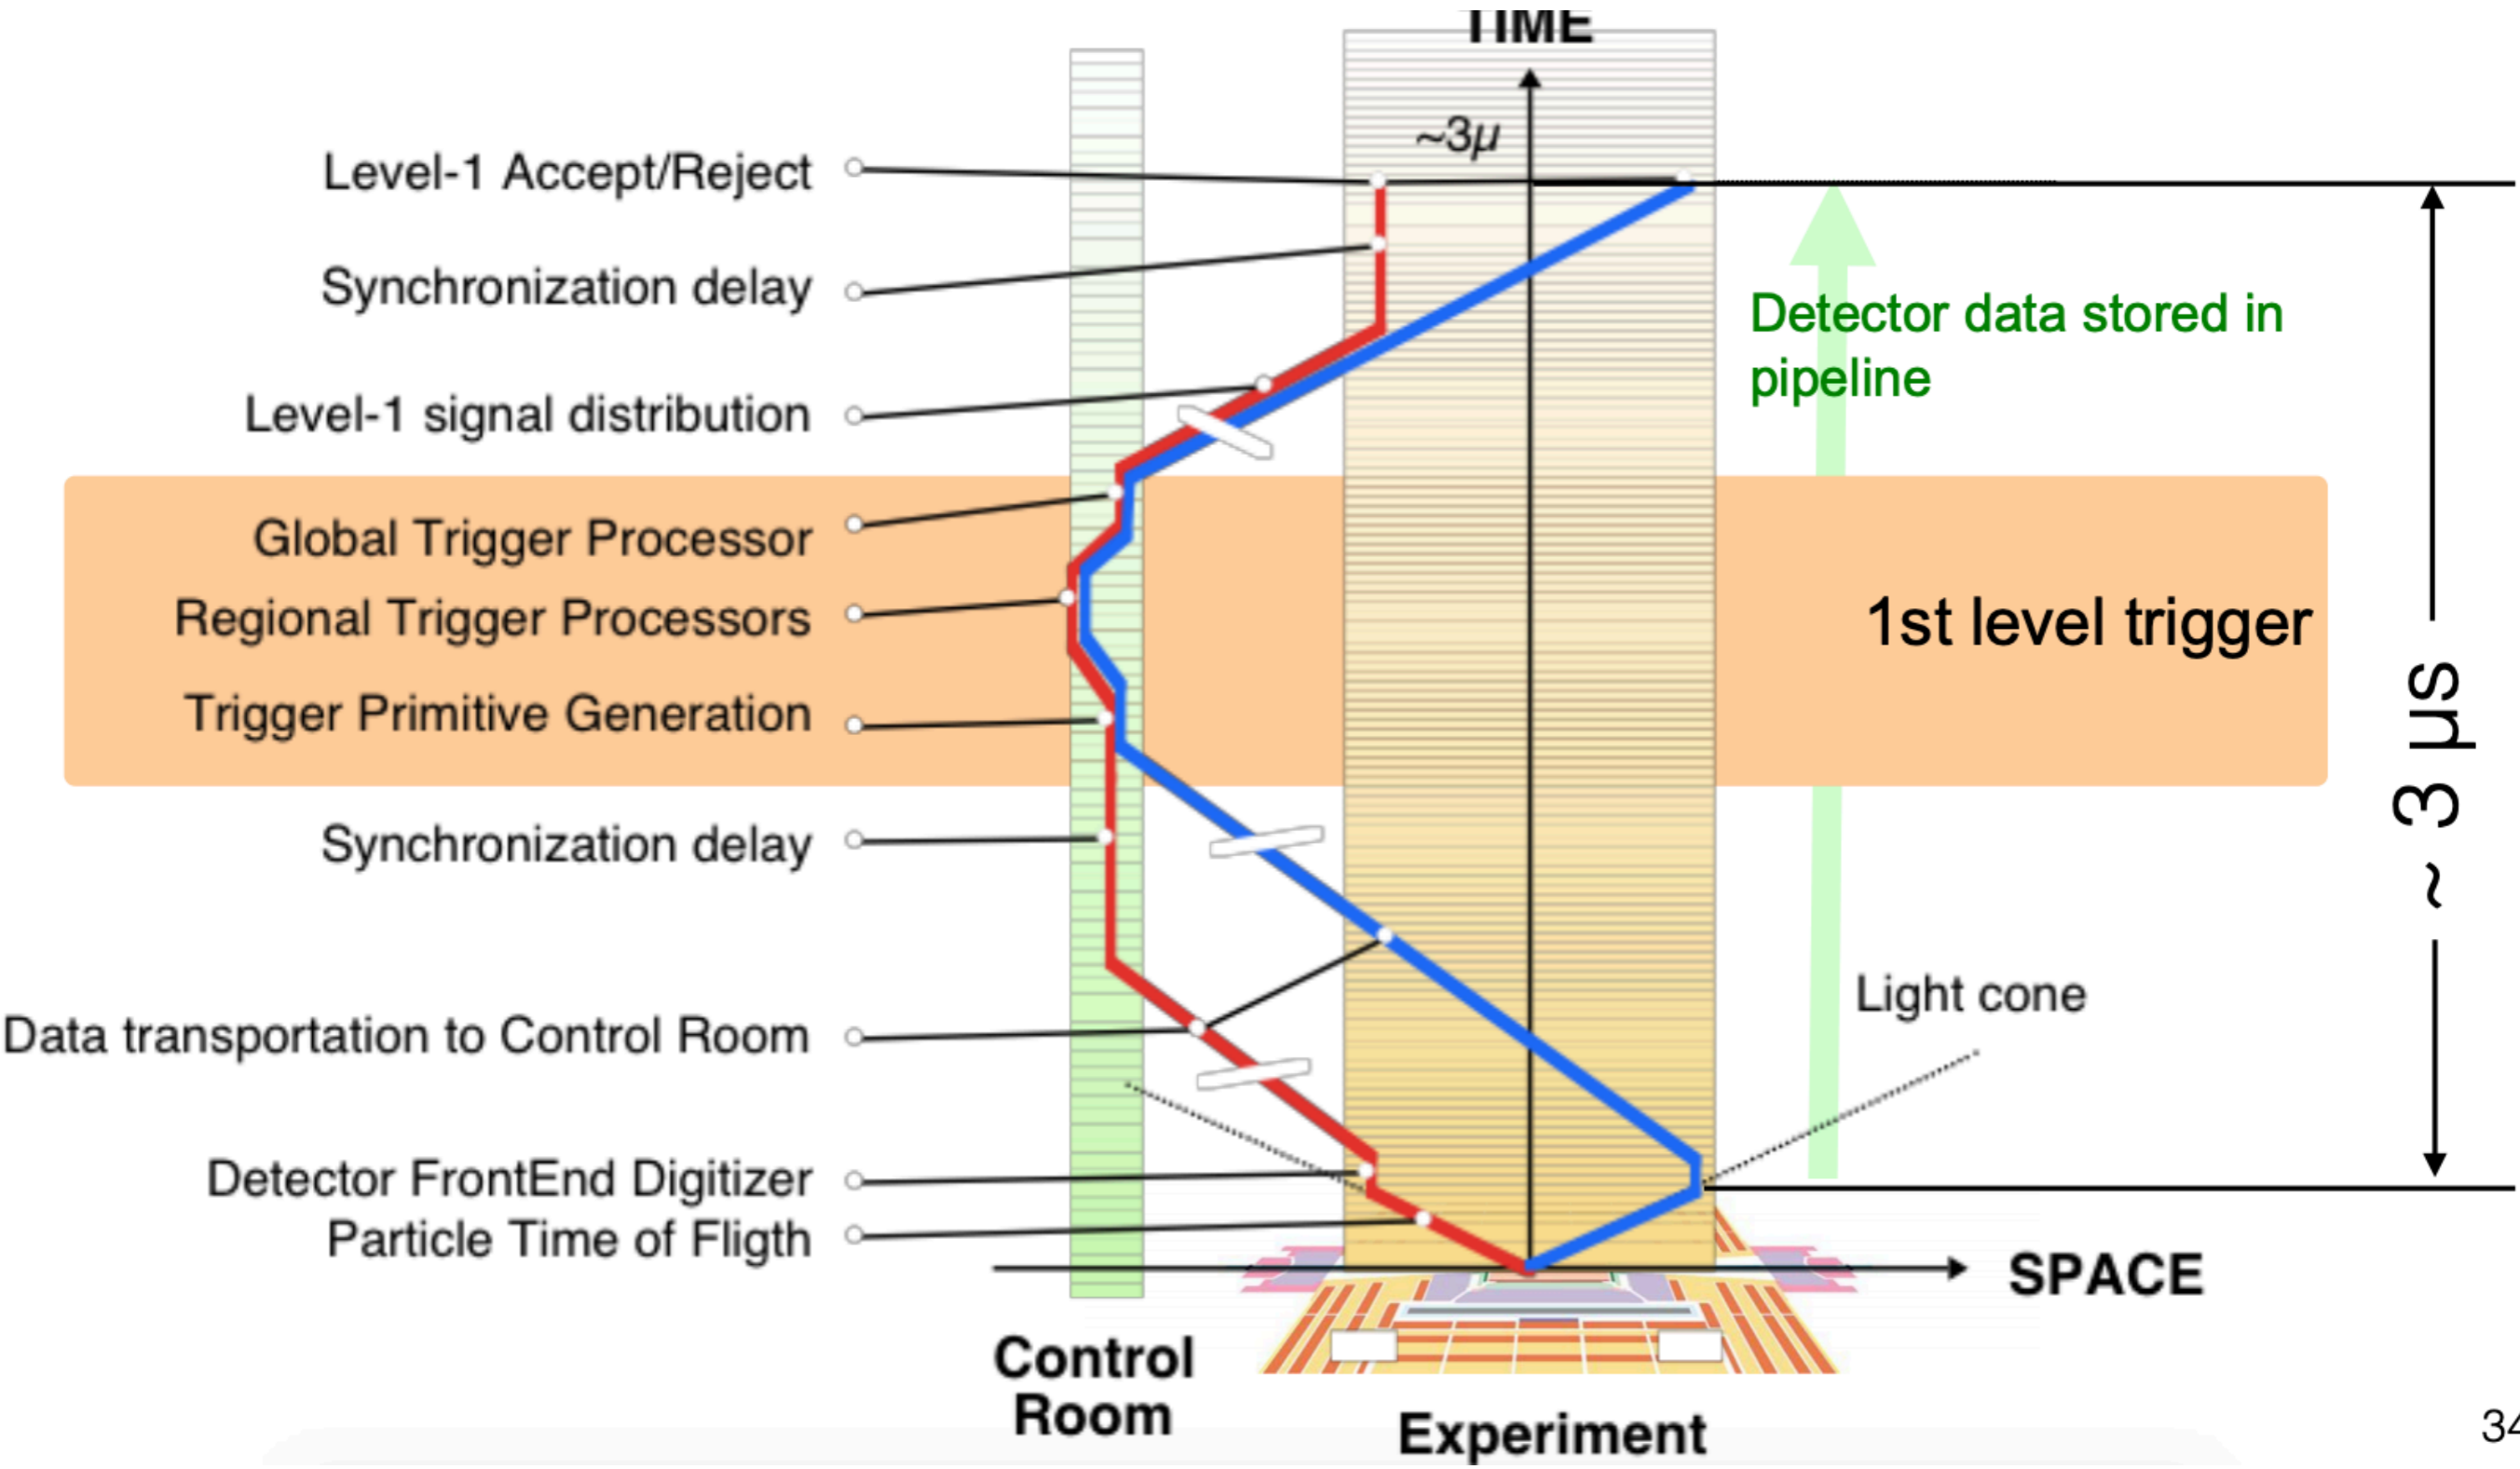
Deadtime and de-randomisation




The *pipelines* buffer the data over the trigger latency.

They also de-randomise the data, reducing *deadtime* due to subsequent processing stages being *busy*.

Trigger latency





What can we do in a few μs with part of the data?



- Simple pattern recognition and *locality*
- Low data rates

Versus what we can't yet do



- High data rates
- Complicated pattern recognition
- Need to link many sub-detectors

Real life example

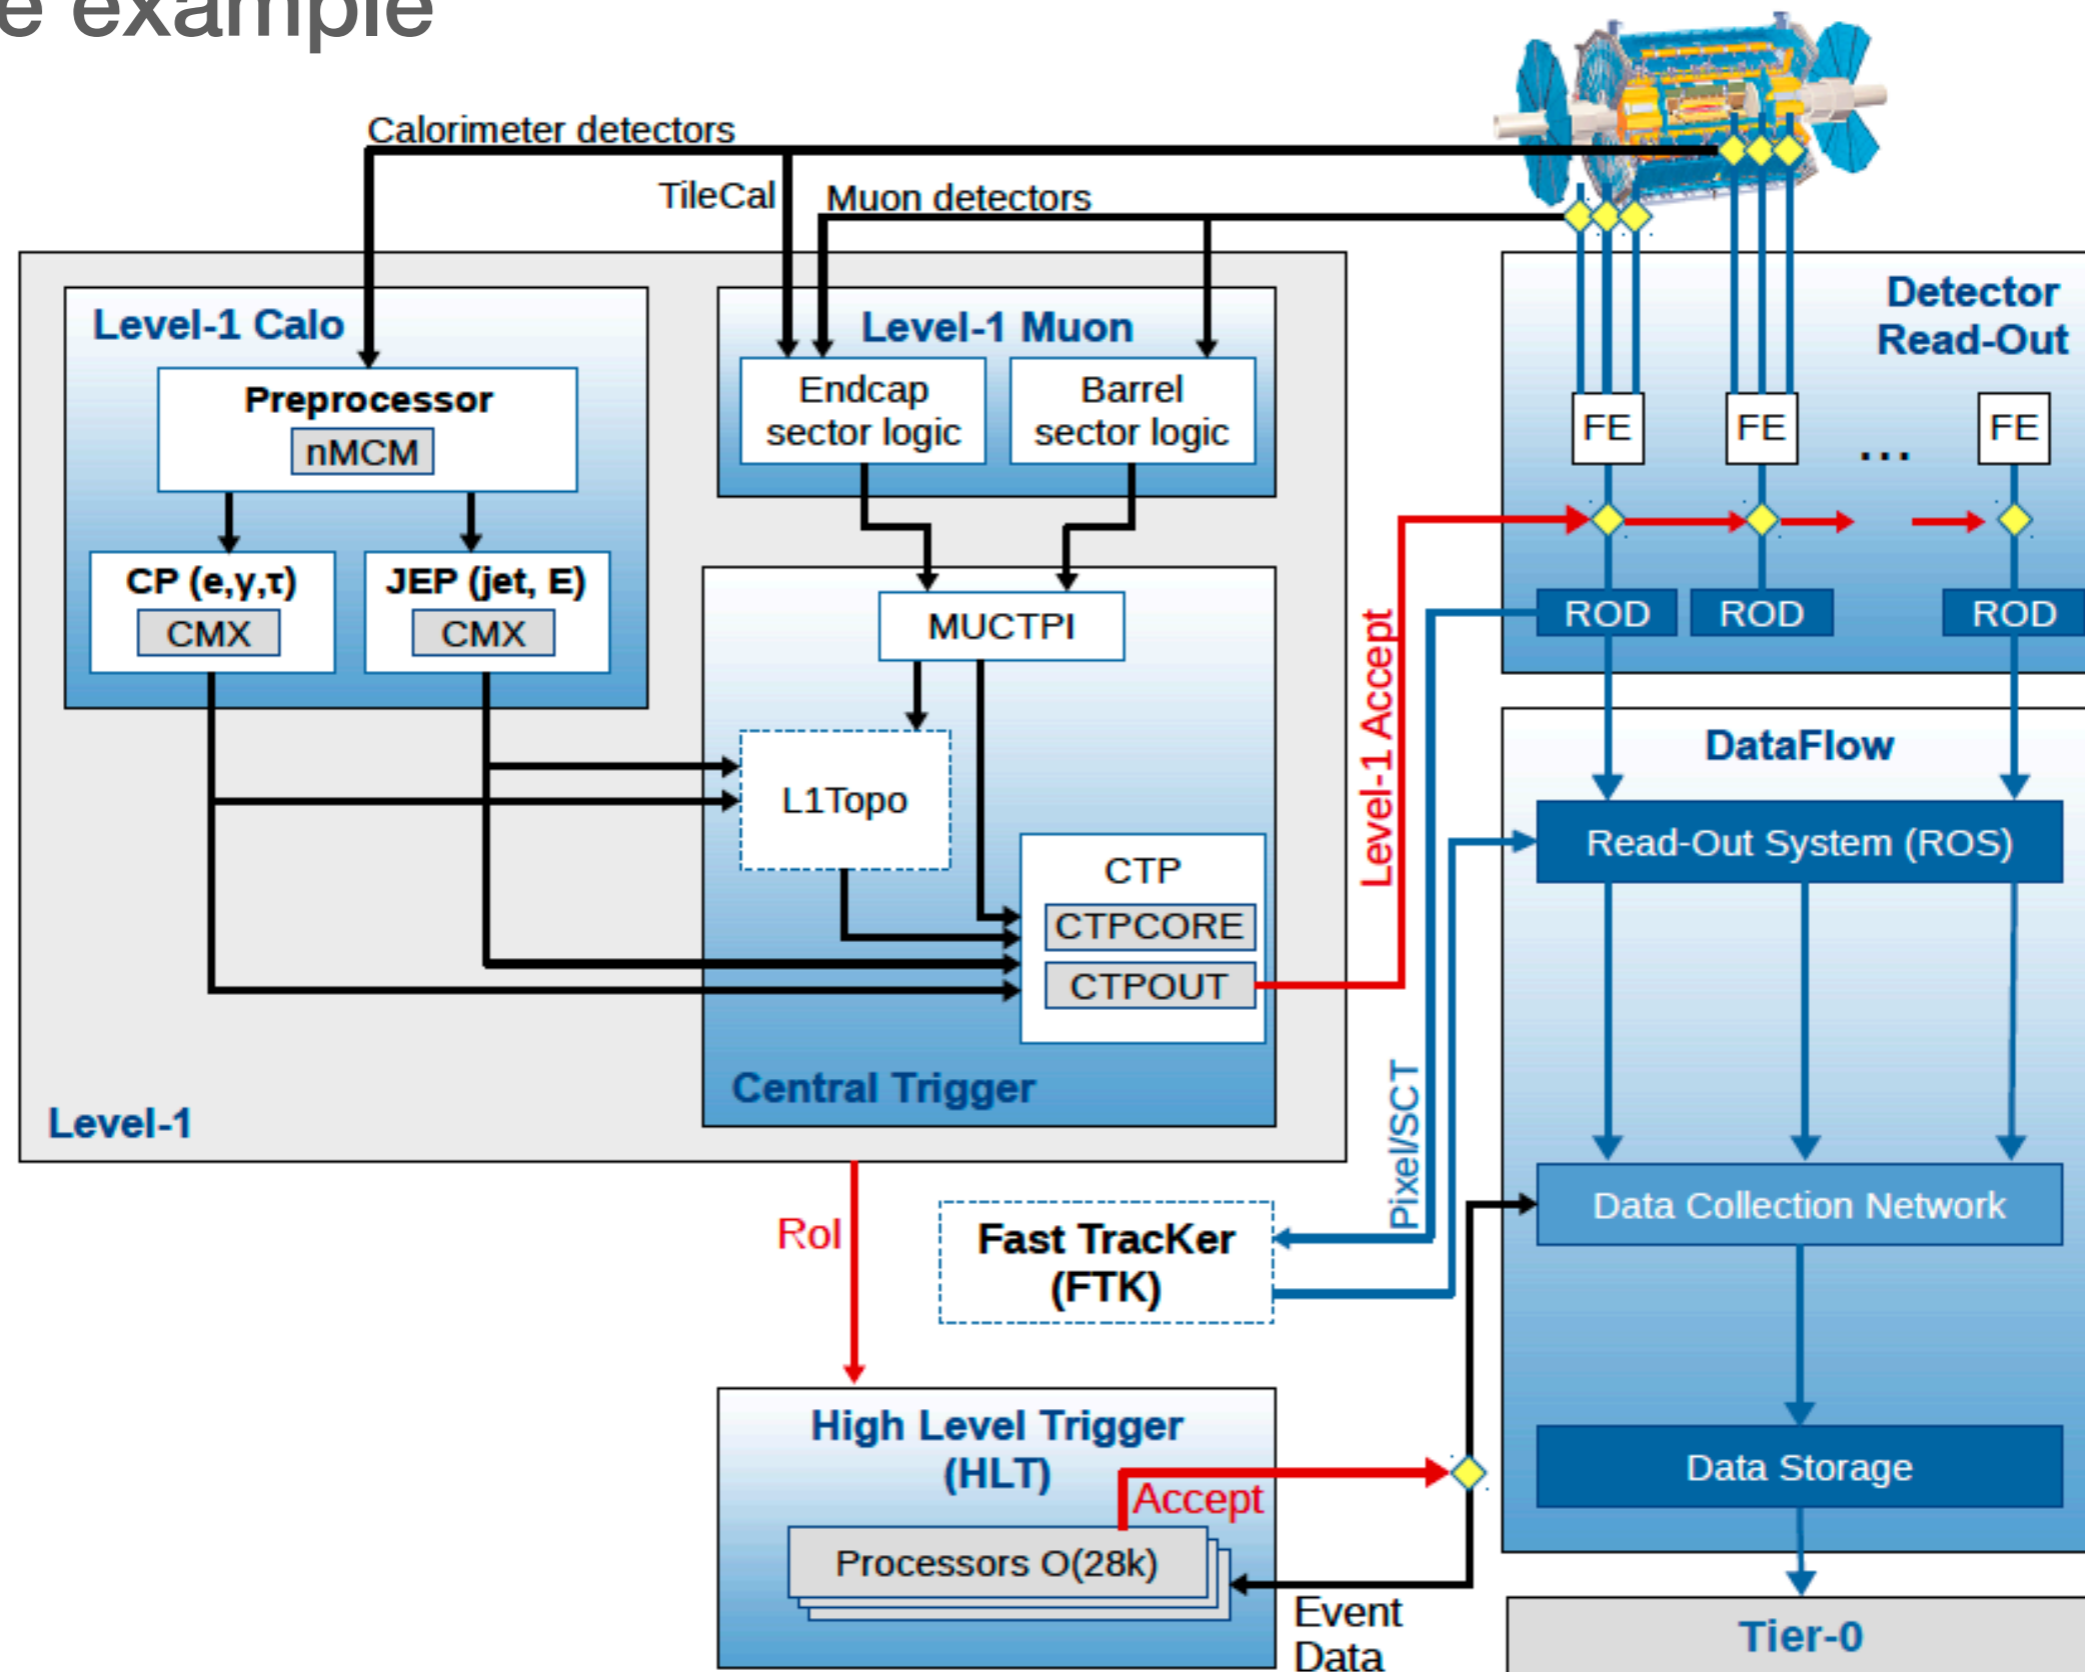



Figure 1: The ATLAS TDAQ system in Run 2 with emphasis on the components relevant for triggering. L1Topo and FTK were being commissioned during 2015 and not used for the results shown here.

E.g., ATLAS thresholds

Table 1: Comparison of selected primary trigger thresholds (in GeV) at the end of Run 1 and during 2015 together with typical offline requirements applied in analyses (the 2012 offline thresholds are not listed but have a similar relationship to the 2012 HLT thresholds). Electron and tau identification are assumed to fulfil the ‘medium’ criteria unless otherwise stated. Photon and b -jet identification (‘b’) are assumed to fulfil the ‘loose’ criteria. Trigger isolation is denoted by ‘i’. The details of these selections are described in Section 6.

Year	2012		2015		
\sqrt{s}	8 TeV		13 TeV		
Peak luminosity	$7.7 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$		$5.0 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$		
Category	p_T threshold [GeV], criteria				
	L1	HLT	L1	HLT	Offline
Single electron	18	24i	20	24	25
Single muon	15	24i	15	20i	21
Single photon	20	120	22i	120	125
Single tau	40	115	60	80	90
Single jet	75	360	100	360	400
Single b -jet	n/a	n/a	100	225	235
E_T^{miss}	40	80	50	70	180
Dielectron	2×10	2×12,loose	2×10	2×12,loose	15
Dimuon	2×10	2×13	2×10	2×10	11
Electron, muon	10, 6	12, 8	15, 10	17, 14	19, 15
Diphoton	16, 12	35, 25	2×15	35, 25	40, 30
Ditau	15i, 11i	27, 18	20i, 12i	35, 25	40, 30
Tau, electron	11i, 14	28i, 18	12i(+jets), 15	25, 17i	30, 19
Tau, muon	8, 10	20, 15	12i(+jets), 10	25, 14	30, 15
Tau, E_T^{miss}	20, 35	38, 40	20, 45(+jets)	35, 70	40, 180
Four jets	4×15	4×80	3×40	4×85	95
Six jets	4×15	6×45	4×15	6×45	55
Two b -jets	75	35b, 145b	100	50b, 150b	60
Four(Two) (b -)jets	4×15	2×35b, 2×35	3×25	2×35b, 2×35	45
B -physics (Dimuon)	6, 4	6, 4	6, 4	6, 4	6, 4

ATLAS L1 bandwidth division



Wed, 4 Sep 2019 (showing first 14 of 22 entries)



[19] [arXiv:1909.00761](https://arxiv.org/abs/1909.00761) [pdf, other]

Performance of electron and photon triggers in ATLAS during LHC Run 2


ATLAS Collaboration

Comments: 55 pages in total, author list starting page 39, 26 figures, 10 tables, submitted to EPJC. All figures including auxiliary figures are available at [this https URL](https://arxiv.org/abs/1909.00761)

Subjects: High Energy Physics – Experiment (hep-ex)




Importance of calibration/resolution




LHCb

$m_B \sim 5 \text{ GeV}$
 $(\gamma_{CT})_B \sim 1 \text{ cm}$




$m_D \sim 2 \text{ GeV}$
 $(\gamma_{CT})_D \sim 4 \text{ mm}$





Not well suited to fast low-level triggers.

LHCb


$m_B \sim 5 \text{ GeV}$
 $(\gamma_{CT})_B \sim 1 \text{ cm}$






$m_D \sim 2 \text{ GeV}$
 $(\gamma_{CT})_D \sim 4 \text{ mm}$







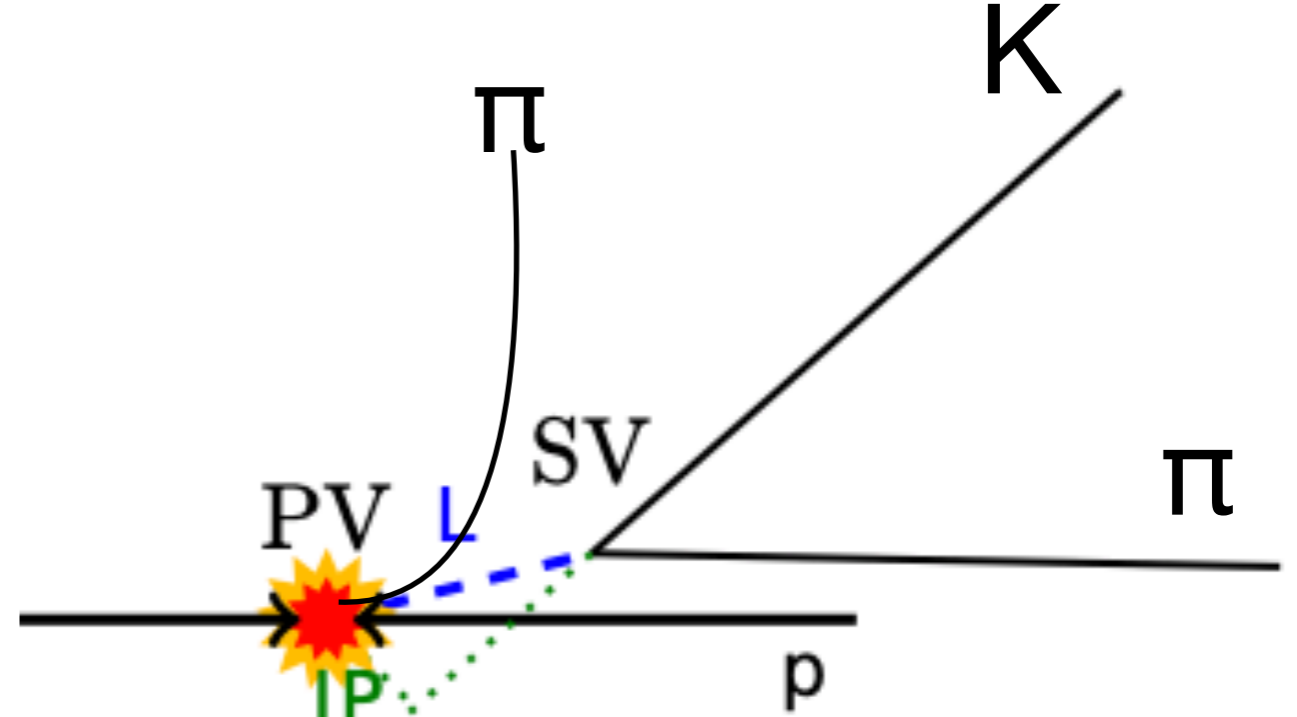
1. Run at a *levelled* luminosity of $4 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$
2. Profit from smaller events to take x10 higher L1 accept rate than ATLAS/CMS.




Classification versus selection



Classification versus selection



Classification versus selection

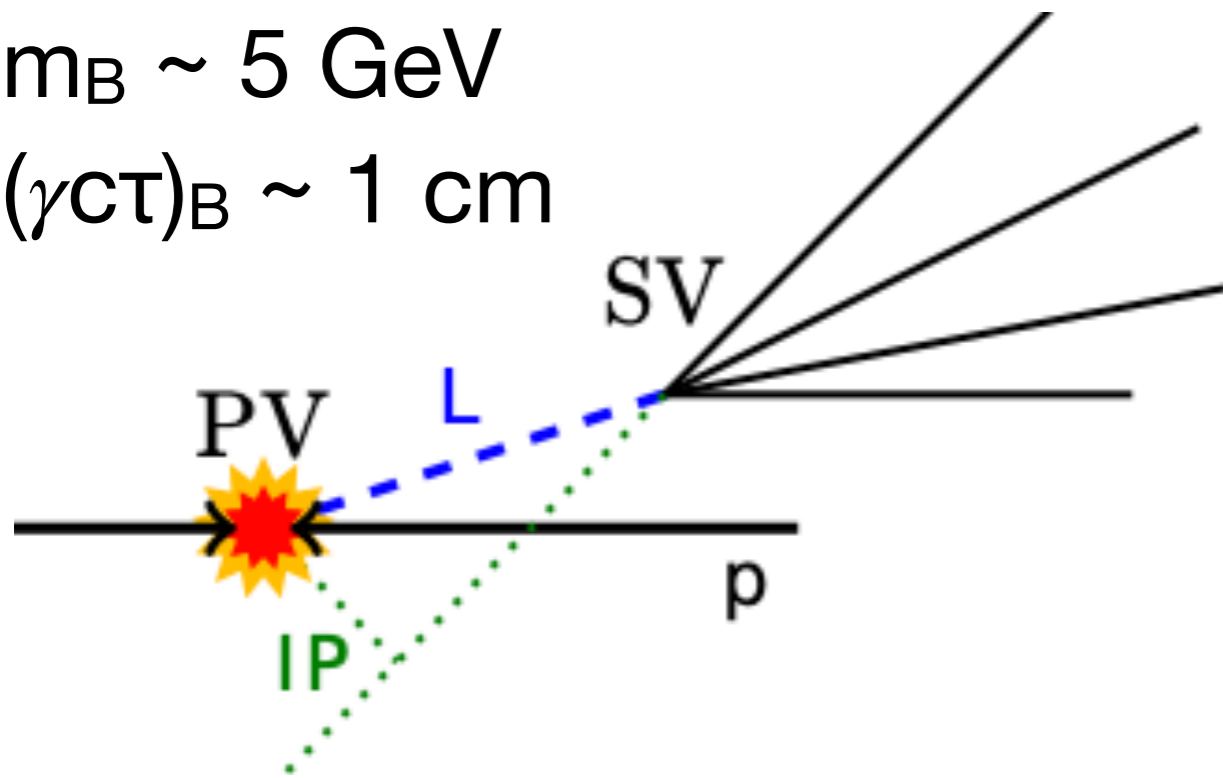


So far we have talked about *inclusive* selections, i.e. based on part of the signal.


What I just showed was an example of where an *exclusive* selection is required to classify the signals. This typically requires full offline quality event reconstruction.

LHCb L0 (1 MHz) trigger

$m_B \sim 5 \text{ GeV}$
 $(\gamma_{CT})_B \sim 1 \text{ cm}$



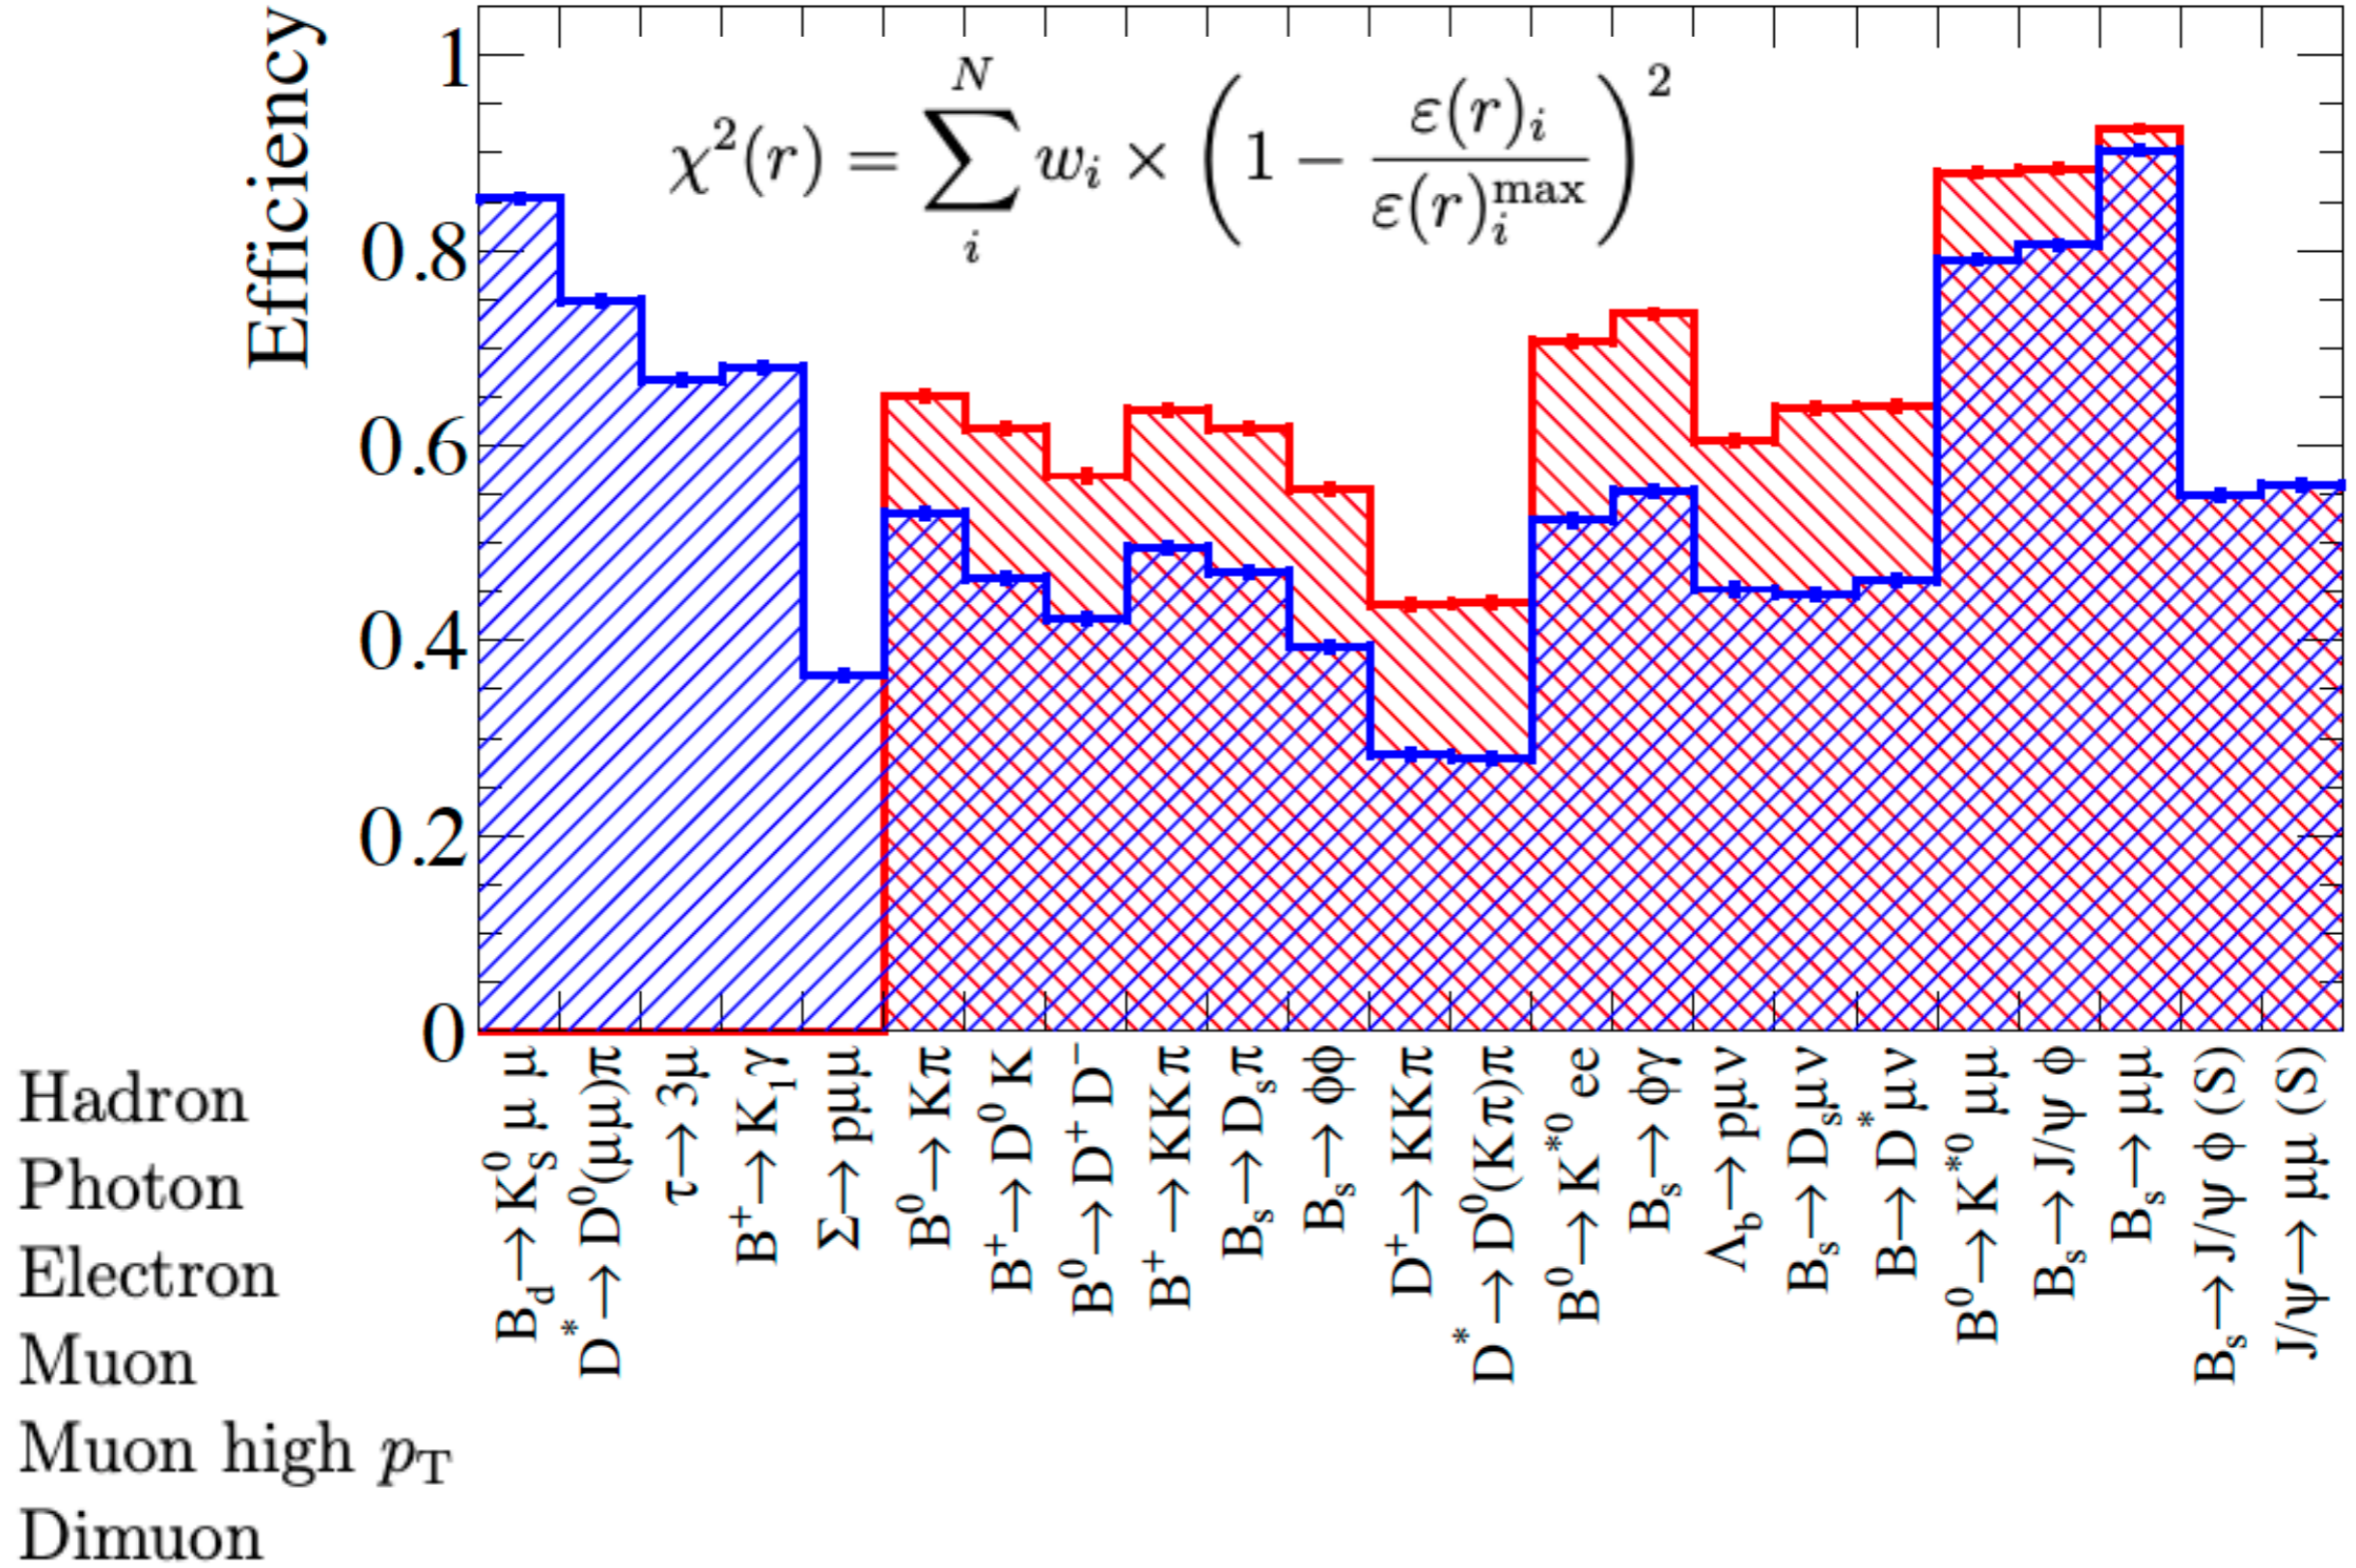
$m_D \sim 2 \text{ GeV}$
 $(\gamma_{CT})_D \sim 4 \text{ mm}$



L0 trigger	E_T/p_T threshold			SPD threshold
	2015	2016	2017	
Hadron	$> 3.6 \text{ GeV}$	$> 3.7 \text{ GeV}$	$> 3.46 \text{ GeV}$	< 450
Photon	$> 2.7 \text{ GeV}$	$> 2.78 \text{ GeV}$	$> 2.47 \text{ GeV}$	< 450
Electron	$> 2.7 \text{ GeV}$	$> 2.4 \text{ GeV}$	$> 2.11 \text{ GeV}$	< 450
Muon	$> 2.8 \text{ GeV}$	$> 1.8 \text{ GeV}$	$> 1.35 \text{ GeV}$	< 450
Muon high p_T	$> 6.0 \text{ GeV}$	$> 6.0 \text{ GeV}$	$> 6.0 \text{ GeV}$	none
Dimuon	$> 1.69 \text{ GeV}^2$	$> 2.25 \text{ GeV}^2$	$> 1.69 \text{ GeV}^2$	< 900

How are these thresholds decided?


LHCb L0 (≈ 1 MHz) bandwidth division



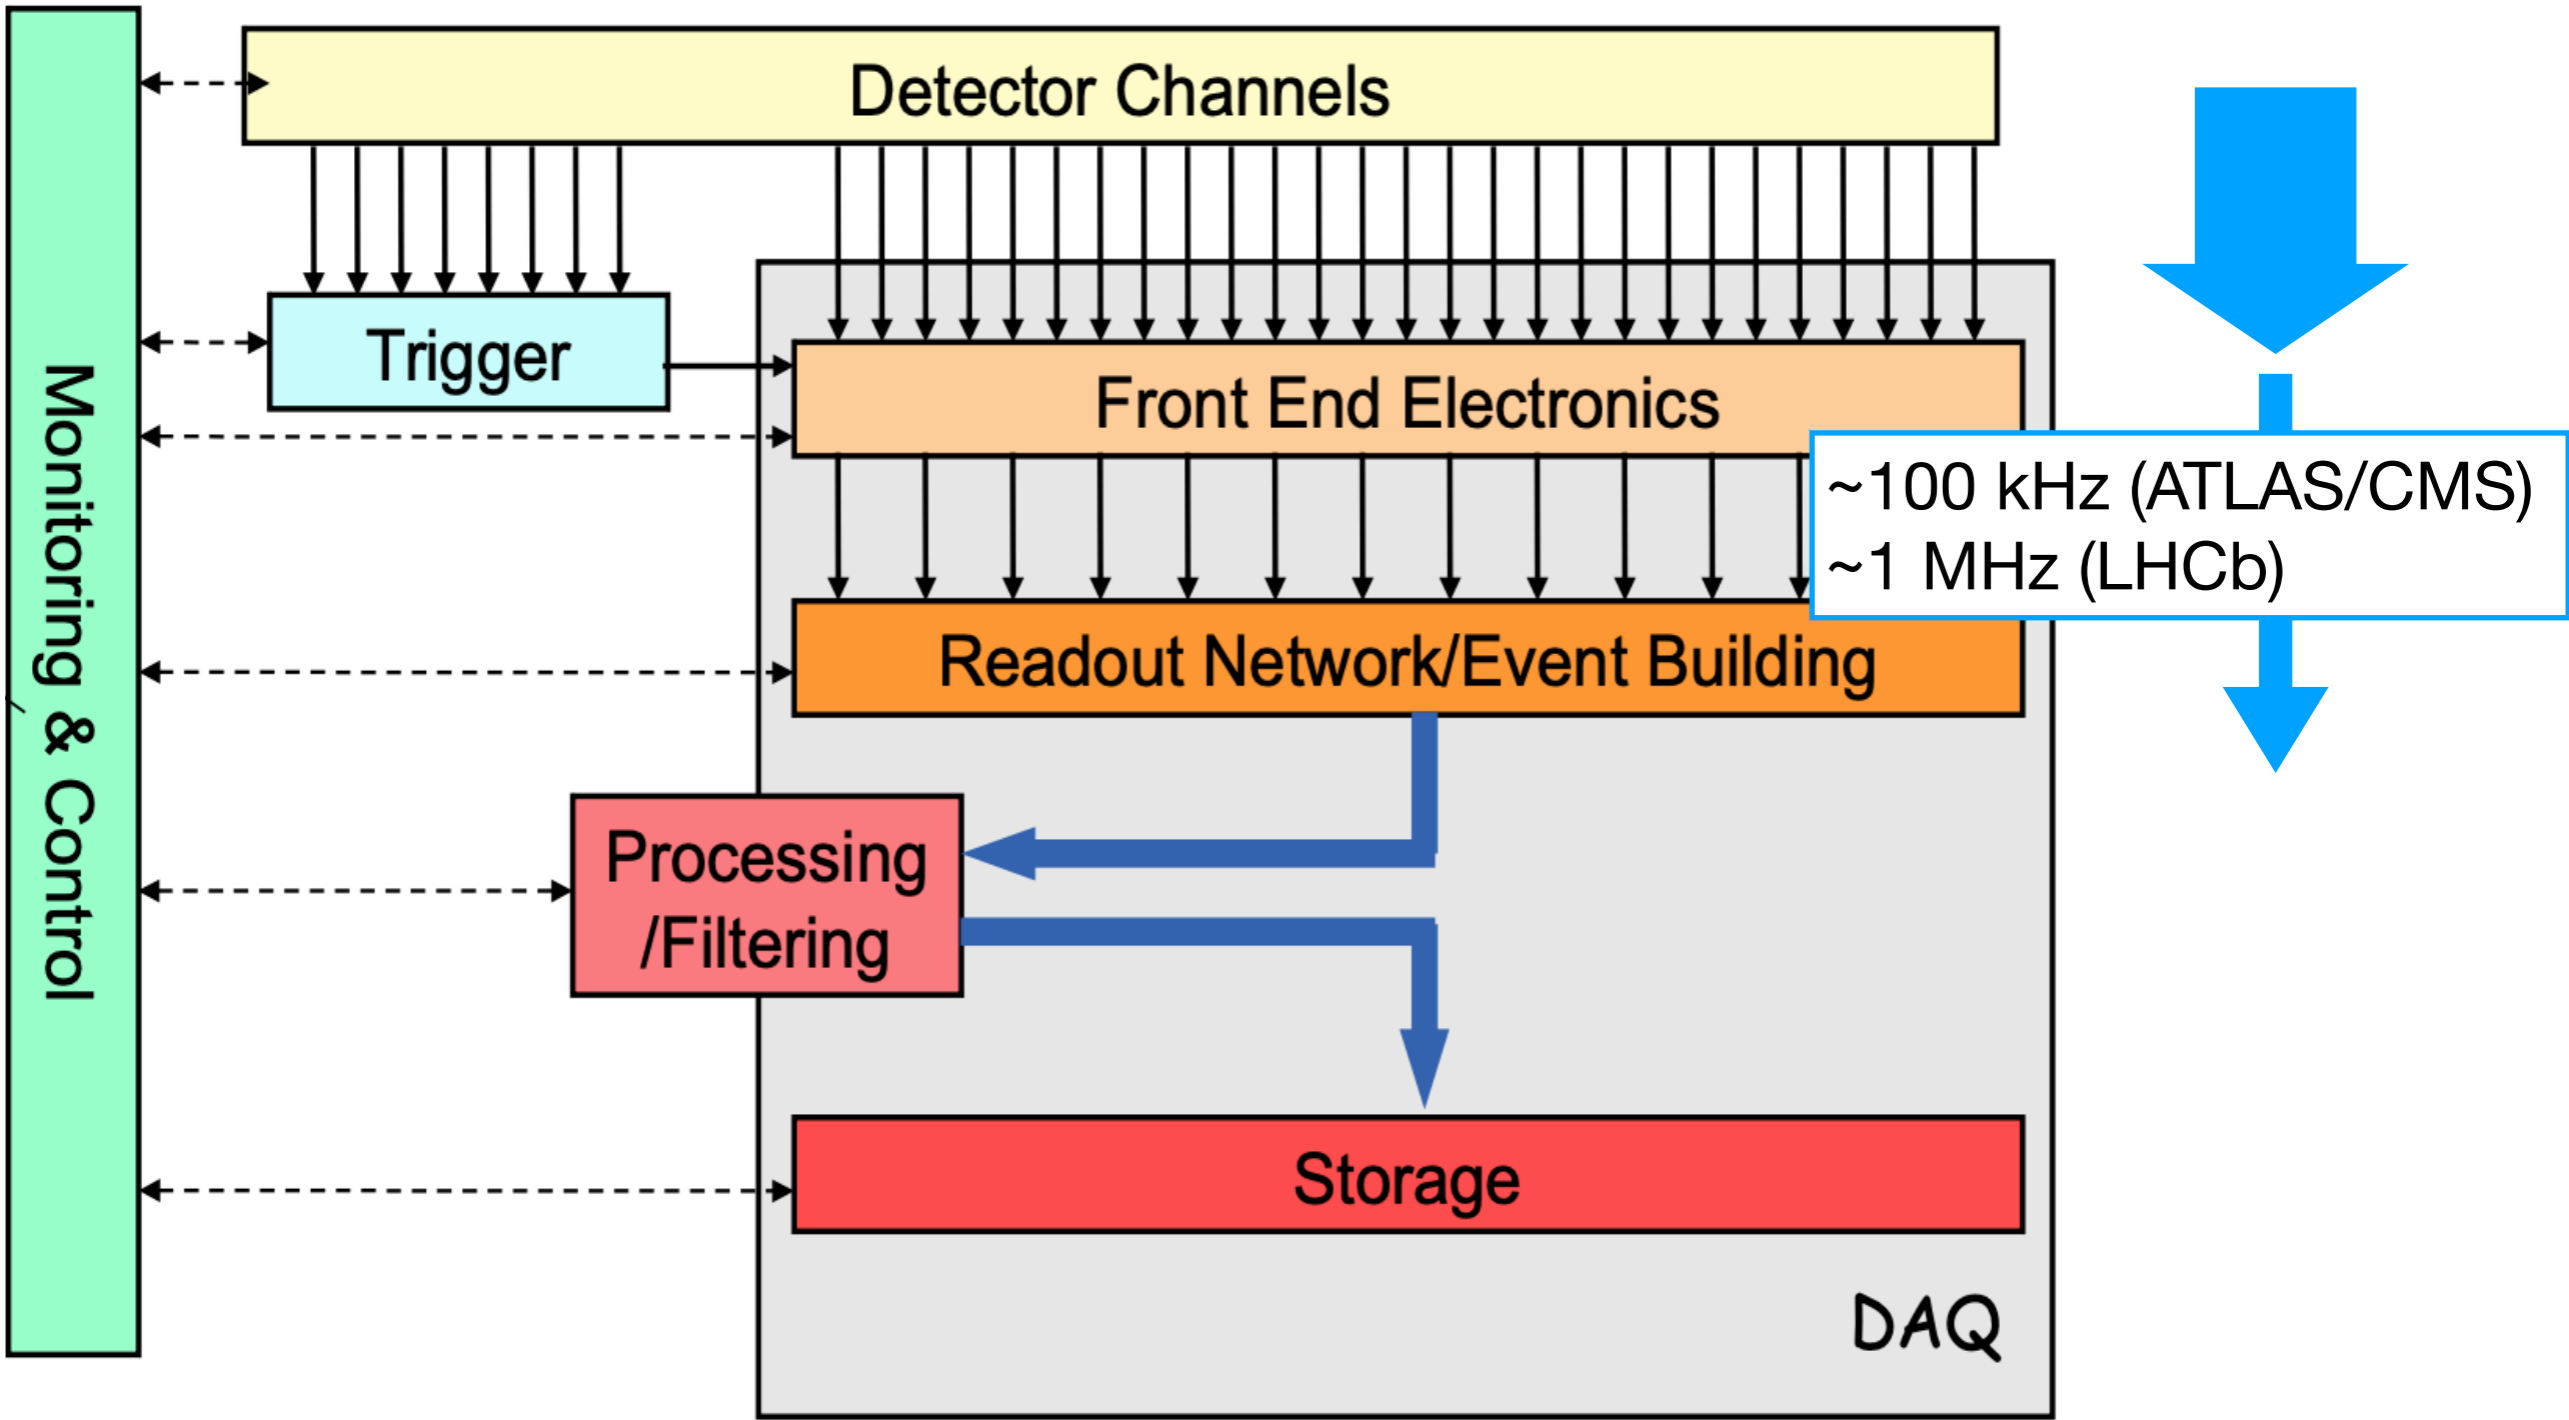
Why don't we just raise the luminosity?

Why don't we just raise the luminosity?

[Letter of Intent for the LHCb Upgrade, CERN-LHCC-2011-001](#)



Going beyond the low-level triggers



LHCb Trigger Run 2

Bunch crossing rate

↓ 40 MHz

L0 Hardware trigger

high p_T/E_T signatures

↓ 1 MHz

Software trigger

High Level Trigger 1
partial event reconstruction

↓ 110 kHz

**10 PB
buffer**

**Alignment &
Calibration**

↓ 110 kHz

High Level Trigger 2

full event reconstruction



↓ 12.5 kHz

Storage

LHCb software trigger farm in Run-II: ~27000 physical cores running ~55000 processes.



The need for offline quality alignment and calibration



The LHCb RICH detectors








FIG. 6. Difference between the measured and expected Cherenkov angle, $\Delta\theta_C$ plotted as a function of the azimuthal angle ϕ and fitted with $\theta_x \cos(\phi) + \theta_y \sin(\phi)$, for one side of the RICH2 detector [6]. The upper plot is prior to alignment, and shows a dependency of the angle θ_C on the angle ϕ . The bottom plot is after the alignment correction, and $\Delta\theta_C$ is uniform in ϕ .

LHCb real time alignment and calibration



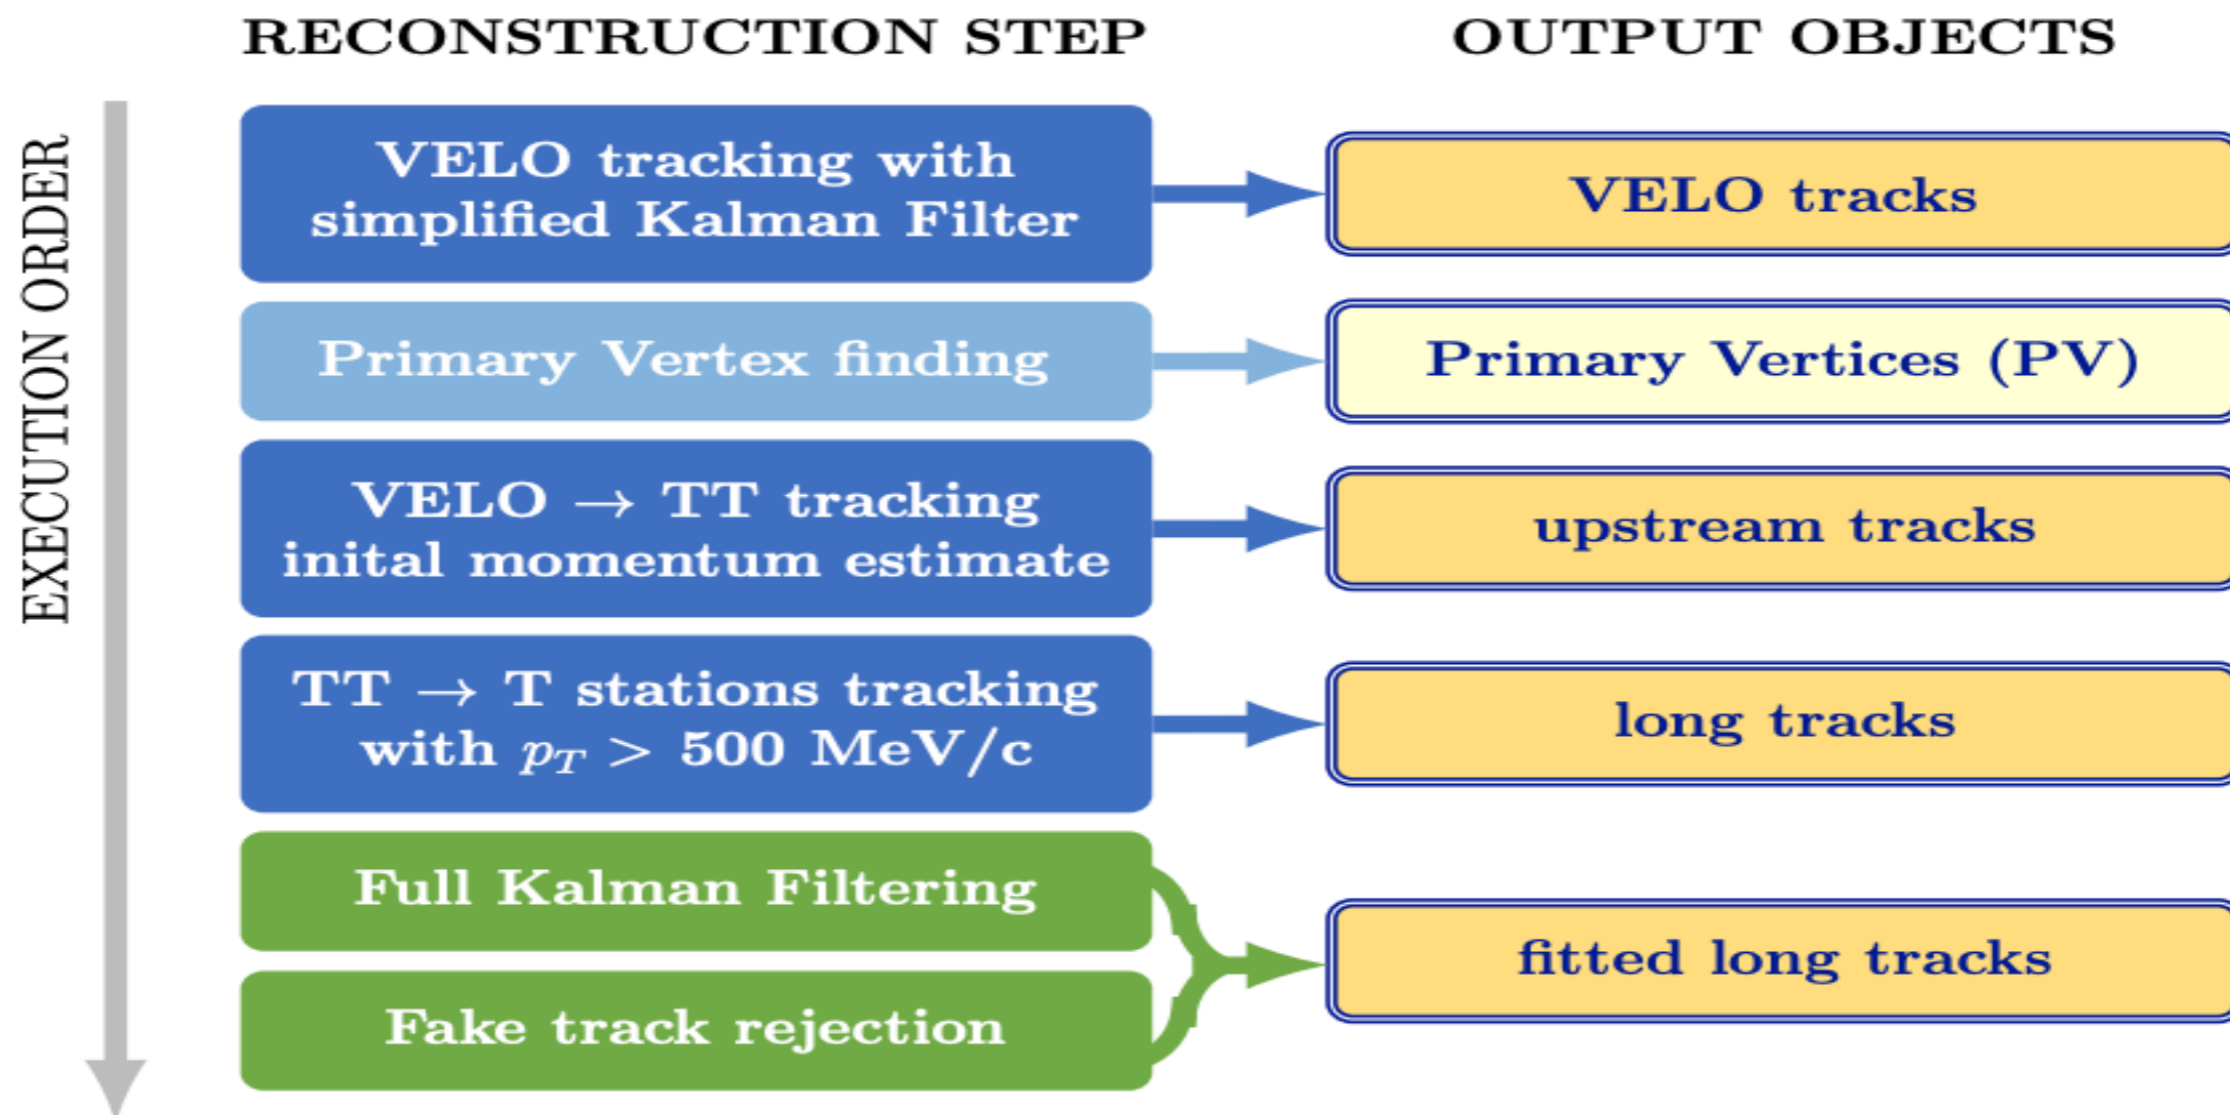

E.g., offline quality RICH PID for HLT2



Performance slightly *better* than the offline version from Run-I.


We'll see tomorrow how RICH PID is a crucial requirement for the *Turbo stream*.

Partial event reconstruction in HLT1





HLT1 selections

$m_B \sim 5 \text{ GeV}$
 $(\gamma_{CT})_B \sim 1 \text{ cm}$










$m_D \sim 2 \text{ GeV}$
 $(\gamma_{CT})_D \sim 4 \text{ mm}$




Most ($\approx 100 \text{ kHz}$) of the rate is taken by inclusive one- and two-track heavy-flavour lines.

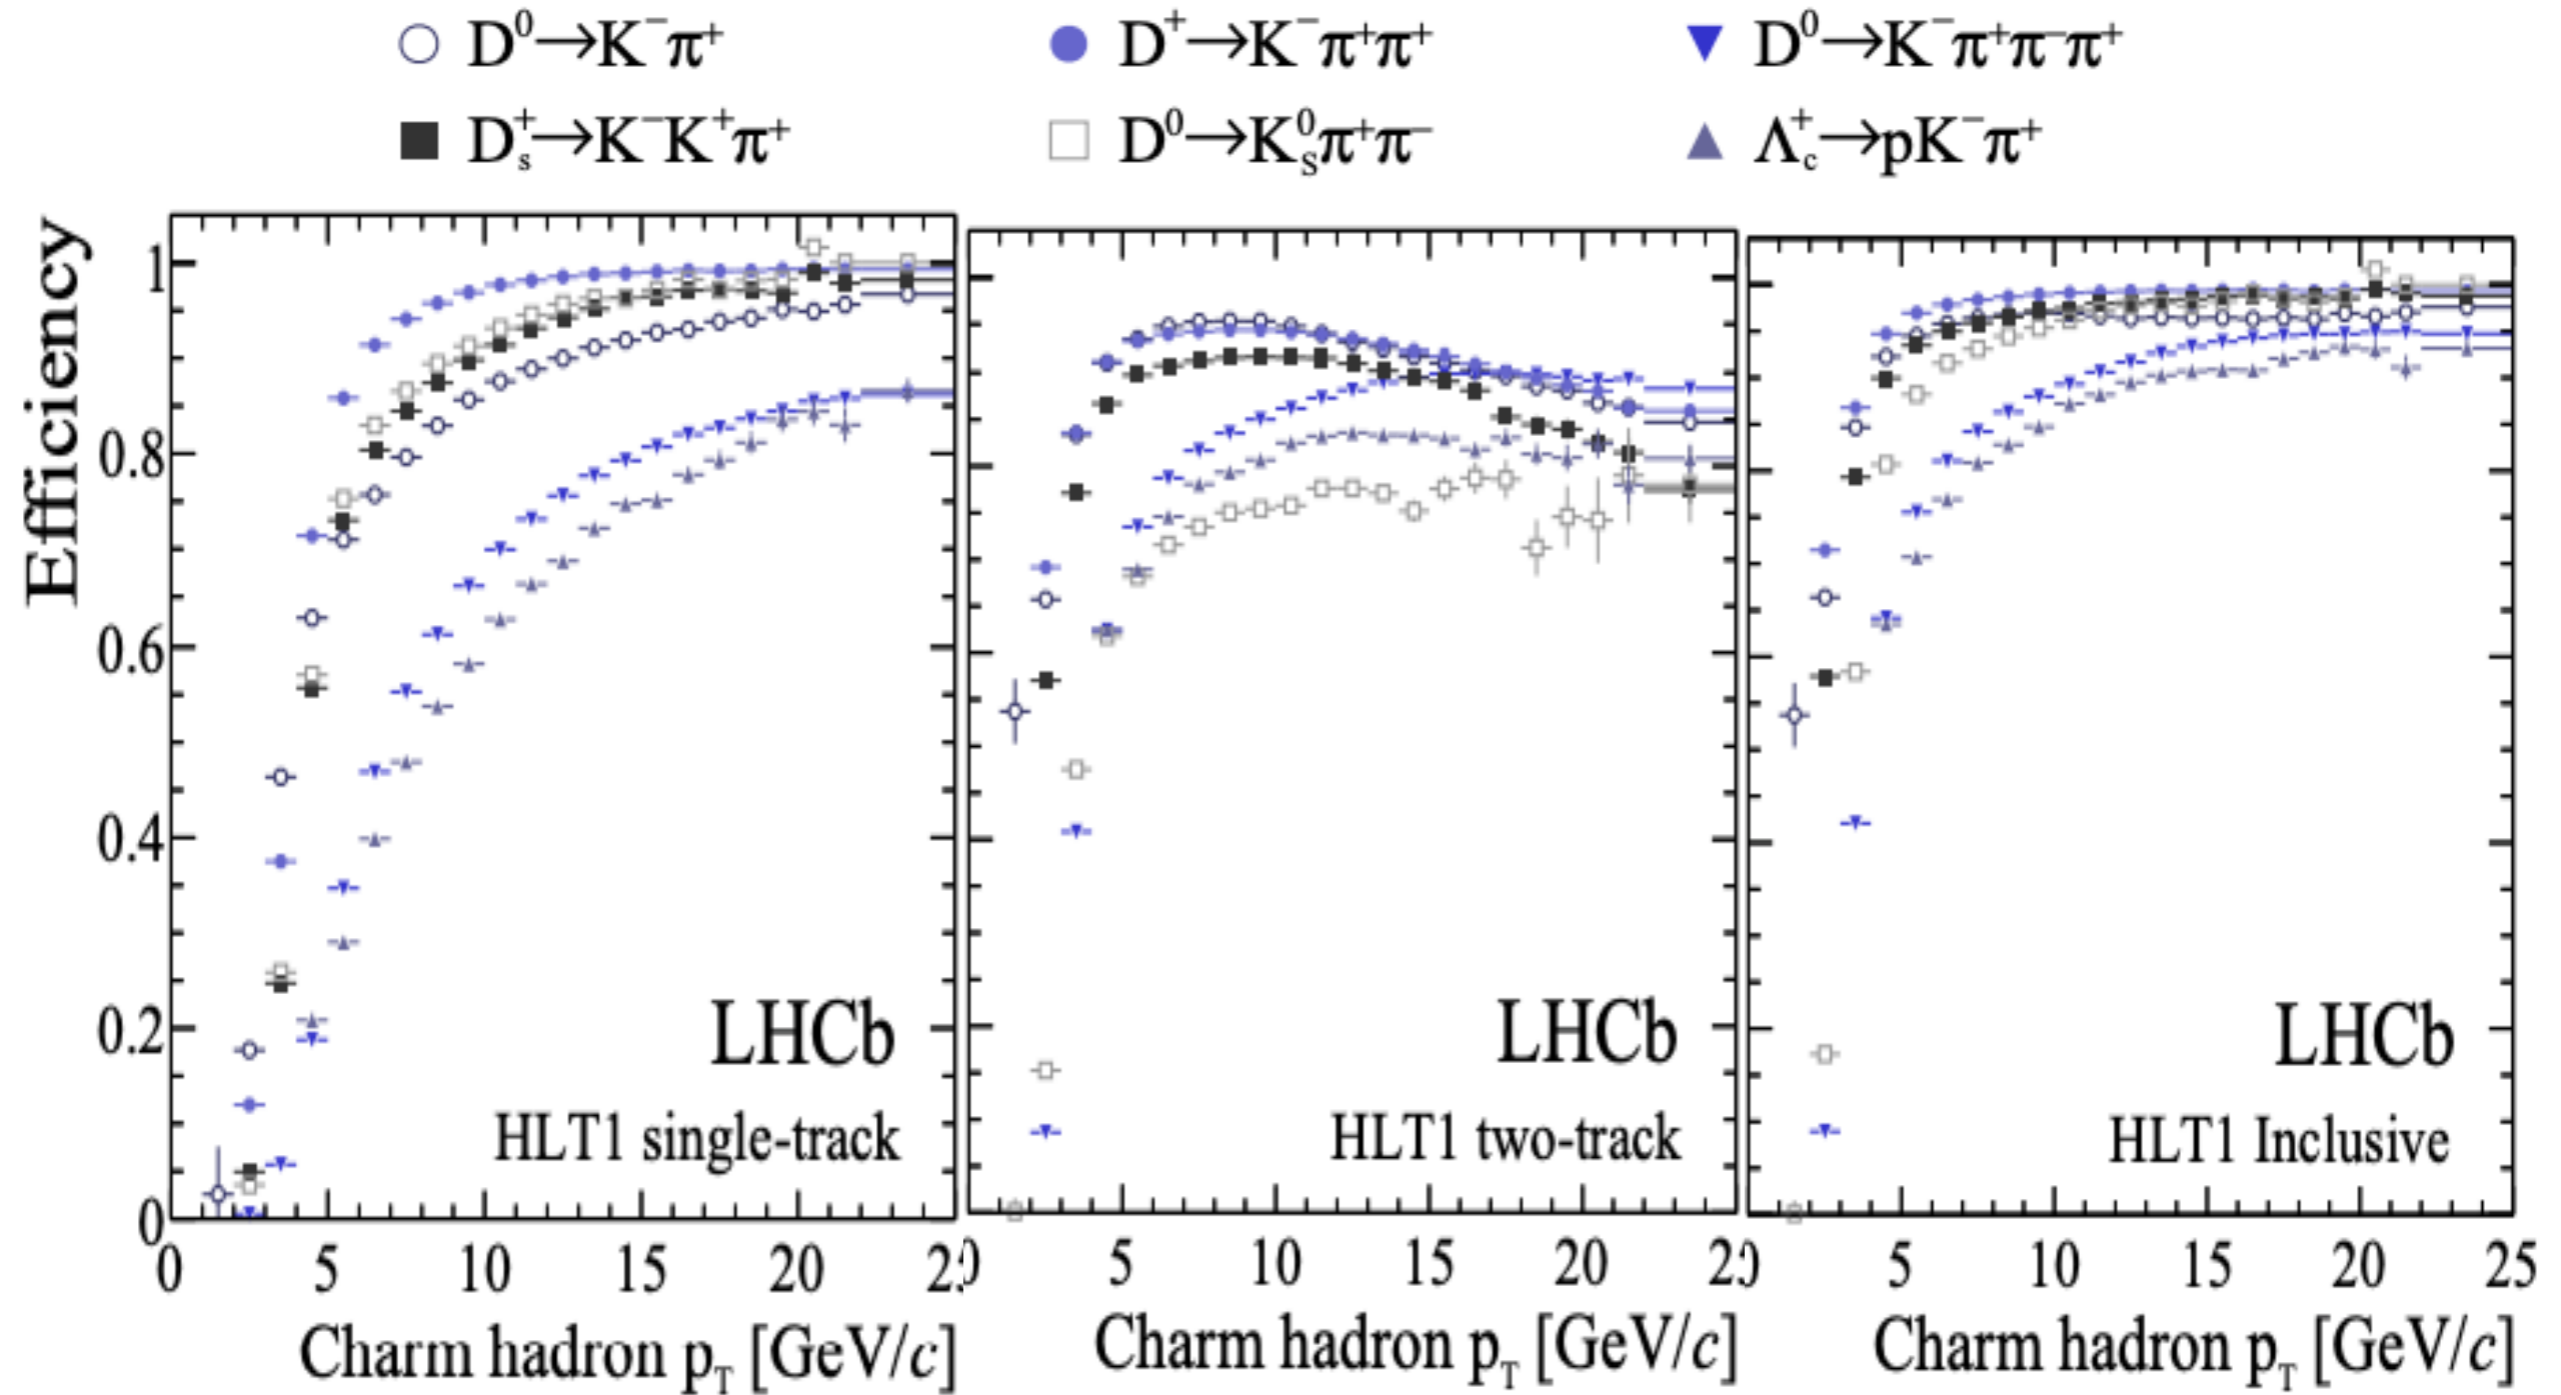
Inclusive HLT1 track line(s) performance



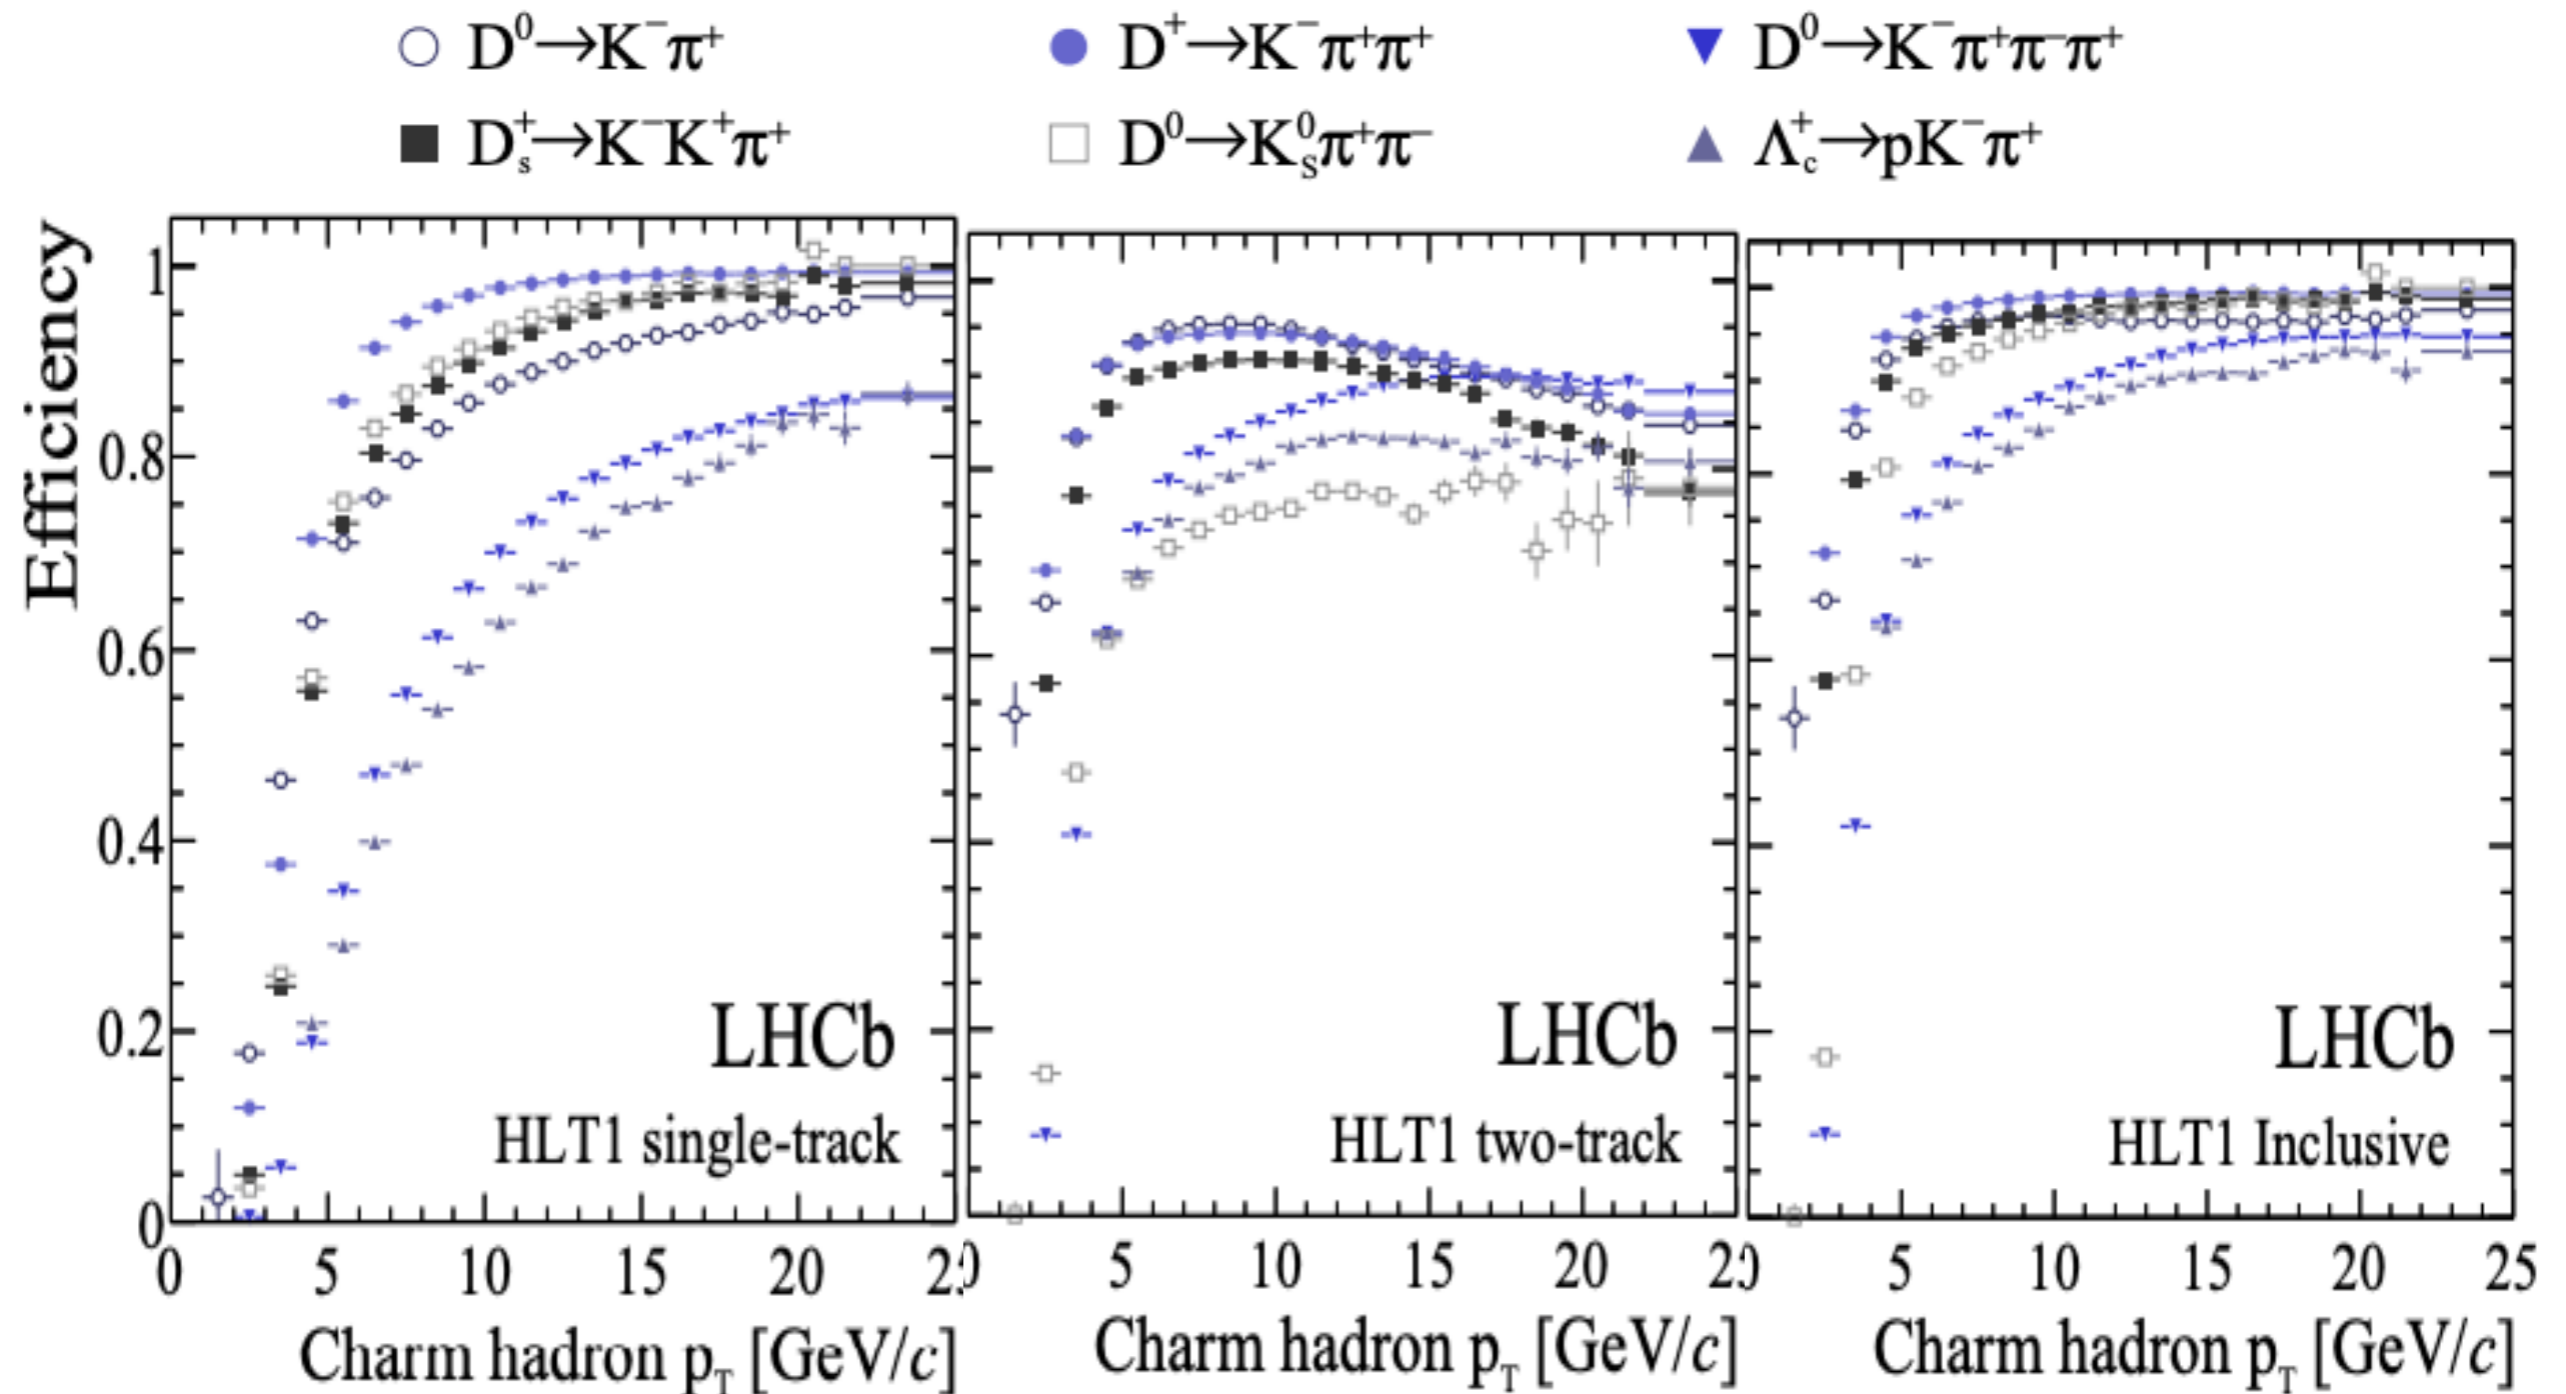
Inclusive HLT1 track line(s) performance



Inclusive HLT1 track line(s) performance



Inclusive HLT1 track line(s) performance



There is an interesting interplay with the reconstruction because, e.g., the two-track line is only really useful if we can reconstruct tracks down to a p_T threshold of ~ 500 MeV.

LHCb Trigger Run 2

Bunch crossing rate

↓ 40 MHz

L0 Hardware trigger
high p_T/E_T signatures

↓ 1 MHz

High Level Trigger 1
partial event reconstruction

↓ 110 kHz

**10 PB
buffer**

**Alignment &
Calibration**

↓ 110 kHz

High Level Trigger 2
full event reconstruction


↓ 12.5 kHz

Storage

Software trigger


The 10 PB disk buffer in, e.g., 2017

- HLT1 output rate ~ 150 kHz
- HLT2 throughput ~ 80 kHz out-of-fill (and ~ 30 kHz in fill).
- Average machine efficiency 30-50%.




The 10 PB disk buffer in, e.g., 2017

- HLT1 output rate ~ 150 kHz
- HLT2 throughput ~ 80 kHz out-of-fill (and ~ 30 kHz in fill).
- Average machine efficiency 30-50%.



HLT2 selections

About 500 HLT2 “lines” by the end of Run-II




The zoo of b hadron decay modes

<http://pdglive.lbl.gov/>

Indentation is used to indicate a subchannel of a previous reaction. All resonant subchannels have been corrected for resonance branching fractions to the final state so the sum of the subchannel branching fractions can exceed that of the final state. For inclusive branching fractions, e. g., $B \rightarrow D^\pm X$, the values usually are multiplicities, not branching fractions. They can be greater than one.




Mode	Fraction (Γ_i / Γ)	Scale Factor/ Conf. Level	P (MeV/c)
Semileptonic and leptonic modes			
Γ_1	$\ell^+ \nu_\ell X$	[1] 84% (10.99 ± 0.28)%	
Γ_2	$e^+ \nu_e X_c$	(10.8 ± 0.4)%	
Γ_3	$D \ell^+ \nu_\ell X$	(9.7 ± 0.7)%	
Γ_4	$D_0^- \ell^+ \nu_\ell$	[1] (2.35 ± 0.09)%	2310
Γ_5	$D_*^- \tau^+ \nu_\tau$	(7.7 ± 2.5) × 10 ⁻³	1911
Γ_6	$D_*^- (2007)^0 \ell^+ \nu_\ell$	[1] (5.66 ± 0.22)%	2258
Γ_7	$D_*^- (2007)^0 \tau^+ \nu_\tau$	(1.88 ± 0.20)%	S=1.0 1839
Γ_8	$D^- \pi^+ \ell^+ \nu_\ell$	(4.4 ± 0.4) × 10 ⁻³	2306
Γ_9	$D_{\emptyset}^- (2420)^0 \ell^+ \nu_\ell, D_{\emptyset 0}^- \rightarrow D^- \pi^+$	(2.5 ± 0.5) × 10 ⁻³	
Γ_{10}	$D_{\emptyset}^- (2460)^0 \ell^+ \nu_\ell, D_{\emptyset}^- \rightarrow D^- \pi^+$	(1.53 ± 0.16) × 10 ⁻³	S=1.0 2065
Γ_{11}		(1.86 ± 0.26)%	
Γ_{12}	$D_*^- \pi^+ \ell^+ \nu_\ell$	(6.0 ± 0.4) × 10 ⁻³	2254
Γ_{13}	$D_{10}^- (2420)^0 \ell^+ \nu_\ell, D_{10}^- \rightarrow D_*^- \pi^+$	(3.03 ± 0.20) × 10 ⁻³	2084

Inclusive B trigger



$$m_{\text{corr}} \equiv \sqrt{m_{\text{inv}}^2 + |P_{T \text{ miss}}|^2 + |P_{T \text{ miss}}|}$$



E.g. for a 4-body B decay:




Used as key variable in with a “Bonzai” BDT [[JINST 8 \(2013\) P02013](#)] to provide a few kHz of \sim pure $b\bar{b}$.

Few kHz of inclusive B lines gives:

● $B^+ \rightarrow \bar{D}^0 \pi^+$ □ $B^0 \rightarrow D^- \pi^+$ ▲ $B^+ \rightarrow J/\psi(e^+e^-)K^+$ △ $B^+ \rightarrow J/\psi(\mu^+\mu^-)K^+$



- Lecture 1 outline:
 - Introduction
 - Trigger and DAQ basics
 - Low-level trigger
 - High-level trigger
- Lecture 2 outline:
 - Real-time analysis
 - Interaction with analysis and efficiencies
 - Upgrades



Backup slides follow from here...