Trigger & DAQ

Gerhard Raven
VU Amsterdam & Nikhef

- constraints
- physics
- architectures
- strategy
- examples

- resource optimization
- special cases
- interaction w. analysis
- upgrades

2015 CERN-Fermilab
Hadron Collider Physics School

Acknowledgements:
Niko Neufeld, Vladimir Gligorov, Paris Sphicas, Brian Dahmes, ISOTDAQ lecturers, and many, many others…

Trigger (particle physics)

From Wikipedia, the free encyclopedia

In particle physics, a **trigger** is a system that uses simple criteria to rapidly decide which events in a particle detector to keep when only a small fraction of the total can be recorded. Trigger systems are necessary due to real-world limitations in data storage capacity and rates. Since experiments are typically searching for "interesting" events (such as decays of rare particles) that occur at a relatively low rate, trigger systems are used to identify the events that should be recorded for later analysis. Current accelerators have event rates greater than 1 MHz and trigger rates that can be below 10 Hz. The ratio of the trigger rate to the event rate is referred to as the selectivity of the trigger. For example, the Large Hadron Collider (LHC) has an event rate of 20 MHz ($2 \cdot 10^7$ Hz), and the Higgs boson is expected to be produced there at a rate of roughly 1 Hz. The LHC detectors can manage a readout of a few hundred events per second. Therefore the minimum selectivity required is 10^{-5}, with much stricter requirements for the data analysis afterwards.[1]

- Simple criteria
- Rapidly decide
- Selective
- Small fraction can be recorded

Small Fraction

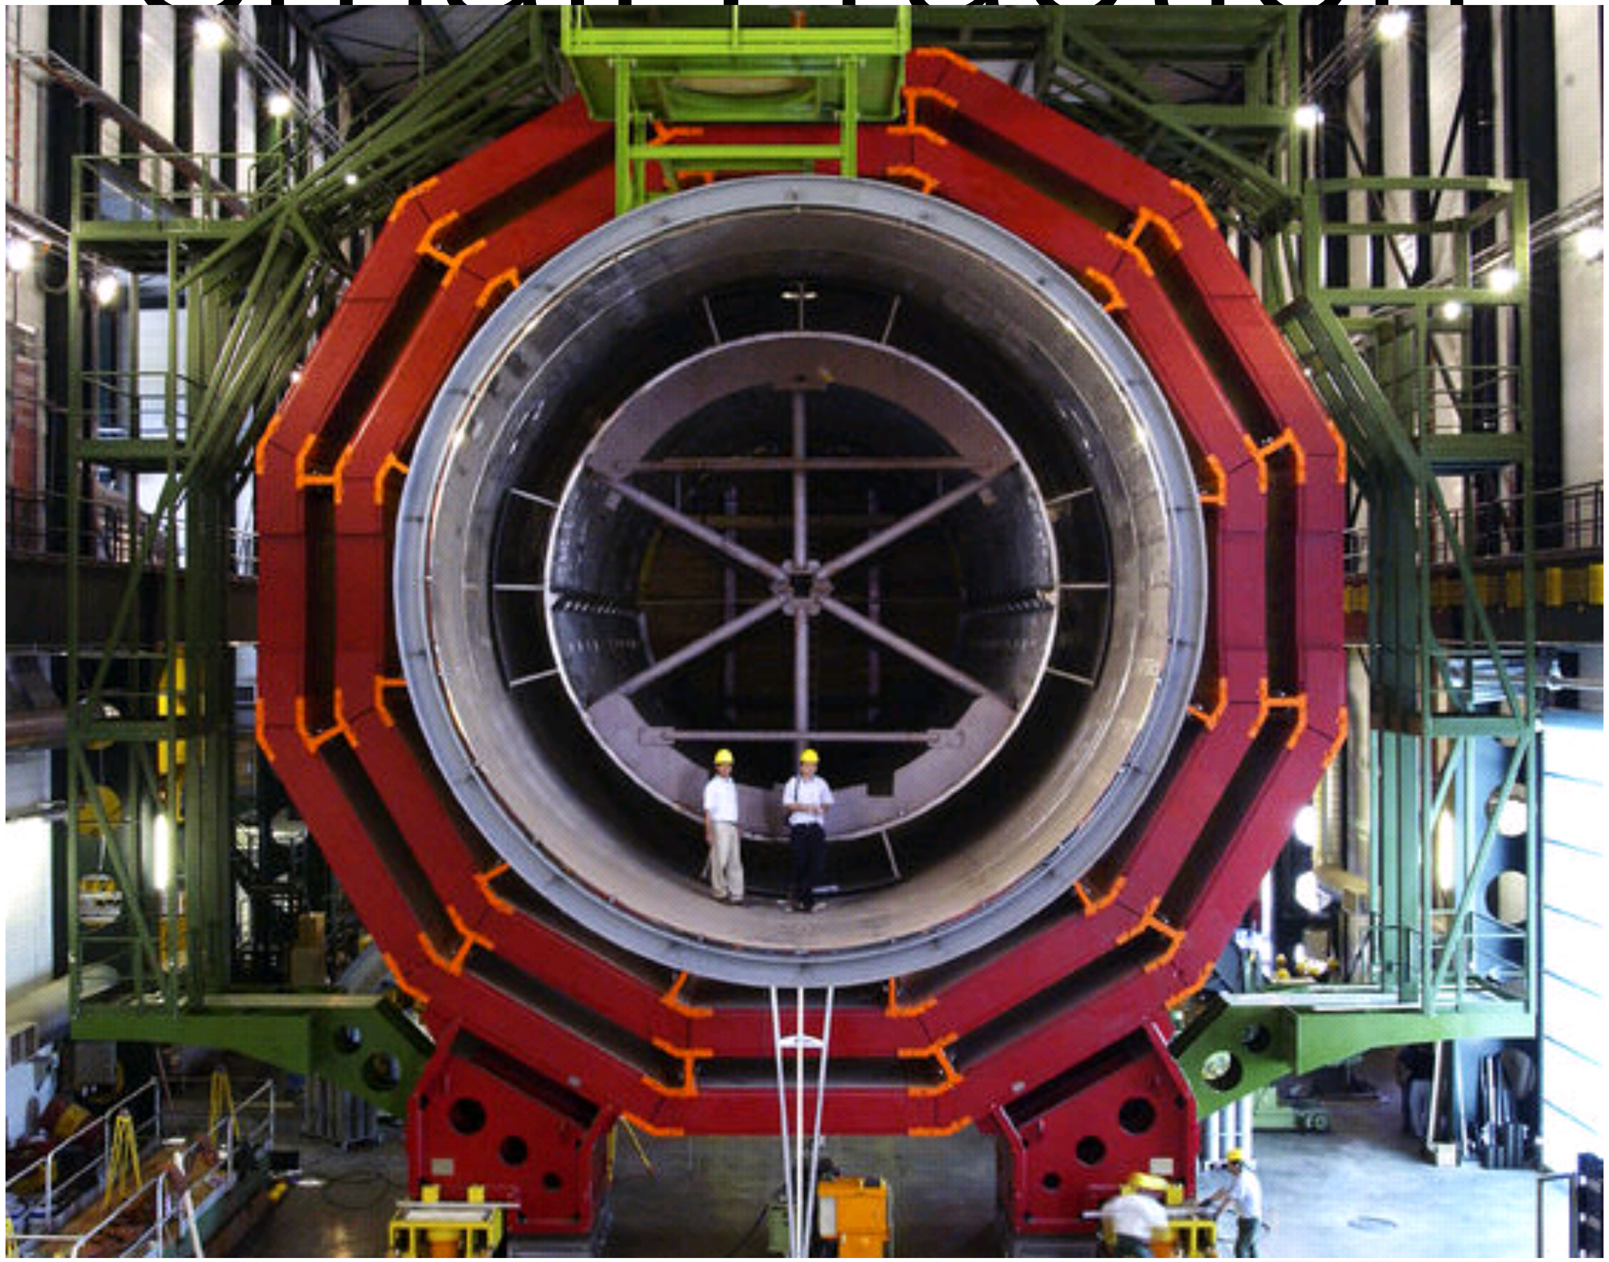

The CMS Detector

Small Fraction

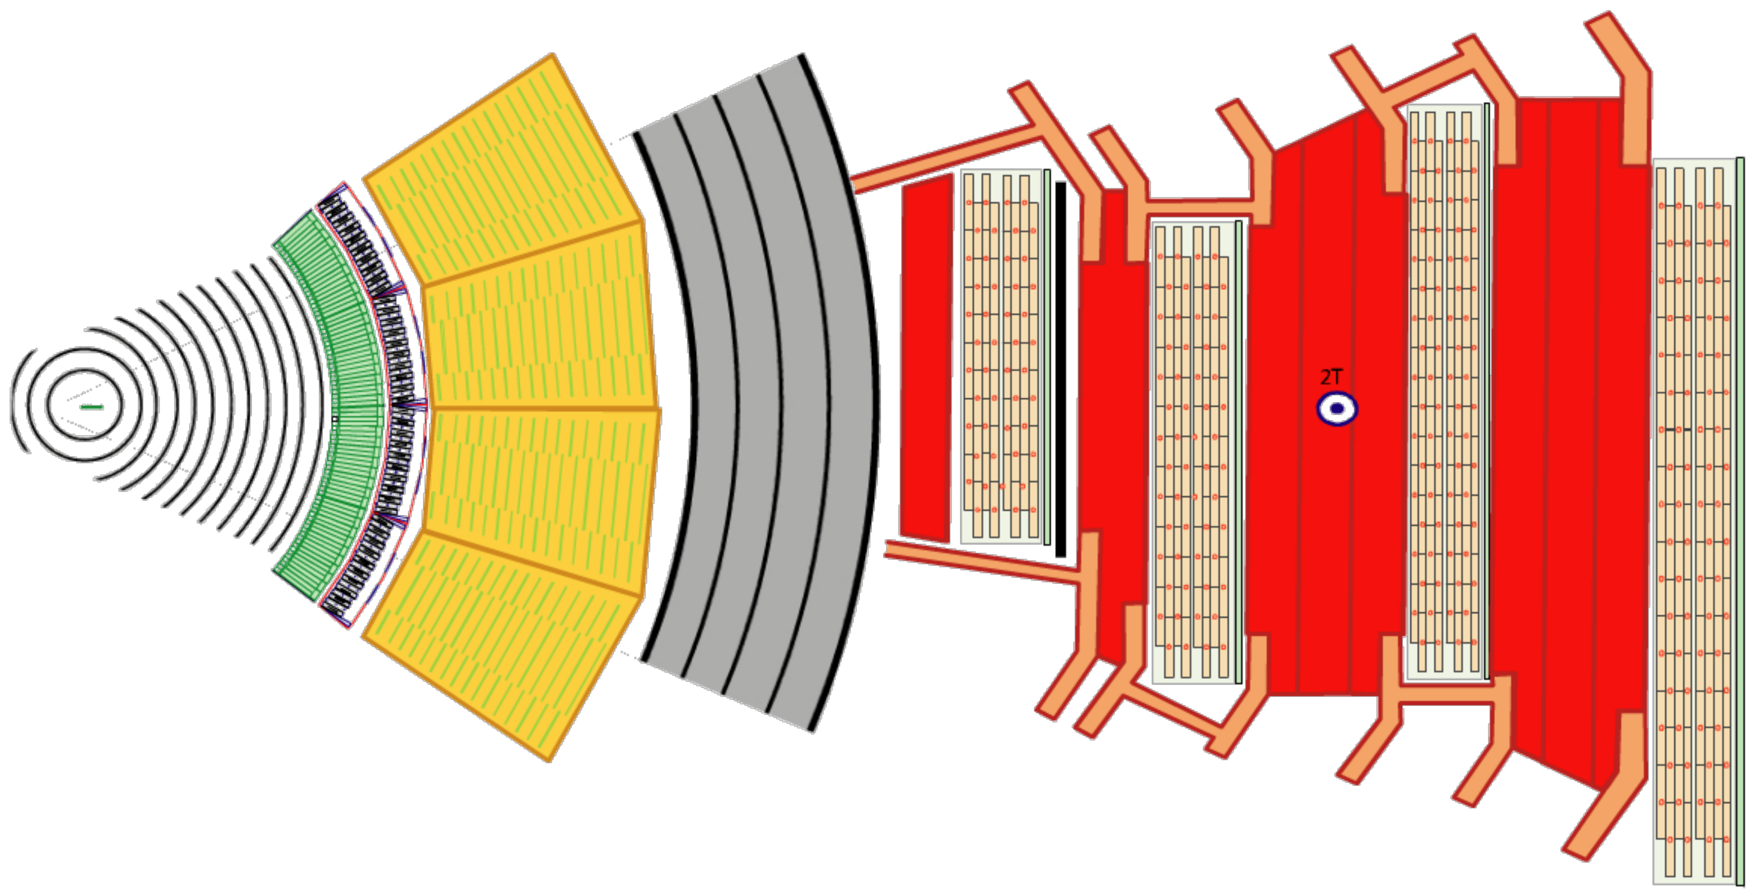

The CMS Detector

Small Fraction

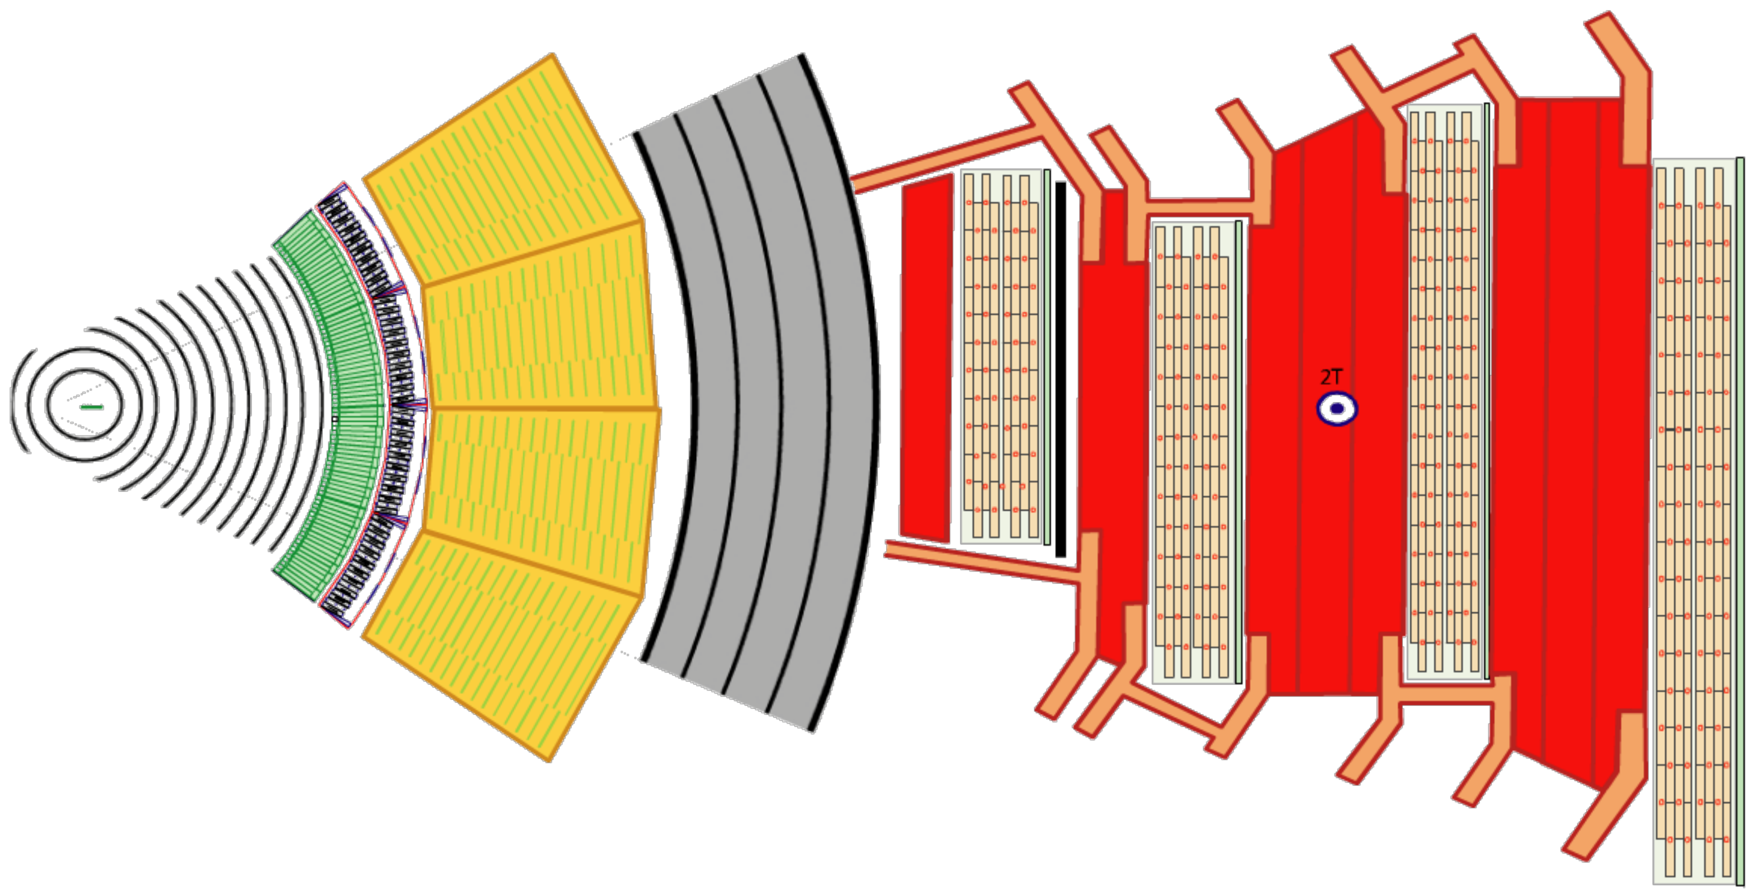

Small Fraction

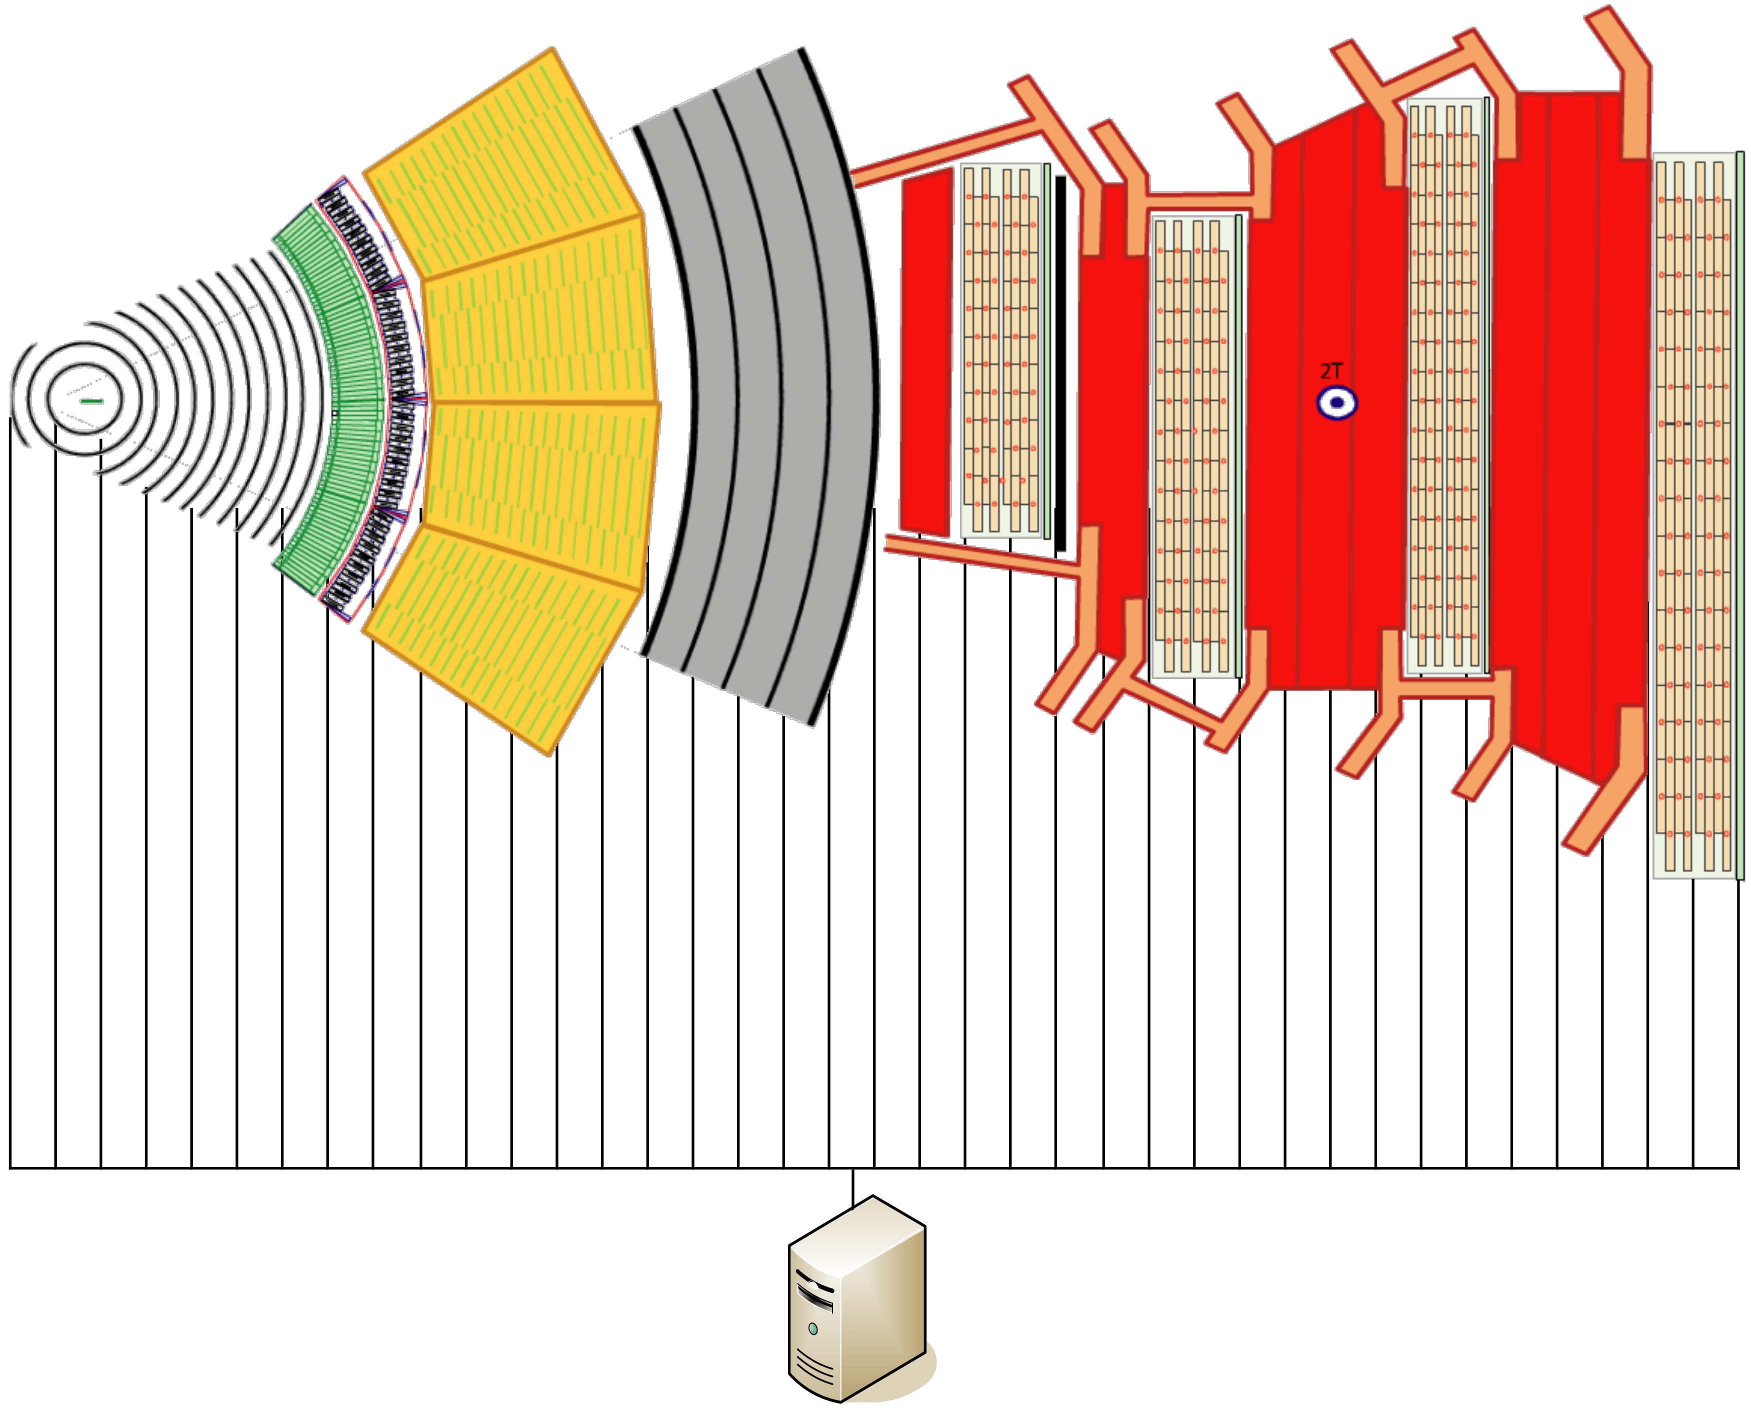

Small Fraction

- CMS tracker alone: 70 Mpixels + 10M strips @ 40 MHz ≥ 3200 Tb / s
- No (affordable) DAQ system can read this out
- If it could, think about the cables required…
- … and the power, and cooling, …

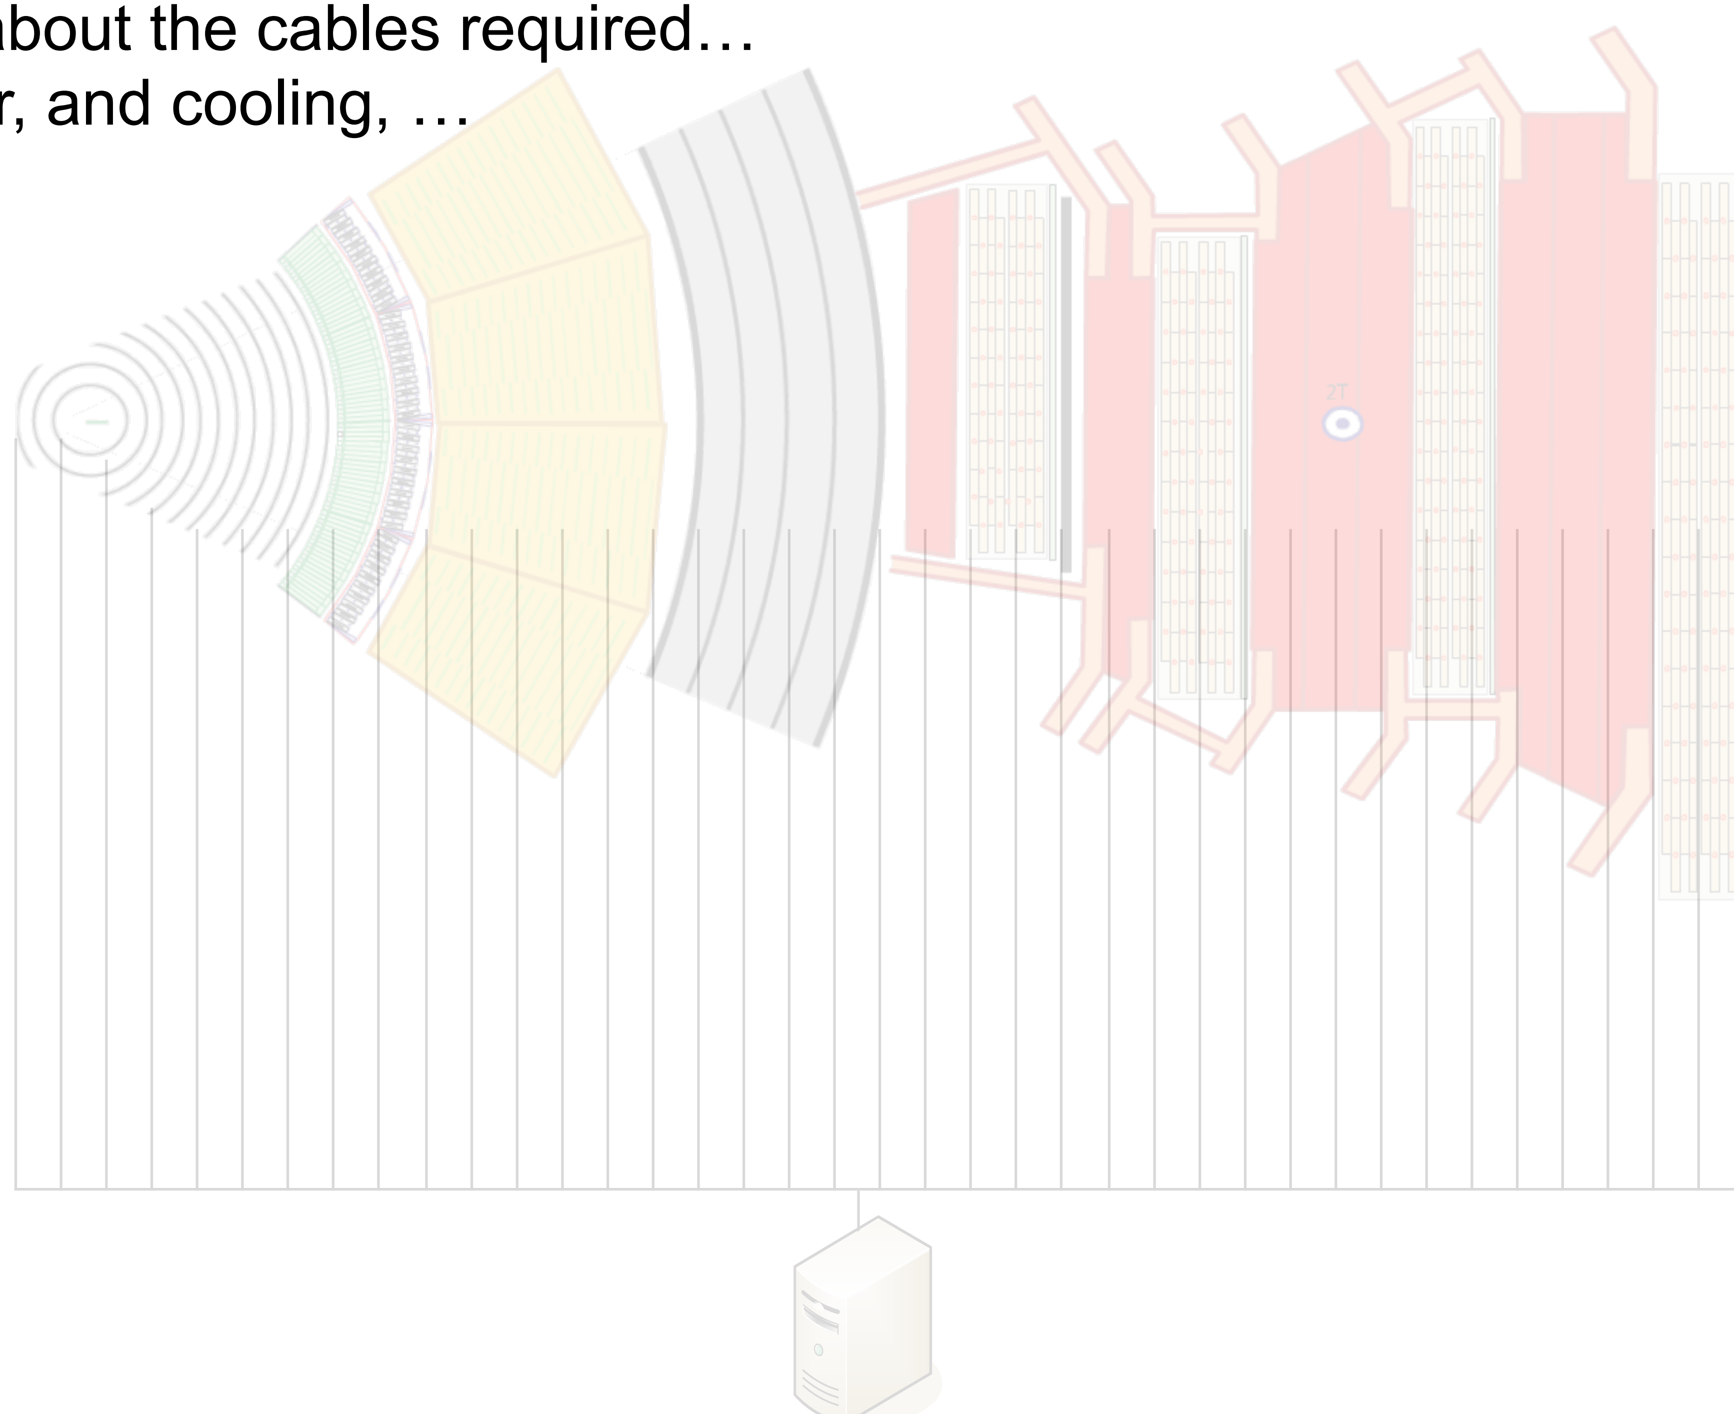

Small Fraction

- CMS tracker alone: 70 Mpixels + 10M strips @ 40 MHz ≥ 3200 Tb / s
- No (affordable) DAQ system can read this out
- If it could, think about the cables required…
- … and the power, and cooling, …

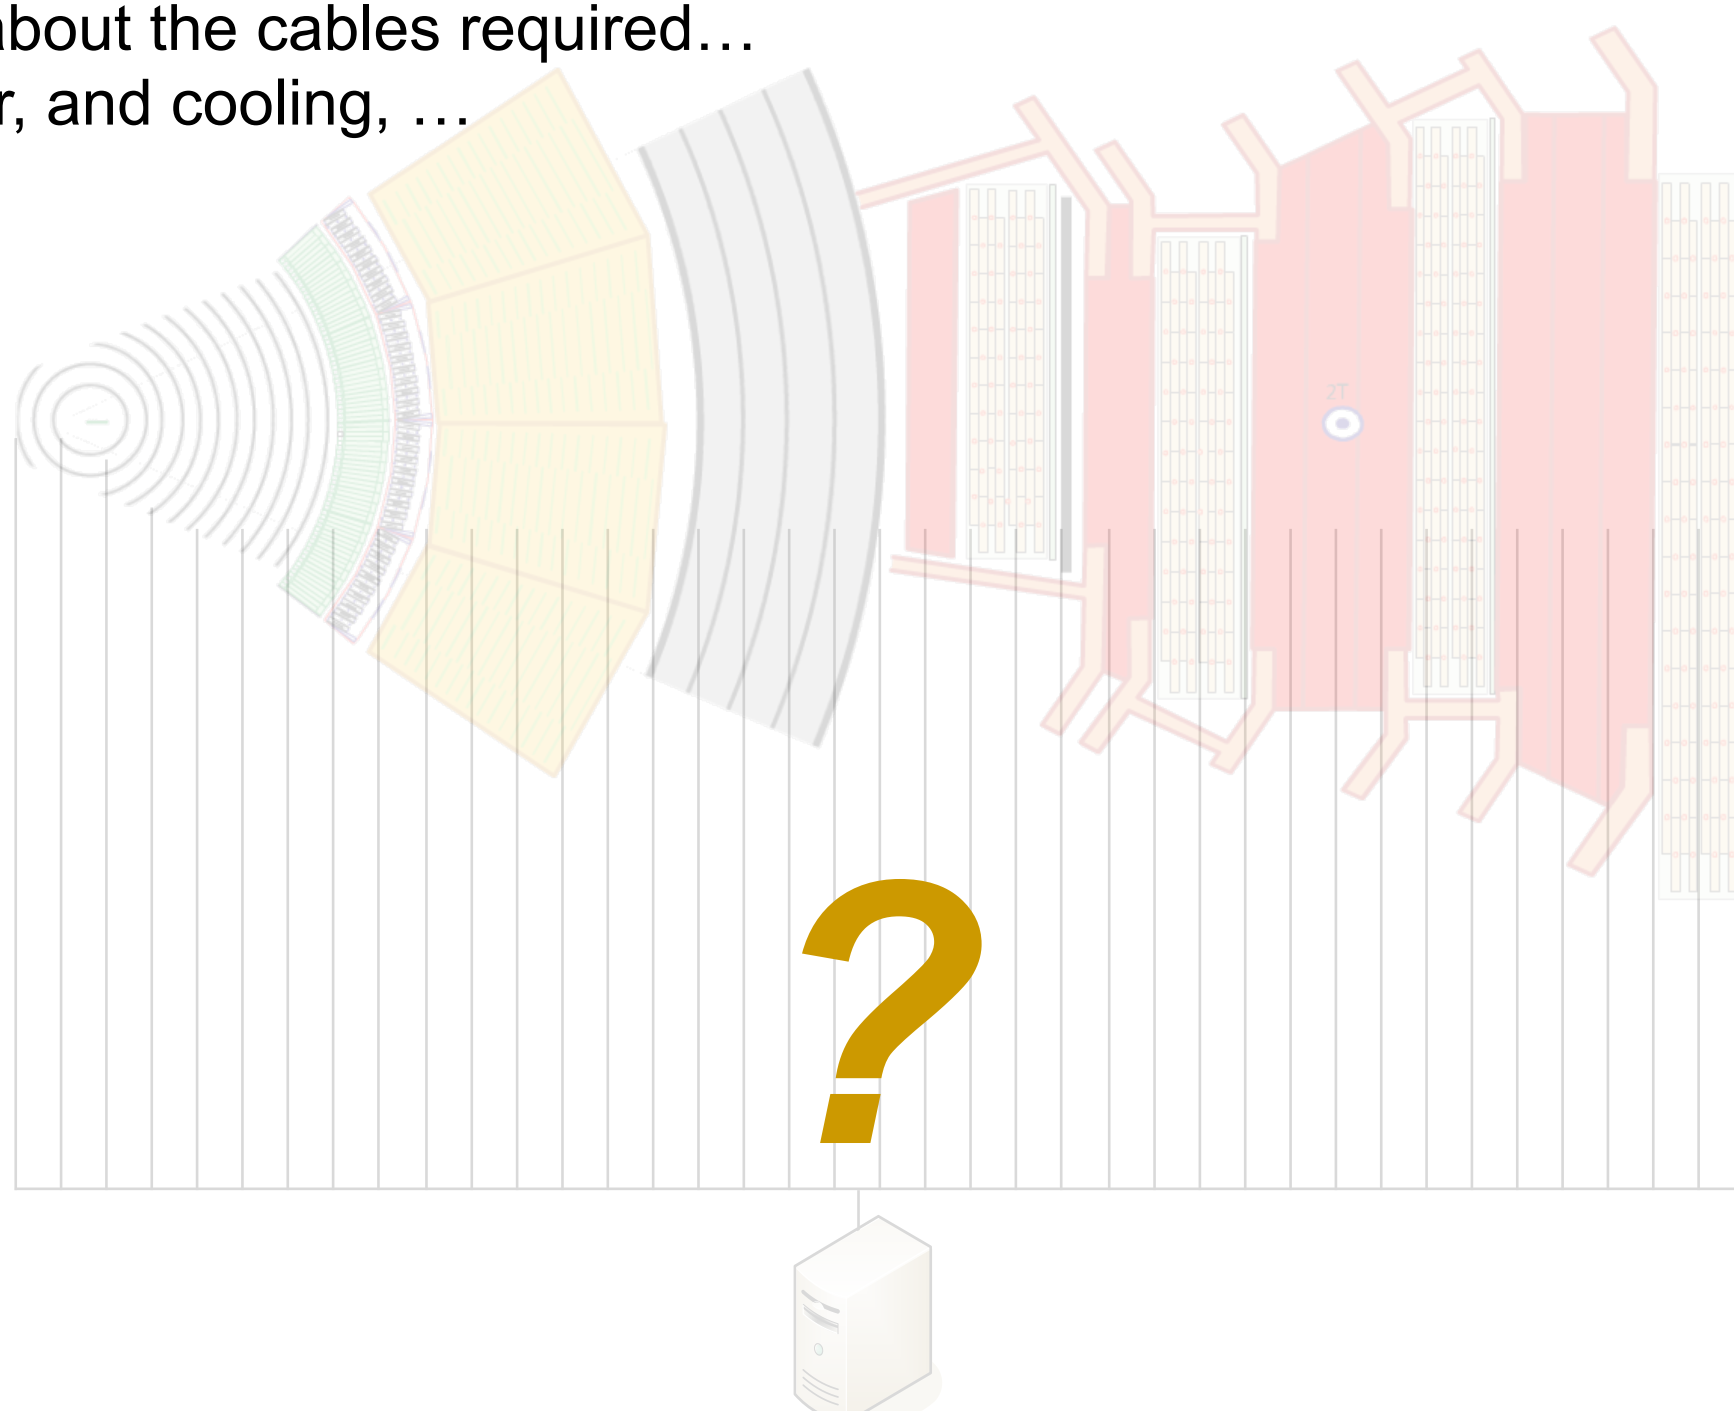

Bottlenecks!

digitization →

data transfer →

event building →

processing →

permanent storage →

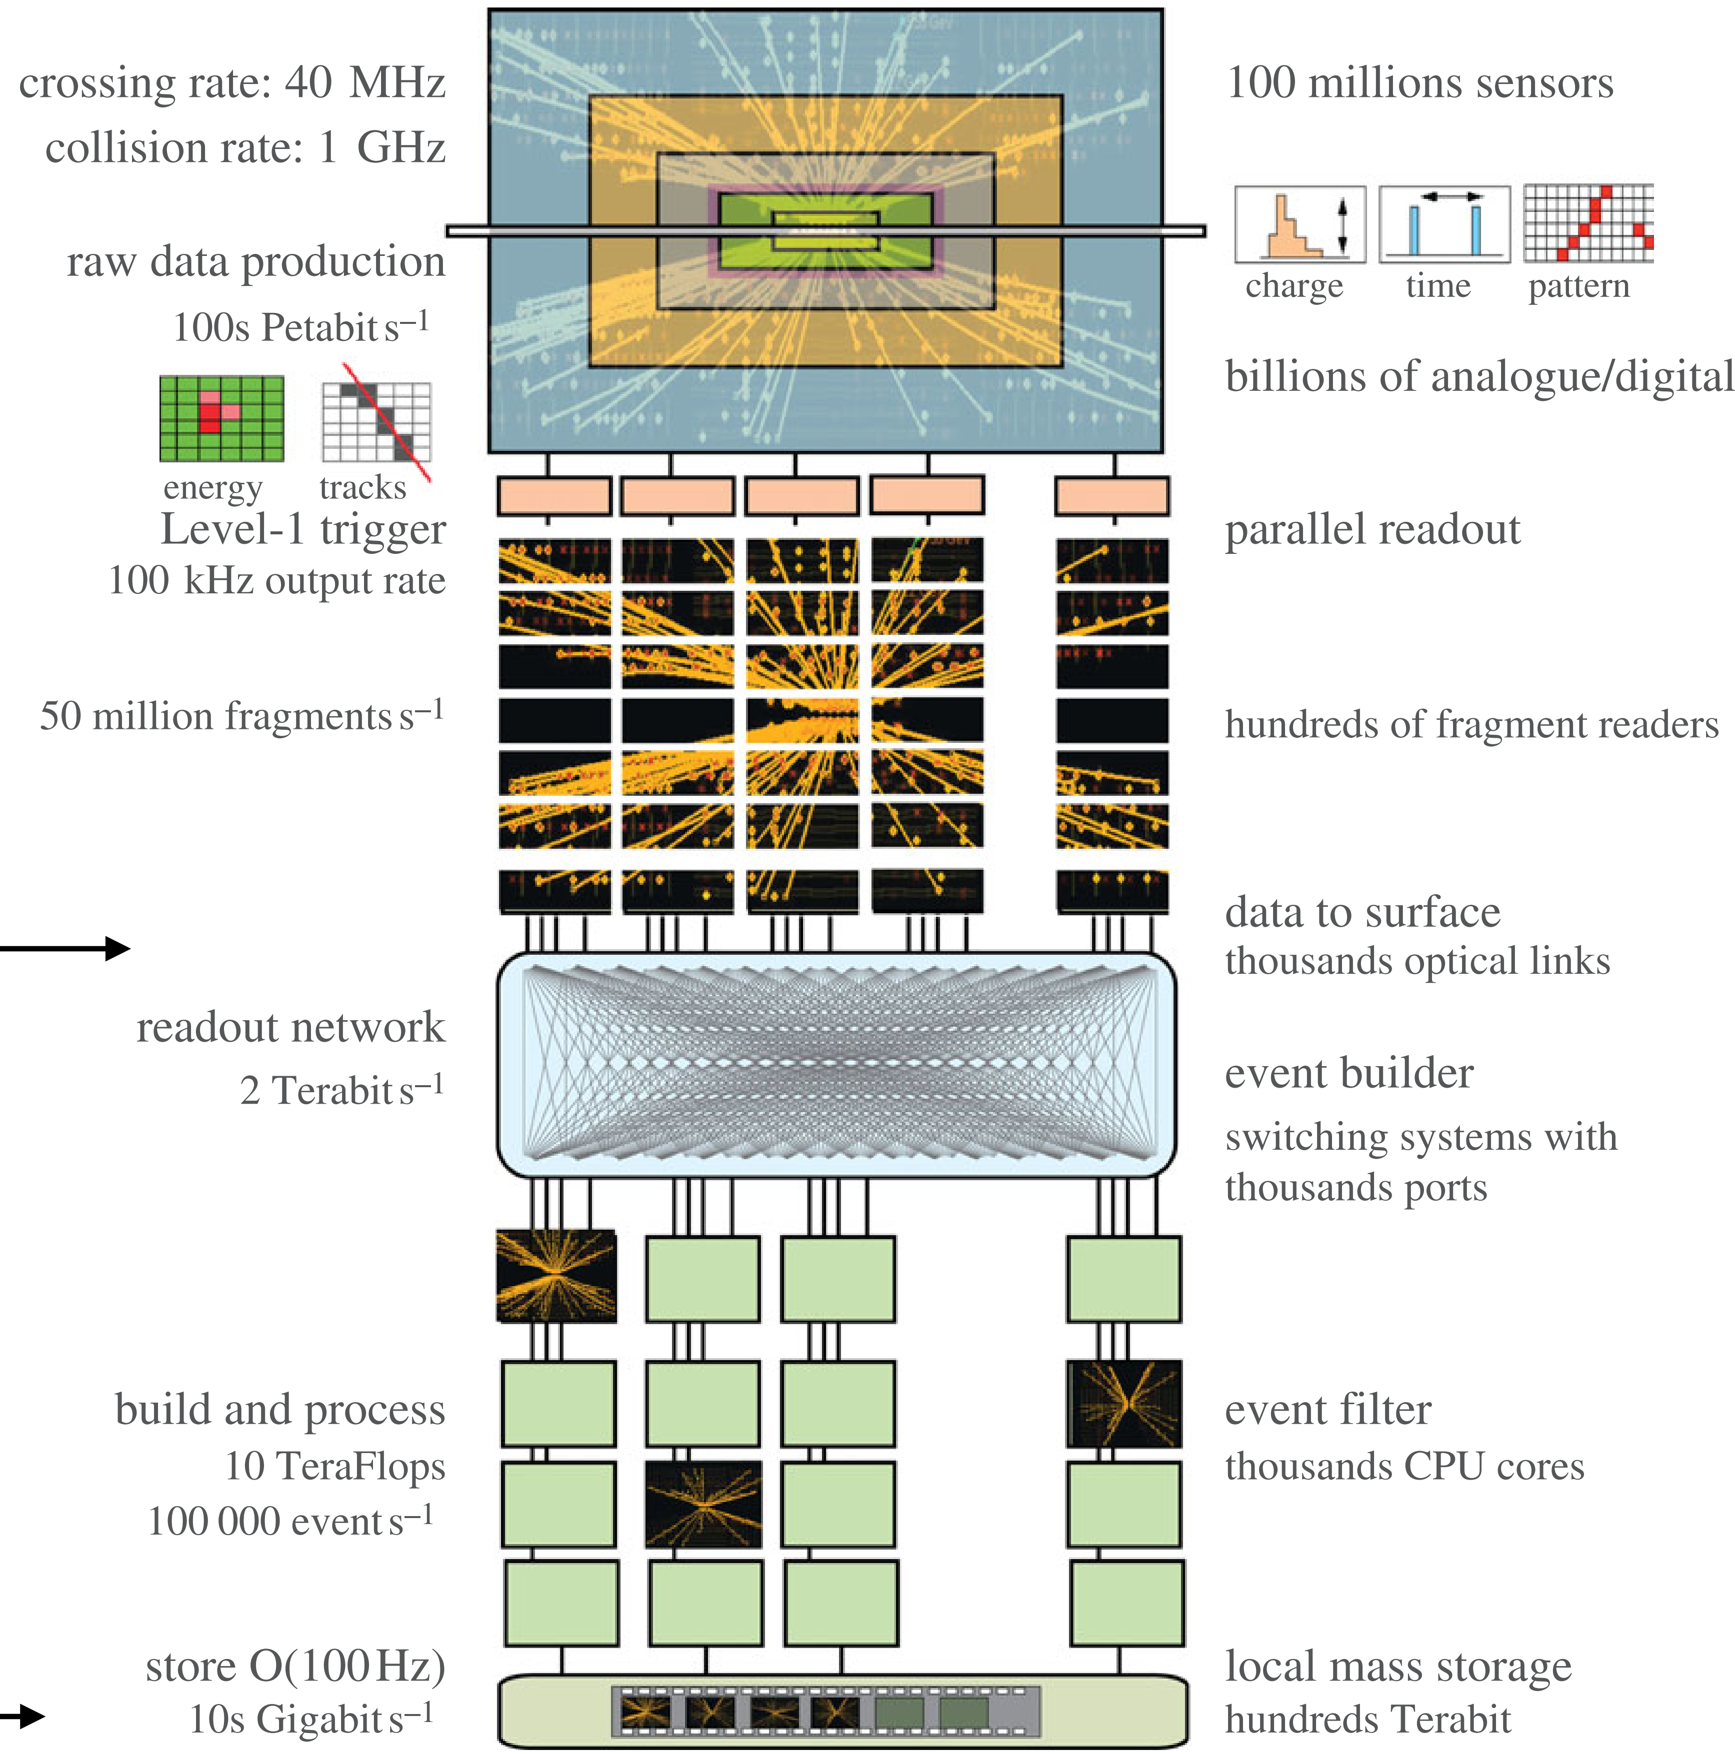

Be *Selective*

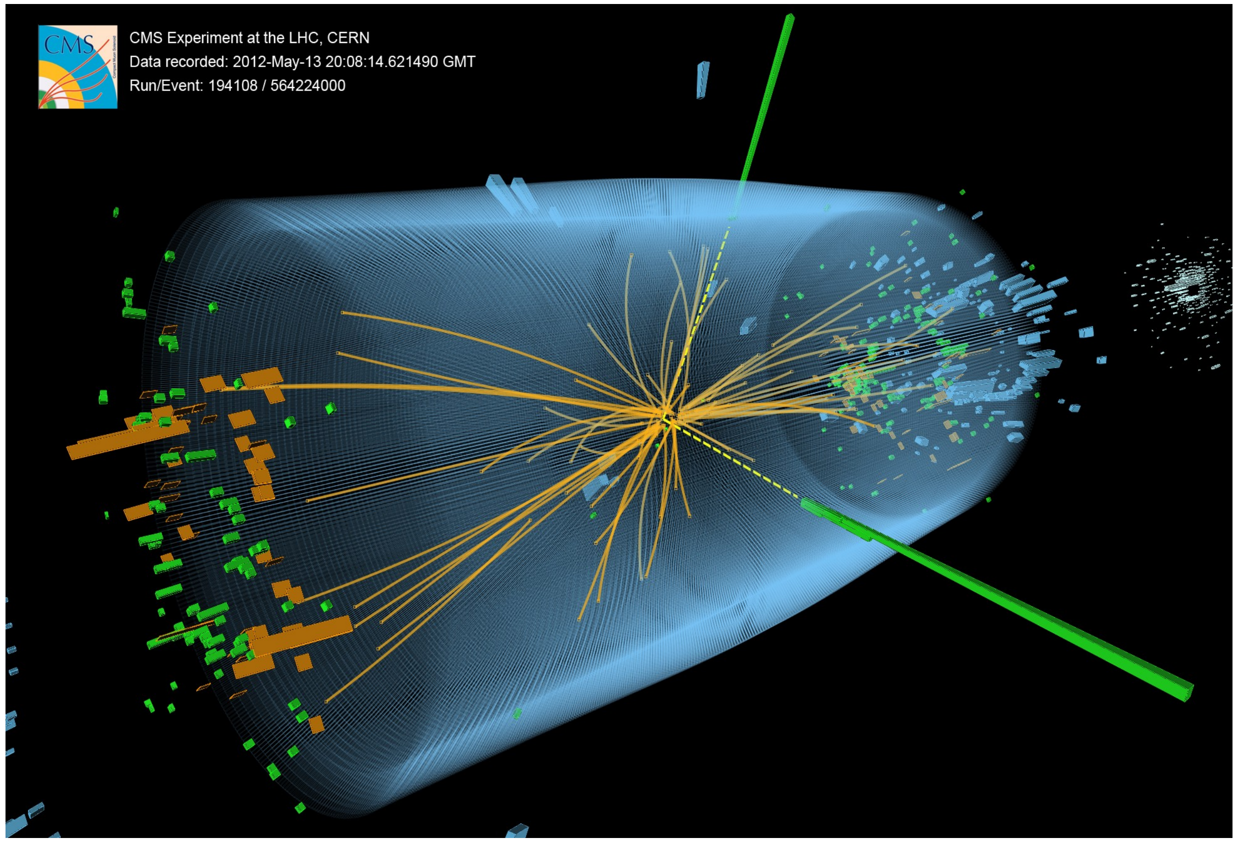

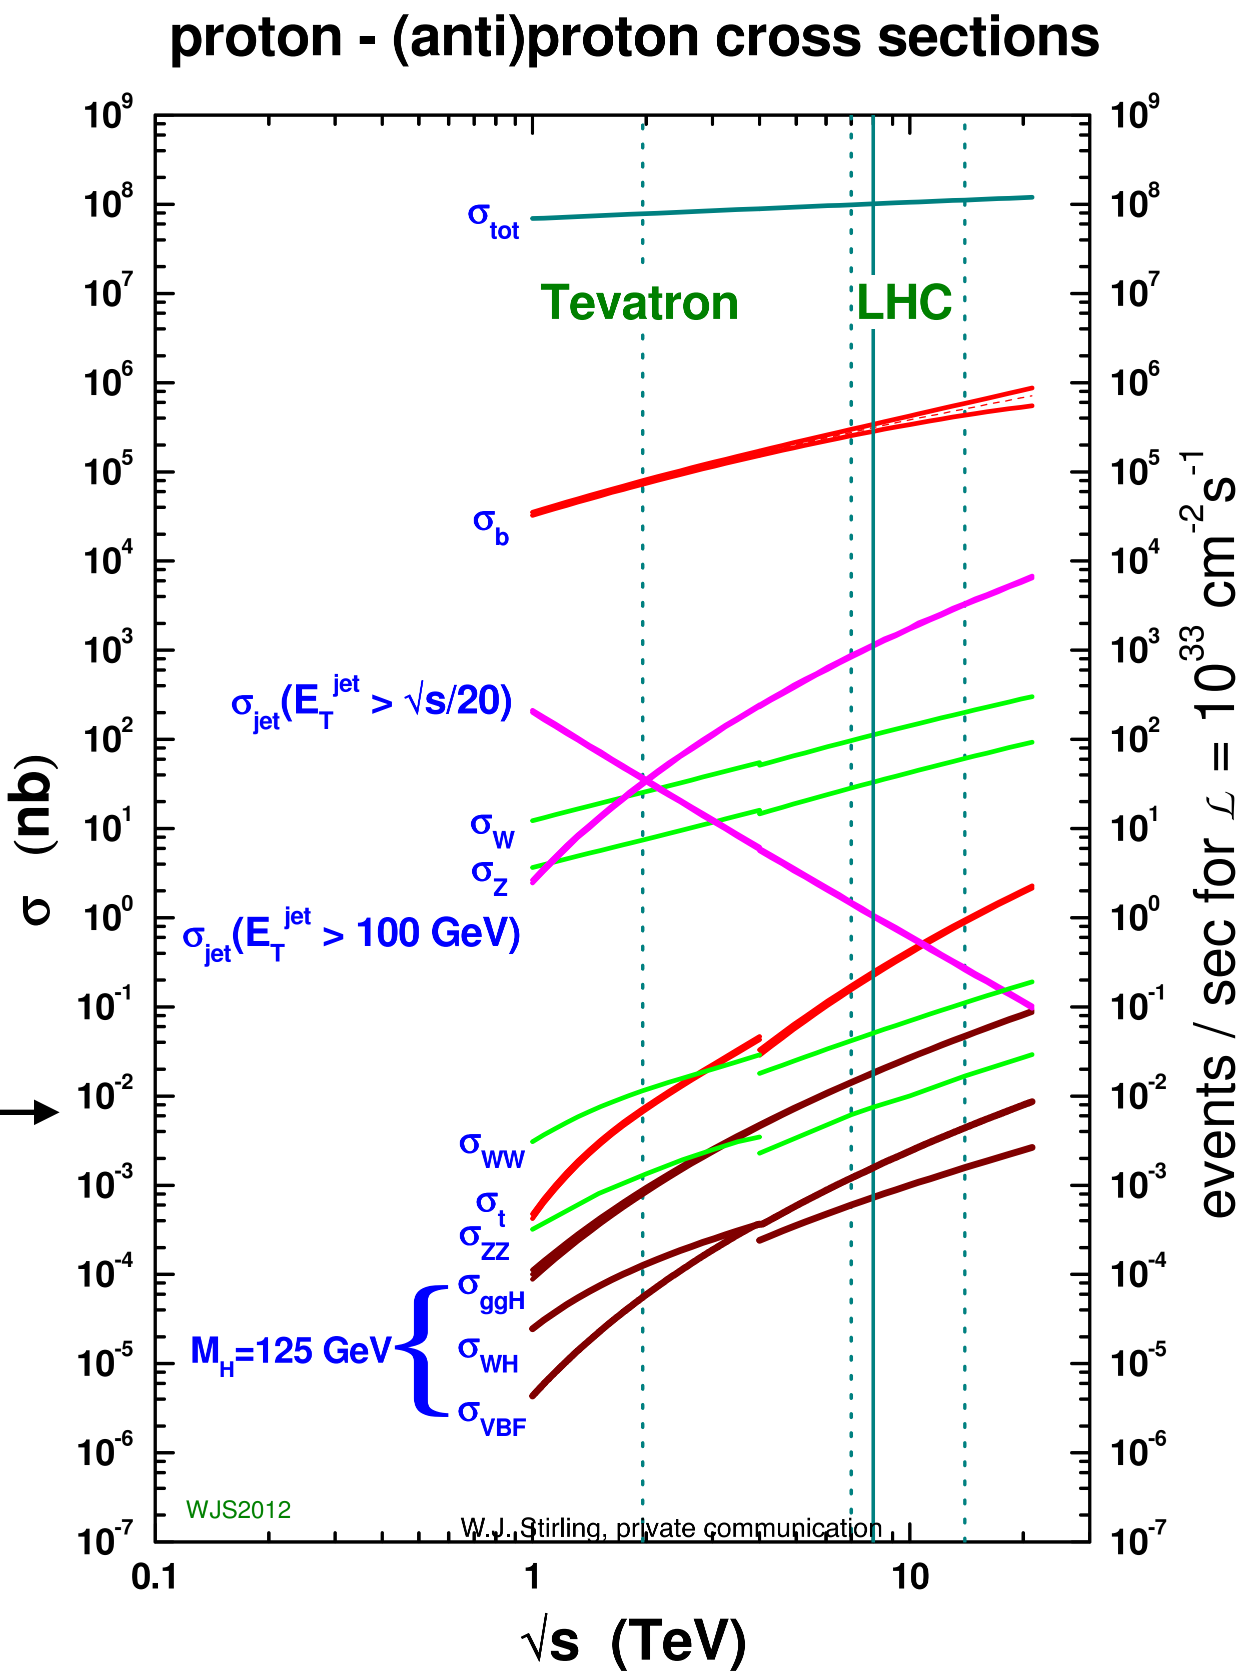

One (125 GeV/c^2) Higgs for every 10^{10} pp interactions

Selective in Perspective

drop: 0.1 ml = 10^{-7} m^3

olympic swimming pool: 50x25x2 m^3

Higgs : min bias ~ 2 drops : olympic swimming pool

Some things are not so rare…

PRL **110,** 101802 (2013)   Selected for a Viewpoint in *Physics*   PHYSICAL REVIEW LETTERS   week ending 8 MARCH 2013

Observation of $D^0 - \bar{D}^0$ Oscillations

R. Aaij *et al.**
(LHCb Collaboration)

(Received 6 November 2012; published 5 March 2013)

We report a measurement of the time-dependent ratio of $D^0 \to K^+\pi^-$ to $D^0 \to K^-\pi^+$ decay rates in D^{*+}-tagged events using 1.0 fb^{-1} of integrated luminosity recorded by the LHCb experiment. We measure the mixing parameters $x'^2 = (-0.9 \pm 1.3) \times 10^{-4}$, $y' = (7.2 \pm 2.4) \times 10^{-3}$, and the ratio of doubly-Cabibbo-suppressed to Cabibbo-favored decay rates $R_D = (3.52 \pm 0.15) \times 10^{-3}$, where the uncertainties include statistical and systematic sources. The result excludes the no-mixing hypothesis with a probability corresponding to 9.1 standard deviations and represents the first observation of $D^0 - \bar{D}^0$ oscillations from a single measurement.

DOI: 10.1103/PhysRevLett.110.101802   PACS numbers: 12.15.Ff, 13.25.Ft, 14.40.Lb

8.4M events

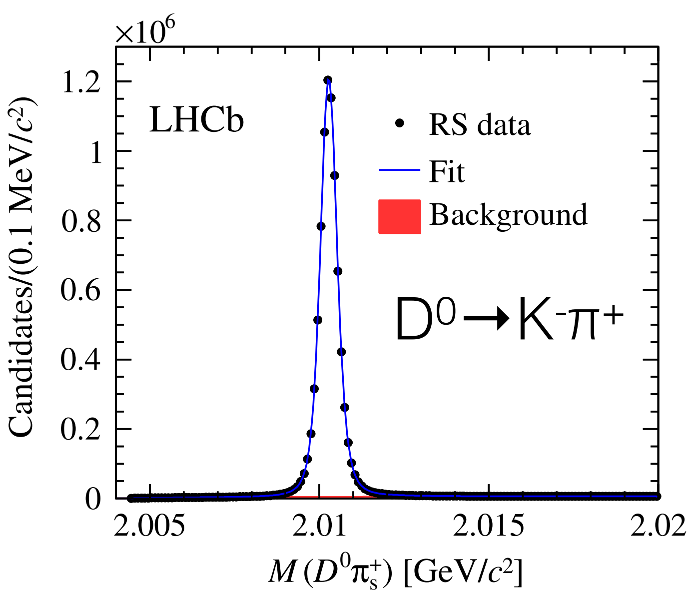

3.6K events

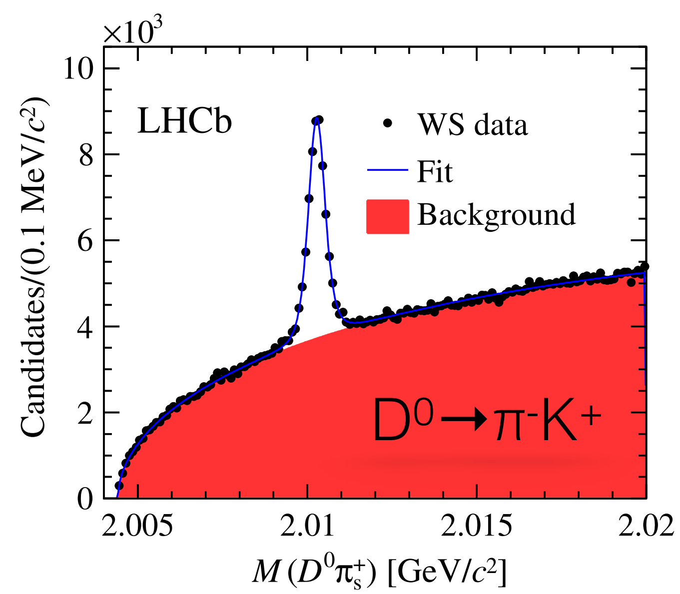

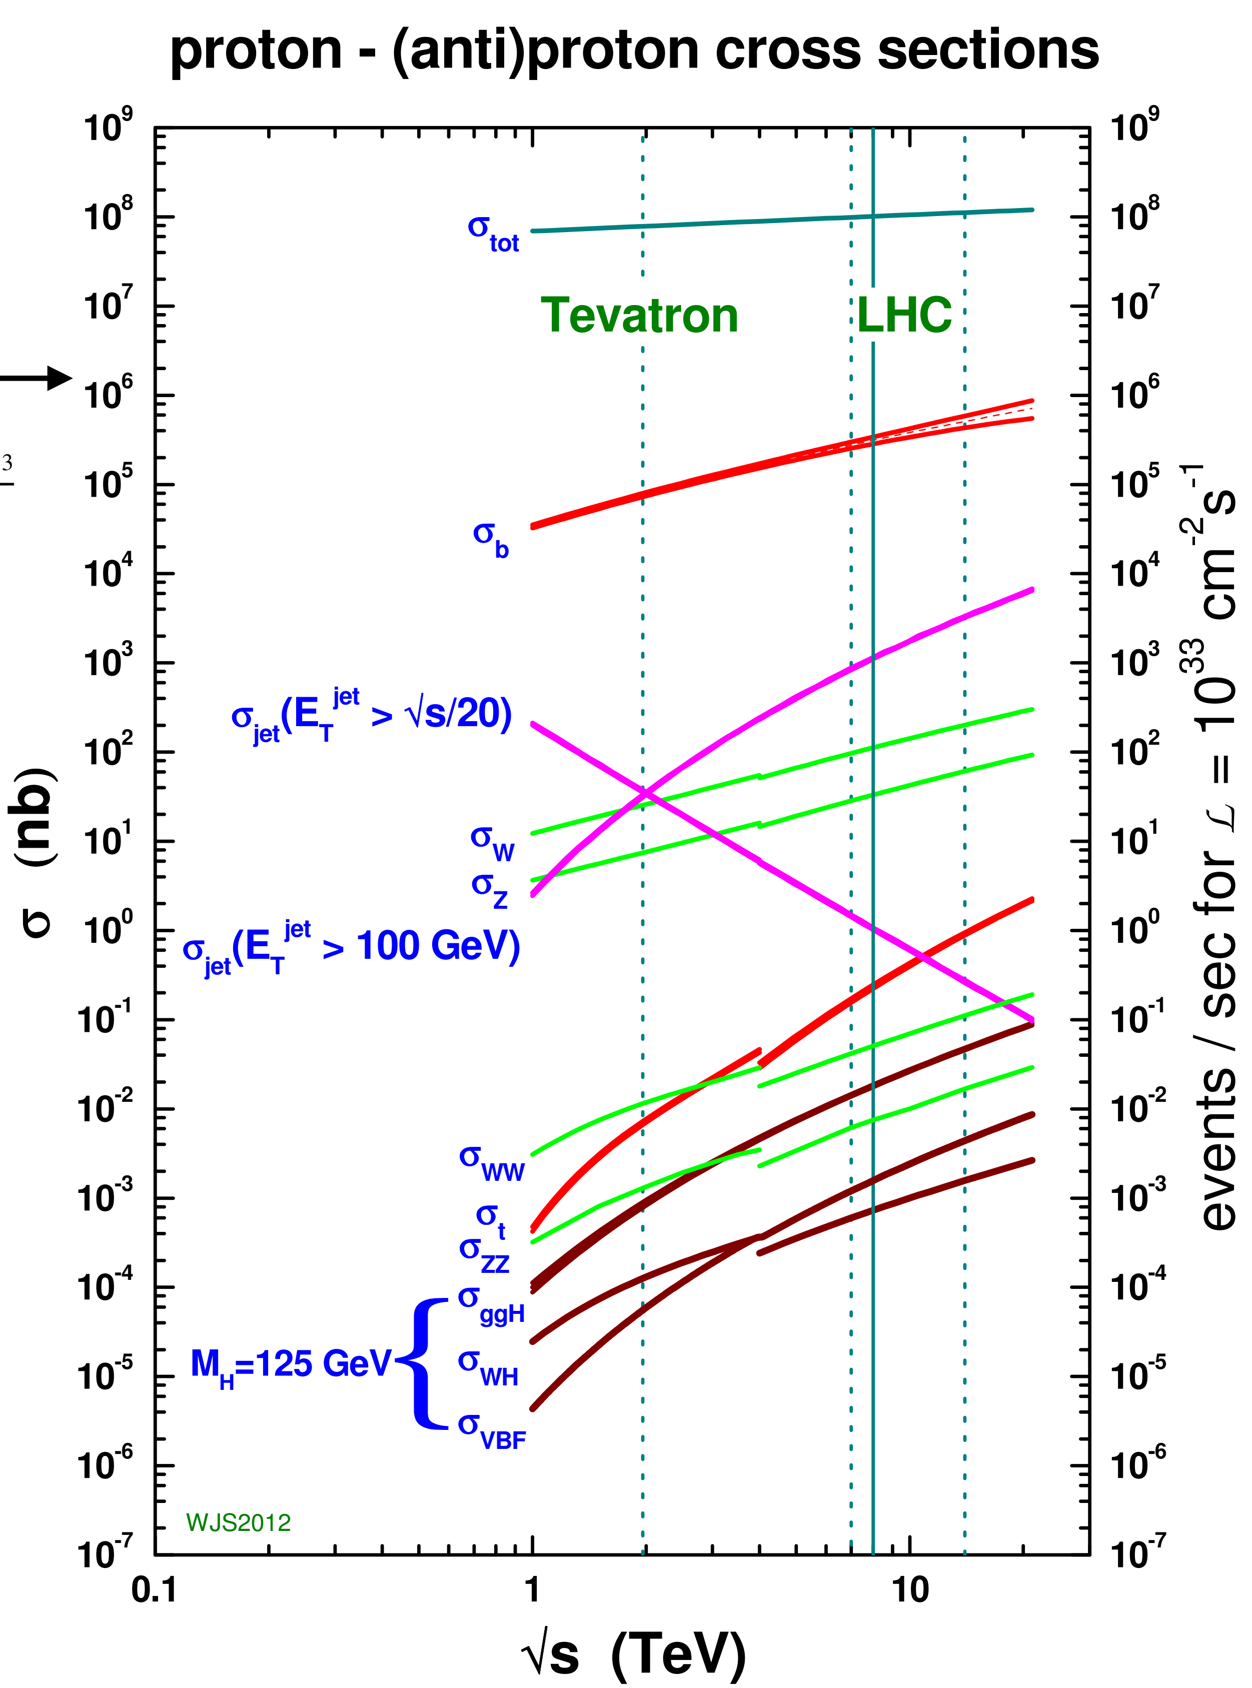

Some things are not so rare…

Selected for a Viewpoint in *Physics*
PRL **110,** 101802 (2013) PHYSICAL REVIEW LETTERS week ending 8 MARCH 2013

Observation of $D^0 - \bar{D}^0$ Oscillations

R. Aaij *et al.**
(LHCb Collaboration)

(Received 6 November 2012; published 5 March 2013)

We report a measurement of the time-dependent ratio of $D^0 \to K^+\pi^-$ to $D^0 \to K^-\pi^+$ decay rates in D^{*+}-tagged events using 1.0 fb^{-1} of integrated luminosity recorded by the LHCb experiment. We measure the mixing parameters $x'^2 = (-0.9 \pm 1.3) \times 10^{-4}$, $y' = (7.2 \pm 2.4) \times 10^{-3}$, and the ratio of doubly-Cabibbo-suppressed to Cabibbo-favored decay rates $R_D = (3.52 \pm 0.15) \times 10^{-3}$, where the uncertainties include statistical and systematic sources. The result excludes the no-mixing hypothesis with a probability corresponding to 9.1 standard deviations and represents the first observation of $D^0 - \bar{D}^0$ oscillations from a single measurement.

DOI: 10.1103/PhysRevLett.110.101802 PACS numbers: 12.15.Ff, 13.25.Ft, 14.40.Lb

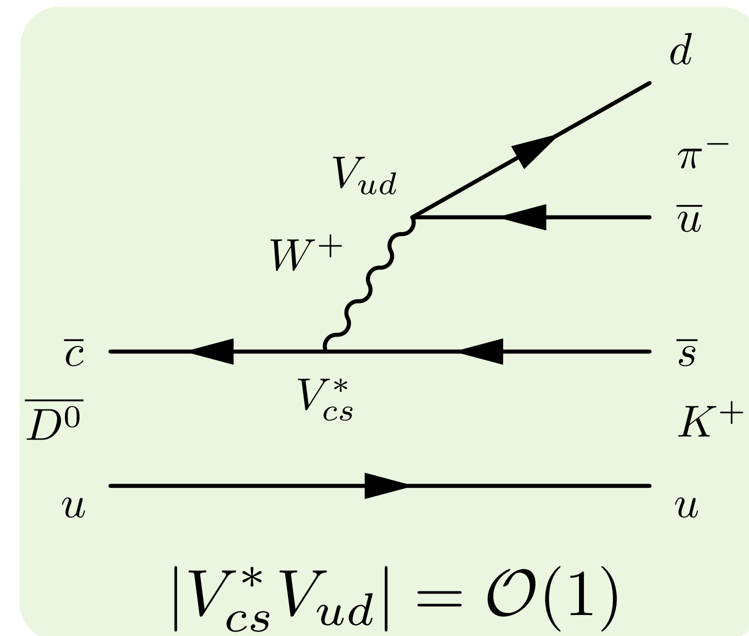

8.4M events

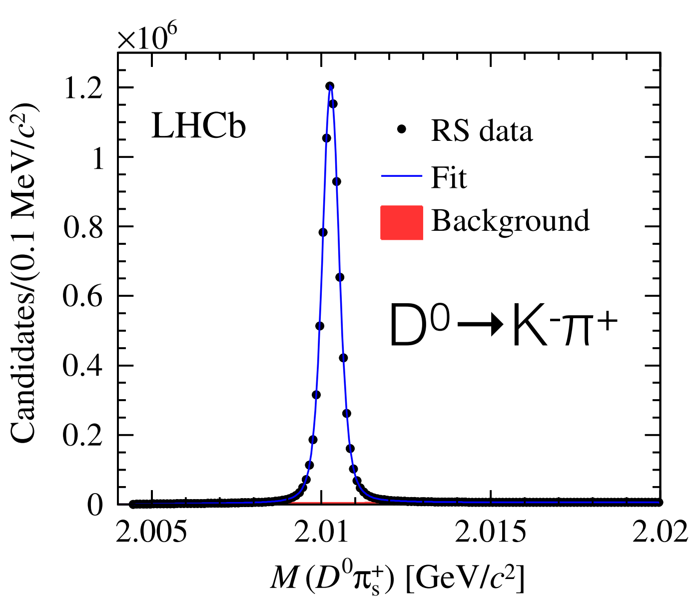

3.6K events

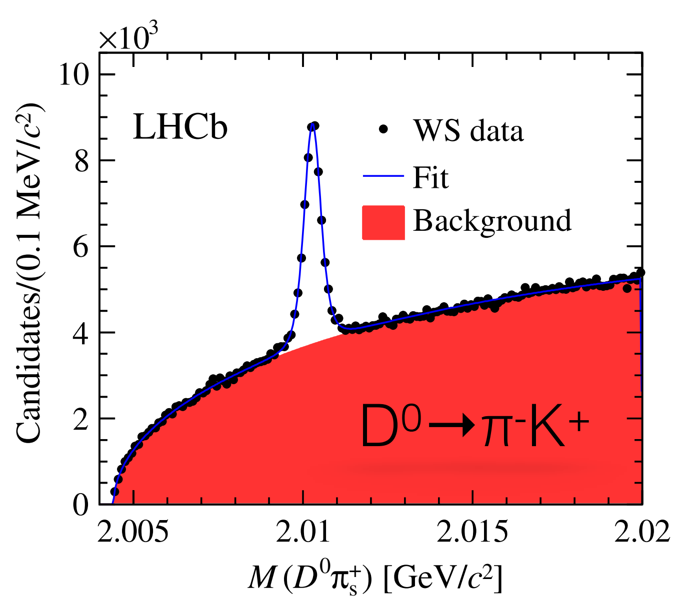

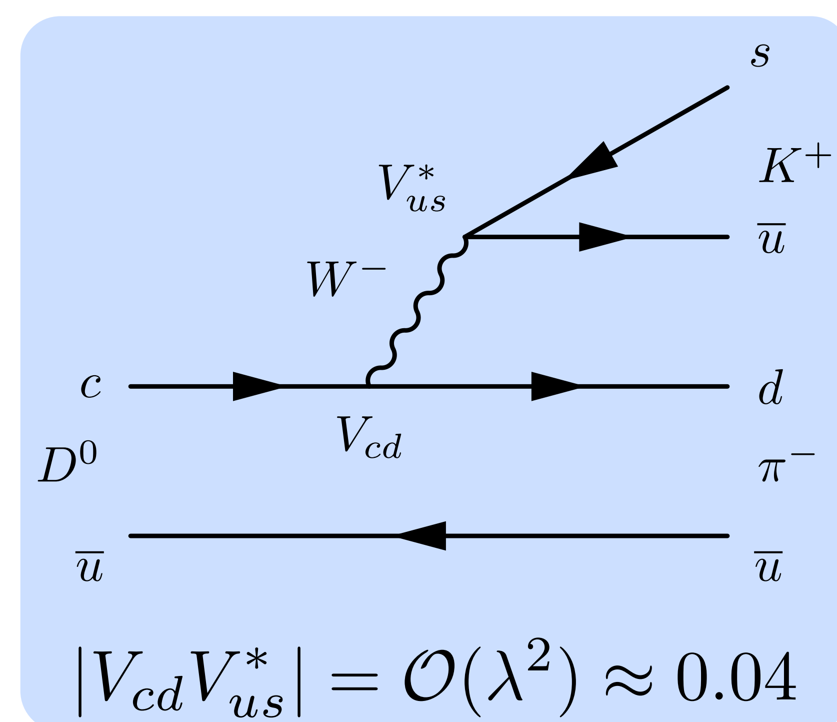

Rapidly Decide — where/when?

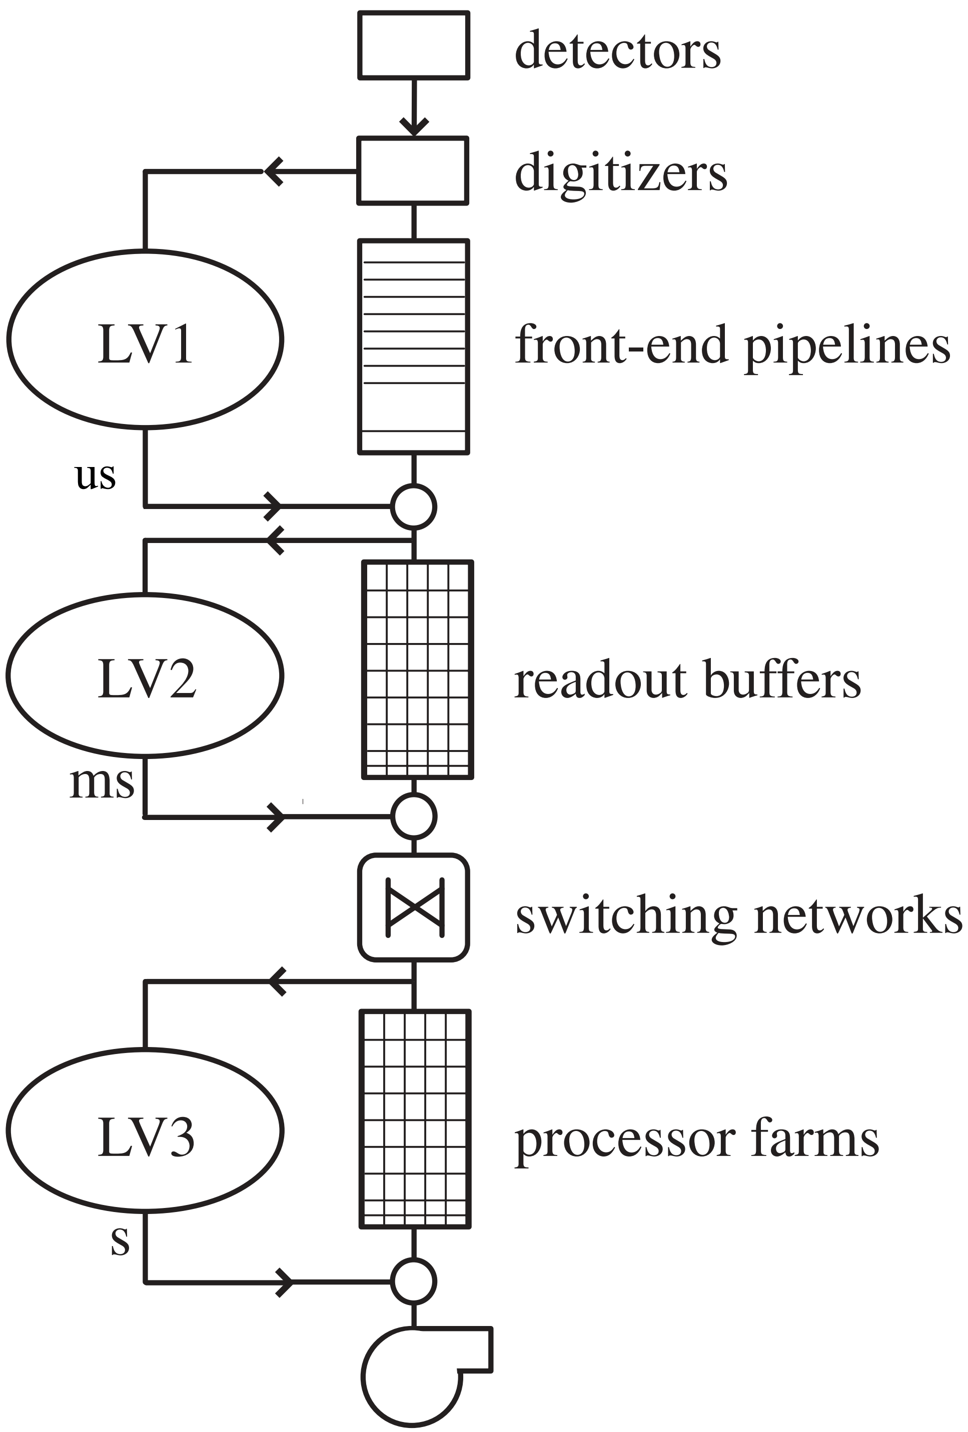

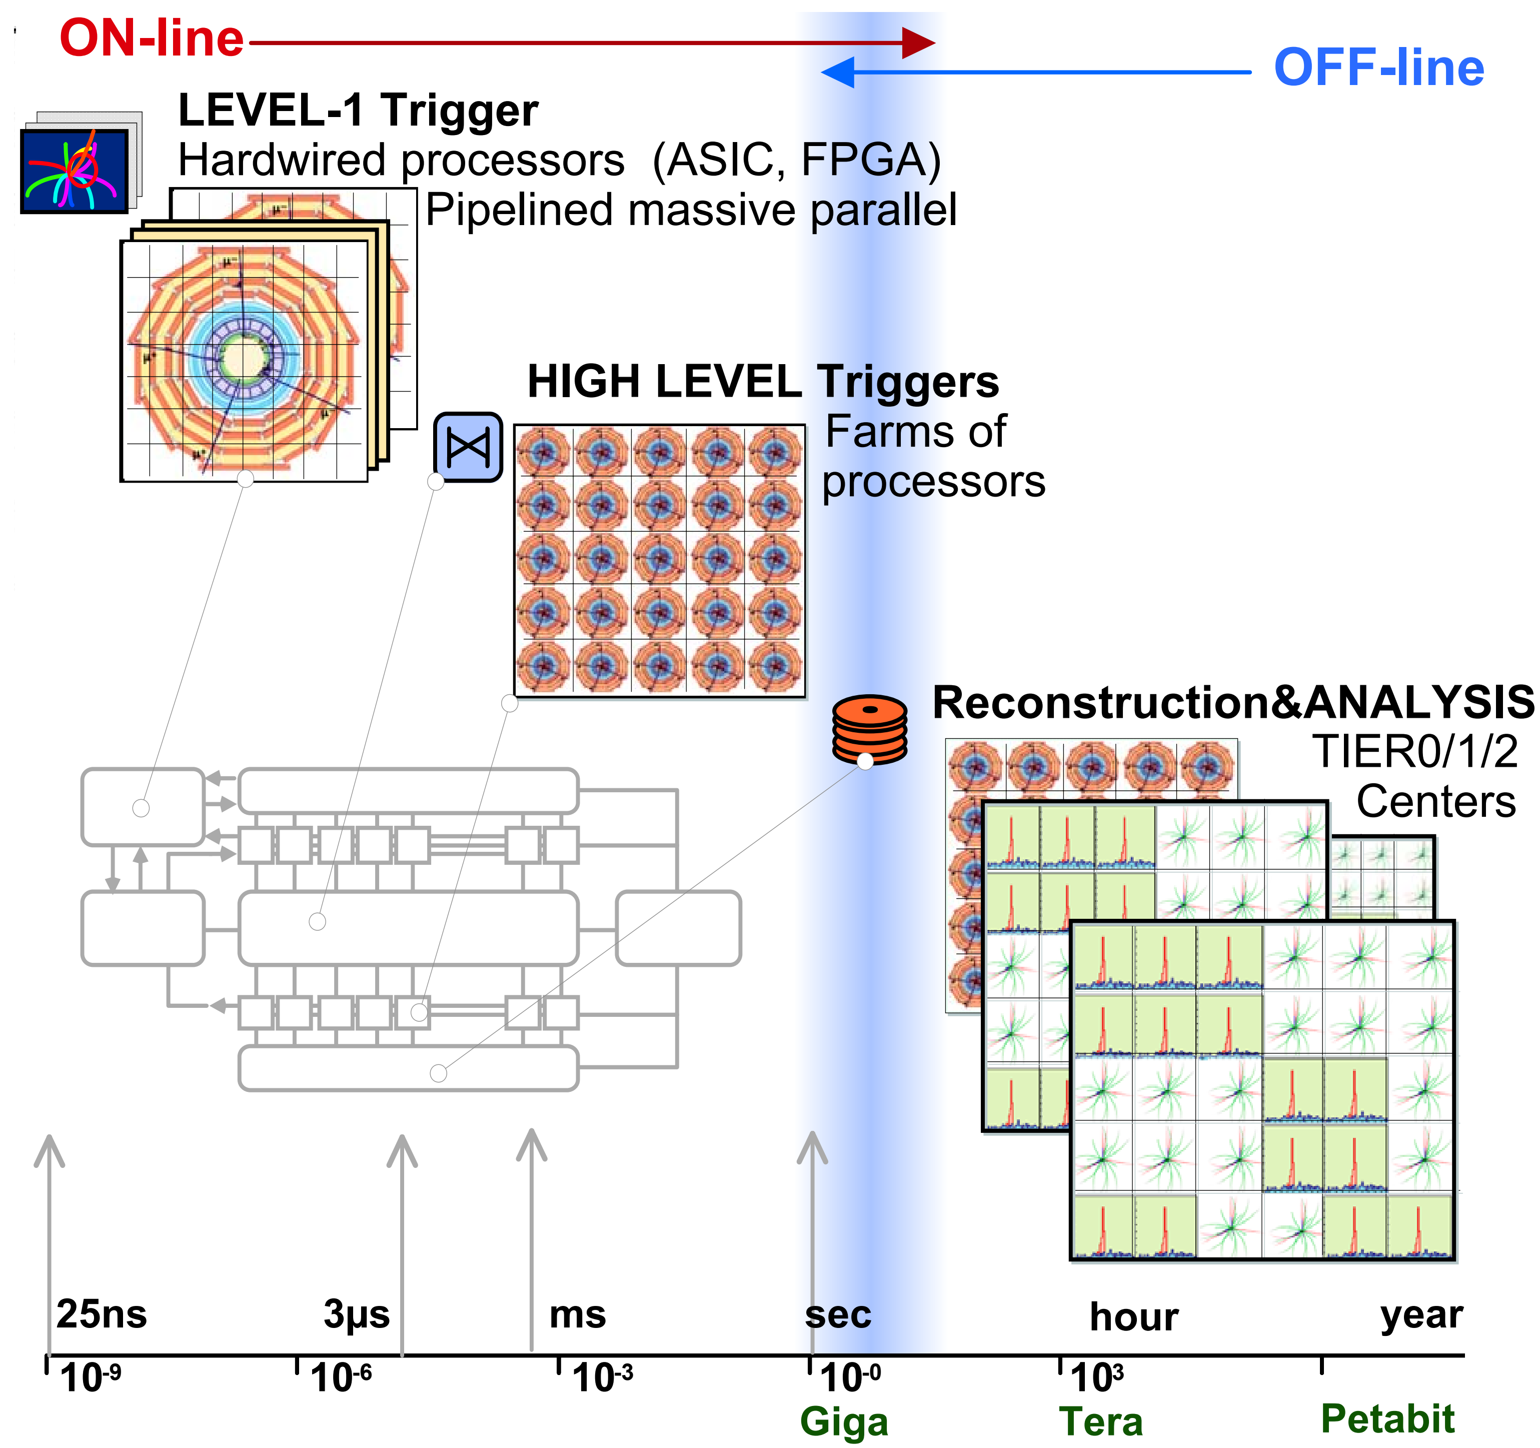

Trivial DAQ

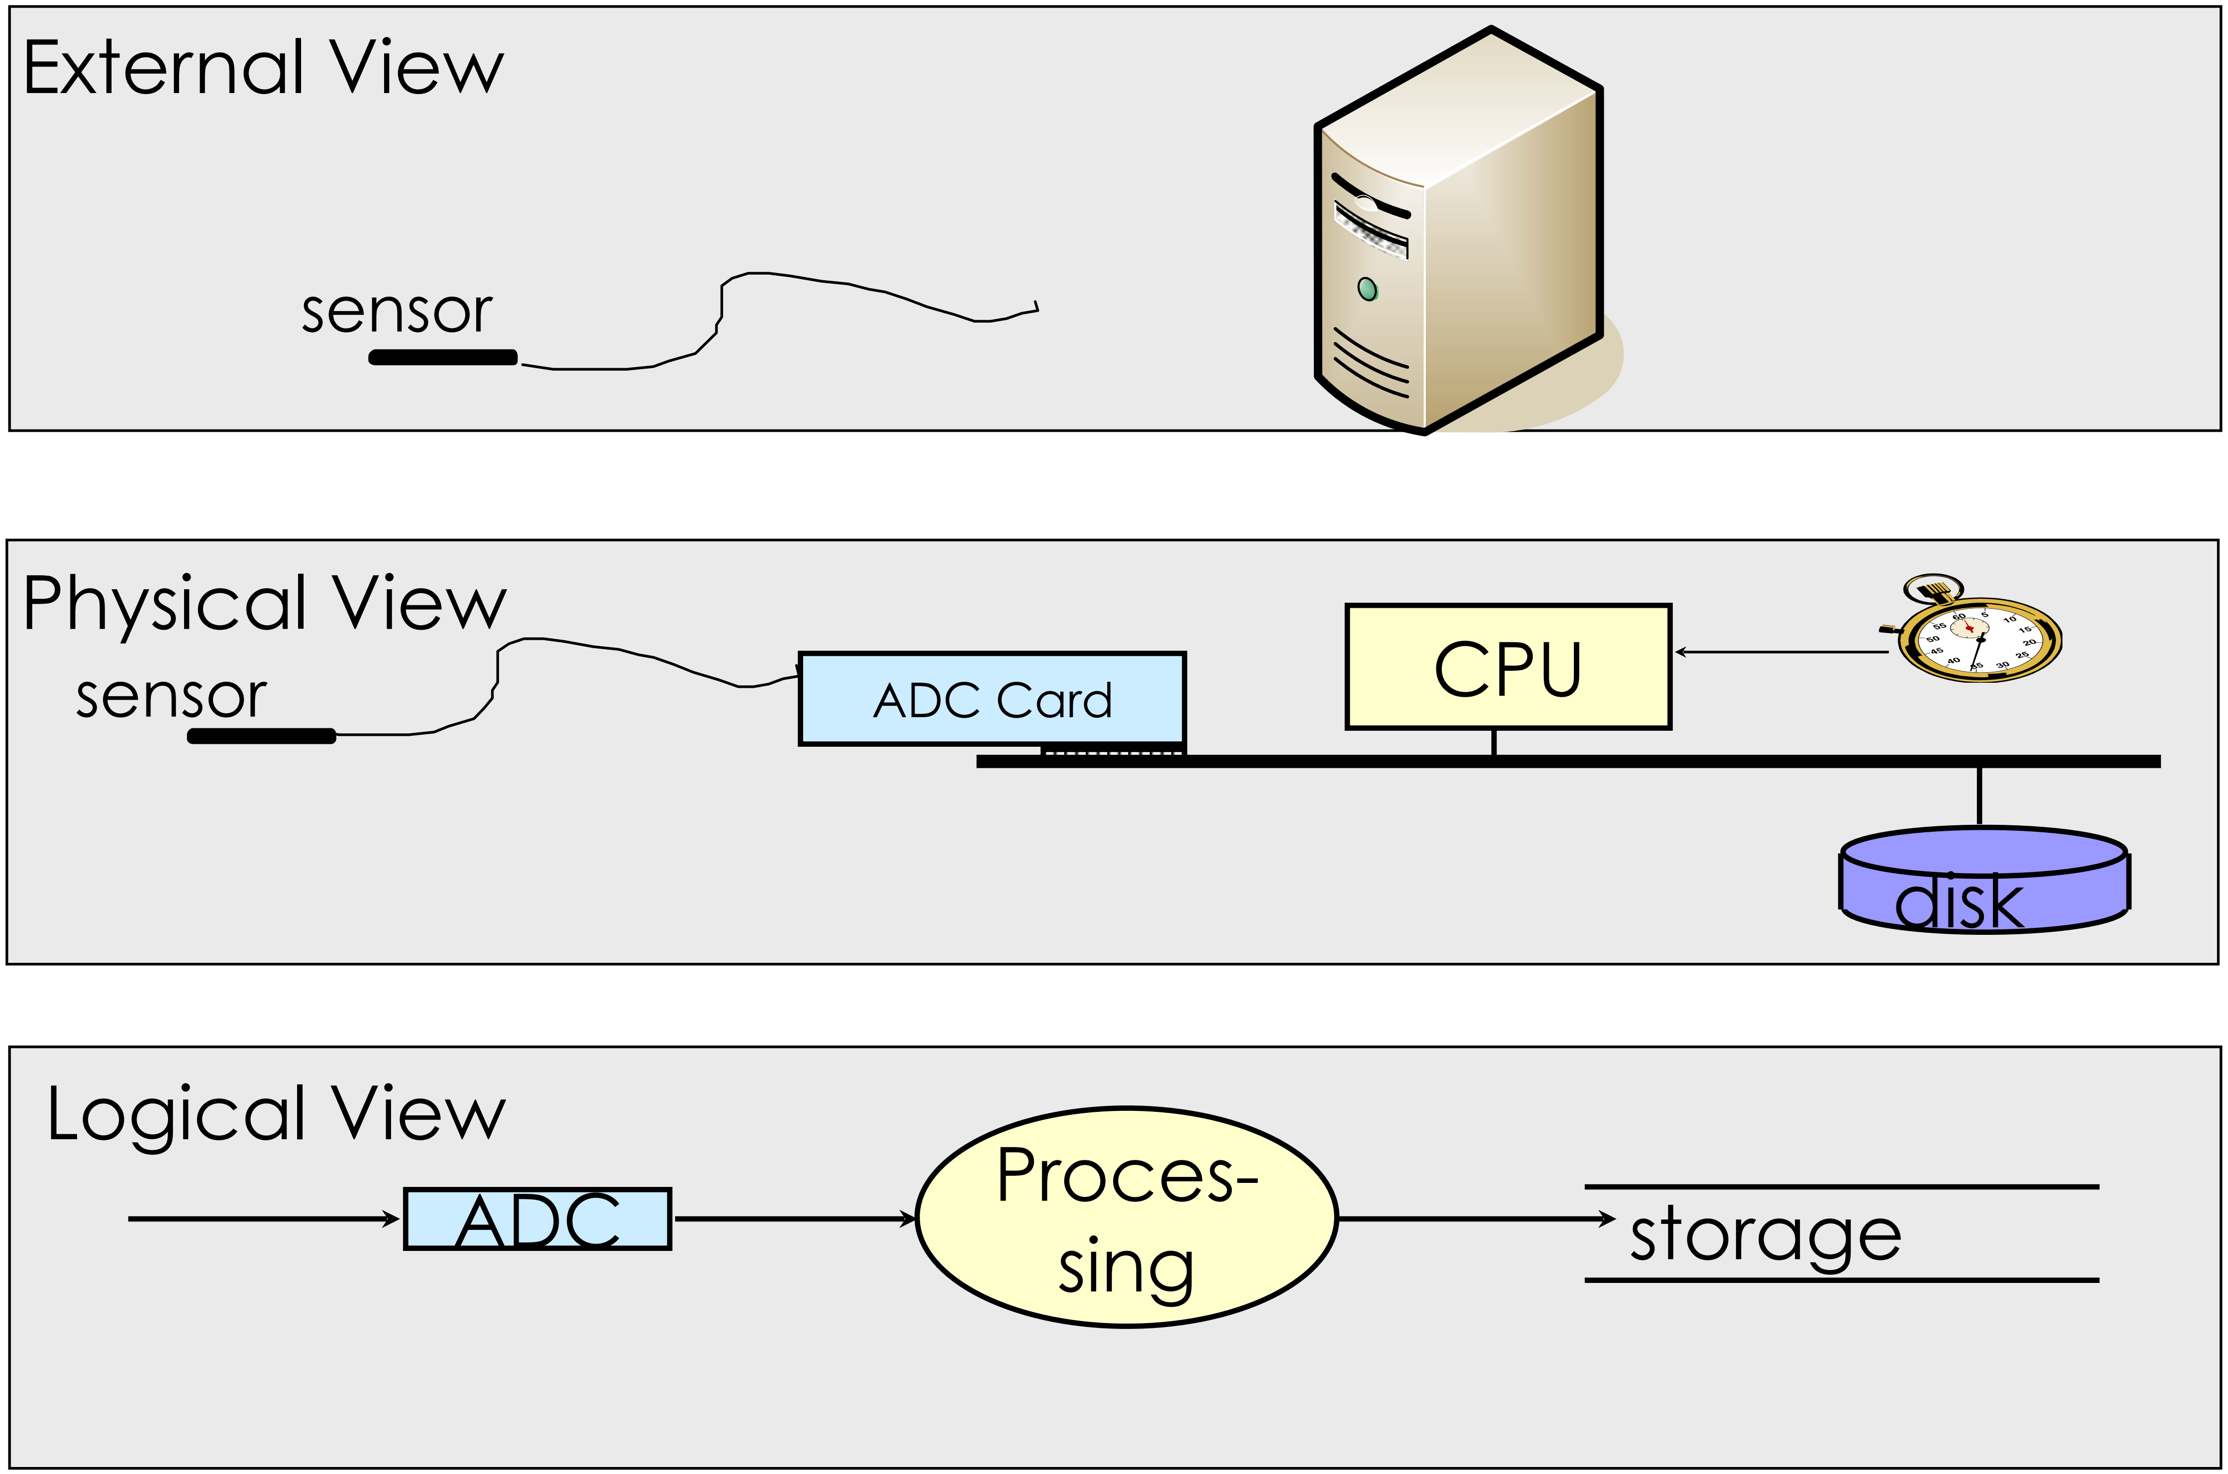

see eg. https://indico.cern.ch/event/349821/session/27/material/0/0.pdf

Trivial DAQ with a trigger

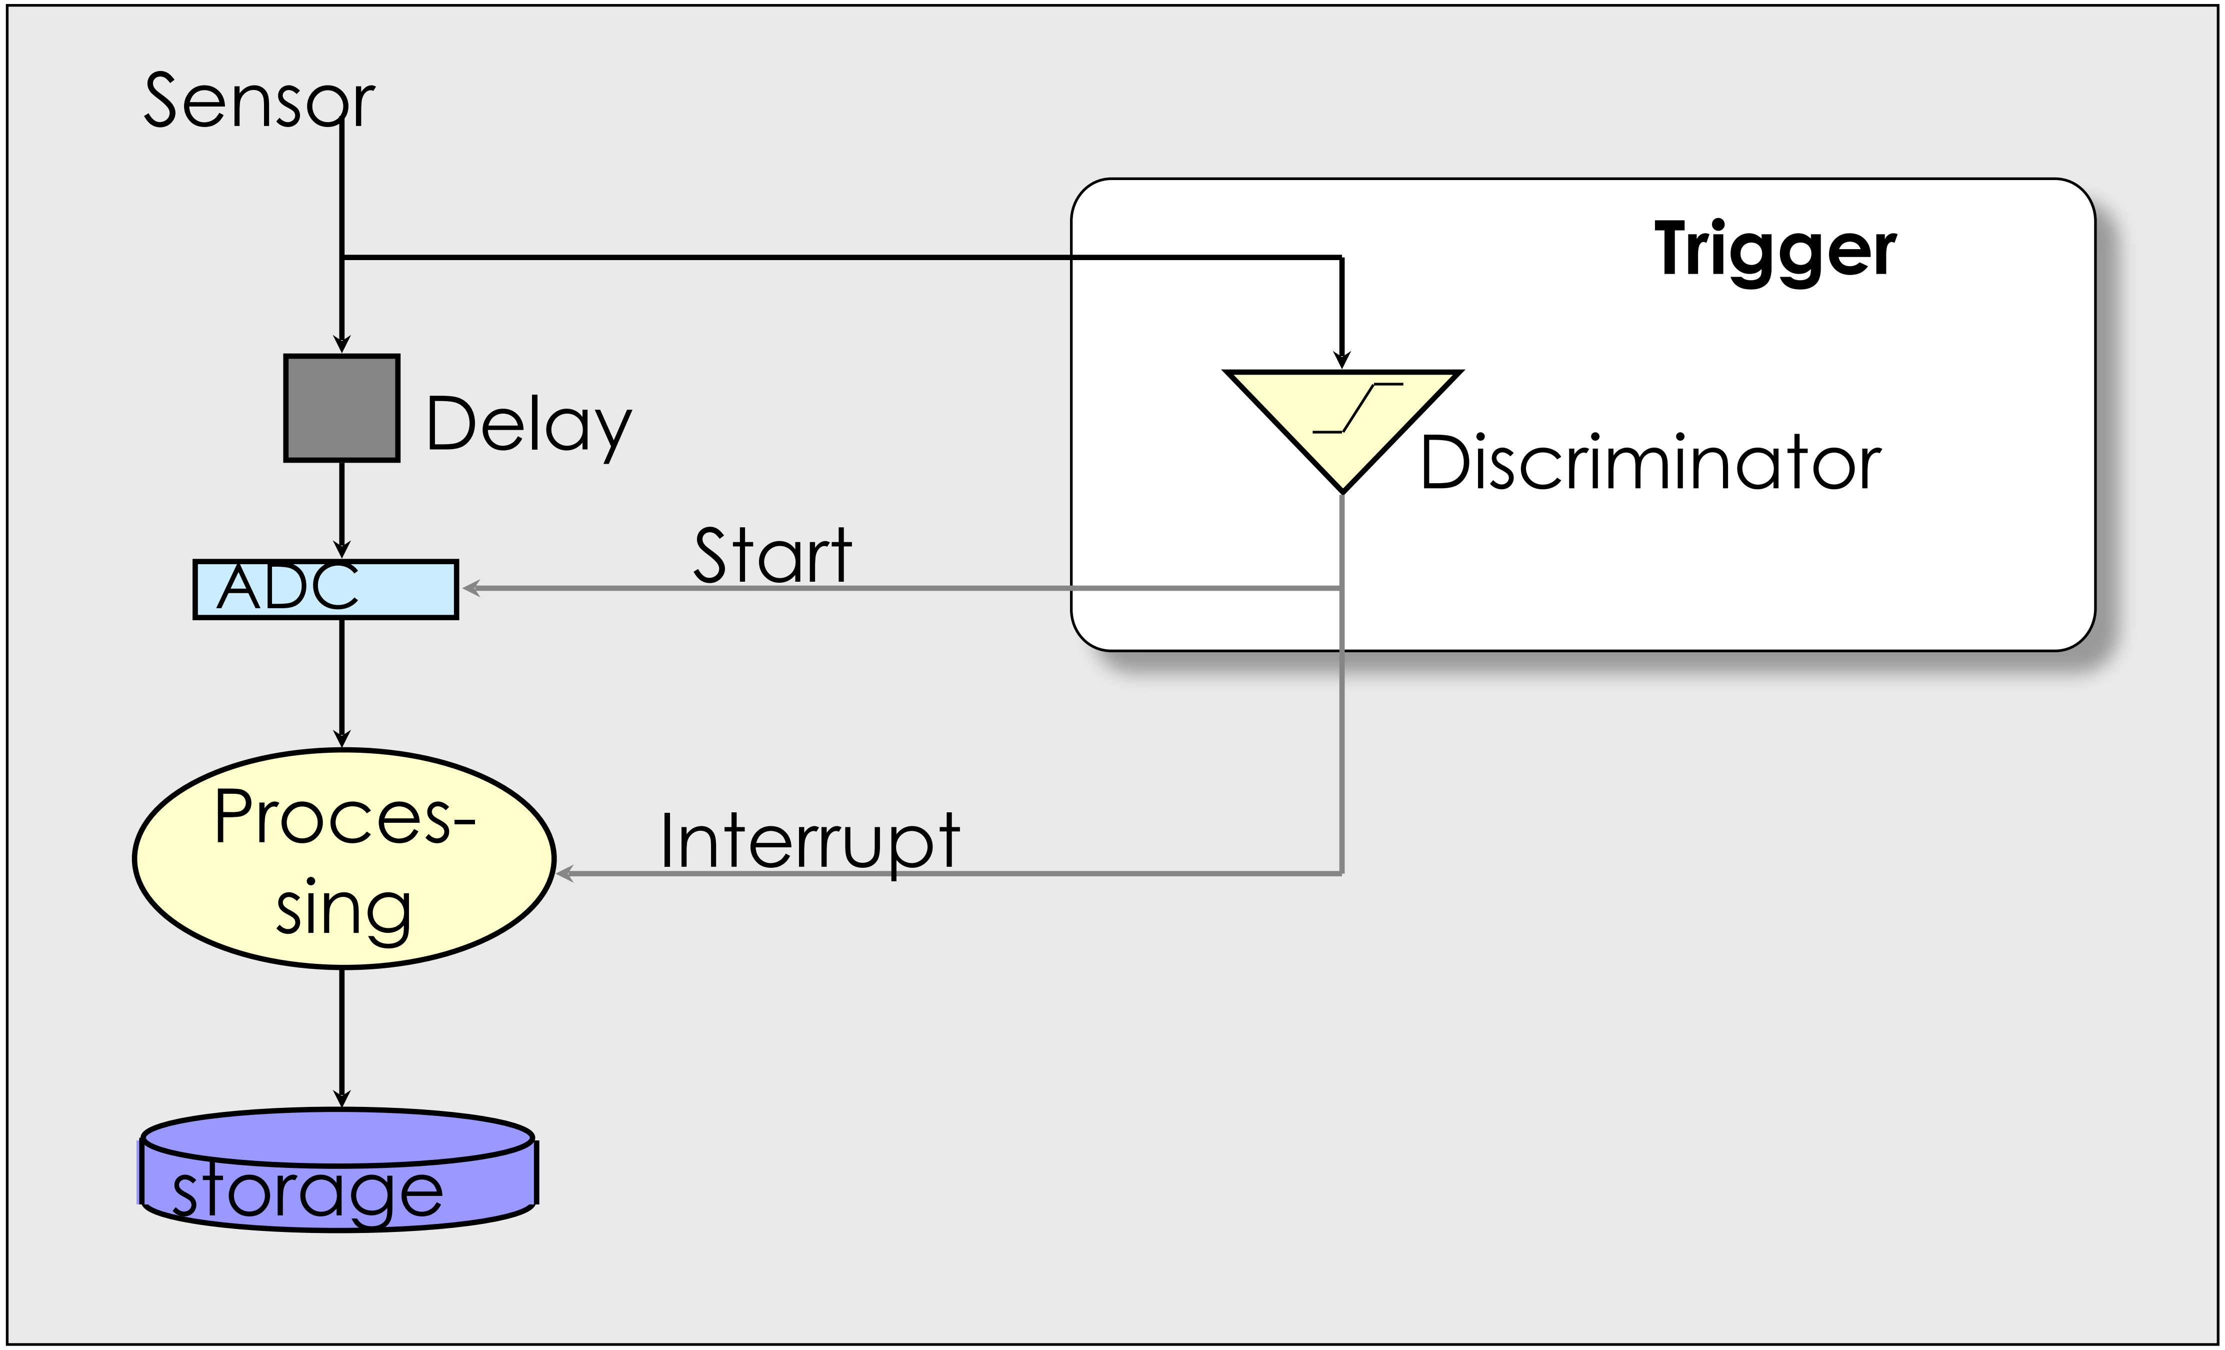

What if a trigger is produced when the *ADC* or *processing* is busy?

see eg. https://indico.cern.ch/event/349821/session/27/material/0/0.pdf

Trivial DAQ with a trigger

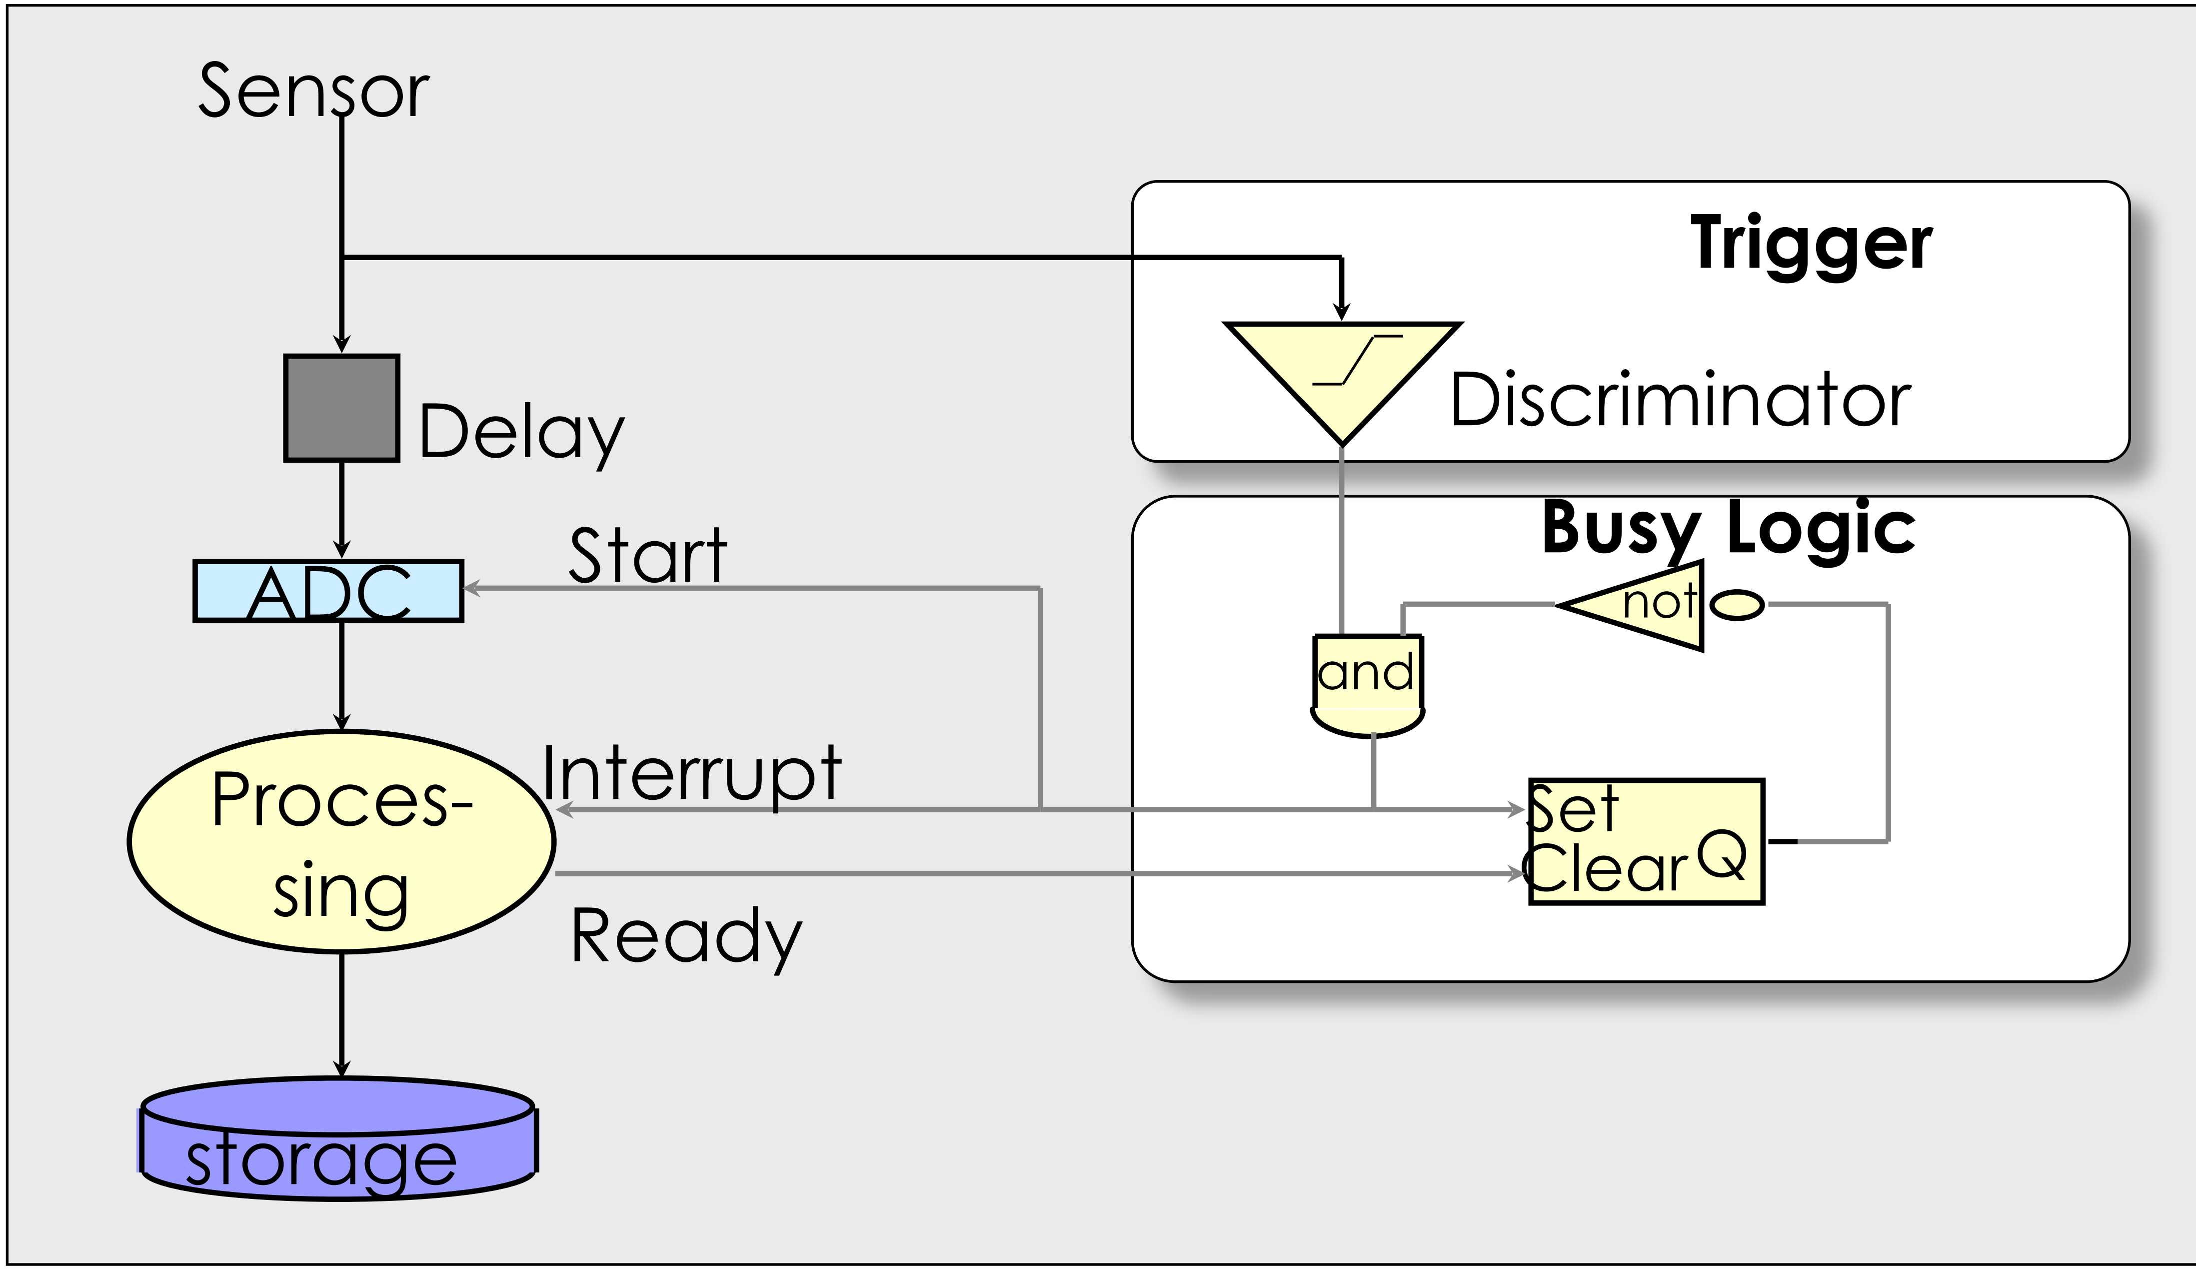

Deadtime (%) is the ratio between the time the DAQ is *busy* and the total time.

see eg. https://indico.cern.ch/event/349821/session/27/material/0/0.pdf

Trivial DAQ with a trigger

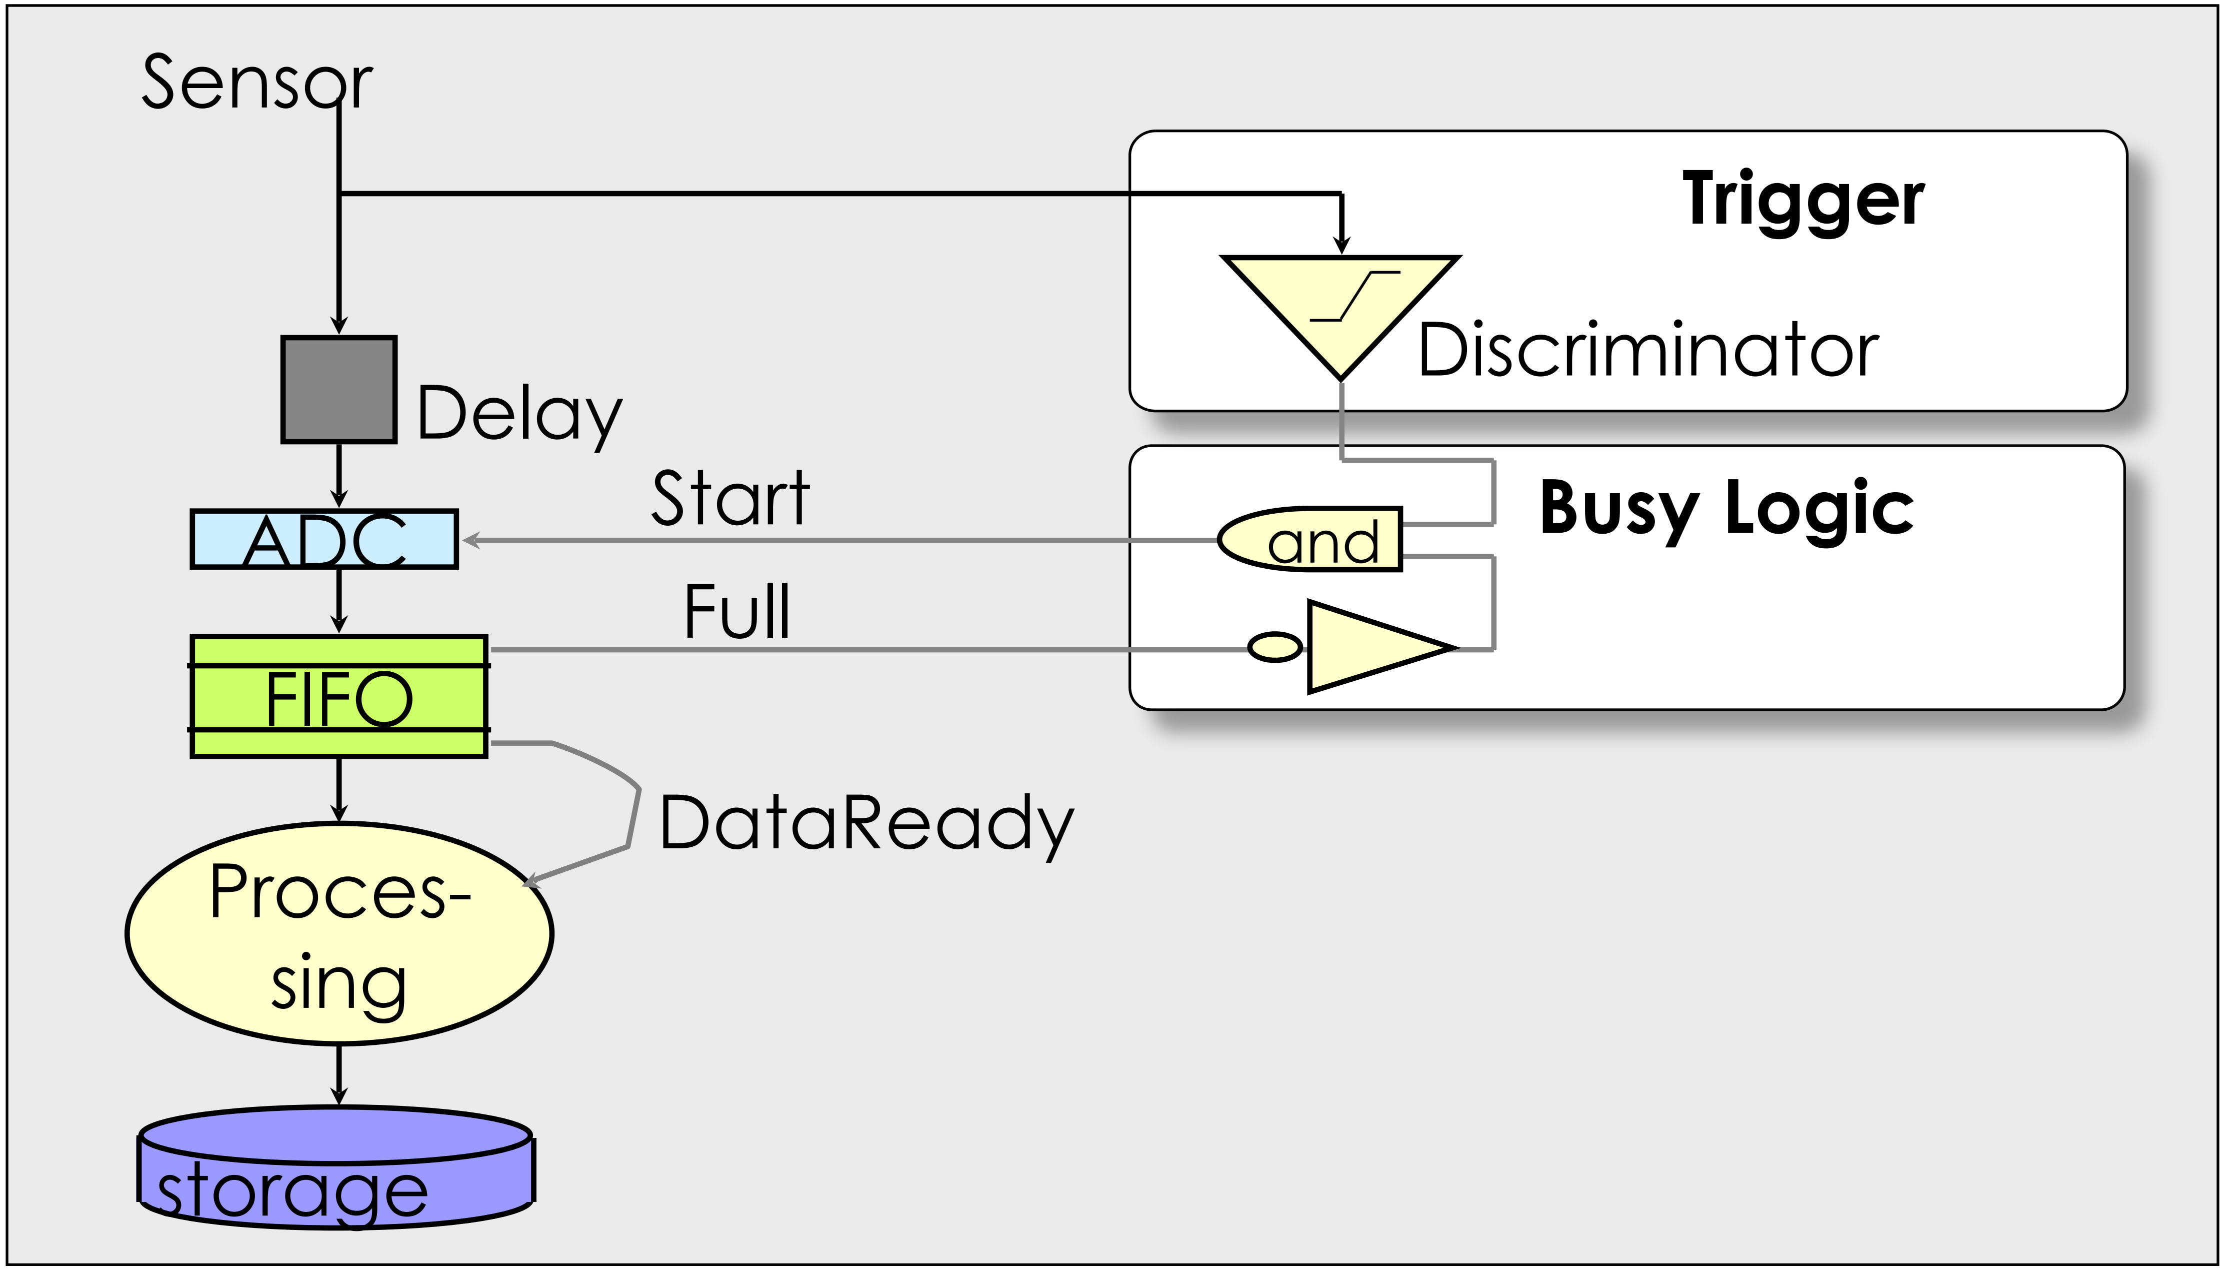

Buffers are introduced to de-randomize data: decouple the data production from the data consumption. **Better performance**.

see eg. https://indico.cern.ch/event/349821/session/27/material/0/0.pdf

Real case is less trivial...

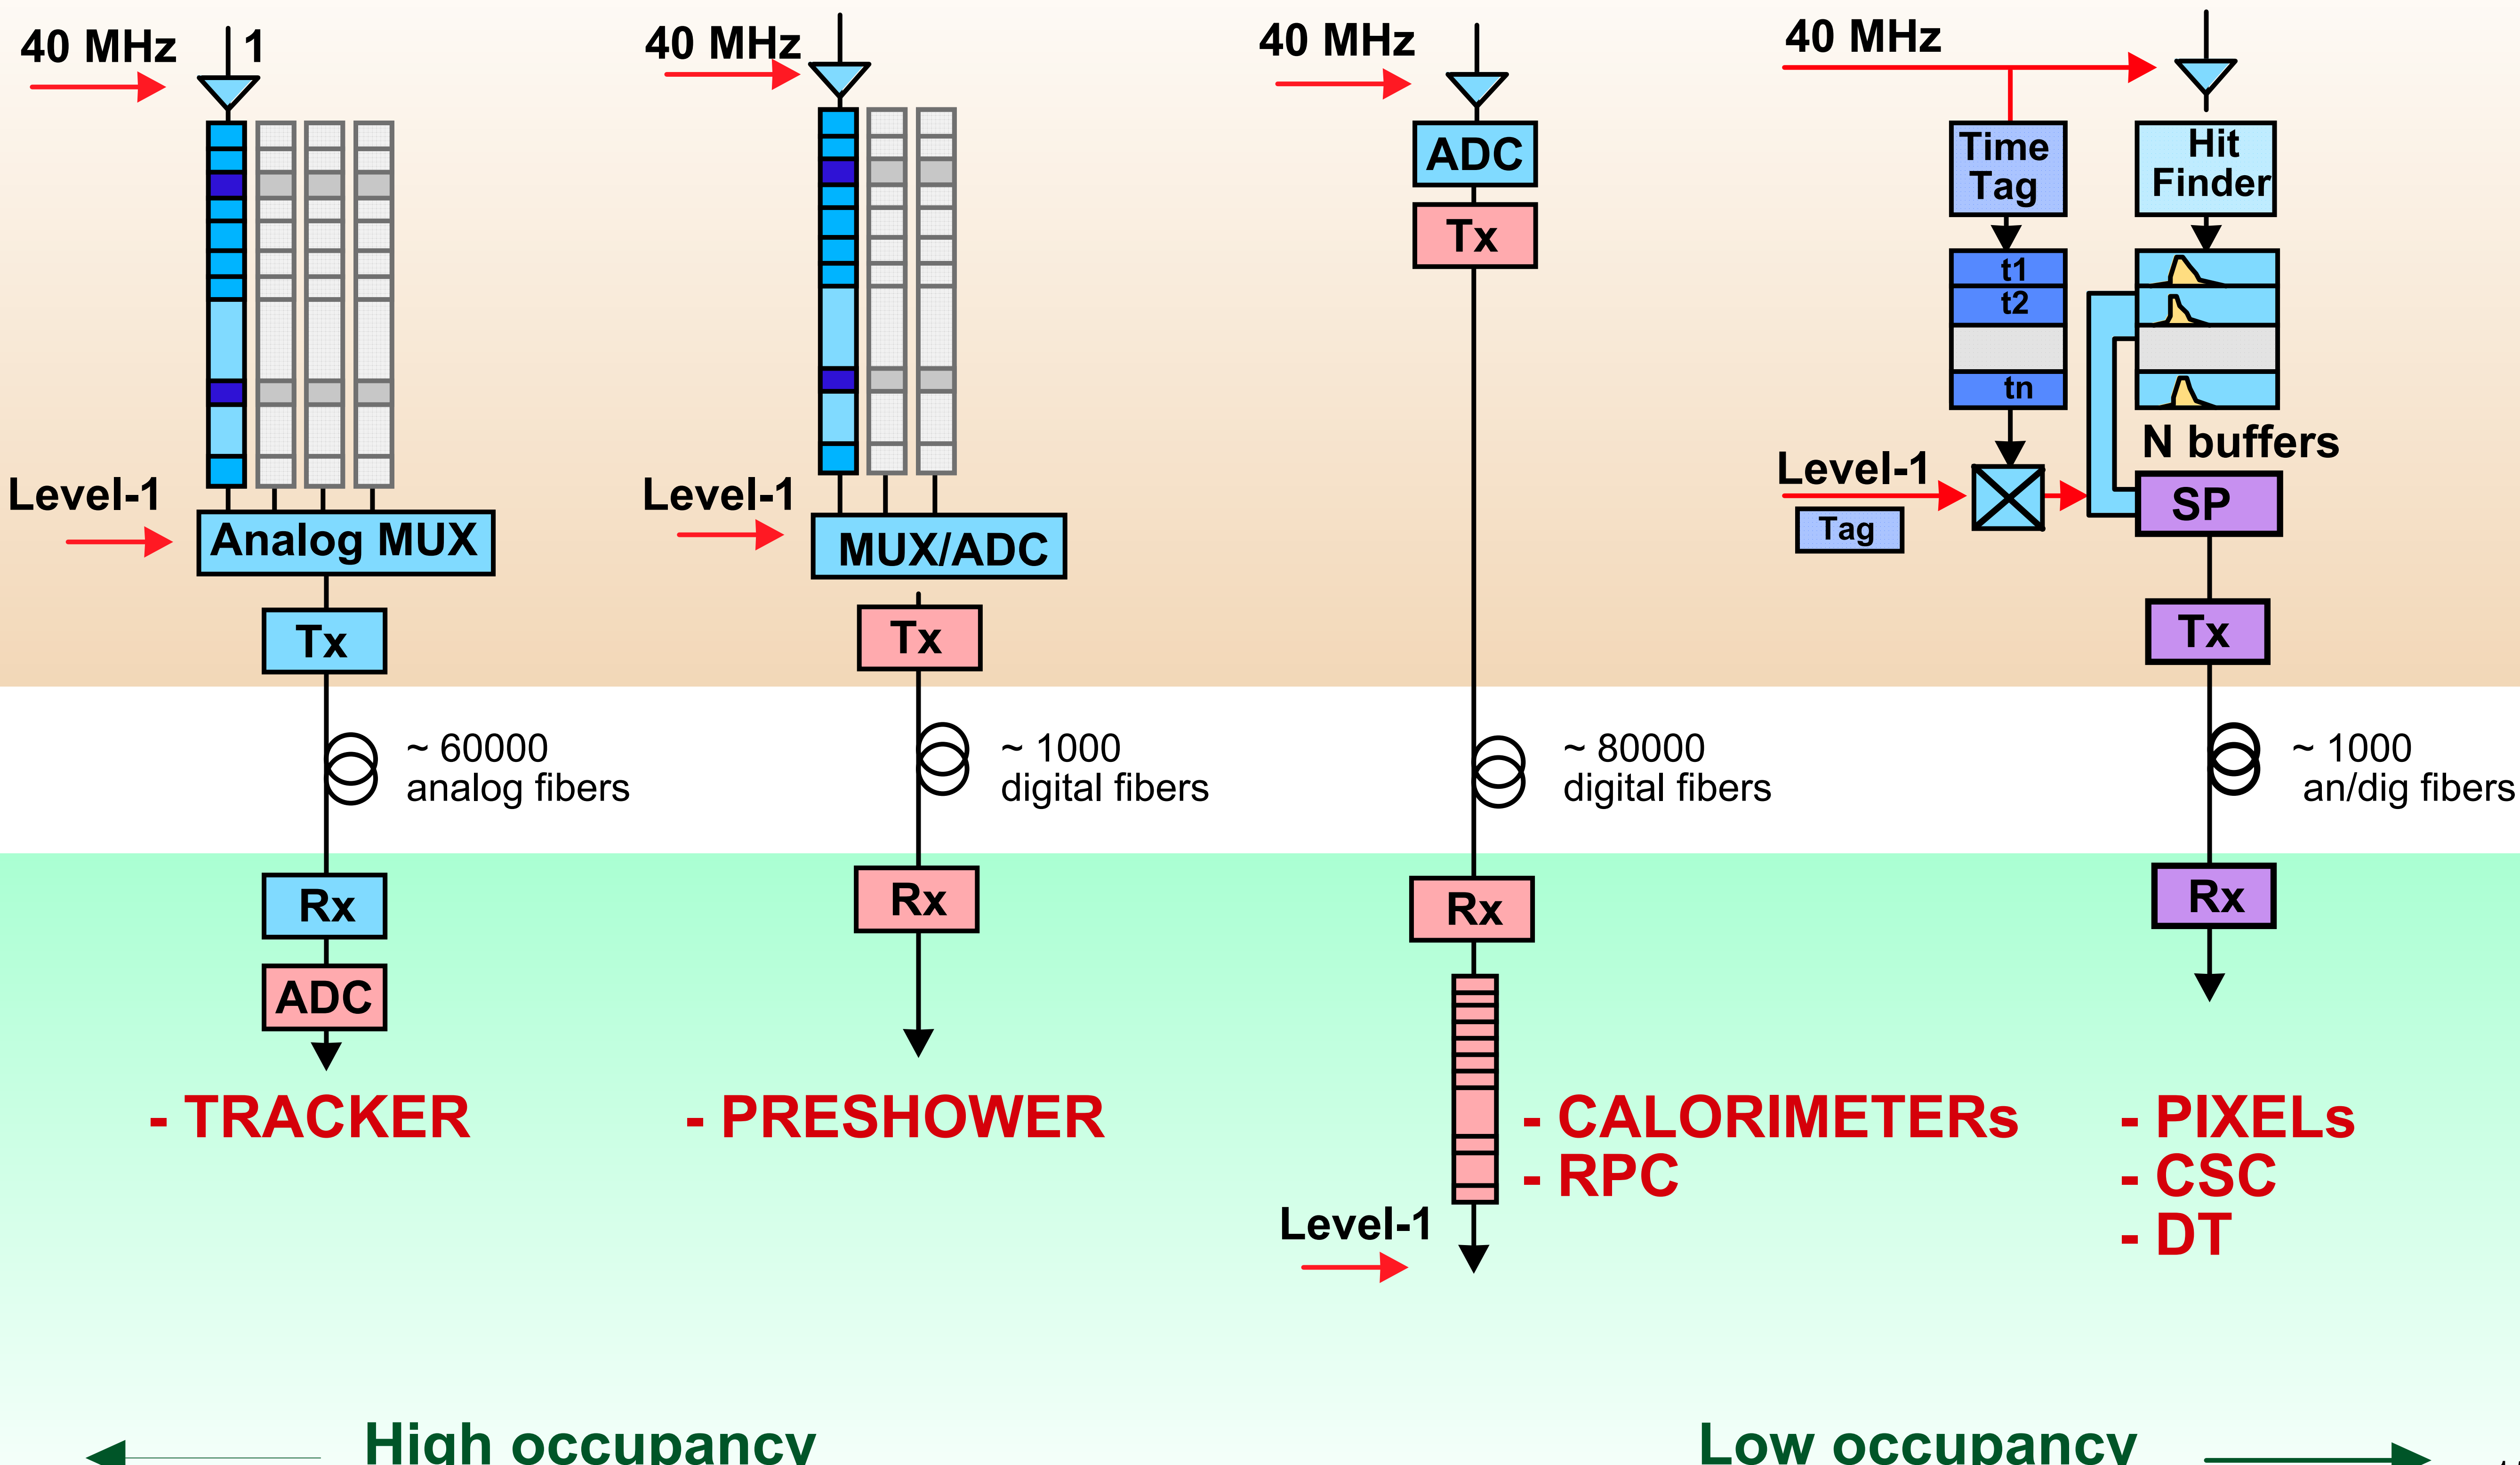

DAQ with trigger & event builder

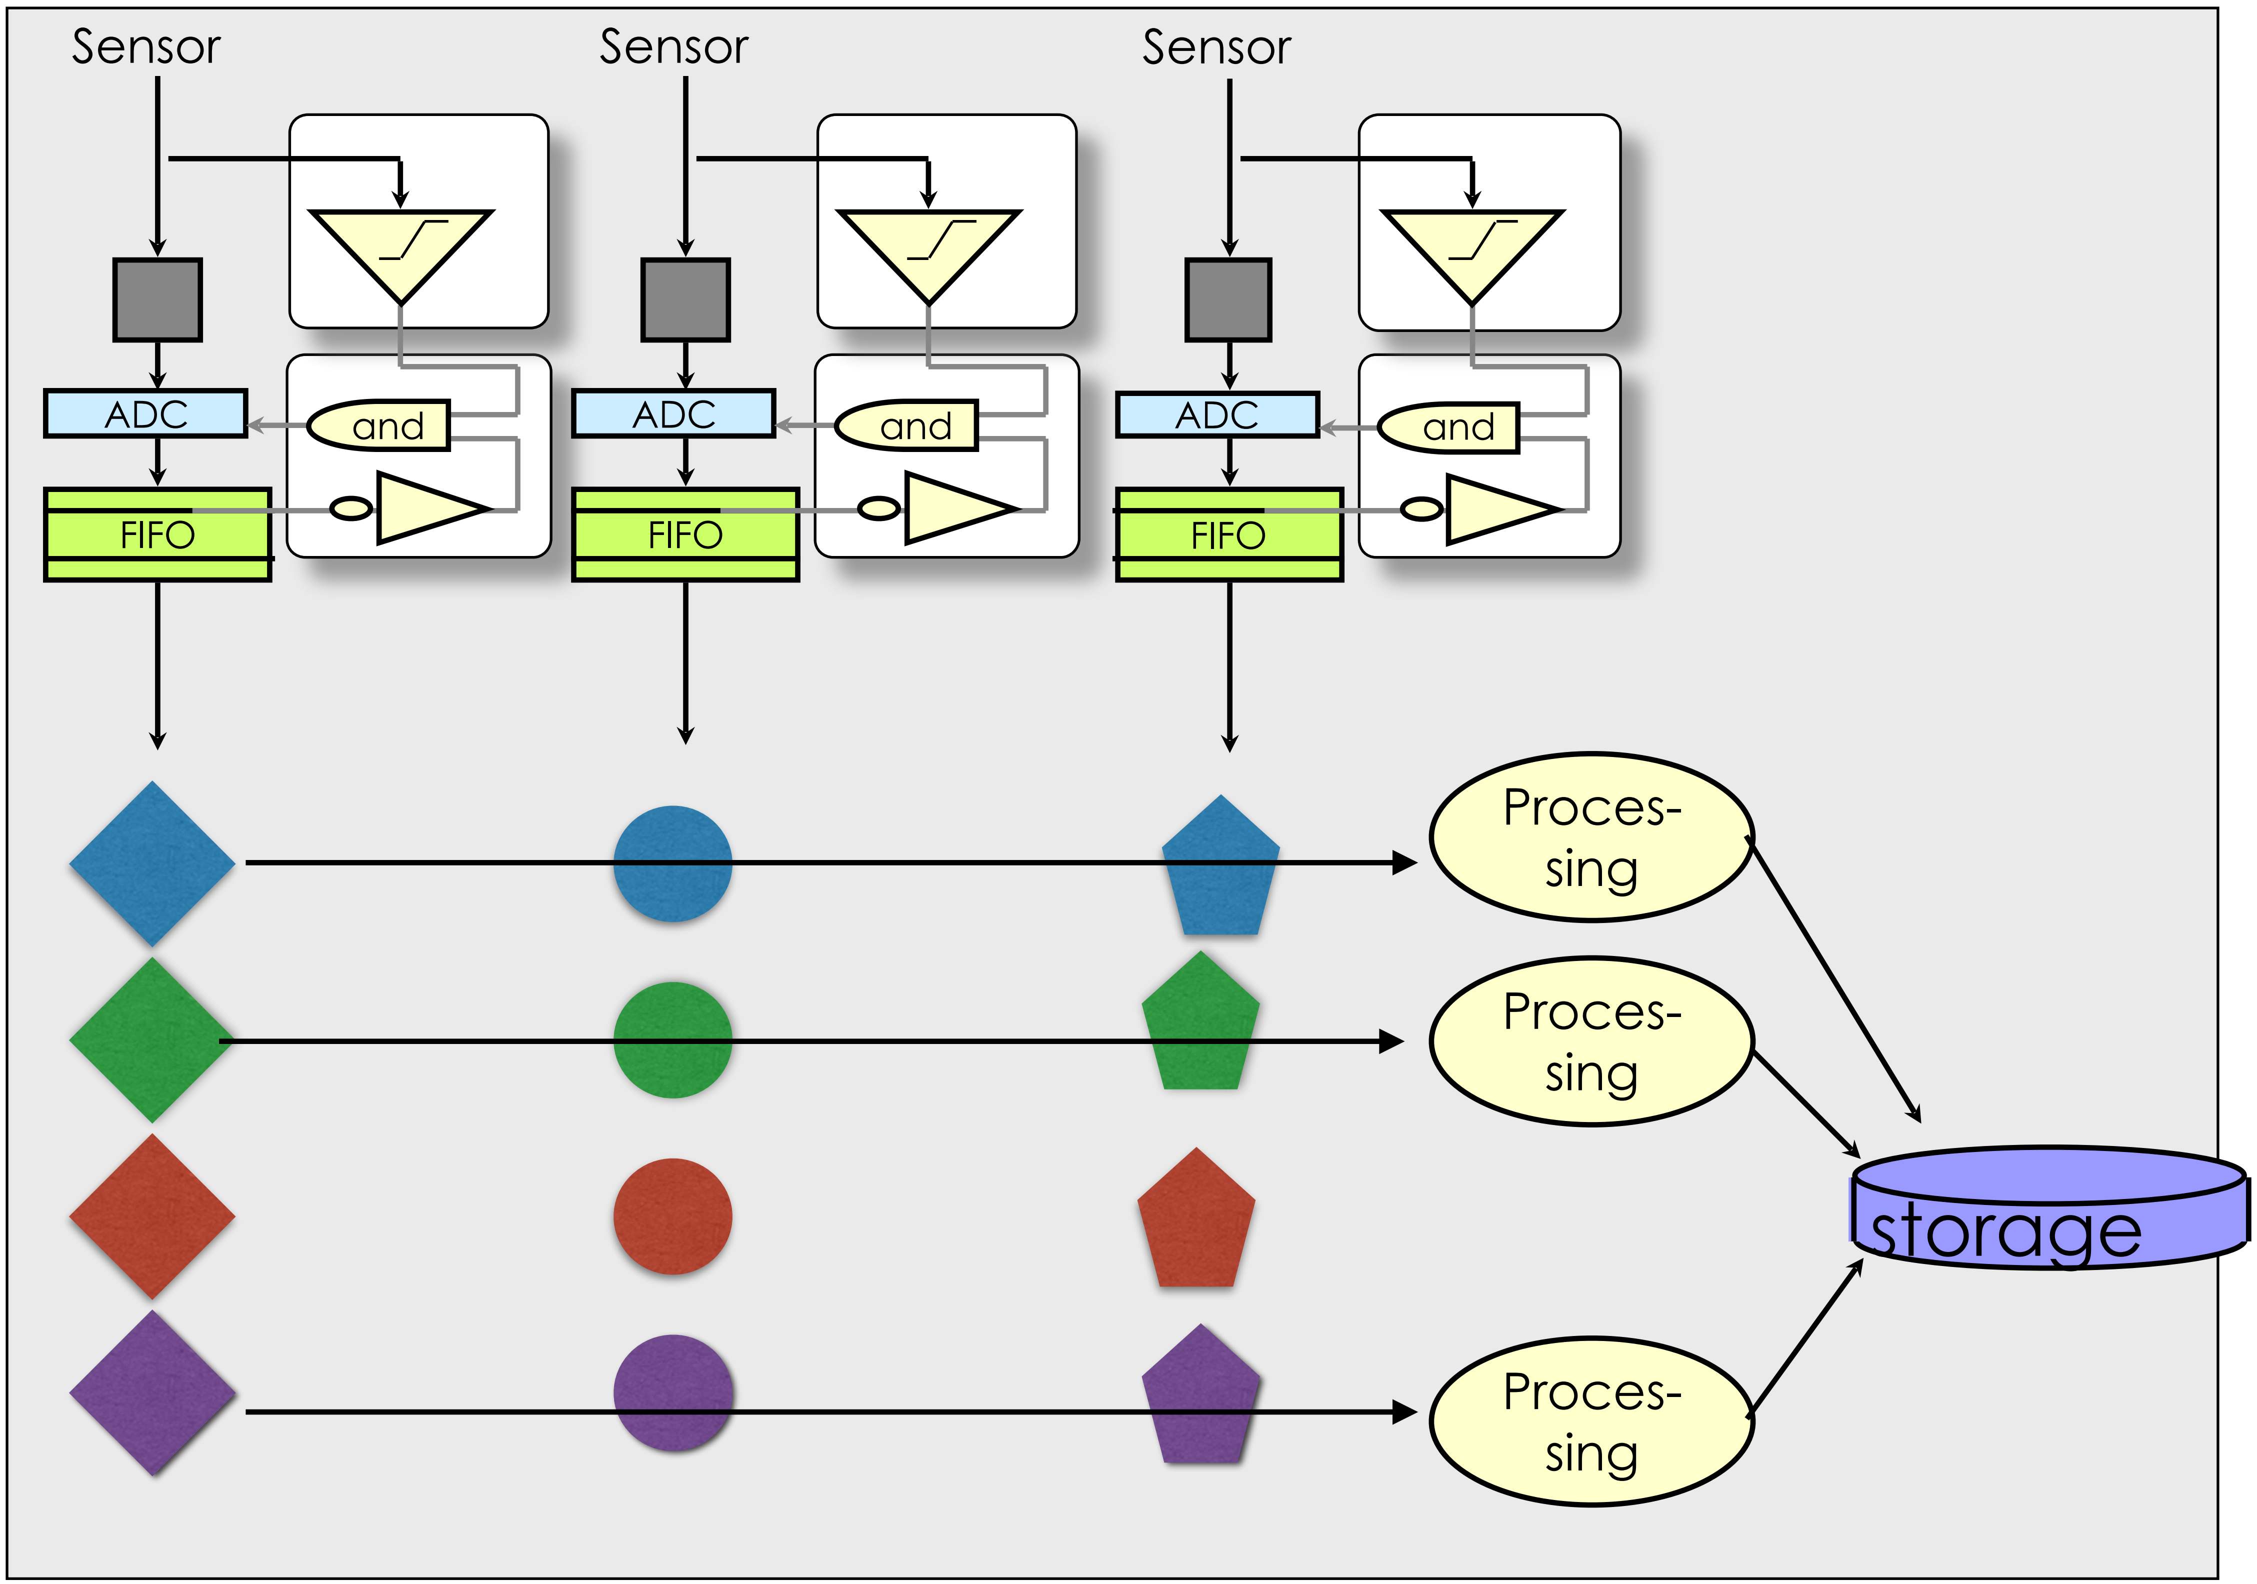

Reduction in stages…

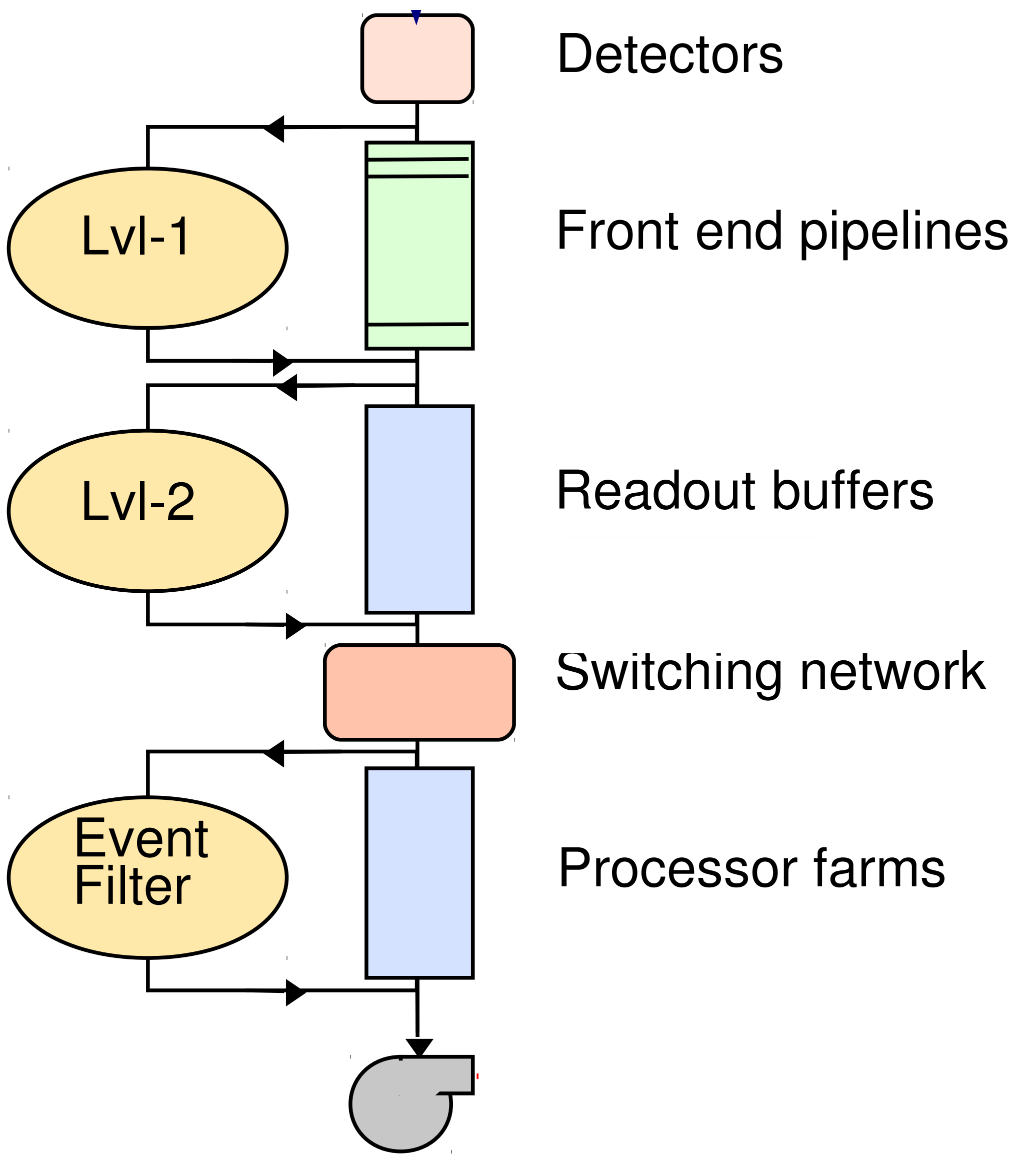

Reduction in stages…

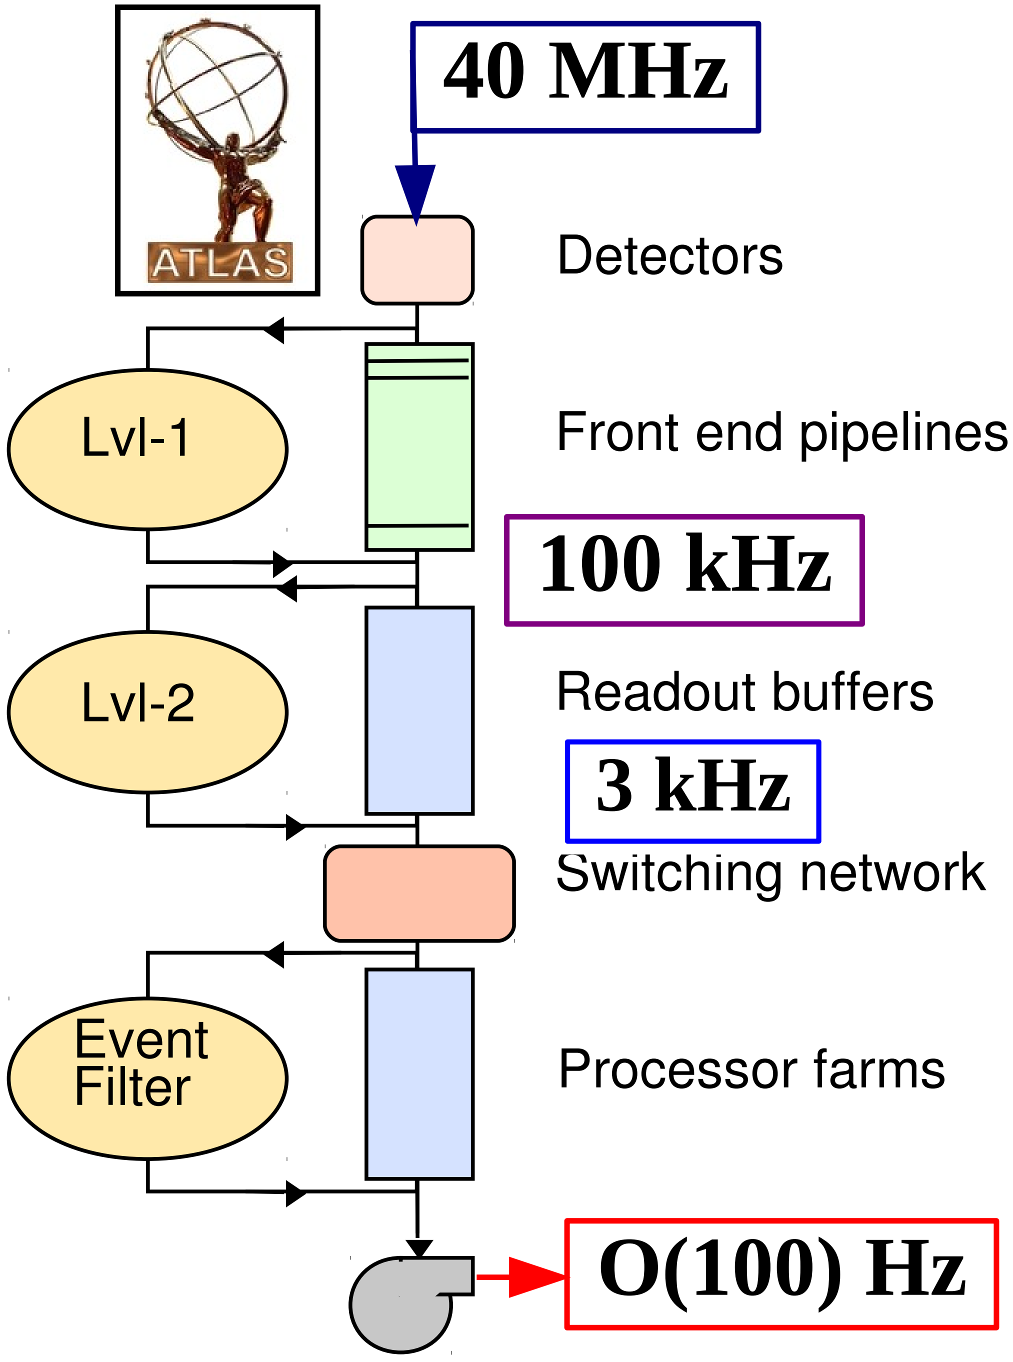

What about CMS?

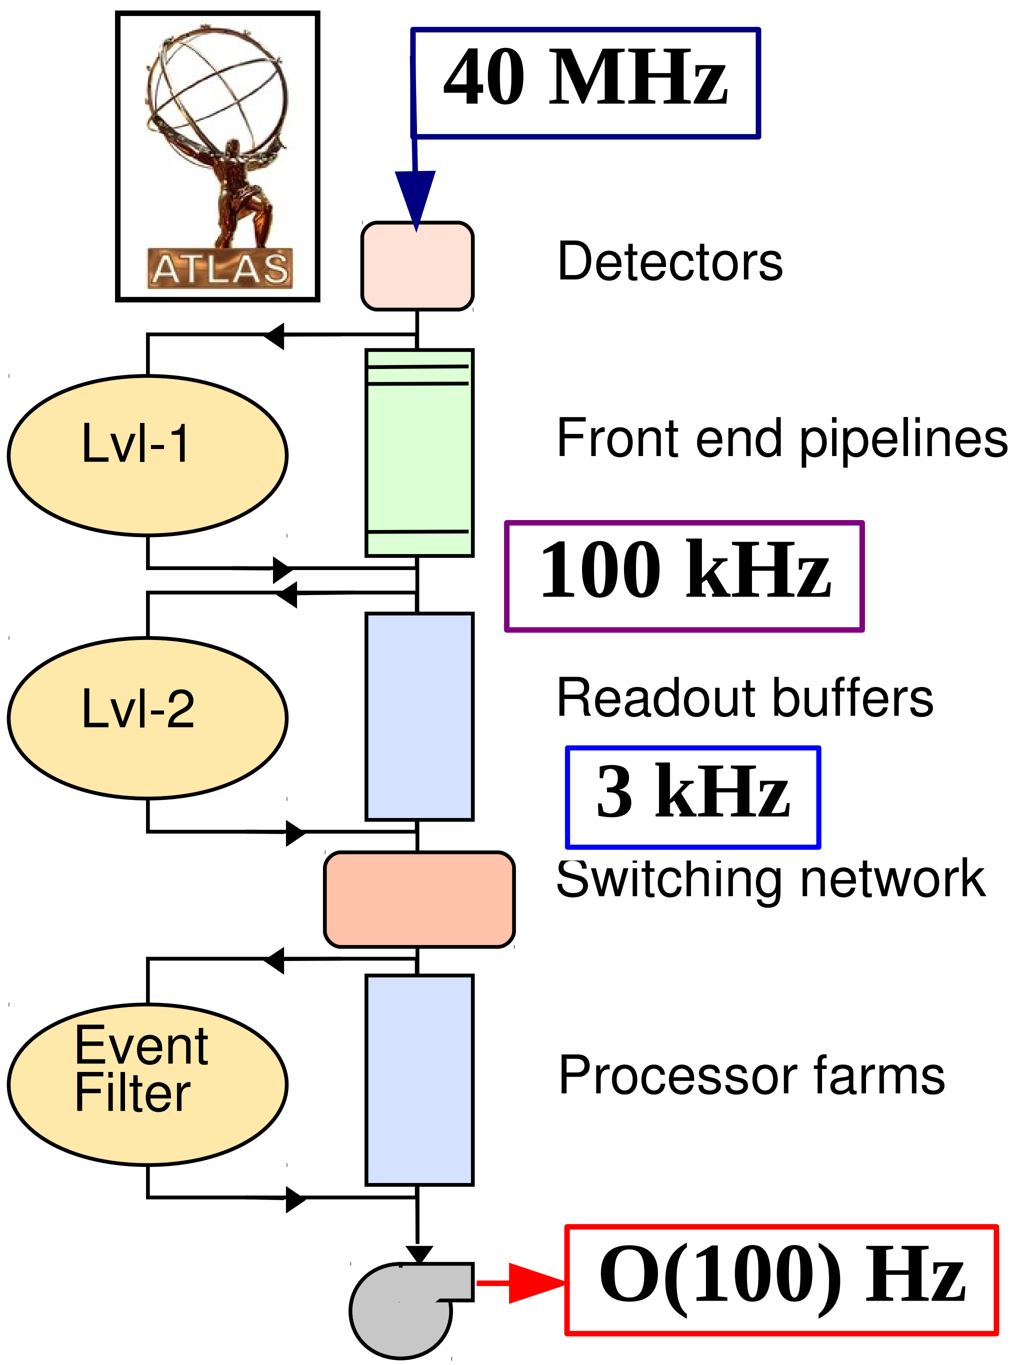

At CMS, TriDAS (Trigger and Data Acquisition System) Project Manager Sergio Cittolin, TriDAS Institution Board Chair Paris Sphicas and Smith decided not to have a second level trigger. They would take the output of L1 straight to the computer farm for software processing. The main reason for doing this was that the L2 hardware was too restrictive. It was not fully programmable, and was only used at the time because there was no telecom switch that could convey the full L1 output of 100 kHz of 1 MB events to the farm.

However, Cittolin, studying technology trends and extrapolating world-wide computing network infrastructure, was convinced that a switch with the required bandwidth would be available and affordable by the year 2000. So, when the technical proposal was written in ’94-’95, a plan to go from L1 to the computer farm was laid out.

- from http://cms.web.cern.ch/news/triggering-and-data-acquisition

What about CMS?

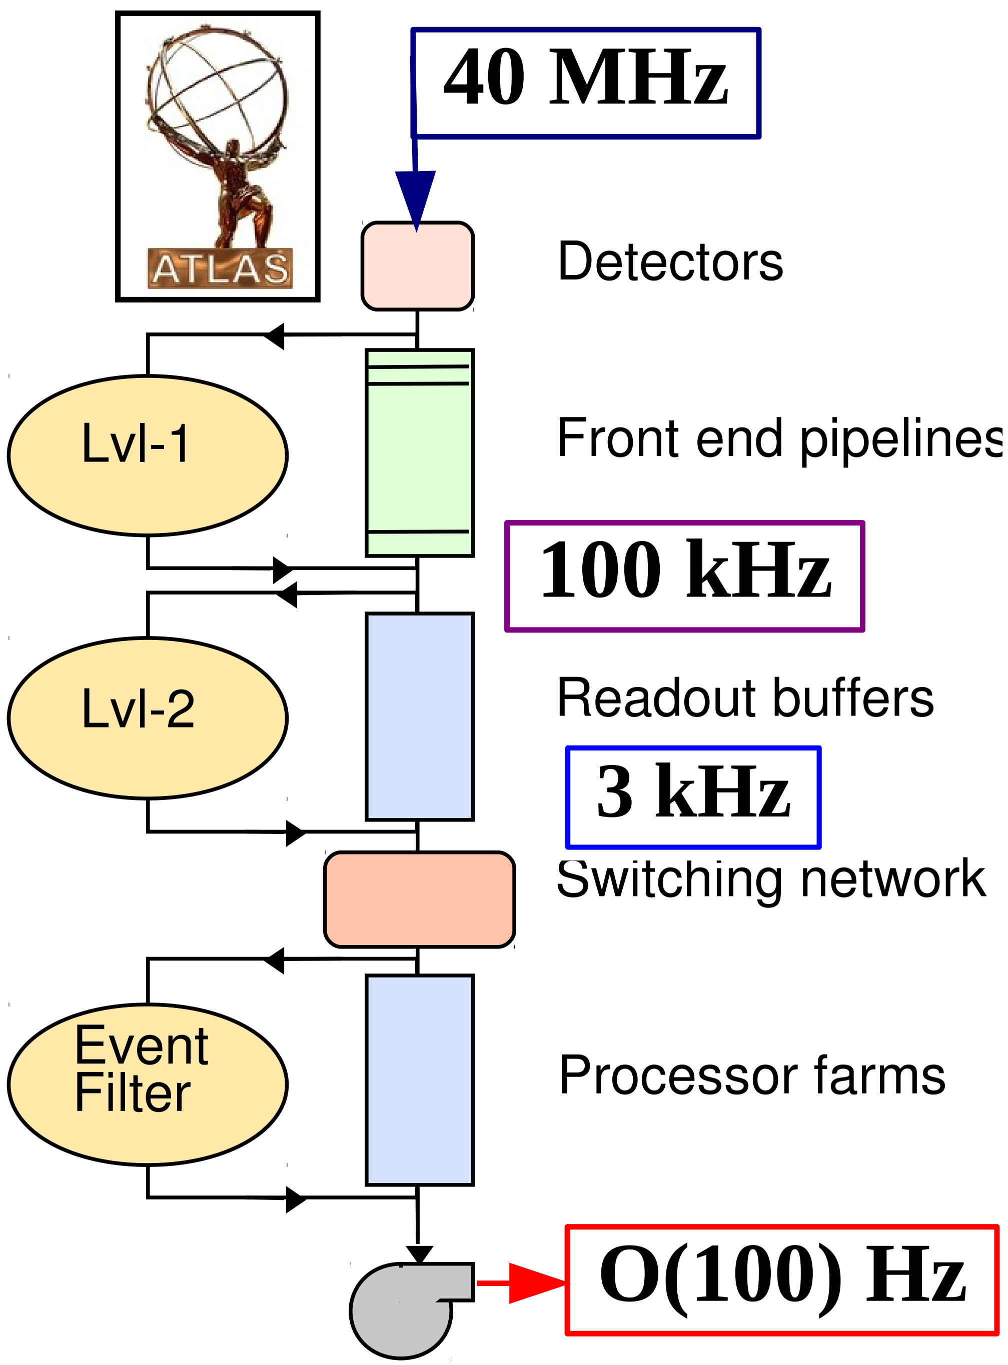

When the technical proposal was presented to the LHC Experiments Committee (LHCC) in 1994, a different approach was adopted to deal with the bandwidth problem, in order to pass the reviews: The proposal said that 10% of the data would go through the switch for processing and would be used as a basis to knock down the rate by a factor of 10 — sufficient for the switches of 1994 to handle! This 10% — including main calorimeter information, summary of tracking information and so on — would then be used to decide whether to keep the event or not. The group hoped, though, that they would not have to face the problem at all.

Wesley Smith adds: “I think it’s been proven that we really built it almost optimally. It was a long extrapolation but I think in retrospect it’s a fairly successful design. The higher level triggers worked out quite well, are well designed, and showed the flexibility of the system. In the end, if you look back at the history, we did make the right decisions. Although at the time, as always, it was not clear.”

from http://cms.web.cern.ch/news/triggering-and-data-acquisition

Reduce in stages…

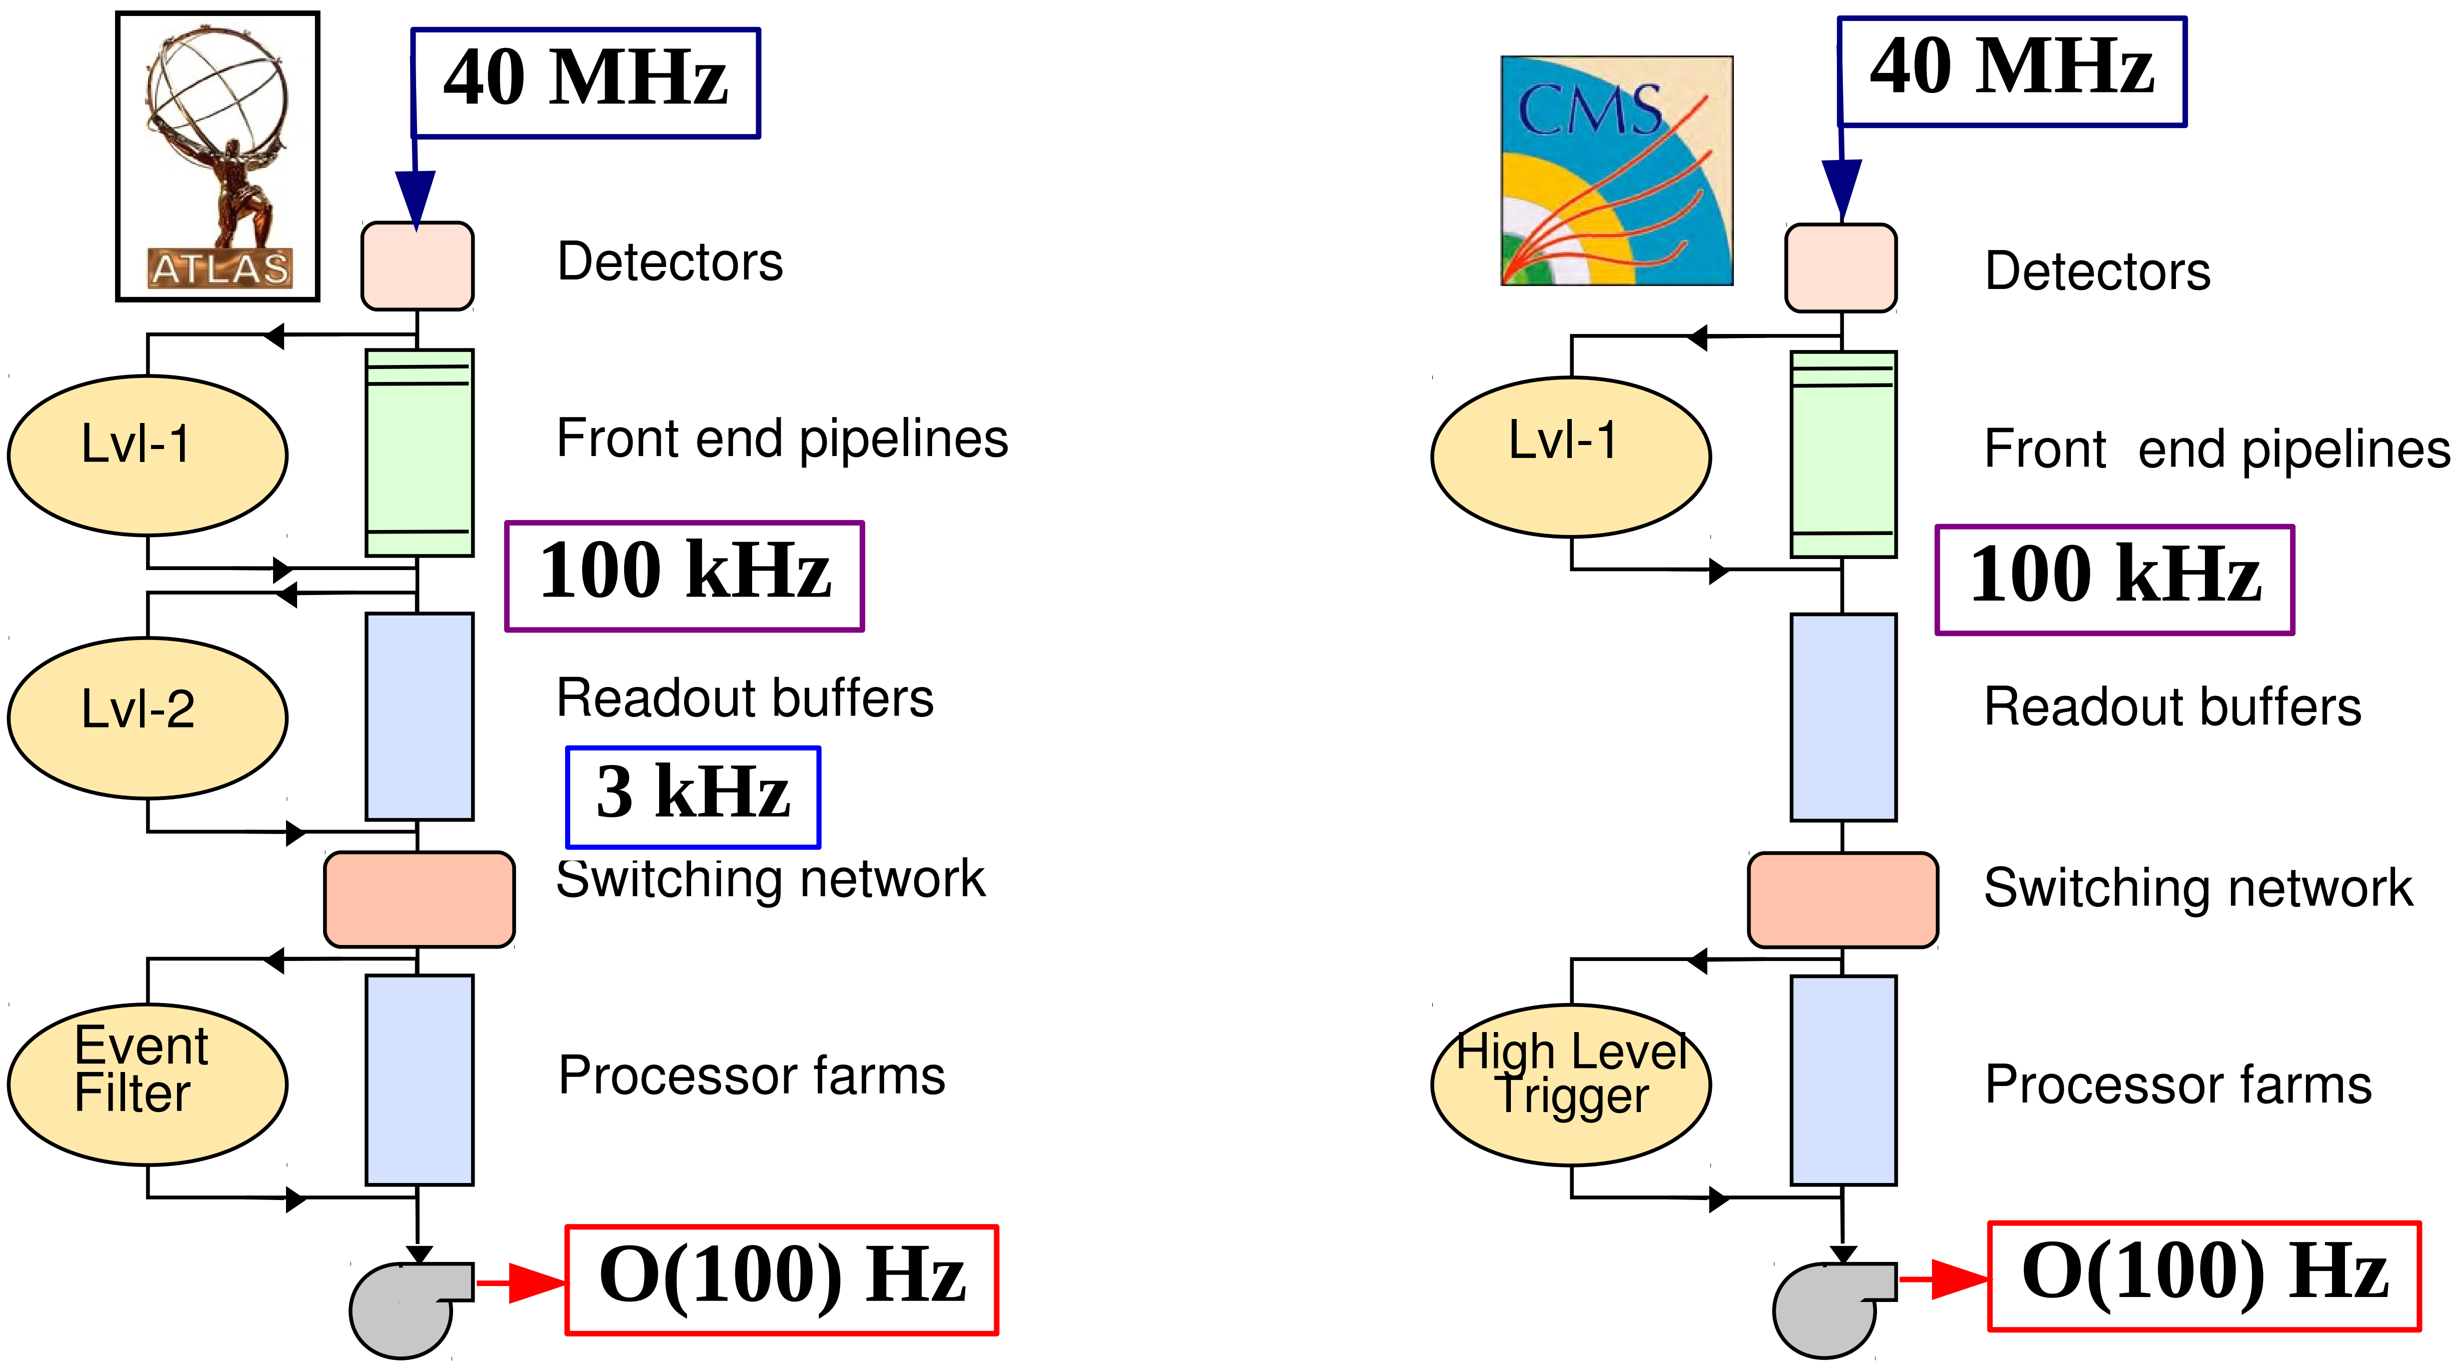

Complexity vs. Capacity

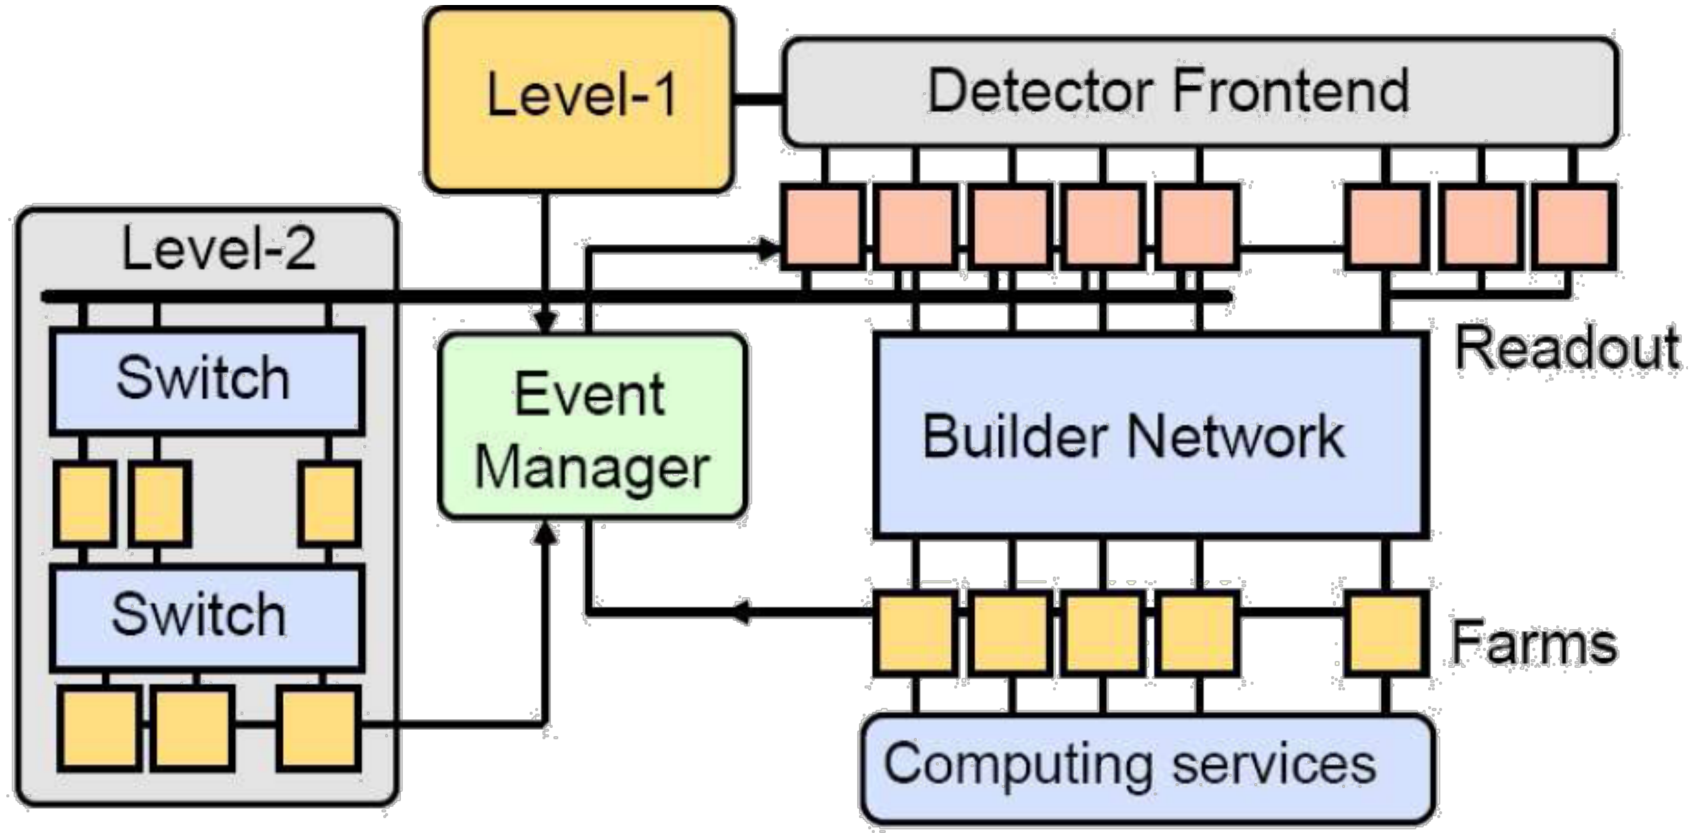

- Send a part, ask a better question; Send everything only if interesting (ATLAS)

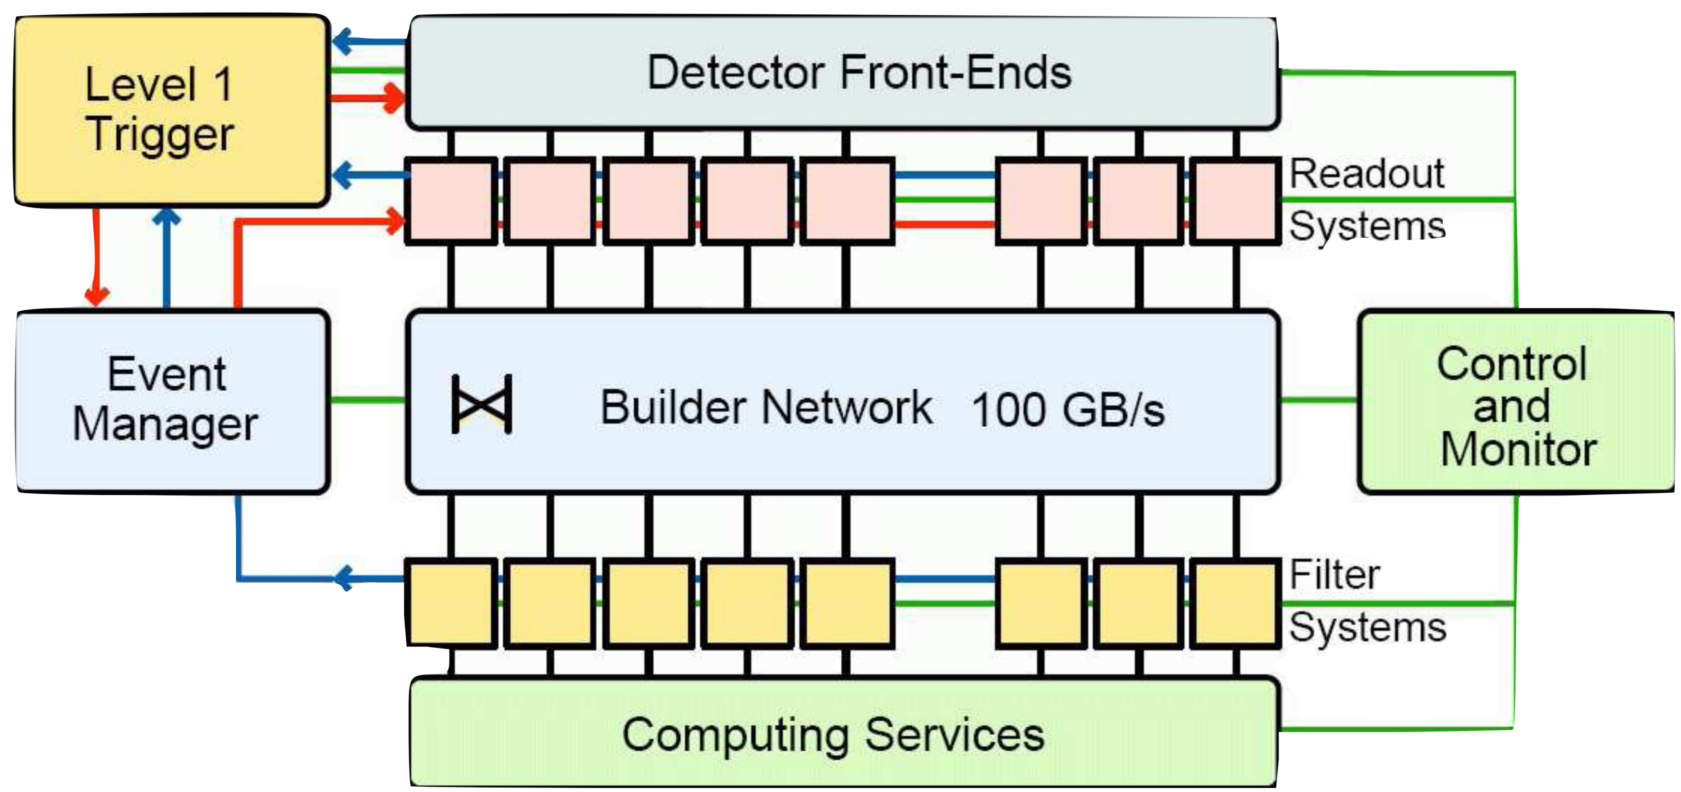

- Send everything, ask questions later (ALICE, CMS, LHCb)

Atlas Run 1

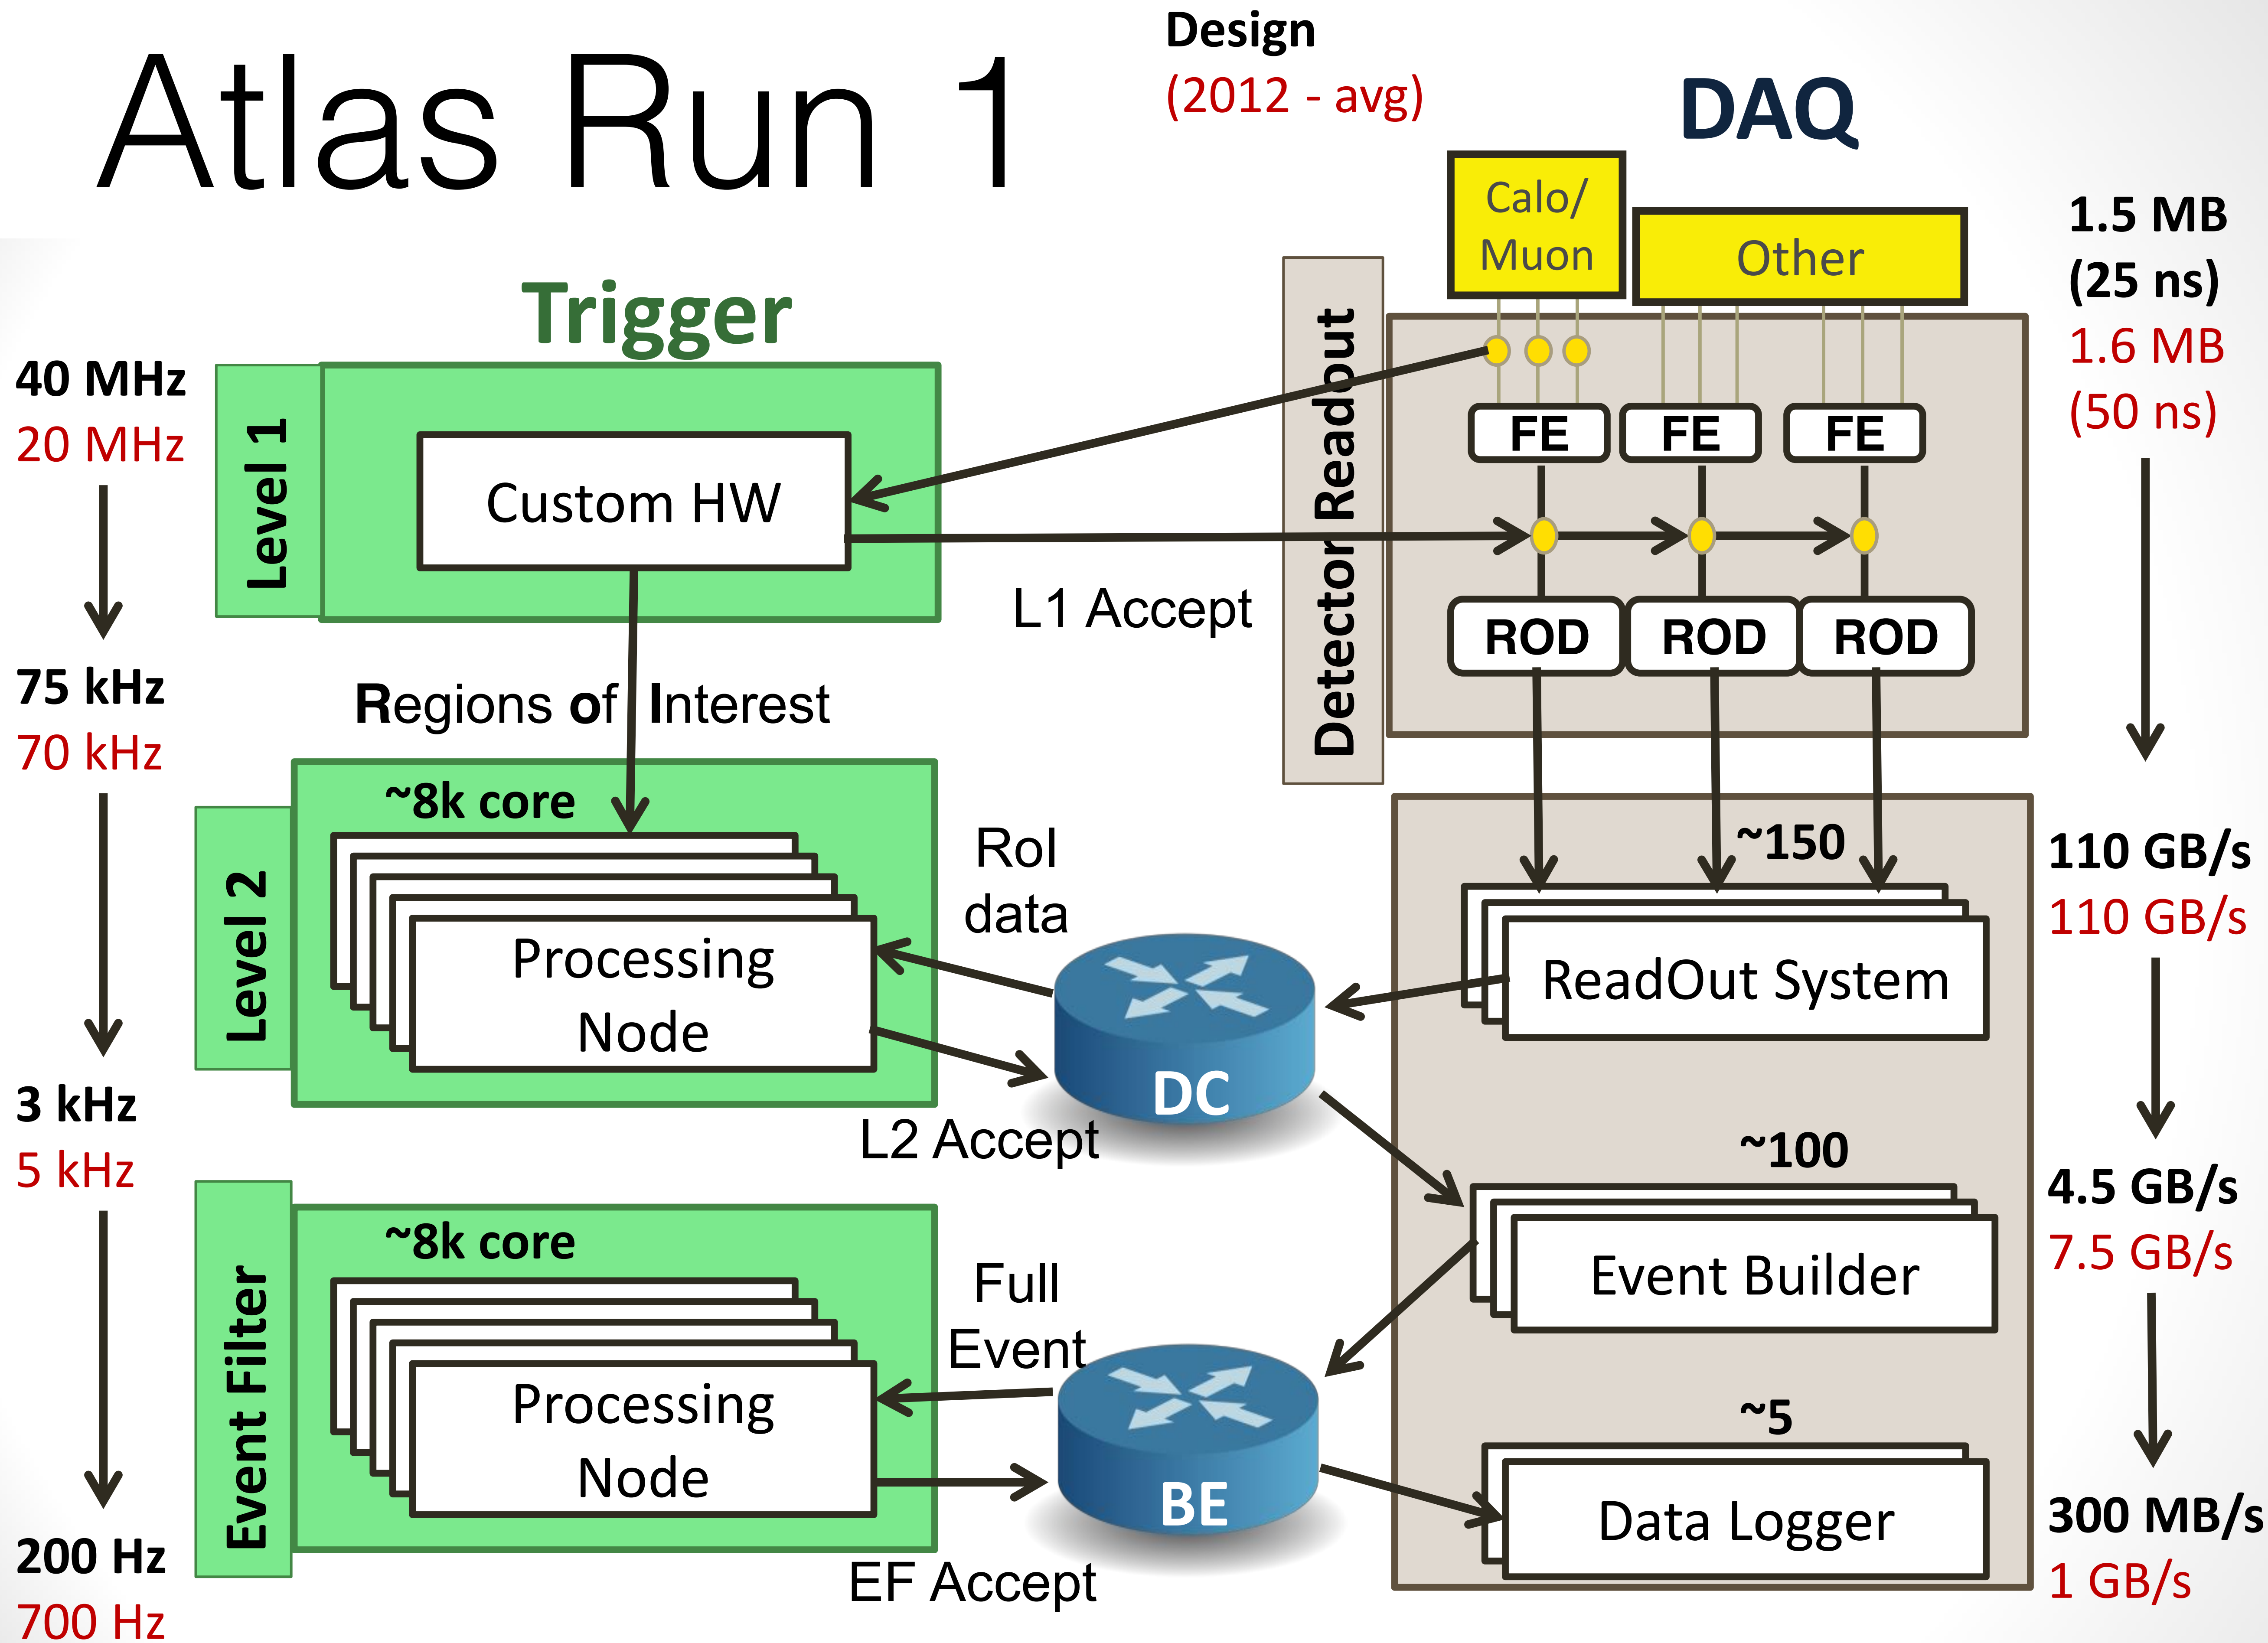

Note: Level 2 (L2) + Event Filter (EF) = High Level Trigger (HLT)

Atlas Run 2

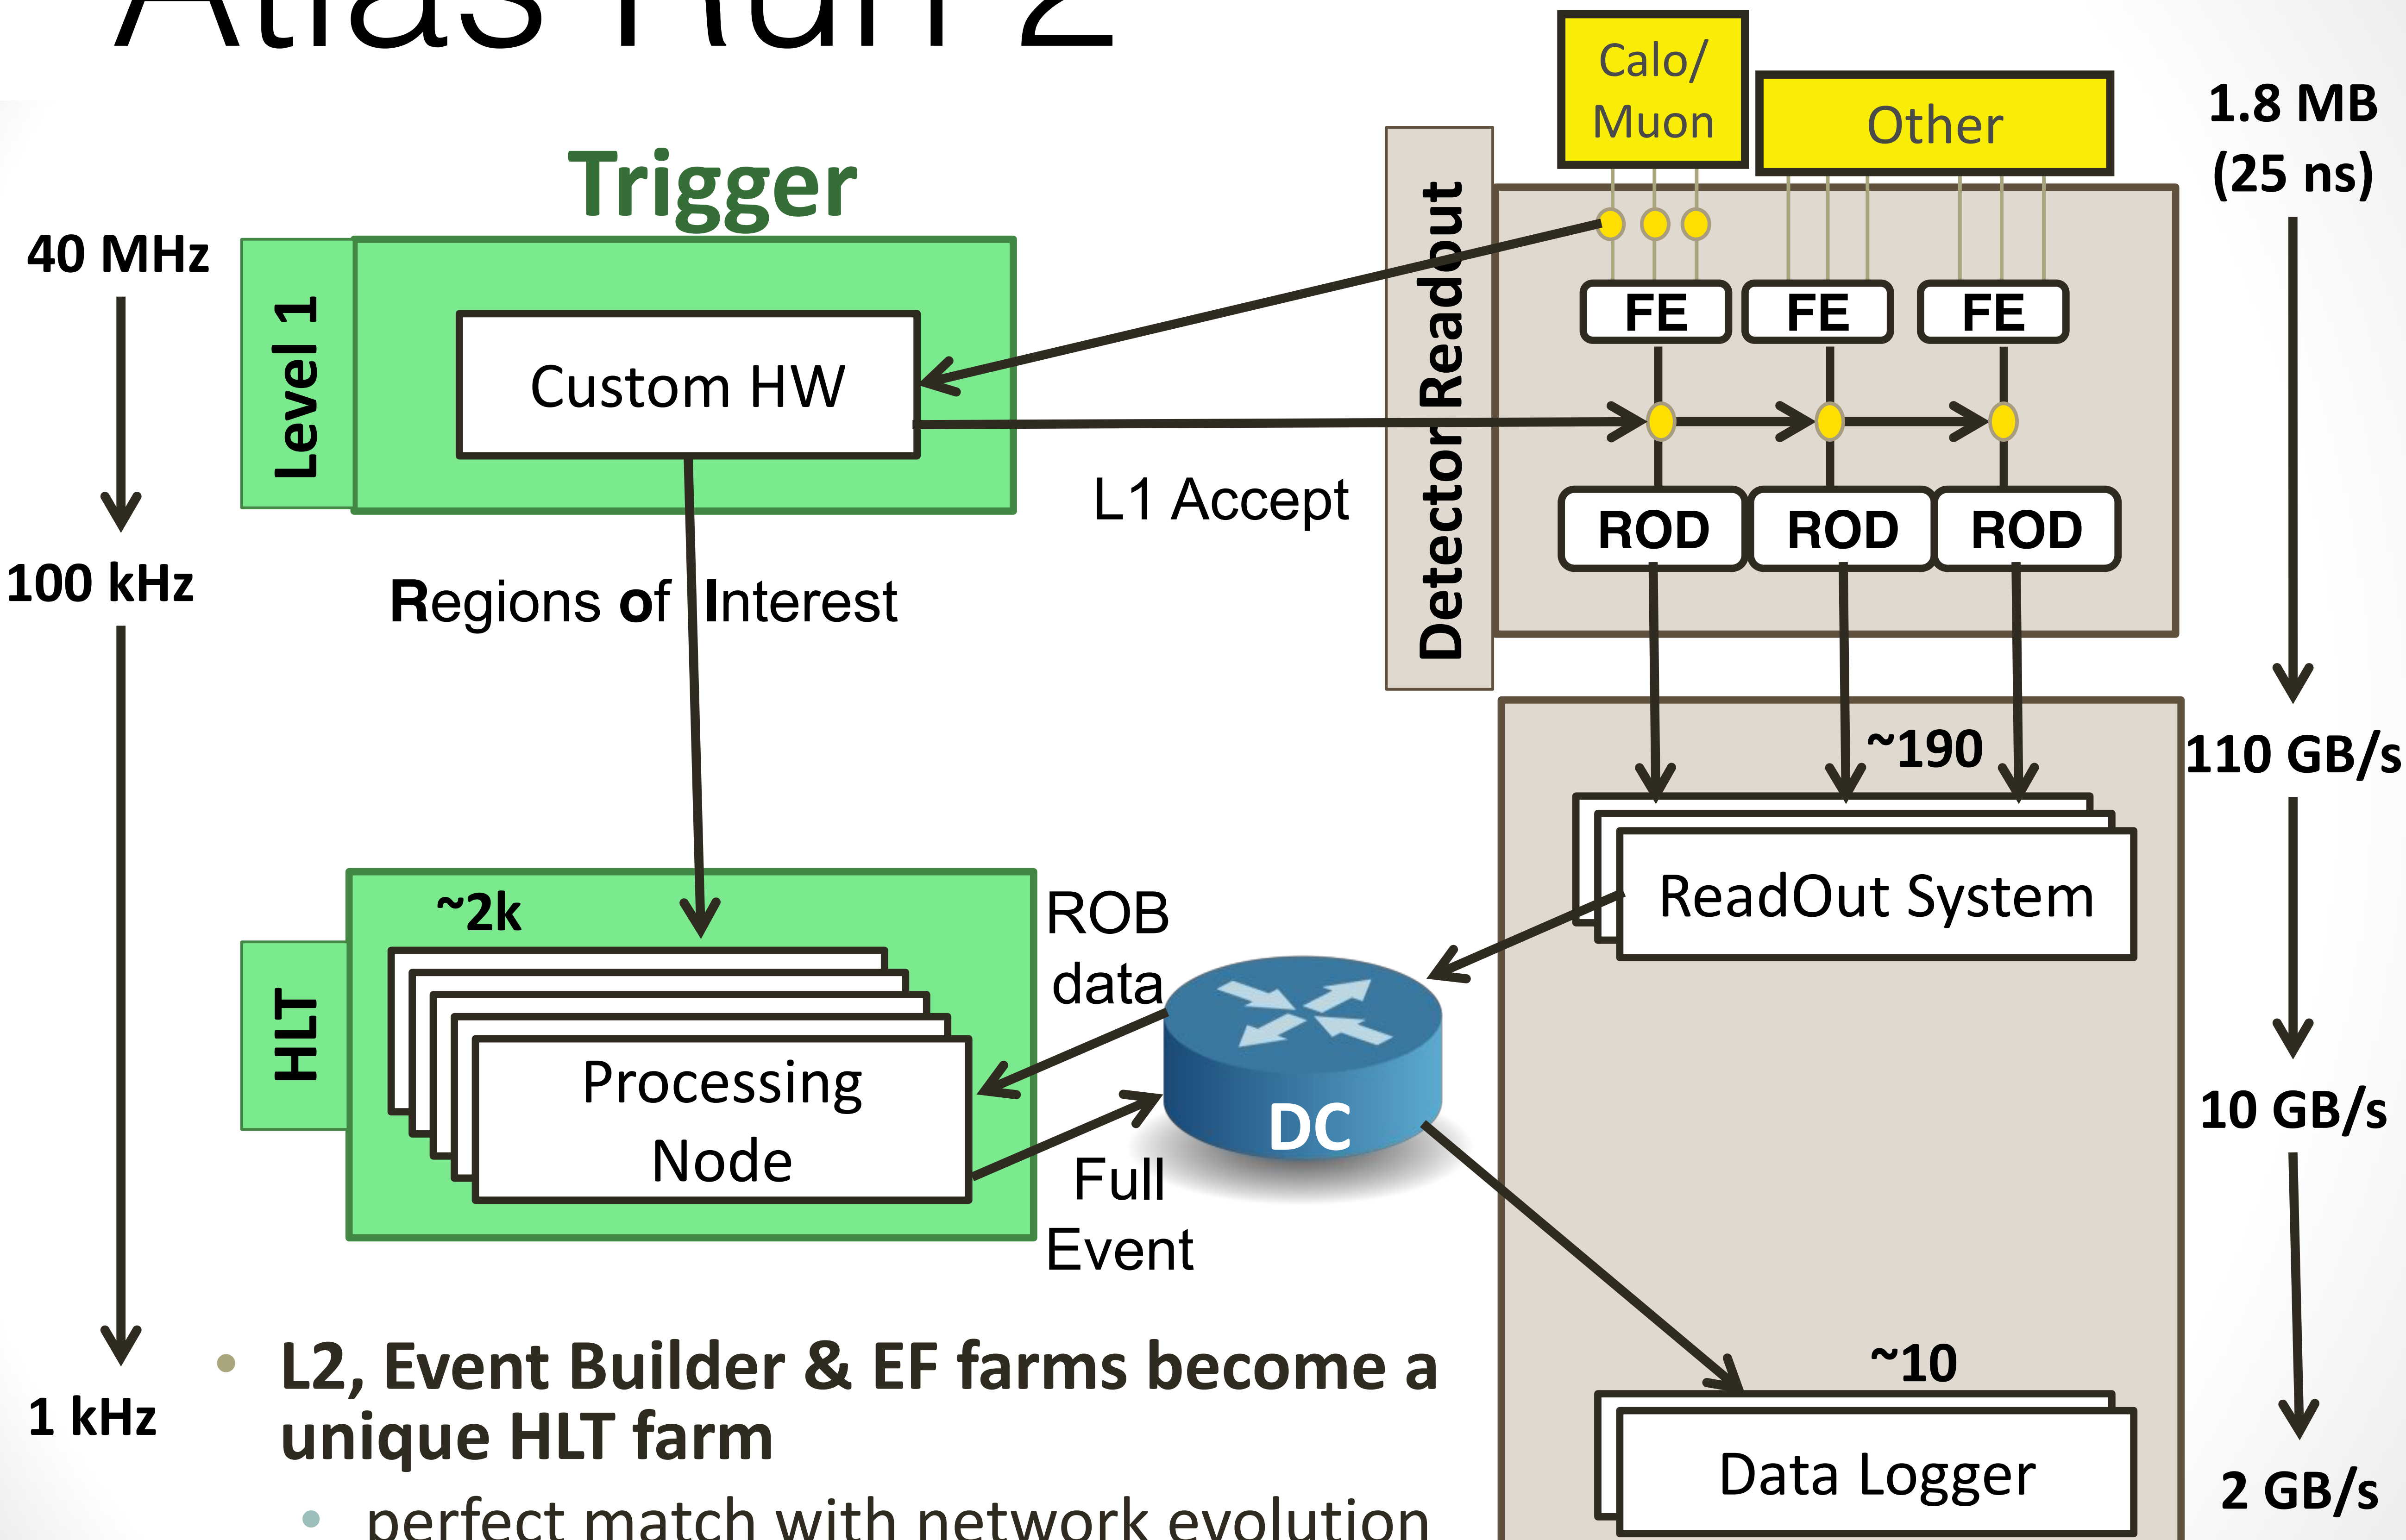

- **L2, Event Builder & EF farms become a unique HLT farm**
  - perfect match with network evolution
  - automatically balanced distribution of computing resources

LHCb

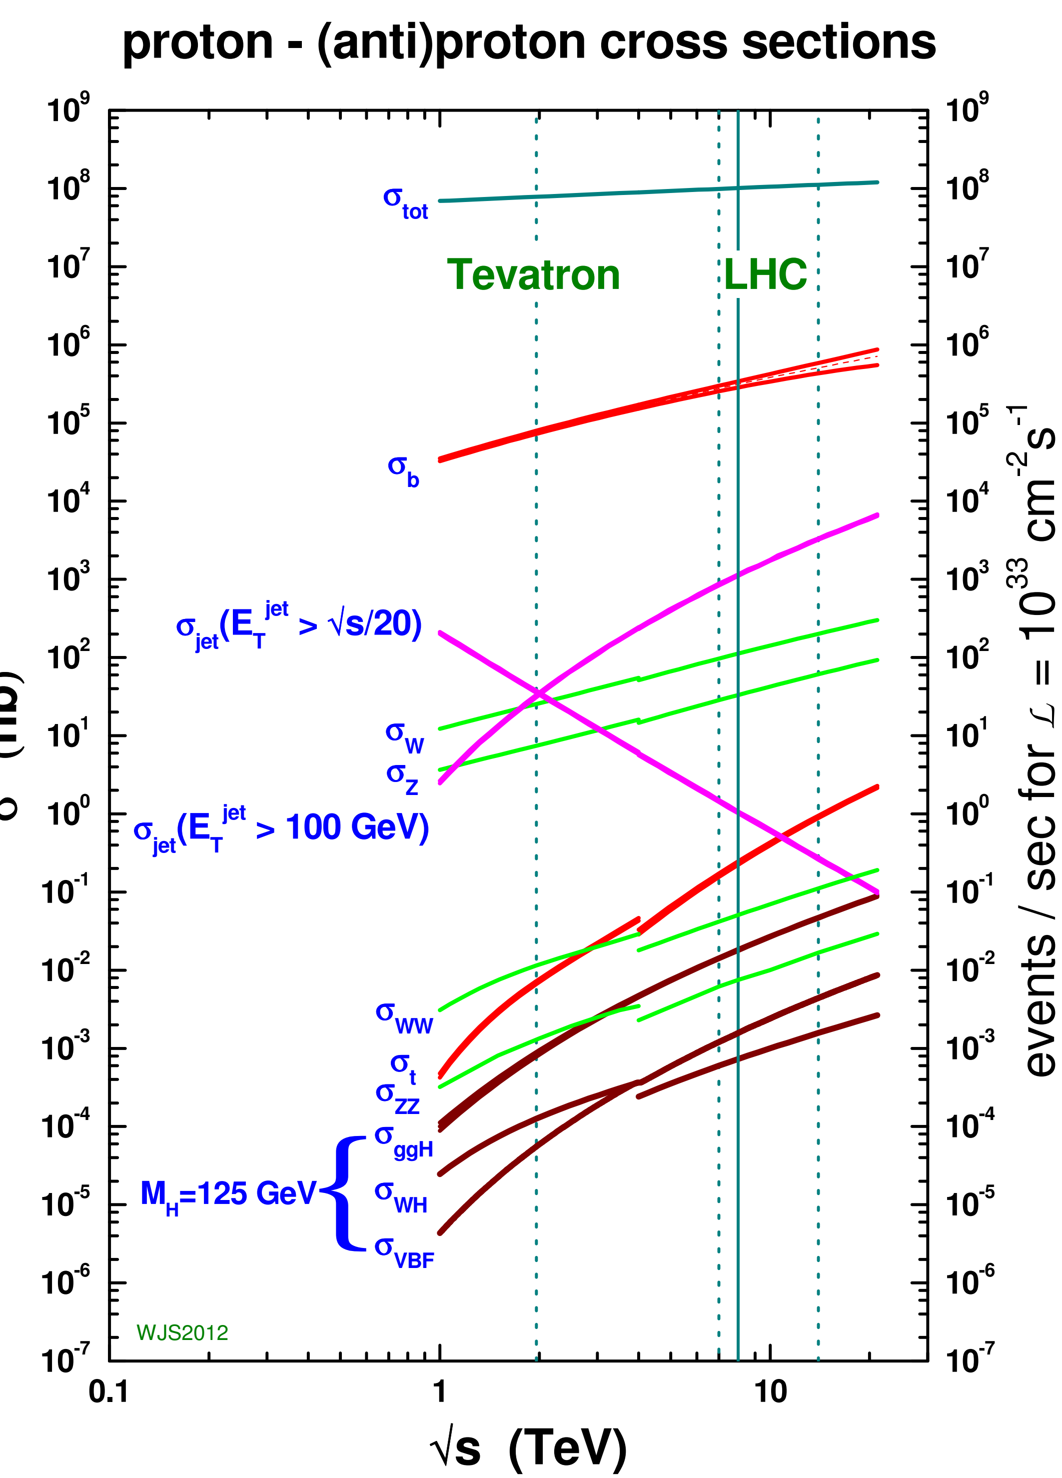

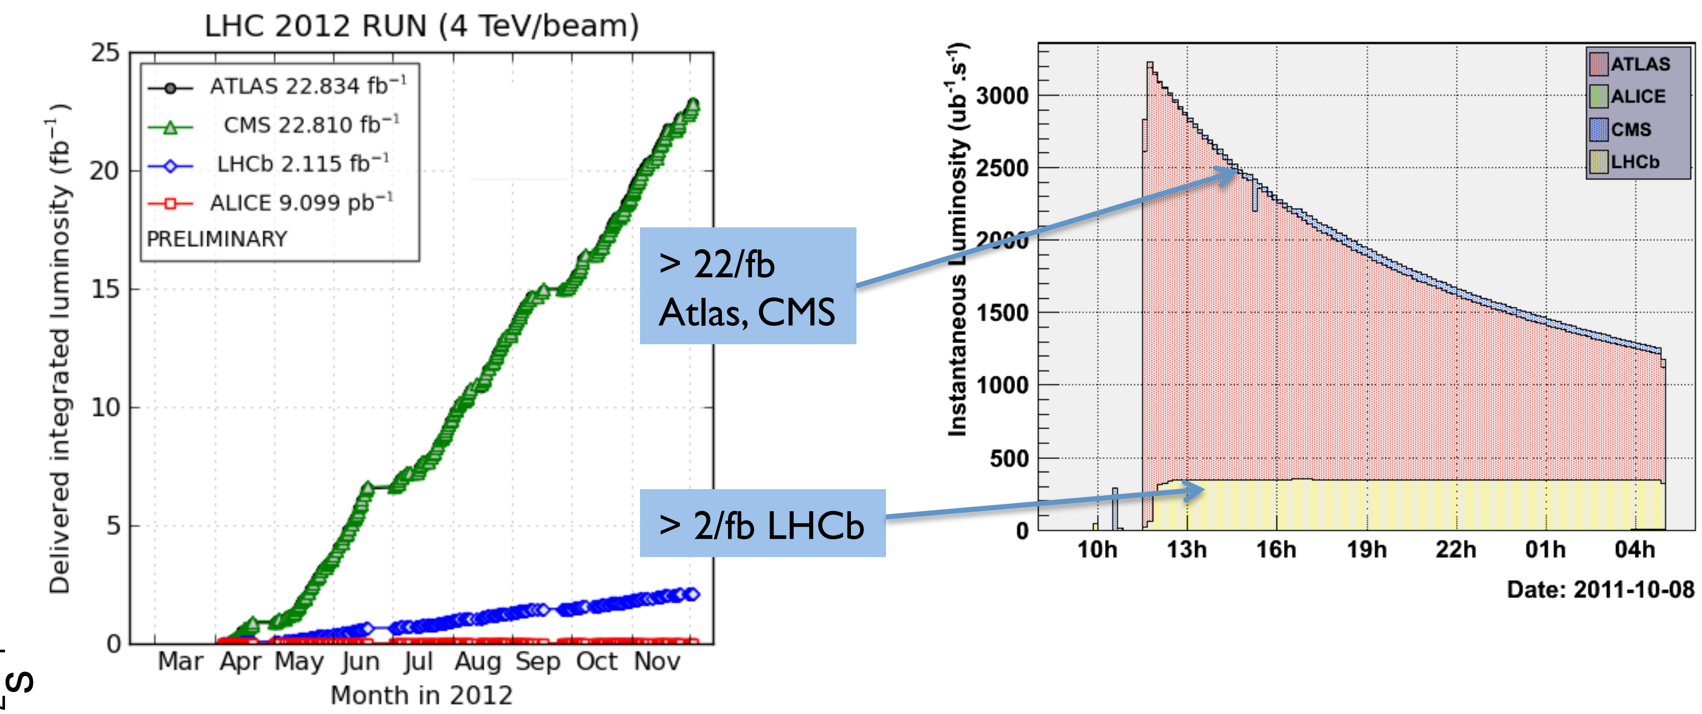

lumi-leveled @ 4 10^{32} $cm^{-2}s^{-1}$

LHCb

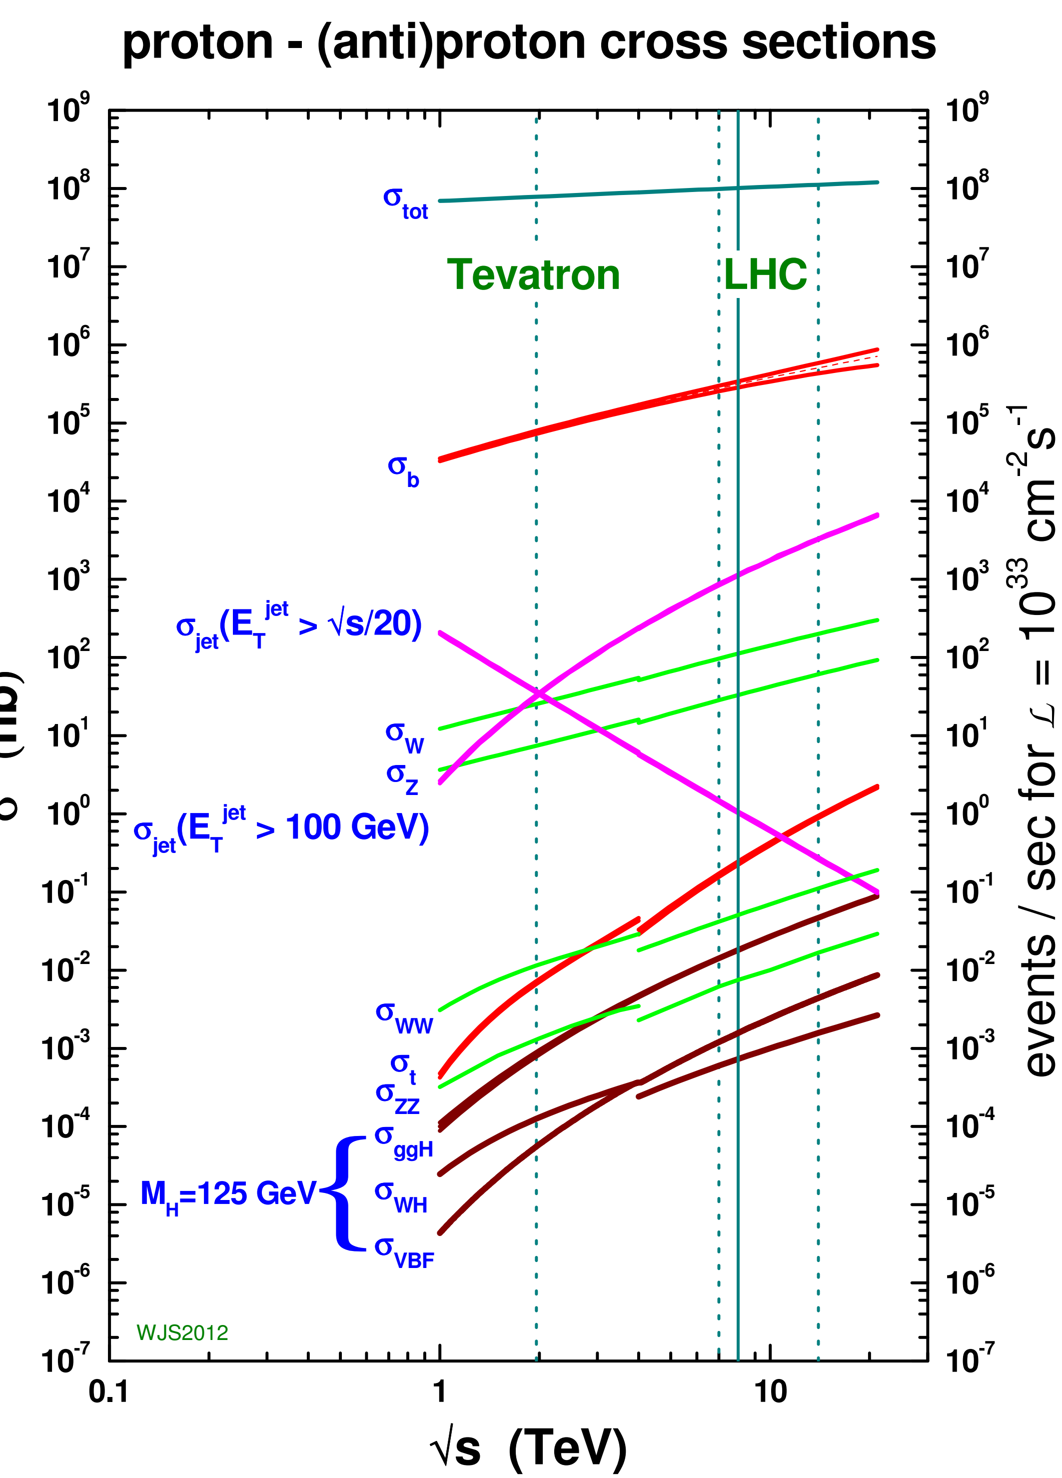
proton - (anti)proton cross sections
Tevatron
LHC
σtot
σb
σjet(ET jet > √s/20)
σW
σZ
σjet(ET jet > 100 GeV)
σWW
σt
σZZ
σggH
σWH
σVBF
MH=125 GeV
WJS2012
√s (TeV)
events / sec for ℒ = 10^33 cm^-2 s^-1

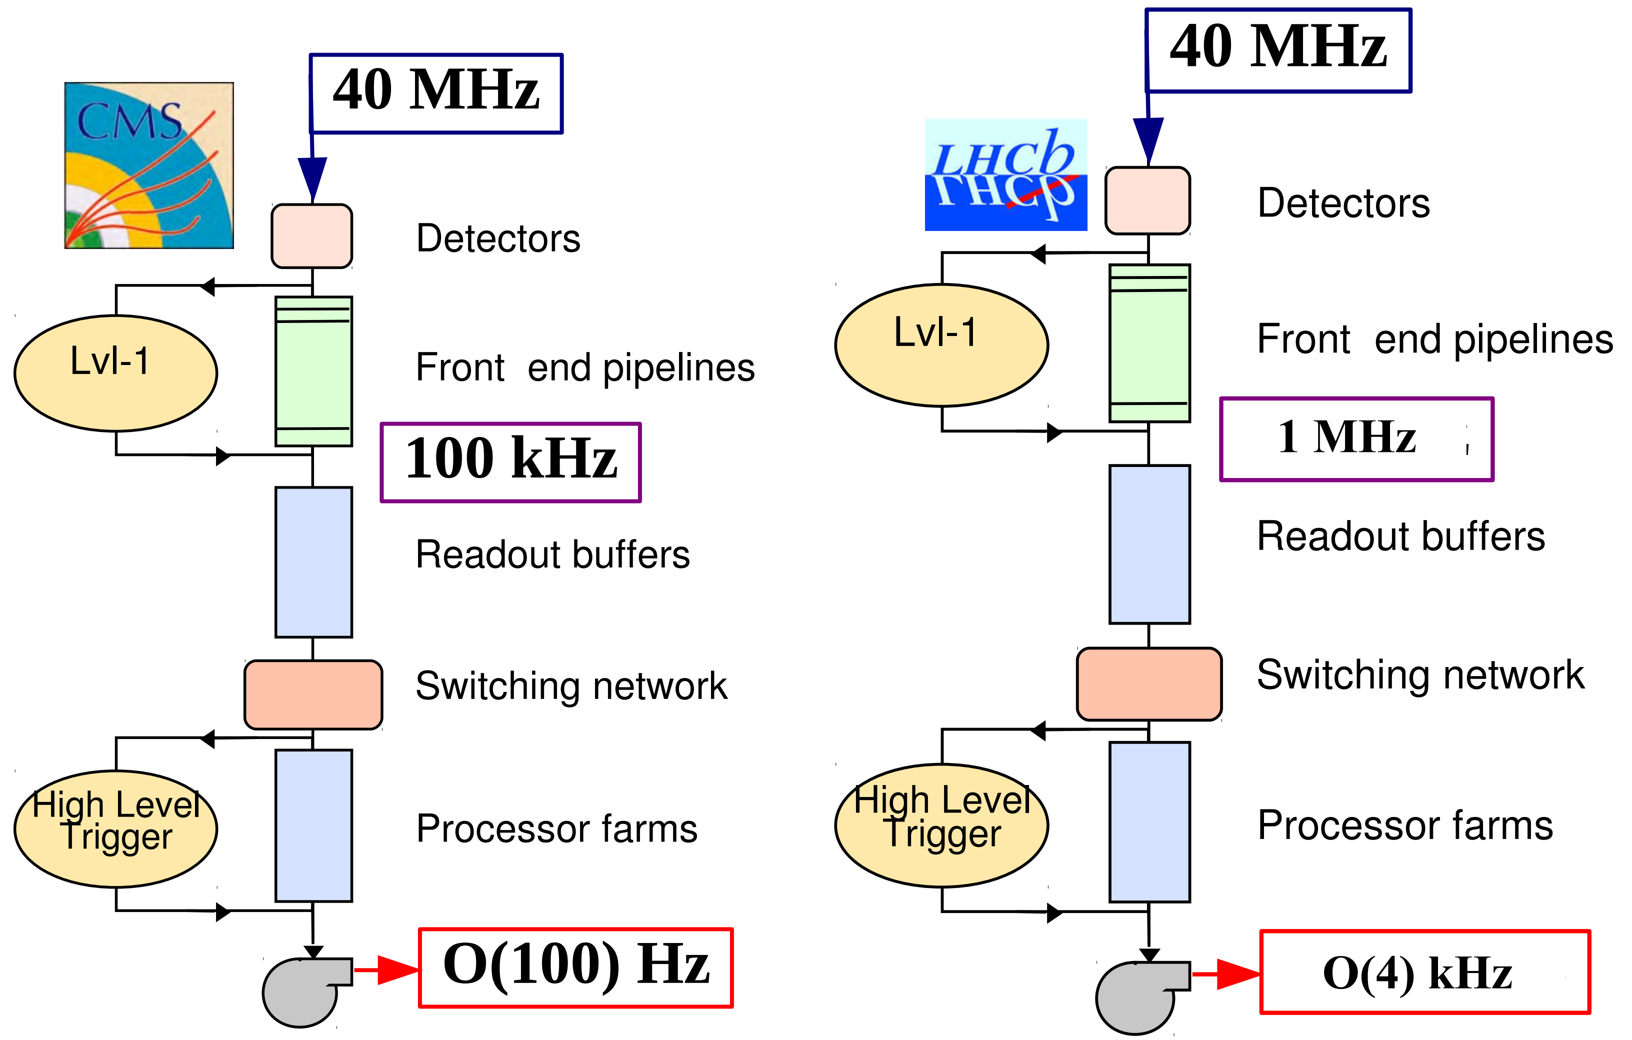
CMS
40 MHz
Detectors
Lvl-1
Front end pipelines
100 kHz
Readout buffers
Switching network
High Level Trigger
Processor farms
O(100) Hz
LHCb
40 MHz
Detectors
Lvl-1
Front end pipelines
1 MHz
Readout buffers
Switching network
High Level Trigger
Processor farms
O(4) kHz

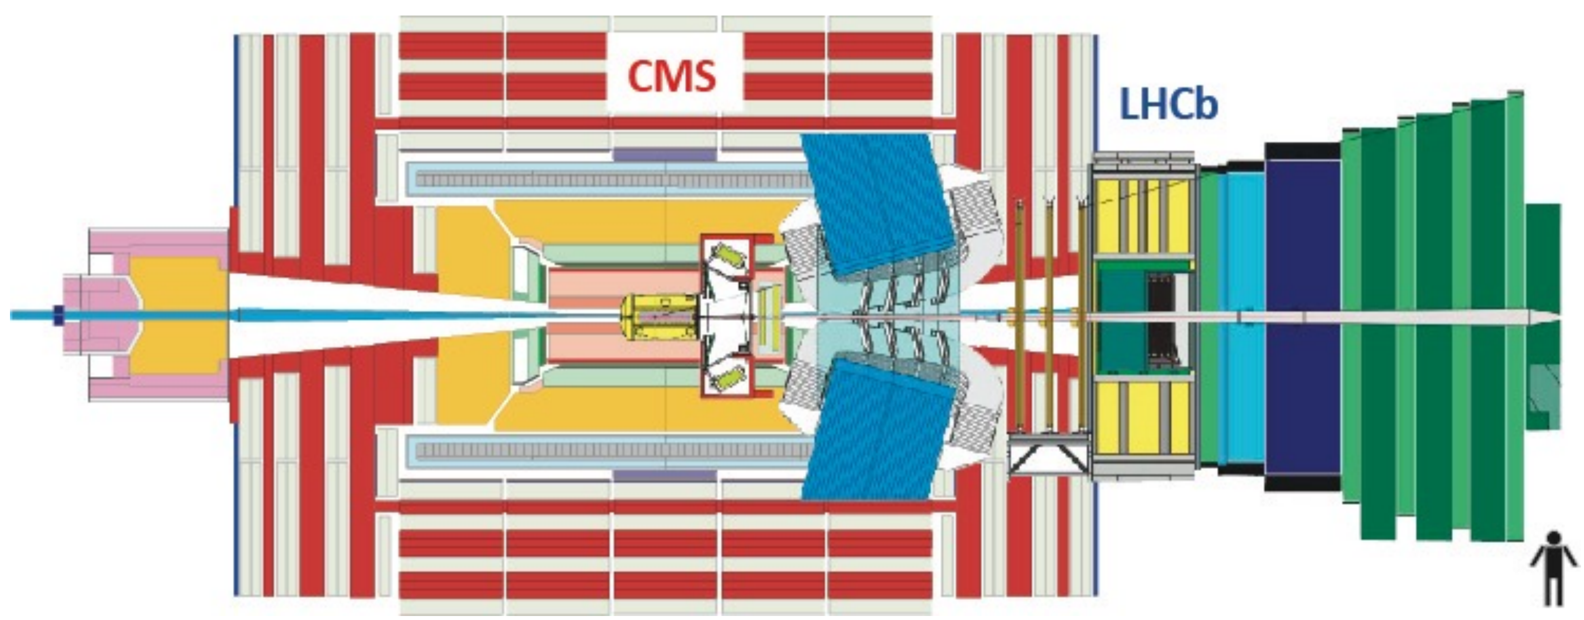
CMS
LHCb

Data Rates

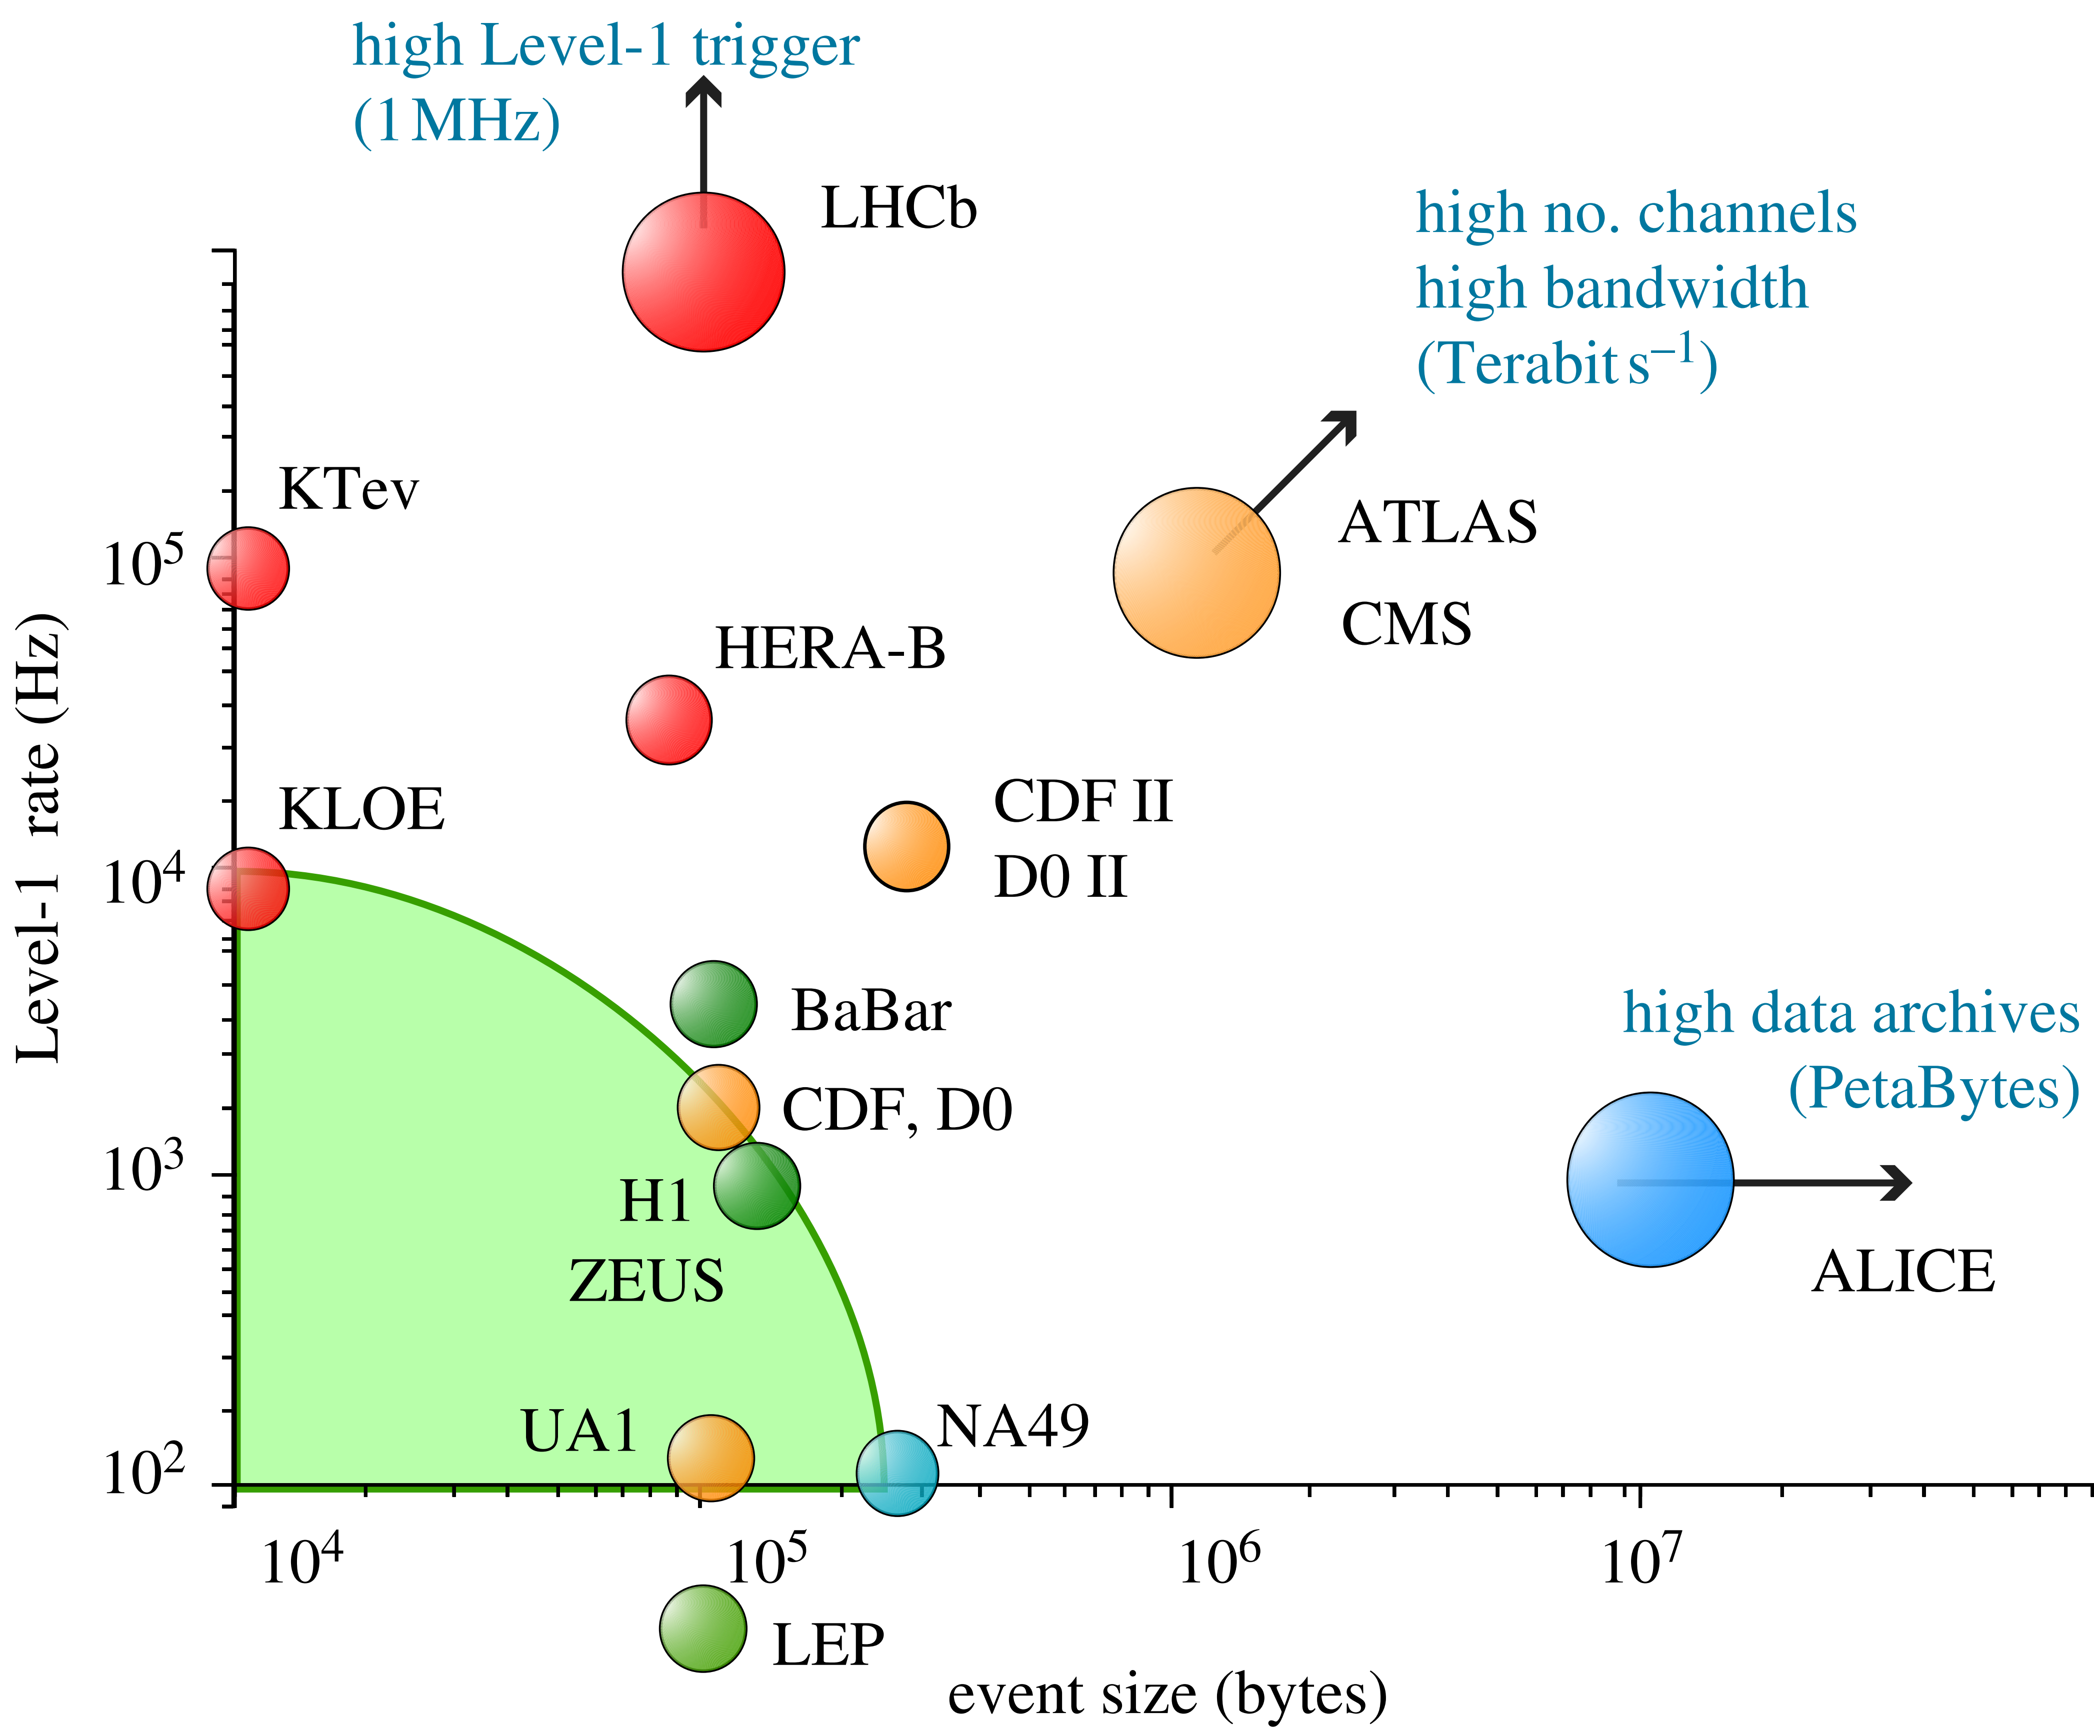

high Level-1 trigger
(1 MHz)
LHCb
high no. channels
high bandwidth
(Terabit s^{-1})
KTev
ATLAS
CMS
HERA-B
CDF II
D0 II
KLOE
BaBar
high data archives
(PetaBytes)
CDF, D0
H1
ZEUS
ALICE
UA1
NA49
LEP
Level-1 rate (Hz)
10^2
10^3
10^4
10^5
10^4
10^5
10^6
10^7
event size (bytes)

Data Rates

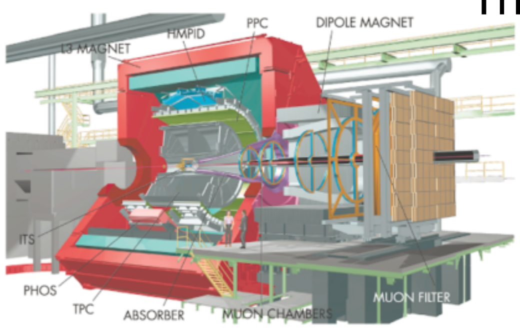

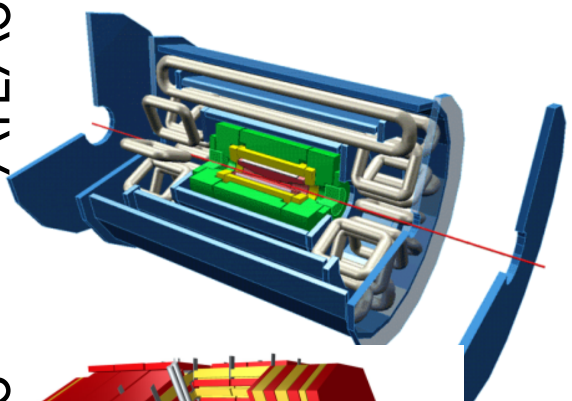

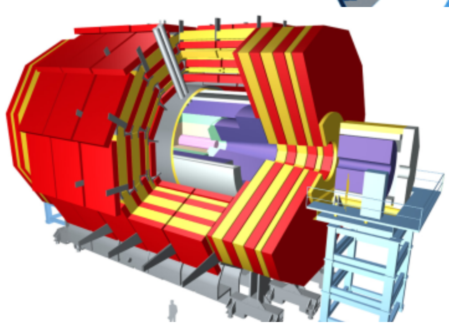

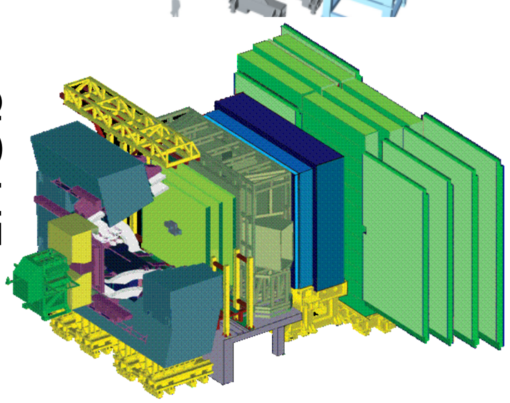

	No.Levels Trigger	**Level-0,1,2** Rate (Hz)	**Event** Size (Byte)	**Readout** Bandw.(B/s)	**HLT Out** B/s (Event/s)
ALICE	**4**	Pb-Pb **500** p-p **1K**	**50M** **2M**	**25G**	**1250M**(0.1K) **200M** (0.1K)
ATLAS	**3**	LV-1 **75K** LV-2 **3K**	**1.5M**	**4.5G**	**900M** (0.6K)
CMS	**2**	LV-1 **75K**	**1M**	**100G**	**1000M** (1K)
LHCb	**2**	LV-0 **1000K**	**0.05M**	**50G**	**200M** (4K)

The background…

- So what do those 70 mb of ‘uninteresting’ minimum bias events look like?

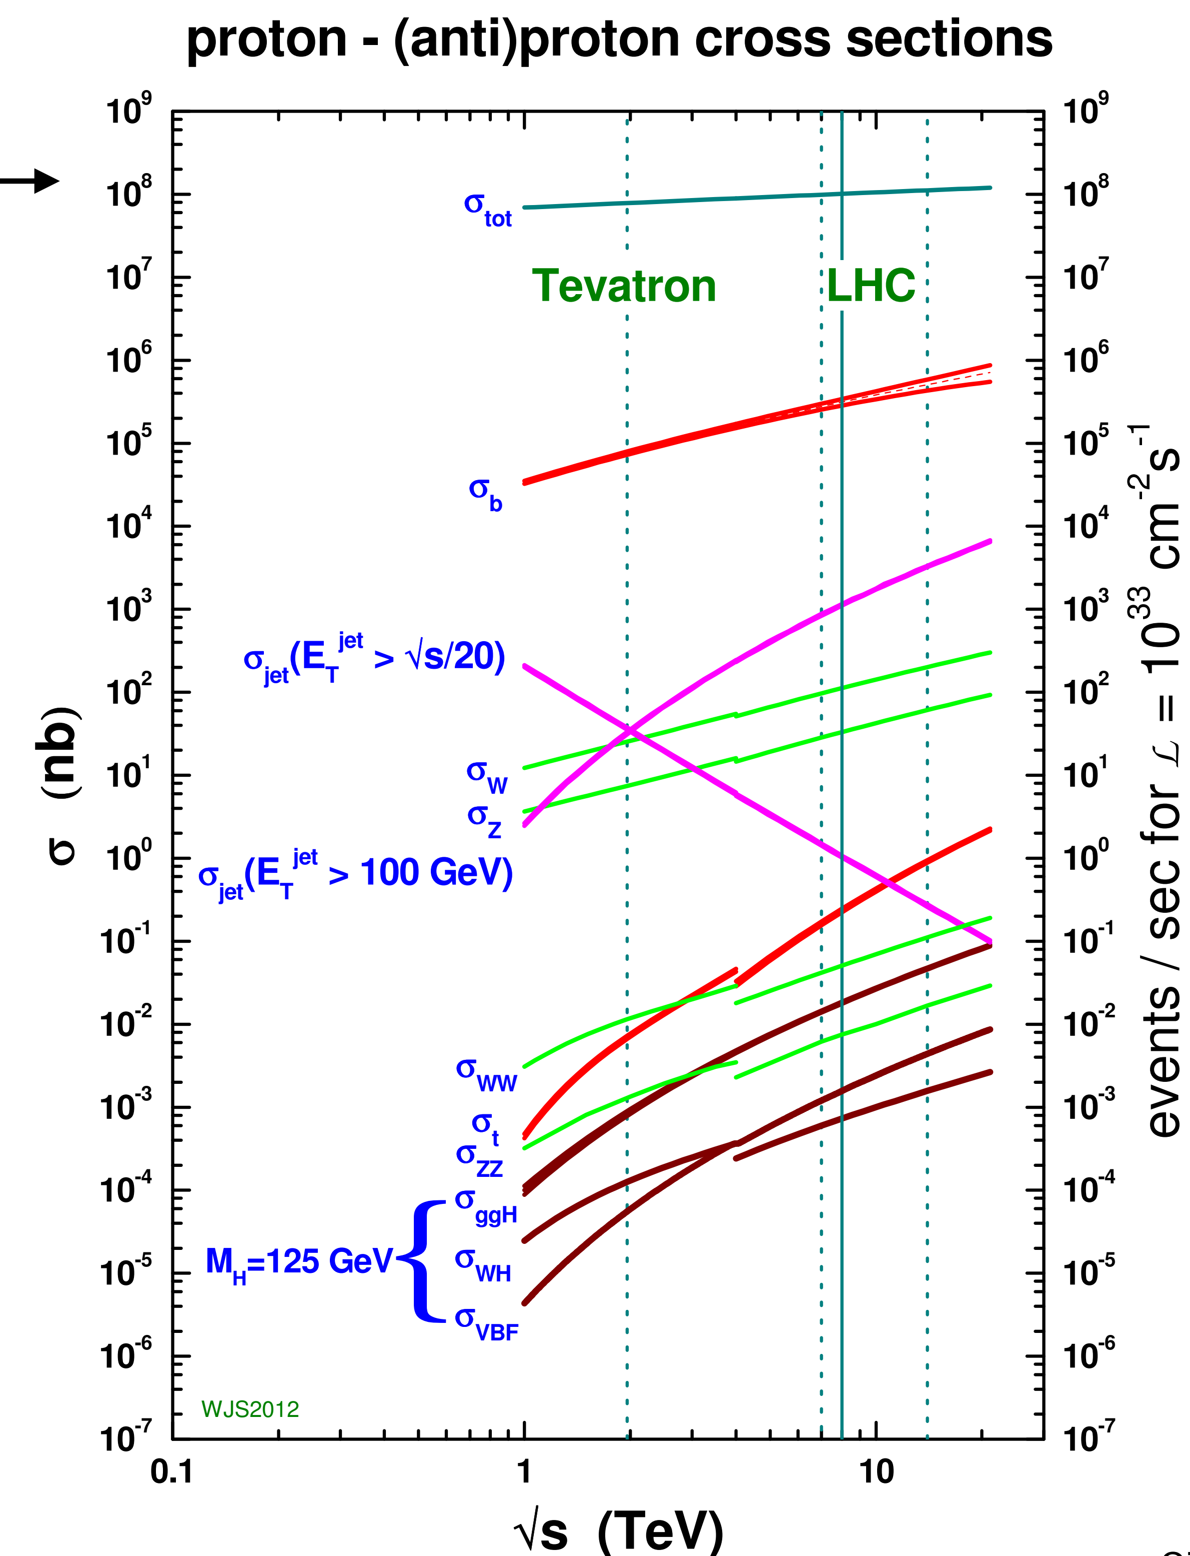

Know Your background

- $\sigma(pp)$ = 70 mb; @ 10^{34} $cm^{-2}s^{-1}$ = 700 MHz
- @ 40 MHz: pileup of O(20)

- 1 min bias event @ 14 TeV:
  - 5 charged particles (dominantly pions) per unit rapidity
  - isospin: $\pi^0 : \pi^\pm$ = 1:2

- within $|\eta|<2.5$: 25 $\pi^\pm$ + 25 photons per min bias
- # tracks > 3 GeV/c: 0.8% x 25 = 0.2 per min bias event

- @ 20 pileup: ~4 min bias tracks with P_T > 3.0 GeV.

P_T tracks in minimum bias events

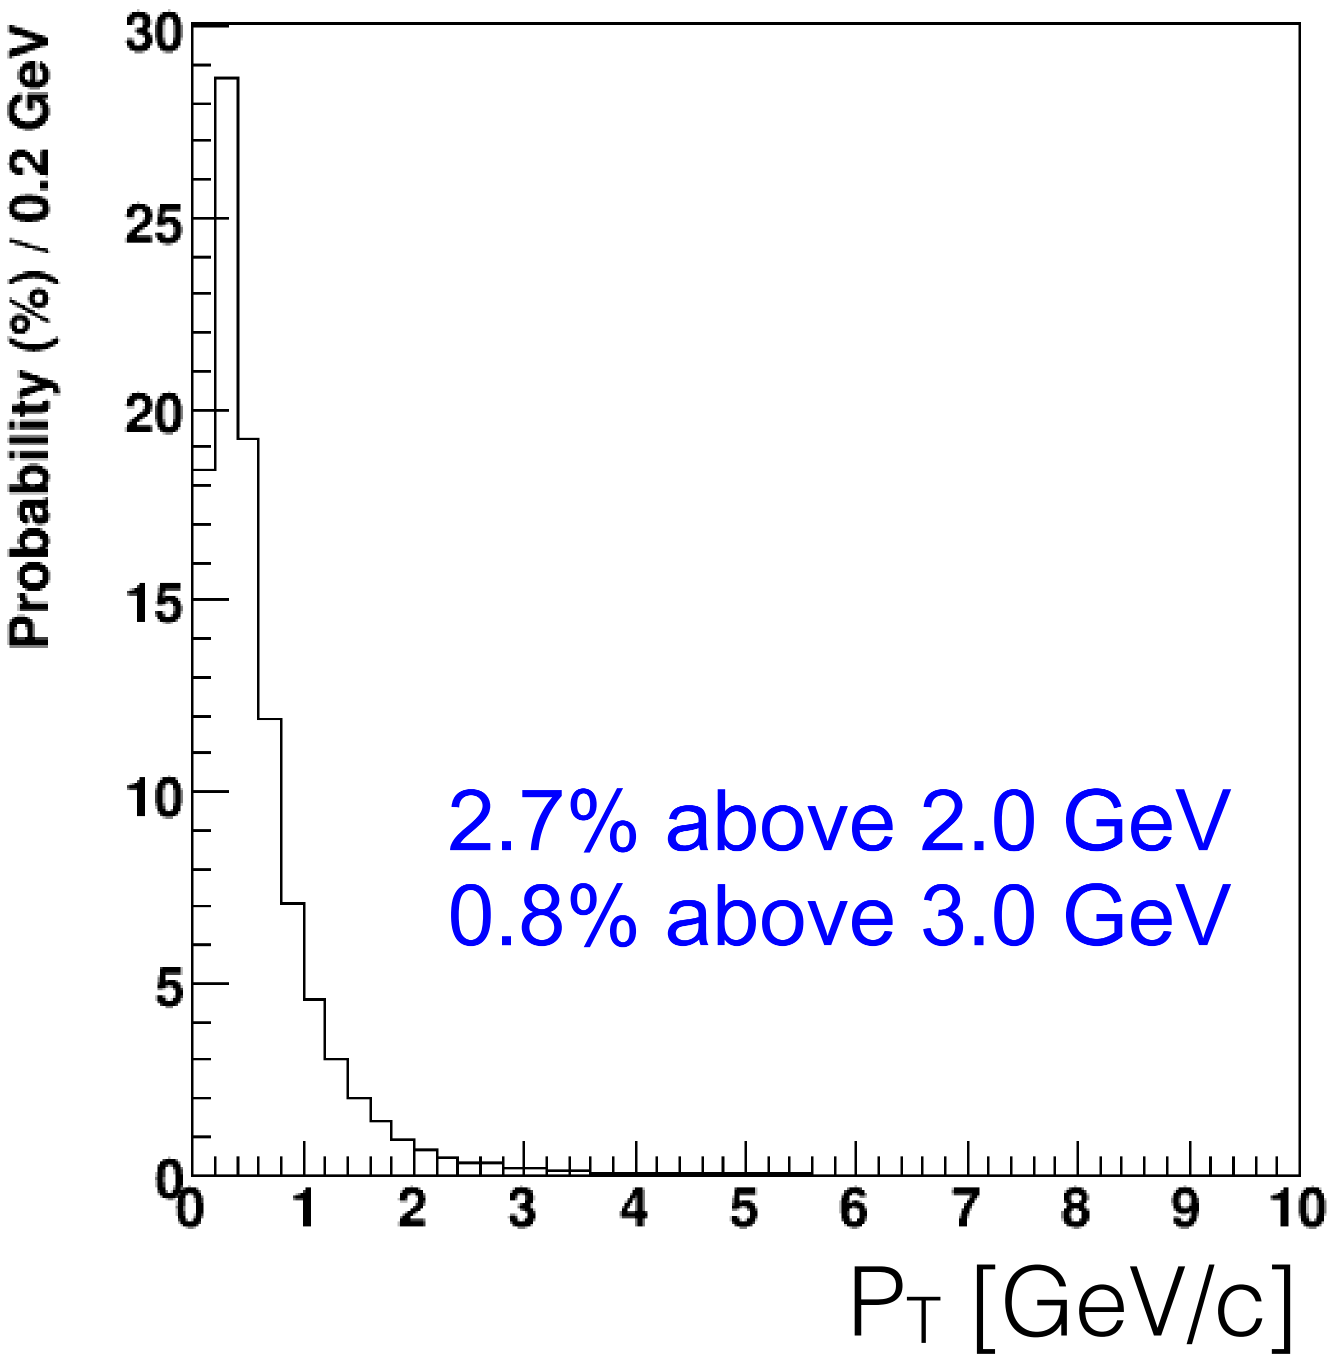

Know Your Enemy

this repeats every 25 ns…

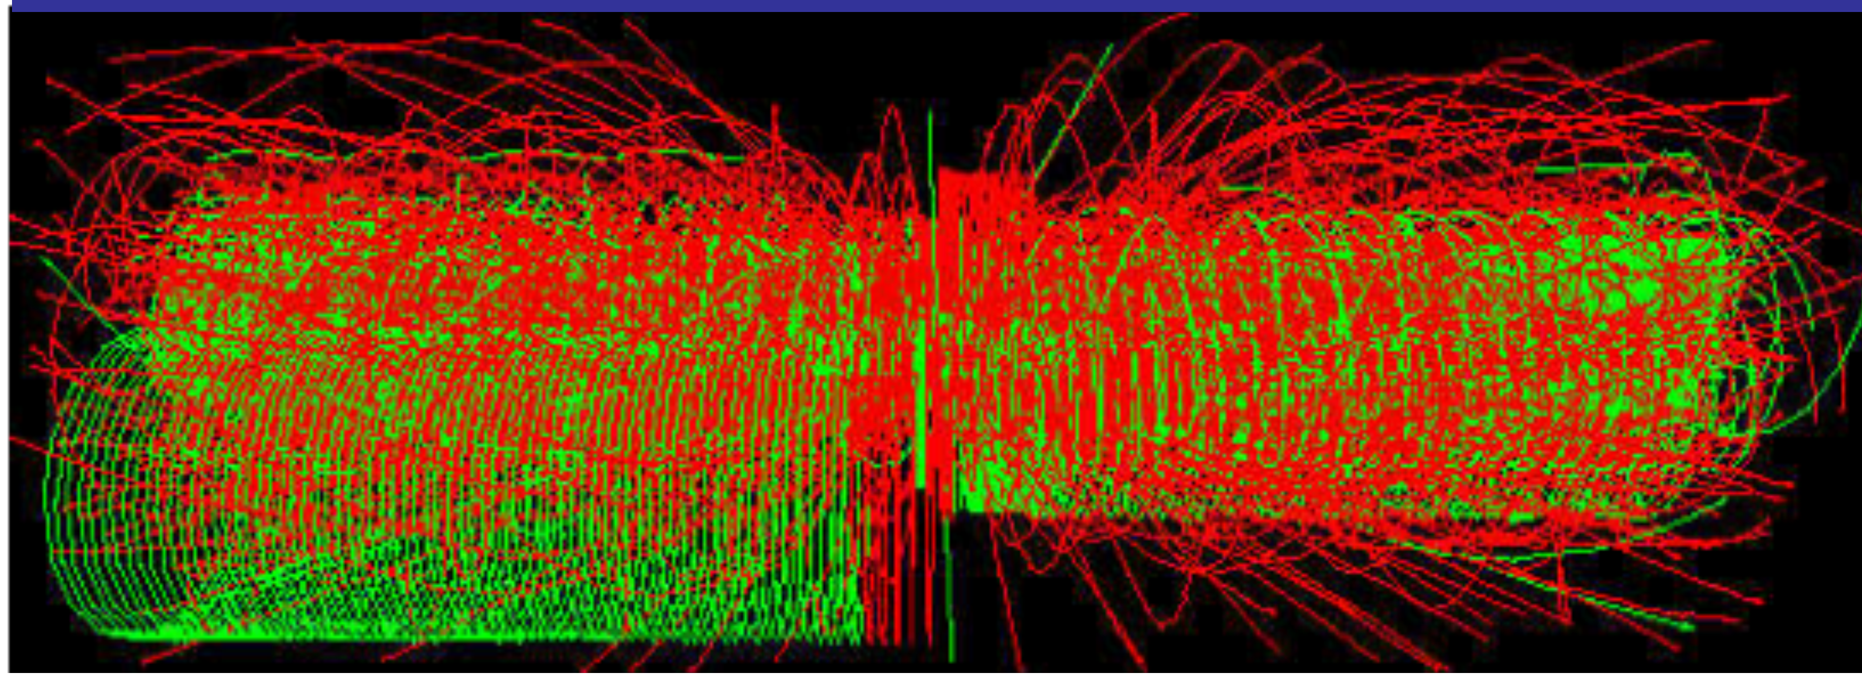

- $\sigma(pp)$ = 70 mb; @ 10^{34} $cm^{-2}s^{-1}$ = 700 MHz
- O(20) min bias events will overlap

Know Your Enemy

- $\sigma(pp)$ = 70 mb; @ 10^{34} $cm^{-2}s^{-1}$ = 700 MHz
- O(20) min bias events will overlap
- $H \rightarrow ZZ^* \rightarrow \mu\mu\mu\mu$

this repeats every 25 ns…

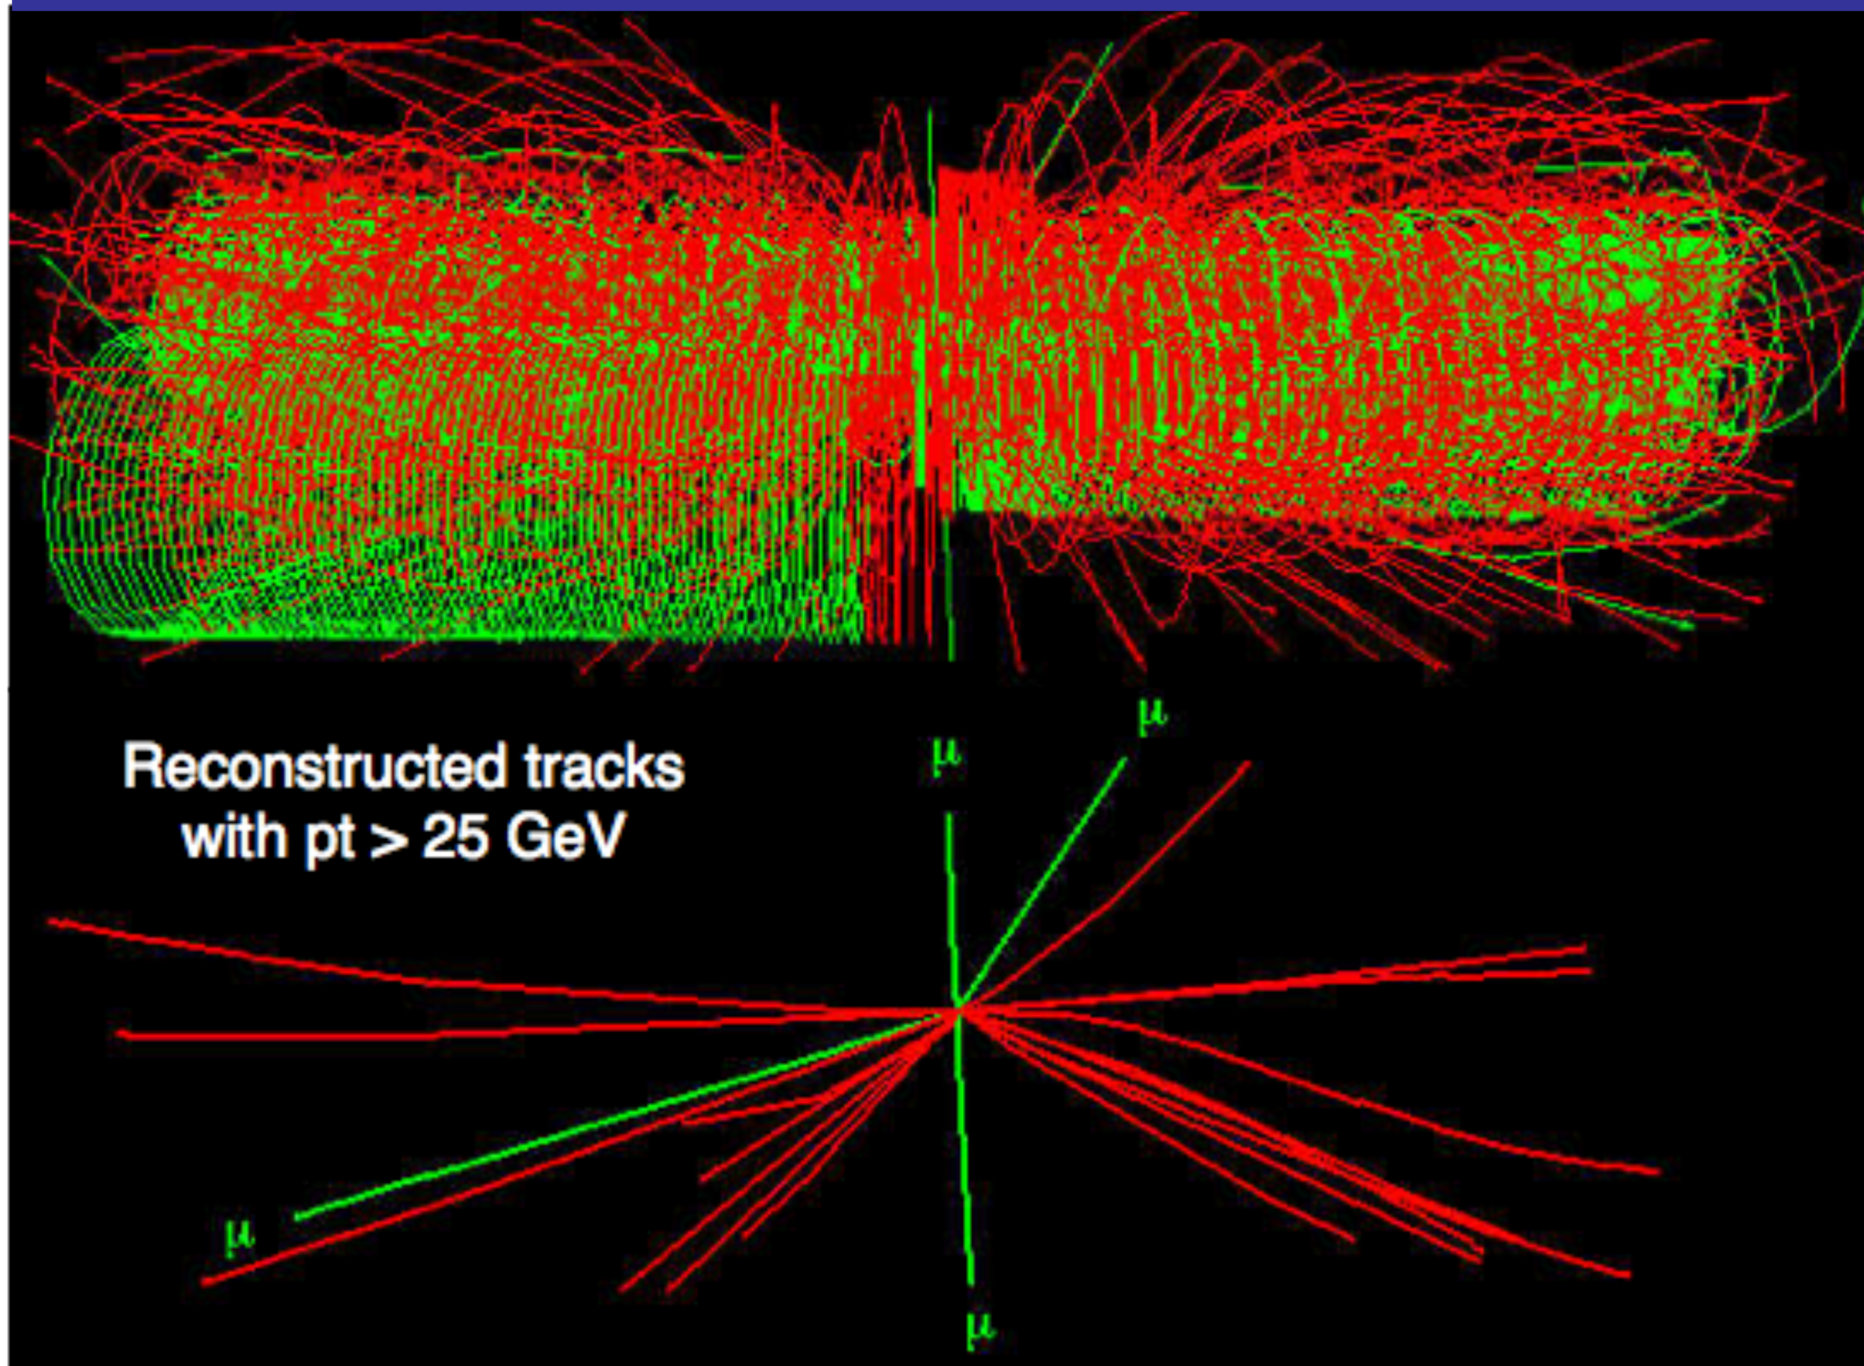

and sometimes this is 'hidden' inside

Know your signal

- pp collisions produce typically hadrons with P_T of O(1) GeV/c
- “interesting” physics has leptons & hadrons with “large” pt:
  - $W \rightarrow e\nu$: M(W)=80Gev/c^2, P_T(e) ~30-40 GeV/c
  - $H \rightarrow \gamma\gamma$: M(H)=125 GeV/$c^2$; $P_T(\gamma)$ ~ 50-60 GeV/c
  - $B_s \rightarrow \mu\mu$; M(B) = 5 GeV/c^2; P_T(mu) ~ 2 GeV/c

Know your signal

- pp collisions produce typically hadrons with P_T of O(1) GeV/c
- "interesting" physics has leptons & hadrons with "large" pt:
    - W→eν: M(W)=80Gev/c^2, P_T(e) ~30-40 GeV/c
    - H→γγ: M(H)=125 GeV/c^2; P_T(γ) ~ 50-60 GeV/c

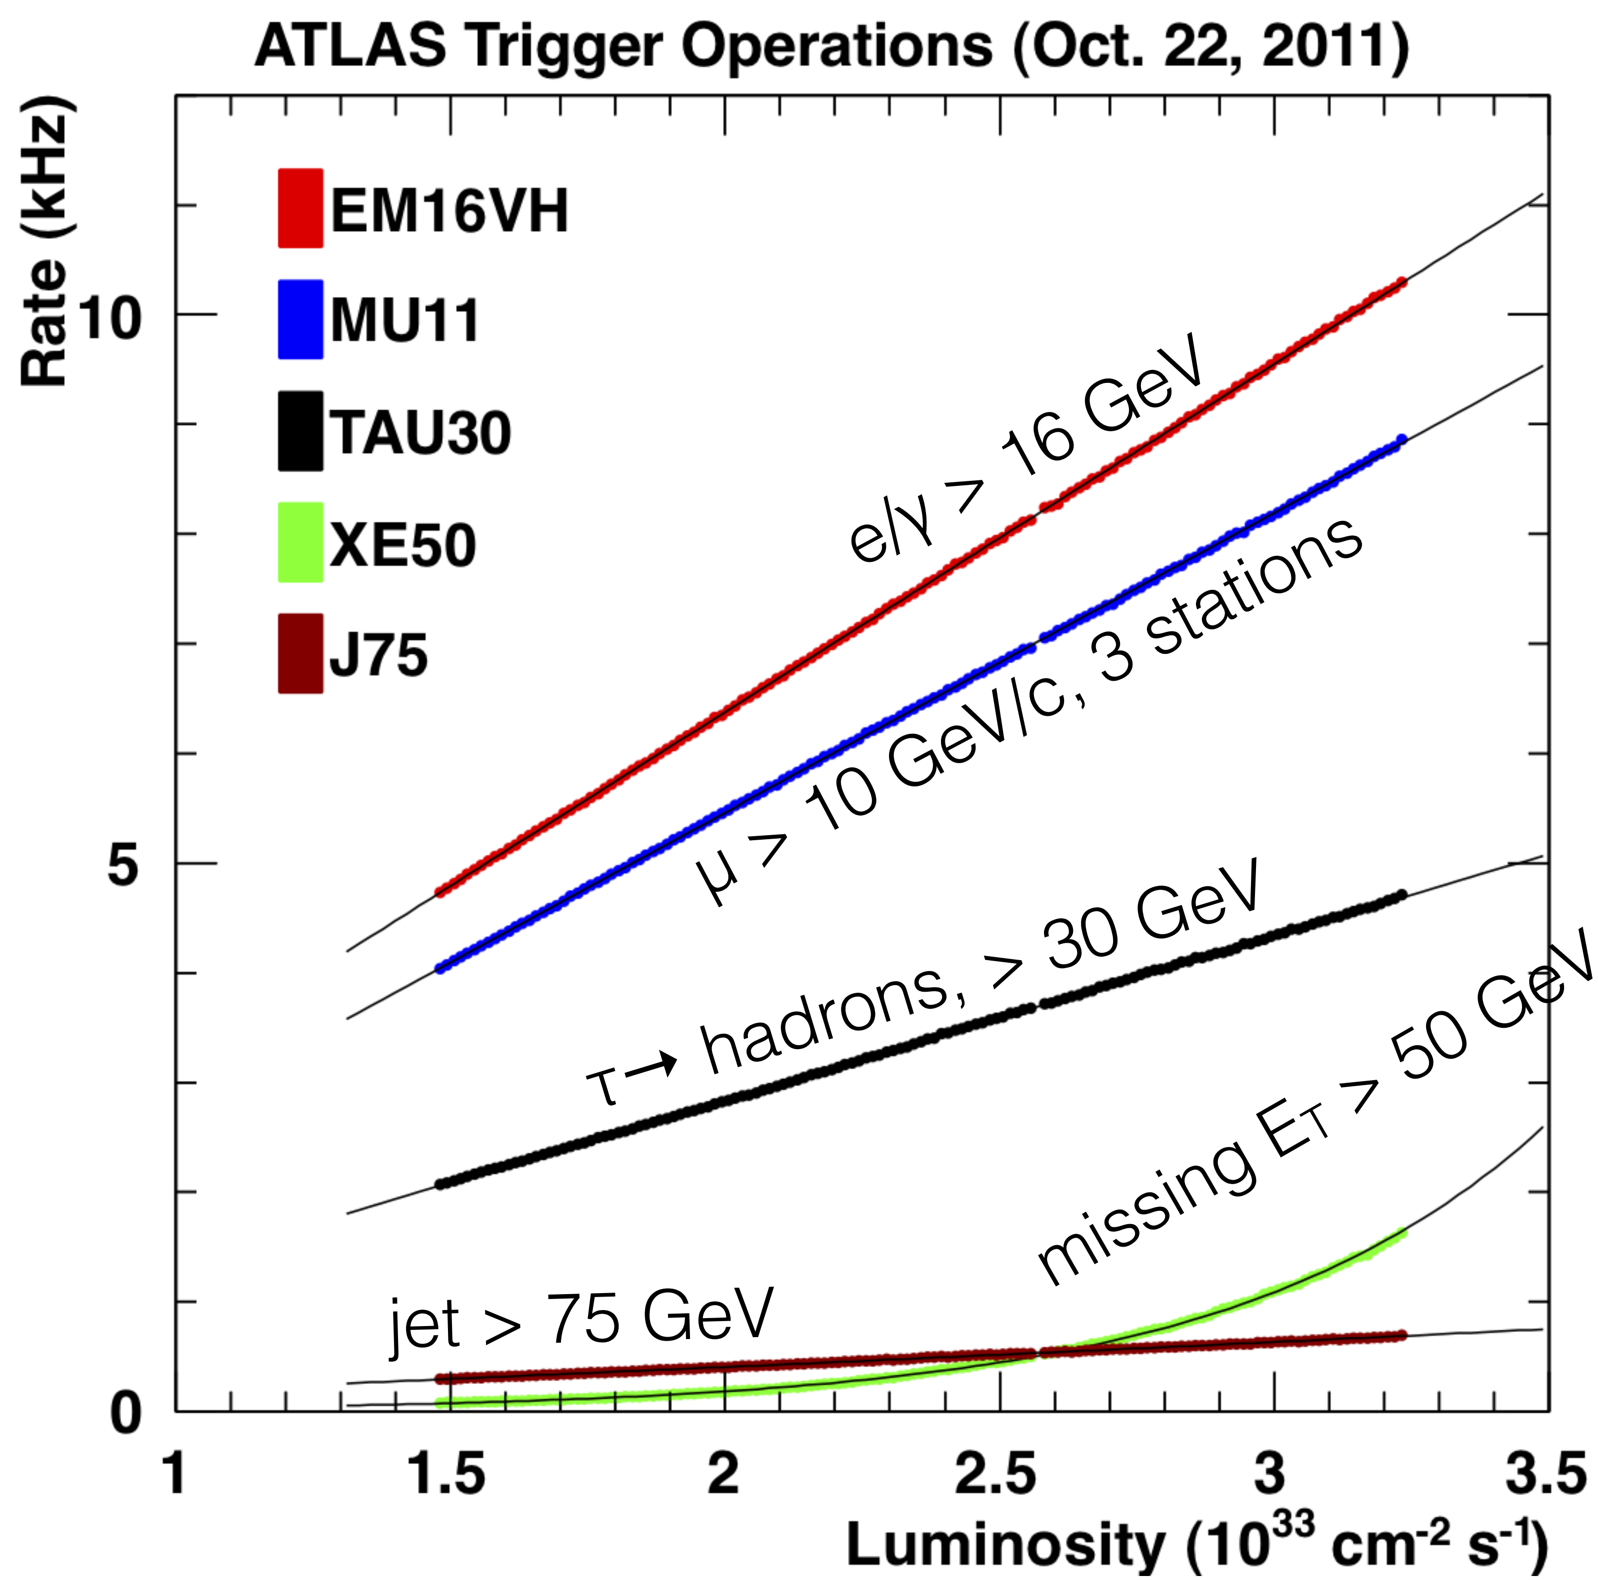

- impose high P_T requirement on identified electrons, muons, "jets"…
- beyond that: more complicated algorithms (eg. missing E_T) & coincidences

Identified High P_T objects

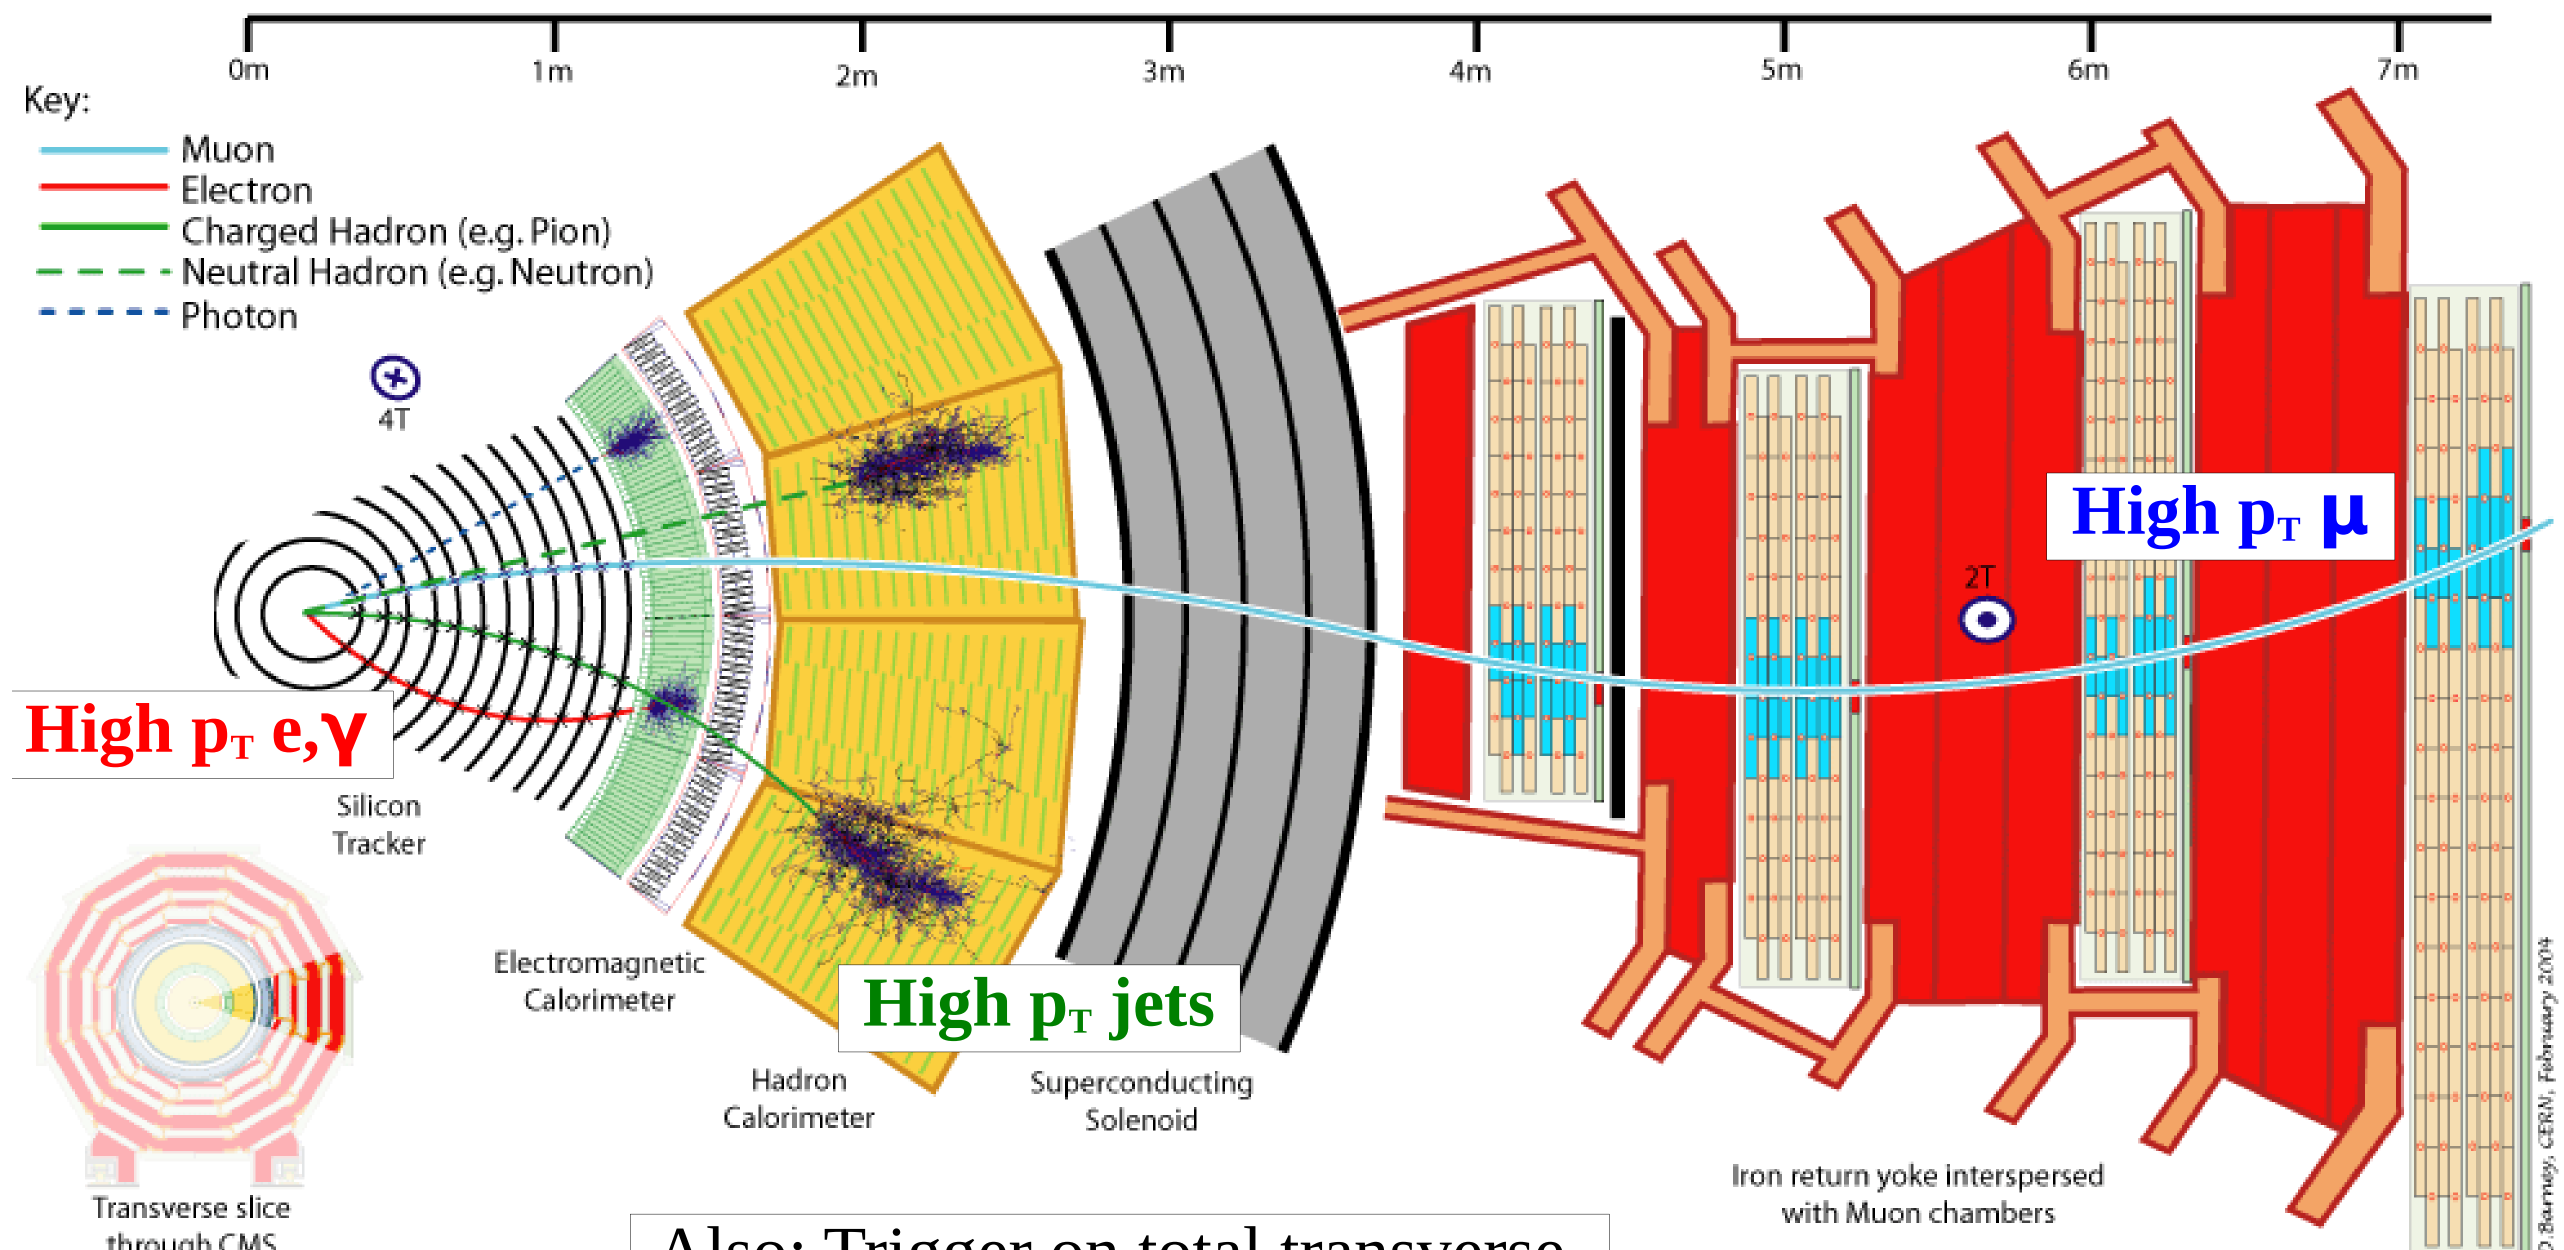

Also: Trigger on total transverse
(or missing transverse) energy

To track or not to track?

- **Pattern recognition much faster/easier**

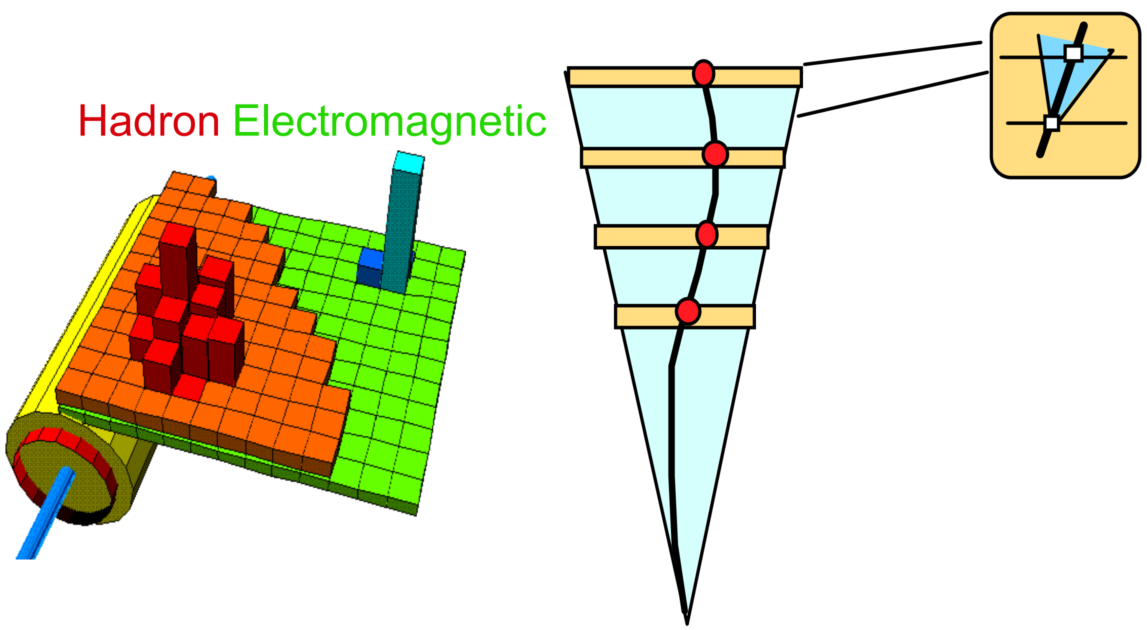

- **Simple algorithms**
- **Small amounts of data**
- **Local decisions**

- **Compare to tracker info**

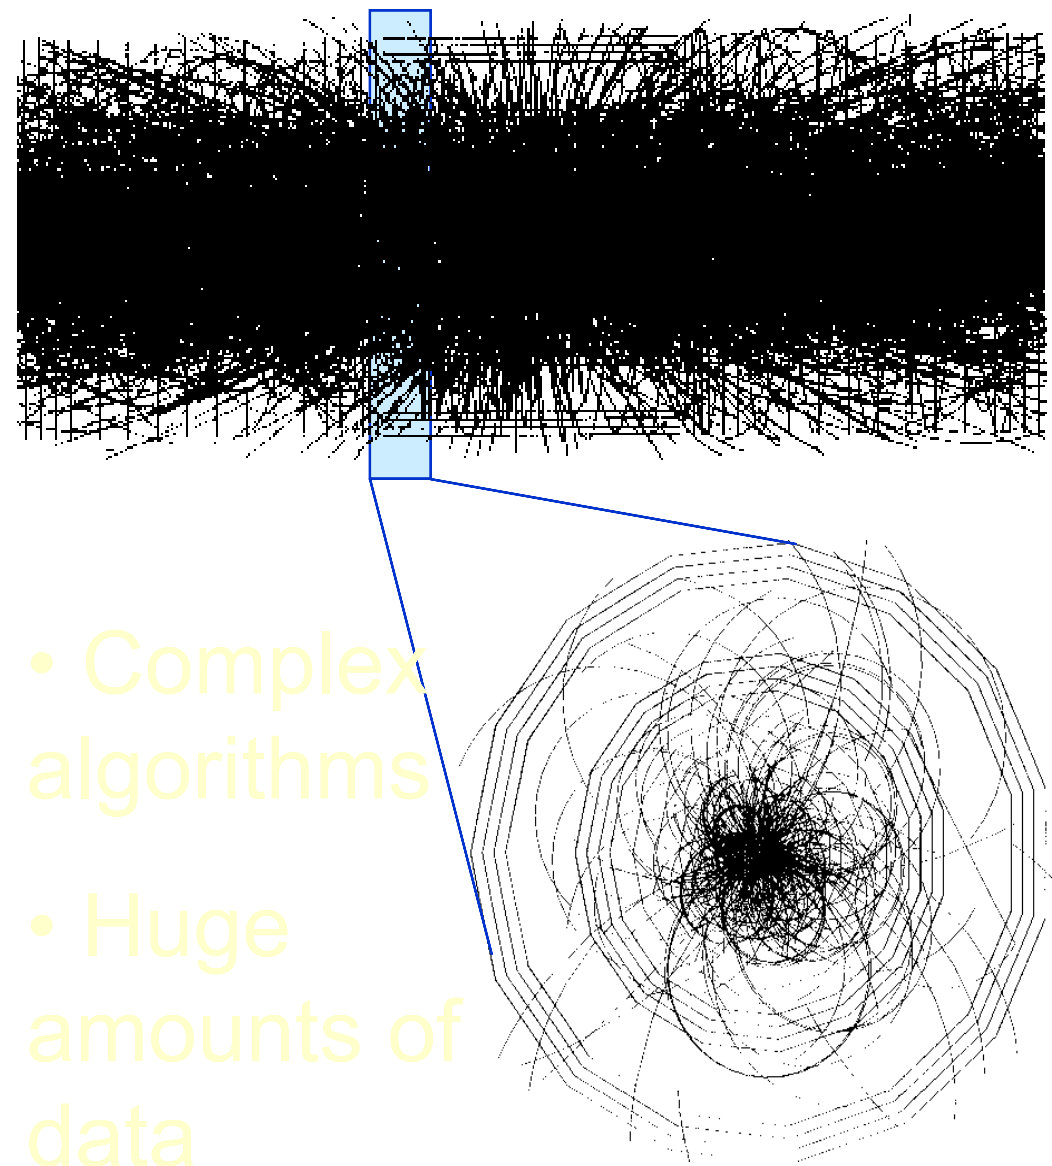

- Complex algorithms
- Huge amounts of data
- Need to link sub-detectors

Keep it local…

Use prompt data (calorimetry and muons) to identify:
High p_t electron, muon, jets, missing E_T

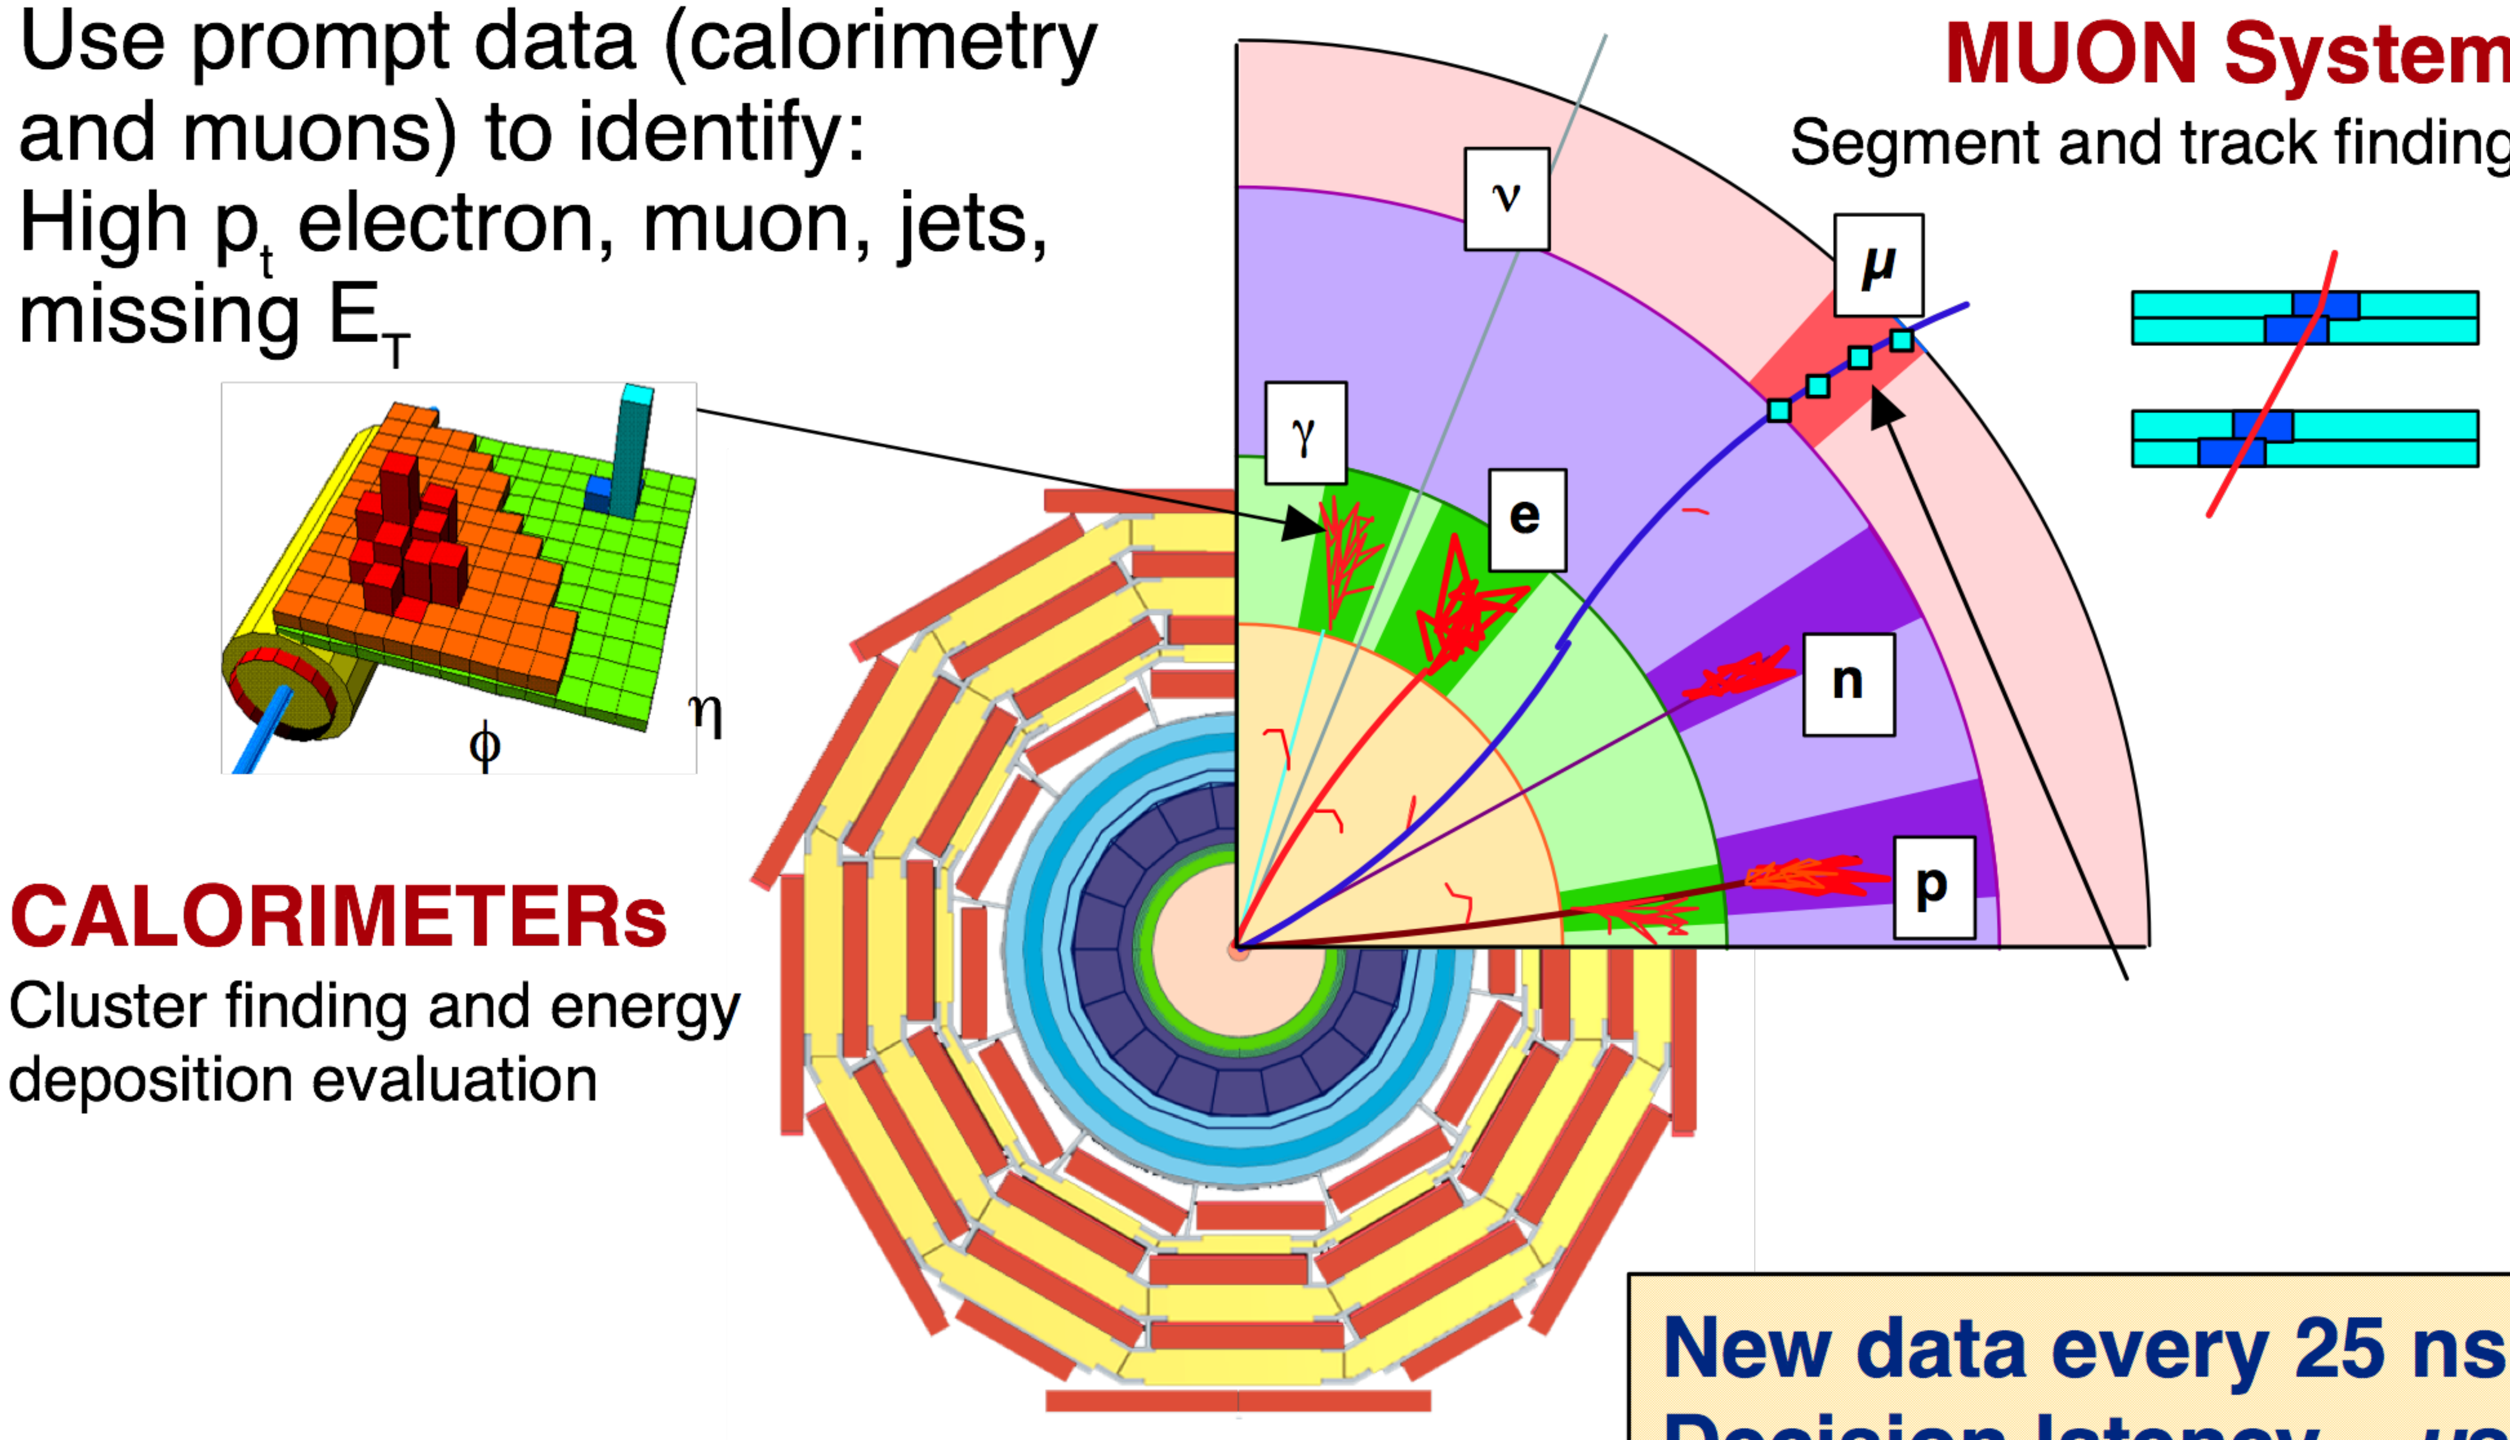

New data every 25 ns
Decision latency ~ μs

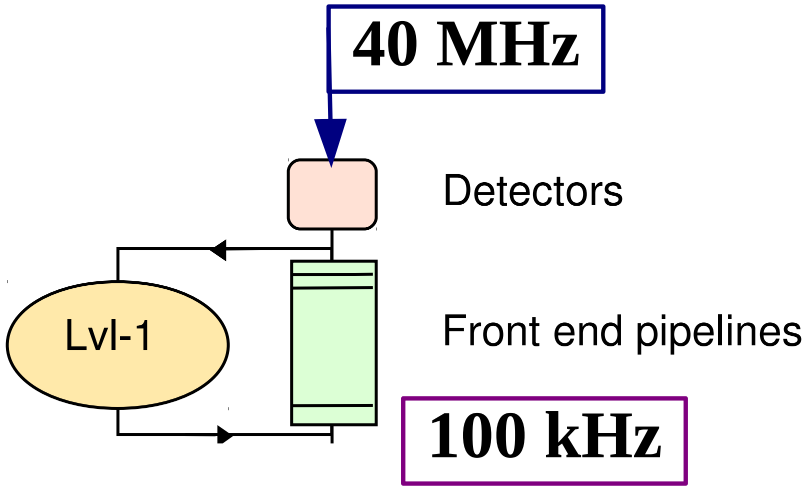

From local objects to global decisions

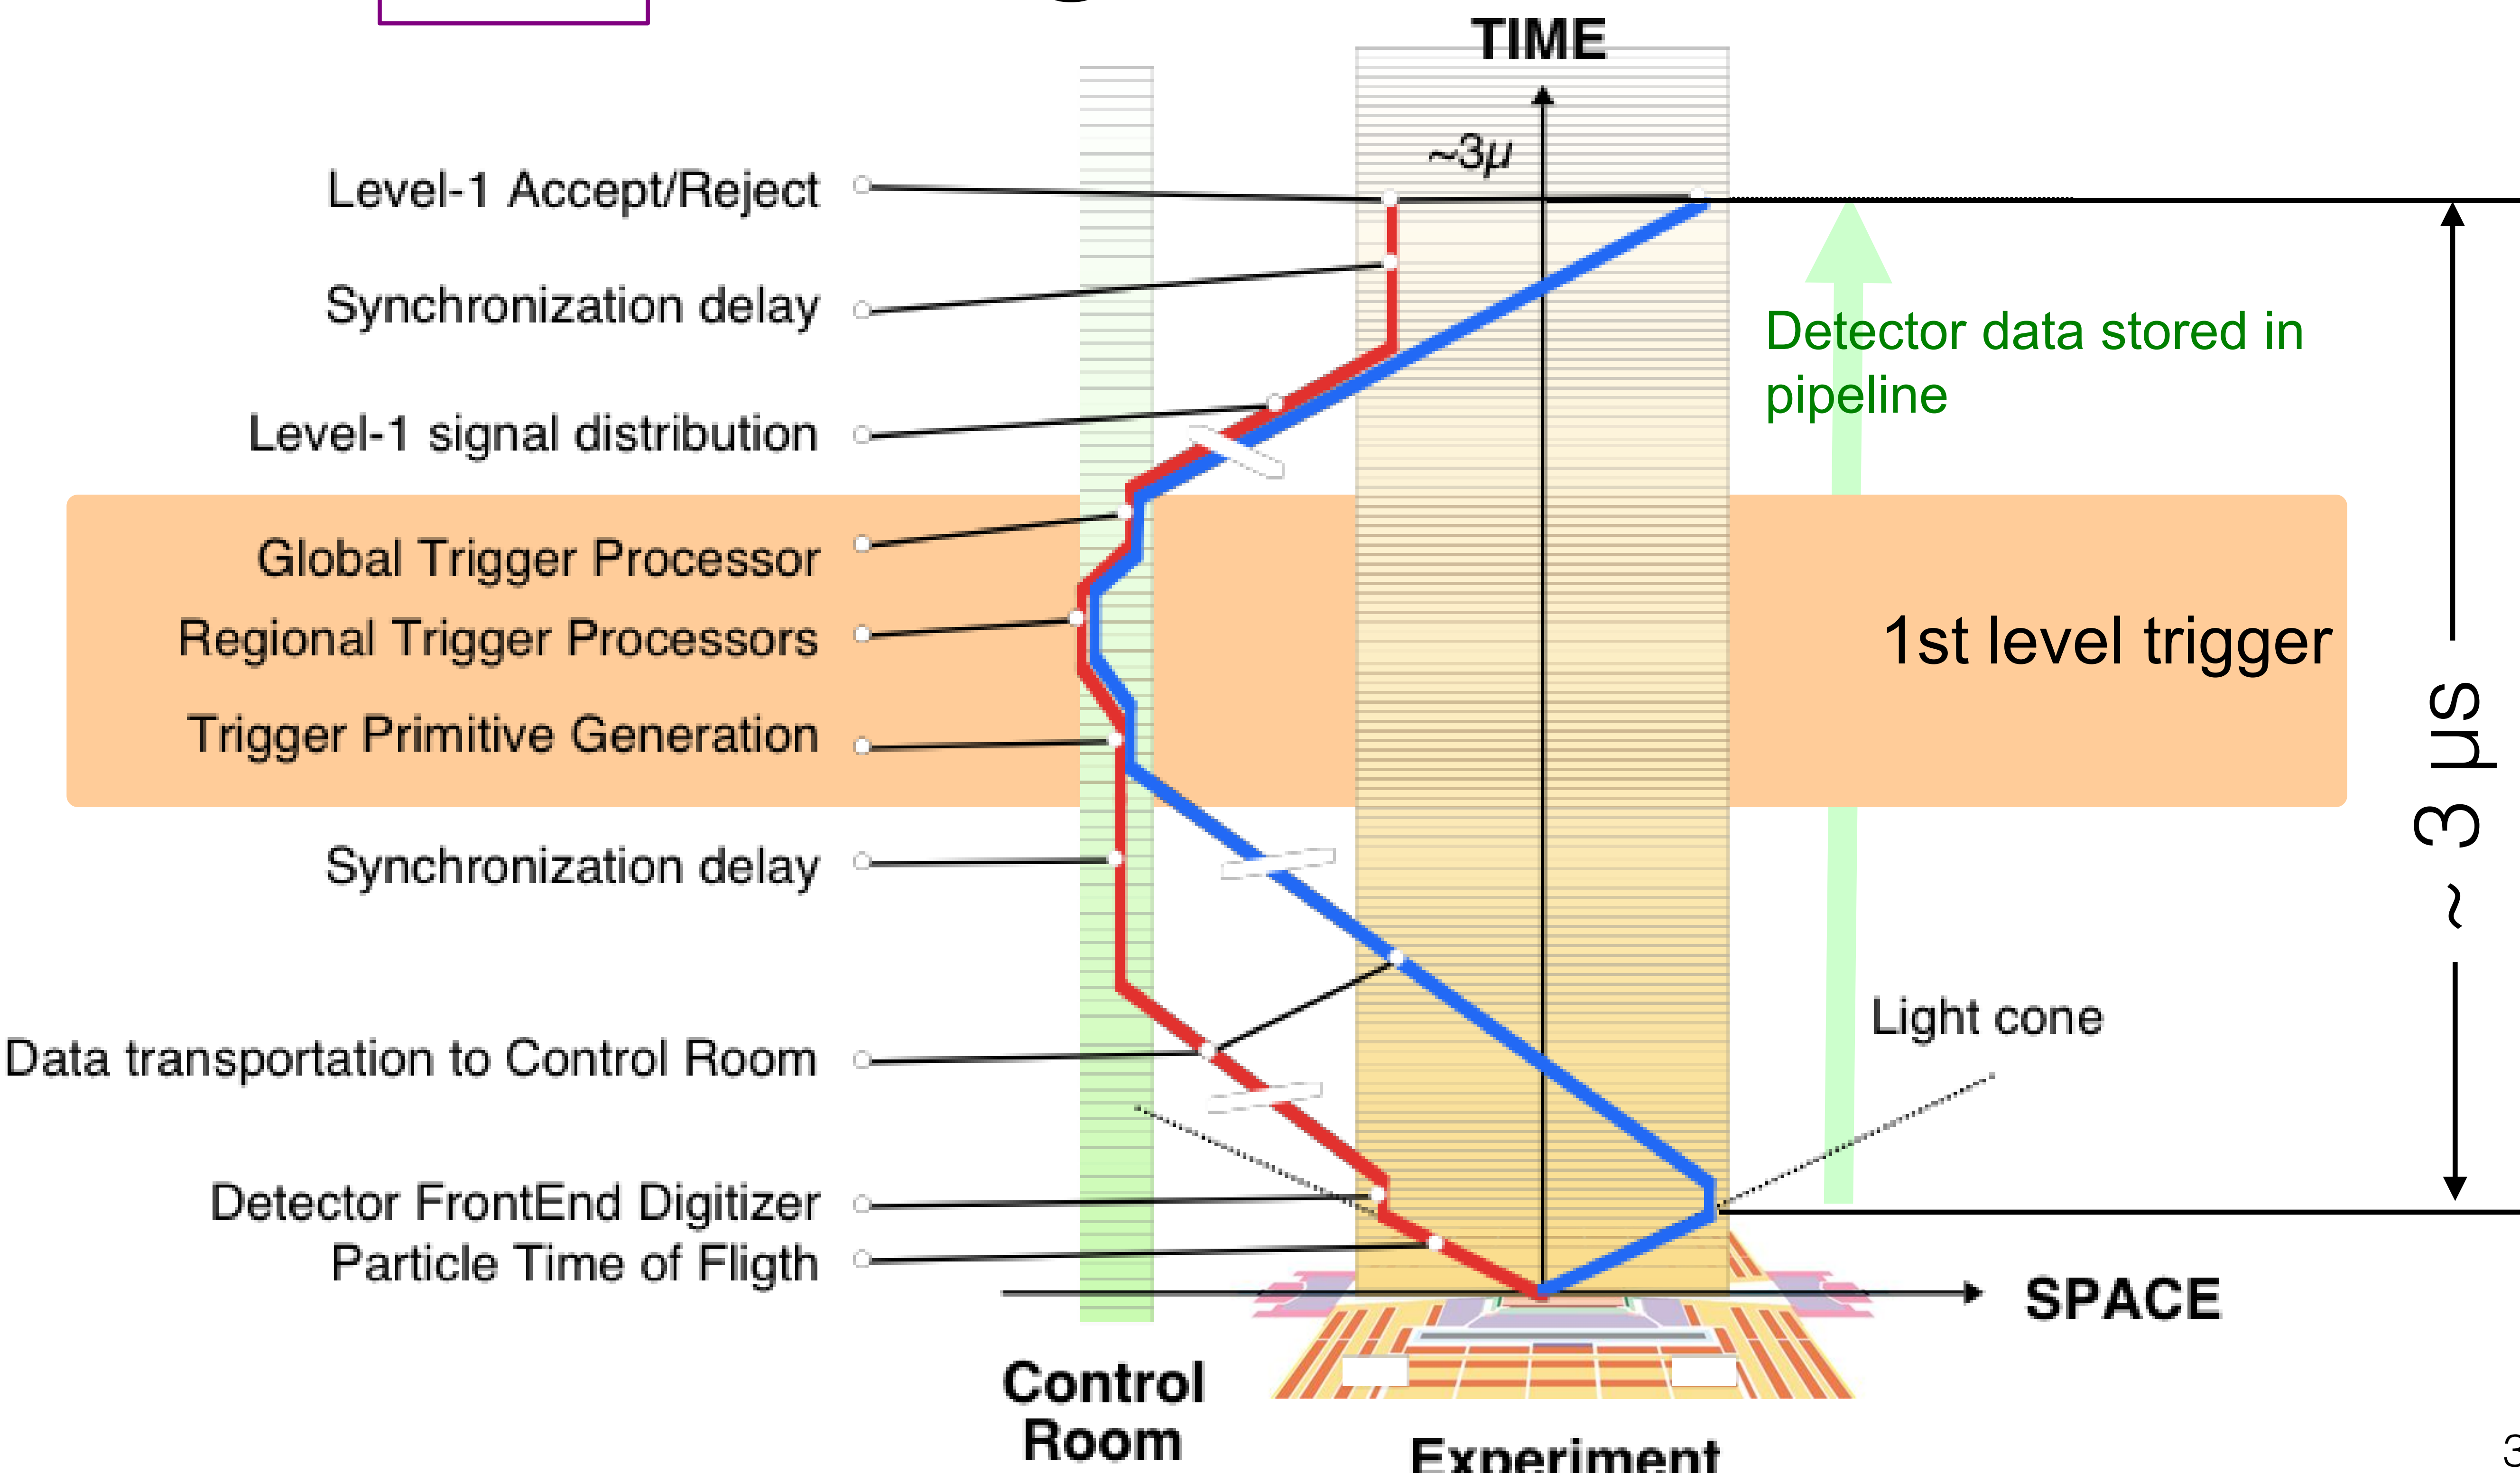

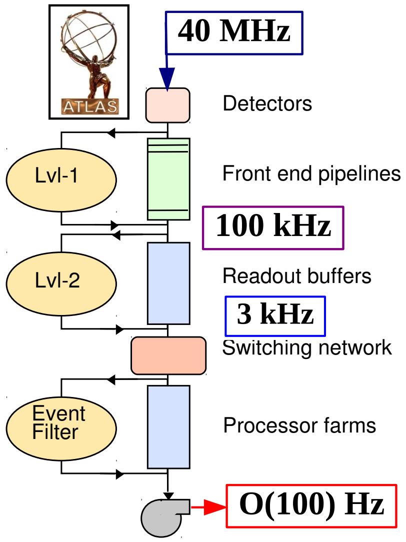

Life After L1

- Where are we:
  - Atlas / CMS: 75 - 100 kHz; event size 1 - 2 MB
  - LHCb: 1 MHz; event size 60 KB
  - Alice: O(kHz) & O(GB)
- Ideally: run ‘full-blown’ reconstruction and selection
  - this takes O(s) / event —> your Intel sales rep will be extremely happy to sell you 0.1M-1M CPUs
- Reality: “staged” reconstruction and selection
  - start by looking at part of the event data
  - Atlas: region of interest from L1: < 40 ms / event, reduce by factor ~30
  - LHCb: 1st stage HLT partial tracking: < 10 ms / event, reduce by factor ~15
- Then do the “expensive” parts

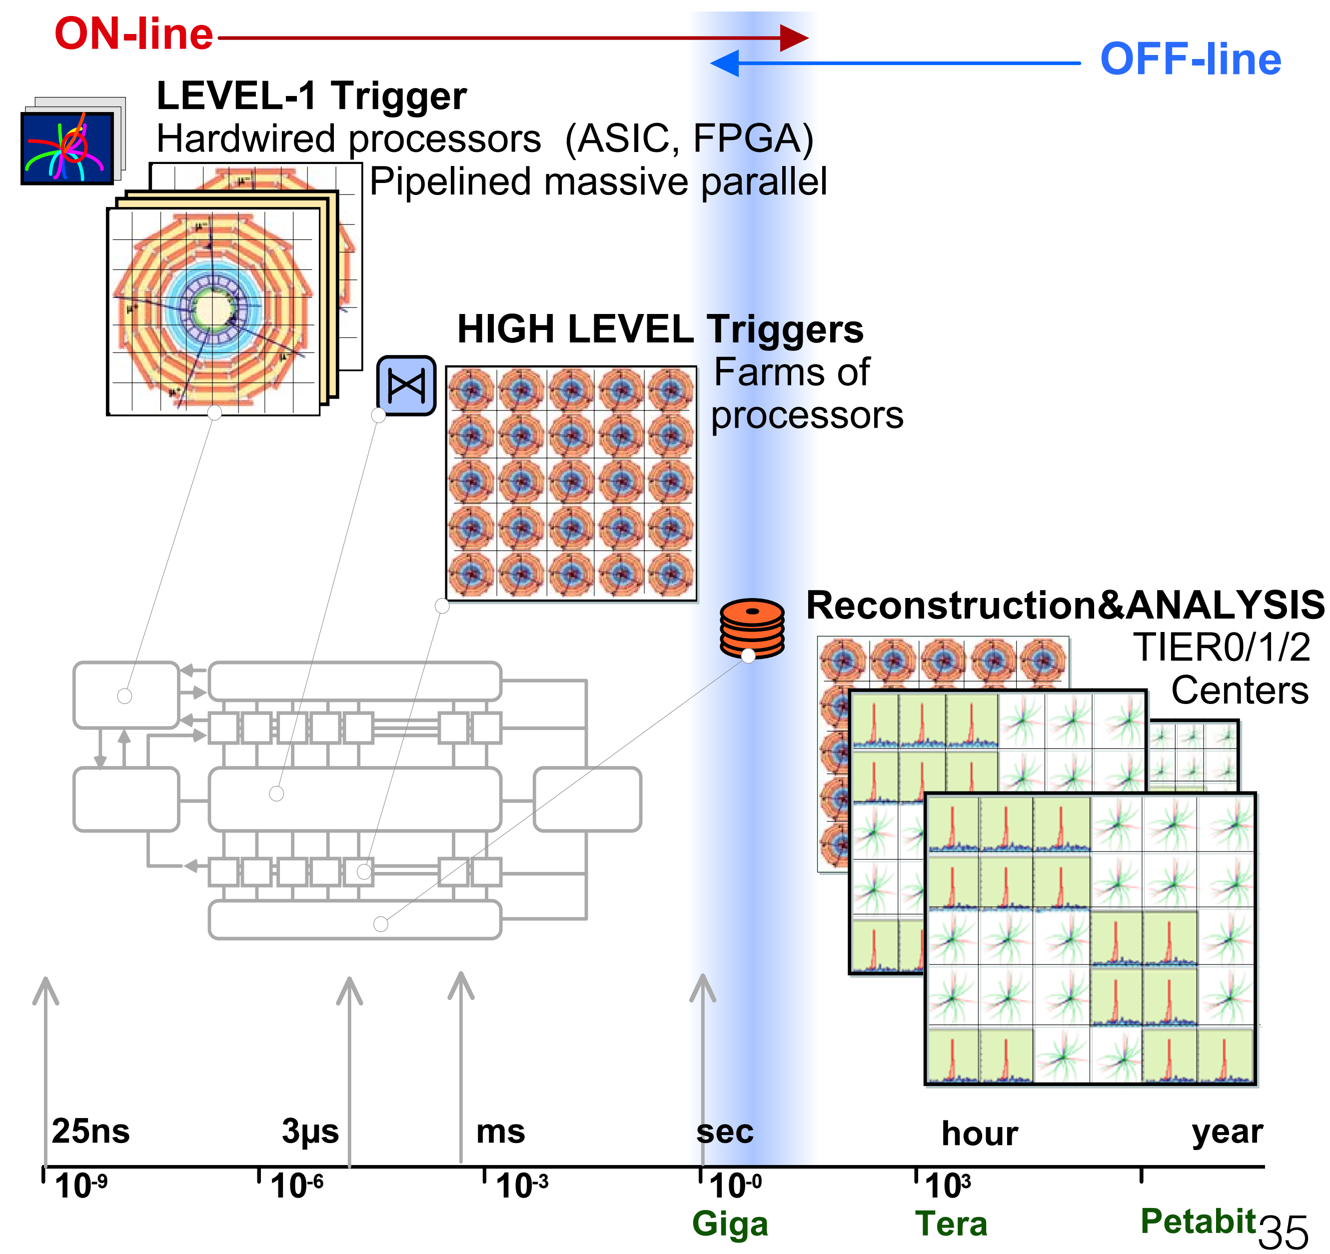

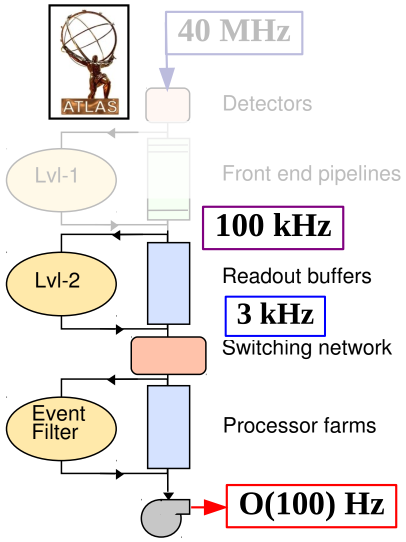

High Level Trigger

- Software running on "of the shelf" compute servers
- Atlas: 17K CPUs split between LVL2 and Event Filter (run1)
- CMS: 13K CPU cores (run 1)
- LHCb: 29K CPU cores (run 1)
- "same" software (framework) as offline, but "tuned" for the online case
- Take "a few hundred" different and *independent* decisions

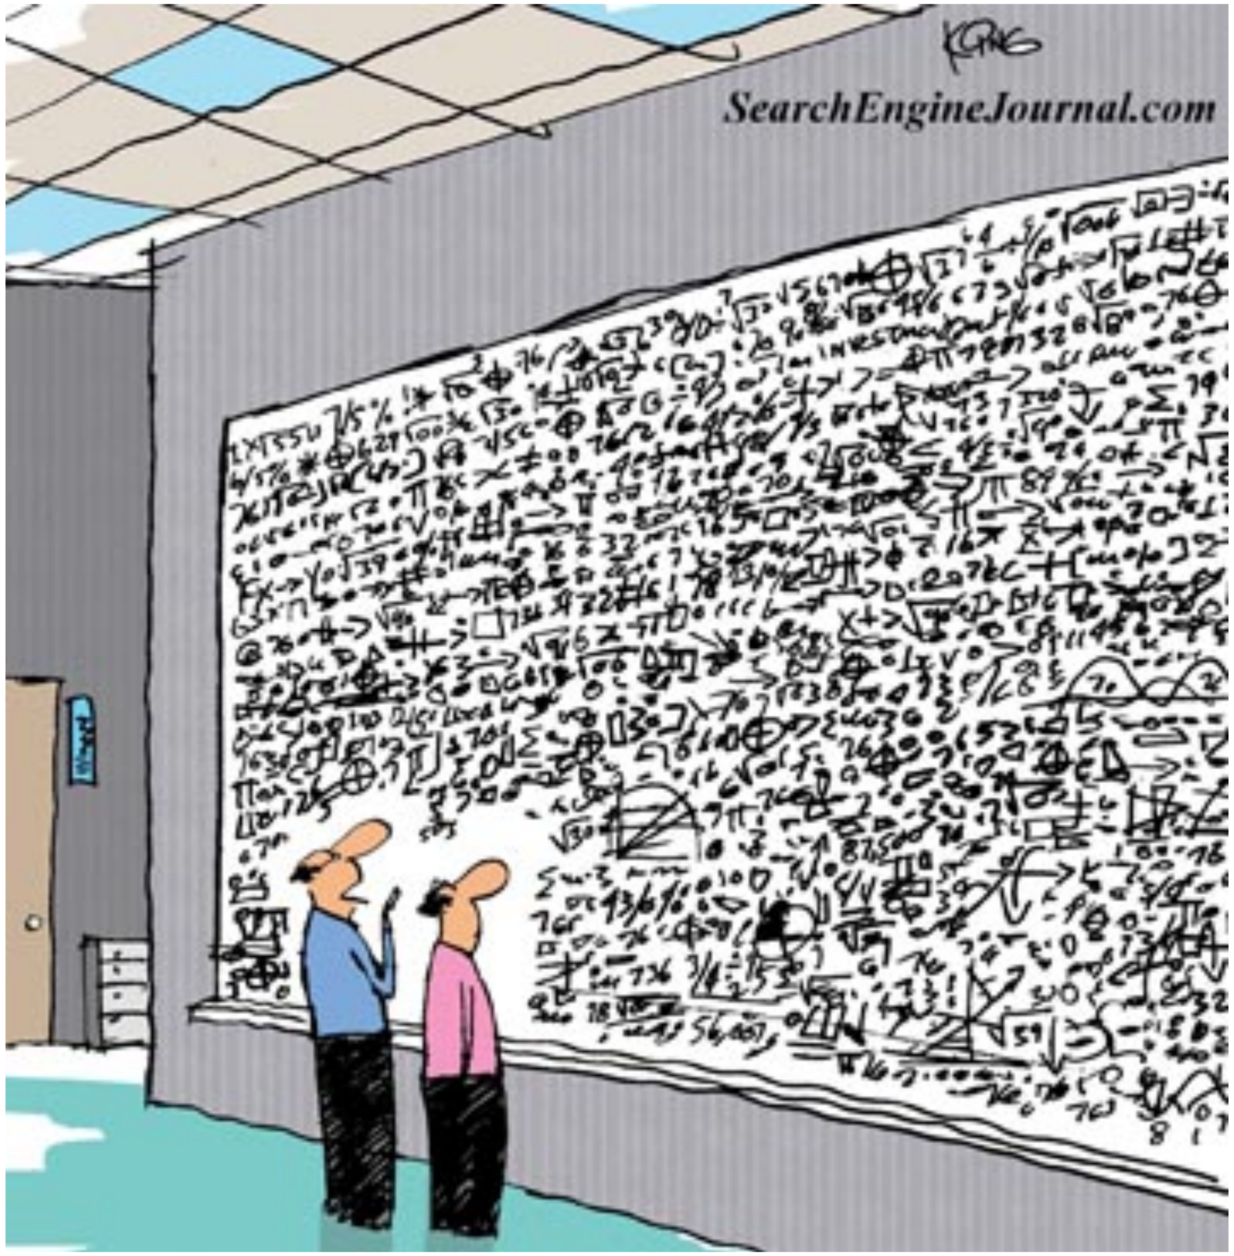

And that, in simple terms, is what we do in the High Level Trigger

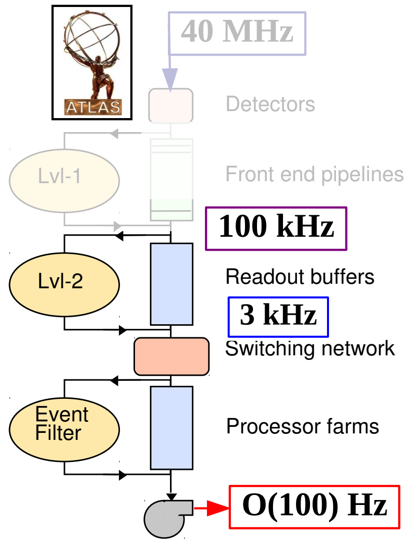

High Level Trigger

- Software running on “of the shelf” compute servers
- Atlas: 17K CPUs split between LVL2 and Event Filter (run1)
- CMS: 13K CPU cores (run 1)
- LHCb: 29K CPU cores (run 1)
- “same” software (framework) as offline, but “tuned” for the online case
- Take “a few hundred” different and *independent* decisions

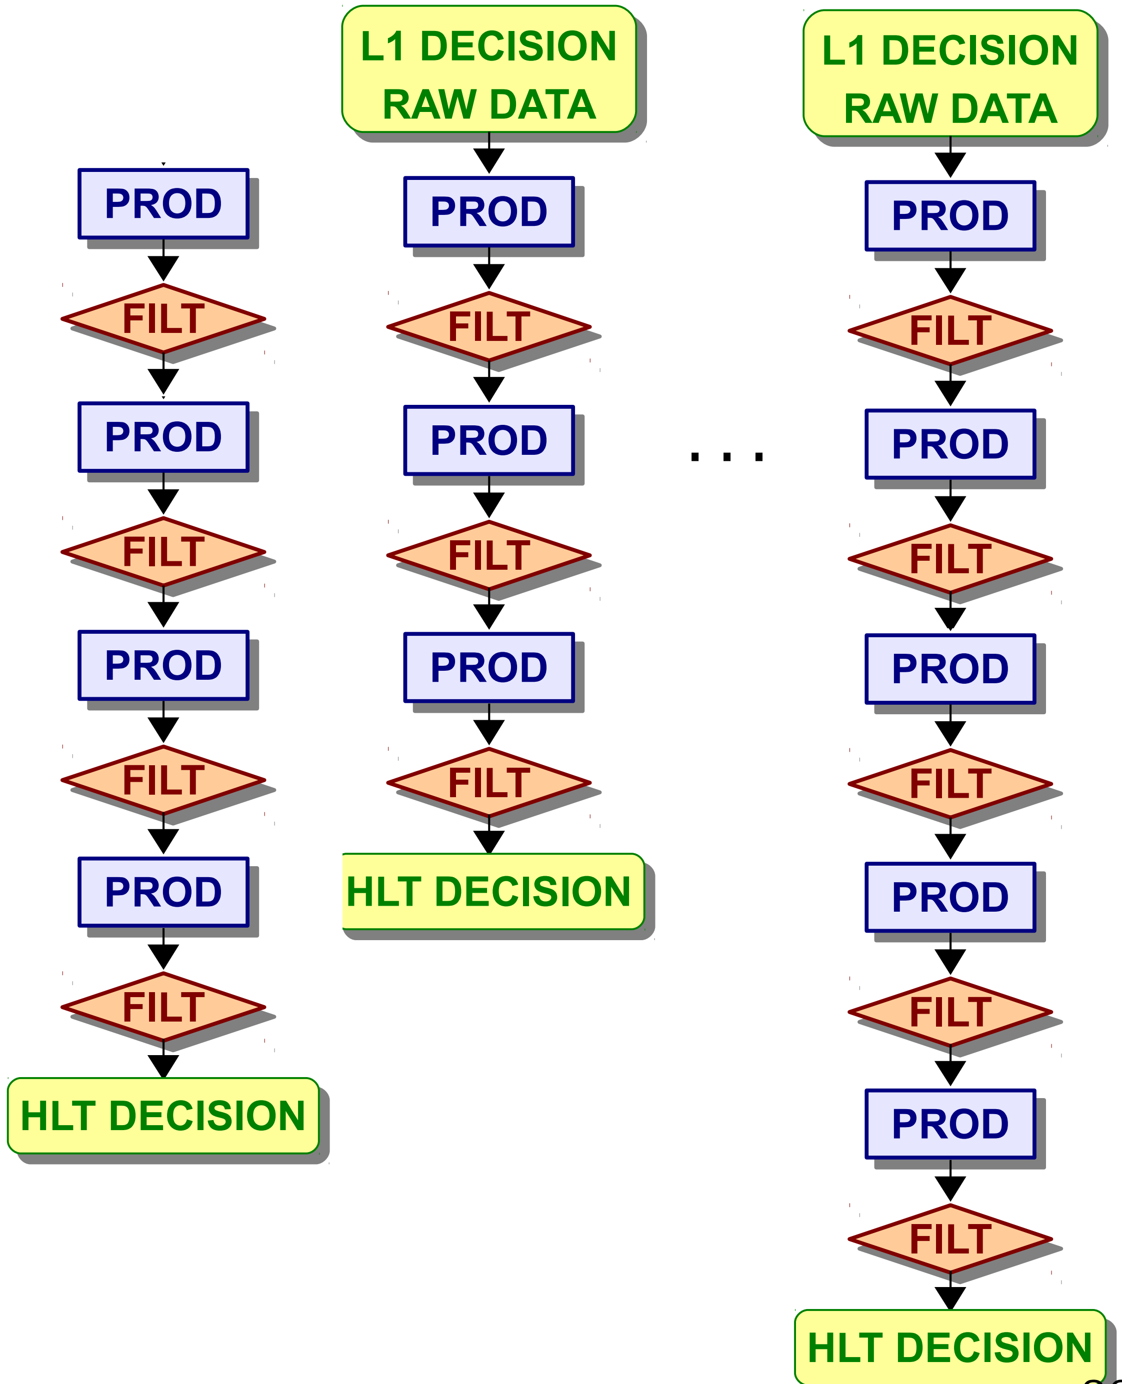

HLT “paths”/“chains”/“lines”

- HLT accept is an ‘OR’ of 100-1000 different decisions…
- paths are (logically) independent
- but may share common parts
- each path is typically alternates reconstruction & filtering
  - typically do ‘fast’/‘easy’ things first
  - can use the L1 information as a ‘seed’
- remember: most events are rejected
  - “early accept” is not only a bad idea, it doesn’t help either…
  - early abort, within a path is a good idea

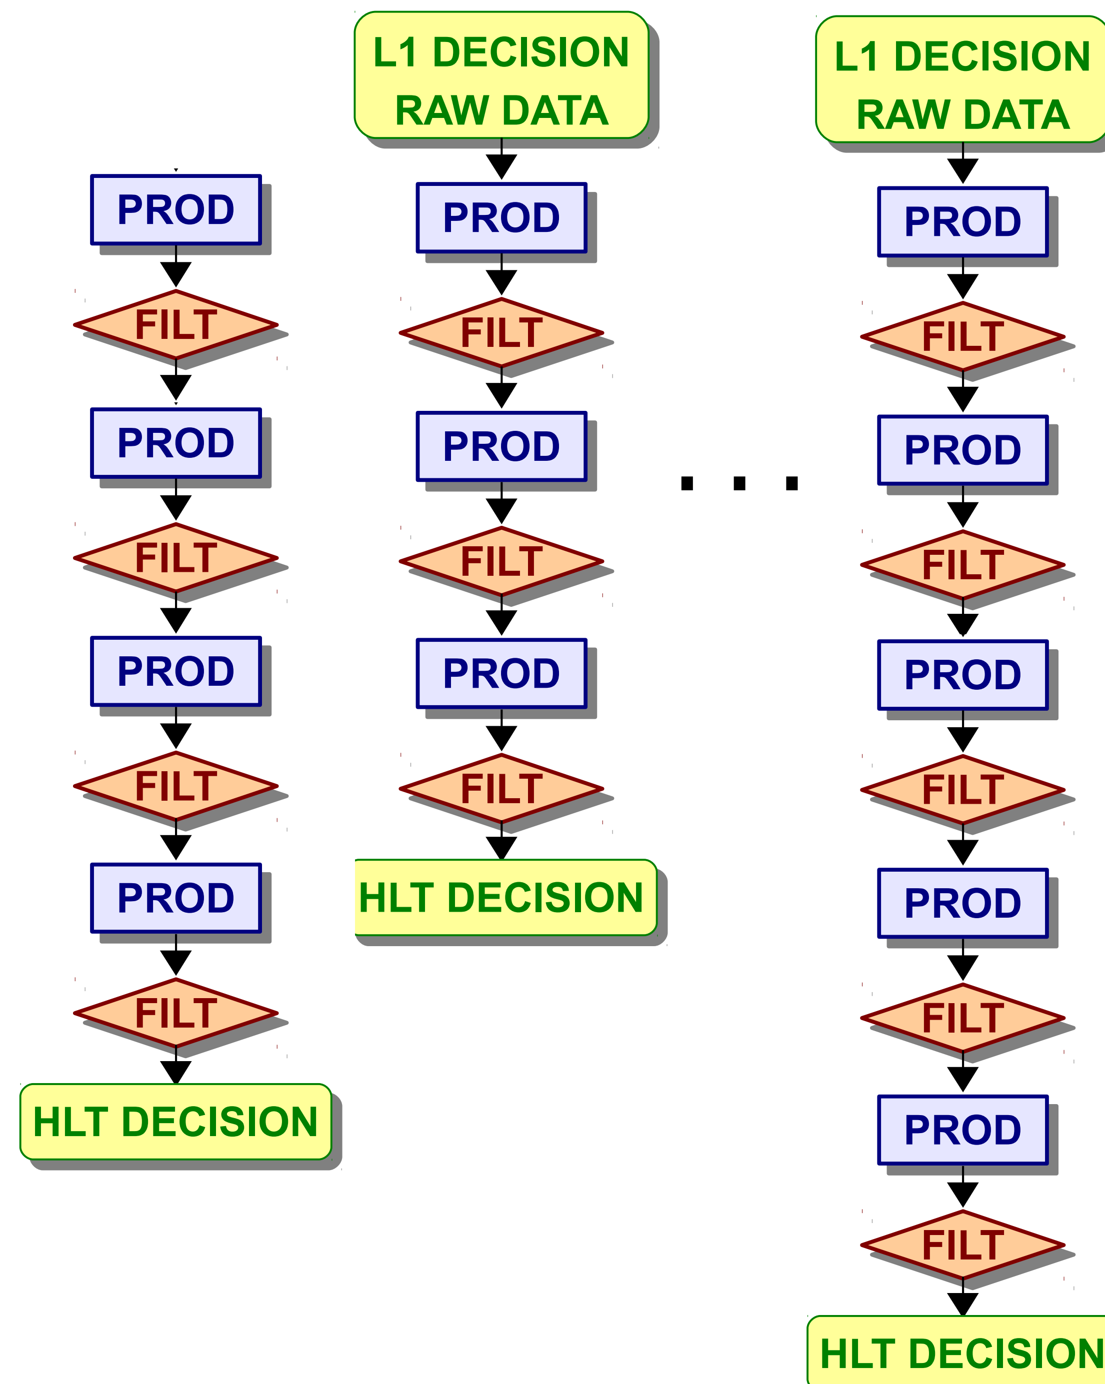

CMS HLT Muon Workflow

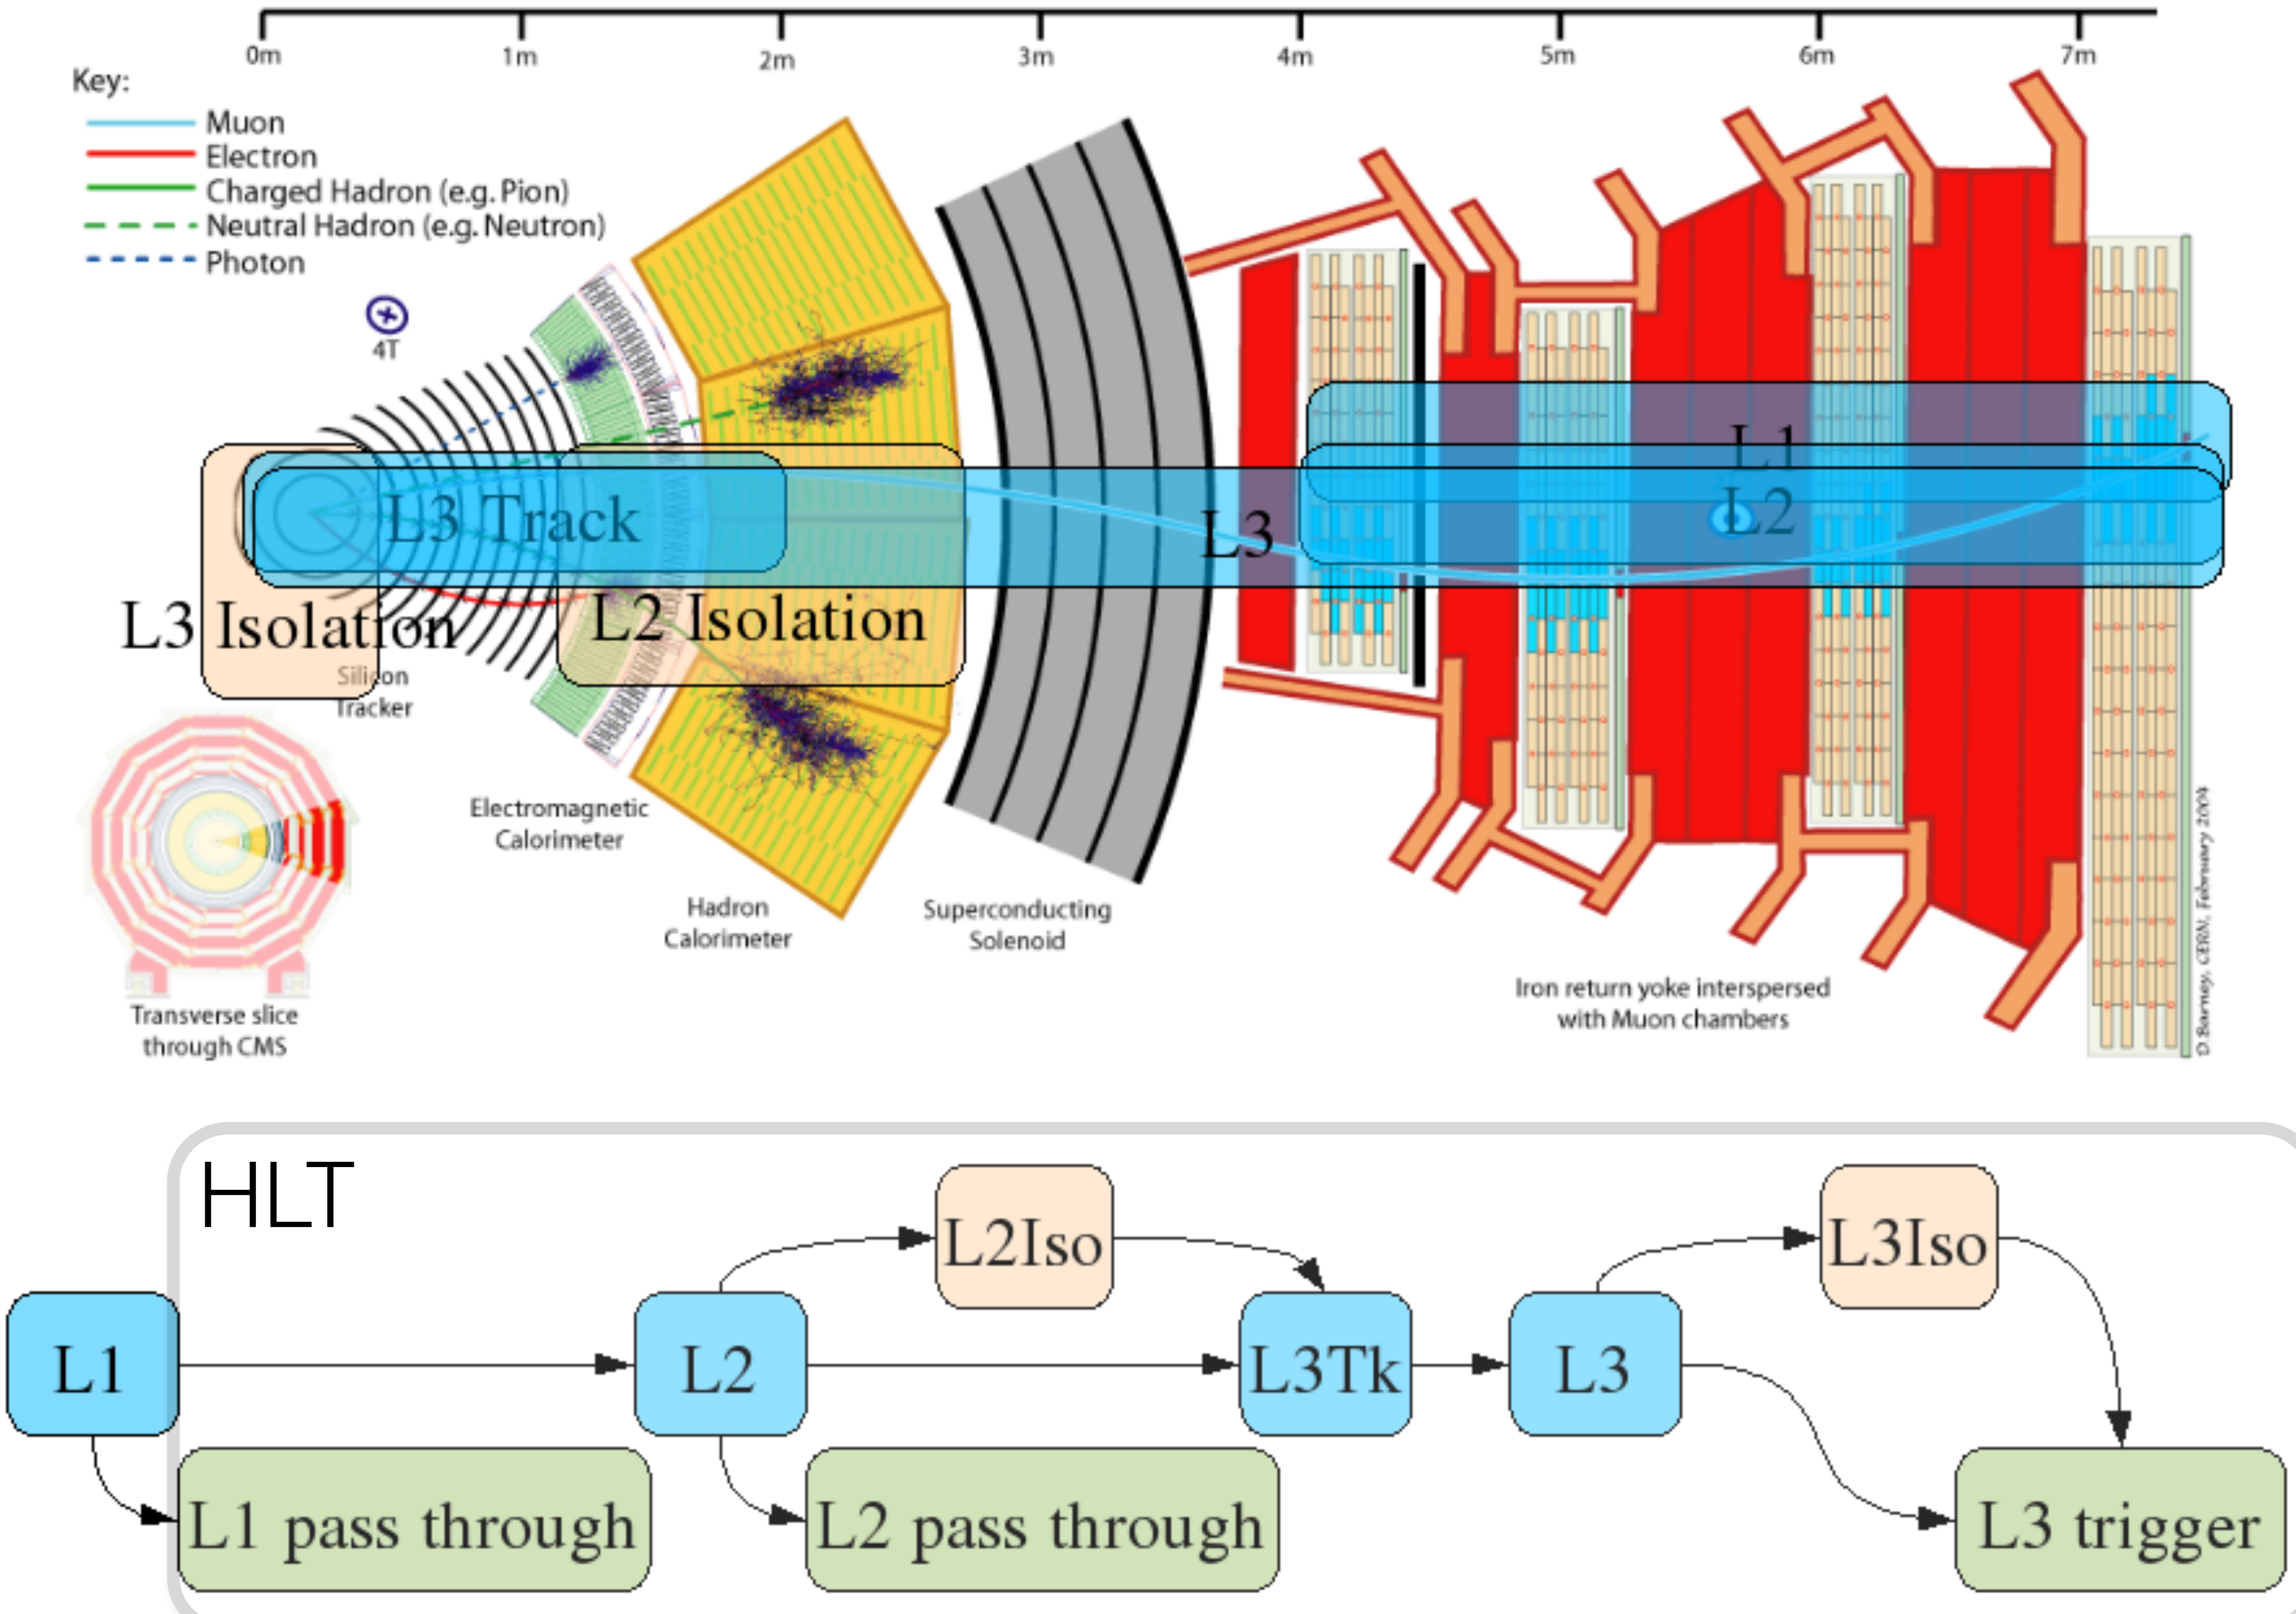

Trigger “menus”

(Unprescaled) Object	Trigger Threshold (GeV)	Rate (Hz)	Physics
Single Muon	40	21	Searches
Single Isolated muon	24	43	Standard Model
Double muon	(17, 8) [13, 8 for parked data]	20 [30]	Standard Model / Higgs
Single Electron	80	8	Searches
Single Isolated Electron	27	59	Standard Model
Double Electron	(17, 8)	8	Standard Model / Higgs
Single Photon	150	5	Searches
Double Photon	(36, 22)	7	Higgs
Muon + Ele x-trigger	(17, 8), (5, 5, 8), (8, 8, 8)	3	Standard Model / Higgs
Single PFJet	320	9	Standard Model
QuadJet	80 [50 for parked data]	8[100]	Standard Model /Searches
Six Jet	(6 x 45), (4 x 60, 2 x 20)	3	Searches
MET	120	4	Searches
HT	750	6	Searches

Object breakdown for LHC Run 1 instantaneous luminosities of (nearly 7 x 10^{33} $s^{-1}cm^{-2}$ at the start of a fill)

Menu details

- Primary chains:
    - main source of signal
- Backup chains:
    - increased threshold — in case of ‘trouble’ (malfunction, or increased luminosity ;-)
- As luminosity decays, capacity can be used — add additional chains
    - eg. lower threshold of primary chains supporting chains
- monitoring chains:
    - used for eg. data-quality
- supporting chains:
    - extract eg. trigger efficiencies, background control regions, …

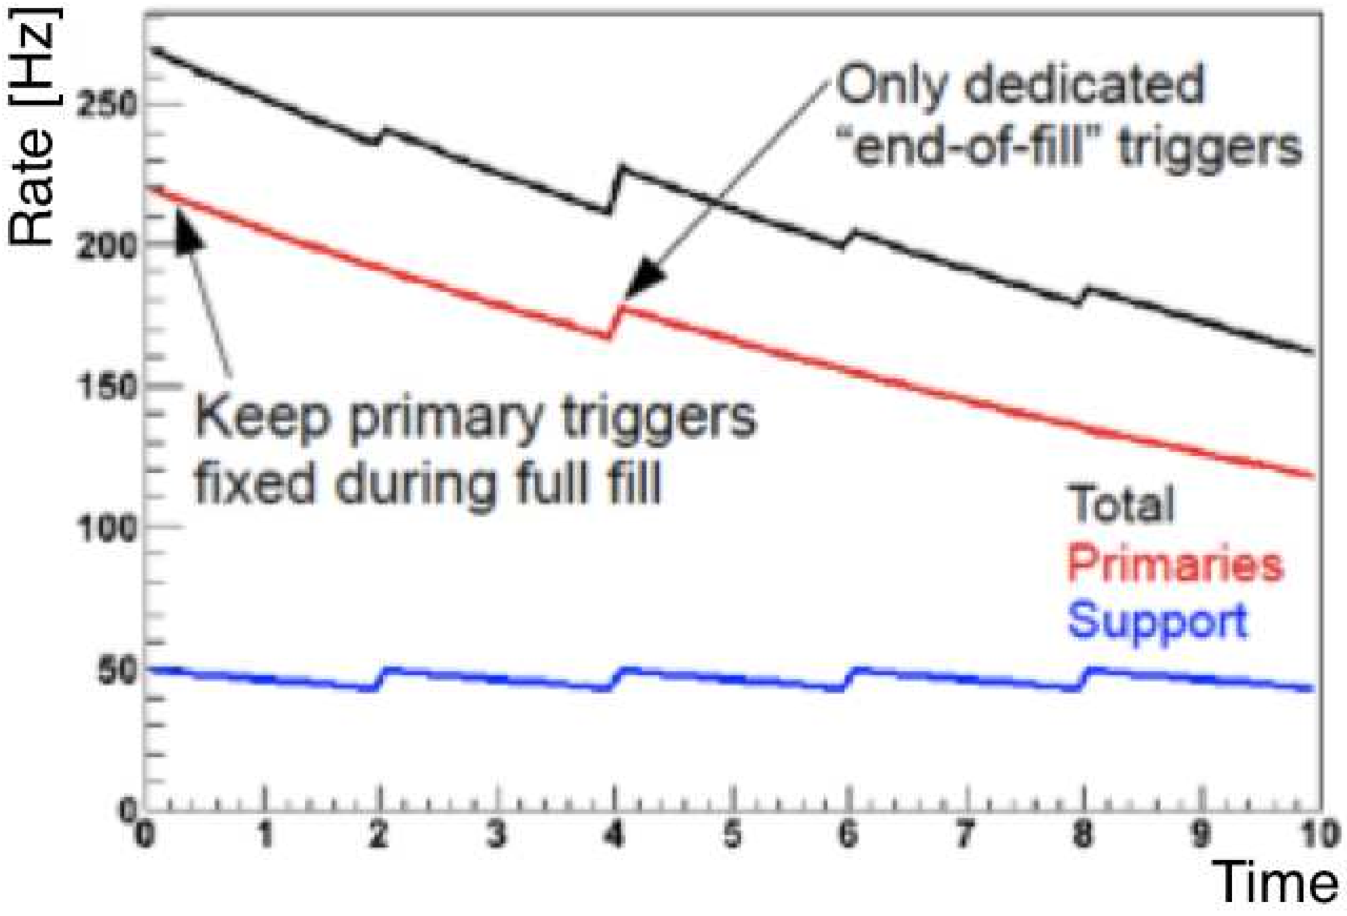

Rates changed
by ‘prescale’ sets

Menu 'evolution'

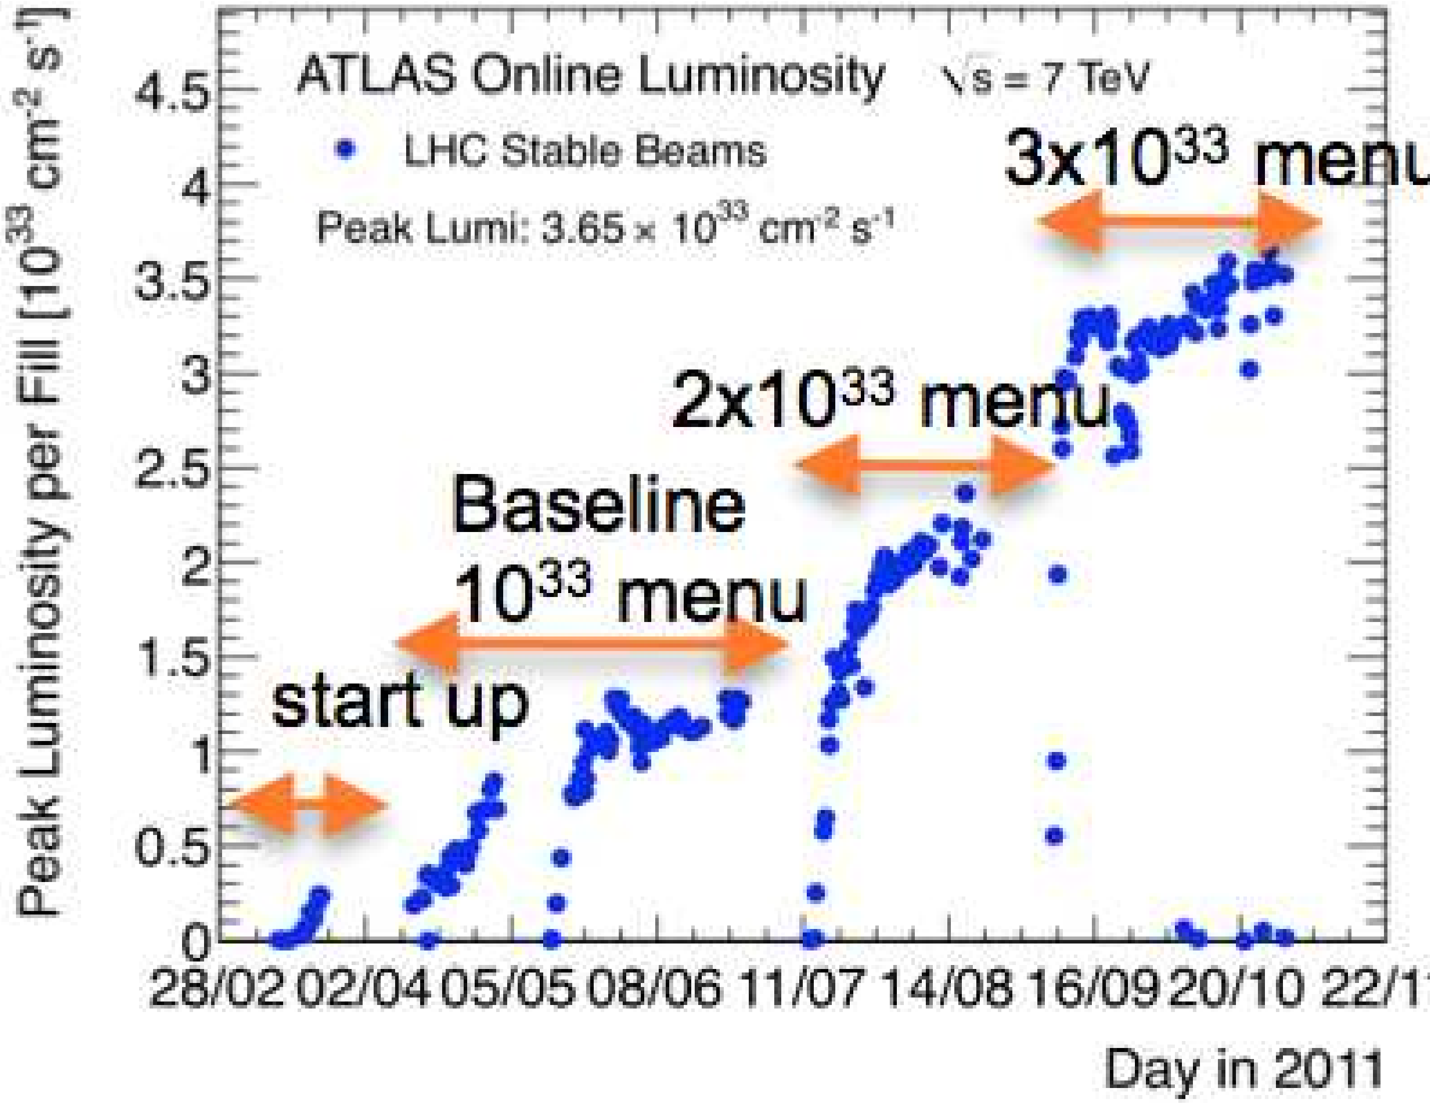

- Output bandwidth is a finite, shared resource
- How to decide what gets how much?
- Depends on the 'physics goals' of the experiment…

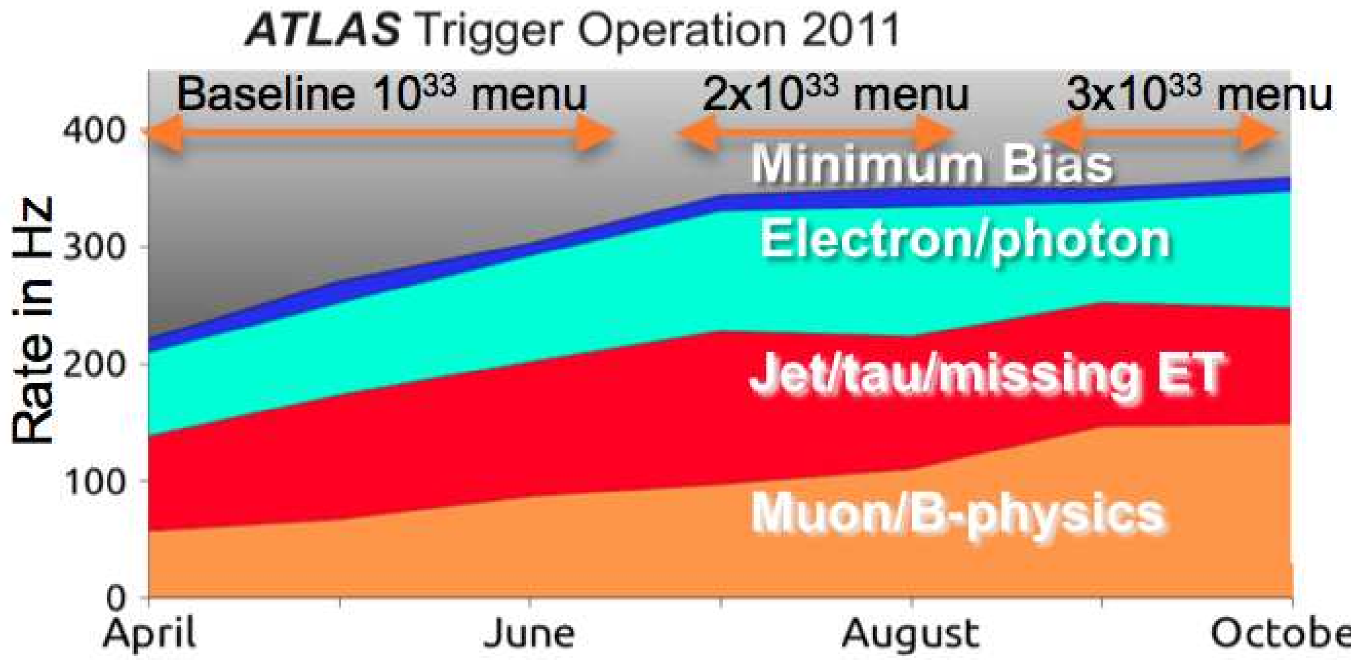

Interaction with Analysis

analysis preparation:

- setup/ optimize a trigger for your physics signal
  - define a trigger strategy (based on the available resources)
  - convert to trigger chain (already existing?)
  - determine rates and efficiencies from MC
- define a monitoring strategy
  - define trigger chain to be used for monitoring of your physics trigger (efficiency from data)
  - rates of the monitoring trigger (pre-scales?)
- integrate this in the overall trigger menu (done by Trigger Menu Coordination for online running)

not OK

threshold?
more exclusive?
pre-scaling?
more conditions?

OK

- use the trigger online (take data)
- monitor trigger quality
- determine trigger eff. (from data)
- correct your measurement

LHCb: 2011-2012 trigger

Three stage architecture:

- Level 0 (L0) — near detector hardware, fixed 4 μs latency
- Higher Level Trigger (HLT1 and HLT2) — software running on 29,000 cores

JINST 8 (2013) P04022, and arXiv:1310.8544

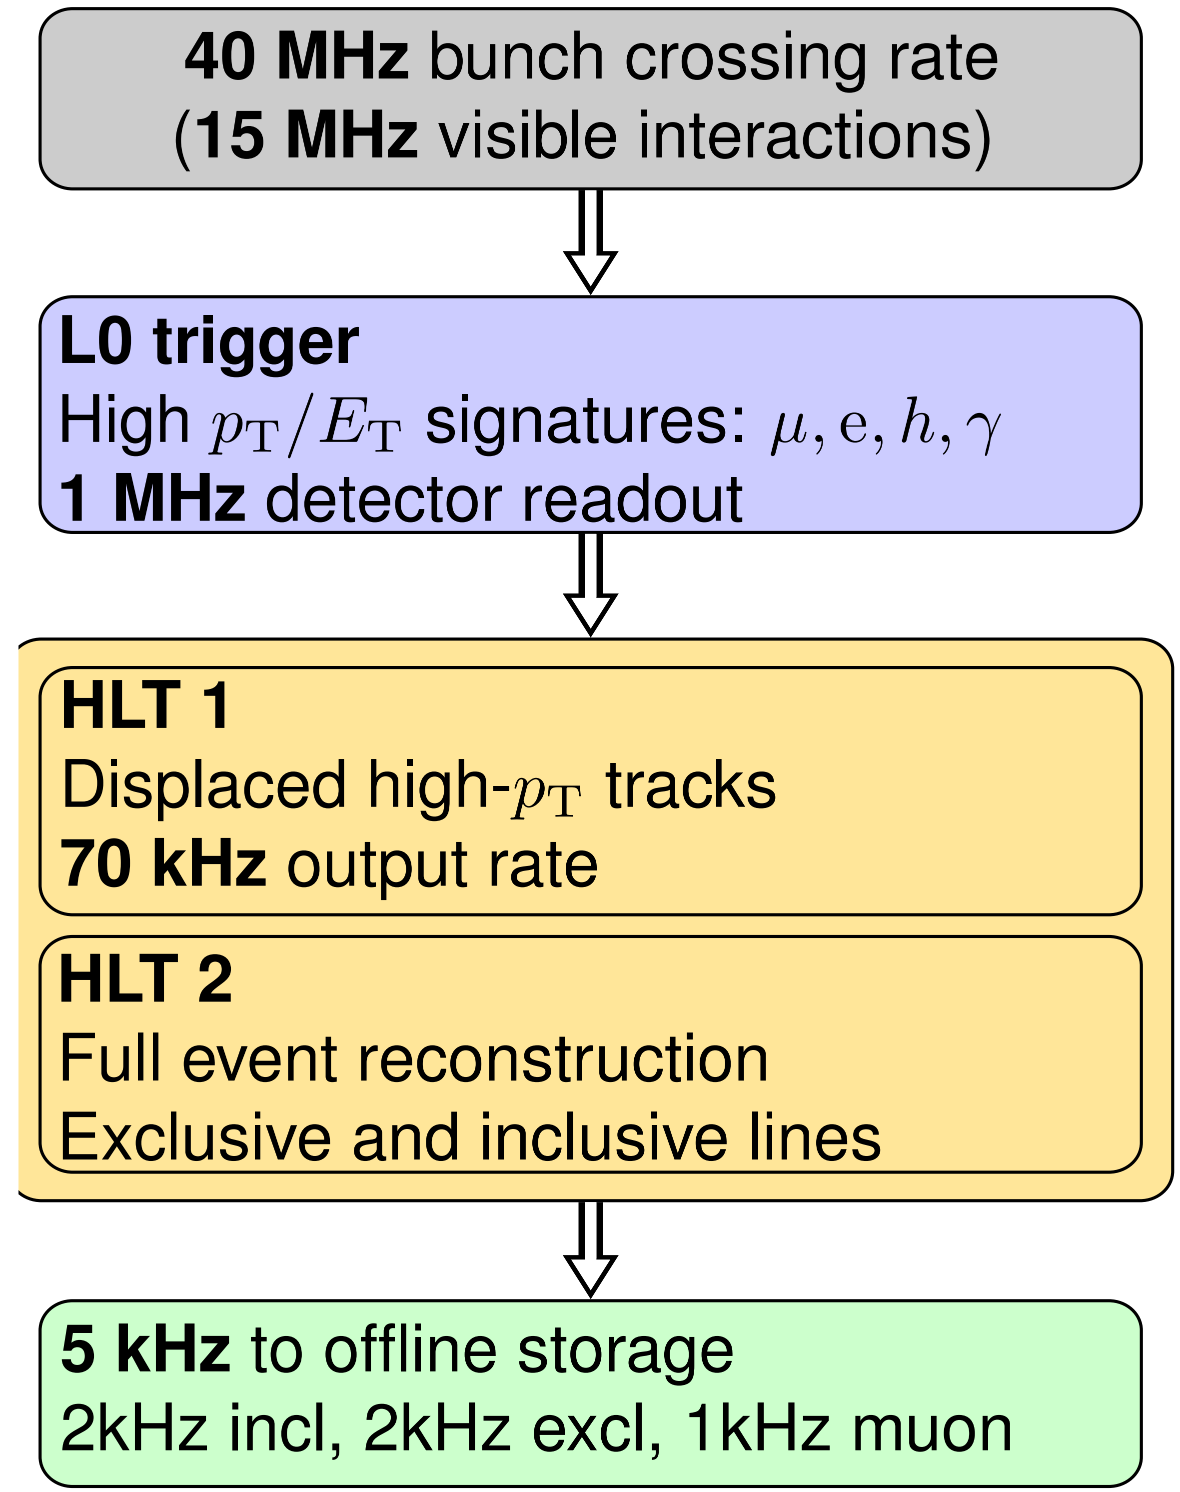

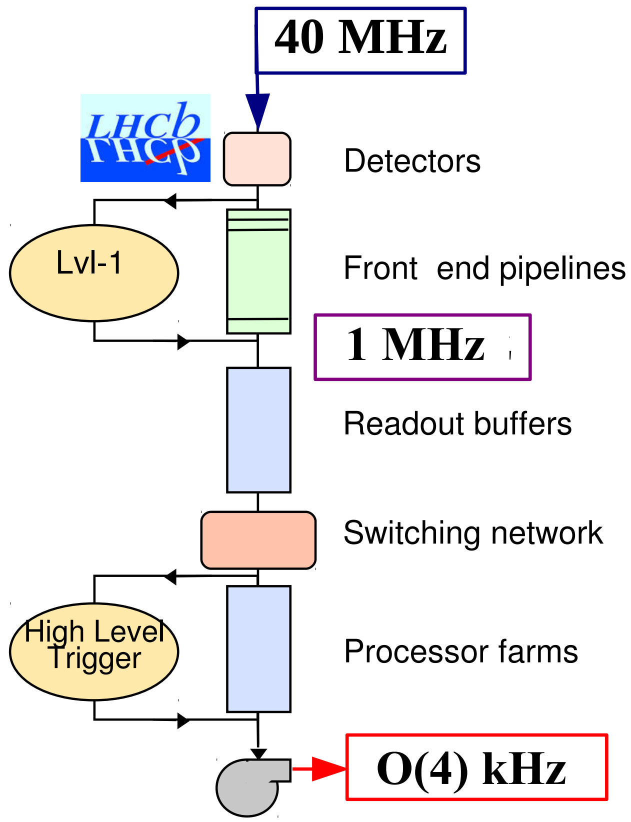

LHCb: Inclusive B triggers

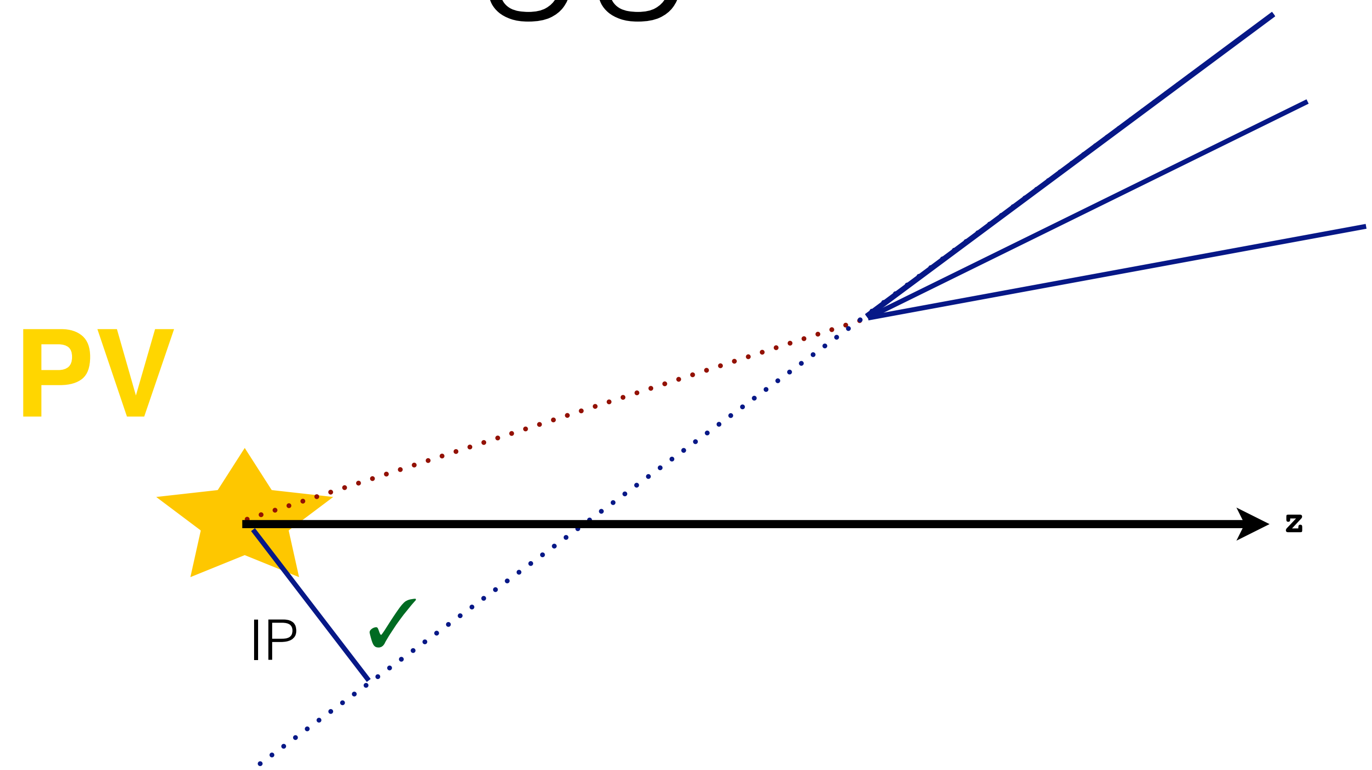

- B momentum O(100) GeV/c
- B mass O(5) GeV/c^2
- B lifetime ~1.5 ps

1. Typically (at least) one daughter has a "large" P_T
2. O(10%) of the time has a muon as daughter
3. Typically (at least) one daughter has "high impact parameter" (IP) —> requires 'track trigger

HLT1: Partial Tracking

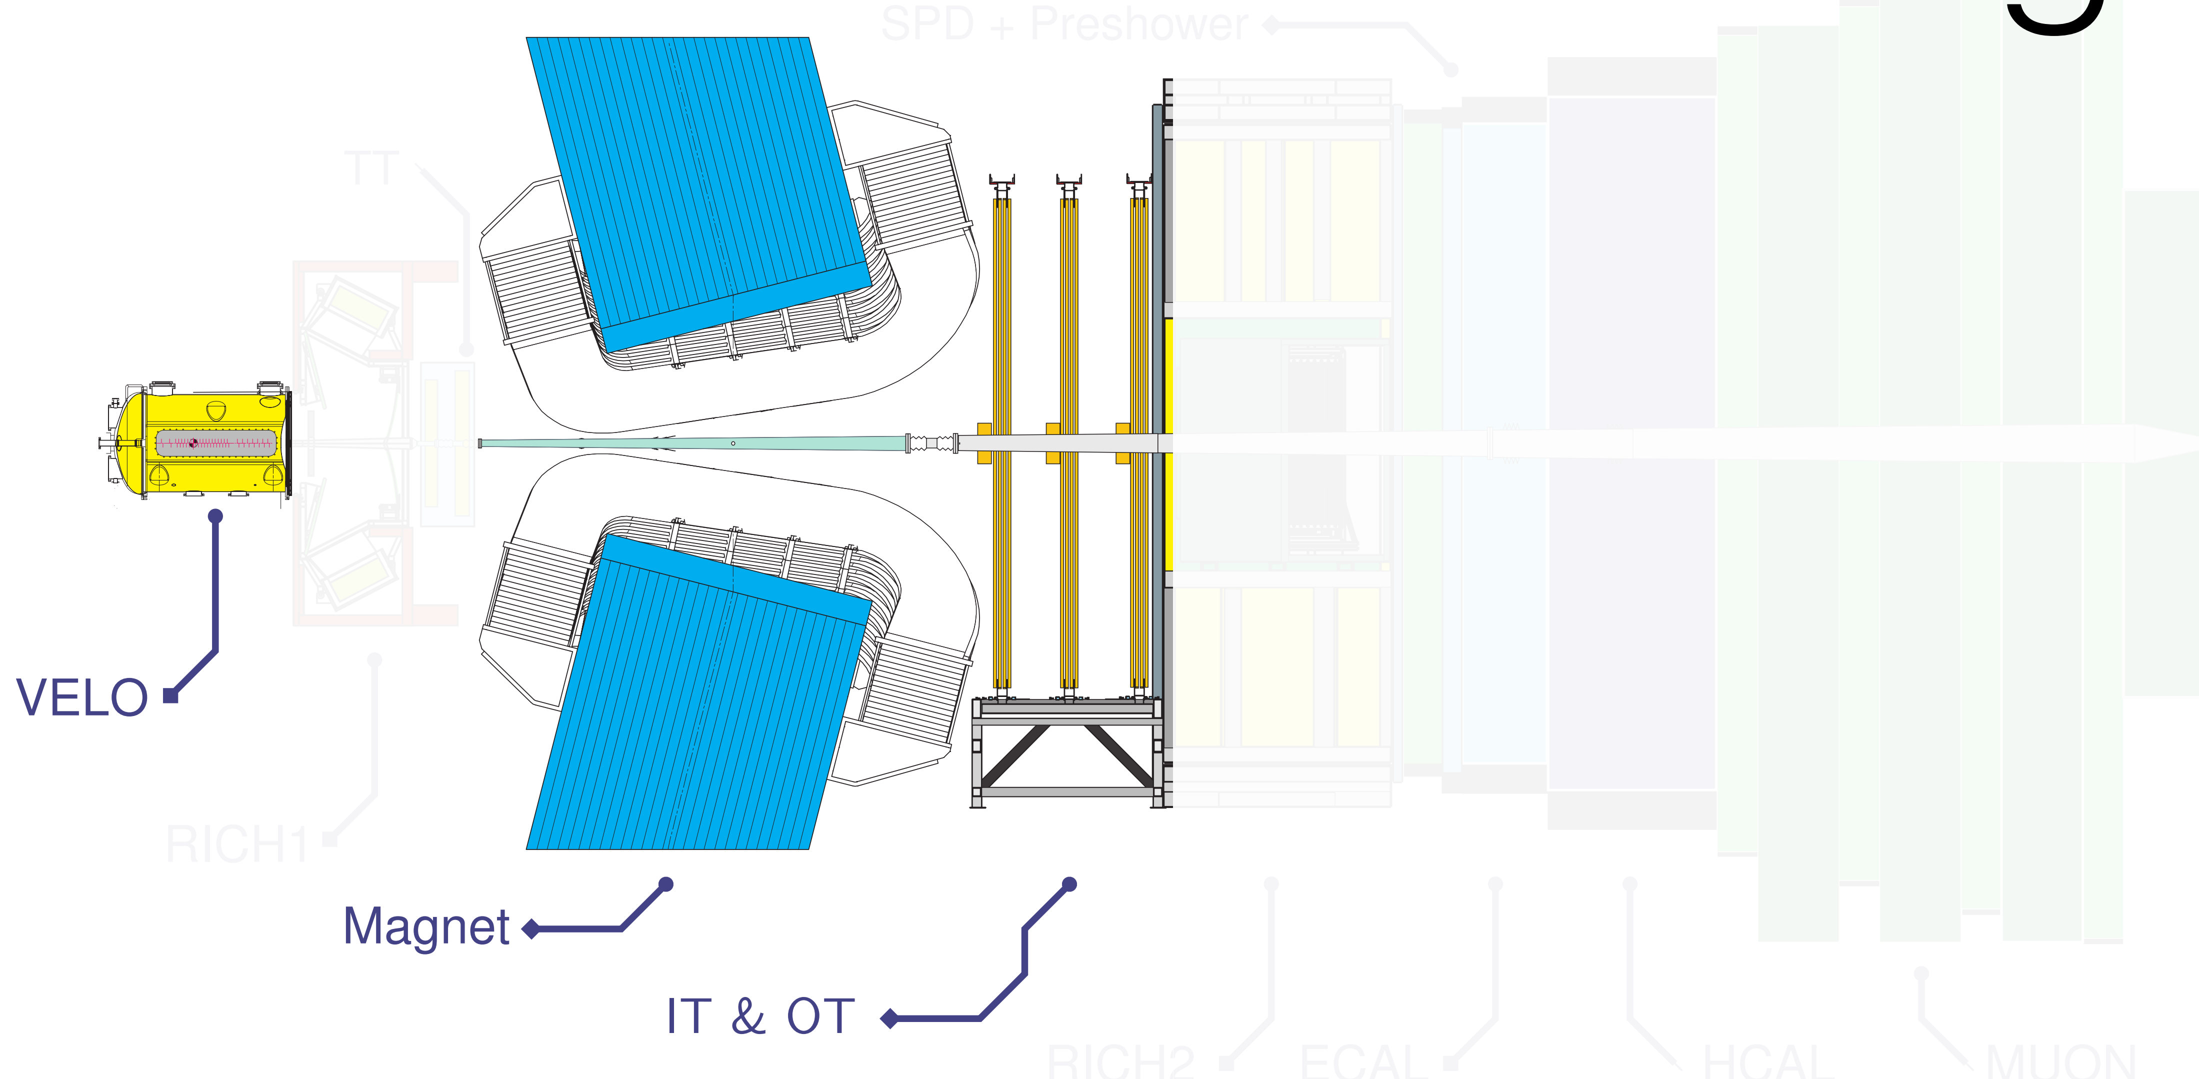

- HLT1 adds tracking in VErtex LOcator (VELO) and primary vertex reconstruction

- VELO tracks, either matched to muon hits, or with large IP are extended through the magnet

  - P_T dependent search windows:

track	μ	$\mu\ \mu$	other
min. p_T [GeV]	1.0	0.5	1.6

Hlt1Track trigger

- First reconstruct all tracks in vertex detector
- Compute primary vertices
- Select displaced track(s)
- Determine momentum of displaced track(s)
    - reconstruct in 'region of interest' — defined by minimum req. P / PT (3 / 1.6 GeV in 2012)

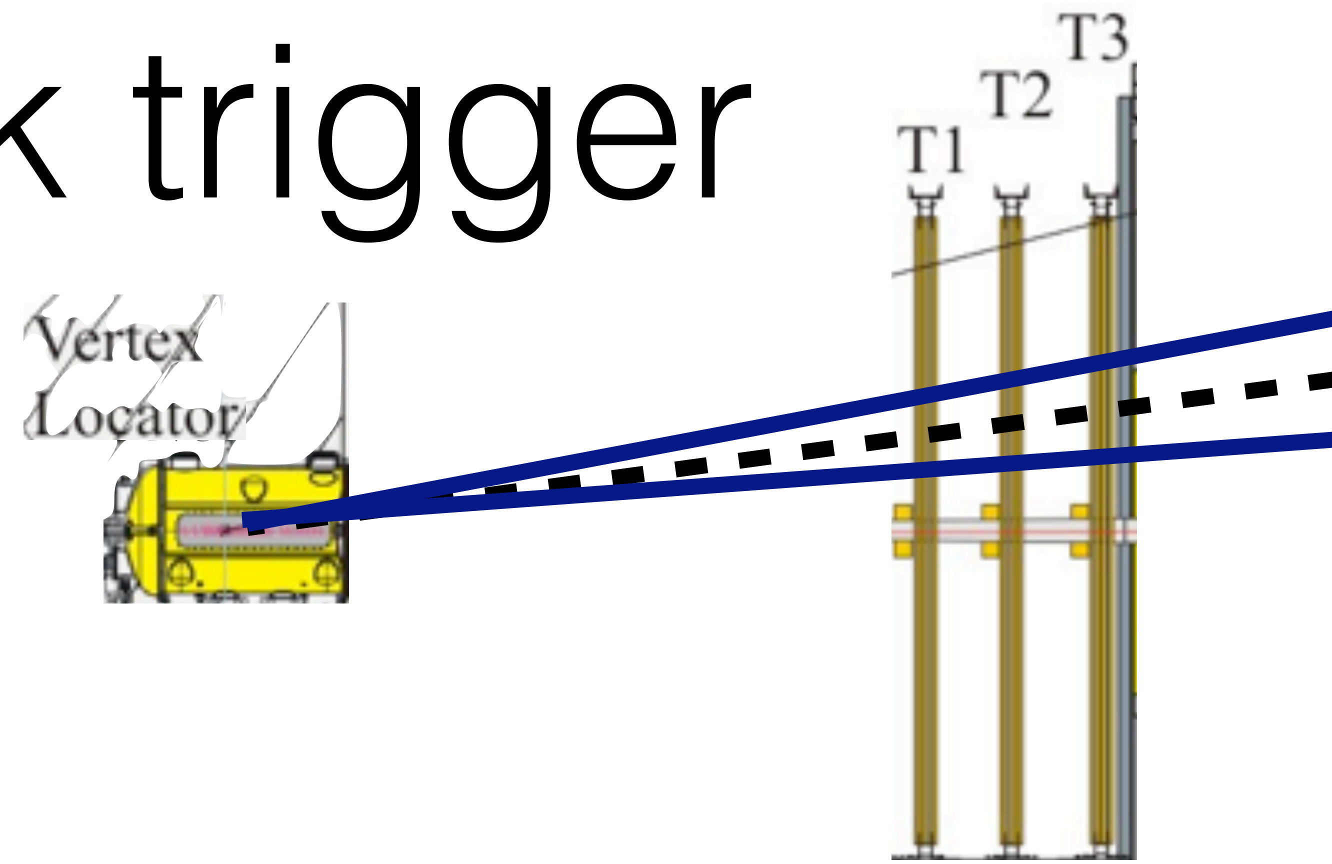

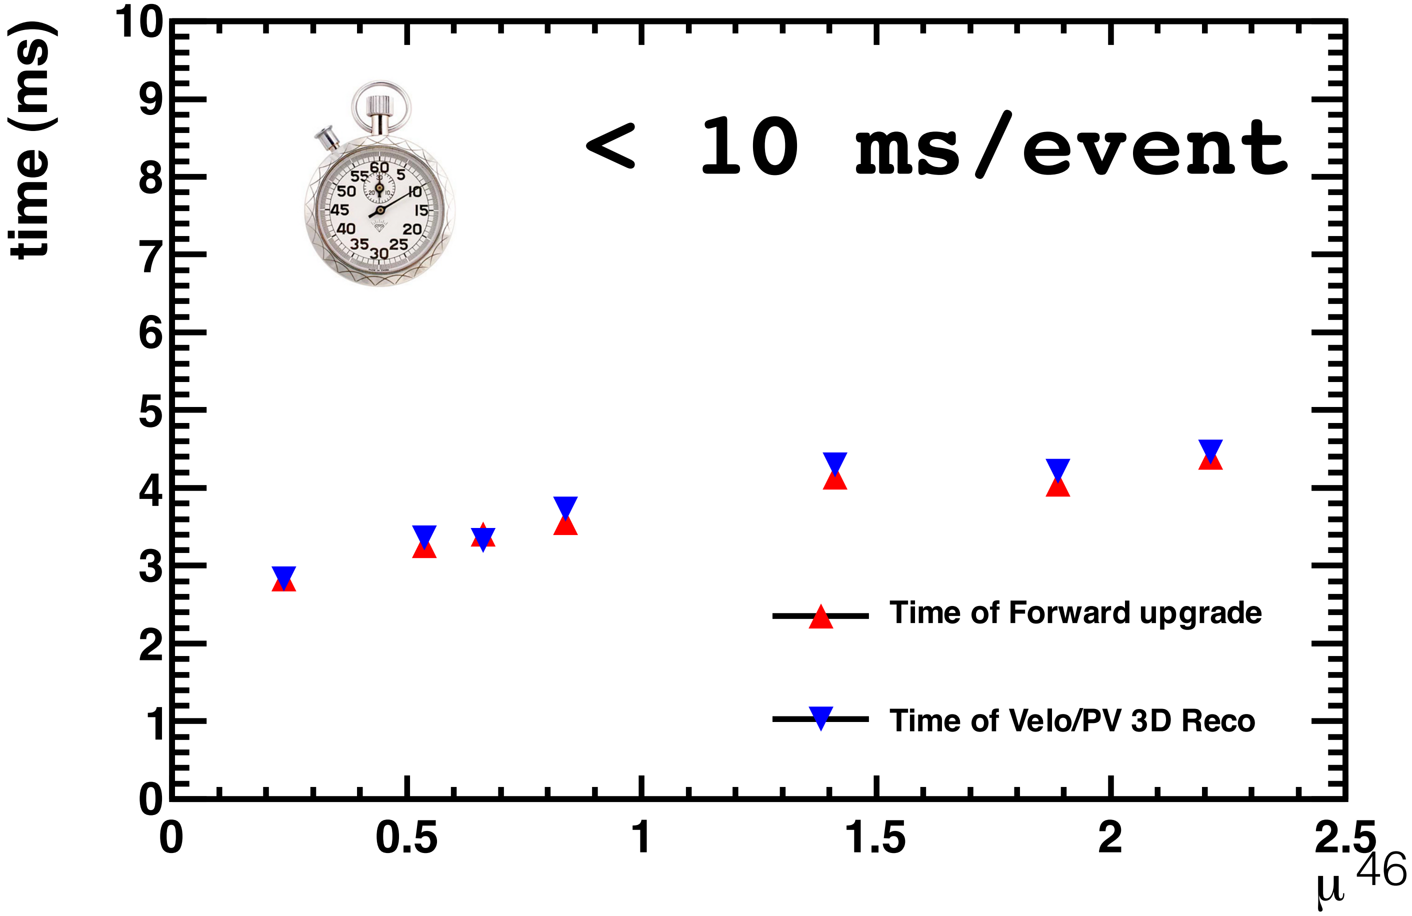

Single Muon trigger

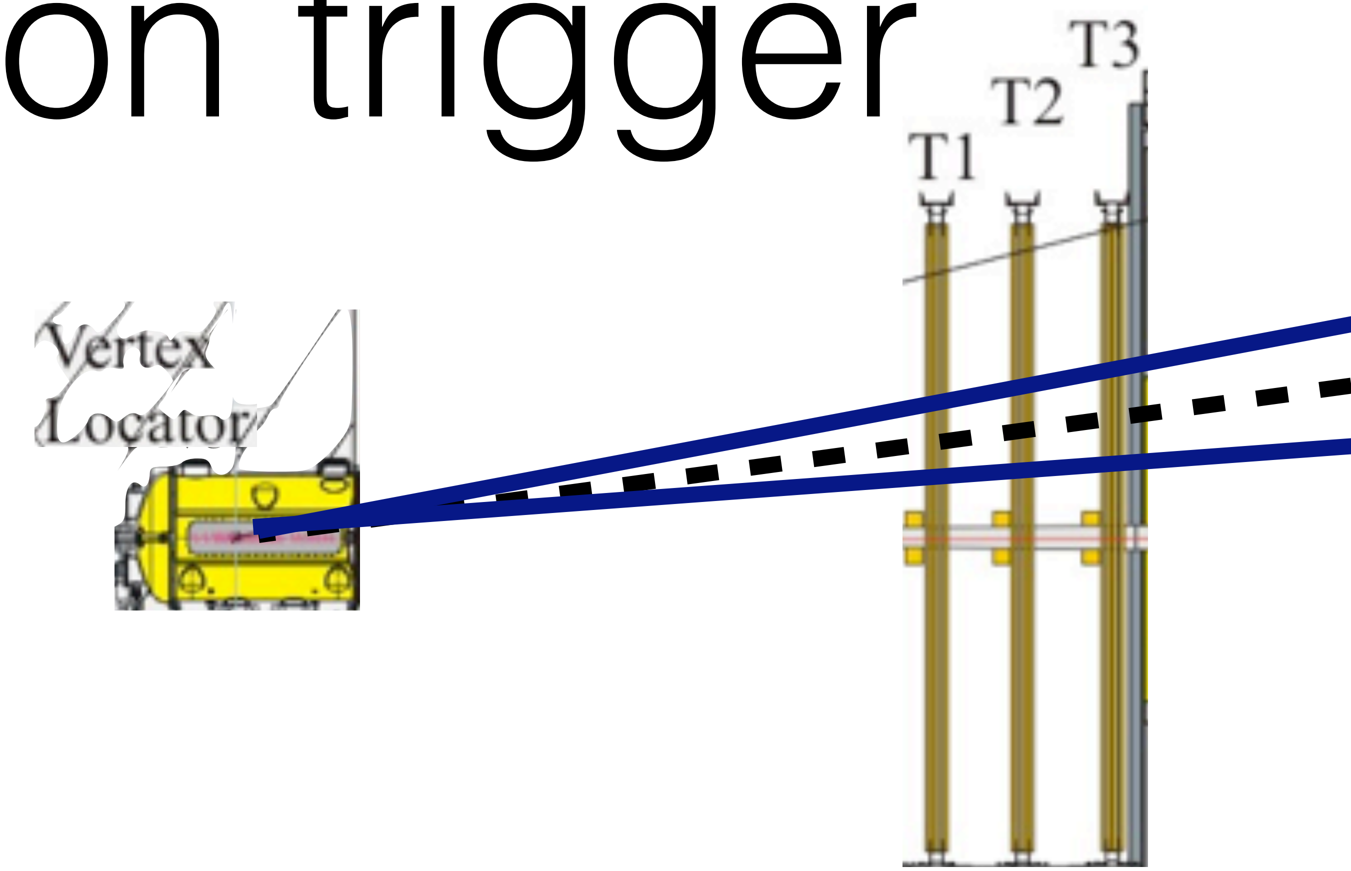

- First reconstruct all tracks in vertex detector
- Select muon matches
- Determine momentum of muon-matched track(s)
  - reconstruct in 'region of interest' — defined by minimum req. P / PT (3 / 0.5 GeV in 2012)
  - lower then 1Track because muon matching more 'selective' than IP
- Guarantee "independence", but avoid "duplicate work"

HLT1: Performance

- Muon lines
  - Track matched to muon hits
  - Either high P_T or large IP
  - ~ 14 kHz
- Inclusive lines
  - Single track with large IP and high P_T
  - ~56 kHz
- 1 MHz input → 70 kHz output
  - Tuned to HLT2 CPU usage
  - At this point, ~half the CPU capacity has been used — so HLT2 can take 14x longer

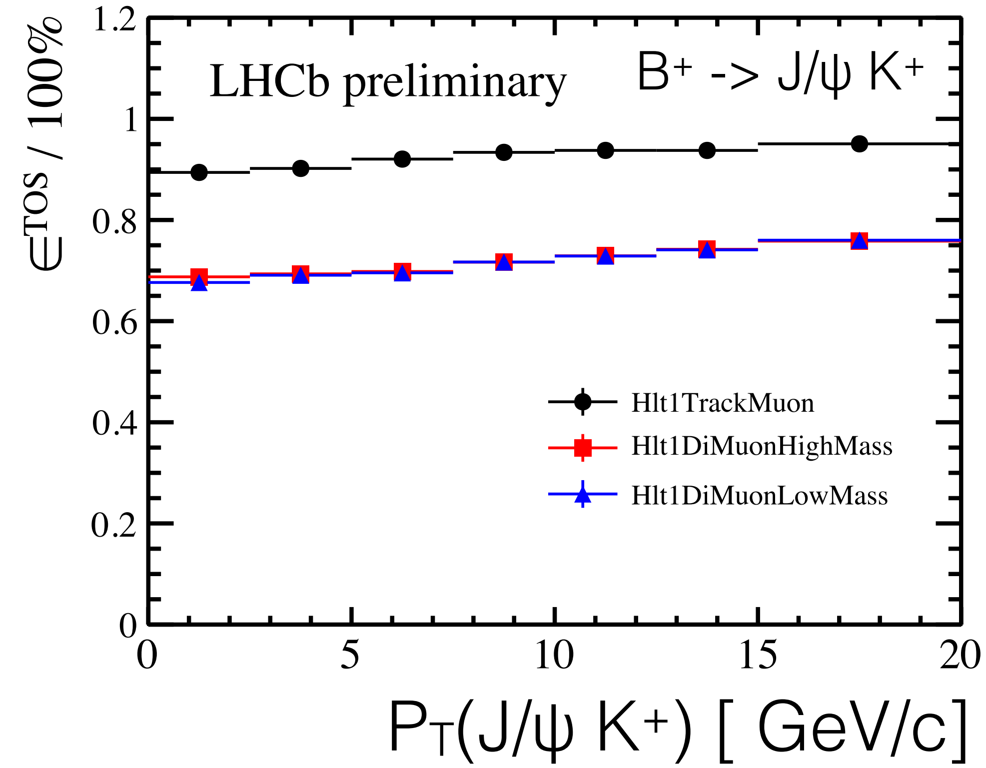

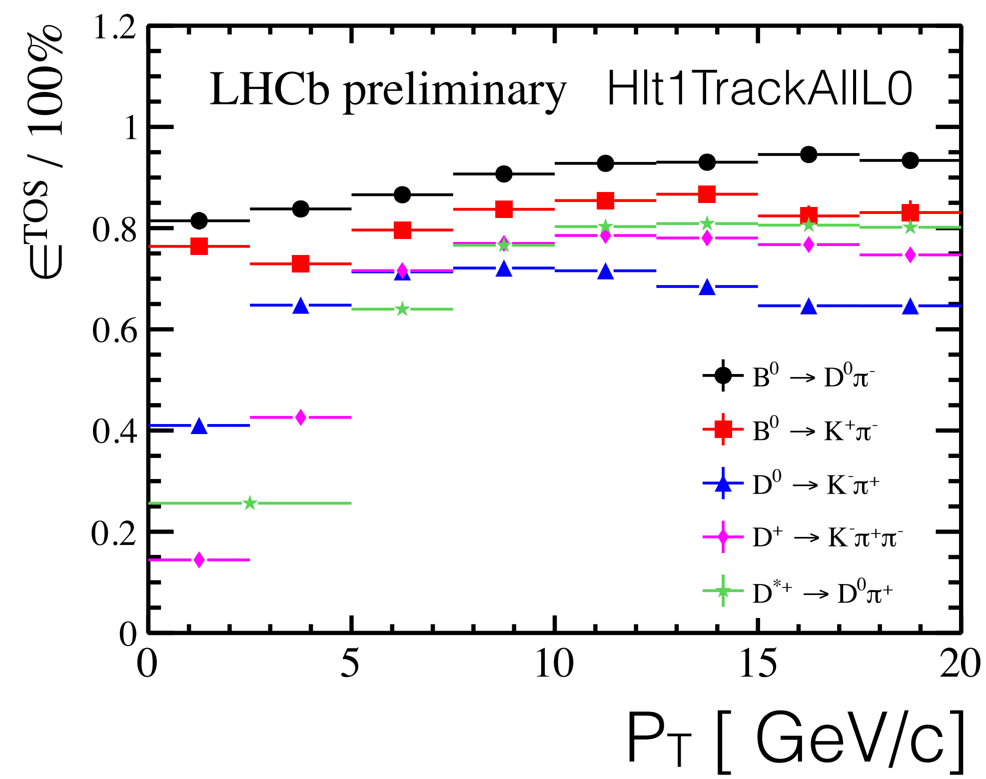

HLT2: Full Reconstruction

- Tuned versions of offline reconstruction algorithms
    - eg. $P_T > 300$ MeV/c
- Combination of inclusive and exclusive trigger decisions
- Flexible software environment
    - supports eg. dedicated MVA-based selections

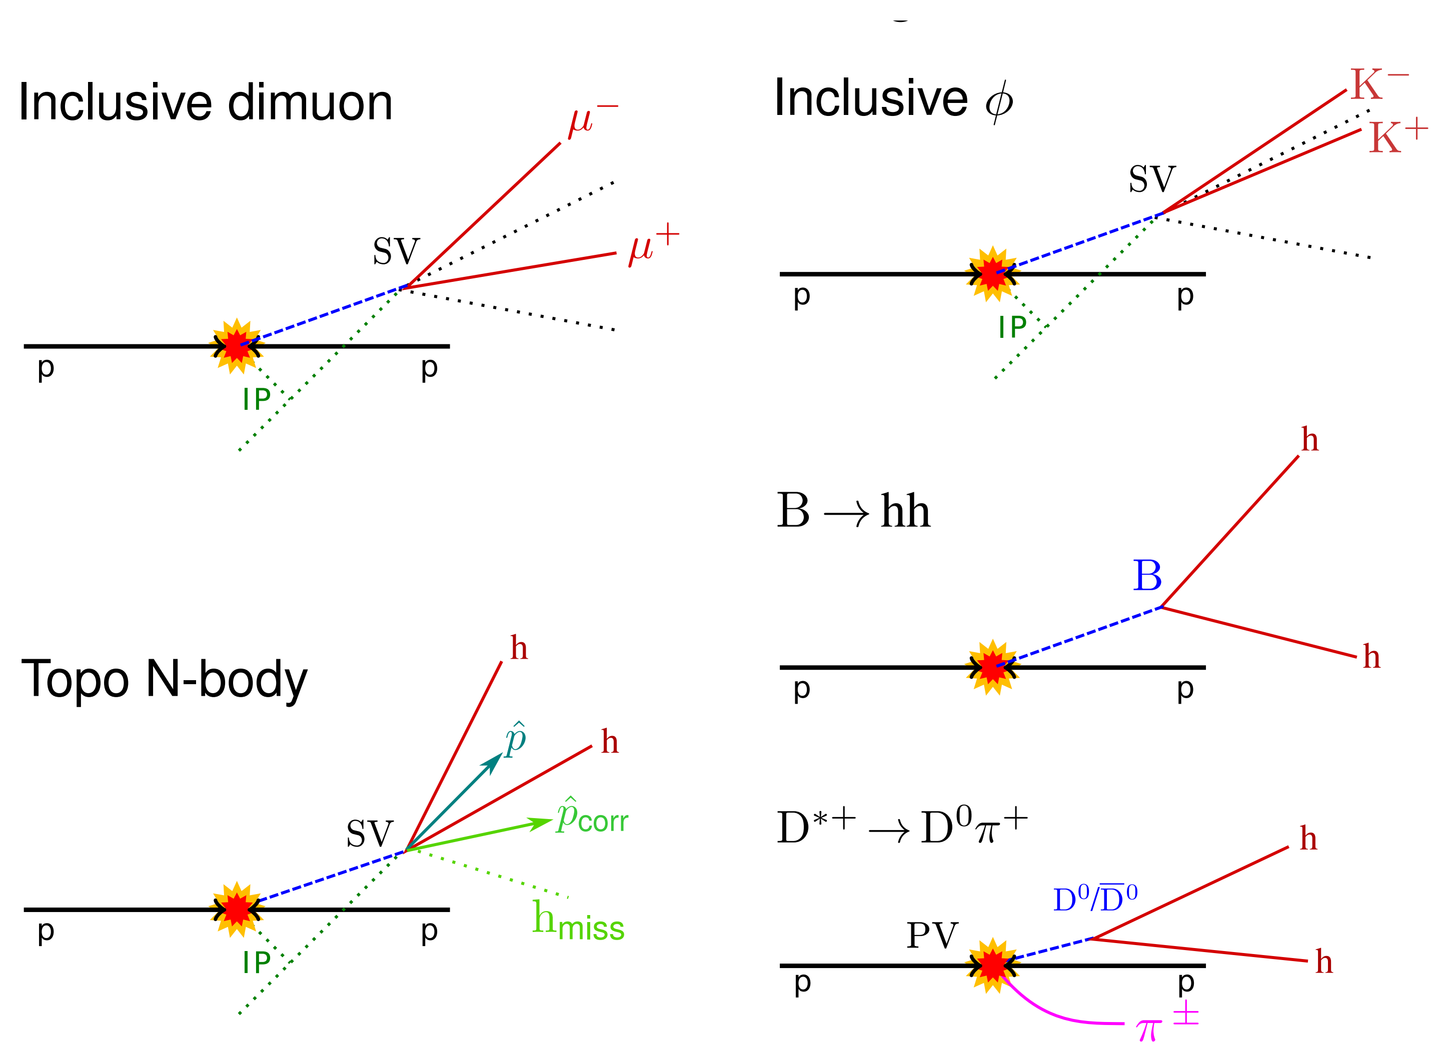

Topological N-body Triggers

- Utilizes excellent vertex and momentum resolution to compute:

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

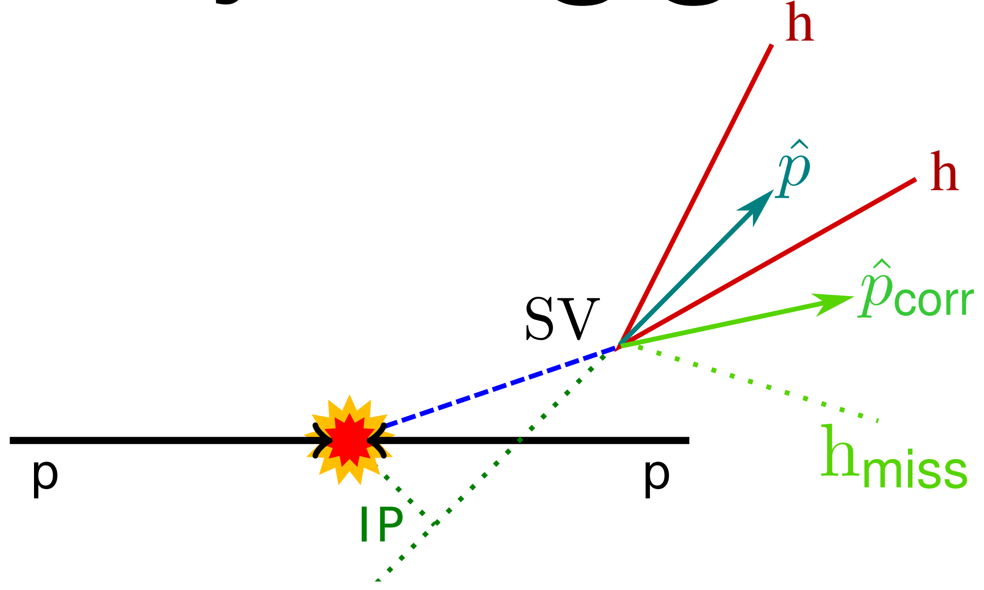

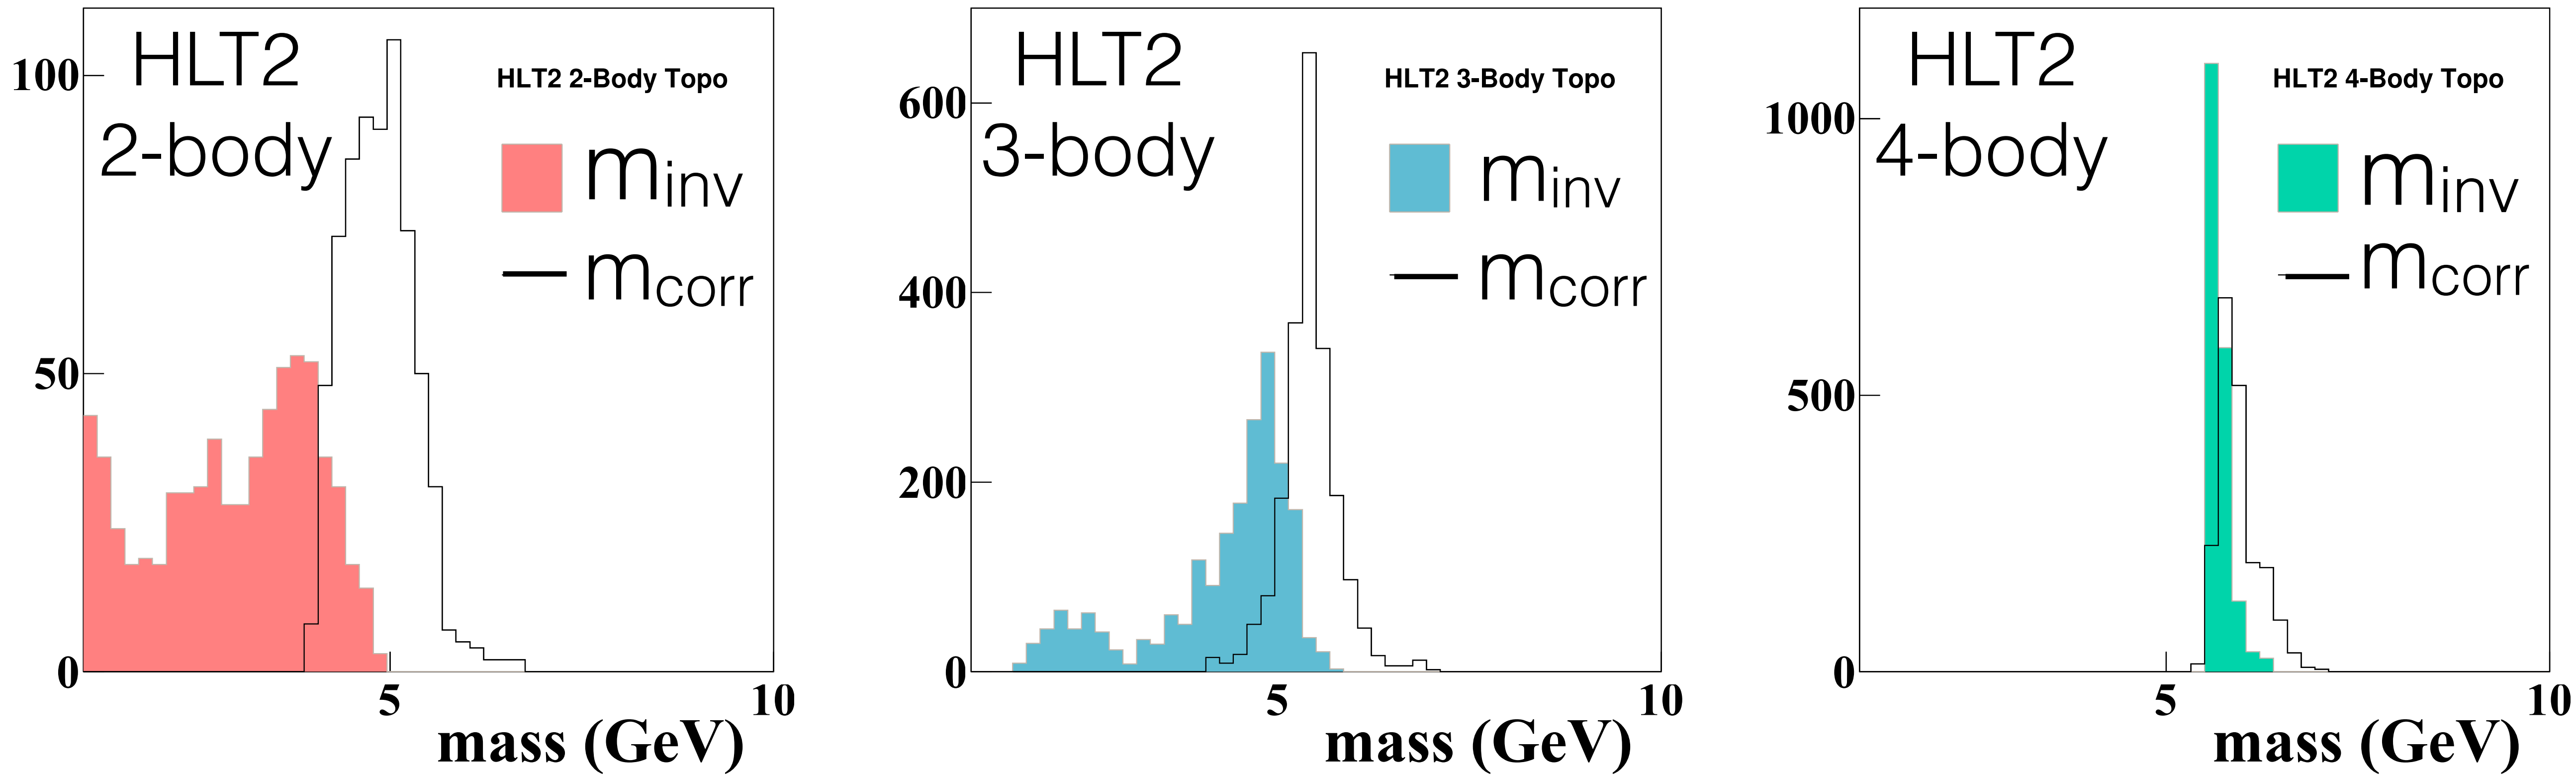

Example: 4-body B decay, m_{inv} and m_{corr} for 2, 3 and 4 body selections

Topological N-body Triggers

- Utilizes excellent vertex and momentum resolution to compute:

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

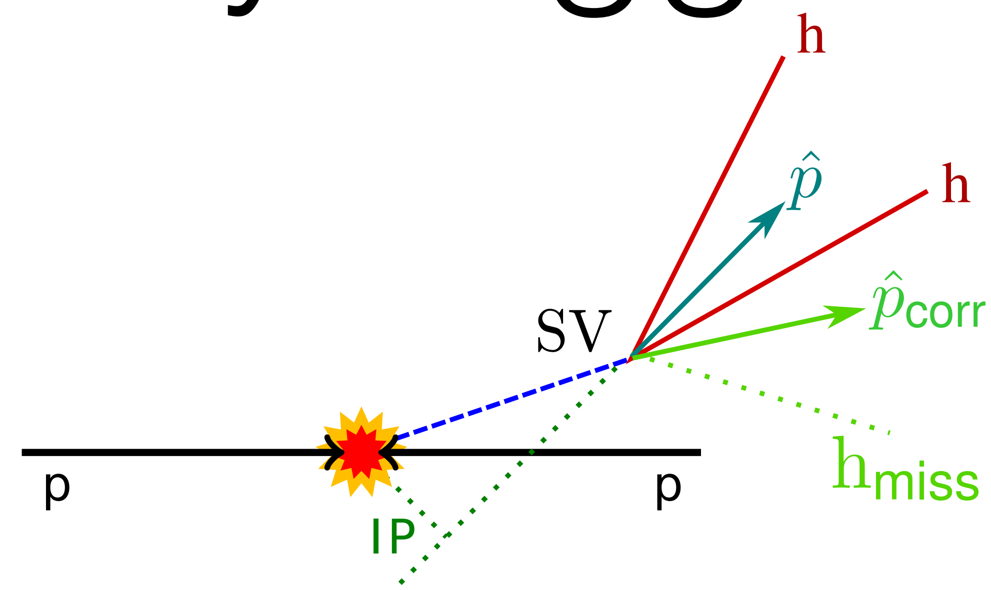

- Uses a dedicated "Bonzai" Boosted Decision Tree [JINST 8 (2013) P02013] with
  - P_T, $IP\chi^2$, $FD\chi^2$, m_{inv}, m_{corr}
- Capable of filling its allotted bandwidth with ~100% pure generic $b\bar{b}$ events

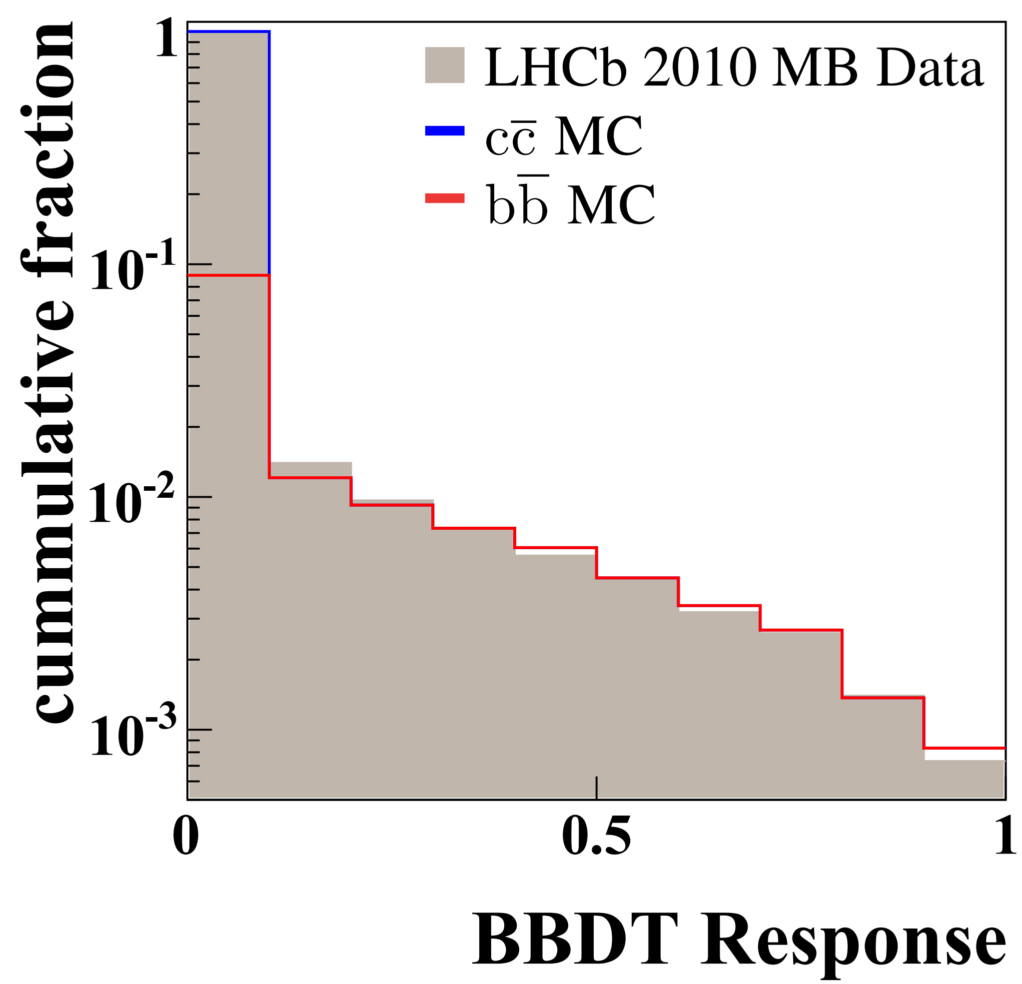

ENTR'ACTE

Trigger & DAQ

Gerhard Raven
VU Amsterdam & Nikhef

- constraints
- physics
- architectures
- strategy
- examples

- resource optimization
- special cases
- interaction w. analysis
- upgrades

Acknowledgements:
Niko Neufeld, Vladimir Gligorov, Paris Sphicas, Brian Dahmes, ISOTDAQ lecturers, and many, many others…

2015 CERN-Fermilab Hadron Collider Physics School

Bottlenecks!

digitization →

data transfer →

event building →

processing →

permanent storage →

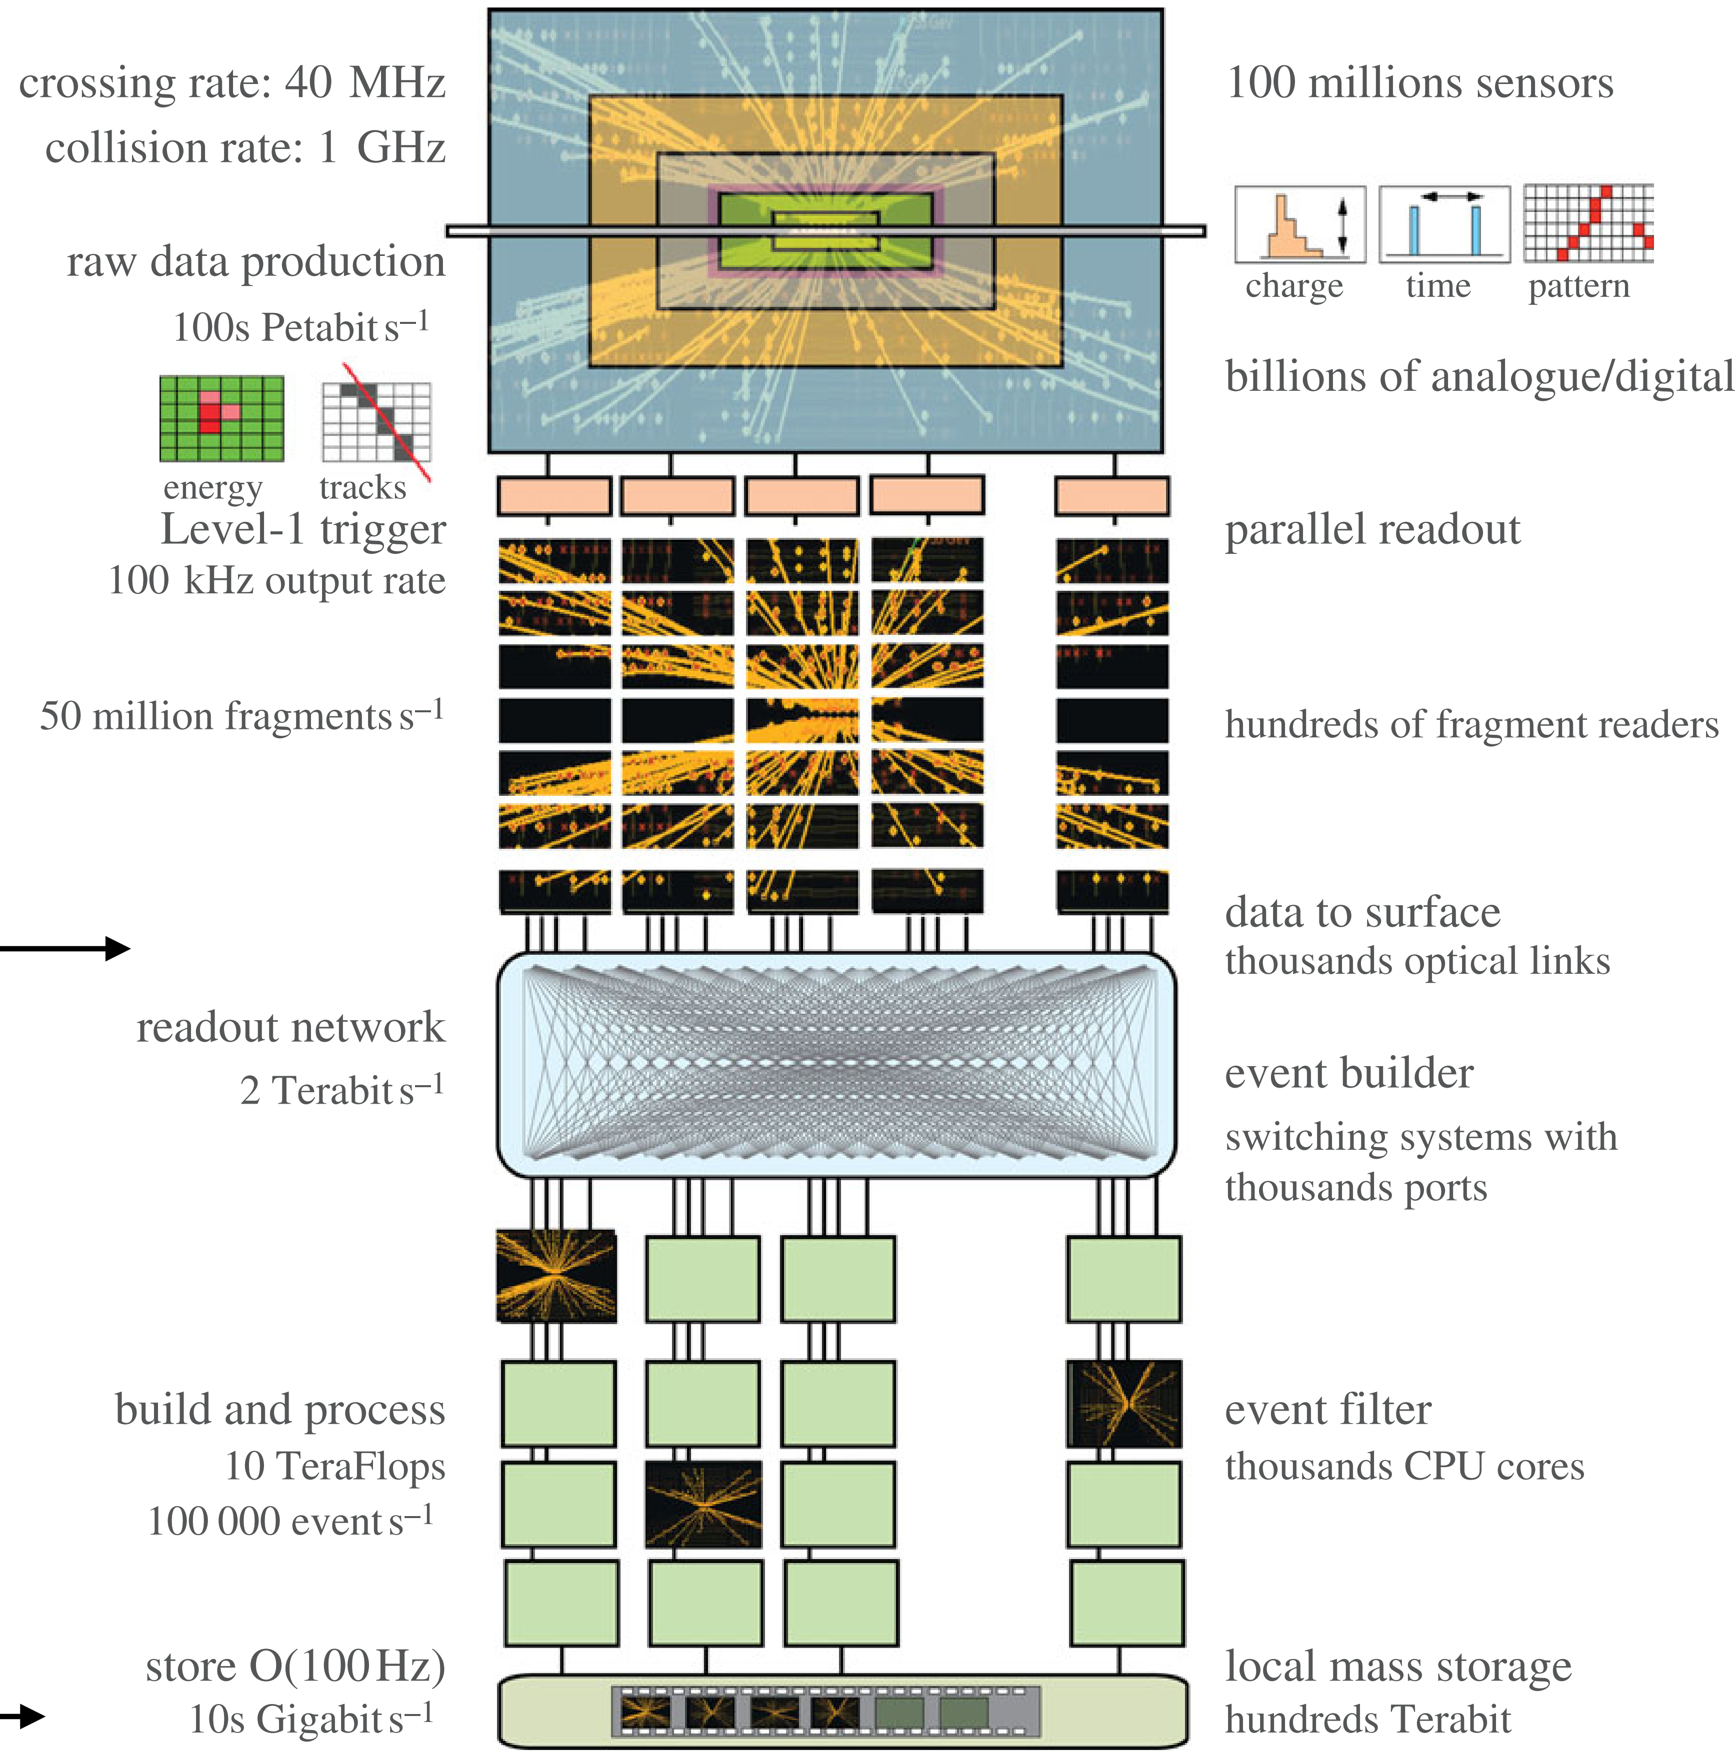

Over-committing

- During 2012, both CMS and Atlas ‘overcommitted’ their output rate
- The excess is ‘parked’ (delayed) and only processed offline eg. during shutdowns
- Allows to eg. increase statistics for precision measurements

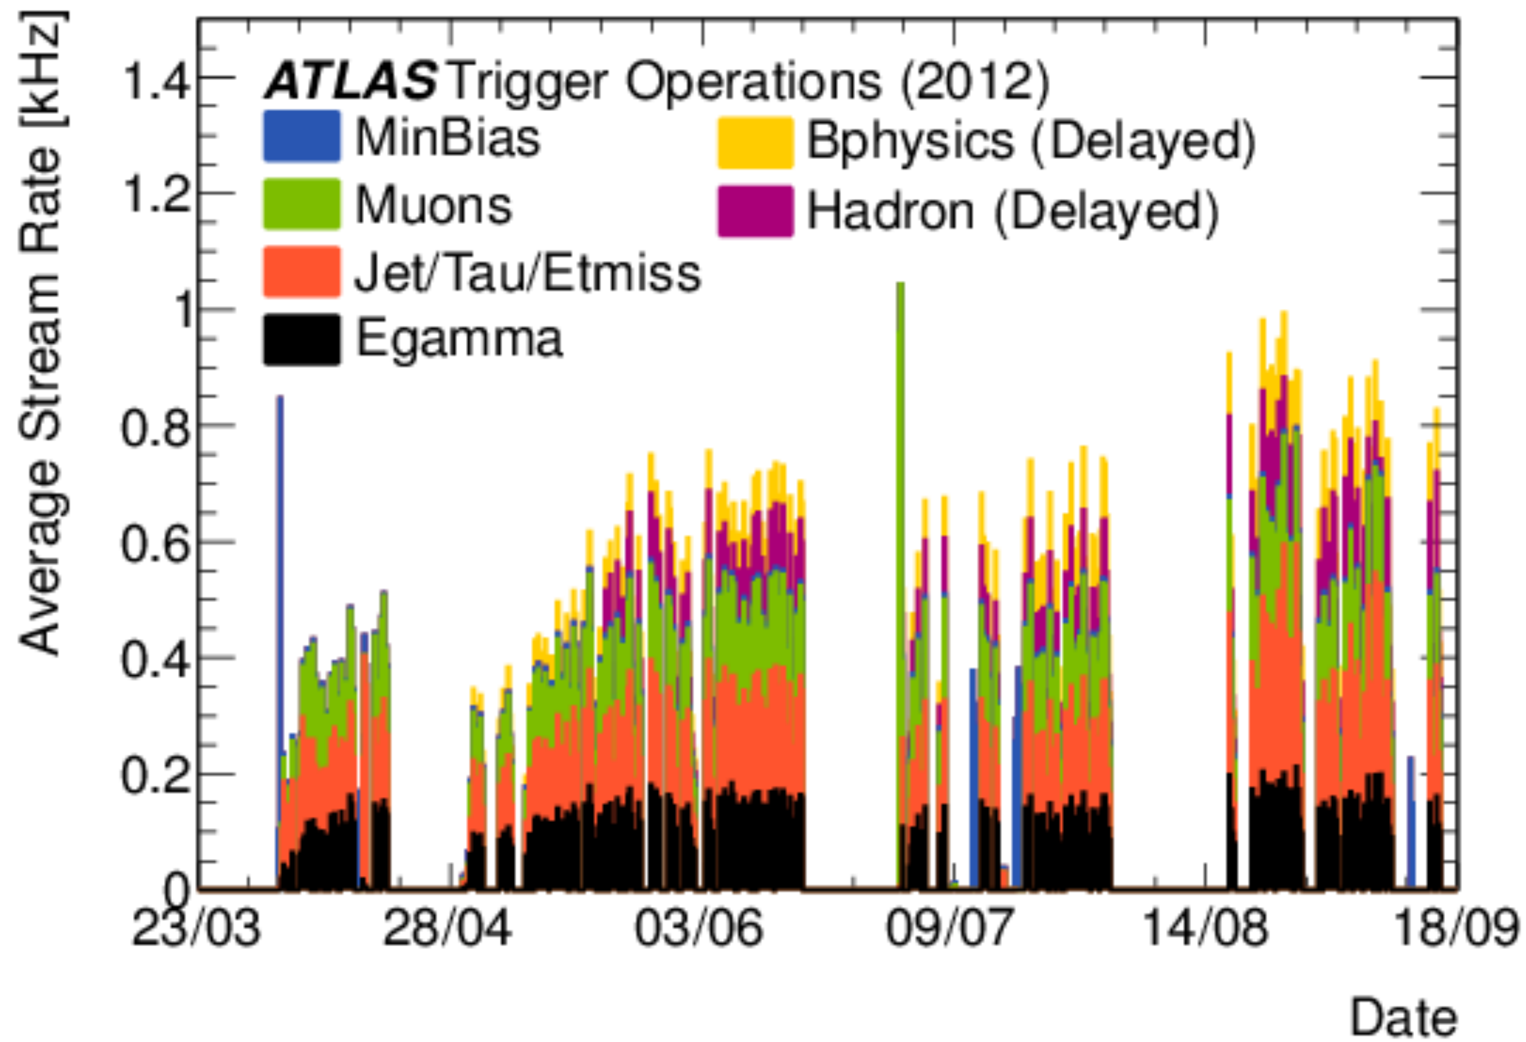

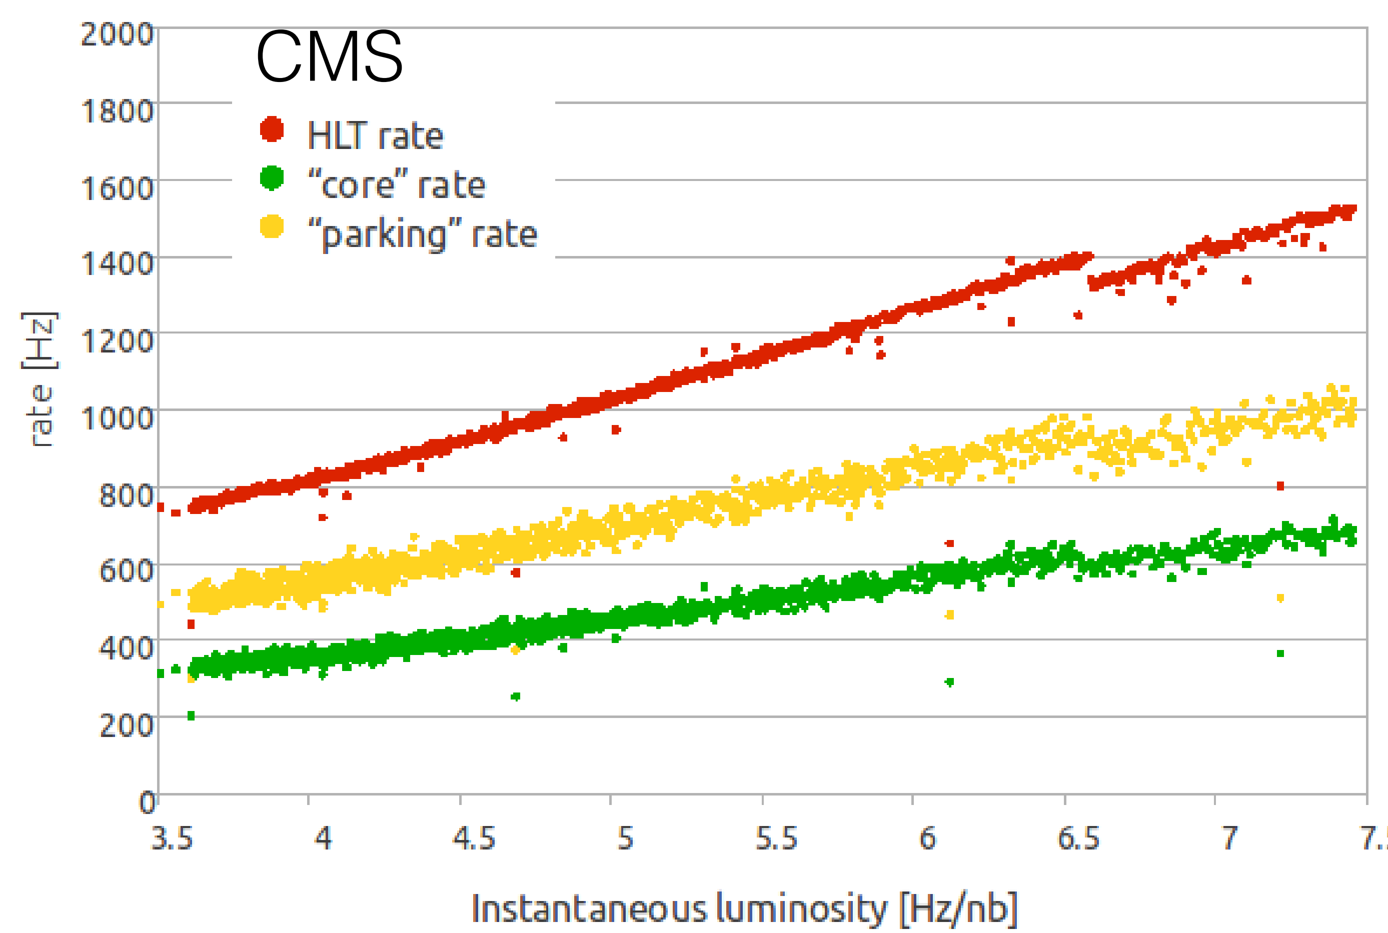

LHCb 2012: Deferred Trigger

- LHC "only" delivers collisions ~35% of the time
  - trigger farm idle ~65% of the time!

Table 5. LHC availability 2012

Mode	% of scheduled time
Access	14%
Setup	28%
Beam in	15%
Ramp and squeeze	8%
Stable beams	36%

LHCb 2012: Deferred Trigger

- LHC "only" delivers collisions ~35% of the time
  - trigger farm idle ~65% of the time!
- "Over commit" CPU resources, buffer overflow to local disk & catch up in between fills
  - 20% of L0 triggers are "deferred"
  - 25% extra CPU capacity!
  - allows decrease of Hlt2 tracking thresholds P_T> 500 MeV/c→ P_T>300 MeV/c

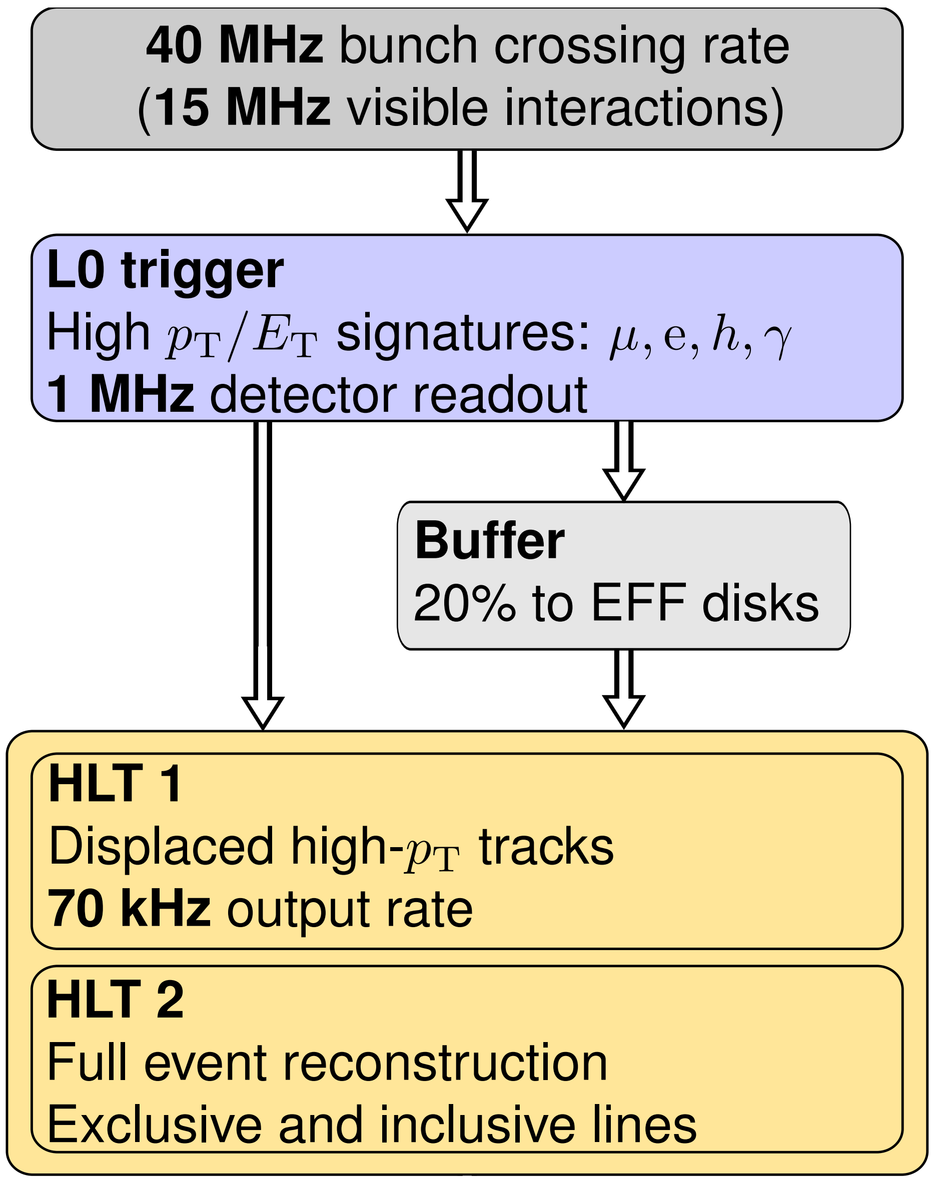

LHCb 2012: Deferred Trigger

- LHC "only" delivers collisions ~35% of the time
    - trigger farm idle ~65% of the time!
- "Over commit" CPU resources, buffer overflow to local disk & catch up in between fills
    - 20% of L0 triggers are "deferred"
    - 25% extra CPU capacity!
    - allows decrease of Hlt2 tracking thresholds P_T> 500 MeV/c→ P_T>300 MeV/c
    - Peak disk usage in 2012: 88%

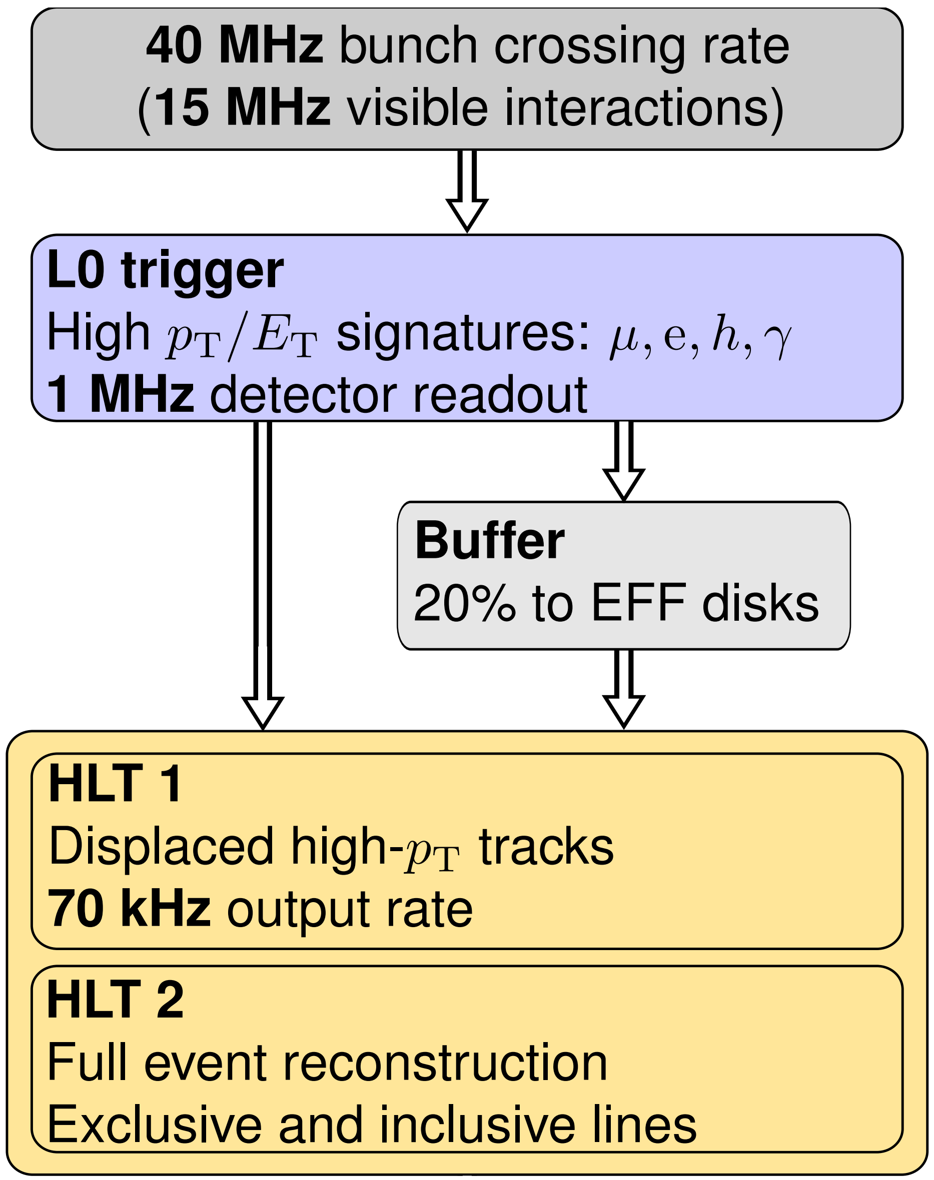

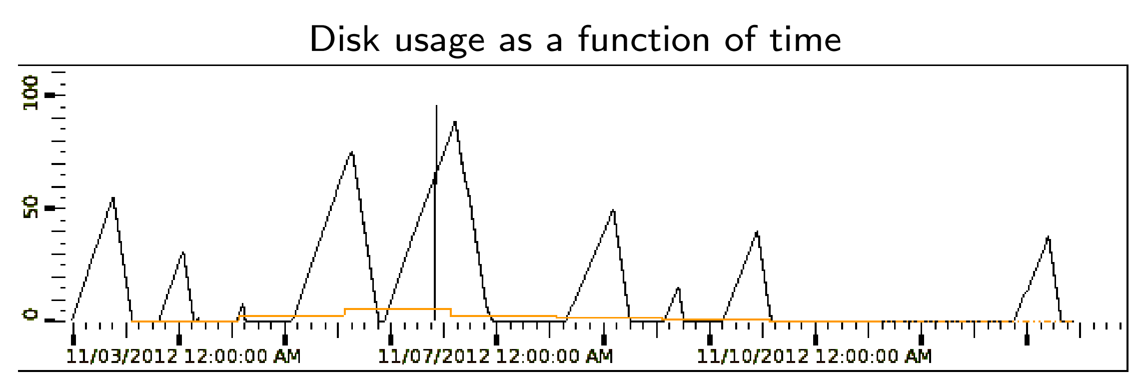

Reduced events

CMS ‘scouting’:

- lower thresholds -> higher rate
- compensate by reducing event size
  - eg only HLT calorimeter jets — O(kB)
- If something shows up, incorporate into core physics / parked triggers

data scouting extended the dijets searches below 1 TeV

CHEP 2013

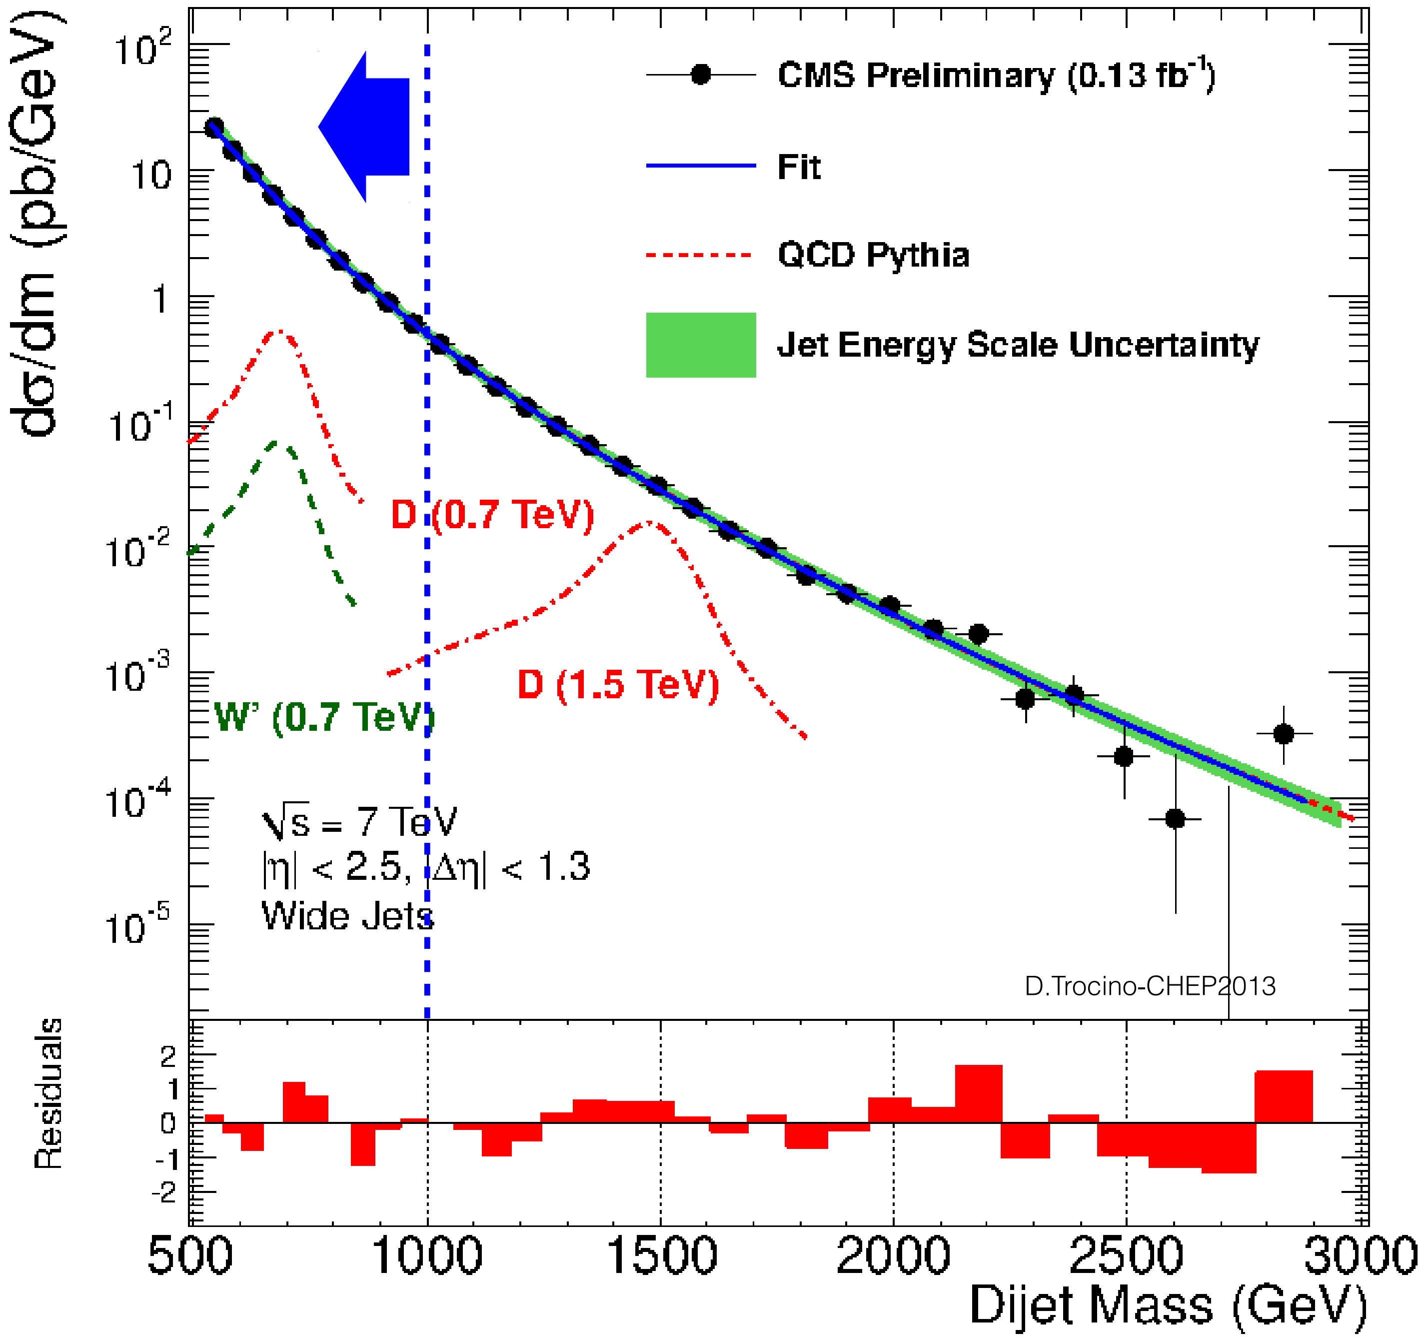

LHCb Lumi Trigger

- reduced event — only tens of bytes
    - stores only ‘luminosity counters’:
        - # of vertices, # of tracks, various other multiplicities
- Generated at a known, fixed rate, O(1 kHz)
    - rate is ‘relative prime’ wrt. LHC revolution frequency
- kept with all data streams at all times
- used as ‘heartbeat’ — allows *rate limited* lines in a distributed system (without assumptions on # of CPUs, CPUs with different speeds, …)

LHCb luminosity calibration

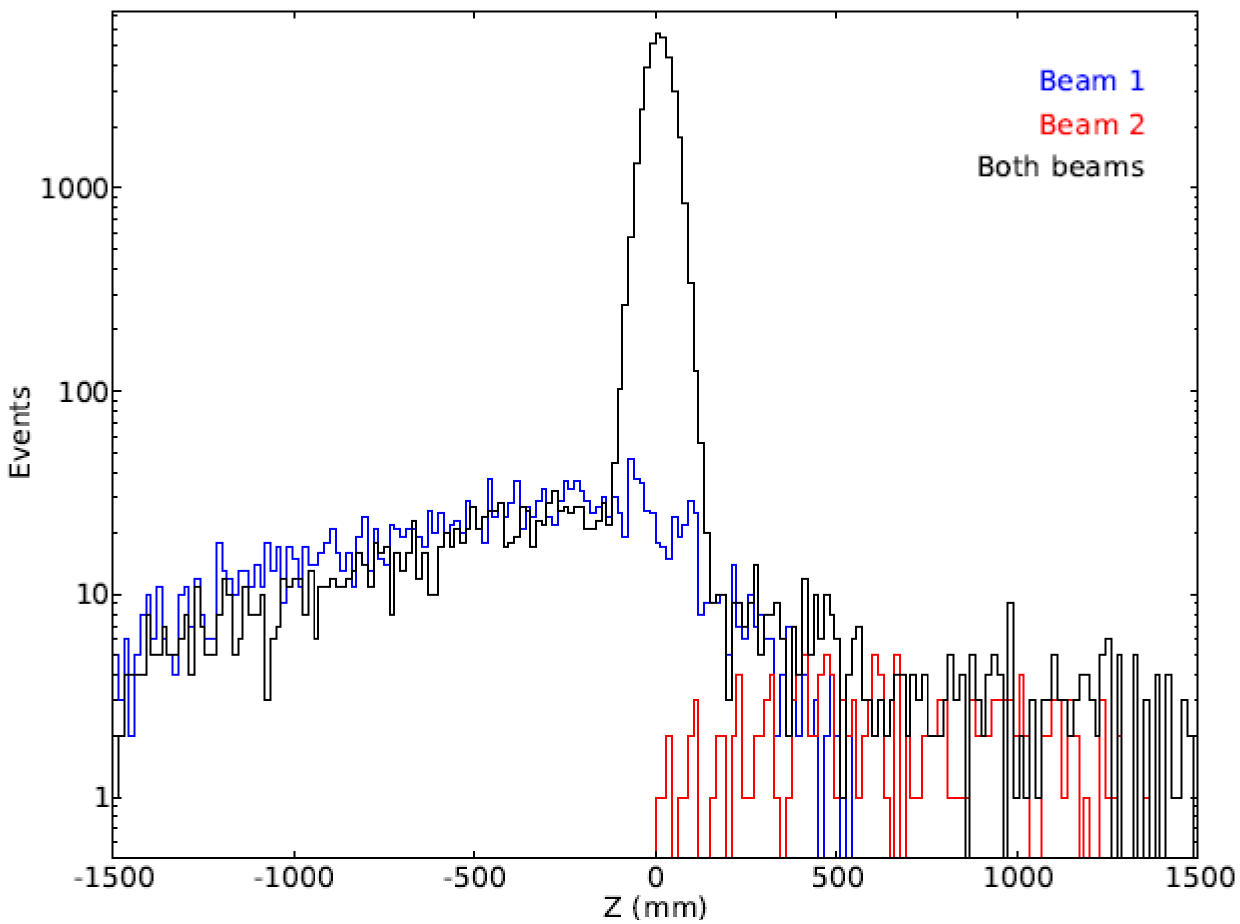

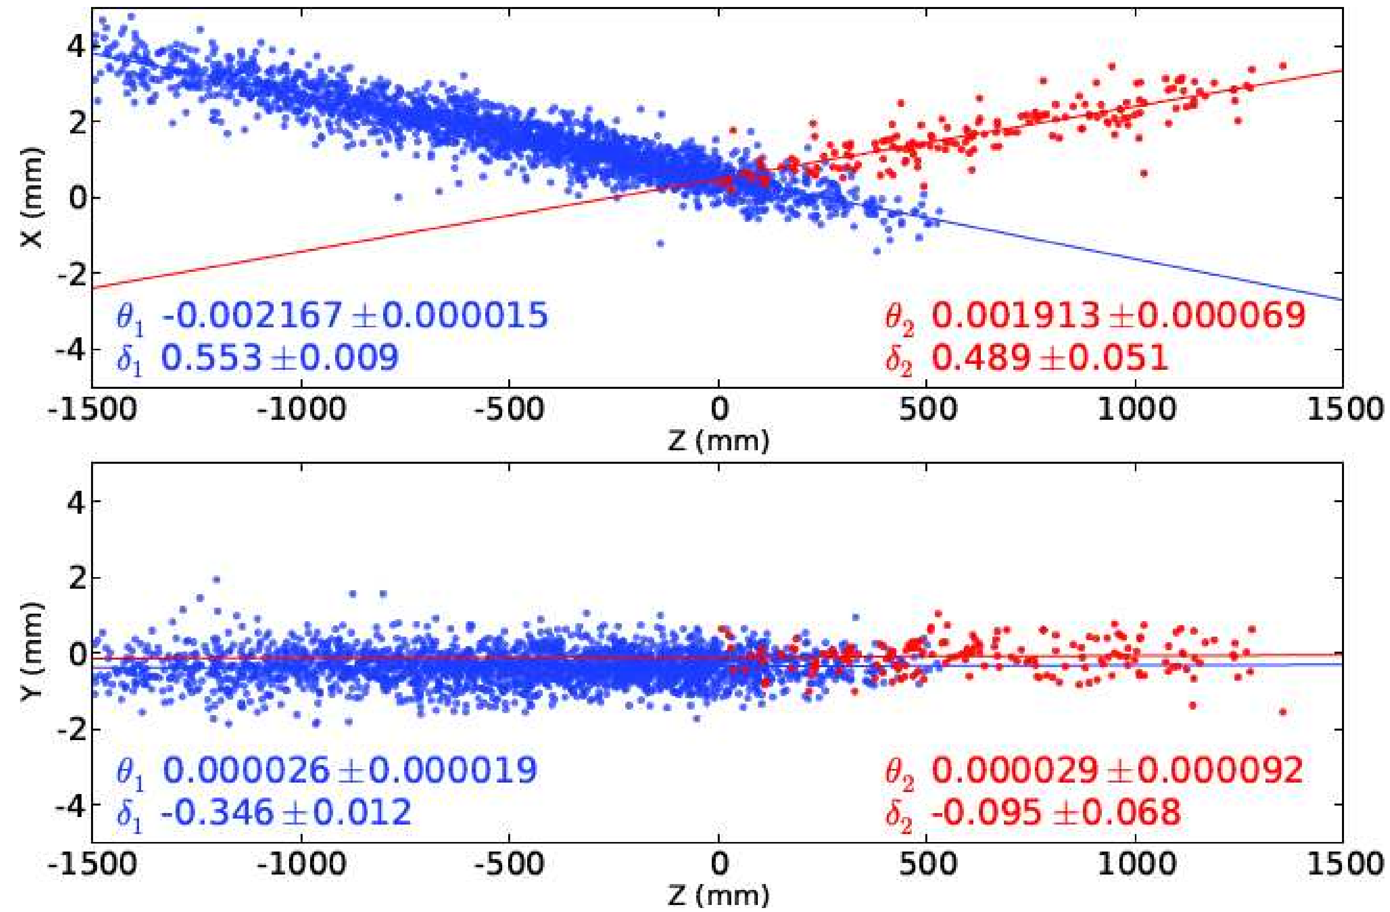

	25 Apr 2010	**15 Oct 2010**
LHC fill number	1059	1422
$N_{1,2}$ (10^{10} protons)	1	7–8
$\beta^\star$ (m)	2	3.5
$n_{\rm coll}/n_{\rm tot}$	1/2	12/16
$\mu_{\rm vis}^{\rm max}$	0.03	1
Trigger	minimum bias	22.5 kHz random $\sim$130 Hz minimum bias beam-gas
$\tau_{N_1 N_2}$ (h)	950	700
τ_L (h)	30	46

Use dedicated configuration:

- HLT lumi counters
- vertex detector raw data,

to increase rate by ~10x

LHCb luminosity calibration

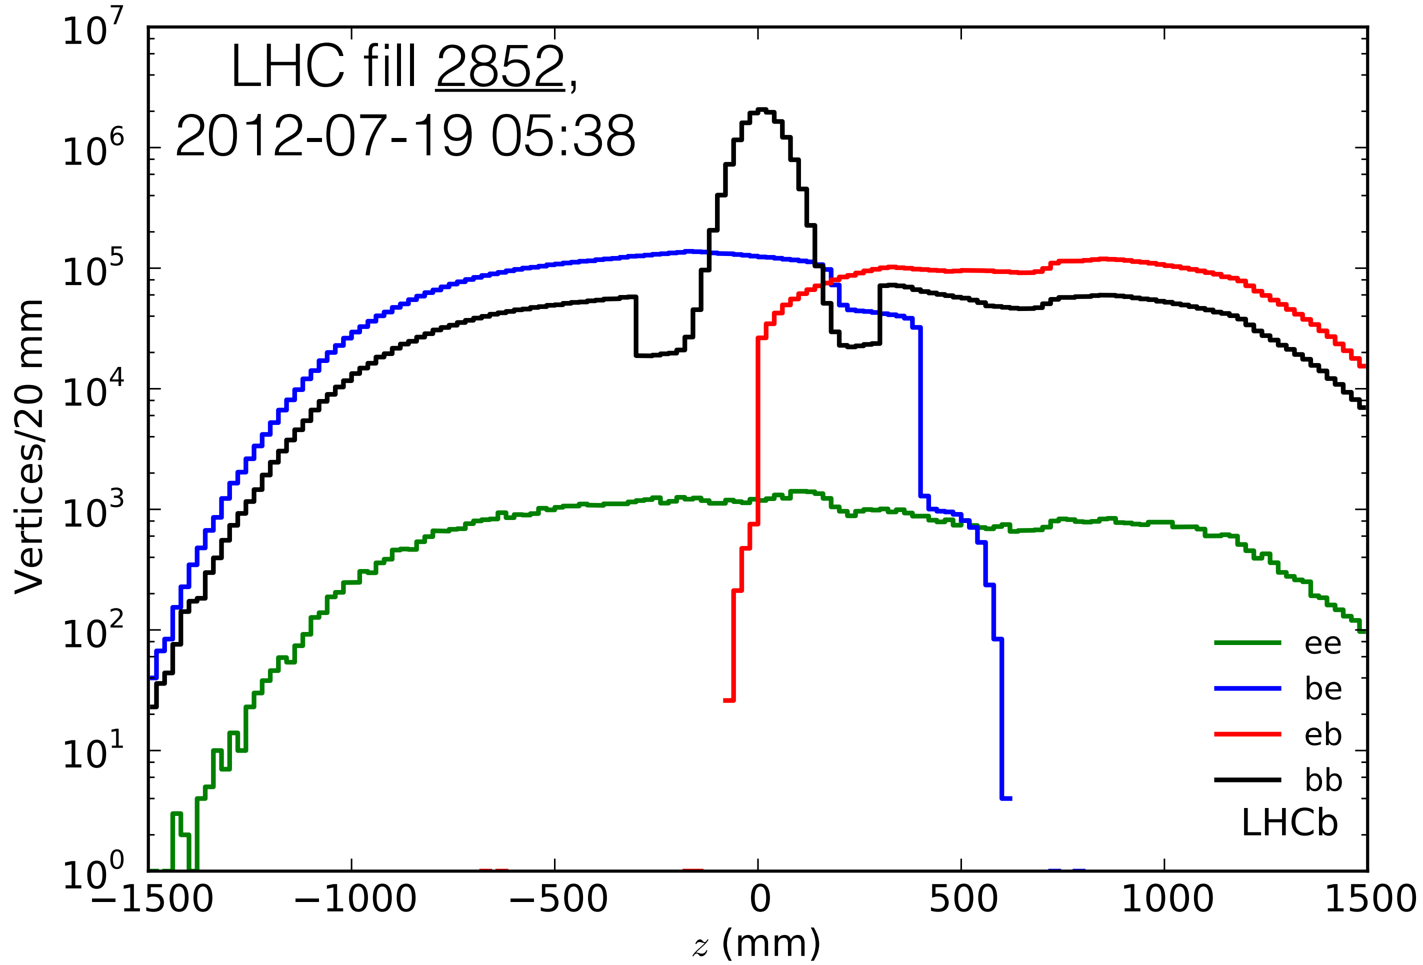

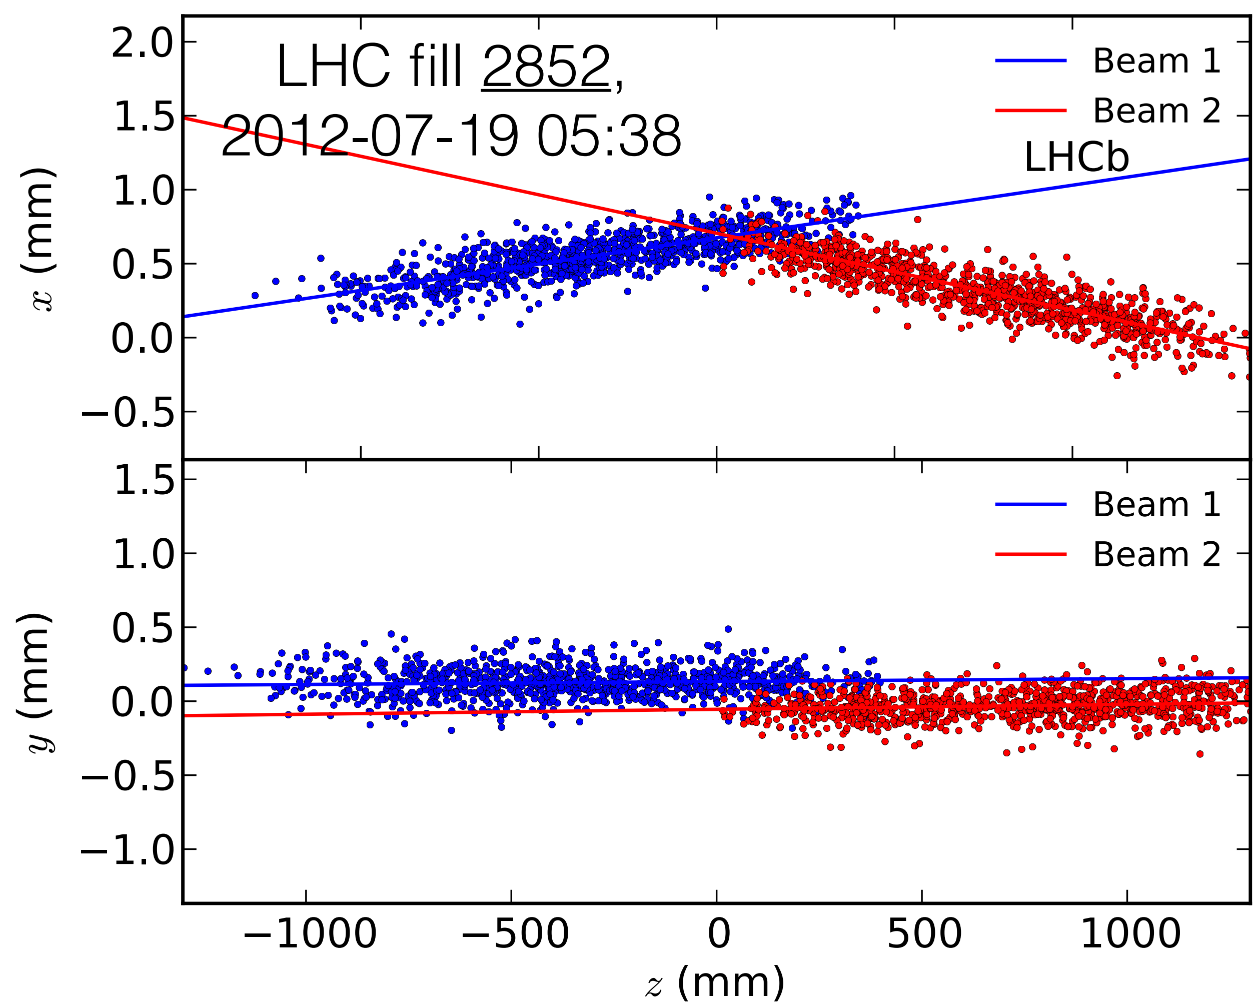

40 min of data taking

)

first 1000 vertices

Data is used to determine beam shape, size & overlap

==> provides a 1.4% calibration of luminosity

Charm @ LHCb

- Charm important part of LHCb physics:
  - Observation of D0-D0bar oscillations [PRL 110 (2013) 101802]
  - Measurement of D0-D0bar mixing parameters [PRL 111 (2013) 251801]
- High production rate, 600 kHz in 2012, requires *exclusive* selections.
- Exception: $D^{*+} \rightarrow D^0 \pi^+$
  - use D*-D0 mass difference to select $D^0 \rightarrow h^+h'^-$
  - Cabibbo favored ($D^0 \rightarrow K^-\pi^+$) rate is 300x suppressed rate ($D^0 \rightarrow \pi^-K^+$)

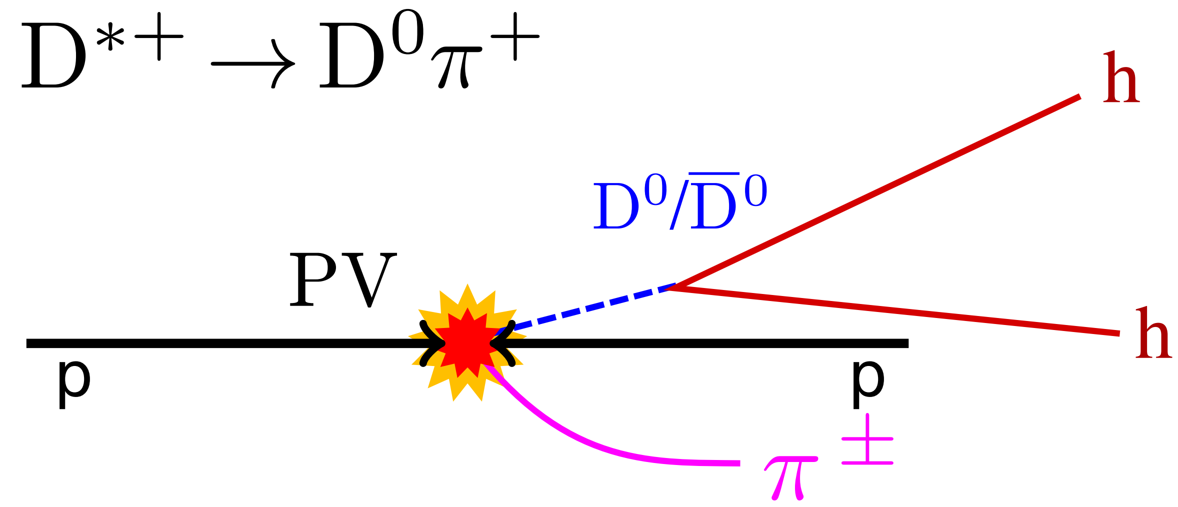

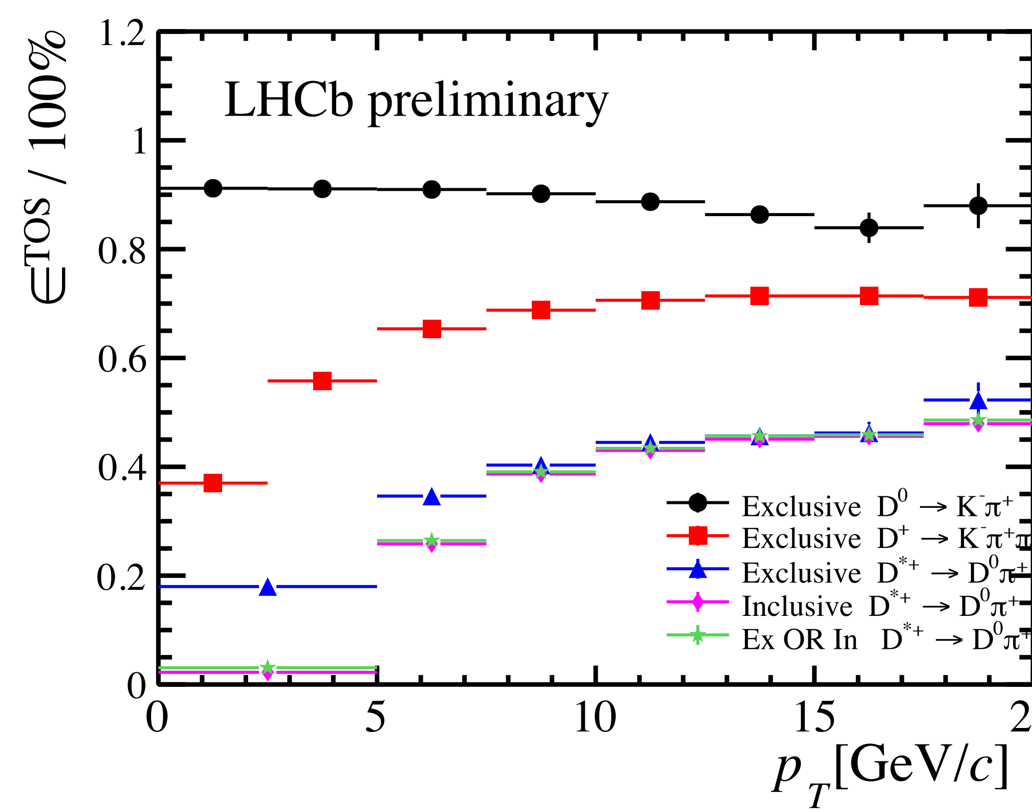

Interaction with Analysis

analysis preparation:

- setup/ optimize a trigger for your physics signal
  - define a trigger strategy (based on the available resources)
  - convert to trigger chain (already existing?)
  - determine rates and efficiencies from MC
- define a monitoring strategy
  - define trigger chain to be used for monitoring of your physics trigger (efficiency from data)
  - rates of the monitoring trigger (pre-scales?)
- integrate this in the overall trigger menu (done by Trigger Menu Coordination for online running)

threshold?
more exclusive?
pre-scaling?
more conditions?

not OK

OK

- use the trigger online (take data)
- monitor trigger quality
- determine trigger eff. (from data)
- correct your measurement

Triggering on Charm

many (excited) charmed hadrons, huge production rate (20xsigma_b), small mass

but typically smallish *interesting* branching ratios

—> go exclusive!

- implies low momentum tracking & particle ID
- which is expensive
- and requires good calibrations

but if you go down this road:

- can reject large fraction of favored decays while keeping the suppressed ones
- use trigger candidates for analysis —> turbo!
- and drop offline processing!

8.4M events

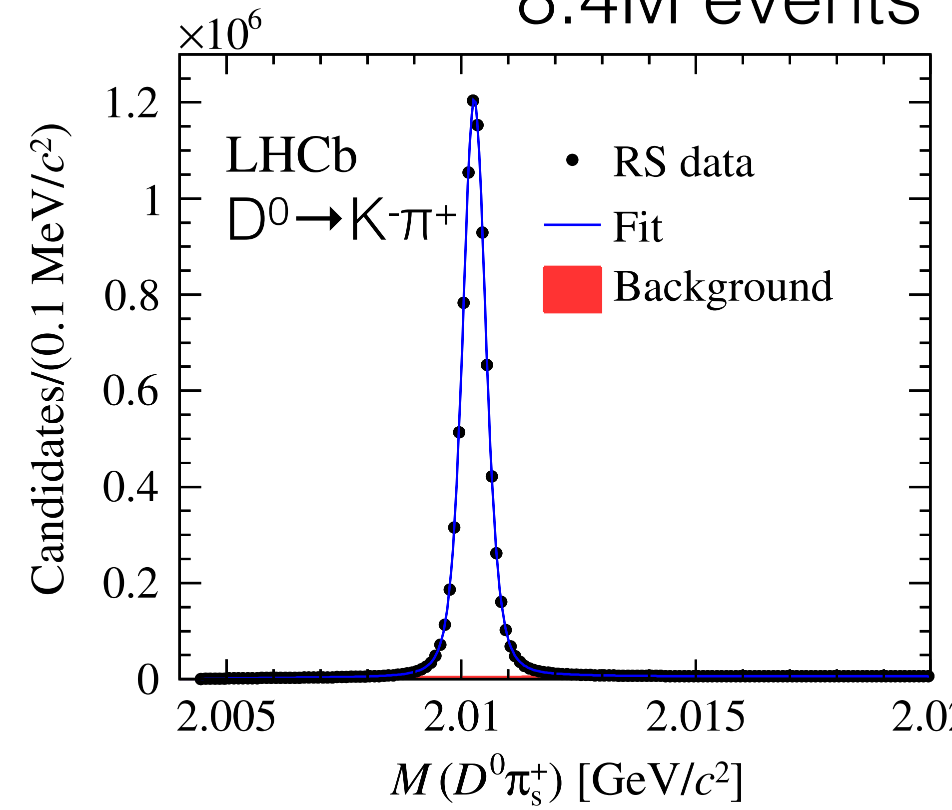

3.6K events

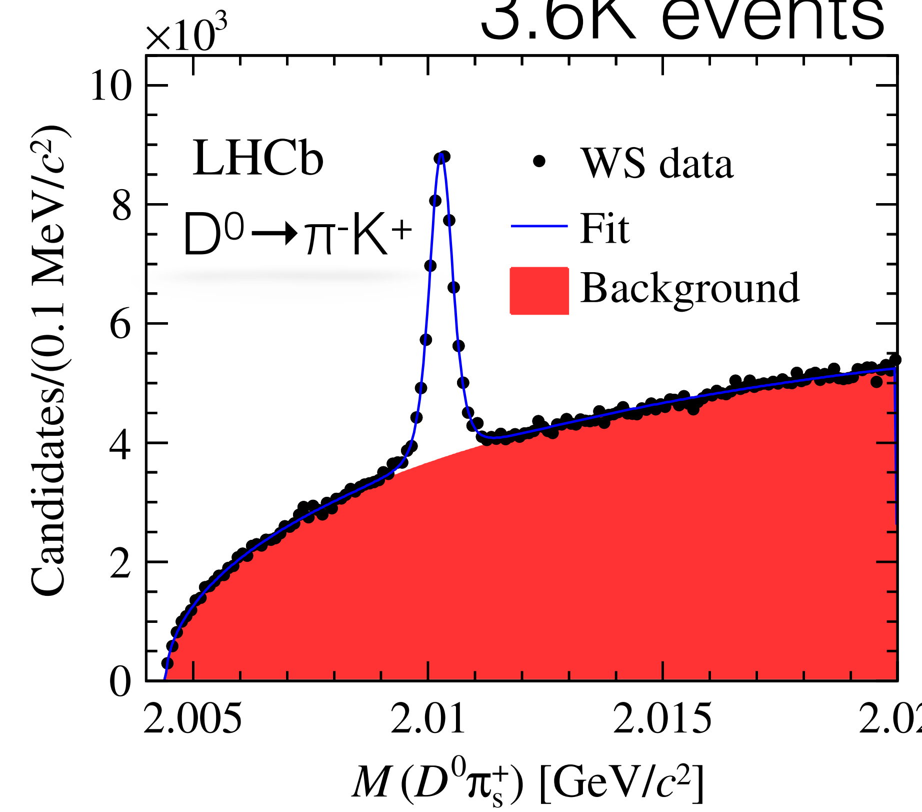

Enabling online calibrations

LHCb 2012

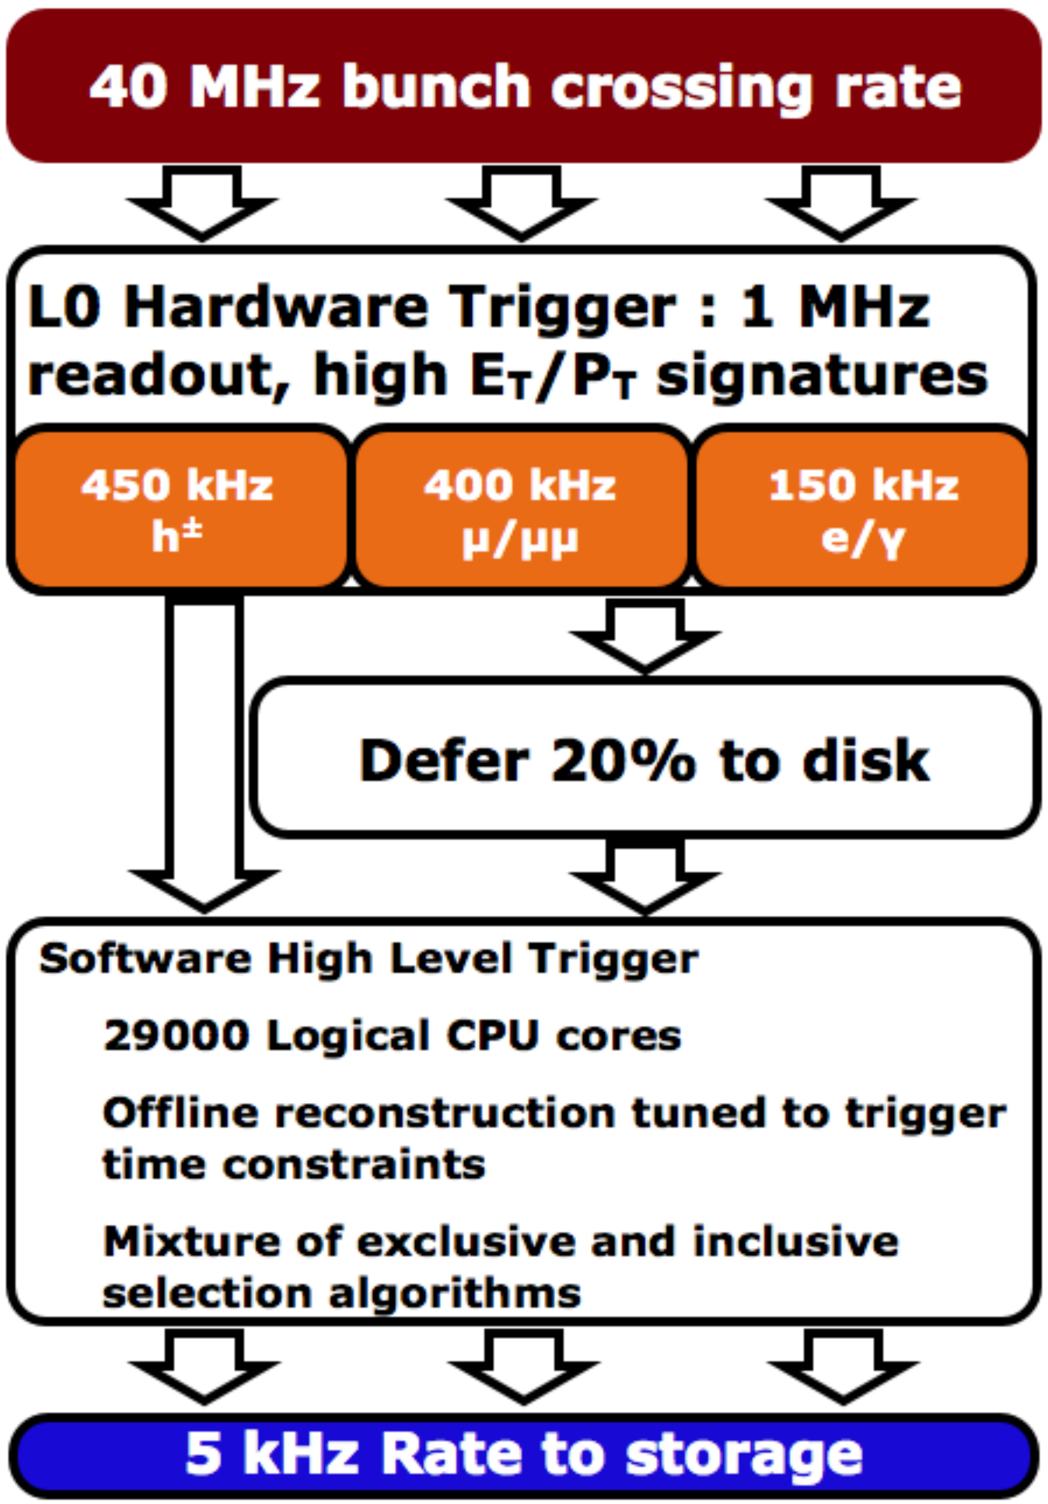

LHCb 2015 Trigger Diagram

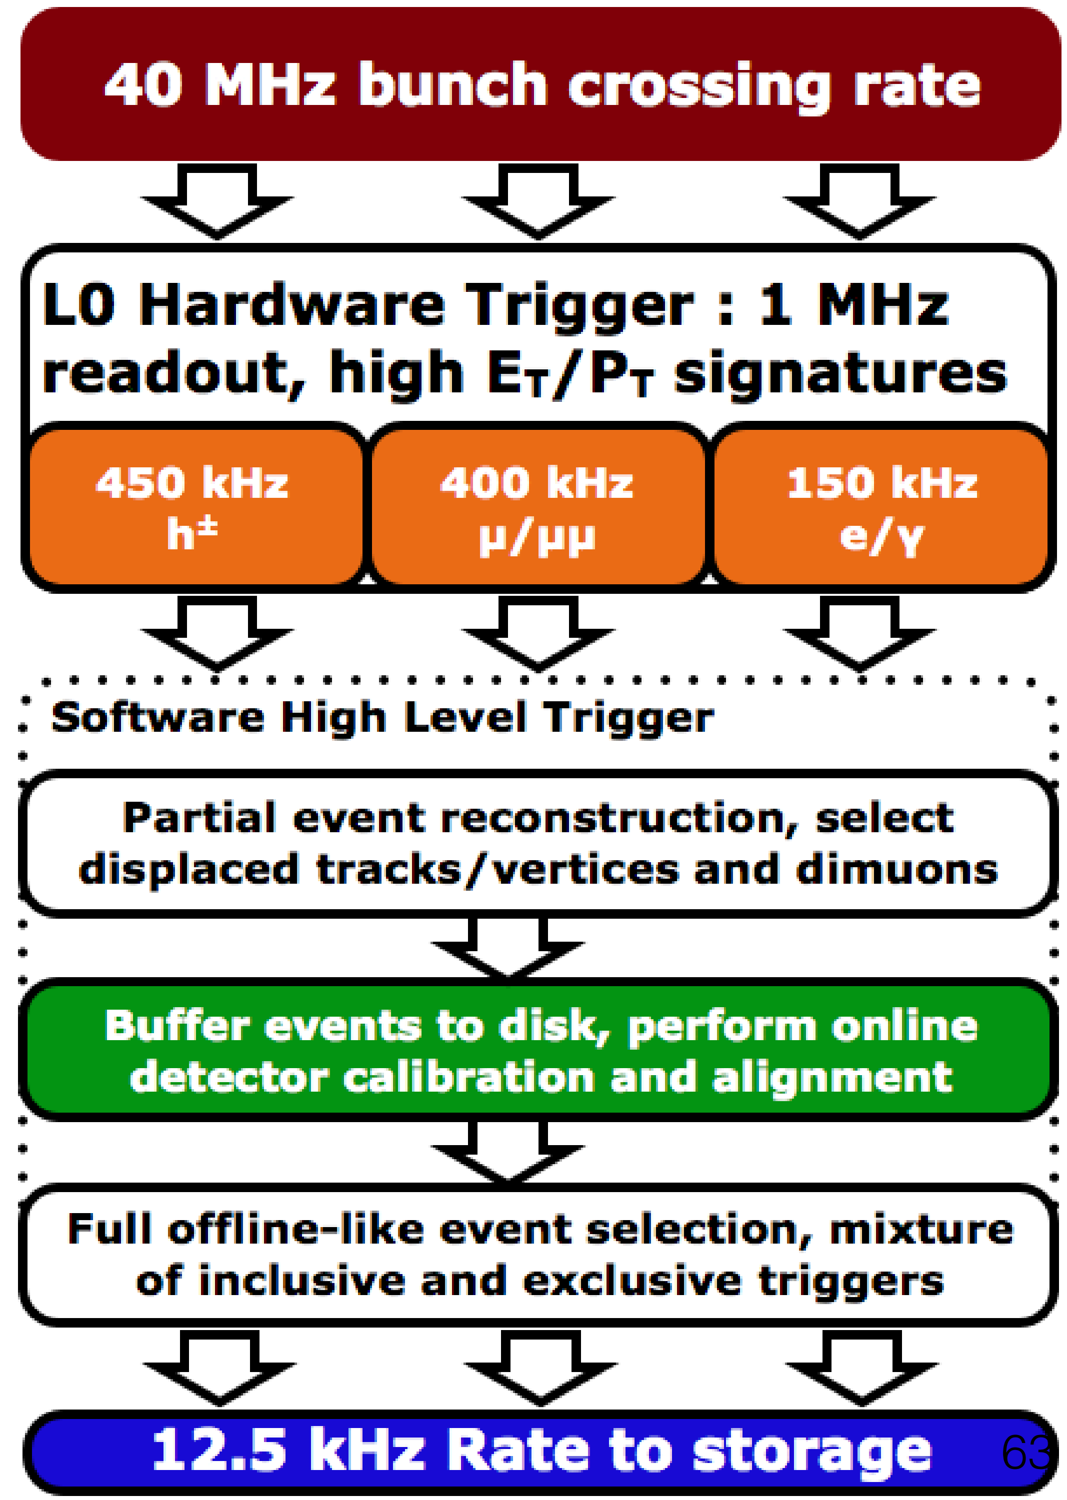

Real-time calibrations

- minimizes online-offline differences
- more stable performance
- allows ‘sharper’ HLT selections
- enables physics analysis directly on HLT output —> turbo!

First alignment

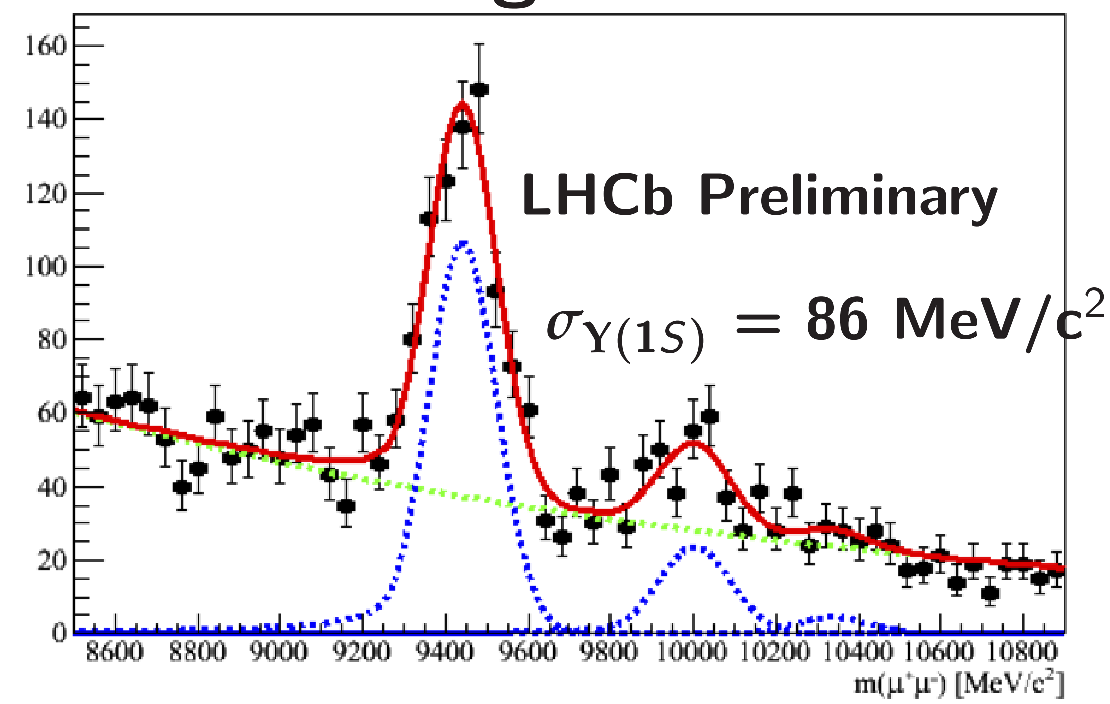

- eliminates one round of offline-reconstruction
  - less offline grid-usage & less storage —> increase HLT output bandwidth!
- calibrations become part of the workflow at the pit
  - better documented, more automated

Improved alignment

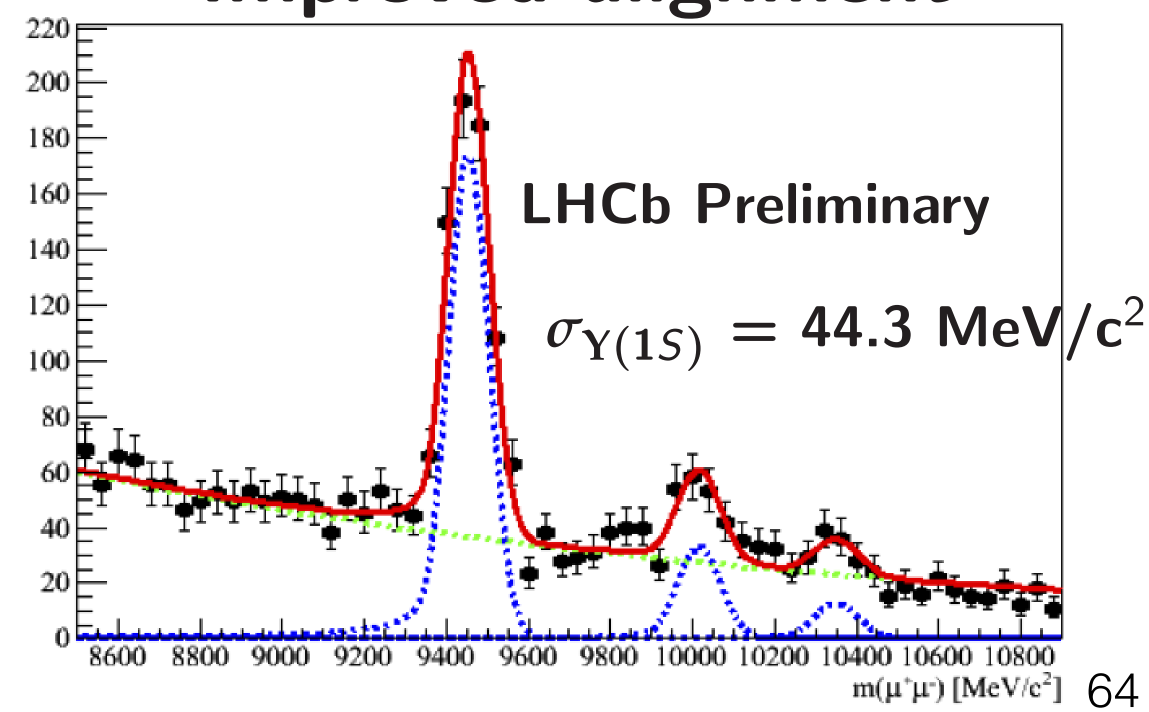

LHCb ‘turbo stream’

$D^0 \rightarrow K^-\pi^+$

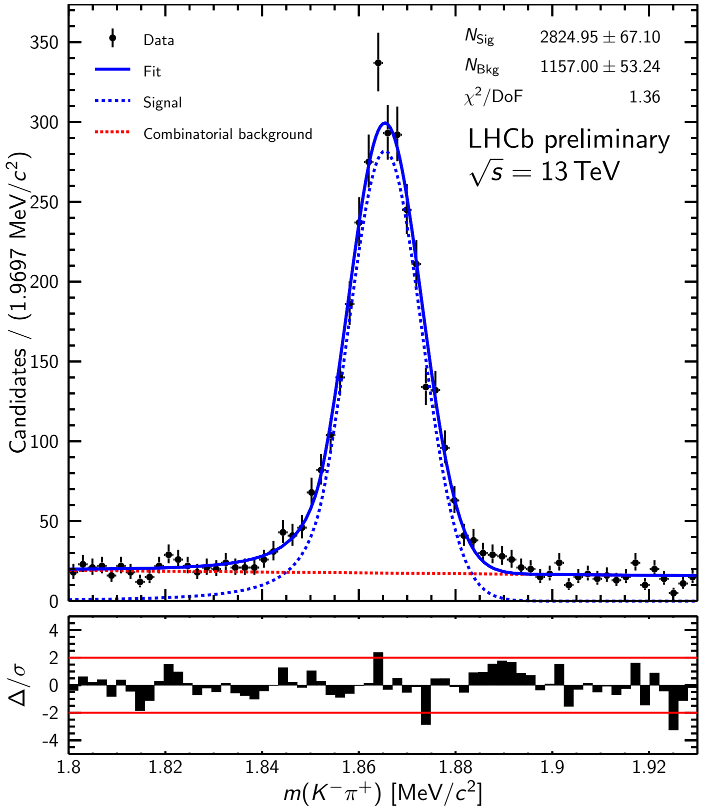

$D^- \rightarrow K^-\pi^+\pi^-$

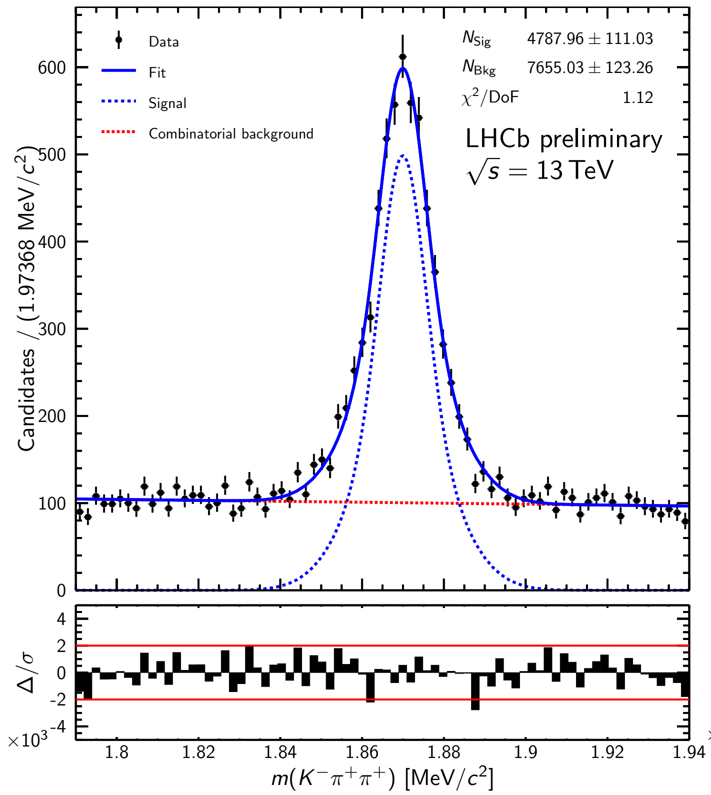

$D^+ \rightarrow K^-K^+\pi^+$

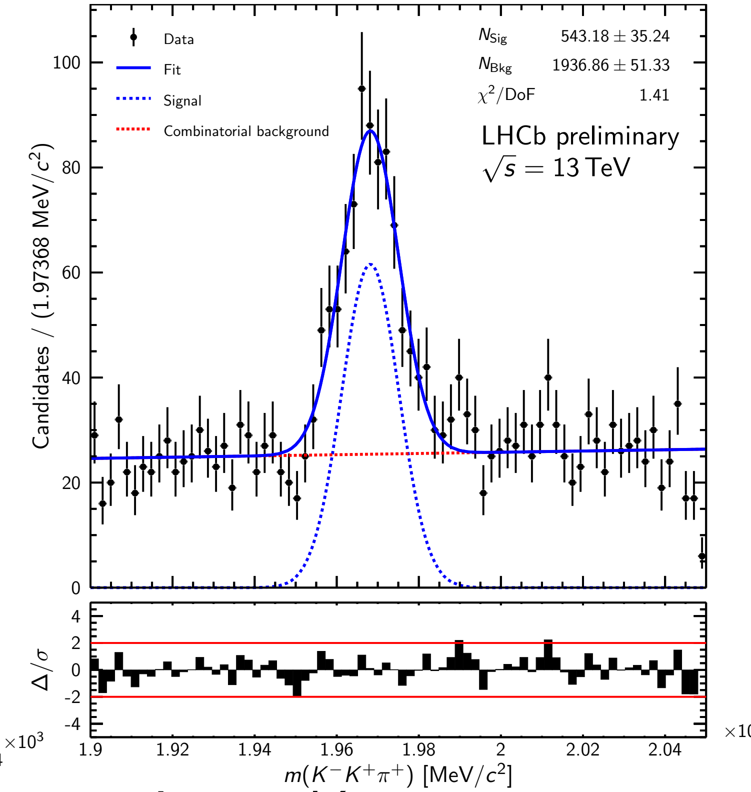

$D^{*+} \rightarrow D^0\pi^+_s \rightarrow (K^-\pi^+)\pi^+_s$

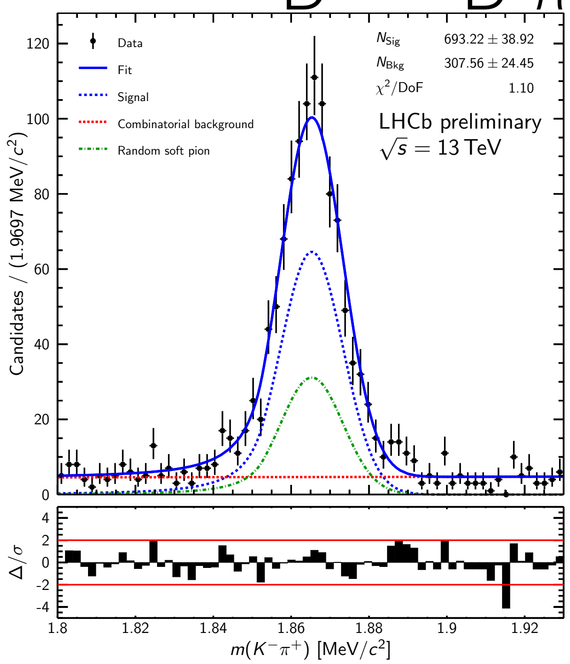

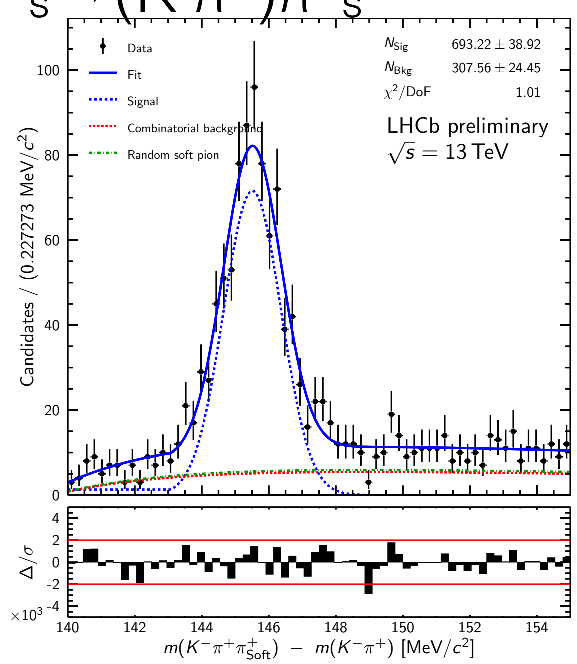

$\Lambda_c \rightarrow pK^+\pi^+$

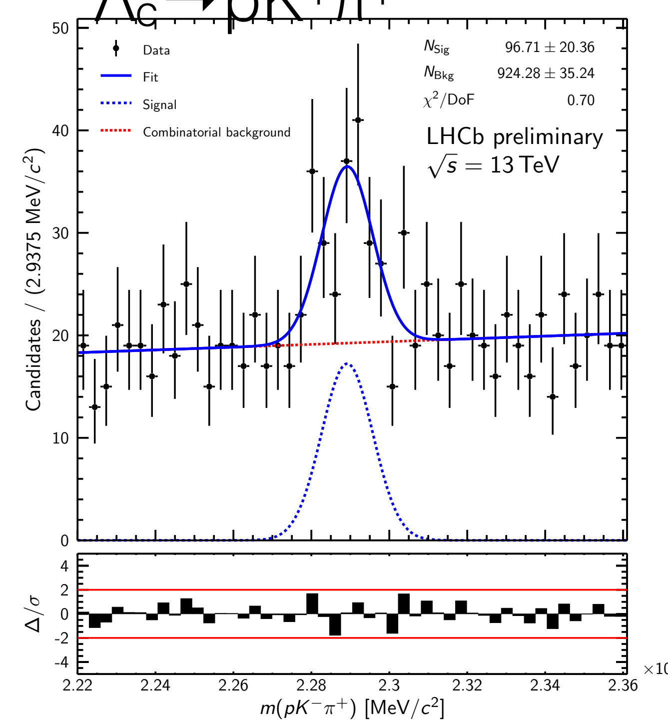

Early June 2015 — clear charm in first “physics” fill @ 13 TeV

LHCb ‘turbo stream’

$D^0 \to K^-\pi^+$

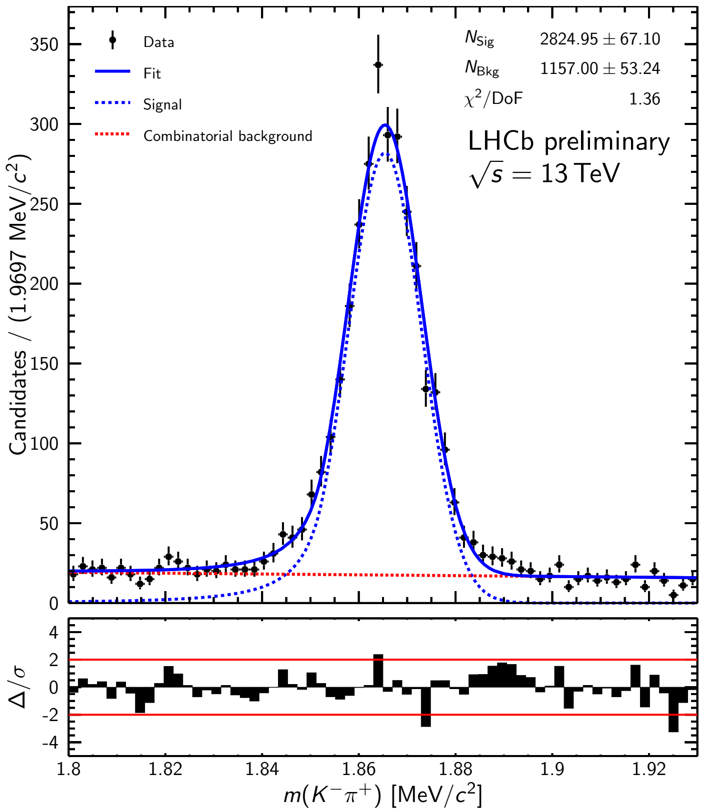

$D^- \to K^-\pi^+\pi^-$

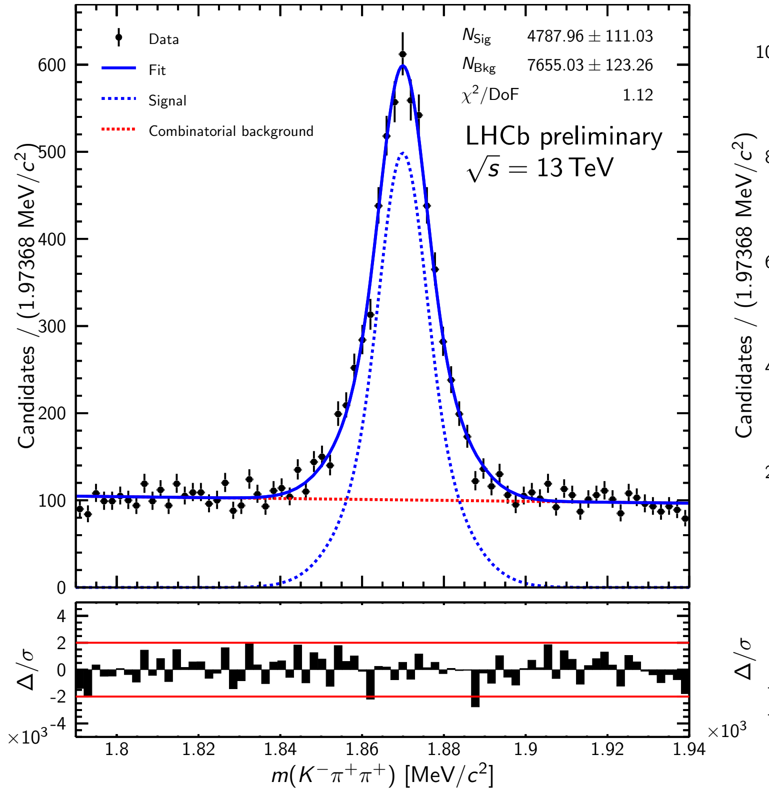

$D^+ \to K^-K^+\pi^+$

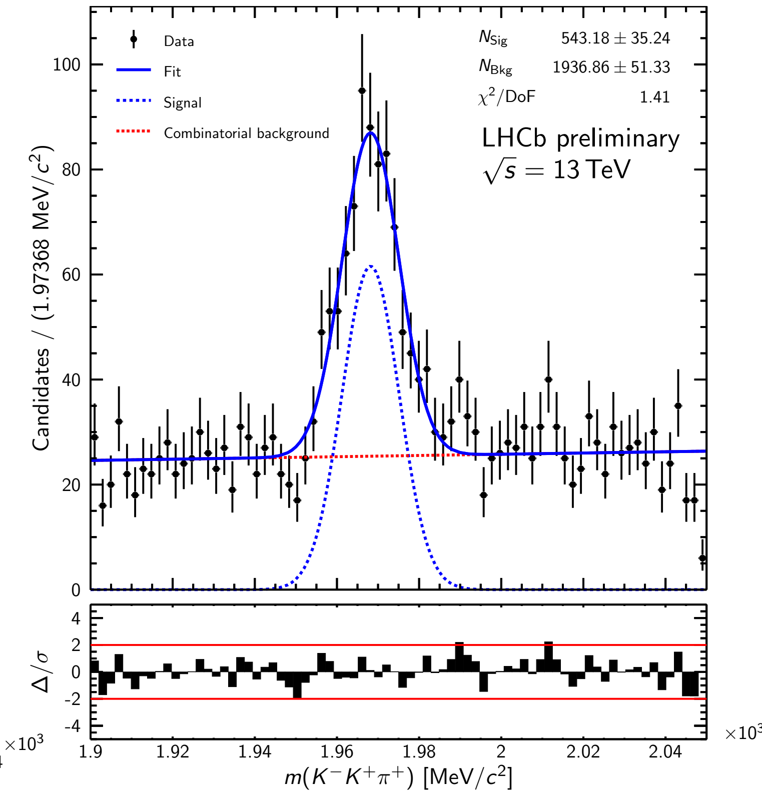

- As *complete* signal is found by the HLT, provide ‘just’ the HLT output for analysis.
  - Must insure *all* information required for analysis is provided by HLT output
  - First iteration ‘now’: let’s see if we can publish a ‘turbo analysis’ or that we missed something…
- Initially keep ‘raw event’ — just in case we forgot something — keep it ‘parked’
  - consider it ‘parking with a preview’
- If it works, only keep HLT output (like the CMS ‘scouting’) ; smaller events —> higher rate.

Interaction with Analysis

analysis preparation:

- setup/ optimize a trigger for your physics signal
  - define a trigger strategy (based on the available resources)
  - convert to trigger chain (already existing?)
  - determine rates and efficiencies from MC
- define a monitoring strategy
  - define trigger chain to be used for monitoring of your physics trigger (efficiency from data)
  - rates of the monitoring trigger (pre-scales?)
- integrate this in the overall trigger menu (done by Trigger Menu Coordination for online running)

threshold?
more exclusive?
pre-scaling?
more conditions?

not OK

OK

- use the trigger online (take data)
- monitor trigger quality
- determine trigger eff. (from data)
- correct your measurement

Determining efficiencies: Tag and Probe”

- Assuming a single-muon trigger
- Try to identify a muon *without* using the muon system
- Reconstruct a dimuon resonance with one identified muon (‘tag’) + one track (‘probe’)
    - require $M_{tag+probe} \sim M_Z$
- Now check if trigger *also* recognized this ‘probe’
- backgrounds can be subtracted using the dimuon invariant mass spectrum

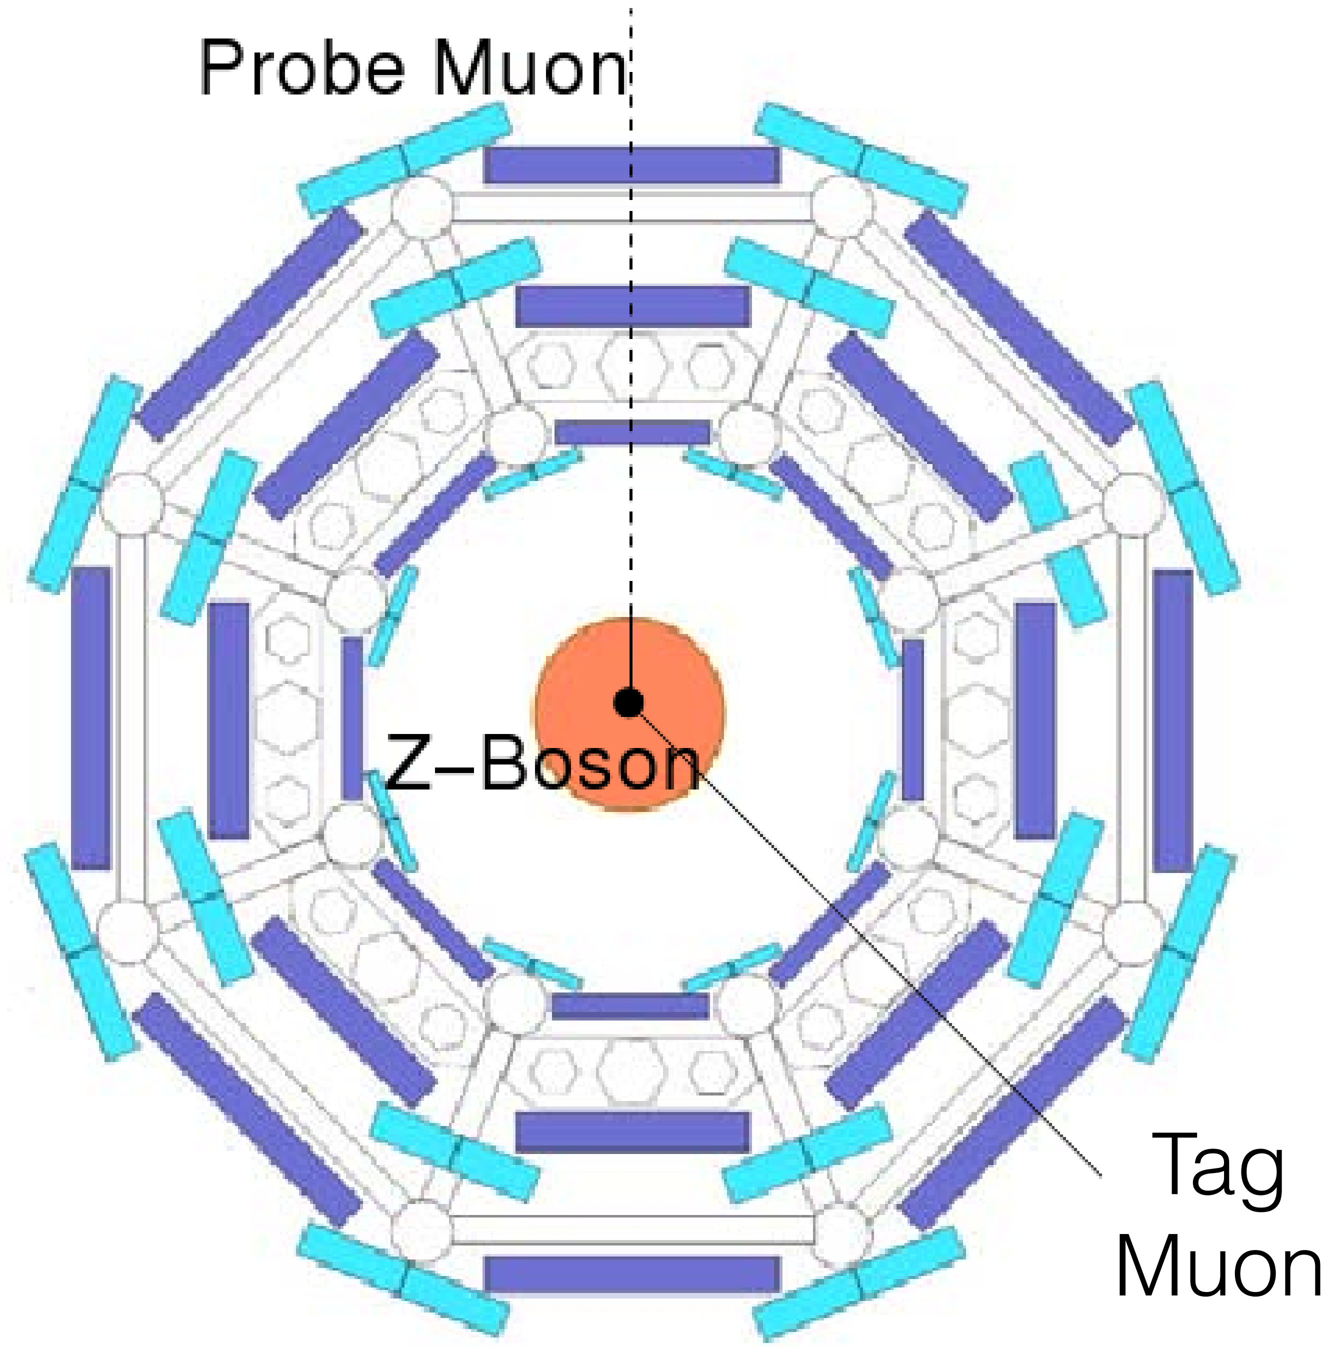

Determining efficiencies: Tag and Probe”

$N_{TT} = N\,\varepsilon^2$

$N_{TP} = N\,2\,\varepsilon\,(1-\varepsilon)$

$\Rightarrow \varepsilon = 2\,N_{TT} / (N_{TP}+2N_{TT})$

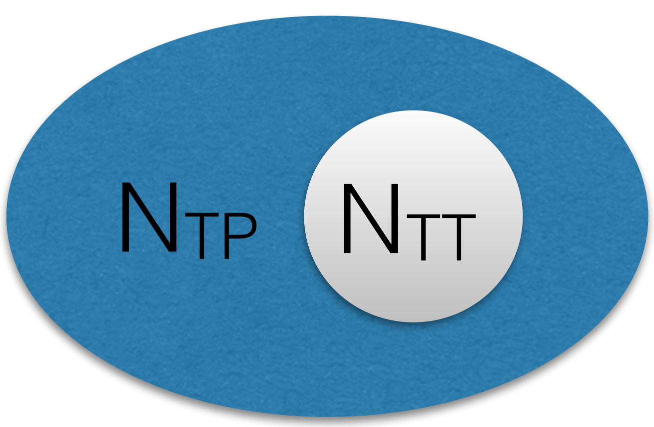

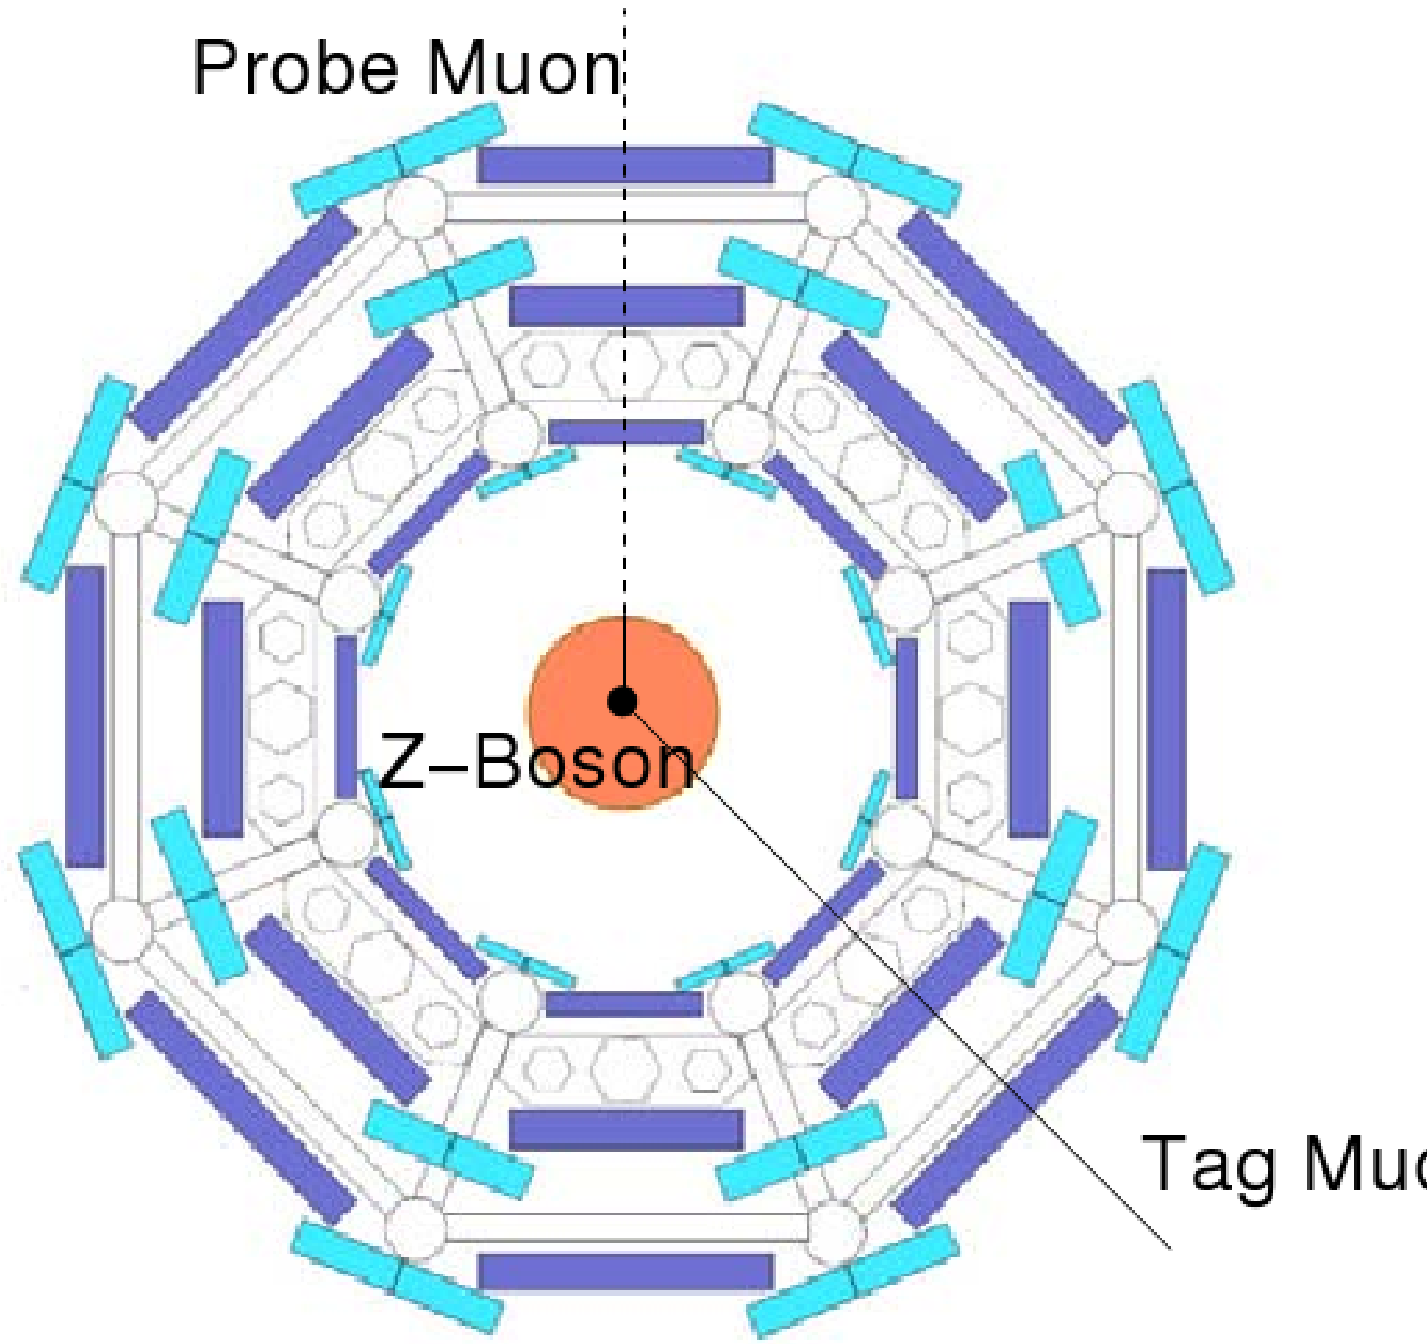

Determining efficiencies: Tag and Probe”

$N_{TT} = N\, \varepsilon^2$

$N_{TP} = N\, 2\, \varepsilon\, (1-\varepsilon)$

$\Rightarrow \varepsilon = 2\, N_{TT} / (N_{TP}+2N_{TT})$

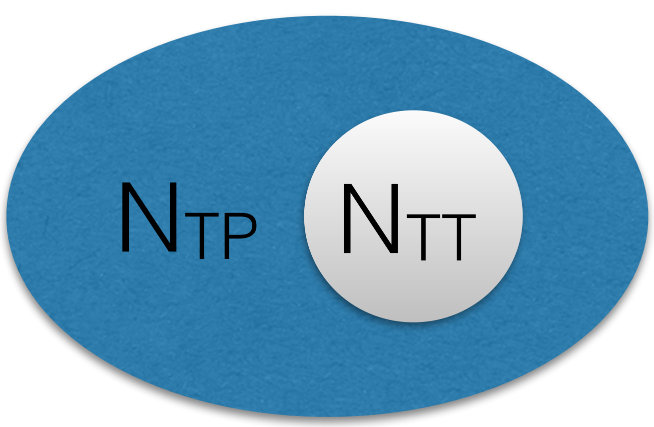

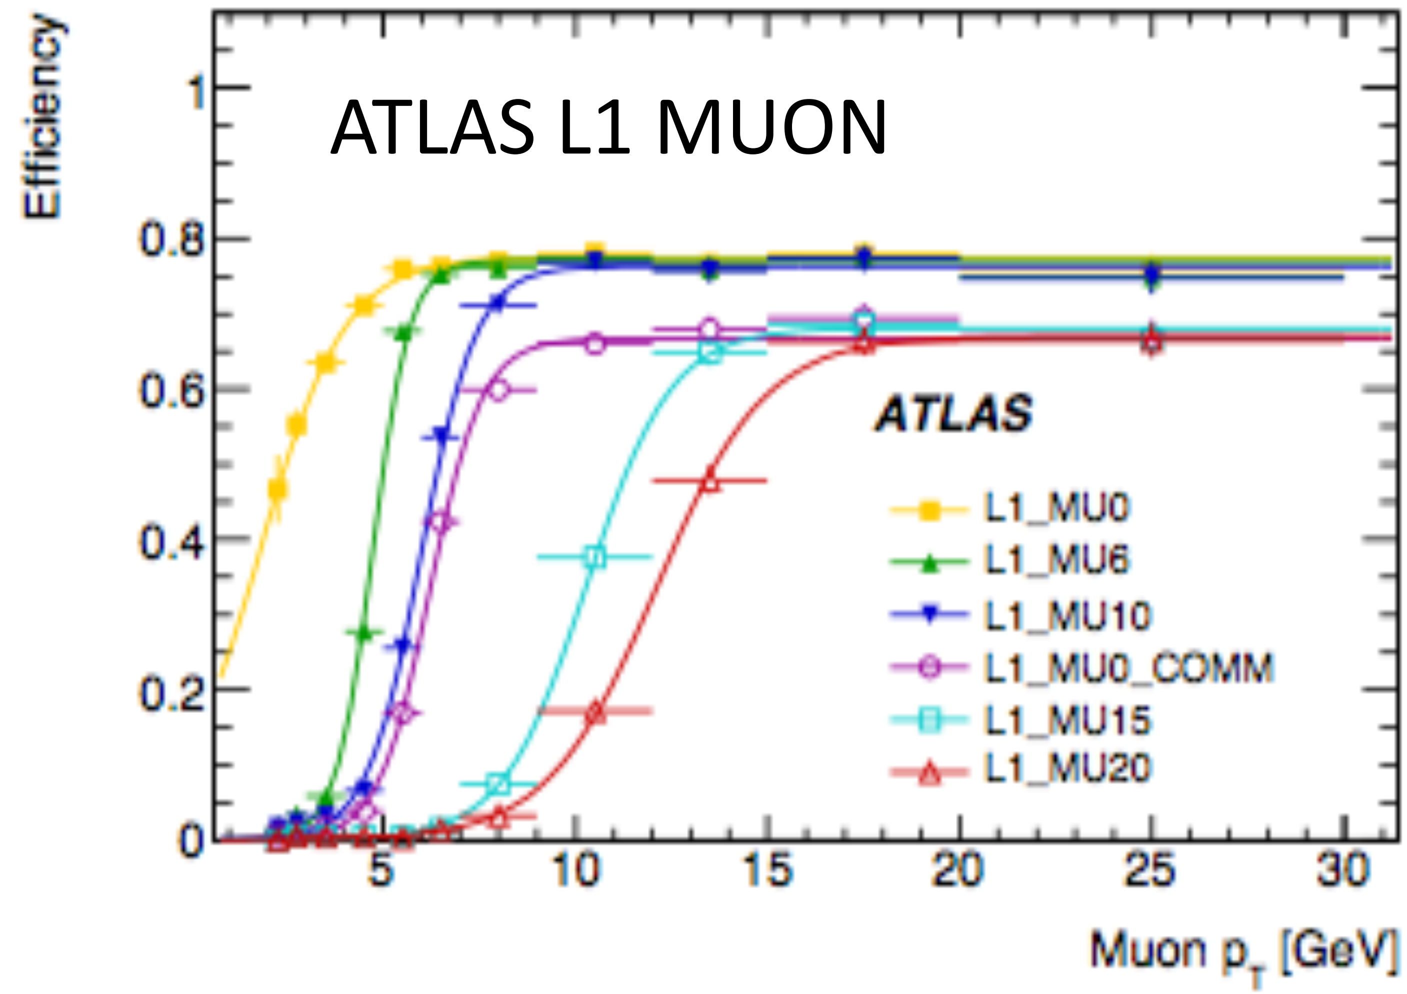

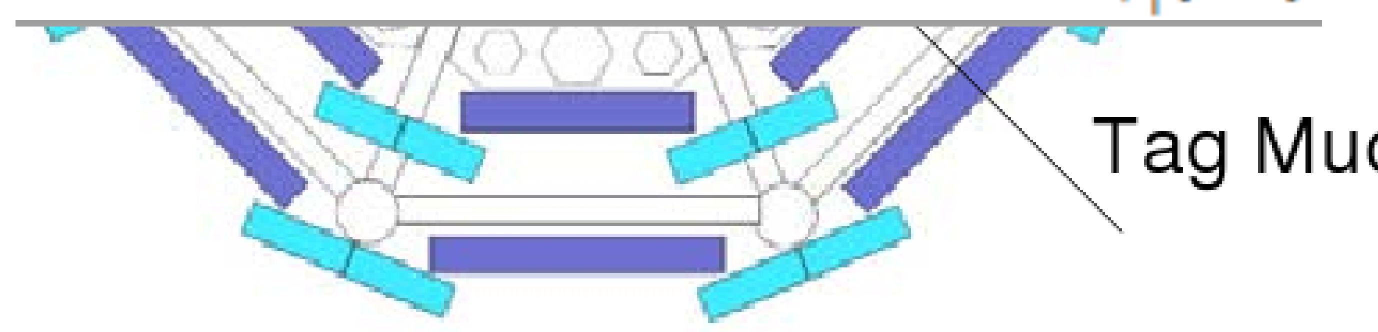

Determining efficiencies: Tag and Probe"

$N_{TT} = N\,\varepsilon^2$
$N_{TP} = N\,2\,\varepsilon\,(1-\varepsilon)$
$\Rightarrow \varepsilon = 2\,N_{TT} / (N_{TP}+2N_{TT})$

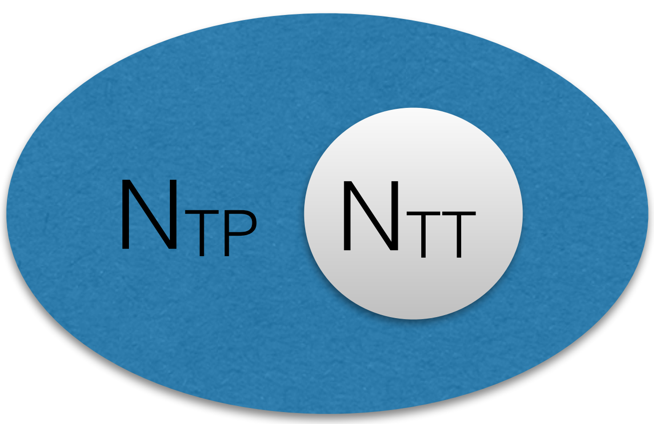

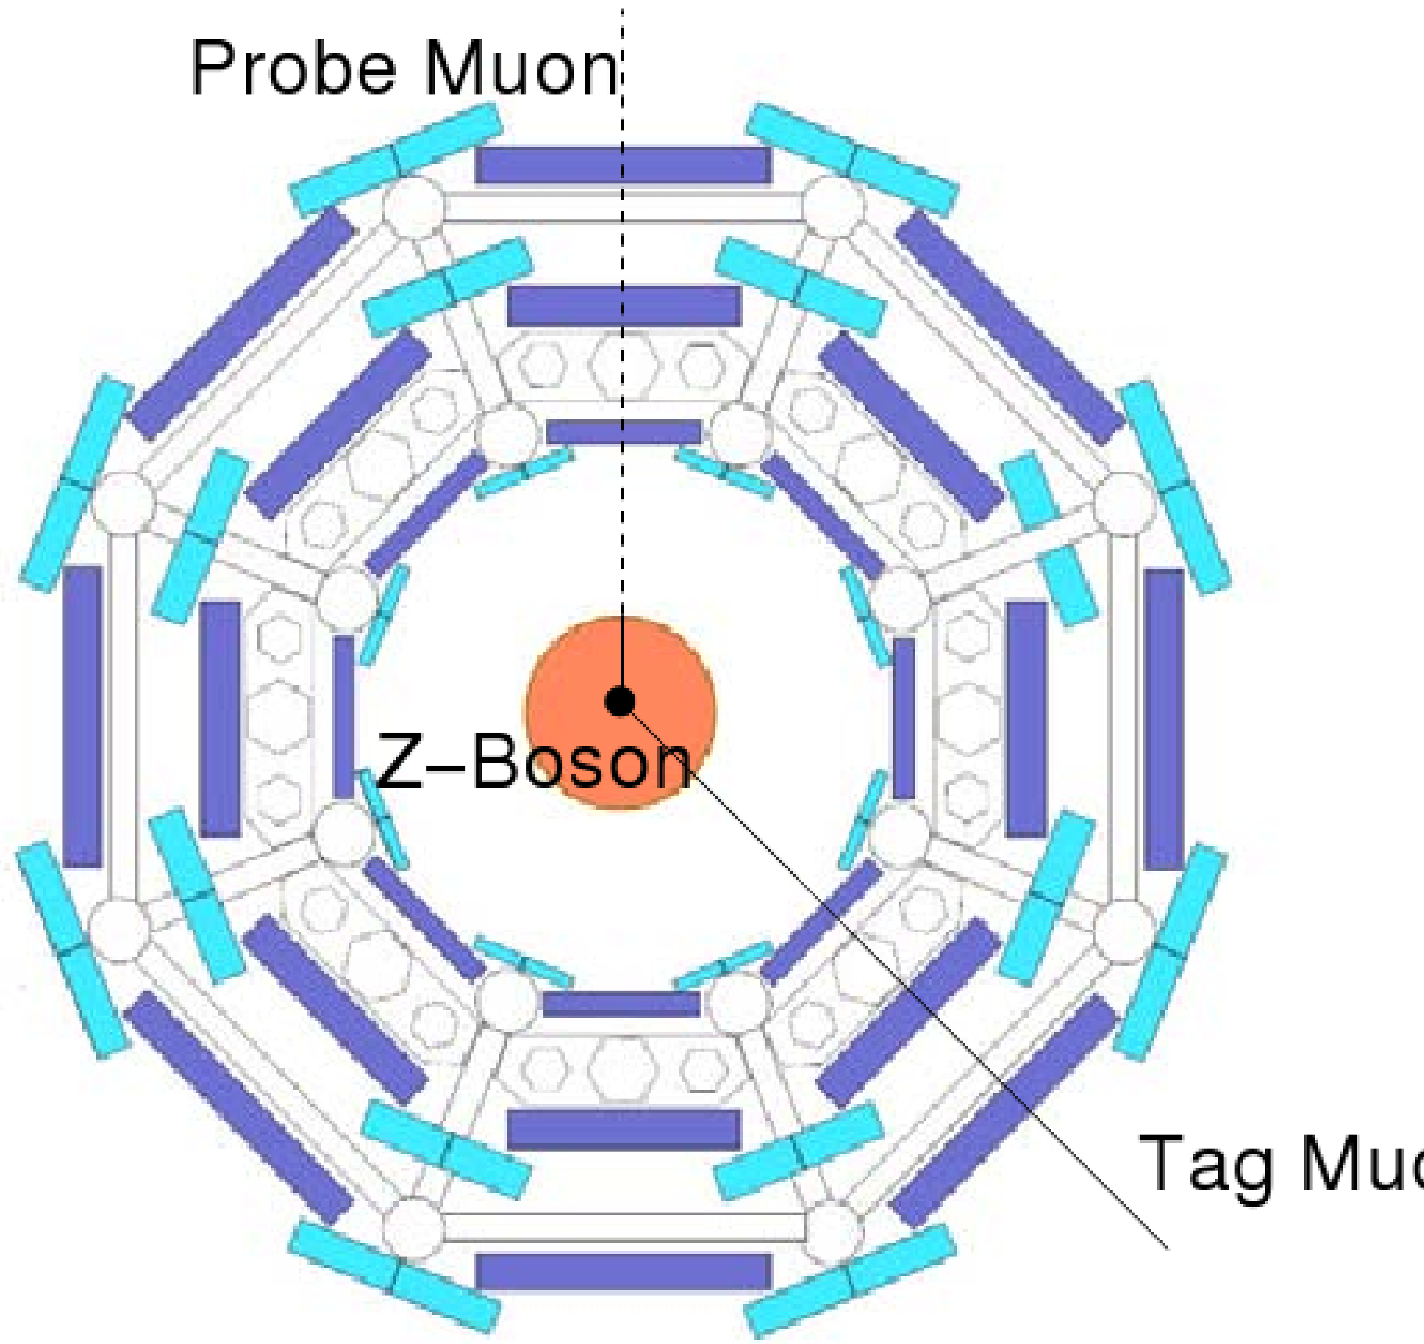

formally: conditional efficiency of getting a 2nd (tag) muon given that at least one was identified (as tag) assuming identifying a 2nd muon is uncorrelated from the 1st, for the given sample of Z→μμ decays

Determining efficiencies: LHCb $B_s \rightarrow J/\psi\ \phi$

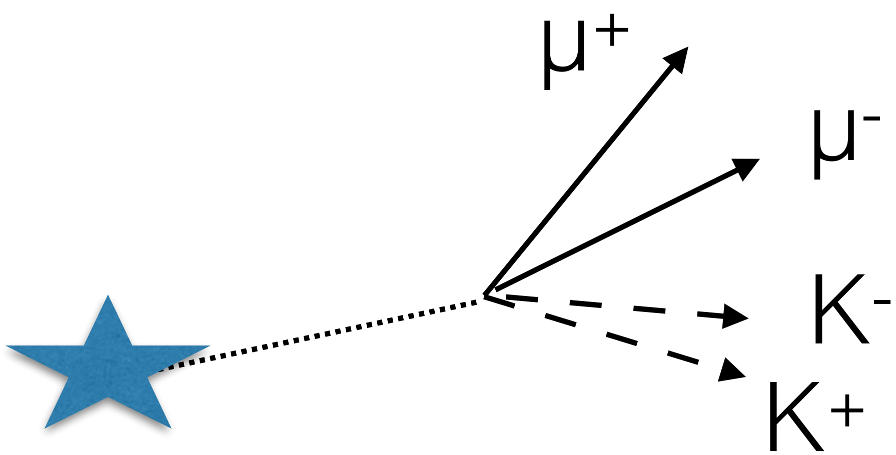

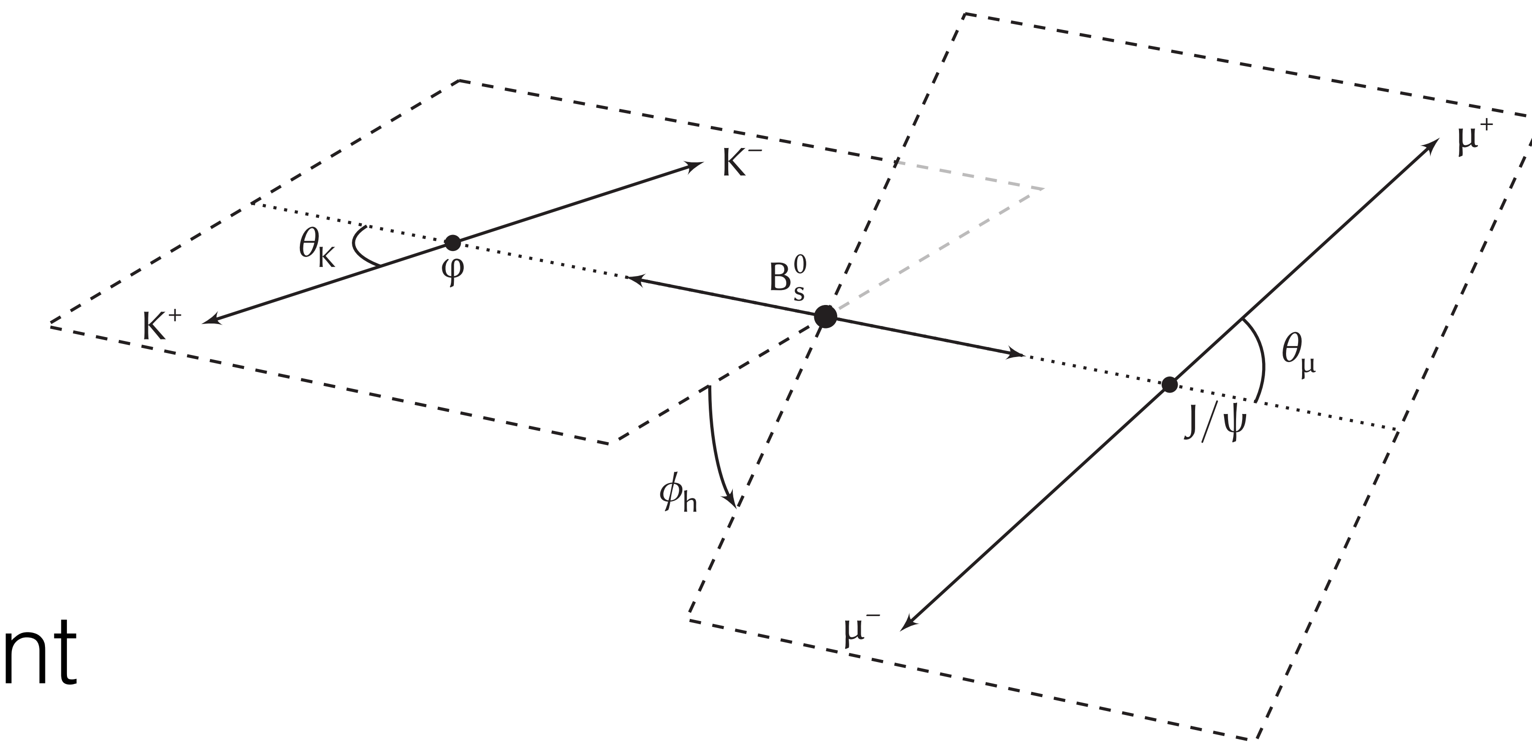

- CP violation measurement
  - must measure (amongst others) proper time, decay angles
- CDF / D0 collected 9k signal events each in ~7 years
- LHCb collected 90K signal events in ~2 years

Interaction with Analysis

analysis preparation:

- setup/ optimize a trigger for your physics signal
  - define a trigger strategy (based on the available resources)
  - convert to trigger chain (already existing?)
  - determine rates and efficiencies from MC
- define a monitoring strategy
  - define trigger chain to be used for monitoring of your physics trigger (efficiency from data)
  - rates of the monitoring trigger (pre-scales?)
- integrate this in the overall trigger menu (done by Trigger Menu Coordination for online running)

threshold?
more exclusive?
pre-scaling?
more conditions?

not OK

OK

- use the trigger online (take data)
- monitor trigger quality
- determine trigger eff. (from data)
- correct your measurement

Triggering $B_s \to J/\psi\ \phi$

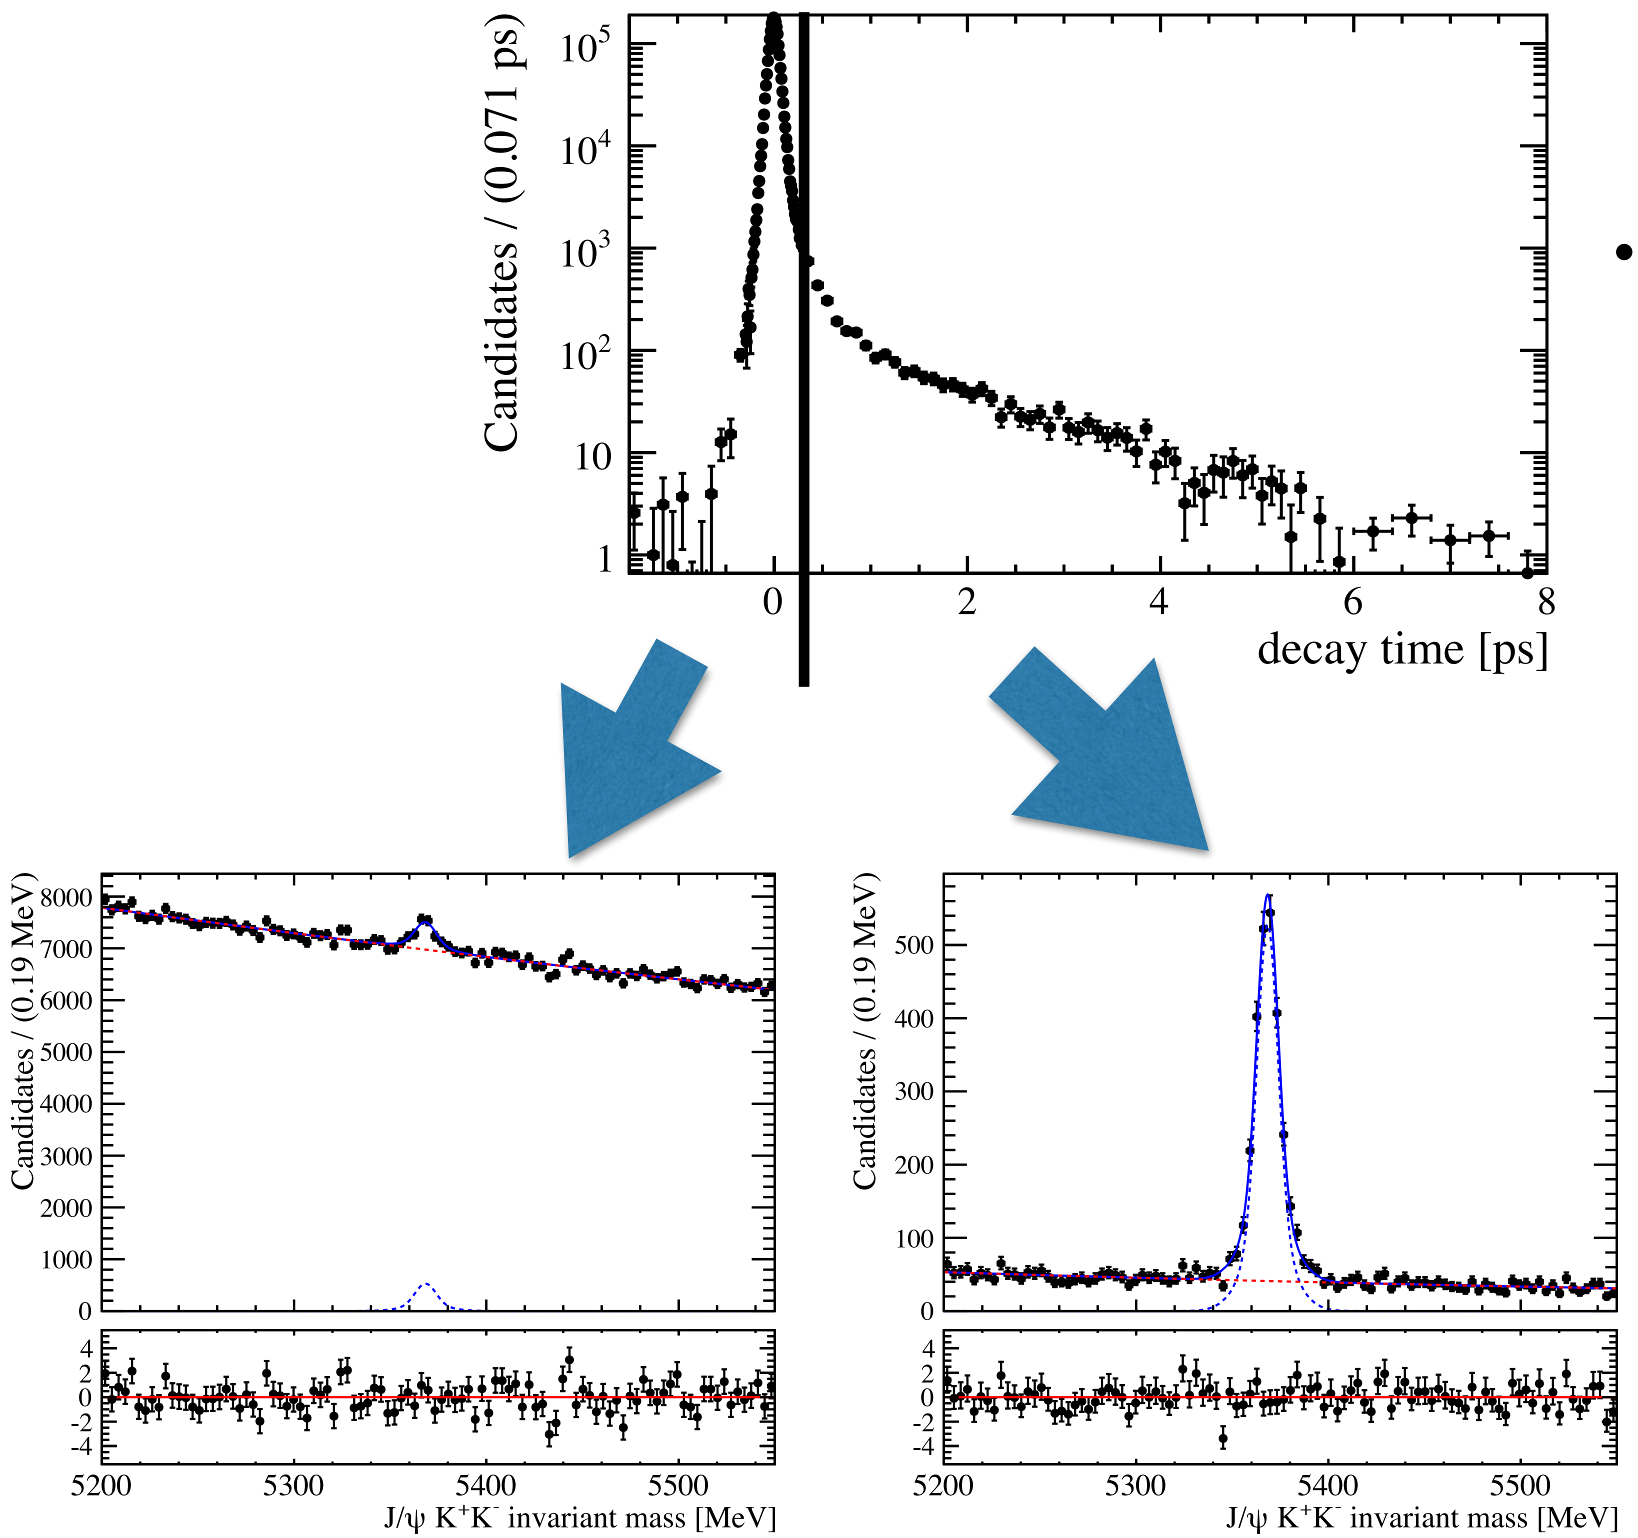

- CDF / D0 used a ‘dimuon trigger’
- does not ‘distort’ the decay time distribution
- In LHCb, lots of $J/\psi \to \mu\mu$ produced—> ‘(di)muon’ trigger consumes significant bandwidth
- How to reduce this bandwidth?
    - Argue in meetings that your analysis is very important, demand larger fraction of bandwidth…
    - prescale…
    - use more powerful (and complicated) methods, eg. ANN, BDT, … — typically requires ‘exclusive’ selection’
    - add P_T cuts on the J/ψ or muons — distorts the angular distribution
    - require ‘detached J/ψ’ — distorts decay time distribution

Interaction with Analysis

analysis preparation:

- setup/ optimize a trigger for your physics signal
  - define a trigger strategy (based on the available resources)
  - convert to trigger chain (already existing?)
  - determine rates and efficiencies from MC
- define a monitoring strategy
  - define trigger chain to be used for monitoring of your physics trigger (efficiency from data)
  - rates of the monitoring trigger (pre-scales?)
- integrate this in the overall trigger menu (done by Trigger Menu Coordination for online running)

threshold?
more exclusive?
pre-scaling?
more conditions?

not OK

OK

- use the trigger online (take data)
- monitor trigger quality
- determine trigger eff. (from data)
- correct your measurement

Use of multiple triggers

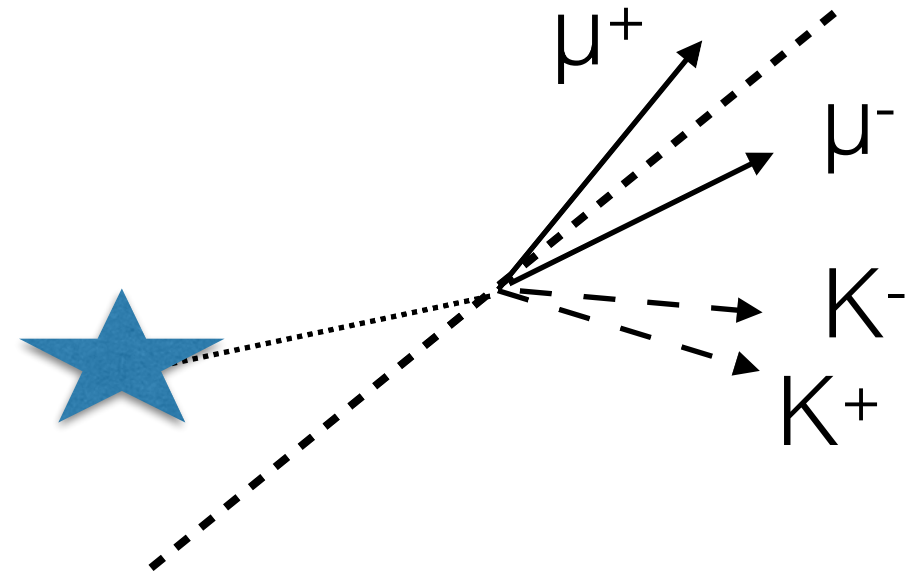

- Take existing dimuon trigger, add prescale → ε_p
- Take existing dimuon trigger, add detached requirement → $\varepsilon_d(t)$

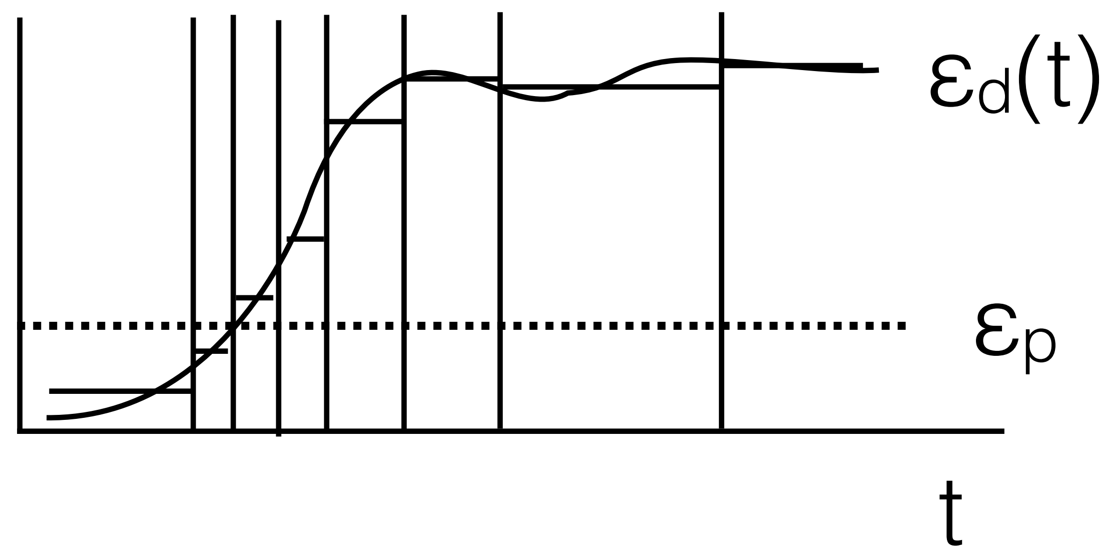

- $\varepsilon_p = N_p / (N_p + N_d)$
- $\varepsilon_d = N_d / (N_p + N_d)$
- actually, $\varepsilon_d = \varepsilon_d$ *given* p, and assuming ε_p is constant — which is it isn't due to non-trivial tracking efficiency…
- actually: handle background
- actually: there are two levels of HLT and multiple lines / level

for each time interval:

	P	D	
P & !D	ε_P	$1-\varepsilon_D(t)$	N_p
!P & D	$1-\varepsilon_P$	$\varepsilon_D(t)$	N_d
P & D	ε_P	$\varepsilon_D(t)$	N_{pd}
!P & !D	$1-\varepsilon_P$	$1-\varepsilon_D(t)$	?

Interaction with Analysis

analysis preparation:

- setup/ optimize a trigger for your physics signal
  - define a trigger strategy (based on the available resources)
  - convert to trigger chain (already existing?)
  - determine rates and efficiencies from MC
- define a monitoring strategy
  - define trigger chain to be used for monitoring of your physics trigger (efficiency from data)
  - rates of the monitoring trigger (pre-scales?)
- integrate this in the overall trigger menu (done by Trigger Menu Coordination for online running)

threshold?
more exclusive?
pre-scaling?
more conditions?

not OK

OK

- use the trigger online (take data)
- monitor trigger quality
- determine trigger eff. (from data)
- correct your measurement

Hlt1Track

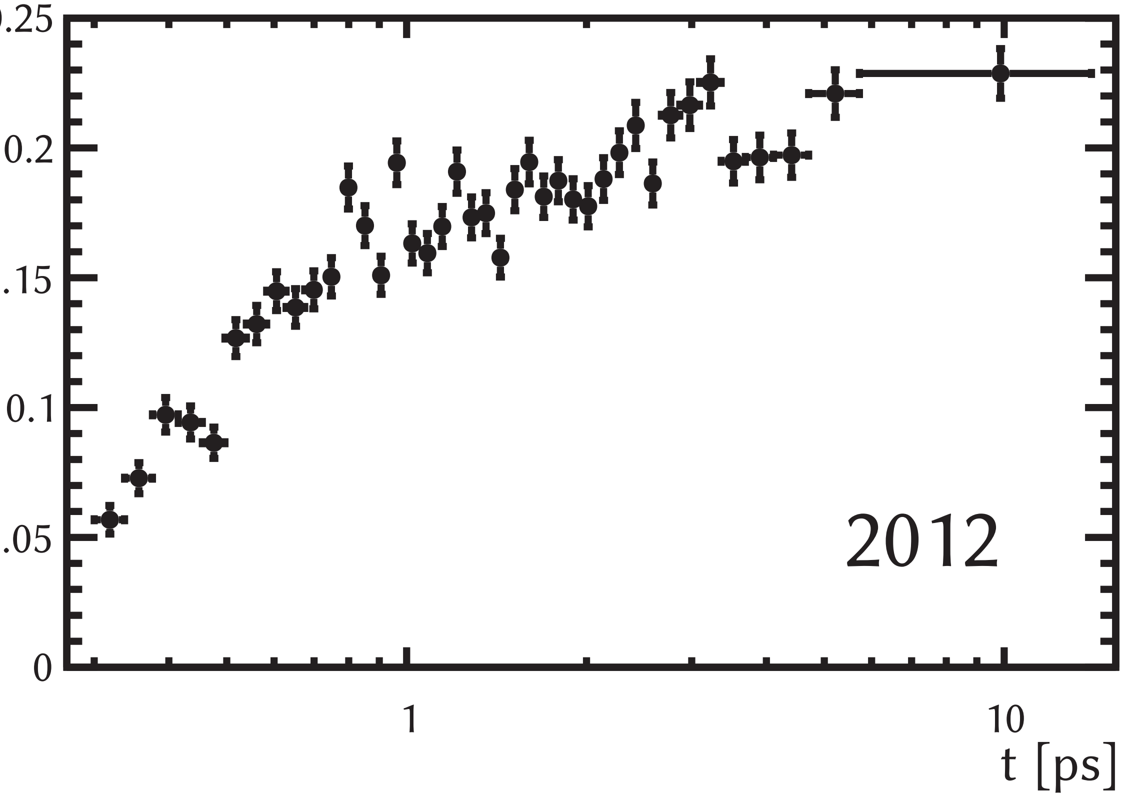

Hlt2DiMuonDisplaced

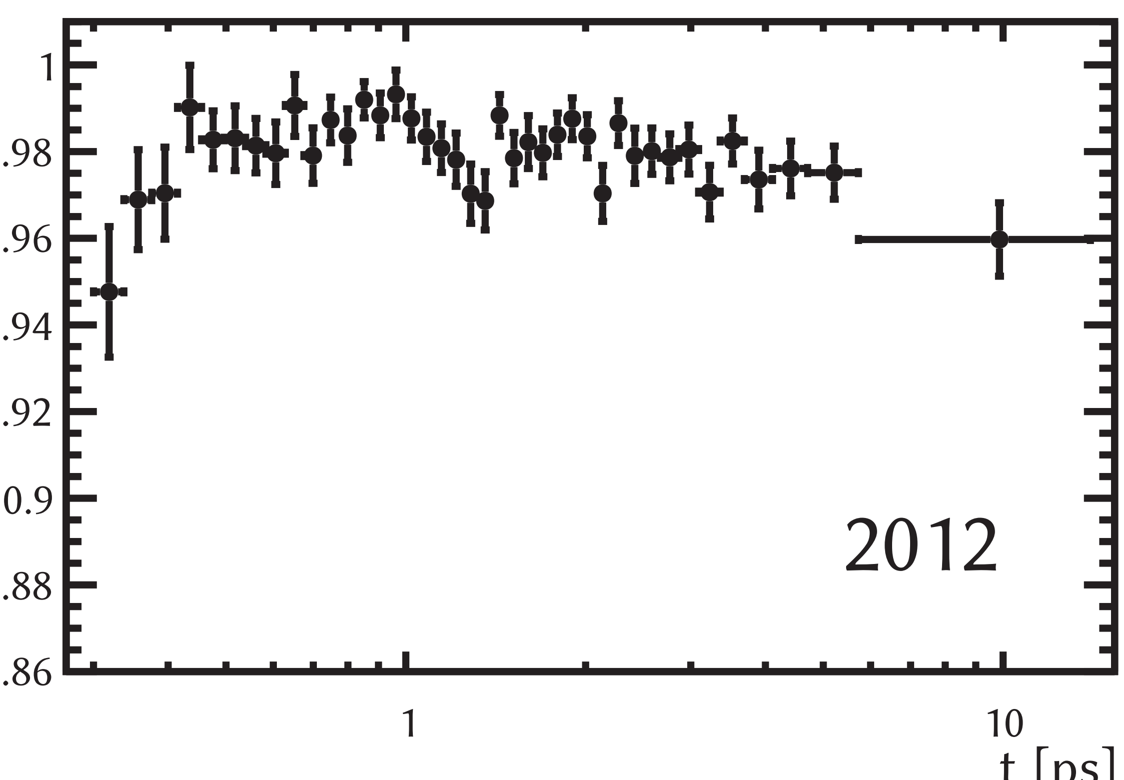

Efficiency of displaced J/ψ triggers

- Measured efficiency *relative* to prescaled 'lifetime unbiased' triggers

- Next one has to include this in the likelihood normalization.

$$\int dt\ \epsilon(t) \int dt' R(t-t')F(t')$$

$$= \sum_i \epsilon_i \int_{t_i}^{t_{i+1}} dt \int dt' R(t-t')F(t')$$

- If you don't like 'discretized' efficiency description, use splines instead (see eg. http://arxiv.org/abs/1407.0748 for relevant integrals)

HL-LHC

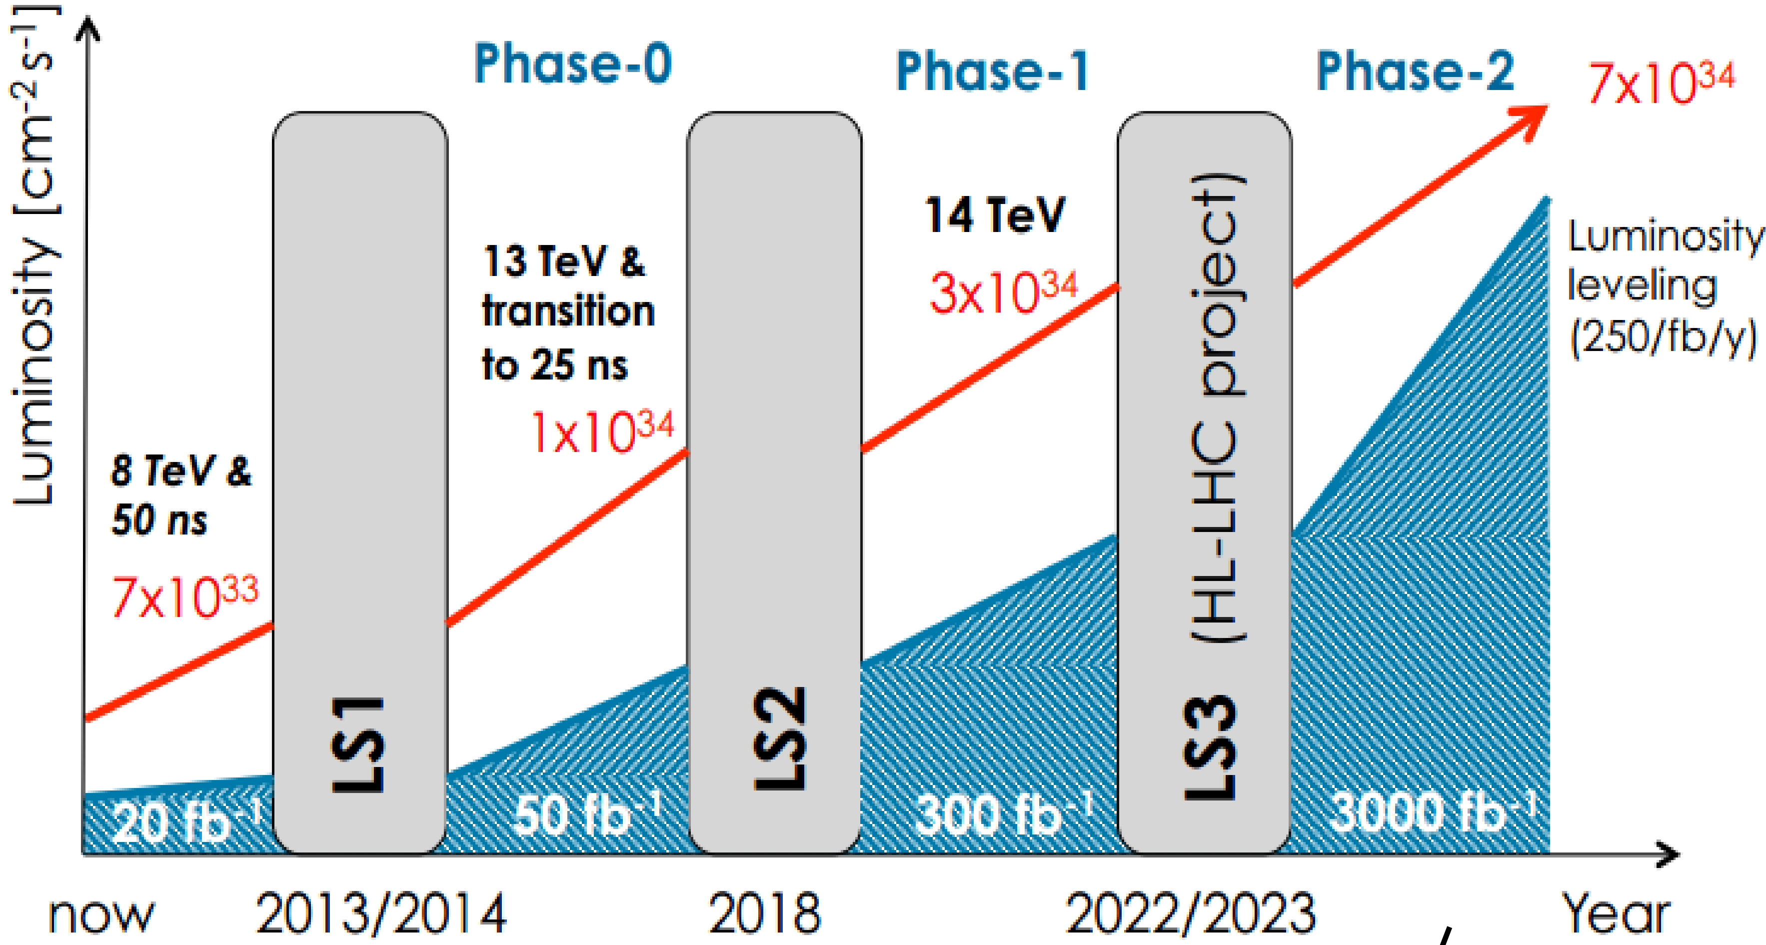

- lumi-leveling @ 5 10^{34} cm^{-2} s^{-1}
- 25 ns
- 14 TeV — minbias = 80 mb
- 140 PU *on average*, in region $\sigma z = 7.5cm/\sqrt{2}$ along beam
  - ~ 0.3 mm between PVs

Processs	σ	Rate (Hz)
Min bias	84 mb	3.4×10^9
ttbar	900 pb	45
$W\rightarrow l\nu$	20 nb	1.0×10^3
SM Higgs	30 pb	1.5

Vertex Consistency

- With O(140) interactions, ‘multi-object’ triggers will fire due to combinations from different PVs
- With tracking, can restrict to combinations from same interaction

CMS event with 40 PVs

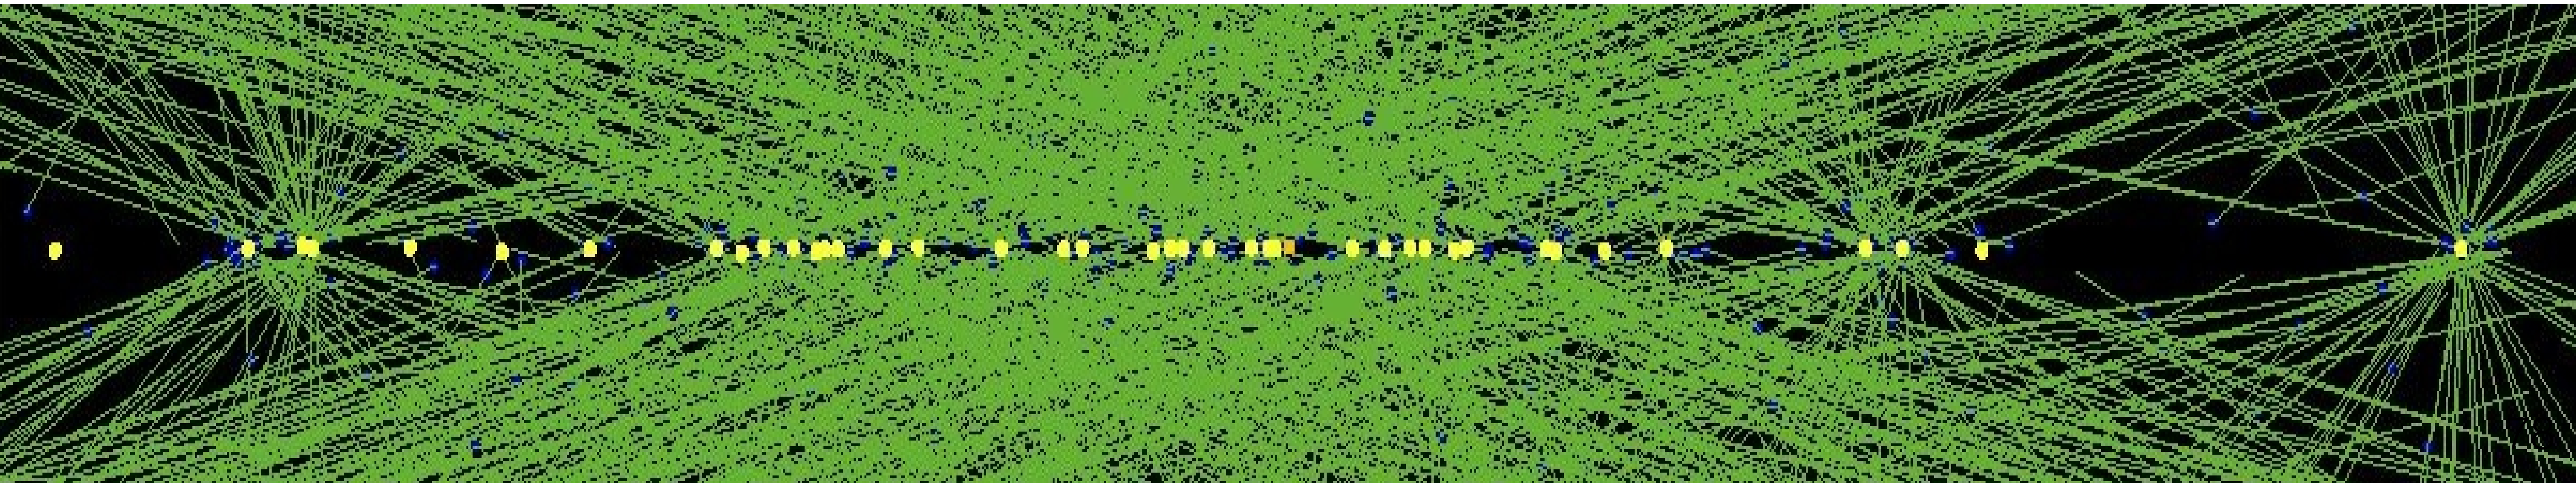

Use of tracking at L1…

- For example, *for CMS*, L1 muons momentum measurement can be significantly improved using the tracker

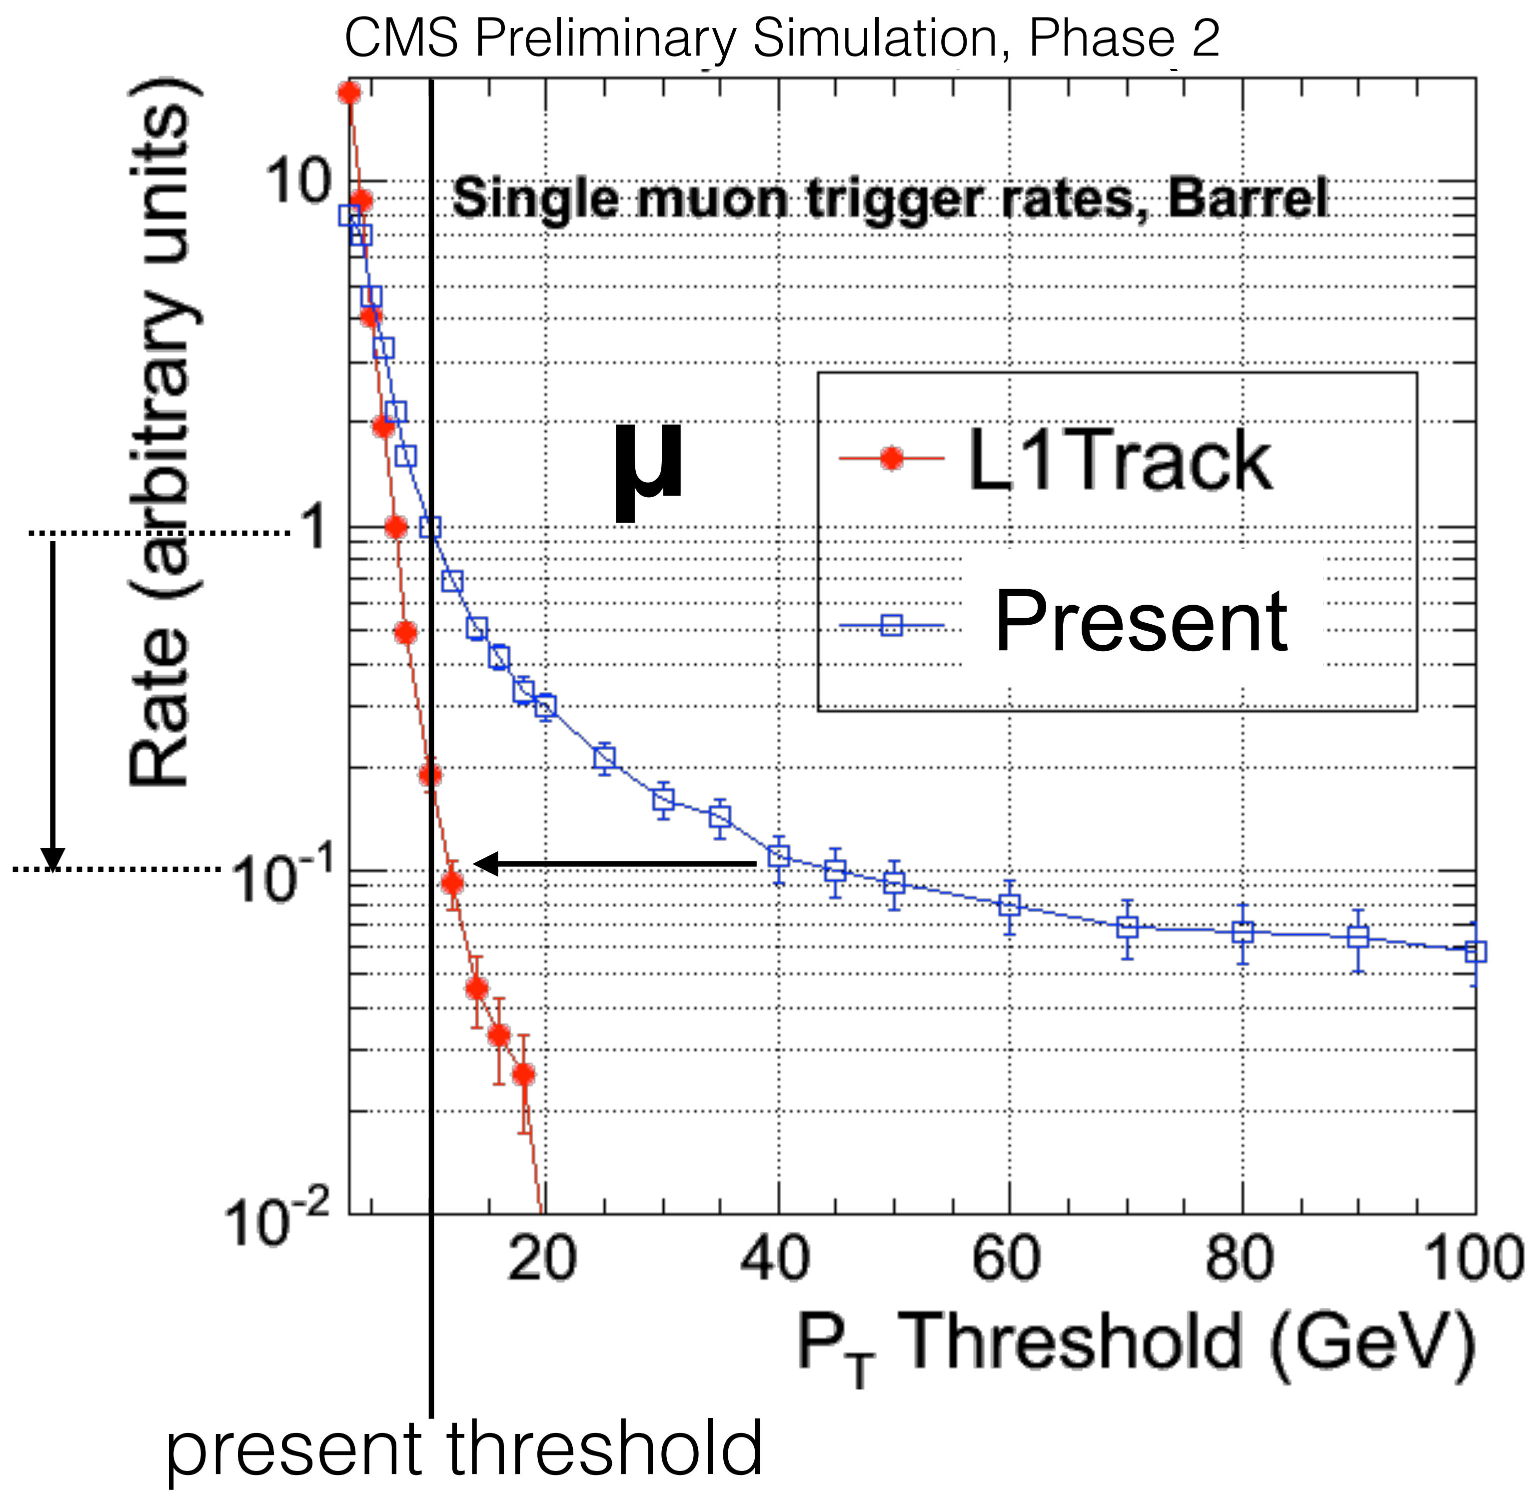

How to track at L1?

All Tracks

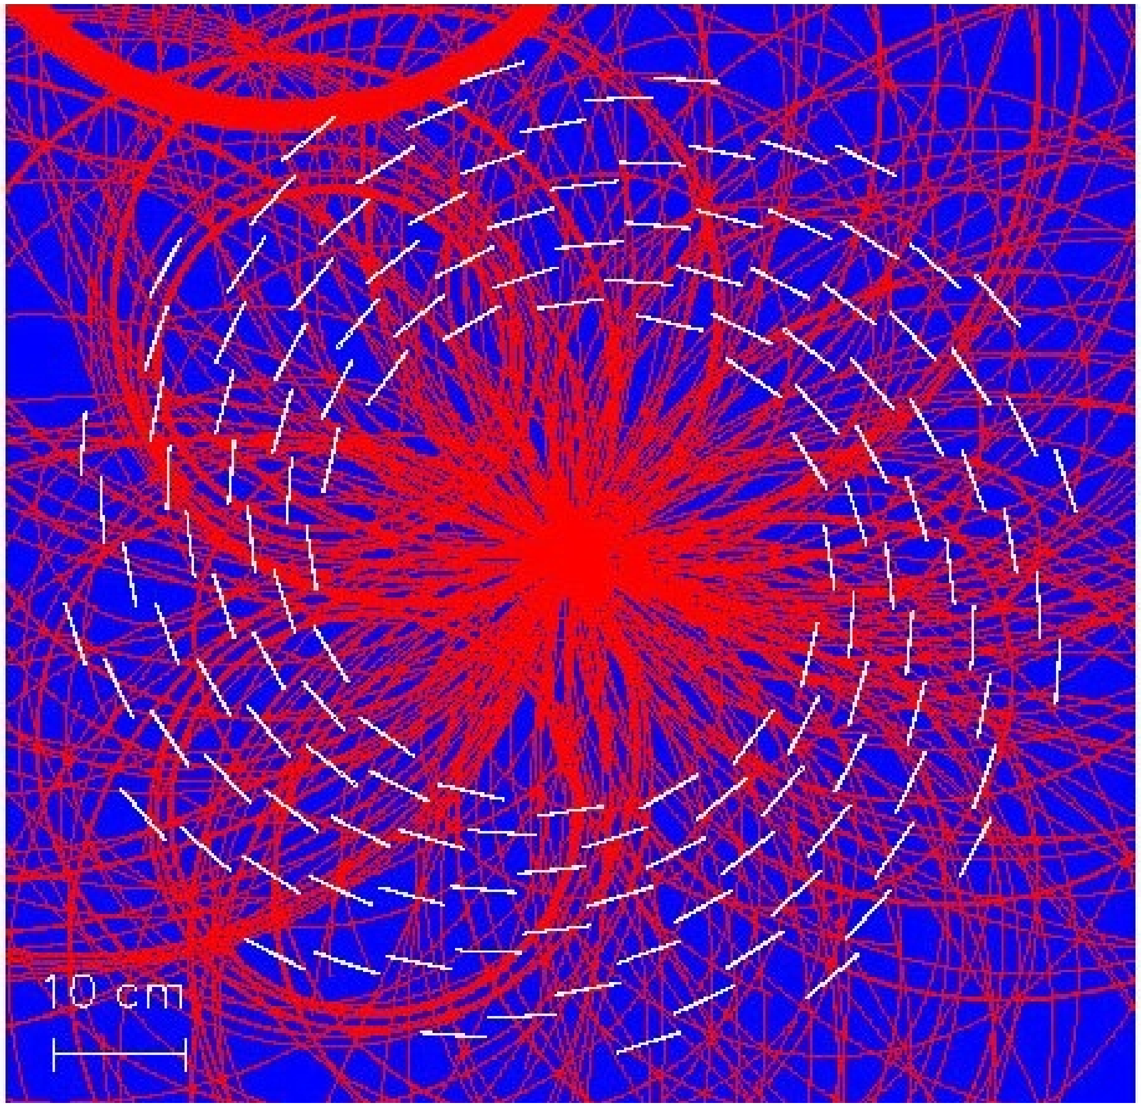

Tracks with P_t>2 GeV

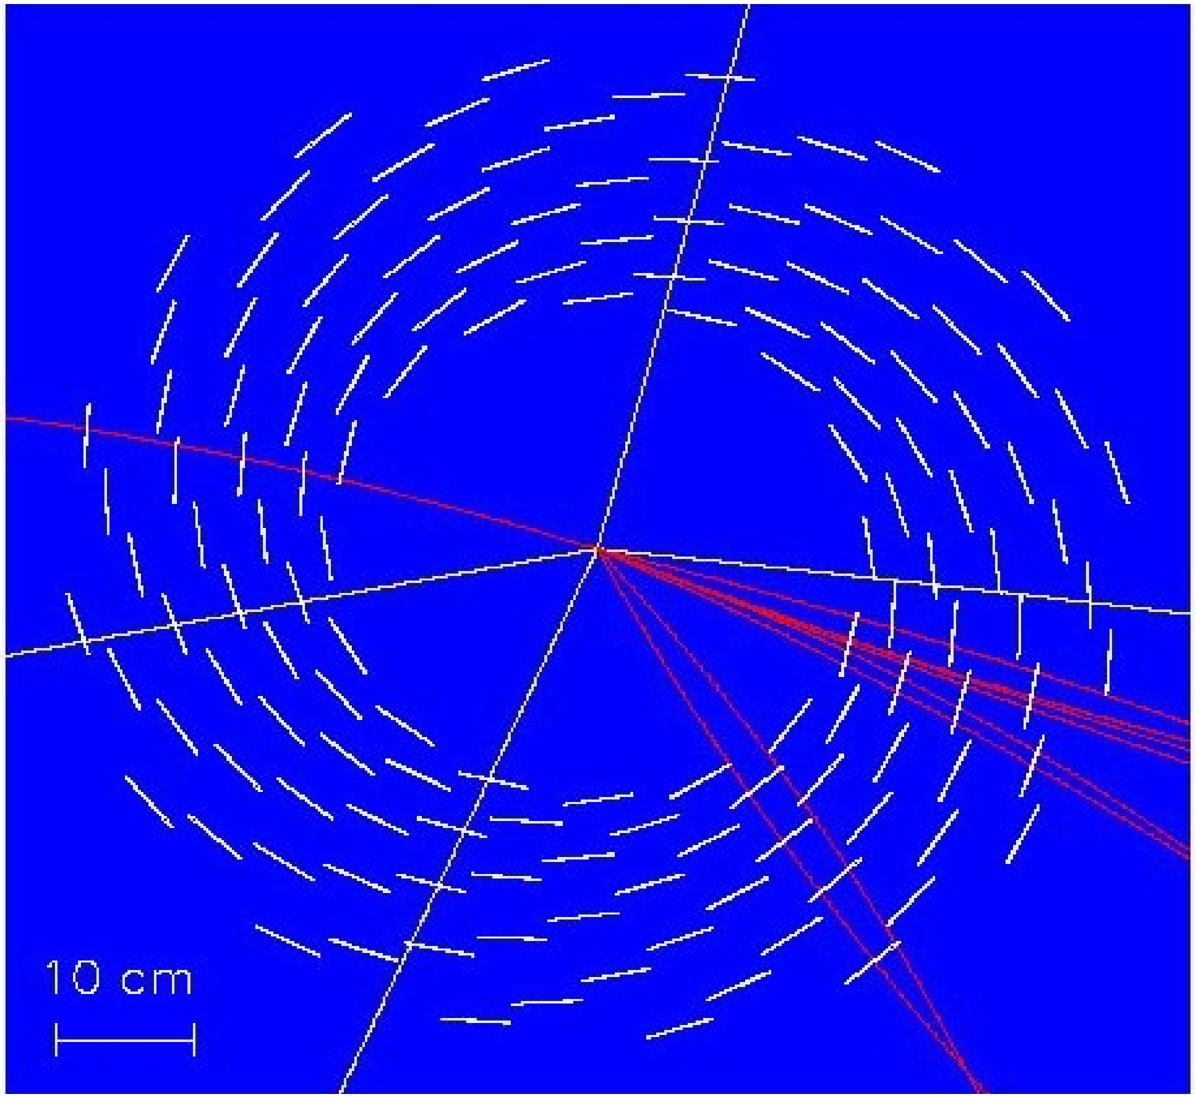

H -> 4mu, with 20 pile-up

How much data?

- Assuming 50 cm^2 modules, need 800 MHz of hits.
- @25 bits / hit —> 20 GB/s *average*.
- To deal with peak rate: O(50 GB/s)
- state-of-the-art *rad-hard*
  - GBLD10, 10 GB/s —> need O(5) links / module!!!
- How to reduce this?

R = 35 cm

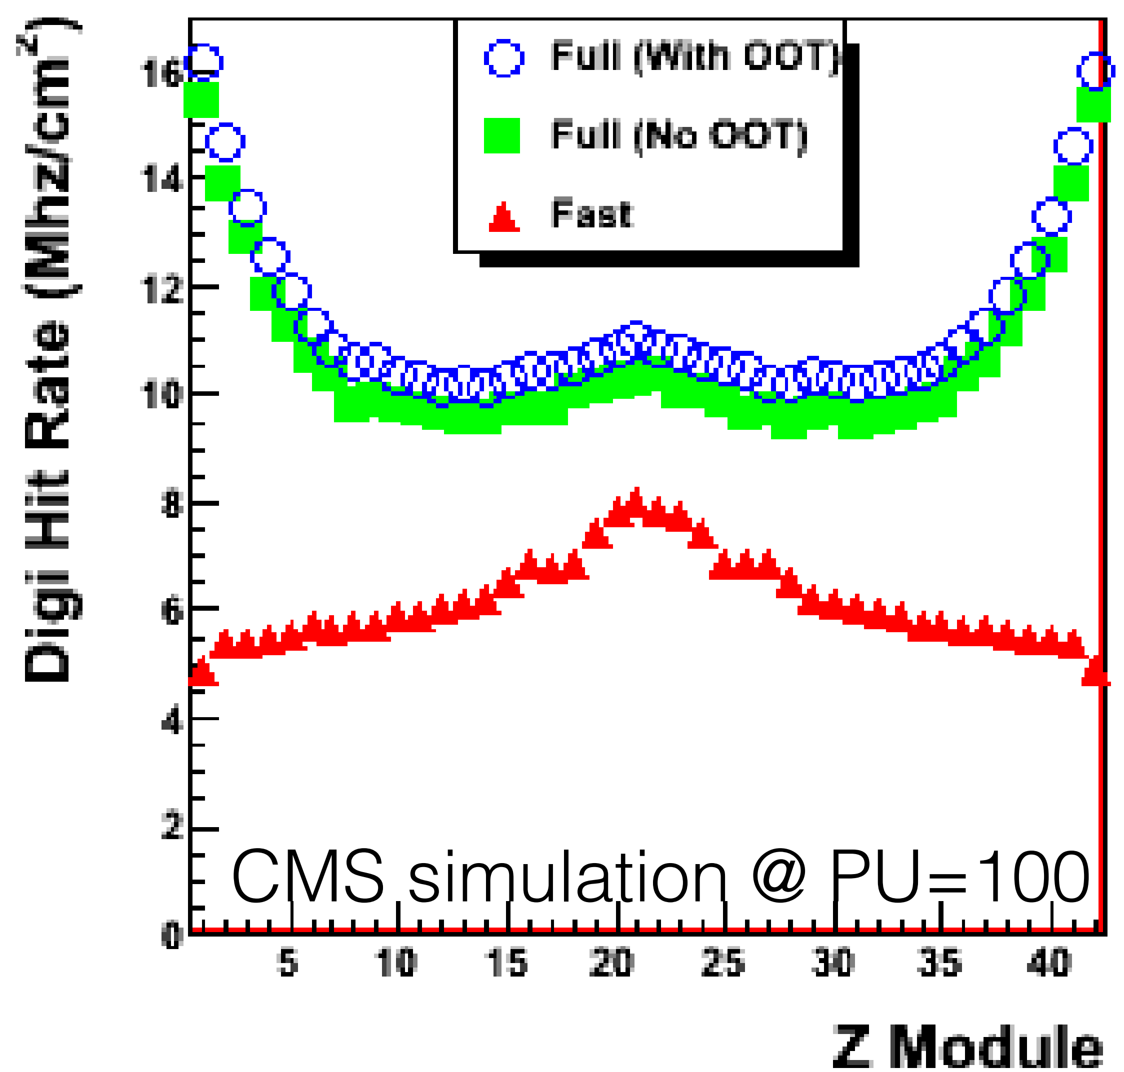

How much data?

- Assuming 50 cm^2 modules, need 800 MHz of hits.
- @25 bits / hit —> 20 GB/s *average*.
- To deal with peak rate: O(50 GB/s)
- state-of-the-art *rad-hard*
  - GBLD10, 10 GB/s —> need O(5) links / module!!!
- How to reduce this?

R = 35 cm

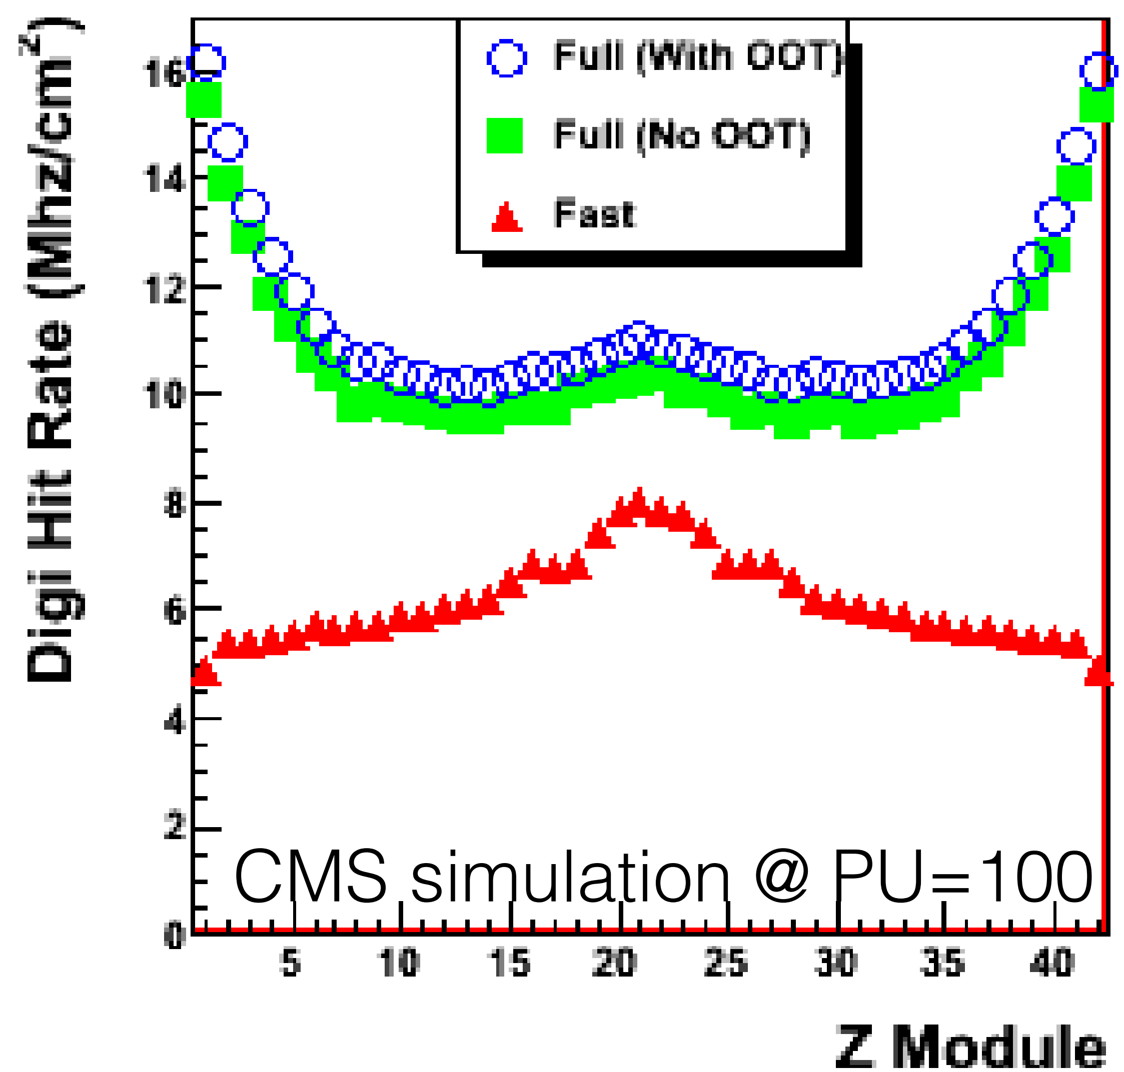

Atlas: use 'LVL0' to provide ROI using calorimeter, muon system & read out at reduced rate

CMS: perform *local* 'self-seeding' at full rate

Colorblind Monkeys…

- … are better at detecting *camouflaged* insects!

ANIMAL BEHAVIOUR, 2007, **73**, 205–214
doi:10.1016/j.anbehav.2006.07.003

ELSEVIER

Available online at www.sciencedirect.com
ScienceDirect

Effects of colour vision phenotype on insect capture by a free-ranging population of white-faced capuchins, *Cebus capucinus*

AMANDA D. MELIN*, LINDA M. FEDIGAN*, CHIHIRO HIRAMATSU†,
COURTNEY L. SENDALL* & SHOJI KAWAMURA†
*Department of Anthropology, University of Calgary
†Department of Integrated Biosciences, Graduate School of Frontier Sciences, University of Tokyo

*(Received 19 April 2006; initial acceptance 30 May 2006;
final acceptance 10 July 2006; published online 29 November 2006; MS. number: A10426)*

Unlike most eutherian mammals, which have dichromatic (two-colour) vision, most platyrrhine primate species have polymorphic colour vision. This unique characteristic is enabled via multiple alleles for a mid- to long-wavelength-sensitive (M/LWS), single-locus opsin gene on the X chromosome. In combination with the autosomal opsin common to most vertebrates, this arrangement provides heterozygous females with trichromatic (three-colour) vision, whereas homozygous females and males are dichromats. Trichromatic vision enables visual differentiation among longer-wavelength colours, such as red, orange, yellow and green. Currently, many researchers attribute the evolution and maintenance of polymorphic colour vision to trichromat (= heterozygote) advantage. However, dichromacy may be more suited for achromatic tasks, such as penetrating colour camouflage, especially under low-light conditions. We evaluated whether dichromatic capuchin monkeys (*Cebus capucinus*) were more efficient than trichromatic monkeys at capturing camouflaged and noncamouflaged insects. Through faecal DNA analysis, we determined the genotypes of the M/LWS opsins for 34 capuchins in two groups inhabiting Santa Rosa National Park, Costa Rica. Dichromatic monkeys were more efficient at detecting camouflaged, surface-dwelling insects, especially under conditions of low ambient light. However, unexpectedly, trichromats were more efficient in extracting embedded, noncamouflaged insects from substrates. To our knowledge, this is the first study to document a foraging advantage to dichromatic monkeys in the wild. Our findings show that there is a lack of heterozygote advantage in foraging for surface-dwelling insects and therefore indicate that this mechanism may not be the sole driving force maintaining polymorphic colour vision in this population.

CMS Stacked Tracker

Use correlation between P_T and direction @ some radius:

1. cluster width depends on angle

2. with two layers close in radius, P_T cut implies geometric correlation

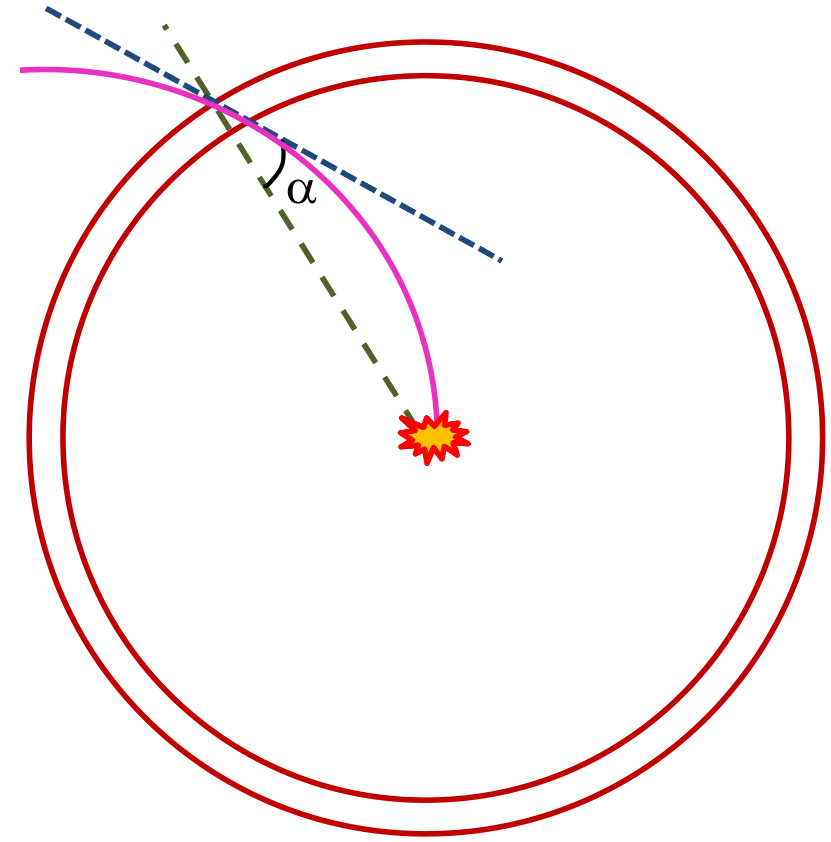

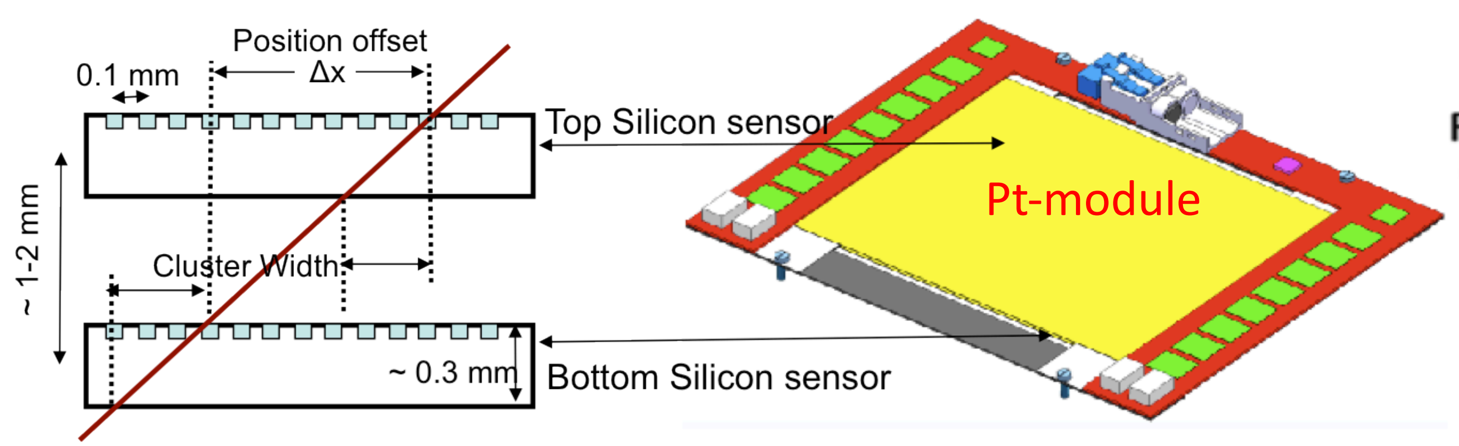

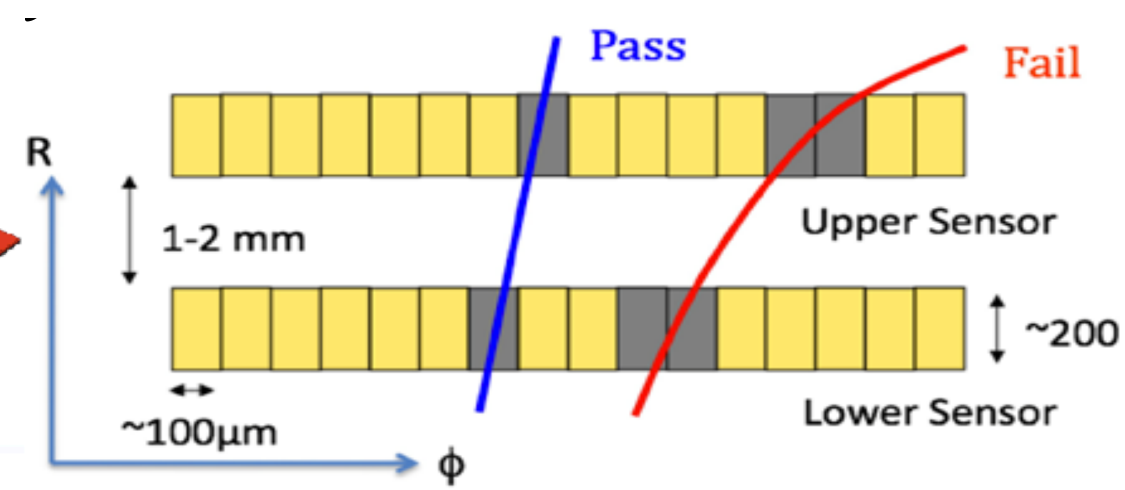

Finding Valid Patterns

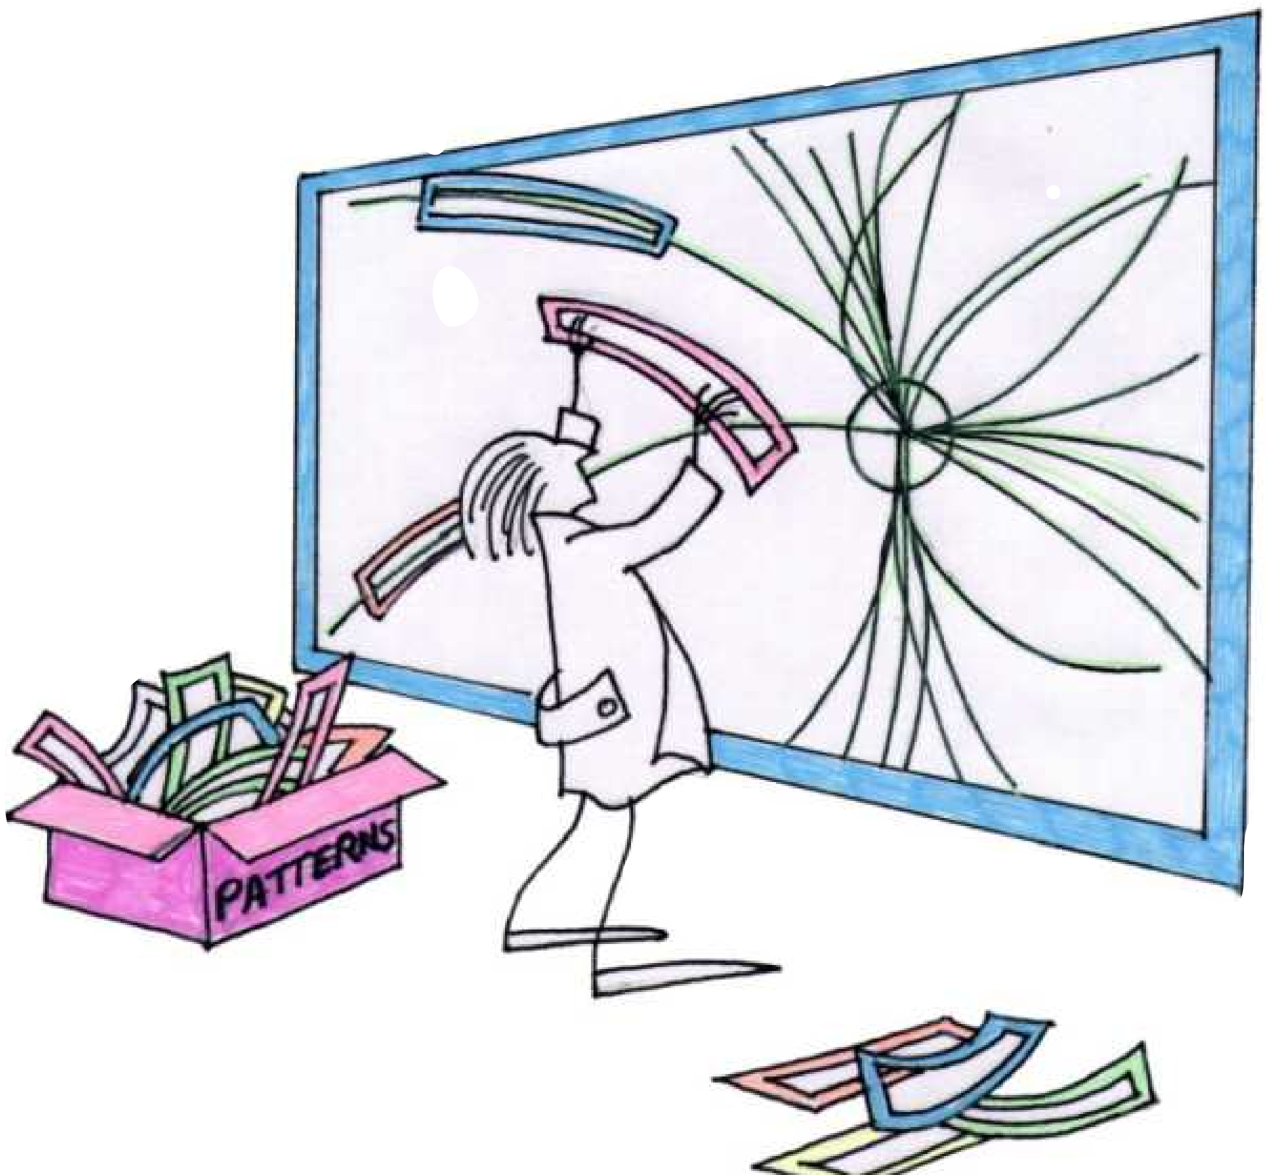

1. Based on P_T requirements, precompute 'all' allowed hit combinations & store in look-up table

Finding Valid Patterns

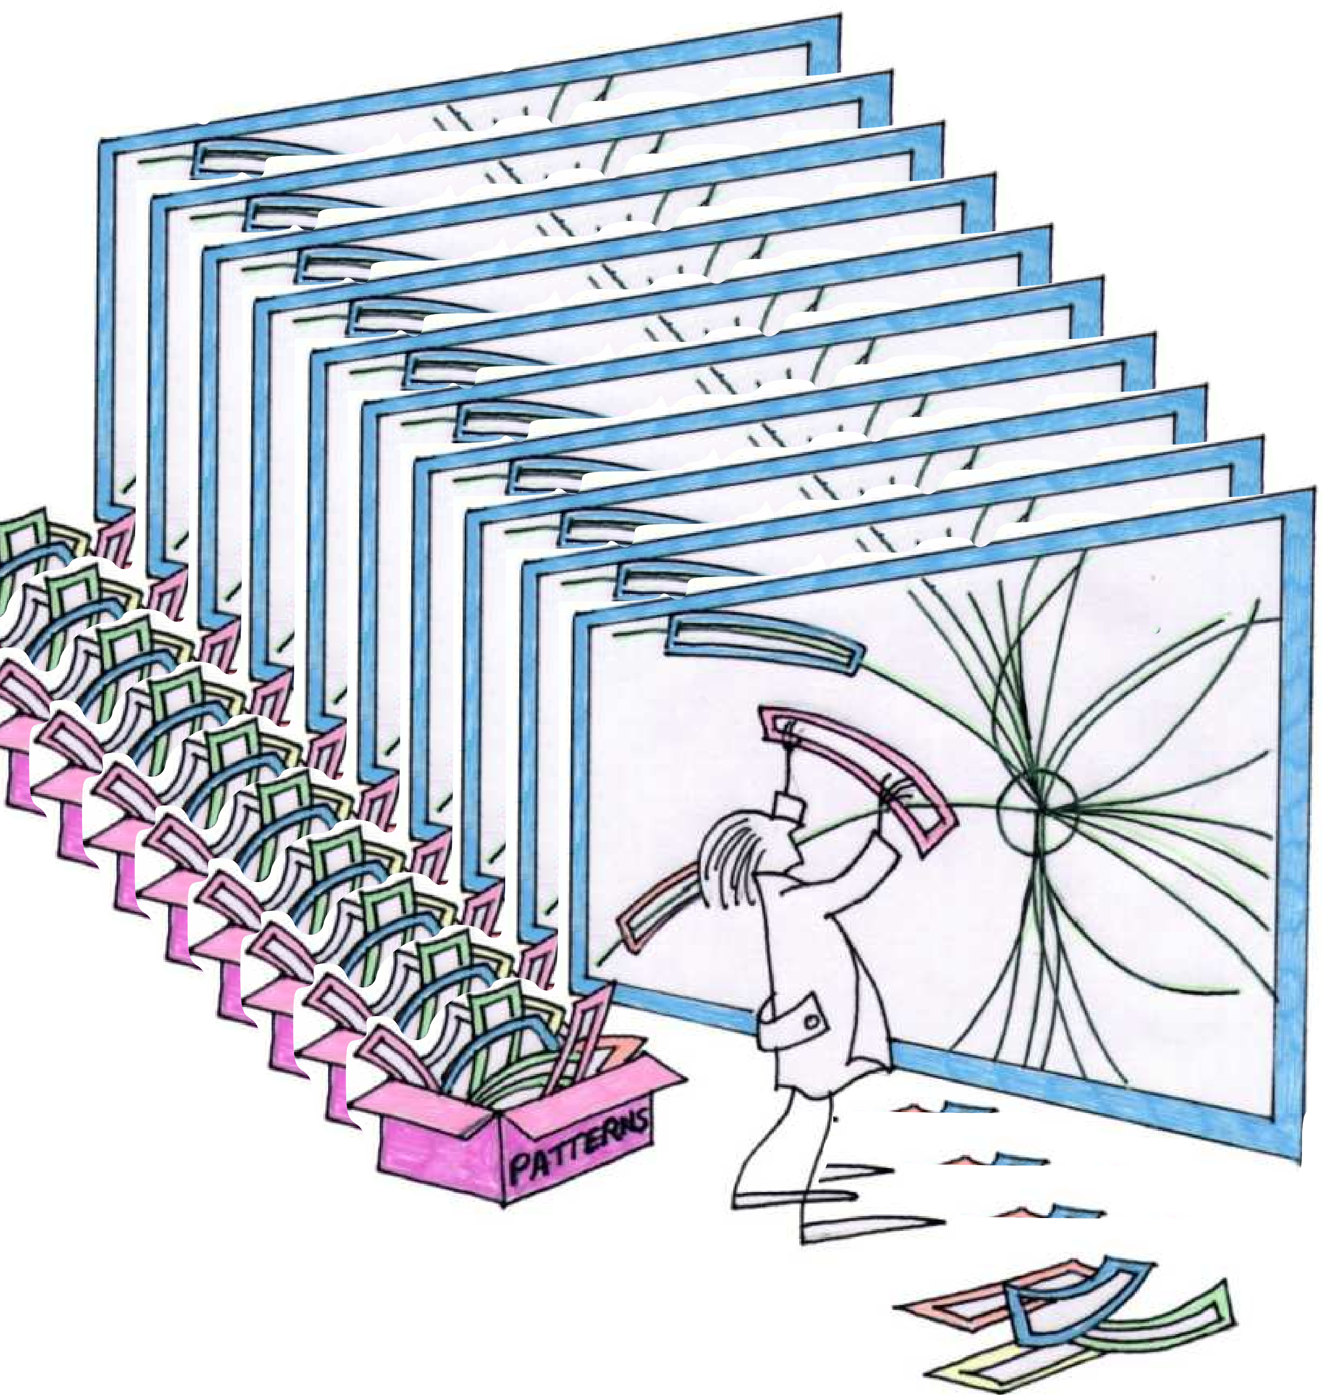

1. Based on P_T requirements, precompute 'all' allowed hit combinations & store in look-up table

Finding Valid Patterns

1. Based on P_T requirements, precompute ‘all’ allowed hit combinations & store in look-up table
2. Use FPGA or eg. so-called ‘Associative Memory’ to match observed pattern with table (“reverse lookup”)

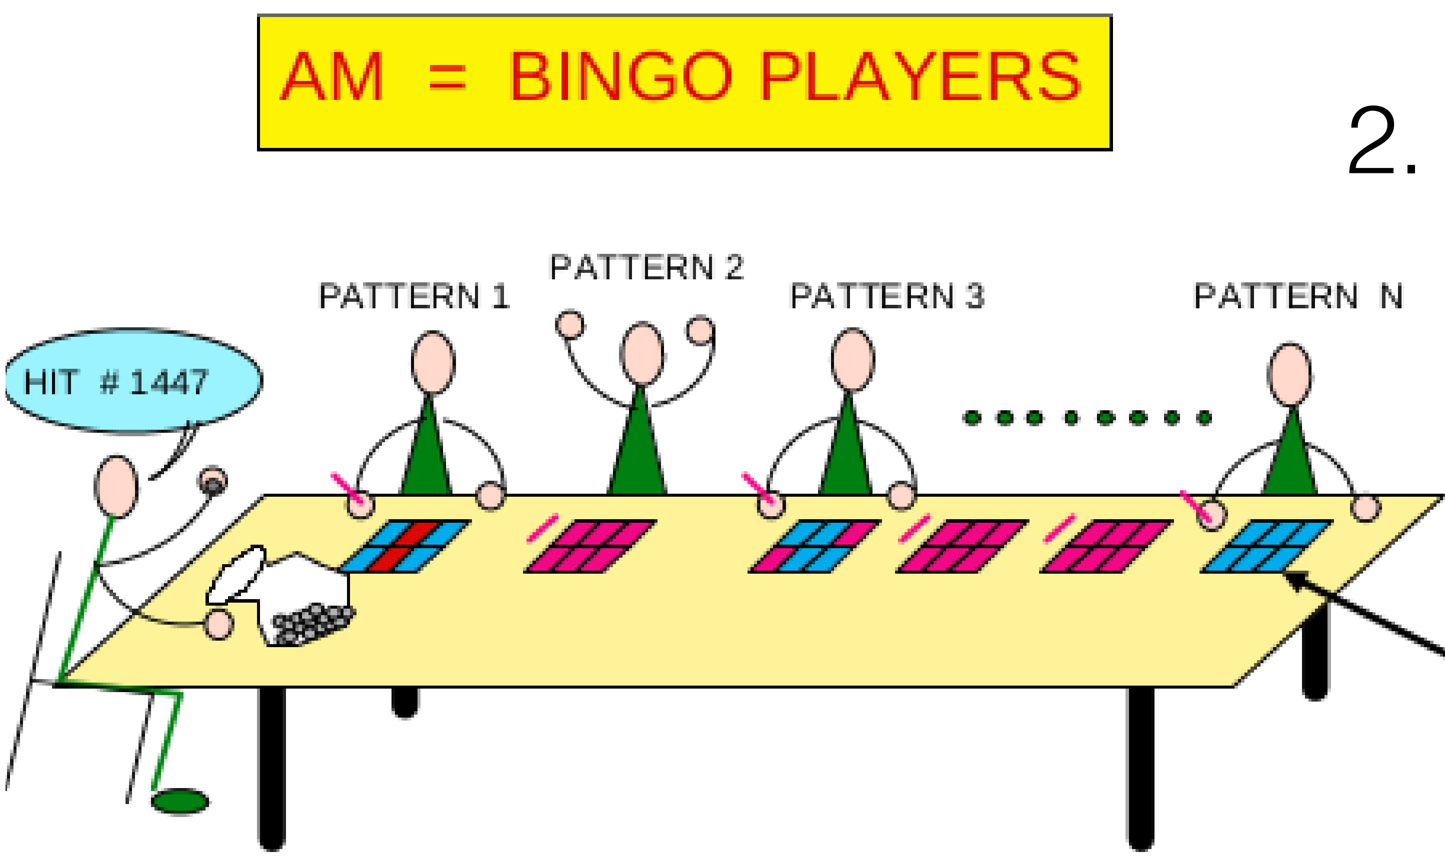

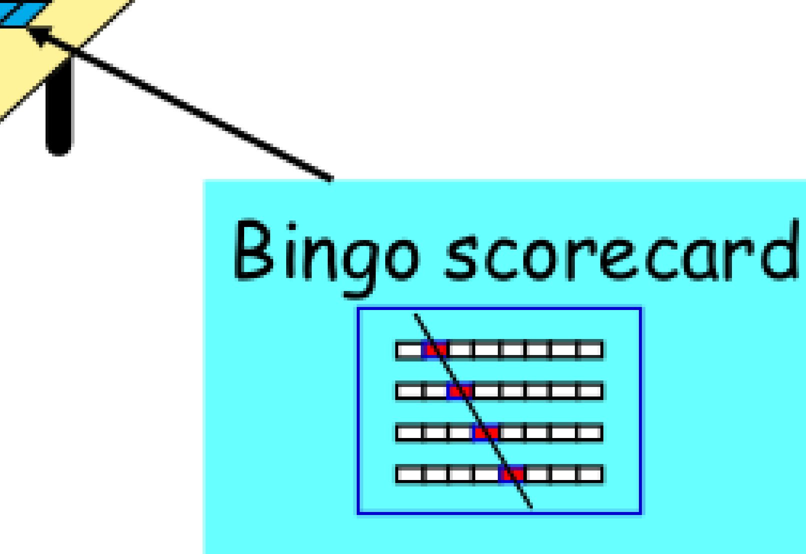

Used in eg. Atlas FTK

M. Dell’Orso, L. Ristori, NIM A 278, 436 (1989)

Atlas FTK

2015: partial installation
2016: initial HLT usage
2017: complete coverage, use for physics

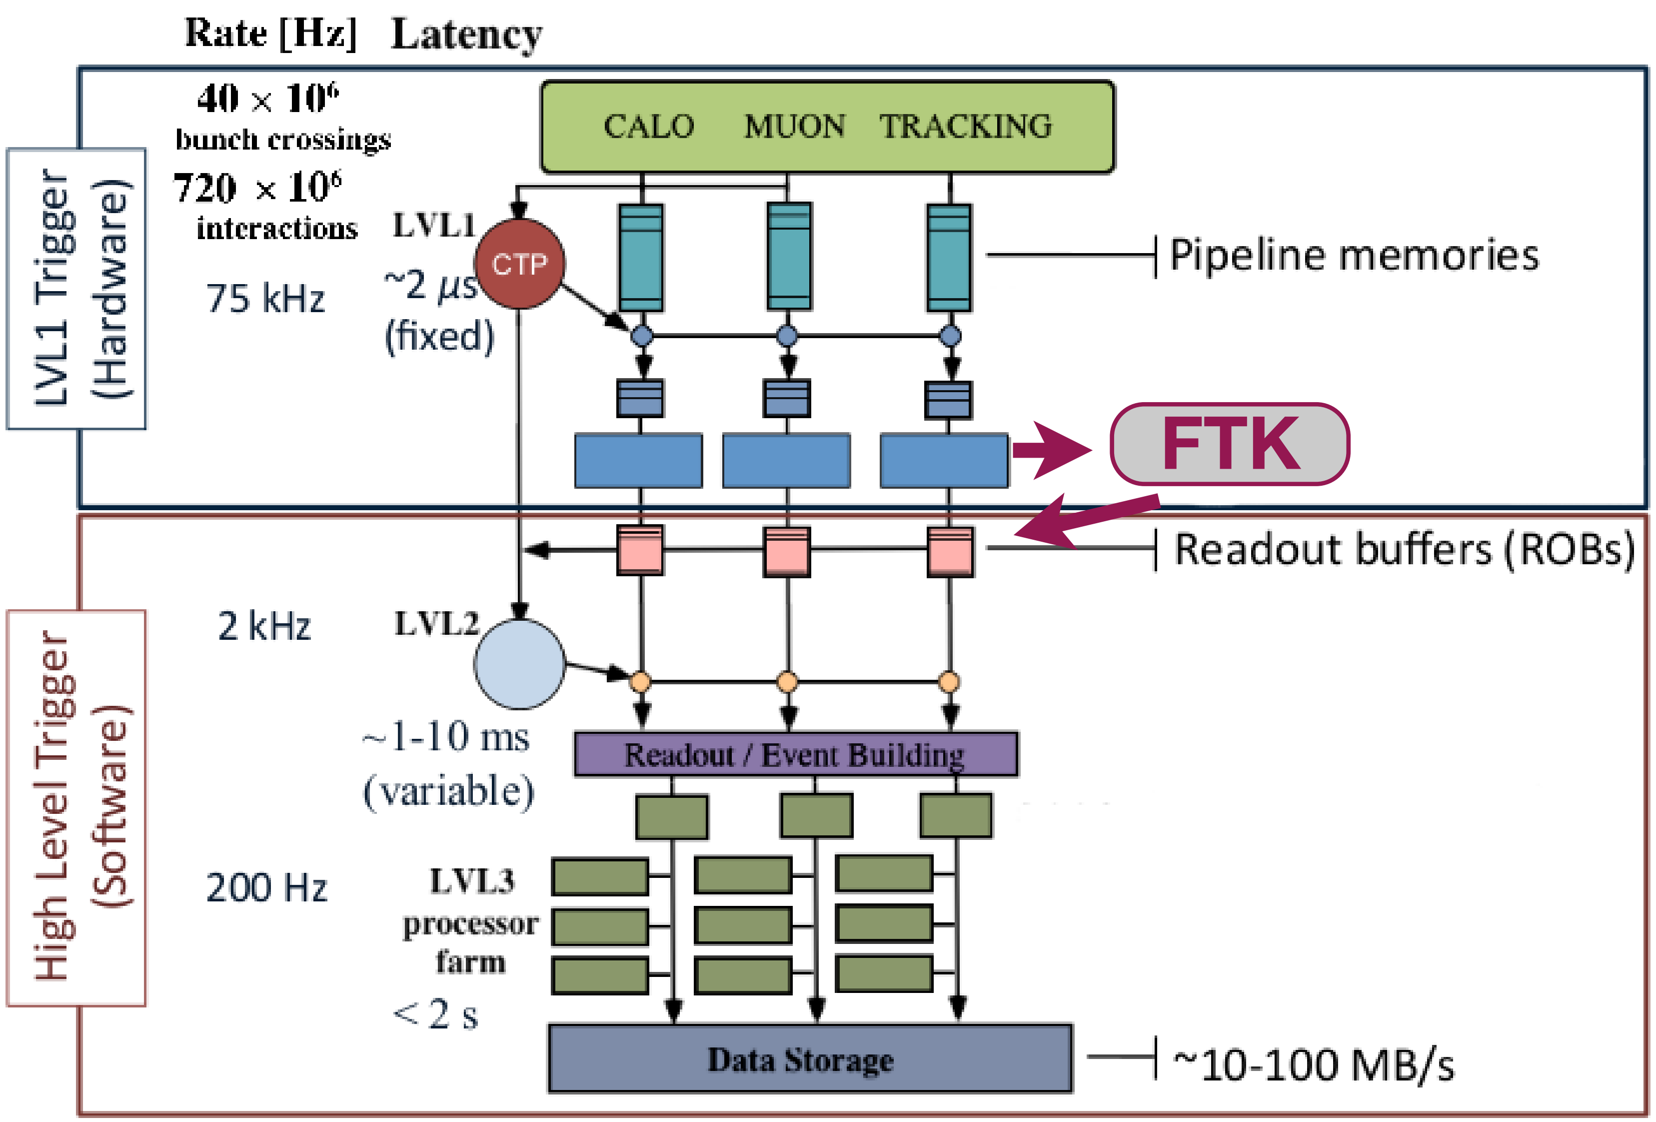

Processes L1 accepted events (100kHz)

Output tracks with P_T>1 GeV/c to HLT.

FTK latency ~100μs

Atlas DAQ after LS3

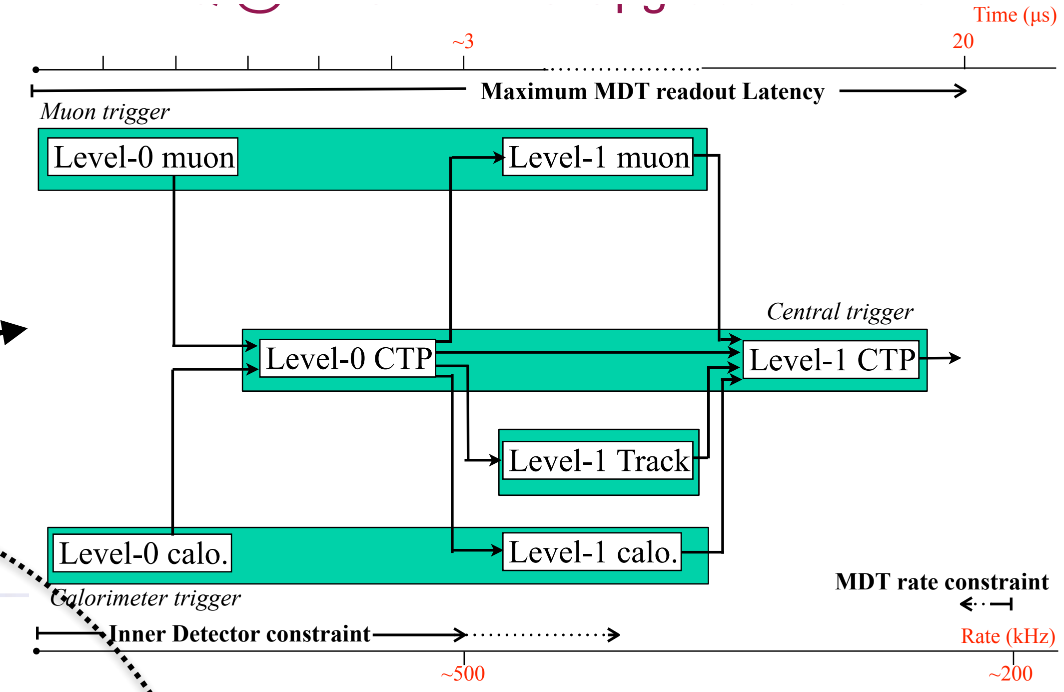

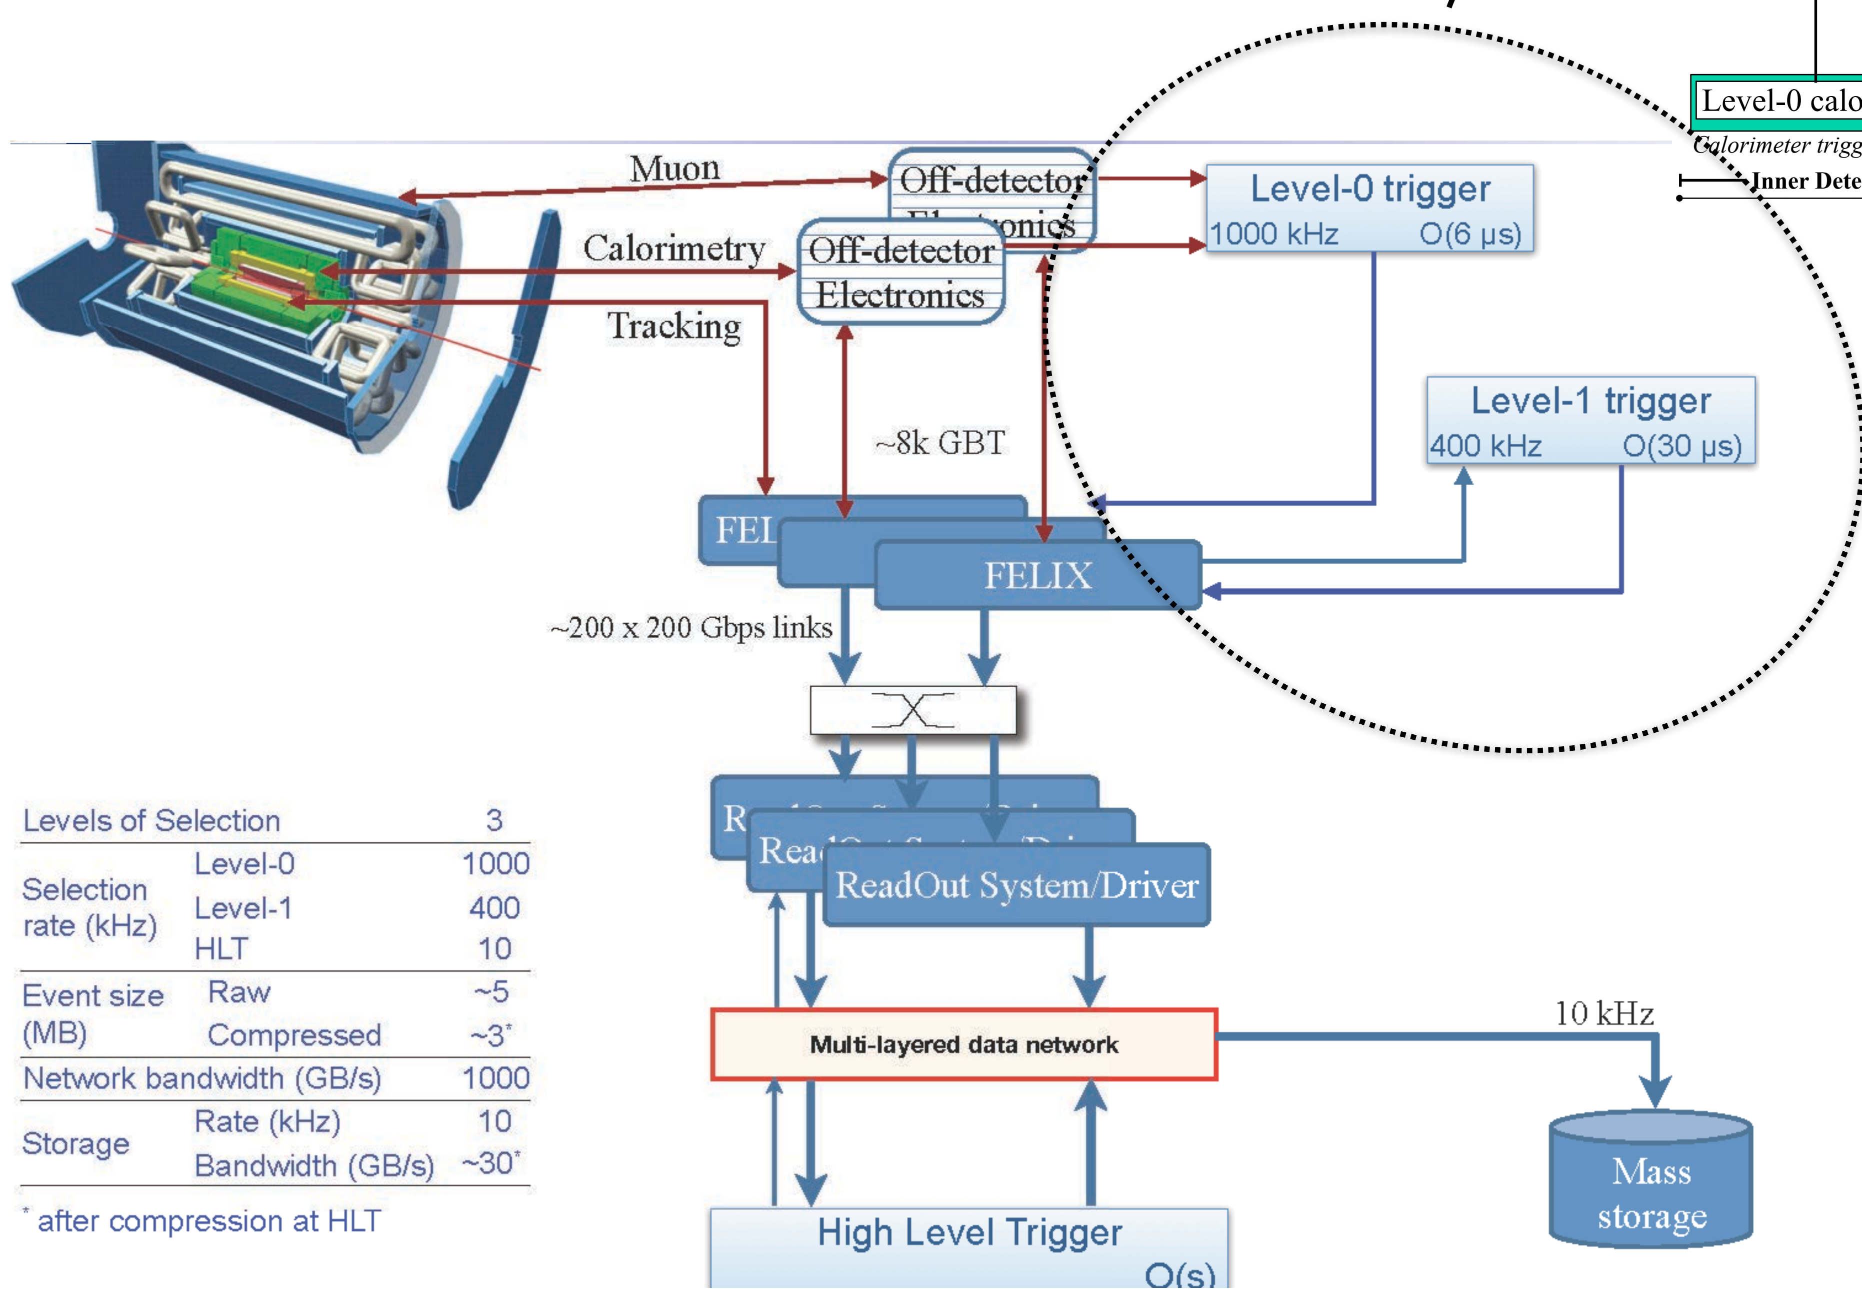

Levels of Selection		3
Selection rate (kHz)	Level-0	1000
	Level-1	400
	HLT	10
Event size (MB)	Raw	~5
	Compressed	~3*
Network bandwidth (GB/s)		1000
Storage	Rate (kHz)	10
	Bandwidth (GB/s)	~30*

* after compression at HLT

CMS DAQ after LS3

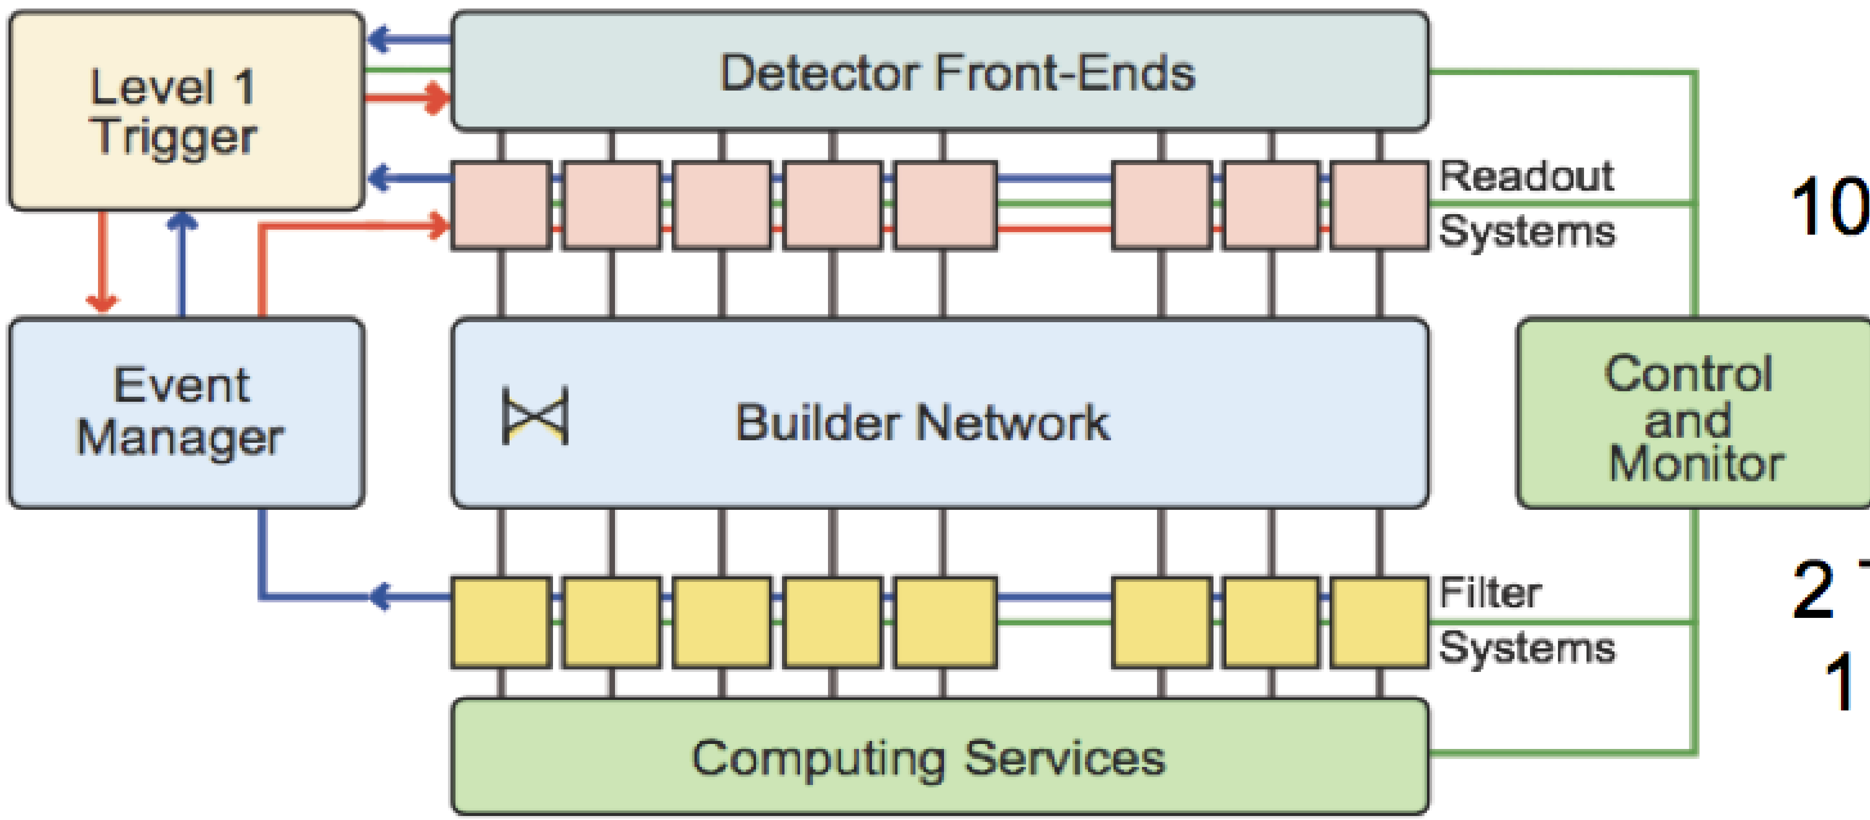

40 MHz
100 kHz → 750 kHz

2 Tbps → ~50 Tbps
1 kHz → ~7.5 kHz

- Track trigger combined with calorimeter & muon triggers regionally, then transmitted to global L1
- L1 latency: 12.5 us

LHC Trigger/DAQ upgrades

	ALICE	LHCb	CMS	ATLAS
Hardware trigger	No	No	Yes	Yes
Software trigger input rate	50 kHz Pb-Pb 200 kHz p-Pb	30 MHz	500/750 kHz for PU 140/200	0.4 MHz
Baseline processing architecture	CPU/GPU/FPGA/ Cloud&Grid	CPU farm (+coprocessors)	CPU farm (+coprocessors)	CPU farm (+coprocessors)
Software trigger output rate	50 kHz Pb-Pb 200 kHz p-Pb	20-100 kHz	5-7.5 kHz	5-10 kHz

HLT: can it scale?

- CMS HLT will have to process x10, with same rejection ÷100 as today
- Scaling of CPU time seems ~ linear with pile-up (non-trivial!)

✓ Need O(10 x 5) todays capacity
— 50 = 30%/year x 10 years

- And will have access to tracks from L1 track trigger

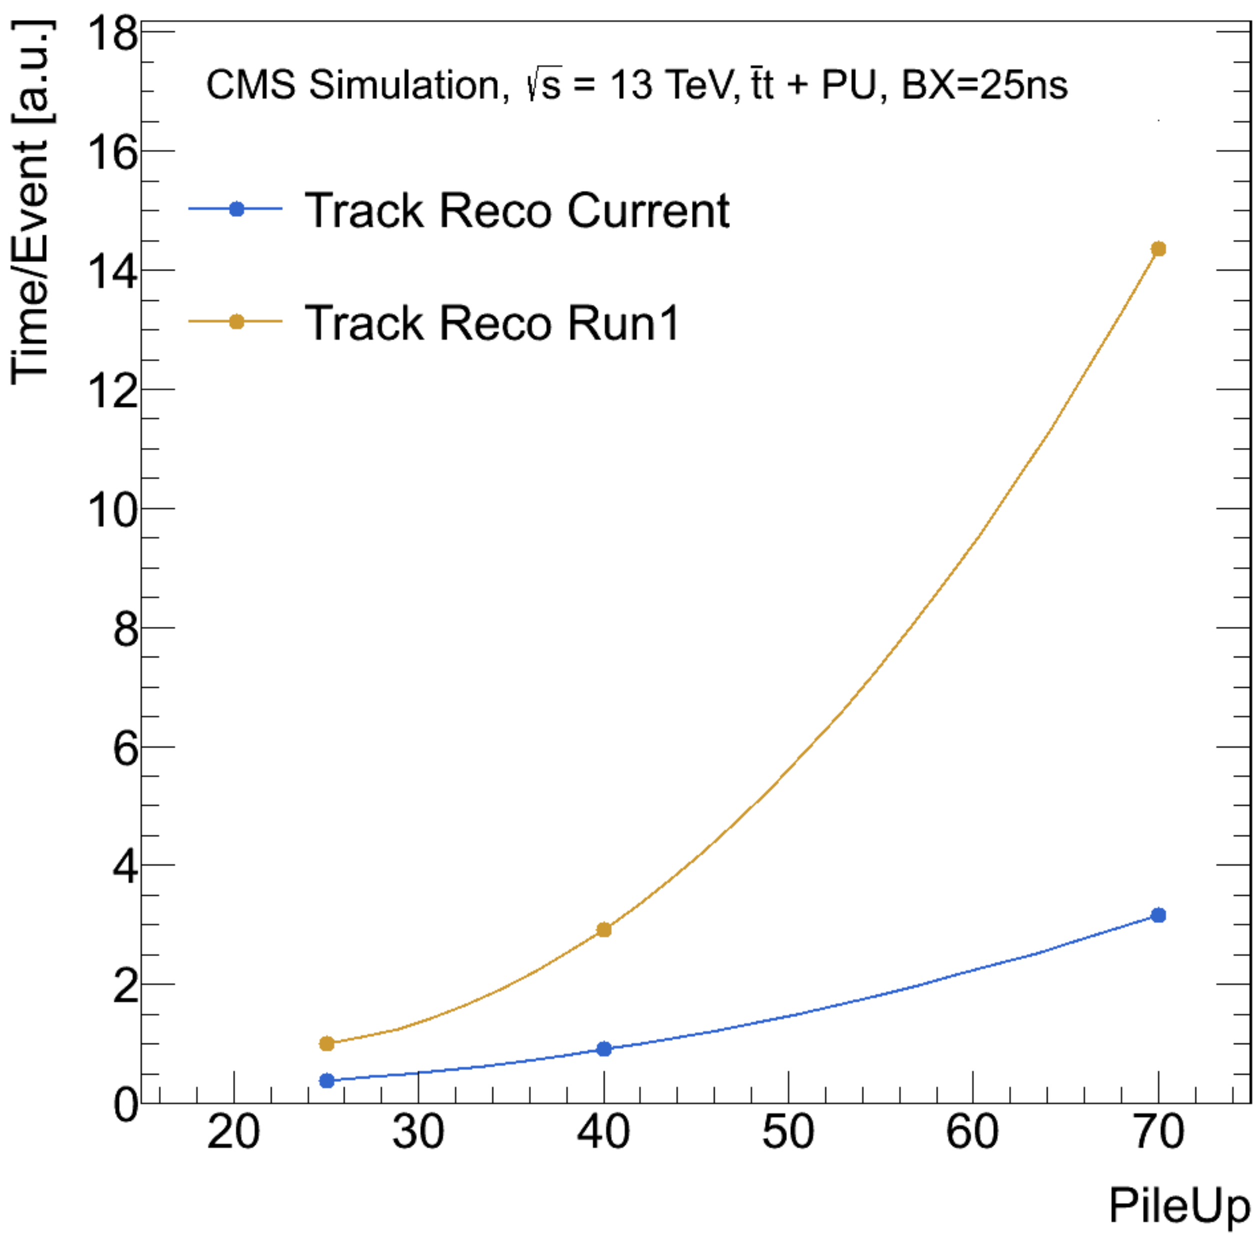

HLT: will it scale?

- Already in 2005 (*) it was realized that “things are about to change”
- First 2 GHz Intel CPU: August 2001
- Days of ever faster “sequential processing” are long gone…

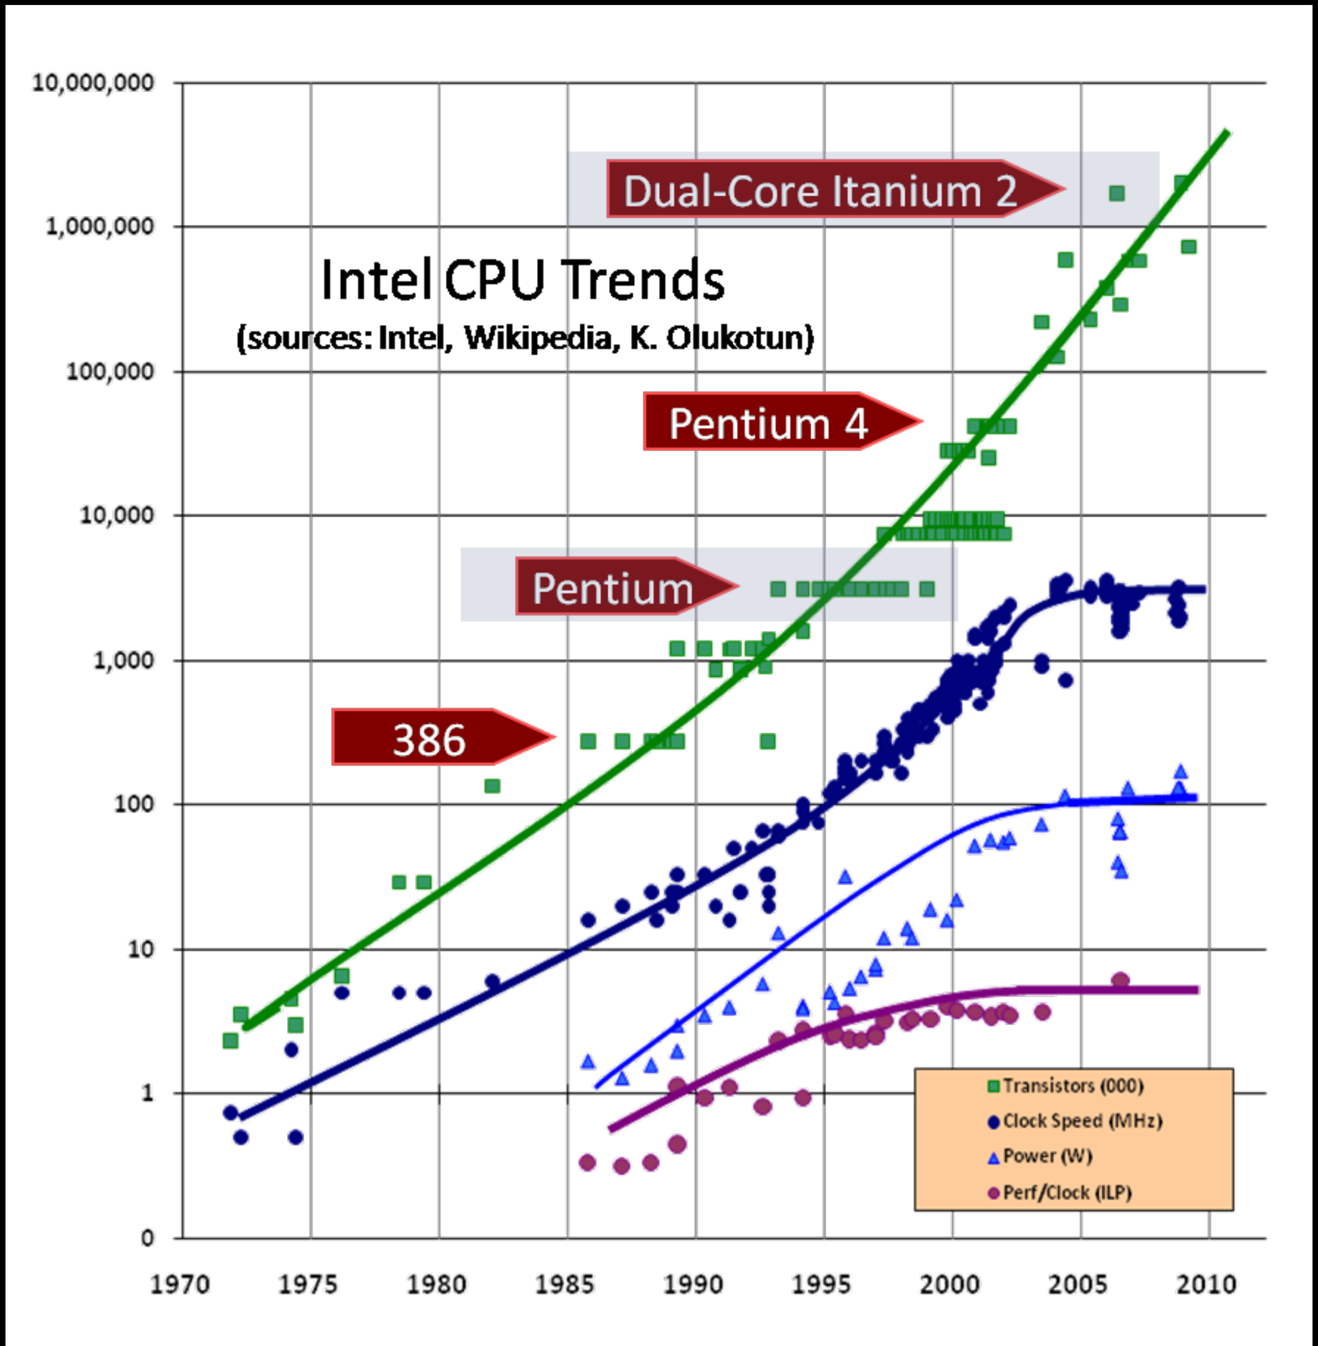

(*) Herb Sutter, Dr. Dobb's Journal, 30(3), March 2005.

HARDWARE LANDSCAPE

"Yesterday"
3.1 million transistors

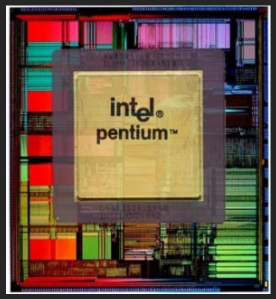

Today
1.4 billion transistors

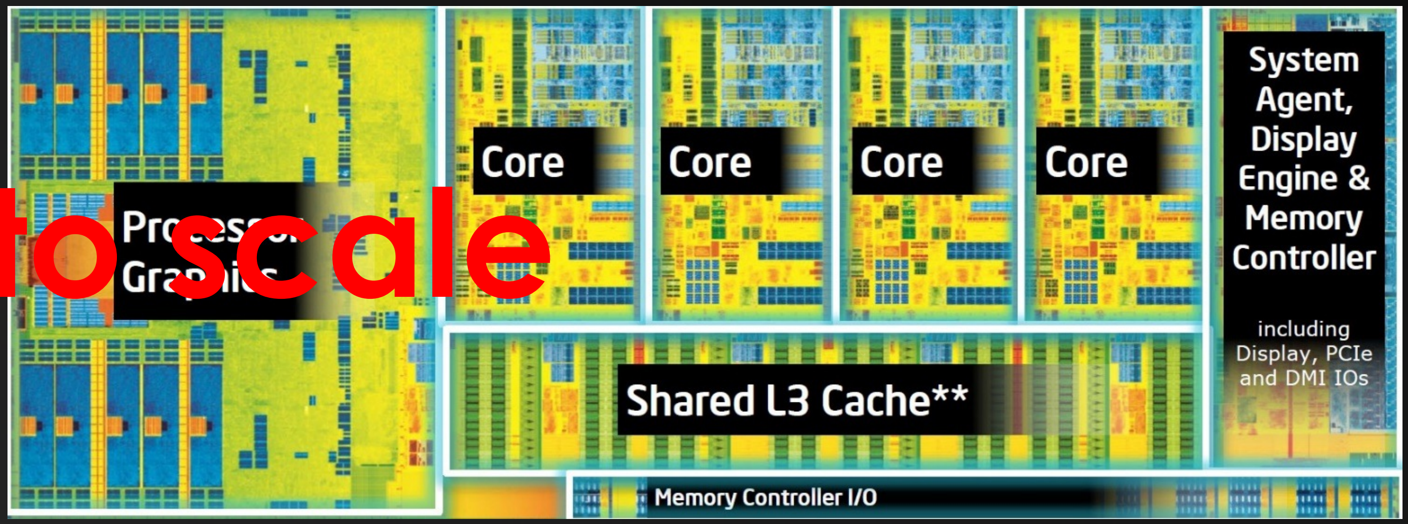

Not to scale

Taken verbatim from
"Compiler Confidential", Eric Brummer @ Going Native 2013

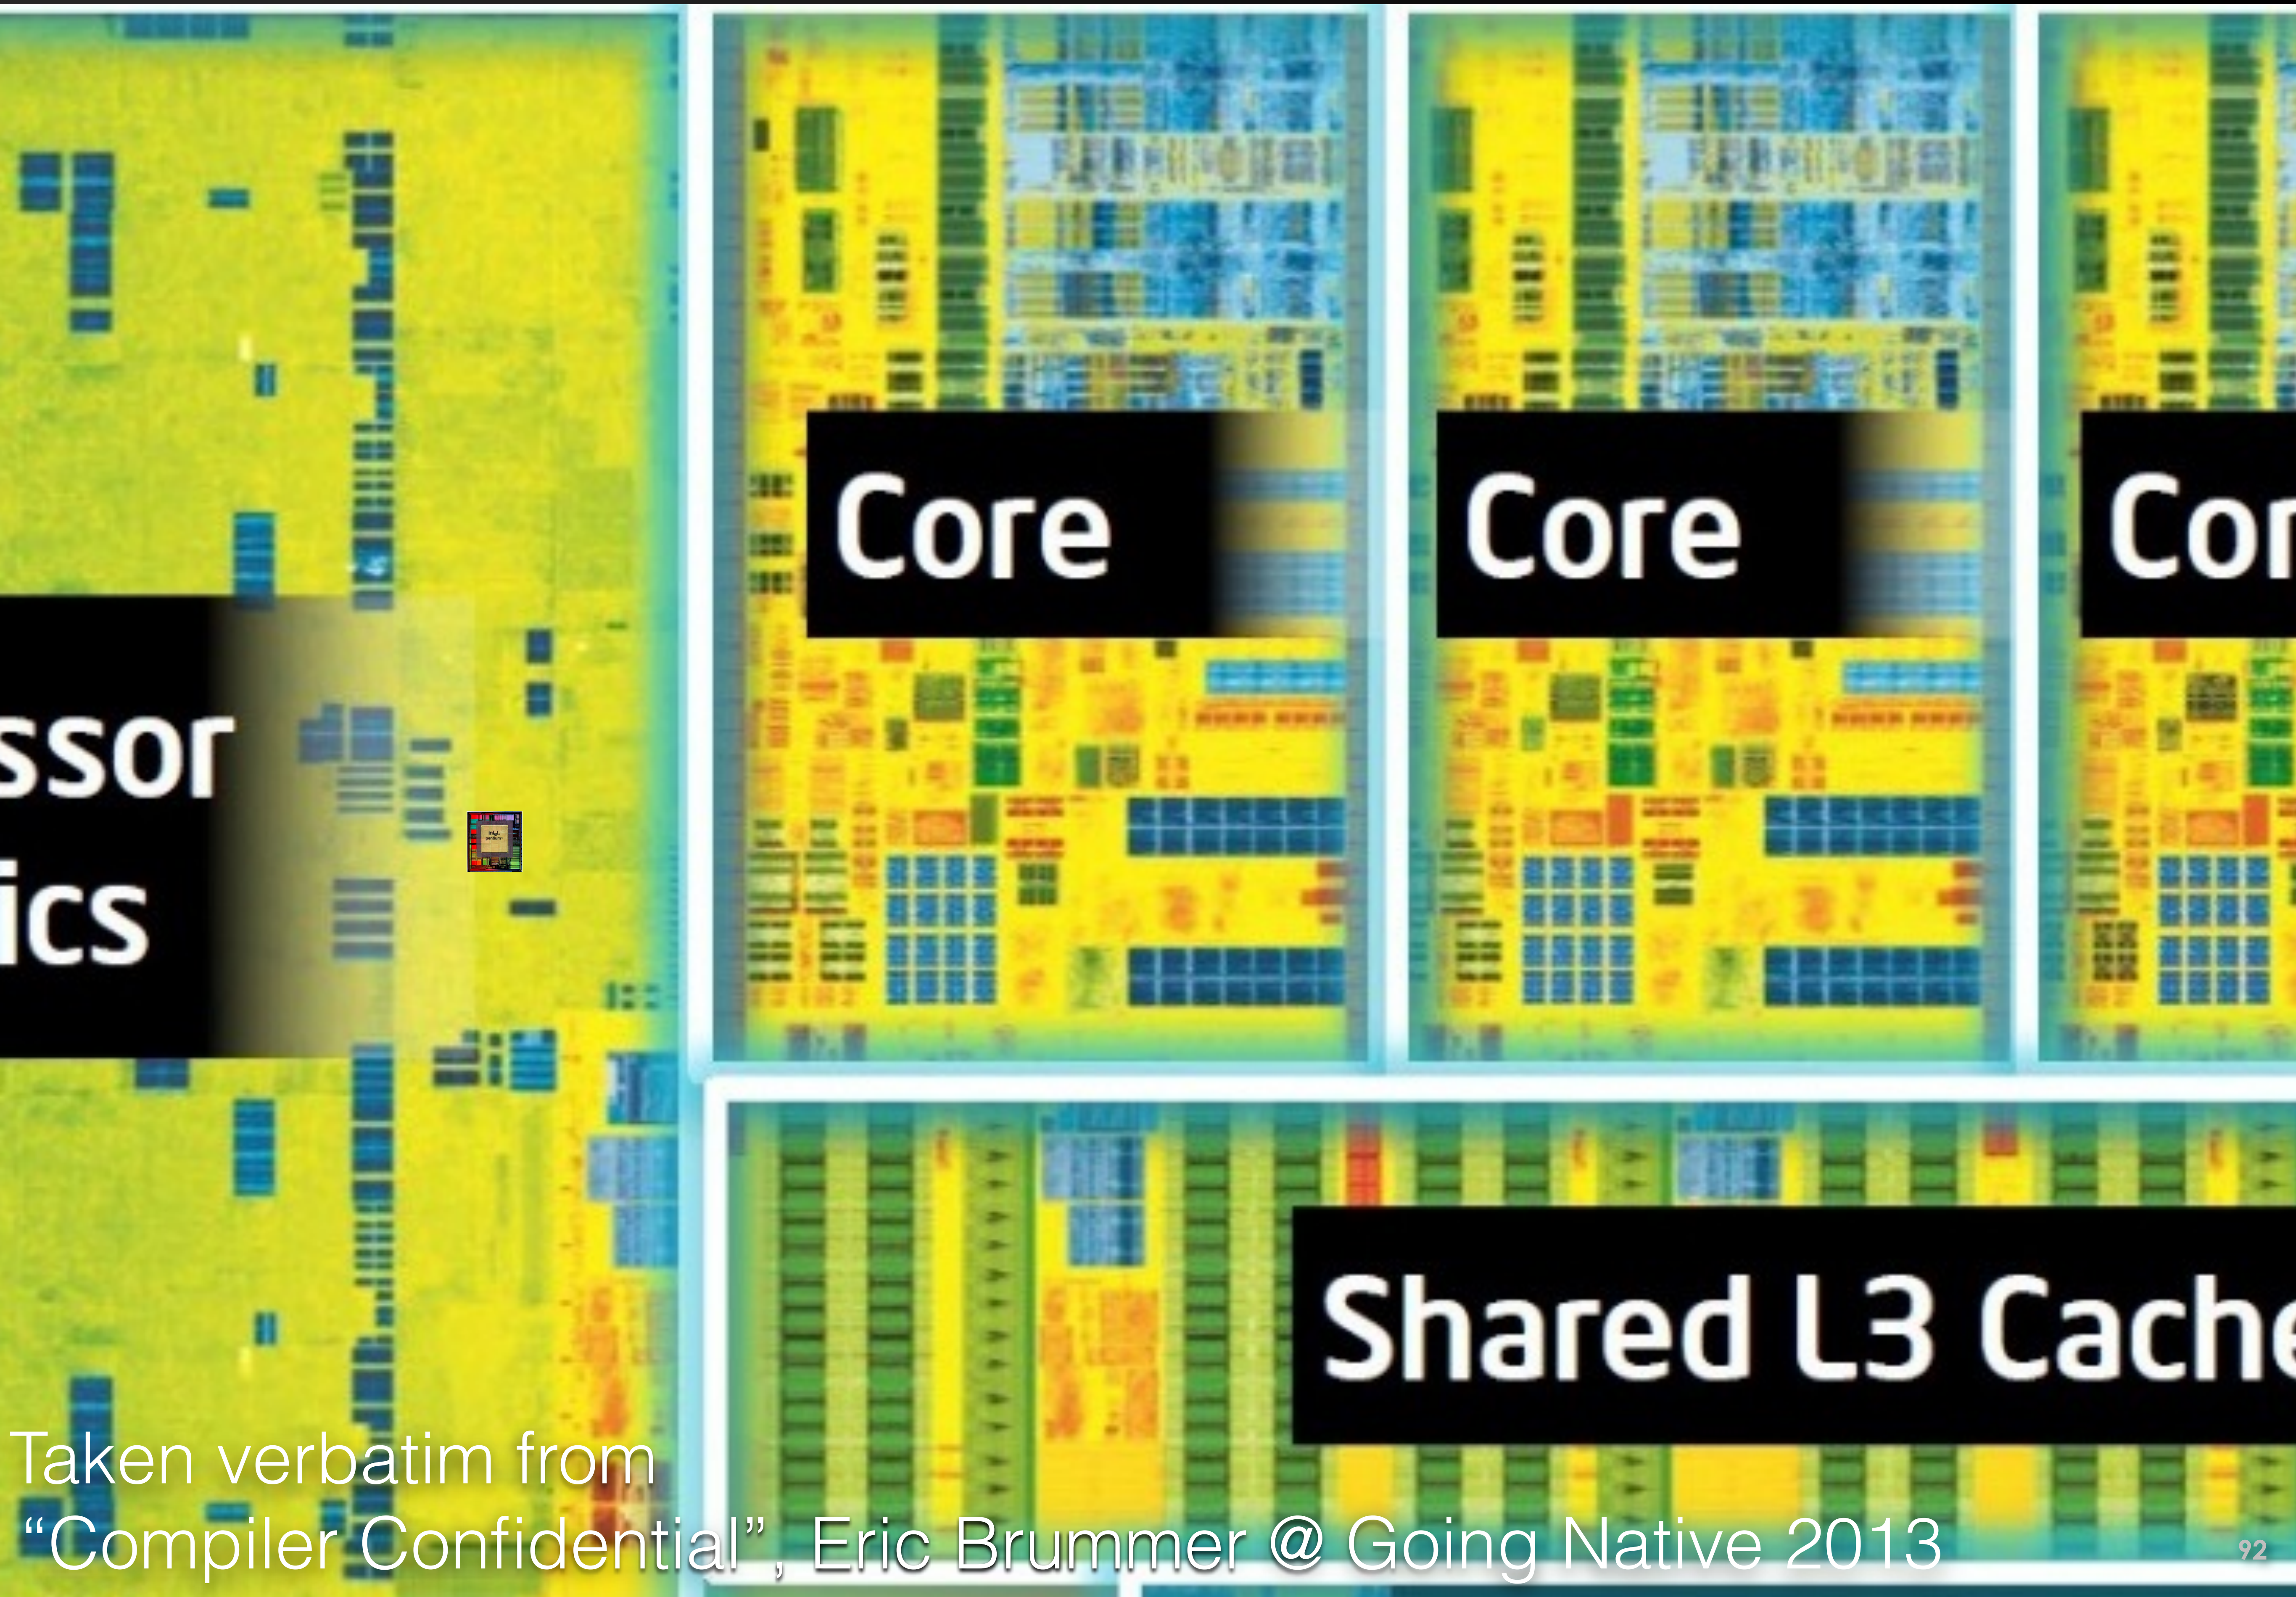

Taken verbatim from
“Compiler Confidential”, Eric Brummer @ Going Native 2013

The LHCb Upgrade: post-LS2

- After LS2, LHCb will run at 5x higher luminosity: $L = 2 \cdot 10^{33}$ $cm^{-2}s^{-1}$

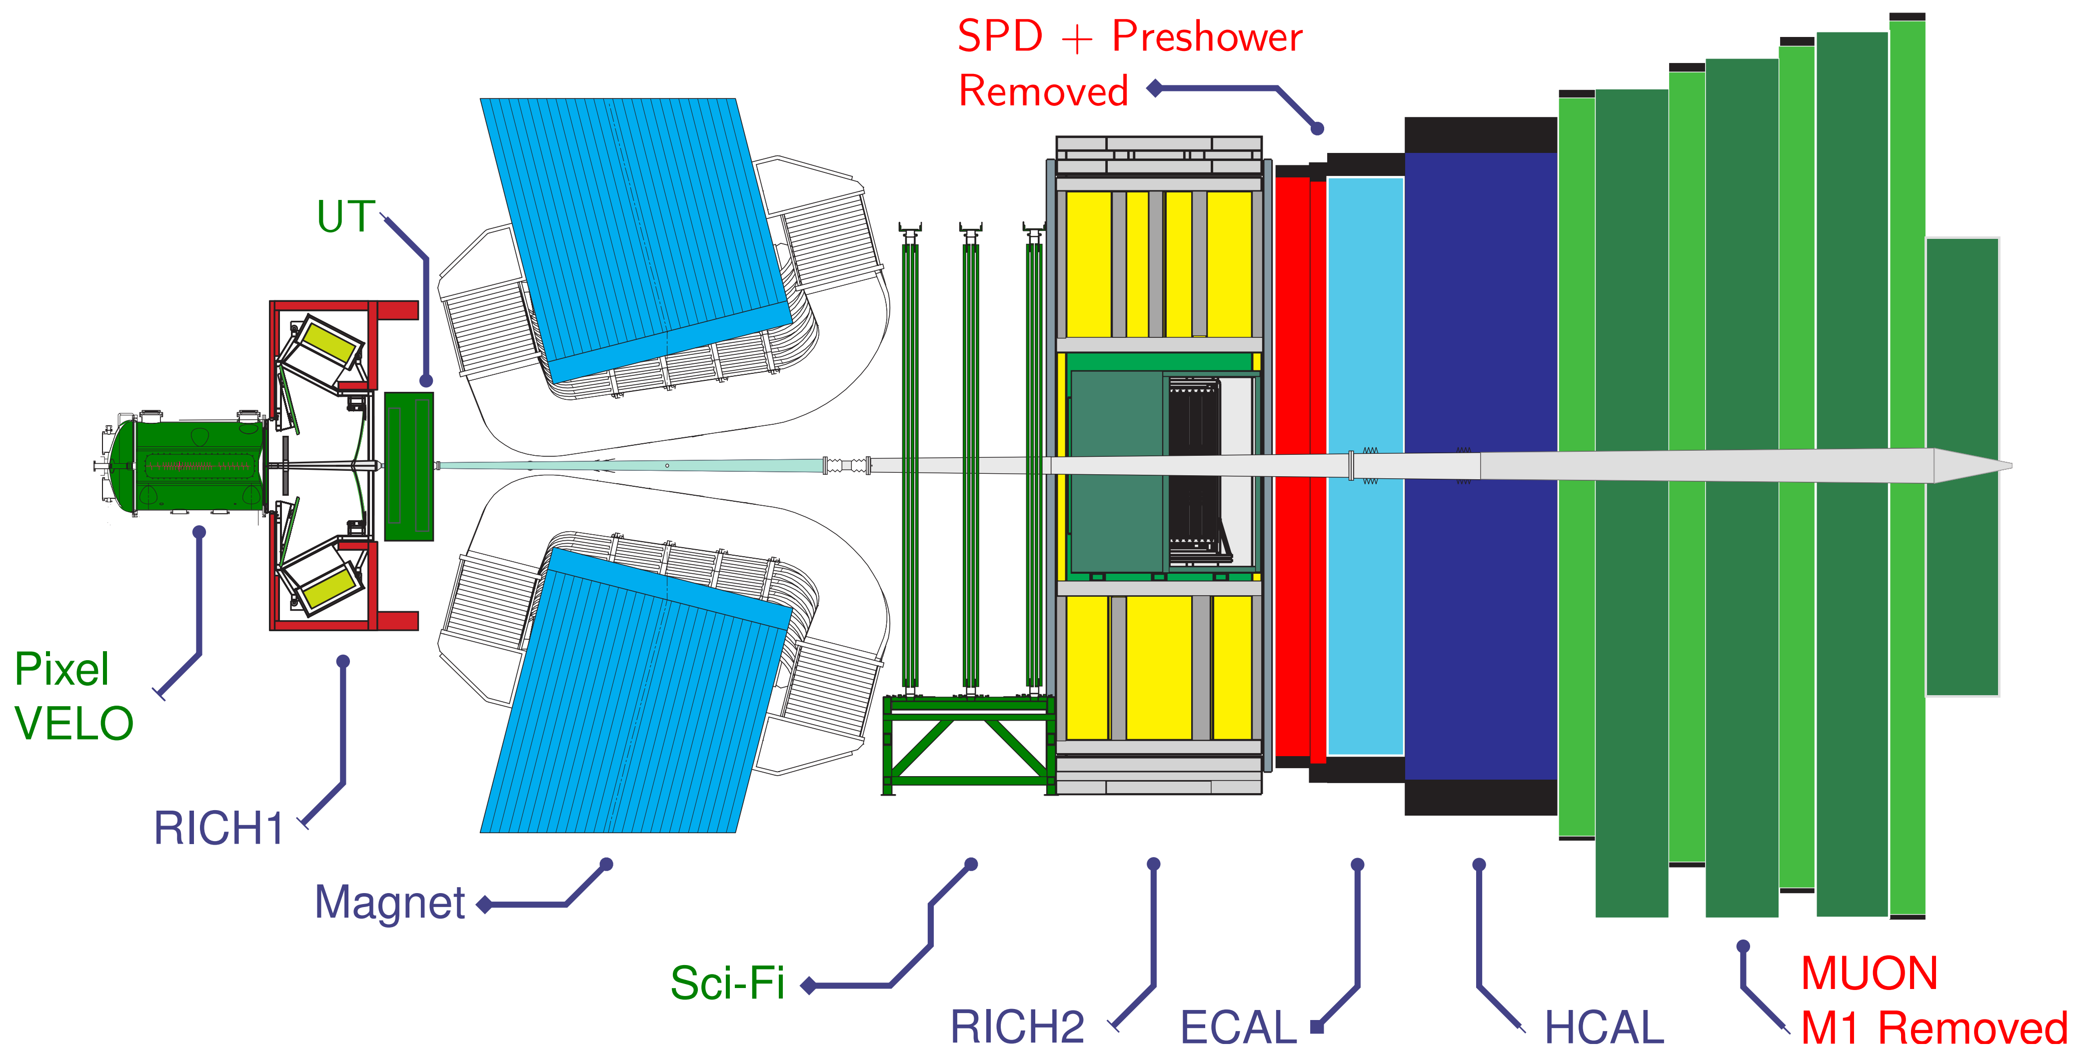

- VELO: r,ϕ strips →pixels
- Trackers: strawtubes → scintillating fibers + silicon microstrips
- RICH: replace photon detectors; CALO: SPD, PRS removed; MUON: M1 removed

LHCb Upgrade Environment

- Average pp collisions per bunch crossing: 1.5→7.6

Run I	Per event	with vertex in VELO	Rate [GB/s]
b-hadrons	0.0258 ± 0.0004	0.0029 ± 0.0001	0.9
c-hadrons	0.297 ± 0.001	0.0422 ± 0.0005	3.3
light, long-lived hadrons	8.04 ± 0.01	0.511 ± 0.002	1.1
Upgrade	**Per event**	**with vertex in VELO**	**Rate [GB/s]**
b-hadrons	0.1572 ± 0.0004	0.01874 ± 0.0001	27
c-hadrons	1.422 ± 0.001	0.2138 ± 0.0005	80
light, long-lived hadrons	33.291 ± 0.006	2.084 ± 0.001	26

- Challenge: must go beyond rejecting background —- classify signal, and choose wisely…

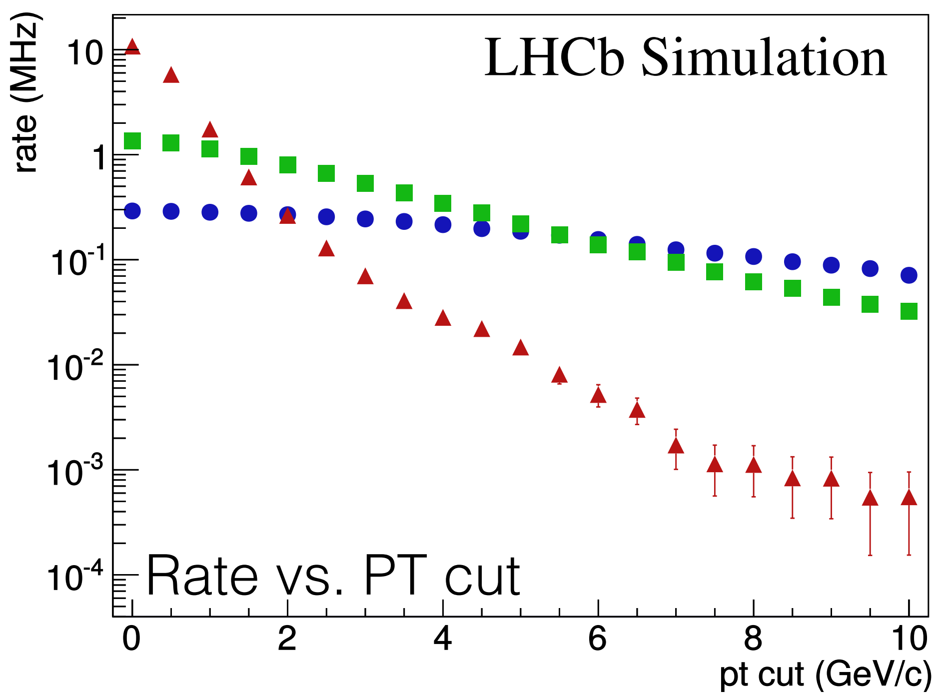

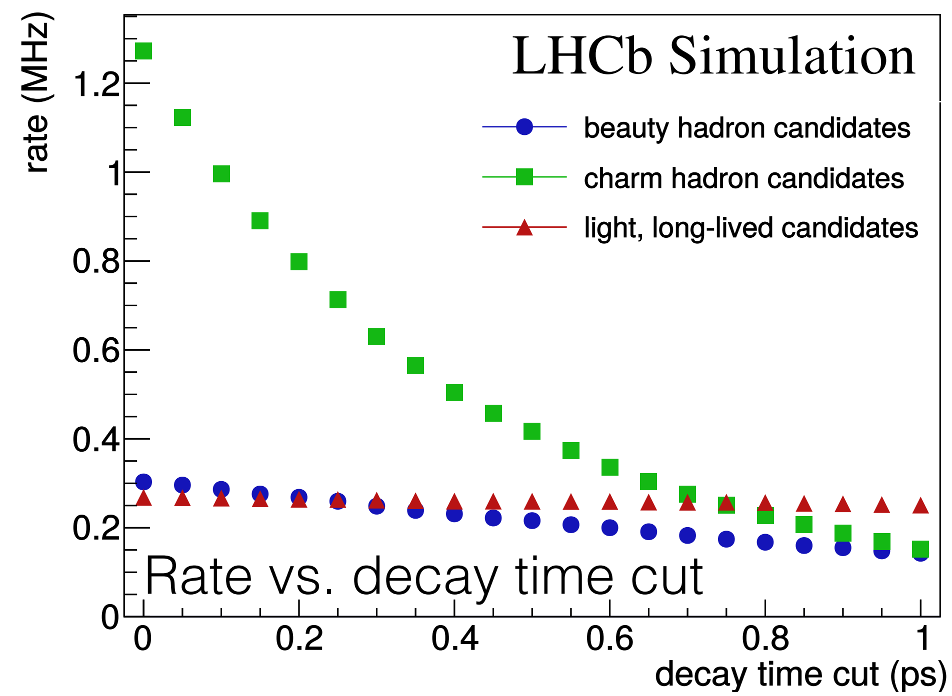

- P_T and IP alone not sufficient to reduce rate: requires *all* available detector information…

LHCb upgrade…

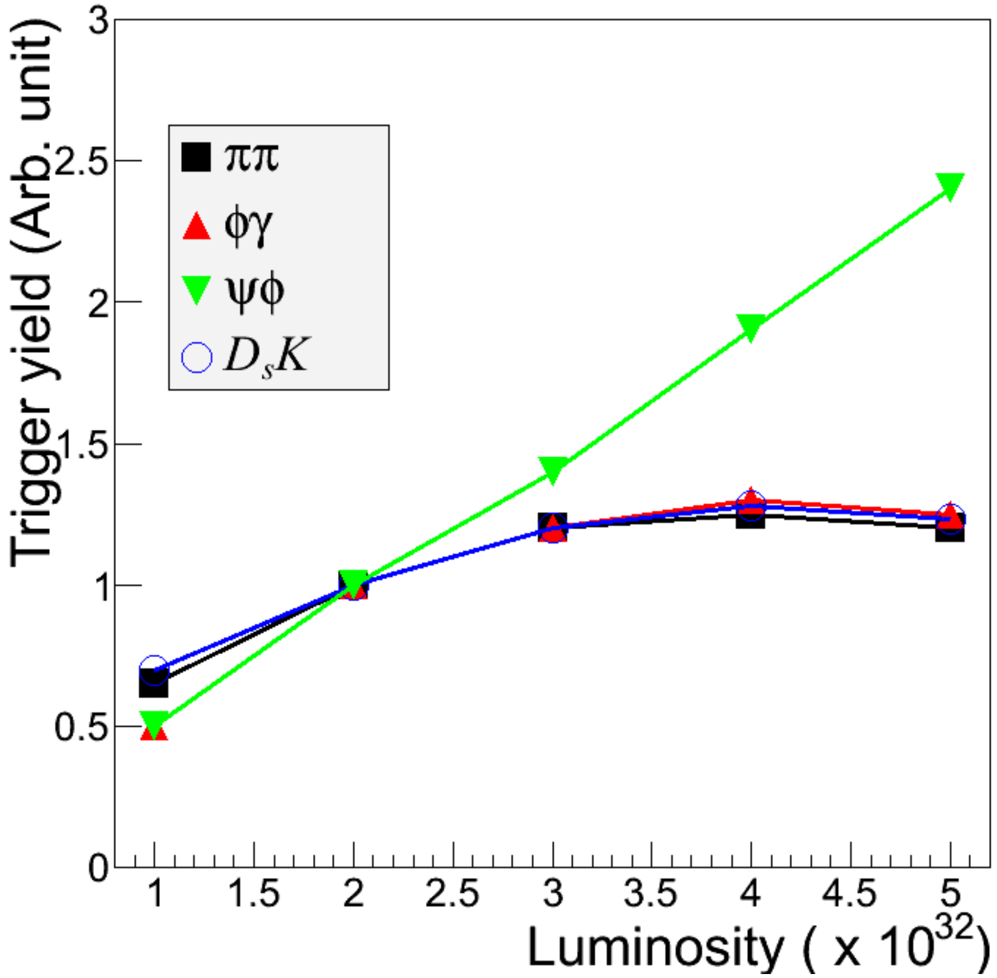

The 1 MHz readout becomes bottleneck:
L0 has insufficient information
to sufficiently purify the signal

1. build a better (and thus much more complicated!) L0
2. or …

Remember CMS?

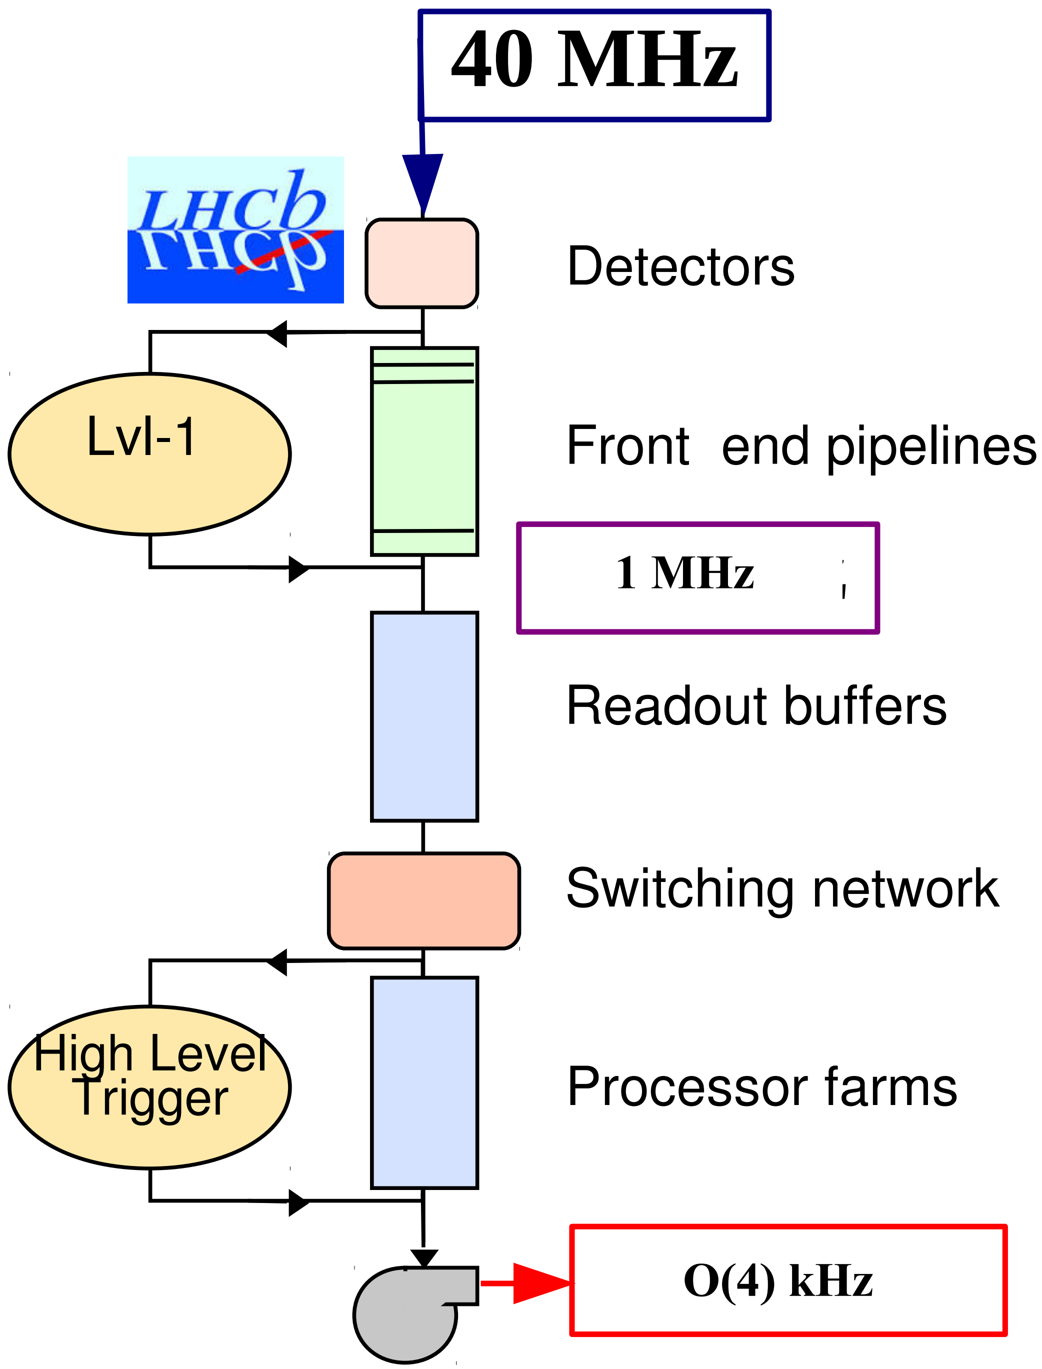

At CMS, TriDAS (Trigger and Data Acquisition System) Project Manager Sergio Cittolin, TriDAS Institution Board Chair Paris Sphicas and Smith decided not to have a second level trigger. They would take the output of L1 straight to the computer farm for software processing. The main reason for doing this was that the L2 hardware was too restrictive. It was not fully programmable, and was only used at the time because there was no telecom switch that could convey the full L1 output of 100 kHz of 1 MB events to the farm.

However, Cittolin, studying technology trends and extrapolating world-wide computing network infrastructure, was convinced that a switch with the required bandwidth would be available and affordable by the year 2000. So, when the technical proposal was written in '94-'95, a plan to go from L1 to the computer farm was laid out.

- from http://cms.web.cern.ch/news/triggering-and-data-acquisition

Remember CMS?

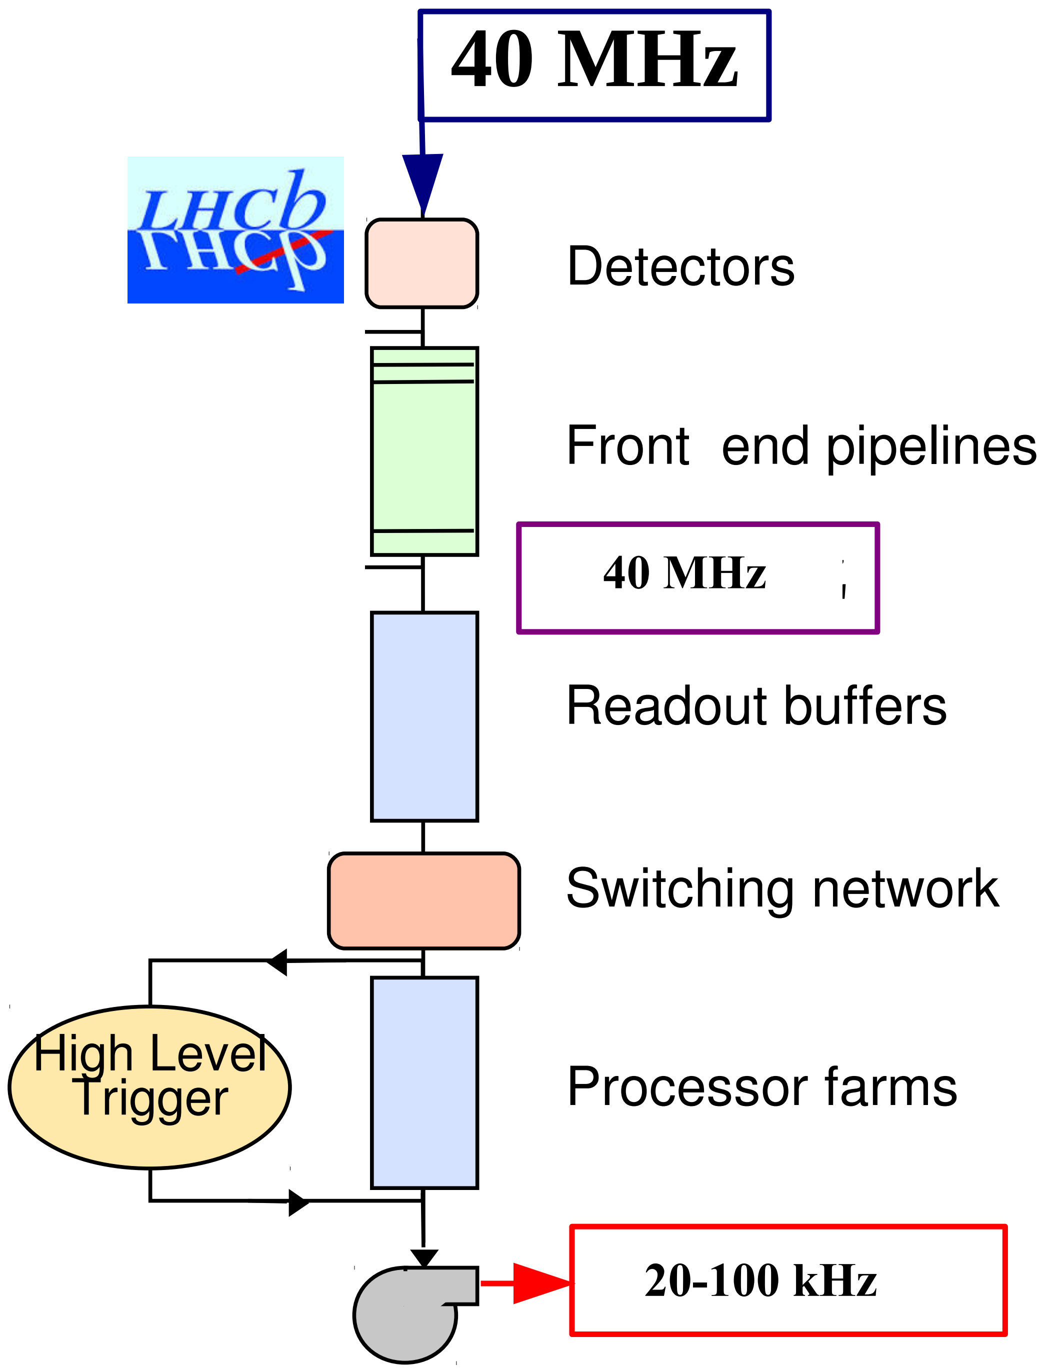

At CMS, TriDAS (Trigger and Data Acquisition System) Project Manager Sergio Cittolin, TriDAS Institution Board Chair Paris Sphicas and Smith decided not to have a second level trigger. They would take the output of L1 straight to the computer farm for software processing. The main reason for doing this was that the L2 hardware was too restrictive. It was not fully programmable, and was only used at the time because there was no telecom switch that could convey the full L1 output of 100 kHz of 1 MB events to the farm.

However, Cittolin, studying technology trends and extrapolating world-wide computing network infrastructure, was convinced that a switch with the required bandwidth would be available and affordable by the year 2000. So, when the technical proposal was written in '94-'95, a plan to go from L1 to the computer farm was laid out.

- from http://cms.web.cern.ch/news/triggering-and-data-acquisition

Triggerless Readout @ 40 MHz

- At L= 2 ·$10^{33}$$cm^{-2}s^{-1}$, the 1 MHz readout limit becomes a bottleneck
  - Signal no longer (sufficiently) easy to identify
- Readout upgraded to 40 MHz
- ⇒ Ship *every* visible pp

  interaction (30 MHz) to a CPU farm running the Higher Level Trigger

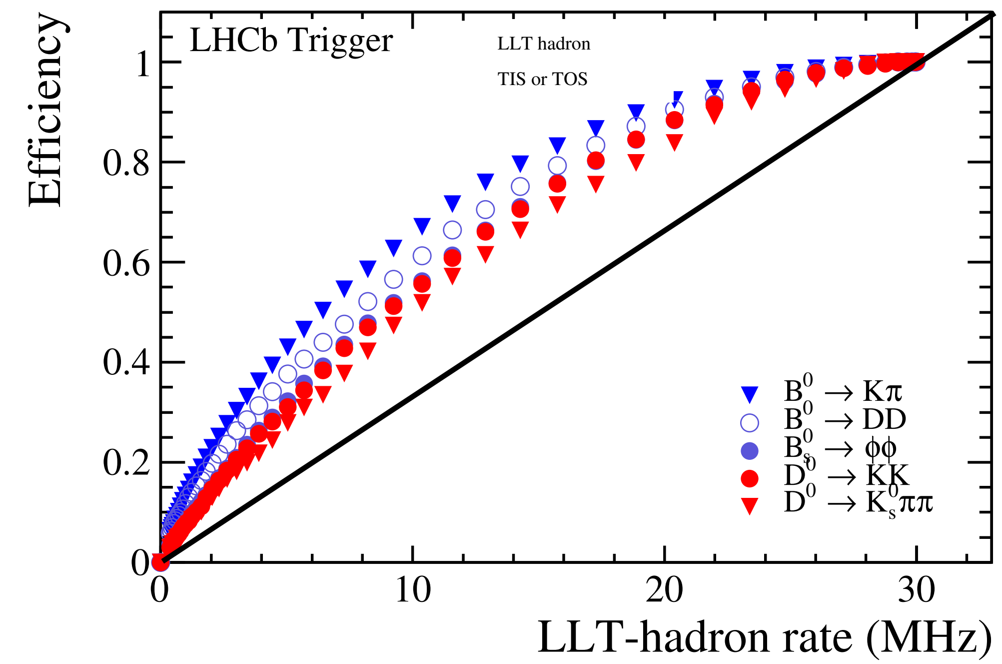

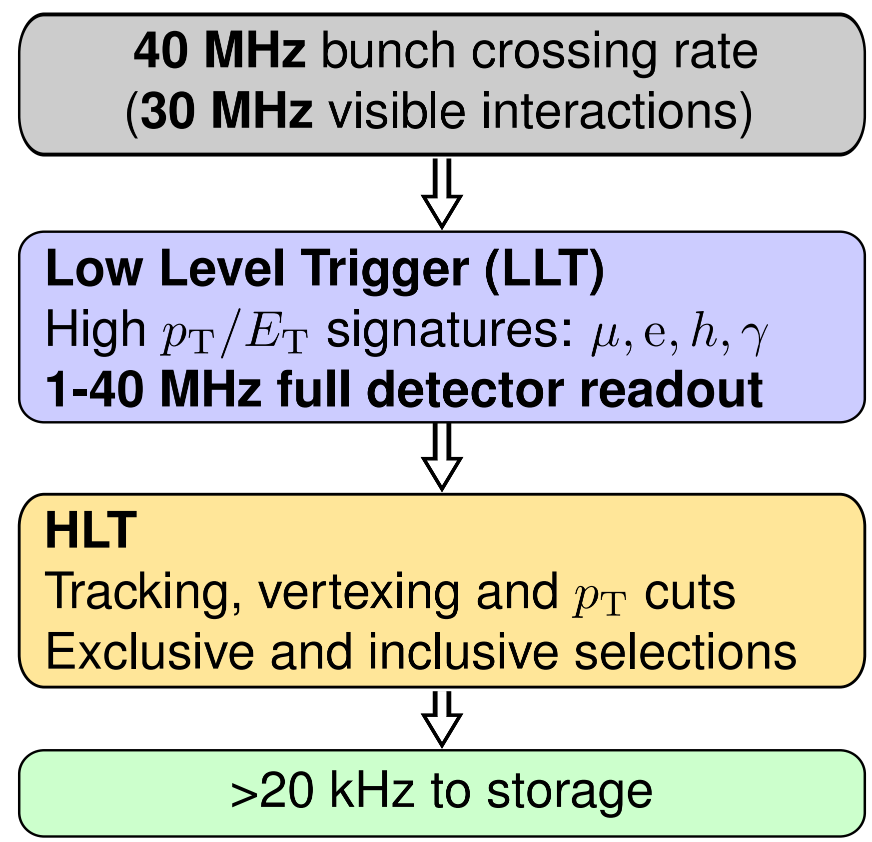

Triggerless Readout @ 40 MHz

- At L= 2 ·10^{33}cm^{-2}s^{-1}, the 1 MHz readout limit becomes a bottleneck
  - Signal no longer (sufficiently) easy to identify
- Readout upgraded to 40 MHz
- ⇒ Ship *every* visible pp

  interaction (30 MHz) to a CPU farm running the Higher Level Trigger

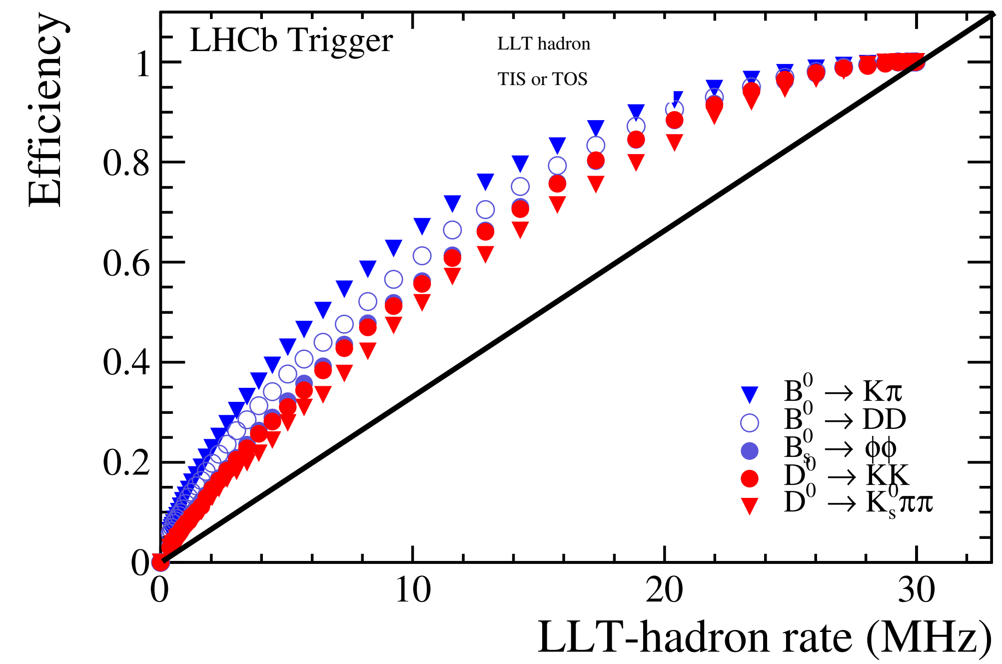

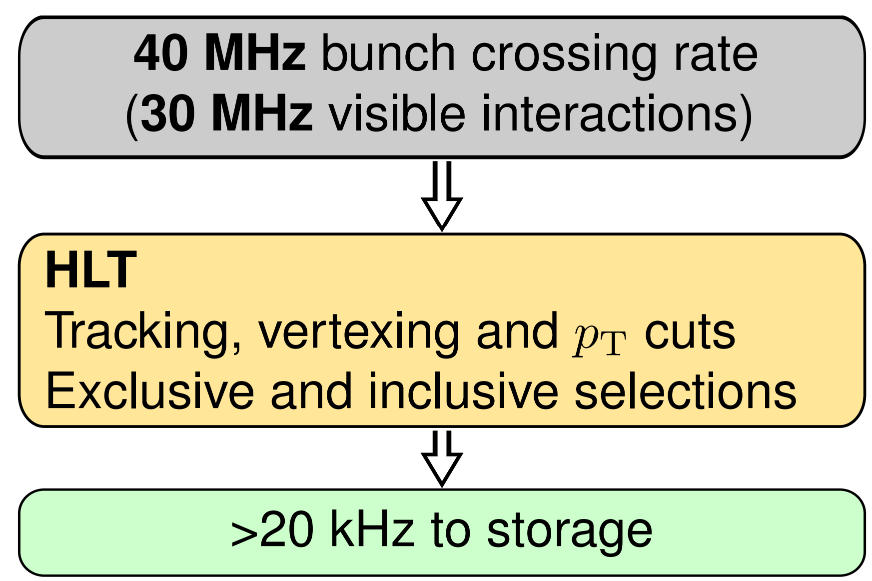

Event Building @ 40 MHz

- 32 Tbit/s
- "All data to the surface"
- Decouple front-end electronics from event builder network
    - Frontend→GBT→PCIe
    - GBT: Rad-hard, integrated into front-end, so no commodity solution possible…
- Buffering in PC memory

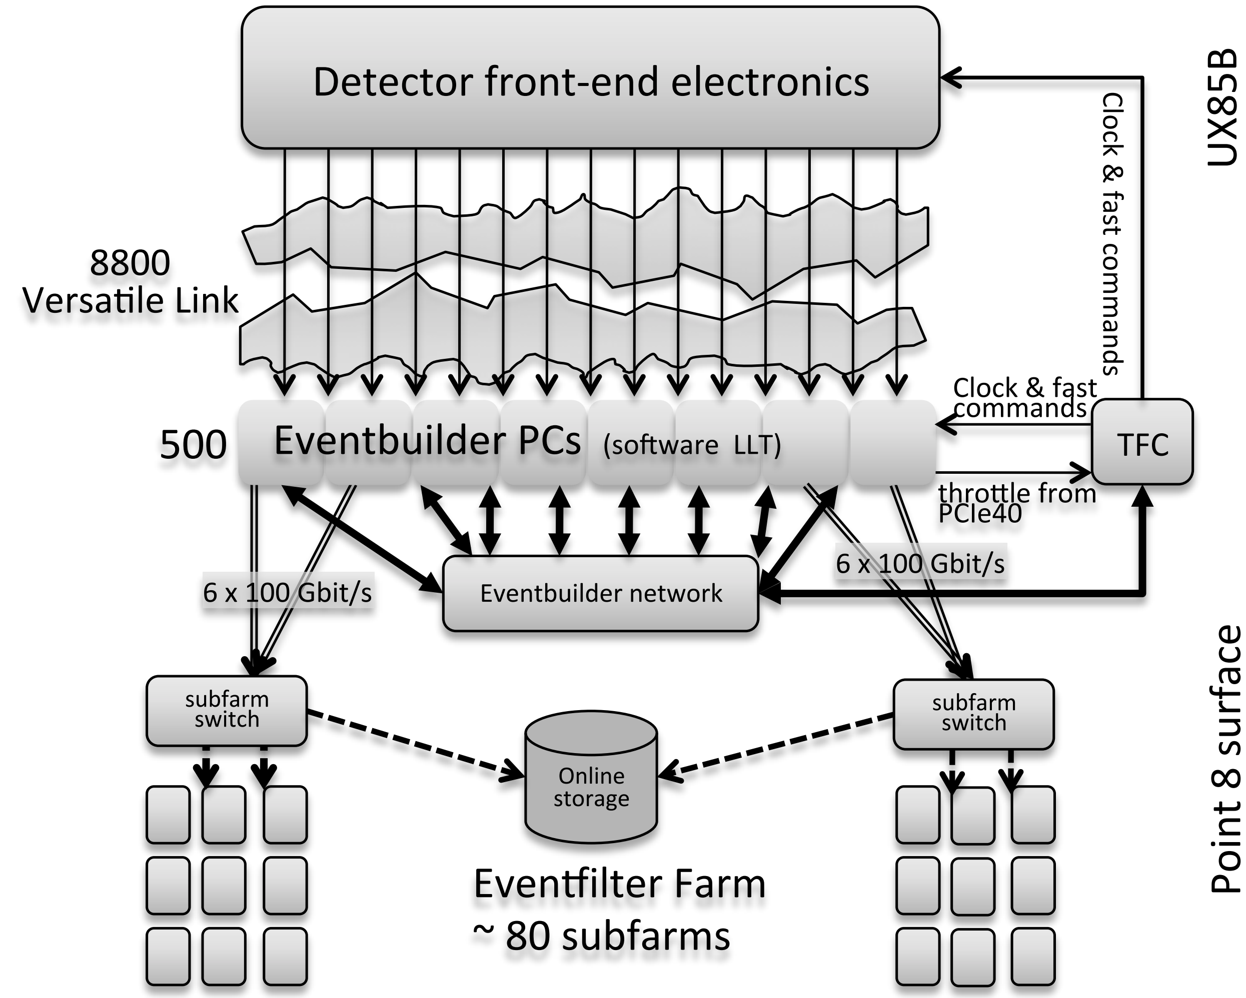

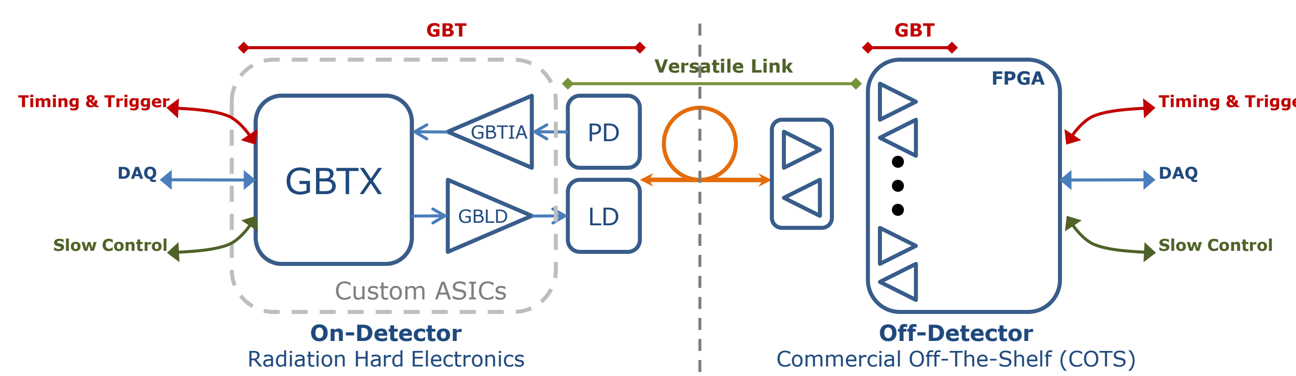

Event Building @ 40 MHz

- “COTS” as soon as possible
- O(500) servers for event building
- “Data Center” (“thin” switch, Infiniband/Ethernet/OmniPath) instead of “Telecom” (ATCA, “fat” switch)
- Event Filter: O(1000) servers

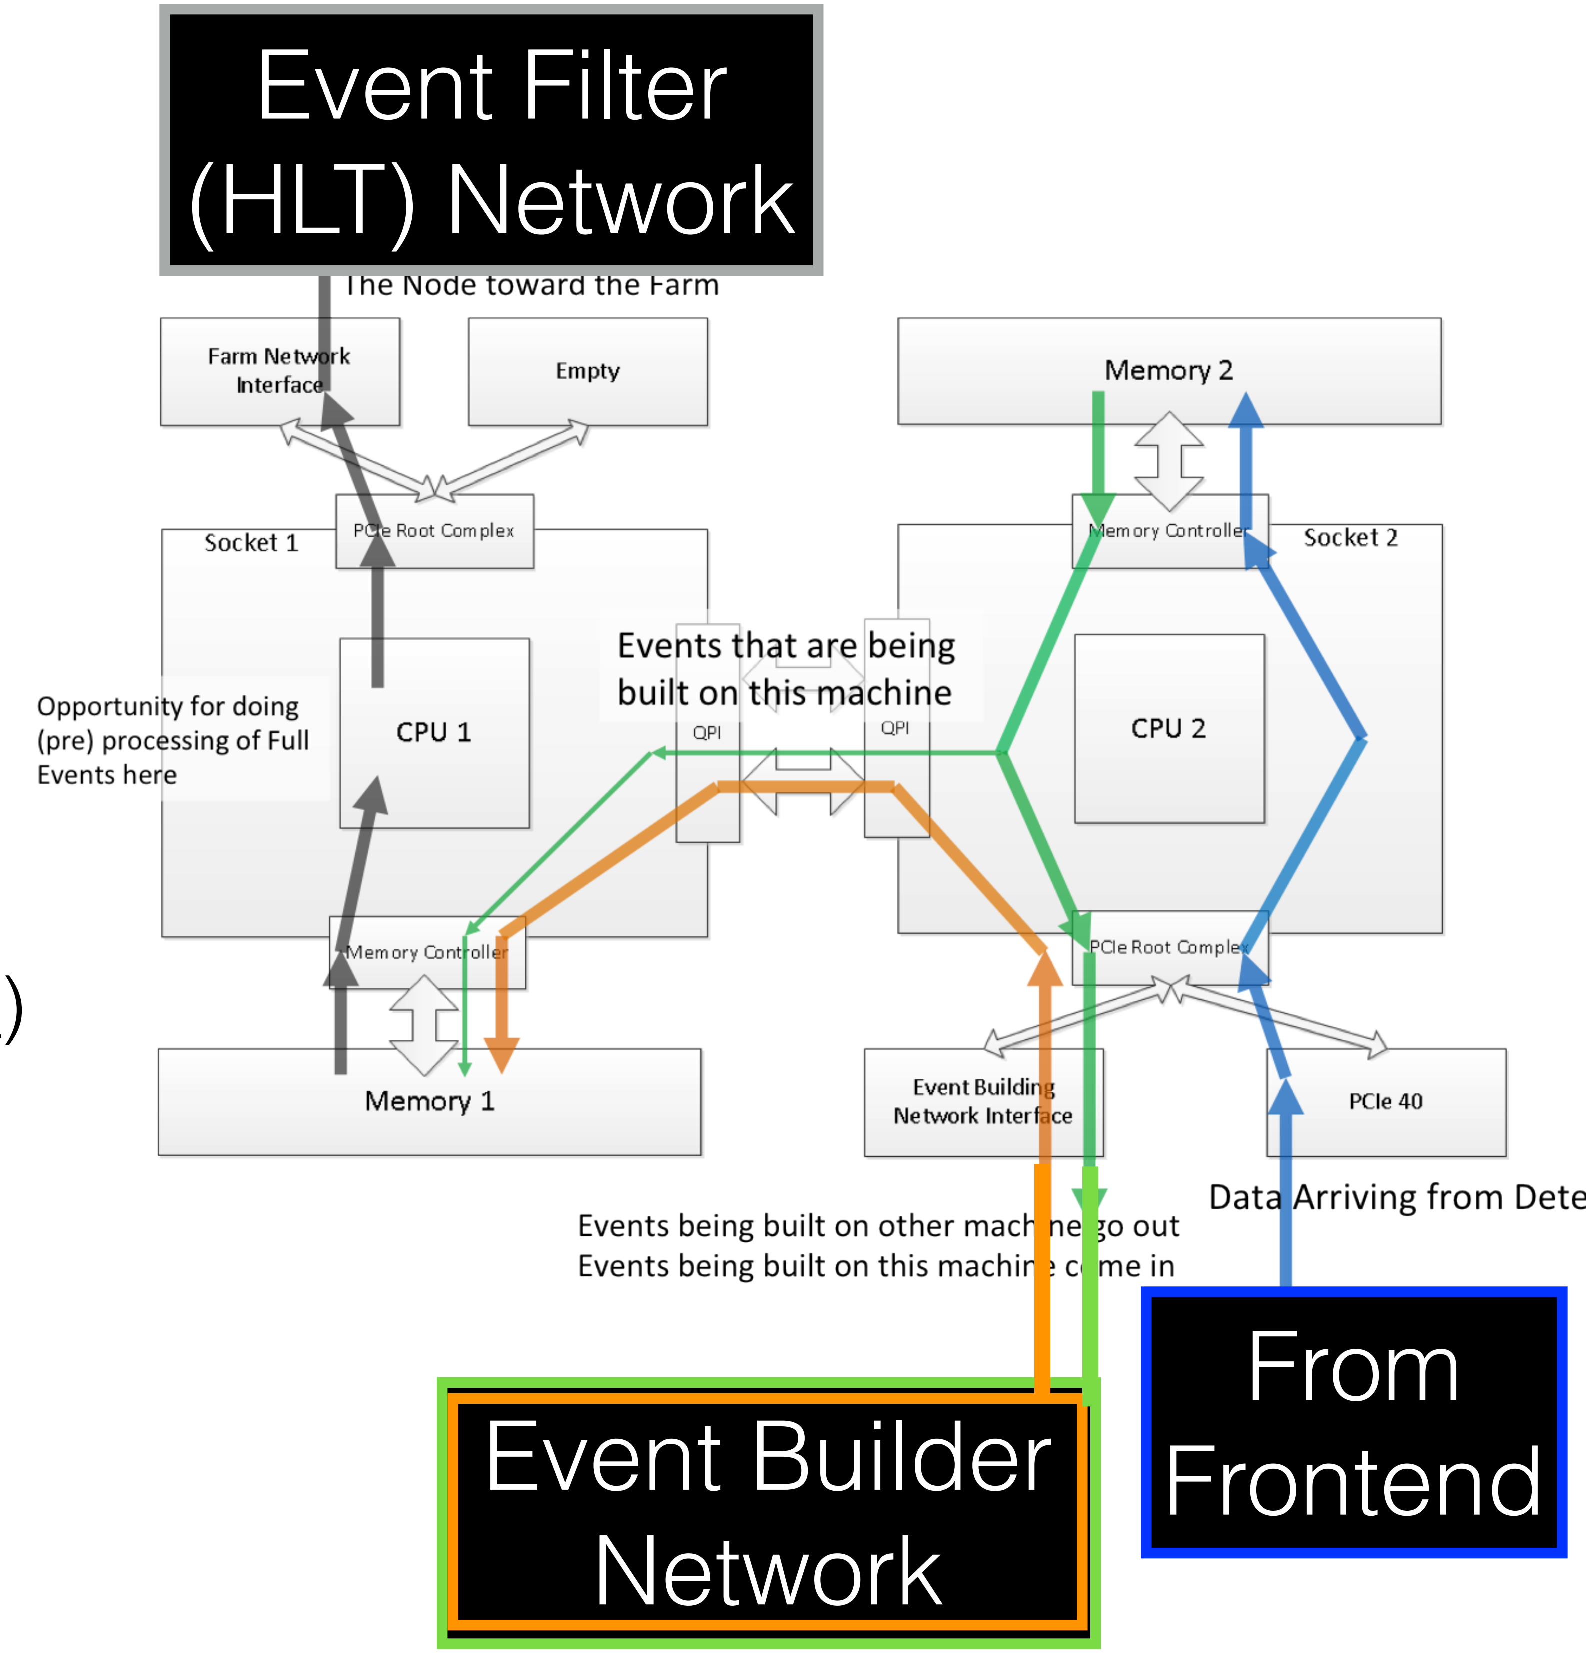

Upgrade Trigger

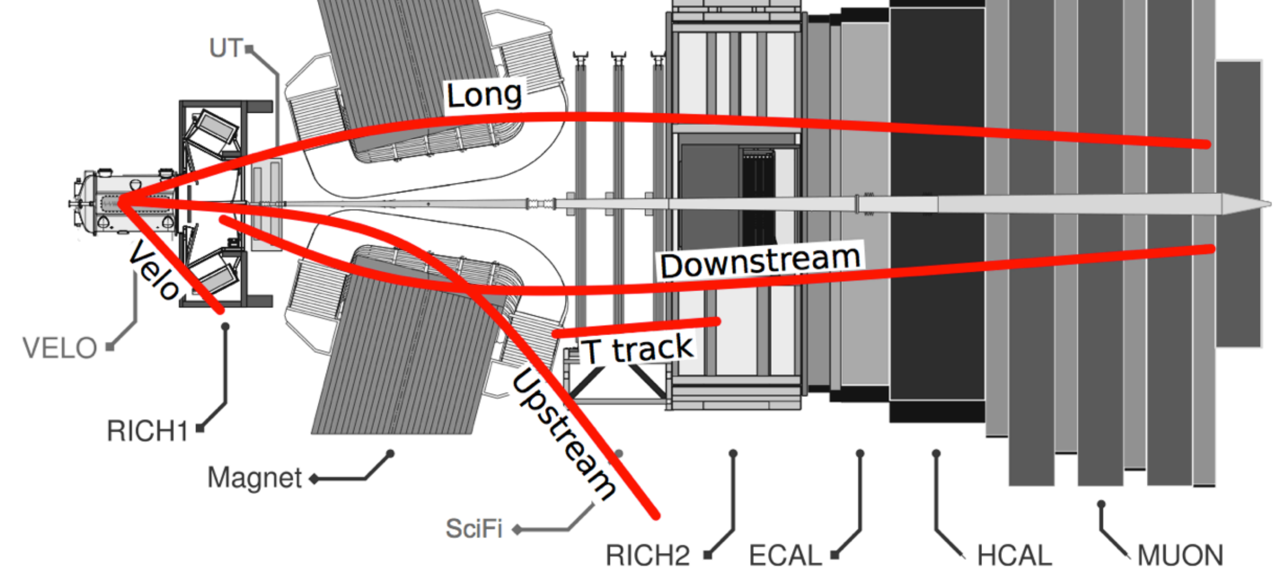

- Offline-quality tracking, *in software*, is possible @ 30 MHz
  - Estimated trigger farm: O(1000) servers
  - Tracking requires 5.4 ms/event, out of an *estimated* budget of 13 ms/event @ 30 MHz (*)
  - Thanks to the upgraded vertex detector & tracker designs!
- Converge online and offline reconstruction.

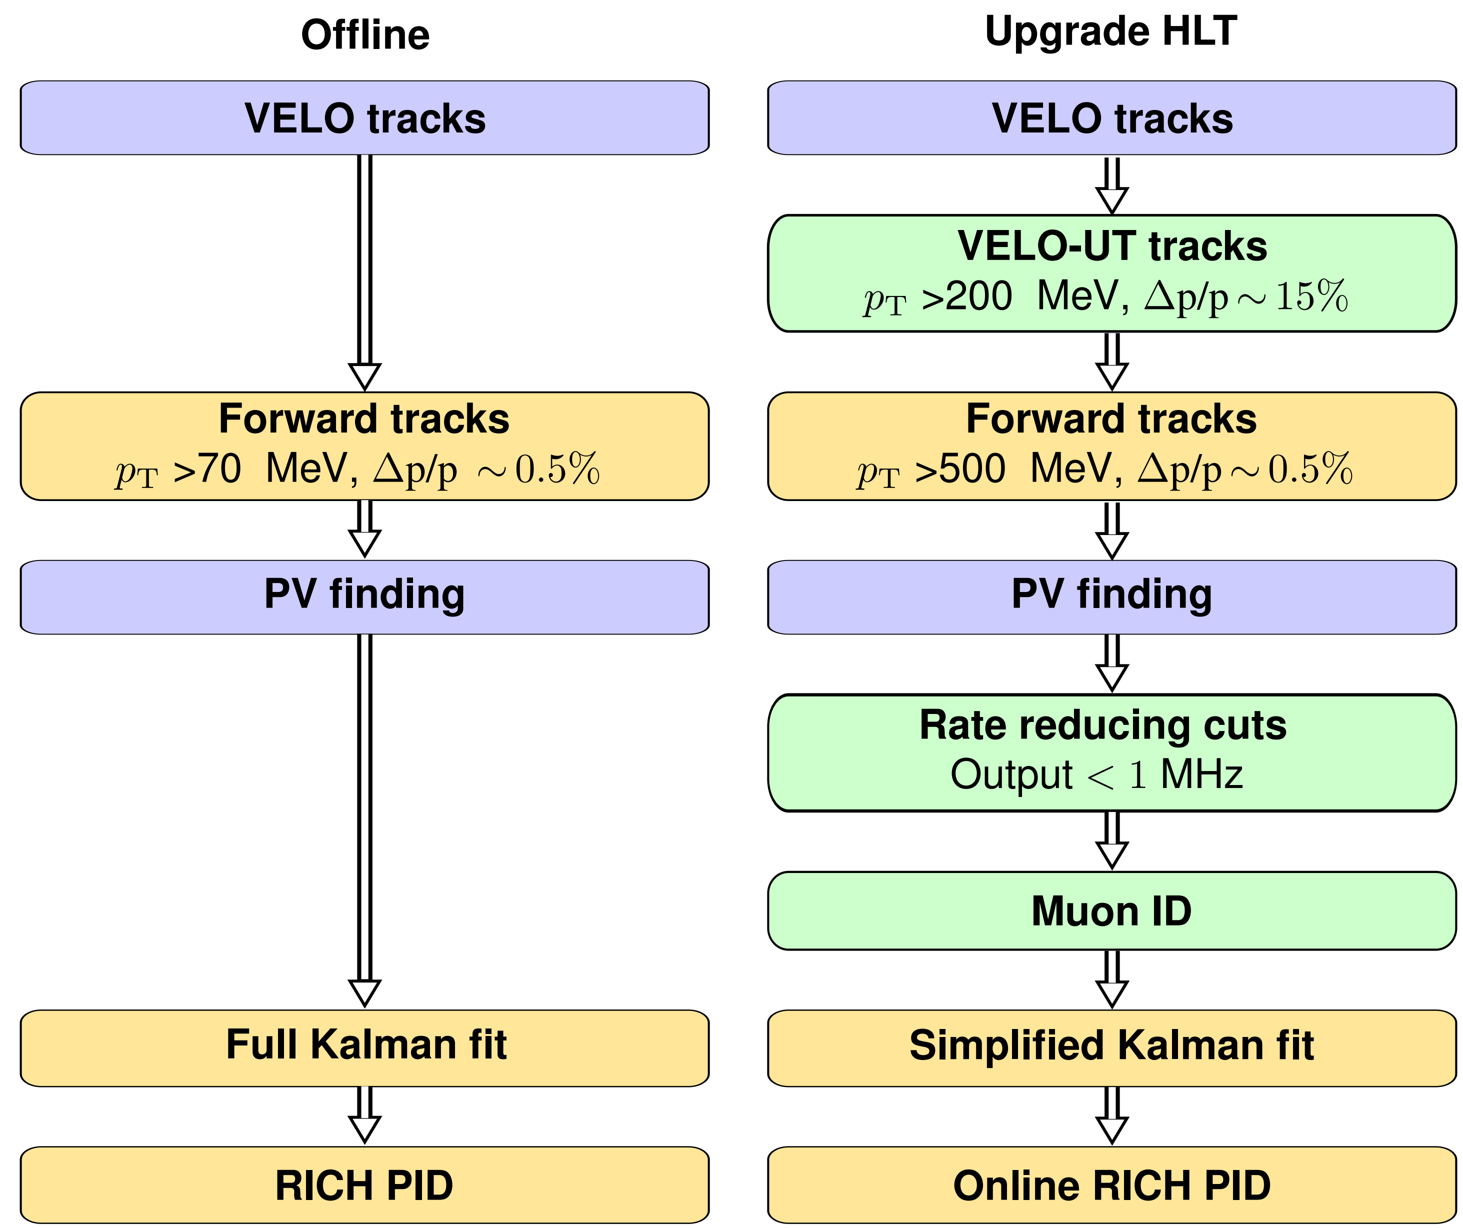

(*) on our 2011 reference machine: Intel X5650 (Westmere) @ 2.67 GHz

Upgrade Topological Trigger

- Same principle as Run 1 : preselect displaced tracks with $\sum P_T$, followed by BBDT
- Timing: <0.1 ms (*)
- At 25-50 kHz output rate, large efficiency gains over Run 1
  - red: run 1 efficiency
  - green: 2x run 1 efficiency
- LHCb-PUB-2014-031

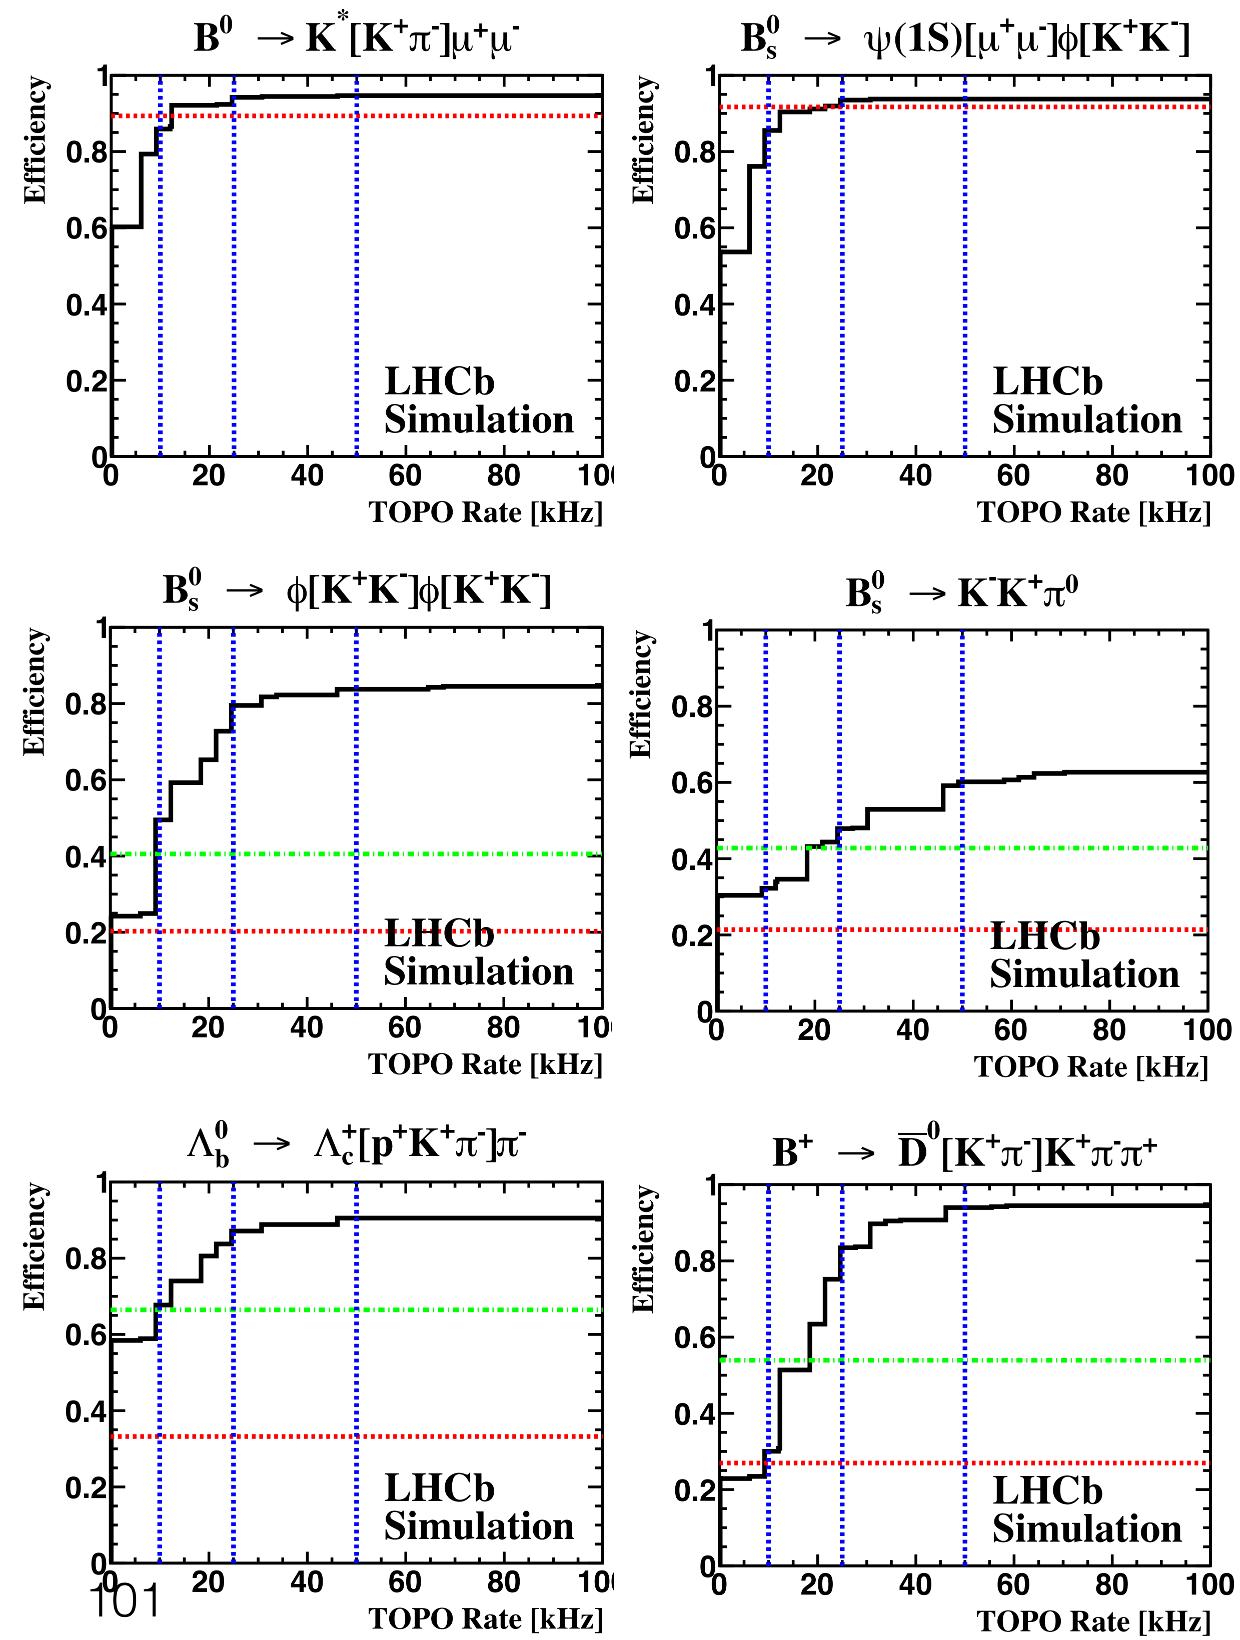

(*) on our 2011 reference machine: Intel X5650 (Westmere) @ 2.67 GHz

Trigger-less GPD(*)?

- for HL-LHC, Atlas/CMS are augmenting their trigger system by (hardware) track triggers
- Would a a 40 MHz trigger-less GPD be possible?
- HLT compute power:
  - 400x (rate) 3x (pileup) = O(1000) x today
  - $1000 >> 1.3^{15} = 50$ (and even that seems far too optimistic)
  - can co-processors help to close the gap?
- Event Builder: 5 MB @ 40 MHz
  - 2500 links @ 400 Gb/s
  - 25x today, in 10 years; maybe..
- Rad-hard Frontend links — not a commercial product!
  - GBT O(5 Gb/s) —> O(200K) links
  - way too much Cu [W. Smith — ECFA HL-LHC workshop 2014] [0.35g /m; P.Moreira, ECFA HL-LHC workshop 2014]

(*) GPD = General Purpose Detector

Trigger-less GPD(*)? Wireless!

- for HL-LHC, Atlas/CMS are augmenting their trigger system by (hardware) track triggers
- Would a a 40 MHz trigger-less GPD be possible?
- HLT compute power:
  - 400x (rate) 3x (pileup) = O(1000) x today
  - 1000 >> 1.3^{15} = 50 (and even that seems far too optimistic)
  - can co-processors help to close the gap?
- Event Builder: 5 MB @ 40 MHz
  - 2500 links @ 400 Gb/s
  - 25x today, in 10 years; maybe..
- Rad-hard Frontend links — not a commercial product!
  - GBT O(5 Gb/s) —> O(200K) links
  - way too much Cu [W. Smith — ECFA HL-LHC workshop 2014] [0.35g /m; P.Moreira, ECFA HL-LHC workshop 2014]
  - But keep dreaming: wireless transmission!

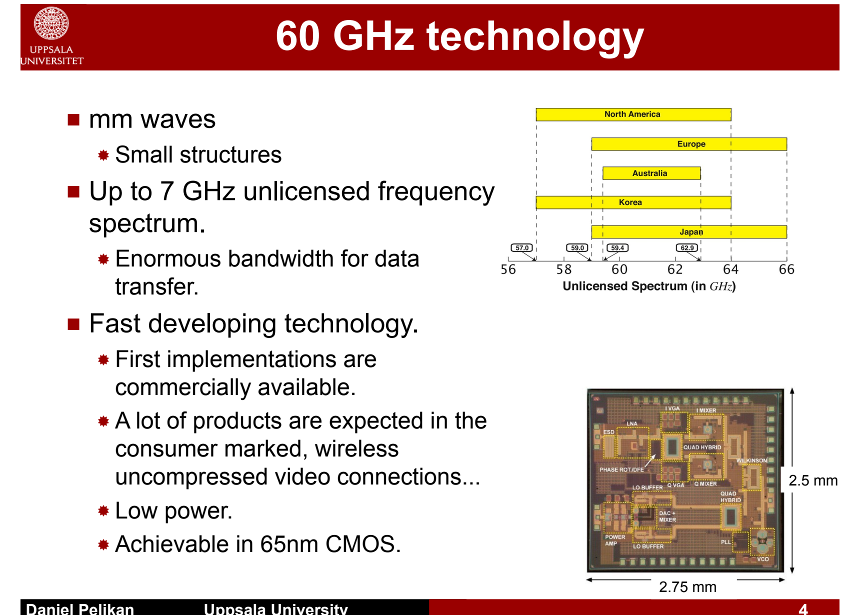

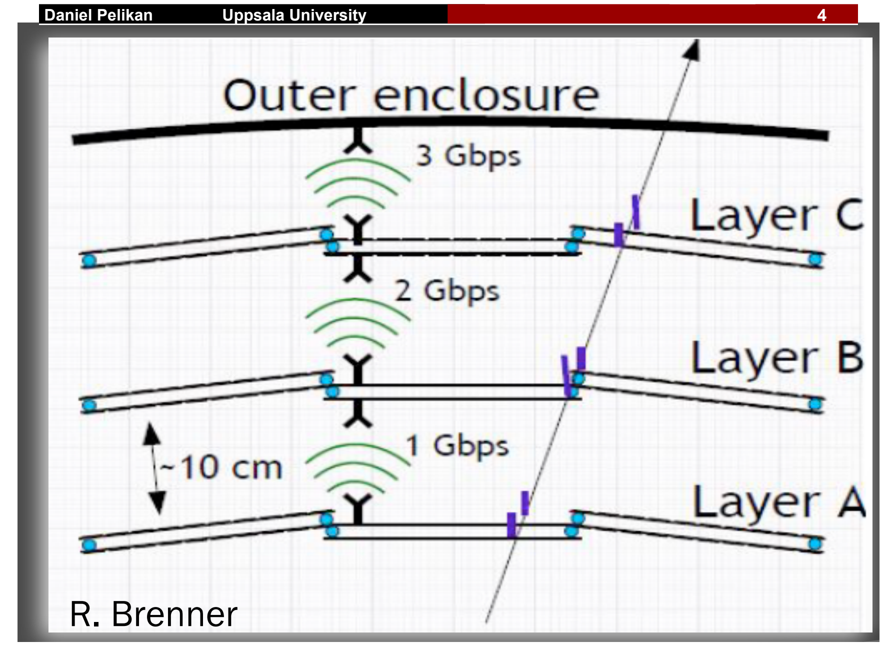

Specifications	Value
Frequency band	57-66 GHz
Bandwidth	9 GHz
Data Rate	4.5 Gbps
Target Power consumption	250 mW
Transmission Range	10 cm (1m)

(*) GPD = General Purpose Detector

The Future

- When possible, go ‘triggerless’, and shift burden to CPUs (+coprocessors) (LHCb, Alice)
- Real-time calibration has arrived (LHCb, Alice)

- Atlas & CMS will need L1 triggers
- Tracking at L1 should allow current thresholds in the HL-LHC era

- Evolution of processing power of (co-)processors crucial for HLT
  - Will we be able to use all the silicon effectively?

- There remains a lot to be done! Join the challenge!

Trigger & DAQ

Gerhard Raven
VU Amsterdam & Nikhef

- constraints
- physics
- architectures
- strategy
- examples

- resource optimization
- special cases
- interaction w. analysis
- upgrades

Acknowledgements:
Niko Neufeld, Vladimir Gligorov, Paris Sphicas, Brian Dahmes, ISOTDAQ lecturers, and many, many others…

2015 CERN-Fermilab
Hadron Collider Physics School

CERN White Paper

- “Data acquisition is where instruments meet IT systems. “
- “Costs and complexity must be reduced by replacing custom electronics with high-performance commodity processors and efficient software.”