

Real time event selection at the LHC

V. Gligorov, CERN

CEA Seminar, Orsay

19th February 2014

Pause for jargon

Real time event selection

==

Triggering

Why do we need triggers at the LHC?

Input data rate of one LHC
experiment = 1.5 TB/second

Why do we need triggers at the LHC?

Input data rate of one LHC
experiment = 1.5 TB/second



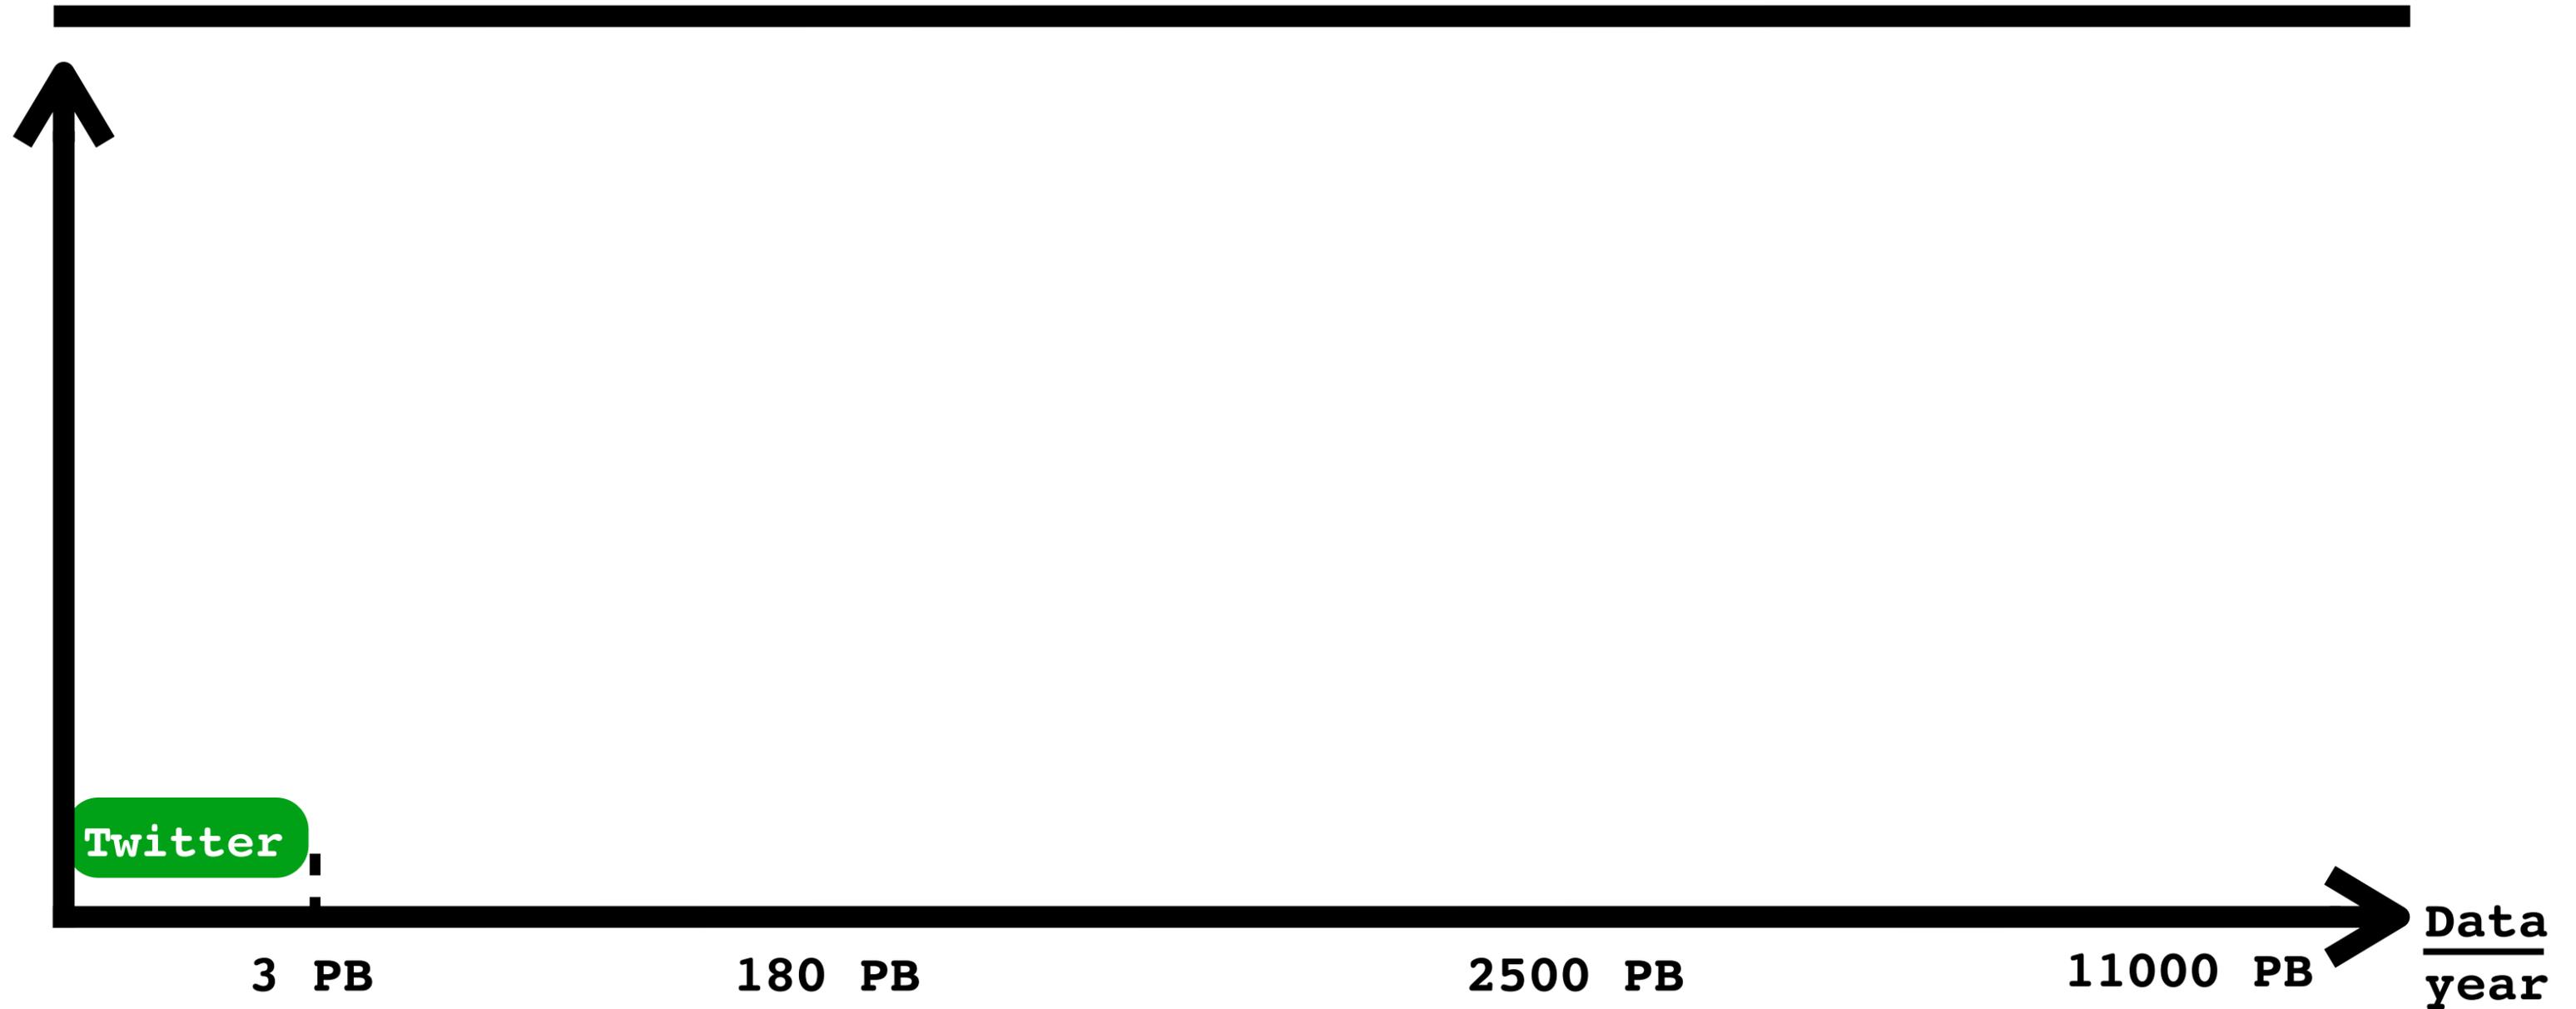
This means about 15000 PB
of data every year

Why do we need triggers at the LHC?

Input data rate of one LHC experiment = 1.5 TB/second



This means about 15000 PB of data every year

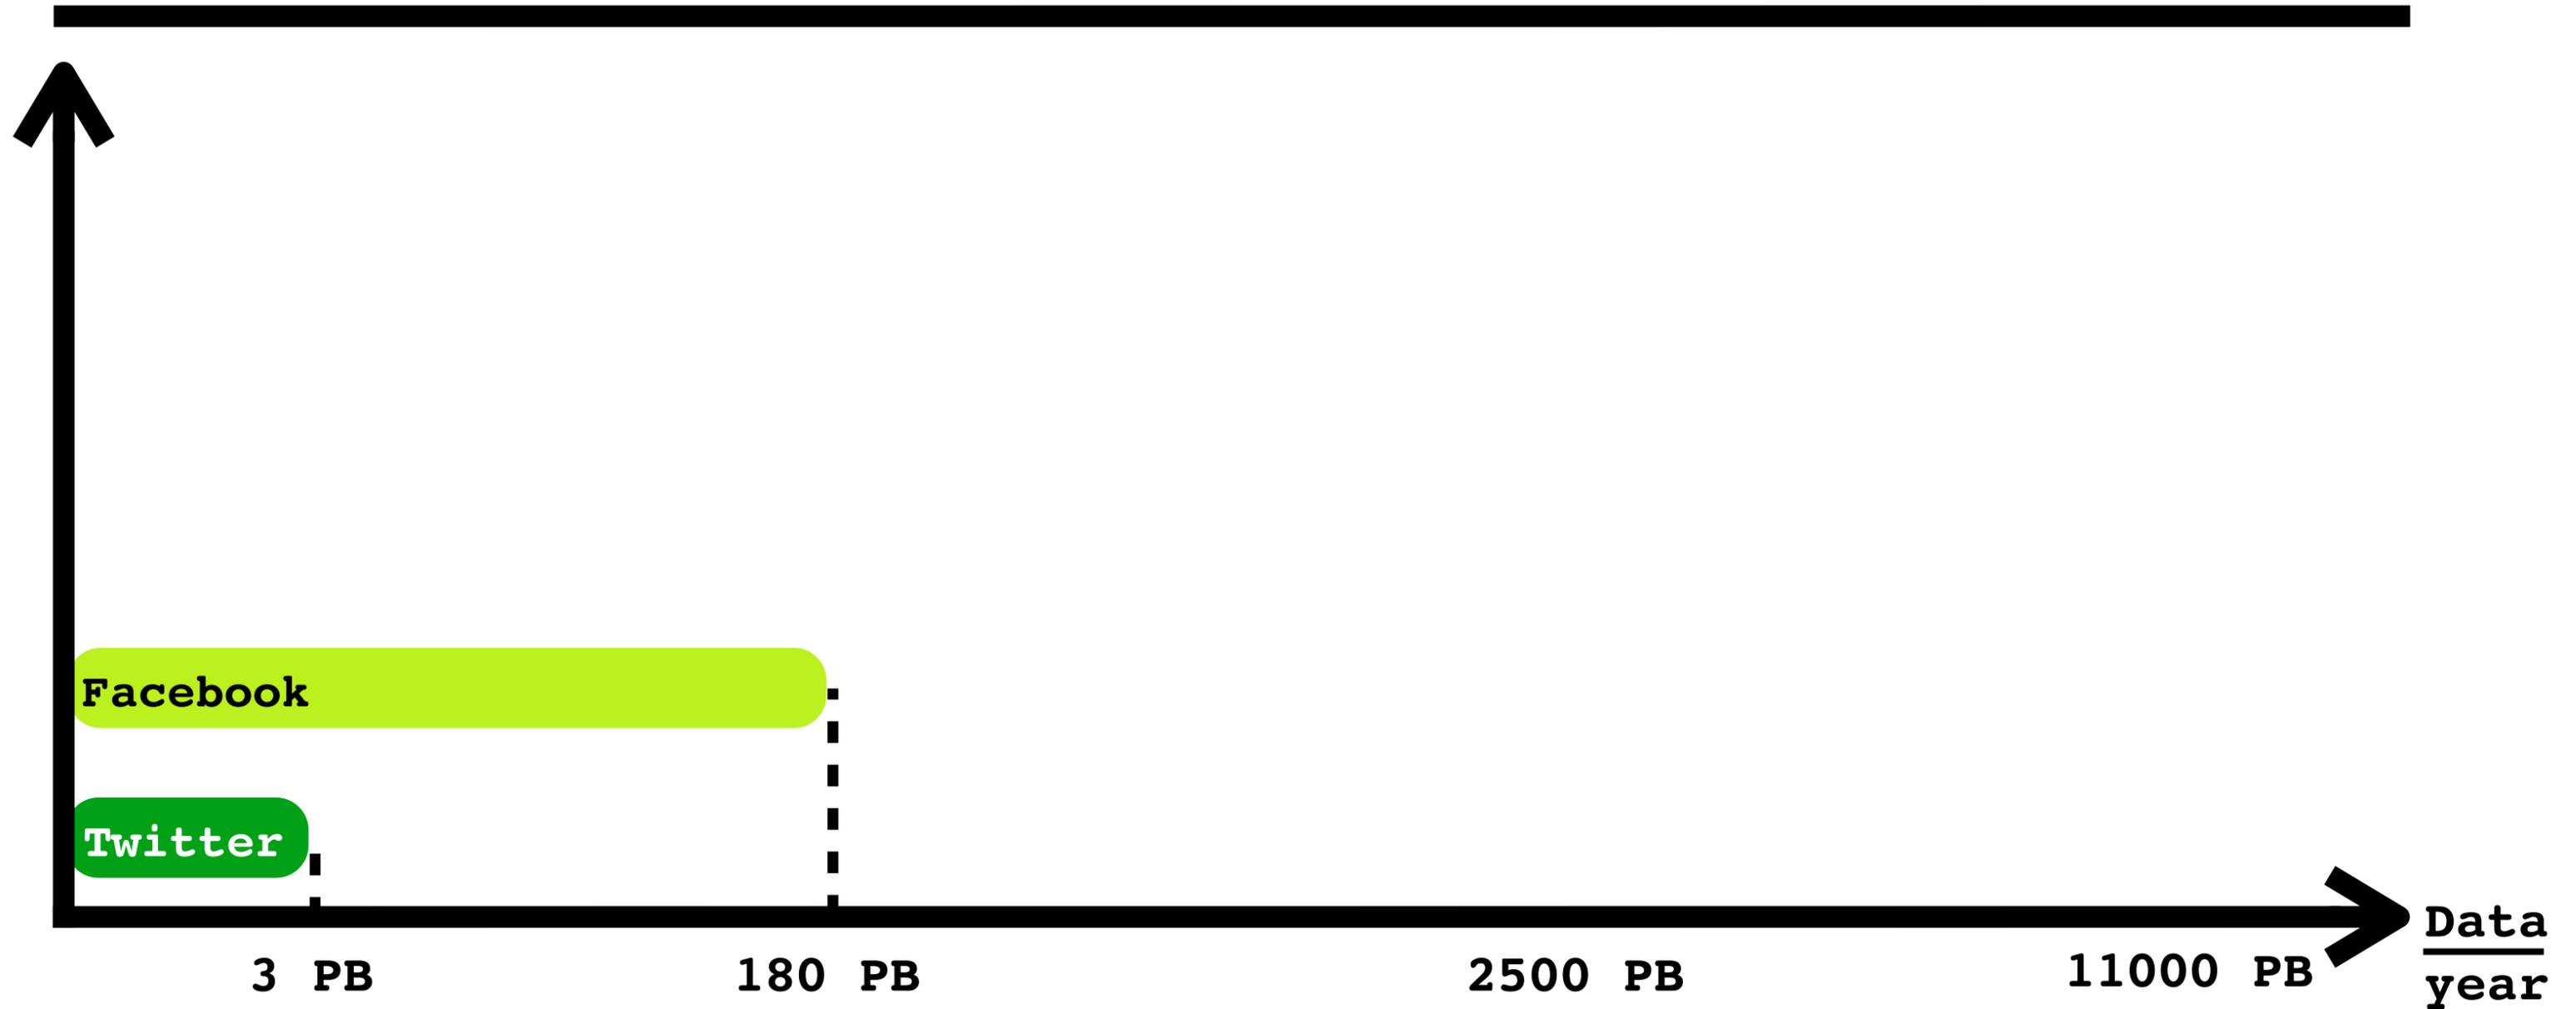


Why do we need triggers at the LHC?

Input data rate of one LHC experiment = 1.5 TB/second



This means about 15000 PB of data every year

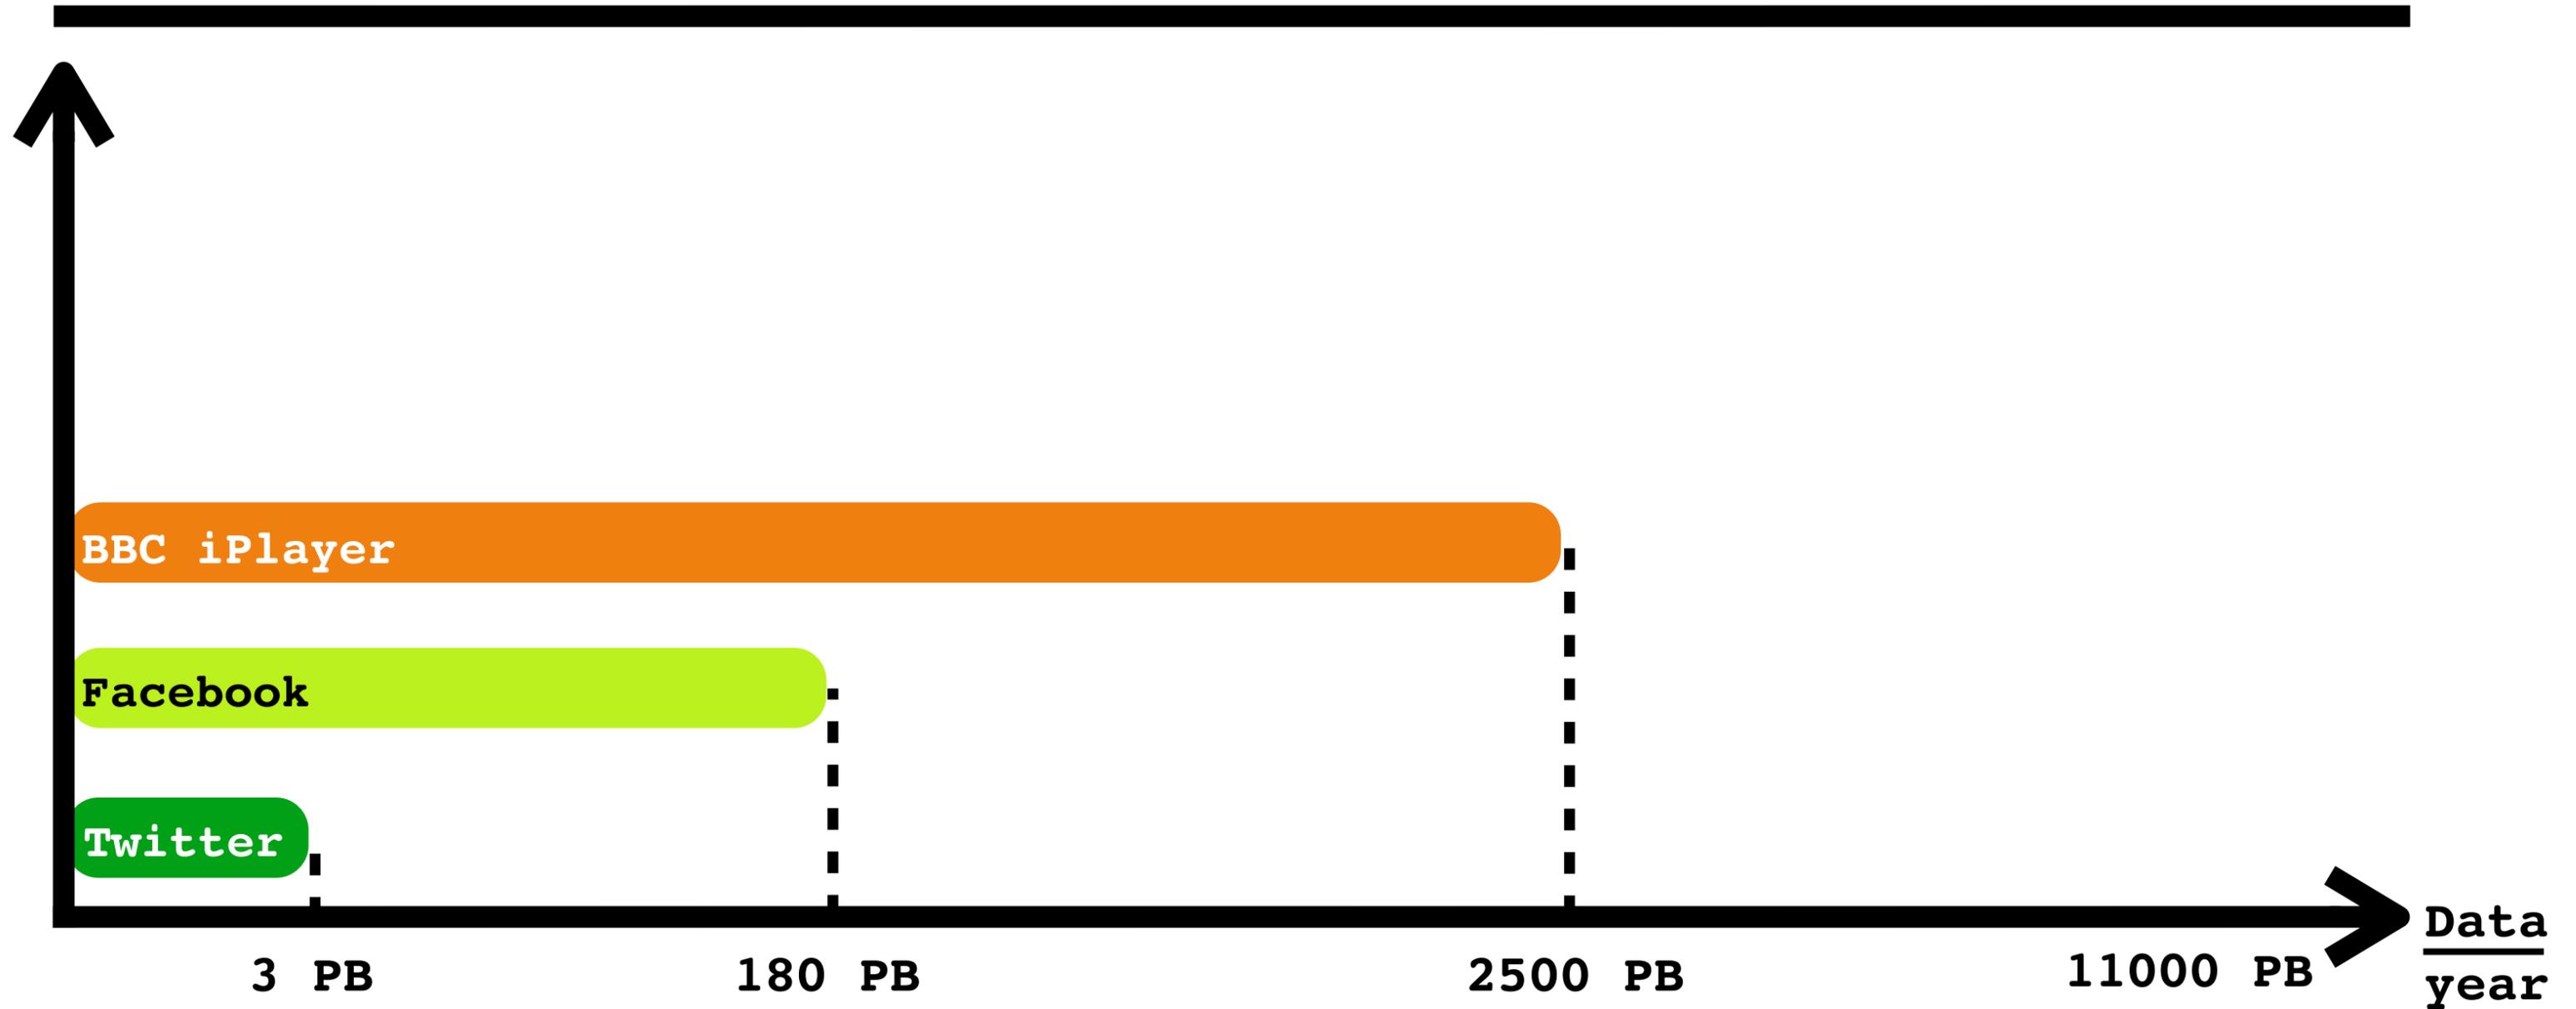


Why do we need triggers at the LHC?

Input data rate of one LHC experiment = 1.5 TB/second



This means about 15000 PB of data every year

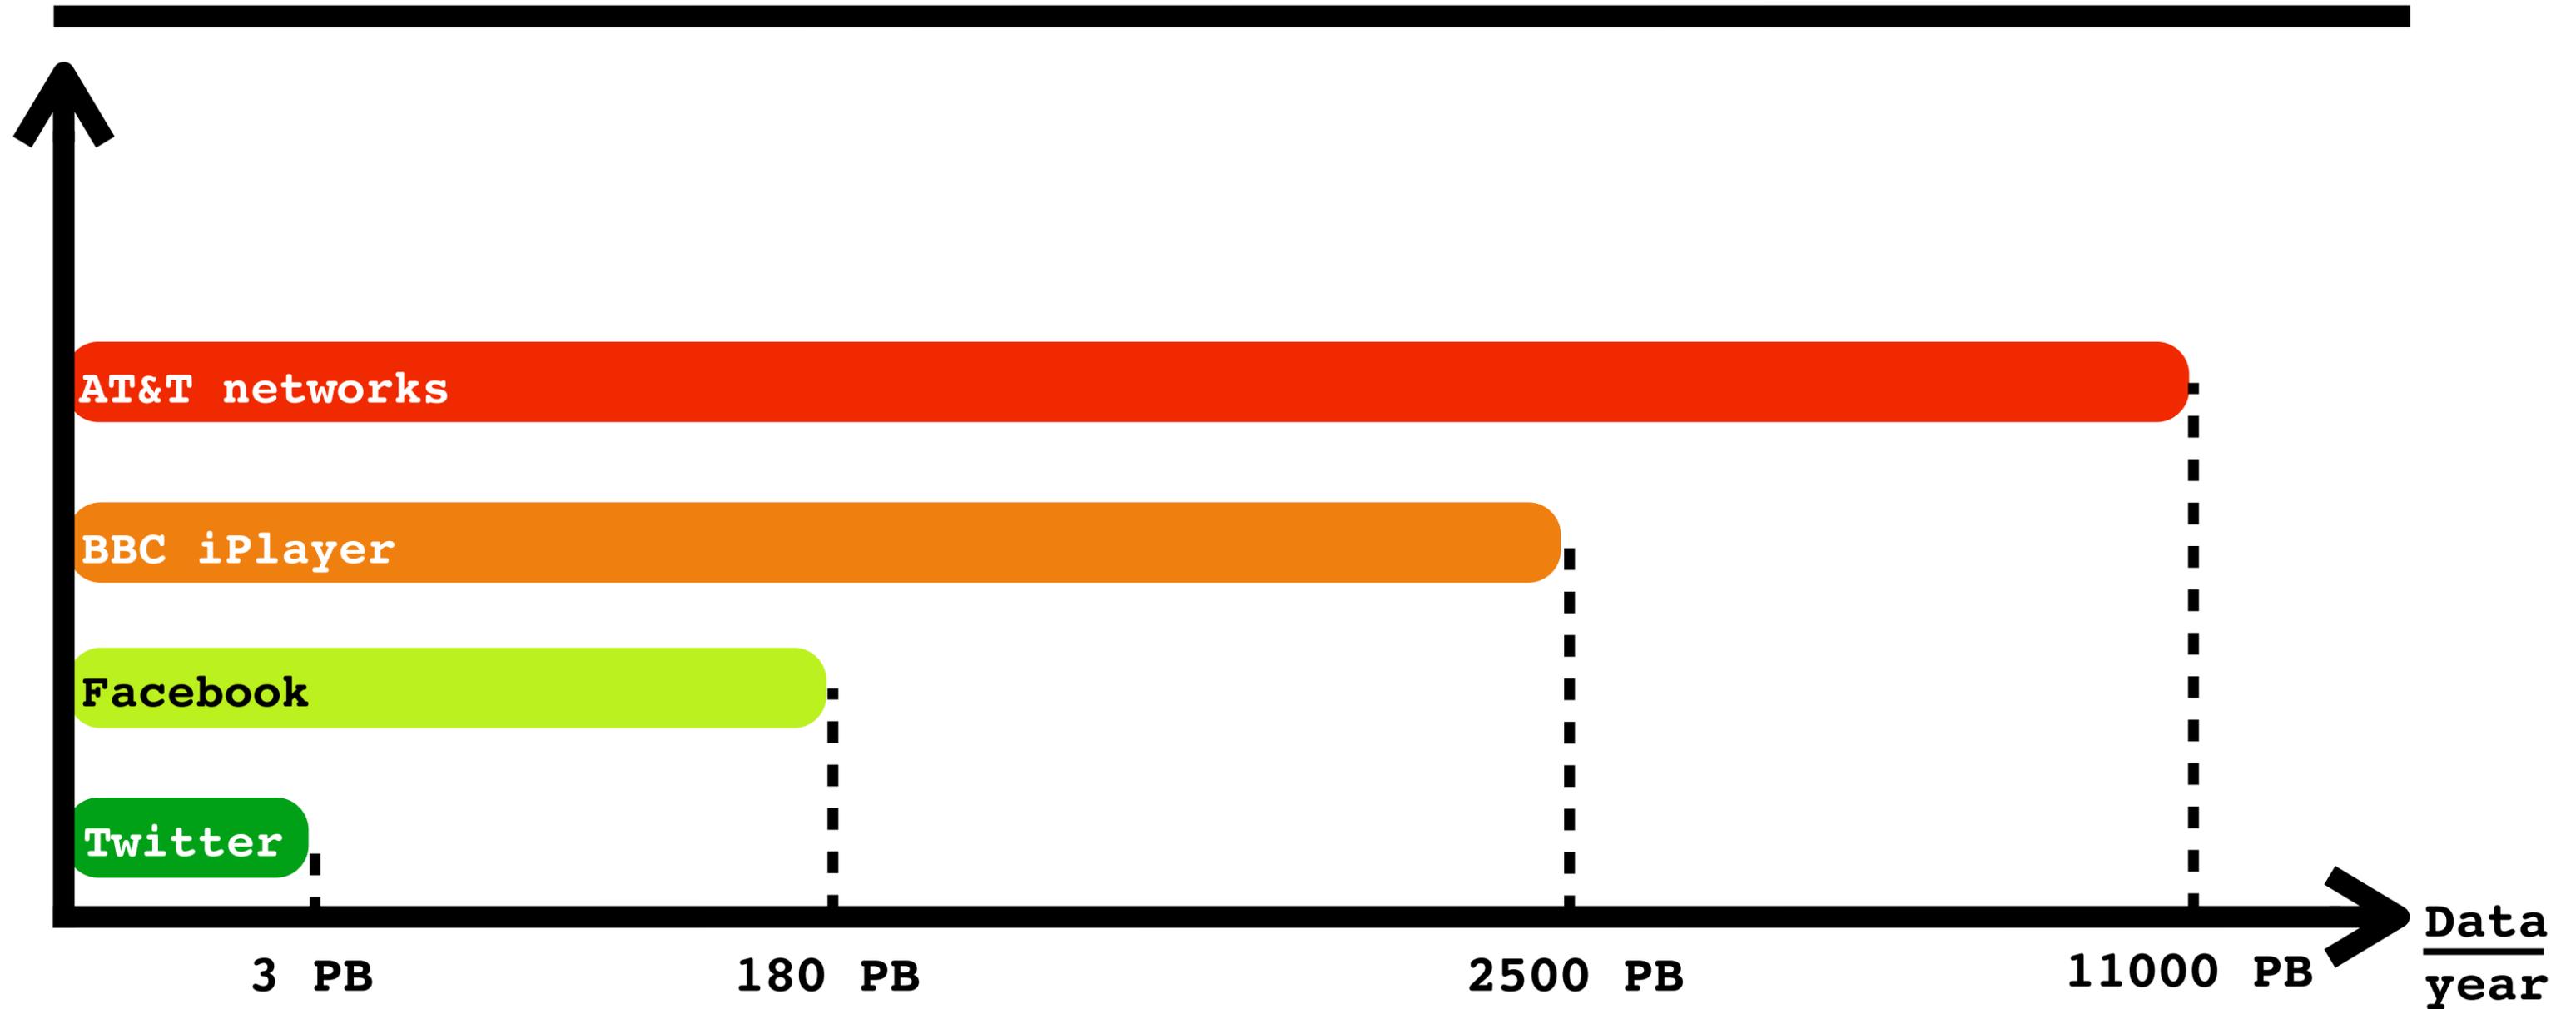


Why do we need triggers at the LHC?

Input data rate of one LHC experiment = 1.5 TB/second



This means about 15000 PB of data every year

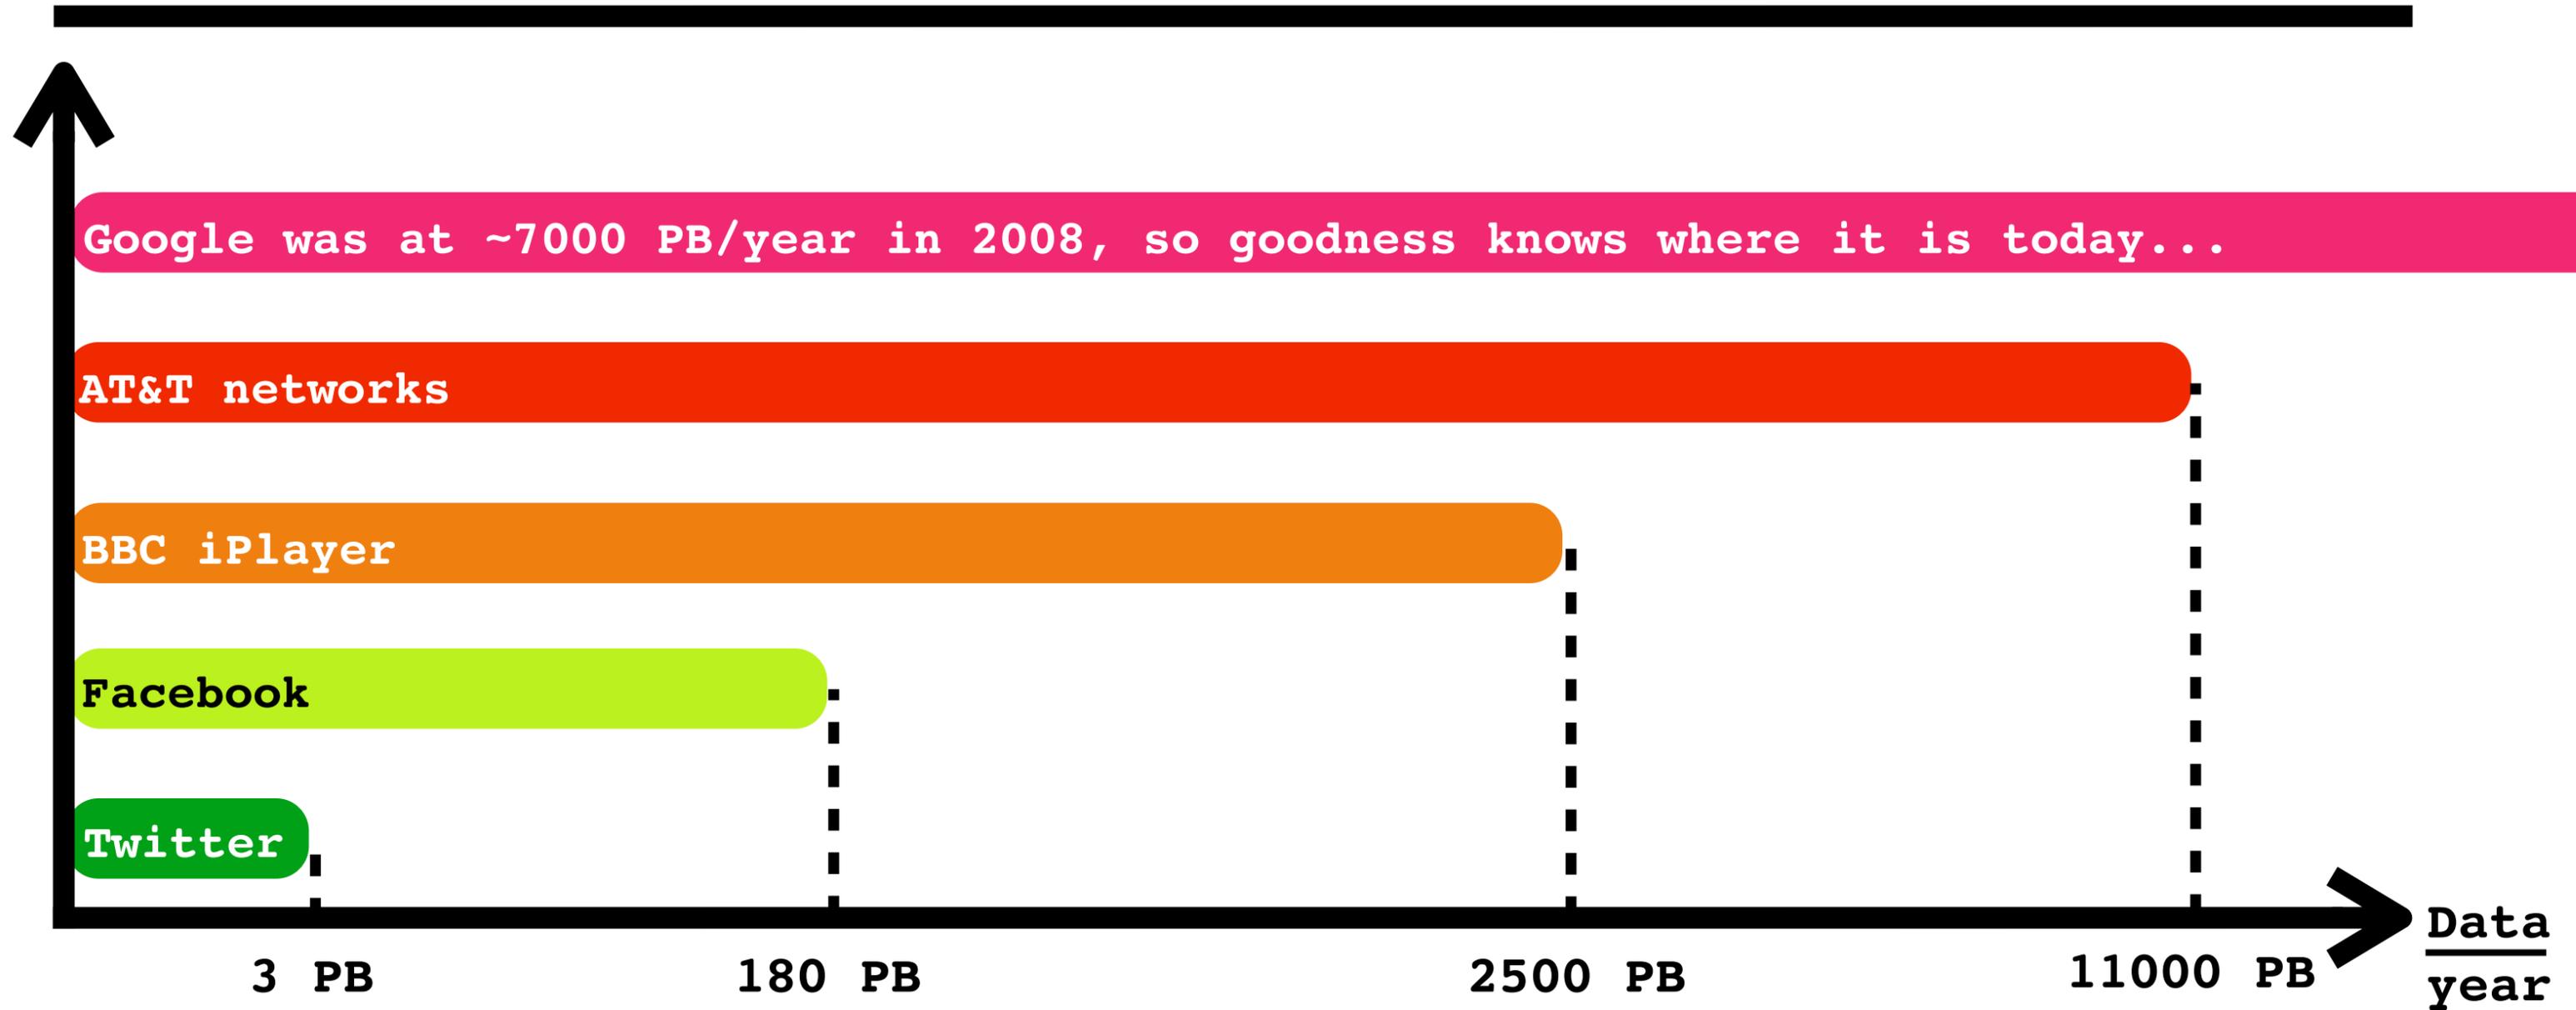


Why do we need triggers at the LHC?

Input data rate of one LHC experiment = 1.5 TB/second



This means about 15000 PB of data every year



It is mostly about the money

Facebook \approx 180 PB/year

It is mostly about the money

Facebook \approx 180 PB/year

Facebook spending on data
 \approx 600 M\$/year (circa 2011)

It is mostly about the money

Facebook \approx 180 PB/year

Facebook spending on data
 \approx 600 M\$/year (circa 2011)

Scaled cost of triggerless storage of one LHC experiment \approx 50 B\$/year

It is mostly about the money

Facebook \approx 180 PB/year

Facebook spending on data
 \approx 600 M\$/year (circa 2011)

Scaled cost of triggerless storage of one LHC experiment \approx 50 B\$/year

Total LHC budget \approx 10 B\$

It is mostly about the money

Facebook \approx 180 PB/year

Facebook spending on data
 \approx 600 M\$/year (circa 2011)

Scaled cost of triggerless storage of one LHC experiment \approx 50 B\$/year

Total LHC budget \approx 10 B\$

Real time event selection

==

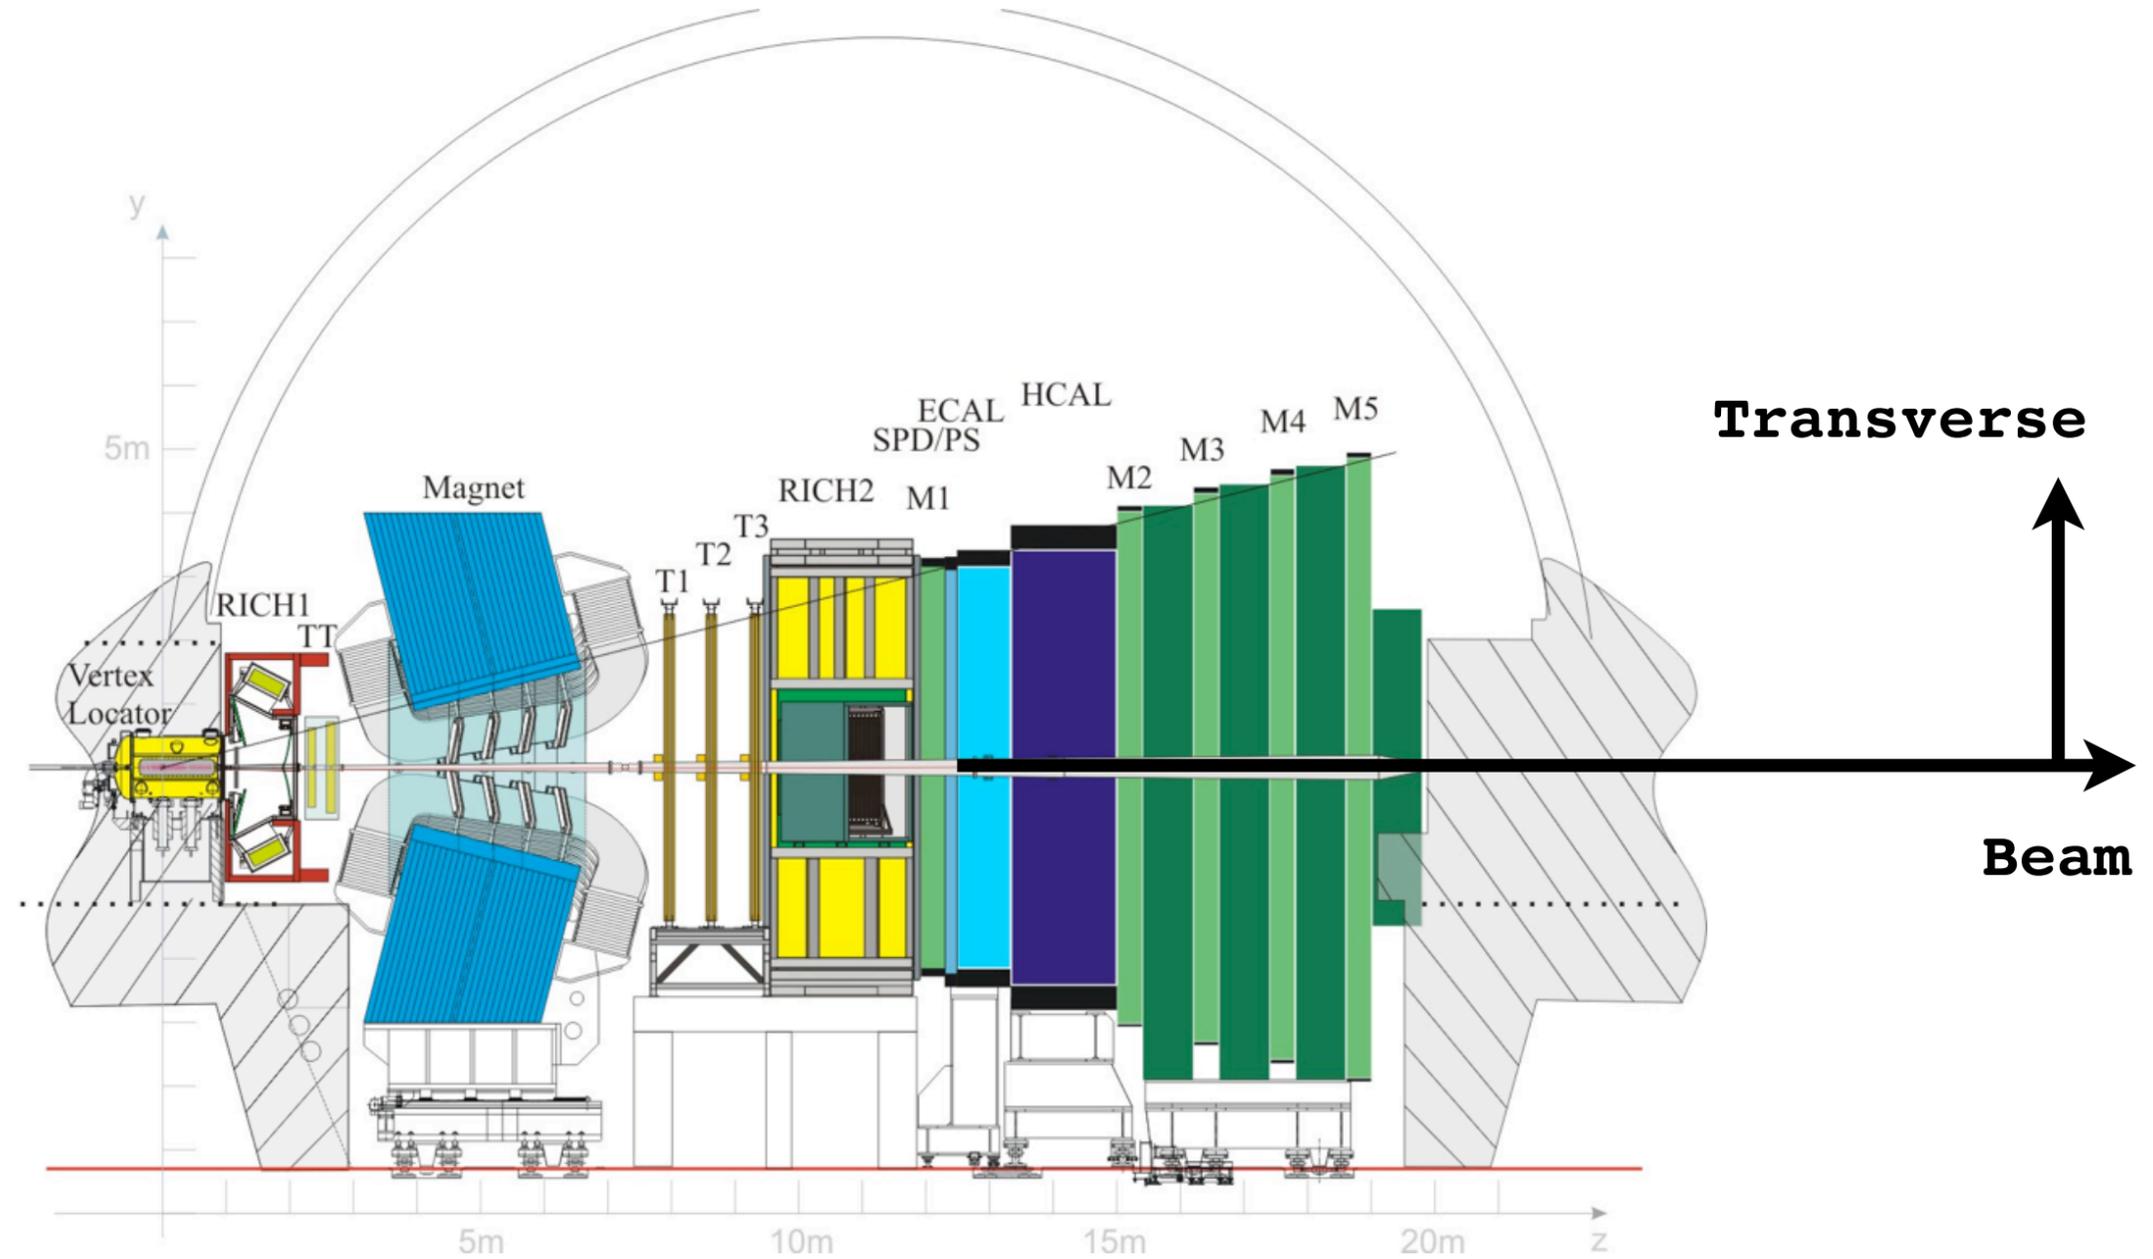
Money

Nota bene : Even assuming infinite resources, transferring 1.5 TB/second from the detectors to storage is not easy... although we are getting to the point where it is possible (more on this later)

Nota bene 2 : Processing data is way cheaper than storing it

The LHCb experiment

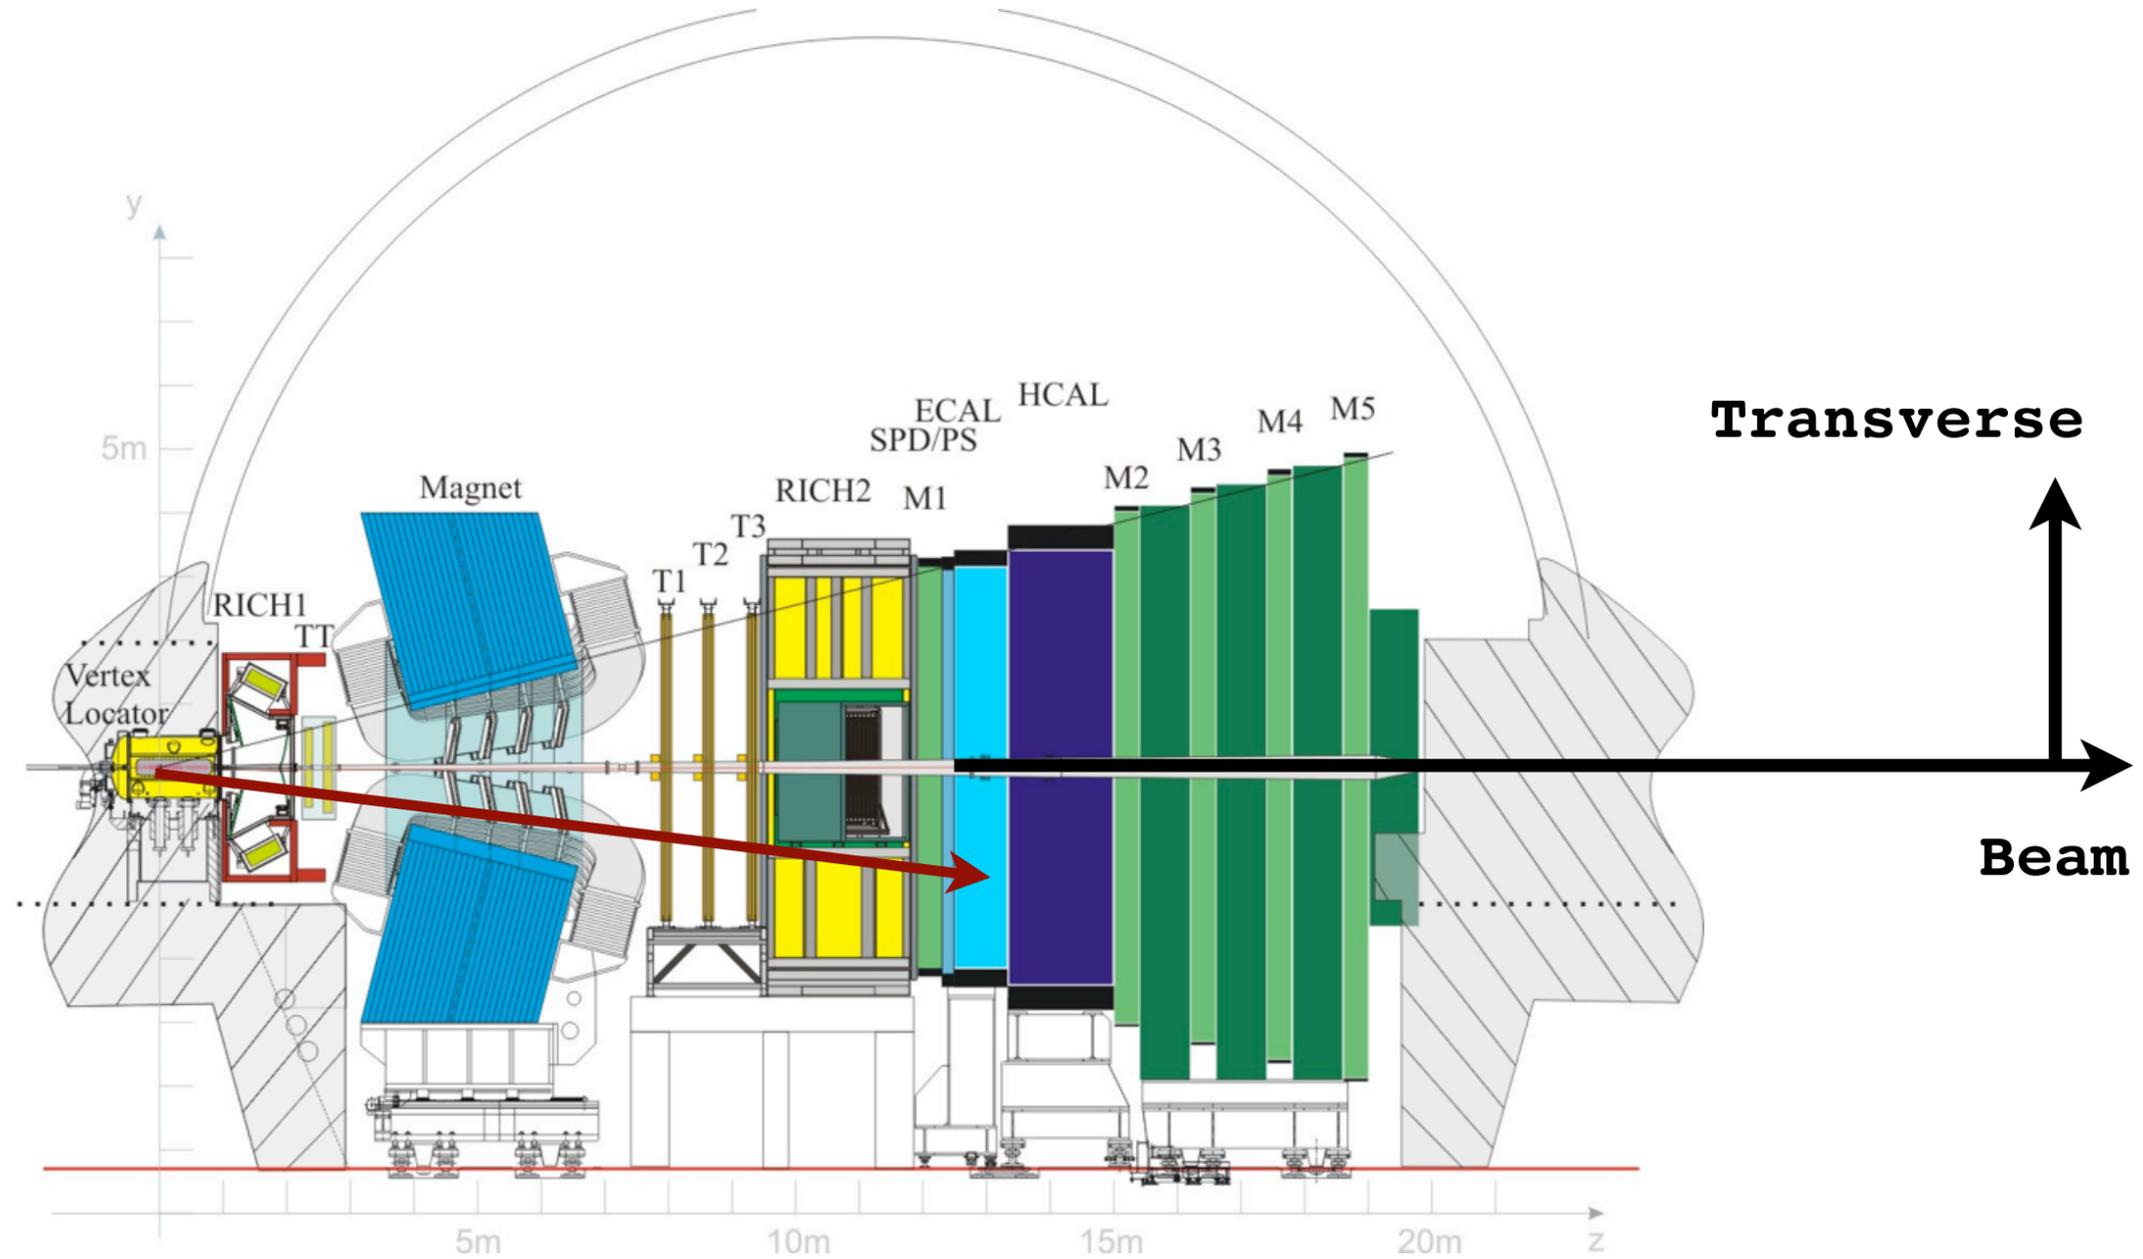
LHCb



p_T = Transverse momentum
 E_T = Transverse energy

LHCb

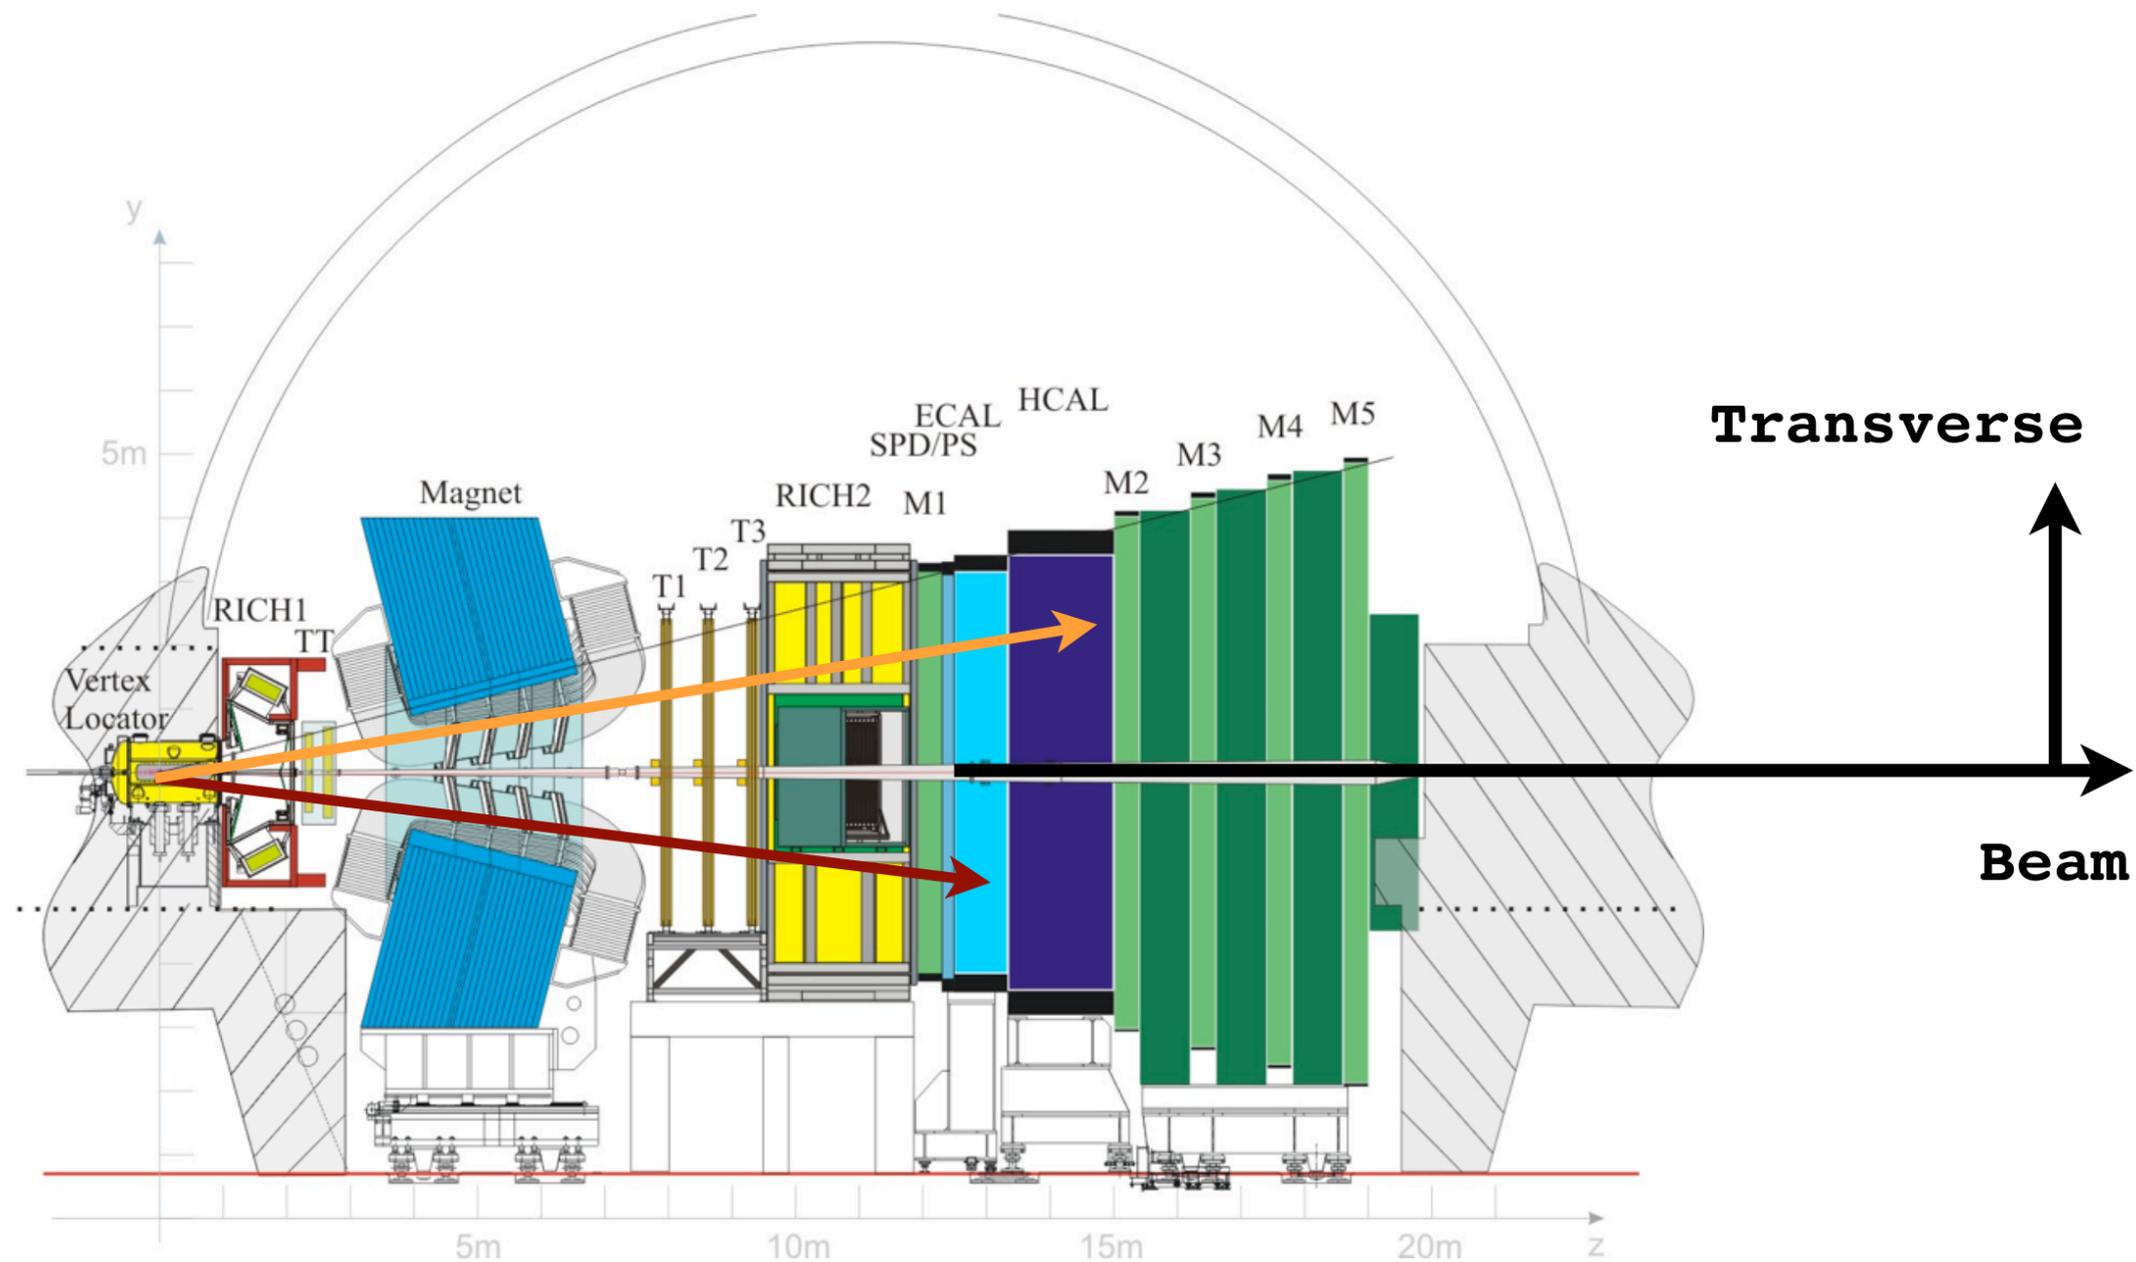
➔ **ELECTRONS**
➔ **PHOTONS**



p_T = Transverse momentum
 E_T = Transverse energy

LHCb

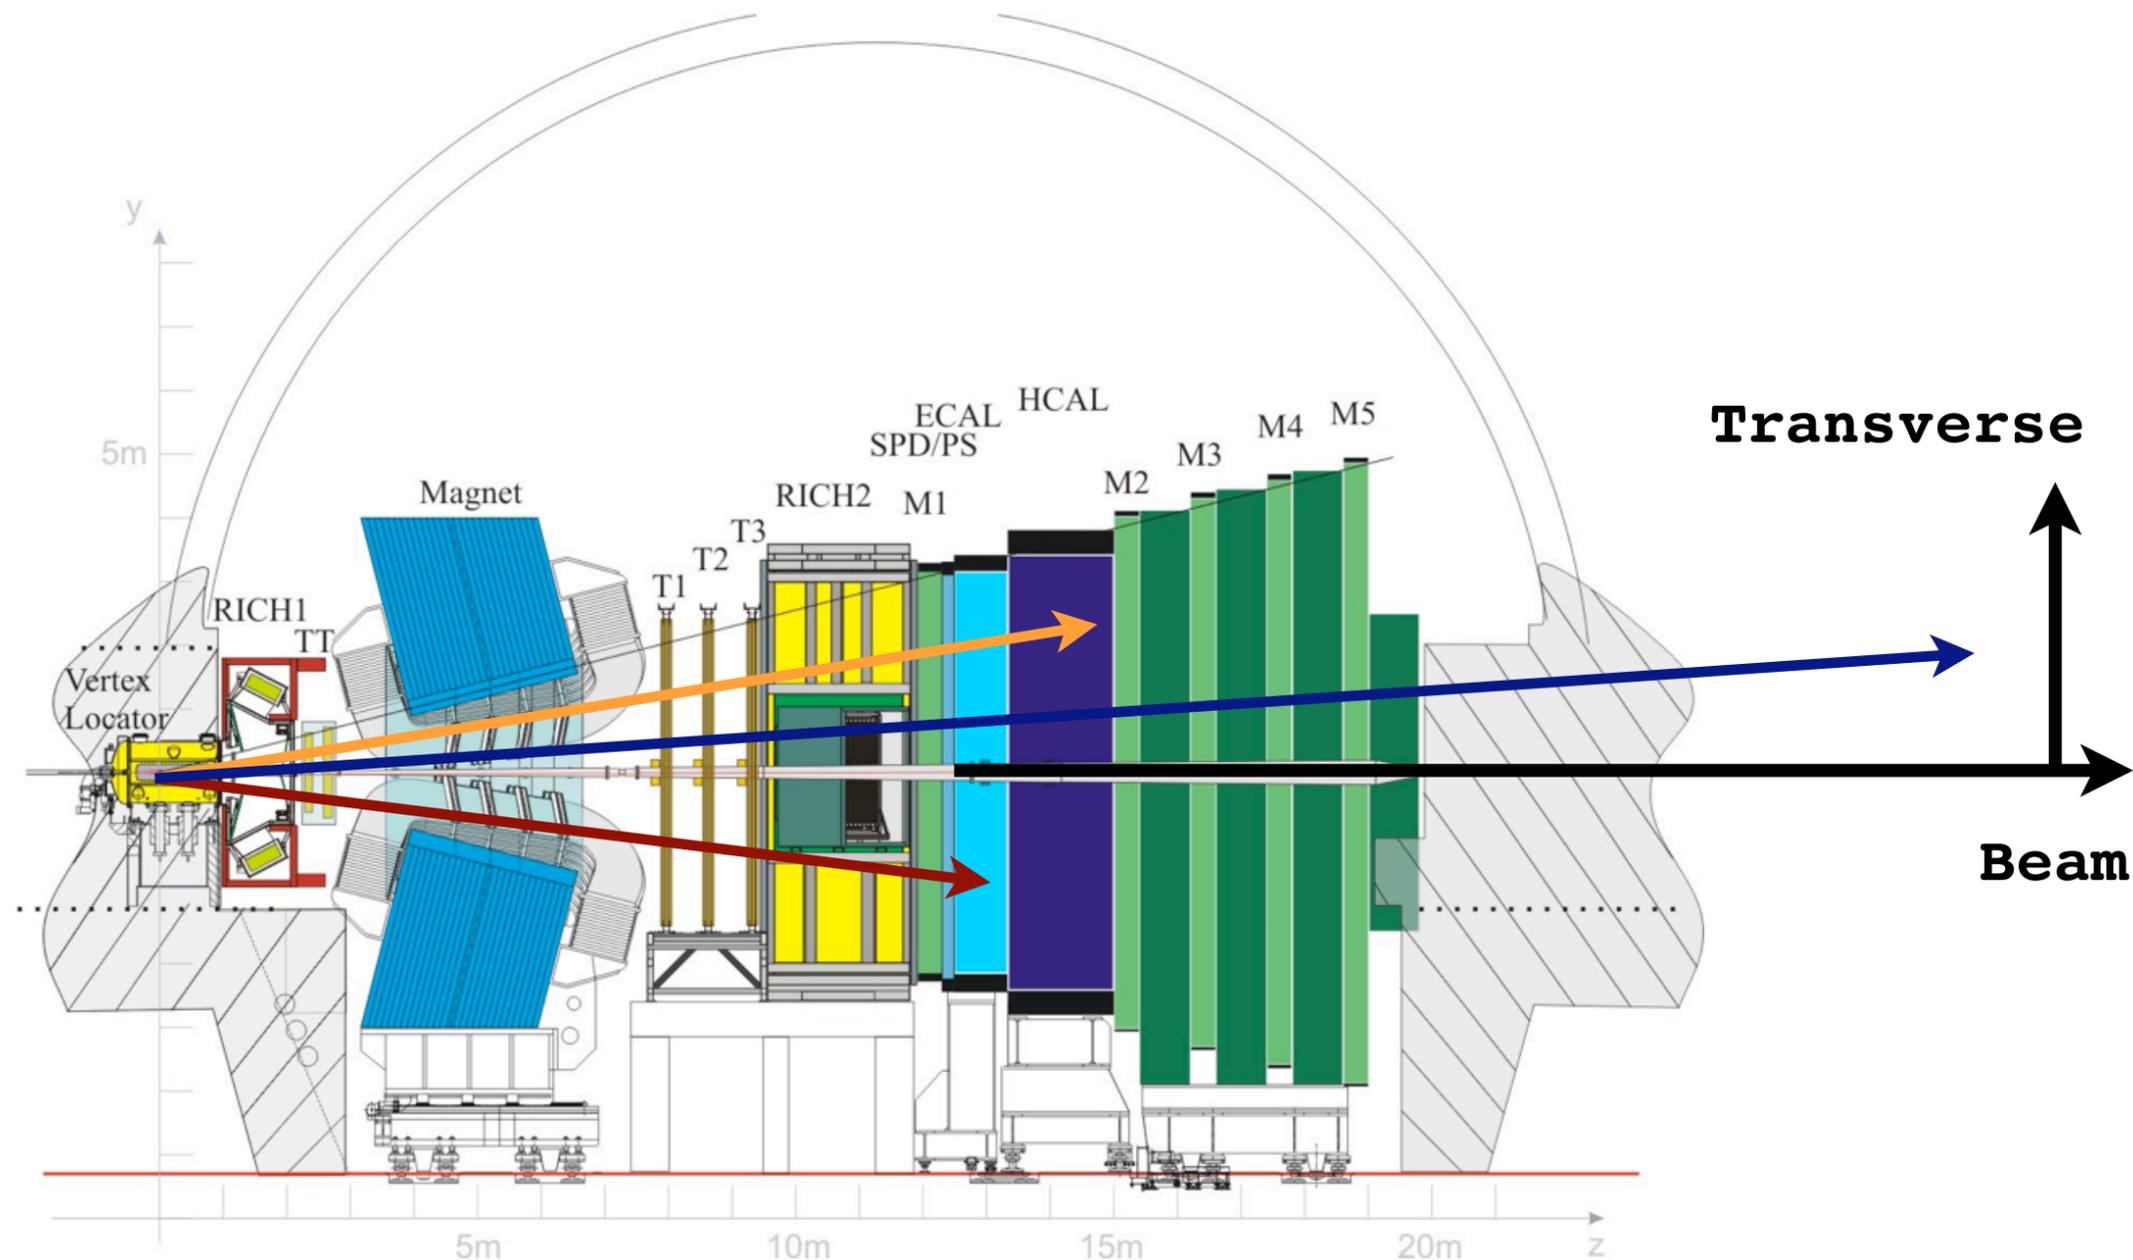
- ➔ **ELECTRONS**
- ➔ **PHOTONS**
- ➔ **HADRONS**



p_T = Transverse momentum
 E_T = Transverse energy

LHCb

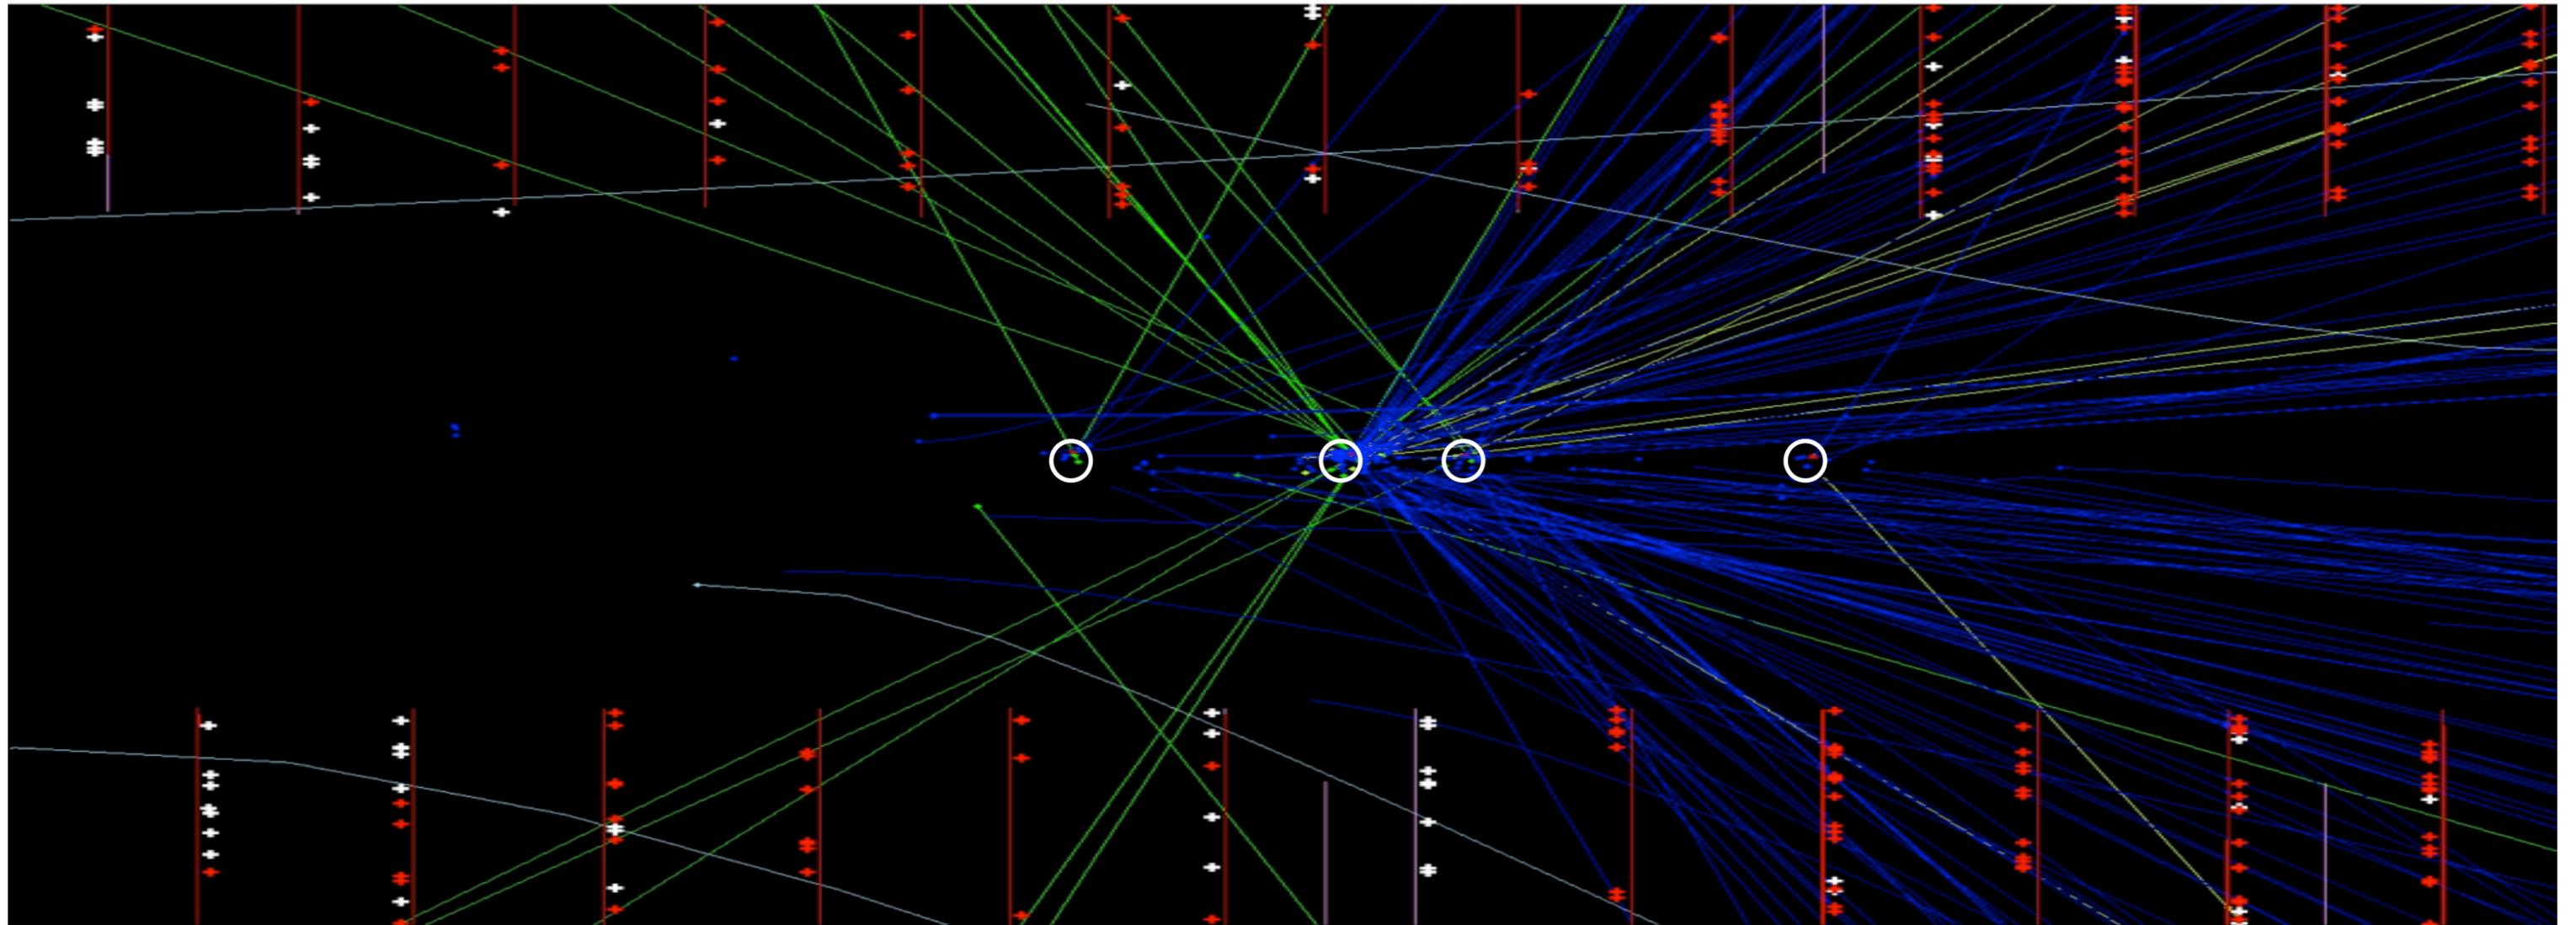
- ➔ ELECTRONS
- ➔ PHOTONS
- ➔ HADRONS
- ➔ MUONS



p_T = Transverse momentum
 E_T = Transverse energy

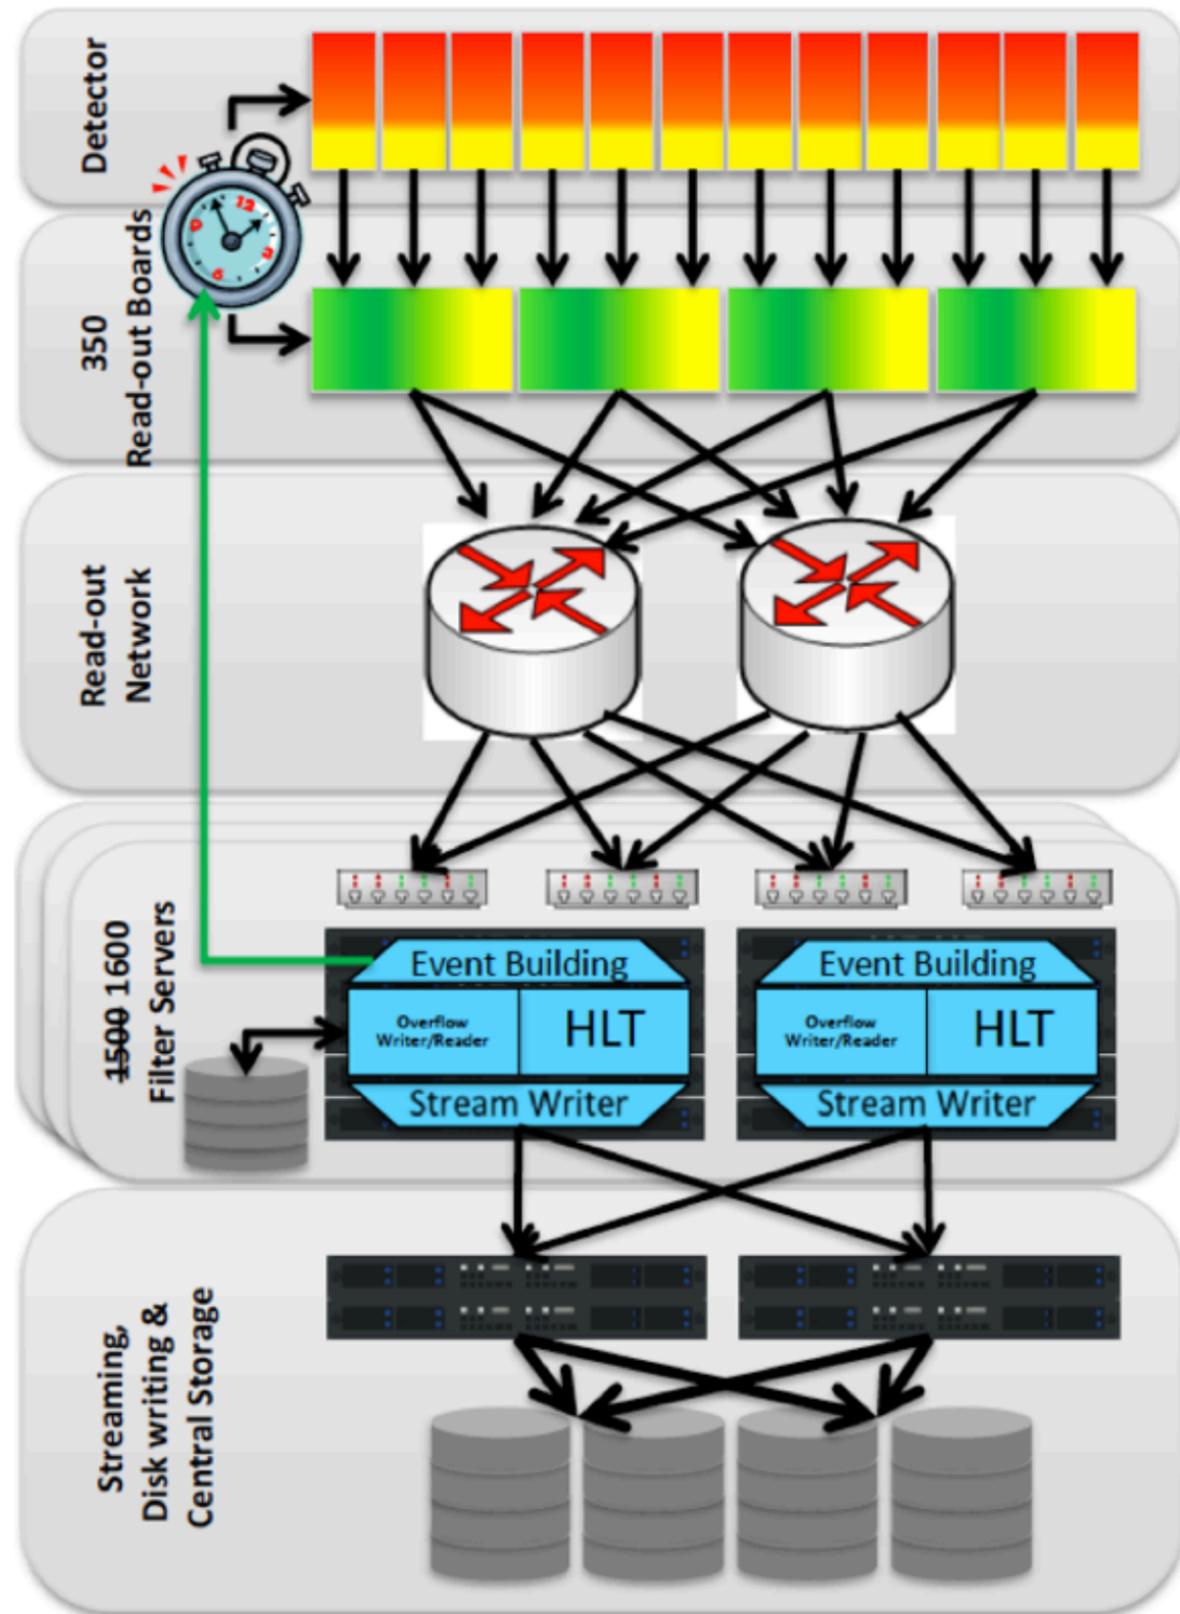
LHC environment

VELO rz view



In 2010–2012 : 15 MHz of bunch crossings, ~1.5 proton-proton interactions per bunch crossing, ~30 particles produced in the detector acceptance per interaction

LHCb trigger : hardware constraints



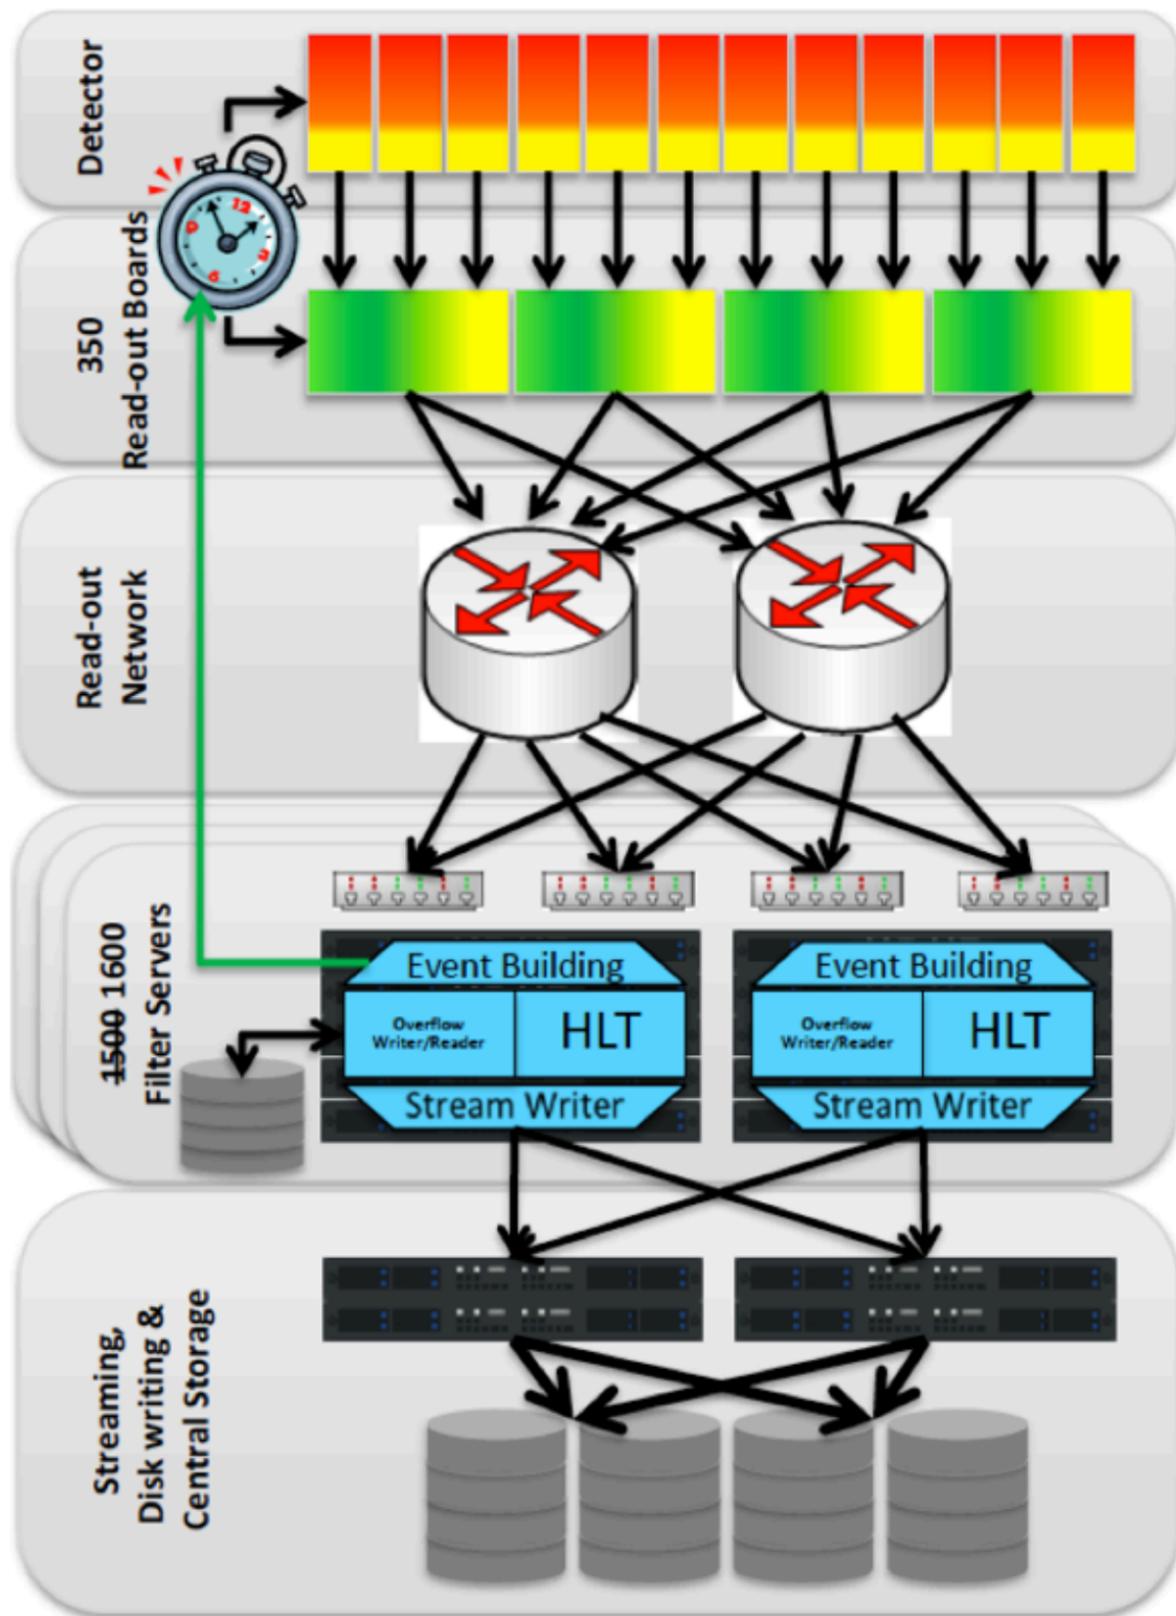
- DAQ

- Readout rate: 1MHz
- Total Event size: 50+kB
- HLT output rate: 5000Hz
- HLT output bandwidth: 250MB/s

- Architecture

- Dual core routers
- Data that can't be processed by the HLT is temporarily stored on local HLT node discs for inter-fill processing

LHCb trigger : hardware constraints



- DAQ

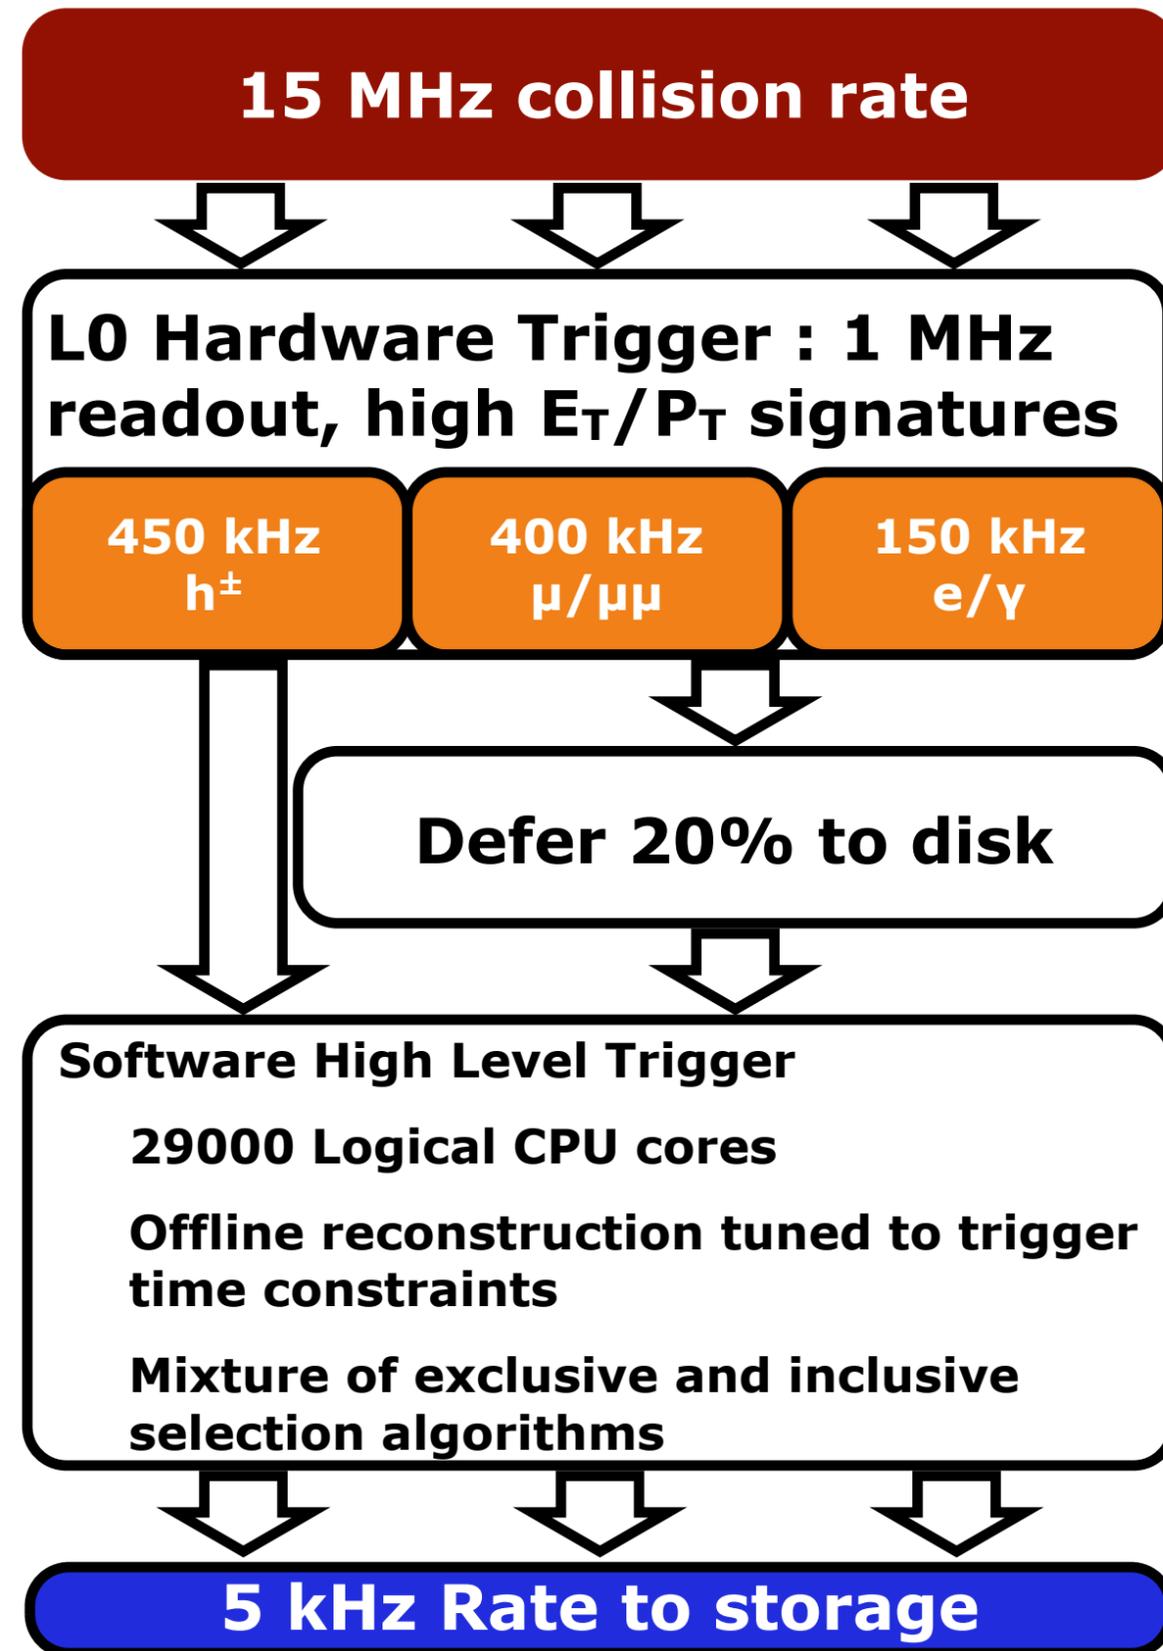
- Readout rate: 1MHz
- Total Event size: 50+kB
- HLT output rate: 5000Hz
- HLT output bandwidth: 250MB/s

Because we cannot read the full LHCb detector at 40 MHz, we need a hardware trigger

Implemented in custom electronics with 4 μ s latency

Triggers on calorimeter clusters and muon station segments

LHCb trigger scheme



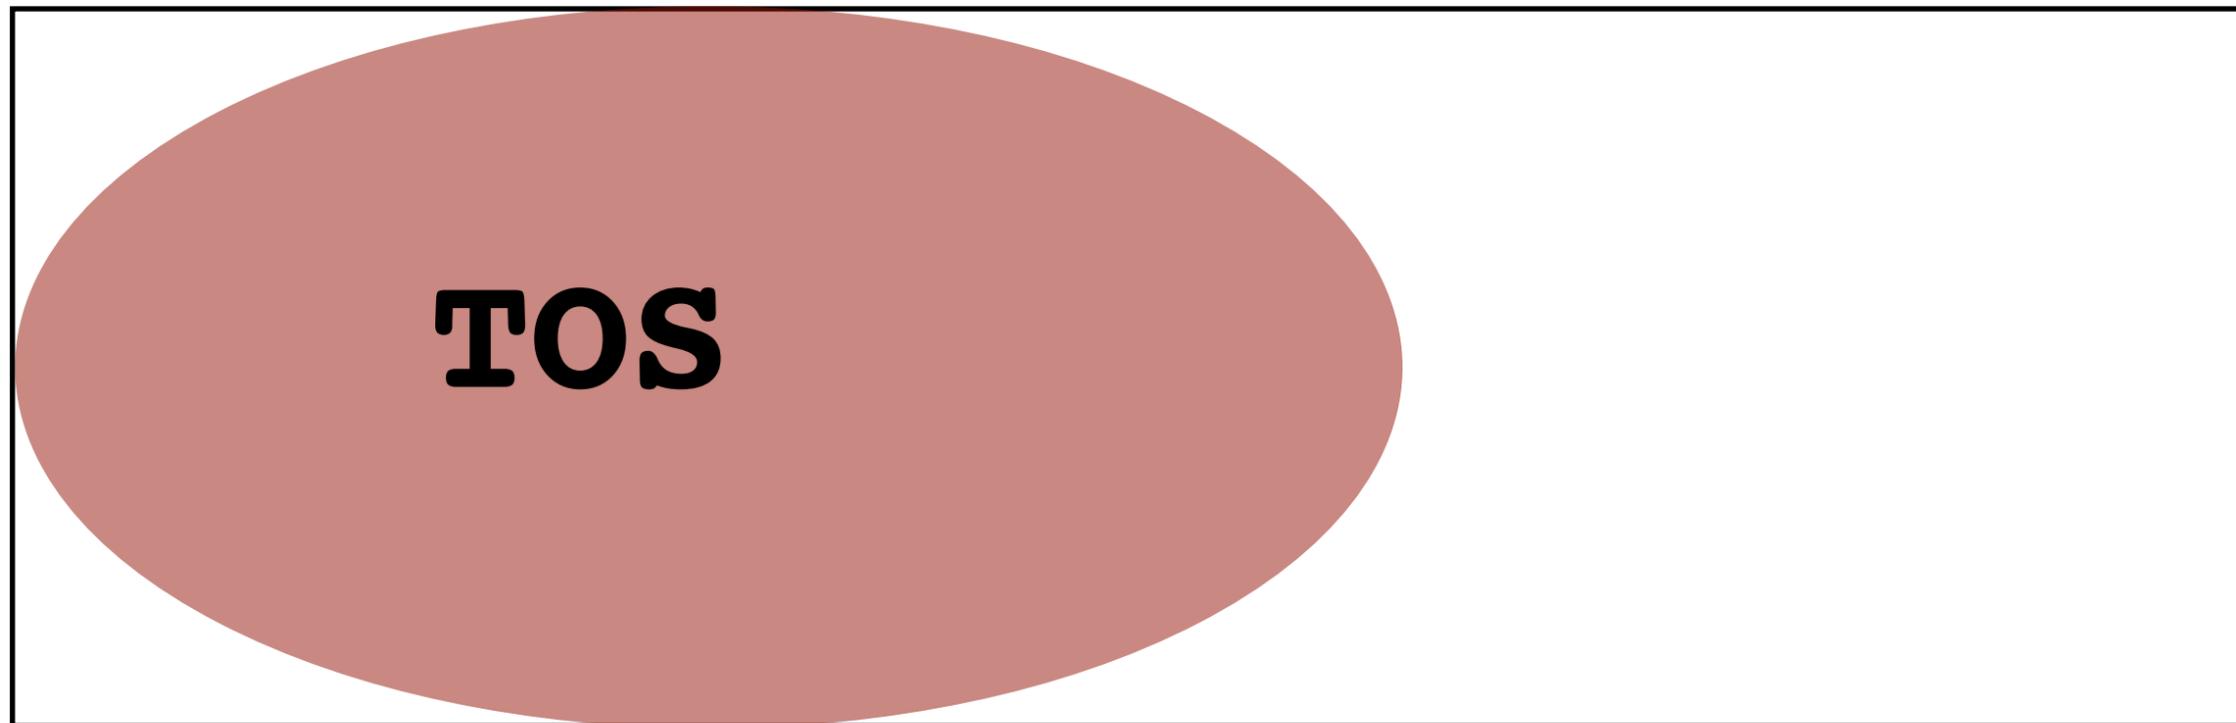
Classifying triggered
events : an interlude

TIS-TOS-TOB



All triggered events can be split into three categories

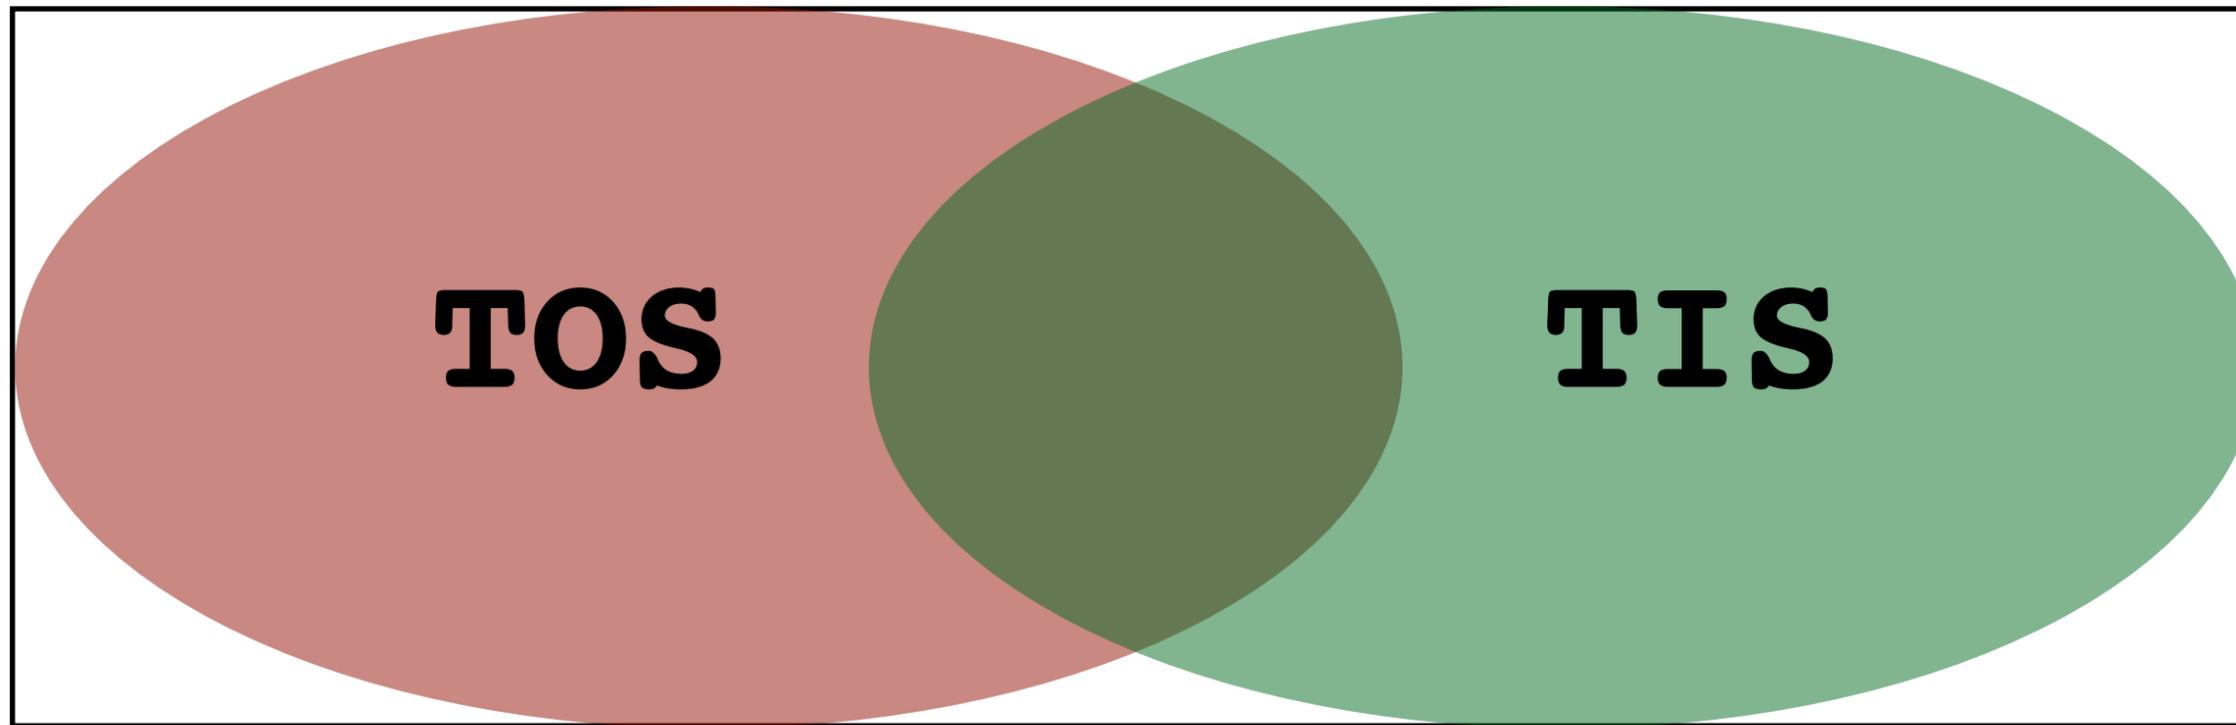
TIS-TOS-TOB



All triggered events can be split into three categories

TOS : The event would have triggered if only the particles belonging to the signal candidate were present.

TIS-TOS-TOB

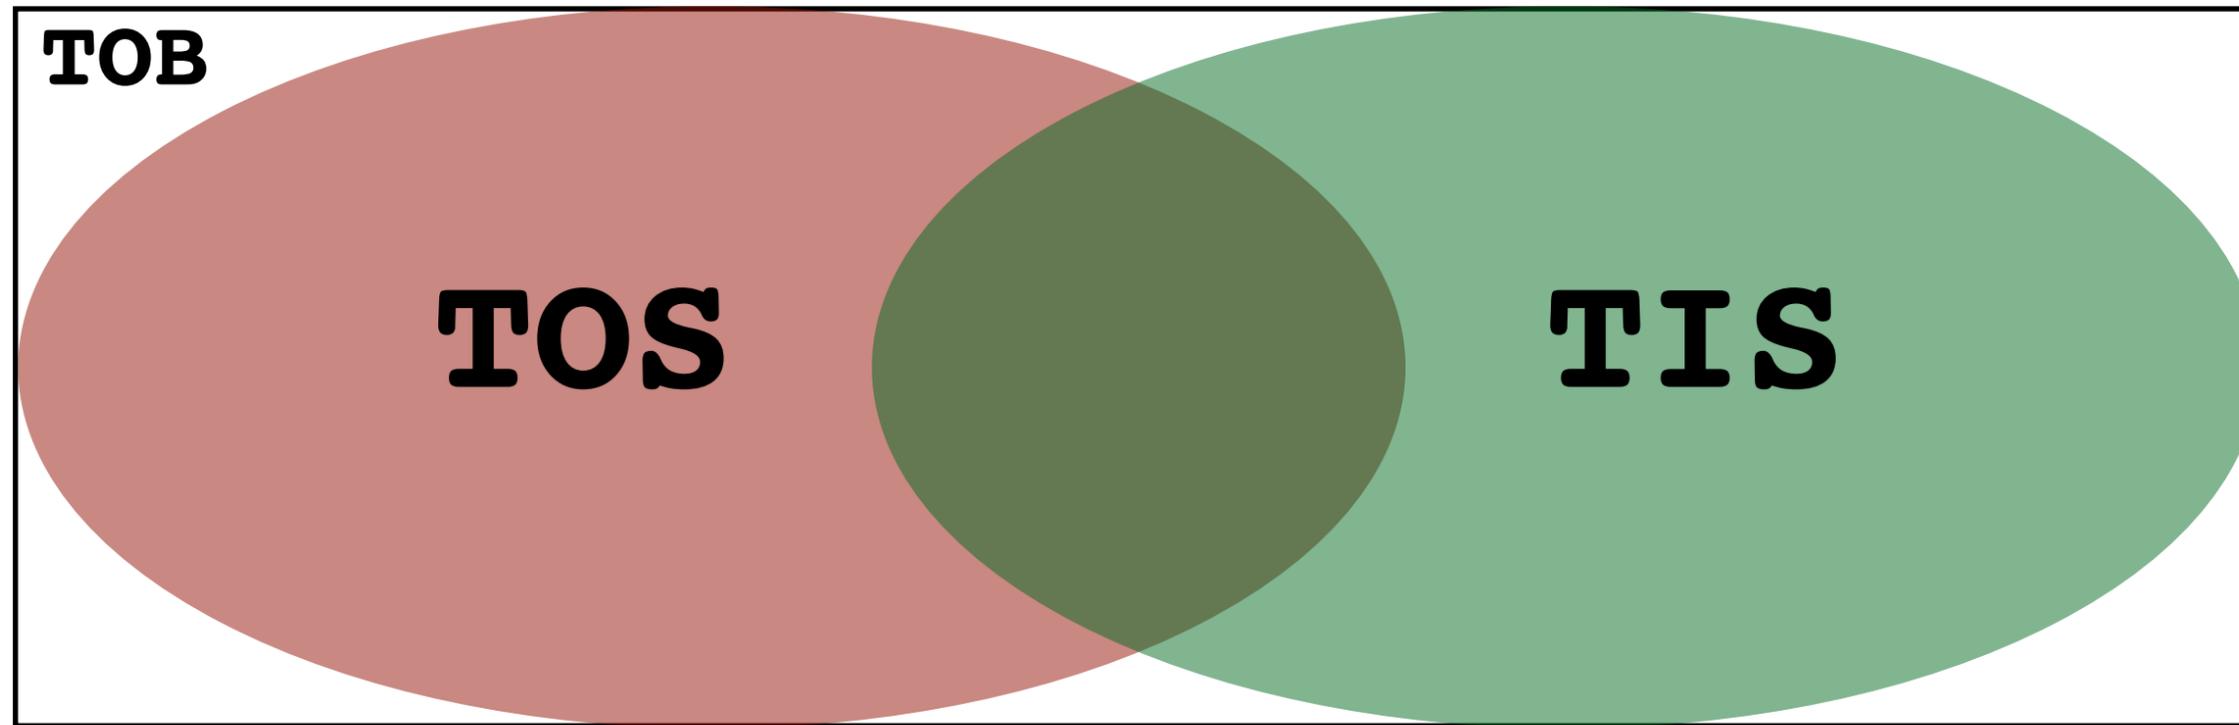


All triggered events can be split into three categories

TOS : The event would have triggered if only the particles belonging to the signal candidate were present.

TIS : The event would have triggered if the particles belonging to the signal candidate were not present.

TIS-TOS-TOB



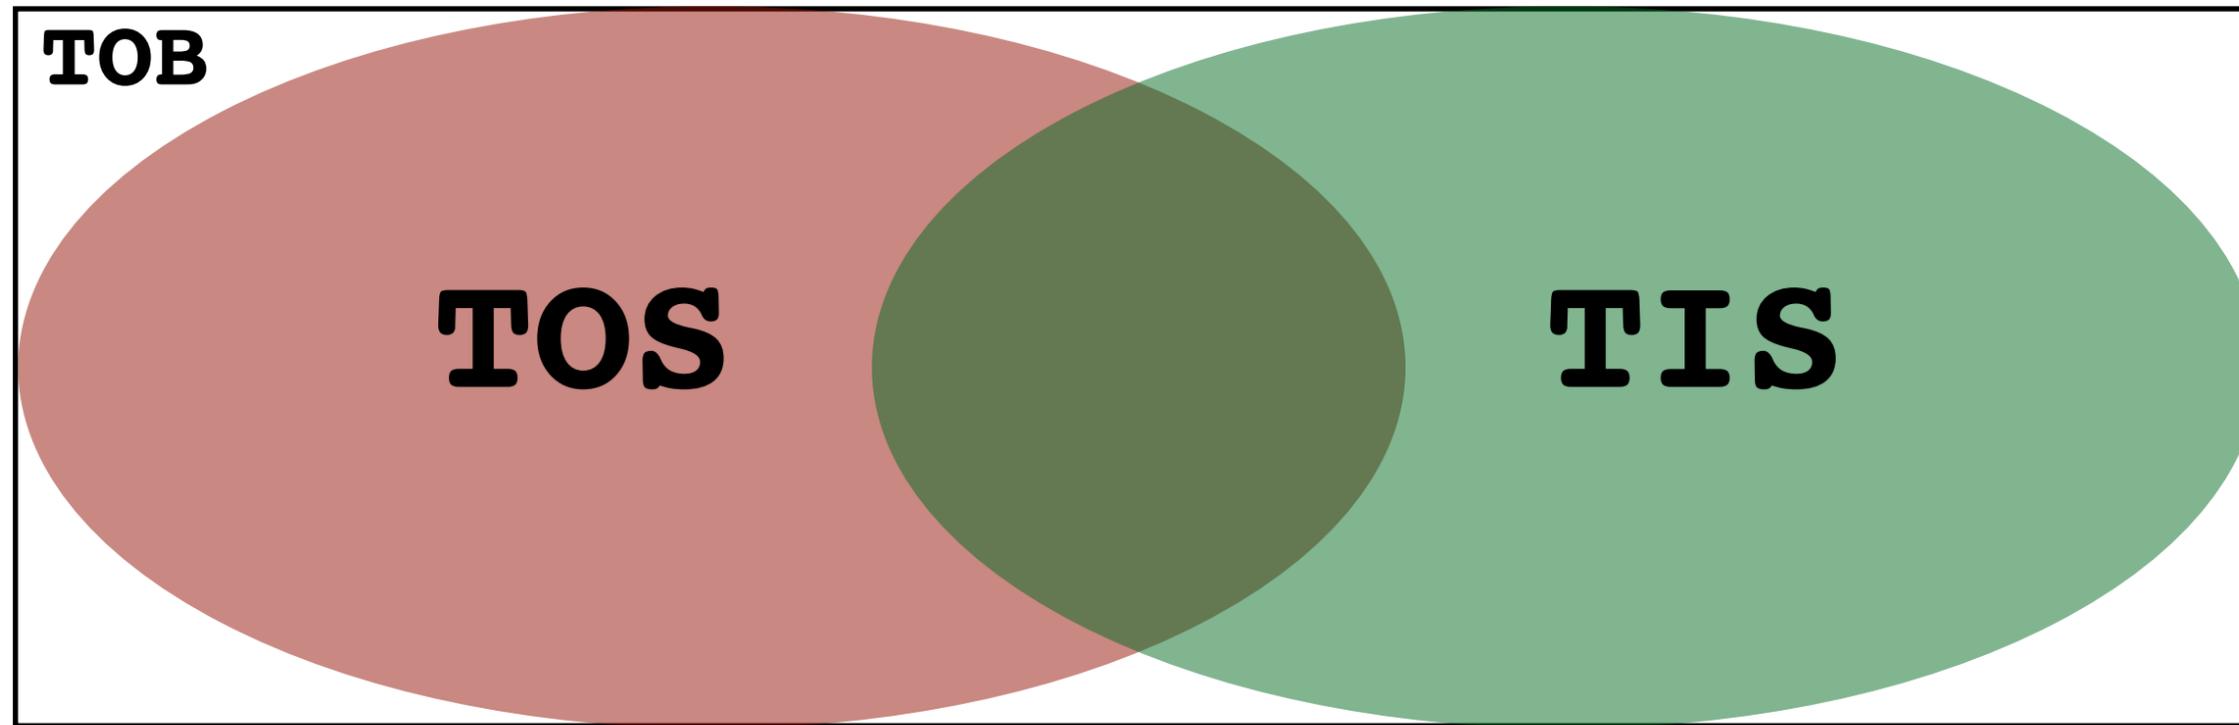
All triggered events can be split into three categories

TOS : The event would have triggered if only the particles belonging to the signal candidate were present.

TIS : The event would have triggered if the particles belonging to the signal candidate were not present.

TOB : neither TIS nor TOS

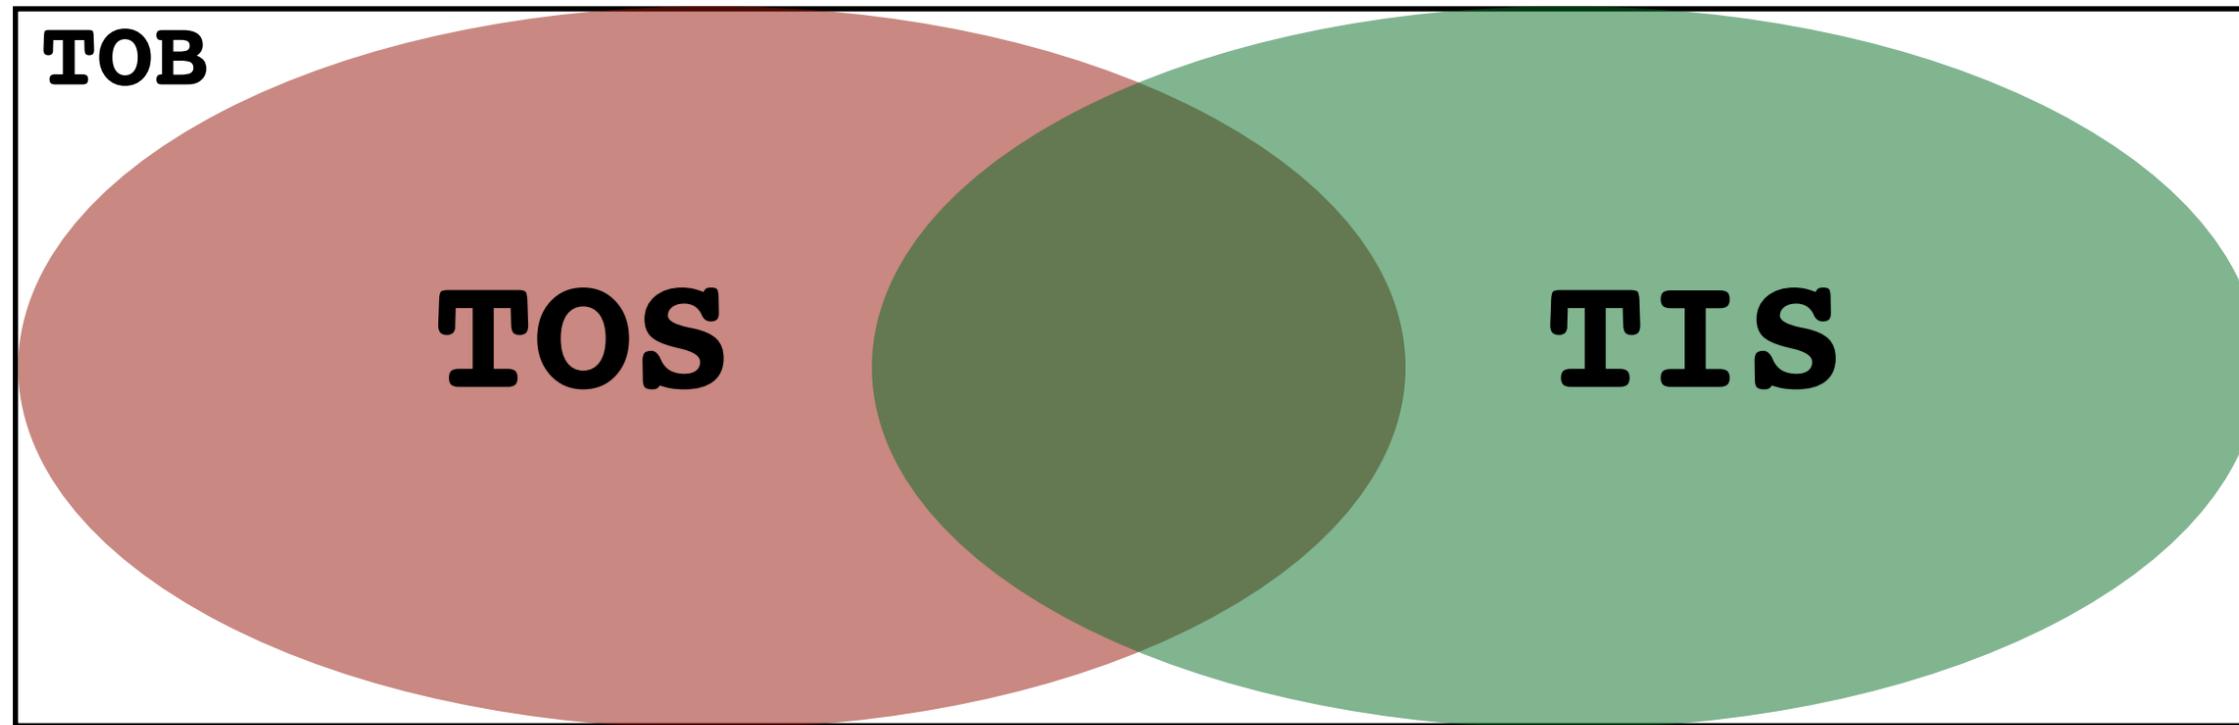
TIS-TOS-TOB



The overlap between TIS and TOS can be used to measure the trigger efficiency on data.

- 1) Select your signal events offline
- 2) Measure the fraction of TIS events which are also TOS of the trigger line which you are interested in
- 3) This gives the TOS efficiency of that line relative to the offline selection. The TIS efficiency can be similarly measured (fraction of TOS which are also TIS)

TIS-TOS-TOB

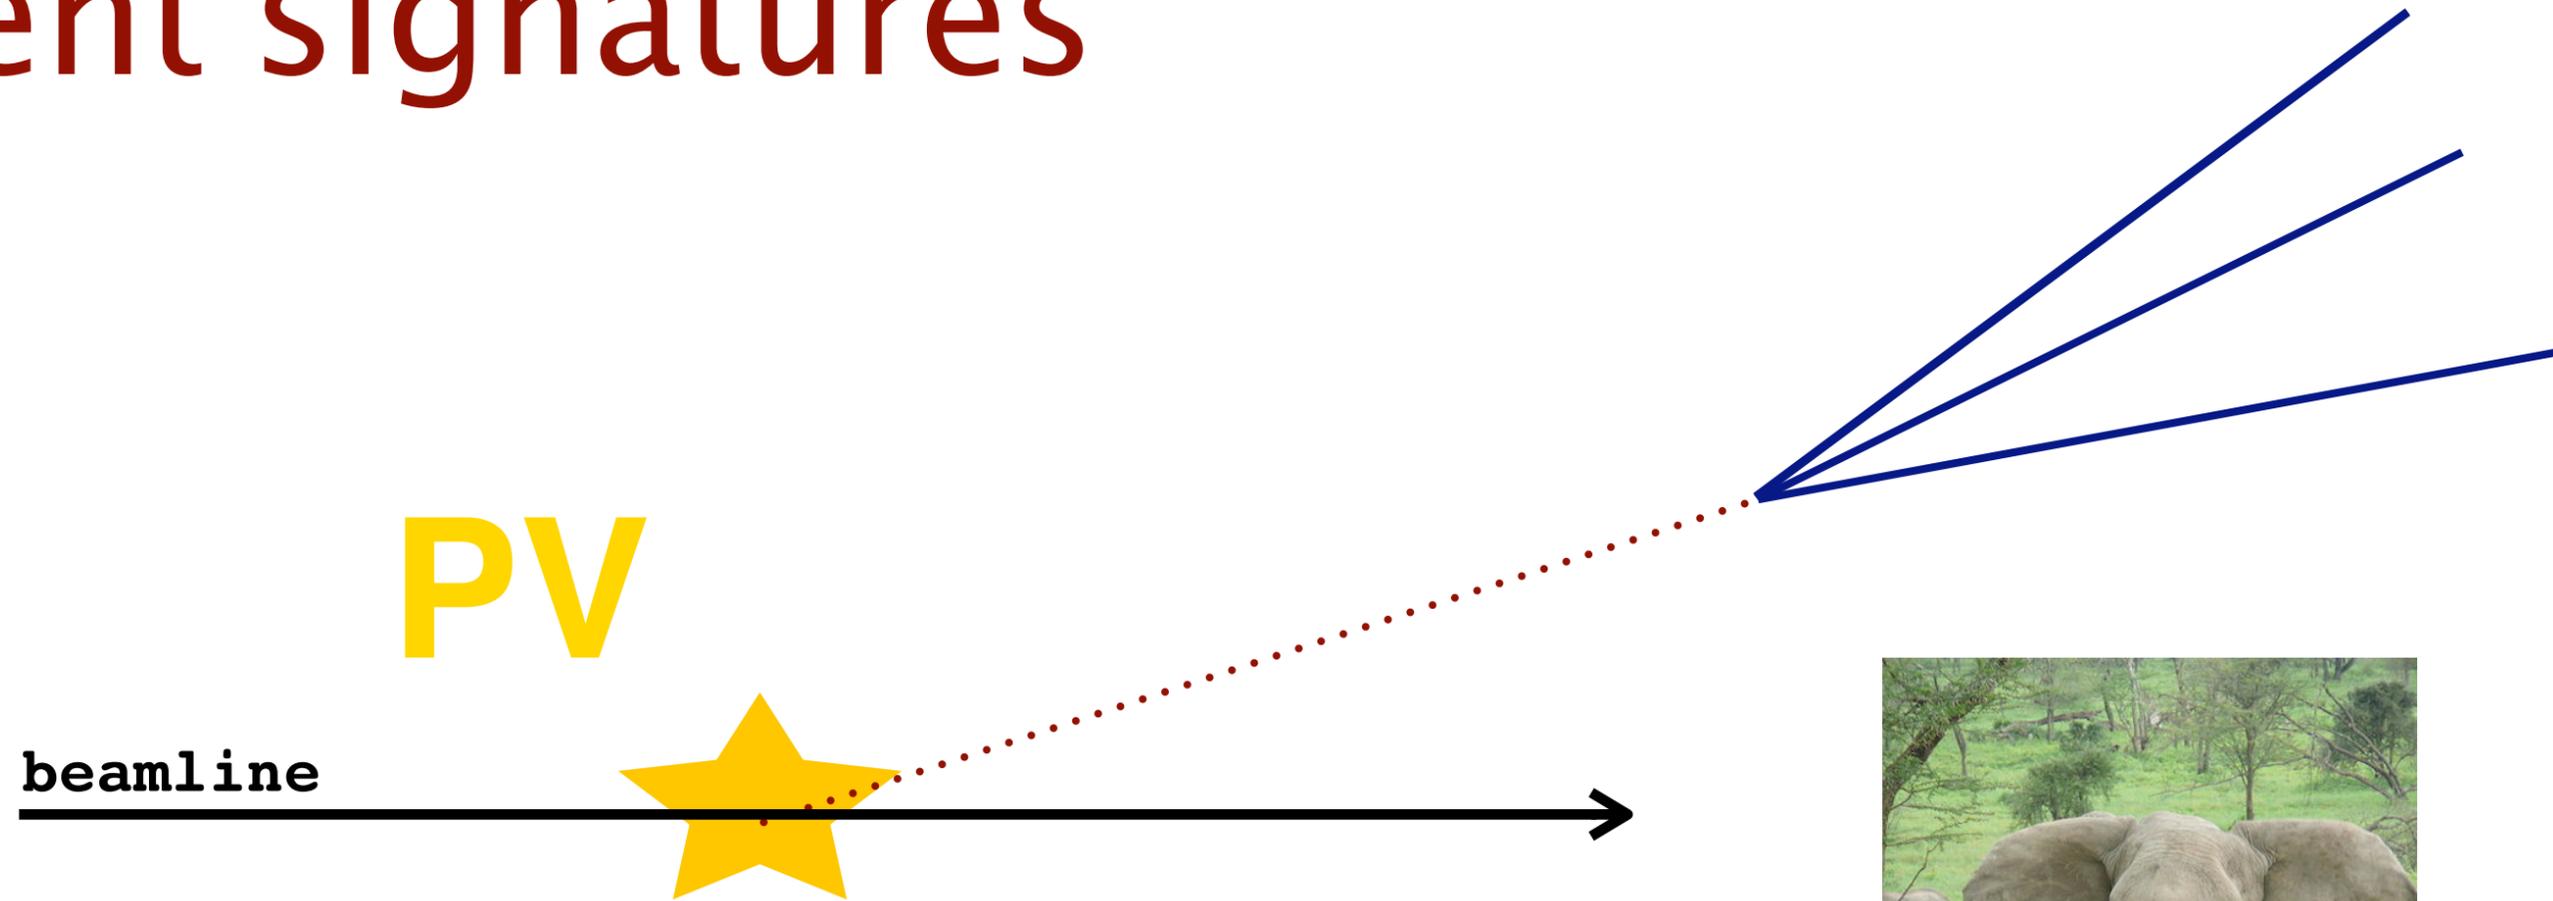


One big caveat in all of this

-- This whole concept relies on the fact that individual trigger decisions are independent of each other!

Triggering b-hadrons at LHCb

B event signatures



B meson signatures :

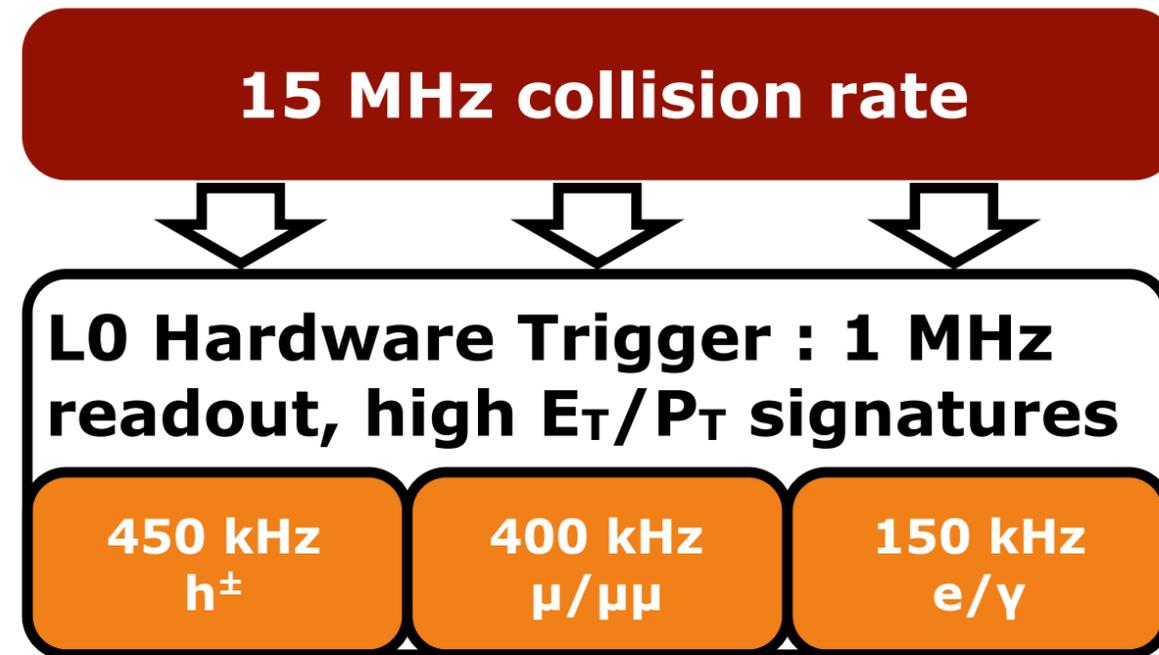
Large child transverse momentum

**Large child impact parameter or
vertex displacement**

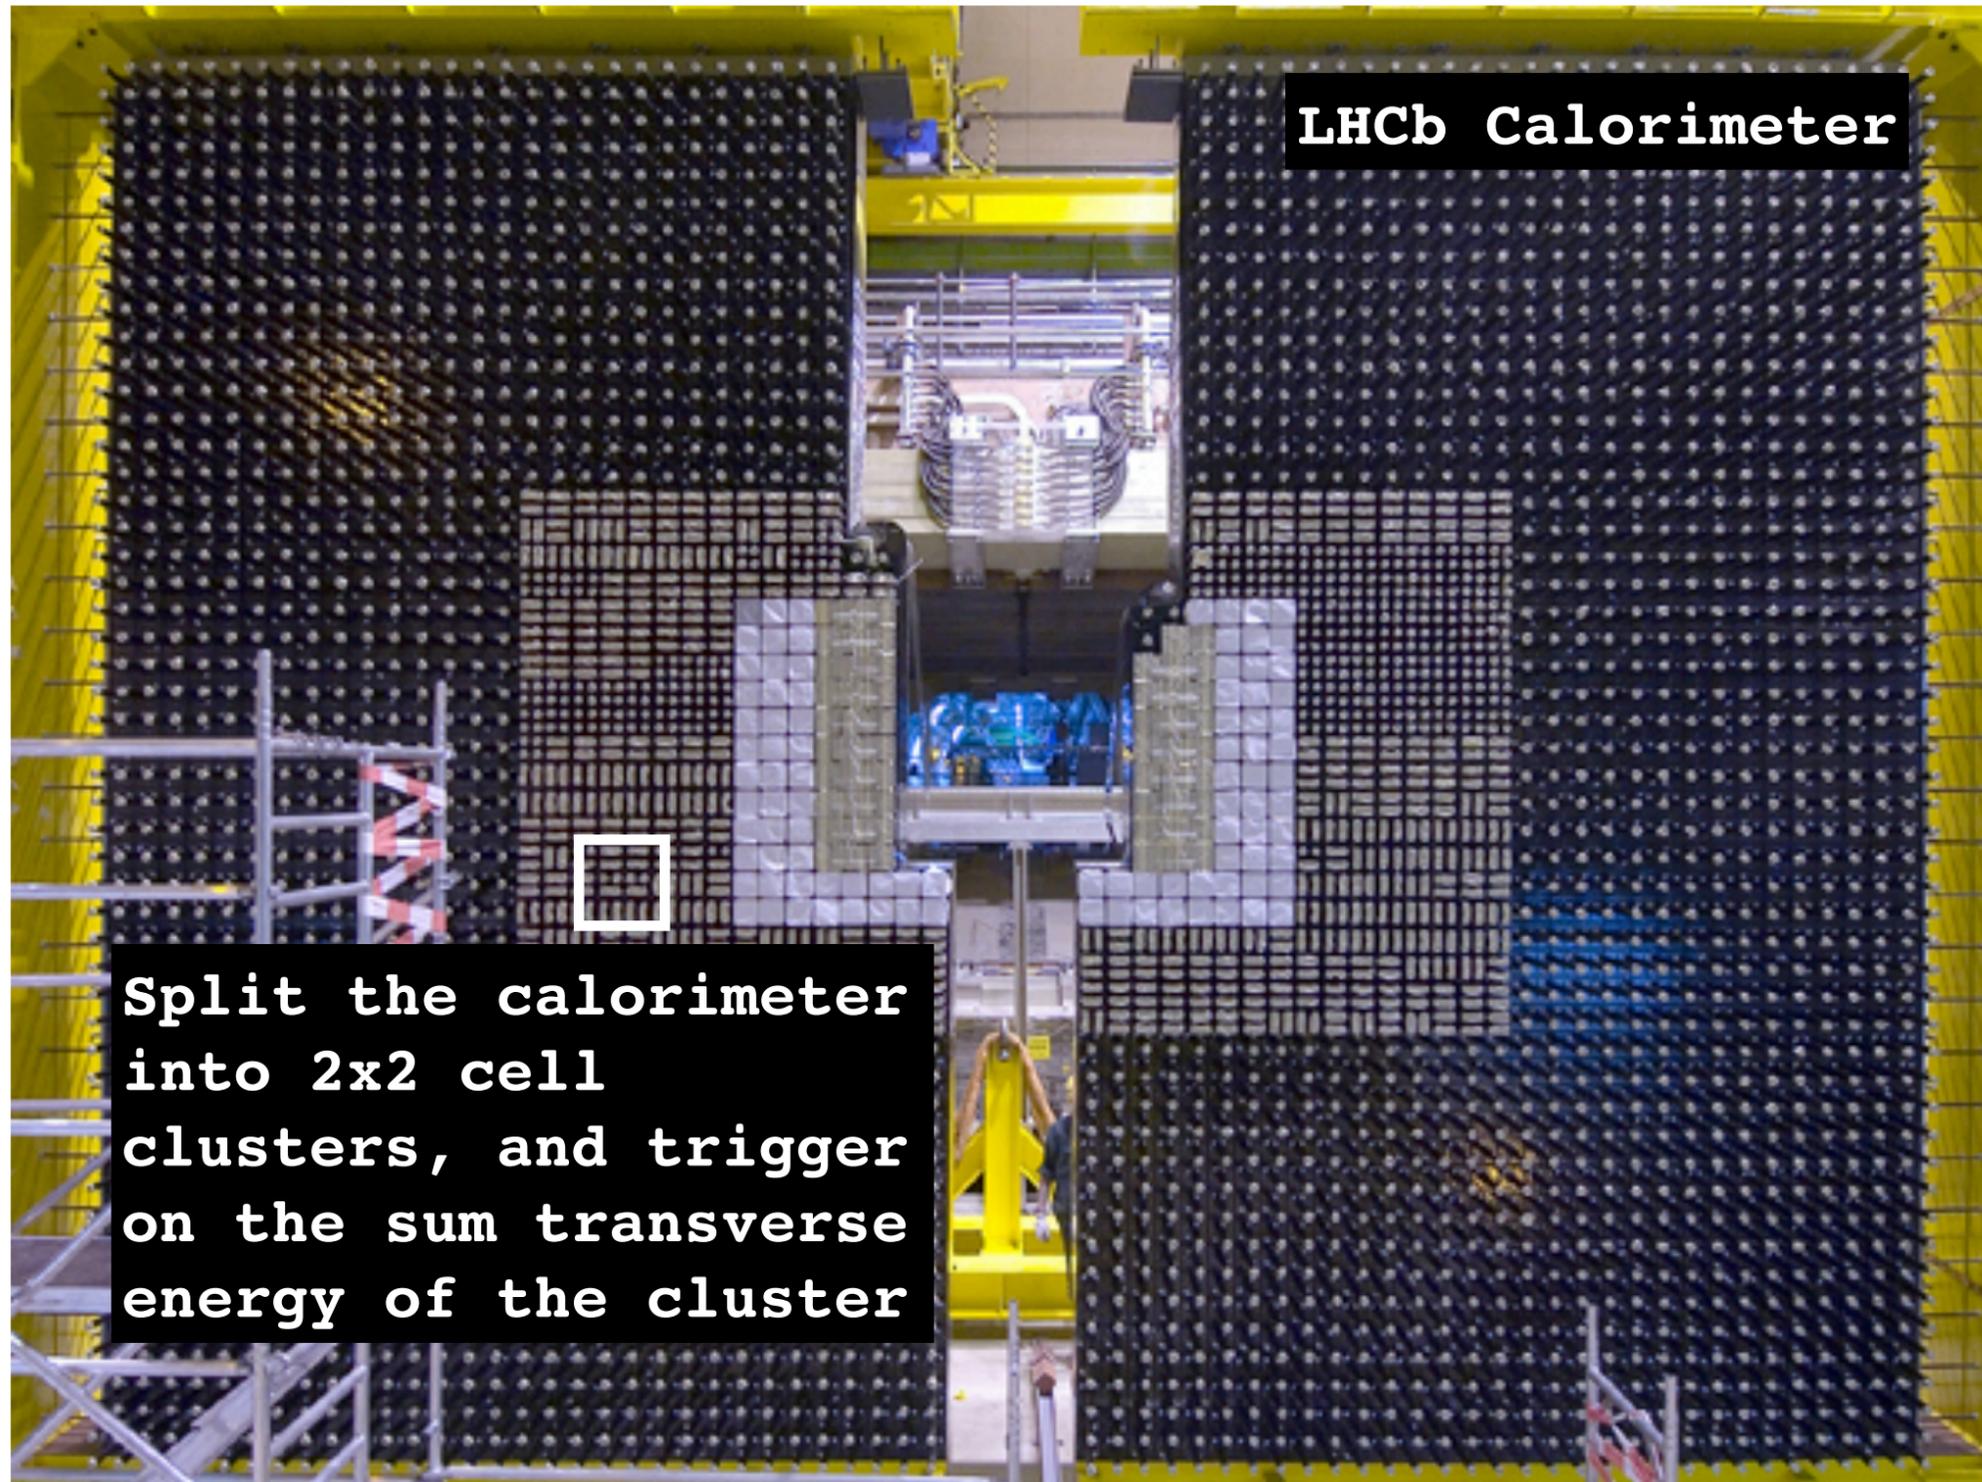


**"A B is the elephant of the particle zoo: it is very heavy and
lives a long time" -- T. Schietinger**

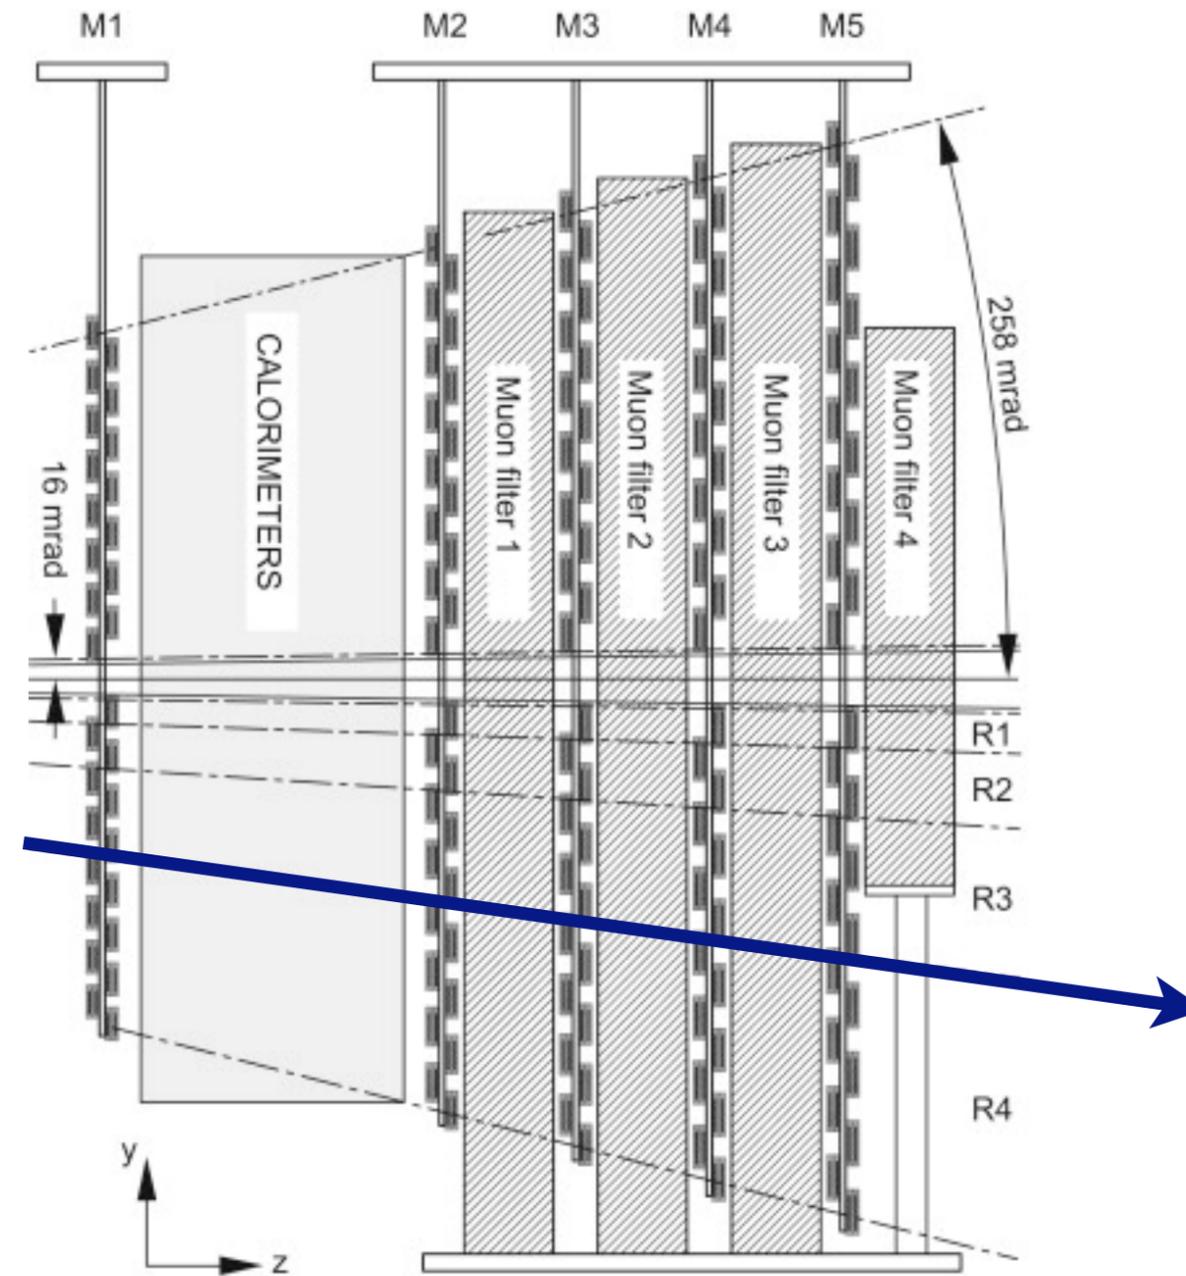
The hardware trigger stage



The calorimeter hardware trigger

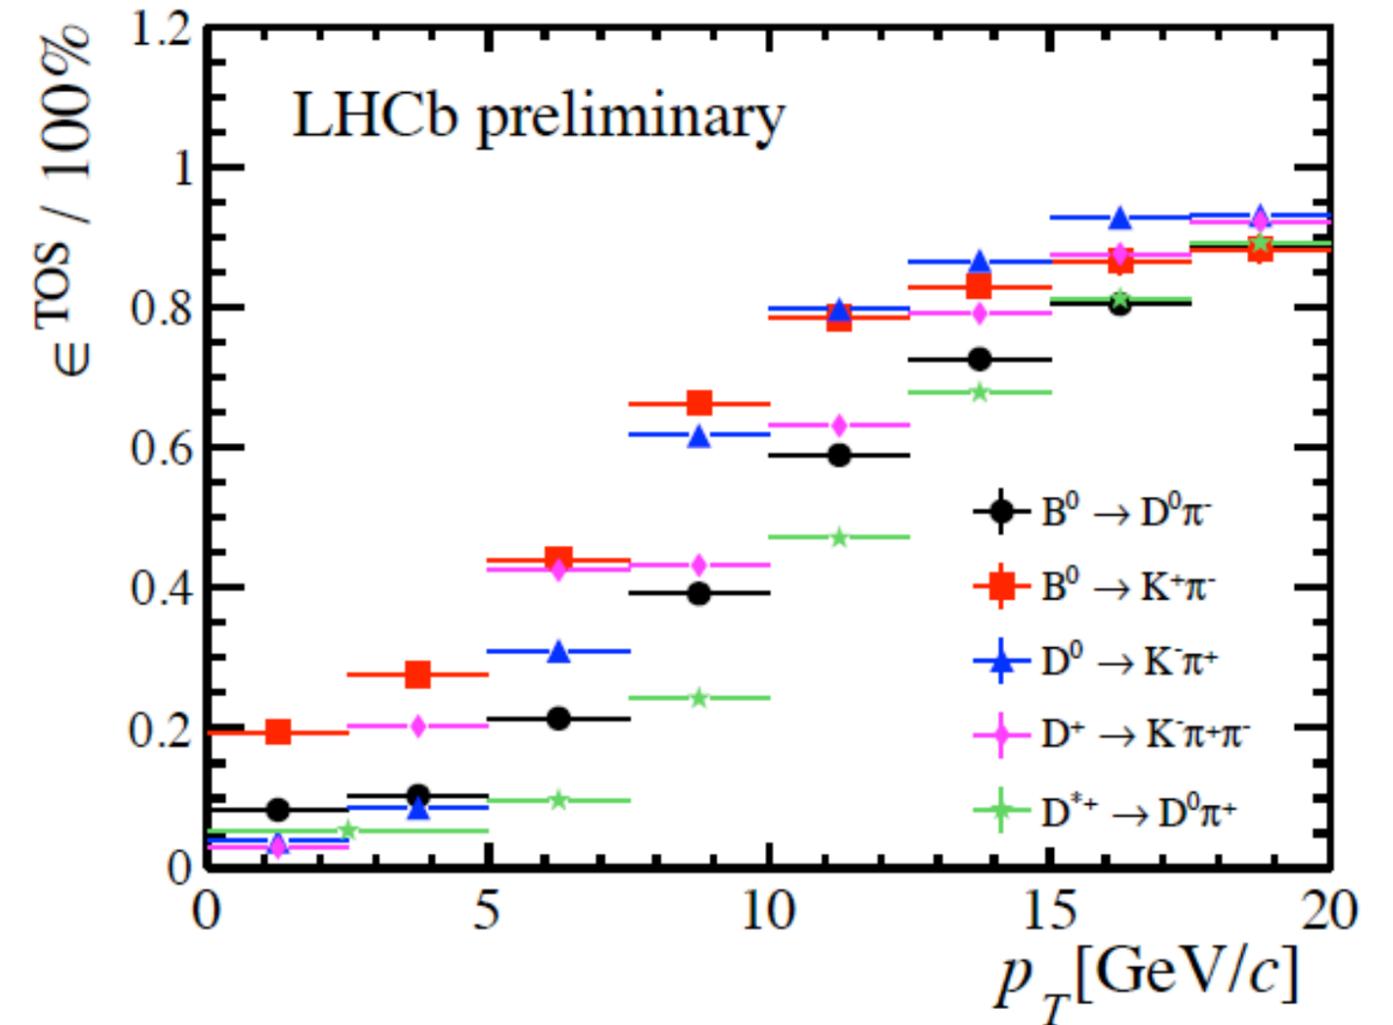
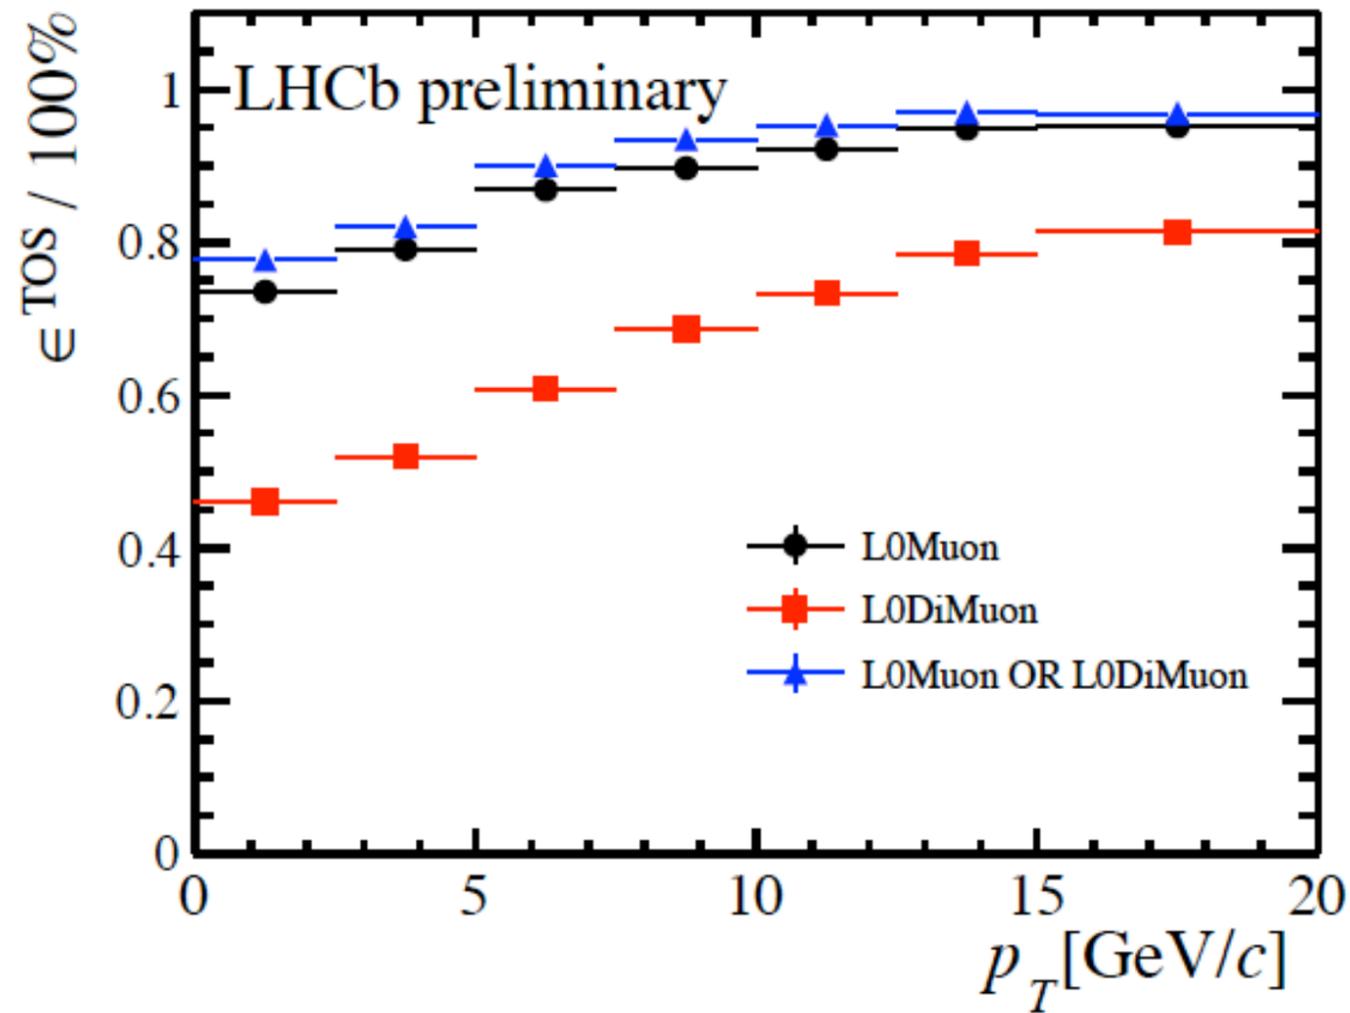


The muon hardware trigger



For muons search for track in first three stations

L0 trigger performance



L0 trigger summary



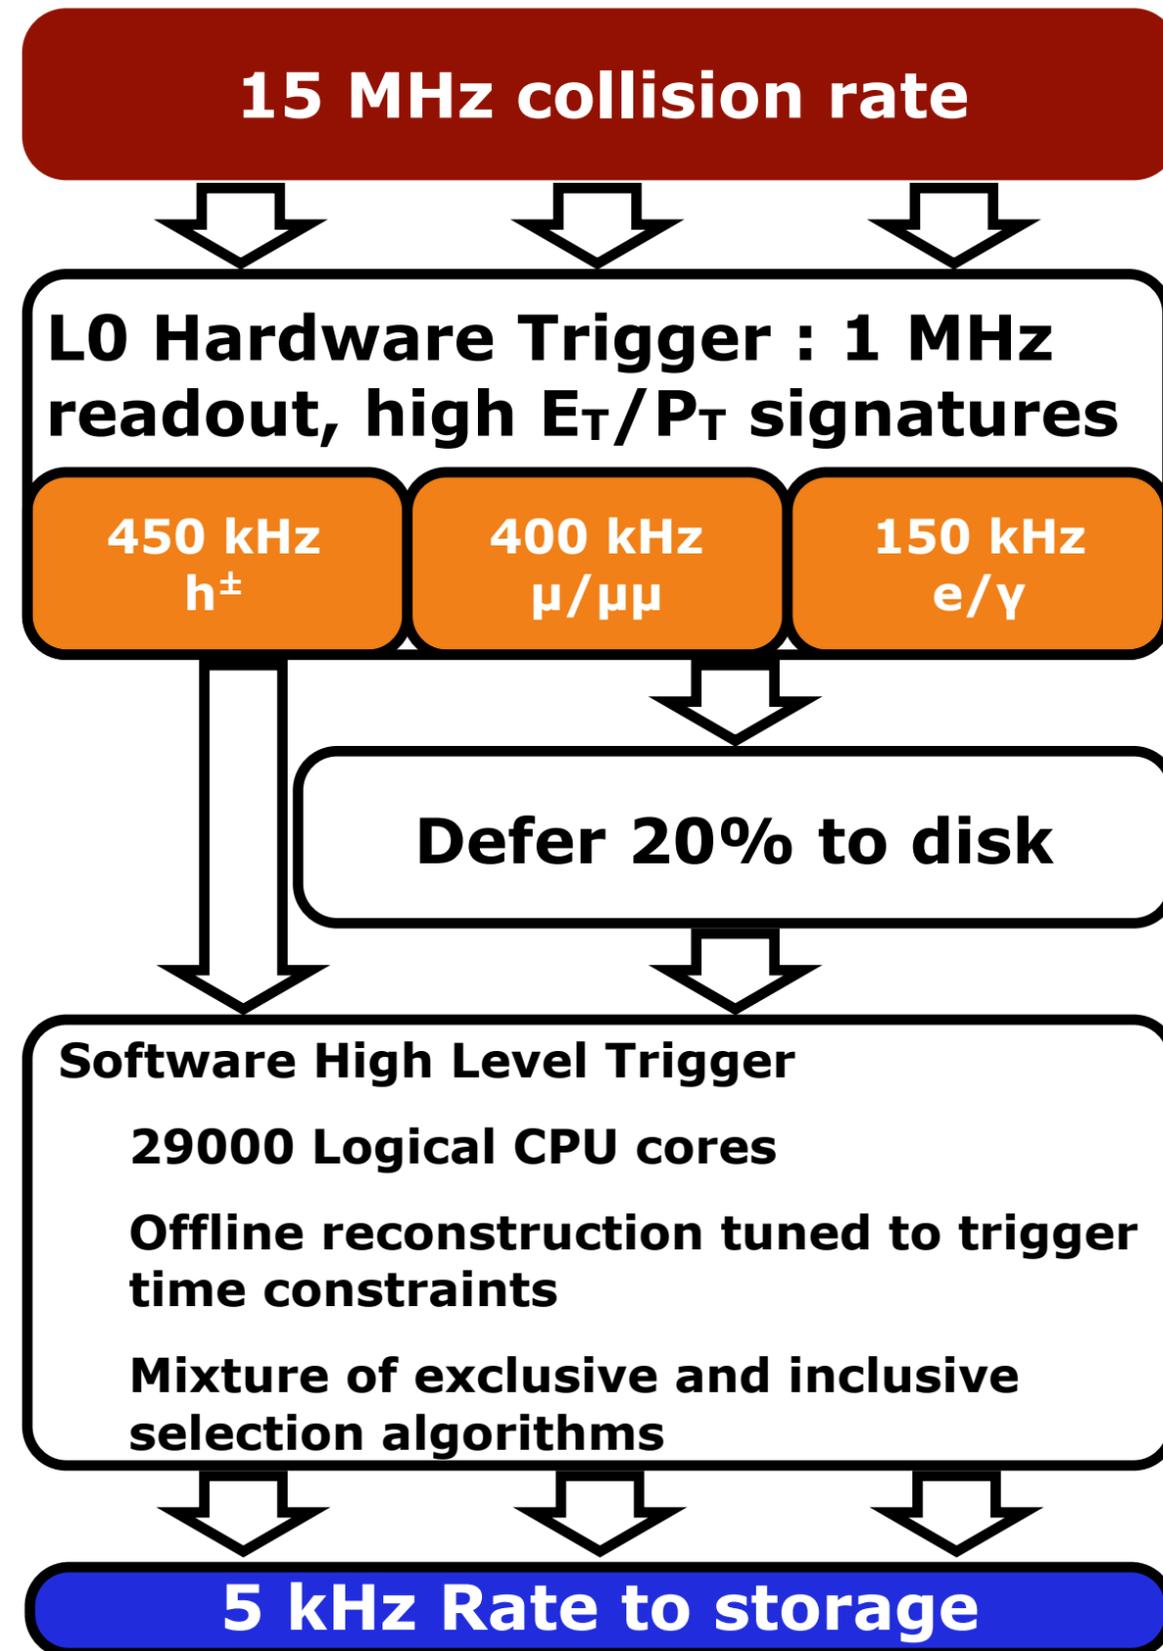
It may be necessary temporarily to accept a lesser evil, but one must never label a necessary evil as good.

(Margaret Mead)

izquotes.com

I'll explain later how the LHCb upgrade will make the L0 trigger unnecessary

The software trigger stage



Selection cascades

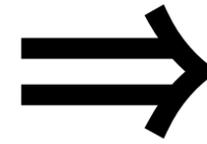
1.

Information gathering
("reconstruction") stage

Selection cascades

1.

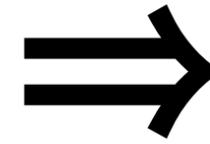
Information gathering
("reconstruction") stage



Selection cascades

1.

Information gathering
("reconstruction") stage



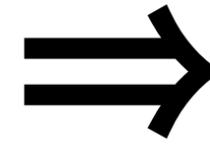
2.

Event selection stage

Selection cascades

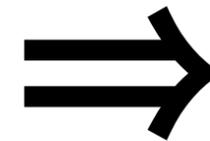
1.

Information gathering
("reconstruction") stage



2.

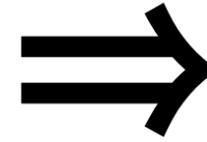
Event selection stage



Selection cascades

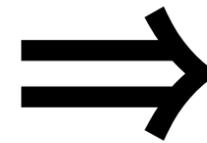
1.

Information gathering
("reconstruction") stage



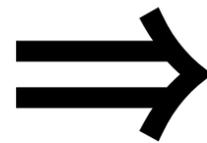
2.

Event selection stage



3.

Next reconstruction stage



A small conceptual moment

Real time event selection

=>

Budgeted learning

**See Killian's
talk after me**

HLT1Track reconstruction

1.

Full reconstruction of tracks in vertex locator

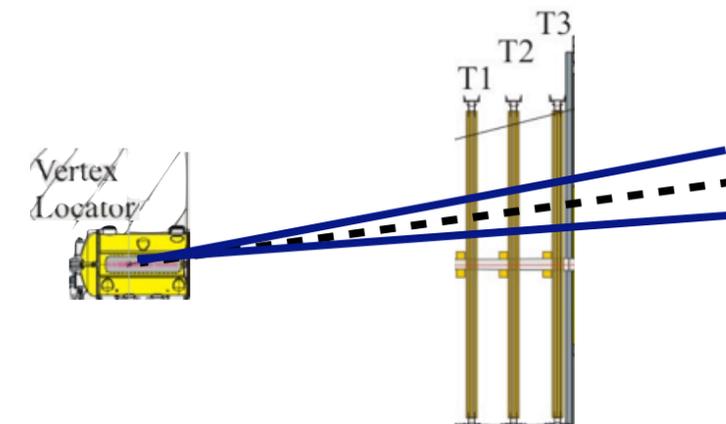
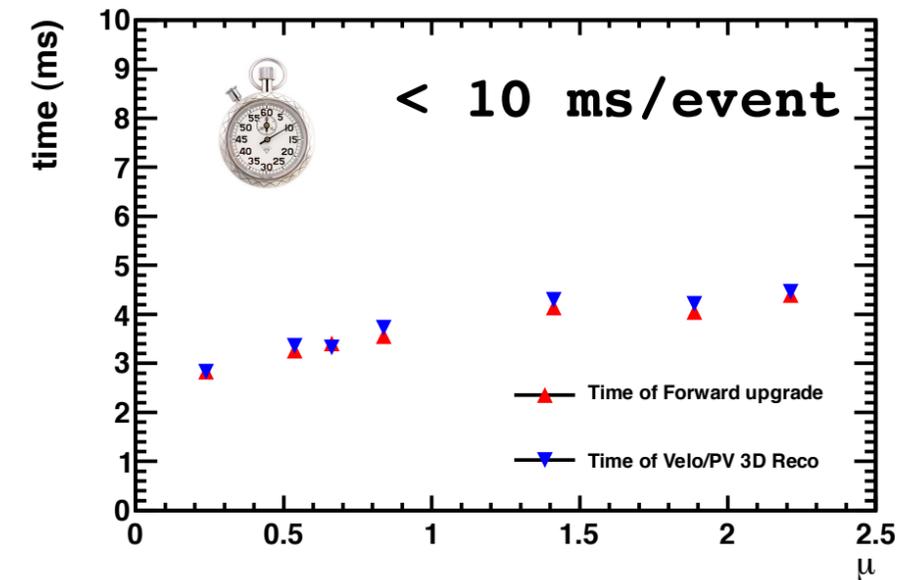


Select displaced tracks



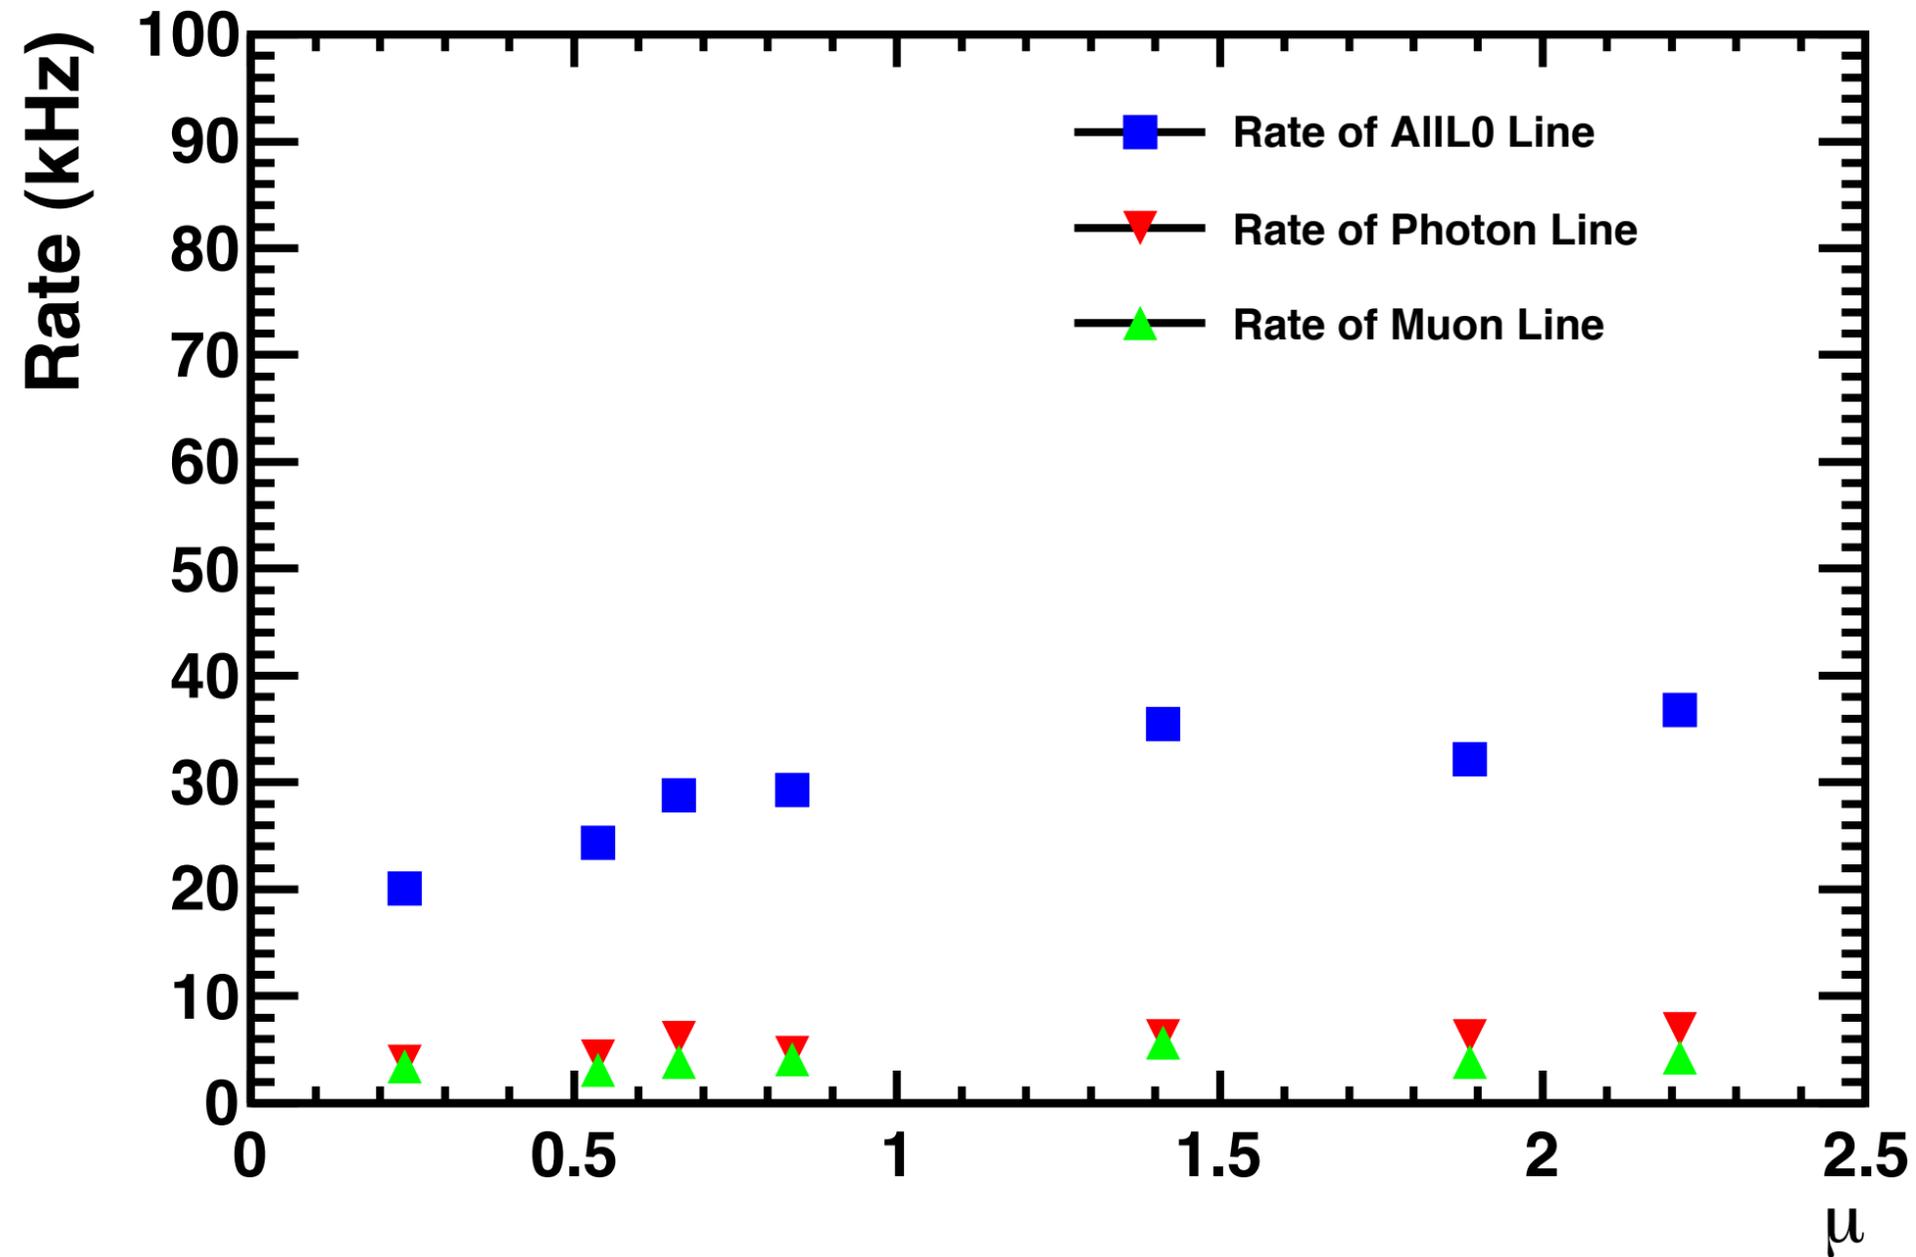
2.

Reconstruction of displaced tracks in regions of interest



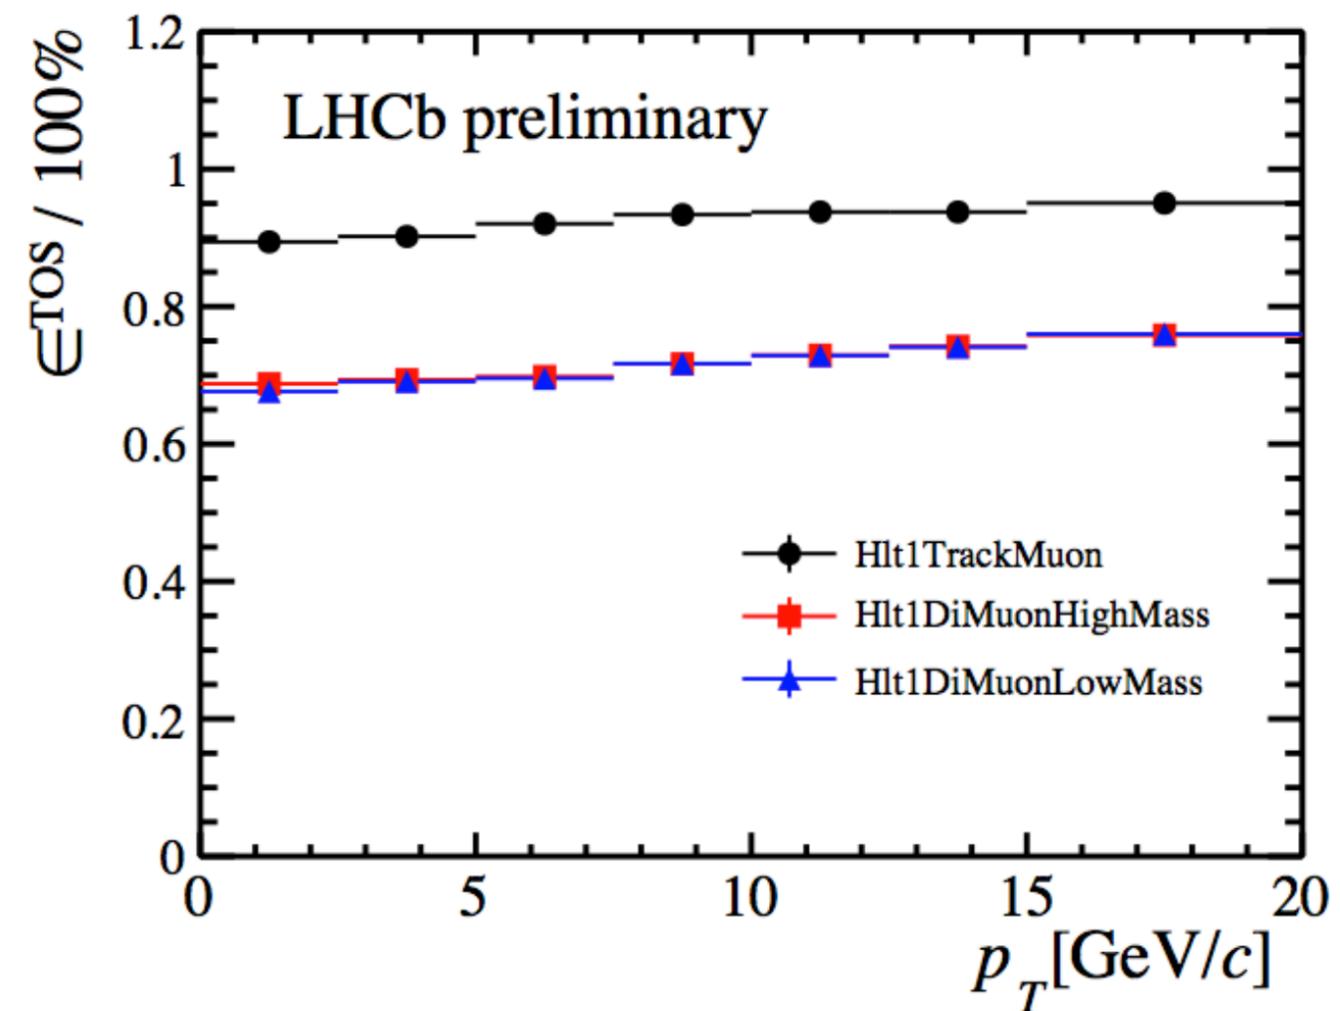
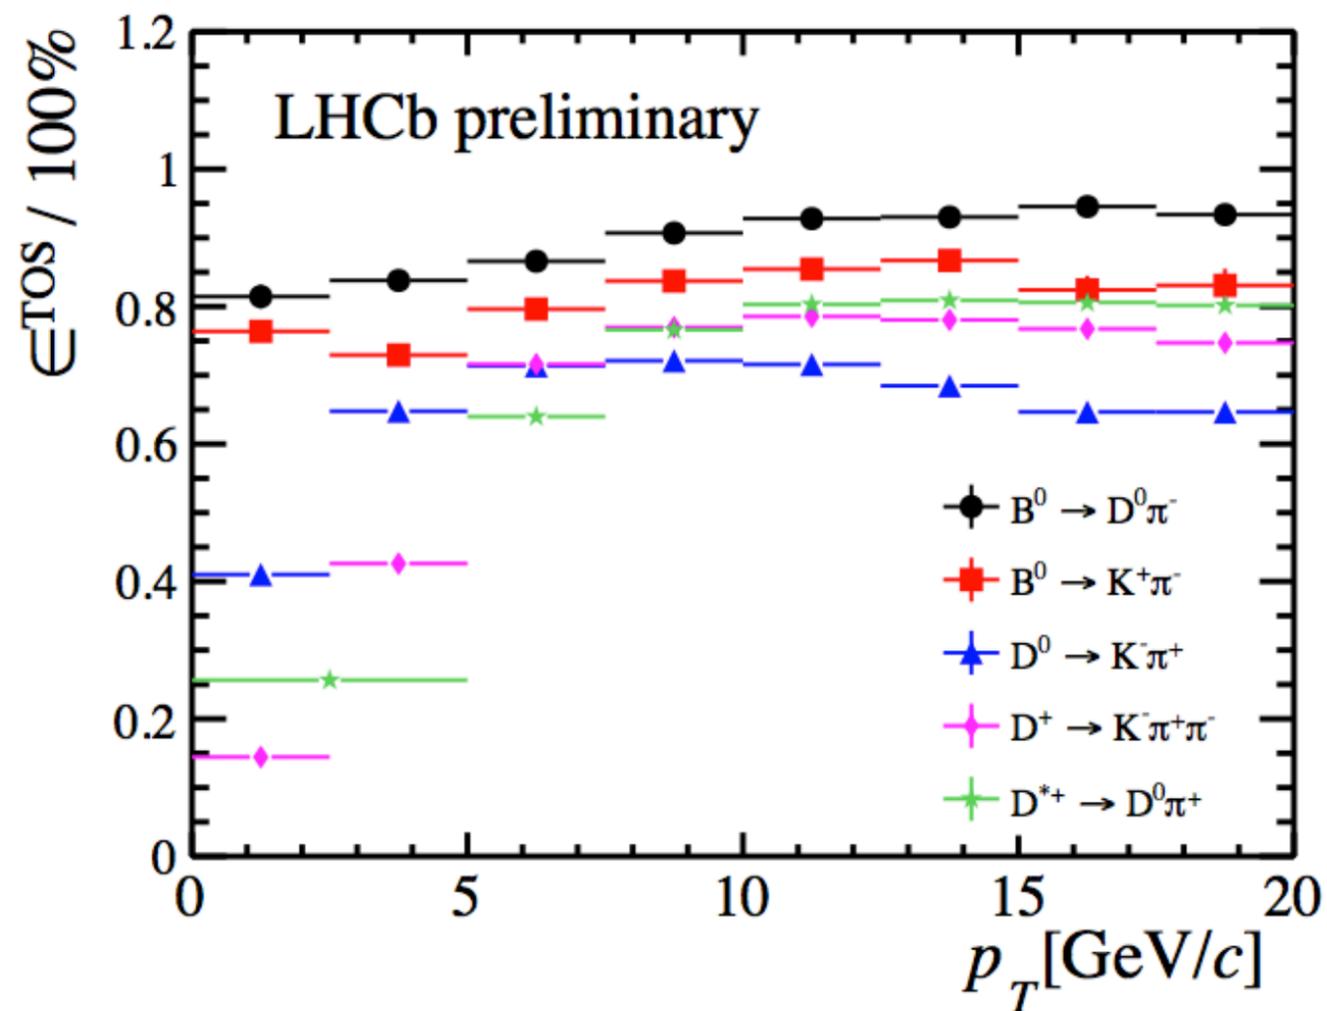
Region of interest defined by assumed track P/P_T , 3/1.6 GeV in 2012

Single track triggering and pileup

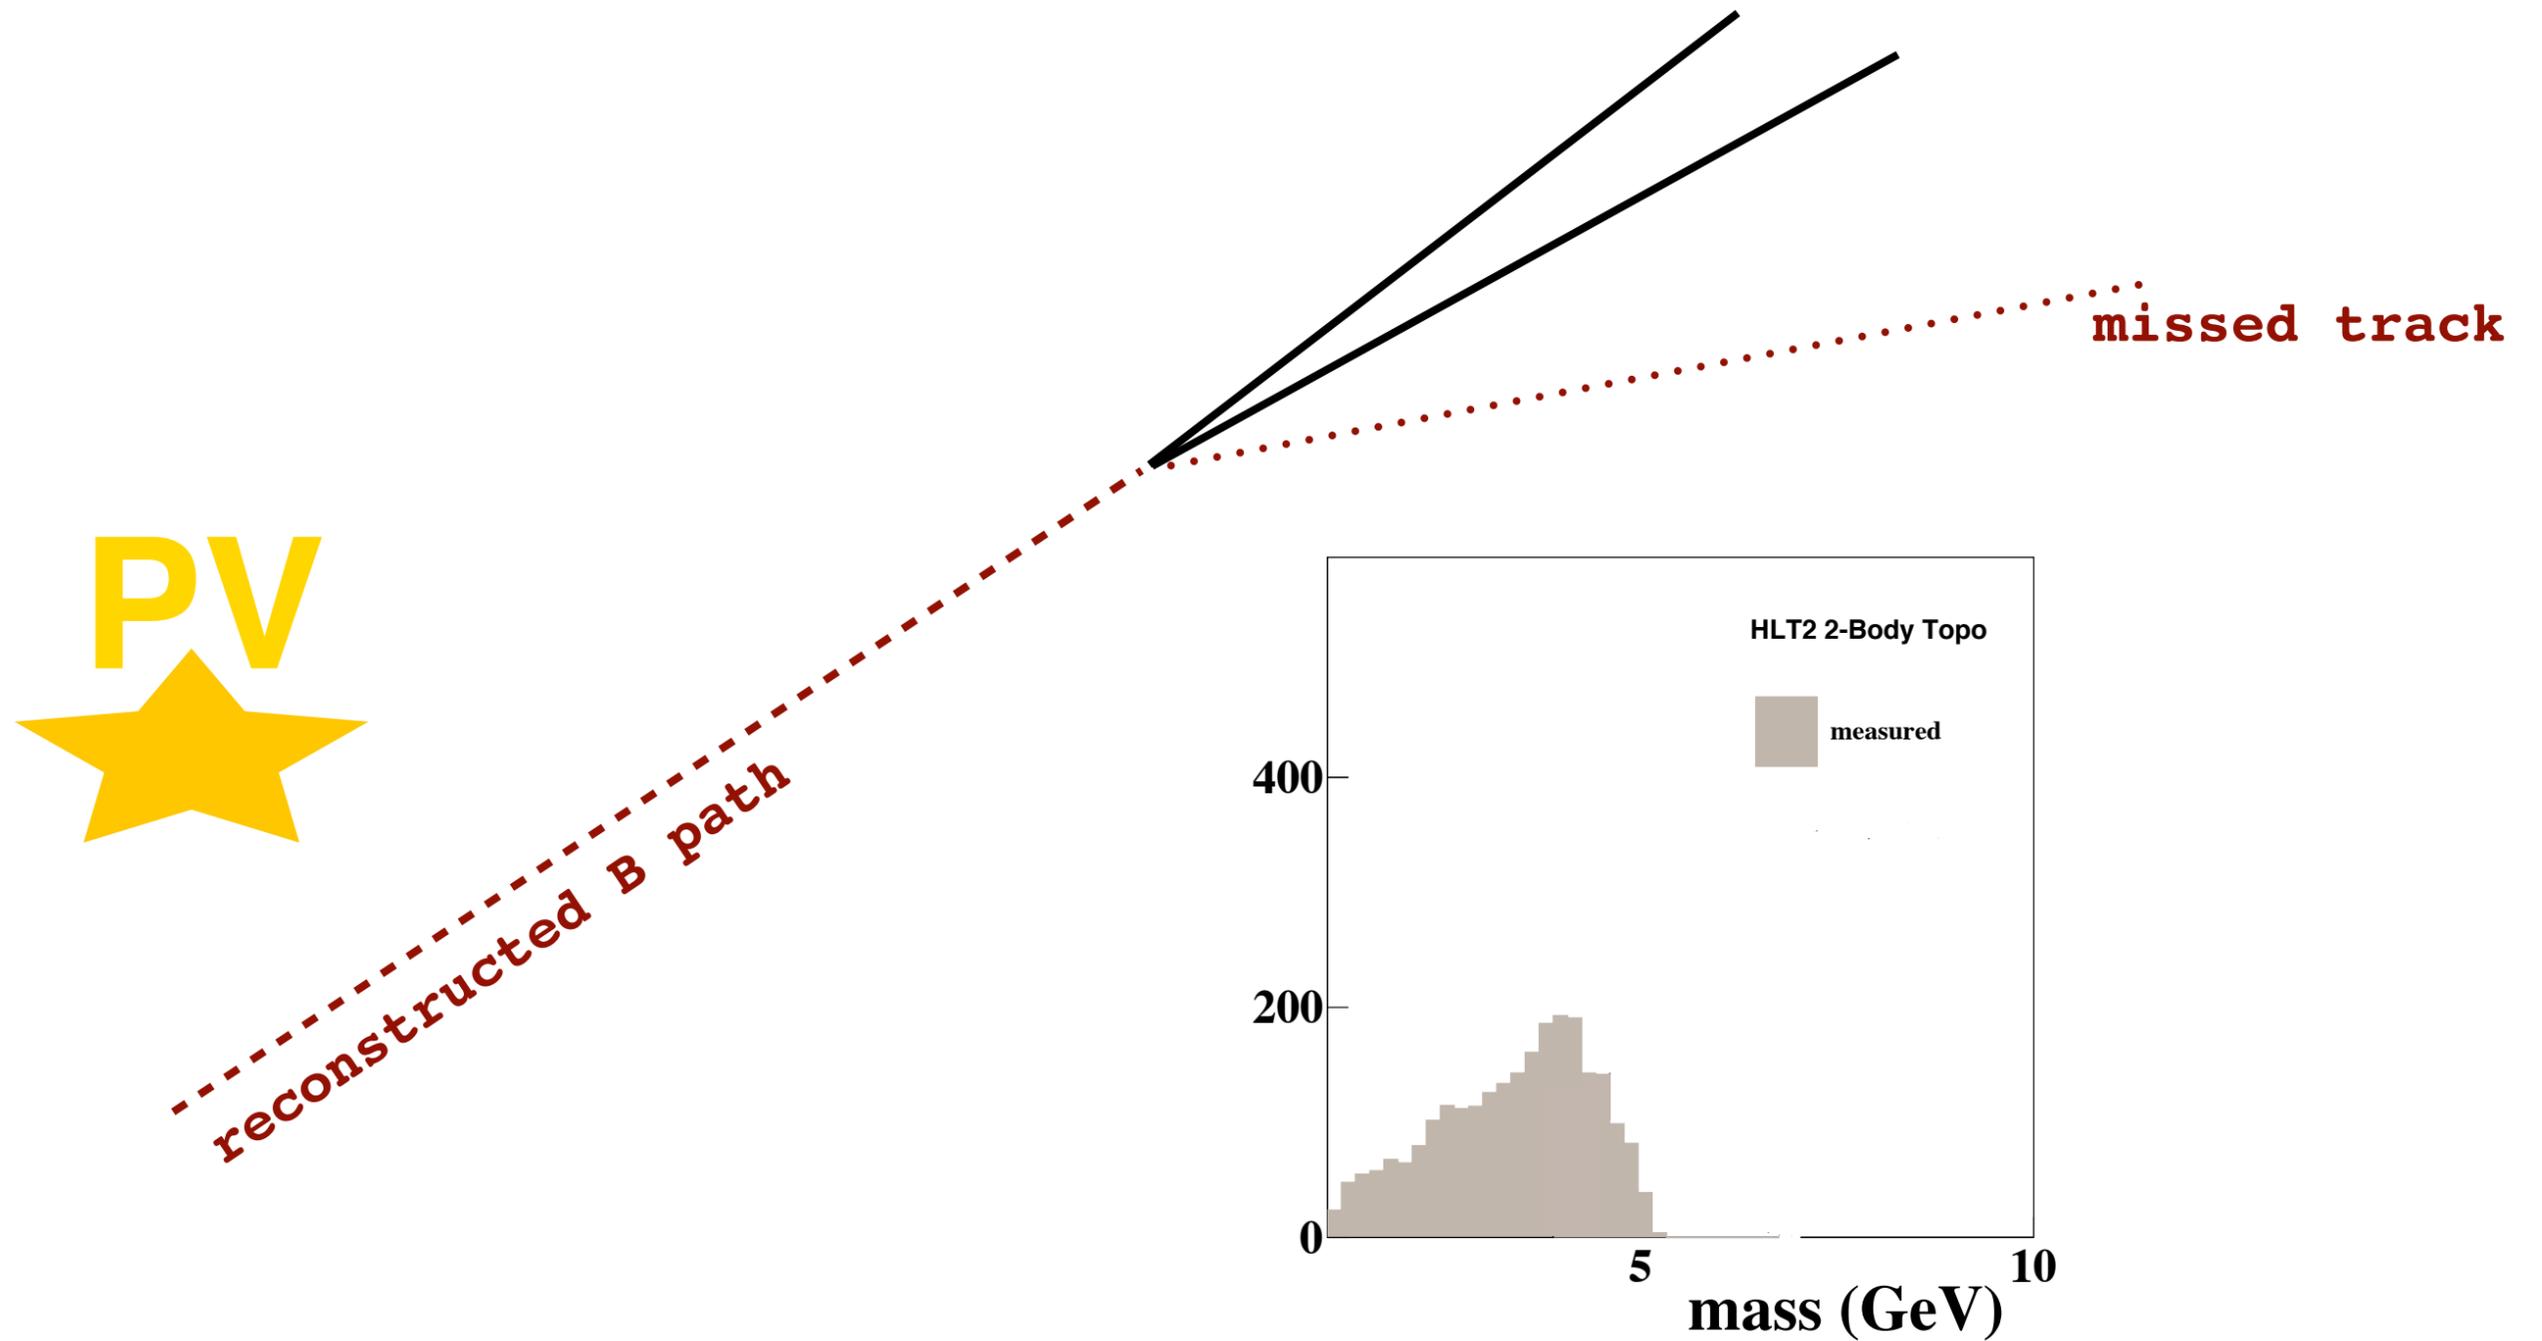


Triggering on a single track makes you basically insensitive to increasing combinatorics, since there is nothing to combine!

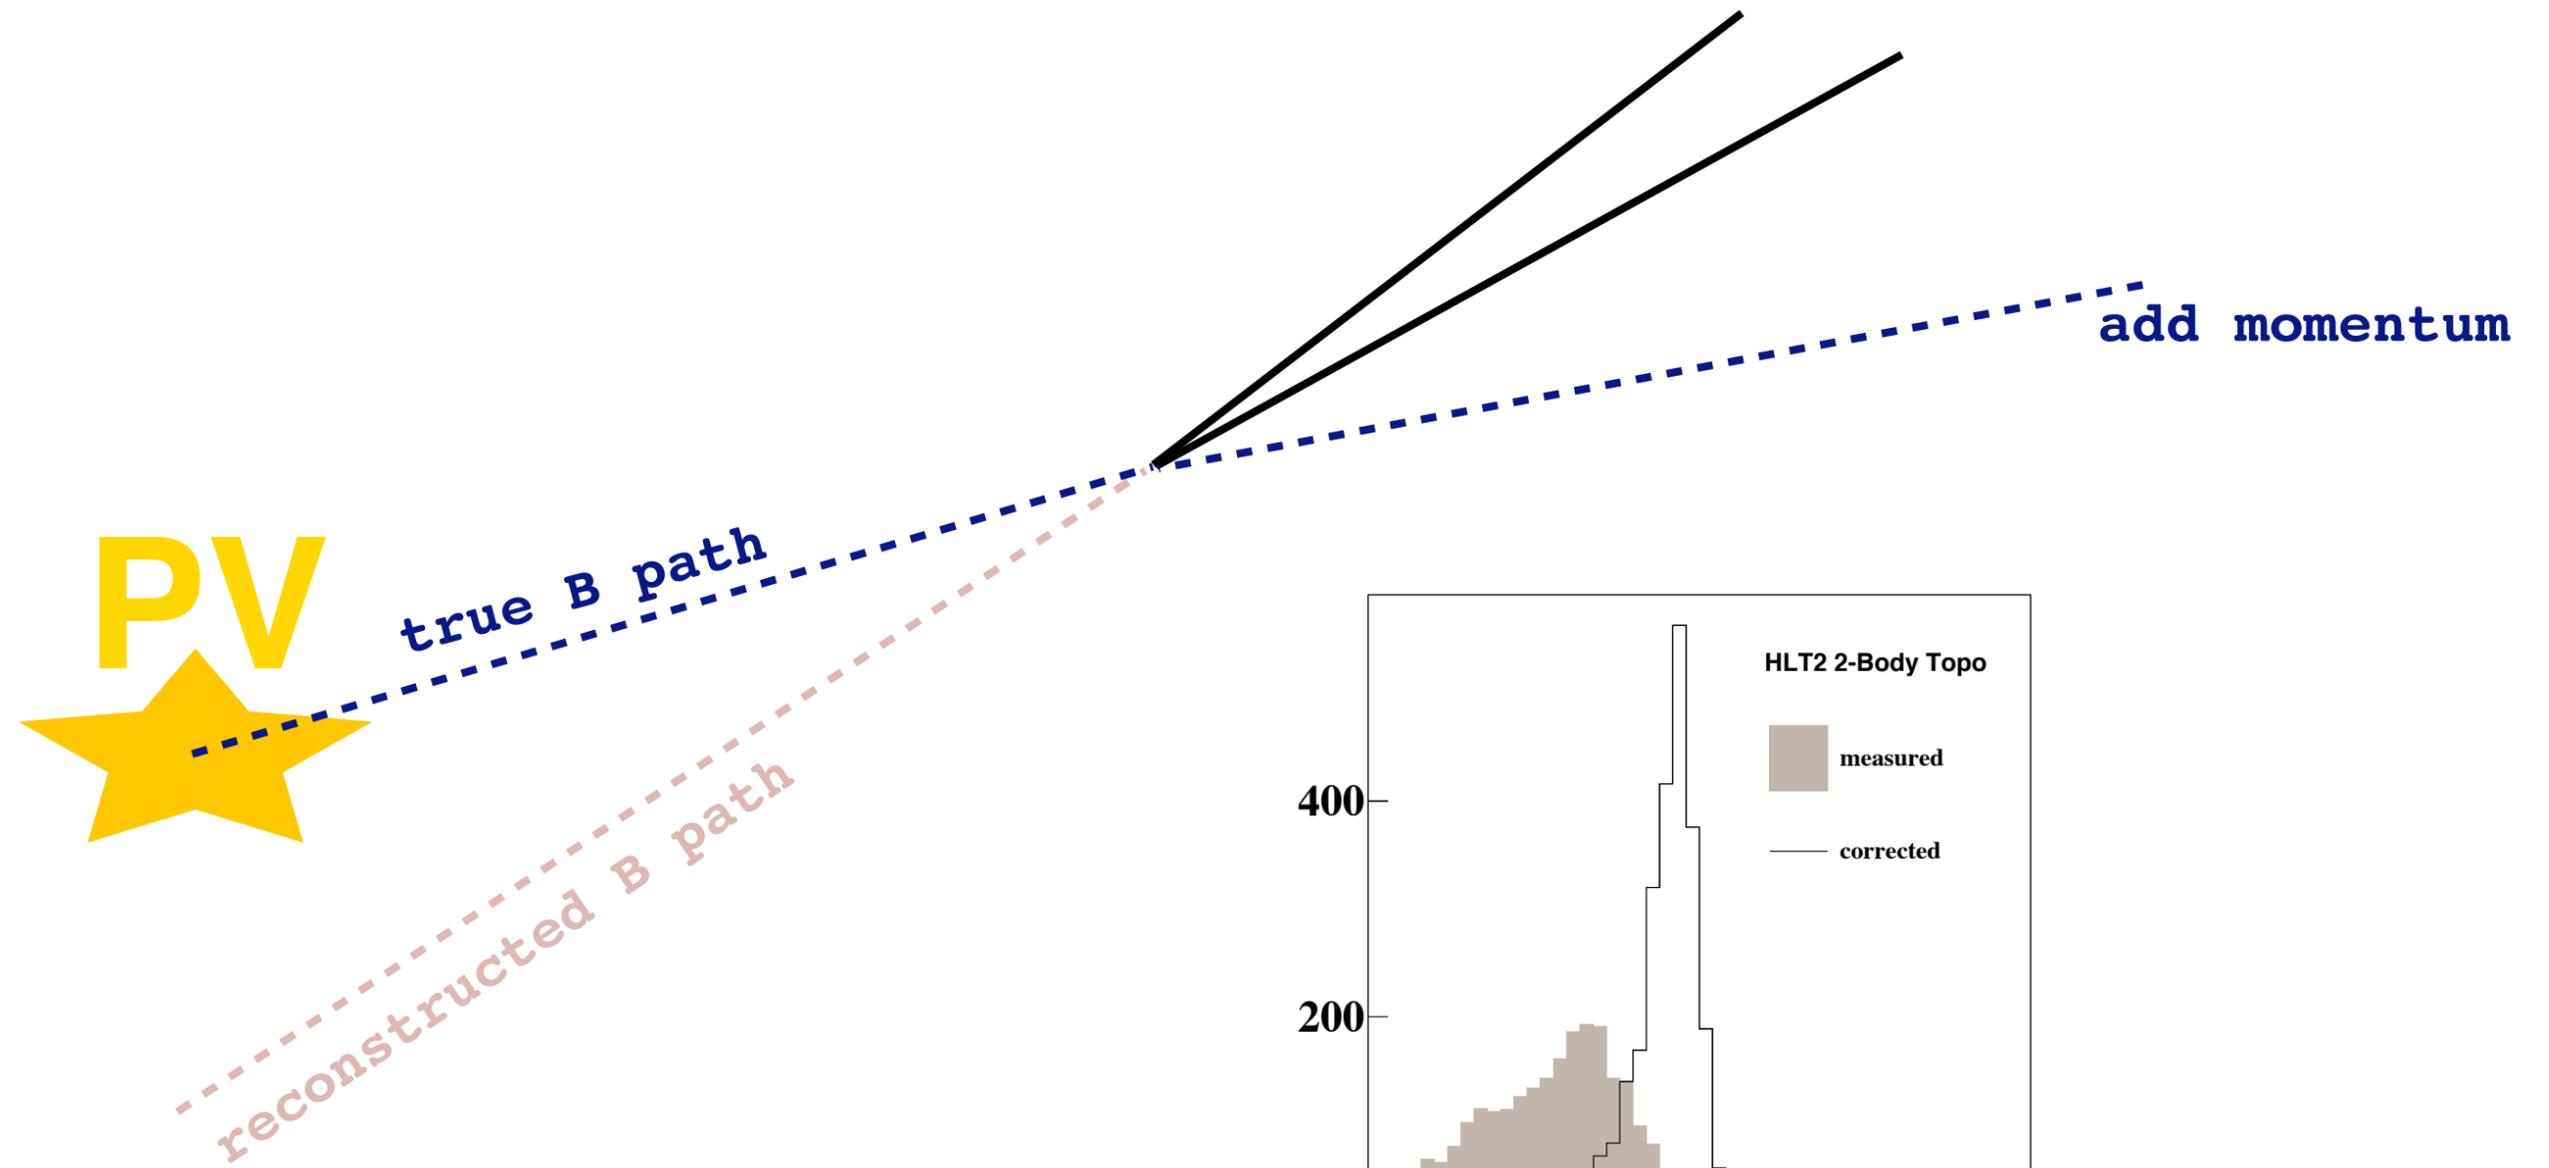
HLT1 signal performance



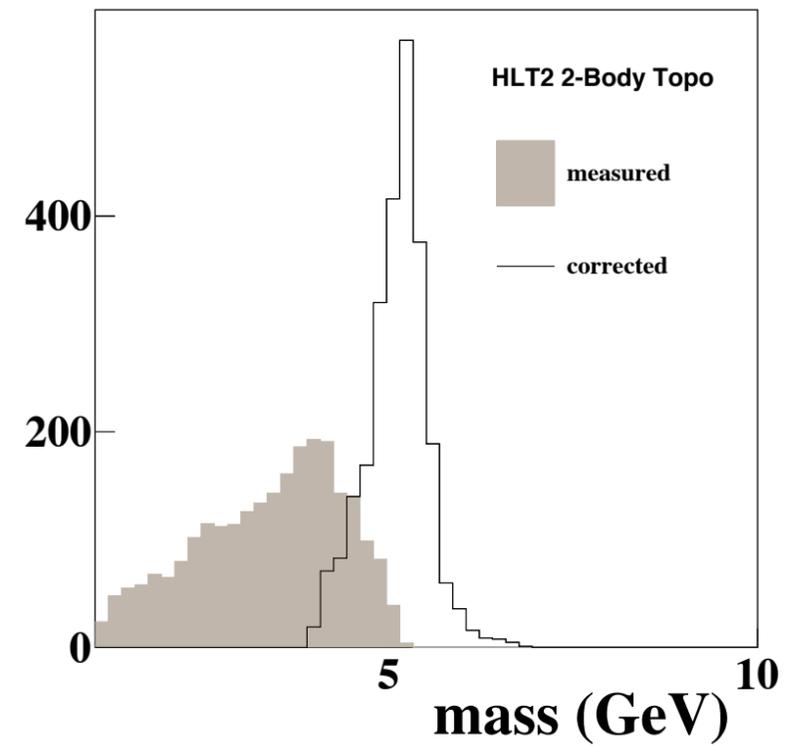
50 → 5 kHz : the topological trigger



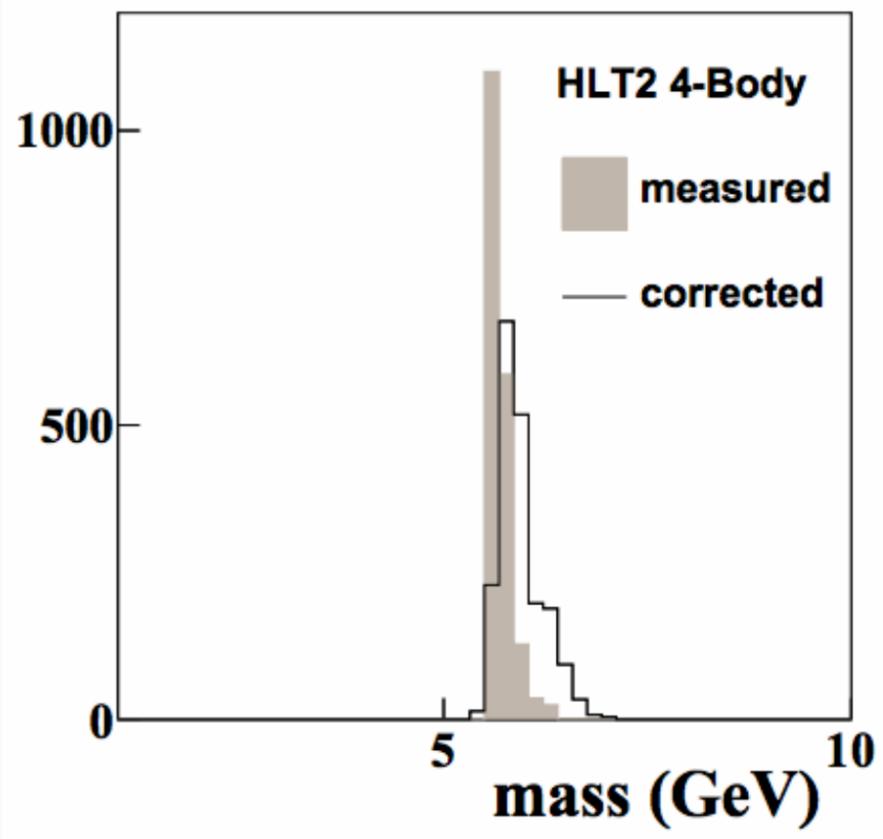
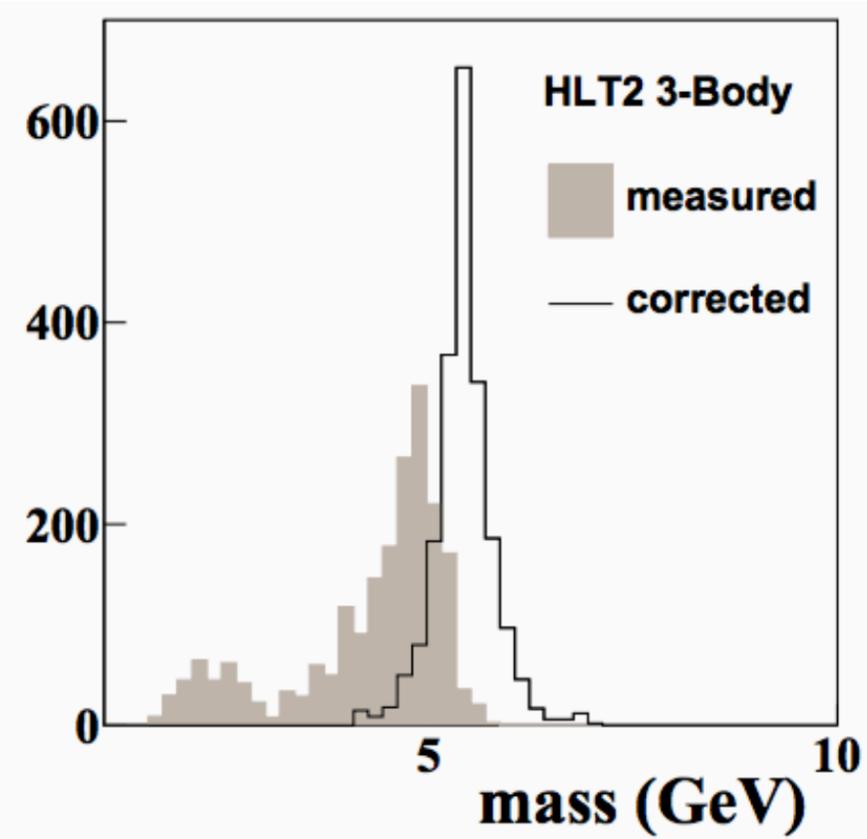
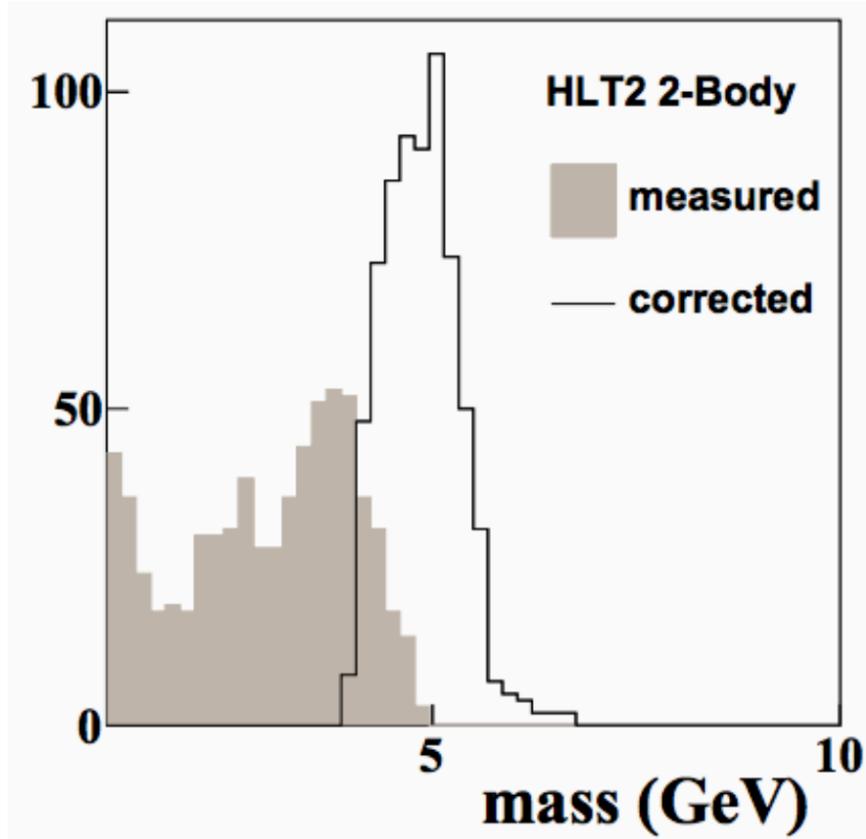
50 → 5 kHz : the topological trigger



$$m_{\text{corrected}} = \sqrt{m^2 + |p'_{T\text{missing}}|^2 + |p'_{T\text{missing}}|}$$

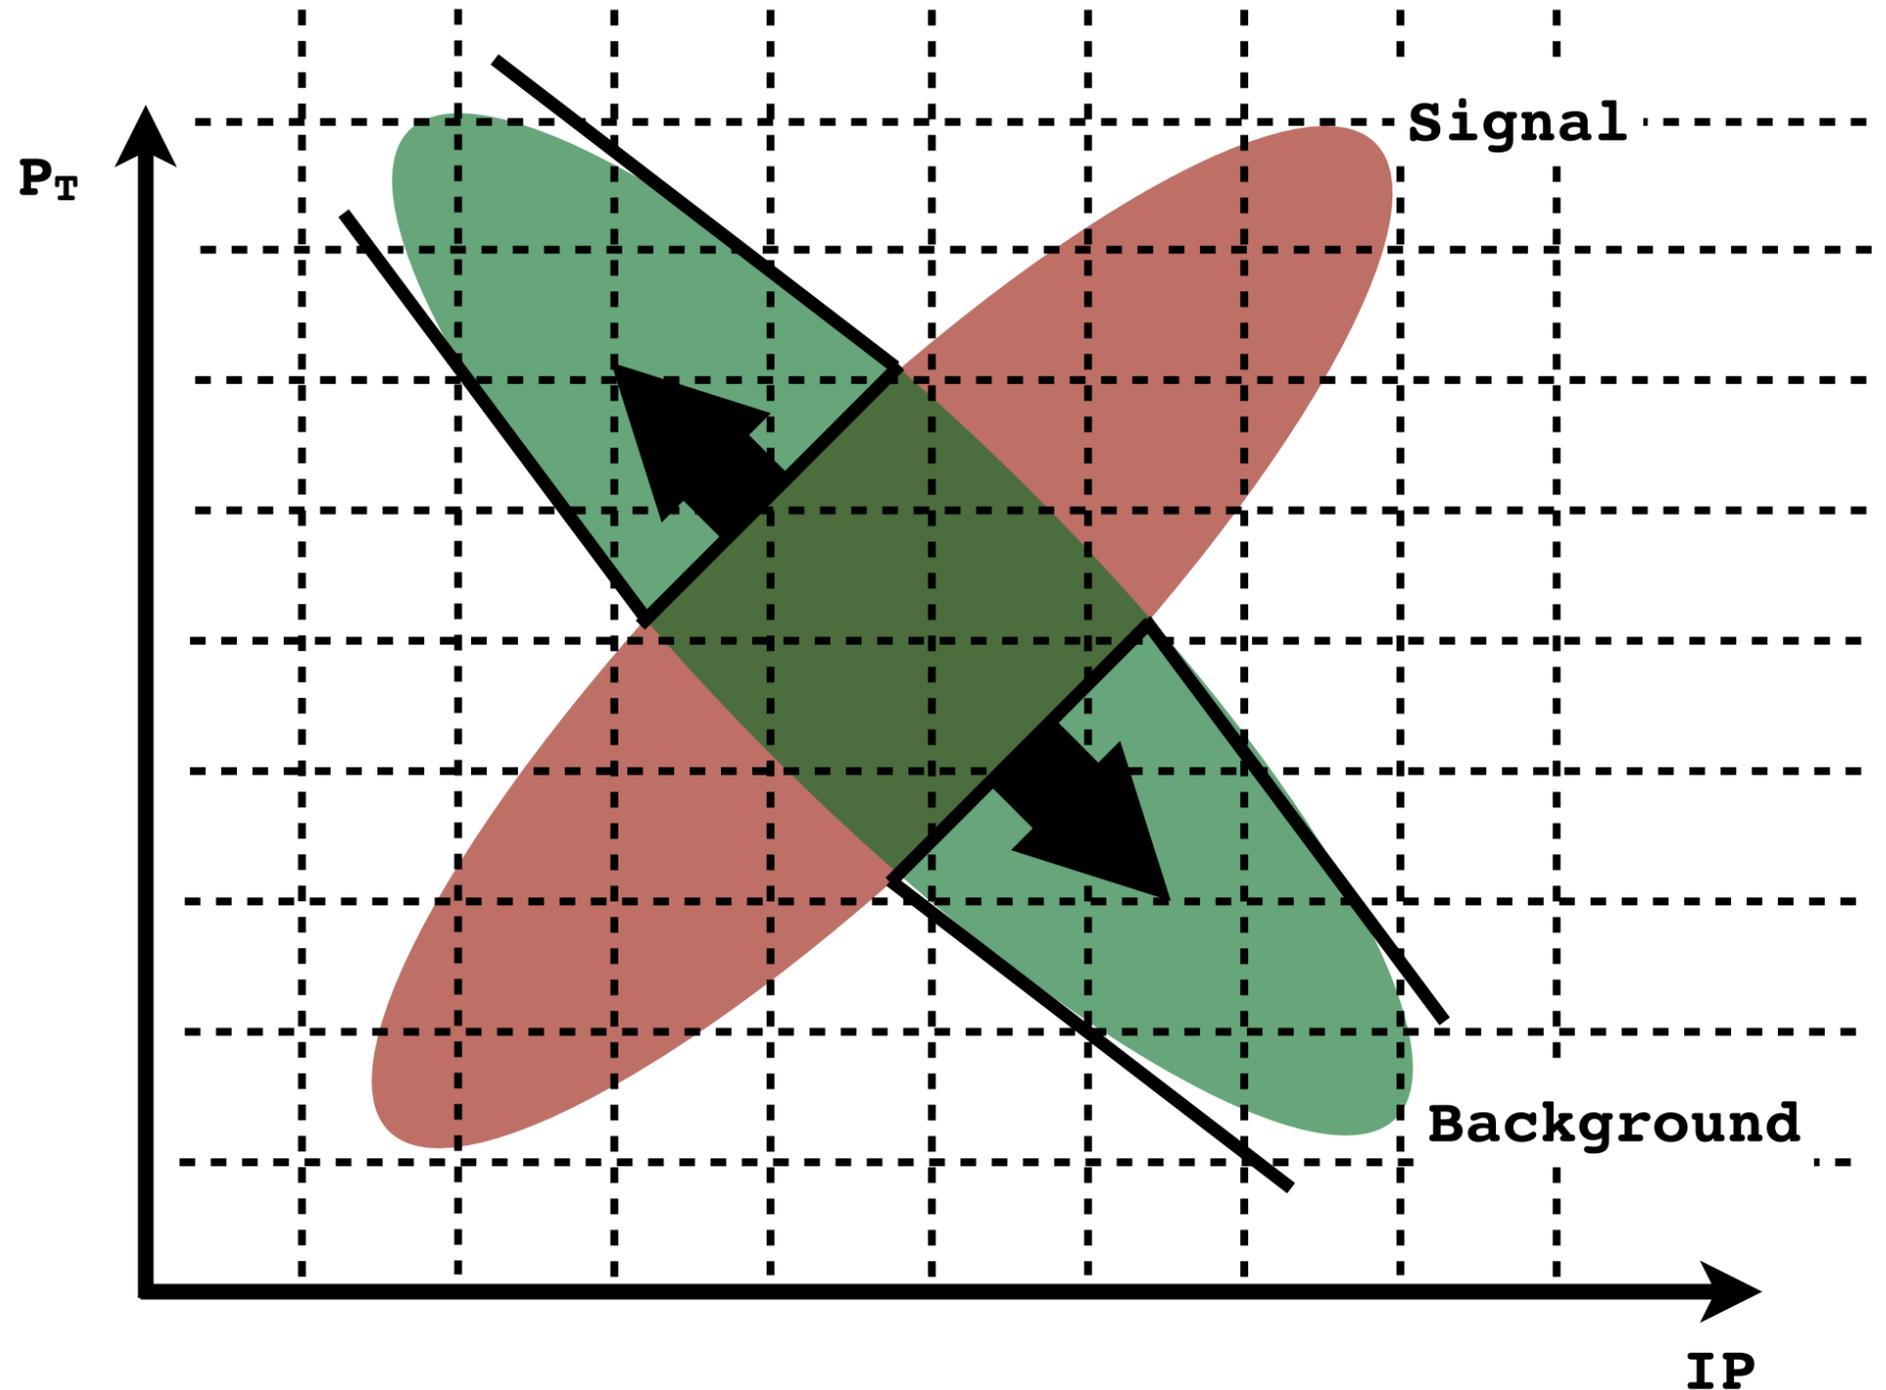


Corrected mass and vtx multiplicity



A bansai boosted decision tree

Consider a two-variable boosted decision tree : this is like a binned selection where the BDT algorithm picks the optimal bin sizes and boundaries



The challenges of running online

- If the keep regions are small relative to the resolution or stability of the detector, the signal could oscillate in and out of the keep regions. This would result in, at best, a less efficient trigger and, at worst, a trigger whose efficiency is very difficult to understand.
- In many cases the signal samples by necessity must come from simulations because the signals have, in fact, not yet been observed in data. In other cases the trigger is meant to be *inclusive, i.e.*, the trigger is meant to select classes of signal types rather than one specific signal channel. In both cases, the signal PDFs might not be completely accurate or even available during the training process.
- Any HLT algorithm must run in the online environment; thus, it must be extremely fast.

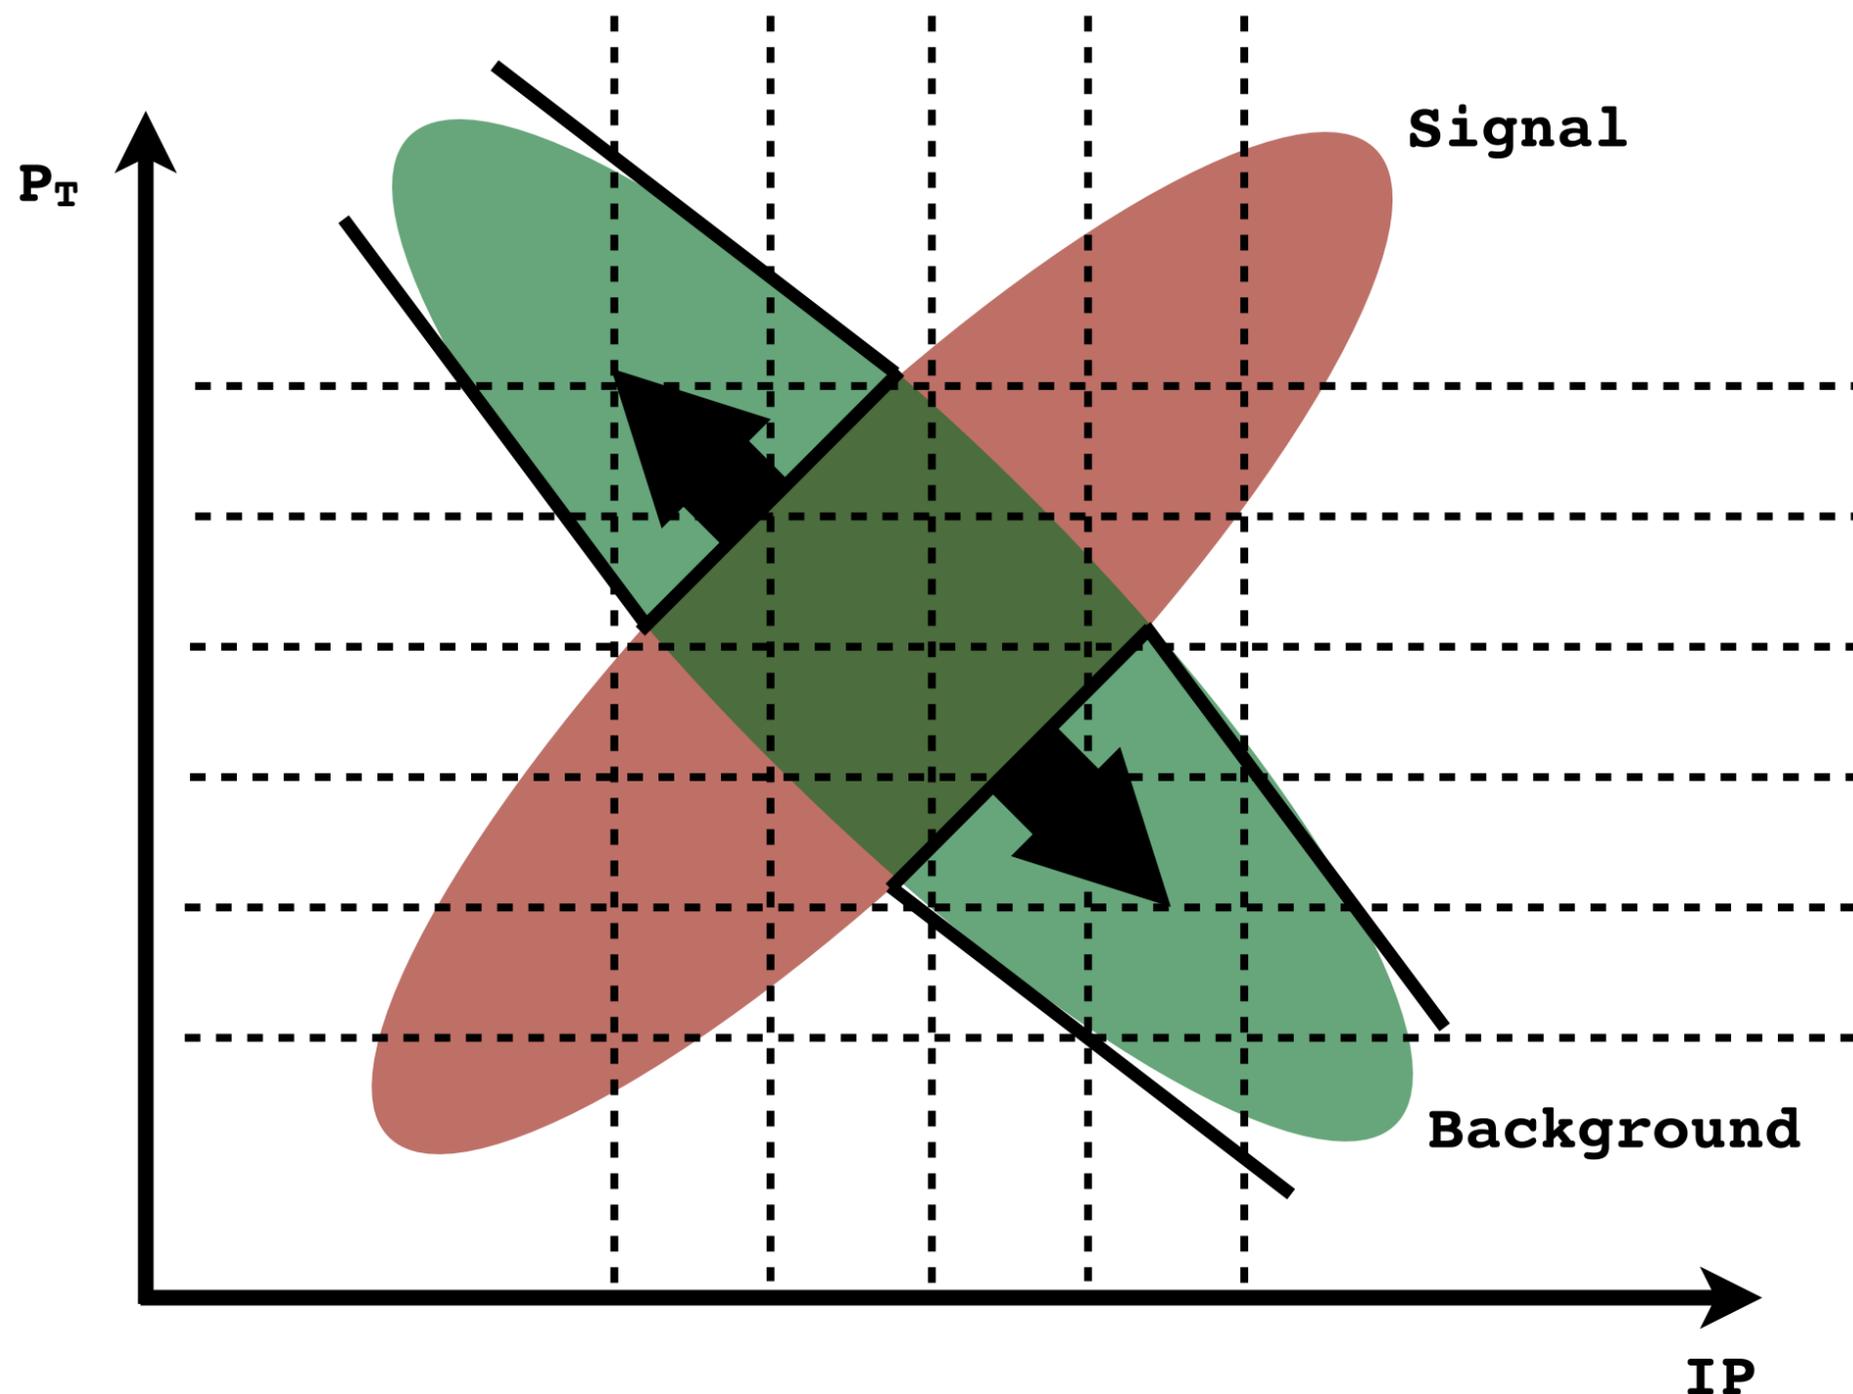
A bansai boosted decision tree

Consider a two-variable boosted decision tree : this is like a binned selection where the BDT algorithm picks the optimal bin sizes and boundaries

THEREFORE : discretize the variables yourself!

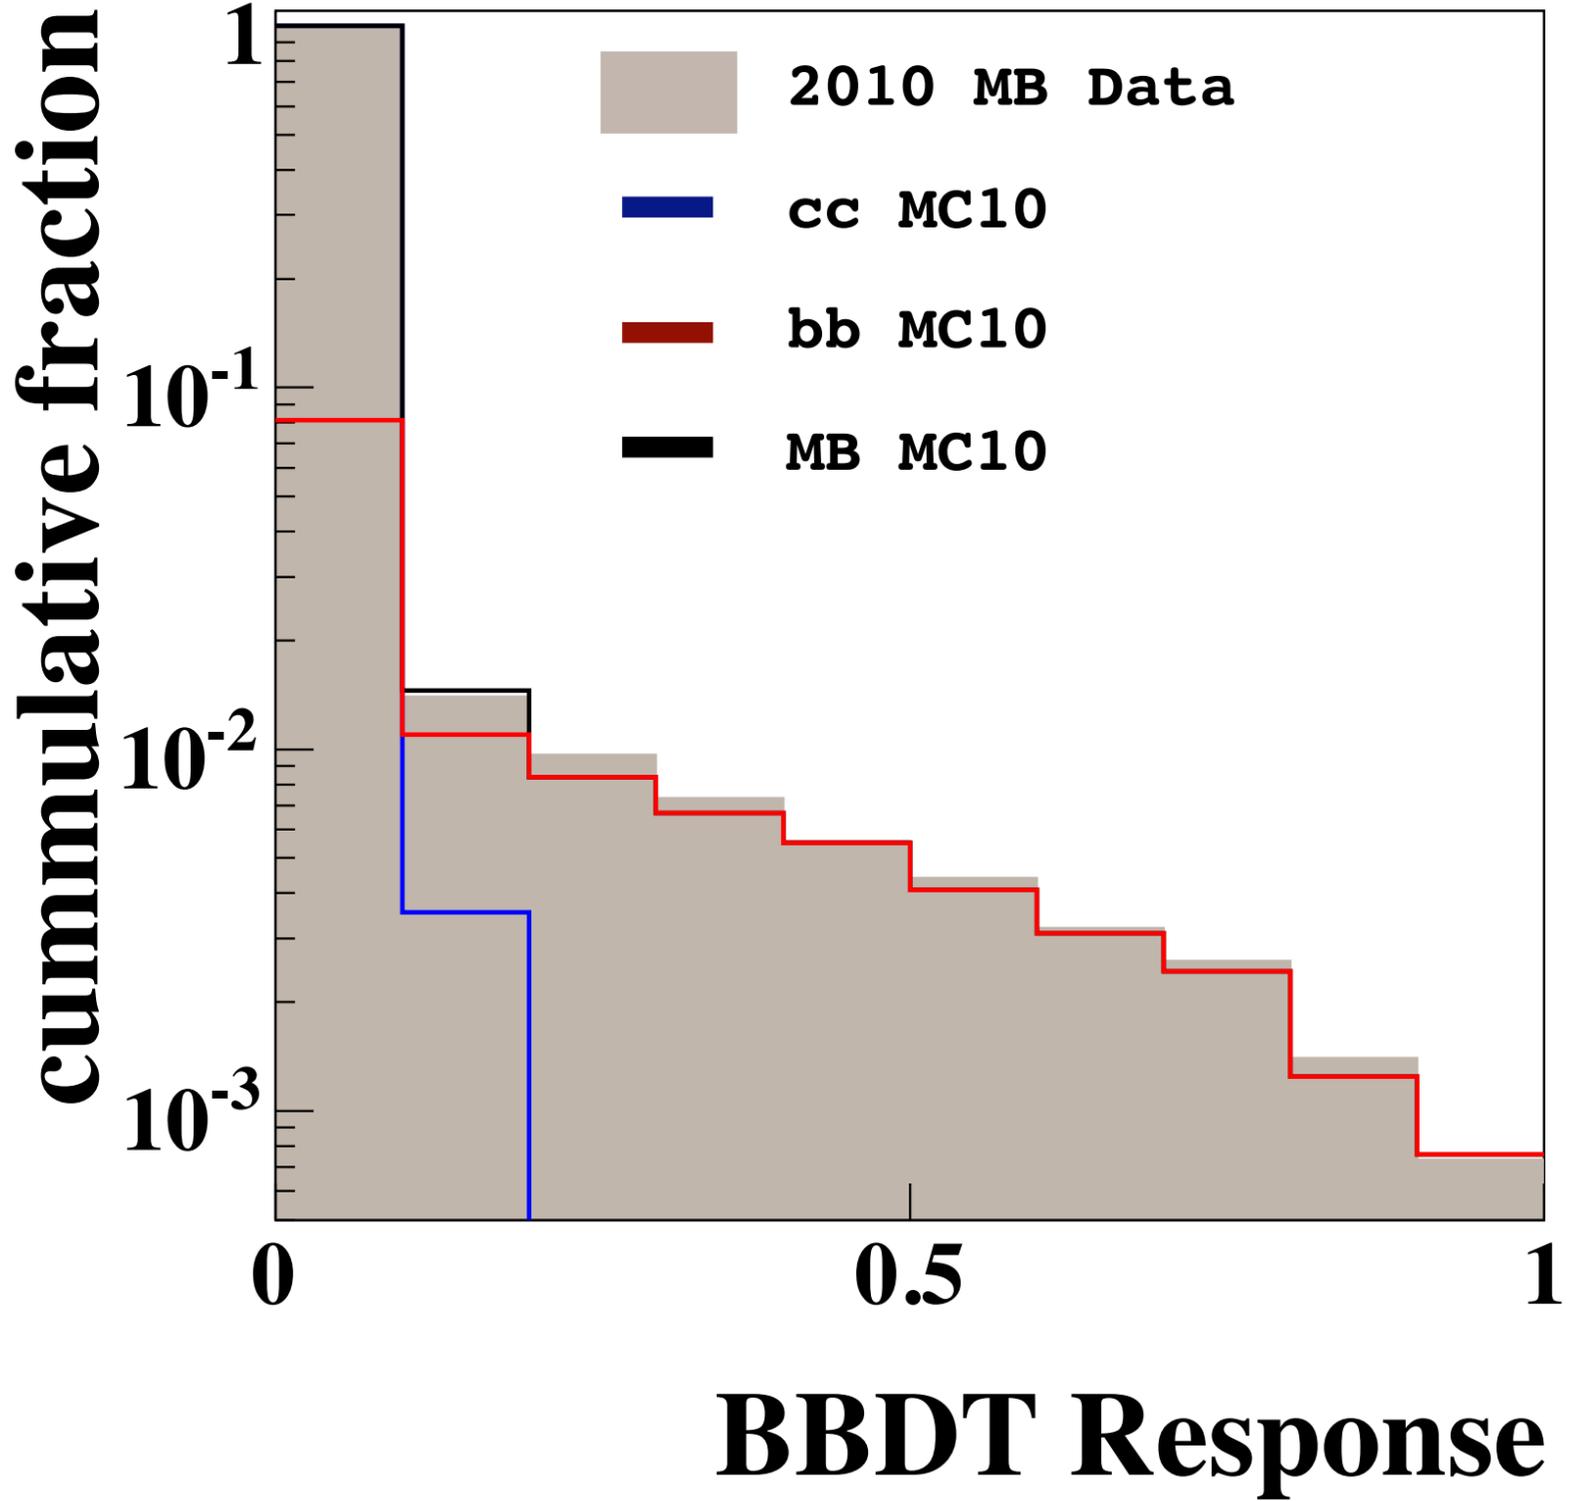
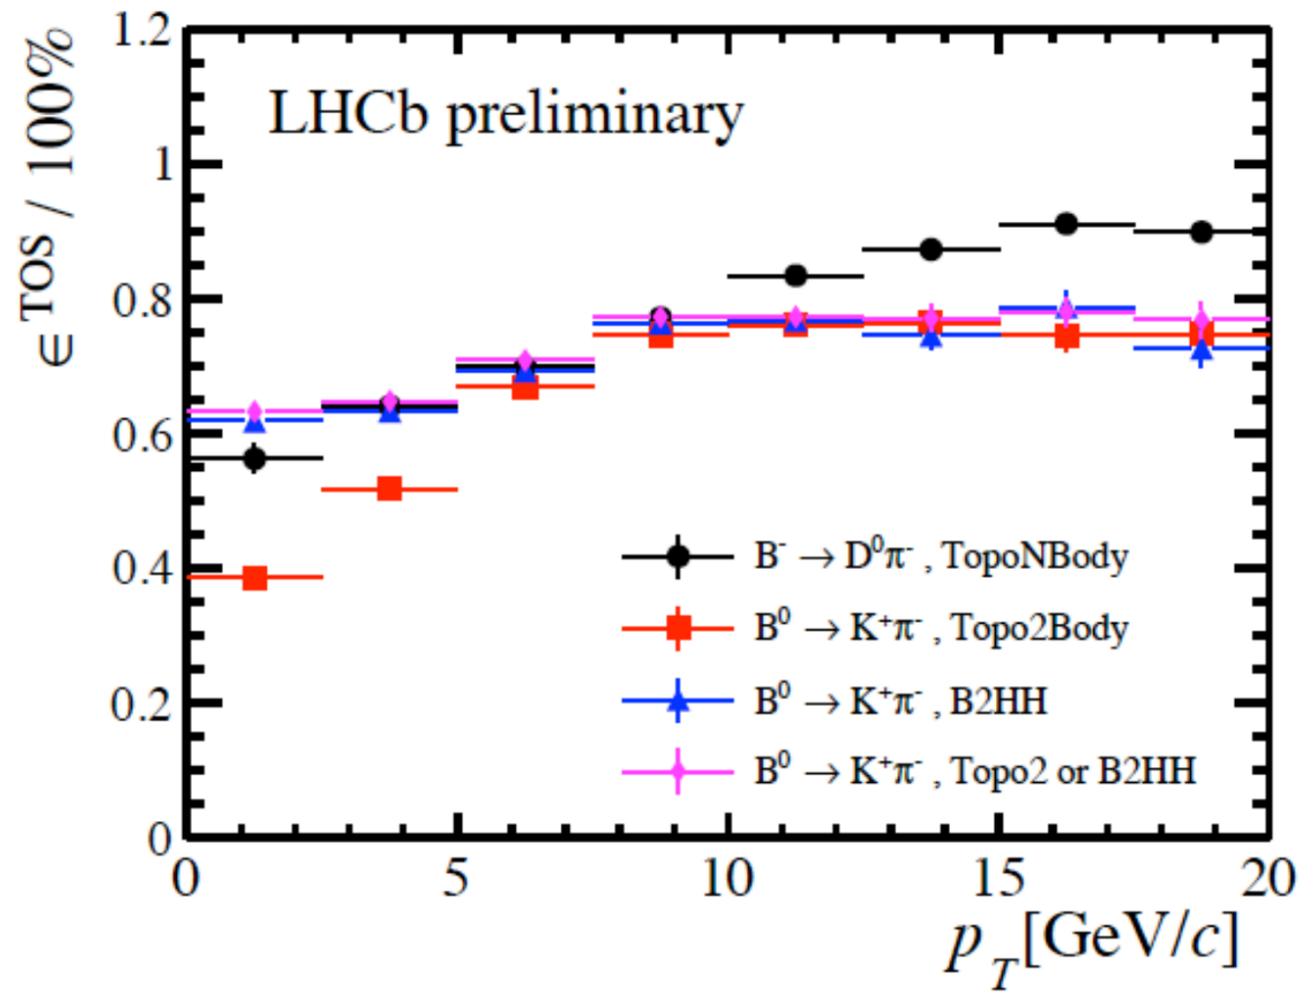
=> Makes sure that the trigger is insensitive to resolution fluctuations.

=> Transforms the trigger into a 1D lookup table making it essentially infinitely fast.



Topological performance

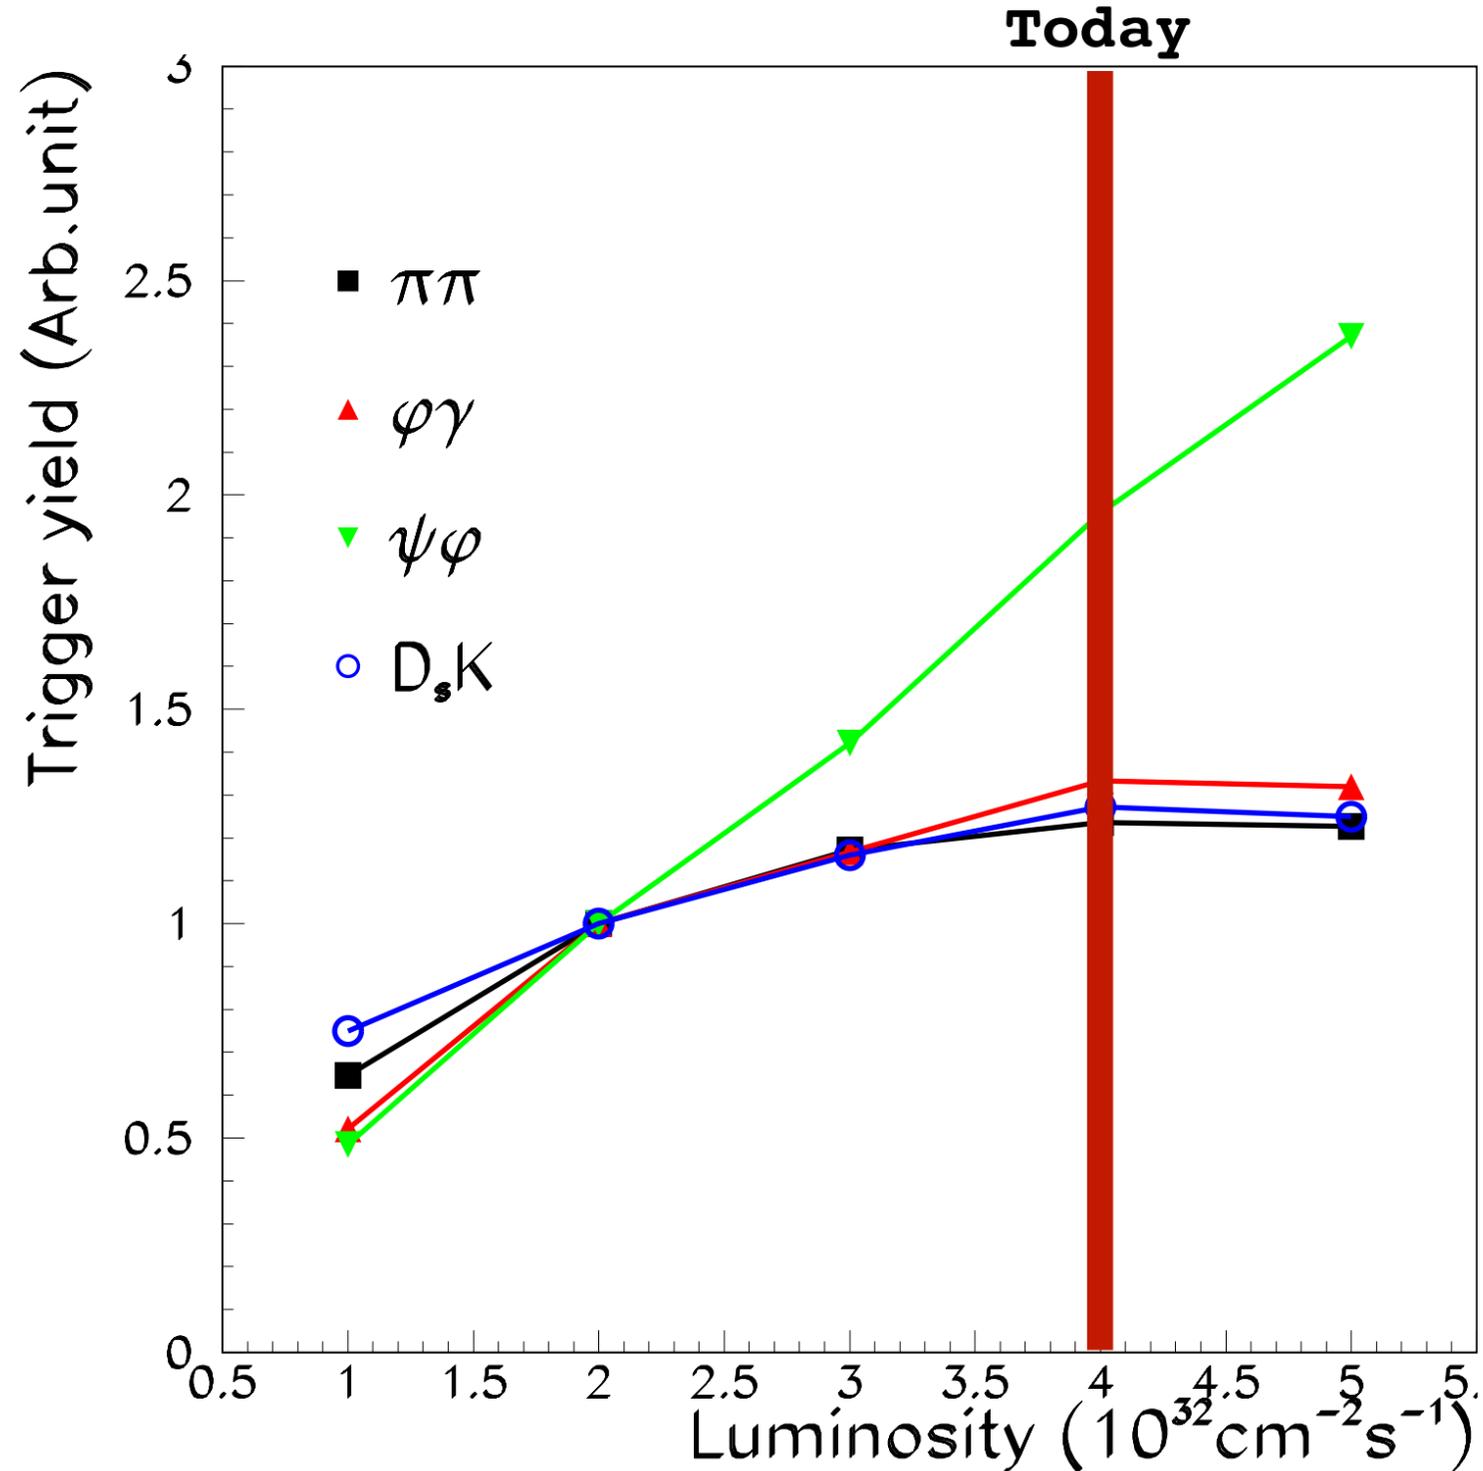
Measured output is almost 100% consistent with $b\bar{b}$ events!



See also LHCb public notes and trigger publications
LHCb-PUB-2011-002,003,016
<http://arxiv.org/abs/1310.8544>
<http://arxiv.org/abs/1211.3055>

The LHCb upgrade and trigger challenges

Why upgrade?



Only being able to read out the full detector at 1 MHz severely limits the event yields for hadronic modes

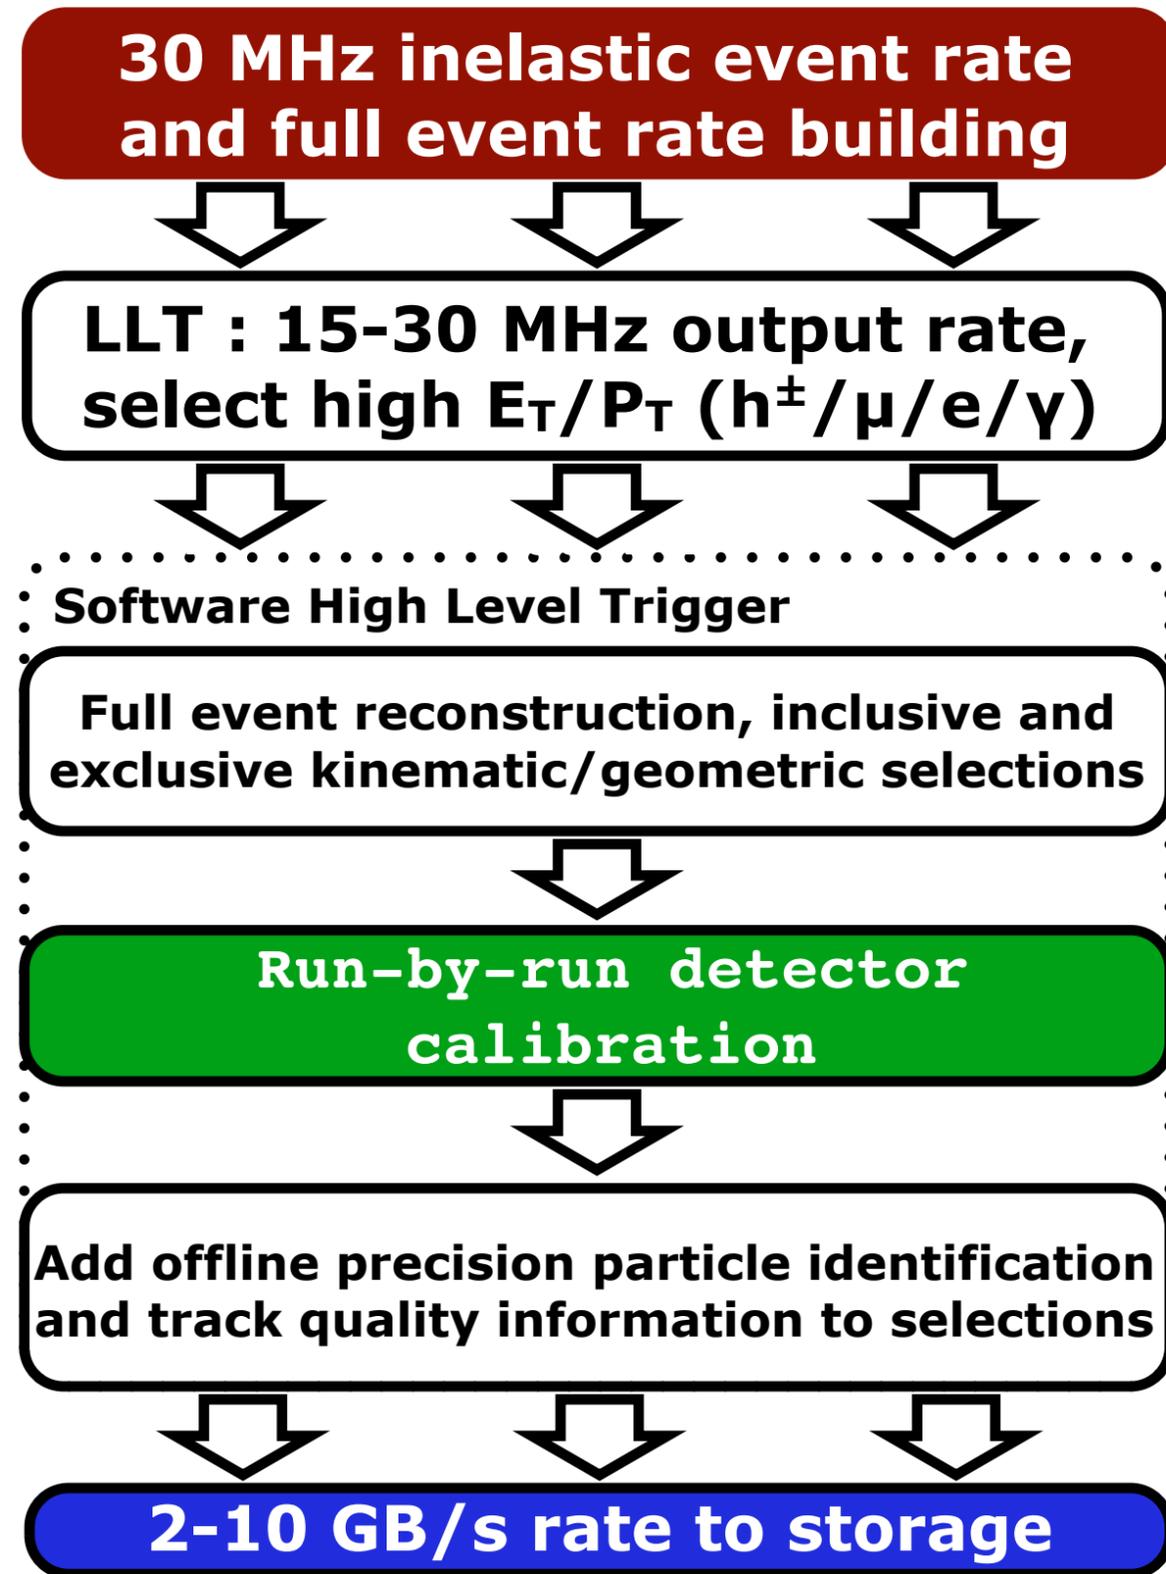
To run at higher luminosity we must remove this bottleneck

=> Full 40 MHz detector readout

=> All software trigger

=> Allows to run at a 5 times bigger proton-proton interaction rate compared to 2010-2012

Upgrade trigger layout



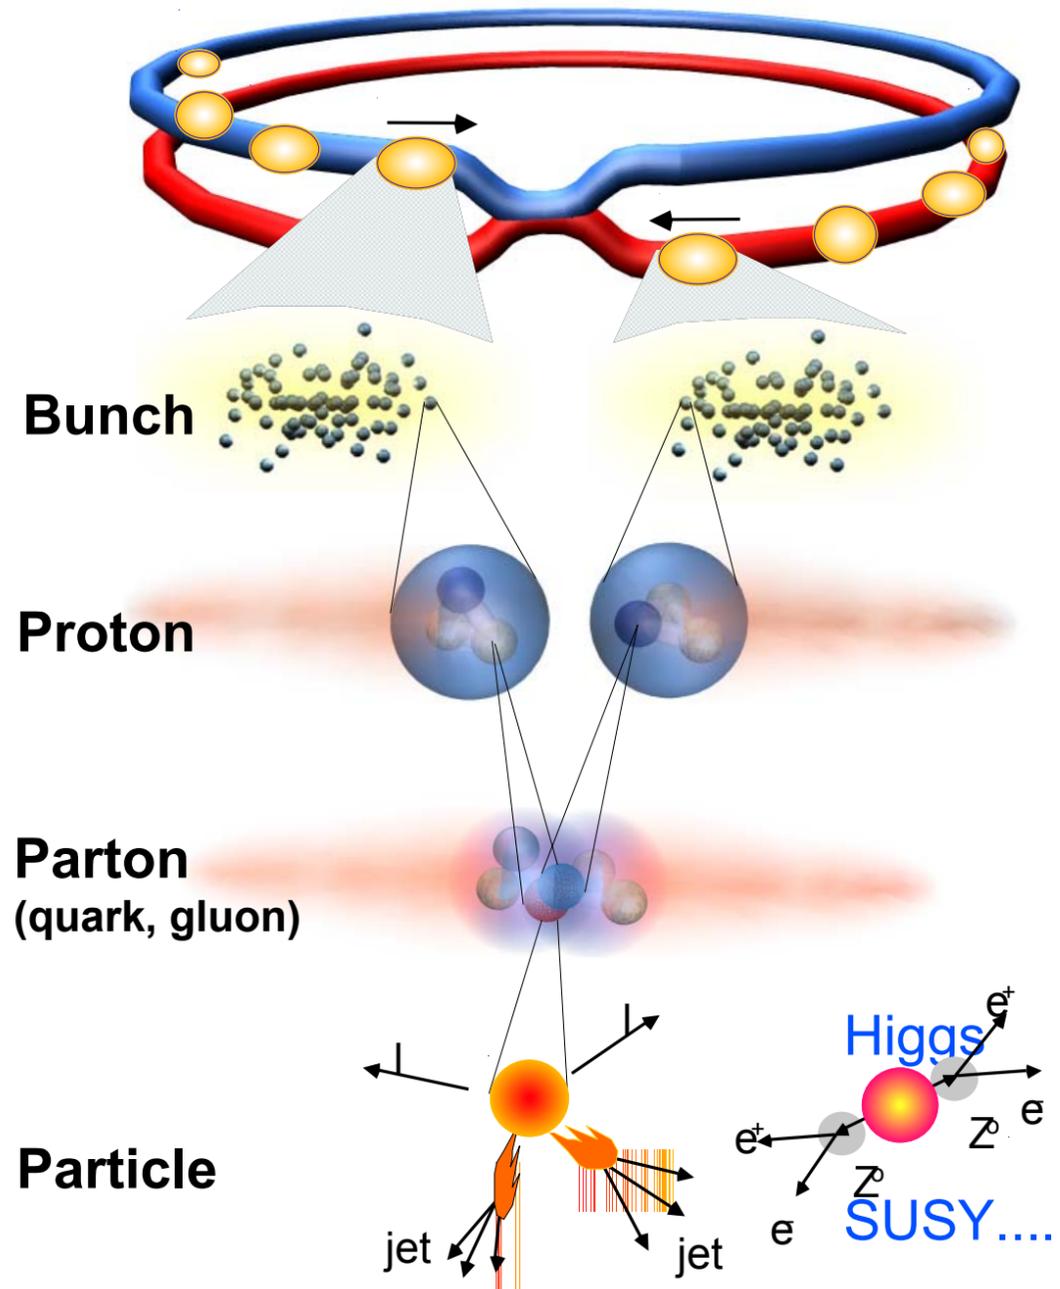
Full burden of rate reduction is now on the software trigger

Rate now given in GB/s to reflect the fact that we want to maximize event rates by storing only partial information about some events, but individual triggers will need to achieve 1/10000-1/1000 reduction

The traditional view of a trigger's job



Collisions at the LHC: summary



Proton - Proton 2804 bunch/beam
Protons/bunch 10^{11}
Beam energy 7 TeV (7×10^{12} eV)
Luminosity $10^{34} \text{cm}^{-2} \text{s}^{-1}$

Crossing rate 40 MHz

Collision rate $\approx 10^7 - 10^9$

New physics rate $\approx .00001$ Hz

Event selection:
1 in 10,000,000,000,000

But their background is our signal...

Question : what fraction of LHCb upgrade events contain a beauty or charm hadron with substantial transverse momentum and decay time which can be partially reconstructed in the detector?

But their background is our signal...

Question : what fraction of LHCb upgrade events contain a beauty or charm hadron with substantial transverse momentum and decay time which can be partially reconstructed in the detector?

Answer : ~3% (!!)

What is the challenge?

The traditional trigger finds rare signals with maximum efficiency

The LHCb upgrade trigger needs to classify different kinds of signal

Real time event selection

==

Signal classification

This is a new kind of problem

Traditionally we use inclusive triggers which partially reconstruct signals in order to

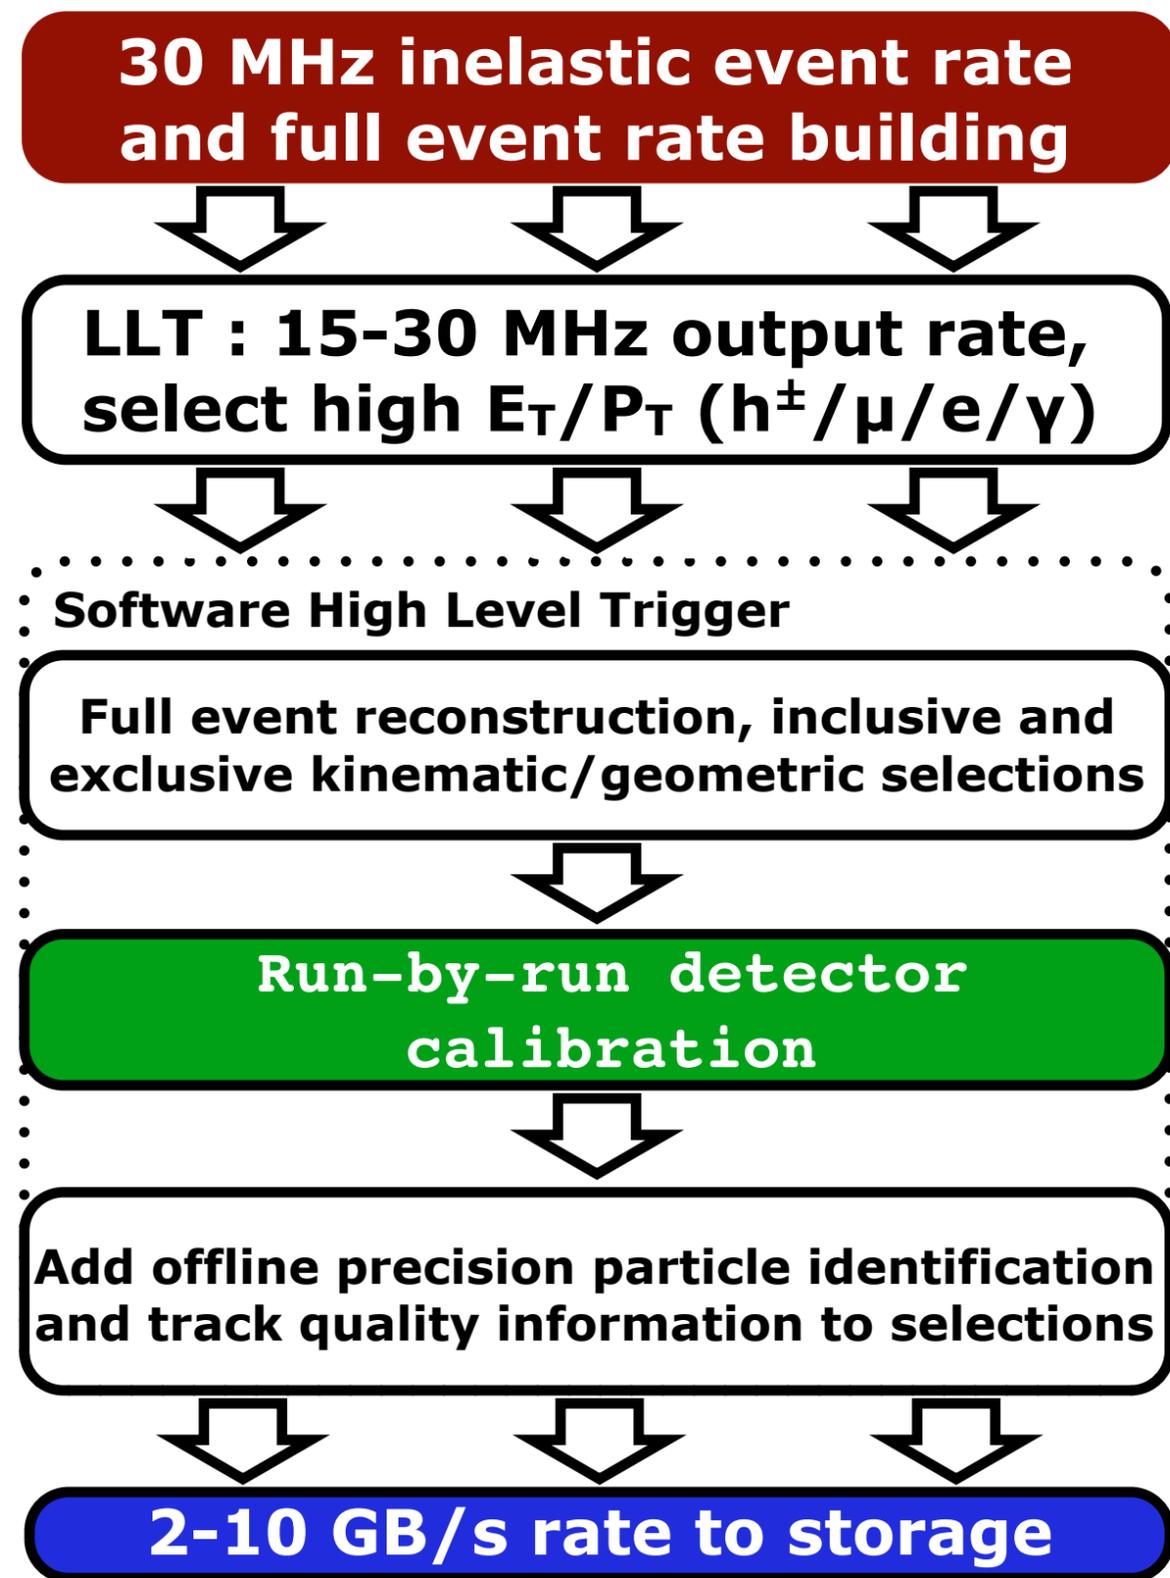
- 1) Select a wide range of signals with a single trigger
- 2) Remain robust against changing detector conditions

By definition this approach will not be optimal for discriminating between different signals

On the other hand, fully reconstructing different signals poses other difficulties

- 1) Many more triggers to maintain
- 2) Each trigger costs time to execute, which adds up
- 3) Need to have a real-time detector calibration and alignment

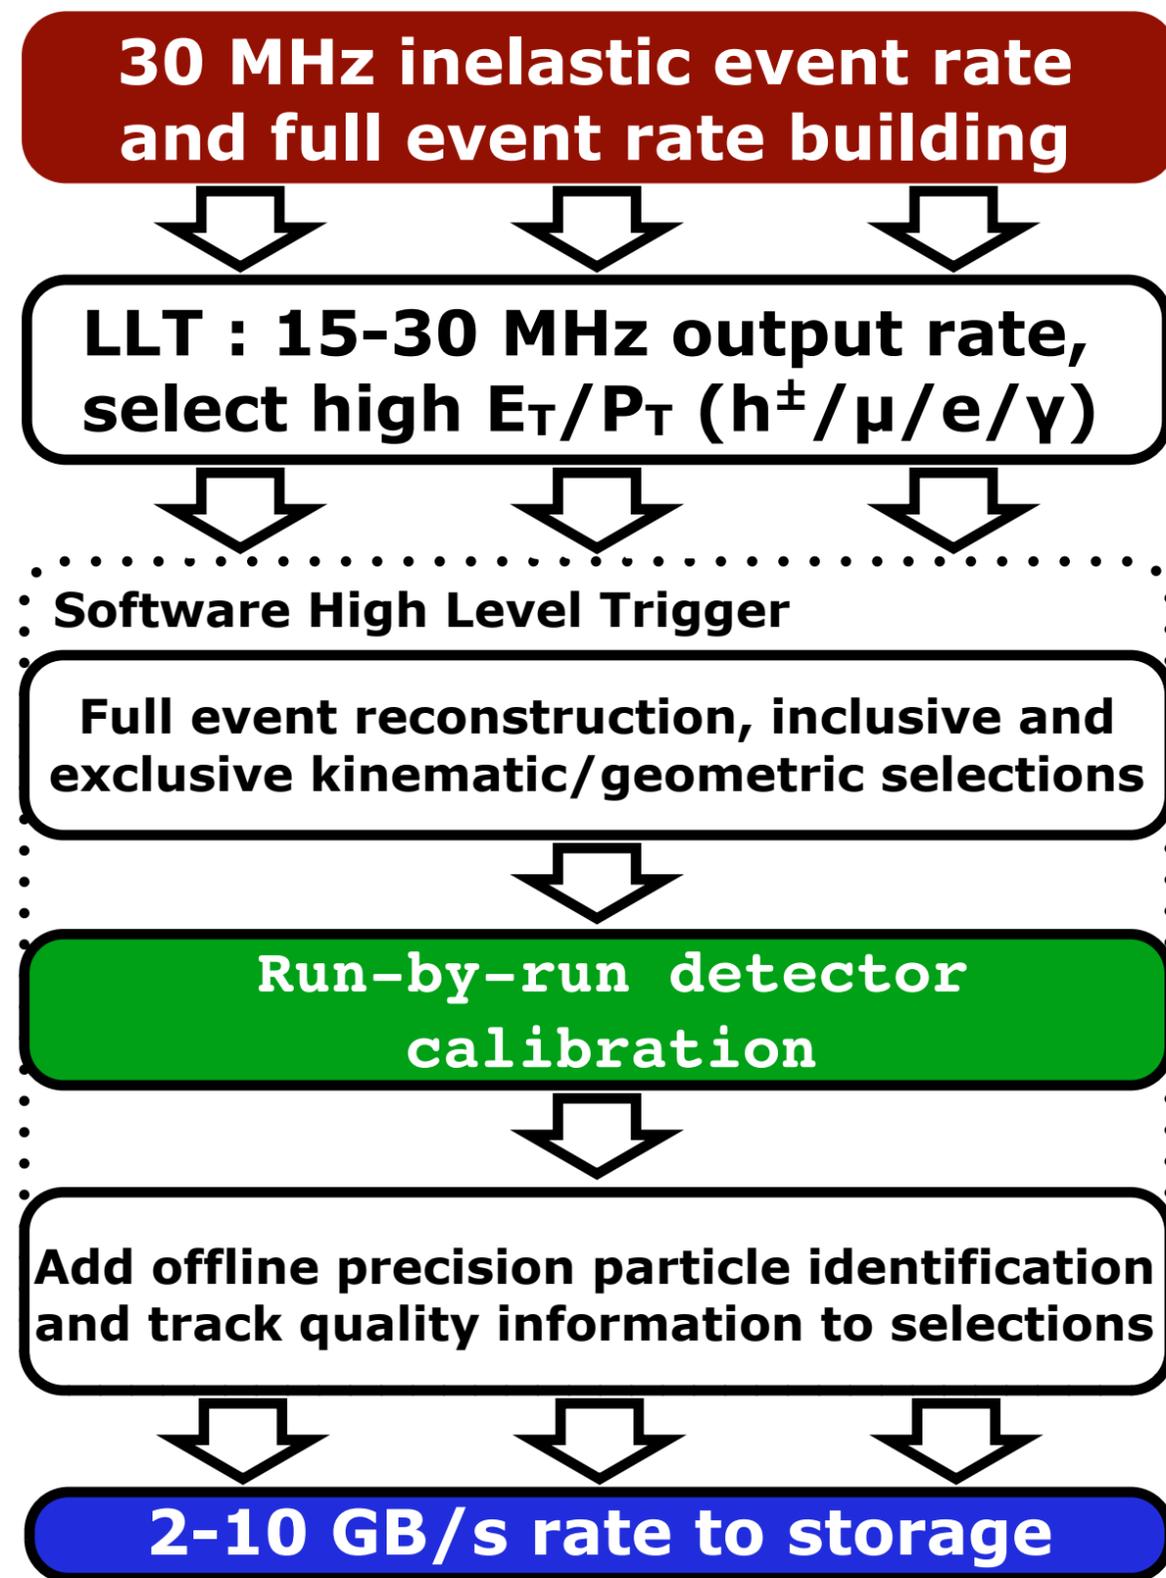
But it is also an opportunity



The run-by-run detector calibration is something which we plan to deploy already in 2015

For the selections, can we be smarter than having a long list of exclusive triggers?

But it is also an opportunity

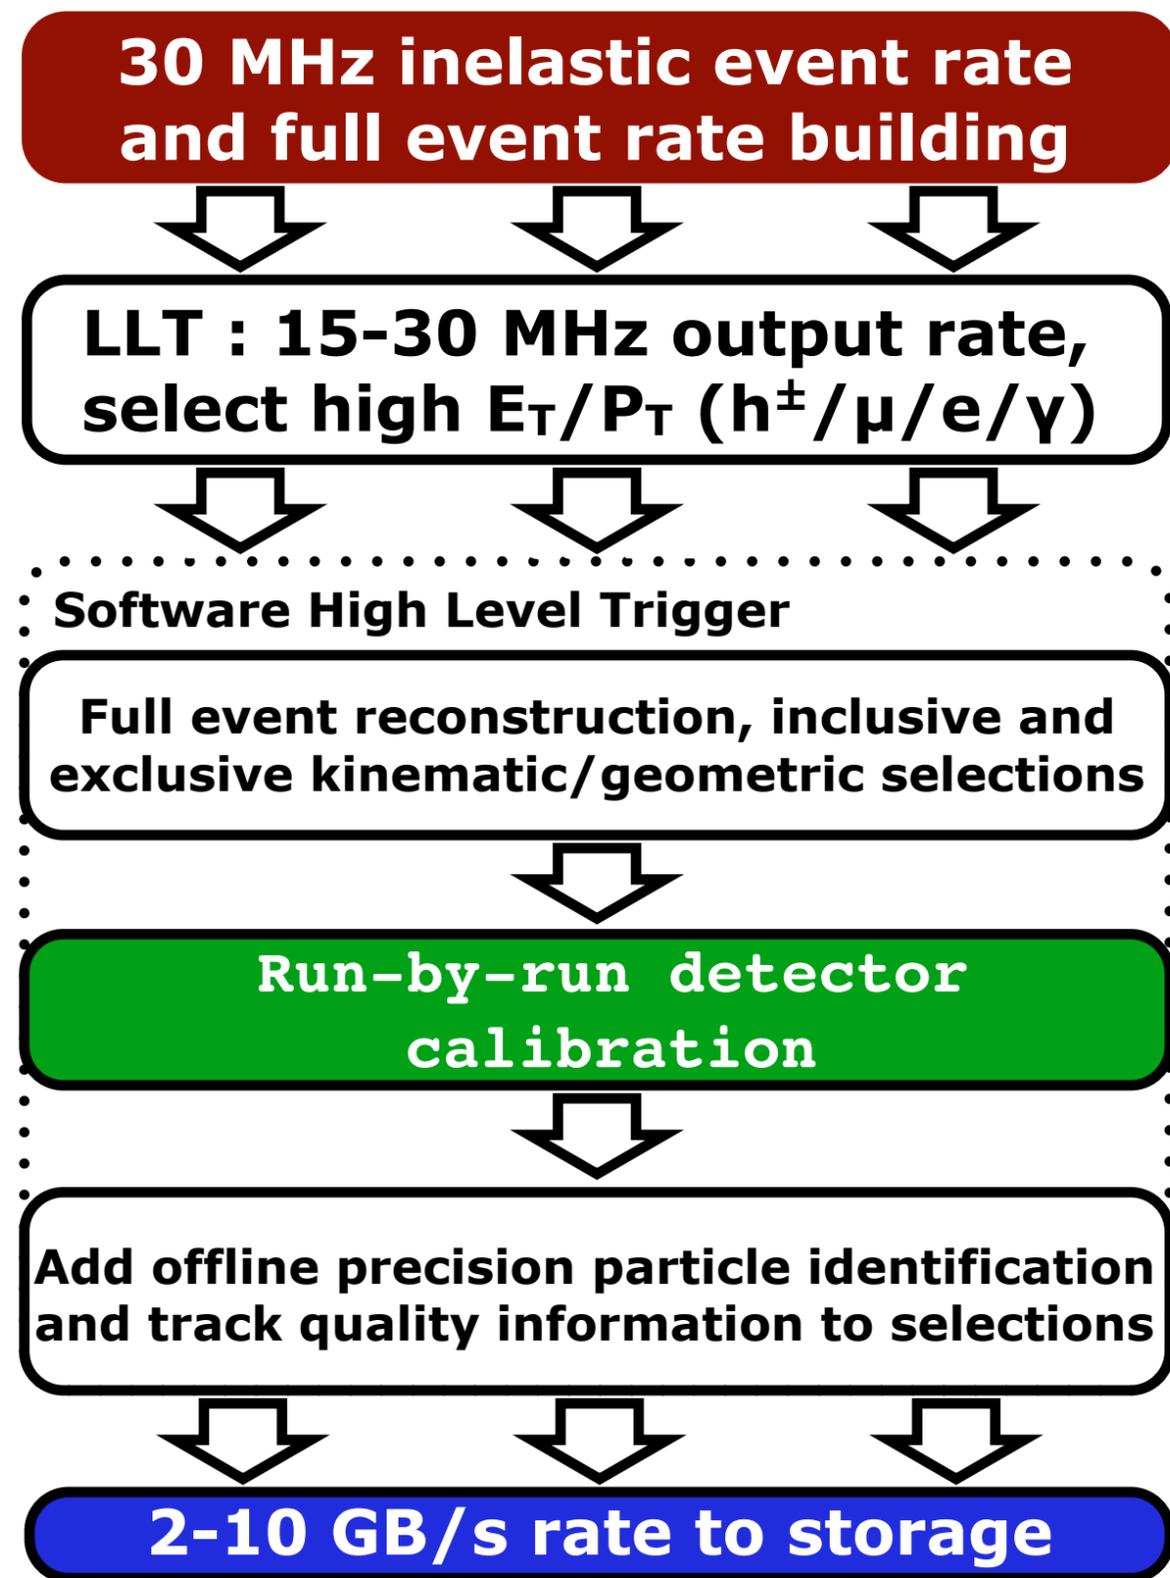


The run-by-run detector calibration is something which we plan to deploy already in 2015

For the selections, can we be smarter than having a long list of exclusive triggers?

E.g. by deploying multivariate multi-class algorithms to assign each event a likelihood of being a certain kind of signal?

But it is also an opportunity



The run-by-run detector calibration is something which we plan to deploy already in 2015

For the selections, can we be smarter than having a long list of exclusive triggers?

E.g. by deploying multivariate multi-class algorithms to assign each event a likelihood of being a certain kind of signal?

E.g. by having budgeted classification algorithms which optimize the time-cost of the information used in the classifier?

Conclusion

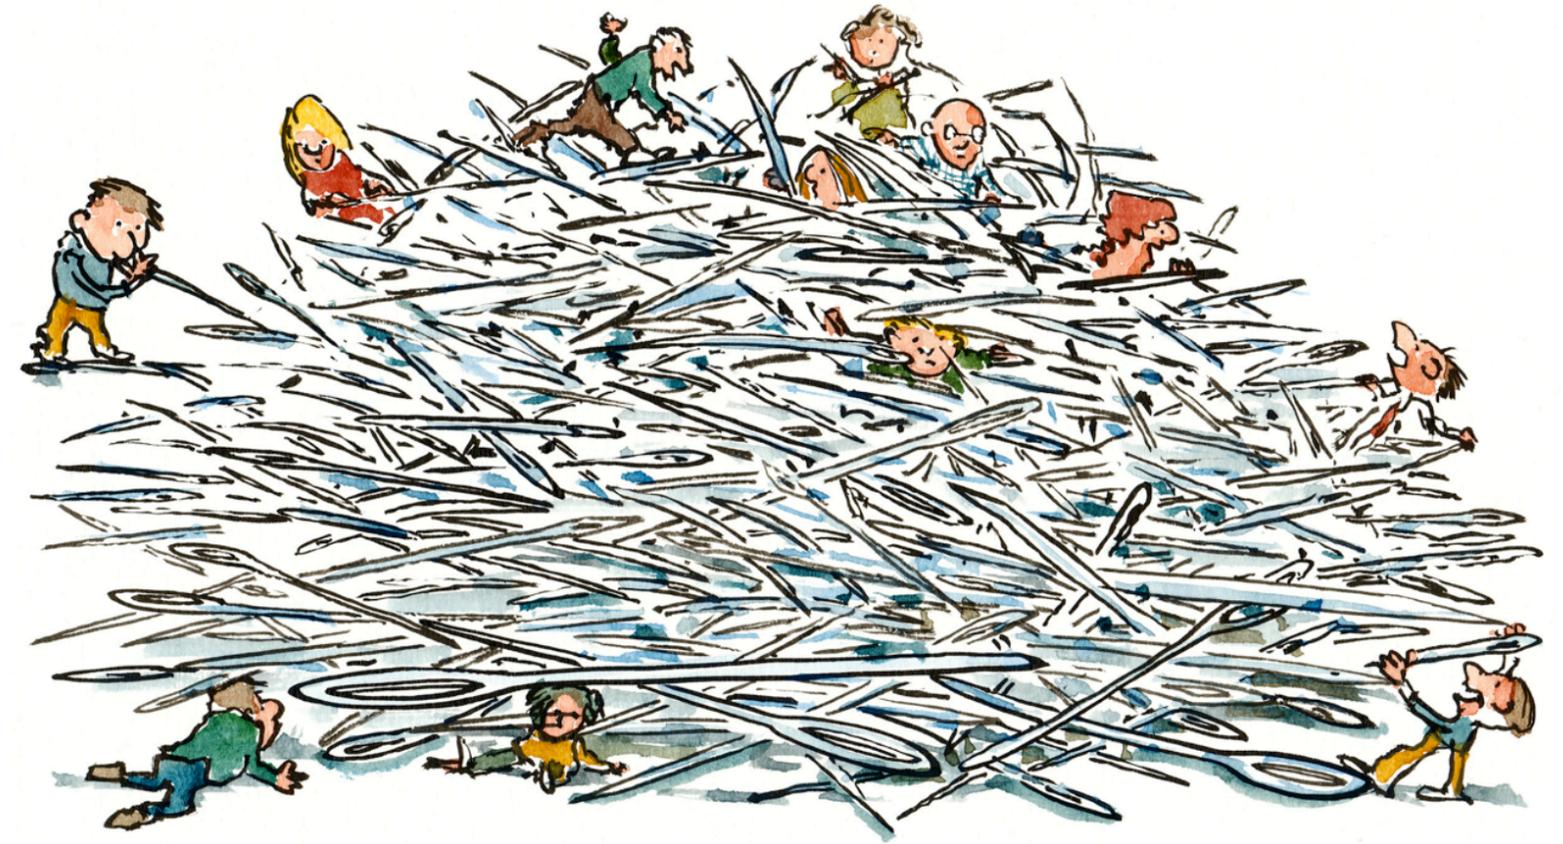


**Triggers
today**

Conclusion



**Triggers
today**



**Triggers
in the future**

Conclusion



The LHCb upgrade trigger is a signal classification problem

We could solve it with many separate signal selections...

...but do multi-class (budgeted) classifiers offer a smarter and more inclusive way forward?



www.jolyon.co.uk

**Triggers
today**



**Triggers
in the future**

Summary

What is a trigger?

trig·ger [↵] (trĭg'ər)
n.

1.
 - a. The lever pressed by the finger to discharge a firearm.
 - b. A similar device used to release or activate a mechanism.
2. An event that precipitates other events.
3. *Electronics* A pulse or circuit that initiates the action of another component.

tr.v. **trig·gered, trig·ger·ing, trig·gers**

1. To set off; initiate: *remarks that triggered bitter debates.*
2. To fire or explode (a weapon or an explosive charge).

What is a trigger?

trig·ger [↵] (trĭg'ər)
n.

1.

a. The lever pressed by the finger to discharge a firearm.

b. A similar device used to release or activate a mechanism.

2. An event that precipitates other events.

3. *Electronics* A pulse or circuit that initiates the action of another component.

tr.v. **trig·gered, trig·ger·ing, trig·gers**

1. To set off; initiate: *remarks that triggered bitter debates.*

2. To fire or explode (a weapon or an explosive charge).

ATLAS/CMS vs. LHCb

	Rate of bunch crossings	Mean interactions per bunch crossing	Mean event size
ATLAS/CMS	15 MHz	~25	1500 kB
LHCb	15 MHz	1.5	50 kB

The data rates at ATLAS and CMS are 30 times greater than at LHCb. This drives a design in which much more work is done by hardware triggers which make their decisions based on information from only a part of the detector.

The confirmation strategy

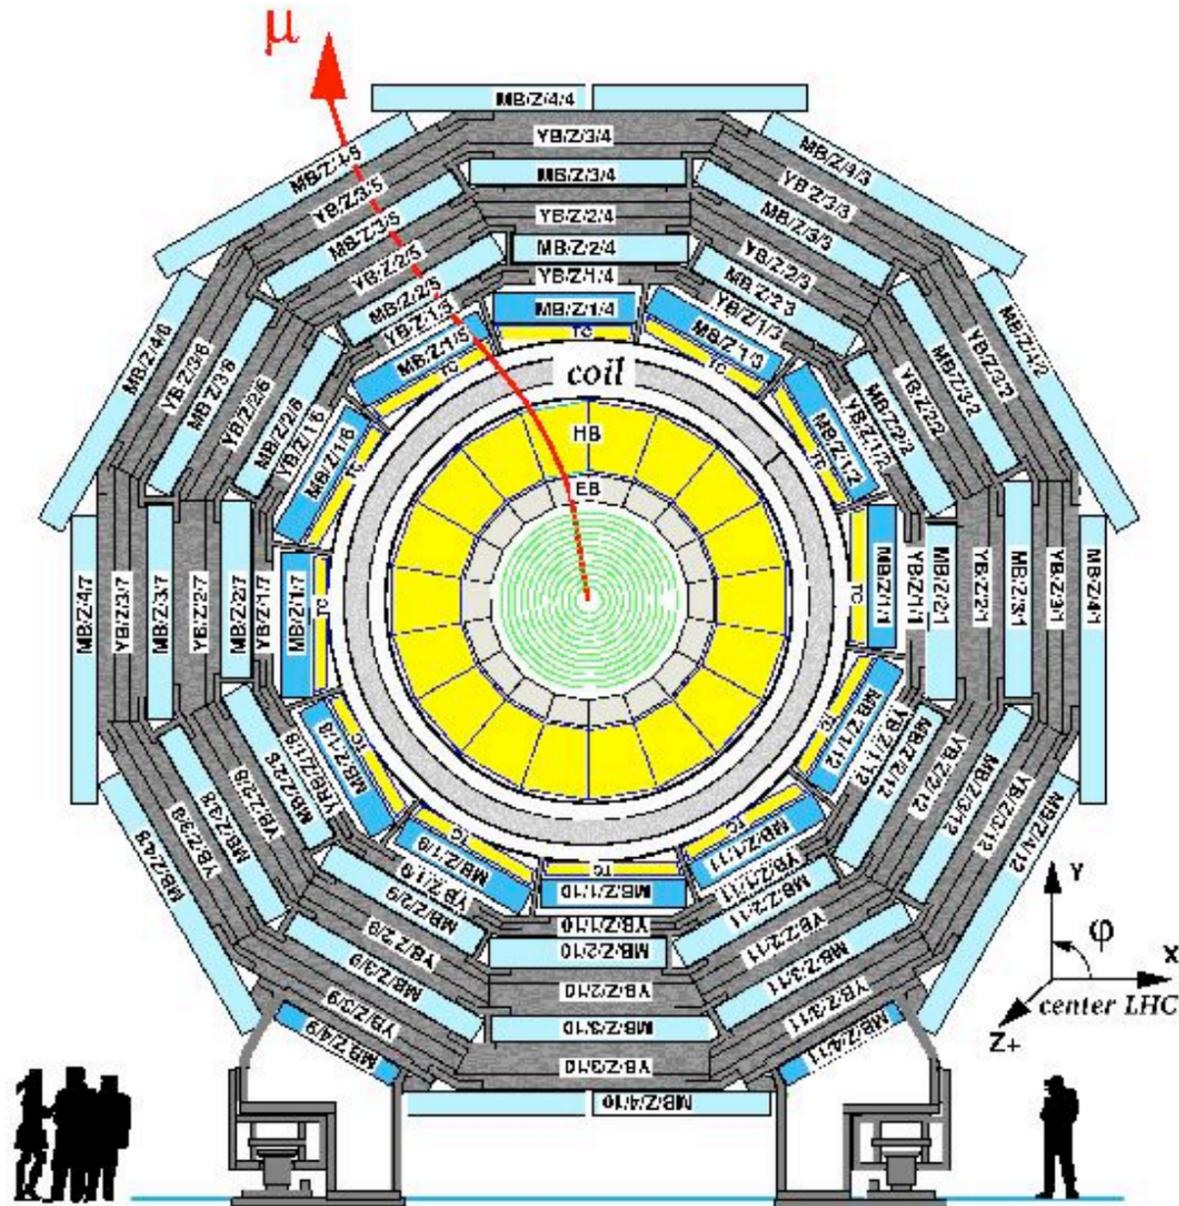
For LHCb, the High Level Trigger ignores the hardware trigger

In ATLAS/CMS, in order to speed up execution, the high level trigger is set up to "confirm" the decision of the hardware trigger

The confirmation strategy

For LHCb, the High Level Trigger ignores the hardware trigger

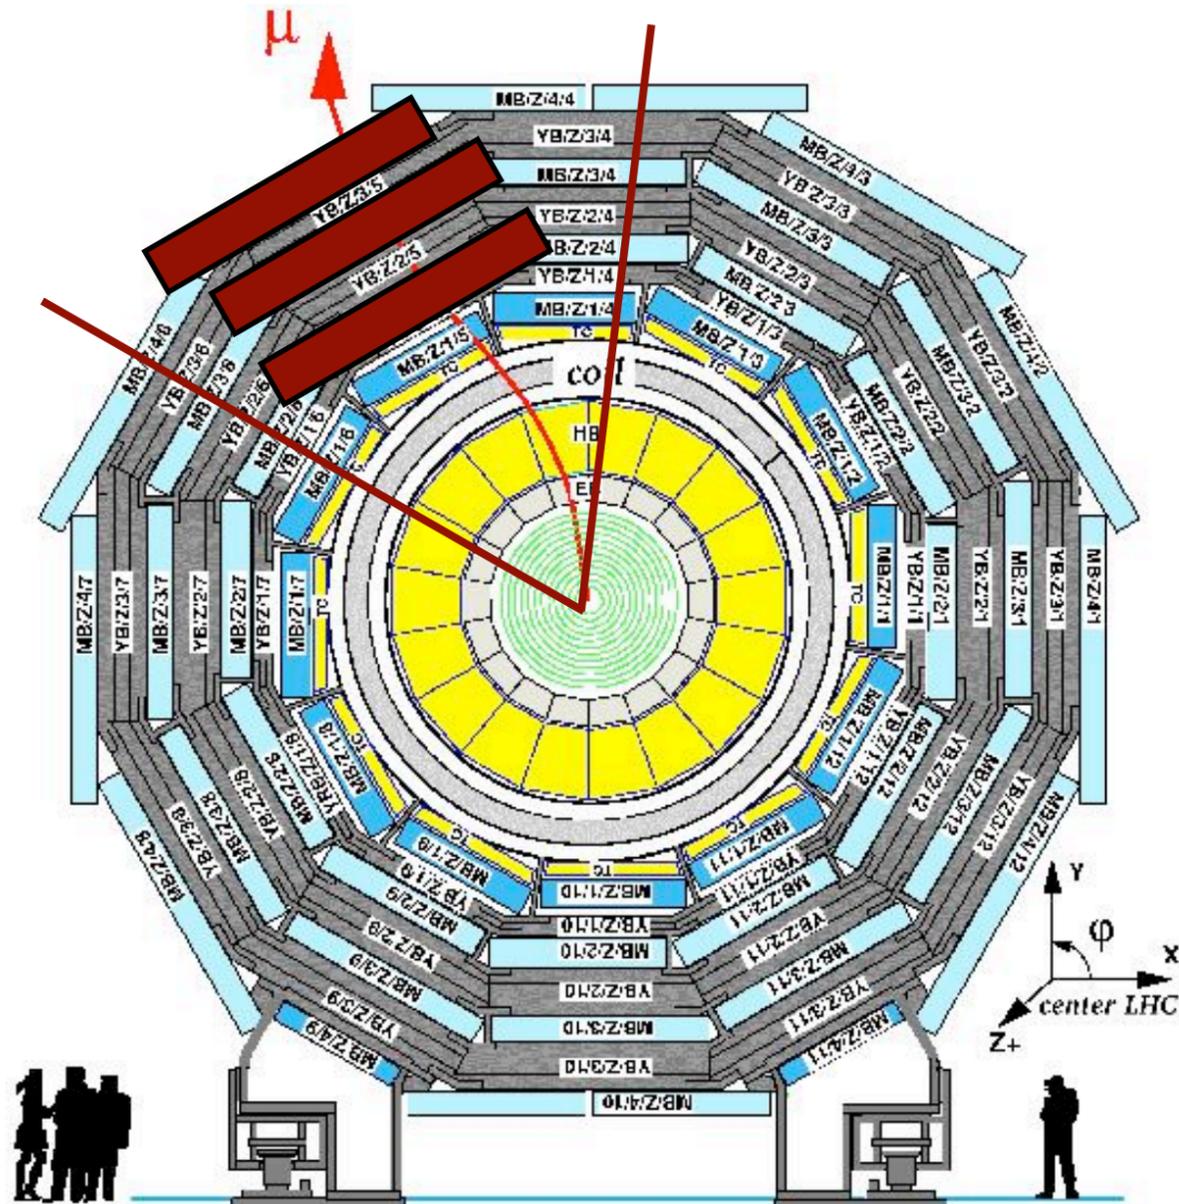
In ATLAS/CMS, in order to speed up execution, the high level trigger is set up to "confirm" the decision of the hardware trigger



The confirmation strategy

For LHCb, the High Level Trigger ignores the hardware trigger

In ATLAS/CMS, in order to speed up execution, the high level trigger is set up to "confirm" the decision of the hardware trigger



The hardware trigger has fired because of a muon identified in the muon system

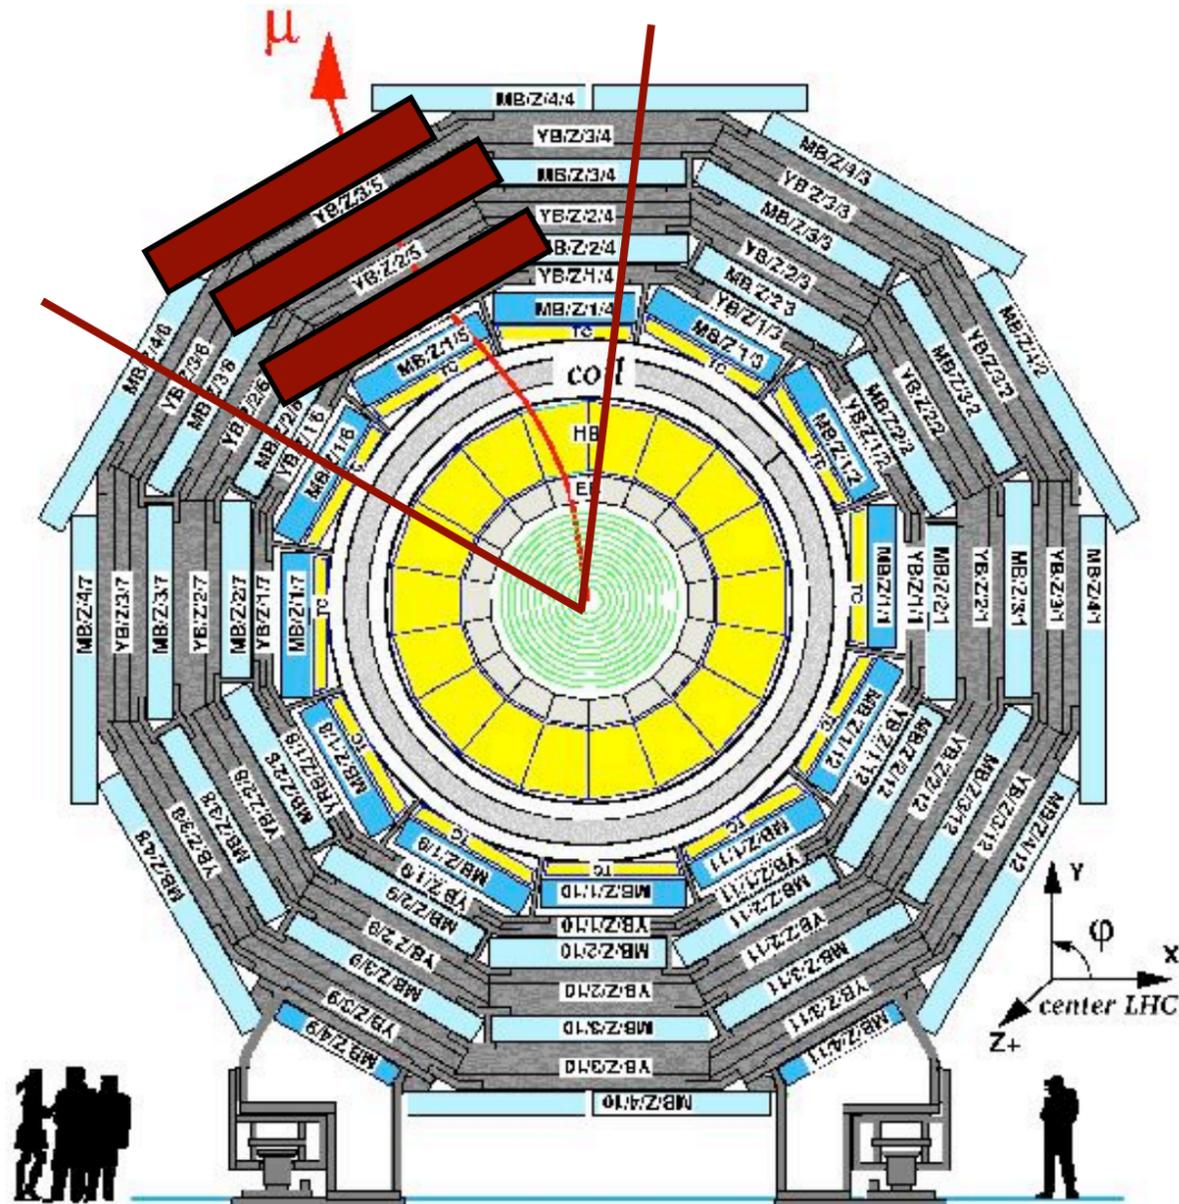
The High Level Trigger "confirms" that this is a muon by finding it in the tracking system as well

The region of interest for this search is defined by the detector geometry and the location of the hardware trigger candidate

The confirmation strategy

For LHCb, the High Level Trigger ignores the hardware trigger

In ATLAS/CMS, in order to speed up execution, the high level trigger is set up to "confirm" the decision of the hardware trigger



The hardware trigger has fired because of a muon identified in the muon system

The High Level Trigger "confirms" that this is a muon by finding it in the tracking system as well

The region of interest for this search is defined by the detector geometry and the location of the hardware trigger candidate

Confirmation is a very common strategy in triggering. It works if the early triggers fire predominantly because of the presence of signal.

If you need to build a trigger...

Two parameters to consider

- 1) The frequency with which events occur**
- 2) The complexity (size) of each event**

Key concepts to remember from this seminar

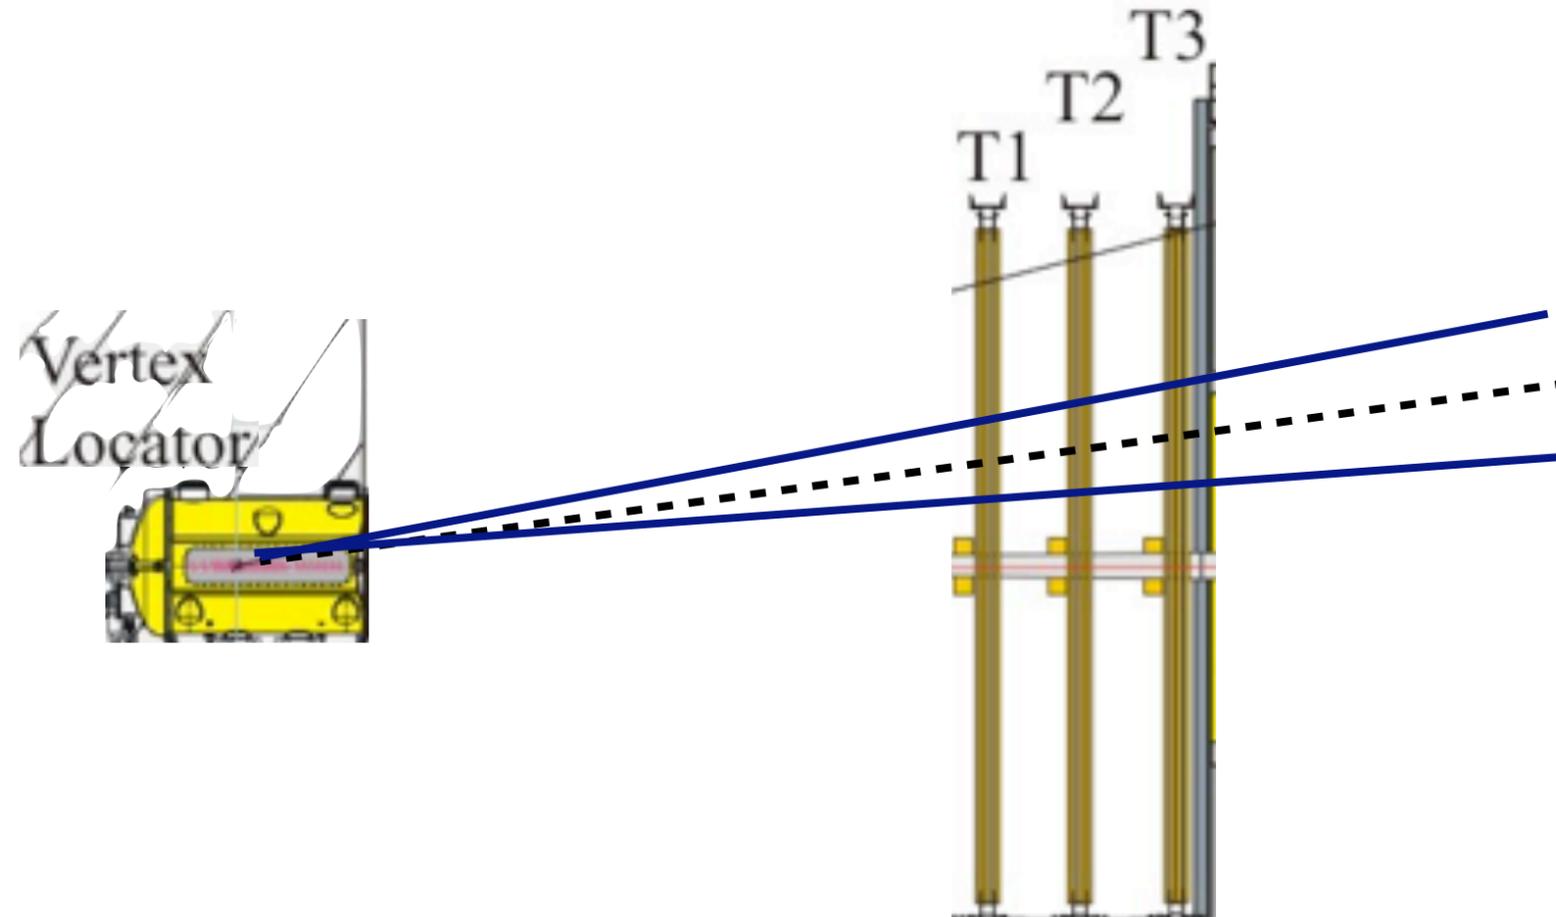
Different triggers make their selection on different criteria, but they must always be independent of each other.

A trigger has a finite time to make its selection, so you need to optimize taking into account the time cost of obtaining information

Multivariate selections are very powerful but usually need a simpler preselection to allow them the time to do their job

Backup

Merging reconstruction and selection



A lot of time is spent in the trigger reconstructing charged particles

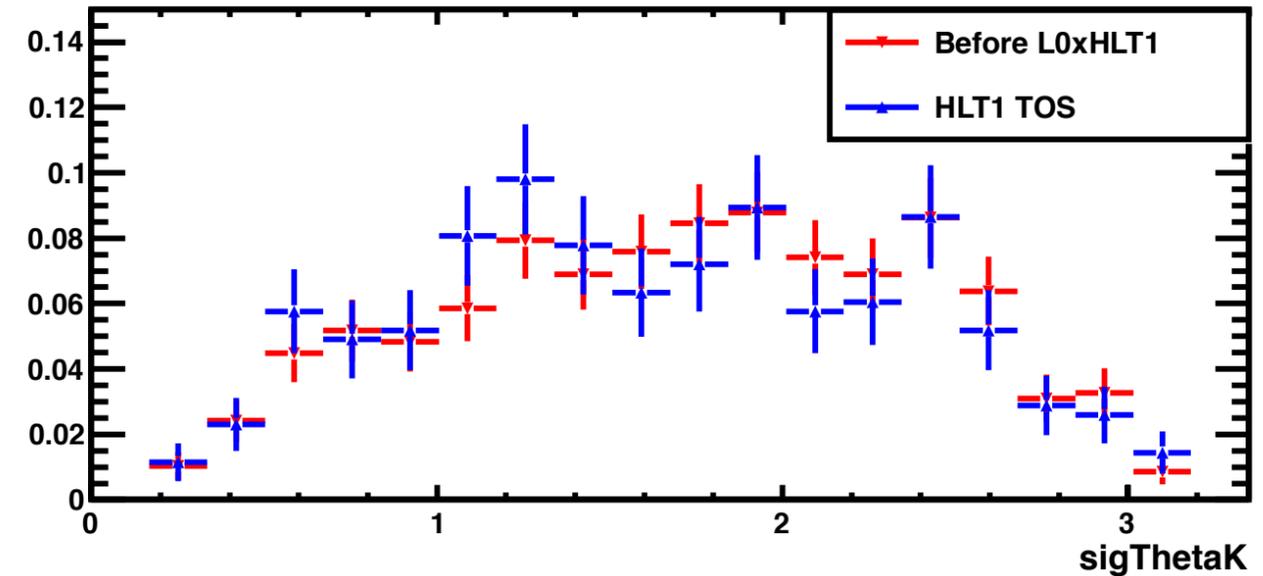
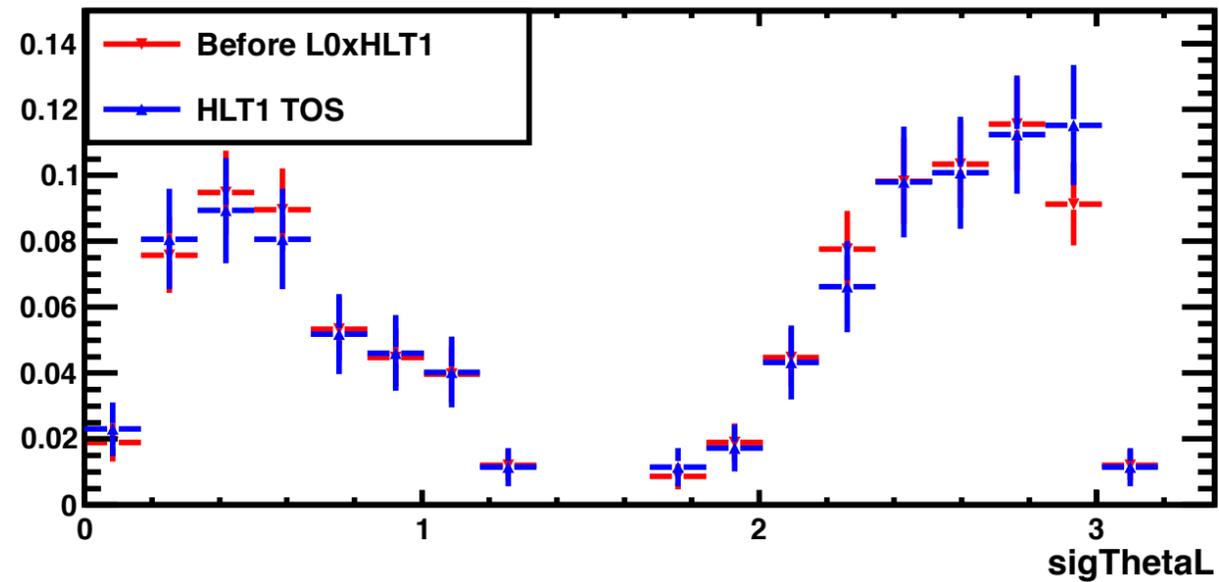
If you know that you will cut on some minimum momentum of these particles, you can build this cut into the reconstruction

The higher the momentum the straighter the charged particle path

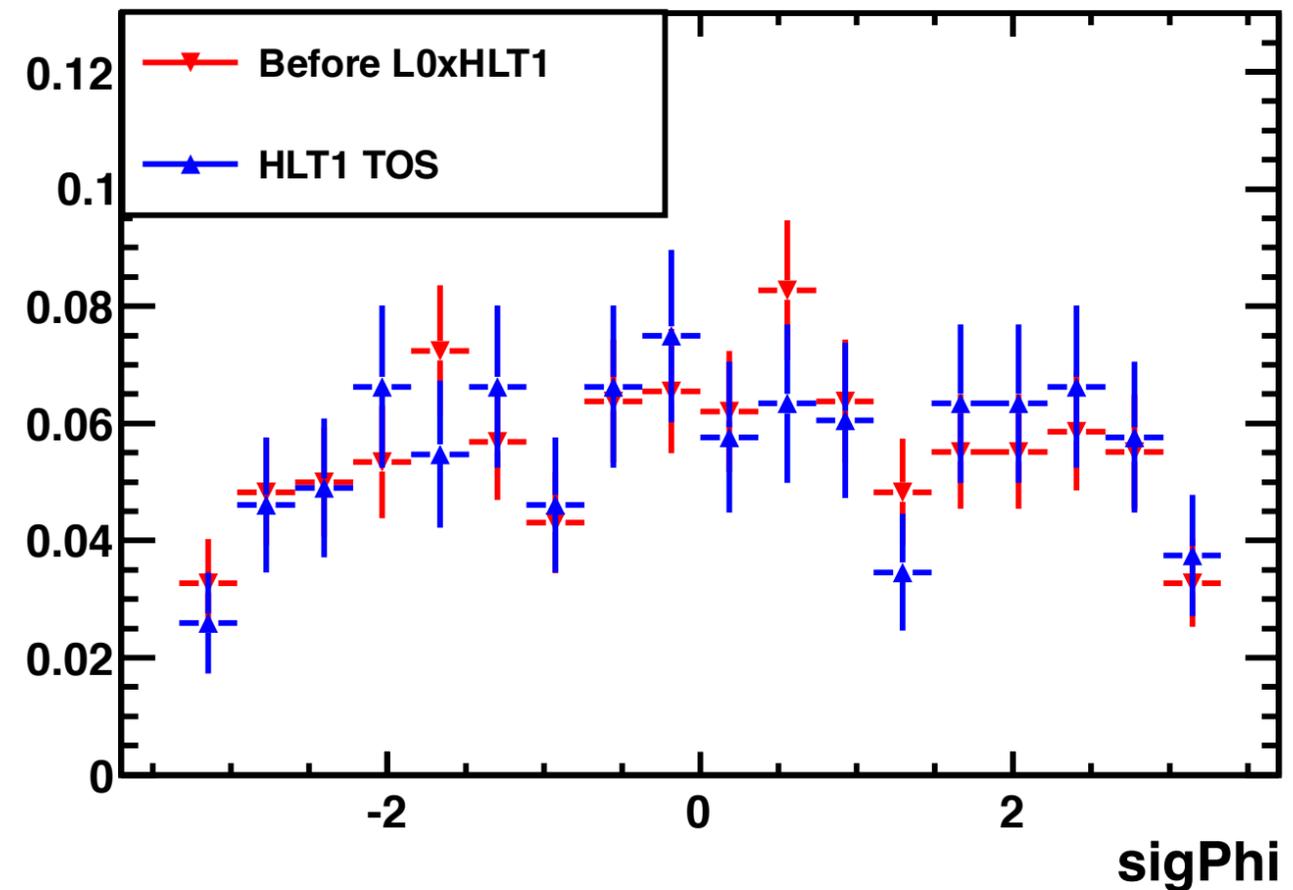
Can define a narrow path depending on momentum; saves a lot of time looking for fake paths

Always look for ways to build a selection into your reconstruction!

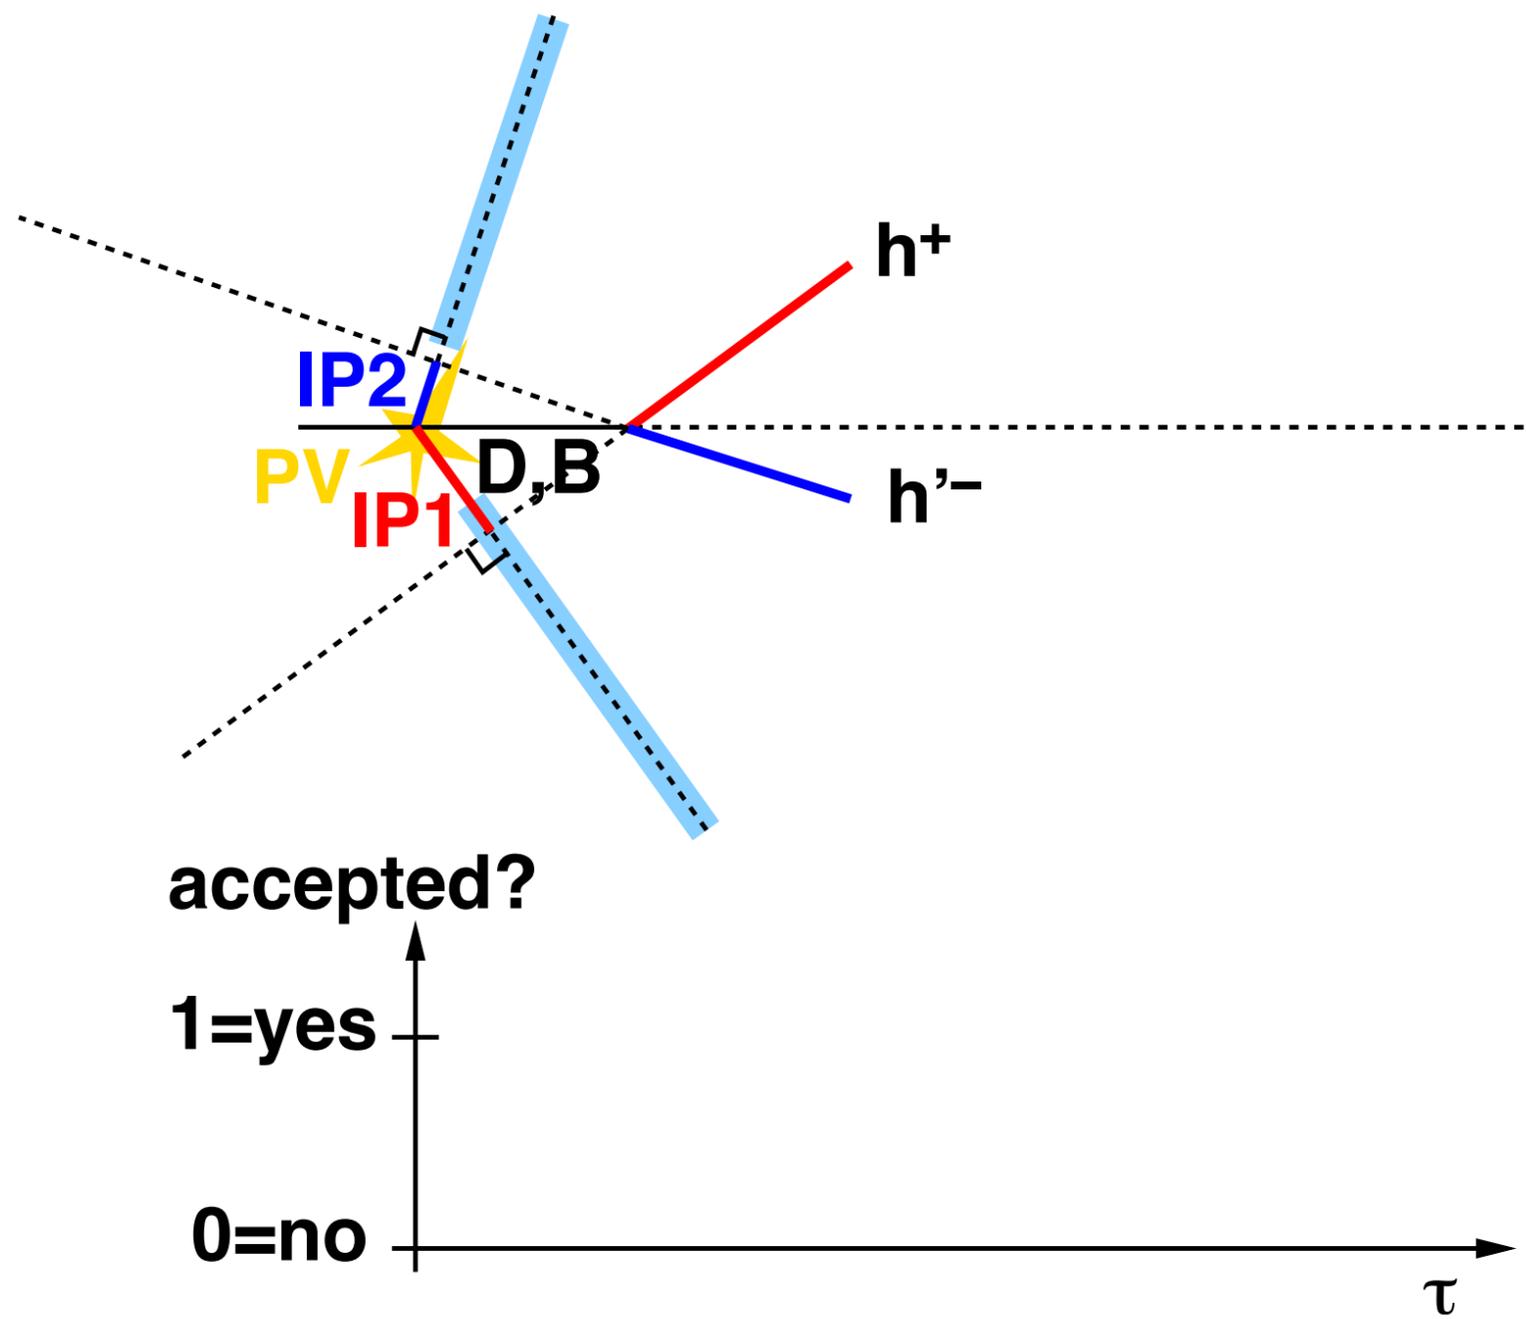
Angular biases?



One of the key advantages of an inclusive trigger is that we minimally bias offline distributions, e.g. angular acceptances in $K^*\mu\mu$, or Dalitz acceptances in $KK\pi$, are kept as flat as possible



Lifetime biases?

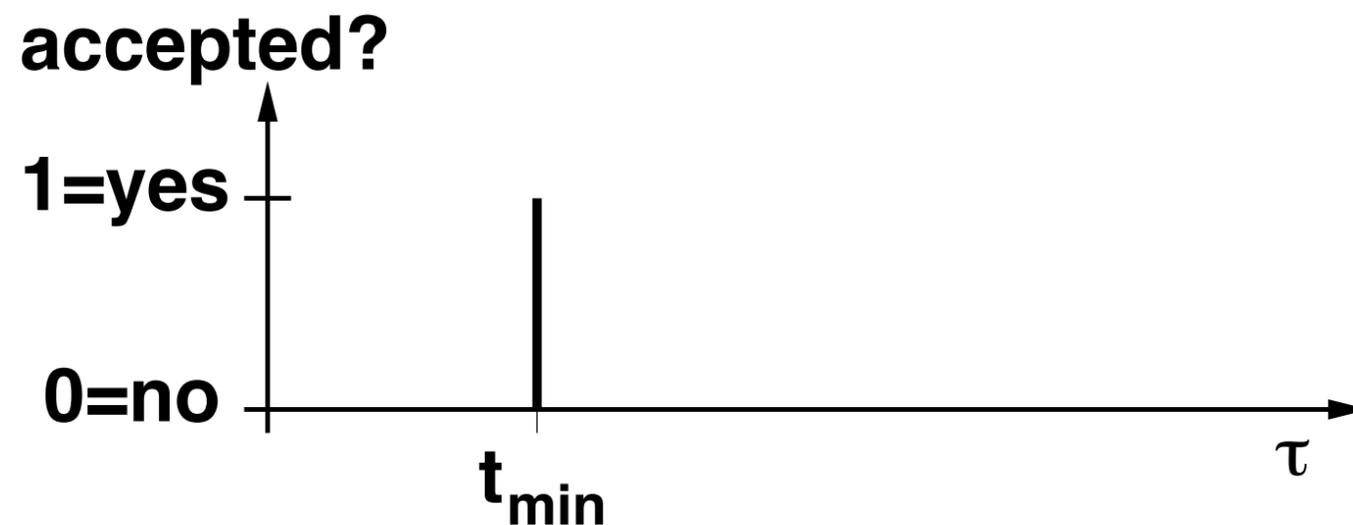
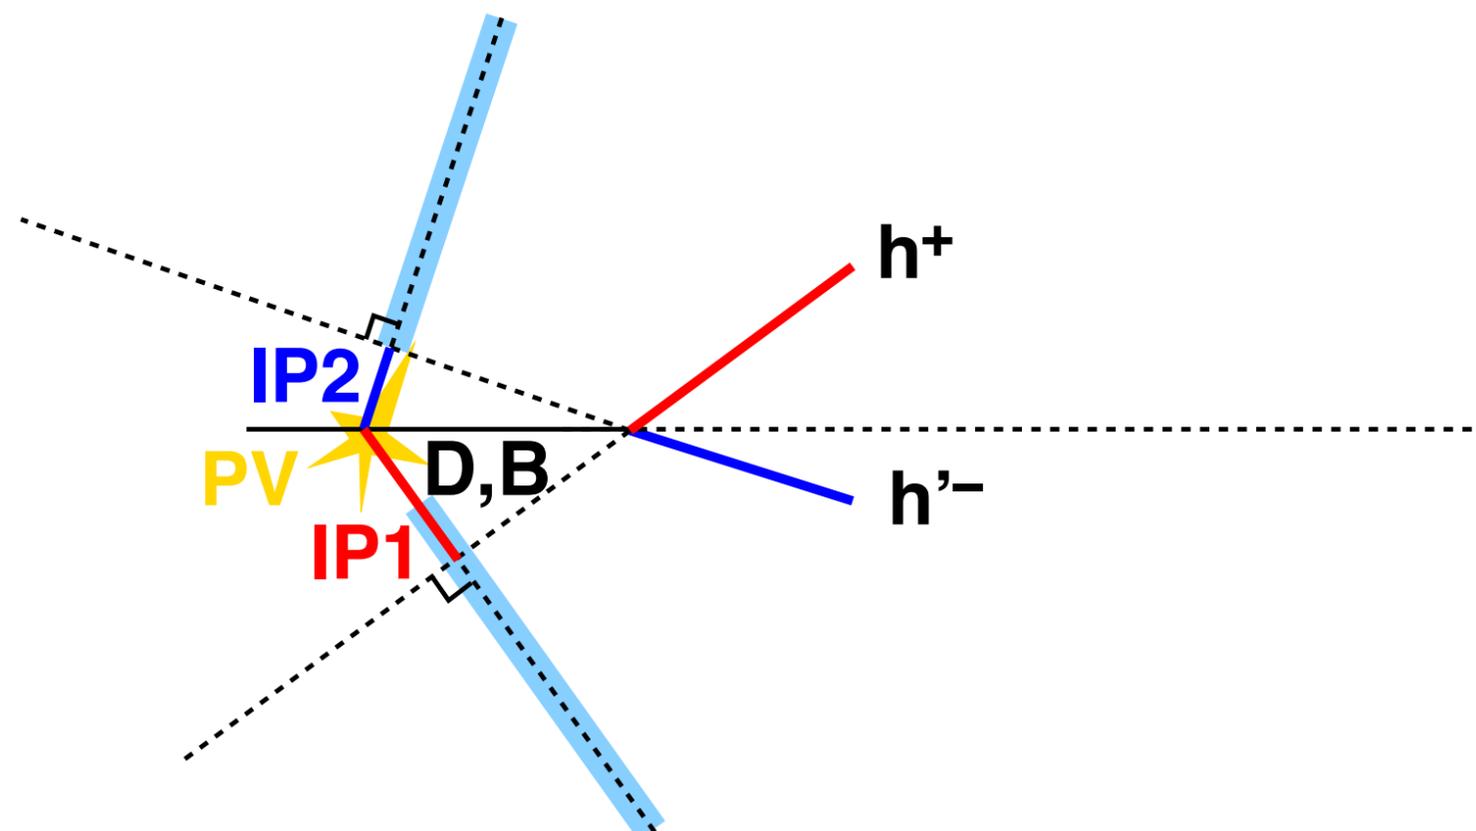


Because we can reproduce the trigger decisions offline, we can measure lifetime biases in a data driven way offline

Get an event-by-event acceptance by replaying the trigger decision for the full range of possible B/D lifetimes

No trigger emulation needed, correct alignment and detector conditions automatically taken into account.

Lifetime biases?

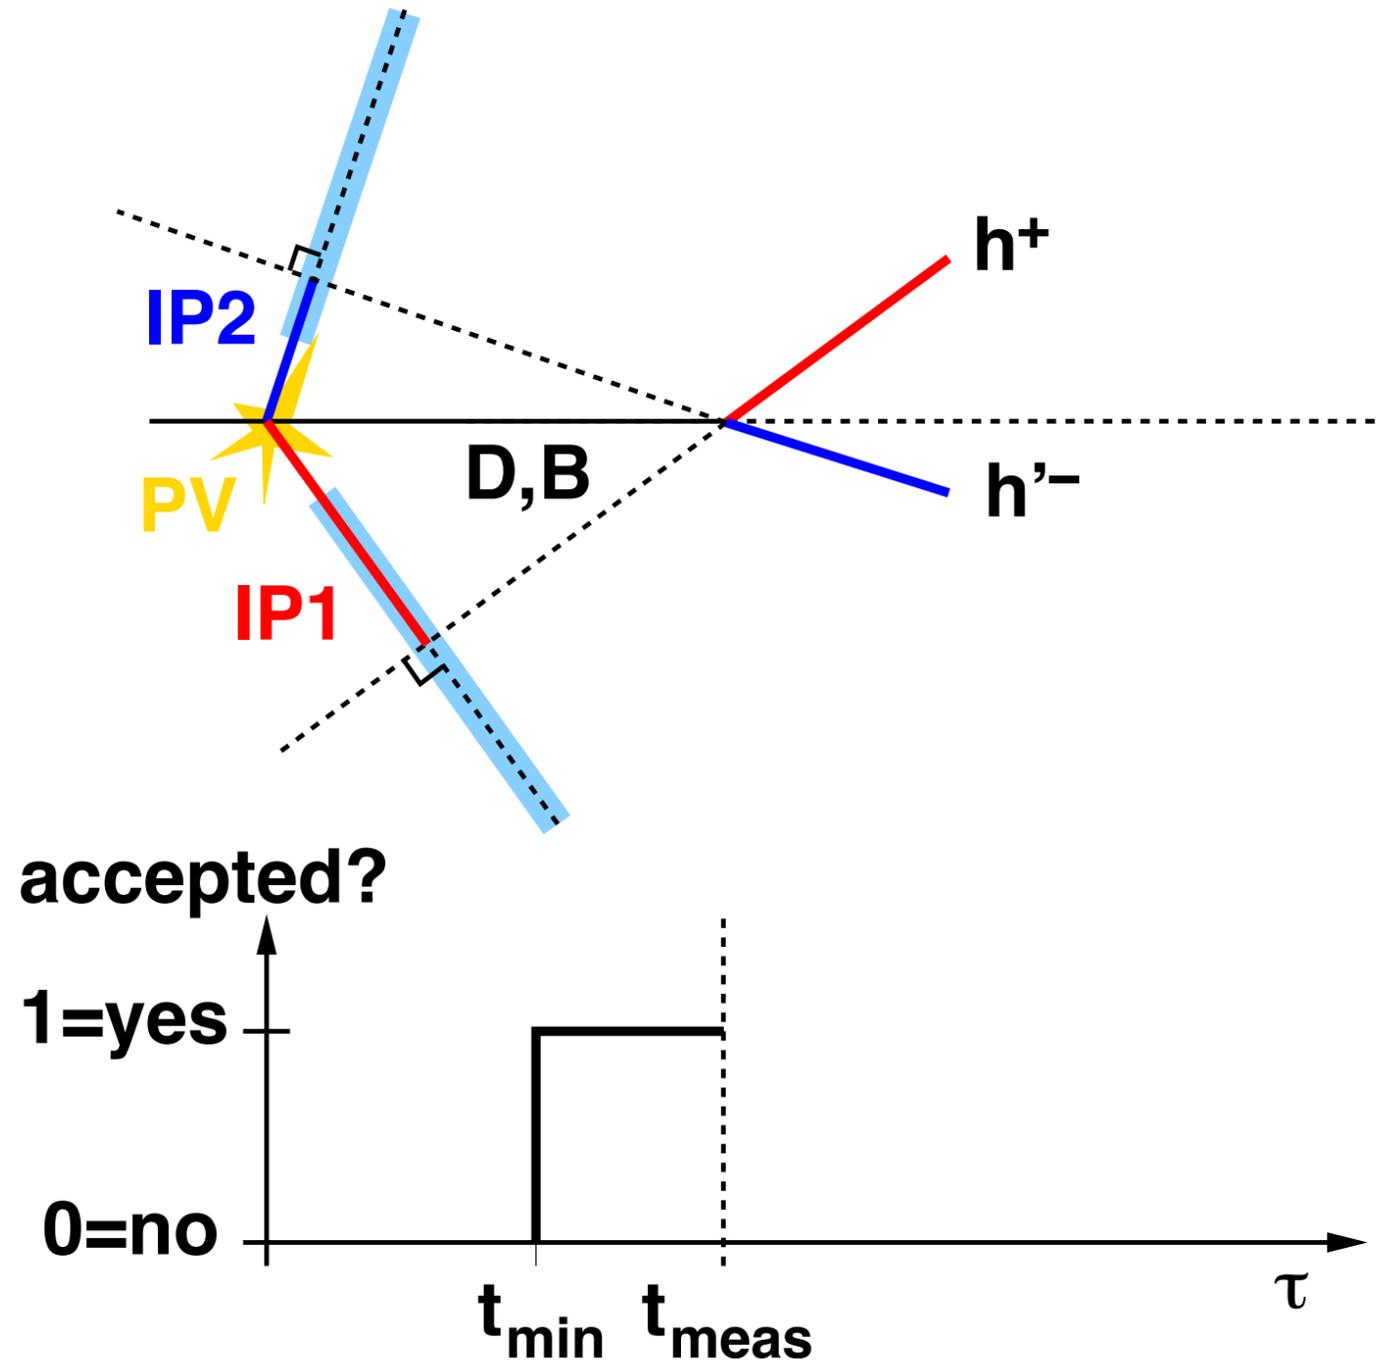


Because we can reproduce the trigger decisions offline, we can measure lifetime biases in a data driven way offline

Get an event-by-event acceptance by replaying the trigger decision for the full range of possible B/D lifetimes

No trigger emulation needed, correct alignment and detector conditions automatically taken into account.

Lifetime biases?

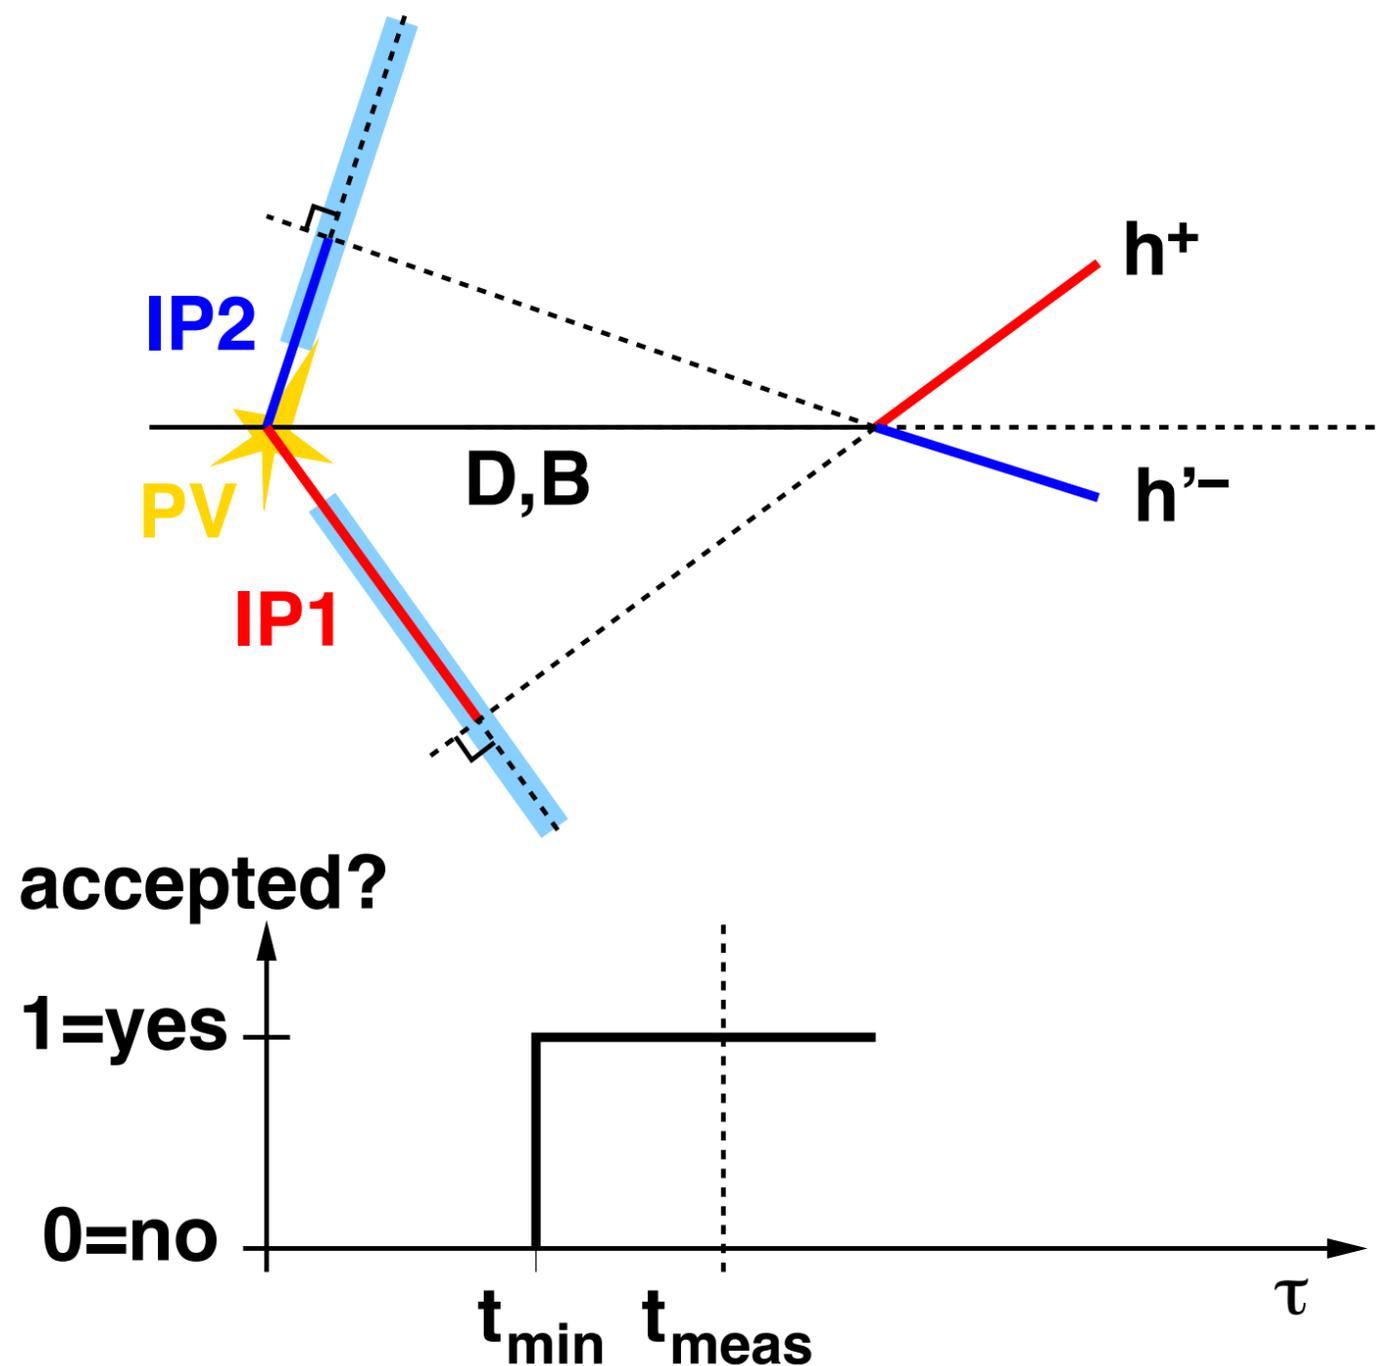


Because we can reproduce the trigger decisions offline, we can measure lifetime biases in a data driven way offline

Get an event-by-event acceptance by replaying the trigger decision for the full range of possible B/D lifetimes

No trigger emulation needed, correct alignment and detector conditions automatically taken into account.

Lifetime biases?

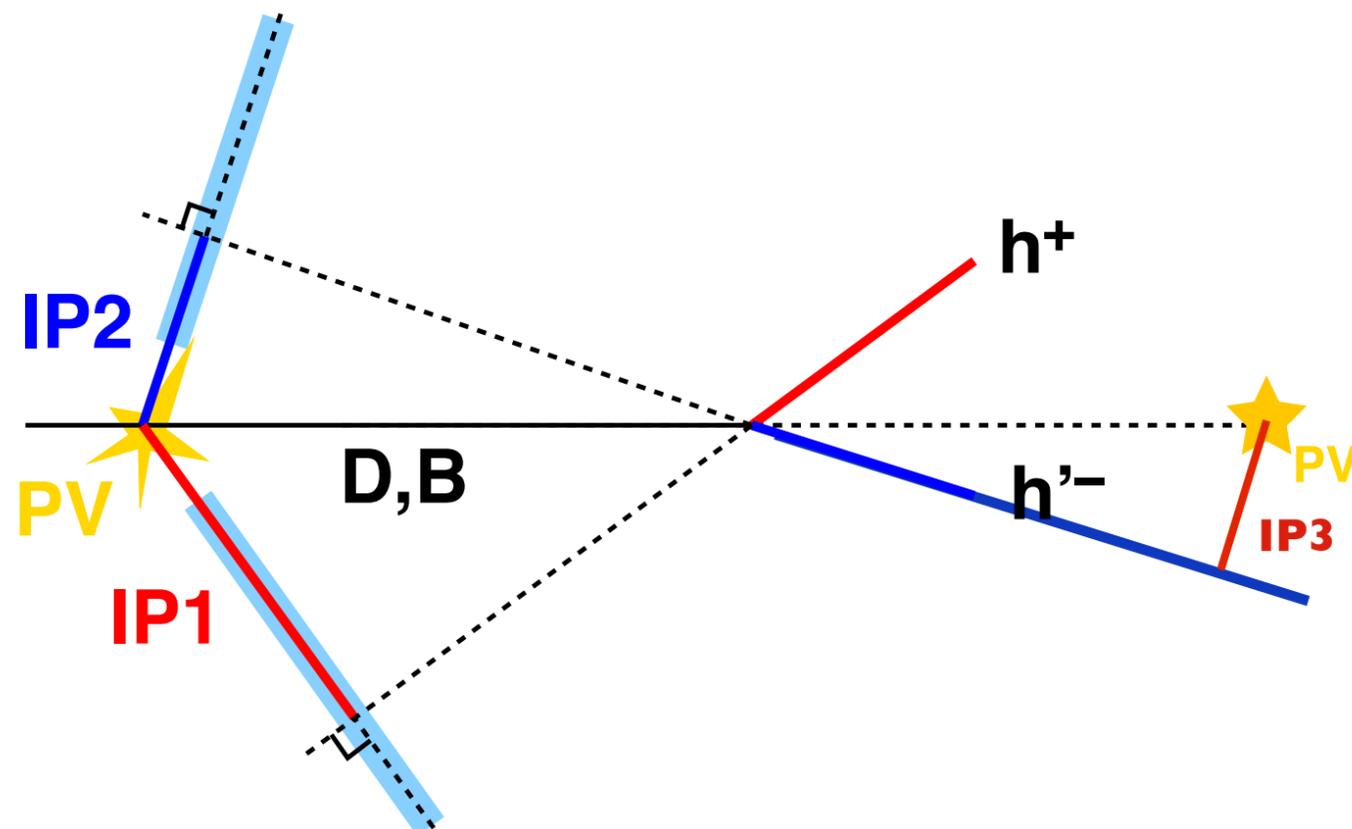


Because we can reproduce the trigger decisions offline, we can measure lifetime biases in a data driven way offline

Get an event-by-event acceptance by replaying the trigger decision for the full range of possible B/D lifetimes

No trigger emulation needed, correct alignment and detector conditions automatically taken into account.

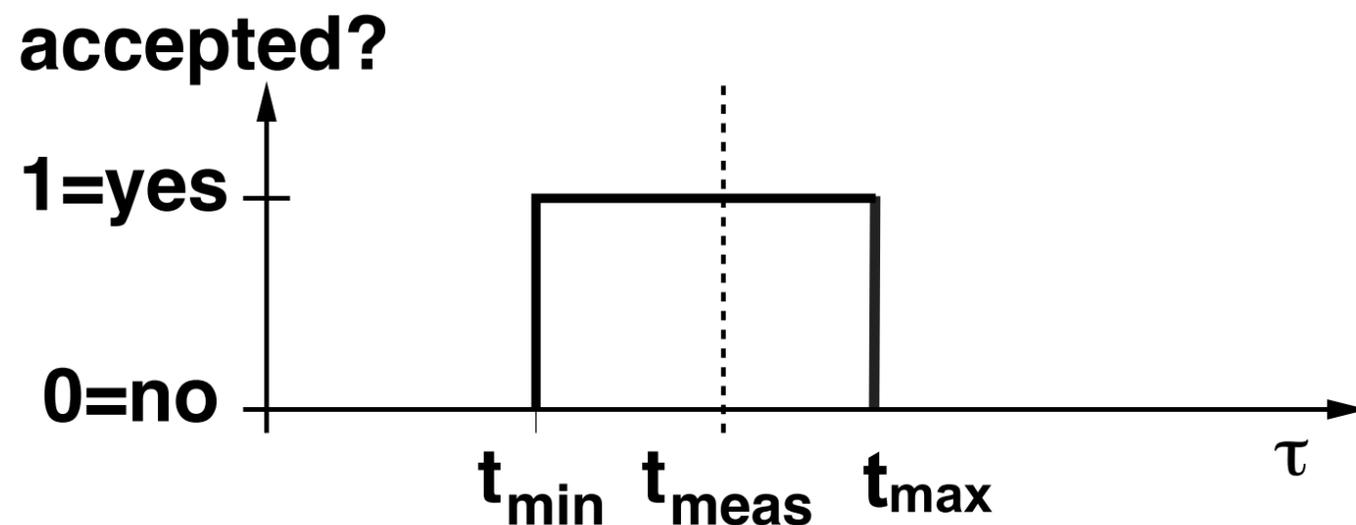
Lifetime biases?



Because we can reproduce the trigger decisions offline, we can measure lifetime biases in a data driven way offline

Get an event-by-event acceptance by replaying the trigger decision for the full range of possible B/D lifetimes

No trigger emulation needed, correct alignment and detector conditions automatically taken into account.



The LHCb physics programme...

Charm Physics

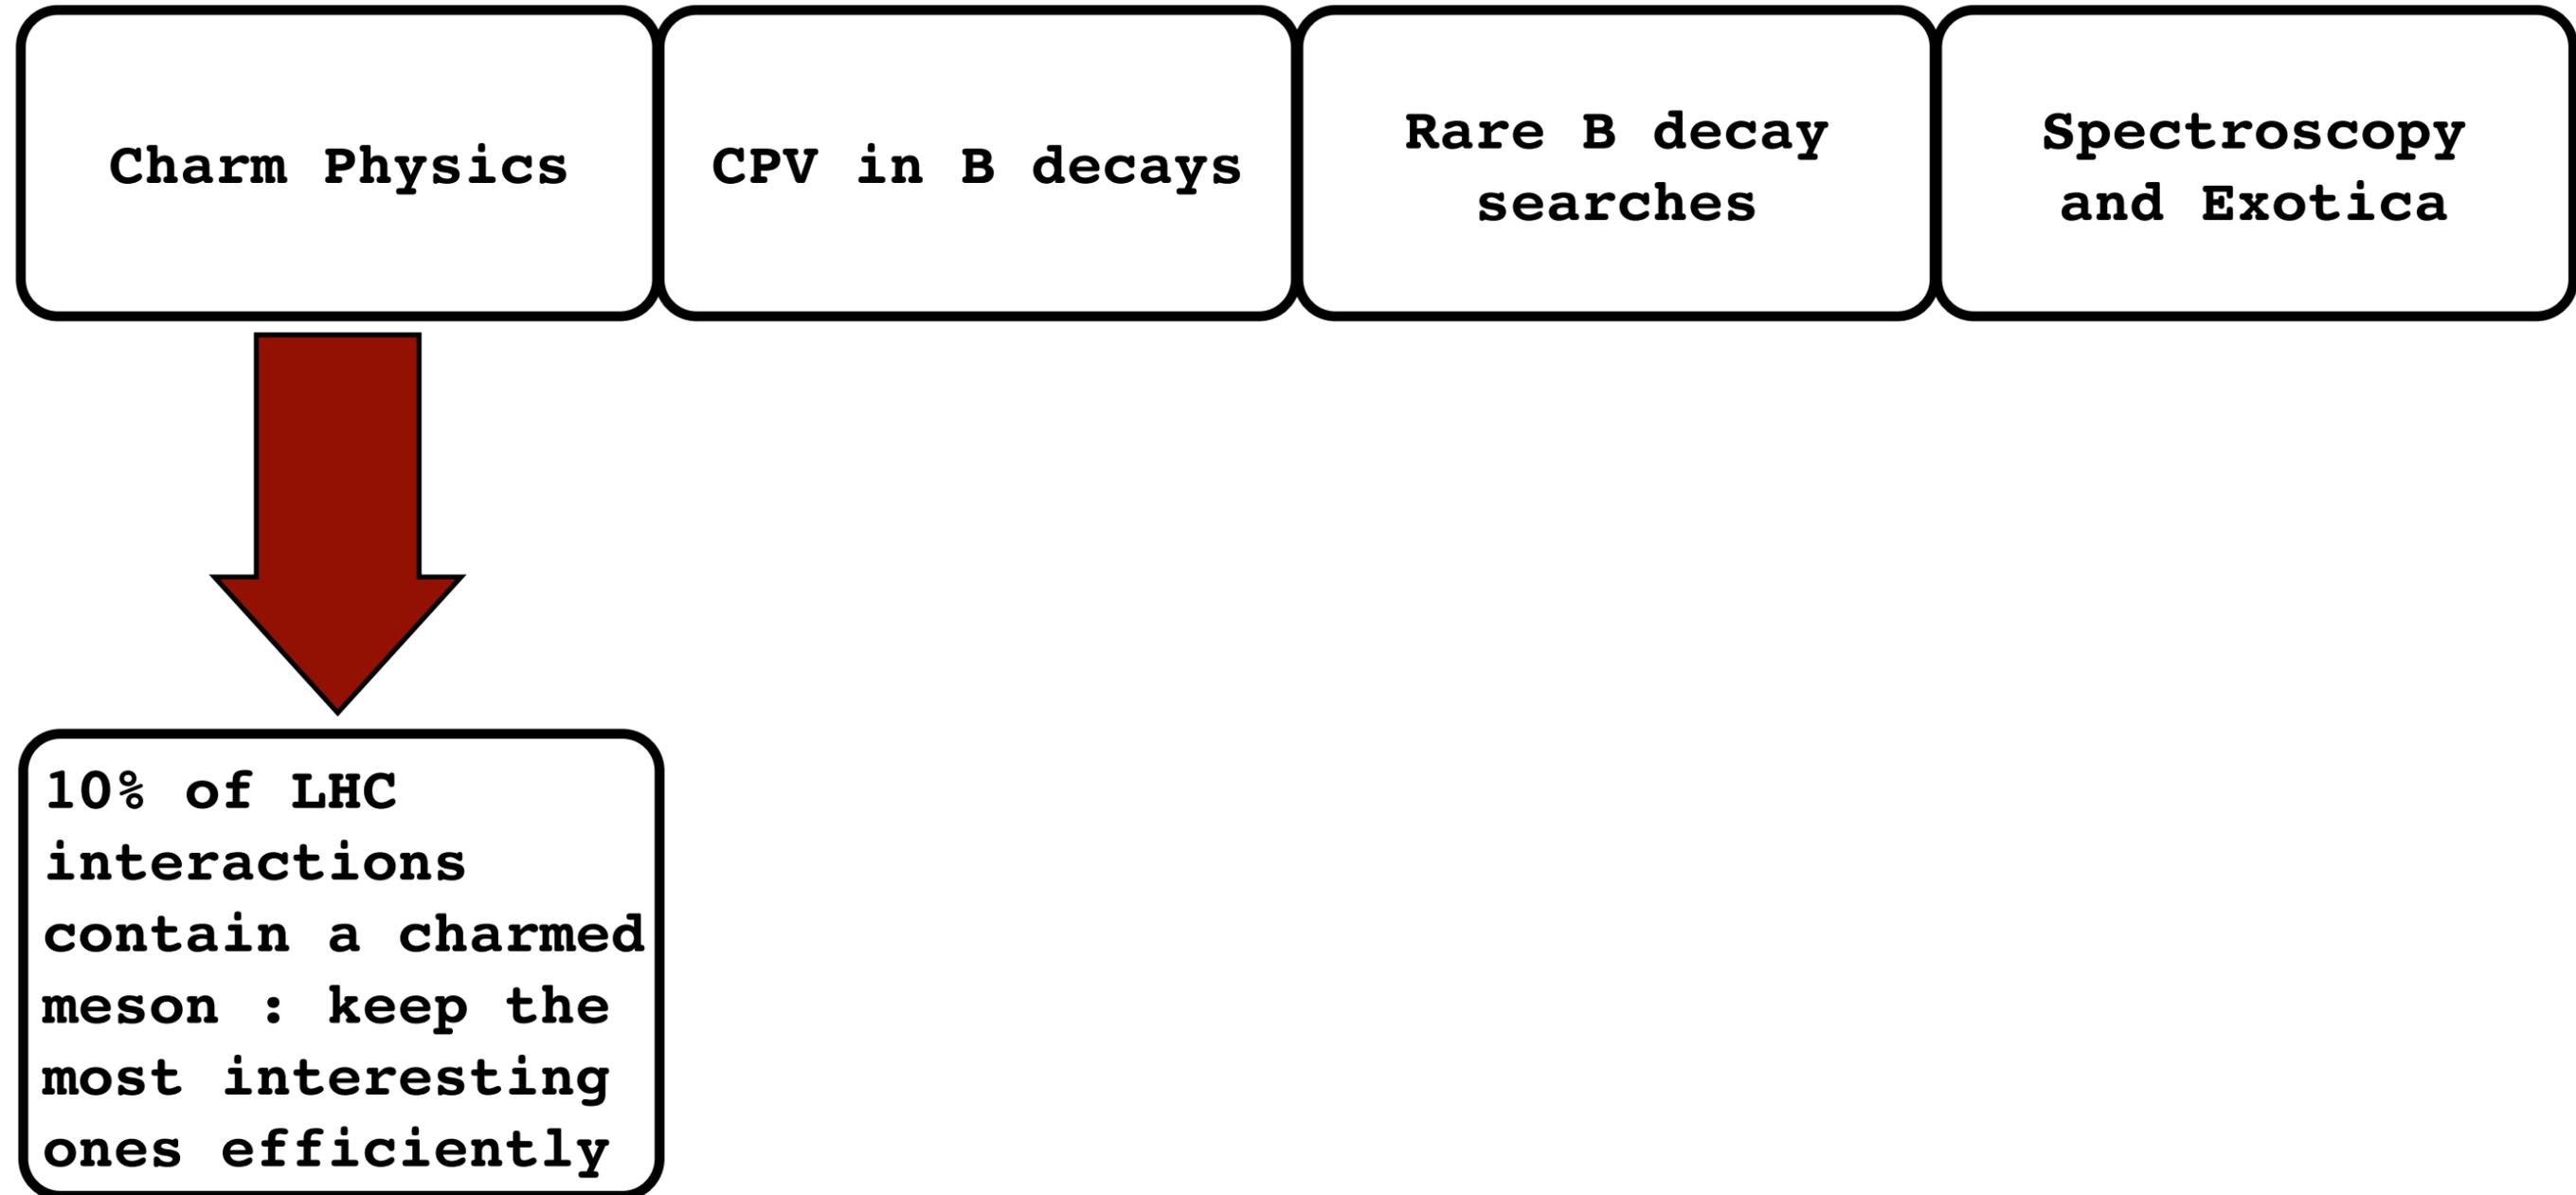
CPV in B decays

**Rare B decay
searches**

**Spectroscopy
and Exotica**

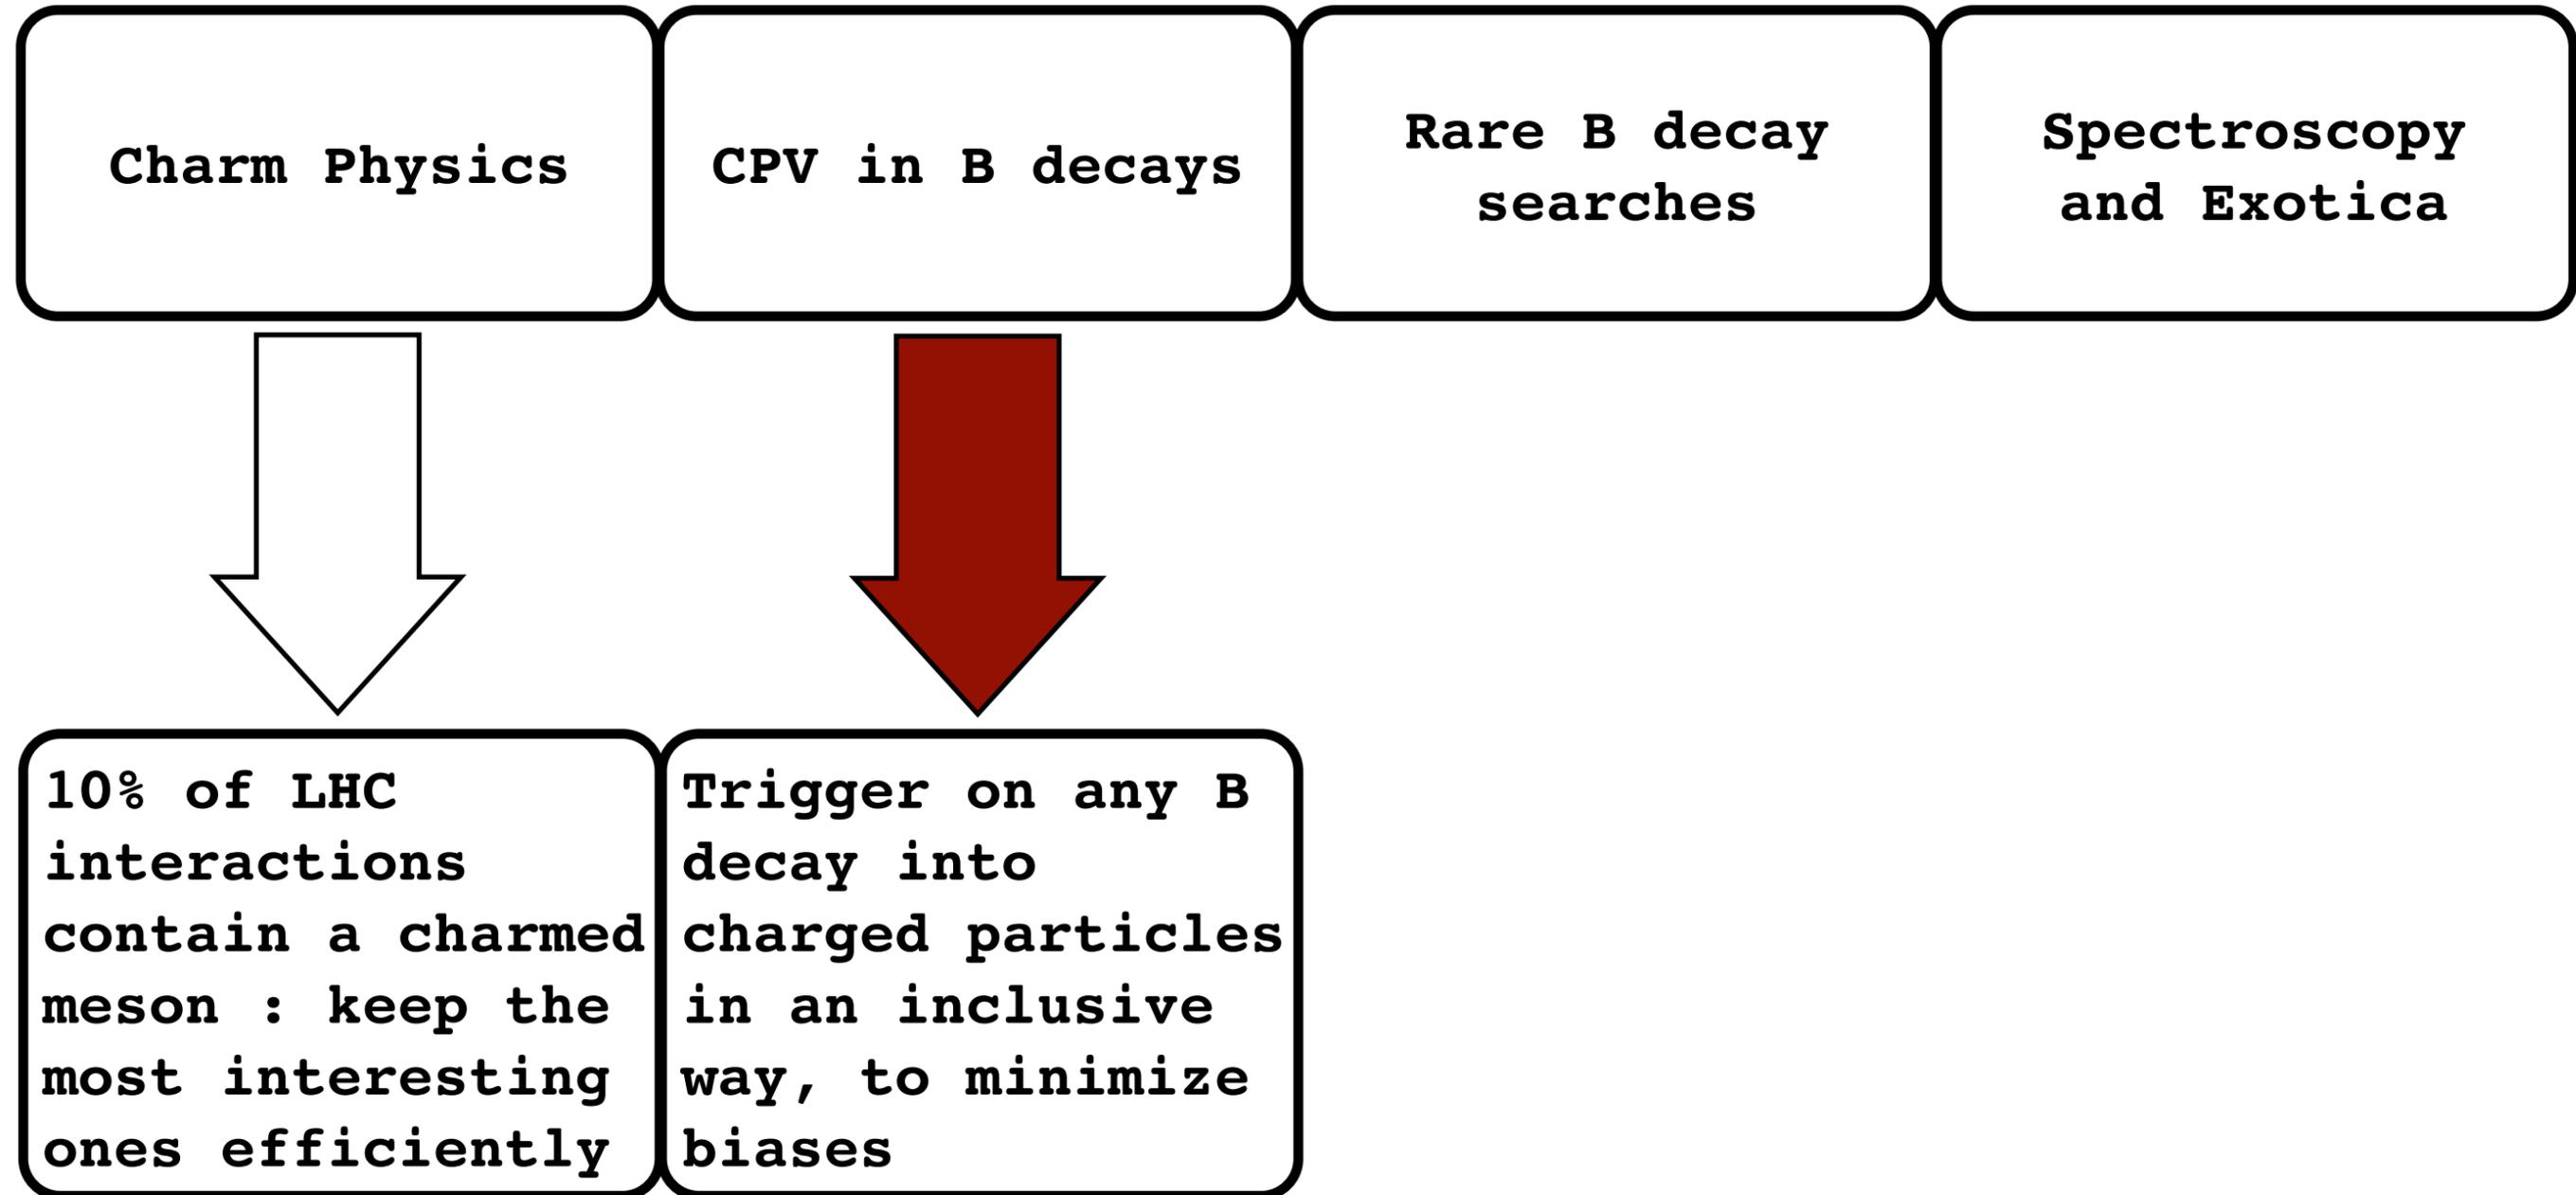
**Note : clearly not the entire physics programme,
see the [LHCb upgrade LOI](#) for more details**

...and its demands on the trigger



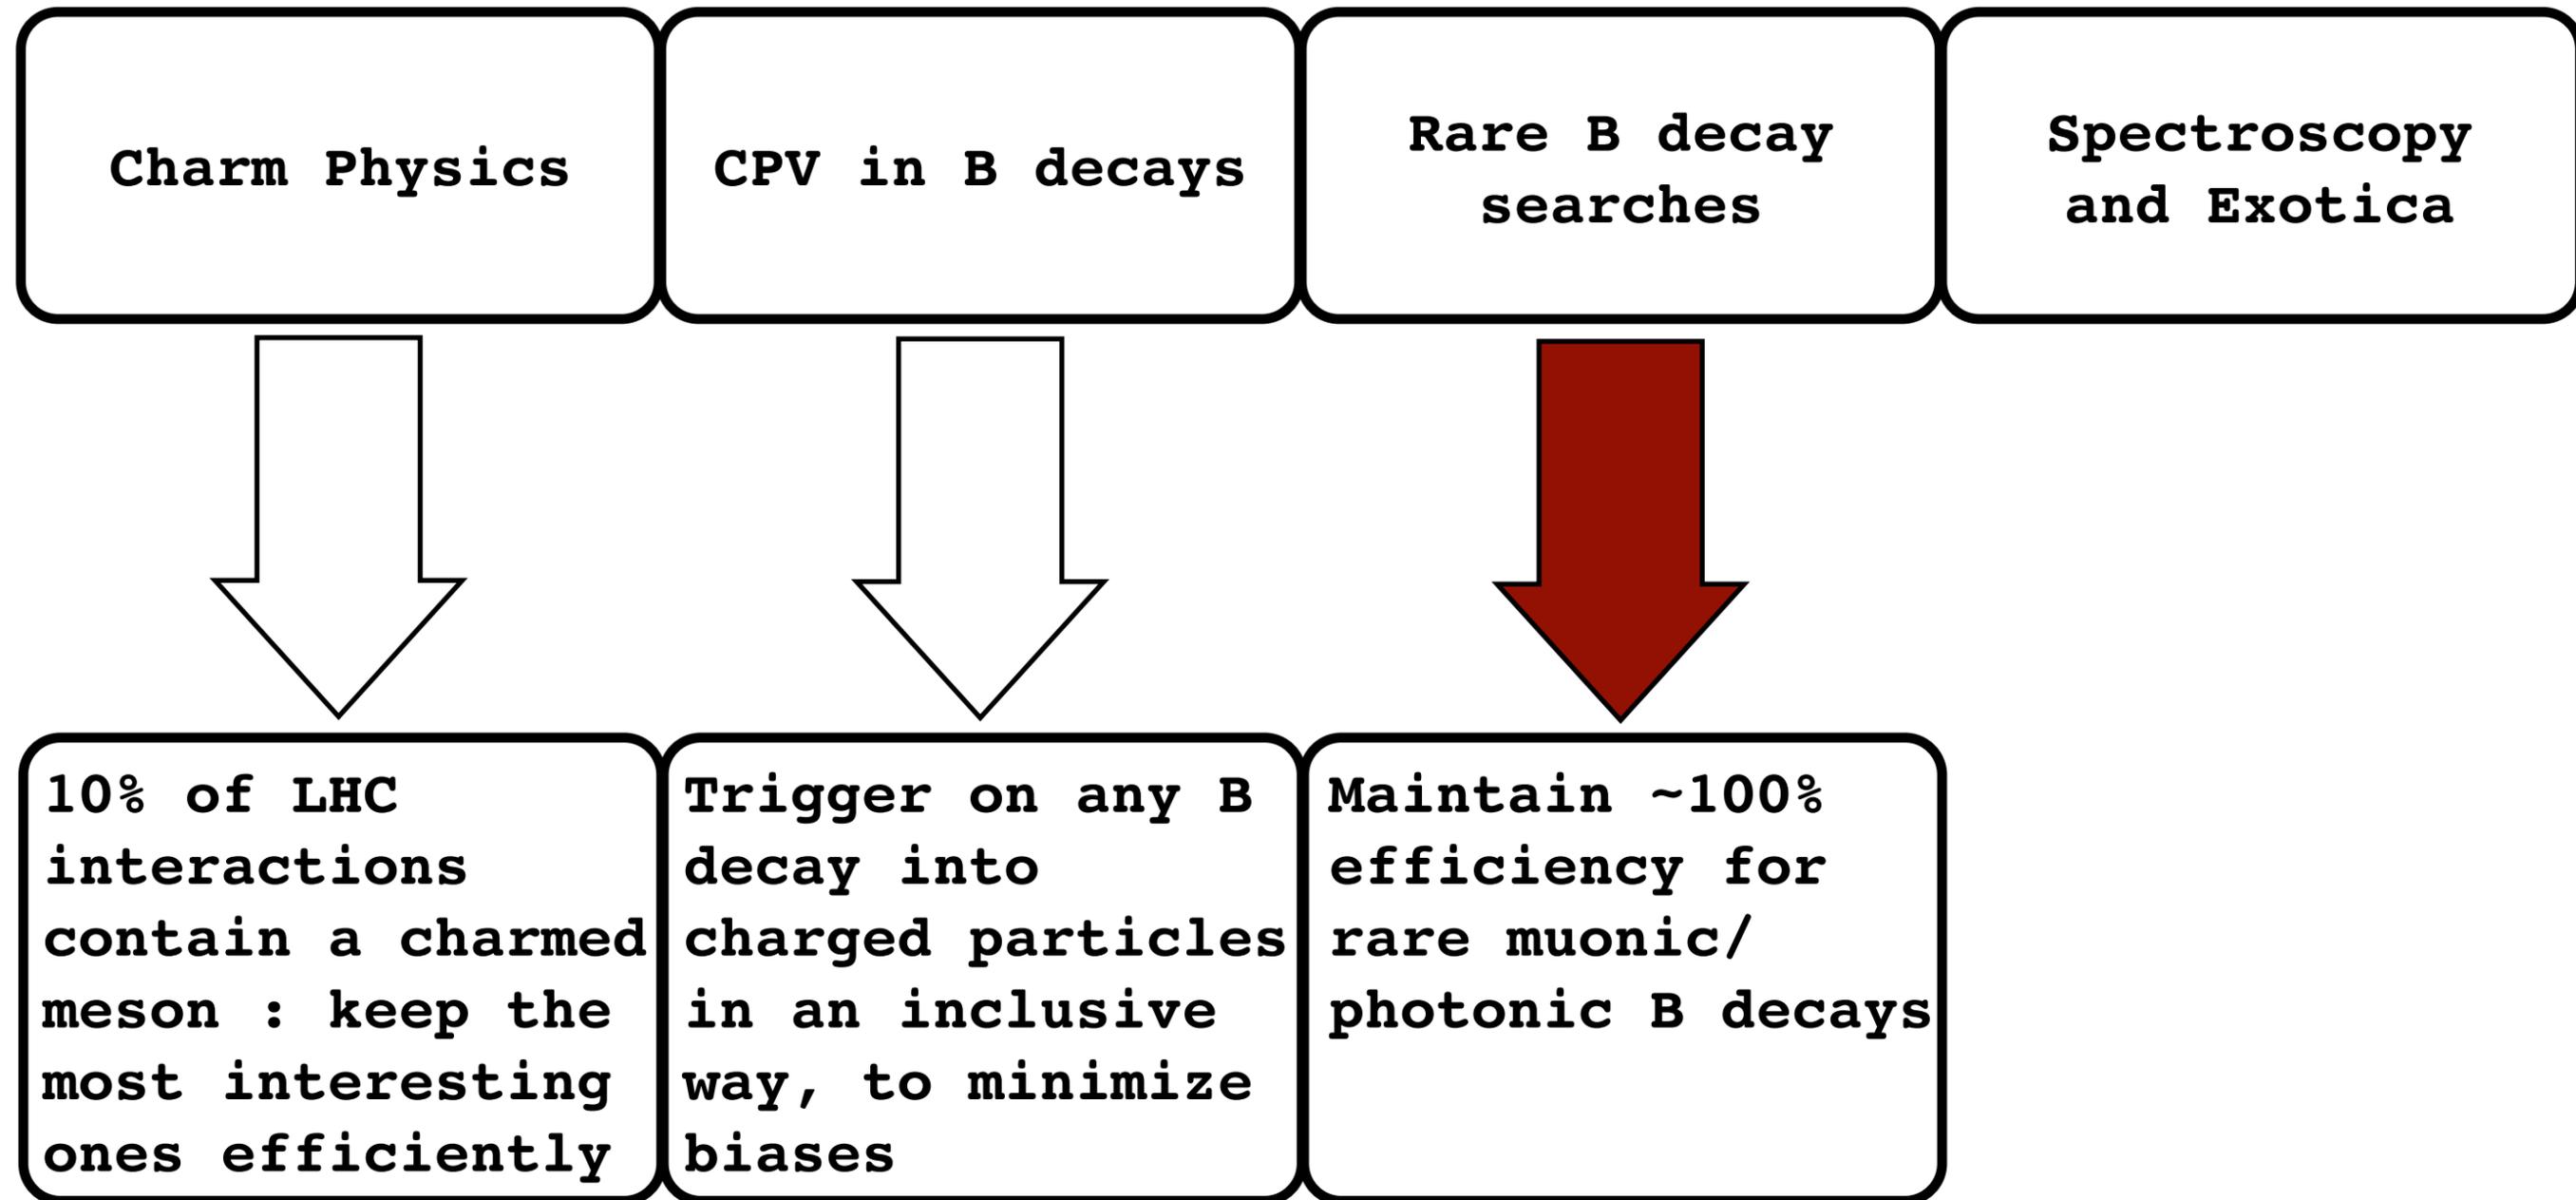
Note : clearly not the entire physics programme, see the [LHCb upgrade LOI](#) for more details

...and its demands on the trigger



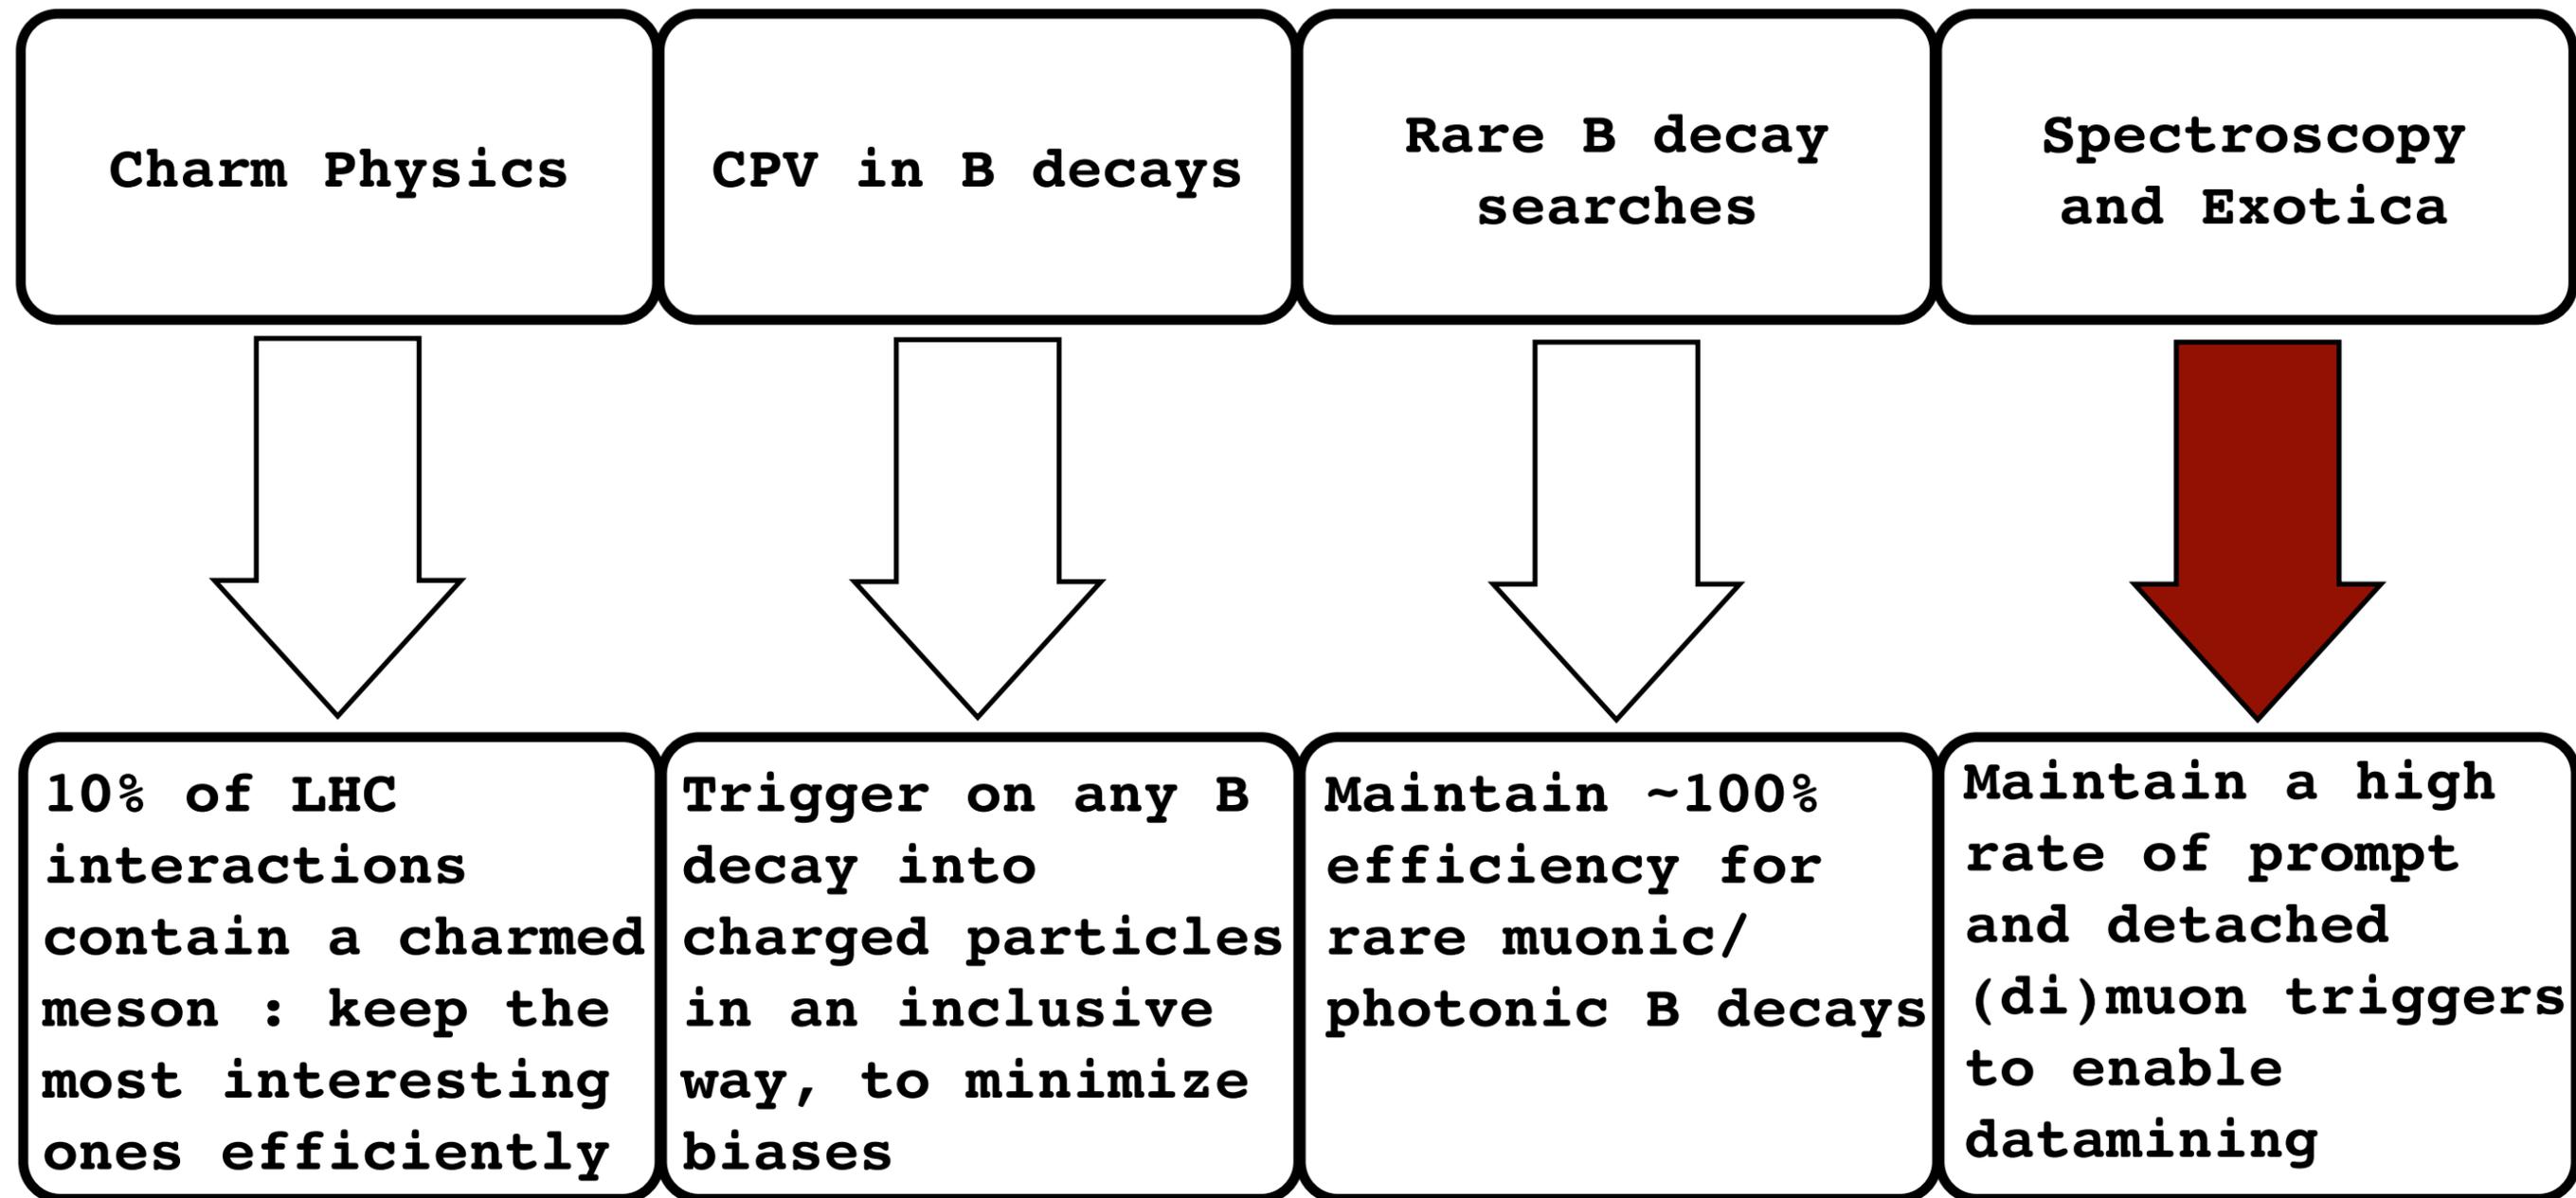
Note : clearly not the entire physics programme, see the [LHCb upgrade LOI](#) for more details

...and its demands on the trigger



Note : clearly not the entire physics programme, see the [LHCb upgrade LOI](#) for more details

...and its demands on the trigger



Note : clearly not the entire physics programme, see the [LHCb upgrade LOI](#) for more details

...and its demands on the trigger

Charm Physics

CPV in B decays

Rare B decay searches

Spectroscopy and Exotica

And all this must fit into an output rate of ~4 kHz!

KEY CHALLENGE : discriminate against prompt charm (300 kHz in the LHCb acceptance) while keeping the most interesting prompt charm!

10% of LHC interactions contain a charmed meson : keep the most interesting ones efficiently

Trigger on any B decay into charged particles in an inclusive way, to minimize biases

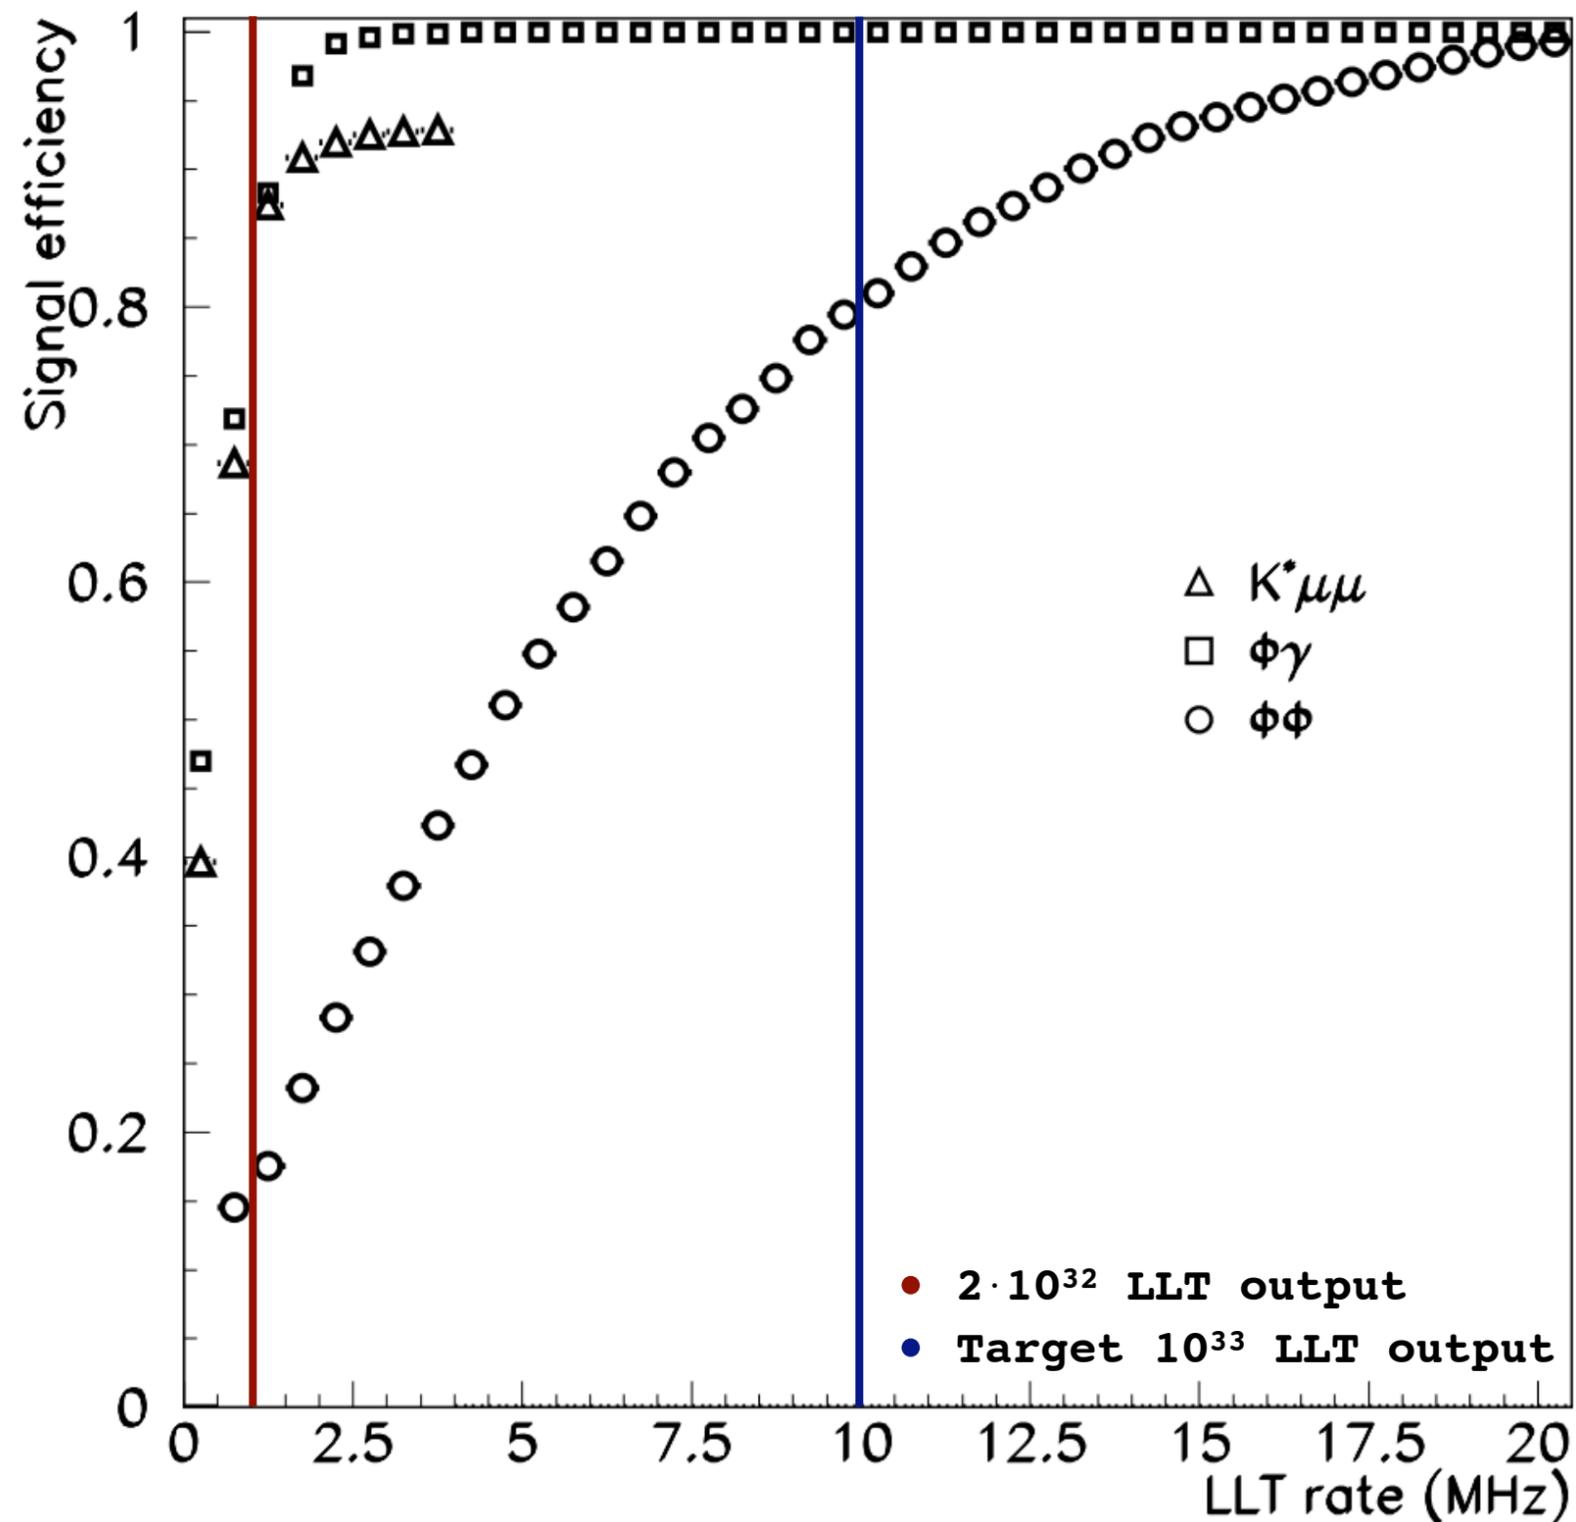
Maintain ~100% efficiency for rare muonic/photonic B decays

Maintain a high rate of prompt and detached (di)muon triggers to enable datamining

Note : clearly not the entire physics programme, see the [LHCb upgrade LOI](#) for more details

Triggers in the
era of the
upgraded LHC

The LHCb trigger upgrade



The 1 MHz detector readout is the bottleneck in the current DAQ chain

Particularly limiting for hadronic decay modes, and would become more limiting as the luminosity rises due to pileup

Therefore LHCb will upgrade all subdetectors to read out at 40 MHz

And then scale the actual detector readout according to the available CPU capacity in the HLT farm

Make the L0 (LLT) trigger less and less important as the upgrade progresses