Triggers and Machine Learning

Mika Vesterinen

27th August 2014

1. Introduction
   1. LHC/trigger generalities
   2. Low level triggers
   3. High level triggers
2. MLHEP and triggers (LHCb HLT)

UNIVERSITÄT
HEIDELBERG
ZUKUNFT
SEIT 1386

Unterstützt von / Supported by
Alexander von Humboldt
Stiftung / Foundation

The LHC

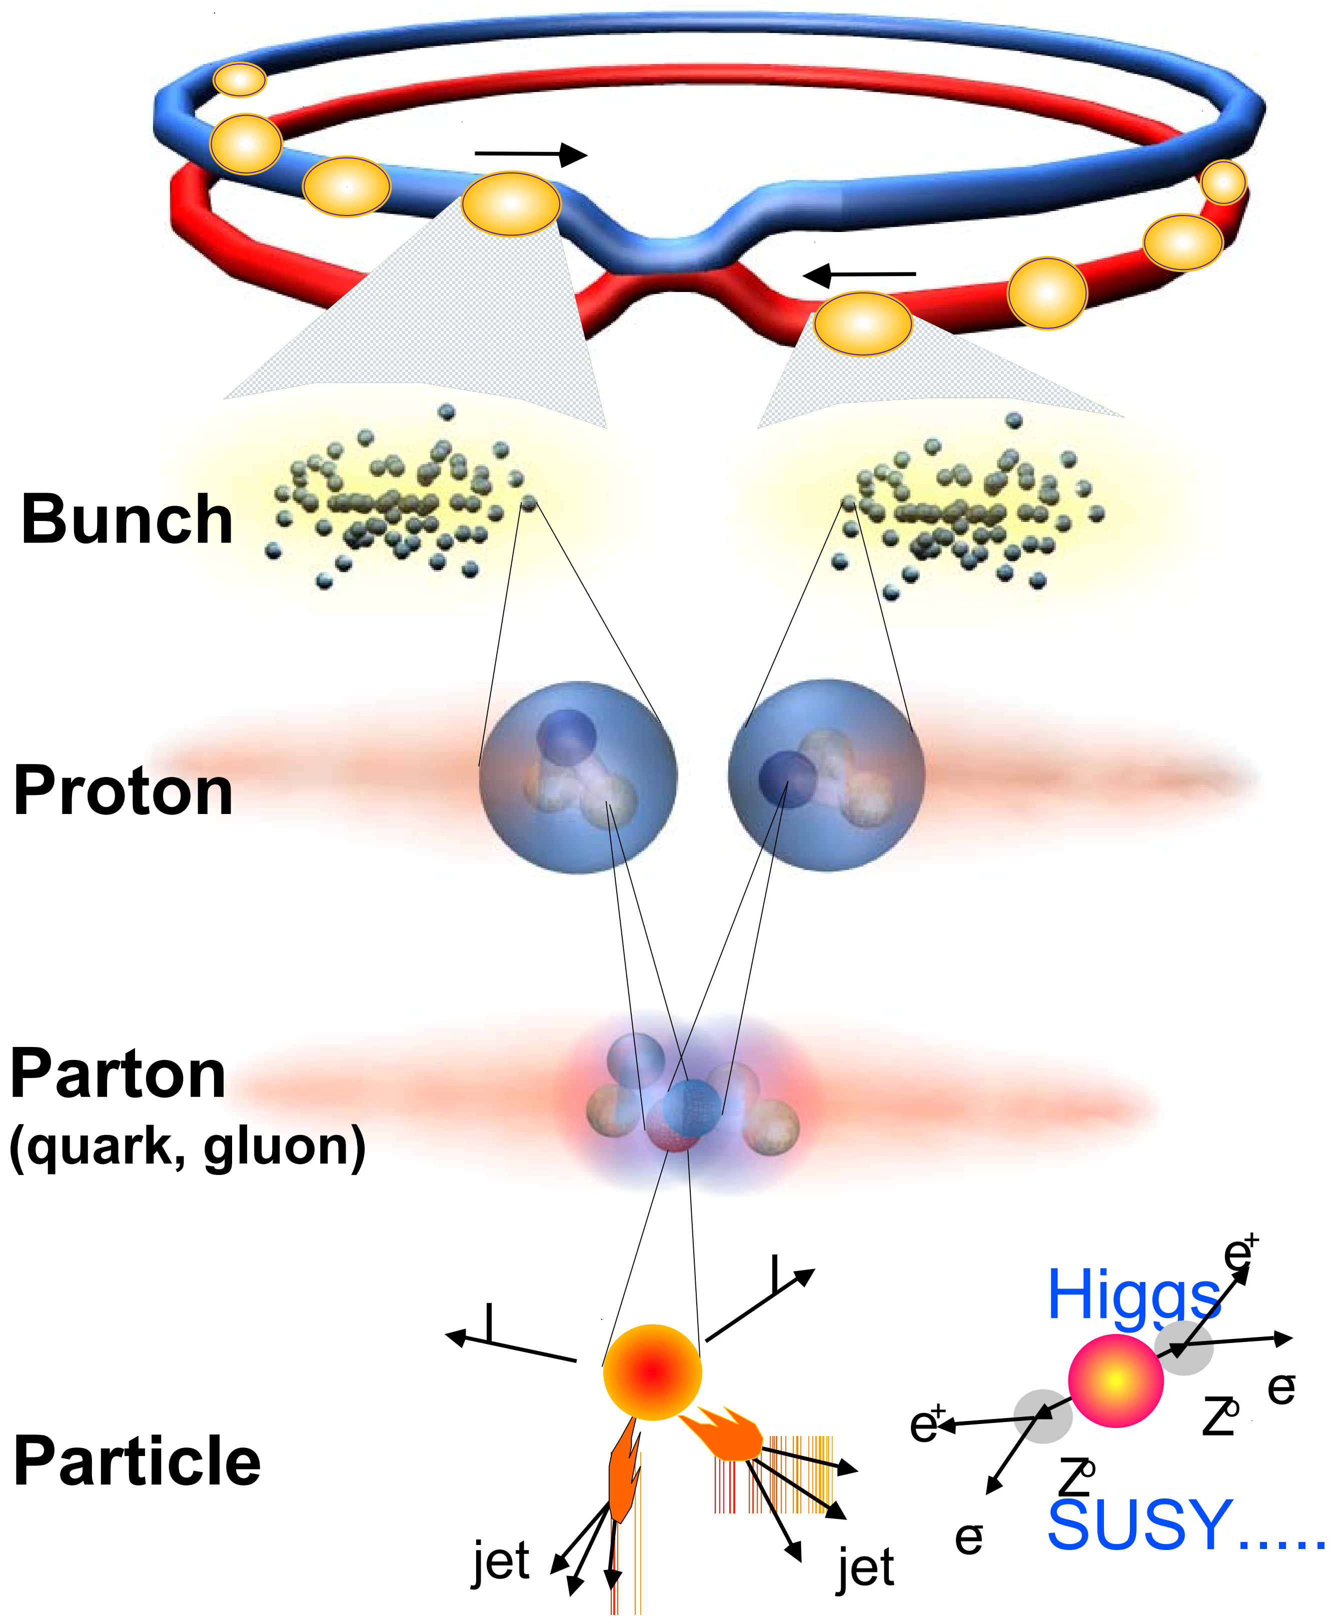

Proton - Proton	2804 bunch/beam
Protons/bunch	10^{11}
Beam energy	7 TeV (7×10^{12} eV)
Luminosity	$10^{34} cm^{-2}s^{-1}$
Crossing rate	40 MHz
Collision rate ≈	10^7-10^9

New physics rate ≈ .00001 Hz

Event selection: 10^{13}

1 in 10,000,000,000,000

The detectors

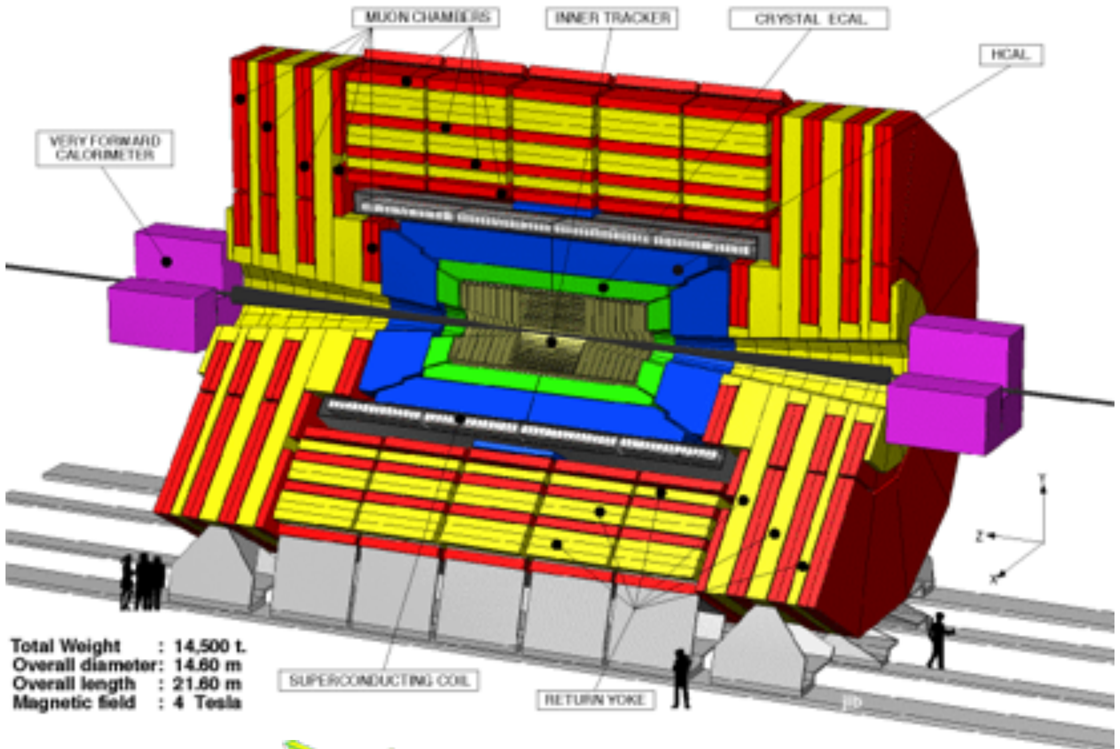

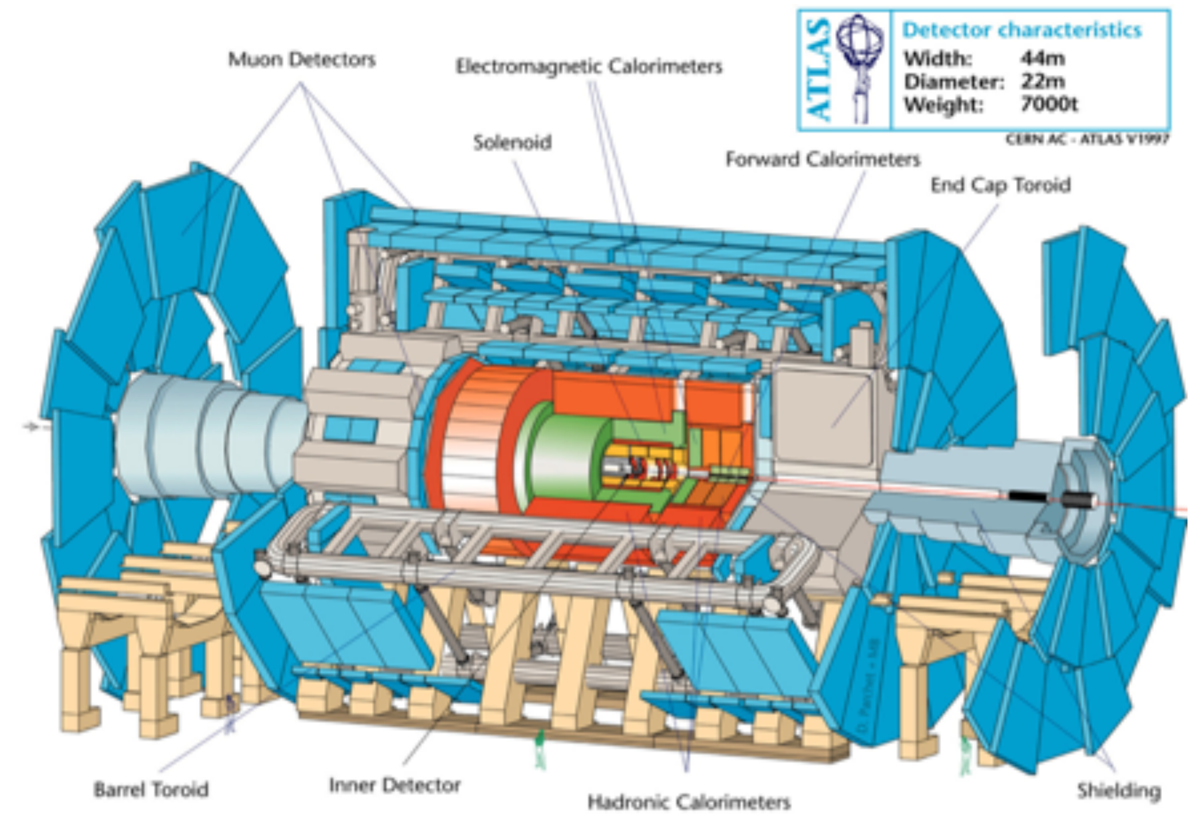

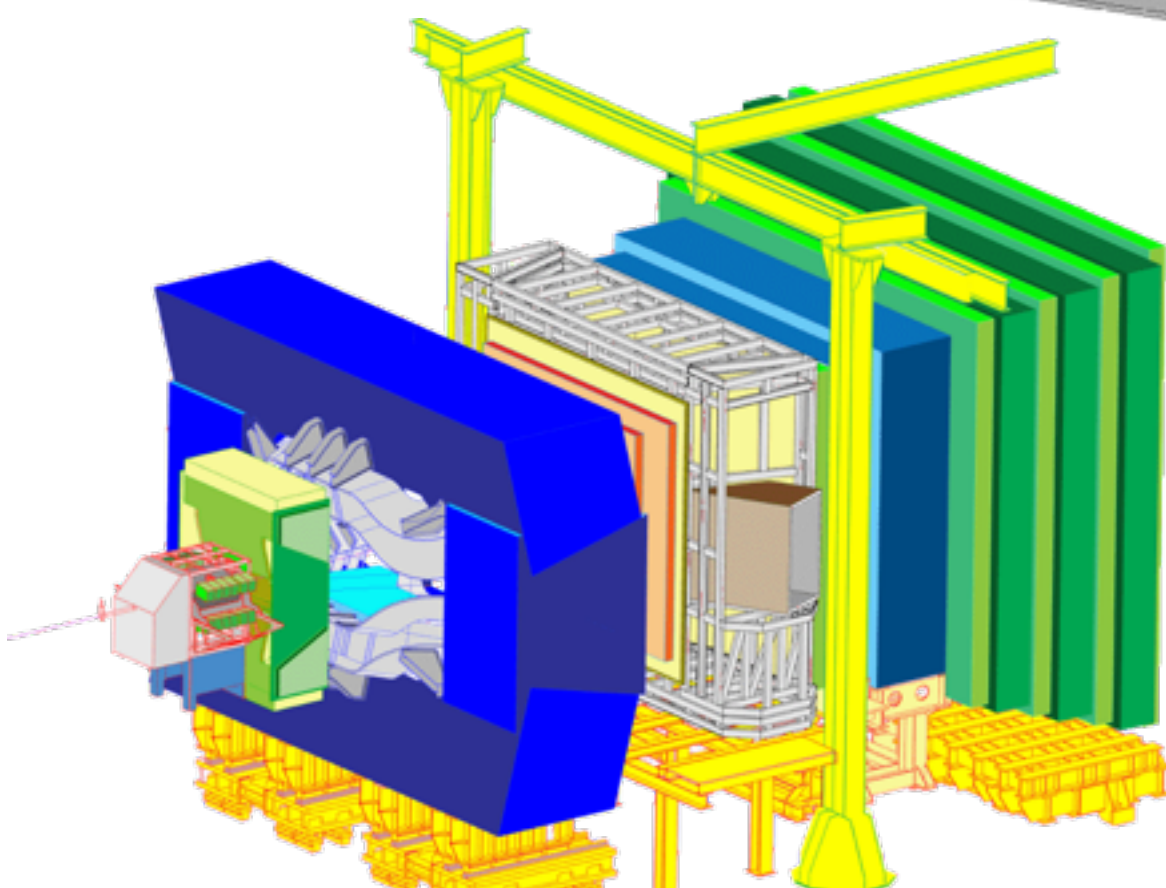

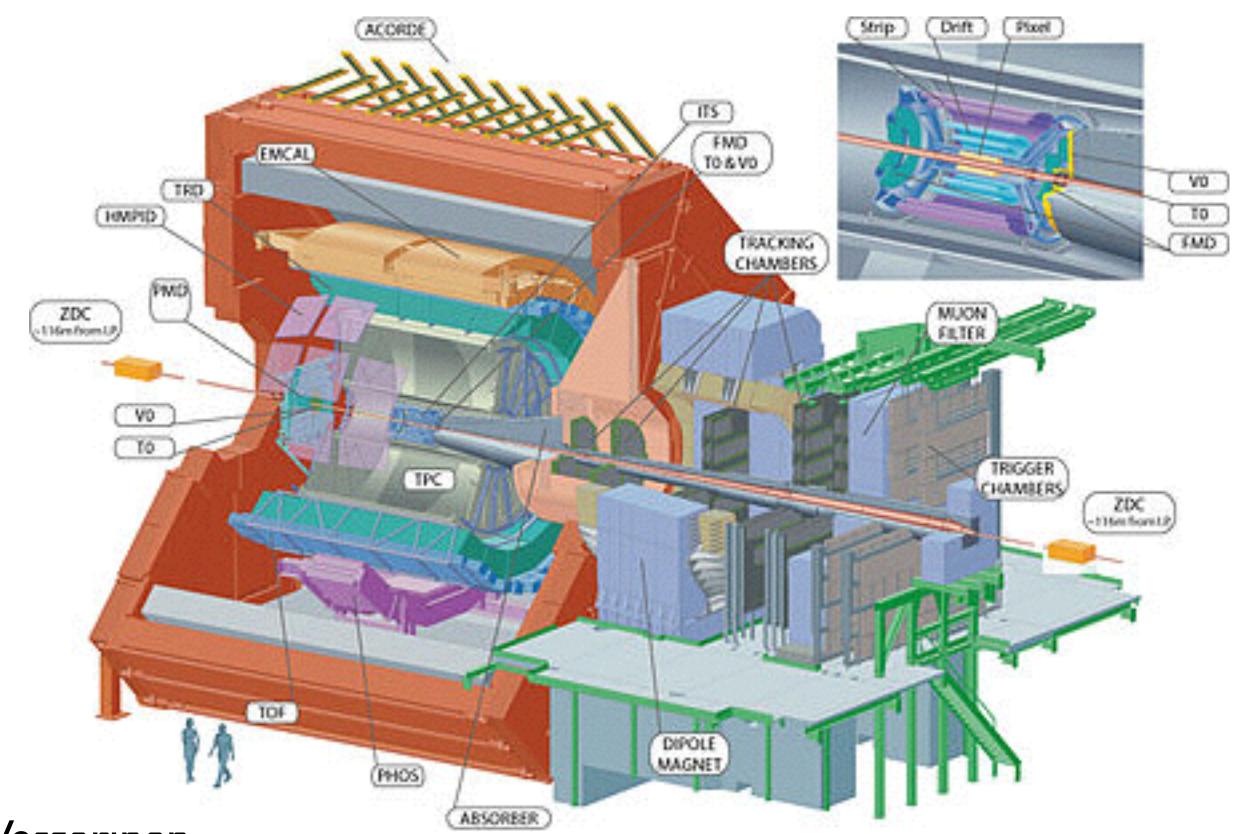

Nomenclature

Event: A single bunch crossing (40 MHz)

(Versus single proton-proton **interaction**)

Nomenclature

Event: A single bunch crossing (40 MHz)

(Versus single proton-proton **interaction**)

$\sigma \sim 70$ mb

$L_{ATLAS/CMS} = 10^{34} cm^{-2} s^{-1}$

Interaction rate = L x σ ~ 700 MHz

~20 "pileup" interactions per event@ 40 MHz

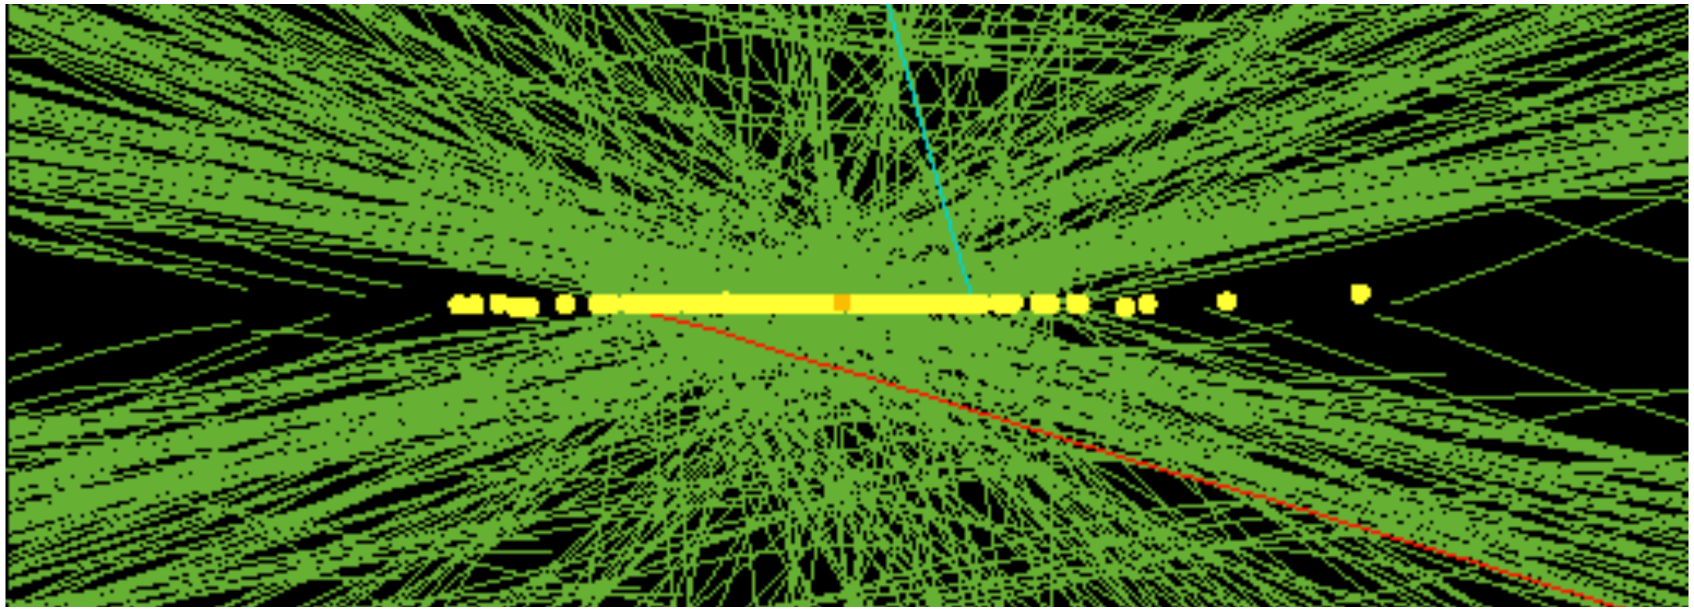

Q: How much would it cost to have a trigger less detector?

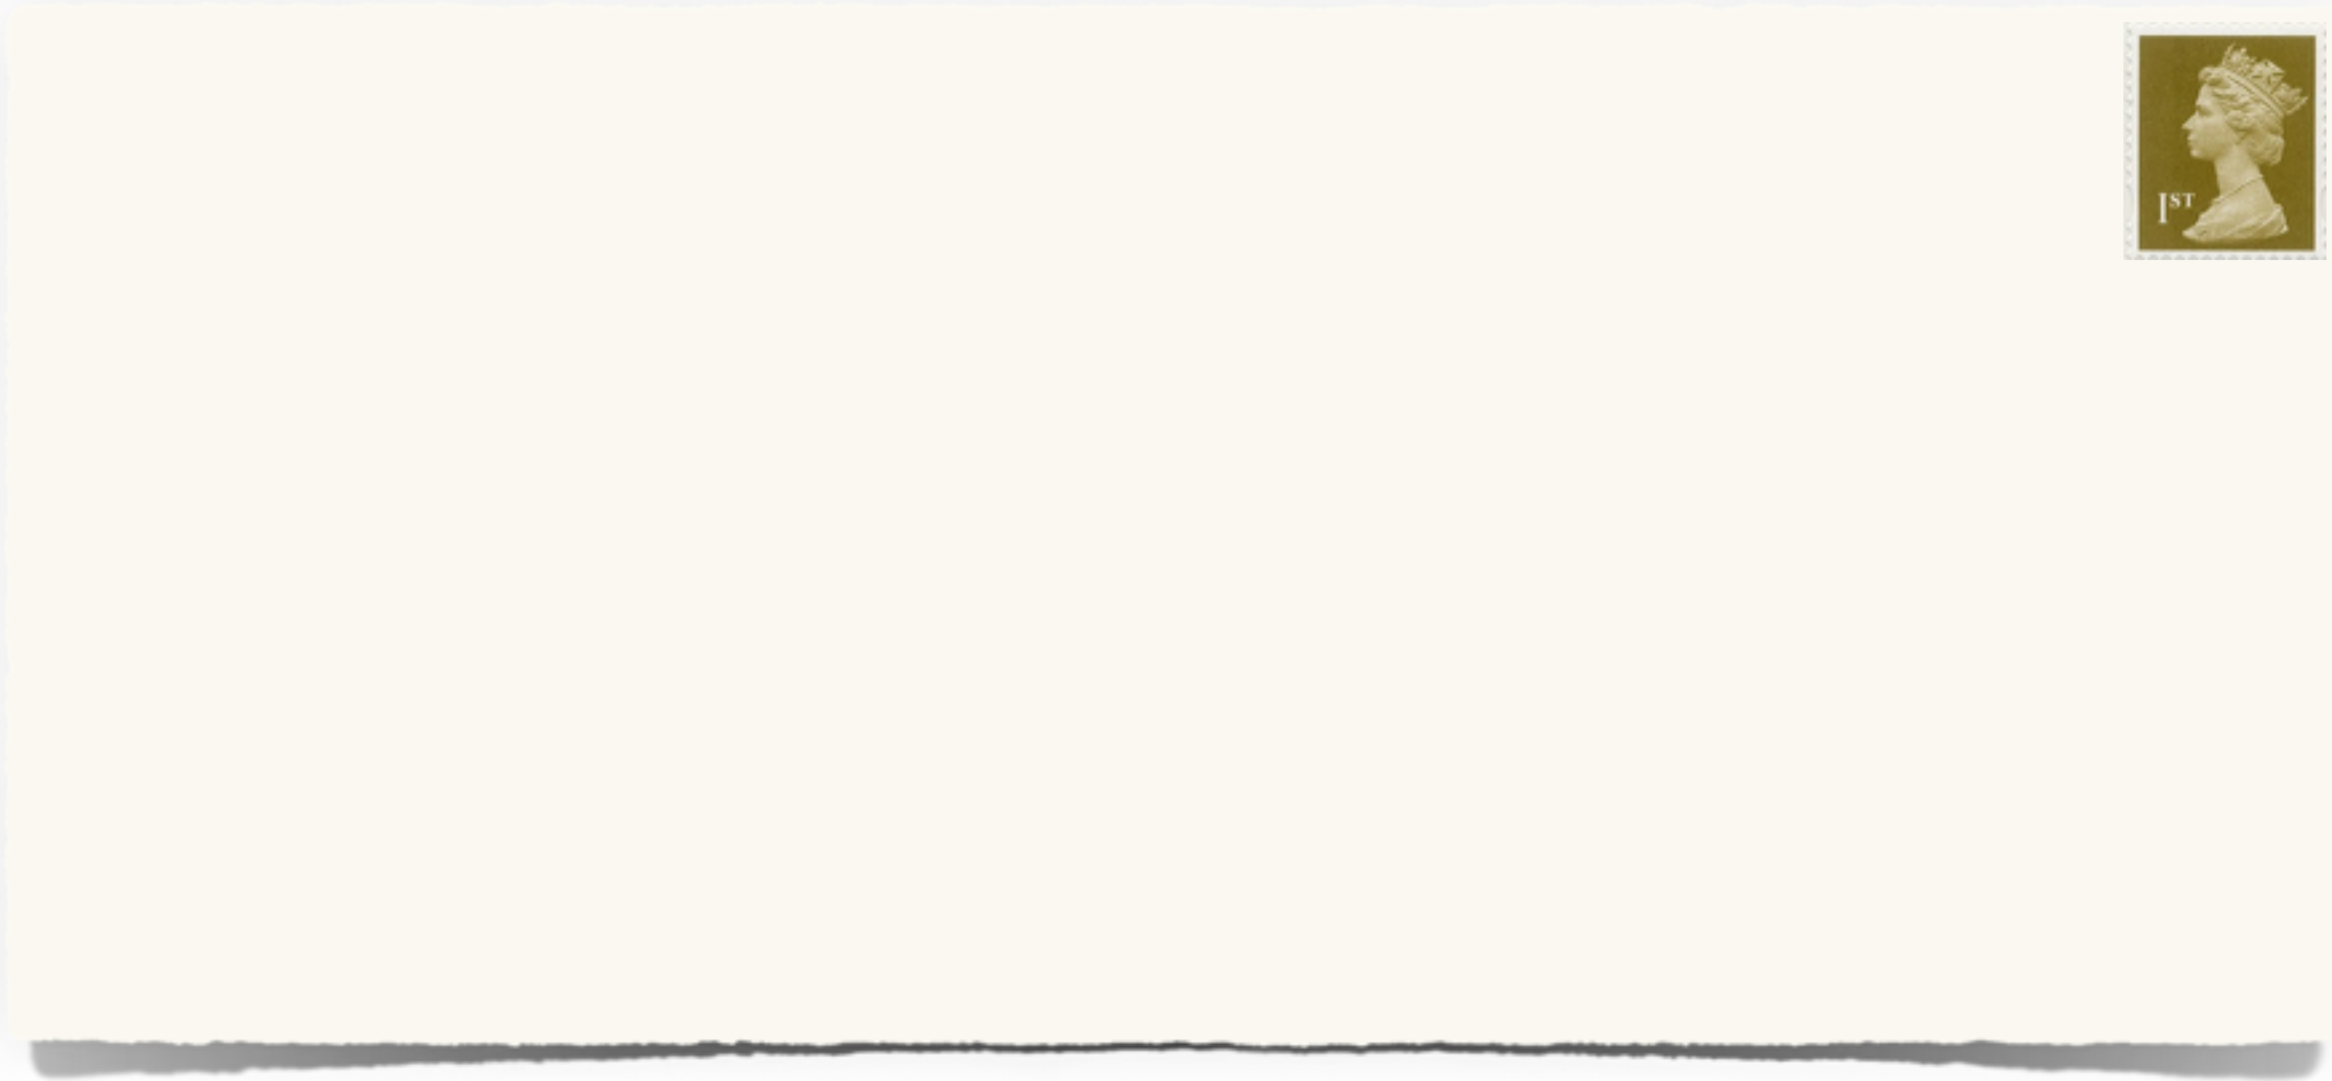
1ST

ATLAS...

10^8 channels: 1 MB /event

30 MHz bunch crossing rate

LHC year: 5×10^6 seconds

ATLAS...

10^8 channels: 1 MB /event

30 MHz bunch crossing rate

LHC year: 5×10^6 seconds

$O(10^5)$ pB per year

Storage cost: €100B

Total budget of LHC €10B

Trigger: reduce rate by 4-5 orders of magnitude in real time.

Roughly the total global internet traffic!

Selectivity

The really interesting stuff is ~13 orders of magnitude smaller.

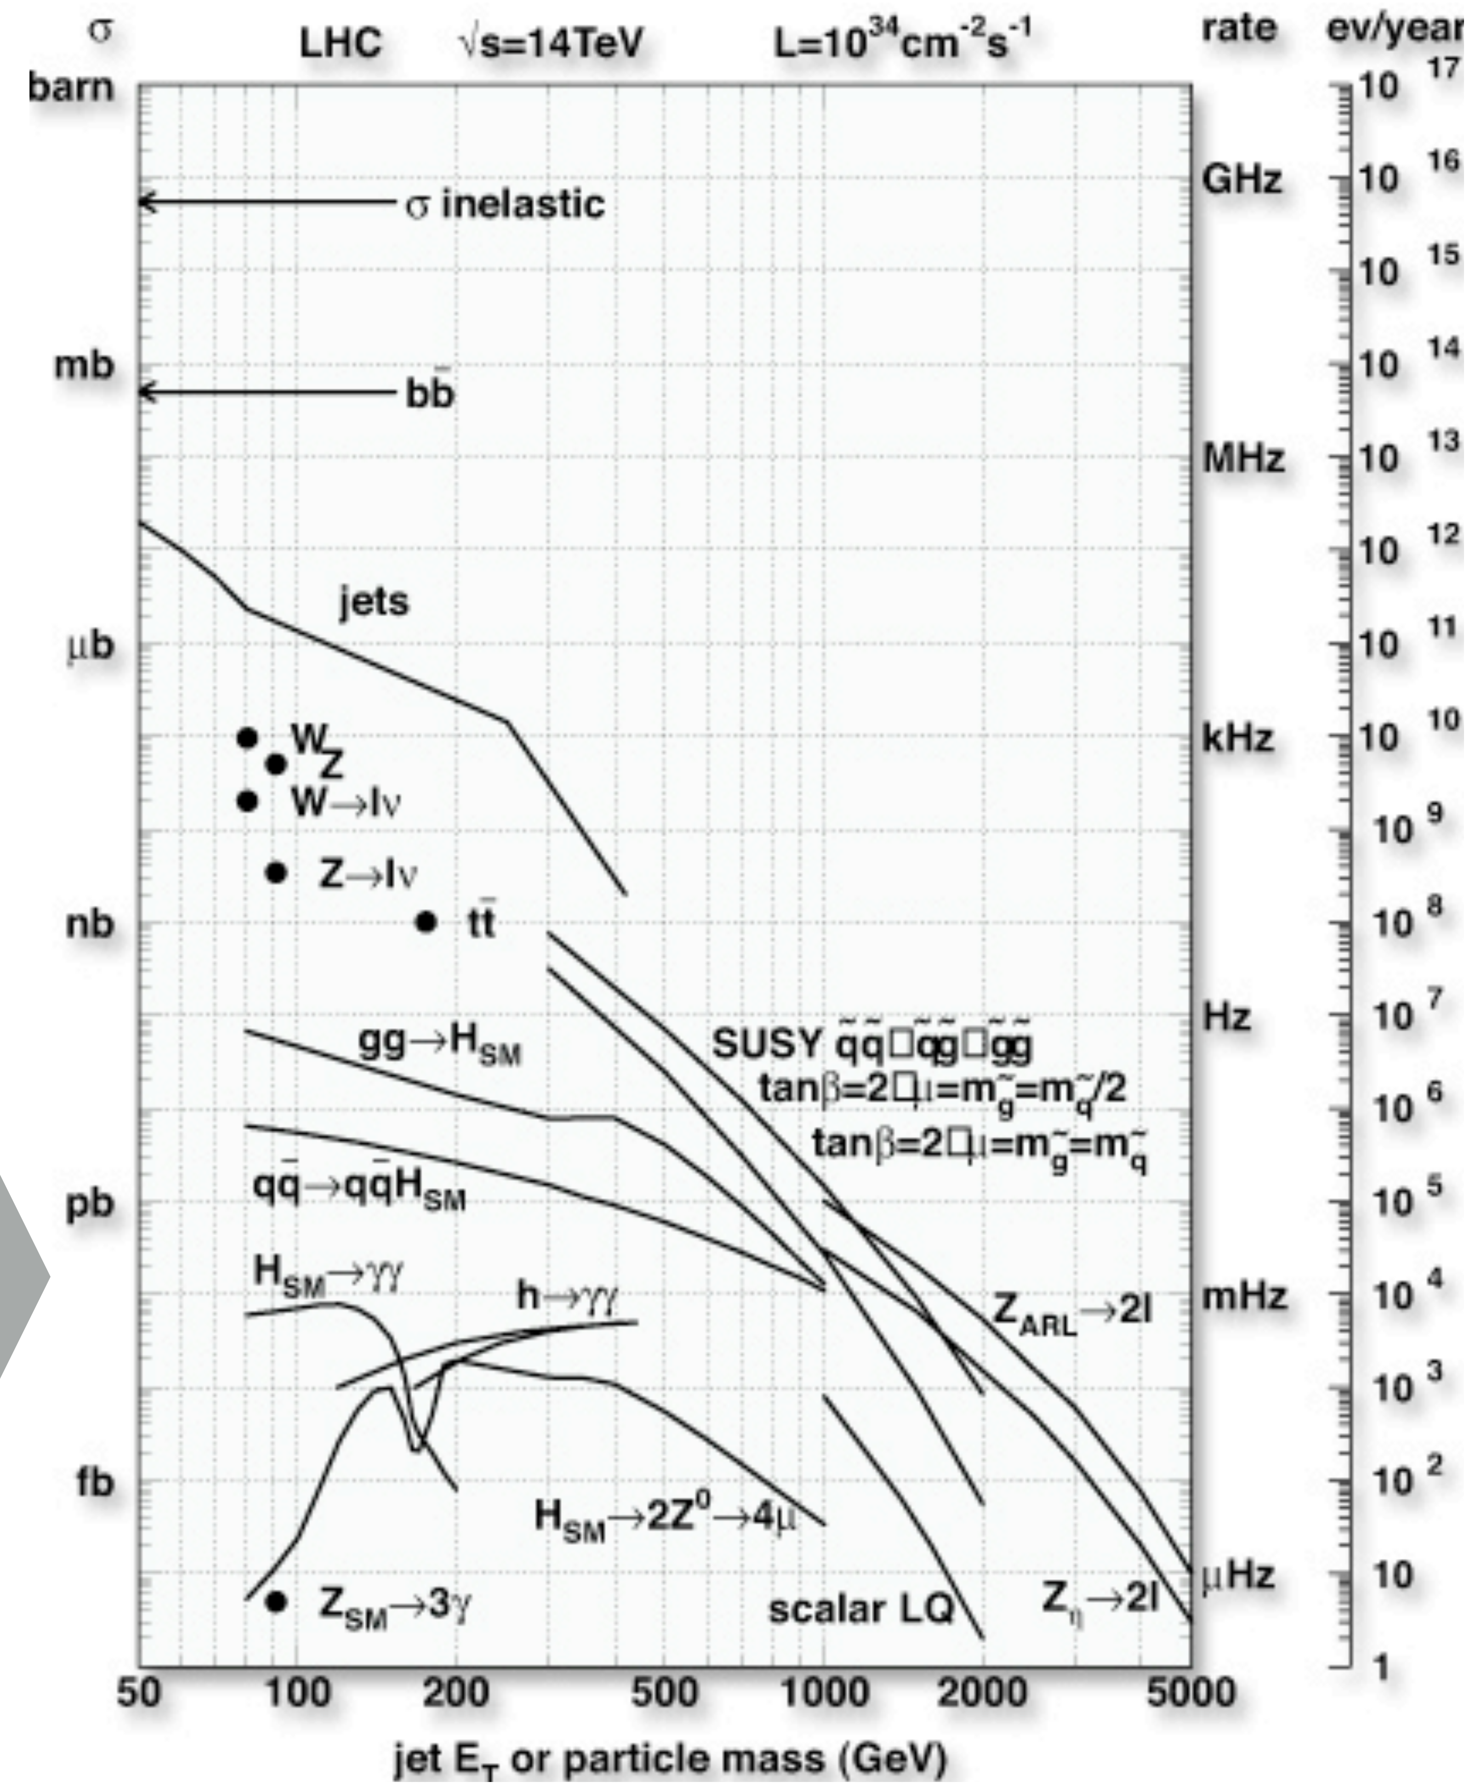

Background is low p_T

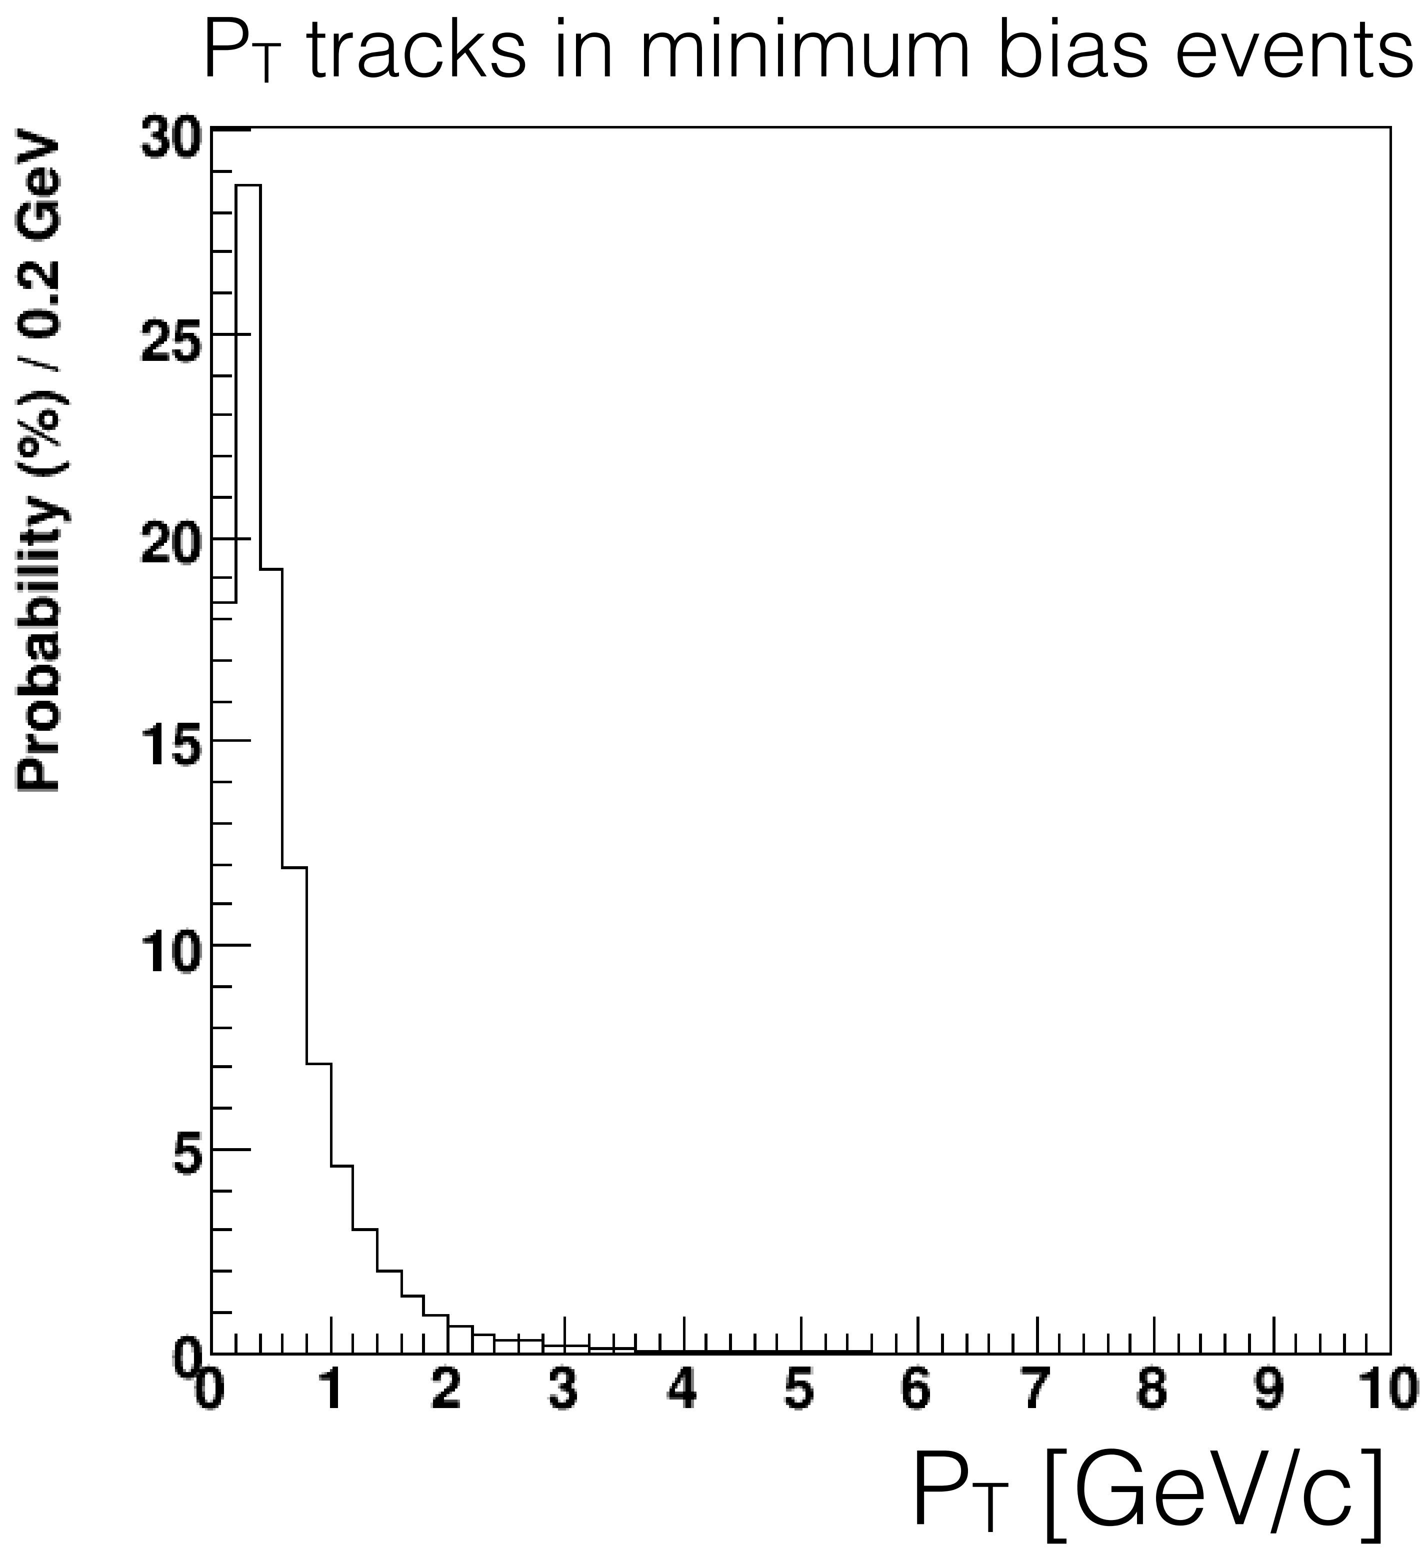

Signal is high p_T

E.g., a 3 TeV Z' decay to top pair

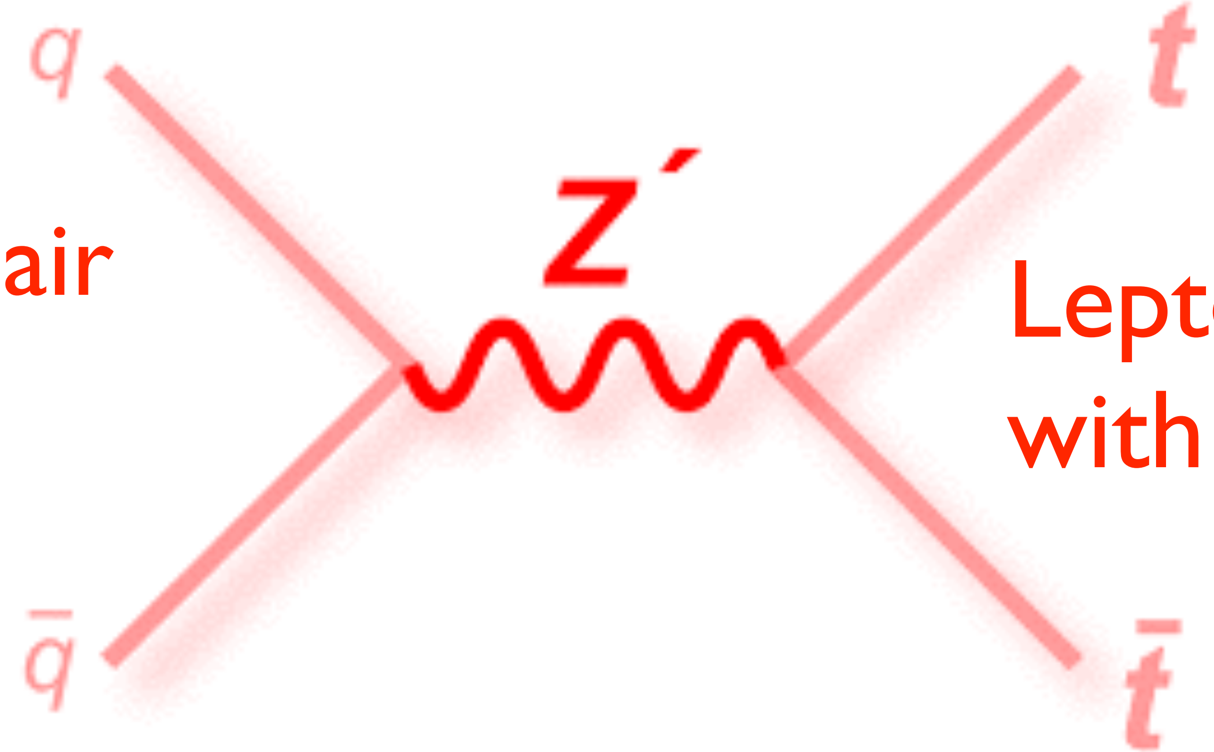

Leptons, jets with $p_T > 100$ GeV,

Can trigger **inclusively** on single high p_T leptons, jets, missing E_T, or simple combinations of these

Only part of the decay is needed to trigger the event, so large physics progamme with small set of triggers.

How to discover new physics at the LHC

Direct approach

Works if M ≪ 14 TeV

Indirect approach

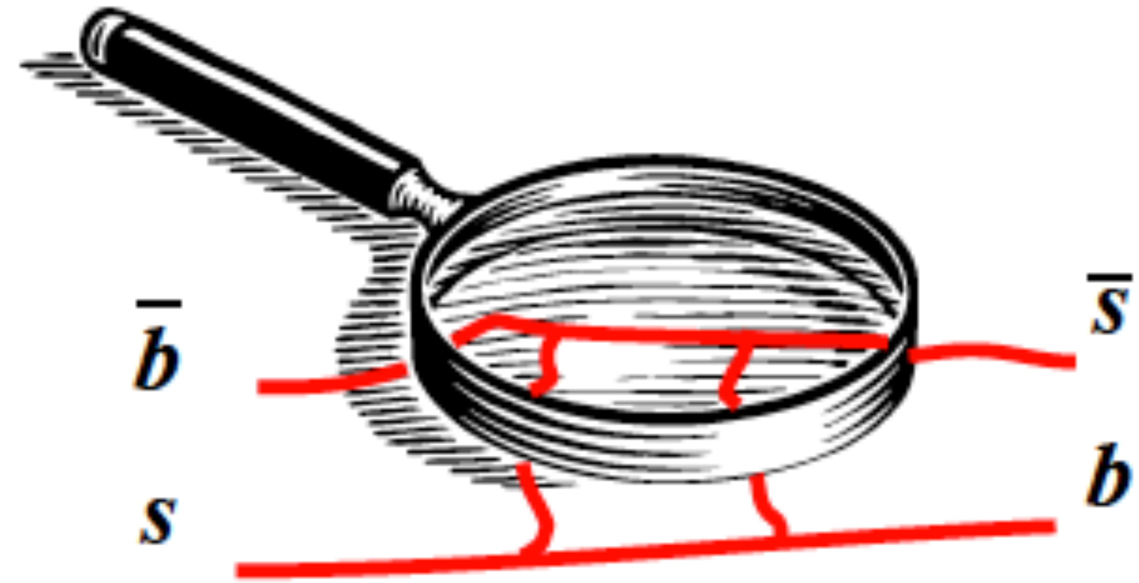

With high enough precision (theory and experiment), can probe any mass scale

How to discover new physics at the LHC

Direct approach

Works if M ≪ 14 TeV

Inclusive triggering on high p_T particles

Indirect approach

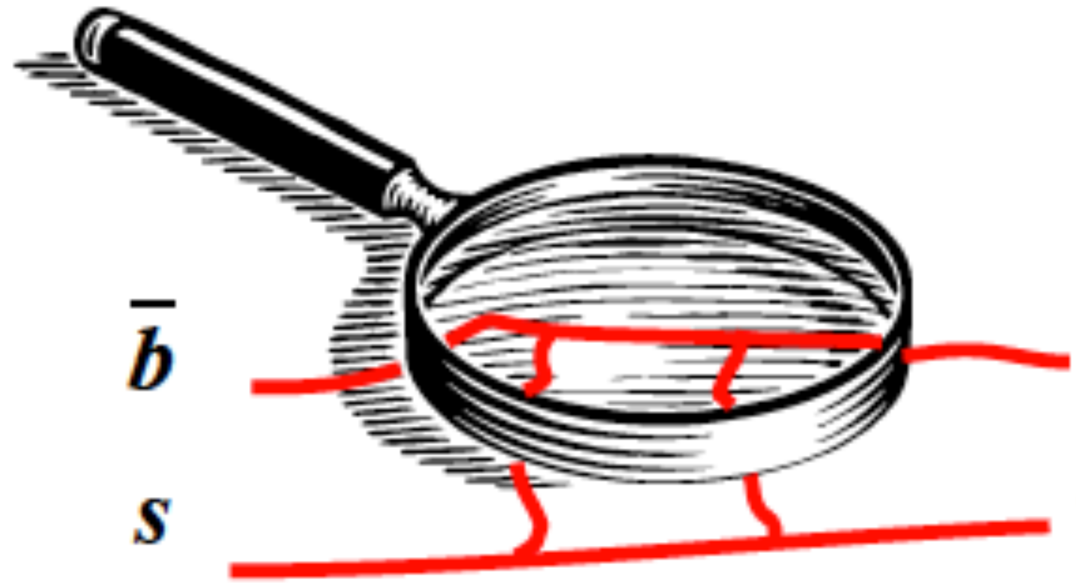

With high enough precision (theory and experiment), can probe any mass scale

Mixture of inclusive and exclusives…

How to discover new physics at the LHC

Direct approach

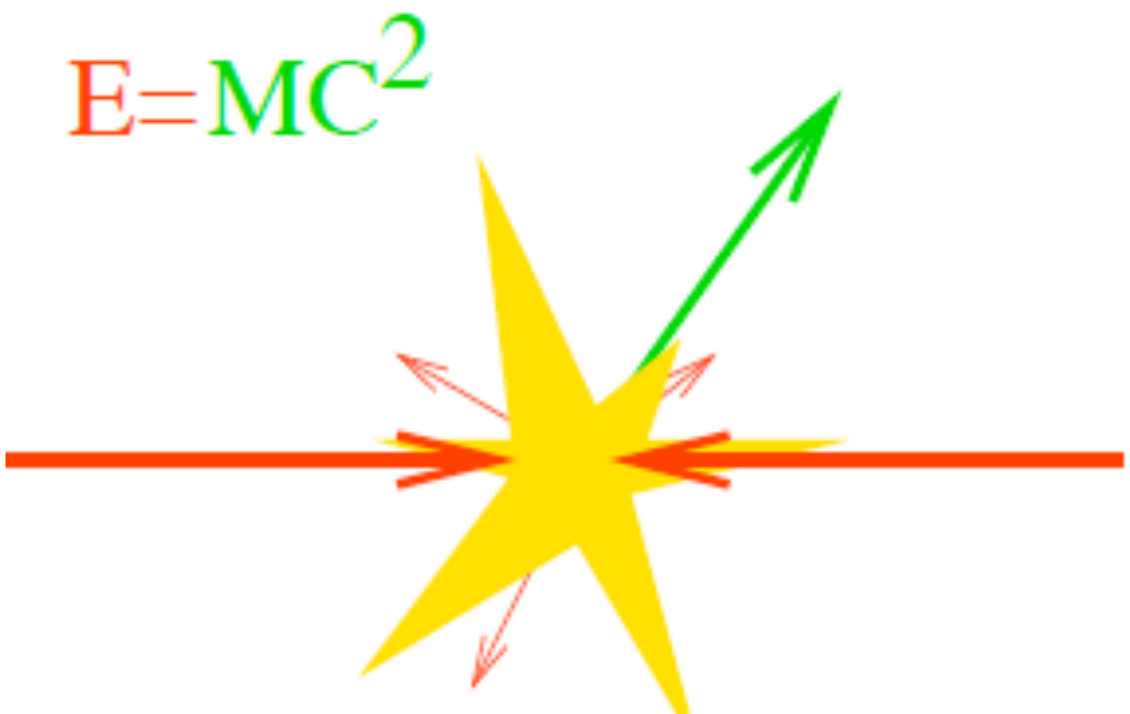

Works if M ≪ 14 TeV

Indirect approach

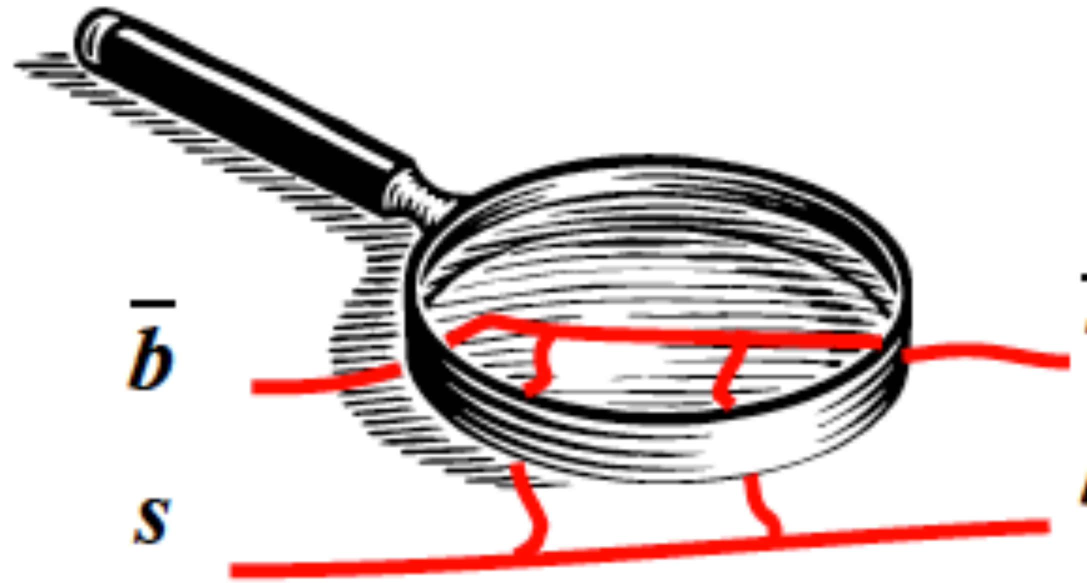

With high enough precision (theory and experiment), can probe any mass scale

Inclusive triggering on high p_T particles

Mixture of inclusive and exclusives…

MLHEP in final analysis

MLHEP in trigger

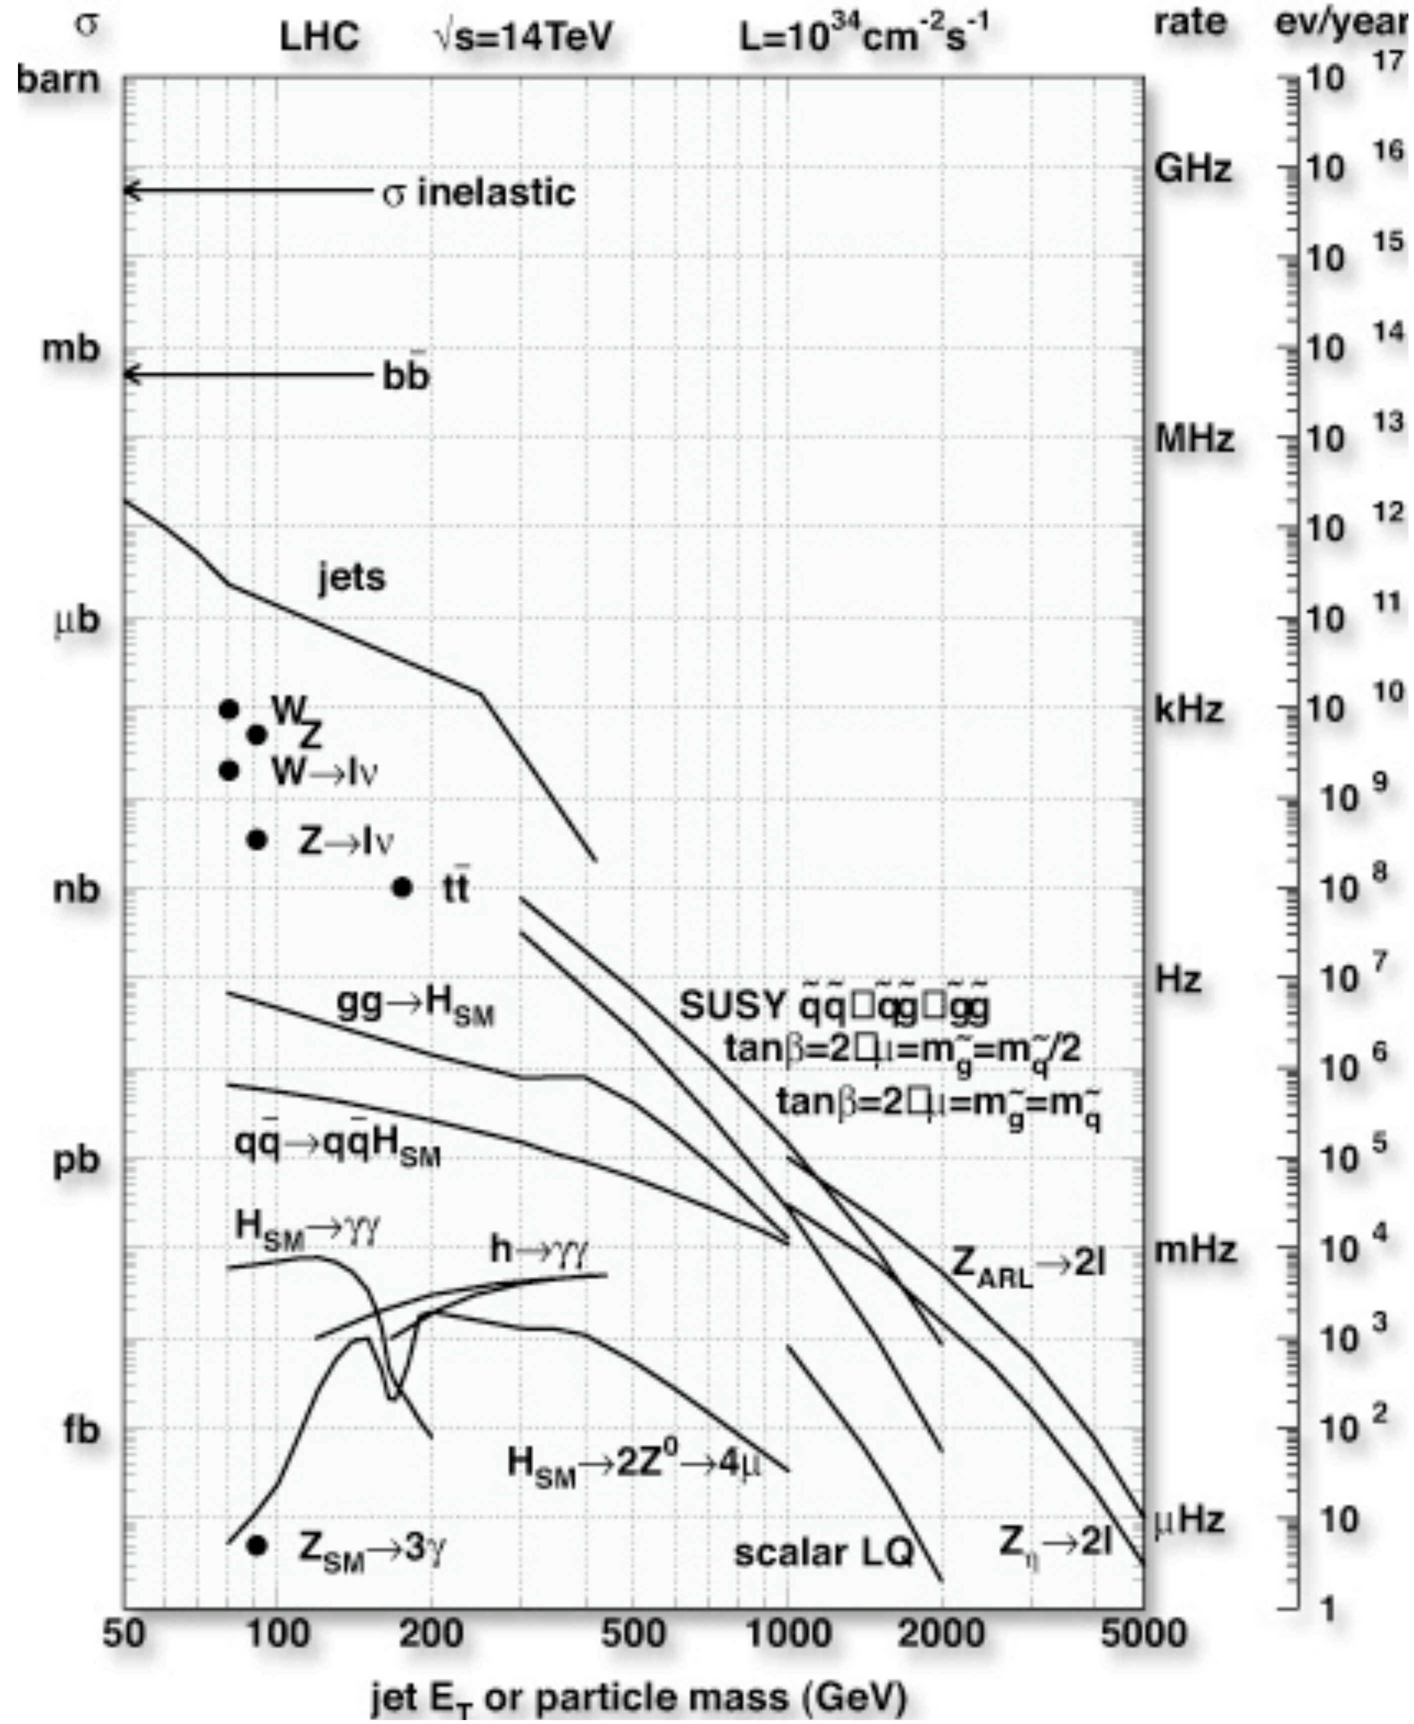
σ
barn
LHC
√s=14TeV
L=10^34cm^-2s^-1
rate
ev/year
σ inelastic
mb
bb̄
jets
μb
W
Z
W→lν
Z→lν
nb
tt̄
gg→H_SM
SUSY q̃q̃ q̃g̃ g̃g̃
tanβ=2 μ=m_g̃=m_q̃/2
tanβ=2 μ=m_g̃=m_q̃
qq̄→qq̄H_SM
pb
H_SM→γγ
h→γγ
Z_ARL→2l
fb
H_SM→2Z^0→4μ
Z_SM→3γ
scalar LQ
Z_η→2l
GHz
MHz
kHz
Hz
mHz
μHz
50
100
200
500
1000
2000
5000
jet E_T or particle mass (GeV)

Charm physics

And it is interesting!

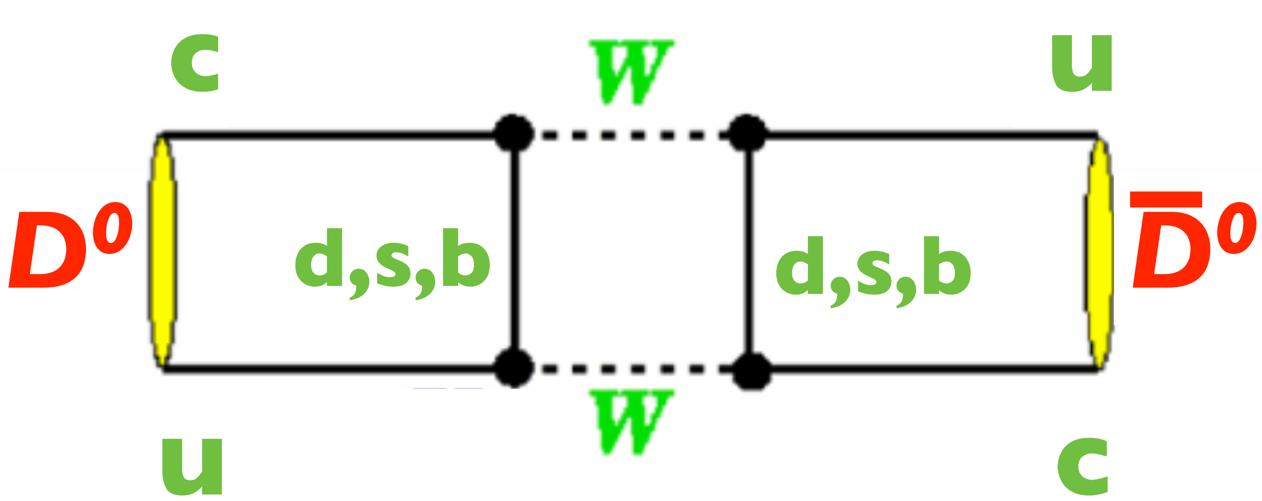

σ_c

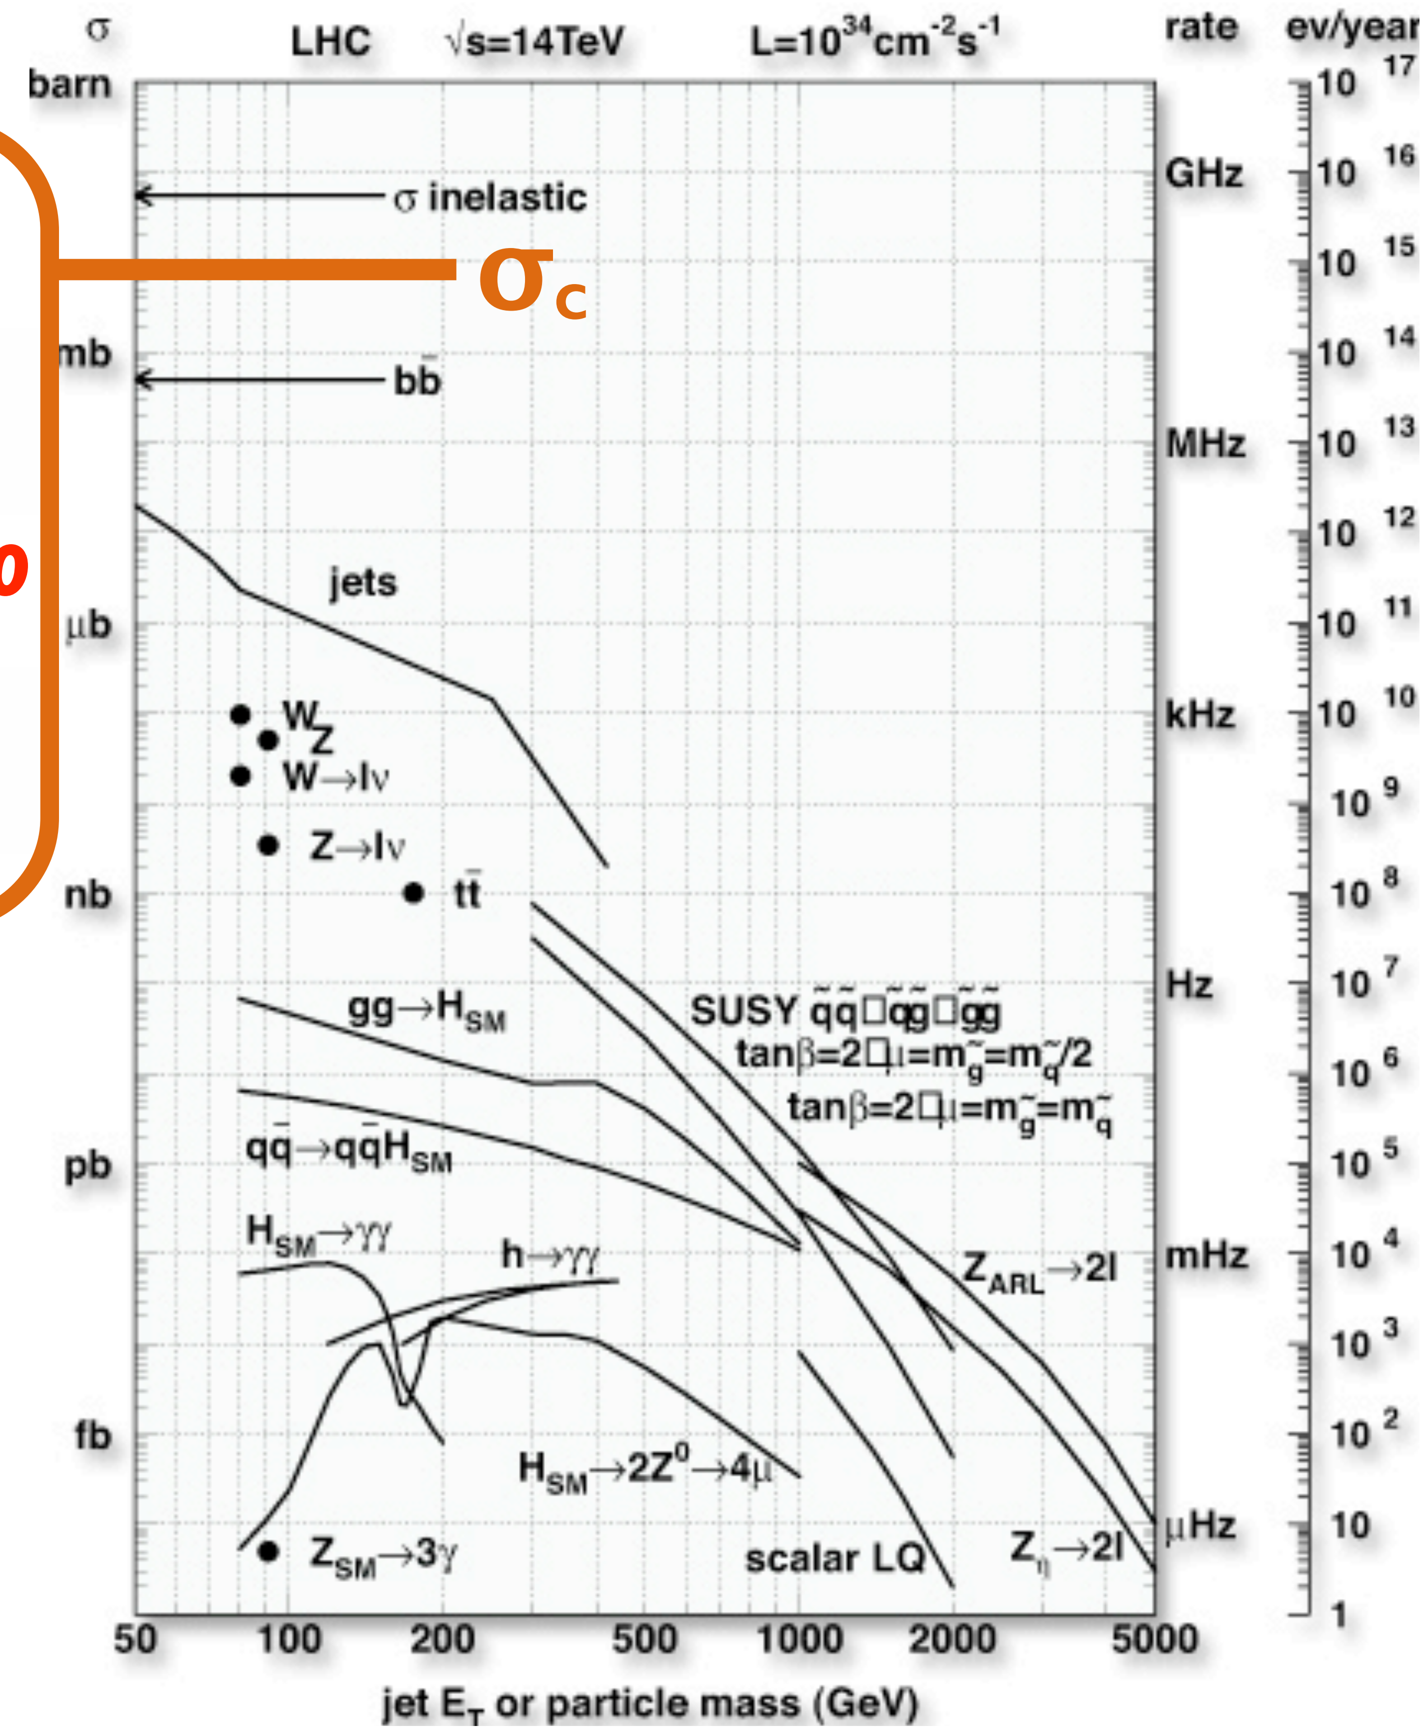

How much charm?

LHCb

σ_c around 2 mb

Lumi: 4 x 10^{32} $cm^{-2}s^{-1}$

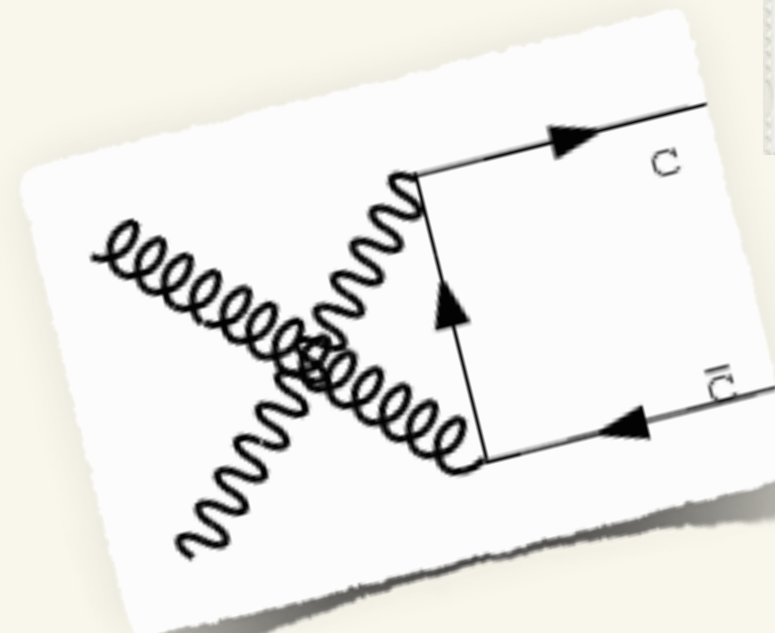

Around 1 MHz of charm!

Not all charm is interesting though…

Example of neutral D meson mixing

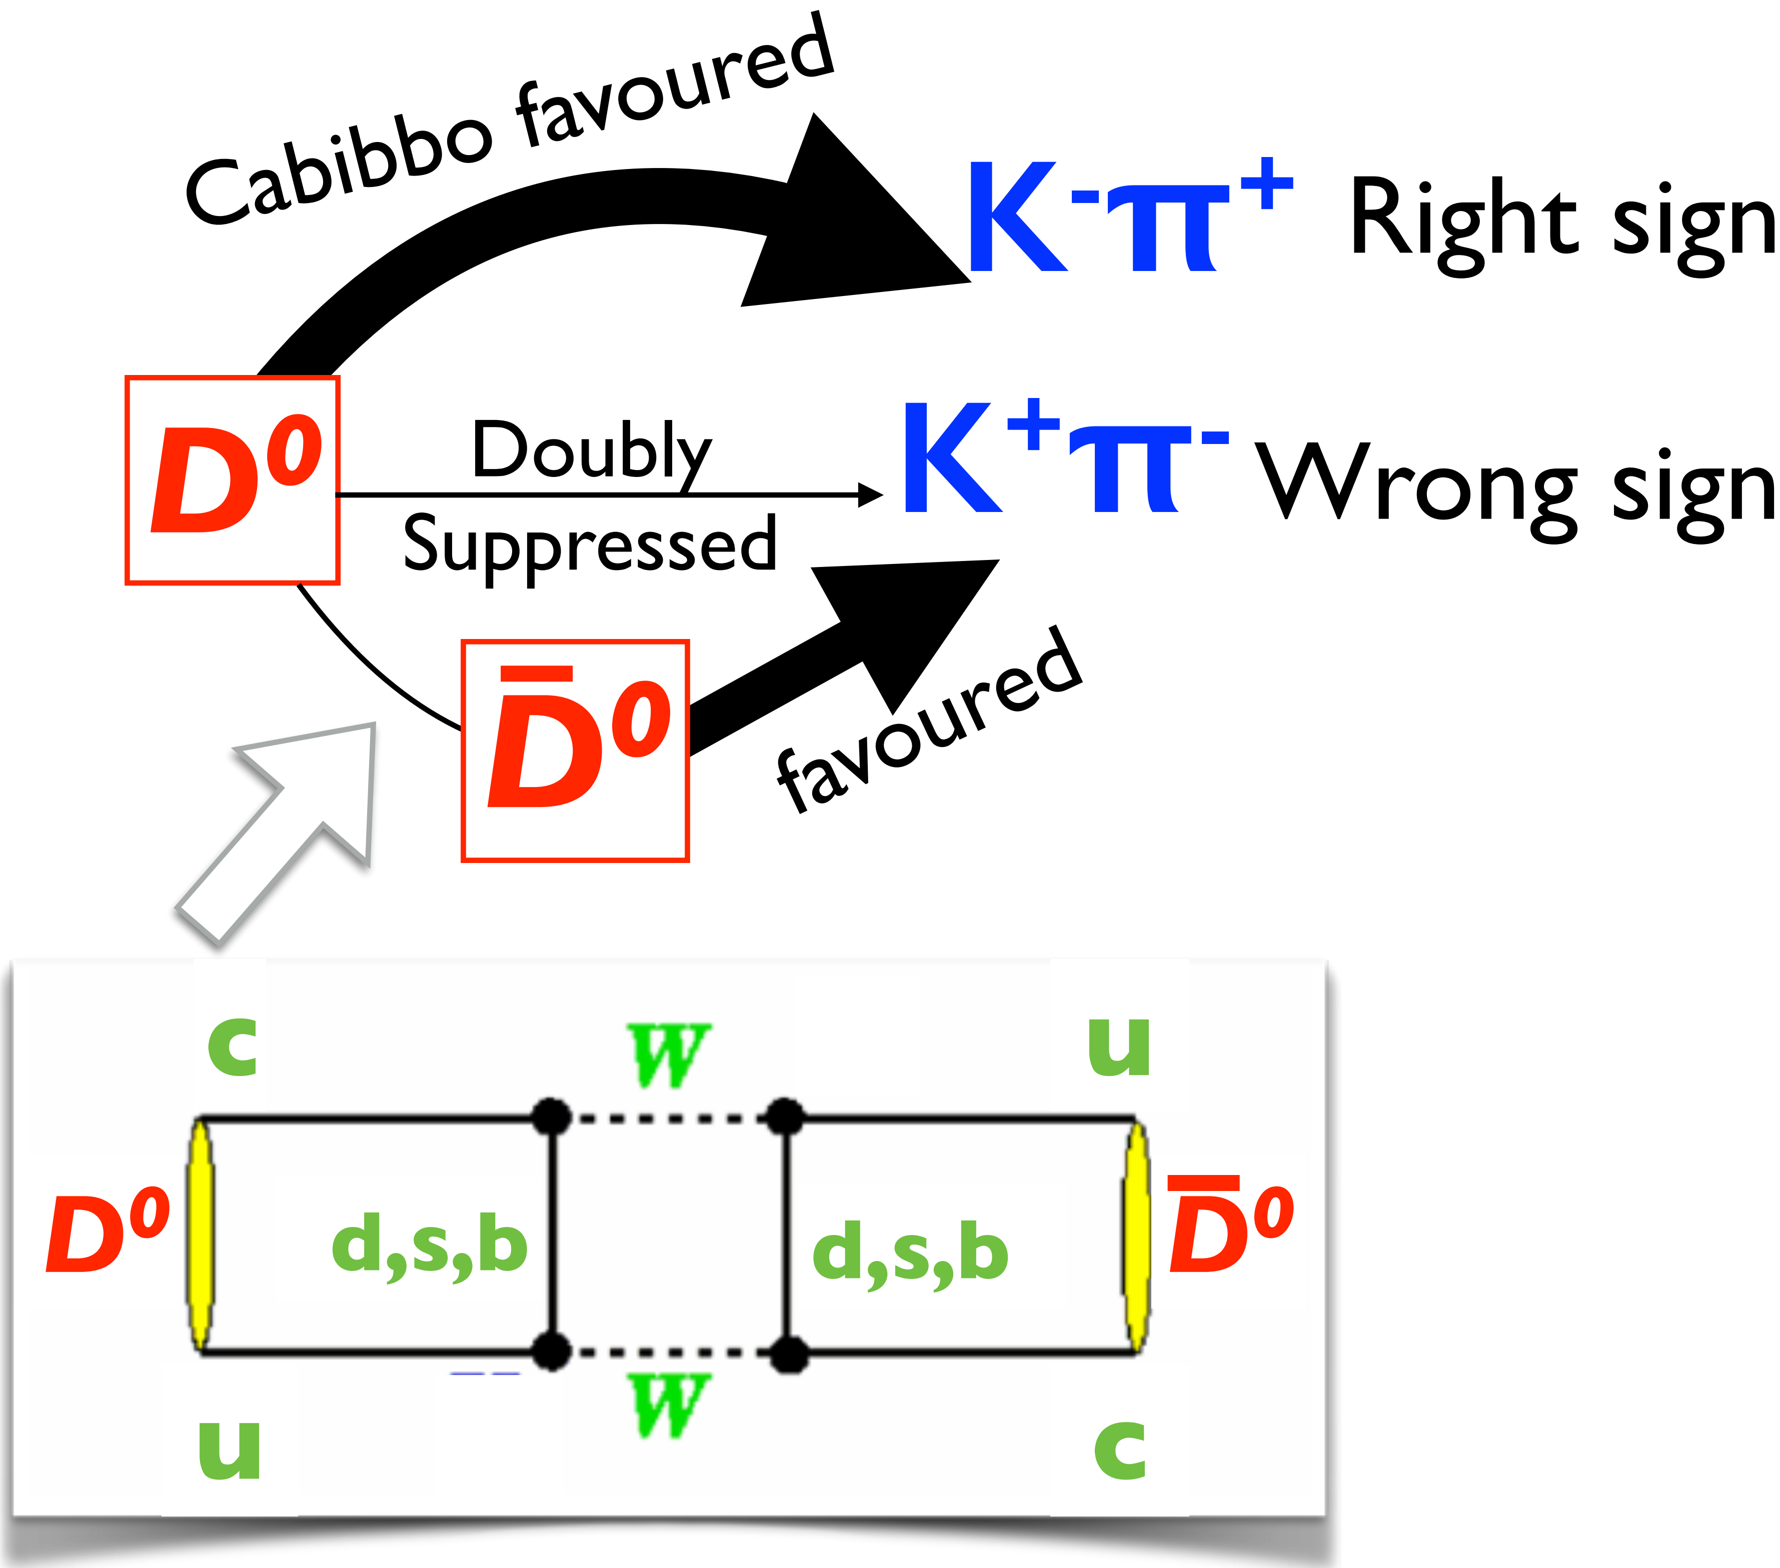

Example of neutral D meson mixing

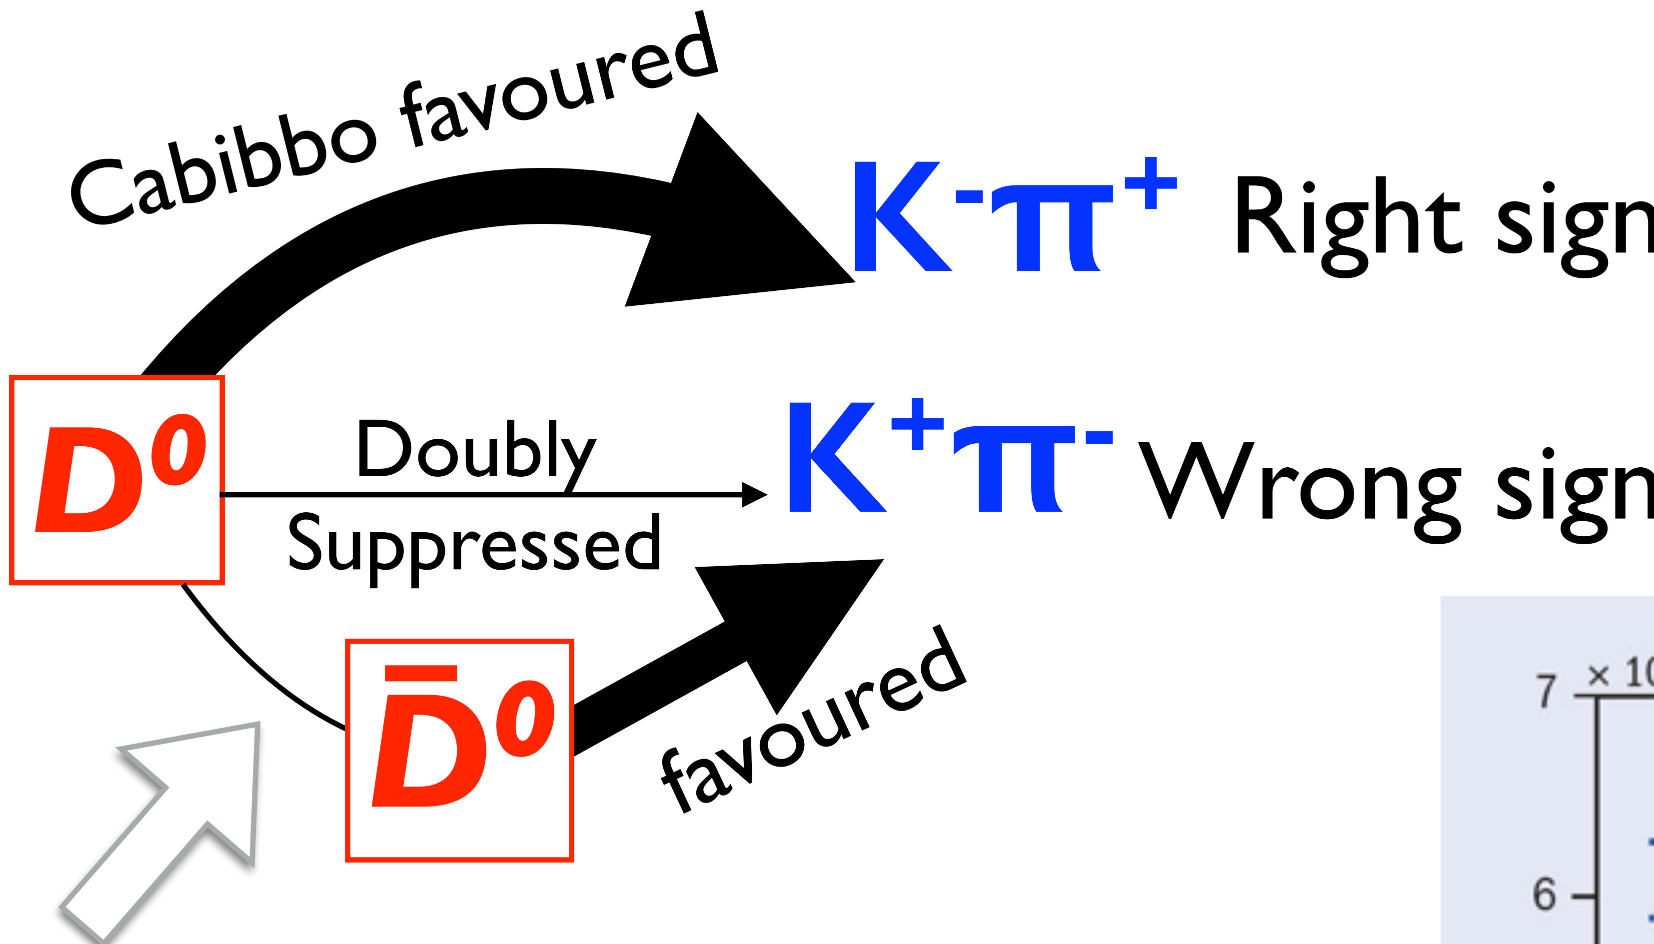

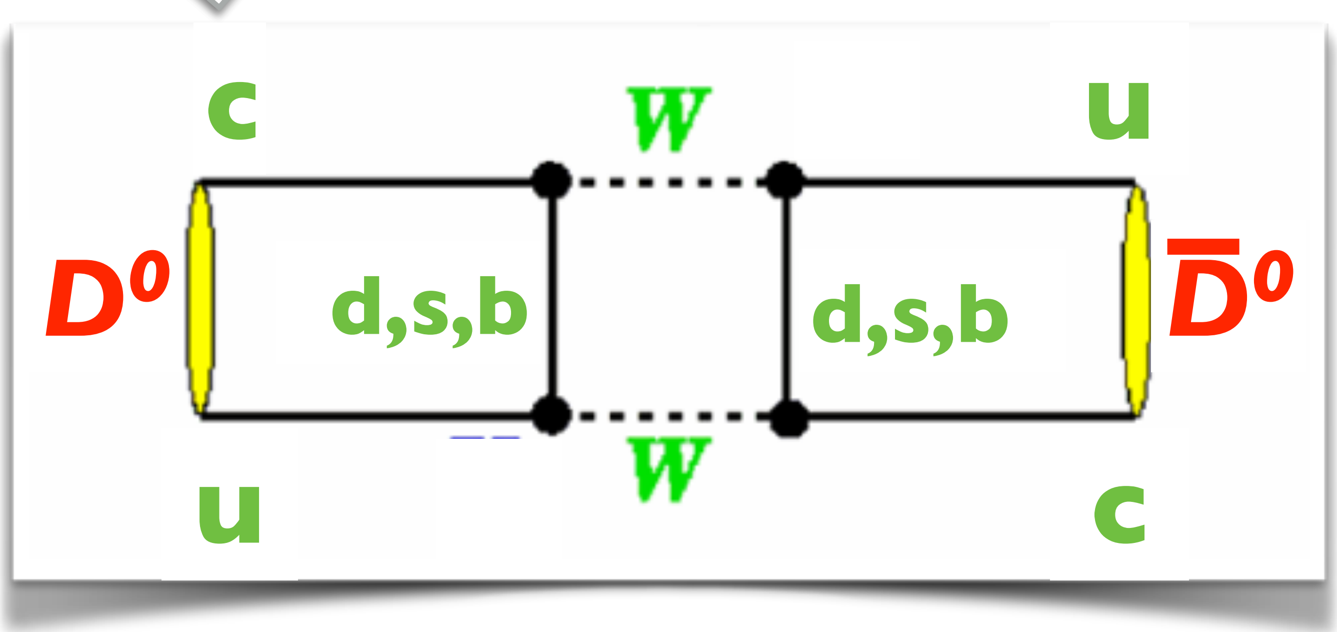

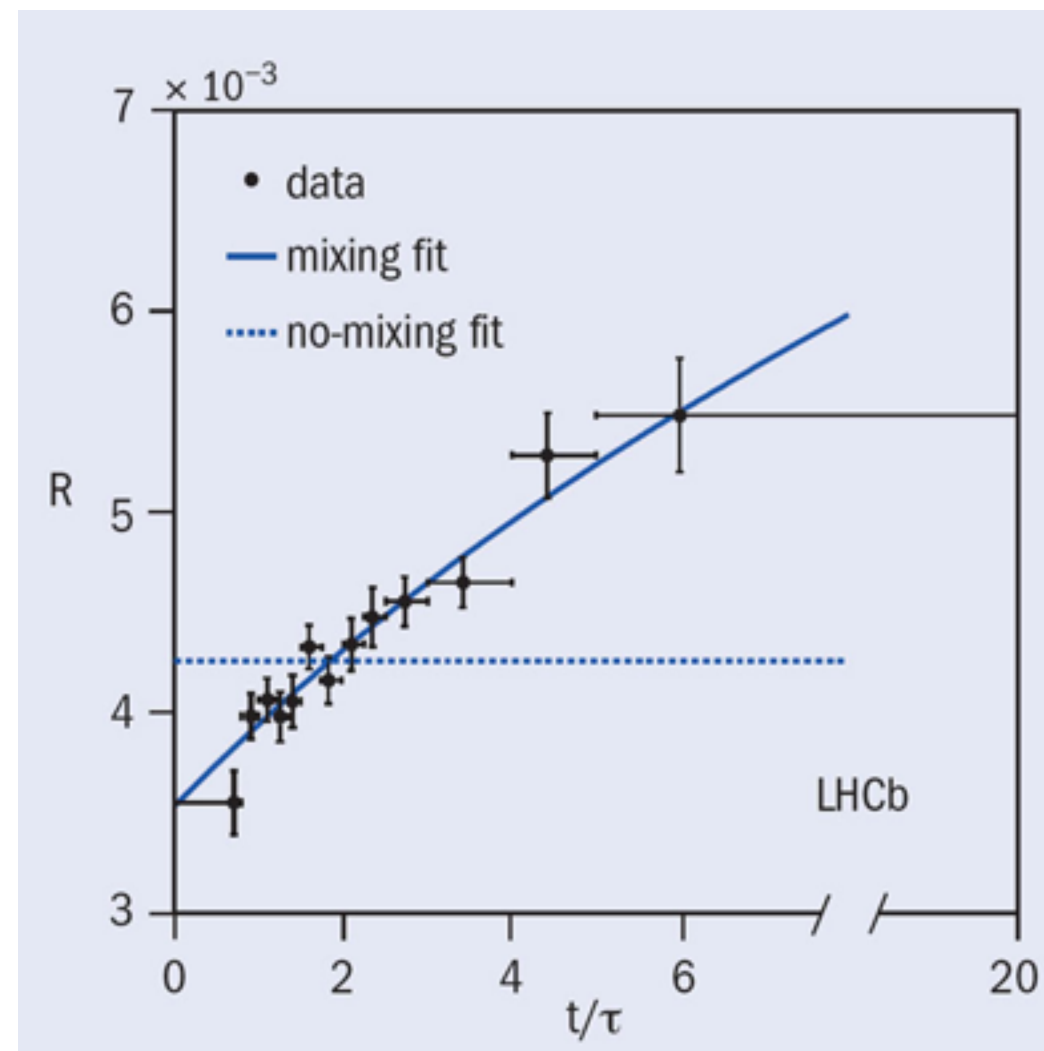

Exclusive approach

Out of our 1 MHz of charm, we needed…

~100 Hz right-sign ~0.1 Hz wrong-sign

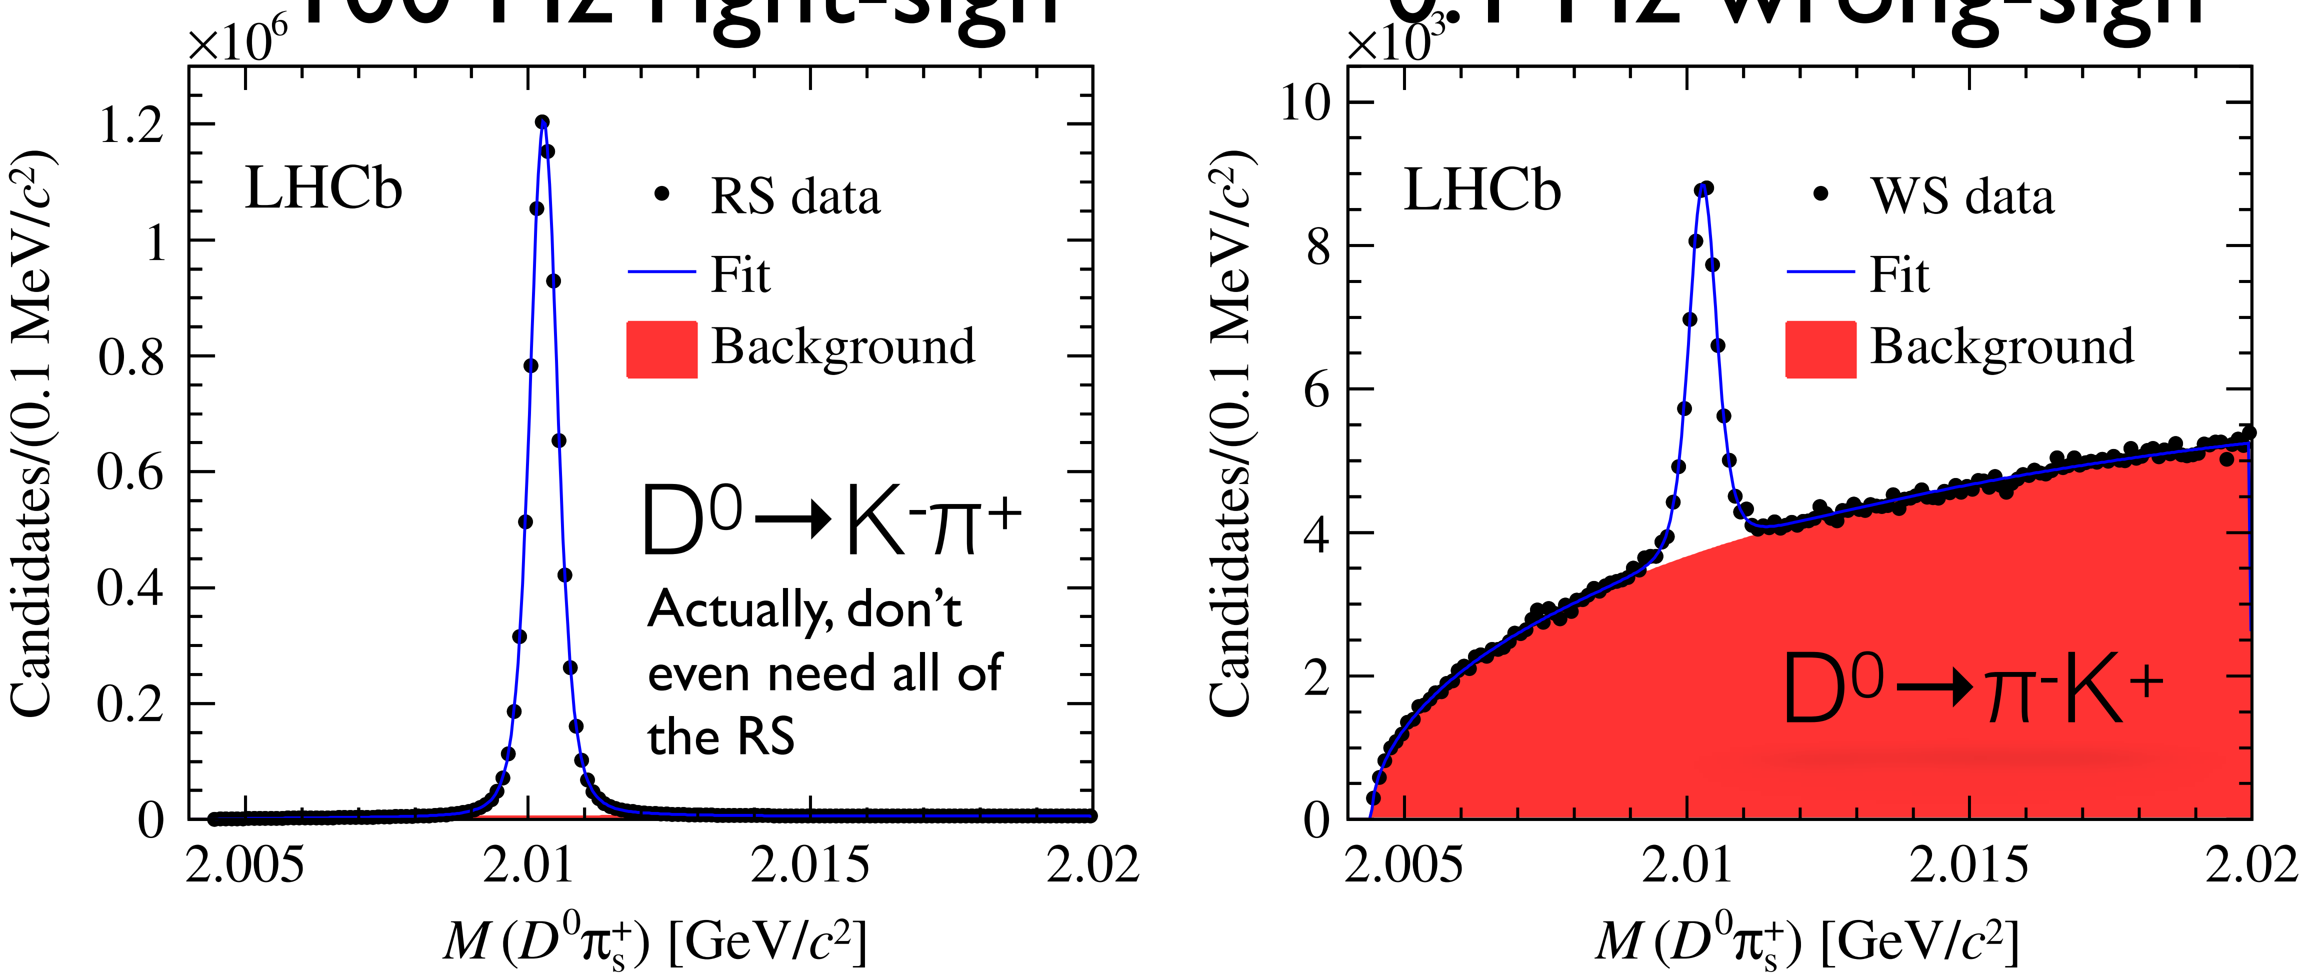

In reality LHCb has ~100 exclusive charm lines for different physics interests with 2-4 kHz

High p_T physics

Charm physics

Low level triggers

Another bottleneck

- Ideally we would wait until we have all of the best possible information to hand, before deciding which events to keep.
- Offline event reconstruction: 1 second with CPU core
- Would need 40 million processes

few billion €

(€1k/PC) * (40e6 procs) / (40 cores)

And imagine moving all that data

- ATLAS/CMS @ 40 MHz means 40 TB/s
- And you need to assemble complete events to send to a single PC.

Forget it!

…until the LHCb (which has relatively small events) upgrade (2019).

Time of flight

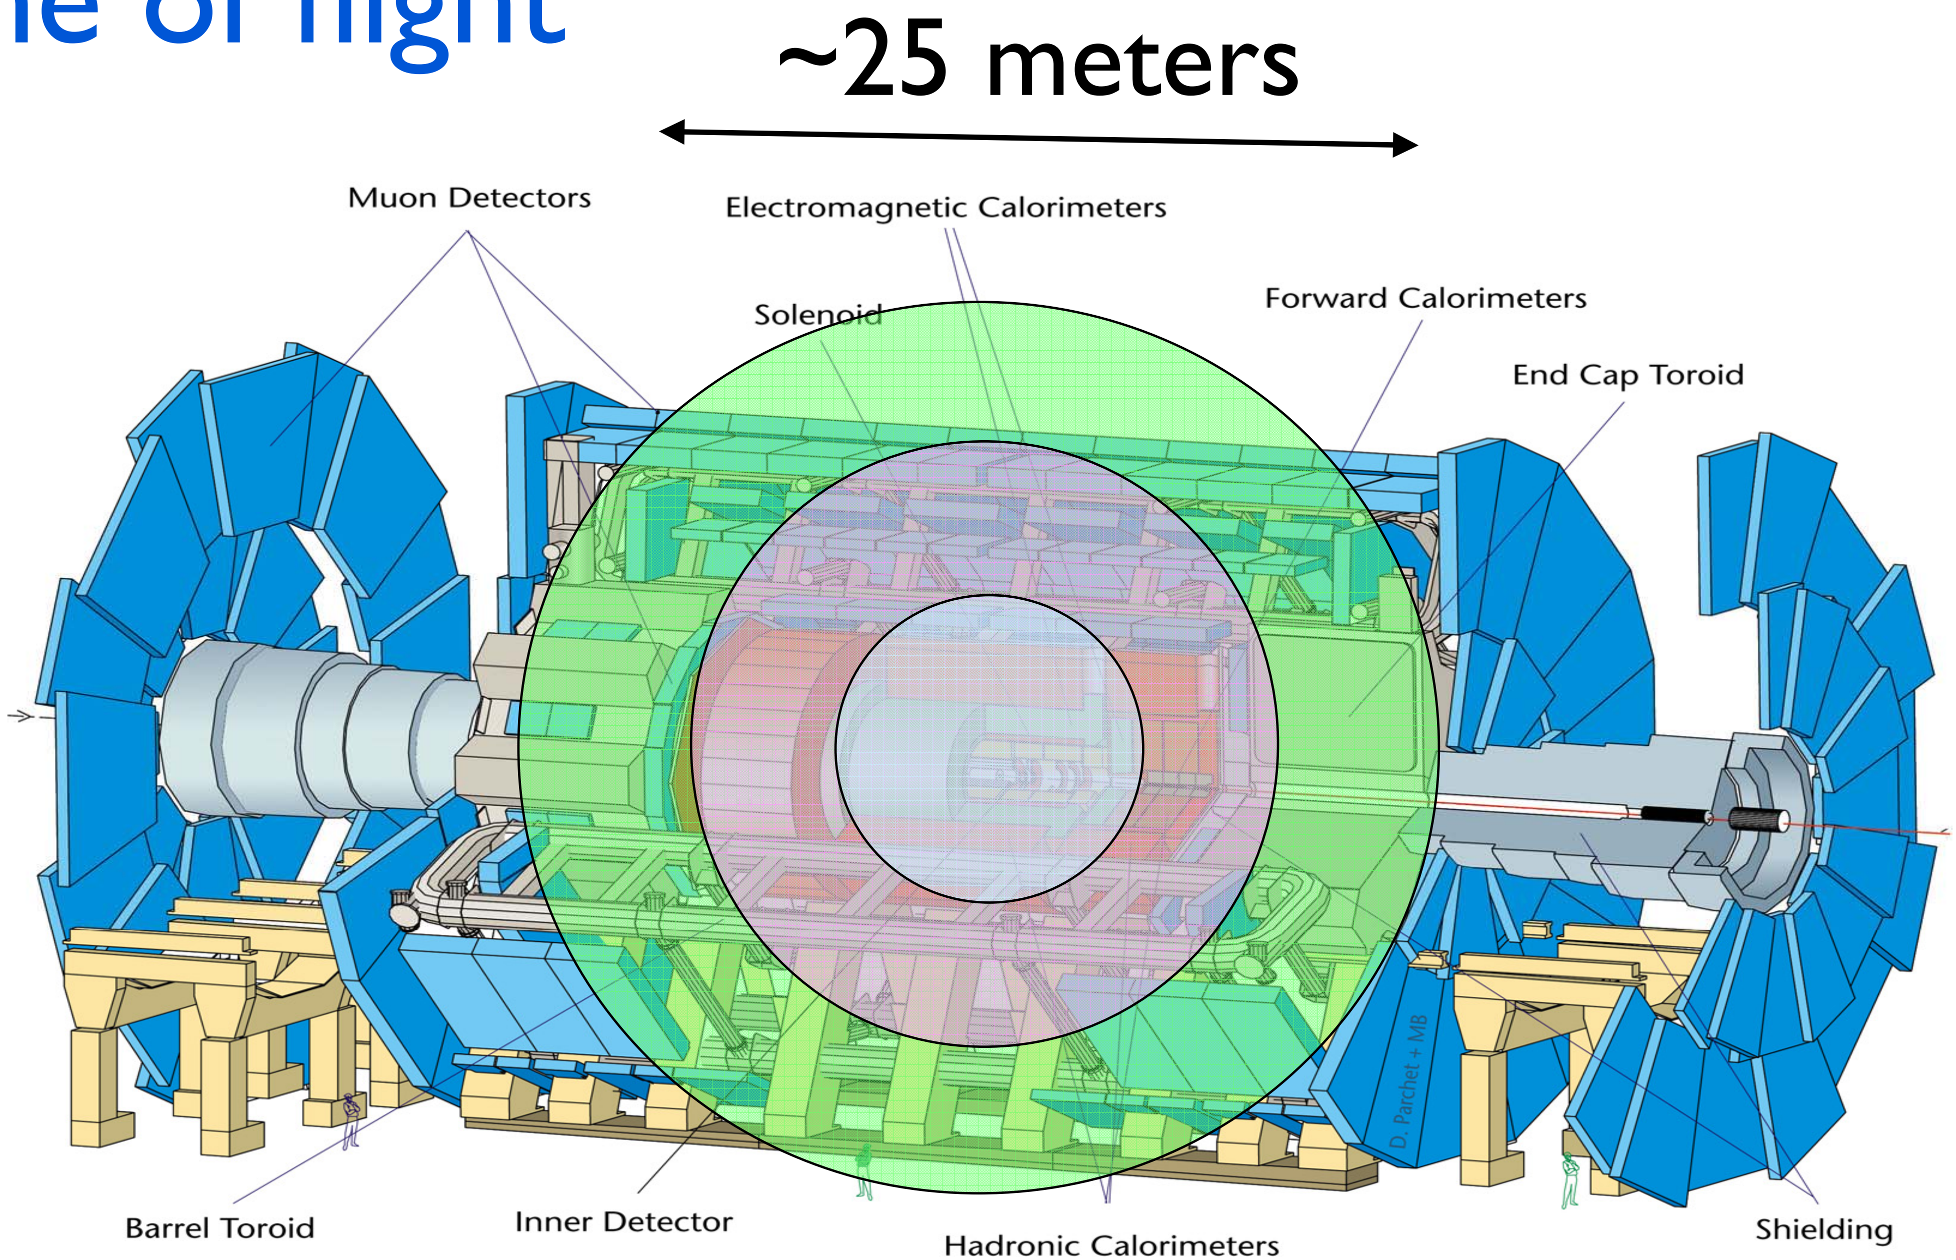

In 25ns, light travels 7.5m

Already means that we need to buffer many bunch crossings while we wait for a decision.

Real trigger and DAQ

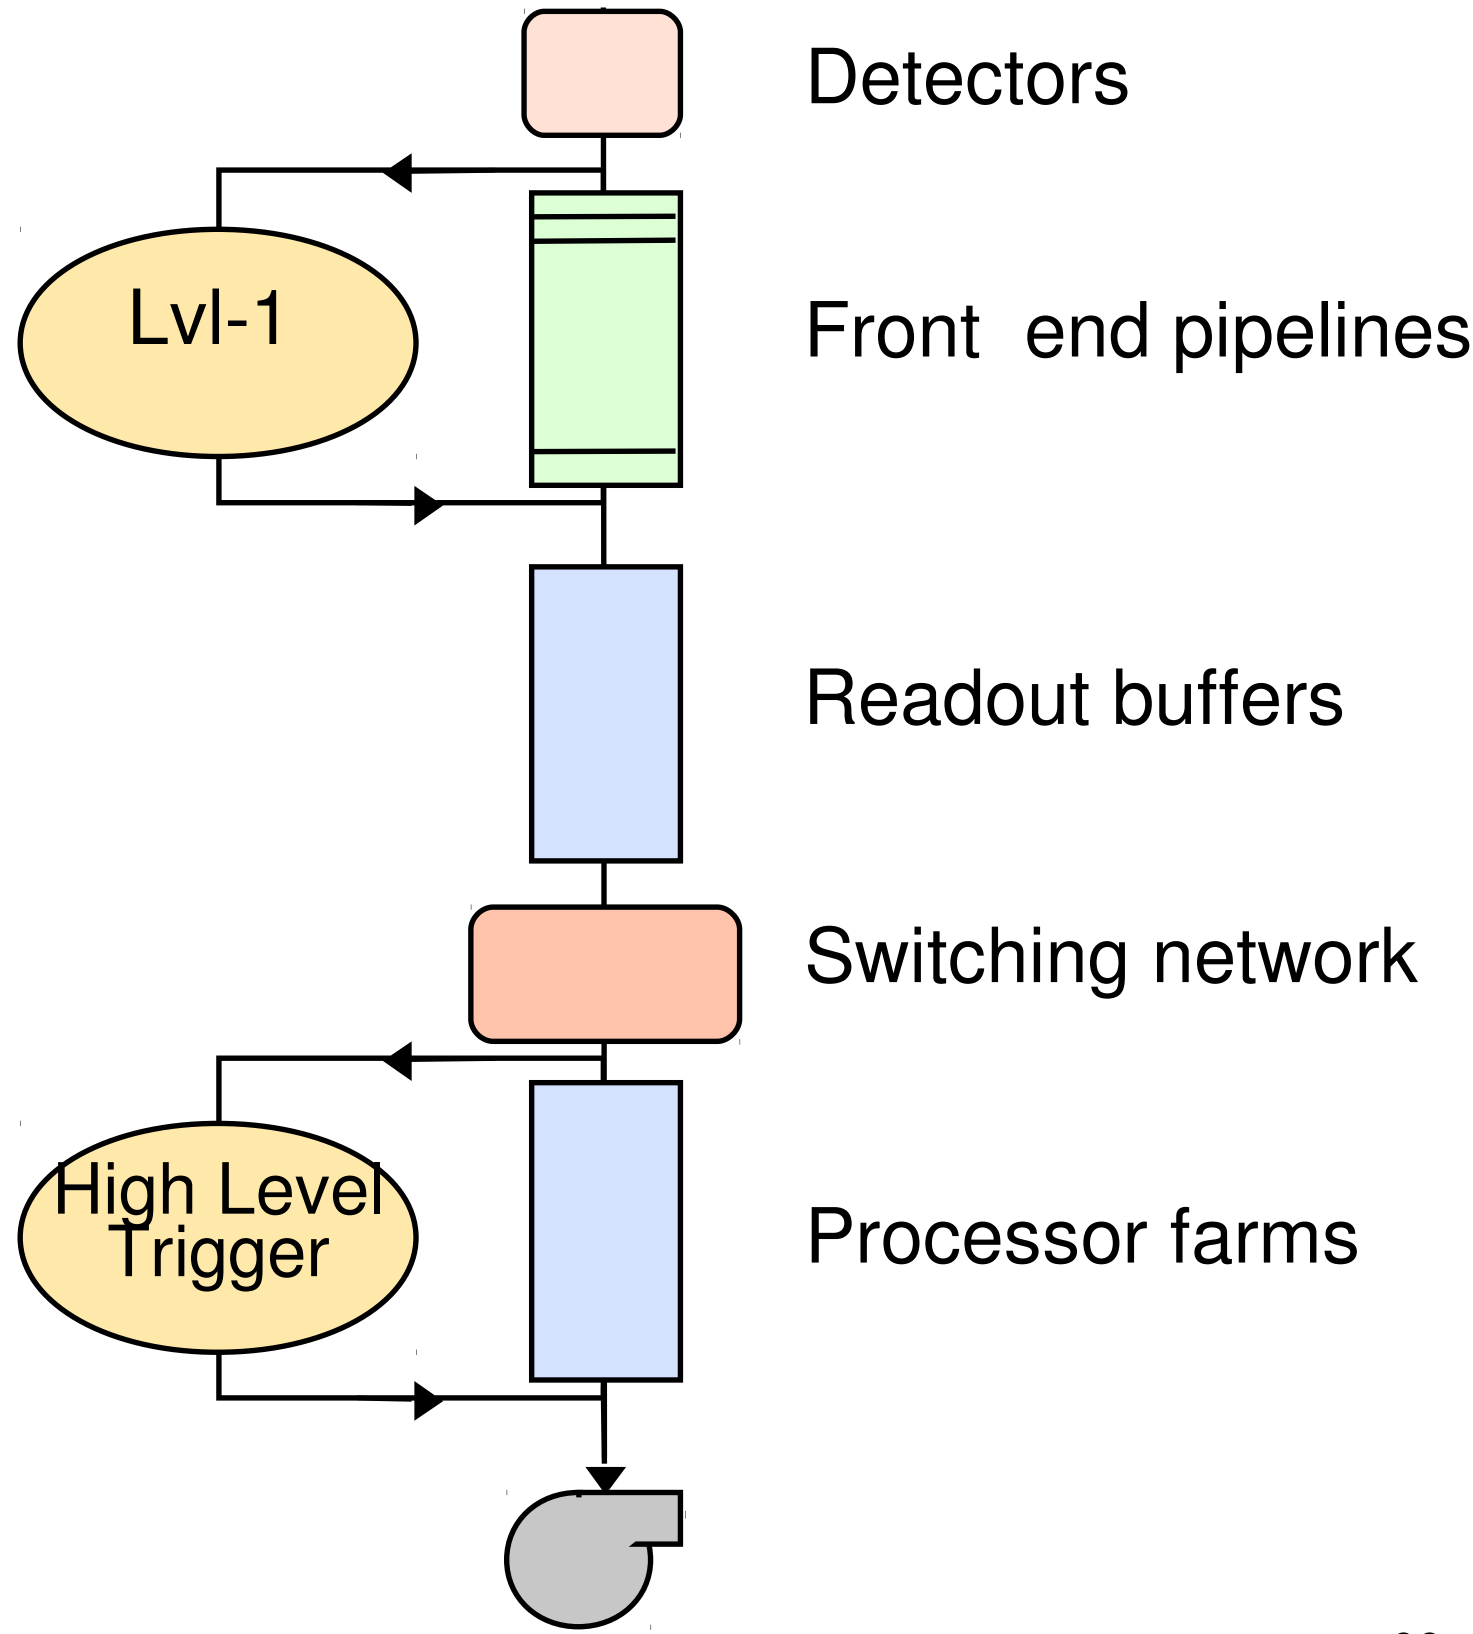

Real trigger and DAQ

- Low level trigger to reduce the event rate to 0.1 - 1 MHz

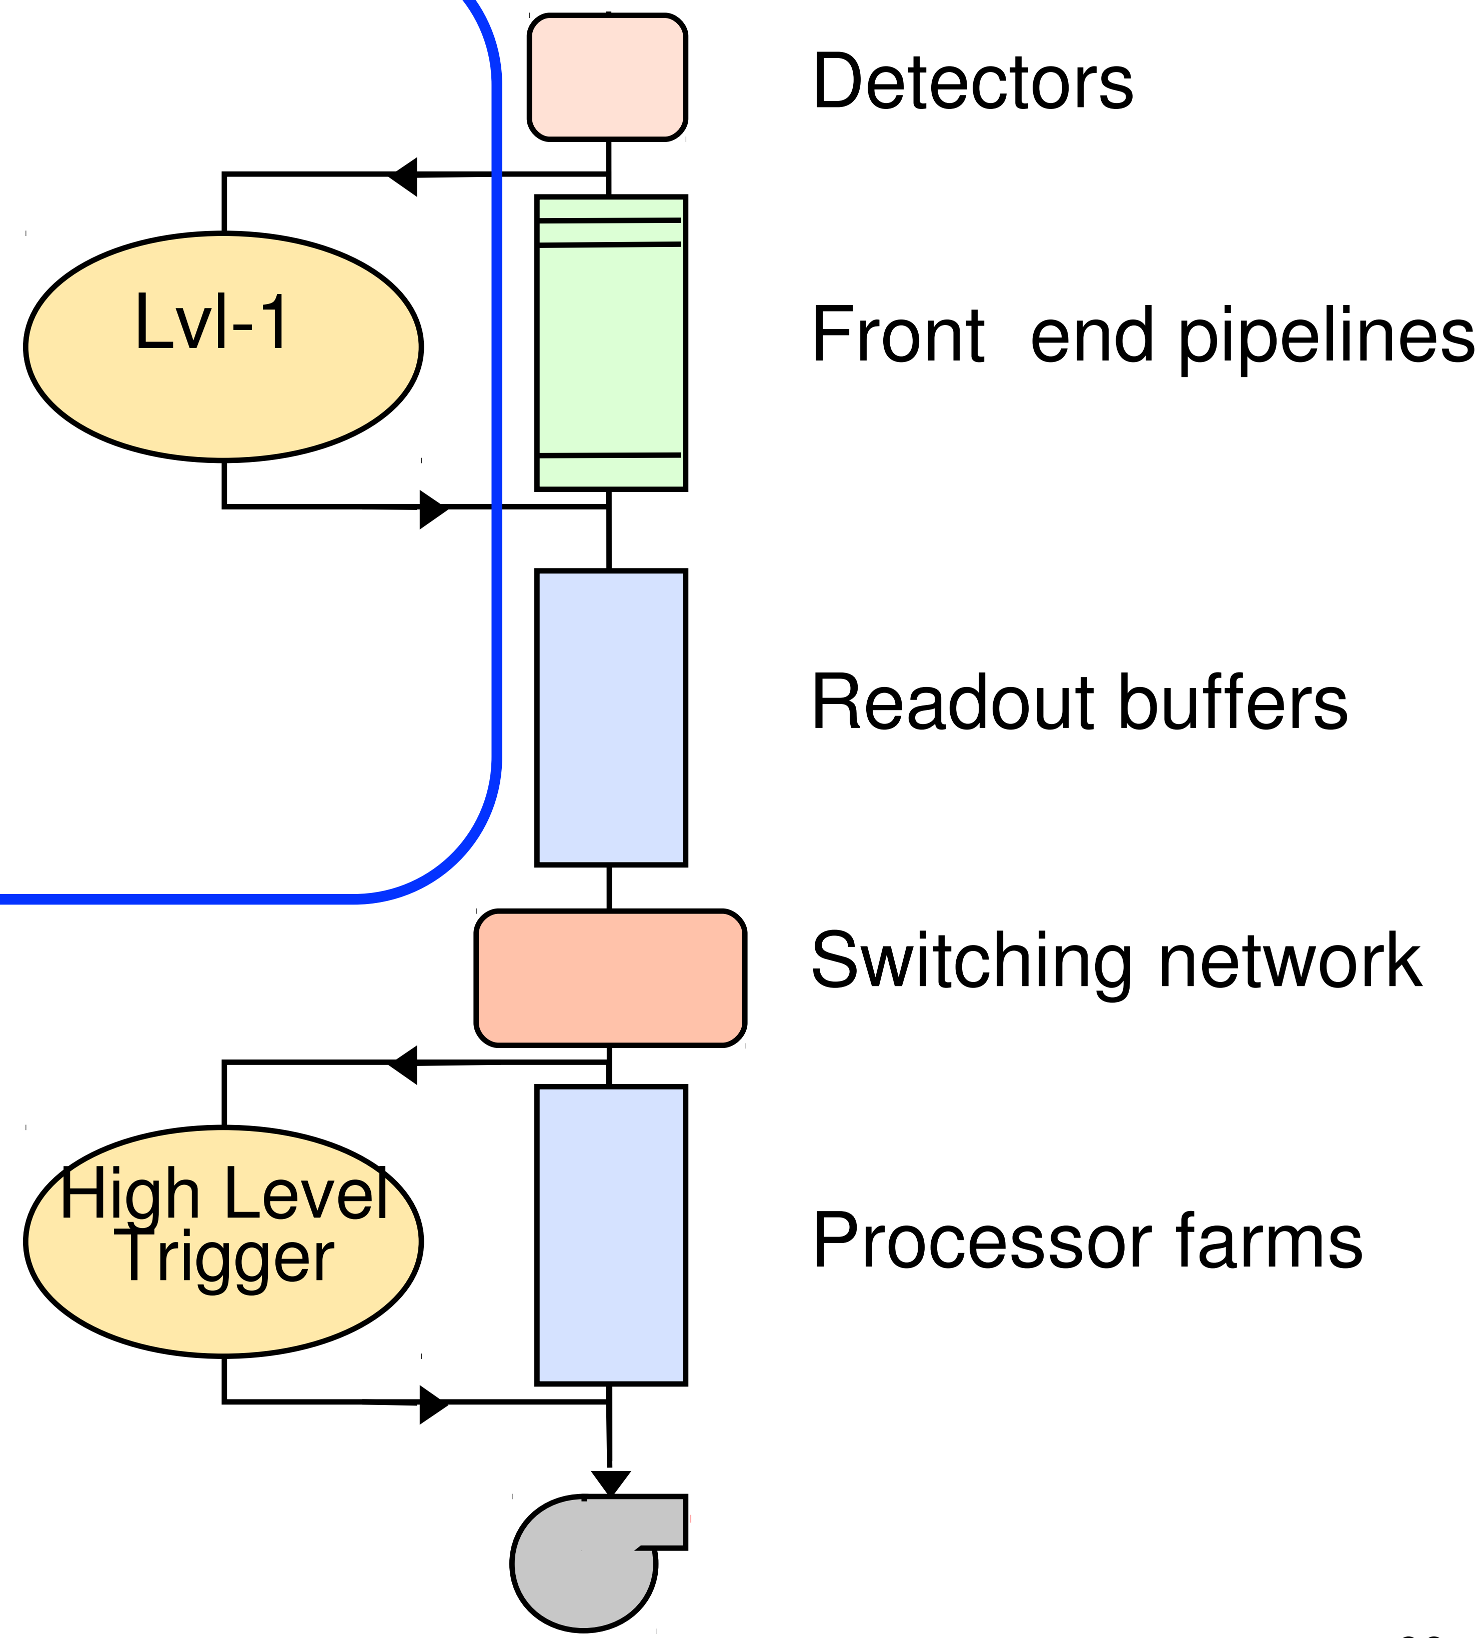

Real trigger and DAQ

- New event every 25 ns
- FE pipelines can store around 100 crossings
- LLT needs to decide with a few μs!

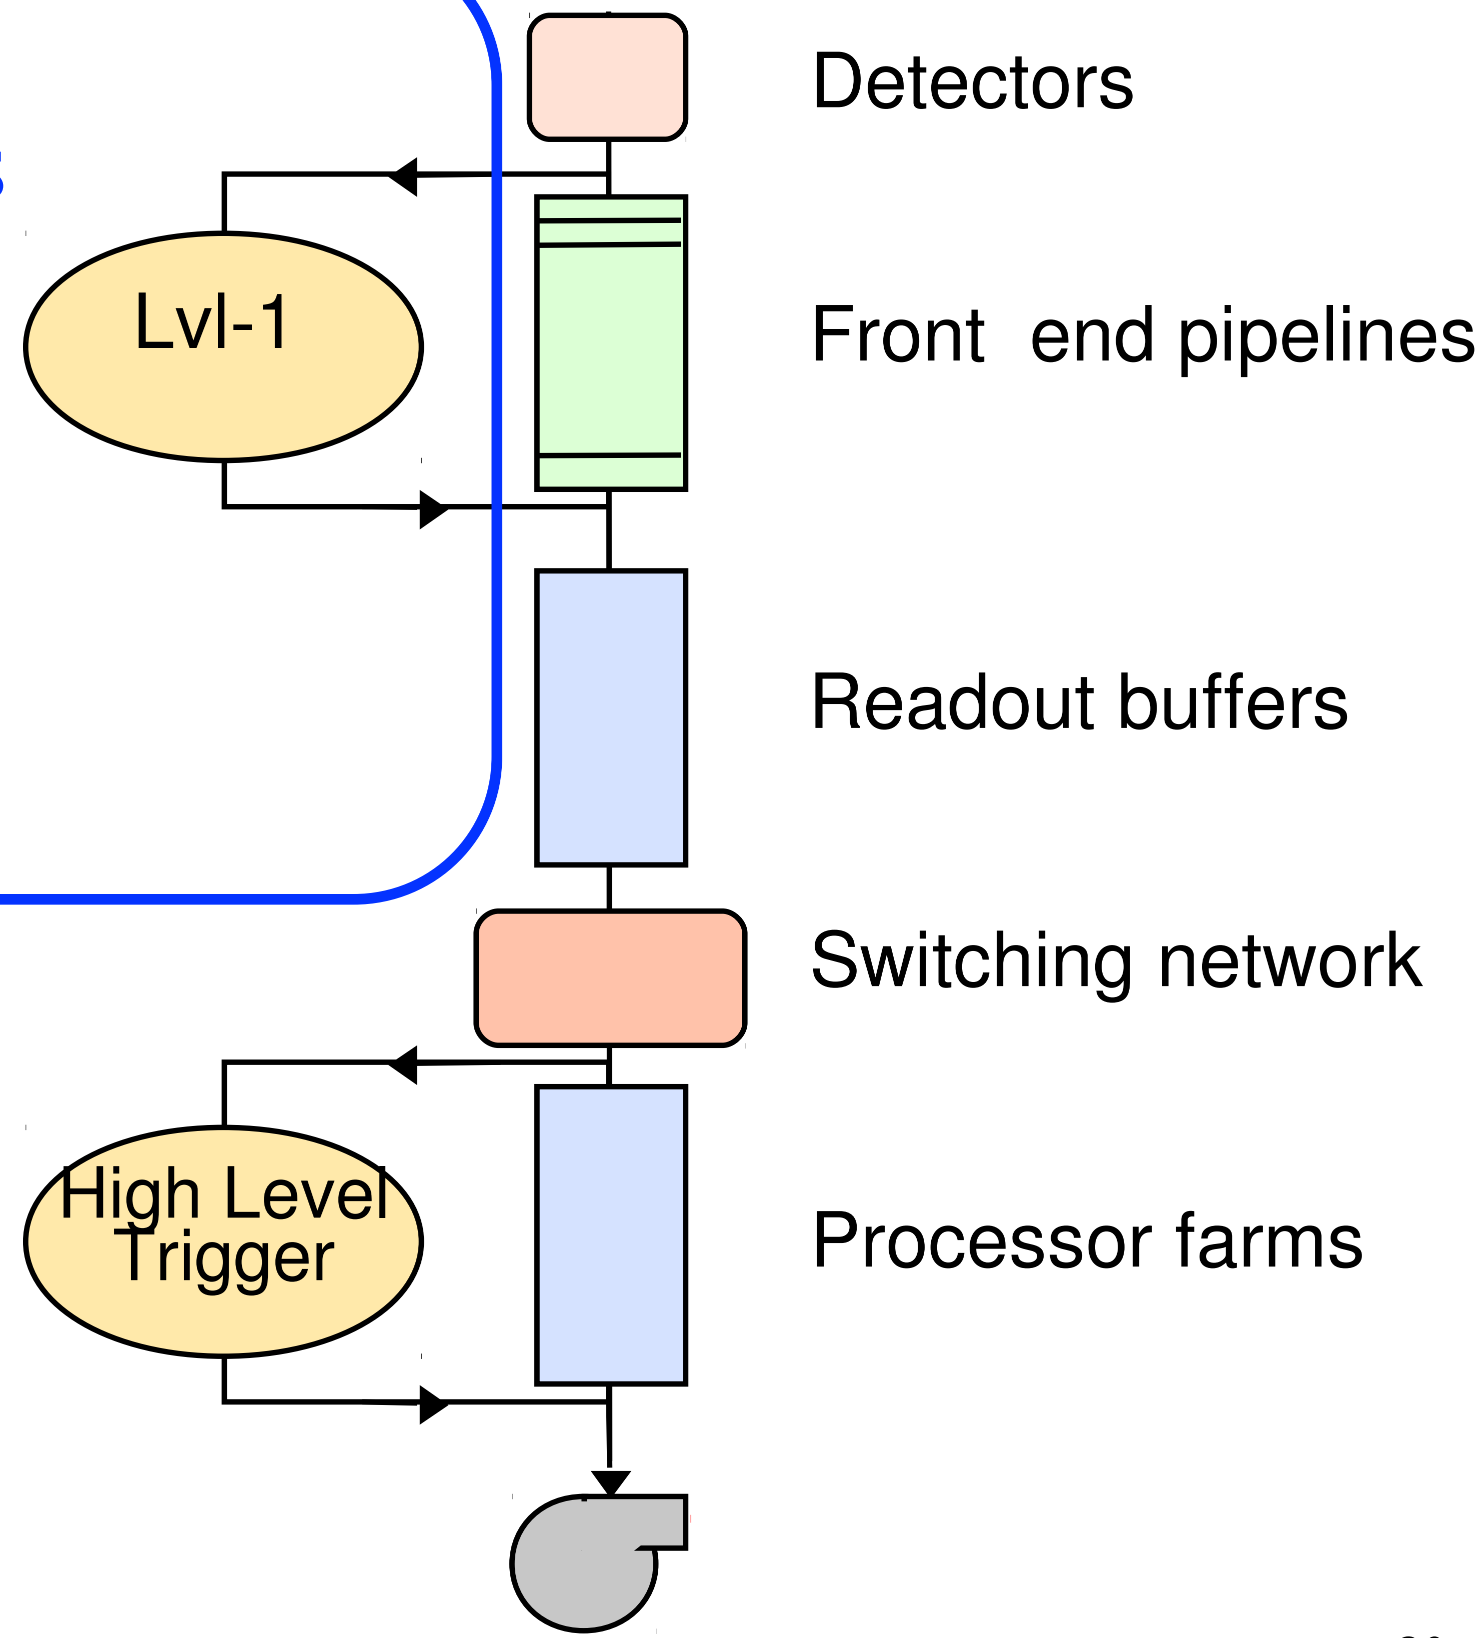

Look at the easy stuff first

CALO

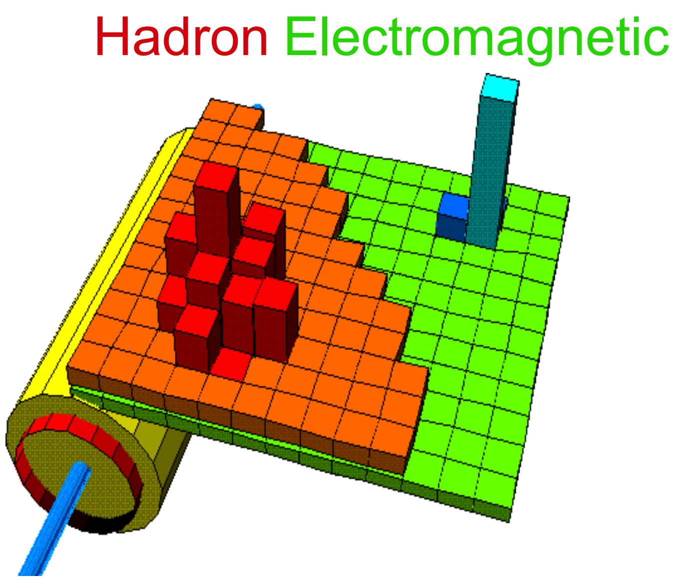

Muon

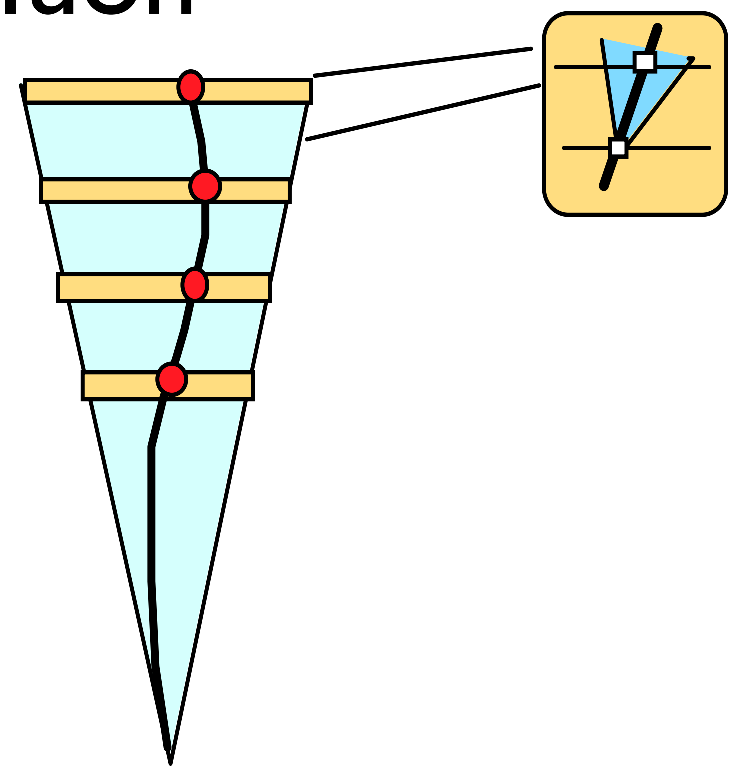

- Small amount of data
- Simple algorithms
- Local

Look at the easy stuff first

CALO

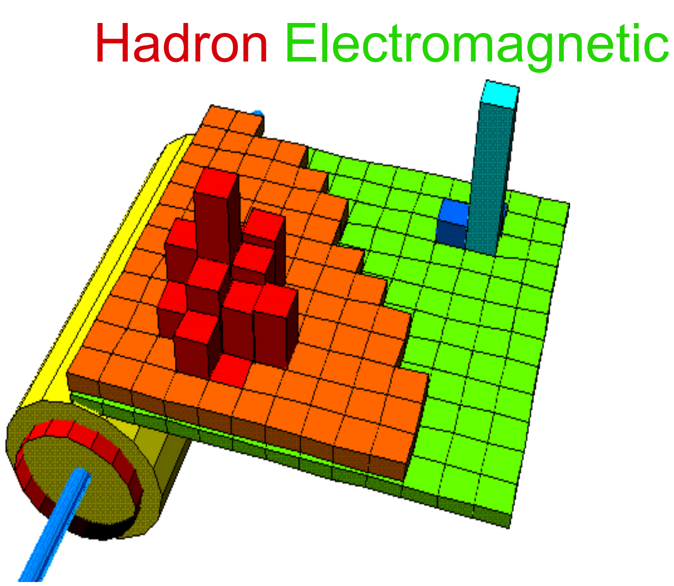

Muon

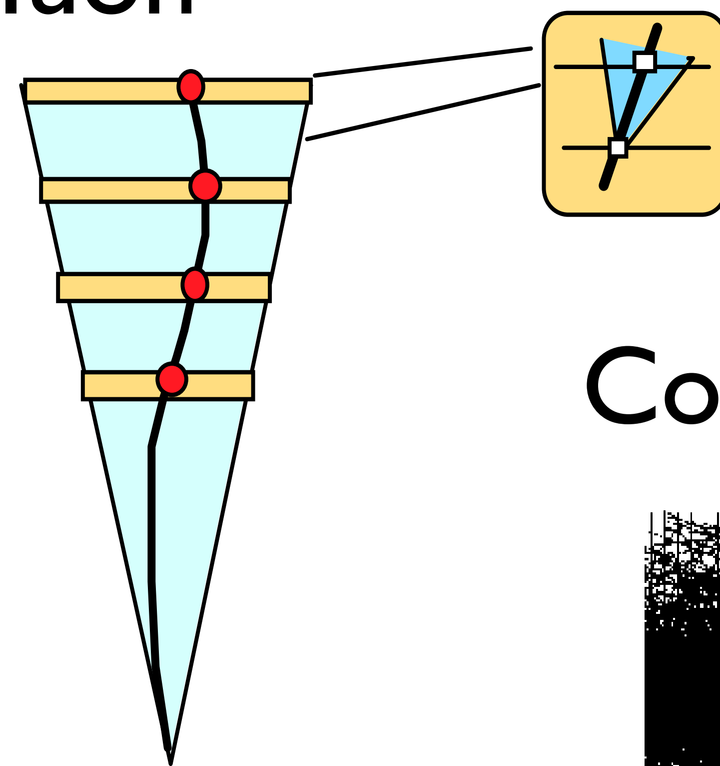

Compare with tracker!

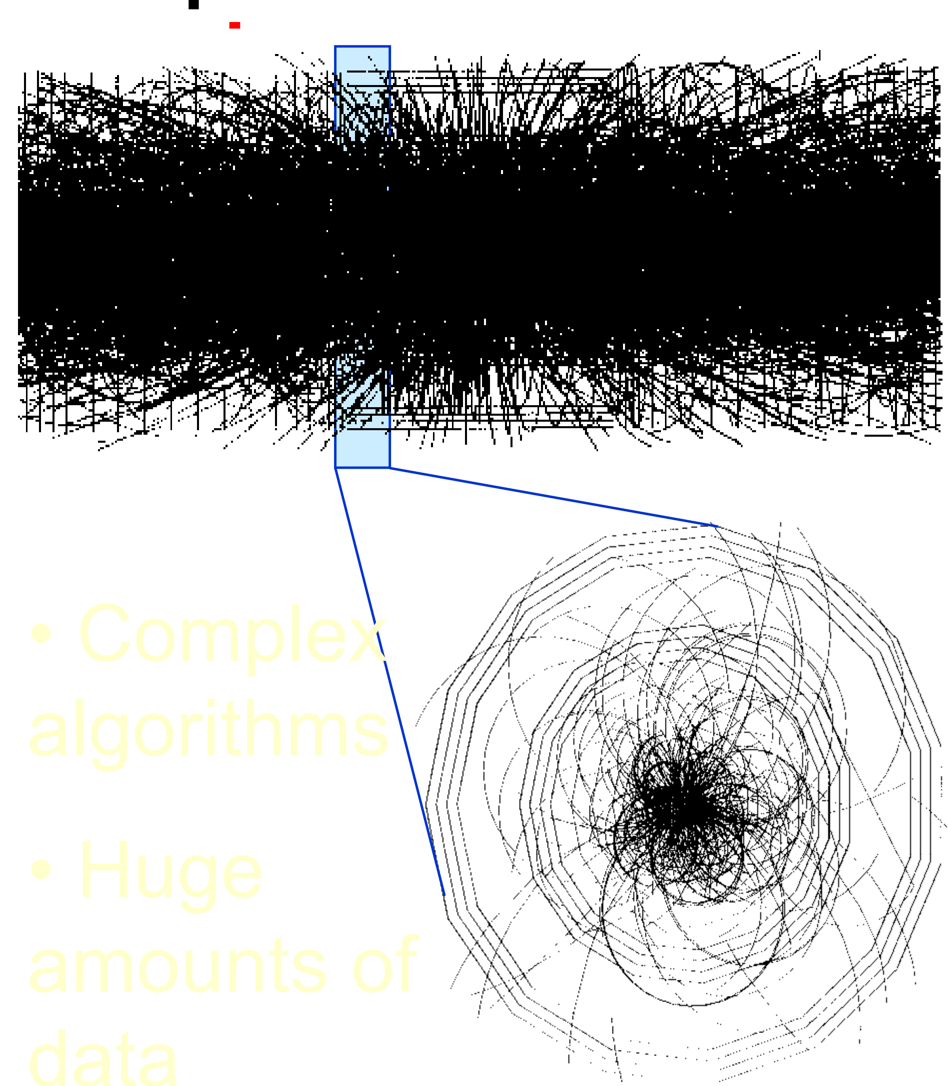

- Small amount of data
- Simple algorithms
- Local

- Complex algorithms
- Huge amounts of data

Inclusive high p_T muon and calo triggers

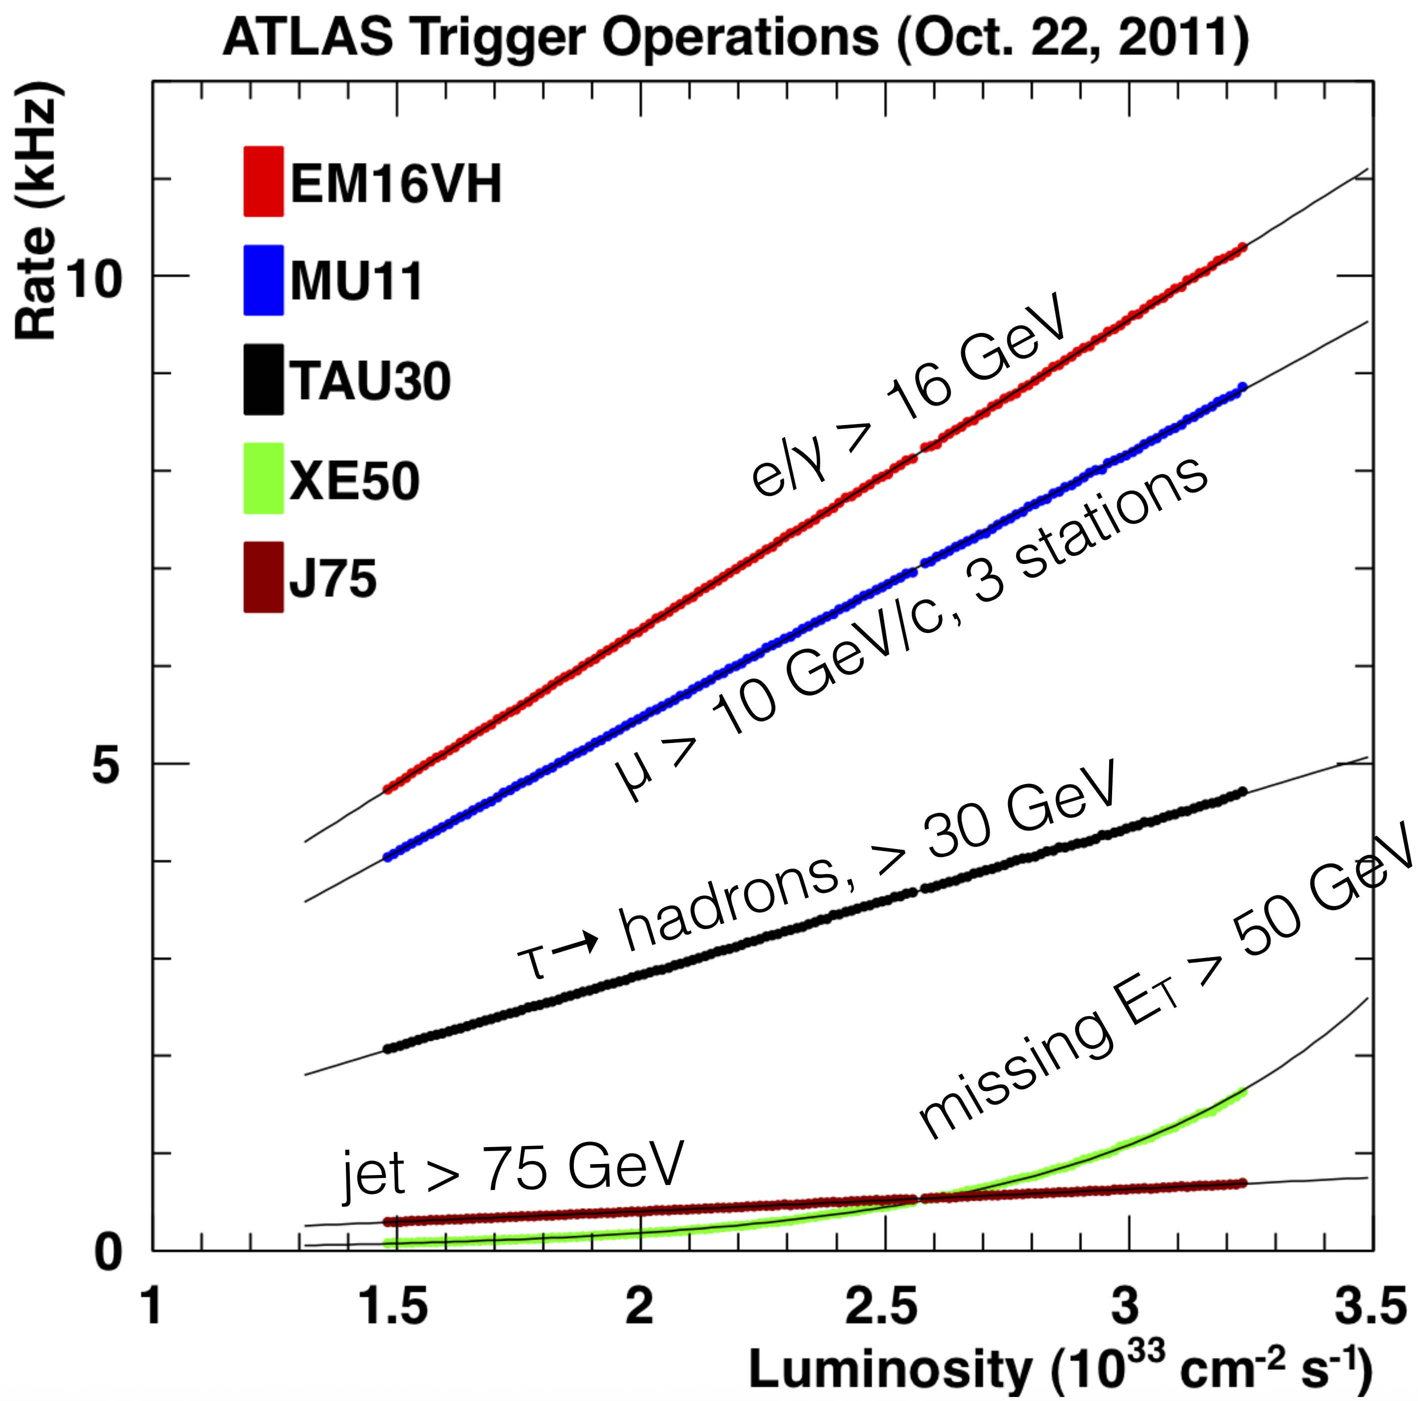

ATLAS and CMS can be efficient on the high p_T physics with ~100 kHz L1 output rate

Low level triggers

- Efficient when asked to find a clear feature that discriminates between signal and background.
- That can be identified with simple information

For example

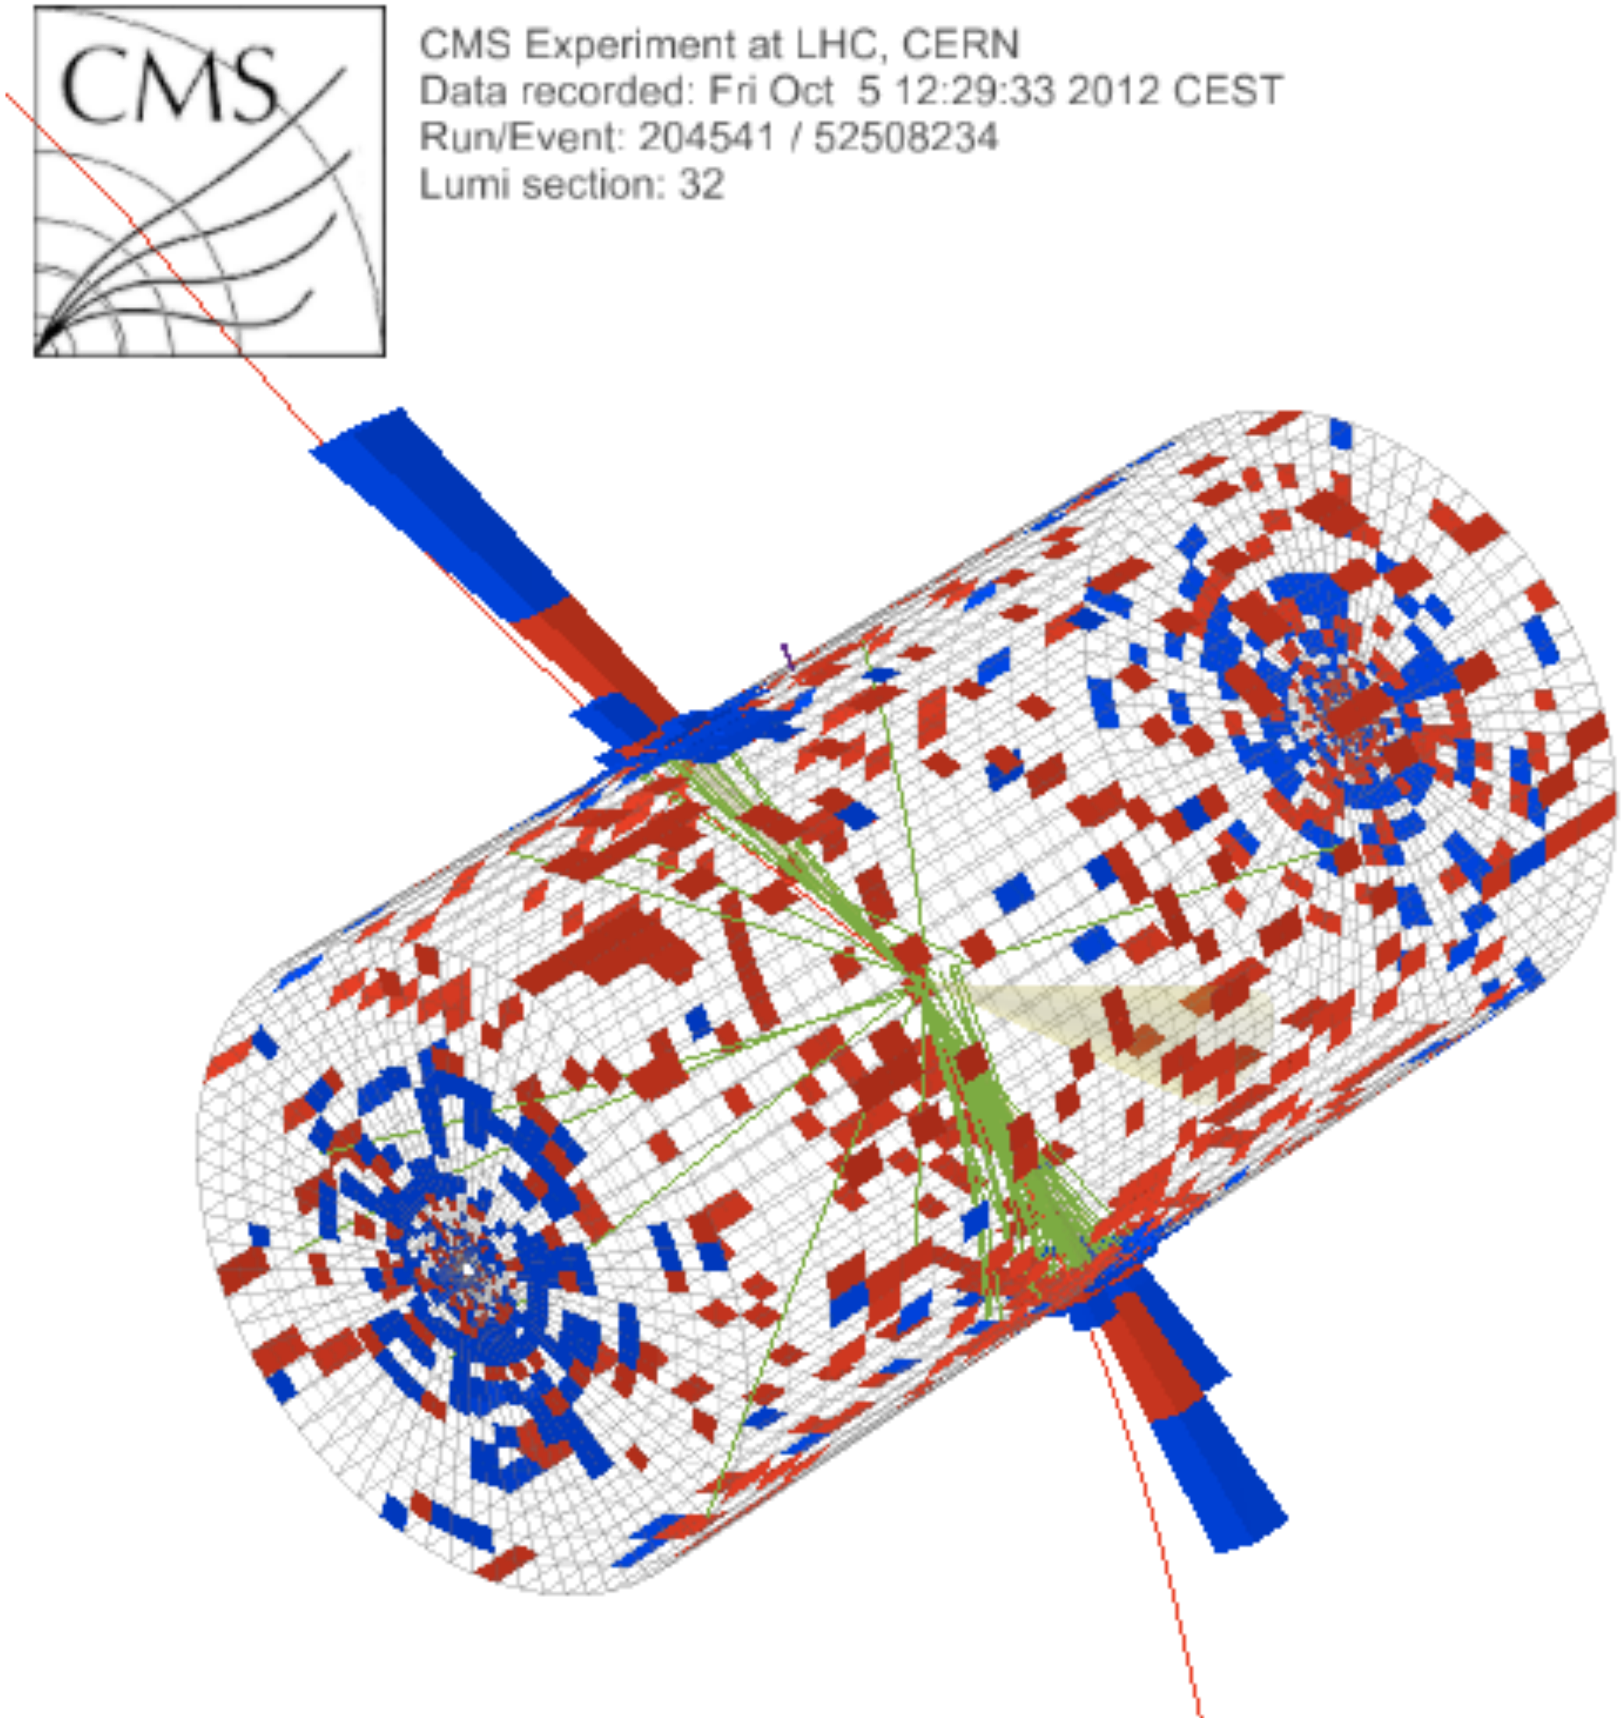
CMS Experiment at LHC, CERN
Data recorded: Fri Oct 5 12:29:33 2012 CEST
Run/Event: 204541 / 52508234
Lumi section: 32

LHCb case a little harder

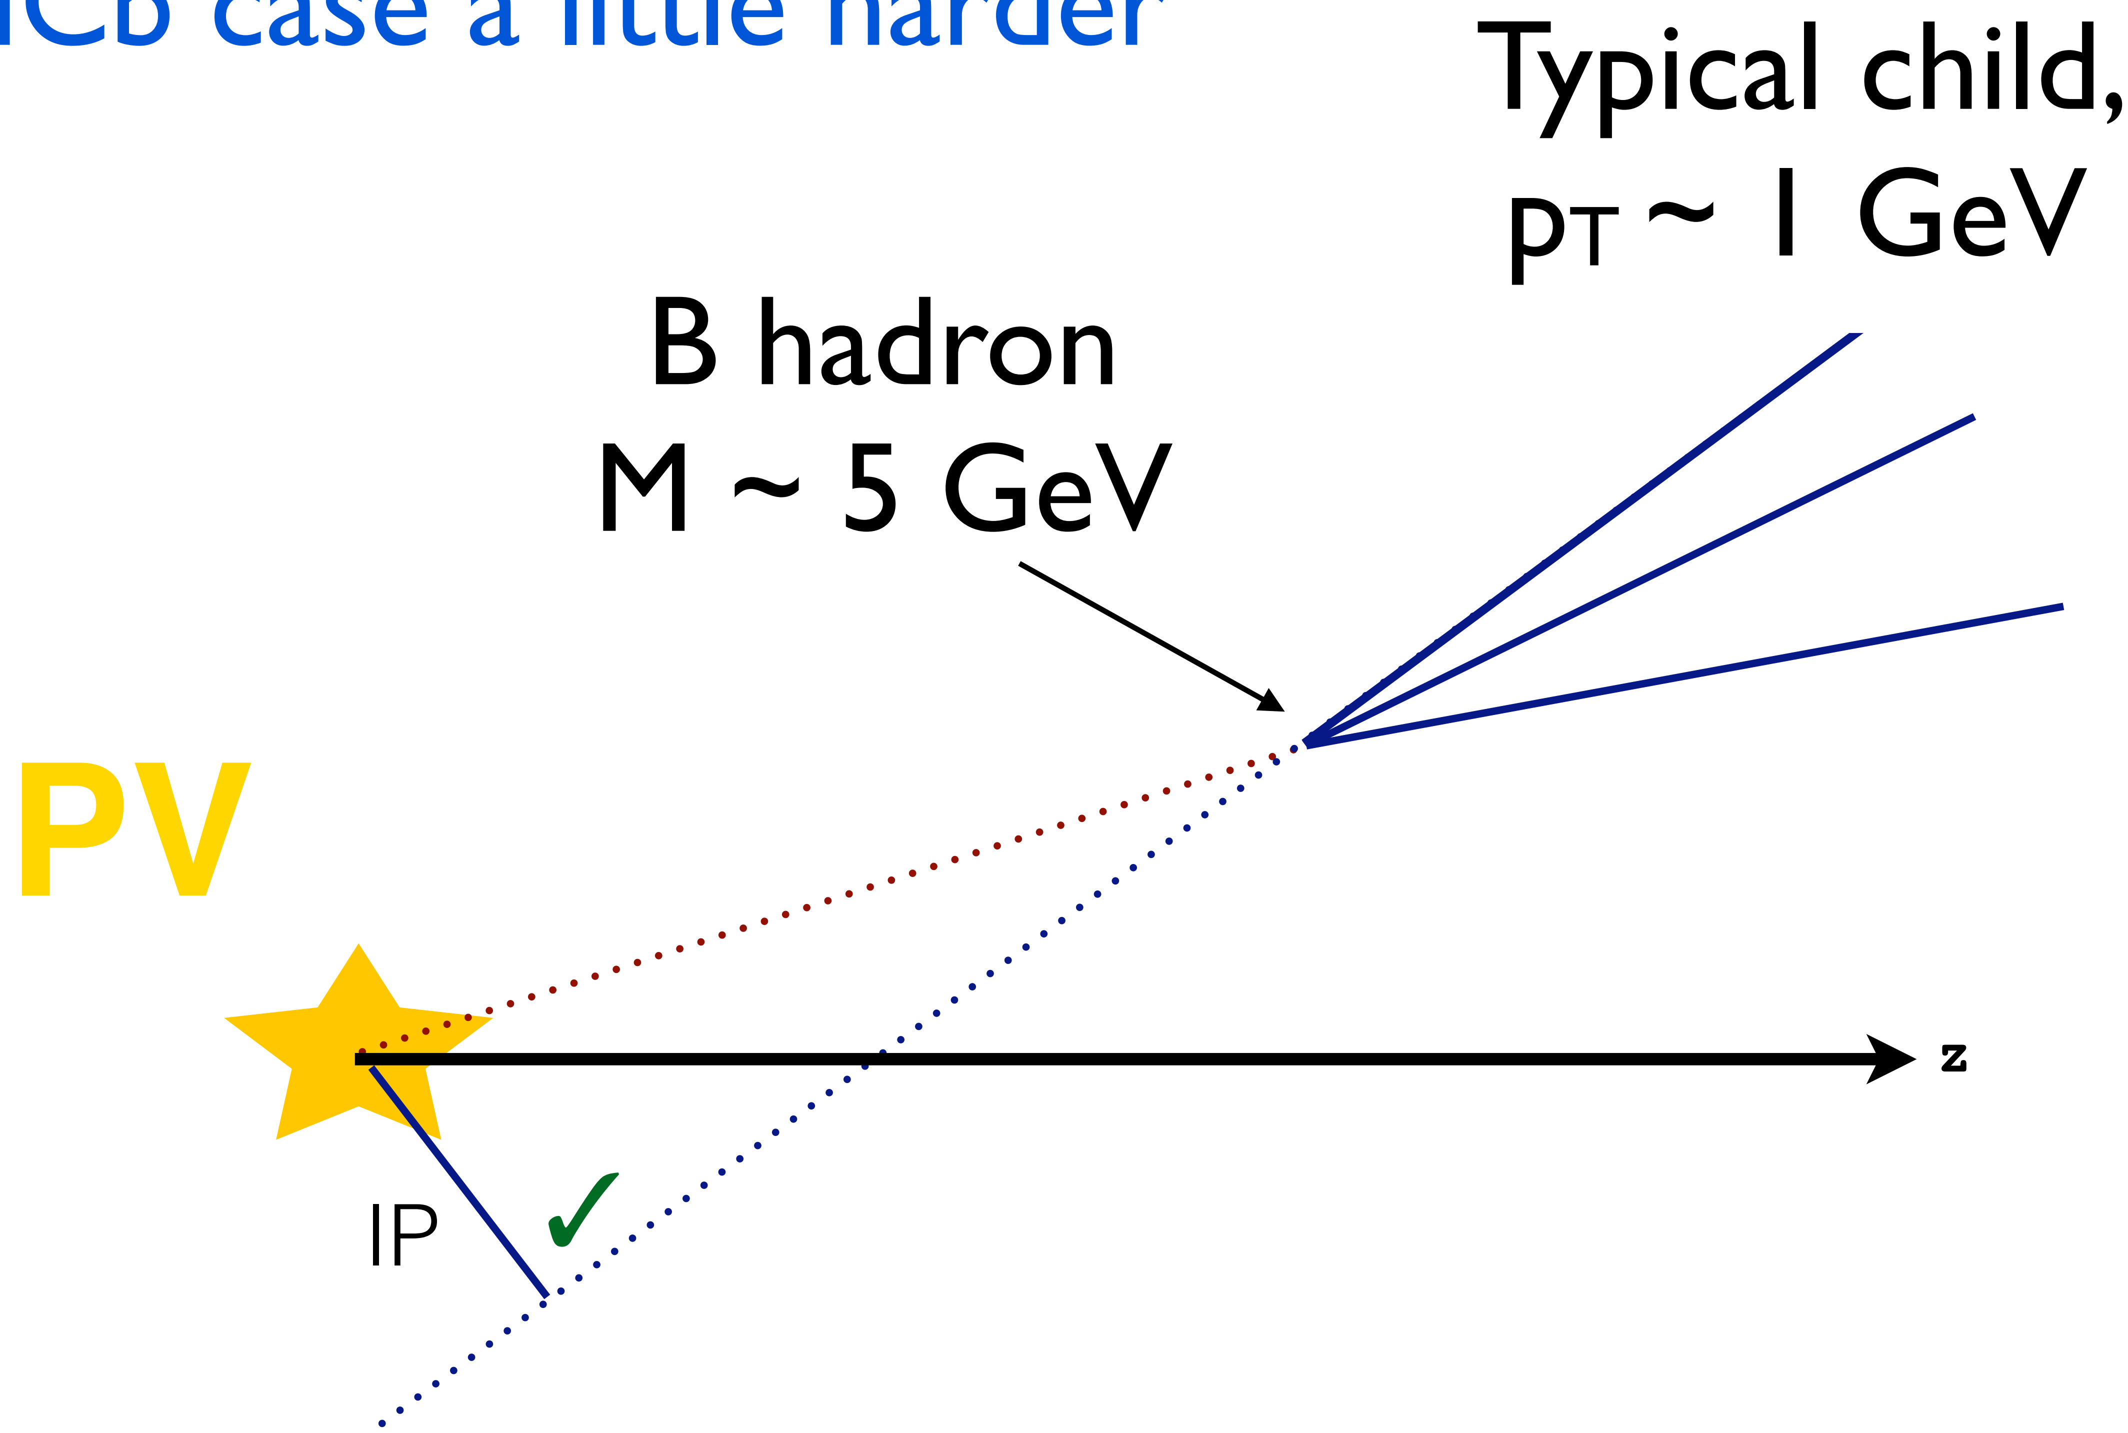

Can't reduce the rate as efficiently with low-level hardware trigger

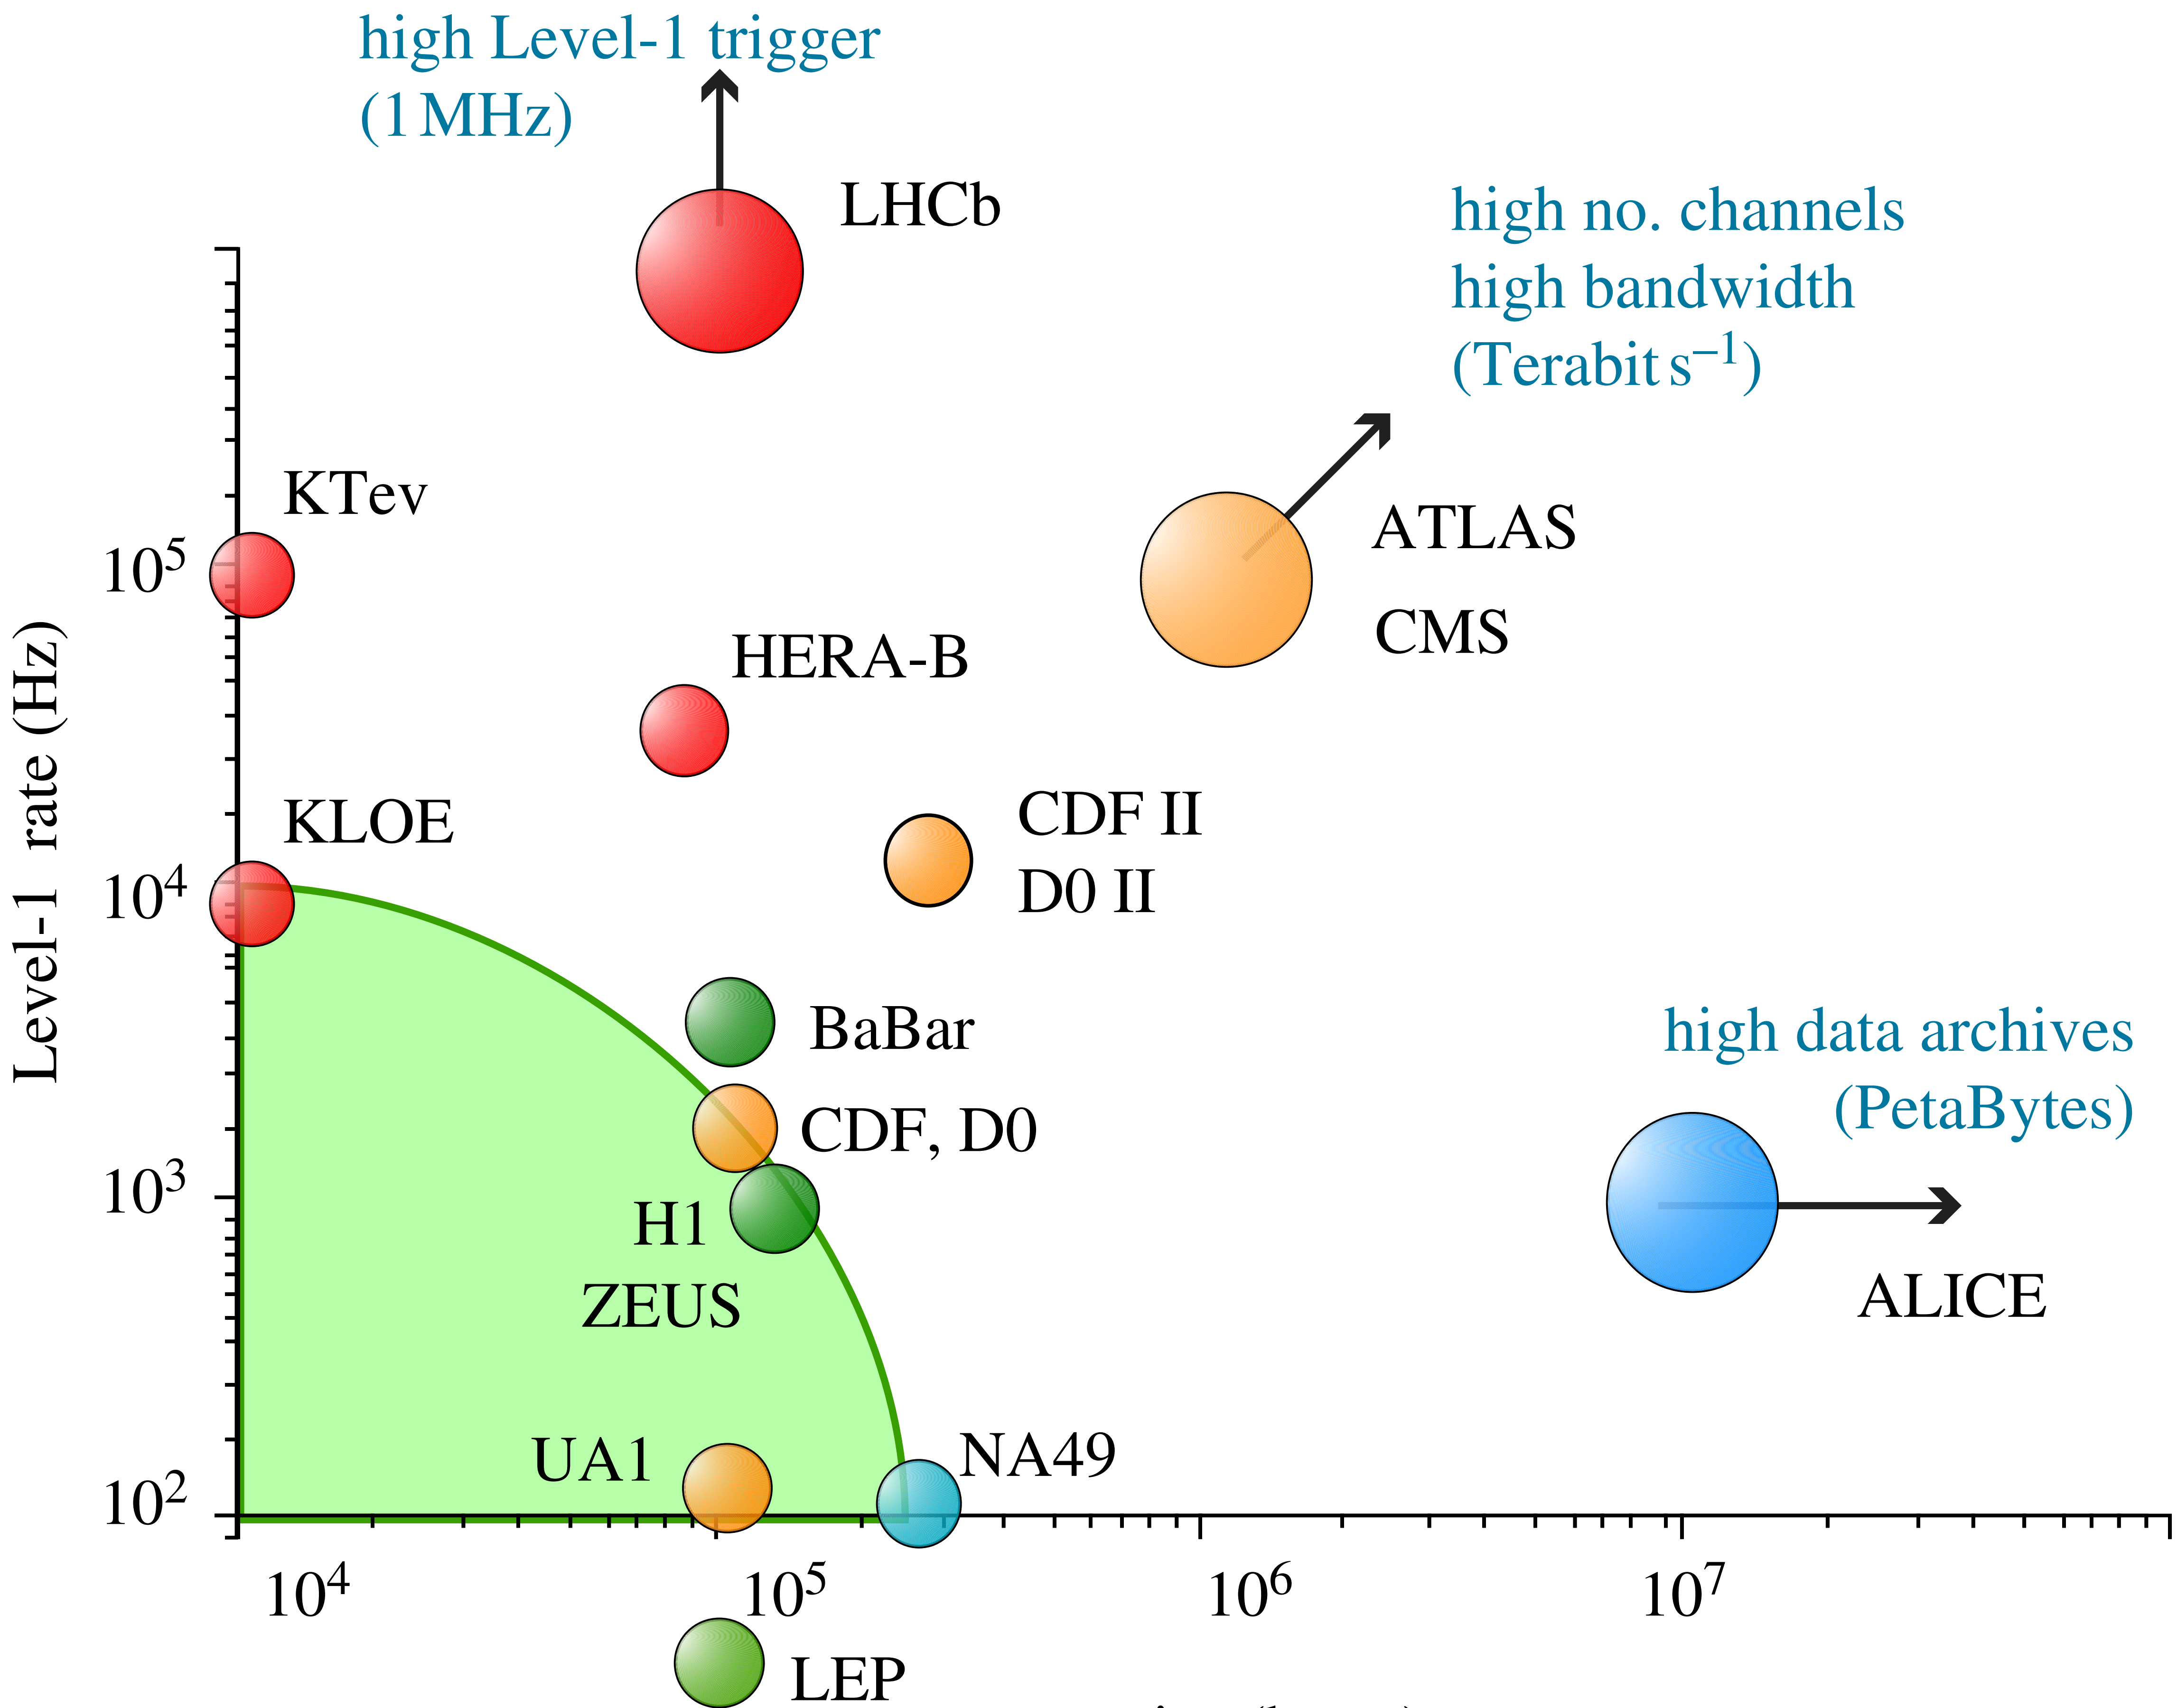

high Level-1 trigger
(1 MHz)
LHCb
high no. channels
high bandwidth
(Terabit s−1)
KTev
ATLAS
CMS
HERA-B
CDF II
D0 II
KLOE
BaBar
high data archives
(PetaBytes)
CDF, D0
H1
ZEUS
ALICE
UA1
NA49
LEP
Level-1 rate (Hz)
event size (bytes)

Lumi levelling

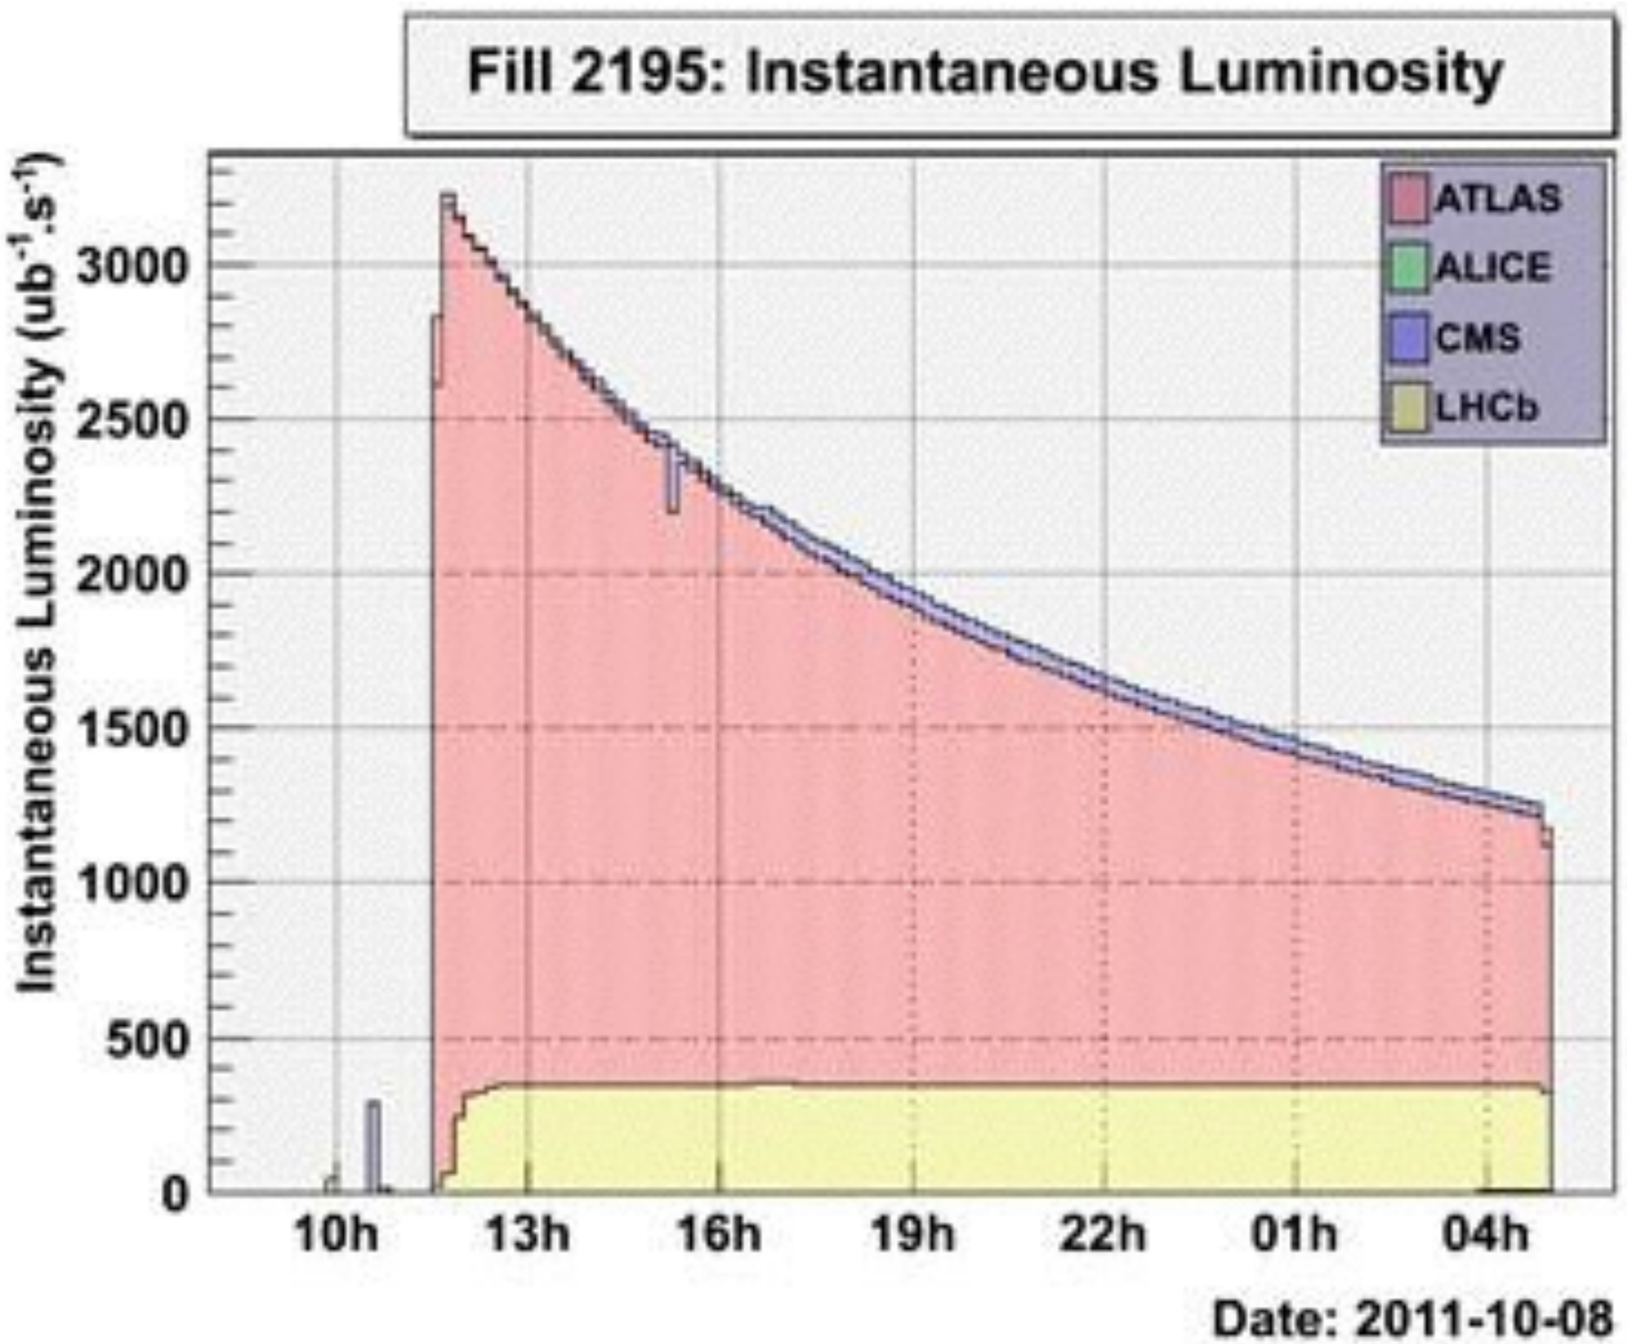

LHCb levelled at 4×10^{32} cm^{-2}s^{-1}

LHCb Run-I trigger

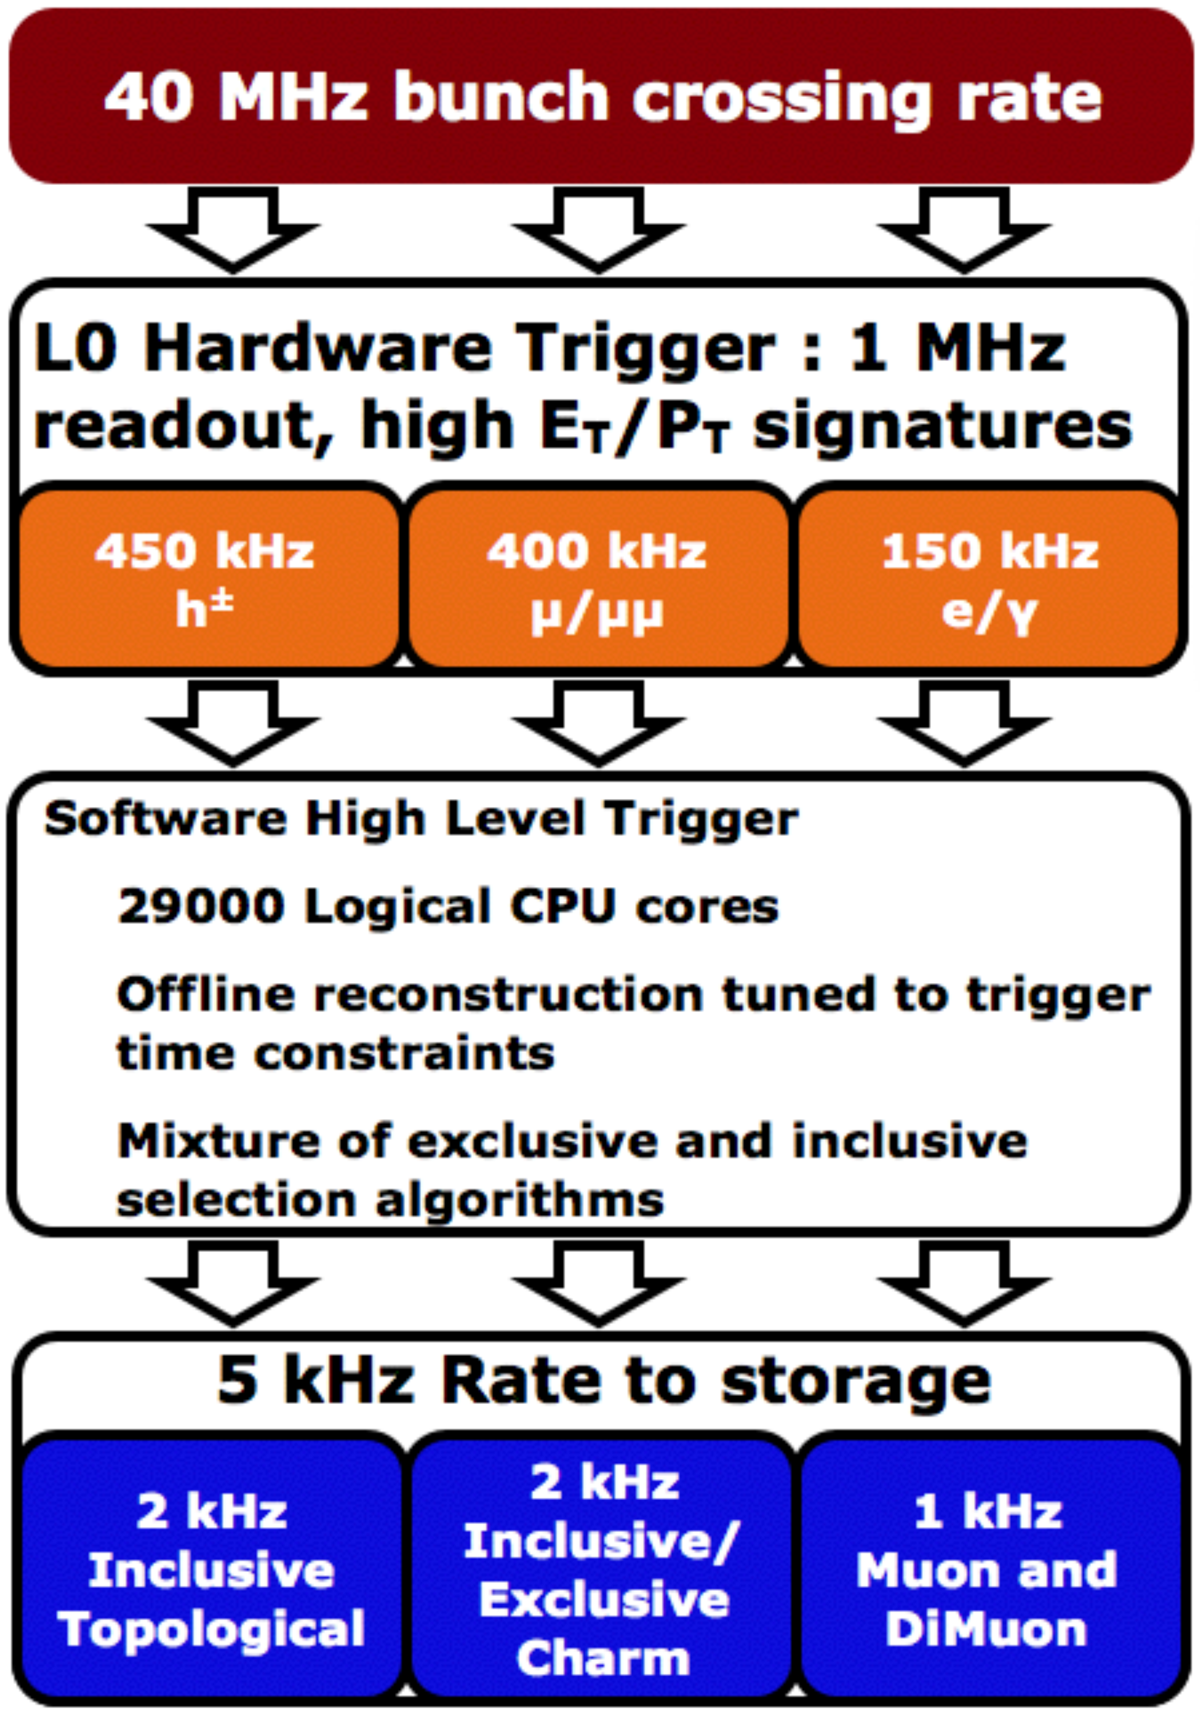

For 50:50 split of the 1 MHz between hadrons and muons:

- Muon p_T > 1.7 GeV
- Hadron E_T > 3.5 GeV

(Run-I)

We need hadronic B for physics

$$\begin{bmatrix} d' \\ s' \\ b' \end{bmatrix} = \begin{bmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{bmatrix} \begin{bmatrix} d \\ s \\ b \end{bmatrix}.$$

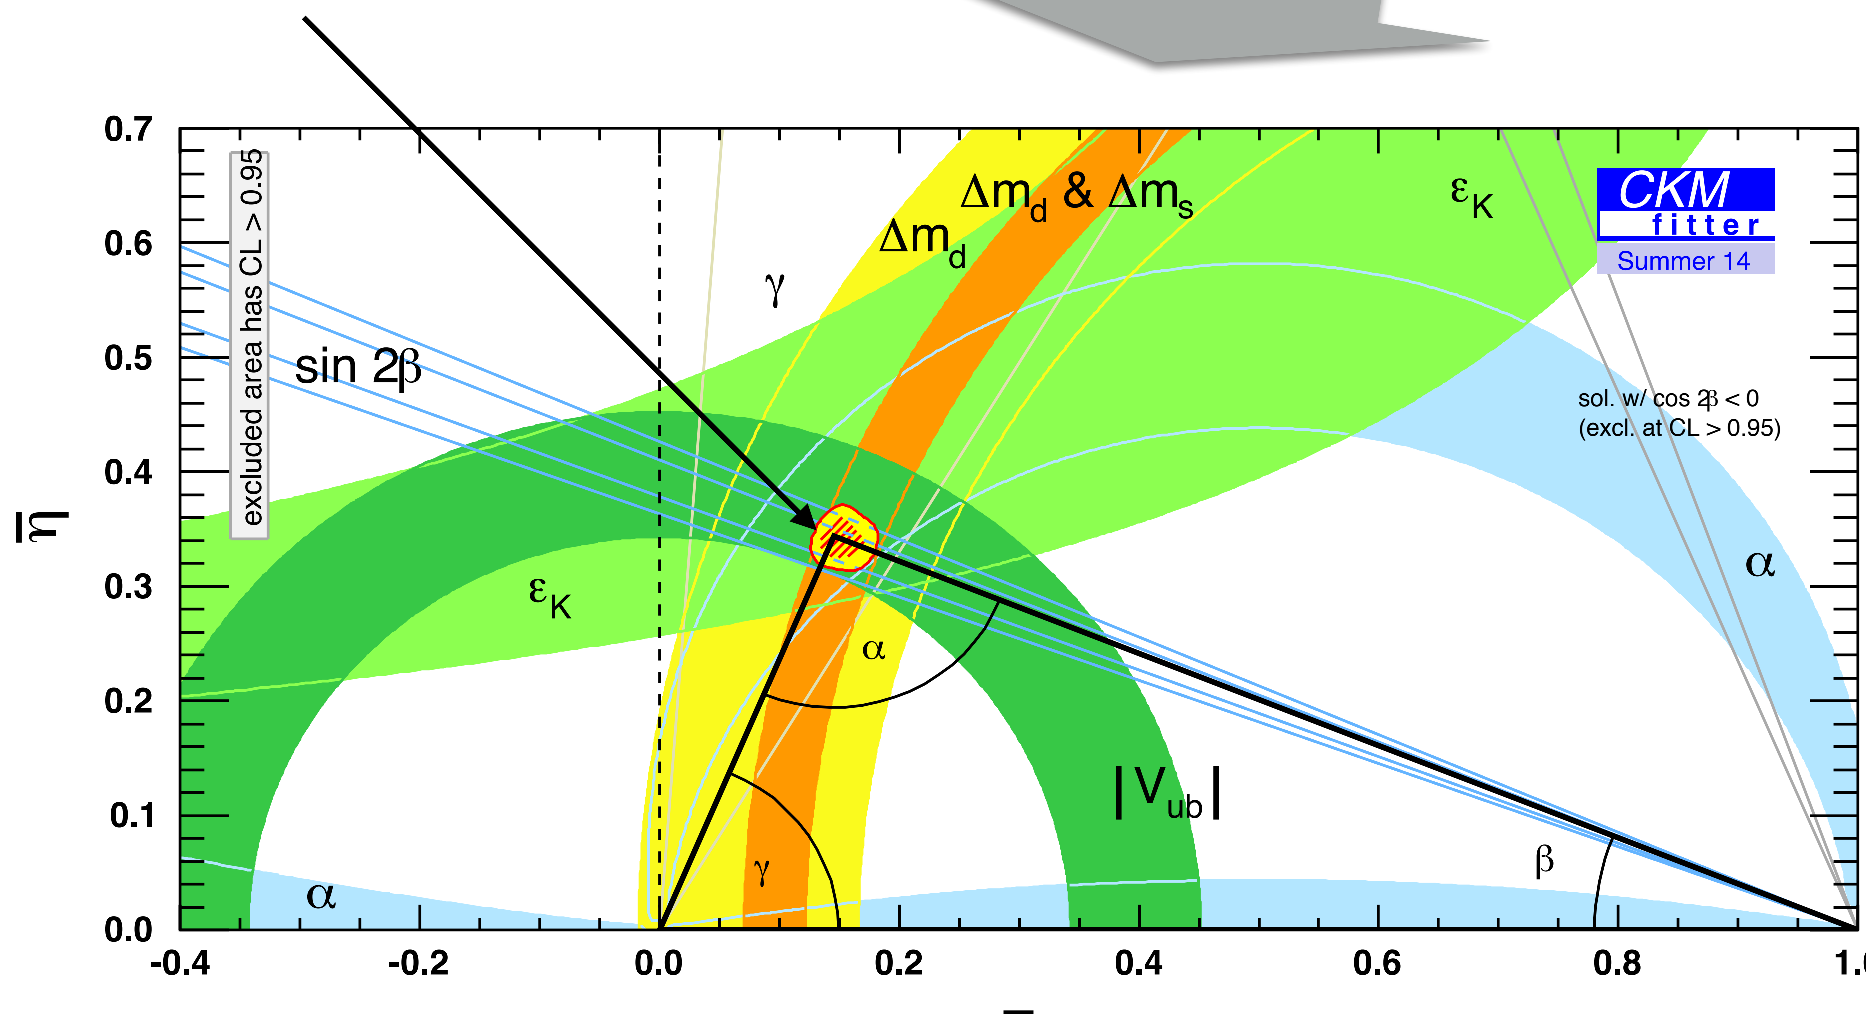

We need hadronic B for physics

$$\begin{bmatrix} d' \\ s' \\ b' \end{bmatrix} = \begin{bmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{bmatrix} \begin{bmatrix} d \\ s \\ b \end{bmatrix}.$$

Unitarity

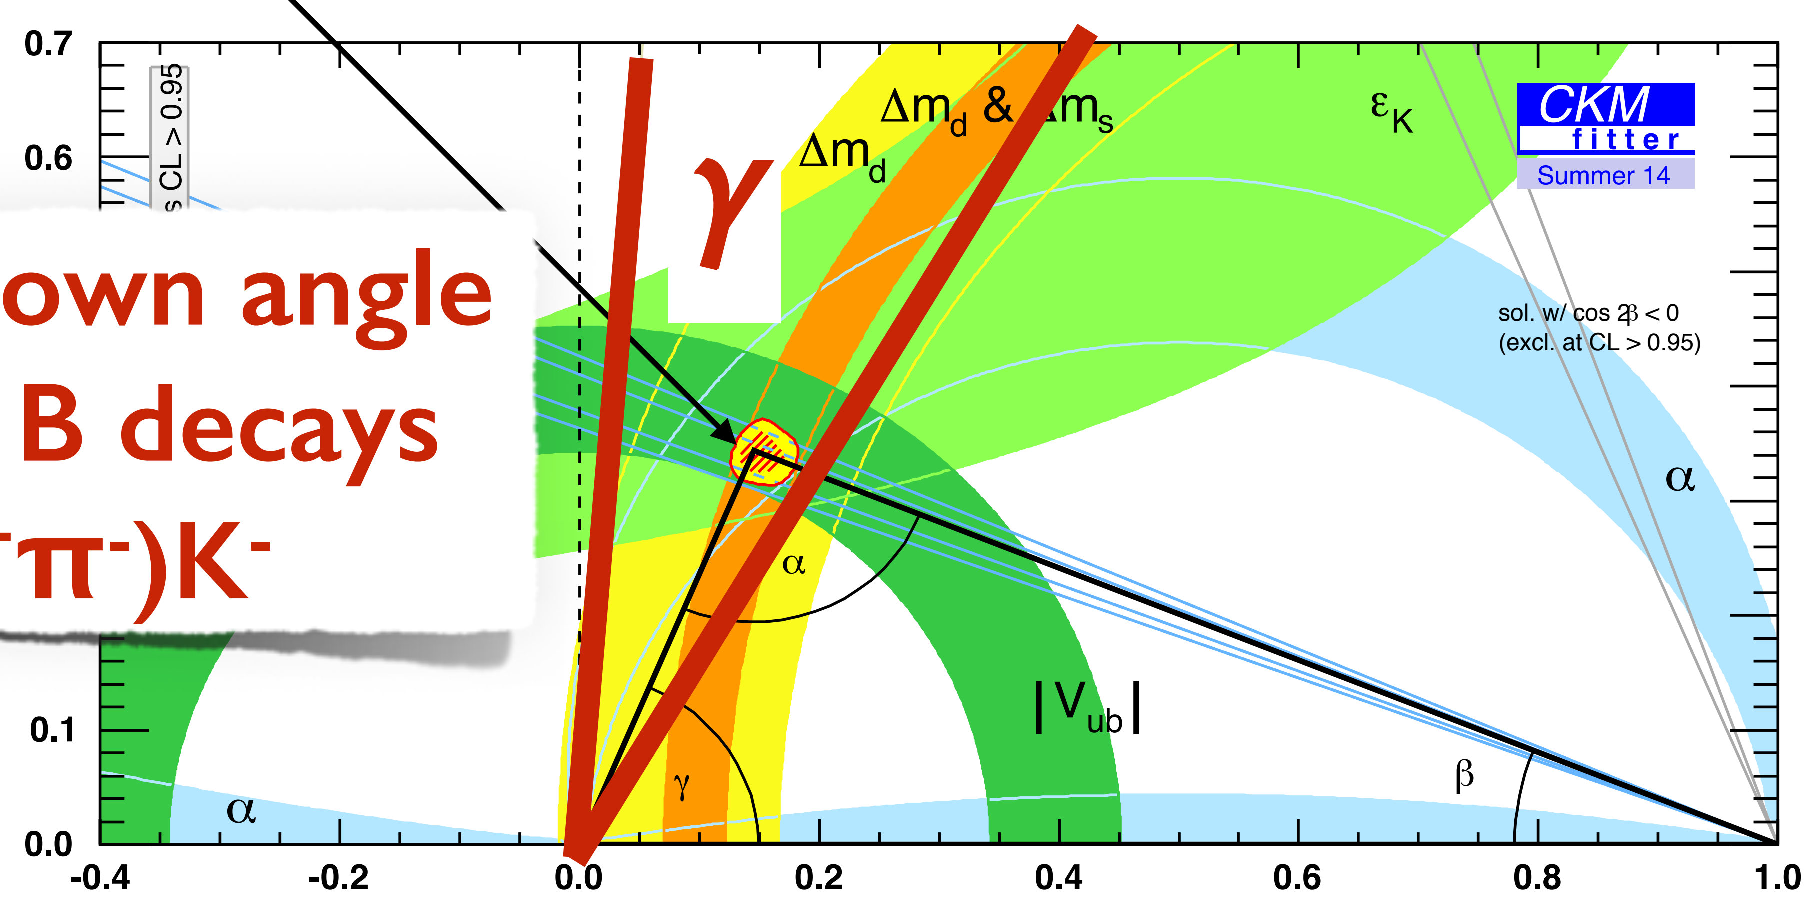

E.g. least well known angle needs hadronic B decays like $D^0(\rightarrow K^+\pi^-)K^-$

The LHCb upgrade (2019)

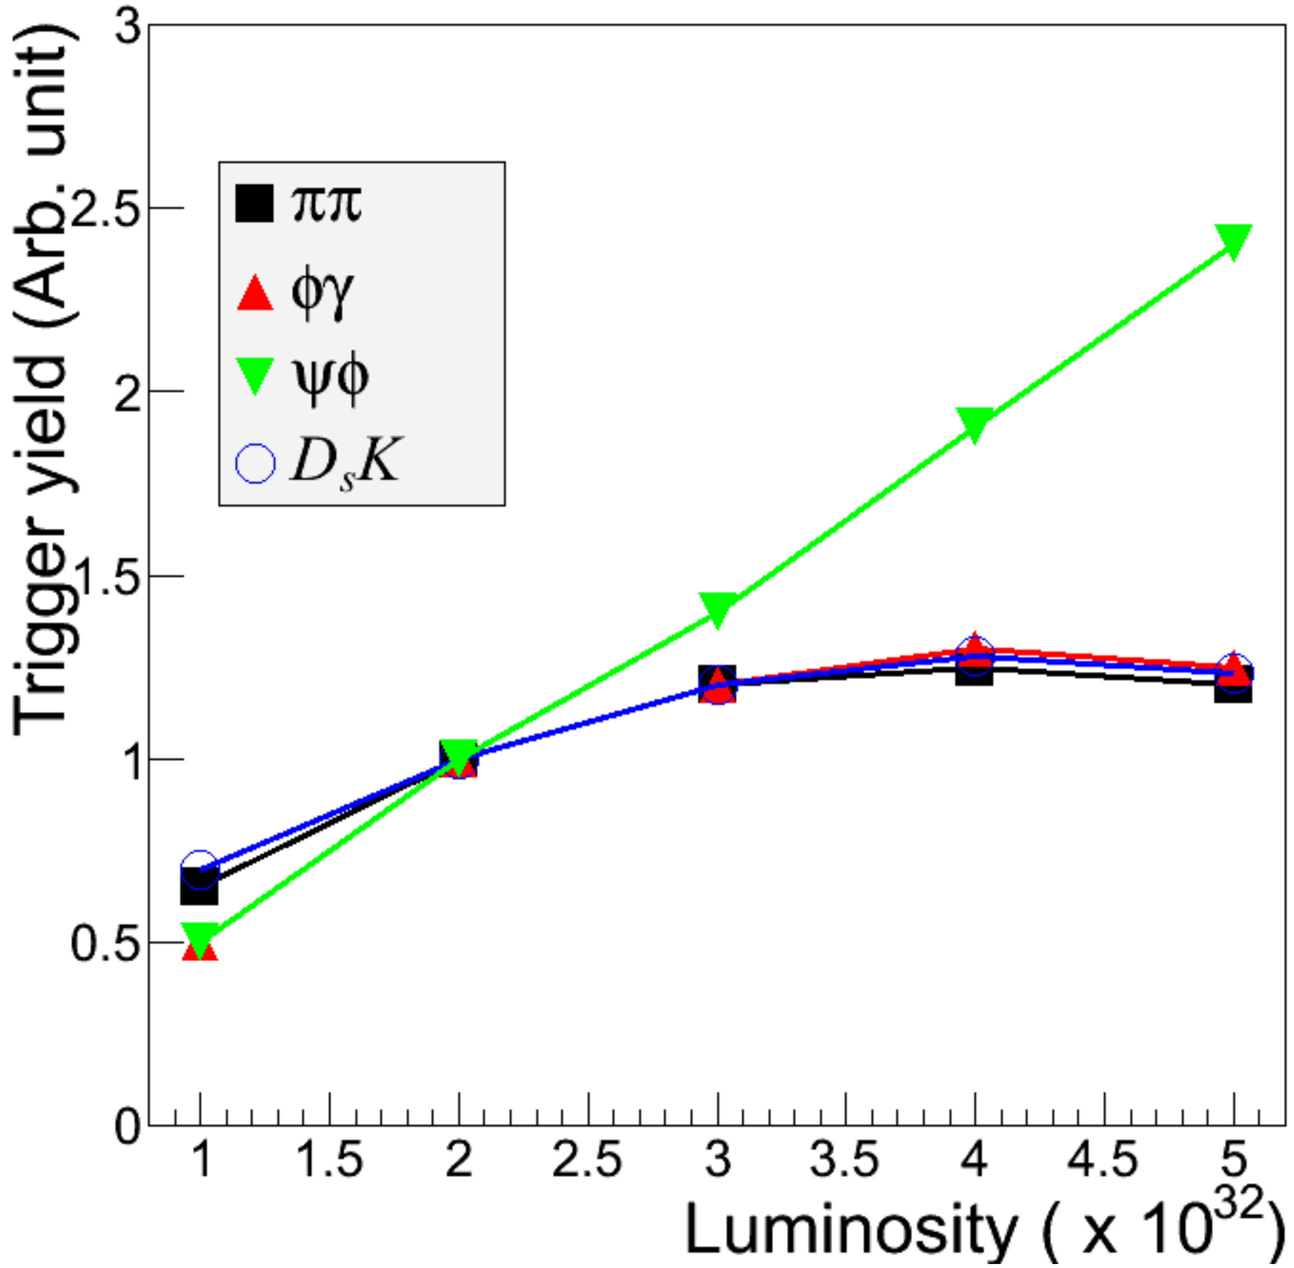

In Run-III (2020), LHCb will have a full software trigger at 30 MHz

High level triggers

How much data can we write out?

Limited by tape and disk storage,
and ability to reconstruct and process the data

Run-I numbers

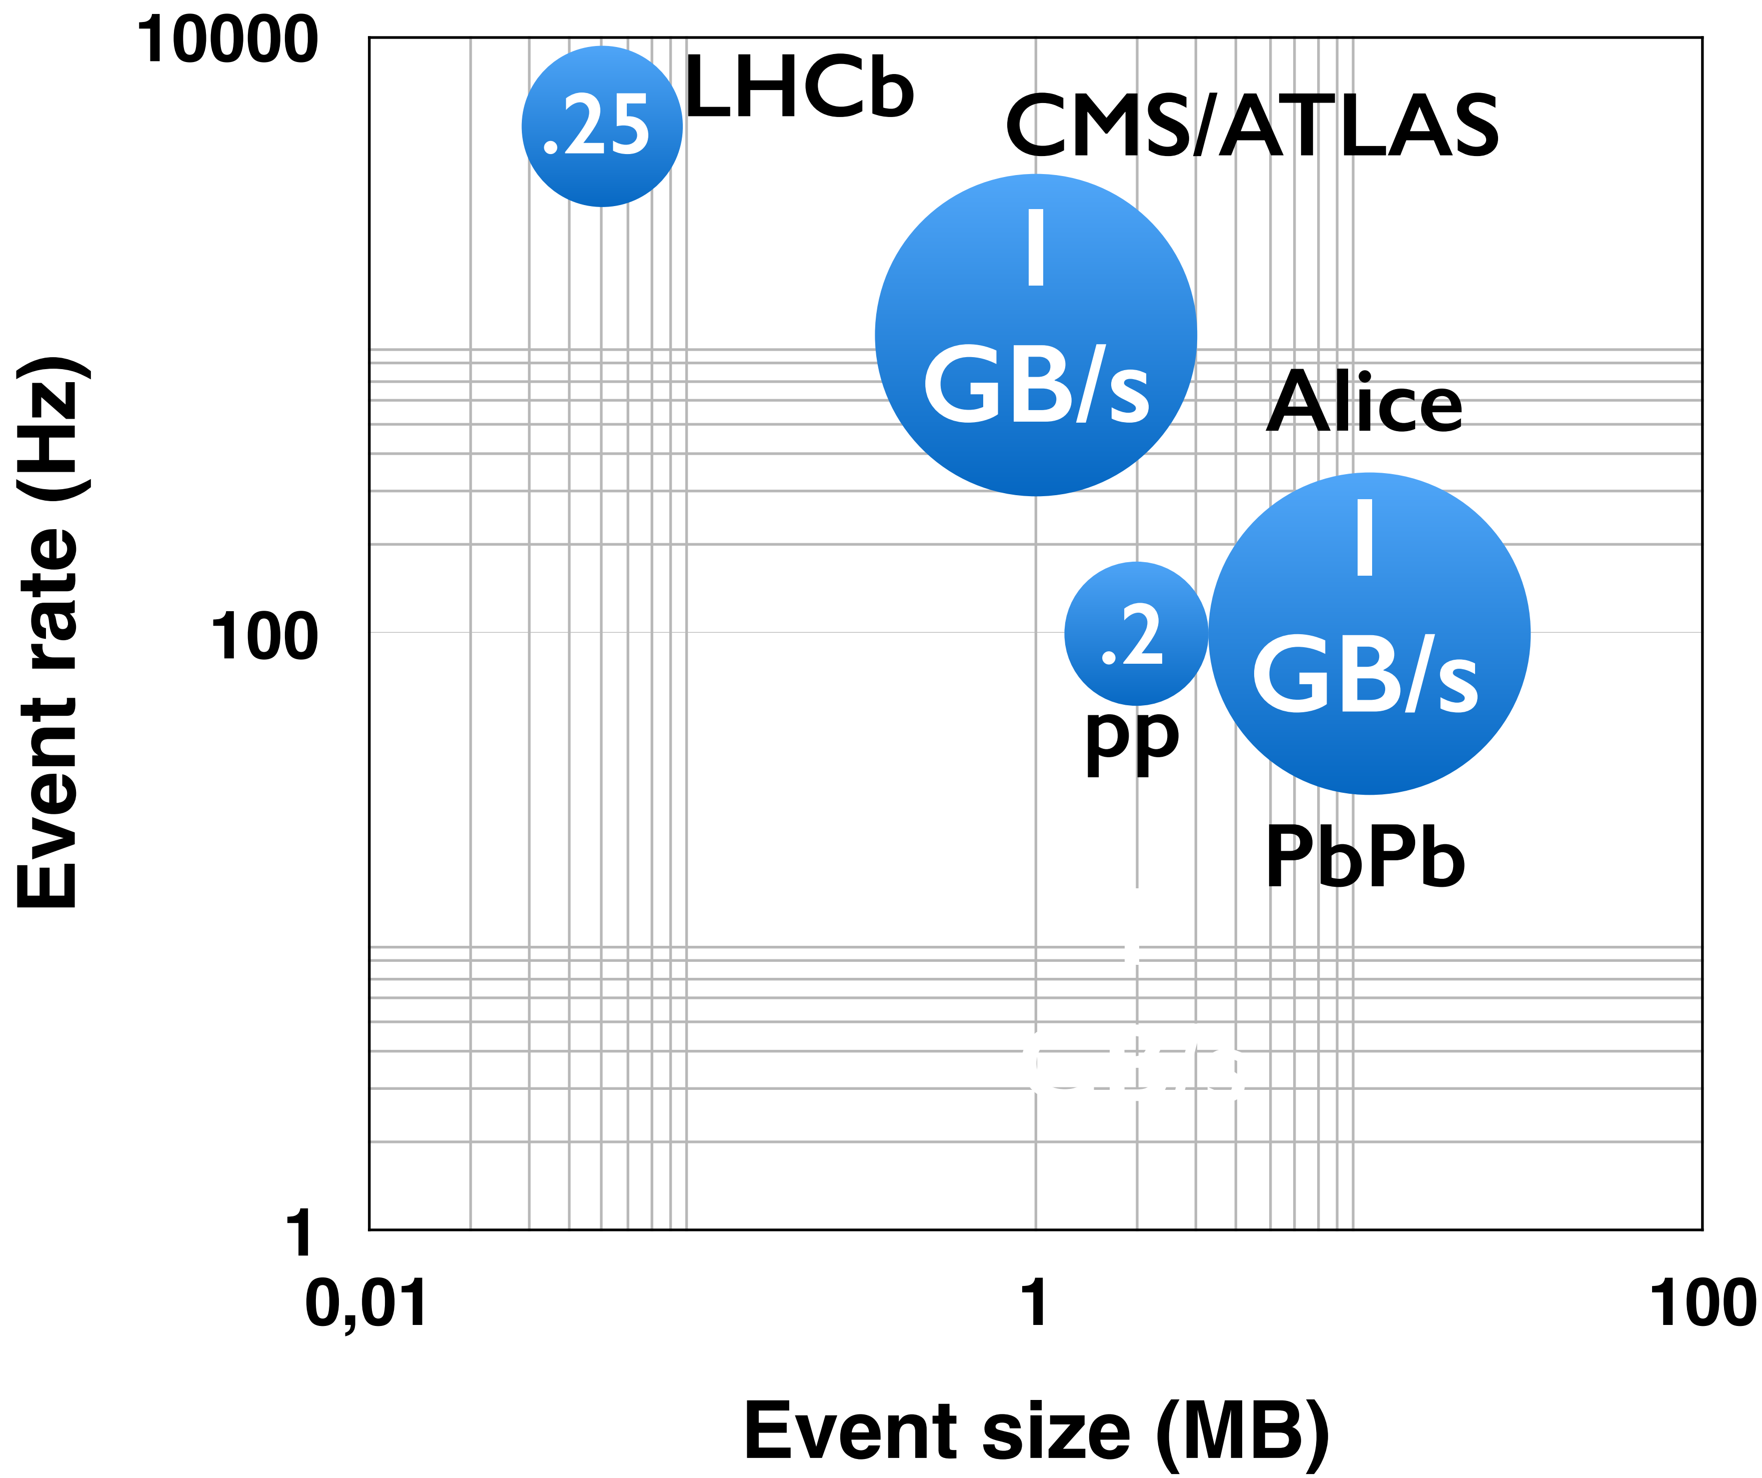

The High Level Triggers

Thanks to the LLT and event builder, we can now perform full event reconstruction with off-the-shelf PCs

Run-I specifications:

ATLAS	17k CPU cores*
CMS	13k CPU cores
LHCB	29k CPU cores

Use the same event reconstruction software as used offline, but tuned/configured differently.

Build an OR of 100-1000 different "lines"

*split between L2 and L3

ATLAS and CMS

"Menus" of mostly inclusive high p_T physics lines

(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

CMS muon example

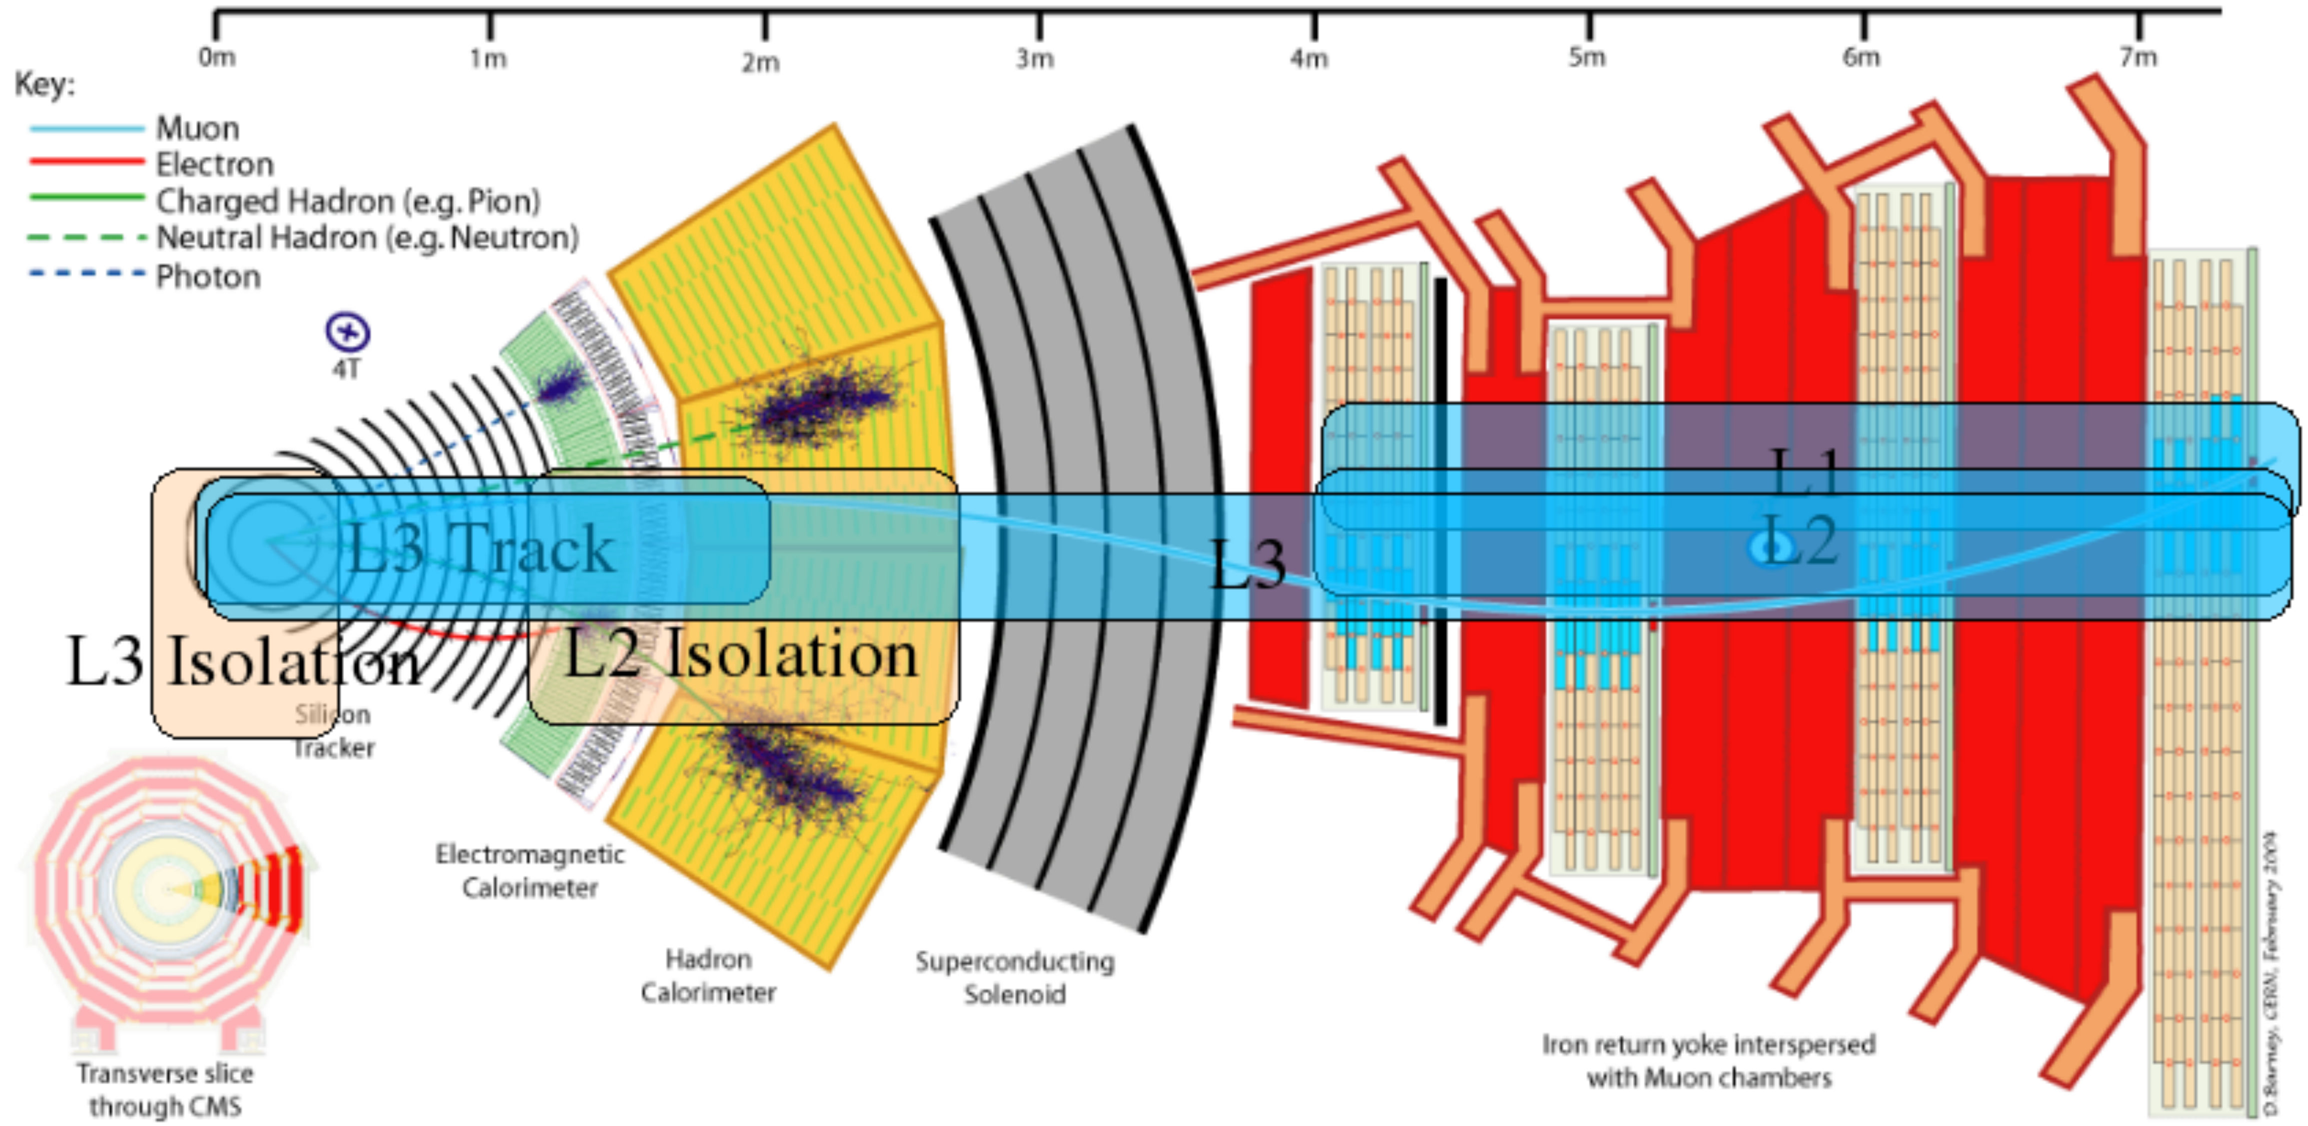

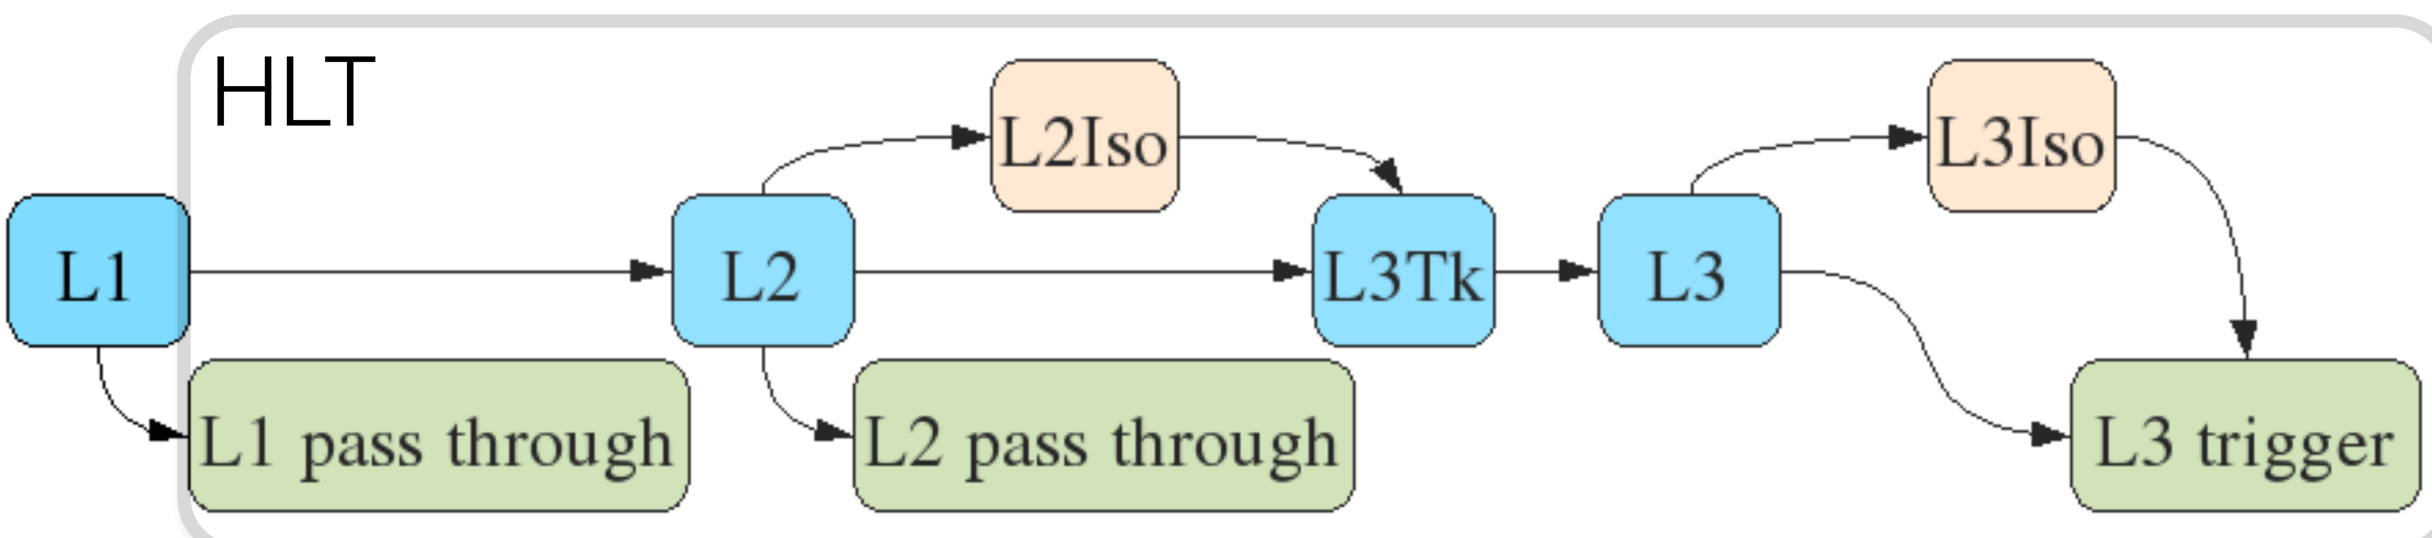

Changing conditions

A fill of the LHC machine lasts ~10 hours.
Exponential decay of the luminosity.

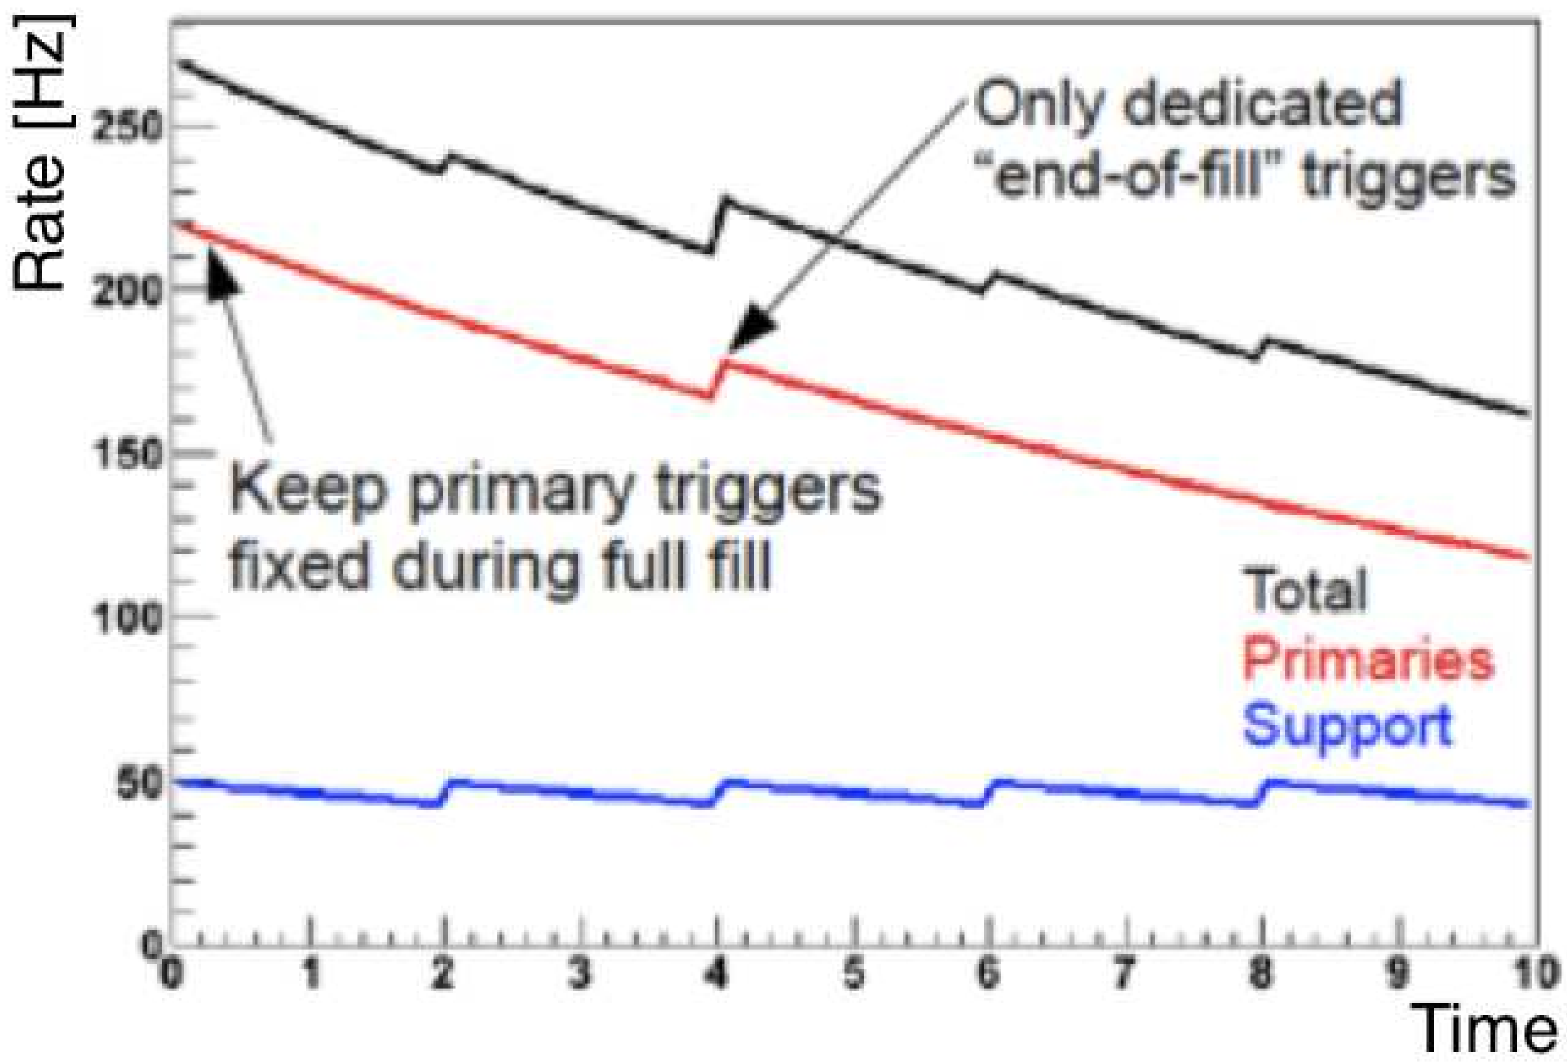

Changing conditions

Output bandwidth is a finite resource.

- The machine evolves
- And the physics priorities evolve too.

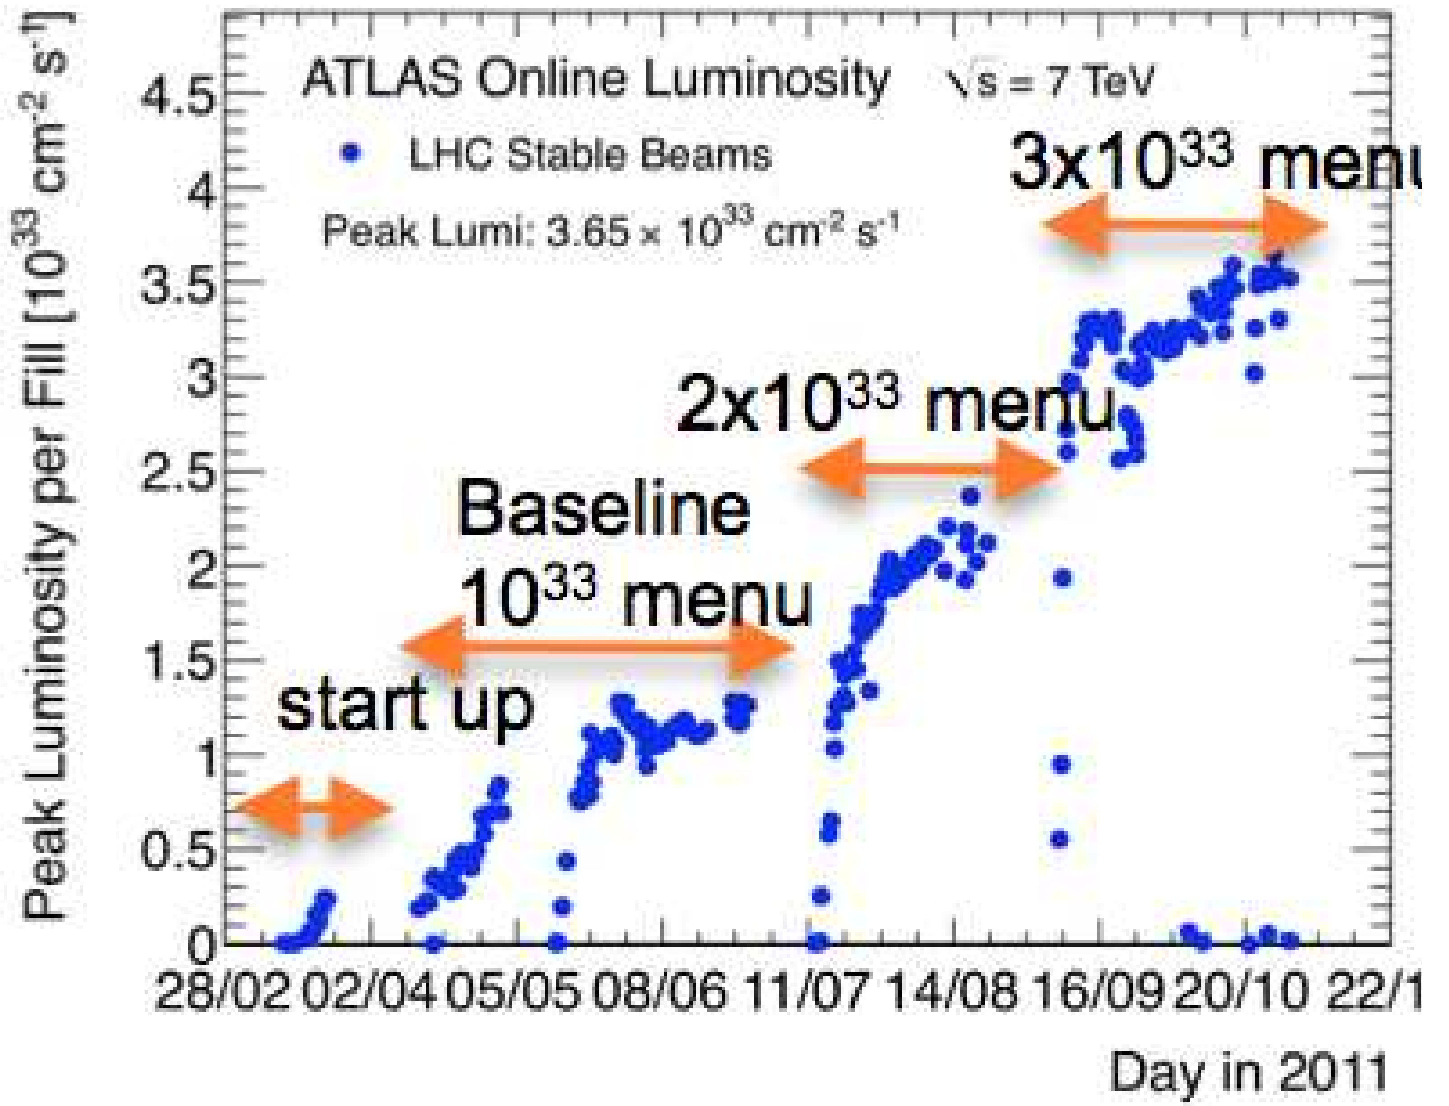

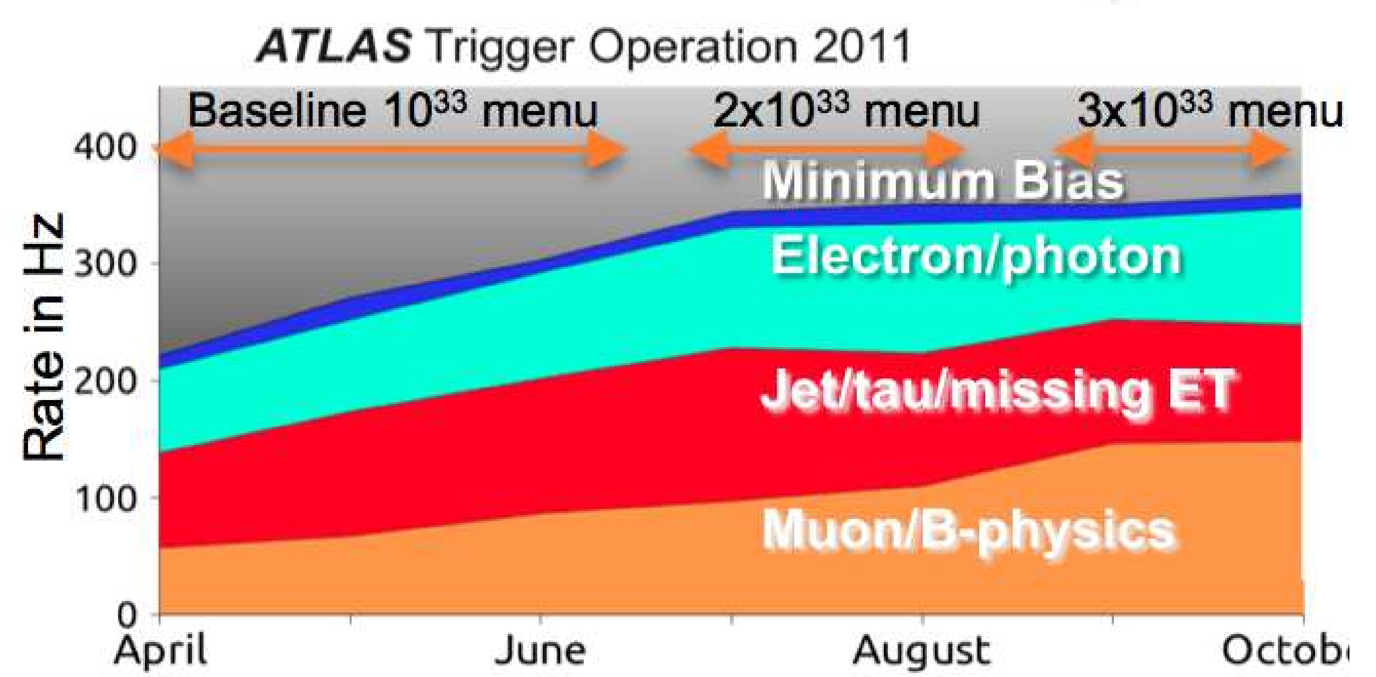

Requires plenty of interaction with analysis, and careful monitoring

MLHEP and triggers

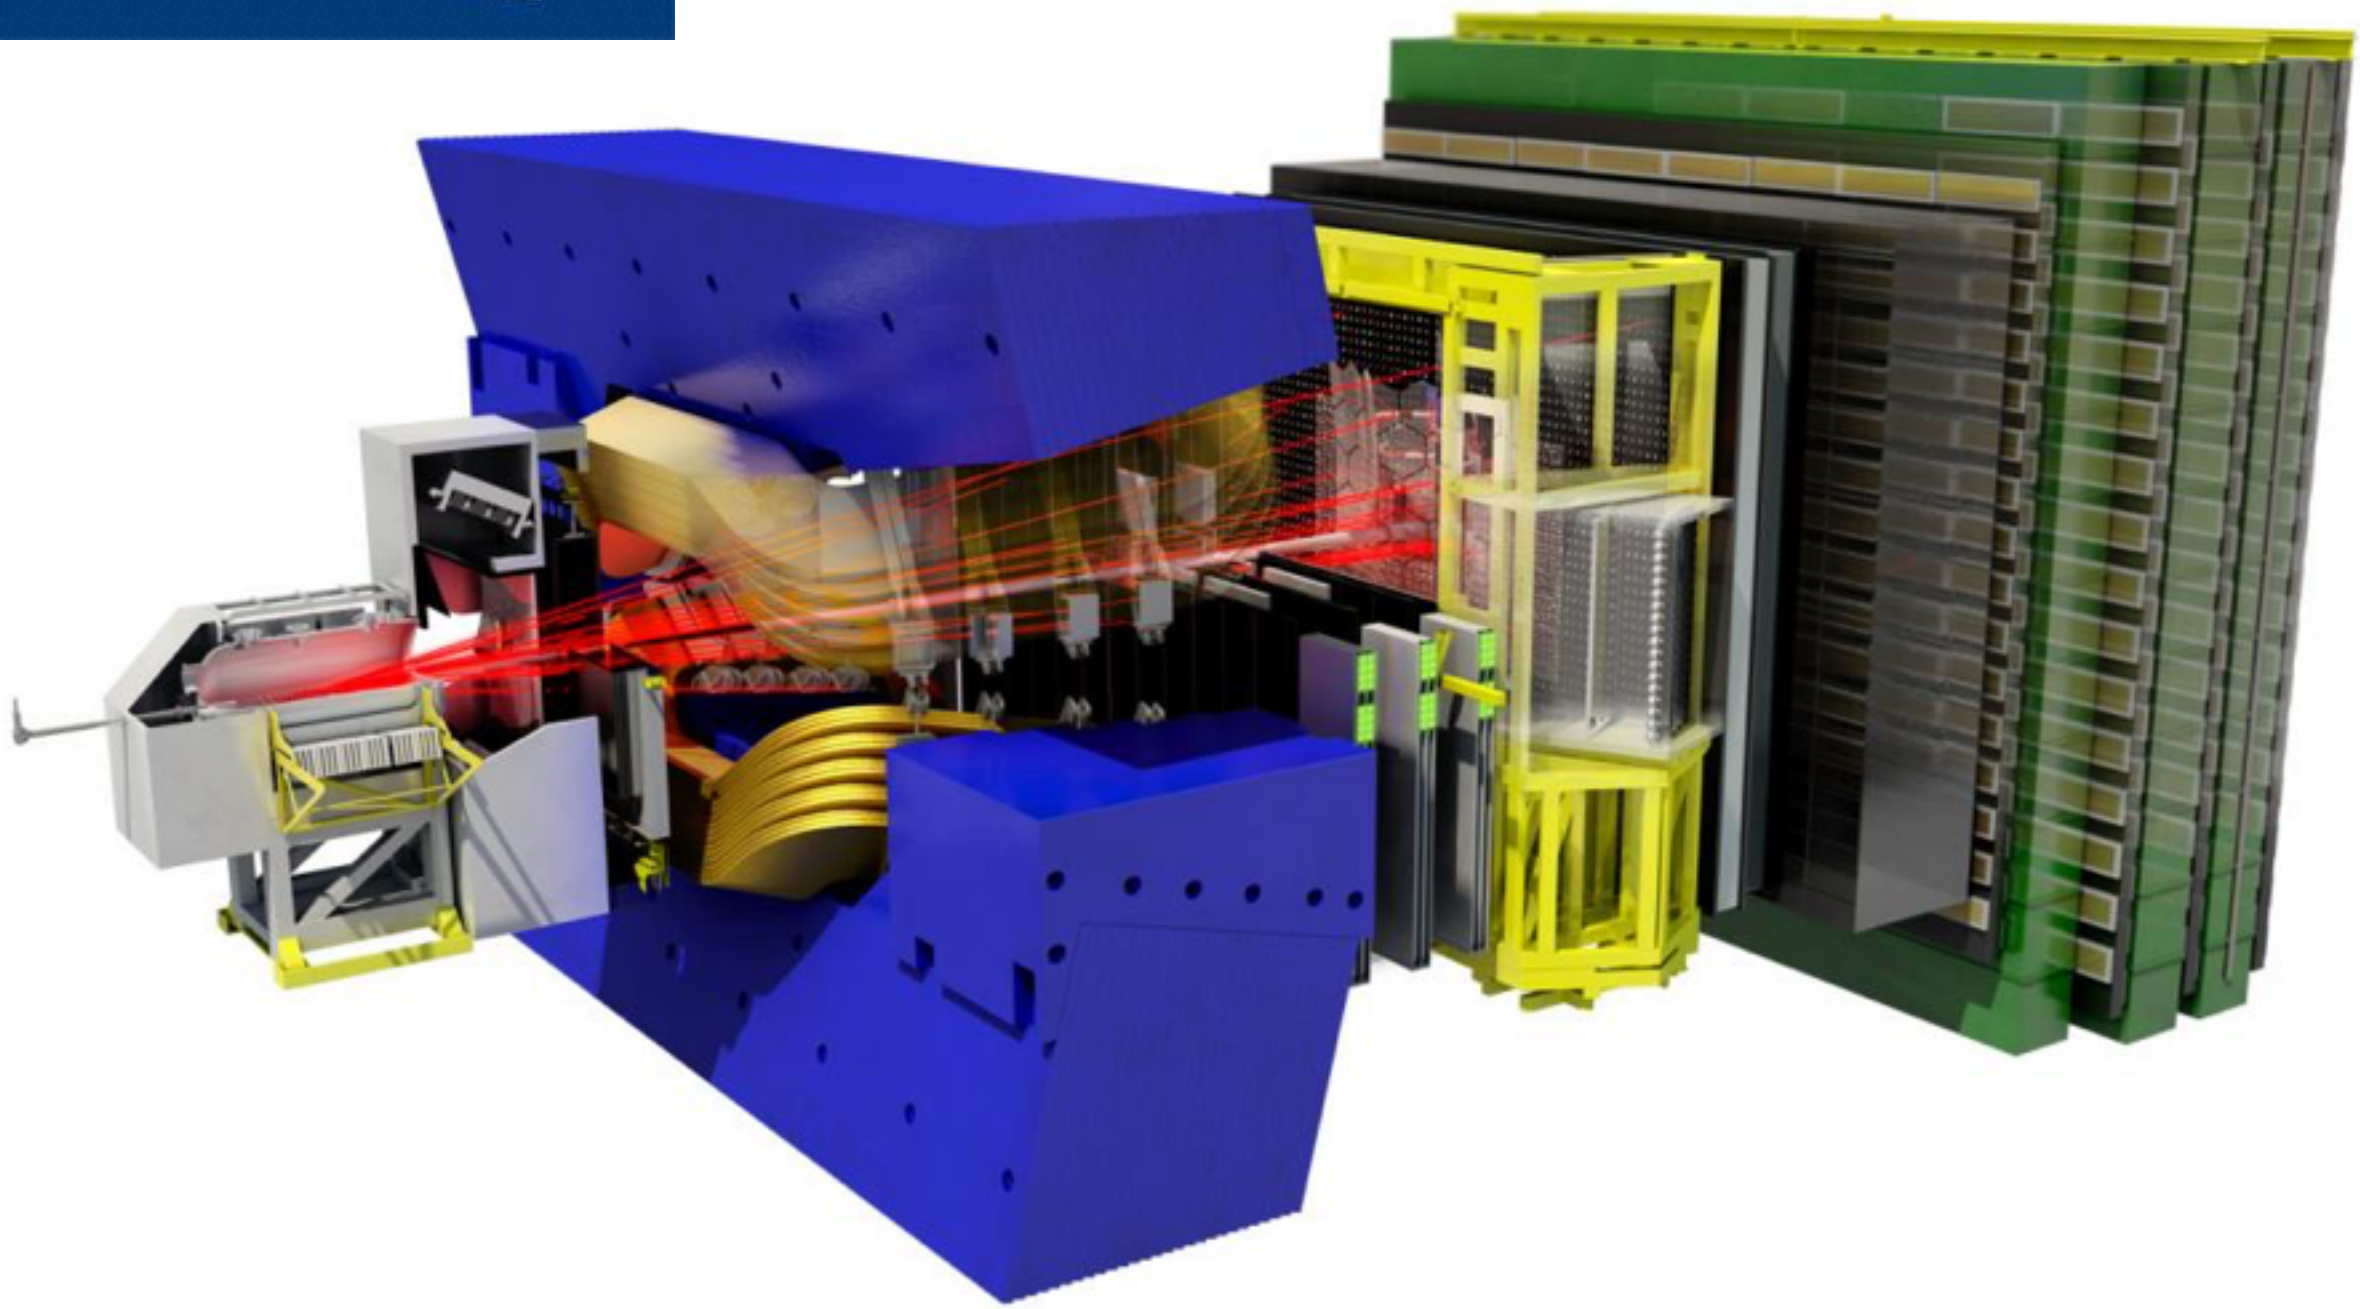

HLT needs to reduce the 1 MHz from L0 to ~10 kHz

~1 MHz in LHCb
(Run-I/II)

Charm

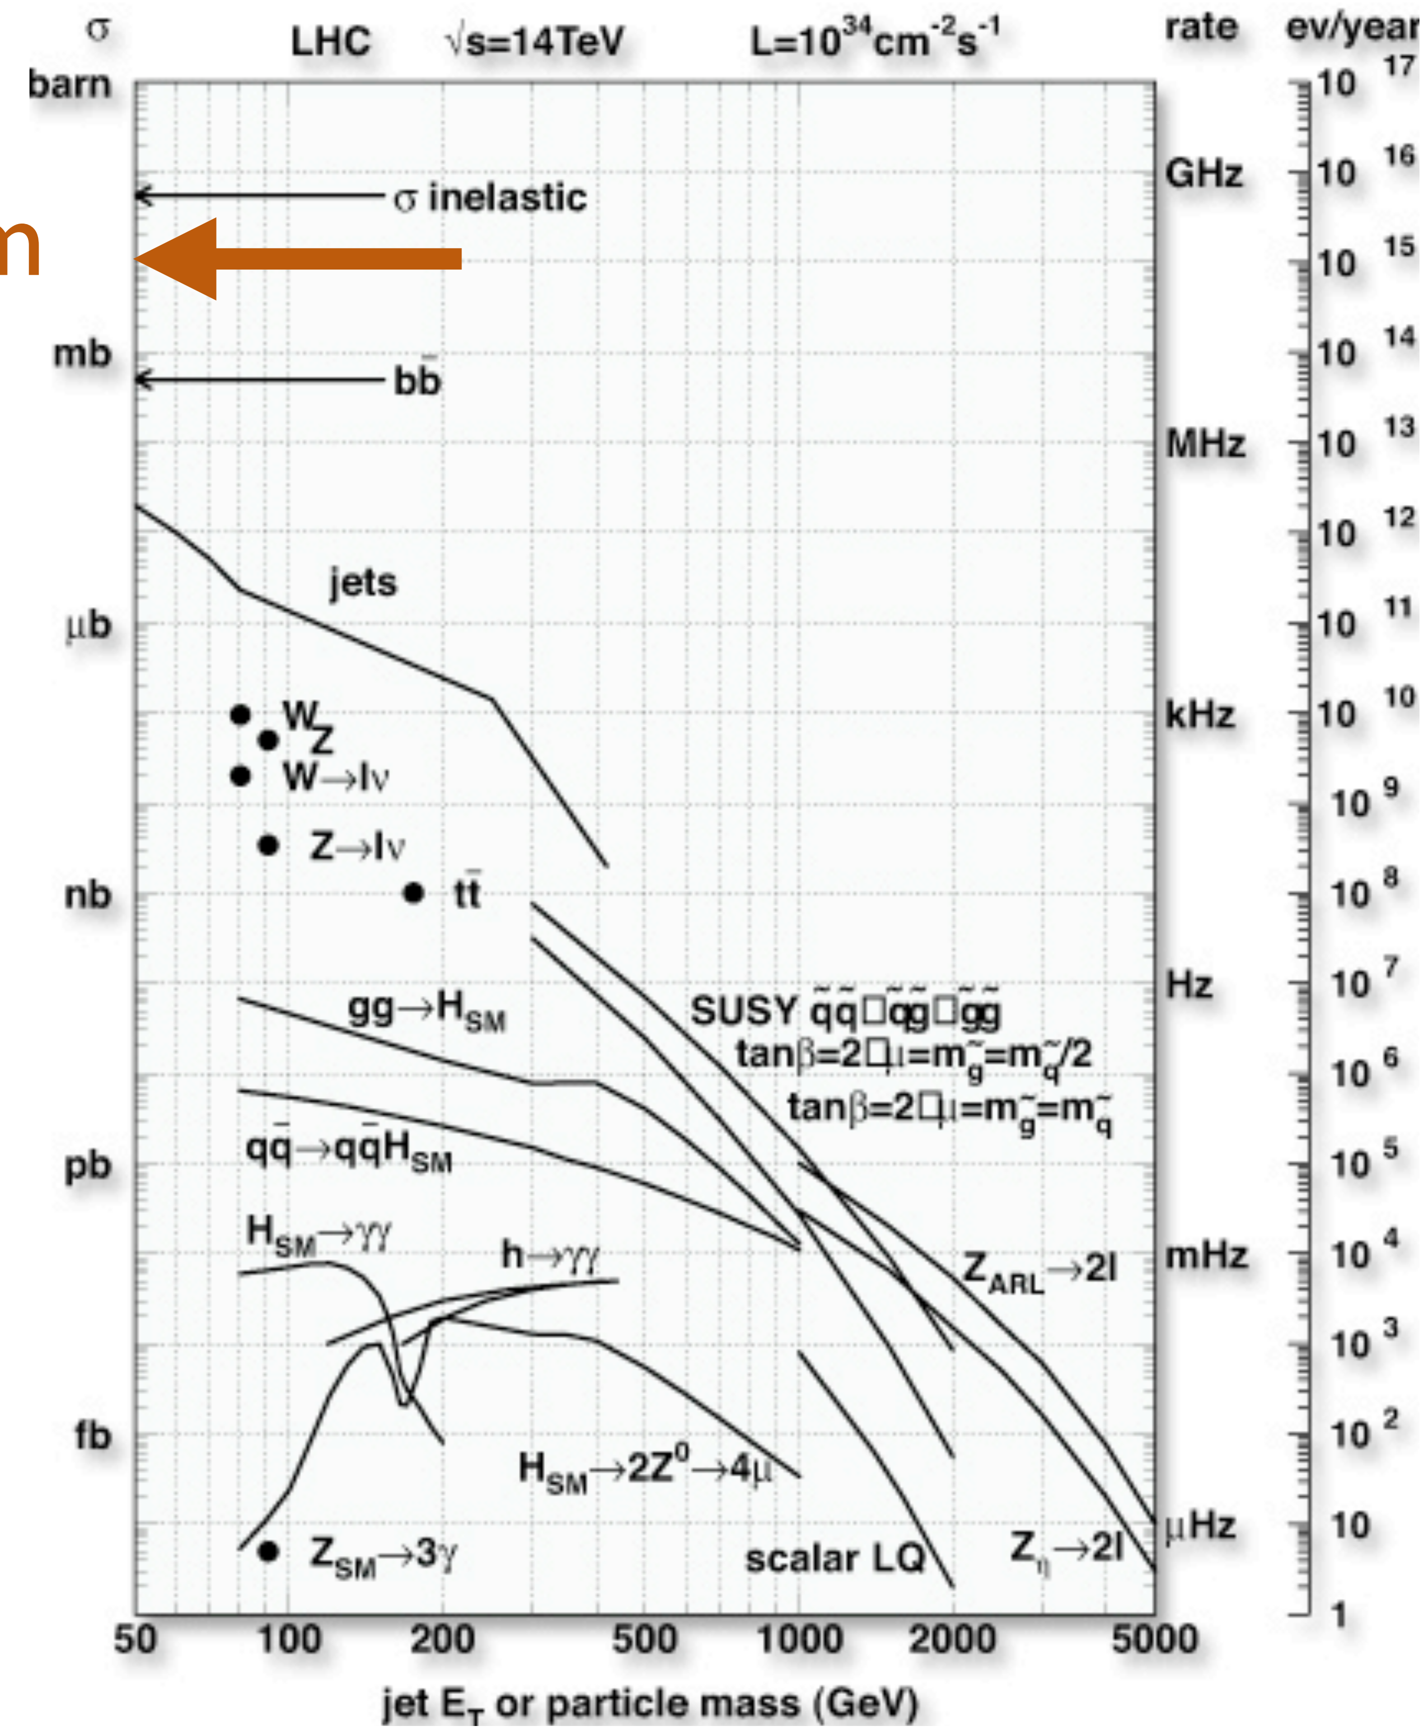

Even if we could inclusively select pure charm, we couldn't take all of it.

Need **exclusive selections** of the decay modes that we are interested in.

Still too much signal (50-100 kHz*)

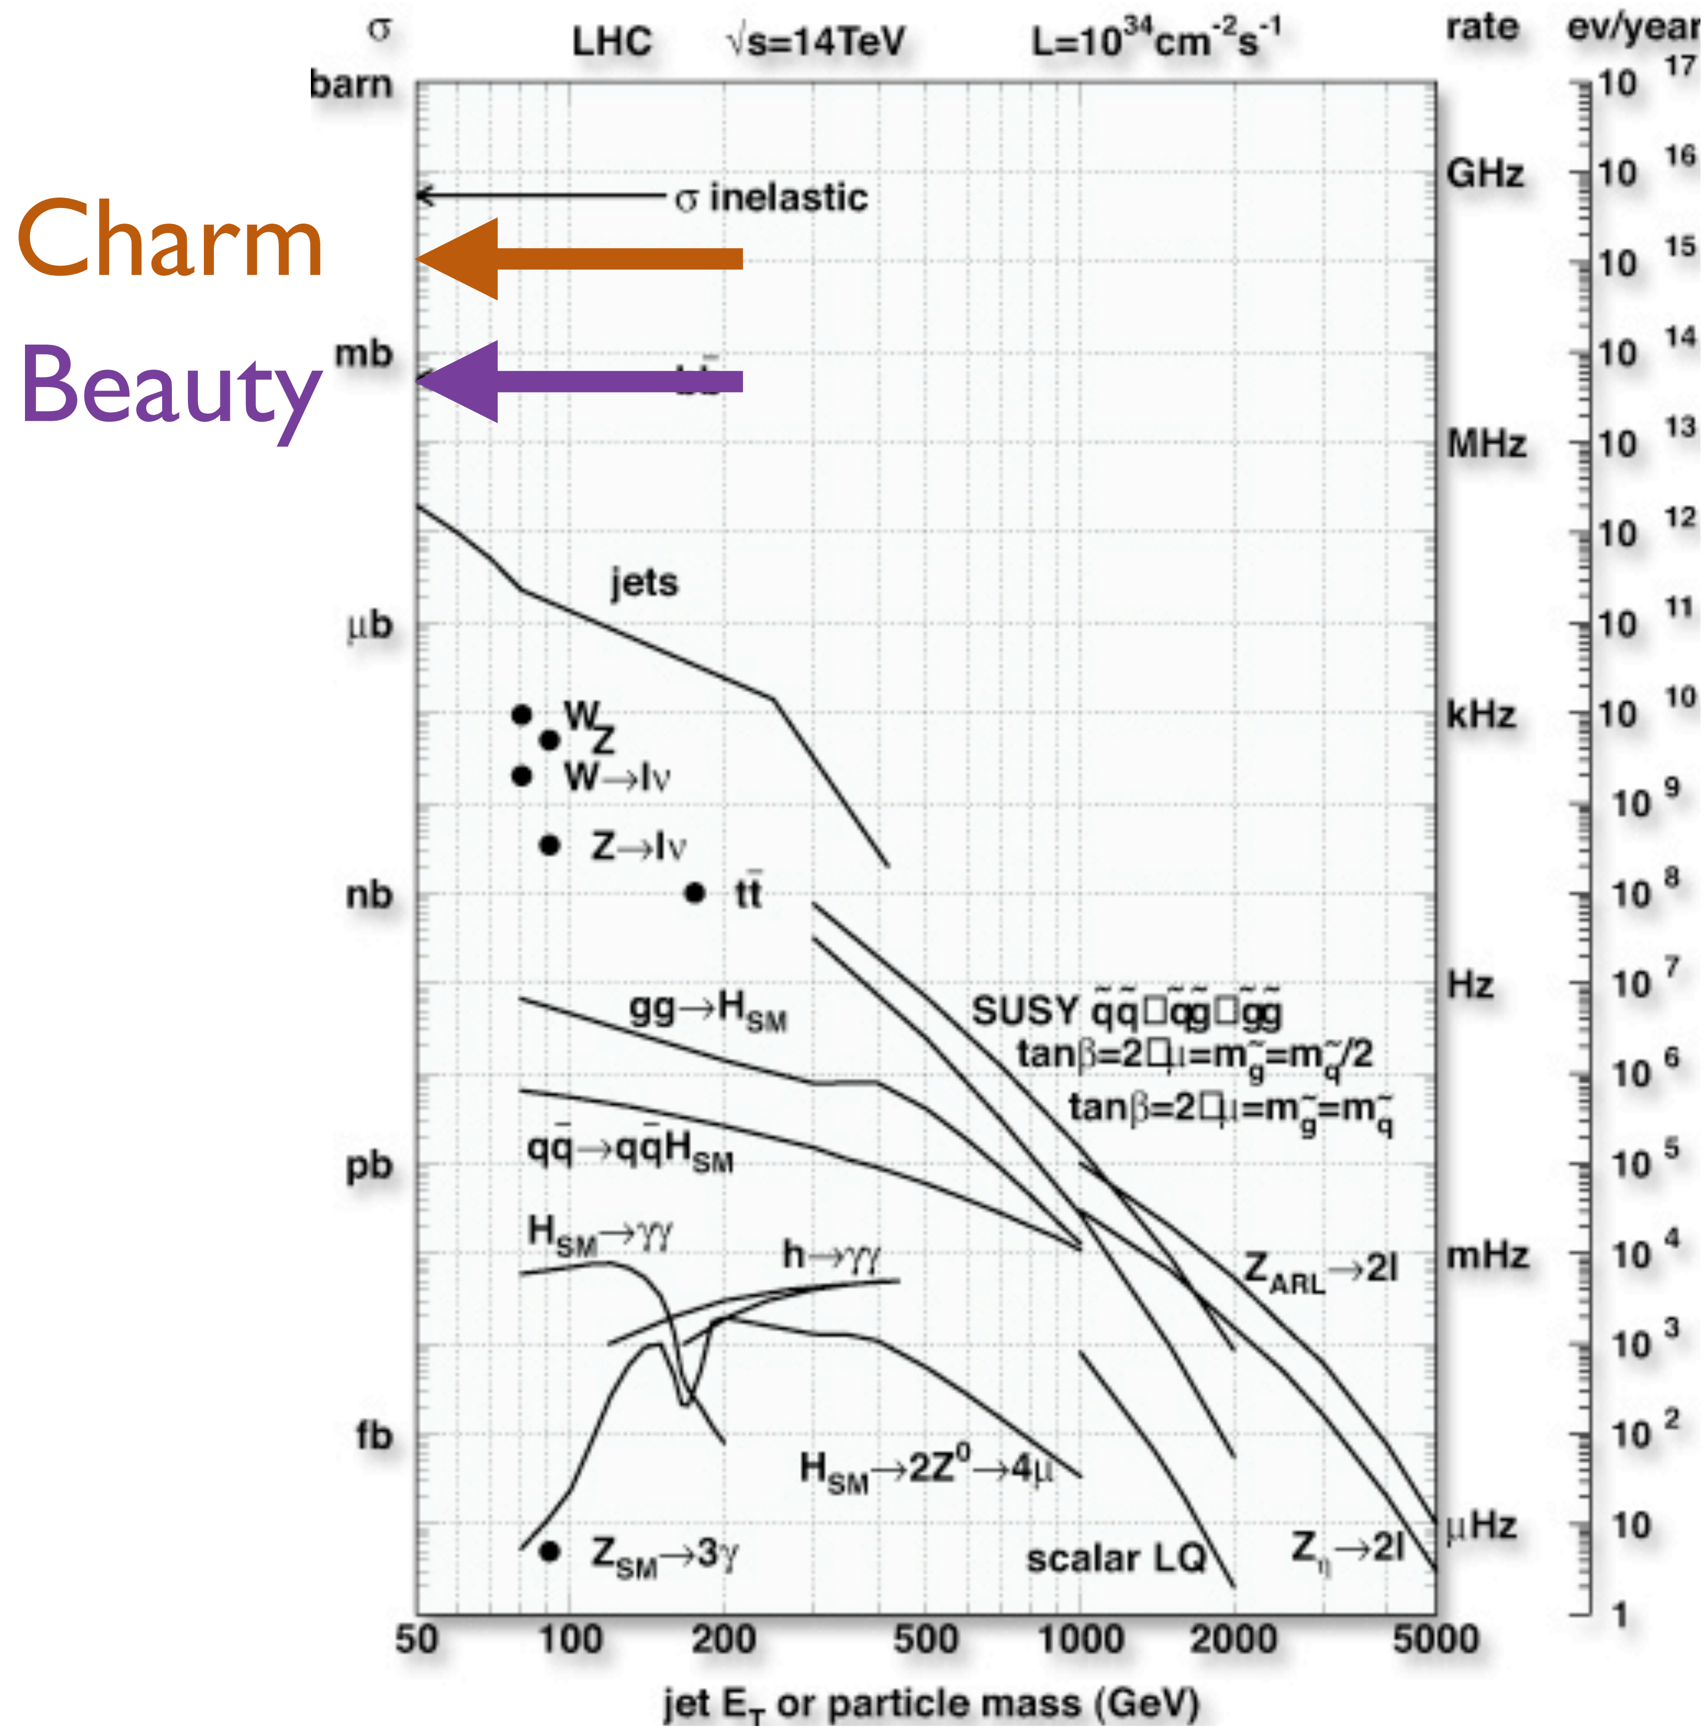

*In Run-I/II, roughly speaking

Charm

Beauty

“Good” beauty

$\tau > 0.2$ ps, $p_T > 2$ GeV

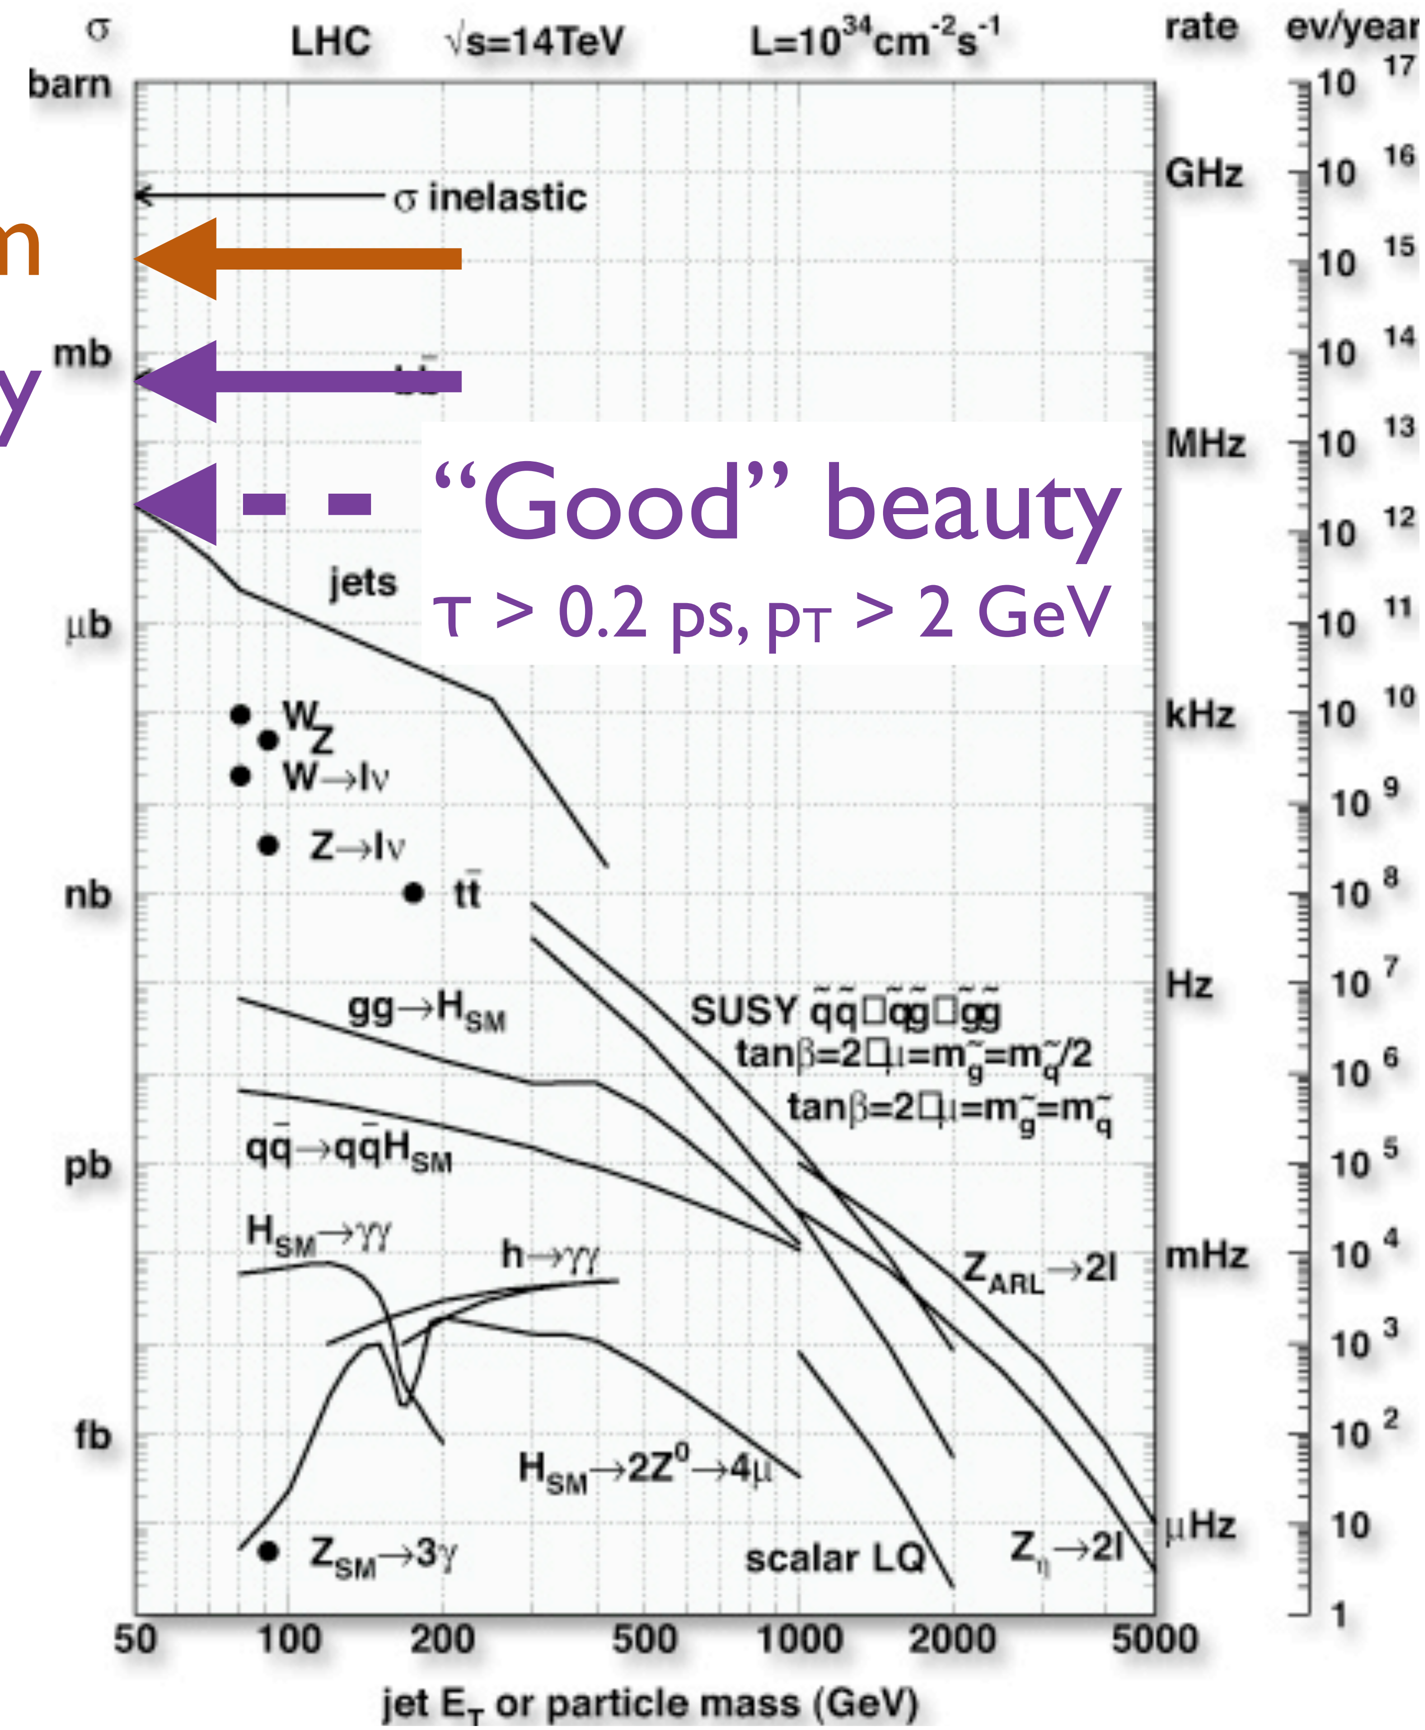

Still too much signal **(50-100 kHz*)**

Most of it is indistinguishable from background. Beauty with reasonable decay time and p_T can fit in a few kHz. Can have an **inclusive beauty trigger** #MLHEP

Run-I

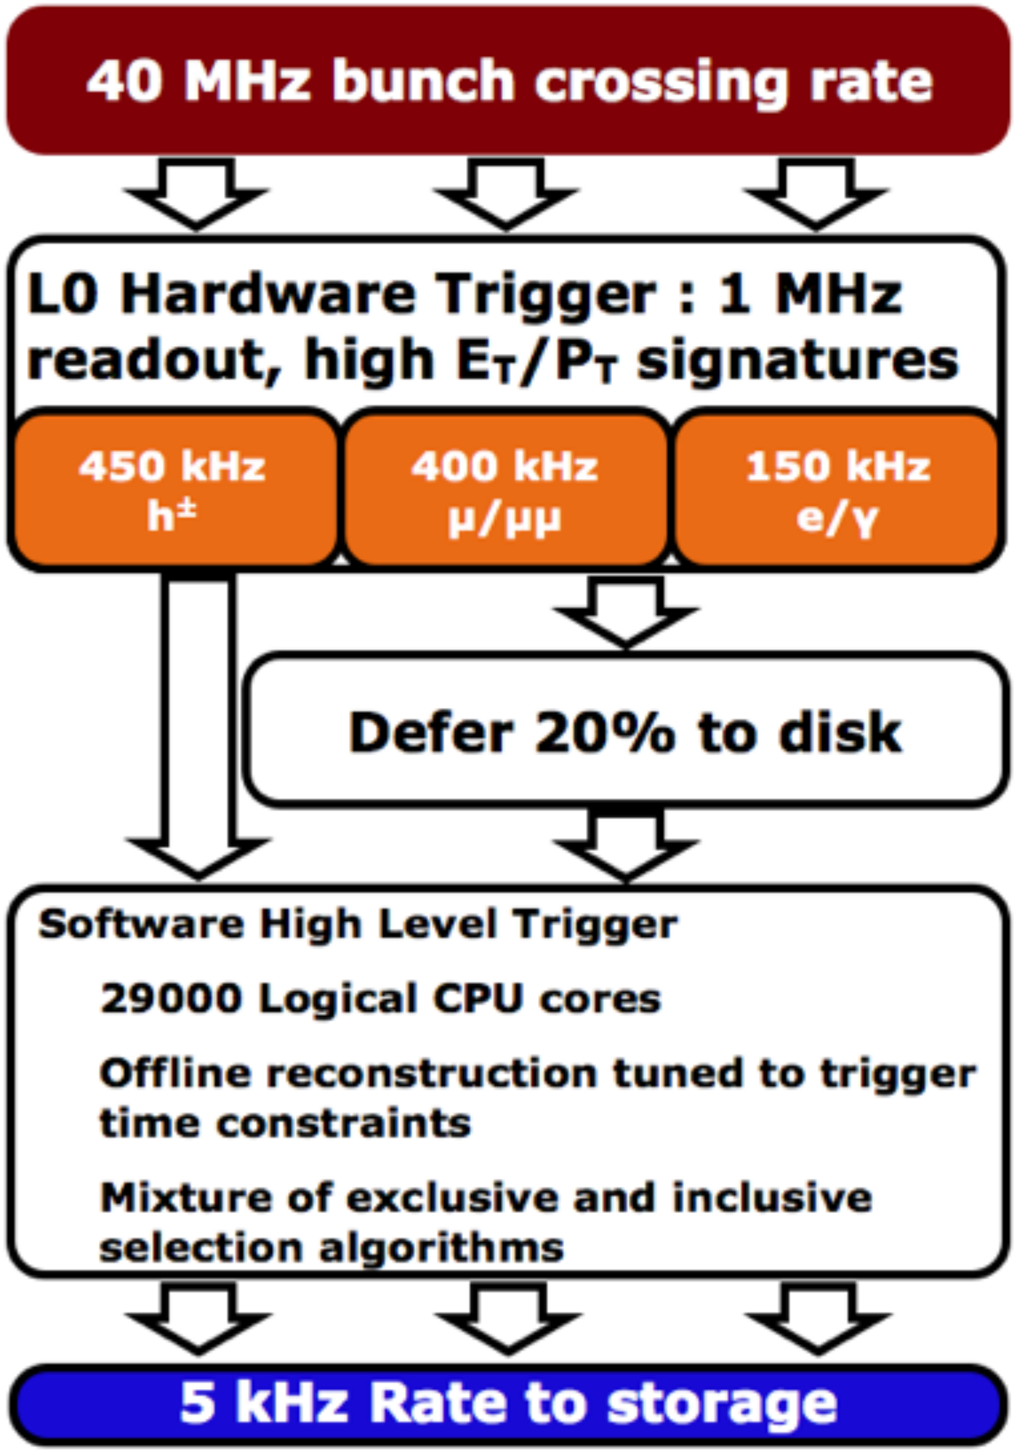

40 MHz bunch crossing rate
L0 Hardware Trigger : 1 MHz readout, high E_T/P_T signatures
450 kHz $h^{\pm}$
400 kHz $\mu/\mu\mu$
150 kHz e/γ
Defer 20% to disk
Software High Level Trigger
29000 Logical CPU cores
Offline reconstruction tuned to trigger time constraints
Mixture of exclusive and inclusive selection algorithms
5 kHz Rate to storage

Enabling real-time calibrations

In **Run-I**, the final alignments could arrive months after the data was taken — to be used in the final offline processing.

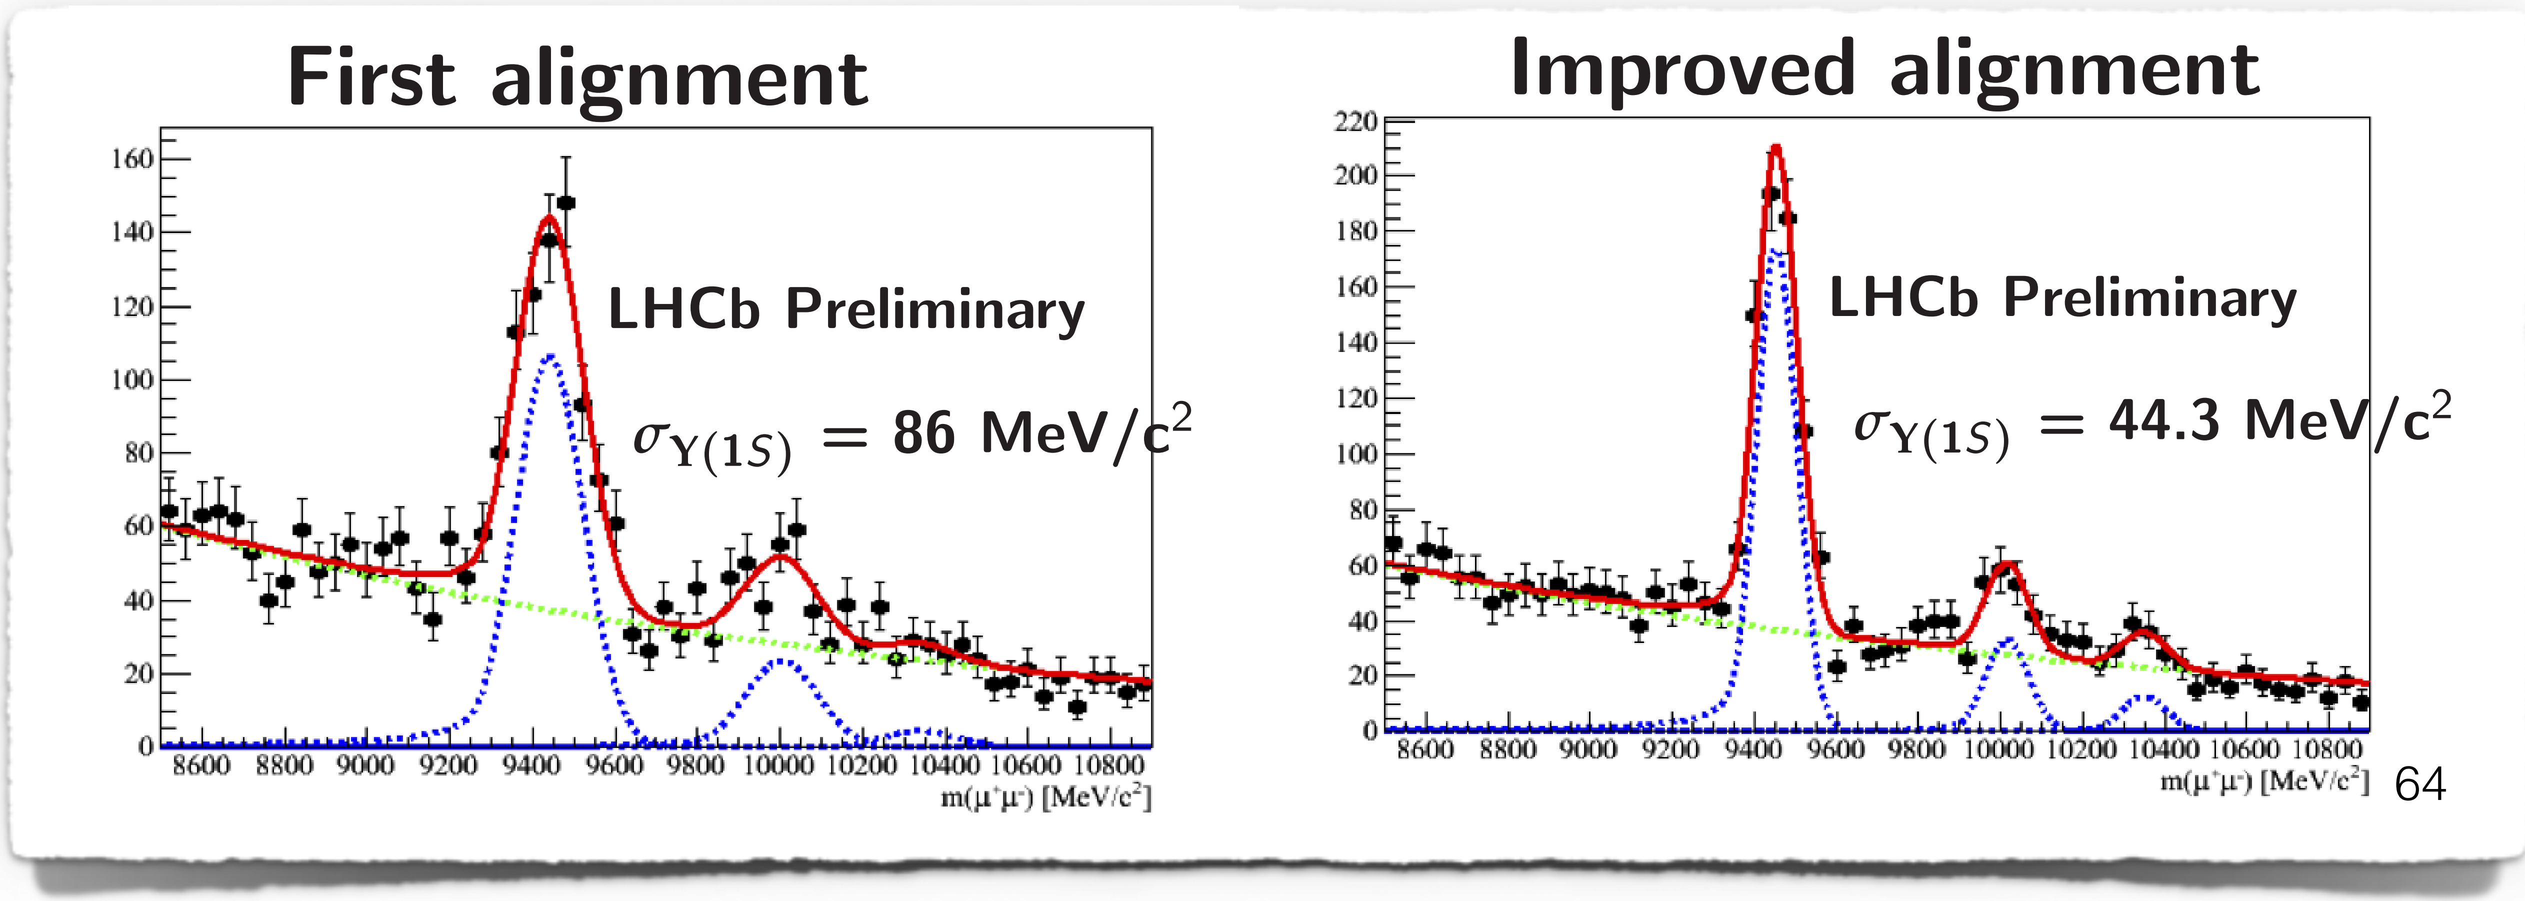

Run-II: alignment and calibration become part of the online workflow.

- Better information = better real-time selections
- And, we can do physics analysis on the HLT output!

Run-I

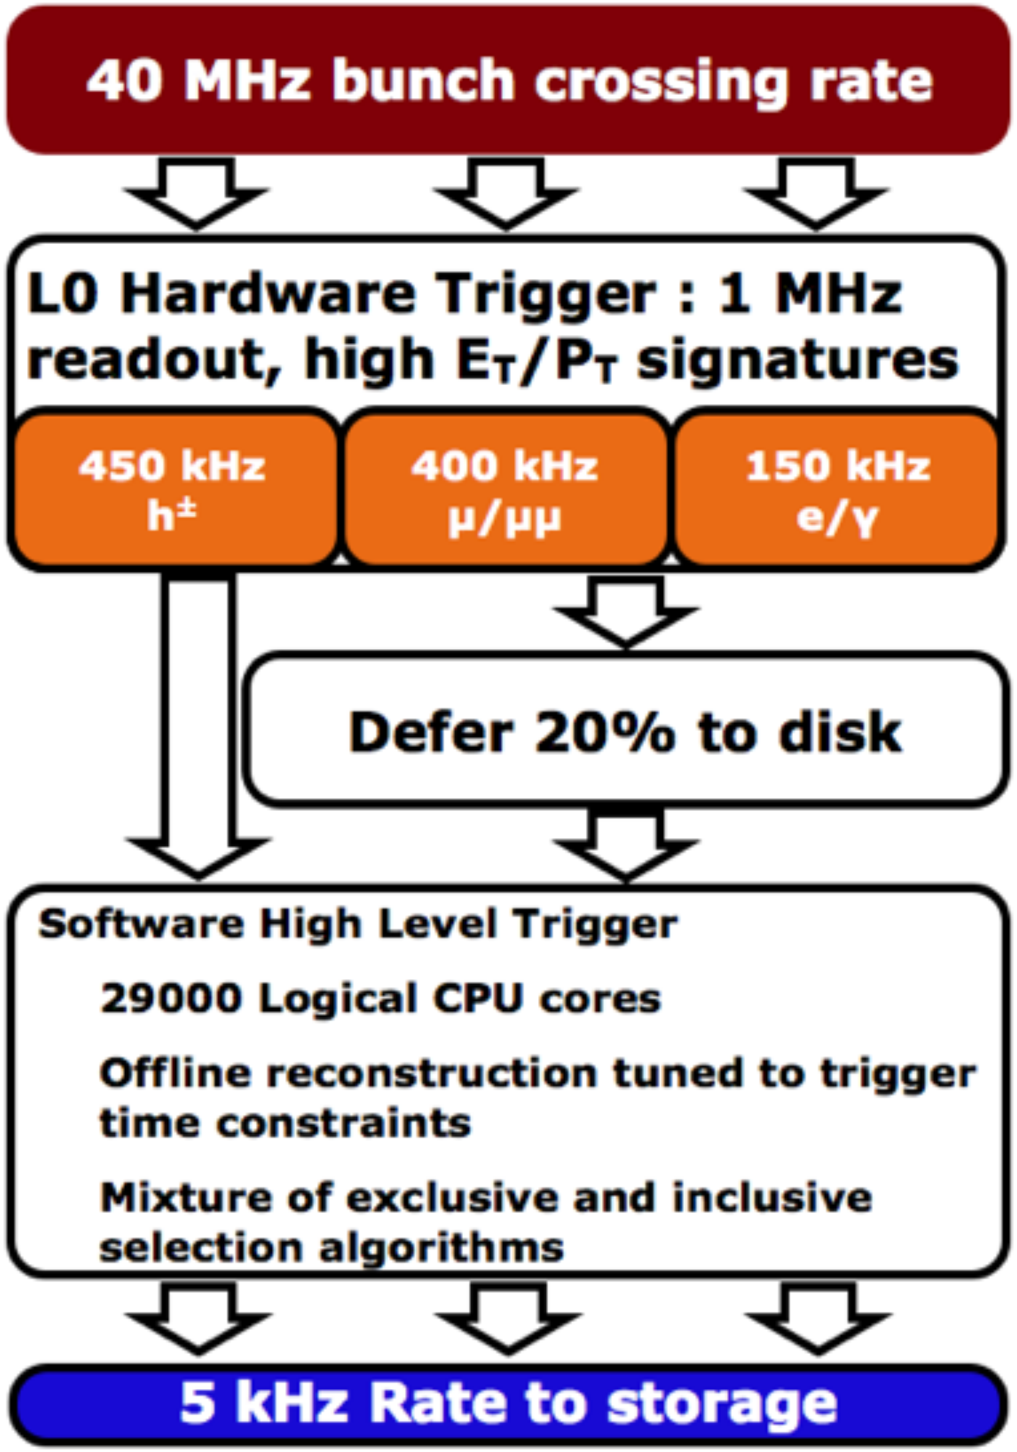
40 MHz bunch crossing rate
L0 Hardware Trigger : 1 MHz
readout, high E_T/P_T signatures
450 kHz
$h^{\pm}$
400 kHz
$\mu/\mu\mu$
150 kHz
e/γ
Defer 20% to disk
Software High Level Trigger
29000 Logical CPU cores
Offline reconstruction tuned to trigger time constraints
Mixture of exclusive and inclusive selection algorithms
5 kHz Rate to storage

Run-II

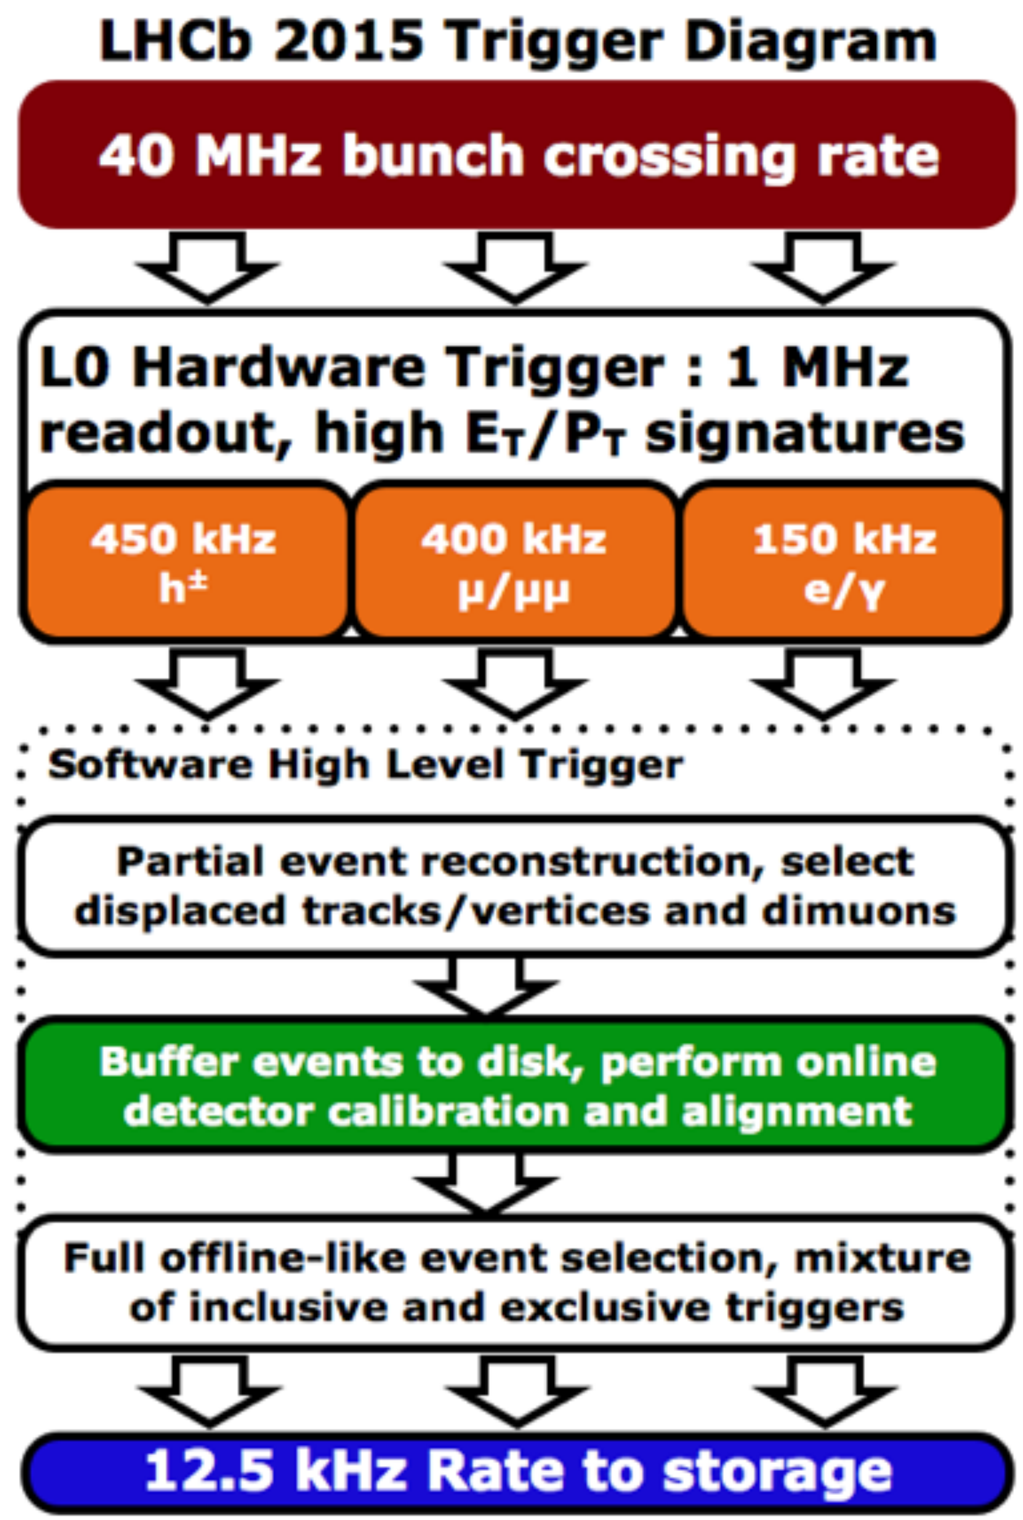
LHCb 2015 Trigger Diagram
40 MHz bunch crossing rate
L0 Hardware Trigger : 1 MHz
readout, high E_T/P_T signatures
450 kHz
$h^{\pm}$
400 kHz
$\mu/\mu\mu$
150 kHz
e/γ
Software High Level Trigger
Partial event reconstruction, select displaced tracks/vertices and dimuons
Buffer events to disk, perform online detector calibration and alignment
Full offline-like event selection, mixture of inclusive and exclusive triggers
12.5 kHz Rate to storage

The LHCb "Turbo" stream

- Bypass offline reconstruction (70 kB/event ➔ 5 kB/event)
- Perfect for the huge rate of charm exclusives.

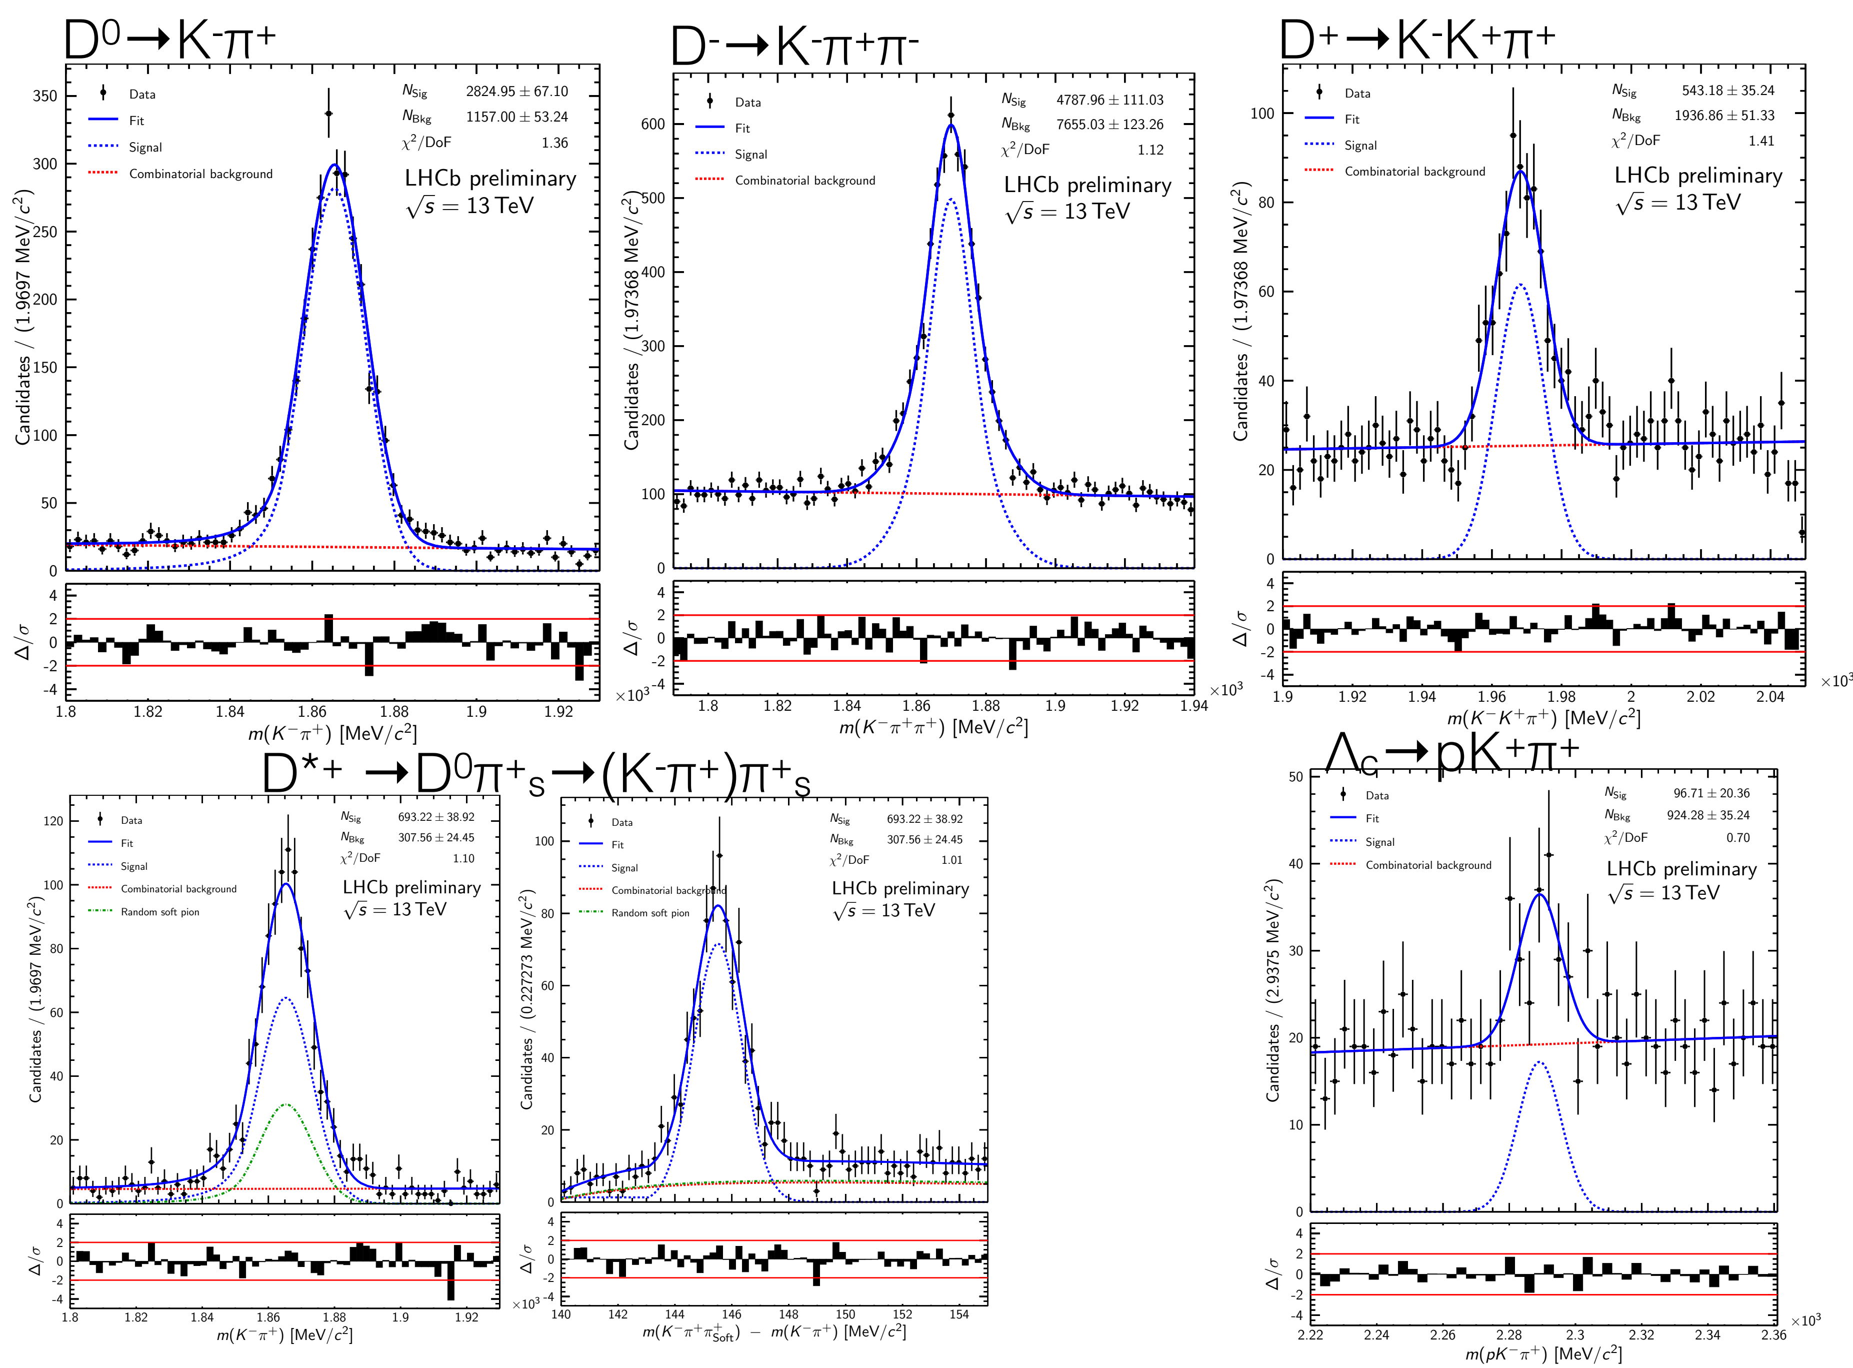

Allocating the 5-10 kHz bandwidth?

Inclusives

Inclusive dimuon

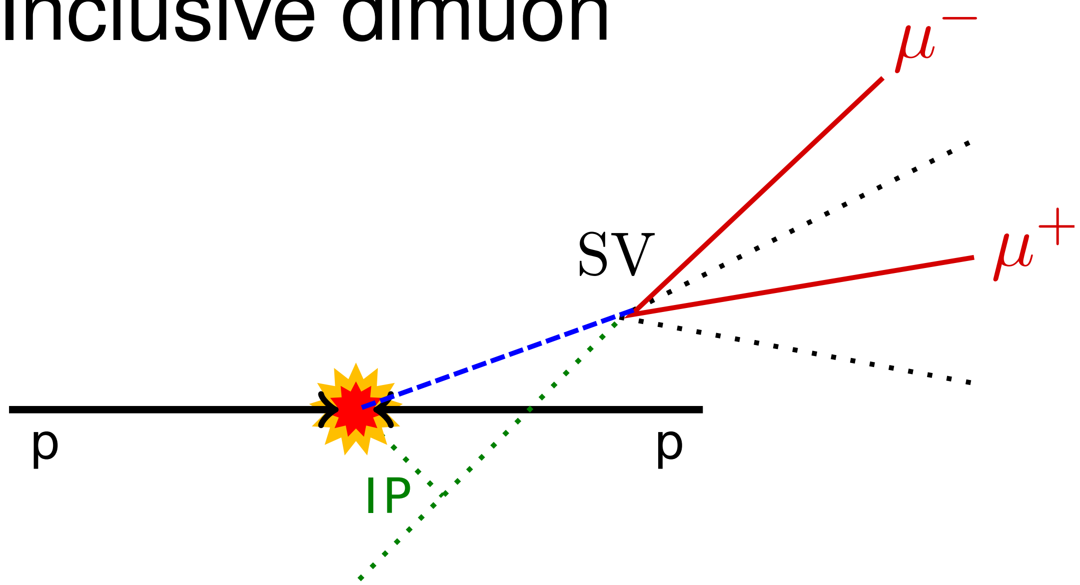

Topo N-body

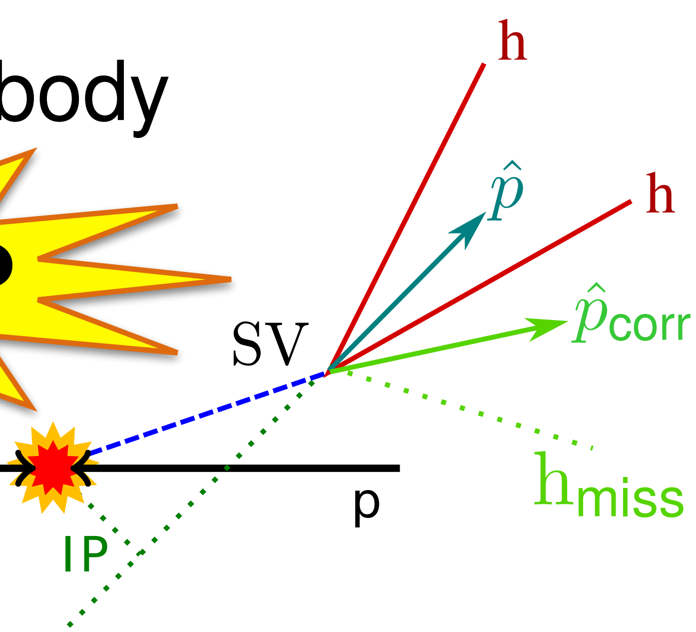

Exclusives

$B \rightarrow hh$

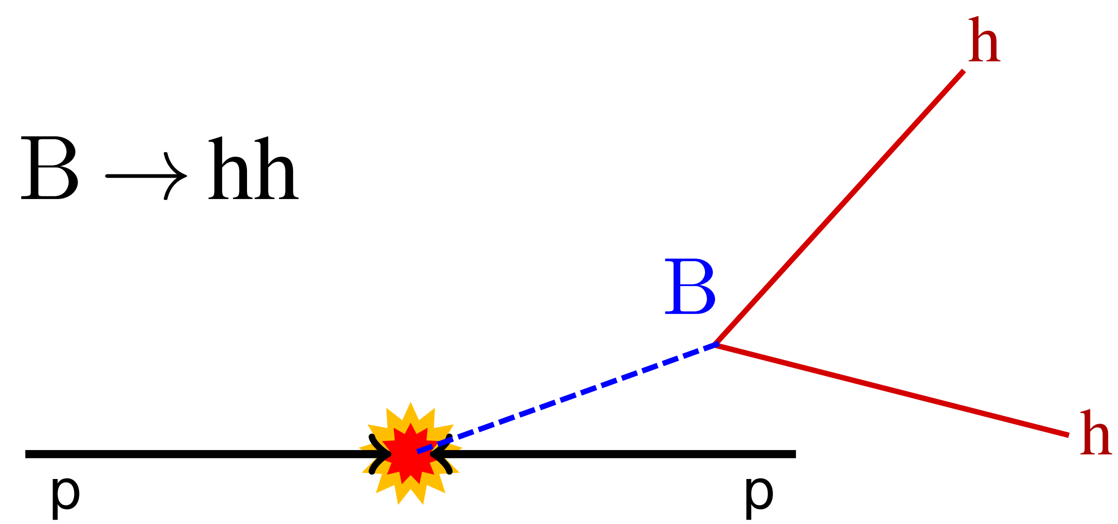

$D^{*+} \rightarrow D^0\pi^+$

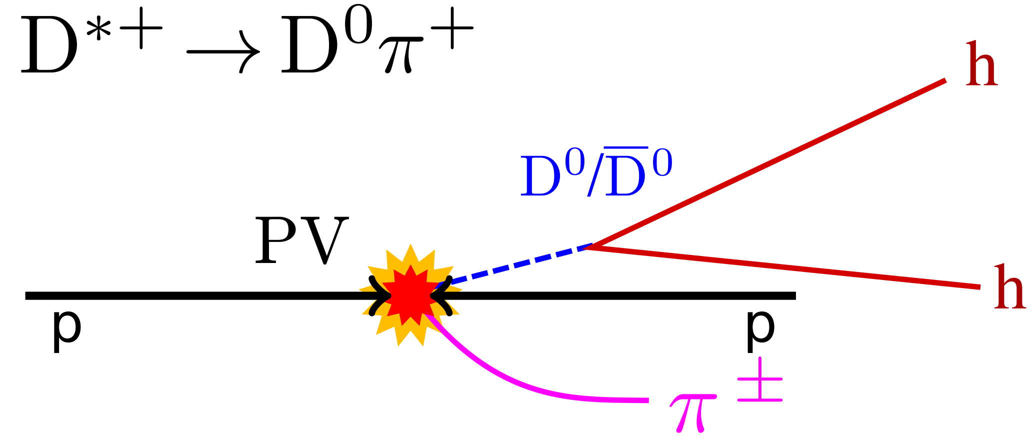

etc.. etc…

Bandwidth division in Run-I (2012)

Rate	**5 kHz**
Lines	**~200**

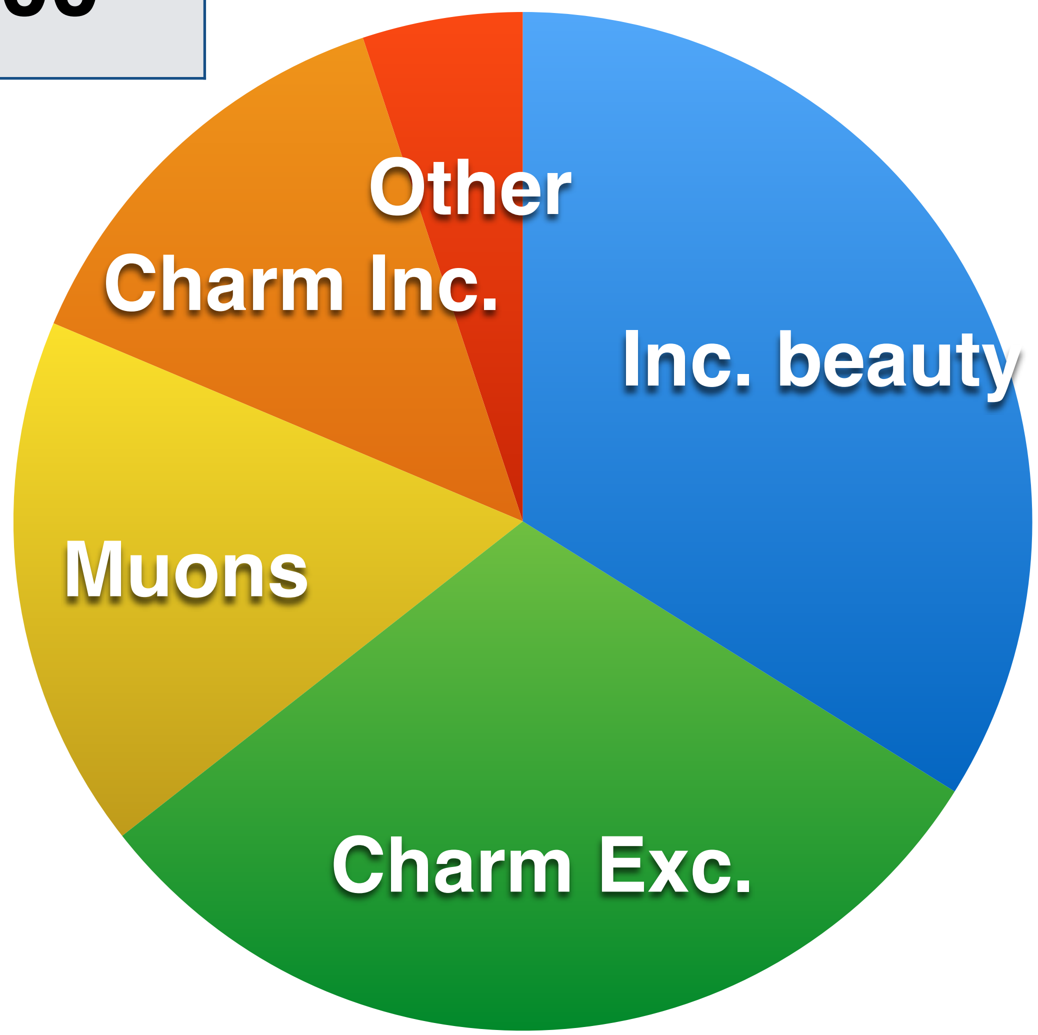

Bandwidth division in Run-II

Rate	12.5 kHz
Lines	~400

10 kHz
“Full” stream

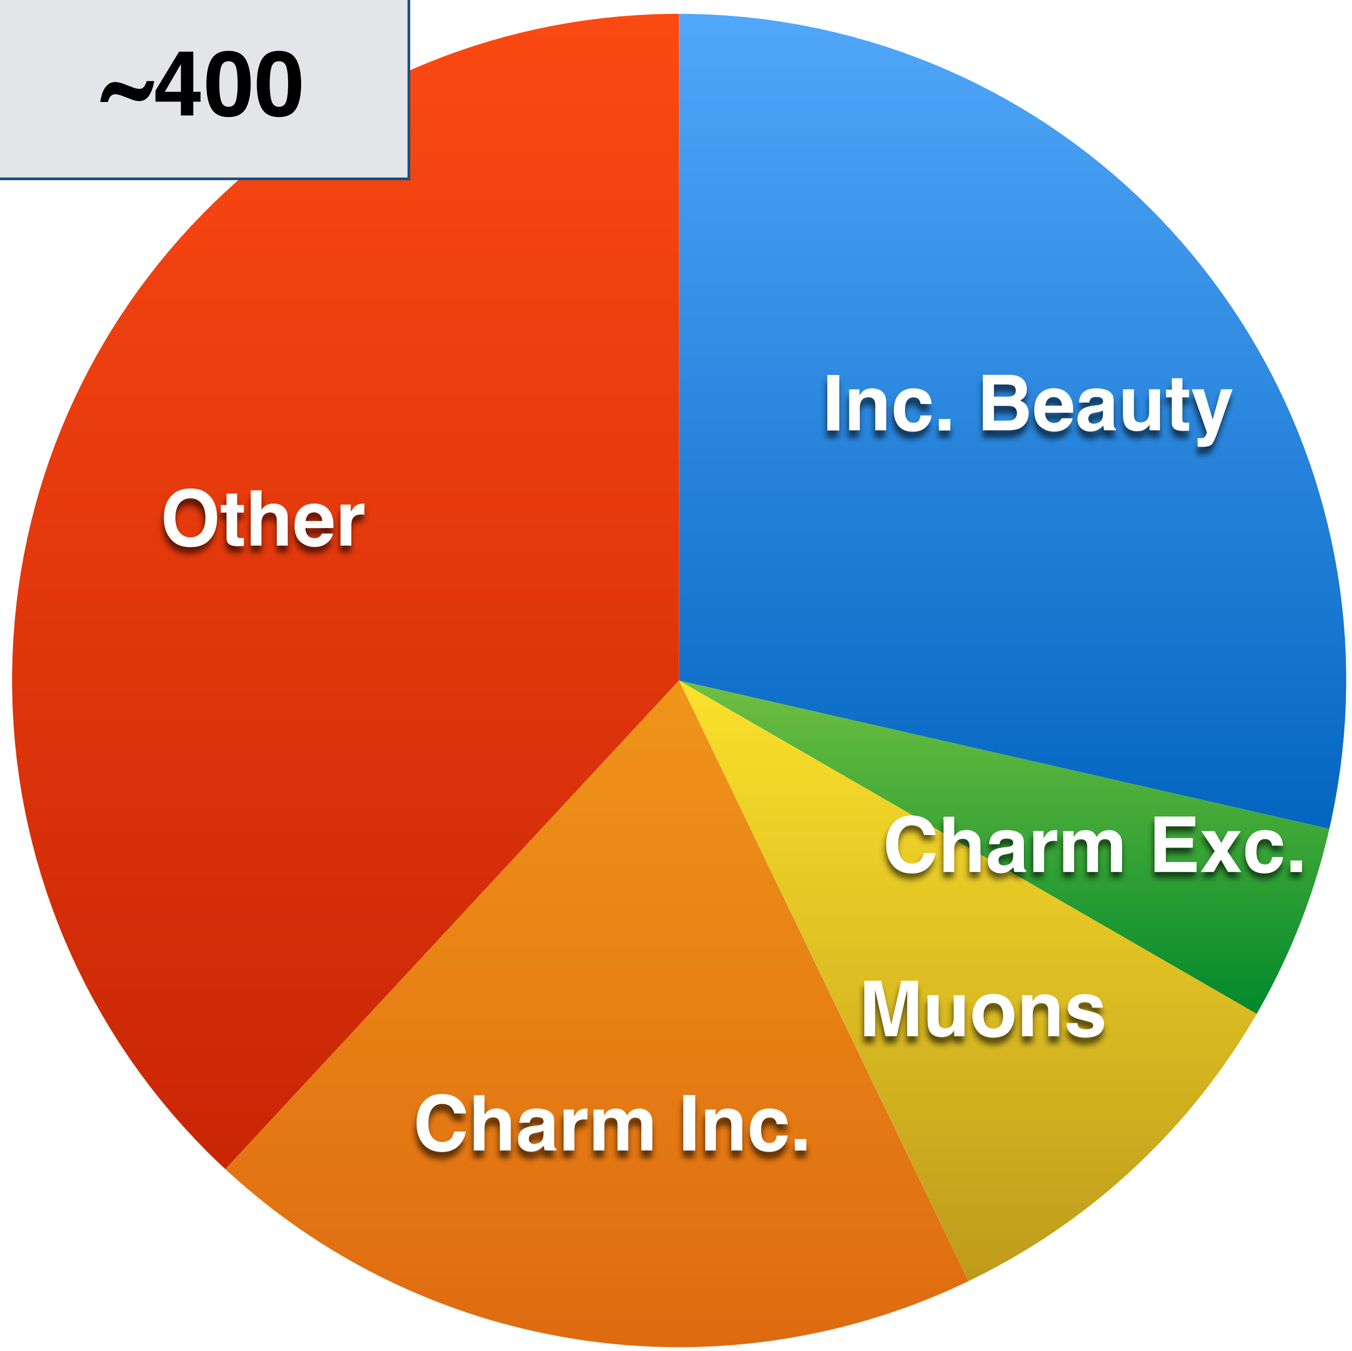

+2.5 kHz
“Turbo” stream

Managing 100s of exclusive lines

(Inclusives prefered)

Inclusive B trigger

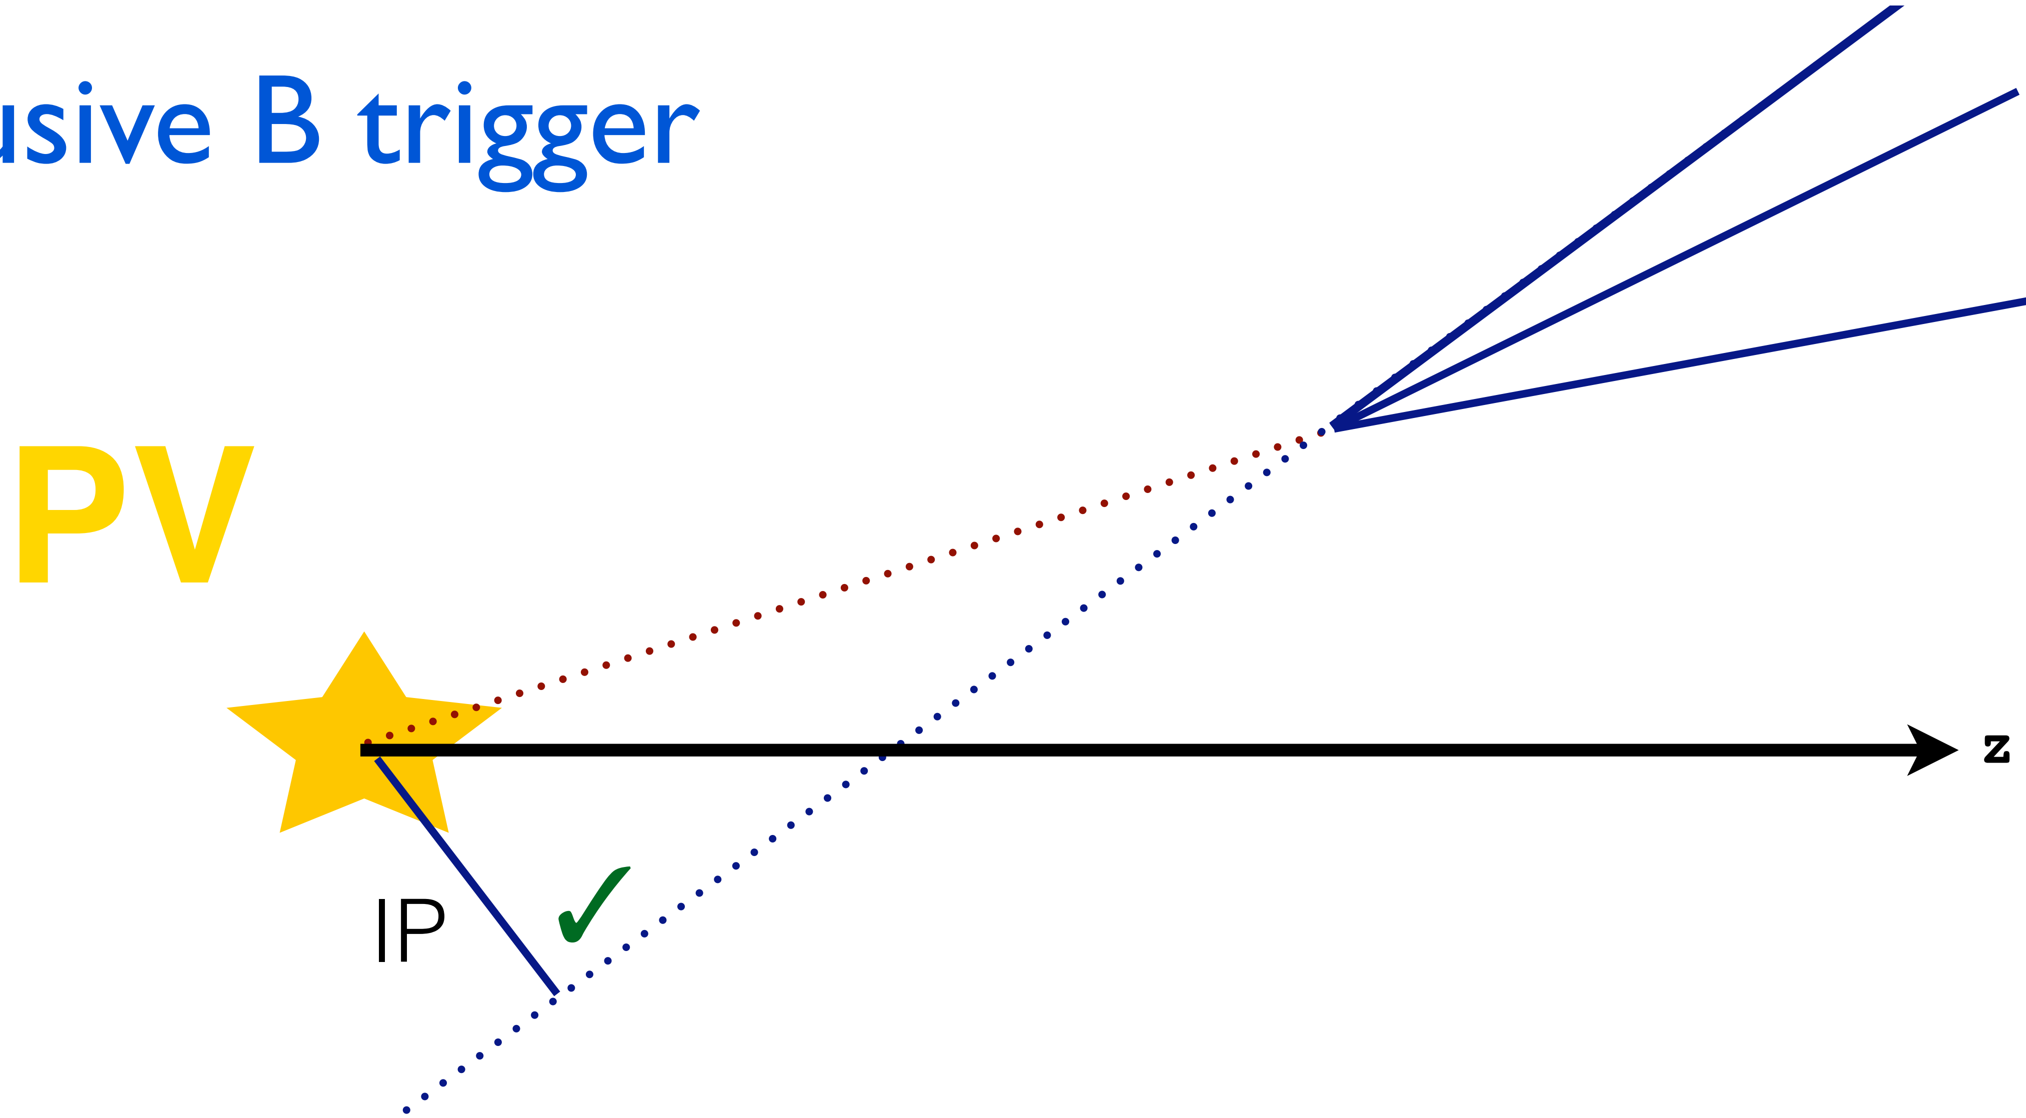

The goal (Run-I)

*A small set of line which select any "reasonable"**
B hadron decay with > 80% efficiency, and with an output rate of 2 kHz.

And fast:
~50k processes @ 1 MHz: ~50 ms to make a decision

*All charged daughters reconstructible and $p_T(B)$ > 2 GeV and $\tau(B)$ > 0.2 ps.

Staged HLT

CPU budget: 50(25) ms in Run-II(I)

Complete offline event reco: ~1 second

Need a staged approach

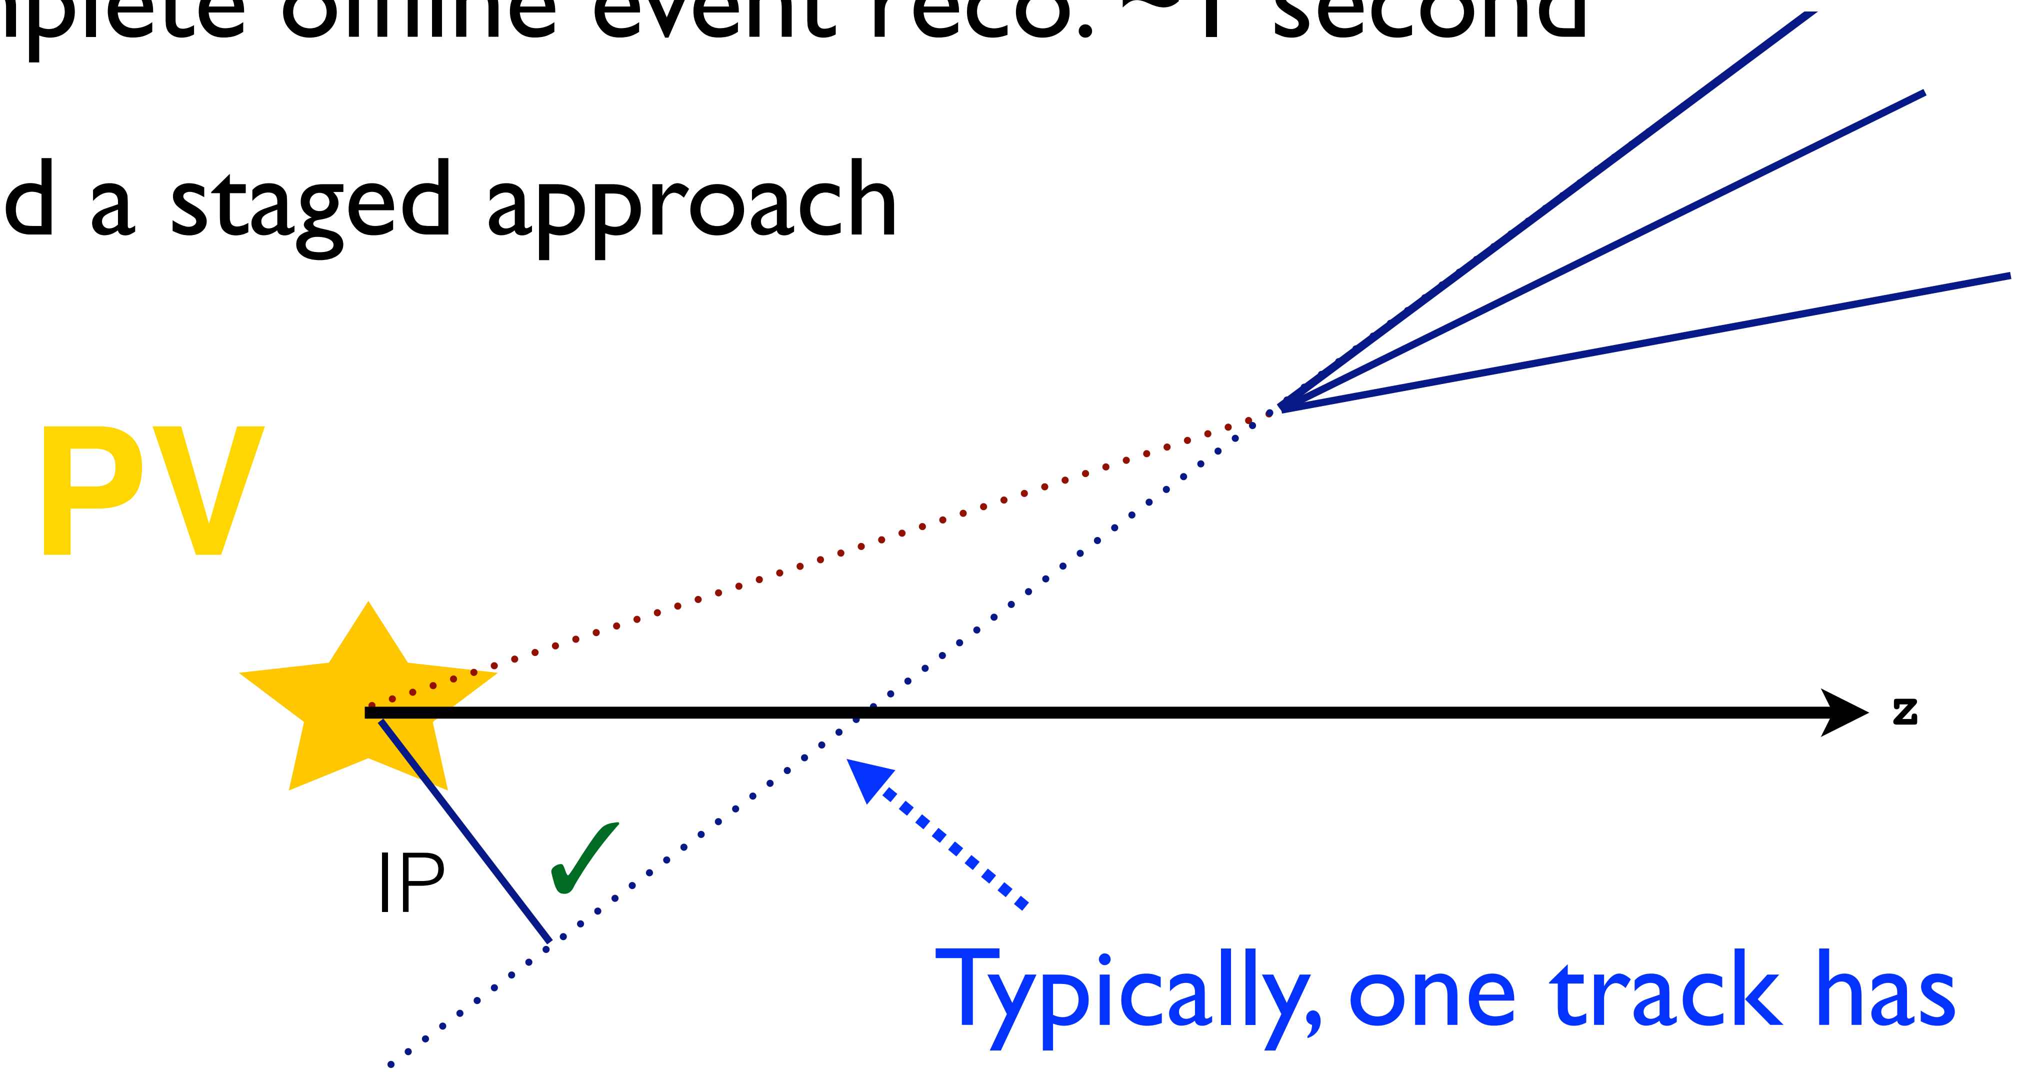

Typically, one track has large p_T and large IP

Start with tracking

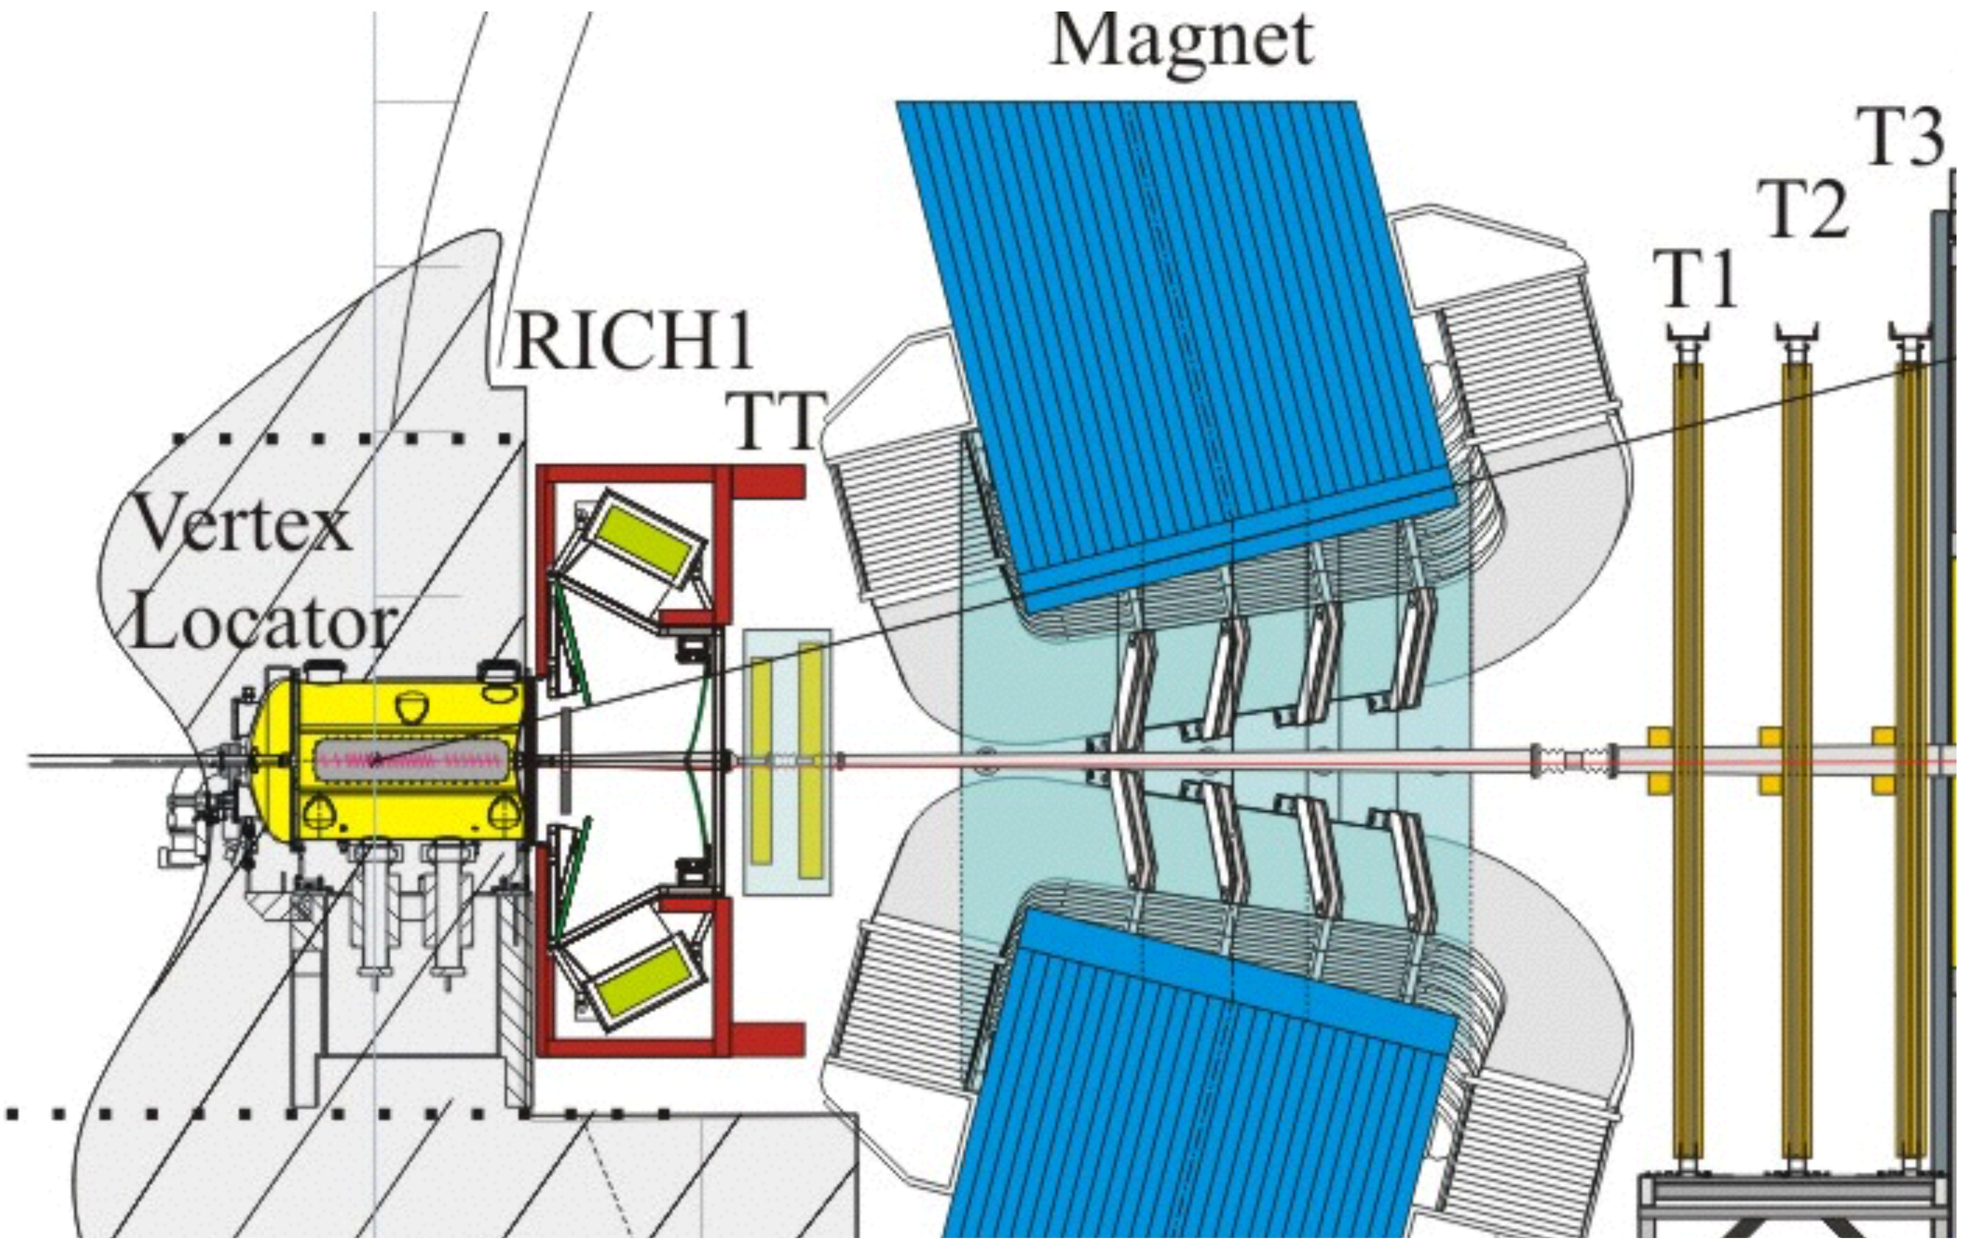

20-200 ms

depending primarily on how low you want to go in p_T
and how well you want to fit the tracks

Run-I

1 MHz

Tracking
p_T > 1250 MeV
Incl. selection

HLT1
15 ms/ev

70 kHz

Tracking
p_T > 300 MeV
Incl. + Excl.

HLT2
200 ms/ev

5 kHz

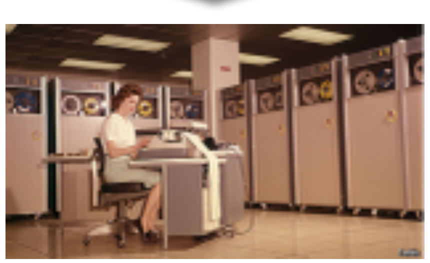

HLT1 inclusive beauty (and charm)

- Single track line with p_T > 1700 MeV and IP > 100 μm
- Little to be gained from looking for other tracks with p_T > 1250 MeV.

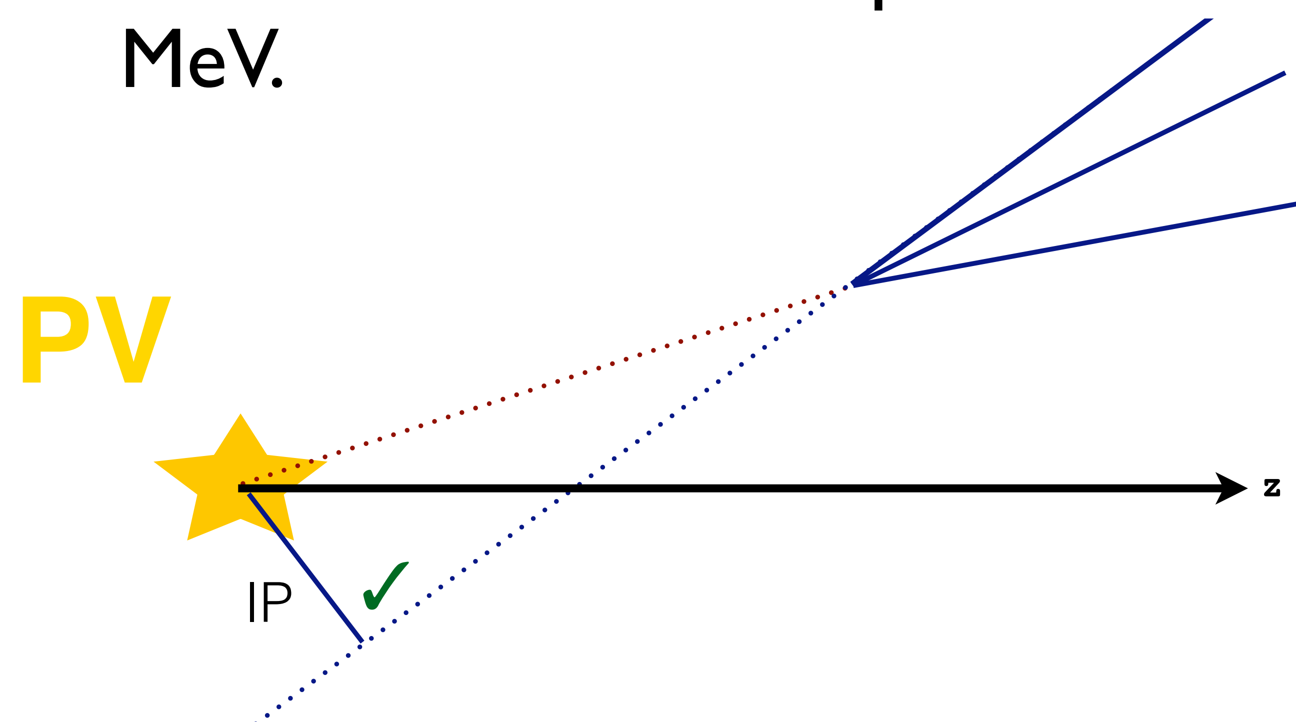

Run-II

1 MHz

Tracking
p_T > 500 MeV

Incl. + excl.

HLT1
35 ms/ev

150 kHz

Tracking
p_T > 70 MeV

Incl. + Excl.

HLT2
500 ms/ev

12.5 kHz

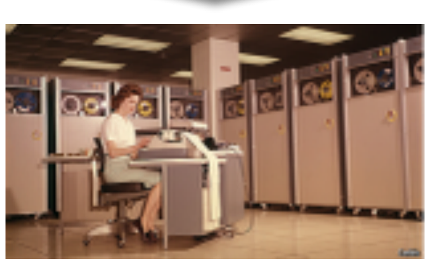

HLT1 inclusive beauty (and charm)

- At 500 MeV, now we start to see more of the beauty daughters
- Start to gain with 2-track triggers

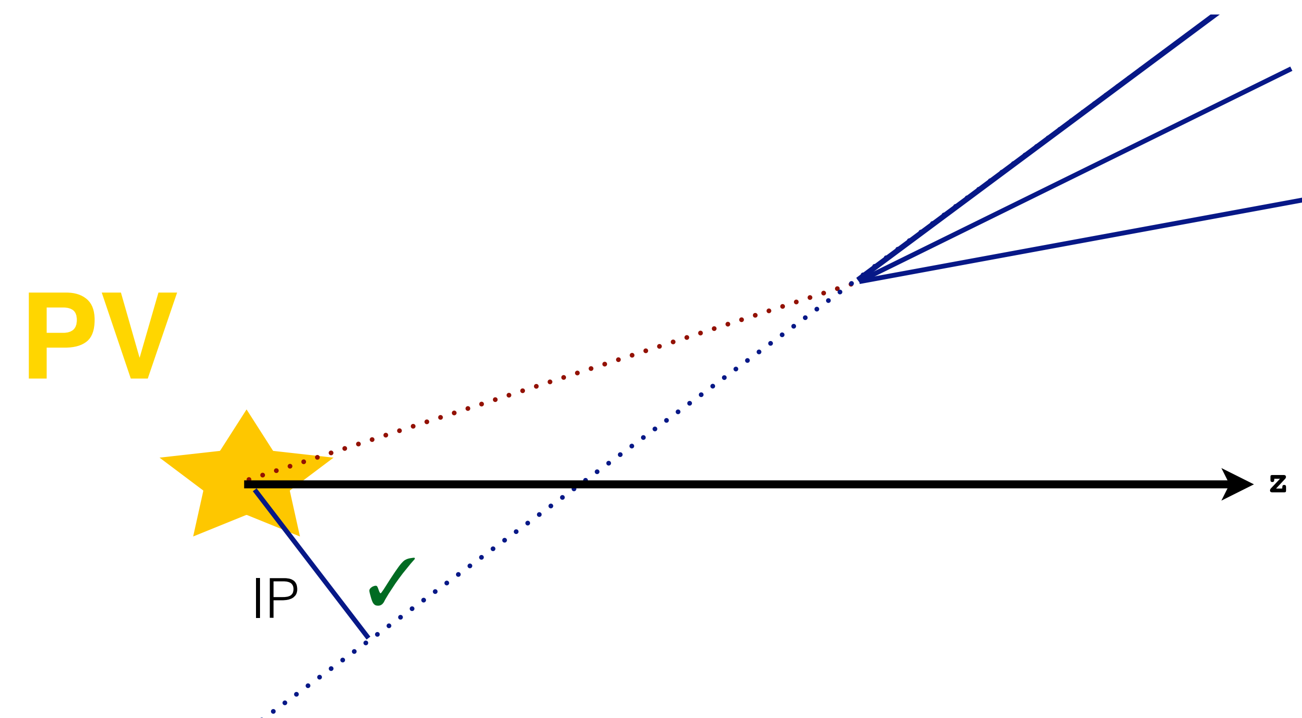

Run-II inclusive 1- and 2-track triggers

MLHEP

1 Track:

- track pt > 500 MeV; IP chi2 > 4; track chi2/ndof < 3;
- BDT using pt, IP chi2;
- Adding cut-based TrackMuon requiring pt > 600 MeV, IP chi2 > 6 adds about 20 kHz to the total.

2-body SV:

- tracks with pt > 500 MeV; IP chi2 > 4, track chi2/ndof < 2.5;
- SV pt > 2 GeV; vertex chi2 < 10; 1 < MCOR < 10 GeV; 2 < eta < 5 (PV to SV).
- BDT using pt, MCOR, vertex chi2, min track pt, FD chi2, N(tracks with IP chi2 < 16).

The HLT2 Topo lines are very similar, except that we have 3- and 4-body versions too.

Variables

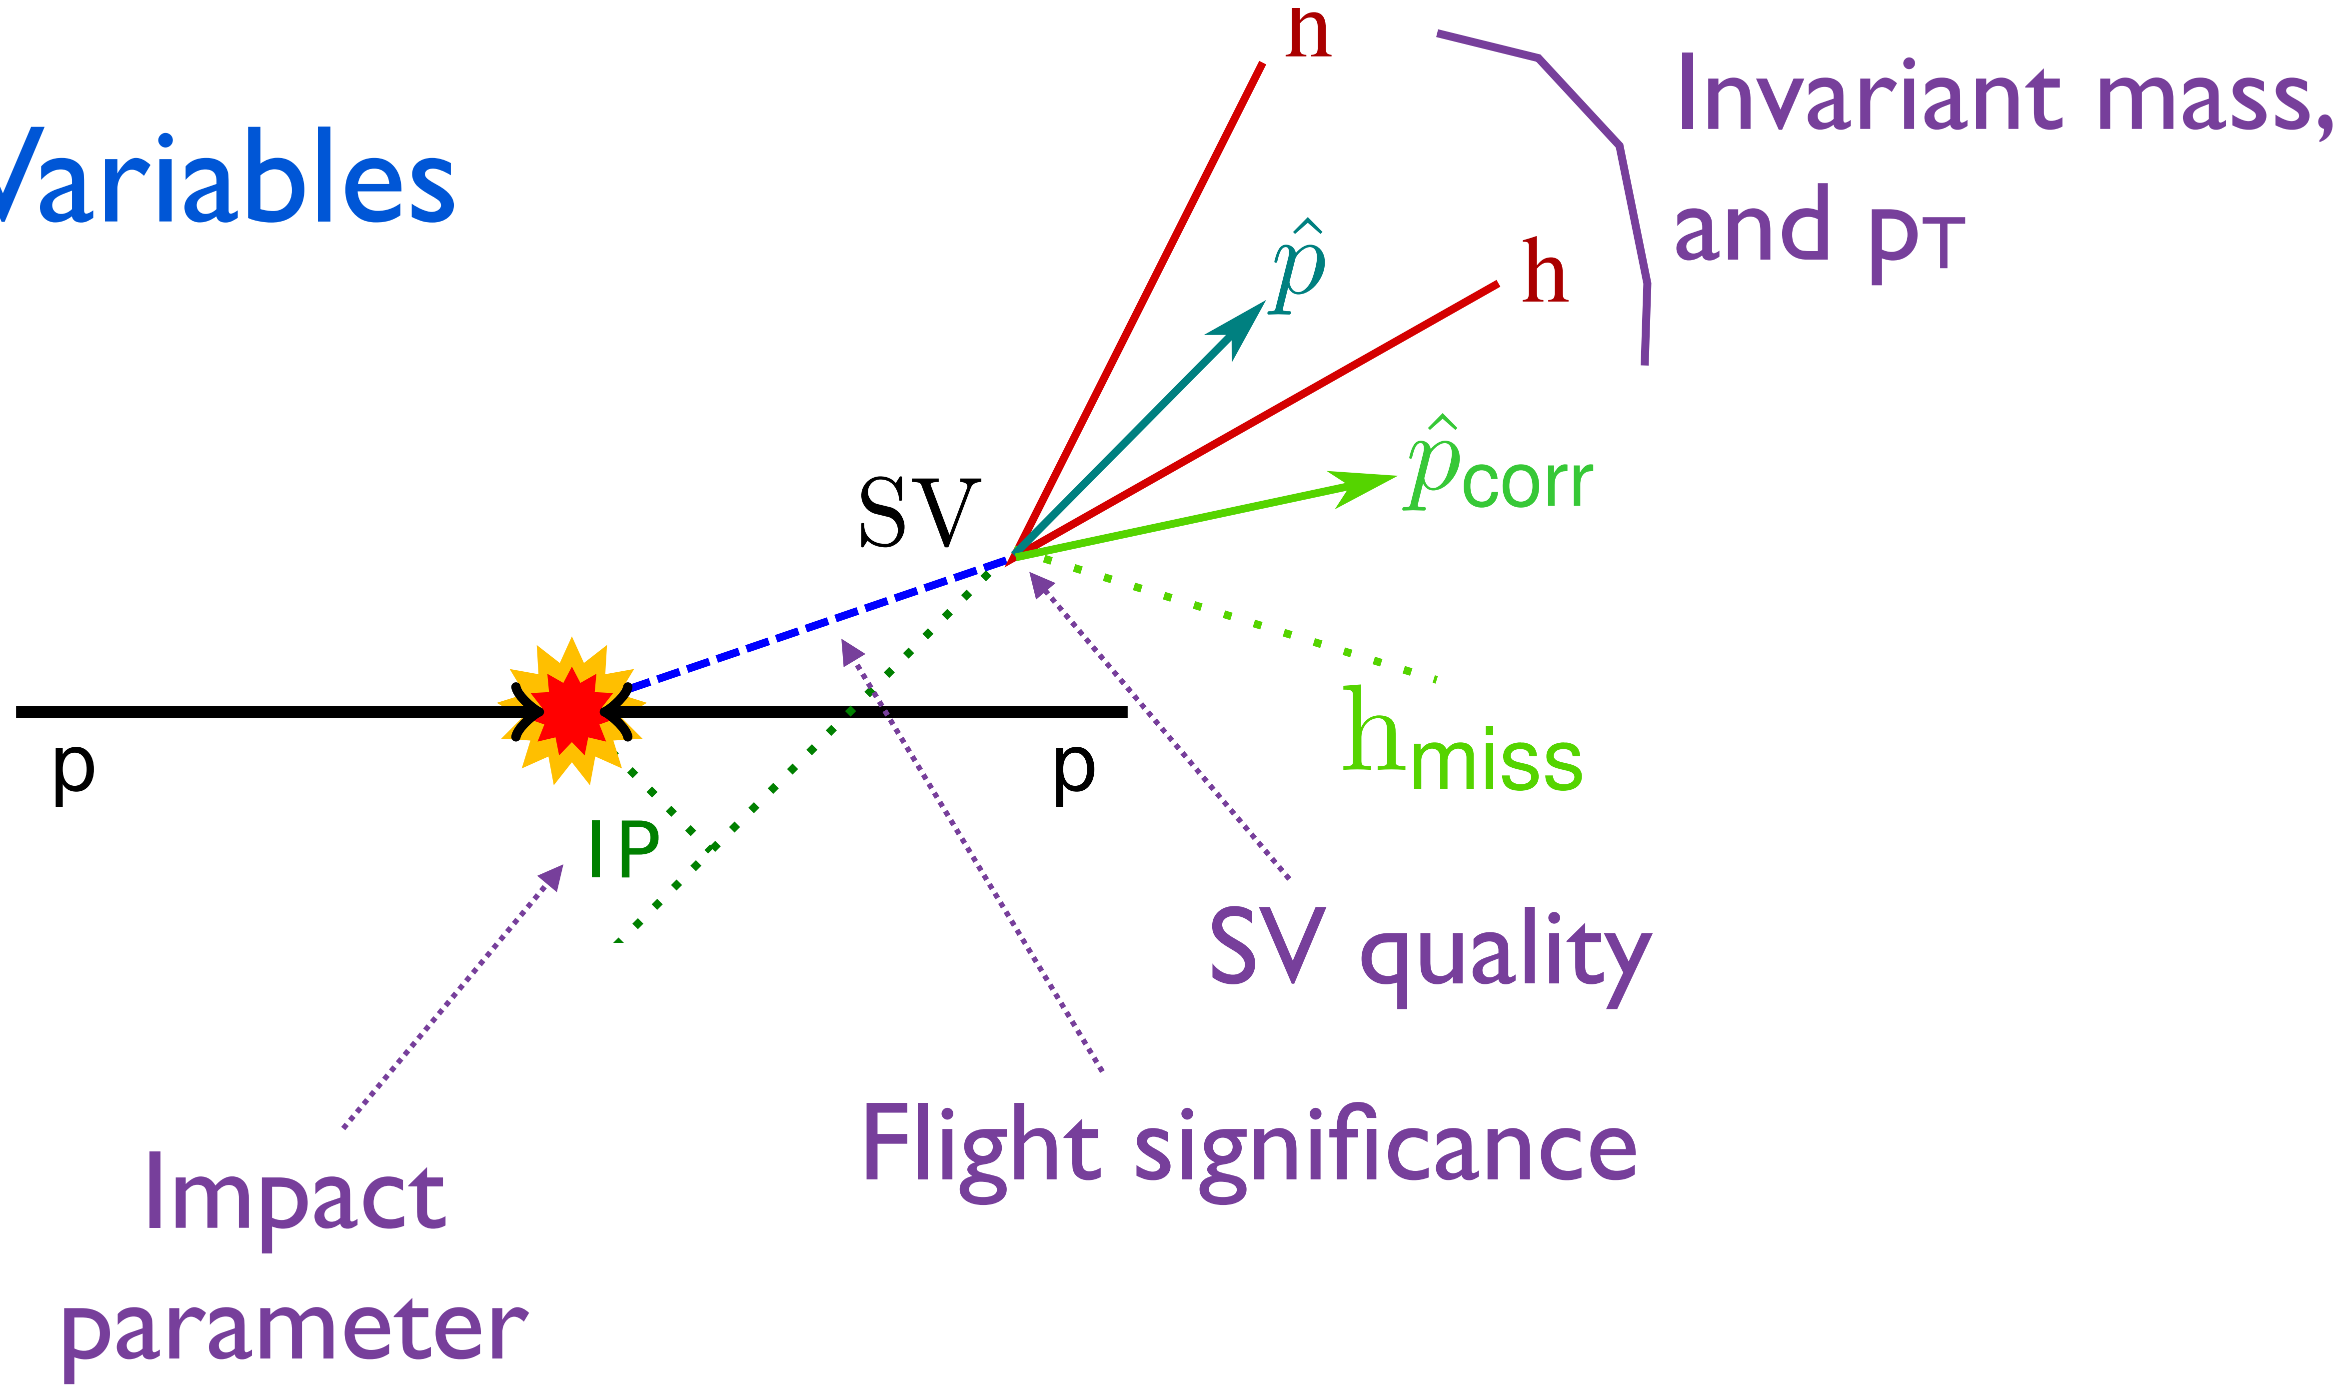
h
Invariant mass,
and pT
$\hat{p}$
h
$\hat{p}$corr
SV
h_{miss}
p
p
IP
SV quality
Flight significance
Impact
parameter

Variables

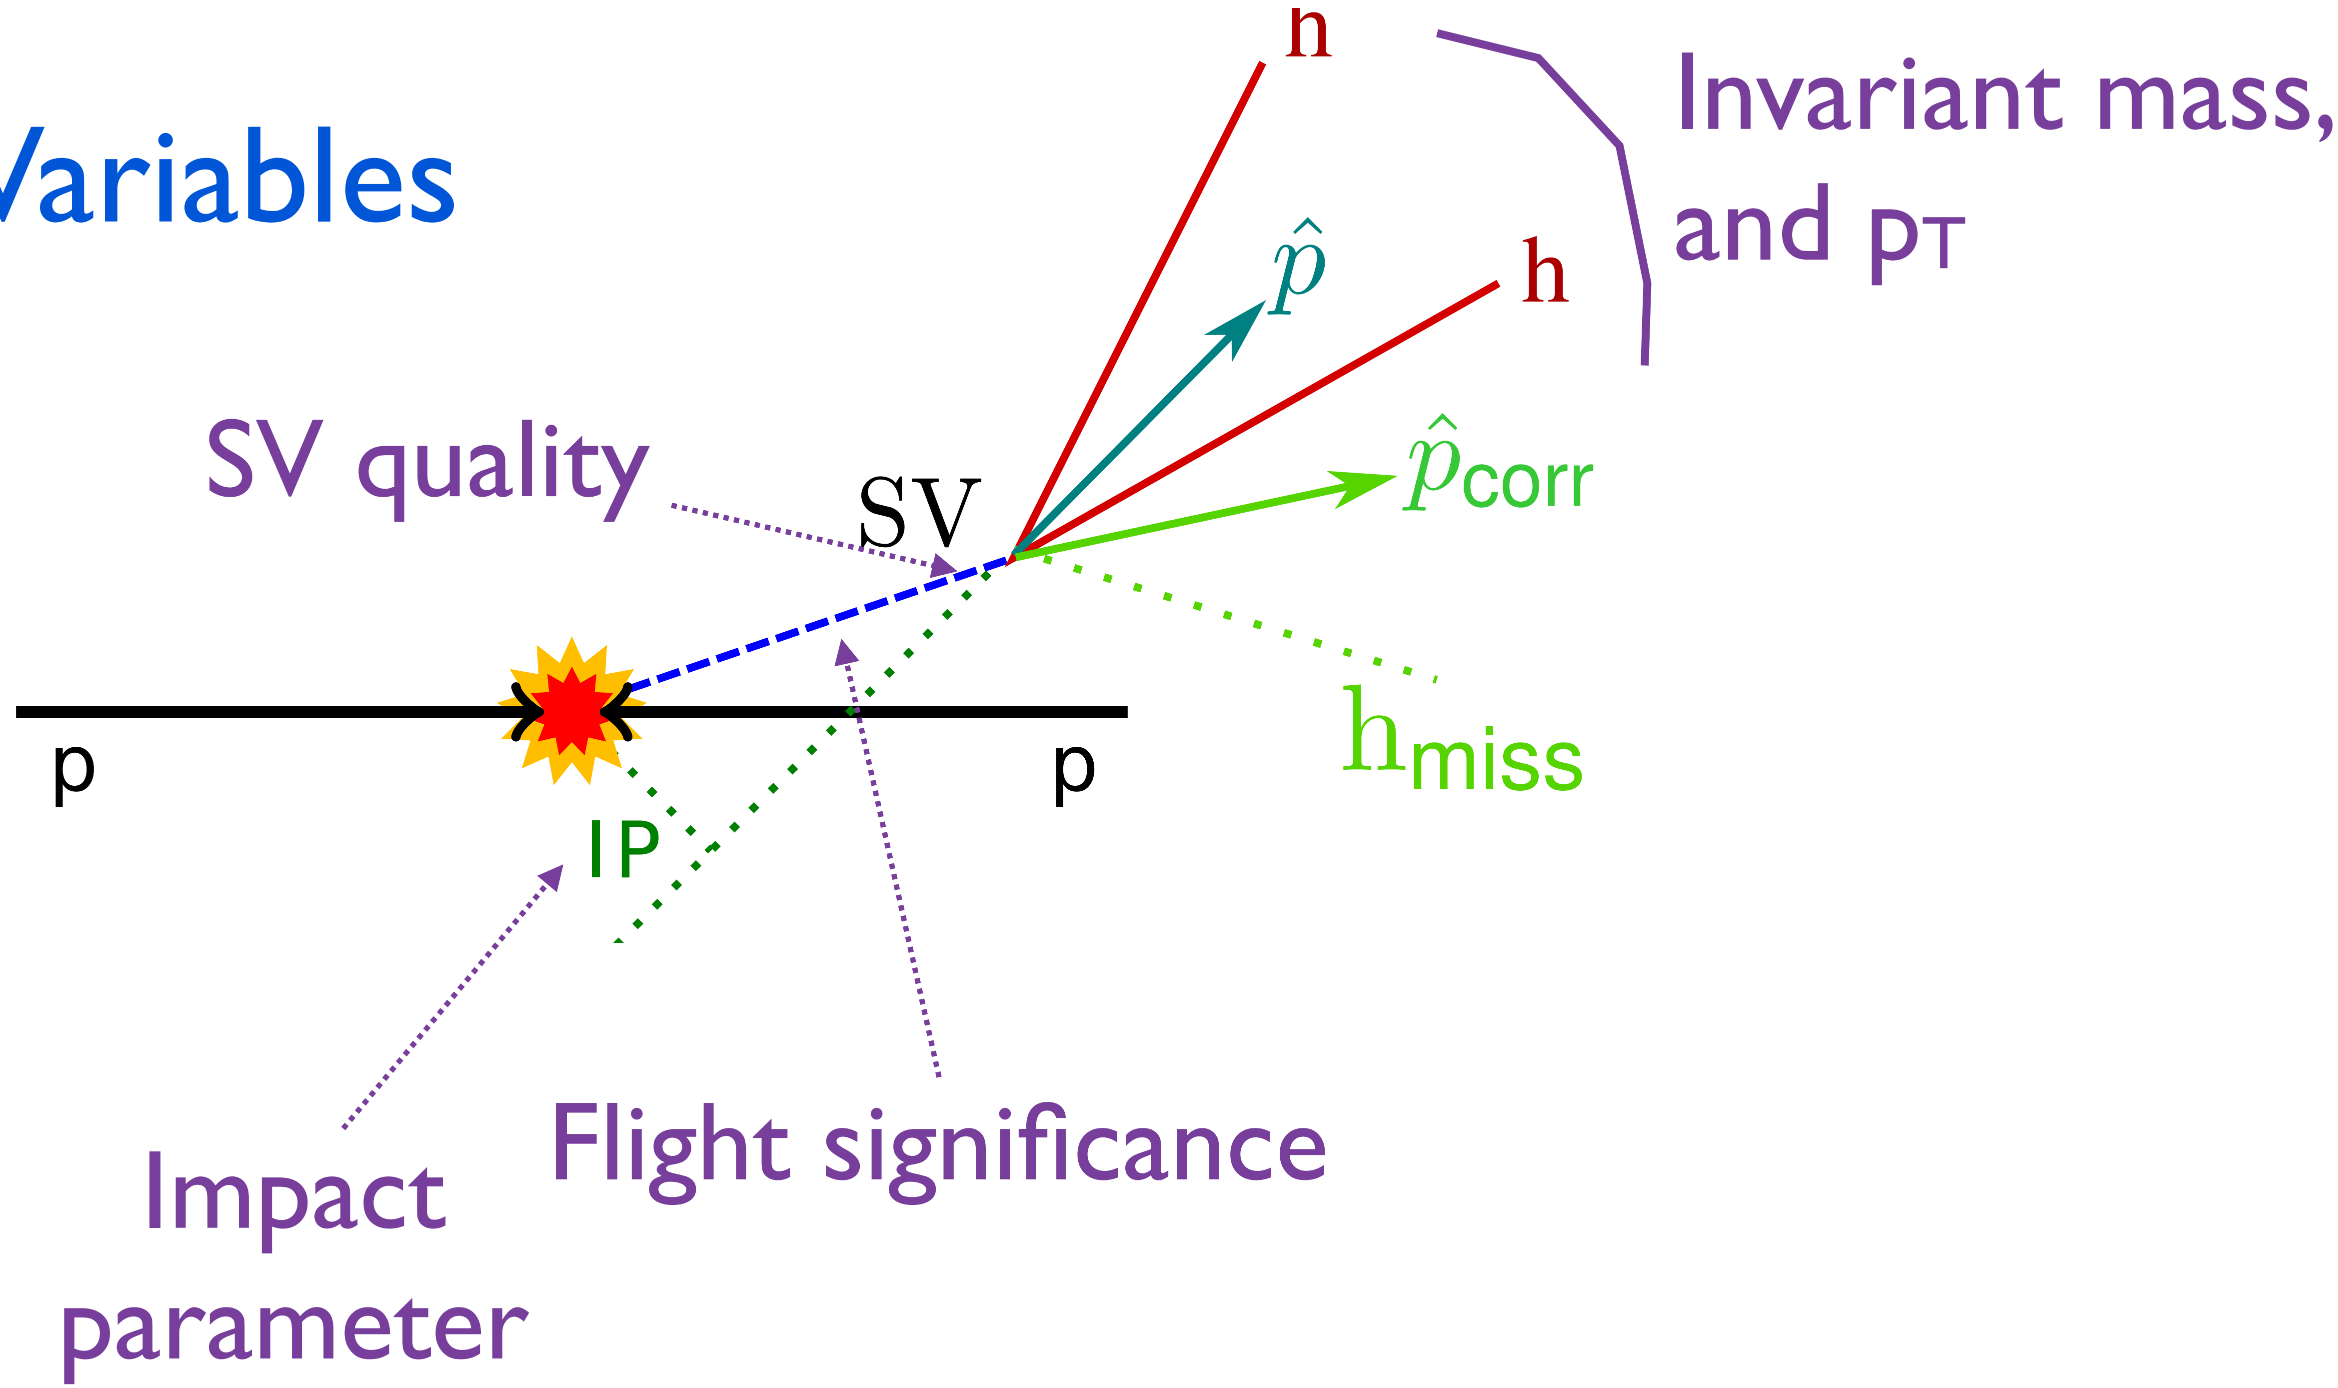
h
Invariant mass,
and p_T
$\hat{p}$
h
SV quality
SV
$\hat{p}_{corr}$
p
IP
p
h_{miss}
Flight significance
Impact
parameter

Variables

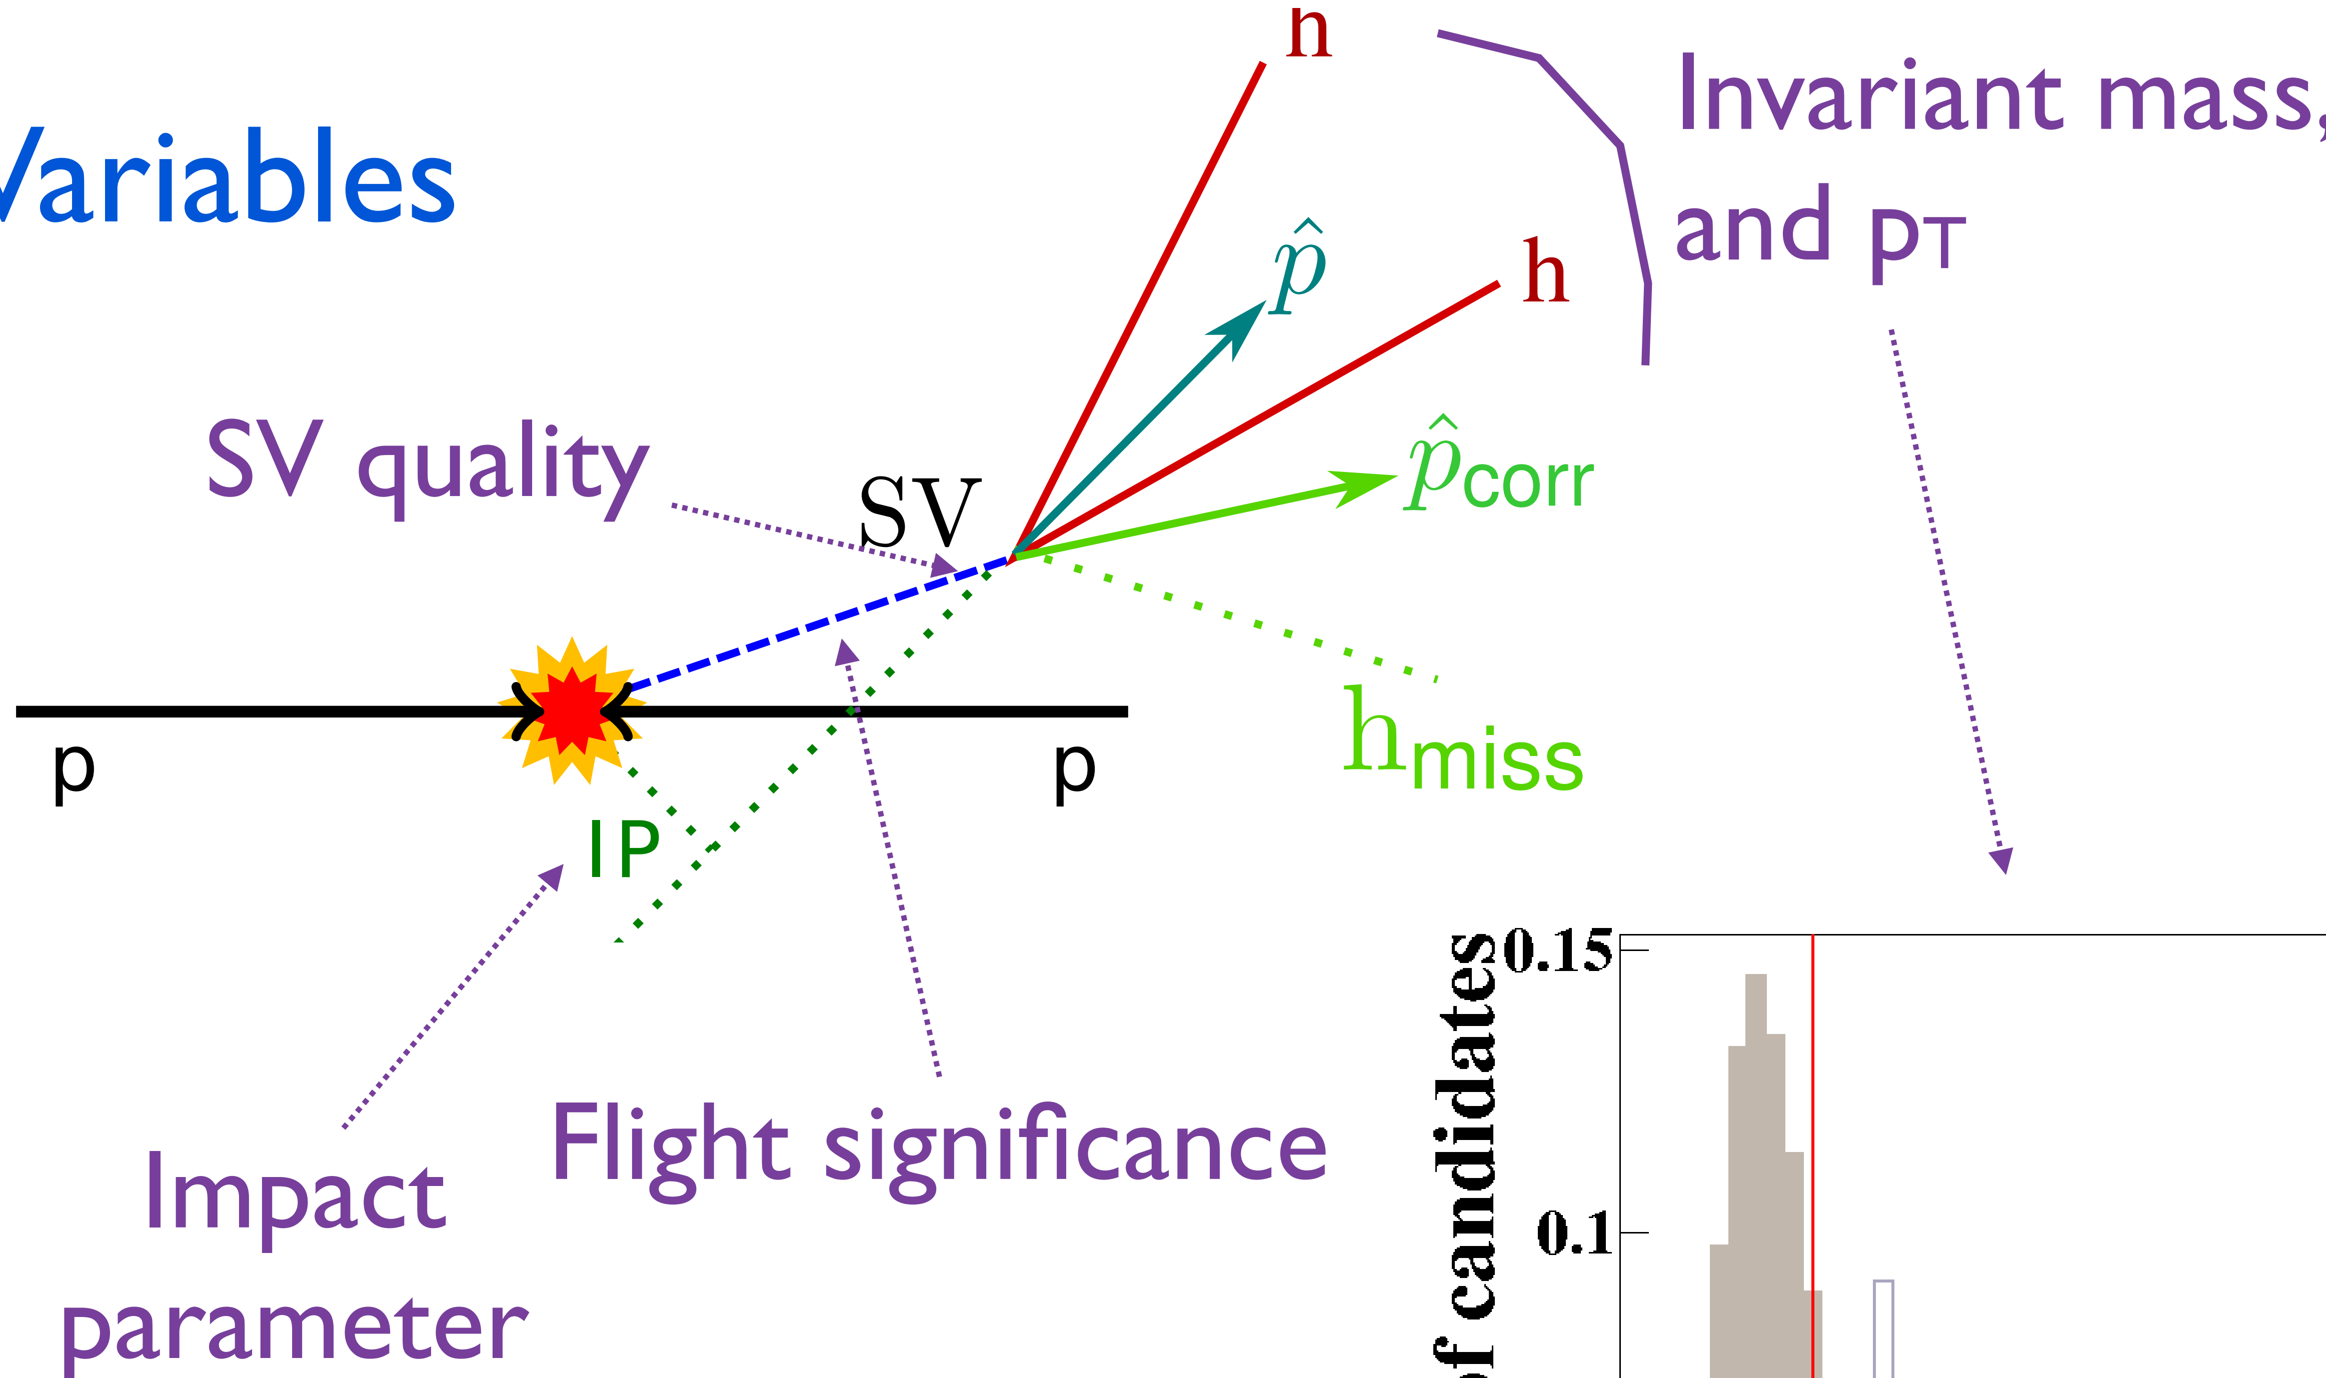
h
$\hat{p}$
h
$\hat{p}_{corr}$
h_{miss}
Invariant mass, and pT
SV quality
SV
p
p
IP
Impact parameter
Flight significance

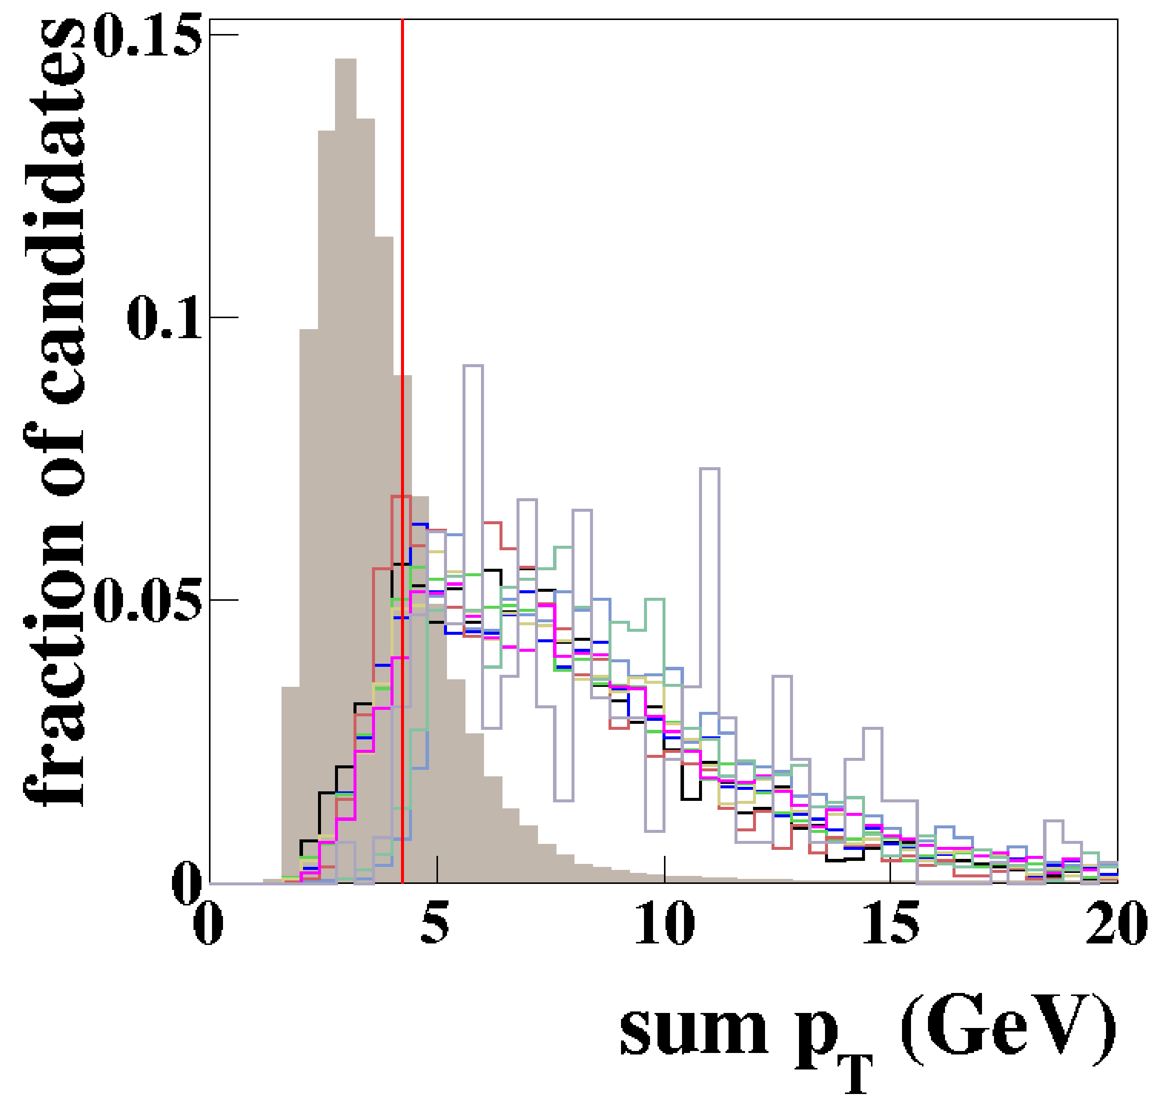
fraction of candidates
0.15
0.1
0.05
0
0
5
10
15
20
sum p_T (GeV)

Variables

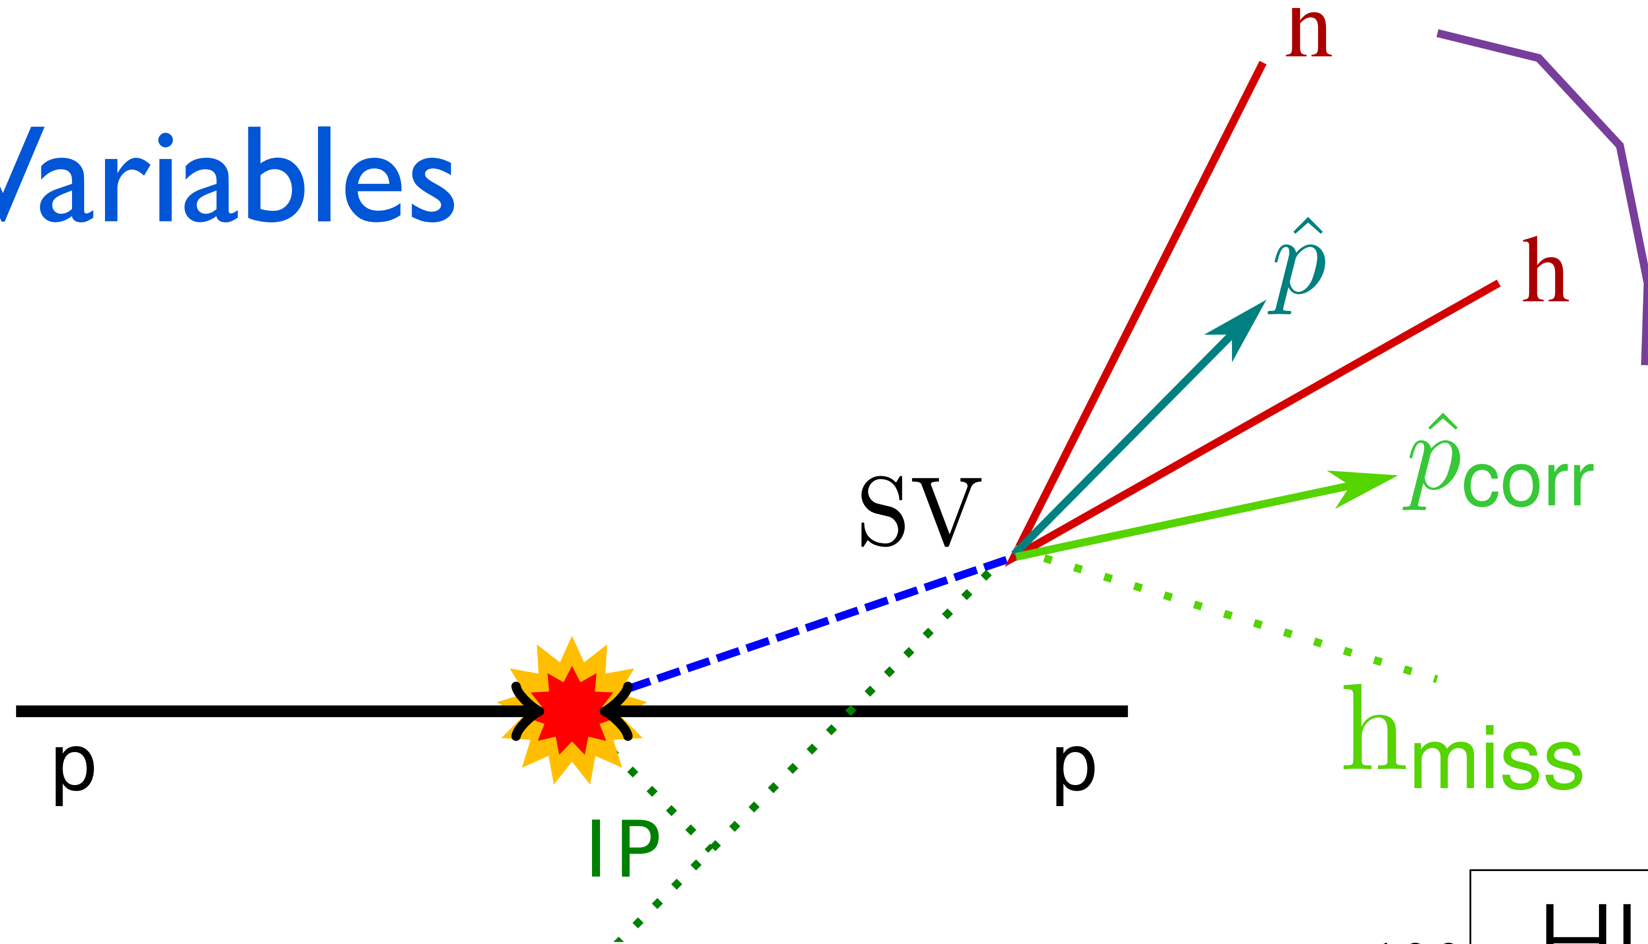

The "corrected mass"

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

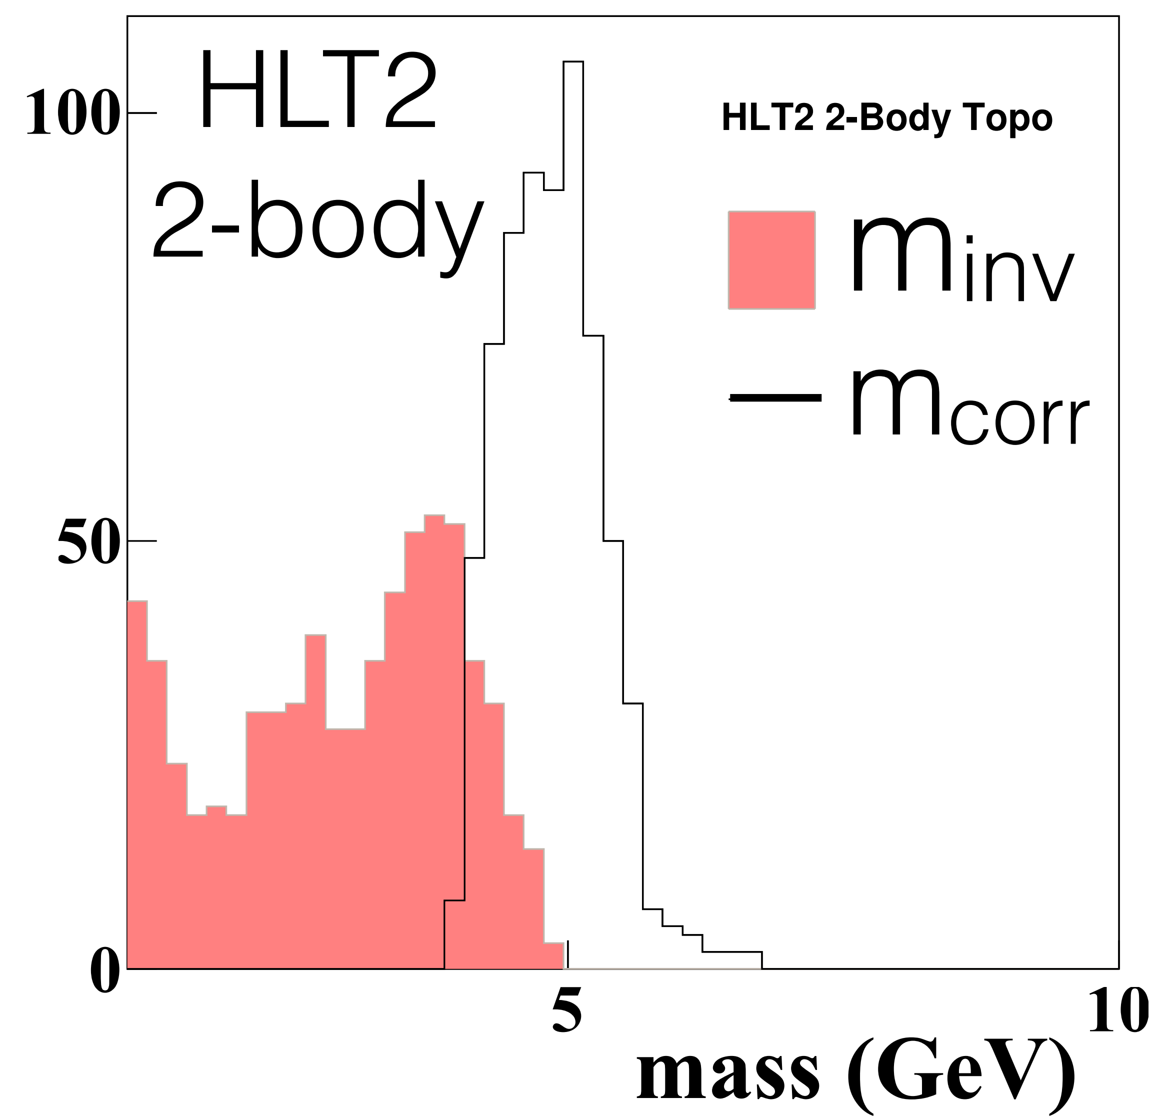

2-, 3-, and 4-body versions

Same 3-body B decay as seen by the 2-, 3- and 4-body Topos

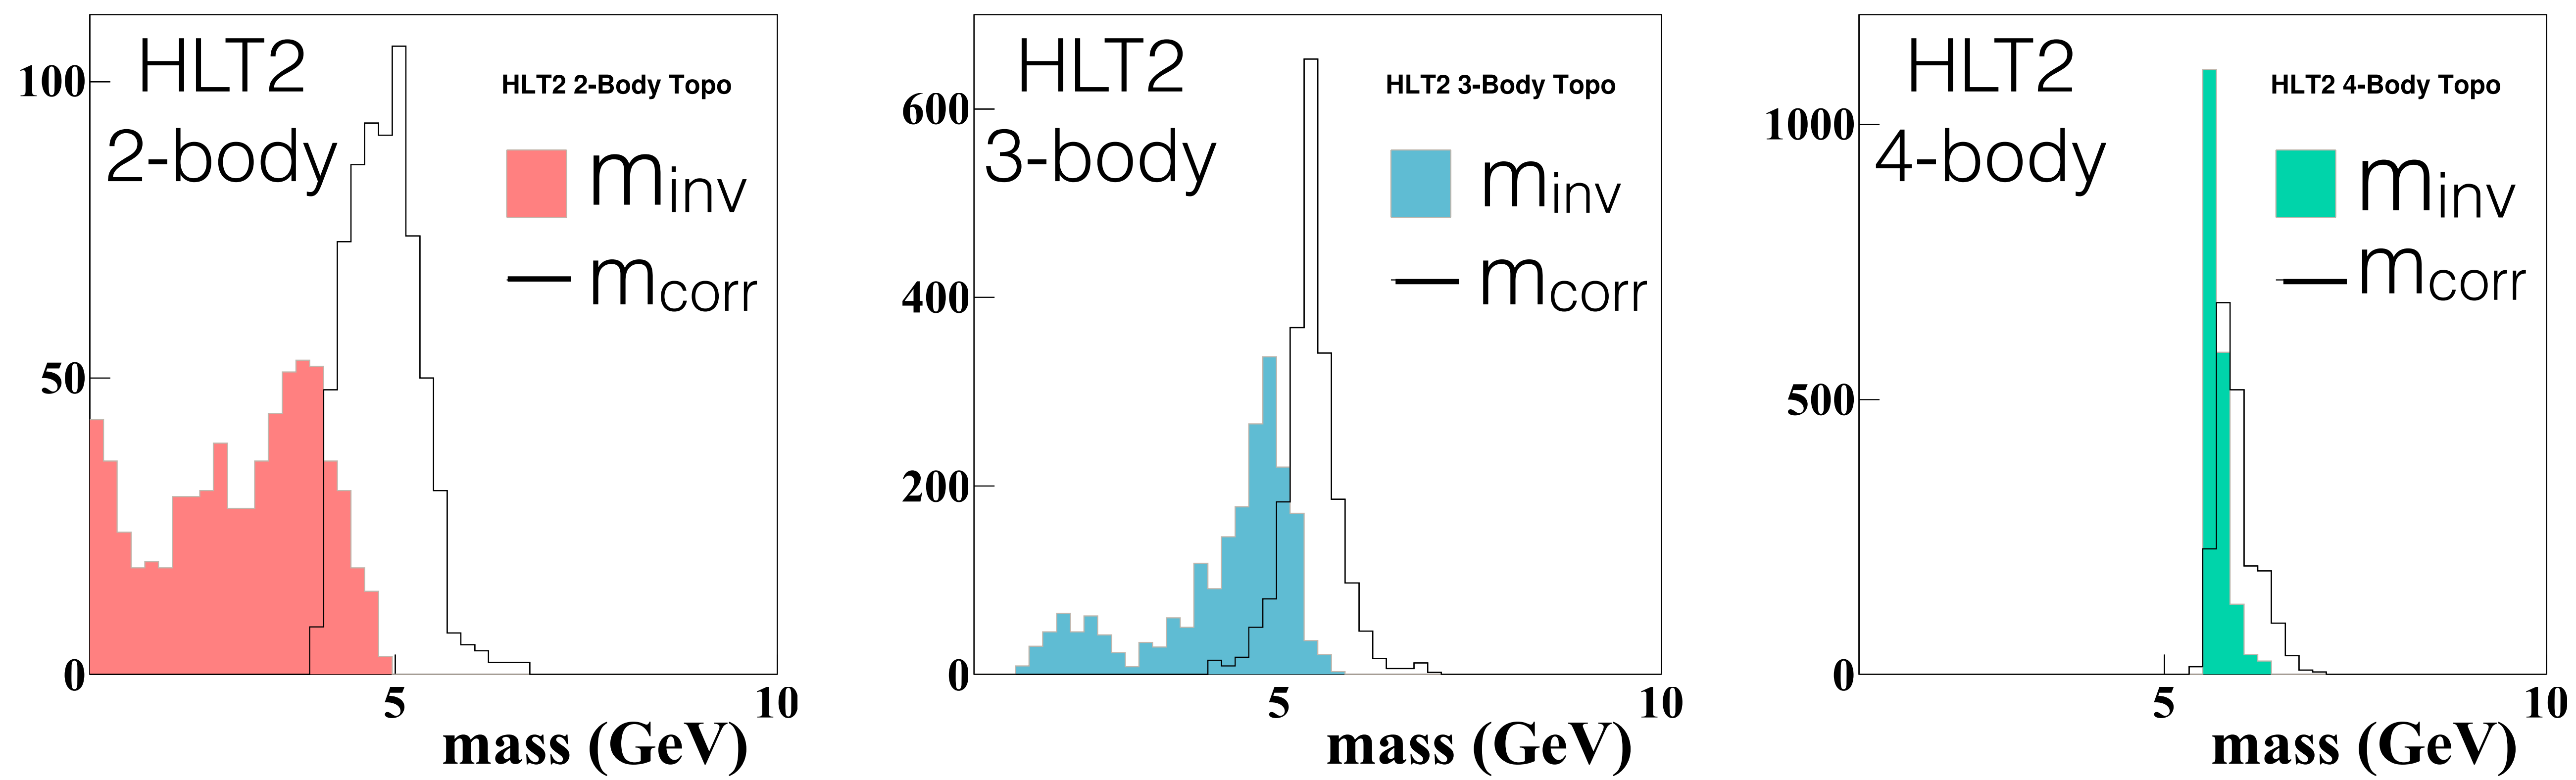

3-body candidates: BKGD = shaded grey, $B \rightarrow$ anything = various colors

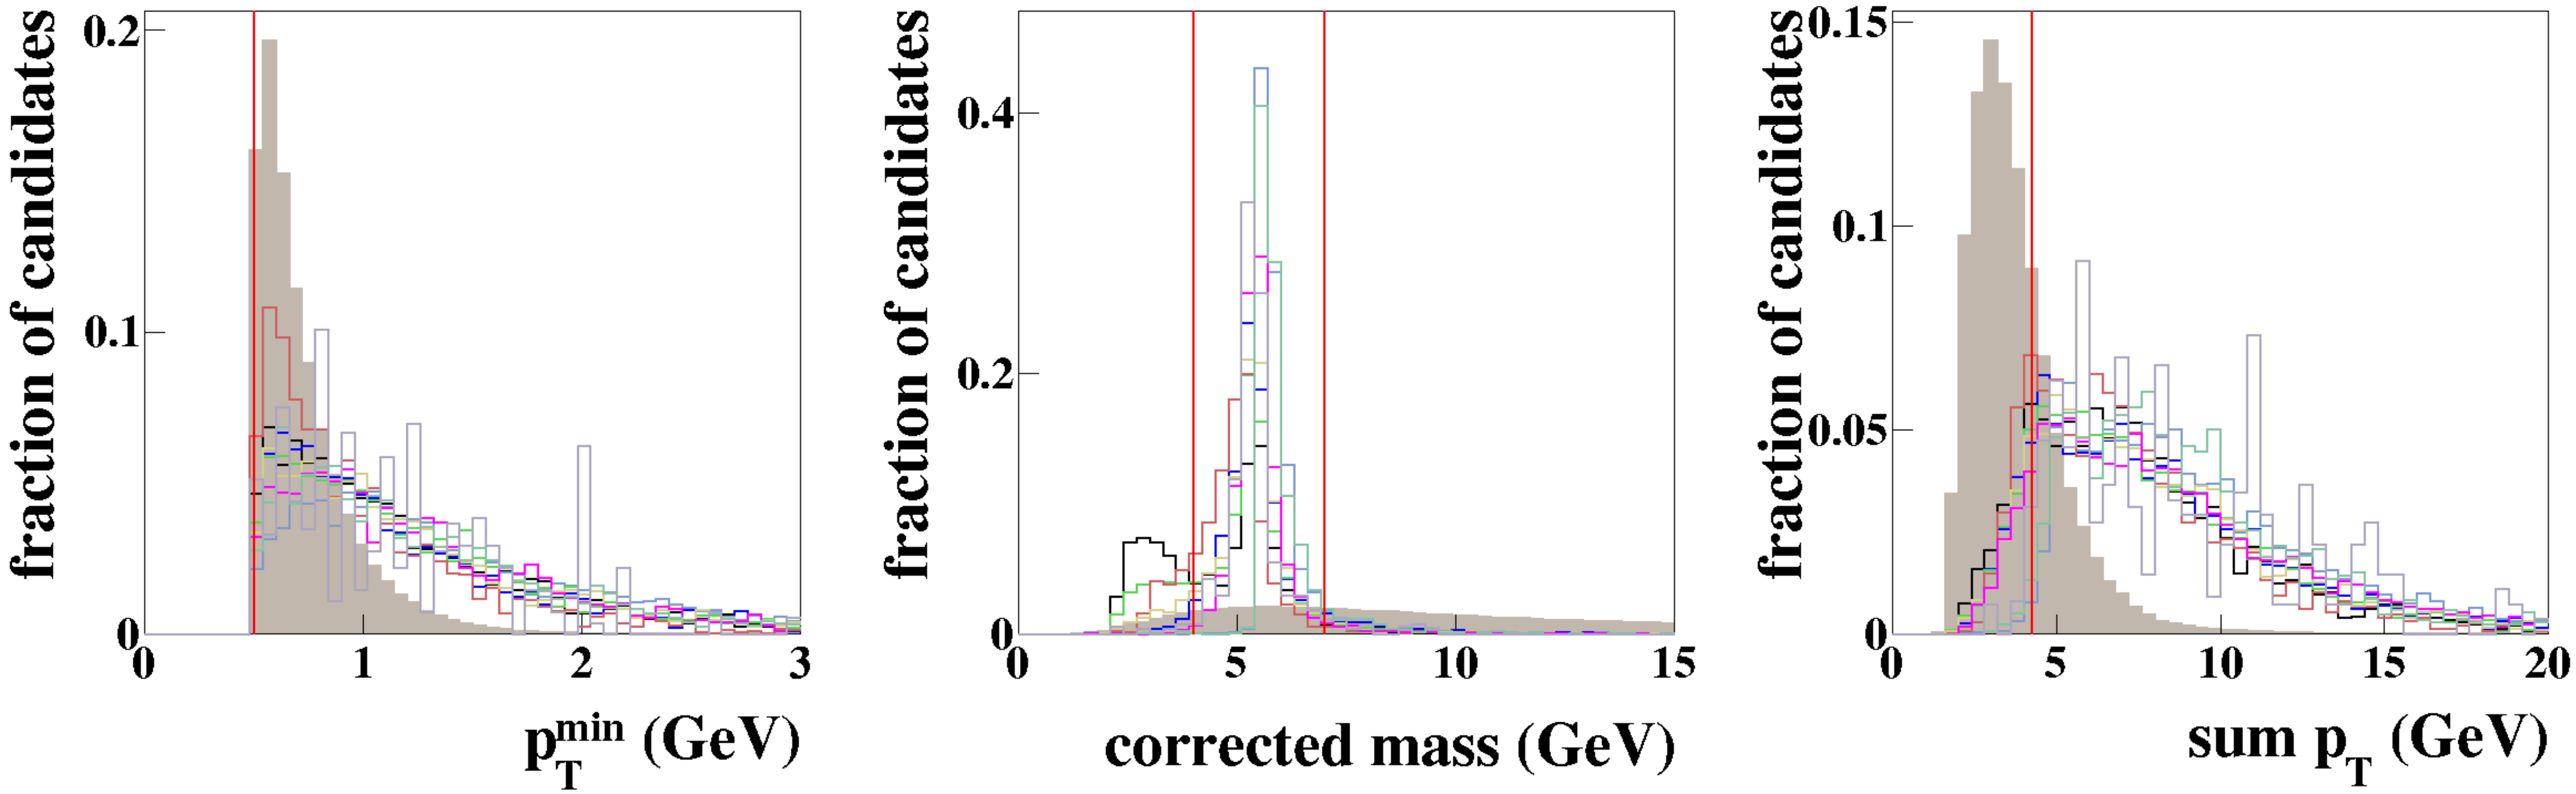

MLHEP

Discussing the selection cuts, pre MVA

The Topo with MLHEP

Signal efficiency

Output rate

Single BDT threshold

Alternative view: MVA in my HLT!

MVA first introduced into the inclusive B triggers in 2011, when we were forced to reduce the background rate by a factor of 3 given the fast ramp up of the luminosity compared to 2010. Some in LHCb had concerns though.

Concerns

- **Efficiency:** needs to be stable and well understood.
- **Generality:** Want to learn generic b decays not just the limited set that we can realistically train on
- **Speed:** Decision in a few ms.

One of the most commonly used classifiers in HEP is the BDT

Concerns

- **Efficiency:** needs to be stable and well understood.
- **Generality:** Want to learn generic b decays not just the limited set that we can realistically train on
- **Speed:** Decision in a few ms.

One of the most commonly used classifiers in HEP is the BDT

Solution: Control what the tree can learn by discretising the variables as we choose.

Bonzai Boosted Decision Tree

Efficient, reliable and fast high-level triggering using a bonsai boosted decision tree.

V. V. Gligorov[1] **and M. Williams**[2,a]

[1] *Organisation Européenne pour la Recherche Nucléaire, Geneva, Switzerland*
[2] *Imperial College London, London SW7 2AZ, UK*
[a] *Current address: Massachusetts Institute of Technology, Cambridge, MA, USA*

ABSTRACT: High-level triggering is a vital component in many modern particle physics experiments. This paper describes a modification to the standard boosted decision tree (BDT) classifier, the so-called *bonsai* BDT, that has the following important properties: it is more efficient than traditional cut-based approaches; it is robust against detector instabilities, and it is very fast. Thus, it is fit-for-purpose for the online running conditions faced by any large-scale data acquisition system.

Discretisation of the variables

Table 4: Allowed split points in the bonsai boosted decision tree.

variable	cuts(2,3,4-body)	allowed splits
$\sum \lvert p_T \rvert$	$> 3,4,4$ GeV	3.5,4,4.5,5,6,7,8,9,10,15,20 (GeV)
mass	< 7 GeV	2.5,4.75 (GeV)
DOCA	< 0.2 mm	0.05,0.1,0.15 (mm)
IPχ^2		20
corrected mass		2,3,4,5,6,7,8,9,10,15 (GeV)
p_T^{min}	> 0.5 GeV	0.6.0.7.0.8,0.9,1,1.25,1.5,1.75,2,2.5,3,4,5,10 (GeV)
FDχ^2	> 100	2,3,4,5,6,7,8,9,10,25,50,100 $\times 100$

Training (run-I)

parent	daughters
$B^{\pm}$	$K\pi\pi$, $D_{[K\pi]}\pi$, $D_{[hhhh]}K$, $D_{[K_S\pi\pi]}K$, $D_{[K\pi\pi]}K\pi$
B^0	$K^*_{[K\pi]}\mu\mu$, $K^*_{[K\pi]}ee$, $D_{[K\pi\pi]}\pi$, $K\pi$, $D_{[K\pi]}K\pi$, $D^*_{[D(K\pi)\pi]}\mu\nu$, $D_{[K\pi\pi]}K\pi\pi$
B_s	$D_{s[KK\pi]}\pi$, $D_{s[KK\pi]}K\pi\pi$, $K^*_{[K\pi]}K^*_{[K\pi]}$
Λ_b	$\Lambda_{c[pK\pi]}\pi$, $\Lambda_{c[pK\pi]}K\pi\pi$

Performance

Stage	Rate	Beauty fraction
L0	1 MHz	1%
Hlt1	70 kHz	9%
Hlt2Topo	2 kHz	100%

3 lines (2-,3-,4- body) which give 2 kHz of ~100% beauty

(plus looser muon and electron versions)

Performance

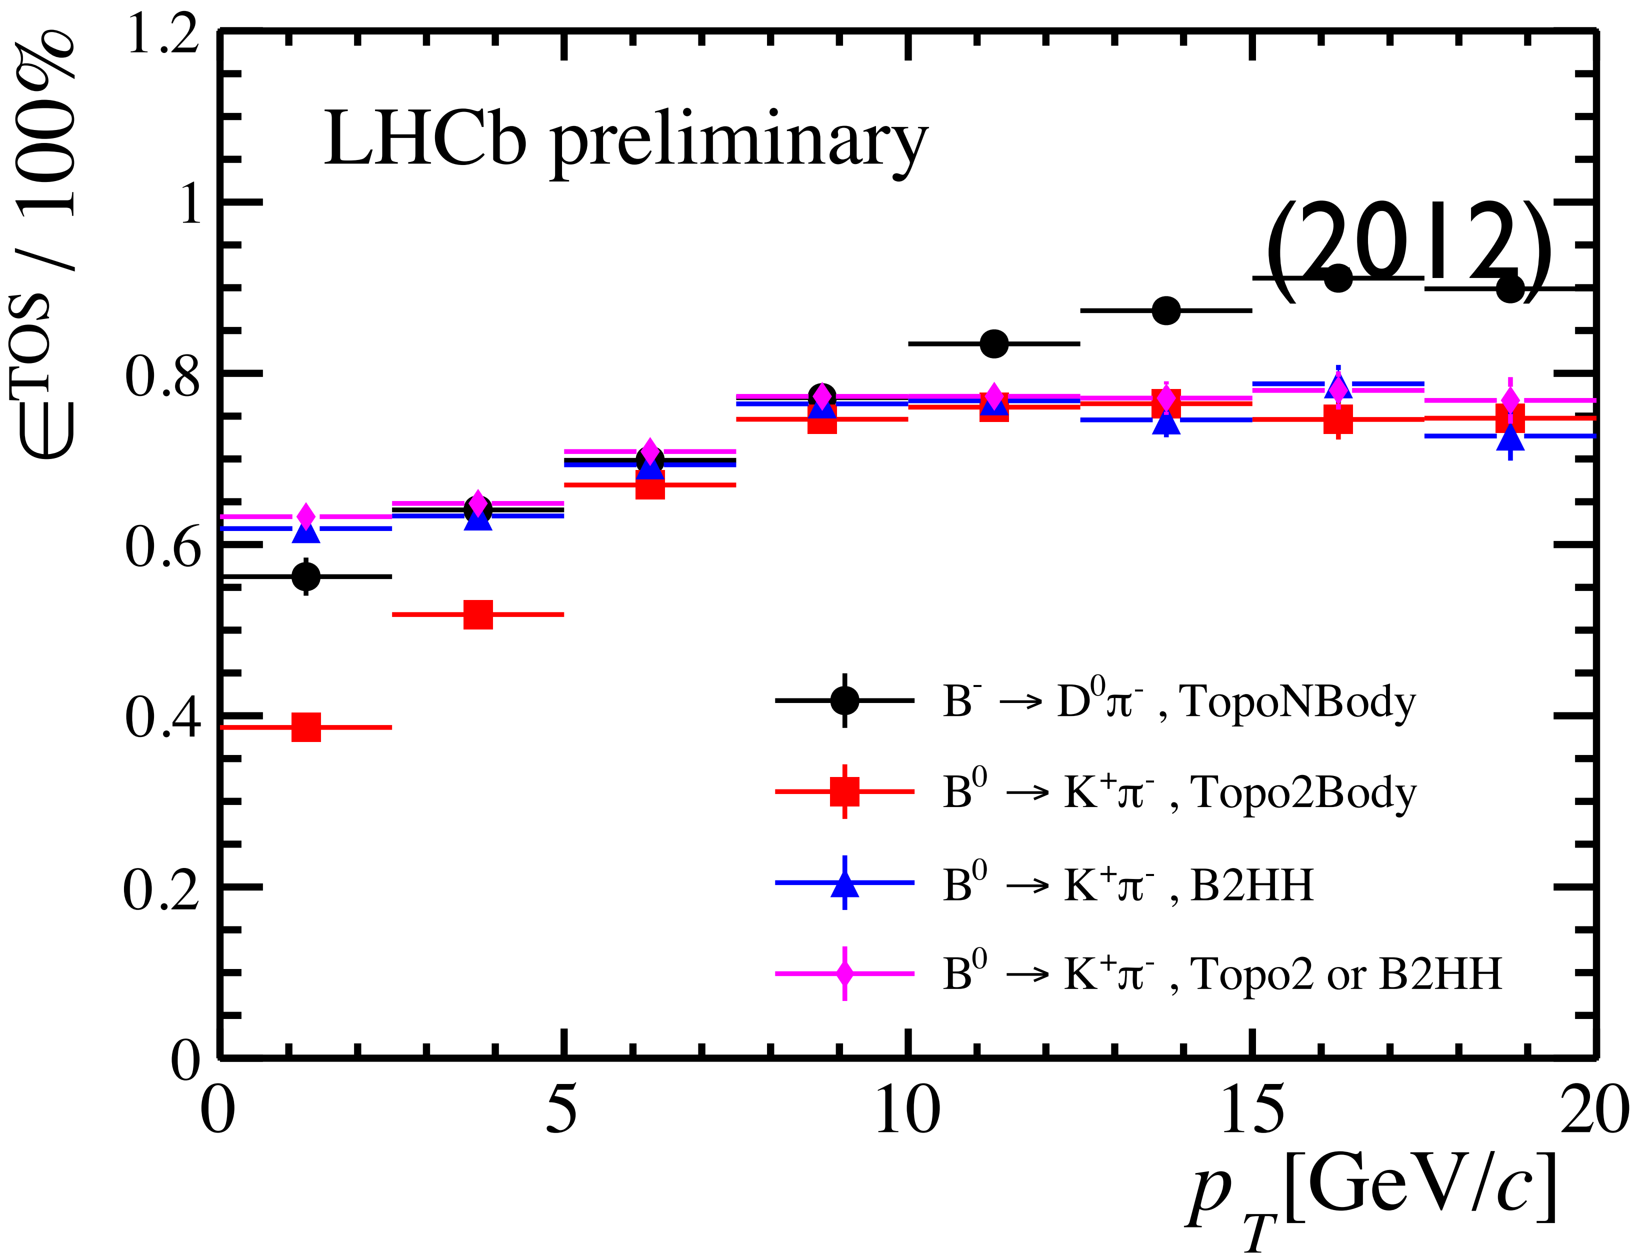

60-80% efficient on various fully hadronic B decays
output rate: 2 kHz

Classifier optimisation for Run-II

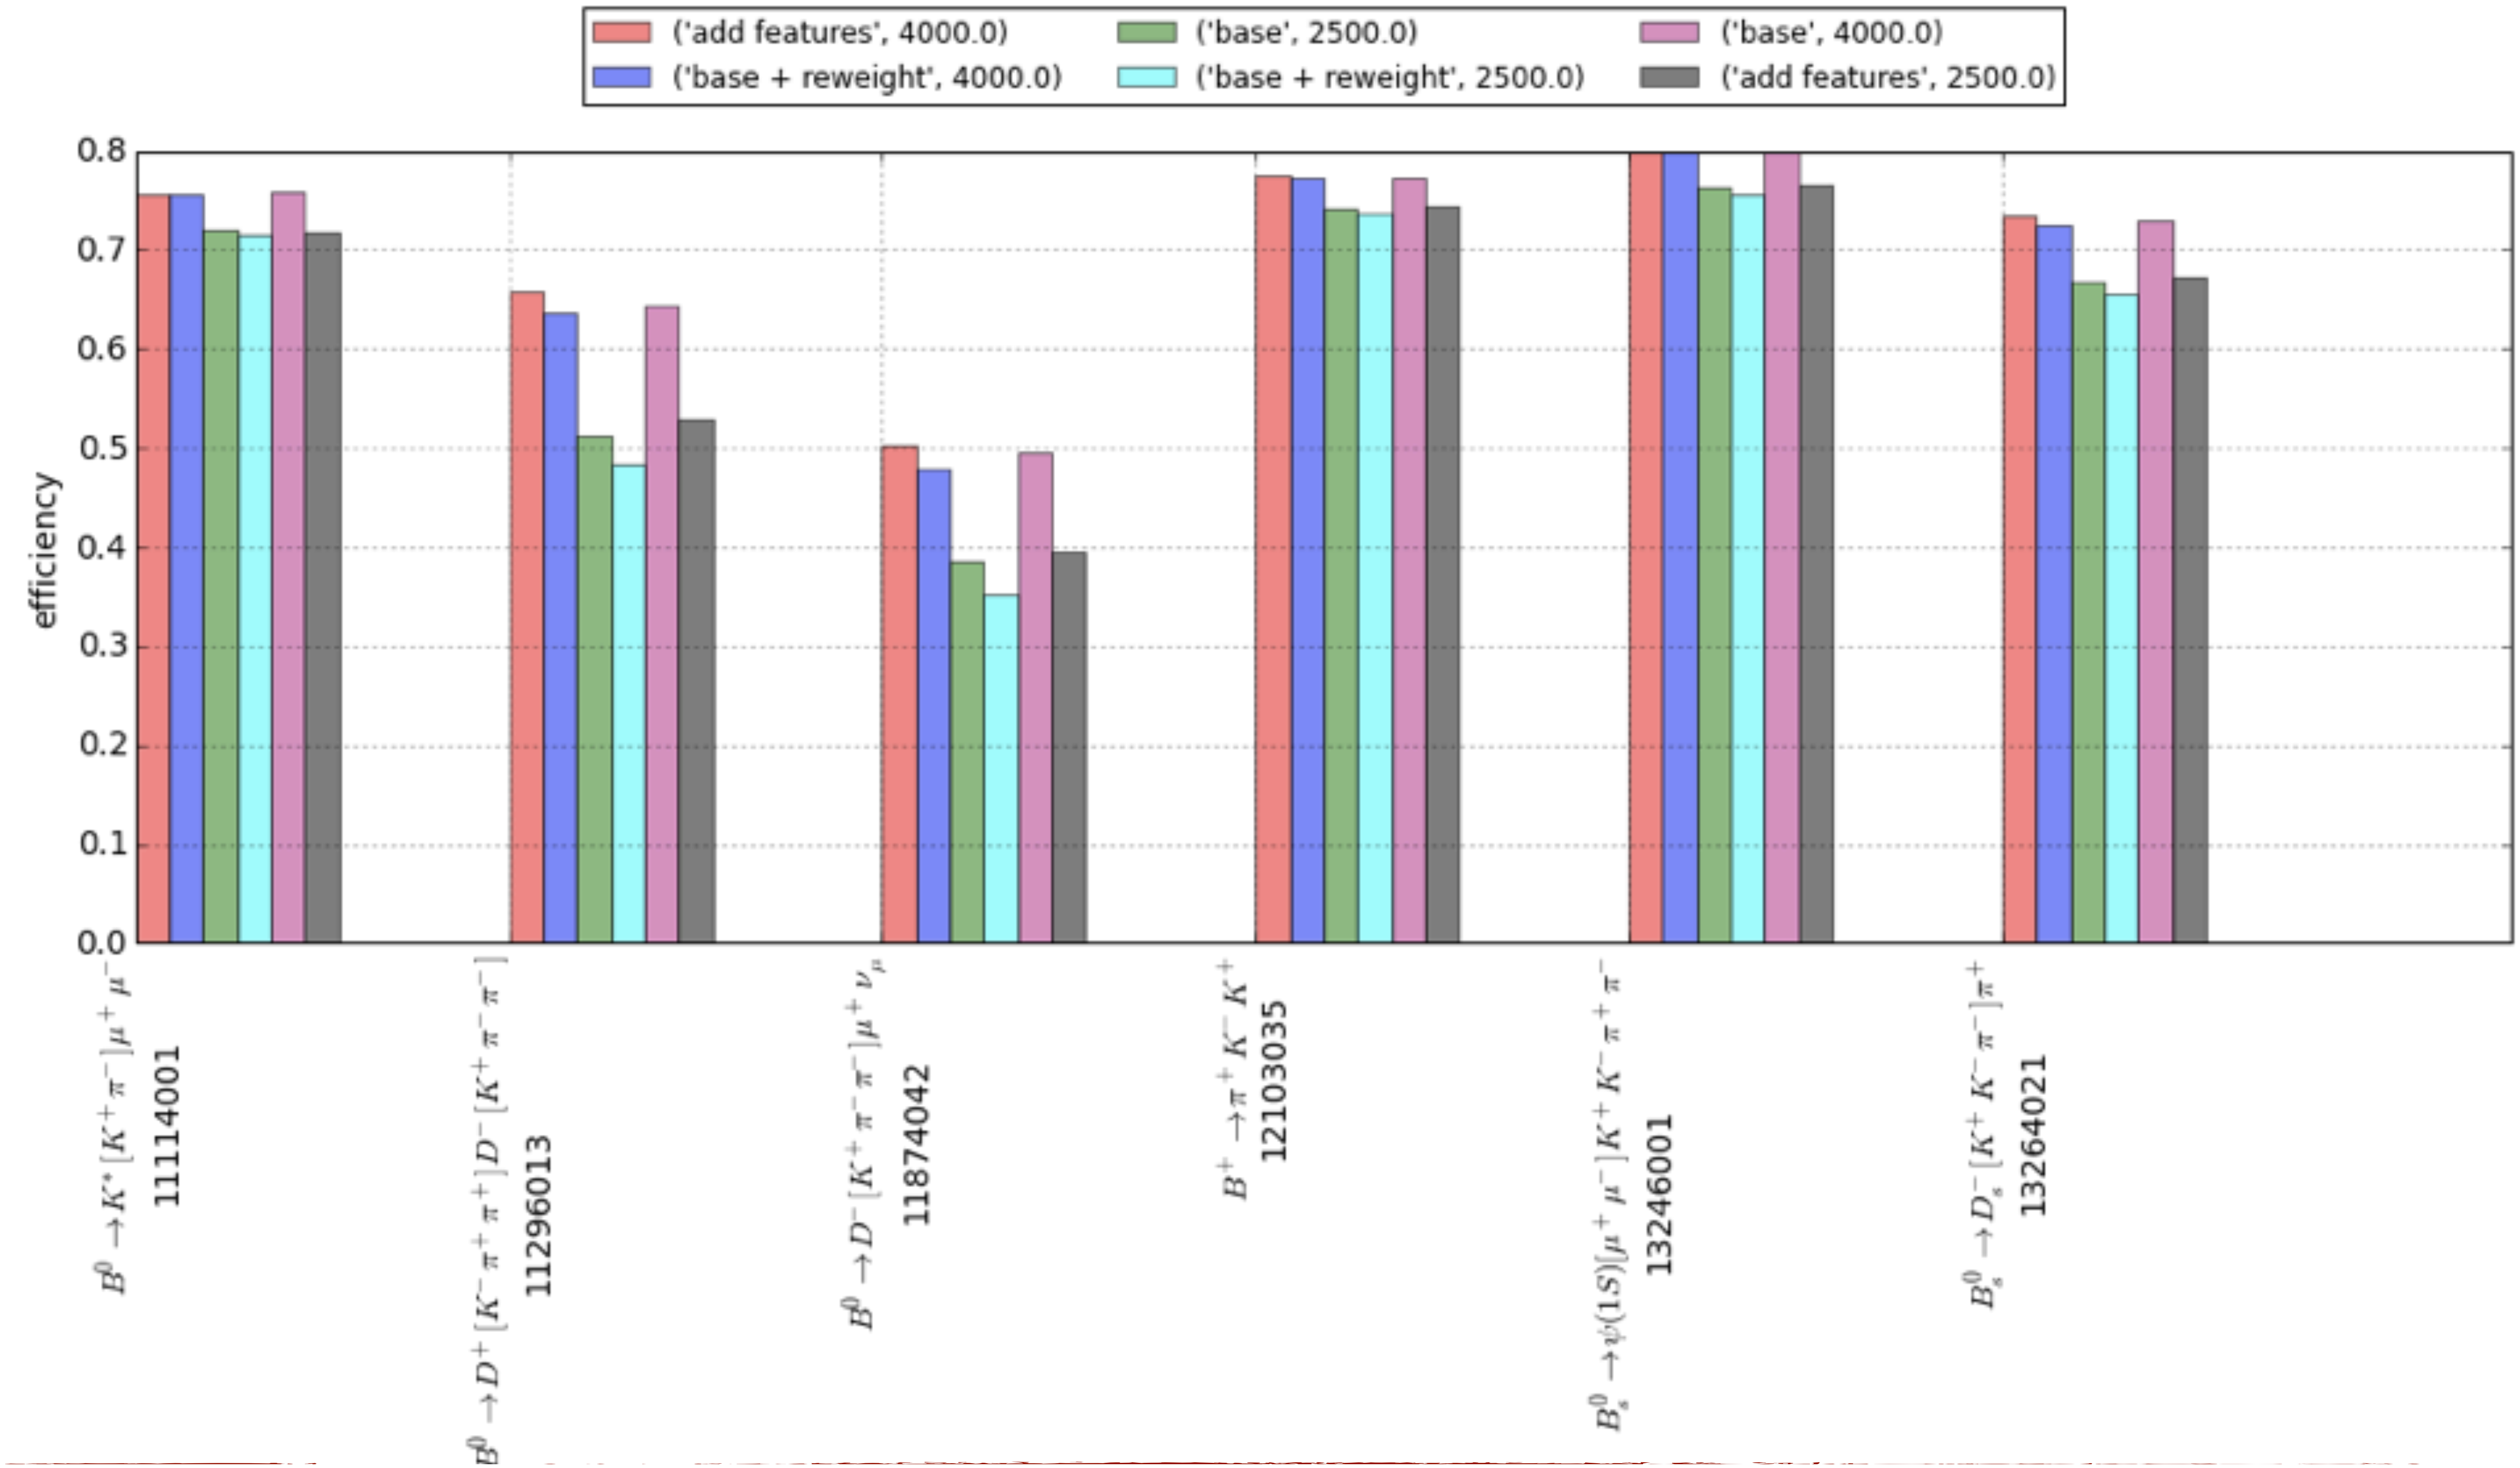

Performance improvements for Run-II

Denominator is now just "reconstructible" and $p_T > 2$ GeV, $\tau > 0.2$ ps

LHCb unofficial

mode	HLT1	HLT2/HLT1 Run 1 ("as was")	HLT2/HLT1 New @ 2.5 kHz	HLT2/HLT1 New @ 4 kHz
$B^0 \to K^*[K^+\pi^-]\mu^+\mu^-$	94	47	77	81
$B^+ \to \pi^+K^-K^+$	95	49	78	81
$B_s^0 \to D_s^-[K^+K^-\pi^-]\mu^+\nu_\mu$	95	36	41	53
$B_s^0 \to \psi(1S)[\mu^+\mu^-]K^+K^-\pi^+\pi^-$	97	48	78	82
$B_s^0 \to D_s^-[K^+K^-\pi^-]\pi^+$	96	48	70	73
$B^0 \to D^+[K^-\pi^+\pi^+]D^-[K^+\pi^-\pi^-]$	98	35	49	65

If we want to be picky though….

Want to select b decays inclusively without biasing the kinematic and decay time distributions.

Want to measure lifetimes of b hadrons, e.g., due to their indirect sensitivity to CPV parameters.

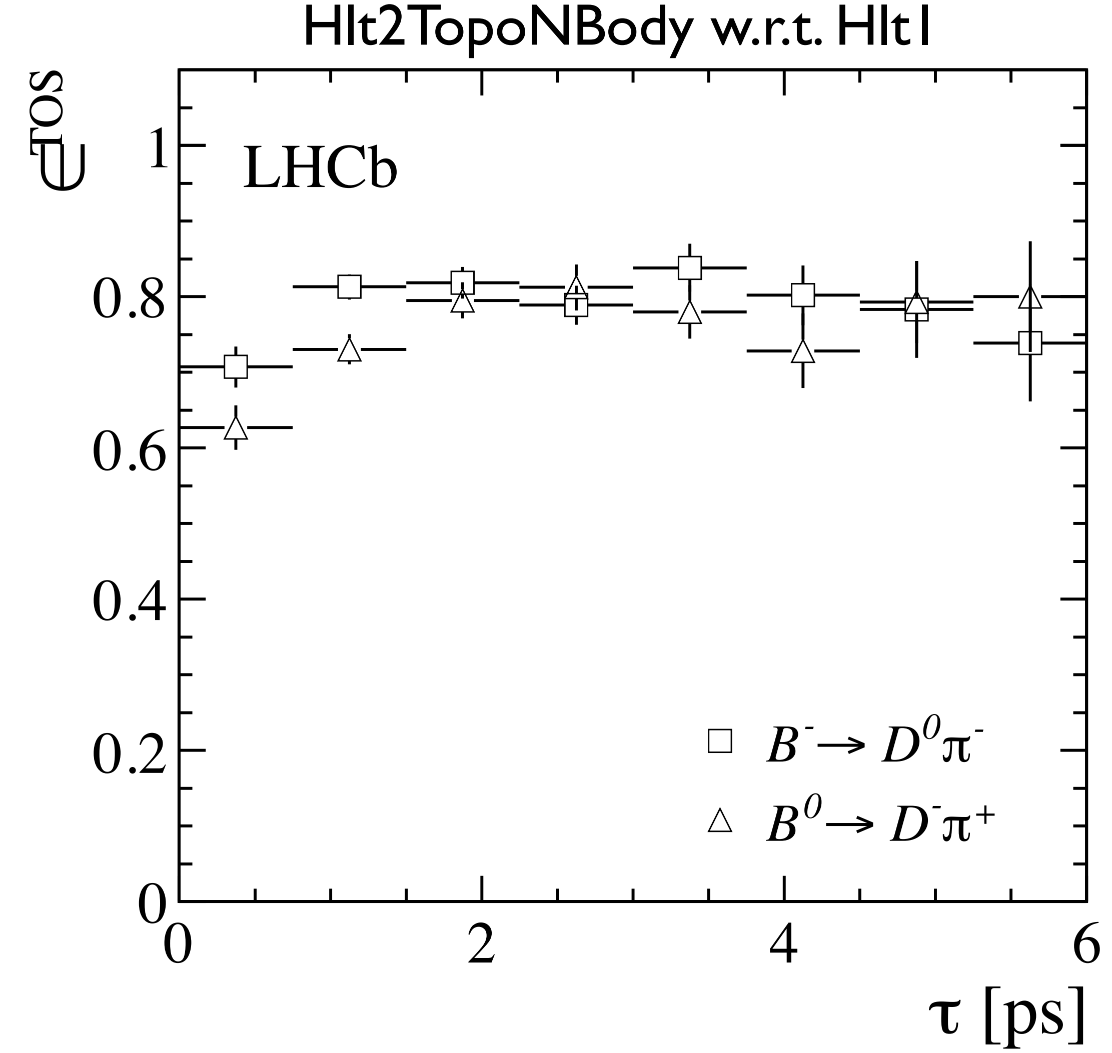

The Dalitz plot

- Useful way to visualise a decay 3-body decay

E.g., $D_s \rightarrow K^+K^-\pi^-$

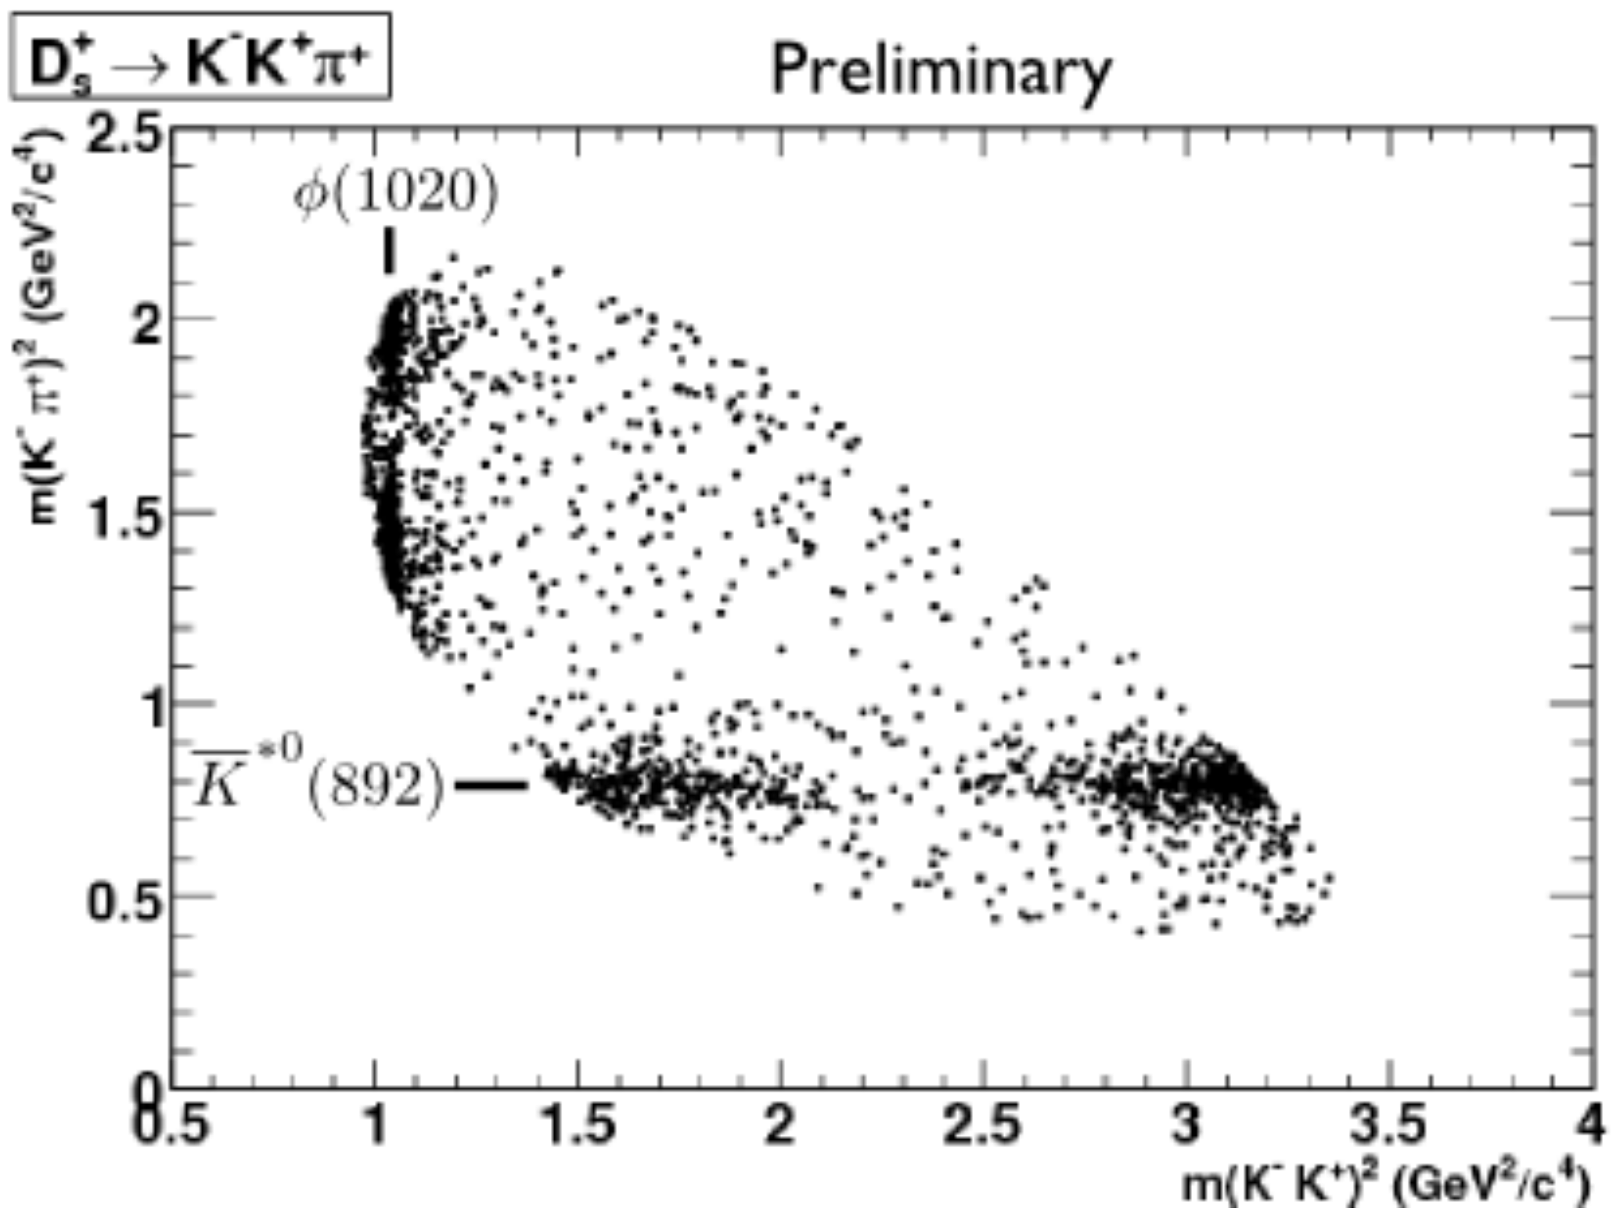

The problem

Toy Dalitz example X → abc

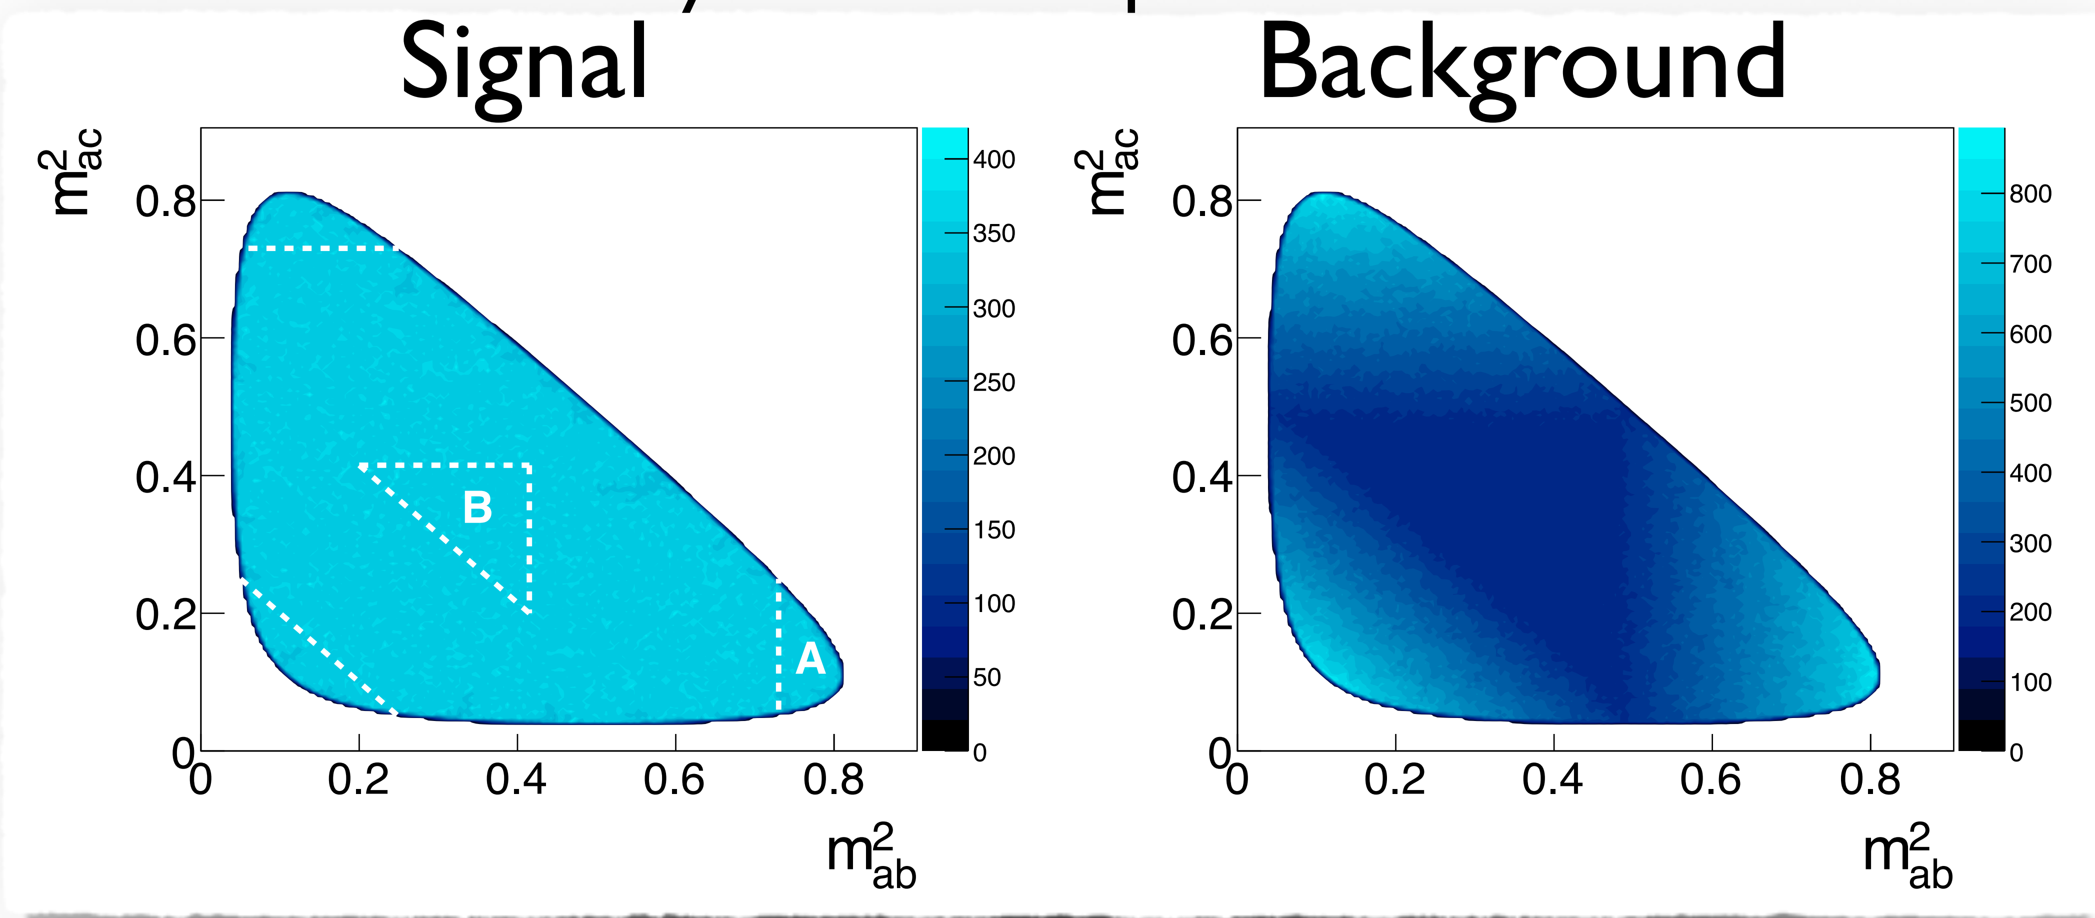

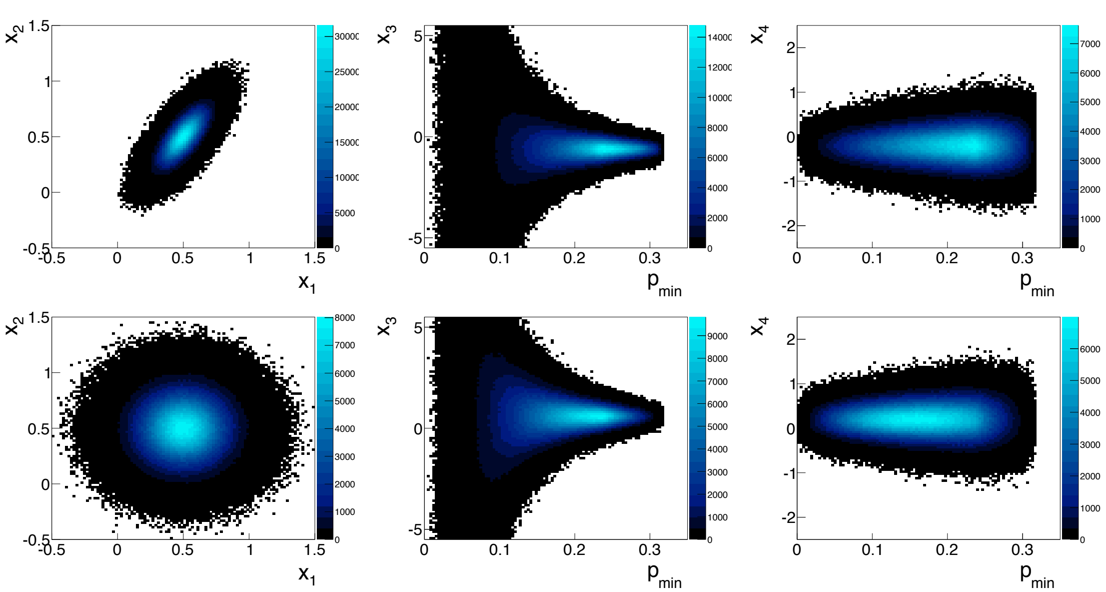

4 variables

- 2 correlated with Dalitz
- 2 uncorrelated

BDT

The problem

Toy Dalitz example X → abc

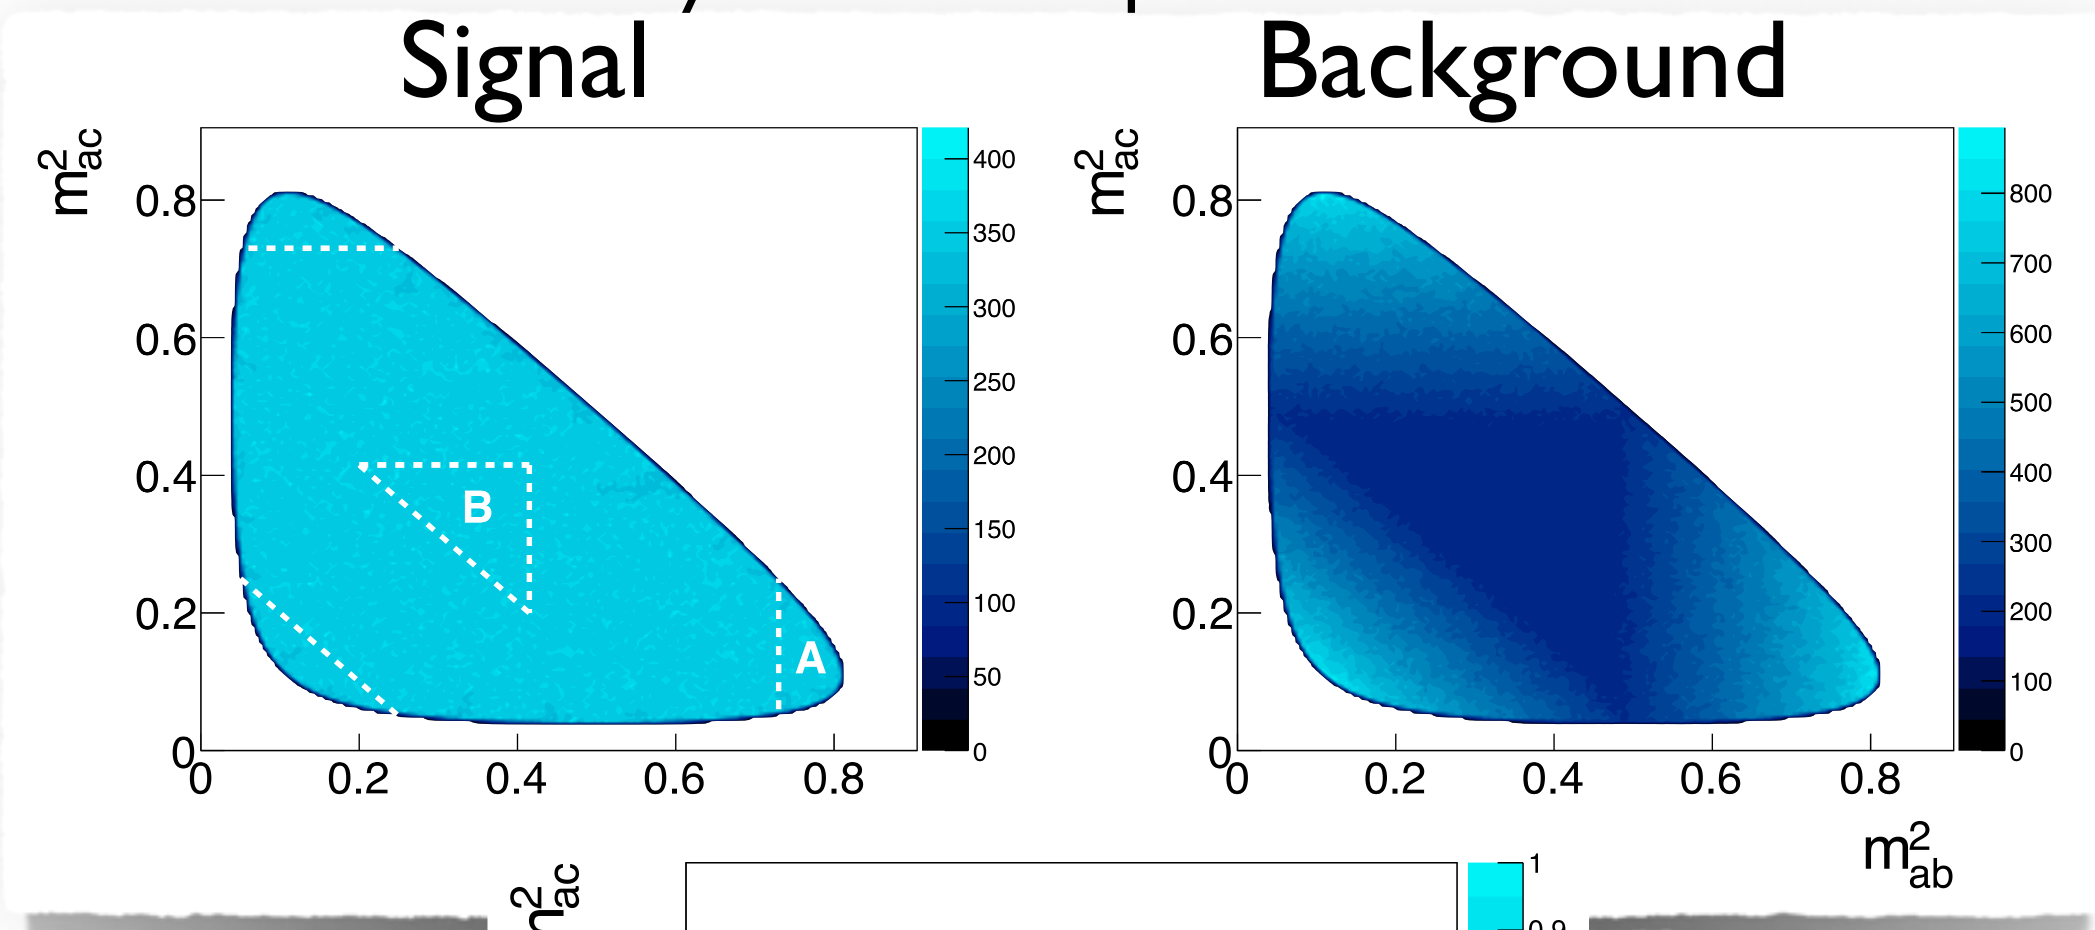

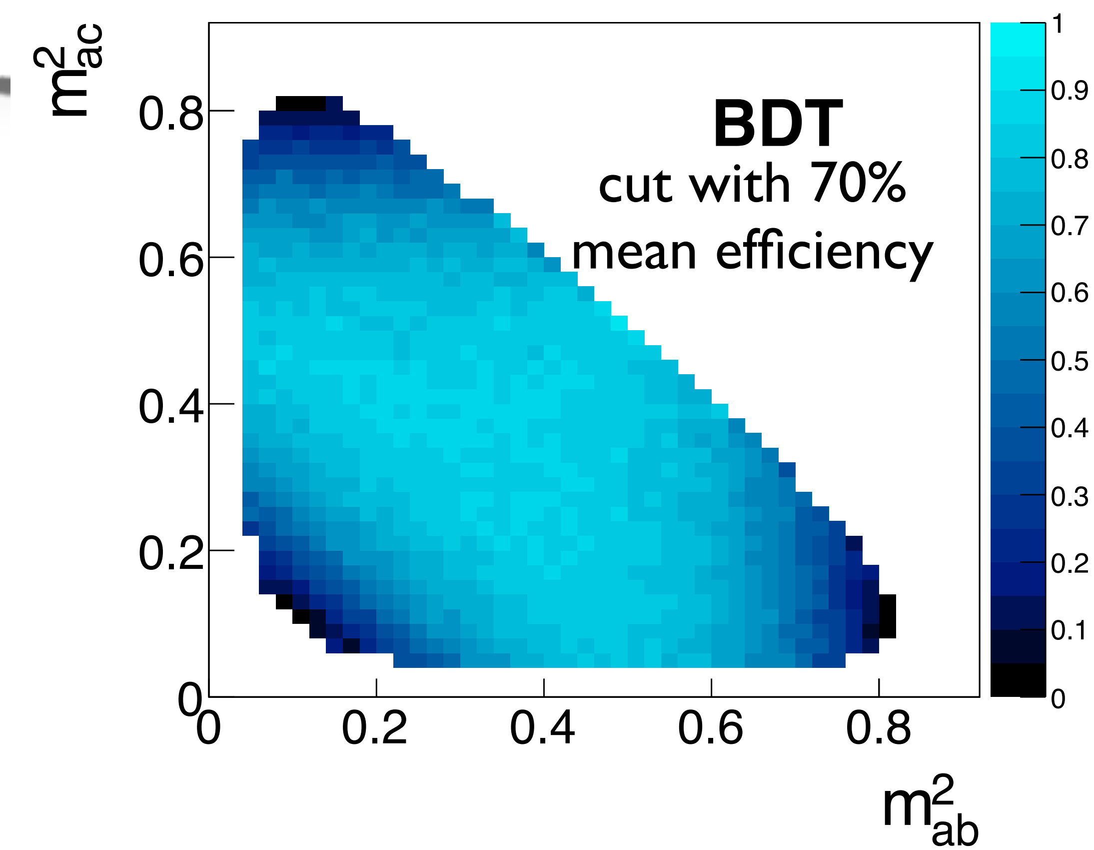

Efficiency

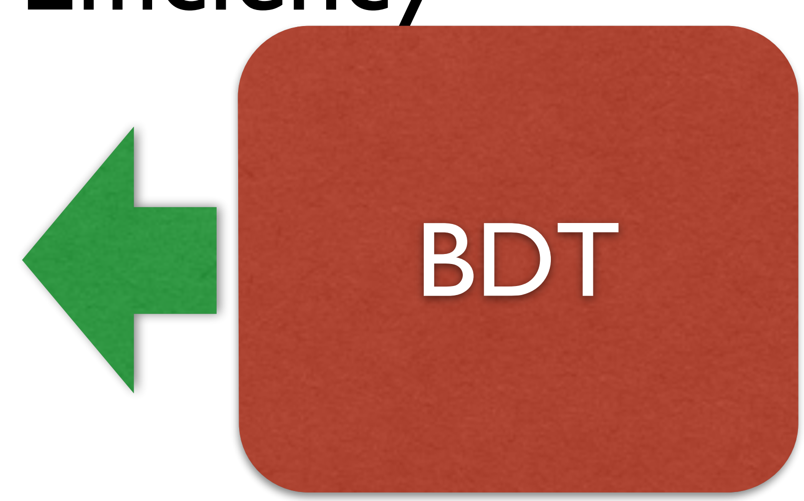

Biased Dalitz distribution!

The uBoost method

arXiv:1305.7248v2 [nucl-ex] 24 Dec 2013

uBoost: A boosting method for producing uniform selection efficiencies from multivariate classifiers.

Justin Stevens and Mike Williams

Massachusetts Institute of Technology, Cambridge, MA, United States

ABSTRACT:
The use of multivariate classifiers, especially neural networks and decision trees, has become commonplace in particle physics. Typically, a series of classifiers is trained rather than just one to enhance the performance; this is known as boosting. This paper presents a novel method of boosting that produces a uniform selection efficiency in a selected multivariate space. Such a technique is well suited for amplitude analyses or other situations where optimizing a single integrated figure of merit is not what is desired.

The uBoost method

- Boosting: train many trees, and assign larger weights to events if they were mis-classified in previous rounds.
- uBoost: assign larger weights to events in which the local efficiency is smaller.

For the same 70% mean efficiency

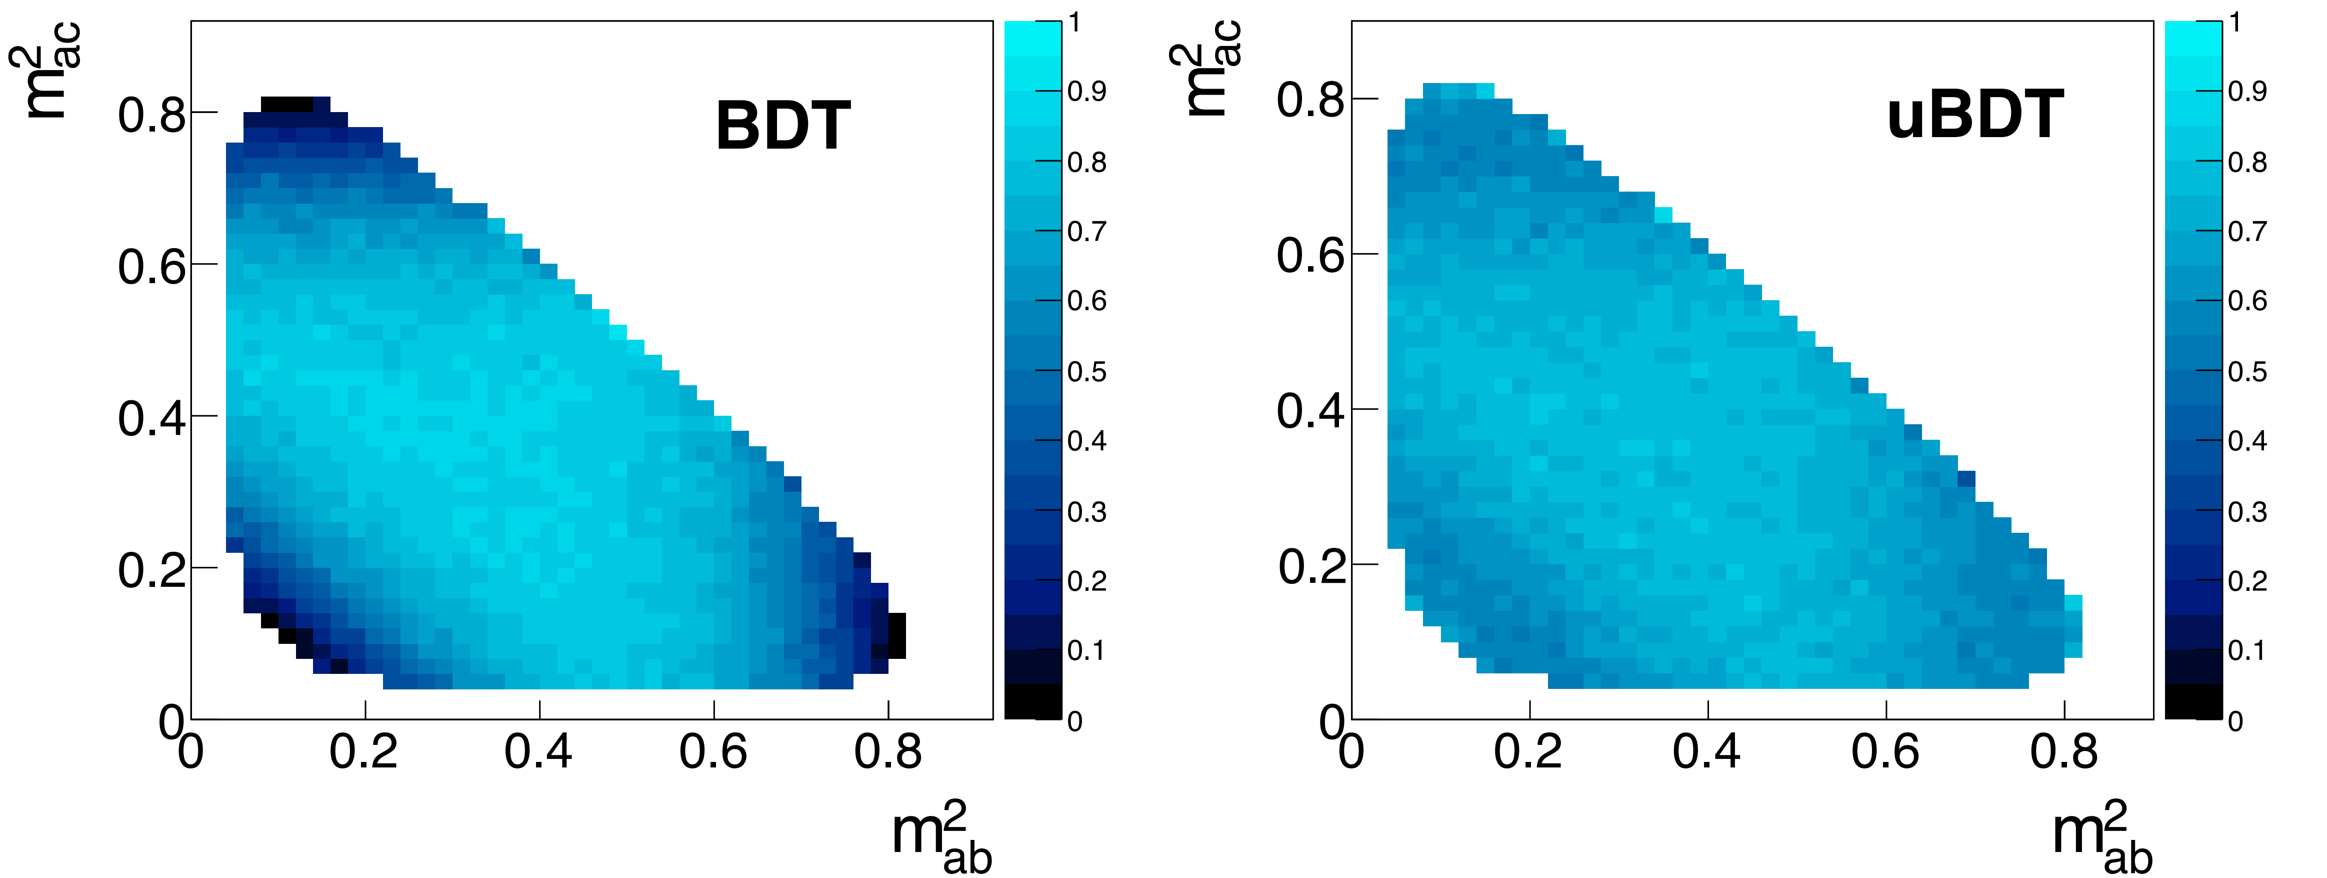

The uBoost method

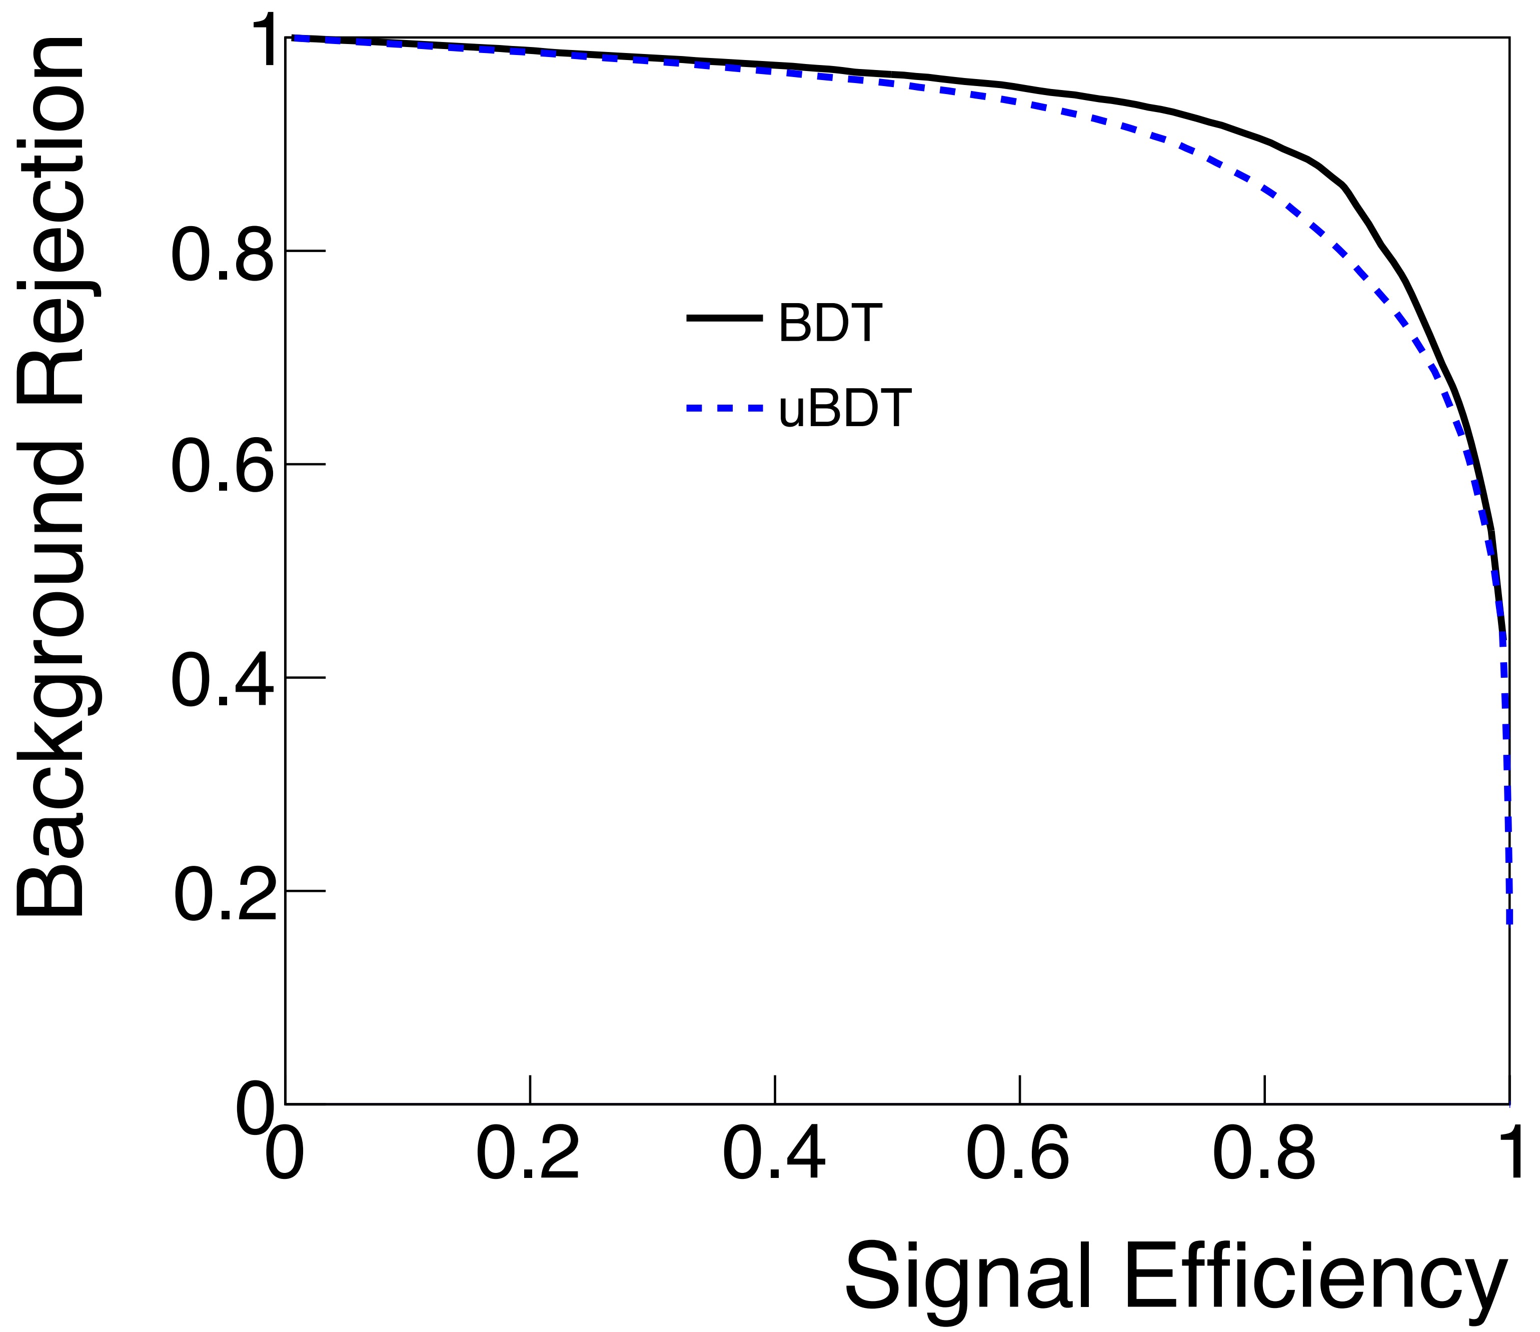

Closing thought

Surely this is only the beginning of MLHEP in triggers.

- Pattern recognition
- Detector and DQ monitoring
- Data popularity / trigger bandwidth division
- Softer multibody signatures in ATLAS/CMS
- More efficient exclusive lines in LHCb
- ….