

New strategies for LLP searches at LHCb: **BuSca** (*Buffer Scanner*)

Arantza Oyanguren (IFIC-Valencia)

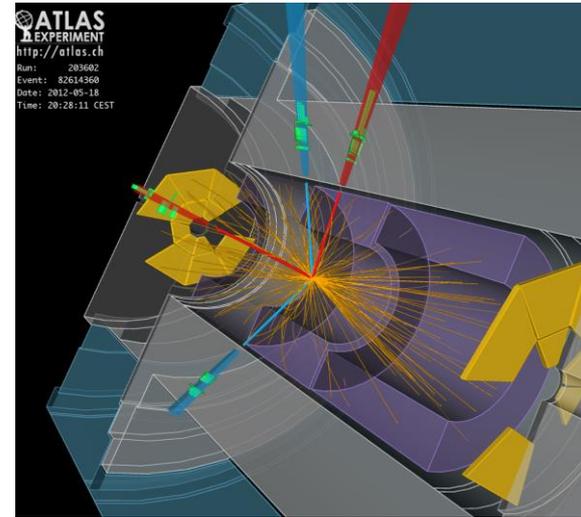
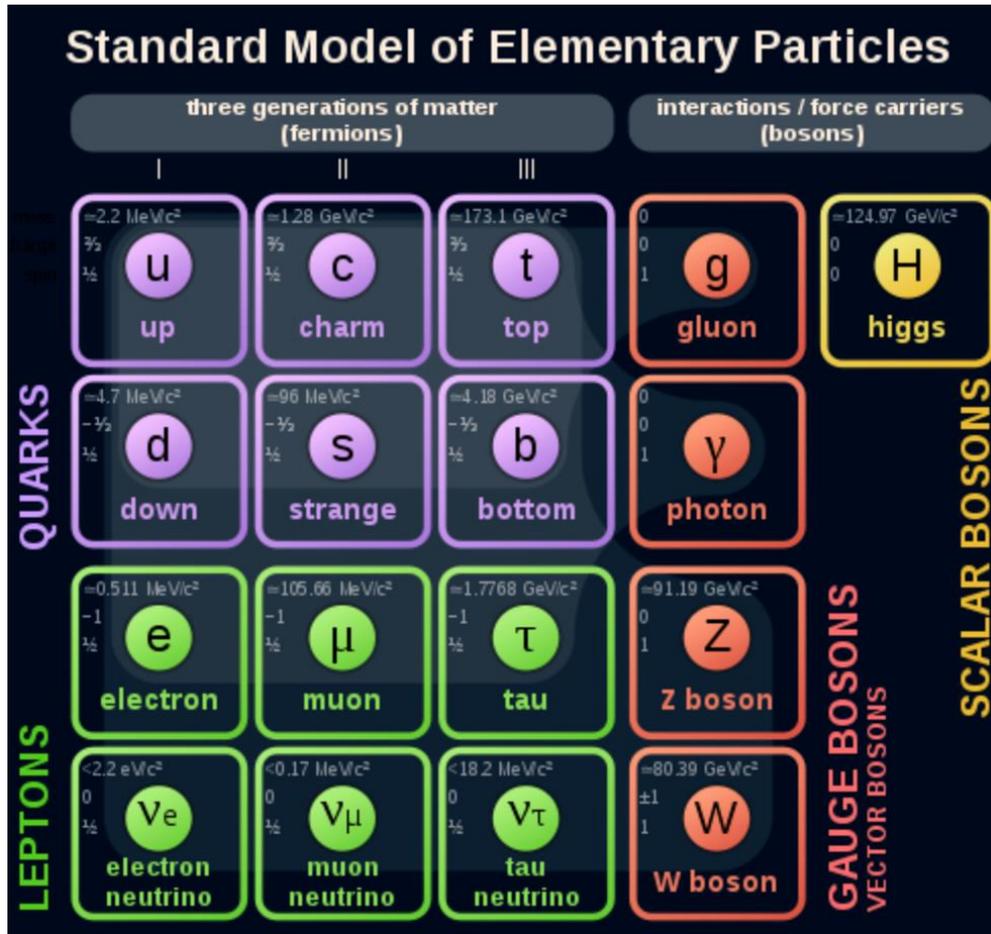
(& J. Zhuo, V. Kholoimov)



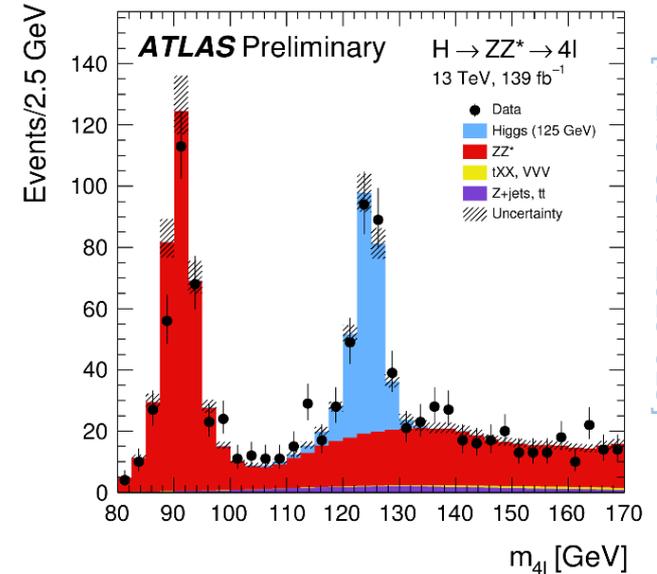
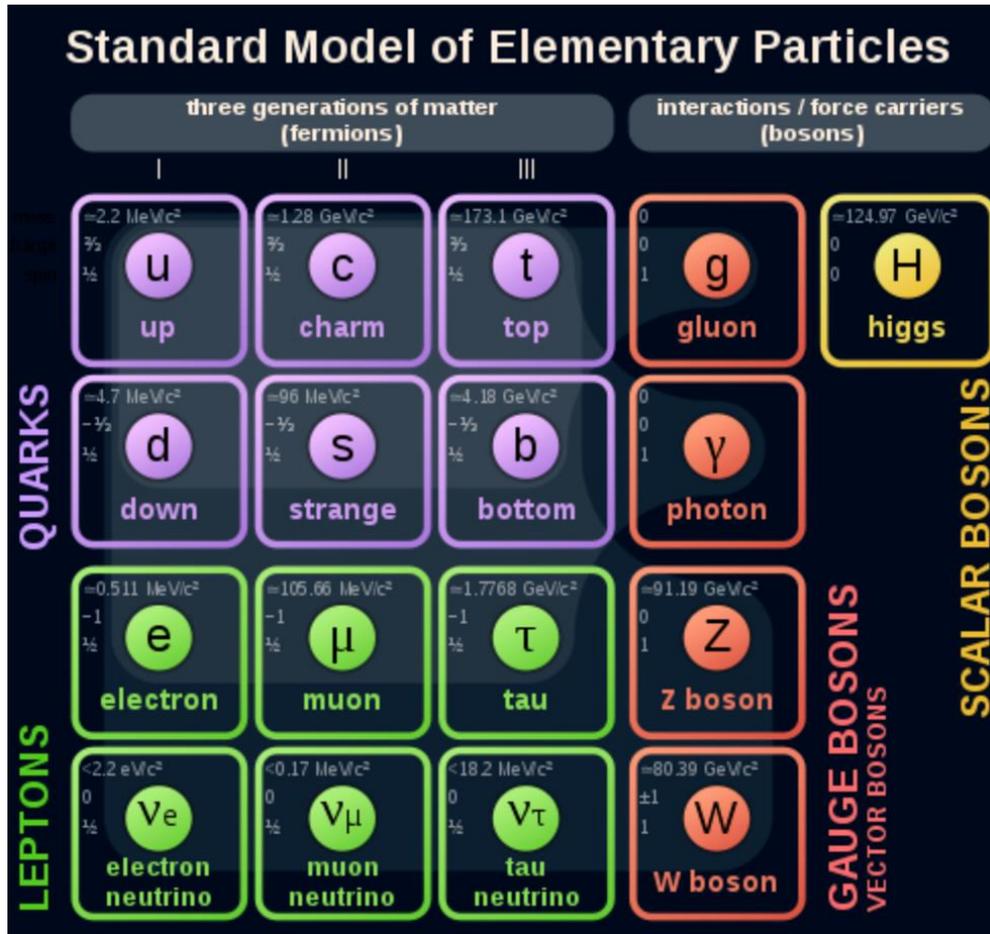
Outline

- Introduction
- The LHCb experiment
- The new LHCb trigger
- Downstream tracks
- BuSca (Buffer Scanner)
- Results with first data
- Conclusions & prospects

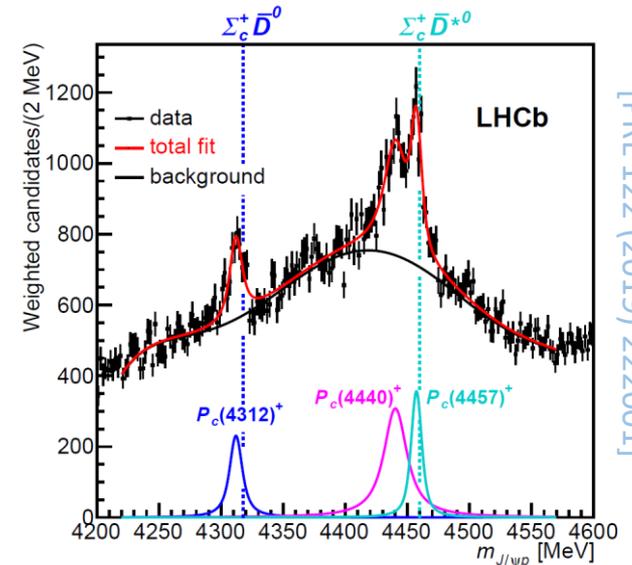
Introduction



Introduction



[ATLAS-CONF-2019-025]



[PRL 122 (2019) 222001]

HISTORY OF THE UNIVERSE

13.8 billion years

Dark energy accelerated expansion

Structure formation

Cosmic Microwave Background radiation is visible

RHIC & LHC heavy ions

Accelerators

LHC protons

High-energy cosmic rays

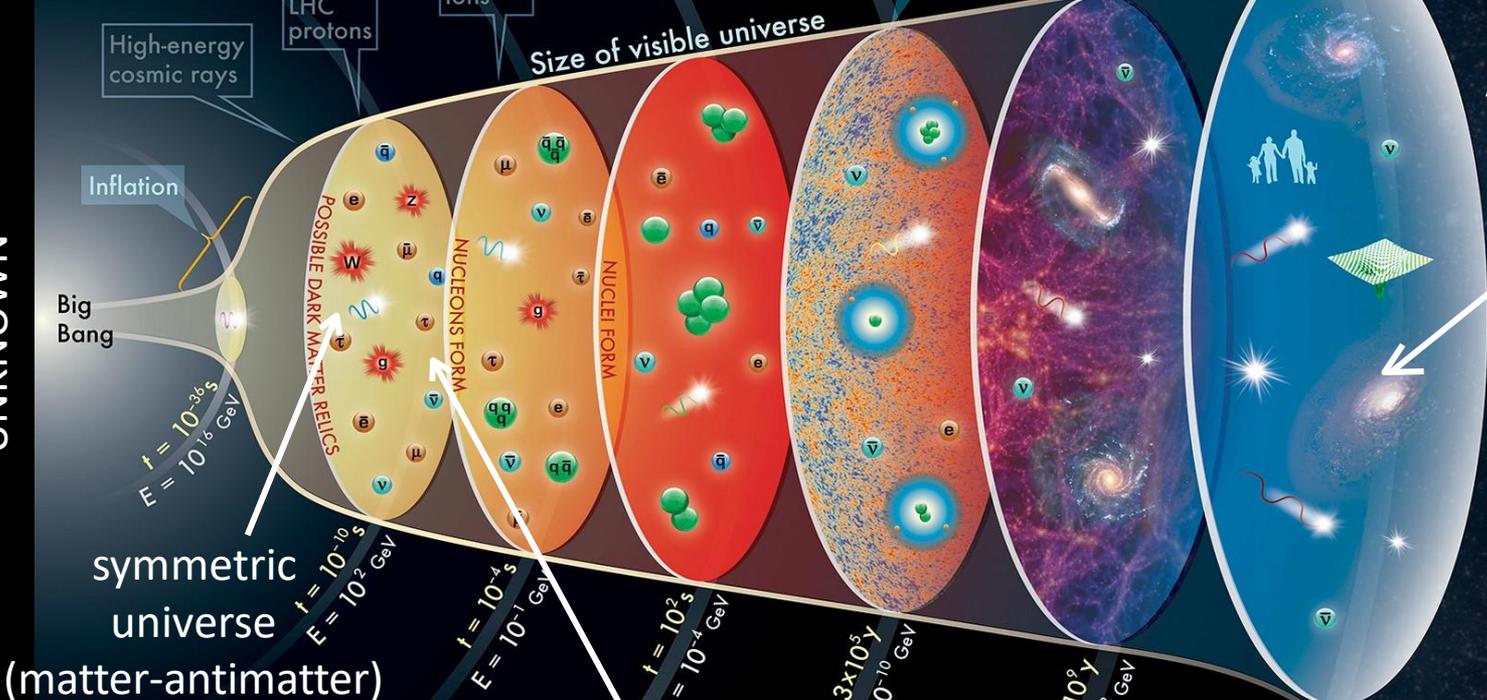
Inflation

Big Bang

TODAY

Anti-symmetric a(matter)

UNKNOWN



symmetric universe (matter-antimatter)

t = Time (seconds, years)
E = Energy (GeV)

Key

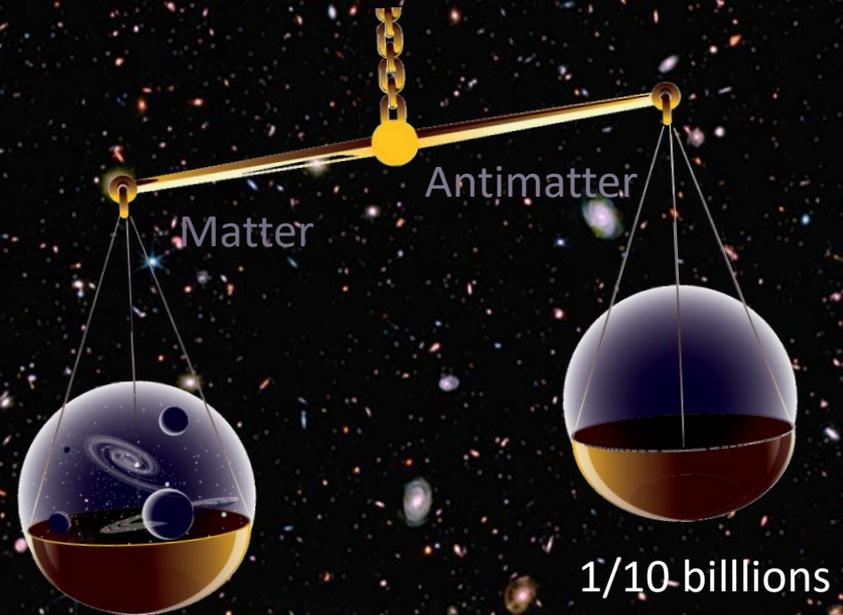
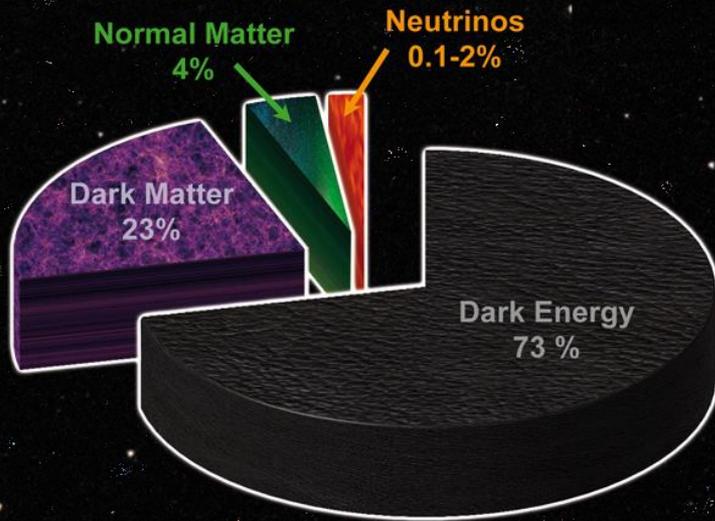
quark	neutrino	ion	star
gluon	bosons	atom	galaxy
electron	meson	photon	black hole
muon	baryon		
tau			

¿What did it happen?



Introduction

Misteries of our Universe...



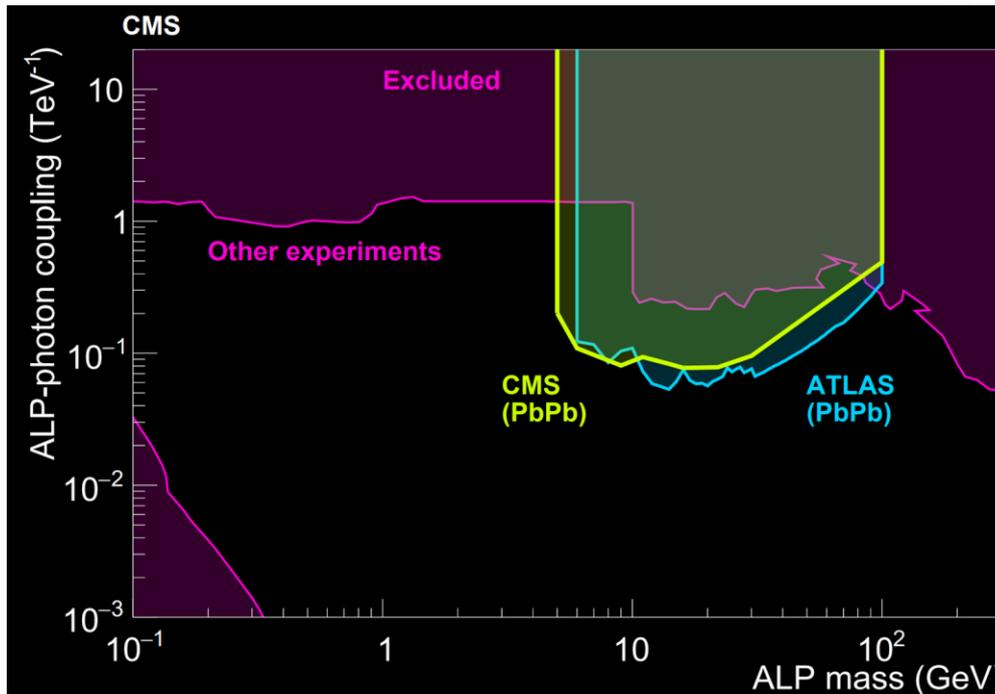
What is dark matter? Where is antimatter? Which particle is describing gravity? 6

Introduction

Dark Bosons

Heavy Neutral Leptons

Composite Higgs



SUSY

Massive Dark Photons

Extra dimensions

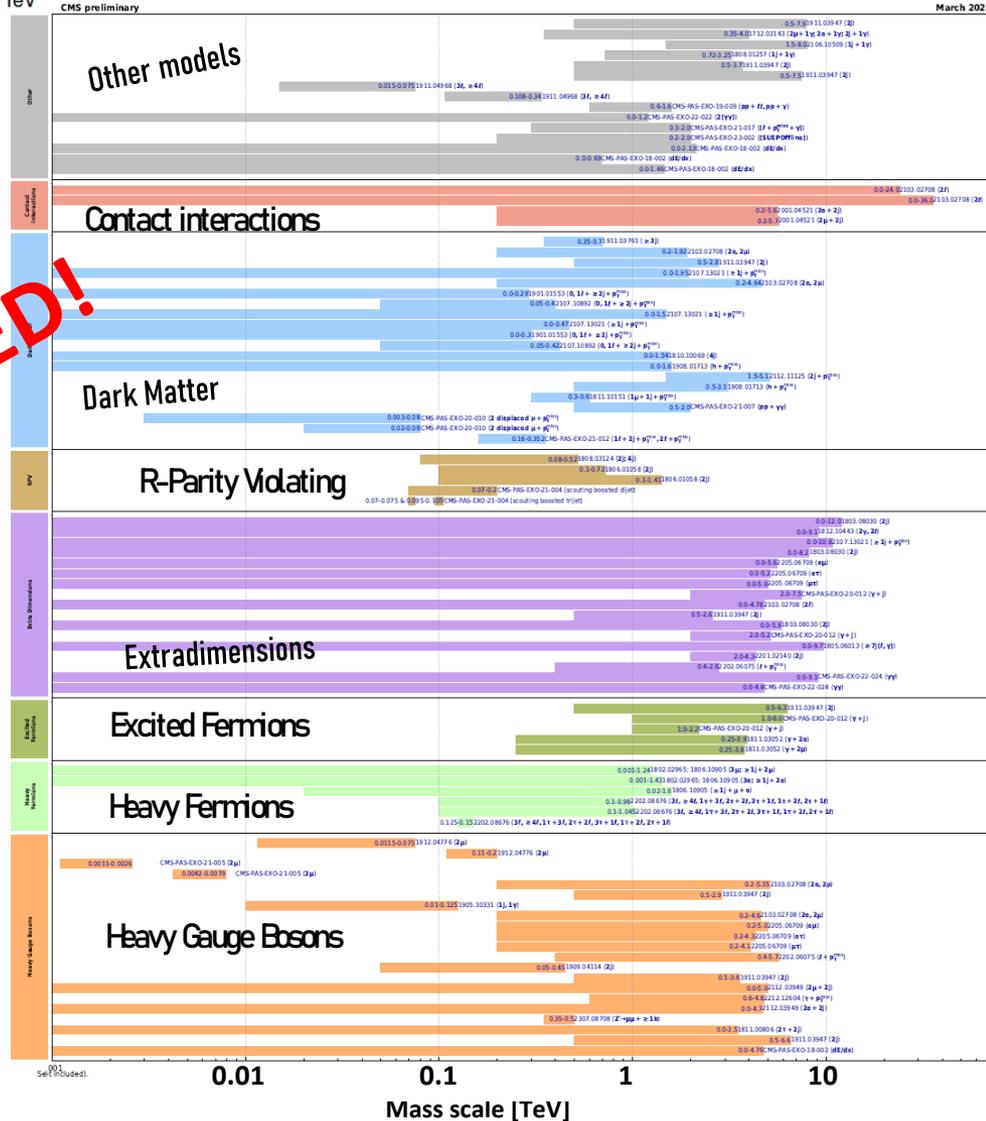
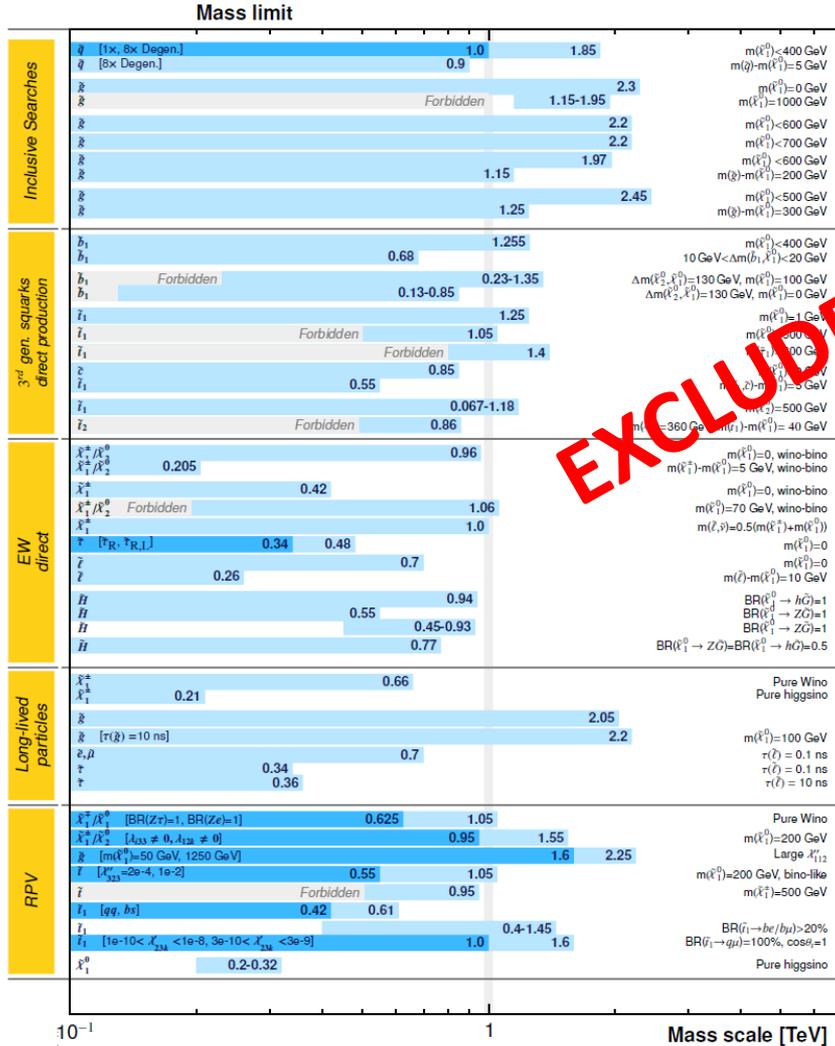
Axion-like particles

Introduction

(Extracted from) [CMS Twiki](#)

ATLAS SUSY Searches* - 95% CL Lower Limits ATLAS Preliminary
August 2023 $\sqrt{s} = 13$ TeV

Overview of CMS EXO results

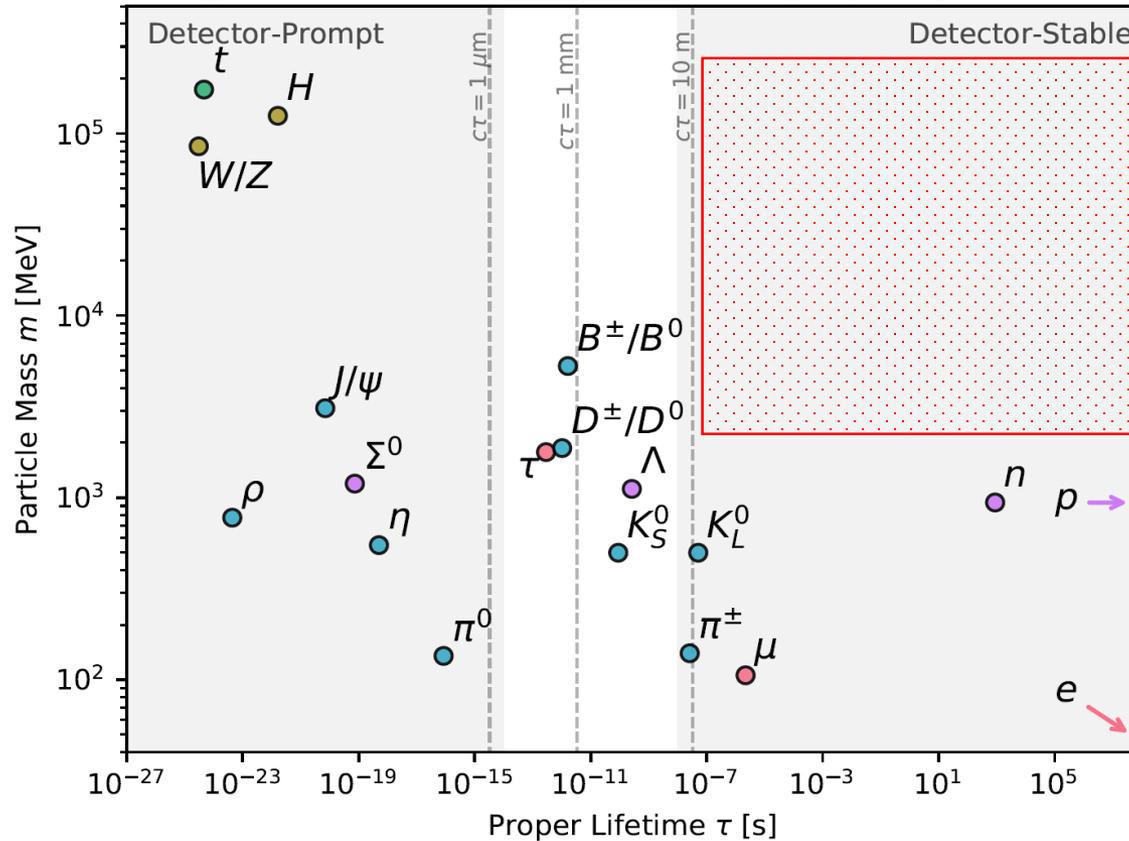


(Extracted from) [ATL-PHYS-PUB-2023-025](#)

Introduction

Long-lived particles in the SM and Beyond

[Prog. in Part. and Nuc. Phys. 106 (19)210]



Unexplored region ?

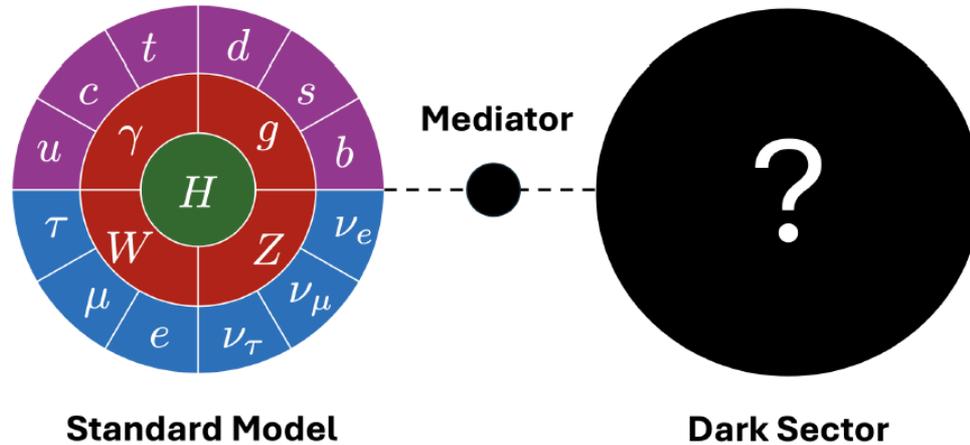
Effect of LHCb trigger in Run2:

[Front. Big Data, '22 Big Data and AI in HEP]

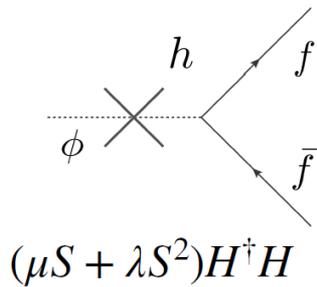
Introduction

Long-lived particles Beyond the SM?

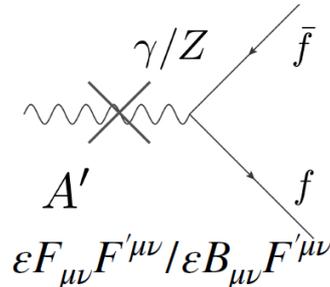
The dark sector:



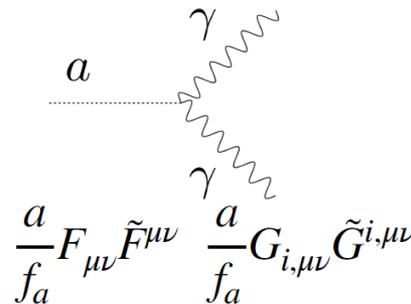
Portal examples:



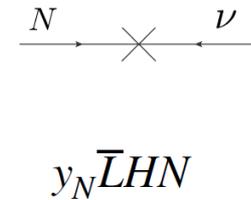
Higgs portal



Vector portal

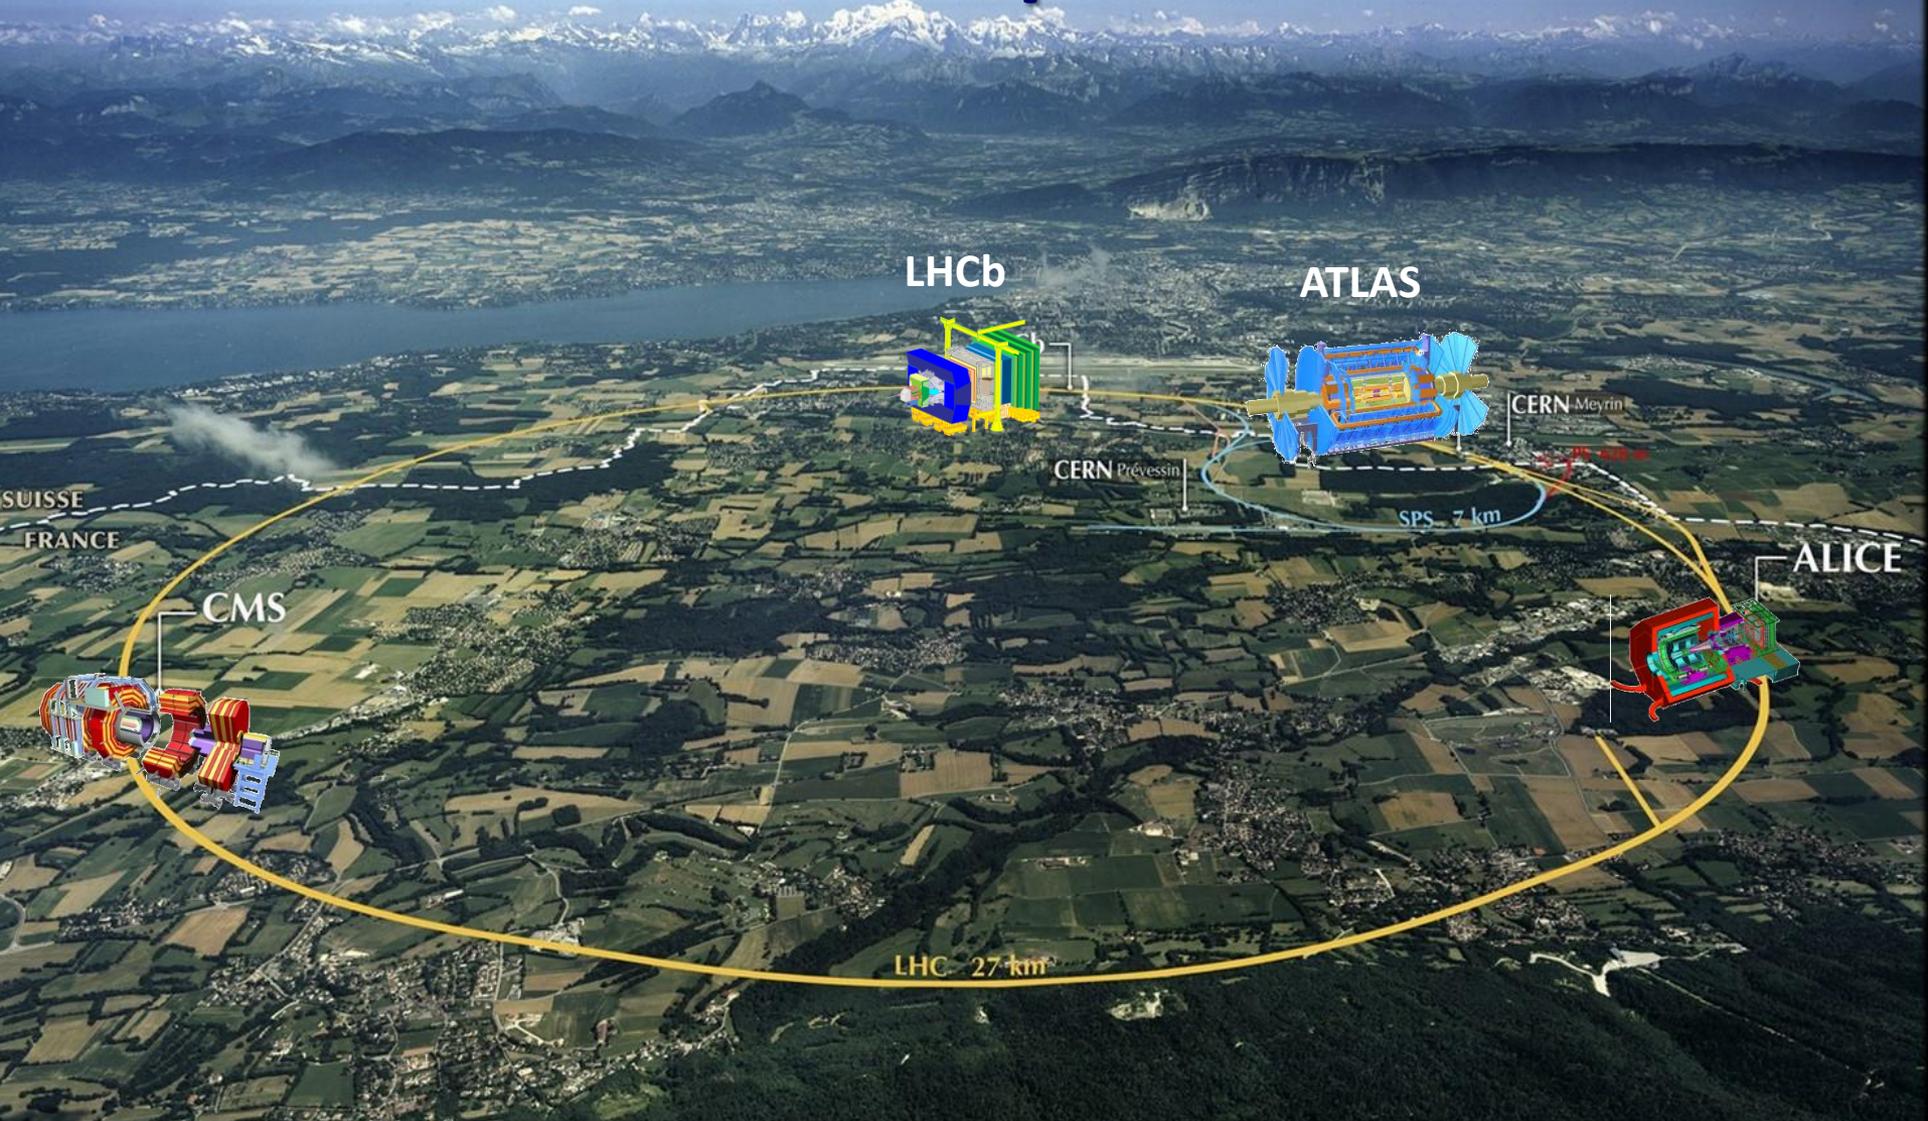


Axion portal



Neutrino portal

The LHCb experiment



LHC: the proton-proton collider at CERN with an energy of 13.6 TeV

The LHCb experiment



The LHCb experiment

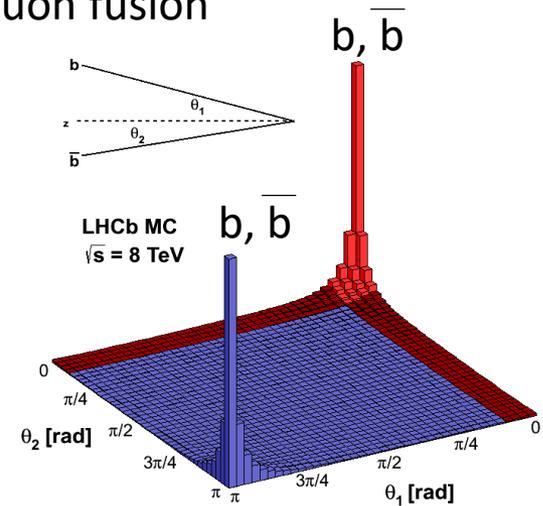
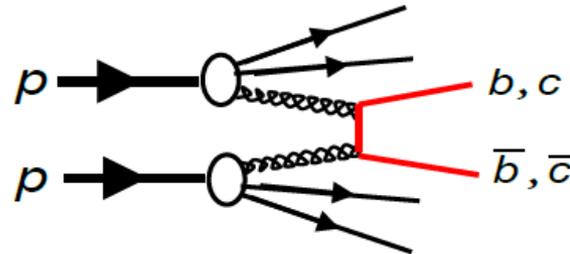
- The $b\bar{b}$ cross section in pp collisions is large, mainly from gluon fusion

~ 300 μb @ $\sqrt{s}=7$ TeV

~ 600 μb @ $\sqrt{s}=13$ TeV

[PRL 118 (2017) 052002]

[JHEP 02 (2021) 023]



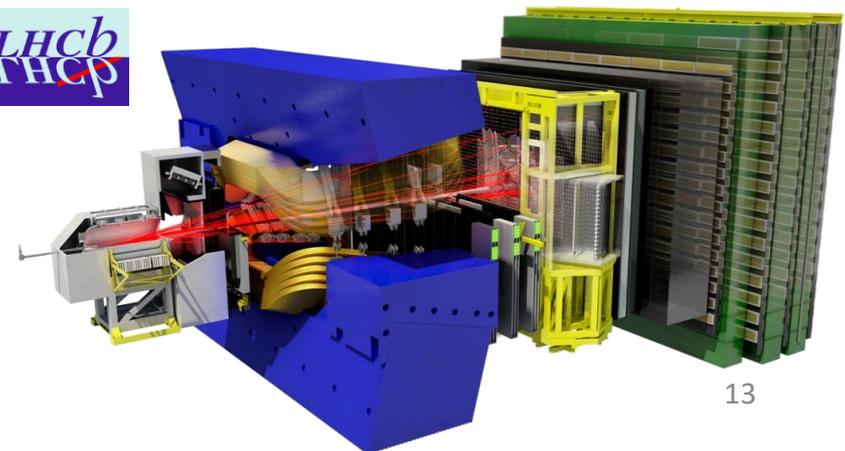
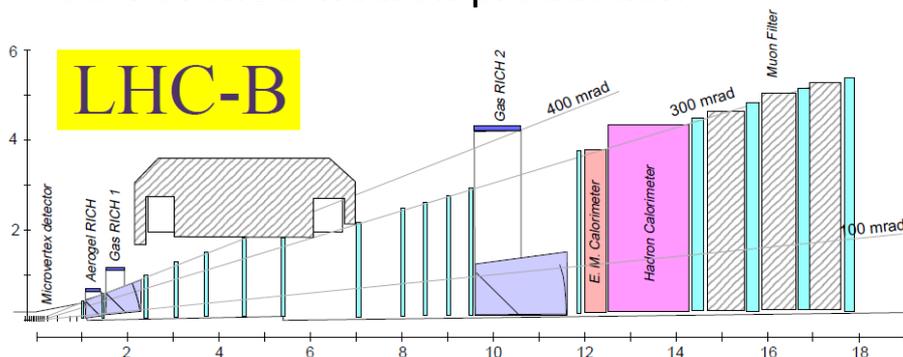
The b quarks hadronize in $B, B_s, B^*_{(s)}, b$ -baryons...

→ average B meson momentum ~ 80 GeV

- The LHCb idea: to build a single-arm forward spectrometer:

~ 4% of the solid angle ($2 < \eta < 5$),

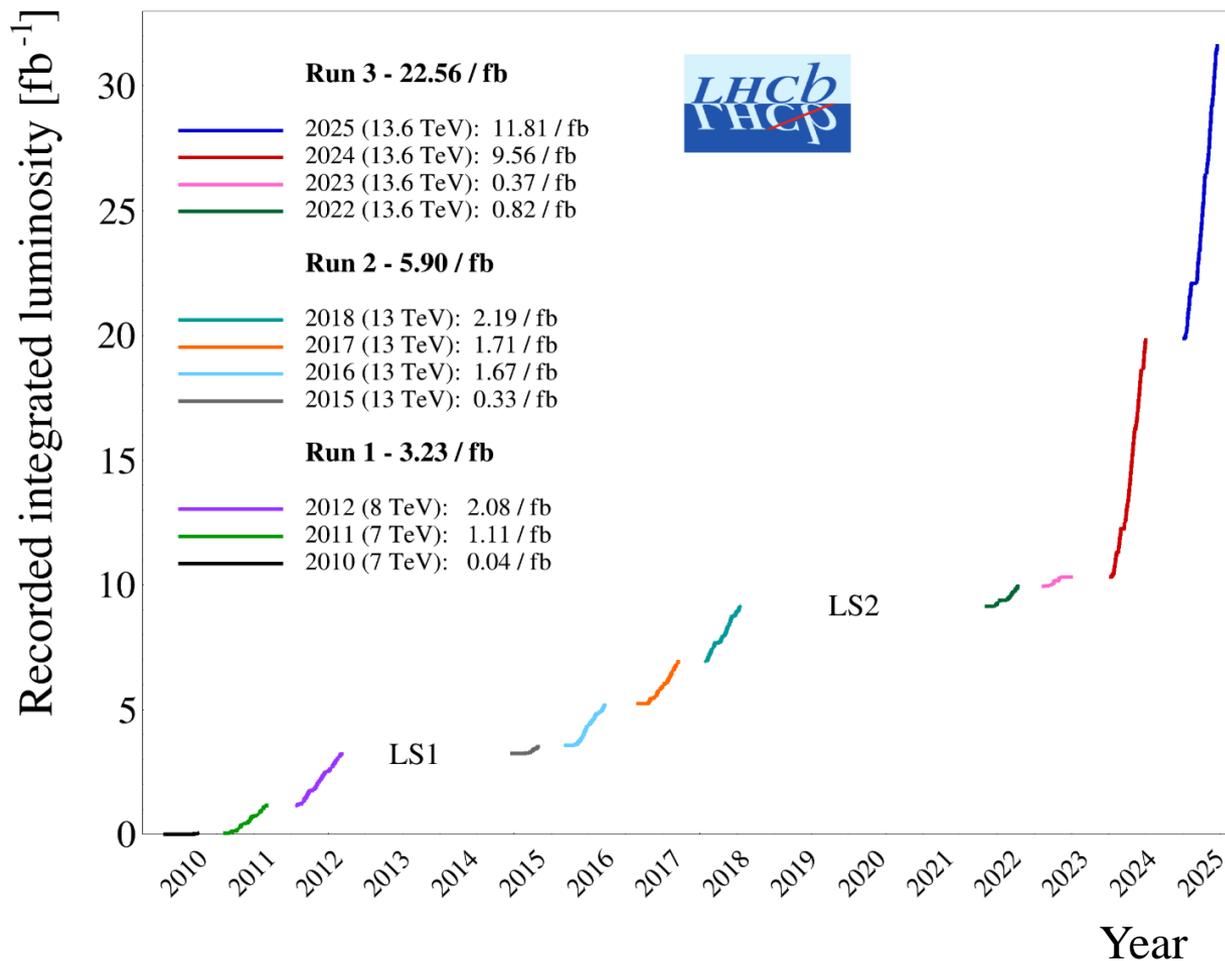
~30% of the b hadron production



The LHCb experiment

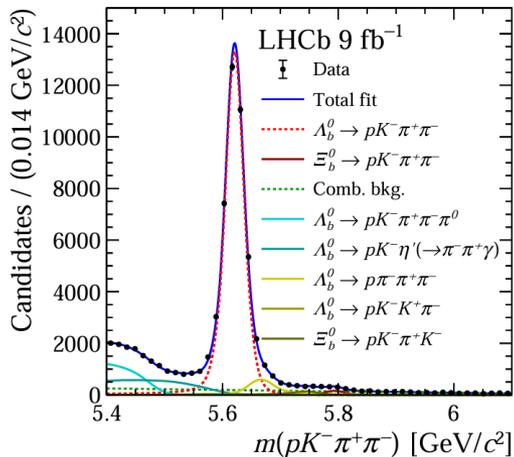
Taking data since 2011:

Total recorded luminosity – pp – 31.7 fb^{-1}

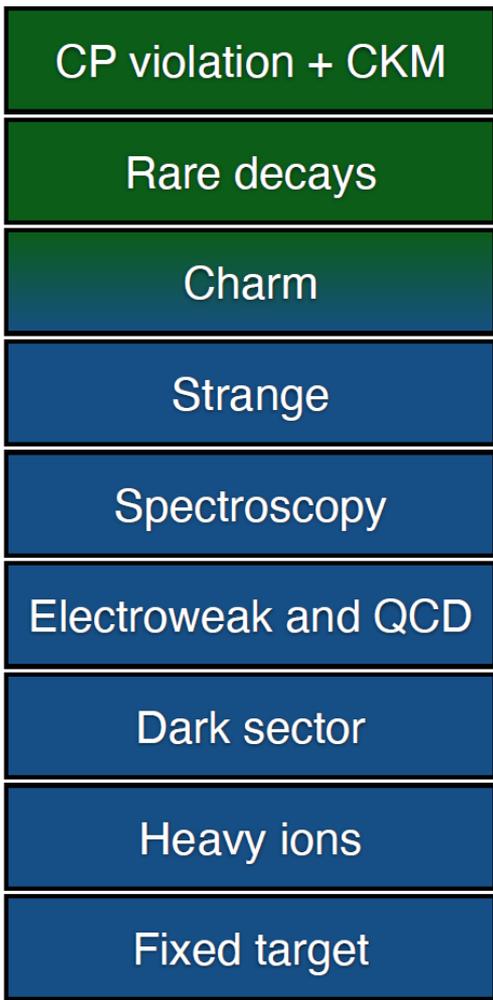


The LHCb experiment

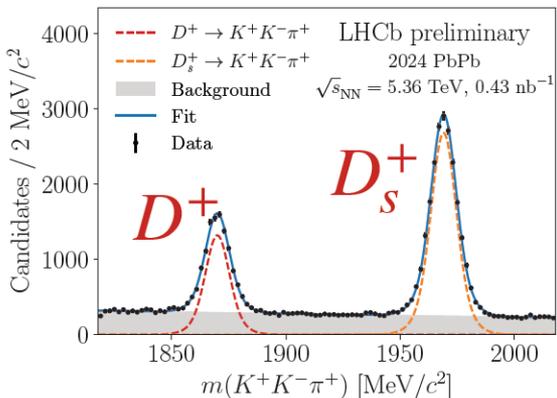
Nature 643 (2025) 1223



original programme

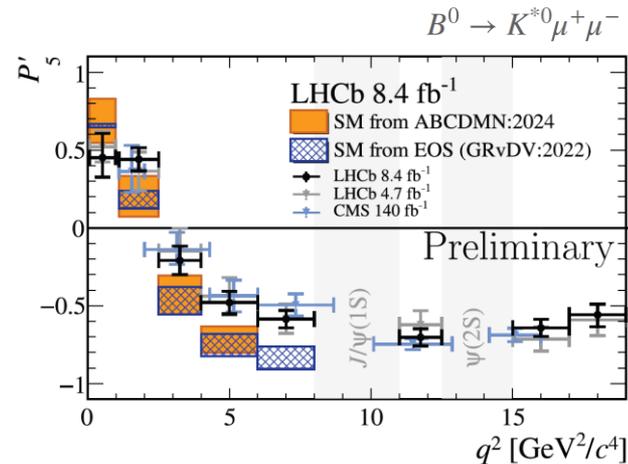


LHCb physics

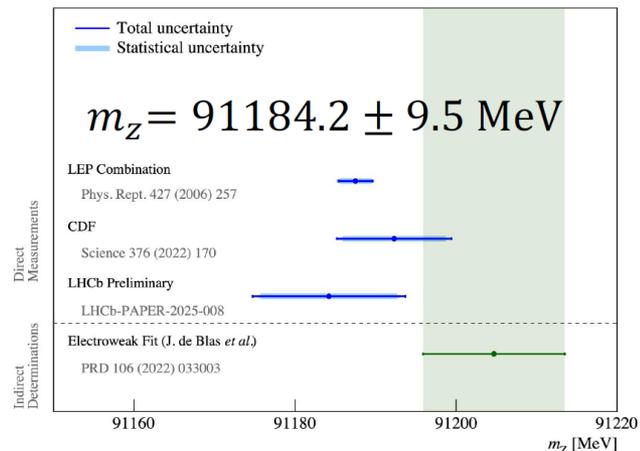


LHCb-FIGURE-2025-004

LHCb-PAPER-2025-041
(in preparation)

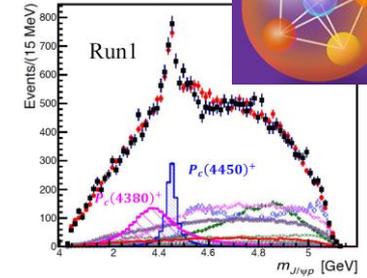
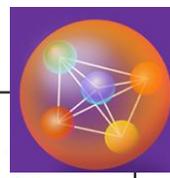


PRL 135 (2025) 161802

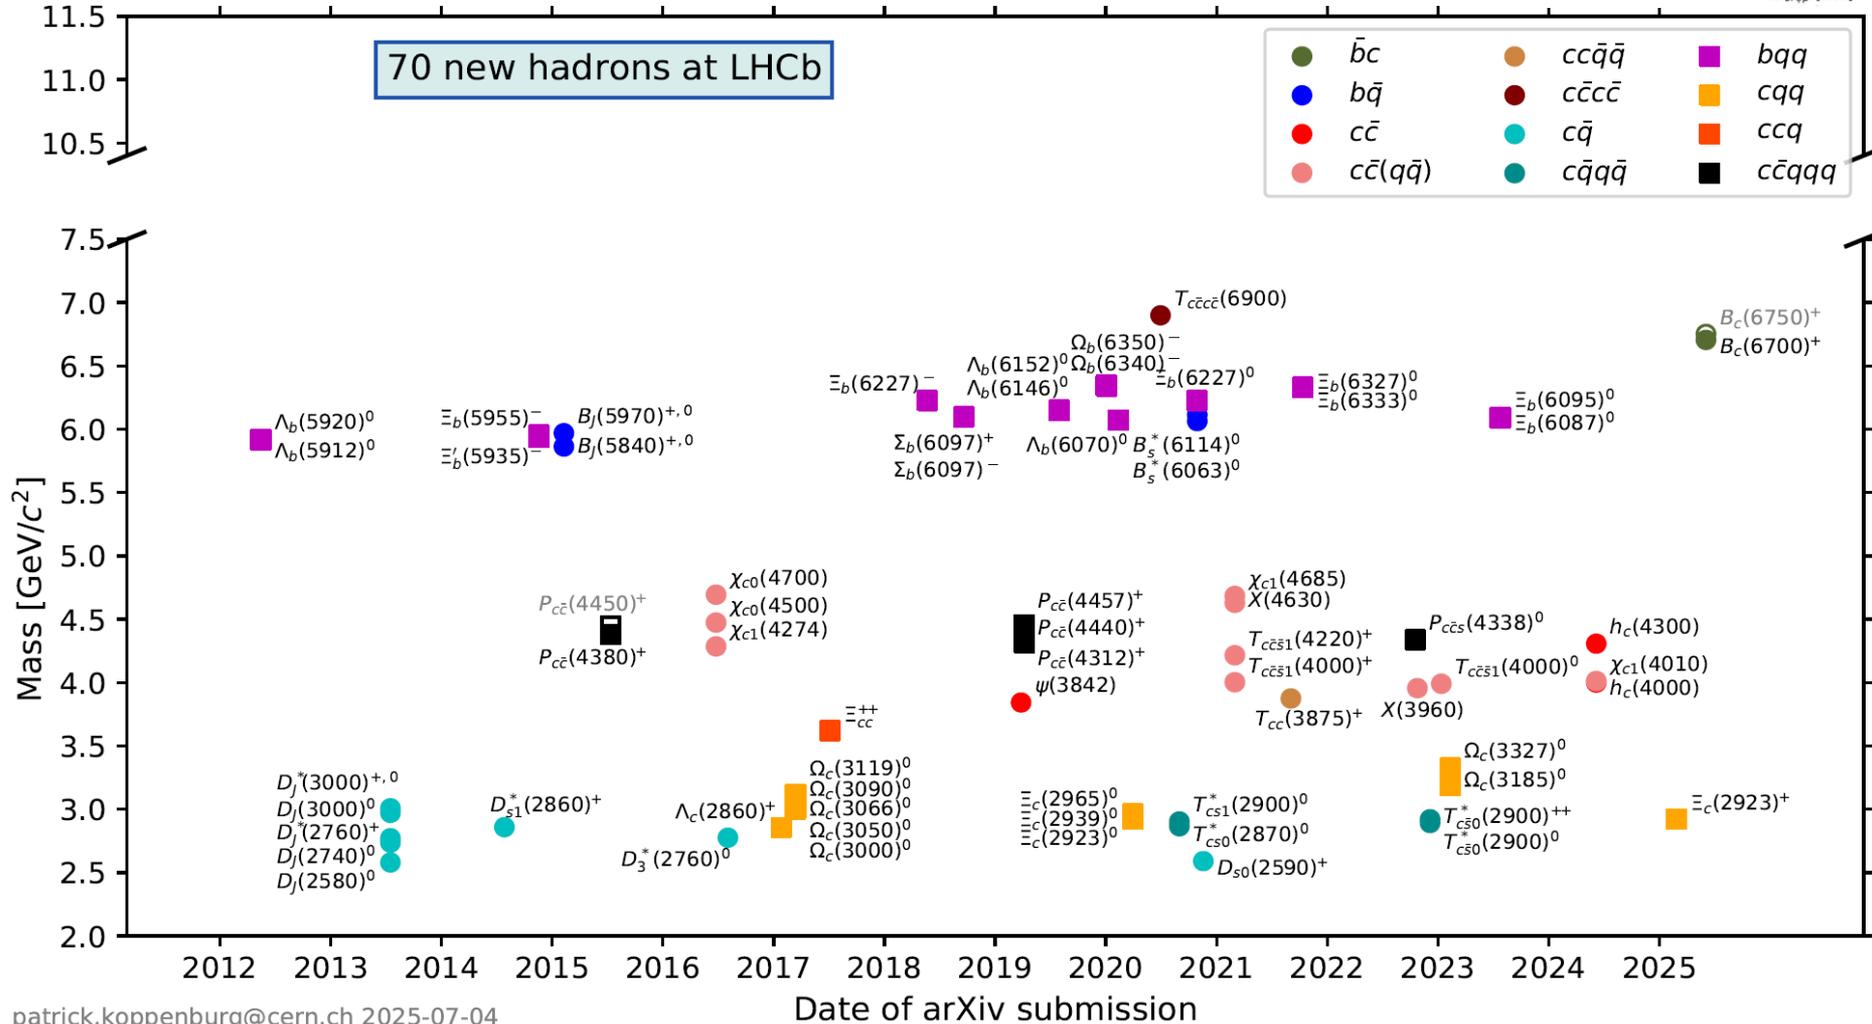


The LHCb experiment

[PRL 115 (2015) 072001]



Spectroscopy



The LHCb experiment

The LHCb upgraded detector (operating since 2023):

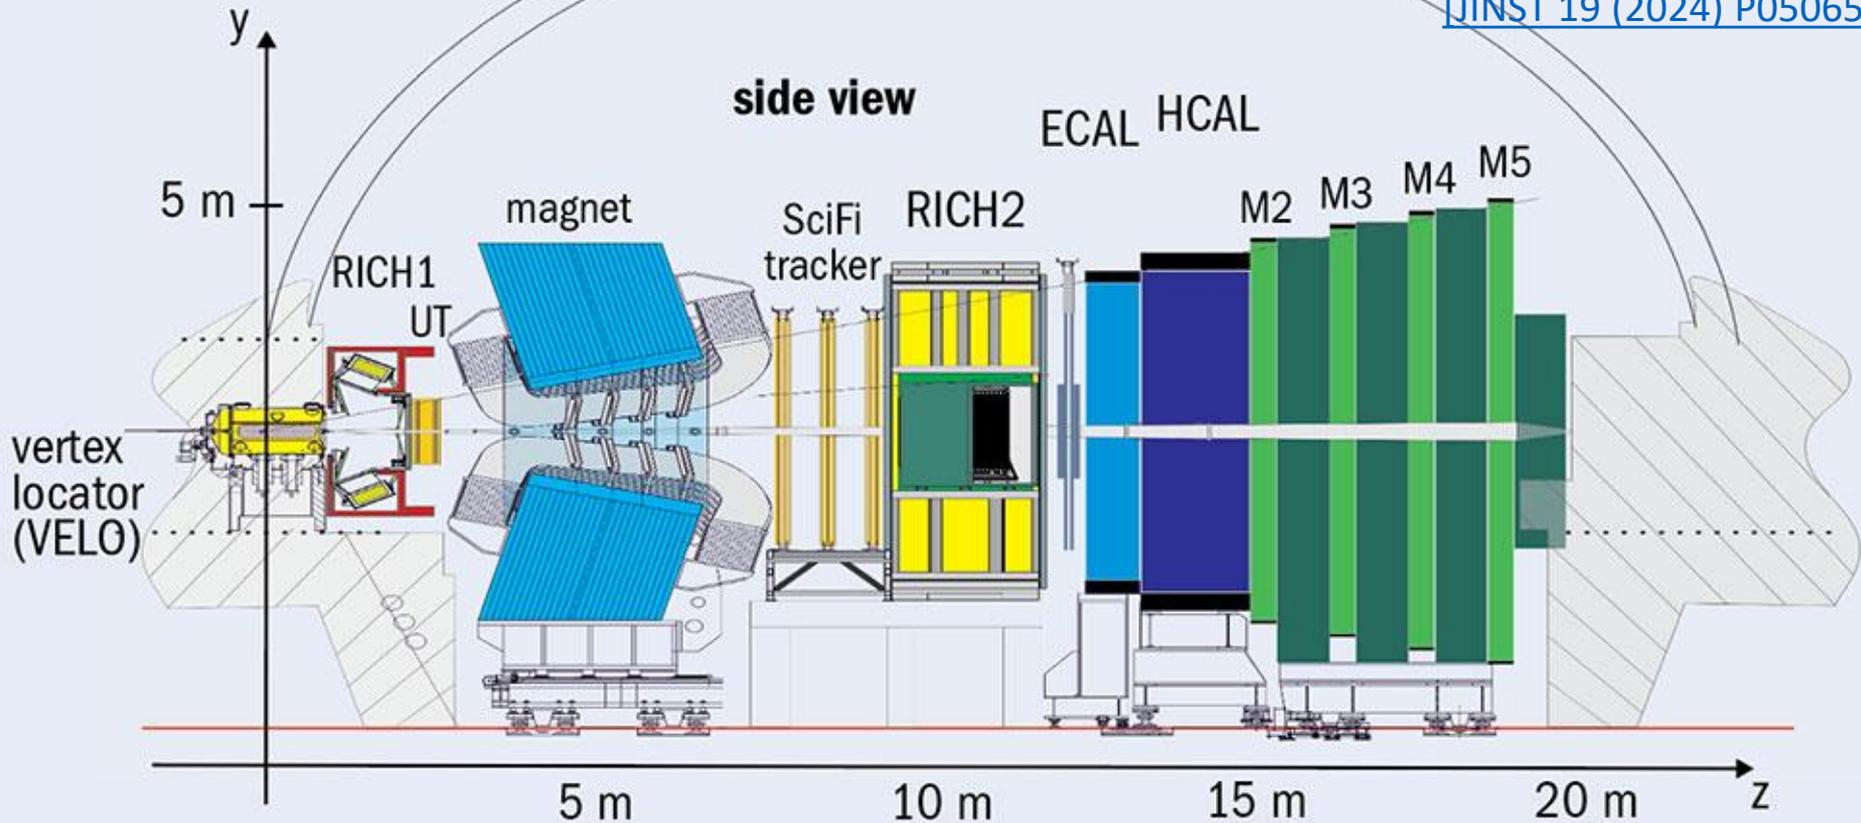
And a new fully software trigger based on GPUs!

Silicon strips

[\[LHCb-TDR-021\]](#)

[\[Com. Softw. for Big Sci. 6, 1 \(2022\)\]](#)

[\[JINST 19 \(2024\) P05065\]](#)

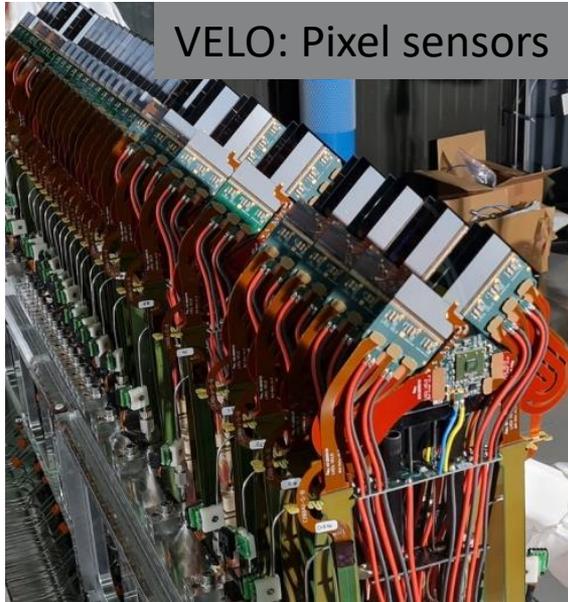


New pixel detector

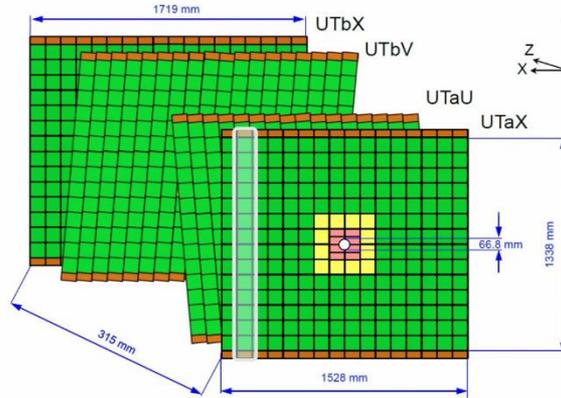
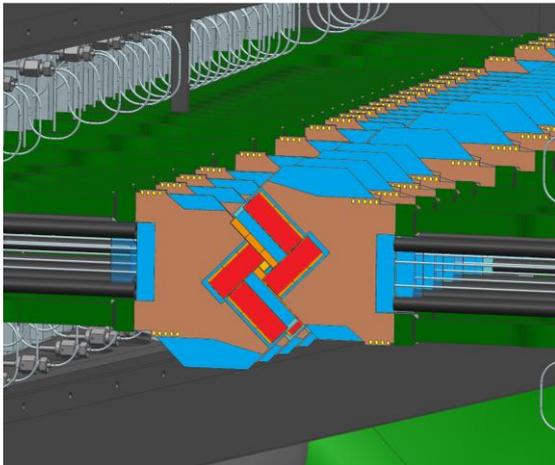
Scintillator fibers

The LHCb experiment

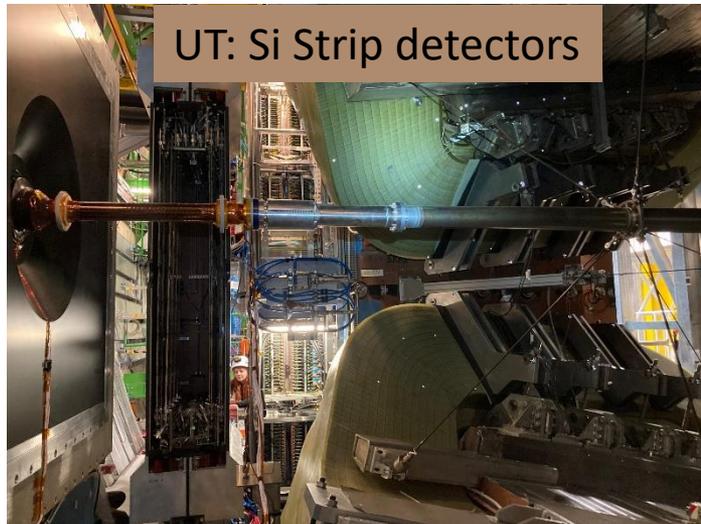
The LHCb trackers:



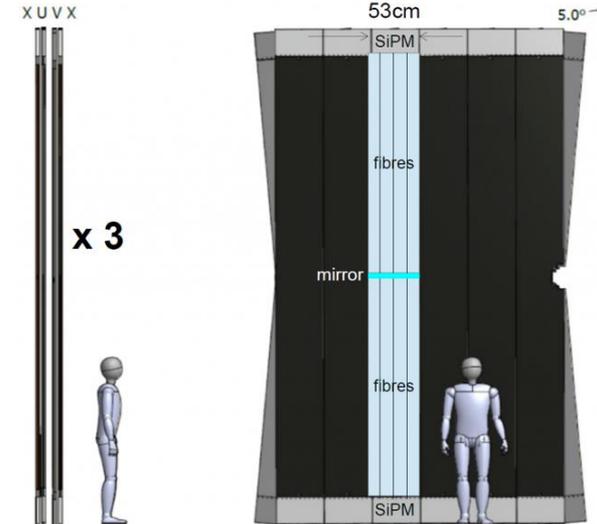
52 modules. Average: 42 hits/module.



4 layers. Average: 350 hits/layer.



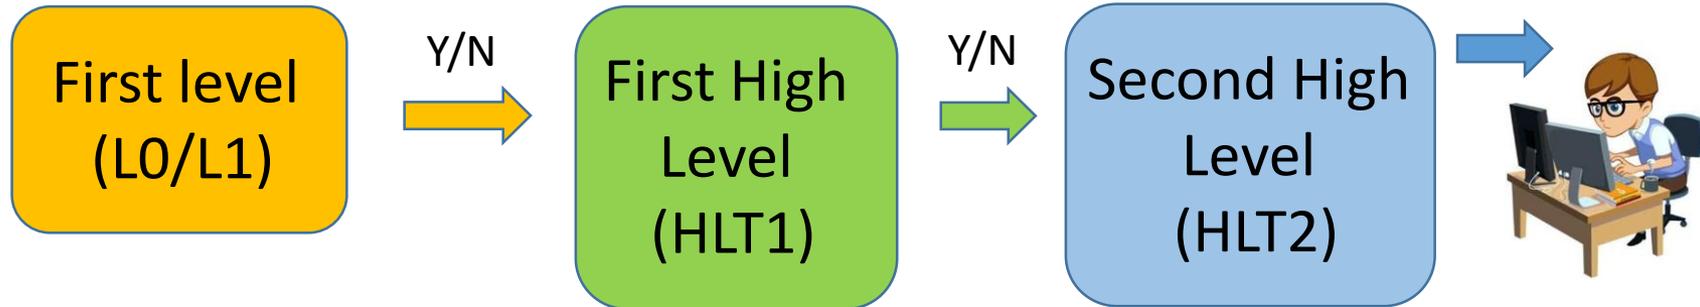
4 layers x 3 stations. Ave: 400 hits/layer



The new LHCb trigger

It is impossible to select all the data: need to select the events of interest

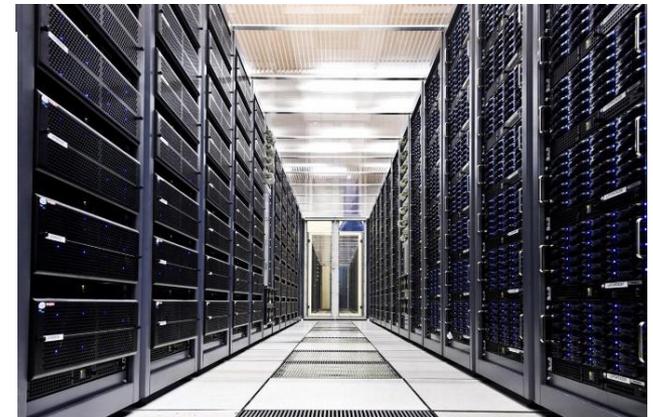
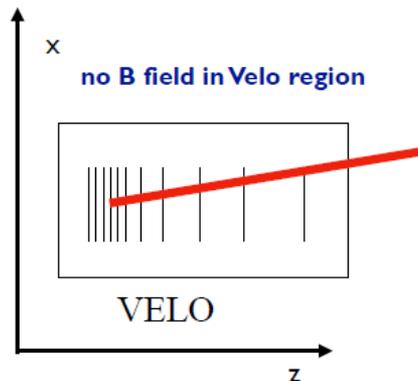
Traditional trigger systems:



Custom electronics (FPGAs),
Information from calorimeters
and muon stations

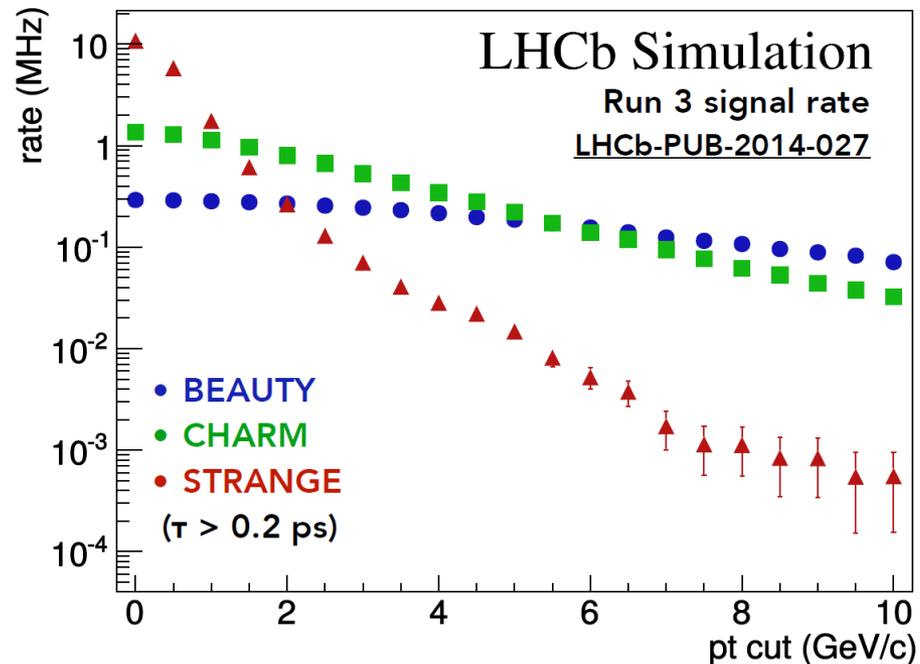
Processors farm,
fast Information
from tracking

Processors farm,
detailed information
to reconstruct the event



The new LHCb trigger

- Aim of LHCb Run3: Increasing the luminosity x 5 as compared to Run2 ($L = 2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$)
- The problem: the trigger rate saturated for b and charm physics



We were already selecting all the possible signals in our Run2 trigger scheme!

The new LHCb trigger

Removing our bottleneck:

~~First level
(LO/L1)~~

Y/N



First High
Level
(HLT1)

Y/N



Second High
Level
(HLT2)



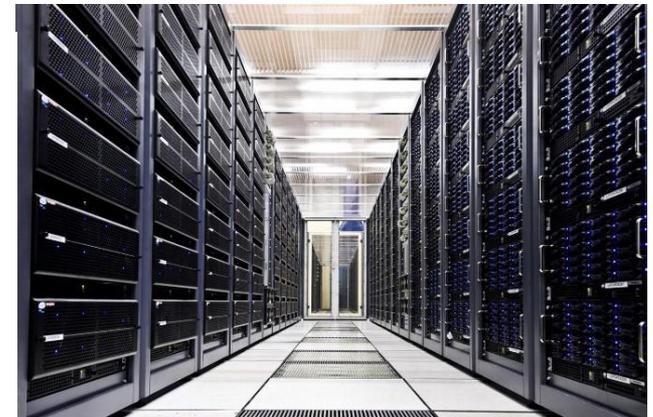
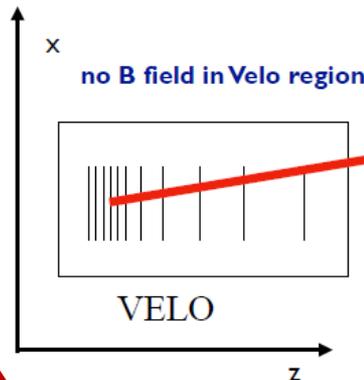
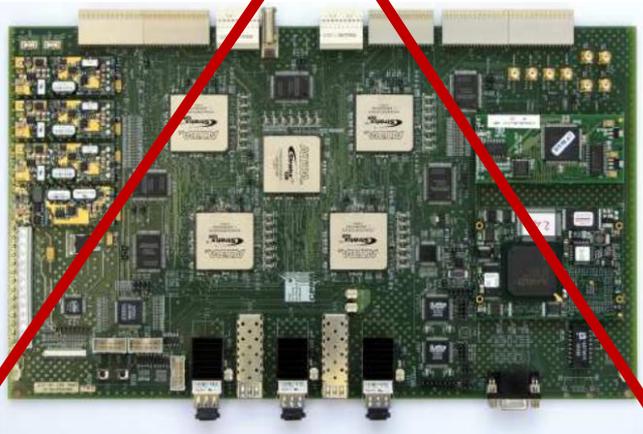
~~Custom electronics (FPGAs),
Informers~~

Removed!

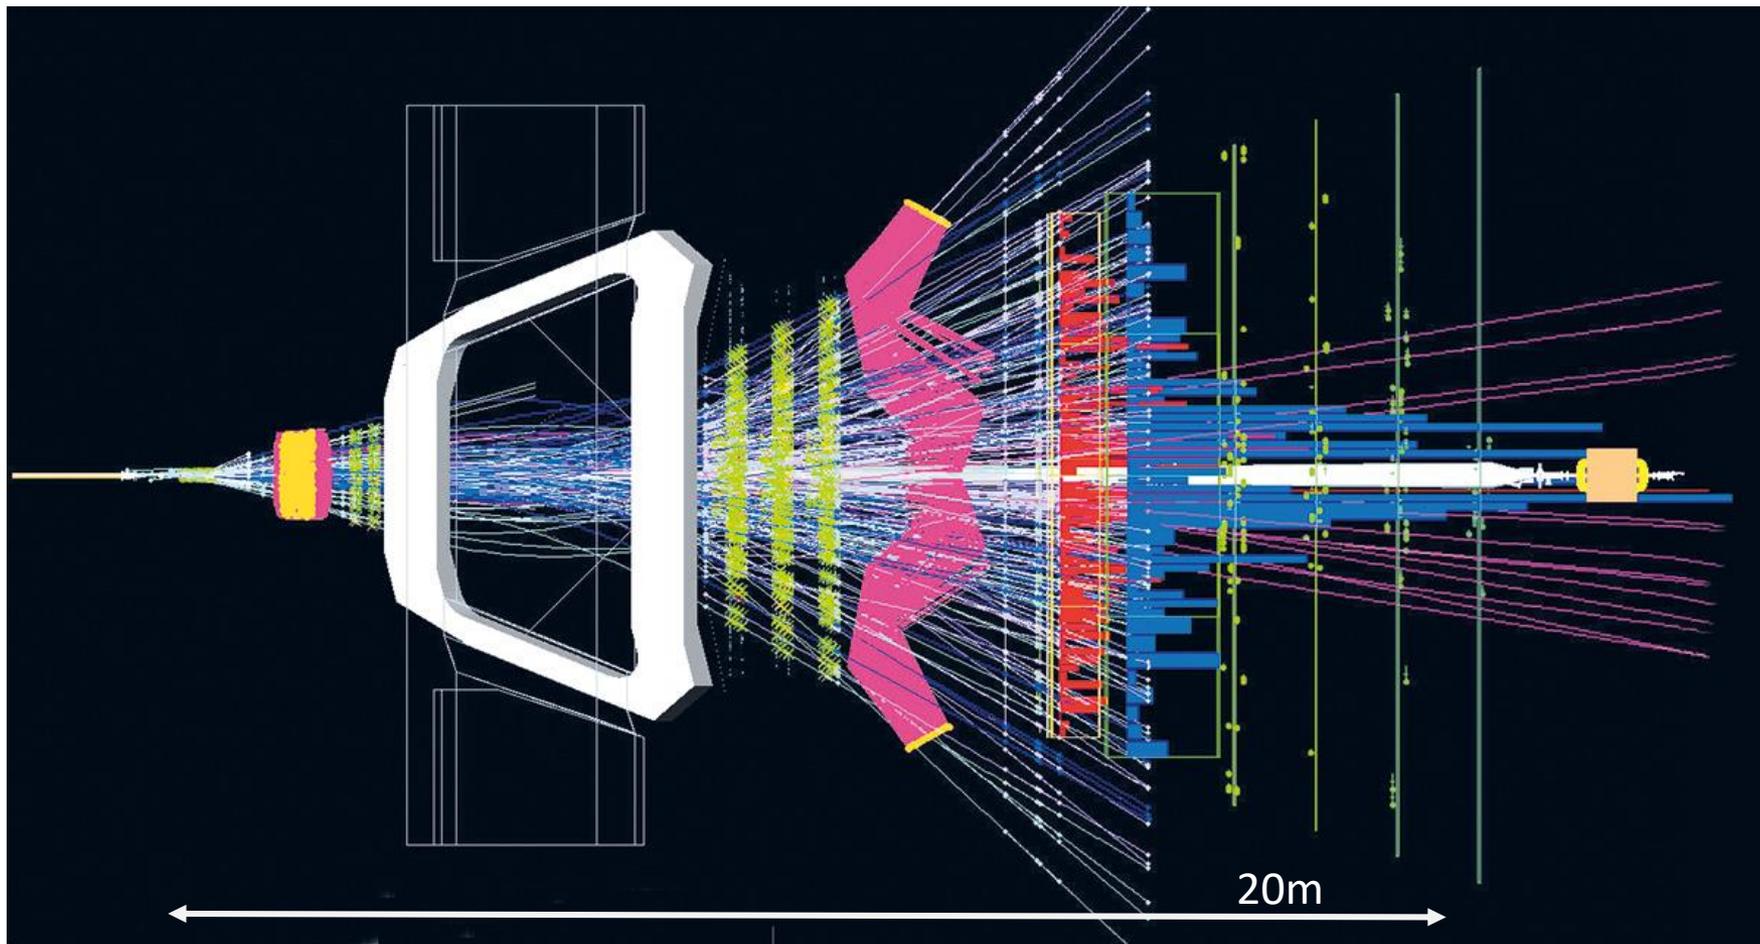
Full software trigger with
detector readout at 30 MHz

farm,
information
to reconstruct the event

from tracking



The new LHCb trigger



	LHCb	ATLAS	CMS	ALICE
$\mathcal{L} [cm^{-2} s^{-1}]$	2×10^{33}	2×10^{34}	2×10^{34}	6×10^{29}
pile-up	5	60	60	1
reconstruction rate	30 MHz	100 kHz	100 kHz	50 kHz
reconstructed tracks/s	1800 M	90 M	90 M	10 M

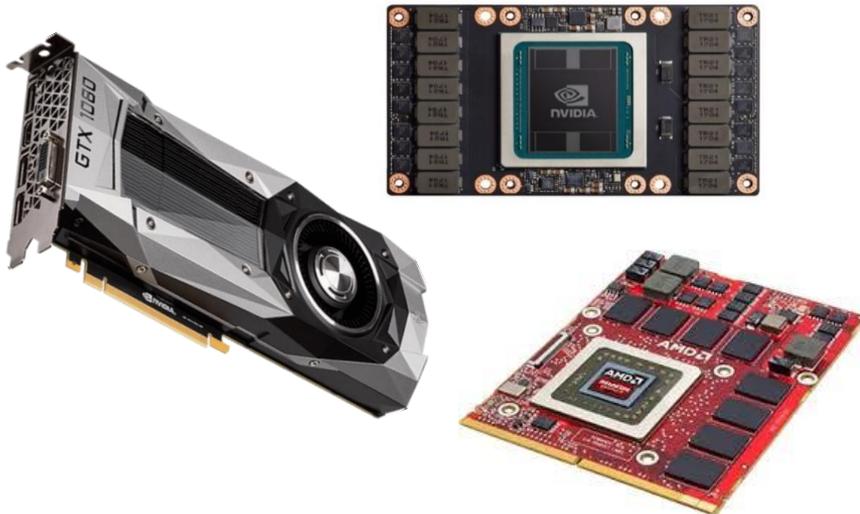
The new LHCb trigger



Exploiting hardware accelerator technologies in event reconstruction:

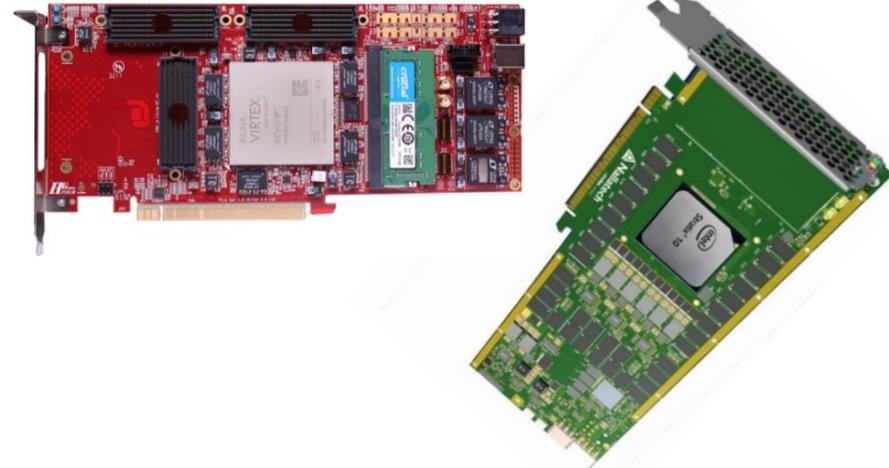
- Use more than one kind of processor or cores to maximize performance or energy efficiency.
- Exploit the high level of parallelism to handle particular tasks.

Graphic Processor Units (GPUs)



- Multicore processors, highly commercial
- High throughput
- Ideal for data-intensive parallelizable applications

Field Programmable Gate Arrays (FPGAs)

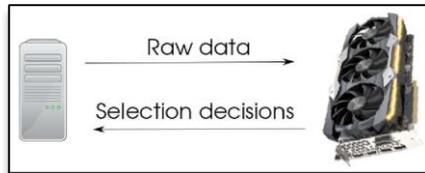


- Programmable and flexible devices
- Low latency
- Low power consumption
- Ideal for compute- and data-intensive workloads

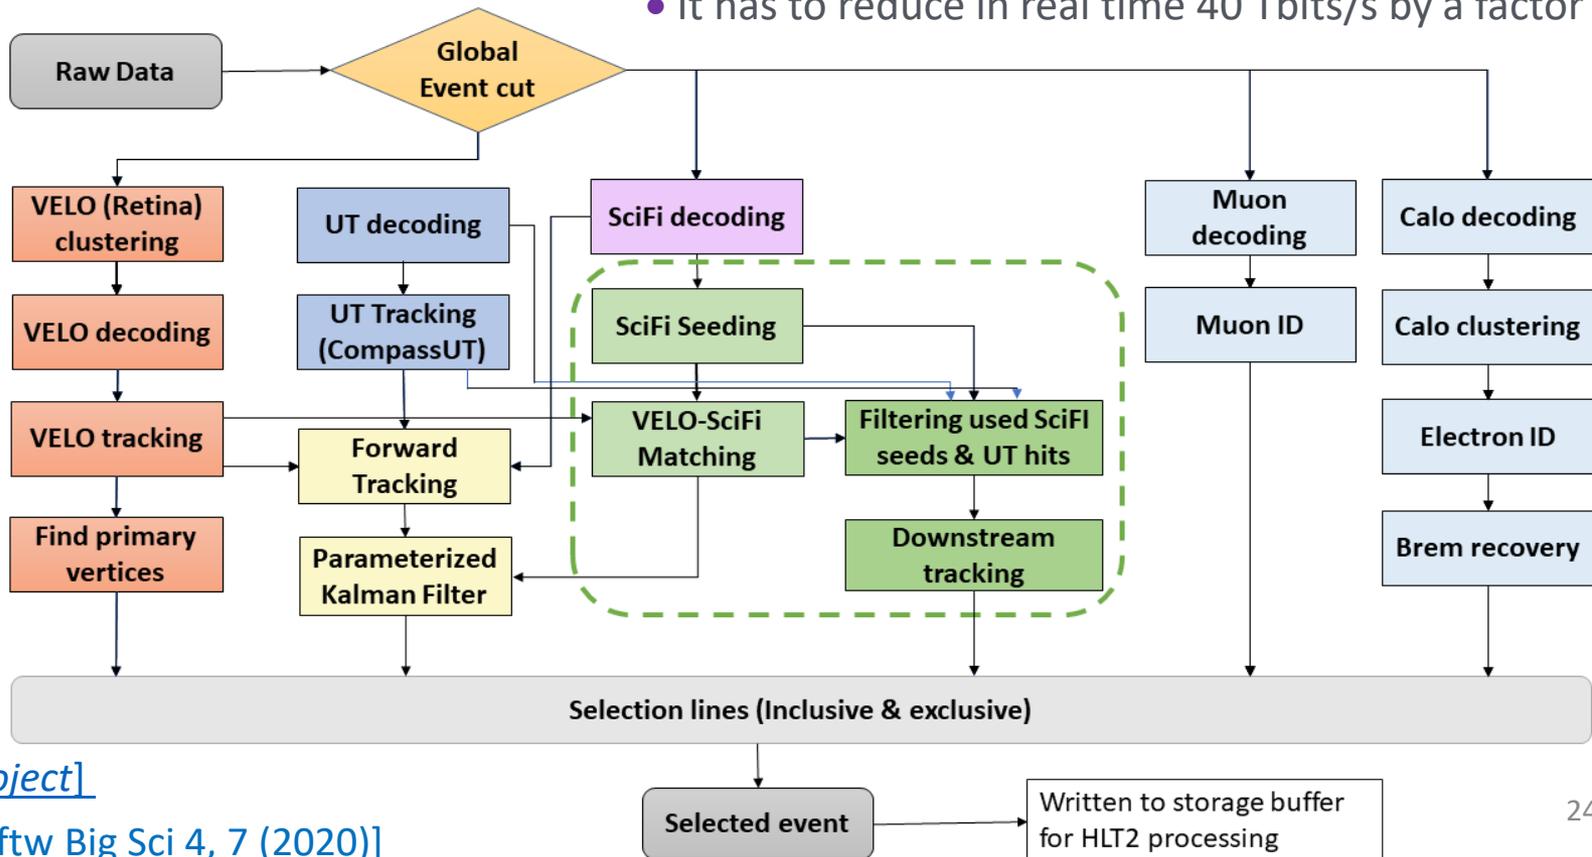
The new LHCb trigger

Allen: the LHCb high-level trigger 1 (HLT1) application on GPUs.

[LHCb-TDR-021] → Fast detector reconstruction in O(500) Nvidia RTX A5000



- Portable: executed on several architectures: CPU, GPU
- Modular: design allows various execution sequences
- Total of approx. 250 algorithms used in data-taking
- It has to reduce in real time 40 Tbits/s by a factor 50



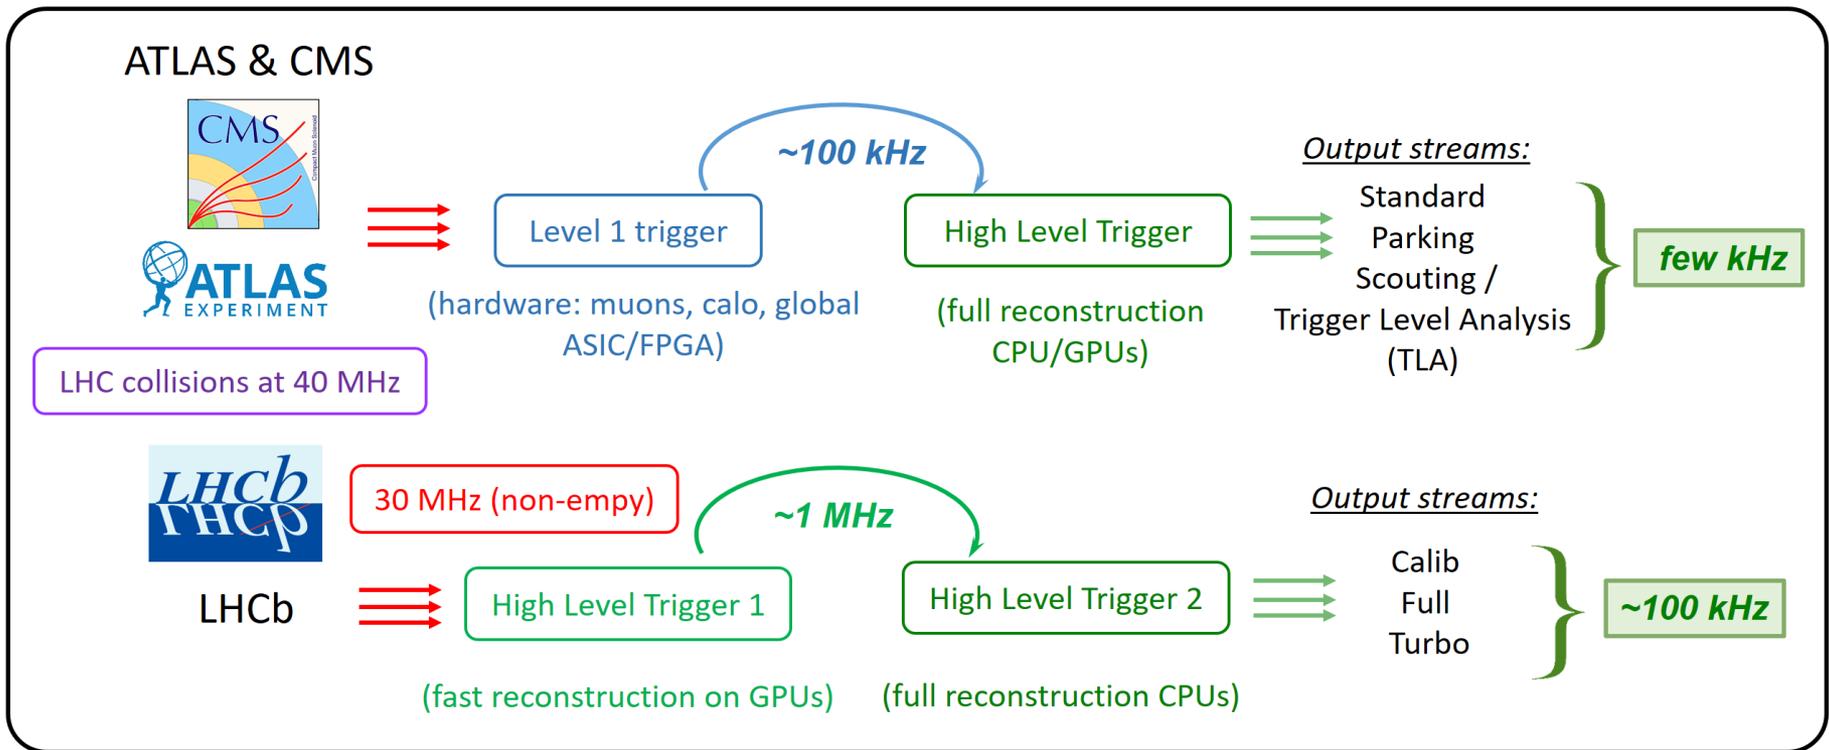
[Allen project]

[Com. Softw Big Sci 4, 7 (2020)]

The new LHCb trigger

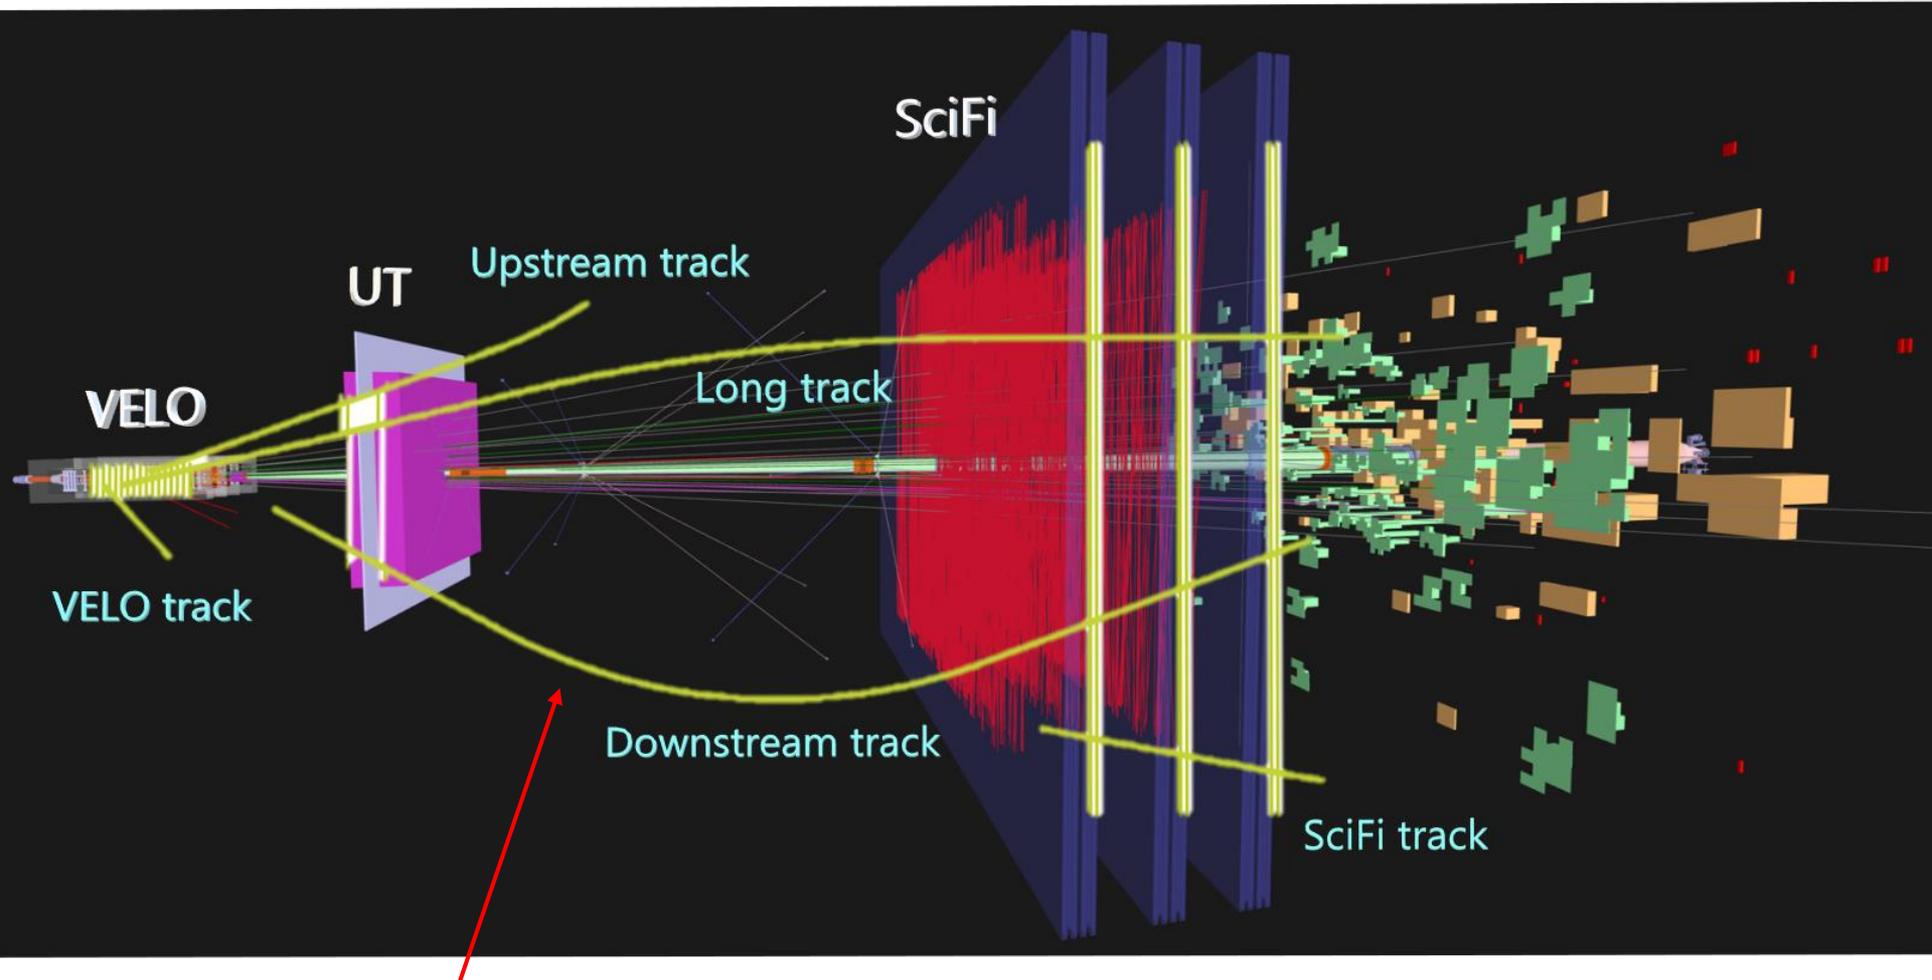
Comparison to other LHC experiments:

Trigger schemes:



Downstream tracks

Track types at LHCb:



Included in HLT1 trigger for 2025 data taking period !

Downstream tracks

- Tracking detectors have very high occupancies:

→ Until now it was impossible to reconstruct and select **downstream** tracks at the first level of the trigger (**very low efficiency for long-lived particles!**).

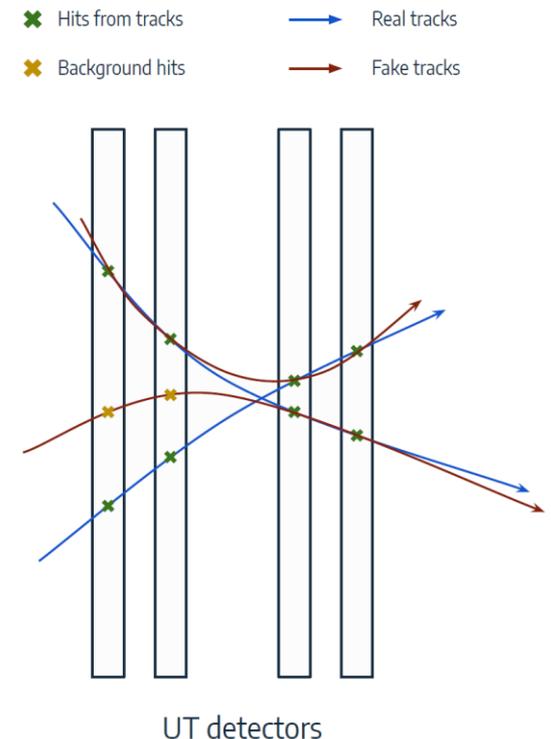
→ The SciFi detector has 3 stations of fibers with an average of 400 hits per layer (**400³ combinations**) in each event.

→ The **UT** detector has 4 layers of strip detectors, and each layer may have up to **1000 hits** per event.

- Tracks reconstructed from incorrect hit combinations are named **fake tracks** (or **ghost tracks**).

Challenges:

- Algorithm design → **very high throughput**
- Fake track suppression → **fast ML approach**



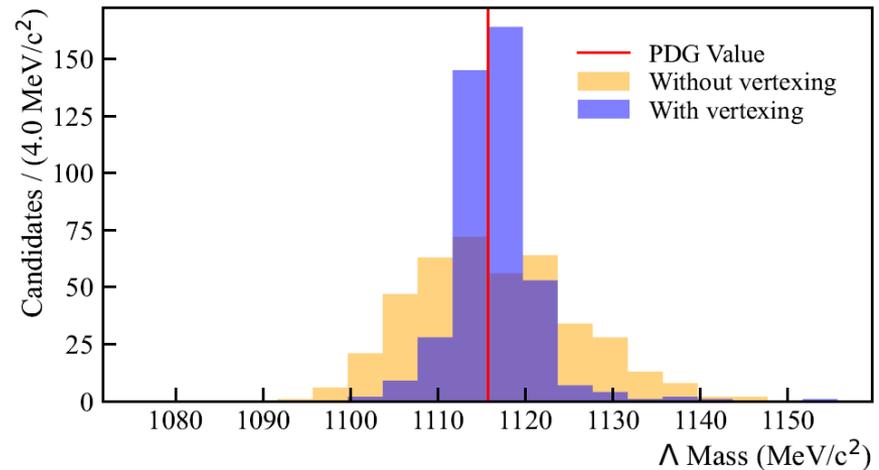
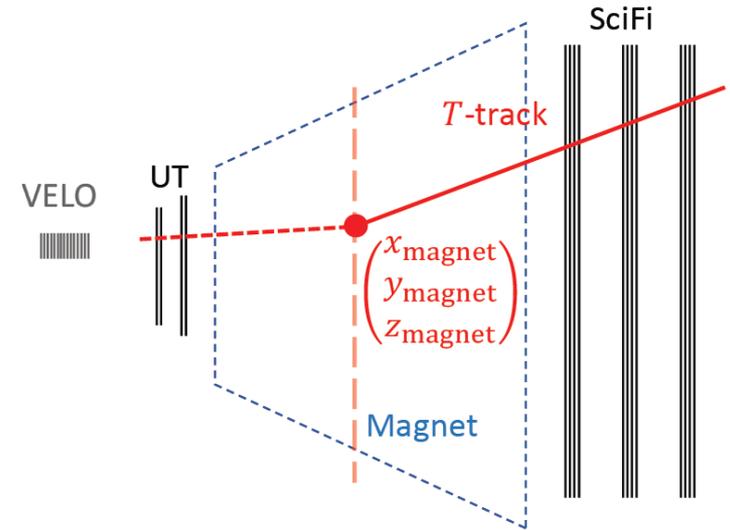
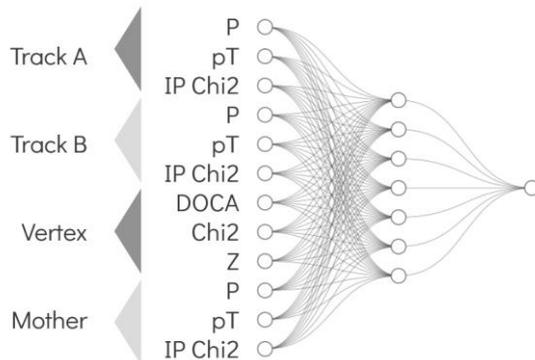
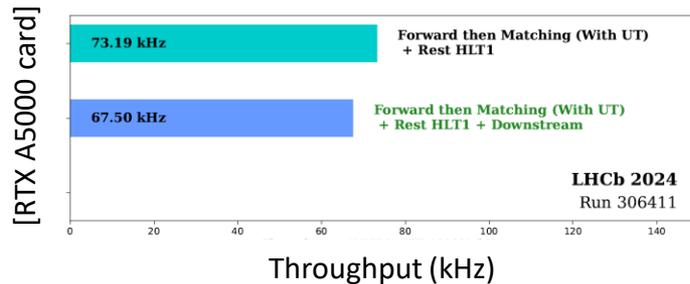
Downstream tracks

The new Downstream algorithm:

Reconstructing and vertexing *downstream* tracks at the first high level trigger (HLT1), taking into account the magnetic field in the UT. These tracks are then paired using a streamlined NN to select the proper vertices.

[\[Comput. Softw. Big Sci. 9 \(2025\) 1, 10\]](#)

Throughput decreases ~ 5 kHz, very fast algorithm!



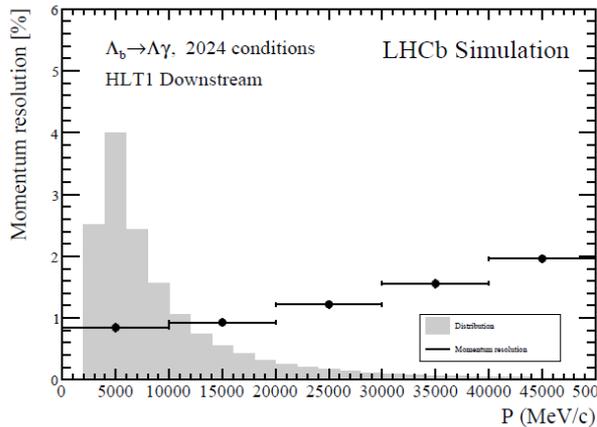
- Bias for Λ decays ~ 1 cm

Downstream tracks

Performance:

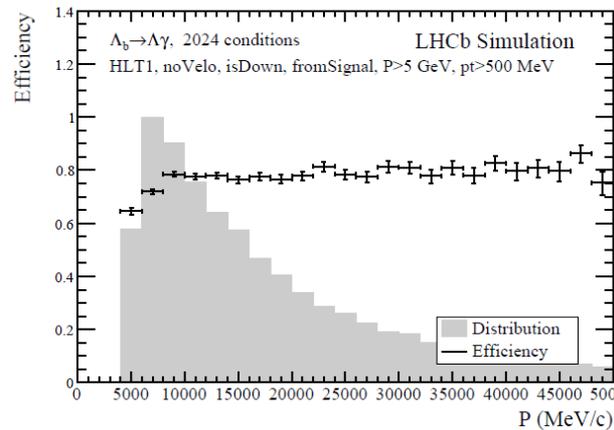
Track momentum resolution:

1-2 %



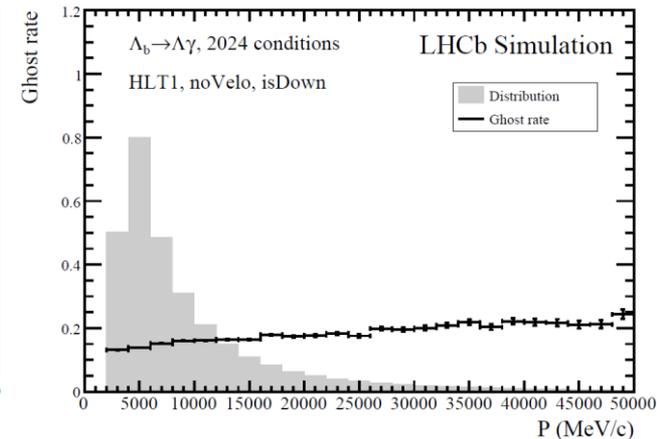
Efficiency:

~ 80 %



Ghost rate:

Less than 20%

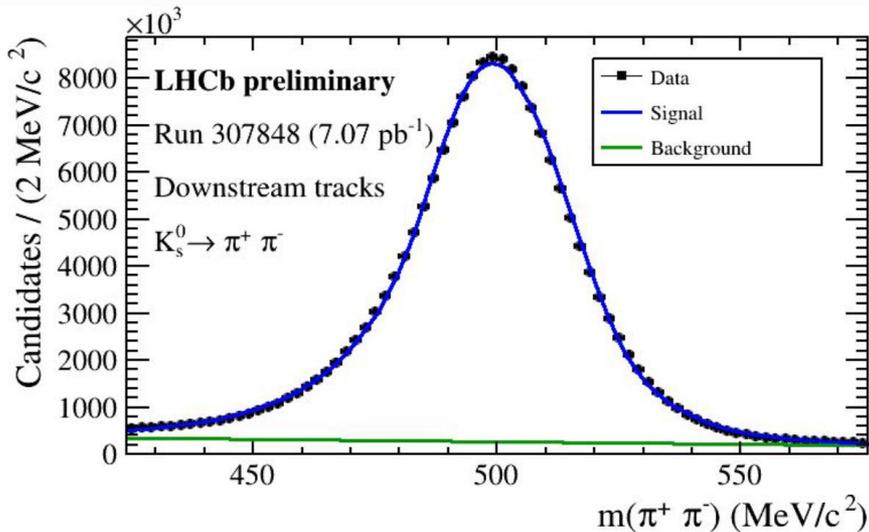


- Independent of the physics channel.
- Efficiency is also flat in other variables such as angular acceptance, track multiplicity, etc..

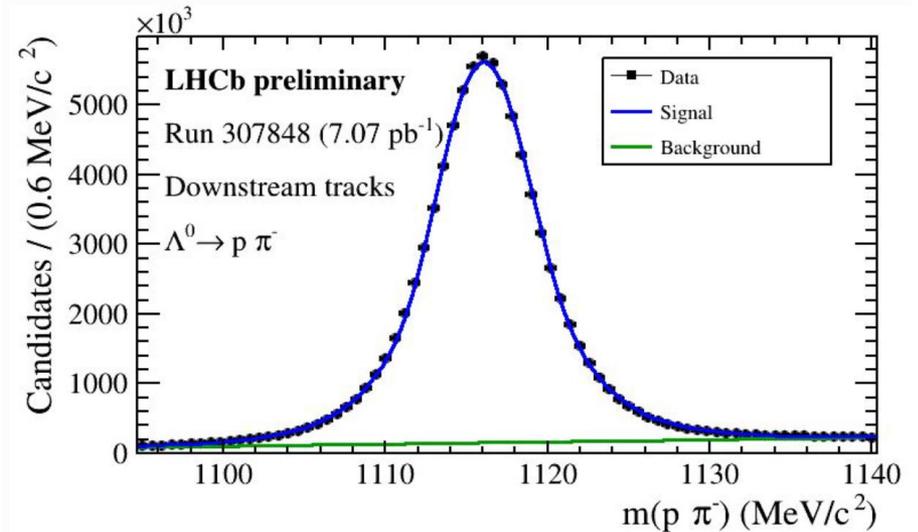
Downstream tracks

Performance:

Mass resolution after pairing two downstream tracks:



$\sigma_M \sim 15$ MeV (K_s)

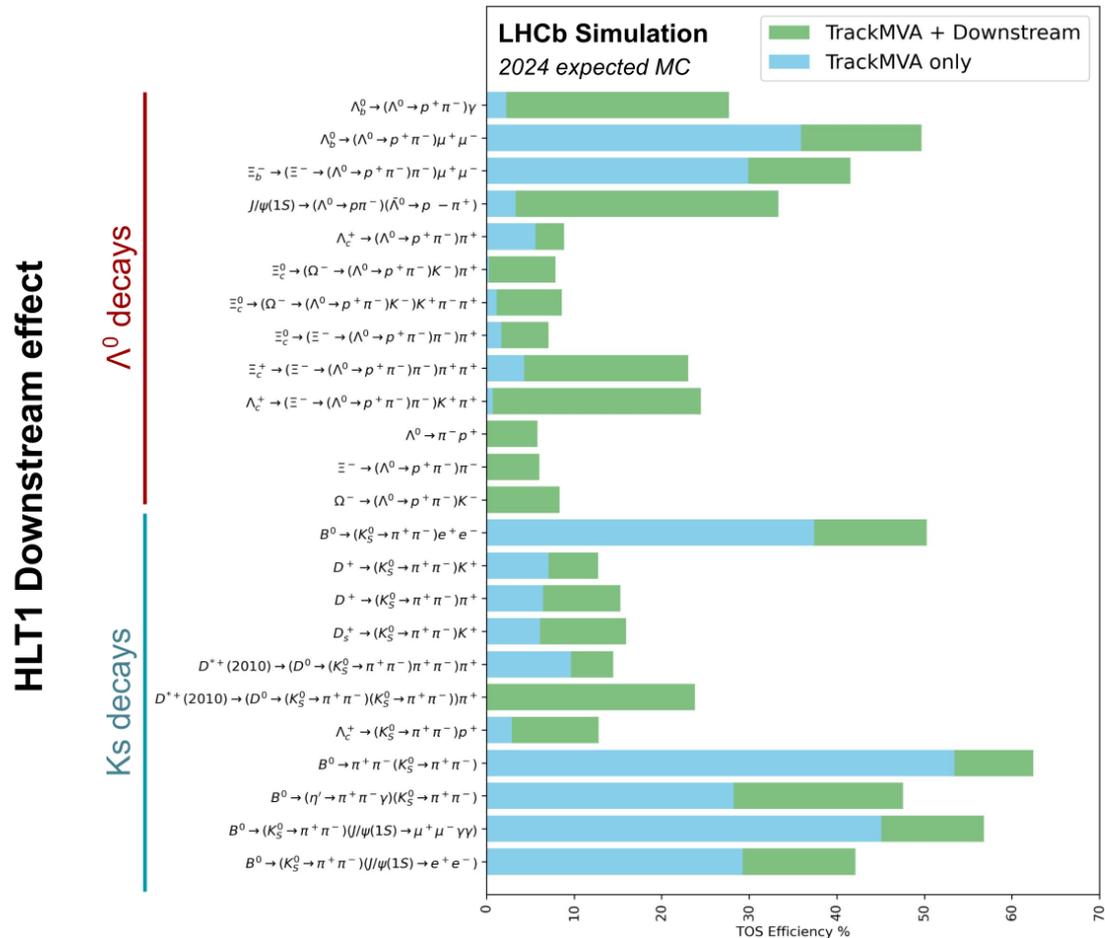


$\sigma_M \sim 4$ MeV (Λ^0)

Enhancing LHCb capabilities

Very high impact on physics with Λ and K_S !

[\[Comput. Softw. Big Sci. 9 \(2025\) 1, 10\]](#)

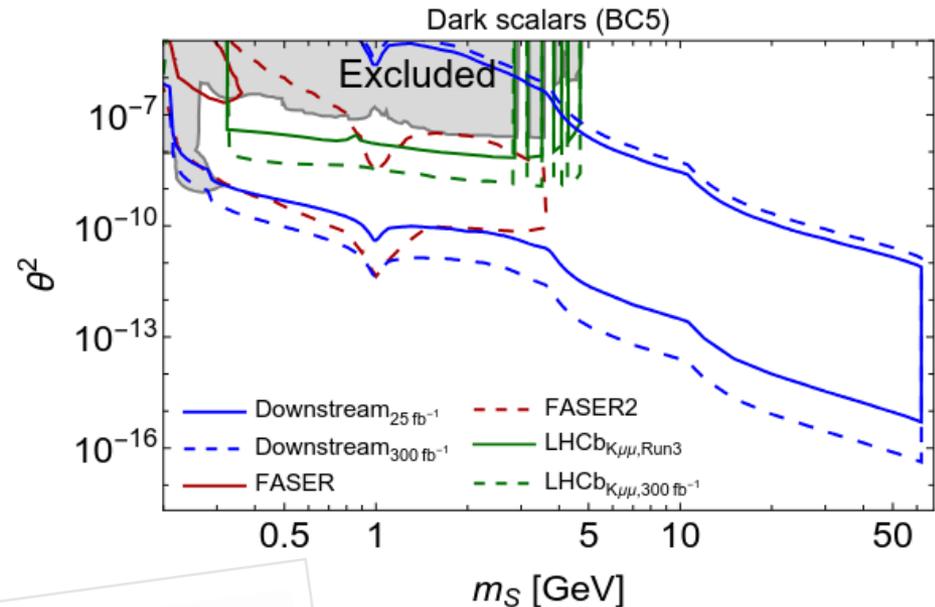
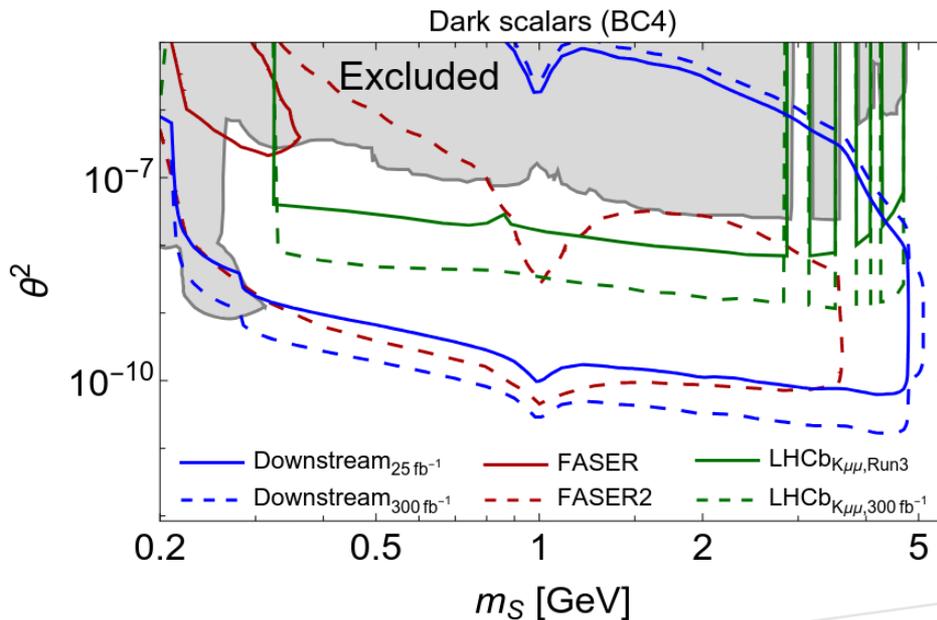


Enhancing LHCb capabilities

Potential of the *Downstream* algorithm for LLP searches:

New scalar boson in the dark sector coming from B or H decays:

[\[Eur.Phys.J.C 84 \(2024\)6, 608\]](https://doi.org/10.1051/epjconf/202466008)



EUROPEAN ORGANIZATION FOR NUCLEAR RESEARCH (CERN)

LHCb-PUB-2025-001
April 1, 2025

Discovery potential of LHCb Upgrade II

Input to the European Particle Physics Strategy Update 2024–26

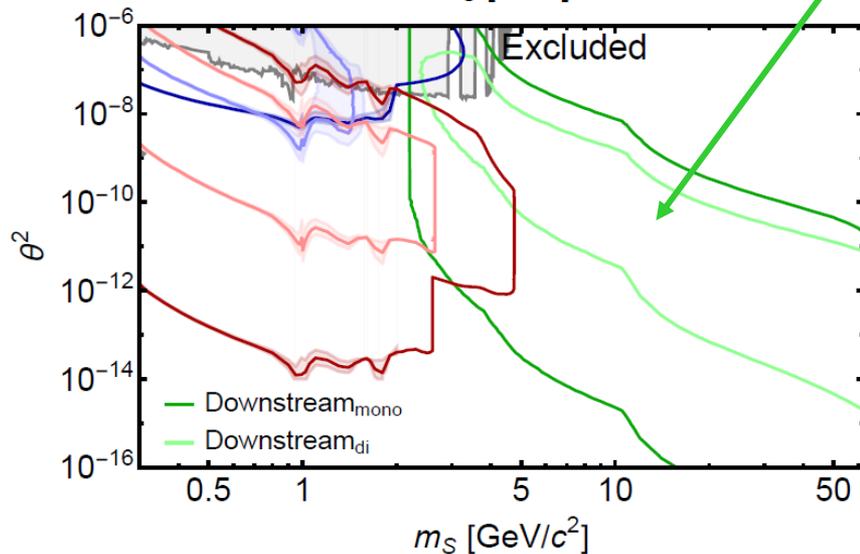
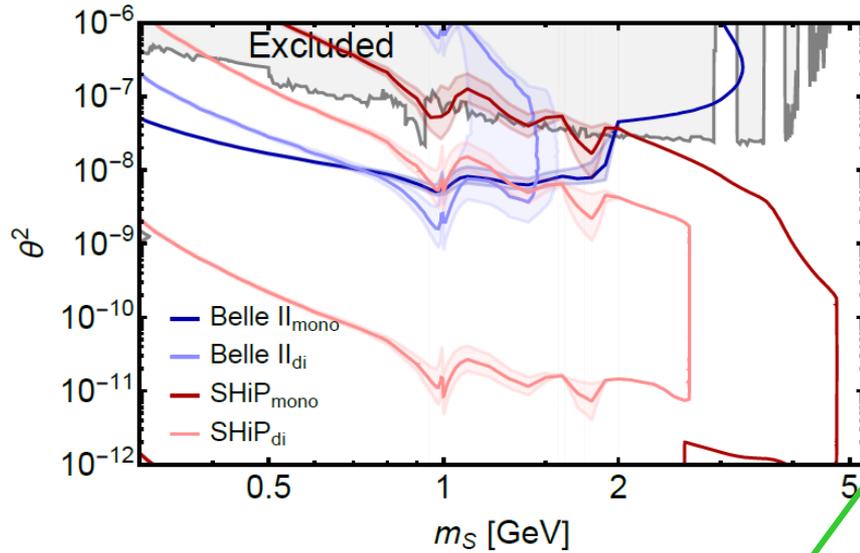
Mar 2025

LHCb

Sensitivity extended
two orders of
magnitude!
(similar for HNL or ALPs)

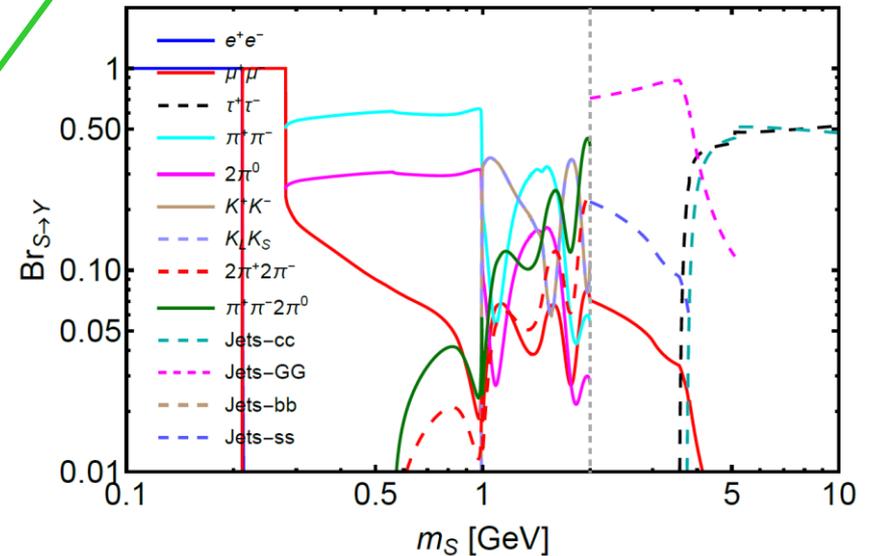
Enhancing LHCb capabilities

Di-decays signatures: B or H decays into a couple of scalars ($B_s \rightarrow SS$, $B \rightarrow SSX$, $H \rightarrow SS$)



Open an unexplored phase-space window for high masses

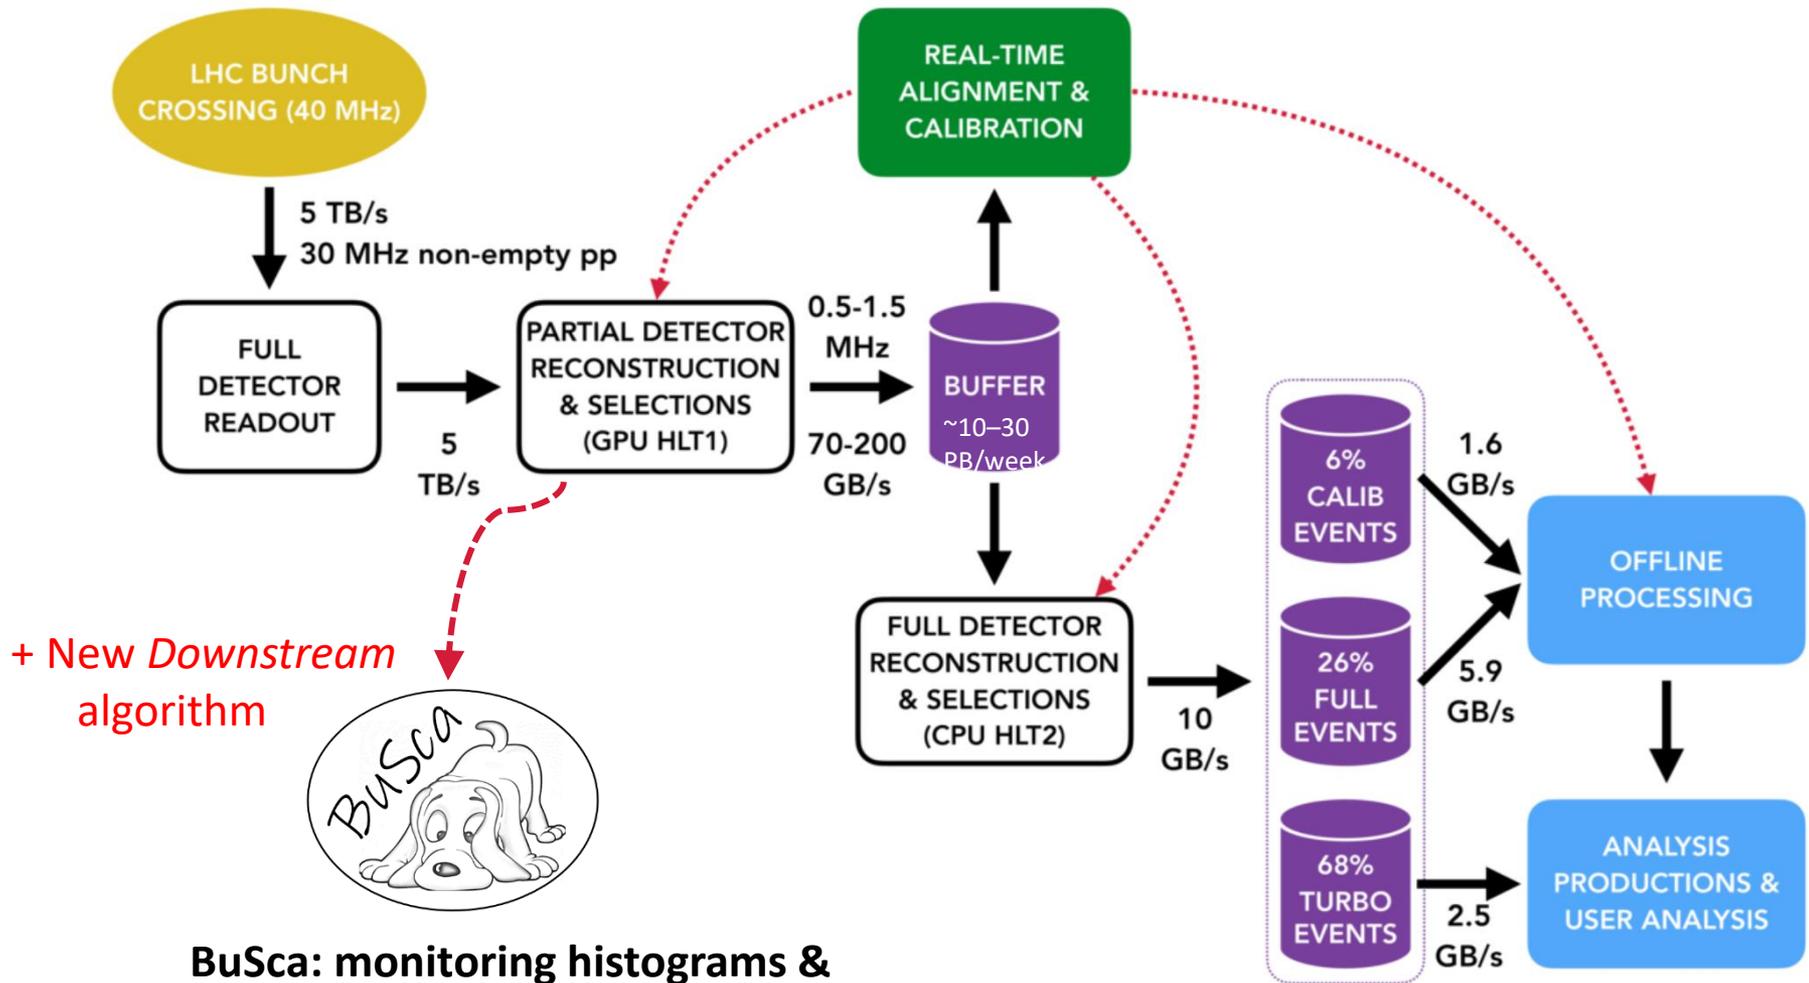
Scalar boson decay rates:



G. Dalla Valle Garcia and M. Ovchinnikov,
[arXiv: 2503.01760 \[hep-ph\]](https://arxiv.org/abs/2503.01760)

BuSca (Buffer Scanner)

The new RTA trigger scheme at LHCb:



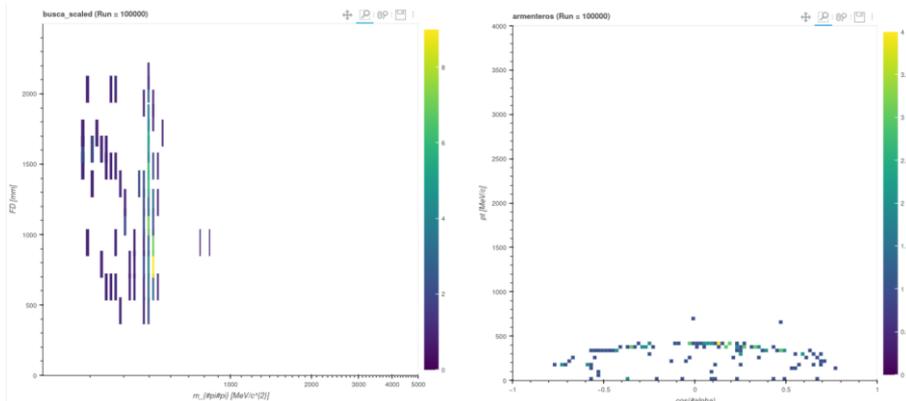
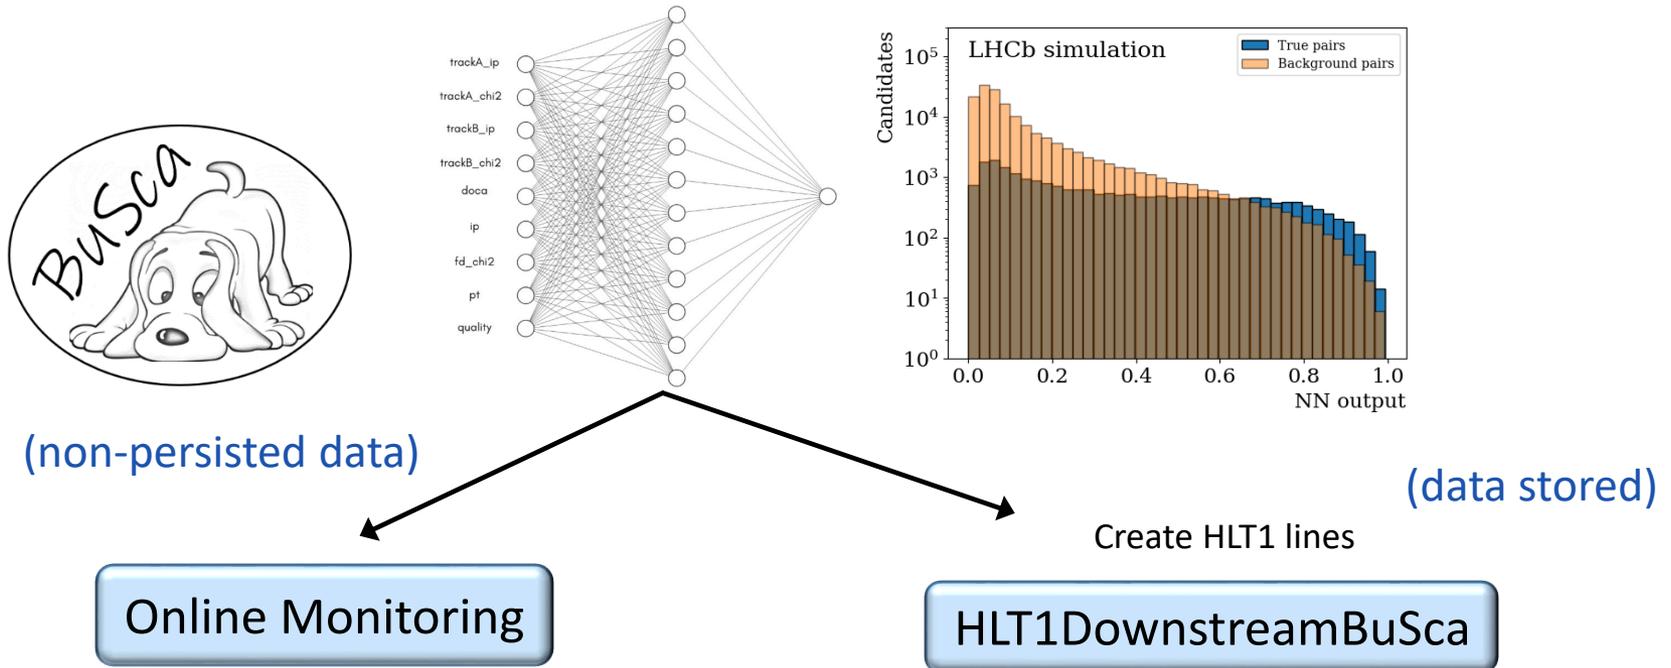
BuSca: monitoring histograms & create HLT1 lines for BSM searches

[[EPJ Web Conf. 337 \(2025\) 01161](#)]

[[V. Kholoimov MSc](#)]

BuSca (Buffer Scanner)

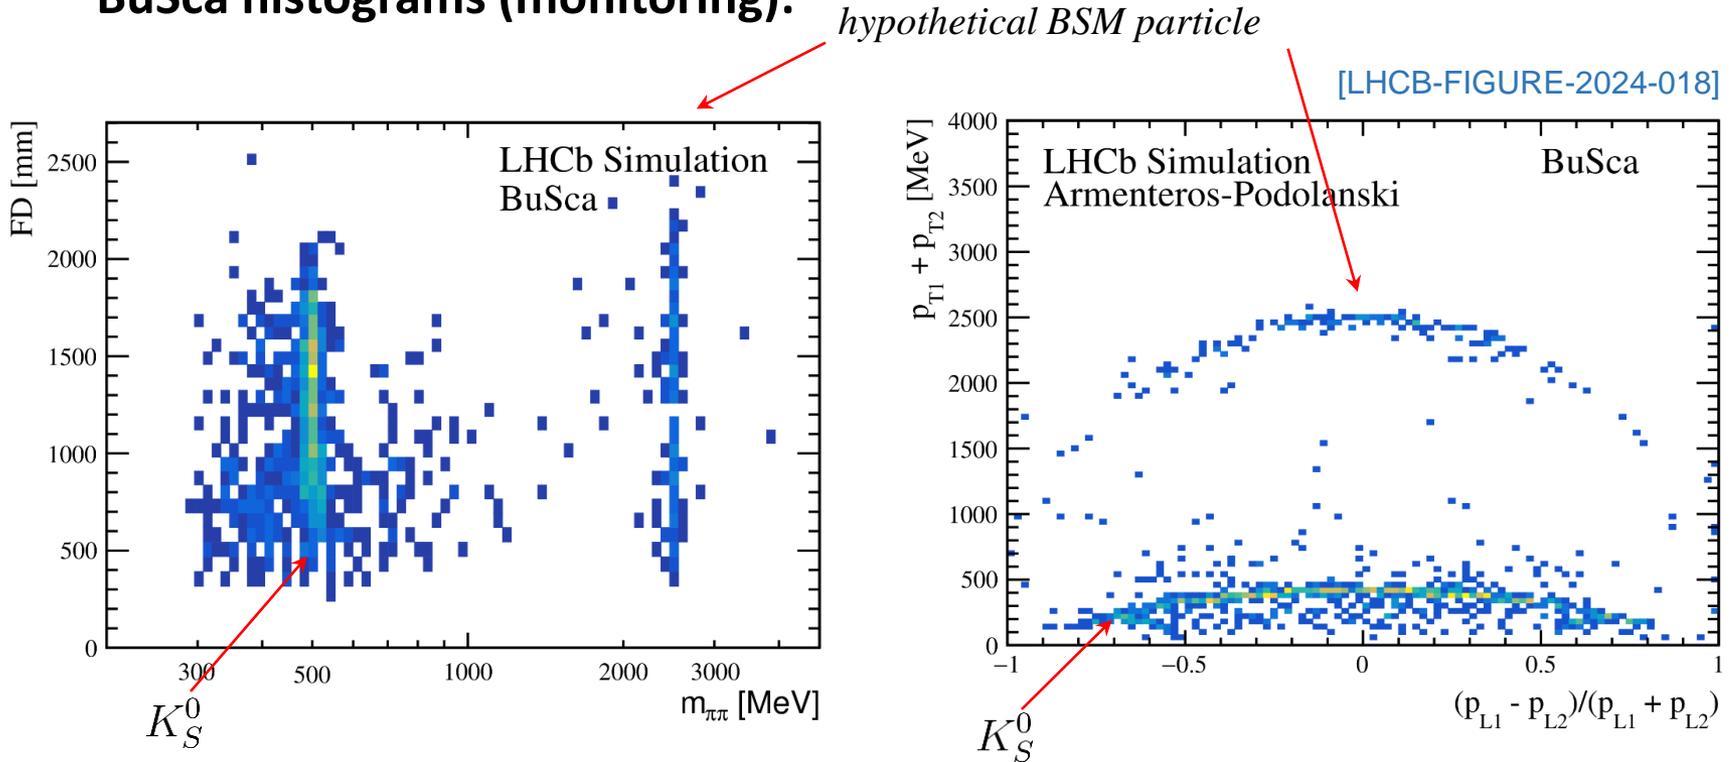
High performance NN to select good track candidates at 30 MHz:



- HLT1_BuSca_Mass
- HLT1_BuSca_PID
- HLT1_BuSca_Multitrack
- HLT1DownstreamBuSca_prescaled (bkg studies)

BuSca (Buffer Scanner)

BuSca histograms (monitoring):



Data is stored in two 2D histograms: **di-pion hypothesis mass vs flight-distance** of the particle and **Armenteros-Podolanski plot*** (now also simplified helicity).

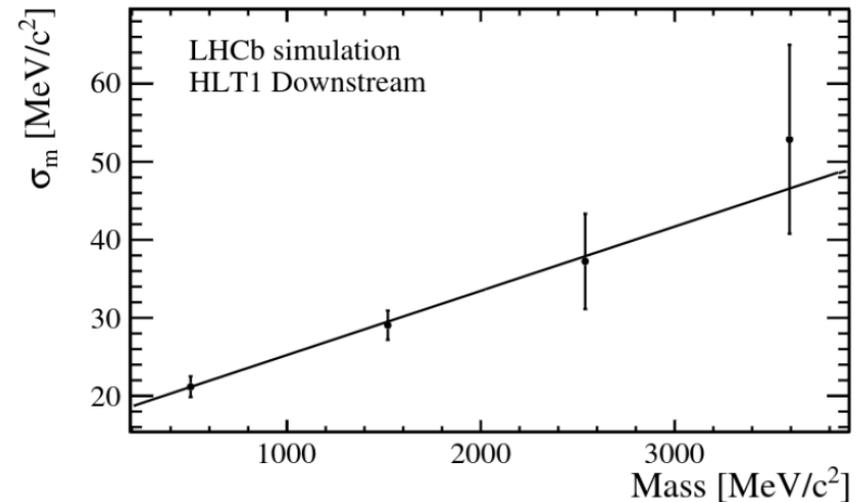
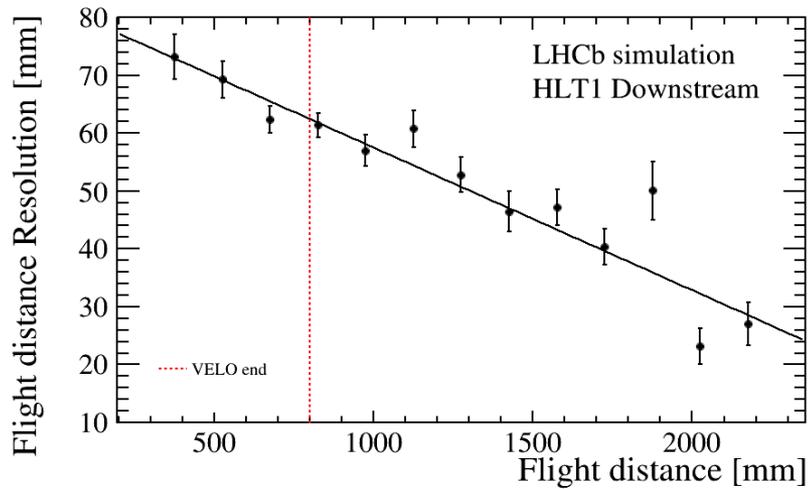
*Transverse momentum versus the longitudinal momentum asymmetry.

BuSca (Buffer Scanner)

BuSca design

Mass and Flight Distance binning schemes:

[LHCB-FIGURE-2024-036]



Binning is adapted according to the expected mass and FD resolution:

$$\sigma_{FD} = 80 - FD \cdot 0.02 \quad [mm]$$

$$\sigma_m = m \cdot 0.02 \quad [MeV/c^2]$$

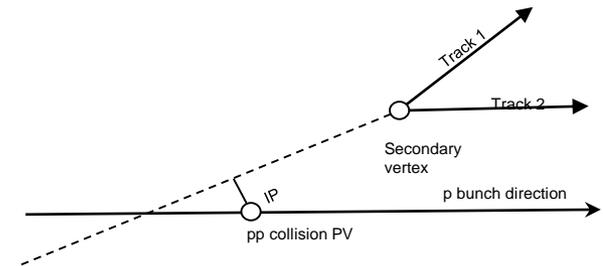
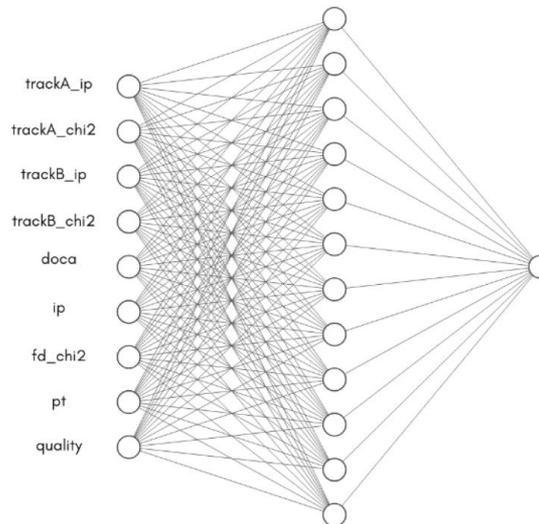
BuSca (Buffer Scanner)

BuSca selection algorithm

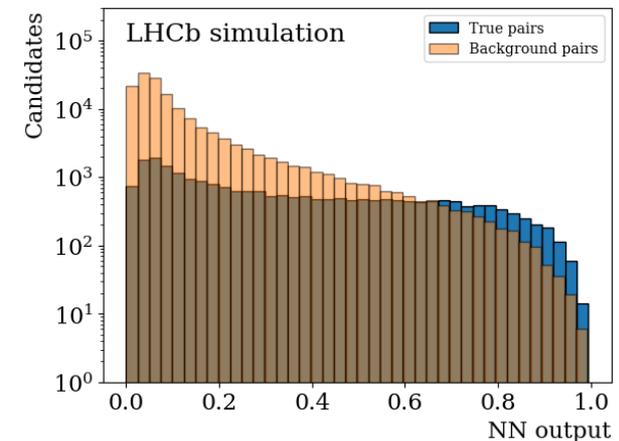
Neural network, trained to select true pairs based on their reconstruction quality

- 9 inputs
- 12 nodes
- $[0, 1]$ output

Trained on MinBias using MC information.



[LHCb-FIGURE-2024-036]

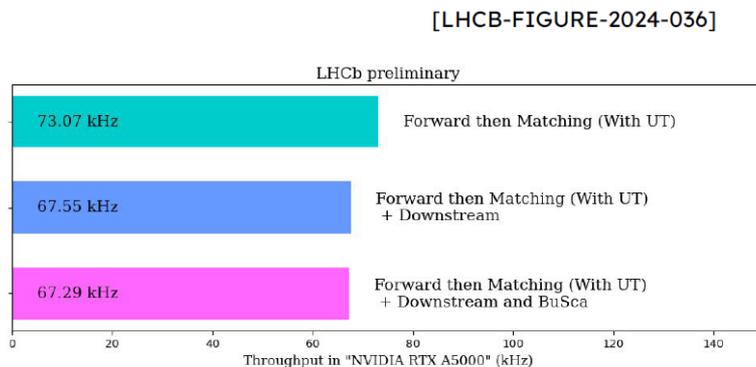


New version: 8 inputs (no kinematic info) and 32 hidden nodes.

BuSca (Buffer Scanner)

Performance:

Throughput (# events/s)



Global throughput impact:

0.38%

HLT1 bandwidth

Three trigger lines are developed:

- diMuon line 0.11 kHz
- diElectron line 0.28 kHz
- diHadronline 0.16 kHz

HLT1 Bandwidth impact:

0.06%



BuSca (Buffer Scanner)

The BuSca framework:

- It allows to monitor LLPs (*downstream track pairs*) at 30 MHz!
 - Mass and flight distance (FD) info: allows the study of material interaction (crucial for bkg. understanding).
 - Armenteros-Podolanski info: PID orientation (important for models).
- ★ Phase-space corners never explored before, it is just based in high quality track pairs.
- ★ Full model independent (not even needed B production, higher mass range for new particles!).
- It allows to create HLT1 lines with specific conditions (PID, mass region, multi-track vertexing, prescaled...).
- It provides a selection at the 2nd High Level Trigger (HLT2) based on BuSca HLT1 lines, allowing the detailed analysis of the BuSca selected data.
- It allows to perform exclusion limits, and it can be adapted for any model.

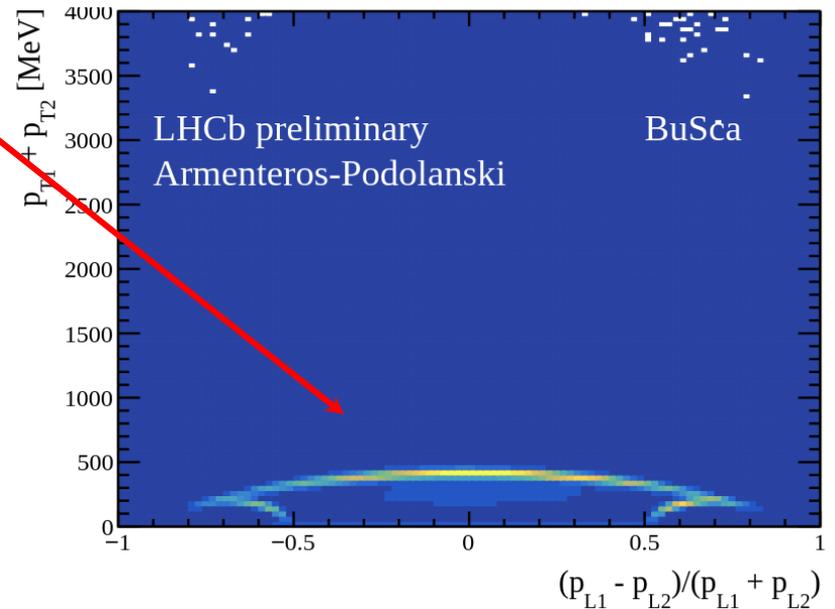
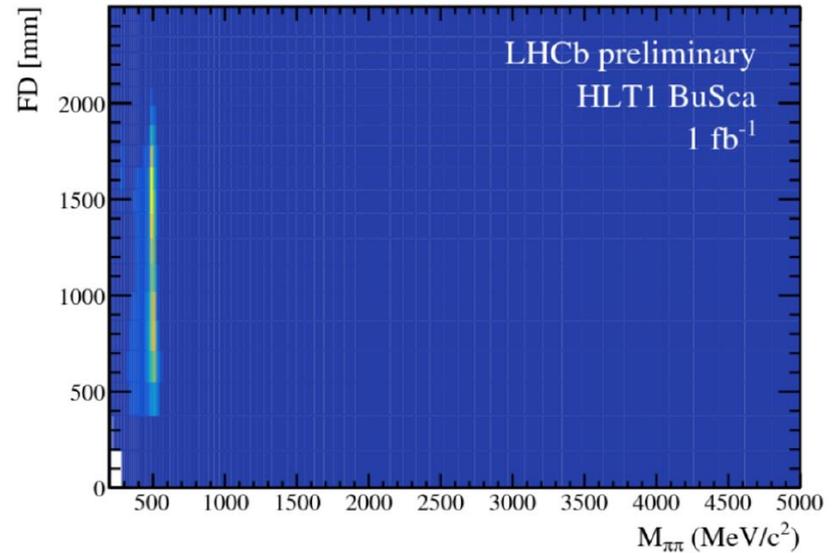
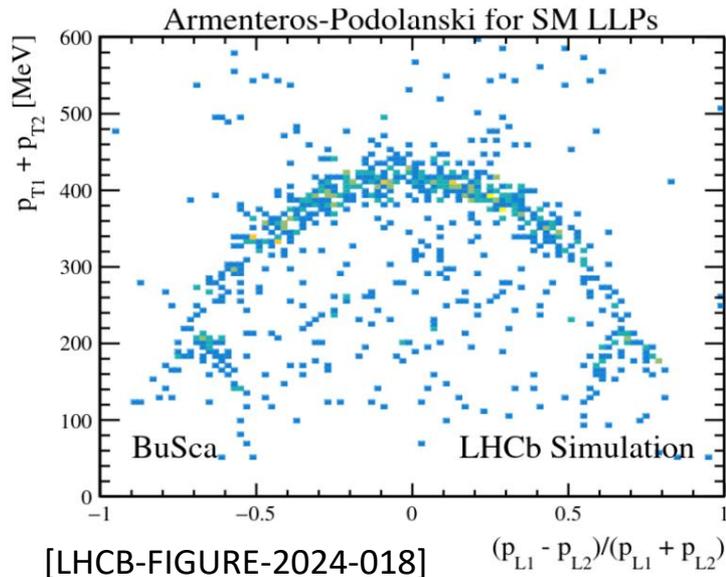
BuSca (Buffer Scanner)

[V. Kholoimov MSc]

First downstream data!

(2 weeks in October '24):

K_S and Λ from pp collisions



$$p_i^L = (p_{ix} \cdot \mathbf{p}_x + p_{iy} \cdot \mathbf{p}_y + p_{iz} \cdot \mathbf{p}_z) / \mathbf{p}$$

$$p_{ix}^T = (p_{iy} \cdot \mathbf{p}_z - p_{iz} \cdot \mathbf{p}_y)$$

$$p_{iy}^T = (p_{iz} \cdot \mathbf{p}_x - p_{ix} \cdot \mathbf{p}_z)$$

$$p_{iz}^T = (p_{ix} \cdot \mathbf{p}_y - p_{iy} \cdot \mathbf{p}_x)$$

$$p_i^T = ((p_{ix}^T)^2 + (p_{iy}^T)^2 + (p_{iz}^T)^2) / \mathbf{p}$$

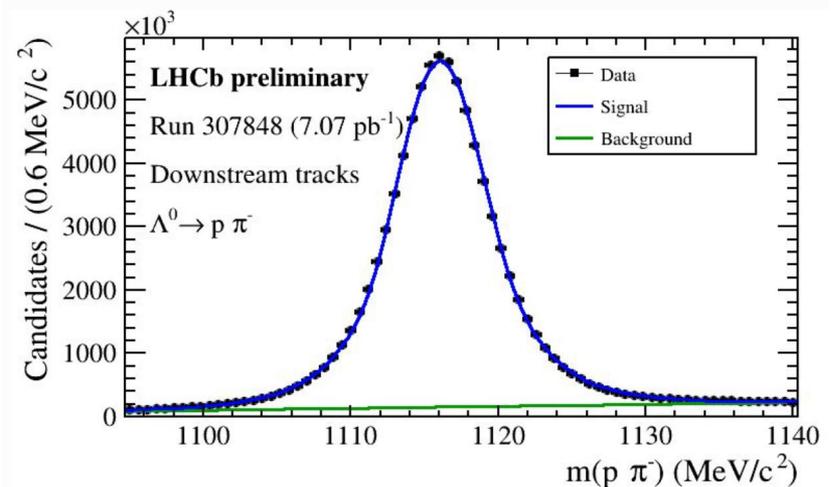
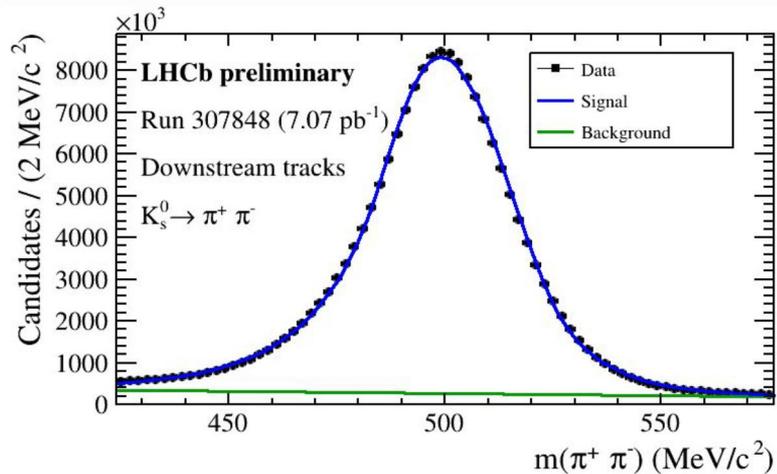
BuSca (Buffer Scanner)

Background events (key for searches!):

1- Strange candidates:

SM particles with large lifetimes (K_s , Λ)

→ *Mass region can be vetoed*

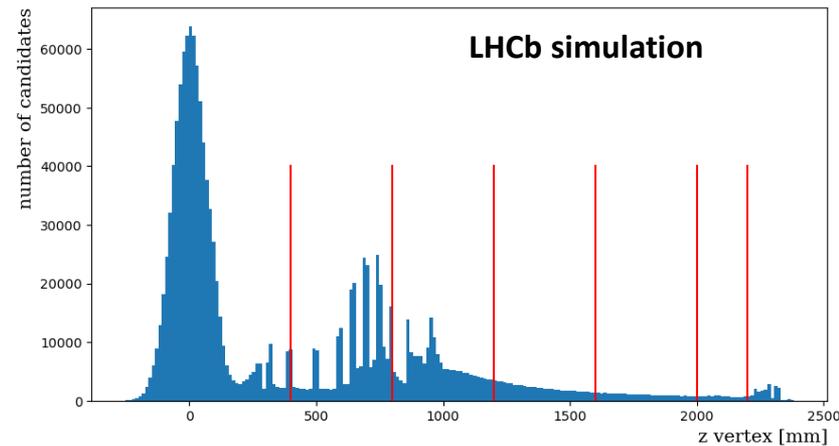
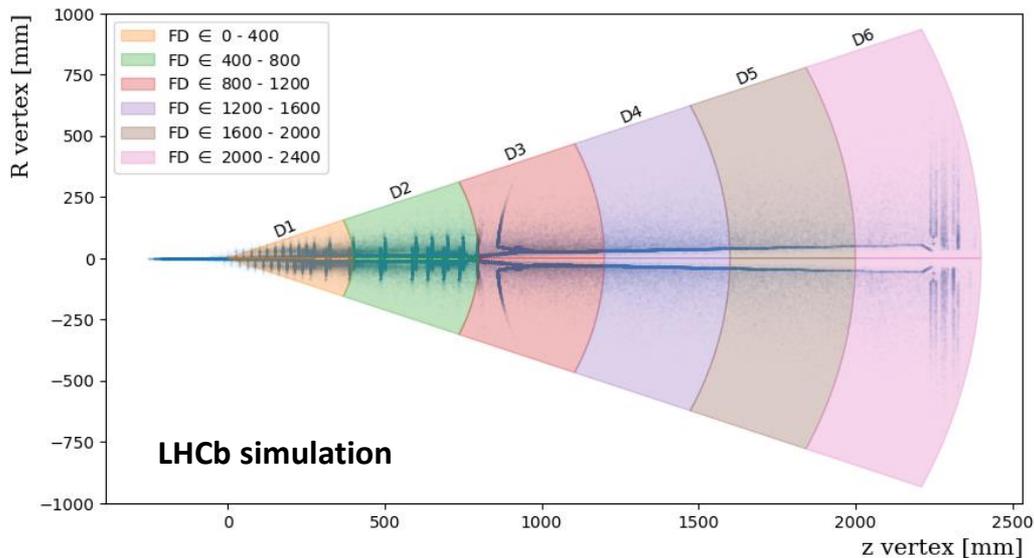


- No SM candidates for masses above 1.2 GeV.
- No other two-prong LLP decays in the SM.

BuSca (Buffer Scanner)

2- Hadronic resonances

► J/ψ , ψ , $\phi(1020)$, $\psi(2S)$, $\psi(3770)$, $\psi(4160)$ decay in the VELO detector ($< 1\text{m}$) and do not contribute to the BuSca background.

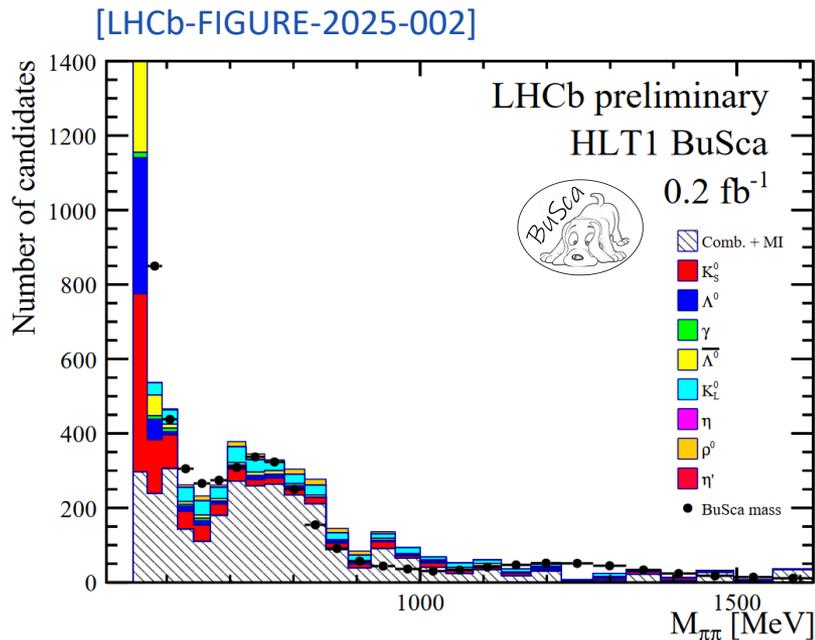


► Light resonance (ρ , f_2 , η , ω ...), can be created from material interaction.

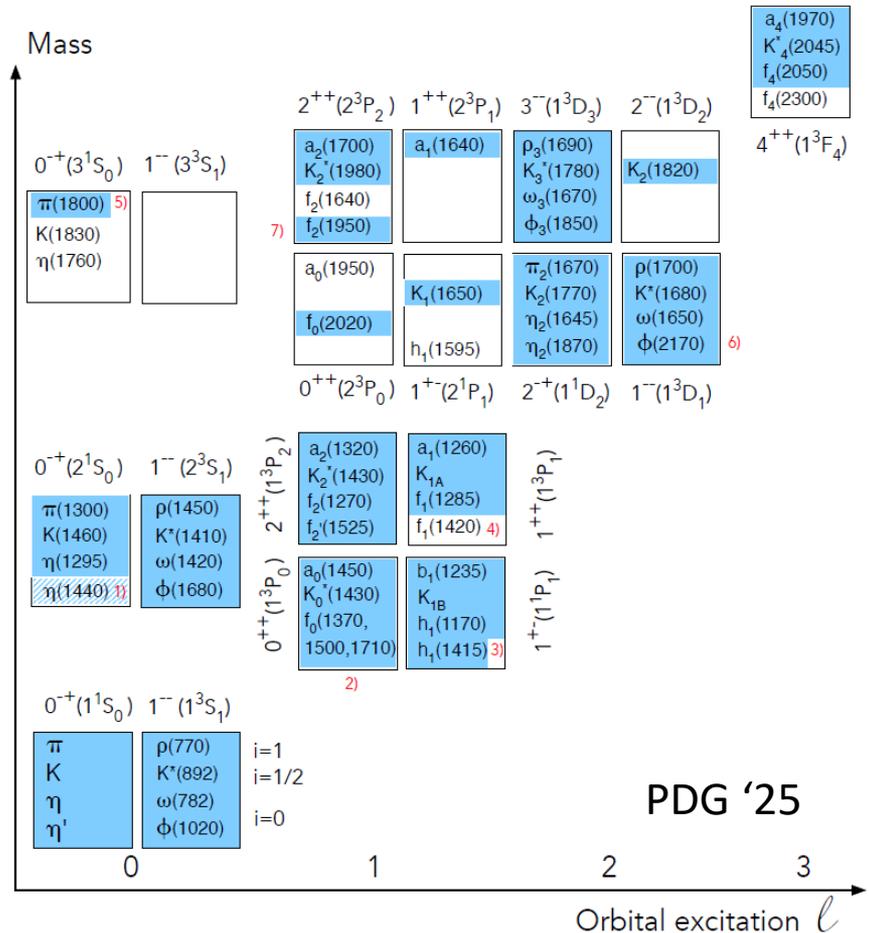


BuSca (Buffer Scanner)

- From simulated events without FD cut (MC Minimum Bias, 10M)
- Data in the region 800-1300 mm



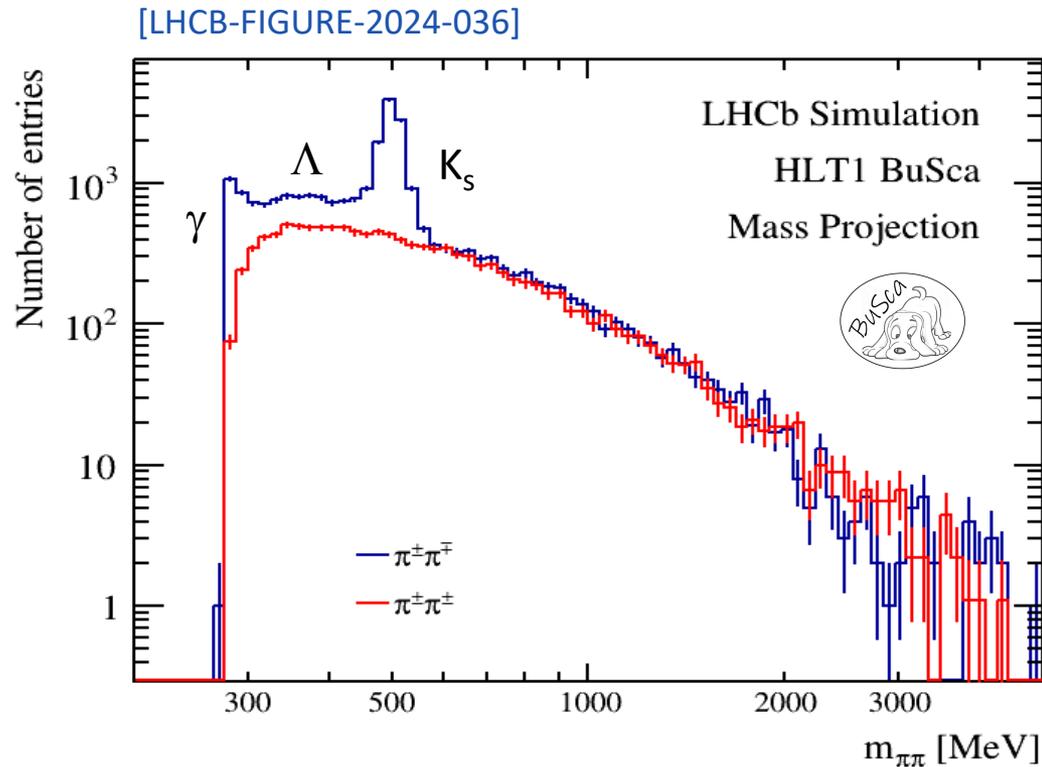
Large contribution of hadronic resonances below 2 GeV: it may be not suitable for BSM searches in hadronic channels.



BuSca (Buffer Scanner)

3- Combinatorial background

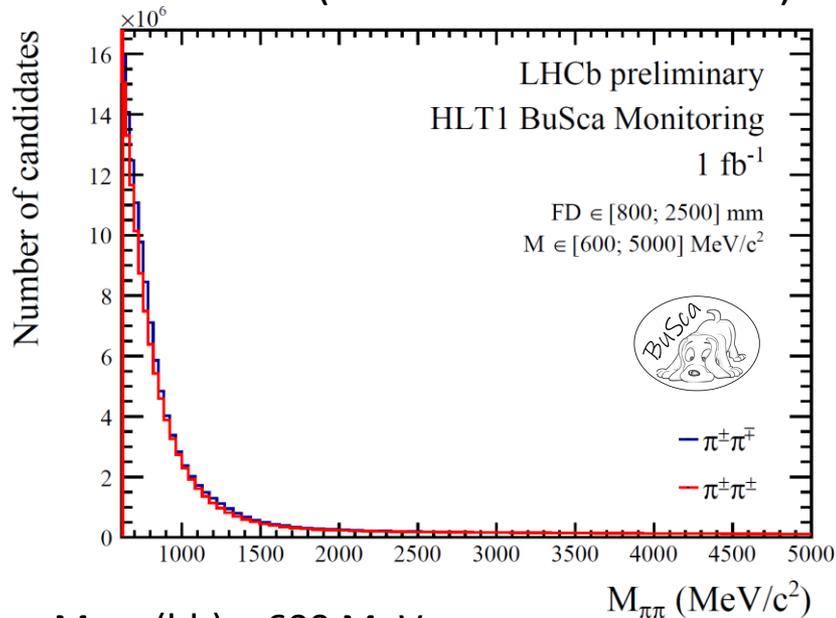
- ▶ Random combinations of two tracks:
It can be controlled using tracks with the same sign.



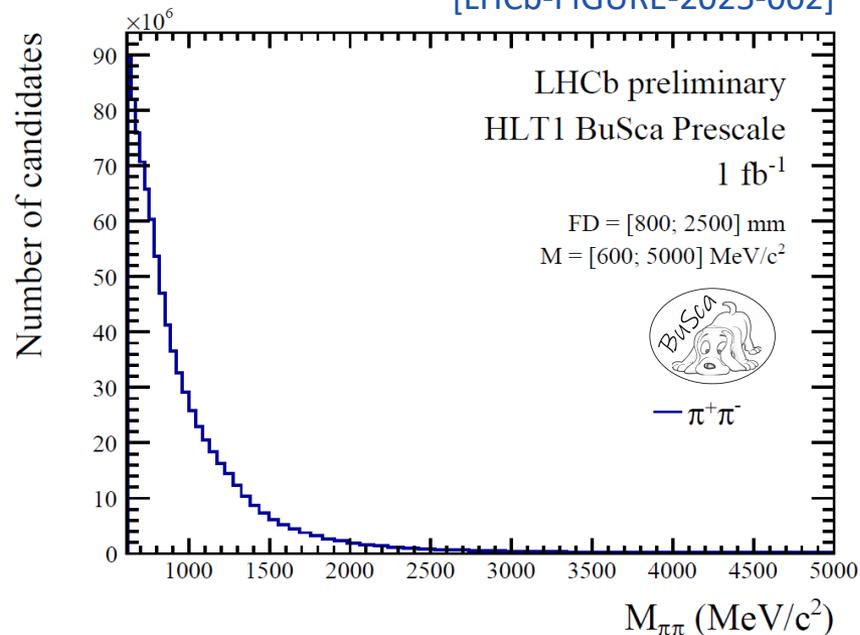
- ▶ Combinatorial background decreases with mass.

BuSca (Buffer Scanner)

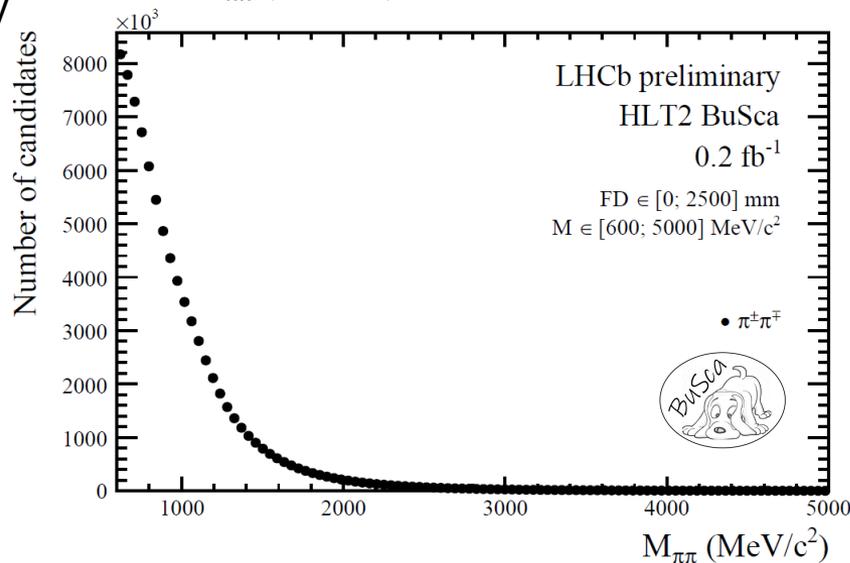
BuSca data (2 weeks in October '24):



[LHCb-FIGURE-2025-002]



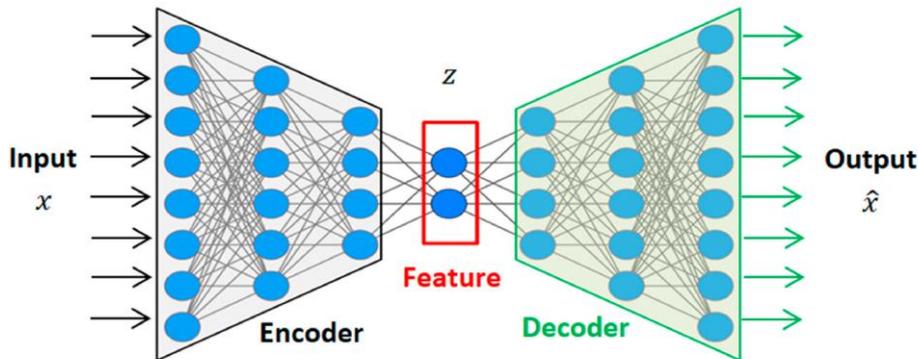
Mass (hh) > 600 MeV



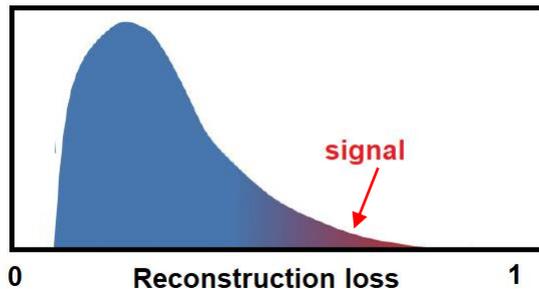
BuSca (Buffer Scanner)

Aiming to suppress the combinatorial background using an autoencoder (or a NN) trained on same-sign track-pairs data.

UNDER DEVELOPMENT



The difference between the original and reconstructed data can be used as an effective anomaly score in monitoring.



Trying to define regions where new physics is more probable: open HLT1 lines to collect and analyse these data.

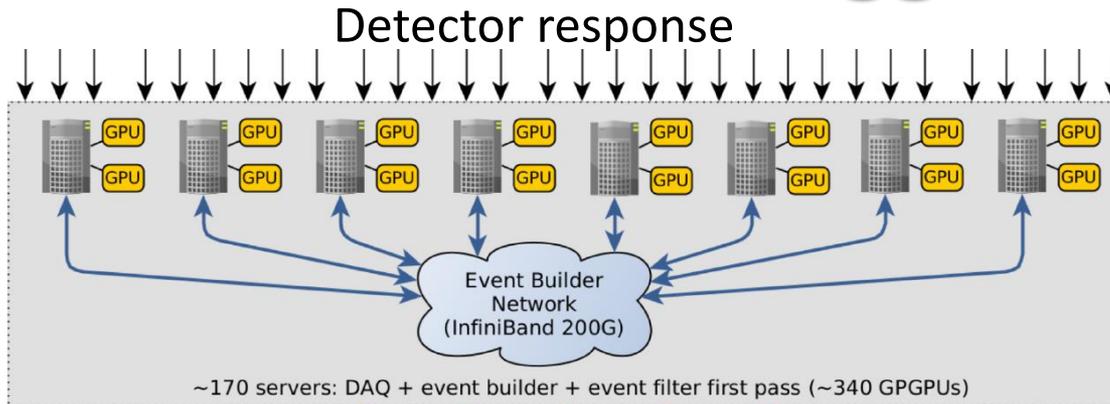
Conclusions & prospects

- **BuSca** is a LHCb model independent LLPs search algorithm at HLT1 level in real time over **30 MHz** of data.
- Based on a new performant algorithm (*downstream*) and possible thanks to the new trigger scheme of LHCb based on GPUs.
- It allows to guide HLT1 & HLT2 trigger lines in **new phase space corners**.
- Already tested with 1fb^{-1} data-taking in October 2024: background studies.
- A lot of developments ongoing (new features to be included!)
- Searching for new observables and combinations of them.
- Stay tuned, **11.8fb^{-1} 2025 data** already taken with BuSca!

Thanks!

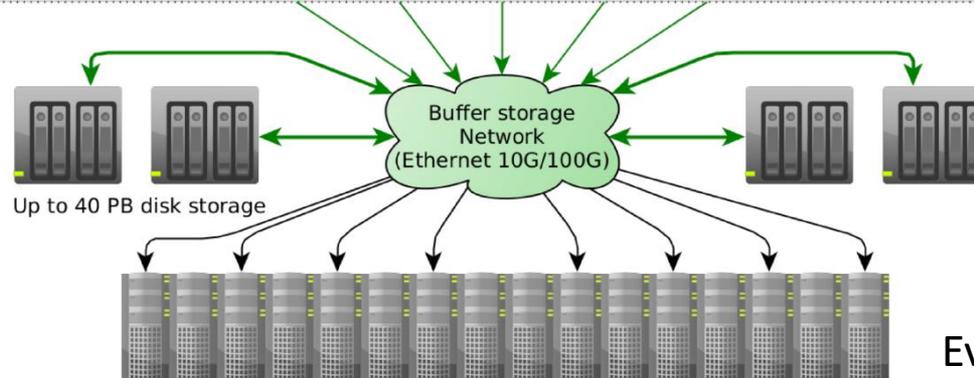
The new LHCb trigger

30 MHz
(non-empty pp)
5 TB/s



Event builder
& event filter 1

1 MHz
100 GB/s



Event filter 2

100 KHz
10 GB/s

Event filter second pass (up to 4000 servers)

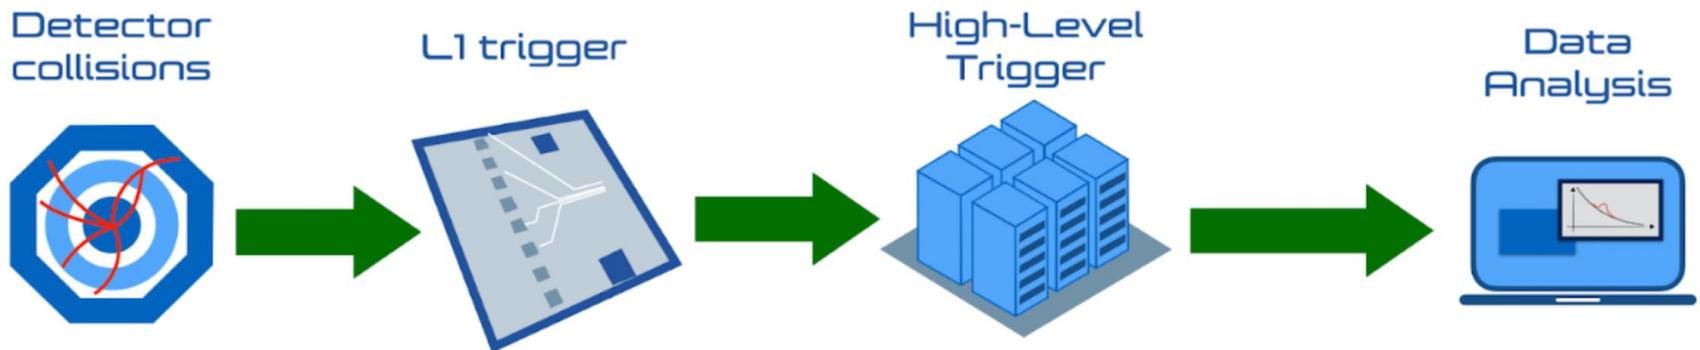


LHCb data
center at Pit 8

The new LHCb trigger

Experiment	Data size	Rate before trigger	Rate after trigger	DAQ throughput	
ALICE now	12000 kB	50 kHz	50 kHz	700 GB/s	
ATLAS now	1500 kB	27000 kHz	90 kHz	55 GB/s	90x
ATLAS 2030	5000 kB	27000 kHz	1 MHz	5000 GB/s	
CMS now	2000 kB	27000 kHz	100 kHz	200 GB/s	32x
CMS 2030	8400 kB	27000 kHz	750 kHz	6300 GB/s	
LHCb now	80 kB	24000 kHz	24000 kHz	2400 GB/s	5x
LHCb 2036	200 kB	24000 kHz	24000 kHz	12000 GB/s	

[T. Colombo]

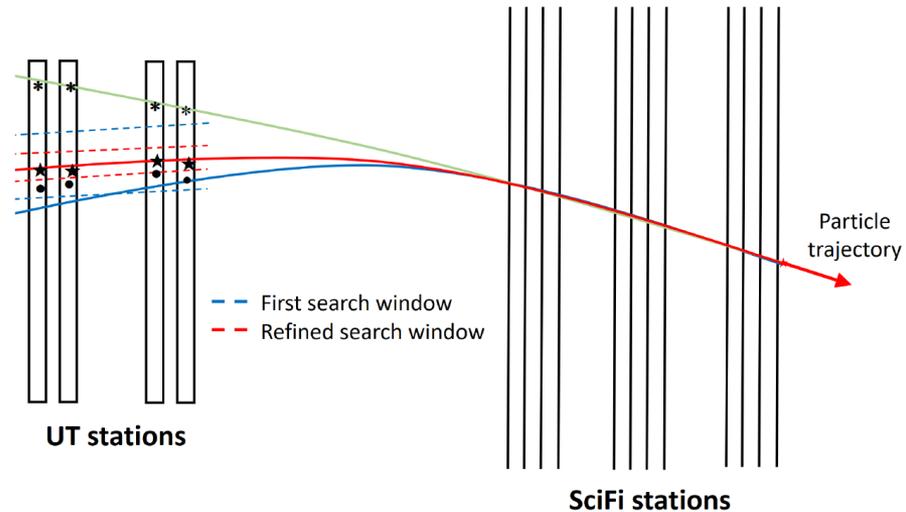
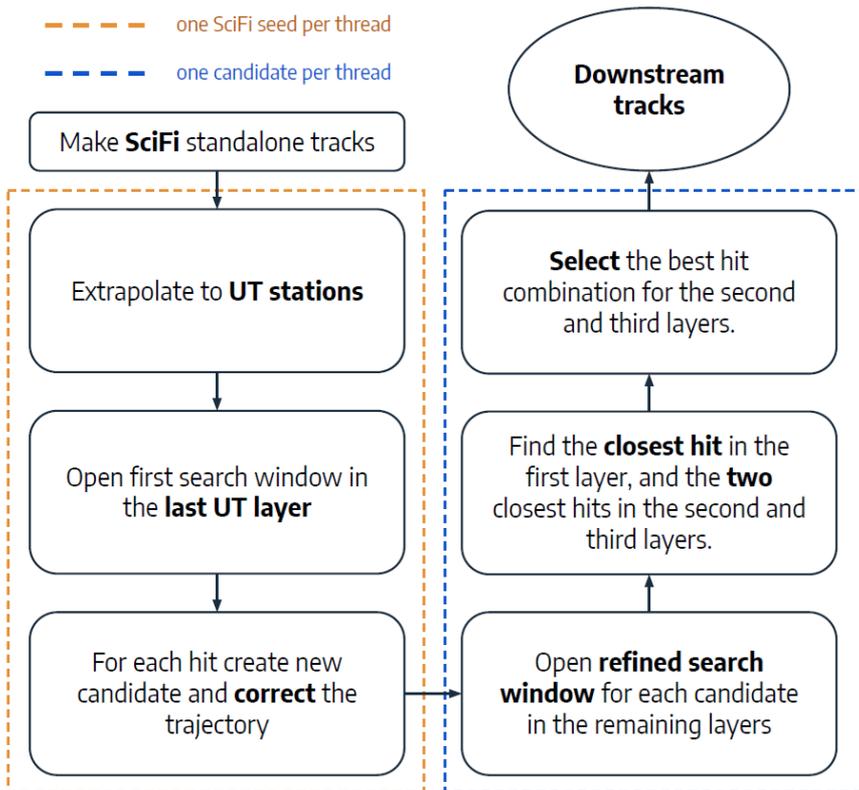


We cannot save the information of all the event!

(Scouting, TLA & Turbo techniques: persisting partial information)

Downstream tracks

Downstream algorithm design:



[\[Comput. Softw. Big Sci. 9 \(2025\) 1, 10\]](#)

C++/CUDA tricks for acceleration:

STATIC STRUCTURES

Fixing the number of input variables and the number of neurons (compile-time optimizations, registers vs global memory).

```

namespace DownstreamGhostKiller {
    namespace Model {
        constexpr unsigned num_node = 14;
        constexpr unsigned num_input = 8;
    }
}
  
```

LOOP UNROLLING

Expanding for-loops using the NVCC-specific `#pragma unroll` directive (replicates the loop body multiple times).

```

// First Layer
DownstreamHelpers::ominded, Model::num_node){&[int i] {
    DownstreamHelpers::const_inpd, Model::num_input){&[int j] {
        h1[i] += input[j] * Model::weights[i][j];
    }
};
h1[i] = ActivateFunction::relu(h1[i] + Model::bias[i]);
};
  
```

FAST MATH FUNCTIONS

Such as `_fdividef` and `_expf`, to accelerate floating-point operations.

```

namespace ActivateFunction {
    // rectified linear unit
    __device__ inline float relu(const float x) {
        return x > 0 ? x : 0;
    }
    // sigmoid
    __device__ inline float sigmoid(const float x) {
        return _fdividef(1.0f, 1.0f + _expf(-x));
    }
} // namespace ActivateFunction
  
```