Real-time analysis at the LHCb experiment

Murilo Rangel
on behalf of the LHCb Collaboration

Workshop da RENAFAE - Projetos para o Futuro da Física de Altas Energias no Brasil

Instituto de Física
Universidade Federal do Rio de Janeiro

Caveat:
1) Alguns slides estão em inglês e outros em português
2) O LHC e LHCb não serão introduzidos (detalhes técnicos estão no backup)
3) Essa apresentação é uma tentativa de atingir o objetivo do Workshop

Instituições
UFRJ, PUC and CBPF

Professores
Alberto Reis
Carla Göbel
Erica Polycarpo
Fernando Rodrigues
Helder Lopes
Murilo Rangel
Sandra Amato

Estudantes de doutorado
Lucas Meyer Garcia
Juan Baptista de Souza Leite
Felipe Luan Souza de Almeida

História

2014
Início da participação do grupo no desenvolvimento do HLT1 para GPU

2018 (Dezembro)
Real-time analysis project é aprovado pelo *LHCb Collaboration Board* (CB)

2019
Início do projeto e organização de tarefas entre os institutos

2020 (Maio)
LHCb CB aprova o uso de GPUs no primeiro nível do *trigger* (HLT1)

2021 (Fevereiro)
Assinatura do *collaboration agreement,* excluindo a parte que envolve financiamento (5 kCHF / ano Cat B).

LHCb Upgrade I

CERN-LHCC-2012-007

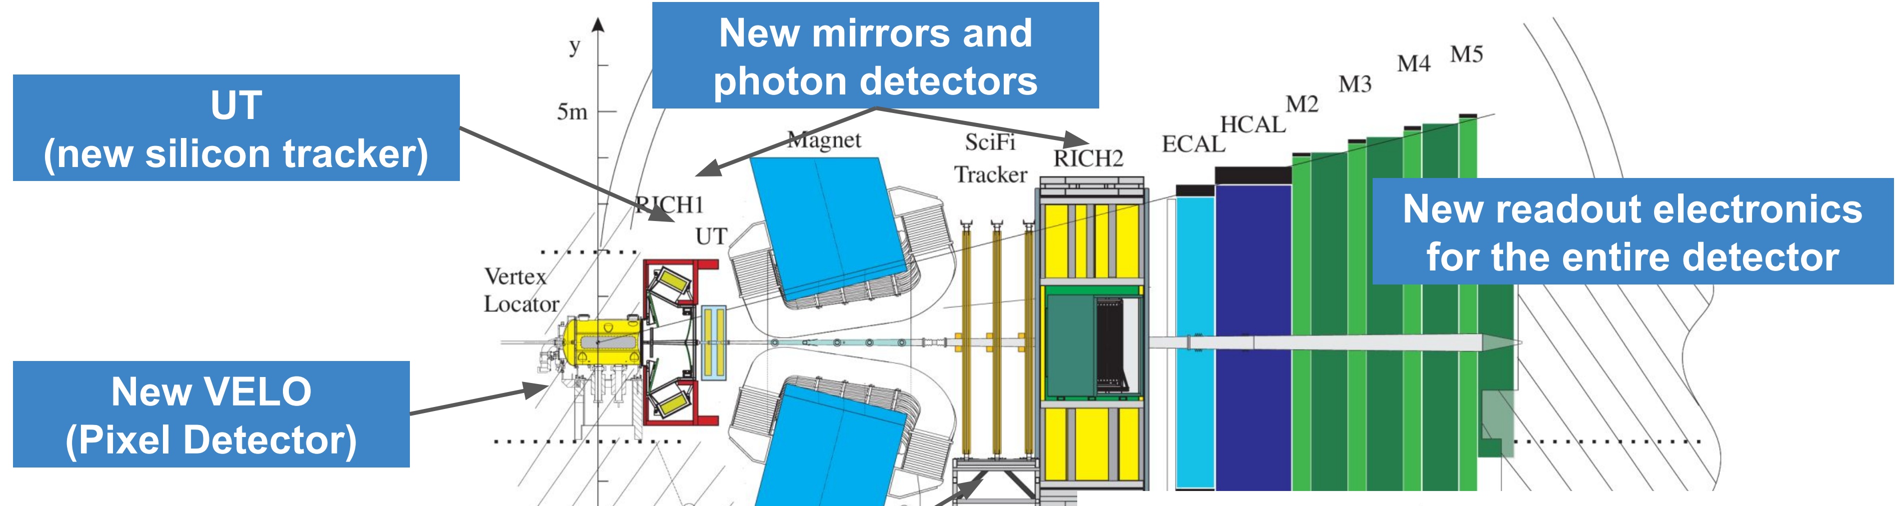

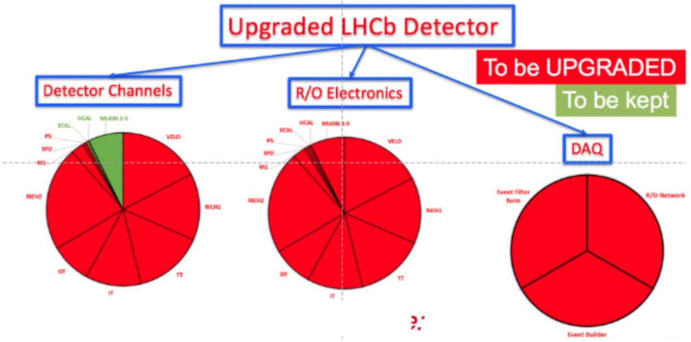

LHCb Upgrade I

CERN-LHCC-2012-007

❈ Increase instantaneous luminosity:
$4 \times 10^{32} \rightarrow 2 \times 10^{33}$ cm^{-2} s^{-1}

❈ New front-end electronics compatible with 30 MHz readout

❈ Remove hardware trigger stage and operate software trigger at 30 MHz input rate with 5 x more pileup than Run 2.

❈ Data volume to be processed in the software stage
- 32 Terabits per second
- similar to ATLAS/CMS for HL-LHC

❈ Upgrade numbers
- HLT1 output: from 100 kHz to 1 MHz
- Disk buffer contingency: from weeks to days
- HLT2 output: from 0.6 GB/s to 10 GB/s

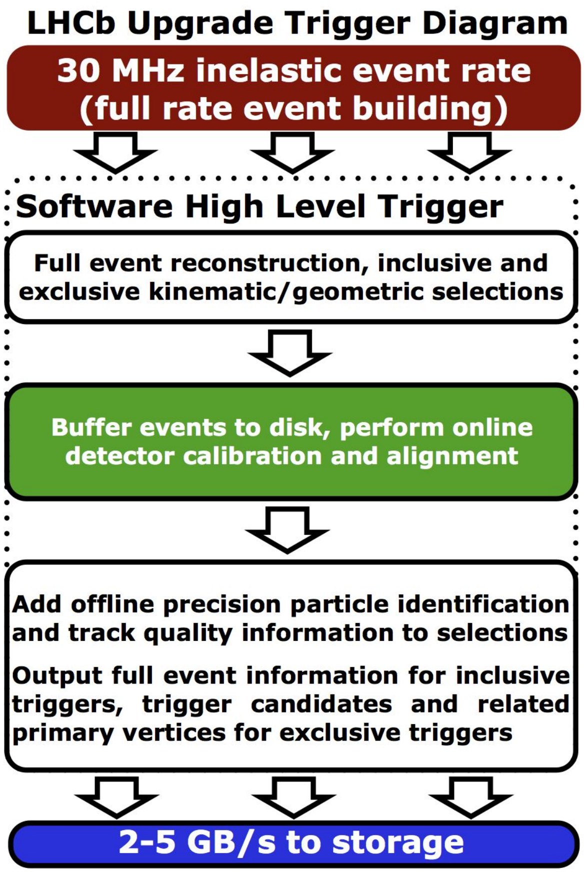

LHCB-PUB-2014-027

GPU in LHCb Trigger

Comput.Softw.Big Sci. 4 (2020) 1
https://arxiv.org/abs/2105.04031

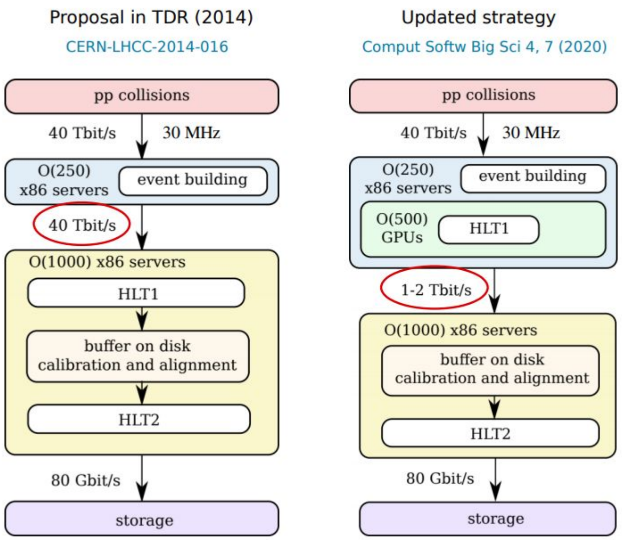

Scenario	Maximum HLT1 output rate (kHz)
CPU-only	1185
Hybrid	1425

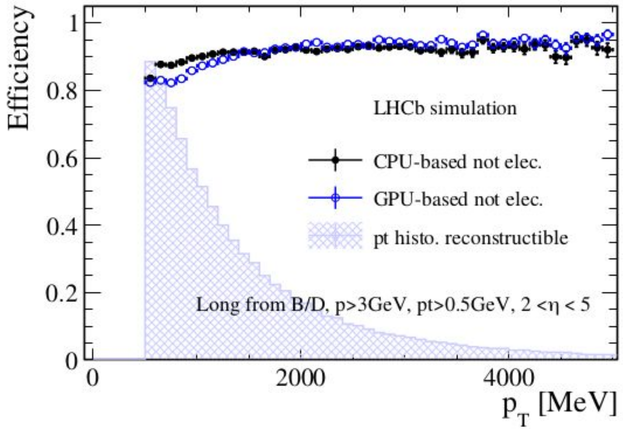

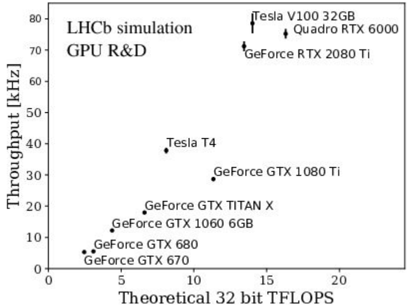

The demonstrated event throughput shows that the full HLT1 sequence can run on about 500 of either one of the RTX 2080 Ti, V100 or Quadro RTX 6000 Nvidia GPU cards.

LHCb Real Time Analysis Project

1 Aim

It is proposed to form a real-time analysis (RTA) project covering software trigger, reconstruction and calibration activities for Upgrade I onwards. The project will start in January 2019, pending ratification in the December LHCb week.

ARTICLE 2 – CONTRIBUTIONS

The RTA project is defined in full analogy to a sub-detector project. The primary way to contribute to the project is by taking responsibility for an RTA deliverable or parts of it. A deliverable designates a piece of software to be used in the trigger sequence during datataking. Maintenance is considered as a component of each deliverable, and an institute which contributes in this way will define upfront the period over which it can take responsibility for the relevant deliverable (or part of) it and for its maintenance.[2]

In addition, each RTA institute may contribute to the RTA project in one or more of the following ways:

- By basing developers at CERN during data-taking. These will be responsible for the "day-to-day operation" of the RTA software, including any necessary time-critical maintenance and development.
- By making developers available during and out of datataking for the "general maintenance" of those software components defined as the direct responsibility of the RTA project in Annex 2.
- By providing local infrastructure which enables RTA software development. Examples of this infrastructure are monitoring algorithms, automatic nightly tests or computing power to develop and test the trigger software.

LHCb Real Time Analysis Project

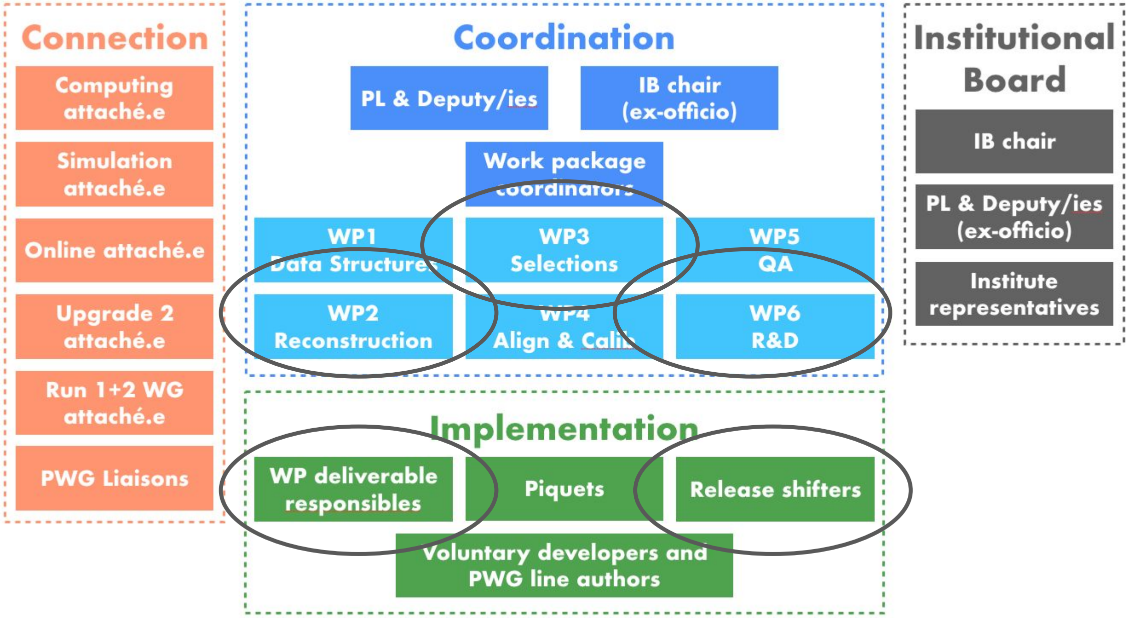

LHCb Real Time Analysis Project

<u>Nossas responsabilidades</u>

Reconstrução de jatos

Implementação da seleção dos *inputs*, clusterização (*FastJet*), jet tagging, calibração, indentificação, ...

Seleção D-meson (3-body)

Implementação de MVA, testes throughput, ..

HLT1 (GPU) reconstrução de clusters do calorímetro e seleção de di-photons

Validação com CPU e HLT2, implementação do algoritmo, ...

Shifts

Manutenção do *LHCb's software*

- **título do projeto**: Análise em tempo real no experimento LHCb
- **instituições/grupos envolvidos no projeto**: UFRJ, PUC e CBPF

- **cronograma aproximado**:
  - 2021-2022 finalizar o desenvolvimento do software para RTA
  - 2022-2024 comissionamento e operação dos *deliverables*
  - 2022-2025 montar infraestrutura para testar software para Run 4?

- **estimativa de custos e possíveis fontes de recursos**
  - Idealmente, precisamos de 5k CHF / ano para assinar o *collaboration agreement*
  - Em 2021, contribuimos 3k CHF para compra de 2 placas GPUs com bolsa de bancada JCNE/FAPERJ
  - Em geral, a infraestrutura local e viagens estão sendo suportadas por bolsas de bancada

- **possíveis sinergias com atividades experimentais de outros grupos/experimentos**
  - Provavelmente, a maior barreira é a diferença entre os *frameworks* dos experimentos
  - Podemos tentar usar infraestrutura já existente para verificar a compatibilidade para projetos futuros

- **possíveis spin-off (aplicações) do projeto**
  - Treinamento de estudantes em computação de alta performance.
  - ?

Refletindo …

High Energy Physics has a large investment in software

--- Critical part of the physics production pipeline: real-time and offline analysis software

--- LHC experiments use about 1M CPU cores every hour of every day
++ 1000 PB of data with 100 PB of data transfers per year (10-100 Gb links)

--- Estimated to be around 50M lines of C++ which would cost more than 500M to develop commercially

--- There are vast challenges ahead of us to support our ongoing physics programme

https://hepsoftwarefoundation.org

Concluindo

Desde 2019, o experimento LHCb possui o projeto *Real-Time Analysis* (RTA) para atender a necessidade de unificar a coordenação do desenvolvimento e manutenção de *software* para reconstrução, alinhamento, calibração e monitoramento em tempo real.

O projeto permite um maior nível de comprometimento nessas tarefas e um nivelamento das estratégias utilizadas, resultando em ganho das análises do experimento.

A possibilidade de um projeto com outros experimentos pode ser estudada com a infraestrutura atual para verificar se diferentes *frameworks* e realidades de cada instituição serão uma barreira.

Obrigado!!!!!

LHCb is a single arm spectrometer fully instrumented in the forward region (2.0<η<5.0)

Designed for heavy flavour physics and also exploited for general purpose physics

[Int. J. Mod. Phys. A 30, 1530022 (2015)]

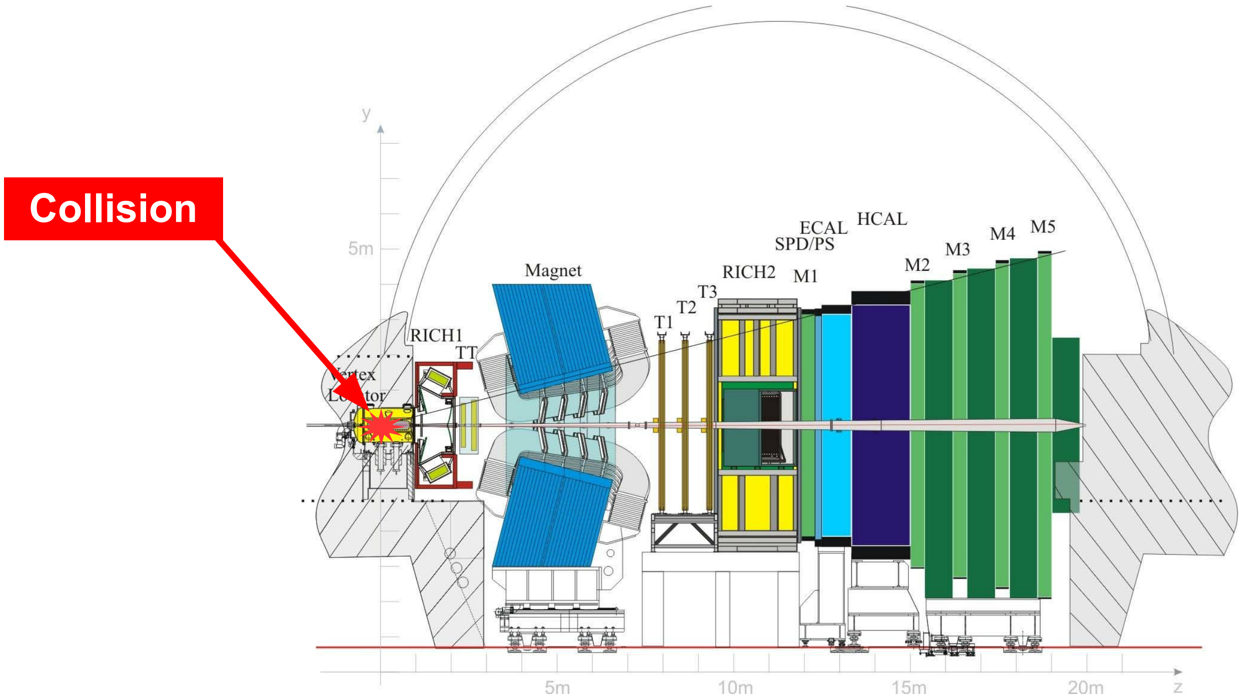

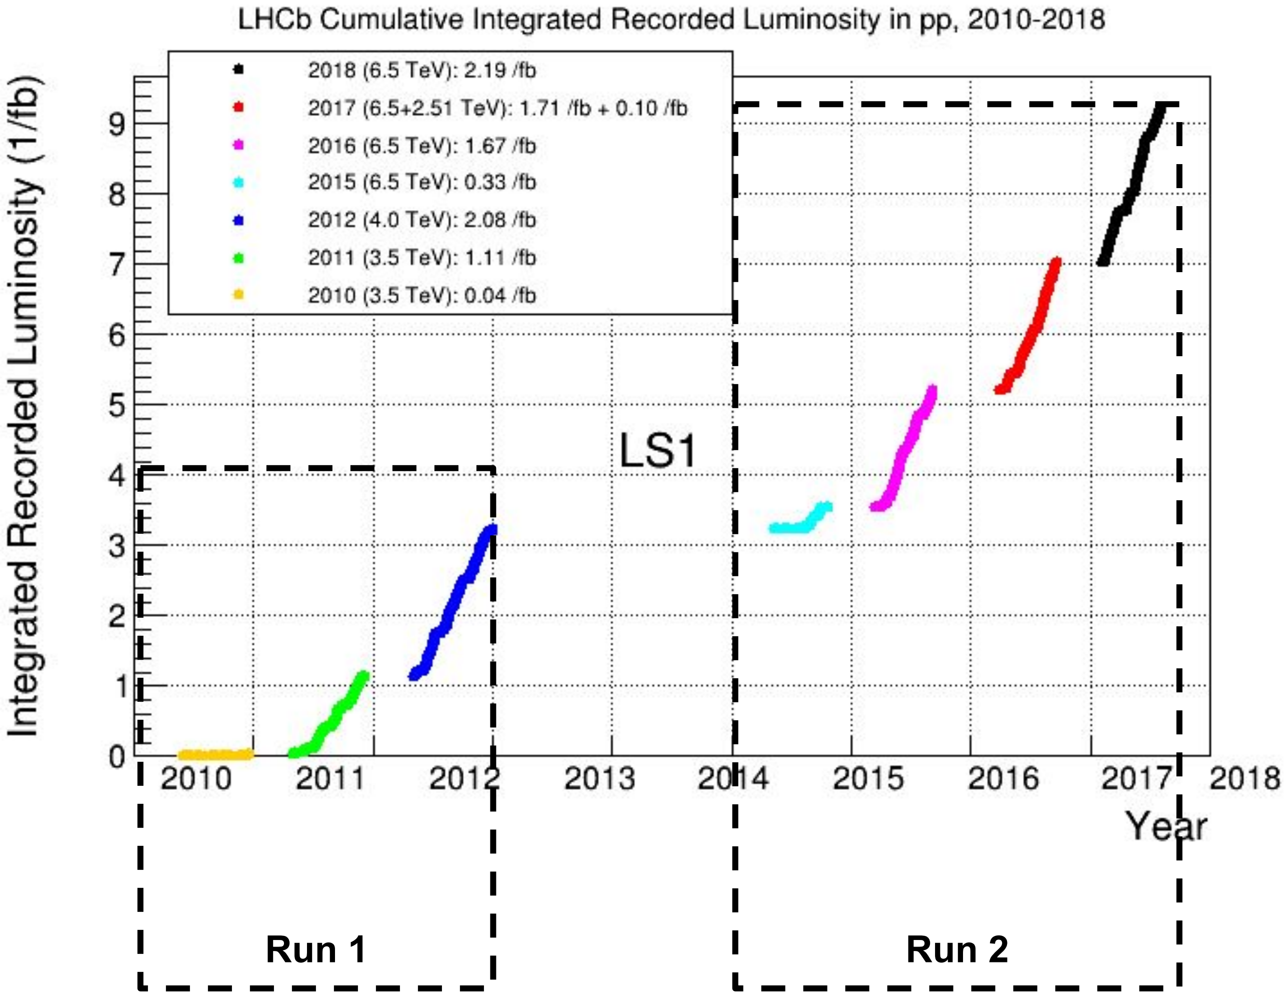

LHCb Cumulative Integrated Recorded Luminosity in pp, 2010-2018
2018 (6.5 TeV): 2.19 /fb
2017 (6.5+2.51 TeV): 1.71 /fb + 0.10 /fb
2016 (6.5 TeV): 1.67 /fb
2015 (6.5 TeV): 0.33 /fb
2012 (4.0 TeV): 2.08 /fb
2011 (3.5 TeV): 1.11 /fb
2010 (3.5 TeV): 0.04 /fb
Integrated Recorded Luminosity (1/fb)
Year
LS1
Run 1
Run 2

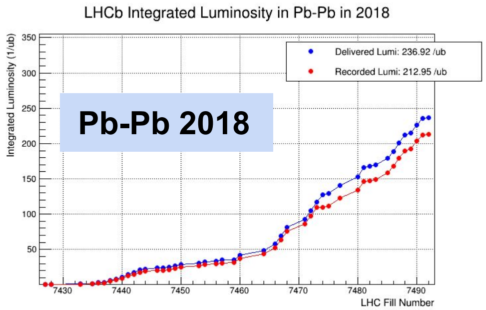

LHCb Integrated Luminosity in Pb-Pb in 2018
Delivered Lumi: 236.92 /ub
Recorded Lumi: 212.95 /ub
Pb-Pb 2018
Integrated Luminosity (1/ub)
LHC Fill Number

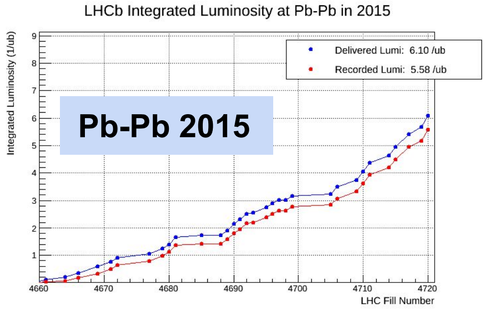

LHCb Integrated Luminosity at Pb-Pb in 2015
Delivered Lumi: 6.10 /ub
Recorded Lumi: 5.58 /ub
Pb-Pb 2015
Integrated Luminosity (1/ub)
LHC Fill Number

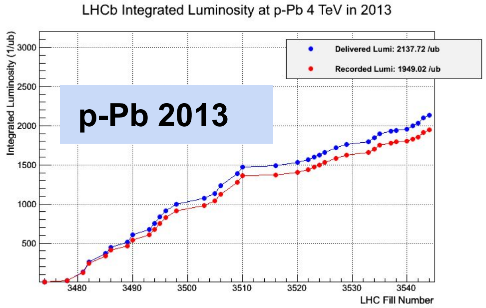

LHCb Integrated Luminosity at p-Pb 4 TeV in 2013
Delivered Lumi: 2137.72 /ub
Recorded Lumi: 1949.02 /ub
p-Pb 2013
Integrated Luminosity (1/ub)
LHC Fill Number

Run 2 trigger

LHCb Run II Trigger Diagram (2015 - 2019)

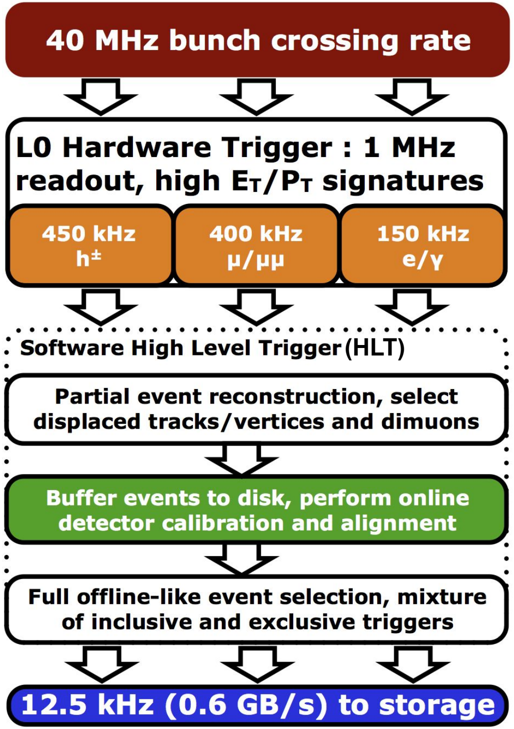

JINST 14 (2019) no.04, P04013

Trigger structure:

- Hardware: energies deposited in calorimeters and muon stations hits are used to bring 40 MHz to 1 MHz

- Software: events built at 1 MHz (~27000 physical cores)
  HLT1: fast tracking and inclusive selections
  1 MHz to 100 kHz
  HLT2: complete event reconstruction and selections

Run 2 trigger

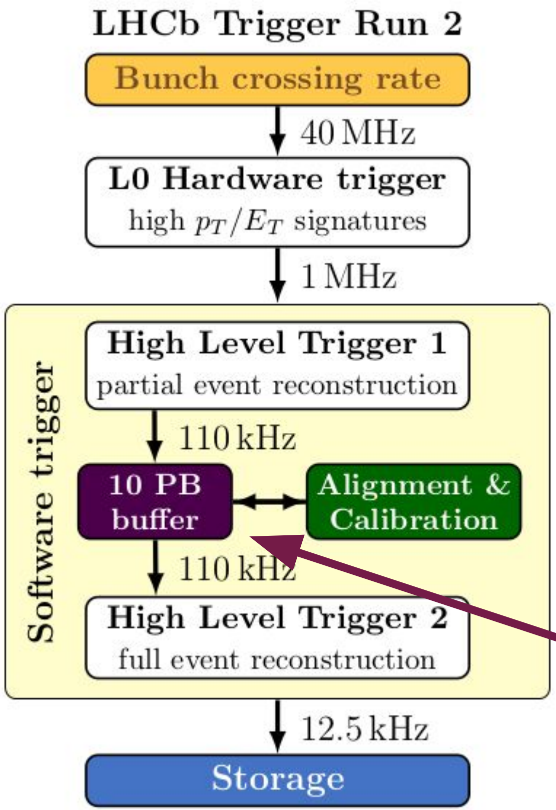

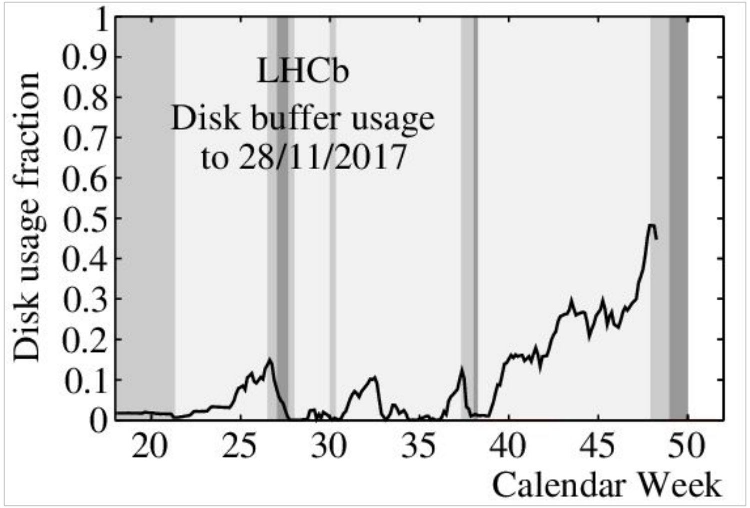

♯ HLT Farm with 10 PB disk space
♯ At an average event size of 55 kB with 100 kHz: up to 2 weeks before HLT2 has to be executed
♯ 2x trigger CPU capacity since Farm is used twice for HLT (excess used for simulation)

JINST 14 (2019) no.04, P04013

Run 2 trigger

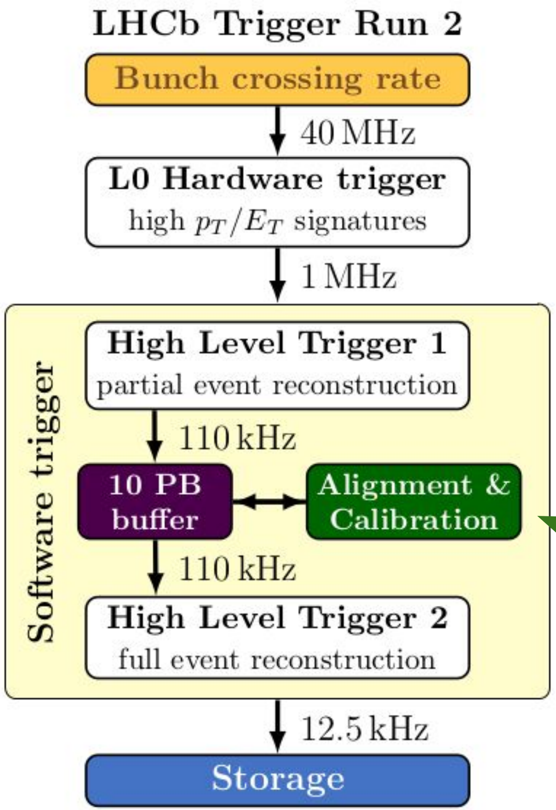

JINST 14 (2019) no.04, P04013

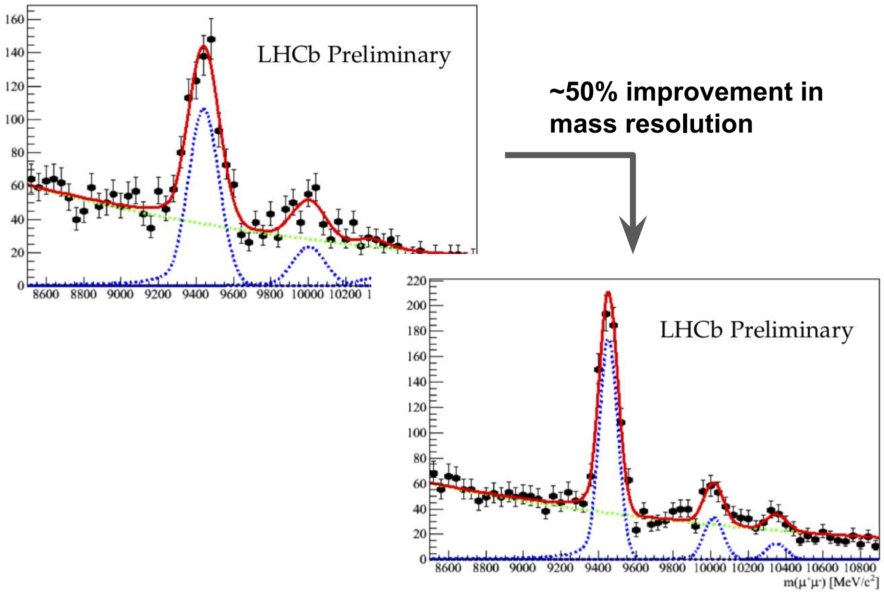

~50% improvement in mass resolution

♯ Real-time alignment and calibration
♯ Dedicated HLT1 trigger lines supply samples for the alignment
♯ Alignment & calibration tasks run in parallel while events are are being processed by HLT1

Run 2 trigger: Turbo

$$\text{Bandwidth}\,[\text{GB}\,\text{s}^{-1}] \propto \text{Trigger output rate}\,[\text{kHz}] \times \boxed{\text{Average event size}\,[\text{kB}]}$$

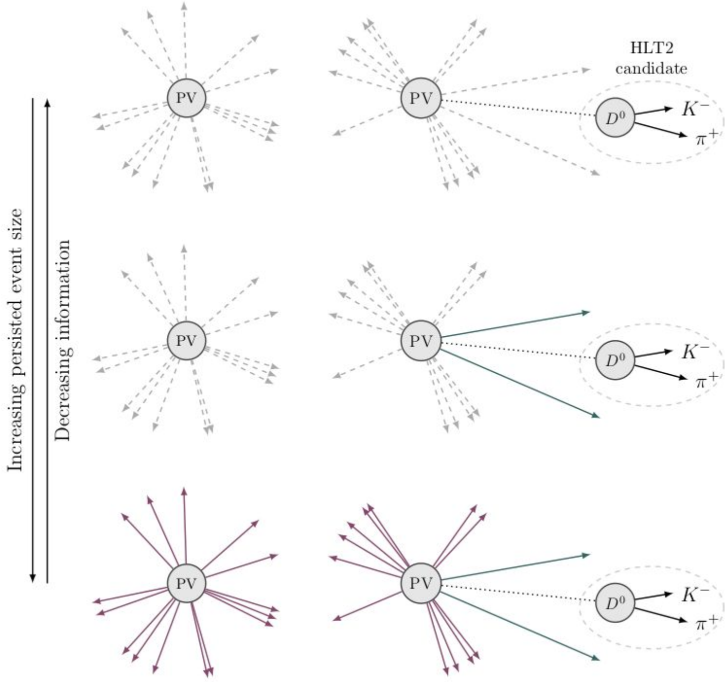

Turbo data processing model

♯ Analyses that can be done using trigger objects can profit of reduced event size and higher trigger rate.

♯ Event size can be reduced from 70 kB to 7 kB depending on the persistence level

♯ Calibration samples increased, reducing systematic uncertainties on efficiency measurements

♯ 50% of HLT2 trigger lines are Turbo counting 10% of the bandwidth

JINST 14 (2019) no.04, P04006

Run 2 Trigger: Turbo Analyses

Study of J/ψ Production in Jets

R. Aaij *et al.* (LHCb Collaboration)
Phys. Rev. Lett. **118**, 192001 – Published 8 May 2017

Physics See Viewpoint: Probing Quarkonium Production in Jets

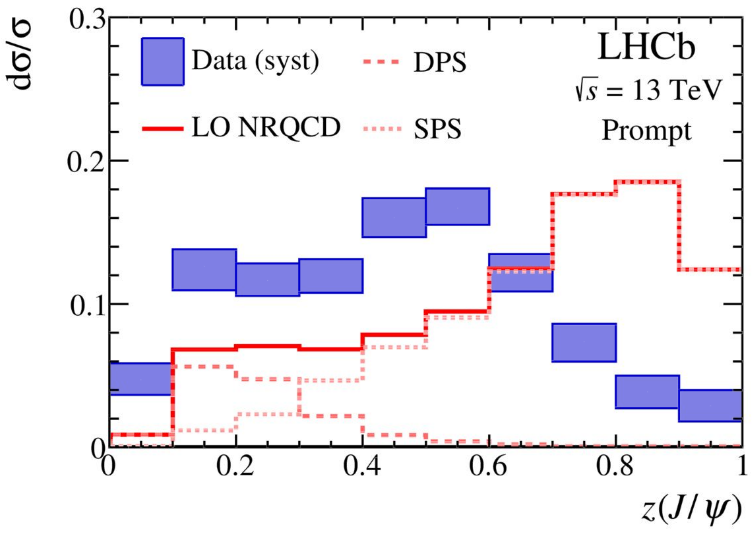

Observation of the Doubly Charmed Baryon Ξ_{cc}^{++}

R. Aaij *et al.* (LHCb Collaboration)
Phys. Rev. Lett. **119**, 112001 – Published 11 September 2017

Physics See Viewpoint: A Doubly Charming Particle

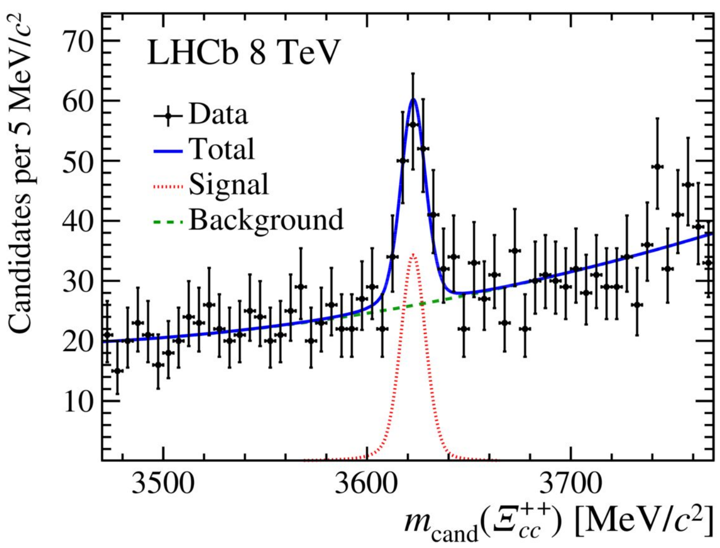

Phys.Rev.Lett. 118 (2017) no.19, 192001

Phys. Rev. Lett. 119, 112001

Run 2 trigger: Efficiencies

$$\epsilon = \frac{N(\text{TOS and TIS})}{N(\text{TIS})}$$

TOS: events triggered on the signal
TIS: events triggered independently of the presence of the signal

If entire L0 bandwidth is granted
If there is bandwidth division

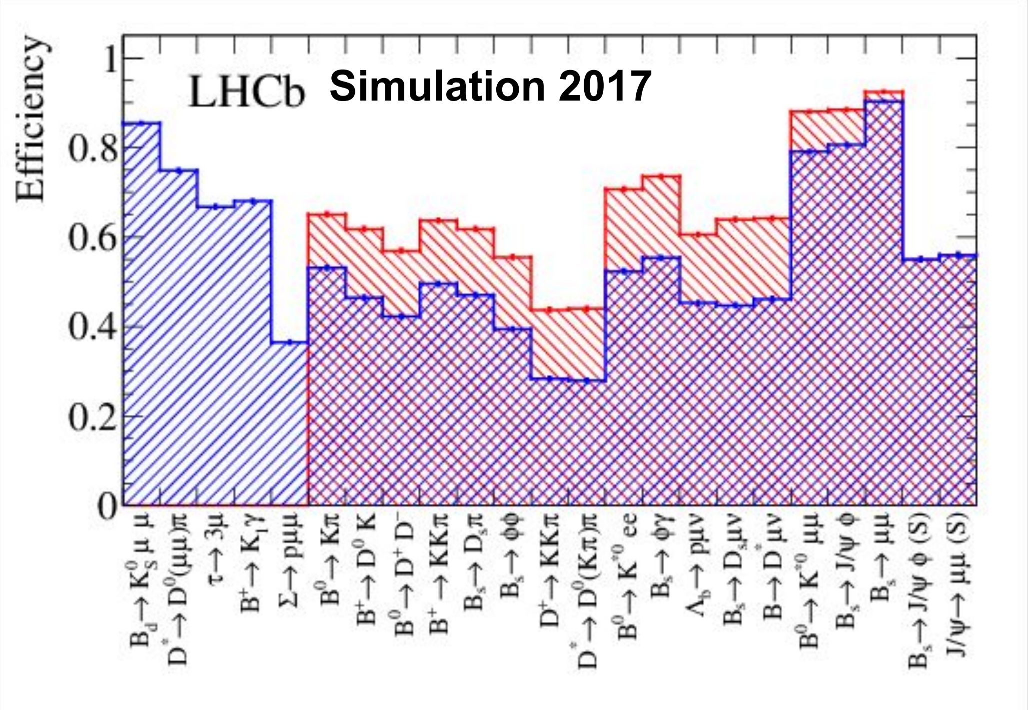

JINST 14 (2019) no.04, P04013

Run 2 trigger: Plots

LHCb-CONF-2016-005

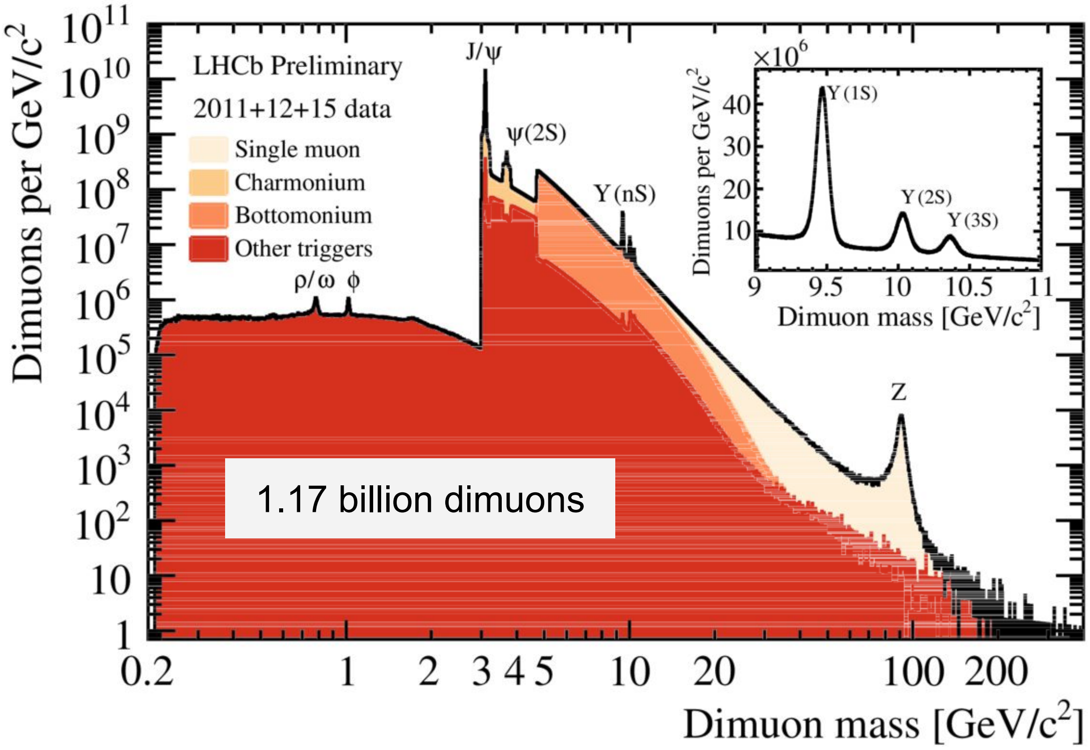

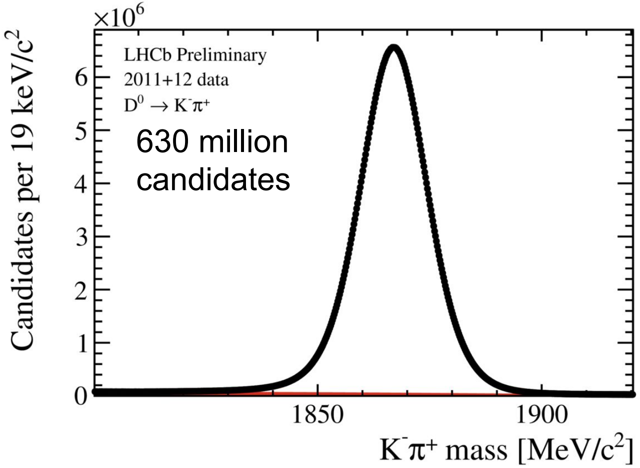

Rare events: high efficiency
Copious production: high purity

Run 3 Trigger

LHCB-PUB-2014-027

Rates as a function of pT cut for part. reco. candidates

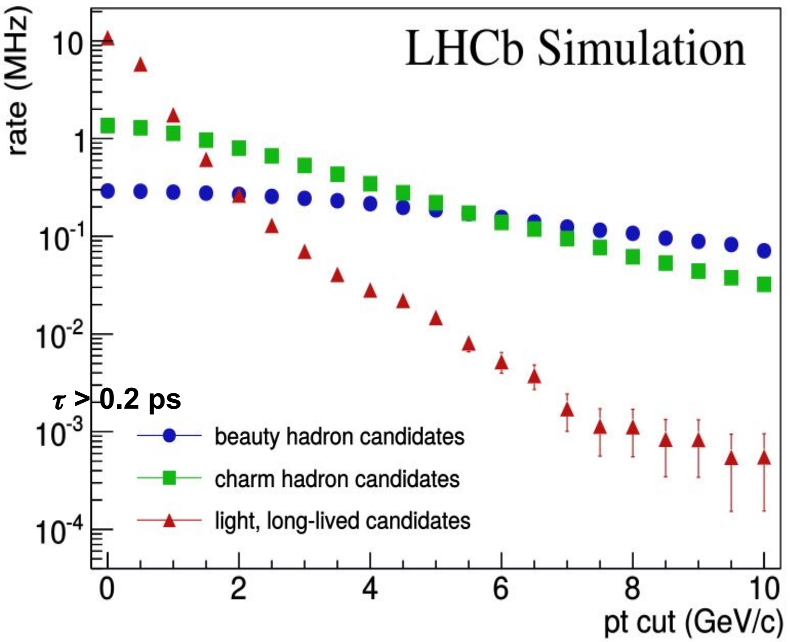

2018

stream	event size (kB)	event rate (kHz)	rate fraction	throughput (GB/s)	bandwidth fraction
FULL	70	7.0	65%	0.49	75%
Turbo	35	3.1	29%	0.11	17%
TurCal	85	0.6	6%	0.05	8%
total	61	10.8	100%	0.65	100%

Upgrade

stream	rate fraction	throughput (GB/s)	bandwidth fraction
FULL	26%	5.9	59%
Turbo	68%	2.5	25%
TurCal	6%	1.6	16%
total	100%	10.0	100%

Run 3 Trigger: HLT1

LHCB-PUB-2017-006

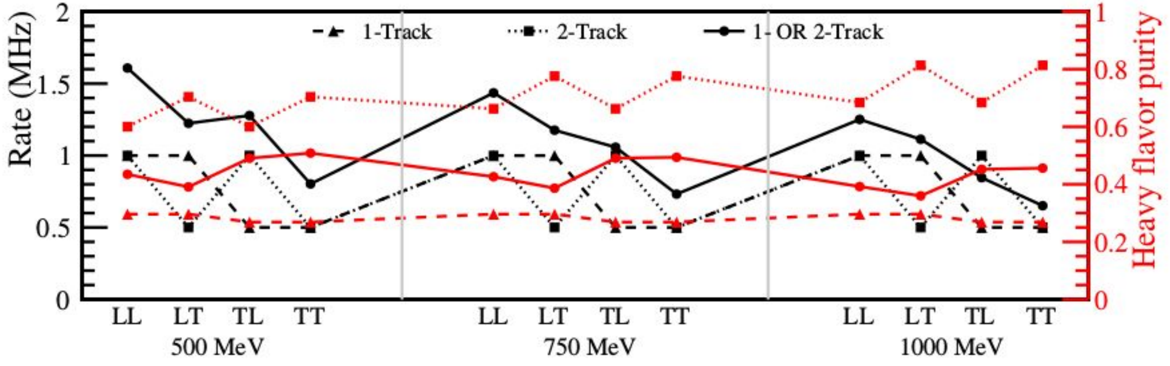

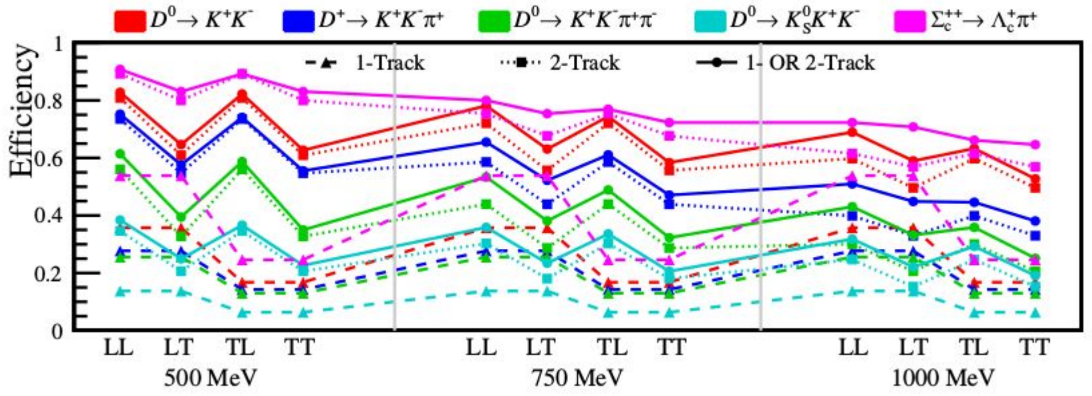

♯ Loose (L) track selection: 1 MHz
♯ Tight (T) track selection: 0.5 MHz

♯ Inclusive selection using combination of Loose (L) and Tight (T) with different p_T thresholds

♯ Optimisation for both selections is expected using more realistic ghost rate

Run 3 Trigger: HLT2

LHCB-PUB-2017-006

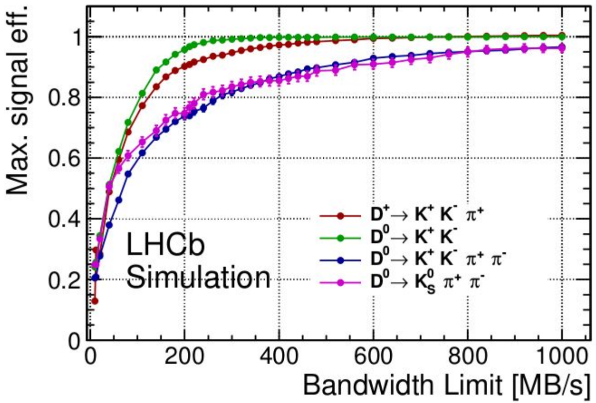

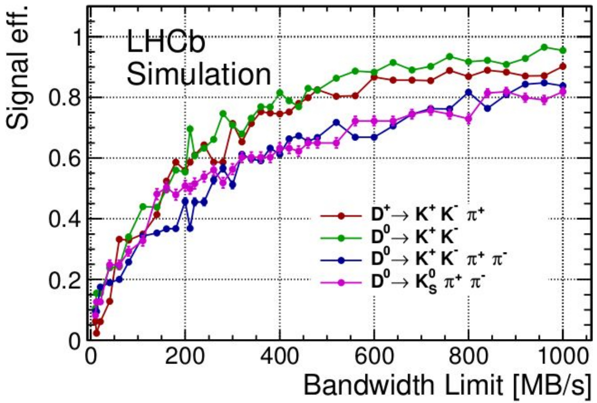

♯ Same strategy as in Run 2 increasing Turbo stream usage

♯ To cope with many (> 500) trigger lines, a genetic-algorithm based on bandwidth division is used.

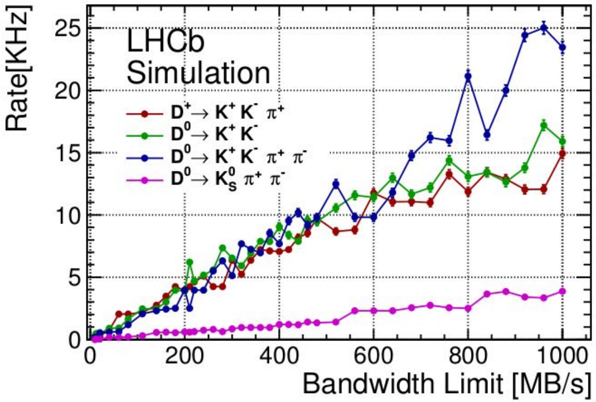

LHCb Upgrade

arXiv:1808.08865 [hep-ex]

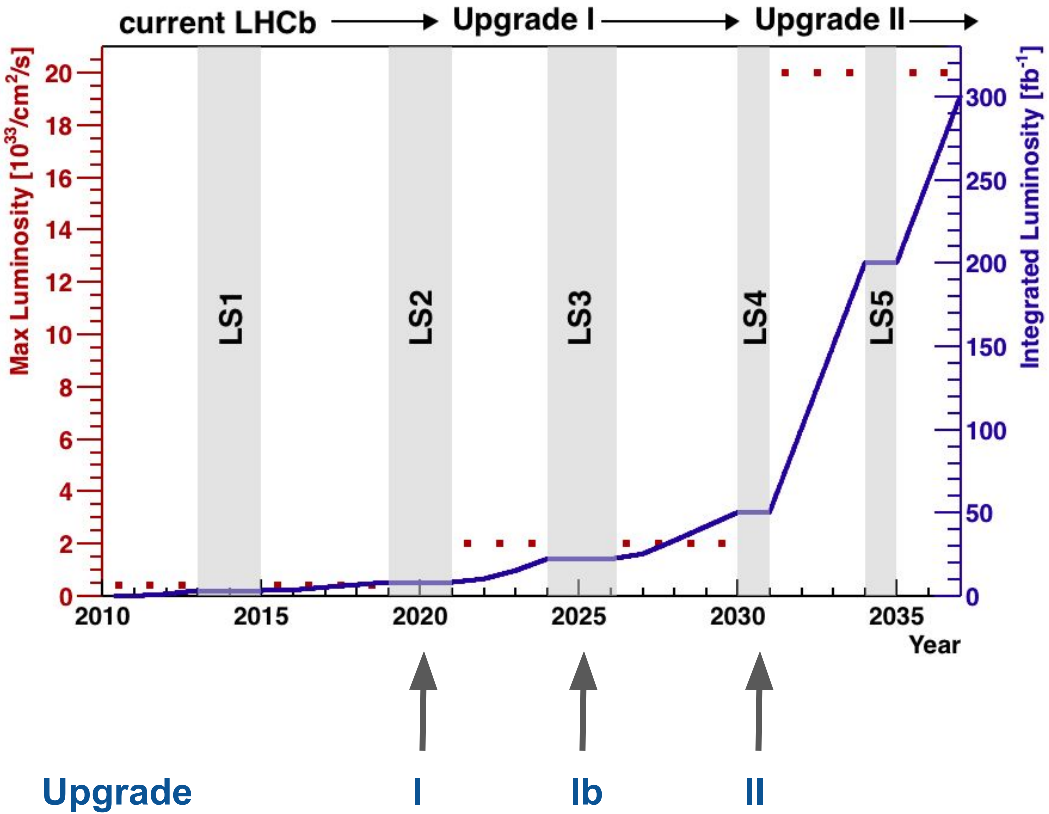

Run 1+Run 2: 9.1/fb
Run 3: 25/fb

LHCb Upgrade

CERN-LHCC-2012-007

UT
(new silicon tracker)

New mirrors and
photon detectors

New VELO
(Pixel Detector)

New readout electronics
for the entire detector

SciFi
(new scintillating fibre tracker)

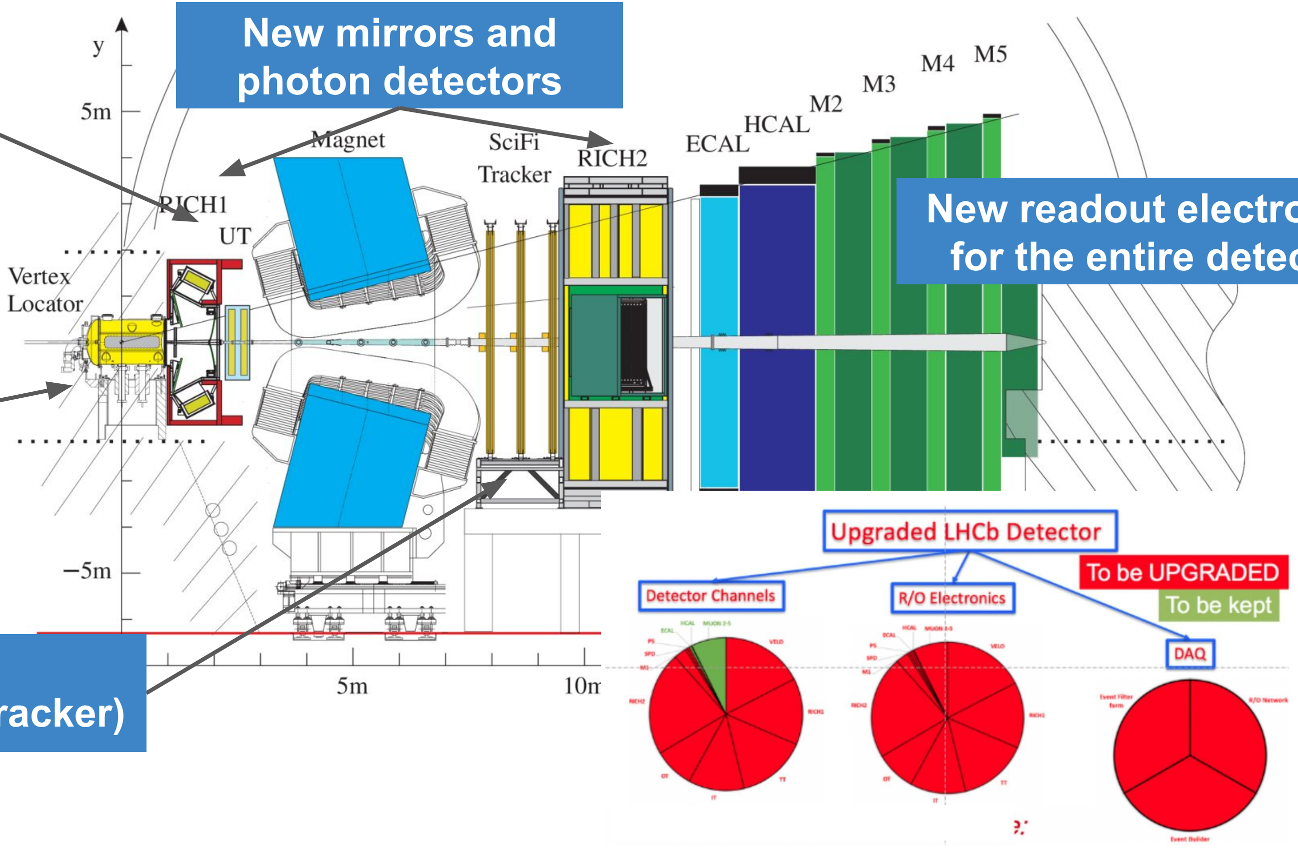

GPU in LHCb Trigger

Comput.Softw.Big Sci. 4 (2020) 1

- GPU will be use in the first trigger level

- One GPU should process events at roughly 60 kHz
  ++Reduce data rate by factor 30 based on single and two-track selections

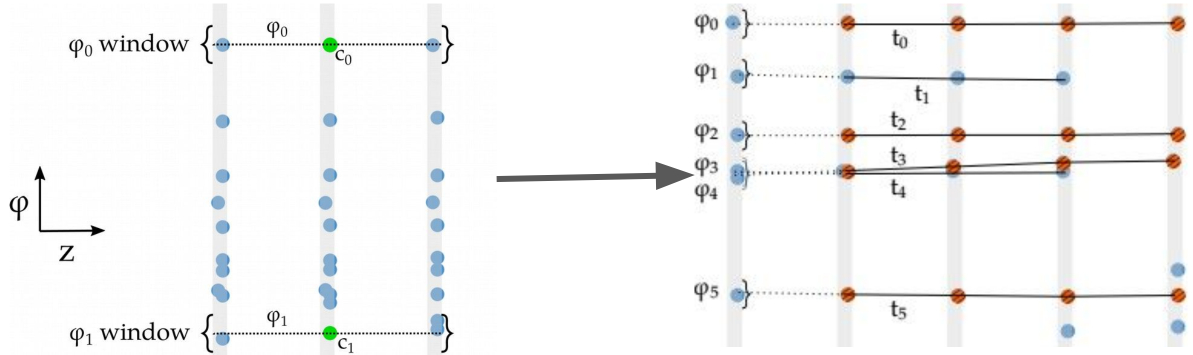

Velo on GPUs
++ 26 planes of silicon pixel detectors
++ Parallel algorithm for cluster finding using bitmasks
++ Pattern recognition finding seed triplets in parallel
++ Primary vertex finding using cluster finding in a histogram

GPU in LHCb Trigger

Comput.Softw.Big Sci. 4 (2020) 1
https://arxiv.org/abs/2105.04031

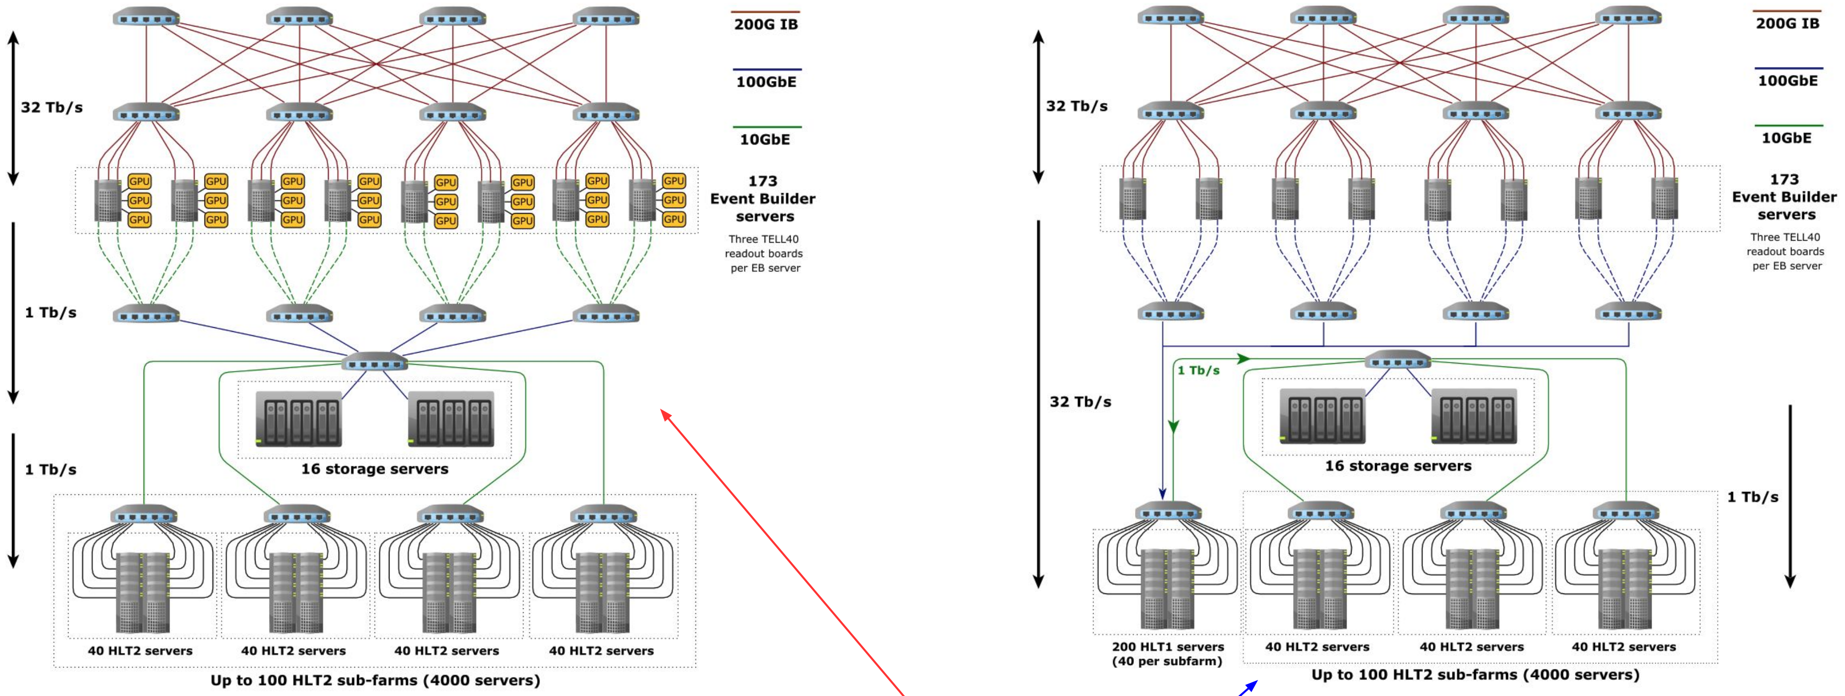

Scenario	Maximum HLT1 output rate (kHz)
CPU-only	1185
Hybrid	1425

Status of software

Allen project

- Implemented on C++, CUDA and python
- Running also for CPU, and HIP (AMD, experimental)
- It can run standalone and on Gaudi
- Still improving:
  v0.9 released last week

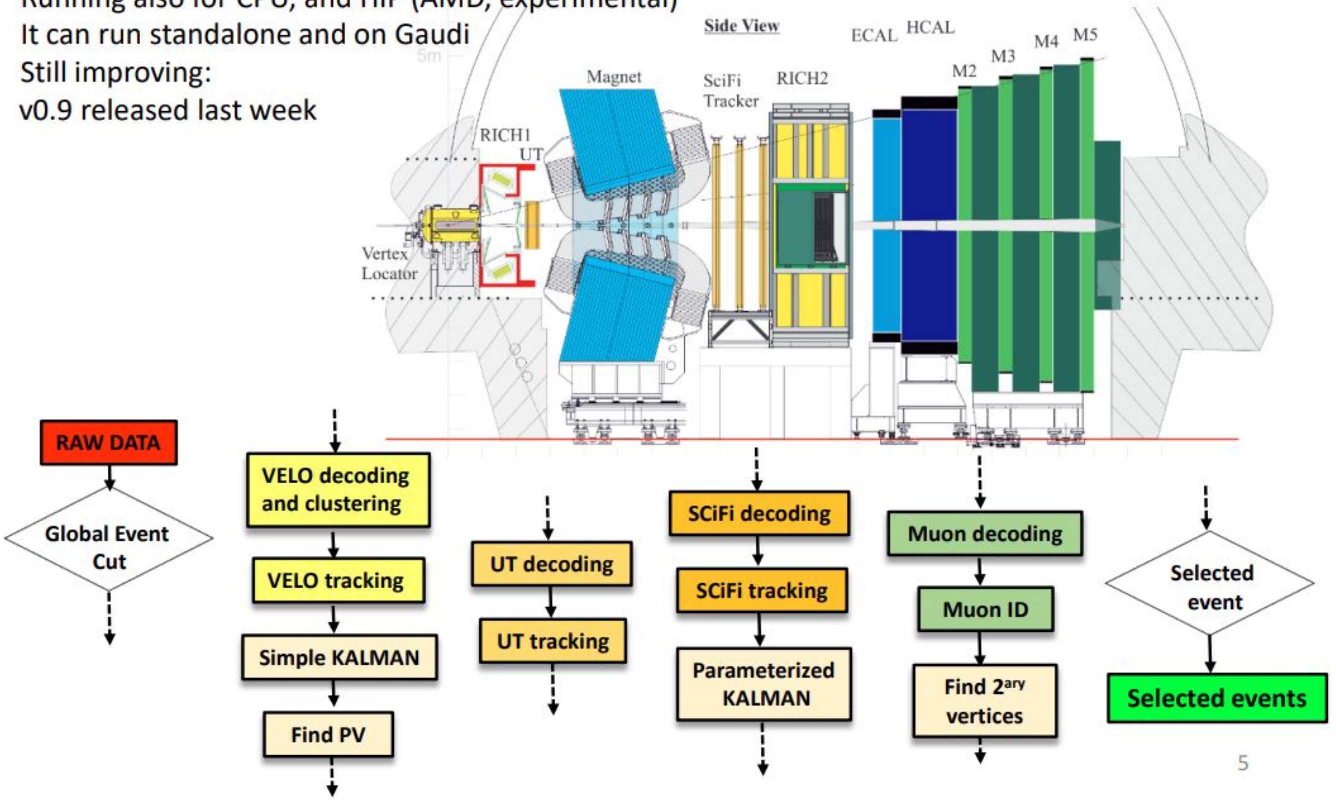

Table 2. Overview of GPU usage for real-time analysis in various HEP experiments.

Experiment	Main task processed on GPU	Event / data rate	Number of GPUs	Types of GPUs tested	Date for employment	References
NA62	RICH ring pattern reconstruction	10 MHz / 2.5 Gbit/s	1	Nvidia K20c, P100	Tested in 2017 & 2018, planned for 2021	[7, 8]
Mu3e	Track- & vertex reconstruction in the pixel tracker, data selection	20 MHz / 32 Gbit/s	$\mathcal{O}$ (10)	Nvidia GTX980, GTX1080, RTX1080Ti	2021	[9]
CMS	Decoding of raw data, clustering, pattern recognition in the pixel detector	100 kHz / -		Nvidia RTX2080, K20	Planned for 2021	[10, 11]
ALICE	Track reconstruction in the TPC	< 500 Hz Pb-Pb or < 2 kHz p-p / < 100 Gbit/s	64	Nvidia GTX480	2010–2013	[12]
ALICE	Track reconstruction in the TPC	< 1 kHz Pb-Pb or < 2 kHz p-p / < 384 Gbit/s	180	AMD S9000	2015–2018	[12]
ALICE	Track reconstruction in three sub-detectors	50 kHz Pb-Pb or <5 MHz p-p / 30 Tbit/s	$\mathcal{O}$ (2000)		2021	[13, 14]
LHCb	Decoding of raw data, clustering, track reconstruction in three sub-detectors, vertex reconstruction, muon identification, inclusive selections	30 MHz / 40 Tbit/s	$\mathcal{O}$ (500)	Nvidia RTX2080Ti, RTX6000, V100	Possibly in 2021	[15]

D. vom Bruch 2020 JINST 15 C06010

Run 2 trigger: Efficiencies

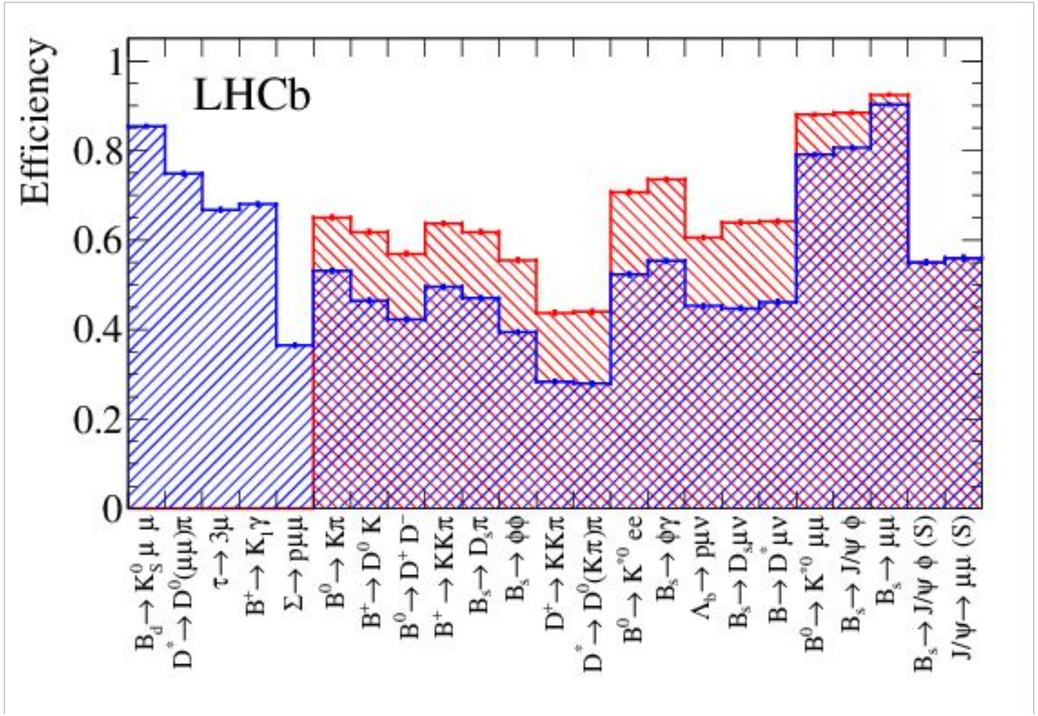

JINST 14 (2019) no.04, P04013

LHCb Run II trigger

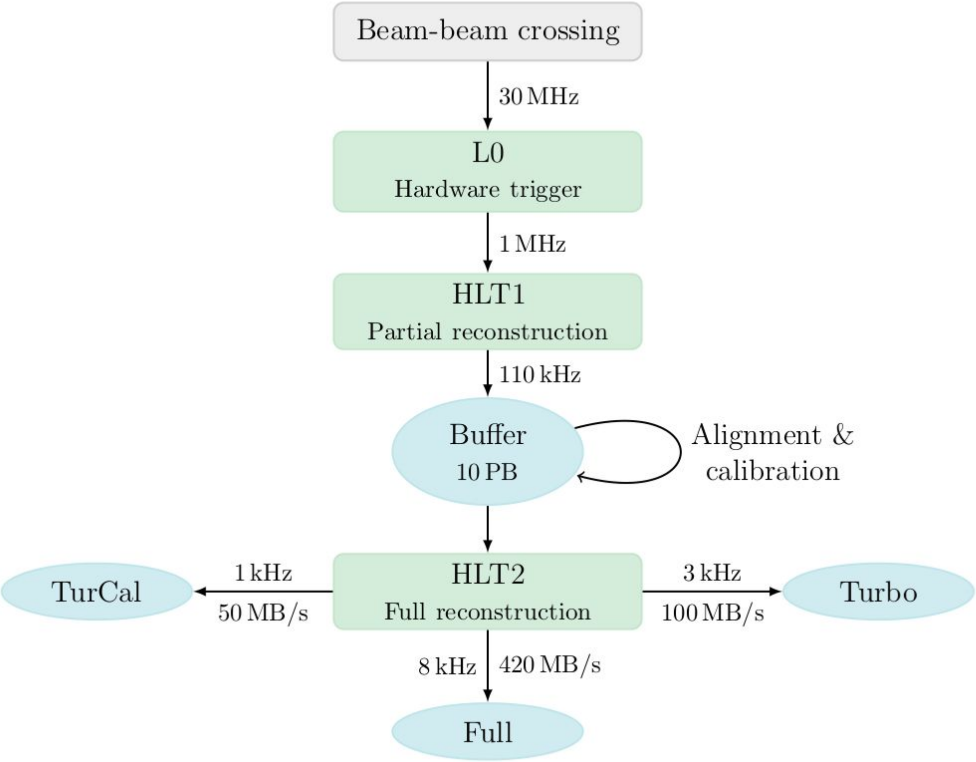

JINST 14 (2019) no.04, P04006

Turbo


2015
2016
2017
Total
Muon+El+EW
Lowmult
Incl.B
Exot+Jet+Strng
Excl.B
Charm
Calib
0
5
10
15
20
Rate [kHz]
stream
CALIBRATION
TURBO
FULL

Run 3 Trigger: HLT1

LHCB-PUB-2017-006

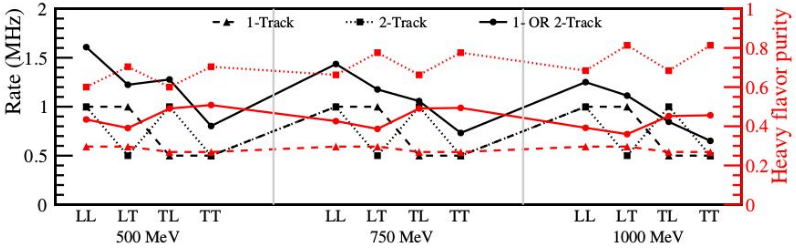

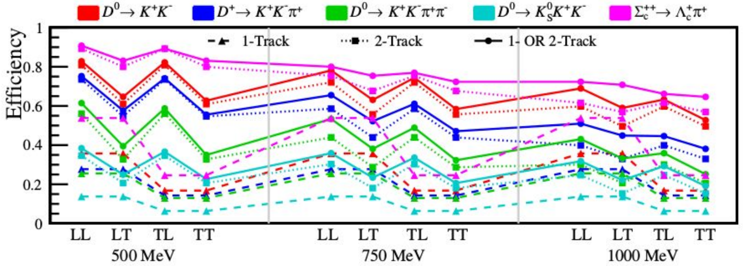

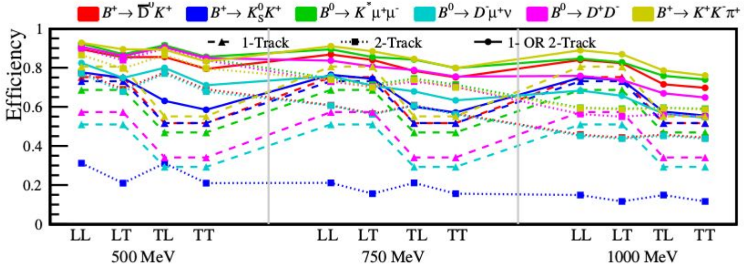

♯ Loose (L) track selection: 1 MHz
♯ Tight (T) track selection: 0.5 MHz

♯ Inclusive selection using combination of Loose (L) and Tight (T) with different p_T thresholds

♯ Optimisation for both selections is expected using more realistic ghost rate

Run 2 trigger: Real-time alignment

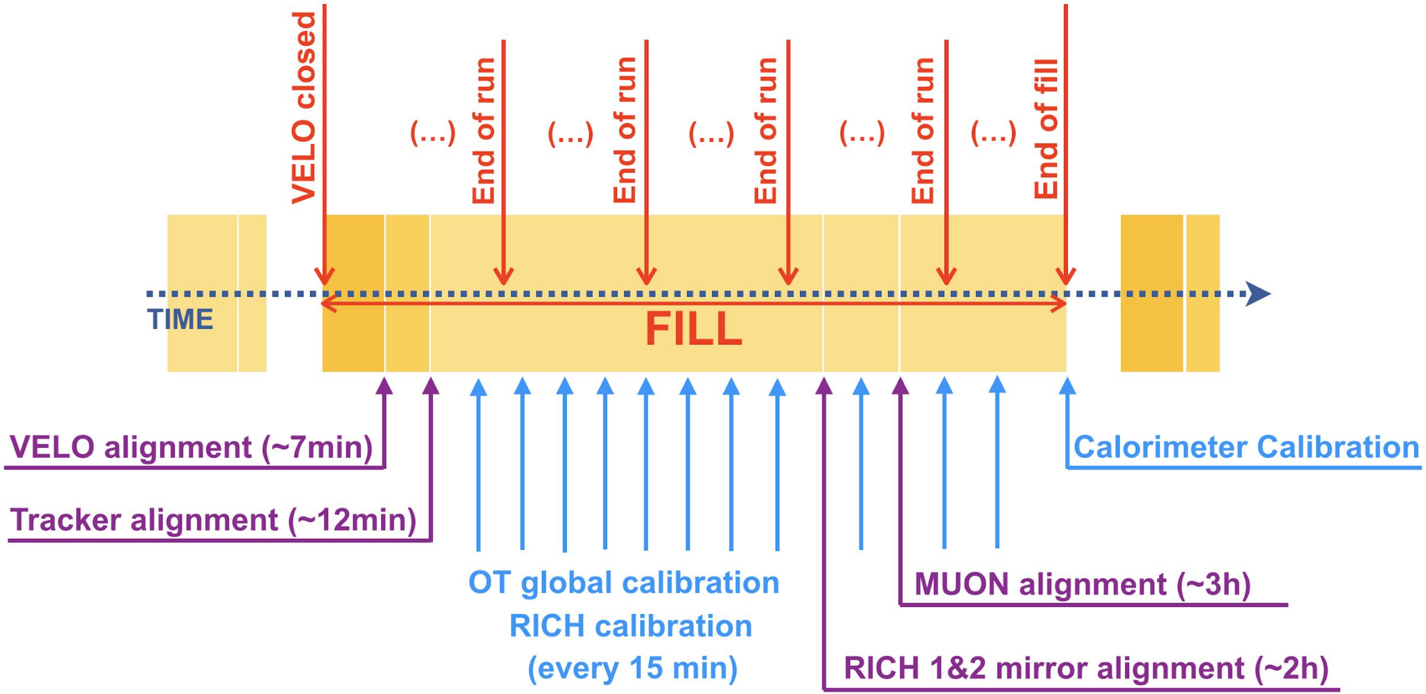

((~7min),(~12min),(~3h),(~2h)) - time needed for both data accumulation and running the task