# Turbo: the flexible reduced data format for real-time analysis at LHCb

Alex Pearce

on behalf of the LHCb collaboration

Saturday 7th July 2018

ICHEP 2018, Seoul, Republic of Korea

# The LHCb experiment

JINST **3** (2008) S08005, Int. J. Mod. Phys. **A30** (2015) 1530022

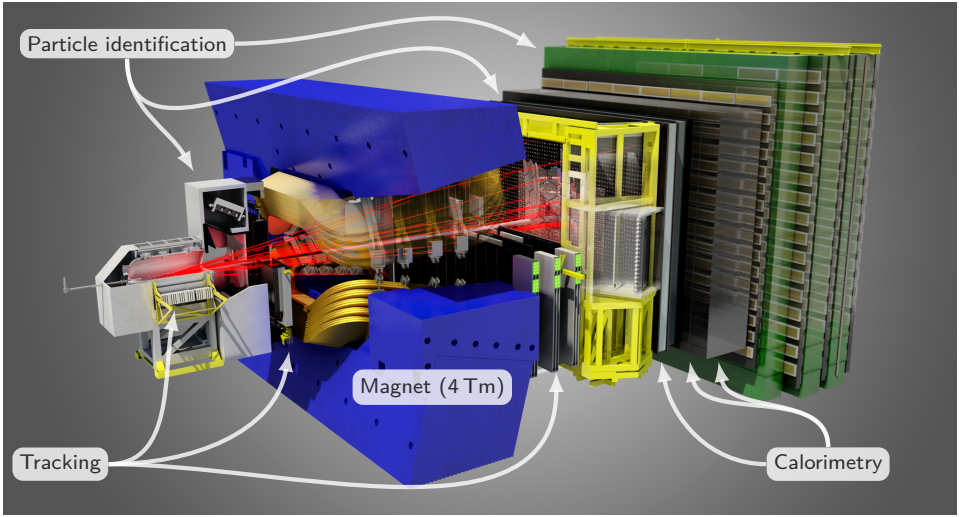


| | |
|---|---|
| VELO | Primary and secondary vertex, impact parameter |
| TT, IT, OT | Momentum of charged particles |
| RICHs | $K^{\pm}$, $\pi^{\pm}$, and $p/\overline{p}$ PID |
| MUON | Trigger on high $p_{\mathrm{T}}$ $\mu^{\pm}$, add PID |
| SPD/PS | Separate $\gamma/e^{\pm}$ and $h^{\pm}/e^{\pm}$ |
| ECAL/HCAL | EM/hadronic energy |

# LHCb is a general-purpose detector

- 'The heavy flavour experiment at the LHC'
- We have a fantastically broad programme
  - Beauty
  - Charm
  - Production (QCD, CEP)
  - Electroweak
  - Exotics
  - Heavy ion
  - Fixed target
  - ...
- Permitted so far by a very flexible two-stage software trigger

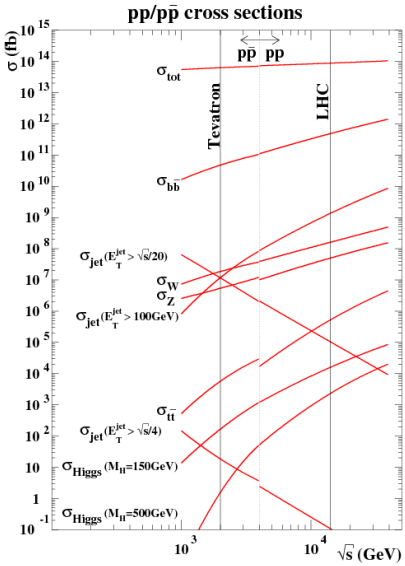


Phys. Rep. **426** (2006) 47-358

# Online data flow overview: Run 1

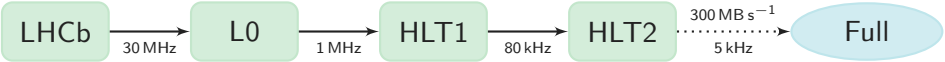


1. Hardware trigger uses information from calorimeters and muon stations
2. HLT1 performs a simplified event reconstruction, confirms L0 decision
3. HLT2 reconstructs the event
   - Primary vertices, full and stub tracks, neutrals...
   - Hundreds of trigger lines, typically fully reconstructing exclusive decay, also a core set of inclusive selections
4. Subsequent reconstruction performed offline to achieve best performance
   - Occasionally re-run to update alignment and calibration constants

# Online data flow overview: Run 1

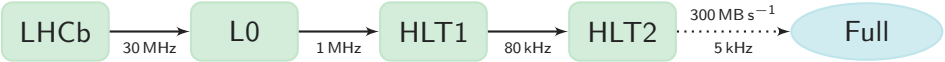


1. Hardware trigger uses information from calorimeters and muon stations
2. HLT1 performs a simplified event reconstruction, confirms L0 decision
3. HLT2 reconstructs the event
   - Primary vertices, full and stub tracks, neutrals...
   - Hundreds of trigger lines, typically fully reconstructing exclusive decay, also a core set of inclusive selections
4. Subsequent reconstruction performed offline to achieve best performance
   - Occasionally re-run to update alignment and calibration constants

Since 2015, LHCb operates an **offline-quality** real-time analysis strategy in HLT2, permitting a reduced event format **without loss** of physics performance

# What is ‘real time’?

> *The HLT [high-level trigger] application reconstructs the particle trajectories of the event in real time, where real time is defined as the interval between the collision in the detector and the moment the data are sent to permanent storage.*

“Tesla”, *Comput. Phys. Commun.* 208 (2016)

# What is ‘real time’?

> *The HLT [high-level trigger] application reconstructs the particle trajectories of the event in real time, where real time is defined as the interval between the collision in the detector and the moment the data are sent to permanent storage.*

“Tesla”, *Comput. Phys. Commun.* 208 (2016)

- $\Rightarrow$ This describes any HEP trigger!
- But HLT2 has a killer feature to permit offline quality

# Online data flow overview: Run 2

- First software trigger stage writes raw output to disk buffer
- Time available for 'real time' defined by buffer size

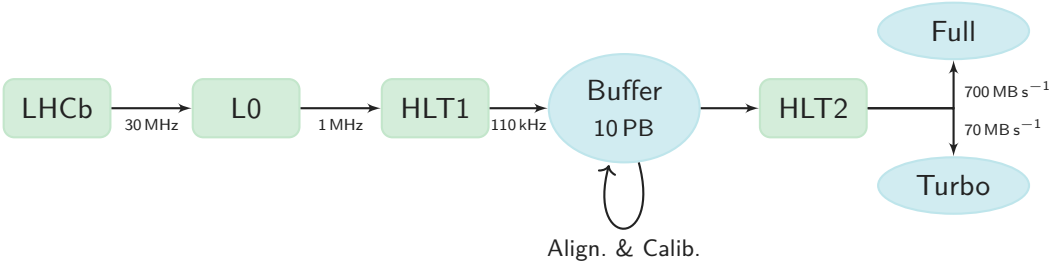


1. Alignment and calibration performed after HLT1, results available for second stage
2. Second stage runs asyncronously to the LHC
   - Increases effective HLT1 and HLT2 compute power
   - Additional time allows for offline reconstruction to be run
   - Can change HLT2 configuration *after* the data are taken!

# Why bother?

- Much smaller disk buffer in Run 1 (2011–2012) used only to increase effective compute power
- Run 1 HLT2 reconstruction was good, though not quite offline-equivalent
- Why did we go to all of this trouble in Run 2 (2015–2018)?

# Run 2 constraints

- We want to (must!) continue collecting the same physics as in Run 1

# Run 2 constraints

- We want to (must!) continue collecting the same physics as in Run 1
  - Plus whatever else we think of

# Run 2 constraints

- We want to (must!) continue collecting the same physics as in Run 1
  - Plus whatever else we think of
- At proton-proton $\sqrt{s} = 13\,\text{TeV}$, this includes our 'core' programme:

$$\text{Charm @ } 1\,\text{MHz}, \;\; \text{Beauty @ } 45\,\text{kHz}$$

# Run 2 constraints

- We want to (must!) continue collecting the same physics as in Run 1
  - Plus whatever else we think of
- At proton-proton $\sqrt{s} = 13\,\text{TeV}$, this includes our 'core' programme:

$$\text{Charm @ } 1\,\text{MHz}, \;\; \text{Beauty @ } 45\,\text{kHz}$$

- This must fit within our computing resources, principally disk space:

$$\text{Trigger bandwidth } [\text{MB}\,\text{s}^{-1}] \propto \text{Trigger rate } [\text{kHz}] \times \text{Average event size } [\text{kB}]$$

# Run 2 constraints

- We want to (must!) continue collecting the same physics as in Run 1
  - Plus whatever else we think of
- At proton-proton $\sqrt{s} = 13\,\text{TeV}$, this includes our 'core' programme:

$$\text{Charm @ } 1\,\text{MHz}, \;\; \text{Beauty @ } 45\,\text{kHz}$$

- This must fit within our computing resources, principally disk space:

$$\text{Trigger bandwidth } [\text{MB}\,\text{s}^{-1}] \propto \text{Trigger rate } [\text{kHz}] \times \text{Average event size } [\text{kB}]$$

- Raw detector information for one event weighs $\sim 60\,\text{kB}$, so:

$$\text{Charm} + \text{Beauty} \approx 60\,\text{GB}\,\text{s}^{-1}$$

# Run 2 constraints

- We want to (must!) continue collecting the same physics as in Run 1
  - Plus whatever else we think of
- At proton-proton $\sqrt{s} = 13\,\text{TeV}$, this includes our 'core' programme:

$$\text{Charm @ } 1\,\text{MHz}, \;\; \text{Beauty @ } 45\,\text{kHz}$$

- This must fit within our computing resources, principally disk space:

$$\text{Trigger bandwidth } [\text{MB}\,\text{s}^{-1}] \propto \text{Trigger rate } [\text{kHz}] \times \text{Average event size } [\text{kB}]$$

- Raw detector information for one event weighs $\sim 60\,\text{kB}$, so:

$$\text{Charm} + \text{Beauty} \approx 60\,\text{GB}\,\text{s}^{-1}$$

- In reality $\epsilon_{\text{Sig.}} < 100\,\%$, however $\epsilon_{\text{Bkg.}} > 0$

# Run 2 constraints

- We want to (must!) continue collecting the same physics as in Run 1
  - Plus whatever else we think of
- At proton-proton $\sqrt{s} = 13\,\text{TeV}$, this includes our 'core' programme:

$$\text{Charm @ } 1\,\text{MHz}, \;\; \text{Beauty @ } 45\,\text{kHz}$$

- This must fit within our computing resources, principally disk space:

$$\text{Trigger bandwidth } [\text{MB}\,\text{s}^{-1}] \propto \text{Trigger rate } [\text{kHz}] \times \text{Average event size } [\text{kB}]$$

- Raw detector information for one event weighs $\sim 60\,\text{kB}$, so:

$$\text{Charm} + \text{Beauty} \approx 60\,\text{GB}\,\text{s}^{-1}$$

- In reality $\epsilon_{\text{Sig.}} < 100\,\%$, however $\epsilon_{\text{Bkg.}} > 0$
- In Run 2, 2015–2018, it seems we **cannot afford** to save every trigger of interest

# Run 2 constraints

- We want to (must!) continue collecting the same physics as in Run 1
  - Plus whatever else we think of
- At proton-proton $\sqrt{s} = 13\,\text{TeV}$, this includes our 'core' programme:

$$\text{Charm @ } 1\,\text{MHz}, \;\; \text{Beauty @ } 45\,\text{kHz}$$

- This must fit within our computing resources, principally disk space:

$$\text{Trigger bandwidth } [\text{MB}\,\text{s}^{-1}] \propto \text{Trigger rate } [\text{kHz}] \times \text{Average event size } [\text{kB}]$$

- Raw detector information for one event weighs $\sim 60\,\text{kB}$, so:

$$\text{Charm} + \text{Beauty} \approx 60\,\text{GB}\,\text{s}^{-1}$$

- In reality $\epsilon_{\text{Sig.}} < 100\,\%$, however $\epsilon_{\text{Bkg.}} > 0$
- In Run 2, 2015–2018, it seems we **cannot afford** to save every trigger of interest
- But there's more!

# The LHCb upgrade (or: even tighter constraints)

- Comprehensive detector, DAQ, and online upgrade for beginning of Run 3, 2021

# The LHCb upgrade (or: even tighter constraints)

- Comprehensive detector, DAQ, and online upgrade for beginning of Run 3, 2021
- Instantaneous luminsosity increasing by a **factor 5**

$$L_{\text{Inst.}} = 2 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$$

# The LHCb upgrade (or: even tighter constraints)

- Comprehensive detector, DAQ, and online upgrade for beginning of Run 3, 2021
- Instantaneous luminsosity increasing by a **factor 5**

$$L_{\text{Inst.}} = 2 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$$

- The heavy flavour rates increase proportionally

$$R_{b\bar{b}} = 200\,\text{kHz}$$

# The LHCb upgrade (or: even tighter constraints)

- Comprehensive detector, DAQ, and online upgrade for beginning of Run 3, 2021
- Instantaneous luminsosity increasing by a **factor 5**

$$L_{\text{Inst.}} = 2 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$$

- The heavy flavour rates increase proportionally

$$R_{b\bar{b}} = 200\,\text{kHz}$$

- The average upgrade raw event size is $\sim 100\,\text{kB}$, and so

$$\begin{aligned}\text{Beauty bandwidth} &\propto \text{HLT output rate} \times \text{Average event size},\\ &\approx 2\,\text{GB}\,\text{s}^{-1}\end{aligned}$$

# The LHCb upgrade (or: even tighter constraints)

- Comprehensive detector, DAQ, and online upgrade for beginning of Run 3, 2021
- Instantaneous luminsosity increasing by a **factor 5**

$$L_{\text{Inst.}} = 2 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$$

- The heavy flavour rates increase proportionally

$$R_{b\bar{b}} = 200\,\text{kHz}$$

- The average upgrade raw event size is $\sim 100\,\text{kB}$, and so

$$\begin{aligned}\text{Beauty bandwidth} &\propto \text{HLT output rate} \times \text{Average event size},\\ &\approx 2\,\text{GB}\,\text{s}^{-1}\end{aligned}$$

- This might *just* fit into available resources, but with no room for anything else!

# The LHCb upgrade (or: even tighter constraints)

- Comprehensive detector, DAQ, and online upgrade for beginning of Run 3, 2021
- Instantaneous luminsosity increasing by a **factor 5**

$$L_{\text{Inst.}} = 2 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$$

- The heavy flavour rates increase proportionally

$$R_{b\bar{b}} = 200\,\text{kHz}$$

- The average upgrade raw event size is $\sim 100\,\text{kB}$, and so

$$\begin{aligned}\text{Beauty bandwidth} &\propto \text{HLT output rate} \times \text{Average event size},\\ &\approx 2\,\text{GB}\,\text{s}^{-1}\end{aligned}$$

- This might *just* fit into available resources, but with no room for anything else!
- In Run 3 it seems we **cannot afford** to have every trigger selection we had in Run 2

# The LHCb upgrade (or: even tighter constraints)

- Comprehensive detector, DAQ, and online upgrade for beginning of Run 3, 2021
- Instantaneous luminsosity increasing by a **factor 5**

$$L_{\text{Inst.}} = 2 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$$

- The heavy flavour rates increase proportionally

$$R_{b\bar{b}} = 200\,\text{kHz}$$

- The average upgrade raw event size is $\sim 100\,\text{kB}$, and so

$$\begin{aligned}\text{Beauty bandwidth} &\propto \text{HLT output rate} \times \text{Average event size},\\ &\approx 2\,\text{GB}\,\text{s}^{-1}\end{aligned}$$

- This might *just* fit into available resources, but with no room for anything else!
- In Run 3 it seems we **cannot afford** to have every trigger selection we had in Run 2
  - Or can we?

# Beyond the raw event

## The idea

- In general, raw event is only needed for the centralised reconstruction
  - Not used in analyses directly
- We could save a **subset** of the information from the event!
- Raw information hard to split usefully, but **reconstructed objects** often independent
  - Can discard raw information after the reconstruction
- We have **offline quality candidates** in HLT2, often fully reconstructed
  - Can perform full physics measurements on these directly
- Extract desired objects from trigger reconstruction, save **only these**
  - Can require careful bookkeeping between (often implicit) object dependencies

# Turbo

Comput. Phys. Commun. **208** (2016)

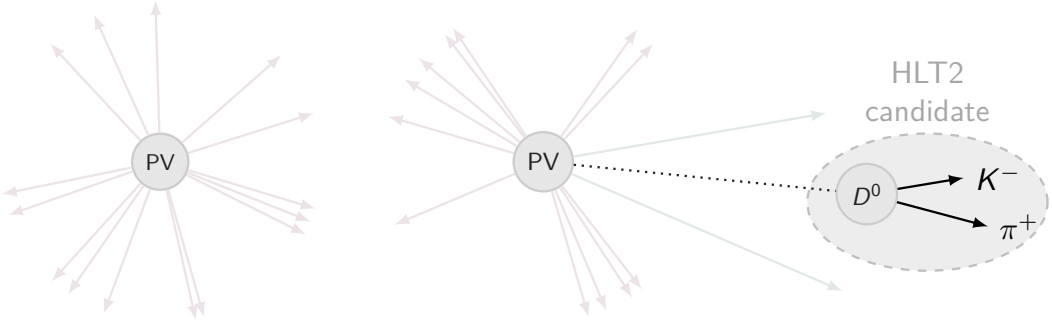


- Turbo (2015): Save only the objects involved in the trigger decision, plus all PVs

# Turbo

Comput. Phys. Commun. **208** (2016)

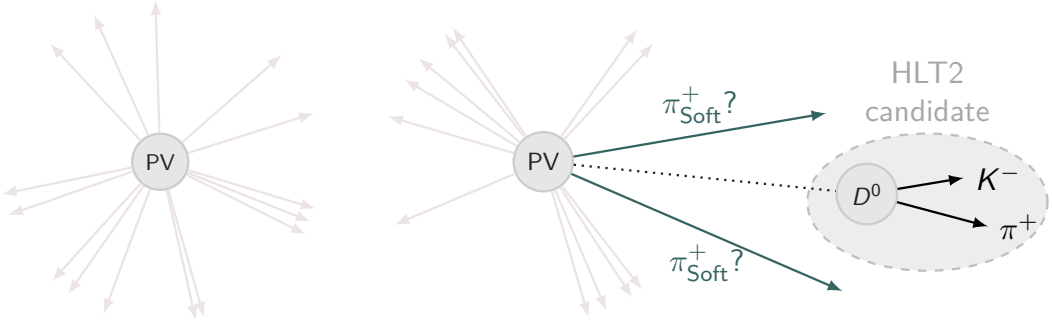


- Turbo (2015): Save only the objects involved in the trigger decision, plus all PVs
- TurboSP (2017): Completely flexible specification of additional objects to save

# Turbo

Comput. Phys. Commun. **208** (2016)

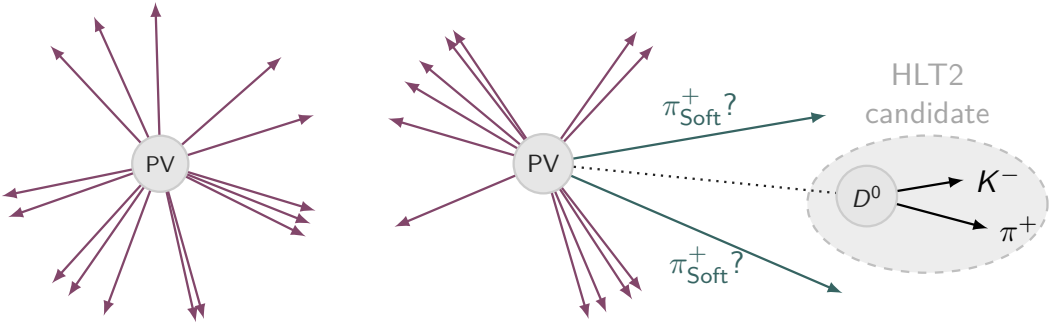


- Turbo (2015): Save only the objects involved in the trigger decision, plus all PVs
- TurboSP (2017): Completely flexible specification of additional objects to save
- Turbo++ (2016): The entire HLT2 reconstruction

# Bandwidth savings

- Typical Turbo event weighs a few kilobytes, Turbo++ around 50 kB
  - TurboSP anywhere in between
  - Raw event around 60 kB
- Today, Turbo lines account for 25 % of the trigger rate, 10 % of bandwidth
  - Allows for 20x increase in rate for a given bandwidth
- Addition of TurboSP in 2017 reduced Turbo bandwidth by 50 %

| Stream | Rate [kHz] | Bandwidth [MB $s^{-1}$] |
|---|---|---|
| Turbo (2016) | 3.5 | 140 |
| Turbo (2017) | 2.5 | 70 |
| Full (2017) | 10 | 700 |

# Bandwidth savings

- Typical Turbo event weighs a few kilobytes, Turbo++ around 50 kB
  - TurboSP anywhere in between
  - Raw event around 60 kB
- Today, Turbo lines account for 25 % of the trigger rate, 10 % of bandwidth
  - Allows for 20x increase in rate for a given bandwidth
- Addition of TurboSP in 2017 reduced Turbo bandwidth by 50 %

| Stream | Rate [kHz] | Bandwidth [MB $s^{-1}$] |
|---|---|---|
| Turbo (2016) | 3.5 | 140 |
| Turbo (2017) | 2.5 | 70 |
| Full (2017) | 10 | 700 |

(This approach also saves money, as post-trigger processing is minimal)

# Physics enabled by Turbo today

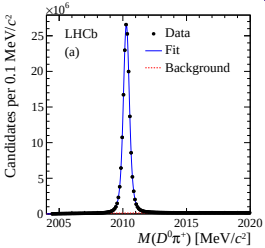


Phys. Rev. D **97**, 031101 (2018)

- Rich charm programme not feasible in Run 2 without Turbo
- Can collect other high-rate physics, e.g. prompt dimuons

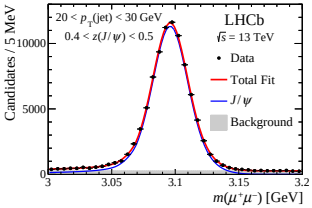


Phys. Rev. Lett. **118**, 192001 (2017)

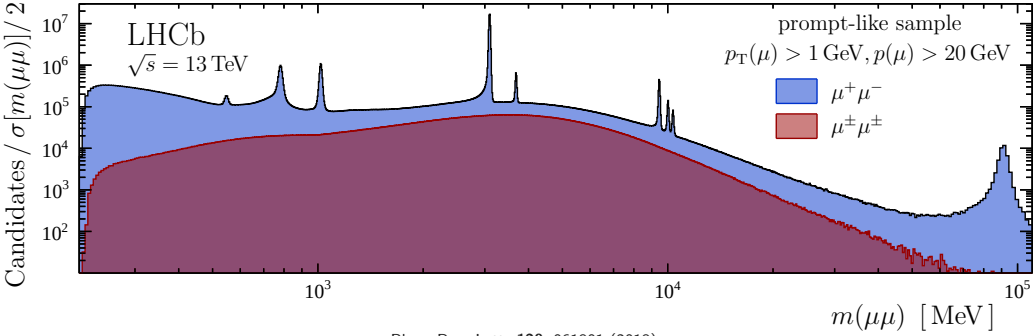


Phys. Rev. Lett. **120**, 061801 (2018)

# Physics enabled by Turbo in Run 3

- Foreseen that most analyses will work with a reduced event format
- Important to study TurboSP as a bridge between exclusive and inclusive triggers

- Promising studies ongoing to select $B$ decay tracks using MVA techniques

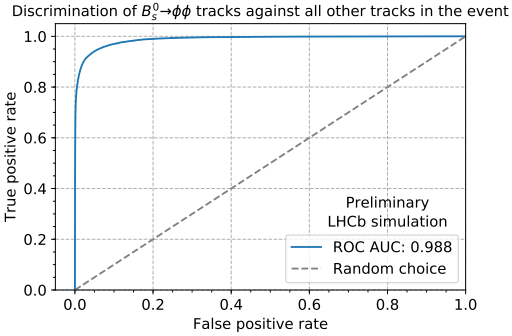

# Summary

- Since 2015, LHCb has offline-quality reconstruction in its final trigger stage
  - Thanks to a disk buffer, a real-time alignment and calibration procedure, and a speed-up of the reconstruction software
- This has enabled trigger objects to be saved directly
  - No loss in physics performance
  - No need for substantial subsequent processing, e.g. a second offline reconstruction
- Enables physics that would otherwise not be permissible within computing resources
- Recent updates to this, TurboSP, allow the model to accommodate full programme
- Crucial for the entire upgrade programme, promising studies now ongoing

# Trigger scheme: Run 1

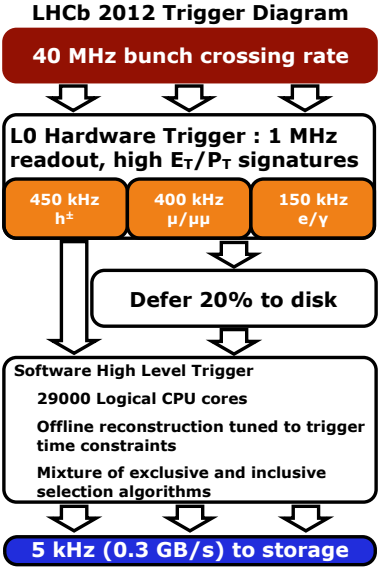

# Trigger scheme: Run 2 (current)

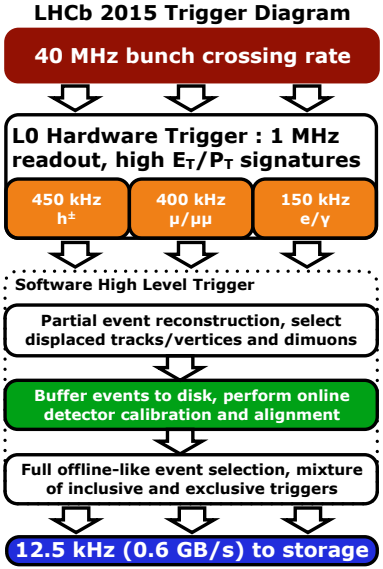

# Trigger scheme: Run 3 (upgrade)

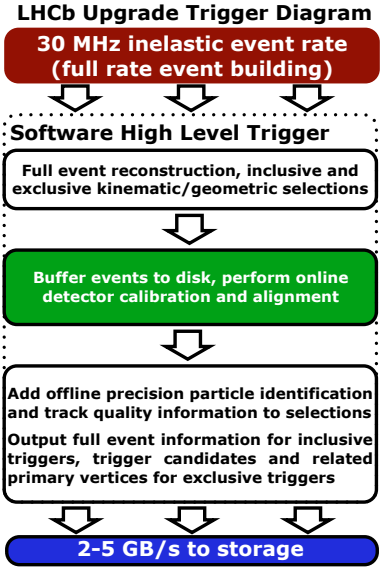

# What about physics performance?

- We care about high rate processes, but also precision
  - Allows for strong background suppression with minimal signal bias
- With Turbo there's no going back, so we must do the best we can upfront
- There is **no loss of physics performance** with Turbo in Run 2

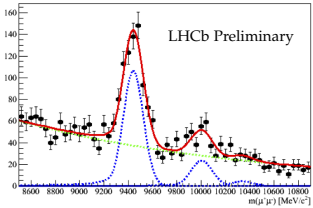


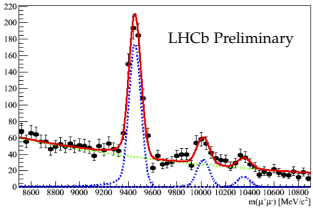


⇒

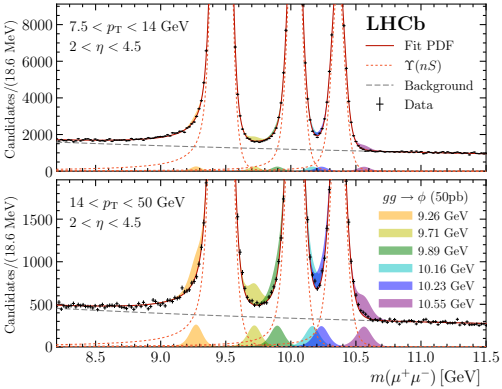


arXiv:1805.09820

# Alignment and calibration

- Full alignment and calibration of detector run every fill
- If an update is required, new constants available for HLT2 immediately
- Completely automated

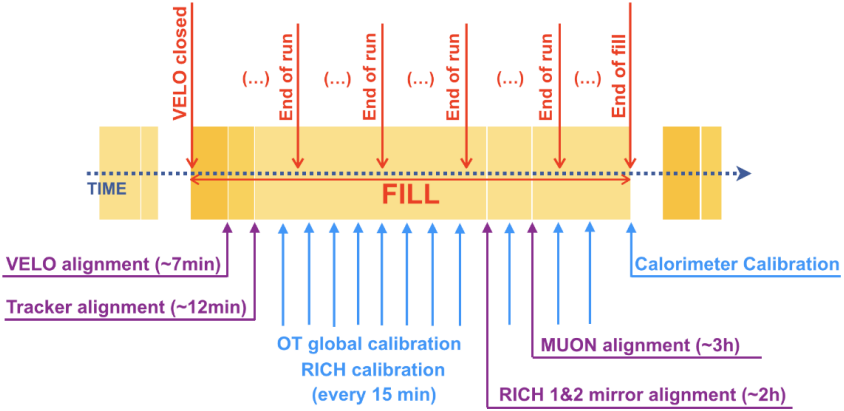


(...) - time needed for both a data accumulation and running the task

# Example: VELO alignment

- VELO centres itself around beam at start of each fill
- Align using a Kalman filter approach, using track hit residuals with PV constraints
- Automatic alignment of VELO halves takes around 10 minutes

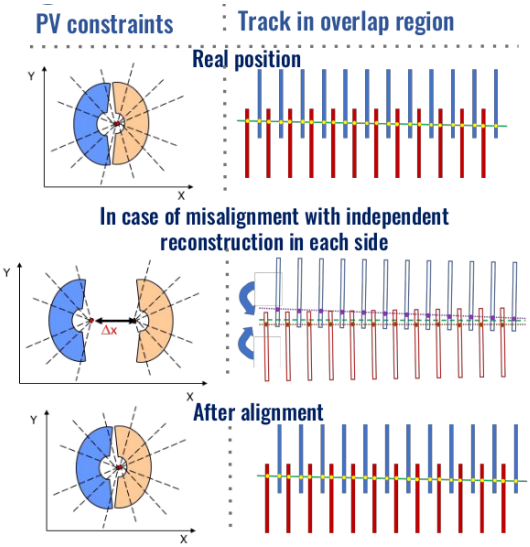

# 2018 operation

- Alignment and calibration procedure operational as in 2017
- Stable conditions with time, only occasional updates needed