# Status of LHCb Upgrade I and future prospects

CERN-RRB-2024-093

LHCb collaboration

29th October 2024

# Contents

# 1 Introduction

The LHCb experiment will have three major phases of detector technology. The original detector completed data taking at the end of LHC Run 2 in 2018. Physics results arising from the analysis of data collected in this phase are described in this document. The LHCb Upgrade I was primarily installed during LHC Long-shutdown 2 (LS2) in 2018-2022. The final subsystem, the Upstream Tracker (UT), was put in place during YETS 2022/23, successfully completing the installation. The LHCb Upgrade II is planned to be primarily installed in LS4, with some smaller modifications foreseen in LS3. Progress on the R&D for this detector is also described in this document.

A total of $10\,\text{fb}^{-1}$ were delivered to LHCb in Run 1 and Run 2 data taking periods, with $9\,\text{fb}^{-1}$ recorded. The LHCb Run-1 and Run-2 dataset comprises $pp$, $p$Pb, and PbPb at various centre-of-mass energies, as well as $p$A (A = He, Ne, Ar) collisions in fixed target mode, using the experiment's unique gas injection system. Exploitation of Run 1 and Run 2 data is still progressing.

The LHCb Upgrade I is arguably the largest CERN particle-physics detector project to begin operation since the completion of the LHC. It has been completed on-budget. All subsystems have been operated during 2023 and their commissioning has successfully proven their functionality. An incident occurred in the LHC vacuum system of the VELO detector during YETS 2022/23 which has damaged the aluminium foil that separates the secondary and primary vacuum volumes. This has prevented the closure of the VELO to its minimum aperture in the 2023 run and the foil has been replaced in the YETS 2023/24.

The Upgrade I detector is based on a novel trigger system able to read out all sub-detectors at 40 MHz and to select physics events of interest by means of a pure software trigger at the bunch crossing rate of the LHC. The software trigger allows the experiment to collect data with high efficiency at a luminosity of up to $2 \times 10^{33}\,\text{cm}^{-2}\text{s}^{-1}$. Flavour-physics measurements are going to be performed with higher precision than was possible with the previous detector and across a wider range of observables. The flexibility inherent in the new trigger scheme will also allow the experiment to further diversify its physics programme into important areas beyond flavour. During 2024, the Upgrade I detector has been operated in nominal conditions, recording more than $9\,\text{fb}^{-1}$ of integrated luminosity and surpassing the combined luminosity of Run 1 and Run 2 in a single year.

The Upgrade I detector was proposed in the Letter of Intent [1] in 2011, and its main components and cost-envelope were defined in the Framework TDR [2] one year later. Technical Design Reports (TDRs) were written for all systems [3–11] and approved by the Research Board. Addenda to the Memorandum of Understanding (MoU) were presented to the RRB in April and October 2014, covering the division of resources and responsibilities for Common Project items [12] and sub-system items [13], respectively. Enhancements to the ECAL and RICH systems are proposed for LS3 and were approved in December 2023. A TDR for additional enhancements to the Online system in LS3 was submitted to the LHCC

in February 2024, and approved in September 2024.

The LHCb Upgrade II detector was proposed in the Expression of Interest [14] in 2017 and the Physics Case described in [15]. Its main components and cost-envelope were defined in the Framework TDR [16] that was approved in 2022. A Scoping Document, describing various scoping scenarios in terms of cost and physics performances, was submitted to the LHCC in September 2024, and is now under review. The content of the Scoping Document will be presented in a dedicated session at this RRB meeting.

# 2 Collaboration matters

Eight new institutions and three new countries (Chile, Kazakhstan and Lithuania) have joined LHCb since the April 2024 RRB. Here is the complete list

- University of Freiburg, as Full Member;
- University of Vilnius, as Associate Member, associated through the University of Freiburg;
- Andres Bello University of Chile, as Associate Member, associated through the University of Zurich;
- GSI, as Technical Associate Member, associated through the University of Oxford;
- Subatech Nantes, as Technical Associate Member, associated through ICJLab;
- Institute of Nuclear Physics of Kazakhstan, as Technical Associate Member, associated through CERN;
- Jagiellonian University, as Technical Associate Member, associated through AGH University;
- University of Naples and INFN, as Technical Associate Member, associated through LNF.

In addition, IRFU CEA Saclay, which was an Associate Member with LLR as host, became a Full Member.

# 3 Project organisation

The management of the experiment is composed of the Spokesperson and two deputies, the Technical Coordinator, the Resource Coordinator and the Software and Computing Coordinator.

The Upgrade I detector is overseen by the LHCb Technical Board. The Technical Board is chaired by the Technical Coordinator and is comprised of all detector

project leaders. Day-by-day detector work is overseen by the technical coordination team. The commissioning and operations of the Upgrade I detector are overseen by the Run coordinator and discussed at Run meetings, with representation from all projects. Physics aspects are overseen by the Physics Coordinator and discussed at the Physics Planning Group, whereas Operational ones are overseen by the Operations Coordinator and discussed at the Operations Planning Group.

The Software and Computing Coordinator chairs the Software and Computing Board, which includes the project leaders of the software and computing projects, and the software coordinators of the various subdetectors.

Because the Run 3 LHCb detector is fully read out into a software trigger, the first level of this trigger (HLT1) plays a critical role in monitoring the data taking performance and selecting events based on technical triggers which are then used for specific subdetector monitoring tasks. Additionally much of the online and offline monitoring and data quality, as well as subdetector-specific monitoring, now relies on specific configurations of the second-level software trigger (HLT2). This has led to a greater degree of interdependence between the systems than had been the case in the previous LHCb experiment. Consequently, a subcommittee of the technical board (TBSC) is in place during 2024 to strengthen communication between the Real Time Analysis project, Online, and the LHCb subdetectors around software tools and components of common interest. An additional deputy Technical Coordinator was appointed for a period of one year with a mandate to chair this subcommittee.

Activities related to LHCb Upgrade II are overseen by the Upgrade II Planning Group, chaired by the Upgrade II Coordinator and composed of several members with specific responsibilities in the various areas of physics performance, hardware and software development.

# 4 Financial matters

According to the financial plan for the M&O Cat. A levels, which was finalised, submitted and approved by the Scrutiny Group and by the RRB in 2017 and 2018, the proposed and subsequently approved budget for M&O Cat. A is 3,070 kCHF for the year 2024. Throughout this year, the full detector has been operating at high efficiency. It started with the well-prepared repair of the VELO detector. This was complemented by important early expenditures, in particular for the Online and technical infrastructure to maximise operational efficiency. We do not expect any large expenditures for the remainder of the year but we again note increasing SLA costs in many areas, which motivates us to ask for a modest budget increase for 2025. The proposed budget for 2025, as discussed with the Scrutiny Group, is therefore increased to 3,120 kCHF, an increase of 50 kCHF (1.6 per cent of the overall budget).

We continue to study all options to face the effects of the war of the Russian Federation against Ukraine on our experiment and its resources. The proposed

resource sharing for 2024 includes contributions from Russian funding agencies and institutes. Following the resolution of December 2023, CERN Council will not extend the International Collaboration Agreement with Russia at its expiration date (November 2024). This decreases the collaboration membership and increases the technical and financial requirements on the remaining collaborating institutes. The effect is partially mitigated by new institutes joining. Thus, the average increase in Member States contributions is evaluated at 9.8 per cent for M&O Cat. A (see also CERN-RRB-2024-099).

The funding requirements for the LHCb Upgrade construction have been defined in detail in Addendum No. 1 to the Memorandum of Understanding (MoU) for Common Projects [12] and in the Addendum No. 2 to the MoU for the Upgrade of the Sub-Detector Systems [13], which refer to the LHCb Upgrade Framework Technical Design Report [2] and the Technical Design Reports [3–8] for all Upgrade subdetectors. At present, the LHCb Upgrade project continues to progress well. All of the sub-detectors have been installed and are in their operational phases. Spending of core funds is essentially over. The final Common Project funds have been spent in 2020-2024 and are now also finished. In order to ensure the optimal value-for-money and performance for the experiment, we have now acquired a new slice of the computing farm.

The Upgrade I construction project is thus ended and has been achieved within the agreed cost envelope. No request for further funds has been put forward and the RRB has agreed to close the books in the meeting held in October 2023.

# 5 Operations

In preparation of the 2024 activities, the Operations Planning Group had established a detailed plan for the main operational activities in 2024. It included milestones and timelines for all activities relevant for data taking, simulation and data processing. This plan has been followed successfully in the last six months which also helped to establish 2024 as a very successful year in terms of data taking for LHCb. When necessary the plan was adapted to the current running conditions. As of October 14, in total a dataset corresponding to a luminosity of 9.39 $\text{fb}^{-1}$ has been recorded and processed by the experiment as shown in Fig. 1. Combined with the improvements in the trigger compared to the Run-1 and Run-2 detector, this dataset will enable measurements with a statistical precision surpassing that of Run 1 and Run 2.

Concerning data collected in 2024, they were processed following the computing model foreseen for the LHCb Upgrade [17]. As reminded in Fig. 2, after data are collected and processed in real-time by the first stage trigger based on GPUs (HLT1), they are stored in a buffer, waiting for an update of the alignment and calibration constants. Subsequently, they are processed by the second stage of trigger based on CPUs (HLT2). The data are then sent offline and undergo the concurrent Sprucing process, i.e. the centralised skimming, trimming and stream-

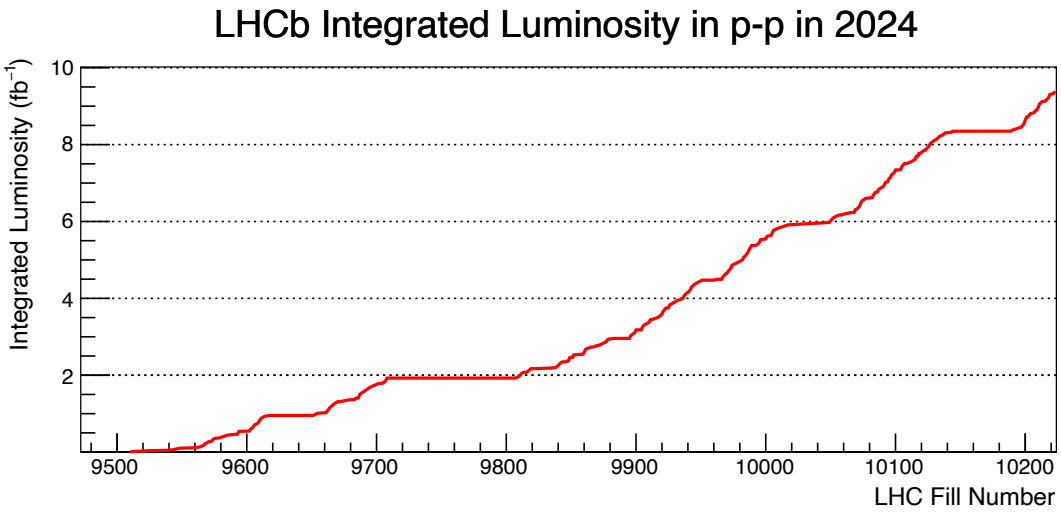


Figure 1: Recorded luminosity by the LHCb experiment in 2024.

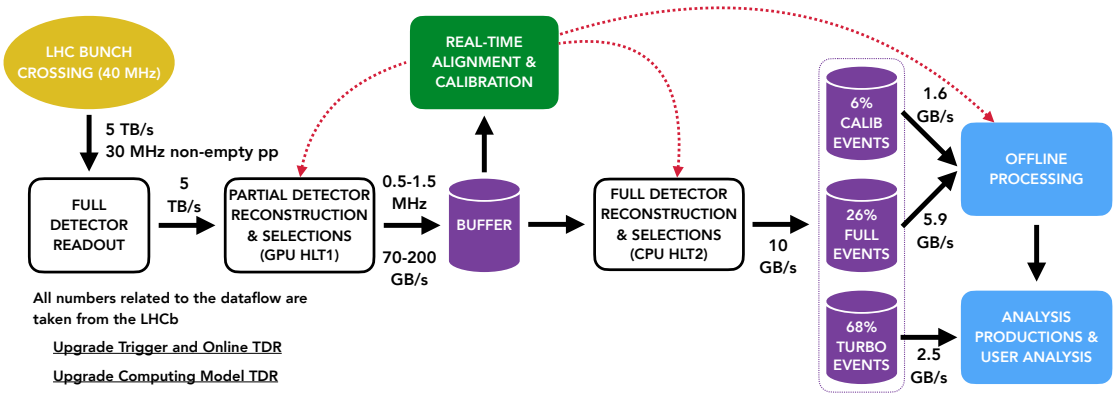


Figure 2: Schematic view of the online processing for the LHCb Upgrade I.

ing of the data, before being made available to the analysts (see Fig. 3).

At the beginning of the year deadlines for the various software releases needed in the data processing chain had been defined and they were followed later. Data recorded by HLT1 were quickly processed by HLT2 and subsequently by the Sprucing. Analysts have been using the new centralised analysis production scheme on spruced data for producing ntuples and performing performance and physics studies. This allowed for quick feedback, sometimes within one week, by analysts on the quality of the data concerning alignment, reconstruction and trigger efficiencies.

Feedback was often provided in the so-called General Performance Meeting, held every week. The meeting, chaired by Operations Coordination, Physics Coordination and other members of the Operations Planning Group, has been established as a meeting to exchange information between analysts, detector, reconstruction and trigger experts in 2024.

The quality of the data is assessed online and offline. The offline data quality

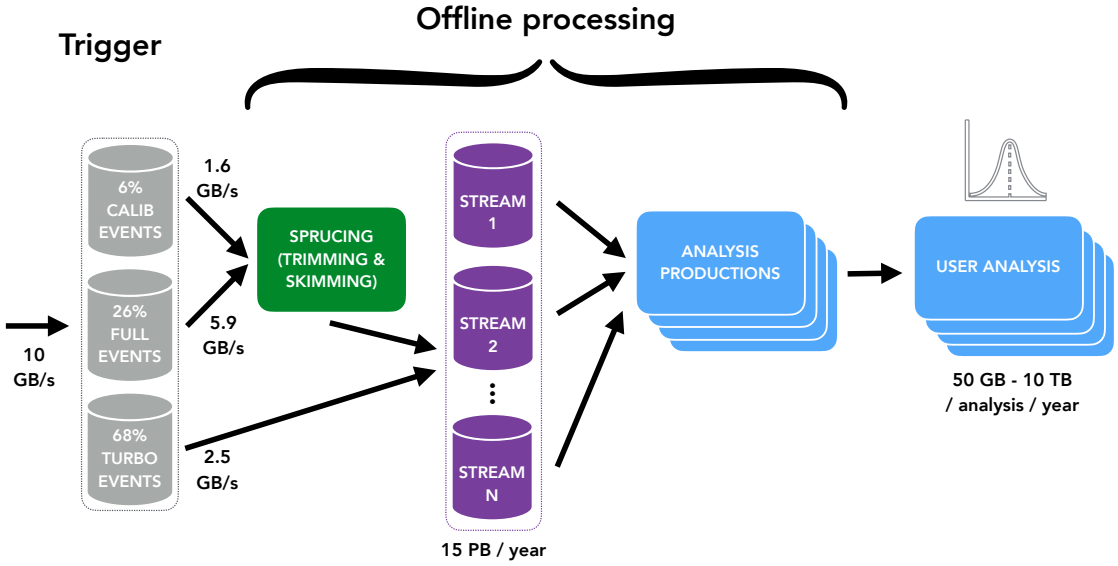


Figure 3: Schematic view of the offline processing for the LHCb Upgrade I.

team, using inputs from online monitoring and the full event reconstruction in HLT2, characterises the data as suitable (good) or not (bad) for physics analyses. At the beginning of the run a small fraction of 2024 data was declared as bad for analyses, but later, more than 99% of the data was rated good. Some data have been intentionally left unchecked as they are not necessarily bad but show features which make them different to other data in that period, *e.g.* fills used to commission running at higher instantaneous luminosity, requiring more thorough scrutiny at a later stage.

The dataset recorded in 2024 can be grouped into different phases, each having a size comparable to a year of running in Run 1 or 2. They are characterised by an increasing instantaneous luminosity from phase to phase, the inclusion of the UT after the June Technical Stop, and improved alignments deployed throughout the year. The increase in instantaneous luminosity also required improvements in the reconstruction and trigger algorithms to cope with the increase. Here the inclusion of the UT in the reconstruction helped to reduce the amount of fake tracks and enabled faster algorithms as the UT can provide a quick momentum estimate. The algorithmic improvements were accompanied by a significant increase of the compute power both in HLT1 and HLT2 after the installation of additional GPUs and CPUs, respectively. More details can be found in subsequent chapters.

One challenge throughout the year was to adapt event selection algorithms in HLT1, HLT2 and Sprucing to the changing conditions while maintaining a high efficiency and respecting the limits defined in Figs. 2 and 3. At several checkpoints which had been defined during the planning stage by the Operations Planning Group the resource usage was reviewed and the necessary actions to tune trigger selections were performed. In addition, the average event size was significantly reduced by reducing the amount of redundant information and implementing a

better compression. On average, the limits were respected throughout the year but there were periods where especially the bandwidth after the Offline processing was too high. Subsequently, better monitoring procedures were defined and implemented.

The improvements in all areas culminated in a data taking at design luminosity with all detectors included, a trigger with a significantly better efficiency compared to Run 2 and mass resolutions approaching the Run 2 performance.

The preparation of the Heavy Ion data taking including a proton-proton reference run at the end of 2024 is mostly concluded. The plan is to record lead-lead collision with minimum bias triggers and in parallel to use the SMOG2 storage cell and record fixed-target lead-neon and lead-argon collisions. In contrast to the proton-proton strategy, the full detector information for hadronic lead-lead collisions will be written out to allow for improvements of the reconstruction of high occupancy events offline. As fixed-target collisions have a higher rate but smaller occupancy, there is no need to make improvements in the event reconstruction offline, and to increase the amount of data recorded and maximise the physics output the raw event will be discarded for SMOG2 events.

The process of validating and tuning the description of the sub-detectors in the simulation framework is actively ongoing. Simulations with improved detector description were made available throughout the year, with more improvements expected towards the end of the year. Options for different trigger and reconstruction configurations have been provided and it is expected that many analyses will request simulated samples in the coming months.

The priorities of the next months are focused on providing the necessary inputs for physics measurements with the 2024 data. These include simulations for the different phases of data taking, a tool for analysts to access the luminosity from a web-based database, the necessary tools to apply data-driven corrections, and the reprocessing of the Full stream of 2024 with additional selections in the Sprucing. In parallel the planning for 2025 will start.

# 6 Status of 2024 run

The LHC schedule was adjusted at the start of the year, shifting four weeks of *pp* running at 13.6 TeV from 2025 to 2024. This change results in 124 days of 25 ns proton physics with the LHC in luminosity production and 23 days dedicated to *pp*-reference and PbPb running.

The bunch intensity was limited to $1.6 \times 10^{11}$ protons per bunch to mitigate risks associated with degraded RF contact fingers in warm modules. This intensity is sufficient to keep LHCb levelled at its nominal upgrade luminosity ($\mathcal{L} = 2 \cdot 10^{33}$ cm$^{-2}$s$^{-1}$) for more than 15 hours. The reduced bunch intensity also provided the benefit of lowering the load on the cryogenic system. The train length, constrained by cryogenic margins, particularly in sector 78, and closely linked to the maximum number of circulating bunches, was set accordingly to the

3x36b scheme. This avoided the hybrid scheme, which would have reduced the number of colliding bunches in LHCb. Under these conditions, the nominal filling scheme delivered by the LHC included 2352 proton bunches per beam and 2133 colliding bunches in LHCb, enabling the achievement of the nominal instantaneous luminosity with a pile-up of 5.3. The total input rate to the Event Builder farm was approximately 27 MHz. The switch from the Standard 25 ns to the Batch, Compression, Merging and Splitting (BCMS) beams, implemented by the LHC during week 21, had no impact on LHCb.

Before entering the luminosity production phase, LHCb adopted a strategy of gradually increasing the pile-up alongside the LHC intensity ramp-up. This schedule was notably compressed in 2024, as the LHC began the luminosity production phase on 14 April, eleven days ahead of schedule. The planning was to ensure that the subsystems were ready to operate efficiently and at their best possible performance at nominal luminosity, to check the stability of the frontend electronics and data acquisition system and to create contingency for any unforeseen issues. During this phase, low-level performance metrics, such as subdetector occupancy, hit or cluster linearity, frontend rates, high-voltage currents, single-event upset rates, and data acquisition rates, were closely monitored by each subdetector to anticipate potential problems at nominal luminosity. No critical issues were identified at this stage.

The VELO fully closed around the interaction region during the first step of the proton-proton intensity ramp-up on 5th of April, confirming the successful recovery achieved during the 2023/2024 Year-End Technical Stop following the LHC vacuum incident at the start of 2023. The automated per-fill closure procedure, managed by the central shifters, has been commissioned and is since then in operation. A residual movement of the C-side VELO half was observed, prompting a modification of the closing procedure to mitigate this drift. Additionally, an automated procedure for correcting this issue through software alignment was implemented and integrated into the controls. The first three fills with stable beams at 13.6 TeV were dedicated to subdetector local calibrations, addressing some performance deficits identified in the 2023 data. Subsequently, global running was carried out to check the HLT1 and HLT2 throughputs, output bandwidths, spatial alignment status, and to collect samples for promptly assessing performance with the improved calibrations of the various subsystems.

Full machine operation was achieved on the 26th of April. During the luminosity production phase prior to the first Technical Stop (TS1), calibration and physics production fills were alternated. Calibration fills were mainly focused on completing the commissioning of the UT and on addressing residual electronics instabilities that were limiting data-taking efficiency at nominal luminosity. This work was coordinated by the Detector Electronics Commissioning Task Force (DECTF). Additionally, HLT1, HLT2 and alignment and calibration procedures were commissioned at nominal luminosity with all subdetectors included, with interactions between subdetectors and RTA on these matters coordinated by the Technical Board Sub-Committee (TBSC). Physics production fills, taken at a pile-up of 3.3

without the UT being included in the global data-taking, enabled the commissioning of the full data processing chain and provided samples for analysts to evaluate and fine-tune detector performance. These samples included data where hydrogen, helium, neon, and argon were injected into the SMOG2 cell. Furthermore, the van der Meer programme was conducted to calibrate online and offline luminometers, and magnet polarity flips from Down to Up and vice-versa were carried out to assess detector asymmetries and performance in both configurations.

By TS1, LHCb was able to operate with all subdetectors except the UT in the global partition, achieving a DAQ efficiency of over 95% at nominal luminosity. Some yield saturation was observed initially for high-occupancy events, which was quickly mitigated owing to further improvements in the tracking sequence. The UT was able to operate up to a pile-up of 4.4 with good tracking performance, but further work was required to ensure stable operation without performance deficits.

The progress made during this phase enabled operation at an increased pile-up of 4.4 since 18th of June, hereafter referred to as post-TS1 period, with the UT included in global data-taking and in the track reconstruction at HLT2 level. From the 26th of August, the UT was also incorporated into the HLT1 track reconstruction. A substantial amount of data, approximately $6\,\text{fb}^{-1}$, was collected during this period, with gases injected into the SMOG2 cell for the majority of the time, and alternating magnet polarities to balance the datasets corresponding to different tracker alignment conditions. Thanks to the additional resources added to the HLT2 farm, the HLT1 output rate was maintained at approximately 1 MHz without the risk of overfilling the HLT1 storage buffer. Additionally, a 800 µm shift of the interaction region along the x-axis was implemented by the LHC to align it with the LHCb coordinate system.

During this phase, physics production was prioritised, with tests for running at nominal luminosity kept to a minimum. Calibration fills were focused on finding the optimal balance between UT frontend electronics and DAQ stability on one side, and tracking efficiency on the other. Additional 163 GPUs were installed during Machine Development 3 (MD3). Between MD3 and MD4, the new GPUs were integrated into the control system, and in parallel an optimal configuration for stable operation of the UT at nominal luminosity was identified.

Since MD4, ended on 1st of October, LHCb has been running at the upgrade design instantaneous luminosity with a data-acquisition efficiency of approximately 95%, consistent with post-TS1 values, as reported in Fig. 4. Additionally, downstream tracking has been incorporated into the HLT1 sequence. HLT1 output rate in these conditions is approximately 1.3 MHz.

An integrated luminosity of about $9.4\,\text{fb}^{-1}$ has been collected as of October 14, surpassing the total recorded luminosity of the combined Run-1 and Run-2 data samples, as shown in Fig. 5, thanks to the excellent LHCb DAQ efficiency and LHC performance in 2024. This was one of the main objectives outlined at the last RRB meeting. Enhancements in the online monitoring infrastructure have ensured full control over data quality, with nearly all runs flagged as suitable for physics analyses. Consequently, the goals of the 2024 *pp* run have been achieved,

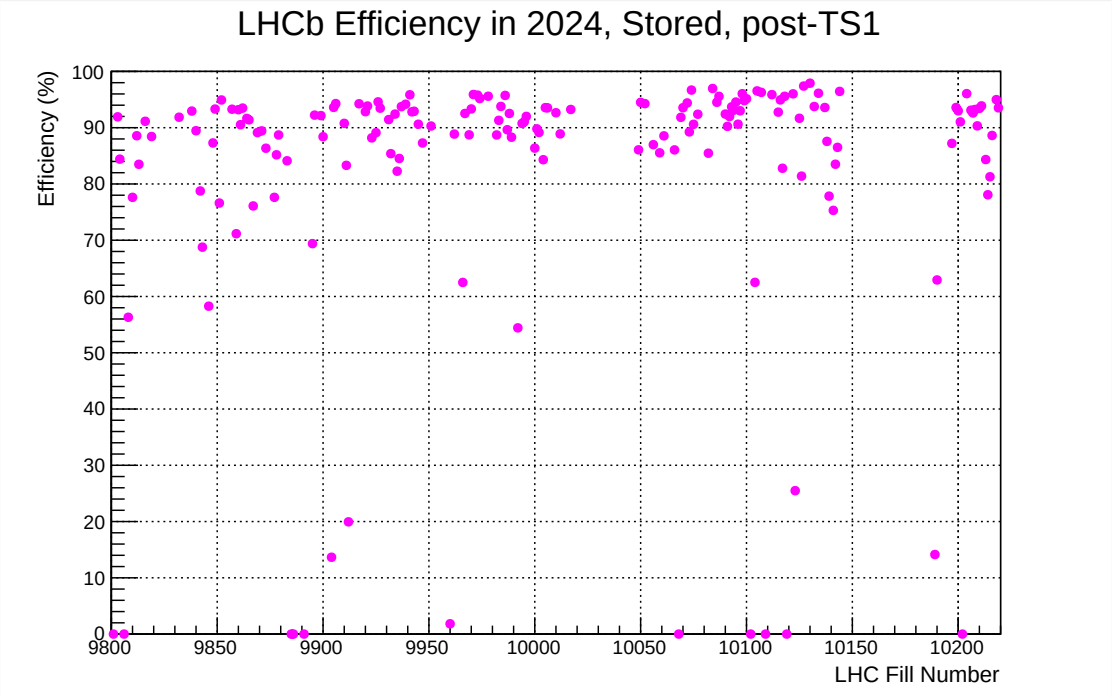


Figure 4: Data acquisition efficiency post-TS1. Lower efficiencies correspond either to isolated, occasional issues, dedicated tests or short LHC fills.

and LHCb is on track to meet and even exceed the Run-3 luminosity target.

The $pp$ run will be followed by a $pp$-reference run at 2.68 TeV, during which LHCb aims to collect around $100\,\text{pb}^{-1}$ of data. This will be followed by the heavy ion run, where LHCb expects to integrate approximately 0.3 $\text{nb}^{-1}$, benefiting from a 40% increase in the number of colliding bunches provided by the LHC and operations with magnet up polarity, which reduces the effective crossing angle. These factors together result in an estimated 70% increase in instantaneous luminosity with respect to 2023.

# 7 Physics

From April 2024 to October 2024 the LHCb collaboration submitted 29 new physics publications. This brings the number of LHCb physics publications to a total of 741 at the time of writing (not including a dozen of collaboration papers on physics performances), of which 723 are already published, Further 25 publications are being processed by the LHCb Editorial Board and are close to submission. All papers relative to this period are listed in Table 1.

In the following, selected results from recent publications are briefly highlighted.

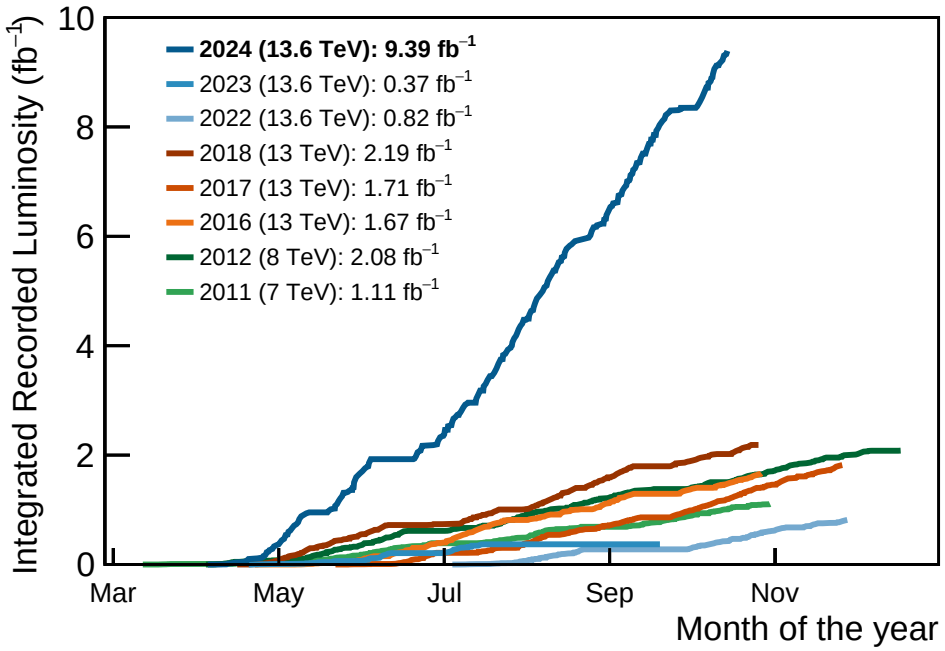


Figure 5: Integrated luminosity collected by the LHCb experiment, separated by data-taking year. Years characterised by negligible integrated luminosities, namely 2010 and 2015, are excluded from the plot.

## 7.1 Exotic hadronic states in beauty hadron decays

The study of $b$-hadron decays provides essential information about the underlying properties of the weak and strong interaction. Combinations of results obtained from multiple decays allow Standard Model parameters to be constrained, and may ultimately reveal inconsistencies possibly due to physics processes not accounted for in the Standard Model (new physics). LHCb has recently produced multiple results that help improve the overall understanding of the weak and strong interactions.

Hadronic states with more than three valence quarks were predicted by the quark model. The first tetraquark state was observed by the Belle collaboration in 2003, and the first pentaquark by the LHCb collaboration in 2015 [47]. The proximity of the observed pentaquark state masses with the open-charm $\Sigma_C \bar{D}^{(*)}$ mass thresholds favours loosely bound pentaquark models. To probe these thresholds, a search was performed by LHCb for four new $\Lambda_b^0 \to \Sigma_c^{(*)++} D^{(*)-} K^-$ decays, which may reveal the presence of pentaquark states in the decay amplitudes [20]. The four decay channels were observed, with $100 - 500$ signal candidates each. Although the yields are statistically insufficient to perform the amplitude analyses necessary to reveal the presence of pentaquarks, these results open the way for a complete analysis using the abundant Run-3 data set.

Table 1: Full list of LHCb publications submitted for publication from April to October 2024.

| | Title | arXiv | Ref. |
|---|---|---|---|
| 1. | Search for $B_s^0 \to \mu^+\mu^-\gamma$ decay | 2404.03375 | [18] |
| 2. | Search for prompt production of pentaquarks in open charm hadron final states | 2404.07131 | [19] |
| 3. | First observation of $\Lambda_b^0 \to \Sigma_c^{(*)++} D^{(*)-} K^-$ decays and measurement of their relative branching fractions | 2404.19510 | [20] |
| 4. | Amplitude Analysis of $B^+ \to D^{*-} D_s^+ \pi^+$ | 2405.00098 | [21] |
| 5. | Search for time-dependent CP violation in $D^0 \to \pi^+\pi^-\pi^0$ decays | 2405.06556 | [22] |
| 6. | Transverse polarisation measurement of $\Lambda$ hyperons in pNe collisions at $\sqrt{s_{NN}}$=68.4 GeV with the LHCb detector | 2405.11324 | [23] |
| 7. | Study of $b-$hadron decays to $\Lambda_c^+ h^- h'^-$ final states | 2405.12688 | [24] |
| 8. | Search for the lepton flavor violating decay $B_s^0 \to \phi\mu\tau$ | 2405.13103 | [25] |
| 9. | Comprehensive analsysis of local and nonlocal amplitudes in the $B^0 \to K^{*0}\mu^+\mu^-$ decay | 2405.17347 | [26] |
| 10. | Amplitude analsysis of the $B_s^0 \to K^+K^-\gamma$ decay | 2406.00235 | [27] |
| 11. | Observation of new charmonium or charmoniumlike states in $B^+ \to D^{*\pm} D^{\mp} K^+$ decays | 2406.03156 | [28] |
| 12. | Measurement of the branching fraction ratios $R(D^+)$ and $R(D^{*+})$ using muonic $\tau$ decays | 2406.03387 | [29] |
| 13. | Precision measurement of the $\Xi_b^-$ baryon lifetime | 2406.12111 | [30] |
| 14. | Probing the nature of the $\chi_{c1}(3872)$ state using radiative decays | 2406.17006 | [31] |
| 15. | Search for the rare decay of charmed baryon $\Lambda_c^+$ into the $p\mu^+\mu^-$ final state | 2407.11474 | [32] |
| 16. | Amplitude analysis of $B^+ \to \psi(2S)K^+\pi^+\pi^-$ decays | 2407.12475 | [33] |
| 17. | Study of charmonium production via the decay to $p\bar{p}$ at $\sqrt{s} = 13$ TeV | 2407.14261 | [34] |
| 18. | Observation of exotic $J/\psi\phi$ resonances in diffractive processes in proton-proton collisions | 2407.14301 | [35] |
| 19. | Measurement of $D^0 - \bar{D}^0$ mixing and search for CP violation with $D^0 \to K^+\pi^-$ decays | 2407.18001 | [36] |
| 20. | Observation of muonic Dalitz decays of $\chi_b$ mesons and precise spectroscopy of hidden beauty | 2408.05134 | [37] |
| 21. | Study of the rare decay $J/\psi \to \mu^+\mu^-\mu^+\mu^-$ | 2408.16646 | [38] |
| 22. | Measurement of CP violation observables in $D^+ \to K^-K^+\pi^+$ decays | 2409.01414 | [39] |
| 23. | Measurement of $\Lambda_b^0$, $\Lambda_c^+$ and $\Lambda$ decay parameters using $\Lambda_b^0 \to \Lambda_c^+ h^-$ decays | 2409.02759 | [40] |
| 24. | Measurement of CP violation in $B^0 \to D^+D^-$ and $B_s^0 \to D_s^+D_s^-$ decays | 2409.03009 | [41] |
| 25. | Measurement of exclusive $J/\psi$ and $\psi(2S)$ production at $\sqrt{s} = 13$ TeV | 2409.03496 | [42] |
| 26. | First determination of the spin-parity of the $\Xi_c(3055)^{+,0}$ baryons | 2409.05440 | [43] |
| 27. | Analysis of $\Lambda_b \to pK^-\mu^+\mu^-$ decays | 2409.12629 | [44] |
| 28. | Search for $B_{(s)}^{*0} \to \mu^+\mu^-$ in $B_c^+ \to \pi^+\mu^+\mu^-$ decays | 2409.17209 | [45] |
| 29. | Measurement of the effective leptonic weak mixing angle | 2410.02502 | [46] |

A search for tetraquark states was performed by LHCb in an amplitude analysis of the $B^+ \to D^{*-} D_s^+ \pi^+$ decay [24]. No evidence for exotic contributions in the $D_S^+\pi^+$ and $D^{*-}D_s^+$ channels was found, and stringent limits are set on the $T^*_{c\bar{s}0}(2900)^{++}$ state, which had been observed in the decay $B^+ \to D^- D_s^+ \pi^+$. Another ambitious amplitude analysis of $B^+ \to \psi(2S)K^+\pi^+\pi^-$ decays revealed new production modes of charged charmonium-like states and confirmed the $T_{c\bar{c}1}(4430)^{\pm}$ resonance with high significance [33]. Hidden-charm exotic states decaying to $\psi(2S)K^+\pi^-$ are observed. Their interpretation is yet unclear, but they could be

tetraquark states with minimal quark content $cc\bar{s}\bar{d}$. This first complete four-body amplitude analysis of a $B$-meson decay with one vector particle in the final state lays the groundwork for investigations into other such decays, which may further help understanding the intricate spectrum of exotic hadrons.

Studies of baryonic $b$-hadron decays to final states of the type $\Lambda_c^+ h^- h'^-$ were also performed, allowing for the measurement of fragmentation fractions and production asymmetries of the $\Xi_b^-$ baryon in $pp$ collisions [24]. Another study led to the most precise measurement of the $\Xi_b^-$ baryon lifetime relative to that of the $\Lambda_b^0$ baryon, improving the precision over the world-average value by a factor two [30]. The results of this analysis are in good agreement with the most recent theoretical predictions.

Several other results based on the spectroscopy of known and exotic states were published in the same time period [19, 28, 31, 35, 37].

## 7.2 *CP* violation and the CKM matrix

The LHCb experiment maintains an important and varied programme of studies of the charge-parity (*CP*) symmetry.

Over the past six months, LHCb published the most precise time-dependent *CP* asymmetries in the decays $B^0 \to D^+D^-$ and $B_s^0 \to D_s^+D_s^-$ [41]. With 5700 signal $B^0 \to D^+D^-$ decays, the hypothesis of *CP* symmetry, where both the cosine and sine time-dependent terms vanish, is excluded at more than six standard deviations. The results are consistent with previous results from LHCb and Babar, but are in tension with the Belle measurement, which is centered outside of the physically allowed region. Similarily, the 13000 signal $B_s^0 \to D_s^+D_s^-$ decays are used to obtained the most precise measurement in this channel, and are consistent with *CP* symmetry.

Several *CP*-violation results were also obtained in the study of charmed-meson decays. A search for *CP* violation was performed using the Run-2 dataset in $D^0 \to K^+\pi^-$ decays produced promptly via the strong process $D^{*+} \to D^0\pi^+$, which allows the $D^0$ flavour at production to be tagged with the charge of the associated pion. The simultaneous study of Cabibbo-favoured and suppressed decays proceeding directly or through $D^0 - \overline{D}^0$ oscillations, allow the most precise measurements to date of both *CP* asymmetry and mixing parameters in this channel to be obtained. The results are consistent with *CP* symmetry [36]. The study of *CP* asymmetries was also performed in three-body decays. A search is performed for localised *CP* violation in the phase space of $D^+ \to K^-K^+\pi^+$ decays [39], and for time-dependent *CP* violation in $D^0 \to \pi^+\pi^-\pi^0$ decays [22], which are both consistent with *CP* symmetry.

An update of the simultaneous determination of the CKM angle $\gamma$ and parameters related to mixing and *CP* violation in the charm sector was published in July 2024 as a note for the ICHEP conference [48]. Improved charm-mixing and *CP*-violation parameters are presented, as well as the most precise determination of the CKM angle $\gamma$ from direct measurements, which is found to be $(64.6 \pm 2.8)^\circ$.

## 7.3 Rare decays

The term *Radiative decays* is often used to describe quark level transitions of the form $b \to s\gamma$. Interest in these decays stems from the $b \to s\gamma$ process being a flavour changing neutral current (FCNC), which implies that these decays can only proceed within the Standard Model via loop or penguin diagrams. Given that there is no restriction as to which particles may run in these loops, FCNC processes define an interesting benchmark to perform precision tests of the Standard Model. Measurements of branching fractions and angular distributions can probe new physics scenarios arising above the weak scale, encoded in the Wilson coefficients of the weak effective theory. The results of several recent analyses by the LHCb collaboration have made remarkable impact in the knowledge of several Wilson coefficients.

A search for the $B_s^0 \to \mu^+\mu^-\gamma$ process has been performed using Run-2 data [18]. The photon in the final state lifts the chiral suppression factor present in the fully leptonic counterpart and gives access to the study of the local form factors describing the $B_s^0 \to \gamma$ transition. The search is performed in four regions of the dimuon invariant mass, resulting in the first limits set with this fully reconstructed final state and providing information on the set of Wilson coefficients $\mathcal{C}_{7,9,10}^{(\prime)}$.

An amplitude analysis of the $B_s^0 \to K^+K^-\gamma$ decay mode was performed using the full Run-1 and Run-2 datasets [27]. This result establishes the first observation of the radiative $B_s^0$ decay to an orbitally excited meson, $B_s^0 \to f_2'(1525)\gamma$, and the second radiative transition observed in the $B_s^0$ sector.

A complementary way of testing FCNC processes uses transitions in which two oppositely charged leptons are produced in the final state, namely, $b \to s\ell^+\ell^-$. The results of a comprehensive analysis of the local and nonlocal contributions in $B^0 \to K^{*0}\mu^+\mu^-$ decays were recently published by the LHCb collaboration [26]. This measurement employs for the first time a model of both one-particle and two-particle nonlocal amplitudes, and utilises the complete dimuon mass spectrum without any veto regions around the narrow charmonium resonances. The results show that interference with nonlocal contributions, although larger than predicted, only has a minor impact on the Wilson coefficients determined from the fit to the data. As a remarkable novelty, the nonlocal contribution due to $\tau^+\tau^-$ loop contributions rescattering into muons was included in the model. Direct measurements of the $\mathcal{C}_{9,10}^{(\prime)}$ Wilson coefficients are obtained, as well as the first direct determination of $\mathcal{C}_{9\tau}$. While the nonlocal contributions play a clear role in the studied angular distributions, these results show that fits to data still prefer a value of $\mathcal{C}_9$ that is shifted from the Standard Model expectation.

Similar studies of the decay rate as a function of the dilepton mass squared in FCNC processes were performed in the baryon sector. Being subject to different theoretical treatments due to the distinct hadronic system, these studies provide complementary information to those reported above. An analysis of $\Lambda_b \to pK^-\mu^+\mu^-$ decays [44] with intermediate $\Lambda$ resonances decaying to the $pK^-$ final state was performed with the full data sample collected with the LHCb de-

tector between 2011 and 2018. The analysis results in a first measurement of the differential branching fraction of the $\Lambda_b \to pK^-\mu^+\mu^-$ process across its entire phase space, in bins of the dihadron and of the dilepton masses. A first measurement of a complete set of angular observables for the $\Lambda$ states with spin less than 5/2 is also reported. The pattern of measurements appears consistent with SM expectations. However, a detailed interpretation of the results requires a more complete understanding of the hadronic system and the different contributing states.

Decays of heavy-flavoured mesons into final states containing more than two leptons are an interesting but rather unexplored probe of the Standard Model. Their rates can be computed in a relatively precise way, but possible new physics that would couple to leptons may alter these predictions. As preparatory works towards these searches, decays of the form $J/\psi \to \mu^+\mu^-\mu^+\mu^-$ have been searched for and studied in a recent LHCb publication [38]. The dimuon invariant-mass distributions in the background-subtracted data agree, within available statistical precision, with the expectations from the QED calculation, confirming that the decay is dominated by the final-state radiation (FSR) process. This validation of the theoretical description of FSR can then be applied to the much rarer four-lepton decays of $B$ mesons.

The breadth of the rare decays programme is further expanded by searches of rare processes, such as the $\Lambda_c^+ \to p\mu^+\mu^-$ decay mode [32], the lepton flavour violating process $B_s^0 \to \phi\mu^\pm\tau^\mp$ [25] or the very rare $B_{(s)}^{*0} \to \mu^+\mu^-$ decay modes in $B_c^+ \to \pi^+\mu^+\mu^-$ decays [45].

## 7.4 Electroweak physics

Using $5.4\,\text{fb}^{-1}$ of data collected during Run 2, LHCb measured the effective leptonic weak mixing angle in the $pp \to Z/\gamma^* \to \mu^+\mu^-$ process [46]. The weak mixing angle $\theta_W$ is a fundamental parameter of the Standard Model, which describes the unified electroweak interaction, and as a consequence controls the couplings of the $Z$ boson. Including higher-order corrections, an effective mixing angle is defined, $\theta^\ell_{\text{eff}}$, which can be measured experimentally.

In the LHCb measurement, the determination of $\theta^\ell_{\text{eff}}$ is performed in two stages. First, the forward-backward asymmetry $A_{\text{FB}}$ is measured as the asymmetry between the number of forward ($\Delta\eta > 0$) and backward ($\Delta\eta < 0$) occurrences of the $Z/\gamma^* \to \mu^+\mu^-$ process, where $\Delta\eta = \eta(\mu^-) - \eta(\mu^+)$ is the difference between the pseudorapidities of the two muons produced in the $Z$ boson decay. In order to maximise the sensitivity to the mixing angle, $A_{\text{FB}}$ is determined in bins of $|\Delta\eta|$.

In a second stage of the analysis, the distribution of the measured forward-backward asymmetries is compared with theoretical templates to determine the effective leptonic weak mixing angle, $\sin^2\theta^\ell_{\text{eff}} = 0.23147 \pm 0.00044 \pm 0.00005 \pm 0.00023$, where the uncertainties are statistical, systematic, and related to the theoretical model. This result improves by more than a factor two on the precision of the previous LHCb measurement, and is consistent with other direct measurements as well as with predictions from the global electroweak fit. This analysis demon-

strates the important role played by the LHCb experiment in core electroweak physics measurements.

# 8 Tracking detectors

## 8.1 Vertex Locator (VELO)

The VELO Upgrade is a pixel detector consisting of 52 modules, each equipped with four hybrid planar pixel tiles, arranged in thin walled radio-frequency (RF) boxes which form secondary vacuum enclosures within the LHC primary vacuum. It is cooled with evaporative $CO_2$ and provides a data push triggerless readout, with the total rate reaching 1.2 Tb/s.

### 8.1.1 Re-commissioning of the VELO

Following the reinstallation of the VELO RF boxes during the YETS 2023/2024, the VELO system required recommissioning. During beam-off periods, the thresholds for the 41 million pixels were re-optimized to reduce noise and improve timing synchronization, ensuring a well-calibrated front-end before beam operations began.

In the initial low-intensity LHC fills, the RF box was imaged by reconstructing particle interactions with its material, confirming that the box shape was as expected after installation, and enabling the full closure of the VELO. This is illustrated on the left of Fig. 6. If the closing procedure itself was quickly commissioned, early closings revealed a slow drift in the C-side VELO modules within the RF box. To mitigate this, adjustments were made to the closing and alignment procedures to prevent any impact on data quality. The source of the drift is being further investigated and will be addressed during the 2024/25 YETS.

As the number of bunches in the LHC increased during the intensity ramp-up, a new noise pattern correlated with the size of the bunch trains was detected, leading to instability in data acquisition at the highest instantaneous luminosity. After further analysis, the issue was resolved using the time-over-threshold filtering capabilities of the VeloPix ASIC. By early May, the VELO was fully commissioned and running under nominal conditions.

### 8.1.2 Operation and performances in 2024

In June, one of the 52 VELO modules became unresponsive, likely due to a faulty connection in the electronic control path. Fortunately, the position of the module and the system's redundancy ensured that tracks in the acceptance are not significantly affected. Efforts to recover the connection will be made during YETS 2024/25. Despite this missing module, 97.4% of the VELO's 624 ASICs remain active, with hit efficiency exceeding 98% for the active ASICs.

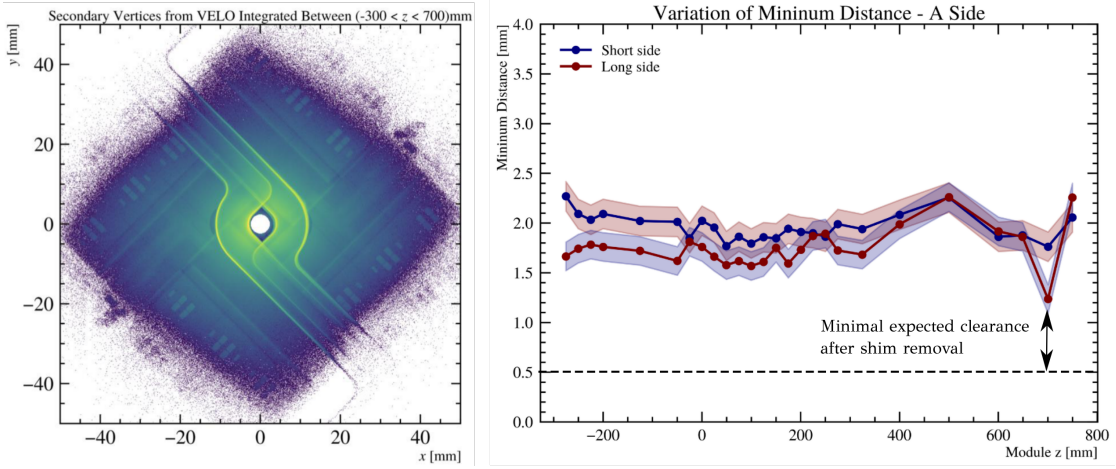


Figure 6: Position of hadronic interaction vertices in the first fills of the year (left). The RF box is well visible in the centre and analysing that dataset allowed to extract the minimal distance between the modules and the box in the various module z positions (right).

Over time, the sources of errors in VELO operations that could lead to data-taking inefficiencies have been identified. Monitoring systems were implemented to automatically address these issues, which helped reduce VELO data-taking inefficiencies to below 1%. Additionally, the time required to fully close the VELO was reduced to under 6 minutes.

In July, the motion system's belts were replaced after a non-conformity was detected in the movement of the VELO halves. Although the belts had been inspected during YETS 2023/24 and no defects were found, the presence of damage to a single tooth, shown on the right of Fig. 7, on both side belts suggests that the issue is likely related to the 2023 vacuum incident.

With approximately 10 $\text{fb}^{-1}$ delivered, radiation damage effects have begun to appear in the sensor performance. Weekly monitoring of leakage current is conducted to track the system's evolution as a function of particle fluence, see the left of Fig. 7. A procedure has been developed and applied to estimate optimal high-voltage settings to mitigate radiation-induced aging. No degradation in cluster efficiency was observed throughout the 2024 data-taking period.

Finally, the efficiency of track reconstruction in the VELO was assessed using a tag-and-probe method as a function of momentum for muons from $J/\psi$ decays, both in data and simulation. The VELO tracking efficiencies in data and simulation (see Fig. 15 later) show remarkable agreement, with remaining discrepancies mainly affecting tracks near the detector's edge, which cross the unresponsive module. These discrepancies are expected to decrease once the cluster efficiencies measured in data are fully integrated into the simulation.

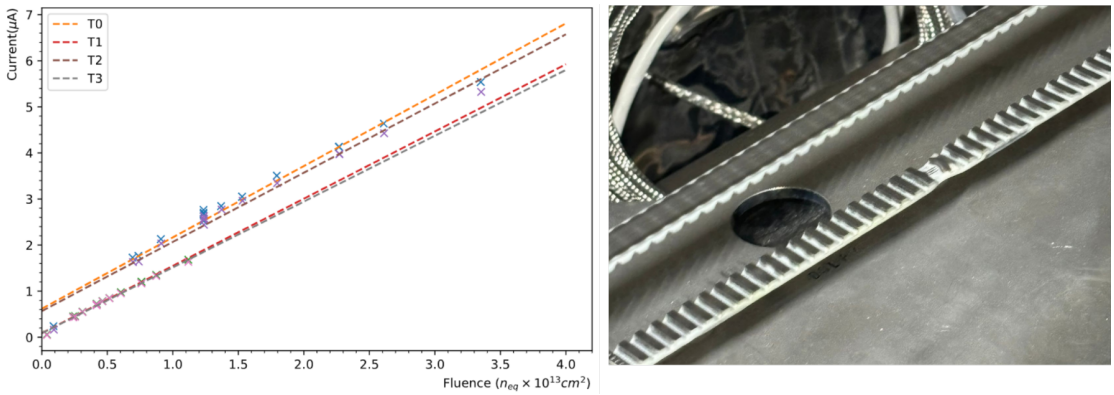


Figure 7: The measured current as a function of the total fluence seen by the four sensors (T0, T1, T2 and T3) of module 38 are evolving linearly as expected (left). Picture of the damaged tooth on the A-side belt (right).

### 8.1.3 Preparation for the 2024/25 YETS

As noted in the previous report, a leak was discovered in the secondary vacuum following the reinstallation of the RF box. It was not possible to repair the leak during the last YETS without causing significant delays to the LHC operation, but it was deemed safe to operate with it temporarily. The procedures and tools required to perform this repair with minimal impact on the detector have since been developed, and the issue will be resolved during the 2024/25 YETS.

In addition to fixing the leak, the C-side half fixation will be improved to further reduce module drift. It will also allow for the removal of the shims that currently keep the VELO modules retracted by 500 μm from their nominal positions. These shims were introduced as a precaution to maintain a safe distance until enough data confirmed that the clearance was sufficient to reposition the modules as intended, despite potential variations in box shape and module alignment. The analysis of material interaction data, shown on the right of Fig. 6, has confirmed that the clearance will indeed be adequate once the shims are removed.

### 8.1.4 SMOG2

As a result of installing the SMOG2 system, a fixed-target gas system able to provide collisions of both proton and heavy-ion beams with nuclei at $\sqrt{s_{NN}}$ ranging between 72 and 115 GeV, beam-gas collisions are collected simultaneously with the beam-beam collisions during the entire data taking period. LHCb collected approximately 15, 2, 0.6, and 1 $\text{pb}^{-1}$ of data using $H_2$, He, Ne and Ar gas targets, respectively. Overall, the system was very stable and reliable. With a pileup of about 4.3, the gas flux was kept at a level that increased the total data throughput by about 4% only. The trigger lines and reconstruction algorithms were optimised, enabling high-efficiency event reconstruction. Another important aspect is the momentum resolution and, consequently, the invariant mass reconstruction,

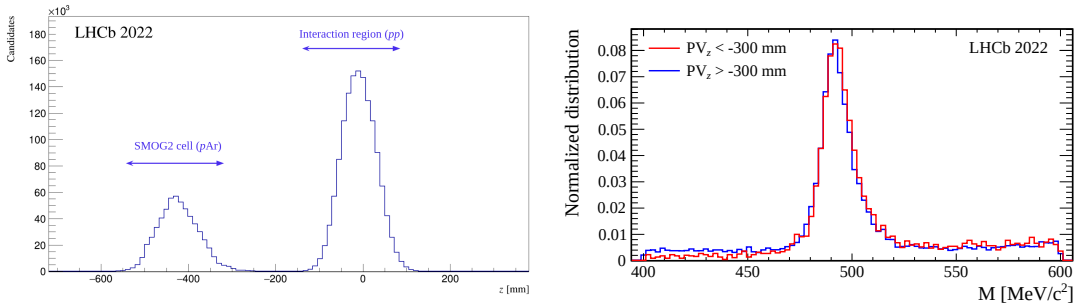


Figure 8: Left: Distributions of the z-coordinate for collision vertices reconstructed in HLT1 during a data-taking with overlapping $pp$ and $p$Ar collisions. Gaussian and a triangular-like shapes are clearly visible. Right: Invariant mass distributions for $K^0_S \to \pi^+\pi^-$ decays, reconstructed in HLT1 from $pp$+$p$Ar collisions. The two overlaid distributions for $p$Ar (red) and $pp$ (blue) have compatible widths.

where LHCb demonstrated the same optimal performance for both the beam-gas and beam-beam events. This proves that the two colliding systems operate simultaneously without interfering with each other. In Fig. 8, the left panel shows the distributions of the $z$ coordinate for collision vertices reconstructed at the first trigger level during a data-taking with simultaneous $pp$ and $p$Ar collisions. Two distinct distributions, one Gaussian and one triangular-like, are clearly visible. On the right of Fig. 8, the invariant mass distributions for $K^0_S \to \pi^+\pi^-$ decays, reconstructed HLT1 from $pp$+$p$Ar collisions, are shown. The two overlaid invariant mass distributions for $p$Ar (red) and $pp$ (blue) have compatible widths. The impact in LHC beam lifetime due to the presence of the gas has been shown to be negligible, producing a reduction of 1/3 of lifetime over 2000 days of continuous running, that must be compared to the typical 10-hour beam lifetime. Gas injections were also used during the van der Meer scans, exploiting the beam-gas imaging (BGI) technique for the luminosity determination and for evaluating ghost charges, not only for LHCb but also for the other LHC experiments. A detailed publication about the SMOG2 system and its initial performance results is in press [49] and has been selected as an Editors' Suggestion. During the next YETS, all the six 1-litre gas reservoirs will be refilled at the standard pressure of 1.5 bar.

## 8.2 Upstream Tracker (UT)

### 8.2.1 Introduction and overview

The UT is a silicon microstrip tracking detector located between RICH1 and the dipole magnet of LHCb. It comprises four planes of vertical columns called staves. Each stave is a carbon-fibre structure with embedded Ti pipes to provide evaporative $CO_2$ cooling that hosts silicon front-end (FE) hybrid modules attached to both sides to achieve full acceptance in the vertical direction. The dedicated (SALT) front-end ASICs, assembled on the 4 ASIC hybrids, feature 128 process-

ing channels, including a low-noise preamplifier and shaper, and a 6 bit analog-to-digital converter (ADC). The digitised information is processed in a complex DSP, performing common-mode suppression and zero suppression for the normal data-taking operation.

The two-way communication between data acquisition boards (SOL40 for experimental control and TELL40 for data processing) is coordinated by data control boards (DCBs) located in crates on the top and bottom of the UT boxes. Regulated power to the DCBs and front-end hybrids is provided by dedicated boards (LVRs) that are based on the radiation hard ASIC LHC4913, hosted in crates located in the service bay area (SBCs).

The four UT detector layers are called UTaX, UTaU, UTbV and UTbX, with X indicating the vertical strip orientation, while U and V indicate $5^o$ stereo angles. The staves are staggered in the $x$ direction to achieve complete coverage of the LHCb acceptance. The staves are mounted on two separate boxes. The two halves can be retracted to allow for beam pipe bake-out and maintenance.

### 8.2.2 2024 commissioning

The UT detector was installed in LHCb during the 2022/23 YETS and closed around the beam pipe in March 2023. The readout system was very unstable throughout 2023. Loss of synchronisation between various readout elements posed a serious obstacle to detector and firmware commissioning efforts. Several improvements in electronic board communications and firmware were implemented over 2023/24 YETS and allowed proper commissioning of the detector on $pp$ collisions in spring 2024. Initial commissioning was performed by taking data in "local" stand-alone mode in which UT was read out separately from the rest of the detector. Coarse time alignment was accomplished in March. Fine time alignment was carried out in April-May. SALT ASIC analog baselines were calibrated to peak near zero, minimising potential for corruption of pedestal values stored in the ASICs. This and other types of SEUs were observed and mitigated by implementation of a watchdog process running on the control computer, which resets ASIC registers if they don't match the template recipe. Zero-suppression ADC thresholds in the ASICs were set low enough to capture most of the ionisation peak from single charged track passing silicon sensor. Subsequently, data with UT in "global" mode were taken together with other LHCb subsystems, not yet including UT in HLT1 triggering, to commission use of UT in track reconstruction offline. The UT sensor positions were spatially aligned with the other tracking detectors through May-June, initially via a stand-alone procedure which fixed internal degrees of freedom, and later by simultaneous optimisation of relative positions of all three tracking devices. Online monitoring was substantially expanded to include UT performance in track reconstruction.

By the end of May, the LHC beam filling scheme reached its full capacity, at which UT DAQ system developed some instabilities. UT TELL40 boards were collapsing due to rare event size fluctuations, which would sometimes overwhelm

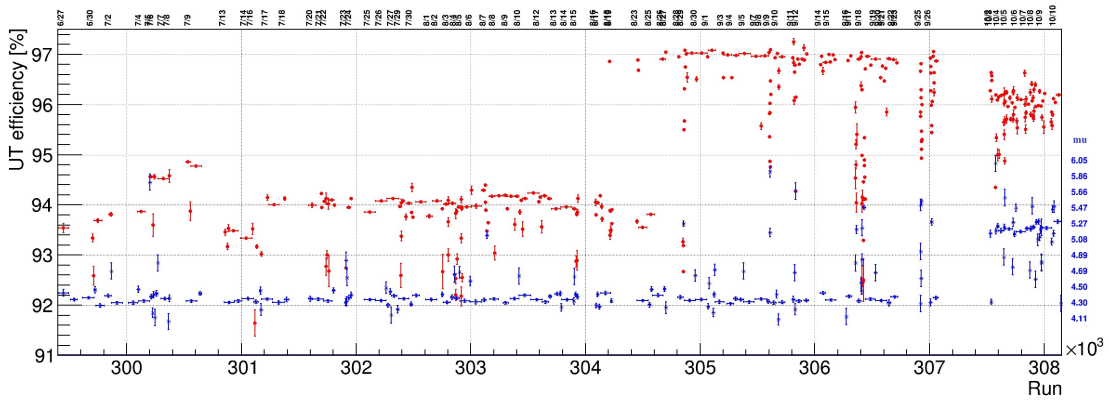


Figure 9: UT efficiency for matching hits in at least three UT layers to long tracks reconstructed between VELO and SciFi as a function of run number after UT joined regular data taking in June 2024 (red points). Average pileup values for data taken with UT are also shown (blue points) with their scale given on the right axis. Dates are indicated above the upper horizontal axis. The initial efficiency variations are due to changes in the track reconstruction unrelated to the UT detector. Significant improvement in UT efficiency at the end of August was due to the changes in FE ASIC configuration, as discussed in the text. Narrow run ranges with a large variation of pileup and UT efficiency reflect the commissioning tests in search for UT FE set-up suitable for high pileup values. Slight decrease in UT efficiency for steady 5.3 pileup running is due to the present DAQ limitations discussed in the text. The efficiency was extracted from the distribution of number of the UT layers with hits matching long tracks within the geometrical acceptance of the UT detector.

internal buffer lengths, limited by the amount of available electronic resources. The firmware was modified to self-guard buffer utilization. When incoming FE hits to be processed reach 80% of the buffer length, a TELL40 board disables FE operations until buffer utilization drops to its normal low level. This mechanism protects against TELL40 board crashes interrupting data taking, but is a source of potential detector inefficiency in the area served by particular TELL40 board. To prevent these inefficiencies, oversized events were truncated locally at an ASIC level by utilizing programmable hit limit, above which no hits are sent to a TELL40. As particle densities decrease substantially with radial distance from the beamline, the limits varied from 4 in the outermost detector to 63 in the innermost ASICs. This setting limited luminosity losses from UT TELL40 crashes to the percent level, which allowed UT to be included in regular physics data taking at the end of June.

## 8.3 2024 data taking

Since raw data are discarded for the majority of events after their reconstruction in HLT2, UT hits were included in HLT2 track reconstruction once UT joined the rest of LHCb physics data taking. The HLT2 algorithm requires three out

of four UT layers to produce hits matching long tracks reconstructed between VELO and SciFi subsystems. Once considered in the track fit, UT hits improve momentum resolution by providing precise track position at the entrance to the dipole magnet. The information about the number of matched UT hits is also used in a ghost suppression neural-net algorithm eliminating wrong VELO - SciFi track-segment combinations. UT triplets are essential for downstream track reconstruction in which UT and SciFi segments are matched to provide reconstruction of charged particles originating beyond the VELO volume. Downstream tracks greatly increase yields of reconstructed $K^0_S \to \pi^+\pi^-$ and $\Lambda^0 \to p\pi^-$ decays. UT efficiency on long-tracks remained relatively steady around 94% for the HLT2-match algorithm for data taking around a pileup of 4.4 through July and August, as illustrated in Fig. 9. With enough physics data accumulated, we could educate the ASIC hit limits better than initially set in June. They were tightened in the inner UT region to 20 and then ADC thresholds were lowered in most of the detector, except for the TELL40 boards serving the transition region from the finely segmented and highly instrumented inner detector to much coarser outer detector. These TELL40s are the most vulnerable to buffer overruns. As a result, UT efficiency on long tracks increased to about 97% (Fig. 9). At the same time, UT was included in the HLT1 two-layer match algorithm, which allowed a significant fraction of background from purely hadronic triggers to be eliminated.

With the pileup 4.4 FE setting, UT DAQ was too unstable for data taking at the design luminosity, corresponding to a pileup of 5.3. To gain ability to take data at such conditions we pushed the ASIC hit limits to a level at which the detector efficiency is reduced by about 1%. Since the particle flux drops quickly with a radial location, the limits were made dependent on ASIC distance from the beamline. With DAQ stabilised, we performed an optimisation of ADC thresholds in a series of short commissioning tests performed at high pileup values. The thresholds were raised in a limited number of areas served by TELL40 boards in the inner and transition regions which were particularly prone to inefficiencies induced by the buffer protection mechanism. At the same time, we were able to lower ADC thresholds in some other areas. As a net effect, UT efficiency on long-tracks is about 96.2% for steady data taking at 5.3 pileup during the last two weeks of the 2024 *pp* run (Fig. 9). In HLT1, use of UT via forward VELO-UT track reconstruction before matching to SciFi segments allowed for faster long-track reconstruction than use of UT in VELO-SciFi segment matching alone, which in turn provided for more efficient triggering at high pileup values. The most recent data taking also includes HLT1 downstream-track trigger which will be very beneficial to signal yields in channels with $K^0_S$ and $\Lambda^0$ particles in the final state.

## 8.4 Plans for the future

The remaining task for data taking this year involves finding a proper setting of FE ASICs for heavy-ion data taking. While heavy-ion collisions can produce

higher multiplicity events than proton-proton interactions, the frequency of such collisions will be small allowing for a quick TELL40 buffer recovery. Therefore, we expect to be able to loosen the ASIC hit limits, which will be essential for providing better UT efficiency in very-high multiplicity events.

Thanks to the quick commissioning on beams, the UT detector was able to catch up with the rest of the LHCb experiment and play its design role in tracking and triggering. Collecting data at the design luminosity pushes the UT DAQ to its limit imposed by the number of TELL40 readout boards available. Further increase in hit rates, for example by increasing pileup value beyond 5.3, increases UT inefficiency via data truncation by TELL40 boards activated by the firmware buffer-protection mechanisms. No further tightening of FE hit limits is possible without efficiency losses due to the data truncation by the hit limits in FE ASICs. Speeding up TELL40 firmware could provide for a more robust high-pileup operations in the future. Firmware simulations will be also used to identify potential improvements.

While the vast majority of SALT ASICs perform with electronic noise levels as expected from the beam tests and with a good stability of analog baseline, some outer ASICs in the central staves exhibit time-dependent baseline distortions for a part of their channels. As these reach a few ADC count level, this behaviour makes it difficult to set low ADC thresholds without flooding the ASICs and DAQ with noise hits. In order to understand this phenomenon, as well as other rare ASIC pathologies, more electronic calibrations and tests are being commissioned, like timing scans with injected charge. Since testing time is limited during data taking period, these activities will intensify over upcoming YETS period. A small number of hybrids had to be deactivated for data taking. Resurrecting them for active data taking is another task for the yearly LHC shutdown period.

A number of hardware interventions are planned for the YETS period: replacement of two DCBs, swapping a broken fibre, investigating high current in one LVR channel. investigating possible $CO_2$ leak, replacing dead nitrogen sensor. Presently the detector is operated at coolant temperature of $-10°$C, which is limited by humidity levels inside the UT box. To improve on them sealing of the box will be improved. This will be followed by tests at lower coolant temperature.

The remaining software development will be focused on improving simulation. We have delivered FE simulations which reproduce the UT efficiency for the two different FE set-ups used in the 4.4 pileup data taking within a percent. Accurate simulation of UT efficiency at 5.3 pileup will likely require simulating average TELL40 inefficiencies and accounting for all sources of hits in the detector because of the tight hit limits.

## 8.5 Scintillating-Fibre Tracker (SciFi)

### 8.5.1 System overview

The technology and design of the SciFi system are described in the LHCb Tracker Upgrade TDR [5]. The active element of the SciFi consists of 250 $\mu$m thick and 2.5 m long scintillating fibres arranged as hexagonally close-packed six-layer mats of 131 mm width. Eight of these mats are joined to form 5 m long and 52 cm wide modules. The fibres are read out from the outer ends by 128-channel arrays of silicon photomultipliers (SiPM), which can be cooled to -40°C to limit the single-pixel dark-count rate after irradiation. It is composed of a total of 524,288 read-out channels.

The fibre modules together with the readout electronics are mounted on 6 m $\times$ 4 m large support frames (C-frames). These frames provide the necessary services (low voltage and bias voltage, optical connections, SiPMs cooling) including water cooling of the electronics. The necessary cooling of the SiPMs is achieved with a separated cooling system operated at temperatures as low as $-50$°C.

The detector is divided in half with 6 C-frames arranged in 3 stations of X-U V-X layers installed on each side of the beam-pipe; there are 12 C-frames in total. A survey system (BCAM) is used to monitor the movements of the C-frames over time in addition to the alignment with tracking data.

We have reached a stable operation of the SciFi Tracker since 2023.

### 8.5.2 Detector operation

SciFi has been running smoothly within the LHCb data acquisition system during the year and data-taking was very efficient. Only 4 out of 512 front-end electronics boards and a few links went faulty during the operation in 2024 for various reasons. They are excluded at this time, resulting in 99.1 % of the detector functional. They will be fixed during the YETS.

Online control and monitoring tools have been enhanced to make diagnosis easier for non-experts. For example, a SciFi dashboard has been designed to monitor the health of the entire detector in real time, and overlay information on hitmaps for known problems is provided. This information can also be used to get the status of the data links and quickly identify a faulty link. Automated alarms have been set up to identify holes appearing on the hitmap or hot links, enabling automatic actions to be taken to remedy them. This procedure allows fast and efficient recovery.

We are using a Light Injection System (LIS) to calibrate all the SiPM channels of SciFi, but unfortunately for some SiPMs (we have 4096 in total), until the middle of 2024, running the result was not optimal and around 11% of the SiPMs showed a slightly reduced number of hit clusters, which had a 2-5% effect on hit efficiency performances on those SiPMs. While it is not fully understood why the LIS is not providing enough light in these region (and investigations are continuing), an alternative strategy has been developed to improve the calibration of these

channels and define new thresholds for them. This new method has been used successfully since August, and has resulted in uniformity in hit clusters across the entire detector.

A first part of the upgrade of the Condensation Prevention System (CPS) has been performed during the 2023/24 YETS. It consisted, to make the system more reliable, in using two Power System Units (PSUs) redundantly per channel with diode boxes to allow automated switching between the PSUs and to separate the PSUs (such that a failure on one should not affect the other one). There was a delay in the delivery of the new PSUs, but everything was ready for use, and they were installed when they arrived at the beginning of July. Since then, no problems have been experienced due to the CPS system.

It was seen previously that some of the thermal transfer oil in the heat-exchanger at the plant, used to cool the SiPMs, was jelling, coating the walls, and reducing the efficiency of the heat-exchanger, such that this was not sufficient to maintain a stable setpoint of -50°C for a long period. To fix this issue, EN-CV, in addition to the yearly maintenance of the cooling plant, has made a change of the oil in the exchangers. Unfortunately, even if the situation has improved significantly, some drifts are still observed over longer periods. A remedy consists of heating the detector every 4-6 weeks, during a couple of hours, until EN-CV conducts another upgrade, likely during LS3.

Over the past few months (and years), we have encountered a handful of problems with our high-voltage power supply system, mainly concerning the modules (channel turning off, in undervoltage condition or large current fluctuations). These modules were sent to the manufacturer for repair, but some of the tests carried out were inconclusive, and some modules were not repaired correctly. Unfortunately, these modules cannot be tested at the CERN Electronics Pool, and their faulty behaviour was only observed when they were installed on the detector. We are investigating what could be the origin of these issues, and we plan to purchase additional spare modules and to design a dedicated test bench, reproducing the detector environment, to test these modules as soon as they return from manufacturer's repairs.

### 8.5.3 Detector performances

In 2023 we reached an average hit detection efficiency of all SiPMs close to 98% and, as expected, a better understanding of the detector allowed us to achieve the highest performances in 2024, as shown in Fig. 10. The average efficiency of a layer (looking at the non-edge regions of the SiPM) varies between 99.0% and 99.5% for the low threshold set, and between 97.8% and 99.0% for the highest thresholds.

With the low set of thresholds, a sizeable amount of spillover is present. It has several origins: low momentum particles arriving late, delayed crosstalk and after-pulsing of the SiPMs, decay time of the scintillator, shaping and integration of the signal by the electronics, etc. On average, it is 29% from the central bunch crossing to the next one, but still also 10% in the subsequent bunch. A fine

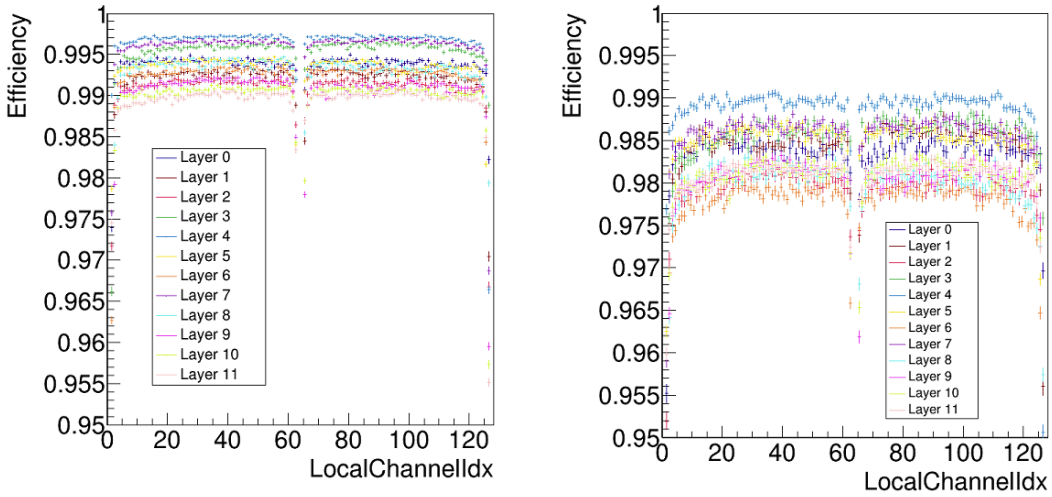


Figure 10: The average hit efficiency per channel of an SiPM for the SiPMs of 12 layer of SciFi. The left plot is for the lowest thresholds in Fill 9659 (1.5, 2.5, 3.5). The right plot is for the higher thresholds (2.5, 3.5, 4.5) in Fill 9662.

tuning of these thresholds leads to 30% (relative to default) lower spillover with comparable efficiency. They had been set slightly lower than necessary, giving more sensitivity to 1 and 2 p.e. signals. In addition, it gives 10% throughput improvement in HLT1 . This was implemented in early September.

Additional information on the performances of the SciFi detector in the tracking system is given in Sec. 11.

### 8.5.4 Activities plan during YETS

There are several activities planned for the 2024/25 YETS. After having upgraded the PSUs system of the upgrade of the CPS, we will perform the second step of its upgrade consisting in replacing the CPS electronic boards installed on the C-frames. This will ensure a more reliable system with finer granularity, such that we can only disconnect a small part of the system/sensor in the unlikely event of a problem. This new system is undergoing intensive testing using the prototype C-frame in the assembly hall at Point 8, before being installed on the detector.

Following a positive evaluation, a different cooling liquid for the SiPM, called NOVEC 7100, with a lower greenhouse warming potential (GWP=300) will be used to replace the current $C_6F_{14}$ (GWP = 10,000).

Radiation damages on SiPMs can be partially recovered by heating up the detector at room temperature (up to 35°C). This will be done during a couple of weeks.

We will investigate the few Front-End boards (FEB) links that are not fully functional, and replace some of them if necessary. Minor repairs to malfunctioning temperature sensors, flow metres, etc. and usual maintenance like the cleaning of

water filters, etc., will also be carried out. The origin of the problems with FEBs could be the use of non-baked VTRx for a large fraction of the detector, which suffer from out gassing effects when reaching higher temperatures. This creates a residue on the surface of the data and control fibres interrupting communication. The prevalence of this out-gassing and its affect on our optical fibre communication to the back-end will need to be investigated more thoroughly on a portion of the detector during next YETS, and the next. In the best case scenario, we should intervene only on a few fibres as they become problematic. In the worst case scenario, it may be necessary to move the C-frames from the cavern to the assembly hall in LS3 to perform a more significant intervention, which will require a year or more of activity and then recommissioning.

# 9 Particle identification detectors

## 9.1 Cherenkov detectors (RICH)

### 9.1.1 System overview

The RICH system of LHCb is composed by two detectors, RICH1 and RICH2, designed to provide charged hadron identification over the momentum range $2.6 - 100\,\text{GeV}$. The upgrade of the RICH system consists of a full new optics and mechanics design for the RICH1 detector, to cope with the increased luminosity, and new opto-electronics chain in both RICH1 and RICH2. The new opto-electronics chain is composed of multi-anode photomultiplier tubes (MaPMTs), as photon detectors, read out by a custom ASIC, the CLARO chip, designed to digitise the signal at 40 MHz. Finally, a FPGA-based digital electronics takes care of the interface with the PCIe40 for both the trigger and fast control and the data transmission. The full installation of the upgraded RICH system was achieved during LS2 so that both detectors were ready to be operated at the start of Run 3. By the end of the fist year of data taking, the performance of the RICH was already very close to target with a clear improvement with respect to Run 2, while operating with a five-fold increase in the instantaneous luminosity.

### 9.1.2 Operations

The RICH system was fully commissioned in 2022 and its performance was tuned in 2023. This year of data taking, after the VELO RF-foil replacement, has been a year of refinement for physics production. From the operations point of view, the focus has been on consolidating the monitoring infrastructure and the automation facility for the data taking. Overall LHCb has reached an excellent efficiency in the data taking, as described in Sec. 6. The RICH contributed to the small inefficiency with percentages of $\sim$0.3% and 0.1% for RICH1 and RICH2, respectively.

The infrastructure to inspect possible ageing of the detector has been put in place, with dedicated campaign of data taking for calibration acquired during the

machine ramp up periods. No ageing of the photon detector has been observed so far, therefore no tuning of the operating parameters will be needed. The next slot for ageing campaign is foreseen for the start of the data taking in 2025.
Since 2023, the RICH detectors have been operated with a signal latching scheme tuned within a window of 6.25 ns to balance the signal efficiency and the background rejection. A campaign dedicated to the optimisation of such window has just been performed to understand possible optimisations for the physics production in 2025 and to provide precious inputs to the future upgrades programme of the RICH system.

### 9.1.3 Performance

This year provided a huge amount of data for performance studies. Two main figures of merit are used to evaluate and tune the performance of the RICH system: the Cherenkov angle resolution and the particle identification efficiency. The resolution of the RICH2 detector is at design level while RICH1, despite being lower if compared to the resolution of Run 2, has still some room for improvements. The constant improvements in tracking and alignment during 2024, has resulted in concurrent improvements in the RICH1 Cherenkov resolution.

The particle identification performance has been studied in depth by the RICH and the RTA teams using different methods and verifying their consistency. As a result, the RICH system is able to provide excellent hadron identification as shown in Fig. 11, where the kaon identification efficiency as a function of the pion mis-identification probability is shown for different bins of primary vertexes: the separation between the two species degrades, as by design, with the increase in the event multiplicity and still results in a much better performance compared to Run 2, in the Upgrade I data taking conditions. Similar studies have been performed also for kaon-proton separation, as well as pion-proton separation, all resulting in excellent particle identification performance, with a large sample ready to be exploited for data analysis in LHCb.

## 9.2 Calorimeters (ECAL and HCAL)

### 9.2.1 System overview

The upgrade of the LHCb calorimeters consisted in replacing the electromagnetic (ECAL) and hadronic (HCAL) calorimeter readout electronics and upgrading the high-voltage, monitoring and calibration systems. The detector modules of the original LHCb were retained. The gain of the photomultipliers (PMT) has been reduced by a factor up to five in order to keep them operational throughout the higher luminosity runs of the upgrade. The new analog electronics partially compensates for the gain reduction by boosting the signals by a factor of 2.5. The remaining factor of two is used to extend the dynamic range of the calorimeter system and thus extend the physics case.

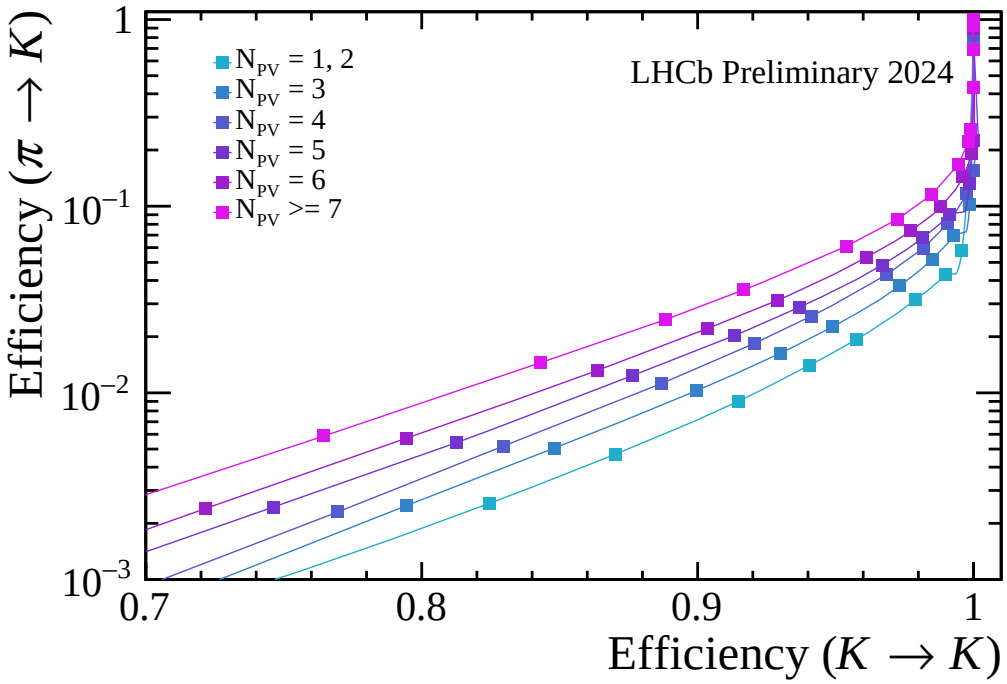


Figure 11: Kaon identification efficiency against pion mis-identifaction probability for Fill 9986, in different bins of number of primary vertices. The average number of primary vertices is 4.7. The average momenta of kaons and pions are both 35 GeV/$c$.

The construction of the electronics, for the front-end, the high-voltage, the monitoring and the calibration of the calorimeters was completed in 2021. After production, the installation started in parallel with the tests of the hardware, such that the installation was achieved in the same year. The final firmware versions to ensure a proper processing of the electronics have been written and loaded also in 2021.

### 9.2.2 Latest activities

The ECAL and HCAL detectors are fully operational and are used routinely in the global LHCb data acquisition since spring 2022. The calorimeters took part in the first high-energy collisions in stable beams in July 2022. Since then we have been routinely and smoothly running the ECAL and HCAL with a very high efficiency.

**Calorimeter issues** During last months, we essentially faced two main problems affecting the calorimeter electronics. The first concerned the configuration of some HCAL crates that could end in an error state. This happened rarely and always occurred after the clock domain of the LHCb experiment changed from the local clock to the LHC one, *i.e.* far before stable beams were declared by the LHC.

A reset of the faulty crates solved the problem. Finally we have been able to reproduce the problem during a Technical Stop, and the original cause was found in the configuration of the control boards that feed the front-end electronics with the 40 MHz clock.

Another issue appeared during summer 2024 affecting randomly a block (4 calorimeter channels) of the front-end electronics. The occurrence of the problem was once a week in June and July, then slowly decreased. We did not observe any such event over last month. We don't have any explanation for the behaviour observed yet. We rapidly took actions so that the problem, if it happens again, can be solved quickly by a reset of the faulty part of the front-end electronics.

**Maintenance and calibration** The ECAL and HCAL energy calibrations are performed routinely and satisfactorily. The HCAL uses a cesium source which got stuck into its garage during the last winter campaign of calibration. After the cleaning of the radioactive capsule by the radioprotection service at CERN, we did several calibration runs without any issue.

The ECAL energy calibration is usually done by reconstructing neutral pions. However, this requires a large sample of data and could not be done at every fill, inducing a noticeable decrease of the response of the ECAL in between calibrations. We have then added a relative calibration, applied at each fill, based on the response of the calorimeter channels to the LED pulses. The combination of the two calibrations, absolute and relative, enables the response of the detector to be stabilised in spite of the high multiplicity reached during the crossings at high pileup.

The time alignment of the ECAL and HCAL channels is now in routine operation. The plan for the future is to perform it only twice a year and rely on predictions of the time adjustments of the channels related to the variations of the HV of the photomultipliers. This method is still under study but would allow the frequency of the time scans to be reduced.

A few channels are dead or noisy at the level of 2 per mil. We benefit from technical stops to regularly cure the corresponding problems. The origin is usually a bad connection on the input connector of the front-end electronics.

**Monitoring and reconstruction** The recent activities focus on the data quality and the monitoring of the detector. Hence, a new alarm procedure has been defined and successfully used during the summer to quickly identify the four-channel issue mentioned earlier.

The reconstruction of the calorimeter data has also been polished and new corrections to the reconstruction of the cluster energy have been calculated. Figure12 shows the performance improvements that were achieved.

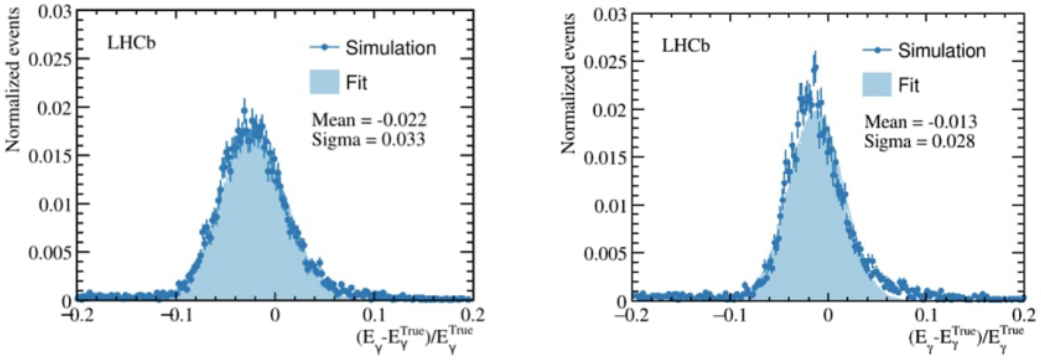


Figure 12: The ECAL energy reconstruction with the old corrections (left) and the new ones (right).

## 9.3 Muon chambers (MUON)

### 9.3.1 System overview

The MUON detector has performed exceptionally well in Run 1 and Run 2. The main changes for Run 3 are the removal of M1, the redesign of the off-detector electronics, and the installation of a new shielding in front of the inner region of M2 to reduce by about a factor two the low energy background rate in this region.

The electronics of the MUON detector upgrade consists of a new readout board (nODE), equipped with four custom ASICs (nSYNC) redesigned to be compliant with the 40 MHz readout of the detector, and of new control boards, the Service Board (nSB) and the Pulse Distribution Module (nPDM), redesigned to be compliant with the new ECS/TFC system.

### 9.3.2 Commissioning and performance

The MUON system is successfully taking part in the Run 3 LHCb data taking. The detector has been working well. The experimental control system (ECS) has been in daily use to operate and control the MUON system. Steady progress has been made in commissioning the detector control system, calibration and alignment using data collected in 2022 and 2023. A lot of effort has been put on the optimisation of the DAQ and in its interaction with the TELL40 firmware. A suite of new ECS panels has been developed to help this optimisation and serve as tools for the piquets and the experts in the management of the system. The new version of the online monitoring has been enriched with new histograms and with the activation of automatic analysis. The MUON TELL40 desynchronisation caused large inefficiencies in 2022. The root cause of the desynchronisation is still being investigated, but many tests allowed a working configuration for power supply, firmware, and software to be found, such that the issue is reduced to a negligible level, no longer affecting either the DAQ or muon identification performances. The MUON system has been aligned to the LHC clock, and the procedure for the

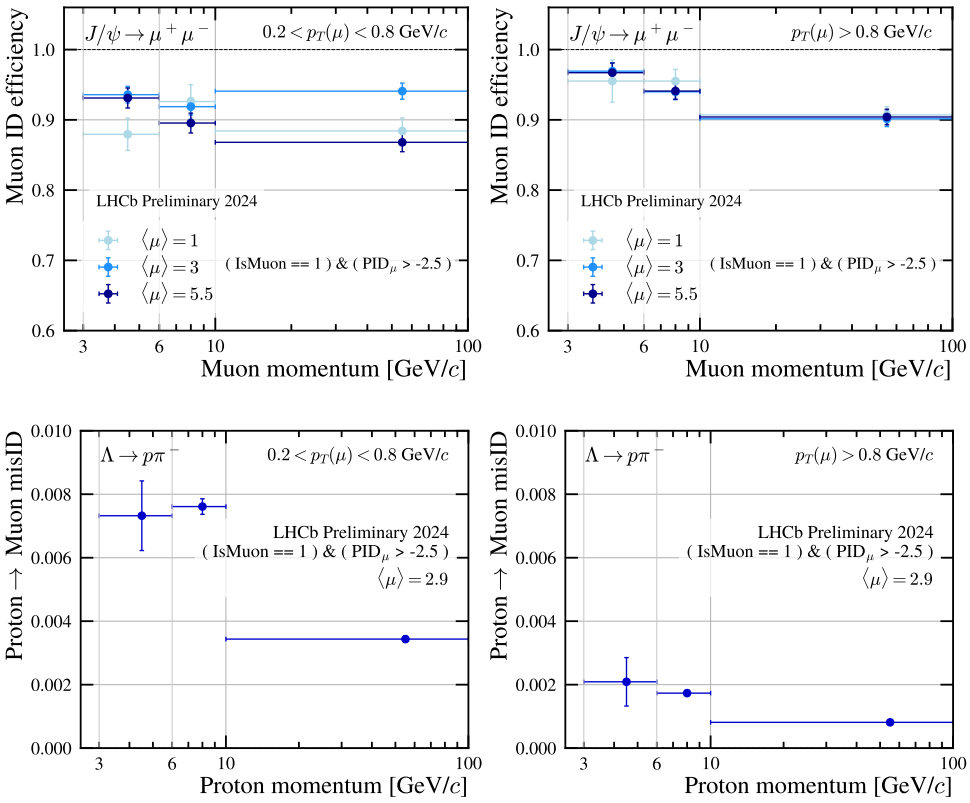


Figure 13: Muon identification efficiency computed on detached $J/\Psi \to \mu^+\mu^-$ (top) and proton to muon misidentification rate (bottom) in 2024 data (9512¡Fill-Number¡9573) as a function of average pileup.

fine time alignment of the electronics within the 25 ns bunch has been completed. To deploy the effective alignment, high quality Time Alignment Events data have been collected at the beginning of the 2024 data taking, allowing near-nominal performance.

A lot of effort has been also put to update the Run-1 and Run-2 MUON software to the Run-3 detector and environment. The readout cabling was changed and that required a new raw data decoding and simulation encoding software. The reconstruction and identification algorithms have been adapted to the tight time constraints of the full software trigger. The improvement of the tracking system and of the muon time-alignment allowed good identification performance to be reached. Figure 13 shows a muon identification working point chosen to have about 90% muon efficiency and low proton misidentification rate. The efficiencies have been evaluated with fit and count on detached $J/\Psi \to \mu^+\mu^-$ and $\Lambda^0 \to p\pi^-$ events in 2024 data.

The removal of M1 and the new shielding required to modify the detector description, and all the geometry has been ported to DD4HEP, while the work to port to DD4HEP all methods is close to be completed. Finally an improved

modelling of the detector response to low energy background and to spillover is ready to be used in the simulation.

# 10 PLUME and luminosity

For Upgrade I, LHCb has developed and equipped its interaction region with dedicated systems whose goals are to continuously monitor luminosity and beam induced background conditions. Moreover, other LHCb sub-detectors started to provide ECS counters to monitor the luminosity, increasing redundancy.

## 10.1 System overview

The luminosity and beam/background monitors at LHCb are composed of a set of hardware- and software-based systems. The PLUME detector is a new detector installed in the LHCb Upgrade, based on quartz sensors and using the calorimeter front-end readout electronics, specifically adapted for Run 3. PLUME is fully dedicated to continuously monitor in real-time the instantaneous luminosity, averaged and per bunch, and to provide a direct measure of the average number of visible $pp$ interactions per bunch crossing ($\mu$). PLUME is synchronised with the LHC clock and in addition it is included in the LHCb Data Acquisition (DAQ) system providing such measurements for offline purposes included in the event bank.

The Beam Conditions Monitor (BCM) directly measures beam losses around the LHCb interaction region and acts as the LHCb primary safety interlock, dumping the beams when such losses are above specified thresholds. The BCM consists of two stations (Upstream and Downstream), each composed of 8 diamond sensors interfaced to current-to-frequency converter (CFC) cards providing continuous running sums (RS) over 40 µs (short RS) and 1280 µs (long RS).

The Radiation Monitoring System (RMS) continuously estimates the amount of beam induced losses observed around the LHCb interaction region. The RMS is composed of 4 stations, each composed of 2 metal foil detectors connected to charge-to-frequency converters. The change in frequency is linearly proportional to the number of minimum ionizing particles (MIPs) providing an estimate of the amount of particles lost by the beams around LHCb.

PLUME, BCM Upstream and RMS are located around 1.5 m from the interaction point in LHCb, in the zone between the VELO and the cavern wall, commonly referred to as the VELO alcove. The BCM Downstream system is located between the Upstream Tracker and the LHCb dipole.

## 10.2 Commissioning, maintenance and operation

The PLUME, BCM and RMS detectors all completed their initial commissioning phase during the first months of data taking in Run 3. Since then, they have been operating continuously. Improvements in stability and robustness were added

throughout the years, and studies on detector efficiency were also performed. The online measurement of instantaneous luminosity showed a non-linear dependence at higher values of pileup. This was demonstrated by dedicated tests where LHCb took data at different values of pileup, from 0.1 to roughly 7. This issue was not problematic during 2022 and 2023, as LHCb took data at a maximum pileup value of roughly 2.5. A solution has been studied during the 2022/23 YETS and has been successfully deployed in time for the first PbPb collisions in 2023. The validation of the new approach for high pileup was performed with dedicated tests in April 2024, demonstrating good linearity up to pileup values higher than those used in design operation.

In addition, PLUME has worked on enhancing the recipes for the front-end boards, improving pedestal subtraction by determining values from empty events directly in the TELL40 readout board through dedicated firmware. This process has been monitored throughout 2024, with no issues observed. Further firmware development has enabled an online measurement of the bunch clock phase shift with a precision of about 100 ps. This is currently under commissioning and is expected to be fully deployed for next year's collisions. PLUME's online luminosity determination has been made nearly independent of the LHCb detector status, allowing it to continue operating even if the LHCb run is stopped. This feature was successfully deployed during the first collisions in March 2024, and no problems were observed during subsequent data taking. The measurement of luminosity per bunch has been implemented in the TELL40 readout board and is currently available.

The BCM system had its readout board updated, exchanging the Tell1 board with a custom one named MIBAD. The operation was successful and after a first validation period in which both boards were used in parallel, the BCM system is now running with the new readout without issues.

During last months, the VELO, SciFi and RICH sub-detectors deployed different ECS counters that have been added to the luminosity panels in the control room. These counters could be used in the future to obtain independent determinations of the online luminosity after a calibration performed with van der Meer scans.

## 10.3 Offline luminosity determination

Two dedicated luminosity lines have been added to the pool of trigger lines contained in the HLT1 framework. The first one, running at a rate of 30 kHz, contains around 60 counters from different sub-systems that will be used for the determination of offline luminosity. The second one, with a rate of 1 kHz, contains additional counters not included in the first due to event size constraints. These lines are propagated to the second trigger level, where additional counters such as those from RICH detectors are added. The output of the 1 kHz line, propagated to HLT2 and processed by Sprucing, is also exploited to count the number of luminosity events contained in physics streams. This information is then propagated

to the analysts' ntuples via the Final State Report (FSR) mechanism. This will allow the integrated luminosity to be determined from the ntuples, by counting the number of events stored in the FSRs and matching it with the offline luminosity determined per run contained in an online server deployed on GitLab and running on a CERN virtual machine.

This mechanism has been first deployed in September 2023 and will be finalised when an estimate of the offline luminosity collected by the experiment will be available. A first preliminary analysis of the data collected during the May 2024 vdM scan is being performed to achieve this goal. First results are expected for winter 2024. This will allow the offline luminosity collected during 2024, and subsequently for other years, to be calibrated with a precision of few percent.

Similar runs to van der Meer scans but shorter in duration, are used to verify that the cross-sections calibrated during van der Meer scans remain stable throughout the year. For the LHCb experiment, these scans are carried out at the end of stable beams, just before scheduled beam dumps. The first emittance scan took place in July 2024, and data are currently under analysis. Additional scans have been conducted during the year, and the schedule for 2025 will be determined based on the results of the data analysis.

# 11 Real time analysis

The Upgrade trigger [6] is fully implemented in software. Collisions are reconstructed in real time with the best possible quality; candidates and event-level information are then selected and written to offline storage. This process is done in two steps. In the first one, HLT1, a fast reconstruction sequence is run on about 500 GPGPUs to reduce the rate to about 1 MHz. Data are then stored in a disk buffer, enabling detector calibrations and derivation of alignment constants. Once ready, the second step, HLT2, performs the full reconstruction using CPUs on about 4000 nodes and either selects full events or individual particle candidates to be analysed further offline. In this scheme, the full reconstruction is performed online and not redone later.

## 11.1 First trigger stage

The first trigger stage (HLT1) has been running in production since the first collisions at all instantaneous luminosities in 2024. The running period can be split in 3 parts. During the first period, from the first collisions to the Machine Development in August, charged particles were reconstructed by a dedicated chain of algorithms. Its main development during that time was the mitigation of a non-linear behaviour at high detector occupancies, leading to a diminished reconstruction efficiency for events with many primary vertices, by improvements in the software. A neural network to reject fake tracks was deployed and updated, which allows the output of HLT1 to have a higher purity. Several so-called bandwidth

divisions were deployed - they serve to maximize the physics output by allocating a certain amount of HLT1 output bandwidth for each of the inclusive and exclusive trigger lines and balancing this allocation with respect to all other trigger lines. The different versions made progressive adaptations to the higher luminosities and to improvements for a more effective rejection of fake tracks.

During the second period, from after the August Machine Development until the end of September, a second, complementary, track reconstruction algorithm for charged particles was put in production, further increasing the track reconstruction efficiency for high-occupancy events, and including the UT detector into the HLT1 reconstruction chain, leading to a further reduction of fake tracks. In addition, HLT1 was running on 3 GPUs per Event Builder server (compared to 2 beforehand), increasing the effective processing power in HLT1 by 50%, thanks to the purchase of an additional sets of Nvidia A5000 GPUs in the first half of 2024.

During the third period, from the beginning of October on, a new track reconstruction algorithm to reconstruct trajectories of long-lived particles was deployed in HLT1, significantly increasing the trigger efficiency of decays such as $D^0 \to K^0_S K^0_S$ or of BSM physics for particles with long lifetimes. For this last running period at an instantaneous luminosity of $2 \cdot 10^{33}$ cm$^{-2}$s$^{-1}$, a new bandwidth division was deployed, running HLT1 at an output rate of 1.25 MHz, i.e. 25% above design value, to maximise the physics reach. This was made possible by the improvement in HLT2 processing speed - it is expected that in 2025 a similar or higher HLT1 output rate will be used in production, further increasing the physics reach.

## 11.2 Alignment and calibration

The work done since the start of the 2024 data taking aimed at understanding and improving the alignment and calibration of the detector using the data collected during 2024. The two halves of the VELO detector are typically aligned at the beginning of each fill. However, the C-side has been showing a slow residual movement after the end of the closing procedure. This is mitigated by frequent realignments of the C-side, triggered by a difference in the PV positions determined by the left- and right-half sensors. The VELO sensors and modules were also realigned, with a larger statistics sample, showing improved residuals.

Two versions of the alignment of the trackers were performed. The first one was released after the June Technical Stop and aimed at improving the internal degrees of freedom of the SciFi and UT detectors. At the same time, a new magnetic field map simulation, fitted to measurements, had been deployed to improve the difference between the mean value of the particle masses and their nominal PDG values. The second iteration took place at the beginning of August and mainly addressed the global alignment of all subdectors. It first used magnet-off data, to disentangle possible correlations between the magnetic field and the alignment, and then magnet-on data. This alignment sequence led to a significantly improved mass resolution, and mitigated most of the variation of the mass position. Fig-

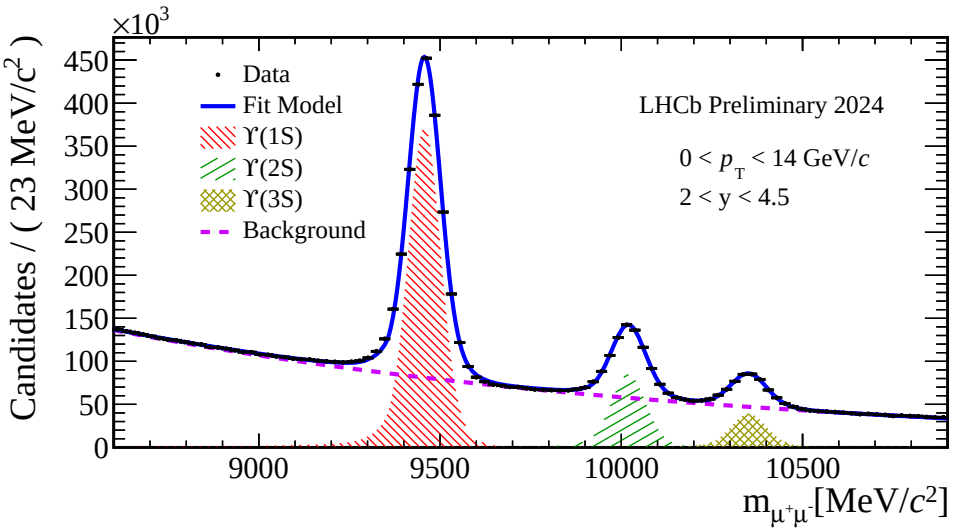


Figure 14: Invariant-mass distribution of a $\Upsilon \to \mu^+\mu^-$ data sample collected by the LHCb experiment in 2024, with a fit projection overlaid. Candidates of $\Upsilon$ states are required to have a transverse momentum within $(0, 14)$ GeV/$c$ and rapidity within $(2, 4.5)$ [50].

ure 14 shows the dimuon invariant mass spectrum in the $\Upsilon(1S), \Upsilon(2S), \Upsilon(3S)$ signal region. The $\Upsilon(1S)$ peak is well separated from the other two $\Upsilon$ states, and the overlap between $\Upsilon(2S), \Upsilon(3S)$ peaks is also small, which indicates a good momentum resolution.

Thanks to the huge amount of 2024 calibration data, the tracking and particle identification performance could be studied in great detail using the tag-and-probe technique. The track reconstruction efficiency is measured with $B \to J/\psi X$, $J/\psi \to \mu\mu$, $J/\psi \to ee$, and $K_S \to \pi^+\pi^-$ decays. With the inclusion of the UT in the reconstruction, the track reconstruction efficiency can be determined with higher precision, thanks to the reduced background and the improved purity of the signal. The global track reconstruction efficiency achieved higher values over time, following the improvements in the alignment of the detectors. The comparison between data and simulation results shows an overall good agreement as reported in Fig. 15, while some discrepancies as a function of kinematic variables are still to be understood.

The high quality performance of the RICH detectors, already observed with 2022 and 2023 data, has been confirmed with the 2024 data set. Figure 16 shows the kaon identification efficiency and pion misidentification probability as a function of momentum.

The statistics of the $B^+ \to J/\psi K^+$ sample, where the $J/\psi$ decays into two electrons, allowed the electron identification performance to be measured for the first time in Run 3. Figure 17 shows the electron identification efficiencies against momentum. The performances are comparable to those of Run 2.

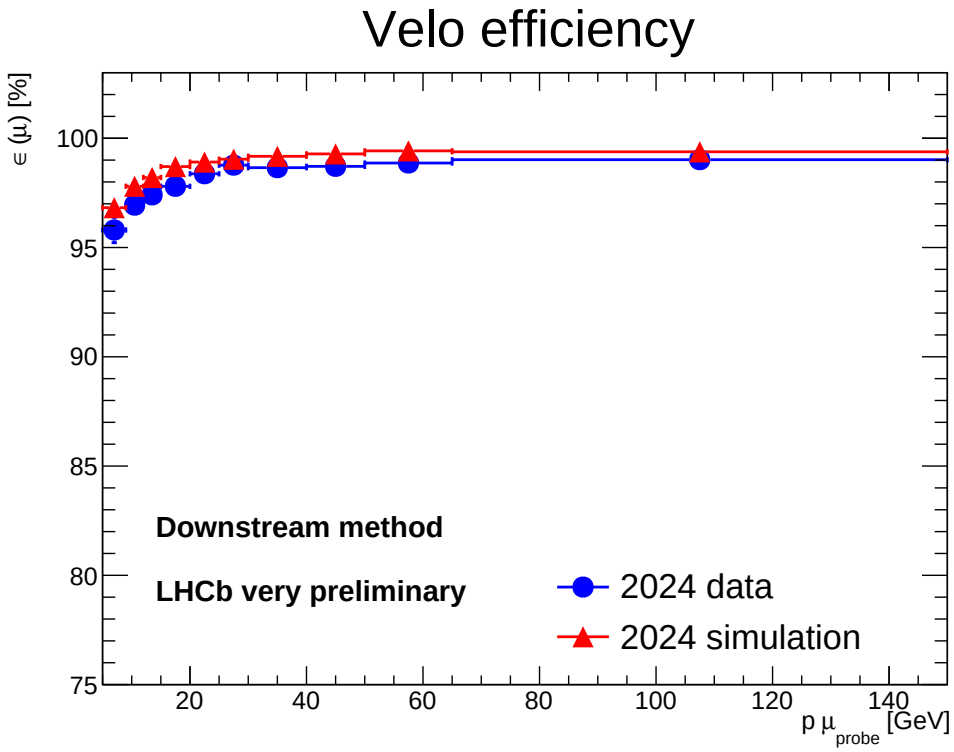


Figure 15: VELO efficiency in bins of momentum in comparison of 2024 data and simulation.

New particle identification variables, combining information from the RICH, Calo and Muon detectors into a neural network, were added to the reconstruction to achieve even better discrimination for the charged particle species. The further improvement of this classifier over the standard particle identification likelihood discriminator is visible in Fig. 18 showing the muon identification efficiency against proton misidentifaction probability.

The calibration procedure, which allows the analysts to determine the efficiency as a function of a particle identification criterion, has been developed and mostly automatised.

## 11.3 Second trigger stage

The second trigger stage (HLT2) has seen progressive improvements over the past running period on the reconstruction and on the selection sides. One key figure for HLT2 is the throughput, which needs to be high enough to empty the disk buffer after HLT1 quickly enough, taking the uptime of the LHC into account. Thanks to the replacement of the CPUs in the event filter farm with newer models, and the addition of 200 high-performance servers, the throughput could be roughly doubled over the course of the year, reaching about 800kHz at the design luminosity of $2 \cdot 10^{33}$ cm$^{-2}$s$^{-1}$, thus allowing the HLT1 output rate to be increased beyond

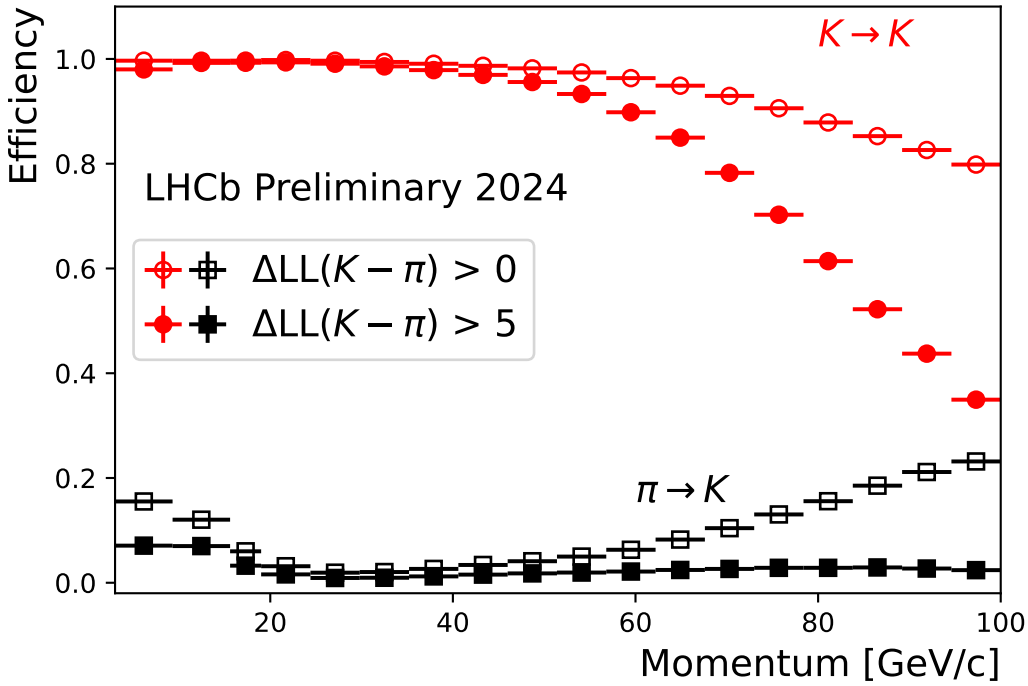


Figure 16: Kaon identification efficiency and pion misidentification probability as a function of momentum for Fill 9986. The average momenta of the kaons and pions are 40 and 39 GeV/$c$, respectively. The average number of primary vertices is 4.86 [51].

its design value of 1 MHz. In addition, several software improvements led to an additional increase in processing speed of about 10 - 20%.

Since the June Technical Stop, HLT2 has been running with the nominal configuration, including all design features. One major step was the inclusion of the UT in the reconstruction chain, which improves the momentum resolution of charged particles, reduces the amount of fake tracks thanks to a dedicated neutral net, and allows the reconstruction of long-lived particles, which decay outside the volume of the VELO detector. In addition, all particle identification variables are available already online, allowing for a precise selection of the decays of interest. In total, about 3000 inclusive and exclusive lines are running in HLT2 as of October 2024.

To monitor the output rate of HLT2, which should not surpass 10 GB/s, and its connected rate to disk (as candidates in the TURBO stream do not undergo any further selection), which should not surpass 3.5 GB/s, regular detailed tests of the data rate for each HLT2 line were performed. While the rates were below these limits before the August Machine Development, they significantly increased afterwards, thanks to the access to higher occupancy events in HLT1, the increased HLT1 output rate, and later also due to the higher instantaneous luminosity. A rapid development cycle between the physics analysts, the line developers and the core RTA team ensured that the output bandwidth could be brought under control

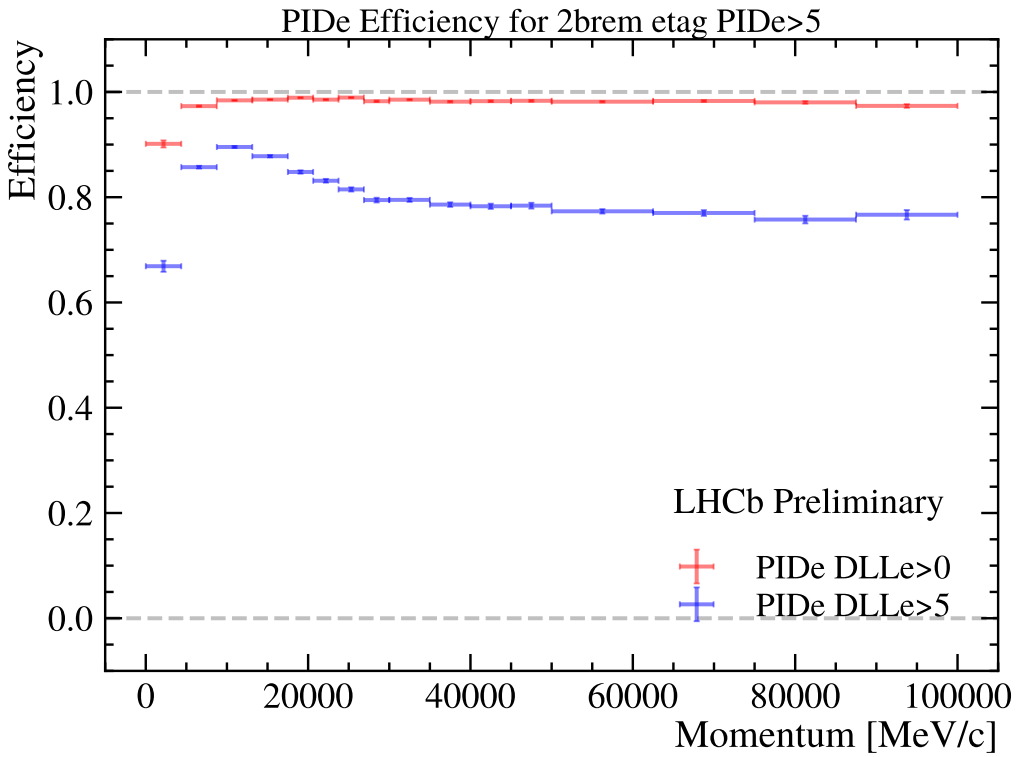


Figure 17: Electron identification efficiencies as a function of momentum, where the bremsstrahlung photons are added to the electrons.

in a short timescale, without compromising the LHCb physics programme.

### 11.4 Trigger efficiency

A major motivation for a pure software trigger was the removal of the hardware trigger (L0) and consequently a significantly higher trigger efficiency for decays with hadronic and electron final states. In order to quantify this improvement, comparisons between the trigger efficiency for L0 and HLT1 in Run 2, and for HLT1 in Run 3 were performed. The main results are shown in Fig. 19, where the left plot shows the comparison for the decay $B^+ \to D^0\pi^+$ and the right plot the comparison for the decay $B^+ \to J/\psi K^+$, with $J/\psi \to e^+e^-$, both showing gains of more than a factor 2.

## 12 Online

The Online system consists of the Experiment Control System (ECS), the Readout and Data Acquisition (DAQ), and the Online Infrastructure. At the center of the DAQ is the event-builder, which assembles the events at a rate of 40 MHz,

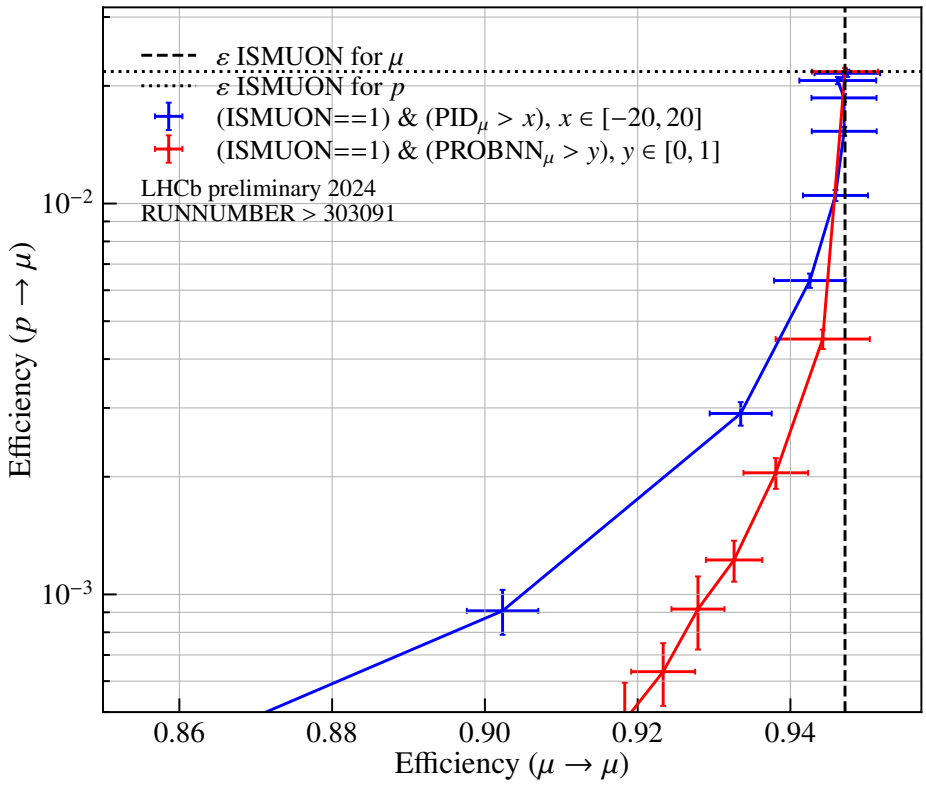


Figure 18: Muon identification efficiency as a function of proton misidentification probability.

aggregating 40 Tb/s. It is composed of 163 PC-servers interconnected through a 200 Gb/s bi-directional network, similar to the networks used to connect large machine learning training clusters. The event-builder is fully connected to the sub-detectors. It is the biggest and highest throughput system of its kind in the world today. There is continuous improvement of the event-builder software, the latest being a re-tuning of the data-flow to accommodate the use of a third GPGPU in each event-builder node.

New readout firmware has been deployed in all PCIe40 boards. It successfully addresses all known issues and has been extensively tested in long, stable runs. Many more monitoring and diagnostic counters and features have been added to facilitate detailed understanding and tuning of front-end electronics and readout.

The ECS[1] is focusing now on improving automation and speeding up certain operations (such as HLT2 processing).

[1]The ECS comprises all control and monitoring in LHCb, including the run-control and slow-controls.

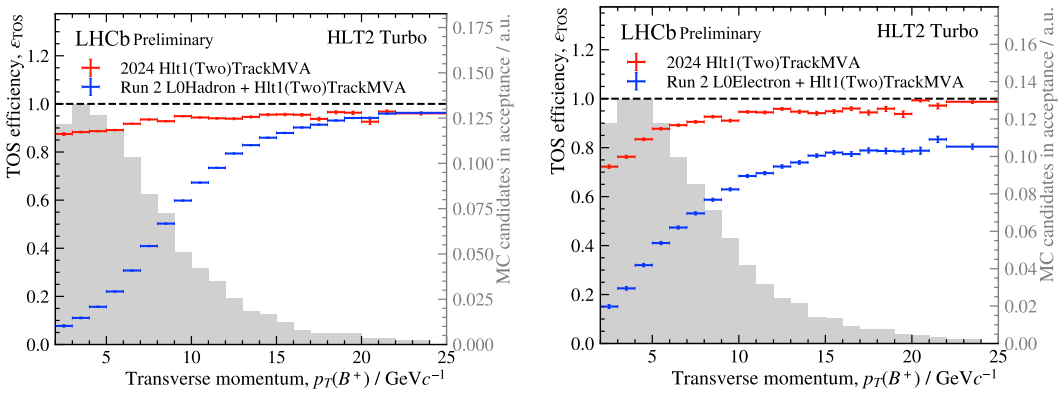


Figure 19: Efficiency of the first trigger stage(s) in (blue) Run 2 and (red) Run 3 (left) for the decay $B^+ \to D^0\pi^+$ and (right) for the decay $B^+ \to J/\psi K^+$, with $J/\psi \to e^+e^-$ [52].

In the LHCb data-centre at Point 8 more than 200 high-performance servers have been installed to increase the capacity of the HLT farm by more than 50%. Also a large fraction of the existing nodes has been upgraded to a more modern CPU. This has avoided obsolescence of the concerned servers and increased their performance by more than 50%. These CPUs were bought on the second-hand market, they would have been prohibitively expensive at the time the servers were originally acquired. The system administration team manually changed and reconfigured more than 3000 servers in only about a single month, without disturbing operations. Combined, the two measures have more than doubled our HLT capacity.

The data flow between HLT2 and the CERN Tier-0 has been made more robust by also adding monitoring for rare cases and is now mostly unattended. Further improvements of the resource usage are under way.

Apart from their important responsibilities in day-to-day operations of LHCb, the Online team is increasingly working on the preparation for the upgrades and improvements foreseen during LS3 for Run-4 and beyond.

No major changes are foreseen in the Online until LS3. All systems are working at or above design specifications.

# 13 Electronics

The Detector Electronics Commissioning Task Force (DECTF) successfully completed its work by the beginning of the summer. The Task Force is still in place, although not active, in case issues regarding detector electronics would come up and need urgent follow-up.

Since the last report, one other issue arose and it was related to the necessity of performing on-the-fly re-synchronisation of de-synchronised Front-End links. This

is due to exceptional instabilities that may be due to Single Event Upsets (SEUs) that can be solved with a simple procedure, issuing a reset and a re-synchronisation command centrally generated via the readout control system. Each sub-system implemented its own specific monitoring of desynchronised links, informing the Experiment Control System and finally triggering such a command. No stop of the run or loss of data taking is associated with such a command and this procedure allowed the LHCb detector to record data, since May, with a Data Acquisition Efficiency well above 90% steadily.

Finally, the Task Force, together with the Electronics Coordination and sub-detector experts will perform a study on SEU effects observed in 2024, running at the nominal instantaneous luminosity for Upgrade I. As of now, no destructive or degrading effects have been observed.

There has been significant progress on the new electronics for the LS3 enhancements. The FastRICH chip for RICH successfully passed a design review and the design is close to completion. It will be submitted for fabrication in Q1 2025. The design of the next-generation readout module, PCIe400, was finished and first prototype boards are expected in November. To prepare for future replacement of ageing power supplies, LHCb joined forces with CMS to launch a joint market survey to identify potential manufacturers.

# 14 Software and computing

## 14.1 Data processing

LHCb data are further processed offline after the HLT2 trigger stage to prepare them for disk storage where it is made available to analysts. This processing is referred to as *Sprucing* and involves a combination of (depending on the HLT2 stream) data slimming, further data selections and a file reformat before being streamed to disk.

The Sprucing of the HLT2 physics, including Turbo, TurCal, FULL and technical streams, NoBias and Lumi, has been running successfully and concurrently with 2024 data taking, by exploiting the WLCG resources in collaboration with the Computing team. So far in 2024 we have spruced more than 35 PB of data. It should be highlighted that due to framework improvements, Sprucing requires a factor five less CPU per event than its Run-2 counterpart known as the stripping.

For the Full and Turcal stream sprucing is essential to reduce the storage footprint to disk whilst maintaining high persistence to tape. In 2024 we have achieved the required factor 7.5 reduction in bandwidth between tape and disk at nominal luminosity. The to-disk bandwidth is monitored "live" in GitLab pipelines using real data as input.

Analysis productions are a centralised way for analysts to create ntuples from the data stored on disk that again exploits the WLCG resources. The adoption of analysis productions has been almost unanimous across Run 3 analyses. Over 1200

analysis productions on Run 3 data have been submitted so far with almost 1000 live productions picking up data as they are spruced concurrently with data taking; in a single day over 14 PB of spruced data was tupled with analysis productions. With the Sprucing running almost in real-time — the offline data transfer and Sprucing of an event is completed within 2-3 days of it being recorded by LHCb — analysts have minimal time-to-insight for the LHCb Upgrade dataset.

Development is ongoing to allow analysis production "sharing" such that multiple analyses can exploit the same production job; this reduces the read load on the Spruced data stored at WLCG sites.

Analysis productions run the DaVinci application which shares the same framework as the trigger and the Sprucing. It is the final application in the central data-processing chain. Comprehensive integration tests are in place to ensure the data processing chain is performing optimally.

The offline operations detailed above are the remit of the LHCb Data Processing and Analysis (DPA) project. Alongside this, the DPA project is responsible for the further exploitation of legacy Run 1 and 2 data, analysis preservation, open data and recently training within LHCb. The training work package is, in particular, charged with the creation of a Run 3 LHCb starterKit; the Run-1 and Run-2 version of which was instrumental in efficient on-boarding of new collaboration members.

## 14.2 Simulation

The simulation software continues to support productions for legacy datasets (Run 1 and Run 2), current datasets as they arrive (2023-2024) and future datasets (2025-2026 and Run 5). In the last six months, efforts have been focused on supporting both Run-3 data taking and ensuring that Run-5 (LHCb Upgrade II) simulation samples are available for the scoping document process and sub-detector TDRs in the coming years.

Special attention is being paid to ensure the simulated samples for 2024 are as accurate as possible and reproduce the varying periods of data taking. The sub-detectors are regularly implementing improvements with updates to the digitisation, allowing the detector response to mimic that of the real detector. To implement such updates, a number of productions rounds have been performed, starting from expected 2024 performance to samples aimed at specific data taking periods. Detailed samples including the high level triggers have been generated for one 2024 data taking period, with additional periods to follow shortly.

Significant attention has again been spent on the DD4hep detector description, focusing on getting it production ready for use with both Run-3 and Run-5 detector geometries. For Run 5, bi-weekly meetings follow the progress of sub-detectors who have now delivered default material descriptions for the Scoping Document and some of the descoping options to be studied further. The Run-3 subdetectors regularly provide feedback and improvements, many based on the data taken during 2024.

Progress in exploiting fast simulation approaches has continued both in using possible machine learning models for the calorimeter response arising from the Fast Calorimeter Simulation Challenge 2022 initiative, and using a point library approach. Studies are well advanced in each case: the internal code review process has started for the point library approach and is well advanced for the machine learning model. Commissioning for production to allow for test usage by selected physics analysts will be the focus of the next 6-12 months.

On the operational side, LbMCSubmit, the tool for submission of simulation production, continues to run smoothly. A new shift role has been introduced this summer, with the so-called Monte Carlo production chief responsible for collecting information about new types of productions that are required. Each production request is followed in a gitlab issue, which is closed once all of the packages required to submit the request have been released. This continues to streamline the submission of simulation requests, and the pool of shifters is expected to gradually increase moving forwards.

## 14.3 Offline computing

The Offline computing team has been focusing its activities on the support for 2024 data taking and data productions, to ensure that the latest, concurrent spruced data are available for analysis as soon as possible, and that the Monte Carlo productions proceed correctly. To achieve this, a smooth operation of the distributed computing infrastructure has been crucial.

The pledged compute resources (CPU, Disk and Tape storage) to support these activities were deployed at the LHCb Grid sites before the start of data taking. Two new Tier-1 centres have been added to the distributed computing infrastructure. NCBJ (Poland), confirmed as Tier-1 site in December 2023, and IHEP (China), confirmed in June 2024, following successful storage upgrades and tests, have been configured and are contributing to the data processing workflow. This brings the number of LHCb Tier-1 sites to a total of eight. In general, these sites have been working very well, although the unplanned downtime of the storage systems at one of the Tier-1 sites has been putting operational strain on the data processing workflows.

The extension of the 2024 data taking period, combined with the excellent performance of the LHCb detector and the high LHC efficiency, has led to a higher than anticipated need for storage capacity. To mitigate this, the collaboration reached out to WLCG and its funding agencies to request an advancement on the 2025 disk storage pledges. An additional 20 PB disk storage has been made available at the Tier-0 and most of the Tier-1 centres, which should allow LHCb the full processing of the data produced during the 2024 running period.

The production of Monte Carlo data sets consumes the bulk of the CPU resources on the Grid. Production has been running smoothly on the distributed computing resources pledged to LHCb. Most resources were used for simulation productions, 5% for physics analysis and 6% to process real data sets.

The tests of the simulation software on ARM hardware at Glasgow revealed an issue in the calorimeter simulation, which will be addressed in the coming months. An enhanced integration of ARM resources into the LHCb DIRAC data management workflow to enable the production use on Grid sites is also foreseen.

A dedicated document [53] has been prepared to present the preliminary requests for the computing resources in 2026. This request will be updated when more details about the 2026 LHC run will become available.

# 15 Infrastructure

LHCb's technical infrastructure (including the detector cooling systems) has been operating stably during most of the 2024 data-taking period. Thanks to the consolidation work performed during the YETS 2023/24, the temperature drift in the SciFi SiPM and VELO $CO_2$ cooling systems has been reduced significantly with respect to 2023.

During several days in summer with hot and humid weather, a temperature increase in the demineralised water circuit cooling the coils of the dipole magnet was observed, reaching values close to the trip threshold. As part of the effort to address this issue, a cleaning of the heat exchanger will be performed during the YETS.

## 15.1 YETS preparations

The schedule for the YETS 2024/25 is currently being finalised. In addition to regular maintenance tasks, the interventions foreseen on the detectors include the removal of the shims in the VELO, the repair of the leak in the VELO secondary vacuum, improvements to the sealing of the UT detector box, and the upgrade of the SciFi condensation prevention system. To further improve the Run-3 magnetic field map, work is underway to prepare additional magnetic field measurements, complementing the measurement campaign performed in 2021. In preparation for LS3, a refurbishment of the UGC1 gallery is planned, which will be used as a temporary storage area (*e.g.*, for calorimeter modules). Work has started, and will continue during the YETS, to connect the low-temperature detector cooling systems to an uninterruptible power source, with backup from the Point-8 diesel generator, in order to improve the resilience of the experiment against power outages.

# 16 Upgrade II status

A second major upgrade of the LHCb detector is necessary to maximise the physics potential of the HL-LHC. It will be installed after the completion of Run 4, by when the current detector will have accumulated its design integrated luminosity of $50\,\text{fb}^{-1}$. As shown in the LHCb Upgrade II Framework Technical Design Report

(FTDR) [16], by equipping subdetectors to provide fast timing information and using radiation-hard technologies, it will be possible to operate at instantaneous luminosity of up to $1.5 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ while maintaining or even exceeding the detector performance in the current experiment. This will enable a total sample exceeding $300\,\mathrm{fb}^{-1}$ to be integrated by the end of HL-LHC operation, and will give unprecedented sensitivity to physics beyond the Standard Model in beauty and charm physics. LHCb Upgrade II will additionally provide several other unique science opportunities, including in hadron spectroscopy, in electroweak precision measurements, in dark sector searches and in studies of heavy ion collisions. The European Strategy for Particle Physics [54] vision of full exploitation of the High Luminosity LHC (HL-LHC), including flavour physics, can only be achieved with LHCb Upgrade II.

A multi-step process to scrutinise and approve the Upgrade II project and secure the needed funding has been proposed by the LHCC at the June 2022 RRB meeting [55]. The next stage in the process is the preparation of a Scoping Document, which will complement the FTDR by adding information on detector scoping options matched to the established cost ranges of approximately 100%, 85% and 70% of the baseline cost outlined in the FTDR. Preliminary scenarios were discussed with the LHCC and presented to the RRB meeting in April 2024. The complete Scoping Document has been submitted to the LHCC for review on 2$^{\mathrm{nd}}$ September 2024, as per the agreed schedule. This document contains information on the financial envelope and physics performance in each of the descoped detector scenarios, on person-power and funding profiles, and on project organisation and milestones, and is briefly summarised in the following.

Since the approval of the framework TDR, significant progress has been made towards the design of an experiment that can be constructed and installed in the available time frame. Further studies in collaboration with the LHC accelerator division have confirmed that LHCb Upgrade II will not significantly impact on ATLAS and CMS operation during the HL-LHC era. Additionally, a flat optics configuration has been developed, which can provide higher integrated luminosity in the LHCb interaction region. Although not yet validated with dynamic aperature studies and dedicated machine tests, this provides a very promising scenario and illustrates robustness in the expected accelerator performance.

On the detector side, a re-evaluation of the baseline choices made in the framework TDR has been carried out in light of ongoing R&D activities. As a result, the FTDR design is largely confirmed. A schematic side-view of the Upgrade II baseline detector is shown in Fig. 20. A reduction of the fraction of the Mighty Tracker area using silicon pixels, relative to the FTDR design, is achieved thanks to improvement in the scintillating fibre technology that covers the remaining area. The area covered by the Upstream Pixel tracker is also reduced, removing outer regions where hits are not required to reconstruct good quality long tracks. The design of the PicoCal is updated to re-use, after refurbishment, modules from the present electromagnetic calorimeter in the outermost regions where the requirements allow this. The cost reduction from these optimisations is compensated

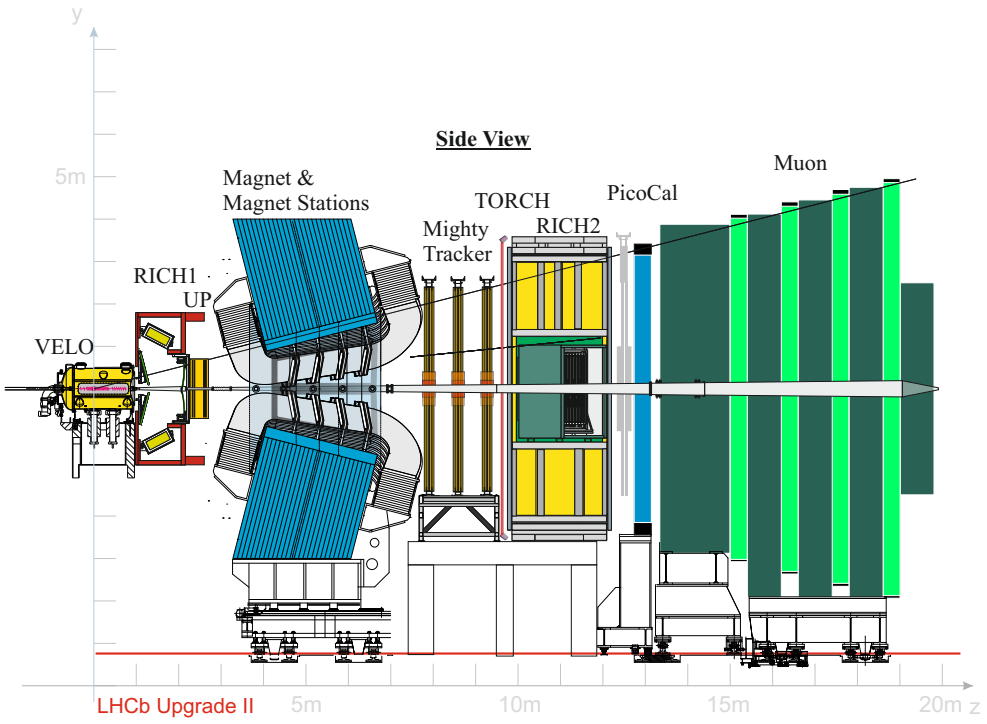


Figure 20: Schematic side-view of the Upgrade II baseline detector.

by an increase in other subdetectors due to a combination of inflation-related effects and design updates. The total Baseline scenario cost is $\sim 182$ MCHF, which exceeds only slightly the value estimated in the FTDR.

In addition to re-evaluation of the Baseline scenario, detector scoping options have been explored, with the aim of establishing scenarios at reduced cost and complexity of the overall project. Two descoped scenarios, Middle and Low, have been identified with cost envelopes estimated at $\sim 156$ MCHF and $\sim 125$ MCHF, respectively. In both cases the peak luminosity is reduced from 1.5 to $1.0 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Due to the longer levelling time at lower peak luminosity, the reduction in total LHCb integrated luminosity is proportionately less than that of the peak value, but is still significant at $\sim 12\%$ (or $\sim 20\%$ with flat optics). However, the corresponding decreases of peak detector occupancy and data throughput lead to important simplifications in the detector design, with associated cost savings. Additionally, the longer levelling time will provide operational stability, which is expected to be important for the high precision measurements at the core of the physics programme.

In the Middle scenario, the area covered by silicon pixels in the Mighty Tracker is further reduced from the Baseline configuration, and significantly less investment is required in the trigger farm. Both the MUON and RICH detectors have reduced

granularity, and the active area of the TORCH detector is reduced. Additional shielding before the Muon stations is not required, reducing the complexity of the infrastructure requirements.

In the Low scenario, the TORCH detector and the Magnet Stations are eliminated, reducing the overall complexity of the project but at significant loss in physics capability. The active areas of both the Upstream Pixel tracker and the Mighty Tracker scintillating fibres are reduced, and some VELO stations are removed, with a direct impact on acceptance. A more conservative VELO module substrate with increased material is used, with equally direct consequences for the impact parameter resolution, which is crucial for the efficient operation of LHCb's real time analysis system. Single rather than double readout is used in the outer region of the PicoCal, and the RICH2 optical system and its vessel are not upgraded. The cost of the trigger farm is further reduced, which may necessitate a narrowing of the physics programme if it is not possible to process all channels of interest with the available resources.

The impact of the different scenarios on physics performance has been studied. The Middle scenario has less redundancy than the Baseline, with the reduction in capability of operating at the highest pile-up reducing the flexibility to compensate for possible changes in the LHC operation such as the number of colliding bunches in LHCb or the number of physics days. Nonetheless, this scenario is still expected to provide robust tracking and PID performance, since the reduction of granularity is balanced by the reduction of peak luminosity. The lower integrated luminosity will reduce the achievable sensitivity across the physics programme. On the other hand, the longer levelling time operating at instantaneous luminosity of $1.0 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ will have some advantages for stable operation of the detector. In the Low scenario, the loss of redundancy in detector features puts at risk the success of some key measurements that require the full detector capabilities. A good programme of measurements is expected to still be possible, but is likely to be less broad in scope. It will also be challenging to achieve the highly precise measurements that are planned due to worse signal-background separation and increased detector material.

Plans for the project organisation and schedule have also been developed since the FTDR. The proposed organisation is based on experience from previous construction projects, both for LHCb and other experiments. Technical Design Reports for the subdetectors are anticipated in 2026, after which the project will move to the construction phase. Each subdetector has developed a plan to complete construction by mid-2032, so that installation can proceed during the two years of Long Shutdown 4 currently scheduled for 2033–34, and be ready for data-taking at the start of Run 5 in 2035. (The dates given here are those in the Scoping Document, not updated following the change in LHC schedule that was made subsequent to the LHCC submission at the start of September 2024. In the revised schedule, Long Shutdown 4 occurs during 2034–35, and the LHCb Upgrade II detector will be ready for data-taking at the start of Run 5 in 2036.) The required human resources and technical infrastructure have also been scru-

tinised. The main risks associated with these schedules have been identified and work is already ongoing to prepare mitigation plans.

To conclude, three detector scenarios, Baseline, Middle and Low, have been considered. These correspond to different cost ranges, and also have implications for the overall complexity of the project and the required human resources and technical infrastructure. Each can be considered plausible depending on assumptions on the support from Funding Agencies and the growth of the collaboration. An exciting and unique physics programme will be possible in each scenario, although the reduction of integrated luminosity and detector features in the downscoped scenarios has a clear impact on the anticipated physics performance. The higher cost of the Baseline and (to some extent) Middle scenarios not only results in improved performance and a stronger science programme, but also provides additional robustness to achieve the goal of full exploitation of the HL-LHC.

# References


[1] LHCb collaboration, *Letter of Intent for the LHCb Upgrade*, CERN-LHCC-2011-001. LHCC-I-018, 2011.

[2] LHCb collaboration, *Framework TDR for the LHCb Upgrade: Technical Design Report*, CERN-LHCC-2012-007, 2012.

[3] LHCb collaboration, *LHCb VELO Upgrade Technical Design Report*, CERN-LHCC-2013-021, 2013.

[4] LHCb collaboration, *LHCb PID Upgrade Technical Design Report*, CERN-LHCC-2013-022, 2013.

[5] LHCb collaboration, *LHCb Tracker Upgrade Technical Design Report*, CERN-LHCC-2014-001, 2014.

[6] LHCb collaboration, *LHCb Trigger and Online Upgrade Technical Design Report*, CERN-LHCC-2014-016, 2014.

[7] LHCb collaboration, *LHCb Upgrade Software and Computing*, CERN-LHCC-2018-007, 2018.

[8] LHCb collaboration, *Computing Model of the Upgrade LHCb experiment*, CERN-LHCC-2018-014, 2018.

[9] LHCb collaboration, *LHCb SMOG Upgrade*, CERN-LHCC-2019-005, 2019.

[10] LHCb collaboration, *LHCb Upgrade GPU High Level Trigger Technical Design Report*, CERN-LHCC-2020-006, 2020.

[11] LHCb collaboration, *LHCb PLUME: Probe for LUminosity MEasurementt*, CERN-LHCC-2021-002, 2021.

[12] LHCb collaboration, *Addendum No. 1 to the Memorandum of Understanding for Collaboration in the Construction of the LHCb Detector. The Upgrade of the LHCb Detector: Common Project items* CERN-RRB-2012-119A, revised April 2014, CERN, Geneva, 2012.

[13] LHCb collaboration, *Addendum No. 2 to the Memorandum of Understanding for Collaboration in the Construction of the LHCb Detector Upgrade of the LHCb detector Sub-Detector Systems*, CERN-RRB-2014-105, 2014.

[14] LHCb collaboration, *Expression of Interest for a Phase-II LHCb Upgrade: Opportunities in flavour physics, and beyond, in the HL-LHC era*, CERN-LHCC-2017-003, 2017.

[15] LHCb collaboration, *Physics case for an LHCb Upgrade II — Opportunities in flavour physics, and beyond, in the HL-LHC era*, arXiv:1808.08865.

[16] LHCb collaboration, *LHCb Framework TDR for the LHCb Upgrade II Opportunities in flavour physics, and beyond, in the HL-LHC era*, CERN-LHCC-2021-012, 2022.

[17] LHCb Collaboration, *Computing model of the Upgrade LHCb experiment*, CERN-LHCC-2018-014. LHCB-TDR-018, CERN, Geneva, 2018.

[18] LHCb collaboration, R. Aaij *et al.*, *Search for the $B_s^0 \to \mu^+\mu^-\gamma$ decay*, JHEP **07** (2024) 101, arXiv:2404.03375.

[19] LHCb collaboration, R. Aaij *et al.*, *Search for prompt production of pentaquarks in open charm hadron final states*, Phys. Rev. **D110** (2024) 032001, arXiv:2404.07131.

[20] LHCb collaboration, R. Aaij *et al.*, *First observation of $\Lambda_b^0 \to \Sigma_c^{(*)++} D^{(*)-} K^-$ decays*, Phys. Rev. **D110** (2024) L031104, arXiv:2404.19510.

[21] LHCb collaboration, R. Aaij *et al.*, *Amplitude analysis and branching fraction measurement of $B^+ \to D^{*-} D_s^+ \pi^+$ decays*, JHEP **08** (2024) 165, arXiv:2405.00098.

[22] LHCb collaboration, R. Aaij *et al.*, *Search for time-dependent $CP$ violation in $D^0 \to \pi^+\pi^-\pi^0$ decays*, Phys. Rev. Lett. **133** (2024) 101803, arXiv:2405.06556.

[23] LHCb collaboration, R. Aaij *et al.*, *Transverse polarisation measurement of $\Lambda$ hyperons in pNe collisions at $\sqrt{s_{NN}}$=68.4 GeV with the LHCb detector*, JHEP **09** (2024) 082, arXiv:2405.11324.

[24] LHCb collaboration, R. Aaij *et al.*, *Study of $b-$hadron decays to $\Lambda_c^+ h^- h'^-$ final states*, JHEP **08** (2024) 132, arXiv:2405.12688.

[25] LHCb collaboration, R. Aaij *et al.*, *Search for the lepton-flavor violating decay* $B_s^0 \to \phi\mu^\pm\tau^\mp$, arXiv:2209.09846, Submitted to Phys. Rev. D.

[26] LHCb collaboration, R. Aaij *et al.*, *Comprehensive analysis of local and nonlocal amplitudes in the* $B^0 \to K^{*0}\mu^+\mu^-$ *decay*, JHEP **09** (2024) 026, arXiv:2405.17347.

[27] LHCb collaboration, R. Aaij *et al.*, *Amplitude analysis of the radiative decay* $B_s^0 \to K^+K^-\gamma$, arXiv:2406.00235, Submitted to JHEP.

[28] LHCb collaboration, R. Aaij *et al.*, *Observation of new charmonium or charmoniumlike states in* $B^+ \to D^{*\pm}D^\mp K^+$ *decays*, Phys. Rev. Lett. **133** (2024) 131902, arXiv:2406.03156.

[29] LHCb collaboration, R. Aaij *et al.*, *Measurement of the branching fraction ratios* $R(D^+)$ *and* $R(D^{*+})$ *using muonic* $\tau$ *decays*, arXiv:2406.03387, Submitted to Phys. Review Letters.

[30] LHCb collaboration, R. Aaij *et al.*, *Precision measurement of the* $\Xi_b^-$ *baryon lifetime*, Phys. Rev. **D110** (2024) , arXiv:2406.12111.

[31] LHCb collaboration, R. Aaij *et al.*, *Probing the nature of the* $\chi_{c1}(3872)$ *state using radiative decays*, arXiv:2406.17006, submitted to JHEP.

[32] LHCb collaboration, R. Aaij *et al.*, *Search for the rare decay of charmed baryon* $\Lambda_c^+$ *into the* $p\mu^+\mu^-$ *final state*, Phys. Rev. **D110** (2024) 052007, arXiv:2407.11474.

[33] LHCb collaboration, R. Aaij *et al.*, *Amplitude analysis of* $B^+ \to \psi(2S)K^+\pi^+\pi^-$ *decays*, arXiv:2407.12475, Submitted to JHEP.

[34] LHCb collaboration, R. Aaij *et al.*, *Study of charmonium production via the decay to* $p\bar{p}$ *at* $\sqrt{s} = 13$ *TeV*, arXiv:2407.14261, Submitted to Eur. Phys. Journal C.

[35] LHCb collaboration, R. Aaij *et al.*, *Observation of exotic* $J/\psi\phi$ *resonances in diffractive processes in proton-proton collisions*, arXiv:2407.14301, Submitted to Phys. Rev. Lett.

[36] LHCb collaboration, R. Aaij *et al.*, *Measurement of* $D^0$–$\overline{D}^0$ *mixing and search for CP violation with* $D^0 \to K^+\pi^-$ *decays*, arXiv:2407.18001, Phys. Rev. D.

[37] LHCb collaboration, R. Aaij *et al.*, *Observation of muonic Dalitz decays of* $\chi_b$ *mesons and precise spectroscopy of hidden beauty*, arXiv:2408.05134, Submitted to JHEP.

[38] LHCb collaboration, R. Aaij *et al.*, *Study of the rare decay* $J/\psi \to \mu^+\mu^-\mu^+\mu^-$, arXiv:2408.16646, Submitted to JHEP.

[39] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP violation observables in $D^+ \to K^-K^+\pi^+$ decays*, arXiv:2409.01414, Submitted to Phys. Rev. Lett.

[40] LHCb collaboration, R. Aaij *et al.*, *Measurement of $\Lambda_b^0$, $\Lambda_c^+$ and $\Lambda$ decay parameters using $\Lambda_b^0 \to \Lambda_c^+ h^-$ decays*, arXiv:2409.02759, Submitted to Phys. Rev. Lett.

[41] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP violation in $B^0 \to D^+D^-$ and $B_s^0 \to D_s^+D_s^-$ decays*, arXiv:2409.03009, Submitted to JHEP.

[42] LHCb collaboration, R. Aaij *et al.*, *Measurement of exclusive $J/\psi$ and $\psi(2S)$ production at $\sqrt{s} = 13$ TeV*, arXiv:2409.03496, Submitted to sciPost. Phys. .

[43] LHCb collaboration, R. Aaij *et al.*, *First determination of the spin-parities of the $\Xi_c(3055)^{+,0}$ baryons*, arXiv:2409.05440, Submitted to Phys. Rev. Lett.

[44] LHCb collaboration, R. Aaij *et al.*, *Analysis of $\Lambda_b \to pK^-\mu^+\mu^-$ decays*, arXiv:2409.12629, Submitted to JHEP.

[45] LHCb collaboration, R. Aaij *et al.*, *Search for $B_{(s)}^{*0} \to \mu^+\mu^-$ in $B_c^+ \to \pi^+\mu^+\mu^-$ decays*, arXiv:2409.17209, Submitted to Eur. Phys. J. C.

[46] LHCb collaboration, R. Aaij *et al.*, *Measurement of the effective leptonic weak mixing angle*, arXiv:2410.02502, Submitted to JHEP.

[47] LHCb collaboration, R. Aaij *et al.*, *Observation of $J/\psi p$ resonances consistent with pentaquark states in $\Lambda_b^0 \to J/\psi pK^-$ decays*, Phys. Rev. Lett. **115** (2015) 072001, arXiv:1507.03414.

[48] LHCb collaboration, *Simultaneous determination of the CKM angle $\gamma$ and parameters related to mixing and CP violation in the charm sector*, LHCb-CONF-2024-004, 2024.

[49] O. B. Garcia *et al.*, *High-density gas target at the LHCb experiment*, Physical Review Accelerators and Beams (in press) (2024) arXiv:2407.14200.

[50] LHCb collaboration, *Upsilon mass peaks in 2024 data*, LHCb-FIGURE-2024-025 (2024).

[51] LHCb collaboration, *Hadron PID Performance in 2024*, LHCb-FIGURE-2024-031 (2024).

[52] LHCb collaboration, *HLT1 trigger efficiencies in 2024 data*, LHCb-FIGURE-2024-030 (2024).

[53] C. Bozzi, *LHCb Computing Resources: preliminary 2026 requests*, LHCb-PUB-2024-005 (2024).

[54] European Strategy Group, *Deliberation document on the 2020 update of the European Strategy for Particle Physics*, CERN-ESU-014, 2020.

[55] Large Hadron Collider Committee, *LHC Experiments Phase IIb Upgrades Approval Process*, CERN-LHCC-2022-012, 2022.