

ABSTRACT

Title of dissertation: Study of the decay $B^+ \rightarrow K^+\pi^0$ at LHCb
and mechanical development for the design
of the Upstream Tracker

Jason Emory Andrews
Doctor of Philosophy, 2018

Dissertation directed by: Abolhassan Jawahery
Department of Physics

The LHCb experiment at the Large Hadron Collider (LHC) is designed to measure the properties of particles containing charm (c) and bottom (b) quarks. This dissertation documents two major studies I have completed, one analyzing data collected by the LHCb detector, and another contributing to the design and development of an extensive upgrade to the detector.

The pattern of CP asymmetry measurements of the $B \rightarrow K\pi$ family of decays deviates from expectations derived from the standard model (SM), a contradiction known as the “ $K\pi$ puzzle.” The present size of the experimental errors are such that more precise measurements in the $B^+ \rightarrow K^+\pi^0$ decay channel are especially important. An analysis of the $B^+ \rightarrow K^+\pi^0$ decay using data collected during Run 1 is performed. Despite low reconstruction and trigger efficiencies and enormous combinatorial backgrounds, a signal is found with a statistical significance of 3.7σ . This achievement has led to the creation of a dedicated $B^+ \rightarrow K^+\pi^0$ trigger, and has inspired the creation of a number of dedicated triggers for decay modes with similar

topologies. A preliminary analysis of data collected during Run 2 demonstrates that the new trigger is a major success, with excellent prospects for making the world's best measurements in the $B^+ \rightarrow K^+\pi^0$ decay channel using the entire Run 2 data set.

Run 2 of the LHC will conclude at the end of 2018, and will be followed by Run 3, scheduled to begin in early 2021. In the interim, the LHCb detector will be upgraded to be read-out in real-time at 40 MHz, and to withstand the radiation damage associated with collecting 50 fb^{-1} of integrated luminosity by the conclusion of Run 4. A key part of this upgrade is the design and construction of a new silicon-strip tracking detector—the upstream tracker (UT). Regions at the periphery of the UT suffer from severe electrical and mechanical constraints, making a high-fidelity CAD model a critical element of the design process. The result is a mechanical integration solution that is entirely non-trivial, and which has had significant influences on the UT design. This solution and the constraints that influence it are shown in detail.

Study of the decay $B^+ \rightarrow K^+ \pi^0$ at LHCb and mechanical
development for the design of the Upstream Tracker

by

Jason Emory Andrews

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2018

Advisory Committee:

Professor Abolhassan Jawahery, Committee Chair/Adviser

Assistant Professor Alberto Belloni

Professor Zackaria Chacko

Professor Nicholas Hadley

Professor Richard Mushotzky, Dean's Representative

© Copyright
Jason Emory Andrews
2018

Acknowledgments

I am deeply indebted to my adviser, Professor Hassan Jawahery, for his unflagging support during my years here at UMD. It has been exciting to be around his enthusiasm, expertise, and passion, and it has been an honor to be a part of the superb team of researchers he has cultivated. His guidance has allowed me to overcome many daunting hardships, and his encouragement has enabled me to perform to the best of my abilities.

To my fellow members of Hassan's research group, I owe so much thanks for having patience with my many questions, and for providing such valuable input and assistance, all while maintaining an atmosphere of good humor and camaraderie.

Being part of the LHCb collaboration has been an honor, and a great privilege. I have received the benefit of the many years of hard work and ingenuity of hundreds of experts and students from many walks of life. The value of having access to the LHCb data and tools, and the LHC computing resources is absolutely priceless. I hope that I have in return contributed to that value for the current and future members of the collaboration.

Finally, I would like to thank and acknowledge the University of Maryland and the National Science Foundation for supporting me during my PhD program.

Table of Contents

Acknowledgements	ii
List of Tables	v
List of Figures	vi
List of Abbreviations	xiv
1 Introduction	1
2 Theory	3
2.1 Ingredients	3
2.2 Gauge theories	4
2.3 Electroweak theory	6
2.4 CP violation	7
2.5 CP violation in meson decays	10
2.6 The $K\pi$ puzzle	13
3 The LHCb experiment	16
3.1 The LHC	16
3.2 The LHCb detector	17
3.2.1 Magnet	21
3.2.2 Tracking	22
3.2.3 Particle identification	29
3.2.4 Trigger	37
3.2.5 Running conditions	40
3.2.6 Computing	41
4 Study of the decay $B^+ \rightarrow K^+\pi^0$	43
4.1 Introduction	43
4.2 Data	46
4.3 Event selection	46
4.3.1 Reconstruction techniques	46

4.3.2	Basic selection requirements	50
4.3.3	Particle identification	53
4.3.4	Mother-trajectory DOCA	54
4.3.5	Isolation techniques	56
4.3.6	Summary of preliminary selection	59
4.3.7	Multivariate selection	60
4.3.8	Optimization of the final selection	68
4.4	Peaking backgrounds	69
4.5	Results of the Run 1 analysis	70
4.6	Implications and future analyses	72
4.7	Dedicated $B^+ \rightarrow K^+\pi^0$ trigger	74
4.7.1	Resources and constraints	74
4.7.2	Event selection	77
4.8	Preliminary Run 2 results	78
4.8.1	Data	78
4.8.2	Event selection and results	80
4.8.3	Prospects	84
5	The LHCb upgrade	87
5.1	Introduction and strategy	87
5.2	Upgrade plans	89
5.2.1	Tracking	89
5.2.2	Particle identification	94
5.2.3	Trigger	96
6	The Upstream Tracker in the LHCb upgrade	98
6.1	Overview and geometry	98
6.2	Sensors and front-end electronics	103
6.3	Peripheral electronics	114
6.4	Mechanical design	118
6.4.1	External and internal constraints	118
6.4.2	Overview	121
6.4.3	Mechanical design of the peripheral interface region	125
6.4.4	Mechanical design of the PEPI	135
6.4.5	Interface of the box halves	147
6.4.6	Summary	148
	Bibliography	149

List of Tables

2.1	Measurements of $\mathcal{A}^{\text{CP}}$ for the $B \rightarrow K\pi$ family of decays	14
2.2	Measurements of $\mathcal{A}^{\text{CP}}$ for the $B \rightarrow K^+\pi$ decay modes	15
3.1	Data-taking conditions at the LHCb IP during Run 1 of the LHC, and the integrated luminosity recorded by the LHCb detector.	40
4.1	Numbers of simulated events generated at the center-of-mass energies of 7 TeV and 8 TeV, corresponding to the 2011 and 2012 data-taking conditions, respectively.	46
4.2	Requirements for candidate $B^+ \rightarrow K^+\pi^0$ decays that define the pre- liminary selection for the 2011(2012) data-taking period, where the merged π^0 algorithm is used to reconstruct π^0 candidates.	59
4.3	Discriminating variables used to create the BDT classifiers. Some variables are transformed by the natural logarithm to create broader distributions, which can result in more efficient training.	63
4.4	Selection requirements for $B^0 \rightarrow K^+\pi^-$ candidates. Requirements added to make this and the $B^+ \rightarrow K^+\pi^0$ decay selections more simi- lar are listed separately from the pre-existing requirements. Generic hadrons are denoted as h	65
4.5	Geometric, reconstruction, selection, and total efficiencies estimated from simulation, with statistical uncertainties.	69
4.6	Requirements for candidate $B^+ \rightarrow K^+\pi^0$ decays that define the se- lection for the Run 2 trigger, where the merged π^0 algorithm is used to reconstruct π^0 candidates.	79
4.7	Measurements of $\mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+\pi^0)$, and the world average. The first uncertainties are statistical, and the second are systematic. For the CLEO measurement, only a statistical error is quoted as the sys- tematic error is estimated to be negligible in comparison.	85
6.1	A stack-up of the DataFlex layers.	109

List of Figures

2.1	A summary diagram of the SM particles and their properties. For each particle, the mass, charge, and spin are indicated in descending order in the upper-left.	4
2.2	The lowest-order box diagrams for the $K^0/\overline{K}^0$, $B^0/\overline{B}^0$, and $B_s^0/\overline{B}_s^0$ mesons, where q_d is a down-type quark and q'_d is a distinct down-type quark. The corresponding diagrams for the $D^0/\overline{D}^0$ mesons can be produced by the substitutions $q_b \rightarrow c$, $q'_b \rightarrow u$, and $(t, c, u) \rightarrow (b, s, d)$	10
2.3	Penguin and color-favored tree diagrams of the $B \rightarrow K\pi$ decays. The spectator quark q and the q' quark are each either a u or a d quark.	13
2.4	Color-suppressed tree diagrams of the $B \rightarrow K\pi^0$ decays. The spectator quark q is either a u or a d quark.	14
3.1	Production rate of $b\bar{b}$ pairs at $\sqrt{s} = 7$ TeV as a function of decay angle. The rate normalization is arbitrary. The LHCb acceptance is indicated in red.	18
3.2	Production rate of $b\bar{b}$ pairs at $\sqrt{s} = 7$ TeV as a function of pseudorapidity. The rate normalization is arbitrary. The LHCb acceptance is indicated by the red box in the forward region, and that of a typical GPD is indicated by the yellow box.	18
3.3	Cross-sectional view of the CMS and LHCb detectors. The images of the detectors are to-scale, and the LHCb IP is placed over that of the CMS detector.	19
3.4	Cross-sectional view of the LHCb detector. The z axis of the LHCb coordinate system points to the right.	20
3.5	Perspective view of the LHCb detector. Obvious features are the magnet in blue, the RICH 2 detector in yellow, and the muon system in green.	20
3.6	Perspective view of the magnet from downstream; the VELO is behind the magnet in this view.	21
3.7	The strength of the magnetic field along the beam-line as a function of z	22

3.8	Schematics of the VELO sensors. Top: top-view showing the relative positions of the sensors along the beam-line. Bottom: Views in the x - y plane showing the two halves of the VELO closed (left), and open (right). Note the overlapping profiles of the sensor halves in the closed position.	23
3.9	Diagram of the planes of the TT, looking downstream. The rotated configuration of the middle u and v planes is shown.	25
3.10	Perspective view of the TT and the T-stations. The TT is shown in purple on the left, the IT and OT stations are shown on the right in purple and teal, respectively.	25
3.11	Perspective view of the configuration of an IT station. The z axis points along the pipe towards the upper right.	26
3.12	Top views of one plane of the OT, showing the staggered two-layer configuration of the straw-tube modules.	27
3.13	Plots demonstrating the performance of the tracking system. Left: momentum resolution $\delta p/p$ as a function of p . Right: IP resolution as a function of $1/p_T$. The typical distributions of B mesons at LHCb are shown below the plots.	28
3.14	An illustration showing different types of tracks superimposed over the LHCb tracking system. The view is in the x - z plane—the bending plane of the magnet.	29
3.15	Schematic of particle signatures in a typical configuration of a modern particle detector experiment.	30
3.16	Segmentations of the calorimetry sub-detectors. Left: segmentation of the SPD, PS, and ECAL. Right: segmentation of the HCAL. . . .	31
3.17	Diagram of an HCAL module, showing the arrangement of the scintillating and absorbing tiles.	32
3.18	Side view of the muon system. Note that the M1 station is positioned in front of the calorimetry sub-detectors.	34
3.19	Plot of the Cherenkov angle θ_C vs. momentum for the RICH radiators.	36
3.20	Background rejection vs. signal efficiency for four of the global PID variables, derived from data sidebands and signal MC in the study of the hypothetical $\Sigma^+ \rightarrow p\mu^+\mu^-$ decay. DLL variables are drawn as black circles, and ProbNN variables are drawn as red stars.	37
4.1	Mass distributions of B^+ mesons from simulated $B^+ \rightarrow K^+\pi^0$ decays, reconstructed using the merged (a) and resolved (b) π^0 algorithms. The dashed red lines in (b) indicate the range of reconstructed B^+ mass values ultimately allowed in the even selection.	48
4.2	Mass distributions of B^+ candidates from a sample of experimental data, reconstructed using the merged (a) and resolved (b) π^0 algorithms.	49

4.3	Signal and background distributions for the mass (a), p (b), and p_T of the π^0 candidates (c). Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.	50
4.4	Signal and background distributions for the p_T (a) and $\ln(\Delta\chi_{\text{IP}}^2)$ (b) of the K^+ candidates. The $\Delta\chi_{\text{IP}}^2$ is transformed by the natural logarithm in order to broaden the sharply peaking distributions, making the discriminating power of the variable more apparent. Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.	51
4.5	Signal and background distributions for the mass (a), p (b), p_T (c), and η of the B^+ candidates (d). Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.	52
4.6	Signal and background distributions for the ProbNNk of the K^+ candidates (a) and the CL of the π^0 candidates (b). Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.	54
4.7	MT-DOCA and the reconstructed $B^+ \rightarrow K^+\pi^0$ decay topology. . . .	55
4.8	Signal and background distributions of $\log_{10}(\chi_{\text{MT-DOCA}}^2)$ for the K^+ candidates. Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.	56
4.9	Signal and background distributions for the $V_{\text{Mult.}}$ (a), $\ln(\min(\Delta\chi_{\text{vtx}}^2))$ (b), and, $\mathcal{A}^{p_T}$ of the B^+ candidates (c). Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.	58
4.10	A diagram of a decision tree. The root node represents the initial data set, and the other nodes represent subsets. At each node, the most discriminating variable (xi, xj, xk) is identified, as well as the criterion ($c1, c2, \dots, c4$) that best segregates the signal (S) and background (B) distributions of that variable. Notice that variables can appear more than once and with different criterion (xj here, with criterion $c1, c2$).	61

4.11	The $K^+\pi^-$ invariant mass spectrum. The data sets are drawn as black points. The total fit models are drawn as solid blue lines, with $B^0 \rightarrow K^+\pi^-$ “signal” components drawn in dashed blue, $B_s^0 \rightarrow K^+K^-$ components drawn in dashed gold, $B_s^0 \rightarrow \pi^+K^-$ components drawn in dashed green, and combinatorial backgrounds drawn in dashed red.	65
4.12	Isolation variable distributions from simulation both before and after correction weights are applied. The raw simulation distributions are drawn in black, and the weighted distributions are drawn in red.	67
4.13	Distributions of the response variables generated by the BDT classifiers. The distributions are continuous, in contrast to that of a binary decision tree. Signal distributions are drawn in blue, and background distributions are drawn in red.	67
4.14	The figure of merit $N_S/\sqrt{N_S+N_B}$ as a function of cutting on the combined BDT classifier response.	68
4.15	Mass distribution of the B^+ candidates, selected from the entire Run 1 data set. The data are drawn as black points. The total fit model is drawn as a solid blue line. The signal component is drawn in dashed blue, the combinatorial background is drawn in dashed purple, and the low-mass background is drawn in dashed magenta.	71
4.16	A plot over time of the integrated luminosity of proton-proton collisions recorded by the LHCb detector during the years 2010–2017. Beam energies and integrated luminosities by year are given in the figure. An additional 1.7 fb^{-1} at a beam energy of 6.5 TeV is anticipated to be recorded in the final year of Run 2: 2018.	80
4.17	Mass distribution of the B^+ candidates, selected from the 2016 data set. The data are drawn as black points. The total fit model is drawn as a solid black line. The signal component is drawn in blue, the combinatorial background is drawn in in green, and the partially reconstructed backgrounds are drawn together in dashed red. The component representing partially reconstructed backgrounds where only one decay product is not reconstructed is drawn in yellow, while that representing cases where two or more decay products are not reconstructed is drawn in red.	82
5.1	The upgraded LHCb detector. Changes that are immediately apparent are the replacements of the VELO with a pixel-based system, the replacements of the TT and T-stations by the UT and SciFi trackers, respectively, and the removal of the PS/SPD detectors and the M1 muon station.	89
5.2	A perspective rendering of the upgraded VELO detector without the RF foils. The sensors are shown in red, the VeloPix ASICs are shown in yellow, and the cooling substrate is shown in blue. Notice how the backsides of the ASICs are visible underneath the sensors for tiles positioned on the opposite module face.	90

5.3	Side view of the upgraded VELO module, showing the arrangement of the VeloPix, the sensors, and the silicon substrate with embedded micro-channels. Units are in μm	91
5.4	Side and front views of a SciFi Tracker station with a representation of a human figure for scale.	92
5.5	Trigger and tracking schemes of the upgrade.	96
6.1	Plots of the fluence (left) and dose (right) deposited by 50 fb^{-1} of integrated luminosity as a function of the y coordinate and with $x = 0$, as estimated from simulation. The z coordinates plotted here correspond to the extents of the UT active area.	99
6.2	A diagram of the UT detector planes with names, as seen from the VELO. Different colored cells represent different types of sensors. The grey rectangle highlights an example of a mechanical stave.	100
6.3	An example of a stave with various elements labeled. The general structure shown here is common to all staves, though this particular example is of a central stave, as evidenced by the semi-circular cutout for the beam-pipe.	101
6.4	A layer of the carbon-fiber foam stave core. The titanium CO_2 cooling tube is shown in purple, superimposed over the core. In assembly it is sandwiched in a milled channel between this and another layer. . .	103
6.5	An illustration of the different types of silicon-strip sensors, superimposed to-scale over the raw silicon wafers. Color-codes in this diagram match that of Fig. 6.2.	104
6.6	An exploded view of the module diagramming the main components. The anchor tabs of the stiffener extend beyond the profile of the hybrid and the sensor, allowing the module to be removed and replaced by breaking the stiffener at the tabs.	105
6.7	A block diagram of the SALT chip. The 128 analog channels are represented on the left. The data output and TFC input SLVS e-ports are on the right.	107
6.8	An example of the signal and power layers of a DataFlex cable. The LV power planes are drawn in green, and the signal and HV traces are drawn in red. The large green rectangular element on the left end of the cable is a temporary stiffener for mounting the BGA connector. The location of the BGA connector is clearly indicated by the compact array of bonding pads. The large pads at the left-most edge are for the HV connectors. The HV traces are routed along the perimeter of the LV and signal routings.	110
6.9	An oscilloscope reading from a test of a prototype DataFlex cable. The lower half of the readout shows an eye diagram. That the eye is open indicates that the flex cable has good signal integrity.	111

6.10	The geometries of the three DataFlex variations, including positions of the detector modules. The BGA and HV connectors are depicted at the top. The length of the longest cable will ultimately be determined by the final geometry of the Type D sensor.	112
6.11	A photograph of the end of the DataFlex nearest the beam-pipe. The wire bonding pads of the hybrid are visible at the middle of the edge of the cable. High-voltage bonding pads are visible on the right. The module positioned here is mounted directly on the stave—entirely off of the DataFlex. The cut-outs shown here accommodate the anchor tabs of that module, to prevent it from otherwise being mounted at a slight angle due to the thickness of the flex cable.	113
6.12	A photograph of the end of the DataFlex nearest the detector periphery. Mounting pads for the 400-position BGA connector are visible in the center of the image, and those of the HV connectors are visible on the lower-right. A temporary stiffener for mounting the connectors is also visible under the flex cable.	113
6.13	A diagram of the PEPI electronics, and their interface to the front-end SALT ASICs. Dominating the diagram are the data control boards (DCBs), which host the ECS, TFC, and data concentrator functions. The dense interconnect backplanes are entirely passive.	115
6.14	A rendering of the backplane. The BGA connectors along the bottom of the backplane are pigtail connectors. The larger BGA connectors across the middle of the backplane are DCB connectors. Two DCBs are shown for reference. The VTTx and VTRx connectors are visible at the edge of the DCBs. LV breakout boards are shown connected along the top of the backplane.	116
6.15	A view of the TT station from the Access-side (A-side) of the cavern. The TT is visible in the center in its closed position. The rails and cable chains are above and below the TT. The blue RICH 1 is to the right of the TT, and the yellow coils of the magnet are visible on the left. The service bays are in front of the magnet yoke.	119
6.16	A detail of the PEPI region of the UT corresponding to the view in Fig. 6.15. The RICH 1 and hardware supporting the rails is shown in yellow on the right. The arms used to support the beam-pipe are shown in red on the left. Note the “I” shaped cross section of the blue I-beam between the two PEPI chassis.	120
6.17	A perspective view of the UT station. The UT is in the center in the closed position. The yellow rails are above and below the UT, and the cable chains are in front of them. At the ends of the chains are the service bays. The cable chains and service bays shown here are those of the TT, acting as stand-ins.	122
6.18	Renderings of the polymer gaskets that interface with the beam-pipe. The gaskets are shown in green and the detector box is shown in blue. Only one half of the detector box and gaskets are shown for clarity. .	124

6.19	The region at the ends of the staves viewed from the front (a) and side (b). For clarity, the side view is also drawn in perspective (c). The I-beam is rendered in violet, and the detector box is rendered in pale blue. The front and back faces of the detector box are omitted in (a) and (c). All of the PEPI electronics are omitted for clarity. The remaining components will be identified in the figures that follow.	126
6.20	The region at the ends of the staves, where only the I-beam, detector box, and stave frames are included. The stave frames are rendered in red. The individual frames are attached to rails, which are in turn attached to the inner face of the box. The white vertical shapes conspicuous in the side view are fixtures that are embedded in the side of the detector box through which the HV cables pass.	128
6.21	A rendering including the HV cables in teal. A single cable for each stave enters through the side of the box, and breaks out into individual insulated wires to attach directly to the DataFlex cables.	129
6.22	A rendering including the cooling tubes and manifolds in grey. The tubes pass over the HV cables and through the openings in the stave frames to reach the staves.	130
6.23	A close-up view showing the mechanical interfaces of the CO ₂ tubes, stave frames, and HV cables.	131
6.24	A rendering of a single cooling tube. The circular winding allows the tube to flex during the installation of the staves. There are three mechanical variations of the part shown here for a total of four geometries to reach the different stave positions.	131
6.25	A rendering including the staves attached to the stave frames. Bolts with stand-offs establish the staggering of adjacent staves in z	133
6.26	A rendering including the pigtails in blue. Serpentine geometry of the sections inside the box provide extra flexibility.	134
6.27	A rendering of the PEPI chassis region showing only the detector box, C-frame, and backplane base plates.	135
6.28	A rendering including the three backplanes attached to the base plates.	136
6.29	A rendering including the pigtails connected to the backplanes. . . .	137
6.30	A rendering of the pigtails that connect the stereo-rotated staves to the backplanes. They are shown here as they are manufactured, before they are deformed into their installed shapes.	138
6.31	A rendering of the pigtails that connect the non-rotated staves to the backplanes. They are shown here as they are manufactured, before they are deformed into their installed shapes.	139
6.32	Tracks that constrain the pigtails to their corresponding position on the backplane. They provide for travel along the z axis, allowing the pigtails to be connected and disconnected from the backplane while fixed in place in the x - y plane.	139
6.33	Photographs of a 3D-printed mock-up of the pigtails.	141
6.34	A photograph of a prototype pigtail subjected to deformations similar to the planned geometry of pigtails installed in the UT.	142

6.35	A rendering including the DCB cooling pipes and support structures.	143
6.36	An example of an optical MPO-to-LC harness that breaks out into twelve LC connectors.	144
6.37	A rendering including the MPO-to-LC optical harnesses connecting the DCBs to the end-cap. The optical cables will be routed either above or in front of the DCBs.	145
6.38	An illustration of the mechanical coupling of the two PEPI chassis at the opposite ends of a stave.	146
6.39	A rendering of a detector box half showing the shape of the box-halves interface.	147
6.40	A view of the interface of the two detector halves where the two halves are partially retracted. The top plane of the I-beam is omitted for clarity.	148

List of Abbreviations

AdaBoost	Adaptive Boosting
ADC	Analog-to-Digital Converter
ALICE	A Large Ion Collider Experiment
ASIC	Application-Specific Integrated Circuit
ATLAS	A Toroidal LHC Apparatus
BDT	Boosted Decision Tree
BGA	Ball Grid Array
CDF	Collider Detector at Fermilab
CERN	European Organization for Nuclear Research
CKM	Cabibbo Kobayashi Maskawa
CMS	Compact Muon Solenoid
DCB	Data Control Board
DLL	Delta Log Likelihood
DOCA	Distance Of Closest Approach
DSP	Digital Signal Processing
ECAL	Electromagnetic Calorimeter
ECS	Experiment Control System
EFF	Event Filter Farm
GANGA	Gaudi/Athena and Grid Alliance
GEM	Gas Electron Multiplier
GPD	General Purpose Detector
HCAL	Hadronic Calorimeter
HLT	High Level Trigger
HPD	Hybrid Photon Detector
HV	High Voltage
I2C	Inter-Integrated Circuit
IT	Inner Tracker
IP	Impact Parameter
IP	Intersection Point
L0	Level-0 (zero)
LCG	LHC Computing Grid
LEP	Large Electron-Positron Collider
LHC	Large Hadron Collider
LS2	Long Shutdown 2
LV	Low Voltage
MaPMT	Multianode Photomultiplier
MT-DOCA	Mother Trajectory DOCA
MWPC	Multi-wire Proportional Chamber
OT	Outer Tracker

PCB	Printed Circuit Board
PDG	Particle Data Group
PEPI	Peripheral Electronics Processing Interface
PID	Particle Identification
PMT	Photomultiplier Tube
PS	Preshower Detector
PS	Proton Synchrotron
RICH	Ring Imaging Cherenkov
SALT	Silicon ASIC for LHCb Tracker
SciFi	Scintillating Fiber
SiPM	Silicon Photomultiplier
SLVS	Scalable Low-Voltage Signaling
SM	Standard Model
SPD	Scintillating Pad Detector
SPS	Super Proton Synchrotron
SSB	Spontaneous Symmetry Breaking
SV	Secondary Vertex
TFC	Timing and Fast Control
TOS	Triggered On Signal
TIS	Triggered Independent of Signal
TT	Tracker Turicensis
UT	Upstream Tracker
QCD	Quantum Chromodynamics
QED	Quantum Electrodynamics
QFT	Quantum Field Theory
VELO	Vertex Locator
VEV	Vacuum Expectation Value

Chapter 1: Introduction

The field of particle physics was born in the 20th century, and tremendous advancements were made both in experimental and in theoretical development, culminating in what we know today as the Standard Model of particle physics (SM). The SM is an exceptionally successful theory, but it is an incomplete description of the universe. It does not describe gravity, and has no explanation for dark matter or dark energy.

Another shortcoming is the amount of CP violation predicted by the SM. The present-day universe is matter-dominated, but studies of the early universe indicate a high degree of balance between matter and antimatter in the first few moments of the universe's history [1–3]. The evolution of this early balanced state to the current matter-dominated state is called baryogenesis, and it requires three ingredients: thermal non-equilibrium, baryon number violation, and CP violation [4]. The SM includes CP violation, but the amount it provides is too small by several orders of magnitude to account for baryogenesis [5–7]. This is a strong hint that CP violation studies may be a way towards finding physics beyond the SM.

The LHCb experiment is built to stringently test the SM in several important areas, especially CP violation. This dissertation documents work that I have done as part of the LHCb collaboration. After establishing some theoretical background in Chapter 2 and describing the LHCb experiment in Chapter 3, I describe two major studies that I have completed. The first is a novel analysis of experimental data which investigates an inconsistency in measurements of CP violation in B meson decays from expectations derived from SM. This is covered in Chapter 4, and is followed by a description of the upgrade of the LHCb detector in Chapter 5. The second project is a contribution to the mechanical design of a new tracking detector for the experiment, described in Chapter 6.

The entirety of Chapter 4 documents my original work with two exceptions. Firstly, the $\Lambda_b \rightarrow \Lambda \gamma$, $\Xi_b^- \rightarrow \Xi^- \gamma$, and $\Omega_b \rightarrow \Omega \gamma$ triggers referred to in Sections 4.1 and 4.6 were created by Albert Puig at the University of Zurich. Secondly, the preliminary Run 2 results presented in Section 4.8 is the work of Will C. Parker, one of my fellow researchers at the University of Maryland.

Regarding the work described in Chapter 6, Section 6.4 presents integration designs that I was instrumental in creating, with the exception of the beam-pipe gasket in Fig. 6.18. Many of the figures in this section show models which contain elements that are the original creations of others, especially Figures 6.19 through 6.26, and Fig. 6.35. The overall integration design that these models embody is my primary contribution to the work described in Chapter 6.

Chapter 2: Theory

2.1 Ingredients

The SM is a relativistic quantum field theory (QFT) that describes matter particles known as fermions, their interaction via three of the four known fundamental forces, and the Higgs particle. Fermions are spin-1/2 particles, and the fundamental fermions we find in nature are the three generations of leptons and the three generations of quarks. Each of these twelve fermions also has an anti-matter counterpart. The three generations are distinguished by their increasing masses, except for the extremely low-mass neutrinos. The nature of neutrino mass is an open research question, and it may turn out that they have a different mass hierarchy than that of their corresponding leptons. The fundamental forces are represented by gauge bosons. The massless photon is the gauge boson of the electromagnetic force, and eight massless gauge bosons called gluons mediate the strong nuclear force. Three gauge bosons known as the W^+ , W^- , and Z^0 mediate the weak nuclear force, and unlike the photon and the gluons, they are massive. Finally, the Higgs boson is a special ingredient responsible for the mass of the weak bosons and the fermions. A summary of these particles and their properties are given in Fig. [2.1](#).

Standard Model of Elementary Particles




Figure 2.1: A summary diagram of the SM particles and their properties [8]. For each particle, the mass, charge, and spin are indicated in descending order in the upper-left.

2.2 Gauge theories

When attempting to describe nature, it pays to notice that it has a number of fundamental symmetries. Emmy Noether discovered that when a theory has a symmetry, there is a corresponding conserved quantity [9]. For example, we know that nature has certain fundamental “laws” like the conservation of energy and the conservation of momentum. But if we realize that these are consequences of

invariance with respect to time and position translations, then we do not have to worry about making those laws distinct postulates. We can, in general, just restrict ourselves to those theories that have the underlying symmetries of spacetime translations.

Gauge symmetries in particular turn out to be very important. As an example, consider the Lagrangian of a free spin-1/2 particle, given by

$$\mathcal{L} = \bar{\psi}(i\not{\partial} - m)\psi. \quad (2.1)$$

This Lagrangian is invariant under the global gauge transformation $\psi \rightarrow e^{iq\theta}\psi$, where q and θ are constants. Now suppose we are inspired to say that nature has a local gauge symmetry—*i.e.*, $\theta \rightarrow \theta(x)$. We can accomplish this for massless spin-1/2 particles by introducing a spin 1-particle A_μ that interacts with the spin-1/2 particle as $q\bar{\psi}\gamma^\mu\psi A_\mu$. The resulting Lagrangian, including the free field of the newly introduced boson is

$$\mathcal{L} = i\bar{\psi}\not{\partial}\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} - q\bar{\psi}\gamma^\mu\psi A_\mu, \quad (2.2)$$

and the symmetry transformation is

$$\psi \rightarrow e^{iq\theta}\psi, \quad A_\mu \rightarrow A_\mu + \partial_\mu\theta. \quad (2.3)$$

With knowledge of spin-1 and spin-1/2 fields, and the application of symmetry, we have arrived at a striking result: Equation 2.2 is the theory of quantum electrodynamics (QED).

That a Lagrangian can be given local gauge symmetry by introducing new vector fields is known as the “gauge principle,” and the new particles as “gauge

bosons.” The gauge principle described here can be extended to non-Abelian gauge transformations, and it can be applied elegantly through group theory; we can simply say that QED is a gauge theory with $U(1)$ symmetry. The explicit gauge transformations and the interaction terms in the Lagrangian follow from this simple statement. Likewise, quantum chromodynamics (QCD) can be said to be a gauge theory with $SU(3)$ symmetry. This not only accounts for the 8 massless gluons, but also that the gluons carry color charge and hence interact with themselves.

2.3 Electroweak theory

The gauge principle can guide us in creating a theory of weak interactions, though there are some complications. Suppose we start with an $SU(2)$ model. This introduces the requisite number of 3 gauge bosons, but they are massless. Yet it is possible to generate pure mass terms for gauge bosons using a phenomenon known as spontaneous symmetry breaking (SSB). We introduce a new scalar field that does indeed have local gauge symmetry, but also has degenerate vacuum expectation values (VEVs). The physical system modeled by the theory must have some specific VEV, and this breaks the symmetry present in the dynamics of the theory. This symmetry breaking, however, preserves the gauge symmetry of the system. Symmetry breaking of this form is said to be “spontaneous” when there is no external mechanism dictating the choice of one VEV over another.

In the SM, we start with a model having $SU(2) \otimes U(1)$ symmetry and then implement SSB by introducing the Higgs field. This reduces the $SU(2) \otimes U(1)$ sym-

metry to a new $U(1)$ symmetry, namely, the $U(1)_Q$ symmetry of QED, where Q stands for the conserved electric charge arising from the symmetry. The $SU(2)$ bosons are known as W_1 , W_2 , and W_3 weak isospin bosons, and the $U(1)$ boson is known as the B hypercharge boson. After symmetry breaking, the W_1 and W_2 bosons combine to make the massive $W^\pm$ bosons, and the W_3 and B combine to make the massive Z^0 and the photon γ , which remains massless.

Furthermore, the Higgs mechanism can explain the mass of the fermions. The Lagrangian of Equation 2.1 is not, in fact, compatible with local gauge symmetry; the mass term $m\bar{\psi}\psi$ breaks the symmetry, consequentially requiring the fermions to be massless to preserve the symmetry. Through the SSB mechanism, the fermion interaction terms are generated, and these Yukawa interactions produce mass terms for the fermions, without breaking the gauge symmetry. The Higgs mechanism is thus very compelling, and was a central prediction of the SM, finally vindicated by the announcement of the discovery of the Higgs particle in 2012 [10, 11].

2.4 CP violation

The parity transformation P inverts the three spatial coordinates, leaving scalars unchanged, inverting spatial vectors, and so on. It is equivalent to a rotation by π radians followed by a reflection through the axis of rotation. This mirror reflection exchanges the left- and right-chiral handedness of particles, and it is the concept at the heart of the parity transformation.

The weak force is known to “maximally” violate P because it couples exclu-

sively with left-handed fermions. It was assumed for some time that the combined transformation of P with charge conjugation, CP, was a good symmetry of the weak interactions, but even CP has been shown to be violated in the K , B , and B_s meson systems [12–15]. CP violation is therefore a necessary feature of the SM, and it is in fact accommodated when there are at least three quark generations.

In the electroweak unification scheme, the weak eigenstates of the quarks and their mass eigenstates do not have a common basis. They are related through the 3×3 Cabibbo-Kobayashi-Maskawa (CKM) matrix V :

$$d'_i = V_{ij}d_j, \quad (2.4)$$

where the d'_i are the weak eigenstates and the d_j are the mass eigenstates. For illustration, consider the charged-current interaction terms such as

$$\frac{g}{\sqrt{2}}\overline{u_{iL}}V_{ij}\gamma^\mu W^-_\mu d_{jL} + \frac{g}{\sqrt{2}}\overline{d_{iL}}V_{ij}^*\gamma^\mu W^+_\mu u_{iL} \quad (2.5)$$

We see that V_{ij} modulates the transition from the i -th down-type quark to the j -th up-type quark via W^- emission, while up-type to down-type via W^+ emission is modulated by V_{ij}^* .

As a probability conserving change-of-basis transformation, the CKM matrix is necessarily unitary. A general 3×3 unitary matrix has nine real parameters, but the six quark fields are each invariant under an arbitrary phase redefinition. Redefining all six fields simultaneously results in a physically meaningless overall phase that can be factored out of all the quark terms, so five quark phase redefinitions can be used to eliminate five V_{CKM} parameters. The CKM matrix can thus be fully defined by only four parameters:

$$V_{\text{CKM}} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta} & c_{23}c_{13} \end{pmatrix}, \quad (2.6)$$

where c_{12} is short for $\cos\theta_{12}$, etc. The three parameters $\theta_{12}, \theta_{13}, \theta_{23}$, are called mixing angles and the complex phase δ is the CP-violating parameter.

To see that non-zero values of δ cause CP-violation, consider again the charged-current interactions of Equation 2.5. Hermiticity requires such couplings to appear in Hermitian-conjugate pairs, but the CP operation exchanges the operators without modifying the coefficients, *e.g.*,

$$\begin{aligned} CP \Big[\frac{g}{\sqrt{2}} \overline{u_{iL}} V_{ij} \gamma^\mu W_\mu^- d_{jL} + \frac{g}{\sqrt{2}} \overline{d_{iL}} V_{ij}^* \gamma^\mu W_\mu^+ u_{iL} \Big] \\ = \frac{g}{\sqrt{2}} \overline{d_{iL}} V_{ij} \gamma^\mu W_\mu^+ u_{jL} + \frac{g}{\sqrt{2}} \overline{u_{iL}} V_{ij}^* \gamma^\mu W_\mu^- d_{iL}. \end{aligned} \quad (2.7)$$

Hence the CP transformation leaves such pairs unchanged only when $V_{ij} = V_{ij}^*$. All of the other SM couplings are real, so it is the complex nature of V_{ij} alone which generates CP violation.

The smallness of CP violation in the SM is not due to the smallness of the parameter δ , but rather the smallness of the off-diagonal V_{CKM} elements. This is shown more clearly in the Wolfenstein parameterization—an approximate parameterization in terms of parameters (A, η, λ, ρ) , where $\lambda \equiv |V_{us}| \approx 0.225$ acts as a smallness parameter:

$$|V_{\text{CKM}}| \approx \begin{pmatrix} 1 & \lambda & \lambda^3 \\ \lambda & 1 & \lambda^2 \\ \lambda^3 & \lambda^2 & 1 \end{pmatrix}. \quad (2.8)$$

The structure suggested by Equation 2.8 and its connection (if any) to the quark mass hierarchy is not currently understood.

2.5 CP violation in meson decays

Neutral CP conjugate mesons M^0 and $\bar{M}^0$ can transform into each other through “box” diagrams, illustrated in Fig. 2.2. Because of this mixing, the flavor eigenstates $M^0/\bar{M}^0$ are not mass eigenstates. The mass eigenstates are written M_L and M_H for “light” and “heavy” respectively, and are expressed in the flavor basis as

$$|M_L\rangle = p|M^0\rangle + q|\bar{M}^0\rangle, \quad |M_H\rangle = p|M^0\rangle - q|\bar{M}^0\rangle, \quad (2.9)$$

where p and q are complex numbers that normalize the mass eigenstates.

Mixing was first observed in the neutral kaon system [16], and has since been observed in the neutral D , B_d , and B_s systems [17–19]. The dynamics of meson




Figure 2.2: The lowest-order box diagrams for the $K^0/\bar{K}^0$, $B^0/\bar{B}^0$, and $B_s^0/\bar{B}_s^0$ mesons, where q_d is a down-type quark and q_d' is a distinct down-type quark. The corresponding diagrams for the $D^0/\bar{D}^0$ mesons can be produced by the substitutions $q_b \rightarrow c$, $q_b' \rightarrow u$, and $(t, c, u) \rightarrow (b, s, d)$.

mixing is characterized by the differences in the decay widths $\Delta\Gamma$ and masses Δm of the light and heavy states. In particular, the probability to mix depends sinusoidally on Δm .

CP violation can occur in the mixing of neutral mesons, in the interference between mixing and decay diagrams, or directly in the decay. In the SM, CP violation in B^0 and B_s^0 mixing is expected to be very small [20]. CP violation parameters in the $D^0/\overline{D}^0$ system are expected to be even smaller, and in addition are difficult to calculate [21]. The only system in which CP violation in mixing has actually been observed is the $K^0/\overline{K}^0$ system. This type of violation occurs when $|p/q| \neq 1$, which is traditionally expressed in the neutral kaon system through the parameter

$$\epsilon \equiv \frac{1 - \frac{q}{p}}{1 + \frac{q}{p}}. \quad (2.10)$$

The very first observation of CP violation [12] was the measurement of a non-zero value of ϵ . The best measurement of this parameter has been made by the KLOE collaboration, which found that [22]

$$\text{Re}(\epsilon) = (164.9 \pm 2.5) \times 10^{-5}. \quad (2.11)$$

CP violation in the interference between mixing and decay diagrams can occur when both M^0 and $\overline{M}^0$ can decay to the same final state f , for instance, in the decays $B^0/\overline{B}^0 \rightarrow J/\psi K_S^0$. The first observation of CP violation in the B^0 system was of this type, and was made by the BaBar and Belle experiments [13, 14], where they measured the CKM angles

$$\beta \equiv \arg\left(-\frac{V_{cd}V_{cb}^*}{V_{td}V_{tb}^*}\right) \quad \text{and} \quad \beta_s \equiv \arg\left(-\frac{V_{cb}V_{cs}^*}{V_{tb}V_{ts}^*}\right). \quad (2.12)$$

The LHCb collaboration has also recently measured $\sin 2\beta$ through the decay-time dependent asymmetry

$$\mathcal{A}(t) \equiv \frac{\Gamma(\overline{B^0}(t) \rightarrow J/\psi K_S^0) - \Gamma(B^0(t) \rightarrow J/\psi K_S^0)}{\Gamma(\overline{B^0}(t) \rightarrow J/\psi K_S^0) + \Gamma(B^0(t) \rightarrow J/\psi K_S^0)} \approx \sin 2\beta \sin \Delta m t, \quad (2.13)$$

finding that $\sin 2\beta = 0.731 \pm 0.035(\text{stat}) \pm 0.020(\text{syst})$ [23]. The SM predictions for this large CP violation are consistent with the experimental results. The SM prediction for the CP violating phase ϕ_s , on the other hand, is very small, making measurements of this parameter an important test of the SM. Analyses of the decays $B_s^0 \rightarrow J/\psi K^- K^+$ and $B_s^0 \rightarrow J/\psi \pi^- \pi^+$ made by the LHCb collaboration give a combined result of $\phi_s = -0.010 \pm 0.039$, consistent with both the SM and CP conservation [24, 25].

Direct CP violation occurs when

$$\mathcal{A}^{\text{CP}} \equiv \frac{A_f - \overline{A_{\bar{f}}}}{A_f + \overline{A_{\bar{f}}}} \neq 0, \quad (2.14)$$

where the decay amplitudes of CP conjugate mesons $M/\overline{M}$ to CP conjugate final states $f/\overline{f}$ via the weak interaction Hamiltonian $\mathcal{H}$ are

$$A_f = \langle f | \mathcal{H} | M \rangle, \quad \overline{A_{\bar{f}}} = \langle \bar{f} | \mathcal{H} | \overline{M} \rangle. \quad (2.15)$$

This is caused by interference between different amplitudes contributing to the total decay amplitude, and it has been established in the K , B , and B_s^0 systems. This is the only type of CP violation observed thus far in the B_s^0 system, having been established solely in the decay $B_s^0 \rightarrow K^- \pi^+$ [26].

2.6 The $K\pi$ puzzle

Direct CP violation in the B meson system was first established in the decay $B^0 \rightarrow K^+\pi^-$ [27]. This decay as well as the similar decay $B^+ \rightarrow K^+\pi^0$ receives significant contributions from radiative loop—or “penguin”—diagrams (Fig. 2.3a), as well as tree diagrams (Fig. 2.3b). Since the amplitudes are similar in magnitude, large CP violation is possible for these two decay modes. The other two decays in the $B \rightarrow K\pi$ family, $B^0 \rightarrow K^0\pi^0$ and $B^+ \rightarrow K^0\pi^+$ are dominated by penguin diagrams. The $B \rightarrow K\pi^0$ decays also receive contributions from color-suppressed tree diagrams (Fig. 2.4), but this is expected to be very small, thus no significant direct CP violation is predicted in this channel.

Measurements of $\mathcal{A}^{\text{CP}}$ made by the BaBar and Belle collaborations are shown in Table 2.1. As expected, no significant direct CP violation is found in the $B^0 \rightarrow K^0\pi^0$ or $B^+ \rightarrow K^0\pi^+$ decay modes, although the experimental uncertainties are large. The $B^0 \rightarrow K^+\pi^-$ decay, however, exhibits a large CP asymmetry, indicating penguin-




Figure 2.3: Penguin and color-favored tree diagrams of the $B \rightarrow K\pi$ decays. The spectator quark q and the q' quark are each either a u or a d quark.




Figure 2.4: Color-suppressed tree diagrams of the $B \rightarrow K\pi^0$ decays. The spectator quark q is either a u or a d quark.

tree interference. Simply exchanging the spectator d quark for a u quark in the penguin and tree diagrams of the $B^0 \rightarrow K^+\pi^-$ mode results in the diagrams for the $B^+ \rightarrow K^+\pi^0$ mode. Based on this and on the large CP violation measured in the $B^0 \rightarrow K^+\pi^-$ decay mode, one naïvely expects to find large CP violation in the $B^+ \rightarrow K^+\pi^0$ mode as well. Contrary to expectations, the measured $\mathcal{A}^{\text{CP}}$ in this mode is small and has the opposite sign, as shown in Table 2.1.

The LHCb collaboration has made the most precise measurement of $\mathcal{A}^{\text{CP}}$ for the $B^0 \rightarrow K^+\pi^-$ decay to date [28]. This measurement, as well as the $\mathcal{A}^{\text{CP}}$ world averages [29] for this mode and the $B^+ \rightarrow K^+\pi^0$ decay mode are shown in Table 2.2. The LHCb measurement is included in the world average of $\mathcal{A}^{\text{CP}}(B^0 \rightarrow K^+\pi^-)$.

Table 2.1: Measurements of $\mathcal{A}^{\text{CP}}$ for the $B \rightarrow K\pi$ family of decays

	BaBar	Belle
$B^0 \rightarrow K^0\pi^0$	$+0.13 \pm 0.13 \pm 0.03$ [30]	$+0.14 \pm 0.13 \pm 0.06$ [31]
$B^+ \rightarrow K^0\pi^+$	$-0.029 \pm 0.039 \pm 0.010$ [32]	$-0.011 \pm 0.021 \pm 0.006$ [33]
$B^0 \rightarrow K^+\pi^-$	$-0.107 \pm 0.016^{+0.006}_{-0.004}$ [34]	$-0.069 \pm 0.014 \pm 0.007$ [33]
$B^+ \rightarrow K^+\pi^0$	$+0.030 \pm 0.039 \pm 0.010$ [35]	$+0.043 \pm 0.024 \pm 0.002$ [33]

Table 2.2: Measurements of $\mathcal{A}^{\text{CP}}$ and world averages for the $B \rightarrow K^+\pi$ decay modes

	$B^0 \rightarrow K^+\pi^-$	$B^+ \rightarrow K^+\pi^0$
LHCb	$-0.088 \pm 0.011 \pm 0.008$	n/a
World average	-0.082 ± 0.006	0.040 ± 0.021

The disparity of the $\mathcal{A}^{\text{CP}}$ measurements in these two modes is characterized by the quantity

$$\Delta\mathcal{A}^{\text{CP}}(K\pi) = \mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+\pi^0) - \mathcal{A}^{\text{CP}}(B^0 \rightarrow K^+\pi^-). \quad (2.16)$$

Using the values in Table 2.2, we have $\Delta\mathcal{A}^{\text{CP}}(K\pi) = 0.122 \pm 0.022$. That this value disagrees with zero by more than 5σ is referred to as the “ $K\pi$ puzzle”. Note that the error on $\Delta\mathcal{A}^{\text{CP}}(K\pi)$ is dominated by the error on the world average of $\mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+\pi^0)$. While this anomaly may have SM origins, it may also be an indication of physics beyond the SM. A more precise measurement of $\mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+\pi^0)$ is therefore an important step in advancing experimental investigations of the $K\pi$ puzzle.

Chapter 3: The LHCb experiment

3.1 The LHC

The Large Hadron Collider (LHC) at CERN is the largest particle collider ever constructed. It is primarily a proton-proton collider, operating at center-of-mass energies of 7 and 8 TeV during 2011 and 2012, and most recently 13 TeV during the 2015–2018 operating period. Specialized heavy-ion runs have also been performed for short periods, at center-of-mass energies of up to 5 TeV. These energies are the highest ever reached by a particle collider, both for hadron collisions—proton-proton and heavy-ion—and for electron-positron collisions.

The LHC is built in the circular tunnel originally constructed for the Large Electron-Positron Collider (LEP). The tunnel is 27 km in circumference, and positioned approximately 100 m underground. This facility was built underground primarily because this was much cheaper than the extensive land purchases that would have been necessary to build it on the surface [36]. The surrounding rock also provides good radiation shielding.

Bare protons are produced by using a magnetic field to ionize hydrogen gas. The protons are transferred through pre-existing CERN accelerators—the Proton Synchrotron (PS) and the Super Proton Synchrotron (SPS)—before being injected

into the LHC at 450 GeV. Finally, the LHC itself accelerates the beams up to their designed colliding energies. The beam orbits are controlled by 1,232 superconducting dipole magnets, and are made to intersect each other at specified intersection points (IPs) once the beams reach their peak energy. The various particle experiments stationed along the LHC ring are located at these IPs.

The LHCb experiment is one of the four large particle detector experiments at the LHC, together with the ATLAS, CMS, and ALICE experiments. The ALICE detector is designed to study quark-gluon plasma produced in heavy-ion collisions. The ATLAS and CMS detectors are general-purpose particle detectors (GPDs) with nearly full 4π steradian coverage. It is these experiments that were able to monumentally discover the Higgs boson in 2012.

3.2 The LHCb detector

A succinct description of the LHCb detector is that it is a single-arm forward spectrometer covering the pseudorapidity range $2 < \eta < 5$. What is meant is that the LHCb detector is roughly a cone shape centered closely around the LHC beam-pipe with the IP at the apex, and can discriminate the types of particles that pass through it with excellent precision.

Pairs of b and $\bar{b}$ quarks are produced dominantly along the beam-line in proton-proton collisions at the LHC, as shown in Fig. 3.1. The LHCb acceptance is designed to exploit this, albeit on only one side of the IP. The LHCb detector is thus often referred to as a forward detector, and it is very complementary to the hermetic




Figure 3.1: Production rate of $b\bar{b}$ pairs at $\sqrt{s} = 7$ TeV as a function of decay angle. The rate normalization is arbitrary. The LHCb acceptance is indicated in red.

CMS and ATLAS detectors. This can be seen in Fig. 3.2, which shows the $b\bar{b}$ production rate at the LHC as a function of pseudorapidity, as well as the acceptance of the LHCb detector and that of a typical GPD. For comparison, a to-scale cross-sectional view of the LHCb and CMS detectors is shown in Fig. 3.3. The LHCb IP is superimposed over that of the CMS detector.




Figure 3.2: Production rate of $b\bar{b}$ pairs at $\sqrt{s} = 7$ TeV as a function of pseudorapidity. The rate normalization is arbitrary. The LHCb acceptance is indicated by the red box in the forward region, and that of a typical GPD is indicated by the yellow box.




Figure 3.3: Cross-sectional view of the CMS and LHCb detectors. The images of the detectors are to-scale, and the LHCb IP is placed over that of the CMS detector.

The LHCb acceptance is a pyramid shape, as this is far more practical to instrument than a conical shape, and the acceptance pyramid is rectangular in accord with the bending plane of the magnet. A cross-sectional view identifying different sub-detectors is shown in Fig. 3.4, and an overall perspective view of the LHCb detector is shown in Fig. 3.5. The z axis of the right-handed LHCb coordinate system points down the beam-pipe from the VELO, with the y axis pointing up, leaving the x axis to point left when looking downstream.

A dipole magnet separates charged particles according to their charge and momentum, a tracking system provides precise geometric information and momentum measurements of charged particles, and a particle identification (PID) system distinguishes between different types of charged and neutral particles as well as providing energy measurements for neutral particles.




Figure 3.4: Cross-sectional view of the LHCb detector. The z axis of the LHCb coordinate system points to the right.




Figure 3.5: Perspective view of the LHCb detector. Obvious features are the magnet in blue, the RICH 2 detector in yellow, and the muon system in green.

3.2.1 Magnet

Momentum measurements of charged particles is made possible by the dipole magnet, shown in perspective in Fig. 3.6. The two coils of the magnet are saddle-shaped, and they are each mounted at an angle such that the assembly of the coils makes a roughly conical shape to match the LHCb acceptance. The coil assembly is framed by a 1500 ton iron yoke to both contain and shape the magnetic field induced by the coils. A mapping of the magnetic field along z is shown in Fig. 3.7. The magnetic field has a bending power of about 4 Tm.

During operation, the magnetic field tends to push positively charged particles to one half of the detector, and negatively charged particles to the other half. To




Figure 3.6: Perspective view of the magnet from downstream; the VELO is behind the magnet in this view.

reduce experimental bias that could be introduced by this condition, the polarity of the magnet is inverted periodically between data-taking runs, the goal being to collect equally sized “magnet-up” and “magnet-down” data sets.




Figure 3.7: The strength of the magnetic field along the beam-line as a function of z .

3.2.2 Tracking

Three major sub-detectors make up the LHCb tracking system: the Vertex Locator (VELO), the Tracker Turicensis (TT), and the downstream tracking stations collectively referred to as the T-stations.

The VELO is a compact silicon-strip detector situated around the LHCb IP. It is designed to reconstruct primary proton-proton interaction vertices (PVs), secondary decay vertices of weakly decaying particles such as b and c hadrons, and to precisely resolve these secondary vertices and the PVs they originate from.

The VELO is made up of 23 silicon-strip stations positioned along the beam-line over a distance of about 1 m, as shown in Fig. 3.8. The stations are positioned such that any particle within the LHCb acceptance will pass through at least three of them. Each station is made of a pair of left- and right-half modules, and each




Figure 3.8: Schematics of the VELO sensors. Top: top-view showing the relative positions of the sensors along the beam-line. Bottom: Views in the x - y plane showing the two halves of the VELO closed (left), and open (right). Note the overlapping profiles of the sensor halves in the closed position.

module is made of two back-to-back silicon strip sensor assemblies, one with strips oriented to measure the radial coordinate r , and the other with strips oriented to measure the angular coordinate ϕ . The two halves are staggered along the z axis to allow them to overlap in the x - y plane, providing full acceptance coverage. The pitch of the silicon strips ranges from 38 to 102 μm , depending on the distance of the sensor from the beam-line. The raw hit resolution ranges from 10 to 28 μm , depending on the pitch of the sensor. During the injection of protons into the LHC, the beams have a relatively broad aperture, and are only later focused into a tighter geometric cross section for collisions. To avoid excessive radiation damage during injection, the two halves of the VELO each retract 3 cm away from the beam-line

as indicated in Fig. 3.8. The VELO halves can also make fine translations along the y axis, allowing the two halves to be moved to the actual position of the stable beams, which can vary between fills by ± 5 mm in both x and y . This articulation allows the sensors to be positioned with a precision of $10\text{ }\mu\text{m}$. In order to reduce the material that particles traverse before reaching the sensors, the VELO halves are separated from the beam-pipe vacuum by aluminum structures 0.5 mm thick. The VELO sensors are also in vacuum to avoid deforming these relatively delicate structures.

The TT is located downstream from the VELO, immediately before the magnet. It consists of four planes of silicon-strip sensors with a strip-pitch of $183\text{ }\mu\text{m}$, and a hit resolution of about $50\text{ }\mu\text{m}$ is ultimately achieved. The planes are roughly 150 by 130 cm , and are arranged in an x - u - v - x configuration, meaning that strips of the x planes are aligned vertically, while those of the u and v planes are rotated by $\pm 5^\circ$ about the z axis, respectively. This configuration is shown in Fig. 3.9, and the TT along with the downstream T-stations is shown in Fig. 3.10. One might intuit that the u and v planes ought to be rotated by 90° relative to the x planes, but the near-vertical alignment of $\pm 5^\circ$ is dramatically simpler to read out given that read-out hardware will already be present for the vertical x planes. This arrangement also reduces the number of ambiguities in the positions of particles when multiple particles leave signals in the same sensor at the same time. This in turn reduces the number of fake (ghost) tracks created during reconstruction.

Immediately downstream of the magnet are the T-stations, also shown in Fig. 3.10. They are three constructions of a common design, spaced about 70 cm




Figure 3.9: Diagram of the planes of the TT, looking downstream. The rotated configuration of the middle u and v planes is shown.




Figure 3.10: Perspective view of the TT and the T-stations. The TT is shown in purple on the left, the IT and OT stations are shown on the right in purple and teal, respectively.

apart in the z direction. A T-station itself is made up of two sub-detectors: the the Inner Tracker (IT), made of silicon-strip sensors, and the Outer Tracker (OT), made of straw drift-tubes. The IT covers a cross-shaped region of 120 by 40 cm, while the OT extends to the periphery of the LHCb acceptance and measures roughly 6 by 5 m. Charged particle densities for minimum-bias events are nearly an order of magnitude higher near the beam-pipe than in the outer regions of the detector, motivating the presence of the IT to complement the OT.

The IT stations themselves are furthermore composed of four assemblies arranged around the beam-pipe as shown in Fig. 3.11. Each of the IT assemblies, as well as the OT stations, are made up of four planes in the same x - u - v - x configuration as the TT. The strip-pitch of the IT sensors are $198\,\mu\text{m}$, and as with the TT, a spatial resolution of about $50\,\mu\text{m}$ is achieved.




Figure 3.11: Perspective view of the configuration of an IT station. The z axis points along the pipe towards the upper right.

The OT detector planes are each made of two layers of tubes. The layers are staggered in the x - z plane, shown in Fig. 3.12, in order to provide full acceptance coverage in x . A 7:3 ratio of Argon and CO_2 fills the tubes, which provides a drift-coordinate resolution of 200 μm .

The tracking system, together with the magnet, results in a momentum resolution of $\delta p/p = 0.35\%$ to 0.55% , depending on the momentum of the particle. In terms of transverse momentum (p_T), the impact parameter (IP) resolution is $\sigma_{\text{IP}} = 14 \mu\text{m} + 35 \mu\text{m}/p_T$ with p_T in GeV. These are plotted in Fig. 3.13. A Monte Carlo (MC) simulated data set of $B^0 \rightarrow J/\psi K_S^0$ was used to obtain these results. Fig. 3.13 also shows the p and p_T spectra of the B^0 mesons for reference. A track is considered reconstructible if it creates a minimum number of hits in the tracking detectors. If a minimum of 70% of hits belonging to a track originate from a single MC particle, then this track is considered to be successfully reconstructed. Given these definitions, the efficiency to reconstruct a charged particle traversing the entire tracking system is 94%.




Figure 3.12: Top views of one plane of the OT, showing the staggered two-layer configuration of the straw-tube modules.




Figure 3.13: Plots demonstrating the performance of the tracking system. Left: momentum resolution $\delta p/p$ as a function of p . Right: IP resolution as a function of $1/p_T$. The typical distributions of B mesons at LHCb are shown below the plots.

Pattern-recognition algorithms create tracks by searching among hits in the tracking detectors, producing different types of tracks depending on which tracking detectors are involved. These track types are diagrammed in Fig. 3.14, which shows a view in the bending plane of the magnet. Tracks from particles that traverse only the VELO and the TT before being bent out of the detector acceptance by the magnet are called Upstream tracks. Similarly, particles that traverse the TT and the T-stations without leaving hits in the VELO can produce Downstream tracks. Long tracks are produced by particles passing through the entire tracking system, and it is these particles for which the LHCb detector achieves its best momentum-resolution performance. Indeed, most physics analyses are performed using only Long tracks, while Downstream tracks are used to reconstruct candidate K_S^0 and Λ particles. Upstream tracks are primarily used to study backgrounds in the PID




Figure 3.14: An illustration showing different types of tracks superimposed over the LHCb tracking system. The view is in the x - z plane—the bending plane of the magnet.

algorithms of the RICH 1 detector. Finally, tracks that leave hits only in the VELO or the T-stations are called VELO tracks and T-tracks, respectively. VELO tracks contribute to the reconstruction of PVs, while T-tracks are used in the RICH 2 PID algorithm. After tracks have been created, they are improved using a Kalman-filter-based track fit where additional tracking information is added in an attempt to account for multiple scattering and energy loss through interactions with the detector material [37].

3.2.3 Particle identification

The LHCb PID system is comprised of two ring-imaging Cherenkov detectors (RICH 1 and 2), an electromagnetic (ECAL) and a hadronic (HCAL) calorimeter, and a muon system. Different types of particles interact with the tracking and PID

systems in different ways—if at all, which creates an overall signature for particle types. This is illustrated in Fig. 3.15, where, for example, the neutral photon bypasses the tracking system and is stopped by the ECAL, while charged hadrons leave tracking hits, interact with the ECAL, and pass through to deposit much more energy in HCAL.




Figure 3.15: Schematic of particle signatures in a typical configuration of a modern particle detector experiment [38].

The ECAL is the only LHCb sub-detector sensitive to photons from hadronic decays. Thus, in order for the LHCb experiment to have an acceptable energy resolution for reconstructed photons, the ECAL must contain the entire electromagnetic shower generated by them. The thickness of the ECAL is therefore chosen to be 25 radiation lengths (X_0). The hit density at LHCb varies by nearly two orders of magnitude between the region near the beam-line to the periphery of the detector acceptance. To account for this, the ECAL (as well as the other calorimetry instruments) is segmented laterally into sections of varying cell size, shown in Fig. 3.16.




Figure 3.16: Segmentations of the calorimetry sub-detectors. Left: segmentation of the SPD, PS, and ECAL. Right: segmentation of the HCAL.

The ECAL is a shashlik-type calorimeter, wherein alternating tiles of scintillator and absorber are stacked and penetrated longitudinally by optical fibers, not unlike a shish kebab. Polystyrene scintillator and lead absorber materials are used to construct detector modules. Photomultiplier tubes (PMTs) read out the optical fibers which traverse the length of a module twice by looping out of and back into the front of a module. An energy resolution of $\sigma_E/E = (8.5\text{--}9.5\%)/\sqrt{E} \oplus 0.8\%$ is achieved, where E is in units of GeV. The resolution depends on the type of module (inner, middle, or outer). This corresponds, for example, to a B mass resolution of 65 MeV for the $B^0 \rightarrow K^{*0}\gamma$ decay, and 75 MeV for the $B^0 \rightarrow \rho\pi$ decays.

In contrast with the ECAL, it is not critical for the HCAL to produce precise energy measurements because the tracking system already provides high-precision momentum measurements. Performance requirements for the HCAL are in fact primarily dictated by its role in the Level-0 (L0) hardware trigger, discussed in the next section. Simulation studies have shown that a thickness of 5.6 interaction lengths (λ_I) is sufficient without taking up excessive space in the confines of the LHCb cavern. Due to the larger transverse size of hadronic showers, the HCAL is

separated into only two sections of varying cell size, shown previously in Fig. 3.16. The HCAL modules are made of tiles with faces aligned parallel to the beam-line, illustrated in Fig. 3.17. Steel is used as an absorbing material. Optical fibers are routed longitudinally along the sides of the modules and are read out by PMTs. Each fiber is in optical contact with three scintillators. The energy resolution of the modules has been measured in beam tests to be $\sigma_E/E = (65 \pm 5)\%/\sqrt{E} \oplus (9 \pm 2)\%$ in units of GeV.




Figure 3.17: Diagram of an HCAL module, showing the arrangement of the scintillating and absorbing tiles.

The PID performance of the ECAL is augmented by a preshower detector (PS) and a scintillating pad detector (SPD). This is motivated by the L0 electron trigger, which only uses calorimetry information due to the fast timing requirements of the L0 trigger system. The thickness of the ECAL represents $1.2 \lambda_I$, hence the overall

ECAL response is not limited to only electrons and photons, providing a need to identify charged hadronic backgrounds. The L0 electron trigger must furthermore discriminate between electrons and neutral pions. These requirements are met by the principle of reading out individually the first scintillator layer and the first absorber layer of the basic sampling calorimeter design. The SPD is the fully instrumented scintillator layer, and PS is the fully instrumented absorber layer, which is in turn made of a layer of lead followed by a scintillator layer for read-out. The SPD/PS system is laterally segmented identically to the ECAL segmentation. Both the SPD and PS are read out by PMTs mounted along the top and bottom of the detectors.

The Muon system consists of five stations (M1–M5). Muons lose very little energy as they traverse the LHCb detector, and the muon sub-detector system is therefore placed the furthest downstream of all the sub-detectors, behind the calorimeters. The exception is the M1 station, which is placed in front of the calorimeters as shown in Fig. 3.18. This placement improves the fast p_T measurement that the muon system provides for the L0 triggers. Lead filters are placed between M2–M5 to mitigate any penetrating hadronic background. 99% of the muon system is instrumented with multi-wire proportional chambers (MWPC). In the inner region of M1, the rate of particles is so high that MWPCs would age beyond acceptable performance over the lifetime of the experiment. This region is therefore instrumented with triple-GEM (gas electron multiplier) detectors, which are more resistant to aging. Both detector types use a gas mixture of Ar/CO₂/CF₄ in a ratio of (40:55:5) for the MWPCs, and (45:15:40) for the triple-GEMs. The primary role of the muon system is to provide efficient muon identification. In particular, the muon system alone provides a




Figure 3.18: Side view of the muon system. Note that the M1 station is positioned in front of the calorimetry sub-detectors.

binary identification decision, named IsMuon. Data-driven methods have been used to evaluate the muon identification performance, and for tracks with p_T greater than 0.8 GeV, IsMuon correctly identifies muons with 97–99% efficiency [39]. The same studies show that hadrons are misidentified as muons by IsMuon at a rate of 1%.

It is critical to the LHCb physics program to distinguish between charged pions, kaons, electrons, and protons. These leave the same signature in the scheme shown previously in Fig. 3.15, so specialized detectors are needed to distinguish them by mass. The RICH detectors induce charged particles to radiate Cherenkov light, the pattern of which depends on the velocity of the particle. This, together with

a momentum measurement from the tracking system provides a mass measurement for the particle. The RICH 1 detector is positioned between the VELO and the TT. It has full coverage of the LHCb acceptance, and it is optimized for the particle momentum range of $\sim 1\text{--}60$ GeV. Both aerogel and C_4F_{10} gas are used as radiators. The RICH 2 detector is optimized for the momentum range of $15\text{--}100$ GeV and higher, and is positioned behind the T-stations, downstream of the magnet. After the magnet, high-momentum particles are still concentrated near the beam-line. For this reason, the RICH acceptance is ± 15 mrad to ± 120 mrad horizontally and ± 100 mrad vertically, rather than the full peripheral LHCb acceptance of ± 300 mrad horizontally and ± 250 mrad vertically. CF_4 gas is used as a radiator in RICH 2.

The fitness of this arrangement of radiators and active areas is demonstrated in Fig. 3.19, which plots the Cherenkov angle θ_C vs. particle momentum for the RICH radiators. The Cherenkov photons are collected by arrays of pixel hybrid photon detectors (pixel-HPDs) developed by LHCb scientists in close collaboration with industry [40]. In the pixel-HPD, a photon induces a photoelectron, which is accelerated onto a silicon detector composed of a square matrix of 1024 pixels, each $500 \times 500 \mu\text{m}$ in size. In order to reduce the material budget, the HPDs are placed outside of the LHCb acceptance, and the Cherenkov light is directed onto them with spherical and planar mirrors. The performance of the RICH detector configuration was evaluated in beam tests, which determined the θ_C resolution to be 1.66 ± 0.03 mrad, in agreement with the simulation expectation of 1.64 mrad.

The information from all of the PID sub-detectors is combined into global PID variables using two different techniques. The first is the more familiar delta-




Figure 3.19: Plot of the Cherenkov angle θ_C vs. momentum for the RICH radiators.

log-likelihood (DLL) method where the variable $\Delta\log\mathcal{L}_{\text{comb}}(X - \pi)$ is formed by simply adding the likelihood information from each sub-detector linearly. Here, X represents an electron, muon, kaon, or proton mass hypothesis. The second method is a more sophisticated artificial-neural-net (ANN) approach, where correlations between the different sub-detectors are taken into account. This approach produces a probability variable for each particle hypothesis, ProbNNX, where X again stands for the different particle types in question. As an example of the DLL performance, the tag-and-probe method was used to extract efficiencies in the $B^+ \rightarrow J/\psi(e^-e^+)K^+$ decay. At 90% efficiency for positively identifying electrons, the pion misidentification rate is 0.6% [41]. It is known that the ProbNNX performance is overestimated in simulation. Nevertheless, Fig. 3.20 provides well-approximated examples of the performance of ProbNN μ and ProbNN p , and comparisons of the performance relative to the corresponding DLL variables. These results were derived from data sidebands and signal MC in the study of the hypothetical $\Sigma^+ \rightarrow p\mu^+\mu^-$ decay.




Figure 3.20: Background rejection vs. signal efficiency for four of the global PID variables, derived from data sidebands and signal MC in the study of the hypothetical $\Sigma^+ \rightarrow p\mu^+\mu^-$ decay. DLL variables are drawn as black circles, and ProbNN variables are drawn as red stars.

3.2.4 Trigger

As with many modern particle detectors, it is overwhelmingly impractical to record the total output of the LHCb detector during physics runs of the LHC, due both to storage and processing bandwidth limitations. Furthermore, the b -hadron decay modes of prime importance in the LHCb physics program typically have branching fractions on the order of 10^{-3} to 10^{-5} or smaller. A trigger system is therefore necessary to quickly assess whether the output of an event should be saved indefinitely for further analysis, while also greatly enhancing the purity of b -hadron decays in the recorded data sets.

The LHCb trigger system is split into a hardware stage and a software stage. The hardware stage is called L0, as mentioned in the previous section, and the software stage is referred to as the high-level trigger (HLT). The L0 hardware trigger uses information from the calorimetry sub-detectors and the muon system to reduce the 40 MHz bunch-crossing rate of the LHC to a 1 MHz output. The hardware of

the L0 system, including the sub-detectors taken as input, operates synchronously with the LHC bunch-crossing rate, making trigger decisions in at most $4\,\mu\text{s}$. The calorimetry system calculates the transverse energy E_T in clusters of 2×2 cells, and makes decisions based on different thresholds for different particle types. For example, the `L0Photon` trigger checks the highest E_T of ECAL clusters that have hits in the PS, but not in the corresponding SPD cells; a signature consistent with a photon. If such a cluster has $E_T > 2.5\,\text{GeV}$, the event is accepted by the trigger. For comparison, the threshold for the similar HCAL-based `L0Hadron` trigger is $3.5\,\text{GeV}$. Details on the other L0 trigger lines and their thresholds can be found in Ref. [42].

The software HLT runs on a farm of PCs called the event filter farm (EFF), located in the Point-8 surface facilities directly above the LHCb cavern. The entire LHCb detector can be read out at the L0 output rate, allowing the HLT to perform more sophisticated reconstruction and selection algorithms. The HLT is furthermore divided into two stages, HLT1 and HLT2. A partial reconstruction is performed for the HLT1, which then uses this information to reduce the L0 rate to $\sim 40\,\text{kHz}$. Trigger lines in the HLT1 are primarily tracking-based, with p_T thresholds dependent on the L0 trigger that selected the event. For example, the threshold for tracks with the highest p_T is lower for events that pass the `L0Photon` trigger than for events that pass the `L0Hadron` trigger, a concession afforded by the `L0Photon` rate being lower than the `L0Hadron` rate.

A more thorough tracking reconstruction is performed for the HLT2, with a tracking performance nearly as good as the offline reconstruction. In addition to fast muon identification, tracks are also matched with ECAL clusters for electron

identification. The primary output bandwidth of the HLT2 is allocated to inclusive “topological” trigger selections. These triggers reconstruct secondary vertices using different numbers of tracks, and use information such as the vertex fit quality and the displacement of the secondary vertices from the PVs in an attempt to identify partially-reconstructed b hadron decays. A specialized implementation of a boosted decision tree (BDT) multivariate classifier [43, 44] is also used to improve the efficiency of the topological trigger selections [45].

There are also exclusive HLT2 triggers that are designed to enhance the efficiency of a single decay mode, such as the $B_s^0 \rightarrow K^+ K^-$ decay. The trigger for this decay mode is designed without using information correlated with the B_s^0 lifetime in order to avoid biasing lifetime measurements in off-line analyses. This implies a higher combinatorial background rate, which is compensated by tight requirements on allowed values of the reconstructed B_s^0 mass. This is possible because the signal is fully reconstructed, in contrast with the topological triggers.

Trigger efficiencies vary substantially depending on the decay mode in question. Typical efficiencies for the L0Hadron trigger are $\sim 40\%$, while the efficiency of the L0Photon trigger can be as high as 80% for a particularly ideal decay mode such as $B^+ \rightarrow K^+ \pi^0$. The HLT performs well in general, with typical efficiencies from 60—90% for the HLT1 triggers, and 40—90% for the topological HLT2 triggers.

3.2.5 Running conditions

The LHCb experiment is designed for an instantaneous luminosity nearly two orders of magnitude smaller than that of the ATLAS and CMS experiments, and the ALICE experiment is designed for an instantaneous luminosity a further two orders of magnitude smaller than that of LHCb. The LHC provides these hugely different conditions to the experiments simultaneously by using a “luminosity-leveling” technique: the intersection of the profiles of the proton beams in the plane transverse to the beam-line is continually adjusted to maintain a constant luminosity and constant number of visible proton-proton interactions—also called pile-up, at the ALICE and LHCb IPs. As the proton beams are depleted through collisions during data-taking periods, the intersections of the beams at the LHCb and ALICE IPs are increased. Beyond providing stable and optimized conditions for data-taking, this mechanism is critical to avoid damaging the active elements of the detectors.

The conditions provided by the LHC at the LHCb IP during Run 1 are given in Table 3.1, along with the integrated luminosity recorded by the LHCb detector. The center-of-mass energy is denoted by $\sqrt{s}$, luminosity by $\mathcal{L}$, and integrated luminosity by $\int \mathcal{L}$.

Table 3.1: Data-taking conditions at the LHCb IP during Run 1 of the LHC, and the integrated luminosity recorded by the LHCb detector.

Year	$\sqrt{s}$ (TeV)	$\mathcal{L}$ ($10^{32}\text{cm}^{-2}\text{s}^{-1}$)	Pile-up	$\int \mathcal{L}$ (fb^{-1})
2011	7	3.5	1.4	1.1
2012	8	4.0	1.8	2.1

3.2.6 Computing

Large quantities of high-quality simulated data are critical for the understanding and calibration of the LHCb detector, and for the extensive offline analysis of the data collected by it. The creation of simulated data starts with the Gauss software package [46], using PYTHIA [47, 48] to generate proton-proton collision events with a specific LHCb configuration [49], and EVTGEN [50] together with PHOTOS [51] to simulate hadronic particle decays and final-state radiation. A description of the detector and the magnetic field is used with GEANT4 [52, 53] to simulate the interaction of particles with the detector material. The BOOLE software package is then used to digitize the fully simulated data, resulting in the identical data format as that produced by the LHCb detector. Simulation of trigger decisions is performed by processing this data with the MOORE software package, which is the same software used to run the HLT during data collection. Past this point in the data processing pipeline, real and simulated data are treated identically.

The BRUNEL software package performs the full offline event reconstruction, creating tracks and generating PID information for them, as well as creating neutral candidate particles and corresponding PID information. Finally, the DAVINCI software package is used by analysts to reconstruct candidate particle decays, calculate additional variables, and to define preliminary event selections for physics analyses. The LHCb-specific software packages listed here are described in detail in Ref. [54].

A global network of computing centers called the LHC Computing Grid (LCG) provides distributed computing resources for physics analyses, as well as high-

capacity distributed storage, which in turn allows for enhanced data integrity through the replication of data collected by the LHC experiments to several remote sites [55]. Analysts in the ATLAS and LHCb collaborations interface with the computing resources of the Grid and the data stored on it using the Gaudi/Athena and Grid Alliance (GANGA) front-end software package [56].

Chapter 4: Study of the decay $B^+ \rightarrow K^+\pi^0$

4.1 Introduction

Rare decays of heavy-flavored hadrons that primarily proceed through penguin (loop) diagrams are among the most powerful probes for the effects of new physics beyond the SM. The family of $B \rightarrow K\pi$ decays consists of the B^0 meson decays to the $K^+\pi^-$ and $K^0\pi^0$ final states, and the B^+ meson decays to the $K^+\pi^0$ and $K^0\pi^+$ final states.¹ In the SM, these decays are dominated by penguin topologies, and are therefore sensitive to new physics contributing inside the loop. Measurements of the $B \rightarrow K\pi$ decays at the B factories, CDF, and LHCb have revealed significant deviations from the expected pattern of CP violating asymmetries in these channels—the “ $K\pi$ puzzle” described in 2.6. While the precision of the world average of the $\mathcal{A}^{\text{CP}}(B^0 \rightarrow K^+\pi^-)$ measurements is dominated by the LHCb contribution, the resolution of the puzzle requires a comprehensive study of the entire decay family, including improved measurements of final states containing π^0 mesons.

Analyses involving the $H_b \rightarrow h\pi^0$ decay topology face substantial challenges in the LHCb environment. Here, H_b denotes a weakly-decaying hadron containing a b quark, and h denotes a lighter hadron, such as a K meson or a Λ baryon. These

¹The inclusion of charge-conjugated decays is implied throughout.

decays are reconstructed with low efficiency due to the presence of the π^0 meson. Furthermore, this decay topology lacks a secondary vertex. Events are said to be triggered on-signal (TOS) with respect to a particular trigger when the constituents of a reconstructed signal candidate are sufficient to satisfy the requirements of that trigger. Likewise, events are said to be triggered independent-of-signal (TIS) with respect to a trigger when the rest of the event satisfies that trigger’s requirements. Every HLT2 trigger algorithm used during Run 1 data-taking requires a reconstructed secondary vertex, hence signal events in this channel cannot be TOS with respect to any HLT2 trigger. Events containing an $H_b \rightarrow h\pi^0$ decay must therefore trigger based on the presence of an additional decay in the event that does have a reconstructable secondary vertex, and the trigger efficiency for these decay modes will consequentially be very low. There are also large numbers of π^0 candidates in each event, particularly for candidates reconstructed from two resolved photon-candidate clusters in the ECAL. Finally, the branching fractions of these decays are typically on the order of 10^{-5} . Altogether, these characteristics result in very low overall signal yields and extremely high combinatorial backgrounds.

The present study of the $B^+ \rightarrow K^+\pi^0$ decay led to an analysis strategy for extracting a signal from the current LHCb data, while developing techniques applicable to analyses of decays with similar topologies. The success of this analysis has resulted in the creation of a dedicated HLT2 trigger for each of the $B \rightarrow K\pi^0$ decays, discussed further in Sect. 4.6. The ongoing analysis of the Run 2 data, discussed in Sect. 4.8, has shown a significant signal in the $B^+ \rightarrow K^+\pi^0$ channel, confirming the effectiveness of the dedicated HLT2 trigger and the analysis technique.

The success of the present analysis has also resulted in the creation of dedicated HLT2 triggers for the as yet unobserved $\Lambda_b \rightarrow \Lambda\gamma$, $\Xi_b^- \rightarrow \Xi^-\gamma$, and $\Omega_b \rightarrow \Omega\gamma$ decays. These decays involve the $b \rightarrow s\gamma$ transition, with the photon being produced dominantly in the left-handed polarization state in the SM. New physics models can enhance the right-handed polarization component, making these decays an important area in which to search for new physics effects. Angular analyses of the $\Lambda_b \rightarrow \Lambda\gamma$ decay in particular can serve as an effective probe for new physics [57].

A brief outline of the analysis strategy is as follows. Candidate B^+ mesons are formed by adding the four-momentum of a K^+ candidate and a π^0 candidate. A vertex isolation procedure is then performed wherein each track in the event is combined with the K^+ candidate, one at a time, in an attempt to form good two-track secondary vertices. Vertex isolation variables calculated with this procedure are used to suppress events consistent with an H_b decay involving a secondary vertex. A set of track isolation variables independent of vertex isolation is also used to select candidates with a K^+ that is isolated from other tracks in the event. These isolation variables help to reduce partially reconstructed backgrounds, as well as the combinatorial background since they tend to suppress events with many tracks. To exploit the significant flight distance of the B^+ meson in the absence of a secondary vertex, a variable named the mother-trajectory distance-of-closest-approach (MT-DOCA) is used. The MT-DOCA is explained in detail in Sect. 4.3. A multivariate technique is employed to obtain the optimal selection criteria. Finally, a fit to the invariant mass distribution of the reconstructed B^+ candidates that pass the final selection criteria is performed to estimate the signal yield.

4.2 Data

The present analysis examines the full data set recorded by the LHCb experiment during Run 1 of the LHC. This corresponds to approximately 3 fb^{-1} of integrated luminosity, recorded at center-of-mass energies of 7 TeV and 8 TeV in 2011 and 2012, respectively.

Dedicated MC simulated $B^+ \rightarrow K^+\pi^0$ data sets have been produced for this analysis. Additional simulated data sets previously produced for other analyses are used for background studies, and for the simulation corrections described in Sect. 4.3. The size of the simulated data sets are summarized in Table 4.1.

Table 4.1: Numbers of simulated events generated at the center-of-mass energies of 7 TeV and 8 TeV, corresponding to the 2011 and 2012 data-taking conditions, respectively.

	Decay mode	Number of events	
		7 TeV	8 TeV
Signal mode	$B^+ \rightarrow K^+\pi^0$	10087420	10119943
Background studies	$B^0 \rightarrow K^{*0}(K^+\pi^-)\gamma$	4650262	6053226
For use with simulation corrections	$B^0 \rightarrow K^+\pi^-$	268999	531995
	$B^0 \rightarrow \pi^+\pi^-$	1527244	3067742
	$B_s^0 \rightarrow K^+K^-$	1532248	3052242
	$B_s^0 \rightarrow \pi^+K^-$	1514494	3071739

4.3 Event selection

4.3.1 Reconstruction techniques

Candidate π^0 mesons are reconstructed from ECAL clusters that are associated with neutral particles by the lack of corresponding activity in the SPD. Neutral

pion candidates reconstructed from two clusters are referred to as “resolved” π^0 candidates. At higher energies, the opening angle of the π^0 decay products shrinks, causing the two clusters to overlap with respect to the segmentation of the ECAL cells. These “merged” π^0 candidates are reconstructed using a dedicated algorithm which attempts to subdivide clusters using the two largest energy deposits in the original cluster. It is possible for some π^0 mesons to be reconstructed with either technique, but those with transverse energy above 3.1 GeV/c are predominantly reconstructed by the merged π^0 algorithm.

In each collision event, there is an abundance of low-energy ECAL clusters, mainly due to low-energy π^0 mesons produced in the primary proton-proton interaction. This ultimately leads to very large combinatorial backgrounds in reconstructed $B^+ \rightarrow K^+ \pi^0$ candidate decays. Candidate decays reconstructed with a resolved π^0 candidate suffer from an extremely large combinatorial background in particular, due both to the prompt π^0 mesons as well as random combinations of neutral ECAL clusters inherent to the resolved π^0 reconstruction technique. The combinatorial background for these candidates is in fact an order of magnitude larger than that of candidate $B^+ \rightarrow K^+ \pi^0$ decays reconstructed with a merged π^0 candidate. Dedicated selections have been optimized for B^+ candidates reconstructed with each π^0 reconstruction technique in order to account for the different background levels.

Plots of the mass distributions of reconstructed B^+ mesons from simulated $B^+ \rightarrow K^+ \pi^0$ decays are shown in Fig. 4.1. These plots are drawn on the same scale for comparison. The mass distribution of B^+ mesons reconstructed with the merged π^0 algorithm exhibits a long, low-mass tail due to energy leakage in the ECAL for




Figure 4.1: Mass distributions of B^+ mesons from simulated $B^+ \rightarrow K^+\pi^0$ decays, reconstructed using the merged (a) and resolved (b) π^0 algorithms. The dashed red lines in (b) indicate the range of reconstructed B^+ mass values ultimately allowed in the even selection.

π^0 mesons that have very high energy. When π^0 mesons are reconstructed from resolved ECAL clusters that actually have a slight overlap, the energy recorded in the overlapping cells is double-counted, resulting in a small high-mass tail of the reconstructed B^+ mesons, which can be seen in Fig. 4.1b. Fits to these mass distributions using Crystal Ball functions [58] give mass resolution estimates of $\sigma = 95 \pm 4 \text{ MeV}/c^2$ and $\sigma = 97 \pm 1 \text{ MeV}/c^2$ for B^+ mesons reconstructed with merged and resolved π^0 candidates, respectively. These resolutions are dominated by the energy resolution of the ECAL. For comparison, the mass resolution of reconstructed $B^0 \rightarrow K^+\pi^-$ decays is typically about $22 \text{ MeV}/c^2$.

Plots of the mass distributions of candidate B^+ mesons reconstructed from experimental data are shown in Fig. 4.2. These plots represent about only 10% of the total data set. For the data sets shown here, about 260 signal $B^+ \rightarrow K^+\pi^0$ decays are expected among the candidates reconstructed using the merged π^0 algorithm, and about 90 are expected among those reconstructed using the resolved π^0




Figure 4.2: Mass distributions of B^+ candidates from a sample of experimental data, reconstructed using the merged (a) and resolved (b) π^0 algorithms.

algorithm. Note that unlike the plots from simulation shown previously in Fig. 4.1, the mass range of the plots shown here are not on the same scale. The range of B^+ candidate mass values allowed in the selection that uses resolved π^0 candidates has to be made very narrow in order to avoid creating an excessive number of B^+ candidates: $4.8 < m < 5.8 \text{ GeV}/c^2$, where m is the reconstructed mass. This range is indicated with dashed red lines in Fig. 4.1b, and it is the same range shown in Fig. 4.2b. Given the resolution of these candidates, the narrow mass window ultimately makes it unfeasible to estimate the background levels and to distinguish the B^+ mesons from the background shape in a fit to the B^+ candidate mass spectrum. Due to this major constraint, and the enormous combinatorial background that remains, the B^+ candidates ultimately used in this analysis are reconstructed using only merged π^0 candidates.

4.3.2 Basic selection requirements

Candidate π^0 mesons are required to have p_T greater than 2.6 GeV/c, p greater than 10 GeV/c, and a reconstructed mass within ± 60 MeV/c² of the nominal π^0 mass. Plots of the signal and background distributions of these variables are shown in Fig. 4.3. In these plots, and in the similar plots that follow, the signal distributions come from simulated $B^+ \rightarrow K^+ \pi^0$ decays, and the background distributions are taken from recorded data with a reconstructed candidate B^+ mass greater than 5.8 GeV/c², also referred to as the mass sideband.




Figure 4.3: Signal and background distributions for the mass (a), p (b), and p_T of the π^0 candidates (c). Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.

Reconstructed charged particles with p_T greater than 250 MeV/c are used as K^+ candidates. Decay products from H_b hadron decays have trajectories that in general do not point back to the primary vertex (PV). Candidate K^+ mesons are therefore required to have a large impact parameter (IP) with respect to the PV by requiring that the $\Delta\chi_{IP}^2$ of the candidate is greater than 49. The variable “ $\Delta\chi_{IP}^2$ ” is defined as the difference between the χ^2 of a PV fit with and without the K^+ candidate track. This is motivated by the fact that a displaced track not originating from the PV should significantly worsen the quality of the vertex fit when it is included. This technique has the advantage of incorporating the resolution information of both the PV and the track momentum. The signal and background distributions of these variables are shown in Fig. 4.4.




Figure 4.4: Signal and background distributions for the p_T (a) and $\ln(\Delta\chi_{IP}^2)$ (b) of the K^+ candidates. The $\Delta\chi_{IP}^2$ is transformed by the natural logarithm in order to broaden the sharply peaking distributions, making the discriminating power of the variable more apparent. Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.

The momentum vector of the π^0 candidate is calculated using the energy and position information from the ECAL, and by taking the LHCb interaction point to be the origin. The four-momentum of the K^+ and π^0 candidates are then added to form B^+ candidates. In addition to the reconstructed mass range of $4.00 < m < 6.05 \text{ GeV}/c^2$, each B^+ candidate is required to have p_T greater than $1.5 \text{ GeV}/c$. The signal and background distributions of these variables are shown in Fig. 4.5. The pseudorapidity η and p of the B^+ candidate are also shown in this figure, as they are used in the multivariate selection described later in Sect. 4.3.7.




Figure 4.5: Signal and background distributions for the mass (a), p (b), p_T (c), and η of the B^+ candidates (d). Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.

4.3.3 Particle identification

The K^+ candidates are distinguished from other charged particles using PID information from the RICH and calorimeter systems. The neural-net-based ProbNNk variable described in Sect. 3.2.3 is used for this purpose, given its superior performance over that of the more traditional DLL method described in the same section. In general there are far more charged pions than kaons in a given collision event, making the kaon identification an important criterion for reducing the combinatorial backgrounds composed with charged pions. This is especially important for the candidate selection of the dedicated trigger, discussed later in Sect. 4.6.

Neutral PID methods are also very important due to the large amount of activity in the ECAL in each collision event. First of all, a track-matching algorithm is used to reject clusters that can be significantly associated with charged particles. After this initial track veto, a number of discriminating techniques are used to calculate a DLL for merged π^0 identification. Overlapping ECAL clusters that are produced by π^0 mesons are in general transversely broader than those produced by single photons. The shower shapes in both the PS and ECAL are therefore used to help distinguish merged π^0 candidates from single photon candidates. In addition, the patterns of activity in the PS, ECAL, and HCAL are used to help discriminate π^0 candidates from neutral hadrons such as neutrons or K_L^0 mesons. These particles pass through the PS and ECAL yet can still induce a small shower in the ECAL that could be consistent with the lighter π^0 meson, if not for the correlation with large HCAL clusters, and the anti-correlation with PS activity. The DLL variable

created from these attributes is transformed into a so-called confidence level (CL) variable ranging between zero and one.

Candidate K^+ mesons are required to have a ProbNNk greater than 0.8, and π^0 candidates are required to have a CL greater than 0.9. Plots of the signal and background distributions of these variables are shown in Fig. 4.6.




Figure 4.6: Signal and background distributions for the ProbNNk of the K^+ candidates (a) and the CL of the π^0 candidates (b). Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.

4.3.4 Mother-trajectory DOCA

The lack of a secondary vertex in the $H_b \rightarrow h\pi^0$ decay topology is not a mere inconvenience with respect to the pre-existing LHCb trigger strategy; it results in a truly enormous combinatorial background in the proton-proton collision environment of the LHC. It was found that this background is ultimately overwhelming if the discriminating power of the H_b meson lifetime is not somehow exploited in the

event selection. This is accomplished in the $B^+ \rightarrow K^+\pi^0$ decay mode, for example, by making a trajectory from the PV closest to the daughter K^+ candidate along the direction of the reconstructed B^+ candidate momentum. The variable called MT-DOCA is defined as the DOCA between the K^+ candidate and this trajectory. An illustration of the geometry involved is shown in Fig. 4.7.




Figure 4.7: MT-DOCA and the reconstructed $B^+ \rightarrow K^+\pi^0$ decay topology.

For well-reconstructed signal events, the MT-DOCA is small, whereas the distribution of random combinations of final-state particles has a very long tail at large values. A χ^2 for the MT-DOCA is constructed using the covariance matrix of the reconstructed PV, and is shown in Fig. 4.8. In some cases, the K^+ meson from a signal decay is associated with the wrong PV. These signal events resemble the background, as can be see in Fig. 4.8, but the effect on the performance is small. For the selection in the present analysis, the $\chi^2_{\text{MT-DOCA}}$ is required to be less than 10 for B^+ candidates reconstructed from 2011 data, and less than 9 for those

reconstructed from 2012 data. The requirements for the two data-taking periods are optimized separately to allow for different background levels at different center-of-mass energies, though the difference in performance of the requirements is ultimately negligible.




Figure 4.8: Signal and background distributions of $\log_{10}(\chi^2_{\text{MT-DOCA}})$ for the K^+ candidates. Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.

4.3.5 Isolation techniques

The discriminating power provided by variables that characterize the isolation of a candidate relative to other tracks in the event is critical for reducing the overall background level. Vertex isolation variables are calculated by combining another track in the event with the K^+ candidate to form a two-track secondary vertex. This procedure is performed for all tracks in the event. The multiplicity of two-track secondary vertices with $\chi^2_{\text{vtx}} < 9$ is recorded as the variable $V_{\text{Mult.}}$, where χ^2_{vtx}

is the χ^2 of the vertex fit. Of these vertices, the one with the smallest χ_{vtx}^2 is built on to form a four-track vertex. The two additional tracks used are those that result in the smallest χ_{vtx}^2 after being added to the vertex. The variable $\min(\Delta\chi_{\text{vtx}}^2)$ is then given by the change in the χ_{vtx}^2 upon adding these two tracks.

The p_T asymmetry of the reconstructed B^+ candidate is defined as

$$\mathcal{A}^{p_T} = \frac{p_{TB} - p_{T\text{cone}}}{p_{TB} + p_{T\text{cone}}}, \quad (4.1)$$

where p_{TB} is the transverse momentum of the reconstructed B^+ signal candidate, and $p_{T\text{cone}}$ is the magnitude of the vector sum of the transverse momentum of the charged particles near the reconstructed B^+ candidate. The $\mathcal{A}^{p_T}$ is used to isolate the reconstructed candidate from nearby tracks. To determine whether a track is “near” the signal candidate, the quantity $\Delta R = \sqrt{(\Delta\phi)^2 + (\Delta\eta)^2}$ is required not to exceed 1.2, where $\Delta\phi$ is the difference between the azimuthal angle of the momentum of the reconstructed candidate and the track, and $\Delta\eta$ is the difference between their pseudorapidities. The cone size, $\Delta R = 1.2$, is optimized using simulated data for the signal, and a subset of experimental data for the background. This optimization is performed by evaluating the discriminating power $D_{\Delta R}$ at different values of ΔR . The discriminating power is defined as

$$D_{\Delta R} = \sum_{i=0}^n \frac{(S(i) - B(i))^2}{S(i) + B(i)}, \quad (4.2)$$

where S and B are normalized signal and background histograms for a given ΔR , with i representing the i -th bin of the n -bin histograms. The discriminating power thus defined is zero for identical signal and background distributions, and positive

for dissimilar distributions, with the maximally discriminating value of ΔR giving the largest relative value of $D_{\Delta R}$.

The signal and background distributions of the isolation variables discussed here are shown in Fig. 4.9. All reconstructed B^+ candidates are required to have $\mathcal{A}^{p_T}$ greater than -0.6 . Candidates reconstructed from 2011 data are required to have $V_{\text{Mult.}}$ less than 8, and those reconstructed from 2012 data are required to have $V_{\text{Mult.}}$ less than 6. The requirements for the two data-taking periods are optimized separately to allow for different performances at different center-of-mass energies.




Figure 4.9: Signal and background distributions for the $V_{\text{Mult.}}$ (a), $\ln(\min(\Delta\chi^2_{\text{vtx.}}))$ (b), and, $\mathcal{A}^{p_T}$ of the B^+ candidates (c). Signal distributions from simulation are drawn in blue, and background distributions from the mass sideband of the reconstructed B^+ candidates are drawn in red. The vertical axes are arbitrary, and the two samples are not drawn to-scale.

In particular, the multiplicity of tracks in each collision event is somewhat greater at 8 TeV than at 7 TeV. No explicit requirements are made for the $\min(\Delta\chi_{\text{vtx}}^2)$ variable, but this as well as the $\mathcal{A}^{p_T}$ and $V_{\text{Mult.}}$ variables are used as inputs for the multivariate selection described later in Sect. 4.3.7.

4.3.6 Summary of preliminary selection

The selection requirements described thus far are summarized in Table 4.2. Given the various distributions of the discriminating variables shown previously, it is apparent that many of these requirements are aggressive. The consequential sacrifice of signal selection efficiency is necessary in the presence of the enormous backgrounds in order to avoid retaining an excessive number of signal candidates in the computing stages leading up to the final selection procedure described in the following sections.

Table 4.2: Requirements for candidate $B^+ \rightarrow K^+\pi^0$ decays that define the preliminary selection for the 2011(2012) data-taking period, where the merged π^0 algorithm is used to reconstruct π^0 candidates.

Particle	Variable	Requirement
π^0	p	$> 10 \text{ GeV}/c$
	p_T	$> 2.6 \text{ GeV}/c$
	mass	$79.6 < m < 199.6 \text{ MeV}/c^2$
	CL	> 0.9
K^+	p_T	$> 250 \text{ MeV}/c$
	$\Delta\chi_{\text{IP}}^2$	> 49
	$\chi_{\text{MT-DOCA}}^2$	$< 9(10)$
	ProbNNk	> 0.8
B^+	p_T	$1.5 \text{ GeV}/c^2$
	mass	$4.00 < m < 6.05 \text{ GeV}/c^2$
	$\mathcal{A}^{p_T}$	> -0.6
	$V_{\text{Mult.}}$	$8(6)$

4.3.7 Multivariate selection

The preliminary selection greatly reduces the size of the overall background—albeit at the expense of signal selection efficiency—such that the number of candidate $B^+ \rightarrow K^+\pi^0$ decays created and saved to disk does not excessively tax the available computing resources. Yet the number of candidates created is more than two orders of magnitude greater than the number of $B^+ \rightarrow K^+\pi^0$ decays that are expected to be in the selected data set. In order to further reduce the background without excessively diminishing the signal, a multivariate analysis is developed using boosted decision tree (BDT) classifiers. The large gain in discriminating power created by these classifiers is critical to the overall selection strategy for the $H_b \rightarrow h\pi^0$ decay modes.

A BDT classifier is a machine-learning-based algorithm designed to solve classification problems. Data sets representing the signal and background are used to create (or “train”) a BDT classifier which can then be used to evaluate how signal- or background-like a candidate is. The central concept is that of a decision tree, illustrated in Fig. 4.10. A tree is created by subdividing the training data sequentially until subsets have been defined that contain a sufficient purity of signal or background events. To achieve this, a set of discriminating variables (x_i, x_j, x_k) is first defined. At each step, or node, in the subdivision sequence, the most discriminating variable is identified, as well as the criterion $(c1, c2, \dots, c4)$ that best segregates the signal and background distributions of that variable, at that node. The sequences made up of the decision nodes constitute the tree, and any candidate will have




Figure 4.10: A diagram of a decision tree [59]. The root node represents the initial data set, and the other nodes represent subsets. At each node, the most discriminating variable (x_i, x_j, x_k) is identified, as well as the criterion ($c_1, c_2, \dots, c_4$) that best segregates the signal (S) and background (B) distributions of that variable. Notice that variables can appear more than once and with different criterion (x_j here, with criterion c_1, c_2).

a single path down the tree and be unambiguously classified. Decision trees are easy to use because variables with low discriminating power are naturally ignored, and because they make efficient use of the available statistics; candidates are not removed during the training procedure, which would otherwise destroy information.

However, decision trees are very susceptible to statistical fluctuations in the distributions of the discriminating variables, which can result in a pitfall referred to as over-training. This is circumvented with a general method known as boosting, where the training data is statistically modified with weights (or “boosted”) during the creation of the classifier. In the context of decision trees, events that were

classified incorrectly by the tree are assigned a higher weight, and another tree is created using the weighted data set. This procedure is applied again and again to create many trees, each of which are assigned a score. Classifications are then made by taking the weighted average decision of the trees, where the score of a tree is used as its weight. While a decision tree gives a binary classification for each candidate, the creation of a boosted tree results in a classification variable with a continuous distribution reflecting the margin of classification of each candidate, *e.g.*, a candidate that a weighted 90% of trees classifies as signal is much more signal-like than one classified as such by only 55% of the trees.

Different classification techniques and different boosting methods have been tested with the TMVA package [59] in order to determine the ideal classifier for this analysis. BDTs using the AdaBoost [60] (Adaptive Boosting) algorithm have been found to perform the best. The parameters of the training procedure have been carefully optimized in order to maximize performance while avoiding over-training. The classifiers are trained and tested using experimental data to represent the background, and simulated $B^+ \rightarrow K^+\pi^0$ data that are corrected using event weights described below. These simulated events have undergone the same reconstruction and selection as the real data. The background data is taken from the $5.7 < m < 6.0 \text{ GeV}/c^2$ range of the reconstructed B^+ mass, which is well separated from the mass peak of the B^+ meson (refer back to Fig. 4.1).

To avoid introducing bias in the training procedure, a double classifier technique is employed. The background data set is randomly divided into two disjoint subsets that are used to train and test two independent BDT classifiers, referred to

as BDT1 and BDT2. When the entire experimental data set is processed to classify the candidates, one of the two classifiers is chosen at random for each candidate, to classify that candidate. The seeds used to generate random numbers are recorded for reproducibility. This technique has the advantage that the entire background data set is used to train and test the classifiers, enhancing their performance. For the same reason, the entire simulated data set is used in the training and testing of the classifiers, though in this case the simulated data is not divided into two disjoint subsets for use with the two independent classifiers. Rather, the entire data set is utilized for each classifier.

The discriminating variables used to create the BDT classifiers are listed in Table 4.3. Among them are the isolation variables $\mathcal{A}^{p_T}$, $V_{\text{Mult.}}$, and $\min(\Delta\chi_{\text{vtx}}^2)$, which depend on the multiplicity of non-signal tracks in the event. To ensure that the simulated $B^+ \rightarrow K^+\pi^0$ distributions of these variables are correct, event weights are generated from the ratios of these distributions between experimental and simulated $B^0 \rightarrow K^+\pi^-$ events. These corrections are important because the discriminating power of the isolation variables tends to be overestimated in simulation. Classifiers

Table 4.3: Discriminating variables used to create the BDT classifiers. Some variables are transformed by the natural logarithm to create broader distributions, which can result in more efficient training.

B^+	K^+	π^0
p	p	p
p_T	p_T	p_T
η	$\ln(\chi_{\text{IP}}^2)$	
$\mathcal{A}^{p_T}$	$\ln(\chi_{\text{MT-DOCA}}^2)$	
vtxMult		
$\ln(\min(\Delta\chi_{\text{vtx}}^2))$		

trained with uncorrected simulated data could therefore under-perform and even suppress the true $B^+ \rightarrow K^+\pi^0$ signal decays.

The isolation variable distributions of the $B^0 \rightarrow K^+\pi^-$ decay are extracted from the experimental data using the *sPlot* technique [61]. For the $B^0 \rightarrow K^+\pi^-$ events, vertex isolation variables are generated by attempting to form vertices using only the K^+ candidate track, as opposed to the reconstructed $K^+\pi^-$ vertex. This ensures that the variables are defined the same as in the $B^+ \rightarrow K^+\pi^0$ mode. The selection requirements used to define the data sets for the $B^0 \rightarrow K^+\pi^-$ decay are given in Table 4.4. To make economic use of computing resources, it is necessary to re-use candidates created for previous $B^0 \rightarrow K^+\pi^-$ analyses, and the requirements for that pre-existing selection are listed under the “Prior” column. The additional requirements listed in the table are made in order to make the selection of $B^0 \rightarrow K^+\pi^-$ candidates as similar as possible to that of the $B^+ \rightarrow K^+\pi^0$ candidates.

The $K^+\pi^-$ invariant mass distributions of the selected events are shown in Fig. 4.11. Background contributions to the $K^+\pi^-$ distributions are expected from the $B^0 \rightarrow \pi^+\pi^-$ and $B_s^0 \rightarrow K^+K^-$ decays due to misidentification of a π or K meson candidate, but the $B^0 \rightarrow \pi^+\pi^-$ component is determined from simulation to be negligible. A contribution from the $B_s^0 \rightarrow \pi^+K^-$ decay is expected as well. Fits to the invariant mass distributions are used to generate event weights, with separate fits performed on the 2011 and 2012 data sets corresponding to the 7 TeV and 8 TeV center-of-mass energies. The invariant mass line shape of the $B^+ \rightarrow K^+\pi^-$ decay has been studied in detail [62–64], but a simplified approach is taken here. The resulting fit models, while not appropriate for precise parameter estimations, are sufficient for

Table 4.4: Selection requirements for $B^0 \rightarrow K^+\pi^-$ candidates. Requirements added to make this and the $B^+ \rightarrow K^+\pi^0$ decay selections more similar are listed separately from the pre-existing requirements. Generic hadrons are denoted as h .

Particle	Variable	Prior	Additional
$h^\pm$	p_T	$> 1.1 \text{ GeV}/c$	
	IP	$> 150 \mu\text{m}$	
	χ_{IP}^2	> 100	
n/a	$\max(h^+p_T, h^-p_T)$	$> 2.7 \text{ GeV}/c$	
	$\max(h^+\text{IP}, h^-\text{IP})$	$> 270 \mu\text{m}$	
	$\max(h^+\chi_{\text{IP}}^2, h^-\chi_{\text{IP}}^2)$	> 200	
K^+	ProbNNk		> 0.8
	ProbNNpi		> 0.2
π^-	p		$> 10 \text{ GeV}/c$
	p_T		$> 2.6 \text{ GeV}/c$
B^+	p_T	$> 1.5 \text{ GeV}/c$	
	IP	$< 80 \mu\text{m}$	
	χ_{IP}^2	< 12	
	proper lifetime	$> 0.6 \text{ ps}$	
	mass	$4.8 < m < 5.8 \text{ GeV}/c^2$	
	χ_{vtx}^2		< 9
	DIRA ²		> 0.9999




Figure 4.11: The $K^+\pi^-$ invariant mass spectrum. The data sets are drawn as black points. The total fit models are drawn as solid blue lines, with $B^0 \rightarrow K^+\pi^-$ “signal” components drawn in dashed blue, $B_s^0 \rightarrow K^+K^-$ components drawn in dashed gold, $B_s^0 \rightarrow \pi^+K^-$ components drawn in dashed green, and combinatorial backgrounds drawn in dashed red.

the generation of event weights. While important, the simulation corrections generated with this procedure are small, hence enhancements of the fit models would produce second-order effects in the sense that there would be corrections of corrections that are already small. In the fit, the $B^0 \rightarrow K^+\pi^-$ component is modeled with a double-sided Crystal Ball function with a central Gaussian function, both sharing a common mean. Both sides of the Crystal Ball function share a common width. The shape of the $B^0 \rightarrow K^+\pi^-$ model is fixed from simulation. Gaussian functions are used to model the $B_s^0 \rightarrow \pi^+K^-$ and $B_s^0 \rightarrow \pi^+\pi^-$ components, with their means and widths fixed from simulation. The combinatorial background is modeled with an exponential function, and its shape is fixed to a preliminary fit in the high mass region. The relative yields of the fit components are determined from a fit to the full mass range. The effects of the weights derived from this procedure can be seen in Fig. 4.12, which shows the distributions of the isolation variables in $B^+ \rightarrow K^+\pi^0$ simulated events both before and after the correction weights are applied.

Classifier response distributions generated by the BDTs are shown in Fig. 4.13, with test samples superimposed over the training samples. The consistency of the distributions of the BDT1 and BDT2 responses indicates that it is unlikely that biases have been created in the BDT classification of $B^+ \rightarrow K^+\pi^0$ signal candidates. The compatibility between the training and testing distributions demonstrates that no significant over-training has taken place.

²The DIRA is the cosine of the angle between the momentum vector of the reconstructed candidate and the displacement vector of its decay vertex from the associated PV.




Figure 4.12: Isolation variable distributions from simulation both before and after correction weights are applied. The raw simulation distributions are drawn in black, and the weighted distributions are drawn in red.




Figure 4.13: Distributions of the response variables generated by the BDT classifiers. The distributions are continuous, in contrast to that of a binary decision tree. Signal distributions are drawn in blue, and background distributions are drawn in red.

4.3.8 Optimization of the final selection

The optimal requirement placed on the classifier response variable is found by maximizing the figure of merit $N_S/\sqrt{N_S + N_B}$, where N_S is the number of signal events and N_B is the number of background events. A plot of the figure of merit as a function of cutting on the classifier response is shown in Fig. 4.14. As annotated in the figure, the final data set is formed by rejecting candidates with a classifier response less than 0.47. The efficiency of this requirement is defined as

$$\varepsilon_{\text{Final}} = N_{\text{Final}}/N_{\text{Prelim.}}, \quad (4.3)$$

where $N_{\text{Prelim.}}$ and N_{Final} are the numbers of $B^+ \rightarrow K^+\pi^0$ decays expected in the data sets defined by the preliminary and final requirements, respectively. This efficiency is estimated from simulation to be about 7%. While optimal, this requirement is very strict as a consequence of the enormous combinatorial background present even after the preliminary selection.




Figure 4.14: The figure of merit $N_S/\sqrt{N_S + N_B}$ as a function of cutting on the combined BDT classifier response.

The total selection efficiency is defined as

$$\varepsilon_{\text{Total}} = \varepsilon_{\text{Geo.}} \cdot \varepsilon_{\text{Reco.}} \cdot \varepsilon_{\text{Prelim.}} \cdot \varepsilon_{\text{Final}}, \quad (4.4)$$

where $\varepsilon_{\text{Geo.}}$ is the efficiency due to the geometric acceptance of the LHCb detector, $\varepsilon_{\text{Reco.}}$ is the reconstruction efficiency, and $\varepsilon_{\text{Prelim.}}$ is the efficiency of the preliminary selection, including the TIS trigger efficiency. Estimations of these efficiencies from simulation are given in Table 4.5. The geometric acceptance is determined at the particle-generation level of the event simulation process. The reconstruction efficiency is estimated by performing a reconstruction using the merged π^0 algorithm, and then imposing very mild selection requirements in comparison to the preliminary selection defined previously. For comparison, the LHCb analysis that discovered the $B^0 \rightarrow K^+ K^-$ decay [65] achieved a total efficiency of 0.68% for selecting the $B^0 \rightarrow K^+ \pi^-$ decay used as a normalization channel.³

Table 4.5: Geometric, reconstruction, selection, and total efficiencies estimated from simulation, with statistical uncertainties.

Efficiency	Value (%)	
$\varepsilon_{\text{Geo.}}$	19.1	± 0.1
$\varepsilon_{\text{Reco.}}$	10.9	± 0.1
$\varepsilon_{\text{Prelim.}}$	2.0	± 0.1
$\varepsilon_{\text{Final}}$	7.0	± 0.2
$\varepsilon_{\text{Total}}$	0.003 ± 0.001	

4.4 Peaking backgrounds

Overall, no backgrounds with a peaking shape similar to that of the signal are found to make a significant contribution to the $K^+ \pi^0$ invariant mass spectrum.

³This efficiency is quoted from an internal LHCb document rather than the paper cited here.

The decays $B^0 \rightarrow K^{*0}\pi^0$ and $B \rightarrow K^+\rho$ might hypothetically appear near the signal peak as partially reconstructed backgrounds, but only negligible contributions are expected since the branching fractions of these decays are an order of magnitude smaller than that of the signal decay. The decays $B \rightarrow K^*\gamma$, with $K^{*0} \rightarrow K^+\pi^-$ or $K^{*+} \rightarrow K^+\pi^0$, could be present as partially reconstructed backgrounds in cases where a photon is misidentified as a π^0 meson. Such a misidentification can occur when candidate π^0 mesons are reconstructed assuming that clusters with an oblong shape in the ECAL are overlapping clusters from two photons. The decay channel that proceeds as $K^{*0} \rightarrow K^+\pi^-$ has been studied in simulation, and is found to make a negligible contribution to the final data set. The complementary channel where the K^* decays as $K^{*+} \rightarrow K^+\pi^0$ is further suppressed by the same low trigger efficiency as that of the signal decay.

4.5 Results of the Run 1 analysis

The mass distribution of reconstructed candidates, after all selection requirements, is shown in Fig. 4.15. To extract the number of signal events, a fit is made to this distribution using a model that consists of an exponential function for the combinatorial background; the tail of a Gaussian function for additional background in the low-mass region; and the sum of two Crystal Ball functions for the signal, to account for the tails on both sides of the signal peak. The Crystal Ball functions share a common mean and width, and their tail shapes are taken from simulation. The background remaining in the final selection consists primarily of partially recon-




Figure 4.15: Mass distribution of the B^+ candidates, selected from the entire Run 1 data set. The data are drawn as black points. The total fit model is drawn as a solid blue line. The signal component is drawn in dashed blue, the combinatorial background is drawn in dashed purple, and the low-mass background is drawn in dashed magenta.

structed $B \rightarrow K^+\pi^0 X$ decays, or B^+ candidates reconstructed with misidentified π^0 candidates. The shape of the combinatorial background is determined from a preliminary fit to the high-mass region where $5.7 < m < 6.0 \text{ GeV}/c^2$. Normalizations of the individual components are allowed to float in the fit, as well as the means and widths of the signal and low-mass background shapes.

A signal yield of 72 ± 26 candidates is found. The statistical significance of the signal is determined from the change in the fit likelihood, with and without a signal component, to be 3.7σ . The signal width of $99 \pm 38 \text{ MeV}/c^2$ is consistent with

the expected resolution from simulation, which is estimated to be $95 \pm 4 \text{ MeV}/c^2$. Several variations of the fit are performed with different shapes used to model the low-mass background. Variations are also performed with and without the signal width fixed from simulation, and with the signal mean fixed to the known value of the B^+ mass. All variations are found to be compatible with the fit shown here, and a systematic uncertainty of the signal yield due to the choice of fit model is estimated to be 10 events.

4.6 Implications and future analyses

Despite a low TIS-only trigger efficiency and a very low selection efficiency after suppressing the combinatorial background, evidence for a signal is found for the $B^+ \rightarrow K^+\pi^0$ decay mode for the first time in the challenging environment of a hadron collider. As a proof-of-concept for studies involving the $H_b \rightarrow h\pi^0$ decay topology, discussed previously in Sect. 4.1, this study can be considered a success. An obvious improvement for future analyses is the implementation of a dedicated HLT2 trigger. Trigger algorithms have indeed been put in place for each of the $B \rightarrow K\pi^0$ decays, as well as the as yet unobserved $\Lambda_b \rightarrow \Lambda\gamma$, $\Xi_b^- \rightarrow \Xi^-\gamma$, and $\Omega_b \rightarrow \Omega\gamma$ decays.

Through these triggers, the present analysis has enabled the possibility of observing these radiative baryonic decay modes for the first time. Depending on the presence of signals and on their yields, analyses of these decay modes could represent new experimental avenues for testing the SM in investigations of the photon polar-

ization in $b \rightarrow s\gamma$ transitions, especially in the $\Lambda_b \rightarrow \Lambda\gamma$ decay mode, as discussed previously in Sect. 4.1.

As will be shown in Sect. 4.8, the new $B^+ \rightarrow K^+\pi^0$ trigger has enabled the unambiguous observation of this decay mode for the first time at a hadron collider, with excellent prospects for a competitive $\mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+\pi^0)$ measurement—one of the most relevant inputs for the $K\pi$ puzzle given the current size of experimental uncertainties. It has recently been suggested that existing measurements of the $B^0 \rightarrow K^0\pi^0$ decay mode contain a previously unrecognized tension with predictions of the SM, providing a new aspect to the $K\pi$ puzzle. The same paper goes on to point out that the CP observables in this decay mode are correlated with new physics models that are particularly intriguing in light of experimental anomalies in the $B \rightarrow K^{(*)}\ell^+\ell^-$ decays [66]. The new $B^0 \rightarrow K^0\pi^0$ trigger removes an experimental barrier in making measurements that are therefore quite intriguing, and an analysis of the Run 2 data collected with this trigger is underway.

Finally, an observation of the $B^0 \rightarrow K^0\pi^0$ decay at LHCb could justify the creation of an exclusive trigger for the $B^0 \rightarrow K^0\pi^0\gamma$ decay mode, which is sensitive to investigations of the photon polarization through a time-integrated CP violation analysis. Studies of this decay mode are even more experimentally challenging than those with an $H_b \rightarrow h\pi^0$ decay topology due to the presence of the additional combinatorial background that comes from adding another neutral ECAL cluster—the γ candidate. The higher trigger rate that would consequentially be necessary might nevertheless be justified due to the discovery potential of photon polarization measurements.

In summary, the present analysis has created excellent prospects for relevant and competitive measurements in the $B^+ \rightarrow K^+\pi^0$ decay mode, and has been responsible for the elimination of barriers in the LHCb trigger for a number of important measurements, adding new possibilities to the ever-expanding LHCb physics program. Details of the trigger for the $B^+ \rightarrow K^+\pi^0$ decay are presented in the next section, followed by the preliminary results using data taken during Run 2 of the LHC.

4.7 Dedicated $B^+ \rightarrow K^+\pi^0$ trigger

The implementation of a trigger algorithm for collecting collision events containing a $B^+ \rightarrow K^+\pi^0$ decay is not trivial, even given the successful selection strategy developed for the Run 1 analysis; restrictions on both the rate at which events are selected by the trigger and the CPU time required per-event to execute the trigger algorithm provide additional constraints.

4.7.1 Resources and constraints

The number of computing nodes in the event filter farm (EFF) has been expanded for Run 2, nearly doubling its processing and storage capacity relative to Run 1. In addition to this hardware upgrade, an extensive upgrade and optimization campaign of the HLT software has resulted in an HLT that runs nearly 40% faster. These enhancements have enabled a full event reconstruction to be run in the HLT2, including tracking, the reconstruction of neutral particle candidates, and PID algo-

rithms. As stated in Sect. 3.2.4, the HLT2 reconstruction during Run 1 consisted of tracking and the creation of rudimentary PID information from matching tracks to calorimeter and muon system clusters. In contrast, the Run 2 HLT2 reconstruction includes the full neutral particle candidate reconstruction sequence—notably including the merged π^0 algorithm, as well as the full PID sequence, including the RICH detectors reconstruction and both the DLL and ProbNN algorithms. The creation of a practical $B^+ \rightarrow K^+\pi^0$ trigger algorithm would not be possible without the presence of the merged π^0 and kaon PID information that the EFF and HLT enhancements have provided for Run 2 data-taking.

The CPU time consumed by the HLT in Run 2 is dominated by the HLT2 event reconstruction, which takes about 650 ms per-event on an EFF node. In contrast, the time spent forming a particle decay candidate in a trigger algorithm is typically on the order of hundreds of nano seconds. Yet the HLT2 configuration in Run 2 includes nearly 400 discrete trigger algorithms, and so an explicit restriction of 250 ns has been established for the CPU time consumed by each trigger algorithm. This is easily achieved for the typical particle candidate selections at LHCb, which feature secondary and even tertiary vertices. But the situation is not so trivial for the $B^+ \rightarrow K^+\pi^0$ selection, where the combinatorial background has the potential to cause the consumed CPU time to skyrocket, and to do so in an unstable way, ranging across several orders of magnitude depending on the multiplicity of tracks and ECAL clusters in the collision event. The timing restriction in fact has a direct influence on the strategy of the event selection of the $B^+ \rightarrow K^+\pi^0$ trigger, discussed in the next section.

Given the processing and storage capacity of the EFF, and the CPU time consumed per-event by the HLT, a total rate of 12.5 kHz has been budgeted for the HLT2 during Run 2. The bulk of the budget is allocated to inclusive triggers, including the topological triggers discussed in Sect. 3.2.4, triggers optimized for selecting baryons, and triggers designed to select hadrons containing charm quarks, such as a generic trigger for the $D^* \rightarrow (D \rightarrow X)\pi$ category of decays. About 2.5 kHz is allocated to “turbo” trigger algorithms designed for particle decays with low backgrounds, where it is practical to save to disk only reconstructed candidates instead of the entire event, substantially reducing the bandwidth of the HLT2.

The remainder of the HLT2 rate budget is allocated to exclusive triggers, each dedicated to selecting events containing a particular particle decay. There are about 200 such triggers in place for Run 2, with output rates and processing time consumption varying substantially among them. A strong majority of them have output rates of 10 Hz or less, which is a necessary scenario given the large number of them. The limited selection power and broad reconstructed B^+ mass peak of the $B^+ \rightarrow K^+\pi^0$ decay, however, has the potential to cause a very large output rate for an exclusive trigger. To account for such cases, a strict upper limit of 50 Hz has been established for any exclusive trigger in general. This restriction effectively puts an upper limit on the signal efficiency of the selection, and the optimization of the trigger algorithm represents an attempt to develop a selection that reaches this limit while remaining within the required rate. Ultimately, a rate of about 48 Hz is achieved for the $B^+ \rightarrow K^+\pi^0$ trigger.

4.7.2 Event selection

The event selection of the dedicated $B^+ \rightarrow K^+\pi^0$ trigger is similar to the preliminary selection described in Sect. 4.3, though the requirements have been entirely re-optimized to account for the increased center-of-mass energy of the LHC collisions, as well as the CPU consumption and output rate restrictions for exclusive HLT2 trigger algorithms. Candidate π^0 mesons are again reconstructed using the merged π^0 algorithm, and they are required to have p_T greater than 3.5 GeV/c, a p greater than 5 GeV/c, and a reconstructed mass within ± 60 MeV/c² of the nominal π^0 meson mass.

To contend with the issue of high CPU-time consumption, the π^0 candidates are required to be TOS with respect to the L0 trigger stage. Similarly, the K^+ candidates are required to be TOS with respect to the HLT1 trigger stage. This strategy significantly reduces the combinatorics in a single collision event, which in turn greatly reduces the consumed CPU time to an acceptable and stable level.

Candidate K^+ mesons are required to have p_T greater than 1.2 GeV/c, p greater than 12 GeV/c, and $\Delta\chi_{\text{IP}}^2$ greater than 50. They are distinguished from other charged particles using the DLL-based kaon PID variable. This PID requirement in particular greatly reduces the combinatorial background and therefore the trigger rate, due to the large number of charged pions that would otherwise contaminate the selection.

The four-momentum of the K^+ and π^0 candidates are added to form B^+ candidates. Each B^+ candidate is required to have p_T greater than 6.5 GeV/c, and

a reconstructed mass within the range $4.0 < m < 6.2 \text{ GeV}/c^2$. Note that this mass range extends $150 \text{ MeV}/c^2$ higher than that of the Run 1 selection. The purpose of this requirement is to provide a larger sample of purely combinatorial background events, which benefits the training multivariate classifiers, as well as fits to the data, but at the cost of a slightly higher trigger rate. The $\chi_{\text{MT-DOCA}}^2$ is required to be less than 8 for all B^+ candidates. This variable is a key ingredient in reducing the combinatorial background and hence the output rate of the trigger, and is in fact critical for practical implementations of any of the $H_b \rightarrow h\pi^0$ family of triggers.

The requirements described above for the trigger selection are summarized in Table 4.6. The isolation variables described in Sect. 4.3 provide excellent discriminating power, but they are computationally expensive to calculate on the time-scale of HLT2 trigger decisions since they involve iterating over all reconstructed tracks in collision events. They are therefore not included in the event selection of the trigger, though they are still utilized in the multivariate selection of the offline analysis described in the next section.

4.8 Preliminary Run 2 results

4.8.1 Data

A plot showing the integrated luminosity of proton-proton collisions recorded by the LHCb detector at the time of this writing is shown in Fig. 4.16. This figure covers all of Run 1 (2010–2012), and the extent of Run 2 that has been completed

⁴In contrast with the Run 1 analysis, the MT-DOCA used here has been implemented in the official LHCb software as a property of the mother (B^+) rather than the daughter particle (K^+).

Table 4.6: Requirements for candidate $B^+ \rightarrow K^+ \pi^0$ decays that define the selection for the Run 2 trigger, where the merged π^0 algorithm is used to reconstruct π^0 candidates.

Particle	Variable	Requirement
π^0	p	$> 5.0 \text{ GeV}/c$
	p_T	$> 3.5 \text{ GeV}/c$
	mass	$79.6 < m < 199.6 \text{ MeV}/c^2$
	n/a	L0 TOS
K^+	p	$> 12.0 \text{ GeV}/c$
	p_T	$> 1.2 \text{ GeV}/c$
	$\Delta\chi_{\text{IP}}^2$	> 50
	n/a	HLT1 TOS
B^+	p_T	$6.5 \text{ GeV}/c^2$
	$\chi_{\text{MT-DOCA}}^2$ ⁴	< 8
	mass	$4.00 < m < 6.2 \text{ GeV}/c^2$

to-date (2015–2017). Beam energies and total integrated luminosities by year are given in the legend of the figure. As can be deduced from the legend, about 3 fb^{-1} were recorded in Run 1, and nearly 4 fb^{-1} have been recorded so far during Run 2. Proton-proton collisions are scheduled to resume in April of 2018, extending through November, with Run 2 concluding four weeks later after a final run of heavy ion collisions. In this final year of Run 2 collisions, the LHC is anticipated to perform similarly as during 2017. The LHCb detector is likewise expected to maintain the performance of 2017, recording an additional integrated luminosity of about 1.7 fb^{-1} of proton-proton collisions. This will result in a final 13 TeV data set larger than 5 fb^{-1} , and a total data set of about 8.5 fb^{-1} .

When comparing the sizes of the Run 1 and 2 data sets, it is important to note that at the 13 TeV center-of-mass energy of Run 2, the production cross section of $b\bar{b}$ quark pairs in proton-proton collisions is about twice as large as at the 7–8 TeV center-of-energy of Run 1 [67].




Figure 4.16: A plot over time of the integrated luminosity of proton-proton collisions recorded by the LHCb detector during the years 2010–2017. Beam energies and integrated luminosities by year are given in the figure. An additional 1.7 fb^{-1} at a beam energy of 6.5 TeV is anticipated to be recorded in the final year of Run 2: 2018.

An analysis of data taken during Run 2 using the $B^+ \rightarrow K^+ \pi^0$ trigger is in progress. The preliminary results of this analysis are used to both validate the trigger and to assess the precision of a measurement of $\mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+ \pi^0)$. These results are based on an examination of data collected during 2016, corresponding to approximately 1.6 fb^{-1} of integrated luminosity recorded at the center-of-mass energy of 13 TeV.

4.8.2 Event selection and results

The structure of the Run 2 analysis differs from that made on the Run 1 data by the use of the new HLT2 trigger algorithm, as well as a modification of the

multivariate selection strategy. The entire preliminary selection of the Run 1 analysis is replaced by the event selection of the new $B^+ \rightarrow K^+\pi^0$ trigger. The discriminating power of the isolation variables used in the Run 1 preliminary selection that cannot be used in the Run 2 trigger is still exploited by including these variables as inputs during the creation of the multivariate classifiers.

The multivariate classifiers are again BDTs, and the simulated data used to represent distributions of the $B^+ \rightarrow K^+\pi^0$ decay have been corrected with new weights derived from the $B^0 \rightarrow K^+\pi^-$ decay. These weights are generated with the identical method described in 4.3.7, albeit with the appropriate Run 2 simulated and real data sets, and an improved fit to the $(K^+\pi^-)$ mass spectrum. The multivariate selection strategy differs in that an additional set of BDT classifiers is created using real data from the $(K^+\pi^0)$ mass region with $4.0 < m < 4.9$ GeV/c². These data are used to represent background events, and include additional background components beyond just combinatorial background, such as partially reconstructed H_b decays. This is in contrast with the initial set of classifiers created using real data from the $5.7 < m < 6.2$ GeV/c² range—a range which consists of pure combinatorial background. To make the final selection, new optimal requirements on the classifier response variables are found with a two-dimensional optimization using the $\epsilon_S/\sqrt{N_S + N_B}$ figure-of-merit, where ϵ_S is the efficiency to select $B^+ \rightarrow K^+\pi^0$ decays as estimated from simulation.

After the requirements on the BDT classifiers responses are applied, a fit is made to the $(K^+\pi^0)$ mass distribution, shown in Fig. 4.17. The signal shape is modeled with the sum of two Crystal Ball functions which share a common mean and

width. Their tail parameters and common width are fixed to values determined from simulation. An exponential function is used to model the combinatorial background.

Due to the substantially increased statistics remaining in the final selection relative to that of the Run 1 analysis, the model used to represent the background below the signal peak is more sophisticated. It is assumed that this background consists of partially reconstructed H_b decays, and that the misidentification of a




Figure 4.17: Mass distribution of the B^+ candidates, selected from the 2016 data set. The data are drawn as black points. The total fit model is drawn as a solid black line. The signal component is drawn in blue, the combinatorial background is drawn in green, and the partially reconstructed backgrounds are drawn together in dashed red. The component representing partially reconstructed backgrounds where only one decay product is not reconstructed is drawn in yellow, while that representing cases where two or more decay products are not reconstructed is drawn in red.

photon as the merged π^0 candidate might also be involved. These contributions are modeled with two Argus functions [68]; one to represent cases where only one H_b decay product is not reconstructed, and one to provide a more general shape for contributions from decays where two or more decay products are not reconstructed. For the shape involving only one missing decay product, the cutoff of the corresponding Argus function is fixed to $m(B^+) - m(\pi^0)$: the difference between the nominal mass of the B^+ and π^0 mesons. It is possible that this shape could model contributions from the $B^+ \rightarrow K^{*+}\gamma$ decay, where the K^{*+} decays as $K^{*+} \rightarrow K^+\pi^0$, the corresponding π^0 meson is not reconstructed, and the photon is misidentified as the π^0 candidate of the $B^+ \rightarrow K^+\pi^0$ candidate. Although no significant contribution from this type of background was expected in the limited statistic of the Run 1 analysis, it could in fact be efficiently selected by the new $B^+ \rightarrow K^+\pi^0$ trigger. Investigations into this possibility are ongoing in case the fit to the data can be further refined. Each of the Argus functions are convolved with a Gaussian function to model the experimental resolution of the $(K^+\pi^0)$ mass, and this resolution is fixed to that of the signal shape, itself being fixed from simulation. The curvatures of the Argus functions are allowed to float in the fit.

A signal yield of 5526 ± 255 candidates is found. A comparison with the previous result of 72 ± 26 events is readily made by noting that the data set analyzed here is approximately half as large as the Run 1 data set, while being recorded at a $b\bar{b}$ cross-section that is roughly twice as large, as reported in Sect 4.8.1. The

efficiency to select $B^+ \rightarrow K^+\pi^0$ decays has therefore improved by a factor of about

$$\left(\text{Size of data set} : \frac{1}{2}\right) \times \left(b\bar{b} \text{ cross-section} : 2\right) \times \frac{\text{Signal} : 5526}{\text{Run 1 signal} : 72} = \frac{5526}{72} \approx 77.$$

That the selection efficiency for the signal can improve by a factor as large as 77 is in accord with estimations derived from simulated data prior to the start of Run 2 data taking.

4.8.3 Prospects

Recall that direct CP violation in a particle decay to a final state f , represented by the amplitude A_f is characterized by the CP asymmetry $\mathcal{A}^{\text{CP}}$, defined as

$$\mathcal{A}^{\text{CP}} \equiv \frac{A_f - \bar{A}_{\bar{f}}}{A_f + \bar{A}_{\bar{f}}}, \quad (4.5)$$

where $\bar{A}_{\bar{f}}$ is the CP conjugate decay to the CP conjugate final state. For experimental purposes, it is useful to write $\mathcal{A}^{\text{CP}}$ in terms of $N_{\pm}$, where N_+ and N_- are the number of observed $B^+ \rightarrow K^+\pi^0$ and $B^- \rightarrow K^-\pi^0$ decays, respectively:

$$\mathcal{A}^{\text{CP}} = \frac{N_+ - N_-}{N_+ + N_-}. \quad (4.6)$$

The statistical uncertainty of a measurement of $\mathcal{A}^{\text{CP}}$ thus defined can be estimated as

$$\sigma(\mathcal{A}^{\text{CP}}) \approx \sqrt{N_{\text{Tot.}}^{\text{Cand.}}}/(N_+ + N_-), \quad (4.7)$$

where $N_{\text{Tot.}}^{\text{Cand.}}$ is the total number of selected candidates, including backgrounds.

Using the results of the fit shown in Fig. 4.17 to estimate $N_{\pm}$, the statistical uncertainty of an $\mathcal{A}^{\text{CP}}$ measurement from the present data is estimated to be

$\sigma(\mathcal{A}^{\text{CP}}) \approx 0.033$. For comparison, prior measurements of $\mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+\pi^0)$ and the world average are given in Table 4.7. The precision of the present estimation of an LHCb measurement is competitive with that of the measurement from the BaBar collaboration. Furthermore, by the conclusion of Run 2, the total data set collected at LHCb by the dedicated $B^+ \rightarrow K^+\pi^0$ trigger is expected to be about 3 times larger than the 2016 data set examined here, implying that the statistical uncertainty of LHCb measurements derived from the entire Run 2 data set can be estimated to be $\sigma(\mathcal{A}^{\text{CP}}) \approx 0.019$, which is more precise than the world average.

Table 4.7: Measurements of $\mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+\pi^0)$, and the world average. The first uncertainties are statistical, and the second are systematic. For the CLEO measurement, only a statistical error is quoted as the systematic error is estimated to be negligible in comparison.

	$\mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+\pi^0)$
World average	0.037 ± 0.021 [69]
BaBar	$0.030 \pm 0.039 \pm 0.010$ [35]
Belle	$0.043 \pm 0.024 \pm 0.002$ [33]
CLEO	-0.29 ± 0.23 [70]

To summarize, we have embarked on the study of the rare decay $B^+ \rightarrow K^+\pi^0$, a decay that is sensitive to new physics effects. An analysis strategy for extracting this decay from the LHCb data of Run 1 has been developed and validated, even in the absence of a dedicated trigger. The success of this analysis has justified the creation of a new dedicated trigger, and has inspired dedicated triggers for additional modes with similar experimental topologies, all of which have been deployed for Run 2 data taking. An analysis of Run 2 data collected at 13 TeV has been developed and optimized. The preliminary result of this analysis has again validated the

analysis procedure, and has validated the new trigger, showing enormous gain in the efficiency to reconstruct and select $B^+ \rightarrow K^+ \pi^0$ decays, with excellent prospects to make the world's best measurement of the CP asymmetry $\mathcal{A}^{\text{CP}}(B^+ \rightarrow K^+ \pi^0)$ using the entire Run 2 data set.

Chapter 5: The LHCb upgrade

5.1 Introduction and strategy

Run 2 of the LHC will conclude at the end of 2018. The following period will be the second long shutdown (LS2) of the LHC, which will extend until the first quarter of 2021, marking the beginning of Run 3. For the collider, the primary purpose of LS2 is for upgrades to the injection system—the system of accelerators mentioned in Sect. 3.1. These upgrades will enable the LHC to operate at double the present instantaneous luminosity, increasing to $2 \times 10^{34} \text{ cm}^{-2}\text{s}^{-1}$. Furthermore, the ultimate design value of the LHC center-of-mass energy of $\sqrt{s} = 14 \text{ TeV}$ will finally be reached.

During LS2, the LHC experiments will undergo detector maintenance and upgrades, especially the ALICE and LHCb experiments [71–73]. In Run 1, the LHCb experiment collected 3 fb^{-1} of integrated luminosity. By the end of the ongoing Run 2, it is anticipated that nearly 5 fb^{-1} of additional data will be collected. Since statistical uncertainties in a counting experiment are proportional to $1/\sqrt{N}$, where N is the number of events collected, doubling this collected data set over then next decade would increase the precision of LHCb physics by only a factor of $\sqrt{2} \approx 1.4$. To more effectively advance the experimental frontier in the LHCb physics

program, the LHCb experiment will operate at a leveled instantaneous luminosity of $2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$, an increase over that of previous runs by a factor of 5, with plans to collect a total of 50 fb^{-1} of integrated luminosity by the conclusion of Run 4 in 2028.

However, the current LHCb detector would have to operate with significantly increased thresholds in the L0 hardware trigger in order to cope with the increased running conditions. Aside from decay modes that are selected by the L0 muon triggers, this results in extremely diminished efficiencies due to the limited discriminating power of the L0 trigger system, such that there is almost no gain in signal yields with increasing luminosity. The bottleneck at the root of this situation is the 1 MHz read-out frequency of the full LHCb detector. The solution chosen by the LHCb collaboration, therefore, is to eliminate the hardware trigger and read out the entire detector in real-time at 40 MHz, allowing trigger decisions to be entirely performed with a full event reconstruction in a new software trigger system.

This strategy requires extensive upgrades of the LHCb detector; several sub-detectors must be re-designed and others modified to enable 40 MHz readout. Most of the read-out electronics therefore need to be replaced. Several sub-detector elements need to be upgraded to cope with the increased occupancy at the increased instantaneous luminosity, and to tolerate the radiation damage associated with collecting 50 fb^{-1} of integrated luminosity. Most substantially, all of the silicon detectors need to be replaced, as well as the drift-tube OT. This configuration of the upgraded detector is shown in cross section in Fig. [5.1](#).




Figure 5.1: The upgraded LHCb detector. Changes that are immediately apparent are the replacements of the VELO with a pixel-based system, the replacements of the TT and T-stations by the UT and SciFi trackers, respectively, and the removal of the PS/SPD detectors and the M1 muon station.

5.2 Upgrade plans

5.2.1 Tracking

The VELO sub-detector is being upgraded with entirely new silicon pixel sensors and electronics, as well as the needful mechanics to support the new sensor modules [74]. New RF foils are also being produced to accommodate a new sensor geometry, while the remaining bulk of the mechanics is being reused from the existing detector. Dedicated front-end pixel ASICs named VeloPix have been developed to be sufficiently radiation-hard for the duration of Runs 3 and 4, to handle the increased occupancy, and to be capable of the requisite 40 MHz read-out. They are

square chips composed of $55 \times 55 \mu\text{m}$ pixels arranged in a 256×256 grid. Three of them are bump-bonded side-by-side onto a rectangular sensor to create a tile. Four tiles are arranged in an “L” shape to create detector modules which cover one half of the LHCb acceptance, with a pair of left- and right-half modules making up a station, as with the current VELO. This configuration is illustrated in Fig. 5.2. The tiles




Figure 5.2: A perspective rendering of the upgraded VELO detector without the RF foils. The sensors are shown in red, the VeloPix ASICs are shown in yellow, and the cooling substrate is shown in blue. Notice how the backsides of the ASICs are visible underneath the sensors for tiles positioned on the opposite module face.

are mounted on opposite sides of a silicon substrate, allowing them to overlap for sufficient acceptance coverage. A series of channels $200 \mu\text{m}$ wide by $120 \mu\text{m}$ deep are etched between the two layers of the silicon substrate, allowing evaporative CO_2 cooling to be routed directly under the high-speed VeloPix chips. Fig. 5.3 shows a cross section of a module to more clearly illustrate this configuration.

Using simulated data, the performance of the upgraded VELO has been studied in detail, as well as that of the current VELO for comparison. Performances in both the 2011 and upgrade running conditions have been evaluated. The upgraded




Figure 5.3: Side view of the upgraded VELO module, showing the arrangement of the VeloPix, the sensors, and the silicon substrate with embedded micro-channels. Units are in μm . Details of the other depicted elements are available in Ref. [74].

VELO is found to have superior performance in every critical metric, from track reconstruction efficiency and ghost rejection, to PV and IP resolutions, to decay time resolution. Details of these studies can be found in Chapter 5 of Ref. [74].

The TT is being completely replaced by the Upstream Tracker (UT)—a silicon-strip detector with 40 MHz read-out capable of withstanding 50 fb^{-1} worth of radiation [75]. While the primary utility of the TT is in offline reconstruction, the UT replacement is a critical element in the upgrade trigger strategy. Its position in the fringe field of the LHCb magnet allows for quick momentum measurements, which in turn allows low-momentum tracks to be eliminated before the full track reconstruction algorithm is run. This position also helps to reduce the number of ghosts that are reconstructed. In the full reconstruction, the use of the fast UT momentum measurements increases the speed of reconstruction of high momentum tracks. Overall, the presence of the UT increases the speed of the complete tracking sequence in the trigger by a factor of 3. Furthermore, the silicon sensors overlap in

the UT design to provide complete coverage of the LHCb acceptance, an improvement over the small gaps between adjacent sensors of the TT. The UT detector will be discussed in much more detail in Chapter 6.

Each of the downstream T-stations are being replaced by Scintillating Fiber Trackers (SciFi Trackers) [75]. The SciFi trackers alone cover the full LHCb acceptance in contrast with the present configuration of drift-tube OTs supplemented with silicon-strip ITs. As required, they are capable of withstanding 50 fb^{-1} worth of radiation, and they are read out at 40 MHz. Side and front views of a SciFi station is shown in Fig. 5.4.

As with the OT, each SciFi station is made of four layers in an $x-u-v-x$ configuration. Each layer is made of 12 modules which measure 52 cm wide by 5 m high. Each of these modules is made of eight fiber mats. The mats are 13 cm wide and




Figure 5.4: Side and front views of a SciFi Tracker station with a representation of a human figure for scale.

are half the height of a module at 2.5 m. Each mat is in turn made of five layers of 250 μm -diameter scintillating fibers, except for those in modules next to the beam-pipe where six layers are used to mitigate performance degradation due to radiation damage. Silicon photomultipliers (SiPMs) are used to efficiently read the low light output of the long fibers. These instruments are particularly susceptible to aging in the LHCb environment, but this is mostly averted by the use of neutron shielding, and any remaining effects are suppressed by operating them at -40°C . They are positioned at the ends of the modules, across the tops and bottoms of the detector planes. Light that travels towards the middle of the modules longitudinally is reflected back to the instrumented ends of the modules by aluminized Mylar film that is at most 100 μm thick after being glued in place.

The performance of the UT and SciFi trackers is represented by evaluations of the performance of the complete upgraded tracking system. The precision of IP measurements and vertex positions is dominated by the performance of the VELO—discussed previously—while the UT and SciFi trackers primarily contribute to overall track reconstruction, and dominate the precision of momentum measurements. Simulations of $B_s^0 \rightarrow \phi\phi$ and $D^* \rightarrow D^0(\rightarrow K_S^0\pi\pi)\pi$ decays have been produced for these evaluations. It is found that the momentum resolution of the upgraded detector is about 10% better than that of the present detector, due mostly to reduced material in the tracking system and therefore reduced multiple-scattering rates. Reconstruction efficiencies, however, are anywhere from 2 to 10% lower for the upgraded detector in the upgrade running conditions than those of the present detector in the 2011 running conditions, depending on the tracking algorithm being

evaluated. Ghost rates have also suffered in the same respect. These performance reductions are primarily due to the harsher post-upgrade running conditions. Indeed, the performance of the present detector is found to be inferior to that of the upgraded detector when both are evaluated in the post-upgrade conditions. More details of these studies can be found in Ref. [75]. Tracking algorithms are still being optimized for the upgraded detector in an effort to recover reconstruction efficiency and ghost rejection performances.

5.2.2 Particle identification

Modifications to the PID system of LHCb are in general much less extensive than those of the tracking and trigger systems [76]. The most significant modification to the calorimeter system is the removal of the PS and SPD. As discussed in Chapter 3, the primary utility of the PS/SPD configuration is discrimination between photons and electrons for the L0 trigger. This role does not exist in the upgraded trigger system where tracking information is used to reliably identify charged clusters in the ECAL. The removal of these redundant detectors will improve the energy resolution of the ECAL, and will simplify the ECAL and HCAL calibration procedures.

Although the calorimeter system is already read out at 40 MHz, the read-out electronics will be replaced to ensure durability under radiation damage, and to comply with the common protocol of the upgrade LHCb electronics. For the same reasons, the read-out electronics of the muon system are being replaced. In addition

to this electronics upgrade, the muon system is being modified with the removal of the M1 station. The purpose of this station is to increase the precision of the fast momentum measurement of the muon system for the L0 trigger. As with the PS/SPD, this function is now unnecessary. The removal of the M1 station also provides additional room for new sub-detectors being researched for future upgrades during LS3. Finally, shielding is being placed about the beam-pipe in front of M2 in order to reduce the occupancy in this area under the increased luminosity.

Several different configurations of the RICH detectors have been studied, involving extending the vertical size of the RICH 1, or combining the two detector functions into one new detector located at the current position of RICH 2 [76]. Results of performance studies as well as practical considerations have lead to the decision to retain the present RICH detectors, with upgrades enabling 40 MHz read-out capabilities. The 1 MHz read-out electronics of the HPDs are contained within the instruments, necessitating their complete replacement. Multianode photomultipliers (MaPMTs) will be used as several options are available commercially, all with external read-out electronics. In addition to these sensor upgrades, the RICH 1 aerogel will be removed because it gives greatly reduced output at the higher upgrade luminosity. This results in a small reduction in PID performance that is outweighed by the advantage of removing the $\sim 3.5\%$ X_0 that it represents. The removal of the aerogel increases the photodetector space available for imaging the rings produced by the C_4F_{10} gas. The RICH 1 optics have therefore been re-optimized to spread the gas rings over the full photodetector area.

5.2.3 Trigger

As mentioned in Sect. 3.2.4, the LHCb software trigger runs on the EFF, a farm of PCs located in the surface Point 8 facilities above the detector. The processing capacity of the EFF will be increased to run the full software trigger of the upgraded experiment [77]. While the upgraded detector is read out at the bunch-crossing rate of 40 MHz, the software trigger only needs to have an input capacity equal to the proton-proton inelastic collision rate of 30 MHz. This is reflected in Fig. 5.5a, which illustrates the main processing steps of the upgrade trigger. Aside from the absence of an L0 hardware stage, the major characteristics of the trigger are an “online” calibration step, and a reconstruction that is very nearly the same as in the offline data processing. In the calibration step, fully reconstructed events are buffered to disk while the output of dedicated calibration triggers are used




Figure 5.5: Trigger and tracking schemes of the upgrade.

to calibrate the PID system and to align the tracking system. This allows the trigger algorithms to employ nearly offline-quality selections, enhancing efficiencies and reducing systematic uncertainties relative to strategies that require substantially different online and offline reconstructions. This online calibration and alignment technique has been successfully implemented in the ongoing data-taking of Run 2.

The differences in reconstruction that do exist in the upgrade trigger scheme can be seen in Fig. 5.5b, which shows the online and offline tracking sequences. As mentioned previously in Sect. 5.2.1, quickly-produced momentum measurements from the UT is used to eliminate slow tracks in online tracking. The online track finding algorithms that follow place more aggressive requirements on low- p_T tracks than the offline counterparts. The remaining substantial difference is that a simplified Kalman fit is used to improve the quality of tracks produced by the pattern-recognition algorithms. What is meant by “Kalman fit” is the Kalman-filter-based track fit mentioned in Sect. 3.2.2. This involves matrix algebra and detailed modeling of the detector that is computationally costly, taking 60% of the CPU time spent in the offline tracking sequence [78]. The use of a simplified Kalman fit and the other time-saving steps mentioned here results in a tracking sequence that runs in less than 10 ms per event.

Chapter 6: The Upstream Tracker in the LHCb upgrade

A key element of the charged-particle tracking system of the LHCb detector is the Tracker Turicensis (TT), which is made up of four planes of silicon strip detectors located just upstream of the LHCb magnet. In the upgraded LHCb detector, the TT will be replaced by the new Upstream Tracker (UT), which is also comprised of four planes of silicon strip detectors. Its function extends beyond that of the TT, being a critical component of the online operation of the detector. As a replacement, it meets the read-out speed and radiation-hardness requirements of the LHCb upgrade design. It has increased instrumented area relative to the TT, covering the full LHCb acceptance by eliminating gaps that exist between sensors in the layout of the TT.

At the time of this writing, the design of the UT is nearly completed, and several key components have entered production. While details for some elements have yet to be finalized, the information presented here is anticipated to very closely match the ultimate design of the detector.

6.1 Overview and geometry

In addition to 40 MHz read-out capability required by the LHCb upgrade strategy, the UT electronics must withstand the radiation damage that corresponds to

being exposed to 50 fb^{-1} of integrated luminosity—a condition five times greater than that of the TT. This requirement can be deduced from the plots of expected fluence and dose shown in Fig. 6.1. These plots are generated from simulation, evaluated at $x = 0$, and at z coordinates that correspond to the extents of the UT active area. Taking a safety factor of four, the plots show that the components near the beam-line need to withstand an exposure of 40 MRad. A substantial amount of the UT electronics are located at the periphery of the active area above and below the detector. The vertical (*i.e.*, y) extent of the UT active area is about 140 cm from the beam-line. The peripheral electronics therefore need to withstand an exposure of about 100 kRad.

The active area of the UT is arranged in the familiar x - u - v - x configuration of four detector planes as shown in Fig. 6.2. The sensors in this diagram are color-coded according to their segmentation and geometry. Sensors closer to the beam-line have




Figure 6.1: Plots of the fluence (left) and dose (right) deposited by 50 fb^{-1} of integrated luminosity as a function of the y coordinate and with $x = 0$, as estimated from simulation. The z coordinates plotted here correspond to the extents of the UT active area.

finer segmentation, both with respect to the silicon-strip pitch and the size of the sensors, in order to account for the progressively higher occupancies in this region.

Each detector plane is made up of staves—mechanically independent assemblies of single columns of sensors, indicated in Fig. 6.2 and shown explicitly in Fig. 6.3. This mechanical segmentation allows the sensors to overlap along the x axis by staggering the z positions of adjacent staves. Staves are 10 cm wide and 160 cm high, corresponding roughly to the width of a sensor and the full height of the LHCb acceptance, plus additional height for mechanical, electrical, and cooling interfaces. The upstream “UTa” detector planes are composed of 16 staves, while the “UTb” planes are composed of 18 staves to cover the expanded acceptance roughly




Figure 6.2: A diagram of the UT detector planes with names, as seen from the VELO. Different colored cells represent different types of sensors. The grey rectangle highlights an example of a mechanical staff.

26 cm further downstream. Staves adjacent to the beam-pipe have a semi-circular cutout to fit closely around the pipe. This is a further improvement over the acceptance coverage of the TT, which has a square gap that is intrinsically larger than a circular gap by a factor of $4/\pi$, or about 27% larger. The absolute area gained is not large, but the performance gain is not insignificant given that the multiplicity of tracks is highest near the beam-line.

Sensors are rendered in green in Fig. 6.3. Only half of the sensors are visible in this view because the other half are mounted on the opposite side, allowing them to overlap along the y axis in aggregate. Front-end electronics are positioned at




Figure 6.3: An example of a stave with various elements labeled. The general structure shown here is common to all staves, though this particular example is of a central stave, as evidenced by the semi-circular cutout for the beam-pipe.

the edge of the sensors, which reduces electronic noise relative to propagating the analog signals to the periphery of the detector acceptance. For a given sensor, these electronics are hosted on a printed circuit board (PBC) referred to as a hybrid which, together with a sensor, makes up a detector module.

Modules are mounted on thin Kapton PCB cables referred to as DataFlex cables. These flex cables route all electrical signals to the ends of the stave, including low-voltage (LV) for the front-end ASICs, and high-voltage (HV) for the silicon sensors. Each flex cable supports half of the modules of one side of a stave, so together with the elements on the other side, the flex cables make up four electrically independent assemblies per stave. The flex cables add mass to the detector, but the sensors are thinner than those of the TT (320 vs. 500 μm), resulting in about the same total radiation length as the TT.

Flex cables are laminated with epoxy onto a carbon-fiber foam core which provides the mechanical integrity of the stave. Evaporative CO_2 cooling is provided by a titanium tube sandwiched in a milled channel between layers of the foam core, as shown in Fig. 6.4. The tube traverses the entire length of the stave and is routed through a series of 90° bends directly under all of the front-end ASICs, which dominate the heat load of the modules. The total thickness of an assembled stave is about 7 mm, and the stagger between adjacent staves in the z direction is 15 mm.




Figure 6.4: A layer of the carbon-fiber foam stave core. The titanium CO₂ cooling tube is shown in purple, superimposed over the core. In assembly it is sandwiched in a milled channel between this and another layer.

6.2 Sensors and front-end electronics

The functionality of a silicon particle detector is based on the behavior of the p - n junction—the interface between positively and negatively doped silicon. This junction acts as a diode, and has an intrinsic electric field which creates a volume devoid of free charge carriers known as the depletion zone. A junction is said to be reverse-biased when an external voltage is applied by connecting the positive lead to the n material and the negative lead to the p material. By using a high enough voltage in the reverse-bias configuration, the depletion zone can be increased to encompass the entire volume of the silicon. High-energy charged particles that pass through the silicon ionize atoms, creating a current in the depletion zone that can be amplified and measured, thus signaling the presence of the traversing particle. The full-depletion voltage of the UT sensors is about 180 V, though they will be

operated at higher voltages over the lifetime of the experiment to compensate for noise caused by accumulating radiation dose.

There are four different types of sensors, illustrated in Fig. 6.5. Types A and B measure 97.5 mm wide by 99.5 mm high, with an active area of 96.0×97.5 mm. Types C and D are about half as high at 49 mm. The circular cutout of the Type D sensors has a nominal radius of 35 mm, though this could increase slightly depending on the insulation needs of the beam-pipe. As noted in the figure, Type A sensors have 512 strips with 190 μm pitch. They are “*p-in-n*” sensors—meaning that the strips are p^+ implants in an *n*-type bulk. Types B, C, and D have 1024 strips with a finer 95 μm pitch, and use an *n-in-p* configuration for its enhanced radiation hardness. More than 90% of the sensors are of Type A, and so the more conventional *p-in-n* technology is used as it is less expensive than the more recently developed *n-in-p* technology. The UT sensors are novel in that the pitch adapters for matching the 80 μm ASIC pitch are embedded in the sensor itself, and that the bias voltage is applied through the top of the sensor with wire bonds to reduce the risk of high-voltage shorts.




Figure 6.5: An illustration of the different types of silicon-strip sensors, superimposed to-scale over the raw silicon wafers. Color-codes in this diagram match that of Fig. 6.2.

A stack-up of the components that make up a detector module is shown in Fig. 6.6. The sensor is wire-bonded to a low-mass Kapton-based hybrid which is in turn wire-bonded to the DataFlex. Front-end ASICs are mounted on the hybrid. The ASICs are of a dedicated design named SALT (Silicon ASIC for LHCb Tracker), which supports 128 readout channels. Type A sensors therefore require 4 ASICs for read-out, while the other Types require 8, necessitating two types of hybrids to host them. Both hybrid designs are essentially identical up to the number of ASICs supported. The sensor and the hybrid are glued with epoxy onto a ceramic stiffener, which supports the sensor along only two edges to avoid over-constraining the thin tile of silicon under varying thermal conditions. Three tabs on the stiffener extend




Figure 6.6: An exploded view of the module diagramming the main components. The anchor tabs of the stiffener extend beyond the profile of the hybrid and the sensor, allowing the module to be removed and replaced by breaking the stiffener at the tabs.

beyond the profile of the sensor and the hybrid, and the hybrid is attached with epoxy to the DataFlex at these three points. This provides a contingency for the removal of a module by breaking the stiffener at the anchor tabs. A tab-less module can then serve as a replacement, albeit with a more permanent mounting with epoxy under the main body of the stiffener.

A fundamental requirement following from the LHCb upgrade strategy is that the sub-detector front-end electronics will perform digitization and digital signal processing (DSP)—particularly zero-suppression—dramatically reducing the total bandwidth needed to read out the entire detector at 40 MHz. The SALT chip is a custom radiation-hard ASIC that achieves this requirement for the UT. It is a 1.2 V, high-speed, 11.0×4.5 mm integrated circuit designed in CMOS 130 nm technology. It performs digitization and the crucial DSP functions, and outputs signals using the scalable low-voltage signaling (SLVS) standard.

A block diagram of the SALT chip is shown in Fig. 6.7. Each of the 128 channels is equipped with a preamplifier, shaper, and a 6-bit analog-to-digital converter (ADC). This is followed by the single DSP block, common to all the 128 channels. The DSP block performs a sequence of algorithms, starting with the masking of any bad or noisy channels. This is followed by an algorithm that removes the small voltage offsets in each channel that are present even when no signals are present—a procedure known as pedestal subtraction. Next, a common-mode-suppression algorithm uses all un-masked channels to estimate the electronic noise common to all of them. Finally, the zero-suppression algorithm applies previously determined thresholds for each channel to reject false hits. The DSP output is then serialized




Figure 6.7: A block diagram of the SALT chip. The 128 analog channels are represented on the left. The data output and TFC input SLVS e-ports are on the right.

and sent out as SLVS signals through up to 5 serial ports referred to as e-ports, each with a bandwidth of 320 Mbit/s. The location of the sensor determines the number of active e-ports; the channels of sensors closer to the beam-line have higher occupancy, requiring greater total bandwidth to fully read out. Very nearly all of the Type A sensors require a bandwidth corresponding to SALTs with only 3 e-ports, so the inactive ports are disabled on all SALTs in general in order to save power. An additional e-port is used to receive timing and fast control (TFC) signals. It is the TFC interface that provides the common 40 MHz clock of the LHC. An additional control interface to the SALT called the experiment control system (ECS) provides the common monitoring and control of the entire upgraded LHCb detector. An industry standard I2C bus (Inter-Integrated Circuit) operating at 400 kHz is used to carry the ECS commands.

The DataFlex is a large Kapton-based PCB. It is designed to be low-mass, with three layers of signal and power, adding up to a total thickness of 404 μm . As a passive interconnect cable, it carries the 320 Mbit/s SLVS signals, all of the monitoring and control channels, and both the LV and HV power lines. A stack-up of the electrical layers and substrates is given in Table [6.1](#).

The SLVS signal lines are differential pairs—a scheme where the electrical signal is the voltage difference between a pair of copper traces with a common current flowing in opposite directions. This configuration has good resistance to noise that could be induced by nearby signals or external sources, a quality especially important in the flex cables given their lengths of about 60 to 80 cm.

Table 6.1: A stack-up of the DataFlex layers.

Layer name	Thickness (μm)	Purpose
Kapton (coverlay or solder mask)	25.0	Substrate
Glue	25.0	Substrate
Top Copper Layer	22.5	Signal and HV
Kapton	75.0	Substrate
Glue	25.0	Substrate
Inner Copper Layer	34.0	LV and HV
Kapton	50.0	Substrate
Glue	25.0	Substrate
Kapton	50.0	Substrate
Bottom Copper Layer	22.5	Signal and HV
Glue	25.0	Substrate
Kapton (protective layer)	25.0	Substrate

An example of the layout of the flex cable routings is given in Fig. 6.8. In this figure, the LV power planes are drawn in green, and the signal and HV traces are drawn in red. The large green rectangular element visible on the left end of the cable is a temporary stiffener for mounting the BGA connector. The large pads at the left-most edge are for the HV connectors. The HV lines are routed along the perimeter of the LV and signal routings. They are designed to withstand 1000 V between adjacent lines, and have been tested up to 2000 V. There are up to 240 signal traces (120 differential pairs) in the data flex cables, including the timing and control signals. They have an impedance of about 107Ω , and a round-trip voltage drop less than 300 mV.




Figure 6.8: An example of the signal and power layers of a DataFlex cable. The LV power planes are drawn in green, and the signal and HV traces are drawn in red. The large green rectangular element on the left end of the cable is a temporary stiffener for mounting the BGA connector. The location of the BGA connector is clearly indicated by the compact array of bonding pads. The large pads at the left-most edge are for the HV connectors. The HV traces are routed along the perimeter of the LV and signal routings.

An “eye diagram” is shown in the lower half of the oscilloscope reading displayed in Fig. 6.9. This diagram was made by overlaying many synchronized oscilloscope sweeps from a data stream measured from a prototype DataFlex cable. The rise and fall of different bit sequences produces the central shape that resembles an eye. Such a diagram is intended to characterize the signal integrity of the circuit being tested. That the eye is open in Fig. 6.9 indicates that the high-speed bit sequences in the DataFlex cable can indeed be resolved as digital bits.



Figure 6.9: An oscilloscope reading from a test of a prototype DataFlex cable. The lower half of the readout shows an eye diagram. That the eye is open indicates that the flex cable has good signal integrity.

Due to the different numbers and positions of detector modules, there are three DataFlex designs, the geometries of which are shown in Fig. 6.10. While differing

mechanically, these are essentially three variations of a common design. Fig. 6.11 shows the end of the DataFlex nearest the beam-line, where the wire-bonding pads for a detector module are visible. Dedicated wire-to-board connectors are mounted at the other end of the flex cable for the sensor HV lines. All other channels are interfaced through a high-density, 400 pin ball grid array (BGA) connector. The soldering pads for these connectors are shown in Fig. 6.12. All three of the DataFlex variants share the same BGA pinout.




Figure 6.10: The geometries of the three DataFlex variations, including positions of the detector modules. The BGA and HV connectors are depicted at the top. The length of the longest cable will ultimately be determined by the final geometry of the Type D sensor.




Figure 6.11: A photograph of the end of the DataFlex nearest the beam-pipe. The wire bonding pads of the hybrid are visible at the middle of the edge of the cable. High-voltage bonding pads are visible on the right. The module positioned here is mounted directly on the stave—entirely off of the DataFlex. The cut-outs shown here accommodate the anchor tabs of that module, to prevent it from otherwise being mounted at a slight angle due to the thickness of the flex cable.




Figure 6.12: A photograph of the end of the DataFlex nearest the detector periphery. Mounting pads for the 400-position BGA connector are visible in the center of the image, and those of the HV connectors are visible on the lower-right. A temporary stiffener for mounting the connectors is also visible under the flex cable.

6.3 Peripheral electronics

The staves are entirely contained in a thermally insulated and gas tight box referred to as the detector box. The electronics located immediately outside of the detector box are referred to as the peripheral electronics processing interface (PEPI). Flexible PCBs called pigtails connect the DataFlex cables to the PEPI electronics using BGA connectors at both ends. HV for the sensors is provided by discrete cables independent of the pigtails. The PEPI electronics are positioned on the top and bottom of the detector box, and it is the electrically passive pigtails that interface mechanically with the box. As the name suggests, the geometries of the pigtails are not trivial, and will be discussed later in greater detail.

A diagram of the PEPI electronics is shown in Fig. 6.13. All of the active PEPI electronics are located on data control boards (DCBs). Each DCB hosts TFC and ECS control chips, optoelectronic components that enable high-speed links to the surface counting room, and Giga-Bit Transceiver (GBT) chips that serialize the SLVS signals from the SALT chips for transmission on the optical links. “GBT” in general refers to a family of radiation-hard ICs developed at CERN [79]. The serializers mounted on the DCB are named GBTx, which group data from the e-links of multiple SALT chips into the GBT data frame for transmission at 4.8 Gbit/s. A single GBTx chip processes up to 12 SALT e-links, corresponding to 2 to 4 SALT chips. The optical transmitting and receiving components were also developed at CERN, originating from a joint project by the ATLAS and CMS collaborations [80]. They are the radiation-resistant, magnetic field tolerant Versatile Link




Figure 6.13: A diagram of the PEPI electronics, and their interface to the front-end SALT ASICs. Dominating the diagram are the data control boards (DCBs), which host the ECS, TFC, and data concentrator functions. The dense interconnect backplanes are entirely passive.

Transceiver (VTRx) and Twin-Transmitter (VTTx), operating at 2.5 V. Each DCB is equipped with one VTRx for control commands, and has the capacity to host six GBTx chips and three VTTx units—one for each pair of GBTx chips.

Interface boards in the counting room titled SOL40 and TELL40 send control signals and receive event data, respectively. They are the common control and read-out boards of the upgraded LHCb detector. Specifically, the TELL40 boards align the fragmented event data from the various GBTx chips of each sub-detector for

event building, and finally for trigger filtering and storage through the event filter farm (EFF).

DCBs are mounted on backplanes, as depicted in Fig. 6.14. The backplane is an electrically passive PCB which interconnects DCBs, pigtails, and LV power. Its physical size is 30 by 28 cm, and it is extremely dense, requiring 24 copper layers to perform its signal and power routing functions. Each backplane supports twelve DataFlex cables through pigtails, and hosts up to twelve DCBs. The pigtail is a totally passive element that provides an essentially one-to-one mapping between the common BGA pinout of the DataFlex cables and the BGA pinout on the backplane. All of the pigtail connectors on the backplane therefore share a common pinout.

The various LV channels (1.2, 1.5, and 2.5 V) are routed into the backplane through nine connectors positioned at the top of the board. LV regulator boards




Figure 6.14: A rendering of the backplane. The BGA connectors along the bottom of the backplane are pigtail connectors. The larger BGA connectors across the middle of the backplane are DCB connectors. Two DCBs are shown for reference. The VTTx and VTRx connectors are visible at the edge of the DCBs. LV breakout boards are shown connected along the top of the backplane.

positioned in service bays about 7 m from the UT detector are the origin of these regulated DC power lines.

A total of 24 backplanes are required to support all of the DataFlex cables of the UT detector. Due to the different number of SALT ASICs on different detector modules and the different number of active e-links among those ASICs, the routings of the 24 backplanes are not necessarily similar. A mapping and depopulation scheme has been developed that nevertheless results in only two master designs for all of the backplanes. This is possible because all of the DataFlex cables share a common BGA pinout and depopulation scheme, and because there is a single master DCB design, compatible with any DCB socket on the backplane. Backplanes in positions that support DataFlex cables with significant e-link depopulations require only 8 or 10 DCBs to be installed. Furthermore, some DCB sockets are mapped to a number of e-links which, under this depopulation scenario, corresponds to only 4 GBTx chips. In these cases a DCB is used that has only 4 GBTx chips and only 2 VTTx units installed in order to save power, cost, and fiber-routing space.

6.4 Mechanical design

Much of the electronics systems described previously were designed independent of mechanical integration concerns. But for the region at the ends of the staves outwards, the number and severity of mechanical and electronics constraints results in a complete mechanical integration solution that is entirely non-trivial. The mechanical design process in fact significantly influenced the design of the PEPI electronics and the detector box, and introduced the pigtail into the signal and power transmission scheme.

It would not have been possible to arrive at the solution described here without the development of a high-fidelity CAD model during the early stages of the designs of the PEPI electronics, the detector box, and all of the mechanical and electrical interfaces in the region at the ends of the staves.

6.4.1 External and internal constraints

Fig. 6.15 shows a view from the A-side of the TT station in the pit. The TT is in the center of the image, and immediately upstream and in close proximity is the RICH 1. The most obvious structure downstream from the TT is the magnet, but between the TT and the magnet are arms that help support the beam-pipe. These arms are in fact very close to the volumes to be occupied by the PEPI electronics. The proximities of these arms and of the RICH 1 were not an issue for the design of the TT, where the electronics and cabling for reading out the analog signals required only a modest amount of space. The PEPI electronics, on the other hand, take up far




Figure 6.15: A view of the TT station from the A-side of the cavern. The TT (A) is visible in the center in its closed position. The rails (B) and cable chains (C) are above and below the TT. The blue RICH 1 (D) is to the right of the TT, and the yellow coils of the magnet (E) are visible on the left. The service bays (F) are in front of the magnet yoke (G).

greater space and require dedicated cooling. Fitting the PEPI chassis into the space designed for the analog TT electronics therefore makes the proximities of the RICH and the beam-pipe arms significant mechanical constraints, as shown in Fig. 6.16. To alleviate this, the total space available for the PEPI chassis has been enlarged somewhat by increasing the height of the I-beam relative to that of the TT.




Figure 6.16: A detail of the PEPI region of the UT corresponding to the view in Fig. 6.15. The RICH 1 and hardware supporting the rails is shown in yellow on the right. The arms used to support the beam-pipe are shown in red on the left. Note the “T” shaped cross section of the blue I-beam between the two PEPI chassis.

A direct consequence of increasing the volume allotted to the PEPI chassis is that there is less space overall inside the detector box. Yet several mechanical interfaces exist in close proximity at the ends of the staves, and space available for these interfaces is an asset. The division of the total available vertical space into allocations inside and outside of the detector box must therefore be carefully balanced. Furthermore, the exact positioning along the z axis of the detector planes is influenced by the particular solutions for the mechanical interfaces located at the ends of the staves. The positions of the detector planes has a direct influence on the shapes of the pigtails, which again influence the mechanical composition of the PEPI electronics outside of the box. Altogether, these conditions mean that the mechanical designs of all of the detector components from the ends of the staves outward are in fact coupled into a single mechanical design with many internal constraints.

6.4.2 Overview

A view of the UT and supporting structures is shown in Fig. 6.17, with the LHCb coordinate system indicated in the lower-left corner. The box that encloses the detector planes is visible in the center of the image, and one half-face of the box is removed to show the staves inside. The detector box is made of woven carbon fiber sheets laminated over a low-mass Airex foam core, with a thickness of 20 mm on the front and back panels, and 25 mm on the top, bottom, and side walls. A layer of copper mesh is also included in the laminate to provide electrical shielding.




Figure 6.17: A perspective view of the UT station. The UT (A) is in the center in the closed position. The yellow rails (B) are above and below the UT, and the cable chains (C) are in front of them. At the ends of the chains are the service bays (D). The cable chains and service bays shown here are those of the TT, acting as stand-ins.

To prevent condensation on the sensors or the beam-pipe, an ambient temperature of $-10\text{ }^{\circ}\text{C}$ is maintained inside the box, and it is filled with dry N_2 gas.

The detector is divided into two half-assemblies which travel on precision rails, allowing the halves to open and close independently about the beam-pipe. This functionality decouples the the installation and maintenance of the detector from that of the beam-pipe. Insulation is applied directly to the beam-pipe, and the front and back faces of the detector box have semi-circular polymer gaskets that interface with the beam-pipe. This creates an insulated, gas-tight seal when the detector is closed, which reduces the amount of matter near the beam-line relative to the half-cylinder inserts that are a part of the TT box structure. The UT beam-pipe segment and the plugs are shown in Fig. 6.18.

Before the installation of the beam-pipe segment that intersects the UT, the C-side of the detector is mounted on the rails from the A-side, and then moved to the far C-side end of the rails. Here, A- and C- side stand for the access and cryogenics side of the LHCb detector, respectively. The A-side of the detector is exposed to the large volume of the cavern, while the C-side is relatively obstructed by the cavern walls and the LHC cryogenics. After the C-side of the UT detector is mounted on the rails, the beam-pipe is installed, followed by the mounting of the A-side on the rails. The UT detector uses the same precision rails already installed for the TT.

The box halves are held by C-frame structures, so named because of their profile in the x - y plane. The C-frames are mounted on the rails, and are each constructed from three aluminum I-beams, which take their name from their “I”-shaped




Figure 6.18: Renderings of the polymer gaskets that interface with the beam-pipe. The gaskets are shown in green and the detector box is shown in blue. Only one half of the detector box and gaskets are shown for clarity.


transverse cross section. The PEPI electronics are contained in eight electrically independent assemblies referred to as chassis. To make maximum use of the available space, the PEPI chassis extend into the concave volumes of the I-beam, and integrate directly with the I-beam itself. Four chassis are visible in the overview shown previously in Fig. 6.17, one on each of the top and bottom of each detector half, and the other four are on the opposite side of the detector in the corresponding positions. The four service bays housing the LV regulator boards are located about 7 m away from the detector, near the ends of the rails.

Cable chains carry the LV lines from the service bays to the PEPI chassis end-caps. Optical links from the counting room to the end-caps are also routed through


the cable chains. Similar to a bicycle chain, the cable chains articulate when the detector opens and closes, avoiding the major intervention of disconnecting the entire UT power and signal lines when the detector or beam-pipe are serviced. The cable chains used for the TT are pictured in Fig. 6.17 as placeholders. There are about four times as many read-out chips in the UT as there are in the TT. Furthermore, the optical-link serialization hardware of the TT is located in the service bays, while the UT counterparts—the PEPI electronics—are located at the edges of the detector on the other end of the cable chains. This results in a far greater number of LV lines being routed through cable chains to the UT detector, necessitating the replacement of the TT chains with high-capacity, heavy-duty hardware.

6.4.3 Mechanical design of the peripheral interface region


Fig. 6.19 shows the section inside of the box at the ends of the staves, and the figures that follow step through the various components to show each of them more clearly and to demonstrate their integration.



(a) Front



(b) Side



(c) Side in perspective

Figure 6.19: The region at the ends of the staves viewed from the front (a) and side (b). For clarity, the side view is also drawn in perspective (c). The I-beam is rendered in violet, and the detector box is rendered in pale blue. The front and back faces of the detector box are omitted in (a) and (c). All of the PEPI electronics are omitted for clarity. The remaining components will be identified in the figures that follow.


Fig. 6.20 shows the components that make up the stave frames in red, which provide mechanical support for the staves. Stainless steel rails are attached with screws to the top and bottom inner faces of the detector box. The frames themselves are segmented into individual stainless steel frames for each stave, aiding in the assembly of the detector and the alignment of the staves. Each frame can be positioned independently on the rails, providing fine-tuning of the positions of the staves adjacent to the beam pipe in particular. Openings in the frames allow HV wires and cooling tubes to pass through, and also avoid over-constraining the shapes of the pigtails.

HV cables and routing hardware are in Fig. 6.21. One cable for each stave containing up to 16 insulated wires enters through the side of the detector box and travels along the stave frames, until finally splitting into individual wires for attachment to the DataFlex cables.


The stainless steel CO₂ cooling tubes and manifolds are shown in Fig. 6.22. A winding in each tube provides flexibility during its connection to the stave. Tubes are shaped into very particular configurations in order to fit over the HV cables and between the stave frames and pigtails. A closer view of the integration of the CO₂ tubes, stave frames, and HV cables is shown in Fig. 6.23. A single tube is shown in Fig. 6.24. There are three mechanical variations of the tube shown here for a total of four geometries to reach the different stave positions.



(a) Front




(b) Side




(c) Side in perspective


Figure 6.20: The region at the ends of the staves, where only the I-beam, detector box, and stave frames are included. The stave frames are rendered in red. The individual frames are attached to rails, which are in turn attached to the inner face of the box. The white vertical shapes conspicuous in the side view are fixtures that are embedded in the side of the detector box through which the HV cables pass.



(a) Front




(b) Side




(c) Side in perspective


Figure 6.21: A rendering including the HV cables in teal. A single cable for each stave enters through the side of the box, and breaks out into individual insulated wires to attach directly to the DataFlex cables.



(a) Front



(b) Side



(c) Side in perspective

Figure 6.22: A rendering including the cooling tubes and manifolds in grey. The tubes pass over the HV cables and through the openings in the stave frames to reach the staves.





Figure 6.23: A close-up view showing the mechanical interfaces of the CO₂ tubes, stave frames, and HV cables.




Figure 6.24: A rendering of a single cooling tube. The circular winding allows the tube to flex during the installation of the staves. There are three mechanical variations of the part shown here for a total of four geometries to reach the different stave positions.

Fig. 6.25 shows the staves in place. They are attached to the stave frames with stand-off bolts that set the stave z coordinates, including the stagger of the staves. Staves are fixed in position at the top, and are allowed to slide along the length of the stave at the bottom to avoid stresses induced by thermal contraction during operation of the embedded stave cooling.


Finally, Fig. 6.26 shows the pigtails connected to the staves and fixed into the detector box. As can be seen from the side view, the pigtails have a serpentine shape on the inside of the detector box, which provides flexibility during installation and during thermal contraction. Pigtails fit over the HV connectors on the DataFlex cables, and between the stave cooling tubes.



(a) Front




(b) Side




(c) Side in perspective


Figure 6.25: A rendering including the staves attached to the stave frames. Bolts with stand-offs establish the staggering of adjacent staves in z .



(a) Front



(b) Side



(c) Side in perspective

Figure 6.26: A rendering including the pigtails in blue. Serpentine geometry of the sections inside the box provide extra flexibility.

6.4.4 Mechanical design of the PEPI

A view of the region of one PEPI chassis is shown in Fig. 6.27, where the pigtails are not yet included for clarity. Aluminum base plates with stand-off bars for each backplane are attached to the I-beam. The stand-off bars provide space for components on the back side of the backplane while also stiffening the backplane directly behind the DCB and pigtail connectors.




Figure 6.27: A rendering of the PEPI chassis region showing only the detector box, C-frame, and backplane base plates.

The backplanes are included next in Fig. 6.28. Each chassis includes three backplanes, all of the same common design. They will either be clamped in place along their edges or glued in place with removable epoxy. Both cases avoid the need for through-holes in the backplane—a valuable condition given the extreme density of electrical traces in the backplane.




Figure 6.28: A rendering including the three backplanes attached to the base plates.

The segments of the pigtails on the outside of the box are shown in Fig. 6.29. Pigtails connect to the backplane with a right-angle BGA connector. The use of a right-angle connector allows the segments of the pigtails inside the box to have relatively simple shapes, a condition ultimately critical to having any common solution at all to the mechanical interfaces at the ends of the staves. Consequentially, the shapes of the pigtails outside the box are not at all trivial.




Figure 6.29: A rendering including the pigtails connected to the backplanes.

The undeformed pigtails as they are manufactured are shown in Figs. 6.30 and 6.31. Since the DataFlex and backplane BGA connectors all share a common pinout, each of the pigtails are electrically identical aside from the path-lengths of the copper traces. Mechanically, however, there are four variations of pigtails due to the different positions of the detector planes and the fact that similar pigtail geometries connect to opposite faces of a stave.

Tracks that constrain and help shape the pigtails are shown in Fig. 6.32. They are attached to the I-beam, and within a moderate tolerance, they hold the pigtails in the x - y plane over their corresponding backplane connectors. These mating connectors have guide-posts to further align the connector pairs as they are brought together. The tracks provide for travel along the z axis, allowing the pigtails to be connected and disconnected from the backplane while fixed in place transversely.




Figure 6.30: A rendering of the pigtails that connect the stereo-rotated staves to the backplanes. They are shown here as they are manufactured, before they are deformed into their installed shapes.




Figure 6.31: A rendering of the pigtails that connect the non-rotated staves to the backplanes. They are shown here as they are manufactured, before they are deformed into their installed shapes.




Figure 6.32: Tracks that constrain the pigtails to their corresponding position on the backplane. They provide for travel along the z axis, allowing the pigtails to be connected and disconnected from the backplane while fixed in place in the x - y plane.

Not only are the pigtails deformed into rather convoluted shapes outside of the detector box, but they are tightly packed together, and route between and around each other to reach the backplane. It is not self-evident from a CAD model that manufactured pigtails can actually be deformed into the final planned shapes, or that they can be installed with reasonable effort and success in the presence of adjacent pigtails that have already been installed. Therefore, as a first-order validation of the presented concept, a 3D-printed mock-up of a set of the four mechanically unique pigtails has been created, together with adjacent pigtails, pigtail tracks, and a segment of the detector box and backplane. Photographs of the mock-up are shown in Fig. [6.33](#).




Figure 6.33: Photographs of a 3D-printed mock-up of the pigtails.

Each pigtail is made up of three layers that are mechanically independent, except where they converge at the BGA connectors. This allows them to be more flexible than if they were made of a single mechanical layer, where the sharp bends of their installed geometries would tend to shear the PCB laminate apart. Fig. 6.34 shows a prototype pigtail that has been subjected to deformations similar to those of the the installed geometries. Electrical tests have shown that while the differential impedance of the signal lines of the deformed pigtail does increase relative to the manufactured shape, the performance is still within the design requirements.




Figure 6.34: A photograph of a prototype pigtail subjected to deformations similar to the planned geometry of pigtails installed in the UT.

Enough heat is generated by the GBT chips and optical transceivers that the PEPI chassis require active cooling. To facilitate this, each DCB has an aluminum heat spreader glued onto it that has about the same profile as the DCB itself. Heat pipes that travel the length of the DCB are clamped onto each long edge, providing both a thermal interface and a mechanical support interface. These heat pipes are then clamped onto stainless steel pipes that traverse the length of the PEPI chassis, as shown in Fig. 6.35. Demineralized water with a temperature of 20 °C circulates through the pipes to provide cooling, and the framing of the pipes themselves provides mechanical support for the DCBs. The PEPI chassis are continually flushed with dry air to prevent condensation both on the active cooling elements, as well as the pigtails that exit the cold detector box.




Figure 6.35: A rendering including the DCB cooling pipes and support structures.

The VTT/Rx transceivers on the DCBs use standard LC optical connectors. MPO-to-LC harnesses such as the one shown in Fig. 6.36 are used to link these to MPO ports on the PEPI end-cap. Since a single MPO cable contains 12 optical fibers, this saves a significant amount of space and routing complexity.




Figure 6.36: An example of an optical MPO-to-LC harness that breaks out into twelve LC connectors.

The optical fibers are routed across the DCBs, as illustrated in Fig. 6.37. A common MPO-to-LC harness length is used for all cables within the PEPI chassis. Since the DCBs are positioned at many different lengths from the end-cap, the various extra lengths of the common MPO cables are taken up in loops at the end of the chassis, as can be seen in Fig. 6.37. Additional long-distance MPO cables are used to route the optical signals through the cable chains and all the way out to patch-panels in the counting room.




Figure 6.37: A rendering including the MPO-to-LC optical harnesses connecting the DCBs to the end-cap. The optical cables will be routed either above or in front of the DCBs.

PEPI chassis are coupled to only two of the four DataFlex cables of a given stave. This means that each stave is mechanically coupled to two PEPI chassis—one at each end of the stave, as illustrated in Fig. 6.38. Consequentially, two electrically independent PEPI chassis are assembled simultaneously during the installation of the stave.




Figure 6.38: An illustration of the mechanical coupling of the two PEPI chassis at the opposite ends of a stave.

6.4.5 Interface of the box halves

Critical to the mechanical integration of the various elements at the end of the staves is the simple shape of the pigtails inside the detector box, where a pigtail simply continues along the longitudinal axis of the stave it is connected to. This leads to openings in the box that are shifted along the x axis for stereo-rotated staves relative to the openings for the non-rotated staves. Consequentially, the top and bottom faces of the detector box halves have protrusions and corresponding recessions to accommodate the fixation of the pigtails, as shown in Fig. 6.39.




Figure 6.39: A rendering of a detector box half showing the shape of the box-halves interface.

The shape and length of the pigtail segments outside of the detector box determines the positions of the PEPI chassis along the x direction. This results in the chassis being shifted relative to the edges of the box halves, with half of them

extending over the edge and the other half being recessed from the edge. Fig. 6.40 shows the interface of adjacent PEPI chassis, where the detector halves are partially retracted.




Figure 6.40: A view of the interface of the two detector halves where the two halves are partially retracted. The top plane of the I-beam is omitted for clarity.

6.4.6 Summary

The extreme mechanical and electrical constraints of the LHCb upstream tracker have necessitated the development of a detailed model of the detector, in particular the interface region between the detector box and the peripheral electronics. By developing a high fidelity CAD model of the system, a solution was found and proposed, which has now been deployed in the UT construction, including the production and testing of several key elements of the detector.

Bibliography

- [1] A.A. Starobinsky. A new type of isotropic cosmological models without singularity. *Physics Letters B*, 91(1):99 – 102, 1980.
- [2] Alan H. Guth. Inflationary universe: A possible solution to the horizon and flatness problems. *Phys. Rev. D*, 23:347–356, Jan 1981.
- [3] A.D. Linde. A new inflationary universe scenario: A possible solution of the horizon, flatness, homogeneity, isotropy and primordial monopole problems. *Physics Letters B*, 108(6):389 – 393, 1982.
- [4] A. D. Sakharov. Violation of CP Invariance, C Asymmetry, and Baryon Asymmetry of the Universe. *Pisma Zh. Eksp. Teor. Fiz.*, 5:32–35, 1967.
- [5] M. B. Gavela, M. Lozano, J. Orloff, and O. Pene. Standard model CP violation and baryon asymmetry. Part 1: Zero temperature. *Nucl. Phys.*, B430:345–381, 1994.
- [6] Patrick Huet and Eric Sather. Electroweak baryogenesis and standard model CP violation. *Phys. Rev.*, D51:379–394, 1995.
- [7] Glennys R. Farrar and M. E. Shaposhnikov. Baryon asymmetry of the universe in the standard electroweak theory. *Phys. Rev.*, D50:774, 1994.
- [8] MissMJ [CC BY 3.0 (<http://creativecommons.org/licenses/by/3.0>) or Public domain], via Wikimedia Commons. Standard model of elementary particles, June 2006.
- [9] Emmy Noether. Invariant Variation Problems. *Gott. Nachr.*, 1918:235–257, 1918. [Transp. Theory Statist. Phys.1,186(1971)].
- [10] Serguei Chatrchyan et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Phys. Lett.*, B716:30–61, 2012.

- [11] Georges Aad et al. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Phys. Lett.*, B716:1–29, 2012.
- [12] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay. Evidence for the 2π Decay of the K_2^0 Meson. *Phys. Rev. Lett.*, 13:138–140, 1964.
- [13] Bernard Aubert et al. Observation of CP violation in the B^0 meson system. *Phys. Rev. Lett.*, 87:091801, 2001.
- [14] Kazuo Abe et al. Observation of large CP violation in the neutral B meson system. *Phys. Rev. Lett.*, 87:091802, 2001.
- [15] R. Aaij et al. First observation of CP violation in the decays of B_s^0 mesons. *Phys. Rev. Lett.*, 110(22):221601, 2013.
- [16] K. Lande et al. Observation of long-lived neutral v particles. *Phys. Rev.*, 103:1901–1904, Sep 1956.
- [17] R. Aaij et al. Observation of $D^0 - \bar{D}^0$ oscillations. *Phys. Rev. Lett.*, 110(10):101802, 2013.
- [18] H. Albrecht et al. Observation of B^0 - anti- B^0 Mixing. *Phys. Lett.*, B192:245–252, 1987.
- [19] A. Abulencia et al. Observation of $B_s^0 - \bar{B}_s^0$ oscillations. *Phys. Rev. Lett.*, 97:242003, Dec 2006.
- [20] I. I. Bigi. CP violation. volume 3 of *Adv. Series on Direct. in High Energy Phys.*, pages 175–248. World Scientific, Singapore, 1989.
- [21] Alexey A Petrov. Charm physics: Theoretical review. In *Flavor physics and CP violation. Proceedings, 2nd International Conference, FPCP 2003, Paris, France, June 3-6, 2003*, page MEC05, 2003.
- [22] F. Ambrosino et al. Determination of CP and CPT violation parameters in the neutral kaon system using the Bell-Steinberger relation and data from the KLOE experiment. *JHEP*, 12:011, 2006.
- [23] R. Aaij et al. Measurement of CP violation in $B^0 \rightarrow J/\psi K_S^0$ decays. *Phys. Rev. Lett.*, 115:031601, 2015.
- [24] R. Aaij et al. Measurement of the CP-violating phase ϕ_s in $\bar{B}_s^0 \rightarrow J/\psi \pi^+ \pi^-$ decays. *Phys. Lett.*, B736:186, 2014.
- [25] R. Aaij et al. Precision measurement of CP violation in $B_s^0 \rightarrow J/\psi K^+ K^-$ decays. *Phys. Rev. Lett.*, 114:041801, 2015.
- [26] R. Aaij et al. First observation of CP violation in the decays of B_s^0 mesons. *Phys. Rev. Lett.*, 110:221601, 2013.

- [27] Bernard Aubert et al. Observation of direct CP violation in $B^0 \rightarrow K^+\pi^-$ decays. *Phys. Rev. Lett.*, 93:131801, 2004.
- [28] R. Aaij et al. First evidence of direct CP violation in charmless two-body decays of B_s^0 mesons. *Phys. Rev. Lett.*, 108:201601, 2012.
- [29] Y. Amhis et al. Averages of b -hadron, c -hadron, and τ -lepton properties as of summer 2016. *The European Physical Journal C*, 77(12):895, Dec 2017.
- [30] Bernard Aubert et al. Measurement of time dependent CP asymmetry parameters in B^0 meson decays to ωK_S^0 , $\eta' K^0$, and $\pi^0 K_S^0$. *Phys.Rev.*, D79:052003, 2009.
- [31] M. Fujikawa et al. Measurement of CP asymmetries in $B^0 \rightarrow K^0\pi^0$ decays. *Phys.Rev.*, D81:011101, 2010.
- [32] Bernard Aubert et al. Observation of $B^+ \rightarrow \bar{K}^0 K^+$ and $B^0 \rightarrow K^0 \bar{K}^0$. *Phys.Rev.Lett.*, 97:171805, 2006.
- [33] Y.-T. Duh et al. Measurements of Branching Fractions and Direct CP Asymmetries for $B \rightarrow K\pi$, $B \rightarrow \pi\pi$ and $B \rightarrow KK$ Decays. *Phys.Rev.*, D87:031103, 2013.
- [34] J.P. Lees. Measurement of CP Asymmetries and Branching Fractions in Charmless Two-Body B -Meson Decays to Pions and Kaons. *Phys.Rev.*, D87(5):052009, 2013.
- [35] Bernard Aubert et al. Study of $B^0 \rightarrow \pi^0\pi^0$, $B^\pm \rightarrow \pi^\pm\pi^0$, and $B^\pm \rightarrow K^\pm\pi^0$ Decays, and Isospin Analysis of $B \rightarrow \pi\pi$ Decays. *Phys.Rev.*, D76:091102, 2007.
- [36] Lyndon Evans and Philip Bryant. LHC Machine. *Journal of Instrumentation*, 3(08):S08001, 2008.
- [37] R. Frhwrth. Application of kalman filtering to track and vertex fitting. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 262(2):444 – 450, 1987.
- [38] Christian Lippmann. Particle identification. *Nucl. Instrum. Meth.*, A666:148–172, 2012.
- [39] F. Archilli et al. Performance of the muon identification at LHCb. *JINST*, 8:P10020, 2013.
- [40] M Alemi et al. First operation of a hybrid photon detector prototype with electrostatic cross-focussing and integrated silicon pixel readout. *Nucl. Instrum. Methods Phys. Res., A*, 449(1-2), Jul 1999.
- [41] Roel Aaij et al. LHCb Detector Performance. *Int. J. Mod. Phys.*, A30(07):1530022, 2015.

- [42] R Aaij et al. The LHCb Trigger and its Performance in 2011. *JINST*, 8:P04022, 2013.
- [43] L. Breiman, J. H. Friedman, R. A. Olshen, and C. J. Stone. *Classification and regression trees*. Wadsworth international group, Belmont, California, USA, 1984.
- [44] B. P. Roe, H.-J. Yang, J. Zhu, Y. Liu, I. Stancu, and G. McGregor. Boosted decision trees as an alternative to artificial neural networks for particle identification. *Nucl.Instrum.Meth.*, A543:577–584, 2005.
- [45] V. V. Gligorov and M. Williams. Efficient, reliable and fast high-level triggering using a bonsai boosted decision tree. *JINST*, 8:P02013, 2013.
- [46] M Clemencic et al. The LHCb simulation application, GAUSS: design, evolution and experience. *J. Phys. Conf. Ser.*, 331:032023, 2011.
- [47] Sjöstrand, Torbjörn and Mrenna, Stephen and Skands, Peter. PYTHIA 6.4 physics and manual. *JHEP*, 05:026, 2006.
- [48] Sjöstrand, Torbjörn and Mrenna, Stephen and Skands, Peter. A brief introduction to PYTHIA 8.1. *Comput.Phys.Commun.*, 178:852–867, 2008.
- [49] I. Belyaev et al. Handling of the generation of primary events in GAUSS, the LHCb simulation framework. *Nuclear Science Symposium Conference Record (NSS/MIC)*, IEEE:1155, 2010.
- [50] D. J. Lange. The EvtGen particle decay simulation package. *Nucl. Instrum. Meth.*, A462:152–155, 2001.
- [51] Piotr Golonka and Zbigniew Was. PHOTOS Monte Carlo: a precision tool for QED corrections in Z and W decays. *Eur.Phys.J.*, C45:97–107, 2006.
- [52] John Allison, K. Amako, J. Apostolakis, H. Araujo, P.A. Dubois, et al. Geant4 developments and applications. *IEEE Trans.Nucl.Sci.*, 53:270, 2006.
- [53] S. Agostinelli et al. Geant4: a simulation toolkit. *Nucl. Instrum. Meth.*, A506:250, 2003.
- [54] The LHCb collaboration. LHCb computing: Technical Design Report. Technical Report CERN-LHCC-2005-019, CERN, Geneva, 2005. LHCb-TDR-011.
- [55] Christoph Eck et al. *LHC computing Grid: Technical Design Report. Version 1.06 (20 Jun 2005)*. Technical Design Report LCG. CERN, Geneva, 2005.
- [56] K Harrison et al. GANGA: a user-Grid interface for Atlas and LHCb. (cs.SE/0306085):9 p, Jun 2003.
- [57] Gudrun Hiller and Alex Kagan. Probing for new physics in polarized Λ_b decays at the Z . *Phys. Rev.*, D65:074038, 2002.

- [58] Tomasz Skwarnicki. *A study of the radiative cascade transitions between the Upsilon-prime and Upsilon resonances*. PhD thesis, Institute of Nuclear Physics, Krakow, 1986. DESY-F31-86-02.
- [59] Andreas Hocker et al. TMVA - Toolkit for Multivariate Data Analysis. *PoS, ACAT:040*, 2007.
- [60] Robert E. Schapire and Yoav Freund. A decision-theoretic generalization of on-line learning and an application to boosting. *Jour. Comp. and Syst. Sc.*, 55:119, 1997.
- [61] Muriel Pivk and Francois R. Le Diberder. sPlot: a statistical tool to unfold data distributions. *Nucl.Instrum.Meth.*, A555:356–369, 2005.
- [62] Elisabetta Baracchini and Gino Isidori. Electromagnetic corrections to non-leptonic two-body b and d decays. *Physics Letters B*, 633(23):309 – 313, 2006.
- [63] Angelo Carbone et al. Invariant mass line shape of $B \rightarrow PP$ decays at LHCb. Technical Report LHCb-PUB-2009-031. CERN-LHCb-PUB-2009-031, CERN, Geneva, Jan 2010.
- [64] The LHCb Collaboration. Measurement of direct $\mathcal{CP}$ violation in charmless charged two-body B decays at LHCb. Apr 2011. LHCb-ANA-2011-023.
- [65] Roel Aaij et al. Observation of the annihilation decay mode $B^0 \rightarrow K^+ K^-$. *Phys. Rev. Lett.*, 118(LHCB-PAPER-2016-036. CERN-EP-2016-255.):081801. 9 p, Oct 2016.
- [66] Robert Fleischer, Ruben Jaarsma, and K. Keri Vos. Towards New Frontiers with $B \rightarrow \pi K$ Decays. 2017. arXiv:1712.02323.
- [67] Roel Aaij et al. Measurement of the b -quark production cross-section in 7 and 13 TeV pp collisions. *Phys. Rev. Lett.*, 118(5):052002, 2017. [Erratum: *Phys. Rev. Lett.* 119, no. 16, 169901 (2017)].
- [68] H. Albrecht et al. Search for hadronic $b \rightarrow u$ decays. *Physics Letters B*, 241(2):278 – 282, 1990.
- [69] C. Patrignani et al. Review of Particle Physics. *Chin. Phys.*, C40(10):100001, 2016.
- [70] S. Chen et al. Measurement of charge asymmetries in charmless hadronic b meson decays. *Phys. Rev. Lett.*, 85:525–529, 2000.
- [71] B Abelev et al. Upgrade of the ALICE Experiment: Letter of Intent. Technical Report CERN-LHCC-2012-012. LHCC-I-022. ALICE-UG-002, CERN, Geneva, Aug 2012.
- [72] The LHCb Collaboration. Letter of Intent for the LHCb Upgrade. Technical Report CERN-LHCC-2011-001. LHCC-I-018, CERN, Geneva, Mar 2011.

- [73] The LHCb collaboration. Framework TDR for the LHCb Upgrade: Technical Design Report. Technical Report CERN-LHCC-2012-007. LHCb-TDR-012, CERN, Geneva, 2012.
- [74] The LHCb collaboration. LHCb VELO Upgrade Technical Design Report. Technical Report CERN-LHCC-2013-021. LHCb-TDR-013, CERN, Geneva, 2013.
- [75] The LHCb collaboration. LHCb Tracker Upgrade Technical Design Report. Technical Report CERN-LHCC-2014-001. LHCb-TDR-015, CERN, Geneva, 2014.
- [76] The LHCb collaboration. LHCb PID Upgrade Technical Design Report. Technical Report CERN-LHCC-2013-022. LHCb-TDR-014, CERN, Geneva, 2013.
- [77] The LHCb collaboration. LHCb Trigger and Online Upgrade Technical Design Report. Technical Report CERN-LHCC-2014-016. LHCb-TDR-016, CERN, Geneva, 2014.
- [78] Roel Aaij et al. Upgrade trigger: Biannual performance update. Technical Report LHCb-PUB-2017-005. CERN-LHCb-PUB-2017-005, CERN, Geneva, Feb 2017.
- [79] Paulo Moreira et al. The GBT Project. *Topical workshop on electronics for particle physics*, pages 342–346, 2009.
- [80] Luis Amaral et al. The versatile link, a common project for super-LHC. *JINST*, 4:P12003, 2009.