RICH Particle Identification at the Present and Future LHCb Experiment

Graham Young-Min Kim

A thesis submitted in fulfilment of the requirements
for the degree of Doctor of Philosophy to the
University of Edinburgh

2011

Abstract

LHCb is a high-precision experiment for the study of CP violation and rare decays in B physics. It has two Ring Imaging Cherenkov (RICH) counters for particle identification. The RICH system uses Hybrid Photon Detectors (HPD) for single photon detection. This thesis describes tests performed on individual HPDs for quantum efficiency and vacuum quality. It then presents work done monitoring HPD vacuum quality through regular ion feedback measurements after they were mounted into the RICH. Also discussed is HPD vacuum degradation and how replacement and repair of deteriorated HPDs was implemented using data from the vacuum monitoring. Preparations for an upgrade to LHCb have started. The upgrade will extend the discovery potential of the LHCb by increasing the rate of data collection. This thesis presents particle identification and flavour tagging studies using upgrade simulations. Flavour tagging using kaons was then reoptimised for upgrade conditions and this work is described with discussions of the results.

Declaration

Except where otherwise stated, the research undertaken in this thesis was the unaided work of the author. Where the work was done in collaboration with others, a significant contribution was made by the author.

Graham Kim
January 2011

Acknowledgements

First, I must thank my parents - without them I would not have even been anywhere near the start of the road to a PhD. They have my undying gratitude.

I am grateful to the members of the Edinburgh Particle Physics Experiments group for help and support throughout my PhD. In particular, I must thank my supervisor Franz Muheim and also Stephan Eisenhardt for guiding me from the start of my PhD to the end. Greig Cowan also has my gratitude for helping me learn how to use the distributed computing resources available to me.

I wish to also extend my thanks to the members of the LHCb RICH collaboration, in particular Thierry Gys for looking after me when I was working on the HPDs and Chris Jones for showing me how to use LHCb software. My gratitude also goes to Hans Dijkstra and Marco Musy for guiding me when I started my upgrade analysis.

I thank STFC for financial support throughout my PhD and also for giving me the opportunity to work at CERN for 15 months. My thanks is extended in particular to the members of the STFC UK Liaison office at CERN for helping me settle and live in Geneva while working at CERN.

Finally, I give my thanks to my friends Bradley Wilton and Peter Cibula for keeping me sane while I was on this road and for non-specialist alternate perspectives.

Contents

Abstract i

Declaration ii

Acknowledgements iv

Contents vi

List of Figures xi

List of Tables xxiv

1 Theoretical Background and Motivation 1

- 1.1 Introduction 1
- 1.2 Standard Model 2
  - 1.2.1 Constituent Particles 2
  - 1.2.2 The Strong Force 4
  - 1.2.3 The Electroweak Force 5
- 1.3 CKM Mechanism 6
- 1.4 C, P and CP Violation 9
- 1.5 CP Violation in the B-meson system 10
  - 1.5.1 CP Violation in Decay 10
  - 1.5.2 CP Violation in Mixing 11
  - 1.5.3 CP Violation through Interference between Mixing and Decay 13
- 1.6 CP Violation in $B_s \to \phi\phi$ decays 15
  - 1.6.1 Standard Model Prediction 15
  - 1.6.2 Experimental Prospects at LHCb 17

2 The LHCb Detector 19

2.1 Overview . . . 19

2.2 Charged Particle Tracking . . . 21

2.2.1 Vertex Locator . . . 21

2.2.2 Magnet . . . 24

2.2.3 Silicon Tracker . . . 25

2.2.4 Outer Tracker . . . 27

2.3 Particle Identification . . . 29

2.3.1 RICH Detectors . . . 29

2.3.2 Calorimeters . . . 32

2.3.3 Muon Chambers . . . 35

2.4 Trigger . . . 37

2.4.1 Level-0 Trigger . . . 38

2.4.2 High Level Trigger . . . 39

2.5 Upgrade Considerations . . . 40

3 Quality Testing Hybrid Photon Detectors 43

3.1 Introduction . . . 43

3.1.1 Photon Detector Test Facilities . . . 45

3.2 Standard Tests . . . 46

3.2.1 Overview of Results . . . 47

3.2.2 Ion Feedback . . . 47

3.3 QE Measurements . . . 49

3.3.1 Methodology . . . 49

3.3.2 Repeatability & Aging Tests . . . 51

3.3.3 Comparison with Manufacturer's QE Measurements . . . 52

3.3.4 QE IV Curves . . . 54

3.4 Special Cases . . . 57

3.4.1 Pre-Series and Early Batch HPDs . . . 57

3.4.2 Recovered HPDs . . . 59

3.5 Conclusions . . . 59

4 Monitoring the LHCb RICH2 Sub-Detector 63

4.1 HPD Monitoring . . . 64

4.2 Ion Feedback Measurements . . . 66

4.2.1 Continuous Wave vs Strobescan IFB . . . 67

4.3 HPD Failure at High IFB Rates 71
    4.3.1 The Glow Effect 71
    4.3.2 Threshold for Risk of Failure 73
4.4 Timeline of HPD Vacuum Degradation 75
    4.4.1 Methodology 75
    4.4.2 Vacuum Quality Recovery Cases 77
    4.4.3 Timelines of Failed HPDs 78
    4.4.4 Incorporating the PDTF IFB measurements 79
    4.4.5 Distribution of Timeline Linear Fit Parameters 81
4.5 HPD Lifetime Assessments 84
    4.5.1 Predictions from the Timeline Fits 84
    4.5.2 Failure Susceptibility of Early Batch HPDs 85
4.6 Conclusions 87

5 Flavour Tagging Studies in the LHCb Upgrade using $B_s \to \phi\phi$ 89

5.1 Motivation 89
    5.1.1 Flavour Tagging 91
5.2 Simulating the LHCb Upgrade 93
    5.2.1 Minimal Upgrade Layout 93
    5.2.2 Higher Luminosity Conditions 94
5.3 Selection and Particle Identification 95
    5.3.1 $B_s \to \phi\phi$ Selection 95
    5.3.2 Kaon Particle Identification Methodology 96
    5.3.3 Kaon Particle Identification at Upgrade 99
5.4 Flavour Tagging Performance vs Luminosity 102
    5.4.1 Combined Tagging Performance at Higher Luminosities . . 102
    5.4.2 Track Candidate Selection for Single Particle Taggers . . . 102
    5.4.3 Individual Tagger Performance 104
5.5 Aerogel in the Upgrade 106
    5.5.1 Aerogel Effects on the Selection 107
    5.5.2 Kaon Identification without the Aerogel 107
    5.5.3 Flavour Tagging without Aerogel 110
    5.5.4 Kaon Taggers and Aerogel 112
5.6 Conclusions 113

6 Reoptimising Flavour Tagging 115
6.1 Tagging Track Candidates . . . 116
6.1.1 Candidates Passing Pre-Selection . . . 116
6.1.2 Candidates Passing Selection Cuts . . . 116
6.1.3 Events with Multiple Candidates Passing Selection . . . 117
6.2 Reoptimisation Methodology . . . 121
6.2.1 Preliminary Studies . . . 121
6.2.2 Recursive Optimisation . . . 124
6.3 Reoptimisation of Kaon Tagger Cuts . . . 126
6.3.1 Secondary Pre-selection of Candidates . . . 126
6.3.2 Reoptimisation Progress . . . 127
6.3.3 Tagger Selection Cuts After Reoptimisation . . . 129
6.4 Reoptimised Kaon Tagger Performance . . . 129
6.5 Conclusions . . . 134

7 Conclusions 135

Bibliography 136

List of Figures

1.1 Two of the six Unitarity Triangles of the SM CKM matrix. The lengths are based on the Wolfenstein parameterisation of each triangle's orthogonality relations. Lengths labelled on the left triangle have taken out the common factor of Aλ^3. 7

1.2 Latest results (2010) from the CKM fitter group [8]. The red shaded area represents the 68% CL region where the apex could be. 8

1.3 Feynman diagrams of B_s^0-$\overline{\mathrm{B}}_\mathrm{s}^0$ mixing. The quark and anti-quark of the meson swap flavours by exchanging charged W bosons and temporarily becoming an up-type flavour. The dominant CKM elements for the vertices have been labelled. 11

1.4 Dominant Feynman diagram of the $B_s \to \phi\phi$ decay. The tree level diagram is forbidden in the standard model because this is a flavour changing neutral current decay. 16

2.1 Sideview layout of the LHCb detector, a forward-arm spectrometer. z-axis is defined as the beam axis. 20

2.2 **Left:** polar angles of b and $\bar{b}$-hadrons simulated by PYTHIA the event generator. **Right:** Pion momentum distributions from simulated $B_s^0 \to D_s^- \pi^+$ events where both pions are in LHCb angular acceptance. . . . 21

2.3 Decay length distribution for simulated $B_d^0 \to \pi^+\pi^-$ events where both pions are within LHCb angular acceptance. 22

2.4 **Top:** A top-down view of the VELO in the horizontal plane the nominal interaction point lies. **Bottom:** A front view of one VELO station, illustrating the 2° overlap of both halves when closed and the 6 cm gap made when open for when the LHC is tuning its beam. 23

2.5 B$_z$ measurements of the LHCb magnet (thick line) vs simulated predictions (thin line). For reference: the magnet is centred at $z \simeq 500\,\mathrm{cm}$, RICH1 is centred at $z \simeq 200\,\mathrm{cm}$ and $z = 1000\,\mathrm{cm}$ is beyond the last tracking station. 25

2.6 The third detection layer of the TT, where each module is rotated by +5° to the y-axis. The innermost modules have three readout sectors, modules getting less occupancy are grouped in two sectors for readout. 27

2.7 An IT detection layer. The regions above and below the beam pipe hold one sensor per module, and the regions to the side hold two per module. 27

2.8 **Top:** Diagram of the layout of OT stations (blue) in relation to IT and TT stations (purple). **Bottom:** Cross-view schematic of a straw-tube module. All lengths are in mm. 28

2.9 **Left:** Simulated $B_d^0 \to \pi^+\pi^-$ and background channels, plotted on the same axes in a no-RICH scenario. In particular, the $B_d^0 \to \pi K$ peak is virtually indistinguishable from the signal peak without RICH Particle Identification. **Right:** Polar angle θ distribution of simulated $B_d^0 \to \pi^+\pi^-$ events. The angular and momentum acceptance regions of both RICHes are indicated by the dashed lines. 29

2.10 **Left:** Schematic side-view of RICH1, also indicating the working principle of the RICH detectors. **Right:** Schematic top-view of RICH2. 30

2.11 Lateral segmentation of the SPD/PS and ECAL (left) and the HCAL (right). One quadrant of the detector face is shown. The black area is the area not covered because of proximity to the beam pipe. The left (right) figure shows cell dimensions for the ECAL (HCAL). 33

2.12 View of the HCAL structure. Two master plates with spacers attached are shown exploded with scintillators mounted in the gaps. Wavelength-shifting fibres collect the scintillator light for the PMTs to readout. . . 34

2.13 Front view of a quadrant of a muon station, showing how the detector is divided into regions. Each rectangle represents one chamber. 36

2.14 **Left:** Exploded view of an MWPC muon chamber. The 4 layers are grouped together into pairs by the logical OR. Charged particles passing through the gas layers create ions that drift to the wires to be read out. **Right:** Exploded view of an M1R1 triple-GEM detector. Electrons liberated by the ionising signal particle are accelerated to the readout anode, multiplying at each GEM layer. 37

3.1 A photo of a finished HPD, without the μ-metal shield and a simplified diagram of how an HPD works. It stands 120 mm tall (without the lens cap) and is 87 mm in external diameter. The photocathode is charged to -20kV in comparison to the Anode. Photoelectrons from the photocathode are accelerated towards the anode, where silicon sensors are connected to a readout chip. 44

3.2 Early setup of a PDTF station at Edinburgh. It is largely unchanged even today. From left to right: darkbox where HPD is mounted, HV supply and electronics, MB & FB board for interfacing with the readout chip, other power supplies and picoammeter, all controlled by the PC through the JTAG controller. The oscilloscope checks the blackout conditions using a photomultiplier inside the darkbox. 46

3.3 Summary of HPD test results at PDTF [25]. A-grade HPDs meet all specifications while B-grade HPDs are still usable despite minor shortcomings. 48

3.4 **Left:** Example HPD result for ion feedback, displaying hits and clusters at different time delays between signal light pulse and readout trigger. The red curve is multiplied by a factor of 100 to make it visible. **Right:** Distribution of ion feedback results for all HPDs. Both figures come from [26]. 49

3.5 Photograph of the darkbox for the QE setup. The LED light source is easily repositioned, and a separate light source has been attached to the left. The photodiode provides calibration data before and after HPD measurements. 50

3.6 **Left:** Repeat measurements of QE taken over a year of HPD H527009. All measurements agree within errors. **Right:** QE measurements from the manufacturer (DEP) and at CERN before and after aging the HPD by 10 years' equivalent illumination [28]. As with PDTF, the QE at different wavelengths is plotted. 51

3.7 Correlation between Photonis and PDTF measurements of QE, across the seven wavelengths PDTF measured over. PDTF data points are the first QE measurements at 100 V bias. The pink $y = x$ line is for reference. Photonis did not measure QE at 340 nm directly, so an average between 320 nm and 360 nm was used for this plot. 53

3.8 **Left:** An example of a normal HPD with good agreement between PDTF and Photonis' QE measurements. **Middle:** An old, pre-series HPD with much higher QE from PDTF's measurements than the manufacturer's values. **Right:** The same pre-series HPD with QE measured at 22 V bias instead of the usual 100 V. 53

3.9 **Left:** Summary of QE measurements by Photonis, organised by groups of production batches. Minimum and Typical specs from Table 3.1 are also plotted for reference. **Right:** PDTF values of QE at 270 nm averaged by HPD batch. The running average is also plotted. Errors are RMS widths. 54

3.10 Quantum Efficiency I-V measurements for three example Hybrid Photon Detectors (HPDs) of different vacuum quality. The x-axis plots the bias voltage between the HPD photocathode and anode. The y-axis plots the current drawn from the HPD. Each HPD has two plots: the left shows the total current read out and the right shows the dark current and the noise-corrected photocurrent readings. Note the different y-axis scales. 56

3.11 **Left:** Photon and dark currents of a QE IV test on a degraded HPD, both during ramp up and ramp down. An hour separates the two sets of measuremnts. **Right:** Photon and dark currents of the same HPD a year later, now severely degraded. Note the differnet y-axis scales. . 57

3.12 Correlation between QE IV results and PDTF ion feedback measurements. The ratio of photocurrents measured at 22 V and 100 V is compared with PDTF results. HPDs from early production batches are highlighted in a different colour. 58

3.13 Ratios of QE-IV photocurrent ratios for HPDs recovered from missing pixel columns. The ratios are between 100 V & 22 V measurements on the left and 500 V & 22 V measurements on the right. 60

3.14 QE of an example HPD recovered from missing pixel columns but having vacuum degradation. This HPD passed PDTF standard tests with no problems. 60

4.1 Pixel Hitmap of the 288 HPDs in RICH2. The data are from 3 million events collected with a continuous wave light source in Run 23288 (23 May 2008). Both halves of RICH2 are shown in this hitmap: each RICH2 half has nine columns, labelled C0–C8 and A0–A8 respectively. There are 16 HPDs in each column, labelled 00–15 counting from the top of each column. Logarithmic z-axis represents how many total hits each individual $500\,\mu\text{m} \times 62.5\,\mu\text{m}$ pixel accumulated over the run. . . . 65

4.2 Pixel hitmaps from two runs, filtered in two ways. **Top:** Continuous wave laser illuminated - Run 23288 (23 May 2008) **Bottom:** Dark count - Run 23232 (22 May 2008) **Left:** Only pixel hits in clusters of size $\leq$4 are displayed. **Right:** Only pixel hits in IFB clusters ($\geq$5 hits) are displayed. 68

4.3 Strobescan measurements of the rate of ion feedback (IFB) of HPD A2-06 in RICH2, using the pulsed laser. This HPD has a medium IFB rate. The IFB rate on the y-axis is corrected for dark count. The x-axis is the time delay between the laser pulse and the data acquisition trigger. This set of measurements was taken from inside RICH2 during commissioning. 69

4.4 A comparison between the two methods of measuring Ion Feedback (IFB) rate used in RICH2. The results were all taken on the same day and have been corrected for dark count noise. The x-axis plots the results from the continuous wave diode light source. The y-axis plots the sum of all IFB rates measured in a set of strobe scan runs using the pulsed laser light source. The strobe scan runs measured IFB with a range of delays from 100 ns to 400 ns between the light pulse and the data acquisition trigger. The pink line is a $y = x$ reference line. The inset on the top left is a zoom of the plot closer to the origin. 70

4.5 Ratio comparison between the two methods of measuring Ion Feedback (IFB) rate. The x-axis plots the IFB rate as measured with the Continuous Wave (CW) diode. The y-axis plots the ratio between the sum of all strobe scan IFB rates in a set of runs and the CW diode's IFB rate measurement. The strobe scan measurements set has delays ranging from 100∼400 ns between the laser pulse and the data acquisition trigger. 71

4.6 Light from a glowing HPD, visually observed in a dark room. The image is superimposed with a photograph of an HPD with external lighting to show the location of the light source. The image has been colour edited for clarity. This figure comes from [31]. 72

4.7 Correlation between two measurements of ion feedback rate using the continuous wave light source. The runs are three and a half months apart and have 3 million events each. The data has not been corrected for dark count. The pink line is $y = x$ and is drawn to guide the eye. The eight HPDs glowing in 5 March 2008 are circled in red. 73

4.8 Examples of ion feedback (IFB) development timelines from RICH2 Hybrid Photon Detectors (HPDs). These IFB rates plotted on the y-axis were measured with the continous wave light source without dark count correction. The age of each HPD plotted on the x-axis is measured in days since they were tested at the photon detector test facilities. The top, middle and bottom pair of plots are sample HPDs with low, medium and high rates of IFB respectively. A linear fit has been made, with χ^2 values printed at the top right of each plot. Note the different y-axis scales. 76

4.9 Ion feedback (IFB) development timeline of a sample Hybrid Photon Detector (HPD) showing gradual reduction in its rate of IFB. These IFB measurements plotted on the y-axis were measured with the continous wave light source and have not been corrected for dark count. The age of each HPD plotted on the x-axis is measured in days since they underwent the standard tests at the photon detector test facilities. . . 78

4.10 Examples of ion feedback (IFB) development timelines from glowing Hybrid Photon Detectors (HPDs). These IFB measurements plotted on the y-axis used the continous wave light source and have not been corrected for dark count. The age of each HPD plotted on the x-axis is measured in days since they underwent the standard tests at the photon detector test facilities. The left plot is for an HPD that was glowing from the start of IFB monitoring. The HPD in the right plot started glowing halfway through the timeline data points. 79

4.11 Correlation between linear fits of HPD timeline data and the Ion Feedback (IFB) measurement done at the Photon Detector Test Facilities (PDTF). The x-axis plots the rate of IFB as measured at the PDTF during standard tests. The y-axis plots the intercept c from $y = mx + c$ linear fits on HPD timeline data points. c is an estimate of what the rate of IFB should have been at the time of the PDTF measurement according to the linear model. Not all data points on this graph are shown as the axes are zoomed in to show the low IFB range where most of the data are. 80

4.12 Examples of ion feedback (IFB) development timeline plots incorporating the IFB measurements from the photon detector test facilities. These IFB measurements plotted on the y-axis were measured with the continous wave light source and have not been corrected for dark count. The age of each HPD plotted on the x-axis is measured in days since they underwent the standard tests at the photon detector test facilities. 81

4.13 Distributions of linear fit parameters for ion feedback (IFB) development timelines for RICH2 Hybrid Photon Detectors (HPDs). The x-axis and y-axis plot parameters from fitting without and with the PDTF measurement of rate of IFB respectively. The top and bottom pair of plots show the distribution of the gradients and intercepts found from linear fits to timeline plots of all RICH2 HPDs respectively. The right side plots of each pair is a zoomed up version of the original plot shown on the left. 83

4.14 Predictions of how many HPDs in RICH2 will start to be at risk of glowing in each year. These predictions were made using the linear fits to the timeline plots. 85

4.15 Rate of ion feedback (IFB) against Hybrid Photon Detector (HPD) age. These IFB measurements used the continous wave light source and have not been corrected for dark count. The age of each HPD is measured in days since they underwent the standard tests at the photon detector test facilities. The top plot shows IFB data of RICH2 HPDs at the end of November 2007. The bottom plot shows IFB data almost seven months later of both RICH1 and RICH2 HPDs, drawn in red and blue respectively. 86

5.1 An example $B_s \rightarrow \phi\phi$ decay event with various possible methods to deduce the initial flavour of the signal B-meson. Diagram based on [33]. 91

5.2 Data on the number of 'hard', inelastic proton-proton collisions per event, N_{hard} vs luminosity. Four different luminosities were simulated. These events come from signal datasets so there will always be at least one hard collision per event. 95

5.3 Flavour tagging performance vs number of hard collisions per event, N_{hard}. The events are from different luminosity datasets in the minimal upgrade layout, where the $B_s \rightarrow \phi\phi$ was successfully selected. The x-axis is first divided up into N_{hard}. Within each N_{hard} category, events are further sorted by which luminosity dataset they came from: 2, 5, 10 or 20 × $10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Datasets with an insignificant number of events in a given N_{hard} category are not plotted. 96

5.4 Selection efficiency vs luminosity. A control dataset with no changes to the existing detector layout is plotted in red. The inset plot zooms into the two 2 × $10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ data points. 97

5.5 Momentum distribution of all Monte Carlo truth-confirmed kaon tracks from 30k events where the $B_s \rightarrow \phi\phi$ was successfully reconstructed. The blue (red) curve is from the nominal (upgrade) luminosity dataset. The results have been filtered to only show tracks with information from all tracking stations. Note the logarithmic y-axis. 100

5.6 Kaon Particle Identification (PID) performance vs track momentum at nominal and upgrade luminosity. All tracks are plotted from the $B_s \to \phi\phi$ signal datasets in the minimal upgrade layout. All tracks have transverse momentum, p_T >500 MeV and have left a hit on at least one RICH radiator. Tracks which the PID algorithm could not make a decision on have been filtered out. The blue and green curves are the PID efficiency of correctly identifying a true kaon as a kaon at nominal and upgrade luminosity respectively. The red and purple curves are respectively the 2 and $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ rate of misidentifying a true pion as a kaon. 100

5.7 $\pi \to K$ mis-identification rate vs kaon particle identification efficiency, for various luminosities and $\Delta \ln \mathcal{L}(K - \pi)$ (DLL) cuts. The DLL cuts specify how certain the PID likelihood analysis should be to label a track as a kaon instead of a pion, with DLL > 0 being the default. These are all Monte Carlo truth-matched kaon and pion tracks from events where the $B_s \to \phi\phi$ was successfully reconstructed and selected by the LHCb analysis software. The different data points on each luminosity curve shows results from different DLL cuts, ranging from very loose at the top right to very tight on the bottom left. The y-axis (x-axis) is the fraction of true pions (kaons) satisfying the DLL cut. Note the y-axis is logarithmic. 101

5.8 Flavour tagging performance vs luminosity with results from 500k Monte Carlo signal events. The vertical line separates simulations in existing detector geometry (left) and in minimal upgraded layout (right). Luminosities of 2, 5, 10 and $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ were simulated in the upgrade layout. This is the performance of the combined tagging system that takes into account input from all applicable taggers. . . . 103

5.9 Flavour tagging performance of individual taggers vs luminosity, using 500k Monte Carlo signal events. The vertical line separates simulations in existing detector geometry (left) and in minimal upgraded layout (right). Luminosities of 2, 5, 10 and $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ were simulated in the upgrade layout. 105

5.10 Selection efficiency of $B_s \rightarrow \phi\phi$ when aerogel is removed. The left pair of data points are at $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and the right pair are the 20 $\times\ 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ luminosity results. Each pair of data points compares simulations with aerogel in (left data points) vs aerogel removed (right data points) in otherwise identical MC signal datasets in the minimal upgrade layout. The inset plots show the data zoomed up to clearly show the error bars. 108

5.11 Kaon particle identification $\Delta \ln \mathcal{L}$ distributions. The data are from events where the $B_s \rightarrow \phi\phi$ was selected from the Monte Carlo (MC) signal datasets at nominal luminosity in the minimal upgrade layout. Only MC truth matched, long-type kaon tracks are displayed, both signal and background. Blue plots have aerogel in and red plots are simulations without aerogel. (a) is the $\Delta \ln \mathcal{L}(K - \pi)$ distribution. (b) is the $\Delta \ln \mathcal{L}(K - \mathrm{p})$ distribution and (c) is a zoomed up version of the same plot. 109

5.12 Tagging performance in different environments using signal, minimal upgrade layout datasets. The left two sets of data points are at nominal luminosity and the right two are at upgrade luminosity. For each pair of sets, the left data points are when aerogel is in the simulations and the right data points have removed aerogel radiator. 111

5.13 Momentum distribution of pre-selected candidate tracks for the Other Side Kaon (OS K) tagger. These tracks are from events at nominal luminosity in the minimal upgrade layout where the $B_s \rightarrow \phi\phi$ has been selected. Results with aerogel are plotted in blue and results in the absence of aerogel are in red. 112

5.14 $\Delta \ln \mathcal{L}(K - \mathrm{p})$ distribution for candidate tracks of the other side kaon tagger. The events are from nominal luminosity Monte Carlo datasets in the minimal upgrade layout. The tracks have passed pre-selection and the OS K tagger's default $\Delta \ln \mathcal{L}(K - \pi) > 0$ cut. The blue curve is from the dataset with aerogel intact and the red distribution was simulated without aerogel. 113

6.1 Distributions of the number of tagger candidate tracks passing preselection in each event. The samples come from simulated signal datasets at 2 and $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ in the minimal upgrade layout with aerogel , plotted in blue and red respectively. Only the first 40k events where the $B_s \to \phi\phi$ was successfully reconstructed from each sample are shown. 117

6.2 Distributions of number of tagging candidate tracks that have passed the selection cuts of a kaon tagger. These tracks come from the first 40k $B_s \to \phi\phi$ selected events in minimal upgrade layout datasets. Nominal luminosity data is plotted in blue and red for retaining and removing aerogel respectively. Green and purple represent retaining and removing aerogel at upgrade luminosity. Note the logarithmic y-axes. 118

6.3 Effective tagging efficiency for different methods of selecting the kaon tagging track when many get through the tagger selection cuts. The different methods are labelled as follows: 0 for highest impact parameter pile-up (opposite side kaon only); 1 for highest $\Delta \ln \mathcal{L}(K-\pi)$; 2 for highest $\Delta \ln \mathcal{L}(K-\mathrm{p})$; 3 for highest momentum; 4 for highest transverse momentum and 5 for lowest $|\mathrm{IP}_{\mathrm{Sig}}|$ (same side kaon only). 119

6.4 Distributions of track candidates for the opposite side kaon tagger. These candidates have passed all the selection criteria and the tagger will pick one of them to make a tagging decision with. The data come from the first 40k $B_s \to \phi\phi$ selected events from minimal upgrade layout datasets with aerogel at nominal and upgrade luminosities plotted on the left and right respectively. The top two plots compare how many such candidates are in each event against how many of them would give the right tagging answer if selected by the tagger. The bottom two plots show the fraction of candidates giving the correct tagging answer against the number of selected candidates in each event. 120

6.5 Probability for the opposite side kaon tagger to tag the signal B-meson correctly for different numbers of candidate tracks passing all tagger selection cuts. The data comes from the first 40k $B_s \to \phi\phi$ selected events in the minimal upgrade layout with aerogel intact, for nominal and upgrade luminosity. The red dotted lines are the threshold where events were considered untagged if there were too many candidates passing the tagger selection cuts. 121

6.6 Opposite side kaon tagger performance when its $\Delta \ln \mathcal{L}(K - \pi)$ cut is varied. All other selection cuts were unchanged. The data comes from 480k $B_s \to \phi\phi$ signal events from minimal upgrade layout datasets. The top two plots show nominal luminosity results and the bottom two plots are for upgrade luminosity. The left two plots show tagging efficiency, ϵ_{tag} and mistag rage ω. The right two plots show effective efficiency, ϵ_{eff}. All plots show results with aerogel intact in blue and with aerogel removed in red. 122

6.7 Opposite side kaon tagger performance when its $\Delta \ln \mathcal{L}(K - \mathrm{p})$ cut is varied. All other selection cuts were unchanged. The data comes from 480k $B_s \to \phi\phi$ signal events from minimal upgrade layout datasets. The top two plots show nominal luminosity results and the bottom two plots are for upgrade luminosity. The left two plots show tagging efficiency, ϵ_{tag} and mistag rage ω. The right two plots show effective efficiency, ϵ_{eff}. All plots show results with aerogel intact in blue and with aerogel removed in red. 123

6.8 Effective tagging efficiency, ϵ_{eff} of the kaon taggers during the recursive reoptimisation of their tagger candidate selection cuts. ϵ_{eff} values shown were the highest found at the end of each reoptimisation round. The same side kaon tagger is plotted in blue and the opposite side kaon tagger is in red. The data is from a minimal upgrade layout $B_s \to \phi\phi$ signal dataset at upgrade luminosity and with aerogel removed. Round 0 is the original value of ϵ_{eff} before any reoptimisation. 128

6.9 Values of effective tagging efficiency, ϵ_{eff} obtained from the recursive reoptimisation overtraining check. The dataset was split into two halves: a test and a control sample, plotted in blue and red respectively. The blue data points show the highest ϵ_{eff} found at the end of each reoptimisation round. The red data points are the ϵ_{eff} obtained from using the cut values at the end of the test sample's reoptimisation rounds. The first data points at round 0 is the ϵ_{eff} of each sample without any reoptimisation. 128

6.10 Opposite side kaon tagger reoptimised performance when its $\Delta \ln \mathcal{L}(K-\pi)$ cut is varied. The data comes from 480k $B_s \to \phi\phi$ signal events from minimal upgrade layout datasets at upgrade luminosity. The left two plots show tagging efficiency, ϵ_{tag} and mistag rage ω. The right two plots show effective efficiency, ϵ_{eff}. The performance before and after reoptimisation in the top and bottom pair of plots respectively. All plots show results with aerogel intact in blue and with aerogel removed in red. The dotted lines show where the $\Delta \ln \mathcal{L}(K-\pi)$ cut is applied. After reoptimisation, the cut value is different for whether aerogel is present or not. 131

6.11 Opposite side kaon tagger reoptimised performance when its $\Delta \ln \mathcal{L}(K-\mathrm{p})$ cut is varied. The data comes from 480k $B_s \to \phi\phi$ signal events from minimal upgrade layout datasets at upgrade luminosity. The left two plots show tagging efficiency, ϵ_{tag} and mistag rage ω. The right two plots show effective efficiency, ϵ_{eff}. The performance before and after reoptimisation in the top and bottom pair of plots respectively. All plots show results with aerogel intact in blue and with aerogel removed in red. The dotted lines show where the $\Delta \ln \mathcal{L}(K-\mathrm{p})$ cut is applied. After reoptimisation, the cut value is different for whether aerogel is present or not. 132

6.12 Comparison of effective tagging efficiency for the kaon taggers before and after their reoptimisation, plotted in blue and red respectively. The recursive reoptimisation was applied to 480k events from minimal upgrade layout $B_s \to \phi\phi$ signal datasets. 133

List of Tables

1.1 Quarks in the Standard Model, listed with charge (relative to the magnitude of an electron's charge) and the latest mass measurements [4]. 3

1.2 Leptons in the Standard Model (SM), listed with charge (relative to the magnitude of an electron's charge) and mass [4]. 3

1.3 Gauge bosons, the force propagators in the Standard Model [4]. Gluons have colour instead of charge. Due to the finite mass of its bosons, the weak force has limited range. 4

2.1 Table of selected Silicon Tracker design parameters 26

3.1 Table of selected HPD specifications 45

4.1 Rate of Ion Feedback (IFB) as measured at the end of March 2008 for the five Hybrid PhotonDetectors (HPDs) that began to glow two months later. The IFB rate was measured with the continuous wave diode and is not corrected for noise. 74

4.2 Hybrid PhotonDetectors (HPDs) that began to glow at the end of 2008. Also shown are the dates that the timeline linear fits predicted these HPDs would be at risk according to the glow threshold hypothesis. 84

5.1 List of selection cuts for the $\mathrm{B_s} \to \phi\phi$ channel. 98

5.2 Cuts used by each single particle tagger to select a track to make a tagging decision with. Track candidates have been pre-selected to remove tracks in the signal channel decay. dQ is the mass difference between the track and the B mass. LCS is the track's χ^2 per degree of freedom. Q_{VELO} is the dE/dx charge from the VELO system. 104

6.1 Cuts applied in the secondary pre-selection tailored for each kaon tagger. LCS is track χ^2 per degree of freedom. 126

6.2 Table of how the same side kaon tagger's candidate track selection cuts changed throughout six rounds of recursive reoptimisation. The reoptimisation was performed on a minimal upgrade layout, $B_s \to \phi\phi$ selected signal dataset at upgrade luminosity without aerogel. These cut values are at the beginning of each round. LCS is track χ^2 per degree of freedom. 127

6.3 The reoptimised tagger selection cuts after the recursive iterations finished. Each luminosity and aerogel condition was separately optimised for. LCS is track χ^2 per degree of freedom. 130

Chapter 1

Theoretical Background and Motivation

1.1 Introduction

Above the question “how did life originate?” arguably comes the question “why was there any matter left for life to be formed from?”. The Universe started with the Big Bang which created equal amounts of matter and anti-matter. Yet humans have only known of anti-matter’s existence for less than a century. That awareness is due to anti-matter being created for fleeting moments, not from an observation of a primoridal quantity that survived since time began. So the question remains: why did matter and anti-matter not annhilate each other with no trace left if they were created equally?

In 1967, Andrei Sakharov proposed three conditions [1] for a process to reproduce this unbalance of matter versus anti-matter. One condition is CP violation, first observed in 1964 in K-mesons [2]. The phenomenon occurs at much larger rates in B-mesons. Dedicated particle accelerators (so called ‘B factories’) were built to study this. The amount of CP violation observed so far is several orders of magnitude below what’s required for the amount of matter in the universe today to have survived after all the anti-matter was annhilated [3].

In particle physics, large data samples are necessary to accumulate enough statistics to reveal any new discoveries with certainty. Only a fraction of events produce B-mesons and many decay channels of interest to CP violation studies are rare decays of B-mesons. Decay channels benefit from more statistics as this leads to better, more precise measurements of physics parameters. Therefore particle accelerator experiments strive for faster rates of data collection to increase data sample sizes.

The LHCb experiment at the Large Hadron Collider (LHC) is the latest B-physics oriented particle detector. It is currently being calibrated with real beam collision data. LHCb will collect a much larger data sample than previous B factory experiments. Even as the LHCb begins to collect data, work is being done to prepare for an upgrade to the detector so that data can be collected at a faster rate in the future.

The remainder of this chapter will introduce the theoretical framework relevant to the work presented in this thesis. The early sections of this chapter aims to be accessible to a broad audience, while increasingly specialised knowledge is assumed later in the chapter. Chapter 2 describes the LHCb detector and gives an overview of the plans for the LHCb upgrade. The LHCb experiment uses Ring Imaging Cherenkov (RICH) counters for particle identification. Chapter 3 presents quantum efficiency characterisation of the RICH photodetectors. Chapter 4 describes the monitoring of vacuum quality in RICH2 photodetectors. LHCb analysis software will need to be reoptimised when LHCb is upgraded. Chapter 5 describes the flavour tagging algorithms used at present and how they perform on simulated $B_s \to \phi\phi$ events under upgrade conditions. Chapter 6 presents how flavour tagging performance changes when reoptimised for upgrade conditions.

1.2 Standard Model

The Standard Model (SM) is the suite of theories that describes the foundations of matter and anti-matter and how they interact. These fundamental building blocks are divided up into leptons and quarks. The forces that act between and within these two categories of building blocks are mediated by gauge bosons. Gravity is not included in the SM, so clearly it is not a complete 'Theory of Everything' but it is the closest attempt with very successful experimental confirmation and proven predictions.

1.2.1 Constituent Particles

Tables 1.1 and 1.2 list the quarks and leptons of the SM, both of which are organised into doublets (pairs) called 'generations'. All of them are spin $^1/_2$

Generation	Name	Charge/$\lvert e \rvert$	Mass
1st	u (up)	$+\frac{2}{3}$	1.7 to 3.3 MeVc^{-2}
	d (down)	$-\frac{1}{3}$	4.1 to 5.8 MeVc^{-2}
2nd	c (charm)	$+\frac{2}{3}$	$1.27^{+0.07}_{-0.09}$ GeVc^{-2}
	s (strange)	$-\frac{1}{3}$	101^{+29}_{-21} MeVc^{-2}
3rd	t (top)	$+\frac{2}{3}$	$172.0 \pm 0.9 \pm 1.3$ GeVc^{-2}
	b (bottom)	$-\frac{1}{3}$	$4.19^{+0.18}_{-0.06}$ GeVc^{-2}

Table 1.1: Quarks in the Standard Model, listed with charge (relative to the magnitude of an electron's charge) and the latest mass measurements [4].

Generation	Name	Charge/$\lvert e \rvert$	Mass / MeVc^{-2}
1st	e (electron)	-1	0.51
	ν_e	0^{SM}	0^{SM}
2nd	μ (muon)	-1	105.66
	ν_μ	0^{SM}	0^{SM}
3rd	τ (tau)	-1	1776.82 ± 0.16
	ν_τ	0^{SM}	0^{SM}

Table 1.2: Leptons in the Standard Model (SM), listed with charge (relative to the magnitude of an electron's charge) and mass [4].

fermions. Theoretically there is no reason why there are only three generations, though the Ward identity [5] shows the number of quark and lepton generations are the same[1]. Experimentally, only three generations have been observed. The 'flavour' of a quark is what type it is: whether it is a u, s or c-quark, etc. The neutrinos, electrons and 1st generation quarks are stable, all other particles are short-lived, e.g. the μ has a mean lifetime of 2.197×10^{-6} s. All these particles have anti-particle counterparts of opposite charge.

Table 1.3 lists the gauge bosons which propagate force in the SM[2]. The interaction between these spin 1 bosons and particles dictates how the fundamental forces behave. Photons mediate the Electromagnetic (EM) force between all charged particles. The strong force uses gluons, which interact with the quantum

---

[1] This is due to the so-called triangle anomaly, in which a Feynman diagram has a triangular loop from a virtual charged fermion. This would result in a divergence. However the contributions from all possible elementary charged fermions cancel each other out when summed over. As the amplitude is proportional to the electrical charge, the sum of charges of leptons and all colours of quarks from all generations.

[2] Neutrinos are massless and chargeless in the SM. However, neutrinos have been observed to oscillate in flavour. This can only happen if neutrinos have mass.

Force	Name	Charge/$\lvert e \rvert$	Mass / GeVc^{-2}
Electromagnetic	γ (photon)	0	0
Weak	$W^{\pm}$	± 1	80.399 ± 0.023
	Z^0	0	91.1876 ± 0.0021
Strong	g (gluon)	0	0

Table 1.3: Gauge bosons, the force propagators in the Standard Model [4]. Gluons have colour instead of charge. Due to the finite mass of its bosons, the weak force has limited range.

number ‘colour’ instead of electric charge. Only quarks and no leptons are subject to the strong force but gluons can interact with each other as well as quarks. The $W^{\pm}$ and Z^0 bosons govern the weak force. This is the only force that can interact with neutrinos in addition to charged leptons and quarks. $W^{\pm}$ bosons only act on particles with “left-handed” helicity and “right-handed” anti-particles. Helicity is discussed in Section 1.4.

1.2.2 The Strong Force

‘Colour’, the charge of the strong force was originally suggested as an extra quantum number to explain particles such as the Δ^{++} (uuu) or the Ω^{-} (sss). Since all quarks are fermions, they obey the Pauli Exclusion Principle: no two fermions may share the exact same quantum state. For the well-known case of atomic electrons occupying different shells in atomic orbits, two electrons can occupy the same subshell on a given atom because their quantum spin can have two discrete values. For three quarks to occupy the same space with the same flavour, there must be another quantum number - another property of the state - that can keep them different. This was labelled ‘colour’, which can take three values: red, green and blue (r, g, b). These values have anti-colour counterparts ($\bar{\mathrm{r}}, \bar{\mathrm{g}}, \bar{\mathrm{b}}$) for anti-quarks, just as a positron has the opposite electrical charge to an electron.

Individual, free quarks or gluons are never observed because of the strong force. Quarks are always confined inside hadrons, with a net colour of zero (‘white’): quark-antiquark pairings of colour + anti-colour are called mesons and 3-quark groupings of the three different colours are called baryons.

When quarks or gluons fly apart, the field lines of the strong force build up potential energy, like tension in a rubber string. New quarks are formed from this built up energy, which travel in similar directions to the original quarks or gluons. This leads to further strong force potential energy from which more quarks are created. Thus instead of individual quarks or gluons, we observe jets of hadrons trailing after the initial quark. The u and d-quarks form the common proton (uud) and neutron (udd). These become the building blocks for nuclei that electrons orbit around to make what we call atoms. The strong force within a hadron is not felt by other hadrons until they are within $\sim 10^{-15}\,\mathrm{m}$ due to screening from virtual gluons. Despite its limited range, this is the force that prevents a nucleus made only from positive protons and neutrons from repelling itself apart through the electromagnetic force.

1.2.3 The Electroweak Force

The Electromagnetic (EM) force is mediated by photons. Quarks and charged leptons can emit and absorb virtual photons to change their momentum and energy. These photons are off-shell, i.e. they do not obey $E^2 = |\vec{p}|^2c^2 + m^2c^4$. Heisenberg's Uncertainty Principle dictates how long off-shell particles can exist for in relation to how far off they are from the mass shell: $\Delta E \Delta t \sim \hbar$. The EM force has almost infinite range as virtual photons can be very close to the mass shell, i.e. $E \sim p \sim 0$ which results in a very long lifetime.

The weak force has both charged and neutral bosons. An interaction with $W^{\pm}$ bosons is called a Charged Current (CC) whereas when the Z^0 is involved it is called a Neutral Current (NC) process. Unlike the other forces, the weak force bosons have mass which means they have a finite range of $\mathcal{O}(10^{-18}\,\mathrm{m})$. Decays via the weak interaction have long lifetimes and small cross-sections. To first order (so called 'tree-level'), only CC interactions can change quark flavour. Since the W bosons have electric charge they couple with the EM propagator, the photon.

At first, the weak force description only contained the $W^{\pm}$ bosons and was valid only at low energies. To extend the description to higher energies, the NC was proposed and later discovered. Furthermore, the model which proposed the existence of the Z^0 boson also predicted that it mixed with the photon - that the electromagnetic force and weak force were different manifestations of a unified electroweak interaction. This too was proven and the mass of the Z^0 boson was correctly predicted by the model and precisely measured at the Large Electron-Positron (LEP) collider.

1.3 CKM Mechanism

The weak interaction can change quark flavour, usually staying within the same quark generation but occasionally a quark can jump between generations. This is "flavour mixing", which occurs because the flavour eigenstates that the Electroweak bosons interact with are not the same as the mass eigenstates. A quark property's eigenstate is a quantum state of the quark with a definite value for that property, which the quark's wavefunction collapses to when a measurement is made of that property. An eigenstate is only definite for a single property and the wavefunction will evolve into an uncertain state again.

In the SM, flavour mixing across all three quark generations is modelled by the Cabbibo-Kobayashi-Maskawa (CKM) matrix [6] as follows:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \hat{V}_{\mathrm{CKM}} \begin{pmatrix} d \\ s \\ b \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \tag{1.1}$$

where q' refers to flavour eigenstates and q refers to mass eigenstates of the quarks. The CKM matrix is a 3×3 unitary matrix mapping the mass eigenstates to flavour eigenstates. Each matrix element relates to the probability of a flavour changing process:

$$P(i \to j) = |V_{ij}|^2 \tag{1.2}$$

where i is one of the up-type quark flavours and j is one of the down-type flavours. The CKM matrix is unitary in the SM. The diagonal elements of the CKM matrix are close to 1 and the other elements get smaller the further away from the diagonal they are. This represents the relative ease to change flavours within a generation compared to changing generations. The Wolfenstein parameterisation [7] is an approximation of the CKM matrix which clearly shows this hierarchy:

$$\hat{V}_{\mathrm{CKM}} = \begin{pmatrix} 1 - \frac{1}{2}\lambda^2 & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \frac{1}{2}\lambda^2 & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix} + \mathcal{O}(\lambda^4) \tag{1.3}$$

where $\lambda = 0.22 \simeq |V_{us}|$ and ρ, A, η are real parameters of order 1. Note that only the η term is multiplied by the imaginary unit of i. η parameterises the

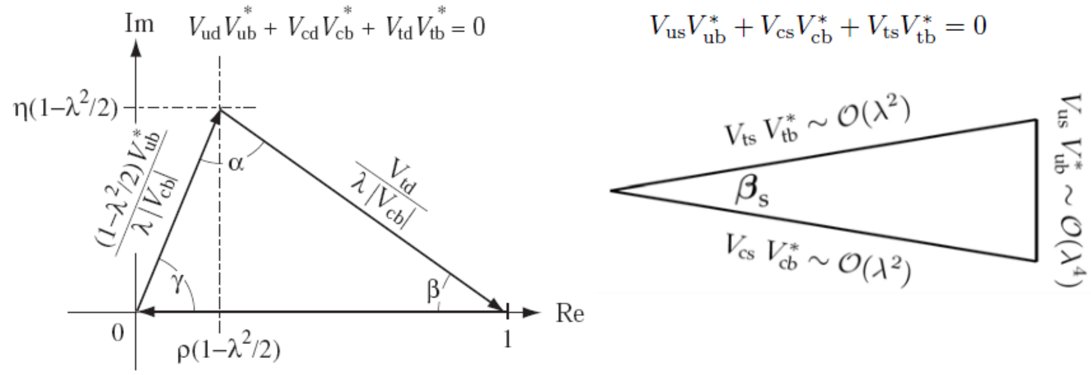

Figure 1.1: Two of the six Unitarity Triangles of the SM CKM matrix. The lengths are based on the Wolfenstein parameterisation of each triangle's orthogonality relations. Lengths labelled on the left triangle have taken out the common factor of $A\lambda^3$.

complex phase of the CKM matrix which accommodates CP violation in the SM. CP violation will be discussed in Section 1.4.

The unitarity of the CKM matrix - i.e. its property of $VV^\dagger = I$, the identity matrix - means its three rows and three columns are orthogonal, leading to a set of six orthogonality relations. Of major interest are the following two:

$$\begin{aligned} &V_{ud}V_{ub}^* + V_{cd}V_{cb}^* + V_{td}V_{tb}^* = 0 \\ =&(\rho + i\eta)A\lambda^3 - A\lambda^3 + (1 - \rho - i\eta)A\lambda^3 \end{aligned} \tag{1.4}$$

from the column orthogonality relations and:

$$\begin{aligned} &V_{ud}^*V_{td} + V_{us}^*V_{ts} + V_{ub}^*V_{tb} = 0 \\ =&(1 - \rho - i\eta)A\lambda^3 - A\lambda^3 + (\rho + i\eta)A\lambda^3 \end{aligned} \tag{1.5}$$

from the rows. In both cases the three terms are roughly equal in magnitude at $\mathcal{O}(\lambda^3)$. These orthogonality relations can be plotted on as triangles in the complex plane using the Wolfenstein parameterisation. Each term is taken as a vector describing a length of a triangle, examples of which are shown in Figure 1.1. The relations demand these three vectors return to the origin when added together. While the other orthogonality relations besides the above two can also be plotted as triangles, in each case one of the terms is much shorter than the others resulting in very 'squashed' triangles, as shown in Figure 1.1 (right).

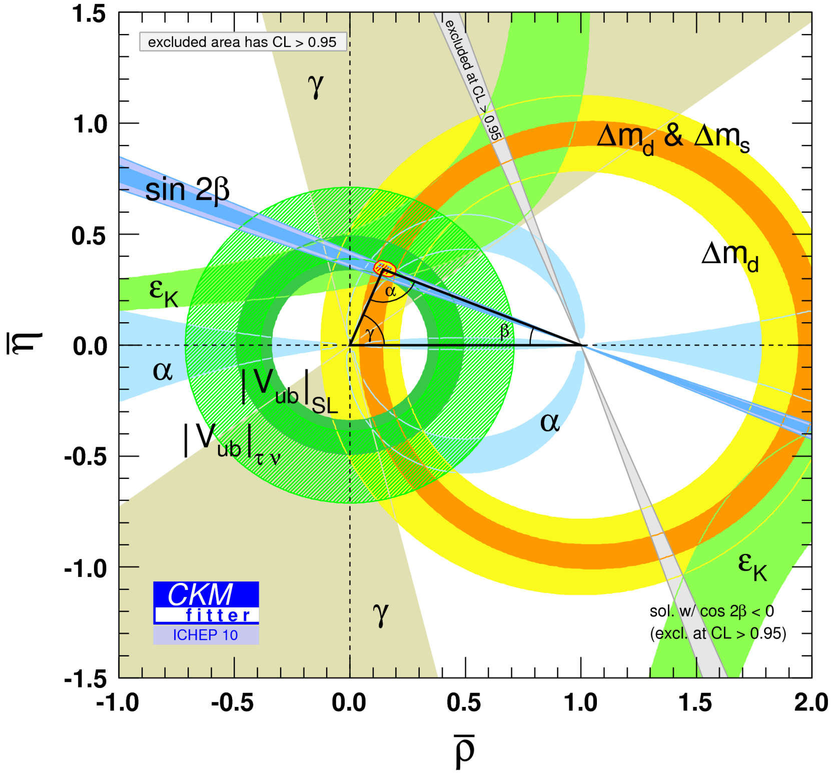

Figure 1.2: Latest results (2010) from the CKM fitter group [8]. The red shaded area represents the 68% CL region where the apex could be.

This assertion of unitarity in the CKM matrix is being experimentally tested. New Physics (NP) beyond the SM could change the nature of this mixing matrix. The current results from over-constraining the CKM unitarity triangle is presented in Figure 1.2. Often these efforts are discussed in terms of measuring the angles of the unitarity triangles:

$$\begin{aligned}
\alpha &= arg(-\frac{V_{td}V_{tb}^*}{V_{ud}V_{ub}^*}) \\
\beta &= arg(-\frac{V_{cd}V_{cb}^*}{V_{td}V_{tb}^*}) \\
\gamma &= arg(-\frac{V_{ud}V_{ub}^*}{V_{cd}V_{cb}^*}) \\
\beta_s &= arg(-\frac{V_{ts}V_{tb}^*}{V_{cs}V_{cb}^*}) \qquad (1.6)
\end{aligned}$$

where the first three angles are shown on Figure 1.1 (left). β_s is the angle relevant for B_s^0 decay channels and is shown in Figure 1.1 (right).

1.4 C, P and CP Violation

The charge-conjugation operator (C) replaces particles for their anti-particle counterparts and vice versa. The strong and EM forces make no distinction after the operation - they are symmetric with respect to this operator. However, the charged weak force only interacts with particles of left-handed helicity and anti-particles of right-handed helicity. The handedness of a particle refers to $\vec{S} \cdot \hat{p}$, the projection of its spin operator, $\vec{S}$ in the direction of its momentum, $\vec{p}$: left-handed particles project the spin in the opposite direction to the momentum. Because of this distinction, the charged weak force maximally violates C-symmetry since a left-handed particle becomes a left-handed anti-particle after the C operator acts on it.

The parity operator (P) flips all spatial coordinates about the origin: $x, y, z \rightarrow -x, -y, -z$. This has the effect of reversing all polar vectors describing the system such as momentum and acceleration, but not axial vectors such as angular momenta and spin. The weak interaction also violates P-symmetry while all other forces preserve it, first observed in 1957 [9].

For a time it was believed that the combined act of both the C and the P operator would have a symmetry even the weak interaction would respect. A left-handed particle would become a right-handed anti-particle, for example:

$$CP|\nu_{L,R}\rangle \rightarrow |\overline{\nu}_{R,L}\rangle$$

The weak interaction should make no distinction of this combined change. However, in 1964 CP-symmetry was found to be violated in neutral Kaon mesons [2]. These mesons are observed in two distinct varieties: a short-lived form, K_S and long-lived form, K_L. Under CP-symmetry, only the K_S mesons should decay to $\pi^+\pi^-$ but 1 in 500 K_L decayed into the same final states. This was because the two CP eigenstates of the neutral Kaon are different from the two lifetime eigenstates that the weak force interacts with. Each lifetime eigenstate is a mixture of the two CP eigenstates. This phenomenon is further described in Section 1.5.2.

1.5 CP Violation in the B-meson system

At present, the largest sources of CP violation in particle physics experiments are found in the B-meson system. B-mesons contain a $\bar{b}$-quark (the anti-particle) and a light quark (u,d or s). $\overline{B}$-mesons have a b-quark and a light anti-quark. Formally, these are flavour eigenstates: they are particle states with flavour content definitively known. They shall be labeled as $|B\rangle$ and $|\overline{B}\rangle$ respectively.

Neutral B-mesons can exhibit CP violation in three ways, described below. Quantum mechanical descriptions involve complex phases which are completely arbitrary in general. Any ket state may be rephased with no measurable consequence. Therefore the observables of CP violation must be invariant under rephasings. Note that the theoretical framework in the following descriptions is applicable to other neutral meson systems too.

1.5.1 CP Violation in Decay

Consider the case $B \to f$ where a $|B\rangle$ state decays to a many-particle final state f. The CP conjugate of this decay is $\overline{B} \to \overline{f}$. Their amplitudes are defined as:

$$\begin{aligned} A_f &\equiv \langle f|\mathcal{H}|B\rangle, \\ \overline{A}_{\overline{f}} &\equiv \langle \overline{f}|\mathcal{H}|\overline{B}\rangle \end{aligned} \tag{1.7}$$

where $\mathcal{H}$ is the Hamiltonian governing weak interactions. The decay rates for these processes are proportional to the square of their amplitudes. Note that for some final states $f = \overline{f}$, such as $B \to \pi^+\pi^-$.

CP violation in decay, also known as direct CP violation, arises when:

$$|A_f| \neq |\overline{A}_{\overline{f}}| \tag{1.8}$$

For example, the current measurement of direct CP violation in the channel $\overline{B}^0 \to K^-\pi^+$ is [4]:

$$\begin{aligned} \mathcal{A}_{K^\mp\pi^\pm} &= \frac{\Gamma(B^0 \to K^+\pi^-) - \Gamma(\overline{B}^0 \to K^-\pi^+)}{\Gamma(B^0 \to K^+\pi^-) + \Gamma(\overline{B}^0 \to K^-\pi^+)} \\ &= -0.098 \pm 0.013 \end{aligned} \tag{1.9}$$

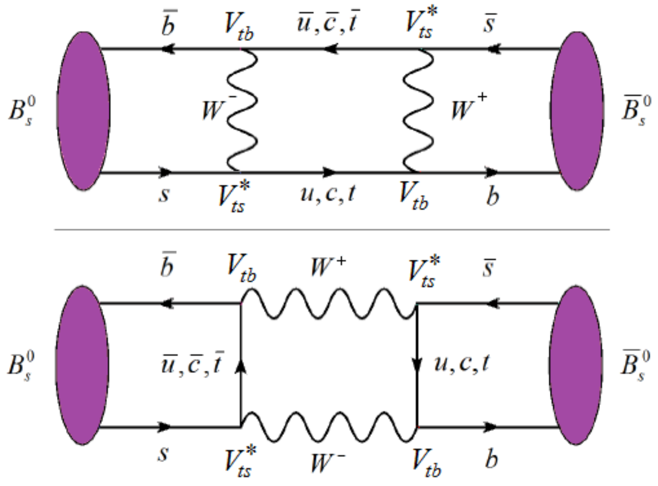

Figure 1.3: Feynman diagrams of B_s^0-$\overline{B}_s^0$ mixing. The quark and anti-quark of the meson swap flavours by exchanging charged W bosons and temporarily becoming an up-type flavour. The dominant CKM elements for the vertices have been labelled.

where Γ is the partial decay width. For charged hadrons, this is the only source of CP violation.

1.5.2 CP Violation in Mixing

Flavour eigenstates of neutral mesons are not physical eigenstates. The eigenstates of the Hamiltonian have well defined mass and decay width. It is these mass eigenstates that propagate through space-time and not the flavour eigenstates. The flavour content of the mass eigenstates oscillates between the two flavour eigenstates as time passes. This phenomenon is known as mixing. It occurs via flavour changing neutral currents, examples of which are shown in Figure 1.3.

There are different conventions for labelling the mass eigenstates for different neutral meson systems. Section 1.4 mentioned the convention for neutral kaons: K_S and K_L are labels referring to the physical eigenstates by their lifetimes. For neutral B-mesons, the physical eigenstates are referred to by mass: B_H and B_L are the heavy and light mass eigenstates of the B-meson system. They are linear combinations of the flavour eigenstates, described by:

$$\begin{aligned} |B_L^0\rangle &= p|B^0\rangle + q|\overline{B}^0\rangle \\ |B_H^0\rangle &= p|B^0\rangle - q|\overline{B}^0\rangle \end{aligned} \tag{1.10}$$

where p and q are complex coefficients normalised such that:

$$|p|^2 + |q|^2 = 1 \tag{1.11}$$

An initially pure flavour eigenstate will evolve with time as:

$$\begin{aligned} |B^0(t)\rangle &= g_+(t)|B^0\rangle + \frac{q}{p}g_-(t)|\overline{B}^0\rangle \\ |\overline{B}^0(t)\rangle &= g_+(t)|\overline{B}^0\rangle + \frac{p}{q}g_-(t)|B^0\rangle \end{aligned} \tag{1.12}$$

where

$$g_\pm(t) = \frac{1}{2}\left[e^{-i(m_L-\frac{i}{2}\Gamma_L)t} \pm e^{-i(m_H-\frac{i}{2}\Gamma_H)t}\right] \tag{1.13}$$

$$\Rightarrow |g_\pm(t)|^2 = \frac{e^{-\Gamma t}}{2}\left[cosh(\frac{\Delta\Gamma t}{2}) \pm cos(\Delta m t)\right] \tag{1.14}$$

expresses the time-dependent probabilities for being one or the other flavour eigenstate. Note that charged mesons cannot oscillate due to electric-charge conservation.

CP violation in mixing arises when:

$$|\frac{p}{q}| \neq 1 \tag{1.15}$$

which would introduce a bias in Equations 1.12. In other words, if Equation 1.15 is true then the rate of a B going to a $\overline{B}$ is different from the reverse process.

This type of CP violation can be measured in semileptonic decays of neutral mesons. For example, the $B^0 \rightarrow D^\mp\ \mu^\pm\ \nu_\mu$ channel. In the absence of any mixing, a μ^+ is emitted when the B^0 decays. With mixing, it is possible to have $B^0 \rightarrow \overline{B}^0 \rightarrow D^+\ \mu^-\ \nu_\mu$ and therefore the μ emitted is the 'wrong sign' compared to the initial flavour when the B meson was first created. Flavour tagging identifies the initial B meson flavour and the charge of the semileptonic decay products characterises the B meson flavour at decay. Therefore it can be discerned whether a B meson changed flavour by mixing between creation and decay. $|p/q|$ can be observed by measuring the asymmetry in the rate of these 'wrong sign' semileptonic decays:

$$\mathcal{A}_{sl}(t) = \frac{d\Gamma/dt(\overline{B}^0 \to B^0 \to D^-\mu^+\nu_\mu) - d\Gamma/dt(B^0 \to \overline{B}^0 \to D^+\mu^-\overline{\nu}_\mu)}{d\Gamma/dt(\overline{B}^0 \to B^0 \to D^-\mu^+\nu_\mu) + d\Gamma/dt(B^0 \to \overline{B}^0 \to D^+\mu^-\overline{\nu}_\mu)} \tag{1.16}$$

The SM predicts the B-meson system will have very little CP violation in mixing. Current measurements at the time of writing does not refute this.

1.5.3 CP Violation through Interference between Mixing and Decay

CP violation observables must be invariant under rephasings. So far we have used amplitudes of parameters as observables. However, it is possible to have CP violation in a process where $|p/q| = 1$ and $|A_f| = |\overline{A}_{\overline{f}}|$, i.e. there is another CP violation observable. Consider a decay channel where both B^0 and $\overline{B}^0$ decay into the same final state f, i.e. $f = \overline{f}$. Let us define the complex parameter λ_f [10]:

$$\lambda_f \equiv \frac{q}{p}\frac{\overline{A}_f}{A_f} \tag{1.17}$$

We are free to change the phase of both the B-meson initial states and f:

$$\begin{aligned} |B^0\rangle &\to e^{i\gamma}|B^0\rangle \\ |\overline{B}^0\rangle &\to e^{i\overline{\gamma}}|\overline{B}^0\rangle \\ |f\rangle &\to e^{i\gamma_f}|f\rangle \end{aligned} \tag{1.18}$$

where the γ terms are arbitrary phases with no physical meaning when considered individually. Under these phase rotations,

$$\begin{aligned} A_f &\to e^{i(\gamma-\gamma_f)}A_f \\ \overline{A}_f &\to e^{i(\overline{\gamma}-\gamma_f)}\overline{A}_f \\ \frac{q}{p} &\to e^{i(\gamma-\overline{\gamma})}\frac{q}{p} \\ \lambda_f &\to e^{i\{(\gamma-\overline{\gamma})-(\gamma-\gamma_f)+(\overline{\gamma}-\gamma_f)\}}\lambda_f = \lambda_f \end{aligned} \tag{1.19}$$

The phase of λ_f does not change no matter what the γ terms are. It is invariant under the arbitrary phase changes and therefore is an observable. A_f, $\overline{A}_f$ and $\frac{q}{p}$ are not observables since their phases depend upon the γ convention. Thus CP violation can occur through the interference between the mixing phases and decay phases. The condition for this is expressed as:

$$\begin{aligned} arg\lambda_f &\neq 0 \\ \Rightarrow \mathfrak{Im}(\lambda_f) &\neq 0 \end{aligned} \tag{1.20}$$

If true, then there will be an asymmetry between decays where the B flavour at decay is the same as the initial flavour ($B \rightarrow f$) and decays where the flavour changed from mixing before decay ($B \rightarrow \overline{B} \rightarrow f$). Using the $B \rightarrow \pi^+ \pi^-$ channel as an example, the asymmetry is measured as:

$$\mathcal{A}_{f(CP)}(t) = \frac{d\Gamma/dt(\overline{B}^0 \rightarrow \pi^+\pi^-) - d\Gamma/dt(B^0 \rightarrow \pi^+\pi^-)}{d\Gamma/dt(\overline{B}^0 \rightarrow \pi^+\pi^-) + d\Gamma/dt(B^0 \rightarrow \pi^+\pi^-)} \tag{1.21}$$

In the B_d-meson system, the approximations $\Delta\Gamma_d = 0$ and $|q/p| = 1$ can be made. This simplifies the above expression to

$$\mathcal{A}_{f(CP)}(t) = S\,sin(\Delta mt) - C\,cos(\Delta mt) \tag{1.22}$$

where

$$S \equiv \frac{2\,\mathfrak{Im}(\lambda_f)}{1+|\lambda_f|^2} \quad , \quad C \equiv \frac{1-|\lambda_f|^2}{1+|\lambda_f|^2} \tag{1.23}$$

S describes CP violation through interference between decay and mixing. C describes contributions from direct CP violation.

The measurement of CP violation can help constrain the CKM triangle. For example, the $B \rightarrow J/\psi K_S$ channel is dominated by a single tree diagram which leads to

$$\begin{aligned}
\lambda_f &= B(\text{mixing}) \cdot \text{B(decay)} \cdot \text{K(mixing)} \\
&= \frac{V_{tb}^* V_{td}}{V_{tb} V_{td}^*} \cdot \frac{V_{cb} V_{cs}^*}{V_{cb}^* V_{cs}} \cdot \frac{V_{cs} V_{cd}^*}{V_{cs}^* V_{cd}}
\end{aligned} \tag{1.24}$$

The K mixing term is necessary as the $J/\psi K_S$ final state cannot be reached without $K^0 - \overline{K}^0$ mixing after the $B \to J/\psi K^0$ decay and its CP conjugate.

Substituting Equation 1.24 into Equation 1.23,

$$S = -sin(2\beta) \quad , \quad C = 0 \tag{1.25}$$

Thus from measuring $\mathcal{A}_{f(CP)}(t)$ we get a constraint on β. Similarly, a measurement of $\mathcal{A}_{f(CP)}(t)$ in the $B_s^0 \to J/\psi\phi$ channel yields a constraint on β_s.

1.6 CP Violation in $B_s \to \phi\phi$ decays

1.6.1 Standard Model Prediction

To consider the CP violation from interference in $B_s \to \phi\phi$, let us expand the left hand side of Equation 1.20 and define some new terms:

$$\begin{aligned}
\phi_s &\equiv arg(\lambda_f) \\
\phi_M &\equiv arg(\frac{q}{p}) \\
\phi_D &\equiv arg(\frac{\overline{A}_f}{A_f}) \\
\Rightarrow \phi_s &= \phi_M + \phi_D
\end{aligned} \tag{1.26}$$

The box diagrams for $B_s^0 - \overline{B}_s^0$ mixing in the SM are shown in Figure 1.3. The t-quark contribution dominates this process due to its large mass. Thus the mixing phase, ϕ_M for B_s^0 channels can be calculated by reading off the CKM elements of the dominant box diagrams:

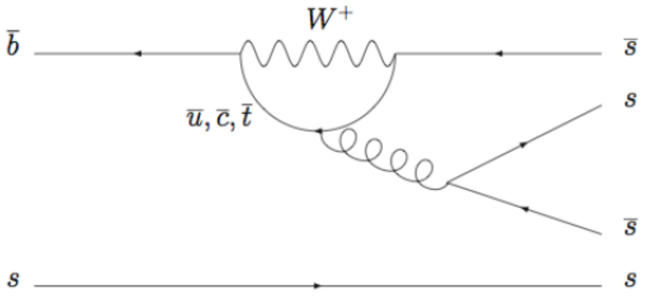

Figure 1.4: Dominant Feynman diagram of the $B_s \to \phi\phi$ decay. The tree level diagram is forbidden in the standard model because this is a flavour changing neutral current decay.

$$\phi_M = arg(V_{tb}V_{ts}^*V_{tb}V_{ts}^*) = 2\,arg(V_{tb}V_{ts}^*) \tag{1.27}$$

The dominant Feynman diagram for the $B_s \to \phi\phi$ channel in the SM is shown in Figure 1.4. As with the mixing diagrams, the t-quark contribution dominates. Looking at this diagram, the CKM contribution in the decay can be determined:

$$\begin{aligned} A_f &\propto V_{tb}V_{ts}^* \\ \overline{A}_f &\propto V_{tb}^*V_{ts} \end{aligned} \tag{1.28}$$

which leads to a decay phase, ϕ_D of:

$$\begin{aligned} \phi_D &= arg(\frac{V_{tb}^*V_{ts}}{V_{tb}V_{ts}^*}) \\ &= arg(V_{tb}^*V_{ts}V_{tb}^*V_{ts}) \\ &= -2\,arg(V_{tb}V_{ts}^*) \end{aligned} \tag{1.29}$$

Therefore in the SM, taking only t-quarks in the loop:

$$\begin{aligned} \phi_s &= \phi_M + \phi_D \\ &= 2\,arg(V_{tb}V_{ts}^*) - 2\,arg(V_{tb}V_{ts}^*) = 0 \end{aligned} \tag{1.30}$$

A full calculation at next-to-leading order approximation is given in [11], which results in $S < \sim 0.02$ where S is as defined in 1.23. Thus an observation of significant CP violation from interference in this channel is a sign of new physics beyond the standard model.

1.6.2 Experimental Prospects at LHCb

The $B_s \to \phi\phi$ channel was first discovered by the CDF experiment [12]. 8 signal candidate events were found from a data sample of $180\,\mathrm{pb}^{-1}$. This was estimated to be a 4.7 σ significant signal. A Monte Carlo toy study [13] estimates LHCb will measure ϕ_s with a statistical uncertainty of $\sigma(\phi_s) = 0.06$ rad for $10\,\mathrm{fb}^{-1}$ of data, corresponding to five years of data taking. This result was based on the recent estimate by the CDF experiment of the $B_s \to \phi\phi$ branching ratio [14].

The toy study also showed the systematic errors would be small compared to this statistical uncertainty. The irreducible theoretical error is also small compared to this statistical limit. Therefore with $10\,\mathrm{fb}^{-1}$ of data the SM will not have been tested to its precision. LHCb will be upgraded to meet this demand with a faster rate of data collection [15]. After five years of running after the upgrade, a data sample of about 400k $B_s \to \phi\phi$ events is expected. This improves $\sigma(\phi_s)$ to the order of the irreducible theoretical uncertainty.

Chapter 2

The LHCb Detector

2.1 Overview

The LHCb experiment is for precision studies of B decays. Figure 2.1 is a side view schematic of the LHCb detector with a right-handed coordinate system: the interaction point is defined as the centre, the z-axis is defined as the beam axis, the y-axis is pointing upwards. LHCb is a single-arm spectrometer covering a polar angle, θ of 10 mrad to 300 (250) mrad in the horizontal (vertical) plane. This geometry was chosen because of the angular distribution of B decays: both b and $\bar{b}$-hadrons produced in each event predominantly evolve in the same forward cone at high energies. This is illustrated in Figure 2.2 (left), which plots the polar angles of these hadrons with respect to the beam axis in the proton-proton (pp) centre-of-mass reference frame. This figure was made from Monte Carlo simulations generated in the PYTHIA event generator [16].

LHCb is designed to operate at a nominal luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, independent to the luminosity of other LHC experiments [17]. This is achieved by reducing the focus on the beams at the interaction point of LHCb. The nominal luminosity is dominated by single pp interactions per bunch crossing. This makes data analysis possible and reduces the rate of dilapidation on the detector with low occupancy and radiation damage. LHCb was designed to collect $\sim 2\,\mathrm{fb}^{-1}$ of data per year with a $b\bar{b}$-quark production cross section of $500\,\mu\mathrm{b}$ at 14 TeV centre-of-mass energy collisions. With these specifications, LHCb will collect much larger data samples of B mesons than any previous particle accelerator experiment.

The $B_s^0 \to D_s^- \pi^+$ channel was studied to determine the momentum range of

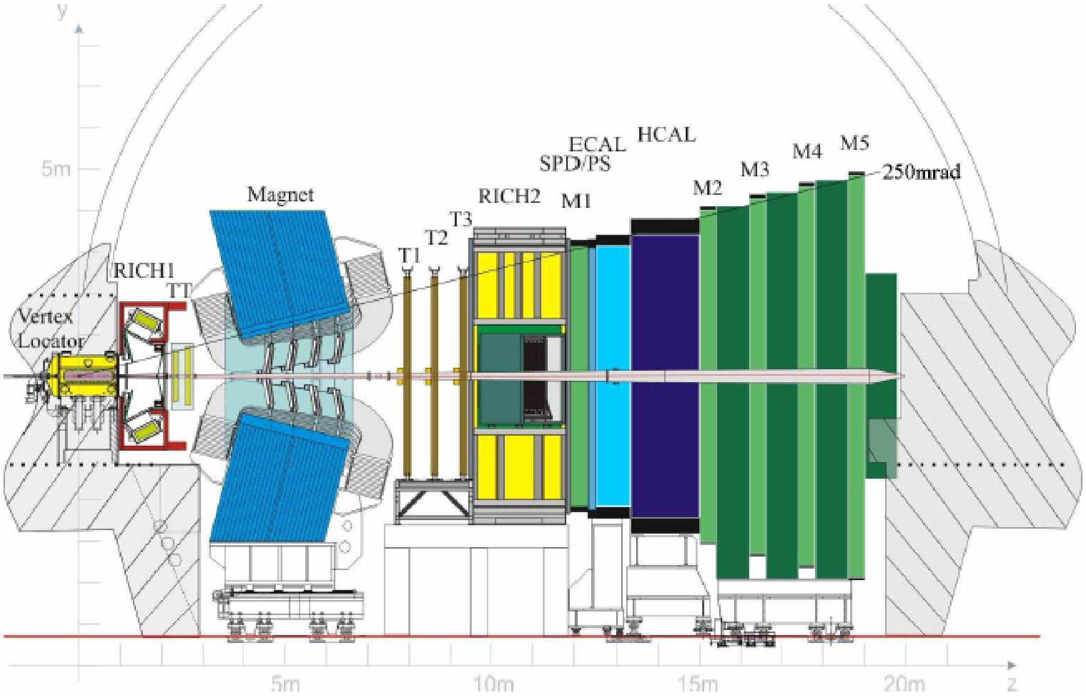

Figure 2.1: Sideview layout of the LHCb detector, a forward-arm spectrometer. z-axis is defined as the beam axis.

LHCb. The π^+ from the B decay has a harder momentum spectrum than the π^- from the D_s^- decay. Figure 2.2 (right) plots momentum distributions of the pions in simulated events of this channel, where both pions were in the detector acceptance. The π^+ momentum distribution indicated LHCb would need to cope with track momenta up to 150 GeV/c.

The LHCb experiment consists of a Vertex Locator (VELO), two Ring Imaging Cherenkov (RICH) detectors, a tracking system, a magnet bending in the horizontal plane, electromagnetic and hadron calorimeters and a muon detector. These subdetectors are used for the trigger, reconstructing decay products and flavour tagging. The LHCb subdetectors used in tracking particles are described in Section 2.2. The subdetectors used in particle identification are discussed in Section 2.3. LHCb will operate at a pp bunch-crossing rate too high to record and reconstruct every single event. Triggering is the multi-tiered procedure to reduce the data rate from 40 MHz to 2 kHz. After the uninteresting events are removed the remaining events are recorded for later analysis. This process is described in Section 2.4. Flavour tagging is the process of identifying whether a signal B decay was a b or $\bar{b}$-hadron and will be described in detail in Section 5.1.1. Finally,

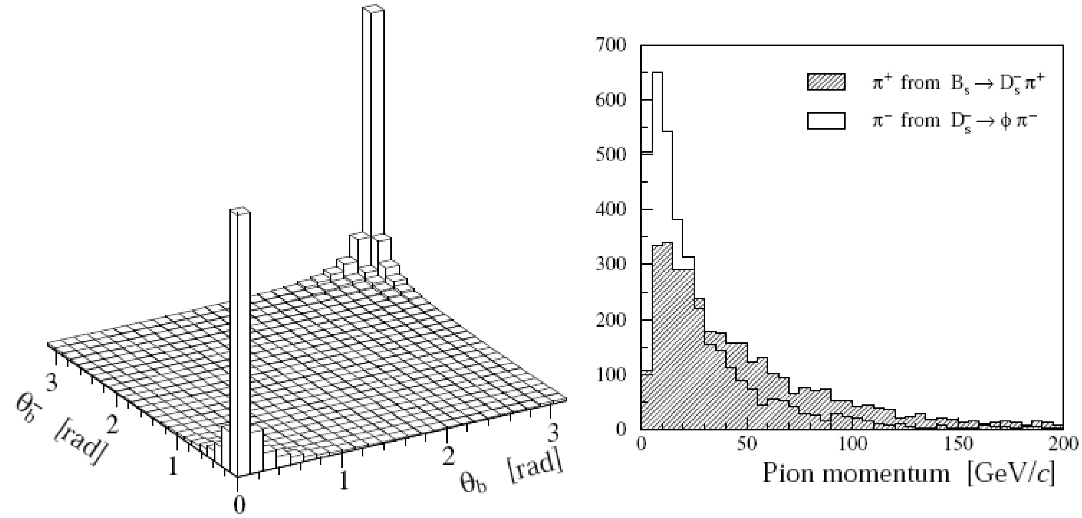

Figure 2.2: **Left:** polar angles of b and $\bar{b}$-hadrons simulated by PYTHIA the event generator. **Right:** Pion momentum distributions from simulated $B_s^0 \to D_s^- \pi^+$ events where both pions are in LHCb angular acceptance.

the planning going into upgrading LHCb is presented in Section 2.5.

2.2 Charged Particle Tracking

Tracking is the process of recording and reconstructing the flight trajectory of charged particles as they travel through the detector. Charged tracks are deflected by the magnetic field. Track momenta are deduced by comparing the trajectories before and after the magnet. A vertex detector makes precision measurements of tracks around the interaction vertex to reconstruct their origin. One tracking station before and three stations after the magnet record further space points to reconstruct tracks with. These tracking stations use two different technologies out of consideration for the different track occupancies in different regions of the detector acceptance.

2.2.1 Vertex Locator

B mesons have lifetimes of the order $\mathcal{O}(1\,\mathrm{ps})$. At the momenta produced at the LHC, this leads to the B mesons decaying at a displaced secondary vertex. It is important to separate the tracks coming from the B decay from those originating at the primary vertex. The Vertex Locator (VELO) is used to make precise

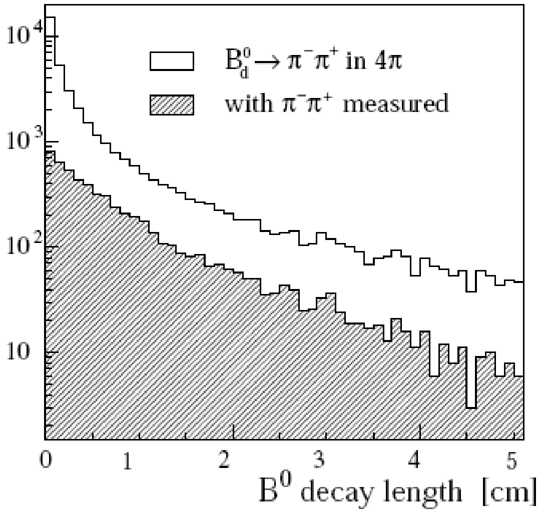

Figure 2.3: Decay length distribution for simulated $B^0_d \rightarrow \pi^+\pi^-$ events where both pions are within LHCb angular acceptance.

measurents of the track positions near the interaction region. From this data, production and decay vertices of *b* and *c*-hadrons are reconstructed and the impact parameters of tracks are measured. This in turn is used to calculate the lifetimes of the hadrons and for flavour tagging.

Figure 2.3 shows the B^0_d decay lengths of simulated $B^0_d \rightarrow \pi^+\pi^-$ events where both pions are within the LHCb detector acceptance. The average decay length is about 1.0 cm, illustrating the degree of accuracy demanded of this subdetector. The VELO's sensitive region starts at 8 mm radius from the beam-axis. However this is too close for beam injection. The VELO is made of two halves mounted on a rail system. This lets the VELO be retracted by 30 mm for the extra aperture required while the *pp* beam is injected, ramped and stabilised. The rail system can quickly bring back the two halves when the beam is declared stable.

The two VELO halves are each enclosed inside a thin-walled box made of aluminium. This protects the vacuum of the LHC beam pipe from outgassing from the VELO sensor modules. Vacuums are also maintained inside these boxes. The vacuum quality in the beam pipe and the VELO vacuum boxes is carefully balanced. The sides of the boxes facing the beam act as Radio Frequency (RF) shields, protecting the VELO against RF pickup from the LHC beams.

The RF shields have a corrugated structure to minimise the material budget and to accommodate the partially-overlapping arrangement of the sensors, explained below. The inside of the vacuum boxes have a polyamide-imide coating.

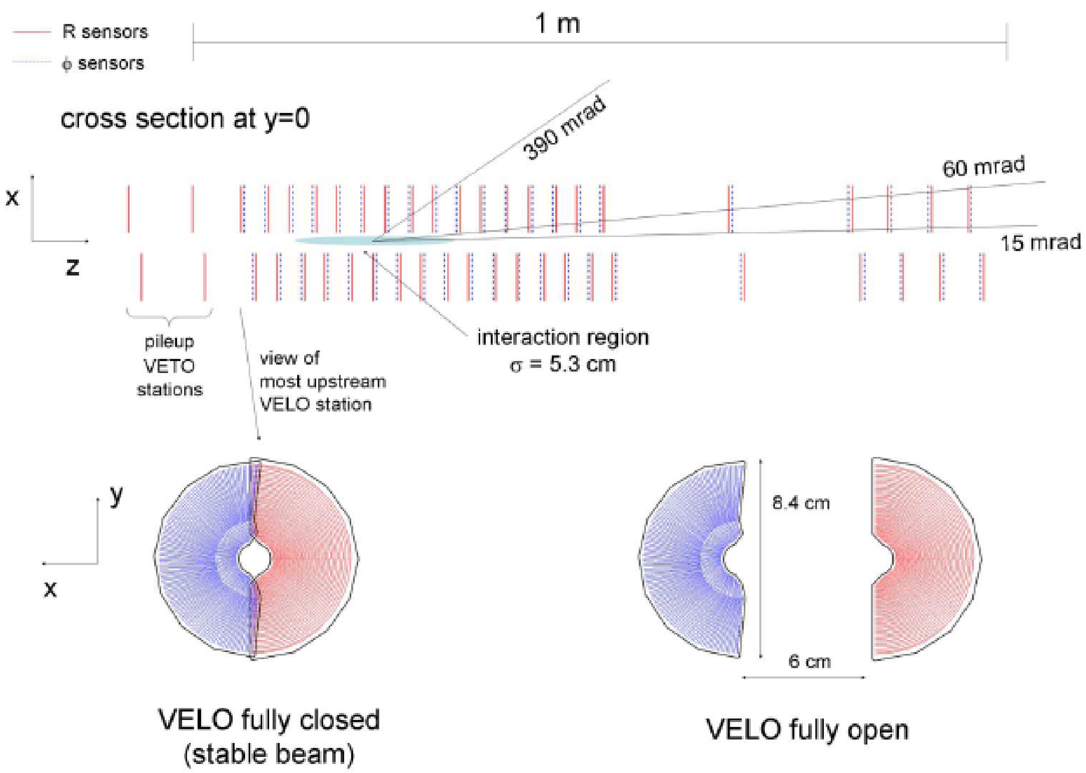

Figure 2.4: **Top:** A top-down view of the VELO in the horizontal plane the nominal interaction point lies. **Bottom:** A front view of one VELO station, illustrating the 2° overlap of both halves when closed and the 6 cm gap made when open for when the LHC is tuning its beam.

This is a highly radiation resistant material that electrically isolates the silicon sensors from the aluminium. The coating also helps dissipate heat generated in the foil from beam-induced effects.

The VELO uses two types of silicon sensors: one for radius measurements (R) and the other for azimuthal measurements (ϕ). VELO modules rigidly hold one of each type of sensor in a known position. VELO stations are defined as a pair of modules, one from each half of the VELO. Stations are positioned from z = -18 cm to +80 cm, as shown in Figure 2.4. The two veto stations in the opposite direction to the LHCb spectrometer are used in the trigger, as explained in Section 2.4.1.

Modules from the two VELO halves need to overlap for full azimuthal coverage and to facilitate relative alignment after the VELO halves are closed. So each station has the right-side module placed 1.5 cm behind the left-side module on the z-axis. The most downstream station positions are needed to guarantee that even tracks with a scattering angle, θ as low as 15 mrad pass at least three stations. The density of stations positioned near the nominal interaction region is derived from the same minimum requirement on space points for tracks with a θ up to 390 mrad [18].

Both two types of VELO silicon sensors are 220 μm thick and semicircular in shape, as shown in Figure 2.4. They span 182° with an inner radius of 8 mm and outer radius of 42 mm. The sensors are designed to optimise their radiation hardness and to minimise the channel occupancy. Analogue instead of digital readout is used as highly non-uniform radiation damage is expected. Analogue readout allows for a gradual adjustment of readout thresholds. The readout electronics and cooling for the sensors is routed through the VELO modules. Due to radiation, the sensors will need to be replaced within three years. The VELO is built to allow for easy access and replacement of the sensors.

2.2.2 Magnet

Precision B measurements require excellent charged particle momentum resolution of about 0.4%. This in turn requires a dipole magnet with an integrated bending power of 4 Tm for tracks up to 200 GeV/c [19]. This magnetic field has many constraints: it needs to be applied over the path segments measured by the tracking stations; it must be minimal for the VELO and RICH1; and the muon trigger demands that the field is uniform in the x-axis. In addition the magnet needs an aperture large enough to let pass all tracks within LHCb angular acceptance. Finally, to eliminate systematic errors from detector asymmetry when measuring CP Violation asymmetry, the magnetic field needs to be flipped regularly, e.g. every few weeks. The LHCb detector meets these requirements with a warm magnet made with a low carbon steel yoke and aluminium coils, which is shown in Figure 2.1. The magnet opening is wedge shaped to fit the LHCb angular acceptance. A superconducting magnet turned out to be unfeasible and disadvantageous in comparison.

The momentum resolution requirement for charged particles demands the magnetic field integral be measured with a relative precision of $\mathcal{O}(10^{-4})$ and the position of the B-field peak be known to within a few mm. Results from magnetic field mapping along the y-axis is plotted on Figure 2.5, showing excellent agreement with simulations. B measurement errors were about 4×10^{-4} within the tracking volume, completely dominated by the precision the measurement

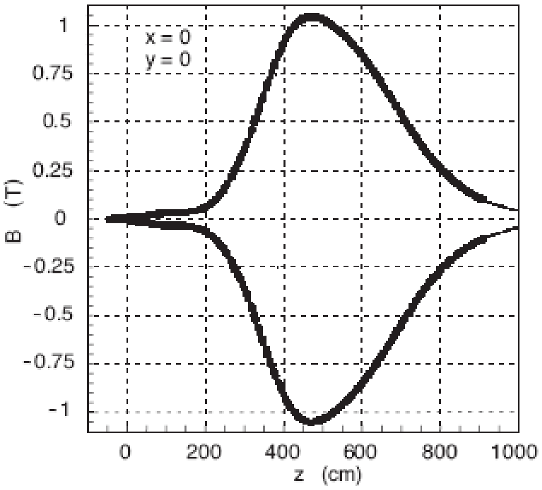

Figure 2.5: B_z measurements of the LHCb magnet (thick line) vs simulated predictions (thin line). For reference: the magnet is centred at $z \simeq 500$ cm, RICH1 is centred at $z \simeq 200$ cm and $z = 1000$ cm is beyond the last tracking station.

probes could be calibrated to. Numerous measurement campaigns with dipole field flipped show the magnet to be reliable and hysteresis effects to have negligible impact on the reproducibility of the field.

2.2.3 Silicon Tracker

Two detectors make up the Silicon Tracker (ST): the Tracker Turicensis (TT) and the Inner Tracker (IT). The TT is a 150 cm width × 130 cm height planar tracking station placed between RICH1 and the magnet. The IT occupies a cross-shaped region at the centre of each of the 3 tracking stations of dimensions [20] as illustrated in Figure 2.7.

Very high single-hit efficiency for minimum ionising particles is demanded to reconstruct high multiplicity B decays. Using silicon microstrip sensors both TT and IT have a detection efficiency of >99.8% provided the signal-to-noise ratio is still above 10:1. The sensor thickness was chosen to provide signal-to-noise ratios in excess of 12:1 even after 10 years of radiation damage from nominal luminosity operation.

Pile-up of consecutive events was avoided using fast front-end amplifiers with signal shaping time of the order of 25 ns to match the bunch crossing interval. Simulation studies showed that in the track reconstruction pile-up is avoided if the previous event's signal peak is already reduced to 50% (30%) in the TT (IT)

Silicon Tracker specifications	
Single-hit resolution	$\sim$50 μm
Relative positioning accuracy of sensor strips	<10 μm R.M.S.
Temperature maintained by cooling	<5°C

	TT	**IT**
Active Area	8.4 m^2	4.0 m^2
Silicon sensor dimensions	9.64 cm × 9.44 cm	7.6 cm × 11 cm
Sensor strip pitch	183 μm	198 μm
Sensor thickness	500 μm	320 μm & 410μm
Hit Occupancy at innermost regions	3.5%	2.5%
Hit Occupancy at outermost regions	0.35%	0.5%
Material Budget per station	0.04 X_0	0.035 X_0

Table 2.1: Table of selected Silicon Tracker design parameters

by the time the next event was sampled.

Table.2.1 lists a selection of other design parameters. The ST detector sensors are housed in light-tight boxes and are thermally and electrically insulated. Cooling to <5°C is required to suppress radiation damage induced leakage currents. The sensor volumes are continuously flushed with nitrogen to avoid condensation on the cold surfaces. For the TT, all dead material from cooling, mechanical supports and cables are located outside LHCb acceptance. For the IT this is impossible to achieve. Its peak material budget is almost 0.30 X_0 dominated by the contribution of the cooling rods.

Each ST station is made of four detection layers arranged in a (y-u-v-y) orientation: the first and last layers have vertical silicon strips and strips in the second and third layers are rotated by $-5°$ and $+5°$ to the y-axis respectively. In the TT, a gap in z of $\sim$27 cm separates the (y-u) layers from the (v-y) layers to aid track reconstruction. See Figure 2.6 for the layout of the v layer in TT.

TT detection layers are made from 7-sensor long half modules that cover half the height of the LHCb acceptance. Each half module groups its silicon sensors into either two or three readout sectors, which are read out separately by readout hybrids at the outer edge of the half module. Two thin rails made from fibreglass and carbon fibre provide the mechanical support to hold the half module together. A full module is made up from two half-modules, with services located at the top and bottom respectively. A detector layer in turn is made of seven (for the y-u

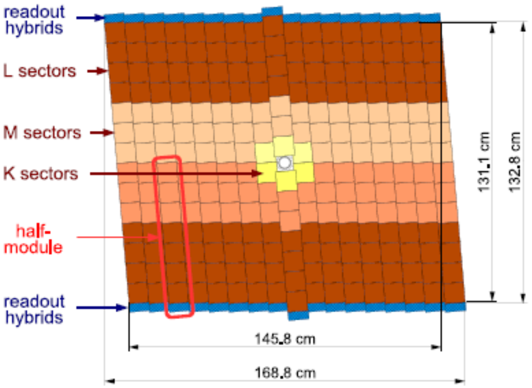

Figure 2.6: The third detection layer of the TT, where each module is rotated by +5° to the y-axis. The innermost modules have three readout sectors, modules getting less occupancy are grouped in two sectors for readout.

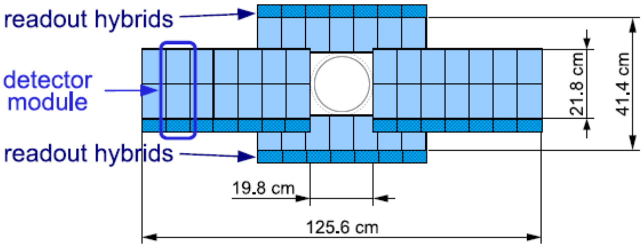

Figure 2.7: An IT detection layer. The regions above and below the beam pipe hold one sensor per module, and the regions to the side hold two per module.

layers) or eight (for the v-y layers) full modules each on the left and right side of the beam pipe. Adjacent modules within a detection layer are staggered by 1 cm in z and overlap by a few mm horizontally to prevent gaps in the acceptance.

Each of the three IT stations are made from four detector boxes: one above, one below and two to the side the beam pipe. Each box houses four detection layers. An IT detection layer is made of seven detector modules, as shown in Figure 2.7. Detection modules have one (two) silicon sensor(s) mounted for boxes above and below (to the sides) of the beam pipe. As with the TT, adjacent modules within a detection layer are staggered by 4 mm in z and overlap by 3 mm horizontally.

2.2.4 Outer Tracker

The Outer Tracker (OT) covers the remainder of the area of the three tracking stations around the IT, as shown in Figure 2.8 (top). It is a drift-time detector

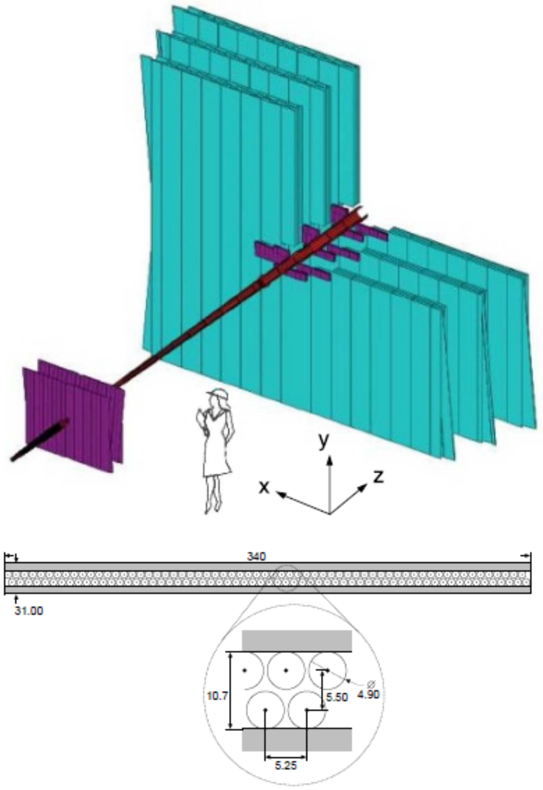

Figure 2.8: **Top:** Diagram of the layout of OT stations (blue) in relation to IT and TT stations (purple). **Bottom:** Cross-view schematic of a straw-tube module. All lengths are in mm.

consisting of an array of individual, gas-tight modules. Each module contains two staggered layers of straw-tubes of inner diameter 4.9 mm, as shown in Figure 2.8 (bottom). A mixture of Argon (70%) and CO_2 (30%) is used as the counting gas, featuring fast drift time (<50 ns) and drift-coordinate resolution of 200 μm. As with the IT, each OT station consists of four detection layers in a (y-u-v-y) orientation with the same $\pm5°$ rotations for the u and v layers. The total active area of a station is $5971 \times 4850\,\text{mm}^2$ and a material budget of 3.2% X_0 on average per station.

Similar to IT's two different module types, the OT has short modules for above and below the pipe and long modules for the sides. A detection layer is built from 14 long and 8 short modules. A short module contains 128 straw tubes, arranged in two staggered layers. Long modules are twice as long with twice as many straw-tubes. Straw-tubes are orientated vertically and are read out from the top or bottom of the modules. The complete OT detector with all three stations comprises ∼55000 single straw-tube channels.

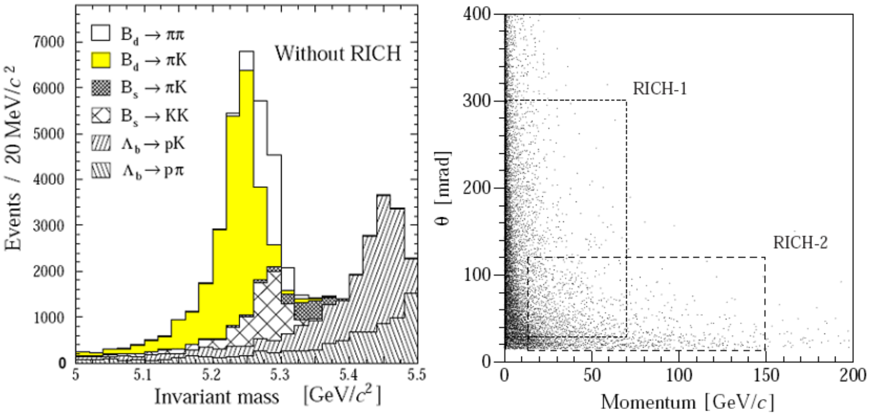

Figure 2.9: **Left:** Simulated $B_d^0 \to \pi^+\pi^-$ and background channels, plotted on the same axes in a no-RICH scenario. In particular, the $B_d^0 \to \pi K$ peak is virtually indistinguishable from the signal peak without RICH Particle Identification. **Right:** Polar angle θ distribution of simulated $B_d^0 \to \pi^+\pi^-$ events. The angular and momentum acceptance regions of both RICHes are indicated by the dashed lines.

The straw-tubes were produced by winding together two strips of thin foils. First a carbon doped polyimide (Kapton-XC) forms the 40 μm cathode on the inside and this is covered with a laminate (Kapton-aluminium) made of 25 μm polyimide (for gas tightness) and 12.5 μm aluminium (for fast signal transmission and shielding to avoid noise and crosstalk). The anode wire is made from gold-plated tungsten of 25.4 μm diameter and strung to a tension of 0.7 N. The straw tubes are positioned with a precision of 100 (500) μm in the x (z) direction and within each, the anode wires are centred to within 50 μm over the full length.

2.3 Particle Identification

2.3.1 RICH Detectors

The ability to separate pions from kaons is essential to distinguish signal from background in key B decays such as $B_d^0 \to \pi^+\pi^-$ [21]. This is illustrated in Figure 2.9 (left). The polar angle θ-momentum phase space is plotted with simulated events in Figure 2.9 (right). Based on this distribution, LHCb uses two Ring Imaging Cherenkov (RICH) detectors to perform pion-kaon separation over the wide phase space.

Schematics of the RICH detectors are shown in Figure 2.10. RICH1 is located between the VELO and TT. It uses two radiators, aerogel and C_4F_{10} gas to cover

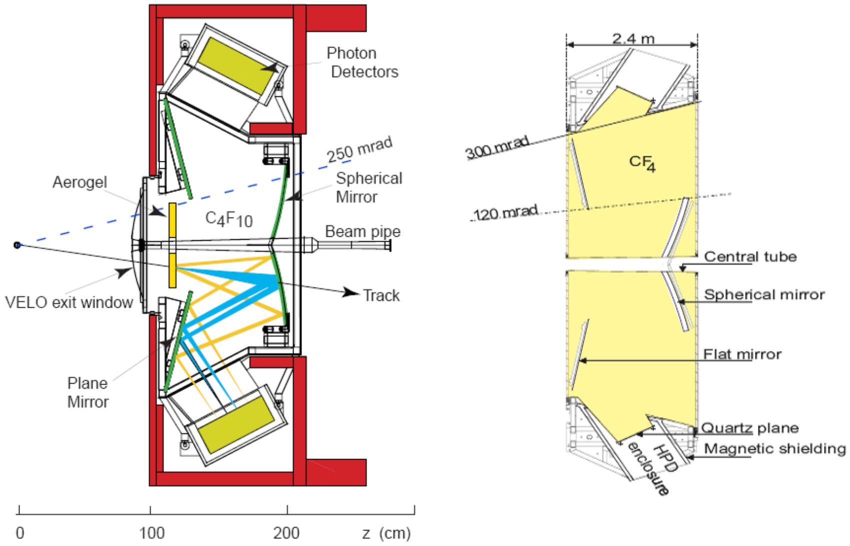

Figure 2.10: **Left:** Schematic side-view of RICH1, also indicating the working principle of the RICH detectors. **Right:** Schematic top-view of RICH2.

a lower momentum range of ∼1-60 GeV/c and the full LHCb angular acceptance. RICH2 is placed between the last tracking station and the first muon station. It has an angular acceptance of 15 mrad to 120 (100) mrad in the horizontal (vertical) plane. RICH2 covers a momentum range of 15-150 GeV/c using CF4 gas as a radiator. There is a 90° rotation between the two RICH detectors.

Figure 2.10 (left) also shows the working principle of the RICH detectors. A charged particle travels through the radiator faster than the speed of light in that medium. The particle emits a ring of Cherenkov photons at a specific angle dependent on the its speed. The particle passes through the detector towards the tracking stations. The Cherenkov photons are collected by mirrors onto the photon detector planes. The radius of the ring of Cherenkov photons observed by the photon detectors is used to deduce the speed of the particle that emitted them. Together with momentum information, the mass of the particle can be deduced and kaons can be separated from pions.

The mirrors direct and focus the Cherenkov light emitted in the RICH radiators to outside of the detector acceptance where the photon detectors are located. This minimises the material budget. Cherenkov rings may be split by the spherical mirrors and detected on both HPD planes on either side of the beam pipe. The spherical mirrors are located within detector acceptance so they were built to minimise the material budget. The total radiation length of RICH1 (RICH2) is ∼8% (∼15%) X_0.

Hybrid Photon Detectors (HPDs) have been chosen as the technology to detect the Cherenkov photons. These are described in Section 3. To summarise, a HPD combines two technologies: vacuum phototubes and pixelised silicon sensors. It uses the photoelectric effect to convert signal photons into photoelectrons that are accelerated to the anode by electrostatic fields.

Both RICHes are exposed to fringe fields from the LHCb magnet and the HPDs need shielding to attenuate the external B-field by a factor of > 20 (15) in RICH1 (RICH2). The detector planes in whole are housed in iron boxes of thickness 100 (60) mm at RICH1 (RICH2) providing a global shielding. In addition, each HPD tube is equipped with a cylindrical MuMetal (a high permeability alloy) shield.

In each RICH detector plane the HPDs are mounted in a hexagonally close packed arrangement, with the HPD centres separated by 89.5 mm. This achieves an active-to-total area ratio of 64%. The two RICH1 detector planes have seven horizontally orientated rows of 14 HPDs. In RICH2 9 columns of 16 HPDs are mounted in each plane, with the columns being vertically orientated. The HPD planes are placed so that the Cherenkov photons strike the entry windows at normal incidence on average.

As radiators RICH1 employs silica aerogel and C_4F_{10} gas, while RICH2 uses CF_4 gas. At $\lambda = 400$ nm, 0°C and 101325 Pa, their refractive indices are n = 1.03, 1.0014 and 1.0005 respectively. Aerogel provides effective particle identification in the momentum range of $\sim$2-10 GeV/c. The C_4F_{10} gas extends the particle identification sensitivity to $\sim$60 GeV/c for the full acceptance of LHCb. In RICH1 particle tracks pass through 5 cm of aerogel and 95 cm of C_4F_{10} on average. Thus for a saturated track ($\beta = 1$), $\sim$6.5 and $\sim$30 photoelectrons are expected to be detected. For RICH2 the tracks pass on average 180 cm of CF_4 gas and a saturated track is expected to yield $\sim$22 photoelectrons. These expectations for the photon yields were derived from test-beam data. They match results from Monte-Carlo simulations and are close to what can theoretically be achieved.

An assumption must be made by the algorithms that reconstruct the ring images from hits on the photon detectors. They assume the particle emitted all its Cherenkov light from the middle point of its trajectory through a radiator. Particles are continually radiating Cherenkov photons while travelling through the RICH. Therefore the reconstructed Cherenkov angle, and in turn the velocity, of a track differs from its true value. The RMS of the distribution of this smearing is referred to as the emission point error. The RICH optics was designed to reduce this emission point error to a point not larger than other sources of finite angular resolution, such as HPD pixel size and chromatic dispersion of the radiator.

2.3.2 Calorimeters

The electromagnetic calorimeter (ECAL) and hadronic calorimeter (HCAL) are used to measure the total and transverse energy of hadrons, electrons and photons. The calorimeters are also used for particle identification. The most demanding task is the identification of electrons fast enough to be used in the trigger. Two major sources of background need to be accounted for before this is possible: rejecting charged pions requires a preshower detector (PS) in front of the ECAL and rejecting neutral pions with high transverse energy, E_T requires a scintillator pad detector (SPD) in front of the PS to select charged particles. These two forward detectors are nearly identical and sandwich a 15 mm thick lead converter.

The SPD/PS detector has an active area of 7.6 m width × 6.2 m height and fits into the 180 mm gap between the first muon station and the ECAL. The ECAL is placed at $z = 12.5$ m and covers the LHCb angular acceptance down to 25 mrad. This cutoff point is due to substantial radiation dose level limiting ECAL acceptance around the beam pipe. The HCAL is located at $z = 13.33$ m. It is 8.4 m high × 6.8 m wide × 1.65 m in depth and weighs a total of ~500 tons.

Optimal energy resolution requires full containment of the showers from high energy photons. Therefore the ECAL thickness is 25 X_0 in addition to the 2.5 X_0 from the SPD/PS. Such stringent containment is not required of the HCAL in the trigger, so its thickness is 5.6 interaction lengths, λ_I due to space limitations. This is in addition to the 1.2 λ_I of the ECAL.

These four detectors all vary the size of their cells based on proximity to the beam pipe: HCAL is segmented into two sections and the other three are split into three sections [22], as shown in Figure 2.11. The SPD, PS and ECAL all have a one-to-one projective correspondence. Each of their planes are subdivided into square cells of length 4, 6 and 12 cm for the inner, middle and outer sections respectively. The HCAL's inner section uses square cells of size 131.3 mm and the outer section's square cells are double in length

The four calorimeters share the same basic principle: scintillation light is transmitted via wavelength-shifting (WLS) fibres to a Photo-Multiplier Tube

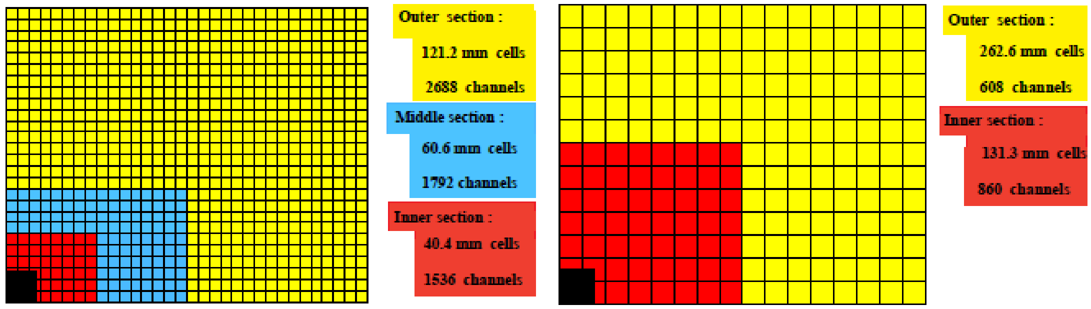

Figure 2.11: Lateral segmentation of the SPD/PS and ECAL (left) and the HCAL (right). One quadrant of the detector face is shown. The black area is the area not covered because of proximity to the beam pipe. The left (right) figure shows cell dimensions for the ECAL (HCAL).

(PMT). The SPD/PS cells are read out with Multianode Photo-Multiplier Tubes (MAPMTs) located at the detector periphery inside magnetic shielding, whereas the ECAL and HCAL send light through fibre bunches to individual phototubes located directly on the modules. The scintillator material uses polystyrene as the base component which is then treated with WLS dopants in concentrations tailored to each detector. The readout electronics make use of the speed of the scintillator and WLS based calorimeter system and avoid pileup signals by being capable of reading out consecutive bunch crossing events.

The ECAL uses 'Shashlik' caliorimeter technology: a sampling scintillator/lead structure readout with plastic WLS fibres for readout. The 1.2 mm diameter readout fibres penetrate the scintillator plates orthogonally. ECAL modules are built from alternating layers of 2 mm thick lead, 120 μm thick, reflecting TYVEK paper and 4 mm thick scintillator tiles. Overall, the ECAL has 66 lead and scintillator layers with a total depth of 42 cm corresponding to 25 X_0. This stack is made light-tight.

The ECAL's energy resolution is $\sigma_E/E = (9\%/\sqrt{E}) \oplus 0.8\%$ (E in GeV) as shown by test beam data. It does not deteriorate due to energy absorbed in the PS in front of the ECAL if the energy measured in the PS is used for corrections. Aging tests showed the innermost ECAL modules would increase the constant term of their energy resolution from 0.8% to 1.5% after 8 years of radiation damage from nominal luminosity. While this degradation is expected and acceptable, the ECAL was designed to let these innermost modules be replaced if necessary.

The HCAL uses iron as its absorber and scintillating tiles as its active material.

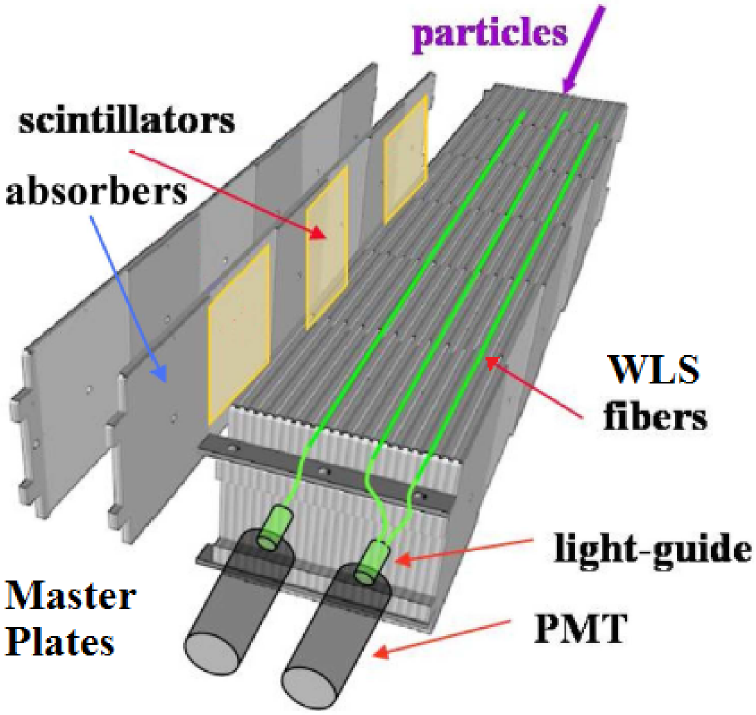

Figure 2.12: View of the HCAL structure. Two master plates with spacers attached are shown exploded with scintillators mounted in the gaps. Wavelength-shifting fibres collect the scintillator light for the PMTs to readout.

Unlike the other calorimeters, the HCAL orientates its scintillaters parallel to the beam-axis, interspersed with 1 cm of iron. In the z-axis, the length of tiles and iron spacers are based on the hadron interaction length λ_I in steel. Active to passive material ratio is 0.18.

HCAL modules repeat 216 times with identical periods of 20 mm thickness. Their layout is illustrated in Figure 2.12. One period contains two 6 mm thick master plates of 1283 mm length × 260 mm height. 4 mm thick spacers of variable length are glued to each master plate to fill up the period's 20 mm thickness. Gaps are left in the spacers where the 3 mm thick scintillator tiles and 1.2 mm diameter WLS fibres are placed. The 1.6 m long WLS fibres collect and transmit the scintillator light to PMTs housed in magnetic shielding at the backside. The full detector is built from 52 modules. Test beam data measured the HCAL resolution to be $\sigma_E/E = (69 \pm 5)\%/\sqrt{E} \oplus (9 \pm 2)\%$ (E in GeV).

2.3.3 Muon Chambers

Triggering with and offline identification of muons are an essential part of the LHCb's B physics program. The muon system has five stations (M1-M5). The geometry of the stations are projective, making their transverse dimensions proportional to their position on the z-axis. The angular acceptance of the muon system covers the LHCb acceptance down to a minimum of 20 (16) mrad in the bending (non-bending) plane respectively. This gives an acceptance of about 20% for muons from inclusive b semi-leptonic decays.

The first level of the trigger demands transverse momentum, p_T to be measured quickly. The first station, M1 is placed in front of the calorimeters to improve the muon p_T measurement. Stations M2-M5 are downstream of the HCAL. The first three stations have high spatial resolution in the bending plane to define track direction and calculate p_T to a resolution of 20%. The last two stations are for identifying penetrating particles and thus have limited spatial resolution.

Stations M2-M5 have 80 cm thick iron absorbers inbetween each adjacent pair of stations. The total absorber thickness is $\sim$20 interaction lengths including the calorimeters. Only muon tracks that have coincident hits on all five stations are accepted by the trigger. Therefore the minimum muon momentum that can be detected is 6 GeV/c. Each station has a detection efficiency of at least 95%.

As with the tracker stations and calorimeters, the muon stations have different particle fluxes depending on distance from the beam axis. Each muon station is divided into four regions (R1-R4) defined from centre outwards [23], as shown in Figure 2.13. The regions have linear dimensions and segmentations scaling in the ratio 1:2:4:8. This is expected to roughly even out the channel occupancy over the four regions of a given station. Different detection technologies are used for regions with different track occupancies.

Multi-Wire Proportional Chambers (MWPC) are used in all the muon stations, except for the central region of the first station, M1R1. An exploded schematic view of a MWPC is shown in Figure 2.14 (left). The chambers have four (two in M1) gas gaps, filled with a gas mixture of Ar/CO_2/CF_4 (40:55:5). The gaps are grouped together into pairs by logical OR. This double gap has an efficiency of >95% in a 20 ns window. In stations M2-M5 the two double gap layers are operated and read out independently.

Different regions use different readout electrodes for their MWPCs. The

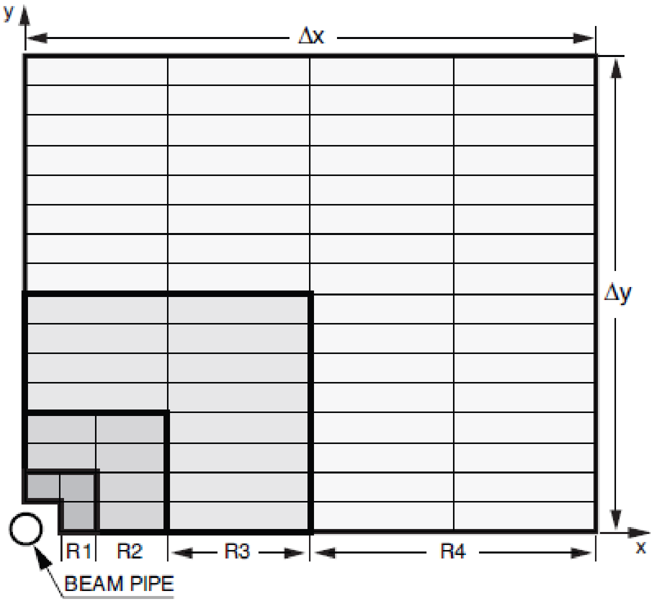

Figure 2.13: Front view of a quadrant of a muon station, showing how the detector is divided into regions. Each rectangle represents one chamber.

outermost regions of all stations use wire pads to meet the modest spatial resolution requirements. Most other regions use cathode pads. Regions R1-R2 in stations M2-M3 use a hybrid solution where y-resolution is defined by a cathode pad but narrow wire-strips read out the x-resolution. This hybrid readout has greater precision than cathode pads alone can practically achieve.

The central region of the first station, M1R1 uses triple Gas Electron Multiplier (GEM) detectors to cope with the higher particle flux. The region has 12 triple-GEM chambers of active area $20 \times 24\,\mathrm{cm}^2$. Each chamber has 2 triple-GEM detectors superimposed to form 2 sensitive layers. These layers are then grouped together with logical OR like the double gaps in the MWPCs.

Triple-GEM detectors consist of three GEM foils are sandwiched between a single anode and cathode planes, as shown in Figure 2.14 (right). A gas mixture of $\mathrm{Ar}/\mathrm{CO}_2/\mathrm{CF}_4$ (45:15:40) surround the foils. Tracks passing through ionise electrons in the drift gap between the cathode and the first GEM foil. These electrons are accelerated by electric fields through the GEM foils. At each GEM foil they are multiplied and the resulting cascade drifts to the anode giving rise to an induced current signal. A time resolution of $\sim$7 ns is achieved.

All muon chambers are segmented into physical pads, each read out by one front-end (FE) electronics channel based on custom radiation-hard chips. The

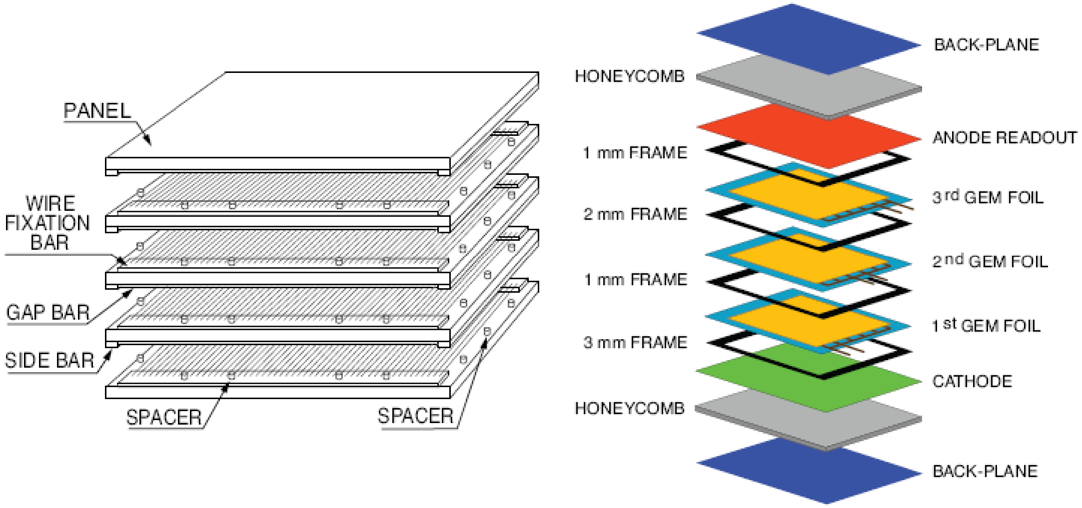

Figure 2.14: **Left:** Exploded view of an MWPC muon chamber. The 4 layers are grouped together into pairs by the logical OR. Charged particles passing through the gas layers create ions that drift to the wires to be read out. **Right:** Exploded view of an M1R1 triple-GEM detector. Electrons liberated by the ionising signal particle are accelerated to the readout anode, multiplying at each GEM layer.

electronics include flexible logical units to perform OR operations on a variable number of FE channels. In most chambers, the size of the physical pads must be smaller than required by spatial resolution to maintain acceptable levels of noise and FE channel dead-time by limiting the rate and electrical capacitance of pads. For such cases, up to four adjacent physical pads are combined into a logical unit by an OR operation done in the FE electronics.

2.4 Trigger

The LHC has a bunch crossing frequency of 40 MHz. This is too fast to store the data from every event. Most of these bunch crossings will be uninteresting for B physics. The role of the trigger is to reject such events before they are stored for full reconstruction and physics analysis.

At the nominal luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, from the 40 MHz of bunch crossings there will be about 100 kHz of reconstructable events containing $b\bar{b}$-pairs [24]. Of these, only about 15% will have all the decay products of at least one B meson be contained in the detector acceptance. Furthermore, many channels of interest are rare decays with branching ratios typically less than 10^{-3}. The trigger was optimised to maximise the efficiency of physics analysis event selections while rejecting the background events as strongly as possible.

The trigger is divided into two tiers: Level-0 (L0) and the High Level Trigger (HLT). The latter is further split into two phases called HLT1 and HLT2. Both the L0 and the HLT can be emulated fully on stored data to monitor trigger performance.

2.4.1 Level-0 Trigger

The Level-0 trigger reduces data rate from 40 MHz to $\leq$1 MHz using limited reconstruction information. The large mass of B meson decays often produce particles with large transverse momentum, p_T and energy E_T. The L0 trigger bases its decision on three sources of information: the calorimeter triggers, the muon trigger and the pile-up system in the VELO. These three systems pass information to the L0 Decision Unit (DU) which decides whether an event should be rejected or be sent to the HLT.

The calorimeter trigger system divides the ECAL and HCAL into 2×2 cells. This zone size is large enough to contain most of the energy from a single particle's shower and is small enough to avoid overlap of multiple particles. In each cell it sums the E_T of showers contained in it. Clusters are identified as originating from an electron, photon or hadron based on information from all four calorimeter detectors. The highest E_T electron, photon and hadron are identified and their information is sent to the DU. The total number of SPD cells with a hit is counted and also given to the DU. This gives an estimate of the charged track multiplicity in the event.

The muon chambers can perform a stand-alone muon reconstruction. This is done by searching for hits defining a straight line through all five stations that point back to the interaction point. The position of a track in the first two statoins allows its p_T to be determined to a resolution of $\sim$20%. For each quadrant of the muon detector, the two highest p_T muons are selected for the DU.

The pile-up system uses four R-sensors in the VELO located on the opposite side of the interaction point compared to the LHCb spectrometer. It distinguishes between bunch crossings with single and multiple pp interactions. The system provides the DU with the position of candidates for primary vertices and a measure of the total backward charged track multiplicity.

Events that the DU approves for the HLT are written to the Event Filter Farm (EFF), which contains about 2000 computing nodes. The information used by the DU to make this decision is also written to the EFF. For events to be approved, they must satisfy one of these criteria: highest $E_T^{hadron} > 3.5\,\text{GeV}$; highest $E_T^{e,\gamma,\pi^0} > 2.5\,\text{GeV}$; at least one muon candidate has $p_T > 1.2\,\text{GeV}$; or for the two muons with largest p_T, $p_T^{\mu_1} + p_T^{\mu_2} > 1.0\,\text{GeV}$. Note these cuts can be changed.

The L0 trigger is fully synchronous with the 40 MHz bunch crossing frequency. All L0 electronics is implemented in custom-designed boards to achieve the necessary speed. The time between a *pp* interaction and the arrival of the DU's final decision at the front-end electronics is fixed at $4\,\mu$s. This time includes the time-of-flight of the particles and delays in electronics and data transmission cables, leaving $2\,\mu$s for the DU to process the L0 data to make a decision. The $4\,\mu$s delay is unaffected by the occupancy nor the bunch crossing history.

2.4.2 High Level Trigger

The EFF serves as a buffer for the HLT to examine events before they are permanently stored for physics analysis. The HLT has access to all data in events that pass the L0 trigger, but seeks to reject uninteresting events using only part of that data. The HLT is fully implemented in software and thus is very flexible. It is evolving with knowledge from early beam collisions and the physics priorities of LHCb.

HLT1 reduces the 1 MHz output rate of the L0 down to ∼30 kHz. The L0 DU notes down why it decided an event was approved. HLT1 examines this reason with progressively more information from the VELO or the tracking stations. For each different criteria that an event can pass the L0 trigger, HLT uses a separate procedure to make its decision. These progressive confirmation procedures are called alleys.

About ∼15% of L0 passing events are selected by multiple triggers and thus, will be examined by more than one alley. While the alleys operate independently, precautions are made to avoid wasting CPU power by having to reconstruct the same track or primary vertex twice.

Events that pass any alley in HLT1 are fed into HLT2, where track reconstruction is performed almost as fully as in physics analysis. First, a set of tracks are selected based on very loose momentum and impact parameter cuts. Composite particles are formed from these tracks, such as $\phi \to K^+K^-$ and $J/\psi \to \mu^+\mu^-$. These composite particles are then used for all HLT2 selections to avoid duplication in creation of final states.

The trigger up until now has not used invariant mass data or precise cuts on pointing of the B momentum towards the primary vertex. HLT2 has two groups of cuts: inclusive and exclusive selections. Inclusive triggers select generic, partial B decays e.g. $B \to \phi X$ or $B \to J/\psi X$. They are less dependent on the HLT2 reconstruction than the exclusive triggers which try to find specific channels. However, exclusive selection of channels produces a smaller rate. HLT2 approves of an event for permanent storage if any of the inclusive or exclusive selections gives a positive result. The data rate is reduced from 30 kHz from the output of HLT1 to 2 kHz after the HLT2.

2.5 Upgrade Considerations

After gathering $\sim 10\,\mathrm{fb}^{-1}$ of data over about five years, continued running of the LHCb at the same luminosity will not be very profitable. The statistical precision on measurements increases very slowly after this point so the LHCb will be upgraded to work with a higher luminosity of $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. R&D is ongoing to adapt each LHCb subdetector to cope with the upgrade luminosity and improve the readout rate of all subdetectors from 1 to 40 MHz [15].

For the VELO, radiation hardness is the primary concern. Even at nominal luminosity, the silicon sensors are foreseen to be replaced once after $6\,\mathrm{fb}^{-1}$ of data taking. New sensors with short strips or pixels have been designed and are being studied alongside ways ot implement 40 MHz readout. Potential imporvements to the RF-foil is also considered for reducing the material budget.

The IT does not expect significant radiation problems, but endeavours to reduce the radiation length of the electronics while increasing readout rate to 40 MHz. The OT also faces a likely complete redesign of its front-end electronics. In addition, the main concern of the OT is the large increase in occupancy at upgrade luminosities. A new drift time vs spatial resolution optimisation study is needed to determine what tracking technology to use in the upgraded OT.

The RICH is doing R&D to study various photodetector technology choices for the upgrade environment. Work must also be done on new front-end electronics to adapt to the 40 MHz readout. This means the photon detectors of the RICH detectors needs to be replaced because they have encapsulated electronics. Furthermore, considerations are made to removing RICH1 entirely and relying on a super RICH detector. Also being studied is a new time-of-flight RICH detector (TORCH) which will be installed in the space the M1 station currently resides.

Preliminary ECAL studies show signal selection efficiency does not significantly degrade with the existing ECAL but an upgrade is not excluded yet. The radiation hardness of the inner section is also a critical concern. At upgrade luminosity, it is foreseen that about half the inner section modules will need replacing after 3 years of operation. The HCAL does not foresee any major deterioration to the hadron trigger performance if the current HCAL is used in the upgrade. As with other subdetectors, the front-end electronics of the calorimeters need to be replaced for 40 MHz readout.

In the upgraded LHCb, the M1 station will no longer be needed. In the present L0-μ trigger, this station improves momentum resolution by $\sim$30% but the tracking system will determine the momentum of muon candidates in the upgrade. Being the most shielded subdetector, the higher particle flux is less of an issue but for the M2R1 MWPC detectors ageing is still a concern. After experience with running at nominal luminosity, an assessment will be made whether the M2R1 region should be replaced with triple-GEM technology like the current M1R1 region. The muon system is already reading out at 40 MHz.

Chapter 3

Quality Testing Hybrid Photon Detectors

3.1 Introduction

For the LHCb RICH detectors to perform as expected, tight requirements were placed on their photon detectors. A total area of 3.3 m^2 needed to be covered, with ∼70% of active area at a reasonable cost. High sensitivity to single photons in the 200 nm–600 nm wavelength region was required at a spatial resolution of 2.5 mm×2.5 mm. The LHC bunch-crossing rate imposed the requirement of the front-end chip operating at 40 MHz. The photo-detectors will be exposed to a maximum of 3 kRad radiation dosage per year and cope with a magnetic fringe field of 10–25 Gauss from the LHCb dipole magnet.

After beam tests and laboratory evaluations, the technology chosen to fit these requirements was the Hybrid Photon Detector (HPD). Many companies were involved in the manufacturing process of the HPD, with the final assembly and encapsulation being done at Photonis[1] in Holland. Figure 3.1 shows the main functional parts of an HPD and a photograph of a finished HPD.

The quartz window has a multi-alkali photocathode deposited on the inside, converting incoming photons into photoelectrons. A high voltage supply biases the photocathode to -20 kV relative to the anode. This bias accelerates and cross-focusses the photoelectrons through the vacuum of the HPD tube body, onto the silicon sensor on the anode. The electron-optics inverts and de-magnifies the photocathode image by a factor of 5, allowing for the sensor and readout

---

[1] formerly known as DEP

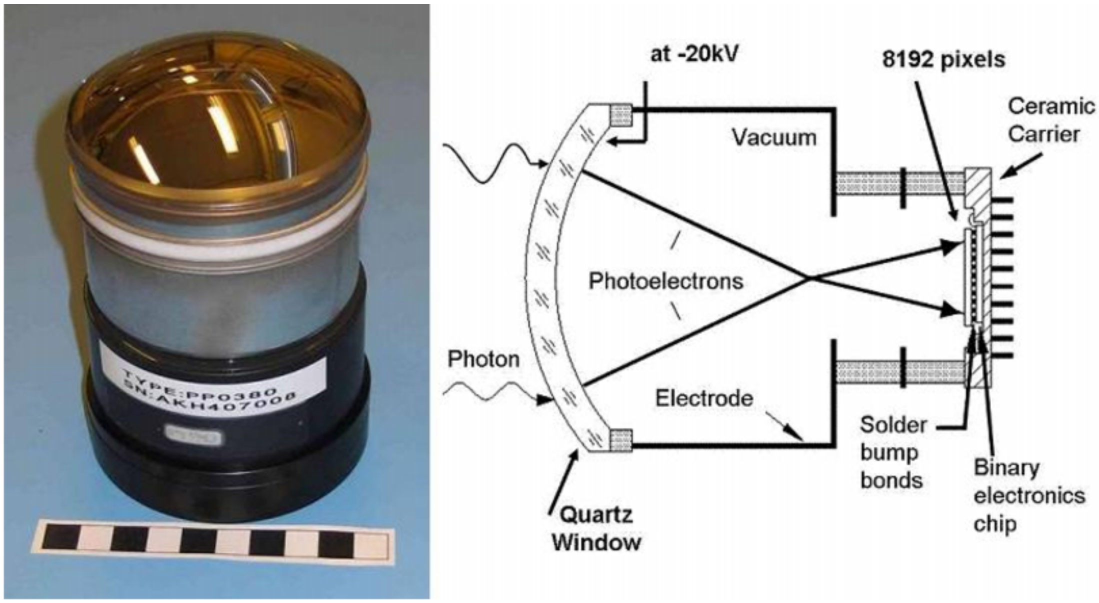

Figure 3.1: A photo of a finished HPD, without the μ-metal shield and a simplified diagram of how an HPD works. It stands 120 mm tall (without the lens cap) and is 87 mm in external diameter. The photocathode is charged to -20kV in comparison to the Anode. Photoelectrons from the photocathode are accelerated towards the anode, where silicon sensors are connected to a readout chip.

electronics to have a smaller active area than the photocathode.

The readout chip was a joint design between LHCb and ALICE. The silicon sensor has 8192 pixels of size $500\,\mu\text{m} \times 62.5\,\mu\text{m}$. Each pixel is bump-bonded to its own channel on the readout chip. LHCb groups 8 of these pixels with an OR operation for its data taking. When a photoelectron hits the sensor, $\sim$5000 electron-hole pairs are released[2]. A Low Voltage (LV) supply biases the silicon sensor to 80 V which draws the released charge carriers to the readout terminal.

In addition to these features, further requirements were set out for quality control. Selected HPD specifications are listed in Table 3.1. Ion feedback and quantum efficiency are explained in Sections 3.2.2 and 3.3 respectively. Photonis would perform quality checks on all HPDs they assemble, and LHCb would independently confirm these tests. To do this independent quality control, the Photon Detector Test Facilities (PDTF) group was formed.

---

[2] The photoelectron's energy is 20 keV and 4 eV are required to produce one electron-hole pair in silicon.

RICH requirements for its Photon Detectors	
Total area of photodetectors	$3.3\,\mathrm{m}^2$
Active area	~70% at reasonable cost
Sensitivity	Single photons
Wavelength range	200 nm–600 nm
Spatial resolution	2.5 mm×2.5 mm
Time resolution	25 ns
Radiation tolerance	3 kRad per year
Magnetic field tolerance	10–25 Gauss
De-magnification factor from electron optics	~5
Quantum Efficiency at 270 nm	$> 20.0\%$
Quantum Efficiency at 400 nm	$> 15.7\%$
Dark count rate at 20°C	$< 5\,\mathrm{kHzcm}^2$
Ion Feedback probability	$< 1\%$
Lifetime (in which output does not drop $> 10\%$ to inital response)	10 years

Table 3.1: Table of selected HPD specifications

3.1.1 Photon Detector Test Facilities

The LHCb RICH detectors require 484 HPDs to populate its photon detector planes. Photonis were contracted to produce 550 HPDs so as to have spares over the ~10 year lifetime of the experiment. To independently test and categorise all these HPDs, the PDTF group set up 2 test stations each at Edinburgh and Glasgow. Each site tests one HPD per working day, with Photonis delivering HPDs in batches at an equal rate. 10% of all produced HPDs receive further tests for Quantum Efficiency as described in Section 3.3. Afterwards, HPDs were shipped to CERN from the PDTF sites, being immersed in dry air while in storage.

Figure 3.2 shows a PDTF test station, which consists of a dark box, power supplies, interaction electronics and a desktop PC. The dark box has a photomultiplier installed inside for detecting light leaks and a hole on one side for the HPD to be mounted onto. Readout instruments and control circuits are connected to the mounted HPDs chip. An optical fibre cable feeds the light source from the LED into the dark box. Separate power supplies are used to bias the HPDs silicon detector (LV) and the photocathode (HV). All this is controlled by the desktop PC running Labview, automatically performing tests and taking

Figure 3.2: Early setup of a PDTF station at Edinburgh. It is largely unchanged even today. From left to right: darkbox where HPD is mounted, HV supply and electronics, MB & FB board for interfacing with the readout chip, other power supplies and picoammeter, all controlled by the PC through the JTAG controller. The oscilloscope checks the blackout conditions using a photomultiplier inside the darkbox.

measurements. Results would be displayed back in Labview and the calls the software made to Cygwin via Windows batch files were transparent to the user.

3.2 Standard Tests

The large amount of data acquisition and analysis required for each HPD test was automated as far as possible to minimise human error. A 2-phase routine was worked out that could be carried out by a trained technician or student, making it possible to find manpower to do the tests.

When the HPD is first taken out of its box, its HV cables and ZIF socket pins are visually inspected for defects. Utmost care is taken to reduce the risk of HV discharge by using unpowdered latex gloves to avoid fingerprints and teflon tape covers for uneven surfaces on the indium seal. Furthermore, a grounding strap is used to prevent static electrodischarges while handling the HPD. Electrical connectivity of each HV cable is checked for with a digital multimeter. The tube body and ZIF socket are cleaned in such a way to avoid causing damage. A mechanical jig with illuminating LEDs is used to check that the mechanical dimensions of the HPD fit within tolerances of ± 0.1 mm.

For the automated second phase, the HPD was mounted in the dark box. After mounting the HPD into the dark box and checking its light-tightness, a test programme implemented in LabView is carried out. First it checks that the electronics are functioning properly, then characterises the HPD anode sensor. On human approval of results, HV is ramped up slowly in steps with the user keeping an eye out for persistent microdischarges or unexpected increases in photoelectron hits and/or anode current.

After the HV ramp-up, the PC will characterise the photon detection of the HPD. Dark count noise, Ion Feedback (IFB) behaviour and image distortion due to demagnification are recorded. Any noisy or dead pixels are noted. Performance is plotted against a range of HV and LV values. IFB will be further discussed in Section 3.2.2.

3.2.1 Overview of Results

HPDs were divided into 5 categories based on their performance in the PDTF Standard Tests: A+, A, B, E, F. A+ HPDs exceeded design specifications, A-grade cases met specifications and B-grade HPDs were still usable in the RICH despite not meeting all expectations. E-grade was given to HPDs not meeting specifications but still usable as contingency spares. HPDs clearly failing the tests were given an F and shipped back to Photonis for replacement or repair.

Figure 3.3 summarises the test results for all 561 HPDs (includes replacements and remade HPDs) tested at PDTF [25]. 97.1% of HPDs passed the tests with 89.7% attaining a B-grade or better. This grading system guided the RICH commissioning when selecting the best HPDs for the highest occupancy positions in the RICH detector planes.

3.2.2 Ion Feedback

Ion Feedback (IFB) occurs in HPDs with a degraded vacuum. As photoelectrons travel from the photocathode to the anode, they ionise remnant gas particles they collide with. These positive ions are accelerated by the HPD's electric field back to the photocathode, where they release many secondary photoelectrons. At the anode, these secondary photoelectrons are detected about 300 ns after the primary signal photoelectron in a tight cluster of pixels. If the HPD vacuum is very contaminated, then the secondary photoelectrons will further ionise gas particles, resulting in a chain reaction and making the HPD unusable for data taking.

IFB measurements are defined by counting clusters of hits. A single

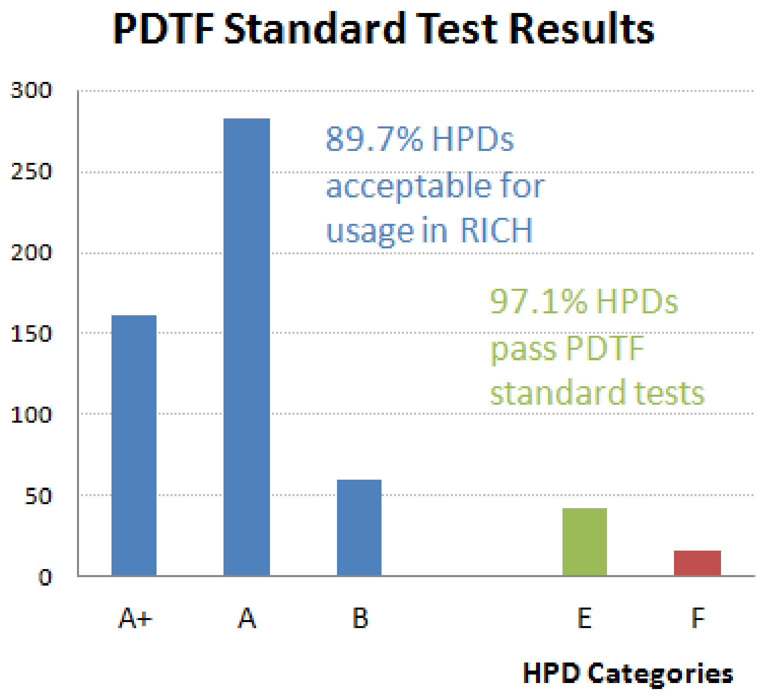

Figure 3.3: Summary of HPD test results at PDTF [25]. A-grade HPDs meet all specifications while B-grade HPDs are still usable despite minor shortcomings.

photoelectron from a signal photon can conceivably register hits in up to four neighbouring pixels if it lands in a corner. Therefore incidents where ≥ 5 adjacent pixels are hit in the same event are treated as IFB. The rate of IFB is defined as the ratio of these large clusters to all clusters accumulated over all events. See Section 4.2 for further discussion including compensating for noise.

To determine the time delay between signal photoelectron and secondary IFB photoelectrons, PDTF measured IFB by a strobescan (SS) method. The light source was an LED at wavelength 470 nm, firing a single pulse of light for each event. The HPD would readout any hits in a 50 ns window when it received a Data Acquisition (DAQ) trigger. The time delay between the light pulse and DAQ trigger was varied. Compare to the SS method later used in RICH commissioning in Section 4.2.1.

Figure 3.4 shows an example HPD IFB measurement result on the left and the IFB results of all HPDs on the right [26]. IFB photoelectrons are predominantly observed ~250 ns after the signal light pulse. This is the time it takes for an ionised gas particle to travel back to the photocathode after travelling from the most likely location for ionisation. The majority of HPDs had very low IFB when

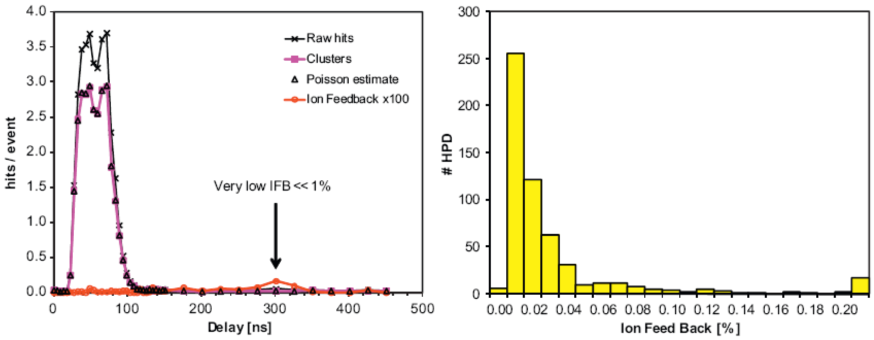

Figure 3.4: **Left:** Example HPD result for ion feedback, displaying hits and clusters at different time delays between signal light pulse and readout trigger. The red curve is multiplied by a factor of 100 to make it visible. **Right:** Distribution of ion feedback results for all HPDs. Both figures come from [26].

measured at PDTF in agreement with Photonis' measurements. The average IFB was 0.04% at PDTF. Compare to the <1% IFB specification in Table 3.1. How these IFB values developed over time is presented in Section 4.4.

3.3 QE Measurements

Quantum Efficiency (QE) is defined as the probability that an incoming photon of a given wavelength produces a photoelectron that gets detected. This factors in inefficiencies from reflections at the Air-Quartz and Quartz-Cathode interfaces as well as the thickness, absorptive qualities and work function of the photocathode. One PDTF station at Edinburgh was adapted to measure QE by comparing ratios of current readings between a well known, calibrated photodiode and the HPD from a stable light source:

$$\eta_q^{(hpd)}(\lambda) = \eta_q^{(pd)}(\lambda)\frac{I_{hpd}(\lambda)}{I_{pd}(\lambda)} \tag{3.1}$$

where η_q is the quantum efficiency.

3.3.1 Methodology

The requirement of a stable, well-known light source that can be configured for different wavelength filters meant the existing LED and LabView setup for standard tests could not be used. Figure 3.5 shows how the PDTF station was adapted for QE tests. An external housing was made for a Quartz-Tungsten-

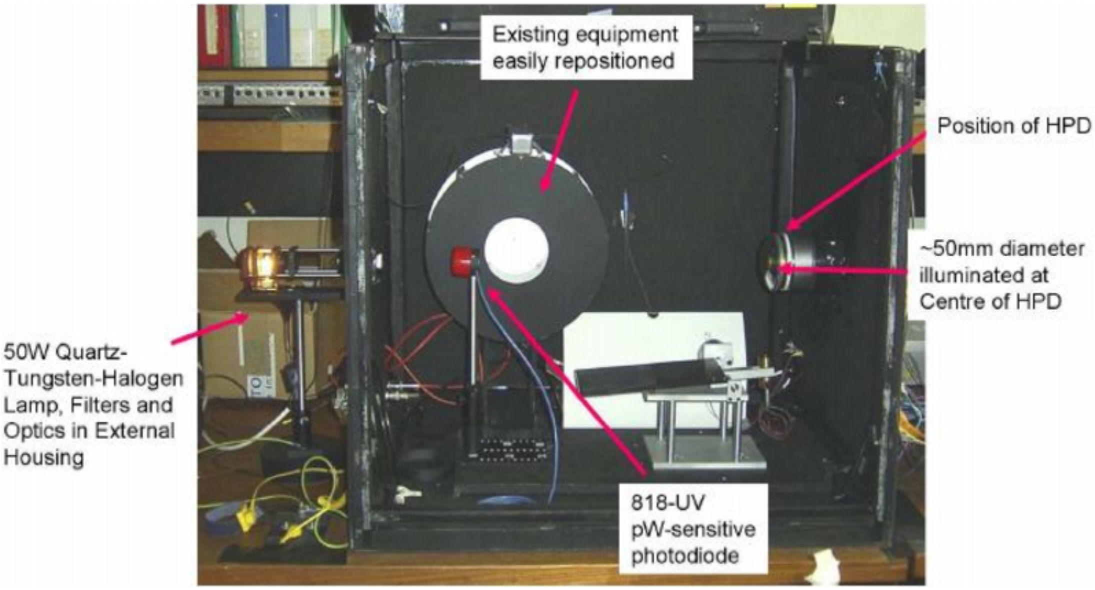

Figure 3.5: Photograph of the darkbox for the QE setup. The LED light source is easily repositioned, and a separate light source has been attached to the left. The photodiode provides calibration data before and after HPD measurements.

Halogen light bulb, filters and the necessary optics to focus the light appropriately. The dark box had a hole made in the side in which the final lens of the focussing optics was placed. The LED light source 'wheel' was temporarily placed aside to make way for the light source's rays. The test station could still be used for standard tests as light-tightness was not compromised by adding the external housing.

The 20 kV HV power supply was inappropriate for the manually performed QE tests: user safety was a concern and the picoampere meter to take HPD current readings with could not operate under this high voltage. A 100 V power supply was used instead and the silicon sensor was left unbiased. The wiring was changed so photoelectrons falling on any surface inside the HPD except the photocathode were read out using the picoampere meter. This includes the unbiased silicon sensor with its internal pixelisation rendered inactive. This test station configuration could be reverted to standard test mode when required.

QE measurements were done at seven wavelengths ranging from 270 nm to 800 nm. The light bulb has an emission spectrum of 220 nm to $>2\,\mu$m. Specific wavelengths were chosen with 10 nm bandpass filters, with other filters in front to block the harmonic multiples (notably the IR spectrum) and to bring down light intensity to HPD-safe levels. The light was focussed via the lens onto a 50 ± 2 mm

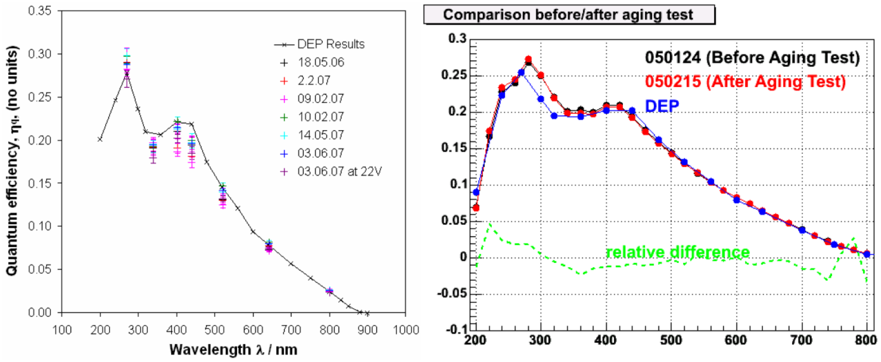

Figure 3.6: **Left:** Repeat measurements of QE taken over a year of HPD H527009. All measurements agree within errors. **Right:** QE measurements from the manufacturer (DEP) and at CERN before and after aging the HPD by 10 years' equivalent illumination [28]. As with PDTF, the QE at different wavelengths is plotted.

diameter spot at the centre of the HPD quartz window, i.e. covering ~50 % of the photocathode area. A set of measurements across the range of wavelengths was done with the photodiode before and after the HPD measurements, mounted in the focus of the beam to catch all the photons from the light source without changing the lens optics. The HPD was measured at 100 V and later, also at 22 V bias as will be explained in Section 3.3.3.

3.3.2 Repeatability & Aging Tests

Figure 3.6 shows the repeatability of PDTF QE measurements on the left. The numerous measurements taken over a year agree with each other and with the original manufacturer's values within experimental errors [27]. This compliments an early accelerated aging test performed at CERN on pre-series HPD H407008 [28].

Aging was simulated by exposing the HPD to a light-intensity 120 times larger than what is expected in the RICH detectors. In addition, the HPD was heated to ~50°C to enhance any out-gassing effect. This was done until the HPD had received illumination equivalent to 2000 days of LHCb operation. The QE measurements before and after this aging test is shown on Figure 3.6, right for a range of signal light wavelengths. There is no difference in QE after the aging simulation and the results match the manufacturer's measurements except for a small disagreement around 300 nm. This showed good promise for the stability of HPD performance over the course of LHCb's operational lifetime.

3.3.3 Comparison with Manufacturer's QE Measurements

The manufacturer tested for QE on every HPD at 20 different wavelengths from 200–900 nm. Like PDTF, a white light source and 10 nm band pass filters were used but without an IR filter to block out harmonic multiples. The HPD was biased to 900 V and was illuminated by a 25 mm diameter light spot at its centre [27].

Figure 3.7 plots Photonis and PDTF measures of QE. It compiles the first QE measurements of all HPDs that PDTF has tested at 100 V, including the special cases covered in Section 3.4. Most HPDs show good agreement between the two sets of measurements. A few HPDs had disagreements between Photonis and PDTF QE values, even with repeat measurements from PDTF. These were considered bookkeeping errors on the manufacturer's part.

Some HPDs had a higher PDTF QE measurement than the original Photonis QE value. This increase in QE was clearly seen when an old prototype HPD which was produced a year earlier was tested for QE, shown in Figure 3.8, middle. This difference in QE depending on bias voltage was puzzling at first.

A hypothesis was formed that IFB may be inflating the QE measurements. QE is calculated using the total current drawn from the HPD in response to continuous wave illumination, compared to a calibrated photodiode. If IFB is occuring in an HPD, secondary photoelectrons from IFB would contribute to the current readout from the HPD and increase the QE values.

At the time, helium was suspected to be the main vacuum contaminant. There are two possible causes for vacuum contamination: internal outgassing or diffusion of gas molecules through the quartz window, imperfections in the HPD body or its vacuum seal. The aging tests described in Section 3.3.2 made outgassing unlikely to be the cause.

Assuming diffusion is the cause for vacuum degradation, the most likely candidates are helium and hydrogen due to being the most diffusive gases. The HPDs were stored in nitrogen and dry air, where there is much less hydrogen than helium. Thus it was assumed the residual gas inside degraded HPDs was helium [27]. Evidence from more recent studies disagree with helium being the contaminant and is presented in Section 4.3.1. Note the study described here does not depend on the source of the vacuum degradation. The consequences of IFB are investigated, not the cause.

Helium's first ionisation energy is 24.6 V so QE measurements were made at

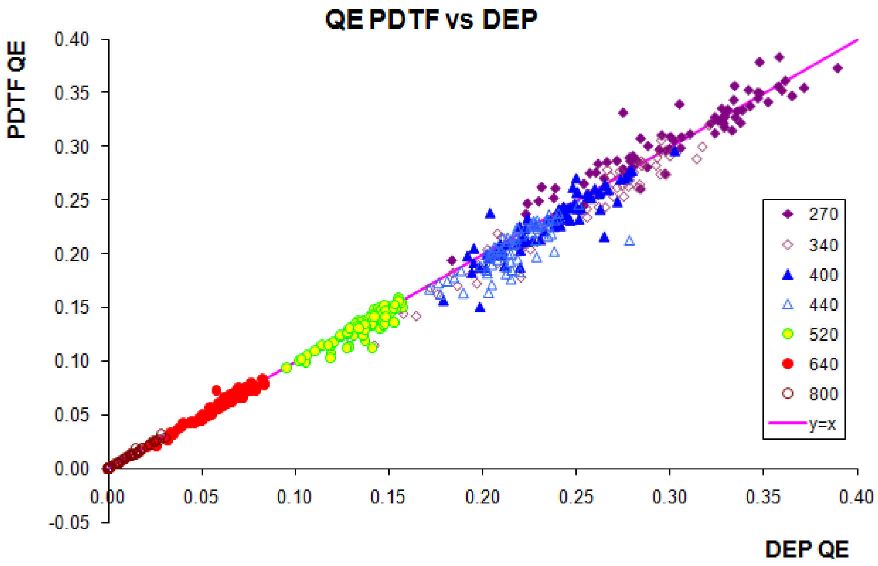

Figure 3.7: Correlation between Photonis and PDTF measurements of QE, across the seven wavelengths PDTF measured over. PDTF data points are the first QE measurements at 100 V bias. The pink $y = x$ line is for reference. Photonis did not measure QE at 340 nm directly, so an average between 320 nm and 360 nm was used for this plot.

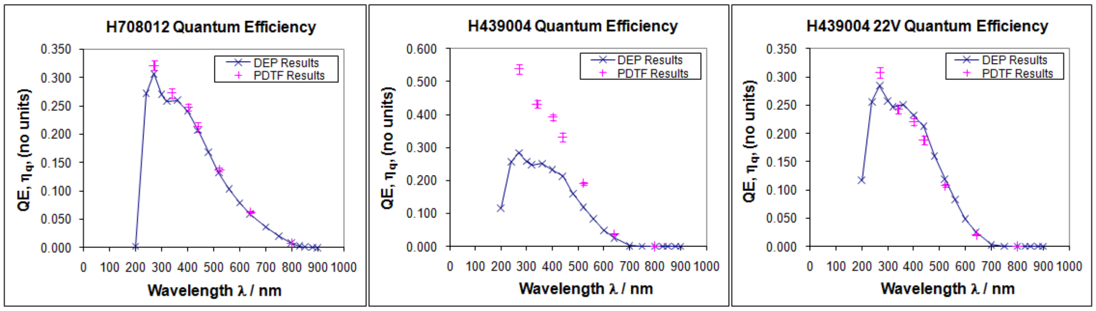

Figure 3.8: **Left:** An example of a normal HPD with good agreement between PDTF and Photonis' QE measurements. **Middle:** An old, pre-series HPD with much higher QE from PDTF's measurements than the manufacturer's values. **Right:** The same pre-series HPD with QE measured at 22 V bias instead of the usual 100 V.

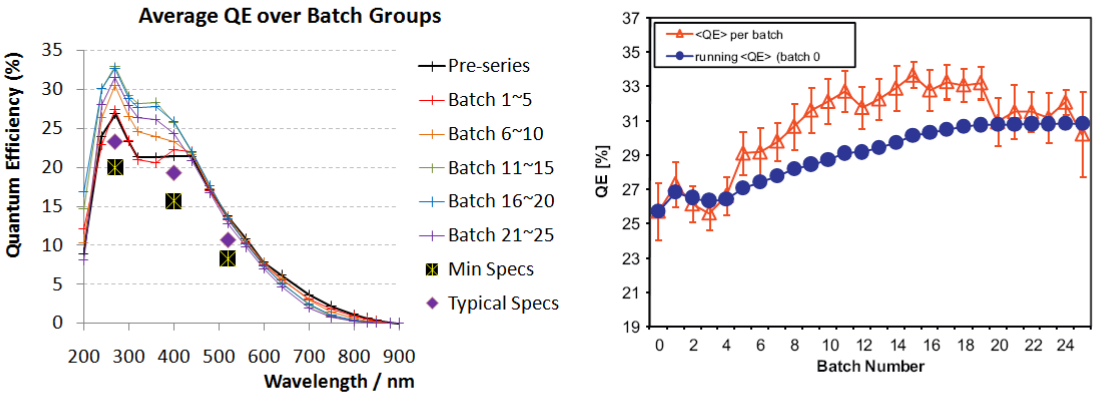

Figure 3.9: **Left:** Summary of QE measurements by Photonis, organised by groups of production batches. Minimum and Typical specs from Table 3.1 are also plotted for reference. **Right:** PDTF values of QE at 270 nm averaged by HPD batch. The running average is also plotted. Errors are RMS widths.

22 V bias to the prototype HPD. The QE results agreed with the manufacturer's data again, as shown in Figure 3.8, right. A new test procedure was setup to explore these effects which will be described in Section 3.3.4. Old prototype HPDs such as this one is discussed in Section 3.4.1.

The manufacturer's QE measurements of all HPDs are presented in Figure 3.9, left. The 25 production batches have been split into five groups and QE has been averaged across these groups. This shows the improvements to the HPD assembly procedure as production went on: QE was improved in the signal range of 200∼600 nm and infra-red noise sensitivity was reduced. In all cases, the HPDs exceeded the QE specifications.

On Figure 3.9, right the PDTF measurements of QE are presented, averaged by batch. The same trend of improvement until batch 20 can be seen. The reason for the apparent drop in performance in batches 20∼25 is because they include faulty HPDs from the earlier batches that underwent a recovery procedure. These recovered HPDs had their photocathode deposited much earlier than their delivery batch would suggest. Recovered HPDs are discussed later in Section 3.4.2.

3.3.4 QE IV Curves

Inspired by the difference in QE at 22 V and 100 V in Figure 3.8, a second QE investigation procedure was developed: the QE Current-Voltage (IV) test. It measures the current drawn from the HPD anode while varying the bias voltage.

The light source was the QE lamp housing as before, but with the filters kept unchanged throughout the test using a constant light intensity at a wavelength of 440 nm. A voltage range of 0∼500 V was used, the upper limit being decided by the maximum safe operating level of the apparatus. Measurements were taken both when ramping up and when ramping down with an hour's wait inbetween. The photocurrent is derived by subtracting the total current from the dark current.

Figure 3.10 shows the three broad categories of QE IV results: normal HPDs, degraded HPDs and severely degraded HPDs. The photocurrent increases sharply from 0∼5 V then flattens out in a plateau for normal HPDs: most of the QE of an HPD is attained by ∼5 V bias between photocathode and readout anode. These HPDs have negligible dark count.

Vacuum degraded HPDs such as the example shown in Figure 3.10b still reach a plateau by ∼5 V. However at ∼22 V the photocurrent increases with the bias voltage. As discussed in Section 3.3.3, this shows IFB starts to occur around helium's first ionisation energy. The HPD in this example has a dark count that increases significantly with the HPD bias voltage. Not all degraded HPDs exhibit high dark count as in this example: the photocurrent is more sensitive to vacuum degradation than the dark count. The current readings are different between ramping up and ramping down the bias voltage, both for dark count and for total current.

For severely degraded HPDs, the dark current becomes dominant and suppresses the photocurrent as shown in Figure 3.10c. The phenomenon occuring in these HPDs is discussed in Section 4.3. Such HPDs cannot be used in the RICH detectors. Different HPDs in this category had a different threshold in bias voltage for when the dark current became significant.

The example QE-IV measurements in Figure 3.10c waited for one hour between ramp up and ramp down, with the HPD kept at 500 V and with no illumination. The dark current settles over this hour, but still dominates the current reading. This is further explored with a second example shown in Figure 3.11. The currents in a QE-IV test of a degraded HPD is shown on the left plot and the same HPD after it became severely degraded is on the right. Currents at both ramp up and ramp down are plotted. An hour separates the two sets of measurements. As vacuum degradation gets severe, the dark current readings differ greatly between ramp up and ramp down. In all HPDs where this

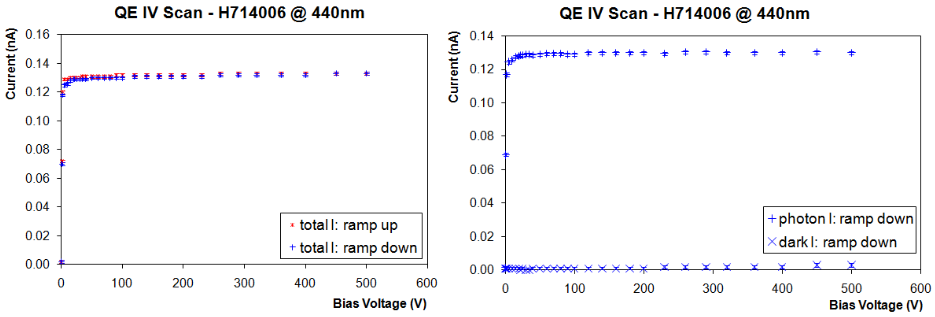

(a) A typical HPD.

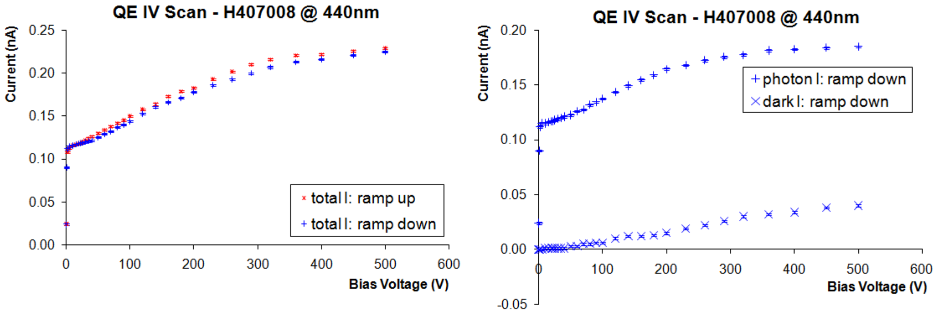

(b) A degraded HPD.

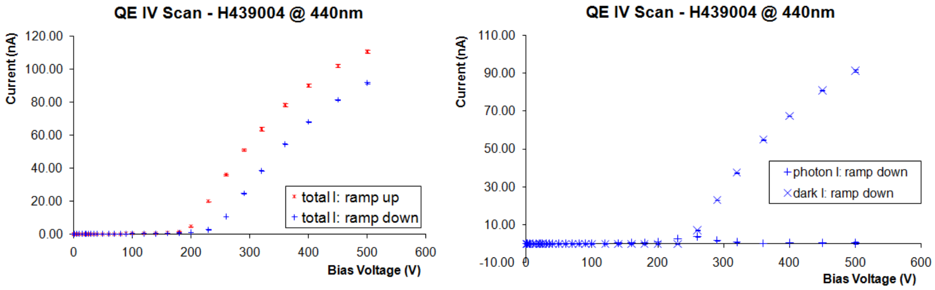

(c) A severely degraded HPD.

Figure 3.10: Quantum Efficiency I-V measurements for three example Hybrid Photon Detectors (HPDs) of different vacuum quality. The x-axis plots the bias voltage between the HPD photocathode and anode. The y-axis plots the current drawn from the HPD. Each HPD has two plots: the left shows the total current read out and the right shows the dark current and the noise-corrected photocurrent readings. Note the different y-axis scales.

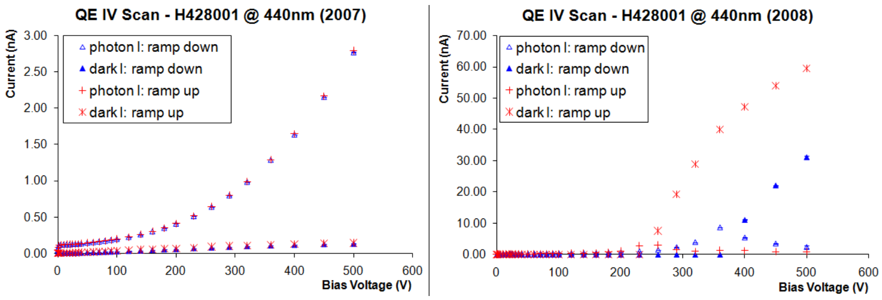

Figure 3.11: **Left:** Photon and dark currents of a QE IV test on a degraded HPD, both during ramp up and ramp down. An hour separates the two sets of measuremnts. **Right:** Photon and dark currents of the same HPD a year later, now severely degraded. Note the differnet y-axis scales.

difference is visible, the dark current decreased during the hour the HPD was kept at 500 V bias and with no illumination. The vacuum is 'annealed' by being kept at high voltage: the IFB effect subsides but is never completely nullified. This is further discussed in Section 4.4.2.

All HPDs underwent an IFB measurement as part of the PDTF standard tests. The QE-IV tests were developed in response to seeing IFB affect QE measurements. Therefore a correlation was expected between the two tests. The ratio of QE measurements at 100 V and 22 V bias is plotted against the PDTF IFB measurement in Figure 3.12. The IFB measurements were taken on or near the day of the QE-IV measurements. In this figure, 'Early Batch' refers to batch 5 or earlier (including prototypes) and 'Late Batch' HPDs come from the last 3 batches of HPD production. A positive correlation is seen between the two quantities plotted, as expected.

3.4 Special Cases

3.4.1 Pre-Series and Early Batch HPDs

Before HPDs were routinely manufactured, nine prototypes were produced to assess the suitability of HPD technology. These HPDs were referred to as 'pre-series' HPDs or 'batch 0'. Once mainstream HPDs started coming through, these pre-series samples were kept aside as references. A few HPDs from the earliest of production batches were also kept behind as references.

One such pre-series HPD was QE tested over two years after it first arrived at PDTF. This was the degraded HPD shown in Figures 3.8 and 3.10c, where almost

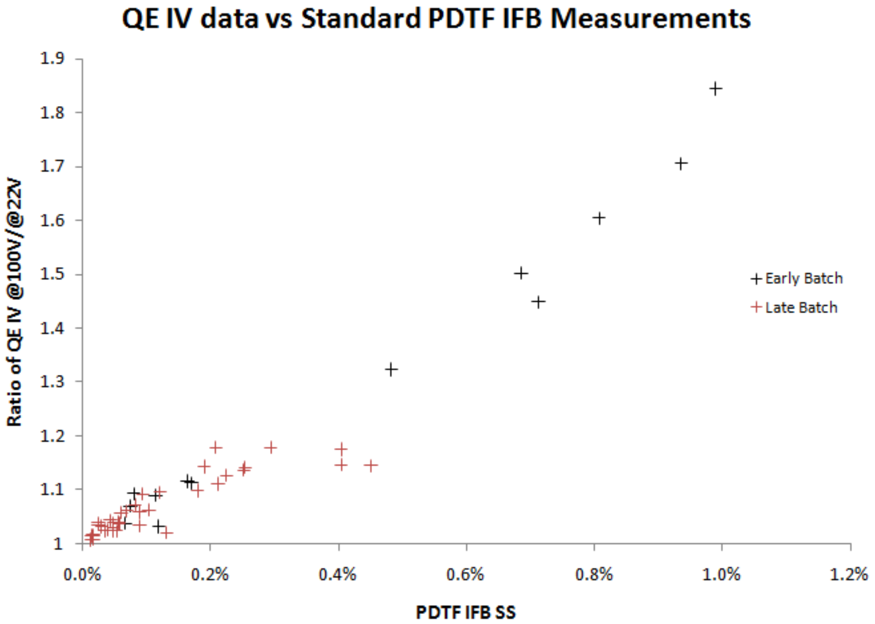

Figure 3.12: Correlation between QE IV results and PDTF ion feedback measurements. The ratio of photocurrents measured at 22 V and 100 V is compared with PDTF results. HPDs from early production batches are highlighted in a different colour.

3 years separate the manufacturer's measurements and PDTF's QE values. This led to all pre-series HPDs and all reference HPDs from early batches to be QE and QE-IV tested. By 2008 all nine pre-series HPDs had degraded of which seven had degraded severely as defined by the measurement of a dominant dark current at 500 V.

Of interest is that H407008, the HPD used in the aging tests described in Section 3.3.2, had also degraded: it is the example shown in Figure 3.10b. The aging tests had heated the HPD up to accelerate any outgassing sources and had found no increase in QE at the time. This supports the theory that the vacuum is being degraded from residual gases that diffuse into the HPD from outside. These pre-series HPDs spent 6 months in ordinary atmosphere before PDTF began storing HPDs in dry air. This could be a possible reason for their degradation. Further investigation of HPD vacuum contamination is presented in Section 4.3.1.

3.4.2 Recovered HPDs

A number of HPDs had a dead column of pixels at the edge of the readout chip. At the time, the number of available anodes for assembling HPDs was running low. The manufacturer devised a procedure to recover these dead columns while recycling the anode. This involved heating the HPD without breaking the vacuum seal. CERN agreed to accept these ‘recovered HPDs’ after testing each HPD twice at the PDTF centres, one month between the tests. Out of 15 recovered HPDs, only one failed the PDTF standard tests. No significant deterioration was observed when these HPDs were re-tested a month later. QE and QE IV tests were done for all these recovered HPDs.

The ratios of photocurrents measured in the QE-IV tests of recovered HPDs are presented in Figure 3.13. The distributions for two ratios are shown: photocurrents measured at 100 V & 22 V on the left and between 500 V & 22 V measurements on the right. Four of the 15 recovered HPDs have high values for these ratios, indicating compromised vacuums. One of these four was the recovered HPD that failed PDTF standard tests. The QE results of an example of the other three is shown in Figure 3.14. The manufacturer’s QE measurements agree with PDTF’s 100 V measurements instead of the QE at 22 V. Thus the degraded vacuum was present when they were recovered and measured by Photonis.

3.5 Conclusions

Over two years, the manufacturer produced and quality-tested 561 HPDs (including repairs and replacements) for the LHCb RICH detectors. PDTF cross-tested and characterised all of these HPDs, with QE tests performed on 10% of all produced HPDs. 97.1% of HPDs passed the PDTF tests and the remainder were repaired or replaced to meet the 550 HPDs ordered from CERN. Out of the $\sim$ 60 HPDs that PDTF tested for QE, only two showed a disagreement with the manufacturer’s values. These anomalies have been assessed to be bookkeeping errors.

Prototype and early production batch HPDs were observed to have unusually high QE. This was due to vacuum degradation causing IFB to occur in these HPDs, which increases the current drawn from the HPD. The increase in measured QE due to IFB was not seen in more recently manufactured HPDs at

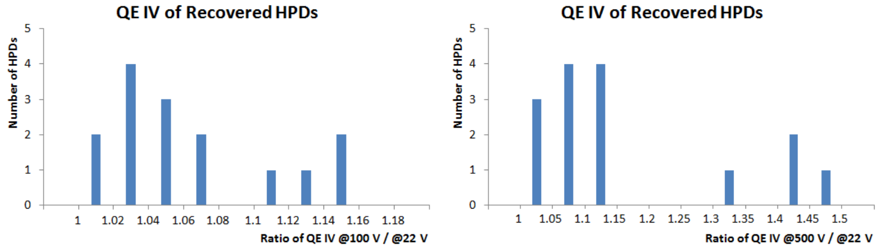

Figure 3.13: Ratios of QE-IV photocurrent ratios for HPDs recovered from missing pixel columns. The ratios are between 100 V & 22 V measurements on the left and 500 V & 22 V measurements on the right.

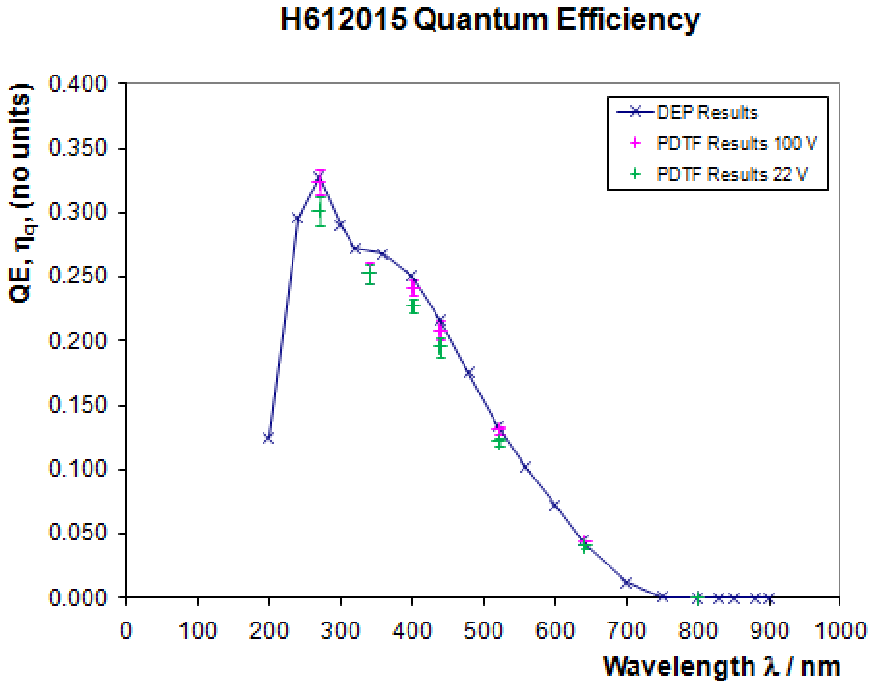

Figure 3.14: QE of an example HPD recovered from missing pixel columns but having vacuum degradation. This HPD passed PDTF standard tests with no problems.

the time of standard testing. The QE-IV test procedure was developed to further investigate this phenomenon. A correlation has been seen from the results of this new test with the existing IFB measurments in the PDTF standard tests. The QE-IV test procedure became a sensitive tool for revealing the degree of HPD vacuum degradation.

A number of HPDs had missing readout pixel columns. The manufacturer developed a process to recover performance of these missing pixels. PDTF tested these recovered HPDs with extra attention to assess the manufacturer's new process. 14 out of 15 recovered HPDs passed the PDTF standard tests and were accepted. The QE-IV tests found three of these HPDs to have vacuum degradation.

Chapter 4

Monitoring the LHCb RICH2 Sub-Detector

The Ring Imaging Cherenkov (RICH) sub-detectors are constantly monitored by various methods to ensure performance meets the strict requirements of the LHCb experiment. The Magnetic Distortion Monitoring System (MDMS) collects data to make corrections for the residual magnetic field getting past the Hybrid Photon Detector (HPD) shielding box and μ-metal cylindrical covers. Mirror alignment is ascertained by the Laser Alignment Measurement System (LAMS) [29]. In addition, HPD response, light leak checks, radiator gas composition and voltage supply conditions are all monitored and logged. This work began during the commissioning phase when the sub-detectors of the LHCb were being installed, before LHC started circulating beams.

The HPD monitoring procedures were first developed when RICH2 HPDs finished mounting, which will be described in Section 4.1. Monitoring the HPD vacuums via ion feedback monitoring is described in Section 4.2. A number of HPDs had severe vacuum degradation, leading to a glow effect. The nature of this issue and why these HPDs needed to be replaced is discussed in Section 4.3. A model to describe the rate of HPD vacuum degradation was developed to help plan for replacing severely degraded HPDs. How this model was made based on monitoring data is presented in Section 4.4. The predictions made by this model for when HPDs are expected to be at risk of needing replacement is discussed in Section 4.5.

4.1 HPD Monitoring

After individual testing at the Photon Detector Test Facilities (PDTF) as described in Chapter 3, the Hybrid Photon Detectors (HPDs) were delivered to CERN to be mounted into the RICH detectors. RICH2 was fitted with HPDs first, finishing installation in early 2007. The HPDs are mounted onto columns integrated with front-end electronics, power supplies and cooling. RICH1 columns contain 14 HPDs and RICH2 columns, 16 HPDs. The columns are staggered inside the RICH detectors to achieve hexagonal close-packing of HPDs. Each RICH1 (RICH2) detector half contains seven (nine) columns. Each column was fitted with HPDs, tested as a group and installed one-by-one over several months. Most HPDs received no further monitoring once inside RICH2 until all columns were mounted and the Experiment Control System (ECS) was ready.

Two light sources are available for monitoring the mounted HPDs: a Continuous Wave (CW) diode emitting at ∼635 nm and a 633 nm pulsed laser for Strobe Scans (SS). Both sources are installed behind the shielding where the high level readout electronics are located, thus do not need to be radiation hard. Optical fibres transport the light to illuminate the RICH detector planes.

Readout of the HPDs is synchronised with the Data Acquisition (DAQ) trigger. When the DAQ trigger is given, all pixels readout any signal they receive in a 25 ns period which forms one event. This 25 ns length is to accommodate the 40 MHz collision frequency of the LHC. For SS monitoring runs, the DAQ trigger can be synchronised to be sent a fixed time before or after the laser's pulse in each event. This captures snapshots of HPD response a set time after each pulse of signal photons. Time alignment of HPDs was monitored using this setup, to adjust for differences in response time between the HPDs.

Figure 4.1 shows a pixel hitmap of RICH2 HPDs and labels the nomenclature for HPD positions. The data comes from a CW light source run. This plot is not an accurate representation of the physical locations of photon hits: the hexagonal close packing of HPDs is not shown and photons can only reach the inscribed circular area of the square silicon sensor area. Most hits beyond the circle area stem from back-scattered photoelectrons which hit the silicon sensor a second time. The display convention is as if the observer was standing behind the RICH2 detector plane and was seeing incoming photons hit the photocathode while looking towards the interaction point. Photoelectrons would be crossfocussed by the electron-optics afterwards but is not shown on this hitmap. The two

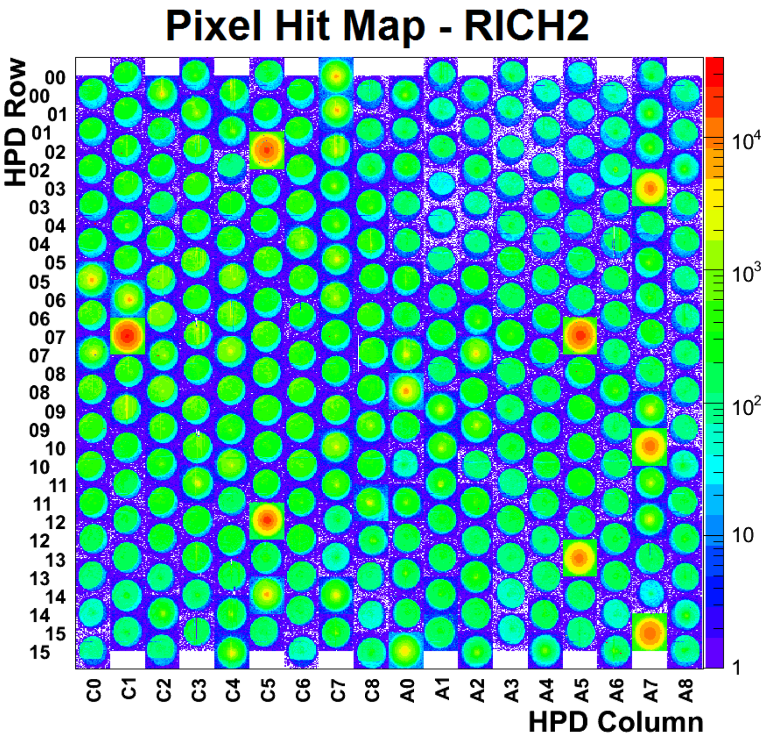

Figure 4.1: Pixel Hitmap of the 288 HPDs in RICH2. The data are from 3 million events collected with a continuous wave light source in Run 23288 (23 May 2008). Both halves of RICH2 are shown in this hitmap: each RICH2 half has nine columns, labelled C0–C8 and A0–A8 respectively. There are 16 HPDs in each column, labelled 00–15 counting from the top of each column. Logarithmic z-axis represents how many total hits each individual 500 μm$\times$ 62.5 μm pixel accumulated over the run.

halves of RICH2 have been brought together for convenience of viewing. It is stressed: a separate pixel viewing mode is used for ring finding. This plot is for error-spotting and HPD quality monitoring.

A handful of HPDs can be seen to have unusually high activity. These are HPDs with severely degraded vacuums like the third category of QE IV results presented in Section 3.3.4. They will be further discussed in Section 4.3. The HPDs generally exhibit a crescent-shaped shadow, especially towards the top. This is due to the fixed light source having a shadow cast by the cylinder of the μ-metal shielding of each individual HPD.

4.2 Ion Feedback Measurements

The vacuum quality of HPDs is monitored by making regular Ion Feedback (IFB) measurements. As described in Section 3.2.2, increased IFB occurs in HPDs with degraded vacuums.

Central to defining IFB measurements is a hit cluster: a group of adjacent pixels (diagonals included) that all got hit in the same event. Each of the 8192 pixels of an HPD has its own readout for IFB analysis as opposed to physics analysis which groups eight pixels with an OR. A single photoelectron from a clean signal should not hit more than four pixels. Clusters of $\geq$5 hits are treated as IFB incidents. The rate of IFB is defined as:

$$R_{IFB,raw} = \frac{N_{clusLarge}}{N_{clusAll}} \tag{4.1}$$

where $N_{clusLarge}$ is the number of clusters with $\geq$5 pixels and $N_{clusAll}$, the number of clusters of any size (including single hits) over an entire data collection run.

The above value needs to be corrected for dark count. The analysis software cannot distinguish whether hits are from noise or signal: only cluster sizes are factored into the R_{IFB} calculation. Therefore an IFB data collection run using a light source must be accompanied by a dark count IFB run taken within a few days to keep conditions similar between the two runs. Ideally the two runs should be back to back but detector commissioning activities made this impractical. The corrected rate of IFB is then:

$$R_{IFB,corr} = \frac{N_{clusLarge} - M_{clusLarge}}{N_{clusAll} - M_{clusAll}} \tag{4.2}$$

where M represents number of clusters measured from a dark count IFB run and N counts clusters in the illuminated IFB run. M and N come from runs with the same number of events.

This is visually represented in Figure 4.2, where the top two hitmaps are from the same run as Figure 4.1 but filtered to only show pixel hits from IFB and non-IFB classified clusters. The top-right plot of Figure 4.2 shows clusters counted in $N_{clusLarge}$. $N_{clusAll}$ counts clusters in both the top plots. The bottom two hitmaps of this figure plot data from a dark count run taken a day before the laser signal run. $M_{clusAll}$ counts clusters from both bottom hitmaps, of which those in the bottom-right plot are also counted by $M_{clusLarge}$.

All HPDs show a degree of IFB as every HPD has at least some clusters plotted on the top-right plot of Figure 4.2. Aside from HPDs with high R_{IFB}, the response to the laser light source is uniform across the circular photocathode regions as seen in the top-left plot. HPDs with high R_{IFB} are clearly visible on both the top-right and top-left plots. This is an indication that a few secondary photoelectrons from a single IFB incident can be registered as separate small (≤ 4) clusters because there are no neighbouring hits joining them to the main cluster.

The bottom-left plot of Figure 4.2 has some HPDs with exceptionally low noise. None of these HPDs show up in the bottom-right plot as expected. All activity in the bottom-right plot is confined to the circular photocathode regions of the pixel chips[1]. This shows IFB in HPDs with degraded vacuums continues to occur in the absence of a light source. As in the top plots, these HPDs have both small and large pixel cluster activity. Compared to the top-left plot, these HPDs have much more square distributions of hits in the bottom-left plot. This shows noise in high R_{IFB} HPDs affects the corners of the pixel chips unlike signal photoelectrons, although the centre pixels remain more likely to be affected. This phenomenon is further discussed in Section 4.3.

4.2.1 Continuous Wave vs Strobescan IFB

IFB was measured with both light sources described in Section 4.1: the CW diode and the pulsed laser. The latter was used for SS IFB monitoring similar to the IFB measurements performed by PDTF as described in Section 3.2.2. Hereafter, the dark-count corrected rate of IFB as measured by these two light sources will be referred to as R_{CWIFB} and R_{SSIFB}. CW IFB measurement runs are much quicker to perform and analyse, which made them easier to fit in during the commissioning and after LHC becomes operational. However, the PDTF IFB data taken for every HPD during quality cross-checking was measured with a SS procedure. Thus it was important to establish the relationship between R_{CWIFB} and R_{SSIFB} as measured inside the RICH to justify using CW runs for regular monitoring. Comparisons with the PDTF IFB measurements is discussed in Section 4.4.4.

Data collection for a SS IFB measurement requires multiple runs, each with a different DAQ trigger delay after signal. Let t be the delay between the laser

---

[1] The exception to this are the HPDs at A1-14 and A1-15, where the noise readout threshold settings needed adjustment to better reject noise activity.

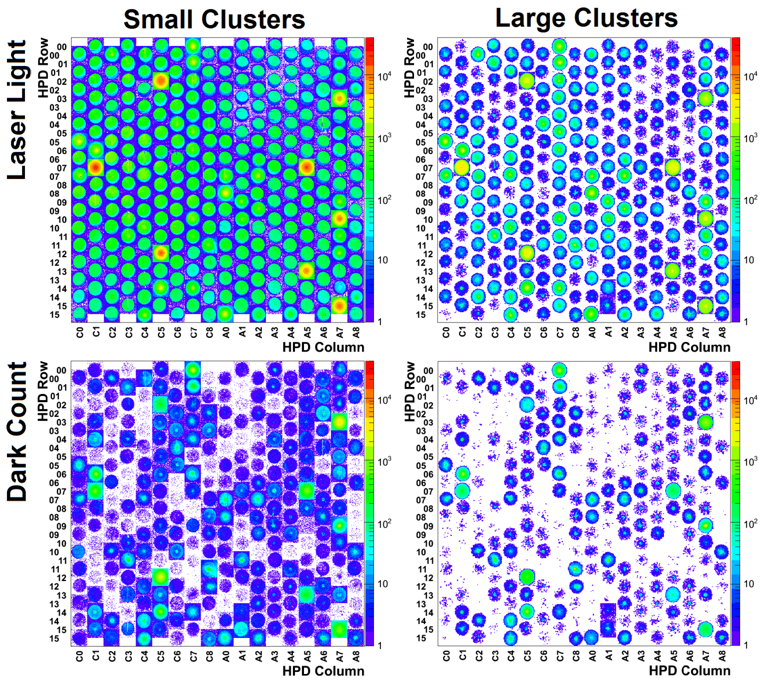

Figure 4.2: Pixel hitmaps from two runs, filtered in two ways. **Top:** Continuous wave laser illuminated - Run 23288 (23 May 2008) **Bottom:** Dark count - Run 23232 (22 May 2008) **Left:** Only pixel hits in clusters of size ≤ 4 are displayed. **Right:** Only pixel hits in IFB clusters (≥ 5 hits) are displayed.

pulse and the DAQ trigger. To measure the corrected IFB at a given t, two control test groups are required: a run with $t = -50\,\text{ns}$ measures the dark count noise and a run with $t = 0\,\text{ns}$ becomes the reference for counting primary, non-IFB photoelectron hit clusters. The modification to Equation 4.2 to calculate $R_{IFB,corr}$ in a set of SS measurements becomes:

$$R_{SSIFB,t} = \frac{N_{clusLarge,t} - N_{clusLarge,t=-50}}{N_{clusAll,t=0} - N_{clusAll,t=-50}} \tag{4.3}$$

where N counts the number of clusters from runs with the same number of events and t is in units of ns.

An example of a set of such measurements with 3 million events per run for one HPD is shown in Figure 4.3. The peak in R_{SSIFB} at 250 ns matches the timing shown in the left plot of Figure 3.4 in Section 3.2.2 and is typical for most HPDs. The dark count was taken before and after the SS runs to check for system stability. Note the R_{SSIFB} is lower at a 300 ns delay compared to 100 ns delay. The time delay between a residual gas particle getting ionised by the signal photoelectron and the gas ion reaching the photocathode to release secondary

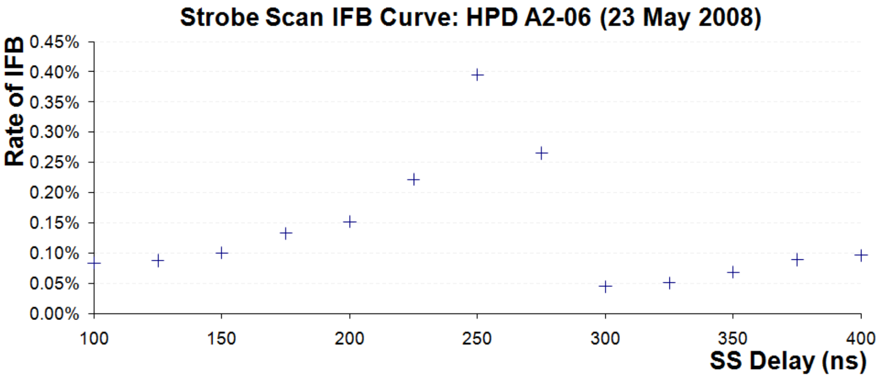

Figure 4.3: Strobescan measurements of the rate of ion feedback (IFB) of HPD A2-06 in RICH2, using the pulsed laser. This HPD has a medium IFB rate. The IFB rate on the y-axis is corrected for dark count. The x-axis is the time delay between the laser pulse and the data acquisition trigger. This set of measurements was taken from inside RICH2 during commissioning.

photoelectrons is consistent with the time difference between primary signal and secondary IFB readout. 250 ns corresponds to the most probable location for ionisation but it can happen earlier, leading to a fraction of IFB having a much shorter delay. R_{SSIFB} gradually increases as the delay is extended from 300 ns to 400 ns. This is an observation of a tertiary response, where an IFB photoelectron causes further ionisation.

IFB measurements with the diode use constant illumination and the DAQ trigger running at 40 MHz. The continuous light signal causes continuous IFB at random phases to the trigger. Thus CW IFB measurements give a time-integrated rate whereas SS IFB measurements give 25 ns wide samples of the IFB time profile. For comparison, a CW IFB measurement was taken on the same day as the SS IFB runs plotted in Figure 4.3. For the HPD shown in the figure, R_{CWIFB} was measured to be 1.78%. Summing up all the R_{SSIFB} values of this HPD from delays of 100 ns to 400 ns, the time-integrated rate of IFB, $\sum R_{SSIFB} = 1.54\%$.

$R_{CWIFB} > \sum R_{SSIFB}$ is expected for all HPDs since CW IFB runs integrate over all IFB responses. This is the case for most HPDs, as shown in Figure 4.4 which plots these two quantities against each other. At low IFB, the two quantities are well correlated and $\sum R_{SSIFB}$ falls behind at high R_{CWIFB}. A few HPDs seem anomalous by having $\sum R_{SSIFB} > R_{CWIFB}$: these are hyperactive HPDs such as A7-03 seen in Figure 4.1. They are discussed in Section 4.3.

To explore the relationship between R_{CWIFB} and $\sum R_{SSIFB,100<t<400}$ further, the ratio of the two, $\frac{\sum R_{SSIFB}}{R_{CWIFB}}$ is plotted against R_{CWIFB} in Figure 4.5. This

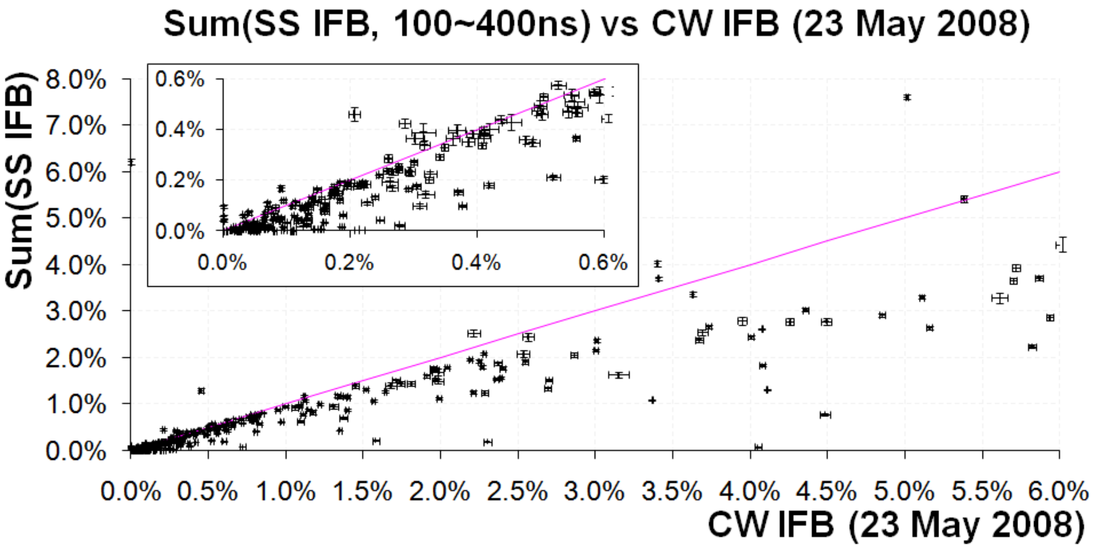

Figure 4.4: A comparison between the two methods of measuring Ion Feedback (IFB) rate used in RICH2. The results were all taken on the same day and have been corrected for dark count noise. The x-axis plots the results from the continuous wave diode light source. The y-axis plots the sum of all IFB rates measured in a set of strobe scan runs using the pulsed laser light source. The strobe scan runs measured IFB with a range of delays from 100 ns to 400 ns between the light pulse and the data acquisition trigger. The pink line is a $y = x$ reference line. The inset on the top left is a zoom of the plot closer to the origin.

plot can be categorised into three regions on the x-axis: $R_{CWIFB} < 0.3\%$, $0.3\% < R_{CWIFB} < 1.0\%$ and $R_{CWIFB} > 1.0\%$. In the first R_{CWIFB} region, $R_{IFB,corr}$ becomes much more sensitive to noise clusters, hence the scatter in the y-axis. In the middle, most of the $RIFB, corr$ seen in the CW run is observed in the SS measurements when summing over $t = 100 \sim 400\,\text{ns}$. As R_{CWIFB} gets higher, a smaller fraction of it is seen in $\sum R_{SSIFB}$ over this range of t. Cascade IFB effects where photoelectrons from IFB themselves cause further IFB will be observed at $t > 400\,\text{ns}$ so are only seen in the CW result.

Since IFB measurements were made for monitoring HPD vacuum quality, the CW diode was established as the preferable light source for the task. Besides the difficulty of performing a full set of SS IFB measurement runs in a busy commissioning schedule, R_{CWIFB} is more reliable at high $R_{IFB,corr}$: CW measurements include cascade IFB effects which the SS methodology would require increasingly more runs at later t to observe. The slight discrepency seen in some very low $R_{IFB,corr}$ HPDs between R_{CWIFB} and R_{SSIFB} is not a concern as both methods are still clear on the vacuum quality being good.

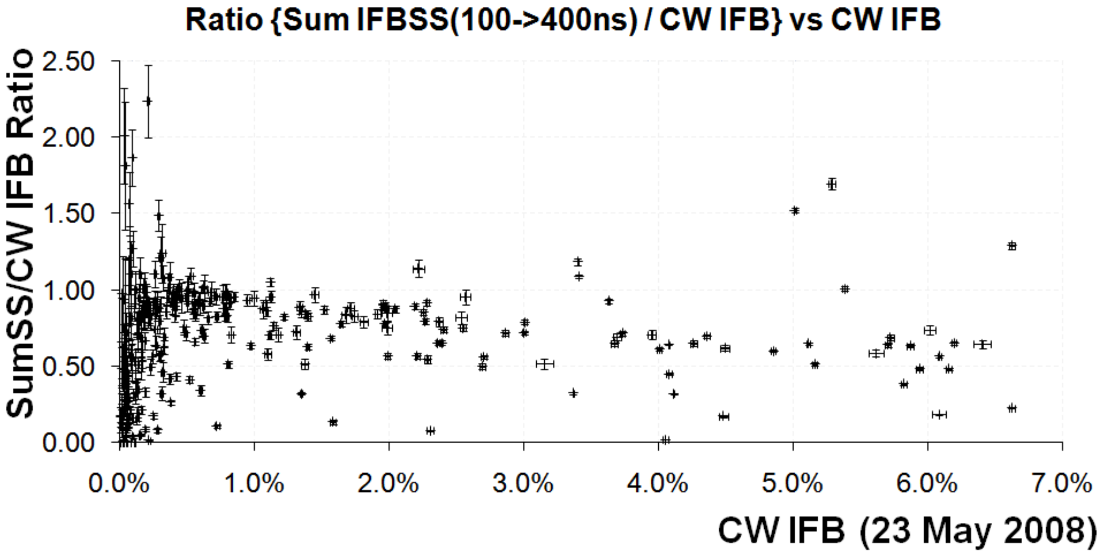

Figure 4.5: Ratio comparison between the two methods of measuring Ion Feedback (IFB) rate. The x-axis plots the IFB rate as measured with the Continuous Wave (CW) diode. The y-axis plots the ratio between the sum of all strobe scan IFB rates in a set of runs and the CW diode's IFB rate measurement. The strobe scan measurements set has delays ranging from 100∼400 ns between the laser pulse and the data acquisition trigger.

4.3 HPD Failure at High IFB Rates

4.3.1 The Glow Effect

When RICH2 was first tested with HPDs installed, a number of HPDs behaved unacceptably. In Figure 4.1, HPDs C1-07 and A5-07 are examples of such problematic HPDs. This high activity was confirmed in the laboratory to be HPDs producing light inside the tube bodies, detectable with external photon detectors. They are "glowing" HPDs.

This phenomenon is a consequence of severe HPD vacuum degradation. All vacuums degrade over time, but some HPDs were degrading much faster than expected. When HPDs were chosen as the technology for the RICH photon detectors, various risks were identified but fast vacuum degradation was not one of them [30]. IFB occurs in vacuum degraded HPDs and as was discussed in Section 3.2.2, IFB becomes self-sustaining in severely degraded HPD vacuums. 'Glow light' is emitted from self-sustaining IFB. A spectroscopy on this light [31] gave results compatible with the presence of hydrogen. There was no evidence for helium, disproving the hypothesis described in Section 3.3.3. The cause of the

Figure 4.6: Light from a glowing HPD, visually observed in a dark room. The image is superimposed with a photograph of an HPD with external lighting to show the location of the light source. The image has been colour edited for clarity. This figure comes from [31].

vacuum leaks leading to glowing HPDs are not understood.

At first, an HPD will start to glow when High Voltage (HV) is being ramped up from a few kV and stop glowing as HV goes past 10∼15 kV but exhibit significant noise during operation. As the vacuum further degrades, the threshold where glowing stops continues to rise until the HPD is glowing even at the nominal operational HV bias of 18 kV.

When glowing, an HPD emits light from inside which liberates photoelectrons from the photocathode just as signal photons would. The HPD's electric field then accelerates these photoelectrons to the silicon sensor anode. The light emitted by glowing HPDs is visible to the naked eye in a dark room, as shown in Figure 4.6. The number of photoelectrons generated by the intensity of glow light overloads the silicon sensor. This draws a large current from the silicon bias power supply. A series resistor limits how much current can be drawn for the nominal operational voltage of 80 V. As the glowing gets worse, the current drawn from the silicon bias is saturated and the HPD's silicon detector becomes under-depleted. Since all HPDs in a column share the same silicon bias power supply, it is essential that glowing HPDs are removed before they compromise their column [31].

Not all high R_{IFB} HPDs could be replaced with spares. 550 HPDs were manufactured of which 484 are mounted inside the RICH detectors and the remainder were spares. Of the remaining 56 HPDs, nine were pre-series

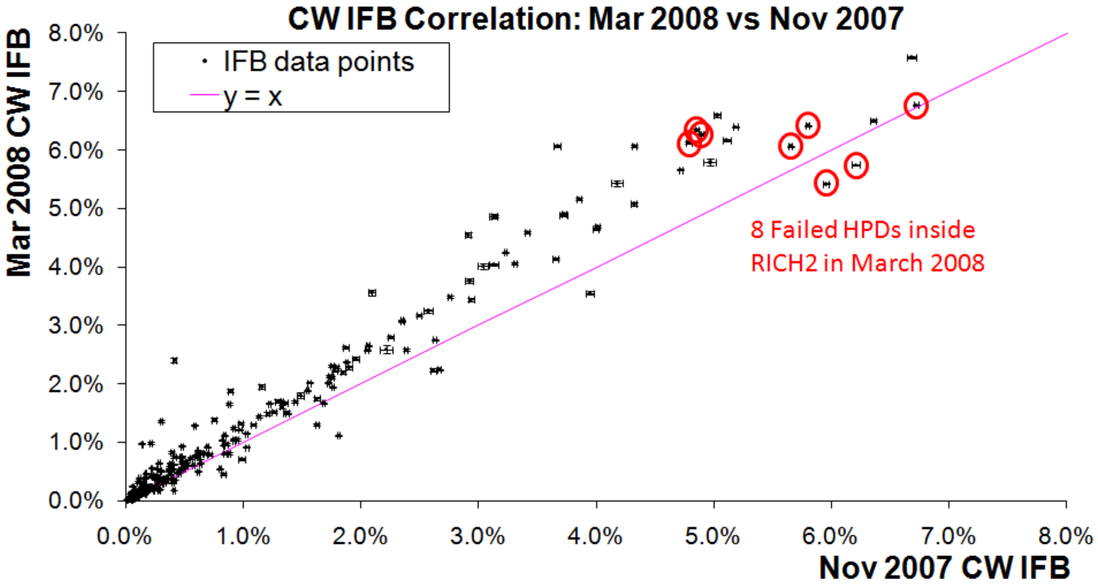

Figure 4.7: Correlation between two measurements of ion feedback rate using the continuous wave light source. The runs are three and a half months apart and have 3 million events each. The data has not been corrected for dark count. The pink line is $y = x$ and is drawn to guide the eye. The eight HPDs glowing in 5 March 2008 are circled in red.

prototypes which are not used; nine were kept behind at PDTF as reference tubes to continue doing periodic tests on over the long-term; and a few others had degraded before they could be used. A full breakdown of the numbers is given in [31]. The number of HPDs with high R_{IFB} exceeded the number of spares to replace them with. Preparing additional HPDs by repairing existing ones in the glow regime was taking a long time to be negotiated.

Physical access to the RICH detectors is limited, further constrained during beam collisions due to radiation safety. Maintenance shutdowns are planned to be annual. In addition, RICH1 is much harder to open up to make HPD replacements after installation. Therefore a model to predict when HPDs would start glowing was in demand. Less urgent HPDs needed to be identified which could be safely kept inside the RICH for the short-term while replacement HPDs became available after thorough PDTF testing.

4.3.2 Threshold for Risk of Failure

Figure 4.7 is from an initial study of the problem, where two sets of IFB measurements three and a half months apart are plotted against each other.

HPD Serial Number	Position in RICH2	Rate of IFB
H521008	A0-15	6.70% ± 0.021%
H525015	A0-08	5.82% ± 0.015%
H527003	A5-13	6.03% ± 0.016%
H539003	A7-10	6.04% ± 0.012%
H542012	A4-15	5.44% ± 0.038%

Table 4.1: Rate of Ion Feedback (IFB) as measured at the end of March 2008 for the five Hybrid PhotonDetectors (HPDs) that began to glow two months later. The IFB rate was measured with the continuous wave diode and is not corrected for noise.

Due to a lack of suitable dark count runs to accompany the CW runs, the data in the plot is not corrected for noise. When the two $R_{CWIFB,raw}$ measurements are compared, HPDs starting at high $R_{CWIFB,raw}$ are observed to degrade more over this time period. The error bars are purely statistical errors and the outliers suggest the main source of errors is differences in the environment. This is inevitable with commissioning constantly bringing the detector closer to running condition.

The eight glowing HPDs at the time of the second measurement in Figure 4.7 are all very high $R_{CWIFB,raw}$ cases. Glowing HPDs already replaced at the time were similarly high in $R_{CWIFB,raw}$. A hypothesis was proposed that all HPDs which start to glow have $R_{CWIFB,raw} > 4.8\%$. Hereafter it will be referred to as the 'glow threshold hypothesis'.

Within three months of this preliminary study, a further five HPDs started to glow. The $R_{CWIFB,raw}$ of these HPDs as measured at the end of March 2008 is shown in Table 4.1. These HPD failures supported the glow threshold hypothesis. Subsequent HPDs that began to glow had all been marked as being at risk by the above condition. HPDs may not glow as soon as they reach the 4.8% threshold, but that is the point where they are at risk to start glowing. None of the HPDs entering glow regime were below the threshold at the time of failure.

This preliminary study could only show that there was a clear increase in $R_{CWIFB,raw}$ over a few months. Whether the rate of vacuum degradation was linear or exponential was a concern. This motivated more frequent IFB measurements to be fitted into the commissioning schedule.

4.4 Timeline of HPD Vacuum Degradation

4.4.1 Methodology

The glow threshold hypothesis was formed when IFB measurement runs were scarce: back then there was no dark count run to use for noise corrections. In the following months, more frequent IFB measurement runs were performed. The urgent demand for a model to predict when HPDs would start glowing required all available IFB data to be used. This included the earliest IFB measurement runs where no dark count correction data are available. Therefore the model was developed using $R_{CWIFB,raw}$ as there was insufficient time to collect many $R_{CWIFB,corr}$ measurements.

For each HPD, $R_{CWIFB,raw}$ measurements were plotted against time to see how its vacuum quality changed over the course of months. These plots shall be referred to as 'timelines'. Examples of timeline plots are shown in Figure 4.8 for HPDs with vacuum quality ranging from very good to degrading. Time was measured as the number of days between the IFB measurement and when the HPD first arrived at the PDTF centres. The HPDs were delivered to the PDTF after being manufactured so this date is a good approximation to their age. As the HPDs were delivered in batches, the 288 HPDs in RICH2 have different ages for the same IFB measurement run. Using the PDTF measurement of IFB in the timeline plots is discussed in the next section.

The timeline plots put to rest the concern that the rate of vacuum degradation might be exponential. As shown in Figure 4.8, the $R_{CWIFB,raw}$ data points increase linearly over time. Most HPDs had a timeline plot which were compatible with a linear fit.

The earliest IFB measurement runs early in the commissioning had a slow rate of data recording. The necessary computing infrastructure to take data at a faster DAQ trigger rate only became available for the later runs. The 3 million events in these early runs shown in Figure 4.7 took a long time to gather. The timeline plots showed 3 million was sufficient for insignificant statistical errors compared to systematic errors. Therefore further data points were measured with the same 3 million event length.

The χ^2 values in these sample plots indicate the IFB measurements are dominated by systematic uncertainties. There are many sources for systematic errors ranging from software configurations, electronics conditions and differences

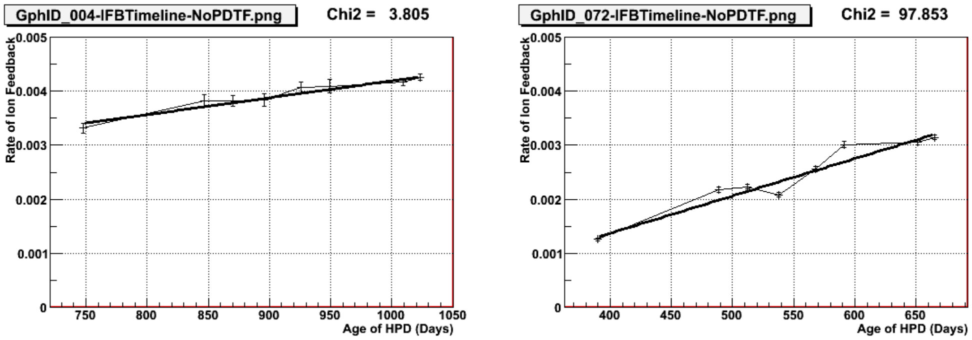

(a) HPD A0-3

(b) HPD A4-7

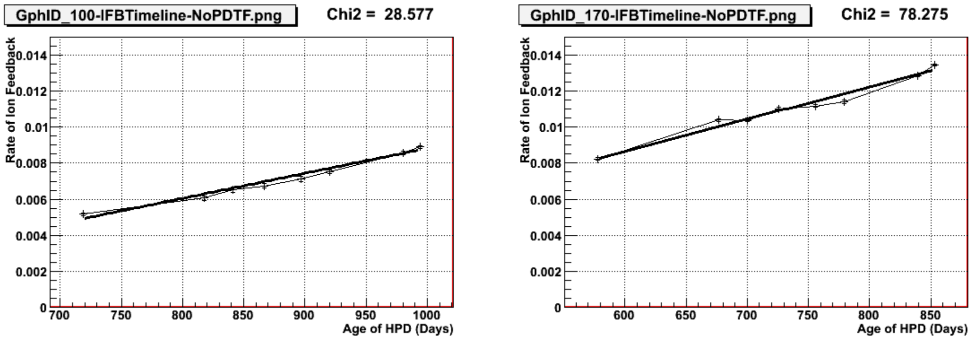

(c) HPD A6-3

(d) HPD C1-9

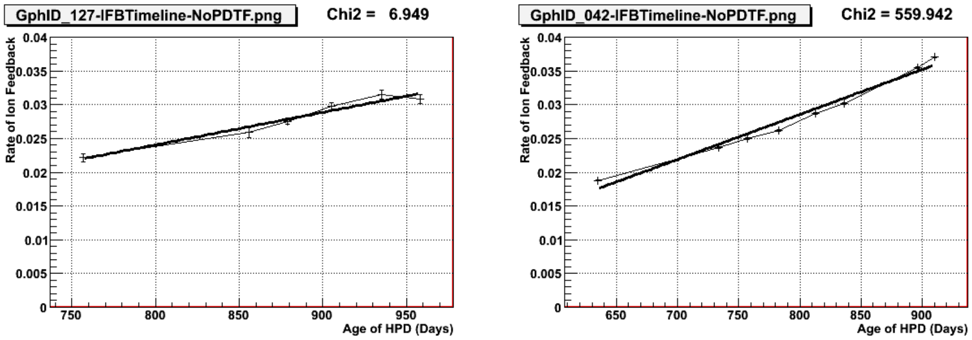

(e) HPD A7-14

(f) HPD A2-9

Figure 4.8: Examples of ion feedback (IFB) development timelines from RICH2 Hybrid Photon Detectors (HPDs). These IFB rates plotted on the y-axis were measured with the continous wave light source without dark count correction. The age of each HPD plotted on the x-axis is measured in days since they were tested at the photon detector test facilities. The top, middle and bottom pair of plots are sample HPDs with low, medium and high rates of IFB respectively. A linear fit has been made, with χ^2 values printed at the top right of each plot. Note the different y-axis scales.

in the environment over the months. As commissioning of the RICH, the other LHCb subdetectors and the detector control software had to continue, methodically eliminating the systematic errors was unfeasible. The distribution of χ^2 values of linear fits in the RICH2 HPD timeline plots is further discussed in Section 4.4.4.

4.4.2 Vacuum Quality Recovery Cases

Separate to the IFB monitoring for the HPDs mounted in the RICH, there was a long-term IFB study at CERN on two test HPDs [31]. The study observed a decrease in R_{IFB} over time when the HPD was subject to high illumination while under constant HV operation. This is a consequence of each residual gas particle getting involved in only a single IFB incident: after getting ionised by signal photoelectrons, the gas particles travel back to the photocathode where they release many secondary photoelectrons. These gas particles are then no longer available to be ionised again by photoelectrons. Thus this is a vacuum quality improving factor which counteracts against the vacuum degradation source.

The test HPD shown in [31] required a certain level of illumination before a decrease in R_{IFB} was observed. This level of illumination is comparable to what is expected in some regions of the RICH photon detector planes under nominal operating conditions. Therefore this vacuum quality improving phenomenon is a significant factor in how vacuum quality changes over time when the LHC is colliding beams.

Furthermore, the long-term IFB study observed vacuum degradation to depend on whether the test HPD was kept biased with HV or not inbetween measurements. The QE-IV results of degraded HPDs presented in Section 3.3.4 back this up: severely degraded vacuums were ‘annealed’ during the hour between the two sets of QE-IV measurements when the HPD was biased at 500 V. This conclusion motivated the policy during commissioning to keep the HPDs in the RICH detectors operated with HV whenever possible to keep vacuum degradation in check. There were periods where this was not possible, therefore this is one source of systematic uncertainties in the timeline plots.

A few HPD timeline plots observed a decrease in $R_{CWIFB,raw}$ over time, such as the example shown in Figure 4.9. The HPD in this example had a very good vacuum, therefore was sensitive to the vacuum quality improvements from keeping the HPD HV supplies operational. The figure includes the result of a linear fit to

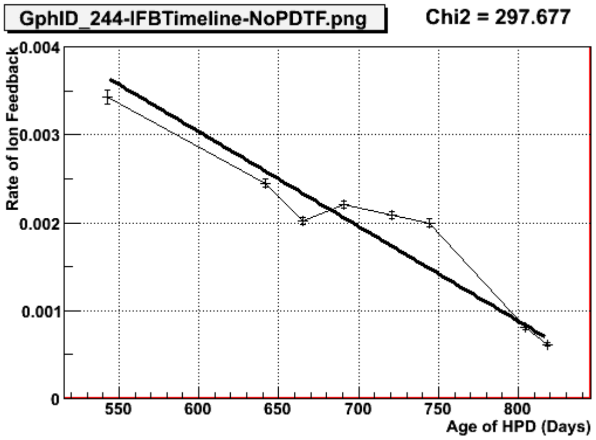

Figure 4.9: Ion feedback (IFB) development timeline of a sample Hybrid Photon Detector (HPD) showing gradual reduction in its rate of IFB. These IFB measurements plotted on the y-axis were measured with the continous wave light source and have not been corrected for dark count. The age of each HPD plotted on the x-axis is measured in days since they underwent the standard tests at the photon detector test facilities.

the timeline data points. The χ^2 value of the fit reminds us that the statistical error bars are small compared to the systematic uncertainties. However, the consistent decrease in $R_{CWIFB,raw}$ over almost 300 days seen in a number of very low $R_{CWIFB,raw}$ HPDs adds weight to this conclusion.

4.4.3 Timelines of Failed HPDs

Timeline plots of two sample glowing HPDs are shown in Figure 4.10. The HPDs are in positions A7-3 and A7-12 in RICH2. The former was observed to start glowing in February 2008, after the first timeline data point. The latter started glowing six months later, towards the end of the timeline as shown in Figure 4.10b. The RICH2 pixel hitmap shown in Figure 4.1 was taken after this first HPD started glowing but before the A7-12 HPD started glowing. Both examples satisfy the glow threshold hypothesis by having $R_{CWIFB,raw} > 4.8\%$ before glowing.

Both example timeline plots show a decrease in $R_{CWIFB,raw}$ after the HPD starts glowing. This is clearer in Figure 4.10a where the trend has enough data points to be visible despite systematic uncertainties disrupting the picture. Glowing HPDs do show ‘annealing’ under certain circumstances, one of which is being kept operational under nominal conditions and high, external illumination [31]. The annealing process reduces the glow effect, which the drop

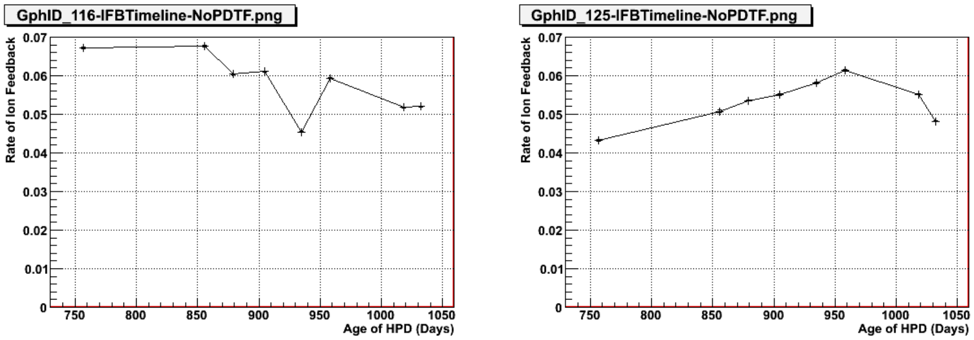

(a) HPD A7-3

(b) HPD A7-12

Figure 4.10: Examples of ion feedback (IFB) development timelines from glowing Hybrid Photon Detectors (HPDs). These IFB measurements plotted on the y-axis used the continous wave light source and have not been corrected for dark count. The age of each HPD plotted on the x-axis is measured in days since they underwent the standard tests at the photon detector test facilities. The left plot is for an HPD that was glowing from the start of IFB monitoring. The HPD in the right plot started glowing halfway through the timeline data points.

in $R_{CWIFB,raw}$ shows. However, this comes at the expense of highly accelerated degradation of the central region of the HPD photocathode.

4.4.4 Incorporating the PDTF IFB measurements

The timeline plots shown thus far have only used $R_{CWIFB,raw}$ data measured during RICH commissioning. Every HPD has a R_{IFB} measurement from when they underwent standard testing at PDTF. This measurement would be a data point at the very start of each HPD's age if incorporated into the timeline plots. However there are differences between how the R_{IFB} was measured at PDTF and in the RICH IFB monitoring. These differences needed to be compensated for to use the PDTF data points.

PDTF used a strobescan method to measure IFB, whereas the timeline plots have data that used the continuous wave light source. All HPDs mounted in the RICH had a low R_{IFB} measurement at PDTF as this was a test criteria to let the HPD be accepted. As discussed in Section 4.2.1, comparisons between R_{CWIFB} and R_{SSIFB} become difficult at very low R_{IFB}. The DAQ trigger was 50 ns wide at PDTF as opposed to the 25 ns width used in the RICH detector. The light source at PDTF was a 470 nm LED, which the HPDs are more sensitive to than the ~630 nm light sources available for IFB monitoring in the RICH. As

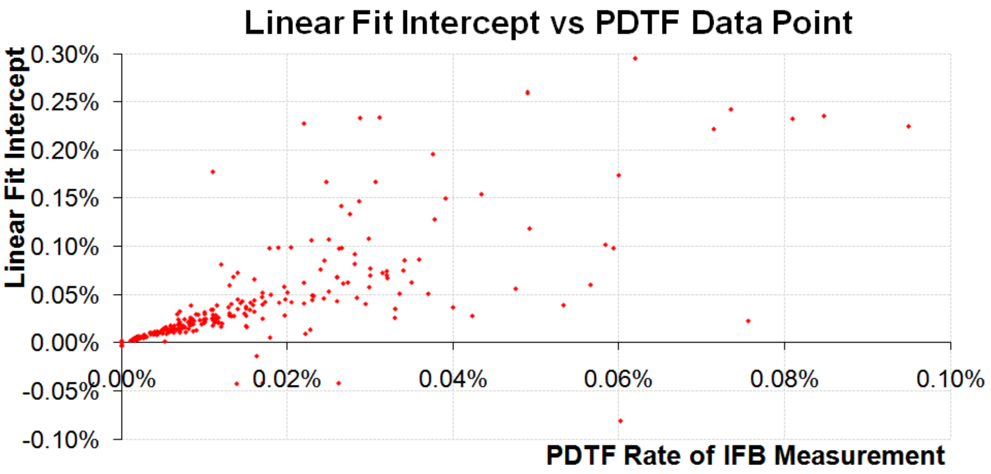

Figure 4.11: Correlation between linear fits of HPD timeline data and the Ion Feedback (IFB) measurement done at the Photon Detector Test Facilities (PDTF). The x-axis plots the rate of IFB as measured at the PDTF during standard tests. The y-axis plots the intercept c from $y = mx + c$ linear fits on HPD timeline data points. c is an estimate of what the rate of IFB should have been at the time of the PDTF measurement according to the linear model. Not all data points on this graph are shown as the axes are zoomed in to show the low IFB range where most of the data are.

presented in Section 3.3, the wavelength of the signal photons has a large impact on the quantum efficiency of the HPD. Finally, the PDTF R_{IFB} data are quotes of the maximum R_{SSIFB} measurement out of all the tested delays between the signal pulse and the DAQ trigger.

As a preliminary study of this issue, the intercepts from linear fits to HPD timeline plots was compared with the PDTF IFB measurements. The correlation between the two is shown in Figure 4.11. This plot was made early in the timeline analysis with a crude linear fitting algorithm that did not take error bar sizes into account.

The majority of HPDs had $< 0.2\%$ IFB rate when tested at PDTF. For very low R_{IFB}, there is a strong correlation between the intercepts and PDTF data points where a linear fit would give a gradient of $\sim$2.5. Given the discussions in Section 4.2.1, the fact the gradient is > 1 is as expected: the PDTF R_{IFB} value is calculated from the IFB activity in a single 50 ns wide bin from the Strobescan (SS) method of measuring IFB. As was shown in Figure 4.4, the SS method of R_{IFB} measurment gave comparable results to the CW method when all IFB activity from 100 ns to 400 ns delay times were summed up. Higher R_{IFB} HPDs are scattered around Figure 4.11 with no correlation between the two

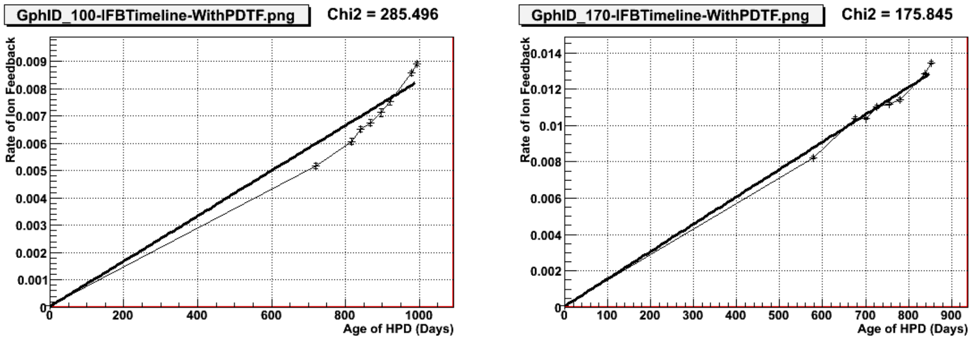

(a) HPD A6-3

(b) HPD C1-9

Figure 4.12: Examples of ion feedback (IFB) development timeline plots incorporating the IFB measurements from the photon detector test facilities. These IFB measurements plotted on the y-axis were measured with the continous wave light source and have not been corrected for dark count. The age of each HPD plotted on the x-axis is measured in days since they underwent the standard tests at the photon detector test facilities.

plotted quantities. There is more room for systematic uncertainties to distort the $R_{CWIFB,raw}$ measurements of such HPDs in the RICH.

Based on Figure 4.11, PDTF IFB measurements were incorporated into the timeline plots by multiplying their quoted R_{IFB} by 2.5. PDTF did not record errors for their IFB measurements. The timeline plots assumed a relative error of 10%. For HPDs with extremely low PDTF IFB measurements, they were assumed to have an absolute error of 10^{-5}. Examples of the resulting timeline plots are presented in Figure 4.12. These examples are the same HPDs as shown in Figures 4.8c and 4.8d. The first example shows how strongly weighted the PDTF data point is when the fit is performed, due to the small error bars on it. The χ^2 value of the timeline linear fit with the PDTF data point is ten times higher than when it was omitted in this example. How timeline plots were affected by including the PDTF data point is further discussed in the next section.

4.4.5 Distribution of Timeline Linear Fit Parameters

Plots comparing the effect of introducing the PDTF data point to the timeline plots are presented in Figure 4.13. The three pairs of plots show the distributions of the gradients, intercepts and χ^2 values of linear fits on the timeline plots of all RICH2 HPDs. Each of these parameters is shown in two plots: one presents the full distribution and the other is a close up of where most of the data points are.

In this section, the gradients and intercepts of the linear fits with and without the PDTF data point will be referred to as $m_{WithPDTF,NoPDTF}$ and $c_{WithPDTF,NoPDTF}$ respectively.

A general trend is seen across all three linear fit parameters when the PDTF data point is introduced. As expected, m_{NoPDTF} and $m_{WithPDTF}$ are very similar. Many HPDs had $|c_{NoPDTF}|$ already close to the origin and $|c_{WithPDTF}| < |c_{NoPDTF}|$. For all RICH2 HPDs, $\chi^2_{NoPDTF} < \chi^2_{WithPDTF}$. Though a lot of HPDs have a small χ^2 value, there is a wide distribution of values which gets wider with the PDTF data point. However there are many HPDs which do not fit all the above trends. Cuts are applied to separate out different cases for a more complete description.

131 out of the 288 RICH2 HPDs have $m_{NoPDTF} > m_{WithPDTF}$. These HPDs have a c_{NoPDTF} distribution with a mean of -8.13×10^{-3} and a Root Mean Square (RMS) of 1.47×10^{-2}. The mean and RMS of the $c_{WithPDTF}$ distribution are 4.75×10^{-4} and 9.75×10^{-4} respectively. This indicates that the example HPD shown in Figure 4.12a is representative of almost half the HPDs in RICH2: the linear fit had a negative intercept before the PDTF data point was introduced to anchor the fit closer to the origin. Therefore the gradient becomes more shallow when the PDTF data point is included.

91 out of the 288 RICH2 HPDs have $m_{NoPDTF} < 0$. Out of these, the majority have shallow gradients. These are HPDs similar to the example shown in Figure 4.9, where the $R_{CWIFB,raw}$ is so low they are sensitive to the vacuum quality improvement process. 12 of the 92 HPDs in this subgroup also have $m_{WithPDTF} > 2.0 \times 10^{-5}$. The same 12 HPDs remain after a cut of $c_{NoPDTF} > 0.05$ and they all have $\chi^2_{WithPDTF} > 10^5$. These are glowing HPDs like the example shown in Figure 4.10a. They undergo the annealing process for the glow effect so their $R_{CWIFB,raw}$ gradually decrease over time compared to when they first start glowing. For these HPDs, the PDTF data point with its small error bars force the linear fit to switch from a negative gradient to a steep positive gradient. This leaves a compromised intercept far away from the origin and an extreme χ^2 value.

29 out of 288 RICH2 HPDs have $|c_{NoPDTF}| < c_{WithPDTF}$. These HPDs had timeline data points which happened to give a linear fit with very small intercept. The PDTF data point is not at the origin so with low R_{IFB} HPDs the intercept is anchored higher than before. An extreme case of this is the outlier with

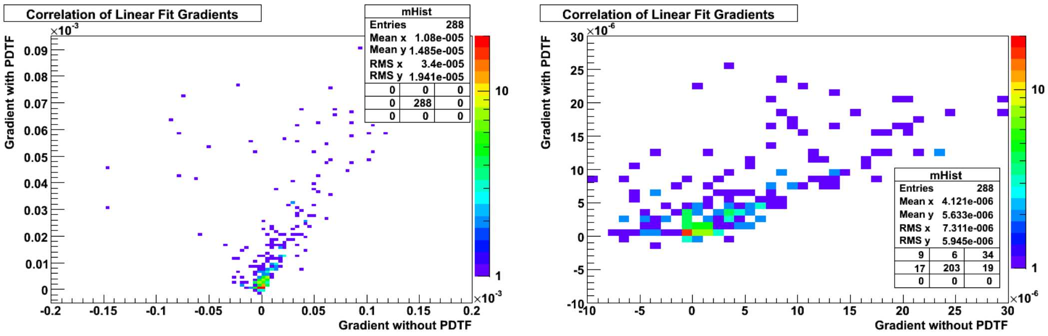

(a) Gradient

(b) Gradient (zoomed in)

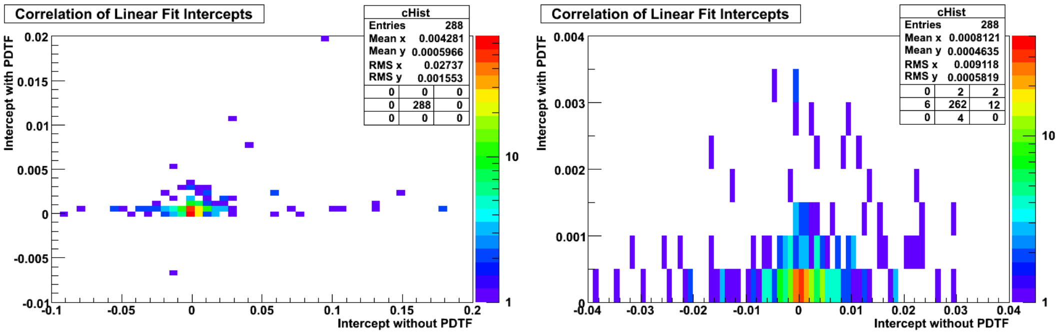

(c) Intercept

(d) Intercept (zoomed in)

Figure 4.13: Distributions of linear fit parameters for ion feedback (IFB) development timelines for RICH2 Hybrid Photon Detectors (HPDs). The x-axis and y-axis plot parameters from fitting without and with the PDTF measurement of rate of IFB respectively. The top and bottom pair of plots show the distribution of the gradients and intercepts found from linear fits to timeline plots of all RICH2 HPDs respectively. The right side plots of each pair is a zoomed up version of the original plot shown on the left.

Serial Number	Location	Glow Risk Date	Glow Start Date
H545001	A8-14	May 2008	Oct 2008
H539007	C2-0	Sept 2007	Nov 2008
H525014	A0-0	May 2008	Nov 2008
H524003	A7-11	Oct 2008	Jan 2009
H610002	C1-6	Dec 2008	Jan 2009

Table 4.2: Hybrid PhotonDetectors (HPDs) that began to glow at the end of 2008. Also shown are the dates that the timeline linear fits predicted these HPDs would be at risk according to the glow threshold hypothesis.

$c_{WithPDTF} = -6.5 \times 10^{-3}$. This particular HPD had a PDTF R_{IFB} measurement of 0.77%, many times higher than the average for RICH2 HPDs but still accepted as it was within tolerance of the PDTF standard tests. The HPD had eight noisy pixels and 29 always registered hits on every event during the PDTF tests. Since the timeline plots multiply the PDTF R_{IFB} measurement by a factor of 2.5 before performing the fit, this HPD had an inappropriately high data point as an anchor.

These observations show the variety of cases there are for HPD vacuum degradation in RICH2. Some of the most notable outliers have been identified to be glowing HPDs or have an exceptional detail. Many HPDs have very low R_{IFB} and many others have linear fits that are compatible with the PDTF data point. The rate of vacuum degradation is varied across the RICH2 HPD population.

4.5 HPD Lifetime Assessments

4.5.1 Predictions from the Timeline Fits

Using the glow threshold hypothesis and the timeline linear fits, HPD glow predictions were made. These predictions will be referred to as ‘HPD lifetimes’: HPDs older than their lifetimes are at risk of glowing. Lifetimes were calculated using the gradient and intercept of the timeline linear fit with the PDTF data point included.

How these lifetime predictions compared to actual dates when HPDs started glowing is presented in Table 4.2. The lifetimes shown in this table were extrapolated from IFB measurements up to August 2008. This model has had no false positives thus far, though the HPD may be deemed at risk for many months before starting to glow.

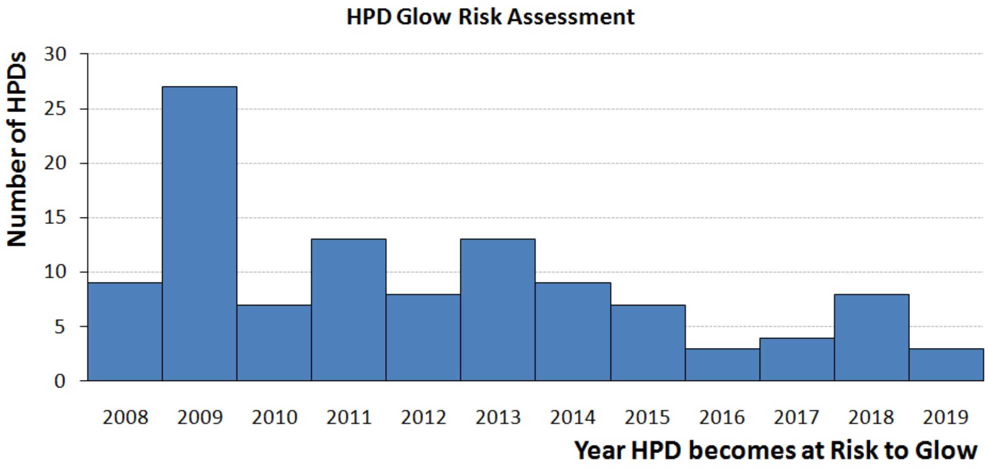

Figure 4.14: Predictions of how many HPDs in RICH2 will start to be at risk of glowing in each year. These predictions were made using the linear fits to the timeline plots.

A bar chart of how many HPDs each year reach the end of their lifetime is shown in Figure 4.14. This lifetime assessment was performed when 18 glowing HPDs had already been replaced, which are not included in the figure. A further 100 of the 288 RICH2 HPDs will be at risk of glowing by the end of 2017. Provided these predictions do not have any false positives, RICH2 will need on average 10 spare HPDs per year available for replacing new glowing HPDs over the operational lifetime of the LHCb. The concentration of HPDs expected to be at risk of glowing by 2009 is discussed in the next section.

4.5.2 Failure Susceptibility of Early Batch HPDs

The expectation that 36 HPDs will be at risk of glowing by the end of 2009 was shown in Figure 4.14. This is in addition to the 18 glowing HPDs already replaced beforehand. The cause for the bias in the distribution of lifetimes was seen when comparing HPD age to glow risk. Two plots of $R_{CWIFB,raw}$ against HPD age are shown in Figure 4.15, with age defined as in the timeline plots. The data for the two plots come from IFB measurement runs 201 days apart. Since HPDs were delivered in batches, their age is quantised. RICH1 was commissioned around June 2008 and the first $R_{CWIFB,raw}$ measurements are included in Figure 4.15b.

HPDs with very high $R_{CWIFB,raw}$ are among the oldest manufactured. In Figure 4.15a, they tend to have an age of >600 days. In Figure 4.15b this boundary has moved to the right by 200 days, staying constant relative to the

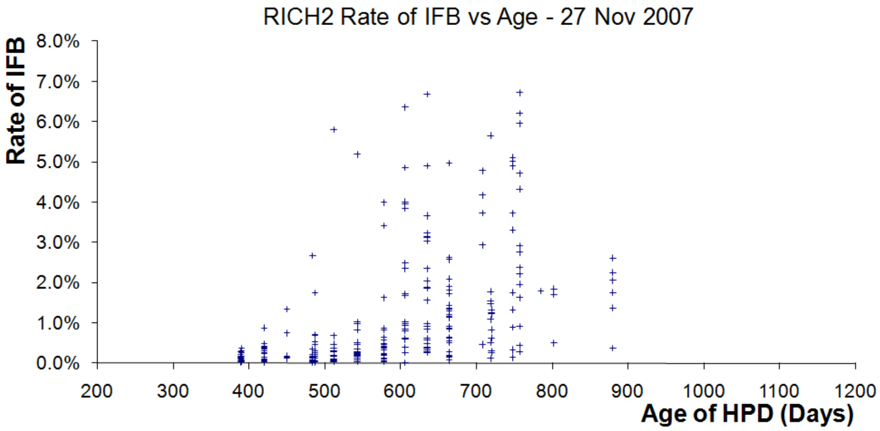

(a) 27 Nov 2007

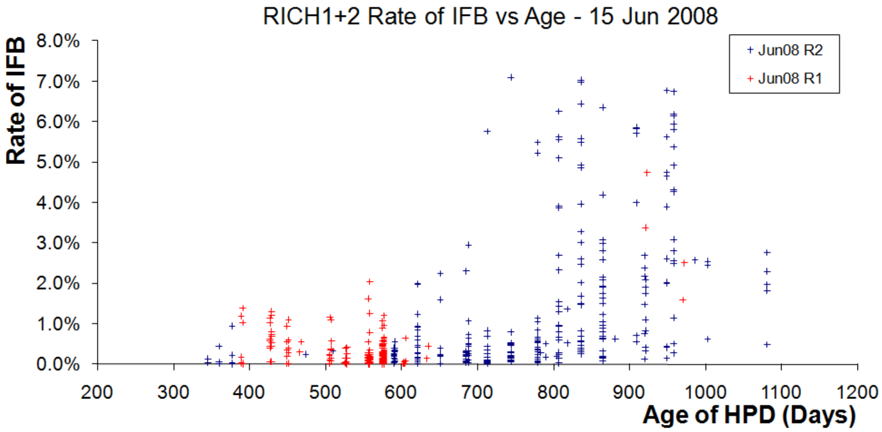

(b) 15 Jun 2008

Figure 4.15: Rate of ion feedback (IFB) against Hybrid Photon Detector (HPD) age. These IFB measurements used the continous wave light source and have not been corrected for dark count. The age of each HPD is measured in days since they underwent the standard tests at the photon detector test facilities. The top plot shows IFB data of RICH2 HPDs at the end of November 2007. The bottom plot shows IFB data almost seven months later of both RICH1 and RICH2 HPDs, drawn in red and blue respectively.

distribution of data points. This division corresponds to production batch 8∼9. It is known that around this time, the manufacturer switched their quartz window supplier. There could be a difference in residual gas leakage rates in the quartz windows from the different suppliers. Unfortunately, Photonis did not keep track of which supplier each quartz window originates from.

RICH1 was commissioned with newer HPDs so there were very few HPDs at risk of glowing. This is fortunate as replacing HPDs in RICH1 is much more difficult than RICH2 due to accessibility.

4.6 Conclusions

HPD vacuum monitoring procedures via IFB measurements were developed and assessed during RICH2 commissioning. A number of HPDs were found to suffer rapid vacuum degradation, becoming unusable due to the glow light effect. These HPDs disrupt the operation of over a dozen other HPDs that share the same column inside the RICH. Data from the IFB monitoring were used together with PDTF IFB measurements to develop a model for each HPD's rate of vacuum degradation. This model made predictions of when HPDs would be at risk of failure and so far has had no false positives. This has guided planning for replacing high risk HPDs with limited spares. The model's results were used in negotations with the manufacturer to acquire additional spares by repairs and new orders.

HPDs prone to fast vacuum degradation are all from old production batches. The manufacturer does not maintain sufficient bookkeeping to track down the exact source of failure. Work is ongoing to determine the cause of vacuum degradation. It is now known that the glow light has spectral lines matching that of Hydrogen, suggesting the presense of water vapour. Both RICH detectors will continue to be monitored with laser light runs, even during the physics programme of the LHCb.

Chapter 5

Flavour Tagging Studies in the LHCb Upgrade using $B_s \to \phi\phi$

5.1 Motivation

As discussed in Sections 1.6.2 and 2.5, after LHCb will have collected ~10 fb^{-1} of data over about five years further improvements to statistical precision would take increasingly exorbitant time to achieve when staying nominal luminosity. There will still be demand on LHCb for more statistics after this point. For example, CP violation in the rare decay $B_s \to \phi\phi$ would not have been measured to satisfactory precision as explained in Section 1.6.2. To meet this demand within a reasonable time of operation, LHCb plans to operate at a higher luminosity of $20 \times 10^{32}\,\text{cm}^{-2}\,\text{s}^{-1}$. A hardware and software upgrade is necessary to run at this upgrade luminosity. R&D for this upgrade will take years and thus is starting now even as the first data are produced at LHC. Whether the reconstruction and analysis software could even function under higher luminosity conditions needed demonstrating.

$B_s \to \phi\phi$ is a flavour changing neutral current decay that is forbidden at the tree-level in the Standard Model (SM). The final decay products are two pairs of charged kaons for both the B_s and its charge conjugate $\overline{B}_s \to \phi\phi$. Therefore CP violation via interference between mixing and decay can be observed in this channel, as explained in Section 1.5. The CP violating weak phase is expected to be zero in the SM, as explained in Section 1.6.1. The penguin loop diagram of this channel is sensitive to new physics beyond the SM which would affect the CP violation observed.

To study CP violation with neutral B decays such as with $B_s \to \phi\phi$, the initial flavour of the B-meson must be determined. Identifying the initial flavour of a signal decay is done through the process of flavour tagging, described in Section 5.1.1. Flavour tagging has been optimised at the nominal luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. To continue studying the $B_s \to \phi\phi$ channel, flavour tagging needs to be adapted for upgrade luminosity.

One of the issues being studied is the aerogel radiator of the Ring Imaging Cherenkov (RICH) detector. The present RICH1 relies on aerogel for Particle Identification (PID) of low momentum tracks. However the aerogel adds to the material budget of the detector which degrades the calorimeter performance. A proposal has been made to rely on a new time-of-flight RICH detector (TORCH) for the low momentum region and replace the aerogel radiator. Another proposal is to remove all of RICH1 and increase the size of RICH2 to have a super RICH detector. Studies were required urgently to establish early in the R&D cycle what the RICH systems should aim for.

This thesis presents studies on how flavour tagging is affected by the LHCb upgrade, with Monte Carlo (MC) $B_s \to \phi\phi$ signal datasets at different luminosities. The rare decay is of great interest as a probe for new physics. Furthermore, it is very dependent on kaon PID due to its final decay products. Thus the channel is also good for assessing the performance of the RICH systems under different proposed arrangements for the detectors.

Note these studies were finishing when the LHCb started taking data. Based on experience with LHC beams, the upgrade plans are changing. Data collection sometimes occurred at higher multiplicities than originally planned. These high multiplicity runs validated the results of these MC studies. At the time of writing, the LHCb upgrade may happen after only $4\,\mathrm{fb}^{-1}$ of data. This increases the importance of the upgrade. Current plans are to lower the upgrade luminosity to $10 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ at the time of writing. The studies in this and the next chapter are written with $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ as the upgrade luminosity. The change in upgrade luminosity plans do not invalidate the results of the studies.

Section 5.2 discusses the technical details of the upgrade simulations used for the studies on how flavour tagging is affected by luminosity. The performance at upgrade of event selection and particle identification algorithms necessary to study this channel is presented in Section 5.3. How flavour tagging is affected by increasing luminosity is discussed in Section 5.4. Finally, the impact of removing

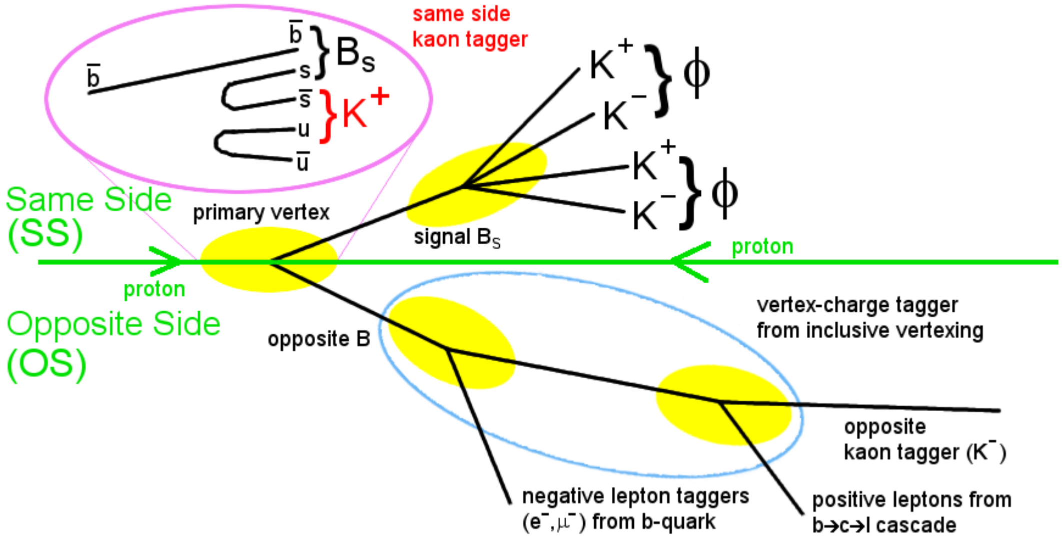

Figure 5.1: An example $B_s \to \phi\phi$ decay event with various possible methods to deduce the initial flavour of the signal B-meson. Diagram based on [33].

the aerogel in RICH1 is explored in Section 5.5.

5.1.1 Flavour Tagging

b-quarks are produced in pairs at the LHC. One of the b quarks forms a B-meson which decays in the channel of interest, $B_s \to \phi\phi$. Flavour tagging is the process of identifying whether a signal decay was produced as a B or a $\overline{B}$ meson at creation. There are different methods - taggers - for deducing the flavour of the signal B-meson [32]. An example of how a $B_s \to \phi\phi$ event can be flavour tagged is shown in Figure 5.1.

There are Same Side (SS) taggers which directly measure the flavour of the signal B-meson and Opposite Side (OS) taggers which deal with the other b-quark produced in the $b\bar{b}$ pair. The flavour of this b-quark is the opposite of the B-meson flavour at creation. B_s mesons oscillate rapidly on the order of $\mathcal{O}(1\,\mathrm{ps})$. The probability for the flavour of the signal B_s meson at creation to be the same at decay is a function of lifetime.

There are two SS taggers: the SS Kaon (SS K) tagger is used on B_s events and the SS Pion tagger acts on B^0 or B^+ mesons. When a B^0_s $(\bar{b}s)$ is formed, an extra $\bar{s}$ is left over which becomes a charged kaon ~50% of the time. If this occurs, the charge of this kaon indicates the flavour of the signal B_s-meson. The same principle holds for the SS pion tag which looks for a pion associated with the signal B.

OS taggers are divided into single particle taggers and the vertex charge tagger. These taggers can be used on any B event regardless of what the signal B is. The OS electron (OS e) and muon (OS μ) taggers use the charge of the lepton from a semi-leptonic b decay to tag the other B-meson. The OS kaon (OS K) tagger looks for the charge of the kaon originating from the $b \to c \to s$ decay chain. The vertex charge tagger tries to determine the net charge of a charged opposite B-meson's decay products to tag its flavour. The charge of all tracks associated with the opposite B-meson's decay vertex (a secondary vertex) are summed together using the transverse momentum, p_T to weight each track.

The individual taggers independently make tagging decisions: whether an event was a B, an $\overline{B}$ or if it could not be decided. Events where a tagger could not decide are called untagged events. To combine the outcome of each tagger for an overall decision on the B-meson flavour, a Neural Net (NN) is used [33]. The NN analyses the kinematic properties of the tracks used by the taggers to make their decisions. For each tagger, it assigns a probability for its decision being correct based on the NN's training. The product of these probabilities is then used to determine a combined probability on the B-meson flavour. During development, the NN was trained on simulation data where the only MC information it required was the true flavour of the signal B-meson. For real data, the NN is trained on self-tagging control channels such as $B^+ \to J/\psi K^+$ where the B-meson flavour is obvious from the decay products.

Flavour tagging performance is characterised with the tagging efficiency, ϵ_{tag}, and the mistag rate, ω. ϵ_{tag} is the fraction of events where a tagging decision was made and ω is the fraction of such events where the tagging decision was incorrect. These two quantities are combined into an effective efficiency, ϵ_{eff}:

$$\epsilon_{eff} = \epsilon_{tag}(1 - 2\omega)^2 \tag{5.1}$$

which is a measure of the equivalent tagging efficiency for an ideal flavour tagging algorithm that never tags incorrectly.

Most taggers rely on finding a single track in the event to make a tagging decision. Therefore they are likely to be sensitive to the higher track multiplicity introducing more dilution at the upgrade luminosity. The algorithm used by single track taggers to choose their tagging track is outlined in Section 5.4.2.

5.2 Simulating the LHCb Upgrade

The LHCb software framework for data simulation and analysis is split up into separate modules. First, Gauss uses the PYTHIA event generator to simulate particles flying through the detector after a proton-proton bunch crossing. Boole then simulates how each of LHCb's subdetectors would respond to these MC particle tracks. Brunel reconstructs tracks from the detector hits. It can either take Boole's simulated results or data from actual beam collisions. Finally, DaVinci analyses the tracks and selects events where the collision produced a decay in the user-specified signal channel. Information from Gauss on what actually was generated in the original event is called MC truth. This can be compared with how Brunel reconstructed the event to assess data analysis techniques in preparation for real data where there is no information on what actually happened, only what we reconstruct.

The MC datasets used in this chapter's upgrade studies were generated by a centrally managed production team. These datasets are available on the LHCb bookkeeping system. They are signal datasets where Gauss was configured so that all events contain a $B_s \to \phi\phi$ decay. The remainder of this section describes the detector geometry simulated in these datasets and discusses what happens when luminosity is increased in steps towards the upgrade target of $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

5.2.1 Minimal Upgrade Layout

While the details of much of the detector upgrade are still under research, some initial changes have been decided upon as feasible. A Minimual Upgrade Layout (MUL) detector geometry was agreed on for preliminary simulation studies of upgrade conditions. The differences compared to the existing detector layout are as follows:

- The Radiation Frequency (RF) foil in the VErtex LOcator (VELO) has been made thinner without changing the geometry.
- The beam-pipe support in the magnet had its materials changed to carbon fiber and titanium to make it lighter.
- The Inner Tracker (IT) plane is moved behind the Outer Tracker (OT) plane in each T-station on the z-axis to reduce occupancy in OT.

- Spill-over suppression for 40 MHz readout is implemented in the VELO, Trigger Tracker (TT) and IT.
- RICH1's aerogel radiator is removed.

Control samples were generated by leaving the aerogel in but with all the other changes still applied. It is assumed the LHC will operate at 7 TeV per beam. These data samples will also be referred to as MUL. Comparisons of performance between MUL simulations and existing detector geometry are presented in Sections 5.3 and 5.4.

5.2.2 Higher Luminosity Conditions

MUL simulations of nominal luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ correspond to an average of 0.7 interactions per proton-proton (pp) bunch crossing. This rate includes elastic interactions. Higher luminosity is about having more interactions per proton-proton bunch crossing. Only 'hard', inelastic, head-on pp collisions are of interest for physics analysis as opposed to 'soft', elastic, interactions. The number of such hard collisions per event, N_{hard} at each luminosity in the MUL datasets is illustrated in Figure 5.2. These plots used events guaranteed to have a $B_s \to \phi\phi$ decay which originates from a hard pp collision. Therefore this distribution is biased as every event has $N_{hard} \geq 1$. Real bunch-crossings have a finite probability for $N_{hard} = 0$ events. Nevertheless, the distribution shown in Figure 5.2a highlights the challenges of higher luminosity. From 2 to $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, not only does $< N_{hard} >$ increase from 1.4 to 4.7, the N_{hard} distribution gets wider.

In MC data samples, events originating from hard collisions can be distinguished from soft collisions by the PYTHIA ID. The only difference between low and high luminosity data samples should be the rate of hard collisions. To test this, events from the different luminosity datasets where the signal $B_s \to \phi\phi$ was successfully selected were filtered for N_{hard}. They were checked for any other differences within each N_{hard} category by plotting their flavour tagging performance as shown in Figure 5.3. Flavour tagging is a process taking in many variables into account as explained in Section 5.4. Within each N_{hard} category, flavour tagging efficiency and mistag rate are consistent within statistical errors as luminosity is varied. Data points with large error bars correspond to luminosities where a given N_{hard} has low statistics as shown in Figure 5.2a. Based on these

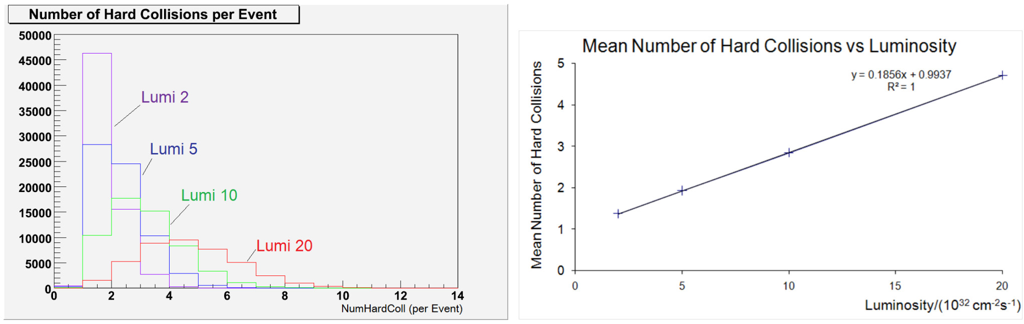

(a) The N_{hard} distribution (b) Linear fit of $< N_{hard} >$ vs luminosity

Figure 5.2: Data on the number of 'hard', inelastic proton-proton collisions per event, N_{hard} vs luminosity. Four different luminosities were simulated. These events come from signal datasets so there will always be at least one hard collision per event.

results, the different luminosity datasets do indeed only differ by the N_{hard} distribution and the track multiplicities that this change implies.

5.3 Selection and Particle Identification

Event selection is the process of picking out signal events and throwing away background events. This is done by filtering on a list of parameters, discussed in Section 5.3.1 in the context of increasing luminosity. For $B_s \to \phi\phi$ where the end products are all kaons, good PID of kaons is essential to selecting the signal events. Furthermore, kaons are used in flavour tagging: all B events can use the OS K tagger and B_s events also use the SS K tagger. Thus the kaon PID performance is important to this study. How kaon PID is accomplished is described in Section 5.3.2 and its performance under upgrade luminosity conditions is discussed in Section 5.3.3.

5.3.1 $B_s \to \phi\phi$ Selection

The pre-selection and selection cuts used for the $B_s \to \phi\phi$ channel were configured in [13]. They were optimised for nominal luminosity conditions and are listed in Table 5.1. How these cuts perform in the MUL and at higher luminosities on signal events is presented in Figure 5.4. First, a small improvement is seen in the selection efficiency going from the existing detector geometry to the MUL at nominal luminosity. This supports the few initial detector changes planned for in the MUL. As luminosity is increased, the selection efficiency degrades as

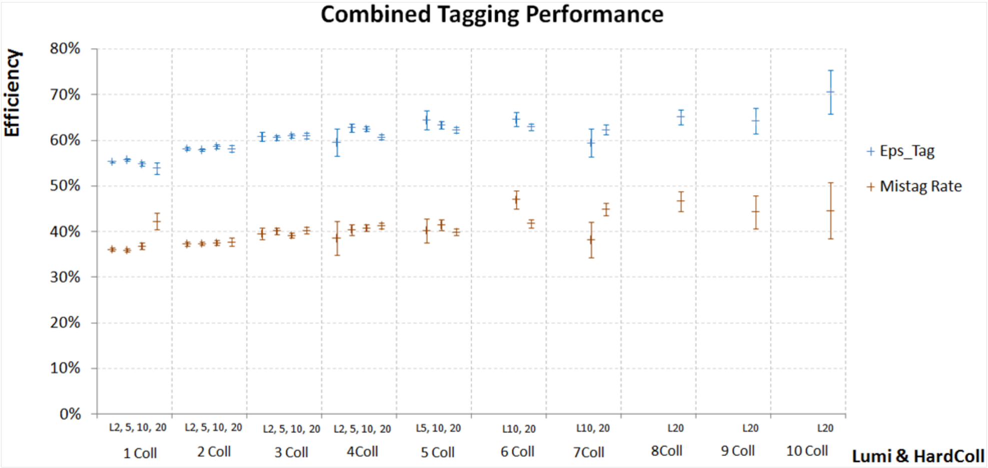

Figure 5.3: Flavour tagging performance vs number of hard collisions per event, N_{hard}. The events are from different luminosity datasets in the minimal upgrade layout, where the $B_s \to \phi\phi$ was successfully selected. The x-axis is first divided up into N_{hard}. Within each N_{hard} category, events are further sorted by which luminosity dataset they came from: 2, 5, 10 or 20 $\times$ $10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Datasets with an insignificant number of events in a given N_{hard} category are not plotted.

expected since the same selection cuts are being applied to increasingly different environments. This decrease is more or less linear with luminosity.

5.3.2 Kaon Particle Identification Methodology

The RICH counters are the only LHCb subdetectors that provide kaon PID. Without them, the LHCb detector cannot distinguish between kaons and pions as explained in Section 2.3.1. There are five particle types that the RICH counters need to separate on a regular bases: kaons, pions, electrons, muons and protons.

Conceptually, the principle behind how RICH works is explained with the ideal case of a single charged particle track giving off Cherenkov photons at a measurable angle determined by its velocity. In practice the RICH has to deal with tens of tracks each event, with each charged track radiating several Cherenkov photons plus noise in its photon detectors. RICH1 further complicates reconstruction by having two radiators and collecting their Cherenkov photons in the same detector plane. Instead of matching each photon to a track and then fitting for Cherenkov angles, RICH PID algorithms do a likelihood analysis starting from the tracks instead of the photon hits [34].

RICH PID uses trajectory information from the tracking system. All tracks are assumed to have radiated Cherenkov light at the centre of their flight path

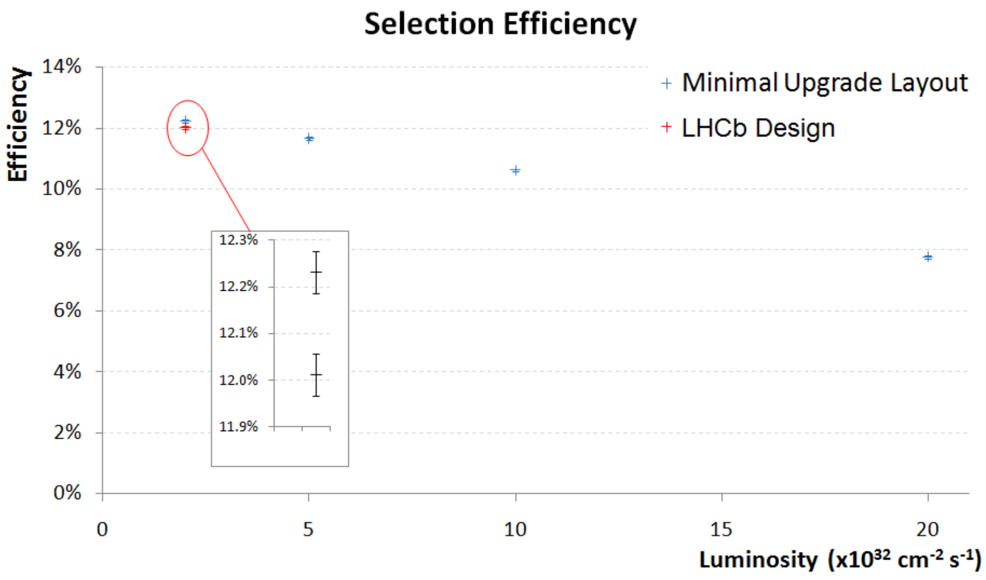

Figure 5.4: Selection efficiency vs luminosity. A control dataset with no changes to the existing detector layout is plotted in red. The inset plot zooms into the two $2 \times 10^{32}\,\text{cm}^{-2}\,\text{s}^{-1}$ data points.

through a radiator. With this information, the expected pattern of hits on the RICH photon detector plane can be worked out for a given event hypothesis of what particle type each track was. The event hypothesis also includes background noise considerations and how they would add to the expected pattern of signal hits. This expected pattern can then be compared to the actual observed pattern of hits to calculate an event likelihood, $\mathcal{L}$ for that hypothesis. Conceptually, the likelihoods from all possible hypotheses can be compared with each other to determine the most likely hypothesis of what happened in the event.

The above conceptual methodology is impractical: the algorithm must process the event quickly and be ready for the next triggered event. With tens of tracks in each event, checking every possible event hypothesis takes too long. The actual strategy used is based on the fact that pions are the most common particle tracks in the LHCb experiment. First, all tracks are assumed to be pions and the event hypothesis is calculated for this starting point. One track is then changed in the hypothesis and the likelihood of the new event hypothesis is compared to the starting point. After seeing the likelihood differences when that track is a proton, muon, electron or kaon the track reverts back to a pion and a new track is selected. When all tracks have been individually tested this way, the scenario giving the highest event likelihood is selected. The process now reiterates with

$\phi \to$ KK Selection	
ϕ mass window	$\pm\,20\,\mathrm{MeV}$
p_T of K	$> 400\,\mathrm{MeV}$
Minimum IP χ^2 of K from Primary Vertex	> 4
χ^2 cut to ϕ vertex	< 25 per Vertex degree of freedom
$\mathrm{B_s} \to \phi\phi$ Selection	
$\mathrm{B_s^0}$ mass window	$\pm\,50\,\mathrm{MeV}$
p_T of ϕ	$> 1200\,\mathrm{MeV}$
Minimum IP χ^2 of ϕ from Primary Vertex	> 4
δ_z distance between ϕ end vertex and B_s vertex	> 0
Angle θ between ϕ momentum and vector between ϕ and B_s vertices	$\cos\theta > 0.999$

Table 5.1: List of selection cuts for the $\mathrm{B_s} \to \phi\phi$ channel.

this new hypothesis as the reference point, until a maximum is found.

The event hypothesis with the maximum $\mathcal{L}$ is used as a reference to record PID information on each track for data analysis. The PID information for a given track is calculated as $\ln\mathcal{L}(i)$ if that track was of particle type i. This information is normalised such that $\ln\mathcal{L}(\pi) = 0$. These $\ln\mathcal{L}(i)$ values can then be compared to get a measure of how certain the PID software is that a track is one particle type instead of another. Such comparisons are expressed as $\Delta\ln\mathcal{L}$ (DLL). For example, $\Delta\ln\mathcal{L}(K-\pi) > 10$ is a harsh cut filtering for tracks which the PID algorithm is very certain of being a kaon compared to being a pion. $\Delta\ln\mathcal{L}(K-\mathrm{p}) = 0$ means the PID cannot distinguish whether this track is a kaon or a proton.

As this PID strategy depends on trajectory knowledge, tracks are also categorised according to how much tracking information is available. Tracks registering hits on all tracking stations including the VELO are labelled 'long tracks'. Downstream tracks have hits in the TT and the IT/OT stations but are missing VELO information. Upstream tracks have VELO and TT hits but nothing in either the IT or OT detectors. There are other track types to cover cases with even less tracking information but they are of little use for PID.

PID performance is measured with two ratios: efficiency is the fraction of true particles of a given type that were correctly identified, and the mis-identification (misID) rate is the probability of tracks to be erroneously selected as a particle type they are not. For example when identifying kaons, efficiency is lowered by failing to identify true kaons and the $\pi \to K$ misID rate is the fraction of pions that are misidentified as kaons.

5.3.3 Kaon Particle Identification at Upgrade

The momentum distribution of MC true, long-type kaon tracks is presented in Figure 5.5. These tracks are from events where the $B_s \to \phi\phi$ in the signal dataset was selected. Both signal and background tracks are included in this figure. The momentum spectrum mostly remains the same despite the difference in luminosity. Track momenta are gathered around the lower region of sensitivity of the RICH system.

Kaon PID performance is plotted against track momentum at nominal and upgrade luminosity in Figure 5.6. A $\Delta \ln \mathcal{L}(K - \pi) > 0$ cut was used to select kaons. This is the default DLL cut for the opposite side kaon tagger when filtering kaons from pions. At nominal luminosity between ~ 4 and $60\,\mathrm{GeV}$, the efficiency for selecting true kaons is $\sim 98\%$ and the misidentification rate of identifying a true pion as a kaon is 5%.

There are fewer tracks with momentum above $60\,\mathrm{GeV}$ as shown in Figure 5.5. Therefore the statistical errors are larger in this momentum range in Figure 5.6. Despite this, it is clear that in RICH2 the PID performance deteriorates towards the upper limit of its momentum range but is still usable for the few kaon tracks with such high momentum. Below $4\,\mathrm{GeV}$, kaon tracks are still identified with good efficiency down to the threshold of the aerogel radiator, but the $\pi \to K$ misID rate rises steeply. For this reason, flavour tagging using kaons does not use tracks below $4\,\mathrm{GeV}$.

Increasing the luminosity does not degrade the kaon PID efficiency by much until the upper end of RICH2's momentum range. The rate of pion misidentification has a large relative increase. If a certain threshold of $\pi \to K$ misID rate is desired, the DLL cut must be made tighter. This is shown in Figure 5.7, which is a plot of pion misidentification rate against kaon PID efficiency for different DLL cut values at different luminosities. These results were taken without any attempt to adapt the PID algorithm to the higher luminosities.

For the default $\Delta \ln \mathcal{L}(K - \pi) > 0$ cut, increasing the luminosity degrades the $\pi \to K$ misID rate faster than the kaon PID efficiency decreases. At very loose

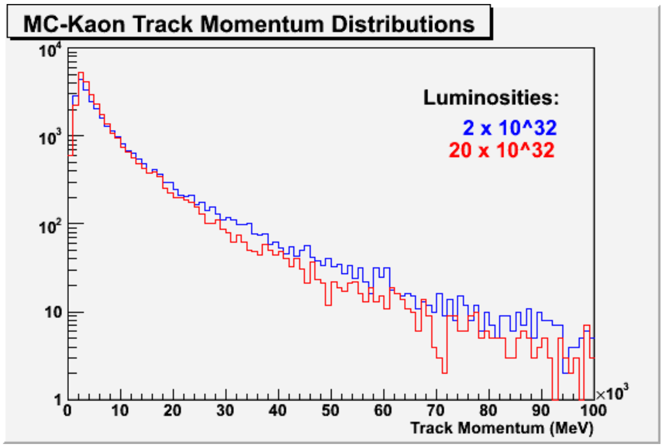

Figure 5.5: Momentum distribution of all Monte Carlo truth-confirmed kaon tracks from 30k events where the $B_s \to \phi\phi$ was successfully reconstructed. The blue (red) curve is from the nominal (upgrade) luminosity dataset. The results have been filtered to only show tracks with information from all tracking stations. Note the logarithmic y-axis.

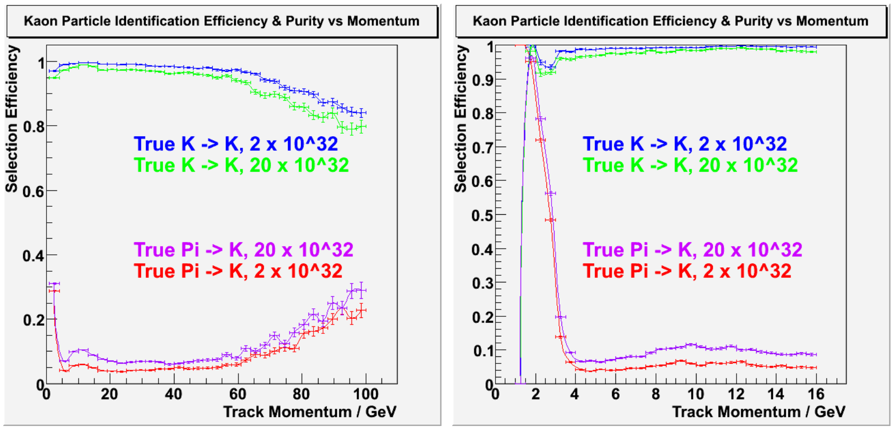

(a) $p = 2 \sim 100\,\text{GeV}$.

(b) $p = 2 \sim 16\,\text{GeV}$.

Figure 5.6: Kaon Particle Identification (PID) performance vs track momentum at nominal and upgrade luminosity. All tracks are plotted from the $B_s \to \phi\phi$ signal datasets in the minimal upgrade layout. All tracks have transverse momentum, p_T >500 MeV and have left a hit on at least one RICH radiator. Tracks which the PID algorithm could not make a decision on have been filtered out. The blue and green curves are the PID efficiency of correctly identifying a true kaon as a kaon at nominal and upgrade luminosity respectively. The red and purple curves are respectively the 2 and $20 \times 10^{32}\,\text{cm}^{-2}\,\text{s}^{-1}$ rate of misidentifying a true pion as a kaon.

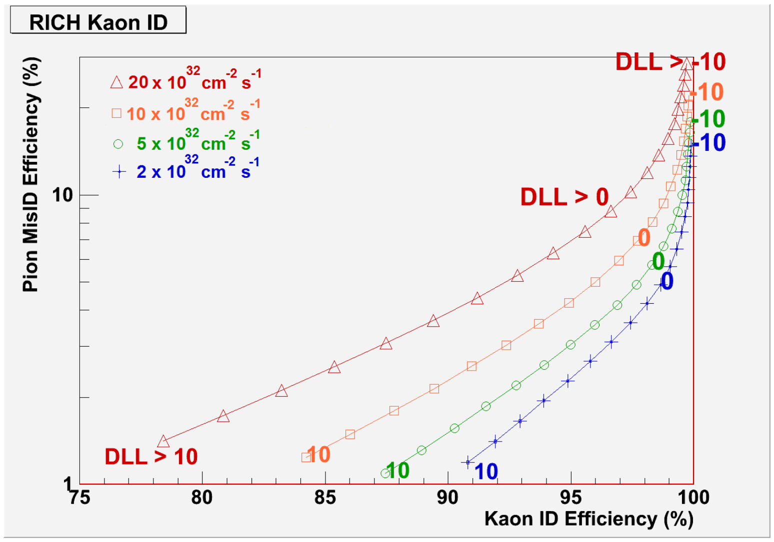

Figure 5.7: $\pi \rightarrow K$ mis-identification rate vs kaon particle identification efficiency, for various luminosities and $\Delta \ln \mathcal{L}(K - \pi)$ (DLL) cuts. The DLL cuts specify how certain the PID likelihood analysis should be to label a track as a kaon instead of a pion, with DLL > 0 being the default. These are all Monte Carlo truth-matched kaon and pion tracks from events where the $B_s \rightarrow \phi\phi$ was successfully reconstructed and selected by the LHCb analysis software. The different data points on each luminosity curve shows results from different DLL cuts, ranging from very loose at the top right to very tight on the bottom left. The y-axis (x-axis) is the fraction of true pions (kaons) satisfying the DLL cut. Note the y-axis is logarithmic.

DLL cuts, the efficiency is close to 100% for all luminosities but the misID rate gets very high, more so at higher luminosities. At very tight DLL cuts, the misID rate is excellent at all luminosities but a large fraction of true kaons do not pass the cut, again more so at higher luminosities. To maintain the current $\pi \rightarrow K$ misID rate of 5% at upgrade luminosities, the PID efficiency drops from ∼98% to ∼92% by making the DLL cut harsher.

To conclude, the current LHCb software for selection and particle identification still function at the higher track multiplicities of the upgrade luminosity. However, without adapting the algorithms the performance degrades. In the case of kaon PID, if the RICH likelihood analysis is not adapted to cope with the higher track multiplicity environment, the default DLL cuts to select kaons will need to be tightened.

5.4 Flavour Tagging Performance vs Luminosity

5.4.1 Combined Tagging Performance at Higher Luminosities

The flavour tagging neural net described in Section 5.1.1 was configured for nominal luminosity. A simulation of how flavour tagging would degrade with increasing luminosity if both the neural net and the individual taggers were left unchanged is presented in Figure 5.8. This shows that at nominal luminosity, there is no statistically significant change in tagging performance between the existing detector geometry and the MUL.

As luminosity is increased, both tagging efficiency and mistag rate increase as shown in Figure 5.8a. Since there are more tracks in each event at higher luminosity, the taggers have more chance of finding a suitable track candidate to make a decision with. This improves the tagging efficiency, but the increase in tracks introduces more wrong tags which the taggers and neural net have not been adjusted to cope with. The mistag rate degrades and the effective efficiency drops as it is more sensitive to mistags than the increased fraction of tagging decisions.

5.4.2 Track Candidate Selection for Single Particle Taggers

The single particle taggers defined in Section 5.3.2 select a candidate track to base their tagging decision by the following procedure. Before any individual tagger is run, a pre-selection is performed on all tracks in an event selected by physics analysis as containing a $B_s \to \phi\phi$ decay:

- $2.0\,\text{GeV} < p < 200\,\text{GeV}$
- Angle of momentum vector wrt z-axis > 0.012 rad
- Track is a charged particle
- Candidate is either a long track or an upstream track
- $p_T < 10\,\text{GeV}$
- Candidate is not a signal track

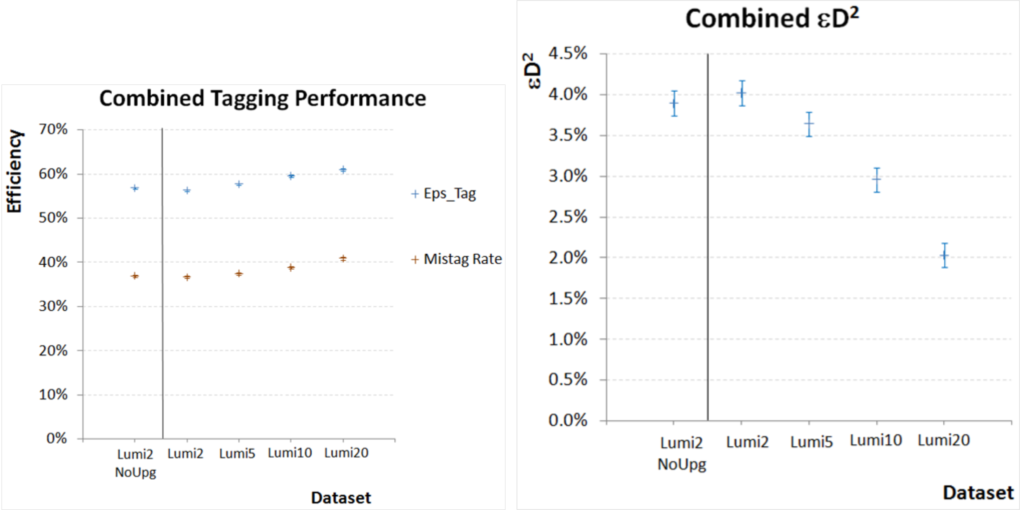

(a) Tagging efficiency and mistag rate.

(b) Effective efficiency.

Figure 5.8: Flavour tagging performance vs luminosity with results from 500k Monte Carlo signal events. The vertical line separates simulations in existing detector geometry (left) and in minimal upgraded layout (right). Luminosities of 2, 5, 10 and $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ were simulated in the upgrade layout. This is the performance of the combined tagging system that takes into account input from all applicable taggers.

- $\phi_{daugh} > 0.005$
- $\mathrm{I}_{PileUp} > 3.0$

ϕ_{daugh} is the lowest difference in angle ϕ between the B candidate momentum vector and all its daughter momenta vectors, i.e. this cut rejects B decay events where any of its daughter tracks do not change direction from its own momentum vector. I_{PileUp} is the lowest Impact Parameter significance (IP_{sig}) to all non-signal PVs, i.e. the last cut is to remove tracks originating from other PVs in events where there were multiple ‘hard’ proton-proton collisions.

This removes tracks that are not useful to tagging early on, saving computation time. The list of pre-selected track candidates is then passed over to every tagger including the vertex charge tagger. The single particle taggers then apply their own cuts to this candidate list, as described in Table 5.2.

In events where more than one track passes all cuts for a tagger, the candidate with the highest p_T is selected to use for tagging. This final selection procedure is the same for all single particle taggers and how flavour tagging is affected if it was changed is discussed in Section 6.1.3.

Cuts applied to select tagging track			
SS K	OS K	OS e	OS μ
$\Delta \ln \mathcal{L}(K - \pi) > 1$	$\Delta \ln \mathcal{L}(K - \pi) > 0$	identified as e	identified as μ
$\Delta \ln \mathcal{L}(K - \mathrm{p}) > -1$	$\Delta \ln \mathcal{L}(K - \mathrm{p}) > -1$	seen in ECAL	
$p > 4\,\mathrm{GeV}$	$p > 4\,\mathrm{GeV}$	$p > 0\,\mathrm{GeV}$	$p > 0\,\mathrm{GeV}$
$p_T > 450\,\mathrm{MeV}$	$p_T > 400\,\mathrm{MeV}$	$p_T > 1.2\,\mathrm{GeV}$	$p_T > 1.1\,\mathrm{GeV}$
$\lvert\mathrm{IP_{sig}}\rvert < 3.0$	$\mathrm{IP_{sig}} > 3.8$		
$\mathrm{d}\eta < 1.0$	$\lvert\mathrm{IP}\rvert$ (mm) < 1.5	Long or Upstream	
$\mathrm{d}\phi < 1.1$		$0.0 < Q_{VELO}$ $Q_{VELO} < 1.3$	
$\mathrm{dQ} < 1600\,\mathrm{MeV}$	$\mathrm{LCS}_{Long} < 2.5$	$E/p > 0.85$	
$\mathrm{LCS} < 2.0$	$\mathrm{LCS}_{Upstream} < 2.0$	$\mathrm{LCS} < 3.0$	$\mathrm{LCS} < 2.0$

Table 5.2: Cuts used by each single particle tagger to select a track to make a tagging decision with. Track candidates have been pre-selected to remove tracks in the signal channel decay. dQ is the mass difference between the track and the B mass. LCS is the track's χ^2 per degree of freedom. Q_{VELO} is the $\mathrm{d}E/\mathrm{d}x$ charge from the VELO system.

5.4.3 Individual Tagger Performance

The effect of luminosity on the tagging performance of each tagger is presented in Figure 5.9 and discussed here. As with the combined tagging performance, there is little significant difference between the dataset with no upgrades and the MUL dataset at $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

The vertex charge tagger is the most sensitive to the luminosity: both ϵ_{tag} and ω increase as luminosity increases. Unlike the other taggers which look for a single, non-signal track to base its tagging decision on, the vertex charge tagger bases its decision on all tracks originating from the opposite b-hadron decay vertex. Hence why it is the most sensitive to the higher track multiplicity of the higher luminosities.

The OS e and OS μ taggers show consistent tagging efficiency across the five datasets. The track candidate selection procedure for these lepton taggers do not tag more often despite an increase in track multiplicity from the luminosity change. This is because the semi-leptonic decays that produce the electrons and muons for these taggers have a low branching ratio. Since ϵ_{tag} is so low, there are less statistics to calculate ω with, hence the larger error bars for them in Figure 5.9b. Ideally, a larger dataset would discern whether ω and the effective

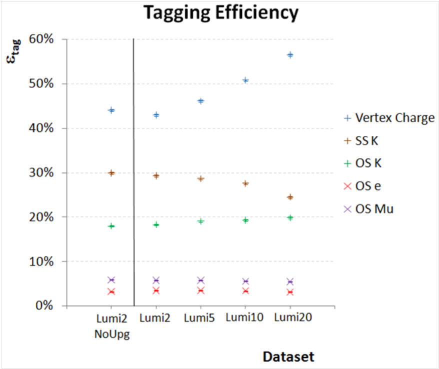

(a) Tagging efficiency.

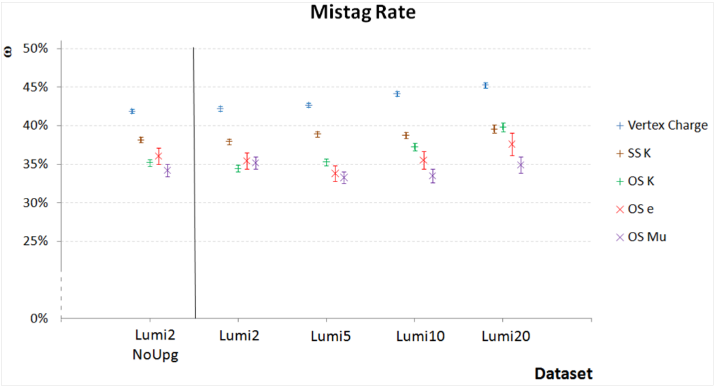

(b) Mistag rate.

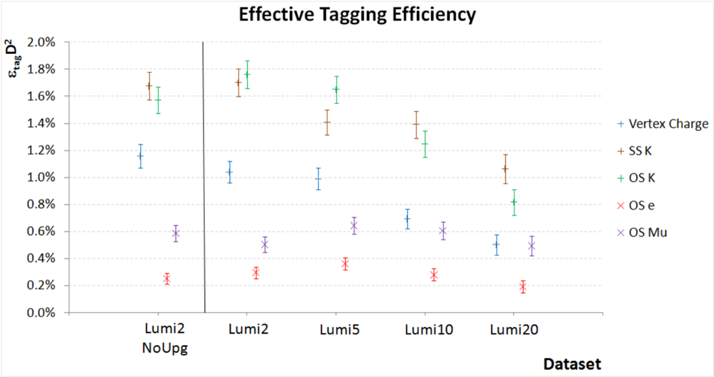

(c) Effective efficiency.

Figure 5.9: Flavour tagging performance of individual taggers vs luminosity, using 500k Monte Carlo signal events. The vertical line separates simulations in existing detector geometry (left) and in minimal upgraded layout (right). Luminosities of 2, 5, 10 and 20 $\times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ were simulated in the upgrade layout.

efficiency of these two taggers are affected much by the luminosity increase. However, priority was given to the study of the kaon taggers.

The kaon taggers have the best effective tagging performance out of the five taggers for $B_s \to \phi\phi$, as shown in Figure 5.9c. They also degrade the most in performance as luminosity increases. Both kaon taggers see an increase in ω as a result of kaon PID degradation at higher luminosity. This was shown back in Figure 5.7. The kaon taggers apply DLL cuts during their track candidate selection, as was listed back in Table 5.2. The analysis shown in Figure 5.9c did not change these cuts. For a given DLL cut, both the $\pi \to K$ mis-identification rate and kaon particle identification efficiency deteriorate at higher luminosity.

The higher track occupancy at upgrade luminosity affects the two kaon taggers differently. It is harder for the SS K tagger to find the kaon that accompanies the signal B decay, resulting in a reduced frequency in decision making. Thus this tagger is the only one to lose tagging efficiency at higher luminosity. The track occupancy of higher luminosities makes it easier for the OS K tagger to find a tagging candidate track. However, with the degraded kaon PID from the luminosity increase the probability for that track to have been misidentified as a kaon is higher. Hence the OS K tagger has a larger increase in ω compared to the SS K tagger: the SS K tagger has more stringent cuts to select its candidate tracks from the signal B decay vertex and not the opposite b-hadron decay vertex.

Overall, this study shows that our existing flavour tagging algorithms can still function under upgrade conditions but at reduced performance. Kaon taggers have the best effective efficiency under nominal luminosity but degrade the most at higher luminosity.

5.5 Aerogel in the Upgrade

The RICH detector uses Cherenkov radiation to identify particles. The angle, θ that Cherenkov photons are emitted at is given by $\cos\theta = \frac{1}{n\beta}$, where n is the refractive index of the radiator medium and $\beta = v/c$. Thus there is a momentum threshold for Cherenkov radiation: if the particle's momentum is low enough such that $\frac{1}{n\beta} > 1$, there are no physical solutions to $\cos\theta$ and no Cherenkov photons are emitted. The RICH uses aerogel as its low momentum radiator as the C_4F_{10} gas will not radiate for kaon tracks below 9.3 GeV momentum [21].

However, the aerogel adds to the material budget of the detector. While the momentum threshold of the C_4F_{10} gas radiator for kaon tracks is 9.3 GeV, pion tracks still radiate in C_4F_{10} gas down to 2.6 GeV. The aerogel radiator has momentum thresholds of ∼2 GeV for kaons and 0.6 GeV for pions. RICH PID algorithms have a veto mode for tracks with momenta between kaon and pion thresholds: if there is no Cherenkov radiation for these tracks, the RICH software will assume it was a kaon as a pion would have radiated.

The low momentum thresholds of the aerogel mean the veto mode is not required often for physics analysis. But upgrade luminosity makes a lower material budget desirable. Can the RICH function satisfactorily by relying on the veto mode if the aerogel is removed? This section seeks to answer this question, using MUL datasets with and without the aerogel in RICH1. Datasets with no aerogel changed the preselection to exclude tracks of momentum below 2.6 GeV as the RICH gets no Cherenkov photons for either pions or kaons in this regime without aerogel.

5.5.1 Aerogel Effects on the Selection

Selecting $B_s \to \phi\phi$ events requires kaon identification. The effect of removing aerogel on the selection efficiency is shown in Figure 5.10. As before in Section 5.3.1, the selection cuts are unchanged and the datasets are MC signal datasets in the MUL. Removing the aerogel has no statistically significant effect on the selection efficiency. Despite many tracks being in the 2.6 GeV to 9.3 GeV range as shown in Figure 5.5, the selection efficiency is not compromised by the removal of aerogel due to the veto mode.

These results show the following studies have similar statistics at each luminosity for comparing aerogel to no aerogel scenarios. Investigating how kaon PID performs under veto mode without aerogel is presented in the next section.

5.5.2 Kaon Identification without the Aerogel

Aerogel is the low momentum range RICH radiator. A significant fraction of kaon tracks in $B_s \to \phi\phi$ events are in the aerogel's momentum range of $2 \sim 10$ GeV as shown in Figure 5.5. How kaon PID is affected at nominal luminosity by removing the aerogel and relying on veto mode is illustrated in Figure 5.11.

The $\Delta \ln \mathcal{L}(K-\pi)$ distribution does not change much when aerogel is removed in Figure 5.11a. Of note is that the peak at 0 is significantly higher without

Figure 5.10: Selection efficiency of $B_s \to \phi\phi$ when aerogel is removed. The left pair of data points are at $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ and the right pair are the $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ luminosity results. Each pair of data points compares simulations with aerogel in (left data points) vs aerogel removed (right data points) in otherwise identical MC signal datasets in the minimal upgrade layout. The inset plots show the data zoomed up to clearly show the error bars.

aerogel but the curve shapes are similar. The RICH PID veto mode successfully differentiates pions from kaons in the low momentum range when the aerogel is removed. Some tracks become indistinguishable while in veto mode, but given a large fraction of tracks have momentum $< 10\,\mathrm{GeV}$ it is remarkable how veto mode can maintain the same DLL distribution shape.

Kaon-proton separation has a very noticable difference when aerogel is removed, shown in Figure 5.11b. In the region $-10 < \Delta\ln\mathcal{L}(K - \mathrm{p}) < 20$, about half the true kaon tracks which would have had a PID decision end up indistinguishable when the aerogel is removed. This is the importance of having aerogel. Kaon-pion separation can manage without it by using veto mode. But kaon-proton separation suffers because low momentum tracks below the $\mathrm{C_4F_{10}}$ radiator's kaon threshold are also below the proton threshold, since the proton has larger mass than the kaon. In flavour tagging, it is important to tell apart

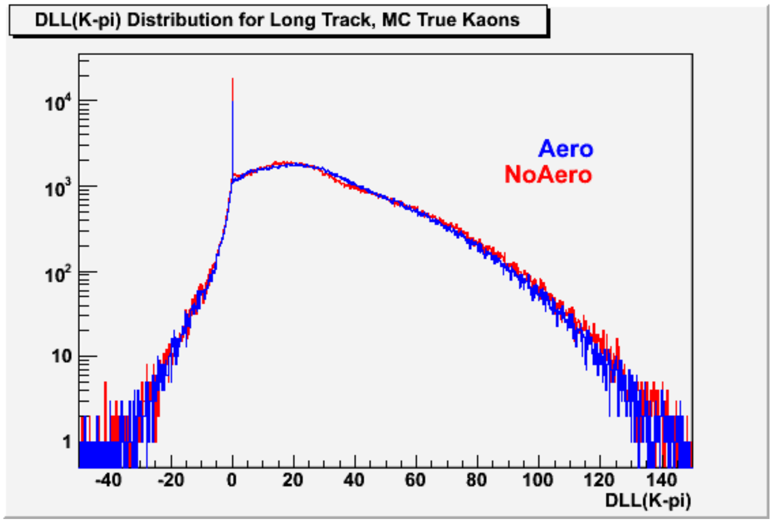

(a) $\Delta \ln \mathcal{L}(K-\pi)$, logarithmic y-axis

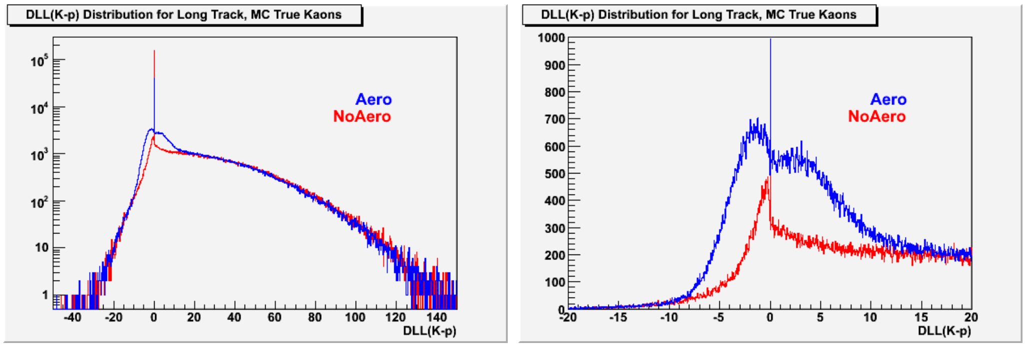

(b) $\Delta \ln \mathcal{L}(K-\mathrm{p})$, logarithmic y-axis (c) $\Delta \ln \mathcal{L}(K-\mathrm{p})$ zoomed up, linear y-axis.

Figure 5.11: Kaon particle identification $\Delta \ln \mathcal{L}$ distributions. The data are from events where the $B_s \to \phi\phi$ was selected from the Monte Carlo (MC) signal datasets at nominal luminosity in the minimal upgrade layout. Only MC truth matched, long-type kaon tracks are displayed, both signal and background. Blue plots have aerogel in and red plots are simulations without aerogel. (a) is the $\Delta \ln \mathcal{L}(K-\pi)$ distribution. (b) is the $\Delta \ln \mathcal{L}(K-\mathrm{p})$ distribution and (c) is a zoomed up version of the same plot.

kaons from protons when using the kaon taggers. How the kaon taggers are affected by removing aerogel is discussed in the next section.

5.5.3 Flavour Tagging without Aerogel

How the individual taggers are affected by the removal of aerogel in the MUL simulations is shown in Figure 5.12. The kaon taggers tag more often when aerogel is removed, both at nominal and upgrade luminosities. The cause of this is discussed in Section 5.5.4. The other taggers see no significant difference in tagging efficiency from the aerogel removal. This is expected as these other taggers are not concerned with the kaon PID change in the absence of aerogel.

After the mistag rate is taken into account the effective efficiency, ϵ_{eff} of individual taggers is shown in Figure 5.12b. At nominal luminosity, the kaon taggers degrade in performance despite tagging more often because they are making the wrong decisions more often too. Since effective efficiency is dependent on $(1-2\omega)^2$, the mistag rate increase brings down tagging performance despite making more frequent decisions. The error bars of the tagging efficiency and mistag rate for the other taggers are too large to draw any strong conclusions of how they are affected by the aerogel removal. Somewhat surprisingly, the kaon taggers see no difference in ϵ_{eff} when aerogel is removed in the upgrade scenario. The lepton taggers seem to have an increase in ϵ_{eff} by removing aerogel at upgrade luminosity but the error bars are too large to be conclusive.

The combined effective efficiency from the NN tagging decisions is presented in Figure 5.12c. The degradation of kaon taggers and the improvement of the other taggers when aerogel is removed result in no net change to tagging performance at nominal luminosity. At upgrade luminosity, the effective efficiency has potentially increased by removing aerogel though the statistical errors are too large to draw a conclusion.

These results were obtained without any changes to the existing taggers or the NN. Within this scope, the absence of aerogel is not detrimental to flavour tagging as a whole. The question of how these results would change if the tagging algorithms were adapted for the absence of aerogel and the higher luminosity is explored in the next chapter.

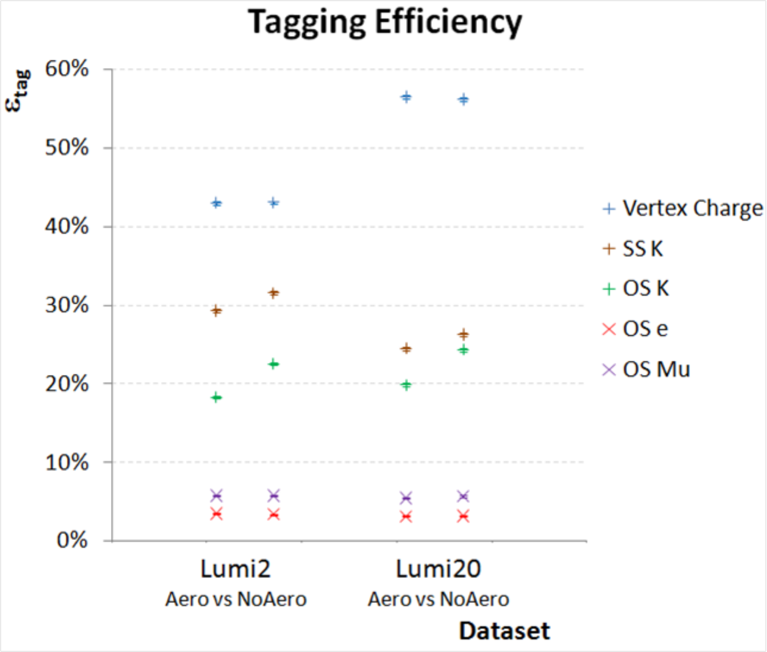

(a) Tagging efficiency for individual taggers

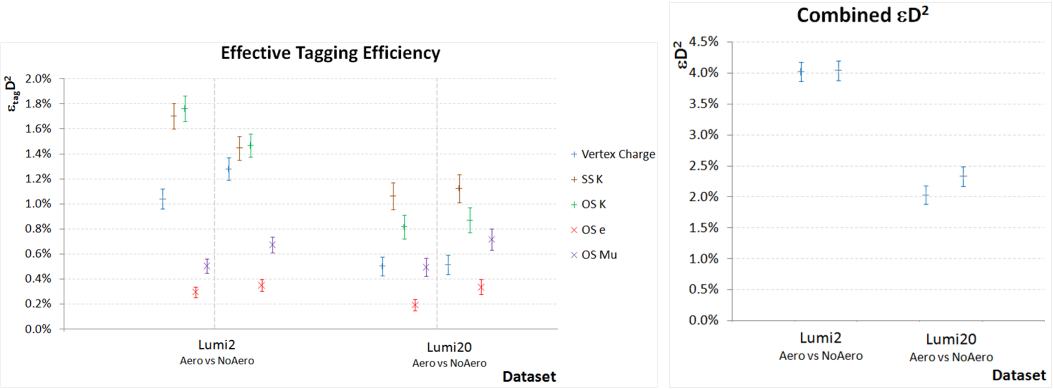

(b) Effective efficiency for individual taggers

(c) Effective efficiency of the combined tagging algorithms

Figure 5.12: Tagging performance in different environments using signal, minimal upgrade layout datasets. The left two sets of data points are at nominal luminosity and the right two are at upgrade luminosity. For each pair of sets, the left data points are when aerogel is in the simulations and the right data points have removed aerogel radiator.

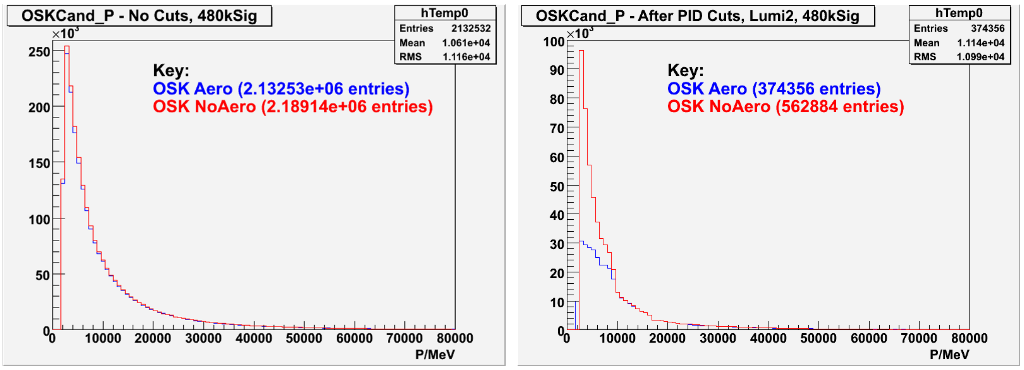

(a) Before any cuts by OS K (b) After PID cuts by OS K

Figure 5.13: Momentum distribution of pre-selected candidate tracks for the Other Side Kaon (OS K) tagger. These tracks are from events at nominal luminosity in the minimal upgrade layout where the $B_s \to \phi\phi$ has been selected. Results with aerogel are plotted in blue and results in the absence of aerogel are in red.

5.5.4 Kaon Taggers and Aerogel

As seen in Figure 5.12a, both kaon taggers are tagging more often and with an increased rate of making wrong tags. The cause of the higher tagging efficiency is shown in Figure 5.13 using candidate tracks for OS K as an example. The pre-selected list of candidate tracks remains unchanged by removing aerogel. The slightly higher statistics in Figure 5.13a for the dataset without aerogel is due to the slightly higher selection efficiency seen in Figure 5.10 resulting in a higher number of selected events being available for this plot.

The momentum distribution of these candidate tracks after the OS K PID cuts have been applied are shown in Figure 5.13b. The PID cuts are the first two cuts listed in Table 5.2. A similar distribution is seen for SS K candidate tracks after the PID cuts. Above 10 GeV, there is no difference between having and not having aerogel since this is beyond the momentum range of aerogel. Within the low momentum range where veto mode is used in the absence of aerogel, there are many more candidates passing the PID cuts.

The $\Delta \ln \mathcal{L}(K - \mathrm{p})$ distribution of the OS K candidate tracks passing the kaon-pion separation cut is shown in Figure 5.14. As was seen in Figure 5.11c, by removing aerogel a lot of tracks that otherwise would have had a kaon-proton distinction end up inconclusive in the PID algorithm. However in this particular distribution of pre-selected, non-signal tracks passing the $\Delta \ln \mathcal{L}(K - \pi) > 0$ cut, the majority of these tracks that are kaon-proton indistinguishable without

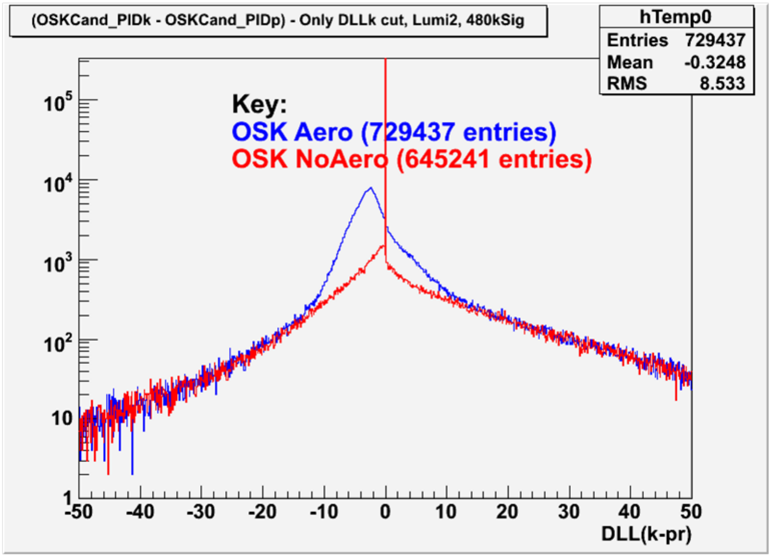

Figure 5.14: $\Delta \ln \mathcal{L}(K - \mathrm{p})$ distribution for candidate tracks of the other side kaon tagger. The events are from nominal luminosity Monte Carlo datasets in the minimal upgrade layout. The tracks have passed pre-selection and the OS K tagger's default $\Delta \ln \mathcal{L}(K - \pi) > 0$ cut. The blue curve is from the dataset with aerogel intact and the red distribution was simulated without aerogel.

aerogel would have been identified as protons.

Both kaon taggers have a default cut of $\Delta \ln \mathcal{L}(K - \mathrm{p}) > -1$ for kaon-proton separation. This cut keeps the large spike of $\Delta \ln \mathcal{L}(K - \mathrm{p}) = 0$ tracks. When aerogel is removed, this means many tracks that would have been rejected as they were identified as protons are now being retained as candidates for the kaon taggers. Thus the increased frequency that the kaon taggers find a track which passes all their candidate selection cuts raises their respective tagging efficiencies as shown in Figure 5.12a. Since these extra track candidates are likely to be protons that are being accepted while indistinguishable without aerogel, the mistag rate also increases. As a result, the effective efficiency of kaon taggers degrades at nominal luminosity by removing aerogel as shown in Figure 5.12b.

5.6 Conclusions

$B_s \to \phi\phi$ selection, kaon particle identification and flavour tagging as currently implemented for nominal luminosity still function in simulations of higher luminosity. But without adapting the algorithms for the new conditions the selection loses a third of its efficiency, kaon identification must find a new compromise between the mis-identification rate and efficiency while flavour tagging performance is halved at $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Simulations of the aerogel being removed showed RICH particle identification could still distinguish kaons from pions in the aerogel's momentum range by using veto mode. However, kaon-proton separation suffers with many low momentum tracks ending up indistinguishable without the aerogel radiator. The selection efficiency and lepton taggers stand to benefit a little from the reduced material budget by removing aerogel. At nominal luminosity, the degradation of kaon-proton separation is detrimental to kaon tagger performance, but this is not the case at upgrade luminosity with current algorithms. The net result is that flavour tagging is unchanged without the aerogel at nominal luminosity without any adaptations. At upgrade luminosity there is a potential improvement in the overall effective efficiency within statistical errors.

The conclusions on the aerogel removal studies are somewhat surprising. Further studies on this subject are presented in the next chapter, where some preliminary changes to the algorithms are explored.

Chapter 6

Reoptimising Flavour Tagging

As discussed in Section 5.1, an upgrade is planned for the LHCb experiment to run at $20 \times 10^{32}\,\mathrm{cm^{-2}\,s^{-1}}$ in the future. Chapter 5 presented how flavour tagging works and how its performance degrades at the upgrade luminosity if its algorithms are not adapted for the new conditions. The importance of kaon Particle Identification (PID) was shown and how removing the RICH aerogel radiator affects it.

There was urgent need for guidance on which layout the RICH should adopt for the upgrade: there were proposals to replace the aerogel radiator with a new time-of-flight RICH detector (TORCH) or even replace the whole of RICH1 with a larger RICH2 detector. Thus this thesis chose to further study kaon tagging by reoptimising it for upgrade scenarios. The event selection cuts, kaon Particle Identification (PID) algorithms and the neural net for combining all the individual tagger decisions in an event into a single tagging outcome were all left unchanged while the individual kaon taggers were reoptimised.

Only the nominal luminosity and $20 \times 10^{32}\,\mathrm{cm^{-2}\,s^{-1}}$ are considered in this chapter. However, at the time of writing after the studies presented in this chapter were complete, the LHCb collaboration is leaning towards lowering the upgrade luminosity to $10 \times 10^{32}\,\mathrm{cm^{-2}\,s^{-1}}$. Furthermore, the upgrade may happen with less than the $10\,\mathrm{fb^{-1}}$ target of data collected. The analysis presented in this chapter should be repeated at $10 \times 10^{32}\,\mathrm{cm^{-2}\,s^{-1}}$ to help assess the new upgrade plans, but are still valid to study detector performance at higher multiplicity.

The majority of the kaon tagging process involves taking the list of pre-selected tagger candidate tracks for the event and applying cuts to choose which track use for tagging. Section 6.1 studies other areas of the process subject to optimisation for upgrade luminosity. Section 6.2 discusses the methodology used to reoptimise the cuts for selecting a tagger candidate track. An account of how the cuts were adjusted at each step of the reoptimisation is described in Section 6.3 for one example scenario. The change in performance for the kaon taggers after the reoptimisation is presented in Section 6.4.

6.1 Tagging Track Candidates

6.1.1 Candidates Passing Pre-Selection

The pre-selection of tagging candidate tracks was described in Section 5.4.2. The distribution of how many pre-selected candidate tracks there are in events is shown in Figure 6.1. The tracks shown in the plots are from minimal upgrade layout $B_s \to \phi\phi$ signal datasets with aerogel at nominal and upgrade luminosities. Removing the aerogel has minimal effect as the loose cuts of the pre-selection listed in Table 5.4.2 do not involve PID checks.

As expected, for an equal number of $B_s \to \phi\phi$ selected events, there are more tracks passing pre-selection per event at upgrade luminosity. At nominal luminosity the mean of this distribution is 34.89 with an RMS of 20.78. At upgrade luminosity, the mean and RMS is 74.89 and 56.03 respectively. This is a consequence of higher track occupancy when luminosity is increased.

6.1.2 Candidates Passing Selection Cuts

Each tagger applies its own track selection cuts to its copy of the list of pre-selected candidates for each event. As explained in Section 5.4.2, when a single particle tagger has $N_{pass} > 1$ it uses the highest transverse momentum (p_T) candidate to make a tagging decision. Let N_{pass} be the number of candidate tracks that pass all selection cuts of a given single particle tagger in an event. The distribution of N_{pass} is illustrated in Figure 6.2 and is discussed below. A total of eight different scenarios are plotted based on two luminosities, whether aerogel is present or not and for OS & SS kaon taggers.

When a tagger finds no candidate tracks passing all its selection cuts, it does not make a tagging decision for that event. There are 30k out of 40k events with $N_{pass} = 0$ under all scenarios. This agrees with the tagging efficiency of kaon taggers shown back in Figure 5.9a. Of the events where a tagging decision is made,

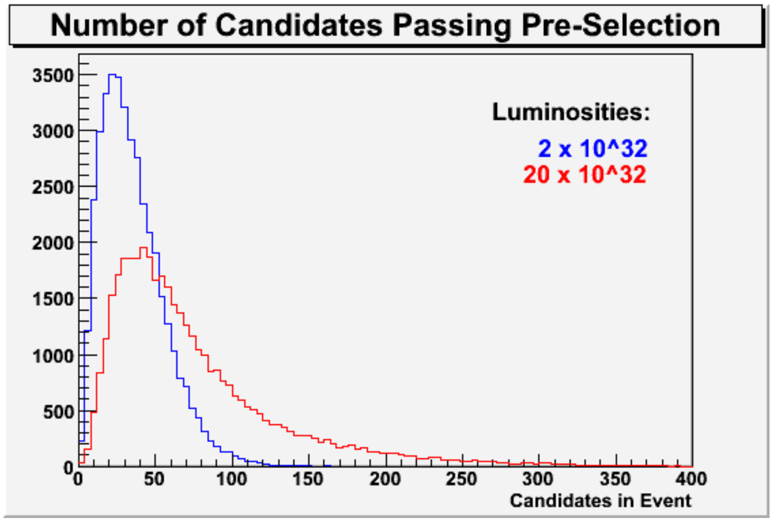

Figure 6.1: Distributions of the number of tagger candidate tracks passing pre-selection in each event. The samples come from simulated signal datasets at 2 and $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ in the minimal upgrade layout with aerogel , plotted in blue and red respectively. Only the first 40k events where the $B_s \to \phi\phi$ was successfully reconstructed from each sample are shown.

the majority have $N_{pass} = 1$. However a significant fraction of events have $N_{pass} > 1$. At nominal luminosity, both kaon taggers see an increase in N_{pass} when aerogel is removed. This agrees with the conclusions drawn in Section 5.5.4: without aerogel, kaon-proton separation cannot be done at low momentum. Candidate tracks that would have been rejected by the $\Delta \ln \mathcal{L}(K - \mathrm{p})$ cut get accepted as their $\Delta \ln \mathcal{L}(K - \mathrm{p})$ value increases to 0, the default value for indistinguishable tracks. These kaon-proton indistinguishable tracks increment N_{pass}.

The increased track multiplicity at upgrade luminosity has different effects on each kaon tagger. For the OSK tagger, N_{pass} increases compared to nominal luminosity. Removing aerogel increases the chance of getting a few candidates passing all cuts but decreases the chance of getting several candidates through the cuts. In the SSK case, increasing the luminosity has decreased the chance of finding candidates passing all cuts and removing aerogel further reduces this chance.

6.1.3 Events with Multiple Candidates Passing Selection

For events with $N_{pass} > 1$, the single particle taggers need to select one candidate. This ‘best’ tagging candidate is chosen as the track with highest transverse

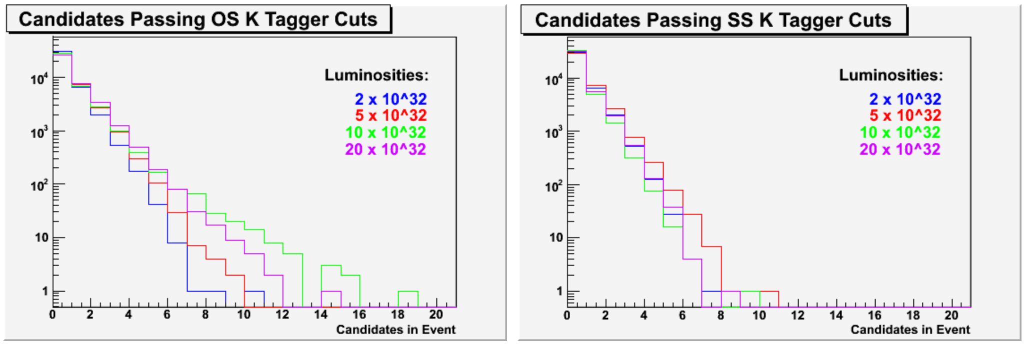

(a) Opposite side kaon tagger.

(b) Same side kaon tagger.

Figure 6.2: Distributions of number of tagging candidate tracks that have passed the selection cuts of a kaon tagger. These tracks come from the first 40k $B_s \to \phi\phi$ selected events in minimal upgrade layout datasets. Nominal luminosity data is plotted in blue and red for retaining and removing aerogel respectively. Green and purple represent retaining and removing aerogel at upgrade luminosity. Note the logarithmic y-axes.

momentum, p_T. The validity of this last step was tested for nominal and upgrade luminosity. This is important for the upgrade as $N_{pass} > 1$ events are more frequent with the higher track multiplicity.

Different methods for selecting the best tagging candidate in $N_{pass} > 1$ events were studied. The methods all rely on finding the track with the highest or lowest value of the following properties: impact parameter pile-up significance, I_{PileUp}; PID $\Delta \ln \mathcal{L}$; momentum, p and absolute impact parameter significance. These are described in Section 5.4.2.

These alternate methods only affect how the best tagging candidate is chosen, they do not change the tagger selection cuts themselves. Hence N_{pass} and therefore ϵ_{tag} remain unchanged regardless of which method is used, only the mistag rate, ω is affected. How ϵ_{eff} changed as a result of using these different methods is plotted in Figure 6.3. There were no statistically significant differences in ϵ_{eff} between the different methods. Therefore the existing method of choosing the highest p_T track as the best tagging candidate was kept.

For candidate tracks passing all tagger selection cuts, Monte Carlo truth was used to infer which candidates would give the correct tagging answer if selected by the tagger. Let N_{corr} be the number of such tracks in an event for a given tagger. Plots studying N_{corr} are presented in Figure 6.4 using the OSK tagger. Both the distribution of N_{corr} and the effective tagging efficiency, ϵ_{eff} for given values of N_{pass} are plotted, at nominal and upgrade luminosity. The N_{corr} distribution

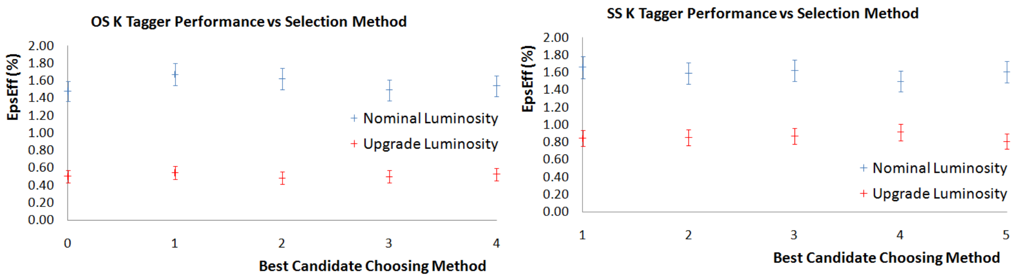

(a) Opposite side kaon tagger.

(b) Same side kaon tagger.

Figure 6.3: Effective tagging efficiency for different methods of selecting the kaon tagging track when many get through the tagger selection cuts. The different methods are labelled as follows: 0 for highest impact parameter pile-up (opposite side kaon only); 1 for highest $\Delta \ln \mathcal{L}(K - \pi)$; 2 for highest $\Delta \ln \mathcal{L}(K - \mathrm{p})$; 3 for highest momentum; 4 for highest transverse momentum and 5 for lowest $|\mathrm{IP}_{\mathrm{Sig}}|$ (same side kaon only).

plots, Figures 6.4a and 6.4b, do not show $N_{pass} < 2$ events as they dominate the statistics. For events with a given N_{pass}, the N_{corr} distribution is slightly biased towards being half of N_{pass} but is otherwise evenly spread at both luminosities.

Figure 6.4c is a nominal luminosity plot of the fraction N_{corr}/N_{pass} for events at each given N_{pass} value. There are few events with $N_{pass} > 5$, hence the rapidly increasing statistical errors beyond that threshold. The fraction of candidates passing all OSK cuts which then give the right tagging answer decreases significantly as N_{pass} increases. In Figure 6.4d, the upgrade luminosity plot, this fraction decreases much more slowly.

After the selection cuts, the taggers pick the highest p_T candidate out of N_{pass}. The probability of the OS K tagger to correctly tag the signal B-meson on events with a given N_{pass} is shown in Figure 6.5. At $N_{pass} = 1$, this is just $1 - \omega$. A deterioration of tagging performance at nominal luminosity is shown in Figure 6.5a when $N_{pass} > 2$. At upgrade luminosity, the performance is more stable as N_{pass} increases. Both plots suffer from reduced statistics at higher N_{pass} so for Figure 6.5b it is uncertain where exactly the performance starts to drop off.

Based on these observations, the kaon taggers were modified to leave events with $N_{pass} > 3$ as untagged. This threshold was chosen as a compromise between the two luminosities. Figure 6.2 shows that up to 2% of events are above this threshold for the OS K tagger cuts and <1% of events have $N_{pass} > 3$ for the SS K tagger cuts. The exact number of events treated as untagged depends on whether the dataset is at nominal or upgrade luminosity, and whether the

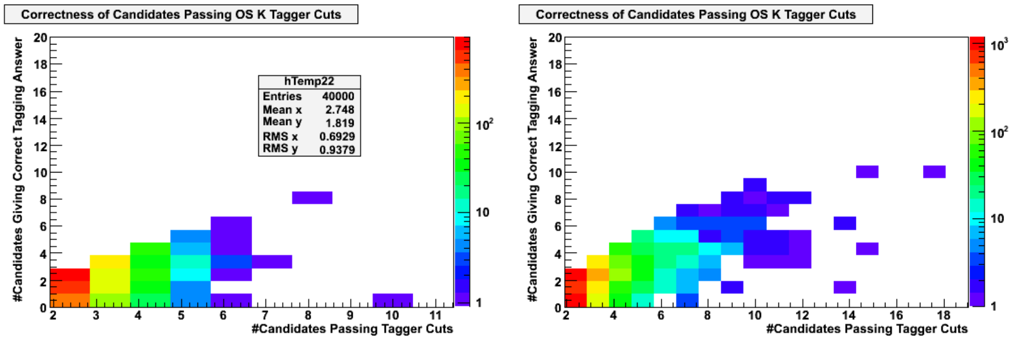

(a) Nominal Luminosity. (b) Upgrade Luminosity.

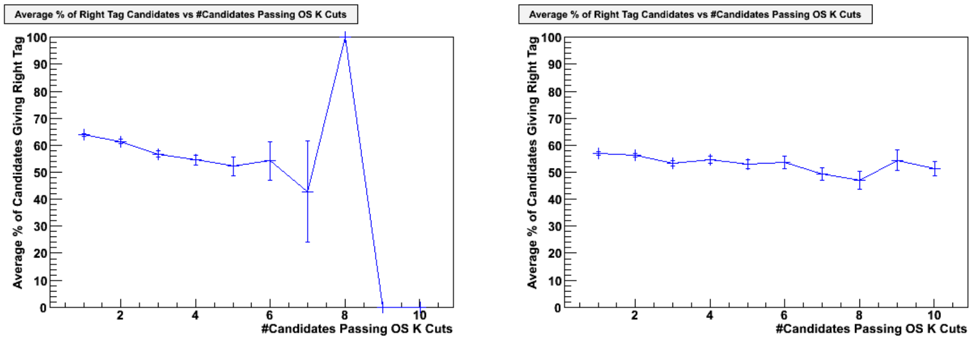

(c) Nominal Luminosity. (d) Upgrade Luminosity.

Figure 6.4: Distributions of track candidates for the opposite side kaon tagger. These candidates have passed all the selection criteria and the tagger will pick one of them to make a tagging decision with. The data come from the first 40k $B_s \to \phi\phi$ selected events from minimal upgrade layout datasets with aerogel at nominal and upgrade luminosities plotted on the left and right respectively. The top two plots compare how many such candidates are in each event against how many of them would give the right tagging answer if selected by the tagger. The bottom two plots show the fraction of candidates giving the correct tagging answer against the number of selected candidates in each event.

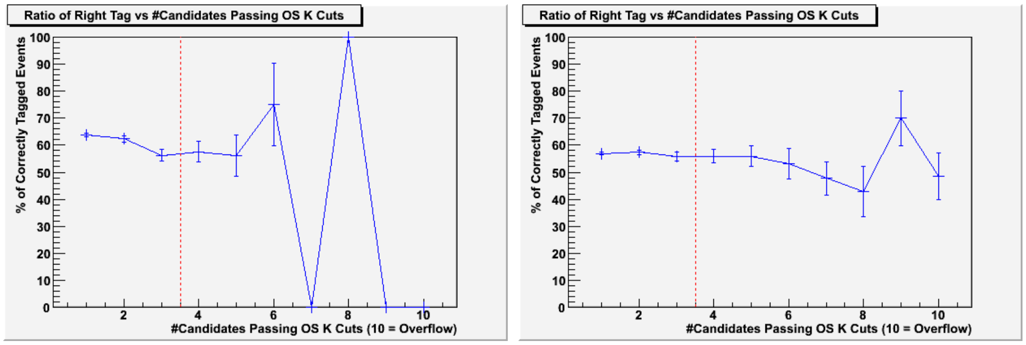

(a) Nominal luminosity.

(b) Upgrade luminosity.

Figure 6.5: Probability for the opposite side kaon tagger to tag the signal B-meson correctly for different numbers of candidate tracks passing all tagger selection cuts. The data comes from the first 40k $B_s \to \phi\phi$ selected events in the minimal upgrade layout with aerogel intact, for nominal and upgrade luminosity. The red dotted lines are the threshold where events were considered untagged if there were too many candidates passing the tagger selection cuts.

aerogel is removed or not. However on average, this $N_{pass} > 3$ threshold results in a relative loss in ϵ_{tag} of ~30% and ~25% for the OS K and SS K taggers respectively.

6.2 Reoptimisation Methodology

6.2.1 Preliminary Studies

After changing the kaon taggers to ignore events where $N_{pass} > 3$, individual tagger selection cuts were studied. Back in Section 5.5.4 it was shown how the $\Delta \ln \mathcal{L}(K - \mathrm{p})$ distribution of the tagger candidate tracks affected flavour tagging when aerogel was removed at nominal luminosity. How the OSK tagger's ϵ_{tag} , ω and ϵ_{eff} are affected by changing the $\Delta \ln \mathcal{L}(K - \pi)$ and $\Delta \ln \mathcal{L}(K - \mathrm{p})$ selection cuts are shown in Figures 6.6 and 6.7 respectively. The plots show data both from simulations with and without aerogel. The SSK tagger shows similar behaviour in ϵ_{eff} for the same changes.

ϵ_{eff} increases when either the $\Delta \ln \mathcal{L}(K - \pi)$ or $\Delta \ln \mathcal{L}(K - \mathrm{p})$ cut are made tighter than the current default value. Other cuts were also examined and some showed similar improvements in performance for both kaon taggers. As mentioned before, the existing flavour tagging algorithms were optimised for the $\mathrm{B}_s \to \mathrm{J}/\psi\phi$ channel. It is not surprising that the selection cuts for kaon taggers are not

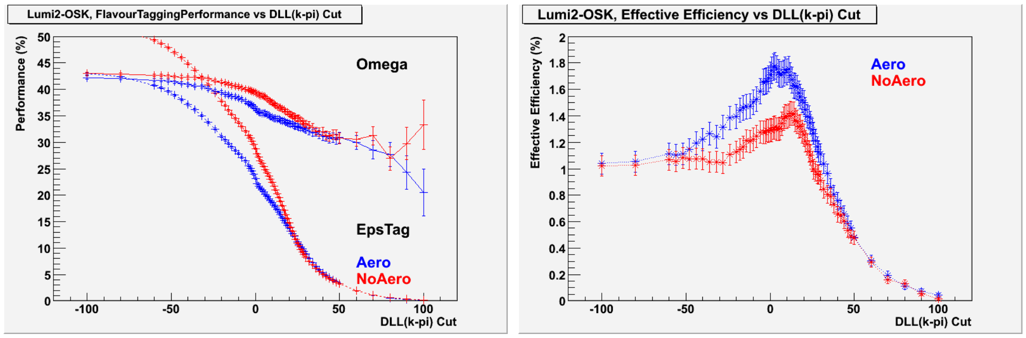

(a) Nominal luminosity, ω and ϵ_{tag}.

(b) Nominal luminosity, ϵ_{eff}.

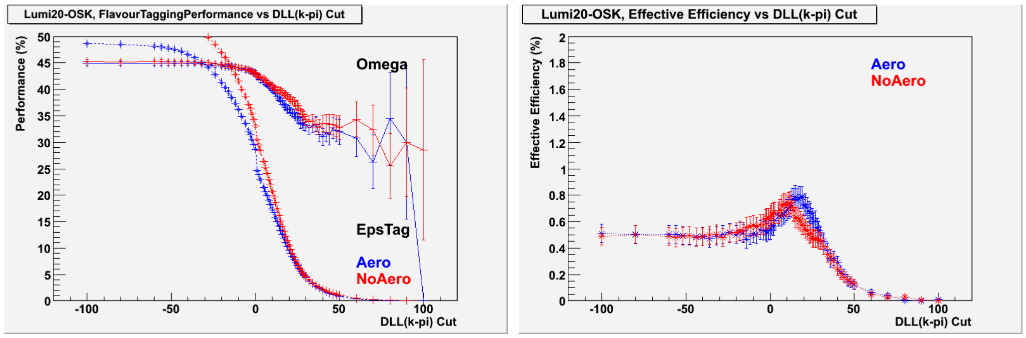

(c) Upgrade luminosity, ω and ϵ_{tag}.

(d) Upgrade luminosity, ϵ_{eff}.

Figure 6.6: Opposite side kaon tagger performance when its $\Delta \ln \mathcal{L}(K-\pi)$ cut is varied. All other selection cuts were unchanged. The data comes from 480k $B_s \to \phi\phi$ signal events from minimal upgrade layout datasets. The top two plots show nominal luminosity results and the bottom two plots are for upgrade luminosity. The left two plots show tagging efficiency, ϵ_{tag} and mistag rage ω. The right two plots show effective efficiency, ϵ_{eff}. All plots show results with aerogel intact in blue and with aerogel removed in red.

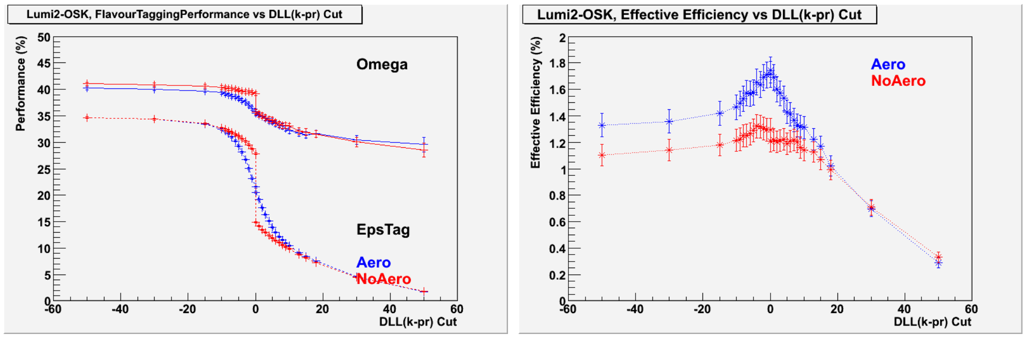

(a) Nominal luminosity, ω and ϵ_{tag}.

(b) Nominal luminosity, ϵ_{eff}.

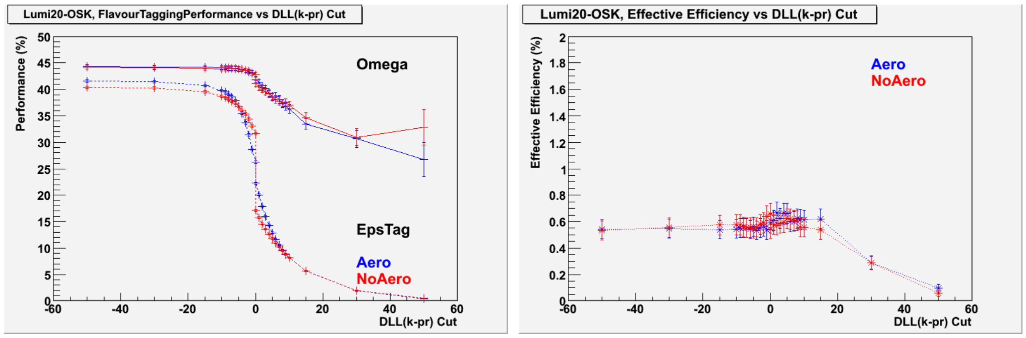

(c) Upgrade luminosity, ω and ϵ_{tag}.

(d) Upgrade luminosity, ϵ_{eff}.

Figure 6.7: Opposite side kaon tagger performance when its $\Delta \ln \mathcal{L}(K - \mathrm{p})$ cut is varied. All other selection cuts were unchanged. The data comes from 480k $B_s \to \phi\phi$ signal events from minimal upgrade layout datasets. The top two plots show nominal luminosity results and the bottom two plots are for upgrade luminosity. The left two plots show tagging efficiency, ϵ_{tag} and mistag rage ω. The right two plots show effective efficiency, ϵ_{eff}. All plots show results with aerogel intact in blue and with aerogel removed in red.

optimised even at nominal luminosity.

At nominal luminosity there is a significant relative decrease in ϵ_{eff} when aerogel is removed. At upgrade luminosity this relative difference is not observed. The optimum $\Delta \ln \mathcal{L}(K-\pi)$ and $\Delta \ln \mathcal{L}(K-\mathrm{p})$ cuts to maximise ϵ_{eff} are different between having or not having aerogel and nominal or upgrade luminosity.

Tracks that are indistinguishable have a $\Delta \ln \mathcal{L}$ value of 0. There is a large discontinuity in the curve for the aerogel removed scenario in Figure 6.7a at $\Delta \ln \mathcal{L}(K-\mathrm{p}) = 0$. No such discontinuity is observed at $\Delta \ln \mathcal{L}(K-\pi) = 0$ in Figure 6.6a. This reinforces the discussion back in Section 5.5.4: without aerogel the RICH cannot distinguish between kaons and protons for low momentum tracks. Kaons and pions can still be distinguished using veto mode as the momentum threshold for pions to emit Cherenkov light is much lower than for kaons and protons.

Just before the $\Delta \ln \mathcal{L}(K-\mathrm{p}) = 0$ threshold in Figure 6.7a, ϵ_{tag} and ω are both higher in the aerogel removed scenario. This is due to the cut retaining these kaon-proton indistinguishable tracks. When the $\Delta \ln \mathcal{L}(K-\mathrm{p})$ cut is moved to just above the threshold, ω drops to match the scenario with aerogel. ϵ_{tag} is actually lower in this region when aerogel is removed because some tracks that would have a low but positive $\Delta \ln \mathcal{L}(K-\mathrm{p})$ value also become indistinguishable so are not seen without aerogel.

In Figure 6.7a the discontinuities in ϵ_{tag} and ω at $\Delta \ln \mathcal{L}(K-\mathrm{p}) = 0$ only occur when aerogel is removed. At upgrade luminosity, as shown in Figure 6.7c, such discontinuities are less pronounced but occur for both scenarios with and without aerogel. Similar discontinuities in ϵ_{tag} are seen at $\Delta \ln \mathcal{L}(K-\pi) = 0$ in Figure 6.6c regardless of the aerogel scenario. This suggests that while the aerogel radiator is useful at nominal luminosity, it is less important at upgrade luminosities. The PID algorithms fail to identify some tracks even with the aerogel radiator under upgrade conditions.

These observations served as motivation to reoptimise as much as possible: each different simulation scenario is treated as a separate case to reoptimise for.

6.2.2 Recursive Optimisation

The strategy for reoptimising the kaon selection cuts was inspired by a recursive optimisation tool called CROP [35]. The tool took a signal and a background dataset and would optimise a group of user-specified cuts to maximise signal while minimising background retention. It would then optimise again and again until no further improvement in performance was found to within a user-specified resolution limit. The tool assumes there is only one candidate that needs to be selected per event: which is true in physics analysis when e.g. a B-meson with specific decay products are what is being searched for. With kaon tagger cut optimisation there can be events with many candidates after all the cuts, i.e. $N_{pass} > 1$. CROP was not designed to do the post-selection procedure in $N_{pass} > 1$ events to find the best candidate for tagging. Therefore a script was written for kaon tagger cut optimisation, based on the principle behind CROP as described below.

All the cuts in a given kaon tagger are optimised together. First, one cut is chosen and is varied with the other cuts at their default values. The cut is varied until a maximum ϵ_{eff} is found and the cut value at this point is remembered. Another cut is selected and individually optimised with the previous cut reset to its default value. This is repeated for all cuts in the tagger and then the cuts are set to each of their maximum ϵ_{eff} points as their new default values. This concludes a single round of the recursive optimisation. The process is then repeated for several rounds until no further gain in ϵ_{eff} is found.

The script optimises one kaon tagger at a time. For each given set of tagger cut values, the script runs over all the events in the input dataset. For each candidate track passing the secondary pre-selection in an event, all the tagger cuts are applied. If more than one track passes all the cuts, the script picks the highest transverse momentum track as the best tagging candidate. MC truth is consulted to determine if a tagging candidate found in an event would give the right tagging answer or not. The number of events with a right, wrong or no tagging decision are summed up to calculate ϵ_{tag} , ω and ϵ_{eff} after all the events are processed. This whole procedure is then repeated to evaluate another set of tagger cut values.

Due to the length of computation time involved for the recursive optimisation, the script was designed with flexibility in mind. Instead of sampling different cut values at fixed intervals over a wide range, the user would specify sample values manually. This allowed the user the option to go back and try out more sample cut values before moving onto the next round of optimisation. The end point of the recursive process was also left to the user to decide. The studies presented below stopped reoptimising when cuts were being sampled at a precision where

Cut Variable	L2Aero	L2NoAero	L20Aero	L20NoAero
Opposite Side Kaon Tagger				
$\Delta \ln \mathcal{L}(K-\pi) >$	1	7	11	6.5
p (GeV) $>$	3.6	2.6	6.0	4.65
p_T (GeV) $>$	0.4	0.45	0.65	0.58
|IP| (mm) $<$	2.6	1.3	1.4	1.4
LCS $<$	2.5	2	2.8	2.4
Same Side Kaon Tagger				
$\mathrm{I}_{PileUp} >$	3	3	5.5	4
$\Delta \ln \mathcal{L}(K-\pi) >$	4	5	3.5	8
p (GeV) $>$	2.2	4.0	4.0	4.3
p_T (GeV) $>$	0.4	0.5	0.7	0.48
$\lvert \mathrm{IP}_{\mathrm{sig}} \rvert <$	3.5	3.7	4.8	3.1
dη $<$	1.2	1.0	1.0	0.93
dϕ $<$	1.0	1.0	1.0	0.82
dQ (MeV) $<$	1800	2000	2000	2030
LCS $<$	3.5	3.0	4.0	3.9

Table 6.1: Cuts applied in the secondary pre-selection tailored for each kaon tagger. LCS is track χ^2 per degree of freedom.

the differences were on the order of changing the tagging outcome of one event.

6.3 Reoptimisation of Kaon Tagger Cuts

6.3.1 Secondary Pre-selection of Candidates

As the recursive reoptimisation was performed, it became apparent the length of computing time to reoptimise was too long to run the whole tagging process for every candidate track in each event. A secondary pre-selection was designed for each kaon tagger and simulation condition. The cuts of this second pre-selection were chosen by looking at the ϵ_{eff} distributions from early tagger cut reoptimisation results. These cuts are presented in Table 6.1 and are applied to the general pre-selected tagger candidate tracks of $B_s \to \phi\phi$ selected events. For the upgrade luminosity dataset with aerogel removed, this reduced the number of tracks to be processed by two orders of magnitude.

Cut Variable	Default	R1	R2	R3	R4	R5	R6
$\mathrm{I}_{PileUp} >$	3	3	4.5	4.5	4.5	4.2	4.3
$\Delta \ln \mathcal{L}(K-\pi) >$	1	4	4	9	5	9	5
$\Delta \ln \mathcal{L}(K-\mathrm{p}) >$	-1	-1	-2	-1	-1.2	-1.8	-1.9
p (GeV) $>$	4.0	4.0	4.5	4.5	4.5	4.5	4.5
p_T (GeV) $>$	450	450	450	500	500	500	500
$\lvert\mathrm{IP}_{\mathrm{sig}}\rvert <$	3.0	3.0	3.2	3.0	3.0	3.0	3.0
$\mathrm{d}\eta <$	1	0.7	0.9	0.9	0.9	0.9	0.89
$\mathrm{d}\phi <$	1	0.7	0.8	0.8	0.8	0.77	0.79
dQ (GeV) $<$	1.6	1.6	2.0	1.9	2.0	1.95	1.97
LCS $<$	2	2	3.5	3.3	3.8	3.5	3.7

Table 6.2: Table of how the same side kaon tagger's candidate track selection cuts changed throughout six rounds of recursive reoptimisation. The reoptimisation was performed on a minimal upgrade layout, $B_s \to \phi\phi$ selected signal dataset at upgrade luminosity without aerogel. These cut values are at the beginning of each round. LCS is track χ^2 per degree of freedom.

6.3.2 Reoptimisation Progress

Throughout the recursive optimisation, no single cut made a large difference in tagging performance. As an example, the history of the changing cut values of the SSK tagger is listed in Table 6.2 for the upgrade luminosity case with aerogel removed for each round of recursive optimisation.

The cuts found in the first round were identified by looking at the ϵ_{eff} distributions such as the example in Figure 6.6d. Subsequent rounds locked on any increase in ϵ_{eff} no matter how insignificant compared to the statistical errors. This meant that for more than one local maxima in the distribution, the other cuts could change which peak was highest resulting in the optimised cut oscillating between rounds as is the case for $\Delta \ln \mathcal{L}(K-\pi)$ from round 3 onwards. The sample range for later optimisation rounds got narrower to increase the resolution in finding ϵ_{eff} maxima. This is demonstrated in, for example, the $\mathrm{d}\phi$ cut values increasing in precision towards the end of the recursive process.

The gradual increase in ϵ_{eff} as a result of the cut changes shown in Table 6.2 is shown in Figure 6.8. After the recursive reoptimisation, the SS K tagger has a relative ϵ_{eff} increase of ~70% for this dataset. Data from the OS K tagger of the same dataset is also shown, which also shows a similarly large relative ϵ_{eff} increase. Note the figure plots the highest ϵ_{eff} value found at the end of each reoptimisation round. Some rounds saw a higher increase than others, making it

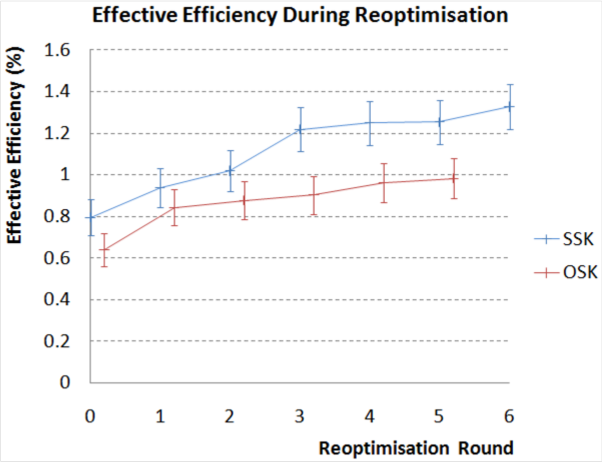

Figure 6.8: Effective tagging efficiency, ϵ_{eff} of the kaon taggers during the recursive reoptimisation of their tagger candidate selection cuts. ϵ_{eff} values shown were the highest found at the end of each reoptimisation round. The same side kaon tagger is plotted in blue and the opposite side kaon tagger is in red. The data is from a minimal upgrade layout $B_s \to \phi\phi$ signal dataset at upgrade luminosity and with aerogel removed. Round 0 is the original value of ϵ_{eff} before any reoptimisation.

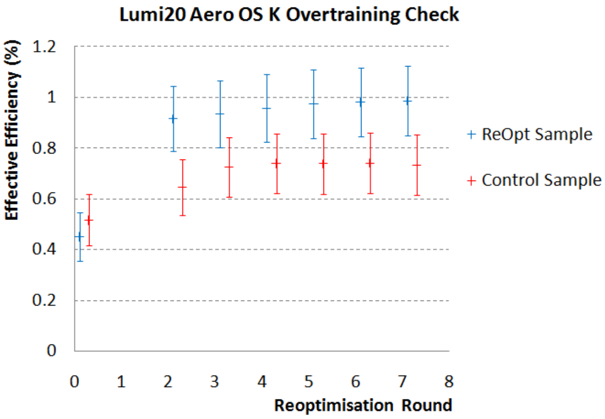

Figure 6.9: Values of effective tagging efficiency, ϵ_{eff} obtained from the recursive reoptimisation overtraining check. The dataset was split into two halves: a test and a control sample, plotted in blue and red respectively. The blue data points show the highest ϵ_{eff} found at the end of each reoptimisation round. The red data points are the ϵ_{eff} obtained from using the cut values at the end of the test sample's reoptimisation rounds. The first data points at round 0 is the ϵ_{eff} of each sample without any reoptimisation.

hard to use this plot to judge when to stop the recursive process.

The resolution for finding ϵ_{eff} maxima during reoptimisation was made tighter in the later rounds. The recursive procedure was stopped when the cut value changes were on the order of changing the tagging outcome of one or two events. A check was made to see how much overtraining this introduced. The upgrade luminosity dataset with aerogel was used for this check. This dataset was split into two halves: a test sample to reoptimise on and a control sample. After each round of reoptimising the OS K tagger's cut values on the test sample, those cuts were applied to the control sample. This check started with the cuts at the start of the second round of reoptimisation to speed up the process.

The results of this overtraining check are shown in Figure 6.9. The figure shows the reoptimisation progress in terms of ϵ_{eff} of both samples. As the ϵ_{eff} of the test sample increases throughout the reoptimisation, the control sample does not lose ϵ_{eff} at any round. Thus there is no overtraining observed by this check. However, the limited statistics in the samples lead to large error bars. The test and control samples each have 20k $B_s \to \phi\phi$ selected events. Larger datasets are recommended to better determine whether the recursive reoptimisation should be stopped earlier to avoid overtraining.

6.3.3 Tagger Selection Cuts After Reoptimisation

The cut values after the reoptimisation for each kaon tagger under each simulation scenario is presented in Table 6.3. The I_{PileUp} cut being changed is of interest as this used to be part of the original pre-selection cuts applied to tracks before the taggers received the list of candidates. Between two taggers, two luminosities and two aerogel scenarios some cuts can have remarkably different values for specific circumstances. For example, the OS K tagger has a positive $\Delta \ln \mathcal{L}(K - \mathrm{p})$ cut value at upgrade luminosity with aerogel in whereas every other case is more lenient towards protons being accepted.

6.4 Reoptimised Kaon Tagger Performance

As an example of how much ϵ_{tag} , ω and ϵ_{eff} changed, a comparison between the OSK tagger performance against the $\Delta \ln \mathcal{L}(K - \pi)$ and $\Delta \ln \mathcal{L}(K - \mathrm{p})$ cut distributions for the original selection cuts and the reoptimised cuts are presented in Figures 6.10 and 6.11 at upgrade luminosity. The top two plots of each figure

Cut Variable	L2Aero	L2NoAero	L20Aero	L20NoAero
Opposite Side Kaon Tagger				
$\mathrm{I}_{PileUp} >$	7.10	12.30	7.00	6.88
$\Delta \ln \mathcal{L}(K-\pi) >$	2.41	8.50	12.30	7.035
$\Delta \ln \mathcal{L}(K-\mathrm{p}) >$	-0.271	-3.605	2.055	-1.214
p (GeV) $>$	3.985	2.600	6.670	4.803
p_T (GeV) $>$	0.450	0.5024	0.709	0.5982
$\mathrm{IP}_{sig} <$	4.246	4.285	4.205	3.69
$\lvert\mathrm{IP}\rvert$ (mm) $<$	2.02	1.148	1.239	1.28
LCS $<$	1.875	1.775	2.467	2.00
Same Side Kaon Tagger				
$\mathrm{I}_{PileUp} >$	4.0	3.92	7.75	4.36
$\Delta \ln \mathcal{L}(K-\pi) >$	6.00	7.02	3.94	9.20
$\Delta \ln \mathcal{L}(K-\mathrm{p}) >$	-1.24	-5.22	-0.99	-1.77
p (GeV) $>$	2.40	4.800	5.0	4.52
p_T (GeV) $>$	0.468	0.55512	0.755	0.502
$\lvert\mathrm{IP}_{\mathrm{sig}}\rvert <$	3.133	3.506	4.445	3.00
$\mathrm{d}\eta <$	1.162	0.798	0.974	0.887
$\mathrm{d}\phi <$	0.950	0.896	0.901	0.795
dQ (MeV) $<$	1450	1704.4	1960	1960
LCS $<$	2.91	2.91	3.94	3.71

Table 6.3: The reoptimised tagger selection cuts after the recursive iterations finished. Each luminosity and aerogel condition was separately optimised for. LCS is track χ^2 per degree of freedom.

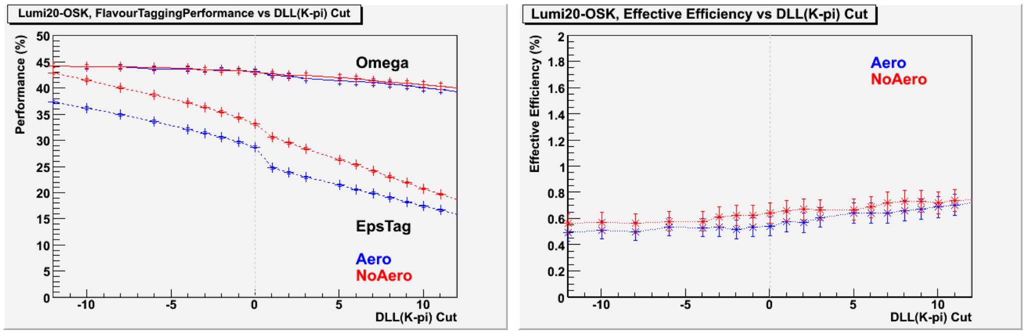

(a) Before reoptimisation, ω and ϵ_{tag}.
(b) Before reoptimisation, ϵ_{eff}.

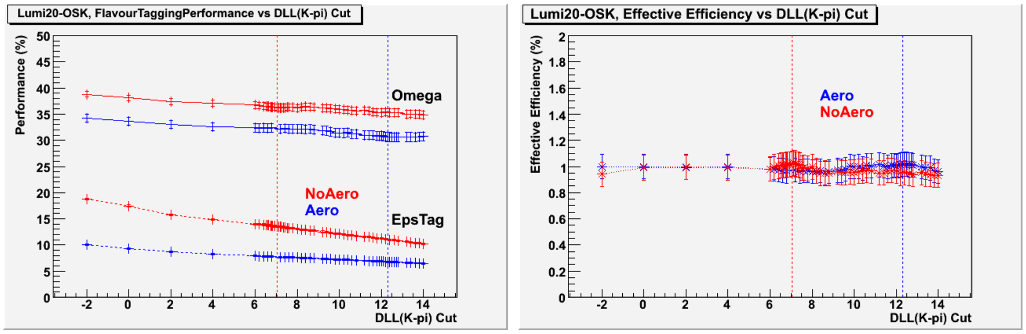

(c) After reoptimisation, ω and ϵ_{tag}.
(d) After reoptimisation, ϵ_{eff}.

Figure 6.10: Opposite side kaon tagger reoptimised performance when its $\Delta \ln \mathcal{L}(K-\pi)$ cut is varied. The data comes from 480k $B_s \to \phi\phi$ signal events from minimal upgrade layout datasets at upgrade luminosity. The left two plots show tagging efficiency, ϵ_{tag} and mistag rage ω. The right two plots show effective efficiency, ϵ_{eff}. The performance before and after reoptimisation in the top and bottom pair of plots respectively. All plots show results with aerogel intact in blue and with aerogel removed in red. The dotted lines show where the $\Delta \ln \mathcal{L}(K-\pi)$ cut is applied. After reoptimisation, the cut value is different for whether aerogel is present or not.

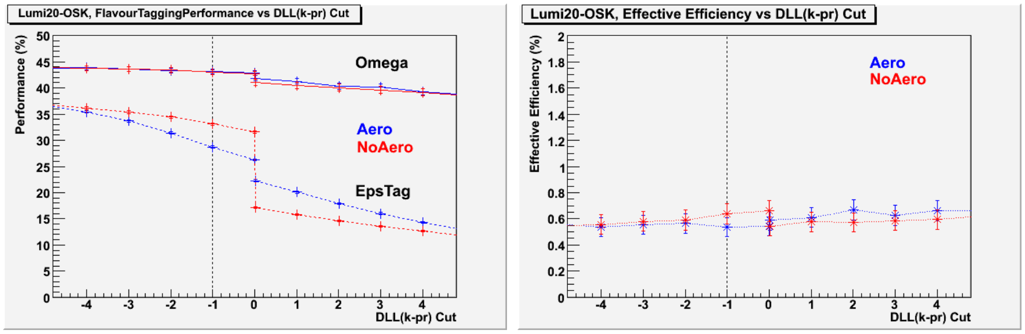

(a) Before reoptimisation, ω and ϵ_{tag}. (b) Before reoptimisation, ϵ_{eff}.

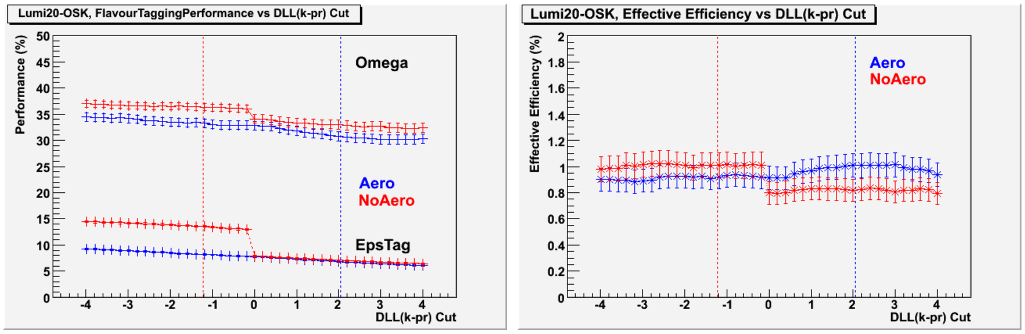

(c) After reoptimisation, ω and ϵ_{tag}. (d) After reoptimisation, ϵ_{eff}.

Figure 6.11: Opposite side kaon tagger reoptimised performance when its $\Delta \ln \mathcal{L}(K - \mathrm{p})$ cut is varied. The data comes from 480k $B_s \to \phi\phi$ signal events from minimal upgrade layout datasets at upgrade luminosity. The left two plots show tagging efficiency, ϵ_{tag} and mistag rage ω. The right two plots show effective efficiency, ϵ_{eff}. The performance before and after reoptimisation in the top and bottom pair of plots respectively. All plots show results with aerogel intact in blue and with aerogel removed in red. The dotted lines show where the $\Delta \ln \mathcal{L}(K - \mathrm{p})$ cut is applied. After reoptimisation, the cut value is different for whether aerogel is present or not.

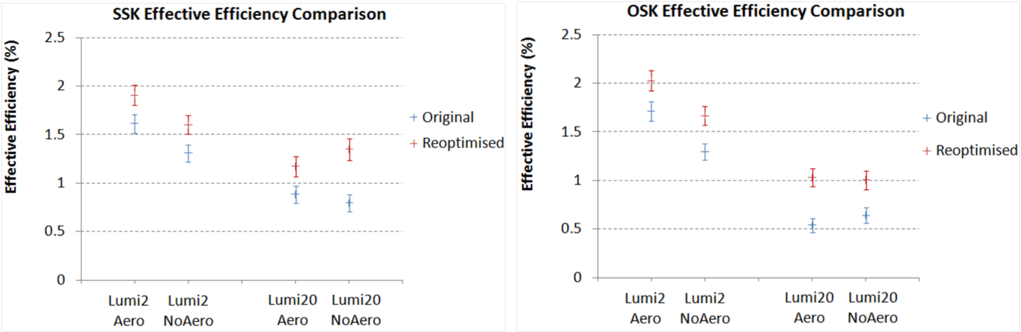

(a) Same side kaon tagger. (b) Opposite side kaon tagger.

Figure 6.12: Comparison of effective tagging efficiency for the kaon taggers before and after their reoptimisation, plotted in blue and red respectively. The recursive reoptimisation was applied to 480k events from minimal upgrade layout $B_s \to \phi\phi$ signal datasets.

are a zoomed in versions of the upgrade luminosity plots in Figures 6.6 and 6.7 respectively, with an indication of the original cut value. The bottom two plots of each figure show how tagger performance is affected when the $\Delta \ln \mathcal{L}(K-\pi)$ or $\Delta \ln \mathcal{L}(K-\mathrm{p})$ cut value is varied and the other cuts are at their reoptimised values.

After reoptimisation, there is still a trend of both ϵ_{tag} and ω decreasing as the $\Delta \ln \mathcal{L}(K-\pi)$ and $\Delta \ln \mathcal{L}(K-\mathrm{p})$ cut values are increased. In the scenario with aerogel, there is no longer a discontinuity at $\Delta \ln \mathcal{L}(K-\mathrm{p}) = 0$ and it is much smaller when aerogel is removed. The ϵ_{eff} distributions become flatter after reoptimisation compared to before, as seen by comparing Figure 6.6d and 6.10d. This shows how the reoptimisation is an accumulation of small increases in ϵ_{eff} over many cuts: no individual cut dominates the ϵ_{eff} optimisation.

How much the ϵ_{eff} of each kaon tagger was increased by the recursive optimisation is presented in Figure 6.12. The ϵ_{eff} of each kaon tagger is plotted before and after reoptimisation at both luminosities and both aerogel scenarios. The improvement at nominal luminosity was noted in Section 6.2.1 as the result of reoptimising in a different decay channel than the one originally used to optimise flavour tagging. Removing the aerogel is detrimental to kaon tagging even after the taggers are reoptimised at this luminosity.

At upgrade luminosity after reoptimisation, the SSK tagger might see an improvement without aerogel but the difference is not statistically significant. It is probable however, that removing aerogel will not deteriorate kaon tagging at $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. It is also notable that, especially for the OSK tagger's case, there is a large relative drop in ϵ_{eff} for kaon tagging despite the recursive reoptimisation efforts by going to upgrade luminosity.

6.5 Conclusions

The process of flavour tagging with kaons has been studied and reoptimised using $B_s \to \phi\phi$ signal datasets simulated in a minimal upgrade layout. Besides the changes to the tagger selection cuts, a new cut was implemented to discard events where too many track candidates pass all the selection cuts. The tagger selection cuts were optimised using a recursive routine. Both luminosities, with and without aerogel see improvement in effective tagging efficiency for both kaon taggers. This is partially due to these taggers having originally been optimised with the $\mathrm{B}_s \to \mathrm{J}/\psi\phi$ channel. The cuts have yet to be checked if they are over-trained on the limited statstics using an independent dataset. The reoptimised results suggest aerogel has no significant effect on kaon tagging at upgrade luminosity though is clearly beneficial at nominal luminosity. Of concern is the large relative drop in performance by increasing luminosity from 2 to $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ despite this reoptimisation.

There have also been a few data collection runs at the present LHC where luminosity slipped higher than the designed nominal threshold. Experience from trying to analyse data from these runs combined with other upgrade studies suggest $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ is too high a goal for the upgrade and a more modest $10 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ is being considered. On the other hand, the alternate layouts for the LHCb RICH system may help recover more of the kaon tagging performance lost by increasing luminosity, such as the time-of-flight RICH detector (TORCH) helping out with low momentum track identification. Studies incorporating these different designs will be necessary before a firm conclusion can be drawn on this issue. The reoptimisation study presented in this chapter should also be done at $10 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ to help understand what the optimal luminosity for the upgrade is. It would be of great interest to extend the study to include the other single particle taggers and other decay channels.

Chapter 7

Conclusions

The LHCb experiment finally started collecting data in early 2010 as the LHC began colliding beams. LHCb is now gathering data to push the boundaries of precision B-physics measurements. With its startup, the search for answers to the matter-antimatter asymmetry in the Universe has picked up pace.

Part of the reason LHCb is now the best precision B-physics experiment is its RICH counters. These detectors are used for particle identification, allowing for better signal selection amidst charged particle backgrounds in many decay channels. The RICH uses HPDs to detect the Cherenkov photon rings emitted by charged tracks passing through its radiators. 561 HPDs were assembled and each was characterised to cross-examine the manufacturer's quality assurance. Over 97% met or exceeded their specifications. QE was one of the specifications that exceeded expectations and was further improved throughout the manufacturing period.

Monitoring routines were set up for the RICH detectors. HPD vacuum quality was monitored by measuring IFB. During the commissioning of the RICH detectors, some HPDs were observed to degrade faster than others over the months. QE tests on these degraded showed increased QE values due to IFB signals contributing to the total current read out from the HPD. A few HPDs had severely degraded vacuums and started glowing. These HPDs could no longer be used in the RICH and needed to be replaced. IFB data from the HPD monitoring and from the PDTF tests were used to develop a model for each HPD's rate of vacuum degradation. It predicted when HPDs would be at risk of failure and the model has had no false positives. This has guided the judicious use of limited spare HPDs and was used in negotiations with the manufacturer for repair and replacements.

Though LHCb will gather more statistics in B-physics measurements than any previous experiment, there is demand for even more. For example, CP violation in $B_s \to \phi\phi$ cannot be measured precisely enough to test the SM prediction within a reasonable time frame if the LHCb remains at $2 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. Therefore the LHCb detector will be upgraded to run at a higher luminosity. R&D is ongoing to prepare for this upgrade even as LHCb is collecting its first data. One major subject was whether the RICH detectors should change their layout.

How flavour tagging and kaon particle identification were affected by higher luminosity was studied. MC $B_s \to \phi\phi$ signal datasets were used in full-detector simulations at different luminosities. Thee existing analysis software could still flavour tag at higher luminosities but with degraded performance. The kaon taggers were recursively reoptimised to see how much performance they could recover at upgrade luminosity. While there were large relative gains in effective tagging efficiency, the kaon taggers could not perform as well as at nominal luminosity. A simulated scenario without the RICH1 aerogel radiator showed the aerogel was responsible for a large relative gain in kaon tagger performance at nominal luminosity. However at $20 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, the aerogel had no effect on flavour tagging performance. This did not change even after reoptimisation.

Bibliography

[1] A. Sakharov. *Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe*. Journal of Experimental and Theoretical Physics Letters, 5, 24-27, 1967.

[2] J. Christenson, J. Cronin, V. Fitch, R. Turlay. *Evidence for the 2π decay of the K_2^0 meson*. Physical Review Letters, 13, 138-140, 1964.

[3] M. Gavela, P. Hernández, J. Orloff, O. Pène. *Standard model CP-violation and baryon asymmetry*. Nuclear Physics B, 430, 1994.

[4] K. Nakamura *et al.* (Particle Data Group). Journal of Physics G 37, 075021 (2010)

[5] J. C. Ward. *An Identity in Quantum Electrodynamics*. Physical Review Letters, 78, 182-182, 1950.

[6] M. Kobayashi and T. Maskawa. *CP-violation in the renormalizable theory of weak interaction*. Progress of Theoretical Physics, 49, 652-657, 1973.

[7] L. Wolfenstein. *Parameterization of the Kobayashi-Maskawa matrix*. Physical Review Letters, 51, 1945-1947, 1983.

[8] `http://ckmfitter.in2p3.fr`

[9] C. Wu, E. Ambler, R. Hayward, D. Hoppes, R. Hudson. *Experimental test of parity conservation in beta decay*. Physical Review, 105, 1413-1415, 1957.

[10] G. Branco, L. Lavoura, J. Silva. *CP Violation*. Oxford Science Publications. `ISBN 0 19 850399 7`

[11] M. Bartsch, G. Buchalla, C. Kraus. *$B \to V_L V_L$ Decays at Next-to-Leading Order in QCD*. `arXiv:0810.0249v1`.

[12] M. Rescigno *et al.* (the CDF collaboration). *First evidence for* $B_s \to \phi\phi$ *and penguin B decays at CDF.* Proceedings of the 32nd International Conference on High Energy Physics, 1113-1117. World Scientific, 2004.

[13] N. Styles *et al. Measuring the weak phase in the decay* $B_s \to \phi\phi$ *at the LHCb experiment.* LHCb-PUB-2009-025.

[14] D. Horn, *Study of* $B_s^0 \to D_s^{+(*)} D_s^{-(*)}$ *and* $B_s^0 \to \phi\phi$ *decays at CDF II.* Presented at the 2009 Europhysics Conference on High Energy Physics.

[15] LHCb Collaboration. *Expression of Interest for an LHCb Upgrade.* CERN-LHCC-2008-007.

[16] LHCb Collaboration. *LHCb reoptimized detector design and performance.* CERN-LHCC-2003-030.

[17] A. Augusto Alves Jr *et al.* (the LHCb Collaboration). *The LHCb Detector at the LHC.* Journal of Instrumentation, 3, S08005, 2008.

[18] LHCb Collaboration. *LHCb VELO technical design report.* CERN-LHCC-2001-011.

[19] LHCb Collaboration. *LHCb Magnet technical design report.* CERN-LHCC-2000-007.

[20] LHCb Collaboration. *LHCb Inner Tracker technical design report.* CERN-LHCC-2002-029.

[21] LHCb Collaboration. *LHCb RICH technical design report.* CERN-LHCC-2000-037.

[22] LHCb Collaboration. *LHCb Calorimeters technical design report.* CERN-LHCC-2000-036.

[23] LHCb Collaboration. *LHCb Muon System technical design report.* CERN-LHCC-2001-010.

[24] LHCb Collaboration. *LHCb Trigger System technical design report.* CERN-LHCC-2003-031.

[25] J. McCarron. *PhD Thesis, School of Physics, University of Edinburgh.* 2007.

[26] S. Eisenhardt. *Production and tests of Hybrid Photon Detectors for the LHCb RICH detectors*. Nucl. Instr. and Meth. A 595 (2008) 142–145.

[27] R. Lambert, Y.M. Kim[1], *et al. Quantum Efficiency of LHCb RICH Pixel Hybrid Photon Detectors*. LHCb-note-2008-025.

[28] T. Gys, N. Kanaya, D. Piedigrossi. *Quantum efficiency measurements of LHCb RICH pre-series pixel HPDs*. LHCb-RICH-2007-136.

[29] C. Blanks, F. Soomro, J. Nardulli. *The Laser Alignment and Monitoring System for the LHCb RICH detectors*. LHCb-INT-2010-039.

[30] LHCb Collaboration. *Photon Detectors for the LHCb RICH: addendum to the Technical Design Report*. CERN-LHCC-2003-059.

[31] T. Gys. *Report on the HPD status*. LHCb-2009-016.

[32] M. Calvi, O. Dormond, M. Musy. *LHCb Flavour Tagging Performance*. LHCb-2003-115.

[33] M. Calvi, O. Leroy, M. Musy. *Flavour Tagging Algorithms and Performances in LHCb*. LHCb-2007-058.

[34] R. Forty and O. Schneider. *RICH pattern recognition*. LHCB-98-040.

[35] C. Fitzpatrick. *simpletools: Handy command line tools for ntuple manipulation and analysis*. LHCb-INT-2009-029.

---

[1] legally changed name to G.Kim in 2010