A Study of Central Exclusive Production at LHCb

Scott R. Stevenson
St. Hugh's College, Oxford

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2015

A Study of Central Exclusive Production at LHCb

Scott R. Stevenson
St. Hugh's College, Oxford

A thesis submitted in fulfilment of the requirements for the degree of Doctor of Philosophy at the University of Oxford

Trinity Term, 2015

Abstract

This thesis presents an analysis of the central exclusive production (CEP) of $\chi_c(1P)$ mesons in proton-proton collisions at a centre-of-mass energy of 8 TeV, using data collected with the LHCb detector corresponding to an effective integrated luminosity of $(126.3 \pm 4.3)\,\text{pb}^{-1}$. Candidate $\chi_c(1P)$ decays are reconstructed in the K^+K^- and $\pi^+\pi^-$ final states, where the final state particles have pseudorapidities between 2.5 and 4.5. The products of the cross sections and branching fractions are measured as

$$\sigma \times \mathcal{B}(\chi_{c0}(1P) \to K^+K^-) = 28.0 \pm 2.8 \pm 13.9\,\text{pb},$$
$$\sigma \times \mathcal{B}(\chi_{c0}(1P) \to \pi^+\pi^-) = 20.3 \pm 2.4 \pm 8.1\,\text{pb},$$
$$\sigma \times \mathcal{B}(\chi_{c2}(1P) \to K^+K^-) < 12.1\,\text{pb},$$
$$\sigma \times \mathcal{B}(\chi_{c2}(1P) \to \pi^+\pi^-) < 10.5\,\text{pb},$$

where the first uncertainties are statistical and the second systematic, and limits are set at the 95 % confidence level. A major source of systematic uncertainty is the determination of the fraction of the selected data which is CEP rather than inelastic backgrounds.

Also described is HERSCHEL, a new subdetector installed at LHCb during the first long shutdown of the LHC. This is a system of forward shower counters which will provide CEP analyses with event-by-event rejection of inelastic backgrounds. The pseudorapidity coverage and detection efficiency of HERSCHEL are determined using Monte Carlo simulations. The pseudorapidity coverage extends to $|\eta| \geq 10$ due to the showering of collision products at the aperture limit. With a photoelectron threshold yield of 0.2 times the yield in the mean minimum bias event, the detection efficiency for single diffractive events is predicted to be 89 % in the forward direction and 91 % in the backward direction.

In memory of Leslie Scheftsik

Acknowledgements

The journey I've undertaken since arriving in Oxford which culminates in this thesis has not followed a direct course, and there are a number of people I wish to thank for their support and guidance along the way.

It is conventional that one first thanks their supervisor, and here entirely appropriate. I am indebted to Guy Wilkinson for his expert supervision, and for the opportunity to join the LHCb group. His comprehensive guidance and feedback on all aspects of my work have been invaluable, and I've learnt a lot about how to conduct research from him.

The analysis presented within benefited greatly from the assistance of the other members of the LHCb CEP working group. I'm grateful to Kazu Akiba, Victor Coco, and Paula Collins for the ideas and suggestions they offered, always with cheer. Chris Thomas ensured I hit the ground running in the initial stages of the analysis, and helped me gain familiarity with the LHCb software quickly. I was fortunate to work with Heinrich Schindler in the HERSCHEL group, and enjoyed our collaboration.

I thank the staff of the UK Liaison Office at CERN, and Sue Geddes and Kim Proudfoot in the Particle Physics secretariat in Oxford, for their assistance on many occasions.

My fellow graduate students ensured the office always provided a pleasant working environment. I especially enjoyed discussing a myriad of obscure shared interests with Nazim Hussain and Donal Hill. In the wider Department, I am grateful for the friendship of Neven Blaskovic Kraljevic, Amy Cottle, Shaun Gupta, James Henderson, Krishanu Majumdar, and others. I particularly appreciate the assistance of Neven and Krish whilst completing this thesis. At St. Hugh's, I thank Fenner, Joe, Katharina, and Lottie for helping keep things in perspective, and for our adventures over the years.

Finally, I thank my parents and brother for the continued support they have given me in innumerable ways; I wouldn't have reached here without them.

Contents

List of Figures

List of Tables

1

Introduction

Particle Physics was born from the desire to understand the Universe at the most fundamental level. In practice it is the study of the elementary particles which constitute matter, and the fundamental interactions between them. The culmination of these studies within the past century has been the formulation and experimental validation of the Standard Model, a quantum field theory which classifies all the known elementary particles and describes three of the fundamental interactions between them.

Its development to the present day has been a collaborative effort by the theoretical and experimental communities, with major advances made on both sides. The last elementary particle of the Standard Model to be observed, the Higgs boson, was predicted in 1964 in a now famous set of papers published in *Physical Review Letters*, but experimental evidence of its existence would not come until its observation at the Large Hadron Collider (LHC) in 2012. In contrast, the discovery of a new particle in electron-positron annihilation in 1975, which would later be named the tau lepton, could not be explained by the conventional theory of the time.

As in the cases of the Higgs boson and tau lepton, most experimental discoveries have been made using colliders. Two beams of charged particles are accelerated to high energy and brought together into violent collision in the centre of a detector. The energy of the collision allows new particles, perhaps hitherto unknown, to be created, which exist for a fraction of a second before they decay into other particles. The detector makes measurements of these decay products as they move away from the collision point, from which the particle originally produced in the collision is reconstructed, and its properties measured.

Most studies at hadron colliders are of inelastic collisions, where the original particles in the opposing beams do not survive the collision, and a large number of new particles are produced. However, another type of collision is possible in which the particles in the opposing beams remain intact after a glancing collision. The colliding particles interact by the exchange of mediator particles between them, and a small number of new particles may be produced in the interactions of these mediators. This process has the unique experimental signature of a handful of particles in an otherwise empty detector, and is named central exclusive production (CEP).

The details of CEP give rise to several characteristic properties which are not replicated in inelastic collisions. Among these are spin selection rules, which mean that simply observing a particle to be produced via CEP is sufficient to determine several of its intrinsic quantum numbers. CEP has therefore seen attention in recent years as a potential way of searching for, and studying the quantum numbers of, hypothetical new particles, and this has culminated in the development of a theoretical description known as the Durham model.

The validity of the Durham model can be tested, and its predictions for new particles benchmarked, by measuring the CEP of established Standard Model particles such as charmonia. With this motivation, this thesis presents a study of the CEP of $\chi_c(1P)$ mesons reconstructed in the K^+K^- and $\pi^+\pi^-$ final states using the LHCb detector at the LHC.

Chapter 2 reviews the Standard Model, focusing on the strong interaction, before introducing diffraction in particle collisions and the specific case of CEP. The Durham model is described and its predictions for $\chi_c(1P)$ production at the LHC presented. The LHCb experiment is described in Chapter 3, including the detector itself in addition to the hardware and software triggers, and event simulation, reconstruction, and analysis software.

The $\chi_c(1P)$ analysis is then presented in detail, with Chapter 4 discussing the selection used to identify $\chi_c(1P) \rightarrow h^+h^-$ decays where h is kaon or pion. The selection efficiency is evaluated using simulated events, with data-driven corrections where necessary. Chapter 5 first describes the invariant mass fit used to measure the signal yields in data, and then the determinations of the exclusive purity of the selected data set and the effective integrated luminosity. Systematic uncertainties are evaluated, and the cross sections calculated and compared to the predictions of the Durham model and previous experimental measurements.

The predominant backgrounds to the $\chi_c(1P)$ measurement, and indeed all studies of CEP carried out at LHCb thus far, are events similar to CEP but where one of the colliding particles does not remain intact or where there is additional final state gluon radiation. In these cases, the spin selection rules which make CEP an attractive process with which to study the production of new particles no longer apply. The large uncertainty on CEP measurements introduced by these backgrounds has motivated the development and installation of a new LHCb subdetector named HERSCHEL, designed to provide event-by-event rejection of inelastic backgrounds.

The motivations for HERSCHEL are described in Chapter 6, along with the identification of suitable locations for the installation of the detectors and their mechanical and electronic design. The development of a geometric model of the new detectors and the surrounding LHC tunnel is presented, for use in simulating their performance. Chapter 7 describes single particle and full event simulations carried out to model the expected behaviour of HERSCHEL in realistic operating conditions, and the inelastic background rejection efficiency attainable in operation is predicted.

Finally, Chapter 8 concludes and provides an outlook on the future of CEP and other forward physics measurements at LHCb during Run 2 of the LHC.

2

Theoretical overview

This chapter reviews the Standard Model and the elementary particles whose interactions it describes. The formulation of the Standard Model as a quantum field theory is presented, focusing on the strong interaction. The process of diffraction in high energy particle collisions, mediated by the strong interaction, is described, and the origins of Regge theory and the Pomeron introduced. Finally, central exclusive production is discussed. The selection rules which apply to CEP are described, and the Durham model of CEP and its predictions for $\chi_c(1P)$ production at the LHC are presented.

2.1 The Standard Model

The Standard Model is a theory describing the elementary constituents of matter and three of the fundamental interactions between them: the electromagnetic, weak, and strong interactions. It is a quantum field theory unifying special relativity and quantum mechanics, and describes the fundamental interactions in terms of the exchange between matter particles of force carrying mediators [1]. A particle participates in an interaction only if it carries the charge of that interaction: electric charge for the electromagnetic interaction, weak isospin for the weak interaction, and colour charge for the strong interaction.

2.1.1 The elementary particles

The elementary particles are classified into two groups based on their intrinsic angular momentum, as given by their spin quantum number [2]. Each particle has an antiparticle with opposite charges; some particles are their own antiparticle. The bosons have integer spin and comprise the gauge bosons and the Higgs boson, whilst the fermions have half-integer spin, therefore observing the Pauli exclusion principle, and are the constituents of ordinary matter.

The bosons are further divided into the gauge bosons with spin 1, and the Higgs boson with spin 0. The gauge bosons are listed in Table 2.1 and are the photon (γ), which mediates the electromagnetic interaction, the massive weak bosons ($W^{\pm}$ and Z), which mediate the weak interaction, and the gluons (g), which mediate the strong interaction. The photon is the familiar quanta of light, and is massless and electrically neutral. The W^{+} and W^{-} bosons are the antiparticles of one another and have electric charge $Q = 1$ and -1 in units of the elementary charge respectively, whilst the Z boson is electrically neutral. The gluons are massless and electrically neutral but carry colour charge and so exhibit self-interaction. There are eight gluons, corresponding to the octet of linearly independent combinations of the three types of colour charge: red, green, and blue.

Table 2.1: The gauge bosons of the Standard Model.

Name	Symbol	Spin	$Q\ [e]$	I_3	Colour	Mediates
Photon	γ	1	0	0	None	EM interaction
W boson	$W^{\pm}$	1	±1	±1	None	Weak interaction
Z boson	Z	1	0	0	None	Weak interaction
Gluon	g	1	0	0	8 coloured states	Strong interaction

The Higgs boson (H) is the only elementary scalar particle and is postulated to couple to all particles with mass, including itself. It is a manifestation of the Higgs mechanism which serves to give mass to the fermions, and explains why the weak gauge bosons are massive whilst the photon is massless in terms of the spontaneous breaking of the symmetry of the unified electromagnetic and weak interactions [2].

The fermions are composed of the quarks which carry colour charge, and the leptons which do not. There are three charged leptons, the electron (e), muon (μ), and tau lepton (τ), and three electrically neutral leptons, the electron neutrino (ν_e), muon neutrino (ν_μ), and tau neutrino (ν_τ). The electron, muon, and tau lepton have $Q = -1$. Left handed leptons carry weak isospin and are therefore weakly interacting, whereas the right handed leptons do not. The leptons are listed in Table 2.2.

Table 2.2: The leptons of the Standard Model.

Name	Symbol	Spin	Q [e]	I_3^{LH}	I_3^{RH}	Colour
Electron	e	1/2	−1	−1/2	0	None
Electron neutrino	ν_e	1/2	0	1/2	0	None
Muon	μ	1/2	−1	−1/2	0	None
Muon neutrino	ν_μ	1/2	0	1/2	0	None
Tau lepton	τ	1/2	−1	−1/2	0	None
Tau neutrino	ν_τ	1/2	0	1/2	0	None

There are six quarks: three up-type called up (u), charm (c), and top (t), and three down-type called down (d), strange (s), and bottom (b, also known as beauty), as listed in Table 2.3. Each of the quarks carries electric charge and participates in both the electromagnetic and weak interactions. The up-type quarks have $Q = \frac{2}{3}$ and the down-type quarks $Q = -\frac{1}{3}$. Left handed quarks carry weak isospin but right handed quarks do not. The quarks are also strongly interacting, each carrying red, green, or blue colour charge, and are therefore the only elementary particles to participate in all of the fundamental interactions.

Unlike leptons, quarks are not observed as free particles in nature due to a property of the strong interaction called colour confinement [2]. Instead, they bind together in composite colour neutral objects called hadrons. As a pair of quarks are separated, the potential energy increases with separation until it is enough to produce a new quark-antiquark pair from the vacuum [3]. This pair combines with the original pair to form two hadrons, in a process

Table 2.3: The quarks of the Standard Model. Each quark may be of colour red (R), green (G), or blue (B).

Name	Symbol	Spin	$Q\ [e]$	I_3^{LH}	I_3^{RH}	Colour
Up	u	1/2	2/3	1/2	0	R/G/B
Down	d	1/2	−1/3	−1/2	0	R/G/B
Charm	c	1/2	2/3	1/2	0	R/G/B
Strange	s	1/2	−1/3	−1/2	0	R/G/B
Top	t	1/2	2/3	1/2	0	R/G/B
Bottom/Beauty	b	1/2	−1/3	−1/2	0	R/G/B

known as hadronisation. There are two conventional types of hadrons: mesons, which are composed of a quark and antiquark, and baryons, which are composed of three quarks. Recently tetraquarks such as the $\mathrm{Z}(4430)^-$, composed of four quarks, have been observed [4]. In addition, a resonance consistent with being composed of five quarks, a pentaquark, has been observed by the LHCb collaboration [5].

2.1.2 Quantum field theory

The Standard Model is a quantum field theory in which the elementary particles are excitations of quantum fields pervading all of space-time. The dynamics are specified by a Lagrangian density (henceforth Lagrangian) which includes terms describing each of the interactions. The Lagrangian is required to be invariant under a continuous group of local transformations, where a local transformation is one in which different transformations apply at different points in space-time. This symmetry is called gauge invariance [1], and is intrinsic to the theory. The set of gauge transformations under which the Standard Model Lagrangian is invariant is given by the symmetry group

$$\mathrm{SU}(3)_{\mathrm{C}} \times \mathrm{SU}(2)_{\mathrm{L}} \times \mathrm{U}(1)_{Y}, \tag{2.1}$$

where the unitary group $U(n)$ is the group of all $n \times n$ unitary matrices, and the special unitary group $SU(n)$ is the subgroup of $U(n)$ with unit determinant [6].

The group $SU(3)_C$ is associated with the strong interaction and the subscript C denotes that particles with colour charge transform under $SU(3)$. The product group $SU(2)_L \times U(1)_Y$ is associated with the electroweak interaction, a unification of the electromagnetic and weak interactions. The subscript L indicates that only left handed particles carry weak isospin and so transform with respect to $SU(2)$, and the subscript Y refers to weak hypercharge, the charge carried by particles which transform with respect to $U(1)$, defined by the relation $Q = I_3 + \frac{Y}{2}$. The Higgs mechanism spontaneously breaks the electroweak symmetry $SU(2)_L \times U(1)_Y$ to $U(1)_{EM}$, where the subscript EM indicates that $U(1)_{EM}$ is associated with the electromagnetic interaction, acting on electrically charged particles.

2.2 Quantum chromodynamics

Quantum chromodynamics (QCD) is a non-Abelian gauge theory describing the strong interaction, in which the symmetry group $SU(3)_C$ acts on the colour charge carried by quarks and gluons.

2.2.1 The Lagrangian

The QCD Lagrangian is

$$\mathcal{L}_{QCD} = -\frac{1}{4} F^a_{\mu\nu} F^{\mu\nu a} + \sum_q \overline{q}(i\gamma^\mu D_\mu - m_q)q, \tag{2.2}$$

where q are the quark fields for quark flavours $q = \{u, d, s, c, b, t\}$ of mass m_q and γ^μ are the gamma matrices [7]. The gluon field strengths $F^a_{\mu\nu}$ are given by

$$F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu - g_s f_{abc} A^b_\mu A^c_\nu, \tag{2.3}$$

where A_ν^a are the gluon potentials for $a = \{1, \ldots, 8\}$, and g_s determines the strength of the interaction. f_{abc} are the structure constants of SU(3). The indices a, b, and c are over the eight colour degrees of freedom. That the gluons are massless is explained by the impossibility of including a mass term such as $m^2 A^\mu A_\mu$ in Eq. (2.2) whilst maintaining gauge invariance [8].

The covariant derivative of the quark fields is

$$D_\mu q = (\partial_\mu + i g_s A_\nu^a T_a) q, \tag{2.4}$$

where T_a are the linearly independent generators of SU(3), whose fundamental representation is given by

$$T_a = \frac{1}{2}\lambda_a, \tag{2.5}$$

where λ_a are the eight 3×3 Gell-Mann matrices. QCD is a non-Abelian gauge theory, meaning that the generators of SU(3) do not commute: they obey the commutation relation

$$[T_a, T_b] = i f_{abc} T_c, \tag{2.6}$$

which is the origin of the gluon self-interaction term $-g_s f_{abc} A_\mu^b A_\nu^c$ in Eq. (2.3) [8], giving both cubic and quartic gluon vertices.

2.2.2 Asymptotic freedom

The strong coupling constant $\alpha_s = g_s^2/4\pi$ determines the strength of the strong interaction. It is not in fact constant, but a function of the momentum scale Q^2 and is therefore said to be a *running* coupling constant. As the momentum scale increases the coupling decreases logarithmically as

$$\alpha_s(Q^2) = \frac{12\pi}{(33 - 2n_f)\ln(Q^2/\Lambda^2)}, \tag{2.7}$$

where n_f is the number of quark flavours and Λ is a free parameter which must be measured experimentally [3].

The result of the running of the coupling constant is that at small momentum scales, α_s is large and quarks are tightly bound within hadrons, whereas at large momentum scales, α_s is small and the strong interaction between quarks and gluons is weak. This property is called asymptotic freedom and makes QCD calculations tractable at high energies by allowing the interaction to be treated as a perturbation of free fields [3] and represented as a series expansion in powers of α_s.

The origin of asymptotic freedom is in the mathematics of renormalisation, a technique for dealing with divergent integrals in quantum field theory, however this behaviour can be understood using a dielectric analogy [3,8]. Consider two quarks next to one another, each carrying red colour charge. Both quarks constantly emit and absorb virtual gluons, such that they are surrounded by a cloud of colour charge, which is preferentially of the same colour as the quark. As the separation between the quarks is reduced, equivalent to increasing the momentum scale, the colour charge distributions overlap. As the overlap increases, the colour charge of each quark as seen by the other quark reduces and the strength of interaction becomes weaker.

2.3 Diffraction

Diffraction refers to a class of scattering processes between strongly interacting hadrons at high energy where the quantum numbers of the vacuum are exchanged [9]. In the following it is assumed both hadrons are protons, the case applicable at the LHC.

2.3.1 Types of diffraction

The simplest form of diffraction is elastic diffraction. In this process, the protons scatter elastically through a small angle and remain intact throughout the interaction. There is no net flow of colour charge between them, so the central rapidity between them is devoid of additional particles. This particle-free region is called a *rapidity gap*, and is a key characteristic of diffraction [10]. In contrast, in non-diffractive inelastic scattering, colour flows between the protons and they break up in the collision. The rapidity range between the protons contains additional particles from the hadronisation of the proton remnants and gluon radiation, and no gap exists.

Single diffraction proceeds in the same manner as elastic diffraction, however one of the protons dissociates into its elementary constituents which then hadronise. As the exchange is still of the quantum numbers of the vacuum, these particles have net quantum numbers identical to those of the initial state proton, and the central rapidity gap remains [9]. In double diffraction, both protons dissociate. Figure 2.1 illustrates these three types of diffraction, where the mediator of the interaction is the Pomeron which will be described in Section 2.3.2.

Diffraction of two protons p_a and p_b can be written most generally as

$$\mathrm{p}_a + \mathrm{p}_b \rightarrow \mathrm{p}_a^* + \mathrm{p}_b^*, \tag{2.8}$$

where "+" indicates a rapidity gap, and p_a^* and p_b^* are a particle or ensemble of particles with the same net quantum numbers as p_a and p_b respectively.

The momentum exchanged between the scattered protons in diffraction with large rapidity gaps is typically low; the processes are said to be *soft*. The strong coupling constant is large and this restricts the application of perturbation theory to QCD, since terms of higher order in α_s cannot be neglected. An alternative formalism to QCD is therefore required for diffraction in the soft regime.

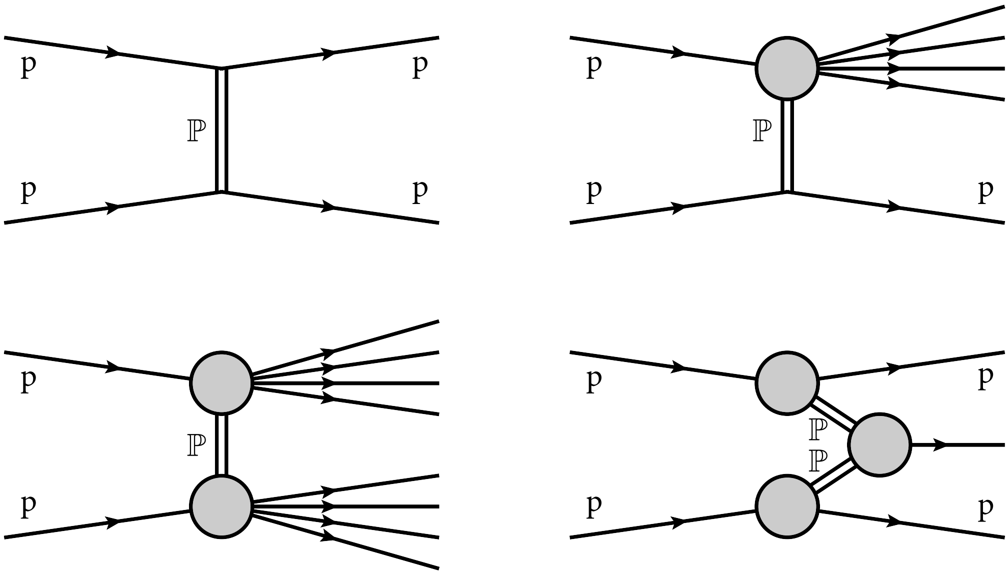

Figure 2.1: Different diffractive processes in the case of two colliding protons, in the Regge theory description where the Pomeron, indicated by the double lines labelled $\mathbb{P}$, is exchanged. Shown are elastic diffraction (top left), single diffraction (top right), double diffraction (bottom left), and double Pomeron exchange (bottom right).

2.3.2 Regge theory and the Pomeron

The framework provided by Regge theory is an alternative to QCD in the regime where perturbation theory cannot be applied. It regards angular momentum as a complex variable [11, 12] and originates from the properties of the scattering matrix, the unitary matrix which relates the initial free particle states at time $-\infty$ to the final free particle states at time $+\infty$ in a scattering interaction at time zero [13]. Regge theory was pursued as a theory of the strong interaction before QCD became established [14].

Regge theory describes the strong interaction in terms of the exchange of an object which has angular momentum $\alpha(t)$ that is a function of t, the square of the momentum transfer. This object is called the *Reggeon* [13]. Due to this functional dependence of the angular momentum, the Reggeon is not identified as a single physical particle. Instead, it is associated with the superposition of physical particles which lie on the same *Regge trajectory* [13]

$$\alpha(t) = \alpha_0 + \alpha' t. \tag{2.9}$$

The slope and intercept were measured to be $\alpha' = 0.86\,\mathrm{GeV}^{-2}$ and $\alpha_0 = 0.55$ by fitting the spin against squared mass of light mesons [14].

The theory states that the total scattering cross section σ_{tot} scales with squared centre-of-mass energy s as

$$\sigma_{\mathrm{tot}} \propto s^{\alpha_0 - 1}. \tag{2.10}$$

This is the Pomeranchuk theorem [13, 14]. Therefore with a measured value for α_0 of 0.55, it was predicted that the scattering cross section would vanish asymptotically at high energy. However, this was found to be inconsistent with experimental measurements of proton-proton and proton-antiproton scattering, which show the cross section to increase with energy [14].

To reconcile the Pomeranchuk theorem with the observed energy dependence of the total scattering cross section, a new Regge trajectory was introduced which necessarily has $\alpha_0 > 1$ and carries the quantum numbers of the vacuum. This trajectory is called the *Pomeron*, and is the object described by Regge theory as being exchanged in diffractive processes [13, 14]. Unlike the Reggeon, no known particles have been identified to lie on the Pomeron trajectory, although the hypothetical glueball would have the correct quantum numbers. In QCD, the simplest exchange of the quantum numbers of the vacuum is of a pair of gluons in a colour singlet state, and this QCD representation of the Pomeron is used in Section 2.4 which follows.

2.4 Central exclusive production

Central exclusive production is a special case of diffraction in which a single object is created at central rapidity—that is with pseudorapidity much closer to zero than the intact protons—and separated from the protons by a rapidity gap on either side. It can be written as

$$\mathrm{p} + \mathrm{p} \to \mathrm{p} + X + \mathrm{p}, \tag{2.11}$$

where X is the centrally exclusively produced state and "+" indicates a rapidity gap [15]. The protons are only slightly perturbed by the interaction.

There are three similar mechanisms which constitute CEP: diphoton exchange, photoproduction, and double Pomeron exchange, which are illustrated in Fig. 2.2. In diphoton exchange, each proton radiates a photon and the central system is produced by the fusion of the two photons. Diphoton exchange results in non-resonant production of pairs of muons. Photoproduction involves the emission of a photon from one proton and a Pomeron (which again is considered as a colour singlet gluon pair) from the other proton and can produce charmonium resonances such as the J/ψ and $\psi(2S)$. Double Pomeron exchange (DPE) is the process where each proton emits a Pomeron, which fuse to create the central system. Pomeron mediated processes dominate over photon mediated processes due to the Pomeron being strongly interacting [15], and DPE is responsible for the production of $\chi_c(1P)$ mesons as studied in this thesis.

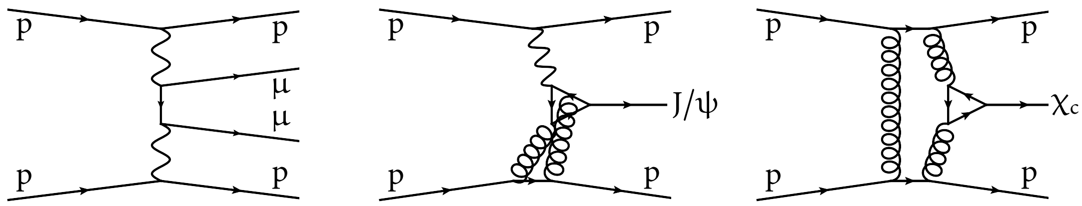

Figure 2.2: The three CEP mechanisms with representative final states: diphoton exchange producing a pair of muons (left), photoproduction of J/ψ (centre), and double Pomeron exchange producing $\chi_c(1P)$ (right). The Pomeron is represented in the QCD scenario where it is considered as two gluons in a colour singlet state.

CEP has received attention from the theoretical community due to the potential it offers at the LHC for searching for and studying the quantum numbers of objects such as the Higgs boson [16–19], and hypothetical particles such as glueballs [20–23] and particles predicted by beyond-the-Standard-Model (BSM) supersymmetric theories [24–27]. Measurement of the CEP of Standard Model mesons such as the J/ψ, ψ(2S), and $\chi_c(1P)$ can be used to benchmark these BSM predictions, and check and correct theoretical models of CEP [28,29]. These charmonia states are theoretically interesting as their CEP also allows several aspects of QCD to be tested, such as the predictions of effective field theories, and can be used to investigate discrepancies between next-to-leading order predictions and experimental measurements [30].

2.4.1 Spin selection rules

The fact that the colliding protons do not dissociate in CEP leads to several important properties which also make it an interesting process to study experimentally, and which are not replicated in inclusive production processes.

There is no angular momentum transfer between the protons, so in the limit that they scatter at zero angle the projection of the total angular momentum of the centrally produced object onto the beam direction J_z is zero by conservation of angular momentum [16]. In the case that the protons scatter at a non-zero but small angle, angular momentum is transferred and so other values of J_z are possible, but they are greatly suppressed and the J_z rule still holds approximately. Additionally, as the Pomeron carries the quantum numbers of the vacuum, the centrally produced object must be a colour singlet which is even under charge conjugation ($C = +1$) and parity transformation ($P = +1$). Together these conditions constitute the spin selection rules $J_z^{PC} = 0^{++}$ [27] for Pomeron-Pomeron mediated CEP, and so simply observing a given object to be produced in CEP is sufficient to determine several of its quantum numbers.

From an experimental perspective, the fact that the protons remain intact and no additional particles are produced except the single centrally produced system provides a clear low multiplicity signature to search for in the detector. Additionally, the mass of the centrally produced object can be measured with high precision if the momenta of the outgoing protons can be measured. The protons do not enter the acceptance of a central detector such as LHCb, but this measurement can be made using dedicated forward detectors, typically housed within *Roman pots*, where available [31].

2.4.2 The Durham model

Quantitative predictions of CEP which can be tested at high energy colliders have been made using the Durham model of DPE [32, 33, and references therein]. This model describes DPE as a perturbative two gluon exchange with non-perturbative corrections for soft rescattering, as shown in Fig. 2.3, and can be used to calculate the CEP cross section for any even C-parity particle X which has mass M_X large enough to allow the amplitude at the production vertex to be calculated perturbatively.

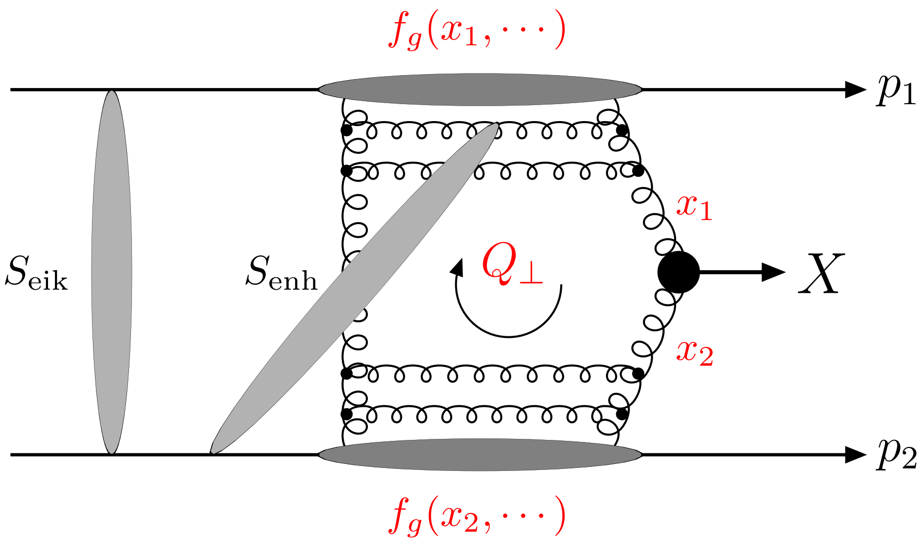

Figure 2.3: A diagram showing the Durham model for the production of an object X via double Pomeron exchange [32]. $f_g(x_i, \ldots)$ are the skewed unintegrated gluon distribution functions, and S_{eik} and S_{enh} are graphical representations of the rescattering effects accounted for by the eikonal and enhanced Reggeon survival factors respectively.

Following Ref. [33], the perturbative amplitude is given by

$$T(\mathbf{p}_{1\perp}, \mathbf{p}_{2\perp}) = \pi^2 \int \frac{d^2\mathbf{Q}_\perp \mathcal{M}}{\mathbf{Q}_\perp^2 (\mathbf{Q}_\perp - \mathbf{p}_{1\perp})^2 (\mathbf{Q}_\perp + \mathbf{p}_{2\perp})^2} \times f_g(x_1, x_1', Q_1^2, \mu^2) f_g(x_2, x_2', Q_2^2, \mu^2) F_p(t_1) F_p(t_2), \tag{2.12}$$

where $Q_\perp$ is the transverse momentum in the gluon loop (indicated in Fig. 2.3), $t_i \approx -\mathbf{p}_{i\perp}^2$ are the proton transverse momentum transfers, f_g are the skewed unintegrated gluon densities of the protons for hard scale $\mu \sim M_X$, and $F_p(t_i)$ are the elastic proton form factors. The proton transverse momenta must be small to ensure the protons remain intact. The colour averaged, normalised sub-amplitude for the $gg \to X$ process is

$$\mathcal{M} \equiv \frac{2}{M_X^2} \frac{1}{N_C^2 - 1} \sum_{a,b} \delta^{a,b} q_{1\perp}^{\mu} q_{2\perp}^{\mu} V_{\mu\nu}^{ab}, \tag{2.13}$$

where a and b are indices over colour, $q_{i\perp}$ are the gluon transverse momenta, and $V_{\mu\nu}^{ab}$ is the $gg \to X$ vertex.

The skewed unintegrated gluon densities give the transverse momentum distribution for the gluons, without additional gluon radiation. They are a complicated function of both the conventional parton distribution functions, and a Sudakov factor. The Sudakov factor is the probability that the gluons in the hard process do not radiate, which would fill the otherwise empty rapidity gap [34]. It is found that this factor causes the amplitude to be insensitive to the low $Q_\perp$ part of the integral around the gluon loop, which is used to justify the perturbative calculation for the $gg \to X$ vertex [32].

The perturbative amplitude of Eq. (2.12) must then be corrected for non-perturbative effects, collectively called soft survival effects because they prevent the rapidity gaps from "surviving" the interaction. These corrections appear as two factors, the eikonal survival factor (labelled S_{eik} in Fig. 2.3) and the enhanced survival factor (labelled S_{enh} in Fig. 2.3), which account for the two types of rescattering shown in Fig. 2.4.

The eikonal survival factor is the probability that additional soft exchanges between the protons do not create additional particles in the final state. S_{eik} is not a simple multiplicative factor, but instead a distribution in the proton impact parameter with $p_\perp$ averaged values of $\langle S_{\text{eik}}^2 \rangle \sim 0.03$ for $\chi_{c0}(1P)$ production and $\langle S_{\text{eik}}^2 \rangle \sim 0.08$ for $\chi_{c2}(1P)$ production at LHC energies [29]. It is not calculable using perturbative QCD and is instead extracted from experimental measurements of soft hadronic scattering [35, 36]. The eikonal factor is an important term which cannot be ignored, as these $p_\perp$ averaged values show neglecting this correction inflates the predicted CEP cross section by two orders of magnitude.

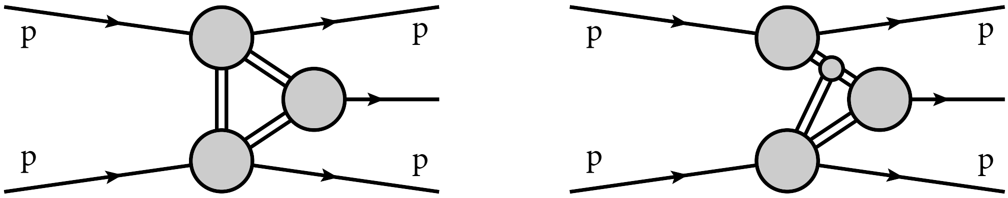

Figure 2.4: Diagrams of the two kinds of soft rescattering accounted for in the Durham model: rescattering between the protons (left), accounted for by the eikonal survival factor $S_{\rm eik}$, and rescattering between a proton and a parton within the loop (right), accounted for by the enhanced survival factor $S_{\rm enh}$.

The enhanced survival factor also accounts for rescattering, but now between a proton and a parton within the loop rather than between the two protons. It is a much smaller correction than the eikonal survival factor [32]. Its value is only approximately known, but it depends predominantly on the parton transverse momenta and only slightly on the proton transverse momenta [37]. A value of $\langle S^2_{\rm enh}\rangle \sim 0.3$ is used in the calculation for $\chi_c(1P)$ production at the LHC [29].

Bringing together the perturbative amplitude and the soft survival factors, the differential cross section for an object X to be central exclusively produced by DPE at rapidity y_X is

$$\frac{d\sigma}{dy_X} = \langle S^2_{\rm enh}\rangle \int d^2\mathbf{p}_{1\perp} d^2\mathbf{p}_{2\perp} \frac{|T(\mathbf{p}_{1\perp}, \mathbf{p}_{2\perp})|^2}{16^2\pi^5} S^2_{\rm eik}(\mathbf{p}_{1\perp}, \mathbf{p}_{2\perp}), \tag{2.14}$$

where $\langle\ldots\rangle$ indicates averaging over $Q_\perp$ [33].

2.4.3 Predictions and prior measurements

Central exclusive $\chi_c(1P)$ production is interesting to study as Eq. (2.14) predicts a cross section of order nanobarns [28], and the cross sections for the three spin states, $\chi_{c0}(1P)$, $\chi_{c1}(1P)$, and $\chi_{c2}(1P)$, are expected to follow a specific hierarchy. Production of $\chi_{c1}(1P)$ and $\chi_{c2}(1P)$ is expected to be suppressed compared to $\chi_{c0}(1P)$, although for different reasons [16]. Production of $\chi_{c1}(1P)$ is suppressed due to the Landau-Yang theorem, which states that the massive spin-1 $\chi_{c1}(1P)$ cannot couple to two massless spin-1 gluons which are on shell. Spin-2 $\chi_{c2}(1P)$ production is suppressed due to the $J_z^{PC} = 0^{++}$ selection rule, but is permitted because the protons scatter at a small but non-zero angle.

The SuperCHIC event generator [28,29,38] calculates the CEP cross section according to the Durham model for a range of final states using the Monte Carlo method to numerically evaluate the integral in Eq. (2.14). Measurements of central exclusive $\chi_c(1P)$ production have been made by the CDF and LHCb collaborations and the measured cross sections compared with the predictions of the Durham model as implemented in SuperCHIC [39,40]. The CDF mass resolution did not allow the ratio of the production of the $\chi_c(1P)$ spin states to be determined; all production was assumed to be of $\chi_{c0}(1P)$ but the measurement of $(76 \pm 10 \pm 10)$ nb is nevertheless consistent with the prediction of 90 nb for a centre-of-mass energy of 1.96 TeV.

An LHCb analysis measured preliminary values of (9.3 ± 4.5) pb, (16.4 ± 7.1) pb, and (28.0 ± 12.3) pb for production of $\chi_{c0}(1P)$, $\chi_{c1}(1P)$, and $\chi_{c2}(1P)$ respectively for the final state $\mu^+\mu^-\gamma$ in the LHCb acceptance, using data collected in 2010 at a centre-of-mass energy of 7 TeV, for which SuperCHIC predicted cross sections of 14 pb, 10 pb, and 3 pb [40]. This measurement was repeated in a preliminary, unpublished analysis using a larger data set collected in 2011 at the same centre-of-mass energy, which measured values of (2.2 ± 3.0) pb, $(4.3^{+7.6}_{-4.4})$ pb, and $(25.0^{+7.1}_{-9.2})$ pb for $\chi_{c0}(1P)$, $\chi_{c1}(1P)$, and $\chi_{c2}(1P)$, within the same acceptance as the 2010 analysis [41].

Therefore there appears to be an enhancement of $\chi_{c2}(1P)$ production compared to the prediction. This could be due to greater inelastic background contributions to the $\chi_{c2}(1P)$ data sample (the $J_z^{PC} = 0^{++}$ selection rule suppressing $\chi_{c2}(1P)$ production does not apply when the protons dissociate), as the sample purity was assumed to be the same for each spin state in the 2010 analysis, and had very large uncertainties in the 2011 analysis, although it was then determined separately for each spin state [32]. Alternatively, this enhancement compared to the prediction may be due to a problem in the theoretical model. It has been suggested the enhancement could be due to additional non-perturbative effects because the $\chi_c(1P)$ is insufficiently massive to fully justify the perturbative calculations used in the Durham model [42].

In order to contribute to resolving this discrepancy between the experimental measurement of the ratio of $\chi_{c0}(1P)$ to $\chi_{c2}(1P)$ production and the theoretical prediction of the Durham model, this thesis presents a measurement of the cross section for central exclusive $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$ production ($p + p \to p + \chi_c(1P) + p$), with the $\chi_c(1P)$ reconstructed in the K^+K^- and $\pi^+\pi^-$ final states using the LHCb detector. The $\chi_{c1}(1P)$ cannot decay into K^+K^- or $\pi^+\pi^-$ and so is inaccessible in this study. These hadronic modes offer improved mass resolution compared to the previous LHCb measurement using the $\mu^+\mu^-\gamma$ mode, allowing the $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$ mass peaks to be distinguished. The SuperCHIC predicted production cross sections at $\sqrt{s} = 8$ TeV are 298 nb for $\chi_{c0}(1P)$ and 4 nb for $\chi_{c2}(1P)$.

3

The LHCb experiment

This chapter describes the LHCb experiment, a dedicated heavy flavour physics experiment located at CERN's Large Hadron Collider, the world's most powerful particle accelerator. The LHCb detector is described as a whole and the design and performance of each sub-detector is discussed. The data processing chain is explained, including the hardware and software triggers, offline event reconstruction, and stripping. Finally, the LHCb simulation applications used to study selection efficiencies and the detector performance are described.

3.1 CERN

CERN, the European Organization for Nuclear Research, is the world's largest physics laboratory, formed in 1954 to foster cooperation between European scientists after the Second World War. The laboratory is situated on the Franco-Swiss border near Geneva and since its founding has been the crucible for several major discoveries.

In 1965 at the Proton Synchrotron the first observation of antinuclei was made, simultaneously with a team at Brookhaven National Laboratory [43, 44]. Electroweak neutral currents were discovered in 1973 by the Gargamelle experiment [45], and the W and Z bosons by the UA1 and UA2 experiments in 1983 [46–49]. In July 2012 the ATLAS and CMS experiments announced the discovery of a new particle consistent with the Standard Model Higgs boson [50, 51], the final particle in the Standard Model to be observed.

3.2 The Large Hadron Collider

The Large Hadron Collider (LHC) [52–54] is the largest and most energetic particle accelerator in the world, with a design centre-of-mass energy of 14 TeV and luminosity of $1 \times 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$. It is a synchrotron located at CERN in the 26.7 km long underground tunnel originally bored for its predecessor, the Large Electron-Positron Collider (LEP), 100 m below the surface.

The power emitted in synchrotron radiation by a charged particle in radial acceleration scales as one over the fourth power of mass. Therefore the LHC accelerates protons, which are 1800 times more massive than the electrons and positrons accelerated by LEP, to achieve the 14 TeV design energy in a ring of the same length.

High density beams are required to reach the design luminosity, which precludes accelerating antiprotons as they are difficult to accumulate. Instead, the LHC accelerates two beams of protons. This necessitates two separate beam pipes, each containing an ultra-high vacuum of 10^{-8} Pa, in which the beams orbit. As viewed from above, the beam that orbits in the clockwise direction is called Beam 1, and the beam that orbits in the counter-clockwise direction Beam 2. The beams each contain up to 2808 bunches, with 1.15×10^{11} protons per bunch, and are brought into collision at four points around the synchrotron.

Each beam is accelerated by eight superconducting radiofrequency cavities operating at 400 MHz and cooled to 4.5 K with superfluid helium. Steering 7 TeV protons around the LHC requires superconducting magnets in order to provide the requisite field strength. The beams are guided by 1232 copper-clad niobium-titanium dipole magnets. At an operating temperature of 3.1 K, again maintained with superfluid helium, the dipoles provide a maximum field of 8.33 T with a current of 11,850 A. Higher order multipole magnets are used for focusing and chromaticity correction.

Protons are obtained for acceleration by stripping the orbital electrons from hydrogen atoms. The protons are injected into Linac 2, a linear accelerator, and accelerated to 50 MeV before injection into a chain of synchrotrons. The LHC injection chain is shown in Fig. 3.1.

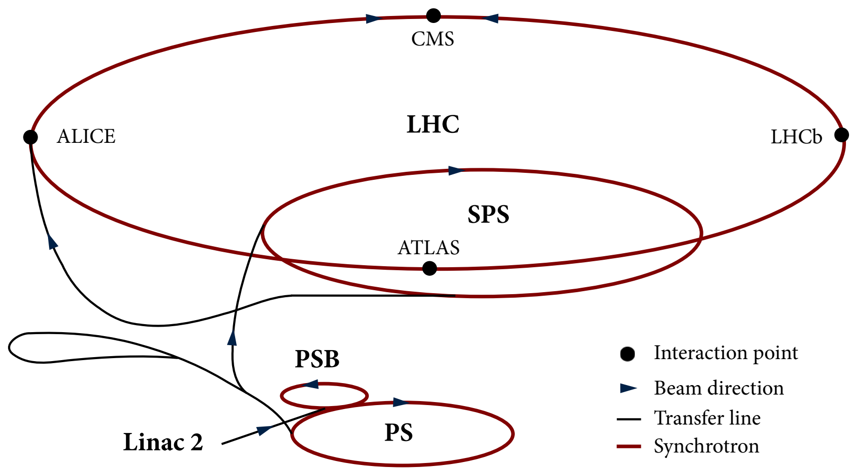

Figure 3.1: A schematic of the LHC injection chain. The beams originate in Linac 2 and follow the path indicated by the blue arrows through the PSB, PS, and SPS to the LHC [55].

The first synchrotron in the chain is the Proton Synchrotron Booster (PSB), a 25 m circumference machine which accelerates the beam to 1.4 GeV. From the PSB the beam is injected into the Proton Synchrotron (PS), a 628 m circumference machine, where it is accelerated to 25 GeV. The beam is then injected into the Super Proton Synchrotron (SPS). This is a 6.9 km circumference ring which accelerates the beam to 450 GeV before it is injected into the LHC.

After the required number of bunches in each beam have been filled, the beams are accelerated from the 450 GeV injection energy to collision energy. Once at collision energy, they are collimated and brought into collision at four interaction points. The LHC experiments begin recording data once the beams are colliding stably. Over time the beam current decays due to depletion from collisions and other losses. After the current has decayed to a level where the luminosity is too low for efficient data taking—typically after around 10 hours—the beams are dumped and the injection cycle is repeated.

There are seven experiments at the LHC, distributed at the four interaction points. Two general purpose detectors (GPDs) are located at opposite sides of the ring: ATLAS (A Toroidal LHC Apparatus) [56] and CMS (Compact Muon Solenoid) [57]. The primary focus of their physics programmes is the search for and study of the Standard Model Higgs boson, in addition to carrying out precision measurements of the Standard Model and searching for beyond the Standard Model physics. The GPDs are hermetic detectors, meaning that the full 4π sr solid angle is instrumented except for the gaps required for the beam pipe to enter the detector volume.

ALICE (A Large Ion Collider Experiment) [58] is a dedicated heavy ion experiment, and is a barrel shaped detector designed to study the quark-gluon plasma during special running periods when the LHC collides lead ions.

LHCb (Large Hadron Collider beauty) [59] is a dedicated experiment for studying hadrons containing b and c quarks. The data analysed in this thesis were recorded with the LHCb detector, and as such it is described in detail in Section 3.3.

TOTEM (Total Elastic and Diffractive Cross Section Measurement) [60], LHCf (Large Hadron Collider forward) [61], and MoEDAL (Monopole and Exotics Detector at the LHC) [62] are smaller specialised experiments which share the interaction points of CMS, ATLAS, and LHCb respectively.

Protons were steered around the entire LHC for the first time in September 2008. The original commissioning plan called for first collisions by the end of the month, with the centre-of-mass energy being increased to 10 TeV by the end of 2008. However, on 19 September an electrical fault caused a superconducting magnet to quench, resulting in serious mechanical damage and necessitating extensive repair.

This delayed operation by 14 months, with beams first circulating again in November 2009. Following the incident it was decided to operate at lower than design energy for the first period of operation, which ended in 2013.

Collisions were achieved again on 23 November 2009 at a centre-of-mass energy of 900 GeV, and on 30 March 2010 collisions were achieved at 7 TeV [63]. The LHC provided 7 TeV collisions during 2010 and 2011 at a bunch crossing rate of 20 MHz, half the design rate. For 2012 the centre-of-mass energy was increased to 8 TeV at the same bunch crossing rate, and the data analysed in this thesis were recorded during 2012. The period of running between 2010 and 2012 is collectively called Run 1.

3.3 The LHCb detector

The LHCb detector is a single arm spectrometer designed for heavy flavour physics. It is optimised for the study of b and c-hadrons in order to make precise measurements of CP violation and search for indirect evidence of New Physics in their rare decays.

In order to carry out this physics programme, the detector must provide excellent momentum and vertex resolution, and identification of electrons, muons, photons, and charged and neutral hadrons (particularly pions and kaons). This must be accompanied by a trigger capable of efficiently selecting a large variety of final states of interest among high multiplicity hadronic backgrounds. These attributes also make the LHCb detector well suited to the study of diffractive processes such as CEP.

The LHCb detector has a forward geometry unique among the LHC experiments. In proton-proton collisions at the LHC, b quarks are predominantly produced in boosted pairs close to the beam axis, with both of the pair moving in the same direction as shown in Fig. 3.2. The conical geometry of LHCb, with an acceptance of 10 mrad to 300 mrad in the bending plane and 10 mrad to 250 mrad in the non-bending plane, maximises the b-hadron acceptance whilst minimising costs. The angular acceptance corresponds to a pseudorapidity coverage of approximately $2 < \eta < 5$, where

$$\eta = -\ln\left(\tan\left(\frac{\theta}{2}\right)\right), \tag{3.1}$$

and θ is the angle with respect to the beam axis. The detector has a forward arm, but no backward arm, in order to fit into the cavern built for the DELPHI experiment in the LEP era.

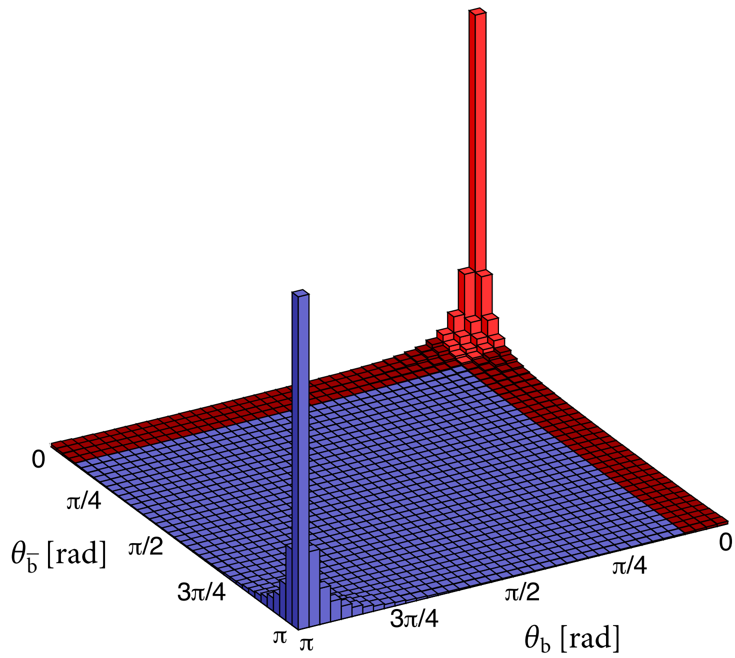

Figure 3.2: The angular distribution of b quark pairs in 14 TeV proton-proton collisions simulated with PYTHIA [64]. The purple region is where both quarks are outside the LHCb acceptance, the dark red region where a single quark is within the acceptance, and the light red region where both quarks are within the acceptance [65].

LHCb was designed to operate at a nominal luminosity of $2 \times 10^{32}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1}$, lower than the LHC's design luminosity of $1 \times 10^{34}\ \mathrm{cm}^{-2}\ \mathrm{s}^{-1}$. This means the majority of bunch crossings result in only a single proton-proton interaction, which simplifies data analysis, minimises radiation damage, and reduces the detector occupancy which would otherwise degrade performance. This is also beneficial for the study of CEP, as tight multiplicity cuts can be applied to bunch crossings containing only a single interaction, in order to select candidate CEP events with only a handful of tracks in the detector. The luminosity is reduced by defocusing the beams so that they only partially overlap at the interaction point. As the beam current is depleted by collisions and other losses over the course of a fill, the luminosity at LHCb is maintained by bringing the beams into greater overlap. This process is called luminosity levelling, and enables LHCb to operate at the chosen luminosity uniformly, in contrast to ATLAS and CMS which experience decreasing instantaneous luminosity over the course of a fill (ALICE also employs luminosity levelling). Due to better than expected detector and trigger performance, the nominal luminosity was increased to $3.5 \times 10^{32}\,\mathrm{cm^{-2}\,s^{-1}}$ in 2011 and again to $4 \times 10^{32}\,\mathrm{cm^{-2}\,s^{-1}}$ in 2012.

The locations of the components of the LHCb detector are specified in a Cartesian coordinate system with its origin at the nominal interaction point. The z-axis points into the spectrometer along the beam direction. The x-axis is horizontal, pointing out of the LHC, and the y-axis points vertically upwards to give a right handed coordinate system.

A cross-section of the detector along the $x = 0$ plane is shown in Fig. 3.3. The Vertex Locator (VELO), a silicon microstrip detector, surrounds the interaction point. This is followed by the first of two Ring-Imaging Cherenkov (RICH) detectors. After the first RICH detector is the Tracker Turicensis (TT), a planar silicon microstrip detector. A large dipole magnet follows, which provides the field necessary for the tracking detectors to measure charged particle momenta. After the magnet are three tracking stations, T1-T3, each containing an Inner Tracker (IT) component and an Outer Tracker (OT) component. Then follows the second RICH detector, after which is the first of five muon detectors. The calorimetry system then follows: a Scintillating Pad Detector (SPD), Preshower detector (PS), Electromagnetic Calorimeter (ECAL), and Hadronic Calorimeter (HCAL). Four additional muon detectors, M2-M5, complete the detector. Each subdetector is now described individually in Sections 3.3.1 to 3.3.9, following Ref. [59].

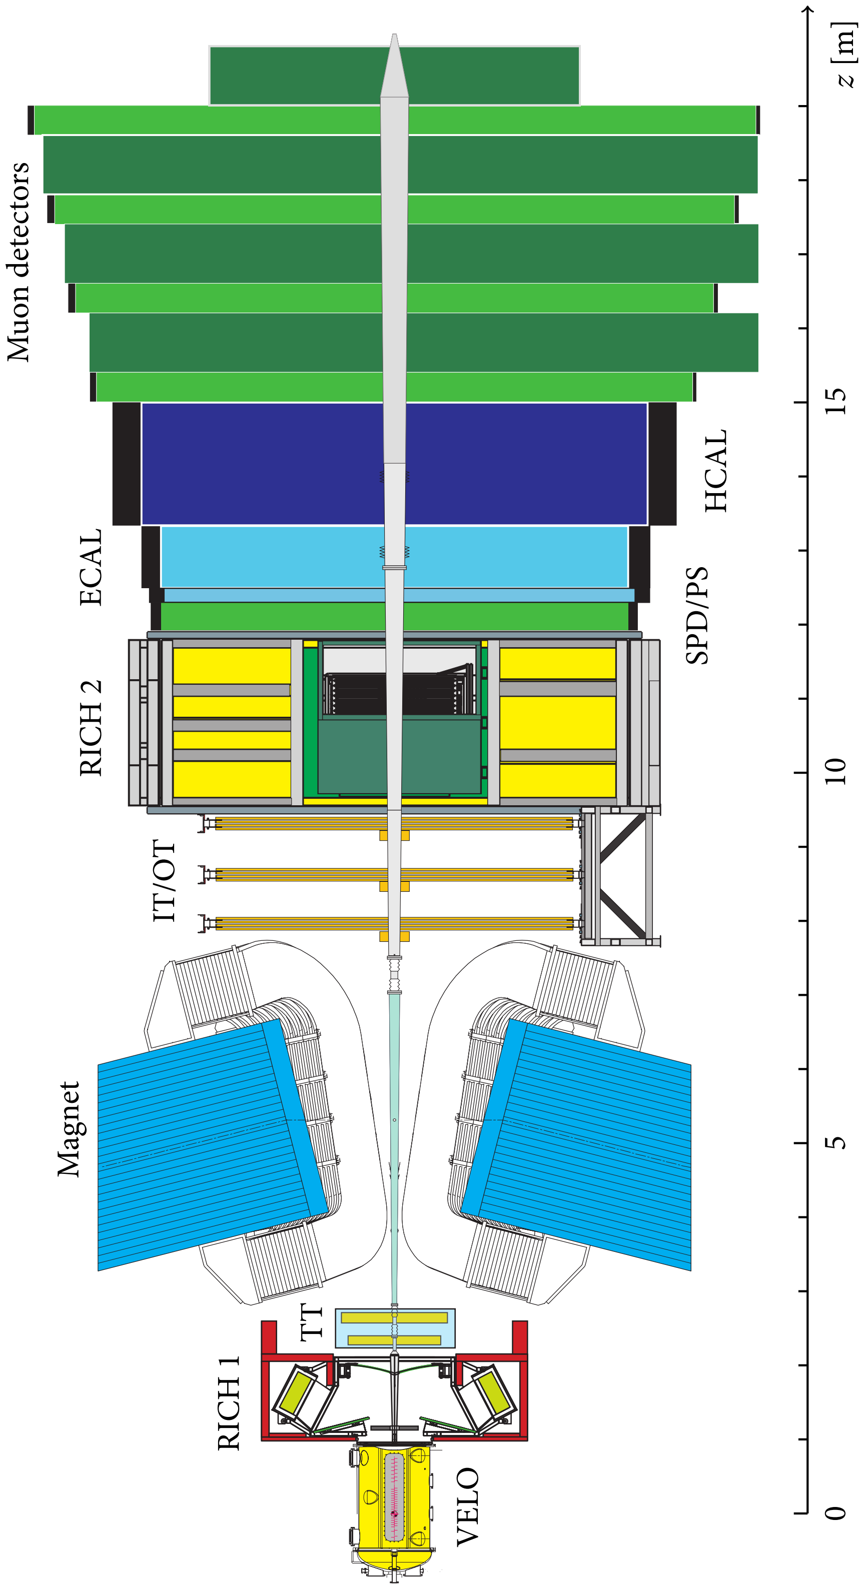

Figure 3.3: A cross-sectional schematic of the LHCb detector along the $x = 0$ plane showing each of the subdetectors described in the text.

3.3.1 Vertex Locator

The VELO [66] is a silicon microstrip detector which measures the coordinates of tracks originating from the interaction region so that production and decay vertices can be reconstructed.

It comprises 21 stations distributed along the z-axis over a distance of approximately 1 m, as shown in Fig. 3.4. Each station has two circular silicon sensors 84 mm in diameter, one of which measures the radial coordinate (r-sensor) and the other the azimuthal coordinate (ϕ-sensor). The longitudinal coordinate is given by the sensor's z position. The number of stations was determined by the requirement that a track with $1.6 < \eta < 4.9$ originating from a vertex with $|z| < 106$ mm must pass through at least three sensors.

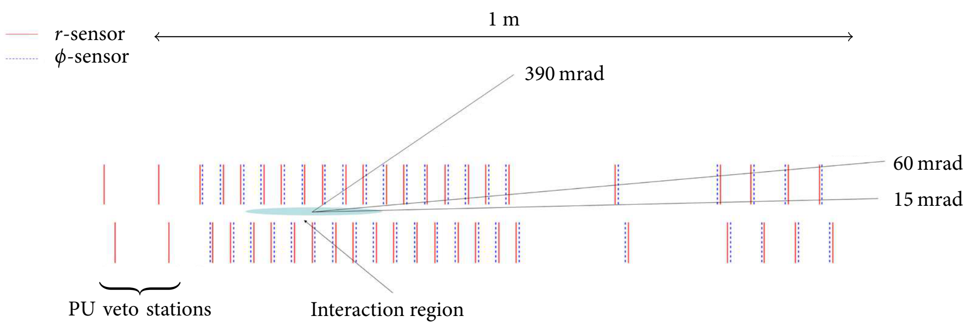

Figure 3.4: A schematic of the VELO and pile-up veto system stations showing their distribution around the interaction region, and the angular coverage of the VELO [59].

The VELO sensors are located in a secondary vacuum inside the beam pipe, separated from the beam vacuum by a 0.3 mm thick corrugated aluminium sheet, shown in Fig. 3.5. This is known as the *RF foil*, as it serves to protect the sensors from the RF effects of the beams. The corrugations allow the two halves of each VELO sensor, which are split vertically with the halves offset longitudinally by 15 mm, to overlap so that there are no acceptance gaps and to aid alignment.

Each sensor has a circular central cutout for the beams to pass through of radius 7 mm, minimised to reduce the extrapolation distance when calculating impact parameters. This

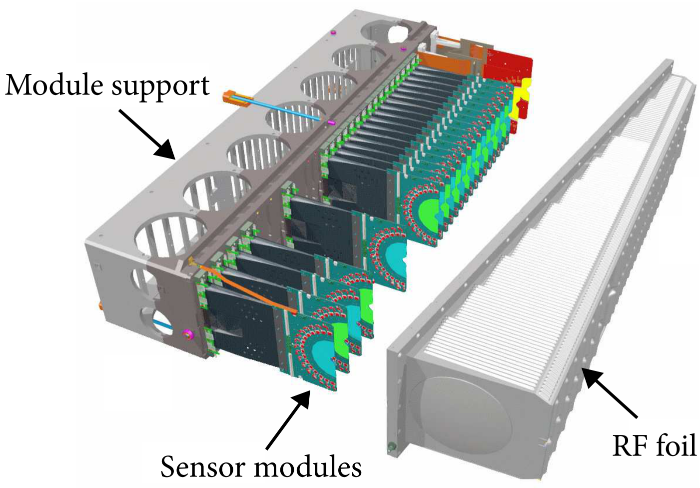

Figure 3.5: One half of the VELO showing the sensor modules on their support structure and the RF foil (shown detached). The red sensors at the far end are part of the pile-up veto system [59].

is smaller than the LHC acceptance during injection, therefore each half of the VELO is retracted 30 mm horizontally for injection and only moved back once the beams are stable.

The microstrips on the r-sensors are arranged in concentric semi-circles, with a pitch of 38 µm at the inner radius increasing linearly to 101.6 µm at the outer radius, whereas the microstrips on the ϕ-sensors are arranged radially as shown in Fig. 3.6.

The ϕ-sensors are split into two concentric regions to prevent having too large a pitch and high an occupancy at the outer radius. The boundary between the regions is at a radius of 17.25 mm, where the pitch of the inner strips is 78.3 µm and that of the outer strips is 39.3 µm. The strips are angled from the radial by 20° in the inner region and 10° in the opposite direction in the outer region, which improves pattern recognition. The angle of skew is reversed in adjacent sensors to improve the rejection of ghost hits.

The VELO meets the design requirement of a spatial hit resolution of 4 µm for the optimal track angle. The measured signal-to-noise ratio in Run 1 was 20, exceeding the design requirement of 14 [67].

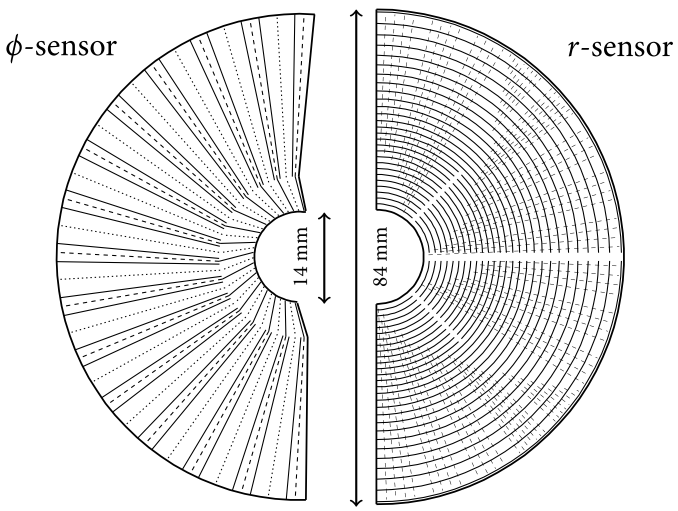

Figure 3.6: A schematic of the microstrip layout on a ϕ-sensor and an r-sensor. Solid lines denote strip boundaries and dashed lines routing lines [66].

Upstream of the primary VELO stations are two further stations containing only r-sensors. These comprise the pile-up veto (PU) system, used to measure the backward track multiplicity and estimate the number of collisions per bunch crossing. The PU system is the subdetector furthest upstream from the interaction point, and is therefore important for selecting CEP candidates where it is used to reject events with backward tracks (tracks with negative pseudorapidity).

3.3.2 Dipole magnet

The dipole magnet [68] provides a field to bend the trajectory of charged particles in the horizontal plane. This enables the momentum of such particles to be measured from their radius of curvature.

The magnet is shown in Fig. 3.7, and comprises two window-frame yokes, one above the beam pipe and one below, which are constructed from 100 mm thick plates of laminated low carbon steel. The inside face of each pole slopes at the angle of the detector acceptance. The magnet is aligned within the detector coordinate system with a precision of ±2 mm.

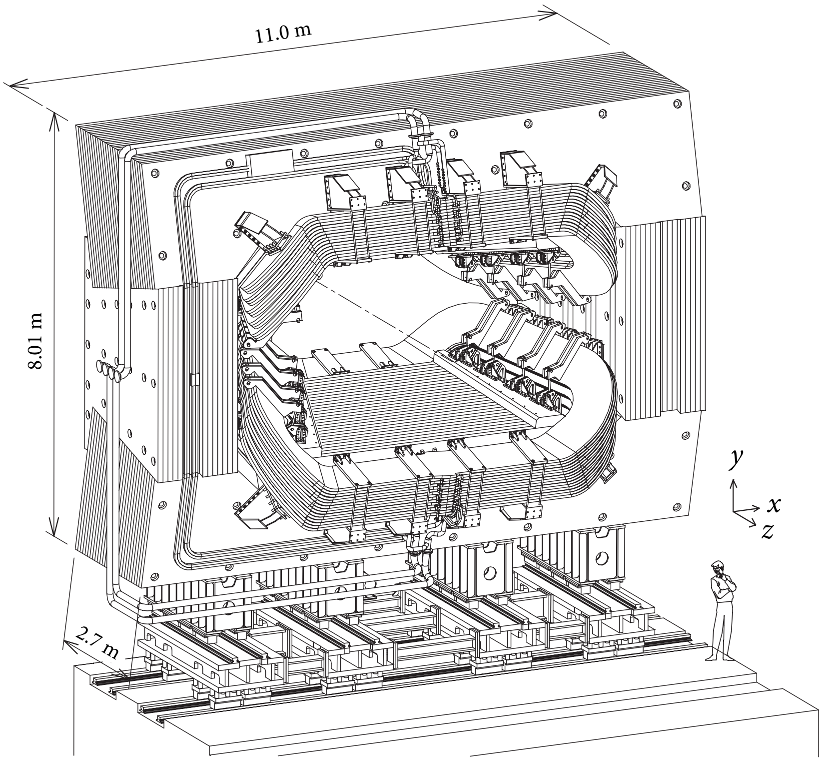

Figure 3.7: The dipole magnet on its mounting with a human figure for scale. [59].

The main design requirements of the magnet were to maximise the field strength between the VELO and TT whilst ensuring it did not exceed 2 mT at the RICH detectors. The field as a function of distance from the interaction point is plotted in Fig. 3.8. With a nominal coil current of 5.85 kA the integrated field is 4 T m for a track length of 10 m. A normal conducting design was chosen to reduce costs and technology risk.

The polarity of the magnet is reversed several times per year during operation to enable potential sources of measurement bias associated with asymmetric detector efficiencies to be studied and controlled.

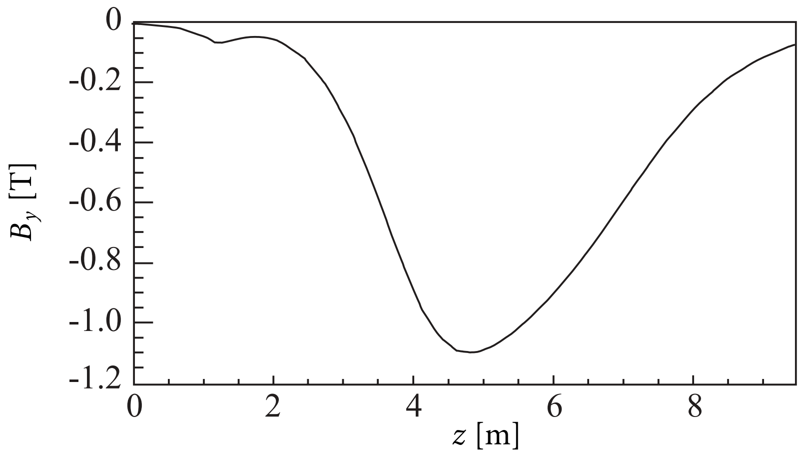

Figure 3.8: The y-component of the magnetic field along the z-axis of the detector [69].

Achieving the required momentum resolution requires knowing the field of the magnet with a precision of 1×10^{-4} and the position of the peak field to a precision of a few millimetres. The field was mapped on an 8 cm × 8 cm × 10 cm grid with an array of Hall probes, aligned with the coordinate system of the experiment with a precision of 1 mm.

3.3.3 Tracker Turicensis

The TT is a silicon microstrip detector located between the first RICH detector and the dipole magnet. It operates in concert with the magnet, VELO, IT, and OT to measure charged particle momenta.

The TT is a planar detector 150 cm wide and 130 cm tall, covering the full detector acceptance with an active area of 8.4 m^2. It is composed of four layers of modules, the third of which is shown in Fig. 3.9, with the first and second layers separated from the third and fourth by 270 mm in z to aid the track reconstruction algorithm.

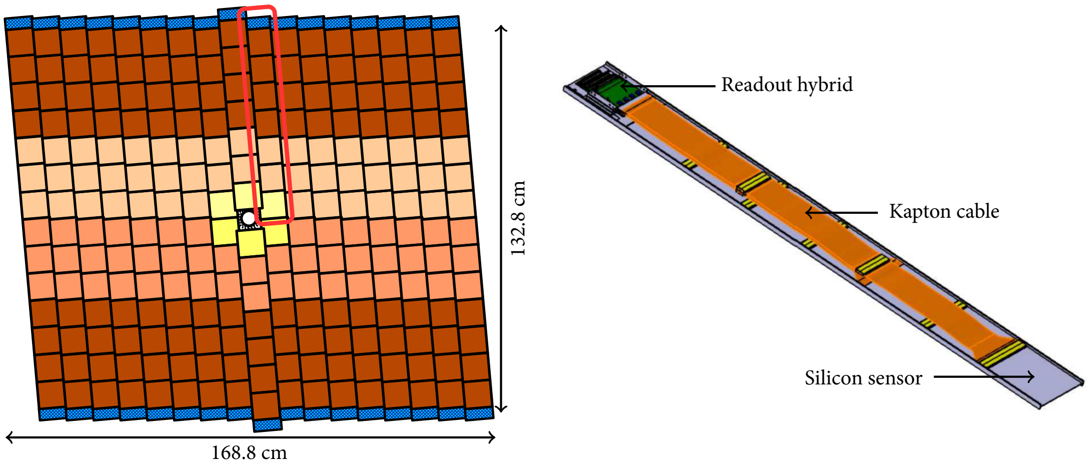

Figure 3.9: A schematic of the third TT layer (left), showing the 5° rotation from vertical, and highlighting one of the constituent half-modules shown individually (right) [59].

A TT layer comprises vertically mounted half-modules, each containing seven silicon sensors, as shown in Fig. 3.9. In the centre of the detector there is one half-module above the beam pipe and one below. To the sides, there are seven rows of pairs of half-modules joined end to end in the first and second layers, and eight rows of similarly joined pairs in the third and fourth layers. Adjacent modules are offset by 10 mm in z and overlap by a few millimetres transversely to aid alignment and ensure there are no gaps in acceptance.

The modules are mounted vertically in the first and fourth layers and at angles of 5° and −5° from vertical in the second and third layers respectively. Each of the silicon sensors in a module is 500 μm thick, and contains 512 strips of 183 μm pitch. These dimensions were chosen to keep the strip occupancy below a few percent and provide a single hit resolution of around 50 μm.

3.3.4 Inner and Outer Trackers

Located downstream of the dipole magnet are three T-stations, T1, T2, and T3, each containing IT and OT modules.

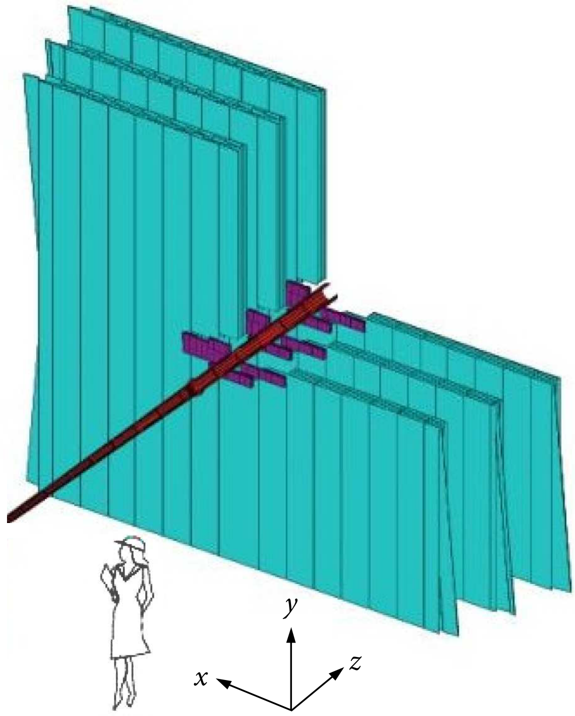

Figure 3.10: A schematic of three quadrants of each of the T-stations showing the beam pipe (red), IT (purple), and OT (turquoise), with a human figure for scale [59].

The IT [70] is a silicon microstrip detector of similar construction to the TT. Each T-station contains four IT boxes arranged in a cross around the beam pipe, 120 cm wide and 40 cm tall. There are four layers in each T-station, with the modules rotated as in the TT. Each layer contains seven modules, with the modules in the boxes above and below the beam pipe having one sensor and those in the boxes to the sides two sensors. Adjacent modules are offset by 4 mm in z and overlap by 3 mm in x.

Each sensor is 11 cm long, 7.6 cm wide, and contains 384 strips of 198 μm pitch. The thickness was determined by the compromise between the signal-to-noise ratio and the material budget, with 320 μm sensors on the single modules and 410 μm sensors on the double modules.

The OT [71] is a straw drift tube detector which has inner acceptance coincident with the outer acceptance of the IT, with the location of the boundary between the IT and OT determined by the OT occupancy limit. The straw tubes in each station are in four layers, with the same rotation of layers as in the TT and IT.

To the left and right of the beam pipe in each layer are long modules (F-type) which are 4850 mm long and contain 256 tubes. Short modules (S-type) are used above and below the beam pipe, which are half the length and number of tubes. In total the OT contains 166 F-type modules and 96 S-type modules. The drift tubes have an internal diameter of 4.9 mm and are filled with a 70:30 mix of Argon and CO_2 for drift times of less than 50 ns and a resolution of 200 μm.

3.3.5 Ring-Imaging Cherenkov detectors

LHCb requires excellent charged hadron identification in order to distinguish between kaons and pions, which is provided by two RICH detectors [72] to achieve the required momentum coverage. RICH detectors measure the angle of emission of Cherenkov photons θ_c as particles pass through a radiator with refractive index $n > 1$ at a velocity greater than the phase velocity of the radiator. The velocity of the particle is given by

$$v = \frac{c}{n \cos \theta_c}, \tag{3.2}$$

which is combined with the momentum measured by the tracking system to calculate the particle mass, and hence identify it. The RICH detectors are crucial for identifying hadronic daughters in CEP processes.

The first RICH detector, RICH 1, is located downstream of the VELO and shown in Fig. 3.12. It covers the entire detector acceptance and the momentum range from 1 GeV/c

to 60 GeV/c. A box surrounding the beam pipe contains aerogel and C_4F_{10} radiators, which have refractive indices of 1.030 and 1.0014, respectively, at 400 nm. The Cherenkov angle in C_4F_{10} is plotted for isolated tracks in Fig. 3.11. Spherical mirrors on the rear face of the box focus the Cherenkov photons onto flat mirrors which reflect them onto two planes of Hybrid Photon Detectors (HPD), one above the box and one below. Each plane consists of seven rows of 14 HPDs, housed in an iron box to shield them from the fringe field of the dipole magnet.

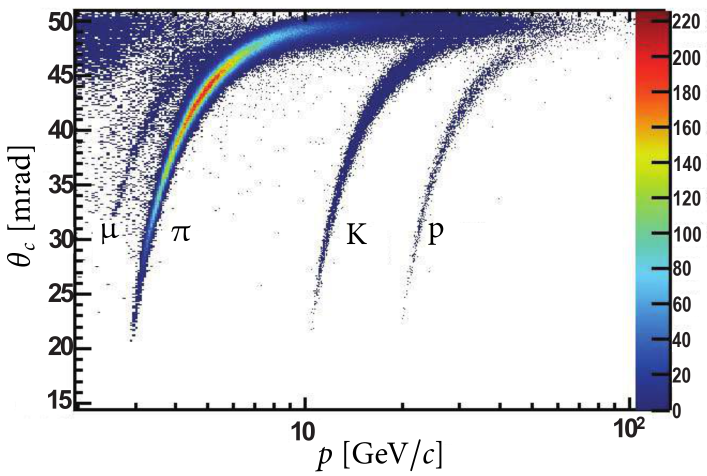

Figure 3.11: The reconstructed Cherenkov angle for isolated tracks in C_4F_{10}, showing the distinguishing power between muons, pions, kaons, and protons [69].

RICH 2 is located downstream of the T-stations and also shown in Fig. 3.12. It has a reduced maximum acceptance of 120 mrad horizontally and 100 mrad vertically, covering the rapidity range of higher momentum particles. A CF_4 radiator with a refractive index of 1.0005 at 400 nm gives momentum coverage from 15 GeV/c to >100 GeV/c. Spherical and flat mirrors focus and reflect Cherenkov photons as in RICH 1, however in RICH 2 the optical system is horizontal rather than vertical, and each plane consists of nine rows of 16 HPDs.

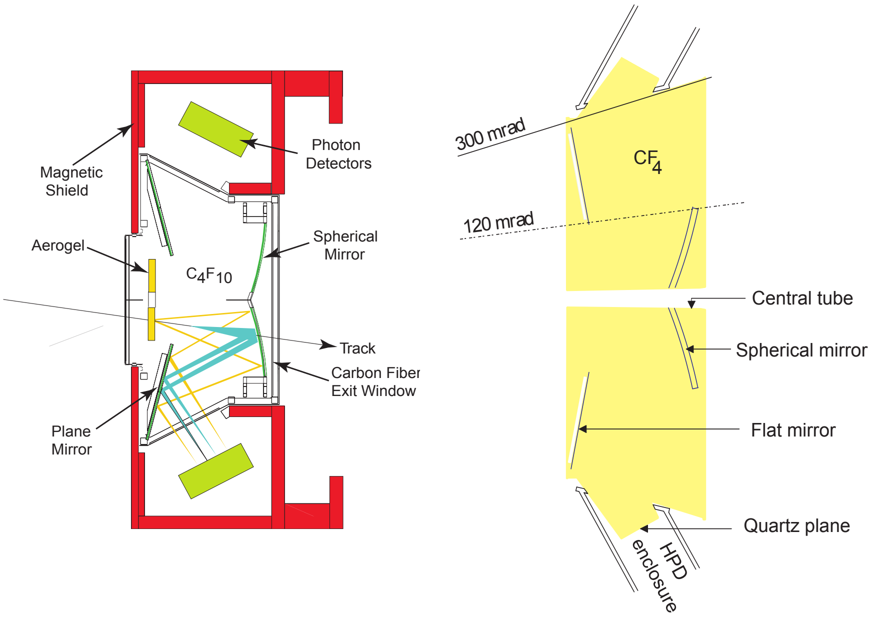

Figure 3.12: Cross-sectional schematics of RICH 1 (left, viewed from the side) and RICH 2 (right, viewed from above). The scale is not shared; RICH 2 is approximately twice as large as RICH 1 [59].

The HPDs are custom photon detectors sensitive to wavelengths from 200 nm to 600 nm. Incident photons are converted to photoelectrons in a photocathode, which are accelerated onto a silicon pixel sensor by a 20 kV potential difference as shown in Fig. 3.13. Each silicon sensor contains 1024 square pixels of side 500 μm, in which a photoelectron will create approximately 5000 electron-hole pairs. The HPDs have an active diameter of 75 mm and are hexagonally close packed in their boxes with a centre-to-centre distance of 89.5 mm.

3.3.6 Scintillating Pad and Preshower Detectors

The SPD and PS [73] are the calorimeters nearest the interaction point and aid the ECAL in electron identification. The SPD identifies charged particles to suppress the neutral pion background, and the PS suppresses the charged pion background. The SPD is additionally

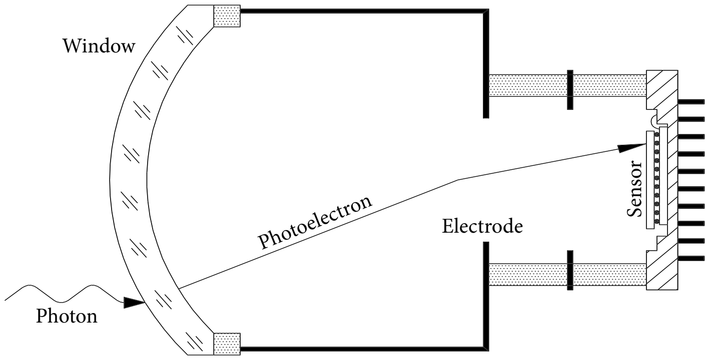

Figure 3.13: A schematic of a RICH HPD with the path of a photoelectron [59].

used to estimate the charged particle multiplicity, in order to reject high multiplicity events when selecting CEP candidates in the trigger.

The SPD and PS are each 7.6 m wide and 6.2 m wide, with the PS 56 mm downstream of the SPD. Between them is a 15 mm thick lead converter to initiate showering before the PS. Each detector consists of a plane of square scintillator cells, which is split into three regions of different cell sizes, as shown in Fig. 3.14, due to the transverse variation in hit density. The transverse dimensions of the SPD are 0.45 % smaller than those of the PS, so that the components are the same size as projected from the interaction point. Each scintillator cell is 15 mm thick and read out by a multi-anode PMT.

3.3.7 Electromagnetic Calorimeter

The ECAL [73] is composed of alternating absorber and scintillator layers, and provides photon and neutral pion information. The design was chosen for its reliability and fast response time, whilst still offering sufficient energy resolution and radiation hardness. The design energy resolution is $\sigma_E/E = 10\,\%/\sqrt{E} \oplus 1\,\%$ for E in GeV. The ECAL is located downstream of the PS and has an outer acceptance of 300 mrad horizontally and 250 mrad

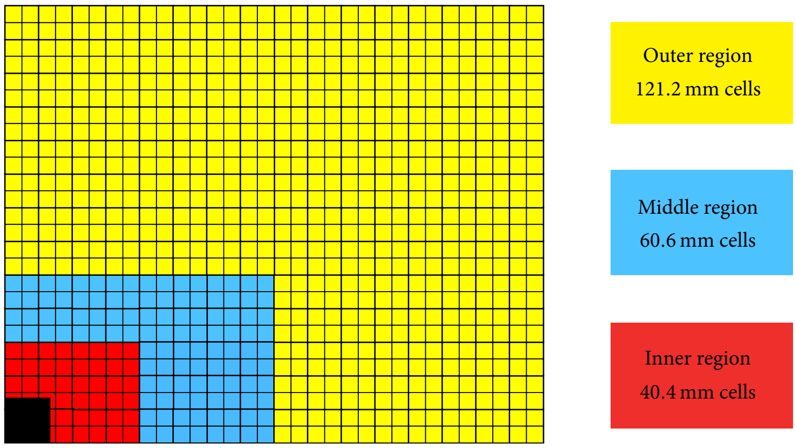

Figure 3.14: A schematic of the upper right quadrants of the SPD, PS, and ECAL, showing the three regions of different cell sizes. The labelled dimensions correspond to the ECAL [59].

vertically. The inner acceptance is limited to 25 mrad by the radiation dose near the beam pipe.

The ECAL is split into three regions transversely in the same manner as the SPD and PS as shown in Fig. 3.14. The detector is 42 cm deep with alternating 4 mm thick layers of lead absorbers and polystyrene scintillator tiles. A 120 μm thick sheet of white reflecting paper between each layer reduces light leakage. The thickness corresponds to 25 radiation lengths, chosen to be sufficient to fully contain a high energy shower, and the Molière radius is 35 mm. The scintillator tiles are read out with Hamamatsu PMTs via wavelength shifting fibres.

3.3.8 Hadronic Calorimeter

The HCAL [73] is located downstream of the ECAL, and has the main purpose of selecting events with high transverse energy in the trigger. It is also a sampling calorimeter, however in contrast to the ECAL the scintillator tiles are mounted in parallel to the beam axis rather than layered perpendicularly to it, as shown in Fig. 3.15.

This orientation was chosen to increase the contact area between the scintillator tiles and optical fibres and so maximise the amount of scintillation light collected. Test beam studies showed the longitudinal orientation provides a light yield two orders of magnitude greater than that expected from a transverse orientation [74].

The scintillator tiles are doped polystyrene, separated laterally by 100 mm thick iron absorber plates. The HCAL is segmented into two regions, rather than three as for the other calorimeters, due to the larger size of hadronic showers. The tile size in each region is given in Fig. 3.16. Wavelength shifting fibres run from the front of a cell to the rear where the PMTs are mounted and each fibre contacts three scintillator tiles. The whole assembly is 8.4 m tall, 6.8 m wide, and 1.65 m deep, corresponding to one interaction length in steel. The measured energy resolution is $\sigma_E/E = (69 \pm 5)\,\%/\sqrt{E} \oplus (9 \pm 2)\,\%$ for E in GeV, where the rather large constant term arises from manufacturing inhomogeneities. This is not detrimental as the primary purpose of the HCAL, selecting events with hadron candidates in the Level-0 trigger, requires only modest energy resolution [59].

3.3.9 Muon detectors

Efficient muon triggering and identification is important at LHCb as many b-hadron decays sensitive to CP violation contain muons in the final state. The muon detectors [75–77] comprise the five stations shown in Fig. 3.17, M1 to M5, with an acceptance of 20 mrad to 306 mrad horizontally and 16 mrad to 258 mrad vertically. All the stations are located downstream of the calorimeters, except for M1 which is located between RICH 2 and the calorimeters to improve the $p_\perp$ measurement in the trigger. M2 to M5 are layered between iron absorbers to improve sensitivity to the most penetrating muons, and the five stations together constitute approximately 20 interaction lengths. Like the SPD and PS, the transverse dimensions of the stations scale projectively with distance from the interaction point. M1 to M3 have the best spatial resolution and measure $p_\perp$ with a resolution of 20 %.

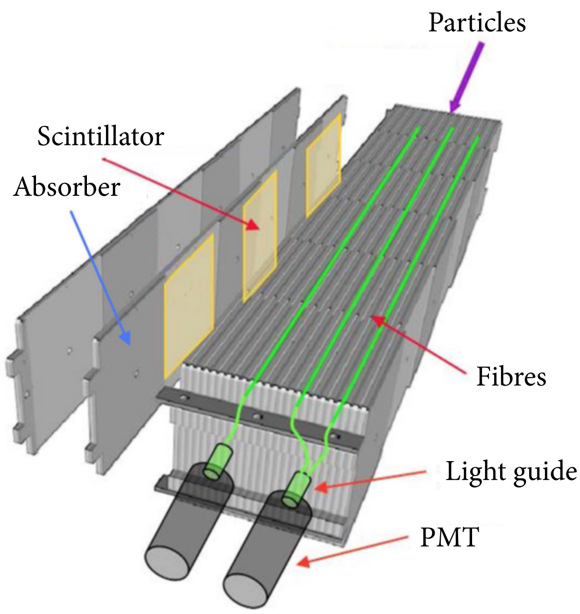

Figure 3.15: A schematic of the periodic structure of the HCAL, showing the alternating absorbers and scintillators, and how the wavelength shifting fibres run from the front to the rear of the detector [59].

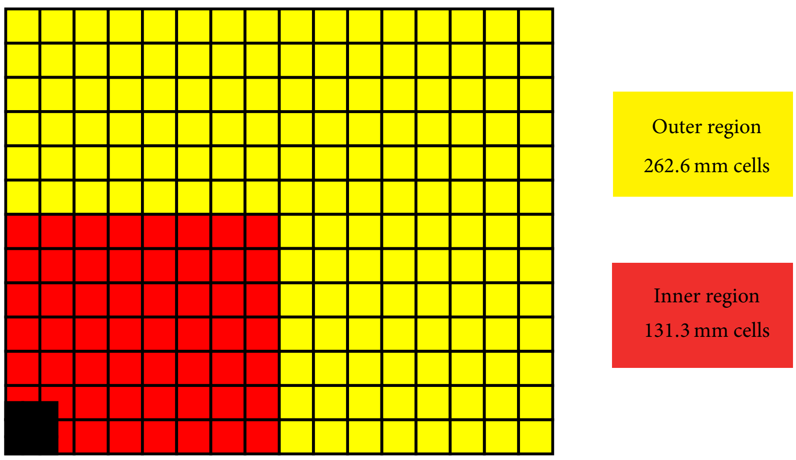

Figure 3.16: A schematic of the upper right quadrant of the HCAL, showing the two regions of different cell sizes [59].

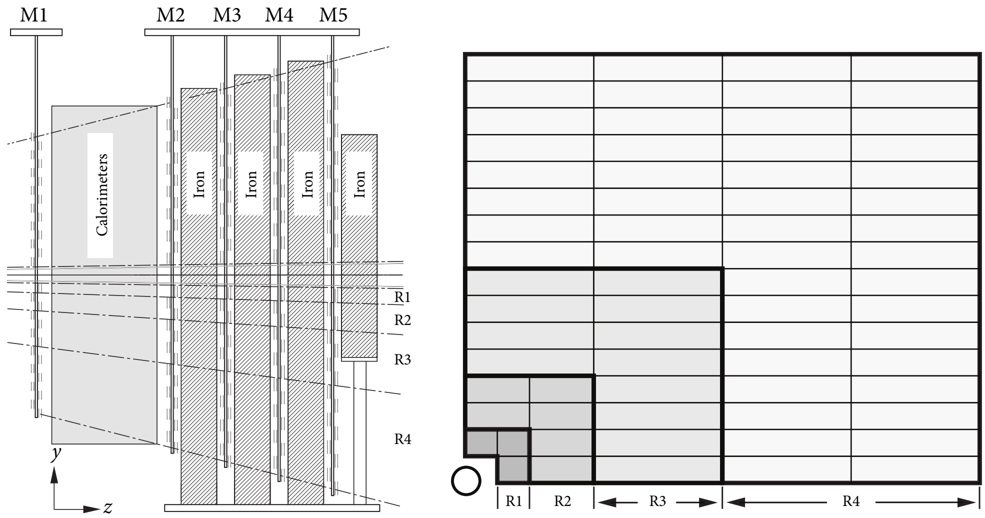

Figure 3.17: A cross-sectional schematic of the muon detectors (left), showing M1 in front of the calorimeters and the remaining detectors behind, and a schematic of the upper right quadrants of M2 and M3 (right) showing the four transverse regions. [59].

Each station is divided into four transverse regions, R1 to R4, with linear dimensions and segmentation in the ratio 1:2:4:8 to keep the occupancy approximately equal in each region. There are 276 chambers in each station, and each chamber is further divided into pads. There are twice as many pads per chamber in each region of M2 and M3 as in the corresponding region of M1, and half as many in the corresponding region of M4 and M5. The pad layout in M2 and M3 is shown in Fig. 3.17.

The sensors in each pad are multi-wire proportional chambers (MWPC) in all but the inner region (R1) of M1, where the hit rate is too high and so gas electron multipliers (GEM) are used instead. The MWPCs are filled with a 40:55:5 mix of $Ar:CO_2:CF_4$ and contain four gas gaps each (two in M1).

The 12 chambers in the R1 region of M1 each contain two superimposed triple-GEMs. Ionisation electrons are created in a drift gap containing a 45:15:40 mixture of $Ar:CO_2:CF_4$. The electrons are accelerated through three 50 μm foils, which contain 50 μm diameter holes.

A large electric field in the holes causes electron multiplication resulting in an avalanche which is detected as a current pulse at the anode.

3.4 Trigger

The design LHC bunch crossing rate is 40 MHz, although the accelerator was operated at 20 MHz in Run 1. There is insufficient bandwidth and processing power to read out the full detector and write every event to permanent storage at these rates. The role of the trigger is to identify and save events interesting for analysis or detector calibration whilst discarding background events [78]. At a luminosity of $4 \times 10^{32}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, as achieved in 2012 when the data analysed in this thesis were recorded, the rate of bunch crossings with at least one inelastic interaction is greater than 11 MHz, which the trigger must reduce to 2–5 kHz to write to permanent storage [69, 79].

The trigger consists of two stages, Level-0 (L0) and the High Level Trigger (HLT). L0 operates at the bunch crossing rate, independent of the detector occupancy, and so uses custom electronic hardware. The HLT processes events passing L0 and is implemented in software running on commercially available processors. L0 and the HLT can be emulated in software in order to study the trigger efficiency using stored or simulated data. A uniquely named sequence of cuts and thresholds used in the trigger to select a specific decay is called a trigger line, and an event is retained only if it passes at least one trigger line in both L0 and the HLT.

3.4.1 Level-0 trigger

The L0 trigger uses information which can be read out quickly from the calorimeters, muon detectors, and pile-up veto system to reduce the event rate to 1 MHz, at which rate the full detector is read out. Cuts are made on global event quantities, such as the calorimeter transverse energy $E_{\perp}$ and the SPD multiplicity, to reject high multiplicity events which would use excessive bandwidth and processor time in the HLT.

Typical events of interest contain high $E_{\perp}$ and $p_{\perp}$ particles, such as the daughters of b and c-hadrons. Therefore the L0 trigger reconstructs the calorimeter clusters with the highest $E_{\perp}$ and the muon candidates with the highest $p_{\perp}$ to build electron, hadron, photon, and muon candidates. These candidates are then evaluated against the L0 trigger lines, and events passing one or more lines forwarded to the HLT. The lines used to select CEP candidates additionally cut tightly on the number of hits in the SPD and the PU system, to reject events that cannot be a single CEP event due to their high multiplicity, and which allows a looser $E_{\perp}$ cut to be applied.

3.4.2 High Level Trigger

The HLT receives L0 triggered events at a rate of 1 MHz, and then has access to the full event information from all subdetectors. The HLT is divided into two stages, HLT1 and HLT2, implemented in software running on 29,000 logical cores in the Event Filter Farm (EFF).

HLT1 carries out a partial event reconstruction at the L0 output rate. Tracks in the VELO are reconstructed and cuts made on the impact parameter and the number of hits in order to identify primary vertices. For events with L0 muon candidates, VELO tracks are extrapolated to the muon detectors and matching muon detector hits searched for. Where matching hits are found, a linear χ^2 fit is carried out and a momentum cut applied to identify muon candidates. The IT and OT hits corresponding to the reconstructed VELO tracks and muon candidates are then identified and fitted with a Kalman filter [80]. Special provision is made for candidate CEP events which allows them to pass directly from L0 to HLT2, bypassing HLT1. This is possible because the smaller size of a candidate CEP event compared to a typical inelastic event reduces the CPU time needed to process it in HLT2.

Events selected by HLT1 are passed to HLT2, at a rate of 43 kHz in 2011 and 80 kHz in 2012. This reduced rate allows HLT2 to carry out a full event reconstruction using pattern recognition algorithms. Tracks are selected with transverse momentum and impact parameter cuts, with a transverse momentum threshold of 100 MeV/c for candidate CEP events. These are looser than those applied in an offline selection to allow improved calibrations to be applied offline. Composite particles are constructed by combining tracks and applying cuts on the mass and vertex fit quality. Events are selected to be written to permanent storage by an inclusive OR of the trigger lines. HLT2 was designed to accept events at a rate of 2 kHz, however a rate of 3 kHz was achieved in 2011. In 2012 the rate was increased to 5 kHz using deferred triggering, whereby 20 % of L0 accepted events were buffered on disk in the EFF and processed in the downtime between LHC fills.

3.5 Offline data processing

Triggered events are buffered in the EFF before transfer to the CERN Tier-0 computing centre. The Tier-0 centre is responsible for reconstructing raw data, which consists of building tracks from hits in the trackers, showers from calorimeter clusters, and identifying particles using data from the RICH detectors, calorimeters, and muon detectors. The offline reconstruction is more precise than that carried out in the trigger.

Tracks are reconstructed between the VELO and calorimeters from hits in the VELO, TT, and T-stations. The reconstruction begins by looking for track segments in the VELO and T-stations, the trackers in the lowest magnetic field, and seeding a Kalman filter with them. Tracks are built by iteratively adding segments from each tracker whilst accounting for multiple scattering and energy loss via material interactions. Figure 3.18 shows an example of a set of tracks resulting from this process, along with the hits from which they were reconstructed.

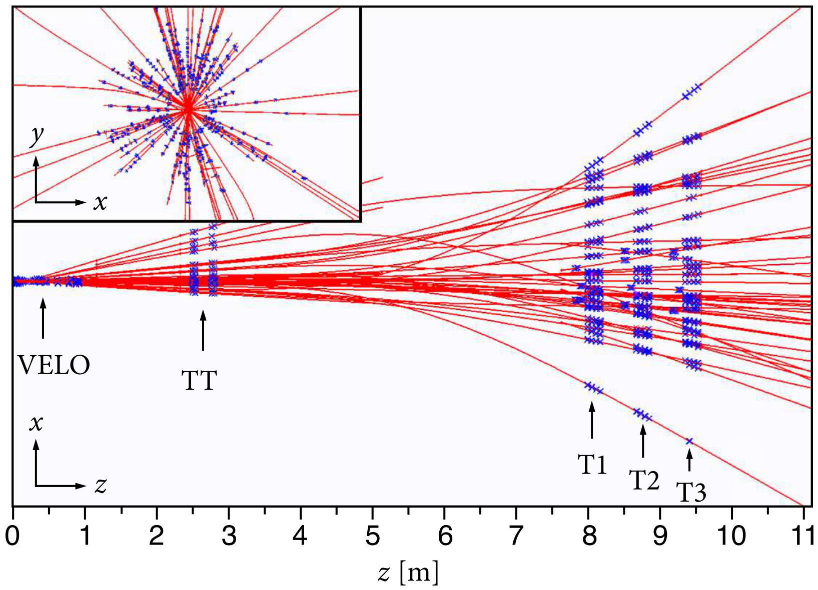

Figure 3.18: Tracks (red lines) in an LHCb event and the tracker hits (blue points) they were reconstructed from [59].

The track reconstruction efficiency is determined from simulation, and is defined as the fraction of reconstructible tracks which are reconstructed, where a reconstructible track has more than a threshold number of hits in the trackers, and a track is reconstructed if at least 70 % of its hits were associated with it. Also calculated is the *ghost rate*; the fraction of reconstructed tracks which are not matched with any simulated track. The track reconstruction efficiency and ghost rate determined from simulation have been validated using collision data [81, 82].

Five different track types are reconstructed at LHCb, as shown in Fig. 3.19. Long tracks are reconstructed from hits in every tracker from the VELO to the T-stations, and provide the best momentum resolution. VELO tracks have hits in the VELO only, and due to their typically low pseudorapidity are important for vertex reconstruction. Upstream tracks have hits in the VELO and TT only, and are typically particles bent out of the detector acceptance by the magnet before reaching the T-stations. Downstream tracks have hits in the TT and

T-stations only and correspond to the daughters of a particle which decayed after transiting the VELO, such as a K^0_S, and finally T tracks have hits only in the T-stations.

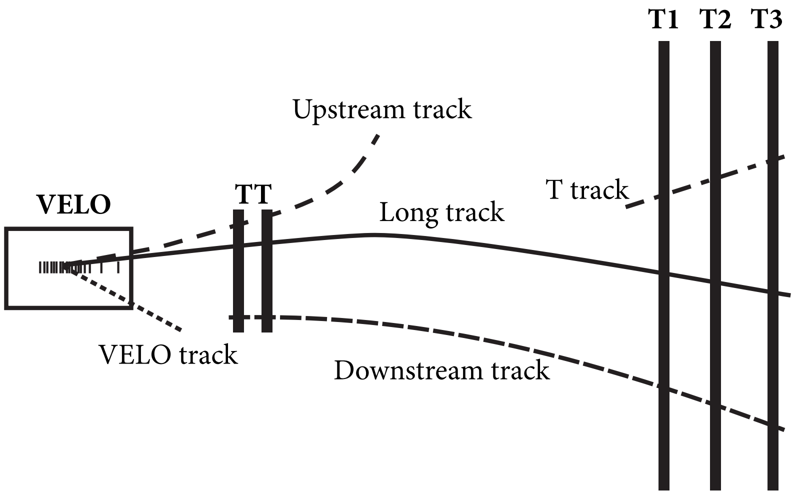

Figure 3.19: The five types of track reconstructed at LHCb, as described in the text [59].

A reconstruction efficiency of ~94 % and ghost rate of ~9 % is found for long tracks with momentum greater than 10 GeV/c, with corresponding values of ~75 % and ~15 % for upstream tracks with momentum greater than 1 GeV/c. An efficiency of 80 % is found for downstream tracks with momentum greater than 5 GeV/c.

Electrons, muons, pions, kaons, and protons are identified using information from the RICH detectors, calorimeters, and muon detectors. Hadrons are identified by comparing the RICH pixel image for each reconstructed track with that expected for each particle identity, and a likelihood calculated for each hypothesis. The hypothesis with the highest likelihood is assigned as the particle identity, quantified by the difference in the log likelihood with the pion hypothesis. For example, $\Delta \ln \mathcal{L}_{K-\pi}$ is the difference between the log likelihood of the kaon and pion hypotheses. The average kaon identification efficiency is ~95 % for momentum from 2 GeV/c to 100 GeV/c, with a pion misidentification rate of ~5 %, with $\Delta \ln \mathcal{L}_{K-\pi}$ cuts of > 0 and > 5 respectively, as shown in Fig. 3.20.

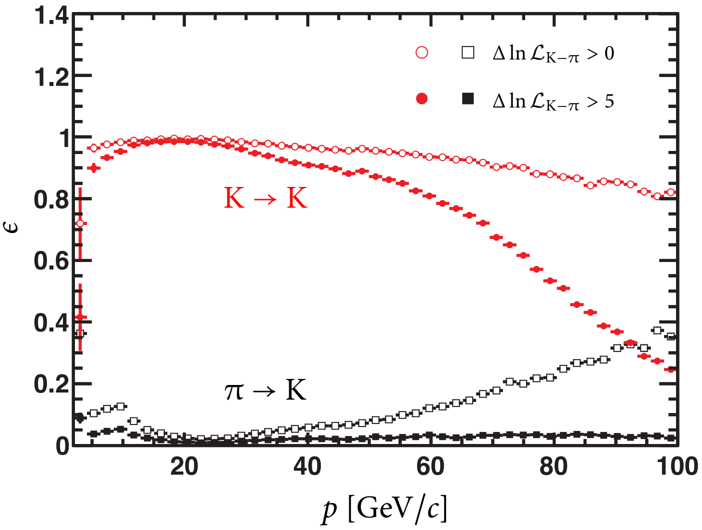

Figure 3.20: The kaon identification efficiency and pion misidentification rate versus track momentum [69].

Electrons are identified as ECAL clusters with the correct ratio of energy to momentum and where the cluster is coincident with an extrapolated track trajectory. The extrapolated track is also used to match neutral clusters from Bremsstrahlung with the electron candidate, in order to make the correct energy measurement. An electron identification efficiency of ~95 % is determined with a pion misidentification rate of 0.7 %.

Photons are identified as neutral ECAL clusters with a measure of coincidence with each reconstructed track above threshold, such that the cluster is not associated with any charged particle. Photons which convert in detector material before reaching the ECAL are identified as clusters with an SPD hit directly in front.

Neutral pions with $p_{\perp} < 3\,\text{GeV}/c$ are reconstructed from pairs of photons. Above 3 GeV/c, the ECAL granularity is insufficient to resolve the individual photons and a single merged cluster is observed. In this case the π^0 is identified using an iterative algorithm which splits the cluster in two and attempts to reconstruct a photon from each sub-cluster.

Muon candidates are identified by extrapolating track trajectories through the muon detectors and searching for hits around the extrapolated trajectory. Hits in two to four stations (depending on momentum) are required to identify a muon candidate. Additional discrimination is provided by requiring that the hits in the muon detectors match the extrapolated trajectory, which gives an identification efficiency of ~93 % and pion misidentification rate of 1 %. The muon identification efficiency is plotted in Fig. 3.21.

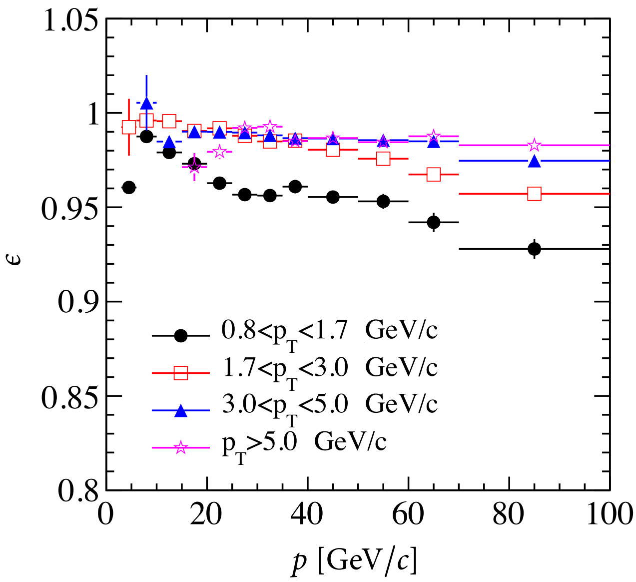

Figure 3.21: The muon identification efficiency versus momentum [83].

The reconstructed data are stored in Data Summary Tape (DST) files which undergo an additional processing offline called stripping, in which selections optimised to select events of interest to specific analyses are applied. The stripping is carried out on the Worldwide LHC Computing Grid, so the workload can be distributed amongst multiple computer centres. Primary and secondary vertices are reconstructed from tracks in the VELO, the four vectors of identified particles are calculated, and composite particles are reconstructed by combining daughter candidates. The stripping is repeated several times per year to enable each selection line to be improved iteratively. The events passing each stripping line are written to DST files and made available to LHCb collaborators for analysis.

Figure 3.22 shows the momentum and mass resolutions of the detector. The relative momentum resolution $\delta p/p$ is approximately 0.5 % for long tracks with momentum less than 20 GeV/c, measured from $J/\psi \to \mu^+\mu^-$ decays. The relative mass resolution $\delta m/m$ is approximately 0.5 % for masses up to ~10 GeV/c^2, measured by fitting the invariant mass distributions for the decays of six resonances (J/ψ, ψ(2S), ϒ(1S), ϒ(2S), ϒ(3S), and Z) to dimuons.

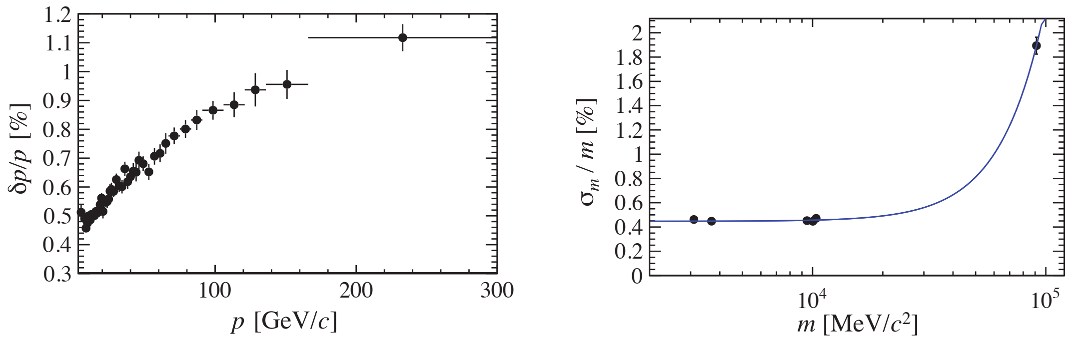

Figure 3.22: The relative momentum resolution for long tracks (left) and the mass resolution (right) measured using dimuon resonances [69].

3.6 Event simulation

The simulation of proton-proton collisions and the detector response is an essential tool for studying selection efficiencies and detector performance. Such simulations are carried out using the LHCb simulation applications Gauss and Boole [84–86], which are written in C++ using Gaudi [87, 88], an object-oriented event processing framework for particle physics experiments.

Gauss runs two main stages: (i) the generation of the primary event and its decay products, and (ii) the transport of these decay products through the detector volume. For most processes, the primary event is generated with PYTHIA [64, 89], and the decays of primary particles with EvtGen [90], using PHOTOS to model final state radiation [91].

Both PYTHIA and EvtGen are integrated into Gauss. Other tools may be substituted for specific purposes, such as particle guns for generating single particles with a given momentum and pseudorapidity, or special purpose event generators such as SuperCHIC for simulating CEP events [28, 29, 38]. Such event generators are interfaced with Gauss using the HepMC event record format [92].

In the second stage, the transport of particles through the detector volume is performed using the Geant4 toolkit [93, 94]. The geometry of the detector is stored in XML in the DDDB (Detector Description Database), which is parsed by Gauss and converted to Geant4's internal geometry format.

After Gauss, Boole is used to model the response of each of the subdetectors. The subdetector outputs are digitised and converted into the same formats as produced by the front-end electronics on the detector. After digitisation, simulated data are passed through the trigger, reconstruction, and stripping in the same manner as for collision data.

4

Exclusive $\chi_c(1P) \to h^+h^-$ candidate selection

In this chapter the experimental characteristics of CEP processes are discussed and the main background processes described. The collision data set and samples of simulated data used to analyse $\chi_c(1P) \to h^+h^-$ CEP are described. The trigger, stripping, and offline selections are presented, and simulated signal events are used to determine the selection efficiencies for each of the modes studied, with data-driven corrections made where necessary.

4.1 Experimental signature and backgrounds

A characteristic feature of CEP processes at LHCb is the small amount of activity observed in the detector. As the initial state protons do not fragment in the collision, the only particles present in the final state are the decay products of the central exclusively produced state (except for any particles originating in pile-up collisions). There are no additional hard interactions and therefore the global event multiplicity is much lower than in the inelastic collisions typically observed at LHCb. An event selection can use this fact to preferentially select CEP events while rejecting inelastic processes by requiring that exactly the number of tracks produced in the mode under study are reconstructed in the detector, and rejecting events with extra tracks.

As the initial state protons remain intact in CEP, their trajectory is only slightly perturbed from the beam axis and they do not enter the detector acceptance, remaining in the beam pipe in the vicinity of the interaction region. This defines two rapidity gaps, regions of the detector which are devoid of any activity, between the central exclusively produced system and each of the final state protons, as illustrated in Fig. 4.1.

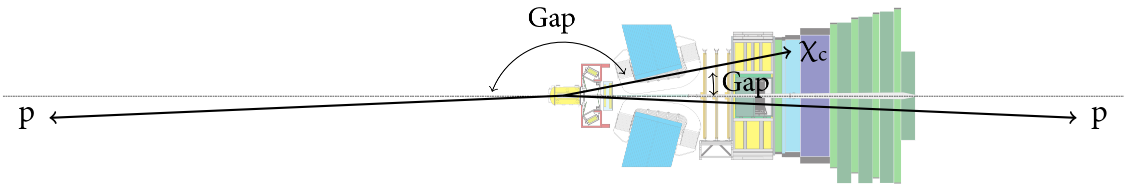

Figure 4.1: A schematic of the rapidity gaps defined by the regions empty of activity between the centrally produced system and the outgoing protons, in the context of $\chi_c(1P)$ CEP at LHCb. The deflections of the protons from the beam axis and the $\chi_c(1P)$ flight distance are exaggerated for illustration.

Requiring that there are two rapidity gaps extending continuously from the central system to each of the outgoing protons imposes a strict requirement on the exclusivity of an event. However, the sizes of the rapidity gaps which can be enforced at LHCb are restricted by the limited acceptance of the detector, especially at negative pseudorapidity. A total gap of 5 units of pseudorapidity can be achieved, corresponding to the VELO coverage of $-3.5 < \eta < -1.5$ in the backward direction and $1.5 < \eta < 4.5$ in the forward direction. Inelastic events may mimic the signature of CEP events if the additional particles produced in the collision fall outside this acceptance and so are not detected. In this case an inelastic event is indistinguishable from CEP based on the multiplicity within the LHCb acceptance alone.

The inelastic processes in which it is most likely that additional particles will fall outside of the detector acceptance are diffractive processes with proton dissociation. These processes proceed in the same manner as CEP, via photon and Pomeron exchange, and may produce identical central systems. However, in the collision one (in single diffraction) or both (in double diffraction) of the protons dissociate. At the LHC the fragments must carry the large momentum of the initial state protons and so leave the interaction region at large absolute rapidity, outside the acceptance of the LHCb detector. These diffractive events therefore closely mimic the signature of central exclusive events in LHCb, and cannot be rejected from the analysis by requiring an exact track multiplicity and two rapidity gaps alone. They therefore form the predominant background in the analysis. Another background to CEP which cannot be excluded only by requiring rapidity gaps and cutting on track multiplicity are inelastic events that would otherwise be CEP except for having additional gluon radiation outside the detector acceptance. Feynman diagrams of CEP, single diffraction, double diffraction, and central production with additional gluon radiation are shown in Fig. 4.2.

However, the larger the momentum transfer between the protons the more probable it is that a proton dissociates [95]. Therefore the transverse momentum of the central system provides additional discrimination between central exclusive and inelastic diffractive events. As a result, a low transverse momentum of the central system may be used in addition to rapidity gaps and the lack of activity in the detector to preferentially select CEP candidates, and this approach is used in this thesis.

Another possible background is feed down, where a higher mass state is central exclusively produced, and decays to a $\chi_c(1P)$ meson and other particles. If the $\chi_c(1P)$ decays to K^+K^- or $\pi^+\pi^-$ and is detected whilst the other decay products are not, the $\chi_c(1P)$ will incorrectly be considered to have itself been produced in CEP.

The main feed down background to $\chi_c(1P)$ production is $\psi(2S) \to \chi_c(1P)\gamma$, as the photon is of low energy and so difficult to be detect [41]. However, based on the $\psi(2S)$ and $\chi_c(1P)$ cross sections previously measured in the $\mu^+\mu^-\gamma$ channel, and the known branching fractions, this is expected to be insignificant compared to the inelastic backgrounds described above [40].

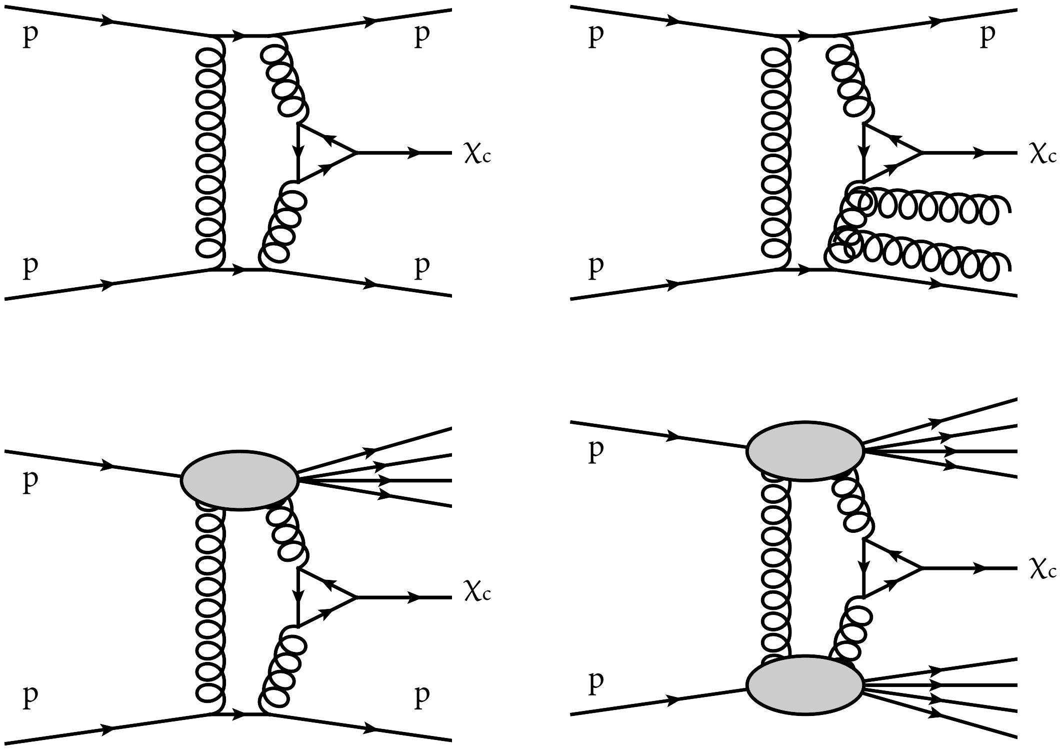

Figure 4.2: Feynman diagrams of CEP (top left) and three backgrounds which cannot be fully distinguished by requiring rapidity gaps and cutting on track multiplicity alone: central production with additional gluon radiation (top right), single diffraction (bottom left), and double diffraction (bottom right).

4.2 Data set

The analysis of $\chi_c(1P) \to h^+h^-$ CEP is carried out using proton-proton collision data collected with the LHCb detector during 2012, at a centre-of-mass energy of 8 TeV. The total data set collected by LHCb in 2012 corresponds to an integrated luminosity of 2 fb^{-1}, half collected with each polarity of the dipole magnet [96]. However, the HLT2 line used to select $\chi_c(1P) \to h^+h^-$ candidates was only commissioned part way through the year, and thus only 34 % of the 2 fb^{-1} was recorded with this line operational and can be used in the analysis. Additionally, in imposing the requirements on activity and rapidity gaps discussed in Section 4.1, bunch crossings with more than a single visible interaction are not selected, even if that bunch crossing contains a genuine CEP event. Due to pile-up, this corresponds to an effective single interaction integrated luminosity used in the cross section measurement that is approximately 18 % of the total integrated luminosity. The precise determination of the effective integrated luminosity given these trigger and pile-up requirements is described in Section 5.4. Only events recorded when all subdetectors were operational and functioning correctly are analysed.

4.3 Event simulation

Simulations of CEP events and their interactions with the LHCb detector are used in the analysis to guide the choice of selection cuts, and to evaluate the trigger, stripping and offline selection efficiencies which are required inputs to a cross section measurement.

The SuperCHIC event generator [28,29,38], implementing the Durham model described in Section 2.4, is used to simulate $\chi_c(1P) \to h^+h^-$ CEP. SuperCHIC was used to simulate each of the four modes studied; that is each combination of $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$ mothers and K^+K^- and $\pi^+\pi^-$ final states. Samples were generated containing 10×10^6 $\chi_{c0}(1P) \to K^+K^-$ events, 10×10^6 $\chi_{c0}(1P) \to \pi^+\pi^-$ events, 5×10^6 $\chi_{c2}(1P) \to K^+K^-$ events, and 5×10^6

$\chi_{c2}(1P) \to \pi^+\pi^-$ events. No acceptance cuts were applied during the generation, such that events are distributed throughout the kinematically accessible phase space.

After generation with SuperCHIC the events are passed through the LHCb simulation application Gauss [84, 85], which models material interactions in the detector, and digitisation application Boole [86], which models the resulting subdetector responses and digitisation of the signals. After digitisation simulated events are passed through the same reconstruction and stripping pipeline as collision data, as described in Section 3.5, including software emulation of the trigger.

4.4 Trigger and stripping selections

To be considered as $\chi_c(1P) \to h^+h^-$ candidates, events are required to pass specific individual lines in the L0 trigger, HLT, and stripping. The L0 and HLT lines are common to all four modes, whereas in the stripping there are dedicated lines for the $\chi_c(1P) \to K^+K^-$ and $\chi_c(1P) \to \pi^+\pi^-$ modes, as listed in Table 4.1. The selection cuts implemented in each line are described in the following sections.

Table 4.1: The lines $\chi_c(1P) \to h^+h^-$ candidates are required to pass at each stage of the trigger and stripping.

Stage	Line
L0	`DiHadron,lowMult`
HLT1	`Hlt1NoPVPassThrough`
HLT2	`Hlt2LowMultChiC2HH`
Stripping	`StrippingLowMultCEP_ChiC2KK_line` or `StrippingLowMultCEP_ChiC2PiPi_line`

4.4.1 Level-0 trigger

In the L0 trigger, candidate events are required to pass the `DiHadron,lowMult` line. This line is designed to select events containing two or more hadronic tracks but having low global event multiplicity, to reject high multiplicity inelastic events. To identify events with hadrons, it requires more than 500 MeV of transverse energy to be deposited in the hadronic calorimeter. The low global multiplicity requirement is evaluated using information from the pile-up veto system and scintillating pad detector, described in Section 3.3.1 and Section 3.3.6 respectively. Comprising the final two sensors in the VELO in the backward direction, the PU system gives an indication of the event multiplicity at negative pseudorapidities, whereas the SPD, mounted in front of the calorimeters, gives a measure of the multiplicity at positive pseudorapidities. There are required to be fewer than three hits in the PU system and fewer than ten hits in the SPD. The PU multiplicity cut is lower than the SPD multiplicity cut as there should be no tracks in the backward direction, but a handful of tracks (the genuine daughters) in the forward direction, in order for a CEP event to be reconstructed in the LHCb detector. A summary of the cuts in the `DiHadron,lowMult` line is given in Table 4.2.

Table 4.2: A summary of the requirements of the `DiHadron,lowMult` L0 trigger line. A pre-scale of 25 % is applied, so the cuts are only evaluated against every fourth event.

Variable	Cut
Hadron $E_\perp$	> 500 MeV
SPD multiplicity	< 10
PU multiplicity	< 3

Due to the high rate at which events must be processed by the L0 trigger (imposed by the 50 ns bunch crossing interval) and the limited readout bandwidth to be shared among all trigger lines, a pre-scale of 25 % is applied to the `DiHadron,lowMult` line. This means that only every fourth event is evaluated against the cuts, with the remaining three in four events discarded immediately in order to keep the time and bandwidth usage of the line within limits.

4.4.2 High Level Trigger

The low multiplicity hadronic events accepted by the `DiHadron,lowMult` L0 line are forwarded to HLT1, where they are evaluated by the `Hlt1NoPVPassThrough` line. As the name suggests, this line is a "pass through" which forwards all the events it receives onto HLT2, without applying any further cuts, and with 100 % efficiency. This is possible because the smaller data size of a CEP event compared to a typical inelastic event, due to the reduced event multiplicity, means reduced time is required to process each event in HLT2. `NoPV` in the name indicates that the events are not required to have a reconstructed primary vertex (PV) to pass the line; a typical HLT1 line for a non-CEP process attempts to reconstruct primary vertices from VELO tracks and only passes those events for which a primary vertex is identified.

In HLT2 a more stringent selection is applied than in L0 or HLT1. Events are required to pass the `Hlt2LowMultChiC2HH` line, which is designed specifically to select $\chi_c(1P) \to h^+h^-$ CEP events with high efficiency. This line was added to the trigger part way through 2012, and therefore only data recorded after the commissioning of the line are analysed. In `Hlt2LowMultChiC2HH` cuts are applied on global quantities such as the number of tracks, on quantities associated with the candidate daughter tracks such as their transverse momentum, and on quantities associated with the mother particle which is reconstructed from the combination of the daughter tracks after a vertex fit, such as the vertex fit quality.

There must be fewer than six VELO tracks in the event, a VELO track being a track formed from hits in the VELO which may or may not have additional hits in the tracking stations. The cut is placed at a maximum of five tracks rather than two tracks, the number expected in the $\chi_c(1P) \to h^+h^-$ modes, so that events with ghost tracks reconstructed from noise in the VELO and T-stations still pass the line. Identification and removal of ghost tracks is carried out in the offline processing. That the cut is placed at fewer than six tracks also allows a sample of events with a similar signature but higher track multiplicity to be collected for use in determining the exclusive purity using the method described in Section 5.3. There must be no backward tracks, which are VELO tracks with pseudorapidity less than zero, to enforce the backward rapidity gap. Each of the two tracks being considered as candidate $\chi_c(1P)$ daughters must have momentum greater than 5 GeV/c, χ^2 per degree of freedom less than three, and transverse momentum greater than 100 MeV/c, imposed by the track fit used in the HLT. Information from the RICH detectors is used to distinguish the K^+K^- final state; kaon candidates must have $\Delta \ln \mathcal{L}_{K-\pi}$ greater than zero, whereas no cut on $\Delta \ln \mathcal{L}_{K-\pi}$ is applied to pion candidates in `Hlt2LowMultChiC2HH`.

The mother candidate reconstructed from the two daughter tracks must have momentum greater than 10 GeV/c, transverse momentum less than 5 GeV/c, and mass between 2850 and 3600 MeV/c^2, forming a window around the $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$ masses at 3414.75 MeV/c^2 and 3556.20 MeV/c^2 respectively [97]. The distance of closest approach (DOCA) between the two tracks must be less than 0.5 mm and the vertex χ^2 per degree of freedom less than 15, to select events with a good quality vertex. A summary of cuts in the `Hlt2LowMultChiC2HH` line is given in Table 4.3. No pre-scale is applied in the HLT, so all events received from `DiHadron,lowMult` are evaluated by `Hlt2LowMultChiC2HH`.

4.4.3 Stripping

Events passing the trigger are saved to disk and processed offline in the stripping. The stripping contains separate lines for the $\chi_c(1P) \to K^+K^-$ and $\chi_c(1P) \to \pi^+\pi^-$ modes, `StrippingLowMultCEP_ChiC2KK_line` and `StrippingLowMultCEP_ChiC2PiPi_line` respectively, whose cuts are aligned with the cuts in the `Hlt2LowMultChiC2HH` HLT line. The description that follows is common to both `StrippingLowMultCEP_ChiC2KK_line` and `StrippingLowMultCEP_ChiC2PiPi_line` except where stated otherwise.

Table 4.3: A summary of cuts in the `Hlt2LowMultChiC2HH` HLT2 line.

Variable	Cut
Global cuts	
Number of VELO tracks	< 6
Number of backward tracks	< 1
Daughter cuts	
p	$> 5.0\,\mathrm{GeV}/c$
$p_\perp$	$> 100.0\,\mathrm{MeV}/c$
Track χ^2/DoF	< 3.0
$\Delta \ln \mathcal{L}_{\mathrm{K}-\pi}$ ($\mathrm{K}^\pm$ only)	> 0.0
Mother cuts	
p	$> 10.0\,\mathrm{GeV}/c$
$p_\perp$	$< 5.0\,\mathrm{GeV}/c$
m	$> 2850.0\,\mathrm{MeV}/c^2$
m	$< 3600.0\,\mathrm{MeV}/c^2$
DOCA	$< 0.5\,\mathrm{mm}$
Vertex χ^2/DoF	< 15.0

There are required to be fewer than six tracks globally, no backward tracks, and two or more long tracks, where long tracks are tracks with hits in the VELO and in the tracking stations both before and after the magnet as described in Section 3.5.

Each daughter track must have momentum greater than $5\,\mathrm{GeV}/c$, transverse momentum greater than $100\,\mathrm{MeV}/c$ and track χ^2 per degree of freedom less than three. The `StrippingLowMultCEP_ChiC2KK_line` line requires the daughter tracks to have $\Delta \ln \mathcal{L}_{\mathrm{K}-\pi}$ greater than zero, whereas `StrippingLowMultCEP_ChiC2PiPi_line` contains no cut on $\Delta \ln \mathcal{L}_{\mathrm{K}-\pi}$.

The mother candidate reconstructed from the two daughter tracks must have momentum greater than $10\,\mathrm{GeV}/c$, transverse momentum less than $5\,\mathrm{GeV}/c$, and mass between 2850 and $4500\,\mathrm{MeV}/c^2$. The distance of closest approach between the two tracks must be less than 0.5 mm and the vertex χ^2 per degree of freedom less than 15.

A summary of the cuts is given in Table 4.4. There is repetition of many cuts from the `Hlt2LowMultChiC2HH` HLT line in the stripping lines, as the stripping lines were originally written to select $\chi_c(1P) \to h^+h^-$ events which had passed a generic low multiplicity event HLT line with less specific cuts than `Hlt2LowMultChiC2HH`.

Table 4.4: A summary of cuts in the `StrippingLowMultCEP_ChiC2KK_line` and `StrippingLowMultCEP_ChiC2PiPi_line` stripping lines.

Variable	Cut
Global cuts	
Number of tracks	< 6
Number of backward tracks	< 1
Number of long tracks	> 1
Daughter cuts	
p	$> 5.0\,\text{GeV}/c$
$p_\perp$	$> 100.0\,\text{MeV}/c$
Track χ^2/DoF	< 3.0
$\Delta \ln \mathcal{L}_{K-\pi}$ (`KK` line only)	> 0.0
Mother cuts	
p	$> 10.0\,\text{GeV}/c$
$p_\perp$	$< 5.0\,\text{GeV}/c$
m	$> 2850.0\,\text{MeV}/c^2$
m	$< 4500.0\,\text{MeV}/c^2$
DOCA	$< 0.5\,\text{mm}$
Vertex χ^2/DoF	< 15.0

4.5 Offline selection

$\chi_c(1P) \to h^+h^-$ CEP candidates selected by one of the stripping lines are processed using DaVinci, the LHCb physics analysis application. During the processing a momentum scale correction is applied to the daughter tracks. The calibration is determined from collision data by the LHCb Tracking and Alignment group using decays such as $B^+ \to J/\psi K^+$. The mass distribution of the resonance is fitted and the offset between the fitted mass and its known value used to calculate a correction factor to apply to the daughter track momenta. The momentum scale correction accounts for effects such as miscalibration of the magnetic field scale, and is parameterised by rapidity and transverse momentum. It improves the absolute mass scale and mass resolution by a few percent [98]. After the correction is applied, the decay vertex is refitted using the scaled track momenta, fitting the entire decay tree simultaneously [99].

After the vertex is refitted with the improved momentum scale calibration, an offline selection is applied which contains three main types of cut: (i) cuts on the mass of the mother particle and the pseudorapidities of the daughter candidates to define the region in which the cross section is measured, (ii) cuts on the event multiplicity based on the known multiplicity of the two body decays to be selected, to reject inelastic and other background processes, and (iii) cuts on the kinematic properties of the daughter tracks and reconstructed mother based on the expected values for signal decays as determined from the simulations described in Section 4.3.

All events are required to have a reconstructed mother mass between 3200 MeV/c^2 and 3600 MeV/c^2, sufficient to select the $\chi_{c0}(1P)$ at 3414.75 MeV/c^2 and $\chi_{c2}(1P)$ at 3556.20 MeV/c^2 whilst providing adequate sidebands for background studies in data. Additionally, the two daughter tracks are required to have pseudorapidity between 2.5 and 4.5. This corresponds to the subset of the LHCb acceptance, shown in Fig. 4.3, where simulation indicates the combined reconstruction, trigger and stripping efficiency is reasonably flat.

As discussed previously, in genuine CEP events the final state should contain only the direct daughters of the central exclusively produced state with no other tracks in the detector. Figure 4.4 shows the number of VELO and long tracks in simulated $\chi_c(1P) \to h^+h^-$ events which pass the trigger and stripping selections described in Sections 4.4.1 to 4.4.3. All the simulated $\chi_c(1P) \to h^+h^-$ events have only two daughters, however after processing more than two tracks may be reconstructed when non-physical ghost tracks are reconstructed. It is observed that in no mode is the fraction of events with more than two reconstructed VELO or long tracks greater than 1 %, and the offline selection requires exactly two VELO

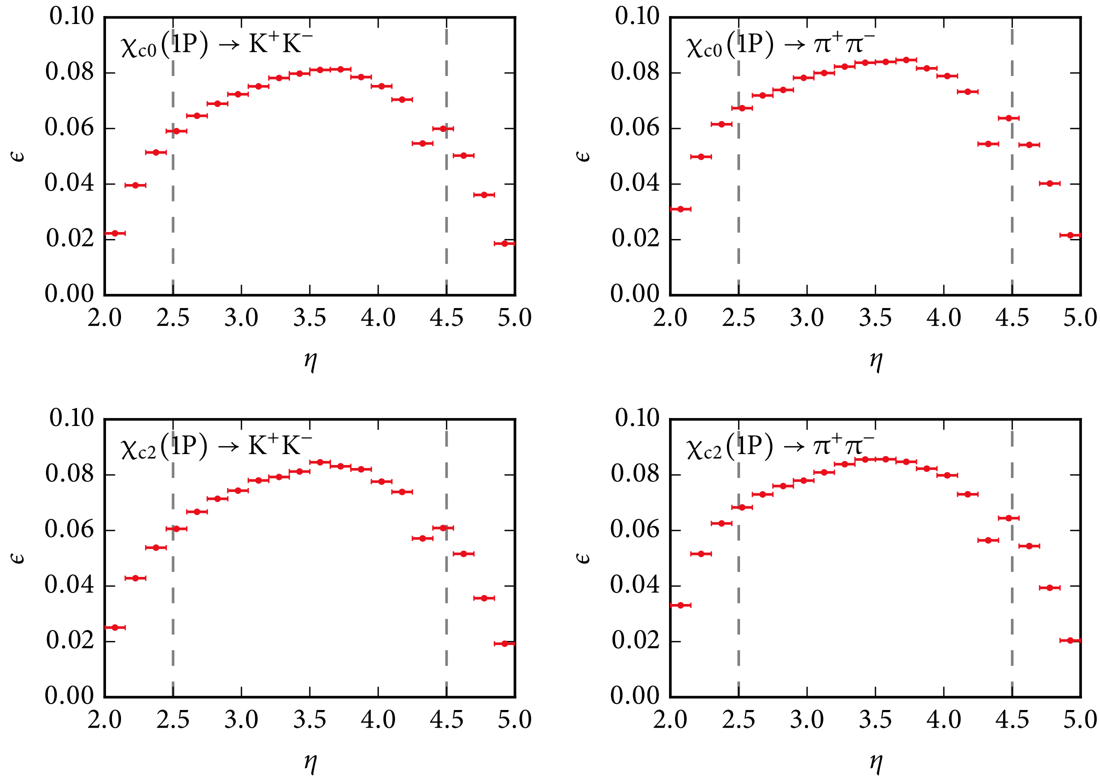

Figure 4.3: The combined reconstruction, trigger and stripping efficiency as a function of the daughter track pseudorapidity over the nominal detector acceptance, as determined from simulated events for $\chi_{c0}(1P) \to K^+K^-$ (top left), $\chi_{c0}(1P) \to \pi^+\pi^-$ (top right), $\chi_{c2}(1P) \to K^+K^-$ (bottom left) and $\chi_{c2}(1P) \to \pi^+\pi^-$ (bottom right) events. The grey dashed vertical lines indicate the range $2.5 < \eta < 4.5$ in which the pseudorapidity of daughter tracks is required to be.

tracks and two long tracks per event. The fraction of ghost tracks in the LHCb simulation has previously been validated against collision data, and good agreement found [82].

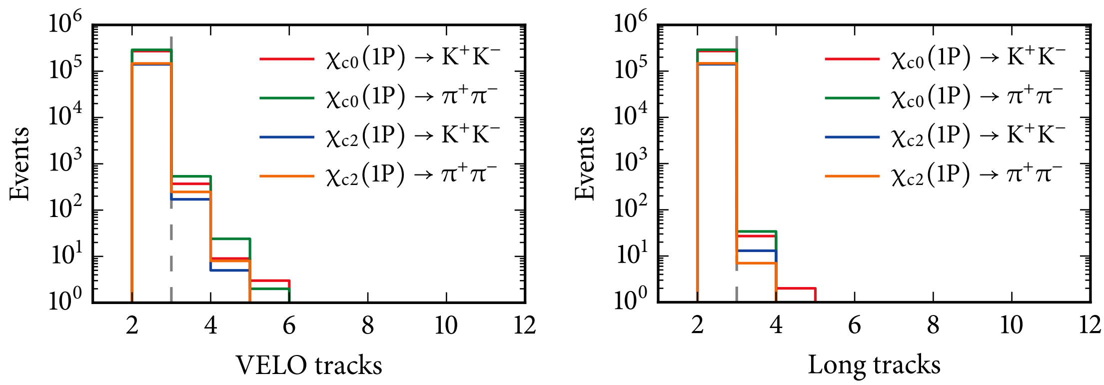

Figure 4.4: The number of reconstructed VELO tracks (left) and long tracks (right) in simulated $\chi_c(1P) \to h^+h^-$ events passing the trigger and stripping selections. Note that the y-axes have logarithmic scales. The grey dashed vertical lines indicate the cut of < 3 applied in the offline selection.

The final set of cuts in the offline selection concerns the kinematics of the daughter tracks and the reconstructed mother. Figure 4.5 shows the daughter track transverse momentum distributions for simulated events from Monte Carlo and collision events from real data, for events passing the trigger and stripping selections. With this selection, the sample of collision data is still dominated by background processes and these are isolated by only plotting events from the mass sidebands in the collision data shown in Fig. 4.6. The sidebands are defined as the regions further than three widths from the $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$ world average masses published by the Particle Data Group [97]. It is seen from Fig. 4.5 that there are a negligible number of daughter tracks in simulation with transverse momentum below $1.0\,\text{GeV}/c$, whereas the distribution (for background events) in collision data extends below this value. Therefore the offline selection requires the daughter track momenta to be greater than $1.0\,\text{GeV}/c$.

Similarly, the daughter $\Delta \ln \mathcal{L}_{K-\pi}$ distributions for simulated events and collision events in the mass sidebands are shown in Fig. 4.7. As already noted, the K^+K^- stripping line

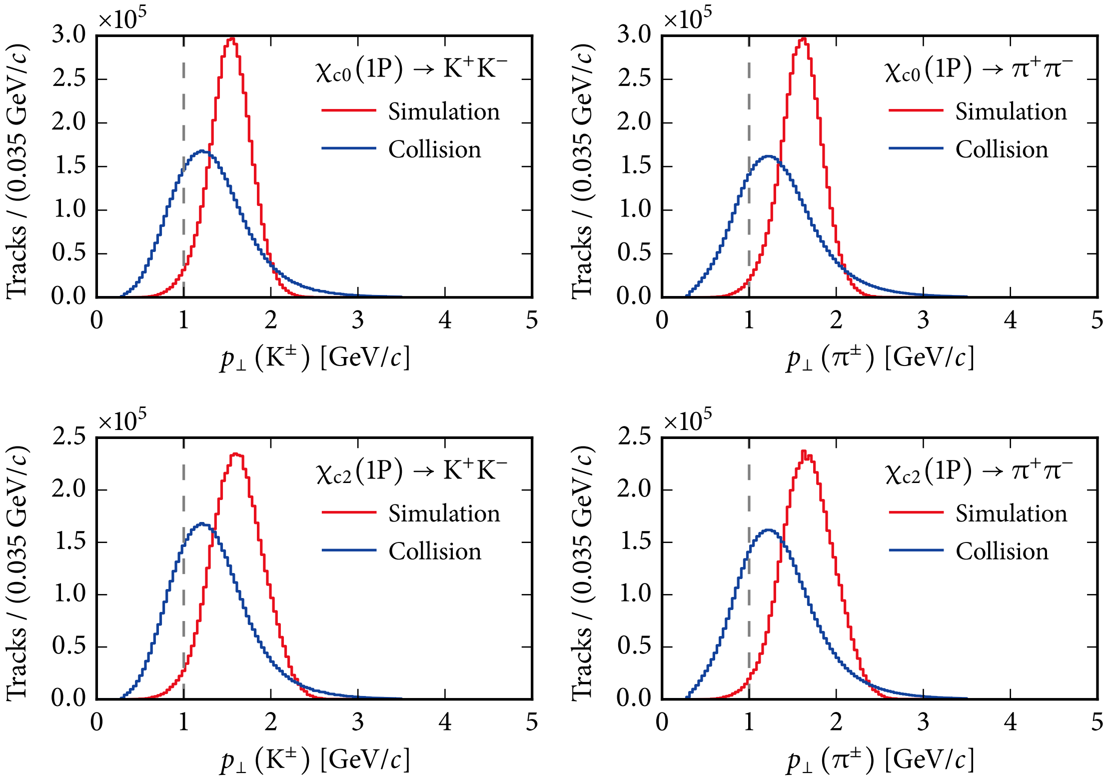

Figure 4.5: The daughter transverse momentum for simulated events for $\chi_{c0}(1P) \rightarrow K^+K^-$ (top left), $\chi_{c0}(1P) \rightarrow \pi^+\pi^-$ (top right), $\chi_{c2}(1P) \rightarrow K^+K^-$ (bottom left) and $\chi_{c2}(1P) \rightarrow \pi^+\pi^-$ (bottom right) events, plotted with the daughter track momentum in candidate $\chi_c(1P) \rightarrow K^+K^-$, $\chi_c(1P) \rightarrow \pi^+\pi^-$, $\chi_c(1P) \rightarrow K^+K^-$ and $\chi_c(1P) \rightarrow \pi^+\pi^-$ events respectively in collision sidebands. Each distribution is normalised to 5×10^6 events. The grey dashed vertical lines indicate the position of the 1.0 GeV/c cut applied in the offline selection.

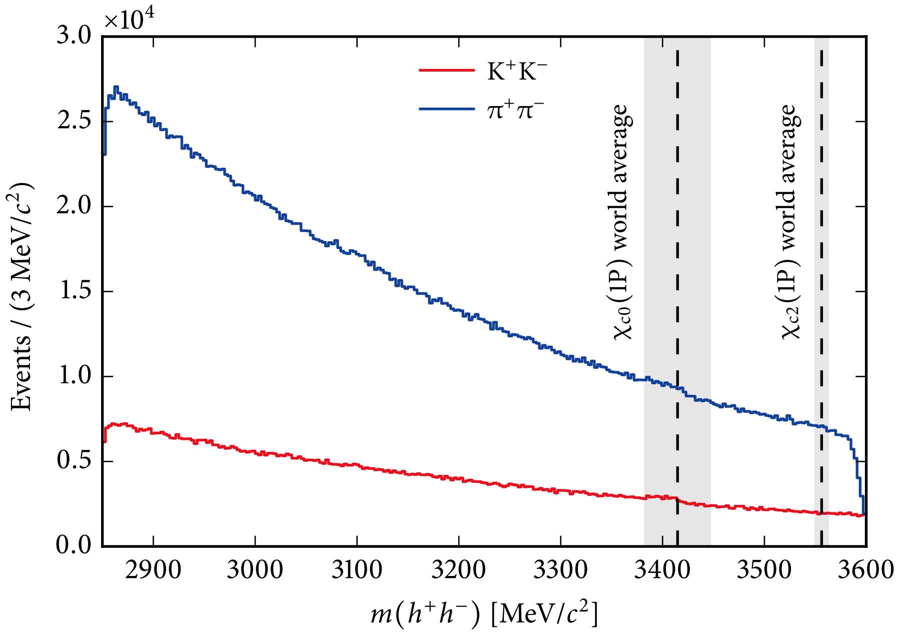

Figure 4.6: The mother mass distribution in collision data for K^+K^- and $\pi^+\pi^-$ events which pass the trigger and stripping selections. The world average masses of the $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$ are indicated by the black vertical dashed lines. The grey shaded regions correspond to masses within three widths of the world average masses. The sidebands are defined as the regions which are not shaded.

requires the daughter candidates to have $\Delta \ln \mathcal{L}_{K-\pi} > 0$. It is observed that the majority of $\chi_{c0}(1P) \to K^+K^-$ and $\chi_{c2}(1P) \to K^+K^-$ events in simulation have daughter $\Delta \ln \mathcal{L}_{K-\pi} > 20$, whilst in the collision data sidebands there is a tail below, and so a cut of $\Delta \ln \mathcal{L}_{K-\pi} > 20$ is applied to $\chi_c(1P) \to K^+K^-$ candidates in the offline selection. In the pion distributions it is seen that in simulation the distribution barely extends above zero, whereas in collision data sidebands there are a considerable number of events with $\Delta \ln \mathcal{L}_{K-\pi}$ greater than zero and hence a cut of $\Delta \ln \mathcal{L}_{K-\pi} < 0$ is applied to pion candidates in the offline selection.

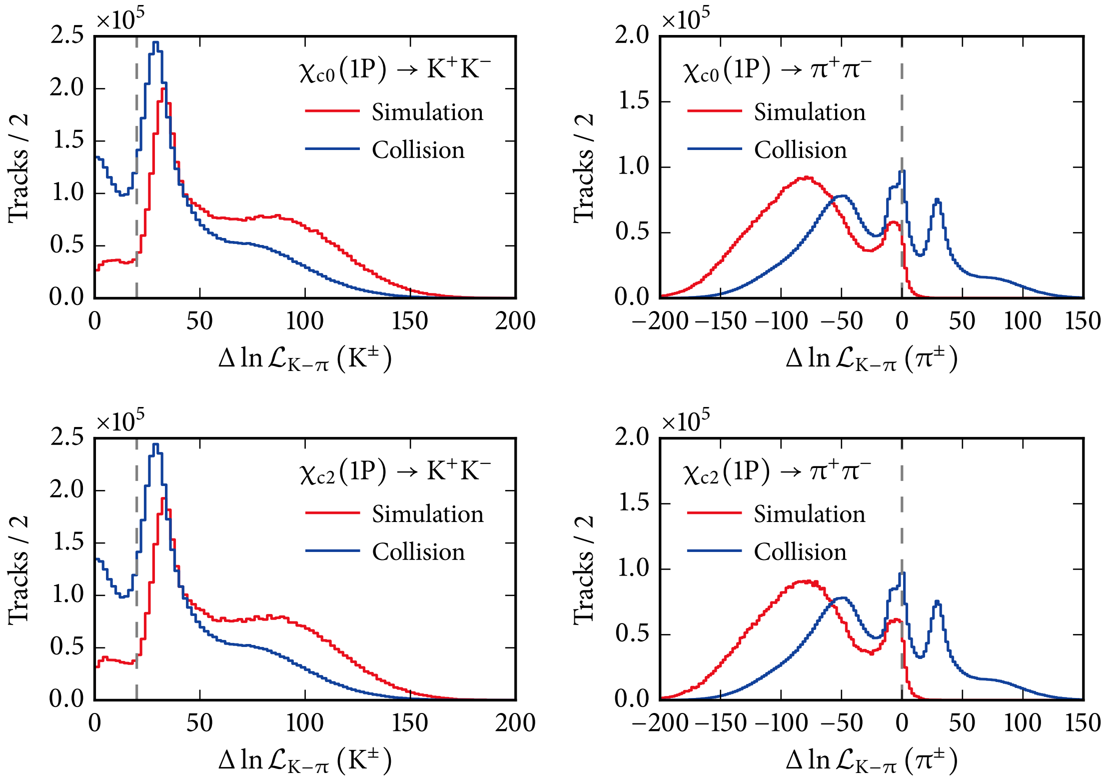

Figure 4.7: The daughter $\Delta \ln \mathcal{L}_{K-\pi}$ for simulated events for $\chi_{c0}(1P) \to K^+K^-$ (top left), $\chi_{c0}(1P) \to \pi^+\pi^-$ (top right), $\chi_{c2}(1P) \to K^+K^-$ (bottom left) and $\chi_{c2}(1P) \to \pi^+\pi^-$ (bottom right) events, plotted with the daughter $\Delta \ln \mathcal{L}_{K-\pi}$ in candidate $\chi_c(1P) \to K^+K^-$, $\chi_c(1P) \to \pi^+\pi^-$, $\chi_c(1P) \to K^+K^-$ and $\chi_c(1P) \to \pi^+\pi^-$ events respectively in collision sidebands. Each distribution is normalised to 5×10^6 events. The grey dashed vertical lines indicate the choice of cut value for each mode, described in the text.

Turning from the daughter tracks to the reconstructed mother, Fig. 4.8 shows the mother transverse momentum in simulated signal events and collision events in mass sidebands. It is seen that the distribution has dropped to a negligible level for both the $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$ in simulation by 1.5 GeV/c, however in the collision data there is a tail out to approximately double this value. As discussed in Section 4.1, proton dissociation is more probable the greater the transverse momentum of the central system, hence the mother transverse momentum is required to be less than 1.5 GeV/c.

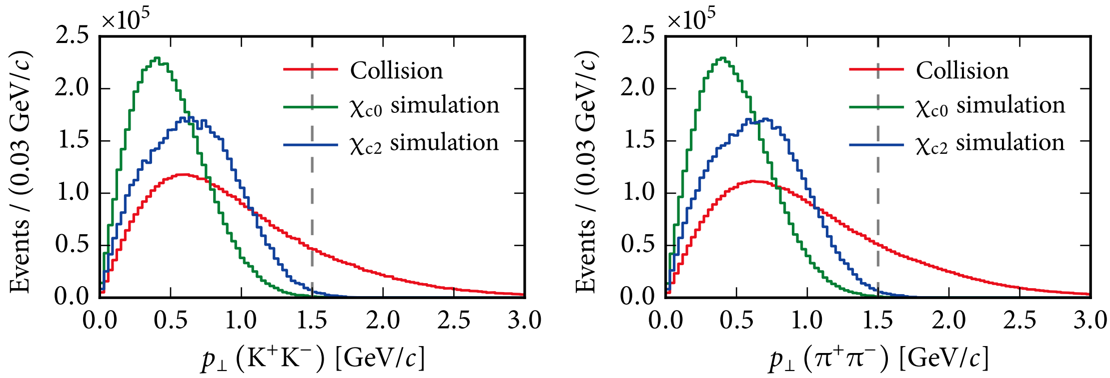

Figure 4.8: The reconstructed mother transverse momentum in mass sideband collision events, simulated $\chi_{c0}(1P)$ events, and simulated $\chi_{c2}(1P)$ events, for the K^+K^- (left) and $\pi^+\pi^-$ (right) final states. Only events which pass the trigger and stripping selections are plotted. The grey dashed vertical lines indicate the position of the 1.5 GeV/c cut applied in the offline selection.

Table 4.5 gives a summary of the cuts comprising the offline selection. After application of these cuts, there are 40,588 selected $\chi_c(1P) \to K^+K^-$ candidates and 65,019 selected $\chi_c(1P) \to \pi^+\pi^-$ candidates in the 2012 data set.

4.6 Selection efficiencies and correction factors

Selection inefficiencies arise when genuine $\chi_c(1P) \to h^+h^-$ CEP events fail any of the selection cuts applied in the L0 trigger, HLT, stripping, and offline selection described above, or are not reconstructed. These efficiencies must be determined so that signal events that

Table 4.5: A summary of cuts in the offline selection.

Variable	Cut
$m(h^+h^-)$	$> 3200.0\,\mathrm{MeV}/c^2$
$m(h^+h^-)$	$< 3600.0\,\mathrm{MeV}/c^2$
$\eta(h^\pm)$	> 2.5
$\eta(h^\pm)$	< 4.5
Number of VELO tracks	$= 2$
Number of long tracks	$= 2$
$p_\perp(h^\pm)$	$> 1.0\,\mathrm{GeV}/c$
$\Delta \ln \mathcal{L}_{\mathrm{K}-\pi}(\mathrm{K}^\pm)$	> 20
$\Delta \ln \mathcal{L}_{\mathrm{K}-\pi}(\pi^\pm)$	< 0
$p_\perp(h^+h^-)$	$< 1.0\,\mathrm{GeV}/c$

are not reconstructed or are excluded by a selection cut are accounted for when calculating the cross section from the measured signal yield.

The selection efficiencies for each stage of the trigger, stripping and offline processing are evaluated using simulated signal events, generated as described in Section 4.3. Only those events which have both daughters with pseudorapidity in the range in which the efficiency is seen in Section 4.5 to be approximately flat, $2.5 < \eta < 4.5$, are used in determining the efficiencies, as this defines the region in which the measurement is carried out.

Each individual efficiency is calculated as the fraction of events evaluated by that stage which pass the selection cuts of that stage. For example, the L0 efficiency is the number of events with daughter pseudorapidities in the range $2.5 < \eta < 4.5$ which pass the `DiHadron,lowMult` line at L0, divided by the number of events generated with daughter pseudorapidities in the same range. The stripping efficiency is the number of events with daughter pseudorapidities in the range $2.5 < \eta < 4.5$ which pass the `StrippingLowMultCEP_ChiC2KK_line` (for the $\mathrm{K}^+\mathrm{K}^-$ final state) or `StrippingLowMultCEP_ChiC2PiPi_line` line (for the $\pi^+\pi^-$ final state) in the stripping, divided by the number of events in the same range which pass the `Hlt2LowMultChiC2HH` line in HLT2. The individual efficiencies factorise such that the total selection efficiency is given by their product.

It is important however that corrections are made to account for the differences between the simulation and signal events in collision data, and these corrections are determined below. This is the case for the L0 trigger efficiency, and those stages of the selection which depend on the RICH response by cutting on $\Delta \ln \mathcal{L}_{K-\pi}$, as these are imperfectly modelled in the LHCb simulation.

4.6.1 Level-0 trigger efficiency correction

As described in Section 4.4.1, the `DiHadron,lowMult` L0 line requires hadron $E_\perp > 500$ MeV, SPD multiplicity < 10, and PU multiplicity < 3. The SPD multiplicity distribution peaks at fewer hits in simulated data than selected collision data, due to noisy channels and cross-talk between cells which is not well modelled. This results in an inaccurate determination of the L0 efficiency from simulation, and therefore a correction is made using collision data.

The SPD multiplicity distributions of simulated $\chi_{c0}(1P) \to K^+K^-$ events and selected $\chi_c(1P) \to K^+K^-$ events in collision data within three widths of the $\chi_{c0}(1P)$ mass are fitted with a Bukin function. This is an asymmetric peaking function which is a modified convolution of a Gaussian function and an exponential function [100]. The distributions are fitted up to nine hits, the maximum multiplicity selected by the `DiHadron,lowMult` cut, using an unbinned maximum likelihood fit, and these fits are shown in Fig. 4.9.

The fitted parameters are used to extend the Bukin function beyond ten hits, modelling the distribution for events failing the L0 cut. The distribution for selected events in collision data includes background events which are not genuine $\chi_c(1P)$ decays, however they are K^+K^- events from the signal region with the same selected kinematics as genuine CEP events, and therefore should have the same trigger efficiency. The efficiency of the cut is determined as the fraction of the area under the Bukin function lying below 10 hits. This is 99.6 % for simulated events and 89.1 % for collision events. A correction of −10.5 % is therefore applied to the L0 efficiencies determined from simulation.

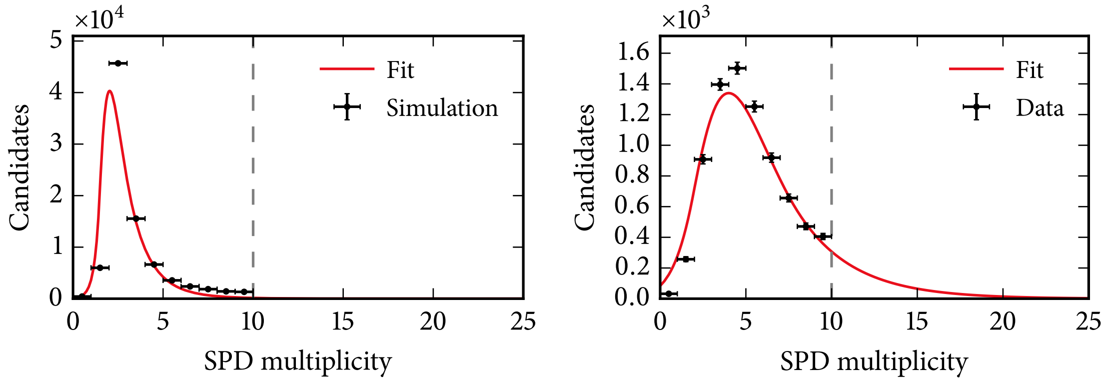

Figure 4.9: Fits to the SPD multiplicity distribution in simulated $\chi_{c0}(1P) \to K^+K^-$ events (left) and selected $\chi_c(1P) \to K^+K^-$ events in collision data within three widths of the $\chi_{c0}(1P)$ mass (right). The grey dashed vertical line shows the cut of < 10 applied in the `DiHadron,lowMult` L0 line.

The same procedure is repeated for the hadron $E_\perp$, fitting the distributions of simulated $\chi_{c0}(1P) \to K^+K^-$ events and selected $\chi_c(1P) \to K^+K^-$ events in collision data within three widths of the $\chi_{c0}(1P)$ mass, as shown in Fig. 4.10. The distributions are found to be best modelled by a Novosibirsk function, an asymmetric peaking function with three parameters: peak location, peak width, and a tail parameter.

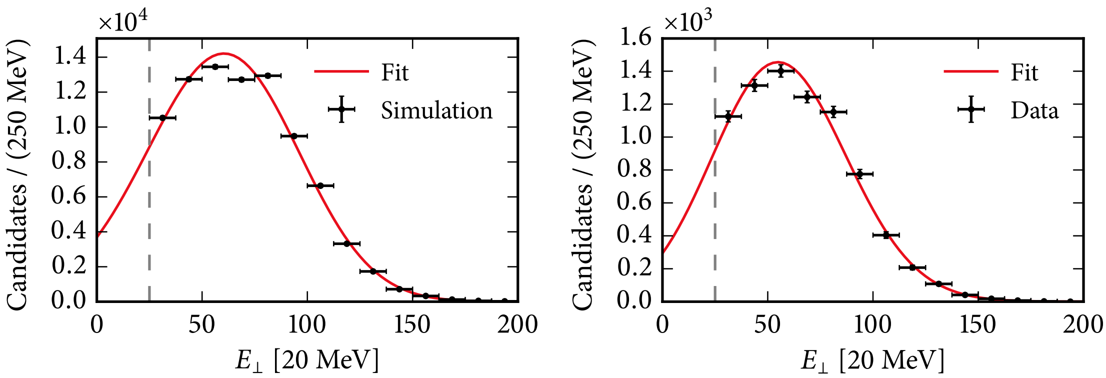

Figure 4.10: Fits to the L0 hadron $E_\perp$ distribution in simulated $\chi_{c0}(1P) \to K^+K^-$ events (left) and selected $\chi_{c0}(1P) \to K^+K^-$ events in collision data within three widths of the $\chi_{c0}(1P)$ mass (right). The grey dashed vertical line shows the cut of > 500 MeV applied in the `DiHadron,lowMult` L0 line.

The hadron $E_\perp$ is better modelled in simulation than the SPD multiplicity. The efficiency, calculated from the fraction of the Novosibirsk function which lies above 500 MeV, is 87.3 % for simulated events and 87.0 % for collision events, and therefore a further correction of −0.3 % is applied to the L0 efficiencies determined from simulation.

The PU multiplicity cut is 100 % efficient for simulated events, with no events having any hits in this subdetector. The distribution for selected events in collision data passing the trigger is also strongly peaked at zero hits, with a small fraction of events with one or two hits. These are assumed to be predominantly from background events in the selected data set, rather than from detector noise. This assumption is validated by inspecting data recorded in 2015 for events without collisions (triggers to select such events were not operational during 2012 data taking), which shows very few events with one or two PU hits and a negligible number with more than two hits. Therefore the efficiency taken from simulation is deemed valid, and no correction to the L0 trigger efficiency for the PU multiplicity is applied.

The net correction applied to the L0 efficiencies, accounting for the hadron $E_\perp$ and SPD multiplicity cuts, is −10.8 %. The correction determined for $\chi_{c0}(1P) \to K^+K^-$ events is also applied to the other three signal modes, as the hadron $E_\perp$ and SPD multiplicity distributions are compatible across the modes. The corrections described here introduce systematic uncertainties from the fitting and extrapolation procedure, and these uncertainties will be evaluated in Section 5.5.

4.6.2 Particle identification efficiency correction

The particle identification (PID) efficiency of the RICH detectors shows a strong dependence on the track multiplicity of an event, with the efficiency degrading with increased occupancy. As discussed previously, the multiplicity of selected $\chi_c(1P) \to h^+h^-$ candidate events is much lower than typical inelastic events. Furthermore, the RICH response is imperfectly modelled in the LHCb simulation. Therefore the efficiencies determined from simulation for the $\Delta \ln \mathcal{L}_{K-\pi}$ cuts applied in the selection are corrected using collision data.

Pion and kaon calibration samples are produced by the LHCb Particle Identification group by selecting exclusive $D^{*+} \to D^0\pi^+$ decays, where the D^0 decays to $K^-\pi^+$, for calibrating the PID performance. The kinematics of this decay allow high purity samples to be selected without using information from the RICH detectors, and the high production cross section provides large samples which minimises statistical uncertainties [101]. The calibration samples are used to evaluate the efficiency of a PID cut by selecting calibration events with similar kinematics to the signal candidates in an analysis, and then measuring the fraction of calibration events which pass the PID cut applied in the signal selection.

Whilst subsets of the calibration samples with similar pseudorapidity and momentum distributions to $\chi_c(1P) \to h^+h^-$ candidates can be selected, the $D^{*+} \to D^0\pi^+$ decays occur in events with much higher track multiplicity than CEP events. Direct use of calibration events from higher multiplicity events would lead to efficiency biases due to the aforementioned dependency on track multiplicity.

Instead, calibration events with similar kinematic distributions to $\chi_c(1P) \to h^+h^-$ candidates are selected by requiring the kaon or pion pseudorapidity to be within $2.5 < \eta < 4.5$ and the transverse momentum to be within $1 < p_\perp < 2.2$ GeV/c. The efficiency of the PID cut applied in the offline selection is then measured in bins of track multiplicity from 15 to 200 tracks. A linear dependence of PID efficiency on track multiplicity is observed, and is fitted using the method of least squares as shown in Fig. 4.11.

The kaon identification efficiency ϵ_K for a cut of $\Delta \ln \mathcal{L}_{K-\pi} > 20$ in an event with track multiplicity n is found to be $\epsilon_K = (-0.1734 \pm 0.0004)n + (91.94 \pm 0.06)$. The pion identification efficiency ϵ_π for a cut of $\Delta \ln \mathcal{L}_{K-\pi} < 0$ is $\epsilon_\pi = (-0.0318 \pm 0.0002)n + (98.14 \pm 0.03)$. The efficiencies for two track events are determined by linear extrapolation, and the values obtained

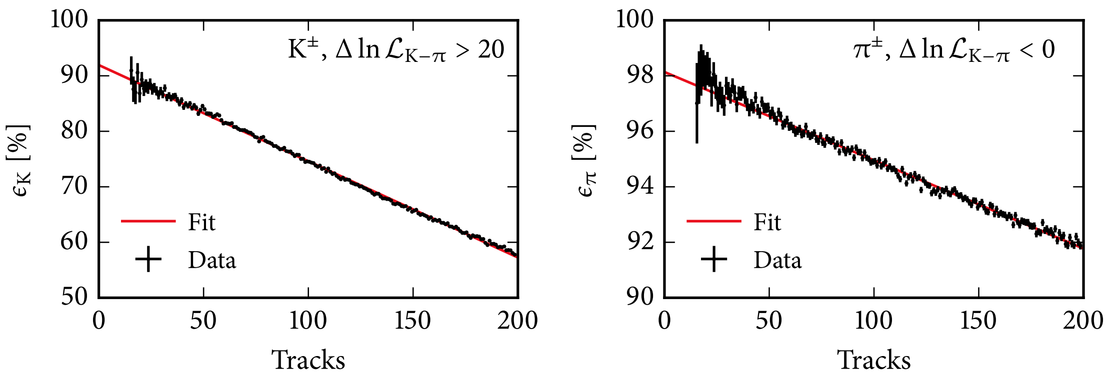

Figure 4.11: The identification efficiencies for kaons (left) and pions (right) with kinematics in the same range as $\chi_c(1P) \to h^+h^-$ daughter candidates, for PID cuts of $\Delta \ln \mathcal{L}_{K-\pi} > 20$ and $\Delta \ln \mathcal{L}_{K-\pi} < 0$ respectively. A linear fit is made as a function of track multiplicity.

are listed in Table 4.6. The efficiencies taken from simulation are therefore corrected to the values taken from the fit to the calibration events.

Table 4.6: Daughter particle identification efficiencies for events with two tracks, from simulated events and extrapolated from higher multiplicity calibration events.

Particle	Cut	Efficiency From simulation [%]	Efficiency From calibration data [%]
$K^\pm$	$\Delta \ln \mathcal{L}_{K-\pi} > 20$	96.2 ± 0.4	91.59 ± 0.06
$\pi^\pm$	$\Delta \ln \mathcal{L}_{K-\pi} < 0$	98.9 ± 0.4	98.07 ± 0.03

In addition to these PID cuts, the HLT, stripping, and offline selections cut on variables determined from the trajectories of the two daughter hadrons, such as momentum. The tracking performance must therefore be well modelled in simulation to ensure the efficiencies of these stages of the selection are accurately determined. The simulated hit efficiency and hit resolution in the trackers have been validated using collision data. A tag-and-probe method was used in which hits were removed from the track and the residual between the hit position and the extrapolated trajectory of the track calculated. Good agreement between the simulation and collision data was observed [69].

The selection efficiencies calculated from simulation with the data-driven corrections described above applied are listed for each decay mode in Table 4.7. The L0 efficiency includes the 25 % pre-scale. The uncertainties quoted in Table 4.7 are statistical, and account only for the finite size of the samples of simulated events.

Table 4.7: The trigger, stripping, and offline selection efficiencies for $\chi_c(1P) \to h^+h^-$ events with both daughter hadrons within $2.4 < \eta < 4.5$, determined from simulation with data-driven corrections. The L0 efficiency includes the pre-scale of 25 %. The offline selection efficiency is lower for the K^+K^- modes than the $\pi^+\pi^-$ modes due to tighter PID requirements for the former. The quoted uncertainties are statistical only. The total selection efficiency is the product of the individual efficiencies.

Mode	L0 [%]	HLT1 [%]	HLT2 [%]	Stripping [%]	Offline selection [%]	Total [%]
$\chi_{c0}(1P) \to K^+K^-$	16.17 ± 0.05	100	71.85 ± 0.25	99.14 ± 0.34	83.29 ± 0.31	9.60 ± 0.06
$\chi_{c0}(1P) \to \pi^+\pi^-$	16.60 ± 0.05	100	73.54 ± 0.25	99.79 ± 0.34	91.91 ± 0.33	11.20 ± 0.08
$\chi_{c2}(1P) \to K^+K^-$	16.46 ± 0.07	100	72.28 ± 0.35	99.0 ± 0.5	72.9 ± 0.4	8.59 ± 0.08
$\chi_{c2}(1P) \to \pi^+\pi^-$	16.81 ± 0.07	100	73.4 ± 0.4	99.8 ± 0.5	81.5 ± 0.4	10.04 ± 0.10

5

Exclusive $\chi_c(1P) \rightarrow h^+h^-$ cross section measurement

In this chapter the fit to the invariant mass distribution of selected $\chi_c(1P) \rightarrow h^+h^-$ candidates is presented. The exclusive purity—the fraction of the fitted yields contributed by genuine CEP events rather than inelastic background events—is evaluated, and the effective integrated luminosity for bunch crossings without multiple interactions is determined. Systematic uncertainties are assigned for each stage of the analysis, and the product of the cross section and branching fraction is measured for $\chi_{c0}(1P) \rightarrow K^+K^-$ and $\chi_{c0}(1P) \rightarrow \pi^+\pi^-$, with upper limits set for $\chi_{c2}(1P) \rightarrow K^+K^-$ and $\chi_{c2}(1P) \rightarrow \pi^+\pi^-$.

5.1 Invariant mass fit

The yield of central exclusively produced $\chi_c(1P)$ candidates is measured by fitting the invariant mass distribution of candidates passing the selection described in Chapter 4 using an unbinned extended maximum likelihood fit. Two independent fits are carried out, one for $\chi_c(1P) \rightarrow K^+K^-$ candidates and one for $\chi_c(1P) \rightarrow \pi^+\pi^-$ candidates. In each fit, the invariant mass distribution is modelled with a probability density function (PDF) which is the sum of three components.

The first component is an exponential function which models the non-resonant background, seen in Fig. 4.6 to fall smoothly with increasing mass. The exponential component has two parameters, the slope c_{exp} and yield N_{exp}, both of which are floated to find the values which maximise the likelihood of the fit.

The other two components of the PDF model the $\chi_{c0}(1\mathrm{P}) \to h^+h^-$ and $\chi_{c2}(1\mathrm{P}) \to h^+h^-$ signal decays. These are each modelled with a Voigt function, which is the convolution of a Breit-Wigner function and a Gaussian function with a common mean. The choice of the Voigt function for the signal decays is physically motivated, as the Breit-Wigner function describes the resonant decay whilst the Gaussian function models the finite resolution of the detector. The $\chi_{c0}(1\mathrm{P})$ and $\chi_{c2}(1\mathrm{P})$ PDF components each have four parameters, which are the mean m_0, the width of the Gaussian function σ, the width of the Breit-Wigner function Γ, and the yield N.

Upon visual inspection of the invariant mass distributions, peaks corresponding to the $\chi_{c0}(1\mathrm{P}) \to \mathrm{K}^+\mathrm{K}^-$ and $\chi_{c0}(1\mathrm{P}) \to \pi^+\pi^-$ decays are visible, whereas no significant excess at the $\chi_{c2}(1\mathrm{P})$ mass is visible. Therefore the mean of the $\chi_{c0}(1\mathrm{P})$ component $m_{0,\chi_{c0}(1\mathrm{P})}$ is floated in the fit whilst the mean of the $\chi_{c2}(1\mathrm{P})$ component $m_{0,\chi_{c2}(1\mathrm{P})}$ is fixed to $3556.20\,\mathrm{MeV}/c^2$, the world average $\chi_{c2}(1\mathrm{P})$ mass published by the Particle Data Group (PDG) [97].

As no significant excess is observed at the $\chi_{c2}(1\mathrm{P})$ mass in either the $\mathrm{K}^+\mathrm{K}^-$ or $\pi^+\pi^-$ invariant mass distributions, the Gaussian and Breit-Wigner widths of the $\chi_{c2}(1\mathrm{P})$ fit component cannot be floated and so are fixed. The Gaussian and Breit-Wigner widths of the $\chi_{c0}(1\mathrm{P})$ components are also fixed, due to structure in the mass distribution below the $\chi_{c0}(1\mathrm{P})$ peaks which will be discussed in Section 5.2. The Breit-Wigner widths of the $\chi_{c0}(1\mathrm{P})$ and $\chi_{c2}(1\mathrm{P})$ components are fixed to the PDG world average value, $10.5\,\mathrm{MeV}/c^2$ for the $\chi_{c0}(1\mathrm{P})$ and $1.95\,\mathrm{MeV}/c^2$ for the $\chi_{c2}(1\mathrm{P})$ [97]. The Gaussian widths are fixed in the fit to collision data, to values which are determined from a fit to simulated signal events in which they are allowed to float.

To determine the expected Gaussian widths to fix in the fit to collision data, a single Voigt function is fitted to each of the simulated signal modes, in independent unbinned maximum likelihood fits. The mean, Gaussian width, and yield are floated whilst the Breit-Wigner width is fixed to the PDG world average value (as in the fit to collision data). These fits are shown in Fig. 5.1, and the values of the fitted parameters listed in Table 5.1.

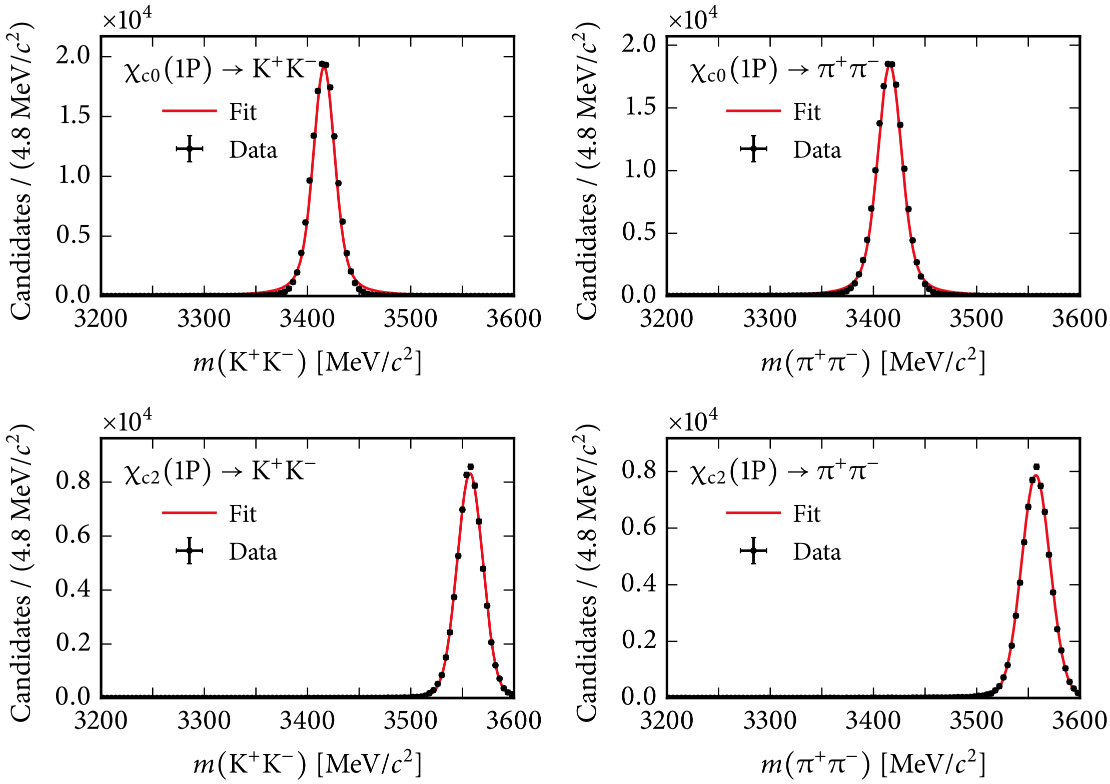

Figure 5.1: Fits to the invariant mass distribution of simulated $\chi_{c0}(1P) \to K^+K^-$ (top left), $\chi_{c0}(1P) \to \pi^+\pi^-$ (top right), $\chi_{c2}(1P) \to K^+K^-$ (bottom left), and $\chi_{c2}(1P) \to \pi^+\pi^-$ (bottom right) events, used to determine the expected values of the Gaussian width of each mode. The Breit-Wigner width is fixed to the PDG world average.

Table 5.1: The values of the parameters of the fits to the invariant mass distributions in simulated data, where the Breit-Wigner width Γ is fixed to the PDG world average.

Mode	m_0 [MeV/c^2]	σ [MeV/c^2]	Γ [MeV/c^2]
$\chi_{c0}(1P) \to K^+K^-$	3415.98 ± 0.03	7.97 ± 0.03	10.5
$\chi_{c0}(1P) \to \pi^+\pi^-$	3415.92 ± 0.04	9.22 ± 0.04	10.5
$\chi_{c2}(1P) \to K^+K^-$	3557.34 ± 0.05	12.04 ± 0.04	1.95
$\chi_{c2}(1P) \to \pi^+\pi^-$	3557.31 ± 0.06	13.49 ± 0.04	1.95

Having determined their expected values from fitting simulated data, the Gaussian widths are fixed to these values in the fits to collision data, which are shown in Fig. 5.2. The fits are carried out in the mass range from 3200 MeV/c^2 to 3586 MeV/c^2, to avoid edge effects near the 3600 MeV/c^2 cut in the HLT. The values of the fitted parameters are listed in Table 5.2.

Table 5.2: The values of the parameters of the fits to the invariant mass distributions in collision data, with widths fixed to the expected values determined from the fits to simulated data, and the $\chi_{c2}(1P)$ mass fixed to the world average value. Fixed parameters are given by values without an associated uncertainty.

Parameter	Units	$\chi_c(1P) \to K^+K^-$	$\chi_c(1P) \to \pi^+\pi^-$
$m_{0,\chi_{c0}(1P)}$	MeV/c^2	3398.7 ± 1.3	3402.3 ± 1.9
$\sigma_{\chi_{c0}(1P)}$	MeV/c^2	7.970	9.22
$\Gamma_{\chi_{c0}(1P)}$	MeV/c^2	10.5	10.5
$N_{\chi_{c0}(1P)}$	-	1039 ± 85	1041 ± 112
$m_{0,\chi_{c2}(1P)}$	MeV/c^2	3556.2	3556.2
$\sigma_{\chi_{c2}(1P)}$	MeV/c^2	12.04	13.49
$\Gamma_{\chi_{c2}(1P)}$	MeV/c^2	1.95	1.95
$N_{\chi_{c2}(1P)}$	-	120 ± 68	187 ± 91
c_{exp}	-	−0.00235 ± 0.00006	−0.00249 ± 0.00004
N_{exp}	-	$(3.941 \pm 0.023) \times 10^4$	$(6.433 \pm 0.030) \times 10^4$

In Fig. 5.2 it is seen that there is additional structure around the $\chi_{c0}(1P)$ peaks which is not well modelled by the fitted $\chi_{c0}(1P)$ shape alone. Apparent in the pull distributions, most clearly for the $\chi_c(1P) \to K^+K^-$ but to a lesser degree also for the $\chi_c(1P) \to \pi^+\pi^-$, is a shoulder at approximately 3360 MeV/c^2 and an accompanying undershoot at approximately 3450 MeV/c^2, as the slope of the exponential component is drawn up to account for the shoulder below the peak. Additionally, a bias in the $\chi_{c0}(1P)$ mass is observed, with fitted values of (3398.7 ± 1.3) MeV/c^2 in the $\chi_{c0}(1P) \to K^+K^-$ fit and (3402.3 ± 1.9) MeV/c^2 in the $\chi_{c0}(1P) \to \pi^+\pi^-$ fit, compared to the world average value of (3414.75 ± 0.31) MeV/c^2.

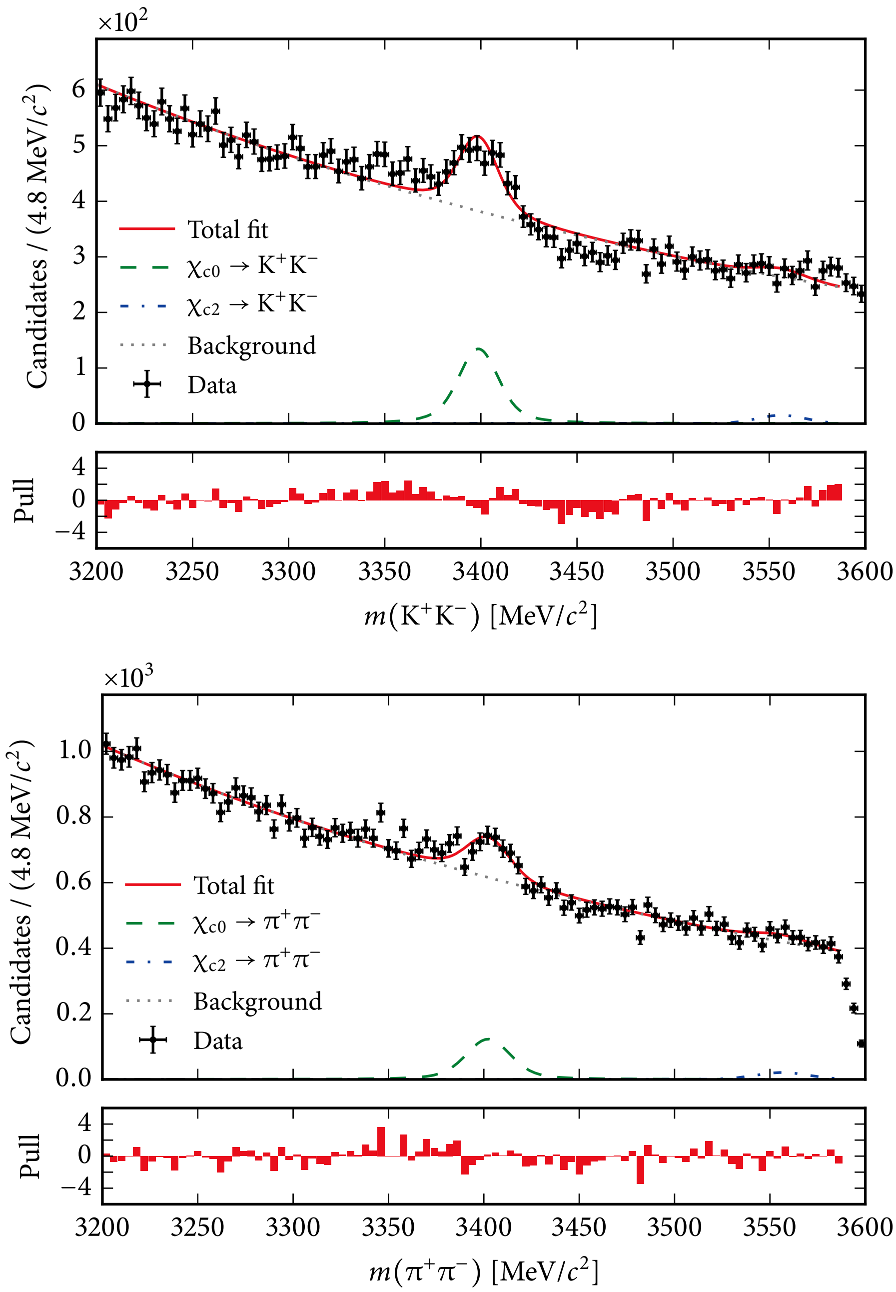

Figure 5.2: Fits to invariant mass distributions for $\chi_c(1P) \to K^+K^-$ (top) and $\chi_c(1P) \to \pi^+\pi^-$ (bottom) candidates in collision data, with widths fixed to the expected values determined from the fits to simulated data, and the $\chi_{c2}(1P)$ mass fixed to the world average value.

5.2 Investigation of low mass structure

This structure in the mass distribution above and below the $\chi_{c0}(1P)$ peak, and the bias in the fitted $\chi_{c0}(1P)$ mass, can be interpreted as originating in an excess of selected candidates with an invariant mass of approximately 3360 MeV/c^2. This interpretation is investigated by adding an additional component to the fit. A low mass component (LMC) is added, described by a Gaussian function with mean $m_{0,\mathrm{LMC}}$, width σ_{LMC}, and yield N_{LMC}. All parameters are allowed to float to the values maximising the fit likelihood. The other three fit components describing the non-resonant background, $\chi_{c0}(1P)$, and $\chi_{c2}(1P)$ respectively, remain unchanged. Fits to the invariant mass distributions of the selected $\chi_c(1P) \to K^+K^-$ and $\chi_c(1P) \to \pi^+\pi^-$ candidates including this low mass component are shown in Fig. 5.3, and the values of the fitted parameters are listed in Table 5.3.

Table 5.3: The values of the parameters of the fits to the invariant mass distributions in collision data, with widths fixed to the expected values determined from fits to simulated data, the $\chi_{c2}(1P)$ mass fixed to the world average value, and an additional Gaussian low mass fit component. Fixed parameters are given by values without an associated uncertainty.

Parameter	Units	$\chi_c(1P) \to K^+K^-$	$\chi_c(1P) \to \pi^+\pi^-$
$m_{0,\chi_{c0}(1P)}$	MeV/c^2	3401.6 ± 1.6	3405.0 ± 1.9
$\sigma_{\chi_{c0}(1P)}$	MeV/c^2	7.970	9.22
$\Gamma_{\chi_{c0}(1P)}$	MeV/c^2	10.5	10.5
$N_{\chi_{c0}(1P)}$	-	905 ± 117	960 ± 163
$m_{0,\chi_{c2}(1P)}$	MeV/c^2	3556.2	3556.2
$\sigma_{\chi_{c2}(1P)}$	MeV/c^2	12.04	13.49
$\Gamma_{\chi_{c2}(1P)}$	MeV/c^2	1.95	1.95
$N_{\chi_{c2}(1P)}$	-	195 ± 69	264 ± 93
$m_{0,\mathrm{LMC}}$	MeV/c^2	3357.3 ± 6.8	3362.1 ± 7.9
σ_{LMC}	MeV/c^2	30.2 ± 5.2	23.9 ± 6.0
N_{LMC}	-	904 ± 170	704 ± 203
c_{exp}	-	-0.00237 ± 0.00006	-0.00250 ± 0.00005
N_{exp}	-	$(3.854 \pm 0.028) \times 10^4$	$(6.361 \pm 0.035) \times 10^4$

An improvement in the pull distributions is immediately apparent, with the excess of candidates below the $\chi_{c0}(1P)$ peak well fitted by the low mass component in both the

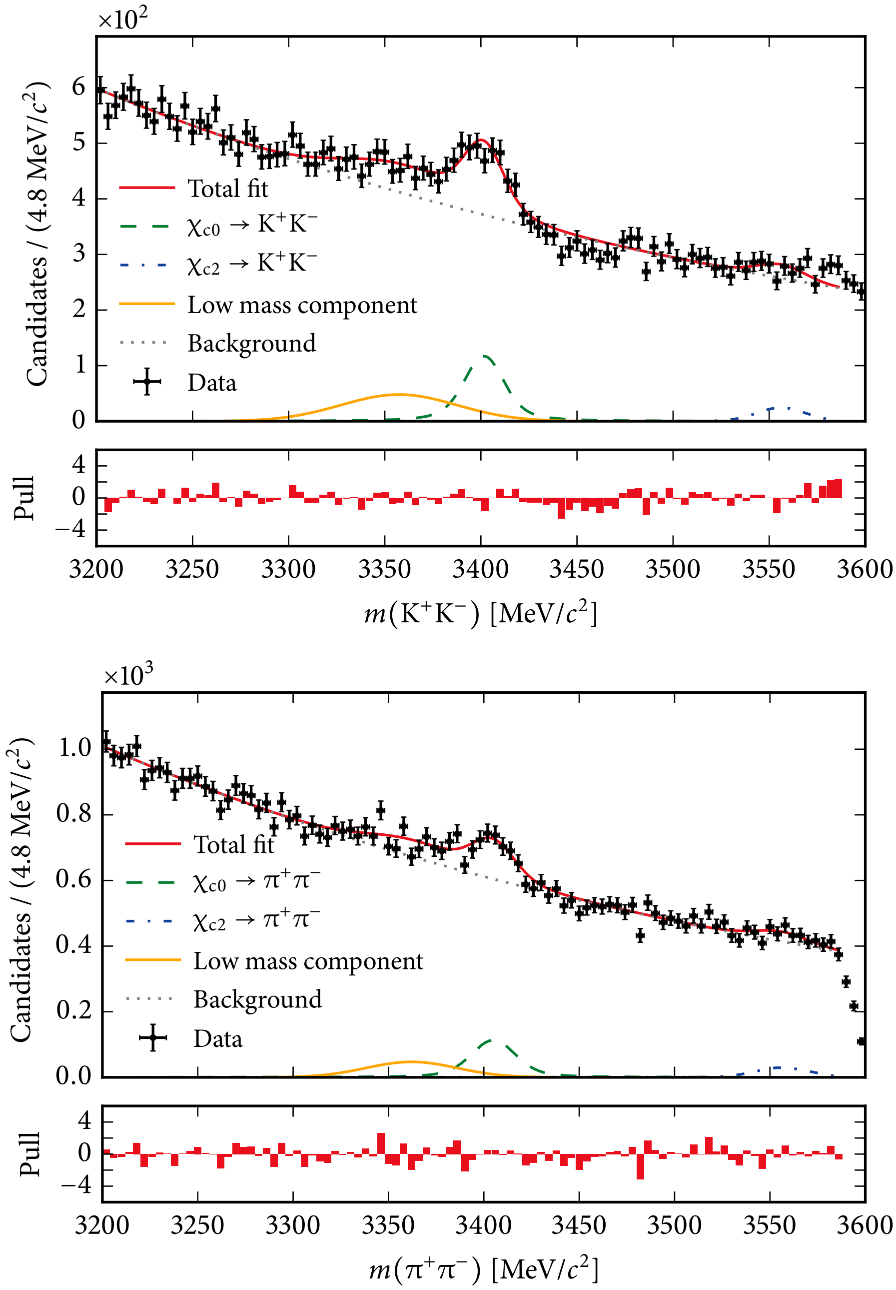

Figure 5.3: Fits to invariant mass distributions for $\chi_c(1P) \to K^+K^-$ (top) and $\chi_c(1P) \to \pi^+\pi^-$ (bottom) candidates in collision data, with widths fixed to the expected values determined from fits to simulated data, the $\chi_{c2}(1P)$ mass fixed to the world average value, and an additional Gaussian low mass fit component.

$\chi_c(1P) \to K^+K^-$ and $\chi_c(1P) \to \pi^+\pi^-$ fits. There is still an undershoot around 3450 MeV/c^2 in the $\chi_c(1P) \to K^+K^-$ distribution, however it is markedly reduced compared to the fit without a low mass component, and is indistinguishable in the $\chi_c(1P) \to \pi^+\pi^-$ distribution. This supports the hypothesis that the origin of the undershoot is the low mass shoulder.

With the addition of the low mass component, the fitted $\chi_{c0}(1P)$ mass increases by 3 MeV/c^2 in both modes. A bias remains, however its magnitude is reduced by 20 % by adding this simple additional component to the fit. This suggests the low mass excess and the $\chi_{c0}(1P)$ mass bias have a common origin in the shoulder below the $\chi_{c0}(1P)$ peak. This effect can be understood as the fit preferring to place the $\chi_{c0}(1P)$ mass at a lower value to absorb some of the low mass excess which otherwise cannot be modelled without additional components.

It is therefore desirable to identify the physical origin of the low mass excess, by determining the decay, or decays, which contribute to it. Identifying the physical origin may allow additional cuts to be placed in the selection to reject the additional candidates that form the excess, or if this is not possible provide understanding which allows a more informed choice of fit component than the simple Gaussian description used above.

It is noted that the fitted width of the empirical low mass component is 30 MeV/c^2 in the case of the $\chi_c(1P) \to K^+K^-$ distribution, which is much wider than the $\chi_{c0}(1P)$ peaks which have a width of less than 10 MeV/c^2. This large width suggests the excess is due to a partially reconstructed decay, where some of the decay products are not reconstructed and therefore the mass resolution is poor. This is supported by the tight track multiplicity cuts used in the selection and described in Section 4.5, which serve to reject events with additional reconstructed tracks.

5.2.1 Inclusive simulation studies

It is investigated whether the low mass excess is due to the partially reconstructed decays of charmonia with masses in the region immediately above 3360 MeV/c^2 using simulated events. It is hypothesised that the low mass candidates passing the multiplicity and rapidity gap cuts imposed in the selection, optimised to select central exclusively produced $\chi_c(1P) \to h^+h^-$ candidates, are themselves produced in CEP or a similar form of diffraction. Therefore the SuperCHIC event generator, as employed to generate the samples of $\chi_c(1P) \to h^+h^-$ signal events, was used to generate $\chi_{c0}(1P)$, $\chi_{c1}(1P)$, and $\chi_{c2}(1P)$ samples, each containing 1×10^6 events. These particles are decayed inclusively to test which, if any, of their decay modes survive the $\chi_c(1P) \to h^+h^-$ selection. SuperCHIC is not capable of generating inclusively decayed events, therefore the mother particles are left undecayed and an alternative approach used.

A two dimensional histogram of transverse momentum $p_\perp$ versus longitudinal momentum p_z is filled for each sample from the momentum four vectors of the generated particles. Examples of such histograms, for the $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$, are shown in Fig. 5.4.

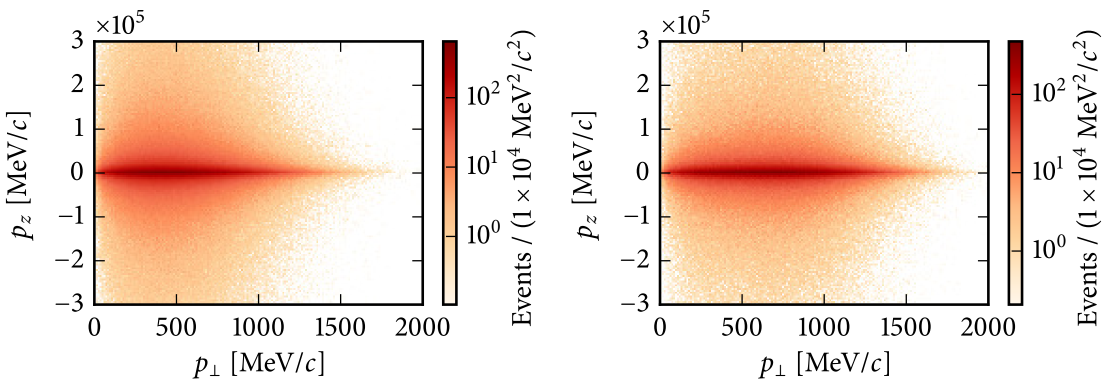

Figure 5.4: The mother particle $p_\perp$-p_z distributions for $\chi_{c0}(1P)$ (left) and $\chi_{c2}(1P)$ (right) as simulated by SuperCHIC.

Inclusive decays of each mother particle are simulated using Gauss, the LHCb simulation application. A so-called particle gun technique is used, where single particles of the desired species are generated with transverse and longitudinal momenta determined by random sampling of the $p_\perp$-p_z histograms. The mothers generated with the particle gun technique are then decayed using the EvtGen event generator [90]. EvtGen is configured to decay the mother particles inclusively, with all decays listed in the *Review of Particle Physics* [97] permitted, except decays to the K^+K^- final state. The number of each decay produced is in proportion to its branching fraction, with a total of 1×10^7 events generated for each mother. The generated inclusive events are then passed through the standard trigger, reconstruction, and stripping pipeline for simulated data, the offline $\chi_c(1P) \to K^+K^-$ candidate selection is applied, and the invariant mass distributions are plotted. The invariant mass distributions of the candidates remaining after application of the offline selection are shown in Fig. 5.3.

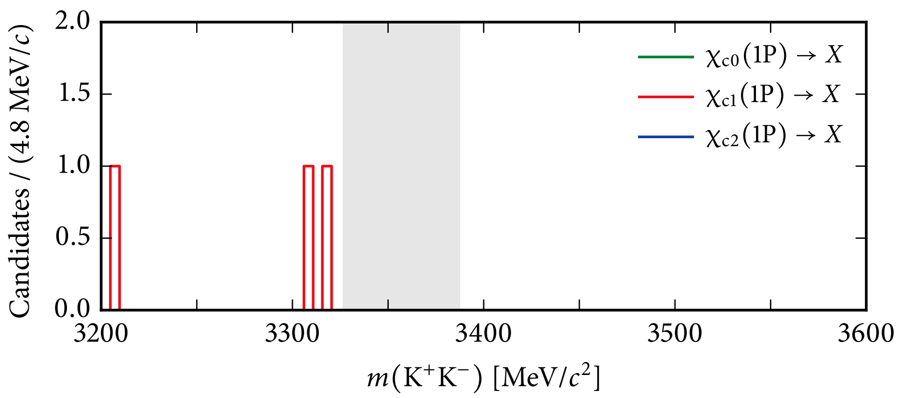

Figure 5.5: The invariant mass distribution of simulated inclusive $\chi_{c0}(1P)$, $\chi_{c1}(1P)$ and $\chi_{c2}(1P)$ events. No $\chi_{c0}(1P)$ or $\chi_{c2}(1P)$ candidates pass the selection. The grey shaded region indicates the width of the Gaussian low mass component fitted to collision data in Fig. 5.3.

No $\chi_{c0}(1P)$ or $\chi_{c2}(1P)$ decays pass the offline selection. Three $\chi_{c1}(1P)$ candidates pass the selection, however these candidates have masses below the structure seen in collision data, and so do not allow any conclusions about the nature of the low mass excess to be drawn.

Additional investigations were carried out by modifying the cuts applied to collision data in the offline selection. A photon veto was applied, requiring candidates with no photon reconstructed in the detector, motivated to reject partially reconstructed radiative decays. The daughter $\Delta \ln \mathcal{L}_{K-\pi}$ and $p_\perp$ cuts were progressively tightened, as were cuts on the track fit quality. However, no cut tested was found to preferentially select or reject candidates contributing to the low mass structure, from which its nature could be identified. Therefore to account for the effect of the low mass excess on distorting the $\chi_{c0}(1P)$ peak and biasing the fitted mass, an additional systematic uncertainty will be assigned to the measured $\chi_{c0}(1P)$ yields using the fitted parameters from the fit with an additional low mass component.

5.3 Exclusive purity determination

The fitted peaks in the invariant mass distributions of selected candidates in Fig. 5.2 contain genuine exclusive decays but also inelastic decays, as inelastic decays may pass the selection when the additional particles produced do not enter the detector acceptance, or otherwise remain undetected. Therefore, in order to make a cross section measurement from the yields determined from the invariant mass fit, the exclusive purity—the fraction of the fitted yield which comprises genuine CEP events as opposed to inelastic events—must be determined.

As discussed in Section 4.1, inelastic background events are expected to have a statistically higher transverse momentum than genuine CEP events [95], and therefore the exclusive purity is determined by fitting the transverse momentum distribution of the reconstructed $\chi_c(1P)$ candidate. This technique is based on that used in previous LHCb studies of CEP using dimuon final states [41, 102]. An unbinned maximum likelihood fit is carried out of the sum of two PDF components, one modelling the CEP contribution and the other the inelastic contribution. The shape of each component is fixed whilst the yield of each component floats. The procedure is carried out for $\chi_{c0}(1P)$ candidates and the $\chi_{c2}(1P)$ purity is taken to be the same as for $\chi_{c0}(1P)$ candidates with the same final state, K^+K^- or $\pi^+\pi^-$, because of the small number of $\chi_{c2}(1P)$ candidates.

5.3.1 CEP component shape

The shape of the fit component for central exclusive events is obtained from simulation. The $p_\perp$ distributions of simulated $\chi_{c0}(1P) \to K^+K^-$ and $\chi_{c0}(1P) \to \pi^+\pi^-$ events passing a modified selection, identical to the offline selection applied to collision events except that no cut on mother $p_\perp$ is applied, are fitted with a Bukin function, which is found to give a better fit quality than any other function tested. It has five parameters, which are the peak location x_p, the peak width σ_p, an asymmetry parameter ξ, and left and right tail parameters ρ_1 and ρ_2. These parameters are all floated in the fits, which are shown in Fig. 5.6. The values of the fitted parameters are listed in Table 5.4, and are used to fix the shape of the signal component in the exclusive purity fit.

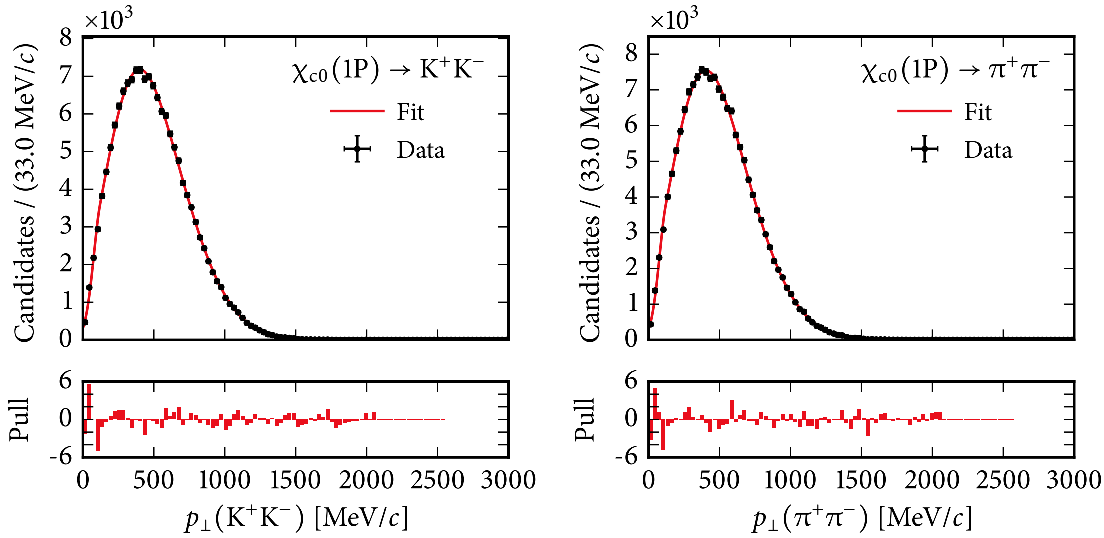

Figure 5.6: Fits of a Bukin function to the mother $p_\perp$ distributions in simulated $\chi_{c0}(1P) \to K^+K^-$ (left) and $\chi_{c0}(1P) \to \pi^+\pi^-$ (right) events, providing the template shape for the CEP signal component in the fit to collision data.

5.3.2 Inelastic component shape

A simulation of the inelastic background processes is not available, so instead the shape of the inelastic component is determined from collision data. A modified selection is applied

Table 5.4: The values of the parameters of the Bukin function fitted to the mother $p_\perp$ distribution in simulated data. These are fixed as the parameters of the signal component in the exclusive purity fit.

Mode	x_p [MeV/c]	ρ_1	ρ_2	σ_p [MeV/c]	ξ
$\chi_{c0}(1P) \to K^+K^-$	399.9 ± 2.1	-9.5 ± 0.6	-0.499 ± 0.019	269.1 ± 1.0	0.118 ± 0.007
$\chi_{c0}(1P) \to \pi^+\pi^-$	404.1 ± 2.1	-10.3 ± 0.7	-0.478 ± 0.018	270.6 ± 1.0	0.111 ± 0.006

which requires three, four, or five tracks, rather than the two tracks required of signal candidates (events with more than five tracks are rejected in the stripping as described in Section 4.4.3, and therefore cannot be selected), and does not include a cut on mother candidate $p_\perp$. Genuine CEP events occurring in isolation contain only two tracks, and so these samples of higher multiplicity events contain predominantly inelastic events. The transverse momentum distributions of three, four, and five track events are independently fitted with a Bukin function in the same manner as used to determine the signal component shape from simulated events, and these fits are shown in Fig. 5.7. The values of the fitted parameters are listed in Table 5.5.

It is observed that the peak transverse momentum increases with the number of tracks. The dependence of the transverse momentum on track multiplicity has not been predicted [103], however this behaviour is consistent with expectations: proton dissociation is more likely with higher $p_\perp$ transfer, and additional radiated gluons will increase the number of tracks in the event whilst also transferring greater $p_\perp$ to the centrally produced state [102]. It is also consistent with the aforementioned dimuon analyses at LHCb.

Therefore the parameters of the fits to events with additional tracks are fitted as a function of track multiplicity with a linear function using the method of least squares. The shape of the inelastic background component in events with two tracks is then estimated by linear extrapolation, and the extrapolated values are given in Table 5.6. These values fix the shape of the inelastic background component in the exclusive purity fit.

The exclusive purity fit is carried out using the shape from simulation for the genuine CEP component and the shape extrapolated from higher multiplicity events for the inelastic background component. The fit is carried out to collision events which pass the offline selection, modified to not cut on the $p_\perp$ of the mother candidate, and have mass of the reconstructed $\chi_c(1P)$ candidate within two widths of the fitted $\chi_{c0}(1P)$ mass. These fits are shown in Fig. 5.8.

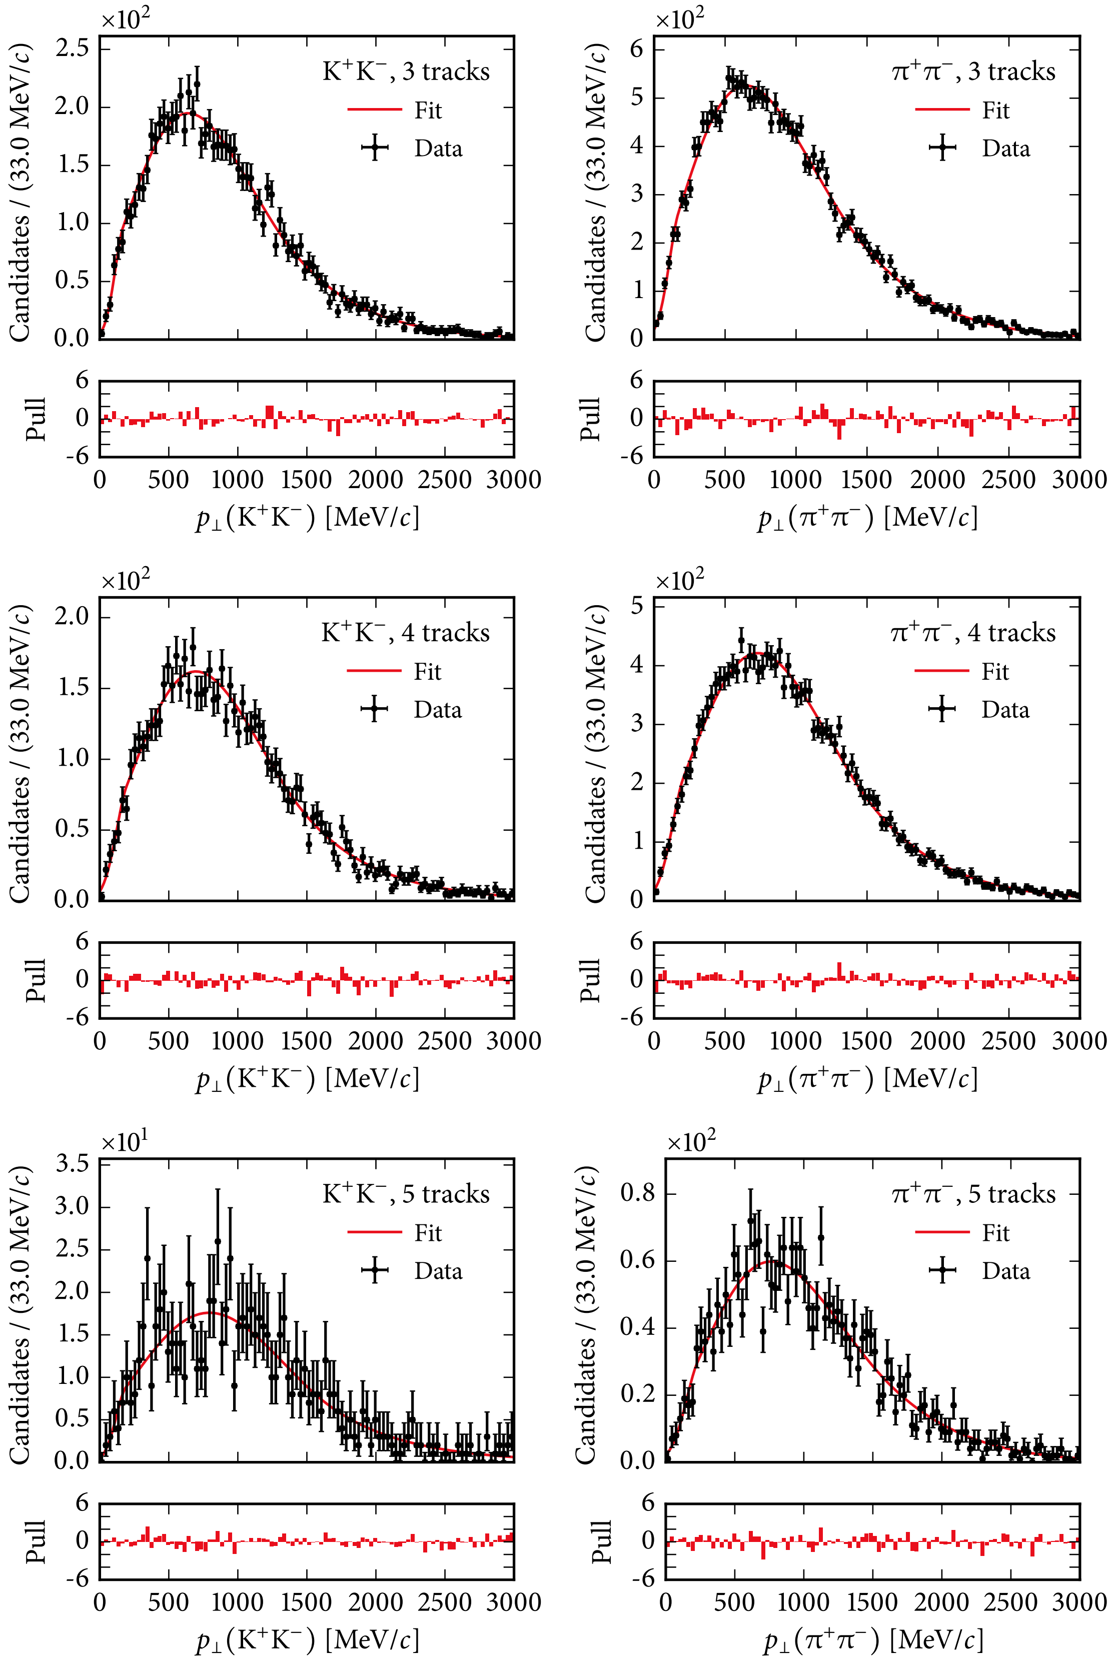

Figure 5.7: Fits to the mother $p_\perp$ distributions in collision events passing the modified offline selections requiring three, four, or five tracks as indicated in the legends.

Table 5.5: The values of the parameters of the Bukin function fitted to the mother $p_\perp$ distribution in collision events passing the modified offline selections requiring three, four, or five tracks.

Mode	Tracks	x_p [MeV/c]	ρ_1	ρ_2	σ_p [MeV/c]	ξ	Yield
$\chi_c(1P) \to K^+K^-$	3	640.0 ± 15.7	-18.5 ± 5.4	-0.06 ± 0.04	466.8 ± 8.5	0.147 ± 0.027	$(7.59 \pm 0.09) \times 10^3$
$\chi_c(1P) \to K^+K^-$	4	696.4 ± 18.4	-12.6 ± 4.0	-0.011 ± 0.035	486.8 ± 10.0	0.125 ± 0.030	$(6.61 \pm 0.08) \times 10^3$
$\chi_c(1P) \to K^+K^-$	5	794 ± 180	-25.0 ± 19.0	-0.00 ± 0.05	560.0 ± 17.5	0.08 ± 0.05	$(8.057 \pm 0.284) \times 10^2$
$\chi_c(1P) \to \pi^+\pi^-$	3	645.4 ± 10.1	-14.8 ± 2.7	-0.095 ± 0.026	483.6 ± 5.2	0.170 ± 0.017	$(2.124 \pm 0.015) \times 10^4$
$\chi_c(1P) \to \pi^+\pi^-$	4	730.9 ± 12.1	-11.7 ± 2.2	-0.073 ± 0.031	504.7 ± 5.7	0.121 ± 0.019	$(1.764 \pm 0.013) \times 10^4$
$\chi_c(1P) \to \pi^+\pi^-$	5	762.8 ± 1.1	-10.3 ± 5.2	-0.153 ± 0.006	528.4 ± 20.3	0.1456 ± 0.0013	$(2.6127 \pm 0.0064) \times 10^3$

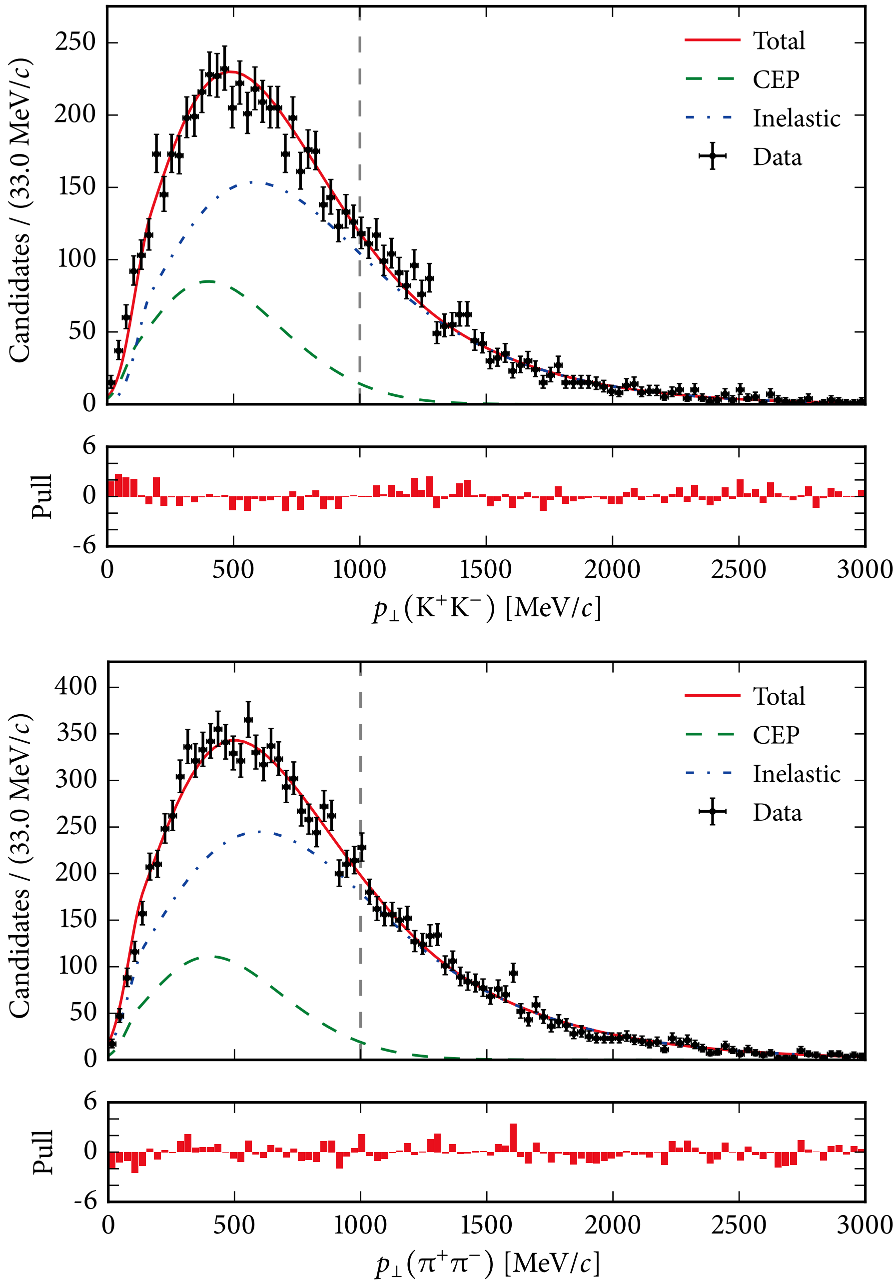

Figure 5.8: Fits to the $p_\perp$ distribution of selected $\chi_{c0}(1P) \to K^+K^-$ (above) and $\chi_{c0}(1P) \to \pi^+\pi^-$ (below) candidates within two widths of the fitted $\chi_{c0}(1P)$ mass. The CEP and inelastic shapes are fixed to predetermined shapes as described in the text. The exclusive purity is calculated for events below the 1 GeV/c cut applied in the offline selection, indicated by the grey dashed vertical line.

Table 5.6: The fixed values of the parameters of the template shape for inelastic background events with two tracks, extrapolated from fits to higher track multiplicity events.

Mode	x_p [MeV/c]	ρ_1	ρ_2	σ_p [MeV/c]	ξ
$\chi_c(1P) \to K^+K^-$	582.3	−20.8	−0.083	423.4	0.180
$\chi_c(1P) \to \pi^+\pi^-$	598.8	−17.1	−0.043	462.0	0.164

The exclusive purity is calculated as the yield of the CEP component below 1 GeV/c (the cut applied in the offline selection) divided by the total fitted yield below 1 GeV/c. The resulting purity is (32.7 ± 1.5) % for $\chi_{c0}(1P) \to K^+K^-$ and (27.6 ± 1.3) % for $\chi_{c0}(1P) \to \pi^+\pi^-$.

5.3.3 Fit validation with pseudo-experiments

The reliability of the fits is studied by carrying out pseudo-experiments in which a large number of toy data sets are generated and fitted. An ensemble of toy data sets is generated by random sampling of the fit PDF after fixing both the parameters and yields to the values favoured by the fit to the collision data. The number of events generated in each toy data set is a Poisson fluctuation of the number of events in the collision data.

Each toy data set is fitted with the same fit model as the real collision data, and the exclusive purity calculated in the same manner. For each pseudo-experiment the pull statistic p is calculated as

$$p = \frac{\lambda_{\text{toy}} - \lambda_{\text{real}}}{\sigma_{\text{toy}}}, \tag{5.1}$$

where λ_{toy} is the exclusive purity from the fit to the toy data set, σ_{toy} is the uncertainty on the purity from the fit to the toy data set, and λ_{real} is the exclusive purity determined for the real data set [104].

The pulls from the pseudo-experiments should be distributed as a Gaussian with a mean of zero and a width of one if the fit is unbiased and the uncertainty is correctly estimated. A deviation of the mean of the pull distribution from zero indicates a bias in the exclusive purity value, whereas a deviation of the width from one indicates a misestimate of the uncertainty.

The pull distributions from ensembles of 30,000 pseudo-experiments are shown in Fig. 5.9, fitted with a Gaussian function.

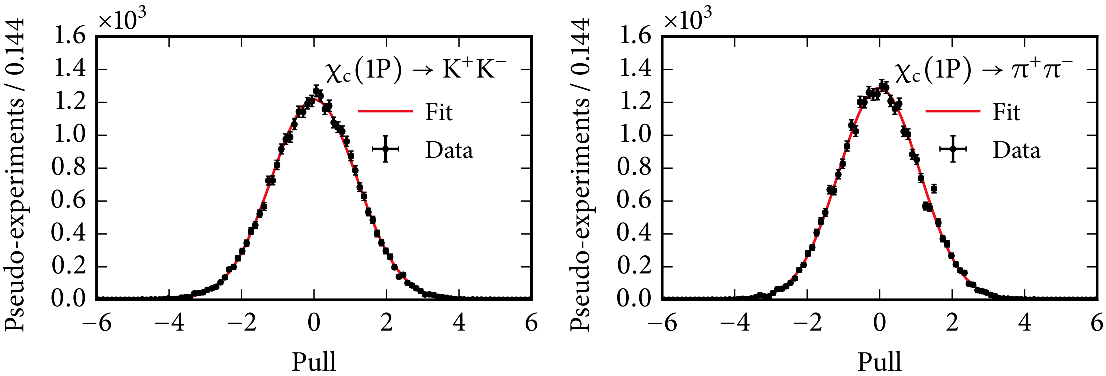

Figure 5.9: The distribution of pulls on the fitted exclusive purity for $\chi_c(1P) \rightarrow K^+K^-$ (left) and $\chi_c(1P) \rightarrow \pi^+\pi^-$ (right) from ensembles of 30,000 pseudo-experiments. The distributions are fitted with a Gaussian function.

The fitted Gaussian function has mean 0.006 ± 0.007 and width 1.180 ± 0.004 for the $\chi_c(1P) \rightarrow K^+K^-$ distribution, and mean -0.009 ± 0.006 and width 1.120 ± 0.005 for the $\chi_c(1P) \rightarrow \pi^+\pi^-$ distribution. This is consistent with a bias of zero, therefore the fit is deemed reliable and no correction to the fitted purity is made. The widths of the pull distributions are 1.18 and 1.12, indicating the uncertainty is underestimated by 18 % in the $\chi_c(1P) \rightarrow K^+K^-$ fit and 12 % in the $\chi_c(1P) \rightarrow \pi^+\pi^-$ fit. The uncertainty quoted on the fitted exclusive purity is therefore increased by this amount, and the corrected purity values are (32.7 ± 1.8) % for $\chi_{c0}(1P) \rightarrow K^+K^-$ and (27.6 ± 1.5) % for $\chi_{c0}(1P) \rightarrow \pi^+\pi^-$.

5.4 Effective integrated luminosity determination

The instantaneous luminosity of a collider L is a function only of the beam parameters, and relates the interaction rate $\mathrm{d}N/\mathrm{d}t$ for a given process to its cross section σ, in accordance with the expression

$$\frac{\mathrm{d}N}{\mathrm{d}t} = L\sigma. \tag{5.2}$$

At a bunched beam collider such as the LHC, the interaction rate is the product of the bunch crossing frequency v and the mean number of interactions per bunch crossing μ, so the instantaneous luminosity can be expressed as

$$L = \frac{v\mu}{\sigma}. \tag{5.3}$$

The selection used in this analysis contains exact cuts on the track multiplicity of each event in order to minimise contamination of the signal sample from inelastic background events, as described in Chapter 4. These cuts reject bunch crossings which contain multiple interactions, even if one of the interactions is a genuine CEP event which would have been selected if it had occurred in isolation in a bunch crossing with no additional interactions. By excluding events in bunch crossings with multiple interactions, the integrated luminosity $\mathcal{L} = \int L\mathrm{d}t$ used in the analysis is reduced from the total integrated luminosity $\mathcal{L}_{\text{tot}}$ to an effective integrated luminosity $\mathcal{L}_{\text{eff}}$, which is the instantaneous luminosity integrated only over bunch crossings without multiple interactions.

The number of interactions per bunch crossing n follows the Poisson distribution

$$P(n;\mu) = \frac{\mu^n e^{-\mu}}{n!}, \tag{5.4}$$

where the mean μ can be expanded as the sum

$$\mu = \sum_{n=0}^{\infty} nP(n). \tag{5.5}$$

This is the mean number of interactions per bunch crossing for the total integrated luminosity. An analogous value for the effective integrated luminosity is defined as the mean number of interactions per bunch crossing in bunch crossings without multiple interactions, denoted by λ. It is found from Eq. (5.5) by evaluating the sum up to $n = 1$:

$$\lambda = \sum_{n=0}^{1} nP\left(n\right) = \mu e^{-\mu}. \tag{5.6}$$

Replacing μ with λ in Eq. (5.3) and substituting in for λ from Eq. (5.6), an expression for the effective integrated luminosity as a function only of $\mathcal{L}_{\text{tot}}$ and μ is obtained:

$$\mathcal{L}_{\text{eff}} = \frac{\nu\lambda}{\sigma} = \frac{\nu\mu}{\sigma}e^{-\mu} = \mathcal{L}_{\text{tot}}e^{-\mu}. \tag{5.7}$$

Therefore the effective integrated luminosity is the total integrated luminosity scaled by a factor of $e^{-\mu}$. Inspection of Eq. (5.4) shows that this factor is the fraction of bunch crossings with zero interactions (henceforth *empty crossings*), $P\left(0\right)$.

The total integrated luminosity is calculated centrally within LHCb, calibrated using van der Meer scans and beam-gas interaction methods. It is calculated for each run, which is a period of continuous data taking with the same detector conditions (this definition of a run as applying to the detector is not the same as the definition of a run applying to the accelerator, which was introduced in Section 3.2).

To calculate the effective integrated luminosity it is then necessary to choose an observable by which a bunch crossing can be classified as empty or not, from which μ is calculated. In this analysis a crossing is classified as empty if fewer than two tracks are reconstructed in the VELO. To be counted, the point of closest approach of a track to the nominal interaction point (IP) must be inside a cylinder centred on the IP of length 600 mm and radius 4 mm. This serves to reject tracks originating from processes other than a proton-proton collision, such as beam-gas interactions or spillover. The number of VELO tracks is used as this quantity had the best stability versus time throughout 2012 of any of the alternative counters.

A threshold of two tracks is used, rather than one, due to the occurrence of ghost tracks, which are spurious tracks wrongly reconstructed from noise hits in the VELO [96].

The fraction of empty crossings per run as determined from the number of VELO tracks in crossings with beam-beam crossings, $P(0)^{\text{b-b}}$, is not however the true fraction of crossings without proton-proton interactions. It is the value for crossings with neither true proton-proton interactions *nor* background interactions,

$$P(0)^{\text{b-b}} = P(0)\, P(0)^{\text{bkg}}. \tag{5.8}$$

This is because a bunch crossing containing no proton-proton interactions will be falsely classified as non-empty if it contains background interactions such as beam-gas or cosmic interactions.

The contribution from background interactions is removed by dividing $P(0)^{\text{b-b}}$ by the fraction of empty crossings in bunch crossings where only the bunches in beam 1 contain protons (beam-empty crossings) and only the bunches in beam 2 contain protons (empty-beam crossings), and so where it is known that all interactions are background interactions [96]. The true value of $P(0)$ is then given as a function of the measured quantities by

$$P(0) = \frac{P(0)^{\text{b-b}}}{P(0)^{\text{b-e}} P(0)^{\text{e-b}}}. \tag{5.9}$$

Substituting Eq. (5.9) into $P(0) = e^{-\mu}$ and rearranging for μ yields

$$\mu = -\ln P(0)^{\text{b-b}} + \ln P(0)^{\text{b-e}} + \ln P(0)^{\text{e-b}}. \tag{5.10}$$

The value of μ calculated for each run in 2012 using this expression is shown in Fig. 5.10, and the total and effective integrated luminosities determined from those values are shown in Fig. 5.11.

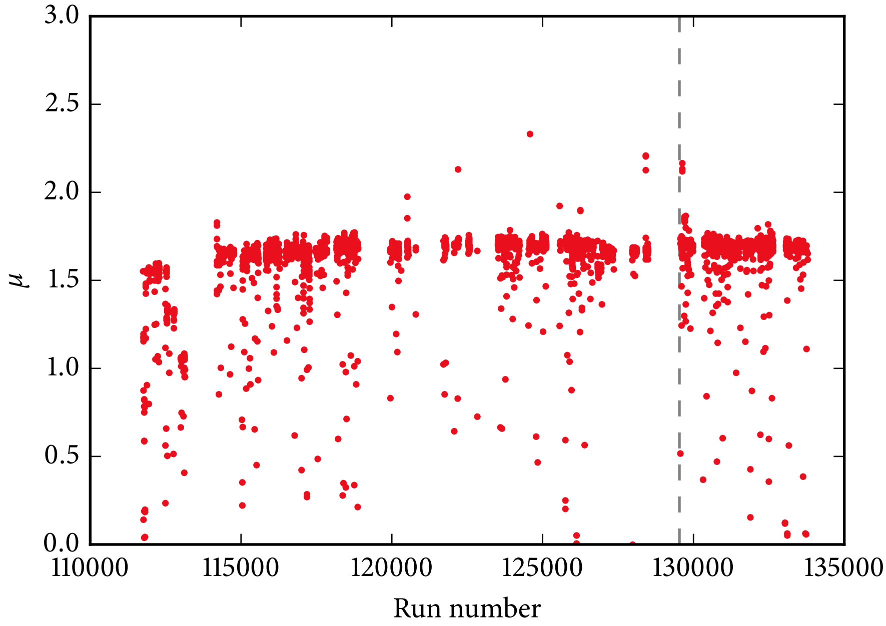

Figure 5.10: The mean number of interactions per bunch crossing for each run in 2012. The $\chi_c(1P) \to h^+h^-$ triggers used in this analysis were operational only from run 129534, indicated by the grey dashed line.

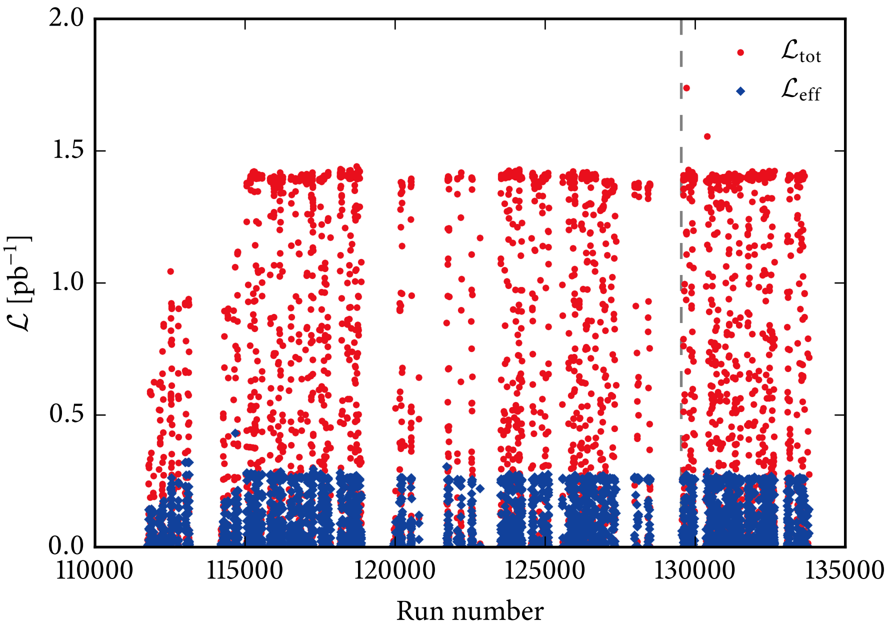

Figure 5.11: The total and effective integrated luminosity per run during 2012. The $\chi_c(1P) \to h^+h^-$ triggers as used in this analysis were operational from run 129534, indicated by the grey dashed line.

The total integrated luminosity collected in 2012 is 1993.3 pb^{-1}, of which 34.4 % was collected during runs with the $\chi_c(1P) \to h^+h^-$ triggers operational, giving $\mathcal{L}_{tot} = 685.3\,pb^{-1}$. Summing over these runs, the effective integrated luminosity for the analysis is $\mathcal{L}_{eff} = 126.3\,pb^{-1}$, which is 18.4 % of $\mathcal{L}_{tot}$.

5.5 Systematic uncertainties

In this section, the sources of systematic uncertainty associated with the selection efficiencies, the invariant mass fit, the exclusive purity determination, and the effective integrated luminosity determination are discussed and evaluated.

5.5.1 Level-0 trigger efficiency

The fitting procedure described in Section 4.6.1, used to determine data-driven corrections to the L0 trigger efficiency, introduces systematic uncertainties. These are from the fit functions imperfectly modelling the SPD multiplicity and hadron $E_\perp$ distributions, the extrapolation of the fitted function to determine the distributions for events failing the trigger cuts, and the assumption of the same correction for all four signal modes. Therefore the change in L0 trigger efficiency introduced by the data-driven corrections, 10.8 %, is assigned as a systematic uncertainty.

5.5.2 Reconstruction efficiency

The track reconstruction efficiency is evaluated by the LHCb Tracking and Alignment group for use by the collaboration, measured with a tag-and-probe method using $J/\psi \to \mu^+\mu^-$ decays. The tag track is fully reconstructed using all the tracking subdetectors, and identified as a muon, whilst the probe track is reconstructed whilst excluding hits in one subdetector. The efficiency is given by the fraction of probe tracks which have a matching fully reconstructed track.

The ratio of the reconstruction efficiency from the LHCb simulation and from the tag-and-probe method is compatible with unity, so no correction to the simulated efficiency is required. A systematic uncertainty of 0.4 % is assigned, which is the maximum change observed when binning the reconstruction efficiency in number of tracks, hits, and primary vertices. An additional systematic uncertainty of 1.4 % is assigned for pions and kaons, to account for additional material interactions compared to the muons used for tag-and-probe [81].

5.5.3 Particle identification efficiency

To account for possible non-linearity in the PID efficiency dependency at low n, the fit in bins of track multiplicity described in Section 4.6.2 is repeated using a second order polynomial rather than a linear function. The extrapolated efficiency for two track events changes to 93.02 % for kaons and 98.46 % for pions, an increase of 1.4 % and 0.4 % respectively. This difference between the extrapolated efficiencies from the linear and quadratic fits is assigned as a systematic uncertainty.

5.5.4 Invariant mass fit

Fixed signal component shapes

A source of systematic uncertainty is introduced in the invariant mass fit by fixing the widths of the signal components to the values obtained by fitting simulated data. To evaluate this, the fits to the mass distribution of simulated events for each of the four modes are repeated, but allowing the Breit-Wigner width to float rather than fixing it to the world average. The invariant mass distribution in collision data is then fitted again with the signal components fixed to the shapes from the floating fit to simulated data. The relative change in the fitted yield compared to the fit with widths fixed to the world average is 6.6 % for $\chi_{c0}(1P) \rightarrow K^+K^-$, 5.3 % for $\chi_{c0}(1P) \rightarrow \pi^+\pi^-$, 5.9 % for $\chi_{c2}(1P) \rightarrow K^+K^-$, and 4.1 % for $\chi_{c2}(1P) \rightarrow \pi^+\pi^-$, and these values are assigned as systematic uncertainties.

Bias in χ_{c0}(1P) mass

The bias in the fitted $\chi_{c0}(1P)$ mass discussed in Sections 5.1 and 5.2 is a significant effect, with a shift in the K^+K^- mode of $-16\,\text{MeV}/c^2$ from the world average. The addition of an empirical low mass component to the invariant mass fit, described in Section 5.2, reduces the mass bias by 20 %, suggesting a common origin of the low mass excess and $\chi_{c0}(1P)$ mass bias.

An improved fit model for the low mass component is expected to reduce the bias further, however visual inspection suggests better modelling of the low mass component alone is insufficient to bring the $\chi_{c0}(1P)$ mass to the expected value, and therefore that there are other underlying causes.

To study this it would be beneficial to improve the exclusive purity of the mass peak, reducing distortion from inelastic background events, and also to remove events forming the low mass excess. If the low mass excess is due to partially reconstructed decays as hypothesised, this could be achieved by detecting the additional particles outside of the LHCb acceptance. This is possible using additional detectors in the LHC tunnel, and such forward detectors will be described in detail in Chapters 6 and 7

These detectors were not operational when the data analysed were recorded, and so the systematic uncertainty introduced by the $\chi_{c0}(1P)$ mass shift is evaluated by repeating the mass fit with the mean of the $\chi_{c0}(1P)$ fit component fixed to $3414.75\,\text{MeV}/c^2$, the world average value, rather than allowing it to float. With this change, the fitted $\chi_{c0}(1P)$ yield is lower in both the $\chi_c(1P) \rightarrow K^+K^-$ and $\chi_c(1P) \rightarrow \pi^+\pi^-$ fits. The relative decrease is 41.6 % and 28.8 % respectively, which are assigned as systematic uncertainties.

Low mass component

The change in the fitted signal yields between the invariant mass fits with and without the additional low mass component introduces an additional systematic uncertainty as described in Section 5.2, as the low mass excess is observed to distort the $\chi_{c0}(1P)$ peak. The systematic uncertainty is taken as the relative change in the signal yields between fits with and without the LMC, which is 12.9 % for $\chi_{c0}(1P) \to K^+K^-$, 7.8 % for $\chi_{c0}(1P) \to \pi^+\pi^-$, 62.2 % for $\chi_{c2}(1P) \to K^+K^-$, and 41.3 % for $\chi_{c2}(1P) \to \pi^+\pi^-$.

5.5.5 Exclusive purity

Four sources of systematic uncertainty in the determination of the exclusive purity are considered: the fixed CEP component shape, the fixed inelastic component shape, the extrapolation of the background component to two track events, and the assignment to the $\chi_{c2}(1P)$ modes of the purity calculated for the $\chi_{c0}(1P)$ modes.

Fixed CEP component shape

The shape of the CEP component is fixed to that obtained by fitting the $p_\perp$ distribution of simulated data. SuperCHIC assumes the squared transverse momentum of the mother particle goes as $\exp(-bp_\perp^2)$, with slope $b = 4\,\mathrm{GeV}^{-2}\,c^2$. An analysis of CEP at LHCb using the $J/\psi \to \mu^+\mu^-$ and $\psi(2S) \to \mu^+\mu^-$ modes was able to determine b from data, and measured a value of $(5.8 \pm 0.8)\,\mathrm{GeV}^{-2}\,c^2$ [102].

A second sample of $\chi_{c0}(1P) \to K^+K^-$ events was generated using SuperCHIC with b set to $5.8\,\mathrm{GeV}^{-2}\,c^2$, and the peak $p_\perp$ of the simulated events is found to shift by $16.7\,\mathrm{MeV}/c$ as a result of the change in b. The exclusive purity fit to collision data is repeated with the fixed CEP component shifted by this value, and the exclusive purity found to change by 2.4 % for $\chi_c(1P) \to K^+K^-$ and 4.0 % for $\chi_c(1P) \to \pi^+\pi^-$. These changes are assigned as systematic uncertainties.

Fixed inelastic component shape

The uncertainty originating from fixing the inelastic component shape is evaluated by repeating the fitting procedure to background events with additional tracks, but with the Bukin function substituted for the Novosibirsk function. The values of the parameters of the Novosibirsk function are extrapolated to the values expected for two track events in the same manner as for the Bukin function, and the fit to collision data repeated. The changes in purity with the alternate extrapolated background shape, which are assigned as systematic uncertainties, are 9.3 % for $\chi_c(1P) \to K^+K^-$ and 6.7 % for $\chi_c(1P) \to \pi^+\pi^-$.

Extrapolation to two track events

The uncertainty from assuming a linear extrapolation of the inelastic component shape in three, four, and five track events down to two track events is evaluated by repeating the exclusive purity fit with a different inelastic component shape. Instead of the extrapolated shape, the shape fitted to three track events is used directly. Given that genuine $\chi_c(1P) \to h^+h^-$ CEP events will contain only two tracks, three track events provide a pure background sample, if the number of genuine $\chi_c(1P) \to h^+h^-$ CEP events with ghost tracks and the number of CEP events with pile-up are negligible. The fitted purity changes by 35.7 % for $\chi_c(1P) \to K^+K^-$ and 44.4 % for $\chi_c(1P) \to \pi^+\pi^-$ when using the alternate shape. Given the strong indication that $p_\perp$ decreases as the number of tracks is reduced, but that with only three bins of track multiplicity it is not sufficient to determine whether the dependency is linear, a systematic uncertainty of half the change in purity when using the three track shape directly is assigned.

Shared purity between $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$

There are an insufficient number of candidates in data to carry out an independent purity fit for the $\chi_{c2}(1P)$ modes, and therefore the $\chi_{c2}(1P)$ purity is assumed to be the same as the corresponding $\chi_{c0}(1P)$ purity. This introduces a systematic uncertainty on the $\chi_{c2}(1P)$ purity. This is estimated by repeating the purity fit to collision data, but taking the CEP component shape from simulated $\chi_{c2}(1P)$ events rather than $\chi_{c0}(1P)$ events. The fit is still carried out to events within two widths of the fitted $\chi_{c0}(1P)$ mass, as no dependence of transverse momentum on mass is observed across the selected mass range. When carrying out this fit, the exclusive purity changes by 18.1 % for $\chi_c(1P) \to K^+K^-$ and 13.2 % for $\chi_c(1P) \to \pi^+\pi^-$, which are assigned as additional systematic uncertainties for the $\chi_{c2}(1P)$ modes.

5.5.6 Effective integrated luminosity

Two sources of systematic uncertainty are considered in the effective integrated luminosity determination: the uncertainty on the total integrated luminosity, and the uncertainty on the determination of the average number of interactions per bunch crossing μ.

The systematic uncertainty on the total integrated luminosity is 1.15 %, evaluated by summing the uncertainty on the total integrated luminosity over those runs where the `Hlt2LowMultChiC2HH` trigger line was operational. The main source of this uncertainty, which is calculated centrally within LHCb for use by all analyses, is the absolute luminosity calibration.

A larger systematic uncertainty arises from the scaling of the total integrated luminosity in the calculation of the effective integrated luminosity. To evaluate the uncertainty arising from the determination of whether a bunch crossing contained an interaction, a separate luminosity study [96] calculated μ for each run during 2012 in the same manner used in this analysis, as described in Section 5.4, with a bunch crossing defined as empty if it contained fewer than two tracks. The calculation was then repeated with the threshold for an event to be considered non-empty lowered from two VELO tracks to one VELO track. The change in the value of μ when lowering the threshold for a non-empty event was found to be 2.4 %, and therefore this is taken as an additional systematic uncertainty.

The systematic uncertainties applying to each mode and their total sum in quadrature are given in Table 5.7.

5.6 Results

To determine the statistical significance of the signal yields measured, a hypothesis test is carried out. Following Ref. [105], the test statistic q is defined as

$$q = -2\ln\frac{\mathcal{L}_{\mathrm{S+B}}}{\mathcal{L}_{\mathrm{B}}}, \tag{5.11}$$

where $\mathcal{L}_{\mathrm{S+B}}$ is the likelihood of the fit to data (the signal plus background hypothesis), and $\mathcal{L}_{\mathrm{B}}$ is the likelihood evaluated when the signal yield is fixed to zero (the background only hypothesis). This choice of statistic is justified by the Neyman-Pearson lemma, which states that q has the greatest power for rejecting an incorrect hypothesis [106]. Larger values of q indicate greater incompatibility between the observed yield and the background only hypothesis. The incompatibility is conventionally quoted as a Gaussian significance Z, which in the case of a one-sided test (the signal yield cannot be negative) is $Z = \sqrt{q}$. If the probability of observing the fitted signal yield if the background only hypothesis is true is p, then Z is the distance from the mean of a Gaussian distribution (in units of sigma) to the point at which the probability under the tail beyond that point is p.

The statistical significance of each of the signal yields calculated in this way is 14.3σ for $\chi_{c0}(1\mathrm{P}) \to \mathrm{K}^+\mathrm{K}^-$, 11.1σ for $\chi_{c0}(1\mathrm{P}) \to \pi^+\pi^-$, 1.7σ for $\chi_{c2}(1\mathrm{P}) \to \mathrm{K}^+\mathrm{K}^-$, and 2.0σ for $\chi_{c2}(1\mathrm{P}) \to \pi^+\pi^-$. The significances of the $\chi_{c2}(1\mathrm{P})$ yields are below the 3σ convention which constitutes evidence of a signal, therefore the $\chi_{c0}(1\mathrm{P}) \to h^+h^-$ cross sections are measured directly whilst upper limits are placed on the $\chi_{c2}(1\mathrm{P}) \to h^+h^-$ cross sections.

Table 5.7: The systematic uncertainties assigned, and their sum in quadrature.

Systematic uncertainty	$\chi_{c0}(1P) \to K^+K^-$ [%]	$\chi_{c0}(1P) \to \pi^+\pi^-$ [%]	$\chi_{c2}(1P) \to K^+K^-$ [%]	$\chi_{c2}(1P) \to \pi^+\pi^-$ [%]
Trigger efficiency: L0	10.8	10.8	10.8	10.8
Reco. eff.: change when binned	0.4	0.4	0.4	0.4
Reco. eff.: material interactions	1.4	1.4	1.4	1.4
PID eff.: extrapolation to two tracks	1.4	0.4	1.4	0.4
Mass fit: signal component shape	6.6	5.3	5.9	4.1
Mass fit: bias in $\chi_{c0}(1P)$ mass	41.6	28.8	–	–
Mass fit: low mass structure	12.9	7.7	62.2	41.3
Purity: CEP component shape	2.4	4.0	2.4	4.0
Purity: inelastic component shape	9.3	6.7	9.3	6.7
Purity: inelastic shape extrapolation	17.9	22.2	17.9	22.2
Purity: shared for $\chi_{c0}(1P)$ and $\chi_{c2}(1P)$	–	–	18.1	13.2
Luminosity: absolute calibration	1.15	1.15	1.15	1.15
Luminosity: determination of μ	2.4	2.4	2.4	2.4
Sum in quadrature	49.8	40.0	69.0	50.8

The product of the cross section and branching fraction $\sigma \times \mathcal{B}$ is given by

$$\sigma \times \mathcal{B} = \frac{pN}{\epsilon \mathcal{L}_{\text{eff}}}, \tag{5.12}$$

where p is the exclusive purity, N is the signal yield, ϵ is the total selection efficiency, and $\mathcal{L}_{\text{eff}}$ is the effective integrated luminosity. The values of these parameters for each signal mode were determined in Sections 5.3, 5.1, 4.6, and 5.4, respectively.

Upper limits on the product of the cross section and branching fraction for $\chi_{c2}(1P) \to K^+K^-$ and $\chi_{c2}(1P) \to \pi^+\pi^-$ are calculated at the 95 % confidence level using the CL_S method [107, 108]. This is a technique for calculating upper limits which uses the test statistic q defined above. The probability the background only hypothesis could give a value of q less than or equal to the observed value is

$$CL_B = p_B(q \le q_{\text{obs}}) = \int_{-\infty}^{q_{\text{obs}}} f(q|H_B)\mathrm{d}q, \tag{5.13}$$

where q_{obs} is the value observed from the data, and $f(q|H_B)$ is the PDF of q for the background only hypothesis, which is given by Wilks' theorem [109]. In the same way, the probability for the signal plus background hypothesis is

$$CL_{S+B} = p_{S+B}(q \le q_{\text{obs}}) = \int_{-\infty}^{q_{\text{obs}}} f(q|H_{S+B})\mathrm{d}q. \tag{5.14}$$

The CL_S statistic is then defined as $CL_S = CL_{S+B}/CL_B$, and the upper limit at the $(1-\alpha)$ confidence level is the signal yield for which $CL_S = \alpha$.

This value is found by scanning through a range of values of the signal yield. For each value, q is calculated from the ratio of likelihoods using Eq. (5.11) and CL_B and CL_{S+B} found by evaluating the integrals in Eqs. (5.13) and (5.14). The CL_S values determined for $\chi_{c2}(1P) \to K^+K^-$ and $\chi_{c2}(1P) \to K^+K^-$ are plotted in Fig. 5.12, expressed as a function of the product of the cross section and branching fraction by application of Eq. (5.12).

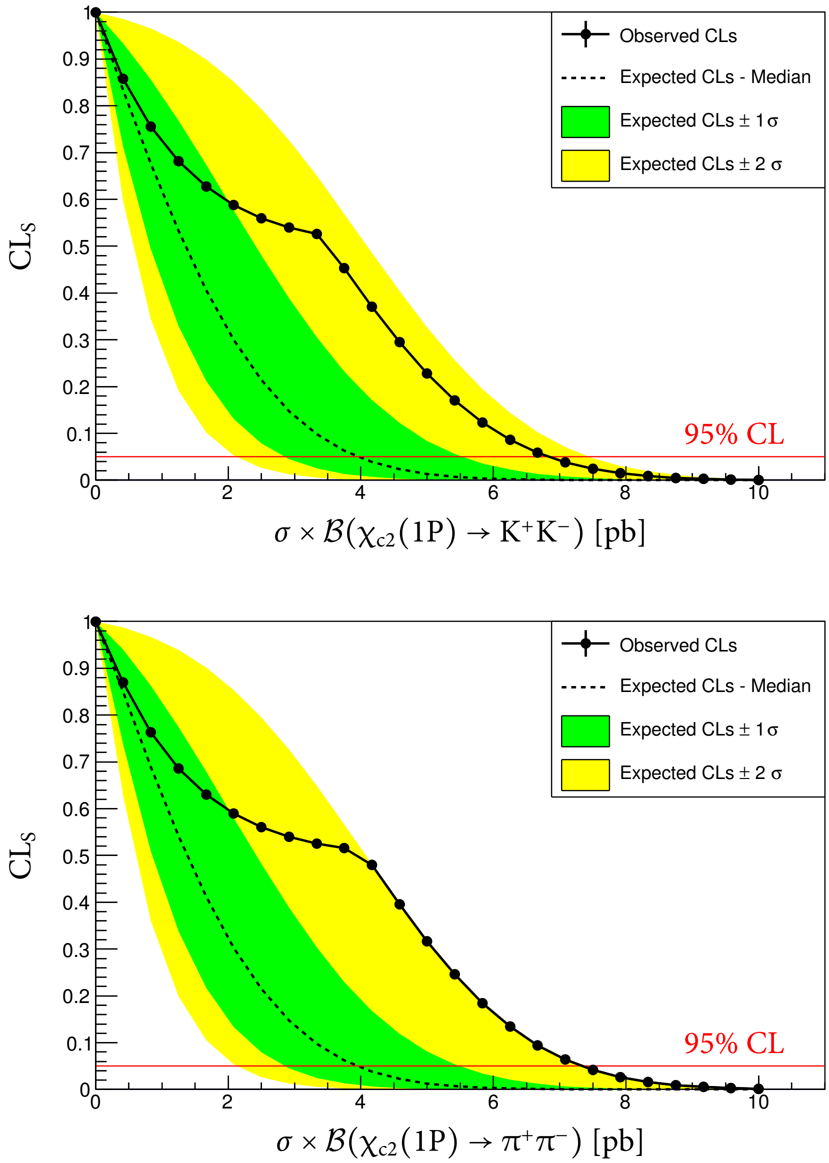

Figure 5.12: $\mathrm{CL_S}$ as a function of the product of the cross section and branching fraction for $\chi_{c2}(1\mathrm{P}) \to \mathrm{K}^+\mathrm{K}^-$ (above) and $\chi_{c2}(1\mathrm{P}) \to \pi^+\pi^-$ (below). The dashed line is the value of $\mathrm{CL_S}$ expected for the background only hypothesis, with 1σ and 2σ intervals shaded in green and yellow respectively. The solid black line with markers is the value of $\mathrm{CL_S}$ for the signal plus background hypothesis. The upper limit on the product of the cross section and branching fraction at the 95 % confidence level is given by the point at which the observed $\mathrm{CL_S}$ reaches 0.05, indicated by the red horizontal line.

The upper limit on the product of the cross section and branching fraction at the 95 % confidence level is given by the point at which the observed $\mathrm{CL_S}$ drops to 0.05. This is 6.8 pb for $\chi_{c2}(1P) \to K^+K^-$ and 7.3 pb for $\chi_{c2}(1P) \to \pi^+\pi^-$. However, these values must be modified to incorporate the systematic uncertainties given in Table 5.7, which weaken the limit. Using the result of Ref. [110], the systematic uncertainties are incorporated by multiplying the upper limit U obtained from Fig. 5.12 by

$$1 + \frac{(U - n)\sigma^2}{2}, \tag{5.15}$$

where n is the value of $\sigma \times \mathcal{B}$ calculated from the central value of the fit and σ is the total relative systematic uncertainty.

The upper limits for the product of the cross section and branching fraction at the 95 % confidence level for $\chi_{c2}(1P) \to K^+K^-$ and $\chi_c(1P) \to \pi^+\pi^-$ with daughters in the range $2.5 < \eta < 4.5$ are listed in Table 5.8, along with the central values for $\chi_{c0}(1P) \to K^+K^-$ and $\chi_{c2}(1P) \to \pi^+\pi^-$ calculated directly from Eq. (5.12).

Table 5.8: The measured product of the cross section and branching fraction for central exclusive $\chi_{c0}(1P) \to h^+h^-$ production for daughters in the pseudorapidity range $2.5 < \eta < 4.5$. For the central values given for $\chi_{c0}(1P) \to K^+K^-$ and $\chi_{c0}(1P) \to \pi^+\pi^-$, the first uncertainty is statistical and the second systematic.

Mode	$\sigma \times \mathcal{B}(2.5 < \eta_{h^\pm} < 4.5)$ [pb]
$\chi_{c0}(1P) \to K^+K^-$	$28.0 \pm 2.8 \pm 13.9$
$\chi_{c0}(1P) \to \pi^+\pi^-$	$20.3 \pm 2.4 \pm 8.1$
$\chi_{c2}(1P) \to K^+K^-$	< 12.1 at 95 % CL
$\chi_{c2}(1P) \to \pi^+\pi^-$	< 10.5 at 95 % CL

The SuperCHIC predictions for the same acceptance cuts are 116 pb for $\chi_{c0}(1P) \to K^+K^-$, 103 pb for $\chi_{c0}(1P) \to \pi^+\pi^-$, 0.27 pb for $\chi_{c2}(1P) \to K^+K^-$, and 0.38 pb for $\chi_{c2}(1P) \to \pi^+\pi^-$. A conservative estimate of the uncertainty on these predictions is quoted by the authors of SuperCHIC as a factor of five, the main source of which is the accuracy of the perturbative approach at the low scale of the $\chi_c(1P)$ mass [28]. Some cancellation occurs when taking the ratio of cross sections for $\chi_{c0}(1P)$ to $\chi_{c2}(1P)$ however, with a quoted uncertainty on the ratio of a factor of two [32]. These predictions are plotted alongside the measured values for $\chi_{c0}(1P) \to K^+K^-$ and $\chi_{c0}(1P) \to \pi^+\pi^-$ in Fig. 5.13.

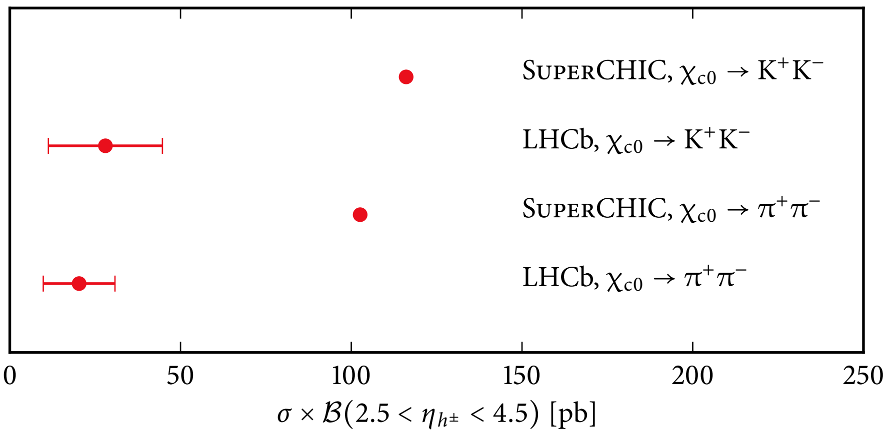

Figure 5.13: Comparison between the measured product of the cross section and branching fraction for $\chi_{c0}(1P) \to K^+K^-$ and $\chi_{c0}(1P) \to \pi^+\pi^-$ and the predictions of the Durham model for $\sqrt{s} = 8$ TeV as implemented in SuperCHIC.

The measured cross sections are significantly below the central values of the predictions, although they are not inconsistent considering the large theoretical uncertainty. The $\chi_{c2}(1P)$ limits are an order of magnitude greater than the predicted values, and so are not sensitive to the predicted cross sections themselves, however they do indicate the ratio of $\chi_{c0}(1P)$ to $\chi_{c2}(1P)$ production cannot be less than of the order of two (SuperCHIC predicts a ratio of the order of 10^2).

To allow comparison with the results of previous preliminary LHCb measurements of $\chi_{c0}(1P)$ CEP using the $\mu^+\mu^-\gamma$ final state, the production cross section for the full solid angle is calculated by first multiplying the measured product of the cross section and branching fraction by the world average branching fraction to obtain the production cross section in the analysis acceptance, and then scaling this by the analysis acceptance efficiency. The acceptance efficiency is calculated from the simulated data generated with SuperCHIC described in Section 4.3 as the fraction of events which have both daughters in $2.5 < \eta < 4.5$, and so introduces a model dependence into the calculated production cross section. The acceptance efficiency is (6.5909 ± 0.0015) % for $\chi_{c0}(1P) \to K^+K^-$ and (6.1999 ± 0.0014) % for $\chi_{c0}(1P) \to \pi^+\pi^-$, where the uncertainty is statistical only.

The production cross section for the full solid angle is determined in the same manner for a preliminary LHCb measurement of $\chi_{c0}(1P) \to \mu^+\mu^-\gamma$ CEP carried out using data collected in 2010 at $\sqrt{s} = 7$ TeV corresponding to an effective integrated luminosity of 3.1 pb^{-1} [40], and an unpublished measurement of the same mode carried out using data collected in 2011 at $\sqrt{s} = 7$ TeV corresponding to an effective integrated luminosity of 222 pb^{-1} [41]. These analyses both required the daughters to be in the range $2.0 < \eta < 4.5$, for which the acceptance efficiency calculated using SuperCHIC is (7.630 ± 0.014) %.

The production cross sections thus determined for the $\chi_{c0}(1P) \to \mu^+\mu^-\gamma$ analyses and $\chi_{c0}(1P) \to h^+h^-$ analysis are given along with the SuperCHIC prediction in Table 5.9, and plotted for comparison in Fig. 5.14. No scaling is carried out to account for the difference in centre-of-mass energy between the $\chi_{c0}(1P) \to \mu^+\mu^-\gamma$ measurements and the $\chi_{c0}(1P) \to h^+h^-$ measurements, as the predicted increase in cross section going from 7 TeV to 8 TeV is only 1.3 %.

Table 5.9: The $\chi_{c0}(1P)$ CEP cross section in the full solid angle as predicted by SuperCHIC and measured using data collected at LHCb. For the K^+K^- and $\pi^+\pi^-$ values, the first uncertainty is statistical and the second systematic.

	$\sqrt{s}$ [TeV]	$\mathcal{L}_{\text{eff}}$ [pb^{-1}]	σ [nb]
SuperCHIC prediction	8	-	298
LHCb, $\mu^+\mu^-\gamma$ (2010)	7	3.1	161.0 ± 130.1
LHCb, $\mu^+\mu^-\gamma$ (2011)	7	222	38.1 ± 52.0
LHCb, K^+K^-	8	126.3	$71.9 \pm 8.1 \pm 35.8$
LHCb, $\pi^+\pi^-$	8	126.3	$58.9 \pm 7.5 \pm 23.5$

The $\chi_{c0}(1P)$ production cross section measured in the K^+K^- and $\pi^+\pi^-$ channels lies significantly below the central value of the SuperCHIC prediction in each case, although remaining consistent given the large uncertainty quoted for the prediction. There is excellent agreement between the cross section measured in the K^+K^- and $\pi^+\pi^-$ channels, although it should be noted that the systematic uncertainties are correlated. There is good agreement between the measured cross section in the hadronic modes and that measured in the $\chi_{c0}(1P) \to \mu^+\mu^-\gamma$ modes in the two prior preliminary analyses, and as such the $\chi_c(1P) \to h^+h^-$ analysis strengthens the conclusions made from the $\chi_{c0}(1P) \to \mu^+\mu^-\gamma$ analyses, that the $\chi_{c0}(1P)$ CEP cross section lies below the value predicted by the Durham model.

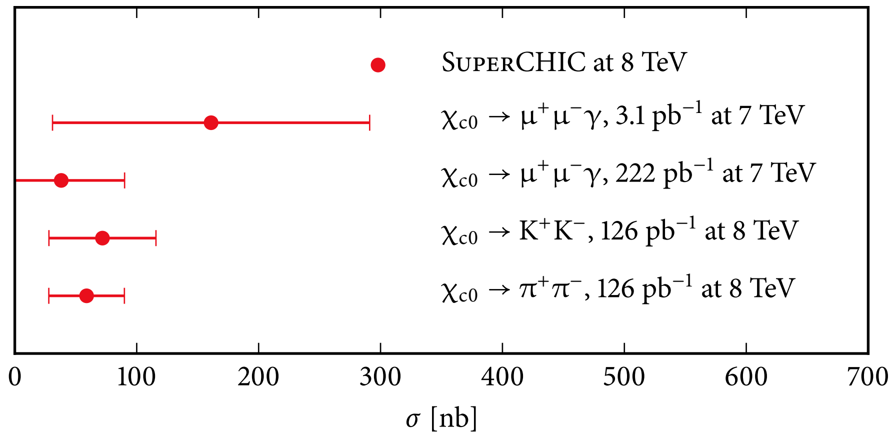

Figure 5.14: The $\chi_{c0}(1P)$ CEP cross section in the full solid angle as predicted by SuperCHIC and measured from LHCb data in the analyses described.

6

HERSCHEL motivations and design

This chapter introduces HERSCHEL (High Rapidity Shower Counters for LHCb), a system of forward shower counters (FSC) installed at LHCb in 2014 and 2015 during the long shutdown following Run 1 of the LHC. The purpose and mode of operation of FSCs is explained, before explaining the motivations for their addition to LHCb and summarising previous studies and installations. The locations chosen for HERSCHEL stations are discussed in light of the physics goals and constraints of the beam line layout, and the design of the stations is presented. Finally, work undertaken to model the beam line geometry and magnetic fields in the LHCb simulation framework is described.

6.1 Motivations

In the analysis of central exclusive $\chi_c(1P)$ production presented in Chapters 4 and 5, the predominant experimental backgrounds are from inelastic processes. As discussed in Section 4.1, the processes most likely to mimic the signature of CEP, and thus contaminate the selected sample, are (i) diffractive processes which would be CEP if it were not for the presence of additional gluon radiation, and (ii) diffractive processes with proton dissociation. Both of these process types proceed similarly to CEP, by photon and Pomeron exchange, and may produce identical final states to CEP in the LHCb detector.

However for type (ii), one incoming proton fragments (single diffraction) or both incoming protons fragment (double diffraction), with the fragments hadronising to form additional particles. This process is sometimes known as central inclusive diffraction [111]. The collision is inelastic and the CEP selection rules do not apply. Due to the large momentum of the colliding protons at the LHC (3.5 TeV/c in 2010 and 2011, and 4 TeV/c in 2012) the remnants from the dissociated proton do not enter the acceptance of the LHCb detector but instead remain within the beam pipe as they leave the vicinity of the interaction region, and are not detected.

It is not possible to determine when proton dissociation has occurred in diffractive events on an event-by-event basis using the current LHCb detector, and so the contribution of these and other inelastic backgrounds to the selected sample is estimated statistically by fitting the $p_{\perp}$ distribution of the reconstructed system with CEP and inelastic background components, as was described in Section 5.3.

The exclusive purity determined using this method is found to be sensitive to (i) the method of determining the shape of the background $p_{\perp}$ distribution for events with the same track multiplicity as signal events, and (ii) the shape of the $p_{\perp}$ distribution for CEP events, which is obtained from simulation. The exclusive purity determination is a dominant systematic uncertainty on the measured cross sections.

The purity of samples of selected CEP candidates at LHCb could be significantly increased if it were possible to veto individual single or double diffractive events by detecting the presence of the additional particles created from the proton fragments. The resulting reduction in systematic uncertainty could improve the precision of a cross section measurement greatly, if good distinction between CEP and inelastic events were achieved. This increased precision would make measurements of CEP at LHCb of greater use in testing theoretical predictions, for example in resolving the discrepancy in the ratio of $\chi_{c0}(1P)$ to $\chi_{c2}(1P)$ production. Additionally, it would be possible to isolate samples of events with proton dissociation from which the residual inelastic background could be studied and better understood.

6.2 Forward shower counters

Whilst the additional particles produced in events with proton dissociation typically have large pseudorapidities and do not enter the acceptance of the LHCb detector, their trajectories do deviate from that of the beam protons. As a result they eventually hit the aperture limit in the LHC tunnel beyond the vicinity of the LHCb detector, such as at a collimator jaw or beam screen.

The collision of proton fragments with the aperture limit initiates particle showers, with high secondary particle multiplicities due to the high momentum of the primary particles. These showers can be used to distinguish events with proton dissociation from CEP events by detecting the secondary particles created in the shower using scintillating detectors. These detectors are called forward shower counters, and are installed inside the accelerator tunnel along the beam line either side of the interaction point.

FSCs are relatively simple devices, composed of a plane of scintillating material oriented perpendicularly to the beam direction. The beam pipe passes through a hole in the scintillator. For optimum sensitivity, shower counter stations should be located close to an aperture limiting accelerator component where showers will be initiated when it is struck by a particle which has travelled from the interaction point. The scintillator absorbs the energy deposited by a particle which traverses it, before emitting that energy as light. The light emitted by the scintillator is detected and amplified with one or more PMTs. In order to detect proton dissociation in both the upstream and downstream directions, and provide the greatest rapidity coverage, multiple stations should be positioned on either side of the interaction point at a range of distances from it. A signal from the shower counters greater than a given threshold value is interpreted as indicating that an inelastic interaction occurred, and that event can then be vetoed from selection as a CEP candidate.

Forward shower counters are only effective in low pile-up conditions. In collisions with significant pile-up, FSCs cannot distinguish between a shower originating from a dissociated proton in a single or double diffractive event and a shower originating from an unrelated inelastic event in the same bunch crossing. However, the expected running conditions at LHCb during Run 2 will provide an ideal low-pileup environment for FSCs. In Run 1 with a 50 ns LHC bunch spacing, luminosity levelling maintained an instantaneous luminosity such that 20 % of bunch crossings contained only a single interaction. In Run 2, the bunch spacing will be reduced to 25 ns, doubling the fraction of bunch crossings with a single interaction for a given instantaneous luminosity. Additionally, the constant instantaneous luminosity throughout each fill provided by luminosity levelling would allow FSCs to operate continuously, rather than only in special low luminosity fills, maximising the amount of data collected.

In addition to vetoing inelastic events such as single and double diffraction, forward shower counters allow the existence of much larger rapidity gaps to be confirmed in CEP events than can be achieved using only a detector at the interaction point with coverage of central pseudorapidities. The maximum pseudorapidity gap that could be required at LHCb during Run 1 was limited to 5 units by the VELO acceptance, which covers pseudorapidities from -3.5 to -1.5 in the upstream direction and from 1.5 to 4.5 in the downstream direction [112]. It will be shown in Section 7.3 that the installation of shower counters at LHCb increases the maximum rapidity coverage to $|\eta| \approx 10$.

Beyond improving the precision of CEP cross section measurements, the installation of shower counters around a central detector also opens access to a broader forward physics programme. For example, they allow single diffractive events to be isolated, enabling the study of low mass single diffractive excitation

$$\mathrm{p} + \mathrm{p} \rightarrow \mathrm{p} + X, \tag{6.1}$$

where X is a state comprising a small number of hadrons with a mass of several GeV. This process cannot be studied with a central detector alone, as X has high momentum but low transverse momentum so is not detected, and no additional central state is produced inside the detector acceptance [111]. Beyond physics studies, shower counters are also of utility for monitoring beam halos, beam losses, and collision rates [113].

6.3 Prior experimental experience

FSCs were installed at the CDF experiment at Fermilab for Run 2 of the Tevatron [113], where the devices were known as beam shower counters (BSC). There were seven BSC stations, four to the West of the CDF interaction point and three to the East, giving a rapidity coverage of $5.5 < |\eta| < 7.5$.

The layout of the stations on the West side is shown in Fig. 6.1. The two stations closest to the interaction point (BSC-1) each contained four counters, whilst the five stations at larger distances (BSC-2, BSC-3, and BSC-4) each contained two counters. Each counter was fitted with a single PMT as shown in Fig. 6.2. The counters were used in analyses of central exclusive dijet, dimuon, $\gamma\gamma$, and $\pi^+\pi^-$ production [39, 114–116].

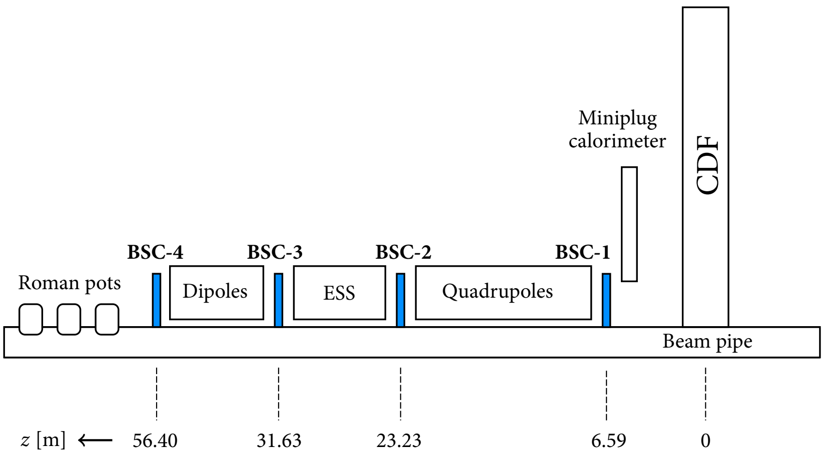

Figure 6.1: A schematic showing the positions of the four shower counter stations BCS-1, BCS-2, BCS-3, and BCS-4 (shaded blue) on the West side of the CDF interaction point with other CDF and Tevatron components [113].

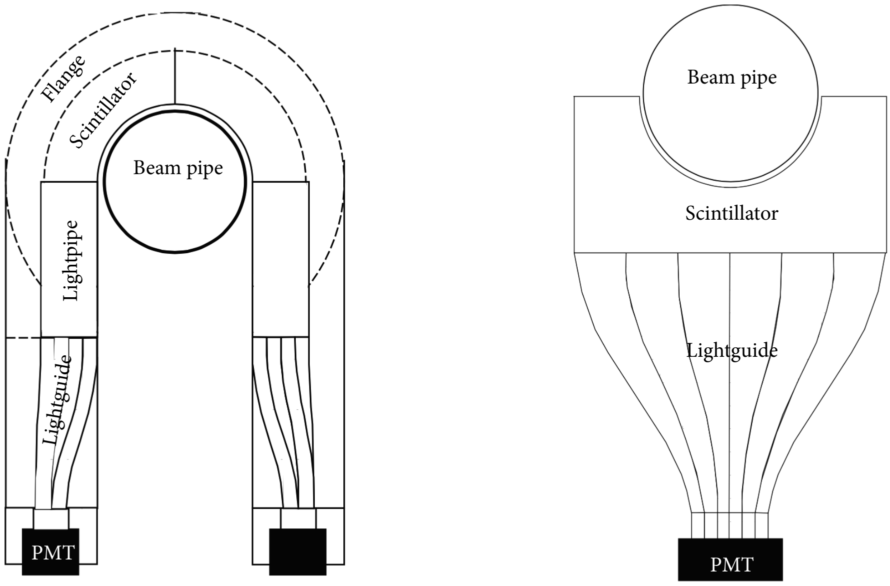

Figure 6.2: Schematics of the two CDF beam shower counter configurations. Left: half of the counter installed at the BSC-1 positions, showing the two scintillator sectors. Right: half of the counter installed at the BCS-2, BCS-3 and BCS-4 positions, showing the rectangular scintillator and semi-circular hole for the beam pipe [113].

The installation of shower counters at LHC experiments to detect diffractive events with rapidity gaps was proposed by Albrow *et al.* in 2009 [111]. As part of the proposal, a simulation study was carried out to determine the benefits to the LHC physics programme which the installation of shower counters could bring. The ability of shower counters to reduce the purity contamination in CEP analyses from inelastic diffractive events was emphasised as a motivating factor. The detection efficiency for single particles and single diffractive events was investigated, and the efficiency with which single diffractive events could be detected using shower counters found to be approximately 75 % for diffractive masses greater than 3 GeV.

The efficiency for detecting central inclusive diffraction was expected to be similar, suggesting that FSCs could provide an effective inelastic event veto in CEP analyses at the LHC. The study used the CMS experiment as a representative example, but the findings are generally applicable to all other central detectors at the LHC [117]. This was confirmed for LHCb in later simulation studies of single particle and single diffractive events [41, 118].

The installation of forward shower counters at CMS was carried out during Run 1 [119, 120]. There are three stations on each side of the central detector, located from 59 m to 140 m from the interaction point to detect showers from particles with pseudorapidity in the range $7.5 < \eta < 10$. They allow detection of single and double diffraction, as well as being used for beam monitoring. Due to the high pile-up encountered at CMS, which operates at higher instantaneous luminosity than LHCb, the counters are most effective during special low pile-up runs with high β^* or at the end of fills in normal running after the bunch current has decayed. The counters were tested in a 12 hour low pile-up run in 2012, and requirements for further commissioning before Run 2 identified, including adding a fourth station on each side and upgrading the PMTs and readout [121].

Motivated by the difficulties experienced in precisely determining the inelastic background contribution in existing CEP analyses at LHCb, in 2013 work began on the HERSCHEL project. This is a system of forward shower counters built and installed at LHCb in the long shutdown between Run 1 and Run 2, to be operational for physics exploitation in 2015.

6.4 Interaction Region 8 layout

The LHCb detector is located in a cavern at the centre of Interaction Region 8 (IR8), one of eight Long Straight Sections (LSS) at the LHC. Along IR8 and extending into the LHC tunnel either side of the LHCb cavern are a set of magnets used for manipulating the beams and bringing them into collision, including a compensator dipole which corrects the orbit bump introduced by the LHCb dipole described in Section 3.3.2. The beam line also contains additional devices for machine protection, beam monitoring, and other purposes.

Dipole magnets are used to guide the beams, steering them from the two horizontally displaced beam pipes in which they orbit the LHC into a single beam pipe for collision at the interaction point, and then separating them again after collision. The beams are focused to a small transverse size by triplets of quadrupole magnets, to provide a high instantaneous luminosity at the interaction point.

Additional devices located along the beam line include collimators and masks, which protect more sensitive components from being damaged by beam halos or beam injection failures (Beam 2, which orbits counter-clockwise, is injected into the LHC from the SPS in the next downstream LSS), and diagnostic devices such as luminosity and beam position monitors.

The positions of the beam line apparatus are specified in a reference coordinate system which has its origin at the nominal interaction point. The z axis is tangential to the beam axis, with the positive direction pointing downstream from LHCb. Negative z is called the left side, and positive z the right side, corresponding to their relative positions if one were to stand at the centre of the LHC and look outwards. The x axis lies horizontally, with the positive x direction pointing out of the LHC ring, and the y axis points vertically upwards to give a right handed coordinate system.

The main beam line apparatus located within 125 m of the interaction point are described in the following paragraphs and are illustrated in Fig. 6.3. Except where otherwise stated, the layout is symmetric around the interaction point. The acronyms in parentheses give the equipment types as specified in the LHC Layout Database [122].

Downstream of the interaction point (positive z coordinate in the LHCb coordinate system) at 3 m is the LHCb dipole magnet (`MBLW`) described in Section 3.3.2. A normal conducting dipole magnet (`MBXWH`) is located at 3.6 m upstream, which has the same strength as the LHCb dipole but opposite polarity to compensate for the orbit deflection introduced

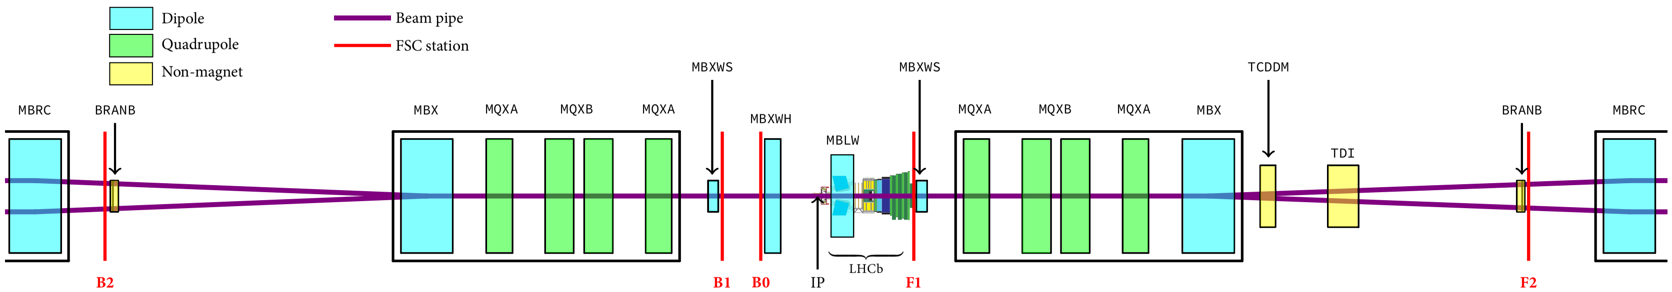

Figure 6.3: A schematic of the beam line layout within 125 m either side of the LHCb interaction point (not to scale), showing the main beam line components and the five positions chosen for FSC stations [52].

by the LHCb dipole. At 20 m on each side is a normal conducting dipole magnet (`MBXWS`), which together close the orbit bump introduced by the LHCb and compensator dipole pair.

Between 23 m and 53 m lies a triplet of superconducting quadrupole magnets (`MQXA`, `MQXB`, and a second `MQXA`), used to provide final focus of the beams before collision. After the quadrupole triplet, at 58 m, is a large dipole separation magnet (`MBX`). This splits the parallel outgoing beams and places them on diverging trajectories, as the beam pipe splits into two. It is followed on the downstream side at 71 m by a fixed mask (`TCDDM`) which protects the separation dipole, and then an injection beam absorber (`TDI`) at 79 m.

After the two beams are clear of one another, a standalone luminosity monitor (`BRANB`) is located between the beam pipes at 114 m. At 122 m is another dipole separation magnet (`MBRC`) which operates in concert with the separating dipole at 58 m. This magnet brings the two beams back onto parallel trajectories within their individual beam pipes after sufficient horizontal separation has been achieved.

The start and end z coordinates of each element described above are listed in Table 6.1, along with their unique names in the LHC Layout Database, which is formed of the equipment type acronym and a location identifier separated by a dot.

6.5 Location of HERSCHEL stations

There are three predominant factors which dictate suitable locations along the beam line for the addition of HERSCHEL stations. Practically, it is necessary to find locations where there is sufficient z-separation between the existing adjacent apparatus that a station can be installed between them. This is the primary restriction at LHCb, as there is a large amount of existing infrastructure around IR8. Secondly, the location should provide close radial access to the vacuum chamber, so that the distance between the beam axis and the edge of the scintillator can be minimised to give maximum rapidity coverage. Finally, each station should be placed near an aperture limiting device—such as a small aperture magnet or

Table 6.1: The major beam line elements within 125 m of the LHCb interaction point. The start and end z coordinates are taken from the FLUKA model of IR8 described in Section 6.7.

Name	Start [m]	End [m]	Description
		Upstream side	
MBXWH.1L8	−3.55	−6.95	Compensating dipole
MBXWS.1L8	−20.375	−21.155	Compensating dipole
MQXA.1L8	−22.965	−29.335	Superconducting focusing quadrupole
MQXB.A2L8	−32.05	−37.55	Superconducting focusing quadrupole
MQXB.B2L8	−38.55	−44.05	Superconducting focusing quadrupole
MQXA.3L8	−46.965	−53.335	Superconducting focusing quadrupole
MBX.4L8	−58.383	−67.833	Single aperture separation dipole
BRANB.4L2	−113.6	−113.75	Standalone luminosity monitor
MBRC.4L8	−121.687	−131.128	Twin aperture separation dipole
		Downstream side	
MBXWS.1R8	20.375	21.155	Compensating dipole
MQXA.1R8	22.965	29.335	Superconducting focusing quadrupole
MQXB.A2R8	32.05	37.55	Superconducting focusing quadrupole
MQXB.B2R8	38.55	44.05	Superconducting focusing quadrupole
MQXA.3R8	46.965	53.335	Superconducting focusing quadrupole
MBX.4R8	58.383	67.833	Single aperture separation dipole
TCDDM.4R8	70.728	71.728	Fixed mask protecting MBX.4R8
TDI.4R8	78.7405	82.9255	Beam absorber for injection
BRANB.4R2	113.6	113.75	Standalone luminosity monitor
MBRC.4R8	121.687	131.128	Twin aperture separation dipole

collimator—which will initiate showers in a direction such that the secondary particles will be incident on the scintillator.

Suitable locations for five stations were identified in the tunnel around LHCb: three upstream of the interaction point and two downstream, as shown in Fig. 6.3. Upstream stations are labelled “B” for backward, and downstream stations “F” for forward. The first upstream station is B0, located at −7.5 m behind the `MBXWH.1L8` dipole. Preliminary single particle simulations showed detection efficiencies of between 75 and 100 % as a function of $p_{\perp}$, with showers originating in the upstream end of `MBXWH.1L8` (the simulations are presented in detail in Chapter 7). A second upstream station, B1, is located at −19.7 m, close to the `MBXWS.1L8` dipole, with simulated detection efficiencies of between 80 and 100 %.

The final upstream station is B2, located at −114 m, upstream of `BRANB.4L2`. BRAN (Beam RAte of Neutrals) detectors are Argon gas ionisation detectors used to monitor the collision rate [123]. As part of the BRAN detector, a solid copper converter is installed between the two beam pipes on the nominal z-axis. The location of B2 behind the BRAN provides high detection efficiency (~100 % from single particle simulations) for high rapidity neutral particles which shower in the converter.

In the downstream direction, the F1 station is located before the `MBXWS.1R8` dipole at 20 m, approximately the same distance from the interaction point as B1 and providing similar detection efficiency. This is the station closest to the interaction point in the downstream direction, due to the asymmetry of the LHCb detector, and is mounted on a transverse shielding wall behind the LHCb muon detectors. The final station is F2, a mirror of B2 located at 114 m behind `BRANB.4R2`.

During preliminary studies another pair of possible locations were identified at ±69 m, behind the MBX separation dipoles. Significant showering was expected in the quadrupole triplet and dipole, however simulations showed that this position is in the “shadow” of the

MBX dipole, resulting in a detection efficiency of less than 20 %. It was therefore decided not to place stations in these locations.

6.6 Station design

The central component of each HERSCHEL station is a square sheet of plastic scintillator, split into four quadrants each of dimensions 300 mm × 300 mm × 10 mm. This size was predominantly determined by mechanical considerations. The scintillator is wrapped in 1 mm thick aluminium sheeting, required by fire safety regulations. The beam pipe passes through a hole in the middle of each station, the size of which is determined by the size of the vacuum chamber at each location. For the B0 and B1 stations, the hole is a circle with a radius of 47 mm to fit around the vacuum chamber's external radius of 42 mm. The location of the F1 station is coincident with a set of vacuum bellows, and so the circular hole is required to have a radius of 61 mm to fit around these bellows. A non-circular hole is used at B2 and F2 to accommodate the dual beam pipes. The holes are rectangles of height 108 mm and width 230 mm, with rounded corners of radius 34 mm.

Figure 6.4 shows a schematic of the layout of the five stations, and Table 6.2 lists the geometric acceptance of each (defined as the pseudorapidities of straight lines traced from the nominal interaction point to the inner and outer edges of the scintillator).

Table 6.2: The geometric acceptance of each HERSCHEL station, evaluated for $x = 0$.

Station	$\eta_{\min}\|_{x=0}$	$\eta_{\max}\|_{x=0}$
B0	−5.77	−3.91
B1	−6.73	−4.88
B2	−8.35	−6.63
F1	4.89	6.49
F2	6.63	8.35

The scintillator is EJ-200 plastic scintillator manufactured by Eljen Technology [124]. This was chosen for its quick response, having rise and decay times of 0.9 ns and 2.1 ns

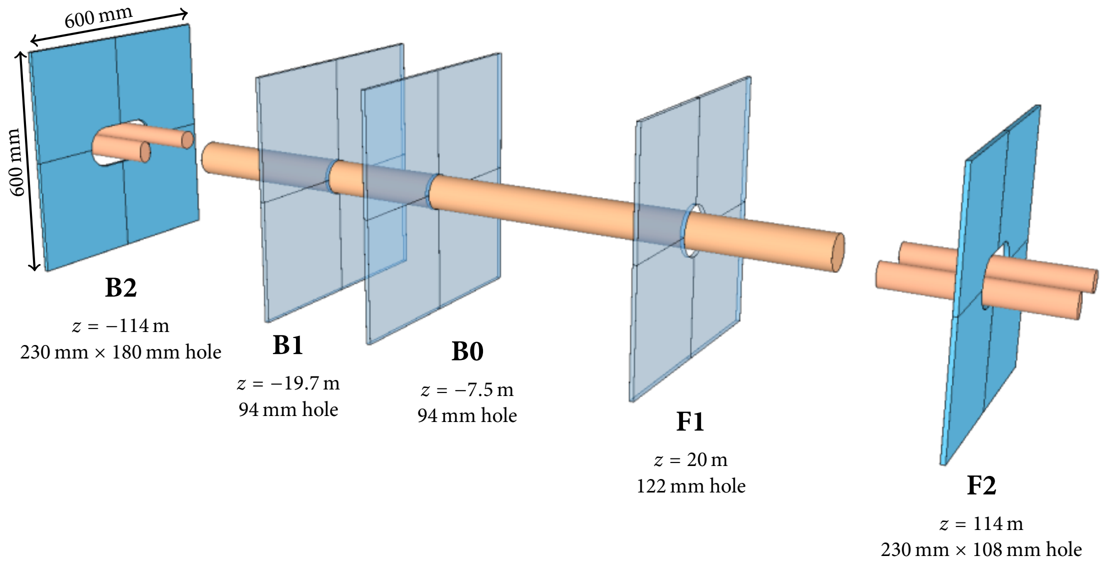

Figure 6.4: A schematic (not to scale) of the layout of the HERSCHEL stations, showing the four scintillator quadrants, circular holes in B0, B1, and F1, and rounded rectangular holes in B2 and F2, through which the beam pipes pass.

respectively. EJ-200 has a density of 1.023 g cm^{-3}, a peak emission wavelength of 425 nm, and a yield of 10^4 photons per 1 MeV electron.

Light guides in the shape of a fishtail are glued to the external horizontal edge of each scintillator quadrant, as shown in Fig. 6.5, and guide the scintillation light to one of four vertically mounted PMTs. The light guides are wrapped in the same aluminium sheeting as the scintillator.

Each PMT is a 51 mm diameter Hamamatsu R1828-01 [125], shown with a disassembled housing in Fig. 6.6 and housed and mounted in place on a station in Fig. 6.7. This model was chosen for its quick response time, with a rise time of 1.3 ns, and its large gain range. A high gain of 10^6 to 10^7 is required for calibration with cosmic muons, but a low gain of 10^3 to 10^4 required in operation to ensure the maximum average anode current of 200 μA is not exceeded in high multiplicity inelastic events.

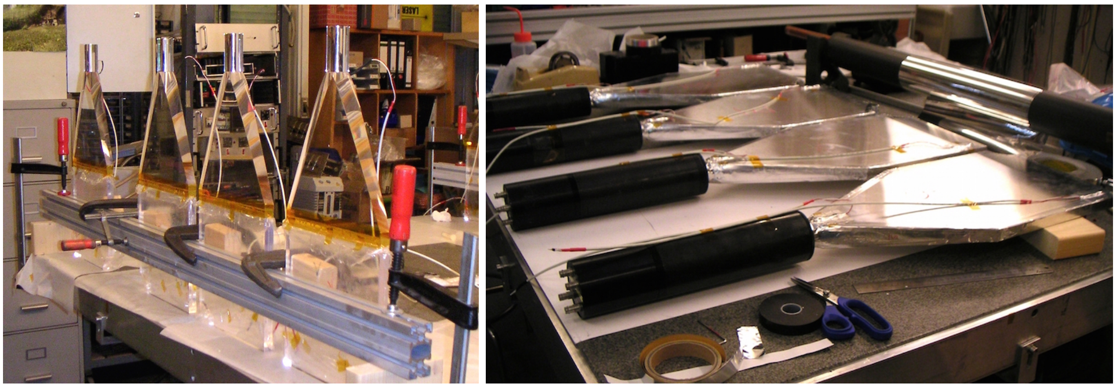

Figure 6.5: Four light guides being glued to scintillator quadrants during construction (left), and after being wrapped with aluminium foil and attached to PMTs (right).

Figure 6.6: An R1828-01 PMT [125] (left) and the disassembled components of a PMT housing with an assembled housing at the rear (right). The device with internal wiring visible and external coaxial connectors is the custom divider.

The divider circuit, which provides the correct distribution of voltages across the PMT dynodes, must be able to cope with the high counting rate expected in operation. A rule of thumb for stable PMT performance is that the divider current should exceed the average anode current by a factor of 10 to 20. Therefore for an average anode current of 200 μA, the divider must drain 2 mA to 4 mA in low gain operation. Calibration however requires a much higher gain, which prevents using a purely resistive divider circuit as the increased gain would lead to excessive heating.

Instead, a custom divider circuit was designed based on experience with LHCb's Beam Loss Scintillator system. An additional bias voltage is applied to the last dynodes, ensuring there is sufficient current drain at the anode whilst restricting the initial divider current. Zener diodes are placed between dynodes to ensure the last dynodes maintain a stable voltage whilst draining a large current [126]. The high voltage supply is delivered to each divider through two 5.2 mm diameter cables running from the LHCb cavern. These are accompanied by 4.8 mm diameter cables, two per station, which power LEDs fitted to aid calibration of the counters.

The signal from each PMT is transferred to the front-end electronics in the LHCb cavern through a 15.1 mm diameter coaxial cable, chosen for its low signal attenuation over the distance from the farthest stations to the cavern. A clipping line is used to contain the signal within 25 ns and thus eliminate spillover into the next bunch crossing. The FE electronics are based on those developed for the LHCb calorimeters. There are two electronics crates, one each for the backward and forward stations, which each contain a front-end Board (FEB) and a Calorimeter Readout Card (CROC). The FEB performs 12 bit digitisation and pedestal subtraction, and transfers the digitised signal to the CROC through the backplane. The CROC collects the data from a single event and transmits it through optical fibres to the LHCb data acquisition system.

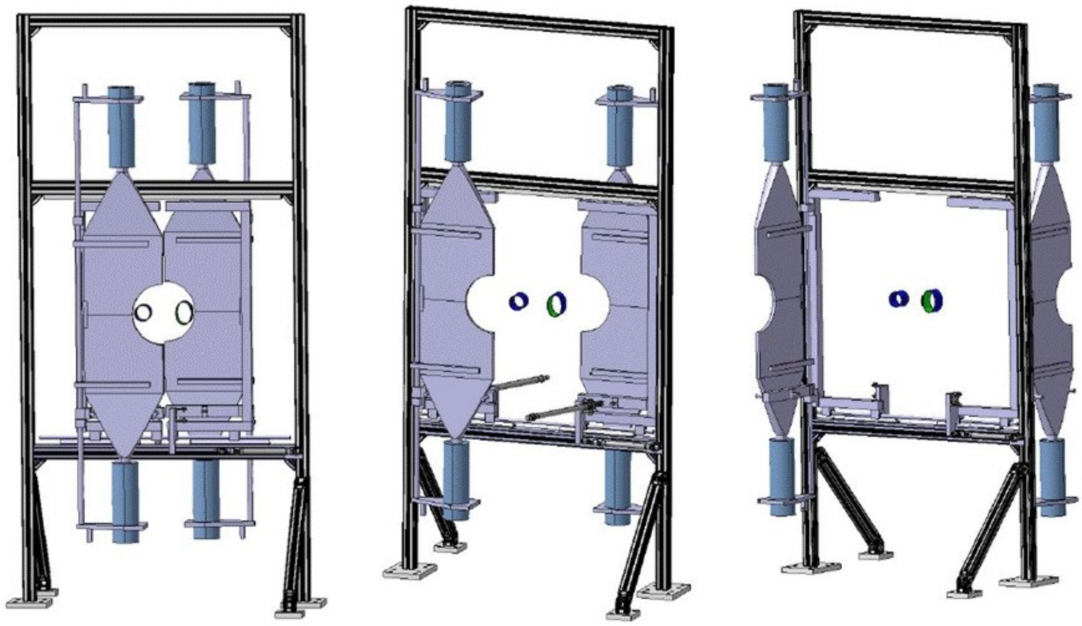

Figure 6.7: A schematic of the far station design representing B2 and F2, showing the operation of the pneumatic retraction mechanism. The station is shown in the data taking position (left), retracted horizontally 150 mm (centre), and retracted horizontally 150 mm and rotated outwards 90° (right). This illustration shows an early circular hole design which was superseded by the rectangular hole design described in the text.

To minimise the dose absorbed by the scintillator when not in use, each station can be retracted from the beam line with a remotely operated pneumatic system. This also allows access to the bellows at F1 if a vacuum intervention is required. The B0, B1, B2, and F2 stations are mounted on frames bolted to the floor through four feet and retract 150 mm horizontally away from the beam pipe before rotating outwards by 90° as shown in Fig. 6.7. As the F1 station is mounted on the shielding wall behind the muon detectors it cannot rotate and instead retracts horizontally 500 mm. Figure 6.8 shows the five stations installed in the tunnel.

6.6.1 Light yield calibration

The light yield of each combined scintillator, light guide and PMT counter assembly was measured using the cosmic calibration stand shown in Fig. 6.9. The counter under test was placed between the blue boxes, above which were two 16 cm diameter circular scintillator counters, and below two 10 cm × 10 cm square scintillator counters. These upper and lower counters were vertically displaced by 32 cm and wired in a coincidence circuit triggering on cosmic muons which passed through the assembly at a rate of 0.21 Hz. Each counter was exposed for 12 hours and the charge collected when struck by a cosmic muon was recorded. The mode of the charge distribution was taken as the charge released by a minimising ionising particle (MIP), and the light yield measured to be 178 photoelectrons per MIP.

6.7 Geometric and magnetic model

Simulation of physics processes and the detector response at LHCb is carried out using Gauss [84,85], the experiment's simulation framework, which uses Geant4 [93,94] to track the passage of the particles produced in collisions and model their interactions with matter.

A geometry description written in XML specifies the position, shape, size and material composition of each component of the detector, from the VELO on the upstream side to

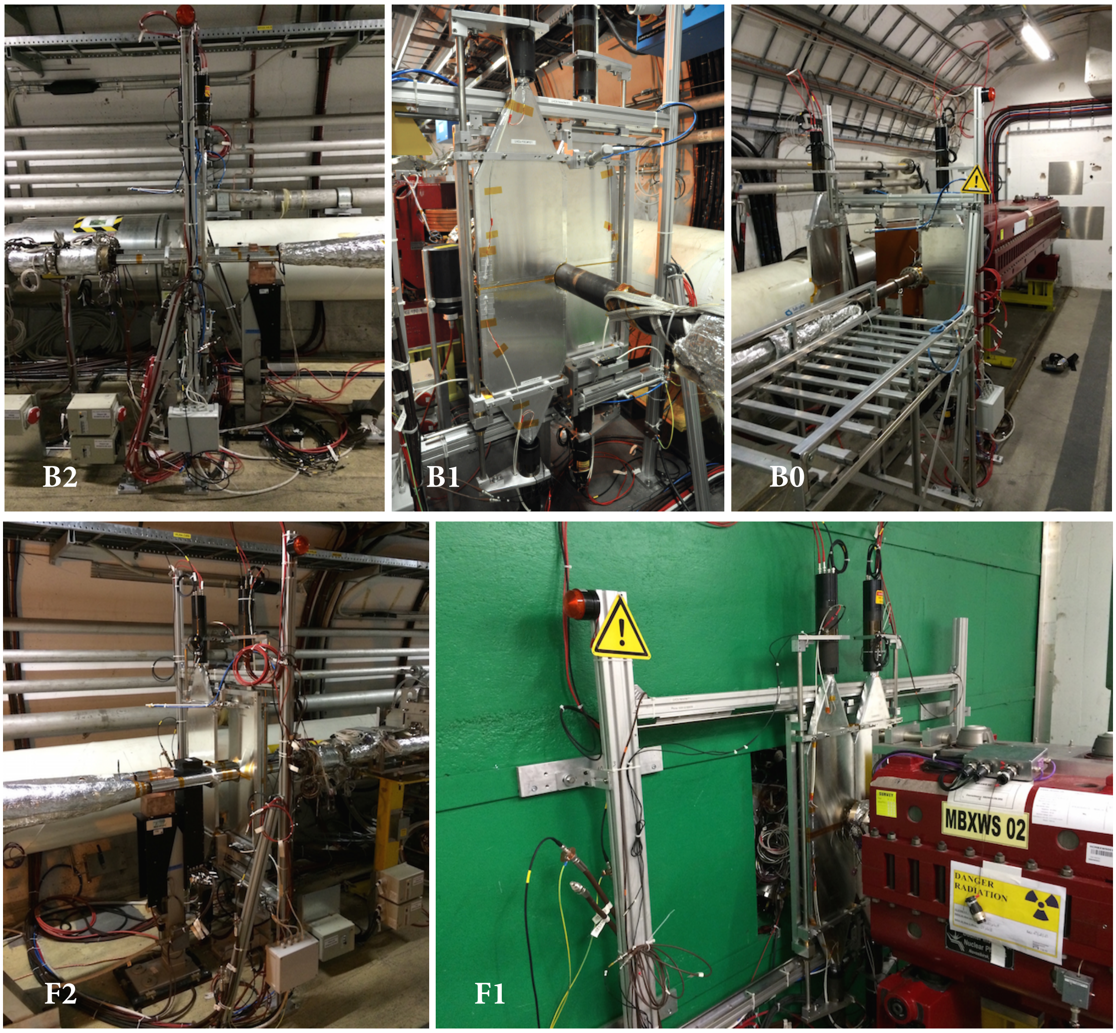

Figure 6.8: The HERSCHEL stations installed in the tunnel. Clockwise from top left: B2 at −114 m; B1 at −19.7 m with the red `MBXWS.1L8` dipole behind it; B0 at −7.5 m (pictured partially retracted) with the red `MBXWH.1L8` dipole behind it; F1 at 20 m mounted on the green shielding wall with the red `MBXWS.1R8` magnet to the right; and F2 at 114 m.

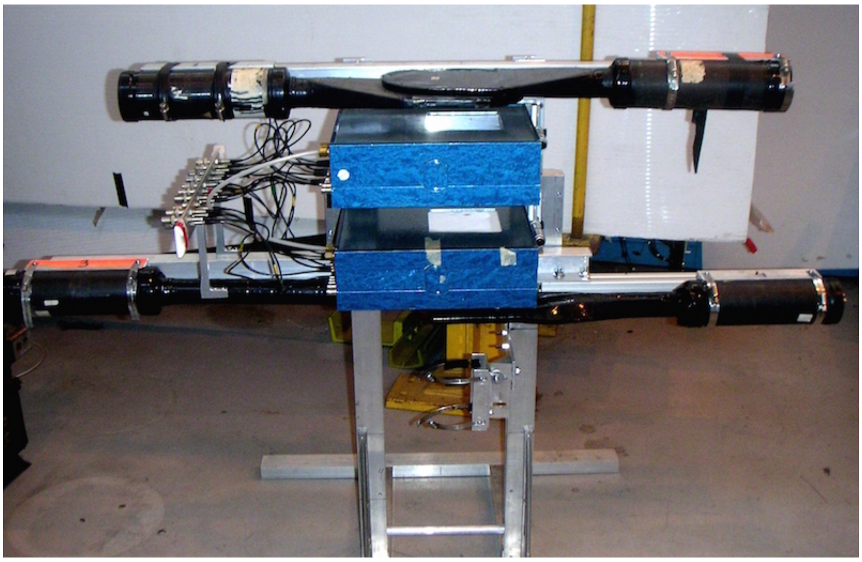

Figure 6.9: The cosmic calibration stand used to measure the light yield of the HERSCHEL counters.

the muon detectors on the downstream side. A separate magnetic field map specifies the field vector at discrete points inside the aperture of the LHCb dipole magnet, for charged particle tracking.

In order to study the expected performance of the HERSCHEL system and optimise the station locations and designs, the existing geometry and magnetic field descriptions were extended to include the beam line apparatus up to 114 m from the interaction point in each direction, so that high rapidity particles originating from the interaction point can be tracked to their impact with the aperture limit and the resulting shower can be simulated.

An extension of the LHCb geometry extending 20 m from the interaction point in each direction, and including field maps for the `MBXWH` and `MBXWS` compensating dipoles, was already available, created in 2010 for beam background studies [127], and was adopted. The remaining geometry to 114 m was implemented based on technical drawings and a FLUKA [128, 129] model created in 2007 for studies of machine induced backgrounds [130]. The FLUKA model provided a comprehensive description of the vacuum chamber, dipole and quadrupole magnets, and tunnel shielding.

In the geometry the shower counters are modelled as 600 mm × 600 mm × 10 mm plastic sheets with central holes for the beam pipes as described in Section 6.6. The EJ-200 scintillator is modelled as polystyrene, $(C_6H_5CHCH_2)_n$, with a density of 1.032 g cm^{-3}. The 1 mm thick aluminium sheeting wrapping the scintillator is included in the model. Figure 6.10 shows the positions of the stations in IR8 and the surrounding beam line apparatus.

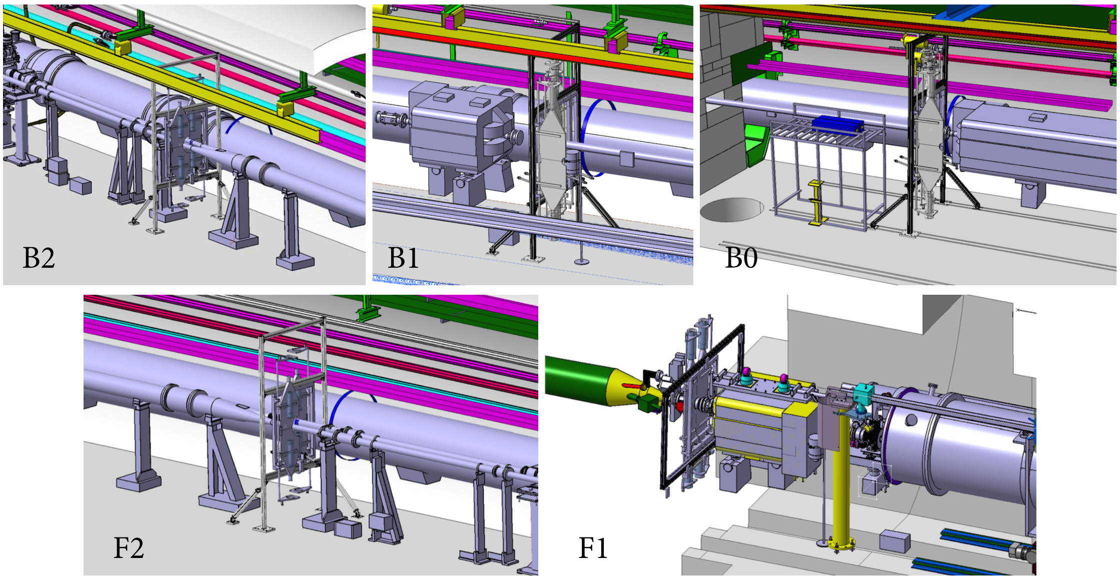

Figure 6.10: Schematics showing the HERSCHEL stations and their surroundings in IR8, including the beam pipe and magnets included in the geometry description. Clockwise from top left: B2 at −114 m; B1 at −19.7 m; B0 at −7.5 m; F1 at 20 m; and F2 at 114 m.

Field maps for the quadrupole triplets and separation dipoles described in Section 6.4 were produced by converting maps produced for use with FLUKA in the aforementioned machine induced background studies to the native file format of Gauss. The FLUKA field maps for the quadrupole triplet were made from measured field maps, whilst the maps for the separation magnets assume an ideal dipole field of nominal strength. The FLUKA maps are for a beam energy of 3.5 TeV, and so the field strength was scaled linearly to account for the increased beam energy in Run 2. Visualisations of the field density in the `MQXA` and `MQXB` quadrupoles made from the final Gauss field maps are shown in Fig. 6.11.

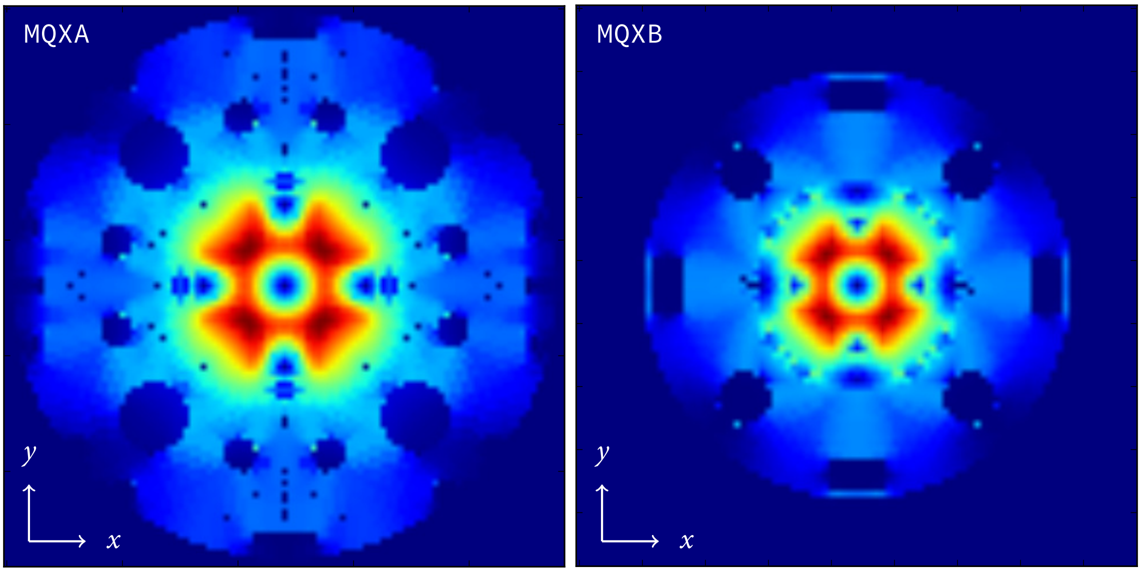

Figure 6.11: Visualisations of the field density in the x-y plane inside the MQXA (left) and MQXB (right) quadrupoles as translated into the Gauss file format from the original FLUKA field maps.

The addition of the geometry and magnetic field descriptions up to the farthest HERSCHEL stations to Gauss enabled a full simulation study of the HERSCHEL detectors' performance to be carried out, which is described in Chapter 7.

7

HERSCHEL simulations and performance projections

This chapter presents simulations of the HERSCHEL detectors and performance projections made using the results of these simulations. Single particle simulations carried out to determine the pseudorapidity coverage of the stations are described, and are used to evaluate the single particle detection efficiency. Simulations of full proton-proton collisions are then presented. Minimum bias events are modelled in order to determine the energy deposition in the scintillators in typical Run 2 conditions, and single diffractive events used to estimate the required sensitivity for detecting diffractive events with proton dissociation. Cosmic calibration data are combined with simulated energy deposition data to predict the photoelectron yield at the PMTs and evaluate the detection efficiency for diffractive events with proton dissociation.

7.1 Motivations

The utility of the HERSCHEL system is quantified by the efficiency with which it can distinguish individually between CEP events and events with proton dissociation or additional gluon radiation, and the size of rapidity gap that can be ensured by requiring that there is no activity in any of the stations.

To evaluate the pseudorapidity coverage and detection efficiency of the stations, a set of simulation studies were carried out using single particles, and full proton-proton collisions modelling the physics processes of interest described in Section 6.1. Simulation studies informed the mechanical design of the stations and the design of the front-end electronics, such as the divider circuit described in Section 6.6, which required knowledge of the expected photoelectron yield. They additionally provide information on where thresholds for detecting signals in the shower counters should be set, in order to detect events with proton dissociation or additional gluon radiation with maximum efficiency whilst minimising the sensitivity to noise due to hits from background sources.

The geometry and magnetic field models described in Section 6.7 were integrated into the LHCb simulation application, Gauss [84, 85]. Gauss is normally configured to track collision products only within the spatial extent of the LHCb detector, however for these studies it was modified to continue tracking particles down the tunnel up to 120 m from the interaction point, in both the backward and forward directions.

Gauss was further extended so that information about particles hitting the scintillators was written to disk, along with pointers to each consecutive mother particle, so that the entire chain of decays from the final hit in the scintillator back through the shower to the initial proton-proton collision could be reconstructed.

7.2 Single and double diffraction

The particles which the HERSCHEL stations are to be optimised to detect are those originating in the dissociation of beam protons in single and double diffraction. In order to determine the typical makeup of an event with proton dissociation and the kinematics of the primary particles in such events, samples of 10,000 single diffractive and 10,000 double diffractive proton-proton collisions were generated using PYTHIA [64] with the standard LHCb configuration [85]. Beam parameters corresponding to expected early Run 2 LHC running at a centre-of-mass energy of 13 TeV were modelled, listed in Table 7.1, as the conditions in which the HERSCHEL stations will be first used. It is important to note that quantities determined from these samples, as plotted in Figs. 7.1 to 7.5, are dependent on the models implemented in PYTHIA.

Table 7.1: Expected Run 2 operating conditions for collisions at a centre-of-mass energy of 13 TeV, modelled in the simulations of single and double diffraction described in the text.

Parameter	Value	Unit
Beam momentum	6.5	TeV/c
β^*	3.0	m
Horizontal crossing angle	−0.105	mrad
Vertical crossing angle	0	mrad
Emittance	2.5	µm
RMS bunch length	90.0	mm
Luminosity per bunch	0.167×10^{30}	$\mathrm{cm^{-2}\,s^{-1}}$
Total cross section	101.5	mb
Mean interactions per crossing	1.5	-

Figure 7.1 shows the number of primary particles in each type of diffractive event. Such events, as modelled by PYTHIA, are found to result typically in approximately 30 primary particles being produced in total, with a mean of 29 in single diffraction and 33 in double diffraction.

The pseudorapidity of primary particles in each simulated sample is plotted in Fig. 7.2, with the acceptance of the VELO ($-3.5 < \eta < -1.5$ in the backward direction and $1.5 < \eta < 4.5$ in the forward direction) indicated by grey shading. The single diffractive sample was generated with half the events having the dissociated proton travelling downstream and half with it travelling upstream, so the distribution is symmetric about $\eta = 0$ as for double diffraction. The distributions for both event types extend out to beyond $|\eta| = 10$, showing most of the ~ 30 particles per event to be produced in the very forward or backward directions. The peaks at $\eta = \pm 10$ for single diffraction correspond to the pseudorapidity of the outgoing beams, and are due to the single proton which remains intact in these events.

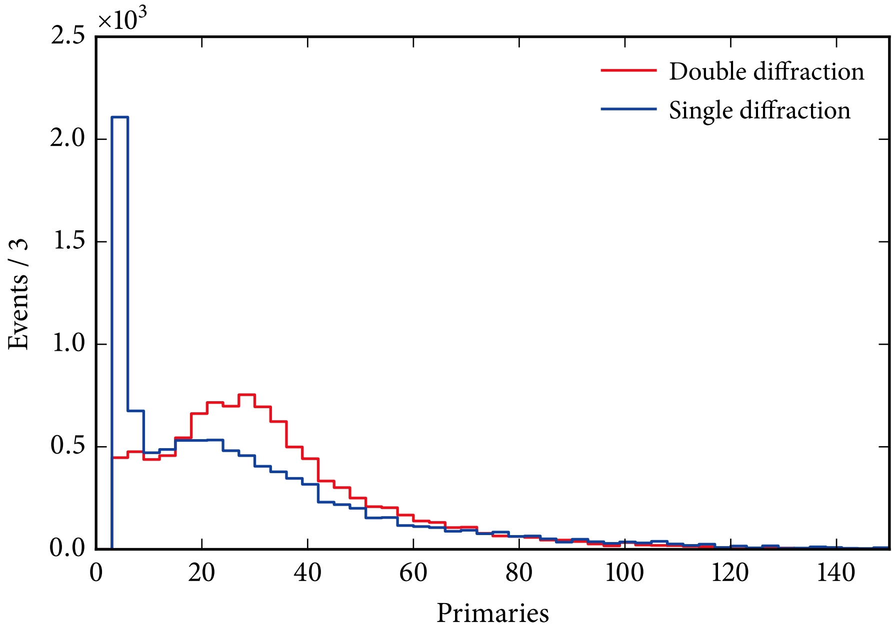

Figure 7.1: The number of primary particles per bunch crossing in single and double diffractive events, as simulated by PYTHIA.

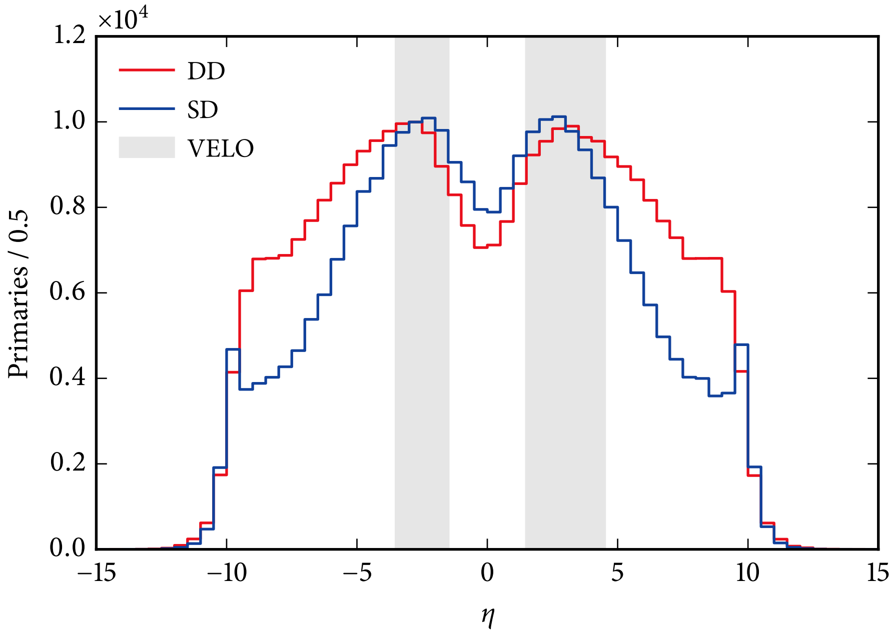

Figure 7.2: The pseudorapidity of primary particles in single diffractive (SD) and double diffractive (DD) events, as simulated by PYTHIA. Most primary particles in each event type are outside the VELO acceptance, which is indicated by the grey shading.

Summing across all simulated collisions, the fraction of primary particles with pseudorapidity within the VELO acceptance is 34 % for single diffraction and 29 % for double diffraction. A larger proportion of the primary particles have pseudorapidity beyond the maximum acceptance of the VELO (*i.e.* $\eta < -3.5$ or $\eta > 4.5$), 48 % for single diffraction and 57 % for double diffraction, and therefore cannot be detected by the LHCb detector. The remaining 18 % of primary particles in single diffraction and 14 % in double diffraction have $|\eta| < 1.5$ and would exit the VELO without passing through enough sensors, if any, to be detected. It is important to note that the VELO is only capable of detecting charged particles, whereas both charged and neutral particles shower upon hitting an aperture limit, which can then be detected with FSCs.

The fraction of primary particles which have pseudorapidity beyond the maximum VELO acceptance in each collision is plotted in Figure 7.3. For single and double diffraction in both the forward ($\eta > 0$) and backward directions ($\eta < 0$) the distribution is strongly peaked at unity, indicating the majority of single and double diffractive events have no primary particles within the VELO, only particles with pseudorapidity beyond its maximum acceptance. This demonstrates that additional forward detectors such as FSCs sensitive to high pseudorapidity particles could identify a large proportion of single and double diffractive events which the central LHCb detector alone cannot.

This is also demonstrated by the momentum and transverse momentum distributions shown in Fig. 7.4. Whilst the momentum is typically of order 1 TeV/c, with the distribution having a tail extending to the 6.5 TeV/c beam momentum (the peak at this value in the single diffraction distribution corresponds to the single intact proton in such events), the transverse momentum is of order 0.1 GeV/c with a tail not extending beyond 3 GeV/c. Particles with such high momentum but small transverse component do not diverge greatly from the beam direction and hence do not enter the VELO. With the exception of the peak in absolute

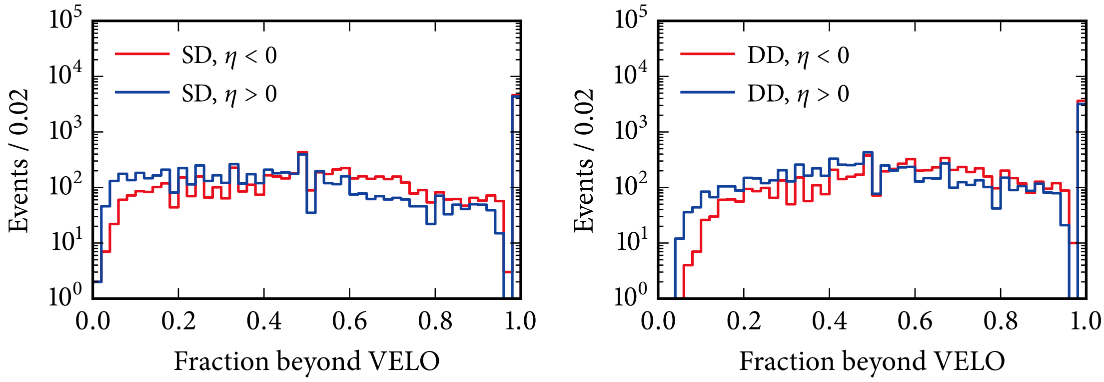

Figure 7.3: The fraction of primary particles in single (left) and double (right) diffraction which have pseudorapidity beyond the maximum VELO acceptance, as simulated by PYTHIA. The distributions peak at unity due to the large number of events with no particles within the VELO.

momentum for single diffraction already described, there is no qualitative difference in the momentum of primary particles in single and double diffraction.

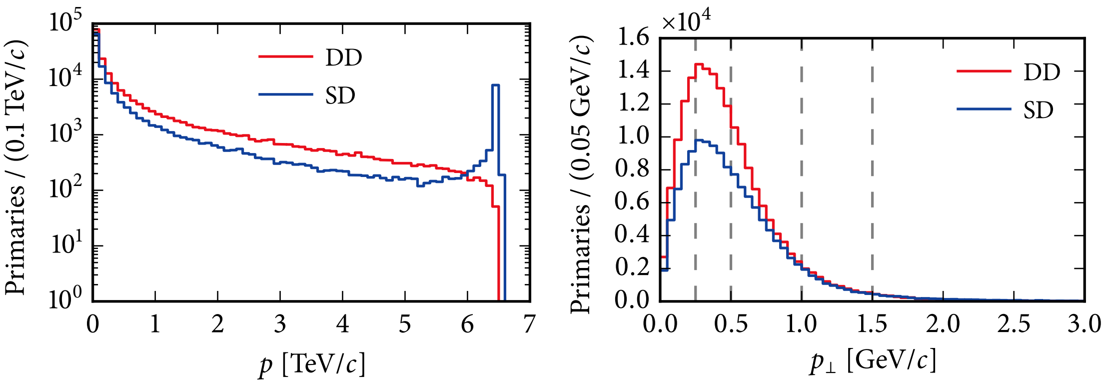

Figure 7.4: The absolute (left) and transverse (right) momentum of primary particles with pseudorapidity beyond the maximum VELO acceptance in single and double diffraction, as simulated by PYTHIA. The vertical grey dashed lines indicate the transverse momentum values modelled in the single particle simulations of Section 7.3: 0.25, 0.5, 1.0, and 1.5 GeV/c.

The species of primary particles in diffractive events where a proton dissociates are shown in Fig. 7.5. They are predominantly hadrons containing up and down quarks, as expected from hadronisation of the light quarks inside a proton. For both single and double diffraction, pions are the most commonly produced particles. Again, the distribution of species is similar in both types of diffractive event and from here no distinction will be made between primary particles in single diffractive events and those in double diffractive events.

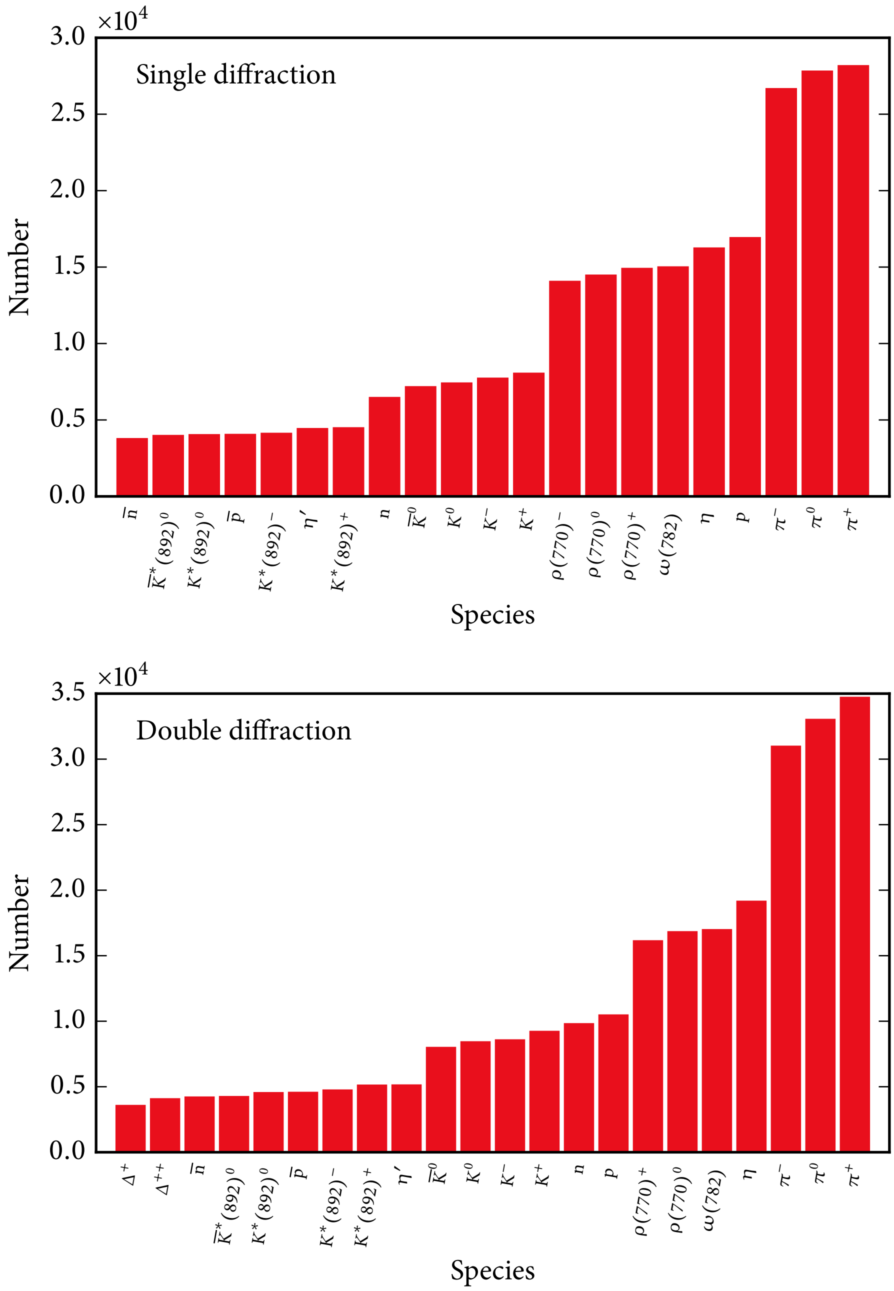

Figure 7.5: The number of each species of primary particle produced in single (above) and double diffraction (below), as simulated by PYTHIA. Species with a multiplicity less than 10 % of that of the most common species, π^+, are not plotted.

7.3 Single particle efficiency and rapidity coverage

The detection efficiency of the stations for individual particles typical of those produced in events with proton dissociation was determined using single particle simulations. A single particle of a given species was created at the nominal interaction point, travelling with fixed pseudorapidity and transverse momentum. Both the primary particle and secondary particles originating from decays and material interactions were tracked through the extended IR8 geometry, and impacts with the scintillator in the HERSCHEL stations recorded as hits.

Only a single particle was created in each event, without any underlying collision, so any activity in the stations is directly attributable to that primary particle. The detection efficiency for the stations—the probability that a particle with given kinematics will lead to some given number of hits—can therefore be evaluated as a function of the pseudorapidity and transverse momentum of the primary particle. For this study an ideal response is assumed such that the efficiency with which an individual hit is detected is 100 %.

The pseudorapidity and transverse momentum values simulated were chosen to sample the ranges found for single and double diffraction in Section 7.2. The pseudorapidity distribution was shown in Fig. 7.2 to extend continuously from the edge of the VELO acceptance, indicated by the grey shaded regions, to $|\eta| \approx 10$. The transverse momentum distribution is seen in Fig. 7.4 to peak at approximately 0.3 GeV/c, with a tail out to beyond 2 GeV/c. Therefore transverse momenta of 0.25, 0.5, 1.0, and 1.5 GeV/c were modelled in the single particle simulations, with pseudorapidities ranging from the maximum VELO acceptance on each side ($\eta = -3.5$ on the backward side and $\eta = 4.5$ on the forward side) to $|\eta| = 10$ in increments of 0.5 units.

Samples of 4000 events were generated for each combination of pseudorapidity, transverse momentum, and particle type from π^+ and π^0, except those combinations with both $p_\perp \geq$ 1.0 GeV/c and $|\eta| > 9.0$ as this corresponds to momentum greater than the 6.5 TeV/c beam momentum and is therefore unphysical. The particles were generated with a uniform distribution in the azimuthal angle.

The z coordinate of the origin vertex of particles hitting each of the five stations is plotted in Fig. 7.6. The majority of particles striking each station are found to originate in the few metres in front of the station, consistent with the hits being from shower secondaries produced when the primary particle hits the aperture limit, as opposed to direct hits from primary particles travelling from the interaction point. Additional hits are seen to originate along the distance between the station and the interaction point for B2 and F2. This can be attributed to secondary particles which travel at sufficiently large rapidity to continue down the beam pipe until striking B2 and F2.

The positions and extent of the beam line components listed in Table 6.1 are indicated with grey shading. Peaks in the distributions are observed near the ends of magnets, as showering occurs when a particle hits the inner aperture of the magnet, but is contained within it if the hit is too far from the end. This can be seen at the end of `MBXWH` magnet at −7 m for B1, and at the end of the `MBX` magnets at −68 m and 68 m for B2 and F2 respectively. A small number of hits are from backsplash on the side of the station facing away from the interaction point. The hits are predominantly electrons and positrons, which constitute 99.1 % of all hits as shown in Fig. 7.7, with a mean momentum of 279 MeV/c, plotted in Fig. 7.8.

The single particle detection efficiency is now considered, assuming the ideal case that individual hits are detected with 100 % efficiency. The single particle efficiency is defined as the fraction of simulated events which result in a number of hits in the stations greater than a given threshold, calculated for a specific primary particle pseudorapidity and transverse momentum. The detection efficiencies thus calculated from the single particle simulations are plotted as a function of pseudorapidity in Figs. 7.9 and 7.10 for π^+ and π^0 mesons respectively. The efficiency is plotted for detection thresholds of one, three, and five hits.

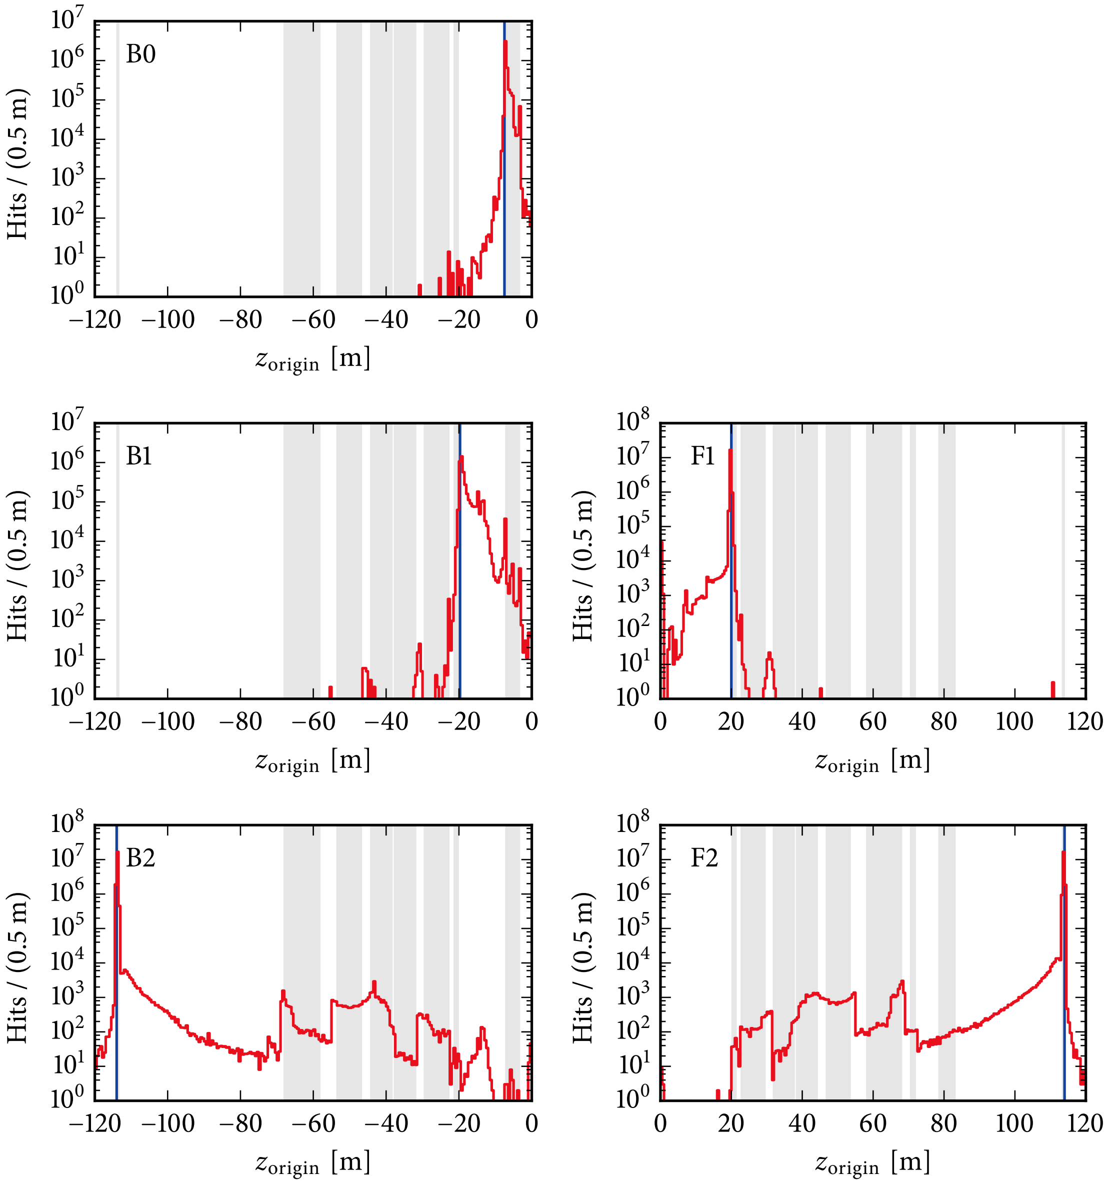

Figure 7.6: The z coordinate of the origin vertex of particles hitting B0 (top left), B1 (centre left), B2 (bottom left), F1 (centre right) and F2 (bottom right), indicating where the particles striking each station originated from. The vertical blue lines indicate the location of each station. The light grey shaded regions show the positions of beam line components listed in Table 6.1.

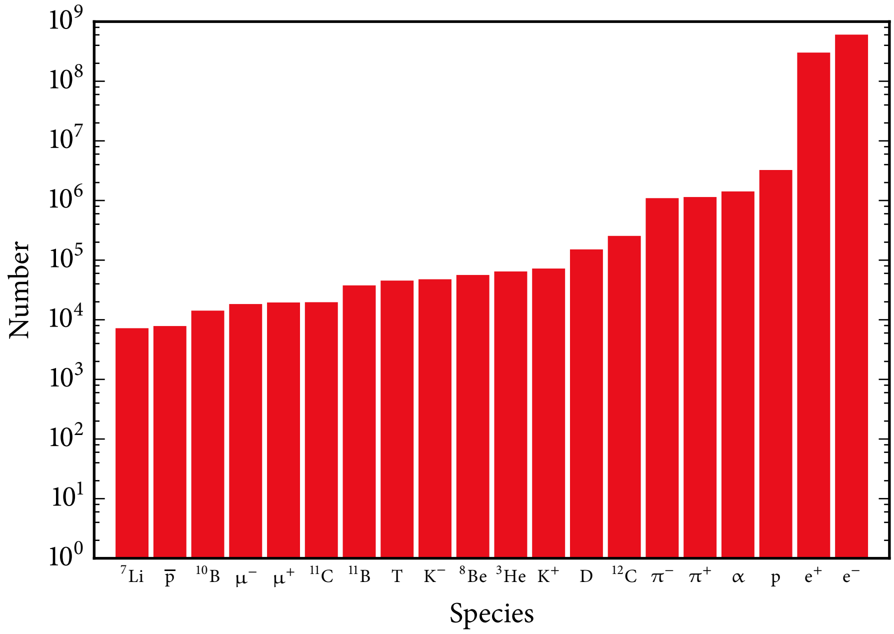

Figure 7.7: The number of each species of particle striking the scintillators in single particle simulations.

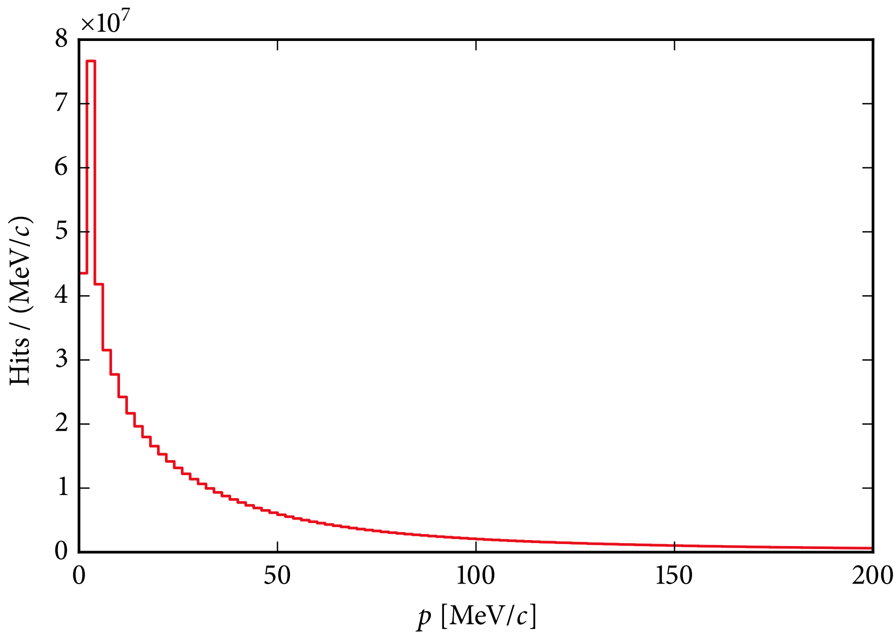

Figure 7.8: The momentum of particles striking the scintillators in single particle simulations.

For negative primary particle pseudorapidity the hits are required to be in B0, B1, or B2, whilst for positive pseudorapidity they are required to be in F1 or F2.

It is confirmed that the HERSCHEL stations increase the pseudorapidity coverage of the LHCb experiment beyond both the acceptance of the central detector, and the simple geometric acceptance of the five stations. The coverage extends to $|\eta| \geq 10$ in both the forward and backward directions, thus reaching the maximum pseudorapidity of primary particles in the simulated single and double diffractive events. This is greater than the geometric acceptance given in Table 6.2, as it is not necessary that the primary particle itself hits a scintillator for it to be detected. Primary particles outside the geometric acceptance of the stations are also detected when the shower resulting from them striking the aperture limit contains secondary particles which are incident upon a scintillator.

The detection efficiency is found to be non-uniform across the pseudorapidity range. In the backward direction (shown in the left column of Figs. 7.9 and 7.10), there is a gap in coverage of approximately 1.5 units between the maximum acceptance of the VELO and the minimum acceptance of the first station, B0, at $\eta \approx -5$. Here the detection efficiency with a single hit threshold is 0.2 to 0.5, depending on transverse momentum, for both π^+ and π^0 mesons. The efficiency then increases with pseudorapidity for both charged and neutral particles, plateauing at approximately 1 between −6 and −7. The efficiency drops to almost zero between −7 and −8, although the effect is less severe for neutral particles and those with lower transverse momentum. This corresponds geometrically to the primary particle striking the aperture limit within the quadrupole triplet, and it is hypothesised that the loss of efficiency is due to containment of the shower within the large `MBX` magnet, which is approximately 10 m long. This is consistent with preliminary simulations which included an additional station at −69 m, just beyond the `MBX` magnet, which exhibited very low hit multiplicities and in which a "shadow" in the shape of the `MBX` was visible in the transverse hit density profile.

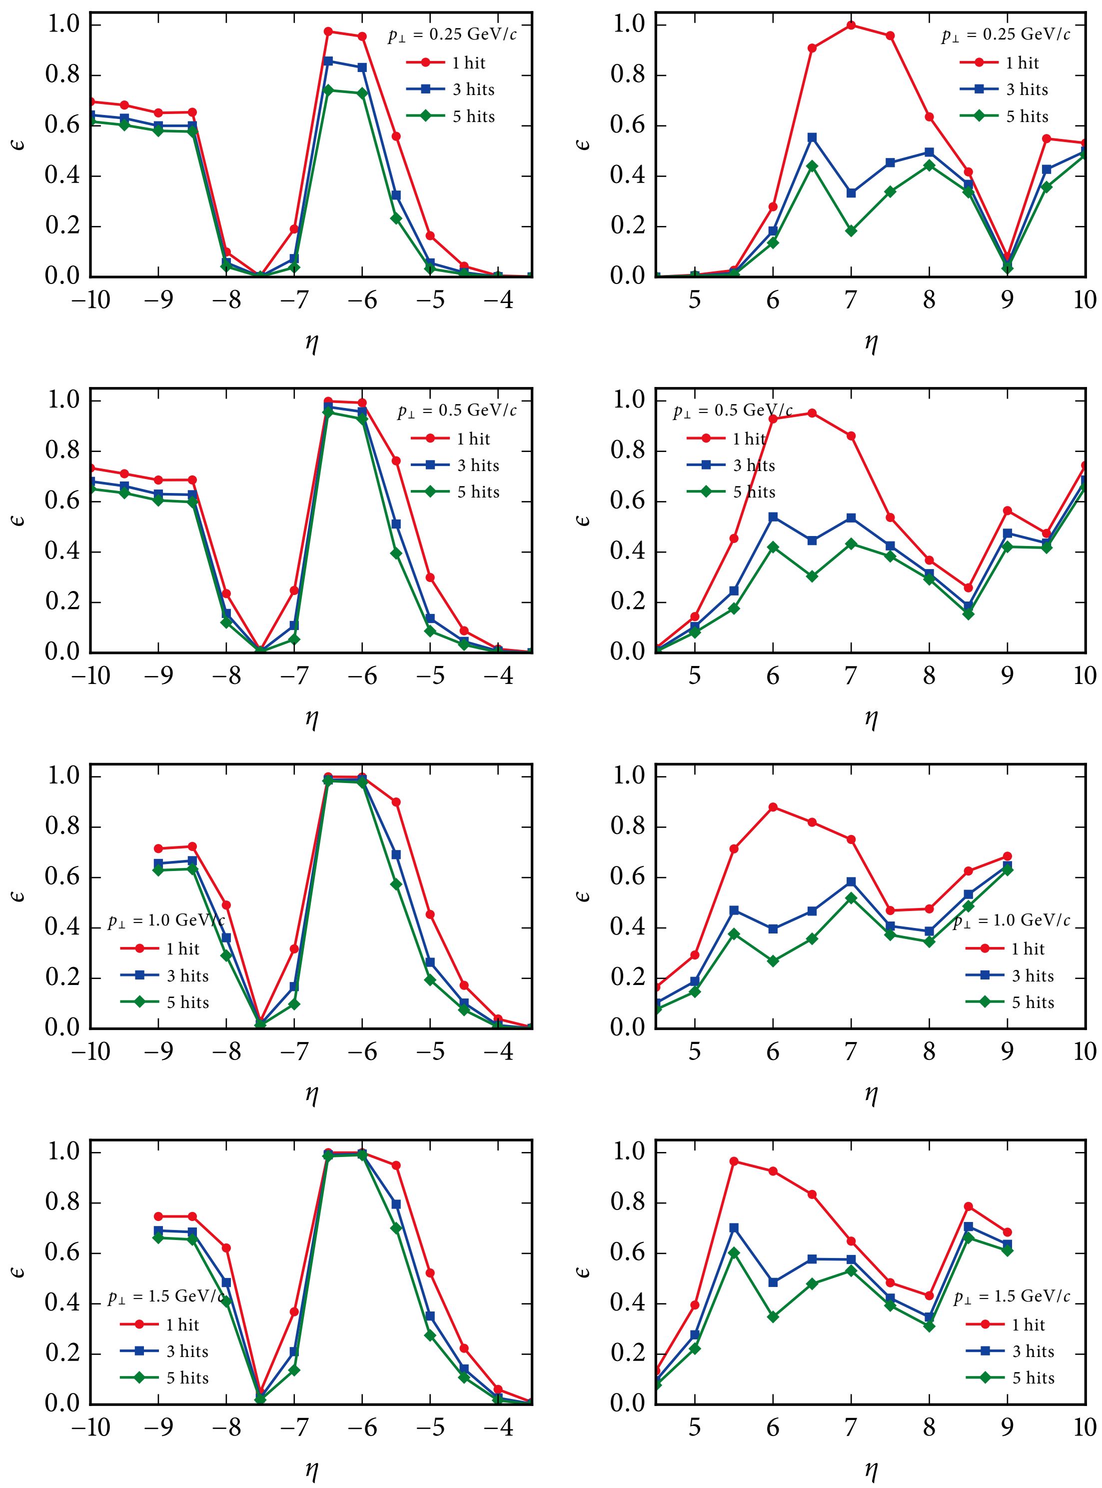

Figure 7.9: The detection efficiency for single π^+ mesons as a function of pseudorapidity assuming 100 % efficiency to detect each individual hit. For $\eta < 0$ (left column) hits are required in B0, B1, or B2, whilst for $\eta > 0$ (right column) hits are required in F1 or F2. Red graphs are for a detection threshold of one hit, blue graphs three hits, and green graphs five hits. Rows correspond to different $p_\perp$ values as the legends indicate.

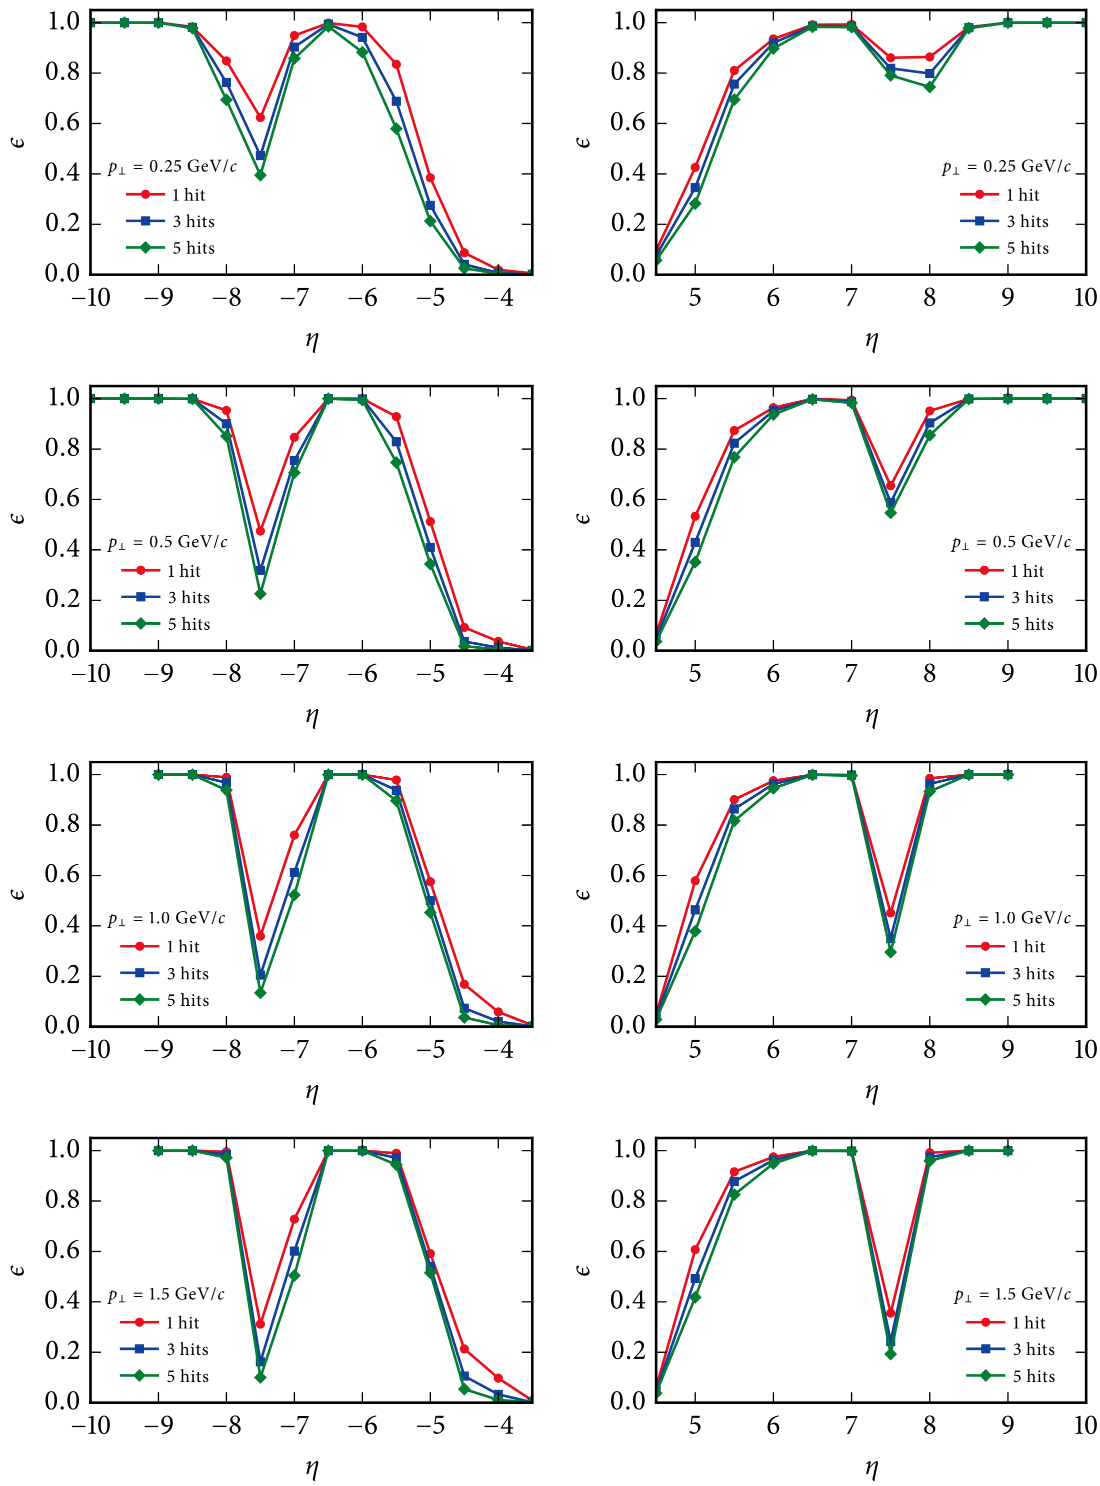

Figure 7.10: The detection efficiency for single π^0 mesons as a function of pion pseudorapidity, assuming 100 % efficiency to detect each individual hit. The layout, hit requirements, and meanings of the colours are the same as for Fig. 7.9.

For pseudorapidities below −8 the efficiency increases quickly again to approximately 0.7 for charged particles and 1 for neutral particles, as the primary particle then strikes the aperture limit outside the triplet and the shower secondaries strike B2. The high efficiency for neutral particles specifically with $\eta < -8.5$ is found to be due to showering in the BRAN converter directly in front of B2, as was intended by placing the station in this position.

The efficiency in the forward direction is qualitatively similar to that in the backward direction, but with quantitative differences resulting from the different beam line arrangement and the one fewer station on the forward side. Despite not having a station in the position mirroring B0, the gap between the maximum forward VELO acceptance and minimum acceptance of F1 is not larger than the gap between the maximum backward VELO acceptance and the minimum acceptance of B0. This is because the maximum VELO acceptance is one unit greater on the forward side than on the backward side, at $\eta = 4.5$ compared to $\eta = -3.5$.

The peak efficiency for charged particles with $5 \lesssim \eta \lesssim 7$ is reduced by there being only one station sensitive in this range on the forward side rather than two, however the acceptance range of the F1 station is larger than the B1 station due to there being no equivalent to the small `MBXWH` aperture which is present on the backward side. The region of $5 \lesssim \eta \lesssim 7$ is the only one where the efficiency is found to depend strongly on detection threshold, and therefore it will be especially important to suppress noise in F1 such that it can be operated with a low detection threshold to maximise the efficiency.

The detection efficiency is greater for particles with greater transverse momentum, which is expected as the total momentum is greater for a given pseudorapidity. The mean detection efficiencies, averaged over each combination of primary particle pseudorapidity and transverse momentum which was simulated, are given in Table 7.2. Additional simulations showed the detection efficiency for negatively charged pions to be the same as for positively charged pions.

Table 7.2: The mean detection efficiency for single π^+ and π^0 mesons for detection thresholds of one, three, and five hits, assuming 100 % efficiency to detect an individual hit. The average is over all values of pseudorapidity and transverse momentum simulated; that is pseudorapidity from the maximum VELO acceptance on each side to $|\eta| = 10$ in increments of 0.5 units, and transverse momenta of 0.25, 0.5, 1.0, and 1.5 GeV/c.

Detection threshold	Backward side		Forward side	
	π^+	π^0	π^+	π^0
One hit	0.47	0.68	0.55	0.81
Three hits	0.40	0.63	0.39	0.78
Five hits	0.36	0.61	0.33	0.76

The number of hits in each station aggregated across all the single particle simulations is plotted in Fig. 7.11, in order to determine the typical number of hits occurring from the showering of a single primary particle. The hit multiplicity increases with distance between the interaction point and the station, with a mean of 7 hits in B0, 8 hits in B1, and 411 hits in B2. On the forward side, the means are 35 hits in F1 and 396 hits in F2. In diffractive events with multiple primary particles, it is therefore expected that a single primary will typically result in around 10 hits in first two backward stations, 35 in F1, and 400 in B2 and F2. Having determined the pseudorapidity coverage of the stations and the hit multiplicities resulting from single primary particles, a more accurate determination of the expected multiplicities in diffractive events is made by simulating full proton-proton collisions.

7.4 Full event studies

Whilst single particle simulations enabled the pseudorapidity coverage of the stations and detection efficiency for single primary particles to be estimated, full event simulations are required to determine the kind of environment in which the detectors must operate. This was important to guide the design of the readout electronics, which required knowledge of the signals expected from the PMTs.

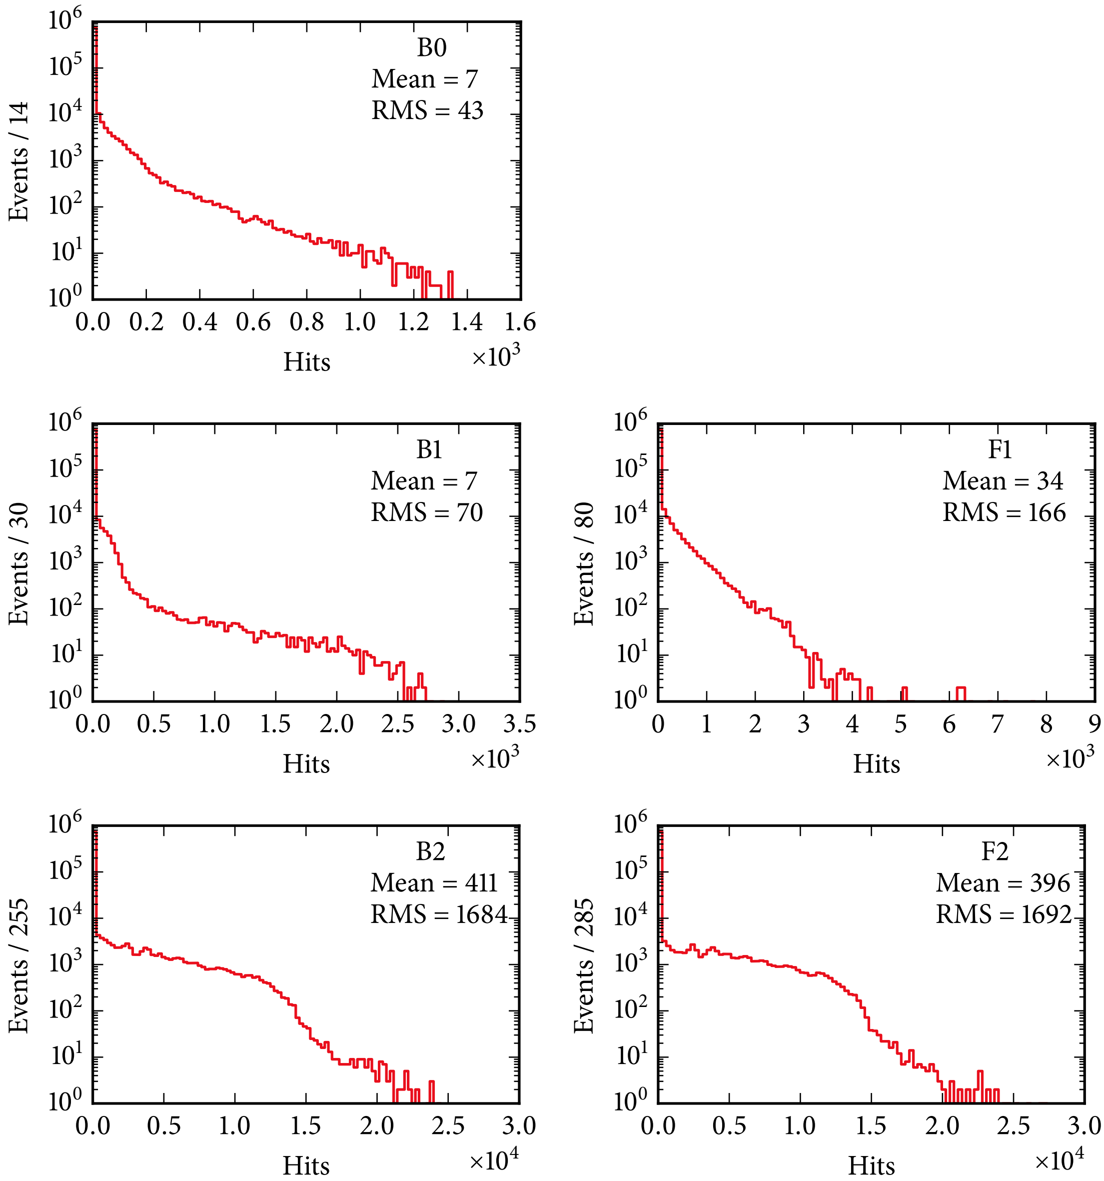

Figure 7.11: The number of hits in each station, plotted in aggregate for the single particle simulations of π^+ and π^0 mesons with the pseudorapidity and transverse momentum values described in the text.

Additionally, simulation of the HERSCHEL response to the type of diffractive events with proton dissociation which it is designed to identify in CEP analyses suggests where detection thresholds must be set to distinguish between the event types of interest—*i.e.* to provide sensitivity between CEP events and diffractive events with proton dissociation—and indicates the detection efficiency which may be expected in operation.

7.4.1 Simulation conditions

In order to carry out these studies, proton-proton collision events were simulated with PYTHIA and processed with the modified Gauss application as for the single particle studies. A sample of minimum bias events was generated using the expected Run 2 beam conditions listed in Table 7.1, containing 20,000 simulated bunch crossings. The sample was generated with pile-up, with the number of interactions per bunch crossing given by a Poisson distribution with a mean of 1.5. These events were used to study typical operating conditions, and determine the hit multiplicity, energy deposition, and resulting photoelectron yield at the PMTs in high multiplicity inelastic events.

In addition, a sample of single diffractive events was generated, split evenly between events with the dissociated proton travelling upstream from the IP and downstream from the IP, containing 20,000 bunch crossings in total. The purpose of this sample was to evaluate the efficiency for detecting events with proton dissociation. The generation of events with the dissociated proton moving both upstream and downstream allowed the sensitivity of both the backward and forward stations to be evaluated.

The number of primary particles with pseudorapidity beyond the maximum VELO acceptance ($\eta < -3.5$ or $\eta > 4.5$) for the minimum bias sample and the subset of the single diffractive sample without pileup is shown in Fig. 7.12. Only single diffractive events without pileup are considered as these are the events of interest which mimic the signature of CEP events in the LHCb detector; events with additional interactions will be rejected by the existing selection cuts described in Section 4.4. The mean minimum bias event as expected in typical Run 2 collisions has 46 primary particles beyond the VELO acceptance, compared to 13 for single diffractive events without pileup.

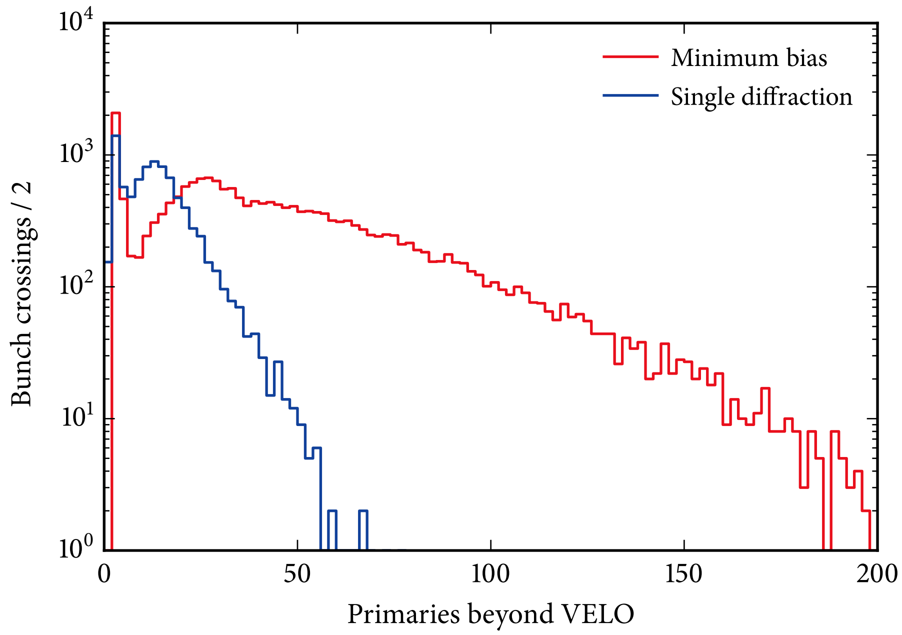

Figure 7.12: The number of primary particles beyond the maximum VELO acceptance in simulated minimum bias events (MB) and single diffractive events without pileup (SD).

The number of hits in the HERSCHEL stations expected in single diffractive events can then be naively estimated by multiplying this by the mean hit multiplicity resulting from a single primary particle, determined in Section 7.3. With 13 primary particles beyond the maximum VELO acceptance, approximately 100 hits are expected in B0 and B1, 500 hits in F1, and 5000 hits in B2 and F2. This estimate assumes an equal distribution of charged and neutral particles interacting in the same manner as pions, and which are uniformly distributed in (i) transverse momentum, and (ii) pseudorapidity across the acceptance of the stations. Recalling Fig. 7.2, it is justified to assume an approximately uniform pseudorapidity distribution out to $|\eta| = 10$, however Fig. 7.4 showed the transverse momentum of primary particles beyond the VELO acceptance to be strongly peaking below 0.5 GeV/c and so these hit multiplicities are likely to be overestimates.

7.4.2 Hit multiplicity

The actual number of hits in each station per bunch crossing taken from the full event simulations is shown in Fig. 7.13. The distribution is plotted for two classes of event: (i) all minimum bias events, labelled “MB”, and (ii) events from the single diffractive sample without pileup, where the dissociated proton is travelling towards the station (*i.e.* travelling upstream from the IP for the B stations and downstream for the F stations), labelled “SD”. Each histogram is normalised to 10,000 bunch crossings for comparison. There are no hits in the stations on the opposite side of the IP to the dissociated proton, as the intact proton on this side remains inside the beam pipe and within the aperture, and so does not shower.

In each of the five stations the hit multiplicity distributions are seen to fall off exponentially, with this behaviour common to both event classes. Despite the qualitatively similar behaviour for each class of events, the two distributions are distinguishable. It is helpful to consider the means and root mean squares for the purpose of comparison, and these are tabulated in Table 7.3. The mean hit multiplicity in each station is approximately twice as large in minimum bias events as in the single diffractive events. This is as expected from the inclusion of pileup in the minimum bias sample, and also the higher primary particle multiplicity at high absolute pseudorapidity seen in Fig. 7.12 in the minimum bias sample.

7.4.3 Energy deposition

Detection efficiencies have so far been evaluated in terms of the number of hits in the stations, assuming individual hits are detected with 100 % efficiency and neglecting readout effects. The efficiency of the stations is not directly dependent on the hit multiplicity however. Rather it is dependent on the energy deposited in the scintillator by those hits, as the scintillation

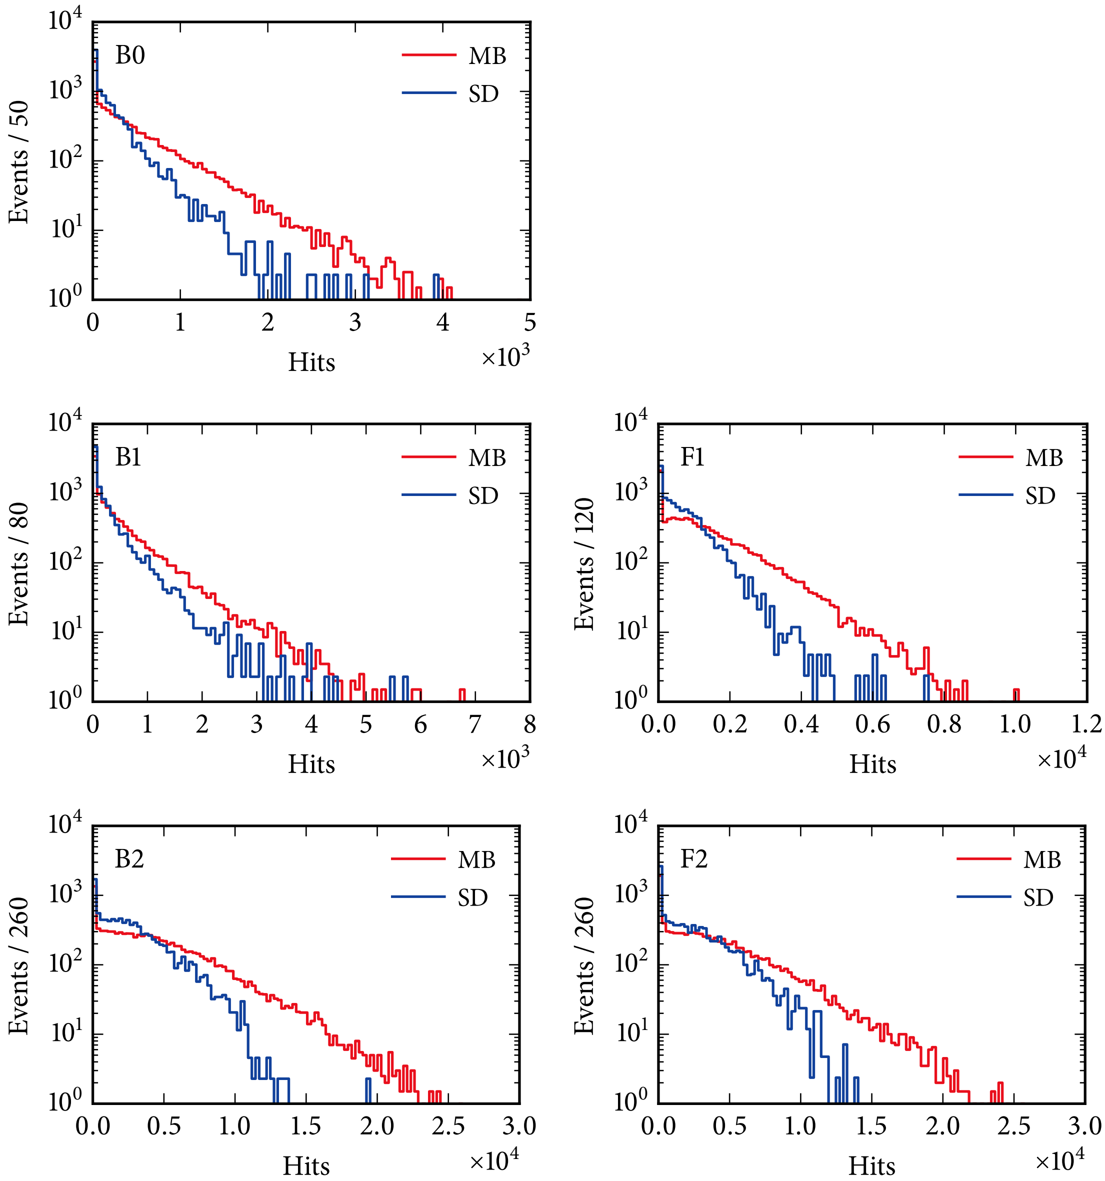

Figure 7.13: The number of hits per bunch crossing in each station. The red histogram (MB) is for minimum bias events and the blue histogram (SD) for all single diffractive events without pileup where the dissociated proton is travelling towards the station (*i.e.* upstream for the B stations and downstream for the F stations).

Table 7.3: The means and root mean squares of the number of hits in each station per bunch crossing plotted in Fig. 7.13.

Station	Event type	Mean	RMS
B0	MB	453	577
	SD	205	303
B1	MB	486	709
	SD	272	460
B2	MB	4399	4004
	SD	2652	2415
F1	MB	1301	1357
	SD	701	744
F2	MB	3905	3890
	SD	2415	2490

light yield is a function of the energy deposited by the particles striking the scintillator. A smaller number of hits which deposit a greater amount of energy results in a larger photoelectron yield than a larger number of hits with smaller total energy deposition.

The energy deposited in each HERSCHEL station per bunch crossing in the full event simulations is measured by summing the energy deposited by each hit (as determined by the interaction models within GEANT4) over all the hits in each station, and is plotted in Fig. 7.14. The deposited energy is plotted for the same event types as the hit multiplicity was in Fig. 7.13, and each histogram normalised to 10,000 bunch crossings.

The distributions are qualitatively similar to the distributions of hit multiplicity shown in Fig. 7.13, indicating similar energy deposition by individual primary particles in minimum bias events and single diffractive events. Like the hit multiplicity, the energy deposited in minimum bias events is approximately double that in single diffractive events, as shown by the means given in Table 7.4.

The optimal choice of PMT gain is constrained by the expected light yield, and thus by dependency the energy deposition, in minimum bias events. The gain must be set such that the maximum PMT anode current is not routinely exceeded by inelastic events with

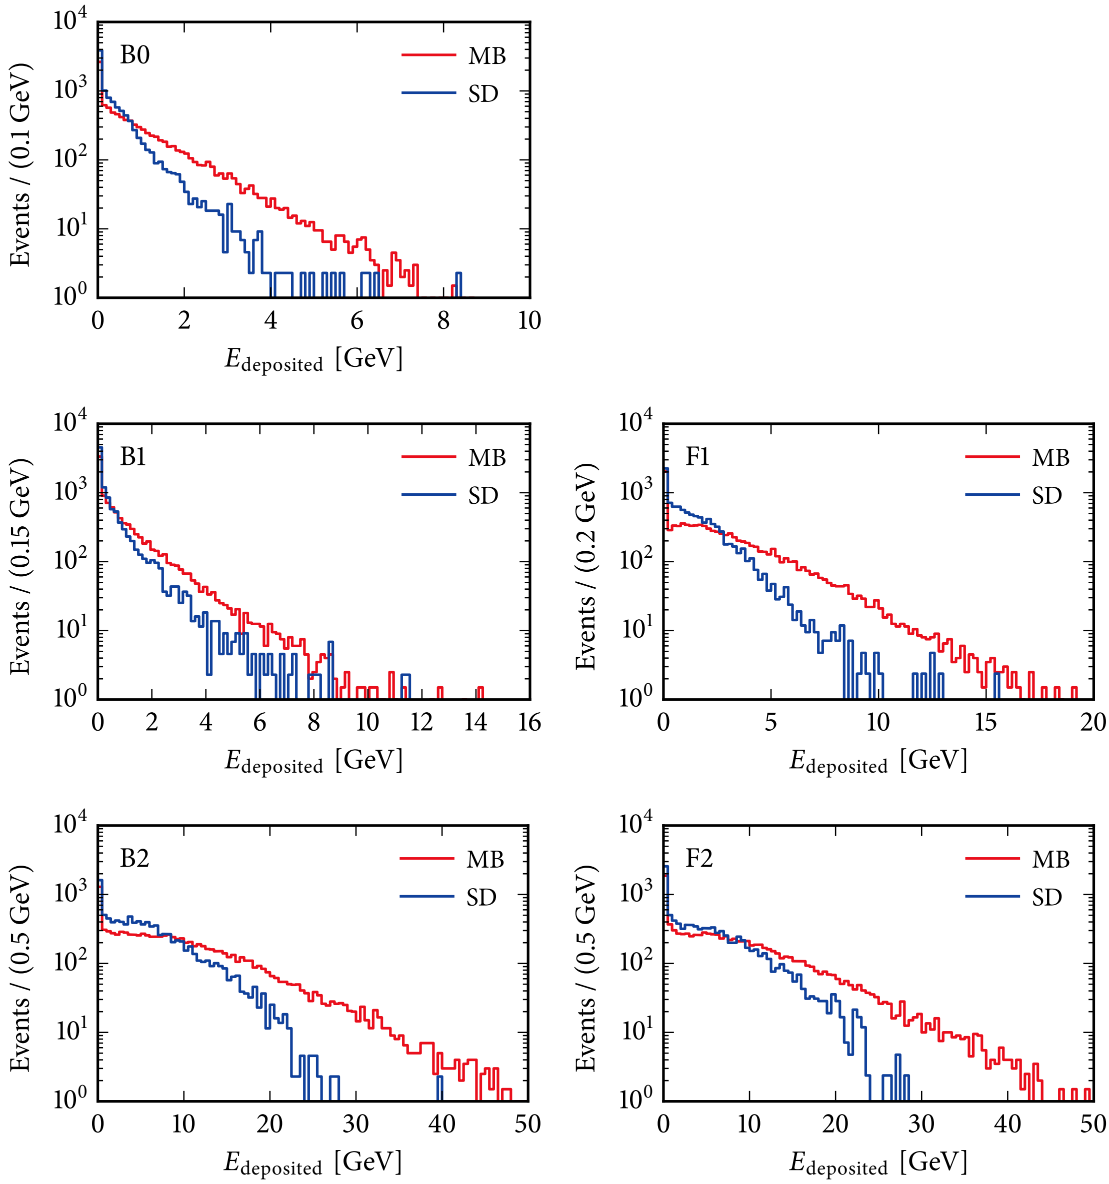

Figure 7.14: The energy deposited per bunch crossing in each station. The meaning of the legend entries is as in Fig. 7.13.

Table 7.4: The means and root mean squares of the energy deposition per bunch crossing in each station for minimum bias events (MB) and single diffractive events without pileup (SD).

Station	Event type	Mean [GeV]	RMS [GeV]
B0	MB	0.97	1.21
	SD	0.44	0.64
B1	MB	0.99	1.45
	SD	0.56	0.94
B2	MB	9.14	8.23
	SD	5.48	4.90
F1	MB	2.75	2.86
	SD	1.49	1.57
F2	MB	8.07	7.97
	SD	4.94	5.04

high primary multiplicity. The larger the difference between the energy deposition in high multiplicity inelastic events and diffractive events with proton dissociation, the less efficient HERSCHEL will be in detecting proton dissociation as the gain must be set low enough to cope with the high multiplicity events, leading to a small signal for proton dissociation.

That the difference between the energy deposited in single diffractive and minimum bias events is a factor of two, rather than a larger value, suggests that with a gain set appropriately to not exceed the maximum anode current in high multiplicity inelastic events, simultaneous sensitivity to very low levels of deposited energy (for example ten times lower, rather than two) is not necessary to detect diffractive events with proton dissociation with acceptable efficiency.

To determine if the light yield is equally distributed between the four PMTs on each station, the transverse density of deposited energy across the four scintillator quadrants is plotted in minimum bias events for each station in Fig. 7.15. The distributions for single diffractive events are of the same profile as those shown for minimum bias events, but with a lower absolute scale as expected from the mean deposited energy values given above.

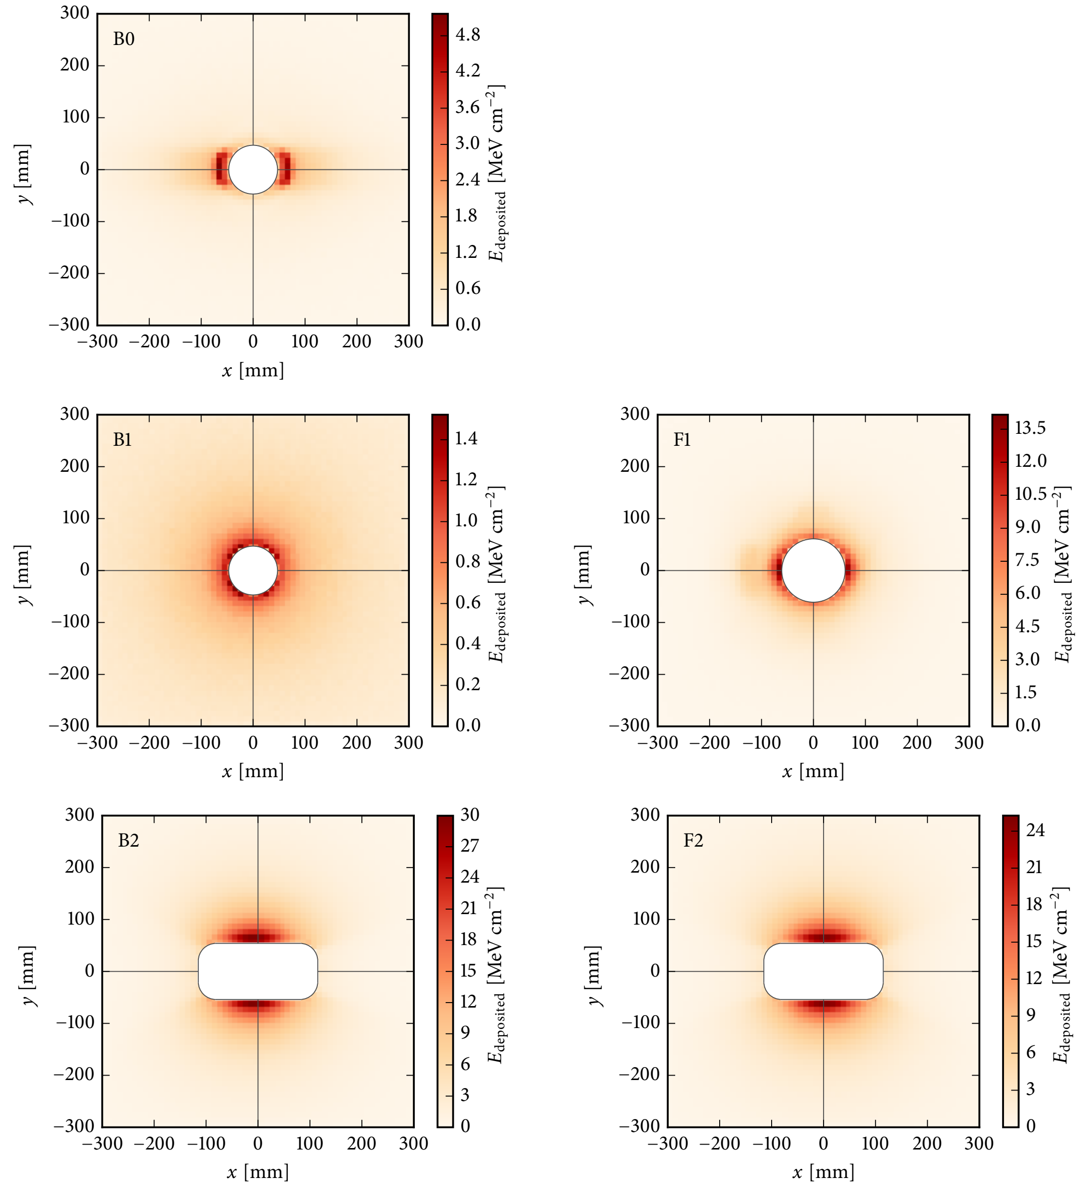

Figure 7.15: The mean energy deposition profile in the x-y plane for each station in simulated minimum bias events.

In all five stations the distribution is highly non-uniform, with most energy deposition concentrated within a 100 mm radius of the beam line centre. A lobe of high energy density can be seen either side of the beam pipe at $|x| \approx 80$ mm in B0, evidence of the horizontal beam crossing angle. The energy deposition is shared approximately equally between the four scintillator quadrants, with the distributions being symmetric around $x = 0$ and $y = 0$. It is expected therefore that each of the four PMTs on each station will collect a quarter of the light yield. The simulated energy deposition in B0 is consistent with measurements of the irradiation in this region of the tunnel, which were made during Run 1 using spare calorimeter modules placed there to study radiation damage to the calorimeter cells.

The greatest energy density is observed at the inner edge of the scintillators in B2 and F2, where the energy deposition reaches 30 MeV cm^{-2} in the mean minimum bias event, and so scintillator ageing will be most severe here.

7.4.4 Proton dissociation detection efficiency

In order to study the proton dissociation detection efficiency that may be achieved, the expected photoelectron yields have been estimated using the full event simulations. The energy deposited in the scintillator by a minimum ionising particle (MIP) is estimated as the modal energy deposited by charged pions traversing the entire thickness of the scintillator. This distribution is plotted in Fig. 7.16, from which the mode is measured as 3.6 MeV. It is therefore estimated that a MIP passing all the way through the station will deposit 3.6 MeV in the scintillator.

This value is then combined with the results of the electronics calibration carried out with cosmic muons described in Section 6.6, which measured the photoelectron yield to be 178 electrons per MIP. Taking the energy deposition of a single MIP to be 3.6 MeV from the simulation, and the MIP yield to be 178 photoelectrons from the calibration, the

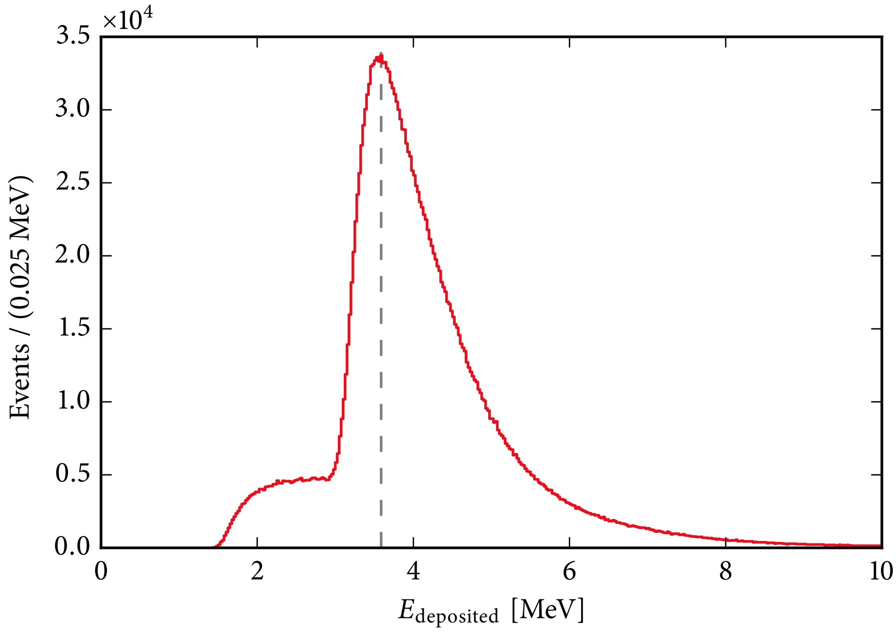

Figure 7.16: The energy deposited per hit aggregated across all stations, where the hit is from a charged pion which traverses the entire thickness of the scintillator. The energy deposited by a MIP is estimated as the mode, 3.6 MeV, indicated by the grey dashed line.

yield is estimated as 178/3.6 = 50 photoelectrons per MeV of energy deposited. A linear dependence of photoelectron yield on energy deposition is assumed.

The energy deposition per bunch crossing is scaled by this factor and the resulting photoelectron yield per bunch crossing in the two classes of events described above plotted in Fig. 7.17, with means and root mean squares listed in Table 7.5 for comparison.

Table 7.5: The means and root mean squares of the photoelectron yield per bunch crossing plotted in Fig. 7.17. The minimum bias values indicate the yield in typical operating conditions, and the single diffraction values the yield in typical events to be vetoed: it is not necessary to distinguish between SD and MB events.

Station	Event type	Mean	RMS
B0	MB	47,891	60,144
	SD	21,798	31,878
B1	MB	49,297	72,162
	SD	34,136	47,445
B2	MB	453,545	408,389
	SD	271,869	243,046
F1	MB	136,665	141,752
	SD	73,983	77,699
F2	MB	400,183	395,351
	SD	245,176	250,307

These distributions were used to estimate the efficiency for detecting the dissociated proton in single diffractive events as a function of the photoelectron threshold, and this is plotted in Fig. 7.18. The detection efficiency is calculated as the fraction of single diffractive events which with a photoelectron yield greater than a given threshold. Note that for the primary application of HERSCHEL in vetoing events with proton dissociation, it is not required to distinguish between single diffractive events and other types of inelastic events contained within the minimum bias sample, and no attempt to do so is made. It is seen in Fig. 7.17 and Table 7.5 that the photoelectron yield is not uniform across the subset of stations on one side of the interaction point, therefore it is sub-optimal to use a constant threshold number of photoelectrons across the stations. Instead, the detection threshold is

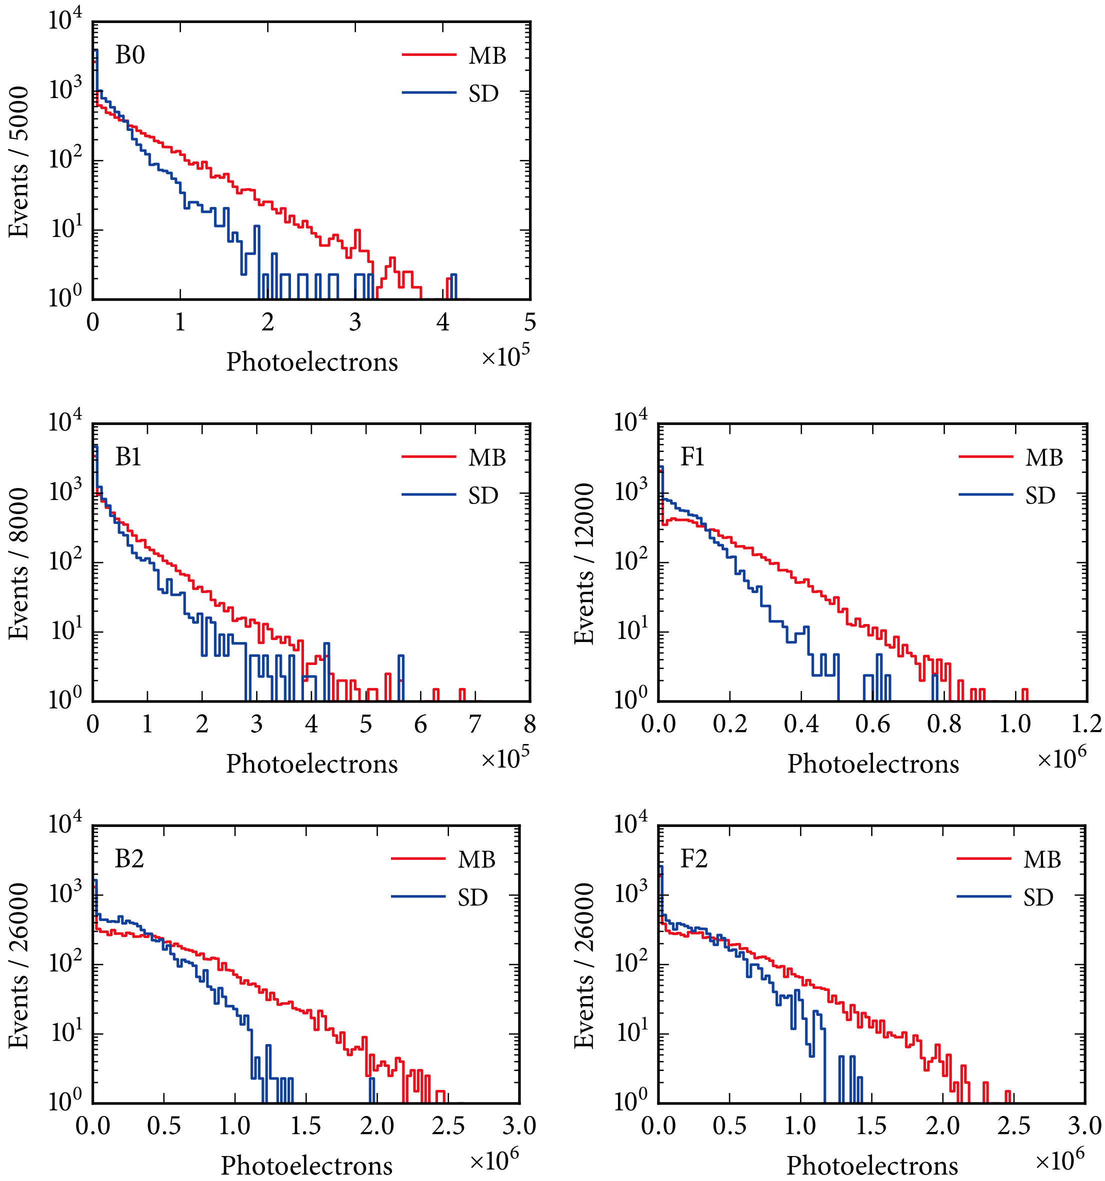

Figure 7.17: The photoelectron yield per bunch crossing for the two event types considered, calculated as described in text using the cosmic MIP calibration and modal simulated energy deposition.

expressed as a scalar factor by which the mean photoelectron yield in minimum bias events in each station, given in Table 7.5, is multiplied.

Consider the detection efficiency for a dissociated proton travelling downstream from the IP, towards the forward stations. The mean number of photoelectrons in minimum bias events is 136,665 in F1 and 400,183 in F2. A threshold factor of 0.5, for example, then corresponds to an absolute threshold of 68,333 photoelectrons in F1 and 200,092 in F2. The single diffractive event is considered detected if the photoelectron yield in either of the F1 and F2 stations exceeds its respective threshold value. From Fig. 7.18 it is estimated that with this threshold 73 % of single diffractive events with proton dissociation on the forward side would be detected with HERSCHEL, and could therefore be vetoed from selection as CEP candidates. In the backward direction, the detection efficiency for single diffractive events is estimated as 76 % for the same threshold factor. If it were possible to set the threshold factor to 0.2, the corresponding single diffraction detection efficiencies are 89 % on the forward side and 91 % on the backward side.

This suggests that the HERSCHEL system will provide an effective veto for diffractive events with proton dissociation. It is essential to consider, however, that these results present an idealised case. In operation, there is certain to be noise which is not modelled, and additional sources of energy deposition that are not included in the simulation, such as spillover and machine induced backgrounds. It is these factors, in addition to those coming from the pulse shape discussed below, which determine how low the threshold should be set. The lower the threshold, the greater the efficiency for detecting background events, however it should not be set so low that genuine CEP events trigger the veto.

It could be attempted to suppress some of these backgrounds, for example by adding shielding on the side of scintillators facing away from the interaction point for machine induced backgrounds, however it is essential that the threshold be calibrated using real data. Additionally, it may be expected that the subset of single diffractive events which

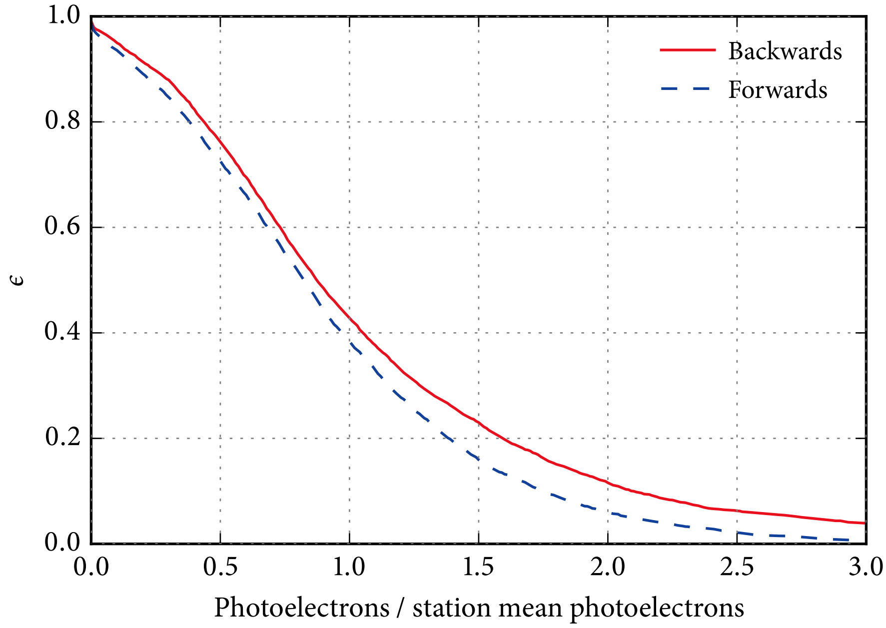

Figure 7.18: The efficiency to detect single diffractive events as a function of the threshold number of photoelectrons produced at the PMTs. The threshold is expressed as a scalar factor by which the mean number of photoelectrons in each station is multiplied as described in the text.

most closely mimic the signature of CEP events within the LHCb detector have a lower energy deposition in the HERSCHEL stations than the average single diffractive event, due to a reduced number of primary particles due to the track multiplicity cuts imposed in the selection, leading to a lower proton dissociation detection efficiency than determined above.

7.4.5 Pulse shape considerations

The detection efficiency is evaluated above in terms of the photoelectron yield rather than the readout signal from the front-end electronics and so does not account for effects such as the shape of the signal pulse, or dispersion in the cables, for example.

The shape of the signal pulse is particularly important due to the reduction in the bunch crossing interval from 50 ns to 25 ns in Run 2. The signal pulse must be contained within a 25 ns window, such that the pulse from a high multiplicity inelastic event does not spillover into the next bunch crossing, which might mask any signal from the following collision and prevent an event with proton dissociation being detected.

Pulse containment within 25 ns is achieved using a clipping line on each PMT, composed of a coaxial cable terminated with an RC circuit. Part of each signal pulse is sent down the clipping line, and inverted. The reflection arrives back at the PMT after a time delay determined by the length of the cable, which was optimised during the cosmic calibrations. The reflected pulse sums with the tail of the original pulse to approximately zero, so "clipping" off the tail of the signal pulse beyond 25 ns. Pulses from a single PMT due to cosmic muons striking the scintillator are shown in Fig. 7.19 both with and without clipping. Without clipping, the pulse lasts more than 30 ns, spilling over into the next event, and the signal amplitude is still at 30 % of its maximum after 25 ns. With the clipping line, the pulse is contained within 25 ns, with a small return-to-baseline effect.

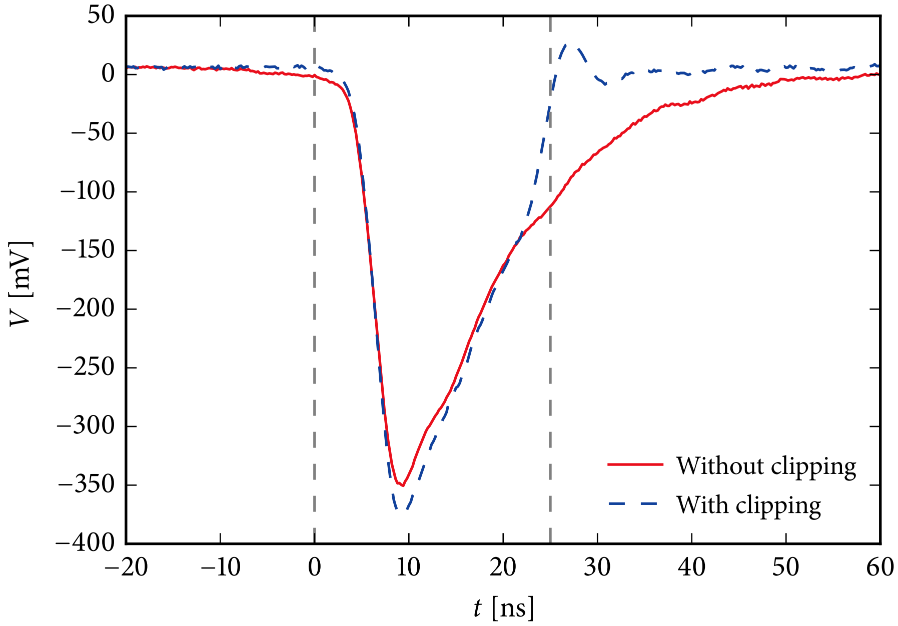

Figure 7.19: PMT signal pulses from cosmic muons with and without clipping. The vertical grey dashed lines indicate the trigger at 0 ns, and the 25 ns mark. It is observed that without clipping the pulse extends beyond 25 ns, whereas with clipping the pulse is contained within 25 ns, with a small return-to-baseline effect.

7.5 Summary

Simulations have been carried out of the performance of the HERSCHEL system, making use of the geometry and magnetic field model of Interaction Region 8 described in Chapter 6. Single particle simulations have been carried out to determine the pseudorapidity coverage of the stations, and map the detection efficiency as a function of pseudorapidity. The efficiency is determined for both charged and neutral pions with kinematics similar to those of primary particles in single and double diffractive events. It is confirmed that hits in the scintillators are due to the showering of primary particles when they reach the aperture limit close to the station in which the hits are detected.

Combining information from the three backward stations and the two forward stations, the pseudorapidity coverage is found to extend from $|\eta| \approx 5$ to $|\eta| \approx 10$. There is however a severe reduction in efficiency in the region $7 \lesssim |\eta| \lesssim 8$ which is hypothesised to be due to the absorption of showers in the `MBX` magnet. This is most severe on the backward side. Averaged over the pseudorapidity and transverse momentum of the primary particles, the efficiency for a detection threshold of a single hit is 50 % to 70 % on the backward side and 55 % to 80 % on the forward side, where the lower value is for charged particles and the higher value for neutral particles.

In addition to the single particle efficiency studies, simulations of full proton-proton collisions were carried out in order to understand the expected operating environment and estimate the efficiency with which diffractive events with proton dissociation might be detected in the ideal case. The hit multiplicity and energy deposition in the scintillators in single diffractive events is found to be on average half that in minimum bias events. The energy deposition is concentrated around the beam pipe and distributed equally between the four scintillator quadrants, with a peak of $30\,\mathrm{MeV\,cm^{-2}}$ in the furthest stations.

The simulated energy deposition is combined with the MIP photoelectron yield measured with cosmic muons to estimate the detection efficiency for diffractive events with a dissociated proton. For an example threshold of 0.5 times the mean photoelectron yield in minimum bias events, the detection efficiency is determined to be greater than 70 % on both the forward and backward sides. Reducing the threshold factor to 0.2 increases the detection efficiency to 90 %. It is important to note, however, that these results are dependent on the PYTHIA models of both diffractive dissociation and the other processes which constitute the minimum bias sample. Additionally, the simulation does not include all sources of background hits nor account for the signal response, which will affect where the detection threshold should be set in operation. The optimum threshold is that which provides the maximum efficiency for detecting diffractive events with proton dissociation whilst not triggering on genuine CEP events.

8

Summary and outlook

This thesis presents a measurement of the cross section for central exclusive $\chi_c(1P)$ production, using data collected with the LHCb detector in 2012 at a centre-of-mass energy of $\sqrt{s} = 8$ TeV. Candidates are reconstructed in the K^+K^- and $\pi^+\pi^-$ final states and are selected by applying tight multiplicity cuts, selecting events without multiple interactions. The efficiency of the selection is determined from simulated events, with corrections made using data-driven methods where the simulation does not adequately model collision data.

As bunch crossings with more than a single interaction are rejected by the selection, the effective integrated luminosity analysed is less than the total integrated luminosity. The effective integrated luminosity is calculated using the fraction of bunch crossings with no interactions, measured using a VELO track counter.

A class of inelastic backgrounds similar to CEP but with proton dissociation or additional gluon radiation in the final state cannot be distinguished from signal events using selection cuts alone. The purity of the selected data set—the fraction corresponding to genuine CEP events rather than these backgrounds—is determined by fitting the transverse momentum distribution of selected $\chi_c(1P)$ candidates. The signal component shape is taken from simulation and the background component shape extrapolated from higher multiplicity events.

The products of the cross section and branching fraction for daughters in the pseudorapidity range $2.5 < \eta < 4.5$ are measured as

$$\sigma \times \mathcal{B}(\chi_{c0}(1P) \to K^+K^-) = 28.0 \pm 2.8 \pm 13.9\,\text{pb},$$
$$\sigma \times \mathcal{B}(\chi_{c0}(1P) \to \pi^+\pi^-) = 20.3 \pm 2.4 \pm 8.1\,\text{pb},$$
$$\sigma \times \mathcal{B}(\chi_{c2}(1P) \to K^+K^-) < 12.1\,\text{pb},$$
$$\sigma \times \mathcal{B}(\chi_{c2}(1P) \to \pi^+\pi^-) < 10.5\,\text{pb},$$

where the first uncertainty is statistical and the second systematic, and limits are set at the 95 % confidence level. The $\chi_{c2}(1P)$ limits are an order of magnitude greater than the theory predictions and so are not sensitive to the central values. The $\chi_{c0}(1P)$ measurements are significantly lower than the theoretical prediction of the Durham model of CEP, although not inconsistent given the factor of five uncertainty on the prediction. They are in good agreement with the results of previous LHCb analyses of $\chi_{c0}(1P) \to \mu^+\mu^-\gamma$ CEP, which also measured $\chi_{c0}(1P)$ cross sections below the predictions. This provides a validation of those measurements with hadronic final states in which the spin states of the $\chi_c(1P)$ are separated in the invariant mass distributions.

A major source of systematic uncertainty in both the h^+h^- and $\mu^+\mu^-\gamma$ measurements is the exclusive purity determination. This is evaluated by fitting the transverse momentum distribution of the $\chi_c(1P)$ candidate, however to make progress in probing the tension between the measured cross sections and the theoretical prediction, a method of eliminating the systematic uncertainty this technique introduces is required.

This need led to the development and installation of HERSCHEL, a system of scintillators located in the LHC tunnel either side of the LHCb detector. The scintillators allow detecting inelastic background events on an event-by-event basis by detecting the particle showers initiated when the remnants of a dissociated proton strike the aperture limit. Suitable locations for the installation of HERSCHEL stations in the LHC tunnel were identified. A geometric model of the LHC tunnel within 125 m of the LHCb detector, including magnetic field maps, was created and integrated into Gauss, the LHCb simulation application.

Single particle simulations to determine the pseudorapidity coverage of the HERSCHEL stations are presented, and it is found that due to showering the coverage extends beyond the geometric acceptance of the stations to $|\eta| \geq 10$. The single particle detection efficiency is determined for various values of the primary particle's pseudorapidity and transverse momentum, for both charged and neutral particles.

Full event studies are described in which minimum bias and single diffractive events are simulated in order to determine the environment in which the detectors must operate, and to predict the efficiency with which diffractive events with proton dissociation can be detected. With a threshold for detection of 0.5 times the number of photoelectrons in a typical minimum bias event, the predicted detection efficiency for single diffractive events is 73 % in the forward direction and 76 % in the backward direction. If the detection threshold is lowered to 0.2, these detection efficiencies increase to 89 % and 91 % respectively.

Once the response of the HERSCHEL detectors has been understood and appropriate detection thresholds set, they are expected to provide an efficient veto of inelastic backgrounds and significantly reduce the systematic uncertainty on CEP measurements. In addition to vetoing inelastic backgrounds in CEP analyses, HERSCHEL makes possible isolating samples of single and double diffractive events for study in their own right, which was not previously possible.

At a later stage information from HERSCHEL could be used in the trigger, allowing the pre-scales to be relaxed. In concert with the move from a 50 ns to 25 ns bunch spacing for Run 2, which doubles the fraction of bunch crossings without multiple interactions for a given luminosity, this will allow a much larger sample of CEP candidates to be collected.

In addition, use of information from HERSCHEL in the trigger may permit simultaneously relaxing the cut on SPD multiplicity, which would allow the selection of higher multiplicity CEP decays, such as the four body decay modes of the $\chi_c(1P)$ which would allow access to the $\chi_{c1}(1P)$ which does not decay to K^+K^- or $\pi^+\pi^-$. Other decays would also be accessible, for example $\psi(3770) \to D\overline{D}$, extending the reach of the LHCb physics programme.

The first CEP dihadron candidates have already been observed in early $\sqrt{s} = 13\,\text{TeV}$ collisions at LHCb, one of which is illustrated in Fig. 8.1. Run 2 of the LHC offers bright prospects for the study of CEP at LHCb, probing a previously unexplored energy regime with much improved precision with the addition of HERSCHEL.

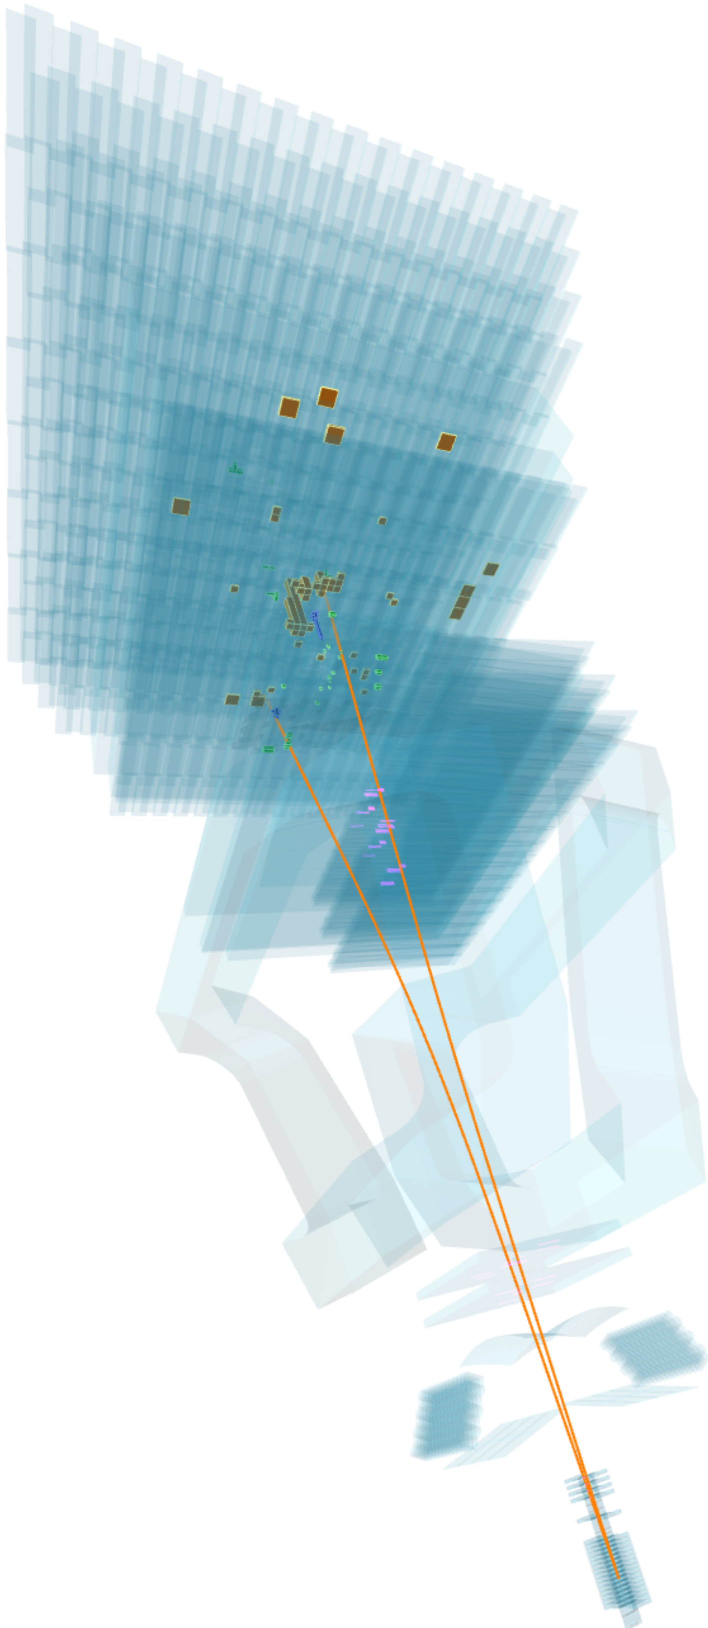

Figure 8.1: A candidate dihadron central exclusive production event reconstructed in the LHCb detector during early Run 2 collisions at $\sqrt{s} = 13$ TeV.

Bibliography

[1] I. J. R. Aitchison and A. J. G. Hey, *Gauge Theories in Particle Physics*, vol. I: From Relativistic Quantum Mechanics to QED of *Graduate Student Series in Physics*, IOP Publishing, 3rd ed., 2003.

[2] B. R. Martin and G. Shaw, *Particle Physics*, Manchester Physics Series, Wiley, 3rd ed., 2008.

[3] F. Halzen and A. D. Martin, *Quarks and Leptons: An Introductory Course in Modern Particle Physics*, Wiley, 1984.

[4] LHCb collaboration, R. Aaij *et al.*, *Observation of the resonant character of the Z(4430)$^-$ state*, Phys. Rev. Lett. **112** (2014) 222002, arXiv:1404.1903.

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

[6] C. Burgess and G. Moore, *The Standard Model: A Primer*, Cambridge University Press, 2007.

[7] S. Donnachie, G. Dosch, P. Landshoff, and O. Nachtmann, *Pomeron Physics and QCD*, Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology, Cambridge University Press, 2002.

[8] R. K. Ellis, W. J. Stirling, and B. R. Webber, *QCD and Collider Physics*, Cambridge Monographs on Particle Physics, Nuclear Physics and Cosmology, Cambridge University Press, 1996.

[9] V. Barone and E. Predazzi, *High-Energy Particle Diffraction*, Theoretical and Mathematical Physics, Springer, 2002.

[10] J. D. Bjorken, *Hard Diffraction and Deep Inelastic Scattering*, in *Proceedings of the International Workshop on Deep Inelastic Scattering and Related Subjects*, 1994.

[11] T. Regge, *Introduction to complex orbital momenta*, Nuovo Cim. **14** (1959) 951.

[12] T. Regge, *Bound states, shadow states and Mandelstam representation*, Nuovo Cim. **18** (1960) 947.

[13] P. D. B. Collins, *An Introduction to Regge Theory and High-Energy Physics*, Cambridge Monographs on Mathematical Physics, Cambridge University Press, 1977.

[14] J. R. Forshaw and D. A. Ross, *Quantum Chromodynamics and the Pomeron*, Cambridge Lecture Notes in Physics, Cambridge University Press, 1997.

[15] M. Albrow, *Central exclusive production issue: Introduction*, Int. J. Mod. Phys. A **29** (2014) 1402006.

[16] V. A. Khoze, A. D. Martin, and M. G. Ryskin, *Double diffractive processes in high resolution missing mass experiments at the Tevatron*, Eur. Phys. J. **C19** (2001) 477, arXiv:hep-ph/0011393.

[17] V. A. Khoze, A. D. Martin, and M. G. Ryskin, *Can the Higgs be seen in rapidity gap events at the Tevatron or the LHC?*, Eur. Phys. J. **C14** (2000) 525, arXiv:hep-ph/0002072.

[18] M. G. Albrow and A. Rostovtsev, *Searching for the Higgs at hadron colliders using the missing mass method*, arXiv:hep-ph/0009336, FERMILAB-PUB-00-173.

[19] A. De Roeck *et al.*, *Ways to detect a light Higgs boson at the LHC*, Eur. Phys. J. **C25** (2002) 391, arXiv:hep-ph/0207042.

[20] D. Robson, *A Basic Guide for the Glueball Spotter*, Nucl. Phys. **B130** (1977) 328.

[21] F. E. Close, *Gluonic Hadrons*, Rept. Prog. Phys. **51** (1988) 833.

[22] AFS collaboration, T. Akesson *et al.*, *A Search for Glueballs and a Study of Double Pomeron Exchange at the CERN Interacting Storage Rings*, Nucl. Phys. **B264** (1986) 154.

[23] F. E. Close and A. Kirk, *A Glueball-$q\bar{q}$ filter in central hadron production*, Phys. Lett. **B397** (1997) 333, arXiv:hep-ph/9701222.

[24] A. B. Kaidalov, V. A. Khoze, A. D. Martin, and M. G. Ryskin, *Central exclusive diffractive production as a spin-parity analyser: From Hadrons to Higgs*, Eur. Phys. J. **C31** (2003) 387, arXiv:hep-ph/0307064.

[25] A. B. Kaidalov, V. A. Khoze, A. D. Martin, and M. G. Ryskin, *Extending the study of the Higgs sector at the LHC by proton tagging*, Eur. Phys. J. **C33** (2004) 261, arXiv:hep-ph/0311023.

[26] S. Heinemeyer *et al.*, *Studying the MSSM Higgs sector by forward proton tagging at the LHC*, Eur. Phys. J. **C53** (2008) 231, arXiv:0708.3052.

[27] V. A. Khoze, A. D. Martin, and M. G. Ryskin, *Prospects for new physics observations in diffractive processes at the LHC and Tevatron*, Eur. Phys. J. **C23** (2002) 311, arXiv:hep-ph/0111078.

[28] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, and W. J. Stirling, *Central exclusive χ_c meson production at the Tevatron revisited*, Eur. Phys. J. **C65** (2010) 433, arXiv:0909.4748.

[29] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, and W. J. Stirling, *Standard candle central exclusive processes at the Tevatron and LHC*, Eur. Phys. J. **C69** (2010) 179, arXiv:1005.0695.

[30] V. A. Khoze, A. D. Martin, M. G. Ryskin, and W. J. Stirling, *Double diffractive χ meson production at the hadron colliders*, Eur. Phys. J. **C35** (2004) 211, arXiv:hep-ph/0403218.

[31] D. G. d'Enterria, *Forward Physics at the LHC*, in *Proceedings of the 15th International Workshop on Deep-inelastic scattering and related subjects*, pp. 1141–1152, 2007. arXiv:0708.0551. doi: 10.3360/dis.2007.221.

[32] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, and W. J. Stirling, *Central exclusive production within the Durham model: a review*, Int. J. Mod. Phys. **A29** (2014) 1430031, arXiv:1405.0018.

[33] L. A. Harland-Lang, V. A. Khoze, and M. G. Ryskin, *Central exclusive production and the Durham model*, Int. J. Mod. Phys. **A29** (2014) 1446004.

[34] L. A. Harland-Lang, *Simple form for the low-x generalized parton distributions in the skewed regime*, Phys. Rev. **D88** (2013), no. 3 034029, arXiv:1306.6661.

[35] V. A. Khoze, A. D. Martin, and M. G. Ryskin, *Diffraction at the LHC*, Eur. Phys. J. **C73** (2013) 2503, arXiv:1306.2149.

[36] V. A. Khoze, A. D. Martin, and M. G. Ryskin, *High Energy Elastic and Diffractive Cross Sections*, Eur. Phys. J. **C74** (2014), no. 2 2756, arXiv:1312.3851.

[37] M. G. Ryskin, A. D. Martin, and V. A. Khoze, *High-energy strong interactions: from 'hard' to 'soft'*, Eur. Phys. J. **C71** (2011) 1617, arXiv:1102.2844.

[38] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, and W. J. Stirling, *Central exclusive meson pair production in the perturbative regime at hadron colliders*, Eur. Phys. J. **C71** (2011) 1714, arXiv:1105.1626.

[39] CDF collaboration, T. Aaltonen *et al.*, *Observation of exclusive charmonium production and $\gamma\gamma \to \mu^+\mu^-$ in $p\overline{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, Phys. Rev. Lett. **102** (2009) 242001, arXiv:0902.1271.

[40] LHCb collaboration, *Central Exclusive Dimuon Production at $\sqrt{s} = 7$ TeV*, LHCb-CONF-2011-022, https://cds.cern.ch/record/1349022, 2011.

[41] P. M. Bjørnstad, *Central Exclusive Production of χ Mesons at LHCb*, PhD thesis, University of Manchester, 2013, https://cds.cern.ch/record/1632249.

[42] L. A. Harland-Lang, V. A. Khoze, and M. G. Ryskin, *A perturbative approach to central exclusive production*, Int. J. Mod. Phys. **A30** (2015) 1542013.

[43] T. Massam *et al.*, *Experimental observation of antideuteron production*, Il Nuovo Cimento **39** (1965), no. 1 10.

[44] D. E. Dorfan *et al.*, *Observation of antideuterons*, Phys. Rev. Lett. **14** (1965) 1003.

[45] F. J. Hasert *et al.*, *Observation of Neutrino Like Interactions without Muon or Electron in the Gargamelle Neutrino Experiment*, Nucl. Phys. **B73** (1974) 1.

[46] UA1 collaboration, G. Arnison *et al.*, *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV*, Phys. Lett. **B122** (1983) 103.

[47] UA2 collaboration, M. Banner *et al.*, *Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN $\bar{p}p$ Collider*, Phys. Lett. **B122** (1983) 476.

[48] UA1 collaboration, G. Arnison *et al.*, *Experimental Observation of Lepton Pairs of Invariant Mass Around 95 GeV/c^2 at the CERN SPS Collider*, Phys. Lett. **B126** (1983) 398.

[49] UA2 collaboration, P. Bagnaia *et al.*, *Evidence for $Z^0 \to e^+e^-$ at the CERN $\bar{p}p$ Collider*, Phys. Lett. **B129** (1983) 130.

[50] ATLAS collaboration, G. Aad *et al.*, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. **B716** (2012) 1, `arXiv:1207.7214`.

[51] CMS collaboration, S. Chatrchyan *et al.*, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys. Lett. **B716** (2012) 30, `arXiv:1207.7235`.

[52] O. S. Bruning *et al.*, *LHC Design Report*, vol. 1: The LHC Main Ring, CERN, Geneva, 2004. doi: 10.5170/CERN-2004-003-V-1.

[53] O. S. Bruning *et al.*, *LHC Design Report*, vol. 2: The LHC infrastructure and general services, CERN, Geneva, 2004. doi: 10.5170/CERN-2004-003-V-2.

[54] M. Benedikt *et al.*, *LHC Design Report*, vol. 3: The LHC injector chain, CERN, Geneva, 2004. doi: 10.5170/CERN-2004-003-V-3.

[55] L. Forthomme, *Map of the CERN accelerator complex*, Wikimedia Commons, `https://commons.wikimedia.org/wiki/File:Cern-accelerator-complex.svg`, 2011.

[56] ATLAS collaboration, G. Aad *et al.*, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003.

[57] CMS collaboration, S. Chatrchyan *et al.*, *The CMS experiment at the CERN LHC*, JINST **3** (2008) S08004.

[58] ALICE collaboration, K. Aamodt *et al.*, *The ALICE experiment at the CERN LHC*, JINST **3** (2008) S08002.

[59] LHCb collaboration, A. A. Alves Jr. *et al.*, *The LHCb Detector at the LHC*, JINST **3** (2008) S08005.

[60] TOTEM collaboration, G. Anelli *et al.*, *The TOTEM experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08007.

[61] LHCf collaboration, O. Adriani *et al.*, *The LHCf detector at the CERN Large Hadron Collider*, JINST **3** (2008) S08006.

[62] MoEDAL collaboration, J. Pinfold *et al.*, *Technical Design Report of the MoEDAL Experiment*, CERN, Geneva, 2009. `https://cds.cern.ch/record/1181486`.

[63] CERN, *The Large Hadron Collider*, CERN timelines, `http://timeline.web.cern.ch/timelines/The-Large-Hadron-Collider/overlay`, 2015.

[64] T. Sjostrand, S. Mrenna, and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*, Comput. Phys. Commun. **178** (2008) 852, `arXiv:0710.3820`.

[65] C. Elsasser, *$\bar{b}b$ production angle plots*, LHCb Speakers' Bureau, `https://lhcb.web.cern.ch/lhcb/speakersbureau/html/bb_ProductionAngles.html`, 2012.

[66] LHCb collaboration, P. R. Barbosa-Marinho *et al.*, *LHCb VELO (VErtex LOcator): Technical Design Report*, CERN, Geneva, 2001.

[67] R. Aaij *et al.*, *Performance of the LHCb Vertex Locator*, JINST **9** (2014) 09007, `arXiv:1405.7808`.

[68] LHCb collaboration, S. Amato *et al.*, *LHCb magnet: Technical Design Report*, CERN, Geneva, 2000.

[69] LHCb collaboration, R. Aaij *et al.*, *LHCb detector performance*, Int. J. Mod. Phys. **A30** (2015) 1530022, `arXiv:1412.6352`.

[70] LHCb collaboration, P. R. Barbosa-Marinho *et al.*, *LHCb inner tracker: Technical Design Report*, CERN, Geneva, 2002.

[71] LHCb collaboration, P. R. Barbosa-Marinho *et al.*, *LHCb outer tracker: Technical Design Report*, CERN, Geneva, 2001.

[72] LHCb collaboration, S. Amato *et al.*, *LHCb RICH: Technical Design Report*, CERN, Geneva, 2000.

[73] LHCb collaboration, S. Amato *et al.*, *LHCb calorimeters: Technical Design Report*, CERN, Geneva, 2000.

[74] R. Dzhelyadin, 2015. Personal communication.

[75] LHCb collaboration, P. R. Barbosa-Marinho *et al.*, *LHCb muon system: Technical Design Report*, CERN, Geneva, 2001.

[76] LHCb collaboration, *LHCb muon system: addendum to the Technical Design Report*, CERN, Geneva, 2003.

[77] LHCb collaboration, *LHCb muon system: second addendum to the Technical Design Report*, CERN, Geneva, 2005.

[78] LHCb collaboration, R. Antunes-Nobrega *et al.*, *LHCb trigger system: Technical Design Report*, CERN, Geneva, 2003.

[79] R. Aaij *et al.*, *The LHCb Trigger and its Performance in 2011*, JINST **8** (2013) P04022, arXiv:1211.3055.

[80] R. Fruhwirth, *Application of Kalman filtering to track and vertex fitting*, Nucl. Instrum. Meth. **A262** (1987) 444.

[81] LHCb collaboration, R. Aaij *et al.*, *Measurement of the track reconstruction efficiency at LHCb*, JINST **10** (2015), no. 02 P02007, arXiv:1408.1251.

[82] LHCb collaboration, R. Aaij *et al.*, *Measurement of charged particle multiplicities and densities in pp collisions at $\sqrt{s} = 7$ TeV in the forward region*, Eur. Phys. J. **C74** (2014), no. 5 2888, arXiv:1402.4430.

[83] F. Archilli *et al.*, *Performance of the Muon Identification at LHCb*, JINST **8** (2013) P10020, arXiv:1306.0249.

[84] LHCb collaboration, M. Clemencic *et al.*, *The LHCb simulation application, Gauss: Design, evolution and experience*, J. Phys. Conf. Ser. **331** (2011) 032023.

[85] LHCb collaboration, I. Belyaev *et al.*, *Handling of the generation of primary events in Gauss, the LHCb simulation framework*, J. Phys. Conf. Ser. **331** (2011) 032047.

[86] G. Corti *et al.*, *Software for the LHCb experiment*, IEEE Trans. Nucl. Sci. **53** (2006) 1323.

[87] G. Barrand *et al.*, *GAUDI - A software architecture and framework for building HEP data processing applications*, Comput. Phys. Commun. **140** (2001) 45.

[88] M. Clemencic *et al.*, *Recent developments in the LHCb software framework Gaudi*, J. Phys. Conf. Ser. **219** (2010) 042006.

[89] T. Sjostrand, S. Mrenna, and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*, JHEP **0605** (2006) 026, arXiv:hep-ph/0603175.

[90] D. J. Lange, *The EvtGen particle decay simulation package*, Nucl. Instrum. Meth. **A462** (2001) 152.

[91] P. Golonka and Z. Was, *PHOTOS Monte Carlo: A Precision tool for QED corrections in Z and W decays*, Eur. Phys. J. **C45** (2006) 97, arXiv:hep-ph/0506026.

[92] M. Dobbs and J. B. Hansen, *The HepMC C++ Monte Carlo event record for High Energy Physics*, Comput. Phys. Commun. **134** (2001) 41.

[93] Geant4 collaboration, S. Agostinelli *et al.*, *Geant4: A simulation toolkit*, Nucl. Instrum. Meth. **A506** (2003) 250.

[94] Geant4 collaboration, J. Allison *et al.*, *Geant4 developments and applications*, IEEE Trans. Nucl. Sci. **53** (2006) 270.

[95] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, and W. J. Stirling, *Latest Results in Central Exclusive Production: A Summary*, in *Proceedings of the 7th International Workshop on Diffraction in High Energy Physics*, 2013. arXiv:1301.2552.

[96] C. Barschel, *Precision luminosity measurement at LHCb with beam-gas imaging*, PhD thesis, RWTH Aachen University, 2014, https://cds.cern.ch/record/1693671.

[97] Particle Data Group, K. A. Olive *et al.*, *Review of Particle Physics*, Chin. Phys. **C38** (2014) 090001.

[98] M. Needham, *Momentum scale calibration using resonances*, CERN-LHCB-2008-037, https://cds.cern.ch/record/1115072, 2008.

[99] W. D. Hulsbergen, *Decay chain fitting with a Kalman filter*, Nucl. Instrum. Meth. **A552** (2005) 566, arXiv:physics/0503191.

[100] A. Bukin, *Fitting function for asymmetric peaks*, arXiv:0711.4449.

[101] LHCb RICH Group, M. Adinolfi *et al.*, *Performance of the LHCb RICH detector at the LHC*, Eur. Phys. J. **C73** (2013) 2431, arXiv:1211.6759.

[102] LHCb collaboration, R. Aaij *et al.*, *Exclusive J/ψ and $\psi(2S)$ production in pp collisions at $\sqrt{s} = 7$ TeV*, J. Phys. **G40** (2013) 045001, arXiv:1301.7084.

[103] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, and W. J. Stirling, *The Phenomenology of Central Exclusive Production at Hadron Colliders*, Eur. Phys. J. **C72** (2012) 2110, arXiv:1204.4803.

[104] T. M. Karbach and M. Schlupp, *Constraints on Yield Parameters in Extended Maximum Likelihood Fits*, arXiv:1210.7141.

[105] G. Cowan, K. Cranmer, E. Gross, and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. **C71** (2011) 1554, arXiv:1007.1727.

[106] J. Neyman and E. S. Pearson, *On the problem of the most efficient tests of statistical hypotheses*, Phil. Trans. A **231** (1933) 289.

[107] T. Junk, *Confidence level computation for combining searches with small statistics*, Nucl. Instrum. Meth. **A434** (1999) 435, arXiv:hep-ex/9902006.

[108] A. L. Read, *Presentation of search results: The CL_S technique*, J. Phys. **G28** (2002) 2693.

[109] S. S. Wilks, *The large-sample distribution of the likelihood ratio for testing composite hypotheses*, Ann. Math. Statist. **9** (1938), no. 1 60.

[110] R. D. Cousins and V. L. Highland, *Incorporating systematic uncertainties into an upper limit*, Nucl. Instrum. Meth. **A320** (1992) 331.

[111] M. Albrow *et al.*, *Forward Physics with Rapidity Gaps at the LHC*, JINST **4** (2009) P10001, arXiv:0811.0120.

[112] R. McNulty, *Results and Future Plans on Central Exclusive Production with the LHCb Detector*, Int. J. Mod. Phys. **A29** (2014), no. 28 1446003, arXiv:1409.8113.

[113] A. Bhatti *et al.*, *The Beam Shower Counters for Run II*, CDF/DOC/CDF/CDFR/5247, http://home.fnal.gov/~tesarek/halo/doc/cdf5247_run2_bsc.ps, 2000.

[114] CDF collaboration, T. Aaltonen *et al.*, *Observation of Exclusive Dijet Production at the Fermilab Tevatron $\overline{p}p$ Collider*, Phys. Rev. **D77** (2008) 052004, arXiv:0712.0604.

[115] CDF collaboration, T. Aaltonen *et al.*, *Observation of Exclusive Gamma Gamma Production in $p\bar{p}$ Collisions at $\sqrt{s}$ = 1.96 TeV*, Phys. Rev. Lett. **108** (2012) 081801, arXiv:1112.0858.

[116] CDF collaboration, M. G. Albrow *et al.*, *Central exclusive production of hadrons in CDF*, AIP Conf. Proc. **1523** (2012) 294.

[117] J. W. Lämsä and R. Orava, *Central Diffraction at ALICE*, JINST **6** (2011) P02010, arXiv:1009.3350.

[118] J. W. Lämsä and R. Orava, *Central Diffraction at the LHCb*, JINST **4** (2009) P11019, arXiv:0907.3847.

[119] A. J. Bell *et al.*, *Physics and Beam Monitoring with Forward Shower Counters (FSC) in CMS*, CMS-NOTE-2010-015, https://cds.cern.ch/record/1306543, 2010.

[120] M. Albrow *et al.*, *Small angle detectors for study diffractive processes with CMS*, JINST **9** (2014), no. 10 C10032.

[121] FSC team, CMS collaboration, and HERSCHEL team, LHCb collaboration, M. Albrow, P. Collins, and A. Penzo, *Forward shower counters for diffractive physics at the LHC*, Int. J. Mod. Phys. **A29** (2014), no. 28 1446018.

[122] S. Chemli *et al.*, *LHC Reference Database: Towards a Mechanical, Optical and Electrical Layout Database*, in *Proceedings of EPAC 2004*, pp. 1882–1884, 2004.

[123] J. Stiller *et al.*, *Simulation of the LHC BRAN Luminosity Monitor for High Luminosity Interaction Regions*, in *Proceedings of IPAC 2010*, pp. 498–500, 2010.

[124] Eljen Technology, *EJ-200 plastic scintillator*, `http://www.eljentechnology.com/images/stories/Data_Sheets/Plastic_Scintillators/EJ200%20data%20sheet.pdf`, 2010.

[125] Hamamatsu Photonics, *Photomultiplier tubes R1828-01*, `http://www.hamamatsu.com/us/en/R1828-01.html`, 2012.

[126] T. Hakamata, ed., *Photomultiplier Tubes: Basics and Applications*, Hamamatsu Photonics K. K., 3rd ed., 2006.

[127] M. Lieng, *Evaluation of the Beam Condition Monitor Sensitivity to local Beam Scraping Scenarios at LHCb*, LHCb-INT-2010-005, `https://cds.cern.ch/record/1236707`, 2010.

[128] A. Ferrari, P. R. Sala, A. Fassò, and J. Ranft, *FLUKA: A multi-particle transport code*, CERN, Geneva, 2005. doi: 10.5170/CERN-2005-010.

[129] T. T. Böhlen *et al.*, *The FLUKA Code: Developments and Challenges for High Energy and Medical Applications*, Nucl. Data Sheets **120** (2014), no. 0 211.

[130] R. B. Appleby and A. Mereghetti, *The FLUKA Model of IR8*, LHC Project Note 427, `https://cds.cern.ch/record/1238789`, 2010.