Central Exclusive Production of χ Mesons at LHCb

A thesis
submitted to The University of Manchester
for the degree of
doctor of philosophy (PhD)
in the faculty of Engineering and Physical Sciences

Pål Marius Bjørnstad
B.Sc. University of Oslo, 2009

Particle Physics Group
School of Physics and Astronomy
2013

CONTENTS

Final word count: 44999

LIST OF FIGURES

LIST OF TABLES

ABSTRACT

The LHCb experiment is a forward spectrometer at the Large Hadron Collider, covering a range of pseudorapidity $1.9 < \eta < 4.9$. It has a very precise vertex detector called the VELO, near the interaction point. The monitoring software for the VELO is discussed in this thesis.

In proton-proton interactions, Central Exclusive Production (CEP) is a process where the protons remain intact after the interaction, and an additional simple central system is produced. CEP processes are selected experimentally by requiring that there is no activity in the detector apart from the central system.

The installation of additional detectors to increase the sensitivity at small angles from the beam axis is discussed. A geometrical description of the region of the LHC up to $100\,\mathrm{m}$ on each side of the LHCb interaction point is developed. Simulated forward shower counters are added to the model, and the efficiencies of the detectors are measured.

The main part of this thesis describes an analysis of CEP of χ_c mesons, reconstructed in the decay $\chi_c \to \gamma J/\psi$, $J/\psi \to \mu^+\mu^-$. The cross sections are measured using the dataset of $1\,\mathrm{fb}^{-1}$ collected by LHCb in 2011. The cross sections times branching ratios for χ_{c0}, χ_{c1} and χ_{c2} are measured to be $2.2 \pm 3.0\,\mathrm{pb}$, $4.3 \pm 7.6\,\mathrm{pb}$ and $25.0 \pm 9.2\,\mathrm{pb}$. This is consistent with theoretical predictions. A search for the CEP of χ_b mesons is also performed. A peak is observed in the mass spectrum, but it cannot be uniquely ascribed to χ_b CEP, because of background contributions which are not precisely determined. A limit is set on the sum of the χ_{b1} and χ_{b2} cross sections times branching ratios. Theoretical predictions for the cross section are at the level of $0.1\,\mathrm{fb}$, so no events are expected, and the data are consistent with this.

Institution	The University of Manchester
Candidate	Pål Marius Bjøernstad
Degree Title	Doctor of Philosophy
Thesis Title	Central Exclusive Production of χ Mesons at LHCb
Date	August 2013

DECLARATION

This work represents the combined efforts of the LHCb collaboration. Some of the content has been published elsewhere and/or presented to several audiences, as detailed later in this thesis. I declare that no portion of this work referred to in this thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Signed: Date: October 24, 2013

Pål Marius Bjørnstad
School of Physics and Astronomy
University of Manchester
Oxford Road
Manchester
M13 9PL
United Kingdom

COPYRIGHT

The author of this thesis (including any appendices and/or schedules to this thesis) owns certain copyright or related rights in it (the "Copyright") and he has given The University of Manchester certain rights to use such Copyright, including for administrative purposes.

Copies of this thesis, either in full or in extracts and whether in hard or electronic copy, may be made only in accordance with the Copyright, Designs and Patents Act 1988 (as amended) and regulations issued under it or, where appropriate, in accordance with licensing agreements which the University has from time to time. This page must form part of any such copies made.

The ownership of certain Copyright, patents, designs, trade marks and other intellectual property (the "Intellectual Property") and any reproductions of copyright works in the thesis, for example graphs and tables ("Reproductions"), which may be described in this thesis, may not be owned by the author and may be owned by third parties. Such Intellectual Property and Reproductions cannot and must not be made available for use without the prior written permission of the owner(s) of the relevant Intellectual Property and/or Reproductions.

Further information on the conditions under which disclosure, publication and commercialisation of this thesis, the Copyright and any Intellectual Property and/or Reproductions described in it may take place is available in the University IP Policy[1], in any relevant Thesis restriction declarations deposited in the University Library, The University Library's regulations[2] and in The University's policy on Presentation of Theses.

---

[1] `http://documents.manchester.ac.uk/DocuInfo.aspx?DocID=487`

[2] `http://www.manchester.ac.uk/library/aboutus/regulations`

ACKNOWLEDGEMENTS

I would like to thank my supervisor George Lafferty for the invaluable insights and support on all the research topics I worked on. I also want to thank him for the help during the process of writing up the thesis.

Dermot Moran worked in parallel on similar topics as I did, and deserves credit for much of the work on Central Exclusive Production at LHCb, as well as for being a great collaborator. Chris Parkes and Eduardo Rodrigues have deep knowledge of the VELO subdetector, and I thank them for the good help on VELO-based analysis and software projects.

Many of my colleagues in the Manchester group contributed with advice and help. Adam Webber and I learned a lot about LHCb together, and we had good discussions about many topics. Grant McGregor was at CERN when I was there and usually had some help when things went wrong. I would like to thank the members of the Manchester group for contributing to a friendly atmosphere, while supporting physical inquiry. Fred Loebinger deserves a special mention in this regard.

My family has been wonderful. I would like to thank my parents Anne and Bjørn for helping me in many ways (though not directly with particle physics), and for keeping me reasonably sane. The visits back to Norway were always something to look forward to, and helped me stay productive.

My friend Zhen has been amazing and made life better while I was in this process, and has been an inspiration in multiple ways.

I would like to thank the Norwegian State Educational Loan Fund for financial support. I would also like to thank the University of Manchester for financial support for travel to CERN and conferences.

CHAPTER

ONE

INTRODUCTION

In Central Exclusive Production (CEP) in proton-proton scattering, the protons remain intact and scatter through small angles, producing a central system with relatively low rapidity. Recent theoretical developments in the field of CEP are based on calculations at the level of quarks and gluons, and there are many predictions of signatures which are observable at hadron colliders. The cross sections of many CEP processes have already been measured in experiments, but more data are required to constrain the models. The study of CEP has intrinsic interest because it helps investigations into the nature of the strong force. Additionally, the dynamics of the production process are approximately independent of what system is produced, so the knowledge gained from measurements of CEP of well-known particles can be used to determine the properties of new and unknown particles.

The LHCb experiment at the Large Hadron Collider is designed to study hadrons with b-quarks to compare the properties of matter and antimatter. The instrumentation is also well suited to measure other physics processes, such as CEP of systems with low mass. The measurement of central exclusive production of χ mesons at LHCb is the main topic of this thesis.

The LHCb detector and all its subdetectors are described in Chapter 2. To identify b-hadrons it is of particular importance to measure the trajectories of decay products very close to the point where the protons collide, to determine the b-hadron decay time. For this reason, LHCb has a precise tracking detector near the proton-proton interaction point, called the Vertex Locator. The software used to monitor the Vertex Locator, to ensure that the data quality remains consistent, is described in Chapter 3.

The theory of CEP is introduced in Chapter 4, preceded by a short review of the relevant concepts in the Standard Model. Chapter 4 also includes a review of relevant CEP results from previous and existing experiments.

The option of installing additional detectors at LHCb to improve CEP and diffraction measurements is discussed in Chapter 5. A system of scintillation counters could be installed along the beam pipe in the LHC tunnel near LHCb. In Chapter 5, such a system is simulated and the performance of the simulated system is characterised.

Chapters 6 and 7 describe an analysis of central exclusive χ_c and χ_b production based on the data collected by LHCb in 2011. The CEP process is selected by only accepting events in which there is no additional activity apart from the χ_c or χ_b meson decay products. The analysis is concluded in Chapter 8, in which the determination of the integrated luminosity is described, and the results are presented and compared to theoretical predictions.

The speed of light is set to unity in equations and quantities when discussing velocities, masses, momenta and energies.

CHAPTER

TWO

THE LHCb EXPERIMENT

LHCb [1] is one of the four major experiments at the Large Hadron Collider (LHC) at CERN. It is a one-arm forward spectrometer with an angular acceptance of $10\,\mathrm{mrad}$–$300\,\mathrm{mrad}$. Its main design goal is to measure CP-violation and to search for rare decays, in processes involving b-quarks, to look for signs of new physics [2]. The physics programme of LHCb has also been expanded beyond b-quark physics, as the unique acceptance and particle identification ability can be used for measurements that are complementary to other experiments.

2.1 The Large Hadron Collider

2.1.1 Overview

The LHC [3–6] is a synchrotron designed to collide protons at a maximum centre of mass energy of $\sqrt{s} = 14\,\mathrm{TeV}$. There are two proton beams, travelling in opposite directions. It is located at CERN, in the same tunnel that was used for the LEP (Large Electron-Positron) collider. The LHC ring has a circumference of $27\,\mathrm{km}$, and is located approximately $100\,\mathrm{m}$ underground. The LHC was designed to explore physics at energy scales up to a few TeV. Two of the main goals are to find and study the Higgs boson, and to search for supersymmetry. It was always envisioned to have a broad physics programme, also including precision measurements of quantities in the Standard Model.

One of the main reasons that protons are chosen as the particle to accelerate is that the energy loss due to synchrotron radiation is smaller than that of electrons by a factor $(m_e/m_p)^4 \sim 10^{-13}$. Using protons makes it possible to have a much higher energy than LEP. Protons can easily be produced in greater amounts than antiprotons, and this is the reason why antiprotons are not used for one of the beams. The LHC also collides heavy ions instead of protons for a few weeks of the year, to explore nuclear physics and quantum chromodynamics at high energy densities.

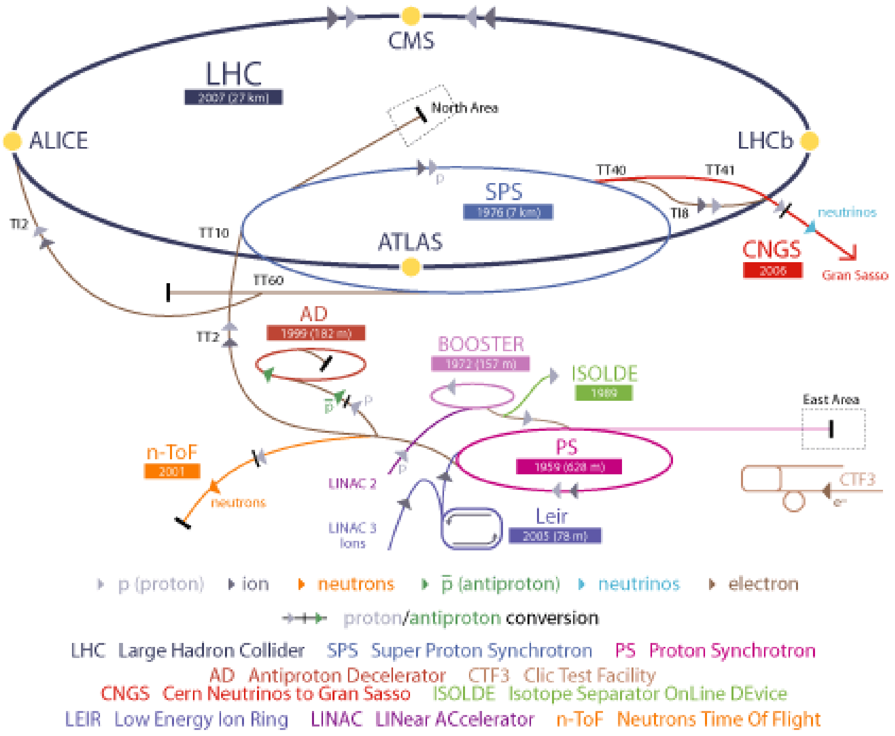

Figure 2.1: The CERN accelerator complex. The rings are not drawn to scale.

Dipole magnets are used to keep the protons in a circular orbit, and quadrupole magnets focus the beams to have a small transverse size at the interaction points. The magnets used in LHC are niobium-titanium (Nb-Ti) superconducting magnets cooled to a temperature of $1.9\,\mathrm{K}$ by superfluid helium. A field of $8.33\,\mathrm{T}$ is required in the dipole (bending) magnets for the maximum beam energy of $7\,\mathrm{TeV}$. The vacuum chambers containing the two beams are only separated by $194\,\mathrm{mm}$, which means that the same magnets have to be used to deflect both beams. The magnets are designed such that the correct field is applied to each beam.

Figure 2.1 shows an overview of the CERN accelerator complex, including the accelerator chain for LHC. The LHC experiments are marked with yellow points. The protons are produced in a duoplasmatron source, and then accelerated by Linac2, the PS booster, the PS, SPS and finally injected into the LHC ring with an energy of $450\,\mathrm{GeV}$ [5]. The proton beams are accelerated to the desired energy by superconducting RF cavities operating at $400\,\mathrm{MHz}$. After acceleration, a period of stable beams commences, in which the beams collide and the detectors can collect data. The fill lasts for a maximum of 10–20 hours, after which it is dumped and the LHC starts with a new fill of protons.

2.1.2 Machine parameters

The protons are not distributed continuously along the ring, but are instead separated into a maximum of 2808 bunches. The maximum intensity is 1.15×10^{11} protons per bunch [3]. The bunches cross to produce collisions at four main experiments, ATLAS, ALICE, CMS, and LHCb. There are also two forward physics experiments, LHCf and TOTEM, located at distances greater than $100\,\mathrm{m}$ from the ATLAS and CMS interaction points. The MoEDAL experiment is installed in the same area as LHCb, and searches for magnetic monopoles.

During nominal operation, bunch crossings will happen at a rate of $40\,\mathrm{MHz}$ at all interaction points. The instantaneous luminosity of a synchrotron can be written as

$$\mathcal{L} = \frac{N_b^2 n_b f_{rev} \gamma_r}{4\pi \epsilon_n \beta^*} F \ , \tag{2.1.1}$$

where N_b is the number of particles per bunch, n_b is the number of bunches in the machine, f_{rev} is the revolution frequency of each bunch, and γ_r is the relativistic gamma factor of the protons. The normalised transverse emittance ϵ_n describes the distribution of beam momentum in the transverse direction. β^* is a parameter of the focussing of the beams, and determines the transverse size of the bunches at the interaction point together with ϵ_n. The beams have to cross at an angle because their trajectories should only intersect at the interaction point. F is the reduction of the luminosity due to the crossing angle. The peak design luminosity is $1.0 \times 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ at the ATLAS and CMS interaction points, and lower at the LHCb and ALICE interaction points.

2.1.3 Operation and performance

The LHC was first started in 2008, but had to be shut down due to a an explosion of one of the magnets caused by a faulty electrical connection. Commissioning with beams was resumed in November 2009 [7], and collisions were achieved on the 23rd November. LHCb got an early start on the commissioning, both before the 2008 incident and before beams were injected in 2009, in tests of the synchronisation of the SPS and LHC, during which beams were dumped on an absorber upstream of LHCb [8]. On the 30th March 2010, the first collisions at $7\,\mathrm{TeV}$ centre of mass energy occurred ($3.5\,\mathrm{TeV}$ per beam), marking the beginning of the LHC physics programme. Throughout 2010, the luminosity was gradually increased by adding more bunches and by increasing the number of protons per bunch. A minimum bunch spacing of $50\,\mathrm{ns}$ was used. The 2010 proton run ended in October, with a subsequent period of proton-lead and lead-lead collisions. The luminosity delivered by the LHC to the experiments is shown in fig. 2.2. Running commenced again in March 2011, with $3.5\,\mathrm{TeV}$ beams, and the peak luminosity was increased throughout the year from about $10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ at the end of March to $40 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ for ATLAS

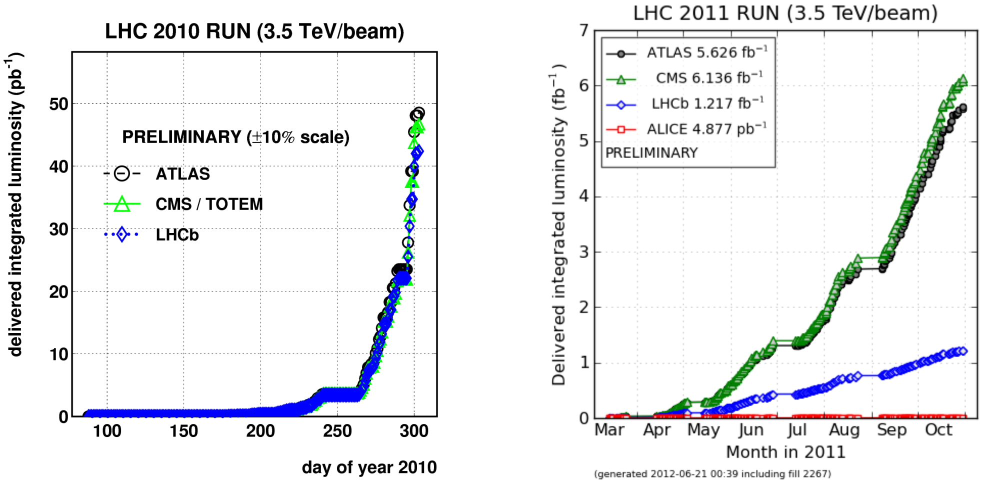

(a) 2010 (b) 2011

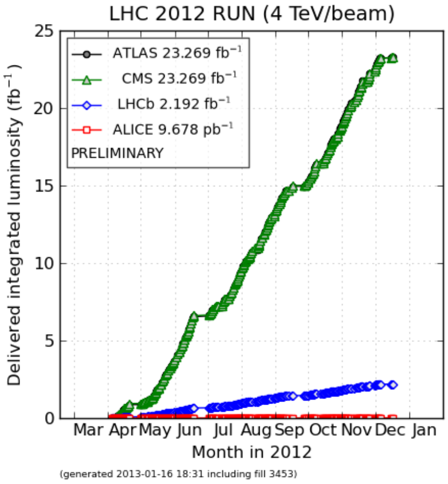

(c) 2012

Figure 2.2: LHC delivered integrated luminosity for the major experiments in the three main periods of running before the first long shutdown.

and CMS at the end of the year. LHCb has a design luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, chosen to take predominantly events with single proton-proton interactions. As LHCb reached its design luminosity in 2011, the beam optics were adjusted to reduce the luminosity in LHCb. The same was done for ALICE, which requires an even lower luminosity. The fast increase of instantaneous luminosity was achieved by quickly scaling the bunch intensity to its nominal value of 1.15×10^{11} protons per bunch. The experiments then got data with a high number of interactions per event (“pile-up”). The proton beams were then shut down at the end of October, for heavy ion beams and then a technical stop. ATLAS and CMS collected $6\,\mathrm{fb}^{-1}$ of integrated luminosity in 2011, and LHCb collected $1\,\mathrm{fb}^{-1}$. The machine was restarted in March 2012, with a beam energy of $4\,\mathrm{TeV}$. The peak luminosity in 2012 was $70 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, and the integrated luminosity was $23\,\mathrm{fb}^{-1}$ for

ATLAS and CMS, and $2\,\mathrm{fb}^{-1}$ for LHCb. After the end of the 2012 running period, there was a period of heavy-ion physics in the beginning of 2013, followed by the first long shutdown.

2.2 Overview of the LHCb experiment

The geometrical layout of LHCb was motivated by the fact that pairs of b-quarks and antiquarks are produced predominantly in a cone in the forward and backward direction. Hadrons with b-quarks have lifetimes of the order of $10^{-12}\,\mathrm{s}$, which means that they travel distances of $\sim 0.5\,\mathrm{mm}$ in the detector when moving close to the speed of light. It is thus important to have a very fine resolution for tracks close to the interaction point, to be able to measure this finite lifetime. The Impact Parameter (IP) is the distance of closest approach between a track and a vertex, i.e. between the measured particle trajectory and the measured position of a particle decay. The IP with respect to the primary vertex, which is the vertex of the interaction of the beam protons, is frequently used in cuts to select tracks from B-decays, functioning as a proxy for the particle lifetime. Flavour physics measurements also require good particle identification, because it is necessary to tell the difference between pions and kaons. A good momentum resolution is also required to measure the mass of particles.

The LHCb coordinate system is defined with the interaction point at coordinates (0,0,0), with the x-axis in the horizontal direction, the y-axis in the vertical direction and the z axis along the beam axis, in the direction of the spectrometer arm. The layout of the subdetectors and the directions y and z of the coordinate system are shown in fig. 2.3. The subdetectors are arranged in an approximately projective geometry, all covering at least the pseudorapidity range $1.9 < \eta < 4.9$ (except the RICH 2). The pseudorapidity is a function of the polar angle θ,

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

The pseudorapidity of a particle can equivalently be written as:

$$\eta = \frac{1}{2}\ln\left(\frac{|p| + p_z}{|p| - p_z}\right) \,, \tag{2.2.1}$$

where $|p|$ is the magnitude of the momentum of the particle and p_z is the longitudinal (parallel to the beam axis) component. The rapidity is defined by replacing $|p|$ with the energy E, and the pseudorapidity is close to the rapidity for relativistic particles.

The tracking system comprises the Vertex Locator (VELO), the Tracker Turicensis (TT), and the Inner Tracker (IT), which are silicon strip detectors, and the Outer Tracker (OT), which is a straw-tube detector. A dipole magnet is used with the tracking detectors

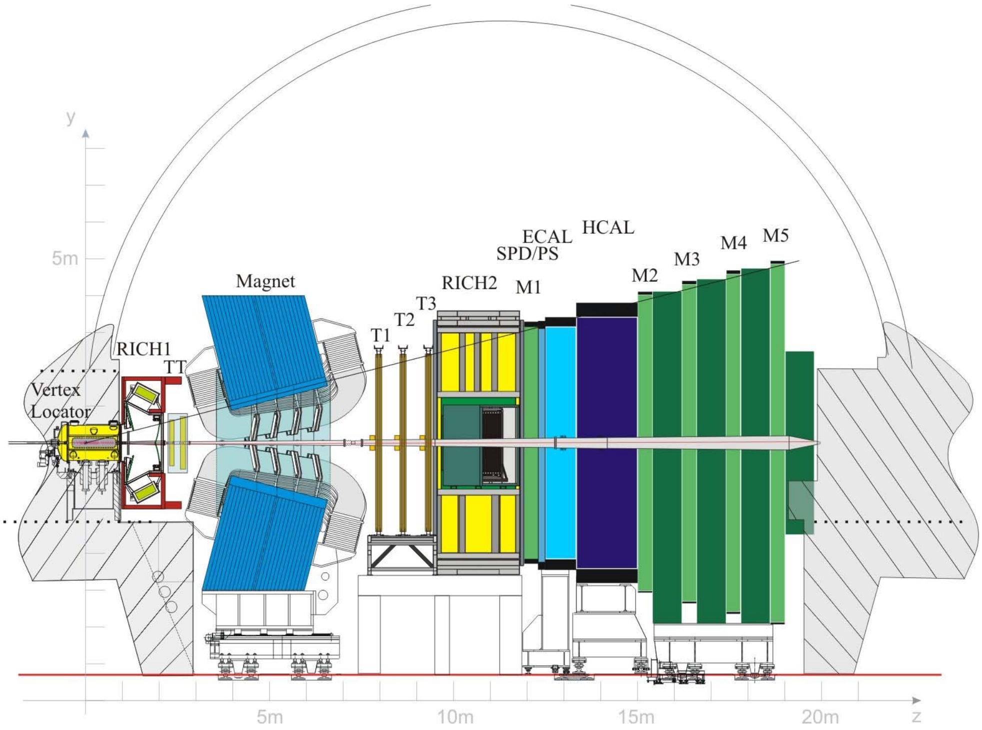

Figure 2.3: LHCb detector view. The interaction point is located inside the Vertex Locator. The beam pipe is the horizontal structure extending throughout the detector at the centre. Most subdetectors have a rectangular shape in the plane perpendicular to the beam axis (not visible in the drawing).

to give charged particles a curved trajectory, so that the momentum can be measured. The magnetic field vector points in the vertical (y) direction, and the polarity (sign) can be changed, in order to cancel certain charge asymmetries. The magnitude of the field is approximately uniform in the x and y coordinates, and has an approximately Gaussian dependence on the z (beam axis) coordinate. The VELO is located near the interaction point, and the TT is further from the interaction point, but before the magnet. Three tracking stations, each with an IT and an OT region, are located after the magnet. The VELO is used to reconstruct tracks and vertices. The VELO must reconstruct primary vertices (PVs) and decay vertices. The TT is used to improve the momentum measurement in track reconstruction. The tracking stations T1, T2 and T3 provide the main momentum measurement for charged particles. All three tracking stations include an IT region close to the beam axis, and an OT region which surrounds the IT and measures particles at large angles from the beam.

LHCb has two Ring Imaging Cherenkov detectors, which are used for particle identification. The RICH 1 and RICH 2 are placed before and after the magnet respectively. An electromagnetic calorimeter (ECAL) measures the energy of photons and electrons in the acceptance. A hadronic calorimeter (HCAL) follows the ECAL, measuring the energy of hadrons. The calorimetry system includes two scintillating pad detectors, the SPD and the PS, which provide particle identification information. The calorimeters are used to trigger on particles with high transverse energy and to reconstruct the energy of neutral particles in analyses. The muon identification system includes five stations, one located before the calorimeters and four after. The muon detectors operate on the principle that muons can penetrate through a large amount of material without destructive interactions, so the particles that make it through all the material can be identified as muons.

LHCb includes a two-level trigger system which reduces the 10–$20\,\mathrm{MHz}$ rate of events with interactions to a $\sim\ 4.5\,\mathrm{kHz}$ rate of physics and calibration events which can be saved for offline analysis. The first trigger level is Level 0 (L0), is implemented on Field Programmable Gate Arrays (FPGAs), and operates at a $40\,\mathrm{MHz}$ input rate, synchronous with the bunch crossings. The algorithms of the L0 trigger are implemented on special-purpose hardware and cannot easily be changed, but there are configurable thresholds and other parameters. The L0 trigger mainly uses information from the calorimeters and the muon detectors.

The next level is the High Level Trigger (HLT), which is subdivided into the HLT1 and HLT2. The HLT is a software program running on commodity computer hardware, and it is thus very flexible and scalable. The HLT has access to the full event data.

2.3 Tracking and vertex reconstruction

Tracking detectors are used to non-destructively measure the trajectories of individual charged particles through the detector. Each tracking station is a plane perpendicular to the beam axis, and measures a single point on a track. The tracking detector allows to measure the momentum of charged particles from the curvature of the tracks in the magnetic field. There are no tracking stations inside the region with the strongest field, so the curvature is measured by matching up a straight VELO track from before the magnet with a track from the tracking stations ("T-stations"). Tracks are slightly curved in the T-stations due to the fringe field of the magnet. A total integrated field of approximately $4\,\mathrm{Tm}$, in the y-direction (vertical) is provided by the non-superconducting dipole magnet. The particles are deflected in the x-direction, and for this reason the tracking system is designed to have greatest resolution in the x-coordinate.

2.3.1 Vertex Locator

The Vertex Locator (VELO) [1, 9] provides tracking information for particles close to the interaction point. The tracks are used to reconstruct primary vertices and decay vertices, by searching for multiple tracks originating from a point. The key VELO measurements are the IP of particles with respect to a vertex, primary vertex reconstruction and decay length. VELO information is also used as an important part of the general

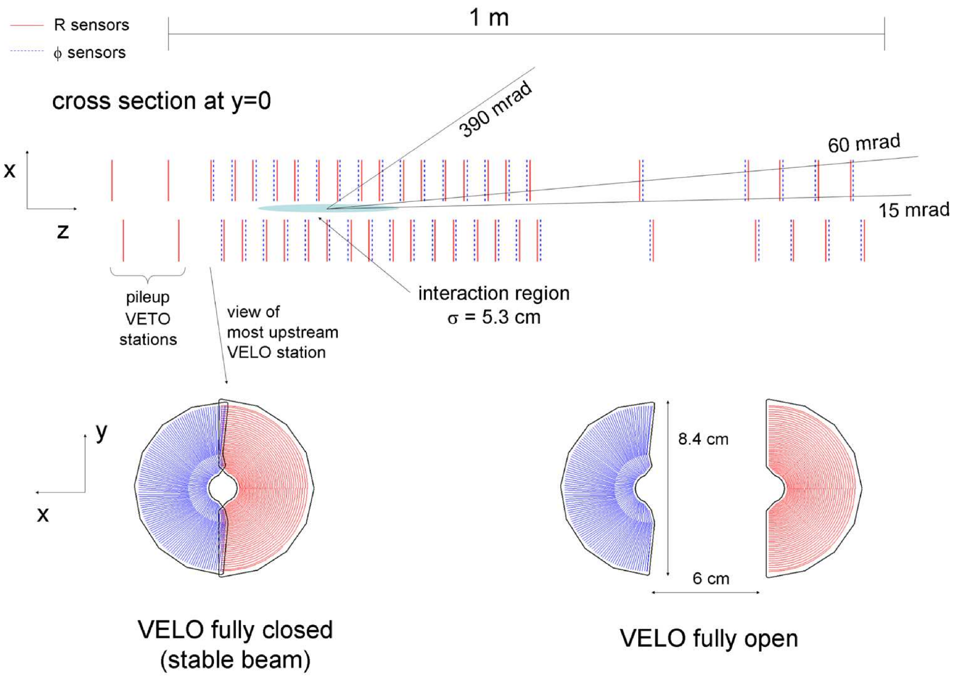

Figure 2.4: Sensors of the VELO. The diagram on the top shows the position of the VELO sensors in the z-direction. The bottom diagrams show the design of a sensor in the xy-plane.

tracking system, and it can significantly impact the mass resolution of particles with a measurable lifetime, like B and D mesons.

The impact parameter resolution depends mainly on three characteristics of the design: the position resolution of single particle hits, the distance from a reference point such as a vertex to the first measurement of a track (extrapolation distance), and the amount of material that a particle has to traverse. The VELO is designed to have an optimal compromise between these parameters, also respecting the external constraints such as the aperture required for the beam. The primary vertex resolution additionally depends on the number of tracks from the primary vertex that are reconstructed, making it beneficial to have a large acceptance.

The VELO consists of 42 “R-sensors” for measuring a radial coordinate r, interleaved with 42 “Phi-sensors” for measuring an azimuthal coordinate ϕ. The layout of the VELO and individual VELO sensors is shown in fig. 2.4. A series of 42 sensors (21 R + 21 Phi) is placed on each side of the beam axis. The VELO sensors are shaped like half discs, and pairs of opposite sensors combine to give a full 360° acceptance in the transverse plane. The sensors are mounted on hybrids, which provide the readout electronics and a thermal interface. The hybrids are cooled by a two-phase $\mathrm{CO_2}$ cooling system. Pairs of R-sensors and Phi-sensors and their hybrids are mounted on modules, which are fastened to a movable support structure. The sensors span approximately $1\,\mathrm{m}$ in the direction of the beam, from $z = -30\,\mathrm{cm}$ to $z = 75\,\mathrm{cm}$. There are four additional R-sensors placed beyond the normal VELO sensors in the opposite (z) direction to the LHCb acceptance, and these are called the pile-up system. The pile-up system sensors were originally meant to be a veto on events with multiple interactions (pile-up), and they are not used for tracking.

The R-sensors have strips arranged in concentric circles around the beam axis, but divided into segments of $45°$. The strip pitch is $40\,\mu\mathrm{m}$ in the inner region of the R-sensors, and increases linearly with r-coordinate to $102\,\mu\mathrm{m}$ at the outer radius, which is $42\,\mathrm{mm}$. The pitch and segmentation are chosen to have an acceptable occupancy, which is also evenly distributed between the strips. The Phi-sensors have strips which are approximately radial, but they are aligned with a stereo angle. The strips are separated into an inner and an outer annulus, where the inner strips have a smaller pitch, but the pitch also naturally increases with r within each annulus.

The VELO sensors are positioned as close to the beam as possible, with the sensitive region only $8.2\,\mathrm{mm}$ away from the beam. When protons are injected into the LHC, their orbits are not completely stable, and it is not safe for the VELO sensors to be placed at such a small distance. It is necessary to move the two halves of the VELO away from the beam axis. During injection the halves are in a position retracted by $3\,\mathrm{cm}$ in the x-direction. When the LHC declares “stable beams” (i.e. that the beams are at the nominal energy and in a stable orbit), the sensors can be moved to their operating positions. The motion is actuated by precise stepper motors, which can position the sensors exactly around the interaction point, also in the y-direction. The repositioning is referred to as “closing” of the VELO, and the VELO is “open” when the sensors are in the retracted position and “closed” when they are in the operating position. The closing can be performed in approximately $210\,\mathrm{s}$ using an automated procedure with multiple intermediate steps and safety checks. The VELO is re-opened when the beams are dumped.

The signal is read out by a front-end chip called the Beetle [10], which amplifies the signal and stores it in an analogue pipeline. After a positive L0 trigger decision, the Beetle multiplexes the signals from 128 strips onto four cables at a maximum trigger rate of $1\,\mathrm{MHz}$. The cables carry the signals to the *counting house*, which is an area $60\,\mathrm{m}$ away from the detector, shielded from radiation. The analogue signals are then digitised and processed by TELL1 boards [11]. The same boards are used by most of the LHCb subdetectors. A suite of signal processing algorithms is run on the TELL1 boards to normalise the signal, with the final result of producing *clusters*. Clusters represent the signal of a particle hitting the detector. The hit position can be reconstructed from the ADC values (normalised output of the Analogue to Digital Converter) with better precision than the strip pitch, using a weighted average of adjacent signals. For most events, only the clusters are stored, and this is referred to as a Zero-Suppressed (ZS) readout. There is a Non Zero-Suppressed (NZS) readout mode which stores the ADC value for every single strip on a sensor, and this runs periodically on all VELO sensors for monitoring purposes.

The VELO is monitored by multiple software systems. Operational parameters such as sensor currents and voltages, pressure and sensor temperatures are monitored by the Experiment Control System. Pressure sensors are located in the vacuum chambers, and many temperature sensors are used, including one on each hybrid. The *online* system analyses events as they arrive, and computes quantities such as the sensor occupancy and minor sensor errors. The online data quality histograms are monitored by an operator in the LHCb control room at all hours. The final level of monitoring is the *offline* monitoring, which is performed by two separate systems. There is a central offline monitoring stack for LHCb which does a comprehensive analysis of the LHCb data, including VELO-related quantities at the level of clusters, tracks, vertices and impact parameters. The detector alignment is also monitored. The other offline monitoring system is dedicated to the VELO, and includes more detailed information. The software is discussed in more detail in Chapter 3.

VELO performance

The VELO has had a generally excellent performance throughout the period of operation [8]. The vacuum, cooling, low and high voltage, and motion subsystems have operated stably. The Data Acquisition (DAQ) chain, and in particular the Beetles and the TELL1 boards, require monitoring and tuning as the signals change with radiation damage. The system has been able to run without major problems throughout the data taking period, and the data quality is consistent. The monitoring software has been continually improved over the period of running, to incorporate new quantities and analyses based on operational experience.

The charge deposited in a strip resulting from the passage of a particle follows a Landau distribution. The signals measured by the ADCs also have a Gaussian resolution component. The distribution of ADC values for clusters on tracks is shown in fig. 2.5a. A Landau distribution convolved with a Gaussian is fitted to the data, giving a good fit up to about 60 ADC counts. There is an excess of clusters with a large number of ADC counts, which is due to $\gamma \rightarrow e^+e^-$ conversions, where the electron and the positron continue on almost overlapping trajectories and thus produce large energy depositions. The signal to noise ratio, shown for a specific sensor in fig. 2.5b, varies across a sensor, but is well above the design value of 15 for all sensors.

The track impact parameter resolution is shown in fig. 2.6, plotted as a function of $1/p_T$. The dependence on $1/p_T$ is approximately linear, with the best resolution at high p_T, low $1/p_T$. The tracks selected for physics analysis, such as the daughters of B-decays, have large p_T, and thus benefit from the good resolution. The slope of the line depends on the amount of material traversed by the particle. The difference between the Monte Carlo (MC) and the data is believed to be due to an incomplete model of the

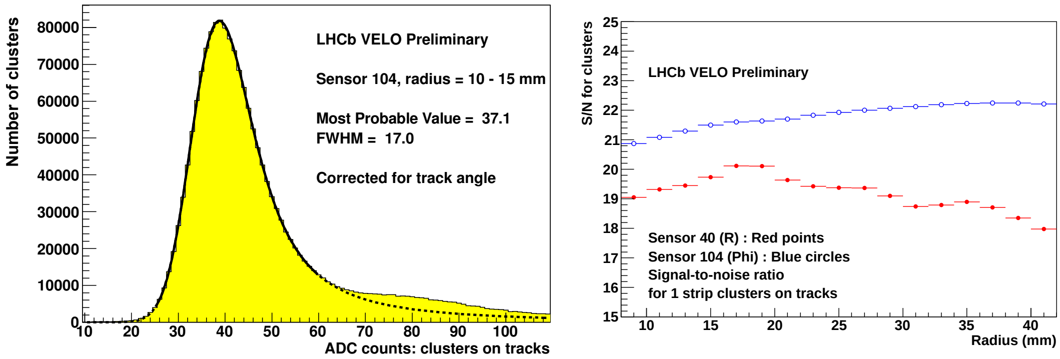

(a) Cluster signal distribution

(b) Signal to noise

Figure 2.5: Cluster ADC values and signal to noise. The solid yellow histogram in (a) shows the distribution of strip signal (ADC) for strips which are on the trajectory of a particle. A fit of a Landau distribution convolved with a Gaussian is shown by the dotted curve. Plot (b) shows the signal to noise ratio for an R-sensor, number 40, and for a Phi-sensor, number 104, as a function of radial coordinate. The signal is extracted from data using fits like the one in (a), and the noise is measured using NZS data.

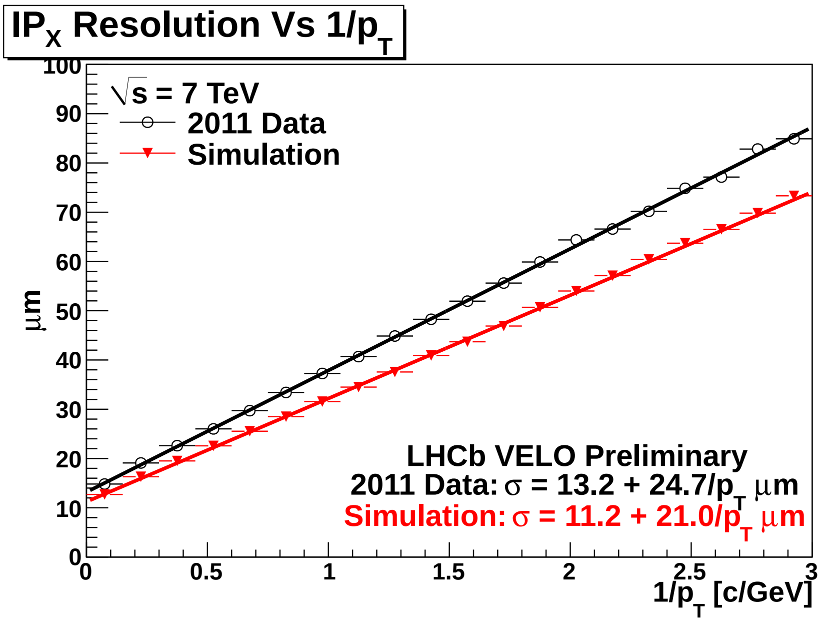

Figure 2.6: Resolution for the x-component of the impact parameter.

shape of an aluminium foil separating the sensors from the main LHC vacuum. Some improvement has been made in the MC description after the plot was made, but there is still a discrepancy, and work is under way to produce a complete description of the material using Computer Aided Design (CAD) tools. Generally, the IP resolution is good and physics analyses are not severely affected by the discrepancy compared to MC.

Because the sensors are close to the interaction point they are irradiated at a very high rate, with a fluence of $45 \times 10^{12}\, n_{\mathrm{eq}}/\mathrm{cm}^2$ (equivalents of $1\,\mathrm{MeV}$ neutrons per unit area) over the first two years of operation [12] (up to the end of 2011, $1.22\,\mathrm{fb}^{-1}$). The fluence falls sharply with radius, and also depends on the sensor z-position. The concern of radiation hardness played a major role in the design of the VELO sensors. The sensors are specified to last for three years of running at nominal conditions, which were assumed to be a maximum fluence per year of $(500 - 1300) \times 10^{12}\, n_{\mathrm{eq}}/\mathrm{cm}^2$ [9]. A complete set of identical copies of the VELO modules has been made, for replacement of the sensors when they receive too much radiation damage.

The silicon bulk material is produced with n-type doping, but the radiation gradually changes the effective doping concentration. In the process referred to as type inversion, radiation damage adds effective p-type doping by displacing atoms in the silicon lattice, and removes n-type doping. Two of the VELO sensors are produced with p-type bulk instead, for trial of this technology as a candidate for upgrades of LHC detectors. To get a high cluster finding efficiency one needs to fully deplete the strips of charge carriers, applying a bias voltage. For a sensor with n-type bulk material the depletion voltage will first decrease, then start to increase after type inversion. The Effective Depletion Voltage (EDV) is monitored in multiple ways, for example by plotting the charge collection efficiency as a function of bias voltage, and declaring as the EDV the voltage at which the cluster finding efficiency has reached $80\,\%$ of its peak value. A comprehensive plot of the EDV as a function of fluence is shown in fig. 2.7, based on five charge collection efficiency scans in 2010 and 2011 for integrated luminosities from zero to $1.22\,\mathrm{fb}^{-1}$ in (a) and for the 2012 data in (b). The fluence is computed from simulation and the knowledge of the luminosity, but it is also checked using a known relationship with the sensor currents. The plot shows the evolution of the EDV from left to right with increasing fluence, and a minimum due to type inversion is clearly visible. The change of EDV with integrated luminosity is in agreement with expectations except in the region where type inversion takes place. The high voltage system is specified to provide bias voltages up to $500\,\mathrm{V}$, but it may be unstable at that point, and it may be desirable to replace the sensors when the depletion voltage reaches $\sim 300\,\mathrm{V}$.

An unexpected loss of efficiency after radiation damage was observed on strips in the outer region of R-sensors [12]. Investigation shows that this is due to a diversion of the charge from the strip to the *second metal layer*, which is a layer of routing lines for reading out the inner strips. Another effect caused by the same process is the production of

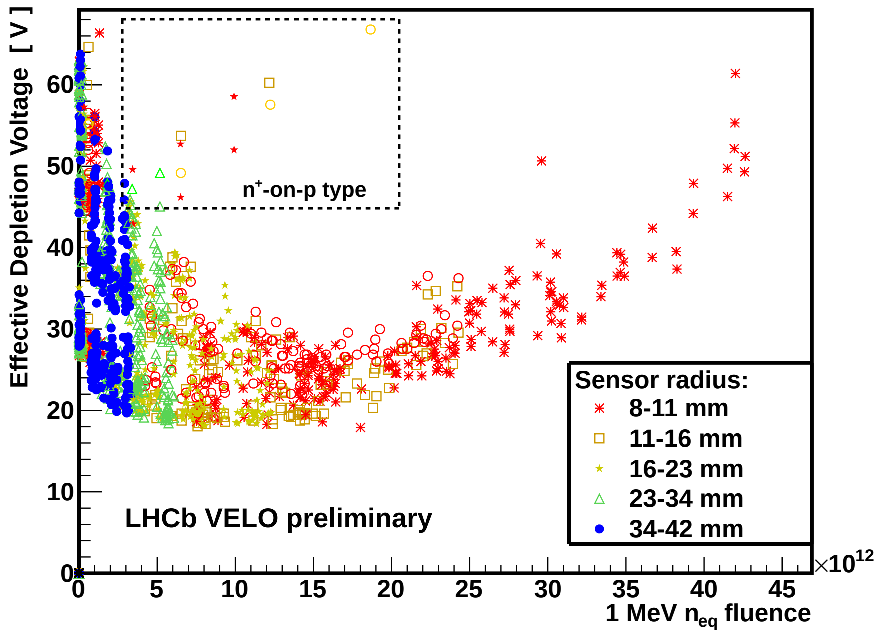

(a) 2010 and 2011

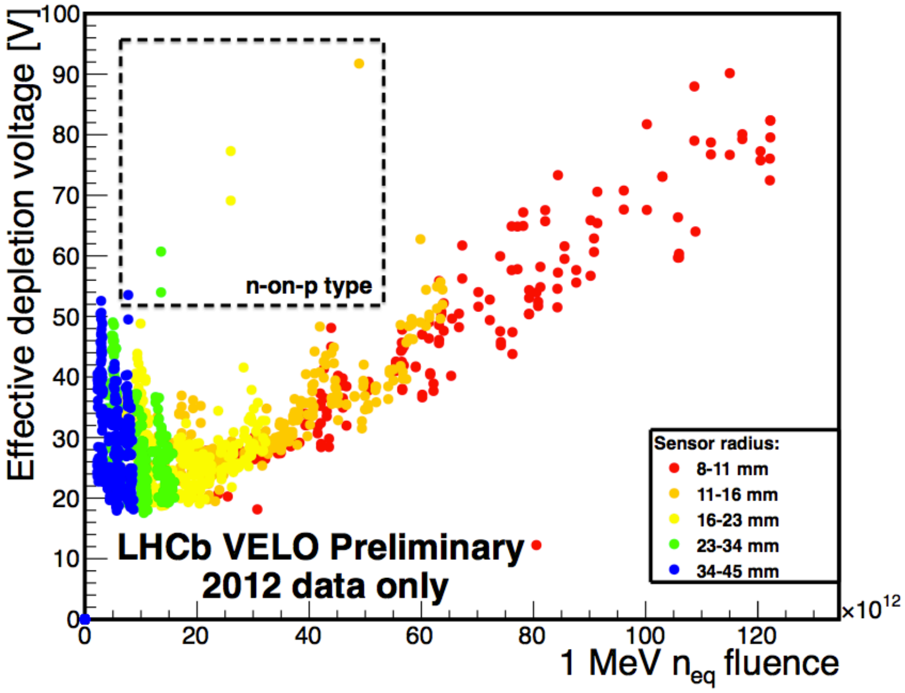

(b) 2012

Figure 2.7: Effective depletion voltage. Sensors are separated into multiple bins of radial coordinate. Each data point corresponds to a bin and a sensor and an EDV scan (a specific point in time). The fluence (from simulation) is shown on the horizontal axis and the effective depletion voltage is shown on the vertical axis.

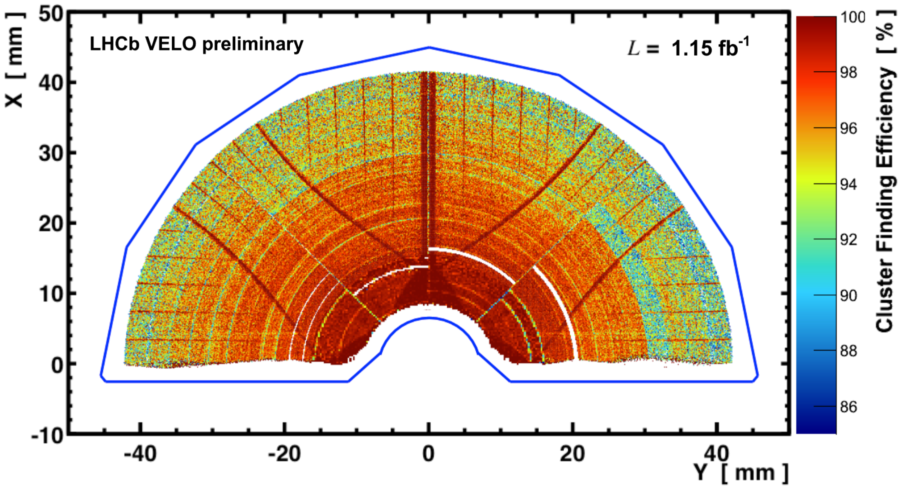

Figure 2.8: Two-dimensional map of sensor 40 after the 2011 data taking, showing a high-resolution scan of the cluster finding efficiency as a function of position on the sensor. The areas without routing lines have a higher efficiency.

fake clusters with low ADC values on the inner strips connected to the routing lines. Neither of the effects are seen in Phi-sensors, because the routing lines are located directly above the outer strips, and run parallel to them, so the routing lines are shielded by the strips. Studies show that the second metal layer effect does not impact the VELO tracking efficiency, within the available precision of $\pm 0.3\,\%$. The effect increases with luminosity (or irradiation), but it appears to level off for some of the most irradiated sensors, indicating that it will not become significantly worse [12]. The origin of the effect is confirmed by plotting the cluster finding efficiency as a function of position on the sensor. Figure 2.8 shows the cluster finding efficiency as it varies across a sensor, and the regions with high efficiency are known to be areas without routing lines.

2.3.2 Silicon Tracker

The TT and IT have many components in common, and are referred to collectively as the Silicon Tracker (ST). The TT is a stand-alone detector, while the IT covers the inner region of the tracking stations.

Tracker Turicensis

The TT [1] is $150\,\mathrm{cm}$ wide and $130\,\mathrm{cm}$ high, and covers the full LHCb acceptance. It is positioned at a distance of $2.5\,\mathrm{m}$ from the interaction point. The hybrids and other readout electronics are placed outside the acceptance. The layout of the TT is shown in fig. 2.9. There are four TT layers, in two pairs measuring coordinates (x, u) and (v, x), where u and v are coordinates rotated by stereo angles of $-5°$ and $+5°$ with respect to the x-coordinate. The two pairs of layers are separated by $27\,\mathrm{cm}$. The TT uses $500\,\mu\mathrm{m}$

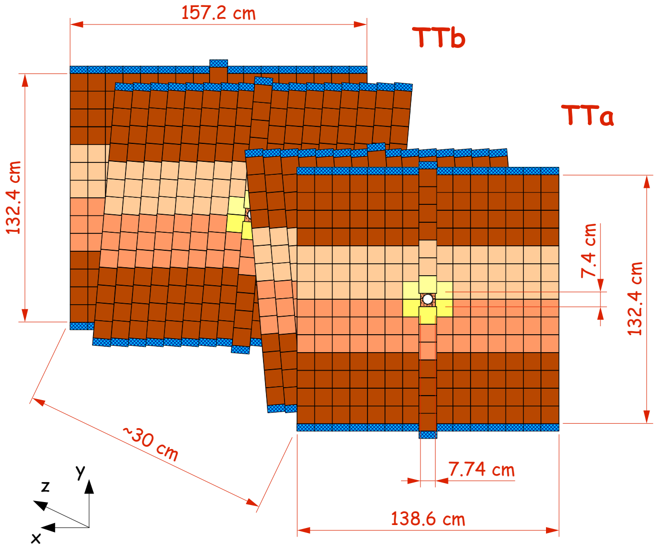

Figure 2.9: Layout and dimensions of the TT, with the layers shown in perspective. The layer in front is closest to the magnet.

thick sensors with 512 strips, with width and length $9.64\,\mathrm{cm} \times 9.44\,\mathrm{cm}$. The sensors are mounted on half-modules of seven sensors in the vertical direction, covering half the LHCb acceptance. Each layer has 17 columns of pairs of adjoining half-modules. In the central column there are two half-modules with a gap between them through which the beam pipe passes. The sensors on a module are not read out individually, as they are bonded together in two or three groups, shown in separate colours in fig. 2.9.

Inner Tracker

The IT covers the area close to the beam pipe for all tracking stations T1, T2 and T3. The outer region of the tracking stations is covered by the OT (Section 2.3.3), but it is necessary to have a finer grained instrumentation near the beam pipe, where the occupancy is much greater. The layout of a layer of an IT station is shown in fig. 2.10. The IT sensors have 384 strips and are $7.6\,\mathrm{cm}$ wide and $11\,\mathrm{cm}$ long. The modules have either one or two sensors, with a thickness of $320\,\mu\mathrm{m}$ or $410\,\mu\mathrm{m}$ respectively. The sensors are bonded together in modules with two sensors. The three IT stations each have four detector boxes, positioned on the left, on the right, above and below the beam pipe. There are four layers in each box, measuring x, u, v and x, as for the TT. The layer shown in fig. 2.10 is an x detection layer. The IT components are designed to have a small material budget, as the hybrids, electronics, cooling system and support structure all have to be in the detector acceptance.

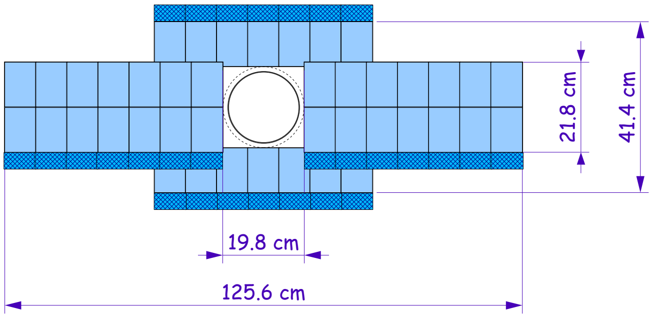

Figure 2.10: A layer of the IT, with dimensions. The sensors are shown in light blue, and the hybrids in dark blue.

Readout

The ST sensors (or groups of sensors) are read out with Beetle chips. Like for the VELO, the signals from the Beetle are forwarded on positive L0 triggers at a maximum rate of $1\,\mathrm{MHz}$. The signals are transferred to service boxes near the detector, but outside the acceptance, where they are digitised and sent over optical fibres to the counting house. Finally, the fibres are connected to TELL1 boards, which produce clusters.

Performance

There have been some minor operational problems with the ST, but the performance has generally been excellent [13,14]. The TT and IT were running with 98 and $99\,\%$ of the channels operational in 2011 [15]. The signal to noise is in the range 12–15 in the TT, and in the IT it is 16.5 for the pairs of sensors bonded together and 17.5 for the single sensors. The TT and IT are aligned to a precision of $16.7\,\mu\mathrm{m}$ and $15.6\,\mu\mathrm{m}$. The radiation damage of the ST sensors is being monitored, by measuring the leakage current and exploiting a known relationship with the fluence.

2.3.3 Outer Tracker

The OT [1] is a straw-tube detector for tracking of particles downstream of the magnet, covering the remainder of the LHCb acceptance outside the IT. The straw tubes have an inner diameter of $4.9\,\mathrm{mm}$, with a $25.4\,\mu\mathrm{m}$ anode wire. The gas used is a mixture of $\mathrm{Ar/CO_2/O_2}$ with components $70\,\%/28.5\,\%/1.5\,\%$ [16]. The tubes are oriented vertically and combined into modules. Each module contains two staggered layers of 64 tubes each, referred to as monolayers. Figure 2.11 shows the layout of a module, as seen from above. There are then four layers of modules (each containing two monolayers) in each station. The OT has the same layer structure as the IT and TT, where each station consists of four layers: x, u, v and x. A precise position measurement requires measuring the time it takes

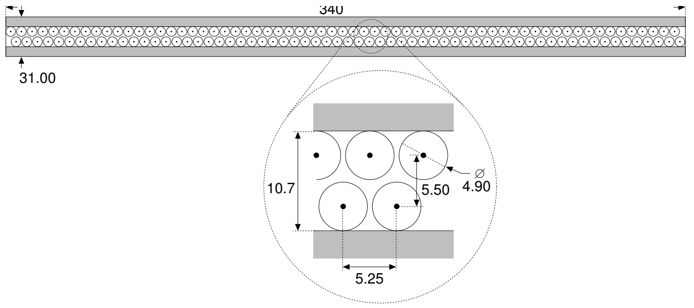

Figure 2.11: A module of the OT, comprising two monolayers. Dimensions are in millimetres.

for the charge to drift from the hit position to the anode wire in the centre of the tube. The OT front end electronics amplify the pulse and then pass it to a time to digital converter. The front-end electronics are specially designed for the OT. The time signals for all of the OT are sent to TELL1 boards over optical links. The OT straw tubes have a position resolution of $220\,\mu\mathrm{m}$ in the drift direction, i.e. the x-coordinate [17]. The momentum resolution for tracks is between 0.3 and $0.5\,\%$, based on early data. The efficiency as determined on data is in the range 99.2–$99.5\,\%$.

Charge is released in the gas when a charged particle passes through the straw tube. The system as a whole acts as an approximately linear amplifier for this signal, and the factor of proportionality between the original charge produced by the particle and the signal at the front end is known as the gain. Tests have shown that the OT modules can withstand an irradiation of $100\,\mathrm{kHz/cm}$ for 10 years of nominal operation without a loss of gain. It was observed, however, that irradiation with a radioactive source caused a loss of gain in a region upstream (with respect to gas flow) of the most intensely irradiated area [16]. This ageing effect was attributed to an insulating deposit on the anode produced by outgassing from the glue used to build the modules. It is believed that the deposition is prevented in the most irradiated region by production of ozone. This effect was not seen in LHCb operation. In fact, a module which had an area affected by this problem was installed in the OT, and over the 2011 period of running, the gain had reverted to its original value, i.e. the problem was cured. The OT gain is monitored periodically using two independent methods, and as of the end of 2011 there was no evidence of ageing; the OT gain appears to be constant.

2.3.4 Track types

A track is the reconstructed trajectory of a charged particle. Different track types, shown in fig. 2.12, are reconstructed to satisfy different physical requirements [1].

- **Long tracks:** Reconstructed from VELO seed tracks and hits in the T-stations. Hits from the TT are added if available, but are not required. Tracks have full momentum

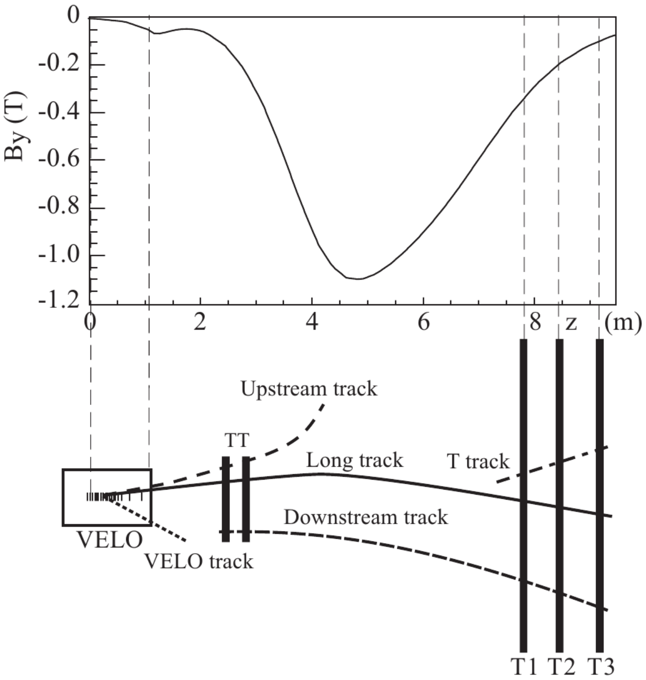

Figure 2.12: Track types in LHCb and magnetic field. The graph shows the approximate y-component of the magnetic field.

information, with good resolution. Long tracks are produced by long lived particles such as electrons, muons and charged pions.

- **Upstream tracks:** Upstream tracks are combinations of VELO seed tracks and hits in the TT. Particles with low momentum will be deflected out of the acceptance while passing the magnet, so they can only be reconstructed as upstream tracks. These tracks are not usually used in physics analyses. Upstream tracks have some momentum information due to the magnetic field between the VELO and the TT, but the resolution is poor.

- **Downstream tracks:** Reconstructed from TT and T-station hits only. Downstream tracks are used for particles which decay after the VELO, but before the TT, such as the K^0_S.

- **VELO tracks:** Tracks with only VELO information, which are primarily used for reconstruction of primary vertices. VELO tracks are usually outside the main LHCb acceptance (because tracks inside the acceptance will also pass through other detectors), and backward tracks are also reconstructed.

- **T tracks:** Tracks from only the T-stations (i.e., the IT and OT). They are not usually used in physics analyses, but the information is used by the RICH 2 calibration.

2.3.5 Tracking performance

The momentum resolution of LHCb is $\Delta p/p = 0.4\,\%$ for tracks with momentum $5\,\mathrm{GeV}$ and increases to $0.6\,\%$ at $100\,\mathrm{GeV}$. The impact parameter resolution is $20\,\mu\mathrm{m}$ for tracks with high p_T. The tracking efficiency depends on the kinematics of the particles, and will be discussed for central exclusive χ_c production in Chapter 6.

2.4 RICH

Ring Imaging Cherenkov (RICH) detectors work on the principle of measuring the Cherenkov light emitted by particles as they propagate through a transparent medium, the radiator. Cherenkov radiation is emitted by particles moving at a speed greater than the phase velocity of light in the medium. The photons are emitted (the light waves interfere constructively) at a fixed angle called the Cherenkov angle relative to the direction of the velocity of the particle. The Cherenkov angle is given by

$$\cos\theta = \frac{1}{n\beta}\ ,$$

where β is the speed of the particle relative to the speed of light in vacuum[1] ($\beta = v/c$), and n is the index of refraction of the medium.

The Cherenkov angle is reconstructed by focussing the cone of light produced by the particle onto a plane of photodetectors using a spherical mirror. This produces a circular pattern with a radius which depends on θ. By measuring the speed using the RICH and the momentum using the tracking system, the particle mass can be computed from $p = \gamma m\beta$ ($\gamma = 1/\sqrt{1-\beta^2}$), and the identity of the particle can be determined.

Two RICH detectors are used [1]: the RICH 1 has sensitivity to low-momentum particles, and the RICH 2 to high-momentum particles. The RICH 1 is located between the VELO and the TT, and covers an angular range from 25 to $300\,\mathrm{mrad}$, limited at the low end by the size of the beam pipe. It has sensitivity to the particle velocity in an approximate momentum range of 2–40 GeV. The RICH 2 only covers the area closest to the beam axis, with an angular acceptance of $\sim$ 15–120 mrad. The reason that this acceptance was seen as sufficient is that the high-momentum particles are concentrated in the area near the beam axis. The RICH 2 momentum range starts from $\sim 15\,\mathrm{GeV}$ and extends beyond $100\,\mathrm{GeV}$.

The Cherenkov radiation is reflected and focussed by spherical mirrors placed in the LCHb acceptance. Flat mirrors are installed to subsequently reflect the light onto the photon detector planes. The photon detectors are Hybrid Photon Detectors (HPDs), con-

---

[1] The speed of light in vacuum is written explicitly as c instead of assuming $c = 1$ in this section, to avoid any ambiguity.

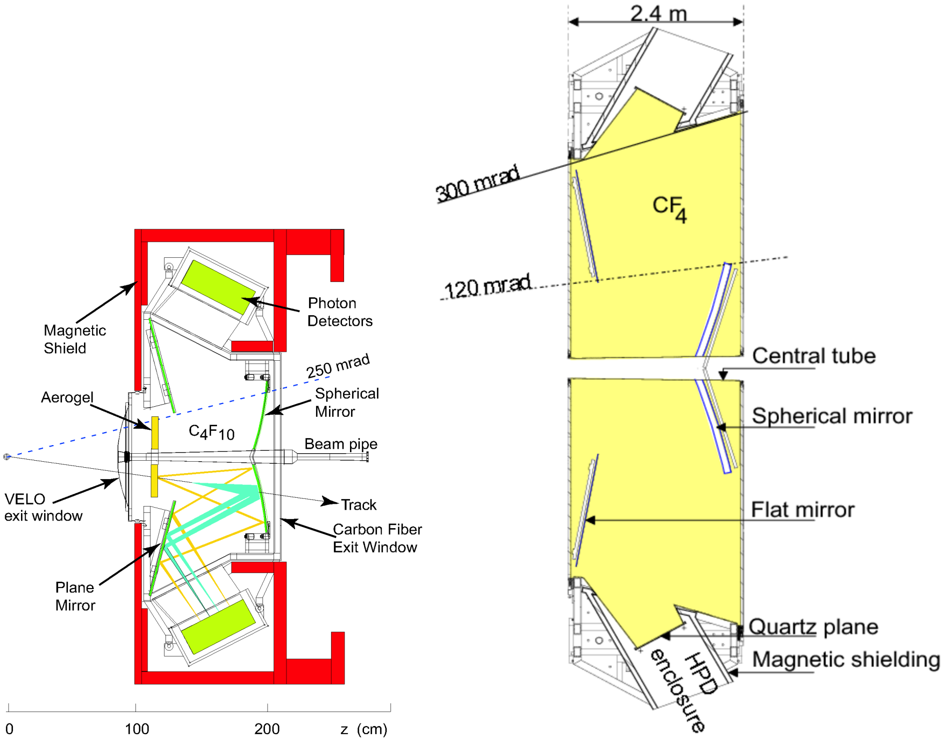

(a) RICH 1

(b) RICH 2

Figure 2.13: Schematic layout of the RICH detectors. The Cherenkov radiation from a particle is also indicated for the RICH 1. The RICH 1 is shown in the yz-plane, and the RICH 2 is shown in the xz-plane (i.e. from above).

sisting of a vacuum tube with an electrostatic accelerating potential, and a silicon sensor with 1024 pixels. The photons interact at the aperture of the tube, producing electrons which are then accelerated over a potential of $20\,\mathrm{kV}$. The electrons are detected by the silicon sensor. The planes of HPDs are enclosed in iron boxes to shield them from the magnetic field, and the individual HPDs are also locally shielded using mu-metal cylinders.

The RICH 1 uses two radiators: silica aerogel and $\mathrm{C_4F_{10}}$. The layout of the RICH 1 is shown in fig. 2.13a. The aerogel is arranged in tiles and placed at the entry point of the gas volume. The aerogel has a refractive index of 1.030 for photons with a wavelength of $\lambda = 400\,\mathrm{nm}$. The $\mathrm{C_4F_{10}}$ gas has a refractive index of 1.0014 at $\lambda = 400\,\mathrm{nm}$ in the conditions used at LHCb [18]. The aerogel gives sensitivity to particles with very low momentum, and studies show that it is possible to reconstruct the rings from two radiators with a single sensor grid. The spherical mirrors are made from Carbon Fibre Reinforced Polymer (CFRP) to have a small material budget. The total amount of material in RICH 1 is $0.08\,X_0$, where X_0 is the radiation length.

The RICH 2 uses a $\mathrm{CF_4}$ radiator with a refractive index of 1.0005. A schematic view of the detector is shown in fig. 2.13b. The photon detectors are located on the left and right side of the detector.

2.4.1 RICH performance

The radii of the Cherenkov rings depend on the position at which the photon is emitted, which is unknown and assumed to be at the centre of the radiator volume. This emission point error causes a loss of precision. The finite spatial resolution of the HPD also causes an uncertainty. The major uncertainty is due to chromatic dispersion in the radiator, which is significantly greater than other errors for the aerogel, and comparable to the others for the $\mathrm{C_4F_{10}}$ and $\mathrm{CF_4}$ radiators. A final uncertainty comes from uncertainty on the parameters of the track. The total single photoelectron angle uncertainties, estimated with high momentum tracks in real data in 2011 [18], are $1.618\,\mathrm{mrad}$ for the $\mathrm{C_4F_{10}}$ in RICH 1 and $0.68\,\mathrm{mrad}$ for RICH 2. The aerogel was found to have an angular resolution of about $5.6\,\mathrm{mrad}$, which is a factor 1.8 worse than expected. This is believed to be due to absorption of $\mathrm{C_4F_{10}}$ in the aerogel, because it is contained in the gas volume. The resolution remained almost unchanged in the gas radiators in 2012, despite the challenges presented by higher luminosity [19]. A new aerogel in a sealed container was installed for the 2012 running period, and this gave a better Gaussian response, with a width of approximately $5.16\,\mathrm{mrad}$. The aerogel provides some proton-pion separation at low momentum. Another important parameter which is monitored is the photon yield per particle. It is lower than expected from MC, but it does not affect the particle-ID performance [18].

The primary function of the RICH detectors is to identify hadrons: $\pi^{\pm}$, $K^{\pm}$, p and $\bar{p}$. The reconstruction is done using a global likelihood model which takes into account all tracks in an event [18]. Initially all tracks are assumed to be pions, because the majority of particles in the event are pions. Different mass hypotheses are then tried for each track. For a given momentum and mass, the expected Cherenkov angle and photon yield are computed. The procedure is then iterated once, with the particle mass set to the best fit in the previous iteration. The end result is a log-likelihood value for each mass hypothesis, relative to the pion hypothesis. The performance is characterised by the positive identification efficiency and the mis-identification rate. For a selected hadron mass, the identification efficiency is the efficiency for correctly identifying hadrons with that mass. The mis-ID rate is the probability to incorrectly identify another particle, i.e. a pion, as the selected particle. The kaon identification efficiency and mis-ID rate are shown in fig. 2.14, for two different cuts on log-likelihood. The mean efficiency and mis-ID rate over the momentum range 2–$100\,\mathrm{GeV}$, for a log-likelihood cut of 0 (i.e. that the likelihood for the kaon hypothesis is greater than that for pion) are $\sim 95\,\%$ and $\sim 10\,\%$ respectively. When instead requiring that the kaon log-likelihood is greater than the pion log-likelihood by five units, the efficiency and mis-ID rate are $\sim 85\,\%$ and $\sim 3\,\%$. The overall particle identification efficiency satisfies the design specifications [18].

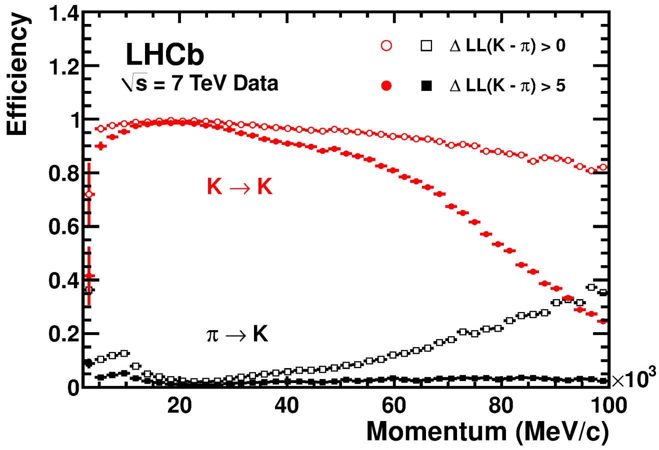

Figure 2.14: Kaon identification efficiency and mis-ID rate on the 2011 dataset (real data) using the RICH detectors.

2.5 Calorimeters

The calorimetry system includes, in the order of increasing z: a preshower system comprising two layers of scintillating pad detectors, an electromagnetic calorimeter (ECAL) and a hadronic calorimeter (HCAL) [1]. The detectors cover the full LHCb acceptance, and each layer is a wall perpendicular to the beam-line. A requirement is that the calorimeters need to identify electrons, photons and hadrons with high E_T for the L0 trigger. This requires the ability to read out the relevant information in less than the L0 latency, which is $4\,\mu$s. The electron trigger requires a high background rejection, demanding a good separation of electrons and π^0. The calorimeters are used for physics analyses to determine the position and energy of particles. The accurate reconstruction of photons and π^0 is especially important for the LHCb physics programme [1].

2.5.1 Scintillating Pad Detector and the Preshower detector

The preshower system includes the Scintillating Pad Detector (SPD) and the PreShower (PS) detector [1], separated by a lead converter, which is $15\,\mathrm{mm}$ $(2.5\,\mathrm{X}_0)$ thick. The purpose of the SPD is to distinguish charged particles from neutral particles, i.e. electrons from photons and π^0. This is done by simply looking for hits in the SPD: if there is a hit in front of the cluster, it was likely a charged particle. A significant inefficiency for photons comes from conversions of photons into electron-positron pairs in the detector material before the SPD, which are then seen as charged particles. For this reason, the SPD veto is not used in the offline photon reconstruction, and analyses which need to trigger on photons may use an OR of the electron and photon L0 trigger. The PS provides separation between photons and π^0 by measuring the transverse shape of the signal. The

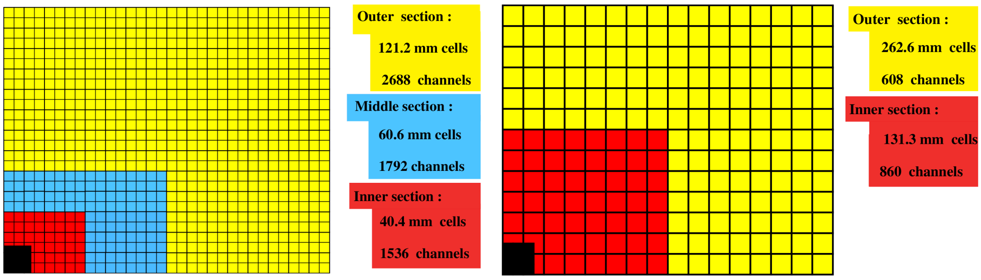

Figure 2.15: Segmentation of the calorimeters and preshower system, shown for a quadrant of the full detectors. The SPD/PS and ECAL segmentation is shown on the left, the HCAL segmentation is shown on the right.

π^0 mesons will produce broader showers, as they decay to two photons (high-energy π^0 cannot be reconstructed as two distinct photons, because they are too close together). The L0 triggers for electrons and photons require that only one or two (up to four in the inner region) cells of the PS were hit [20], to reject the background of π^0.

Both the SPD and the PS are built from scintillating pads, read out via Wavelength Shifting (WLS) optical fibres, connected to Multi-Anode Photomultiplier Tubes. The detectors are almost identical, and the pads are arranged in the same layout as the ECAL cells, in a projective geometry. The segmentation of the calorimeters and the preshower system is shown in fig. 2.15. All dimensions of the SPD are approximately $0.45\,\%$ smaller than in the PS due to the projectivity. The SPD has a binary readout, while the PS records the PMT signal at a 10-bit resolution. The SPD and PS detectors have a total of 12032 channels.

2.5.2 The electromagnetic calorimeter

The ECAL [1] is constructed from alternating layers of scintillator and lead in the longitudinal direction. The scintillation light is collected using WLS fibres, looped inside the scintillator tiles for high light collection. The fibres from each cell are connected to separate PMTs, and the signals are digitised with a resolution of 12 bits. The ECAL is segmented into quadratic cells in the lateral directions (fig. 2.15). Three sections with different cell sizes are used, from $40.4\,\mathrm{mm}$ to $121.2\,\mathrm{mm}$, with the smallest cells near the beam pipe. The ECAL is designed to be able to contain most electron and photon showers, as it has a total depth of $25\,\mathrm{X}_0$.

The gain is tuned to have a uniform E_T (transverse energy) response over the plane of the detector [21]. The energy resolution is measured to be approximately $\sigma_E/E = (8/\sqrt{E[GeV]} \oplus 0.9)\,\%$, where the terms are added in quadrature and $E[GeV]$ is the particle energy. The first term is the stochastic term and the second term is a constant term due to the calibration uncertainty and non-linearity of the readout. This energy resolution

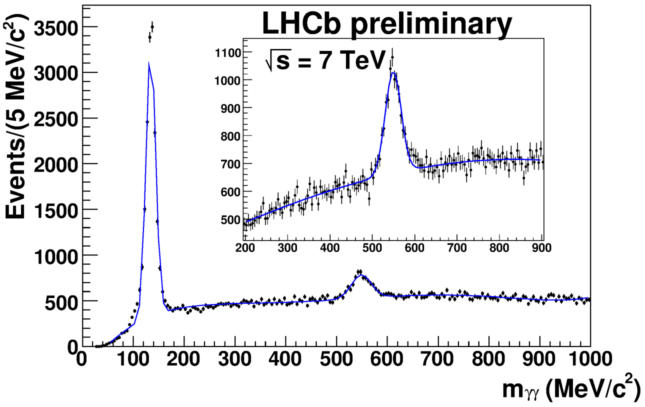

Figure 2.16: Reconstruction of the π^0 and the η using the ECAL in early 2010 data.

satisfies the design requirements.

After installation, the ECAL was first calibrated using Light Emitting Diodes (LEDs) installed near the detectors. When data from real collisions were available, the calorimeter was calibrated using energy flow [21]. The energy in each cell integrated over many minimum bias events is computed, and the detector response is adjusted so that it is a smooth function of position. After more data were available, the π^0 was used to calibrate the response. The peaks from the decays $\pi^0 \to \gamma\gamma$ and $\eta \to \gamma\gamma$ are shown in fig. 2.16, for real data collected early in 2010 (before calibration to the π^0 mass was completed).

2.5.3 The hadronic calorimeter

The hadronic calorimeter [1] is built from iron absorbers and scintillating tiles, running parallel to the longitudinal (beam) direction. A drawing of a HCAL module is shown in fig. 2.17. There are three scintillating pads in the longitudinal direction, which are all read out through the same WLS fibre. The HCAL is constructed from identical periods of transverse size $20\,\mathrm{mm}$. The cells of the HCAL are quadratic and are constructed by connecting multiple columns of tiles to the same PMT. The HCAL contains two regions with different cell size (fig. 2.15). In the inner region the cell size is $131.3\,\mathrm{mm}$ and in the outer region it is $262.6\,\mathrm{mm}$. The HCAL contains $5.6\,\lambda_I$ (interaction lengths) of material, but the ECAL adds another $1.2\,\lambda_I$ before it. This is not sufficient to contain all hadronic showers, but it is enough for the main purposes of the HCAL which are to trigger on high-E_T hadrons and to provide particle identification information.

The HCAL response is calibrated using a radioactive source, brought in proximity with all tiles in turn using a hydraulic system. An absolute calibration was available from the time the detector was installed, but the relative calibration was checked using an energy flow analysis [21].

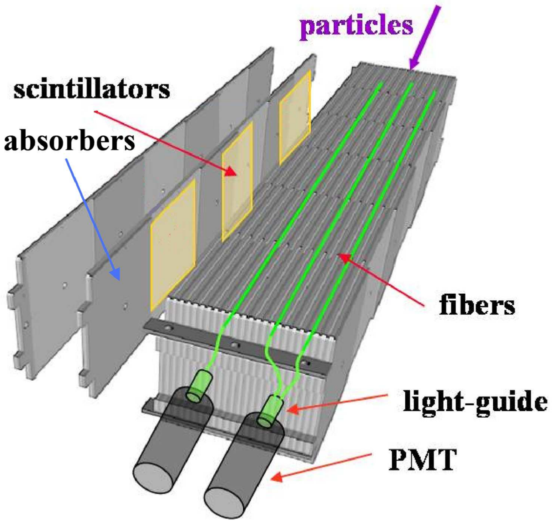

Figure 2.17: HCAL module.

2.6 Muon system

The muon detector system [1] has the important purposes of providing L0 muon triggers, and identifying muons for the high level triggers and physics analyses. It needs to perform an independent track reconstruction for the L0 trigger, because no tracking information is available at that point.

There are five muon stations, named M1 to M5 in order of distance from the interaction point. M1 is located before the calorimeter system, and the other stations are located after the calorimeters. The stations M2-M5 are interleaved with iron absorbers, each with a depth of $80\,\mathrm{cm}$. The total amount of material before M5 is approximately 20 interaction lengths. The positions of the muon stations are shown in fig. 2.18. The muon stations have a projective geometry, such that the angular dimensions as seen from the interaction point are the same for all stations. The detectors are divided into four regions, as shown in fig. 2.19. The region “R1” near the beam has the best resolution, as required by the higher occupancy. The pitch of the individual channels in each region of M1 is shown on the right side of fig. 2.19. The resolution in the inner region is $1\,\mathrm{cm}$ in x and $2.5\,\mathrm{cm}$ in y. Stations M2 and M3 have a resolution improved by a factor two in the x-coordinate, and M4 and M5 have a factor two worse resolution than M1. This is because M1-M3 are used to determine the track direction, and thus the p_T, in the L0 trigger. M4 and M5 are mainly used to confirm the muon identification, selecting for particles that penetrate all the absorbers. The HLT and offline reconstruction do not use information from M1, as they do not need to do a stand-alone tracking in the muon stations.

The detectors are Multi-Wire Proportional Chambers (MWPCs) in all areas except for the inner region (R1) of M1, which uses Gas Electron Multiplier (GEM) detectors.

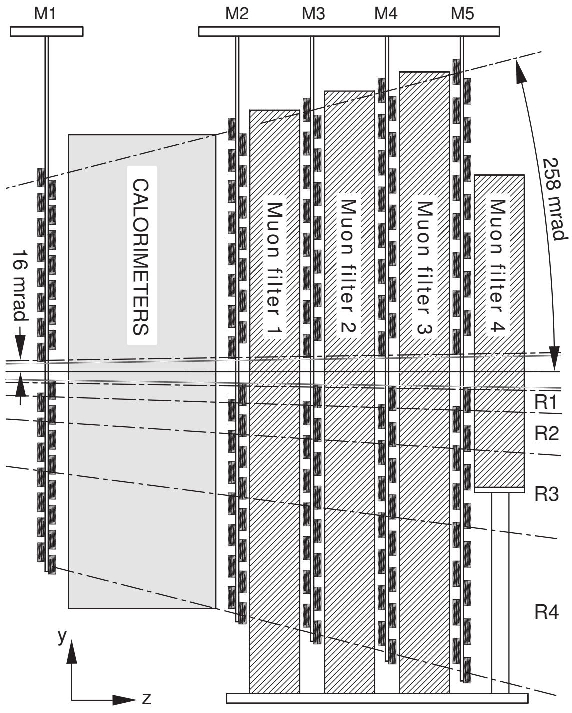

Figure 2.18: Overview of the positions of the muon stations. The beam axis is drawn with a solid black line.

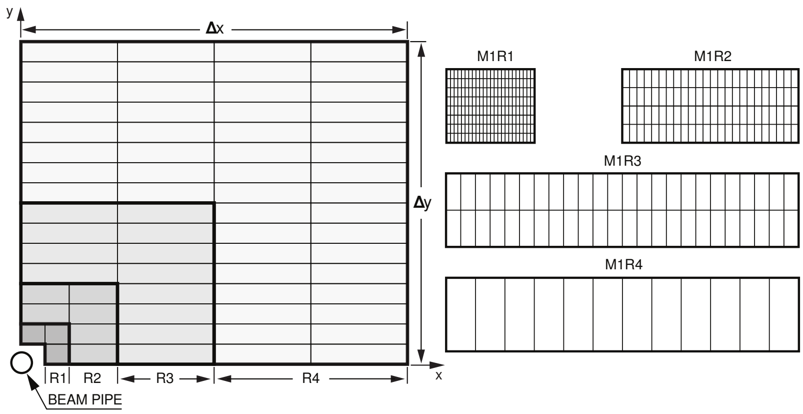

Figure 2.19: Layout of M1. The diagram on the left shows a quadrant of a muon station. The small rectangular boxes represent individual chambers. Each chamber contains multiple logical readout channels, indicated on the right for each region.

The GEM technology is chosen over MWPC because it is able to withstand the high occupancy and harsh radiation environment. The MWPCs in M2-M5 have four redundant gas gaps to guarantee a good efficiency. The material budget of M1 is limited because it is located before the calorimeters, so it only has two redundant gaps. The L0 trigger can be configured to ignore information from M1 in case the efficiency drops below a reasonable level, at the cost of a worse p_T-resolution. This was done for one region in 2010 [22].

2.6.1 Muon system operation and performance

The muon system operated reliably in 2010 [22], with a small number of dead channels, giving an inefficiency of $\sim 1\,\%$. Most dead channels were repaired in the technical stop before the 2011 run. Trips of the high voltage system occurred in single gaps, but thanks to good monitoring and rapid maintenance work, this only caused an effect at the level of $0.1\,\%$ on the efficiency. The GEM detectors of the inner region of M1 have had performance consistent with expectations. The main operational issues with the GEMs have been shorts of a gas layer [23]. This is mitigated by an adjustment of the voltage levels in the GEM. While Ref. [22] is concerned with the early 2010 data, it is noted that the further running experience has been similar.

The muon chamber efficiency was measured [22] using real data in the following ways:

- For the M1, J/ψ candidates were selected and analysed. The events were triggered by a single muon at L0, and a dimuon trigger at HLT.
- For M2-M5, tracks were reconstructed in all muon stations (including M1) except the one "test station" for which the efficiency was to be measured. Hits were searched for at the extrapolated intercept of the track with the plane of test station. The efficiency could then be computed as the number of found hits divided by the number of expected signal hits.

The efficiency is found to be greater than the design requirement of $99\,\%$.

2.7 Trigger

The two levels of the trigger are run in sequence, and are discussed separately in the following.

2.7.1 The Level 0 trigger

The first trigger level is called Level 0, or L0 [1]. It runs synchronously with the $40\,\mathrm{MHz}$ LHC clock. The L0 trigger is required to reduce the event rate to less than

$1.1\,\mathrm{MHz}$, which is the maximum rate at which the full detector can be read out. The L0 trigger is designed to select a sample of events which is enriched in B-decays. The primary discriminator used is the transverse momentum, p_T. Daughters of B-decays typically have high p_T, because the mass of a B-hadron is significantly greater the sum of the masses of its decay products. Most proton-proton interactions ("minimum-bias events") have many particles with low p_T.

The L0 trigger takes information from the Pile-up system, the calorimeters and the muon detectors. The trigger algorithms are implemented on special-purpose circuits and FPGAs, to support the very high event rate. Each subdetector has readout electronics which reconstruct a small number of relevant quantities, which are then provided to the L0 Decision Unit (L0DU).

The information from the calorimeter system is [20]:

- The number of hits in the SPD. This is approximately equal to the number of charged particles in the acceptance. A cut requiring less than 600 SPD hits is used for most triggers to reject "busy" events, as they use a disproportionate amount of computing resources and storage. A much lower cut of 10 SPD hits is used for the L0 *low multiplicity* triggers, which are specifically designed to select events in which few particles are produced, such as central exclusive production. The SPD cut is combined with other conditions.

- The electron and photon candidate from the ECAL with the highest transverse energy, because they are candidates for interesting physics processes. Clusters of 2x2 cells are used as particle candidates.

- The clusters with the highest and second highest E_T in the HCAL, for the same reason as above. If there is a highest-E_T ECAL cluster in front of the HCAL cluster, the energy of the ECAL cluster is added.

- The sum of all the energy in the HCAL is computed, as this can be used as a minimum-bias trigger (i.e. to reject events without visible interactions).

The sensors in the muon stations are arranged in projective towers pointing to the interaction point. The L0 muon trigger reconstructs tracks in the muon stations [1] by looking for towers of hits pointing to the interaction point in the y coordinate. The algorithm starts from hits in M3, and initially assumes that the track points to the interaction point in x as well. A field of interest (FoI) is opened around the extrapolated position in M2, M4 and M5 in the x-direction. If a hit is found in M2, M4 and M5, a candidate track is formed. The extrapolated position in M1 is taken as a straight line through the hits in M2 and M3. The M1 hit closest to the intercept is chosen. The knowledge of the track direction from the hits in M1, M2 and M3 allows the measurement of the momentum. The muon p_T can be determined with a resolution of approximately $25\,\%$ [20]. The algorithm is limited to muons with $p_T \gtrsim 500\,\mathrm{MeV}$. Each of the four quadrants of the muon detectors are processed separately. The processor for each quadrant reports the two muons with the highest p_T for each quadrant. The L0DU can cut on the p_T of the muon with the highest p_T, or the product of the p_T of the muons with the highest and second highest p_T: $p_T^{\text{largest}} \times p_T^{\text{2nd largest}}$.

The Pile-up system trigger processor performs a rudimentary vertex reconstruction based on the radial slopes of tracks in the pile-up sensors, and can provide a veto on events with multiple primary vertices at L0 [1]. It was however not desirable to use such a veto in the running conditions up to the end of 2012 [20].

The L0DU makes decisions based on the data described above, and the cuts are configurable using software in the control room. The L0 decisions are transferred to a Readout Supervisor (RS) board [20], which manages the overall trigger rate, such that the buffers on the front-end electronics and the DAQ system are not overwhelmed. The RS also accepts events randomly at a set rate, in the “no-bias” triggers. A set fraction of no-bias events is recorded with beam-beam crossings, with single beam passages, and without any beams (e.g. the time slots $25\,\mathrm{ns}$ after a collision). The no-bias events are selected purely based on the LHC filling scheme, and no activity is required in the detector.

2.7.2 The High Level Trigger

The HLT is the second level trigger[1], and is a C++ application called Moore. It runs on the Event Filter Farm (EFF), which is a cluster of computers located at CERN. The available EFF resources in 2011 and 2012 were 15440 CPUs, running 26110 copies of the HLT application [20]. The computing resources are sufficient to support a $30\,\mathrm{ms}$ mean processing time for the full HLT (events rejected by HLT1 are much faster, events processed by HLT2 are slower). The HLT receives the full L0 output rate, which was $870\,\mathrm{kHz}$ in 2011 [20], and has to reduce it to a rate which can be handled by the offline system. The first level, HLT1, reduced the rate to approximately $43\,\mathrm{kHz}$, but there is no external requirement on the HLT1 rate. The output rate of HLT2 was $3\,\mathrm{kHz}$ in 2011. In 2012, the mean rate was typically around $4\,\mathrm{kHz}$, and reached $4.5\,\mathrm{kHz}$ [24] for some fills. As the HLT is based on the same software framework as the offline reconstruction application (to be discussed in Section 2.8), it shares many of the algorithms, but it runs simplified versions of the time-consuming algorithms. The most important physics triggers are described in the following, but there are also other triggers for minimum bias, luminosity monitoring and error monitoring.

HLT1 runs the VELO tracking in almost the same configuration as the offline reconstruction [20]. If an event was accepted by an L0 muon or dimuon trigger, a fast

---

[1] Some of the older references refer to a “Level 1”-trigger, but this is not used, and can be considered as equal to the HLT1.

VELO-muon tracking is performed. If there is no L0 muon trigger, VELO tracks are selected for further analysis by a cut on the minimum impact parameter with respect to any primary vertex. The selected VELO tracks with muon information or with a large impact parameter are processed with the forward tracking algorithms, to search for hits in the IT and OT. A smaller window around the extrapolated position is used compared to offline, effectively imposing a momentum cut. The track is fitted if a sufficient number of hits is found in the T-stations. The "track" physics triggers of HLT1 take an inclusive approach, looking for only a single track, and requiring a large IP and p_T, and some other quality cuts. Some complementary triggers ("dimuon") do not require a large IP, and instead look for two muons, with a cut on the mass of the muon pair. These triggers are useful for analyses which cannot bias the lifetime distribution of B-decays, and for triggering on particles which do not have a measurable non-zero lifetime, such as the Z-boson. The low multiplicity triggers at L0 have a low enough rate that a pass-through trigger can be used in HLT1, accepting any event selected by the L0 low multiplicity triggers.

HLT2 runs a more complete reconstruction sequence [20], reconstructing exclusive decay chains for many processes. A large number of B-meson based analyses use instead the *topological trigger* at HLT, which reconstructs inclusive n-body B-candidates using kinematic variables (for $n = 2, 3, 4$). The topological trigger uses a boosted decision tree to make a trigger decision.

There are multiple low-multiplicity selections in HLT2, including muon and hadron final states. These triggers only process events which are accepted by the L0 low multiplicity triggers. The particles from L0 are refined, requiring good-quality tracks. The HLT2 low-multiplicity selection requires that there is no reconstructed primary vertex in the event, and that there are no backward VELO tracks.

Every time the configuration or algorithms of the trigger are changed, the full configuration is stored in a persistent form in a database and assigned a unique number, the Trigger Configuration Key (TCK). The currently active TCK at the time of data taking is saved inside all raw data files. Information is also saved about exactly which hits contributed to the trigger decision, using unique identifiers for each detector element. This makes it possible to determine offline which candidates triggered the event, but it also enables a data-driven efficiency determination [20]. A complete emulation of the L0 trigger is available in software, which can be run on MC for efficiency studies. The HLT can be run directly on MC, as the detector simulation already produces data in a similar format to that seen by the HLT when running online.

The trigger has had excellent performance, and has been able to adapt to the rapidly changing running conditions.

2.8 Computing

Operationally, the computing resources are separated into online and offline systems. Online systems support the day-to-day running and monitoring of the detector, and include the EFF, a monitoring and calibration farm, desktop computers in the control room, a computing cluster and miscellaneous servers. The offline system refers to everything that is done to the data after the triggered events have been saved to files: reconstruction and user analysis. It also includes Monte Carlo (MC) production. Offline systems run asynchronously with the LHCb operation, but some tasks are designed to process data files as they become available. Most of the offline computing and storage is handled by the LHC Computing Grid (or "the grid"), which is a distributed network of computing nodes and storage. The management of grid resources is done by the Dirac software.

The computing model of LHCb [25] is based on the Gaudi framework [26] and the LHC Computing Grid. Gaudi is written in C++ and sets up the basic structure for all applications. Specifically, it introduces these objects:

- Algorithms: objects containing code that is run once for every event.

- Tools: objects with public functions which can be called by algorithms and other tools. Tools typically provide a specific functionality, such as a kinematical calculation. Types of tools are declared with a public interface. Multiple tools can implement the same interface, and the specific implementation of the interface can be selected when the job is started.

- Services: similar to tools, but services satisfy more basic requirements, such as reading data from a file or accessing the detector geometry. Services typically have more internal state variables than tools, and can react to incidents such as advancing to the next event.

- Event data: Data are stored in a virtual hierarchical directory called the Transient Event Store (TES). The TES contains event-specific data, and is repopulated for each event. Algorithms can read from and/or write to the TES. Specific TES locations can be read from, or written to, permanent storage.

- Detector data and conditions: Data about the detector are stored in a database, which is accessed in read-only mode by user jobs. The database is divided into two main sections: The detector description database contains information about the geometry and material of the detector. These data are approximately constant in reality, and only change as the understanding of the detector improves. The other section is the conditions database, which stores data which change over time, such as the magnet polarity and the pedestals of the VELO sensors.

Main applications	
Gauss	MC simulation
Boole	Digitisation; simulation of detector response
Moore	High level trigger, and L0 trigger emulation
Brunel	Event reconstruction
DaVinci	Stripping and physics analysis
Bender	Python-based physics analysis
Panoramix	Visualisation of events and detector geometry
Online	Reporting and monitoring tools running on the online system
Utilities, not LHCb-specific	
Ganga	Distributed computing user interface
Dirac	Distributed computing middle-ware
ROOT	Framework for numerical calculations, statistical analysis, data management and fitting

Table 2.1: Main software applications used for LHCb analyses. Monitoring applications for the subdetectors are not listed.

Gaudi makes use of functionality from, and is integrated with, the ROOT framework [27]. Gaudi includes a facility to configure algorithms and tools using scripts written in the Python programming language. The C++ class declares any properties that should be configured from Python at compile time. For each algorithm, tool or service, a "shadow" object called a Configurable is created. The Python script sets properties of the configurable, which are then copied to class members in the C++ class when the job is started. Gaudi also includes tools for creating ROOT histograms and n-tuples (i.e., tabular data structures). Gaudi manages the input and output of data files.

The LHCb applications are listed in table 2.1, and an overview of the data flow for real data and simulation is shown in fig. 2.20. The chain starts either with real data or with the Gauss MC production software. The real data are saved in "RAW" files, which contain data blocks which are almost identical to the ones output by the TELL1 boards, and information added by the trigger. The RAW files are copied to the primary grid sites. The Brunel reconstruction software is then run, producing a reconstructed data file, in a standard file format managed by Gaudi called the Data Summary Tape (DST). Brunel runs tracking and particle identification algorithms. In addition to the data added by the trigger and by Brunel, the DST also contains the full data of the raw event. The events are then processed by the stripping, described below.

MC production includes similar steps. First, interactions are simulated using Gauss. Gauss includes two phases: generation and simulation. Generation calls a standard event generator, by default PYTHIA 6 [28]. Gauss can work with other generators, and it can also read events from a text file in the HepMC ASCII format. The simulation phase sets

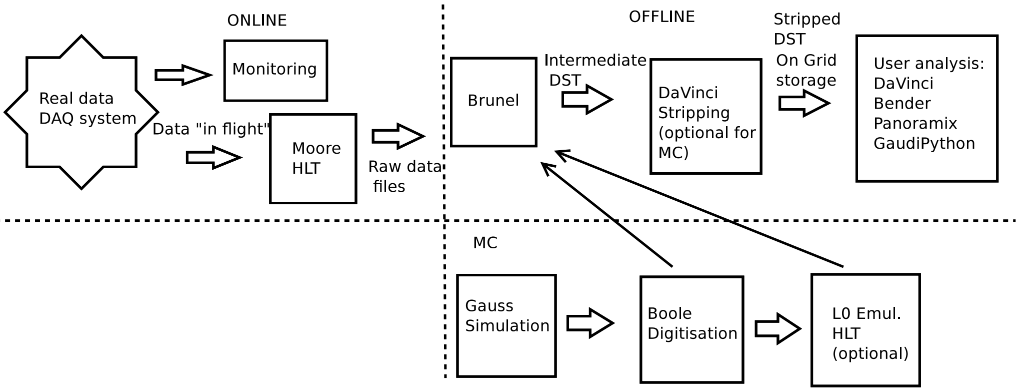

Figure 2.20: LHCb applications and data flow for reconstruction of real data and simulation. The real data processing is indicated in the top row. The raw data files are stored permanently and redundantly in multiple geographical locations. For MC, it is possible to skip the trigger, and the trigger may also be run after the reconstruction step. It is also not necessary to run the stripping on MC.

up a geometrical model of all the material in the LHCb detector. The passage of particles through the detector is simulated using GEANT 4 [29]. The input to GEANT 4 is all the final state particles produced by the event generator. The response of the detector elements, such as the charge in the silicon strips, is subsequently simulated in the application Boole. Boole produces a file which can be read by Moore or Brunel, containing the simulated detector hits and optionally also *MC-truth*-information, i.e. information about the simulated particles. If trigger information is desired, Moore is run on the file before Brunel. Moore includes code to emulate the L0 trigger, and this is run before the main HLT algorithms. Because of the modular structure supported by Gaudi, different applications can share code quite easily. For example, the same tracking algorithms are available for the HLT and in the offline analysis, but they are configured differently using properties.

2.8.1 Stripping

The events accepted by the HLT are saved permanently, but another level of filtering is performed to make more efficient use of the computing resources. The output of the reconstruction software is analysed by the stripping, which runs many selection algorithms, reconstructing specific decay chains. The stripping selections are divided into about 10 “streams”, based on the organisation of working groups for different physics topics. The stripping jobs write separate output files for each stream, creating subsets of data based on different physics interests. There is for example the Electroweak stream, which contains events with $Z \to \mu\mu$, $Z \to ee$, Drell-Yan dimuon production, $W \to \nu_\mu\mu$, central exclusive production / low multiplicity selections, and other selections.

2.8.2 User analysis

There are multiple options for analysing the physics data. The main software is DaVinci. Selection of specific decays is usually implemented using standard algorithms, which allow any cuts and combinations to be specified in the options (properties). After selection, the most common procedure is to use a standard algorithm called DecayTreeTuple to save information about each selected decay in a ROOT n-tuple. Users can opt for the more low-level approach of writing Gaudi algorithms, and DaVinci provides a more convenient interface than the bare Gaudi algorithm interface.

Panoramix is a visualisation application which can display the detector geometry, detector hits such as clusters, tracks and MC-truth information. The geometry information is taken from the same database that is used for simulation.

2.9 Luminosity

LHCb is designed to work at a luminosity of $2 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$ [1], which is lower than the peak luminosity at ATLAS and CMS. This is done to have a lower number of interactions per event μ. Running at lower μ improves the identification and separation of primary vertices and decay vertices. The luminosity at LHCb is primarily controlled by modifying the collision optics, i.e. the field in the magnets that focus the beams near LHCb. This corresponds to modifying β^* in (2.1.1). Another handle for controlling the luminosity is the number of colliding bunches in LHCb, n_b, which is smaller than in ATLAS and CMS, but this has the downside of retaining a high μ. The beams can also be displaced so that they collide off-centre, thus reducing the effective intensity, and in that case (2.1.1) does not apply. It was originally envisioned to run with $\mu \sim 0.5$, but significant work on the trigger and reconstruction software made it possible to run with μ as high as 2.5 at the end of 2010. In 2011, the frequency of collisions was high enough that μ could be reduced, while still increasing the luminosity. The mean μ in 2011 was 1.4. The LHCb luminosity reached its design value in April 2011, but it could be increased beyond that while retaining good trigger and reconstruction performance. At the end of 2011 the detector was running stably at $4 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The luminosity was kept at this value in 2012.

The benefit of using displaced beams to control the luminosity is that the repositioning can be done while the machine is running[1]. This makes it possible to do *luminosity levelling* [30], which is the process of controlling the luminosity and keeping it constant throughout a fill. The beam intensity naturally decays exponentially with time over a fill, but the corresponding decrease of the luminosity can be counteracted with luminosity levelling: the beams are moved closer together in small steps as the intensity decreases.

---

[1] It is also possible to change β^* while the machine is running, but it is more difficult.

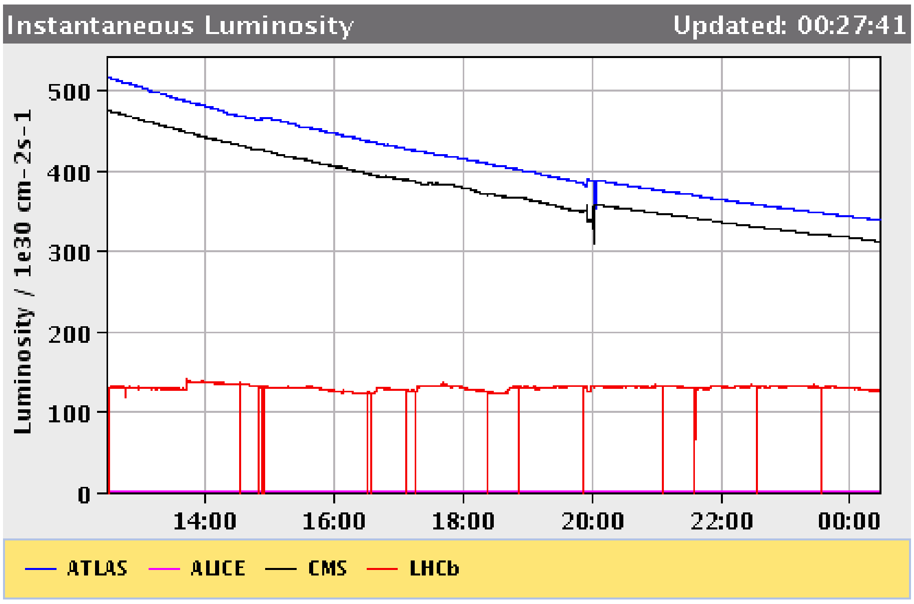

Figure 2.21: Instantaneous luminosity for the four major experiments as it developed during part of a fill in April 2011.

The effect of luminosity levelling is shown for a fill from early 2011 data in fig. 2.21. Instead of having a decaying instantaneous luminosity at LHCb, the optimal luminosity can be kept during the whole fill.

2.10 Conclusion

After three years of operation and at the start of the long LHC shutdown in early 2013, we were in a good position to reflect on the performance of LHCb. The detector was operated at an average luminosity twice as high as originally envisioned (peak luminosity greater by a factor 1.5), with μ at a factor of three higher than the design value. The bunch spacing was kept at $50\,\mathrm{ns}$, twice as high as the nominal value. The energy was about half of the design value, reducing the $b\bar{b}$ cross section by a factor of ~ 0.6 relative to the one at $14\,\mathrm{TeV}$. The yield of $b\bar{b}$ during this first period was expected to be 10^{12}, and this was approximately what was delivered. A total integrated luminosity of $3.3\,\mathrm{fb}^{-1}$ was collected.

The summary of the operation is that all subdetectors performed as expected or better than expectations, and there were no major problems. The VELO IP resolution and the Cherenkov angle resolution of the RICH 1 aerogel were worse than expected, but it did not significantly impact the physics performance. The performance must be seen in the context of the conditions in 2011 and 2012. The additional pile-up made reconstruction more difficult, and increased the demand on the trigger. The greater instantaneous luminosity also demanded a fast trigger. *Spillover* refers to having a residual signal in the detector from the previous event (it is usually an effect in the readout electronics, not from actual particles left in the detector). Because of the $50\,\mathrm{ns}$ bunch spacing, spillover was not a concern. The effects of radiation damage are clearly visible in the VELO, and are progressing at a pace consistent with expectations. The effect of charge loss to the second metal layer in VELO R-sensors was not expected, but it has not had an impact on tracking.

There are many analyses at LHCb which have produced world-class results. LHCb has released a measurement of the branching ratio of $B_s^0 \to \mu\mu$ and a limit on $B^0 \to \mu\mu$ [31]. This result places strong constraints on supersymmetric models for new physics. The CKM angle γ has been measured using decays of $B^\pm \to DK^\pm$ (where D is a D^0 or a $\bar{D}^0$) [32, 33] . A range is given of $\gamma \in [55.1, 79.1]^\circ$ at $68\,\%$ confidence level [33]. LHCb has also measured mixing in the B^0-$\bar{B}^0$ and B_s^0-$\bar{B}_s^0$ systems using only semileptonic B-decays [34]. The hypothesis of no B_s oscillations is rejected with significance 5.8 standard deviations. Because the LHCb acceptance is in the forward region it can provide information about the parton momentum distribution in protons which is complementary to other experiments. The production cross sections of electroweak gauge bosons can provide such information. The Z and W cross sections were measured using the 2011 dataset in final states containing muons [35], tau leptons [36] and electrons [37].

Despite the high precision of current LHCb results, and the even better statistics that will be available by 2018, it may not be sufficient to completely rule out new physics models [38]. The precision for many analyses of B-decays is limited by statistics, so the main challenge is how to scale up the luminosity. The only way to scale the luminosity is to have a more efficient L0 trigger, and this is the proposal for the LHCb upgrade [39]. The upgraded detector will support a readout of the full detector at $40\,\mathrm{MHz}$, and a software trigger at the first level. Having a more selective trigger will allow it to run at a higher luminosity, approximately $10^{33}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

CHAPTER

THREE

VELO MONITORING SOFTWARE

3.1 Monitoring systems

During the running period of 2010–2012 the VELO was monitored constantly to ensure that it operated reliably and had a good efficiency. The purposes of the monitoring systems were to notify the VELO personnel of potential problems before they had an effect on data quality, to urgently alert them of problems which would affect data quality and to prevent damage to the VELO.

The LHCb control system is called the Experiment Control System (ECS). The ECS is used to control the detectors, the data acquisition system, the support infrastructure such as cooling, vacuum pumps and power supplies, and the computing jobs of the trigger and monitoring systems. The ECS is also used to read out sensors for pressure and temperature placed around the detectors, and other environmental factors such as the current and voltage of power supplies. All data are logged. Alarms and automatic actions can be issued if sensor readings go beyond a safe range.

The online monitoring software [40] is the next level of monitoring after the ECS, and is the first level of analysis of the data produced by the detector. The software comprises standard Gaudi jobs running on the monitoring computer farm. It receives a subset of all events as they are accepted by the trigger. Histograms are produced and displayed in the control room after a delay of a few seconds.

The central offline monitoring system produces histograms for all events which are accepted by the trigger, but there is a delay of order hours or days before the plots are available. The offline plots provide more detailed information thanks to the greater statistics available.

As outlined in Section 2.3.1, there is also a separate VELO-specific offline monitoring system, and this is the main focus of this chapter. The VELO offline monitoring is run after each completed run. Runs are the basic unit of data taking for LHCb, and last for up to one hour. The DAQ system copies events accepted by the no-bias trigger to a dedicated

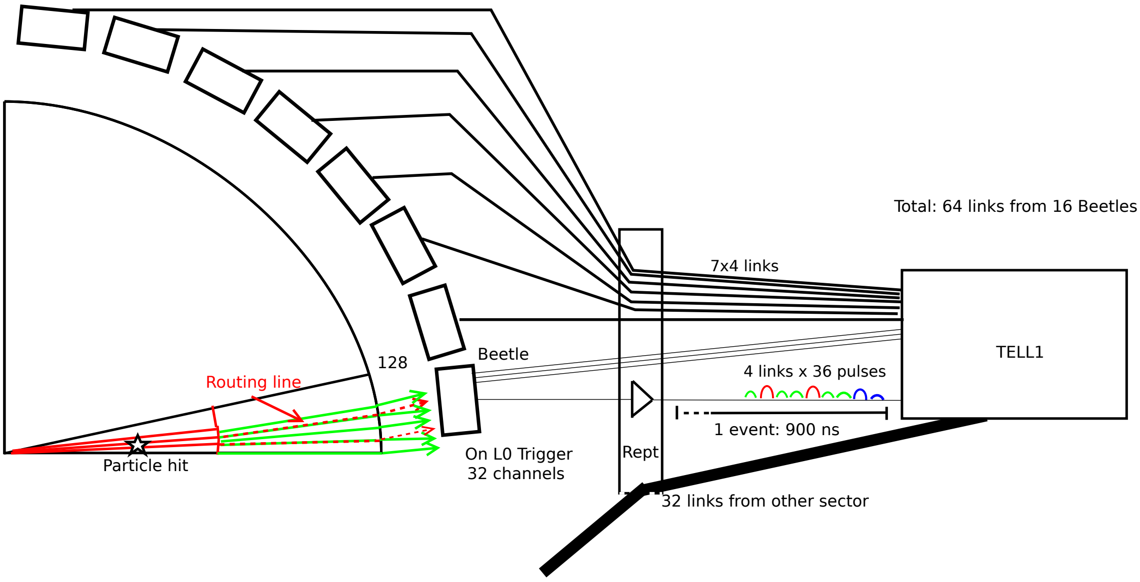

Figure 3.1: Diagram of a half VELO Phi-sensor with Beetle chips, a repeater board and corresponding analogue links. Some strips are shown on the sensor. In red: inner strips. The dotted red lines represent the routing lines, which are located directly above the every other outer strip. The outer strips are shown in green. The size of the strips is exaggerated to make it possible to display them. Groups of links are represented by thick lines, and the four links of the bottom Beetle chip are drawn separately. The pulses of the signals are indicated on one link, showing the order in which they are transmitted. A particle hit is included as an example, with correspondingly higher pulses for the inner strips.

storage area. As was mentioned in Section 2.3.1, the TELL1 [11] boards regularly emit Non Zero-Suppressed (NZS) banks, and these data are also saved in separate files. The VELO offline monitoring software performs an analysis of these files after the run has ended.

3.2 Data acquisition

A slightly more detailed review of the VELO data acquisition system is in order before discussing the monitoring. The Beetle [10] front-end chips sample the charge in every strip every $25\,\mathrm{ns}$. Each sensor has 2048 strips, read out by 16 Beetle chips and multiplexed onto four links per chip. A diagram of the front-end readout for a single Phi-sensor is shown in fig. 3.1. Some strips are drawn, to show the order in which the strips are connected to the Beetle (the Phi-sensors have twice as many outer strips as inner strips).

The signals are amplified and stored in a pipeline, and then transmitted over analogue links if the event is accepted by the L0 trigger. The signals are first sent to a repeater board [1] just outside the VELO vacuum chamber, then over $60\,\mathrm{m}$ long cables to the TELL1 boards. There is one repeater board and one TELL1 board for each sensor. The Beetle chip first sends four digital signals, the header bits, which indicate the position in

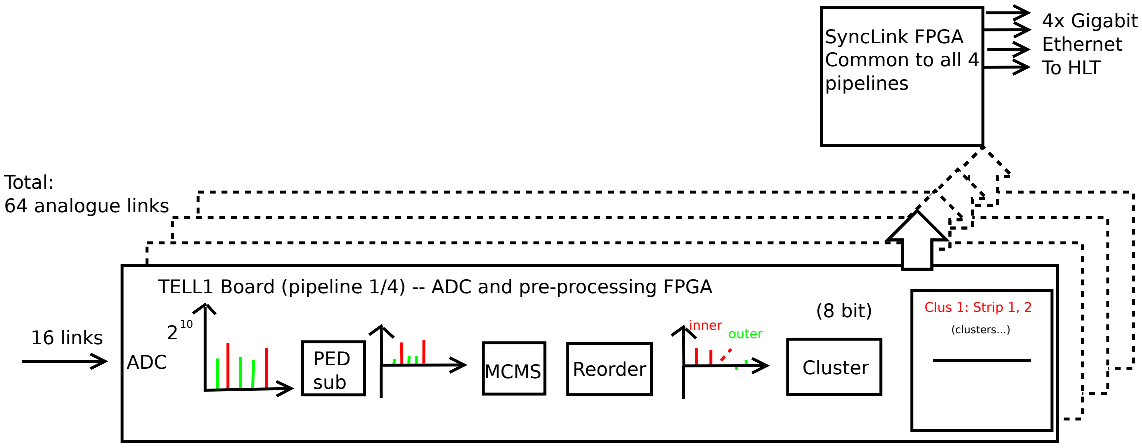

Figure 3.2: Block diagram of the TELL1 data processing, also including a graphical representation of some signals as they pass through the processing chain. The pipeline is replicated in four pre-processing FPGAs, each processing the signals from 512 strips.

the pipeline buffer of the event. It then sends the 32 analogue levels of the strips, with one signal sent every $25\,\mathrm{ns}$.

Four Analogue to Digital Converter (ADC) plug-in cards are mounted on each of the TELL1 boards for the VELO, each with 16 links. The thresholds of the ADCs are set such that a value of approximately 512, the centre of the dynamic range, is read out when no signal is present. The signal from a single minimum ionising particle is about 40 ADC counts, as seen in Section 2.3.1, corresponding to about 2×10^4 electrons at the front end [41]. The ADC value for zero signal is referred to as the pedestal. The pedestal is slightly different for each channel (strip), and it varies more between links, by typically 30 ADC counts.

After digitisation, the signals are processed by algorithms implemented on the FPGAs on the TELL1 board [8]. The algorithms are run in parallel for four groups of 512 channels. A diagram of the processing is shown in fig. 3.2. First a pedestal subtraction algorithm is run. A fixed offset is subtracted from each ADC count. The pedestal values are generally different for every channel, and do not change during normal operation. They are determined from data taken without beam in the machine, and have to be re-evaluated approximately every two months [8]. After the pedestal subtraction, a mean common mode suppression algorithm is run. Common mode noise is noise that coherently affects multiple adjacent channels. It can be produced by noise from the power supply, the Beetle, or RF-pickup in the analogue cables. The algorithm determines the mean ADC for channels without signals in a link, and subtracts this value from each channel. The next step is to reorder the channels such that physically adjacent strips are read out in order. The final step is to look for clusters (clusterisation). The algorithm first searches for strips with signals above a seeding threshold of 10 ADC counts. For each seed strip, it is attempted to expand the cluster by including neighbouring strips with a lower threshold, up to a maximum cluster size of four strips. The data volume is significantly reduced by only saving clusters and not the full ADC values of all the strips. The clusters from all the four processing pipelines are sent to an FPGA which buffers and prepares the data, and emits Ethernet packets to the DAQ system. The VELO TELL1 boards can emit the following data types [42]:

- Zero-Suppressed (ZS): Contains clusters. Each cluster has information on the strip numbers and ADC values of all strips in the cluster.
- Non Zero-Suppressed (NZS): Data are provided directly from the ADCs, without pre-processing or clusterisation.
- Error: Error banks are sent by the TELL1 to report if there are synchronisation errors in the pre-processing FPGAs.
- Pedestal: Pedestal banks report the pedestal for every channel on a TELL1, but these data were not used in normal operation.

The raw data from all subdetectors are combined, and sent to the HLT nodes. The data banks are then decoded to produce C++ objects for clusters (ZS) and channels (NZS). Only the ZS data are used in most analyses and in the HLT.

3.3 VELO offline monitoring and the Vetra application

The offline monitoring application for the VELO and the ST is called Vetra [43]. Its main components are:

- Decoding of NZS data
- Monitoring of NZS and ZS data
- Bit-perfect emulation of the TELL1 algorithms

A large amount of code is also shared with the reconstruction application Brunel and the HLT, such as the decoding of ZS data and the VELO tracking algorithm. Vetra is used for monitoring, calibration of gain and timing, and for testing of modifications to the TELL1 algorithms. Vetra outputs a ROOT file with its results, in the form of histograms, which can be opened directly or processed by other VELO software. Some specific monitoring plots will be discussed in the following sections.

The VELO offline monitoring is performed automatically by a program which is scheduled to run every 15 minutes on the online system. It looks for new runs and starts one ZS and one NZS analysis (Vetra) job for each sufficiently long run with physics data.

3.4 Software testing

A suite of software tests has been developed for the monitoring algorithms of the Vetra project [44]. Tests are small computer programs that interact with other software, to verify that the main software functions correctly. Testing is considered an important part of software development. The Vetra tests are based on a central LHCb testing framework [45], which is again built with the QMTest software[1].

The LHCb software is under continuing development, and changes to the core LHCb software or Gaudi can affect Vetra in unexpected ways. The Vetra tests invoke Vetra on a selection of standard data files, and detect unexpected changes to the monitoring histograms, the log messages written to the standard output and any error messages. Vetra itself is under active development, and the tests have to be updated every time there is an intentional change to the output, such as the addition of a new histogram.

The tests run in three phases:

1. Check of options files. Options files are job steering files written in Python, which create instances of configurables and set properties (defined in Section 2.8). The tests verify that the options files have a valid syntax and that the referenced objects and properties exist.
2. Vetra jobs with no data. These tests verify that the application loads correctly.
3. Monitoring with sample data. These tests run short Vetra jobs on a selection of data files, and compare the output to a set of corresponding references. The references are generated from the output of successful Vetra jobs, in a manual procedure.

All the LHCb source code is contained in a repository, and it is compiled automatically every night. The tests are then run on the compiled software, providing a daily report of possible problems. The report is available through a web page, showing compiler warnings and failed tests. The results are given separately for different hardware architectures, software platforms and compilers.

The Vetra tests have identified a problem which caused an incorrect version of the conditions database to be used, and this was fixed. Multiple configuration problems resulting from the integration of new changes to the LHCb software platform into Vetra have been discovered and corrected.

3.5 The VELO Monitoring Graphical User Interface

The VELO Monitoring Graphical User Interface (VeloMoniGUI) [46] is used to display and analyse information provided by the Vetra monitoring jobs. It is a C++ application written to work with ROOT (but not using Gaudi). It comprises a core Graphical User

---

[1] `http://www.codesourcery.com/qmtest`

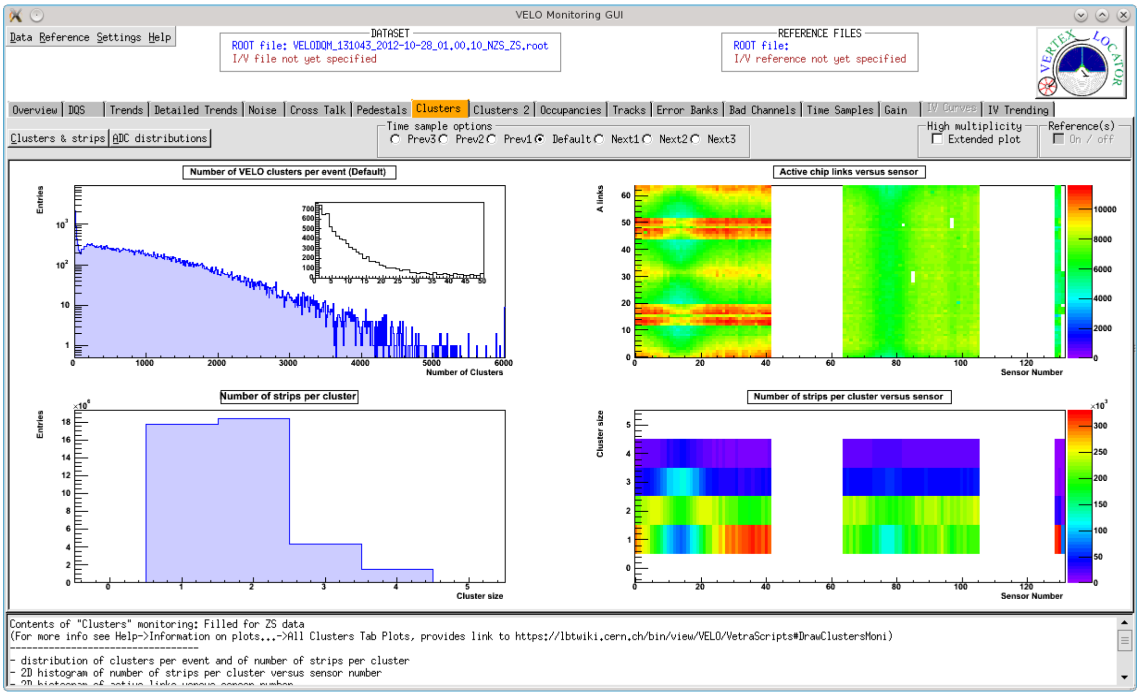

DATASET
ROOT file: VELODQM_131043_2012-10-28_01.00.10_NZS_ZS.root
I/V file not yet specified

REFERENCE FILES
ROOT file:
I/V reference not yet specified

Time sample options
○ Prev3 ○ Prev2 ○ Prev1 ● Default ○ Next1 ○ Next2 ○ Next3

High multiplicity
☐ Extended plot

Reference(s)
☐ On / off

```
Contents of "Clusters" monitoring: Filled for ZS data
(For more info see Help->Information on plots...->All Clusters Tab Plots, provides link to https://lbtwiki.cern.ch/bin/view/VELO/VetraScripts#DrawClustersMoni)
----------------------------------
- distribution of clusters per event and of number of strips per cluster
- 2D histogram of number of strips per cluster versus sensor number
```

Figure 3.3: Main window of the VeloMoniGUI application, with the Clusters tab selected.

Interface (GUI) system, a collection of modules (ROOT macros) which process ROOT files and produce plots, and corresponding adaptor functions which display the plots. The core system creates a single window with a tabbed interface and a menu bar, using the GUI classes provided by ROOT. It lets the user open a ROOT file, which is then used by all the analysis modules. The user can also select another ROOT file as the *reference file*, to compare two runs. An example view of the main window of the VeloMoniGUI application is shown in fig. 3.3, with the tab for cluster-based monitoring activated.

The collection of data analysis modules is quite heterogeneous. The modules are implemented in either C++ or Python. The code is usually invoked using functions which take as parameters the current ROOT file, the reference file and a ROOT `TPad` object. The latter is a representation of the area on the window in which to draw the plots. The modules are written by the experts on the relevant domain, and are updated frequently based on experience with new data. The analysis modules for normal data quality monitoring are:

- Noise monitoring per strip using NZS data.
- Crosstalk: this monitors a small crosstalk effect between the header bits and the first channels on each link.
- Monitoring of the pedestal subtraction using NZS data.
- Cluster-based quantities from ZS data, including multiplicity, size and ADC distribution for clusters, for R- and Phi-sensors separately.

- Cluster ADC distributions per sensor.
- Occupancy, at the level of the full VELO, sensor, r-coordinate, strip and beam bunch.
- Tracks and vertices: spatial distribution of track multiplicity, tracking efficiency, residuals, primary vertex position.
- Error monitoring.
- Identification of bad channels using the occupancy.

Other modules perform analysis of special calibration runs: scans for bad strips, timing scans, gain scans and current-voltage scans.

The VeloMoniGUI produces a Data Quality Summary (DQS) based on the analysis of a file. The DQS is a collection of key numbers for all aspects of VELO data quality. The plots for each run are analysed manually by a VELO operator, and the DQS for each run is saved. The VeloMoniGUI has a function that analyses all DQS files and presents trending graphs, showing the values of the DQS quantities as they change over time.

3.6 Monitoring developments and experience

The VELO monitoring software was in an advanced state when data taking commenced in March 2010. Since then, a significant effort has been expended on improving the software and incorporating new analyses code, based on operational experience.

3.6.1 High-multiplicity events

It was discovered that a small number of events had an unexpectedly large number of VELO clusters. The phenomenon was first noticed due to a problem with the trigger and reconstruction jobs: that the processing of these events took a very long time, and the reconstruction jobs were sometimes terminated due to excessive memory use. It was important to establish the origin of the high-multiplicity events. The clusters could either be produced by an effect of the instrumentation or by real particles. If they were from real particles, the source of the particles had to be determined.

The distribution of the number of clusters per event in the no-bias dataset of 2011 is shown in fig. 3.4. The events are accepted by a random trigger (without any trigger bias), but with a fixed fraction of events with each of four bunch crossing types: events where a bunch from each beam passes through the interaction point (crossing), events where a bunch passes in either direction (beam 1 or beam 2), and events where there are no bunches in LHCb (empty). The beam moving in the direction of the LHCb acceptance,

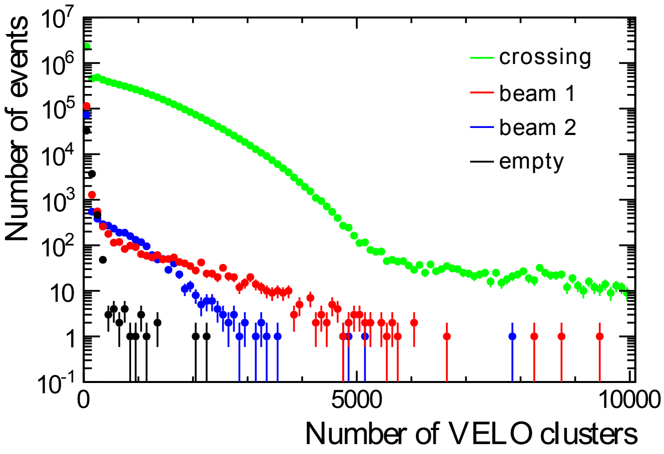

Figure 3.4: Number of VELO clusters per event in no-bias data, on a logarithmic scale. The data are plotted separately for beam-beam crossings, single beam events and events without beam.

i.e. entering in the upstream part of the VELO and exiting by the muon stations, is called beam 1, and the beam moving in the opposite direction is beam 2.

In fig. 3.4 there appears to be one component of beam crossing events which has a decreasing distribution, and does not produce much more than 6000 clusters. Those events are the normal proton-proton interactions. There is also a plateau which falls off much more slowly, and which corresponds to the high-multiplicity events. The second component is not seen in minimum bias MC, supporting the hypothesis that it has a different origin than proton collisions at the interaction point. The plateau is also seen in single-beam events, but there are not enough data to determine if it is present in events without beams. The moderately large number of clusters in some events with a single beam is due to interactions of beam protons with gas molecules in the VELO vacuum chamber or in the beam pipe near the detector. These beam-gas interactions produce showers of particles which can be detected by the VELO, and this is a well understood process.

Events with more than 6000 clusters are selected for further study. The events are written out to a DST file and are then analysed with Vetra and the VeloMoniGUI. A repeating pattern with a period of 128 channels is seen in the distribution of the number of clusters per channel. Figure 3.5 shows this pattern for events with more than 6000 clusters, and for “normal” events for comparison. Because the pattern repeats with the period of the number of channels per Beetle chip, this indicates that it is an instrumentation-related effect. It is checked that the full pattern is seen for single events, not just some of the peaks. Approximately the same pattern is seen in all sensors, including R-sensors and Phi-sensors, but the number of clusters per sensor is not distributed uniformly. One could still imagine that an extremely large number of real particles could cause this effect, because the readout could break down when the occupancy is too large. It is therefore interesting to

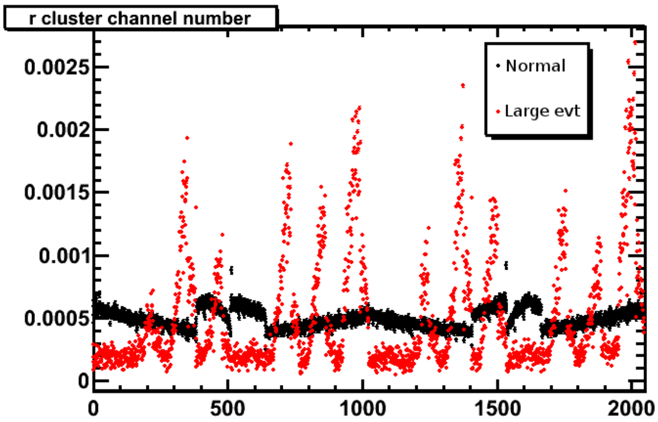

Figure 3.5: Channel number of clusters on R-sensors in normal events and high-multiplicity events. The graphs are scaled to have unit area.

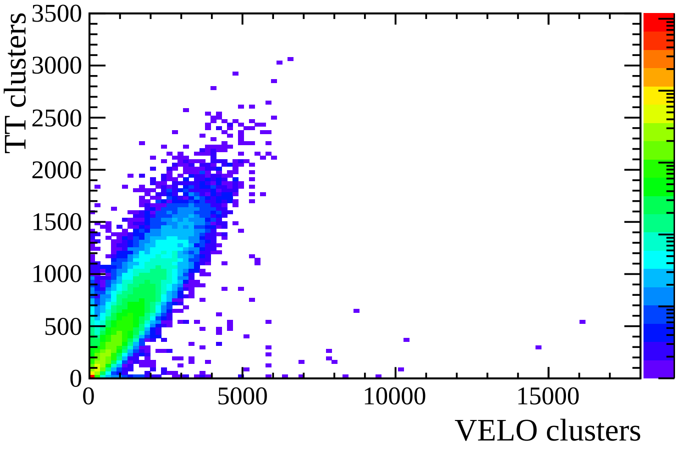

Figure 3.6: Correlation of TT and VELO cluster multiplicity. The plot is made for a specific run. A logarithmic scale is used for the number of events (colour scale), because the number of events varies by five orders of magnitude over the histogram.

compare the VELO cluster multiplicity with the amount of activity in other detectors. The number of clusters in the TT is plotted versus the VELO cluster multiplicity in fig. 3.6, for no-bias data from a run from August 2011. The TT multiplicity is not proportional to the VELO multiplicity for the events considered in this section, with more than 6000 VELO clusters. This indicates that the effect is not due to a large particle shower. It is not completely conclusive as the VELO and the TT acceptance do not overlap in the lateral coordinates (there could be a highly collimated shower of particles that was not seen by the TT).

The no-bias dataset produced for the monitoring jobs contains many more events per run than the no-bias stripping, but it is only temporarily accessible (all the raw data accepted by the HLT are stored permanently, but they are not easily accessible). The no-bias

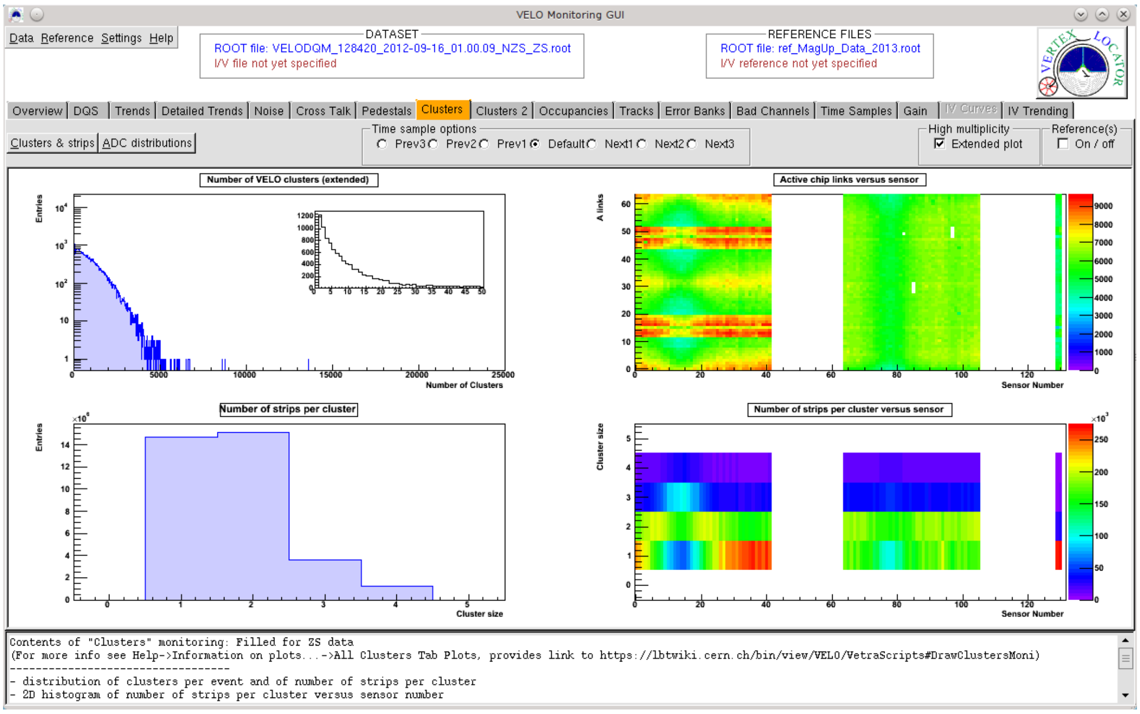

DATASET
ROOT file: VELODQM_128420_2012-09-16_01.00.09_NZS_ZS.root
I/V file not yet specified

REFERENCE FILES
ROOT file: ref_MagUp_Data_2013.root
I/V reference not yet specified

Time sample options
○ Prev3 ○ Prev2 ○ Prev1 ● Default ○ Next1 ○ Next2 ○ Next3

High multiplicity
☑ Extended plot

Reference(s)
☐ On / off

```
Contents of "Clusters" monitoring: Filled for ZS data
(For more info see Help->Information on plots...->All Clusters Tab Plots, provides link to https://lbtwiki.cern.ch/bin/view/VELO/VetraScripts#DrawClustersMoni)
----------------------------------
- distribution of clusters per event and of number of strips per cluster
- 2D histogram of number of strips per cluster versus sensor number
```

Figure 3.7: The VeloMoniGUI clusters page. It includes a function to view an extended cluster multiplicity plot, selected using the option “Extended plot” in the upper right area. The scale of the top left plot is then extended to 2.5×10^4. (The inset plot is zoomed in for *low multiplicity* events, and not related to this specific issue)

monitoring dataset for a run in 2011 is used to map out in which time slots (bunch crossing IDs) the high-multiplicity events occur. The most important observation is an “empty” event which was found to have over 6000 clusters and the pattern discussed in the previous section, and which is separated from beam crossing events by $50\,\mathrm{ns}$. This indicates that the presence of beam is not necessary for the production of high-multiplicity events. It also precludes the possibility that the effect is induced by the next or previous bunch crossing, as there was no beam either $25\,\mathrm{ns}$ before or $25\,\mathrm{ns}$ after this event. The conclusion is that the high-multiplicity events are produced by a form of noise in the VELO which affects all sensors, and is not directly related to the beam.

In practice it is possible to cut on the number of clusters to avoid having to reconstruct these events. A feature was added to the VeloMoniGUI to monitor high-multiplicity events: the scale of cluster multiplicity can be extended up to 25000 clusters, by selecting a check-box “Extended plot”. This is shown in fig. 3.7. The number of high-multiplicity events per 10^4 events is included in the DQS. The trending plot is shown in fig. 3.8. The statistical precision is somewhat low because only about 2×10^4 events per run are analysed in the monitoring jobs, but the rate appears to be stable over the proton runs (the rate is higher for the recent heavy ion runs, but this is expected due to the changed running conditions). The rate of high-multiplicity events is about 1 per 10^4 events. Another

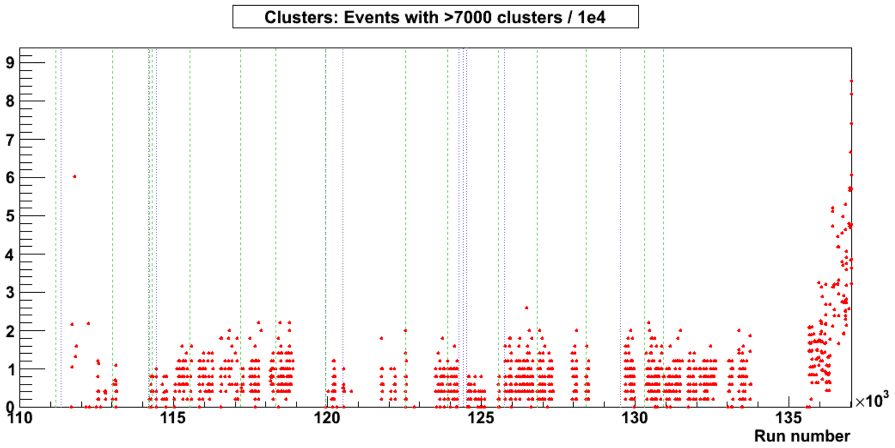

Figure 3.8: Number of events with more than 7000 clusters per 10^4 events. One red dot is shown for each run. The plot includes the runs of 2012 and the heavy ion runs of 2013. The heavy ion data start at approximately run 135×10^3.

source of high multiplicity events is caused by collimated jets of particles coming in approximately parallel to the beam, hitting a small area of the VELO sensors. These events are discussed in Ref. [8].

3.6.2 Support of the investigation of the second metal layer effect

As outlined in the last paragraph of Section 2.3.1, a loss of efficiency is observed in the outer strips of the R-sensors, caused by a diversion of charge from the strips to nearby routing lines (connections between inner strips and Beetles). The effect only manifests itself after irradiation. It results in a lowering of the cluster ADC values on the outer strips, and a corresponding small decrease in efficiency. It also causes production of additional clusters with low ADC values on the inner strips, due to the charge deposited on the routing lines.

There were multiple lines of evidence leading up to the discovery of the effect. Extra clusters with low ADC values were observed on the plot of cluster ADC, as can be seen in fig. 3.9. The upward trend of the ADC distribution towards the threshold of 10 ADC counts is only visible in R-sensors, and appears to be a tail of a large noise component centred at a lower value. We now know that this is from clusters produced by the charge deposited on the routing lines. The GUI tools can also show an ADC distribution for only clusters which are used in tracks. In fig. 3.10 it is seen that the clusters with low ADC counts are only present on the raw distribution, not in the one that shows clusters on tracks. This provided an additional clue that the clusters were not from real particles hitting the strips.

The ADC distributions for clusters on each R-sensor and Phi-sensor are fitted with a Landau function convolved with a Gaussian. The most probable value of the Landau

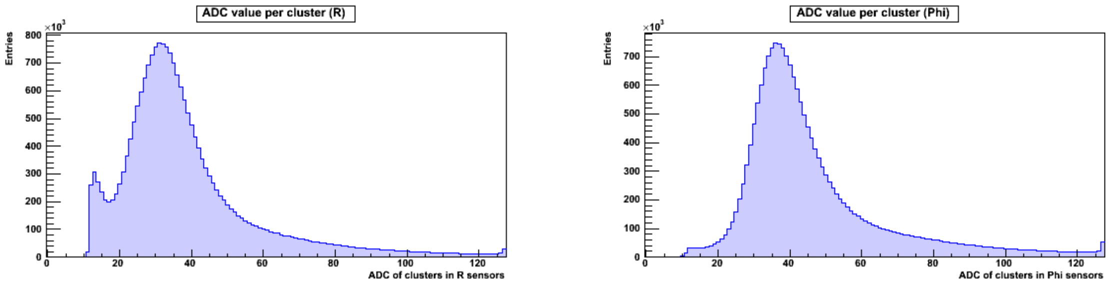

Figure 3.9: ADC distributions for clusters on R-sensors (left) and Phi-sensors (right). These histograms are based on a run from late 2012, but the same features were visible before then.

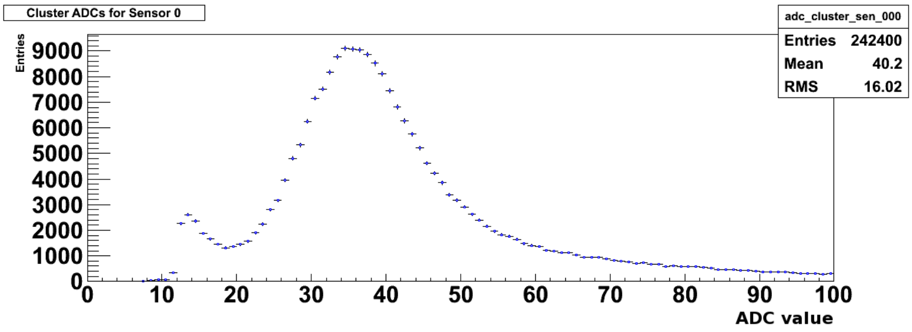

(a) Raw ADC

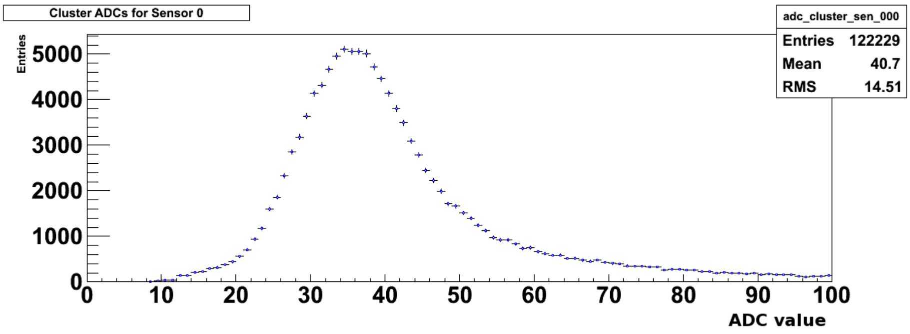

(b) On track

Figure 3.10: ADC distributions for clusters on sensor number 0, which is an R-sensor. The horizontal axis shows the ADC value, and the vertical axis shows the number of clusters. Figure (a) includes all clusters, while figure (b) includes clusters on tracks.

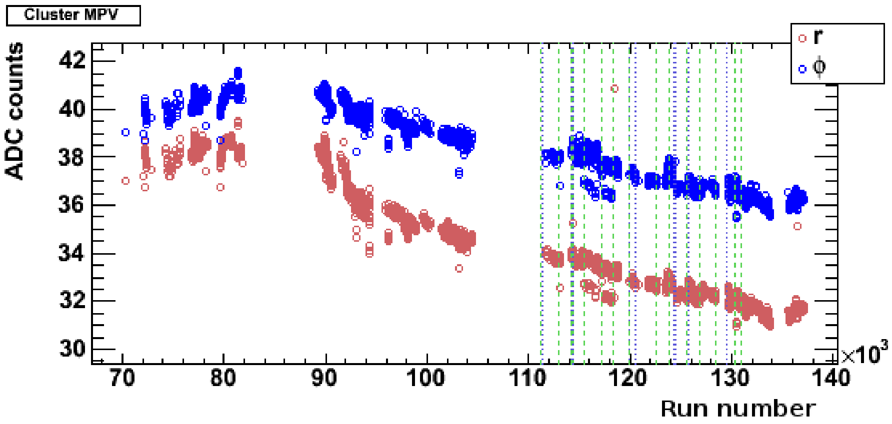

Figure 3.11: Trend plot of the most probable value of the Landau distribution, as a function of run number. The horizontal axis shows the run number. The runs of 2011 start at 87×10^3, the runs of 2012 start at 111×10^3, and the runs above 135×10^3 are heavy-ion runs from 2013 (proton-lead, lead-proton). The blue points correspond to all Phi-sensors, and the red points to all R-sensors. The vertical lines indicate technical stops and changes of the DAQ configuration.

distribution is recorded in the DQS, and the trend is plotted using the VeloMoniGUI. This is shown in fig. 3.11. As the accumulated fluence increases, the ADC distribution for the R-sensors shifts faster towards lower values than for the Phi-sensors, due to the second metal layer effect. This information was however not instrumental in determining the nature of the effect.

While the standard monitoring software provided an important alert and supporting information, it was far from enough to conclude on the issue. A major role was played by the analysis of the cluster finding efficiency and charge collection efficiency, done in the context of assessing radiation damage. It was this analysis which initially saw a drop of the efficiency on the outer strips of the R-sensors. It later provided the “smoking gun” two-dimensional plot of the cluster finding efficiency shown in Section 2.3.1. A further analysis quantifies the dependence of the effect on the distance of clusters from the nearest strip and the nearest routing line [12].

3.7 Conclusion

The VELO monitoring software has operated reliably throughout the running period, and there have been few problems with the VELO itself. The rapid changes of running conditions, such as the luminosity and trigger rate, caused many of the monitored quantities to vary significantly, requiring ongoing evaluation by the experts.

The monitoring functions of the Vetra software were shared with other applications, imposing a strict requirement on the quality of the source code used. The development of the VeloMoniGUI was more experimental and there were many people who needed to add and change functionality, resulting in a higher rate of stability problems (i.e. unexpected program terminations). The interface between Python and C++ source code needed to be changed on multiple occasions. The majority of the VeloMoniGUI was initially run in the ROOT C++ interpreter, but some parts had to be compiled ad-hoc using ROOT commands. In 2011 it was necessary to compile these parts of the VeloMoniGUI source code in the deployment process. Later, all the source code was updated such that the complete VeloMoniGUI software could be compiled, and this had a positive effect on the stability.

CHAPTER
FOUR

THEORETICAL AND EXPERIMENTAL REVIEW

4.1 Introduction

In this chapter the Central Exclusive Production (CEP) process is introduced. The theoretical understanding of the process is based on the Standard Model of particle physics. In Section 4.2 some general concepts in particle physics and the Standard Model are reviewed. The theory of CEP is then discussed in Section 4.3, and various models for the three major processes of CEP are presented. A brief experimental review is then given in Section 4.4, focussing on recent results from hadron colliders, which are comparable to the analysis in Chapters 6–8. The prospects for analyses based on CEP are discussed in Section 4.5.

4.2 The Standard Model of particle physics

The Standard Model of particle physics is a theory of the dynamics of matter and the fundamental forces of nature. It describes the interactions of three of the fundamental forces: electromagnetism, the strong nuclear force and the weak nuclear force. Gravity is also a fundamental force of nature, but it is not described by the Standard Model. For a complete descriptions of Quantum Field Theory and the Standard Model see Refs. [47, 48]. The Standard Model has stood up to increasingly more precise experimental tests, with only minor adjustments necessary. It is possible to make very precise predictions for many processes involving Standard Model particles, and the experiments have generally been in agreement with the predictions. Some important details relevant to the topic of CEP are discussed in the following sections.

The Standard Model is formulated using the Lagrangian formalism of mechanics. It is expressed as a covariant (special relativistic) and quantum mechanical theory. Elementary particles are represented by excitations of fields, with one field for each type of particle. The theory of the electromagnetic force in the Standard Model is called Quantum Electrodynamics (QED), and the theory of the strong force is called Quantum Chromodynamics (QCD).

4.2.1 Particles of the Standard Model

The elementary particles of the Standard Model can most generally be separated into bosons and fermions based on their spin quantum number, which is an intrinsic property of the particles. Bosons have integer spins, and fermions have half-integer spins. For each particle there is also an antiparticle, which has identically the same quantum numbers and mass, except that all charges have the opposite sign. Some neutral particles may be their own antiparticle.

The bosons of the Standard Model are the gauge bosons and the Higgs boson. Gauge bosons have spin $S = 1$, and correspond to the fields of the fundamental forces of nature. The gauge bosons (and their symbols) are the photon (γ), which mediates the electromagnetic force, the $W^{\pm}$ and Z bosons, which mediate the weak force, and eight gluons (g) for the strong force. There are two W bosons, W^{+} and W^{-}, which have respectively positive and negative electric charge, of the same magnitude as the electron charge. All of the other gauge bosons are electrically neutral. The mass of the W bosons is $80.385\,\mathrm{GeV}$, the mass of the Z boson is $91.1876\,\mathrm{GeV}$, and the photon and gluons have zero mass. The photon, the gluons and the Z boson are their own antiparticles. The W^{+} boson is the antiparticle of the W^{-} and vice versa. The Standard Model predicts that there is a single neutral Higgs boson with spin zero, but many (hypothetical) extensions to the Standard Model predict more than one Higgs boson, or bosons with different properties. A Higgs boson has recently been observed at ATLAS [49] and CMS [50], with a mass of $125\,\mathrm{GeV}$. The particle is consistent with the Standard Model Higgs boson.

Fermions are the main constituents of the normal matter in the universe. Some key properties of the elementary fermions, collected by the Particle Data Group [51], are shown in table 4.1. There are two main types of fermions: leptons and quarks. The main difference between these types is that quarks interact with the strong force, and leptons do not. The properties of the strong force are such that isolated strong charges cannot exist, so quarks are always found in bound states.

Leptons interact with the weak and electromagnetic forces. There are three “flavours” of leptons, the electron, muon and the tau particle, and there is a charged lepton and a neutrino for each flavour. The neutrinos are electrically neutral, so they only interact through the weak force.

Quarks interact through the strong force, the weak force and the electromagnetic force. There are three generations of quarks. Each generation has an “up-type” quark with electric charge $+\frac{2}{3}$, and a “down-type” quark with charge $-\frac{1}{3}$. The up-type quarks are named up, charm and top, and the down-type quarks are named down, strange and bottom. These

	Name	Symbol	Q	Mass (MeV)
leptons	electron	e^-	-1	$0.510998928 \pm 0.000000011$
	electron neutrino	ν_e	0	$< 2 \times 10^{-6}$
	muon	μ^-	-1	$105.65836715 \pm 0.0000035$
	muon neutrino	ν_μ	0	$< 2 \times 10^{-6}$
	tau	τ^-	-1	1776.82 ± 0.16
	tau neutrino	ν_τ	0	$< 2 \times 10^{-6}$
quarks	up quark	u	$\frac{2}{3}$	$2.3^{+0.7}_{-0.5}$
	down quark	d	$-\frac{1}{3}$	$4.8^{+0.7}_{-0.3}$
	charm quark	c	$\frac{2}{3}$	$(1.275 \pm 0.025) \times 10^3$
	strange quark	s	$-\frac{1}{3}$	95 ± 5
	top quark	t	$\frac{2}{3}$	$(173.5 \pm 0.6 \pm 0.8) \times 10^3$
	beauty/bottom quark	b	$-\frac{1}{3}$	$(4.65 \pm 0.03) \times 10^3$

Table 4.1: Properties of the fermions of the Standard Model. The electric charge is labelled Q. Both terms “bottom” and “beauty” are used to refer to the b-quark.

different types are known as quark flavours, and there is a corresponding flavour quantum number for each flavour. Quarks are always found in bound states called hadrons, because the strong potential grows with distance, such that the potential energy of a free quark would approach infinity. This phenomenon is known as confinement. The charge of the strong force is called the colour charge. It is different from the electromagnetic charge in that it has three components, or colours, called red, green and blue.

4.2.2 Discrete symmetries

Some of the forces are invariant under certain global transformations of the coordinates or the quantum numbers, which can be applied to a system. Invariance means that the physical results are the same after the transformation has been applied.

- Parity ($\mathcal{P}$): All spatial coordinates are inverted: $\mathbf{x} \to -\mathbf{x}$.
- Charge conjugation ($\mathcal{C}$): The signs of all additive quantum numbers, such as the electromagnetic and strong charges, are inverted.
- Time reversal ($\mathcal{T}$): Inverts the time coordinate $t \to -t$ ($x^0 \to -x^0$ in four-vector notation).

The coordinates or charges return to their original values when these transformations are applied twice, so the only possible eigenvalues of the operators are ± 1.

Elementary and composite particles may be eigenstates of the parity and charge conjugation operators. The intrinsic parity is defined to be positive for vacuum. The quarks are defined to have positive parity and the antiquarks negative parity. A system of two particles with parities P_1 and P_2 in a state with orbital angular momentum L has parity

$$P = P_1 P_2 (-1)^L \ .$$

This follows from the symmetry properties of the angular momentum eigenfunctions (spherical harmonics). The parity of the photon is -1 because the photon corresponds to the classical vector potential, which is a vector quantity, and all vectors transform as the position vector $\mathbf{x} \rightarrow -\mathbf{x}$ under parity transformations.

Only neutral particles can be eigenstates of the charge conjugation operator. The charge conjugation of the photon is negative, as can be derived from electromagnetic theory. The charge conjugation of a pair of a fermion and an antifermion with orbital angular momentum L and total spin S is

$$C = (-1)^{L+S} \ .$$

The charge conjugation operator exchanges the fermion and the antifermion, and the sign of the eigenvalue is determined by the symmetry property of the spin and angular momentum wavefunctions.

The product (or composition) $\mathcal{CPT}$ is an exact symmetry operator of the complete Standard Model. The electromagnetic and probably the strong interaction are symmetrical under all of the transformations individually, but the weak interaction is not. The product $\mathcal{CP}$ is an approximate symmetry of the weak interaction, but the violation of this symmetry is measured accurately in hadrons with an s quark (kaons) and hadrons with a b quark.

4.2.3 Hadrons

Hadrons are systems of quarks bound by the strong force. The known hadrons are

- Systems of one quark and one antiquark, called mesons.
- Systems of three quarks or three antiquarks called baryons.
- Exotic hadrons, which are resonances which are not consistent with the two categories above.

The main constituent quarks of hadrons are called valence quarks. The quantum numbers of hadrons, such as the charge, are determined by the valence quarks. Hadrons also contain pairs of quarks and gluons which are created from the vacuum and which recombine and annihilate after very short times. These ephemeral quarks are called “sea quarks”. The sea quarks and gluons can be probed by high energy scattering experiments.

Mesons contain (as valence quarks) a quark and an antiquark. Mesons with a pair of a quark and an antiquark *of the same type* are special in that they have no net flavour. There are many such states for each quark flavour, with different spins, orbital angular momenta and masses. Mesons with quarks $c\bar{c}$ and $b\bar{b}$ are referred to as charmonium and bottomonium respectively, and collectively as quarkonia.

The J/ψ particle is a $c\bar{c}$ meson with the quantum numbers of the photon, $J^{PC} = 1^{--}$ (that is, total spin 1, negative parity and negative charge conjugation). The mass of the J/ψ is $3096.916\,\mathrm{MeV}$ [51], which is the next to lowest mass for a charmonium meson. The J/ψ cannot decay to mesons with a single charm quark because the lightest pair of such mesons, $D^0\bar{D}^0$, has a higher mass than the J/ψ mass. The J/ψ thus has to decay through either annihilation or weak, flavour-changing interactions (it can also decay by emission of a photon to the η_c, which is the $c\bar{c}$ meson with the lowest mass, but the branching ratio is small). This gives the J/ψ a relatively long lifetime, and a narrowly peaked mass distribution in decays, with a width of $92.9\,\mathrm{keV}$. A family of mesons of particular importance to the studies in this thesis are the χ_{c0}, χ_{c1} and χ_{c2}, which will collectively be referred to as χ_c. Their respective masses are $3414.75 \pm 0.31\,\mathrm{MeV}$, $3510.66 \pm 0.07\,\mathrm{MeV}$ and $3556.20 \pm 0.09\,\mathrm{MeV}$, and their total angular momenta, parities and charge conjugation quantum numbers J^{PC} are 0^{++}, 1^{++} and 2^{++}. The quark spins are parallel, and their orbital angular momentum is $L = 1$, i.e. a P-wave. The combination of the spin and orbital angular momentum J is in a triplet state, and can be either 0, 1 or 2, corresponding to χ_{c0}, χ_{c1} and χ_{c2}. The χ_c mesons can decay to a J/ψ through a radiative decay

$$\chi_c \rightarrow \gamma J/\psi \; .$$

There is a $b\bar{b}$ meson, the χ_b, which has the same spin states as the χ_c, and which can decay to $\gamma\Upsilon(1S)$.

Some hadrons have quantum numbers which are not consistent with the quark model. A famous category is glueballs, which are hypothetical states of only gluons, without valence quarks. The existence of glueballs has not been confirmed in experiments. Another exotic hadron is a resonance called $X(3872)$ [52] with a mass approximately equal to the mass of a pair of charmed mesons $D^{*0}\bar{D}^0$. The width of the mass distribution of a meson so close to the threshold for production of $D^{*0}\bar{D}^0$ meson pairs is expected to be large if the state has quantum numbers compatible with decaying to $D^{*0}\bar{D}^0$. The $X(3872)$ has such quantum numbers, but still has a narrow mass peak, so it is considered likely to be an exotic state. The mass spectrum of $\psi(2S)$ and $X(3872)$ candidates at LHCb, minus the J/ψ mass, is shown in fig. 4.1. Building on results from earlier experiments, LHCb has uniquely determined the quantum numbers to be 1^{++} [53]. The dynamical explanation is still a mystery, but candidate theories [52] are that it could be an "ordinary" $\chi_{c1}(2^3P_1)$ charmonium state, a $D^{*0}\bar{D}^0$ "molecule", a mixture of a molecule and a quarkonium state,

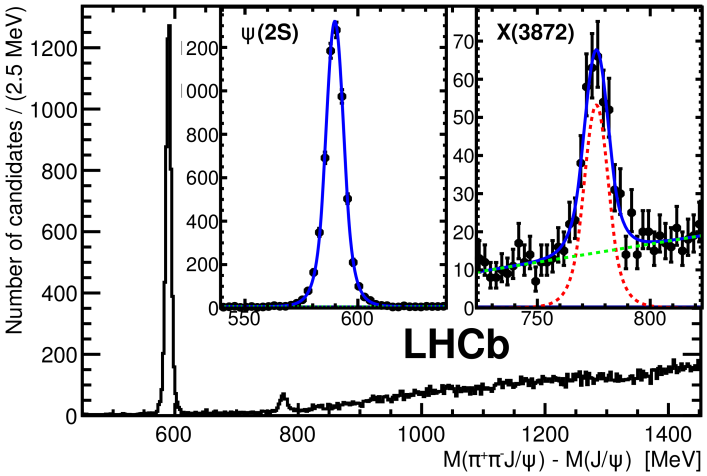

Figure 4.1: Mass difference of $\pi^+\pi^- J/\psi$ to the J/ψ mass for $X(3872)$ candidates reconstructed in the decay to two pions and a J/ψ. The inset plots also show fits for the signal and background components, see Ref. [53].

or a four-quark bound state known as a tetraquark.

4.2.4 Calculation of scattering processes

The probability for interaction is quantified by the cross section, which is a Lorentz-invariant quantity representing the rate of interactions per unit incoming flux per target particle. For a process with two colliding particles in the initial state, and N particles in the final state, the differential cross section is given as

$$d\sigma = \frac{1}{2s} \sum |\mathcal{M}|^2 \, d\Phi_N \tag{4.2.1}$$

where $\mathcal{M}$ is the probability amplitude (or matrix element) and s is the square of the centre of mass energy of the two colliding particles with four-momenta p_1 and p_2,

$$s = (p_1 + p_2)^2 \, .$$

The differential Lorentz invariant phase space $d\Phi_N$ determines the density of available momentum states of the final state particles given their energy and momenta.

One can expand $\mathcal{M}$ for a scattering process as a power series in the coupling constant, a dimensionless number which determines the strength of an interaction. When the coupling constant $\alpha \ll 1$, only the leading terms in α need to be considered when calculating $\mathcal{M}$. This approach is called perturbation theory, because the interaction can be thought of as a perturbation of the free (non-interacting) fields.

Feynman diagrams show the structure of interactions in a two-dimensional representa-

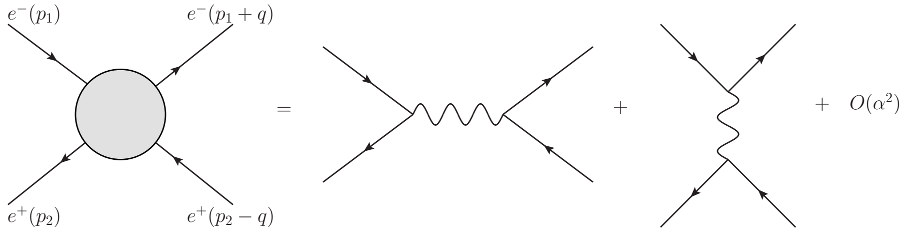

Figure 4.2: Leading order Feynman diagrams for the Bhabha scattering (elastic scattering) process. The first diagram is not a proper Feynman diagram, but it is common to draw a solid area to indicate unspecified dynamics. The two diagrams on the right hand side of the equals sign are the only two processes with exchange of a single photon.

tion of time and and momentum. The time coordinate runs from the left to the right. Each Feynman diagram corresponds to a term in the perturbative expansion of $\mathcal{M}$. For a description of Feynman diagrams see Ref. [48]. The leading order contributions to Bhabha scattering are shown in fig. 4.2, as an example of a Feynman diagram. Bhabha scattering is elastic scattering of an electron and a positron.

4.2.5 Validity of the perturbative expansion

Virtual particles are continually produced and annihilated around a charge, and this produces a screening effect, modifying the effective coupling constant. The effective coupling constant depends on the energy and momentum scale at which the interaction occurs, specifically

$$\mu^2 = |\mathbf{q}^2 - E_q^2| ,$$

where $\mathbf{q}$ is the momentum exchanged in the interaction, and E_q is the energy (μ will sometimes be referred to as the energy scale in the following, for brevity). The distance scale is inversely proportional to the momentum scale. The modification of the coupling constant with the scale μ^2 is known as the running of the coupling constant. The QCD coupling α_s has a strong dependence on μ^2. It is greatest at low energy, and gets smaller at high energy. It is possible to make a perturbative expansion in α_s when the μ^2 scale of the interaction is sufficiently large.

4.2.6 Hadron scattering processes

It is common to factorise the cross section for production of a final state X in a proton-proton interaction into a hard (high-energy) scattering process $ab \to X$, and a phenomenological description of soft (low-energy) hadron physics using Parton Distribution Functions (PDFs)[1] giving the probability to extract partons a and b from the protons.

---

[1] In the following chapters it will be necessary to discuss both the statistical quantity the probability density function, which will be shortened as pdf, and the physical quantity the parton distribution function,

The term parton is used to collectively refer to the particles which interact with the strong force, quarks and gluons, and the PDFs include both sea and valence quarks. The PDF $g_a(x, \mu_F^2)$ gives the probability to find a parton a carrying a fraction x of the proton momentum, and the cross section is given by

$$\sigma = \int_0^1 dx_1 \, g_a(x_1, \mu_F^2) \int_0^1 dx_2 g_b(x_2, \mu_F^2) \hat{\sigma}_{ab\to X} \, . \tag{4.2.2}$$

In this equation, μ_F is the soft-hard factorisation momentum scale, defined such that interactions with momentum scales smaller than μ_F are included in the PDF, while interactions with higher momentum scales than μ_F are included in $\hat{\sigma}_{ab\to X}$. The subprocess cross section $\hat{\sigma}_{ab\to X}$ is the cross section for the hard process, which can be calculated perturbatively. The interesting physical processes usually occur at high energy, so experiments will use the PDFs as input and measure $\hat{\sigma}$. It is also important to perform experiments to constrain the PDFs, as the physics embodied in the PDFs cannot currently be computed analytically or numerically.

The process $ab \to X$ is specified in perturbative QCD with free partons in the initial and final states. The final state partons split into other partons, such that in the end there is a jet of partons, with energy similar to the original parton. This process is referred to as fragmentation. The fragmentation process ends when the partons have low enough relative momentum to form bound states. The quarks finally combine into hadrons, and the gluons fragment into quarks which also form hadrons. These effects are included in the subprocess cross section.

Groups of partons are frequently produced such that they have a large difference in momentum, but their colours are connected. The result is that the colour field between them breaks up and many particles are produced with intermediate momenta. Interactions at hadron colliders thus frequently contain large numbers of hadrons with relatively low energy. The beam protons are generally destroyed in the collision, as they become connected to the hard subprocess by the strong force. There are certain interactions in which the protons remain intact, to be discussed in the following sections.

4.3 The theory of Central Exclusive Production

Central Exclusive Production (CEP) is a class of interactions between two hadrons where both incident hadrons remain intact after the interaction, and an additional simple system is produced. The interacting particles are assumed to be protons in the following theoretical review, because this is the relevant case for the LHC. The momenta of the outgoing protons are similar to the initial state momenta, so the protons continue with almost the same momentum as the beam particles that did not collide. The central system

---

which will be written as PDF.

is typically a single particle or a two-body system, and it is separated from the outgoing protons by large rapidity gaps. The rapidity is a quantity related to the particle velocity, and is defined as

$$y = \frac{1}{2} \ln \left(\frac{E + p_{\parallel}}{E - p_{\parallel}} \right) ,$$

where E is the energy and $p_{\parallel}$ is the momentum component along a defined axis, such as the beam axis. Rapidity gaps are ranges of rapidity where no particles are produced. The rapidity of the central system in CEP is "central" (i.e. close to zero) compared to the protons, but some central systems are produced in the forward region, for example in the LHCb acceptance of $1.9 < \eta < 4.9$ (the pseudorapidity η, defined in (2.2.1), is close to the rapidity y for relativistic particles). The kinematics of CEP are characterised by two hadrons scattering at small angles, and a central system being produced with low p_T. The word "exclusive" in CEP refers to the fact that the full final state is measured and calculated. The outgoing protons can only be measured with special detectors, but their presence can also be inferred from the observation of rapidity gaps, in experiments without such proton taggers.

Properties of the cross sections of hadron interactions can be derived from Regge theory and the physics of the pomeron [54, 55]. Regge theory describes scattering processes in terms of the exchange of mesons. It was seen that the known mesons such as ρ^0 and ω were not sufficient to describe the total cross section at high energy, and it was necessary to introduce the *pomeron*. The pomeron is a neutral object with the quantum numbers of the vacuum, 0^{++}. It is only relevant as an exchange between hadrons, not as a genuine particle. Regge theory and QCD provide complementary descriptions of scattering processes. Regge theory predicts total cross sections and low-energy observables (and some high-energy observables as well), while perturbative QCD can only make predictions at high energy, when $\alpha_s \ll 1$. It is sometimes necessary to invoke Regge theory for CEP calculations when perturbative QCD is not valid.

Lowest-order Feynman diagrams for the most important CEP processes are shown in fig. 4.3. CEP interactions are mediated by systems of two gluons (to leading order in α_s) or by photons. The exchange of two gluons can be thought of as the exchange of a pomeron. The exchanged object cannot carry any electromagnetic or colour charge, because the protons need to remain electrically neutral and colour singlets. As seen in fig. 4.3, different exchanges produce different final states. The fusion of a photon and a pomeron (a) is known as photoproduction, and can produce the mesons J/ψ, $\psi(2S)$, $\Upsilon(1S)$ and others. The exchange of two photons (b) is referred to as the diphoton process, and it produces pairs of muons with a non-peaking mass distribution. The exchange of two pomerons (c) is called Double Pomeron Exchange (DPE), and produces among other particles the χ_c mesons. The CEP processes are generally written as $p + p \rightarrow p + X + p$, where large rapidity gaps are indicated by plus signs and X is the central system. The

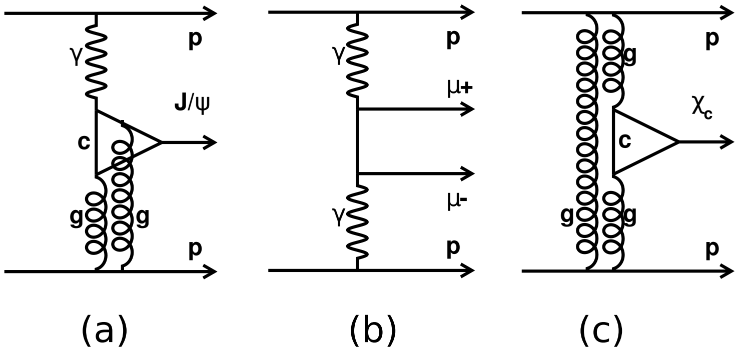

Figure 4.3: Diagrams of the main processes of central exclusive production. Gluons are drawn as helices in Feynman diagrams. The exchange of two gluons in a colour neutral state is equivalent to a pomeron.

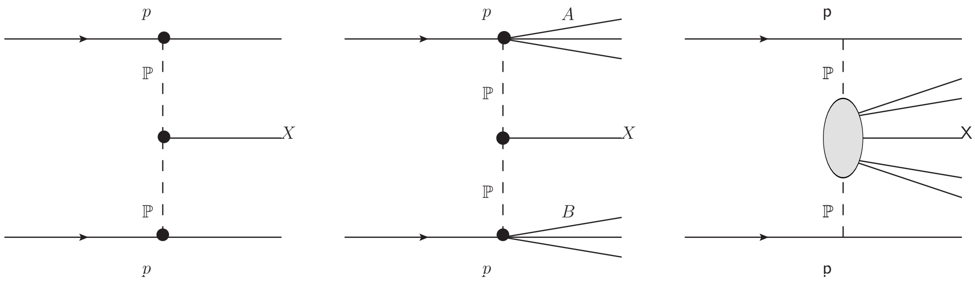

(a) Exclusive (b) Inelastic (c) Inclusive

Figure 4.4: Diagrams for CEP, inelastic production and inclusive production. The pomerons may be replaced by photons.

processes of CEP are described in Sections 4.3.1–4.3.3. The experimental measurement of double pomeron exchange, the process shown in fig. 4.3c, is the main focus of this thesis. The main benefit of measuring CEP processes is that they produce clean final states with constraints on the quantum numbers of the central systems. This will become more clear after a description of the individual processes, and the benefits of CEP measurements are reviewed in Section 4.5.

Two other types of interactions are related to CEP, and may act as backgrounds in CEP analyses, but can also be alternative signal processes. These processes are shown in fig. 4.4. Inelastic production is similar to CEP, but the protons are not required to remain intact. Large rapidity gaps are however still present between the proton remnants and the central system. Inclusive production refers to processes where the central system contains

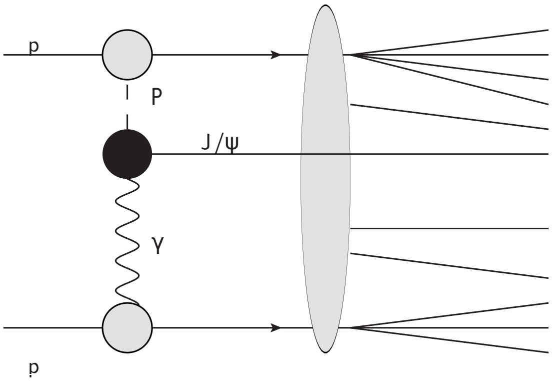

Figure 4.5: Soft rescattering process: the photoproduction interaction occurs as normal through a photon-pomeron exchange, but additional particles are produced by an unrelated interaction between the partons in the protons.

both a specific signal system and any number of additional particles. This process also contains rapidity gaps.

4.3.1 Exclusive photoproduction in hadron interactions

Exclusive photoproduction of vector mesons is shown in fig. 4.3a. It is analogous to a process $e + p \rightarrow e + J/\psi + p$ at the electron-proton collider HERA, with the electron replaced by a proton. Mesons with the quantum numbers of the photon 1^{--} can be produced in this process. In addition to mesons, it is possible to produce the Z^0 boson, but the cross section is low. The cross section for the proton-proton process $p+p \rightarrow p+J/\psi+p$ can be related to that of the electron-proton process by writing the cross section as the product of a photon flux and a photon-proton cross sections. One proton is considered as a source of photons and the other proton is a target. In most experiments it is impossible to determine a priori which proton is the source and which is the target, so the amplitudes for both configurations are added. The quantum-mechanical interference between the configurations when a proton is a source and when it is a target is small for the relevant kinematics and can be ignored [56]. The interaction occurs at a relatively large distance from the proton that emits the photon, because the electromagnetic interaction has a long range.

There may be additional interactions between partons in the two protons, which are not part of the main photoproduction process. An example of such a process is shown in fig. 4.5, where there is a photoproduction process, but an additional unspecified strong interaction causes the protons to disintegrate. The additional interactions are known as rescattering. The probability that no such interaction occurs is quantified by the soft survival probability S^2 [57]. The soft survival probability is also used for the other CEP processes (diphoton exchange and DPE). While S^2 is formally factorised from the hard process, the value of S^2 depends on the overlap of the protons, which again depends on the kinematics of the hard process. It is thus different for photoproduction and for the other processes.

The SuperCHIC MC generator [58], the STARLIGHT MC generator [59] and calculations by other groups use an extrapolation based on HERA data to determine the photon-proton cross section. A power-law function is fitted to the cross section as a function of the photon-proton centre of mass energy W. Motyka and Watt [56] use an impact parameter dependent dipole cross section model and predict a different dependence on W than the power law fit. A model by Schäfer and Szczurek [60] includes a better description of interference effects and the interactions of the other partons which do not participate in the main interaction. The photoproduction cross section is approximately proportional to the gluon density, and the measurement of the cross section as a function of W can be used to constrain the gluon density in protons in a region where there are few experimental data.

Another hypothesised process that can produce the same final state as photoproduction is pomeron-odderon fusion [61]. The photon is replaced with an *odderon*, which is similar to a pomeron, except that it has a negative parity quantum number. The odderon contribution cannot easily be distinguished from photoproduction, but it can be constrained if a cross section consistent with photoproduction alone is observed. Signals consistent with only photoproduction were found by CDF [62] and LHCb [63].

4.3.2 Pair production by diphoton exchange

Non-resonant central exclusive pair production is the process where a particle and an antiparticle are produced with a continuous, non-peaking mass distribution. The main process is the “diphoton” process, mediated through the diagram in fig. 4.3b. The properties of this process are predicted with significantly better precision than the other CEP processes because it is a purely electromagnetic process (unless partons are pair produced), and the use of perturbation theory is justified. The proton form factors enter in the cross section as phenomenological expressions, but these are well known from experiments. The cross section has been computed for dimuon production by several groups [64–66], and has been implemented in the MC generator LPAIR [67]. The muon pair system is produced with low transverse momentum ($p_T(\mu\mu)$) relative to the beam axis, as the momentum transfer from each proton is limited by the elastic form factor, which falls off quickly with momentum. The diphoton interaction may occur at a relatively large distance from the protons, due to the low momentum and the long range of the electromagnetic interaction. This makes the probability for additional strong interactions between spectator partons low, yielding a large survival probability S^2.

It is also possible to produce χ_c mesons by diphoton exchange, by production of a

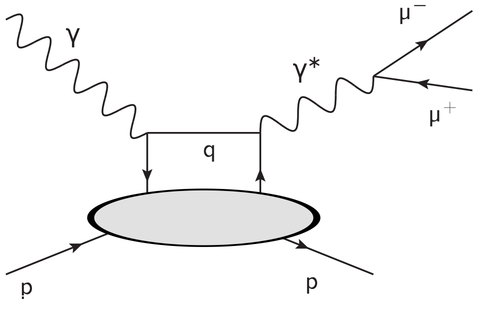

Figure 4.6: The Timelike Compton scattering process. The photon-proton subprocess is shown, the photon flux is given by a proton form factor.

charm quark and an antiquark. The production cross sections for $c\bar{c}$ states and for dimuons are similar because the electromagnetic coupling is similar for leptons and quarks, but the $c\bar{c}$ cross section may be lower due to a smaller phase space. The rate of muon pair production from the decays of $\chi_c \to \gamma J/\psi$, $J/\psi \to \mu^+\mu^-$ is much smaller than in direct dimuon production due to the low branching ratios of the decays.

Another process which produces non-resonant muon pairs is Timelike Compton Scattering (TCS) [68, 69]. It is similar to exclusive photoproduction in that one proton can be considered as a source of photons and the other proton is the target. The reaction, shown in fig. 4.6, is $\gamma p\,(\to \gamma^* p) \to \mu^+\mu^- p$, where an approximately real (zero mass) photon is scattered on a quark and a virtual, time-like photon ($E^2_{\gamma^*} - p^2_{\gamma^*} > 0$) is emitted. The time-like photon fluctuates into a lepton-antilepton pair. It is not known whether the TCS process produces resonances when the dimuon mass is in the region of charmonium ($c\bar{c}$) mesons and down to 2 GeV, but for other masses the spectrum does not have significant resonances. The TCS process has a lower cross section than the diphoton process, and at low $p_T(\mu\mu)$ the diphoton process is expected to dominate.

4.3.3 Double pomeron exchange

Double Pomeron Exchange (DPE) is a pure QCD process, with the Feynman diagram shown in fig. 4.3c. While the central system is shown as a χ_c in the figure, it is also possible to produce other particles and systems with positive parity and charge conjugation $PC = ++$, such as pairs of photons and the Standard Model Higgs boson. The term DPE may refer to processes which are not CEP, in which the protons dissociate or additional particles are produced along with the central system[1]. In the following discussion, DPE will only be used for the CEP process, unless specifically stated otherwise.

The DPE process has been studied in hadron collisions since the earliest experiments. The interest in the process was renewed when it was realised that one could search for

---

[1] The CEP DPE process is sometimes called double diffractive exclusive production, and the similar process where the protons dissociate is called double diffractive inclusive production.

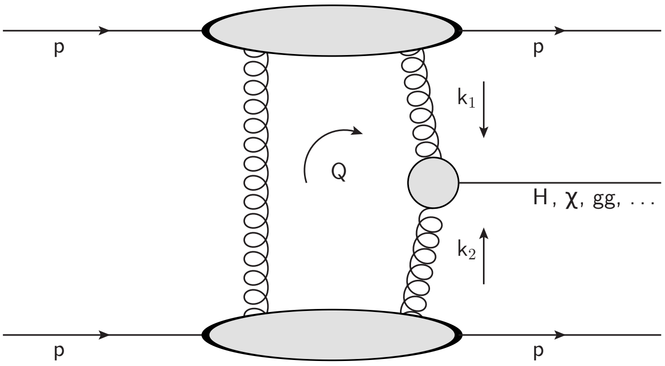

Figure 4.7: Main diagram of the Durham model of CEP. The left gluon screens the colour flow such that no colour is exchanged. The momentum around the loop is Q. The momenta of the gluons which fuse to produce the central system are k_1 and k_2. The large grey blobs represent the unintegrated gluon distribution functions. The small grey circle represents the $gg \to X$ vertex, where X is the central system.

the Higgs boson in the reaction $p + p \to p + H + p$ [70]. A review of the theoretical description of DPE is given in Ref. [57]. The Durham model[1] [70–75] is an important tool for calculation of DPE cross sections and kinematics. Many features of the model are independent of the central system, so the same model can describe central exclusive Higgs production, χ meson production [58, 76], dijet (two back-to-back jets) production, etc., with minor modifications. We will first review the main features of the Durham model and then discuss other models.

The Durham Model of DPE

The DPE process is described as an exchange of two gluons in perturbative QCD, shown in fig. 4.7. Despite the name "DPE" used here, the pomeron does not occur as a distinct term in the mathematical model. It is necessary to exchange two gluons and not just one because there can be no net colour exchanged between the protons. If colour is exchanged over a large range of rapidity, additional particles will necessarily fill the rapidity gap. The gluon on the left in fig. 4.7, which does not couple directly to the central system, is called the "screening" gluon.

The calculation makes use of a factorisation of the probability amplitude into *skewed unintegrated gluon distribution functions*, which describe the proton physics, and the hard subprocess amplitude, giving the amplitude for the fusion of two gluons into the central system. The transition probability amplitude of DPE is given as (following approximately

---

[1] The group responsible for this model is also known as KMR (Khoze, Martin, Ryskin) and later the KRYSTHAL collaboration (Khoze, Ryskin, Stirling, Harland-Lang).

the notation of [58]):

$$T = \pi^2 \int \frac{d^2\mathbf{Q}_\perp V}{\mathbf{Q}_\perp^2(\mathbf{Q}_\perp + \mathbf{p}_{1\perp})^2(\mathbf{Q}_\perp - \mathbf{p}_{2\perp})^2} f_g(x_1, x_1', Q_{1,\perp}^2, \mu^2; t_1) f_g(x_2, x_2', Q_{2,\perp}^2, \mu^2; t_2) \; . \tag{4.3.1}$$

Vectors in three-dimensional space are printed in boldface. There is a loop in the Feynman diagram of fig. 4.7, such that the momenta of the gluons and the intermediate proton lines (blobs in the diagram) are not completely specified by the kinematics of the external particles. The loop momentum transverse to the beam direction is written as $\mathbf{Q}_\perp$. The integral (4.3.1) represents the sum over all values of the loop momentum. The symbols $\mathbf{p}_{i\perp}$ represent the transverse momenta of the outgoing protons.

The fractions of the longitudinal momentum of proton i ($i = 1$ or 2) taken by the screening gluon and the gluon which interacts with the central system are labelled x_i' and x_i. The gluon distribution function f_g gives the probability density for taking a gluon out of the proton with a given momentum fraction x_i and returning a gluon with the same colour with momentum fraction x_i'. The x_i and x_i' are also integrated over to obtain the final amplitude. The other parameters of f_g are the soft-hard factorisation scale μ, the momentum scale Q_i^2 and the four-momentum transfer $t_i = (p_i - p_i')^2$ from each proton (p_i is the initial state momentum of proton i, and p_i' is the final state momentum). When the protons scatter at small angles, the gluon distributions have a factorisation [58]

$$f_g(x_i, x_i', Q_{i,\perp}^2, \mu^2; t_i) = f_g^B(x_i, x_i', Q_{i,\perp}^2, \mu^2) F_N(t_i) \; .$$

The t-dependence is described by a phenomenological form factor

$$F_N(t) = \exp(-b|t|/2) \; .$$

The slope b is extracted from a fit to experimental data. The cross section thus has a Gaussian dependence on transverse momentum, since $|t| \approx |p_T|^2$. This t-dependence enters because the CEP process must take place at small momentum transfers in order for the protons to remain intact. The remaining factor f_g^B can be computed from the ordinary integrated gluon distributions (PDFs) $g(x, Q)$ in the kinematical configuration relevant to CEP.

The subprocess amplitude V is the amplitude for the fusion of two gluons in a colour singlet configuration into the central system, $\hat{\sigma}(gg \to X)$. The coupling to a scalar boson such as the Higgs boson is a constant, independent of the kinematics [73]. The parity and charge conjugation quantum numbers of the central system are given by those of the pomeron, and are both positive, $CP = ++$. Several selection rules apply to the spin structure of the DPE process [73]:

- $J_z = 0$: The component of the spin of the central system along the direction of the proton beam momentum is zero. The rule is a consequence of the kinematics (loop integration). This rule holds in the *forward approximation* (or forward limit), when the protons scatter at zero angle, i.e. when the directions of their momenta are unchanged in the interaction. The selection rule also holds to a good approximation when the outgoing protons are given small transverse momenta.

- The production of gluon jets is favoured over quark jets due to the colour factors and the spin and polarisation of the fusing gluons [73]. Production of a pair of light quarks is forbidden in the $J_z = 0$ state as a result of parity and time reversal symmetry and helicity conservation [73], resulting in a further suppression. The suppression is not as strong for the b quark due to its large mass.

- $J = 1$: The transition of two massless gluons to a spin-1 state is forbidden due to rotation and parity symmetries, a result known as the Landau-Yang theorem [77]. The fusing gluons in the perturbative model are virtual particles and may have non-zero mass $E_g^2 - |\mathbf{p}_g|^2 \neq 0$. The production of χ_{c1} is thus not completely forbidden by the Landau-Yang theorem [78]. Production of $J = 1$ states is also forbidden when the outgoing protons have zero p_T due to symmetry arguments in Regge theory [79].

- Production of χ_{c2} is suppressed [73] because the coupling of two gluons in a $J_z = 0$ state to a bound state of two particles with quantum numbers 2^{++} and spin state ${}^3P_2$ is known to be zero (from experiments with positronium). The suppression is only as valid as the $J_z = 0$ selection rule, and furthermore it relies on the approximation that the relative speed of the produced quarks is non-relativistic ($v \ll 1$).

The soft survival factor S^2 applies to the DPE process, and is smaller than for photoproduction and the "diphoton" process because of the shorter range of pomeron exchange relative to photon exchange. There are also explicit connections between the partons of the pomeron and partons in the other proton, which cause enhanced absorptive corrections [80].

The internal structure of the mesons is taken into account when computing the χ_c subprocess amplitudes [58]. The higher spin states can only be produced at non-zero p_T, so they are suppressed by the proton form factors. This is, however, compensated in experiments by the branching ratio of $\chi_{c0} \to \gamma J/\psi$ which has a small value $1.17\,\%$ compared to those for the χ_{c1} and χ_{c2}, which are $34.4\,\%$ and $19.5\,\%$. The signals from the higher spin states, especially the χ_{c2}, in the $\gamma J/\psi$ channel are thus expected to be observed with a comparable intensity to that of the χ_{c0}.

While the use of QCD perturbation theory is well justified for heavy central systems like the Higgs boson, the χ_c with a mass of $\sim 3.5\,\mathrm{GeV}$ is at the limit of the validity. The gluons typically have to carry an energy of $M_X/2$, where M_X is the mass of the central system, and for a large M_X this ensures that the process occurs at a scale where the strong coupling α_s is small. The "screening" gluon which does not couple directly to the central system (on the left in fig. 4.7) generally has a lower momentum than the gluons which fuse into the central system, making the application of perturbation theory more questionable. The Durham group also calculated the χ_c production cross section using Regge theory [76]. The final results for χ_c production include both a contribution from the Regge calculation in the non-perturbative region, for loop momentum $Q_\perp < 0.85\,\mathrm{GeV}$, and a perturbative contribution. The kinematical distributions of the Regge-based component and the perturbative component are found to be similar, and in the updated paper [58] only the perturbative framework is used.

The Durham group also calculated the cross section for inelastic diffractive production [72, 76]. This is a similar process to CEP, with large rapidity gaps, but where one or both of the protons dissociate into simple multi-particle systems. For heavy central systems, this process has a greater cross section than the CEP process, because the transferred momentum from each proton is not limited by the proton form factor. For lower masses, when there is no pressure to have large values of p_T, the exclusive process is preferred. For the χ_{c0}, the cross section for the inelastic process was originally found to be less than for the exclusive process by approximately a factor 10 [76]. The updated result [58] finds a smaller exclusive χ_{c0} cross section, but the dominance of the exclusive process still holds. There may be greater relative inelastic contributions for the χ_{c1} and χ_{c2}, as the CEP cross sections are suppressed and there is no such suppression in the inelastic process.

The SuperCHIC MC generator implements the Durham model and can produce events for different central systems and collider energies. The χ_c cross sections according to SuperCHIC at the 2011 LHC centre of mass energy $\sqrt{s} = 7\,\mathrm{TeV}$, over the full kinematic range and including the branching ratios of $\chi_c \to \gamma J/\psi$ and $J/\psi \to \mu^+\mu^-$, are $194\,\mathrm{pb}$, $133\,\mathrm{pb}$ and $44\,\mathrm{pb}$ for χ_{c0}, χ_{c1} and χ_{c2} respectively.

Other models of DPE

A summary of the theoretical predictions to be described in this section is given in table 4.2. The double pomeron exchange inclusive (fig. 4.4c) Higgs production cross section was calculated by Bialas and Landshoff [81] using a combination of Regge theory and QCD. In central inclusive production the desired signal particle is produced along with other particles, separated from the protons by large rapidity gaps like in CEP. The protons may remain intact. The main difference with respect to the Durham model is that this model includes a phenomenological non-perturbative description of the gluon, and it does not use the unintegrated gluon distributions [57]. The framework was extended to also describe CEP by Boonekamp et al. [82], though the main focus of Ref. [82] is on

Name	Final state	Comment	Prediction $(14\,\mathrm{TeV})$	Reference
Durham group	H, jj, $\gamma\gamma$, χ, $M\bar{M}$	Perturbative QCD	$\sigma_H = 2.2\,\mathrm{fb}$	[73], [75],
			$\sigma_{\chi_c} = 275\,\mathrm{nb}$	[76], [58], [58]
Bialas, Landshoff	H	Non-perturbative gluons, Inclusive prod.	$\sigma_H(1.96\,\mathrm{TeV}) = 20\,\mathrm{fb}$	[81]
Saclay group (Boonekamp)	H, jj, $\gamma\gamma$, $\ell\ell$	Quasi-exclusive	$\sigma_H = 2\,\mathrm{fb}$	[82]
Bzdak	H, jj, χ_{c0}, χ_{b0}	Non-pert., Sudakov factor	$\sigma_H = 0.015\,\mathrm{fb}$	[83]
			$\sigma_{\chi_{c0}} = 350\,\mathrm{nb}$	[84]
Yuan	χ_{c0}, χ_{b0}	Perturbative, meson wavefunction	$\sigma_{\chi_{c0}} = 735\,\mathrm{nb}$	[85]
Pasechnik et al.	χ_{c0}, χ_{b0}	Various gluon distributions, QMRK	$\sigma_{\chi_{c0}}(1.96\,\mathrm{TeV})$: 55.2–$1127\,\mathrm{pb}$	[86], [87]

Table 4.2: Theoretical predictions cross sections of DPE CEP processes. The cross section for production of Higgs boson with mass $120\,\mathrm{GeV}$ is given as σ_H. The total cross section for χ_{c0} production at a centre of mass energy of $14\,\mathrm{TeV}$ is given as $\sigma_{\chi_{c0}}$. Bialas and Pasechnik do not give the cross section at $14\,\mathrm{TeV}$, so the cross sections at the Tevatron energy are given.

inclusive production. A prediction for "quasi-exlucisve"[1] production of $120\,\mathrm{GeV}$ Higgs bosons is given. This value is compatible with the Durham model, but the agreement is coincidental, and there is a large difference for other M_X [57]. The DPEMC [88] MC generator implements different models of CEP and inelastic diffraction, including the Durham model and the Bialas-Landshoff model (the Durham model is, however, also represented by the SuperCHIC [58] and the ExHuME [89] generators). An alternative extension of the Bialas-Landshoff model to CEP is by Bzdak [83]. A different construction than in Ref. [82] (Sudakov factor) is used to ensure that there are no additional gluons associated with the central system, and a significantly smaller cross section is found (by a factor ~ 100 for Higgs production at the LHC).

The application to χ_c production is considered by the Durham group in Refs. [58, 76]. As mentioned above, it is at the limit of the validity of QCD perturbation theory, and a non-perturbative component is also included in the quoted cross sections. Yuan [85] calculates the χ_c and χ_b CEP cross sections for the Tevatron, in the forward approximation (so, neglecting the χ_{c1} and χ_{c2}). A larger χ_{b0} cross section is found compared to the Durham calculation in Ref. [73], due to a different description of the hard subprocess amplitude, i.e. the coupling of two gluons to the χ_{b0}. The Bialas-Landshoff model is applied to χ meson production by Bzdak [84]. A cross section of $350\,\mathrm{nb}$ is found for χ_{c0} production at $\sqrt{s} = 14\,\mathrm{TeV}$, consistent with the χ_{c0} cross section $275\,\mathrm{nb}$ found by SuperCHIC for the same parameters. Pasechnik, Szczurek and Teryaev performed a calculation of the χ_{c0} CEP cross section [86], using a different approach than the Durham group for the unintegrated gluon distributions f_g, and using Quasi Multi-Regge Kinematics for the subprocess vertex. The cross section is found to have a large dependence on the model parameters and the choice of gluon distributions. There is also a change when the virtual gluons are allowed to have a non-zero mass. A summary and update of the results is given in Ref. [87], which also contains results for the χ_{c1} and χ_{c2} cross sections. The cross sections vary by an order of magnitude when using different unintegrated gluon distributions. The ratio of the χ_{c2} to the χ_{c0} cross section has the lowest value 1.1 using the Durham model, and a maximum value 5.4 for a different unintegrated gluon distribution. A separate computation using pomeron-pomeron fusion was performed, giving a total cross section of the same order of magnitude as the perturbative QCD model, but significantly different kinematical distributions. In the main results, however, the perturbative QCD model is used over the full range of energies, and various continuations of the strong coupling α_s into the non-perturbative (low energy) region are explored [87].

---

[1] Quasi-exclusive production allow for some other particles to be produced in addition to the Higgs boson, but requires that the mass of the system of other particles to be small compared to that of the Higgs boson.

4.4 Experimental review

The study of CEP began when QCD was not well established, and CEP interactions were first described by Regge theory. In this section the results from early hadron colliders will only be briefly summarised. The early hadron colliders had a limited reach in energy, so the focus was on events where the central system had a low mass. While these results are important in their own right, the models described in the previous sections cannot describe this low mass region. The models based on perturbative QCD only apply when there is an energy scale in the process for which α_s is small, and that hard scale is usually provided by the mass of the central system M_X [76, 86]. The results from Tevatron and LHC experiments can be compared to the models, and these results will be described in more detail. A more complete review of the experimental results on CEP (mainly DPE) is given in Ref. [57].

4.4.1 Early hadron experiments

DPE was observed for the first time at the Intersecting Storage Rings (ISR) at CERN, in the reaction $p + p \to p + (\pi^+\pi^-) + p$. The distribution of mass of the central system was studied, and resonance structures were seen, consistent with production of only resonances with parity and charge conjugation $PC = ++$. No definitive evidence of glueballs, i.e. bound states of only gluons, was seen. The understanding of the low-mass region and the existence of glueballs is still not completely known in 2013. The study of CEP was continued by fixed target experiments at the Super Proton Synchrotron (SPS) at CERN, and by the E690 experiment at the Tevatron at Fermilab[1]. The luminosities were increased relative to the ISR. The beam energy may not have been sufficiently large to ensure the dominance of DPE over other strong interactions. The SPS was upgraded to function as a proton-antiproton collider, the Sp$\bar{\text{p}}$S, with energy up to $630\,\text{GeV}$. The physics programme of Sp$\bar{\text{p}}$S was mainly focussed on other areas such as hard QCD, although the experiments UA1 and UA8 performed studies of CEP.

4.4.2 Tevatron

The Tevatron was a proton-antiproton collider at Fermilab, with a maximum centre of mass energy of $1.96\,\text{TeV}$. There were two general purpose detectors at the Tevatron: the Collider Detector at Fermilab (CDF) and the DØ (D Zero) Experiment.

The layout of the CDF detector is shown in fig. 4.8. CDF was well equipped for CEP measurements because it had a large angular coverage, in which it could detect rapidity gaps [91]. The central region is shown in fig. 4.8a. The calorimeters covered pseudorapidities up to 5.1. CDF had scintillator detectors placed along the beam pipe

---

[1] The Fermi National Accelerator Laboratory in Illinois, USA.

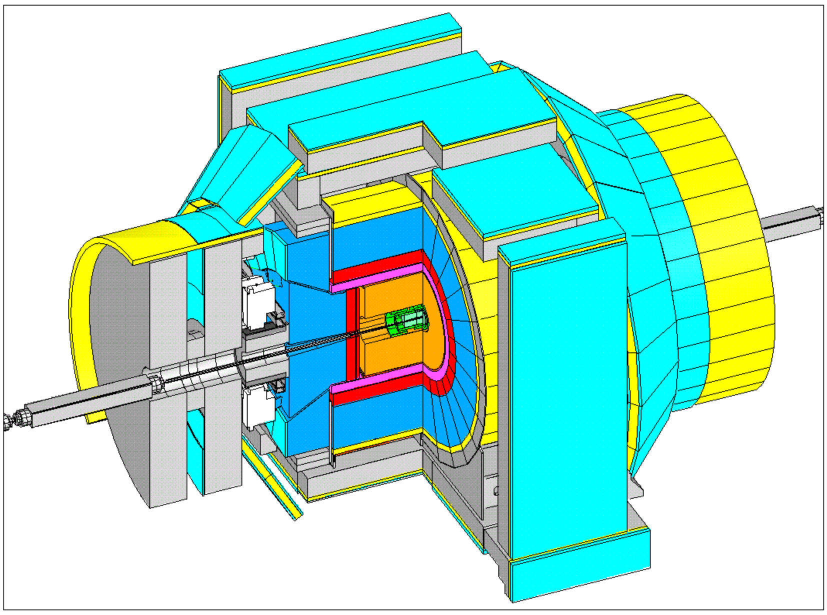

(a) Main central detector

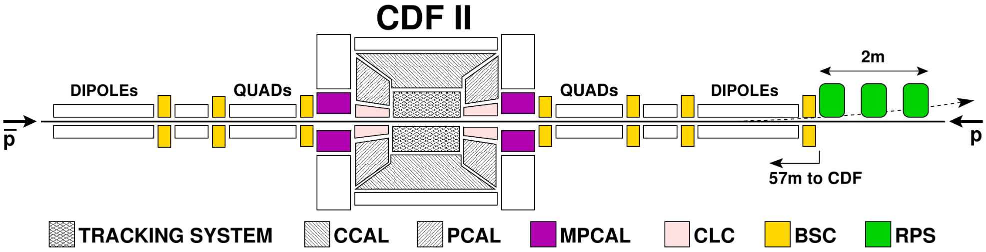

(b) Forward instrumentation

Figure 4.8: The CDF Experiment. Figure (a) shows an isometric drawing of the detector. A section is removed in the drawing to allow visibility into the inner layers. From the centre, the green and orange layers are tracking detectors. The red and blue layers are calorimeters. In (b) a region of about $60\,\mathrm{m}$ on each side of the CDF interaction point is shown, to illustrate the positions of the various forward detectors. The distance scales in the drawing (b) do not represent the true distances. The indicated detectors are CCAL=central calorimeter, PCAL=plug calorimeter, MPCAL=miniplug calorimeter, CLC=Cherenkov luminosity counter, BSC=beam shower counter, RPS=roman pot spectrometer. Figure (a) is from Ref. [90] and Figure (b) is from Ref. [91].

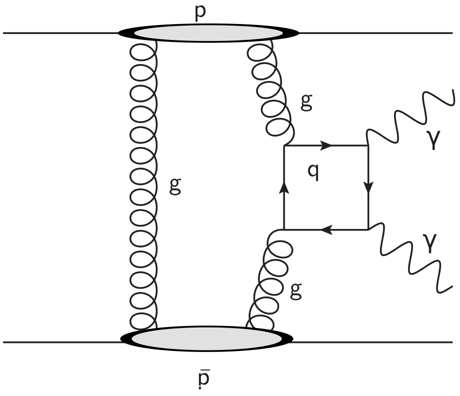

Figure 4.9: CEP of photons through the DPE process.

on both sides of the interaction point, covering $5.4 < |\eta| < 7.4$. These Beam Shower Counters (BSCs) were used to detect particle showers in the forward direction. A proton tagger was operating on one side of the CDF interaction point, measuring the antiprotons emerging from diffractive and CEP interactions. Forward proton taggers work on the principle that the accelerator magnets act as a spectrometer magnet, deflecting protons of different momenta by different amounts. Beam particles which participate in diffractive or CEP interactions lose some momentum, and some of them can thus be detected by the forward proton taggers. The detectors were placed $56\,\mathrm{m}$ from the interaction point, after a series of dipole and quadrupole magnets. The sensors were housed in a secondary vacuum and they were able to move towards and away from the beam.

During Run II (2001–2011), CDF made several measurement of CEP processes [62, 91–95], and also diffraction. An overview of results by CDF and DØ is given in table 4.3. The first CEP measurements were exclusive diphotons [93] and dielectrons [92]. The triggers and analysis methods were very similar for the diphotons and the dielectrons. Photon pairs are produced in DPE, shown in fig. 4.9, in a process very similar to Higgs boson production. Electron-positron pairs are produced in the "diphoton" process discussed in Section 4.3.2 (not to be confused with the diphoton final state). A cut on the transverse energy in the calorimeter of the photon and electron candidates, $E_T > 5\,\mathrm{GeV}$, excluded practically all of the charmonium resonance signal for the dielectrons. The forward proton taggers were not used because they do not have acceptance for the protons when the central system has low mass, and the protons have low longitudinal momentum loss. The BSCs and all the calorimeters were used to veto on particles in the rapidity gaps, giving a combined sensitivity in the range of $|\eta| < 7.4$. The background from events produced in CEP-like processes but with with undetected proton dissociation, i.e. inelastic diffractive production, was estimated from MC.

The dielectron analysis found a cross section consistent with the theoretical predictions. The diphoton analysis found three diphoton candidates, but one of them could have

Expt.	Process	Kinematic region	Veto region	Cross section	Ref.
CDF	$p+\bar{p} \rightarrow p+\gamma\gamma+\bar{p}$	$\lvert\eta\rvert < 1.0$, $E_T > 5\,\mathrm{GeV}$	$\lvert\eta\rvert < 7.4$	$< 410\,\mathrm{fb}$	[93]
CDF	$p+\bar{p} \rightarrow p+\gamma\gamma+\bar{p}$	$\lvert\eta\rvert < 1.0$, $E_T > 2.5\,\mathrm{GeV}$	$\lvert\eta\rvert < 7.4$	$2.48^{+0.40}_{-0.35}\,{}^{+0.40}_{-0.51}\,\mathrm{pb}$	[95]
CDF	$p+\bar{p} \rightarrow p+e^+e^-+\bar{p}$	$\lvert\eta\rvert < 2.0$, $E_T > 5\,\mathrm{GeV}$	$\lvert\eta\rvert < 7.4$	$1.6^{+0.5}_{-0.3} \pm 0.3\,\mathrm{pb}$	[92]
CDF	$p+\bar{p} \rightarrow p+jj+\bar{p}$ (dijet)	$\lvert\eta\rvert < 2.5$, $E_T > 10\,\mathrm{GeV}$	$3.6 < \lvert\eta\rvert < 5.9$, tag $\bar{p}$	$1.10 \pm 0.04^{+1.29}_{-0.54}\,\mathrm{nb}$	[91]
CDF	$p+\bar{p} \rightarrow p+Z+\bar{p}$	multiple	$\lvert\eta\rvert < 7.4$	$< 0.96\,\mathrm{pb}$	[94]
CDF	$p+\bar{p} \rightarrow p+\mu^+\mu^-+\bar{p}$	$\lvert\eta\rvert < 0.6$, $3.0 < M(\mu\mu) < 4.0\,\mathrm{GeV}$	$\lvert\eta\rvert < 7.4$	$2.7 \pm 0.3 \pm 0.4\,\mathrm{pb}$	[62]
CDF	$p+\bar{p} \rightarrow p+J/\psi+\bar{p}$	$\lvert\eta\rvert < 0.6$	$\lvert\eta\rvert < 7.4$	$\frac{d\sigma}{dy}\rvert_{y=0} = 3.92 \pm 0.25 \pm 0.52\,\mathrm{nb}$	[62]
CDF	$p+\bar{p} \rightarrow p+\psi(2S)+\bar{p}$	$\lvert\eta\rvert < 0.6$	$\lvert\eta\rvert < 7.4$	$\frac{d\sigma}{dy}\rvert_{y=0} = 0.53 \pm 0.09 \pm 0.10\,\mathrm{nb}$	[62]
CDF	$p+\bar{p} \rightarrow p+\chi_c+\bar{p}$	$\lvert\eta\rvert < 0.6$	$\lvert\eta\rvert < 7.4$	$\sum \mathcal{B}_i\sigma_i = 8.0 \pm 0.9$ $\pm 0.9\,\mathrm{pb} \quad \mathrm{i} = 0, 1, 2$	[62]
DØ	$p+\bar{p} \rightarrow p+jj+\bar{p}$ (dijet)	$\lvert y\rvert < 0.8$, $M(jj) > 100\,\mathrm{GeV}$	$2.0 < \lvert\eta\rvert < 4.2$	4.1σ evidence	[96]

Table 4.3: Overview of CEP results from Tevatron. The centre of mass energy is $\sqrt{s} = 1.96\,\mathrm{TeV}$. When two errors are given on the cross sections the first is the statistical error and the second is the systematic error. The limits are given at $95\,\%$ confidence level. The kinematic cuts apply to all final state particles listed in the Process column individually except for the protons, unless the system is specified in parentheses.

been from $\pi^0\pi^0$. The diphoton analysis was later followed up with a greater dataset of $1.11\,\mathrm{fb}^{-1}$ [95]. In this analysis the dielectron channel was used as a control channel, to ensure that the analysis was correct. A diphoton cross section of $2.48^{+0.40}_{-0.35}(\text{stat})^{+0.40}_{-0.51}(\text{syst})\,\mathrm{pb}$ was found. The theoretical models have large uncertainties, but the order of magnitude of the CDF result is consistent with the the predictions.

CDF also measured the exclusive dijet cross section. The dijet analysis has the advantage that the cross section is greater than for the diphoton final state, but it is more difficult to impose exclusivity requirements because the final state contains many hadrons in the jets. All the forward detectors up to the first BSC and a region of the central calorimeter were used as vetoes. A tagged antiproton was also required. The CDF analysis measured two different "double pomeron exchange" processes. The first process is a two-step process where a pomeron is first emitted from each proton, and then the pomerons collide. The pomerons are described as hadrons, with parton density functions. This is central *inclusive* production, because there may be pomeron remnants in the final state in addition to the jets from the gluon-gluon collision. The second process is just the ordinary DPE CEP process described above, which is better described as a two-gluon exchange. The ratio of the mass of the dijet system to the total mass of all particles in the detector was used to separate the processes on a statistical basis. CDF found a distribution of this ratio in agreement with a combination of the inclusive DPE process and the MC ExHuME [89], which implements the Durham model. The data were not well described by only inclusive DPE.

CDF set a limit on exclusive Z boson photoproduction [94], but the limit was about 3×10^3 times the Standard Model cross section. The result rules out some theories of physics beyond the Standard Model which predict a large cross section.

CDF also searched for CEP in dimuon final states in the mass range around $M(\mu\mu) \sim 3\,\mathrm{GeV}$ [62]. All the three main CEP processes are represented: the electromagnetic "diphoton" process produces a continuum of muon pairs, the photoproduction process produces J/ψ and $\psi(2S)$, and the DPE process produces χ_c mesons. The analysis was similar to those described above. The signals in the BSCs and calorimeters were required to be consistent with only noise except for the signal particles, giving a rapidity gap veto coverage up to $|\eta| = 7.4$. The J/ψ and $\psi(2S)$ were reconstructed in the decay to $\mu^+\mu^-$, and the χ_c was reconstructed in the decay $\gamma J/\psi$. The calorimeter resolution for the photon measurement was not good enough to resolve the χ_c states in the mass spectrum. In Ref. [62] the χ_c signal was assumed to be due to the lowest spin state, χ_{c0}, but the possibility of contributions from the other χ_c states was explicitly mentioned, and it has later been remarked [58,87] that there could be contributions from χ_{c1} and χ_{c2} of similar magnitude. The continuum and the J/ψ and $\psi(2S)$ signal intensities were extracted using a fit to the dimuon mass spectrum, shown in fig. 4.10. The continuum dimuon cross section was found to be in good agreement with QED predictions (the TCS process was not

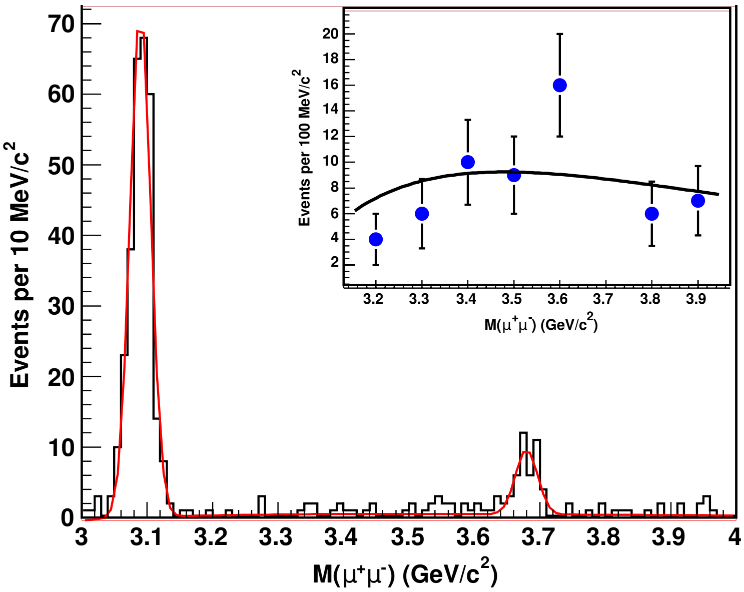

Figure 4.10: Dimuon mass spectrum for CEP candidate events at CDF. Figure from Ref. [62]. The histogram shows the data and the red curve shows the fit. The inset histogram shows the data with the bins with the J/ψ and $\psi(2S)$ resonances removed. The black curve is the distribution of the "diphoton" process, with the shape fixed from a theoretical calculation, but with the normalisation left free.

considered in the paper, but the data are well described by the "diphoton" process alone). The J/ψ and $\psi(2S)$ cross sections were also consistent with theoretical predictions and results from HERA (described in the following section). CDF could thus place a limit on the cross section of the odderon process, because a large odderon contribution would have increased the measured cross section. The odderon is not ruled out, as the predicted odderon cross section is less than the one for photoproduction, and it is consistent with the CDF limit. The χ_c rate was consistent with predictions by e.g. the Durham group.

The DØ Experiment installed Roman pot detectors on both sides of the interaction point during Run II, but no CEP results have been published using those detectors. DØ searched for CEP dijets using the central detectors [96]. The calorimeters were used to measure the jets and to require rapidity gaps, covering the ranges of pseudorapidity $|\eta| < 1.1$ and $1.6 < |\eta| < 4.2$. The CEP dijet selection required that there was no activity in the range $2.0 < |\eta| < 4.2$. CEP dijets were distinguished from the background using a discriminant variable based on the relative amount of energy in the central and the forward regions of the calorimeter. It was not possible to measure the fraction of events with proton dissociation. DØ observed CEP dijets (DPE) with mass greater than $100\,\mathrm{GeV}$ [96] with a 4.1σ significance.

4.4.3 HERA

HERA was an electron-proton collider, located at the DESY facility in Hamburg, which collided protons with energy $920\,\mathrm{GeV}$ with electrons or positrons with energy $27.5\,\mathrm{GeV}$. There were two large detectors, H1 and ZEUS, which both had an excellent instrumentation for measuring diffraction:

- Detectors for measuring the outgoing proton in diffraction and CEP.
- Special detectors which measured the outgoing electron.
- Calorimeters with high forward coverage, up to at least $\eta = 5$ for H1 and ZEUS (after installation of a Forward Plug Calorimeter).
- Scintillators in the forward direction, similar to the BSCs of CDF.

A review of the HERA results is given in Ref. [97]. As the collisions at HERA are electron-proton collisions, the DPE process was not available – electrons cannot emit pomerons.

The main CEP measurements at HERA have been in the photoproduction process. It is practically identical to the proton-proton process in fig. 4.3a, except that the proton which emits the photon is replaced by an electron. The photon and the pomeron interact to produce a central system, and the central system and the proton are separated by a rapidity gap. Both H1 and ZEUS measured exclusive photoproduction of vector mesons, including the J/ψ [98]. The exclusive ρ^0, ω, ϕ and J/ψ cross sections for various photon-proton centre of mass energies are shown in fig. 4.11, for both HERA data and data from the low energy fixed target experiment NMC. The meson cross sections at $W > 10\,\mathrm{GeV}$ are consistent with a power-law dependence on W, but the power is not universal for all mesons. H1 and ZEUS also observed exclusive photoproduced dijets [97]. The data and QCD predictions are consistent within a factor 1.5.

4.4.4 LHC

The LHC was already described in Chapter 2. Results on diffraction physics have been published by ALICE [99] and ATLAS [100]. The TOTEM experiment has released results on the total and elastic proton-proton cross sections. CMS and LHCb have released results on both diffraction and CEP, and will be discussed in more detail. The current CEP analyses at LHC have in common that events are selected based on the absence of particles in the detector apart from the central system, and the final state protons are not tagged. The CEP results from LHC are summarised in table 4.4.

Expt.	Process	Kinematic region	Veto region	Cross section	Ref.
CMS	$p+p \rightarrow p+\mu^+\mu^-+p$	$M(\mu\mu) > 11.5\,\mathrm{GeV}$, $p_T > 4\,\mathrm{GeV}$, $\lvert\eta\rvert < 2.1$	$\lvert\eta\rvert < 2.5$	$3.38^{+0.14}_{-0.13} \pm 0.16 \pm 0.14\,\mathrm{pb}$	[101]
CMS	$p+p \rightarrow p^{(*)}+\gamma\gamma+p^{(*)}$	$E_T > 5.5\,\mathrm{GeV}$, $\lvert\eta\rvert < 2.5$	$\lvert\eta\rvert < 5.2$	$< 1.18\,\mathrm{pb}$	[102]
CMS	$p+p \rightarrow p^{(*)}+e^+e^-+p^{(*)}$	$E_T > 5.5\,\mathrm{GeV}$, $\lvert\eta\rvert < 2.5$	$\lvert\eta\rvert < 5.2$	consistent w. predictions	[102]
CMS	$p+p \rightarrow p^{(*)}+W^+W^-+p^{(*)}$	$p_T(\mu, e) > 20\,\mathrm{GeV}$, $\lvert\eta(\mu, e)\rvert < 2.4$	$\lvert\eta\rvert < 2.4$		
	SM dominated	$p_T(\mu e) < 100\,\mathrm{GeV}$		$\sigma\mathcal{B} < 8.4\,\mathrm{fb}$	[103]
	anomalous dominated	$p_T(\mu e) > 100\,\mathrm{GeV}$		$\sigma\mathcal{B} < 1.9\,\mathrm{fb}$	[103]
LHCb	$p+p \rightarrow p+\mu^+\mu^-+p$	$2.0 < \eta < 4.5$, $M(\mu\mu) > 2.5\,\mathrm{GeV}$, $M(\mu\mu)$ not at J/ψ, $\psi(2S)$ masses	$-4 < \eta < -1.5$, $1.5 < \eta < 5$	$67 \pm 10 \pm 7 \pm 15\,\mathrm{pb}$	[104]
LHCb	$p+p \rightarrow p+J/\psi+p$	$2.0 < \eta < 4.5$	(as above)	$307 \pm 21 \pm 36\,\mathrm{pb}$	[63]
LHCb	$p+p \rightarrow p+\psi(2S)+p$	$2.0 < \eta < 4.5$	(as above)	$7.8 \pm 1.3 \pm 1.0\,\mathrm{pb}$	[63]
LHCb	$p+p \rightarrow p+\chi_{c0}+p$	$2.0 < \eta < 4.5$	(as above)	$9.3 \pm 2.2 \pm 3.5 \pm 1.8\,\mathrm{pb}$	[104]
LHCb	$p+p \rightarrow p+\chi_{c1}+p$	$2.0 < \eta < 4.5$	(as above)	$16.4 \pm 5.4 \pm 5.8 \pm 3.2\,\mathrm{pb}$	[104]
LHCb	$p+p \rightarrow p+\chi_{c2}+p$	$2.0 < \eta < 4.5$	(as above)	$28.0 \pm 5.4 \pm 9.7 \pm 5.4\,\mathrm{pb}$	[104]

Table 4.4: Results from LHC experiments based on proton collisions at $\sqrt{s} = 7\,\mathrm{TeV}$. The first uncertainty on the cross sections is the statistical, the second is the systematic and the third is the luminosity uncertainty.

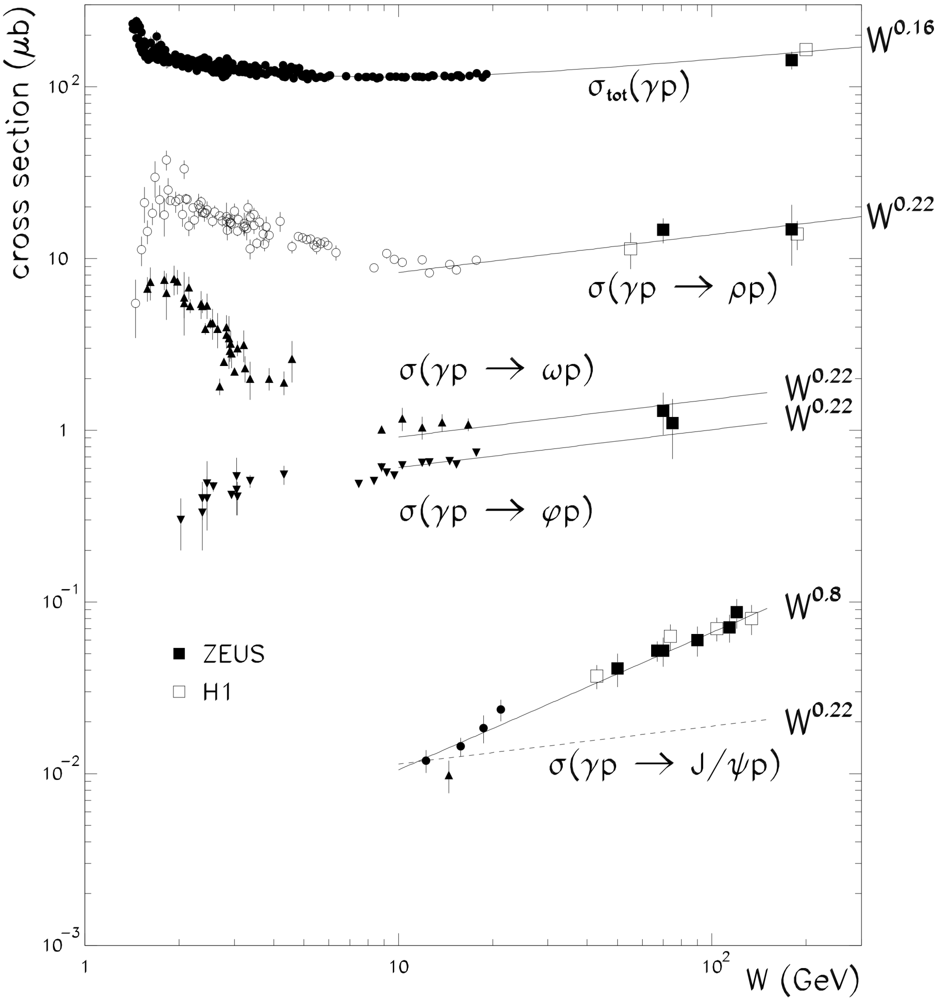

Figure 4.11: Cross sections for exclusive production of vector mesons at fixed target experiments (small points) and at HERA (large points). Lines are drawn to indicate power law fits to the W dependence. Figure from Ref. [98].

CMS

CMS is a general purpose detector [105], covering the central region with tracking detectors, calorimetry and muon detectors. The tracking system covers pseudorapidities up to 2.5. The central calorimeters cover $|\eta| < 3.0$ and a forward calorimeter covers $2.9 < |\eta| < 5.2$. CMS also has specialised detectors in the very forward region; the calorimeter CASTOR covering $-6.6 < \eta < -5.2$ (on one side of the interaction point). The Zero Degree Calorimeter (ZDC), located $140\,\mathrm{m}$ from the interaction point, covers a pseudorapidity $|\eta| \geq 8.3$ on both sides of the interaction point for neutral particles. At this distance, the proton beams have been deflected by the LHC magnets, and the ZDC can thus be positioned in the exact forward direction. The TOTEM detectors can also measure particles from CMS at very forward rapidities. Diffraction was measured using the calorimeters, including CASTOR [106]. The analysis was based on early data from

2010, when the luminosity and pile-up were low. Analyses at CMS which require higher integrated luminosity have to be done in the presence of pile-up.

CMS has released CEP results on final states with dimuons [101], diphotons [102] and W^+W^- pairs [103]. In the dimuon analysis, the "diphoton" process was searched for. Dimuon masses greater than $11.5\,\mathrm{GeV}$ were selected, thus excluding the Υ mesons which are produced by photoproduction. The cross section for photoproduction of Z bosons is negligible compared to the signal process (and no Z candidates were seen). Exclusive events were selected by searching for isolated vertices with only two muons and no other tracks. Proton-proton interaction vertices normally have many more tracks than two, so the signal purity was greatly improved by this selection. The rapidity gap acceptance was defined by the acceptance of the tracking system $|\eta| < 2.5$. The transverse momenta of the muon pairs were analysed to separate the truly exclusive events from processes in which the protons did not remain intact. The components of the p_T distributions were determined using MC simulations. The ratio of the measured cross section to the theoretical prediction was $0.83^{+0.14}_{-0.13}$(stat.) $\pm\, 0.04$(syst.) $\pm\, 0.03$(lumi.), and the measurement is consistent with the theoretical prediction.

CEP diphotons from DPE were searched for in a data sample of $36\,\mathrm{pb}^{-1}$ with low instantaneous luminosity [102]. The analysis also included dielectrons (e^+e^-), as they produce a similar signal in the detector. Exclusive events were selected using global event cuts, requiring that there was no signal above the noise threshold in the tracking detectors and calorimeters, giving an acceptance of $|\eta| < 5.2$ for rapidity gaps. Events with pile-up were rejected. "Semi-exclusive" events with proton dissociation, but no detectable particles in the CMS acceptance, were implicitly included, as it was not possible to distinguish this semi-exclusive and the truly exclusive signal. The e^+e^- result was in good agreement with theoretical predictions. No diphoton candidates passed the selection cuts, so CMS set an upper limit $\sigma(\gamma\gamma) < 1.18\,\mathrm{pb}$ for the kinematical region defined by the selection cuts, consistent with theoretical predictions.

A search for $p + p \rightarrow p + W^+W^- + p$ was done in the presence of pile-up, using the 2011 dataset [103]. The W boson pairs are produced in the diphoton process. The cross section is low in the Standard Model, but it may be enhanced by anomalous gauge boson couplings. The W bosons were reconstructed in the channel where one W boson decays to $\mu^+\nu_\mu$ and the other decays to $e^-\bar{\nu}_e$, or the respective charge conjugates. Most backgrounds were estimated from simulation, but checked extensively using controls samples such as dimuon candidates and candidates with one to six extra tracks (in addition to the two lepton tracks). The probability for proton dissociation was determined from analysis of the $p_T(\mu\mu)$-distribution of dimuon candidates. A histogram of the transverse momentum of the two-lepton system, $p_T(e\mu)$ is shown in fig. 4.12. CMS found a cross section in agreement with the Standard Model for $p_T(e^\pm\mu^\mp) < 100\,\mathrm{GeV}$. The anomalous gauge boson signal is expected to have a tail at transverse momentum higher than $100\,\mathrm{GeV}$.

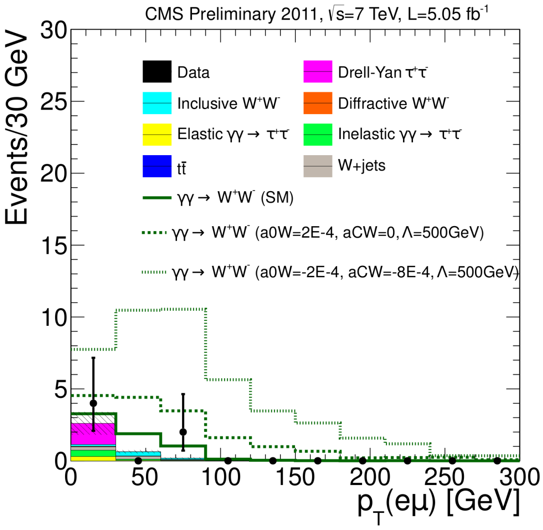

Figure 4.12: Transverse momentum spectrum for the two-lepton system from the CMS measurement of CEP of W^+W^-. The filled histograms show the background, and the open histograms show the signal expectation for the Standard Model and for anomalous gauge couplings. Figure from Ref. [103].

There were no events with $p_T(e^{\pm}\mu^{\mp}) > 100\,\mathrm{GeV}$, and this allowed CMS to place a limit on anomalous gauge couplings.

LHCb

LHCb was described in detail in Chapter 2. It is fully instrumented in the region of pseudorapidity 1.9–4.9 (positive pseudorapidity in the direction of the spectrometer). The VELO has a somewhat greater acceptance, covering a backward acceptance of about $-3.5 < \eta < -1.5$ and a forward acceptance of about $1.5 < \eta < 5.0$.

An analysis of CEP dimuons was performed based on the 2010 dataset [63, 104]. The preliminary result [104] covered non-resonant dimuons (diphoton), photoproduction of J/ψ and $\psi(2S)$ and DPE production of χ_c. The paper [63] focussed on the photoproduction processes, and represents a refinement of the analysis. A subsequent analysis of χ meson production in the 2011 dataset will be the main objective of Chapters 6–8. LHCb has the advantage that it operates at lower instantaneous luminosity than CMS and ATLAS, allowing the use of global event cuts to select exclusive events on all data while keeping about $20\,\%$ of the total luminosity.

The LHCb analyses used a cut requiring no additional tracks or photon candidates apart from the signal in the event. There still remained a background from proton dissociation with undetected particles escaping down the beam pipes. For the diphoton process, the background was estimated from MC. For photoproduction and DPE there does not exist a reliable simulation for the background, so it was estimated by extrapolation from

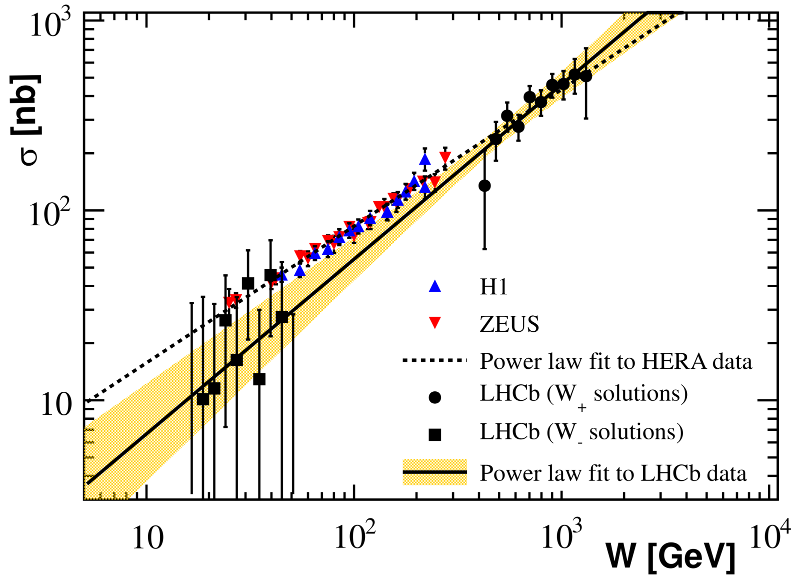

Figure 4.13: Differential cross section for J/ψ production at LHCb as a function of the photon-proton invariant mass W. The black points show the LHCb data. The HERA data are shown in blue and red triangles. Power-law fits to the LHCb data and HERA data are shown.

events with small numbers of additional tracks. The procedure will be discussed in more detail for the χ_c in Section 6.7. The different processes were distinguished by the mass distributions. The χ_c was distinguished from J/ψ by the presence of a photon.

The non-resonant dimuon cross section was in agreement with the prediction for the "diphoton" process. A cut of $p_T(\mu\mu) < 100\,\mathrm{MeV}$ was used to improve the signal purity. The J/ψ and $\psi(2S)$ signals were consistent with predictions from photoproduction, so there was no sign of the odderon. The J/ψ differential cross section was analysed as a function of meson rapidity. The energy in the rest frame of the system of the photon and the proton which emits the pomeron, W, is given by [63]

$$(W_\pm)^2 = m_V \sqrt{s}\,\exp(\pm|y|)\ .$$

There are two solutions for each event because it is not known which proton emitted the photon and which emitted the pomeron. The cross section was reported as a function of W, shown in fig. 4.13. The data are consistent with a power-law behaviour (i.e. $\sigma \propto W^\alpha$ for some α), and are also consistent with a power-law fit to the HERA data. As seen in the figure, the LHCb data probe a previously unexplored region of W. The χ_c data are consistent with the predictions of SuperCHIC, but an excess of χ_{c2} is seen. This excess may be caused by proton dissociation events.

4.5 Prospects for CEP measurements

The benefits of studying CEP are:

- The production process imposes particular constraints on the quantum numbers of the central system.
- The DPE process provides an environment which is dominated by gluons, and where the production of quark jets is suppressed. This is different from normal interactions where there is usually no way to distinguish light quarks and gluons.
- If the outgoing protons are measured, the kinematics are over-constrained and it is possible to measure the mass of the central system with great precision.

CEP can thus be used to learn more about unknown particles, even with some undetectable particles in the final state. It can help with searches for physics beyond the Standard Model and measurements of the dynamics of the strong force, due to its quantum number analysing nature. Now that the Higgs boson has been discovered, it may be useful to study CEP of the Higgs boson in order to determine its properties. In supersymmetric models there are multiple Higgs bosons, which may have similar masses, and CEP can help distinguish these particles. As exemplified by the CMS search for anomalous gauge couplings of photons and W bosons, it is possible to do searches for new physics with CEP.

The production of heavy quarkonia in hadron collisions is not completely understood, because it involves physics at low energy scales where the use of perturbative QCD cannot be justified. A viable model is Non-Relativistic QCD (NRQCD), which is an effective field theory based on QCD (chapter 5 in Ref. [107]). NRQCD provides a factorisation of the high-energy physics of the production of a quark-antiquark pair and the low-energy physics of the quark-antiquark bound state. It is appropriate to use a non-relativistic description because the relative speed of the heavy quarks is small. The CEP processes can provide initial states of definite colour (singlet) and polarisation, which may provide additional constraints on quarkonium production. Data from electron-proton collisions at HERA are already used in this context.

The theory of DPE is consistent, but there are a few parameters such as S^2 which cannot be derived from first principles, and more experimental data are required. There are multiple factors of S^2, which are not independent of the hard subprocess. The predictive power will improve as more data are available, but significant theoretical work has been done [80, 108] and the experimental task will be to check these predictions.

4.5.1 CEP studies at LHCb

LHCb has some advantages over other experiments, particularly for the study of low-mass CEP.

- LHCb can trigger on and reconstruct muons with very low p_T. The complete trigger and reconstruction chain can accept muons with p_T as low as 400 MeV with high efficiency.
- Because LHCb is designed to run at lower luminosity than ATLAS and CMS, the typical number of collisions per event is 1 for LHCb, while it is greater than 20 for ATLAS and CMS. LHCb thus has a much greater fraction of its data with single interactions, allowing the use of global event cuts.
- While LHCb does not have complete angular acceptance, it does have very efficient reconstruction of tracks and neutral particles in the forward region. This results in a good veto power in a region where it is likely that extra particles are produced in diffraction.
- While dimuons have been the final state of choice so far, there is great interest in studying final states with mesons [109]. LHCb has efficient triggers for hadrons of low transverse momentum.

The "hole" in the acceptance at mid pseudorapidity, $|\eta| \lesssim 1.5$, is not particularly troubling. The only backgrounds which could fake a CEP signal by having additional particles in this region are triple pomeron exchanges and central inclusive production. Central inclusive production can be written as

$$p + p \to p + XY + p \ ,$$

where X is the desired signal particle and Y is any number of additional particles. The event would be selected if the desired signal particle was in the acceptance and all other particles were in the "hole" in the acceptance. This is unlikely to happen, because the *veto* region used, which is the full acceptance of the VELO, is greater than the region in which the signal particles are accepted, which starts at $\eta = 2.0$. It would thus require a small rapidity gap within the central system XY, between the signal particle and the other central particles[1], which can occur by chance, but is quite unlikely. Central inclusive production for the Higgs boson was discussed in Ref. [110], and it does have a larger cross section than CEP. The Durham group prediction is that the inclusive and exclusive Higgs cross sections are respectively $50\,\text{fb}$ and $3\,\text{fb}$. These results cannot however be

---

[1] The same argument applies to the rejection of the single diffraction background $p + p \to p + XY$, where XY is produced at large pseudorapidity. The signal is only searched for in $\eta < 4.5$, while the veto acceptance extends up to $\eta \approx 5$.

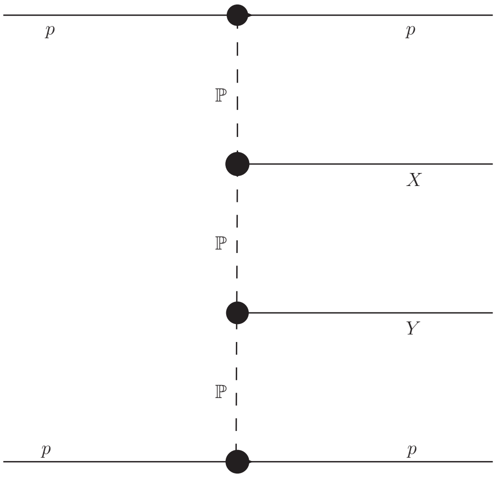

Figure 4.14: Diagram of the triple pomeron exchange process. The protons may dissociate or remain intact. Two isolated systems X and Y are produced, separated by large rapidity gaps from the protons and from each other. The systems X and Y may be one or more particles. As a background to CEP, one of X or Y would be the signal decay and the other would be any other system.

directly applied to central systems of lower mass. Triple pomeron exchange, shown in fig. 4.14, gives rise to three large rapidity gaps, and can fake a CEP signal if one of the central systems is outside the detector acceptance. The contribution from triple pomeron exchange is however expected to be small compared to double pomeron exchange [111].

LHCb can detect central systems in the forward region, approximately $1.9 < \eta < 4.9$. This favours the detection of light central systems, because their pseudorapidity distributions generally extend to larger values. The simulated pseudorapidity distributions generated by SuperCHIC are shown for J/ψ[2], χ_{c1} and $\Upsilon(1S)$ in fig. 4.15. There is a significant signal in the LHCb acceptance $1.9 < \eta < 4.9$ for all processes. The pseudorapidities of the decay products are however generally more central than those of the mother particles.

In Ref. [72] the DPE “luminosity” is derived as a function of various kinematic variables. This luminosity can be multiplied by the subprocess cross sections to give the DPE cross section. The (CEP) DPE luminosity as a function of the rapidity of the central system goes to zero at rapidities greater than 3.9 for a central system with mass $120\,\mathrm{GeV}$, and greater than 2.4 for central systems with mass $500\,\mathrm{GeV}$. Detection of high-mass central systems at LHCb may be possible, but only a small fraction of the total cross section is visible.

---

[2] For the photoproduction process, the signal from the more massive $\Upsilon(1S)$ is further in the forward region than the J/ψ, despite the general observation that heavier systems are produced more centrally.

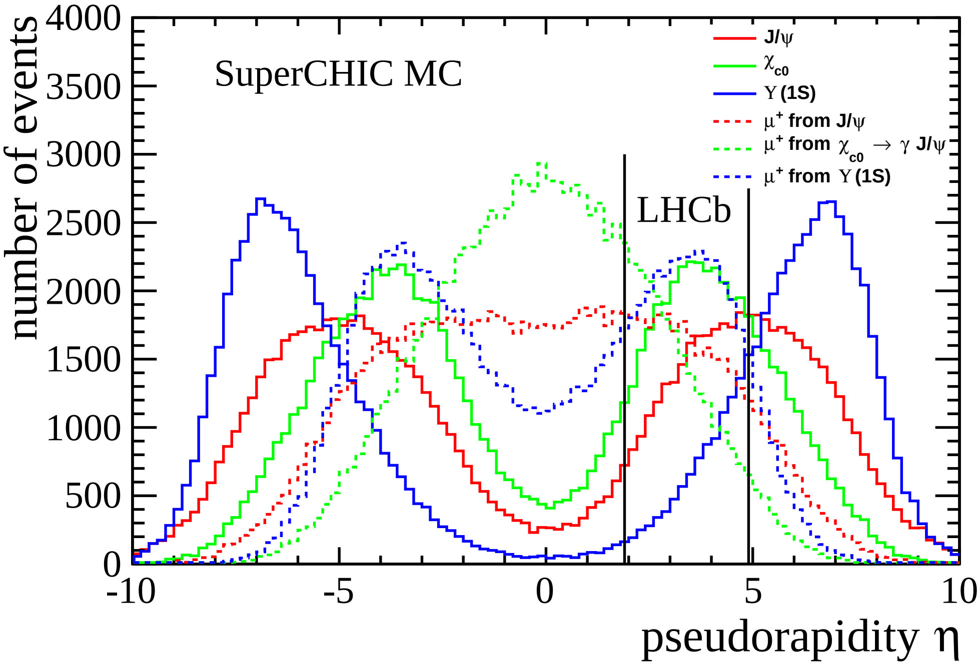

Figure 4.15: The simulated pseudorapidity distribution for CEP J/ψ, χ_{c0} and $\Upsilon(1S)$ is shown with solid lines. The distributions for one of the muons are also shown, with dashed lines. The LHCb acceptance is indicated with vertical lines.

CHAPTER
FIVE

SIMULATION OF FORWARD SHOWER COUNTERS

5.1 Introduction

The purpose of Forward Shower Counters (FSCs) is to measure particle showers produced when particles from the main proton collisions interact with the material in the LHC accelerator system. The particles in the forward direction generally have high energy, so they produce significant showers of secondary particles. FSC stations are typically placed at distances of 20–$120\,\mathrm{m}$ along the beam pipe on both sides of the interaction point, and existing and previous FSC installations cover areas in the lateral direction starting from the outer edge of the beam vacuum chamber out to radial distances of $\sim 20\,\mathrm{cm}$. The angular (or pseudorapidity) acceptance of a collider experiment can be substantially extended by installation of FSCs. Because FSCs only detect secondary showers, they can mainly be used only to veto on, or require the presence of, particles.

The detectors used as FSCs are generally scintillation counters. Scintillation counters comprise a transparent scintillator material which emits light when struck by charged particles, and a photomultiplier tube (PMT) which detects the scintillation light. Solid plastic scintillator materials are typically used. The components of the scintillation counter are enclosed in an opaque container. Scintillation counters are usually simple detectors with coarse grained segmentation and few readout channels (the LHCb SPD is an example of a more complex scintillation detector, with 12032 channels).

The installation of FSCs at LHC experiments was suggested by Albrow et al. [112] in 2009. Simulations were done showing that FSCs are a viable option for identification of diffractive interactions and proton dissociation. The installation of such detectors at LHCb was discussed by Lämsä and Orava [113]. A model was developed for simulation of the propagation of particles in a region about $100\,\mathrm{m}$ on each side of the interaction point. Models of the beam vacuum chambers and the LHC magnets were included.

The bulk of this chapter concerns the modelling of material in the region near LHCb and the simulation of particle showers originated by particles with $5 \lesssim \eta \lesssim 12$, for the determination of the efficiency of a configuration of FSCs. Only the efficiency to detect showers from single particles is given, and the simulation of specific physical processes such as proton dissociation is outside the scope of this study.

5.2 Motivation

Having additional coverage in the forward region allows the classification of interactions as inelastic, diffraction and CEP with or without proton dissociation. For CEP it is of great interest to measure the fraction of events without proton dissociation. This can be done using FSCs as a global event veto. LHCb, and perhaps ALICE, are the only LHC experiments where this is practical, because the other experiments run with high pile-up such that there will almost always be at least one inelastic interaction per event, which would cause signals in the FSCs. As will be seen in the following chapters, the estimation of the inelastic background, or equivalently the signal purity, represents the dominant uncertainty on current χ_c-meson CEP measurements at LHCb. If the acceptance was extended using FSCs, the CEP purity could be improved and the uncertainty could be greatly reduced. Discrimination between events with proton dissociation and true CEP events is important because most theoretical calculations only predict the rate of true CEP. For the diphoton process, the MC generator LPAIR can be configured to generate proton dissociation events, but the uncertainty is greater. The proton dissociation process can be written as $p + p \rightarrow p + X + p^*$, where p^* is an excited nucleon state (also $p^* + X + p^*$ of course). The excited state then decays to a small number of particles, for example a pion and a proton $p^* \rightarrow \pi^0 p$. The decay products are produced in the forward direction because the mass of p^* is generally small compared to the beam energy, and the p^* momentum is close to that of the incoming proton in CEP processes.

Scintillation counters are relatively simple and inexpensive detectors, which provide measurements of the number of particles passing through the scintillator material. The light pulses in the scintillator material are converted into electrical pulses and amplified by the PMT. The integrated signal is an approximately linear function of the number of particles which passed through the detector in a given time interval. In most FSC systems, each station has a single layer of scintillator material in a plane perpendicular to the beam axis, and the stations are placed along the beam pipe at intervals of ~ 20–$50\,\mathrm{m}$. Multiple scintillation counters are usually used on each station, due to mechanical constraints and to improve sensitivity. It is, however, difficult to extract physical information from the lateral shape of the showers, so the signals from the PMTs in a station could be combined, to reduce the number of cables. On the other hand, the segmentation may be useful for calibration. While the stations would primarily function as a veto, the pulse height from each station should be saved for off-line analysis, so that the veto cut can be optimised and calibrated.

It may be necessary to impose timing requirements on the signal, requiring that the signal arrives at a time consistent with travelling at the speed of light from the interaction point. It may not be possible to reconstruct the timing information if the signal pulses are transmitted over long cables. The time-sensitive electronics then have to be implemented at the FSC stations, in radiation-hard front-end chips. It would also be necessary to deliver a clock signal to the front-end chips, but this could be done over twisted pair cables. The front-end would either have to apply a gate signal synchronous with the $40\,\mathrm{MHz}$ bunch clock, with a phase and duration which is configurable remotely, or function as a time to digital converter. None of the front-end chips currently used in LHCb can perform all these functions. If precise timing information is not required (only separation between bunch crossings), the analogue signals may be amplified, shaped and transmitted over shielded twisted pair cables directly, and this is much simpler.

The signals could be digitised at the front end or in the LHCb counting house. The delay due to the propagation of particles to a FSC placed $100\,\mathrm{m}$ from the interaction point, and the subsequent propagation of the electrical signals back to the counting house is $\sim 900\,\mathrm{ns}$, compatible with the $4\,\mu\mathrm{s}$ latency of the L0 trigger. Due to the already low rate of the low multiplicity triggers, however, it may be better to integrate the FSCs with the HLT or only use them in offline analyses, to avoid changes to the L0 trigger.

The integration with the LHC machine is not very complicated, because the FSCs are placed outside the beam vacuum chamber. The scintillator materials, the PMTs and front-end electronics would have to withstand a harsh radiation environment. The radiation in IR8 has been estimated in a simulation [114], showing fluence rates for $20\,\mathrm{MeV}$ hadrons of 10^{10} to $10^{12}\,\mathrm{cm}^{-2}$ per year near the beam axis, integrated over the y coordinate.

5.3 FSC design and placement

There were two types of FSCs at CDF (named Beam Shower Counters) [115]. Circular detectors were used in the two positions closest to the interaction point, and larger square detectors were used in the other positions. The larger of the FSCs, shown in fig. 5.1, had sides of 6 inches, with a hole for the beam in the centre. Each station comprised two scintillation counters which each covered half the acceptance.

CMS has also installed FSCs [116], in the shutdown before the 2012 LHC run. Their use is limited to special runs with low instantaneous luminosity because CMS normally runs with too high pile-up for the detectors to be useful. At CMS, there are three stations on each side, with sensors built from $25\,\mathrm{cm}$ square scintillator tiles. The first stations are positioned at $59.1\,\mathrm{m}$ and $84.8\,\mathrm{m}$ and are built from two tiles stacked vertically. There is a hole for the beam where the tiles are joined, similar to the design in fig. 5.1, but the complete detectors have a rectangular shape. The next stations are at $114.1\,\mathrm{m}$, and have a square shape built from 2×2 tiles.

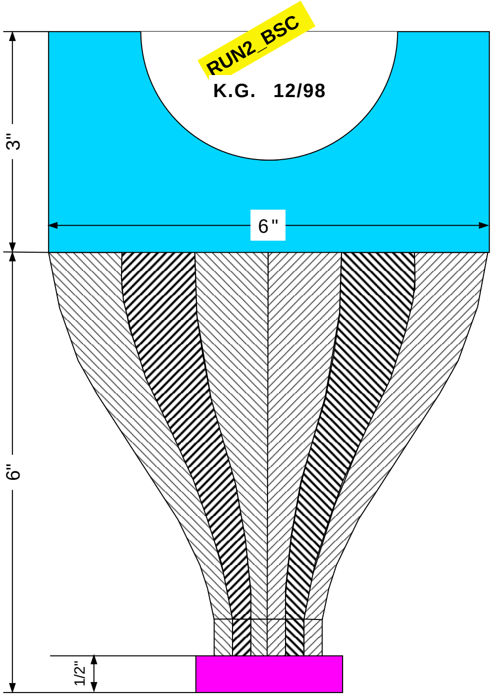

Figure 5.1: Design of BSCs at CDF. The drawing represents one half of a BSC station, covering half the angular acceptance. Figure from Ref. [115].

The presence of the magnets, shielding walls and other infrastructure eliminates certain positions for the FSCs. The details of positioning at LHCb still need to be worked out, but it is confirmed that placement of FSCs near the beam is possible in general. A hermetic coverage is not needed since the FSCs detect secondary showers, not the primary particles. For the simulation, three FSCs are used in the direction of the LHCb spectrometer acceptance, at distances of $23\,\mathrm{m}$, $69\,\mathrm{m}$ and $98\,\mathrm{m}$, and four stations are used in the opposite direction, at $9\,\mathrm{m}$, $23\,\mathrm{m}$, $69\,\mathrm{m}$ and $98\,\mathrm{m}$. These positions are not occupied by elements of the LHC beam optics. The simulation uses square detectors with sides of $800\,\mathrm{mm}$. The holes for the beam pipe have radius $50\,\mathrm{mm}$ for the detectors at $9\,\mathrm{m}$, $23\,\mathrm{m}$ and $69\,\mathrm{m}$, and $130\,\mathrm{mm}$ for the detectors at $98\,\mathrm{m}$.

5.4 Insertion Region 8

The LHC has eight long straight sections, called Insertion Regions (IRs) 1 to 8, which are the locations of experiments and utility services [6]. LHCb is located in IR8, between a betatron cleaning facility in IR7 and the ATLAS experiment at IR1. The layouts of LHC and IR8 are shown in fig. 5.2. The region at negative z is referred to as the left side and the region at positive z is referred to as the right side, corresponding to their positions as

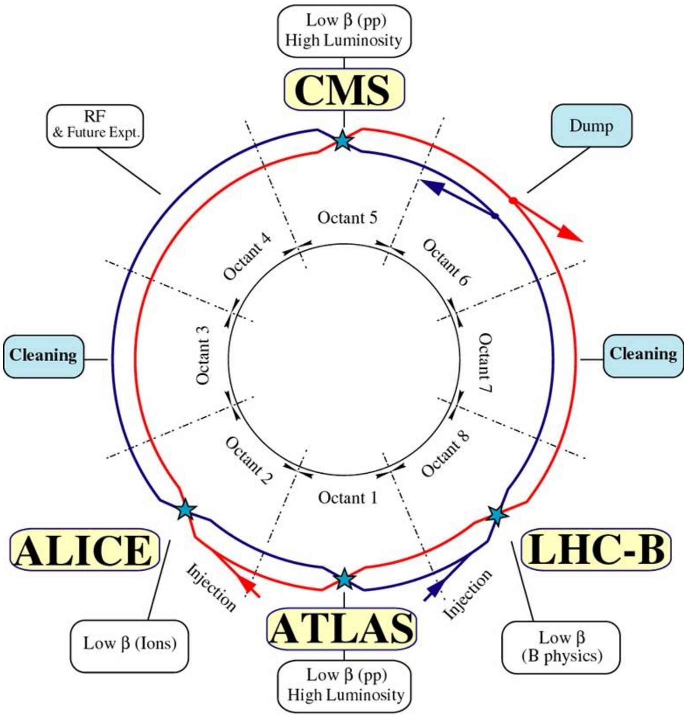

(a) Schematic layout of LHC

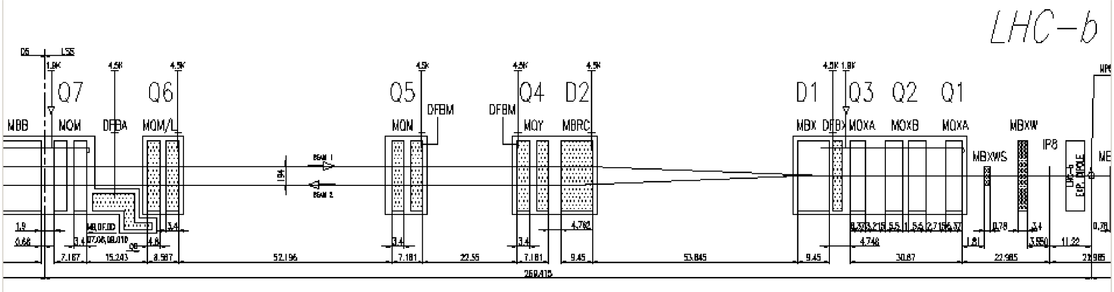

(b) Left side

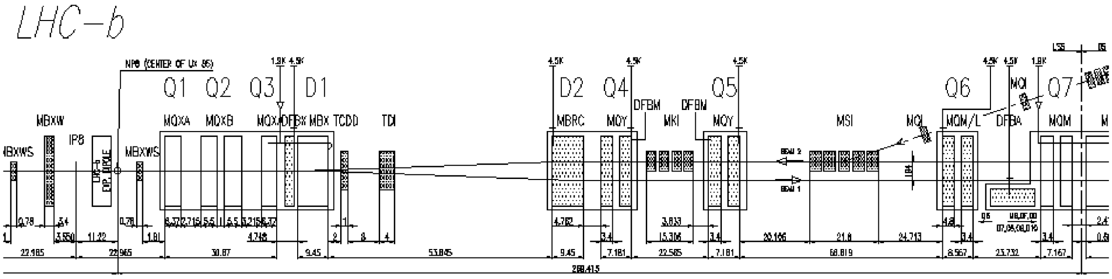

(c) Right side

Figure 5.2: The top diagram shows the LHC experiments and the trajectories of the two beams (the separation between the beams is exaggerated). The middle and bottom drawings show the magnets and collimators (beam optics) of IR8, as seen from the centre of the ring looking outwards. The horizontal coordinate is z. Figures from Ref. [6].

seen from the centre of the LHC ring. On the left side, proton beam 1 is incoming with respect to LHCb, and beam 2 is outgoing, and vice versa on the right side. The right side of the interaction point also contains magnets for accepting incoming bunches for LHC beam 2 from the SPS. LHCb is located in a large cavern at the centre of IR8.

The main components of the IR8 beam optics, moving away from the interaction point up to distances of $\sim 100\,\mathrm{m}$ are [6] (with general names of LHC equipment types in parentheses):

- The LHCb spectrometer dipole magnet (MBLW) on the right side and a corresponding non-superconducting compensator magnet (MBXWH) on the left side. MBXWH provides an integrated field of the same magnitude as the LHCb dipole magnet, with the opposite polarity.
- Additional compensator magnets on both sides (MBXWS), to close the spectrometer dipole bump.
- A magnetic triplet located from $23\,\mathrm{m}$ to $54\,\mathrm{m}$ from IP8. This is a series of three superconducting quadrupole magnets (MQXA, MQXB) which focus the beams to produce the correct beam size at the interaction point. The triplet also includes correction dipole magnets (MCBXV, MCBXH).
- A pair of superconducting separation/recombination dipole magnets (MBX), D1 and D2, for separating the two beams such that each beam can enter its own beam pipe in the arcs following the straight section. D1 immediately follows the triplet, at 58–$68\,\mathrm{m}$ from the interaction point.
- Tertiary collimators (TCTV, TCTH) to protect the inner triplet and D1 from beam halo.
- Machine protection elements related to the injection of beam 2, on the right side of the interaction point: a fixed mask (TCDD) and a movable collimator (TDI).

The TDI and TCTV are vertical collimators, which act on the beam moving towards the interaction point (the beams are separated by $\sim 30\,\mathrm{mm}$ at the position of the TDI, so they can be collimated separately). The collimation is achieved by restricting the vertical aperture using movable collimator jaws. Being vertical collimators, the TDI and TCTV do not significantly restrict the aperture in the horizontal ($y = 0$) plane. The TDI is retracted after injection, so it also has a large vertical aperture under normal operation. The TCDD is fixed and has an approximately rectangular aperture.

The LHCb cavern is shown in fig. 5.3. The detector system has a length of approximately $22\,\mathrm{m}$. To the left of the VELO and to the right of the muon stations, the beams exit the cavern and enter the main LHC tunnel. A visualisation of the standard LHCb detector

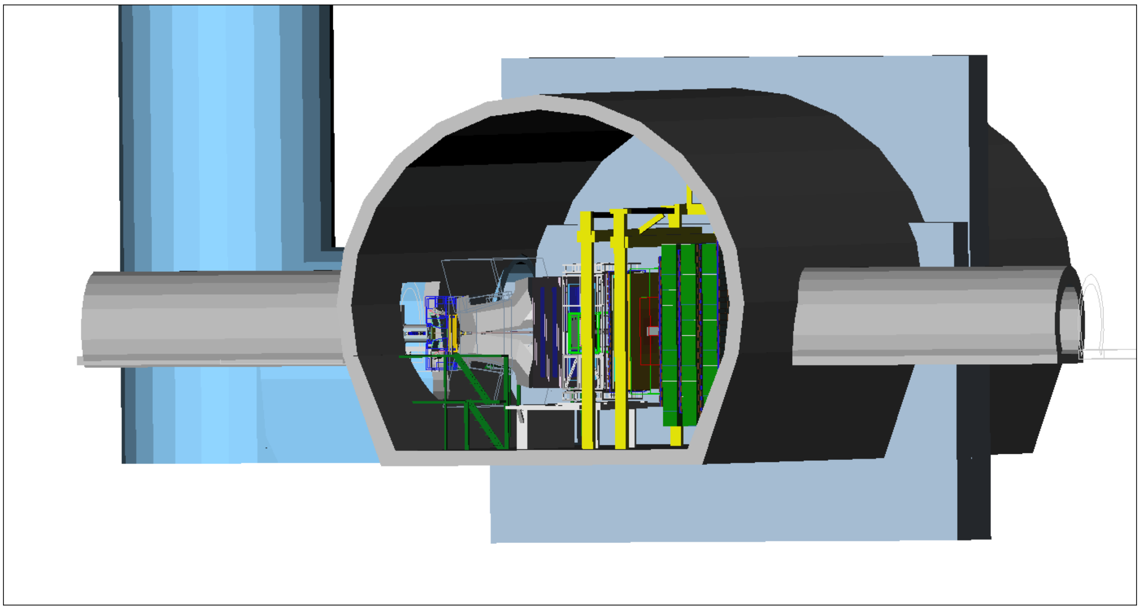

Figure 5.3: LHCb experiment cavern overview, in the standard LHCb detector description. The figure is generated using the visualisation program Panoramix.

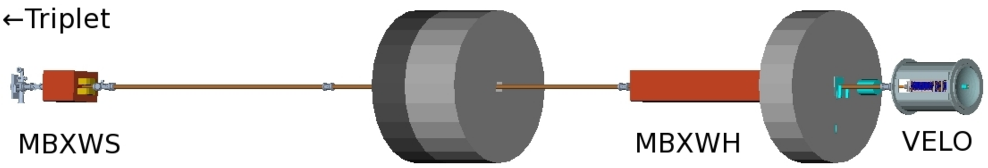

Figure 5.4: Visualisation of the magnets and shielding walls upstream of the VELO, between the interaction point and the triplet.

description in the region beyond the VELO is shown in fig. 5.4. There are two shielding walls upstream of the VELO, which have the purpose of shielding the detector from particles coming from the tunnel (machine induced background). The first wall near LHCb is made of concrete, has a depth of $80\,\mathrm{cm}$, and is located $2.2\,\mathrm{m}$ from the interaction point. The second wall is placed between the compensator magnets, $10.4\,\mathrm{m}$ from the interaction point, and is made of $120\,\mathrm{cm}$ of concrete followed by $80\,\mathrm{cm}$ of iron. The openings for the beam pipe in the shields are square with $20\,\mathrm{cm}$ sides. In the downstream (right) region, the muon filters act as shielding. There is also shielding around the triplet similar to the second wall on the left side, but with a larger opening to accommodate the quadrupole magnet.

The beams need to be separated from each other in order to avoid interactions of

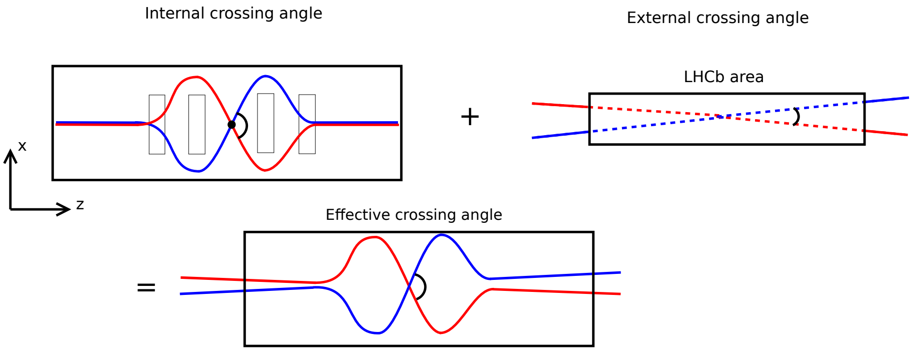

Figure 5.5: Internal and external crossing angles in the original LHC optics. Beam 1 is shown in blue and beam 2 in red. The four boxes in the top left figure represent the MBXWH magnet, two MBXWS magnets and the LHCb spectrometer magnet. The angles are exaggerated.

bunches away from the interaction point. For this reason, the two beams traverse IR8 at a small angle, converging at the interaction point. The angle is known as the external crossing angle. Another angular separation is inevitably produced by the LHCb spectrometer magnet and its corresponding compensator MBXWH. The half crossing angle (angle of each beam with respect to the z axis) caused by the LHCb magnet is $135\,\mu\mathrm{rad}$ in the horizontal plane. The sign of the crossing angle depends on the polarity of the field of the LHCb magnet. In the original design, the external crossing angle was also in the horizontal plane, and had to be added to the internal crossing angle [3]. This is illustrated in fig. 5.5. The sign of the external crossing angle cannot be changed, so in the LHCb "magnet up" configuration a large half external crossing angle of $-210\,\mu\mathrm{rad}$ is required to achieve the necessary separation between the beams. In 2012 a vertical crossing angle was used for the external angle [117], so it was no longer necessary to have a large external crossing angle when the internal crossing angle was positive.

5.5 Geometrical model for simulation

A model of IR8 up to distances $|z| \approx 100\,\mathrm{m}$ from the interaction point was created. Cross-sections of the model in the x-z and y-z planes are shown in fig. 5.6. Table 5.1 shows an inventory of the objects in the model (not including the beam pipe). The model is based on a re-implementation of the Lämsä-Orava (LO) model [113], but some improvements have been made. Simple geometrical shapes are used to model the beam pipe, magnets, collimators, radiation shielding walls and the tunnel. The description of the radiation shielding walls and the tunnel near LHCb was adapted from the LHCb detector description. The description of the LHCb detector is greatly simplified relative to the standard detector description, but this does not affect the results since the relevant particles are outside the LHCb acceptance. A detailed description of the beam pipe in LHCb

Item	Start z (m)	End z (m)
Upstream / Left region		
Left FSC 3	-98.0	-98.0
Collimator TCTV	-74.6	-73.4
Left FSC 2	-69.0	-69.0
Dipole magnet (MBX)	-67.8	-58.4
Quadrupole magnet 3 (MQXA)	-53.3	-47.0
Quadrupole magnet 2B (MQXB)	-44.0	-38.5
Corrector dipole (MCBXH)	-38.2	-37.8
Quadrupole magnet 2A (MQXB)	-37.5	-32.0
Corrector dipole (MCBXH)	-30.1	-29.6
Quadrupole magnet 1 (MQXA)	-29.3	-23.0
Left FSC 1	-22.8	-22.8
LHCb region		
Compensator magnet (MBXWS)	-21.2	-20.4
Upstream shielding 2 (Iron)	-12.4	-11.6
Upstream shielding 2 (Concrete)	-11.6	-10.4
Central FSC 1	-9.0	-9.0
Compensator magnet (MBXWH)	-6.9	-3.6
Upstream shielding (Concrete)	-3.0	-2.2
LHCb dipole magnet	4.0	7.0
LHCb muon filters	16.6	19.8
Compensator magnet (MBXWS)	20.4	21.2
Downstream / Right region		
Right FSC 1	22.8	22.8
Quadrupole magnet 1 (MQXA)	23.0	29.3
Downstream shielding (Concrete)	24.6	25.8
Downstream shielding (Iron)	25.8	26.6
Corrector dipole (MCBXH)	29.6	30.1
Quadrupole magnet 2A (MQXB)	32.0	37.5
Corrector dipole (MCBXH)	37.8	38.2
Quadrupole magnet 2B (MQXB)	38.5	44.0
Quadrupole magnet 3 (MQXA)	47.0	53.3
Dipole magnet (MBX)	58.4	67.8
Right FSC 2	69.0	69.0
Collimator TCDD	69.8	70.8
Collimator TCTV	73.4	74.6
Collimator TDI	79.0	83.0
Right FSC 3	98.0	98.0

Table 5.1: Inventory of FSCs, magnets and collimators in the LO model. The positions in the z coordinate are given.

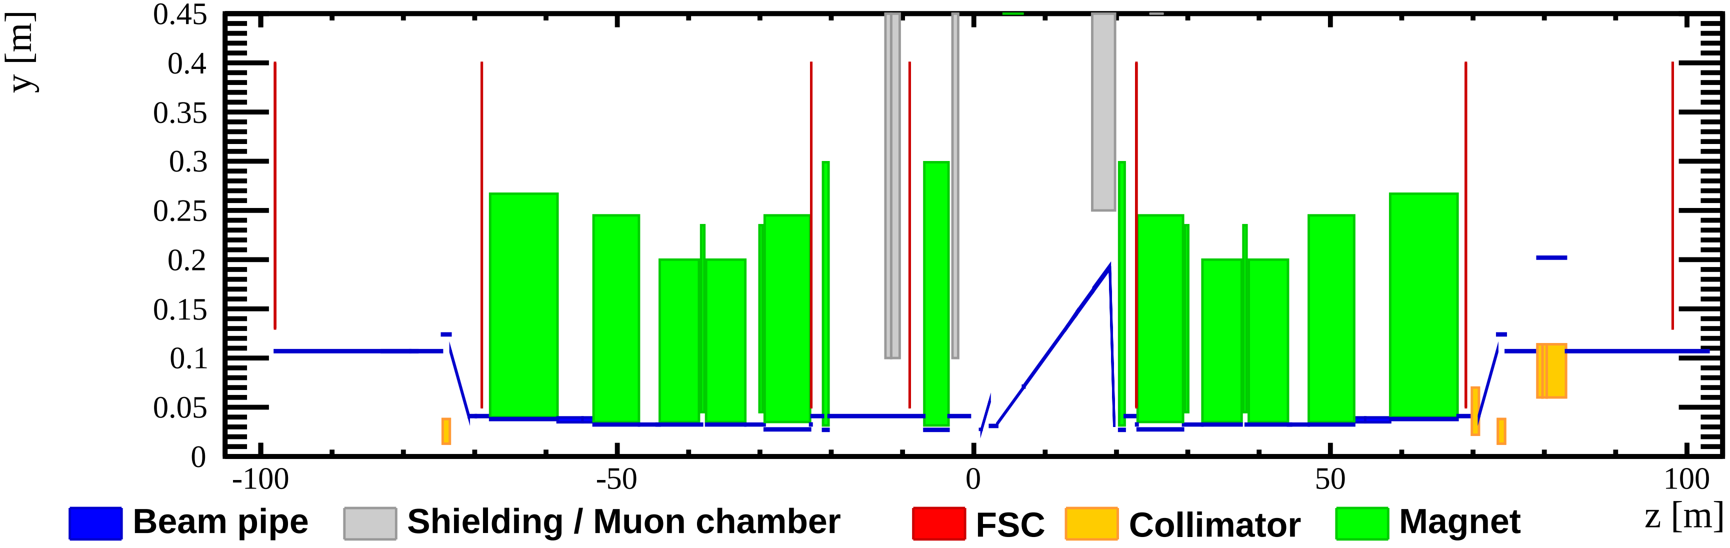

(a) $x = 0$

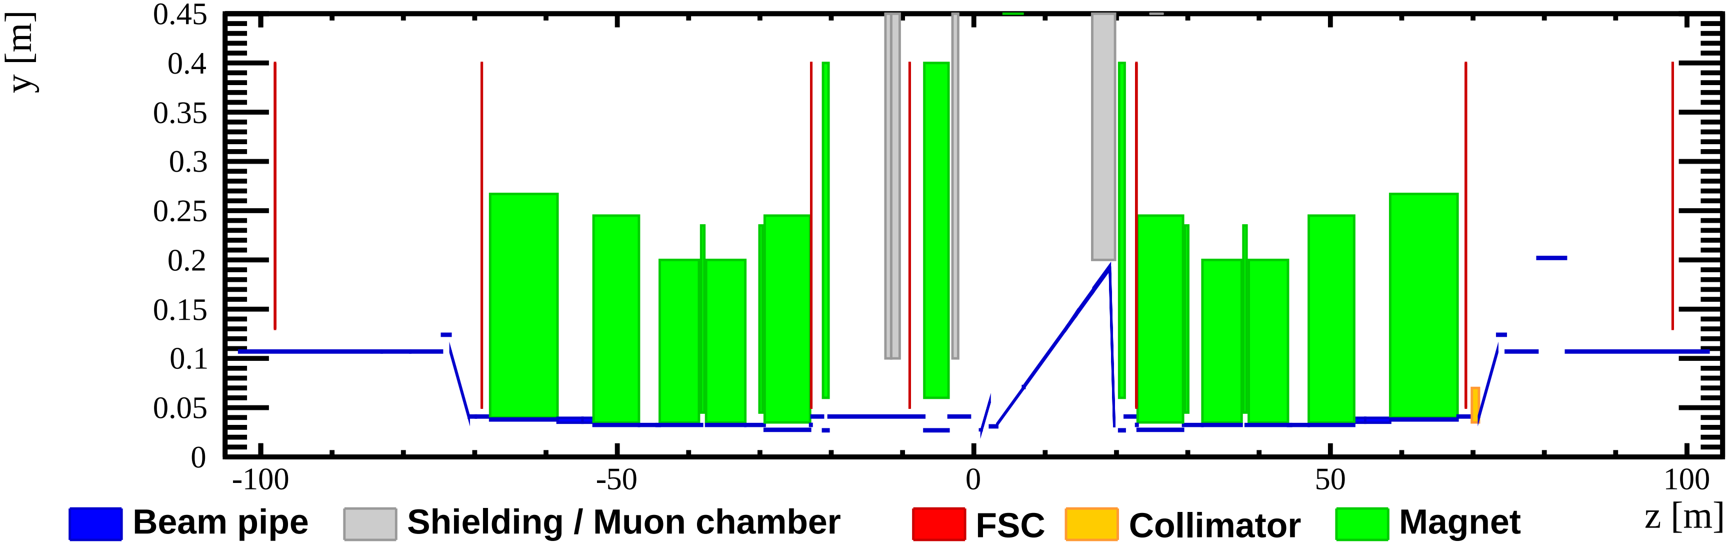

(b) $y = 0$

Figure 5.6: Cross-sectional views of the model at $x = 0$ and $y = 0$, showing most objects in the model. The interaction point is at $z = 0$. The LHCb spectrometer magnet and the right side shielding wall are not shown because their inner radii are greater than $0.45\,\mathrm{m}$.

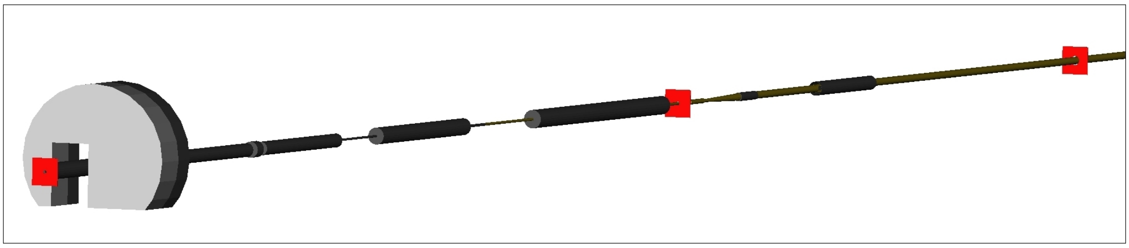

(a) Right side of the model, without the tunnel. The FSCs are shown in red.

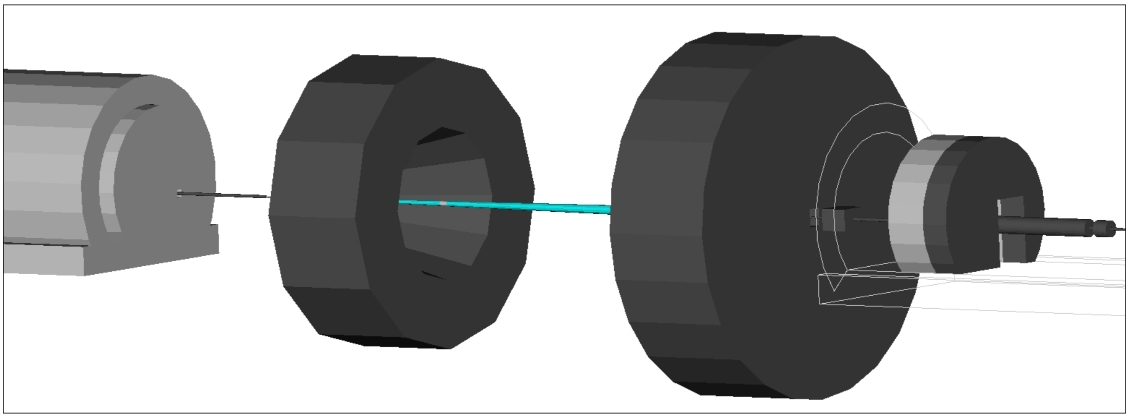

(b) Central region. The left side tunnel and the first shielding wall is visible on the left. The following objects are the LHCb magnet, the muon filters, the MBXWS, the first quadrupole magnet and the downstream shielding wall. An outline of the right side tunnel is shown.

(c) Objects around $z = -23\,\mathrm{m}$, with quadrupole to the left, then an FSC, then MBXWS. Some of the tunnel is not shown, to allow visibility inside.

(d) TDI, viewed from above, with the z axis pointing up and to the right. The TDI vacuum chamber is shown semi-transparently.

Figure 5.7: Visualisations of elements of the geometrical model, generated using the Panoramix event display software.

is included. The geometrical description of the LHC magnets and collimators is based on the LO model, the LHC Design Report [3] and the Layout Database [118]. Visualisations of details of the model are shown in fig. 5.7. The magnetic fields are modelled as ideal fields with magnitudes taken from the LHC optics configuration for "ultimate" energy and intensity ($7\,\mathrm{TeV}$ beams) [119]. The field of the LHCb magnet is simulated as a uniform field in the y-direction with a Gaussian dependence on z. The "magnet down" configuration is used in the simulation, corresponding to a negative y-component of the LHCb magnet field, and a $-200\,\mu\mathrm{rad}$ half effective crossing angle is used.

The compensator magnets, the FSCs and the holes in the shielding walls are modelled with rectangular shapes in the transverse plane, centred on the beam axis. The muon filters are modelled as a large iron cylinder with a square hole for the beam. The LHC tunnel has a cylindrical shape with a flat floor, and the beam axis is offset from the central axis of the tunnel by $(x, y) = (704\,\mathrm{mm}, -300\,\mathrm{mm})$. The models of the beam pipe and the superconducting magnets are cylinders. The magnets are modelled as iron, and the shapes represent the yokes and coils, which are the main masses of the magnets. The beam pipe outside LHCb is made of copper, and the beam pipe in LHCb is made of beryllium and other materials. The beam pipe in LHCb has a double conical shape, designed to minimise the material interactions. The TDI and TCTV are modelled with pairs of boxes located above and below the beam axis representing the collimator jaws. The positions are offset from $x = 0$ because the jaws only affect beam 1. The TDI model is shown in fig. 5.7d. The material of the TCDD mask and the TCTV jaws is copper. The TDI jaws include sections of hexagonal boron-nitride, aluminium and copper.

5.6 Simulation

The simulation is performed using the Gauss application, which calls Geant 4 [29] to simulate the passage of particles through the material. The goal is to map the acceptance of a certain configuration of FSCs, without invoking a model of diffraction or CEP. The particles used for simulation are the pions, and separate runs with π^+, π^0 and π^- are simulated. The momentum of a particle is completely specified by its pseudorapidity η, its transverse momentum p_T and its azimuthal angle ϕ. The simulated transverse momenta are $0.1\,\mathrm{GeV}$, $0.2\,\mathrm{GeV}$, $0.5\,\mathrm{GeV}$ and $1.0\,\mathrm{GeV}$. These values represent typical p_T in low-mass CEP processes, and similar values are expected in inelastic production events. At the beam momentum of $7\,\mathrm{TeV}$, which is the maximum possible momentum of any final state particles, the maximum pseudorapidity for particles with $p_T = 0.1\,\mathrm{GeV}$ is $|\eta| = 11.8$, and for $p_T = 1.0\,\mathrm{GeV}$ it is 9.5. The minimum pseudorapidity (absolute value) in the simulation is set to 4.0. Values of η are generated from a uniform distribution between these minimum and maximum values. After the simulation, the data are binned in η when the efficiency is computed. The efficiency is averaged over the azimuthal angle by generating particles at random ϕ.

A dataset of 4000 events is generated for each pion charge, z-direction and p_T. In the simulation, the FSCs are initially assumed to have perfect efficiency for detecting electrons, muons, charged pions, charged kaons and protons which hit them. Five hits in any single FSC are then required for a signal. It may be necessary to use such a cut on real data, to reduce the effect of single particle backgrounds and detector noise. The cut value is set at a point where the efficiency on signal showers is not greatly reduced, but on real data the cut value must be calibrated for each FSC station. An event display of

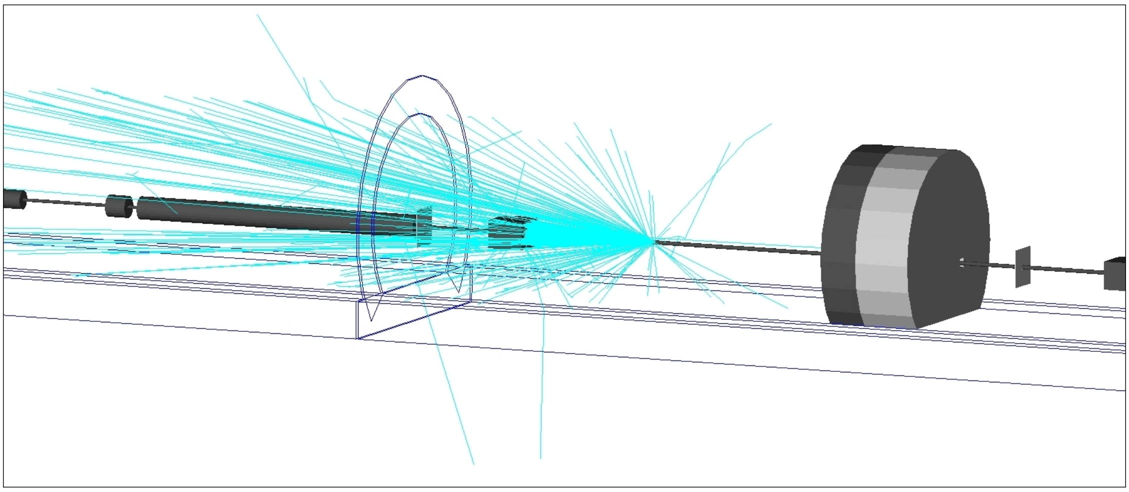

Figure 5.8: Simulated forward shower event. The original pion hits the beam pipe before the MBXWS and produces a large number of secondary particles. Many particles are screened by the MBXWS, but exactly five particles hit the first FSC, which is located between the MBXWS and the superconducting triplet. All particle trajectories are shown in light blue.

a single particle hitting the beam pipe and producing a large particle shower is shown in fig. 5.8.

The magnetic fields in the model are checked by simulating protons emerging from the interaction point with energy 7 TeV, representing beam protons which did not interact at the LHCb interaction point. The simulated beam protons do not interact in the model, as expected, and their trajectories are consistent with expectations. The simulations presented in this chapter represent a preliminary investigation. Geant 4 provides excellent models for interactions of particles with material, but it is not specifically designed to track high energy particles in the magnetic fields of accelerators, and there may be some inaccuracy. If it is decided to continue with an FSC programme, it may be necessary to perform more detailed simulations with software that is more suitable for modelling beam particles, or at least a better model of the optics (as was done for the FP420 [120] research and development project).

5.6.1 Simulation results

The simulation results for the left (negative η) side are shown in fig. 5.9, and the corresponding results for the right side are shown in fig. 5.10. The plots show the detection efficiency as a function of η for each FSC separately, and for the combined FSC system. The efficiency is the fraction of events in which there are five or more hits in a given FSC. The total efficiency is the probability that any FSC has at least five hits. The sub-figures correspond to the efficiencies for π^+, π^0 and π^-, and for different values of p_T. The p_T is used as a parameter because it is related to the mass of the excited system in proton

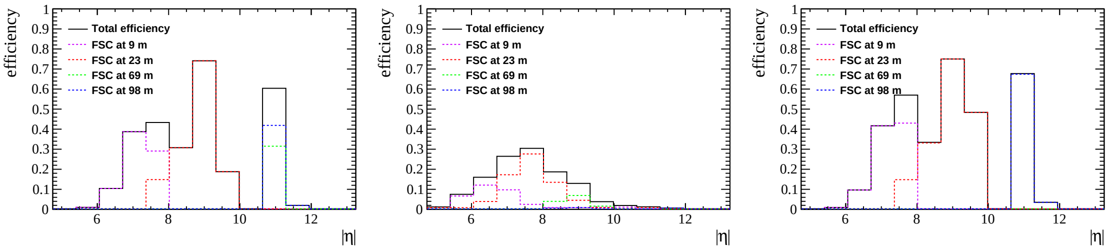

(a) π^+, $p_T = 0.1\,\mathrm{MeV}$ (b) π^0, $p_T = 0.1\,\mathrm{MeV}$ (c) π^-, $p_T = 0.1\,\mathrm{MeV}$

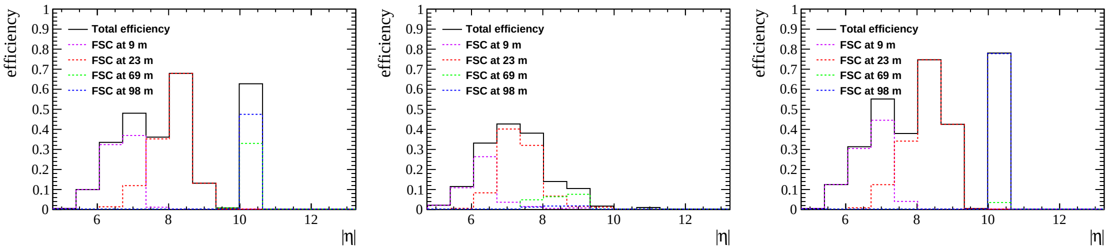

(d) π^+, $p_T = 0.2\,\mathrm{MeV}$ (e) π^0, $p_T = 0.2\,\mathrm{MeV}$ (f) π^-, $p_T = 0.2\,\mathrm{MeV}$

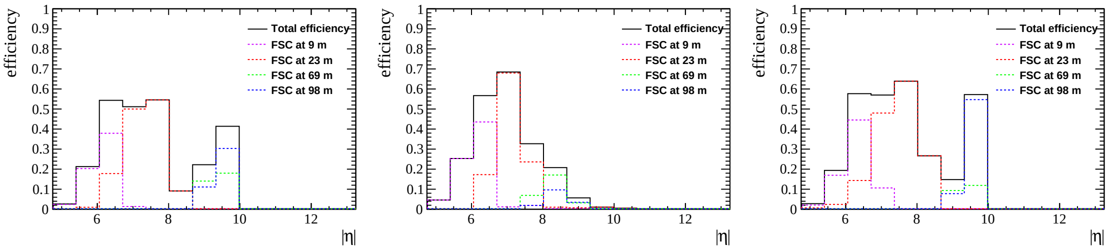

(g) π^+, $p_T = 0.5\,\mathrm{MeV}$ (h) π^0, $p_T = 0.5\,\mathrm{MeV}$ (i) π^-, $p_T = 0.5\,\mathrm{MeV}$

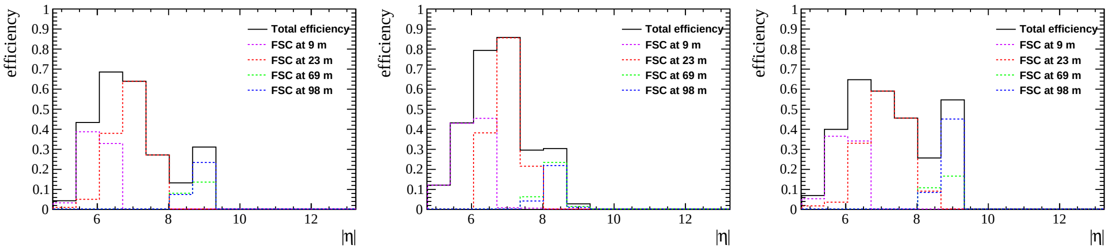

(j) π^+, $p_T = 1.0\,\mathrm{MeV}$ (k) π^0, $p_T = 1.0\,\mathrm{MeV}$ (l) π^-, $p_T = 1.0\,\mathrm{MeV}$

Figure 5.9: Single particle FSC efficiencies for the left side model as functions of η, shown for different fixed values of p_T. The efficiencies for original pions of positive, neutral and negative charge are shown separately.

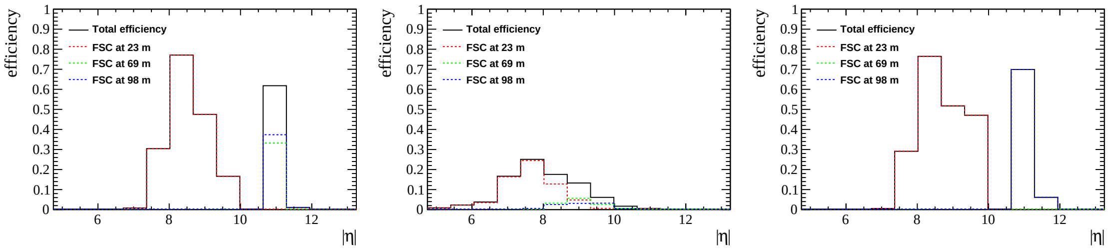

(a) π^+, $p_T = 0.1\,\mathrm{MeV}$ (b) π^0, $p_T = 0.1\,\mathrm{MeV}$ (c) π^-, $p_T = 0.1\,\mathrm{MeV}$

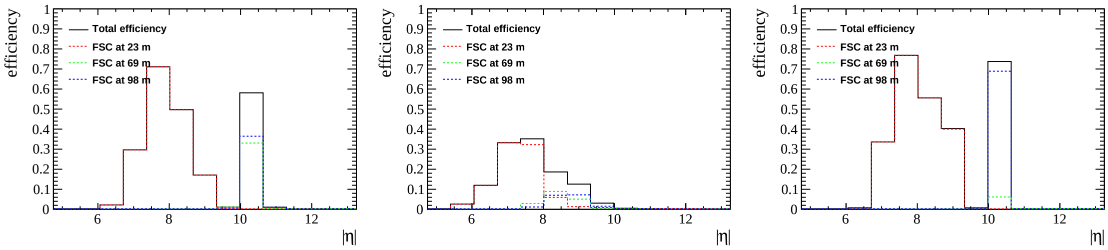

(d) π^+, $p_T = 0.2\,\mathrm{MeV}$ (e) π^0, $p_T = 0.2\,\mathrm{MeV}$ (f) π^-, $p_T = 0.2\,\mathrm{MeV}$

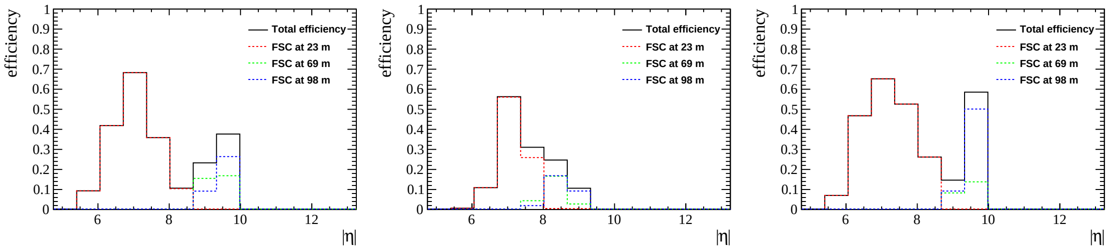

(g) π^+, $p_T = 0.5\,\mathrm{MeV}$ (h) π^0, $p_T = 0.5\,\mathrm{MeV}$ (i) π^-, $p_T = 0.5\,\mathrm{MeV}$

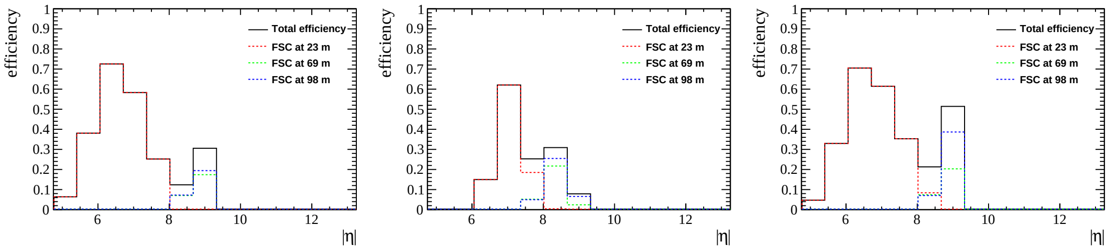

(j) π^+, $p_T = 1.0\,\mathrm{MeV}$ (k) π^0, $p_T = 1.0\,\mathrm{MeV}$ (l) π^-, $p_T = 1.0\,\mathrm{MeV}$

Figure 5.10: Single particle FSC efficiencies for the right side model as functions of η, shown for different fixed values of p_T.

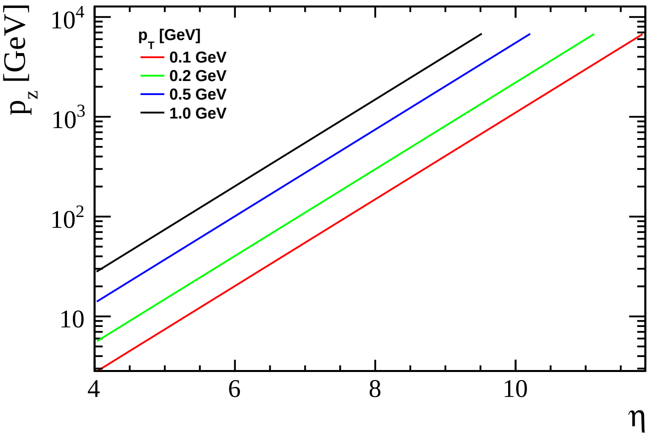

Figure 5.11: Relation between longitudinal momentum and pseudorapidity for different values of p_T. The p_z increases by approximately a factor of 2.7 for each unit of pseudorapidity.

dissociation events. Because each plot corresponds to a fixed p_T, the energy varies with η. The relation between the longitudinal momentum and η for different values of p_T is shown in fig. 5.11.

The FSC at $23\,\mathrm{m}$ provides excellent efficiency for charged pions in most configurations. For low p_T there is a gap in the efficiency corresponding to the un-instrumented area around the triplet and D1. The FSC at $69\,\mathrm{m}$ does not provide much additional efficiency over the one at $98\,\mathrm{m}$, except for at low p_T for positively charged particles (compare the blue and the black histograms). The FSCs at $69\,\mathrm{m}$ and $98\,\mathrm{m}$ provide high efficiency in a small range of η. The efficiency is much lower for neutral pions, which are not affected by the magnetic fields. The π^0 mesons decay quickly, and the branching ratio of the decay to two photons is $98.8\,\%$. A high efficiency is however seen for neutral pions in the datasets with the highest p_T. The average of the efficiencies for all p_T and η and pion charges is $22\,\%$.

It is possible to use the full LHCb detector description in the central region of $|z| < 23\,\mathrm{m}$, instead of the simplified description discussed above. This "hybrid" model contains a slightly more detailed description of the magnets and the beam pipe, and a much more detailed description of the VELO. At the lowest $|\eta|$ values considered here, there may be interactions with the VELO sensors and the RF foil. Ideal magnetic fields are used instead of the default LHCb field map, since the field map does not include the MBXWS and MBXWH fields. Problems with overlapping volumes make it necessary to use a simpler description of a segment of the tunnel and the upstream shielding walls, and this is the reason that the hybrid model is not used by default. The hybrid model is compared with the default model in fig. 5.12. The figure shows the total efficiency averaged over the three pion charges. There is a significant difference at some η values, mainly for low

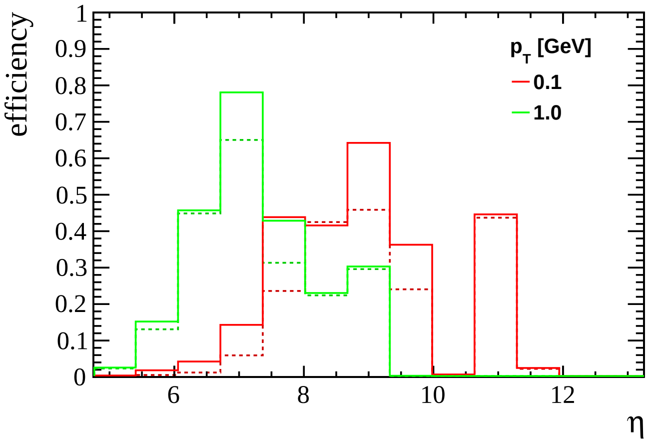

Figure 5.12: Efficiency in the default model and the hybrid model, respectively shown with a dotted line and a solid line. The efficiency is averaged over the left and right side, and over all pion charges. The FSC at $z = -9\,\mathrm{m}$ is not included in this plot, because it is not used in the hybrid model.

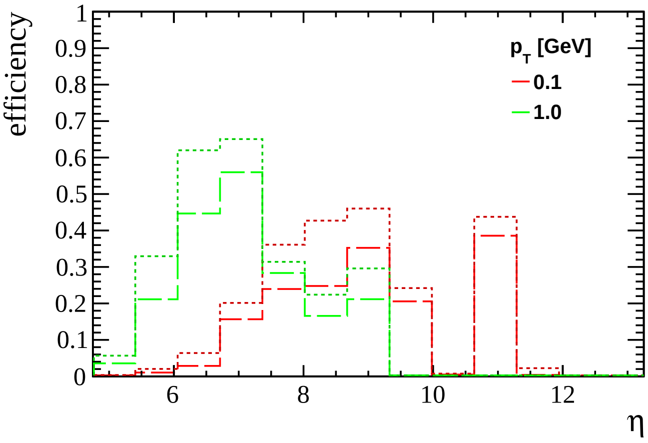

Figure 5.13: Efficiency for FSCs with size $400\,\mathrm{mm}$ is shown with punctured lines, and the default FSC size of $800\,\mathrm{mm}$ is shown with dotted lines.

p_T. The discrepancy appears to be caused by differences in the description of the material outside the beam pipe, in the right side model. The hybrid model may be useful for the determination of the efficiency of the combined system of the standard LHCb detectors and the FSCs.

The simulated scintillators are larger than those used at CDF and CMS. It is thus interesting to see how the efficiency changes when reducing the size of the scintillators. The efficiency for square FSCs with sides of $400\,\mathrm{mm}$, half of the default size, is shown in fig. 5.13 (punctured lines). The efficiency decreases over the whole range of η by up to $\sim 10\,\%$, for both low and high p_T, which is not a large decrease considering that the area is reduced by a factor of four. It is difficult to estimate in advance what the cut on

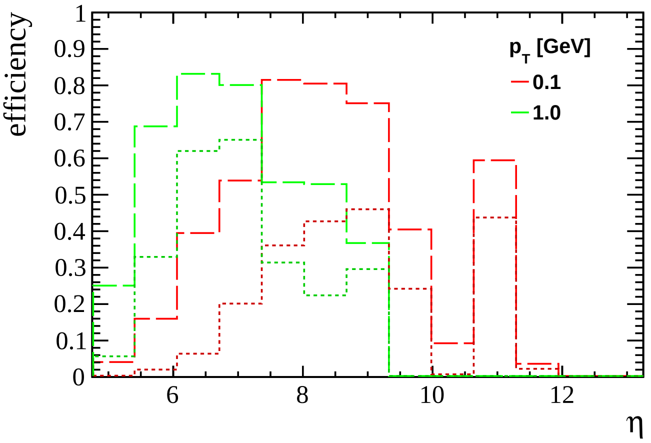

Figure 5.14: Efficiency when counting single particle hits is shown with punctured lines, and the efficiency for the default cut of five hits is shown with dotted lines.

the number of hits to accept a detection should be. The efficiency when counting single particle hits, instead of requiring five hits, is shown in fig. 5.14. The efficiency then increases significantly, showing that the cut should be set as low as possible.

5.7 Backgrounds

The main backgrounds are expected to be from beam-gas interactions in the long straight section and from interactions of beam protons with the collimators. The level of background may be higher in the tunnel than in the LHCb cavern, because the cavern is protected by significant shielding. In the inner region of muon station 5, a per event per area rate of $10^{-5}\,\mathrm{cm}^{-2}$ hits was seen [22] in events without bunch crossings in 2010. This corresponds to a mean number of hits of 0.06 per event for an $80\,\mathrm{cm}$ FSC, but the hits may not be evenly distributed, so the rate could be higher close to the beam. The number of hits cannot be assumed to have a Poisson distribution, because the machine induced background will typically come as showers.

Scintillation counters can be made with a time resolution better than $1\,\mathrm{ns}$ so one can impose timing requirements to reduce the backgrounds from the beams. The simulation predicts that over $86\,\%$ of the shower particles arrive within $1\,\mathrm{ns}$ of the beam protons.

5.8 Conclusions

The rejection of particle showers in the forward region is valuable for CEP studies. Using vetoes at the event level appears to be the only viable model-independent method for measuring a (relatively) pure low-mass CEP signal. Proton taggers can be used to measure the outgoing protons for systems of higher mass ($M_X \gtrsim 40\,\mathrm{GeV}$), and the installation of such detectors at LHCb can also be considered. The proton dissociation background to CEP produces systems of a few particles in the forward direction, and the single particle efficiency given above represents an approximate lower bound on the efficiency for proton dissociation.

The simulated FSC system provides a good efficiency in some ranges of η and p_T. The FSC at $69\,\mathrm{m}$ could be eliminated without a major impact on the efficiency, but it may help with calibration, by searching for events with a coincidence of hits in that station and the one at $98\,\mathrm{m}$. The installation of an FSC on the left side of the interaction point, behind the muon stations, may be difficult because there is little space there. It could also be an option to use hits in muon station M5 to supplement the FSCs. The scintillators at $23\,\mathrm{m}$ would however be indispensable because they provide the bulk of the acceptance. The efficiency is poor in the region in the shadow of the triplet and the following superconducting dipole magnet, because it is not possible to instrument these areas with FSCs. The efficiency then increases significantly for large $|\eta|$. The great dependence on charge, η and p_T (or p_z) of the efficiency may introduce unknown biases in the CEP signal. Inelastic dimuon events produced in the "diphoton" (Section 4.3.2) process may be used to calibrate the efficiency, but the theoretical uncertainty is large. Lämsä and Orava provided results for a p_T-distribution $N \propto \exp(-p_T^2/(0.14\,\mathrm{GeV}^2))$, and the results presented here have a slightly higher efficiency. A model which more closely resembled the LO model was also used, and the results were consistent with the LO results. The installation of FSCs appears to be a good strategy, but it is not straight-forward, and may not give anything close to a perfect efficiency. The detectors would have to be installed in a harsh radiation environment.

CHAPTER

SIX

ANALYSIS OF THE χ_c CEP SIGNAL

6.1 Introduction

The cross sections for Central Exclusive Production (CEP) of χ_c and χ_b mesons at $\sqrt{s} = 7\,\text{TeV}$ are measured using the LHCb dataset from 2011. The χ_c and χ_b mesons are reconstructed in the decays $\chi_c \to \gamma J/\psi$ and $\chi_b \to \gamma \Upsilon(1S)$, with $J/\psi \to \mu^+\mu^-$ and $\Upsilon(1S) \to \mu^+\mu^-$. Most of the analysis effort is concentrated on the χ_c, because the number of candidates is almost three orders of magnitude greater than for χ_b. Many methods are also directly transferable to the χ_b.

CEP events are selected using a cut on the amount of activity in the event. This significantly enriches the signal, because real CEP events contain no other particles than those from the signal decay, while most other interactions produce large numbers of particles in the "underlying event" and fragmentation. A cut-based analysis is used to select the χ_c candidates. The main components of the analysis are the determination of the efficiency of each component of the reconstruction, a fit to the mass distribution to measure the yield, and finally the cross section computation. There is a background process which produces candidates with the same mass distribution as the signal, so an additional analysis is done to distinguish this background from true CEP on a statistical basis. This results in a purity factor which represents the fraction of the candidates which are from genuine CEP interactions.

The cross sections for central exclusive χ_{cJ} production will be computed as

$$\sigma_J = \frac{n_J p}{\epsilon_{tot,J} \mathcal{L}_{eff}} \quad \text{for } J = 0, 1, 2\ , \tag{6.1.1}$$

where n_J is the number of candidates which are reconstructed, selected and fitted, p is the purity, $\epsilon_{tot,J}$ is the total efficiency and $\mathcal{L}_{eff}$ is the effective integrated luminosity. The result is the production cross section times the branching ratios of $\chi_c \to \gamma J/\psi$ and $J/\psi \to \mu^+\mu^-$. It is given for events in which the decay products γ, μ^+, μ^- are all in a fiducial pseudorapidity acceptance of 2.0–4.5. The ratio of the cross sections for the χ_c-states is also computed. The determination of the factors of (6.1.1) is described in this chapter, except for $\mathcal{L}_{eff}$ which is common for χ_c and χ_b, and will be described in Chapter 8. The analysis of χ_b is described in the following chapter. The combination of the factors to give the cross sections is described in Chapter 8.

The backgrounds are discussed in Section 6.2. The MC datasets are described in Section 6.3, and the triggers in Section 6.4. The selection cuts are described in Section 6.5. The number of candidates n_J is determined using a fit to the mass spectrum, described in Section 6.6. Section 6.7 describes the estimation of the signal purity p using a fit to the p_T-distribution of the muon pairs. Section 6.8 describes the efficiency determination. The strategy is to use the simulated events to measure the efficiency, and then to validate the simulation using real data. Most of Section 6.8 is dedicated to the data-driven checks, as the efficiency computation itself is straight-forward.

6.2 Backgrounds

The backgrounds can be divided into two categories:

1. Candidates which are not from a genuine $\chi_c \to \gamma J/\psi$ decay
2. Non-CEP processes which produce the same decay as the signal ($\chi_c \to \gamma J/\psi$) along with other particles, which may not be detected

Backgrounds in category 1 can be distinguished from signal on a statistical basis by a fit to the mass spectrum of the candidates. The muons are usually from a genuine J/ψ, as there is little combinatorial background. The photons can be produced at the same proton-proton interaction as the signal decay or they can come from other interactions, noise or machine induced background. A relevant physical background is the production of $\psi(2S)$ and its decay to $J/\psi\pi^+\pi^-$ and $J/\psi\pi^0\pi^0$, when one of the pions is misidentified as a photon and the other pion is not detected. The J/ψ may be produced either inelastically or in CEP.

The second category of backgrounds is inelastic production of the same particles as the signal ($\chi_c \to \gamma J/\psi$). These backgrounds are processes which produce a χ_c and a few additional particles. When the extra tracks are not detected, the signature is similar to that of the signal process. The relevant inelastic processes are shown in fig. 6.1. They are due to “double pomeron exchange” just like the signal process. The cross sections for such processes are not well known, so the backgrounds have to be estimated using data-driven methods. The p_T of the central system in the signal process (i.e. CEP) is expected from theory to be small. The p_T-spectrum of inelastically produced particles is peaked at higher values than CEP. This can be understood as the signal candidate particle recoiling off the

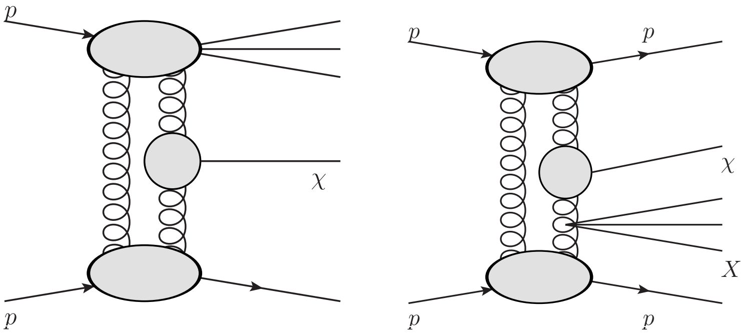

(a) Proton dissociation (b) Soft emissions

Figure 6.1: Inelastic backgrounds to CEP can produce the correct signal particle, but other particles are also produced.

additional particles. The decay $\psi(2S) \rightarrow \gamma\chi_c$ produces a final state which is similar to the CEP final state. The additional photon has low energy due to the small difference between the $\psi(2S)$ and the χ_c masses, and can thus easily escape detection. Based on the number of $\psi(2S) \rightarrow \mu^+\mu^-$ in data and the branching ratios, approximately 500 of these feed-down events are expected for the χ_{c1} and χ_{c2}, and about 20 for the χ_{c0}. This background is expected to be small, based on MC studies.

6.3 Simulation

Simulated CEP interactions from the SuperCHIC [58, 121] Monte Carlo generator are used in the analysis. The generator includes a simulation of the CEP process based on perturbative QCD (the Durham model). The muons and the photon are passed through the LHCb detector simulation using the Gauss application, which calls GEANT 4 [29] to simulate the passage of particles through the detector. The simulated events are then processed with the standard reconstruction software and an emulated trigger, so the simulation and real data can be compared directly. Samples of 4 million events were generated using SuperCHIC with the default options for the χ_{c0}, χ_{c1} and χ_{c2}. An acceptance cut was used at the generator level, requiring that the pseudorapidity was in the range $2.0 < \eta < 4.5$ for the photon and for each of the muons. This left 305193 χ_{c0}, 302974 χ_{c1} and 302640 χ_{c2} decays. Events were also generated with alternative configurations to assess the uncertainty due to the simulated p_T-distribution, as will be discussed in Section 6.7.2.

L0		HLT2	
L0 muon object	$p_T(\mu) > 200\,\mathrm{MeV}$	Muon cuts	Two long tracks
SPD multiplicity	$nSPD < 10$		$p_T(\mu) > 400\,\mathrm{MeV}$
			IsMuon

Stripping	
Backward tracks	$nBack = 0$

Table 6.1: Trigger and stripping selections.

6.4 Trigger

The selection requirements in the trigger and stripping are summarised in table 6.1. The triggers used to select CEP events at L0 are referred to as low-multiplicity triggers, and these triggers generally make decisions based on the logical AND of a signal selection and a low-multiplicity requirement. The trigger L0Muon,lowMult is used for dimuon final states, and this trigger requires that there is a track in the muon detectors with $p_T(\mu) > 200\,\mathrm{MeV}$ and that there are fewer than 10 hits in the SPD. There is approximately one SPD hit for each charged particle in the acceptance. Because the rate of low-multiplicity events is low and the events are fast to reconstruct and take up a small amount of storage, a pass-through trigger is used in HLT1, accepting any event that is accepted by the low multiplicity L0 triggers.

The CEP trigger at HLT2 is called Hlt2diPhotonDiMuon (named after the "diphoton" production process of Section 4.3.2, it does not require any photons). It reconstructs the muons using long tracks and requires hits in the muon detectors consistent with the "IsMuon" algorithm, defined in the following section. Long tracks are tracks with hits in the VELO and the tracking stations. The transverse momentum of each muon is required to be greater than $400\,\mathrm{MeV}$. There are also other independent triggers which are used to measure the efficiency.

An additional level of data reduction is the *stripping*, which is performed after the triggers. The stripping is similar to a fourth trigger level, but events which do not pass any stripping selection are not permanently discarded (though they are not available for the analysis). The dimuon CEP stripping selection applies a cut requiring no *backward tracks* in the event. Backward tracks are VELO tracks that point in the opposite direction to the main LHCb acceptance.

Muon momentum	Muon stations
$3\,\mathrm{GeV} < \mathrm{p} < 6\,\mathrm{GeV}$	M2 + M3
$6\,\mathrm{GeV} < \mathrm{p} < 10\,\mathrm{GeV}$	M2 + M3 + (M4 or M5)
$p > 10\,\mathrm{GeV}$	M2 + M3 + M4 + M5

Table 6.2: Required hits in muon detectors for muon identification.

6.5 Selection

The main selection consists of the signal decay selection and global event cuts. The signal decay selection attempts to select χ_c-decays. It is an ordinary cut-based analysis. The global event cuts are based on global properties of the event, such as the total number of tracks. The purpose of the global event cuts is to select events where only the signal particles are present, i.e. candidates for central exclusive production, and to reject the backgrounds of category 2 in Section 6.2.

6.5.1 Signal decay reconstruction and selection

The offline muon reconstruction starts from long tracks. The long tracks are matched with hits in the muon detectors within a region around the extrapolated position of the particle, referred to as a Field of Interest (FoI) [122]. The size of the FoI depends on the track momentum and the position in the muon detectors. The necessary number of muon stations with hits for a positive muon identification, referred to as IsMuon, is shown in table 6.2. Photon candidates are reconstructed from clusters in the ECAL [123]. Clusters are identified as photons if none of the tracks in the event are consistent with pointing to the cluster.

Candidate χ_c-decays are constructed from two muons of opposite charge and a photon, and then the following selection is applied:

- Muons and photons have pseudorapidity in the range 2.0–4.5
- Both muons have $p_T > 400\,\mathrm{MeV}$
- Dimuon mass in the range 3035–3165 MeV
- Transverse momentum of $\mu^+\mu^-$ less than $900\,\mathrm{MeV}$
- Photon $p_T > 300\,\mathrm{MeV}$

The cuts are used to ensure that the particles are in a kinematic range where the efficiency is well known, except for the cuts on $p_T(\mu\mu)$ and the dimuon mass, which are used to enhance the signal to background ratio. The dimuon mass spectrum for events accepted

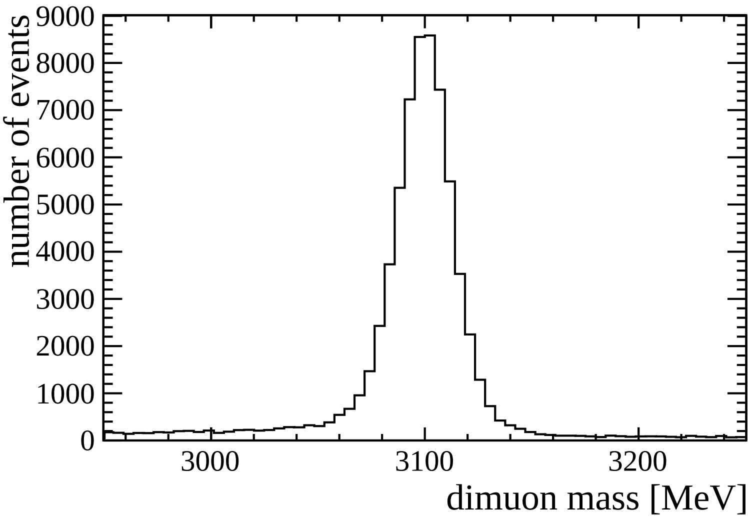

Figure 6.2: Dimuon mass spectrum for events accepted by the trigger, in the range near the J/ψ mass.

by the trigger and stripping, but before any selection cuts, is shown in fig. 6.2. Only events which also have a photon are included. Even at this stage, the J/ψ signal is very pure.

6.5.2 Global event cuts

The aim of the global event cuts is to select events in which only a signal interaction is present, and to reject a significant portion of the events of background category 2 in Section 6.2, i.e. inelastic production of χ_c. This is done by attempting to cut on the particle multiplicity in the event. The total number of particles in the event is not directly accessible, because LHCb has neither a complete angular coverage nor a perfect efficiency. The rejection of additional particles is done by searching for tracks, and we can only demand rapidity gaps with a limited acceptance and efficiency. The component of the selected candidates which are genuine CEP interactions can subsequently be determined in an analysis of the p_T-distribution, and this is described in Section 6.7. A helpful side-effect of the global event cuts is to limit the contributions from background processes in category 1, i.e. not genuine χ_c, which are accepted. For example, the combinatorial background is greatly reduced by requiring that no additional tracks are present in the event. An alternative method for rejecting additional particles will be discussed, using the VELO clusters instead of tracks, but it has some problems which prevents it from being usable.

VELO tracks and long tracks[1] are used for the analysis because they have low noise, and the VELO has a greater acceptance than any other subdetector. VELO tracks are

---

[1] Long tracks are counted separately from VELO-only tracks because the VELO tracks which are used as seeds for long tracks are removed from the track container in the reconstruction software.

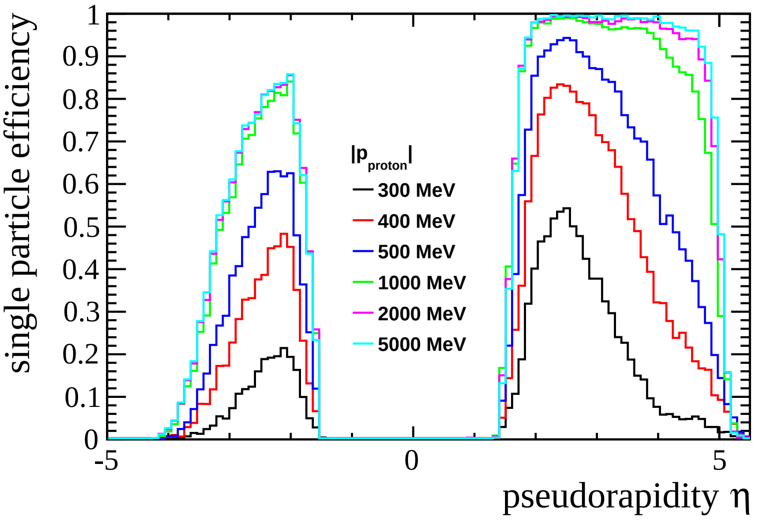

Figure 6.3: Efficiency for reconstructing a single proton originating from the interaction region, with the given momentum and pseudorapidity, as a VELO track or a long track. The sensitivity at negative pseudorapidity is due to reconstruction of backward tracks.

reconstructed using the FastVelo algorithm [124]. Events with exactly two long tracks, from the muons, and no (other) VELO tracks are accepted as CEP candidates. This cut effectively imposes two rapidity gaps: one in the backward direction, and one in the forward direction around the signal decay. The rapidity gap acceptance can be quantified by the probability to reconstruct a single particle at a given momentum and angle (this is different from the tracking efficiency for the signal, which is discussed in Section 6.8.1). The efficiency for reconstructing a VELO track or a long track from a single particle is shown for simulation in fig. 6.3, for six values of particle momentum. The graph is generated by simulating events with single protons which originate at the interaction point, and measuring the fraction of these protons which are reconstructed as tracks. The choice of particle to simulate is in practice arbitrary, as the tracking efficiency is approximately the same for all particles which may be expected to fill the rapidity gap, including protons, pions, electrons and muons. The single particle efficiency shown in fig. 6.3 is not used directly in the analysis, as the purity of the signal can be estimated by other means.

It was considered to use a more sensitive quantity based on the VELO cluster multiplicity instead of only cutting on the number of tracks. The number of *unassociated clusters* is defined as the number of clusters in the event, but with clusters on the tracks of the signal muons excluded from the count. The association of clusters to tracks is done for each cluster in turn. The point of intercept of each muon track with the sensor is computed, and the distance between the cluster and the nearest muon track in the coordinate of the sensor (r for R-sensors, ϕ for Phi-sensors) is computed. If the distance is within a configured tolerance, the cluster is associated to the track and it is not counted as an unassociated cluster. To decide whether an event is a CEP candidate or if it has addi-

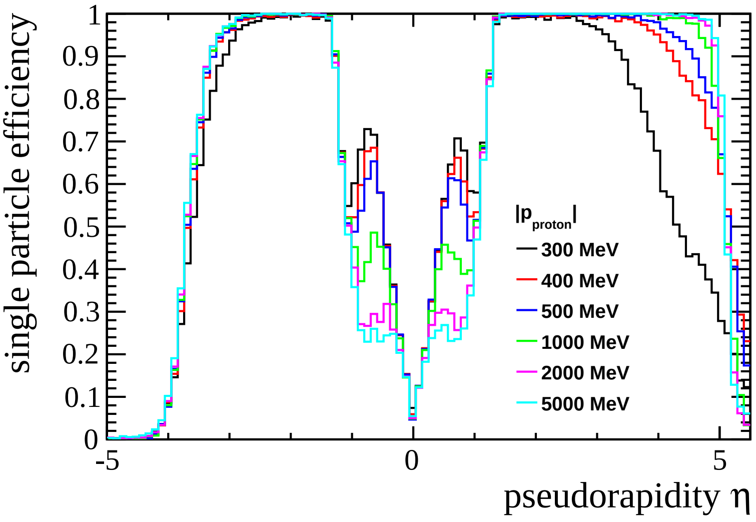

Figure 6.4: Efficiency to detect a single proton. The plot is analogous to fig. 6.3, but instead of requiring a track, events with three or more VELO clusters are accepted as detections. The curves correspond to different proton momenta.

tional particles, a cut on the number of unassociated clusters is used. A non-zero number of unassociated clusters must be allowed even in CEP candidates due to noise. The advantage of looking for unassociated clusters instead of tracks is an improved sensitivity to poorly reconstructed particles and an extended acceptance. The acceptance is greater than that for tracks because the particles do not need to pass through as many sensors as for tracks. A VELO track has to pass through at least three sensors of each type, i.e. six sensors in total. Figure 6.4 shows the probability to detect three or more clusters in simulated events which contain only a single proton, as a function of proton pseudorapidity, for multiple values of the momentum. This is the efficiency for a cut of three unassociated clusters, and it is significantly greater and more uniform than for tracks.

The distributions of the total number of clusters and the number of unassociated clusters in some datasets are shown in fig. 6.5. The number of tracks is also shown for comparison. Some noise may come from the signal decay itself, and this is seen when the number of unassociated clusters is greater than zero in MC (blue graph). The unassociated clusters in the MC are mainly produced by a failure to associate some clusters with the tracks. Other noise from the beams and the detector itself can be estimated from events where only a single beam passes through LHCb. Single-beam events are selected from the no-bias dataset (no signal candidate is present in these events, only noise), and the cluster multiplicity is shown with the red graph. The final dataset, shown in green, is χ_{c2} candidates in real data, without any cut on the number of tracks, but with a mass window and cuts on the SPD hit multiplicity and the number of backward tracks.

The problem is that there is no peak at zero unassociated clusters in the distribution for signal candidates. This is either due to an additional source of noise that is only present

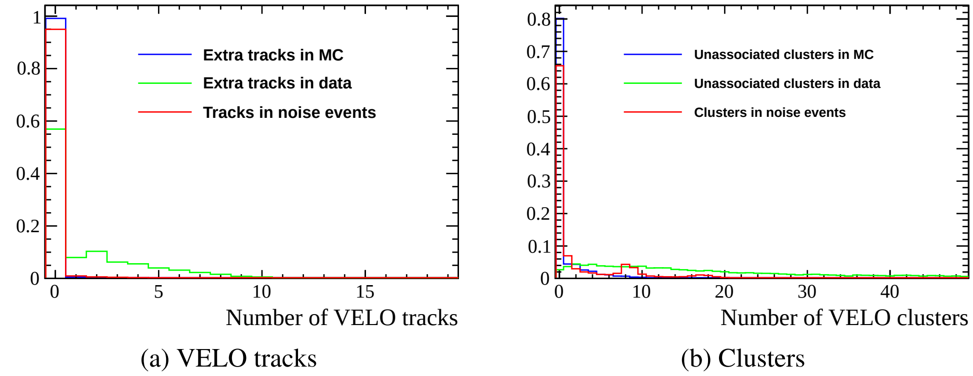

(a) VELO tracks

(b) Clusters

Figure 6.5: Response of the global event cut variables in MC events, χ_{c2} signal candidate events in data, and noise events without beam crossings. The global variables are the VELO track multiplicity (a), and the number of VELO clusters (b). For the first two datasets in (a), the number of extra tracks is shown. Correspondingly, the number of unassociated clusters is shown in figure (b).

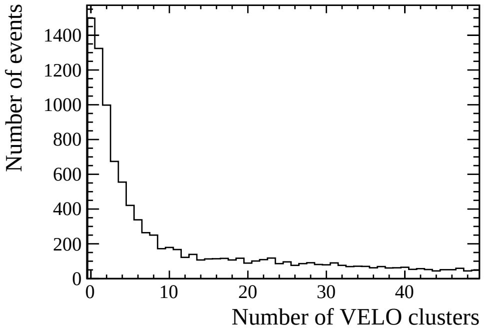

Figure 6.6: Unassociated clusters in χ_{c2} candidates after imposing quality requirements on the clusters. For comparison, the absolute number of χ_{c2} candidates with zero extra tracks is 7233 on the same dataset.

when there are collisions in the detector, or due to spillover. If the clusters are required to have ADC counts in a narrow range around the peak of the Landau distribution, and additional filtering is done to remove noise patterns, a peak can be restored at zero unassociated clusters. This is shown in fig. 6.6. The filtering of clusters cannot however be expected to remove all noise, and as the background models are seen to be inadequate for cluster-based analyses, the efficiency of the cluster-based global event cut on signal cannot be ascertained.

The global event cuts used in the final analysis are: no extra VELO or long tracks, and that the number of reconstructed photons is equal to one (i.e. no extra photons). The cut on the number of photons is used to have some sensitivity to neutral particles, and to reduce the combinatorial background from events with multiple photons. It does not significantly impact the rejection nor the efficiency. A cut on the number of photons was also used in Ref. [63], but then it was primarily used to reject χ_c.

6.6 Fit to the mass spectrum

The yield of each χ_c state is extracted using a fit to the mass spectrum of the candidates in real data. The signal is modelled with a Crystal Ball function [125], which is similar to a Gaussian, but has a power-law mass dependence below a certain threshold mass to account for electromagnetic final state radiation. While the mass distribution for the J/ψ is described well using physically motivated arguments, the photon introduces more complexity. The photon may interact with the detector material, producing an electron-positron pair in a process referred to as a photon conversion. If the conversion happens downstream of the magnet, the full energy of the photon is still measured, and the only effect is a slightly worse energy resolution than for unconverted photons. If the conversion happens before the magnet, the electron and positron are deflected by the magnetic field, and the calorimeter may detect either of these particles or nothing at all. The measured momentum will be wrong if this particle is detected as a photon, and the result is a continuum of candidates with wrong mass.

The mass distributions including the tails are shown for simulation of χ_{c0}, χ_{c1} and χ_{c2} in fig. 6.7, plotted on a log-scale. There is a continuum of candidates from early conversions which is almost constant in candidate mass, giving a tail both below and above the peaks. The mass distribution cuts off at approximately $3330\,\mathrm{MeV}$ because of the selection cut $p_T(\gamma) > 300\,\mathrm{MeV}$, so the tail is mainly visible at high mass. In addition to the conversions there is also a radiative tail at the low end, resulting from final state radiation from the J/ψ.

The mass distribution is modified by complex detector effects which cannot be assumed to be accurately simulated in MC. There is however a good agreement for the size of the main peak for decays involving photons in data and MC, as supported by an analysis of photons from B^+-decays (Section 6.8.4). The signal is modelled with a single Crystal Ball function with a tail at low mass. The high mass tail (continuum) is not included in the signal fit function, so it will instead be included in the background component of the fit. The same effect applies to the photon efficiency, so the calculation as a whole is consistent.

The fit is performed using the fitting software RooFit [126]. The signal is modelled as a sum of three Crystal Ball functions, one for each χ_c state. The functional form of the Crystal Ball function is:

$$f_{CB}(m;\alpha,n,m_0,\sigma) = N \times \begin{cases} \exp\left(-\frac{(m-m_0)^2}{2\sigma^2}\right) & \text{for } \frac{m-m_0}{\sigma} > -\alpha \\ A\left(B - \frac{m-m_0}{\sigma}\right)^{-n} & \text{for } \frac{m-m_0}{\sigma} \leq -\alpha \end{cases} \tag{6.6.1}$$

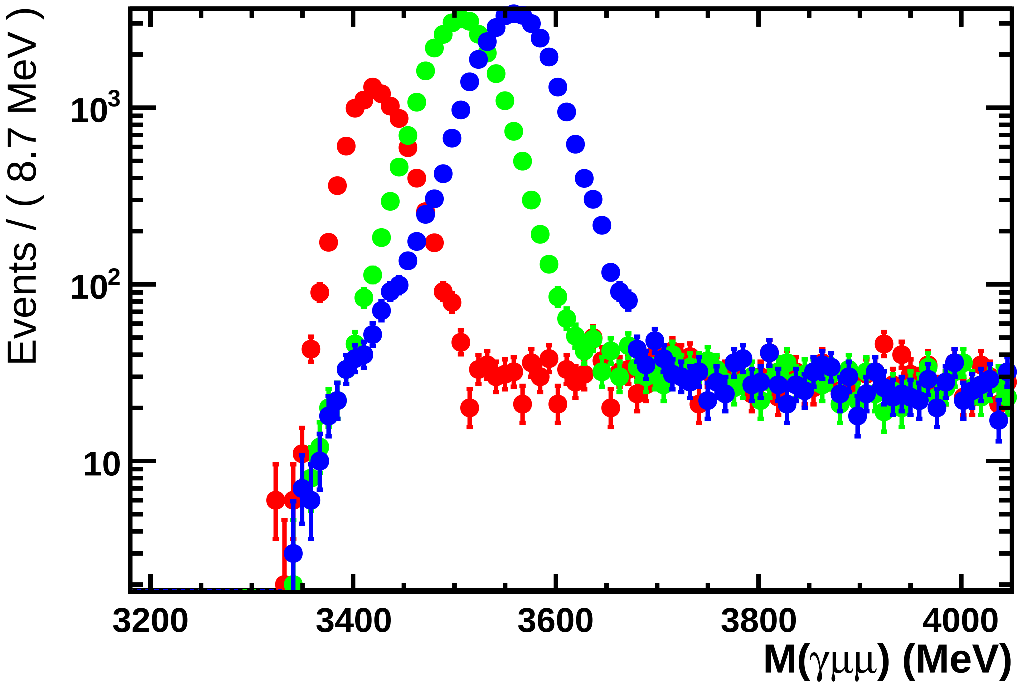

Figure 6.7: Mass distributions for reconstructed simulated events. The χ_{c0}, χ_{c1} and χ_{c2} are shown in red, green and blue respectively. The number of events is plotted on a logarithmic scale to show the non-Gaussian behaviour in the region away from the peak. The peaks correspond to an equal number of true decays, so the sizes of the peaks are proportional to the efficiency.

where

$$A = \left(\frac{n}{|\alpha|}\right)^n \exp\left(-\frac{|\alpha|^2}{2}\right) \qquad \text{and} \qquad B = \frac{n}{|\alpha|} - |\alpha|.$$

Its parameters are α and n, which determine the shape of the tail, σ which is analogous to σ for a Gaussian, and m_0 which is the peak position. The observable m is the candidate mass, and N is a normalisation constant. The available data are not sufficient to determine all the parameters of all three Crystal Ball functions, so the parameters α and n have to be taken from a fit to simulated events, separately for each χ_c state. The fits to MC are shown in fig. 6.8. A large value of $n > 60$ is found for all χ_c states, giving the function an approximately exponential dependence on the mass. This indicates that the tail at low mass is not just from final state radiation; there is also a contribution from poorly reconstructed photons.

The widths σ_0, σ_1 and σ_2 of the peaks, for χ_{c0}, χ_{c1} and χ_{c2} respectively, also cannot be determined from data alone. The relative widths of the mass peaks (σ_1/σ_0, σ_2/σ_0) are taken from simulation, but the overall scale (σ_0) is free. The differences between the peak positions are fixed to the differences between the PDG [51] values. A common offset

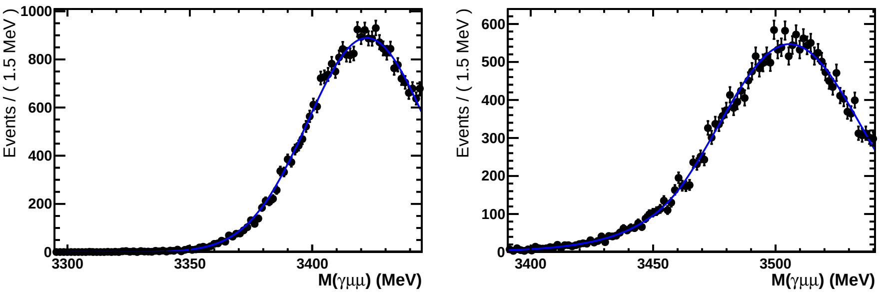

(a) χ_{c0} (b) χ_{c1}

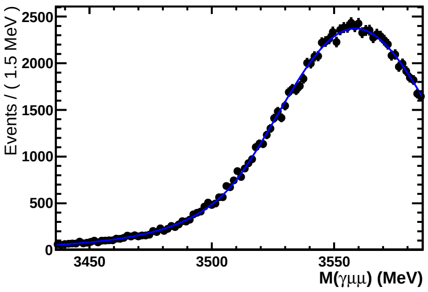

(c) χ_{c2}

Figure 6.8: Fits to mass spectra for simulated events, used to determine the tail parameters of the Crystal Ball functions. Only the lower tail is fitted. The normalisation is arbitrary.

which shifts the peaks to lower or higher mass values is included, to account for incorrect calibration of the calorimeter energy response.

Backgrounds which are not from χ_c, i.e. background from category 1 in Section 6.2, are modelled with a single function. The background is believed to mainly consist of feed-down from $\psi(2S)$ and from events in which J/ψ are produced directly and a photon is randomly associated to form a χ_c-candidate. The shape of the background is investigated using MC studies and control channels in real data, but it is not possible to significantly constrain the shape or the composition.

Two alternative functions are used to model the background: a straight line with an exponential cut-off, and the `RooDstD0BG`, which is a function designed to model a specific kind of combinatorial background. The former function is used as the default, and the `RooDstD0BG` is used only to estimate the uncertainty due to the background model.

The main background model is given by:

$$f_{bg}(m; b, m_{th}, C) = \begin{cases} \left(1 - e^{(m-m_{th})/C}\right)\left(1 + b\,\frac{m-m_{min}}{m_{max}-m_{min}}\right) & \text{if } m > m_{th} \\ 0 & \text{if } m \leq m_{th}\ , \end{cases} \tag{6.6.2}$$

where m is the χ_c candidate mass. The fitted parameters are the slope b of the linear factor, the threshold mass m_{th}, and the cut-off "softness" C which determines the shape of the cut-off. The parameters m_{min} and m_{max} are constants which are set to the low and high end of the fit range, and are not fitted.

The alternative background function `RooDstD0BG` was originally created to model the combinatorial background to the decay $D^{*0} \to D^0\pi$, in the mass difference variable. The difference between the χ_c candidate mass and the world average J/ψ mass, $\Delta m = m - m_\psi$, is used as the observable. The J/ψ mass used is also in good agreement with the J/ψ mass measured at LHCb.

$$
\begin{aligned}
&f_{bg}(\Delta m; A, B, C, \Delta m_{th}) \\
&= \begin{cases} \left(1 - \exp\left[\frac{\Delta m_{th} - \Delta m}{C}\right]\right) \times \left(\frac{\Delta m}{\Delta m_{th}}\right)^A + B\left(\frac{\Delta m}{\Delta m_{th}} - 1\right), & \text{if } \Delta m > \Delta m_{th} \\ 0, & \text{if } \Delta m \leq \Delta m_{th}\ . \end{cases}
\end{aligned}
$$

The parameter Δm_{th} is the mass difference corresponding to the threshold mass, below which the function evaluates to zero. The parameters A, B, C and Δm_{th} are left free in the fit. The mass region below $3350\,\mathrm{MeV}$ is not included in the fit, as we are not interested in the physics in this region, and it is not well described by the model.

The fit for all data, using the exponential-linear background, is shown in fig. 6.9 along with the pull distribution. The pull is the residual divided by the error:

$$
\text{pull} = \frac{N_{data} - N_{model}}{\sqrt{N_{data}}}\ ,
$$

where N_{model} is the expected number of events in a bin from the fit function and $\sqrt{N_{data}}$ is the statistical uncertainty on the number of events in a bin. The fit has most pull values distributed randomly in the range between -1 and 1, indicating a good fit.

6.6.1 Fits to the inelastic background

A fit to the mass spectrum of signal candidates was described in the previous section, using the selection described in Section 6.5. The cut on the number of extra tracks is now changed, keeping all other selection cuts as before (the number of extra tracks is the number of tracks in the event, not counting the two muon tracks from the signal candidate). This is done to investigate the composition of the inelastic background. The data are partitioned into bins of track multiplicity, identified by the number of extra tracks in addition to the two muon tracks. The events with zero extra tracks are the signal candidates that were already considered in the previous section. Events with additional tracks are considered to be mainly from inelastic interactions.

A simultaneous fit is performed to the mass spectra in bins of track multiplicity. The mass spectrum for each track multiplicity is fitted with a separate function, which is a sum of three Crystal Ball functions and the background function. The shapes and positions of

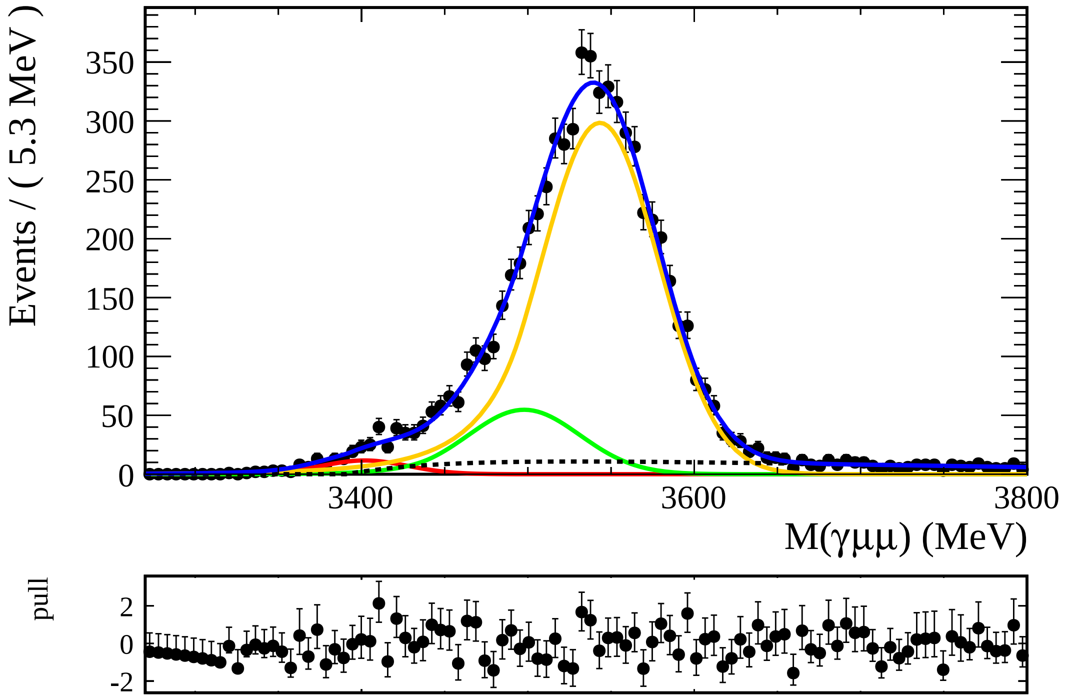

Figure 6.9: Fit to the mass spectrum, to determine the contribution of each χ_c-state in real data. The components of the fit from χ_{c0}, χ_{c1} and χ_{c2} are drawn with solid lines, in red, green and blue respectively. The background is drawn with a dotted black line. The sum of all contributions is drawn with a solid blue line.

the signal peaks, and the background parameters, are constrained to be the same for all track multiplicities. The signal and background yields are fitted individually for each bin. The mass spectra and the fitted functions for zero to five extra tracks are shown in fig. 6.10. The fit is done in the range 3270–3800 MeV. The yields of the signal species and background are shown versus the number of extra tracks in fig. 6.11 .

From the plots one can see that all signal yields are higher for events with zero extra tracks, and this is partially due to CEP. While it is not possible to deduce from the graph, there also appears to be more background with no extra tracks than with a few extra tracks. This can be explained by the background process producing a large number of events where the extra particles are not detected. The high background for events with two extra tracks could be from $\psi(2S)$-decays to a J/ψ and two pions (with a photon produced in a separate process), and MC studies confirm that $\psi(2S)$-decays produce a large number of events with two extra tracks.

6.6.2 Fit results and uncertainties

The parameters obtained from the simultaneous fit are shown in table 6.3, except for the yields in events with more than zero extra tracks (see fig. 6.11). The fit is investigated

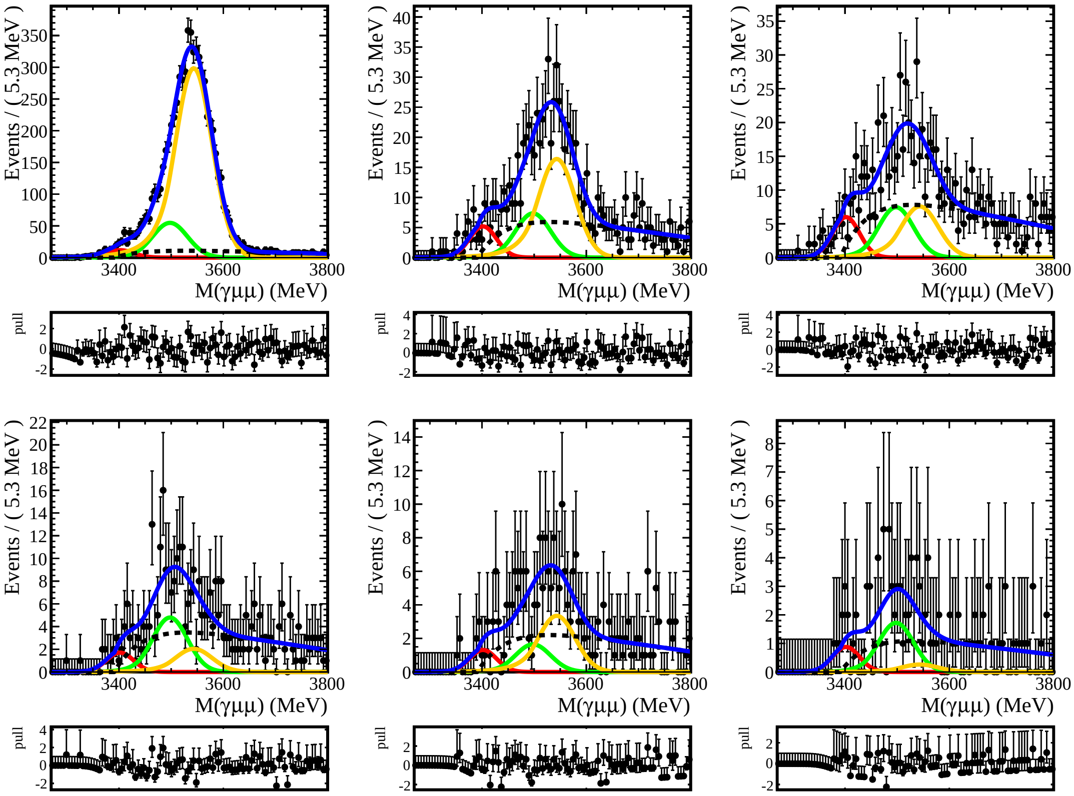

Figure 6.10: Fit to mass distribution for total track multiplicities from 2 (top left), 3 (top centre), to 7 (bottom right), i.e. from no extra tracks to five extra tracks. The same colours as in fig. 6.9 are used for the various components.

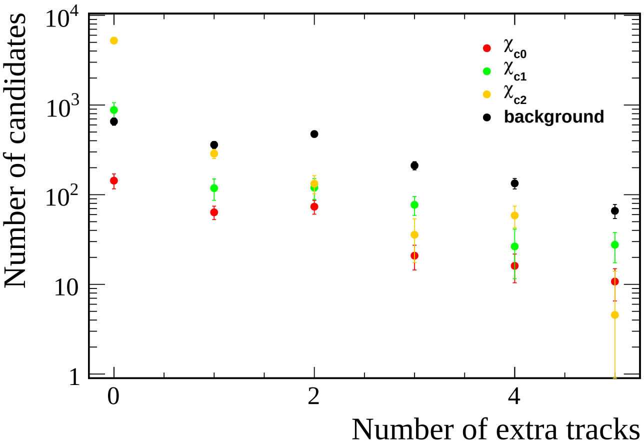

Figure 6.11: Signal and background yield as a function of track multiplicity.

Parameter	Symbol	Value
Background linear slope	b	$-0.663 \pm 0.081\,\mathrm{MeV}^{-1}$
Background threshold mass	m_{th}	$3391.0 \pm 8.3\,\mathrm{MeV}$
Cut-off softness	C	$60 \pm 25\,\mathrm{MeV}$
Signal peak width for χ_{c0}	σ_0	$25.91 \pm 0.53\,\mathrm{MeV}$
Signal peaks position offset	m_{off}	$-12.9 \pm 1.4\,\mathrm{MeV}$
Number of χ_{c0} candidates, tracks bin 0	n_0	143 ± 23
Number of χ_{c1} candidates, tracks bin 0	n_1	880 ± 170
Number of χ_{c2} candidates, tracks bin 0	n_2	5228 ± 186
Number of background candidates, tracks bin 0	n_{BG}	655 ± 60

Table 6.3: Parameters obtained from the simultaneous fit of six functions to the mass spectra corresponding to different track multiplicities. The functions share all parameters except for the numbers of signal candidates n_J and the number of background candidates n_{BG}.

	Comment	χ_{c0}	χ_{c1}	χ_{c2}	Backgd.	χ^2/ndf
1	Standard	143 ±23	880 ±170	5228 ±186	655 ±60	$495.2/571 = 0.87$
2	Fit only zero extra tracks	1.11	1.30	0.92	1.22	$65.3/91 = 0.72$
3	Gaussian signal peaks	1.73	1.73	0.87	0.87	$545.3/571 = 0.96$
4	All m_0 fixed to PDG masses	1.90	2.71	0.71	0.84	$544.6/572 = 0.95$
5	Free width, but $\sigma_0 \leq \sigma_1 \leq \sigma_2$	1.28	1.43	0.91	1.01	$493.7/569 = 0.87$
6	`RooDstD0BG` background pdf	1.00	1.49	0.93	0.92	$515.1/569 = 0.91$

Table 6.4: Signal yields for different fit functions and configurations. The "standard" fit is shown in the first row, and the following rows give the yield relative to this fit. The comment for each fit describes what is different from the standard fit (1).

in more detail by changing the model, with the results shown in table 6.4. The χ^2 per degree of freedom is also given, to indicate the agreement between the model and the data. The modifications are the following:

1. The yields for the "standard" fit are shown in the first row. This fit uses six bins of track multiplicity. The signal peaks are described by Crystal Ball functions with m_0 and σ_0 free, and with α and n fixed to the predictions from MC. The results of this fit are shown in table 6.3.

2. The yields do not change much when fitting events with two tracks only (not a simultaneous fit), because this bin contains the majority of events.

3. If the Crystal Ball functions are replaced with Gaussian functions there is no longer a tail from the χ_{c0} in the low mass region, and the background parameters are changed to compensate for this. A significant number of candidates are moved

from the background to the χ_{c0} peak. While the Crystal Ball function gives a much better fit than the Gaussian on data, the tails of the signal distributions may not be well described in the MC, so there is still an uncertainty in the signal description.

4. The peak positions are then fixed at the world average χ_c masses reported by the PDG [51], which have an uncertainty of $0.31\,\mathrm{MeV}$ for the χ_{c0}, and less than $0.1\,\mathrm{MeV}$ for the χ_{c1} and χ_{c2}. The signal yields change significantly for all χ_c states. The fit is visibly worse and has a higher χ^2/ndf than the standard fit. For comparison, the offsets seen when fitting each peak separately in simulation are $6.53 \pm 0.25\,\mathrm{MeV}$ for the χ_{c0}, $-5.57 \pm 0.16\,\mathrm{MeV}$ for the χ_{c1} and $1.93 \pm 0.24\,\mathrm{MeV}$ for the χ_{c2}. When fixing the peaks to the MC fit values instead of those from the PDG (not shown in the table), the yields change by less than $9\,\%$ relative to the standard fit.

5. When the width is left free, but requiring that the χ_{c0} width is the smallest and the χ_{c2} width is the greatest, the values are changed by at most $43\,\%$.

6. Using the `RooDstD0BG` function for the background results in a larger background contribution at low mass values, and changes the best fit positions of the mass peaks such that the χ_{c1} yield is increased. There are no compelling reasons to prefer either background function over the other, and the choice of background thus gives rise to a systematic uncertainty.

The substantial variations in table 6.4 show that there are large uncertainties on the signal yields, arising because the distributions in the mass variable are not precisely known. The χ_{c0} is highly correlated with the background in the fit, because the components are overlapping and of similar magnitude. An uncertainty of $\pm 90\,\%$ on the fit model is used for the χ_{c0} based on the maximum increase seen in table 6.4. The χ_{c0} yield only fluctuates upward for the alternative models, but there is no reason to rule out lower values than in the standard fit.

The χ_{c1} and χ_{c2} remain more stable, except for when the peak positions are fixed to the PDG values. The goodness of fit (χ^2) is worse, but it is not clear that the preference for a negative displacement of the positions is caused by a real displacement in the reconstruction – the background model could also be incorrect, shifting the fitted signal peaks. MC studies also indicate that the χ_c peak positions could be modified by different amounts, not just a common displacement. The changes in row 4 of table 6.4 are used as the uncertainty of the fit model, as the other changes are small in comparison. A symmetric uncertainty cannot be used for the χ_{c1} as the confidence interval would include negative numbers of events. It is instead assigned an uncertainty of ${}^{+170\,\%}_{-95\,\%}$. The upper limit is taken from row 4. The lower uncertainty is obtained by forcing the displacement m_{off} to $-20\,\mathrm{MeV}$, at which point the χ^2 is the same as when fixing the masses to the PDG values. For the χ_{c2}, the uncertainty on the fit model is set equal to the greatest variation from unity in table 6.4: $29\,\%$ or 1500 events. Because of the large correlation between the χ_{c1} and χ_{c2}, the upper bound on the χ_{c2} yield is limited by the number of χ_{c1}. An asymmetric uncertainty on the number of events, 5228^{+880}_{-1500}, is used for the χ_{c2} in the calculations when this is important.

6.7 Signal purity

After all cuts have been applied, the signal peaks in the fit to the mass spectrum correspond to a mixture of true CEP interactions and inelastic background (category 2 in Section 6.2). The *purity* is the fraction of the signal candidates that are actually from CEP. It is measured using a fit to the dimuon transverse momentum ($p_T(\mu\mu)$)-distribution[1], in which only the number of signal and background candidates are free parameters. The signal $p_T(\mu\mu)$-distribution is predicted by theory to peak at low values, 350–600 MeV. The inelastic background is not well known from theory, but one can argue that the χ_c can recoil against the additional particles, producing a higher p_T on average. A data-driven method is used to get an estimate for the p_T-distribution of the background.

6.7.1 Main analysis

The Novosibirsk function is used to fit the $p_T(\mu\mu)$ distributions for both the signal and inelastic background. The function is a generalisation of a Gaussian function, with an asymmetry parameter (or tail parameter) α which skews it to one side. Its definition is

$$\begin{aligned}&\mathcal{N}(p_T; p_{T_{pk}}, \sigma, \alpha) = \\ &\quad \exp\left\{-\frac{1}{2}\left(\left[\frac{\log(1+\alpha\frac{p_T - p_{T_{pk}}}{\sigma}\sinh(\alpha\sqrt{\log 4})/(\alpha\sqrt{\log 4}))}{\alpha}\right]^2 + \alpha^2\right)\right\}.\end{aligned}$$

The function used to model the signal is the sum of three Novosibirsk functions, with the parameters $p_{T_{pk}}$, σ and α fixed to those obtained in fits to MC. The fits to MC are shown in fig. 6.12. The Novosibirsk function does not describe the distribution perfectly, but it is acceptable to use it when fitting the data, as the available statistics in data are limited. The relative normalisation of the χ_{c0}, χ_{c1} and χ_{c2} components cannot be left to float in the fit to real data, because the data do not sufficiently constrain the parameters. The function used to fit the signal can be written as three terms made up of the Novosibirsk functions $\mathcal{N}_J(p_T)$, the number of total (CEP and inelastic) χ_{cJ} candidates n_J obtained from the fit

---

[1] The dimuon p_T is used because it has a slightly better resolution than the χ_c candidate p_T, but it is essentially the same.

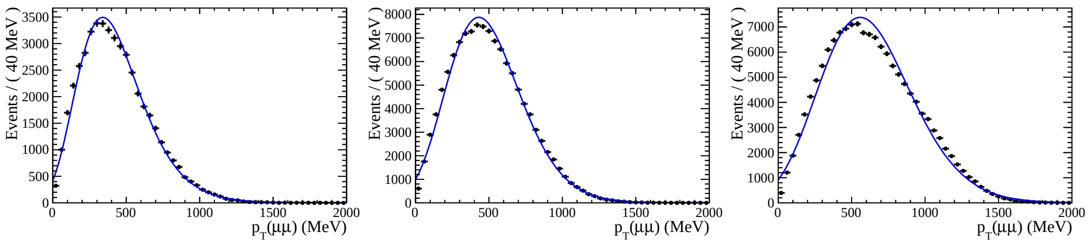

(a) χ_{c0}

(b) χ_{c1}

(c) χ_{c2}

Figure 6.12: Fits to dimuon transverse momentum distributions for simulated events.

of the mass spectrum, and the purity p:

$$f_s(p_T) = p[n_0\mathcal{N}_0(p_T) + n_1\mathcal{N}_1(p_T) + n_2\mathcal{N}_2(p_T)] \ . \tag{6.7.1}$$

Because p is a common factor, this assumes that the purity is the same for all signal particles.

The dataset is partitioned into bins of track multiplicity, labelled by the number of extra tracks in addition to the two muons from the signal decay. Bins with more than zero extra tracks contain almost only inelastic interactions, and are used to model the background of inelastic χ_c-production. The shape of the background p_T-distribution depends on the number of forward tracks, as shown in fig. 6.13a. All candidates within the mass range $3350\,\mathrm{MeV}$ to $3800\,\mathrm{MeV}$ are included, and the cut $p_T(\mu\mu) < 900\,\mathrm{MeV}$ is not applied when fitting the data. The parameters of the Novosibirsk functions are shown as functions of the number of extra tracks in figs. 6.13b–6.13d, and are fitted with straight lines. The values for the background in events with no extra tracks are obtained by extrapolation using the straight line fit, giving a peak position of $0.53 \pm 0.3\,\mathrm{GeV}$, a width of $0.48 \pm 0.02\,\mathrm{GeV}$ and an asymmetry parameter $\alpha = 0.59 \pm 0.36$. This extrapolation is justified based on general kinematical considerations, and by studying the dependence of other simulated processes on the number of tracks [63]. The purity is then estimated from a fit to the p_T-distribution in events with no extra tracks. The normalisations of the background and the signal are left free to vary in the fit. The fit is shown in fig. 6.14, and the fitted purity is then $23.8 \pm 3.3\,\%$. With the cut $p_T < 900\,\mathrm{MeV}$, the purity increases to $29.7 \pm 4.1\,\%$.

6.7.2 Uncertainties on the purity

The possible errors on the purity include the uncertainties on the extrapolated background, the simulated signal kinematics, and the use of the same purity for three χ_c-states.

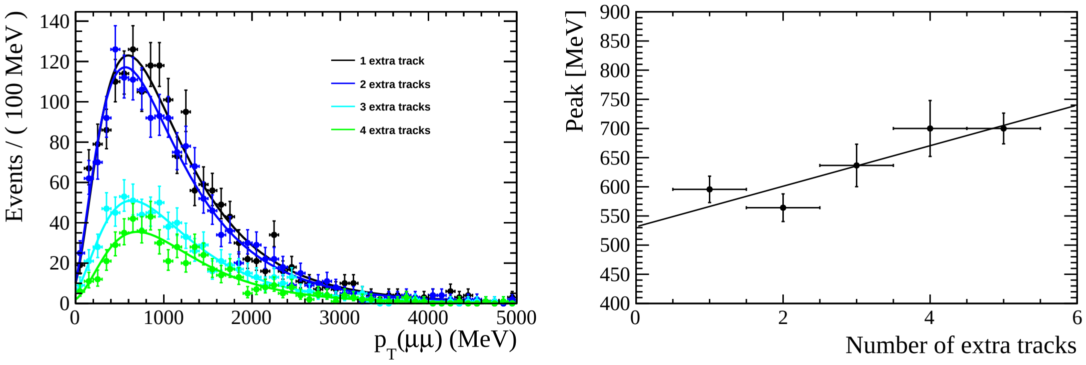

(a) p_T-distribution (b) Peak position

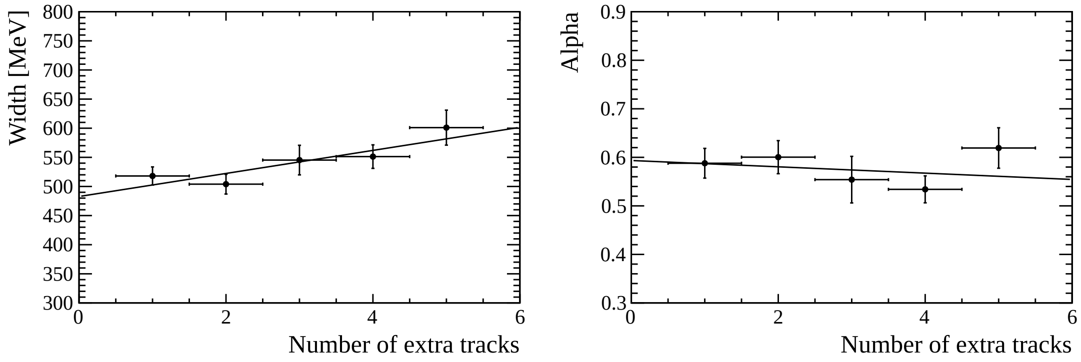

(c) Width (d) Tail parameter

Figure 6.13: Fits to transverse momentum distribution in real data. In (a), the transverse momentum distribution is shown for events with the number of extra tracks equal to one (black), two (blue), three (cyan) and four (green). Events with backward tracks are not included. In (b)–(d) the shape parameters are shown as functions of the number of extra tracks.

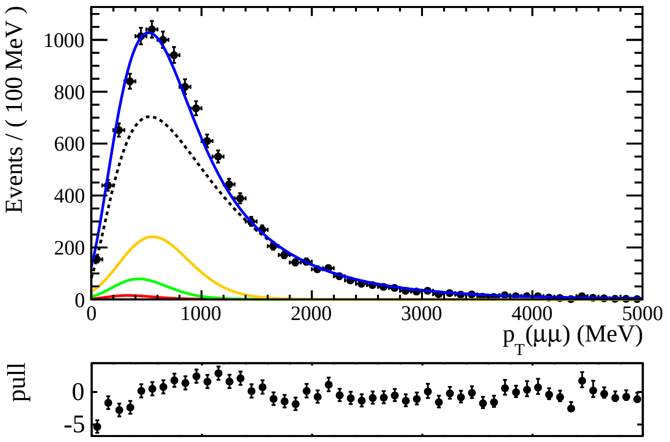

Figure 6.14: Main fit to transverse momentum distribution in data, for the determination of purity. The signal contribution is shown with solid lines in for each of χ_{c0} (red), χ_{c1} (green) and χ_{c2} (yellow). The background is drawn with a dashed black line.

Background extrapolation

An error may occur when tracks from pile-up and noise cause an incorrect classification of events, increasing the number of tracks in the event. This effect is assessed by repeating the background analysis with a modified track multiplicity parameter, counting only tracks with an impact parameter (IP) less than $800\,\mu\mathrm{m}$ between the track and the dimuon vertex. The IP is the distance of closest approach of the track to the vertex. It is peaked at a low value for tracks coming from the same vertex (i.e. tracks from the same proton-proton interaction), and it is large for tracks from other interactions or noise. When using this IP requirement, the purity for candidates with $p_T < 900\,\mathrm{MeV}$ changes to $34\,\%$, which is consistent with the result without an IP cut, and represents a $14\,\%$ relative change.

Simulation

The purity estimate depends on the simulation through the signal p_T-shape. The transverse momentum of the signal is largely controlled by a (QCD) proton form factor

$$F(t) = \exp(-b|t|/2)\ ,$$

where t is the four momentum transfer squared and b is a slope parameter. In SuperCHIC, $b = 4\,\mathrm{GeV}^{-2}$ is used by default [58]. The signal is measured in the forward region at LHCb, i.e. the central system has a p_z significantly different from zero. This means that the momentum transfer x from one proton to the χ_c is significantly larger than from the other proton, such that a region of high x and another region of low x are probed simultaneously. The value of b is not well constrained by other experiments for this kinematical configuration. In the LHCb analysis of exclusive J/ψ and $\psi(2S)$ production [63], a value of b found in an extrapolation based on Regge theory was found to give a better fit to the data than the default in SuperCHIC. To account for the non-zero transverse momentum of the final state protons (which are not observed at LHCb), b is replaced with an "effective" slope parameter b_{eff} in the SuperCHIC MC generator, as a step in the calculation. The values 6.2, 6.3 and 7.6 are used for b_{eff} for the χ_{c0}, χ_{c1} and χ_{c2} respectively. The extrapolations from Regge theory are 10, 11 and 13 for χ_{c0}, χ_{c1} and χ_{c2} respectively.

It is not attempted to extract b_{eff} from data, because there are too many unknowns in the model, for example that the purity of each χ_c-state may be different. The default b_{eff}-values in SuperCHIC are kept, and the Regge-based extrapolation is used as a second benchmark point. A sample of MC is generated with the modified values of b_{eff}, and the purity analysis is repeated using this MC. Half the difference in purity between these points, which is $15\,\%$ relative, is used as a systematic uncertainty on the purity to account for possible errors in the simulated $p_T(\mu\mu)$-distribution.

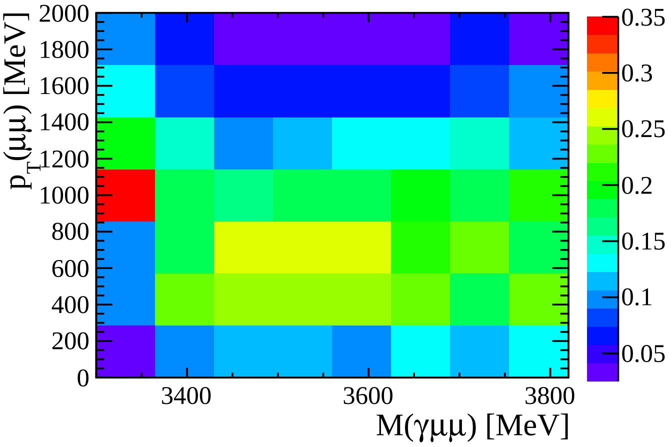

Figure 6.15: The plot shows the p_T-distribution (vertical axis) in bins of candidate mass (horizontal axis). The colour in each bin shows the number of events in that bin divided by the total number of events in that mass range, i.e. each column is normalised to unity.

6.7.3 Analysis using sPlot

It is useful to separate out the $p_T(\mu\mu)$-distribution for each signal species to test the assumption that the purity is the same for each χ_c state. It is not possible to use a cut-based approach because the χ_{c0} peak is of comparable size to the background in that region, and the χ_{c1} and χ_{c2} peaks overlap. The sPlot formalism [127] is used to extract a $p_T(\mu\mu)$-distribution for each χ_c-state based on the fit to the mass spectrum. The values of the probability density functions (pdfs) are used to assign a weight to each event

$$w_n(y) = \frac{\sum_{j=1}^{N_s} \boldsymbol{V}_{nj} f_j(y)}{\sum_{k=1}^{N_s} N_k f_k(y)} , \tag{6.7.2}$$

where $y = (M, \ldots)$ is the vector of observables in an event, $\boldsymbol{V}_{nj}$ is the covariance matrix of the signal yields and $f_j(y)$ are the pdfs. The indices j, k and n run over all the $N_s = 4$ signal and background species. In this analysis, there is only one observable which is used for fitting, the candidate mass, so $y = m$. The distribution in another "control variable" can then be reconstructed by weighting events by $w_n(y)$, and this is used to get the $p_T(\mu\mu)$-distributions for pure samples of χ_{c0}, χ_{c1} and χ_{c2}.

The sPlot method only works when the signal pdfs correctly describe the data and the discriminating variable (mass) is uncorrelated with the control variable $p_T(\mu\mu)$ for each species. The latter requirement is satisfied for the signal peaks, but it is not clear that the background has the same p_T-distribution regardless of mass (in fact, one may expect it to vary significantly). Figure 6.15 shows how the events are distributed in p_T, in bins of mass. The bins at low and high mass correspond to an almost pure background. There are events at low mass with high $p_T(\mu\mu)$, which are presumably from $\psi(2S)$ feed-down

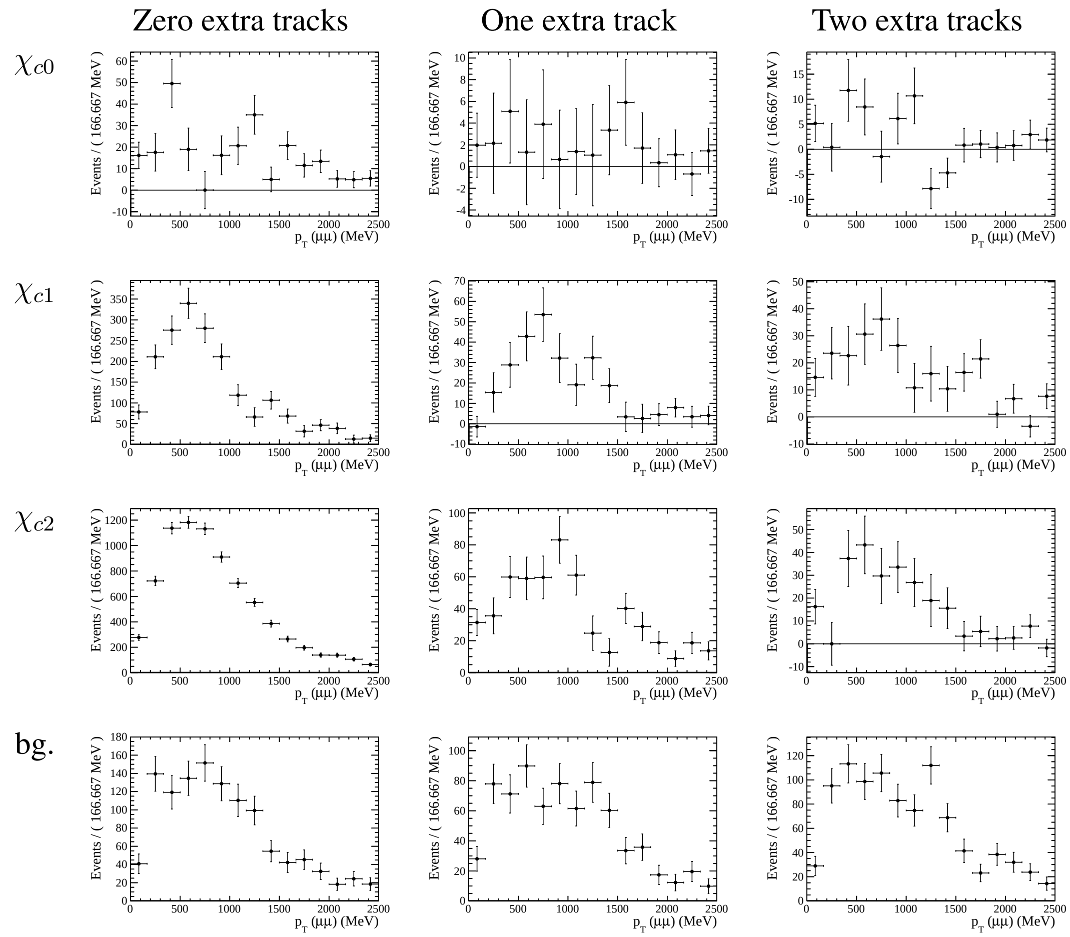

Figure 6.16: Plots of dimuon transverse momentum distributions. The top row shows the χ_{c0}, the second row χ_{c1}, the third row χ_{c3} and the bottom row shows the background component from the fit to the mass spectrum. The left column corresponds to events with no extra tracks, then one track in the middle and two tracks in the right column.

background. The mass range for the fit is therefore limited to $3380–3800\,\mathrm{MeV}$ for the sPlot analysis. A pure background is only visible in the high-mass region, but there it appears to be approximately constant as a function of mass. In an attempt to include the high-mass tail, to estimate its effect on the results, each Crystal Ball function was replaced with a sum of two Crystal Ball functions, with one tail at low mass and one at high mass. This did not affect the sPlots in any other way than to reduce the number of background events (all features of the plots remained, and this would only give a higher efficiency, which would be compensated by a revised efficiency in the MC), so a single Crystal Ball is used in the following. The resulting sPlots are shown for events with zero, one and two extra tracks in fig. 6.16. Some features are worth highlighting:

- The spectrum for χ_{c0} with zero extra tracks has candidates with high $p_T(\mu\mu)$. This may be due to mis-reconstructed $\psi(2S)$ decays which are included in the χ_{c0} peak.
- The χ_{c1} distribution is peaked at a higher value than expected, so it may have a

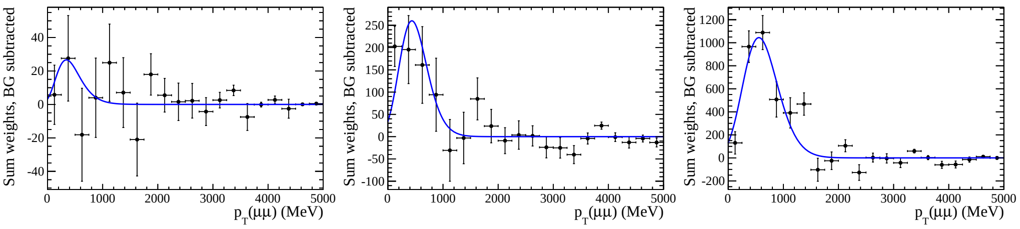

(a) χ_{c0} (b) χ_{c1} (c) χ_{c2}

Figure 6.17: Fits to background-subtracted sPlot data. The data points show data with the background subtraction procedure described in the text. The graph shows the fit of a Novosibirsk function with the parameters fixed to Monte Carlo predictions and the normalisation as a free parameter.

significant inelastic component. The MC has a peak at $450\,\mathrm{MeV}$. It may also be that the fit is selecting χ_{c2} as χ_{c1}, i.e. that the χ_{c1} yield is overestimated.

- The χ_{c2} has a peak at a position consistent with the SuperCHIC value of $560\,\mathrm{MeV}$.

A similar analysis to the one already described is done for each χ_c state separately, by extracting their p_T distribution using the weights from sPlot (sWeights). The Novosibirsk function cannot be fitted well to the inelastic background datasets (events with more than zero extra tracks). The data with a single extra track are instead used directly as a model for the background in events with zero extra tracks, for background subtraction. The data with zero tracks can be considered to have two components, the CEP signal and the inelastic background (the purity is the ratio of these). The data with one extra track are assumed to contain only background.

The normalisations of the background components in the bins with zero and one tracks are estimated by counting the number of events with $p_T(\mu\mu) > 1300\,\mathrm{MeV}$. The fraction of CEP signal events which appear in this range is expected to be less than $3\,\%$, based on the SuperCHIC simulation. If these event counts (sum of weights) are labelled $N_0^{>1300}$ and $N_1^{>1300}$ for events with zero and one tracks respectively, the background-subtracted dataset is obtained by scaling the weights of the dataset with one track by a factor $(N_0^{>1300}/N_1^{>1300})$, and then subtracting it from the dataset with zero tracks. The background-subtracted dataset is fitted with a Novosibirsk function in an extended likelihood fit, with all parameters fixed to the SuperCHIC values except for the number of signal events. The fits are shown in fig. 6.17[1]. The purity is then simply the number of candidates in the fit divided by the number of events in the full dataset with zero extra tracks (not the background subtracted dataset). Only candidates with $p_T(\mu\mu) < 900\,\mathrm{MeV}$

---

[1] The numbers of signal events divided by the purities are not exactly equivalent to the numbers of candidates in the main fit to the mass spectrum, because the fit is slightly different when removing the cut on $p_T(\mu\mu)$. It is necessary to fit the mass spectrum without the p_T-cut in order to obtain sWeights for all candidates.

are counted, because these are the data relevant to the final cross section computation. The numbers of candidates resulting from the fits for χ_{c0}, χ_{c1} and χ_{c2} are 59 ± 63, 648 ± 176 and 3263 ± 365 respectively. The corresponding purities are 0.37 ± 0.40, 0.35 ± 0.10 and 0.41 ± 0.05.

The results are consistent with all χ_c states having the same purity, and can be used to constrain the possibility that they are different. The large uncertainty on the χ_{c0} is caused by the low statistical precision, and the purity cannot be constrained to better than $\pm 100\,\%$ (relative) in relation to the others. The χ_{c1} and χ_{c2} purities can be constrained to be within $10\,\%$ of each other. There is a minor discrepancy between the sPlots, which yield a purity of about $40\,\%$, and the main analysis which gives $30\,\%$. This could be caused by the different models used for the inelastic background. This is investigated by repeating the main analysis using a background function fitted to data with one extra track, instead of doing the extrapolation. The background is then modelled in approximately the same way as for the sPlots. The purity for $p_T(\mu\mu) < 900\,\mathrm{MeV}$ is then $44\,\%$, consistent with the sPlots.

Because the data are dominated by the χ_{c2} signal, its purity is better constrained by the overall fit than by the sPlots. The main fit function (6.7.1) is replaced with a function where the χ_{c0} and χ_{c1} purities are different from the χ_{c2} purity,

$$f'_s(p_T) = (1+\alpha)p[n_0\mathcal{N}_0(p_T) + n_1\mathcal{N}_1(p_T)] + p[n_2\mathcal{N}_2(p_T)] \ .$$

If α is set to -0.1 or 0.1, the fit returns a value of p which is changed by less than $5\,\%$ relative to the original purity. The goal of this section is to assign an uncertainty related to the assumption that the purities for the three χ_c states are the same. The above result indicates that the combined purity is close to the χ_{c2} purity even if the χ_{c0} and χ_{c1} purities are different, so the uncertainty on the χ_{c2} should only be $5\,\%$.

6.7.4 Conclusion on signal purity

The value for the purity found in the main analysis is used for all χ_c-states. The systematic uncertainties which affect all χ_c-states in the same way, and thus cancel when taking the cross section ratios, are the simulation uncertainty of $15\,\%$ (relative) and the background uncertainty of $14\,\%$ (relative). The possibility of having different purities is estimated from the sPlots, giving relative uncertainties of $100\,\%$, $10\,\%$ and $5\,\%$ for χ_{c0}, χ_{c1} and χ_{c2} respectively. The purity has a value of 29.7 % (for $p_T(\mu\mu) < 900\,\mathrm{MeV}$), and the absolute uncertainties on the purities are $30\,\%$, $7\,\%$ and 6.3 %, for χ_{c0}, χ_{c1} and χ_{c2}.

6.8 Efficiency

The total efficiency can be written as a product of the trigger efficiency, the muon and photon reconstruction efficiencies and the selection efficiency:

$$\epsilon = \epsilon_{trg}\epsilon_{rec}\epsilon_{sel} = \frac{N_{trg}}{N_{total}} \times \frac{N_{rec\&trg}}{N_{trg}} \times \frac{N_{sel\&rec\&trg}}{N_{rec\&trg}} ,$$

where N_x is the number of accepted candidates at each stage. The simulation is used to get an estimate for the efficiency, and it is checked wherever possible using data-driven methods.

The reconstruction efficiency is assumed to factorise into independent factors for each particle,

$$\epsilon_{rec} = (\epsilon_\mu)^2\epsilon_\gamma .$$

The muon efficiency ϵ_μ is the probability to reconstruct a true muon and identify it as a muon. Muons are first reconstructed as long tracks in the tracking detectors, and then they are identified as muons by matching the tracks with hits in the muon detectors. The muon efficiency can thus be written as the product of the tracking efficiency and the muon identification efficiency.

6.8.1 Muon tracking efficiency

The tracking efficiency is computed using the tag and probe method using real data, as a check of the MC. A reconstructed muon which passes the HLT2 single muon trigger is used as a “tag”. Other muon candidates in the event are reconstructed using only information from the muon detectors and the TT, and the candidates with opposite charge to the tag are used as “probes”. The mass of each combination of the tag muon and a probe muon is plotted, and the J/ψ mass peak is fitted. Then the Muon-TT (probe) tracks are associated with long tracks by requiring that the Muon-TT track and the long track have at least three hits in common in the muon detectors. The dimuon mass is plotted again only for the probe tracks that have an associated long track. The tracking efficiency ϵ_{trk} is the ratio of the intensity of the J/ψ peak in the second case to that in the first case. This represents the fraction of the basic MuonTT objects which are reconstructed as long tracks.

The data are divided into five bins of probe p_T, and the efficiency is computed separately in each bin, and separately for positive and negative muon charge. Fits to the mass of the J/ψ candidate for a single bin, for negative probe charge, are shown in fig. 6.18.

The plots in fig. 6.19 show a comparison between the tracking efficiency in MC and in real data. The overall tracking efficiency is obtained by taking the mean of the efficiency in each bin weighted by the number of true signal decays in the bin. The efficiency from

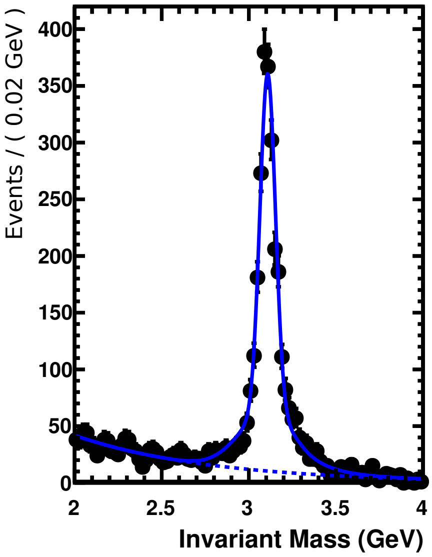

Figure 6.18: Fit to the dimuon mass for tracking efficiency in the transverse momentum bin $1.2\,\mathrm{GeV} < \mathrm{p_T}(\mu_{\mathrm{probe}}) < 1.6\,\mathrm{GeV}$, for negatively charged probe muons. The left plot shows the mass for all combinations of tag and probe muons. Only muons which have an associated long track are shown in the right plot.

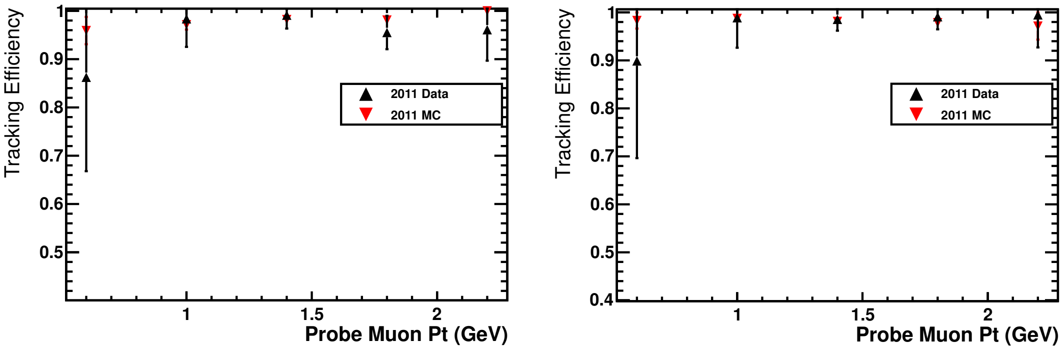

(a) μ^-

(b) μ^+

Figure 6.19: Comparison of tracking efficiency in data and MC for μ^+ and μ^-. Less than $1\,\%$ of the data are in the first and the last bin, so the large errors in these bins are not a problem.

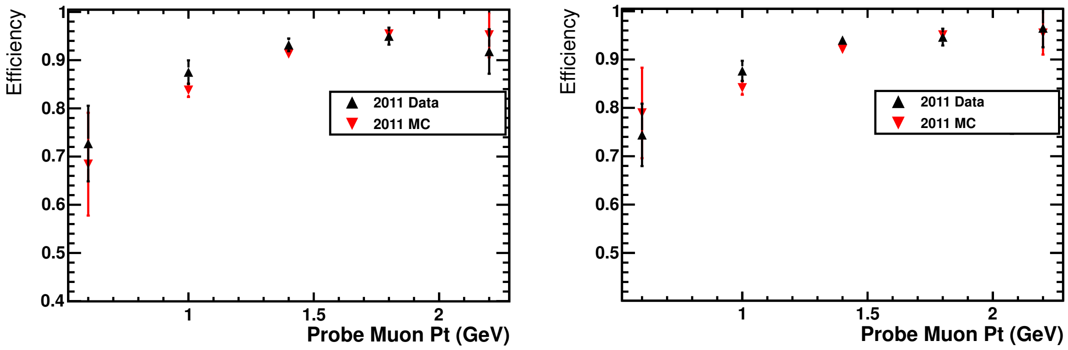

(a) μ^- (b) μ^+

Figure 6.20: Muon identification efficiency in bins of muon p_T, as determined in data using the tag and probe method.

tag and probe is $98 \pm 3\,\%$, and i n MC it is $98.1 \pm 0.3\,\%$, so there is a good agreement. The single muon trigger used in the tag and probe analysis has a cut requiring less than 10 SPD hits, no backward tracks and less than 4 forward VELO tracks. The events used to compute the efficiency are thus similar to the CEP candidate events.

6.8.2 Muon identification efficiency

Muon-ID is the determination of which tracks are from muons, and corresponds to the IsMuon cut in the selection. The muon identification efficiency is measured using a tag and probe method. Events accepted by the same single-muon trigger as for the tracking efficiency analysis are analysed, and the muon which triggered the event is used as a tag. Other long tracks in the event with opposite sign to the tag muon are used as probe muons, and are combined with the tag muon to get J/ψ candidates. Then a muon-ID cut is applied to the probes, and the muon ID efficiency is the ratio of the number of J/ψ with IsMuon on the probe to the number of J/ψ without a cut on IsMuon. A comparison of the muon ID efficiency in data and MC is shown in fig. 6.20. A combined muon identification efficiency of $93 \pm 2\,\%$ is found using tag and probe on data. This is consistent with the efficiencies of 91.8 ± 0.4 and 91.3 ± 0.4 found in the MC for positive and negative muons.

6.8.3 Trigger efficiency

The efficiency of L0Muon,lowMult is the product of the muon trigger efficiency and the efficiency of the SPD multiplicity cut when applied to signal interactions. The effect of multiple proton-proton interactions on the global event cuts, including the SPD cut, will be discussed in the context of the effective luminosity in Section 8.1. It is assumed for now that there are only single interaction events. A separate dimuon trigger without

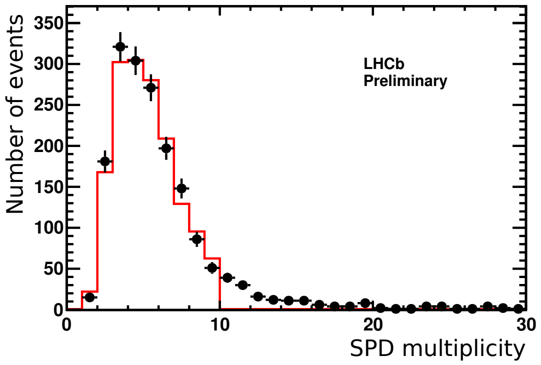

Figure 6.21: SPD hit multiplicity in signal candidate events. The data points in black correspond to events selected by a no-bias trigger, which also are accepted by the offline selection. The red line corresponds to events accepted by the normal low-multiplicity trigger and the offline selection.

an SPD requirement was used to analyse the effect of the SPD multiplicity requirement on signal interactions. The SPD multiplicity distribution is plotted in fig. 6.21 for central exclusive J/ψ production candidates accepted by this L0 trigger and the offline selection. The distribution is also plotted for signal candidates using the normal triggers including the SPD cut, and this distribution is normalised in the range of 0 to 10 SPD hits. The two distributions show a good agreement, so the data from the SPD-unbiased trigger can be used as a model for the signal. The efficiency of the SPD cut is then the number of events with fewer than 10 hits relative to the total number of events, which is $90 \pm 1\,\%$. This is in good agreement with the MC.

The efficiency of the muon triggers is evaluated using a tag and probe method: for the efficiency measurement, the single-muon trigger Hlt2LowMultMuon is used to select events with at least one muon. The muon which triggered the event is used as the tag muon. The LHCb trigger saves a log of what information was used by each specific trigger line. A reconstructed particle is classified as Triggered On Signal (TOS) with respect to a trigger line if detector hits from this particle were sufficient to satisfy the trigger line requirement. The tag muon is required to be TOS with respect to the single-muon trigger, so that no other particles are necessary to trigger the event. Other muons in the event with the opposite sign are reconstructed offline. The tag and the probe muons are combined, and the number of J/ψ is determined using a fit to the dimuon mass spectrum. To get the L0Muon,lowMult trigger efficiency the J/ψ yield is computed when requiring that the probe muon is TOS with respect to L0Muon,lowMult. The efficiency is computed in bins of p_T and separately for positive and negative muons. The overall L0Muon,lowMult muon trigger efficiency (ignoring the SPD cut) is $70 \pm 1\,\%$ per muon. The efficiency is plotted as a function of muon p_T in fig. 6.22. The efficiency for triggering on at least one of the two muons is estimated to be $91 \pm 2\,\%$, which is in good agreement with the

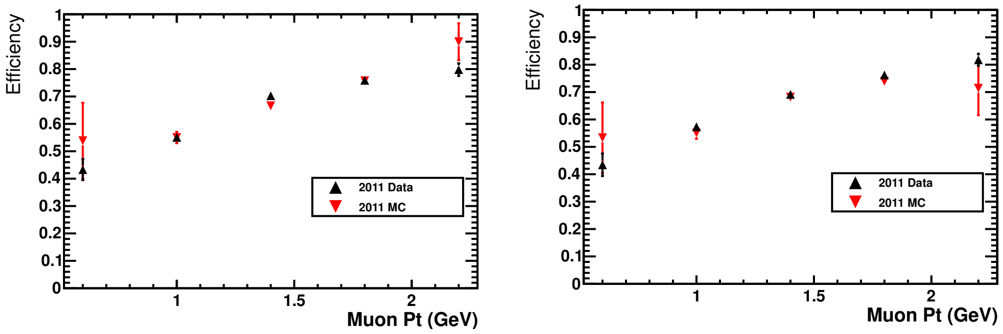

(a) μ^- (b) μ^+

Figure 6.22: L0 efficiency (L0Muon,lowMult) measured using the tag and probe method, for negative and positive muons.

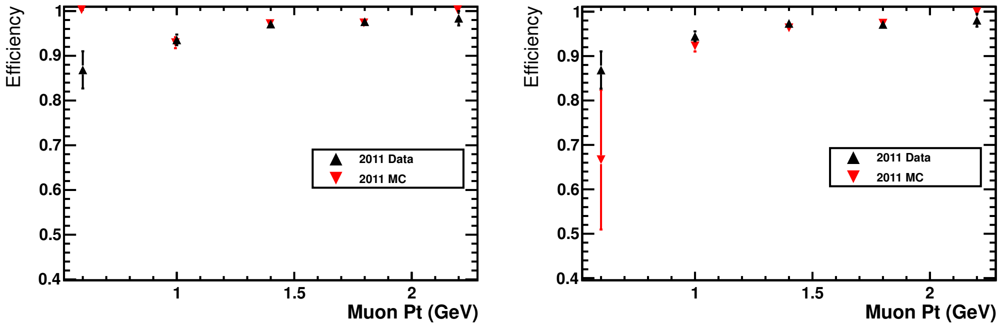

(a) μ^- (b) μ^+

Figure 6.23: Single muon HLT efficiency. Hlt2diPhotonDiMuon requires that both muons are triggered.

corresponding efficiency $90.2 \pm 0.5\,\%$ in the simulation.

The Hlt2LowMultMuon efficiency is needed as an intermediate step to compute the Hlt2diPhotonDiMuon efficiency. It is found using muons triggered by Hlt2LowMultMuon as tag muons and checking what fraction of the probe muons pass the same trigger. The HLT efficiency is plotted in fig. 6.23. An efficiency of $97 \pm 1\,\%$ is found for both positive and negative muons. The selection criterion for the signal trigger Hlt2diPhotonDiMuon is the logical AND of the selection using Hlt2LowMultMuon for the positive and the negative muon[1]. The final selection efficiency for Hlt2diPhotonDiMuon $95 \pm 4\,\%$.

---

[1] The reason that Hlt2LowMultMuon cannot be used directly to select signal decays is that it is prescaled. This means that only a fixed percentage of all events are allowed to be processed by the trigger.

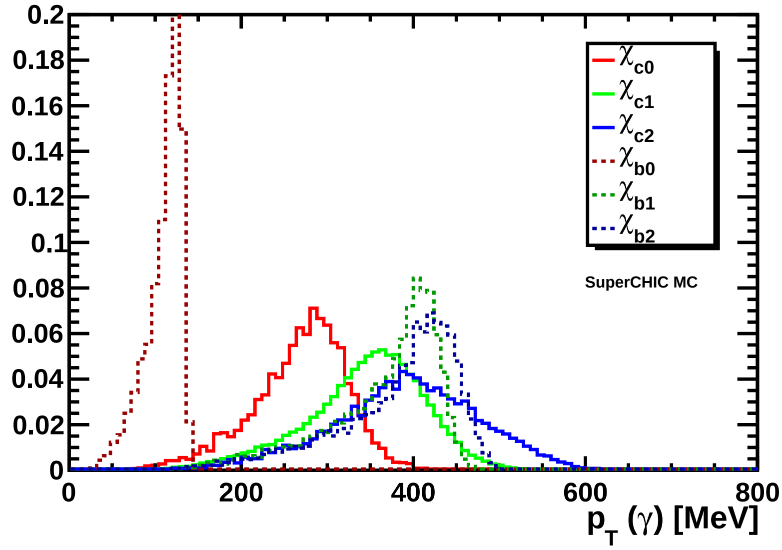

Figure 6.24: Photon transverse momentum distribution directly from the MC truth. Distributions for the χ_c states are shown with solid lines and the χ_b states are shown with dotted lines.

6.8.4 Photon efficiency

The photon efficiency is parametrised as a function of the photon transverse momentum (or transverse energy), because this parameter differs significantly between the χ_c states. The simulated photon transverse momentum spectra for decays of CEP charmonium and bottomonium are shown in fig. 6.24, without any detector simulation or selection. The transverse momentum spectra of the photons are determined by the difference of the χ_c or χ_b masses and the J/ψ or $\Upsilon(1S)$ mass, and by the dynamics of the CEP process. The photons from χ_{b0} decays have a much lower momentum than those from χ_{b1} and χ_{b2} because the χ_{b0} has total spin zero, which means that it can be produced with forward protons, i.e. when the protons scatter at zero angle [58]. With a cut at $300\,\mathrm{MeV}$, the efficiency for χ_{b0} is very low. Several avenues for checking the photon efficiency were explored, but the one giving the best precision was the application of a result by the LHCb calorimetry group.

Efficiency using $B^+ \to J/\psi(K^{*+} \to K^+\pi^0)$ and $B^+ \to J/\psi K^+$

The analysis [128] relies on the reconstruction of the following decays and their charge conjugates in the 2011 dataset:

1. $B^+ \to J/\psi K^{*+}$, with $K^{*+} \to K^+\pi^0$ and $\pi^0 \to \gamma\gamma$
2. $B^+ \to J/\psi K^+$.

$p_T(\gamma)$[MeV]	correction factor
300–350	$0.889 \pm 0.040 \pm 0.022 \pm 0.029$
350–400	$0.973 \pm 0.036 \pm 0.029 \pm 0.032$
400–500	$0.933 \pm 0.021 \pm 0.022 \pm 0.030$
500–600	$0.941 \pm 0.020 \pm 0.028 \pm 0.030$
600–700	$0.980 \pm 0.023 \pm 0.038 \pm 0.032$
700–1000	$1.001 \pm 0.016 \pm 0.028 \pm 0.032$

Table 6.5: Correction factors η_γ^{corr} for the photon efficiency, i.e. the ratio of the efficiency in real data to that in MC. The first error is the statistical error on data. The second error is a systematic error due to the limited MC statistics and the fit to the B^+ mass. The final error is due to the uncertainty on the B^+-decay branching ratios. Data in table are from Ref. [128].

The branching ratios of all these decays are well measured. The neutral pion efficiency can be expressed as a double ratio involving the number of observed events N_1 and N_2 for decay modes (1) and (2), the integrated luminosity $\mathcal{L}$, the B^+ production cross section σ_{B^+}, the kaon and J/ψ efficiencies for each decay mode $\epsilon_{J/\psi;K,1}$ and $\epsilon_{J/\psi,K;2}$, and the branching ratios:

$$\epsilon_{\pi^0} = \frac{N_1/\left[\epsilon_{J/\psi,K;1}\mathcal{L}\sigma_{B^+} \times Br(B^+ \to J/\psi K^{*+}) \times Br(K^{*+} \to K^+\pi^0)\right]}{N_2/\left[\epsilon_{J/\psi,K;2}\mathcal{L}\sigma_{B^+} \times Br(B^+ \to J/\psi K^+)\right]} . \tag{6.8.1}$$

The integrated luminosity and the cross section cancel in the ratio. The kaon and J/ψ efficiencies ($\epsilon_{J/\psi,K;1}$, $\epsilon_{J/\psi,K;2}$) are similar, but not identical, because the momentum spectra of the kaons and the J/ψ are different in the two decay modes. The efficiency is not computed directly with (6.8.1); instead the ratio of the π^0 efficiency in real data to that in MC is computed. For MC, the luminosity and branching ratio are replaced with the number of generated events. The J/ψ and kaon efficiencies then cancel under the assumption that they are equal in real data and MC. The ratio of the efficiency in real data to that in MC is used as a correction factor on the π^0 efficiency:

$$\eta_{\pi^0}^{corr} = \frac{\epsilon_{\pi^0}^{data}}{\epsilon_{\pi^0}^{MC}} .$$

The correction factor for the photon efficiency is then taken to be the square root of the correction factor for the π^0. The results are presented in bins of the transverse momentum $p_T(\gamma)$ of the photon (and $p_T(\pi^0)$ for the π^0 efficiency). The results from Ref. [128] for the photon efficiency correction factors are given in table 6.5. The main uncertainties are the statistical uncertainties on the numbers of events in real data and MC, the uncertainty due to the fit model, and the uncertainties on the branching ratios $Br(B^+ \to J/\psi K^+)$ and $Br(B^+ \to J/\psi K^{*+})$.

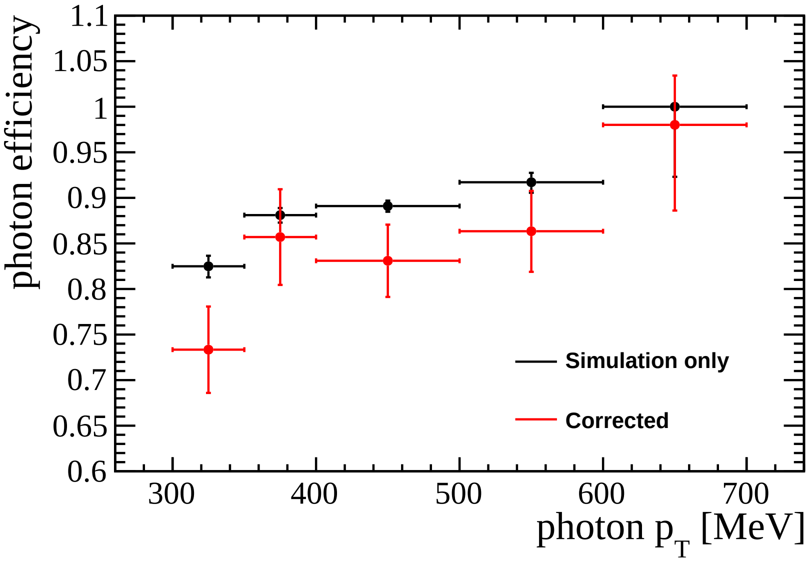

Figure 6.25: The photon reconstruction efficiency for χ_{c2} in bins of $p_T(\gamma)$. The black points show the efficiency directly from MC, and the red points show the MC efficiency after correction, with the full systematic and statistical errors.

The results, which apply generally to photon reconstruction at LHCb, are then applied to the CEP processes using MC. The simulated efficiency is plotted for χ_{c0}, χ_{c1} and χ_{c2} in bins of p_T, and the correction factors are applied separately for each bin. The full efficiency is computed using

$$\epsilon_{MC}^{corr} = \sum_{p_T \text{ bin } i} \frac{N^{rec}(i)\eta^{corr}(i)}{N^{true}(i)} ,$$

where $N^{rec}(i)$ and $N^{true}(i)$ refer to the number of reconstructed candidates and the total number of candidates in a $p_T(\gamma)$ bin. The distribution for χ_{c2} is shown in fig. 6.25, before and after the correction is applied.

The parameter used for binning is the *true* (before detector simulation) transverse momentum of the photon. It has been verified that the reconstructed p_T is reasonably close to the true p_T in simulation, for χ_c candidates with $p_T(\gamma) > 300\,\mathrm{MeV}$. The overall photon efficiencies are $77.5 \pm 4.1\,\%$ for χ_{c0}, $82.1 \pm 3.7\,\%$ for χ_{c1} and $82.5 \pm 3.6\,\%$ for χ_{c2}, including all errors.

6.8.5 Combined reconstruction efficiency in MC

The combined reconstruction and trigger efficiency $\epsilon_{rec}\epsilon_{trg}$ is computed from the number of reconstructed events accepted by the emulated trigger divided by the number of generated events in the acceptance. The numbers are given in table 6.6. The combined uncertainties are obtained, under the assumption that the errors are independent, by adding

	χ_{c0}	χ_{c1}	χ_{c2}
Generated total	4000000	4000000	4000000
In acceptance	305193	302974	302640
Triggered	205637	202263	201380
Reconstructed and triggered	137566	158836	165510
Photon corr. factor	0.90	0.93	0.94
Efficiency in acceptance	$40.6 \pm 3.3\,\%$	$49.0 \pm 3.7\,\%$	$51.8 \pm 3.9\,\%$

Table 6.6: Numbers of events in simulation, and reconstruction efficiencies. The efficiencies are the numbers of reconstructed events divided by the numbers of events in the acceptance, multiplied by the correction factors. The total number of generated events is not used for the efficiency computation.

the relative uncertainties in quadrature.

6.8.6 Clone tracks

A single particle may produce two tracks even if it remains intact while it passes through the tracking system. This is an effect of the tracking algorithms, and the extra tracks are referred to as “clone” tracks. Clone tracks cause an inefficiency because events with just two muons may have three tracks, where one track is a clone. It is unrelated to the effect of noise tracks, to be discussed in Section 8.1, because clone tracks are only produced in association with real tracks.

The clone probability is determined in events with three tracks by analysing the distribution of distance between the track which is not a muon and the nearest muon. The distance is measured in $\eta - \phi$ space. A peak of events with very small $\eta - \phi$ separation is seen, which is attributed to clone tracks. The clone probability is measured to be $0.34 \pm 0.08\,\%$, in good agreement with simulation.

6.8.7 Selection efficiency

The efficiency of the cuts and the fitting procedure is measured using MC by counting the number of events that pass each cut. One starts with events in which the χ_c-decay passes an acceptance cut of $2.0 < \eta < 4.5$ at truth level for all daughters, and is accepted by the emulated triggers. The MC samples contain 137566, 158836 and 165510 events for χ_{c0}, χ_{c1} and χ_{c2} respectively. The cut efficiencies are shown in table 6.7. The low selection efficiency for the χ_{c0} is primarily caused by the cut on $p_T(\gamma)$. As can be seen in fig. 6.24, there is a significant fraction of photons from the χ_{c0} which have lower transverse momentum than $300\,\mathrm{MeV}$.

A single Crystal Ball function is used to fit the mass peaks in simulation, as was discussed in Section 6.6. This gives yields of 28785 ± 184, 87062 ± 314 and 89448 ± 332

Cut (Units in MeV)	Rel. Efficiency (%)		
	χ_{c0}	χ_{c1}	χ_{c2}
VELO+Long tracks = 2	95.7	95.2	94.9
photons = 1	97.0	96.7	96.3
$2.0 < \eta < 4.5$ for γ, μ, μ (see caption)	98.7	98.8	98.8
$\lvert M(\mu\mu) - 3096.9\rvert < 65$	99.7	99.7	99.7
$p_T(\mu\mu) < 900$	92.7	89.8	76.7
$p_T(\gamma) > 300$	34.0	75.3	85.7
Total	28.8	61.4	59.2
Fit range: $3270 < M < 3800$	84.9	93.3	94.0

Table 6.7: Efficiency for each selection cut relative to the previous cut, or to the total number of reconstructed and triggered MC-events for the first row. The statistical precisions are between $0.25\,\%$ and $0.50\,\%$. The η-cut efficiency is computed for events which have already passed the η cut at the generator level.

and combined cut and fitting efficiencies of $21.0 \pm 0.1\,\%$, $54.8 \pm 0.1\,\%$, and $54.0 \pm 0.1\,\%$ for the χ_{c0}, χ_{c1} and χ_{c2}. The fitted yields are significantly smaller than the numbers of events which pass the cuts, but this is partly explained by the range used in the fit: events with incorrectly reconstructed photons have a mass that is very different from the peak value, and are thus not included in the fit. This is supported by the bottom row of table 6.7, which gives the efficiency of the mass cut. After the mass cut, the efficiencies obtained by counting are within $3\,\%$ of the ones obtained from the fit. There is however still a tail in the mass distribution which is not included in the fit, and this explains the final difference.

CHAPTER

SEVEN

ANALYSIS OF THE χ_b CEP SIGNAL

7.1 Introduction

There are three $\chi_b(1P)$ states, just like for the χ_c, but they are only separated by $33\,\mathrm{MeV}$ ($\chi_{b1} - \chi_{b0}$) and $19\,\mathrm{MeV}$ ($\chi_{b2} - \chi_{b1}$), so the peaks cannot be resolved separately. It is not attempted to extract the individual contributions; instead the sum of the χ_{b1} and χ_{b2} cross sections is measured. The signal intensity is much weaker than for the χ_c, but there is also less background. The decay chain $\chi_b \rightarrow \gamma\Upsilon(1S)$, $\Upsilon(1S) \rightarrow \mu^+\mu^-$ is reconstructed.

Similar selection cuts are used for the χ_b as for the χ_c. The efficiencies are evaluated using MC, and the uncertainties on the efficiencies are assumed to be the same as for the χ_c. The trigger and reconstruction efficiencies ($\epsilon_{trg} \times \epsilon_{reco}$) are $\sim\, 3\,\%$ for the χ_{b0}, $67 \pm 5\,\%$ for the χ_{b1} and $68 \pm 5\,\%$ for the χ_{b2}. The low efficiency for χ_{b0} is caused by a cut, $p_T(\gamma) > 200\,\mathrm{MeV}$, in the reconstruction. The cut is expected to reject a large fraction of the χ_b signal, because the simulation predicts that the photons from χ_{b0} have lower transverse momentum than for the other χ_b states, as was shown in fig. 6.24.

The dimuon mass spectrum for events with two muons and a photon, which also pass the global event cuts requiring no extra tracks and just one photon, is shown in fig. 7.1. The dimuon mass is required to be within 9370–9550 MeV, corresponding to a window of $\pm 90\,\mathrm{MeV}$ relative to the $\Upsilon(1S)$ mass given by the PDG [51]. Apart from the dimuon mass cut, the cuts used to select χ_b are the same as those for χ_c (described in Section 6.5.1). The selection efficiencies for χ_{b1} and χ_{b2} are $54\,\%$ and $46\,\%$. The main reason for the difference is that the cut $p_T(\mu\mu) < 900\,\mathrm{MeV}$ has a lower efficiency for the χ_{b2} than for the χ_{b1}. The mean of the selection efficiencies, $50\,\%$, is used to compute the cross section, with an uncertainty of $4\,\%$ due to the unknown χ_{b1} and χ_{b2} contributions. The efficiencies of the $p_T(\mu\mu)$ cut, after application of all other cuts, are $84\,\%$ and $69\,\%$ for χ_{b1} and χ_{b2} respectively.

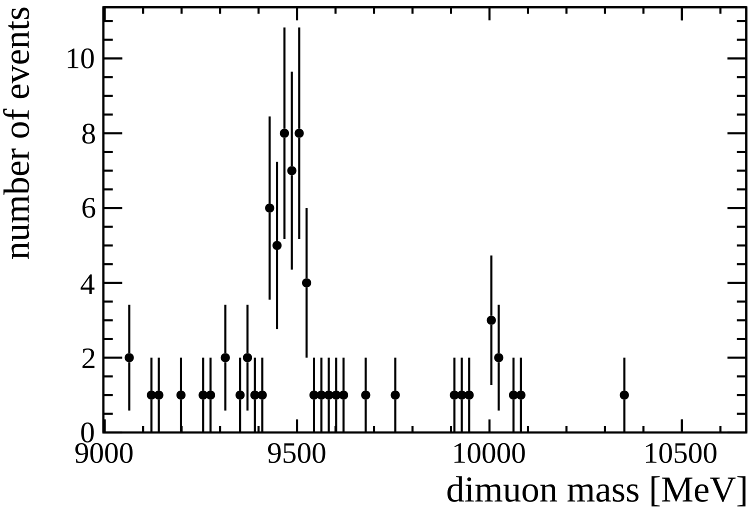

Figure 7.1: Dimuon mass spectrum for events accepted by the trigger, in events with at least two muons and a photon, in the range near the $\Upsilon(1S)$ mass. The $\Upsilon(2S)$ peak is also visible.

7.2 Feed-down background

Higher excitations of bottomonium can decay to $\chi_b(1P)$ mesons, and these events may be identified as χ_b CEP if the additional particles in the decay are not detected. The main sources of feed-down background to χ_b are radiative decays of $\Upsilon(2S)$ and $\Upsilon(3S)$, that is $\Upsilon \rightarrow \gamma\chi_b$. The photon from this decay has a low energy $\sim 100\,\mathrm{MeV}$ in the Υ rest frame, and may not be detected. There are other decays to final states with χ_b, involving pions, but these are expected to produce more activity in the detector, and have branching ratios comparable to, or smaller than, the radiative decays.

The higher Υ states can be produced in photoproduction CEP, and they also decay directly to dimuons like the $\Upsilon(1S)$ in the signal decay chain. The dimuon mass spectrum in the Υ region for CEP dimuon events (without photons) is shown in fig. 7.2. The numbers of candidates are extracted using a fit to the dimuon mass, where the peaks are modelled as Crystal Ball functions and the background is modelled as an exponential. The number of feed-down events can be estimated from the number of events seen in the dimuon channel. No attempt is made to estimate the CEP purity, because the feed-down also comes from inelastic production. The total number of $\Upsilon(2S)$ is obtained by dividing the number of $\Upsilon(2S)$ in the fit by the branching ratio to dimuons. This number is then multiplied by the branching ratio for the feed-down process and the branching ratios for the decay of

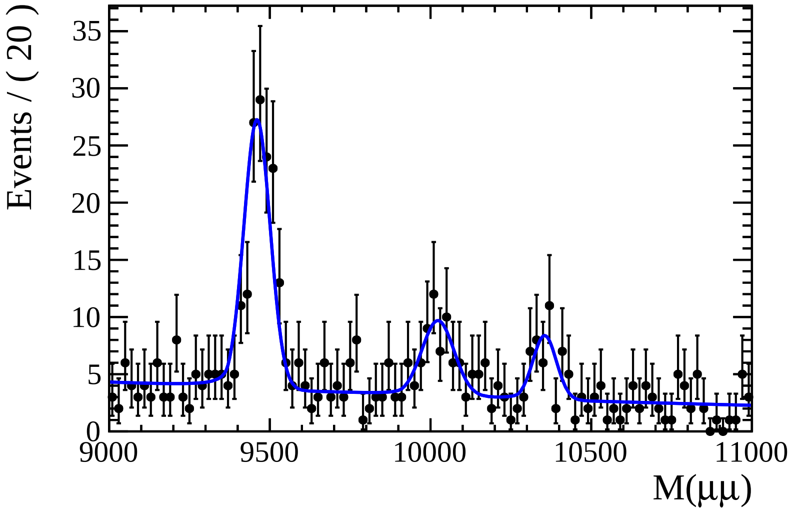

Figure 7.2: Dimuon mass spectrum in the region near the Υ masses for CEP dimuon events.

the χ_b in the analysis channel:

$$N(\Upsilon(2S) \to \gamma\chi_b;\, \chi_b \to \gamma\Upsilon(1S);\, \Upsilon(1S) \to \mu^+\mu^+) = \frac{N(\Upsilon(2S) \to \mu^+\mu^-)}{\mathrm{Br}(\Upsilon(2S) \to \mu^+\mu^-)}$$
$$\times \sum_{J=0}^{2} \mathrm{Br}(\Upsilon(2S) \to \gamma\chi_{bJ}) \times \mathrm{Br}(\chi_{bJ} \to \gamma\Upsilon(1S)) \times \mathrm{Br}(\Upsilon(1S) \to \mu^+\mu^-)\ .$$

The contribution from $\Upsilon(3S)$ is considered analogously. The final result is an expected number of feed-down events from $\Upsilon(2S)$ and $\Upsilon(3S)$ of 4.2, consisting of 0.07 χ_{b0} events, 2.6 χ_{b1} events and 1.5 χ_{b2} events. The contributions from $\Upsilon(2S)$ and $\Upsilon(3S)$ are approximately equal. In some of these events the photon in the $\Upsilon \to \gamma\chi_b$ decay will be detected, but there is a large uncertainty on the photon efficiency. The feed-down background has a large uncertainty, and it will not be subtracted in the following analysis. The main results are limits on χ_b CEP, and the presence of feed-down increases the values of these limits.

7.3 Fits to mass spectra

The mass spectra for the χ_b states in simulation are shown in fig. 7.3. The distribution is similar to that of the χ_c in that there is a tail at high mass. The signal is well described by a single Gaussian in a limited range near the known χ_b mass. The fit is shown in fig. 7.3b. A width of $\sigma = 51.3 \pm 0.9\,\mathrm{MeV}$ is returned by the fit to MC.

The mass spectra for real data are shown in fig. 7.4. In fig. 7.4a the cut on $p_T(\mu\mu)$ is not applied, yielding more events. Three mass spectra are shown, for events with zero, one and two extra tracks. When the $p_T(\mu\mu)$ cut is applied there are no events with one

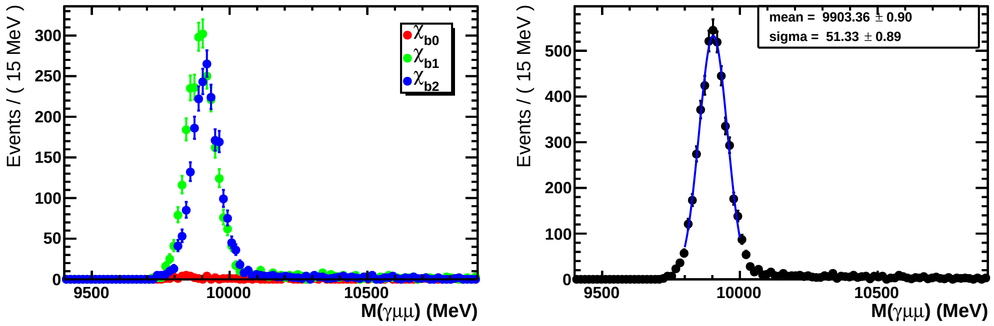

(a) Separate distributions (b) Sum

Figure 7.3: Mass spectra for the χ_b states in simulation, after the simulated trigger, reconstruction and selection have been applied. Figure (a) shows the three contributions separately, with the normalisation of each histogram proportional to the efficiency. The contribution from χ_{b0} is negligible due to the low efficiency. Figure (b) shows the sum of the three contributions. A Gaussian is fitted to the data in a narrow range around the signal peak.

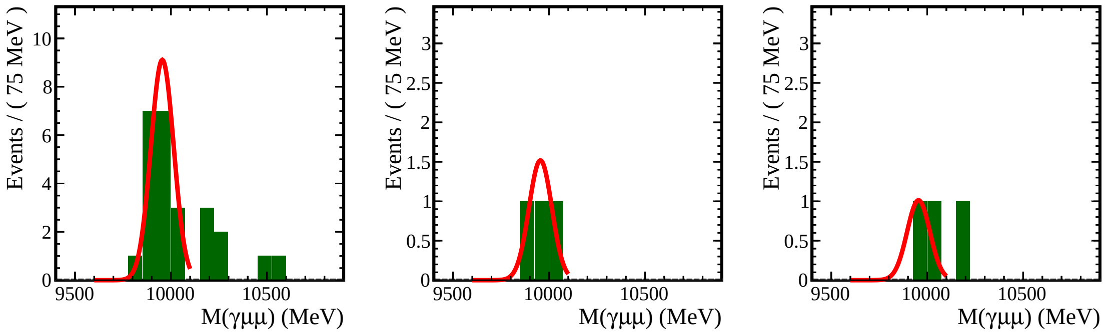

(a) Without $p_T(\mu\mu)$ cut

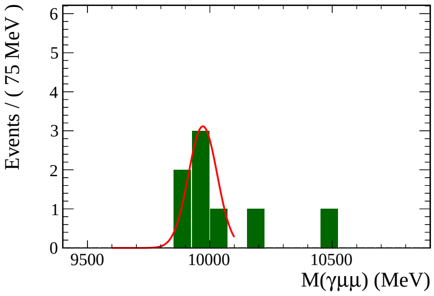

(b) With $p_T(\mu\mu)$ cut

Figure 7.4: Mass spectra in real data. The plots in (a) show the candidate mass distribution in events with zero, one and two extra tracks, from left to right, without the cut $p_T(\mu\mu) < 900\,\mathrm{MeV}$. In (b) the mass spectrum is shown for when the p_T cut is applied, for events with no extra tracks. There are no events with one or two extra tracks which pass the cut.

Symbol	$p_T(\mu\mu)$ [MeV]	Extra tracks	Number of events
$n_<$	< 900	0	6
$n_>$	> 900	0	12
$m_<$	< 900	1 or 2	0
$m_>$	> 900	1 or 2	5

Table 7.1: Four sub-samples are introduced for the combined estimation of the signal yield and the purity. The selection cuts and the numbers of candidates are given in the table.

or two extra tracks. The χ_b signal is fitted with a Gaussian in a narrow range around the known χ_b masses. No background component is included in the fit. The signal yield is alternatively obtained by counting candidates in the range 9750–10060 MeV, and this gives practically identical results to the fitting procedure. The width resulting from the fit is $\sigma = 60 \pm 10\,\mathrm{MeV}$, which is consistent with the MC.

The cut on $p_T(\mu\mu)$ is used to partition the data into a sample with $p_T(\mu\mu) < 900\,\mathrm{MeV}$, which is enriched in CEP signal, and a complementary sample dominated by inelastic production. The data are also partitioned into samples of events with no extra tracks, and samples with one or two extra tracks[1], denoted by n and m respectively. The numbers of candidates in the different samples are given in table 7.1.

7.4 Estimation of the true signal yield and the purity

The estimation of the signal yield N and the purity p are combined, giving the number of true CEP signal events Np, in order to correctly handle the low statistical precision and the asymmetric uncertainties. This is equivalent to the product of the yield and the purity $n_J p$ in (6.1.1), but the sum of the χ_{b1} and χ_{b2} yields, N, is used instead of n_J. A statistical model of the measurement process is constructed. The parameters of the model are the number of genuine CEP events $\bar{s} = Np$, the number of inelastic events with zero extra tracks $\bar{b}$, and the efficiency of the cut $p_T(\mu\mu) < 900\,\mathrm{MeV}$ on signal ϵ_{sig} and on background ϵ_{bg}. The purity estimation is similar to that used for the χ_c in that it makes use of the different p_T-distributions for the signal and the background to estimate the purity, but only two bins are used for the p_T-distribution.

The quantities $\bar{s}$ and $\bar{b}$ are the true mean numbers of signal and background events, which would be obtained if the experiment was repeated many times. The corresponding

---

[1] It may seem arbitrary to stop at two tracks, instead of including any number of tracks when considering the inelastic background. There is a trade-off between the statistical power and the correctness of the background properties, because the background p_T-distribution changes when allowing more tracks. The choice of two extra tracks is based on the fact that there are no candidates at all with three extra tracks, and only two candidates with more than two extra tracks (note that the cuts still require that there are no extra photons, and that the SPD multiplicity is less than ten).

expectation values of the experimental outcomes are:

$$\begin{aligned}
\bar{n}_< &= \bar{s}\epsilon_{sig} + \bar{b}\epsilon_{bg} \\
\bar{n}_> &= \bar{s}(1-\epsilon_{sig}) + \bar{b}(1-\epsilon_{bg}) \\
\bar{m}_< &= \bar{B}\epsilon_{bg} \\
\bar{m}_> &= \bar{B}(1-\epsilon_{bg}) \ ,
\end{aligned}$$

where $\bar{B}$ is the number of background events with one or two extra tracks, and the background efficiency is assumed to be the same for events with zero, one and two extra tracks (this will be discussed later).

A combined pdf is constructed with parameters $(\bar{s}, \bar{b}, \epsilon_{sig}, \epsilon_{bg})$, giving the probability distributions for the number of observed candidates $n'_<$ and $n'_>$. The primed symbols are used to indicate the results of hypothetical experiments, which are generally not the same as the data. The signal efficiency ϵ_{sig} is taken from the SuperCHIC MC. The background efficiency ϵ_{bg} is determined from $m_</(m_< + m_>)$, which is zero, but with a large uncertainty which has to be incorporated in the model. This is done by sampling ϵ_{bg} from a probability distribution. A Bayesian approach is used to get a reasonable distribution representing the statistical uncertainty on $m_</(m_< + m_>)$.

7.4.1 Distribution of the background efficiency

The probability distribution of the background efficiency is determined using information from events with one or two extra tracks ($m_<$, $m_>$). For any given efficiency the probability that $m'_<$ events pass the cut is given by the binomial distribution,

$$P(m'_<|\epsilon_{bg}) = \binom{B}{m'_<} {\epsilon_{bg}}^{m'_<} (1-\epsilon_{bg})^{B-m'_<} \ ,$$

where $B = m'_< + m'_>$ is the total number of events with extra tracks in a given hypothetical experiment. Bayes' theorem can be used to determine the probability distribution for the background efficiency given our data,

$$P(\epsilon_{bg}|m'_<) = \frac{P(m'_<|\epsilon_{bg})P(\epsilon_{bg})}{P(m'_<)} \ .$$

The *prior* $P(\epsilon_{bg})$ is a distribution which represents the knowledge of the background efficiency before the measurement. A flat prior is assumed, i.e. the probability is the same for all efficiencies. The term $P(m'_<)$ is irrelevant because it does not depend on ϵ_{bg},

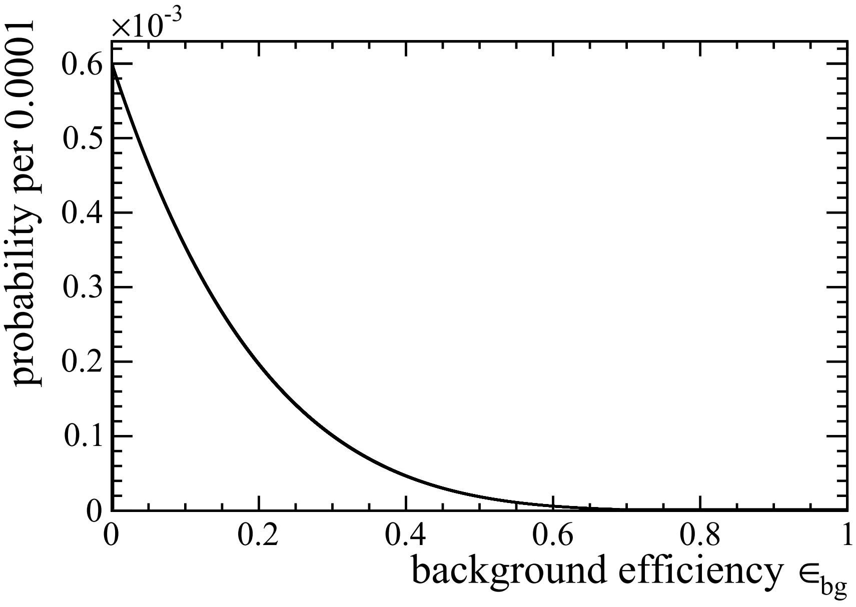

Figure 7.5: Probability distribution for background efficiency computed using Bayes' theorem.

and it is fixed by the normalisation,

$$\int_0^1 d\epsilon_{bg}\, P(\epsilon_{bg}|m'_<) = 1 \ .$$

The pdf for the background efficiency is then given up to a constant factor by the binomial distribution considered as a function of the efficiency, with the total number of events and the number of events passing the cut fixed to the actual experimental outcome. The probability distribution for ϵ_{bg} given the experimental outcome $m_< = 0$ and $m_> = 5$ is shown in fig. 7.5.

As with the sPlot analysis for the χ_c purity (Section 6.7), an approximation is made when events with more than zero extra tracks are used directly to represent the inelastic background, instead of performing an extrapolation of the shape of the p_T-distribution. The χ_c purity was measured both using the extrapolation, yielding the canonical purity of $29.7\,\%$, and using the background shape from events with one extra track, giving the value $44\,\%$. The integral of the χ_c background pdf for $p_T(\mu\mu)$ (as used in the main purity analysis) up to $900\,\mathrm{MeV}$ is 0.53 for the extrapolation and 0.48 for data with one extra track. This integral represents the probability for background events to have p_T lower than $900\,\mathrm{MeV}$, and corresponds to ϵ_{bg}. To assess the uncertainty as it applies to the χ_b, the analysis described in the following paragraphs is repeated with the values e_{bg} sampled from the background efficiency pdf replaced with $e'_{bg} = (0.53/0.48)e_{bg}$. This correction is based on the χ_c, but one can expect a similar effect for the χ_b. The replacement is found to have a negligible effect on the results.

7.4.2 Statistical model for CEP candidate yield

The number of true CEP signal events is determined using a statistical model, using maximum likelihood estimation. There are two parameters to be estimated: the number of genuine CEP events $\bar{s}$, and the number of inelastic production events $\bar{b}$, which is a nuisance parameter. The probability distribution for $n'_<$ and $n'_>$ at a given pair of parameter values $(\bar{s}, \bar{b})$ is determined using a MC method. The number of signal interactions s, the number of inelastic interactions b and the background efficiency e_{bg} are repeatedly sampled from their component pdfs. The pdfs for s and b are taken as Poisson distributions with mean values $\bar{s}$ and $\bar{b}$. The number of signal events which pass the p_T-cut, $s_<$ is then given by the binomial distribution with parameters s and ϵ_{sig}. The number of events which fail the cut is then given by $s_> = s - s_<$. The background is computed in the same way. The efficiency of the p_T-cut on background is sampled from the distribution determined in the previous section. Even though a real experiment will have a fixed ϵ_{bg}, the value can be sampled in every iteration because it does not depend on s nor b, and it should be averaged over (the result would be the same if the analysis was done for fixed values of ϵ_{bg} and the results were then combined with weights determined from the pdf of ϵ_{bg}). For the signal efficiency the mean of the χ_{b1} and χ_{b2} efficiencies in the MC is used, $\epsilon_{sig} = 76.1\,\%$.

The sampling process of selecting s and b from the Poisson distributions, then selecting how many events pass the cut, is repeated a million times for each $(\bar{s}, \bar{b})$. The probability for a given $(n'_<, n'_>)$ is computed as the number of times that result was obtained, divided by the total number of samplings.

7.4.3 Analysis of the likelihood distribution

The probability distribution evaluated at the outcomes seen in the real experiment is interpreted as the likelihood of $(\bar{s}, \bar{b})$,

$$\mathcal{L}(\bar{s}, \bar{b}) = P(n'_< = 6 \cap n'_> = 12 | \bar{s}, \bar{b}) \ .$$

The two-dimensional likelihood distribution is computed by repeating the MC estimation for different values of $\bar{s}$ and $\bar{b}$.

The parameters $(\bar{s}_0, \bar{b}_0) = (6.8, 11.2)$ have the highest likelihood, and are chosen as the best estimate. The natural logarithm of the likelihood is shown for our model in fig. 7.6. The confidence interval is determined using the test statistic

$$D = -2 \log \left(\frac{L(\bar{s}, \bar{b})}{L(\bar{s}_0, \bar{b}_0)} \right) \ ,$$

which is known from statistical theory to have a χ^2 distribution. The lower and upper bounds of the confidence interval are set to the lowest and highest values of $\bar{s}$ on the

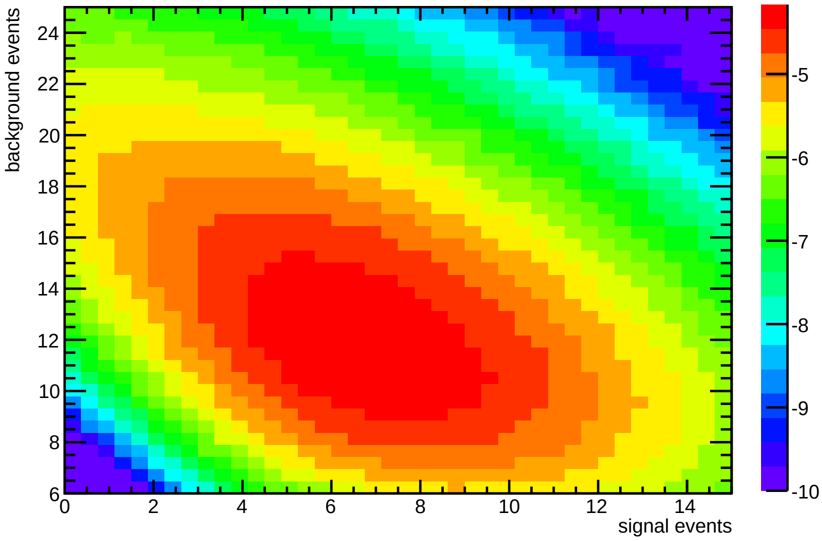

Figure 7.6: Two-dimensional plot of the logarithm of the likelihood as a function of the number of signal events and the number of background events.

contour of $D = 1$. The confidence interval for $\bar{s}_0 = 6.8$ is then 2.6 to 11. This corresponds to an estimate for the true value of the number of signal events passing the cut of

$$\bar{s}_< = 5.1^{+2.9}_{-3.1} \ .$$

7.4.4 Limit on χ_{b1} and χ_{b2} production

The results can also be interpreted as a limit on the number of χ_{b1} and χ_{b2} candidates. The probability distribution $P(n'_<, n'_> | \bar{s}, \bar{b})$ is determined using the same statistical model as above. The limit at the $95\,\%$ confidence level is the least value of $\bar{s}$ for which no more than $5\,\%$ of the generated pseudo-experiments produce values of $n'_<$ or $n'_>$ which are smaller than or equal to the values on real data ($n_<$ and $n_>$). In other words, the limit on $\bar{s}$ is the smallest value for which

$$P(n'_< \leq n_< \cap n'_> \leq n_> | \bar{s}, \bar{b}) < (1 - 0.95) \ ,$$

for all $\bar{b}$. The resulting limit on the number of signal events which pass the cut is

$$s_< < 12.0 \quad \text{at} \quad 95\,\% \text{ C.L.} \ .$$

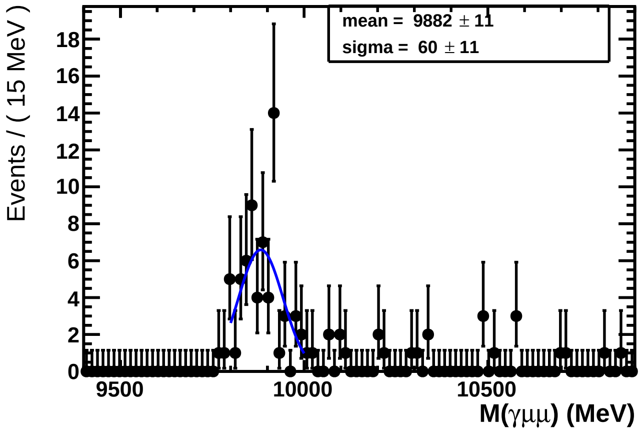

Figure 7.7: Fit to MC with contribution from each χ_b state scaled such that the number of candidates from each state is approximately the same.

7.5 Contribution from χ_{b0}

After trigger, reconstruction and selection efficiencies, the χ_{b1} and χ_{b2} efficiencies are greater than the χ_{b0} efficiency by a factor ~ 70. The low χ_{b0} efficiency is due to the very low $p_T(\gamma)$ seen in the MC for χ_{b0}. The possibility that the efficiency is greater in real data must be considered. Furthermore, the χ_{b0} cross section is expected (from SuperCHIC) to be greater than the χ_{b1} and χ_{b2} cross sections by more than an order of magnitude.

To check if the fit can rule out a χ_{b0} contribution, the MC is re-scaled such that the number of candidates is approximately the same for χ_{b0}, χ_{b1} and χ_{b2}. The fit to these simulated data is shown in fig. 7.7, and it is consistent with the real data, so a χ_{b0} contribution is possible in the data. The candidate mass is not used to constrain the contributions from the χ_b states because there is no independent confirmation of the calibration of the photon energies.

The limit on the χ_{b1} and χ_{b2} yields remains valid in the presence of a significant χ_{b0} signal. It can be interpreted as a limit on the sum of the number of candidates for all χ_b states (before correcting for efficiencies). Because the efficiency for χ_{b0} is much lower than for the higher states, the limit on the cross section is weaker for the χ_{b0}. This is discussed further in Chapter 8.

CHAPTER

EIGHT

LUMINOSITY AND CROSS SECTIONS

The cross sections for central exclusive production of χ_c and χ_b mesons in the LHCb angular acceptance are computed based on the results in Chapter 6 and Chapter 7. The uncertainties are computed and combined. The results are finally compared to theoretical predictions and other experimental results, and their relevance to the understanding of CEP is discussed. The effective integrated luminosity is the final factor required for the cross section computation, and this is the topic of the first section. The cross sections are computed and the results are discussed in the following sections.

8.1 Luminosity

The integrated luminosity is used by all cross section analyses at LHCb so it is computed centrally in the collaboration [129]. The total integrated luminosity cannot be used directly by the CEP analyses, because events with multiple visible interactions are rejected by the global event cuts regardless of whether one of the interactions is CEP. Because the probability to have multiple proton-proton interactions depends on the instantaneous luminosity, this effect is treated as a modified luminosity instead of an efficiency. The method for determining the total luminosity is discussed first, because the notation and the concepts are required to explain the modified luminosity for CEP, the *effective luminosity*.

8.1.1 Total luminosity

The luminosity is a function of the beam parameters, as was discussed in section 2.1.2. The rate of interactions is determined by the instantaneous luminosity $\mathcal{L}$ and the cross section:

$$\frac{dN_{vis}}{dt} = \mathcal{L}\sigma_{vis} \, . \tag{8.1.1}$$

One usually works with the number of *visible* interactions N_{vis} and the visible cross section σ_{vis}, corresponding to processes that produce particles in the LHCb acceptance. Processes which only produce particles outside of the LHCb acceptance need not be considered when computing the luminosity. One is free to define a visible interaction in a way that is suitable for the experiment, but when redefining a visible interaction, the definition of the visible cross section has to be changed accordingly, such that

$$\frac{dN_{vis}/dt}{\sigma_{vis}} = \mathcal{L} \tag{8.1.2}$$

remains constant. The luminosity $\mathcal{L}$ cannot depend on the definition of a visible interaction because it is determined by the beam parameters.

The luminosity at LHCb is computed using a relative and an absolute measurement.

- Any quantity that is proportional to the number of interactions is also proportional to the luminosity, so the relative luminosity is tracked by counting such quantities.
- The absolute luminosity is computed using two separate methods: a Van der Meer scan and by reconstructing beam-gas interaction vertices.

The number of primary vertices is used to track the luminosity on the timescale of a few minutes. Multiple luminosity counters such as the number of VELO tracks and the number of primary vertices are stored in a compact format together with the other data. For the relative measurement it is useful to introduce the number of events, rewriting (8.1.2) as

$$\mathcal{L} = \frac{f_{evt}\mu}{\sigma_{vis}} \, , \tag{8.1.3}$$

where f_{evt} is frequency with which the LHC beams cross in LHCb, and μ is the mean number of visible interactions per event. An event is defined as a bunch crossing, and even bunch crossings with no interactions are counted as events.

The probability that two given protons interact in an event is very small, but there are approximately 10^{11} protons in each bunch, and μ is of order unity. The number of interactions per event has a Poisson distribution:

$$P(n) = \frac{\mu^n e^{-\mu}}{n!} \, . \tag{8.1.4}$$

This is used to get an alternative expression for μ:

$$P(0) = e^{-\mu} \Rightarrow \mu = -\log P(0) \, .$$

The quantity $P(0)$ can be measured in data as the fraction of bunch crossings which have no visible interaction. There is, however, a small probability that background processes are counted as visible interactions. The measured fraction of events with no visible interactions is thus

$$P(0)^{bb} = P(0)P(0)^{BG} \ ,$$

where $P(0)^{BG}$ represents the probability to have no background and $P(0)$ is the probability to have no visible true proton interactions. The superscript *bb* refers to *beam-beam-*events, which are the normal events where two bunches cross in LHCb.

The relevant backgrounds are beam-gas interactions and spillover. Beam-gas interactions are processes where a beam proton interacts with a gas molecule in the LHC vacuum chamber, either in the LHCb beam pipe or outside of LHCb. Spillover refers to a residual signal in the tracking detectors from a track produced in the previous event. These backgrounds can be isolated by analysing events with single beams and with no beams at all in LHCb. The background-corrected μ can be written as

$$\mu = -\log P(0)^{bb} + \log P(0)^{be} + \log P(0)^{eb} - \log P(0)^{ee} \ , \tag{8.1.5}$$

where $P(0)^{be}$, $P(0)^{eb}$ and $P(0)^{ee}$ are the fractions of events without reconstructed PVs in background-only events, where the superscripts refer to the presence (b) or absence (e for empty) of beam 1 and beam 2 respectively. This formula subtracts the backgrounds due to beam 1 and beam 2 explicitly in the two middle terms. The final ee term is included because the no-beam noise is subtracted in both the eb and be terms, so it has to be added back to count it only once. In practice, the spillover and other detector noise are very small because the LHC operated with a $50\,\mathrm{ns}$ bunch spacing, as opposed to the design value of $25\,\mathrm{ns}$. Most ee-events are in the time slots between two collision events ($\pm 25\,\mathrm{ns}$), and these events contain a significant spillover contribution, so if these were naively used in (8.1.5), the spillover would be greatly overestimated. For these two reasons, the ee-contribution is ignored. The proportion of events with no visible interaction is shown in fig. 8.1, separately for each event type. In the plot one can see that the value of $P(0)$ does not change much during the period of data taking. The conditions were intentionally kept stable over the whole period. The inset plot shows that $P(0)$, and thus the luminosity, was stable over a single fill. The luminosity will normally decay exponentially over the fill, but luminosity levelling is used to keep it almost constant.

Application of Eqs. (8.1.3, 8.1.5) and the methods in Ref. [129] yields a total integrated luminosity of $1025 \pm 36\,\mathrm{pb}^{-1}$ for the 2011 run of LHCb. The low multiplicity trigger was not used during the full data taking period, and the luminosity with this trigger present is $953 \pm 33\,\mathrm{pb}^{-1}$, constituting $93\,\%$ (by luminosity) of the full dataset.

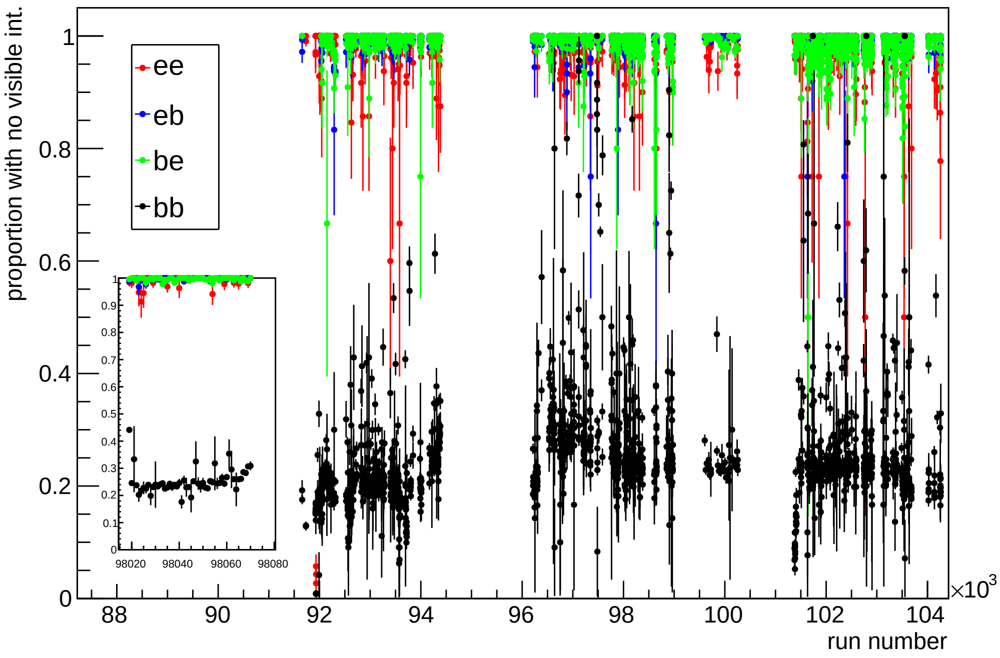

Figure 8.1: Fraction of events with no visible interactions, $P(0)^{bb}$, $P(0)^{eb}$, etc. The graph corresponds to the full dataset, and the inset shows the evolution over a single fill. The run number is used as a time coordinate.

8.1.2 Effective luminosity

When cutting on the number of photons and tracks in the event, for the CEP analysis, events with a genuine CEP interaction and any other visible interaction are rejected. The luminosity corresponding to such events with multiple visible interactions has to be discounted. The effective luminosity is defined as the luminosity obtained when only counting single-interaction events. For the effective luminosity, a *visible interaction* is defined to be a pp-interaction that produces either tracks or photon candidates (i.e., tracks and photons are used as luminosity counters instead of primary vertices).

The definition of μ is the mean number of interactions per event, and the mean of any such discrete distribution can be expanded as a sum,

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

The effective luminosity is obtained by cutting off the sum at $n = 1$,

$$\lambda = 0 \times P(0) + 1 \times P(1) = \mu \exp(-\mu) \ ,$$

where λ is a quantity analogous to μ with the interpretation “mean number of interactions per event in events with no more than one interaction", and the second equality is obtained using the expression for the Poisson distribution. By replacing μ with λ in (8.1.3), the effective luminosity $\mathcal{L}_{eff}$ can be related to the full luminosity:

$$\mathcal{L}_{eff} = \frac{f_{evt}\lambda}{\sigma_{vis}} = \frac{f_{evt}\mu e^{-\mu}}{\sigma_{vis}} = \mathcal{L}e^{-\mu} \; . \tag{8.1.6}$$

The factor $e^{-\mu} = P(0)$ is measured in data by counting events with no visible interactions, where the definition of a visible interaction requires just a photon or a track. The integrated effective luminosity is computed for each run by taking the integrated luminosity $\mathcal{L}_{int}$ from the central LHCb calculation, and computing $P(0)^{bb}$ for each run.

The background could be subtracted in the same way as in (8.1.5), but it turns out that this is not necessary. After the background-corrected effective luminosity had been determined, one would have to take into account the reduction of the efficiency due to backgrounds. This is the case where true single-interaction CEP interactions get rejected because of tracks from background processes. The effective luminosity would then be given as

$$\mathcal{L}_{eff} = \mathcal{L}\, P(0)\, P(0)^{BG} \; .$$

The two last factors are equivalent to $P(0)^{bb}$, so the effective luminosity is instead computed directly as

$$\mathcal{L}_{eff} = \mathcal{L}P(0)^{bb} \; . \tag{8.1.7}$$

This represents the dataset available after correcting for the efficiency of the cuts on track and photon multiplicity due to multiple proton-proton interactions, machine-induced background and noise. An integrated effective luminosity of $222.3\,\mathrm{pb}^{-1}$ is found for the 2011 dataset, $23\,\%$ of the total luminosity.

8.1.3 Uncertainty on the effective luminosity

Luminosity counts are collected at a rate of 1 kHz, so the statistical uncertainty on μ is negligible. The uncertainty on the total luminosity is $3.5\,\%$ and is dominated by the absolute luminosity measurement [129].

The full event information is needed to apply the definition of a visible interaction used to define the effective luminosity. Events from an unbiased trigger, which randomly accepts events, are used to measure $P(0)^{bb}$. While this trigger only has a rate of about $100\,\mathrm{Hz}$, the statistical uncertainty on the integrated effective luminosity due to the determination of $P(0)^{bb}$ is only $\pm 0.3\,\mathrm{pb}^{-1}$. There were about 1000 pairs of colliding bunches and each pair had a different luminosity. The luminosity in the bunches had a spread of about $10\,\%$. The effect of such a spread on the effective luminosity was however found to be negligible.

	χ_{c0}	χ_{c1}	χ_{c2}
Fit (± stat., ± sys.) N	$143 \pm 23 \pm 126$	$880 \pm 170^{+1500}_{-840}$	$5228 \pm 186^{+880}_{-1500}$
Purity p	$29.7 \pm 30\,\%$	$29.7 \pm 7\,\%$	$29.7 \pm 6.3\,\%$
Efficiency #1 $\epsilon_{trg}\epsilon_{rec}$	$40.6 \pm 3.3\,\%$	$49.0 \pm 3.7\,\%$	$51.8 \pm 3.9\,\%$
Selection eff. ϵ_{sel}	$21.0 \pm 0.1\,\%$	$54.8 \pm 0.1\,\%$	$54.0 \pm 0.1\,\%$
Luminosity $\mathcal{L}_{int}$	$222.3 \pm 8.2\,\mathrm{pb}^{-1}$	$222.3 \pm 8.2\,\mathrm{pb}^{-1}$	$222.3 \pm 8.2\,\mathrm{pb}^{-1}$
Cross section $\mathcal{B}_{\psi}\mathcal{B}_{\chi}\sigma$	$2.2 \pm 3.0\,\mathrm{pb}$	$4.3^{+7.6}_{-4.4}\,\mathrm{pb}$	$25.0^{+7.1}_{-9.2}\,\mathrm{pb}$

Table 8.1: Factors for the cross section measurement, as determined in Chapter 6, and the final cross section for each χ_c state. The second row ("Efficiency #1") shows the products of the trigger and reconstruction efficiencies. The given cross sections are the production cross sections multiplied with the branching ratios $\mathcal{B}_{\chi} = \mathrm{Br}(\chi_c \to \gamma J/\psi)$ and $\mathcal{B}_{\psi} = \mathrm{Br}(J/\psi \to \mu^+\mu^-)$ (the value of $\mathcal{B}_{\chi}$ is of course different for the χ_c states).

If background subtraction is performed on $P(0)^{bb}$, by dividing out the contributions from beam-empty and empty-beam events, one will get an estimate for the true $P(0)$, i.e. the probability to have no proton-proton interactions. If this is used in (8.1.7) in place of the experimental value $P(0)^{bb}$, which also contains the probability that there is no background, one gets a larger effective luminosity which includes events with single interactions and background. The difference in effective luminosity between these definitions is $1.1\,\%$, and this is used as an uncertainty on the background, because the nature of the background processes is not completely known. The final integrated effective luminosity is $222.3 \pm 8.2\,\mathrm{pb}^{-1}$.

8.2 Results for χ_c-production

The cross sections are computed using Equation (6.1.1), repeated here:

$$\sigma_J = \frac{n_J p}{\epsilon_{tot,J}\mathcal{L}_{eff,int}} \quad \text{for } J = 0, 1, 2\ . \tag{8.2.1}$$

The formula applies to both the χ_c and the χ_b. The factors have already been discussed, as have their uncertainties. The factors for the cross sections and their uncertainties, and the final absolute cross sections are given in table 8.1. The final cross section is computed using (8.2.1). There is a large uncertainty on the cross sections for the χ_{c0} and χ_{c1}. The uncertainty for the χ_{c0} is dominated by the purity estimate, because it cannot be confirmed that the purity is the same as for the other χ_c states. The uncertainty on the χ_{c1} cross section is dominated by the fit model, and specifically the overlap of the χ_{c1} and χ_{c2} peaks.

The cross section ratios are given separately from the absolute cross sections, because these quantities are theoretically interesting. This would normally produce more precise

	χ_{c0}/χ_{c1}	χ_{c0}/χ_{c2}	χ_{c1}/χ_{c2}
Observed ratio	$0.69^{+0.24}_{-0.61}$	$0.070^{+0.092}_{-0.046}$	$0.067^{+0.067}_{-0.042}$
Corrected for B.R.	20^{+7}_{-18}	$1.19^{+1.56}_{-0.68}$	$0.039^{+0.039}_{-0.024}$

Table 8.2: Cross section ratios for χ_c. The second row includes correction for the branching ratios, and can be interpreted as the ratios of the production cross sections.

results due to cancellations, but in this case the ratios also have large systematic uncertainties from the fitting procedure and the purity estimate. The cross section ratio for two χ_c states labelled i and j is trivially given by

$$\frac{\sigma_i}{\sigma_j} = \left(\frac{n_i}{n_j}\right)\left(\frac{p_i}{p_j}\right)\left(\frac{\epsilon_{tot,j}}{\epsilon_{tot,i}}\right) . \tag{8.2.2}$$

The ratios are computed using a statistical model of the measurement process. The uncertainties are represented by Gaussian distributions. The correlation of the fit uncertainties of χ_{c1} and χ_{c2} is explicitly included. The parameters to be determined are the yield of one χ_c state and the ratio of the cross section of two χ_c states. The best parameter values are determined, using two of the cross sections in table 8.1 as inputs.

The cross section ratios are given in table 8.2. The ratios are reported both with and without a correction for the branching ratios. The uncertainties on the branching ratios of $\chi_c \to \gamma J/\psi$ are a few percent, and can thus be ignored compared to the other errors. The results indicate that the χ_{c0} and χ_{c2} have greater production cross sections than the χ_{c1}, but the uncertainties are large.

8.3 Results for χ_b production

The limit on the χ_b cross sections is computed using (8.2.1), and is based on the limit on the number of signal interactions determined in Chapter 7. The combined relative uncertainty on the efficiencies and the luminosity is $15\,\%$. The limit on the sum of the production cross sections in the fiducial acceptance times the branching ratios is

$$(\sigma_{\chi_{b1}} B_1 + \sigma_{\chi_{b2}} B_2) < 160\,\mathrm{fb} \quad \text{at } 95\,\%\ \text{C.L.} ,$$

with

$$B_J = \mathrm{Br}(\chi_{bJ} \to \gamma\Upsilon(1S)) \times \mathrm{Br}(\Upsilon(1S) \to \mu^+\mu^-) \quad \text{for} \quad J = 1, 2 .$$

The result could also be interpreted as a limit on the sum of all three χ_b cross sections. There is, however, a large uncertainty on the χ_{b0} efficiency, which is difficult to estimate. Because there are no simulated photons from χ_{b0} with true transverse momentum above

$300\,\mathrm{MeV}$ (fig. 6.24), all reconstructed MC χ_{b0} candidates have incorrect photon transverse momentum. Due to the large uncertainty on the efficiency, a reasonable limit cannot be set on the sum of all the cross sections.

The purity factor has a large uncertainty, and it is interesting to determine the "semi-exclusive" cross section, which is the cross section for all CEP χ_b candidates with zero extra tracks, without performing an estimation of the purity. The cut on the number of photons is not used, so all feed-down background is included, but the number of candidates in data does not change when removing this cut. Except for removing the cut on the number of photons, the standard selection is used. This results in a yield of 6 candidates. If χ_{b0} dominance is assumed, i.e. the χ_{b0} efficiency is used, the cross section is $5.2 \pm 2.1 \pm 3.2\,\mathrm{pb}$, where the first uncertainty is statistical and the second is systematic. If instead χ_{b1} and χ_{b2} dominance is assumed, the cross section is $80 \pm 33 \pm 11\,\mathrm{fb}$.

8.4 Conclusion

The aim of the 2011 analysis was to improve the precision of the results with respect to the 2010 analysis [63, 104], with a greater dataset. The improved statistical precision has allowed a better understanding of the signal and background, but it also made clear that there are large uncertainties on the fit and the purity. The χ_{c0}, χ_{c1} and χ_{c2} cross sections were given for the LHCb 2010 dataset in a conference report [104]: respectively $9.3 \pm 4.5\,\mathrm{pb}$, $16.4 \pm 7.1\,\mathrm{pb}$ and $28.0 \pm 12.3\,\mathrm{pb}$. These values can be compared with table 8.1. The relative reduction with respect to Ref. [104] appears to be due to the fitting procedure and the estimation of the effective luminosity.

The cross sections times branching ratios for production in the LHCb acceptance given by SuperCHIC for χ_{c0}, χ_{c1} and χ_{c2} are respectively $14\,\mathrm{pb}$, $9.8\,\mathrm{pb}$ and $3.3\,\mathrm{pb}$ with a theoretical uncertainty of a factor 4–5 [58] (with an acceptance correction factor of 0.076; all quantities are given for final state particles in $2.0 < \eta < 4.5$). As was also seen in Ref. [104], there is more χ_{c2} production in data than expected from SuperCHIC, both in absolute terms and relative to the other χ_c states. By itself, the χ_{c2} cross section is consistent with SuperCHIC when considering the theoretical uncertainties. The ratio of the contributions is however suspect. The enhancement of the χ_{c2} can be caused by the inelastic background processes. The background is enriched in χ_{c2} because it has a greater p_T, and the χ_{c2} cross section increases with $\langle p_\perp^2 \rangle^2$ [75]. This effect is constrained by the fact that the purities are now measured individually for each χ_c state using sPlot, albeit with large uncertainties. There may still be some inelastic background that cannot be identified using the analysis of the $p_T(\mu\mu)$-distribution.

The χ_{c0} cross section is lower than the SuperCHIC prediction, but consistent with the previous LHCb measurements. The discrepancy cannot be explained by errors in the purity estimate and fitting procedure, unless the width of the χ_{c0} peak in real data is much greater than expected. However, if the theoretical uncertainty is considered, the result is consistent with SuperCHIC. The χ_{c1} result is also consistent with the SuperCHIC value. Pasechnik et al. [87] showed that the CEP cross sections and cross section ratios have a large model dependence, of an order of magnitude for the ratios and two orders of magnitude for the cross sections. Ratios which are similar to the current LHCb results are obtained for some models, but the results are for the Tevatron energy and not specific to the LHCb acceptance, so they cannot be compared directly. The χ_c results indicate that further study is necessary in order to eliminate the backgrounds and improve the precision.

The SuperCHIC predictions for the cross sections times branching ratios of χ_{b0}, χ_{b1} and χ_{b2}, for production in the LHCb acceptance, are respectively $0.13\,\mathrm{fb}$, $0.004\,\mathrm{fb}$ and $0.010\,\mathrm{fb}$. No measurable signal is expected with an effective luminosity of $222.3\,\mathrm{pb}^{-1}$. There is a clear peak in the data at the χ_b mass, but it cannot be uniquely ascribed to CEP, because there is a large uncertainty on the CEP purity, and there is potentially also a significant feed-down background. When the data are considered as a cross section measurement, the effective significance is thus just above one standard deviation when including the purity uncertainty, before considering the feed-down background. Because the mass resolution is insufficient to determine the individual contributions of χ_{b0}, χ_{b1} and χ_{b2}, the efficiency corresponding to the observed signal is not known. The efficiency for χ_{b0} is much lower than for χ_{b1} and χ_{b2} (assuming the p_T-distribution from SuperCHIC), so the precision on the χ_{b0} cross section is worse. For this reason, only a limit on the χ_{b1} and χ_{b2} is given, using the mean of the corresponding efficiencies. The limit at $160\,\mathrm{fb}$ is a factor $\sim 10^4$ above the predicted cross section.

The background of CEP $\Upsilon(1S)$ with randomly associated photons cannot be estimated from the mass distribution as for the χ_c, but the mass spectra do not appear to have a significant background contribution. It is possible, as was mentioned for the χ_c, that not all the proton dissociation is distinguishable using the dimuon p_T. This inelastic background is expected to be smaller than the CEP contribution [76], though the χ_c analysis indicates that it may instead be a factor 2 greater. The inelastic background cannot however explain the large signal seen, even if the suppression of χ_{b2} does not apply. The data are consistent with only feed-down background if the probability to detect the second photon is low.

Bzdak predicted a total χ_{b0} cross section of $0.3\,\mathrm{nb}$ [84] at $\sqrt{s} = 14\,\mathrm{TeV}$, which corresponds to a cross section of $B_0\sigma_{\chi_{b0}} = 8\,\mathrm{fb}$ within the acceptance cuts, assuming the same acceptance factor as in SuperCHIC. Yuan predicted a total cross section of $0.88\,\mathrm{nb}$ at the Tevatron [85], and the cross section is not expected to change much when going to $\sqrt{s} = 7\,\mathrm{TeV}$ (or $14\,\mathrm{TeV}$). It was also remarked in Ref. [85] that the higher excited state $\chi_{b0}(2P)$ may be produced with approximately the same cross section as the $\chi_{b0}(1P)$ considered here. The branching ratio of the decay of $\chi_{b0}(2P)$ to $\gamma\Upsilon(1S)$ is however only half that of the $\chi_{b0}(1P)$, which is $1.76\,\%$. There are 5 events which are consistent with

$\chi_{b0}(2P) \to \gamma\Upsilon(1S)$ in the mass spectrum. In Ref. [85], the decay to the higher $\Upsilon(2S)$ was suggested, as this has a larger branching ratio.

The uncertainties on the fits to the mass spectra could be substantially reduced if we had a better understanding of the energy calibration of the LHCb calorimeter, as the positions of the mass peaks could then be fixed. A future analysis in the $\gamma\mu\mu$ channel would have to find a way to confirm that the calorimeter response for the relevant photon kinematics is correctly simulated. For the χ_c, one could also use events where the photon converts into an electron-positron pair, as the momenta of charged particles can be better measured than those of photons. It is also a good strategy to perform an analysis in other decay channels, such as pairs of kaons. This would allow an independent determination of the cross section ratios for the χ_c spin states. There is no similar simple hadronic decay for the χ_b, and it may be quite difficult to measure DPE production of χ_b if the predictions are correct.

The background of proton dissociation and other inelastic processes can be better constrained if FSCs are installed at LHCb. If these backgrounds could be subtracted on a statistical basis, it would be possible to do away with the fits to the p_T-distribution for the purity, and it would also be possible to measure the p_T distribution of both the CEP signal and the inelastic background. The only remaining model dependence would be in the estimates for the efficiency.

BIBLIOGRAPHY

[1] The LHCb Collaboration, A. A. Alves *et al.*, JINST **3**, S08005 (2008).

[2] R. Antunes-Nobrega *et al.*, *LHCb reoptimized detector design and performance: Technical Design Report* (CERN, Geneva, 2003).

[3] O. S. Brüning *et al.*, *LHC Design Report: Volume 1, the LHC Main Ring* (CERN, Geneva, 2004).

[4] O. S. Brüning *et al.*, *LHC Design Report: Volume 2, the LHC Infrastructure and General Services* (CERN, Geneva, 2004).

[5] M. Benedikt, P. Collier, V. Mertens, J. Poole, and K. Schindl, *LHC Design Report: Volume 3, the LHC Injector Chain* (CERN, Geneva, 2004).

[6] L. Evans and P. Bryant, Journal of Instrumentation **3**, S08001 (2008).

[7] CERN, The latest from the LHC: Last period of proton running for 2010, News bulletin, http://cds.cern.ch/journal/CERNBulletin/2010/45/Newsn.

[8] The LHCb Collaboration, Performance of the LHCb Vertex Locator, (In preparation).

[9] The LHCb Collaboration, CERN Report No. LHCC-2001-011, 2001 (unpublished).

[10] S. Löchner and M. Schmelling, CERN Report No. CERN-LHCb-2005-105, 2006 (unpublished).

[11] F. Legger, A. Bay, G. Haefeli, and L. Locatelli, CERN Report No. CERN-LHCb-2004-100. 1-2, 2004 (unpublished).

[12] A. Affolder *et al.*, Radiation damage in the LHCb Vertex Locator, 2013, arXiv:1302.5259. CERN-LHCB-DP-2012-005.

[13] A. Vollhardt, CERN Report No. LHCb-2008-049, 2008 (unpublished).

[14] G. Adeva *et al.*, The LHCb Silicon Tracker: Running experience, in *Proceedings of TWEPP-12*, 2012.

[15] G. Cowan, Nucl.Instrum.Meth. A **699**, 156 (2013).

[16] D. van Eijk *et al.*, Nucl. Instrum. Methods Phys. Res., A **685**, 62 (2012).

[17] B. Storaci, The performance of the Outer Tracker detector at LHCb, in *Proceedings of ICATPP 13*, 2011, LHCb-PROC-2011-058.

[18] M. Adinolfi *et al.*, CERN Report No. CERN-LHCb-DP-2012-003. arXiv:1211.6759, 2012 (unpublished).

[19] N. Torr, Rich operations in 2012, Presentation at the LHCb-UK meeting, 3-4 Jan 2012, Oxford.

[20] R. Aaij *et al.*, Journal of Instrumentation **8**, P04022 (2013).

[21] I. Machikhiliyan, Journal of Physics: Conference Series **293**, 012052 (2011), LHCb-PROC-2010-030.

[22] A. A. Alves Jr *et al.*, CERN Report No. CERN-LHCb-DP-2012-002. arXiv:1211.1346, 2012 (unpublished).

[23] A. Cardini, G. Bencivenni, and P. De Simone, The operational experience of the Triple-GEM detectors of the LHCb muon system: Summary of 2 years of data taking, 2012, LHCb-PROC-2012-060.

[24] LHCb Operations Plots website: `https://lbweb.cern.ch/groups/online/OperationsPlots/index.htm`.

[25] R. Antunes-Nobrega *et al.*, *LHCb computing: Technical Design Report* (CERN, Geneva, 2005), Submitted on 11 May 2005.

[26] G. Barrand *et al.*, Comput.Phys.Commun. **140**, 45 (2001).

[27] R. Brun and F. Rademakers, Root - an object oriented data analysis framework, Proceedings AIHENP'96 Workshop, Lausanne, Sep. 1996, Nucl. Inst. & Meth. in Phys. Res. A 389 (1997) 81-86. `http://root.cern.ch/`.

[28] T. Sjöstrand, S. Mrenna, and P. Skands, Journal of High Energy Physics **2006**, 026 (2006), arXiv:hep-ph/0603175.

[29] GEANT4, S. Agostinelli *et al.*, Nucl.Instrum.Meth. **A506**, 250 (2003).

[30] G. Papotti *et al.*, Experience with offset collisions in the LHC, 2011, CERN-ATS-2011-147.

[31] R. Aaij *et al.*, CERN Report No. CERN-PH-EP-2013-128. LHCB-PAPER-2013-046. arXiv:1307.5024, 2013 (unpublished).

[32] The LHCb Collaboration, Model-independent measurement of CP violation parameters in $B^{\pm} \to (K^0_S h^+ h^-)_D K^{\pm}$ decays, 2013, LHCb-CONF-2013-004.

[33] The LHCb Collaboration, Improved constraints on γ from $B^{\pm} \to DK^{\pm}$ decays including first results on 2012 data, 2013, LHCb-CONF-2013-006.

[34] The LHCb Collaboration, Observation of B^0_s-$\bar{B}^0_s$ mixing and measurement of mixing frequencies using semileptonic B decays, arXiv:1308.1302 [hep-ex].

[35] Measurement of the cross-section for $Z \to \mu\mu$ production with 1 fb^{-1} of pp collisions at $\sqrt{s}$=7 TeV, 2013, LHCb-CONF-2013-007.

[36] R. Aaij *et al.*, J. High Energy Phys. **01**, 111 (2012).

[37] R. Aaij *et al.*, J. High Energy Phys. **02**, 106 (2012).

[38] T. Gershon, CERN Report No. CERN-LHCb-PUB-2012-009, 2012 (unpublished).

[39] Letter of intent for the LHCb upgrade, 2011, CERN-LHCC-2011-001. LHCC-I-018.

[40] LHCb Commissioning Group, First thoughts on online monitoring, CERN EDMS document 698614.

[41] A. Papadelis *et al.*, CERN Report No. LHCb-INT-2009-005, 2009 (unpublished).

[42] T. Szumlak and C. Parkes, Velo event model, 2006, CERN-LHCb-2006-054.

[43] T. Szumlak and C. Parkes, Description of the Vetra project and its application for the VELO detector, 2008, CERN-LHCb-2008-022.

[44] P. M. Bjørnstad and E. Rodrigues, CERN Report No. LHCb-INT-2010-043, 2010 (unpublished).

[45] M. Clemencic, LHCb - automated testing infrastructure for the software framework gaudi, Poster presented at CHEP'09, 2009.

[46] E. Rodrigues, CERN Report No. LHCb-INT-2010-012, 2010 (unpublished).

[47] F. Mandl and G. Shaw, *Quantum Field Theory, Rev.Ed.* (Wiley, 1993).

[48] I. J. R. Aitchison and A. J. G. Hey, *Gauge Thetheories in Particle Physics, (Graduate Student Series in Physics)* (Taylor & Francis, 1989).

[49] G. Aad *et al.*, Physics Letters B **716**, 1 (2012), arXiv:1207.7214 [hep-ex].

[50] S. Chatrchyan *et al.*, Phys. Lett. B **716**, 30 (2012), arXiv:1207.7235.

[51] Particle Data Group, J. Beringer *et al.*, Phys. Rev. D **86**, 010001 (2012).

[52] N. Brambilla *et al.*, EPJ C **71**, 1 (2011).

[53] R. Aaij *et al.*, Determination of the $X(3872)$ meson quantum numbers, arXiv:1302.6269 [hep-ex].

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

[55] 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, 2005).

[56] L. Motyka and G. Watt, Phys. Rev. D **78**, 014023 (2008), arXiv:0805.2113 [hep-ph].

[57] M. Albrow, T. Coughlin, and J. Forshaw, Progress in Particle and Nuclear Physics **65**, 149 (2010), arXiv:0805.2113 [hep-ph].

[58] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, and W. J. Stirling, EPJ **65**, 433 (2010).

[59] S. R. Klein and J. Nystrand, Phys. Rev. Lett. **92**, 142003 (2004), arXiv:hep-ph/0311164.

[60] W. Schäfer and A. Szczurek, Phys. Rev. D **76**, 094014 (2007).

[61] A. Schäfer, L. Mankiewicz, and O. Nachtmann, Physics Letters B **272**, 419 (1991).

[62] CDF Collaboration, J. Aaltonen *et al.*, Phys. Rev. Lett **102** (2009).

[63] The LHCb Collaboration, J. Phys. G **40**, 045001 (2013).

[64] M.-S. Chen, I. J. Muzinich, H. Terazawa, and T. P. Cheng, Phys. Rev. D **7**, 3485 (1973).

[65] B. Schrempp and F. Schrempp, Nucl. Phys. B **182**, 343 (1981).

[66] V. M. Budnev, I. F. Ginzburg, G. V. Meledin, and V. G. Serbo, Zh. Eksp. Theor. Fiz. Pis'ma Red. **12**, 349 (1970).

[67] J. Vermaseren, Nuclear Physics B **229**, 347 (1983).

[68] E. Berger, M. Diehl, and B. Pire, EPJ C **23**, 675 (2002), arXiv:hep-ph/0110062.

[69] B. Pire, L. Szymanowski, and J. Wagner, Phys. Rev. D **79**, 014010 (2009), arXiv:0811.0321 [hep-ph].

[70] V. Khoze, A. Martin, and M. Ryskin, Physics Letters B **401**, 330 (1997), arXiv:hep-ph/9701419.

[71] V. Khoze, A. Martin, and M. Ryskin, EPJ C **14**, 525 (2000), arXiv:hep-ph/0002072.

[72] V. Khoze, A. Martin, and M. Ryskin, EPJ C **23**, 311 (2002), arXiv:hep-ph/0111078.

[73] V. Khoze, A. Martin, and M. Ryskin, EPJ C **19**, 477 (2001), Erratum: C **20** 599 (2001); arXiv:hep-ph/0011393.

[74] V. Khoze, A. Martin, and M. Ryskin, Acta Phys. Polon. B **40**, 1841 (2009), arXiv:0903.2980 [hep-ph].

[75] L. A. Harland-Lang, V. A. Khoze, M. G. Ryskin, and W. J. Stirling, Latest results in central exclusive production: A summary, arXiv:1301.2552 [hep-ph].

[76] V. Khoze, A. Martin, M. Ryskin, and W. Stirling, EPJ C **35**, 211 (2004), arXiv:hep-ph/0403218.

[77] C. Yang, Phys. Rev. **77** (1950).

[78] R. Pasechnik, A. Szczurek, and O. Teryaev, Physics Letters B **680**, 62 (2009), arXiv:0901.4187 [hep-ph].

[79] A. Kaidalov, V. Khoze, A. Martin, and M. Ryskin, EPJ C **31**, 387 (2003), arXiv:hep-ph/0307064.

[80] M. Ryskin, A. Martin, and V. Khoze, EPJ C **60**, 265 (2009), arXiv:0812.2413 [hep-ph].

[81] A. Bialas and P. Landshoff, Physics Letters B **256**, 540 (1991).

[82] M. Boonekamp, R. Peschanski, and C. Royon, Nuclear Physics B **669**, 277 (2003), arXiv:hep-ph/0301244.

[83] A. Bzdak, Physics Letters B **615**, 240 (2005), arXiv:hep-ph/0504086.

[84] A. Bzdak, Physics Letters B **619**, 288 (2005), arXiv:hep-ph/0506101.

[85] F. Yuan, Physics Letters B **510**, 155 (2001), arXiv:hep-ph/0103213.

[86] R. S. Pasechnik, A. Szczurek, and O. V. Teryaev, Phys. Rev. D **78**, 014007 (2008), arXiv:0709.0857 [hep-ph].

[87] R. S. Pasechnik, A. Szczurek, and O. V. Teryaev, Diffractive production of $\chi_c(0^+, 1^+, 2^+)$ mesons at LHC, Tevatron and RHIC, in *PoS EPS-HEP*, 2009, arXiv:0909.4498 [hep-ph].

[88] M. Rangel, G. Alves, J. Barreto, R. Peschanski, and C. Royon, FPMC : a generator for forward physics, arXiv:1102.2531 [hep-ph].

[89] J. Monk and A. Pilkington, Computer Physics Communications **175**, 232 (2006), arXiv:hep-ph/0502077.

[90] CDF Collaboration, CDF non-published documentation website, 2013, `http://www-cdf.fnal.gov/upgrades/upgrades.html`.

[91] CDF Collaboration, T. Aaltonen *et al.*, Phys. Rev. D **77**, 052004 (2008), arXiv:0712.0604 [hep-ex].

[92] CDF Collaboration, A. Abulencia *et al.*, Phys. Rev. Lett. **98**, 112001 (2007).

[93] CDF Collaboration, T. Aaltonen *et al.*, Phys. Rev. Lett. **99**, 242002 (2007), arXiv:0707.2374 [hep-ex].

[94] CDF Collaboration, T. Aaltonen *et al.*, Phys. Rev. Lett. **102**, 222002 (2009), arXiv:0902.2816 [hep-ex].

[95] CDF Collaboration, T. Aaltonen *et al.*, Phys. Rev. Lett. **108**, 081801 (2012).

[96] V. Abazov *et al.*, Physics Letters B **705**, 193 (2011), arXiv:1009.2444 [hep-ex].

[97] G. Wolf, Reports on Progress in Physics **73**, 116202 (2010), arXiv:0907.1217 [hep-ex].

[98] J. Crittenden, Exclusive production of neutral vector mesons at the electron-proton collider HERA, 1997, arXiv:hep-ex/9704009.

[99] ALICE Collaboration, B. Abelev *et al.*, EPJ C (2012), 1208.4968.

[100] G. Aad *et al.*, EPJ C **72**, 1 (2012), arXiv:1201.2808 [hep-ex].

[101] S. Chatrchyan *et al.*, J. High Energy Phys. **01**, 052 (2011), arXiv:1111.5536.

[102] S. Chatrchyan *et al.*, J. High Energy Phys. **11**, 080 (2012), arXiv:1209.1666.

[103] The CMS Collaboration, Study of exclusive two-photon production of W^+W^- pairs in pp collisions at $7\,\mathrm{TeV}$, and constraints on anomalous quartic couplings, 2013, PAS-FWD-12-010; arXiv:1305.5596 [hep-ex].

[104] The LHCb Collaboration, Central exclusive dimuon production at $\sqrt{s} = 7\,\mathrm{TeV}$, LHCb-CONF-2011-022.

[105] The CMS Collaboration, S. Chatrchyan *et al.*, Journal of Instrumentation **3**, S08004 (2008).

[106] The CMS Collaboration, The CMS Collaboration, Measurement of pp diffraction cross sections at 7 TeV, CMS PAS FSQ-12-005, 2012.

[107] N. Brambilla *et al.*, *Heavy Quarkonium Physics* (CERN, Geneva, 2005), arXiv:hep-ph/0412158.

[108] M. Ryskin, A. Martin, and V. Khoze, EPJ C **60**, 249 (2009), arXiv:0812.2407 [hep-ph].

[109] L. Harland-Lang, V. Khoze, M. Ryskin, and W. Stirling, EPJ C **71**, 1 (2011), arXiv:1105.1626 [hep-ph].

[110] V. Khoze, A. Martin, and M. Ryskin, EPJ C **26**, 229 (2002), arXiv:hep-ph/0207313.

[111] L. Harland-Lang, 2012, Private communication.

[112] M. Albrow *et al.*, Journal of Instrumentation **4**, P10001 (2009).

[113] J. W. Lämsä and R. Orava, Journal of Instrumentation **4**, P11019 (2009), arXiv:0907.3847 [physics.acc-ph].

[114] R. B. Appleby and A. Mereghetti, The FLUKA model of IR8, 2010, CERN-LHC-Project-Note-427.

[115] A. Bhatti *et al.*, CDF Internal Note,/CDF/PUB/CDF/5247 **10** (2000).

[116] M. Albrow, Diffraction with forward shower counters FSC, `http://indico.cern.ch/conferenceOtherViews.py?view=standard&confId=91253`, Slides for presentation at LHC physics day “Diffraction”, May 2010.

[117] M. Giovannozzi, Optics options for the 2012 proton run, in *Proceedings of the Chamonix 2012 workshop on LHC Performance*, 2012, CERN-2012-006.

[118] LHC layout database, `http://layout.web.cern.ch/layout/` (accessible from CERN computers).

[119] LHC optics v6.500, `http://proj-lhc-optics-web.web.cern.ch/proj-lhc-optics-web/V6.500/`.

[120] M. G. Albrow *et al.*, Journal of Instrumentation **4**, T10001 (2009).

[121] L. Harland-Lang, V. Khoze, M. Ryskin, and W. Stirling, EPJ C **69**, 179 (2010), arXiv:1005.0695 [hep-ph].

[122] LHCb, E. Polycarpo Macedo *et al.*, Performance of the muon identification in LHCb with 2011 data.

[123] O. Deschamps, F. P. Machefert, M. H. Schune, G. Pakhlova, and I. Belyaev, CERN Report No. LHCb-2003-091, 2003 (unpublished).

[124] O. Callot, CERN Report No. CERN-LHCb-PUB-2011-001, 2011 (unpublished), LHCb.

[125] T. Skwarnicki, *A study of the radiative cascade transitions between the Upsilon-prime and upsilon resonances*, PhD thesis, Institute of Nuclear Physics, Krakow, 1986, DESY-F31-86-02.

[126] W. Verkerke and D. Kirkby, The RooFit toolkit for data modeling, `http://roofit.sourceforge.net`.

[127] M. Pivk, F. R, and L. Diberder, Nucl.Instrum.Meth **555**, 356 (2005), arXiv:physics/0402083 [physics.data-an].

[128] V. Belyaev, V. Egorychev, and D. Golubkov, CERN Report No. LHCb-INT-2012-001. CERN-LHCb-INT-2012-001, 2012 (unpublished).

[129] The LHCb Collaboration, JINST , P01010 (2012), arXiv:1110.2866 [hep-ex].