Chapter 2
The LHCb Detector

2.1 Introduction

The Letter of Intent for the LHCb experiment [1] at the Large Hadron Collider (LHC) was submitted in August 1995, following the approval of the LHC project by the CERN Council in December 1994. The LHC, a circular accelerator, was at the time foreseen to commence operation by colliding protons at a centre of mass energy of 10 TeV, with a later upgrade to 14 TeV. The LHCb detector was designed to make the most of the $b\bar{b}$ cross section, which peaks at high rapidity as shown in Fig. 2.1. A single-arm forward spectrometer with Ring Imagining Cherenkov counters and a silicon microvertex detector was envisaged. The site of this new experiment would be the cavern previously occupied by the DELPHI experiment at the Large Electron-Positron Collider, LEP.

The vision of the first pioneering collaborators was rewarded in 2010 with the accumulation of a substantial dataset for physics measurements. By mid 2011 the LHC machine was colliding protons in LHCb at a luminosity of $3.5 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, compared to the design value of $2 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, and in 2012 data were collected at luminosities of up to $6.0 \times 10^{32}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. The protons were packaged into bunches colliding at a maximum frequency of 20 MHz and a centre-of-mass energy of 7 TeV in 2011 and 8 TeV in 2012. The general structure of the detector is shown in Fig. 2.2.

Before reaching the magnet, the products of pp collisions encounter the silicon strip vertex detector (VELO), the first Ring Imaging Cherenkov detector (RICH1), and the first tracking station, the Tracker Turicensis[1] (TT). The magnet is followed by three further tracking stations, the second Cherenkov detector (RICH2), the first of five muon detectors, the electromagnetic then hadronic calorimeters, and finally the remaining four muon chambers. The instrumented region extends to around 20m in the $+z$ direction, parallel to the proton beams, from the interaction point, which forms the zero of the coordinate system. The $+y$ direction is upwards towards the top of the cavern while the $+x$ direction is to the left for an observer looking along the $+z$ direction.

[1] In Latin, *turicensis* pertains to Zurich.

H. Gordon, *Searches for CP Violation in Charmed Meson Decays*,
Springer Theses, DOI: 10.1007/978-3-319-07067-4_2,

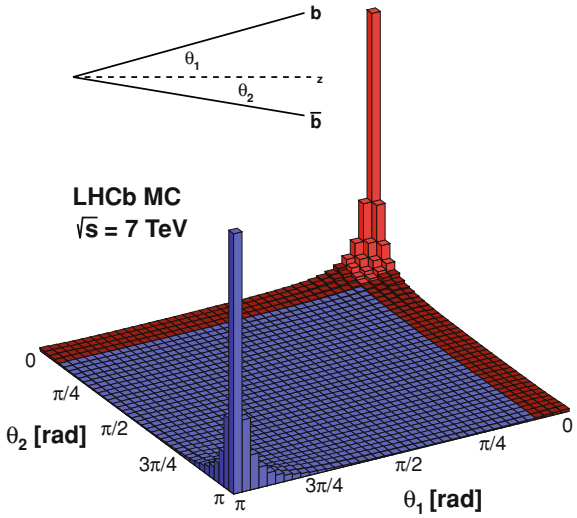

Fig. 2.1 Plot showing a simulation of the relative cross section for the production of b quark pairs as a function of the angles θ_1 and θ_2 they make with the beam axis (z) in pp collisions at $\sqrt{s} = 7$ TeV [2]. The cross section clearly peaks at small angles, which translate to high rapidities

A microvertex detector is essential to isolate the pp collision points from each other and from the displaced decay vertices of B and D decays. The tracking system provides good momentum resolution $\Delta p/p$ of between 0.40 and 0.55 %, which enables precision measurements of particle masses and analyses of Dalitz plots. Without excellent particle identification, decays of particles of similar mass to topologically identical final states would not be distinguishable. To avoid this, the RICHes test long-lived charged particles for consistency with kaon, pion and proton mass hypotheses. They perform well in the momentum range $5 < p < 100\,\text{GeV}/c$. On average, 96 % of kaons are identified by these detectors as kaons at the 95 % confidence level.[2] A hardware trigger uses the calorimeters and muon system to activate the data acquisition system (DAQ) at a rate of up to 1 MHz, and then software triggers reduce the data rate to between 2 and 4 kHz. The selected events are written to tape.

The detector covers an angular acceptance of ± 300 mrad in the bending plane of the magnet (the horizontal or x-direction) and ± 250 mrad in the non-bending plane (the vertical or $y-$direction). The main components are discussed in more detail in the following sections, mainly following the relevant Journal of Instrumentation paper, Ref. [3].

[2] Unfortunately, 95 % confidence is not usually enough for physics analysis, as Fig. 2.8 shows that around 10 % of pions will be identified as kaons at this level.

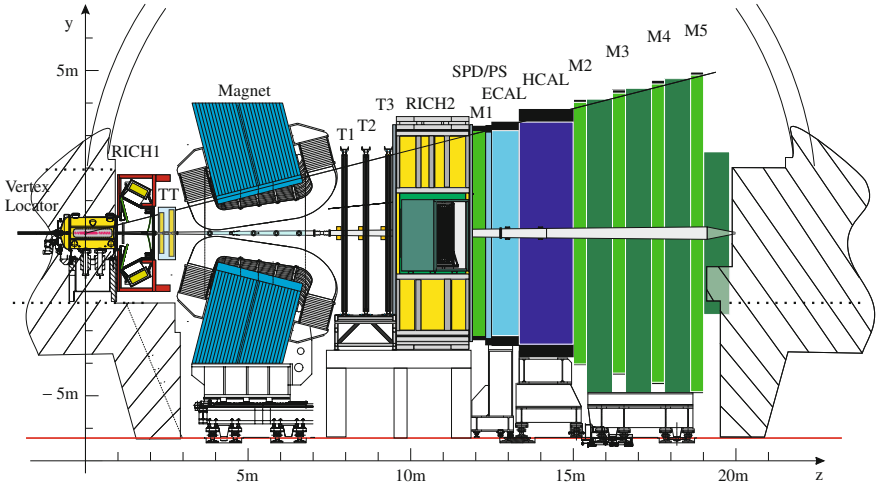

Fig. 2.2 Plan of the LHCb experiment [3]. The labels TT and T1-T3 indicate the tracking stations, M1-M5 the muon detectors, RICH1 and RICH2 the Ring Imaging Cherenkov detectors, SPD/PS the scintillator pad/preshower detector (see Sect. 2.6) and ECAL and HCAL the electromagnetic and hadronic calorimeters respectively

2.2 The LHC Machine

The LHC machine consists of 1232 superconducting dipole magnets in a tunnel 27 km in circumference and 3.7 m in diameter on the Swiss/French border. Protons made by ionising hydrogen atoms are directed through the CERN accelerator complex and are injected into the LHC from the Super Proton Synchrotron with an energy of 450 GeV. The protons are accelerated further in a straight section of the LHC tunnel between the ALICE and CMS experiments. The collisions analysed in this thesis had a centre-of-mass energy of 7 TeV. The bunches of protons cross at the ATLAS, CMS, ALICE and LHCb experiments. Various crucial parts of the LHCb detector, notably the outer tracker and the trigger system, are more effective at lower luminosities than at those used by ATLAS and CMS. Therefore the event multiplicity is limited by "luminosity levelling" to around two pp interactions per bunch crossing event. This is achieved by moving one beam relative to another so that at LHCb the beams do not collide head-on, and the cross sectional area of overlap of the beams is smaller. The effective value of the beam amplitude (β^*) can also be adjusted. The instantaneous luminosity L is given by

$$L = \frac{f N^2 S}{4\epsilon\beta^*} \tag{2.1}$$

where f is the frequency of pp collisions, N is the number of protons in the bunch, S is a geometrical factor, and ϵ is the emittance ($\epsilon\beta^* = \pi\sigma^2$ where $\pi\sigma^2$ is the cross sectional area of the beam). The emittance relates to the beam quality while the

amplitude relates to the optics. The geometrical factor S accounts for the displacement of the beams needed to level the luminosity, and also the slight increase in the size of the interaction region due to the crossing angle of the beams, which is not quite 180°. It is a testament to the good performance of the LHC machine that in 2012, with luminosity levelling in place, the benchmark target of 2fb^{-1} of data collected by LHCb per annum was surpassed, and around 22fb^{-1} of data were collected by ATLAS and CMS.

2.3 The VELO

The LHCb vertex locator is a series of two sets of 21 roughly semicircular silicon modules (or 'stations') distributed around the pp interaction region as shown in Fig. 2.3. The sets move inwards along the x axis to within 8 mm of the beam during data taking, but are retracted to 3.8 cm from the beam axis for protection when the beam is not stable. When the VELO is closed the two halves overlap slightly. Each module contains two silicon strip sensors: an R-sensor, which measures the radial distance from the beam, and a ϕ-sensor, which determines the azimuthal angle in the $x - y$ plane. Tracks must cross at least three VELO stations to be reconstructed. The stations are arranged to ensure that this requirement is usually satisfied for tracks in the LHCb acceptance.

The sensors are essentially p-n junctions, made of silicon enhanced by oxygen to improve the radiation tolerance. Doping the silicon with donor atoms such as phosphorus leads to an n-type material, with an excess of electrons. Conversely, a p-type material is created by doping with electron acceptors such as boron. A depletion region in which there are no free charge carriers is formed at the interface between p and n-type silicon because electrons near the interface tend to diffuse into the p-type silicon, leaving positive ions (holes) in the n-type region. The depletion region is enlarged on one side of the interface by applying a reverse biasing voltage. Electron-hole pairs created by the passage of a charged particle are separated by the electric field and move as an electrical current. The band gap is around 1.1 eV. When a high energy (minimum ionising) particle traverses a 300 μm thick silicon sensor, the Bethe-Bloch Equation [4] suggests that around 117 keV of energy will be lost, and the most probable number of electron-hole pairs generated, from the Landau distribution, is 23400. In the LHCb system, the depletion region is mostly in the n-type material, which comprises the bulk of the material of the sensor. The technology is referred to as oxygenated n^+- on -n. The junction is formed by an implant of n-type material with high dopant concentration from which the charge is read out to aluminium strips, located at the front of the depleted n-bulk, and an implant of p-type material situated at the back.

Within each module, the pitch or segmentation size of both R and ϕ sensors ranges from 38 μm at the inner edge to around 100 μm at the outer edge. As shown on the left of Fig. 2.4, the R sensors use concentric strips while the ϕ sensors use radial strips, with a discontinuity at a radius of 17.25 mm to avoid the pitch becoming

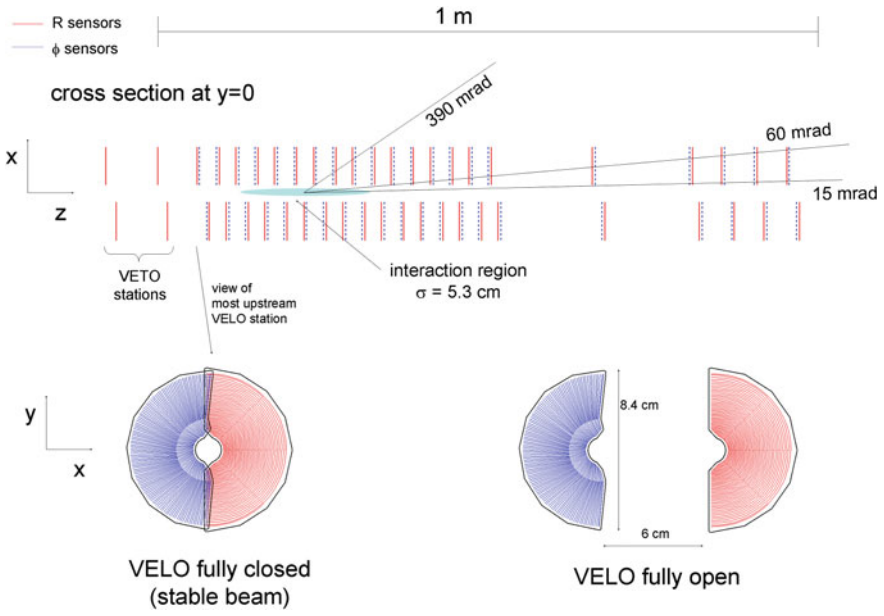

Fig. 2.3 Plan of the LHCb vertex locator, showing the distribution of modules in the z-direction [3]

too large at the outer edge of the modules. Apart from this discontinuity, the pitch increases linearly with radial distance along the ϕ sensors. The sensors are 300 μm thick and are connected with aluminium routing lines to 16 Beetle ASIC chips per sensor (32 per module). The 128-channel Beetle ASICs were designed to sample at the 40 MHz nominal LHC bunch crossing frequency and to read out the modules in a high radiation environment at the 1 MHz maximum rate of the hardware trigger. It takes 900 ns to read out one collision event. The analog front-end is formed by a low-noise charge-sensitive preamplifier, a pulse shaper and a buffer [5]. Events are pipelined, multiplexed, and read out via kapton cables to the data acquisition system (DAQ) via repeater boards.

The modules are enclosed within a vacuum vessel. Each detector half is placed inside a thin-walled aluminium box, and the boxes overlap when the VELO is closed. Aluminium was chosen for its short radiation length and good electrical conductivity. These "RF boxes" separate the primary LHC vacuum at 10^{-8}mbar from the secondary VELO vacuum of 10^{-4} mb, and shield the VELO modules against electromagnetic interference from the LHC beams. The inner face of each box is corrugated to fit around the modules. The structure of the modules is made from carbon fibre and highly thermally conductive pyrolitic graphite, which transfers as much as 32 W to Al blocks cooled by CO_2. This enables the VELO sensors to be kept at or below $-5\,^{\circ}$C at all times. The cooling improves the performance of the electronics inside the vacuum, and permits operation in the high radiation environment. Irradiation can lead to the displacement of atoms from their lattice sites or to charge build-up in

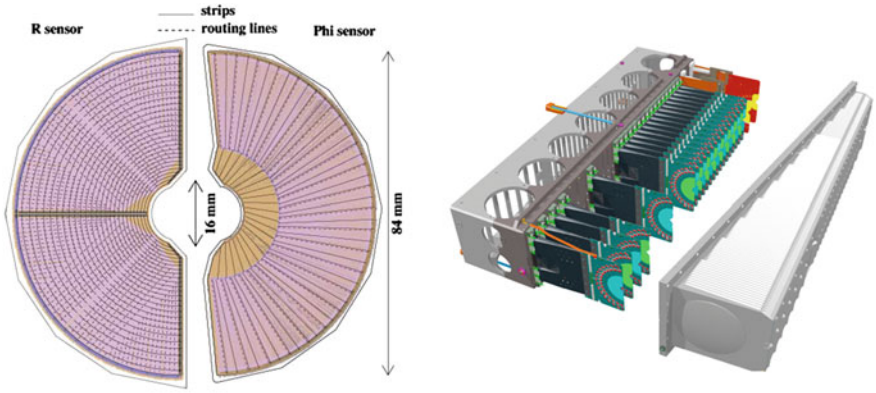

Fig. 2.4 The geometry of the VELO sensors (*left*) and an illustration of the modules that make up one half of the subdetector with the RF box that fits over them (*right*), from Ref. [3]. The two modules with the lowest z-coordinates, shown in red and at the extreme top-right of the figure, form the pile-up system, described in Sect. 2.10

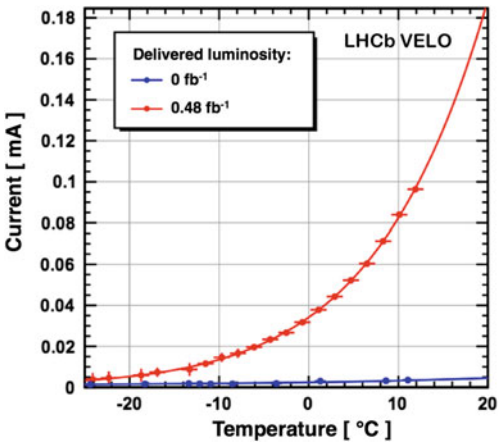

Fig. 2.5 The temperature dependence of the bias current for a typical VELO sensor with the nominal 150 V bias before and after around six months of LHC running, illustrating the effects of radiation damage [6]

surface layers. The vital importance of cooling the VELO after exposure to radiation is illustrated by the plot of the temperature dependence of the leakage current in Fig. 2.5, taken from a recent publication on radiation damage in the sub-detector [6]. The depletion voltage of the sensors changes significantly with irradiation, and the maximum depletion voltage the sensors are designed to cope with, of 500 V, is expected to be needed only after three years of operation.

A key indicator of the performance of the VELO is the resolution with which the impact parameter of a track with respect to a pp collision vertex can be measured. The impact parameter resolution, which is 20 μm for a track with a typical transverse momentum of 3 GeV/c, enables the lifetimes of most B and D mesons to be measured with uncertainties around of 50 fs, which compares favourably to the D^0 lifetime of 410 fs [7].

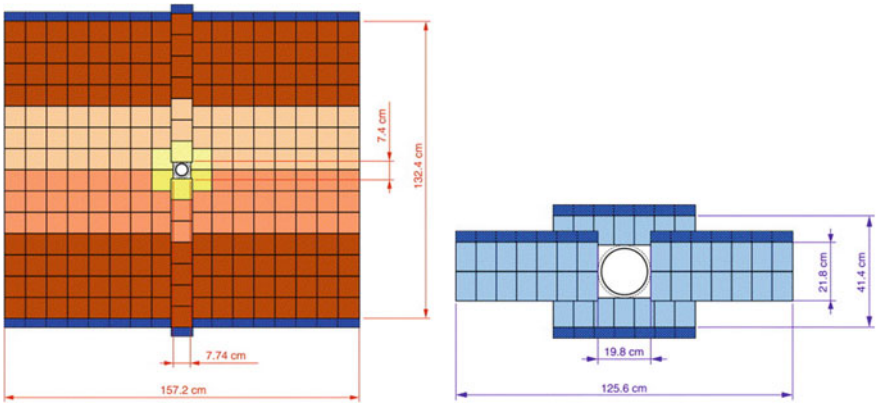

Fig. 2.6 Layout of one of the TT layers, *left*, and an IT module, *right* [8]. The sensors in the half-modules which make up the TT layer are shown

2.4 The Tracking System and Magnet

The tracking system consists of the silicon Tracker Turicensis (TT) located upstream of the magnet (closer to the pp interaction region and VELO), and three tracking stations (T1-T3) immediately downstream of it. The inner part of these three stations (the Inner Tracker, IT) is silicon, while the outer part (the Outer Tracker, OT) is made of straw-tube drift chambers. The silicon detectors were designed to achieve good spatial resolution of around 50 μm and high signal-to-noise ratios while minimising the occupancy, material budget, cost and signal shaping time in a moderately high radiation environment, but one much lower than that of the VELO.

The TT is constructed from four layers of modules that each cover the height of the LHCb acceptance. It contains 143360 readout strips covering an active area of 8.4 m^2. Each half-module consists of a row of seven silicon sensors arranged as shown on the left of Fig. 2.6. The architecture allows all infrastructure (front-end readout hybrids, cooling and support) to be kept outside the detector acceptance. The IT modules consist of either one or two silicon sensors that are connected via a pitch adapter to a kapton front-end readout hybrid. The silicon sensors are single sided p^+-on-n sensors. In the the IT, the sensors are 7.6 cm wide and 11 cm long, and carry 384 readout strips with a strip pitch of 198 mm. For the TT, they are similarly sized and carry 512 readout strips with a strip pitch of 183 mm as shown on the right of Fig. 2.6. The TT and IT are kept below 5 °C by aluminium and copper blocks cooled with C_6F_{14}. The IT is continually flushed with nitrogen to remove condensation.

The design of the OT was driven by the need to achieve momentum resolutions Δ_p/p of close to 0.4 % to resolve the mass of reconstructed B-hadrons to within 10 MeV/c^2. This necessitated a rigid design so that the position of the straw-tubes can be guaranteed to a precision of 100 μm in the x-direction, and good electrical shielding of each tube to ensure low cross-talk and noise. Each OT station consists

of two layers of drift-tubes with inner diameters of 4.9 mm, filled with a mixture of 70 % Ar and 30 % CO_2. The gases are chosen to be easily ionised by charged particles. Each drift tube is made from two strips of thin foils twisted together in a spiral. The inner (cathode) foil is made of carbon-doped polyimide (kapton) while the outer foil is a laminate of polyimide and aluminium. Electrons from the gas drift towards the anode. As the drift velocity v is known ($v = \mu E$ where E is the applied electric field and μ the mobility of the electrons), a measurement of the drift time with respect to the LHC clock allows the distance from the anode to be measured. The position of the ionisation event and hence the charged particle can thereby be resolved. The Al ensures fast signal transmission and good shielding, which is crucial to prevent cross talk and noise, whilst the polyimide is important to ensure the straw is gas-tight. By the end of 2011, typical momentum resolutions were around 0.5 % for B-hadrons [9].

Both TT and IT are read out with Beetle front-end chips. The OT is read out through a series of front-end boards in which analog amplification and shaping is followed by digitisation, which is synchronised with the LHC clock. The data is transferred to the DAQ if a positive decision is received from the hardware trigger. The reconstruction software starts with single hits in one part of a sub-detector, referred to as track seeds. These are extrapolated and combined with hits in other parts of the detector. In general, the tracks are seeded in the VELO or tracking stations where the magnetic field is low. The track candidates are fitted with a Kalman filter which accounts for multiple scattering and corrects for energy losses. The tracking efficiency of the detector is measured using a tag-and-probe technique with $J/\psi \to \mu^+\mu^-$ decays.

Momentum measurement at LHCb depends on a powerful magnetic field: the integrated magnetic field for tracks of 10 m length, from the VELO to the front of RICH2, is 4 Tm. However, it is also necessary to keep the magnetic field lower than 2mT inside the RICHes, because the image recorded by the photon detectors will be distorted if photoelectrons are deflected by the field. The magnet consists of a steel yoke surrounding two identical aluminium coils arranged with mirror symmetry about the horizontal plane as shown in Fig. 2.7. The magnetic field was carefully mapped using Hall probes to a precision of approximately 4×10^{-4} and was found to agree with expectations to within 1 %. The magnet polarity is reversible, so that certain detector asymmetry effects can be reversed and thereby measured and corrected for in physics analyses (see Chap. 3).

2.5 The RICHes

The two Ring Imaging Cherenkov detectors distinguish kaons from pions and protons. This is vital for separating the Cabibbo-favoured and suppressed decays described in Chap. 1, which is of paramount importance both in charm and B-physics, for example in the measurement of the CKM angle γ via the GLW/ADS methods [10]. The subdetectors rely on the Cherenkov radiation emitted at angle θ_c to the direction of motion of a particle travelling at velocity $v > c/n$ in a medium with refractive

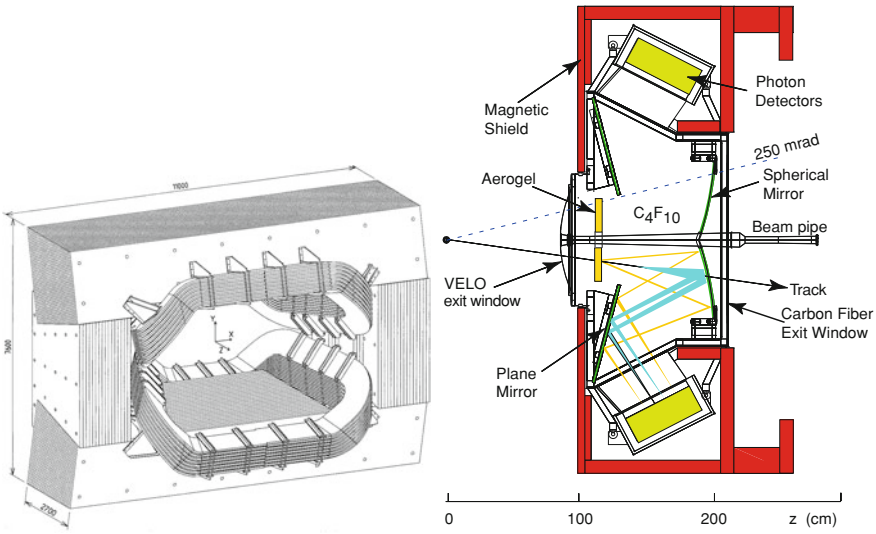

Fig. 2.7 The LHCb magnet, *left*, and the RICH1 detector, *right* [3]

index $n > 1$. The Cherenkov angle θ_c is given by $\cos\theta_c = c/nv$. The ring resolution is then proportional to $\sigma_\theta/\sqrt{N}$ where σ_θ is the uncertainty on θ_c and N is the number of photoelectrons in the ring [7].

The RICH1 detector provides particle identification (PID) in the momentum range $1 < p < 60$ GeV/c while RICH2 operates in a slightly reduced acceptance of ± 15 to ± 120 mrad in the horizontal plane and ± 100 mrad in the vertical plane, and identifies particles with momenta in the range $15 < p < 100$ GeV/c. In both cases, the instrumentation is kept out of the detector acceptance by reflecting the light out of the chamber to hybrid photon detectors (HPDs) with spherical and flat mirrors. HPDs are used instead of photomultiplier tubes due to their superior spatial resolution and response time. To minimise the material budget in RICH1, these mirrors are made from carbon fibre reinforced polymers (CFRP) coated with thin films of Al and MgF_2. In RICH2 they are glass.

RICH1 contains aerogel and C_4F_{10} while RICH2 contains CF_4. Aerogel consists of grains of amorphous silicon dioxide with sizes ranging from 1 to 10 nm linked together in a three-dimensional structure filled by trapped air [11]. It has a refractive index intermediate between that of a gas and a liquid. The CFC gases are chosen for their refractive indices, which are appropriate for the momentum spectrum of the decay products of charm and B particles. Given a reasonable quantum efficiency or number of primary photoelectrons generated per incident photon, a track passing through 5 cm of aerogel with refractive index $n = 1.03$ for light of wavelength 400 nm is expected to yield around 6.5 photoelectrons in a ring from a charged particle with $\beta = 1$. The corresponding yields for 95 cm of C_4F_{10} ($n = 1.0014$ at 400 nm) and 180 cm of CF_4 ($n = 1.0005$ at 400 nm) are 30 and 22 photoelectrons respectively.

A likelihood fit is used to determine whether the ring best matches the expectation from a kaon, pion, or proton hypothesis. A set of likelihood variables describing how consistent a particle is with being a kaon, pion or proton is made available for data analysis, and is described further in Chap. 3.

The gases and aerogel in the RICH detectors are contained within large Al boxes. The RICH1 box is 300 μm thick inside the acceptance on the upstream side, and contains an aperture filled with CFRP in the acceptance on the downstream side to reduce the material budget. As RICH2 is larger, the pressure differential is greater so two 1mm-thick carbon fibre and Al skins are separated by 30 mm of foam in the entrance and exit windows respectively.

The HPDs are arranged in a hexagonal pattern outside the RICHes. They are shielded from the magnetic field by large iron boxes. In RICH2 the stray magnetic field is reduced by these boxes from 15 mT to around 0.5 mT. In the HPDs, photoelectrons are released during the interaction of incident photons with the photocathode. They are accelerated by an applied voltage of 10 to 20 kV onto a reverse-biased, pixellated silicon anode, which results in the creation of an average of one electron-hole pair for every 3.6 eV of deposited energy. The HPDs are read out by integrated pixel chips. When tested, their quantum efficiency was found to be 5–10 % higher than the minimum 20 % required, and the "dark count rate" of fake signals from thermionic emission at the photocathode, ion feedback or electrostatic field emission was also satisfactory.

The silicon anodes of the HPD are bump bonded to the LHCbPIX1 binary readout ASIC and thence to the RICH electronics implemented on FPGAs. A series of algorithms converts the rings seen in the RICH to a likelihood for a kaon, proton, or pion mass hypothesis. The quantum efficiency and number of operational HPDs fluctuated with time during the 2011 run, and this was studied for the asymmetry measurements described in this thesis. The left plot of Fig. 2.8 shows that the performance of the detector slowly degraded through the 2011 data taking period. Furthermore, a difference in the particle identification efficiency between the halves of the RICHes on the left and right hand sides of the detector, when combined with the bending effect of the magnet on oppositely charged particles, produces a charge asymmetry in the particle identification efficiency. This is shown in the right plot of Fig. 2.8. These effects were found not to affect physics results significantly in 2011, largely because they can be neutralised by regularly reversing the polarity of the magnet, but they may cause complications in future. It should also be noted that the particle identification efficiency is strongly dependent on momentum: the kaon-pion discrimination is best in the range $10 < p < 50$ GeV and rapidly degrades outside this interval.

However, for hadrons with typical momenta, the particle identification power is excellent: Fig. 2.8 shows that kaons can be detected with a 95 % efficiency with a misidentification rate of 5 %. This is substantially better than the capability of any other LHC experiment.

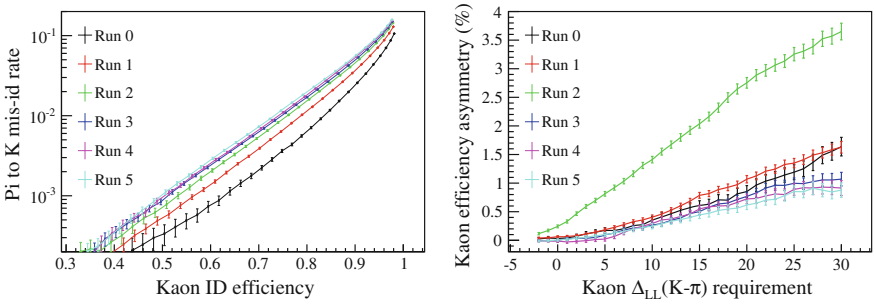

Fig. 2.8 *Left* the kaon ID efficiency plotted against the $\pi \rightarrow K$ misidentification rate for different requirements on the kaon likelihood variable. The highest probability for a kaon hypothesis, corresponding to the strictest likelihood requirement, is at the bottom-left corner of the plot. Each line corresponds to different running periods during 2011 for data taken with the magnet down. *Right* the charge asymmetry in the identification efficiency as a function of the likelihood requirement for the same running periods. A high value on the x axis corresponds to a strict requirement on the likelihood for the particle to be a kaon. The asymmetry in the efficiency is particularly large in Run 2, perhaps due to some faulty HPDs on one side of the detector. The reverse polarity shows the opposite asymmetry effect, indicating that the asymmetry results from left-right differences across the detector. The effect of these asymmetries on the analysis described in Chap. 4 is considered by plotting the charge asymmetries as a function of time in 2011 in Sect. 4.10.10

2.6 The Calorimeters

The calorimeters have the dual role of identifying and reconstructing photons, π^0 mesons and electrons, and measuring the transverse energy $E_T$of electron, photon and charged hadron showers for the hardware trigger. They operate by collecting scintillation light from particle interactions with dense material through optical fibres. There are three components: the preshower/scintillator pad detector (PS/SPD), the electromagnetic calorimeter (ECAL) and the hadronic calorimeter (HCAL). Different particles can be identified because they interact over different length scales, and give showers of light with different characteristic shapes.

The function of the PS/SPD is to distinguish electrons from charged hadrons and neutral pions. It consists of a 15 mm thick lead plate sandwiched between two layers of scintillator pads, before the ECAL. Charged particles, for example electrons, deposit energy in the first scintillator and can thus be distinguished from neutral particles, such as photons, which do not. Both electrons and photons then interact in the thin lead plate and the deposited energy is collected by the second scintillator. Hadrons have a longer interaction length and therefore they pass through without depositing very much energy (although the hits they create are retained). Around 99.6 % of pions do not deposit enough energy to meet the threshold to be identified as an electron, whilst at least 90 % of electrons with momenta above 10 GeV/c pass the threshold and can be detected. A correction is then made in the ECAL for the energy deposited in the lead plate.

The ECAL is a lead sampling scintillator. It is subdivided into inner, middle and outer sections of increasing cell size, a scheme also adopted in the PS/SPD. The total thickness of its 66 layers is 42 cm or 25 radiation lengths, enough for almost all the energy in an electromagnetic shower to be captured. This is required for good energy resolution. Each module is constructed from alternating layers of 2 mm thick lead, 120 μm thick reflecting paper and 4 mm thick scintillator tiles. When excited by the passage of a charged particle, the polystyrene scintillator molecules release a small fraction of the excitation energy as optical photons. The transverse shower size (Molière radius) of the stack of modules is 3.5 cm. The energy resolution $\sigma(E)/E$ was determined using a test beam and improves from 3 % to just over 1 % as the momentum increases from 15 to 100 GeV/c [12]. The hadronic calorimeter has similar structure but the absorber is iron rather than lead. It is divided into square cells of length 131.3 mm in the inner section and 262.6 mm in the outer section. Hadronic triggering does not require such good energy resolution, so to save space for the muon detectors the thickness of the HCAL is only 5.6 interaction lengths. The resolution is much worse than that of the ECAL, with $\sigma(E)/E$ varying as a function of momentum from 23 % at 15 GeV/c to 12 % at 100 GeV/c.

Each cell in both calorimeters is read out to an individual photomultiplier tube (PMT) by plastic wavelength-shifting fibres. The PS/SPD photomultipliers have multiple anodes, each of which covers several cells. The PMTs are encased in iron to shield them from stray magnetic fields. A pulse shaper takes the output of these tubes and removes the tail of any pulses extending beyond 25 ns, so that every LHC bunch crossing can be measured and potentially activate the trigger. The front end board receives the resulting pulses and digitises them. The digital signals are refined slightly by FPGAs to reduce noise and then passed to the calorimeter readout cards in crates above the detector. These cards generate the final L0 trigger signal and pass the information on to the DAQ system.

2.7 The Muon System

Muons are of fundamental importance to LHCb, as the easiest way to trigger the detector is on muons with high transverse momentum p_{T}, and because muons form part or all of the final state in flagship decay channels such as $B_s \to J/\psi\phi$ and $B_s \to \mu^+\mu^-$.

LHCb has five muon stations, one upstream in front of the calorimeters, and four downstream, behind them as shown on the left of Fig. 2.9. The first station is placed upstream to improve the p_{T} resolution. The whole system consists of 1368 multi-wire proportional chambers (MWPCs) and 12 sets of three gas electron multiplier foils (referred to as triple-GEMs) in the region closest to the beam pipe in the most upstream station where the particle flux is highest. The downstream stations alternate with 80 cm thick iron absorbers. The detectors are divided into cells. Each cell provides a binary decision to the trigger, which requires aligned hits above the discriminator threshold in all five stations to fire. This means the efficiency of each

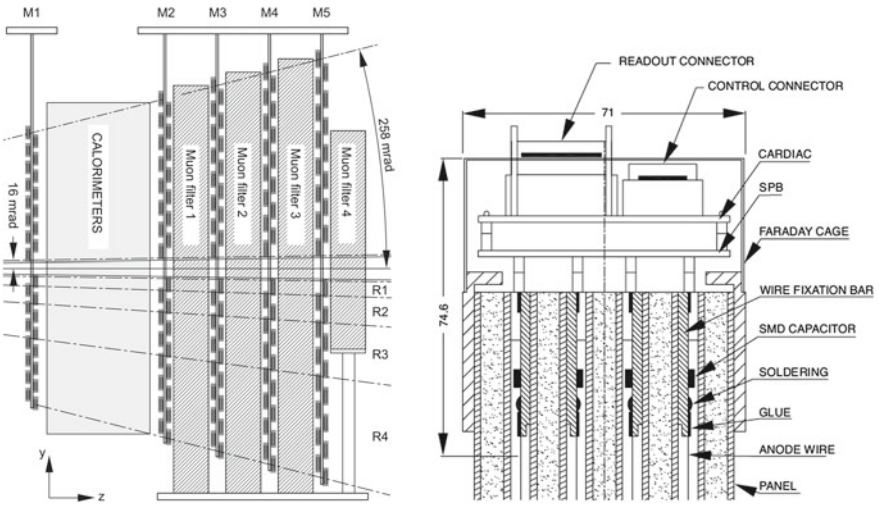

Fig. 2.9 Diagram showing (*left*) the positions of the five LHCb muon stations, the calorimeters and the iron absorbers [13] and (*right*) the cross section through a wire chamber with the four gas gaps and the connection to the readout electronics. The SPB stands for the spark protection board while CARDIAC is the name of the front-end electronics [3]

station must exceed 95 %. The spatial resolution is defined by the cell size, and is most precise for the first three stations in the bending plane of the magnet. The timing resolution is sufficient to distinguish consecutive LHC bunch crossings at 40 MHz.

Each triple-GEM detector consists of three copper-cladded kapton foils sandwiched between an anode and a cathode in a mixture of Ar, CO_2 and CF_4 in ratio 45:15:40. This combination of gases was found to give the detector an excellent time resolution of 3 ns. As the muons pass through the system they ionise the gas, and the ionisation electrons are attracted by the electric field towards the anode. The electrons that enter 70 μm wide holes in the foils are multiplied by the large electric field. The holes are 140 μm apart.

The MWPCs operate on a similar principle to the GEMs, except that a wire mesh is used instead of a perforated foil. They consist of two layers of closely spaced parallel anode wires between two parallel cathode plates as shown on the right of Fig. 2.9. The wires in one layer are perpendicular to those in the other. Ionisation electrons are attracted to the nearest wire, and hence the position of the particle can be resolved in two dimensions. Close to the wires, the electric field increases, which accelerates the electrons and leads them to cause secondary ionisation, a process known as avalanching. This avalanching is proportional to the energy of the incident particle. Thus the MWPC, in principle, allows a measurement of the particle energy, although in LHCb this is not used. Each MWPC in the four most downstream muon stations contains four gas gaps with wires connected in OR, to optimise the efficiency, while the other muon station contains only two such gaps (or two triple-GEMs) in OR to minimise the material budget. In the MWPCs the time resolution is 5 ns. The

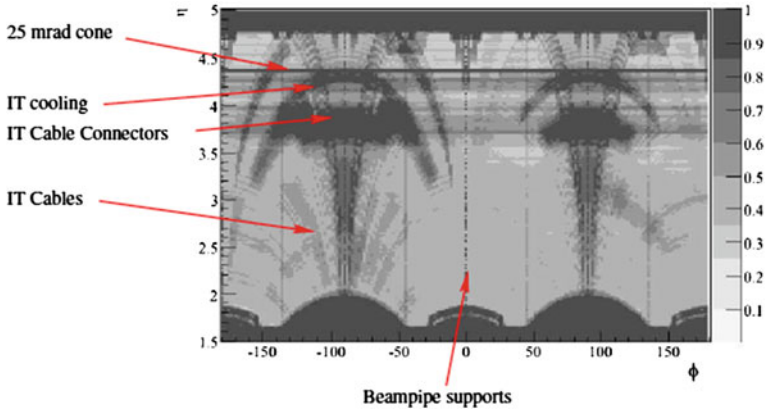

Fig. 2.10 Radiation length in the LHCb detector as a function of pseudorapidity η and azimuthal angle around the beam axis ϕ, integrated up to $z = 930$ cm, with the effect of the IT labelled. The radiation length scale is truncated at one radiation length [14]

same gases are used as in the triple-GEMs, but with ratio 40:55:5. The chambers and front-end electronics are enclosed in a Faraday cage.

The muon system is read out into front-end boards which amplify, time-align and shape the chamber signals. Logical ORs of the channels are taken before the signal is tagged with a code identifying the relevant pp bunch crossing and routed to the DAQ.

2.8 Material Budget

Multiple scattering limits the momentum resolution of the detector, so it is extremely important to minimise the material budget before the calorimeters. The material in the VELO makes up 17.5 % of a radiation length on average, with the largest contribution coming from the RF-foil. There is a strong variation in the amount of material with the azimuthal angle and radial distance from the beam pipe. The material in the detector up to the beginning of RICH2, including the entire tracking system, is shown in Fig 2.10. On average, particles pass through around 60 % of a radiation length and 20 % of an absorption length before entering RICH2 [14].

2.9 The DAQ and Electronics

The LHCb detector is read out from analog front-end chips such as the Beetle or LHCbPIX1. The signals are digitised and sent to standardised readout boards known as TELL1s via optical links. The exception is the VELO, which sends analog

information to TELL1s directly via kapton cables. Four pre-processing FPGAs perform common-mode rejection, zero-suppression and data compression as required. The resulting data fragments are collected by a fifth FPGA and formatted into a raw IP-packet that is subsequently sent to the Event Filter Farm (EFF) via 4Gb Ethernet ports. The EFF is the set of CPUs that perform the processing for the software trigger. The raw data selected by the trigger in the EFF is transferred to the CERN Tier-0 computing centre and subsequently distributed via the Large Computing Grid. The online system is synchronised with the Timing and Fast Control system (TFC) and operated via the Experiment Control System (ECS). The TFC ensures that the LHC machine clock and the trigger decisions are synchronised and that the trigger functions at an acceptable rate. The ECS ensures proper control and monitoring of all parts of the detector and is run by users in the control room with PVSS software.

2.10 The Hardware Trigger

The LHCb trigger operates in three stages. The Level-0 trigger (L0), implemented in custom hardware, operates at the LHC bunch crossing frequency and reduces the rate from 40 to 1 MHz. The High-Level Trigger (HLT1 and HLT2) is software based, and further reduces the rate to 2 or 3 kHz by reconstruction of tracks (in HLT1) and partial or full decay chains (in HLT2). The L0 trigger Decision Unit (L0DU) receives input via 24 high-speed optical links from three subdetectors at 40 MHz: the calorimeters and scintillator pad detectors, which trigger on electrons, photons or hadrons, the muon system, and, in the original design, the pile-up detector. The pile-up system, consisting of four VELO R-sensors upstream of the rest of the VELO, was to be used to locate the z-positions of primary vertices to measure the track multiplicity. This could enable the trigger to infer the probable number of pp interactions in a given bunch crossing and to remove overly complex events, but to date this subsystem has not been used except to measure the luminosity with beam-gas interactions [15, 16]. The L0 trigger decisions from each subdetector are combined in a logical OR to obtain the L0 decision. The L0DU itself is based on a TELL1 board. It instructs the DAQ system whether or not to read out the rest of the detector 4 μs after an interaction.

The calorimeter trigger calculates the E_{T} of 2×2 groups of cells. These usually contain most of the energy deposited by a particle, but are too small to encompass overlapping showers. The highest E_{T}combination is found. The cause of the deposit, be it an electron, photon or hadron, is identified. A hadron should shower only in the HCAL and sometimes the ECAL, a photon should additionally leave a hit in a preshower cell, and an electron should leave hits in both a preshower and an SPD cell [16], as outlined in Sect. 2.6. The energy of this candidate, together with the total E_{T}in the HCAL and the number of hits in the SPD, is sent to the L0DU.

The muon detector is split into four and the muon tracks in each quadrant with the largest and second largest p_{T} are identified. Track finding starts from seed hits in the third muon station, and a track is formed if hits in all five muon stations can be

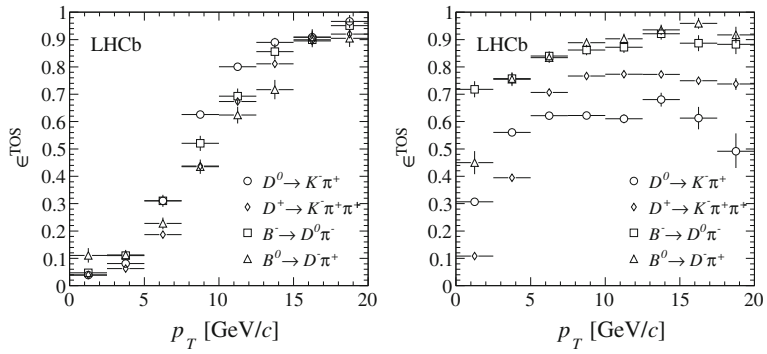

Fig. 2.11 The efficiency of the hadronic level-0 trigger (*left*) and the 1-track inclusive software trigger (*right*) as a function of p_T for important B and D decay channels [16]

aligned. The dedicated processors for the muon system each contain 192 elements (based on a spatial division of each quadrant), and each run 96 tracking algorithms in parallel. Up to eight high-p_T track candidates (up to two from each quadrant) are sent to the L0DU.

It is possible for several trigger systems to give a positive signal on any given event. For example, a B or D meson could decay to several hadrons, of which one or more could pass the threshold to fire the calorimeter trigger, or a muon and some hadrons. Or alternatively two different decays from the same event could each result in a track which fires one of the triggers. If an event is triggered, all of the information about how this happened is saved to disk.

In this thesis, the hadronic Level-0 trigger systems (L0Hadron) are of particular importance, because these can be activated by the signal decay channels. The efficiency of the hadronic calorimeter trigger is shown as a function of p_T for sample decay channels in Fig. 2.11. The poor resolution is apparent: one would expect a fairly clear cut-off with p_T since the trigger makes its decision based on the transverse energy of the particle. However, the resolution is not the only reason for the smearing evident in the plot: the trigger has no information about the charge of the particle. Therefore the measured transverse energy of the particle depends strongly on which way it bends in the magnet. Particles of one charge, pointing into one one side of the detector, will bend outwards for one magnet polarity and inwards for the other. If they bend inwards they are far less likely to have the threshold E_T required to activate the trigger. Therefore this trigger introduces a large left-right asymmetry for particles of a given charge which can be reversed by changing the magnet polarity. This must be considered when measurements of charge asymmetries for *CP* violation studies are attempted.

The threshold E_T for particles to activate the hadron hardware trigger was changed at various points through 2010 and 2011, but was usually around 3.5 GeV. Furthermore, the number of hits in the SPD was generally required to be below 600. Changes in the thresholds during the run reflect variations in the mean number of pp

Table 2.1 Main running periods in 2011 with the magnet polarity and trigger configuration key

Run interval	Main TCK	Moore	Magnet polarity	Int. Lumi (pb^{-1})
87219-87977	0x360032	v12r3	Down	7.4
89333-90207	0x5A0032	v12r5	Up	43.0
90256-92000	0x5A0032	v12r5	Down	40.2
92001-92735	0x6D0032	v12r6p1	Down	68.1
92821-93282	0x730035	v12r6p1	Down	84.3
93398-94386	0x730035	v12r6p1	Up	138.9
95948-97028	0x760037	v12r8	Up	70.3
97114-97992	0x760037	v12r8	Down	74.2
98002-98882	0x760037	v12r8	Down	117.4
98900-100256	0x760037	v12r8	Up	36.7
101373-101862	0x790037	v12r9p1	Up	64.9
101891-102452	0x790038	v12r9p1	Down	97.0
102499-102907	0x790038	v12r9p1	Up	82.9
103049-103686	0x790038	v12r9p1	Down	112.4
103686-104414	0x790038	v12r9p1	Up	46.5

In practice, the trigger configurations used for all of these keys are very similar, or identical. The horizontal lines correspond to the division into running periods described in Sect. 4.10.10, indicated by the vertical lines in Fig. 4.40

interactions per bunch crossing. In principle, this could mean that data taken with one or the other magnet polarity could have a lower average threshold. This would affect how well various detector asymmetries discussed in Chap. 4 cancel when data taken with the two polarities are combined, leading to systematic biases on charge asymmetry measurements. No significant effect is observed, probably because the thresholds were kept constant for the majority of the data taking period. The configurations used in 2011 (where the statistical sensitivities of physics analyses to such issues are higher), and their relation to changes of the magnet polarity are summarised in Table 2.1.

2.11 The Software Trigger

The High Level Trigger (HLT) is a series of algorithms implemented in a C++ software package called Moore which runs on 26110 CPU cores in the Event Filter Farm [16]. In 2011, events were selected by the hardware triggers at an average rate of 870 kHz, slightly below the design rate of 1 MHz . The first stage of the HLT, referred to as HLT1, reconstructs tracks starting from a three-dimensional pattern recognition which uses information from the VELO. The VELO tracks are continued through the tracking system to the OT and IT and a χ^2 for track quality is calculated from the Kalman filter track fit. For events selected by the muon hardware trigger, an attempt is made to match the VELO tracks with hits in the muon chambers. Accurate measurements of the momentum and p_{T} of hadrons and muons are performed: the

invariant mass resolution of $J/\psi \to \mu^+\mu^-$ candidates was measured to be only 3 % larger in the HLT than it is in the offline reconstruction. In most of 2010 and 2011, muonic charm and beauty decays were usually triggered by the presence of a high p_T muon. For most other decays, the inclusive displaced track trigger (Hlt1Track) requires one hadron, muon or electron track with a p_T greater than 1.7 GeV/c, a momentum greater than 10 GeV/c, a track fit χ^2/ndf < 2 and an impact parameter with respect to the PV of at least 0.1mm. This last requirement is highly efficient for the daughters of long-lived B and D^+ mesons, although less so for those of the shorter-lived D^0 and D_s^+.

The second stage of the high level trigger consists of a series of selections designed to reconstruct particular decay modes or groups of decays. The full event information is used. For charm decays, the process is very similar to the offline reconstruction, except that no particle identification is used in the selection of tracks. Sometimes, to reduce the bandwidth, selections are prescaled. This is the case for one of the triggers used in the analysis described in Chap. 3. The prescaling means that the algorithms are only enabled for a fraction of the events for which a trigger decision has to be made.

As the HLT is purely software, it can evolve as more experience of data-taking is acquired and the priorities of the experiment change. In 2011, the HLT had approximately 30 ms to make each decision. Adjustable parameters in the algorithms are controlled by Trigger Configuration Keys (TCKs). More specific information on the trigger for the 2010 run can be found in Chap. 3.

2.12 Event and Detector Simulation

A full Monte Carlo simulation of proton-proton collisions and their interaction with the LHCb detector exists, and it is used principally to optimise event selections and estimate efficiencies. The pp collisions are simulated with PYTHIA 6.4 [17] with a specific LHCb configuration [18]. The most frequent production mechanisms for b and c quarks are the leading order $2 \to 2$ QCD processes reproduced from Ref. [18] in Figs. 2.12a, b. However, other non-negligible processes such as flavour excitation or gluon-gluon splitting are also simulated. To generate standard 'minimum bias' events, diffractive and elastic processes are also activated. In the LHCb configuration, the track multiplicity and fraction of excited states are adjusted in accordance with the results of measurements at LEP and CDF.

Each simulated event is usually forced to contain the particular decay mode under study. Decays of hadronic particles, including, for example, representations of their Dalitz plot distributions and simulations of K, D or B mixing are described by EVTGEN [19]. Selection cuts are applied to the generated decays to save CPU time in later stages of the simulation. Usually, the daughters of B or D hadrons are simply required to be contained inside the detector acceptance. In the next phase, the interactions of the generated particles with the detector and its response are simulated using the GEANT4 toolkit [20, 21] as described in Ref. [22]. An XML geometry

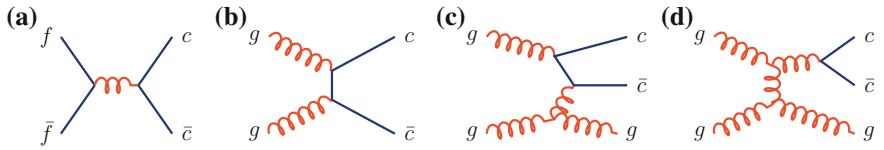

Fig. 2.12 Processes for $c\bar{c}$ production implemented in Pythia 6: tree-level processes at lowest order (**a**, **b**) and higher order (**c**, **d**) [18]

database is used to store the detector description. Together with information on the detector alignment stored in the Conditions Database (also XML), this encodes all of the information needed to model the detector. The simulation is managed by the Gaudi-based [23] Gauss software application. The detector response is converted into the same format as that produced by the detector readout by another application, named Boole, so that the simulated events can be reconstructed in the same way as real pp collisions. The only difference between data and simulation in the context of the software is that in simulated data information about the generated events is also stored and can be matched to reconstructed signal when the data are analysed.

2.13 Reconstruction and Analysis Software

Raw data from the LHCb DAQ system is processed using the Brunel software package. This handles pattern recognition, track fitting, reconstruction of Cherenkov rings in the RICH system, association of reconstructed particles to 'true' information from the generator in the case of simulated events, and performance monitoring. A raw pp collision event is stored as a collection of 'raw banks', each of which contains the information from one TELL1 board. The banks are converted to new data structures stored in the so-called Transient Event Store or TES. This is a dynamic database of event information which serves as the source of particles for the reconstruction of individual decay channels. This reconstruction is handled by the DaVinci analysis framework. Particles satisfying the criteria to be charged pions or kaons, protons, muons, electrons, or photons are taken from the TES and combined to form composite candidates. Selection criteria are applied to isolate signal from background. This is performed first in the HLT2 algorithms, and then again offline in so-called 'stripping' algorithms. More CPU time is available in the stripping step, and therefore a wider range of selection criteria than that available in the trigger can be applied and the vertex fitting can be done more accurately. For example, particle identification information from the RICH detectors is available to the stripping algorithms, but not usually in the HLT. The output of the stripping selections is distributed on the Large Computing Grid for offline analysis. While the raw events selected by the trigger are also stored, they are not made accessible for analysis in order to avoid the need to store many replicas of the same data and to avoid the need for each user to 'strip' the data themselves. For charm decays, to save disk space, the files available

for analysis contain only the tracks of the charm decay selected by the stripping algorithm, instead of the full event. The raw events are periodically 'restripped' to allow data to be reanalysed with different signal selection criteria.

For each candidate that is selected by a stripping algorithm, some information about how the trigger algorithms selected the event is stored by the software. If the event was triggered by a daughter of the candidate that the HLT2 algorithm in question was designed to select, the candidate is labelled as 'TOS', meaning 'triggered on signal'. If, by contrast, the event was triggered by a particle which does not form part of the signal candidate, then the candidate is labelled as 'TIS', meaning 'triggered independently of signal'. Candidates can be both TIS and TOS, if several triggers were activated by the same event, or neither, if they cannot be satisfactorily categorised. If they are neither TIS nor TOS, the trigger was probably activated by a partially reconstructed signal candidate. Such events should not be used in data analysis. Furthermore, 'TIS' candidates are not always fully 'independent of the signal' because the track which activated the trigger can be correlated to the signal candidate. Nevertheless, this categorisation is an invaluable analysis tool, as will become apparent in Chaps. 3 and 4.

The LHCb analysis framework is continually evolving as running conditions change and our understanding of the interaction of particles with detector material advances as more demanding analyses are performed. The relative stability of the trigger in 2011 and 2012 compared to 2010 has enabled closer study of its properties, described in part in Chap. 4. However, the good performance of the detector and the well-developed analysis software also made it possible to produce an impressive catalogue of physics results from the first data taking period in 2010, and one example is presented in the next chapter.

References

1. LHCb collaboration, Nakada T et al (1995) Letter of intent: a dedicated LHC collider beauty experiment for precision measurements of CP violation. In: CERN, LHCC, pp 95–5
2. Elsasser C., Production of *b* hadrons at the LHCb experiment, private communication, 2013
3. LHCb collaboration, Alves AA Jr et al (2008) The LHCb detector at the LHC. JINST 3:S08005
4. Bethe H (1930) Zur theorie des durchgangs schneller korpuskularstrahlen durch materie. Annalen der Physik 397:325
5. Agari M et al (2004) Beetle: r radiation hard readout chip for the LHCb experiment. Nucl Instrum Meth A518:468
6. LHCb VELO collaboration, Affolder A et al (2013) Radiation damage in the LHCb vertex locator. JINST 8 P08002. doi:10.1088/1748-0221/8/08/P08002, arXiv:1302.5259
7. Particle Data Group, Beringer J et al (2012) Review of particle physics. Phys Rev D86:010001
8. Steinkamp O (2005) Silicon strip detectors for the LHCb experiment. Nucl Instrum Meth A541:83
9. LHCb collaboration, Aaij R et al (2013) Implications of LHCb measurements and future prospects. Eur Phys J C73:2373. arXiv:1208.3355
10. LHCb collaboration, Aaij R et al (2012) Observation of CP violation in $B^+ \to DK^+$ decays. Phys Lett B712:203. arXiv:1203.3662
11. Nappi E (1998) Aerogel and its applications to RICH detectors. Nucl Phys Proc Suppl 61B:270

12. Arefev A et al (2008) Beam test results of the LHCb electromagnetic calorimeter, Technical report LHCb-2007-149, CERN, Geneva
13. LHCb collaboration, The LHCb public website, 2013
14. Needham M, Ruf V (2007) Estimation of the material budget of the LHCb detector, Technical report LHCb-2007-025, CERN, Geneva
15. LHCb collaboration, Aaij R et al (2012) Absolute luminosity measurements with the LHCb detector at the LHC, JINST 7, P01010. arXiv:1110.2866
16. Aaij R et al (2013) The LHCb trigger and its performance in 2011. JINST 8:P04022 arXiv:1211.3055
17. Sjöstrand T, Mrenna S, Skands P (2006) PYTHIA 6.4 physics and manual, JHEP 05:026. arXiv:hep-ph/0603175
18. Belyaev I et al (2011) Handling of the generation of primary events in Gauss, the LHCb simulation framework, nuclear science symposium conference record (NSS/MIC). IEEE, New York, pp 1155
19. Lange DJ (2001) The EvtGen particle decay simulation package. Nucl Instrum Meth A462:152
20. GEANT4 collaboration, Allison J et al (2006) Geant4 developments and applications. IEEE Trans Nucl Sci 53:270
21. GEANT4 collaboration, Agostinelli S et al (2003) GEANT4: A simulation toolkit. Nucl Instrum Meth A506:250
22. Clemencic M et al (2011) The LHCb simulation application, Gauss: design, evolution and experience. J Phys Conf Ser 331:032023
23. Barrand G et al (2000) GAUDI : the software architecture and framework for building LHCb data processing applications. in: International conference on computing in high energy and nuclear physics (CHEP 2000), pp 92–95

http://www.springer.com/978-3-319-07066-7

Searches for CP Violation in Charmed Meson Decays
A Study of D+ → K − K+ π+ at the LHCb Experiment
Gordon, H.
2014, XV, 158 p. 102 illus., 21 illus. in color., Hardcover
ISBN: 978-3-319-07066-7