Preprint typeset in JINST style - HYPER VERSION

The LHCb Detector at LHC

LHCb Collaboration

ABSTRACT: This is a description of the LHCb detector......

Contents

---

1. Introduction

1.1 Physics motivations and requirements

LHCb is the experiment dedicated to b physics at the LHC [1],[2]. The aim is the high precision study of CP violation and rare decays. The experiment will extend the B-physics results obtained at B factories and Tevatron: it will overconstrain the unitarity triangle and allow the search for new physics. Although observed CP-violating phenomena are consistent with the CKM mechanism, it cannot be excluded that physics beyond the Standard Model contributes, or even fully accounts for the observed phenomena.

The level of CP violation that can be generated by the Standard Model weak interaction is insufficient to explain the dominance of matter in the universe. This also calls for new sources of CP violation beyond the Standard Model.

Another way to search for physics beyond the Standard Model is to study B-meson decays that are rare or even forbidden in the Standard Model. In addition to investigating CP violation in B-meson decays, the physics programme of the LHCb experiment will include studies of rare B and τ decays, $D-\overline{D}$ oscillations and heavy b-hadrons (e.g. B_c-meson) production and decays.

There are many ways to look for a sign of new physics. In all cases, large numbers of both B^0_s and B^0_d mesons are required, and many different decay modes have to be reconstructed. Compared to other existing accelerators that are in operation the LHC will be by far the most copious source of B mesons, due to the high $b\overline{b}$ cross section, $\sigma_{b\overline{b}}$ = 500 μb and high luminosity. A variety of b-hadrons, such as B_u, B_d, B_s, B_c and b-baryons, will be produced with high rate.

The LHCb experiment plans to operate with an average luminosity of $2 \times 10^{32}\ cm^{-2}\ s^{-1}$, which should be obtained from the beginning of LHC operation. Running at this luminosity has the advantages that the detector occupancy remains low, and radiation damage is reduced. Events are dominated by single pp interactions per bunch crossing that are simpler to analyse. The luminosity at the LHCb interaction point can be kept at its nominal value while the luminosities at the other interaction points will be progressively increased to their design values. This will allow the experiment to collect data for many years under constant conditions. About 10^{12} $b\overline{b}$ pairs are expected to be produced in one year of data taking (expected to correspond to about 2 fb^{-1}).

The LHCb detector is designed to exploit this large number of b-hadrons produced at the LHC in order to make precision studies of CP asymmetries and of rare decays of the b-hadrons. It has a high-performance, efficient and flexible trigger which is robust and optimised to collect B hadrons, based on particles with large transverse momentum and displaced decay vertices.

High quality event reconstruction is necessary, based an excellent particle identification, tracking and vertexing. Particle idetification must allow to distiguish muons, electrons, as well as gammas or π^0, and charged pions and kaons. The latter is necessary for the study of the hadronic decay channels among which $B_d \to \pi^+\pi^-$, $B_s \to D_s^{\pm}K^{\mp}$ and $B_d^0 \to \overline{D}^0 K^{*0}$, D^0K^{*0}, D_1K^{*0}. Excellent tracking and vertexing is also essential in order to achieve the good momentum, energy, mass and proper time resolutions for the study of the rapidly oscillating B_s mesons and their CP asymmetries. Finally, a powerful read out system and online processing are necessary to optimize the physics data acquisition.

1.2 Detector layout

LHCb is a single-arm spectrometer with a forward angular coverage from approximately 10 mrad to 300 (250) mrad in the bending (non-bending) plane. The choice of the detector geometry is justified by the fact that at high energies both the b- and $\overline{\text{b}}$-hadrons are predominantly produced in the same forward cone. The layout of the LHCb spectrometer is shown in figure 1.

Intersection Point 8 of the LHC has been allocated to the experiment. A modification to the LHC optics, displacing the interaction point by 11.25 m from the centre, has permitted maximum use to be made of the existing cavern for the LHCb detector components.

The present paper describes the LHCb experiment, its interface to the machine, the spectrometer magnet, the tracking and the particle identification, as well as the trigger and online systems, including its front end electronics, the data acquisition and the experiment control system. Finally, the expected global performances of LHCb, as deduced from detailed Monte Carlo simulations, are summarized.

The interface with the LHC machine is described in section 2. The description of the detector components is made in the following sequence: (a) the spectrometer magnet which is a warm dipole magnet providing an integrated field of 4 Tm, is described in section 3, (b) the vertex detector system (including a pile-up veto counter), called VErtex LOcator is described in section 4.1, (c) the tracking system made of a Trigger Tracker, Si-μ strip detector, in front of the spectrometer magnet, and three tracking stations behind the magnet, made out of Si-μ strips in the inner parts and of

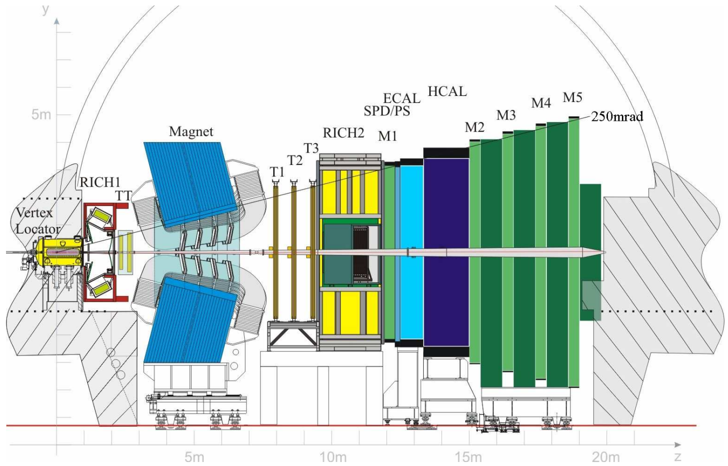

Figure 1. View of the LHCb detector.

Kapton/Al straws for the outer parts is described in sections 4.2 and 4.3, (d) two RICH counters using Aerogel, C_4F_10 and CF_4 radiators, to achieve excellent π–K separation in the momentum range from 2 to 100 GeV/c, and novel Hybrid Photon Detectors are described in Sec. 5.1, (e) the calorimeter system composed of a Scintillator Pad Detector and Preshower (SPD/PS), an electromagnetic (shashlik type) calorimeter and a hadronic calorimeter is described in section 5.2, (f) the μ detection system providing μ identification and contributing to the level-0 trigger of the experiment, composed of MWPC (except in the highest rate, where triple-GEM's are used) is described in section 5.3.
All detector subsystems are assembled in two halves, which can be moved out separately horizontally for assembly and maintenance, as well as to provide access to the beam pipe.
The trigger, the online system, the expected performances of the detector and computing resources are described in sections 6, 7, 8, respectively.

2. Interface to the machine

2.1 Beam pipe

2.1.1 Introduction

The beam pipe design is particularly delicate since the vacuum chamber is located in the high rapidity region of the LHCb detector where the particle density is high, and the number of secondary particles depends on the amount of material seen by incident primary particles. The mass of the beam pipe and the presence of flanges or bellows have direct influence on the occupancy, in particular for the tracking chambers and the RICH detectors. Optimisation of the design and selection of materials were therefore performed in order to maximize the transparency in these critical regions [3], [4].

2.1.2 Layout

The beampipe, schematically represented in Fig. 2 includes the vertex locator (VELO) forward window covering the full LHCb acceptance and four main conical sections, the three closer to the IP being made of beryllium and the one further away of stainless steel.

Beryllium was chosen as the material for 12 m out of the 19 m long LHCb beampipe, for its high transparency to the particles resulting from the collisions. It is the best available material for this application given its high radiation length combined with a modulus of elasticity higher than that of stainless steel. However, its toxicity [5], fragility and cost are drawbacks taken into account in the design, installation and operation phases. Flanges, bellows and the VELO exit window were therefore made of high strength aluminium alloys which provide a suitable compromise between performance and feasibility. The remaining length, situated outside the critical zone in terms of transparency, is made of stainless steel, a material widely used in vacuum chambers because of its good mechanical and vacuum properties. The VELO window, an aluminium 6061-T6 spherically shaped thin shell, is 800 mm in diameter and was machined from a specially forged block down to the 2 mm final thickness. The machining of the block included a four convolution bellows at its smallest radius. The first beampipe section (UX85/1), that traverses RICH1 and TT, is made of 1 mm thick Be, includes a 25 mrad half-angle cone and the transition to the 10 mrad half-angle cone of the three following beampipe sections. In order to avoid having a flange between the VELO window and UX85/1, the two pieces were electron beam welded before installation. Sections UX85/2 (inside the dipole magnet) and UX85/3 (that traverses the Tracker, RICH2, M1 and part of ECAL) are 10 mrad beryllium cones of wall thickness varying from 1 to 2.4 mm as the diameter increases from 65 up to 262 mm. UX85/3 is connected to a stainless steel bellows through a Conflat seal on the larger diameter. The transition between aluminium and stainless steel is ensured by an explosion bonded connection. The three Be beampipes were machined from billets up to 450 mm long and assembled by arc welding with non-consumable electrode under inert gas protection (TIG) to achieve the required length. TIG welding was also used to connect the aluminium flanges at the extremities of the tubes.

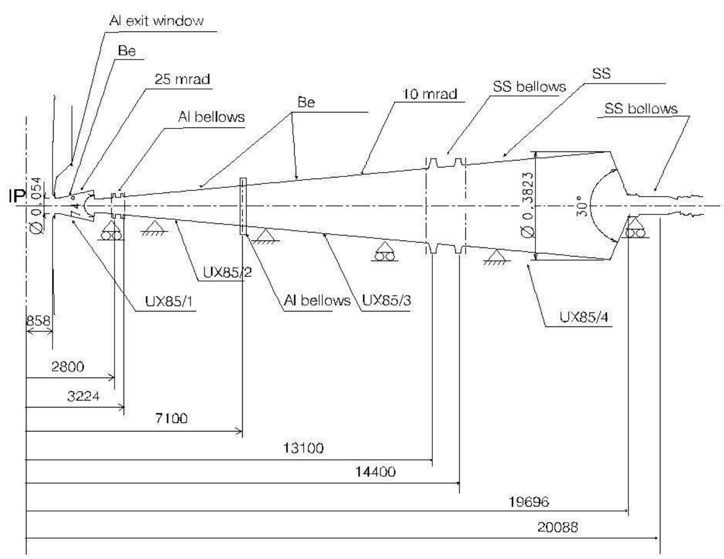

Figure 2. LHCb beampipe schematic layout.

The UX85/4 section completes the 10 mrad cone and includes a 15° half-angle conical extremity that provides a smooth transition down to the 60 mm final aperture. It was manufactured from rolled and welded stainless steel sheet of 4 mm thickness. A copper coating of 100 microns was deposited before assembly on the non-IP end cone to minimise the impedance seen by the beam. The aluminium and stainless steel bellows compensate for thermal expansion during bakeout and provide the necessary flexibility to allow beampipe alignment. Optimised Ultra High Vacuum (UHV) flanges were developed in order to minimise the background contribution from the various connections in the high transparency region [6] . The resulting flange design is based on all-metal Helicoflex seals and high strength AA 2219 aluminium alloy flanges to ensure reliable leak tightness and baking temperatures up to 250°C. A relatively low sealing force allows the use of aluminium and a significant reduction of the overall mass compared to a standard Conflat flange.

Another important source of background is the beam pipe support system [7] . The fixed supports, which must compensate the unbalanced vacuum forces due to the conical shape of the beampipe, are each obtained through a combination of 8 stainless steel cables or rods mounted under tension, pulling in both upstream and downstream directions with an angle to the beam axis (Fig. 3).

Rods are used where a high stiffness is required, as cable stiffness is inferior to that of a rod for the same amount of material. Where a movable support is required to allow thermal expansion, four stainless steel cables are mounted in the plane perpendicular to the beampipe blocking all movements except along the beam axis. All structures from which the cables hold the beampipe are made of aluminium alloy and have minimised mass.

Two sector valves at the extremities allow interventions and commissioning independent of the LHC machine vacuum.

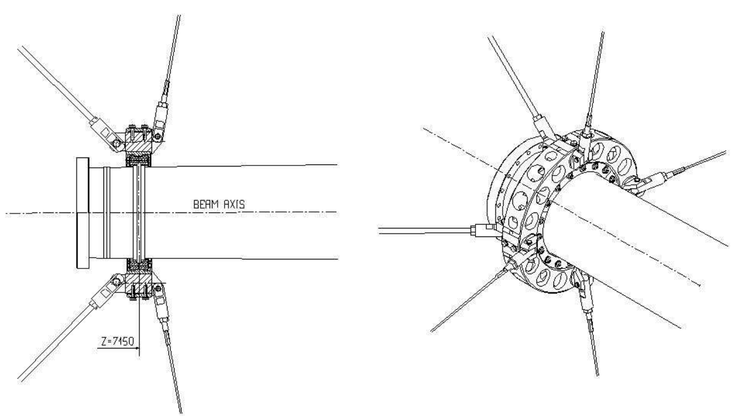

Figure 3. Beampipe fixed support. A system of 8 high resistance rods and cables hold the beampipe through an optimised aluminium collar

2.2 Vacuum

In order to achieve an average total dynamic pressure of 10^{-8} to 10^{-9} mbar with beam passing through, the LHCb beampipe and the VELO RF-boxes are coated with sputtered non-evaporable getter (NEG) [8] . This works as a distributed pump, providing simultaneously low outgassing and desorption from particle interactions with the walls. Another purpose of the NEG coating is to prevent electron multipacting [9] inside the chamber, since the secondary electron emission yield is much lower than for the chamber material. The UHV pumping system is completed by sputter-ion pumps in the VELO vessel and at the opposite end of the beampipe in order to pump non-getterable gases. Once the NEG coating has been saturated, the chamber must be heated periodically (baked out) to 200°C, for 24 hours, in order to recover the NEG pumping capacity. The temperature will have to be gradually increased with the number of activation cycles, however it is limited to 250°C in the optimised flange assemblies for mechanical reasons. Before NEG activation, the vacuum commissioning procedure also includes the bakeout of the non-coated surfaces inside the VELO vacuum vessel to a temperature of 150°C. Removable heating jackets are installed during shutdowns covering the VELO window and the beampipe up to the front of RICH2. From RICH2 to the end of the Muon System a permanent system is installed beacuse of the the difficulty in access. However, the permanent system up to 13 m from the IP requires the use of high transparency materials. The heating power is generated by thin stainless steel strips, mounted between two polyimide, foils and the thermal insulation layer is made of silica aerogel. The beampipe inside the muon filters also has permanent bakeout equipment but as there are no transparency constraints, the insulation is made from a mixture of silica, metal oxides and glass fiber, whilst the heating is provided by standard resistive tapes.

Such an optimised vacuum chamber must not be submitted to external overpressure or shocks while under vacuum due to the risk of implosion. Hence, it must be vented to atmospheric pressure before certain interventions in the surrounding detectors. Saturation of the NEG coating and consequent reactivation after the venting will be avoided by injecting an inert gas not pumped by the NEG. Neon was found to be the most suitable gas for this purpose because of its low mass and the fact that it is not used as a tracer for leak detection, such as helium or argon. However, commercially available Ne must first be purified before injection. A gas injection system installed in the cavern will provide the clean neon to be injected simultaneously into both VELO beam vacuum and detector vacuum volumes, as the pressure difference between the two volumes must be kept lower than 5 mbar to prevent damage to the VELO RF-boxes (c.f. VELO RF-box section 4.1).

3. Magnet

3.1 General Description

A dipole magnet has to be used in the LHCb experiment at LHC to identify and measure momentum of charged particles in the decay of B-mesons produced in p-p collisions. The measurement has to be possible in the large forward volume of space delimited by the ±250 mrad vertical and ±300 mrad horizontal acceptance. The super-conducting magnet originally proposed in the Technical Proposal [1], would have required high investment costs and an unacceptably long construction time. It was replaced by a magnet design with resistive saddle-shaped coils in a window-frame yoke with sloping poles in order to match the required detector acceptance. Details on the design evolution of the magnet project are in the Magnet Design Report and [1],[10],[11],[14]. The design of the magnet with an integrated magnetic field of 4 Tm for tracks of 10 m length, had to accommodate the contrasting needs for a field level inside the RICHs envelope less than 2 mT and a field as high as possible in the regions between the Vertex Locator, and the Trigger Tracker tracking station [13]. The design was also driven by the boundary conditions at P8 of the LHC accelerator, previously occupied by the DELPHI detector: the cavern is equipped with two cranes, each of 40 t lifting capacity and of restricted lateral displacement. This has implied the magnet had to be assembled in a temporary position outside the beam area, by putting together elements relatively light. The DELPHI rail systems and part of the magnet carriages have been used as platform for the LHCb magnet for obvious reasons of cost. Plates 100 mm thick of laminated low carbon steel, having a maximum weight of 25 t, were used to form the identical horizontal bottom and top parts and the two mirror-symmetrical vertical parts (uprights) of the magnet yoke.

The two identical coils are of conical saddle shape and are placed mirror-symmetrically to each other in the magnet yoke. The pancakes themselves are made of pure Al-99.7 hollow conductor in an annealed state which has a central cooling channel of 25 mm diameter and has a specific ohmic resistance below 28 nm at 20°C. It is produced in single-length of about 320 m by rotary extrusion (Holten Conform TM)[1] each one tested for leaks with water up to 50 bars and for extrusion imperfections before being used. The coils were produced in industry by the firm SigmaPhi[2] with some equipment and technical support of CERN. Cast Aluminum clamps are used to hold togheter the triplets making up the coils, and to support and center the coils with respect to the measured mechanical axis of the iron poles with tolerances of some mm. As the main stress on the conductor is the thermal one, the design choice was to leave the pancakes of the coils free to slide upon their supports, with only one coil extremity kept fixed against the iron yoke (where electrical and hydraulic terminations are located). Finite element models (TOSCA, ANSYS) have been extensively used to investigate the coils support system with reference to the effect of the electro-magnetic and thermal stresses on the conductor, and the measured displacement of the coils during magnet operation has matched the predicted value quite well. After rolling the magnet into its nominal position, a final precise adjusting and alignment were carried out in order to follow the 3.601 mrad slope of the LHC machine and its beam. The resolution of the alignment measurements was about 0.2 mm while the magnet could be aligned to its nominal position with a precision of ± 2 mm.

The stand-alone Magnet Control System (MCS) manages the operation of the magnet, from water flow to power supply. The Magnet Safety System (MSS) enforces discharge of the magnet if set-points value (of temperatures, voltages, pressure drop, water leaks) are overcome or if general safety conditions are not fulfilled. The first current in the 1600 ton spectrometer magnet was injected in November 2004, and the nominal current of 5.85 kA was reached soon after, representing the first of the 4 LHC detector magnets to be made fully operational on the beam line. The magnet was safely operated for some time at 13% above the nominal current. Several field map campaigns have been carried out, the first one in November 2004, then another one was measured during June and July 2005 after the RICH's shielding boxes were put in positions, and the final one, including

[1]Holton Machinery, Bournemouth, UK.

[2]SigmaPhi, Vannes, France.

exensive measures in both polarities was completed during November and December 2005. The field in the RICHs boxes, with and without the mu-metal tubes, as well as the fringe field at several locations in the pit, were measured. The main measured magnet parameters are reported in Table I. At today, the magnet has collected many hours of uninterrupted activity and many cycles of operations. All in all the magnet has shown a reliable and stable performance.

3.2 Field Mapping

In order to obtain the necessary high resolution of the charged particle tracks, LHCb needs to know $\int Bdl$ with an momentum uncertainty of a few times 10^{-4} and the position of the B-field peak with a precision of a few mm. This accounts for an absolute precision of the measurements in the order of 10 G. A special measuring machine was designed to enable mapping the LHCb magnetic field with the required precision. To improve data quality and reduce human errors the measurement system has been built with some redundancies. One of these is the possibility to overlap measurements to crosscheck data. A remotely controlled motor system situated outside of the magnet has been used to scan through the dipole longitudinal axis. The motor displaces along the z direction a support holding two adjacent G10 planes each equipped with 30 printed circuit cards distributed over a grid of 80 mm x 80 mm. Every sensor card has mounted on a cube of 4 mm side dimensions three orthogonal and calibrated Hall probes. The support can be placed (manually) orthogonal to the z-axis in the up/down (y-axis) or right/left (x-axis) directions, to allow mapping of different regions. The 3D sensor cards are the result of a joint R&D carried out by CERN and NIKHEF and are calibrated at CERN to a precision of 10^{-4}. To get such a high precision the sensor cards were accurately measured (with NMR) in a constant homogeneous magnetic field B while rotating the cards (which are positioned with a 0.01 mm precision) over two orthogonal axes. The temperature T is also measured to allow taking into account possible effects on the calibration. The Hall-voltage is decomposed in orthogonal functions and the magnetic field parameterized in polynomial coefficients. The calibration process allows corrections for non-linearity, temperatures effects and non-orthogonality. A special calibration machine has been set-up at CERN [4]. After being mounted within the LHCb dipole magnet, the elements of the machine had to be aligned along the LHC beam axis: the rails along which the carriage holding the Hall probes moves and the support itself, were aligned with 0.2 mm relative accuracy, and about 1 mm of absolute precision. Between the aims of the field map campaigns were to measure, inside the tracking volume and for both polarities, the components of the magnetic field and to compare these to simulated data. All that is necessary to produce an accurate and precise field map useful for reconstruction purposes. Also the magnetic field has been measured inside the shielding volumes of RICH1 and RICH2 and near the TT and VELO regions. The machine operations will require a cyclic operation of the magnet; moreover considering that the reversal of the spectrometer dipole field is important for the control of systematic effects in CP asymmetries, it was necessary to fully understand the effect of hysteresis on the repeatability of the magnetic field. At the moment the frequency of polarity reversals is not yet finalised, but expected to be of order of few per month.

3.3 Field Parameters

The level of precision obtained in the final measuring campaigns is shown in figure 4. It is 3.10^{-4}, as it was requested by physics reconstruction needs. In figure 5. is shown the magnetic field

Non-uniformity of $\|B\|$	$\pm 1\%$ in planes xy of 1 m^2 from z=3m to z=8 m
$\int Bdl$ upstream TT region (0-2.5 m)	0.1159 Tm
$\int Bdl$ downstream TT region (2.5 - 7.95 m)	3.615 Tm
Max field at HPD's of RICH1	20 G (14 G with mu-metal)
Max field at HPD's of RICH2	9 G
Electric power dissipation	4.2 MW
Inductance L	1.3 H
Nominal / maximum current in conductor	5.85 kA / 6.6 kA
Total resistance (two coils + bus bars)	R = 130 $m\Omega$ @ 20° C
Total voltage drop (two coils)	730 V
Total number of turns	2 x 225
Total water flow	150 m^3/h
Water Pressure drop	11 bar @ ΔT = 25°C
Overall dimensions H x V x L	11m x 8 m x 5 m
Total weight	1600 tons

Table 1. Measured main parameters of the LHCb magnet

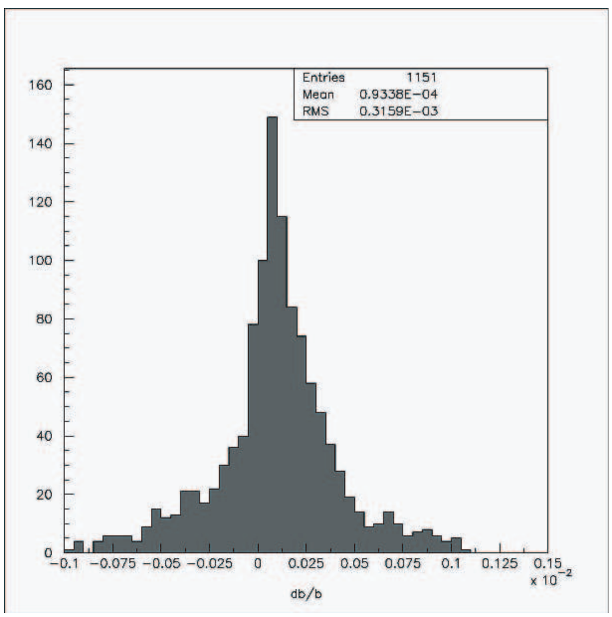

Figure 4. Precision of measurements.

on the axis for both polarities, together with the simulated data. The agreement looks excellent, however a zoom in the RICH1 region, as in figure 6. doesn't show such a perfect match of data with the simulation. Indeed the field integrals, in the region 0 - 250 cm, has been measured to be 115.9 kG∗cm, while the computed one was 120.1 kG∗cm. This 3.5% difference is supposed to be partially an effect of the iron embedded in the concrete (which is very close to the RICH1 shield)

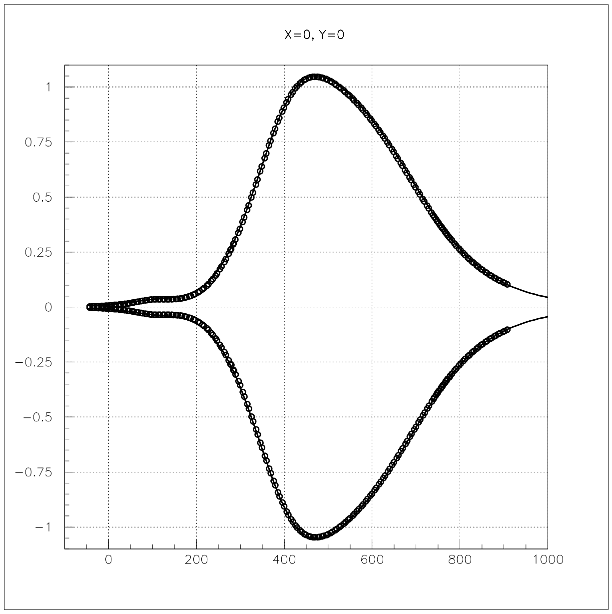

Figure 5. Magnetic field on axis.

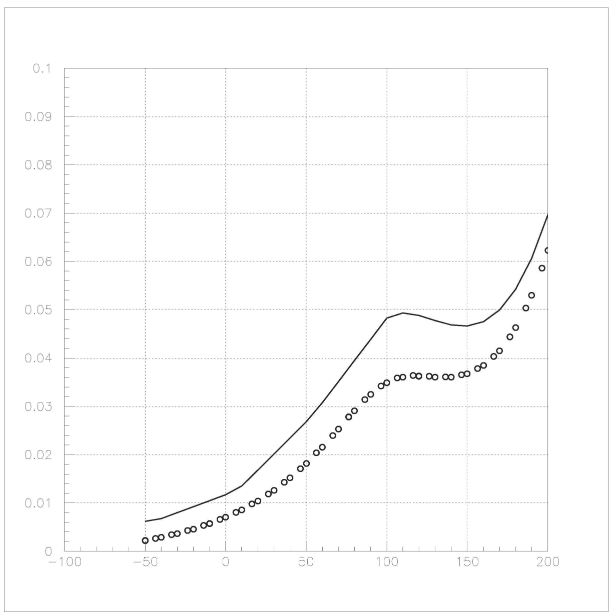

Figure 6. Zoom in RICH1 region - simulation (solid line) and measurements (open circles).

and to the precision of the TOSCA computed model. In all the other regions, and specially inside the tracking volume, the agreement between data and simulation is much better, less than 1%.

In conclusion, Bx, By and Bz have been measured in a fine grid 8cmx8cmx10cm, spanning from the interaction point to the entrance of RICH2 ($\approx 10m$), and covering most of the LHCb acceptance. The measurement of B is reproducible within 0.03%, and the same resolution for

$\Delta B/B$ was measured when changing polarity, once defined the right procedure to demagnetize the magnet.

4. Tracking

4.1 Vertex locator

4.1.1 Introduction

The VErtex LOcator (VELO) provides precise measurements of track coordinates close to the interaction region, which are used to identify the displaced secondary vertices which are a distinctive feature of b-hadron decays. For this, the VELO features a series of silicon modules, each providing a measure of the R and Phi coordinate, arranged along the beam direction (figure 9). They are placed at a radial distance from the beam which is smaller than the aperture required by the LHC during injection and must therefore be retractable. This is achieved by mounting the detectors within a vacuum vessel, separated from the primary vacuum by a thin walled aluminium box, which is corrugated in such a way as to minimise the material before the first measured point, and allow the two halves of the closed detector to overlap. Figure 7 shows a cross section of the VELO vessel, illustrating the separation between the primary vacuum of the beam and the secondary vacuum enclosed by the VELO boxes. Figure 8 shows a zoom with a view from inside one of the boxes, with the sides cut away to show the staggered and overlapping modules of the opposite detector half. The aluminium boxes with their corrugated foils, hereafter referred to as "‘RF-foils"’, provide a number of functions with in the VELO, which are discussed in the following sections.

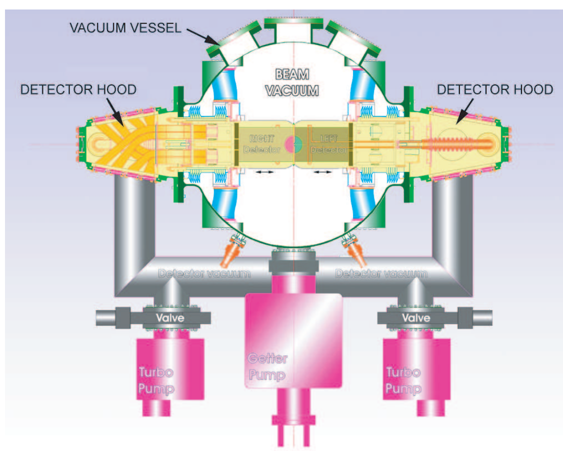

Figure 7. Cross section of the Vertex Locator vacuum vessel, with the detectors in the fully closed position. The routing of the signals via kaptons to vacuum feedthroughs are illustrated. The separation between the primary and secondary vacuua is achieved with thin walled aluminimum boxes enclosing each half.

4.1.2 Requirements and constraints

The ability to reconstruct vertices is fundamental for the LHCb experiment. The track coordinates provided by the **VE**rtex **LO**cator (VELO) are used to reconstruct production and decay vertices of beauty- and charm-hadrons, to provide an accurate measurement of their decay lifetimes and

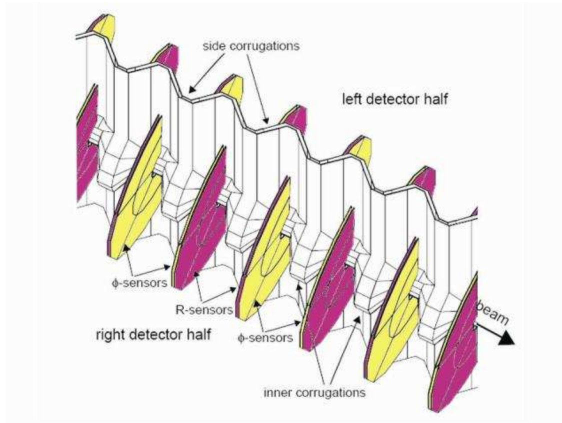

Figure 8. Zoom on the inside of an RF-foil, with the detector halves in the fully closed position. The edges of the box are cut away to show the overlap with the staggered opposing half. The R and Φ sensors are illustrated with alternate shading.

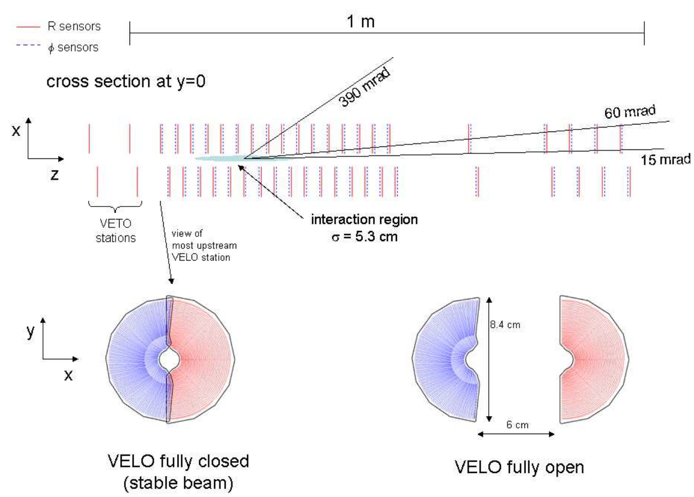

Figure 9. Cross section in z of the VELO silicon sensor positions, at $y = 0$, with the detector in the fully closed position. The front face of the first modules are also illustrated in both the closed and open positions. (Note: this figure has to spread over two columns or it is too small)

to measure the impact parameter of particles used to tag their flavor. These measurements play a vital role in the High Level Trigger (HLT, see section ????), which enriches the b-hadron content of the data written to tape, as well as in the LHCb off-line analysis. This imposed a number of requirements on the VELO design which are briefly described in this section.

Geometry The VELO has to cover the angular acceptance of the downstream detectors, i.e. detect tracks with a pseudorapidity in the range[3] $1.6 < \eta < 4.9$ and emerging from primary vertices

[3]Some coverage of negative pseudorapidity is used to improve the primary vertex reconstruction and, using two special stations, to reduce the number of multiple-interaction events passing the L0 trigger (see section xxx on the Pile-Up System).

in the range $|z| < 10.6\,\mathrm{cm}$. The detector setup was further constrained by the following corollary considerations:

- A polar angle coverage down to 15 mrad for a track emerging at $z = 10.6\,\mathrm{cm}$ downstream from the nominal IP, together with the minimum distance of the sensitive area to the beam axis (8 mm, see below), and the requirement that a track should cross at least three VELO stations, defined the position z_{N-2} of the first of the three most downstream stations: $z_{N-2} \simeq 65\,\mathrm{cm}$.

- A track in the LHCb spectrometer angular acceptance of 300 mrad should cross at least three VELO stations. Given a maximum[4] outer radius of the sensors of about 42 mm, the distance between stations in the central region needed to be smaller than 5 cm. Requiring four stations to be traversed (or allowing for missing hits in one of four stations), imposed a module pitch of at most 3.5 cm. Dense packing of stations near the IP also reduces the average extrapolation distance from the first measured hit to the vertex.

- For covering the full azimuthal acceptance and for alignment issues, the two detector halves were required to overlap. This was achieved by shifting along z the positions of sensors in one half by 1.5 cm relative to sensors in the opposite half.

The use of cylindrical geometry ($R\phi$ coordinates), rather than a simpler rectilinear scheme, was chosen in order to enable fast reconstruction of tracks and vertices in the LHCb trigger. Indeed, simulations showed that 2D (Rz) tracking allows a fast reconstruction in the HLT with sufficient impact parameter resolution to efficiently select events with B-hadrons. Each VELO module was designed to provide the necessary 3D spatial information to reconstruct the tracks and vertices. One of the two sensors, called the phi-measuring sensor, or ϕ sensor, provides information on the azimuthal coordinate around the beam. The other sensor, called the R-measuring sensor, or R-sensor, provides information on the radial distance from the beam axis. The third coordinate is provided by knowledge of the position of each sensor plane within the experiment.

The number of strips for both sensor types needed to satisfy the competing requirements of the LHCb environment and physics and a budgetary limit of about 220,000 channels.

Mechanical accuracy As described in section [xxx], the HLT exploits the fact that (2D) Rz-track reconstruction, and associated Rz-impact parameters, allows good discrimination between b-hadron decay tracks and primary vertex tracks. This imposes alignment requirements on the VELO R-sensors. Ideally, the VELO circular strips should all be perfectly centered around the major axis of the luminous region (sometimes called the beam axis). Simulation studies have shown how the trigger performance would degrade as function of various VELO R-sensors misalignment paremeters [Ruf,Petrie] xxx. The ensuing constraints on the VELO construction and positioning accuracies are summarized in table 2.

Performance The global performance requirements of the detector can be characterized with the following inter-related criteria:

[4] This allowed the use of 10 cm Si wafers for sensor production.

Internal R-sensor RMS spread from nominal	 σ_x and $\sigma_y < 20\,\mu$m $\sigma_z < 100\,\mu$m
External R-sensor average misalignment along z axis	 Δx and $\Delta y < 100\,\mu$m $\Delta\alpha$ and $\Delta\beta < 0.2\,$mrad $\Delta z < 1\,$mm

Table 2. Alignment requirements for the VELO sensors. 'Internal' refers to sensor-to-sensor alignment within a detector half. 'External' refers to alignment of a detector half relative to other objects, such as the luminous region or the opposite detector half. Δx and Δy can be changed online. Other adjustements require an intervention on the VELO mechanics. Here, α and β indicate rotations around x and y.

- Signal to noise ratio (S_tN)[5] In order to ensure efficient trigger performance, the LHCb VELO aimed for an initial signal to noise ratio of greater than 14 [**?**].

- Efficiency: The overall channel efficiency was required to be at least 99% for a signal to noise cut $S_tN > 5$ (or for a signal to noise cut giving about 200 noise hits per event in the whole VELO detector).

- Resolution: A spatial cluster resolution of about 4 μm was aimed at for 100 mrad tracks in the smallest strip pitch region (about 40 μm). Furthermore, it was required that the resolution be not degraded by irradiation nor by any aspect of the sensor design.

A further performance requirement was imposed that affected more the read-out electronics than the sensor design: at 40 MHz read-out speed, the spill-over probability was required to be less than 0.3 in order to keep the number of remnant hits at a level acceptable for the HLT.

Environmental The VELO detector will be operated in an extreme radiation environment with strongly non-uniform fluences. The damage to silicon at the most irradiated area for an accumulated luminosity of $2\,\text{fb}^{-1}$ is equivalent to that of 1 MeV neutrons with a flux of 1.3×10^{14} particles/cm^2 ($= 1.3\times10^{14}$ n_{eq}/cm^2), whereas the irradiation in the outer regions does not exceed a flux of 5×10^{12} n_{eq}/cm^2. The detector was required to sustain 3 years of nominal LHCb operation. In order to evacuate the heat generated in the sensor electronics (in vacuum) and to minimize radiation-induced effects a cooling system was required, capable of maintaining the sensors at a temperature between -10 and 0 C with a heat dissipation of about 24 W per sensor and hybrid. To increase the sensor life time, continuous cooling after irradation was also requested (with the aim to expose the irradiated sensors to room temperature for periods shorter than 1 week per year).

The sensor full depletion voltage is expected to increase with fluence. The ability to increase the operational bias voltage to full depletion during the 3 years life time of the sensors was imposed as a further requirement.

[5]Signal (S) is defined as the most probable value of a cluster due to a minimum-ionizing particle and noise (N) as the RMS value of an individual channel

Machine integration constraint

Mechanical The design of the VELO was constrained by the following boundary conditions:

- A short track extrapolation distance leads to a better impact parameter measurement. Therefore, the innermost radius of the sensors should be as small as possible. In practice, this is limited by the aperture required by the LHC machine. During physics running conditions, the σ of the beams will be less than 100 μm, but for safety reasons, the closest approach allowed to the nominal beam axis is 5 mm (this number is dominated by the yet unknown closed-orbit variations of the LHC and could be reduced in an upgraded detector). To this must be added the thickness of the RF-foil, the clearance between the RF-foil and the sensors, and the need for about 1 mm of guard-ring structures on the silicon. Taking everything into account, the sensitive area can only start at a radius of about 8 mm.

- During injection, the aperture required by the LHC machine increases, necessitating retraction of the two detector halves by 3 cm, which brings the movable parts in the shadow of the LHCb beam pipe (54 mm diameter). Furthermore, the repeatability of the beam positions could not be guaranteed, initially, to be better than a few mm. This imposed that the VELO detectors be mounted on a remote-controllable positioning system, allowing fine adjustment in the x and y directions.

- The need for shielding against RF pickup from the LHC beams, and the need to protect the LHC vacuum from outgassing of the detector modules, required a protection to be placed around the detector modules. This function is carried out by the RF-foil, which represents a major fraction of the VELO material budget in the LHCb acceptance.

Vacuum The required performance of the LHCb VELO demands positioning of the sensitive area of the detectors as close as possible to the beams and with a minimum amount of material in the detector acceptance. This is best accomplished by operating the silicon sensors in vacuum. As a consequence, integration into the LHC machine became a central issue in the design of the VELO.

In this case, the amount of material in front of the silicon detector is mainly determined by the necessity to shield against the RF pickup and the mechanical constraint of building a sufficiently rigid foil. The detectors operate in a secondary vacuum and hance the foils are not required to withstand atmospheric pressure. However, the design of the vacuum system had to ensure that the pressure difference between detector and beam vacuum never be so large as to cause inelastic deformations of the detector box. The VELO surfaces exposed to beam-induced bombardment (secondary electrons, ions, synchrotron radiation) needed to be coated with suitable material in order to maintain beam-induced effects, such as electron multipacting and gas desorption, to levels acceptable for efficient LHC and LHCb operation. The LHC beam vacuum chamber, and therefore also the VELO primary vacuum vessel, were required to be bakeable.

Impedance Beam bunches passing through the VELO structures will generate wake fields which can affect both the VELO system (RF pick-up, losses) and LHC beams (instabilities). These issues have been addressed in detail [**?**, **?**, **?**, **?**] and are further discussed in section **??**. In the design of the

VELO, wake field suppression is achieved by the RF foils, which provide continuous conductive surfaces which guide the mirror charges from one end of the VELO vessel to the other.

General Detector Overview

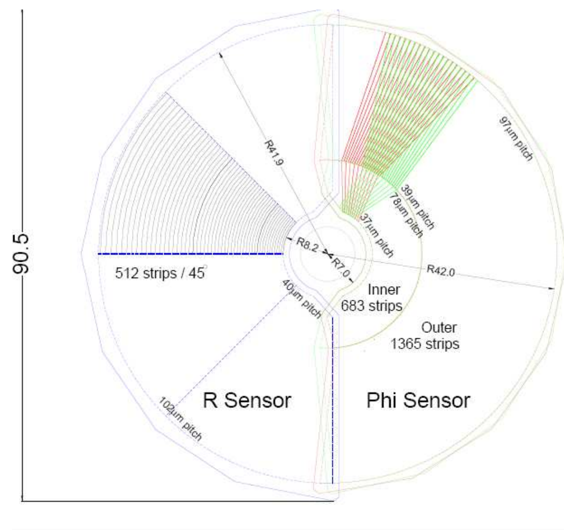

Figure 10. Sketch illustrating the Rϕ geometry of the VELO sensors

Sensors The severe radiation environment at 8mm from the LHC beam required the adoption of radiation tolerant n-implants in n-bulk technology with strip isolation achieved through the use of a p-spray. The minimum pitch achievable by Micron Semiconductor, Ltd, the company chosen to fabricate the LHCb sensors, using this technology was approximately 32 μm(depending on the precise structure of the readout strips). For both the R and ϕ-sensors the minimum pitch is designed to be at the inner radius to optimize vertex resolution.

For the R-sensor the diode implants are concentric semi-circles with their centre at the nominal LHC beam position. In order to minimize occupancy, and hence the chance of spurious ghost hits in the sensors, and the individual strip capacitance each strip is subdivided into four 45 degree regions. The minimum pitch at the innermost radius is 40 μmincreasing linearly to 101.6 μmat 41.9mm. This ensures that hits along the track contribute with roughly equal precision.

The ϕ-sensor is designed to readout the orthogonal coordinate to the R-sensor . In the simplest possible design these strips would run radially from the inner to the outer radius and point at the nominal LHC beam position with the pitch increasing linearly with radius starting with a pitch of 35.5 μm. Given the minimum pitch achievable through fabrication only about 700 strips could be fitted onto such a sensor. In order to improve the segmentation and increase the number of strips to 2048, to match the R-sensor, the phi is subdivided into two regions, inner and outer. The outer region starts at a radius of 17.25 mm and its pitch is set to be half (39.3 μm) that of the inner region (78.3 μm) which ends at the same radius. The design of the strips in the ϕ-sensor are complicated by the introduction of skew to improved pattern recognition. At 8mm from the beam the inner strips have an angle of approximately 20 degrees to the radial whereas the outer strips make an angle of approximately 10 degrees to the radial at 17mm. The skew of inner and outer sections is reversed giving the strips a distinctive "dog-leg" design. The modules are placed so that adjacent phi-sensors have the opposite orientation with respect to the each other. This ensures that adjacent stations are able to distinguish ghost hits from true hits through the use of a traditional stereo view.

The technology utilized in both the r and ϕ sensors is otherwise identical. Both sets of sensors are 300 μmthick. Readout of both r and ϕ-sensors is at the outer radius and requires the use of a second layer of metal (a routing layer or "double metal") isolated from the AC coupled diode strips by approximately 3 μmof chemically vapour deposited (CVD) SiO_2. The second metal layer is connected to the first metal layer by wet etched "vias". The strips are biased using 1MΩ polysilicon resistors and both detectors are protected by an implanted guard ring structure.

	R sensor	Phi-sensor
number of sensors	50 + 4 (VETO)	50
readout channels per sensor	2048	2048
sensor thickness	300 μ m	300 μm
smallest pitch	40 μ m	38 μ m
largest pitch	102 μ m	97 μ m
length of shortest strip	3.8 mm	5.9 mm
length of longest strip	33.8 mm	24.9 mm
inner radius of active area	8.2 mm	8.2 mm
outer radius of active area	42 mm	42 mm
angular coverage	182 deg	$\approx$ 182 deg
stereo angle	-	10-20 deg
double metal layer	yes	yes
average occupancy	1.1 %	1.1/0.7 % inner/outer

Table 3. Principal characteristics of VELO sensors

The pitch as a function of the radius R in μm is given by the following expressions:

$$\mathrm{R}:\ 40+(101.6-40)\times\frac{R-8190}{41949-8190}$$
$$\mathrm{Phi}:\ 37.7+(79.5-37.7)\times\frac{R-8170}{17250-8170}\quad(\mathrm{R}<1.725\ \mathrm{cm})$$
$$\mathrm{Phi}:\ 39.8+(96.9-39.8)\times\frac{R-17250}{42000-17250}\quad(\mathrm{R}>1.725\ \mathrm{cm})$$

The VELO sensors were developed for high radiation tolerance. Early prototype detectors (which here chosen to be n^+n used p-stop isolation. This was later replaced by p-spray isolated detectors which showed much higher resistance to micro-discharges. The n^+n design was compared with an almost geometrically identical p^+n design and again shown to have much better radiation characteristics as measured by charge collection as a function of voltage.

Prototype sensors were also irradiated with non-uniform fluence in order to study the effects of cluster bias due to non-homogenous irradiation. It was shown that the transverse electric fields produce less than 2 μmeffects on the cluster centroid.

A subset of the production sensors were exposed to a high neutron fluence ($1.3\times10^{14}\ \mathrm{n_{eq}/cm^2}$) representing 1 year of operation at nominal luminosity. A strong suppression of surface breakdown effects was demonstrated. The evololution of the depletion voltage was found to correspond to the expectation over LHC operation, as illustrated in figure 12.

Modules The VELO module has three basic functions. Firstly it must hold the sensors in a fixed position relative to the detector base. Secondly it provides and connects the electrical readout to the sensors. Finally it must enable thermal management of the modules which are operating in vacuum.

Each module was designed to hold the sensors in place to better than 50 μmin the plane perpendicular to the beam and within 800 μmalong the direction of the beam. Sensor to sensor alignment (within a module) was designed to be better than 20 μm.

The module was comprised of a substrate, for thermal management and stability, on to which two circuits were laminated. This formed the hybrid. The substrate was approximately 120x170x1mm and fabricated "'in-house" with a core of 400micron thick thermal pyrolytic graphite (TPG), and encapsulated, on each side, with 250 μmon carbon fibre (CF). A CF frame of about 7mm surrounded the TPG and was bonded directly to the CF encapsulation to prevent delamination. The TPG was designed to carry a maximum load of 32W away from the front-end chips. A semicircular hole was cut into the substrate under the region where the detectors were glued. The circuits were fabricated and laminated to the substrates by Stevenage Circuits Ltd. Particular attention in the design and fabrication process was given to trying to produce almost planar hybrids. This was necessary to enable the subsequent module production. Typical non-planarities of order 250μmwere achieved.

The circuits were commercially hand populated, to minimize te exposure of the hybrid to high tempereature and hence the possibility of delamination. The sensor front-end ASIC (Beetle 1.5) were then glued to the circuits. A total of 32 Beetles were used in each module. Kapton pitch adaptors, manufactured by CERN, were glued to the circuits in order to facilitate the wire bonding of the sensors to the Beetle chips. Sensors were glued to the double sided hybrid with an sensor-sensor accuracy of better than 10μm. The sensors were bonded to the electronics using a combination of H&K 710 and K&S 8090 bonding machines with 25 micron Al wire. After bonding and final testing 99.4% of all strips were operational, with no sensor having more than 30/2048 faulty channels.

The final mechanical mounting of the hybrid was to the module base and pedestal. The pedestal was a low mass CF fibre construction designed to hold the hybrid stably. The pedestal is a hollow rigid structure approximately 140x150x10mm. One end was glued to the hybrid and second to a CF base which contains two precision manufacture Invar feet. The design of the base allows repeatable mounting of the module to LHCb with a precision a better than 10 μm. The accuracy of the final assembled modules satisfied the design criteria.

Thermal and electrical connections were made via [this part not yet complete]

4.1.3 Mechanics

Introduction The great challenge for the LHCb vertex detector is an accurate reconstruction of displaced vertices. This requires that the detectors are placed close to the interaction region, with a minimum amount of material between the interaction point and the silicon sensors. The ultra high vacuum requirements of the LHC ring, the necessity for wakefield suppression, the need to shield the detectors from electromagnetic effects induced by the high frequent beam structure, and the necessity to retract the detectors by 30 mm from the interaction region during injection of a new LHC fill, make the design of such a detector demanding. To meet all these constraints, we have opted for a design with two detector halves, each placed inside a thin-walled aluminum box.

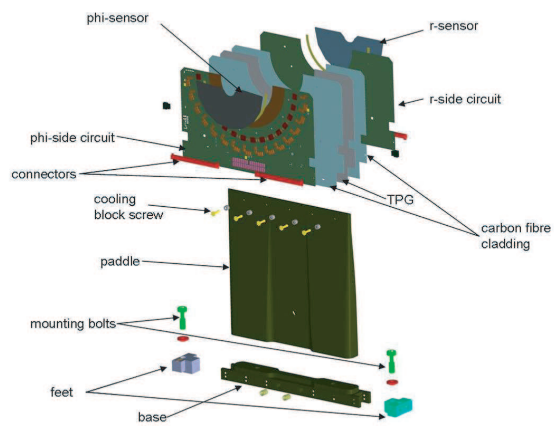

Figure 11. Illustration of the principal components of the VELO module.

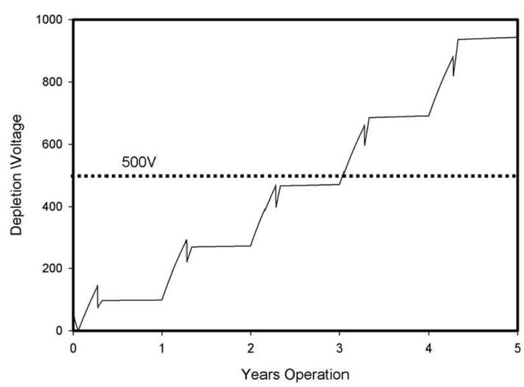

Figure 12. Expected evolution of the full depletion voltage as a function of time. A nominal luminosity of $2 \times 10^{32} cm^{-2} s{-}1$ has been assumed, with an integrated luminosity of 2fb^{-1} per year. After approximately 3 years of nominal operation the maximum deliverable full depletion voltage is reached.

Aluminum was chosen since it has a relatively low Z (resulting in a small radiation length), good electrical conductivity, and can be machined quite easily. The side walls of these boxes are 0.5 mm thick. In order to allow for overlap in the two detector halves, the top surfaces of these vacuum boxes have a corrugated shape and are made from 0.3 mm thick AlMg3 foil (an aluminum alloy with 3 % magnesium). The two detector boxes are placed in a 1.4 m long vacuum vessel with a diameter of 1.1 m. The whole assembly is shown in figure 13. Two rectangular bellows allow for the movement of the detector boxes inside the vacuum system. Each detector support is connected via three spheres on holders placed within circular bellows (one of them can be seen in figure 13) to the movement mechanism that is located outside the vacuum vessel.

To suppress wake field effects, the dimension of the beam pipe as seen by the proton beams has to vary very gradually. To match the beam pipe upstream and downstream from the vertex locator, wakefield suppressors made of 50 μm thick copper-beryllium have been implemented such that both in the open and the closed position there is a good match.

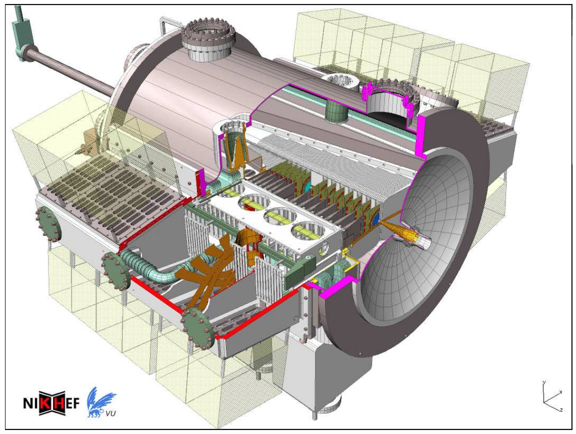

Figure 13. Overview of the vacuum vessel with the vertex locator.

The exit foil of the vessel has been designed and produced at CERN; it consists of a 2 mm thick aluminum window.

Movement system Before the LHC ring is filled, the detectors have to move away from the interaction region by 30 mm in order to allow for beam excursions during injection and ramping. After stable beam conditions have been obtained, the detectors should be placed in to an optimized position centered in x and y around the interaction region. This position is not exactly known beforehand; it may vary over $\pm$ 5 mm in both x and y, even from fill to fill. Therefore, a procedure has been developed to determine the beam position with the detectors not completely moved in, and then move to the optimal position. This is performed with a motion mechanism that can bring the detectors to their position with an accuracy in the order of 10 μm by means of a stepping motor with resolver read-out. Additional potentiometers have been used to verify independently the proper functioning. The motion procedures are controlled by a PLC (Programmable Logic Controller).

Vacuum system In the LHC ring, Ultra High Vacuum conditions are required (better than 10^{-8} mbar in the LHCb area). To maintain these conditions, the beam pipes are equipped with a NEG layer. To maintain the pumping capacity of the NEG coating, the vacuum system will be vented during maintenance with ultra-pure neon.

The thin walled detector boxes will be plastically deformed at a pressure difference of 20 mbar (and above 50 mbar it will even break). Therefore, the detectors also have to be operated under vacuum. Due to outgassing of detectors, hybrids, cables and connectors a vacuum around 10^{-4} mbar is expected. Hence a good separation between the beam and detector vacuum is necessary as exposure to the detector vacuum saturates and poisons the NEG material at the inside of the beam pipe.

An elaborate procedure has been implemented to make sure that during venting and evacuation of the VELO vacuum system the pressure difference between beam and detector vacuum will never exceed 5 mbar. This is obtained by using dedicated valves and restrictions, that are activated by membrane switches that react at a 5 mbar pressure difference. The complete vacuum system between the two sector valves around LHCb is controlled by a PLC.

Cooling system Since the detectors and read-out electronics are operated inside the vacuum system, active cooling is required. Furthermore, in order to limit the effects of radiation damage of the silicon sensors, the irradiated sensors should be operated and kept at temperatures below -5°C at all times. The radiation hard refrigerant in the system is two-phase CO_2 cooled by a conventional freon cooler. The liquid CO_2 is transported via a 60 m long transfer line to the VELO, where it is distributed over 27 capillaries per detector half. Each capillary is thermally connected to five cooling blocks that are attached to each detector module.
A redundant pumping system is incorporated, so that also during maintenance periods cooling capacity is available. The complete cooling system is controlled by a PLC.

4.1.4 Electronics chain

Beetle [a short paragraph will be added here describing the beetle and its parameters relevant for VELO operation]

Kaptons During injection of the proton beams in the ring, the detectors have to be retracted by 30 mm. The total number of data and control signals that run between the hybrids and the feed-through flanges at the vessel exceeds 18000. Kapton cables were opted for as they are thin, flexible and radiation hard. The central part of the cable consists of a 17 μm thick layer with 150 μm wide copper strips. On top and bottom side this layer is covered with a 100 μm thick kapton foil, a 17 μm thick rolled annealed (AR) copper foil which is used to supply power to the Beetle chips, and a 25 μm thick cover kapton foil.
Each cable consists of two parts: a short tail from the hybrid to a fixed connector, and a long cable from this connector to the vacuum feed through on the vessel.

Repeater Boards The repeater board (RPT) is located directly outside of the VELO tank inside repeater crates. The RPT function is mainly a repeater for data differential signals, TFC and FE chips configurations signals. Also, monitoring signals are sent out via the board to the detector slow control system. The RPT carries the voltage regulators required by the FE electronic and the L0 electronic service system. For flexibility in design and mainly for maintenance, the RPT is built as a motherboard hosting several mezzanine cards:

- Four Driver Cards: Four driver cards are mounted in the RPT board as mezzanine cards. Each card contains 16 fully differential analog drivers. Because the data streams are sent to the digitizer card trough a 60m individual shielded twisted pair cable, the drivers include a line equalizer to compensate distortions introduced by the cable.

- One LV card: The low voltage card provides the power for the FE hybrid, the analog driver cards and the ECS card. Eight radhard voltage regulators are mounted on the board. Each voltage is monitored through an amplifier. The card is supplied by three power supplies located at 60m. Sense-lines are used to compensate the voltage drop through the long supply cable.

- One ECS card: The ECS (Experiment Control System) repeats the signal for the I2C configuration bus, control signals and monitors the regulator voltages. Temperature and radiation monitoring signals are multiplexed on this card and sent to the control board located on the balcony

Another part implemented on the board is TFC functionality. These fast lvds signals are signals required by the FE chips and are sent through a LVDS repeater mounted on the board.

Analogue data transmission [a short paragraph will be added here describing the analogue data transmission]

VELO specific TELL1 features The TELL1 boards of LHCb are described elsewhere [**?**]. Specific to the VELO is the digitization of the data on the TELL1 and the complex pre-processing of the data.

Due to the high radiation levels and space constraints the digitization of the data and the use of optical drivers close to the detector were discarded for the VELO. As a consequence, the analogue data are directly transmitted to the TELL1 board and digitized in the analogue receiver cards, the A-Rx [**?**]. Each TELL1 board deals with the data from one sensor, i.e. 64 analogue links, and features 4 A-Rx cards. One A-Rx card provides 16 channels of 10-bit ADC's to sample the analogue data from 4 Beetle chips at 40MHz. In order to compensate for the time skew of the signals resulting from different cable lengths the sampling time can be chosen by phase adjustable programmable clocks, individually for each ADC channel.

After the digitization the TELL1 performs the data processing on the ppFPGAs before sending the zero-suppressed data to the trigger farm, see [21]. Although most of the signal distortion in the long data cable is removed already through the frequency compensation, see section **??**, the first step of the data processing implements an FIR filter acting on the 10bit data coming from the ADC. It takes care of remaining cross talk originating in the sensor as well as in the readout chain behind. The next processing step implements a pedestal follower, the subsequent pedestal subtraction and it offers three choices for limiting the precision to 8bit: by truncating either the two most significant bits, or the two least significant bits or one of each. In both, the R- and the ϕ-measuring sensors, adjacent physical strips are scrambled in the readout chain. For the following processing steps it is essential to bring them back into order. For this reason the channel reordering step was implemented, which uses the 8 bit data as input. The linear common mode suppression corrects for correlated noise pickup. It is implemented as an iterative procedure, where signal channels are masked out and the common mode is modeled as a constant and a slope. The last processing step is the clustering [22]. Strips are selected as seeding strips if they pass a certain seeding threshold. Strips next to the seeding strips are included, if their signal lies above the inclusion threshold cut. A cluster can be formed by a maximum of four strips. A cluster center is calculated with a three bit precision.
The cluster data are formatted and sent for use by the software trigger algorithms. The VELO Raw Data are sent in a format, that allows a fast access of the cluster information by the trigger

by sending the calculated 14bit cluster position. For more refined calculations of the cluster center and offline analyses, the ADC information of all strips is added in another data block, see also [23].

LV system The VELO and PILE-UP low voltage system is based on the multi-channel power supply system from CAEN. All the power supplies are installed in the control room. With 12 fully floating channels, each module can supply 4 repeater cards and each channel has its own sense-line to recover long distance cable voltage drop. [this paragraph will be expanded slightly]

HV System The VELO silicon sensors will be operated under a reverse bias ranging from 100 - 500 V, with the operating voltage being increased as the sensors undergo radiation damage. The high voltage system utilises 6 power supply modules manufactured by Iseg, each controlling up to 16 sensors, which are housed in an uninterruptible power supply crate in the detector counting house. The output is fed via 37 core cables to a patch panel in the counting house. Long 56 core cables connect the counting house to a second patch panel located near the detector. 3 core cables then provide the high voltage, high voltage guard and ground to the repeater board of each module. The high voltage guard connection provides the voltage to a guard trace on the sensor hybrids, which surrounds the detector high voltage trace, thus reducing possibilities of shorting. The high voltage guard line can be connected to the high voltage, to ground or left floating by adjustment of jumper switches in the counting house patch panel. The high voltage system is controlled through PVSS. Voltage and current limits are also set in hardware on the power supply units. The high voltage power supplies are controlled by the VELO hardware interlock system (see section xx).

Grounding and power supply The partitioning of the VELO and the PILE-UP electronics follows the detector topology. Each silicon detector with its hybrid forms a group. There are in total 84 VELO and 4 PILE-UP hybrids with 16 Fe chips each installed inside the VELO tank. The power distribution and grounding scheme will follow this partitioning. The number of group connected to the same power is kept to a minimum. The VELO tank is the main center part of the system and is connected to the LHC machine. The LHC machine earth cables are apart of the cavern network grounding. It must be certain that all metallic devices, such electrical cabinets, cable trays, have ground connection to the main VELO tank support. All electrical devices being a part of the cooling and vacuum system must be grounded for safety reason following the standards. All components connected to the main power network line are equipped with protection and floating devices made of conductive material are not allowed in the system. The analogue FE electronics is the main victim of noise from external noise sources. The effects are minimised by keeping the signal current path as small as possible. In the VELO configuration of the detector and the FE chip, the signal generated by charges traverses a reversed bias diode and is transmitted to the charge preamplifier. The reference of the FE chip is internally connected to the ground pins of the chip. The signal current loop is closed through the bias line. The bias line is AC connected to the reference ground of the FE chip. The silicon detectors and hybrids are sitting close to the RF shield connected to the ground. Potential differences between these two components and are the main noise source. To minimize these differences, the hybrid ground plane is at the same potential as the RF shield. Each hybrid ground plane is tied to the base plate with a grounding strap clamp.

4.1.5 Test Beam Detector Commissioning

In 2006, beam test studies of two different configurations of fully instrumented VELO modules were carried out at the CERN H8 experimental area, using 180 GeV π beams with tuneable intensity and spot size. An external trigger was provided by a scintillator telescope that could be configured to select events including single tracks crossing the detectors without interactions, multiple track event produced by interaction in the detectors or on the entrance window of their enclosure, and, in the second data taking cycle, interactions in targets modules installed in the module array. A total of 4 target planes, mounted in two modules were used. Each plane included a small Cu disk, with a thickness of 300 μm and a 2 mm diameter, centered on the beam axis, as well as 5 mm diameter Cu disks with a radial displacement of 15 mm from the beam axis. The latter targets had the purpose of investigating the vertex reconstruction capabilities of the detector in the retracted position ("open VELO").

The first configuration included 3 fully instrumented half-disk pre-production modules built using 200 μm sensors, enclosed in a box providing accurate positioning on a table including a rotating stage that allowed data taking with the detectors oriented at an angle with respect to the beam direction. Thus the detectors were operated in air, at a typical temperature of 40-50 $^\circ$C. Only 1/4 of the electronics was read out, due to limitations in the available hardware. This data set provided considerable insight on the performance of the real detector modules. Figure 14 shows the measured cluster multiplicity as a function of the strip number, compared with our expectations based on a dedicated Monte Carlo simulation [17]. The agreement between predictions and data is very good. Thus we expect that an optimal charge weighting algorithm will deliver the expected resolution (better than 6 μm for strips at 45 μm pitch, and between 6 and 12 μm for strips. All the parameters affecting the performance of the sensors and the readout electronics were explored, for example we took data with the sensors biased with different high voltages, and we changed the operating parameters of the Beetle chip. In addition, we studied the noise performance both with random triggers, and with electronic calibration runs. These data validate in a multi-detector configuration the laboratory characterization of the individual module components. In addition, the data acquisition and monitoring infrastructure planned for the experiment were used. The system performance was excellent.

The second data set was taken using a system of 10 production modules including 300 μm thick sensors, mounted in the vacuum tank built for the final system, reading out 6 of them in different combinations, depending upon the trigger scheme chosen. Most of the data was taken in vacuum (below $10-3$ mbar) and at a temperature of about -3°C. This test beam cycle provided valuable operating experience with the cooling system built for the experiment and with vacuum implementation and monitoring. Data were takin in air at room temperature, and it was confirmed that the module positions remained stable through the transition between air and vacuum and throughout temperature cycling.

The production detectors operated with a signal to noise ratio of between 17 and 25, depending on strip length, where signal is defined as the most probable value of the charge cluster produced by a minimum ionizing particle and noise is the incoherent noise measured from calibration data upon subtracting the coherent noise component. Their performance proved remarkably stable throughout the 14 days of data taking. Single track and interaction trigger data complement the first data set in

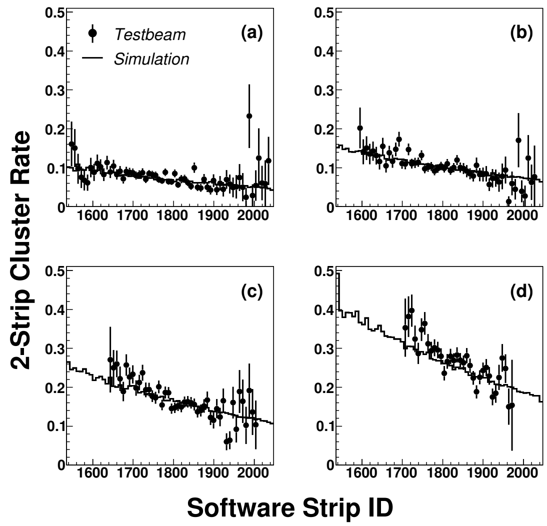

Figure 14. Fraction of two strip clusters in R sensitive detectors, for different beam track angles with respect to the normal to the detector planes: a) 0°, b)5°, c) 10° d) 15°. The horizontal axis represents the strip number, corresponding to a progressive coarser inter-strip pitch.

assessing the hit resolution as a function of angle. In addition, target data produced reconstructed primary vertices from all the targets that will help us in tuning the vertex reconstruction algorithms.

4.1.6 Material Budget, radiation damage

[this paragraph is to come]

4.1.7 Velo Software

Reconstruction Software The TELL1 data processing boards (see section xx) perform a set of processing algorithms on the raw VELO data and identify VELO clusters. The clusters are utilised in the Gaudi framework for use in the trigger and for off-line physics analysis. For use in the trigger the clusters are stored in a compressed form that is optimised for speed and provides 3-bit precision on the inter-strip position of the clusters using a simple weighted pulse height algorithm. For offline use a higher precision calculation, and an estimate of the uncertainty on this position, is provided. This calculation uses the inter-strip pitch and track angle as well as the pulse-height of the strips in the cluster.

In addition to the standard output data format of the TELL1, a number of other output formats are provided for calibration and monitoring purposes and are decoded in the software framework. Notably, these include a raw data format, where the ADC values (at 8 bit precision) are provided for all strips.

A bit-perfect emulation of the full TELL1 processing algorithms is available in the software framework. Using the emulation the raw data can be processed to produce the cluster format. The performance of the TELL1 algorithms can thus be assessed at each stage of the processing. The emulation is also used to tune the optimal settings of the adjustable parameters in the TELL1 emulation, such as the the signal thresholds used in the clustering algorithm. The emulation can also be performed on simulated data.

A simulation framework for the VELO has been provided. This framework describes the material and layout of the VELO detector. The response of the silicon detectors and front-end chip pulse shape to the passage of particles is described, based on physically motivated parametrisations tuned to describe laboratory and test-beam data. The resulting simulated analogue signals are passed through the emulated TELL1 clustering algorithm and the output clusters are stored in the same format as for the real data.

[something on pattern recognition here ? or is there a tracking section of paper ?]

Monitoring The monitoring of the VELO can be divided into two strands: the short term operational checks performed on-line; and the longer-term off-line performance monitoring.

The on-line monitoring is performed using the LHCb monitoring farm. Cluster and track monitoring, including monitoring of residuals for the alignment, is performed using the standard output data. Raw data is also produced for a subset of events at a rate of a fraction of a Hz and, through use of a special calibration trigger, read out to the monitoring farm. This rate allows the performance of individual channels to be accurately assessed on a timescale of one hour. The full information of the TELL1 processing boards is available through monitoring using the TELL1 credit-card PC. Preliminary tests of the VELO on-line monitoring have already been performed in the VELO test-beam.

A critical element of the VELO monitoring is the determination of the beam-position. The alignment framework, reconstruction, pattern recognition, track fitting and vertex-finding algorithms are used to build up a 3 dimensional picture of the beam position. This monitoring process is used to determine the correct step-wise movements that are required to close the VELO halves and centre them around the beam at the start of an LHC fill. The beam stability is then monitored during LHC operation.

The off-line monitoring uses a range of analyses to assess the performance of the VELO and to tune operational parameters including the high voltage applied to each sensor, the TELL1 hit processing parameters, and the cluster resolution model used in the tracking. The analyses include: time alignment studies for beam synchronisation; charge collection efficiency and signal to noise studies; resolution studies as a function of detector pitch and projected angle; cross-coupling, pedestal, noise and common-mode noise studies making particular use of the TELL1 emulation. These studies use the full range of VELO TELL1 output formats for the data and also make use of special calibration trigger and test-pulses generated in the Beetle front-end chip of the VELO modules.

Alignment The alignment of the VELO is reliant on three components: the precision construction and assembly of the hardware; the metrology of the individual modules and assembled system; and the software alignment of the system using tracks.

The construction and assembly of the system is reported elsewhere in this section. The construction precision has tight mechanical tolerances, for example the VELO silicon sensors are nominally located only 1mm from the the aluminium RF foil. However, the driving factor for the required construction tolerance is the successful operation of the LHCb trigger. The VELO pattern recognition algorithm (see section x), used to identify high impact parameter tracks in the trigger, is performed for speed reasons initially in the R-z projection. This requires that the strips on the R sensors accurately describe circles around the beam position.

A survey of the individual VELO modules and of the assembled VELO halves on their base plates was performed. The relative position of the VELO modules and of the R and Phi sensors on a single module were measured to a precison of better than 10 μm. The silicon sensors were found to have no significant curvature: 8 points were measured on the surface of the silicon sensors and the mean rms deviation of the sensors from a plane was found to be 14 μm. The precision survey is an important element of the VELO alignment: not only does it provide the starting position for the VELO software alignment but it also constrains degrees of freedom of the system which it will not be possible to accurately measure with data. For example, the overall z positions of the sensors - an important parameter for particle lifetime measurements - is obtained from the metrology. The sensor alignment parameters obtained from the survey were propagated to the LHCb conditions database.

Whilst the survey was performed at room temperature and pressure, no significant deviations are expected for the final system. The module base plate will be maintained at a constant temperature of 20° C. The deviations of the system under vacuum have been determined from the software alignment in a test-beam using a partially assembled VELO and seen to be typically 10 μm or less, as shown in figure 15.

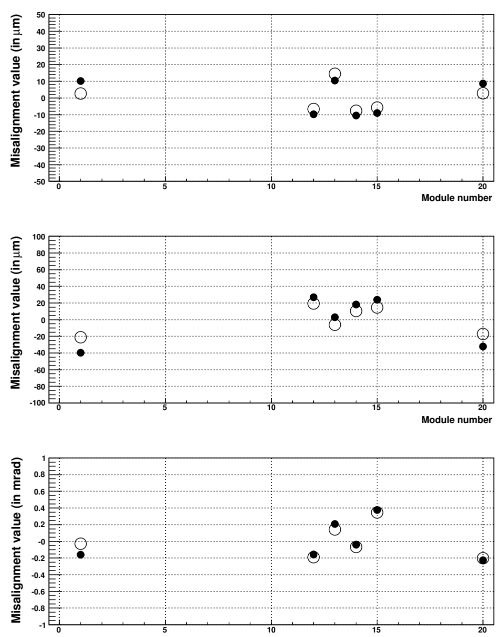

Figure 15. x-axis translation alignment constants of the Velo modules determined by software alignment from data recorded during the VELO test-beam in air and in vacuum. The modules are seen to be stable in the different conditions, including different thermal gradients across the modules.

As previously stated, prior to the LHC establishing stable beams, the VELO is in a retracted position and is brought into its nominal position only after stable beam is established. The position of the VELO halves will be known through the motion control and position measurement system to an accuracy of 10 μm. Combining this information with the relative alignment of the VELO modules obtained from the previous fill is expected to provide sufficient alignment accuracy for operation of the VELO trigger. However, the option to perform a VELO software alignment at the start of each fill remains should this prove necessary.

The software alignment procedure for the detector naturally divides into three distinct parts:

- An internal alignment of the modules within each VELO-half box using the residuals of hits on reconstructed tracks.
- A relative alignment of the two VELO half-boxes with respect to each other principally relying on using tracks passing through the geometrical overlap between the modules in the two VELO half-boxes.
- A global alignment of all sub-detectors relative to each other. This part is reviewed in [20].

Both stages of the VELO alignment are dependent on the same approach, a non-iterative method using matrix inversion. The alignment is based upon a χ^2 function produced from the residuals between the tracks and the measured clusters in the VELO. The track and alignment parameters can be obtained through minimisation of this χ^2 function.

The equations which describe the trajectories of particles are expressed as a linear combination of both the local (track-dependent) parameters and the global (alignment) parameters. All tracks are correlated since the global alignment parameters are common to each track, hence it is necessary to fit all tracks simultaneously.

The χ^2 function can be minimised by solving the set of simultaneous equations given by the derivatives of the χ^2 with respect to the local track parameters and global alignment parameters. This results in a system of equations of a final size, n_{total} given by:

$$n_{total} = n_{local} \times n_{tracks} + n_{global} \tag{4.1}$$

where n_{local} is the number of local parameters per track(four parameters for a straight line in 3D) , n_{tracks} is the number of tracks used for the alignment and n_{global} is the number of alignment constants. Whilst the direct inversion of such large matrices is not computationally practical, the alignment can be handled by inverting the matrix by partition, thus reducing the problem to a $n_{global} \times n_{global}$ matrix inversion. Inversion by partitioning is handled by the Millepede program [19].

The number of tracks required for an effective alignment of the VELO is relatively modest but the alignment is improved by using a mixture of tracks from primary vertex interactions and a complementary tracks set from a source such as beam-halo particles. The CPU requirements of the alignment are also low: of order minutes on a single PC.

Internal Alignment Each VELO half-box contains 21 VELO modules, and each module has three translational and three rotational degrees of freedom. In addition the system is insensitive to some global deformations, eg. an overall translation of the VELO half-box, these limitations are expressed through the introduction of 7 constraint equations.

Tests using data from the VELO beam-test and on Monte-Carlo simulation have demonstrated that x and y-axis translations of the modules can be constrained at the few μm level and rotations around the z axis of order 0.1 mrad, with weaker sensitivity obtained to the other degrees of freedom.

The R and Phi sensors bonded in an individual module have had their relative positions measured during the system metrology and can be cross-checked or further improved by fitting any remaining distortions observed in the residuals across the sensor.

Relative Alignment of VELO half-boxes The VELO modules are contained in the two half-boxes, and the six degrees of fredom that locate one half-box with respect to the other is determined by this stage of the alignment. Once the VELO is fully inserted there is a small overlap between the two VELO halves. The relative alignment of the two half-boxes is primarily constrained by tracks that pass through both VELO halves. However, when the VELO is retracted an alternative technique is required, this relies upon fitting primary vertices using tracks fitted in both halves of the VELO.

Monte-Carlo simulation tests have demonstrated that x and y translations of the half-boxes can be constrained at better than 20 μm and rotations around these axes to better than 0.1 mrad. Rotations around the z-axis are constrained at the 0.2 mrad level.

4.1.8 VELO Performance

The VELO layout has been optimised to minimise the amount of material in the VELO acceptance (see section xx) while providing good geometrical coverage. All tracks inside the LHCb acceptance ($1.6 < \eta < 4.9$) pass through at least three modules, as shown in figure 16.

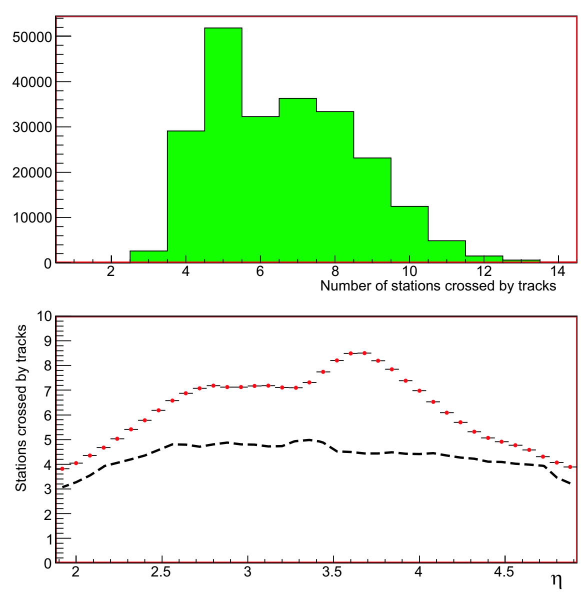

Figure 16. The upper plot shows the number of stations hit per track in the VELO and the lower plot shows the number of hits of a track in the VELO modules as a function of the pseudorapidity of the track η. The dotted line indicates the limit above which 95% of the tracks lie.

The individual hit resolution of the VELO sensors has been determined in a test-beam and is a strong function of the sensor pitch and projected angle, as shown in figure 17. A best resolution of 7 μm was obtained for the raw resolution. The performance is expected to improve when eta corrections and cross talk are fully accounted for.

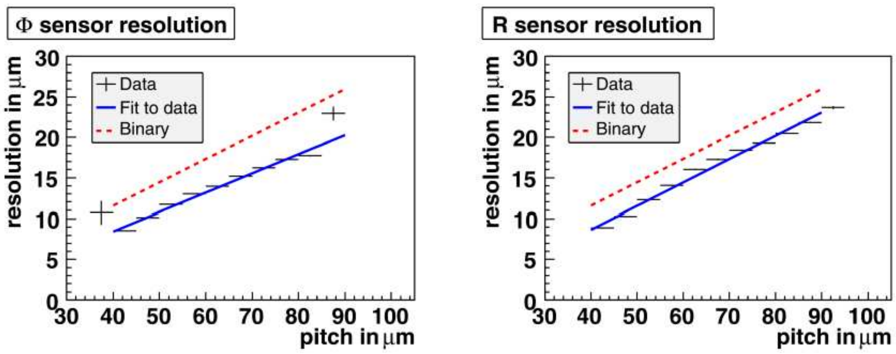

Figure 17. The raw hit resolution as a function of strip pitch as measured in the testbeam. This plot contains the resolution as measured from the weighted centre of the charges on the strips and contains no correction for η or cross talk.

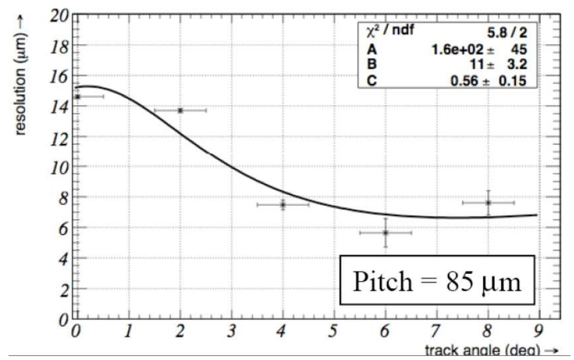

Figure 18. The individual hit resolution of a VELO R sensor as a function of the projected angle for a pitch of 85μm. A correction to account for the η shape has been taken into account, significantly improving the precision.

The geometrical coverage and individual hit resolution combine to give the track overall impact parameter performance shown in figure 19. The optimal 3D impact parameter resolution for tracks of high transverse momentum is yy μm.

4.2 Silicon Tracker

The Silicon Tracker (ST) comprises two detectors that use silicon microstrip sensors with a strip pitch of about 200 μm: the Trigger Tracker (TT) [2, 24], which is a 150 cm wide and 130 cm high planar tracking station located upstream of the LHCb dipole magnet and covering the full acceptance of the experiment, and the Inner Tracker (IT) [25], which covers a 120 cm wide and 40 cm high cross-shaped region in the centre of the three tracking stations downstream of the

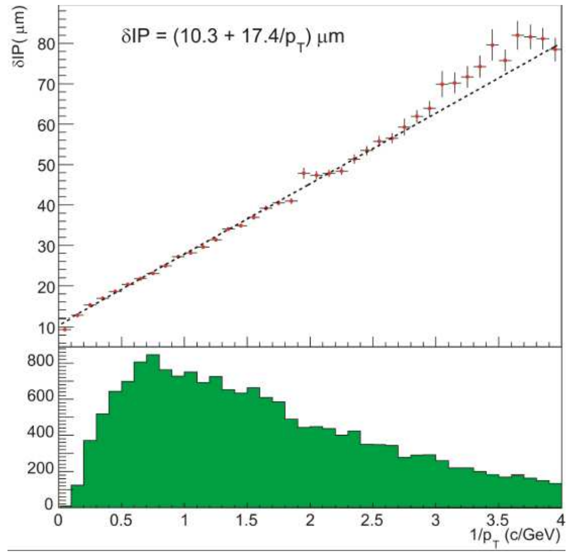

Figure 19. 3D impact parameter resolution as a function of the logarithm of the transverse momentum of the tracks.

magnet. Each of the four ST stations has four detection layers in an (x-u-v-x) arrangement with vertical strips in the first and the last layer and strips rotated by a stereo angle of -5° and +5° in the second and the third layer, respectively. The TT covers a sensitive area of about 8.4 m^2 with 143,360 readout strips of up to 38 cm in length. The IT covers a sensitive area of 4.0 m^2 with 129,024 readout strips of either 11 cm or 22 cm in length.

The main design choices for the Silicon Tracker detectors were largely driven by the following considerations:

Spatial resolution. Monte-Carlo simulation studies have demonstrated that a single-hit resolution of about 50 μm is adequate for both TT and IT, as the momentum resolution of the spectrometer is then dominated by multiple scattering over almost the full range of particle momenta. Readout strip pitches of about 200 μm meet this requirement and were therefore chosen for both detectors.

Hit occupancy. For minimum bias events, charged-particle densities of about 5×10^{-2} per cm^2 for TT and of 1.5×10^{-2} per cm^2 for IT are expected in the hottest regions of the detectors, close to the LHC beam pipe. Towards the outermost regions of the detectors, charged particle densities fall off by almost a factor of ten, to about 5×10^{-4} per cm^2 for TT and 1.2×10^{-2} per cm^2 for IT. Different readout strip lengths were chosen for different regions of the detector to keep maximum strip occupancies at the level of a few percent while minimizing the number of readout channels.

Signal shaping time. In order to avoid “pile-up” of events from consecutive LHC bunch-crossings, fast front-end amplifiers with a shaping time of the order of the bunch-crossing interval of 25 ns have to be used. As a benchmark parameter, the so-called signal remainder was defined as the remaining fraction of the signal amplitude 25 ns after the maximum, i.e. at the signal sampling time corresponding to the subsequent bunch crossing. Simulation studies have shown that signal remainders of 50% for TT and 30% for IT are acceptable for the track reconstruction algorithms.

Single-hit efficiency. Each detection layer should provide full single-hit efficiency for minimum ionising particles while maintaining an acceptably low noise hit rate. The critical parameter is the signal-to-noise ratio, defined as the most probable signal amplitude for a minimum ionising particle, divided by the rms of the single-strip noise distribution. Test-beam studies have shown that the hit efficiency starts to decrease rapidly as the signal-to-noise ratio drops below 10:1. The detector was designed such that a signal-to-noise ratio in excess of 12:1 can be expected, taking into account the expected deterioration from the radiation damage corresponding to ten years of operation at nominal luminosity.

Radiation damage. For ten years of operation at nominal luminosity, 1-MeV neutron equivalent fluences of about 5×10^{14} per cm^2 and 9×10^{12} per cm^2 are expected for the innermost regions of TT and IT, respectively. Basic design rules for radiation-hard sensors were followed to ensure that the detectors will survive these fluences, which are moderate by the standards of modern silicon detectors. The sensors need to be operated at a temperature below 5°C in order to suppress radiation-damage induced leakage currents to a level where shot noise from leakage currents does not significantly deteriorate the signal-to-noise performance of the detector, and where the risk of thermal runaway due to the power dissipated by leakage currents is avoided.

Material budget. As the momentum resolution of the LHCb spectrometer is dominated by multiple scattering, it is important to keep the material budget of the detectors as small as possible. The TT was designed such that all front-end readout electronics and mechanical supports are located outside of the LHCb acceptance. In the case of IT, which is located right in front of the OT detectors, a significant design effort was made to keep the amount of material for mechanical supports and cooling as small as possible.

Number of readout channels. Readout electronics being a major contribution to the overall cost of the detector, the largest readout pitches compatible with the required spatial resolution and the longest readout strips compatible with requirements on occupancy and signal-to-noise performance were chosen in order to minimise the number of readout channels.

Different constraints on the detector geometries resulted in different designs for the detector modules and station mechanics of TT and IT. These are described in Sec. 4.2.1 and 4.2.2, respectively. Common to both parts of the ST are the readout electronics, the power distribution and the detector control and monitor systems. These are the topic of Sec. 4.2.3. Finally, the expected detector performance, based on test-beam measurements and simulation studies, is discussed in Sec. 4.2.4.

4.2.1 Trigger Tracker

All four detection layers of the TT are housed in one large light-tight, thermally and electrically insulated detector volume in which a temperature below 5°C is maintained [26]. To aid track reconstruction algorithms, the four detection layers are arranged in two pairs, (x,u) and (v,x), that are separated by approximately 27 cm along the LHC beam axis.

The layout of one of the detection layers is illustrated in figure 20. Its basic building block is a half-module that covers half the height of the LHCb acceptance and consists of a row of seven silicon sensors, which are electronically split into either two or three readout sectors. A stack of correspondingly two or three readout hybrids is attached at one end of the row of sensors. The regions above and below the LHC beam pipe are covered by one such half-module each. The regions to the sides of the beam pipe are covered by rows of seven (for the first two detection layers) or eight (for the last two detection layers) 14-sensor long full-modules, which cover the

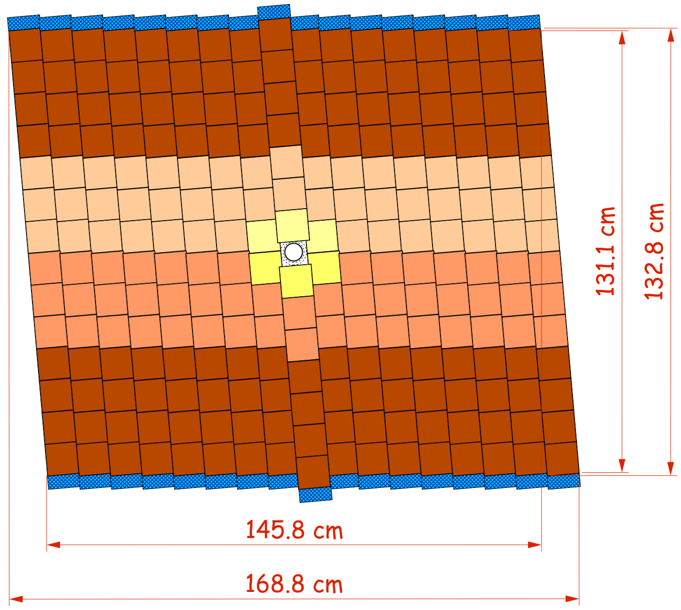

Figure 20. Layout of the third TT detection layer. Different readout sectors are indicated by different shadings.

full height of the LHCb acceptance and are made by joining two half-modules together end-to-end. Adjacent modules within a detection layer are staggered by about 1 cm in z and overlap by a few millimeters in x to avoid acceptance gaps and to facilitate the relative alignment of the modules. In the second and third detection layers, each individual module is rotated by the respective stereo angle of -5° or +5°.

The main advantage of this detector design is that all front-end hybrids are located at the top and bottom ends of the detector, outside of the acceptance of the experiment. This permitted to keep outside of the acceptance all passive material that is invariably associated with the hybrids, cooling, and cables.

TT Detector Modules

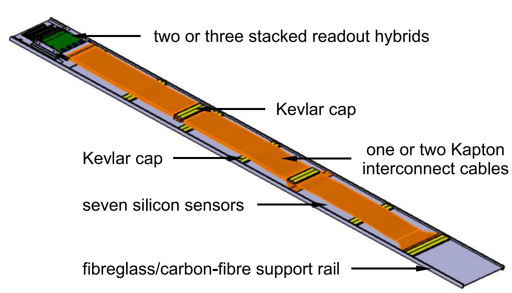

Figure 21. View of a 4-2-1 type TT detector module.

The layout of a half-module is illustrated in figure 21. It consists of a row of seven silicon sensors with a stack of two or three readout hybrids at one end. For half-modules close to the beam pipe, where the expected particle density is highest, the seven sensors are organised into three readout sectors ("4-2-1 type" half-modules). For the other half-modules, the sensors are organised into two readout sectors ("4-3 type" half-modules). In both cases, the first readout sector ("L sector") is formed by the four sensors closest to the readout hybrids and furthest away from the beam. The strips of the four sensors are bonded together and directly connected to the lower-most readout hybrid. For 4-3 type half-modules, the strips of the remaining three sensors are bonded together and form the second readout sector ("M sector"). They are connected via a 39 cm long Kapton flex-cable ("interconnect cable") to a second readout hybrid that is mounted on top of the L hybrid. For 4-2-1 type half-modules, the three remaining sensors are subdivided into an intermediate two-sensor sector ("M sector") and a third sector that consists of the single sensor closest to the beam ("K sector"). The two readout sectors are connected via 39 cm respectively 58 cm long Kapton interconnect cables to two separate front-end hybrids that are mounted on top of the L hybrid. Bias voltage is provided to the sensor backplanes via a thin Kapton flex cable that runs along the back of the half-module. Mechanical stability of the half-module is achieved by glueing two thin fibreglass/carbon-fibre rails along the edges of the L hybrid and the seven silicon sensors.

The *silicon sensors* for TT are 500 μm thick, single-sided p^+-on-n sensors. They are 9.64 cm wide and 9.44 cm long and carry 512 readout strips with a strip pitch of 183 μm. They are identical in design to the OB2 sensors used in the Outer Barrel of the CMS Silicon Tracker and were produced by Hamamatsu Photonics K.K., Japan.

The *Kapton interconnect cables* for the M and K readout sectors were produced using standard plasma-etching technology by Dyconex AG, Switzerland. They carry 512 signal strips and two pairs of bias-voltage and ground strips on a 100 μm thick Kapton substrate. The strips consist of 7 μm thick copper with a 1 μm thick gold plating, are 15 μm wide and have a pitch of 112 μm. The small strip width was required to keep the strip capacitance of the cable small. A short pitch-adaptor section in which the strip pitch widens to 180 μm permits to directly wire-bond the strips on the cable to the silicon sensor strips. A copper-mesh backplane provides a solid ground connection and shielding against pick-up noise. Since the combination of long strips and small strip width led to an unacceptably low production yield for fault-free cables of the required length, it was decided to assemble each cable from either two or three shorter pieces, as illustrated in figure 22. The 39 cm long cables for the M sectors were assembled from a 20 cm long piece that incorporates the pitch-adaptor section and one 19 cm long piece with straight strips. The 58 cm long cables for the K sectors have in addition a second 19 cm long straight piece. Two cable pieces are joined together by glueing them end-to-end onto a common, 1 cm wide and 100 μm thick fibreglass strip. To ensure the electrical connection between the copper-mesh backplanes of the two cable pieces, an electrically conductive adhesive tape is used for glueing them onto the fibreglass strip. The signal, bias voltage and ground strips on the strip-side of the cables are joined together by wire bonds.

Small *Kevlar caps* protect the wire bond rows on the strip-side of the Kapton interconnect cable as well as those in between silicon sensors. These caps are glued onto the surface of the cable or the sensors using an electrically insulating glue (Araldite).

The *front-end readout hybrids* [27] consist of a carrier plate, a pitch adaptor, and a four-layer Kapton flex circuit that carries four Beetle front-end chips [28], some passive SMD components

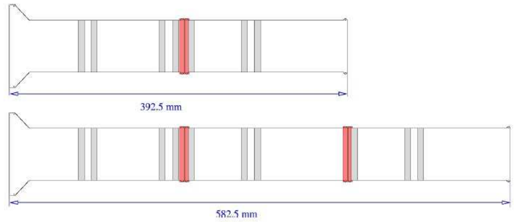

Figure 22. Layout of the two types of Kapton interconnect cables assembled from two or three pieces.

and an 80-pin board-to-board connector through which the analog, multiplexed detector signals are read out and through which control signals, low voltage and bias voltage are provided to the half-module. Of the four conductive layers of the flex circuit, two are used for digital and analog power and ground and the other two carry the signal and control lines. Due to mechanical constraints, three variants of the readout hybrid are used for the three different types of readout sectors. The Kapton flex circuits for all three variants are identical except for a different overall length, the L, M and K hybrids being 67 mm, 57 mm and 46 mm long, respectively. These different lengths are necessary to provide access to the readout connectors on the lower hybrids when the three hybrids are mounted on top of each other. The carrier plates give mechanical stability to the hybrid and act as a heat sink for the heat produced by the Beetle chips. The pitch adaptor matches the input pitch of the Beetle chips to the pitch of the silicon sensors in case of the L hybrid and to the pitch of the Kapton interconnect cable in case of the K and M hybrids. For the K and M hybrids, the carrier plate is made from gold-plated copper and the pitch adaptor is a rectangular piece of alumina (Al_2O_3) with strip lines produced using thin-film technology. The pitch adaptor is glued onto the carrier plate together with the Kapton flex circuit. The carrier plate for the L hybrid is a precisely machined AlN substrate that fulfills multiple purposes. It carries the strip lines of the pitch adaptor, produced using standard thick-film technology. It carries traces and vias to connect the sensor bias voltage from the Kapton flex prints to the backplane of the half-module. Furthermore, it encorporates laser-cut holes that permit to fix and precisely position the half-module inside the detector station, and it defines the thermal interface of the half-module to the cooling plate onto which it is mounted (see below). Finally, the half-module support rails are glued along the sides of this carrier plate. AlN was chosen as a material for this piece for its high thermal conductivity of about 200 W/m·K and its small thermal expansion coefficient of about 4×10^{-6}/K, which is reasonably well matched to that of silicon. The Kapton flex circuits were produced by Optiprint AG, Switzerland. The production of the pitch adaptors and the AlN substrate as well as the assembly and bonding of the hybrids was done by RHe Microsystems, Germany.

The K and M hybrids are mounted onto the AlN substrate of the L hybrid using spacers made of 2.5 mm thick blocks of copper. These copper spacers provide the necessary thermal contact between the K and M hybrids and the AlN substrate of the L hybrid.

The two half-module *support rails* are 5 mm high and 2 mm thick and consist of a 1 mm thick strip of carbon-fibre glued against a 1 mm thick fibreglass strip. A small groove that is milled into the flat side of the fibreglass strip permits to slide the rail over the edge of the seven silicon sensors and the AlN carrier plate of the L hybrid. It is fixed using an electrically insulating two-component glue (Araldite) to ensure the necessary electrical insulation between the strip side of the silicon sensors (which is on ground potential) and their backplane (which carries the bias voltage of up to 500 V). The mechanical rigidity of the rail is mainly given by the carbon-fibre strip. Whereas the fibreglass strips span the full length of the half-module, the carbon-fibre strips extend only from the AlN substrate to the fourth silicon sensor from the hybrids. This permits to join two half-modules to a 14-sensor long full-module by placing them against each other end-to-end and glueing an additional carbon-fibre strip to the free sections of fibreglass strip on both half-modules.

Bias voltage is supplied separately for each of the readout sectors on a half-module through the cable that plugs into the corresponding front-end hybrid. From the hybrid it is connected to aluminium traces on the AlN substrate, using wire bonds in case of the L hybrids and thin copper wires in case of the M and K hybrids. From here, the bias voltages are brought to the back of the half-module through aluminium vias that are embedded in the AlN substrate. Finally, a thin Kapton flex cable, which is glued along the back of the half-module and carries one copper trace per readout sector, provides the bias voltage to the backplanes of the silicon sensors. The electrical connections between the bias voltage pads on the AlN substrate and this cable, and between the cable and the sensor backplanes, are achieved by wire bonds.

TT Detector Station

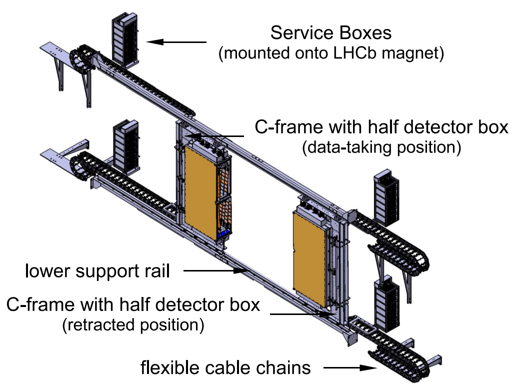

Figure 23. View of the TT station mechanics.

An isometric drawing of the station mechanics is shown in figure 23. Its main structural elements are two large C-shaped aluminium frames, which are mounted onto precision rails and can be retracted for detector maintenance and bake-outs of the beam pipe. Thermal and electrical insulation of the detector volume is provided by walls made of a rigid but light-weight aluminium-clad foam. Mechanical support of the detector modules, as well as cooling of the front-end hybrids and of the detector volume, is provided by so-called cooling plates, which are mounted horizontally near the top and the bottom of the detector volume. They incorporate cooling pipes through which C_6F_{14} at -15°C is circulated as a cooling agent. Additional cooling elements are mounted vertically, close to the side walls of the detector volume. The detector volume is continuously flushed with nitrogen to avoid condensation on the cold surfaces. All electrical signals (detector signals, control signals and supply voltages) are transmitted on Kapton-flex cables through specially designed feed-throughs in the top and bottom walls of the detector box.

The two C-shaped *support frames* are assembled from 15 mm thick aluminium (HABA-35) plates. They rest on a lower precision rail and are guided by an upper precision rail. The two rails are aligned parallel to each other with a precision of better than 100 μm in order to avoid possible distortions of the C-frame during insertion or retraction. The outer walls of the detector box are defined by stiff sandwich plates, consisting of a 30 mm thick aramid honeycomb structure with 1 mm thick aluminium cladding, that are mounted against the inside of the C-frames. For thermal insulation, 40 mm thick plates of polyetherimide foam (Airex R82.60) are mounted against the inside of these honeycomb sandwich plates. On the inside, these Airex plates are laminated with 25 μm of aluminium for electrical insulation and 275 μm of Kevlar for mechanical protection. The front and rear panels of the detector box are made of the same material composition as the Airex plates, except that they are laminated with aluminium and Kevlar on both sides. The front and rear panels are screwed onto the sandwich plates and can be easily removed for the installation and maintanance of detector modules. Around the beam pipe, insulation of the detector volume is achieved by specially machined semi-cylindrical pieces of the same material composition as the front and rear panels. The wall thickness of this beam-pipe insulation piece is 30 mm, except for cut-outs at the positions of the detection layers where the wall thickness is reduced to about 5 mm to minimise the distance between the beam and the innermost detector modules. A clearance of 5 mm between the detector box and the beam pipe is maintained to satisfy LHC safety demands. In data taking position, the insulating elements of the two half stations close off one large light-tight, thermally and electrically insulated volume that houses all detector modules.

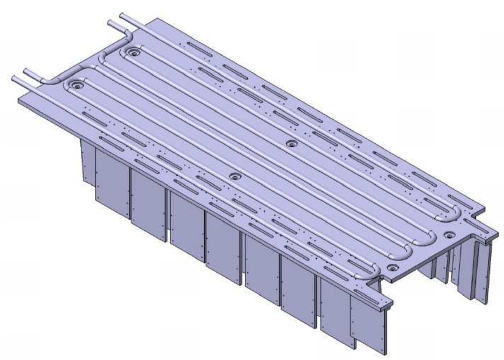

Figure 24. View of one TT cooling plate with mounted cooling balconies.

There are a total of four *cooling plates*, one for each detector quadrant. Each cooling plate is mounted with six pillars made of polyacetal (POM) onto one of the sandwich plates at the top and bottom of the C-frames. A drawing of a cooling plate is shown in figure 24. It is a precisely machined plate of 8 mm thick aluminium that measures 897 mm in x and 348 mm along z. Machined into its outer surface are semi-circular grooves into which two coiled aluminium cooling pipes with an outer diameter of 10 mm and a total length of about 3.5 m are glued. Its inner surface, onto which the detector modules will be mounted, is machined to an overall flatness of better than 100 μm.

The mechanical, thermal and electrical interface between the cooling plates and the half-modules is provided by the *cooling balconies*. They are made from 5 mm thick aluminium and are mounted vertically onto the flat inner surface of the cooling plate where precision pins ensure their accurate positioning. There is one cooling balcony for each half-module. The half-module is screwed onto the flat, vertical surface of the balcony, ensuring a large contact surface and therefore good thermal contact between the balcony and the AlN base plate of the half-module. The correct positioning of the module is ensured by precision positioning pins that are embedded in the balcony. There are two types of balconies: One type for mounting modules vertically (for the x detection layers) and one type for mounting modules under an angle of 5° (for the u and v detection layers). Detector modules in the first two detection layers are mounted onto their balconies from the upstream side of the detector, those for the last two layers from the downstream side of the detector.

The vertical *cooling elements* that are installed at both sides of the detector volume consist of 1 mm thick copper plates onto which long, coiled cooling ducts with a rectangular cross section are soldered.

Detector signals are read out from, and control signals, low voltage and bias voltage supplied to, the detector modules via 50 cm-long **Kapton flex cables** that pass through specially designed feedthroughs in the sandwich cover plates of the C-frames and through dedicated slits in the cooling plate. There is one such Kapton flex cable per readout sector. At one end it plugs directly into the board-to-board connector on the readout hybrid, at the other end it connects to an interface PCB that is mounted onto the outside the detector box. From this patch panel, copper wire cables lead through a flexible cable chain to so-called Service Boxes (see Sec. 4.2.3), that are mounted against the front face of the LHCb dipole magnet and in which the signals are prepared for digital optical transmission to the counting house. There are one flexible cable chain and a stack of six Service Boxes for each quadrant of the detector.

The common electrical ground for all detector modules is defined by the cooling plates. Thin copper wires connect ground pads on each of the Kapton flex prints to the metallic screws that are used to fix the modules onto the cooling balconies.

4.2.2 Inner Tracker

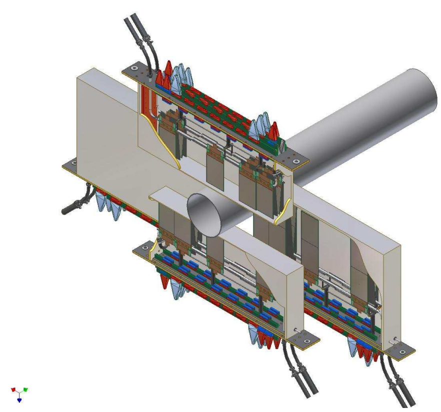

Figure 25. View of the four IT detector boxes arranged around the LHC beam pipe.

Each of the three IT stations consists of four light-tight, thermally and electrically insulated detector boxes that are arranged around the beam pipe as shown in figure 25. They are mounted onto two support frames — one on each side of the beam pipe — that can be retracted for detector maintenance and bake-outs of the beam pipe. Inside the detector boxes, a temperature below 5°C is maintained. Each detector box contains four detection layers and each detection layer consists of seven detector modules. Adjacent modules in a detection layer are staggered by 4 mm in z and overlap by 3 mm in x to avoid acceptance gaps and facilitate the relative alignment of the modules. One-sensor long modules are used in the detector boxes above and below the beam pipe ("top" and "bottom" boxes) and two-sensor long modules are used in the detector boxes to both sides of the beam pipe ("side" boxes). The resulting layout of one IT detection layer is illustrated in figure 26.

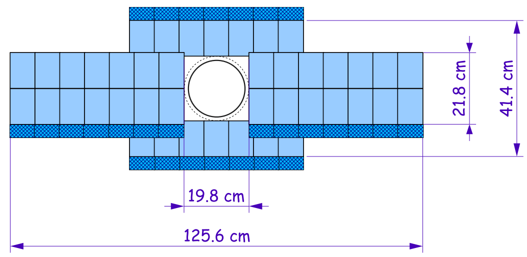

Figure 26. Layout of an x detection layer in the second IT station.

IT Detector Modules

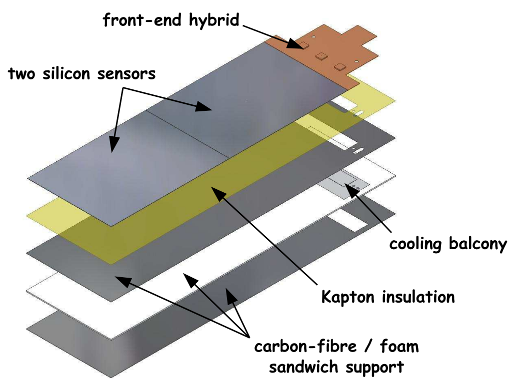

Figure 27. Exploded view of a two-sensor IT module. One-sensor modules are similar except that the support plate is shorter and carries only one sensor.

An exploded view of a detector module is shown in figure 27. A module consists of either one or two silicon sensors that are connected via a pitch adapter to a front-end readout hybrid. The sensor(s) and the readout hybrid are glued onto a flat module support plate. Bias voltage is provided to the sensor backplane from the strip-side via n^+ wells implanted in the n-type silicon bulk. A small aluminium insert ("mini-balcony") that is embedded into the support sandwich at the location of the readout hybrid provides the mechanical and thermal interface of the module to the detector box.

Two types of *silicon sensors* of different thickness, but otherwise identical in design, are used in the IT. They are single-sided p^+-on-n sensors, 7.6 cm wide and 11 cm long, and carry 384 readout strips with a strip pitch of 198 μm. The sensors for one-sensor modules are 320 μm thich, those for two-sensor modules are 410 μm thick. As explained in Sec. 4.2.4 below, these thicknesses were chosen to ensure sufficiently high signal-to-noise ratios for each module type while minimising the material budget of the detector. The sensors were designed and produced by Hamamatsu Photonics K.K., Japan.

The IT *front-end readout hybrids* consist of a four-layer Kapton flex circuit that is very similar in design and routing to that of the TT hybrids, and a pitch adaptor that is similar to those used for the M and K hybrids of TT. The only differences are that the Kapton flex circuit for IT carries only three Beetle front-end chips and that it encorporates an 89 mm long readout tail with straight traces, at the end of which a 60-pin board-to-board connector is mounted. The pitch adaptor is glued onto the Kapton flex circuit and the Kapton flex circuit is glued directly onto the module support plate. The production of flex prints and pitch adaptors as well as the assembly of the readout hybrids took place at the same two companies as for TT.

The *module support plate* consists of a 1 mm thick sheet of polyetherimide foam (Airex) sandwiched in between two 200 μm thick layers of a carbon-fibre composite. The latter are produced from two layers of carbon fibres of high thermal conductivity (Mitsubishi K13D2U) that are oriented at $\pm 10^\circ$ with respect to the module axis. A 25 μm thick Kapton foil is laminated on top of the upper carbon-fibre layer to electrically insulate it from the backplane of the silicon sensors, which carries the sensor bias voltage of up to 500 V. The support plate extends by 1 mm over the edges of the silicon sensors to protect these mechanically. Its edges are sealed with a non-conductive glue (Araldite) to prevent loose fibres from sticking out and touching the sensors, where they might cause a short circuit between strip side and backplane. A rectangular cutout at the location of the readout hybrids permits to insert the *mini-balcony* into the support plate, to which it is fixed using thermally conductive glue (Tra-Duct 2902). The mini-balcony is a 70 mm$\times$15 mm large, precisely machined rectangular piece of aluminium. It is 1.5 mm thick and its flatness is guaranteed to be better than 30 μm. It contains precision holes for the mounting and exact positioning of the module in the detector box and it defines the thermal interface of the module to the cooling rod onto which it is mounted. There are two types of mini-balconies: one for modules that will be mounted vertically (for the x detection layers) and another one for modules that will be mounted under an angle of 5° (for the u and v detection layers). The mini-balcony provides a direct heat path from the Beetle chips to the cooling rod and it brings the cooling rod into thermal contact with the carbon-fibre sheets of the module support plate. These carbon-fibre sheets form large cold surfaces that contribute to the cooling of silicon sensors and detector volume. The mini-balconies are produced by Atelier Mécanique Di Chiara, Switzerland, and the module support plates are produced by Composite Design, Switzerland.

To avoid the risk of mechanical stress, the silicon sensors are glued onto the module support plate using thin strips of non-hardening silicone glue (Dr. Neumann NEE-001-weiss). The hybrids are glued onto the module support plate using a two-component glue (Araldite). Small spots of conductive silver glue are applied at the location of the Beetle chips in order to improve thermal contact and to provide a direct ground connection between the Beetle chips and the mini-balcony.

IT Detector Boxes

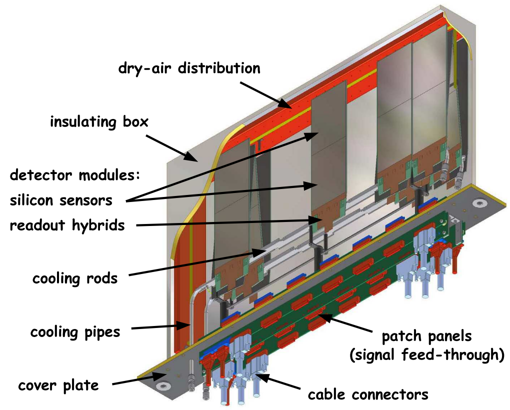

Figure 28. View of an IT side box. Top/bottom boxes are similar except that the box is shorter and contains one-sensor modules.

An isometric view of a detector box is shown in figure 28. Its main structural element is a honeycomb plate (“cover plate”), onto which two cooling rods are mounted. These cooling rods incorporate cooling pipes through which C_6F_{14} at -15°C circulates as a cooling agent. Printed-circuit boards (“feedthrough PCBs”) that are inserted vertically through the cover plate serve to transmit all electrical signals (detector signals, control signals and supply voltages) from and to the detector modules inside the box. The detector volume is closed by an insulating box that is assembled from flat sheets of a light but rigid aluminium-clad foam. The detector volume is continuously flushed with nitrogen to avoid condensation on the cold surfaces.

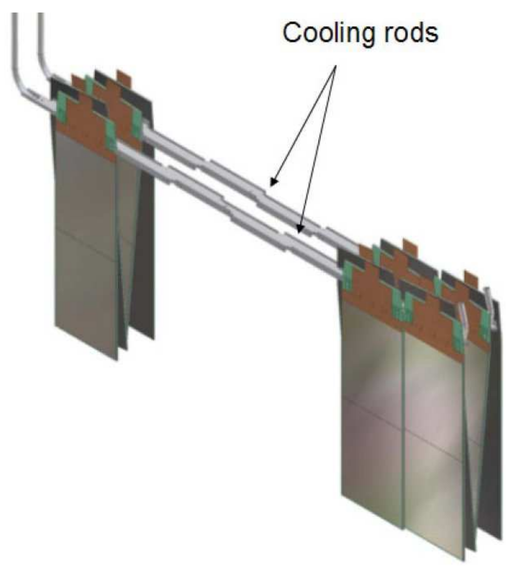

Figure 29. View of an IT cooling rod with a few detector modules.

An isometric view of a *cooling rod* is shown in figure29. It is precisely machined out of a single piece of aluminium and consists of a 3 mm thick central part with 6 mm high and about 70 mm wide mounting surfaces for the detector modules. The accurate positioning of the detector modules is ensured by precision pins in these mounting surfaces. Detector modules are mounted on the cooling rod from both sides, i.e. each cooling rod supports a pair of detection layers, (x,u) or (v,x). The shape of the cooling rod accomodates the staggering of adjacent detector modules in each of the two detection layers. Milled into the lower surface of the central part of the cooling rod is a semi-circular groove into which an aluminium cooling pipe with an outer diameter of 6 mm and a wall thickness of 0.4 mm is glued. The cooling pipes on the two cooling rods in a detector box are connected in series using a short U-shaped rubber hose (Nitril).

The two cooling rods are mounted onto the *cover plate* using carbon-fibre pillars. The cover plate itself is made out of a 14 mm thick polymethacrylimide foam (Rohacell) sandwiched in between two carbon-fibre skins.

Four *feedthrough PCBs*, one for each detection layer, are inserted vertically through the cover plate. They are four-layer printed-circuit boards. The outer layers carry the bias voltage and the analog and digital supply voltage, respectively. The two inners layers are used for the differential signals. At their end inside the detector box, the feedthrough PCBs carry one 60-pin board-to-board connector for each detector module, into which the Kapton tail of the readout hybrids is plugged. On the part outside of the box they are equipped with separate low-mass connectors for signal and bias voltage cables.

The common ground for all modules in a detector box is defined by the cooling rods. The ground connection between module and cooling rod is provided by the conductive glue with which the readout hybrids are glued onto the mini-balconies and by thin copper wires that connect ground pads on the Kapton flex prints to the metallic screws that are used to fix the modules on the cooling rod.

The thermally and electrically insulating box *enclosure* is assembled from 6 mm thick, flat sheets of polyisocyanurat (PIR) foam that are reinforced with a single, 200 μm thick carbon-fibre skin and clad on both sides with 25 μm thick aluminium foil. For the side of the box facing the beam pipe, the wall thickness is reduced to 2 mm to minimise the distance between the beam and the innermost detector modules. Mounted on the inside wall of the enclosure is a distribution channel for the nitrogen with which the box is flushed. Small inserts made out of fibreglass (Stesalit) are embedded in the upper rim of the enclosure and permit to screw it onto the cover plate.

IT Support Frames

Each IT station has two separate support frames, one to each side of the beam pipe. Each support frame holds one side box and either a top or a bottom box. An isometric view of a support frame is shown in figure 30. It hangs from a support rail that is mounted onto the Outer Tracker bridge (see Sec. **??**) and is guided by a lower rail that is mounted onto the LHCb bunker. The inner-most section of these support rails is precision-machined to ensure an accurate positioning of the support frames in data taking position. The support frame itself is assembled from rectangular fibreglass/carbon-fibre rods with a cross section of 70 mm $\times$ 70 mm and a wall thickness of 3 mm, and from flat honeycomb plates made out of aramid clad with carbon-fibre skins. Signal and bias-voltage cables as well as flexible supply lines for the liquid C_6F_{14} coolant and for nitrogen are routed along the support frames from the detector boxes to the lower end of the station. Here, the signal cables are connected to four Service Boxes (see Sec. 4.2.3) that are mounted onto the lower

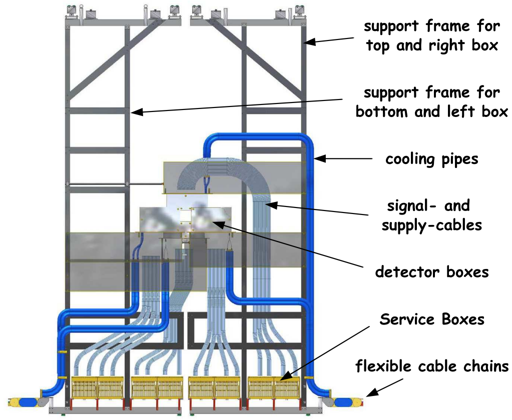

Figure 30. View of an IT support frame.

end of the support frame, just outside of the acceptance of the experiment. Cables and supply lines are further routed through a flexible cable chain that is fixed to the lower end of the support frame at one end and to the edge of the LHCb bunker at the other end.

4.2.3 Electronics

Both TT and IT use the Beetle front-end readout chip [28]. The analog output signals of the Beetle chips are transmitted via copper cables from the detector box to the Service Boxes [29]. Here, they are fed into Digitizer Cards, on which they are digitised, multiplexed, converted to optical. They are further transmitted via 120 m long optical fibres to the counting house, where they are received on the TELL1 board described in Sec. **??**. The Service Box is a custom-made crate that holds several Digitizer Cards, a Control Card that provides the interfaces to the Timing and Fast Control system (TFC) and the Experiment Control System (ECS), and a backplane for the distribution of control signals and low voltage. Low voltage supplies are based on the Maraton system developed for the LHC experiments. Bias voltage for the silicon sensors is provided by commercial high-voltage supplies and is supplied to the detector boxes via dedicated high-voltage cables. Some additional cables are installed for monitoring purposes. The grounding scheme of the Silicon Tracker follows the LHCb grounding rules.

Readout Electronics

Four, respectively three, Beetle chips are located on a TT or IT front-end readout hybrid. Each Beetle chip amplifies and shapes the detector signals of 128 readout strips, samples them at the LHC bunch-crossing frequency of 40 MHz, stores the sampled data in an analog pipeline, and upon a Level-0 trigger accept transmits them multiplexed 32-fold via four output ports. The analog output signals from each readout hybrid are routed to a printed-circuit board on the outside of the detector boxes, as described in Sec. 4.2.1 and 4.2.2. From here, the data from each hybrid are transmitted to the Service Boxes via one approximately 5 m long, shielded 68-wire twisted-pair cable. The wires not used for the transmission of detector signals are used to provide timing and slow-control signals, low voltage and ground from the Service Box to the detector module. Since in the case of the IT these cables run accross the acceptance of the experiment, custom-made cables with reduced shielding braids are employed to minimise the material budget.

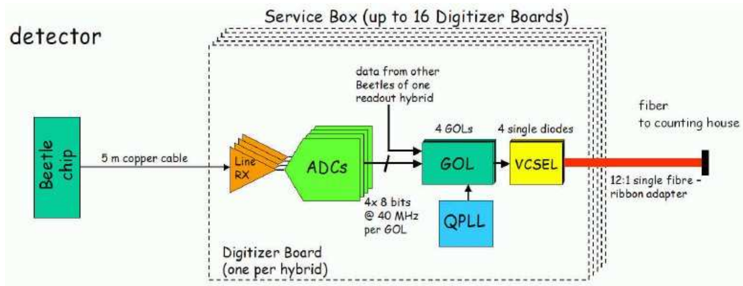

Figure 31. Sketch of the functional block of the Digitiser Card that is used to processed the data from one Beetle chip.

At the Service Box end, each twisted-pair cable plugs into one **Digitizer Card**. Each Digitizer Card therefore processes the data from one readout hybrid. There are two variants of the Digitizer card: A TT version to match the four-chip TT hybrids and an IT version to match the three-chip IT hybrids. The functional block for processing the data from one Beetle chip is illustrated in figure 31. Four differential line receivers convert the signals of the four Beetle output ports from differential to single-sided and match the signal levels to the input range of the four single-channel 8-bit ADC chips that are used to digitise the data. The ADC chips operate at 40 MHz and are phase-locked to the sampling clock of the Beetle chip. The 32-bit wide data stream from the four ADC chips is fed into a single CERN Gigabit Optical Link (GOL) chip [30], which multiplexes and encodes it to a Gigabit Ethernet signal with a data rate of 1.6 Gbit/s. A laser driver integrated into the GOL chip is used to control a 850 nm wavelength VCSEL diode and transmit the digitised optical data into a single optical fibre. This functional block is repeated four times on a TT Digitizer Card and three times on an IT Digitizer Card. In addition, both variants of the card carry a central block for the distribution of the timing and control signals that the Digitizer Card receives from the Service-Box Control Card.

To transmit the data to the counting house, the outputs of 12 VCSEL diodes, corresponding to three TT Digitiser Cards or four IT Digitiser Cards, are connected to one 12-fibre parallel optical cable. This grouping defines a readout partitioning of the detector into groups of three readout sectors for TT and four detector modules for IT. The same partitioning is followed by the low-voltage and bias-voltage distribution, and by the detector control system. A TT Service Box holds 12 (four groups of three), an IT Service Box 16 (four groups of four) Digitizer Cards. Both types of Service Box therefore feed four 12-fibre cables. In the counting house, each TELL1 board receives two 12-fibre cables.

Detector Control and Monitoring

Detector control and monitoring is provided by the *Control Cards* [31], one of which is located in each Service Box.

Each Control Card holds a TTCrq mezzanine (see Sec. **??**), which collects clock, trigger and timing information from the TFC network. These signals are distributed to the Digitizer Boards via impedance-controlled differential traces on the Service-Box backplane. All GOL chips and Beetle chips associated with a single Service Box receive the same clock signal. Care was therefore taken in the design of the backplane to equalise the trace lengths and signal propagation times to each Digitizer Card slot. The same holds for the distribution of the signals on the Digitizer Cards themselves and for the lengths of the signal cables between Digitizer Cards and front-end hybrids.

The interface to the ECS is provided by two SPECS slave mezzanines (see Sec. **??**) that are mounted on the Control Card and provide a total of eight I^2C busses. Four of these are used to control the GOL chips per group of three (TT) or four (IT) Digitizer Cards, the other four are used to control the Beetle chips per group of three (TT) or four (IT) readout hybrids. In addition, 36 I/O control lines permit to individually control radiation-tolerant low-voltage regulators (LHC4913) that provide the power for each readout hybrid and for each group of Digitizer Cards.

The SPECS mezzanines also provide a number of ADC channels that are employed to read out temperature sensors (PT1000) and humidity sensors at various locations in the detector boxes.

Additional ADC chips are located on each Digitizer Board and are used to monitor over-current conditions of the Beetle chips and to read out a PT1000 temperature sensor that is located on the readout hybrid.

Power Distribution and Grounding Scheme

Low voltage levels of 2.5 V, 3.3 V and 5 V, are required for Beetle, GOL and ADC chips, line receivers, and various LVDS drivers on the Digitizer Cards and the Service Box backplane, respectively. They are derived from voltage levels of about 5.5 V and 8 V that are generated by Maraton power supplies in the LHCb cavern. The Maraton supplies are connected to the Service-Box backplanes where they drive radiation-tolerant programmable linear power regulators (L4913, developed by the CERN micro-electronics group). Two separate low-voltage regulators provide each readout hybrid with the digital and analog power (analog and digital power for the Beetle chips are kept separate throughout the system and are connected only on the readout hybrids). Additional regulators provide each group of four (IT) resp. three (TT) Digitizer Cards with the required voltage levels, following the same partitioning as the readout. All power regulators are located on the Service-Box backplane. To remove the heat that they generate, the backplane is connected to the LHCb mixed-water cooling system. The regulators for each readout-hybrid and the regulators for each group of Digitizer Cards can be individually switched off via ECS.

Detector *bias voltage* is supplied by a commercial HV system (CAEN SY1527 crate with A1511B HV modules) located in the counting house. It is connected to the detector boxes via 120 m long HV cables. The HV modules can provide up to 500 V and can deliver a current of 10 mA per channel. With the exception of the innermost readout sectors of the TT, which have individual HV channels to cope with the sensor leakage currents expected after several years of operation, groups of three TT readout sectors and four IT detector modules are connected to one HV channel, following the same partitioning as the readout. Each readout sector and detector module, however, has a separate supply line up to a patch panel that is located in the counting house, close to the HV supplies. On this patch panel, jumpers permit to manually disconnect each individual readout sector and detector module from its HV supply.

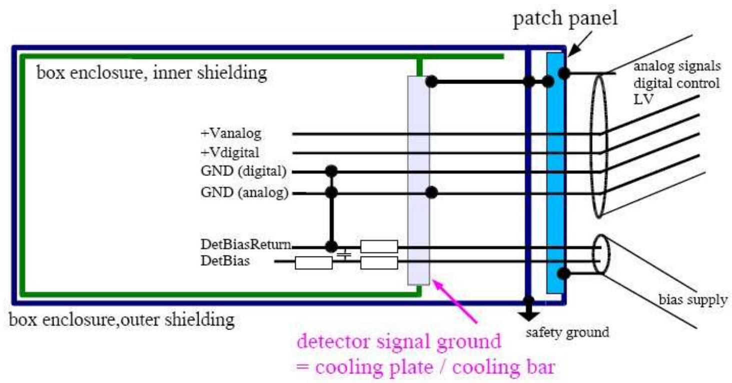

Figure 32. Grounding scheme of the detector boxes.

The *grounding scheme* [32] of the detector boxes is illustrated in figure 32. The common ground for each detector box is defined by the cooling plates (TT) or cooling rods (IT), onto which the detector modules are mounted as described in Sec. 4.2.1 and 4.2.2. The walls of the detector boxes are coated with 25 μm thick aluminium foils on both inside and outside. The inner shielding is connected to the common ground of the detector box, the outer shielding is connected to the LHCb safety ground. The two are connected with each other at one well-defined location in the detector box. The shields of all signal and supply cables entering the detector box are connected to the outer box shielding. The LV power-supply system is kept completely floating and is connected to the LHCb general ground only in the detector boxes. Bias-voltage supply lines are filtered by low-pass RC filters on the hybrids.

4.2.4 Detector performance

An extensive R&D programme, including various tests of prototype detectors in the laboratory and in test-beams as well as simulation studies, has been carried out to validate the detector concept, optimise detector parameters and estimate the expected performance of the detectors in LHCb [33, 34, 35, 36, 37]. In view of the combination of long readout strips and fast pulse shapes employed, the signal-to-noise performance of the detectors was a major concern in these studies. Other studies concerned the expected strip occupancies, which were estimated using samples of simulated events generated using the full Geant4-based LHCb Monte-Carlo, and the material budget of the detector.

Signal-to-noise and efficiency

Various prototype detectors, with effective readout strip lengths of 108 mm up to 324 mm were constructed from silicon sensors of 320 μm to 500 μm in thickness, strip pitches between 183 μm and 228 μm, and ratios of strip implant width over strip pitch of 0.25 to 0.35. Their noise performance, charge collection efficiency and signal-to-noise performance were measured in an infra-red laser test stand and in several test-beams periods in a 120 GeV π^- beam at CERN. Some tested modules included Kapton interconnect cables of the same length as used in the M and K readout sectors of TT. The expected noise performance of the various detector configurations was also investigated in a SPICE simulation, which included a detailed description of the Beetle front-end and in which the readout strips of the detectors were described as an extended LCR network. The results of this simulation agreed with the measurements.

The measured *noise performance* from the test-beams is summarised in figure 33, in which the measured equivalent noise charge (ENC) for the various tested detector configurations is shown as a function of their total strip capacitance. A linear dependence is observed and the slope of a line fitted to the data agrees well with that obtained in test-bench measurements in which discrete load capacitances were attached to the Beetle inputs. Both the measurements and the SPICE simulations demonstrated that the Kapton interconnect cable behaves purely as an additional load capacitance and has no adverse effect on the integrity of the detector signals.

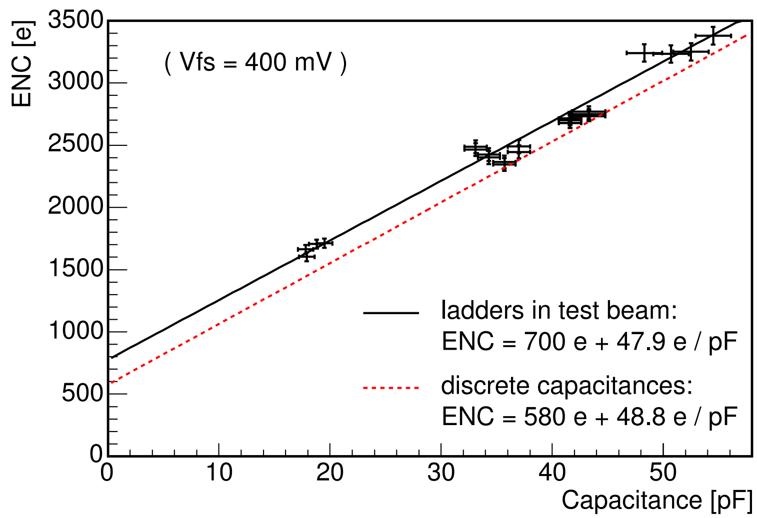

Figure 33. ENC obtained in test-beam measurements as a function of the measured total strip capacitance of the tested prototype modules.

Measurements of the *charge collection efficiency* as a function of the position on the detector showed no observable dependence on the position along the readout strip but a strong dependence on the position orthogonal to the strips. As illustrated in figure 34, a significant drop in charge-collection efficiency in the central region between two readout strips was observed. This charge loss was observed for all tested detector configurations. Its size was found to depend on the readout strip geometry, decreasing with increasing ratio of implant width to strip pitch (w/p) and with increasing sensor thickness (d). In fact, it was found to depend roughly linearly on the ratio $(p-w)/d$. The charge loss could not be reduced by over-biasing the detectors or by increasing the shaping time within the limits allowed by the Beetle chips. It is attributed to charge trapping at the interface between silicon bulk and silicon oxide in between the readout strips.

Full *particle detection efficiency* above 99.8% was measured for all detector configurations as long as the most probable signal-to-noise ratio stayed above 10:1. Below that value, the particle detection efficiency started to decrease rather quickly. With the chosen sensor thicknesses, most probable signal-to-noise ratios in excess of 12:1 can be expected over the full surface of the detectors for both types of IT modules and all four types of TT readout sectors.

Strip occupancies

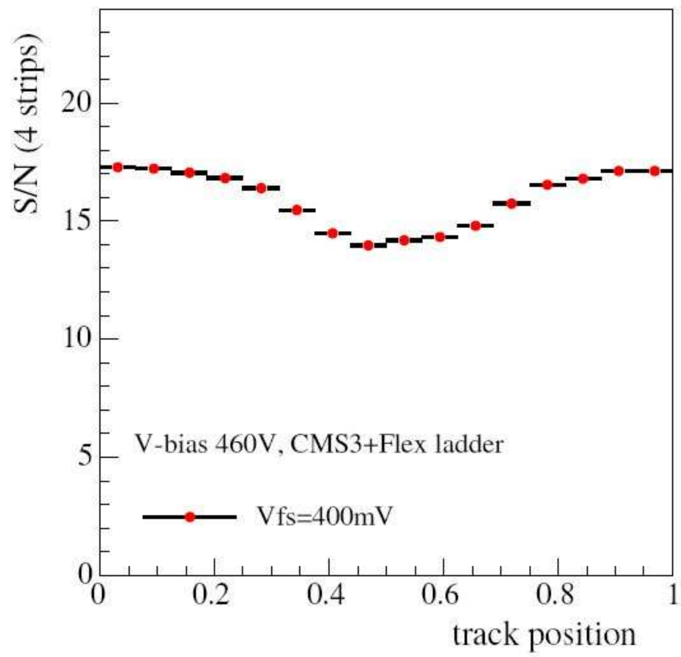

Figure 34. Most probable signal-to-noise ratio as a function of the inter-strip position (0 = centre of left strip, 1 = centre of right strip), measured in the CERN test-beam for a prototype module with the detector geometry of a three-sensor readout sector of TT.

The strip occupancies presented here [38] were obtained from a sample of 5000 $B_d \to J/\psi(\mu^+\mu^-)K_S(\pi^+\pi^-)$ events from the DC06 Monte-Carlo production. The simulation software includes a detailed description of the LHCb detector geometry as well as of the signal collection, amplification and digitisation. The digitisation software was tuned to reproduce results of test-beam measurements performed on prototype detectors. For TT, average strip occupancies of up to about 3.5% are found in the K sectors close to the beam pipe. Occupancies drop to about 0.35% in the outermost L sectors. For the side boxes of the first IT station, average strip occupancies drop from about 2.5% on the strips closest to the beam to about 0.5% on the outermost strips. In the top/bottom boxes, occupancies vary between 0.5% and 0.3%. In the second and the third IT station, occupancies are about 10%, respectively 20% lower than in the first. This is mainly due to the larger distance of these stations from the beam-pipe supports at the exit of the LHCb magnet, which are a prolific source of secondary particles.

Material budget

A careful analysis of all materials located inside the acceptance of the experiment was performed. A detailed description of all active detector elements and a simplified description of the passive components have been implemented [39, 40] in the XML-based LHCb detector description framework.

For TT, where most of the dead material from detector supports, cooling etc. is located outside the acceptance, the material distribution is rather uniform. It corresponds to about 4% X_0, where more than 2% X_0 are due to the active material of the silicon sensors. An increase up to almost 8% X_0 is observed in the very forward region, due to the material of the thermal insulation around the beam pipe. The result of a material scan that was performed using the LHCb transport service is shown in figure 35.

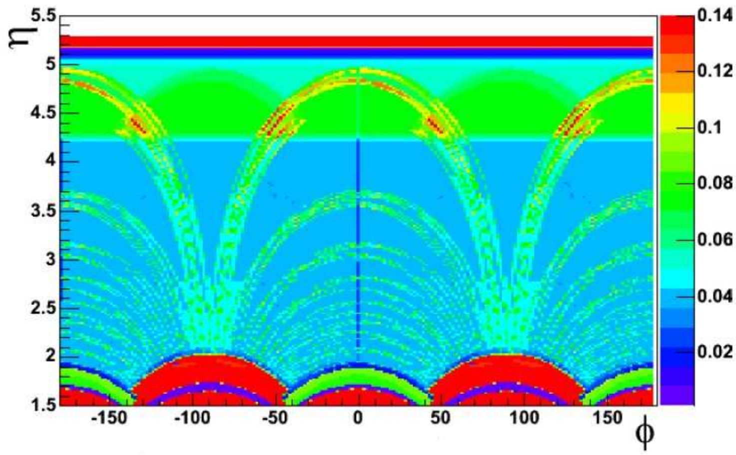

Figure 35. TT radiation length as a function of the pseudo-rapidity η and the azimuthal angle ϕ.

The material distribution for IT is much less uniform, due to the readout hybrids, mechanical supports, cooling pipes and cables that are located inside the LHCb acceptance. In the active region of the detector, close to the beam pipe, the material budget corresponds to about 3.5% X_0 per station, out of which more than 1.5% X_0 are due to the active material of the silicon sensors. The material budget peaks at almost 30% X_0 in the very narrow region of the cooling rod. The result of a material scan for one IT station is shown in figure 36.

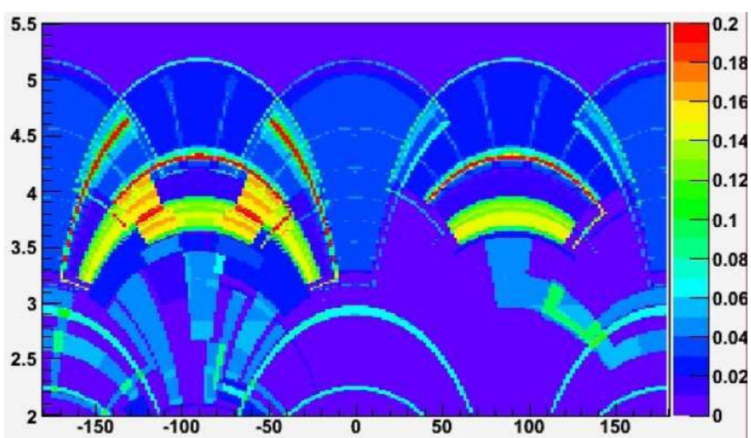

Figure 36. Radiation length for one IT station as a function of the pseudo-rapidity η and the azimuthal angle ϕ.

4.3 Outer Tracker

4.3.1 Introduction and General Layout

4.3.2 Detector Description

The LHCb Outer Tracker (OT) is designed as a modular system. One module contains two staggered layers of 64 straw tubes each, see Fig. 37. The channels in the long F-modules are read out at both ends and are split in the middle of the module. The exact separation point is staggered between the two straw tube layers within the detector module to avoid insensitive regions in the OT. F-modules have a total active length of 4850 mm, see Fig. 38, and contain a total of 256 straws. The short S-modules are about half the length and are located above and below the beampipe. The modules are arranged in three identical stations, each station housing four detector layers. The four layers in one station are arranged in a X-U-V-X geometry. The straw tubes of the X-layers are positioned vertically, whereas the U and V layers are tilted by $\pm\ 5^o$ with respect to the X-layers. The total active area of a station is $5971{\times}4850\text{mm}^2$.

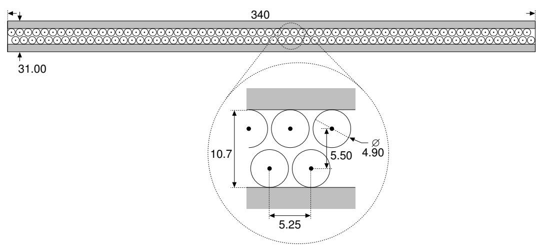

Figure 37. *Module cross section.*

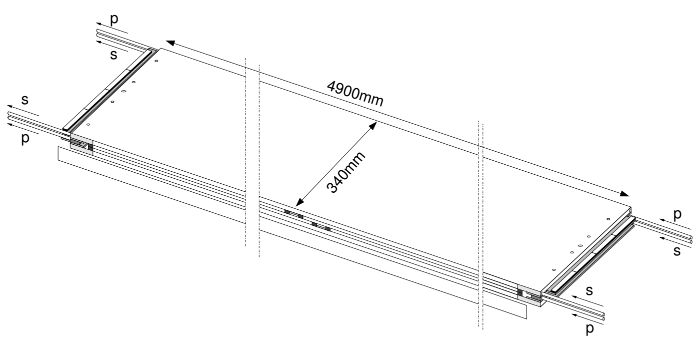

Figure 38. *Drawing of a module.*

The design of the detector modules is based on the concept to construct independent, robust detector modules capable for stand-alone operation, resulting in the following requirements:

- Robustness: The mechanical stability of the detectors must guarantee the position of the straw tubes and the anode wire within a precision of 100(500) μm in the $x(z)$ direction. The wire has to be centered inside the straw tube within 50 μm over the entire length of the straw. It should be sufficiently rigid to withstand an overpressure of 10 mbar. To meet the requirements on the gas quality the leak rate must not exceed $8\times\ 10^{-4}$ l/s [42].

- Material budget: To reach the design momentum resolution the material introduced in the active area of the Outer Tracking system should not exceed 3% of X0 per station.

- Electrical shielding: The straw tubes must be properly shielded from the neighbouring straws to avoid crosstalk. A firm connection to the detector module ground has to be guaranteed. The detector modules form Faraday cages with a firm connection electrical connection to the front-end readout electronics.

- Radiation hardness: The detector modules have to built from non-outgassing and radiation resistant materials. The detector is designed to allow to accumulate an integrated charge up to 2C/cm without degradation of the detector performance.

The straw tubes are winded using two foils as shown in Fig. 39. The inner cathode foil is made of 40 μm carbon doped Kapton-XC. The outer foil is a laminate made of 25 μm electrically conducting Kapton-XC and 12.5 μm aluminium (*GTS-foil*). The usage of the laminate guarantees gas tightness of the straws. The aluminium layer is mandatory for the signal transmission and to minimise crosstalk between adjacent straws. The anode wire is made of gold plated tungsten with a diameter of 25 μm, produced by California Fine Wire. Wires are strung with a load of 70 g. The position of the wire in the centre of the straw is guaranteed by injection molded Noryl wire locators and endpieces. Two wire locators are placed every 80 cm in each straw.

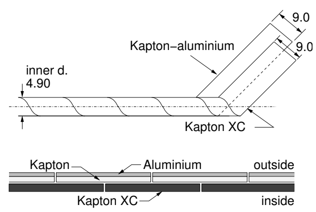

Figure 39. *The straws are winded using two foils, kapton-XC and a laminate of kapton and aluminium.*

The straws are glued to sandwich panels with 120 μm carbon-fibre skins and a 10 mm Rohacell core. An aluminium template was produced to guarantee the required precise positioning of the straws with a precision of better than 50 μm over the entire length of the detector module. Special tools were designed to keep the panels flat with a precision of better than 100 μm during the entire production period of the modules. Finally, the panels are joined by 400 μm thick carbon fibre sidewalls, resulting in a standalone detector module. The panels and the side walls are covered by a laminated foil of 25 μm Kapton and 12.5 μm aluminium to guarantee gas-tightness of the box and to provide a closed Faraday cage. Spacers at the ends of the module separate the two panels apart and pass the gas to the module. All glueing steps are performed using Araldite AY103 with the hardener HY991, cured at room temperature. To enhance the viscosity silica bubbles are added.

SOME TEXT ON SUPPORT (Cframes, bridge, ...)?

Material Budget From the measurements of the weights of the various components, the amount of material inside the LHCb acceptance can be estimated [44]. In particular the separate contributions from the panels, the glue and the straws to the total radiation length is estimated, as can be seen in Table 4. The values for the thickness of the aluminium laminated kapton foil (*GTS foil*), and for the straw layers are taken from the design values. The total amount of glue inside a panel is 140 g, whereas the glue used between the straws and the panel and to close the module amounts to 400 g. The side walls contribute an additional $X/X_0 = 0.048\%$ to the total radiation length. of one detector module. For one entire station, consisting of four module layers (i.e. 8 mono-layers) the radiation length X/X_0 amounts to 3.17%.

Material	Quan.	d(μm)	X_0(cm)	X/X_0(%)
Al	1	12	8.9	0.013
Glue	1	10	(36.1)	0.006
Kapton	1	25	28.6	0.009
GTS Subtotal	1	47	-	0.025
Rohacell	1	10240	1400	0.073
CFC Skin	2	90	(23.3)	0.077
Panel Glue	2	84	(36.1)	0.046
Panel Subtotal		10635		**0.222**
Str.Al	$2\frac{\pi}{2}$	12	8.9	0.042
Str.Glue	$2\frac{\pi}{2}$	10	(36.1)	0.009
Str.Kapton	$2\frac{\pi}{2}$	25	28.6	0.027
Str. GTS Subtotal	$2\frac{\pi}{2}$	47	-	0.079
Str. Glue	$2\frac{\pi}{2}$	10	(36.1)	0.009
Str. Kapton XC	$2\frac{\pi}{2}$	40	28.6	0.044
Straw Subtotal		97		**0.131**
Glue Str-to-Panel	1	70	(36.1)	**0.019**
Total				**0.372**

Table 4. *The various contributions to the radiation length inside the LHCb acceptance are listed for one mono-layer(i.e. one half-module). Note that the circular shape of the straws are taken into account with the factor $\frac{\pi}{2}$. The numbers in brackets are estimates.*

4.3.3 Electronics

The drift-times of the ionization clusters produced by traversing charged particles are measured for every straw tube. The reference for the time measurement is obtained from the beam crossing (BX) signal. The hit times are digitized for every bunch crossing (25 ns) and stored in a digital pipeline to await the Level-0 (L0) decision. On a positive L0 decision, the digitized data is transmitted via optical links to the Level-1 (L1) buffer boards (TELL1).

A schematic of the front-end electronics is shown in Fig. 40, and the elements of the OT specific readout-electronics are described below:

- Feedthrough board: This passive printed-circuit board (PCB) is part of the module. It provides the electrical connections to the straws and it defines the reference ground for electronics and straws.

- High-voltage board: The board distributes the high-voltage to the straw channels and contains capacitors for the signal decoupling from the chamber HV. It is plugged directly to the feed-through board.

- ASDBLR board: The board carries two ASDBLR chips. The ASDBLR ASIC is an 8-channel amplifier-shaper-discriminator with ion-tail cancellation and baseline restoration designed for the ATLAS TRT [46].

- OTIS board: Two ASDBLR boards are plugged to a common board for the time measurement. The board carries the OTIS, a 32-channel TDC-chip which was especially designed for the OT [48, 49]. On a L0 accept the OTIS provides the digital event data on an 8-bit differential output.

- GOL/Aux board: Four OTIS boards are connected to one GOL/Aux board. The GOL/Aux board [51, 52] provides power, fast-control signals (BX clock, resets, L0 trigger) and slow-control signals (I^2C) to the OTIS and ASDBLR boards. The GOL serializer chip is used to mulitplex the 32-bit OTIS information (4×8-bit) for a 1.6 Gbit/s optical link. The data is received by an optical receiver card, used as mezzanine card on the TELL1 board.

The above front-end boards all house inside the front-end electronics box. The electronics boxes are mounted on the upper and lower end of a module. High-voltage and low-voltage cables, fast and slow-control cables as well as the optical fibers are connected to each electronics box individually. One front-end box is the smallest independent readout unit of the OT. Its output is an optical fiber providing the signals of the 128 channels of the module half (upper/lower) to the L1 buffer board.

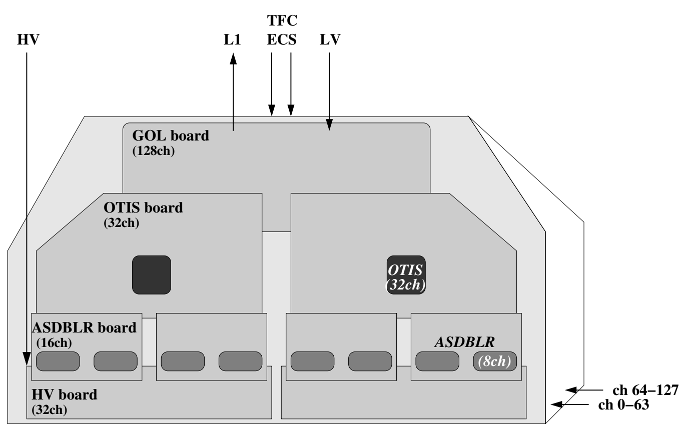

Figure 40. *Schematic of the front-end readout electronics and its boards.*

4.3.4 Test-Beam Results

In order to validate the combination of detector and frontend electronics, four mass-production S-modules were tested with a 6 GeV electron beam at the DESY-II facility in Hamburg [43].

The measurement of the detector parameters, such as efficiency and resolution, requires knowledge on the coordinate at which the beam particle traverses the detector plane. In order to have the possibility to independently measure track parameters, a silicon strip telescope was used. However, the total number of OT modules (4 modules, or 8 monolayers) was sufficient to allow track reconstruction using the OT data only. Both studies were systematically compared, yielding identical results. The trigger signal, which served also as time reference, was produced by coincidence of two scintillator counters installed downstream of the OT modules. The beam illuminated 2–7 straws per OT plane. The analysis to estimate the performance of the OT detector and its readout electronics, consist of the following steps:

1) Attain the predicted distance of closest approach of the particle to the sense wire, obtained either by the silicon telescope or by the OT standalone.

2) Establish the relation between the measured drift time and the predicted distance to the wire, i.e. the rt-relation, see Fig. 41.

3) Convert each measured drift time into position coordinate.

With the measured position coordinate in hand, the resolution of this measurement can be obtained by comparing it to the predicted position. The hit finding efficiency is attained by verifying whether the OT produced a hit at the predicted position. The efficiency as a function of the distance to the wire is shown in Fig. 41.

The performance of the OT has been measured for values of the high voltage ranging between 1200 and 1700 V, and of the amplifier threshold ranging from 500 till 900 mV, corresponding to 1.5 and 5.5 fC, respectively. The resolution is obtained by a single Gaussian fit to the residual distribution, and the required resolution of 200μm is achieved for high voltage values above 1550 V, see Fig. 42. The efficiency is defined as the average cell efficiency, i.e. the probability to find a hit in the straw where the track points to, or in one of its neighbours. As shown in Fig. 42, efficiencies close to 100% and position resolutions below 150 μm can be attained.

The electronic noise level was found to be low: e.g. at a high voltage of 1550 V and amplifier threshold of 800 mV (corresponding to 4 fC), the noise level is lower than 1 kHz. This translates in an occupancy of less than $7.5 \cdot 10^{-4}$, given the high LHC bunch frequency of 40 MHz. Finally, the crosstalk betwe … find a coherent hit in the neighbouring st … V and with a amplifier threshold of 800 m

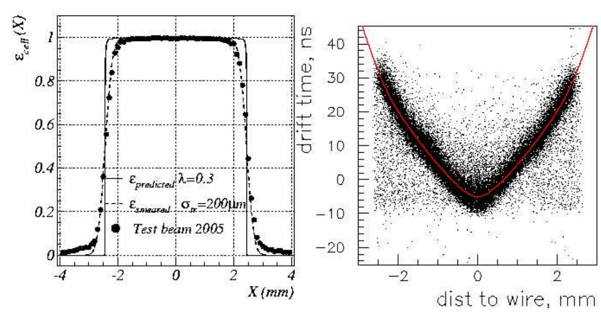

Figure 41. *The efficiency profile inside the straw is given. The efficiency is larger than 99% in the center of the straw, and drops of at the edge of the straw. The relation between the drift-time and distance to the wire is shown on the right. Note the maximum drift-time of approximately 45 ns.*

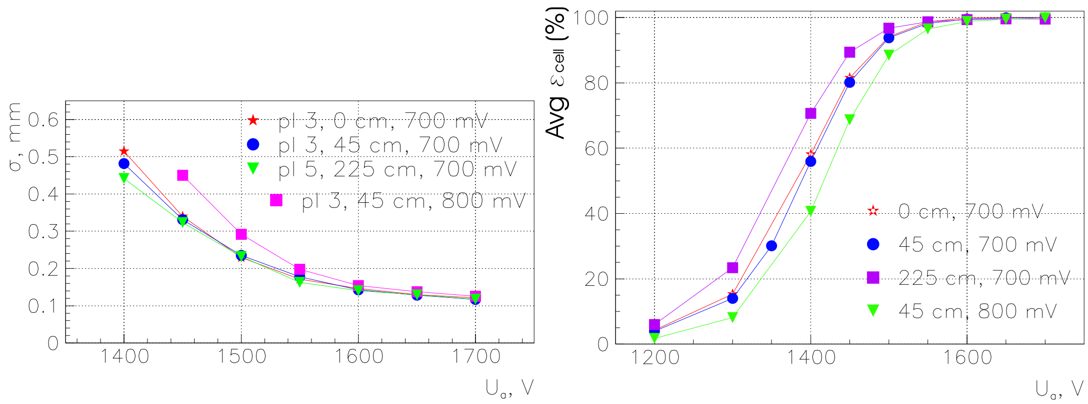

Figure 42. *The detector performance as determined at a beam test. Both the position resolution and the average cell efficiency are shown as a function of the high voltage value. The different curves correspond to different amplifier settings, and different distances of the beam to the front-end electronics.*

4.3.5 Alignment and Monitoring

Errors in the mechanical alignment of the detector components can significantly degrade the momentum measurement originating from the track reconstruction. The alignment of the OT detector modules is assured by several steps. First, the detector is constructed satisfying the stringent requirements as described above [53, 54]. Secondly, during installation the position of the 12 supporting C-frames carying 18 modules each is surveyed and adjusted accordingly to the desired position within 2 mm [55].

The ultimate alignment will be done on the basis of track reconstruction, but to ease the alignment procedure the station alignment is monitored by means of a hardware system that is capable to register relative movements of stations as a function of time. The required precision of the alignment monitoring should be smaller compared to the expected detector resolution for single hits. The single cell resolution of about 200 μm translates in a required envisaged precision of the system below 100 μm. The alignment system is based on the RASNIK monitors developed at NIKHEF [56] and has been adapted to the needs of the Outer Tracker. The basic idea of the RASNIK system is to project a detailed image through a lens onto a CCD camera as shown in Fig. 43. A movement of any of these three elements results in a corresponding movement of the image on the CCD camera. Movements perpendicular to the axis are observed by the change of the image position processed by the CCD camera, whereas movements along the axis are measured by a change in the size of the image. The intrinsic resolution of the system perpendicular(parallel) to the beam axis is better than 10(150) μm.

5. Particle ID

5.1 RICH

5.1.1 Introduction

Particle identification (PID) is a fundamental requirement for LHCb. It is essential for the physics to separate pions from kaons in selected B meson decays. At wide polar angles the momentum

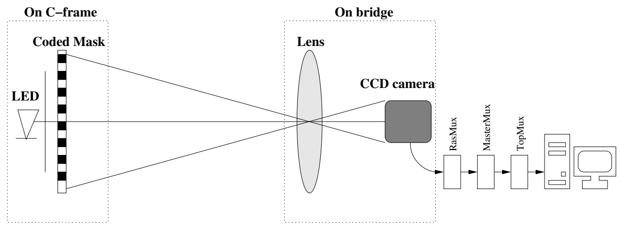

Figure 43. *Schematic of the RASNIK alignment system.*

spectrum is softer; hence the particle identification system consists of two RICH detectors. The upstream detector, RICH 1, covers the low momentum charged particle range ∼1-60 GeV/c using aerogel and C_4F_{10} radiators, while the downstream detector, RICH 2, covers the high momentum range from ∼15 GeV/c up to and beyond 100 GeV/c using a CF_4 radiator (see Fig. 44). RICH 1 has a wide acceptance covering the full LHCb acceptance from ±25 mrad to ±300 mrad (horizontal) and ±250 mrad (vertical) and is located upstream of the magnet to detect the low momentum particles. RICH 2 is located downstream of the magnet and has a limited acceptance of ±25 mrad to ±120 mrad (horizontal) and ±100 mrad (vertical) where there are mainly high momentum particles. In both RICH detectors the focusing of the Cherenkov light is accomplished using a combination of spherical and flat mirrors to reflect the image out of the spectrometer acceptance. RICH 1 has a vertical optical layout symmetry whereas for RICH 2 the symmetry is horizontal. Hybrid Photon Detectors (HPDs) are used to detect the Cherenkov photons in the wavelength range 200-600 nm. The HPDs are surrounded by external iron shields and are placed in Mumetal cylinders to permit operation in magnetic fields up to 50 mT. The RICH detector system including its electronics, monitoring, control, and the performance are described below.

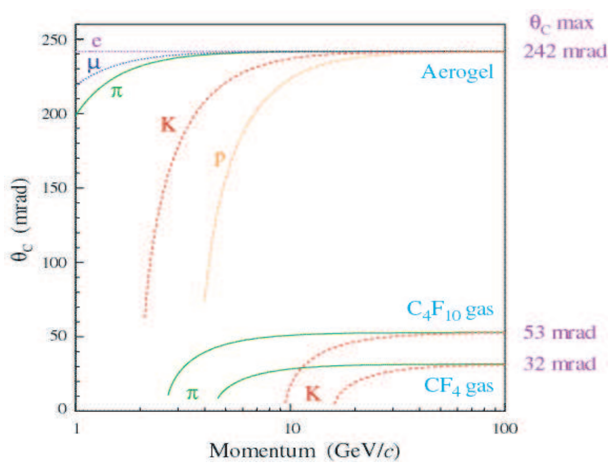

Figure 44. Cherenkov angle versus particle momentum for the RICH radiators.

5.1.2 RICH 1

The RICH 1 detector [57, 58] is located upstream of the LHCb dipole magnet, between the VELO and the Trigger Tracker (see chapter 4). RICH 1 contains aerogel and fluorobutane (C_4F_{10}) gas

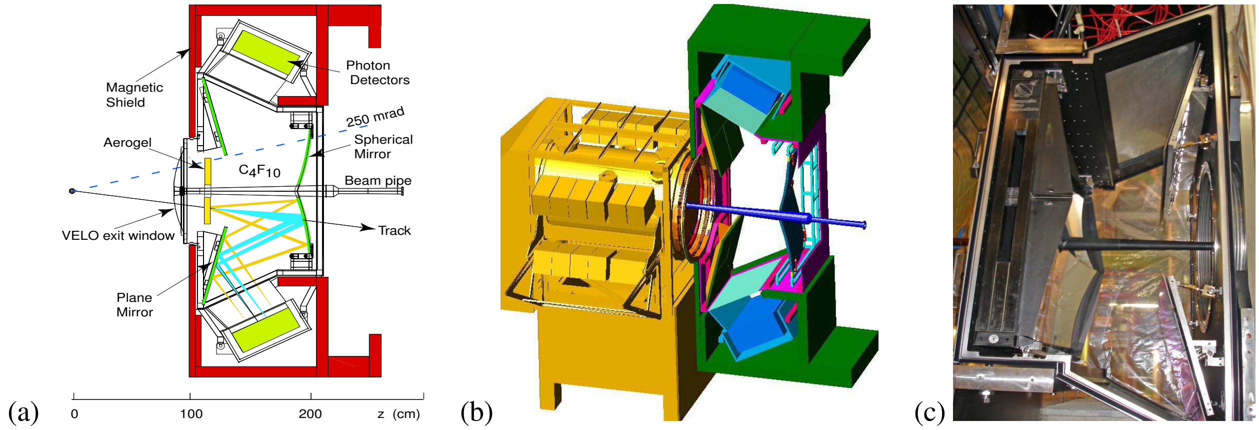

Figure 45. (a) Schematic layout of the RICH 1 detector. (b) Cut-away 3D model of the RICH 1 detector, shown attached by its gas-tight seal to the VELO tank. (c) Photo of the RICH1 gas enclosure containing the flat and spherical mirrors.

radiators, providing PID from approximately $1 - 60\,\mathrm{GeV}/c$ for particles within the polar angle acceptance of ±300 mrad (horizontal) and ±250 mrad (vertical). A schematic, 3D model and photo of the RICH 1 detector is shown in Fig. 45. It is aligned to the LHCb coordinate axes and occupies the region $990 < z < 2165$ mm. The z-axis follows the beamline which is inclined at 3.6 mrad to the horizontal. In addition to the PID requirements the overall design has been motivated by the following:

- Material budget: minimizing the material within the particle acceptance of RICH 1 calls for lightweight spherical mirrors with all other components of the optical system located outside the acceptance. The total radiation length of RICH 1, including the radiators, is $\sim$8% X_0.
- Interface to the LHCb beam pipe: the low angle acceptance of RICH 1 is limited by the 25 mrad section of the LHCb beryllium beam pipe (chapter 2) which passes through the detector. The installation of the beam pipe and the provision of access for its bake-out have motivated several features of the RICH 1 design.
- Magnetic shielding of photon detectors: the HPDs of the RICH detectors, described in section 5.1.6, need to be shielded from the fringe field of the LHCb dipole. Local shields of high-permeability alloy are not alone sufficient so large iron shield boxes are used.

The HPDs are located outside of the LHCb acceptance, above and below the beamline, in a region where the iron shield can be accommodated. Additional planar (flat) mirrors are required to reflect the image from the tilted spherical mirrors onto the photon detector planes.

Optical system The parameters of the RICH 1 optical layout have been optimized with the aid of simulation. Charged particle tracks, originating from the interaction point (IP) are followed through the RICH 1 radiators. Cherenkov photons are generated uniformly along the length of each track in the aerogel and gas radiators, using the appropriate refractive indices for photons within the wavelength acceptance of the HPD photocathodes. These photons are ray-traced through the optical system and their impact points on the planes of HPDs are recorded. The Cherenkov angle at emission is then reconstructed for each photon in turn, assuming that the emission point is midway along the track trajectory through the radiator. As the true emission point is randomly distributed along the track, the tilted mirror geometry causes the reconstructed Cherenkov angle to differ from the true value and results in a smearing of the reconstructed angle. The RMS of the resulting distribution is referred to as the emission point error. The RICH 1 optical system is designed such that the emission point error is not larger than other sources of finite angular resolution, such as the HPD pixel size (0.6 mrad) and the chromatic dispersion of the radiator (0.8/1.6 mrad in C_4F_{10}/aerogel).

In addition to the emission point error, the optical layout determines the required area of coverage of the two HPD planes. The optimization procedure required a 100% acceptance for photons emitted by the gas radiator, while a compromise (cost/PID performance) of 68% acceptance has been chosen for the aerogel photons. The optics parameters are constrained by the limited space available and the tilt of the spherical mirrors must ensure that the flat mirrors lie outside the acceptance of the charged particle trackers. The location of the photon detector planes are determined by the performance of the magnetic shield boxes. To avoid loss of efficiency the HPD planes must be in region where the field does not exceed 3 mT and be tilted to ensure that, on average, Cherenkov photons strike the HPDs at close to normal incidence.

The parameters resulting from the optimization procedure have been adopted for the engineering design of RICH 1. The spherical mirrors have radius of curvature 2700 mm with centres of curvature at $x, y, z = 0, \pm 838, -664$ mm, which defines the mirror tilt with respeact to the horizontal plane. This results in an emission point error of 0.67 mrad for the gas radiator and a value that is negligible compared with other sources of error for the aerogel. The flat mirrors are tilted at an angle 0.250 rad with respect to the y-axis, with horizontal edges lying closest to the beam line located at $y, z = \pm 350, 1310$ mm. The two HPD planes are centred at $x, y, z = 0, \pm 1187, 1502$ mm and tilted at an angle 1.091 rad with respect to the y-axis. They are each covered with 7 rows of 14 HPDs, hexagonally close packed, with centres separated by 89.5 mm. This results in a detector plane of 1302 mm$\times$ 555 mm with an active area fraction of 0.55.

Spherical mirrors The spherical mirrors are located within the LHCb acceptance and are traversed by charged particles and photons. Glass mirrors and their associated supports would contribute about 8% X_0, so a carbon fibre reinforced polymer (CFRP) substrate, with the mirror support outside the acceptance, is adopted to keep the material budget below 2%. The two tilted spherical mirror surfaces, one above the beryllium beampipe, the other below, are each comprised of two CFRP mirrors, making four mirror quadrants in total. Each mirror has dimensions 830 mm$\times$ 630 mm when projected onto the $x-y$ plane. The mirror construction is from two CFRP[6] sheets, moulded over a polished glass mandrel to a spherical surface of 2700 mm radius and separated by a reinforcing matrix of CFRP cylinders, configured as shown in Fig. 46a. The box elements at the two outer edges provide stiffness for the three-point mounting brackets that attach the mirror to a CFRP frame (see Fig. 46c). Each mirror has a circular cut-out of radius 62.5 mm to provide a nominal 10 mm clearance from the beam pipe (see Fig. 46b). The overall structure has an areal density of 6 kg m^{-2}, thus contributing about 1.5% X_0 to the material budget of RICH 1.

[6]Fibres: Toray M46J.
Matrix: Bryte Technologies, EX-1515 cyanate ester resin.
Manufacturer: Composite Mirror Applications Inc.,Tucson, USA.

(a)  (b)  (c) 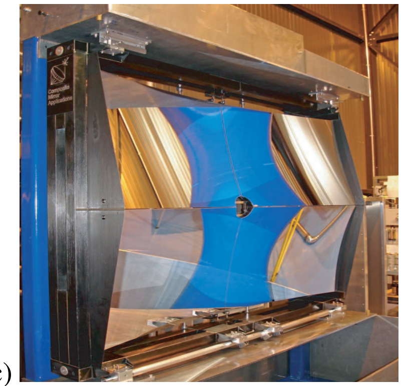

Figure 46. (a) The internal structure of the RICH 1 carbon fibre (CFRP) mirror. (b) Photo of a CFRP production mirror. (c) Photo of the spherical mirror array viewed from the front, mounted onto its CFRP support frame.

The geometrical quality of a spherical mirror is characterized by the variation in the mean radius R of each mirror and the parameter D_0, the diameter of the image (in which is contained 95% of the light) from a point source placed at the centre of curvature. Provided $D_0 < 2.5\,\text{mm}$ and $\Delta R/R < 1\%$, the geometry of the mirrors provides a negligible contribution to the reconstructed Cherenkov angle precision. The manufacturing and subsequent quality assurance ensures that the mirrors satisfy these specifications.

The mirror assembly is made from two CFRP half frames. Each carries an upper and a lower quadrant. The half frames are divided in this way to allow insertion into the RICH 1 gas enclosure (described further on in this section) from either side of the beam pipe, following which the frames are clamped to form a rigid structure around the periphery. The CFRP frame is mounted on rails and supported by an optical alignment rig. Upper and lower mirror pairs are aligned to a common centre of curvature and the CFRP frame is surveyed. This frame is then installed in the gas enclosure using an identical rail system, where the frame can again be surveyed and adjusted if necessary. Simulations [70] have shown that provided all mirror segments are aligned to a precision of about 1 mrad, the alignment can be corrected using data from reconstructed Cherenkov ring images.

The mirrors are coated using a deposition of Al(80nm)+MgF_2(160nm). The reflectivity that can be achieved on the CFRP substrate is the same as that on a glass support, as shown by curve I on Fig. 49 and described in section 5.1.5.

Flat mirrors The flat mirrors are assembled into two planes, one above and one below the beam-line. They are located outside of the acceptance, so glass substrates can be used. Each plane comprises eight rectangular mirrors with dimensions 380 mm (vertical)$\times$ 347.5 mm, fabricated using 6.5 mm-thick Simax[7] glass. As for the spherical mirrors, the flat-mirror geometry is characterized through the parameters R and D_0. All mirrors have $|R| > 600\,\text{m}$ and $D_0 < 2.5\,\text{mm}$. Assuming the deviations from surface sphericity are randomly distributed, this value of D_0 contributes $< 0.2\,\text{mrad}$ to the single photon Cherenkov angle precision. The mirrors are coated using the same

[7] 3.3 borosilicate glass by SKLÁRNY KAVALIER, a.s.
COMPAS, Kinskeho 703, CZ-51101 Turnov.

($Al+SiO_2+HfO_2$) process used for the RICH 2 mirrors (section 5.1.3). Each of the eight mirror segments is connected to a three-point adjustment system via a polycarbonate flange, developed for the RICH 2 mirrors (described in section 5.1.3). The adjustment mechanism is bolted to a rigid plate that is suspended from rails. The eight mirror segments are adjusted in angle to form a single plane, then the angle of this plane is set using the optical alignment rig. Following alignment and survey the plate is mounted on an identical rail system attached to the upstream wall of the RICH 1 gas enclosure, where the angle can again be surveyed and adjusted if necessary.

RICH 1 structure The total weight of the RICH 1 detector is about 16 tonnes, mainly due to the magnetic shielding boxes. The lower box is fixed to the LHCb cavern floor and it supports the gas enclosure and the lower photon detector assembly. The upper box is fixed to the cavern wall and supports the upper photon detector assembly.

Gas enclosure The functions of the gas enclosure are to contain the C_4F_{10} gas radiator and to provide a mechanically stable platform for all optical components. The gas enclosure must be light tight and non-ferromagnetic. It is essentially a six-sided box machined from 30 mm-thick aluminium alloy tooling plate[8] that is welded at the edges to form a 600 kg structure with a total volume of about 3.5 m^3. All six faces have apertures. The boundaries of the upstream and downstream apertures are clear of the RICH 1 acceptance region. The upstream face attaches to a 300 μm-thick stainless steel bellows that provides a gas-tight, mechanically compliant (axial stiffness 37 Nmm^{-1}) seal to the downstream face of the VELO vacuum tank. The downstream face is closed by a low-mass exit window that is sealed to a flange on the beryllium beam pipe using a 1 mm-thick opaque moulded silicone[9] diaphragm. The exit window is manufactured from a sandwich of two 0.5 mm-thick CFRP skins separated by 16 mm-thick polymethacrylimid foam. Its radiation length corresponds to 0.6% X_0. The choice of material thickness is a compromise between material budget and the deflection ($\pm$4 mm) due to the $\pm$300 Pa pressure differential that will be maintained by the C_4F_{10} gas system.

The side faces of the gas enclosure are fully open to maximize access for installation of the optical components and to the beam pipe. They are closed by 10 mm-thick aluminium door panels. With the doors bolted in place the maximum displacement of any part of the gas enclosure due to variations of gas pressure is less than 150 μm. Apertures above and below the beamline are sealed using windows that allow Cherenkov light to reach the HPDs. These windows are 8 mm-thick quartz[10] with dimensions 1360 mm$\times$ 574 mm, fabricated from two equal-size panes, glued together along one edge. The quartz windows are coated with MgF_2 to reduce surface reflection losses from 8% to 4%.

Magnetic shield boxes The HPDs are located in the upstream fringe field of the LHCb dipole. With no magnetic shielding they would experience a B-field of about 60 mT, whereas they can operate at full efficiency using local shielding at B-fields up to a maximum of 3 mT. So the shield boxes need to attenuate the external field by a factor of at least 20, they need to be large enough to accommodate the HPDs and their associated readout, and they must not obstruct Cherenkov

[8] C250, cast using type 5083 alloy.
[9] Dow Corning Sylgard 186, with 5% black pigment added.
[10] HERAEUS Suprasil 2B.

light falling on the HPD photocathodes. In addition, the bending of charged particles in the region between the VELO and the TT station provides a momentum measurement that is crucial for the trigger. Therefore in designing the magnetic shields, due consideration was given to maintaining the field integral in this region.

A schematic of the shielding structure is shown in Fig. 45. It is assembled from 50 and 100 mm-thick high purity ARMCO[11] plates. The magnetic design was optimized using the OPERA/TOSCA[12] finite element modelling software. Measurements made with the shields in place and the LHCb dipole at full field indicate that the maximum B-field at the HPD plane is 2.4 mT, while the field integral between the IP($z = 0$) and the TT($z = 2500$ mm) is 0.12 T·m. Further details of the modelling and measurements are reported in reference [59]. The overall dimensions (x, y, z) of the shield are 1950 mm$\times$ 4000 mm$\times$1175 mm. The weight of each box is about 75 kN and the magnetic forces at full field are about 50 kN. The rigidity of the shielding structure and mounting ensures that any displacement of the HPD assembly is less than 0.5 mm when the LHCb magnet current is switched on.

5.1.3 RICH 2

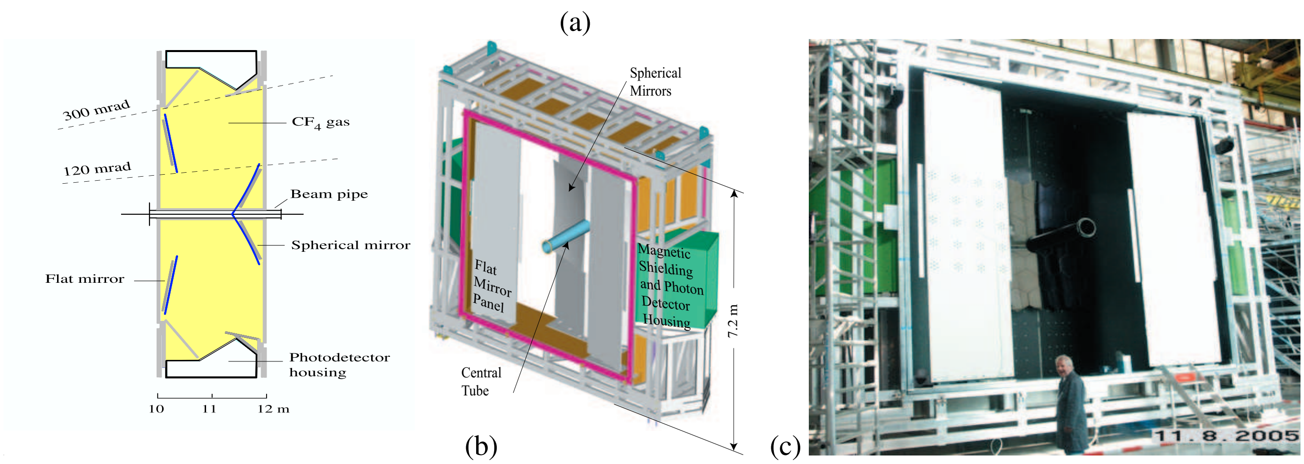

Figure 47. (a) Top view schematic of the RICH 2 detector. (b) A schematic layout of the RICH 2 detector. (c) A photograph of RICH 2 with the entrance window removed.

The RICH 2 detector [60, 61] is located between the last tracking station (chapter 4) and the first muon station (chapter 5). It contains a CF_4 gas radiator, providing PID from approximately 15 to $\geq$ 100 GeV/c for particles within the reduced polar angle acceptance of $\pm$120 mrad (horizontal) and $\pm$100 mrad (vertical). Two schematics and a photograph of the RICH 2 detector is shown in Fig. 47. It is aligned vertically, with its front face positioned at 9500 mm from the interaction point and with a depth of 2332 mm. In addition to the PID requirements the overall design has been motivated by the following:

- Material budget: The supporting structures and the photon detectors need to be placed outside the acceptance of the spectrometer and the HPDs are located left and right of the beam-

[11]ARMCO - Stabilized iron. C$\leq$ 0.01 % S= 0.01 % Mn$\leq$ 0.06 % Si: traces P= 0.01 %.
[12]Vector Fields plc, Oxford, UK.

line ($\pm x$) where the iron shilding is accommodated. To shorten the overall length of the detector, the reflected images from tilted spherical mirrors are reflected again by flat secondary mirrors onto the detector planes. The requirement that the photon detectors are situated outside the full LHCb acceptance defines the lateral dimensions. The total radiation length of RICH 2, including the gas radiator, is ?? X_0.

- Interface to the LHCb beam pipe: The lower angular acceptance of the RICH 2 detector, 25 mrad, is limited by its central tube which runs coaxial to the beam pipe at a distance of 45 mm from it. This distance is required to accommodate the heating jacket and thermal insulation which is required for the bake-out of the vacuum chamber (chapter 2). To gain mechanical stability of RICH 2 and minimize the material in the acceptance of the spectrometer, the detector does not split in two halves along the $x = 0$ plane.
- Magnetic shielding of photon detectors: As for RICH 1, the HPDs need to be shielded from the fringe field of the LHCb dipole and large iron shield boxes are used.

Optical System The final adjustment of the optical layout of RICH 2 has been performed with the aid of simulation, in a similar way to that described in section 5.1.2. This involves defining the position and radius of curvature of the two spherical mirror planes, the position of the two flat mirror planes, and the position of the two photon detector planes. The smearing of the reconstructed Cherenkov angle distribution provides a measure of the quality of the focusing. The RMS of the emission-point error should be small compared to the other contributions to the Cherenkov angle resolution such as the pixelization of the photon detectors and the chromatic dispersion of the radiator. The latter effect is the limiting factor for the resolution in RICH 2, and corresponds to an uncertainty of 0.42 mrad on the Cherenkov angle per photon [60]. The optical elements of RICH 2 must therefore be set such that the emission-point error is small compared to this value.

The parameters resulting from the optimization procedure have been adopted for the engineering design of RICH 2. The spherical mirrors have radius of curvature 8600 mm with centres of curvature at $x, y, z = \pm 3270, 0, 3291$ mm which defines the mirror tilt with respeact to the vertical plane. The flat mirrors are tilted at an angle 0.185 rad with respect to the x-axis, with vertical edges lying closest to the beam line located at $x, z = \pm 1234, 9880$ mm. The two HPD planes are centred at $x, y, z = \pm 3892, 0, 10761$ mm and tilted at an angle 1.065 rad with respect to the x-axis. They are each covered with 9 rows of 16 HPDs, hexagonally close packed, with centres separated by 89.5 mm. This results in a detector plane of 710 mm$\times$ 1477 mm with an active area fraction of 0.56.

Spherical and Flat Mirrors There are two spherical mirror surfaces and two planes of flat mirrors, assembled either side of the beamline. The mirror substrates are made from 6 mm thick Simax glass; the development of these mirrors is described elsewhere [62, 63, 64, 65]. The properties of the production mirrors are shown in Fig. 48. The spherical mirrors are composed of hexagonal mirror elements with a circumscribed diameter of 510 mm, and there are 26 mirrors (or half-mirrors) in each plane. The flat mirror surfaces are formed from 20 rectangular mirror segments in each plane, each 410×380 mm^2 in area. The greatest challenge for the manufacturers was the stability of the thin flat mirror substrates, leading to highly astigmatic or edge deformations. We have therefore chosen to use as flat mirrors, substrates with a finite but large radius of curvature, of about 80 m. The impact on the resolution is discussed in section 5.1.9.

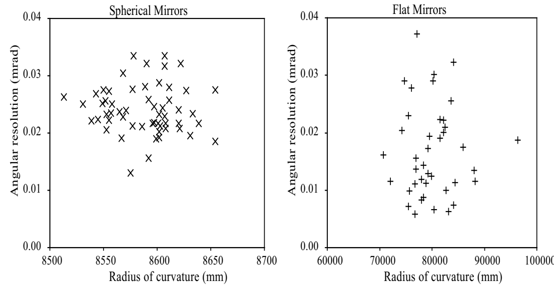

Figure 48. Angular resolution as a function of radius of curvature for the spherical and the flat mirrors in RICH 2. The angular resolution expresses the RMS value of the radius of curvature taken over the mirror surface and is given in radians, under the assumption that the distribution has a Gaussian shape.

Mirror Support and Alignment The mirror supports are the crucial elements that will allow the construction of a near perfect reflective surface from the individual mirror segments; the initial alignments of the mirrors must be better than 1 mrad to have a negligible effect on the Cherenkov ring reconstruction [70]. Each mirror substrate is connected to a three-point adjustment system via polycarbonate flanges and rods [61]. The adjustment mechanism is attached to large aluminium "sandwich" panels which are fastened to the top and the bottom of the superstructure. These panels are made from two 1 mm thick aluminium skins separated by 28 mm aluminium honeycomb[13], giving $\sim 4\%$ X_0. This choice of material is again a trade-off between a long-term stability requirement, reduction of the radiation length and fluorocarbon compatibility. The mirror-support system has been tested in the laboratory for more than one year and, after an initial relaxation period, it is stable to within 100 μrad [63, 65]. The fluctuations are mainly governed by temperature variations.

The mirrors have been installed and aligned in a three-step process [71]. First all the spherical mirror segments have been aligned to within 50 μrad of their common focal point. Then a few flat mirrors were aligned together with the coupled spherical mirrors. The final step was to align the rest of the flat mirrors with respect to these. The total error on the alignment is of the order of 100 μrad. Even though the fully equipped RICH 2 detector was transported by road from the laboratory to the experimental area, no change of the alignment larger than 300 μrad of any mirror has been observed.

RICH 2 Structure. The total weight of the RICH 2 detector is about 30 tons, a large fraction of which, ~ 12 tons, is the overall magnetic shielding structure. The superstructure provides the mechanically stable environment for the optical system, the overall magnetic shielding containing the photon detectors, and a lightweight configuration for the radiator gas enclosure. It is a rectangular box-shaped structure made from welded aluminium alloy rectangular hollow box sections. The deflection of the top of the structure is measured to be < 100 μm under the influence of the magnetic load. The RICH 2 detector is placed and surveyed into position as one unit on the beam line and, after this, the vacuum chamber is then installed.

Gas Enclosure The total volume of the gas enclosure is about 95 m^3, defined by the superstructure, the entrance and the exit windows. The entrance window, constituting a radiation length of

[13] Euro-Composites S.A, Zone Industrielle - B.P. 24, 6401 Echternach.

1.0% X_0, is a low mass composite panel made from two 1 mm thick carbon fibre reinforced epoxy skins, separated by 28 mm thick polymethacrylimide (PMI) foam. The exit window, constituting a radiation length of 2.5% X_0, is similarly made from two 1 mm thick aluminium skins separated by 30 mm PMI foam. The choice of core thicknesses and skin materials is a compromise between radiation length and deflection due to hydrostatic pressure exerted by the Cherenkov gas. The latter will be controlled to within $^{+200}_{-100}$ Pa at the top of the detector. The two windows are clamped and sealed onto the superstructure and are connected to each other by a castellated central tube running coaxial to the vacuum chamber; the tube is made from 2 mm thick carbon fibre epoxy composite. The diameter of the tube at the entrance window is 284.5 mm and 350.5 mm at the exit.

The two planes of photon detectors are separated from the Cherenkov gas by quartz windows on each side. Each window has dimensions of ?? by ?? mm and is made from three quartz panes of 5 mm thickness glued together. The quartz plates have the same antireflective coating as for RICH 1 (section 5.1.2).

Magnetic Shielding Structure The RICH 2 detector is positioned almost halfway between the 1600 ton iron yoke of the dipole and the massive ferromagnetic structure of the hadron calorimeter, shown in Fig. **??**. The maximum stray field in the region where the photon detectors are located is therefore more than 15 mT and is rapidly varying in all directions [69]. The magnetic shield boxes, which are bolted onto both sides of the RICH 2 structure, must provide a mechanically stable and light tight environment for the photon detectors and attenuate the magnetic stray field by a factor of ≥ 15. The shielding, shown in Fig. 47, is a trapezoidal structure made from 60 mm thick ARMCO iron plates. The residual magnetic field measured at the position of the photon detector plane is between 0.2 and 0.6 mT, to be compared to the simulated value of ≤ 0.4 mT.

5.1.4 Radiators

The RICH detectors use three radiator materials for the production of Cherenkov rings to provide particle identification over the momentum range from 1 GeV/c to $\sim$100 GeV/c. In RICH 1 both aerogel and C_4F_{10} gas radiators are used, while RICH 2, which provides for the identification of particles with the highest momenta, CF_4 gas radiator is used.

Aerogel is the ideal radiator to cover the very difficult range of refractive indices between gas and liquid. Silica aerogel is a colloidal form of quartz, solid but with an extremely low density. Its refractive index is tuneable in the range 1.01–1.10, is ideal for the identification of particles with momentum of a few GeV/c. Aerogel has a long-established use in threshold Cherenkov counters. The development of high quality very clear samples [72] has allowed its use in RICH detectors. The dominant cause of photon loss within aerogel is Rayleigh scattering; this leads to the transmission of light with wavelength λ through a block of thickness L being proportional to e^{-CL/λ^4}, where the clarity coefficient C characterizes the transparency of the sample. Large hydrophobic silica aerogel [72] tiles of dimension $20.0 \times 20.0 \times 5.1$ cm^3 have been produced and tested for the LHCb experiment [73]. The refractive index is 1.030 at 400 nm and the clarity is below $C = 0.0054\ \mu\text{m}^4/\text{cm}$. The effect of scattering in the aerogel dominates at high energy, so a thin (0.30 mm) window of glass is placed after the aerogel to absorb the (mostly scattered) photons with $E_\gamma > 3.5$ eV. This also serves to reduce the chromatic aberration.

For a track passing through 5 cm of aerogel with $n = 1.03$, the resulting number of detected photoelectrons in a saturated ring is expected to be approximately 5, calculated by Monte Carlo considering the detailed geometrical setup of the optics of RICH 1 and the wavelength response of the photon detectors. The Monte Carlo was found to describe reasonably well the number of photoelectrons measured in a test beam [73]. Tests have shown [74] that the optical properties will not degrade significantly over the timescale of the LHCb experiment. The aerogel is stable against intense irradiation and shows no significant change in transparency once tested after an accumulated fluence of up to 5.5 x 10^{13}/cm^2 of neutrons or protons, nor after a γ dose of 2.5 $\times$ 10^5 Gy. It is sensitive to water vapor absorption, but its transparency can be restored after a bake-out of the tiles. The volume of aerogel required is modest, $\sim 30\,\ell$, so its replacement, if required, would be relatively straightforward. The refractive index is fairly uniform across a tile, despite the large transverse dimension and thickness. The spread has been measured [75] to be $\sigma_{n-1}/(n-1) \sim 0.76\%$. In RICH 1 the aerogel sits in the C_4F_{10} gas radiator. Tests have shown that C_4F_{10} does not significantly degrade the aerogel performance [77].

The fluorocarbon gases C_4F_{10} (RICH 1) and CF_4 (RICH 2) were chosen because their refractive indices are well matched to the momentum spectrum of particles from *B* decays at LHCb and because they have a low chromatic dispersion. The refractive indices at 0°C and 101325 Pa are parameterized by:
CF_4: (n-1)$\times 10^6$ = 0.12489 /($61.8^{-2} - \lambda^{-2}$)
and
C_4F_{10}: (n-1)$\times 10^6$ = 0.25324 /($73.7^{-2} - \lambda^{-2}$)
where the photon wavelength λ is in nm [76]. For C_4F_{10}, n=1.0014 and for CF_4, n=1.0005 at $\lambda = 400$ nm. The effective radiator lengths are about 95 cm in C_4F_{10} and 180 cm in CF_4. The estimated photo-electron yield is $\sim$16 and $\sim$14 respectively for charged $\beta \approx 1$ particles.

The gas radiators are operated slightly above atmospheric pressure and at ambient temperature. Pressure and temperature are recorded in order to compensate variations in the refractive index (section 5.1.8). These fluorocarbons are transparent well below 150 nm; CO_2 is used as a pressure compensating gas which itself is transparent down to 180 nm. Since the quantum efficiency of the HPDs have iz zero below 190 nm (section 5.1.6), air contamination does not significantly affect the photon yield. O_2 and H_2O contamination are however kept low at about 100 ppm, mainly due to possible radiation-induced formation of HF. The CO_2 fraction is kept constant at $\sim$1 % (section 5.1.8).

5.1.5 Mirror Reflectivity Studies

Several reflectivity coatings are available for RICH mirrors; the choice depends on the Cherenkov photon spectrum, the mean angle of incidence, the long term stability and compatibility to fluorocarbons. Seven different coatings have been tested [62, 66] and the reflectivity measured with a spectrophotometer at an incidence angle of 30°, close to the average angle with which the photons will impinge on the RICH 2 mirrors. The results are shown in Fig. 49.

Simulation studies have been performed by convoluting the reflectivity data with the quantum efficiency of the photon detectors, the Cherenkov photon energy spectra and the transmittance of the CF_4 radiator in RICH 2. The results of the simulation, in terms of number of the relative number of detected photons, are summarised in Table 5.

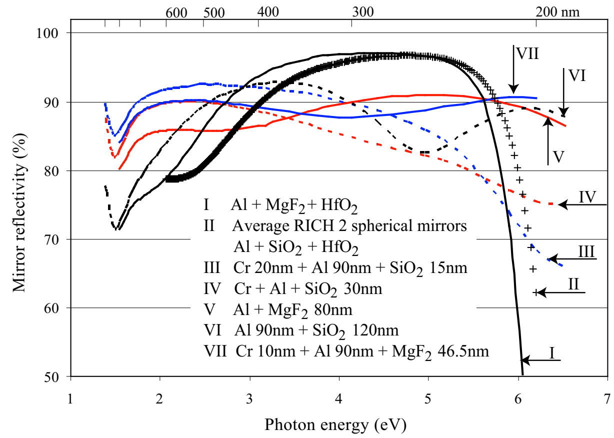

Figure 49. Reflectivity of several mirror coatings on glass as a function of the photon energy. The incidence angle is 30°.

Coating	Photoelectron yield
Cr + Al + SiO_2 (30 nm)	.865
Al + SiO_2	.945
Cr + Al + MgF_2	.947
Al + MgF_2	.960
Cr + Al + SiO_2 (15 nm)	.960
Al + SiO_2 + HfO_2	1.000
Al + MgF_2 + HfO_2	1.000

Table 5. Relative number of detected photons simulated in RICH 2 for different mirror coatings. The values are normalised to the highest photoelectron yield, set to 1.

The reflectivity of the two coatings with a layer of hafnium oxide is very high in the near UV (curve I and II in Fig. 49), where the quantum efficiency of the photon detector is peaked. The simulation shows that good matching, taking into account the two reflections on spherical and flat mirrors, leads to a detected photon yield 5 % higher for this coating compared to other UV extended coatings with magnesium fluoride as the surface layer. Whilst magnesium fluoride coatings have been successfully used in RICH detectors with C_4F_{10} and C_5F_{12} gas radiators in the DELPHI [67] and COMPASS [68] experiments, we have also successfully tested hafnium oxide coatings in C_6F_{14} vapour. Hafnium oxide also provides a very hard and chemically inert protective layer for the mirrors. For these reasons we have chosen the Al + MgF_2 + HfO_2 coating for all RICH 1 and RICH 2 glass mirrors. SiO_2 is used for the middle layer of the multi-layer coating, and not MgF_2, for technical reasons.

5.1.6 Photon Detectors

Pixel Hybrid Photon Detector The RICH detectors utilize Hybrid Photon Detectors (HPDs) to measure the spatial positions of emitted Cherenkov photons. The HPD is a vacuum photon detector in which a photoelectron, released from the conversion in a photocathode of an incident photon, is accelerated by an applied high voltage of typically 10 to 20 kV onto a reverse-biased silicon detector. During the photoelectron energy dissipation process in silicon, electron-hole pairs are created at an average yield of one for every 3.6 eV of deposited energy. Carefully-designed readout electronics and interconnects to the silicon detector result in very high efficiency at detecting single photoelectrons.

A dedicated pixel-HPD has been developed by LHCb, in close collaboration with industry [78]. The specific RICH requirements are a large area coverage ($\sim$3.5 m^2) with high active-to-total area ratio after close-packing (64%), small granularity (2.5$\times$2.5 mm^2 at the photocathode) and high speed (25 ns timing resolution). The pixel-HPD is shown in Figs. 50(a,b). It is based

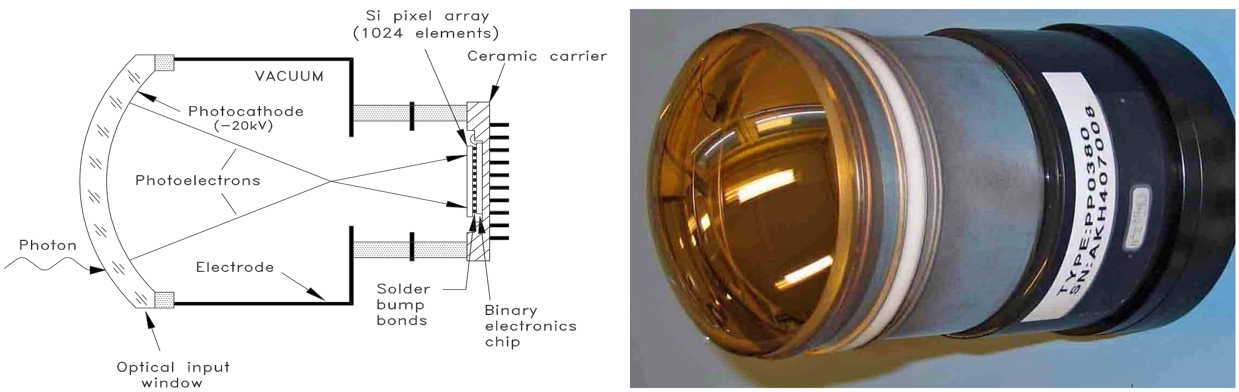

Figure 50. a) A schematic and b) a photograph of the pixel-HPD.

on an electrostatically focussed tube design with a tetrode structure, de-magnifying by a factor of $\sim$5 the photocathode image onto a small silicon detector array. The silicon detector is segmented into 1024 "super" pixels, each 500 μm$\times$500 μm in size and arranged as a matrix of 32 rows and 32 columns. This leads to the required pixel size at the HPD entrance window of 2.5$\times$2.5 mm^2. The nominal operating voltage of the HPD is -20 kV, corresponding to $\sim$5000 electron-hole pairs released in the silicon.

The silicon pixel detector array is bump-bonded to a binary readout chip (section 5.1.7). This flip-chip assembly is mounted and wire-bonded onto a Pin Grid Array (PGA) ceramic carrier, forming the HPD anode. Since all anode parts are encapsulated in vacuum, they must be compatible with the vacuum tube technologies, and must stand high (300°C) bake-out temperatures. In particular, a specific fine-pitch bump-bonding process has been developed for this application. The HPD entrance window is fabricated from quartz and forms a spherical surface, with 7 mm thickness and 55 mm inner radius of curvature. The photocathode is of the "thin-S20" multi-alkali type, deposited on the quartz inner surface. Normally-incident Cherenkov photons can be detected over an active diameter of 75 mm and, since the overall tube diameter is 83 mm, the intrinsic tube active area fraction is $(75/83)^2$=0.817. A total of 484 tubes (196 for RICH1 and 288 for RICH2) are required to cover the four RICH photon detection surfaces.

The de-magnification by 5 of the photo-electron image is achieved by biasing the photocathode at -20 kV, the first electrode at -19.7 kV and the second electrode at 16.4 kV. The values for the point spread function (PSF) standard deviations (at the pixel array) are approximately constant over the tube radius and equal to 80 μm for red light and 180 μm for blue-near UV light, in the absence of magnetic field.

HPD Test Results The HPDs have been fabricated in industry[14] and were then qualified at two test facilities to determine their efficiency and optimum working parameters, before installation at CERN. Each HPD in turn was placed in a light-tight box, illuminated with an LED of wavelength 470 nm, and read out by custom electronics. A selection of test results is presented below. Measurements of the quantum efficiencies (QEs) were carried out after manufacture at Photonis-DEP. Measurements were then repeated at the test facilities for a subsample of 10% of the HPDs, using a calibrated photodiode and a quartz-tungsten halogen lamp. A QE curve for a high-QE tube is shown in Fig. 51. The distribution of peak QE values for all HPDs (484 tubes plus 66 spares) is shown in Fig. 52. The QE curves show an average maximum of 31% at 270 nm, considerably above the specification minimum of 20%.

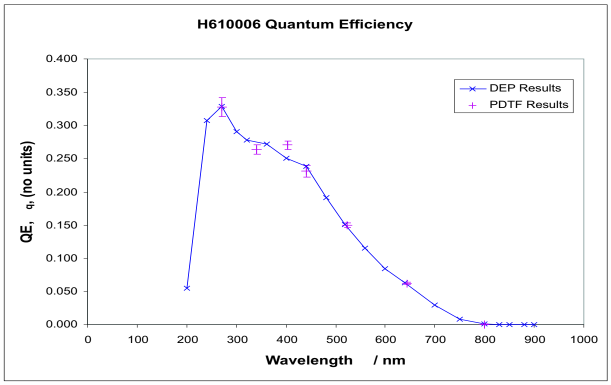

Figure 51. QE measurement for a high-QE HPD.

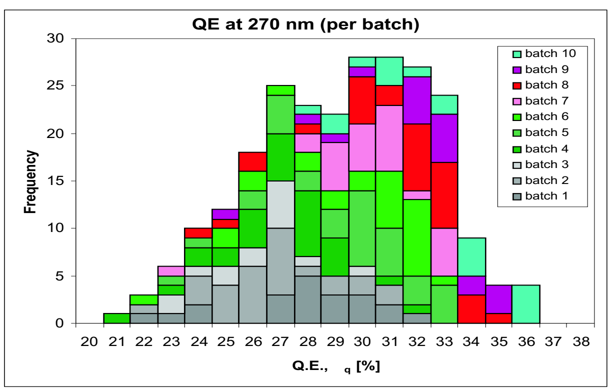

Figure 52. Summary of the peak QE values.

The vacuum quality in the HPD tube is determined by measuring ion feedback. During the acceleration process, photoelectrons may hit residual gas atoms, producing ions. These drift to the photocathode, releasing 10-40 photo-electrons which are detected 200-300 ns after the primary photon signal. Fig. 53 shows that the ion feedback rate is well below the specification maximum of 1%, and indicating that the vacuum quality is excellent.

[14]Photonis-DEP, Roden, Netherlands.

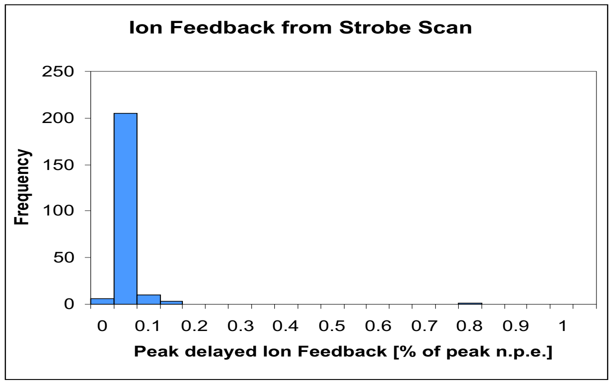

Figure 53. Distribution of the ion-feedback rate.

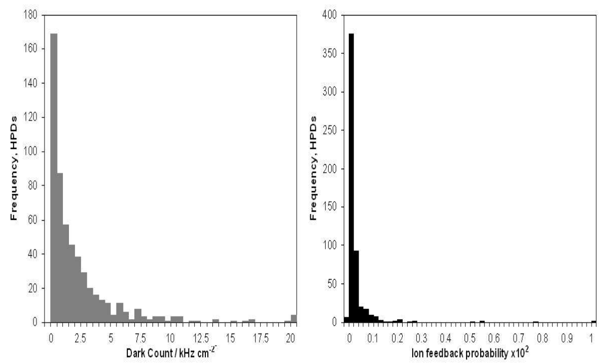

Figure 54. Distribution of the dark-count rates.

An important quality factor of an HPD is low dark-count rate, the main sources being thermionic electron emission at the photocathode, electrostatic field emission and the resulting ion feedback. Fig. 54 shows the measured dark-count rates in the sample of HPDs in a run with no LED. Dark-count rates are typically below the specification value, 5 kHz/cm^2, which in turn is $\sim 10^3$ less than the maximum occupancy expected in the RICH detectors. HPDs with dark-counts greater than the 5 kHz/cm^2 value typically originate from tubes with high quantum efficiency in the red part of the spectrum, causing greater sensitivity to thermionic emission.

HPD testing was carried out at a rate of 30-40 HPDs per month. A total of 97% have been found to meet or exceed the design criteria in key areas.

Photon Detector Integration The HPDs are grouped in four detection planes (two for RICH1 and two for RICH2) and positioned on a hexagonal lattice. The hexagonal close-packing factor is 0.907. Each tube is completely surrounded by a 1 mm-thick cylindrical magnetic shield which protects against stray external B-fields up to 5 mT (the maximum field value within the RICH 1 magnetic shielding has been measured to be 2.4 mT, see section 5.1.2). A tube-to-tube pitch of 89.5 mm in both RICH detectors has been chosen, resulting in a packing factor of $(75/89.5)^2 \times 0.907 = 0.64$.

The HPDs are mounted on columns which are installed in the magnetic shielding boxes of the two RICH detectors. There are 2×7 columns of HPDs in RICH 1 with 14 HPDs per column and 2×9 columns in RICH 2 with 16 HPDs per column. Fig. 55 illustrates the column mounting scheme for RICH 2. The column also contains L0 readout electronics boards (one per pair of HPDs), power supply distribution, cabling and active cooling, all within one mechanical module. The electronics boards and power distribution are described in (section 5.1.7).

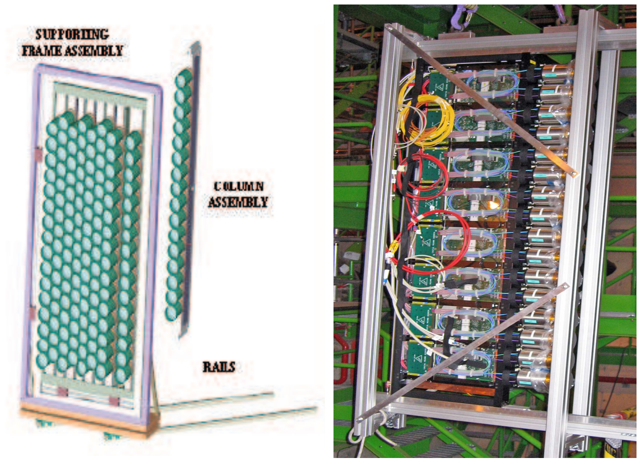

Figure 55. (a) Schematic of the RICH 2 column mounting scheme and b) a photograph of a column.

5.1.7 Electronics

The RICH electronics system reads out data from the HPDs and conforms to the general electronics architecture of LHCb. The electronics system is divided into so-called Level-0 (L0) and Level-1 (L1) regions. The L0 electronics are all located on-detector where they will be exposed to radiation and must therefore contain only radiation-qualified components. The L1 electronics modules are housed in the barracks behind the radiation-shield wall and hence are not required to be radiation tolerant.

Level-0 Electronics The L0 electronics comprises of the pixel chip, ZIFS/kaptons, L0 board and LV/HV distribution, described here below.

Pixel Readout Chip: At the beginning of the electronics readout chain is the LHCBPIX1 pixel chip [79] that forms part of the anode assembly of the HPD (section 5.1.6). The chip has been designed in a commercial 0.25 μm CMOS process using special layout techniques to enhance its resistance to radiation. The chip is connected to a silicon pixel sensor by an array of solder bump-bonds, one per channel. Each pixel measures 62.5 μm by 500 μm and 8192 pixels are arranged as a matrix of 32 columns by 256 rows. Circuitry on the chip logically ORes eight adjacent pixels, thus transforming the matrix into 1024 channels of 32 by 32, each of size 500 μm by 500 μm. Signals from the silicon sensor are amplified, shaped and then compared with a global threshold. Hits are buffered for the duration of the Level-0 trigger latency and events are then stored in a 16-deep multi-event FIFO buffer. The 32 columns are read out in parallel. The analog behaviour of the chip is crucial for a good photo-detection efficiency, in particular low noise and uniform threshold. Measured values are well within the HPD specifications, with typically 160 electrons noise on a signal of 5000 electrons, and a mean threshold of 1200 electrons with RMS spread of 100 electrons.

ZIFs and Kaptons: Signals to and from the HPD pass through the pins of the ceramic pin-grid-array carrier. This is plugged into a 321-pin Zero-Insertion-Force (ZIF) connector mounted on a small circuit board. The board also contains passive components for filtering and line-termination. Signals are then transmitted on two Kapton cables which also carry the low-voltage power to the HPD. These Kaptons allow mechanical flexibility between the mounting of the HPDs and the other Level-0 components.

L0 Board: The L0 board [82] acts as an interface between the HPDs and the ECS(Experiment Control System), TFC (Timing and Fast Control) and data transmission systems. This interfacing is implemented in the Pixel Interface (PINT) antifuse gate array, chosen for its tolerance to single-event effects (SEE). Its main tasks are to receive data from two HPDs, add headers containing event information and data-integrity checks, and transfer events to the data transmission system. A total of 242 L0 boards are used in the RICH detectors, and a photograph of a single board is shown in Fig. 56. The tolerance of the PINT to SEEs was verified by an irradiation program to simulate the expected dose in the LHCb environment. The measured rate of single-event-upsets was negligible, no cases of latch-up were observed and chips survived many times the expected dose of ionising radiation. All other components on the L0 board have been designed in radiation-hardened or tolerant technologies already qualified to levels far in excess of the RICH environment.

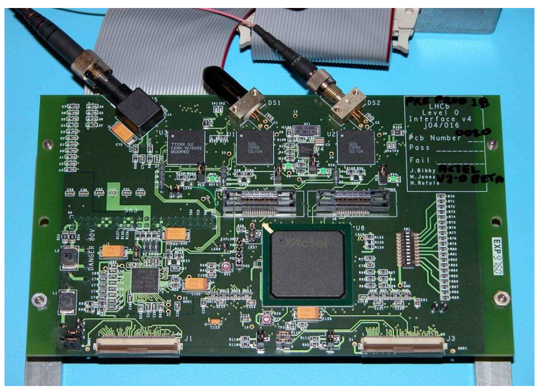

Figure 56. A photograph of the L0 board. The optical receiver and transmitters are visible at the top. The PINT gate array is in the centre of the board.

The optical data transmission consists of Gigabit Optical Link (GOL) chips [80] and Vertical Cavity Surface Emitting Lasers (VCSEL). The GOLs serialise the data into a 1.6 Gbit/s bit stream using 8 b/10 b encoding and the VCSELs convert this into an optical signal of 850 nm wavelength. Two output fibres, one per HPD, transmit the data from each L0 board. At a 1MHz trigger rate, the aggregate output data rate from the RICH photon detectors is in the region of 500 Gbit/s distributed over 484 optical fibres. Also on the board is a TTCrx [81] which generates the 40 MHz clock and trigger and calibration signals. The input to the TTCrx arrives on fibres connected to the global LHCb TFC system. Finally, a chip known as the Analog-Pilot generates reference voltages required by the HPDs and digitises monitoring signals such as the temperature of the board as measured by pt1000 sensors.

LV/HV Distribution: The low voltages required by the HPDs and the L0 boards are provided by the LV distribution system. These voltages are generated locally using radiation-tolerant linear voltage regulators mounted on the LV boards. Each LV board can power two L0 boards independently, and four LV boards are daisy-chained together in an HPD column. The top-most contains a SPECS slave module [83] which provides the interface for the configuration of the HPDs and L0 boards in that column. The SPECS interface is also used for switching on and off the voltage regulators, thus allowing careful control of the powering of the column.

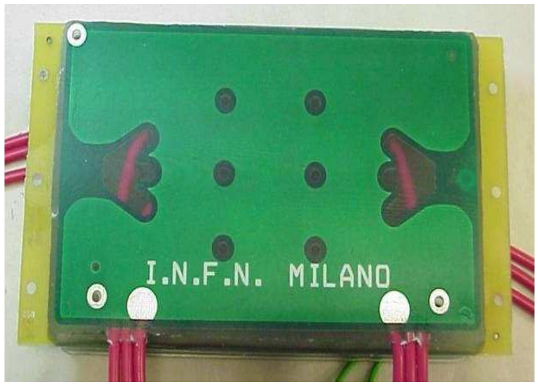

Figure 57. A photograph of an HV board. Around the board is the silicone coating. The cables attach to the HPDs and to the adjacent HV boards.

The high voltages for the HPDs are provided by the HV distribution system [84]. The three voltages are derived from a 20 kV input by means of a resistive divider, one per half-column, and short circuit protection is provided by 1 Gohm resistors in series with each HPD voltage line. Various measures have been taken to minimise the risk of discharge or breakdown. The bare PCB of the HV cards is machined to introduce cuts around critical points which are then filled with dielectric gel to improve insulation. After component mounting and cabling, the entire assembly is encapsulated in silicone rubber. A photograph of a completed board is shown in Fig. 57. A total of 242 boards are used to equip the RICH detectors.

Fig. 58 shows a photograph of all front-end electronics components mounted on a RICH HPD column. On the left are the Kapton cables shown plugged into the L0 boards, then comes a cable tray to carry the optical fibres. Next comes the LV board which has heat sinks mounted on the surface of the voltage regulators. At the far right are the HV boards.

Level-1 Electronics The RICH L1 electronics, located off-detector in the counting room, has been designed to implement data compression and also to serve as the interface between the custom data transmission protocol of the L0 electronics and the industry-standard Gbit Ethernet protocol used by the DAQ network. The off-detector electronics also performs the important function of isolating the DAQ network from errors induced in the L0 data due to radiation induced SEU.

Each incoming serial data channel is first converted from optical to electrical using Agilent HBFR-782 optical receivers. The serial electrical data are then AC coupled to dedicated I/O pins of Virtex2Pro FPGAs where they are deserialised using integrated Gbit transceivers that are com-

Figure 58. A photograph of L0, LV and HV boards mounted on a RICH HPD column.

patible with the data generated by the GOL transmitters in the L0 electronics. All further data processing is done in the FPGA programmable logic.

The L1 electronics modules are controlled by signals broadcast synchronously to all subdetectors by the readout supervisor. These signals are used to control the generation of the data packets sent by the L1 to the DAQ network. The data content of these packets is extracted from the incoming L0 data frames. The generation of the data packets operates autonomously to the arrival of the L0 data and therefore cannot be disrupted by incoming erroneous data. In order to predict the time of arrival of the L0 data frames so that they can be correctly inserted into the generated data packets, the operation of the L0 electronics is emulated in the off-detector modules using the TFC broadcast signals.

Once the incoming data frames have been received, the zero-suppression algorithm is applied in parallel to all streams. The fully pipelined algorithm operates at the 1 MHz input event rate and therefore does not introduce dead-time. Zero-suppressed data are buffered in internal memory in the four ingress FPGAs before being multiplexed into multi-HPD event packets. These packets are then forwarded to the egress FPGA which further multiplexes the data into multi-event packets and transmits them using Gbit Ethernet protocol to the DAQ network using the LHCb quad Gbit Ethernet interface.

The L1 module is configured and monitored via an LHCb Credit Card PC (CCPC) [86] interface. The L1 modules are in 9U format and each can receive data from a maximum of 36 HPDs. The input links are distributed across three 12-channel optical receivers and four Virtex2Pro ingress FPGAs. For 1 MHz operation, the data throughput is expected to be limited by the capacity of the four 1 Gbit/s DAQ links at higher luminosity. Static load balancing is done by physically distributing the input fibre-optic cables across the modules. A total of 26 L1 boards are forseen to read out the data from the RICH detectors, some of whose input channels will be unconnected in order to satisfy the DAQ bandwidth constraint.

5.1.8 Monitoring and Control

The RICH Detector Control System (DCS) [85] monitors the working conditions of both RICH 1 and RICH 2, controls the operating conditions and ensures the integration of RICHes in the LHCb detector control system. It is composed of several parts:

- Power supply control and monitoring (low voltage, silicon detector bias and very high voltage;
- Environment monitoring (temperature, pressure, humidity);
- Gas quality monitoring;
- Magnetic distortion monitoring (and correction);
- Laser alignment monitoring system.

Power Supplies Low voltage (5 V) for the front-end electronic boards is provided by commercial off-the-shelf devices, the Wiener Marathon PL500 power supply. It is radiation and magnetic field tolerant, and can be remotely monitored via a network using standard software interfaces the manufacturer provides (OPC server). Other lower voltages (e.g. 3.3 V) required by the electronics components are generated, regulated and monitored on each front-end LV board. In addition, to improve reliability, temperature sensors placed near hot spots and critical points allow a quick check of the integrity of the electronics. A similar solution has been adopted for biasing the HPD internal silicon pixel detectors, a standard CAEN SY1527 mainframe and plug-in modules power supply has been selected. For the very high voltage of the HPDs (20 kV), no satisfactory commercial solution was found so the system was designed in-house. The system is built around a commercial HV module (ISEG CPn–200 504 24 10) controlled and monitored by a pair of DACs and ADCs connected via an I2C bus interface to a CCPC running the software to regulate the voltage and to connect to the network. The CCPC, using a "bare-bones" version of Linux, runs the control program which, in an endless loop, checks all voltages and currents, and via a DIM protocol [87] reports measurements to the high level control system. In the unlikely case of loss of network connection, the CCPC can still check the working conditions and cut power to the HPDs to avoid damage to personnel and HPDs.

Environment Environment monitoring (temperature, pressure and humidity) is achieved using standard commercial sensors, namely platinum resistors for temperature, diaphragm sensors for pressure and HMX2000-HT sensors (by Hygrometrix) for humidity. All these devices are mounted in the harsh radiation environment hence they must be read by radiation tolerant electronics i.e. the "Embedded Local Monitor Board" (ELMB) [88] developed by the ATLAS-LHC experiment. The ELMB is a general-purpose plug-on I/O module for monitoring and control of subdetector front-end equipment. It is based on the industry-standard CANbus and CANopen and contains 64 analog input ports with programmable gain and with 16 bit resolution plus three byte–wide digital input–output ports.

Gas Quality The purity of the gas radiator is critical to obtain good photon transparency and for this reason a quick analysis tool to spot any contamination is employed. The stability of the gas composition is checked by periodically monitoring the speed of sound of the gas in the vessel. The speed of sound is given by $v_s = \sqrt{\frac{\gamma RT}{M}}$ where $\gamma = \frac{C_p}{C_v}$ is the ratio of specific heats of the gas, R is the universal gas constant, T the absolute temperature and M the average molecular mass. The speed is monitored by measuring the time that a sound pulse takes to propagate from an electrostatic transducer to the end of a cylindrical vessel and then, reflected by the end wall, back to the same transducer, working now as a microphonic sensor. For RICH 1 this vessel is 0.5 m long and the measured time interval is of the order of 10 ms in C_4F_{10}. For RICH 2 the vessel is approximately 0.3 m long and the relative propagation time is of the order of 3.8 ms in CF_4. The whole system is built around a timer/counter provided by a National Instrument acquisition board. The internal counter, running at 20 MHz, provides 50 ns time resolution, more than enough to detect a 1% CO_2 contamination (which will give a 50 μs change in C_4F_{10} and 30 μs change in CF_4 for the total transit time). There are two such systems for each RICH detector, one placed on the inlet pipe and the second on the outlet of the fluid system. A typical measurement is shown in Fig. 59 which shows the change in the propagation time when the percentage of CO_2 changes from 0 to 5%.

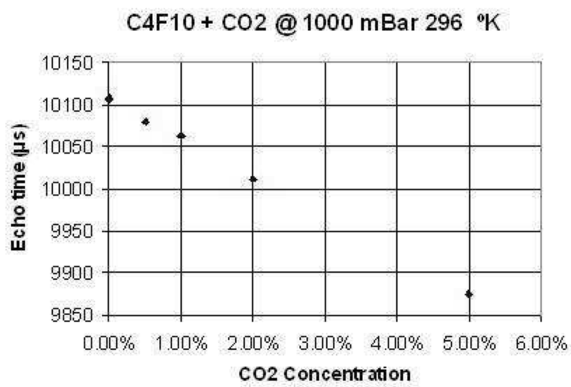

Figure 59. Typical sound propagation time as a function of the CO_2 percentage in C_4F_{10}.

Magnetic Distortions The B–field acts on HPDs modifying the electron trajectories and distorting the image recorded by the pixel chip. Even if HPDs are individually shielded by a μ–metal cylinder and enclosed in an iron box the stray fields are sufficient to affect the required precision. This effect can be clearly seen in Fig. 60 showing a star pattern recorded without B–field and then with a 5 mT field applied perpendicularly to the HPD axis. A correction is needed and considering that data will be taken with both polarities of the B–field to avoid systematics, the effect of the B field on HPDs must be monitored periodically. The system will monitor the distortions projecting (with a commercial device) a fixed dot pattern on the photocathodes and then analysing the data acquired during dedicated calibration runs.

Alignment Mechanical stability of mirrors is of paramount importance to achieve the high angular resolution required to identify particles. This is not easy to achieve due to the size of the mirrors and the difficulty of securing them without introducing mechanical distortions on the reflecting surface. A CCD camera is used to monitor and correct for mirror movements. On its focal plane the image of two light spots is formed: one spot coming directly from a laser source and the second one reflected back from the monitored mirror. Any change in the relative position of the two spots is an indication of a mechanical misalignment and will provide information on how to correct for it off–line. It can be shown that there is a linear transformation law between mirror tilt $(\Delta\Theta_x, \Delta\Theta_y)$

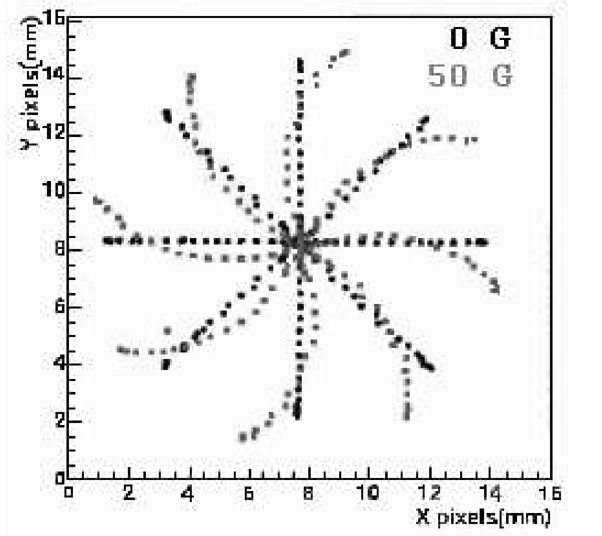

Figure 60. Image of a cross recorded on a HPD without B–field and with a field of 5 mT perpendicular to the HPD axis.

vs. movement of spot on CCD (Δ_x, Δ_y) :

$$\Delta\Theta_x = A\Delta_x + B\Delta_y$$
$$\Delta\Theta_y = C\Delta_x + D\Delta_y$$

The mirror tilts can be recovered by inverting this transformation after observing the spot movement. The system tracks the displacement between the two beam spots, even if the spots move together, with an accuracy of better than 0.01 mrads. The schematic of the system is shown in Fig. 61. The light source is a laser that, thanks to a bundle of optical fibers, generates all the needed light rays (8 mirrors will be monitored in RICH 1, 16 in RICH 2). Rays are then aimed at the beam splitters to obtain the reference beams directed to the cameras and the secondary ones directed to the mirrors.

5.1.9 RICH Performance

Testbeam Performance Regular testbeam campaigns have allowed the RICH photodetectors, readout electronics and radiators to be evaluated with actual Cherenkov radiation. Of particular importance from these campaigns has been the comparison between the expected and observed photoelectron yields and Cherenkov angle resolutions. Because the predictions of these quantities for the testbeam are made using a simulation with identical assumptions to those of the full LHCb Monte Carlo, agreement between prediction and observation is an important indicator that the RICH system will perform to specification.

In 2004 a testbeam was conducted with the first pre-production photon detectors and prototypes of the associated RICH electronics [89]. The tests were performed at the CERN-PS in the T9 facility using 10 GeV/c pions together with an N_2 gas radiator to generate Cherenkov light. The HPDs were arranged in the close-packing arrangement that will be used in the final experiment and were readout at the full LHC speed of 40 MHz. Having included corrections for the asynchronous

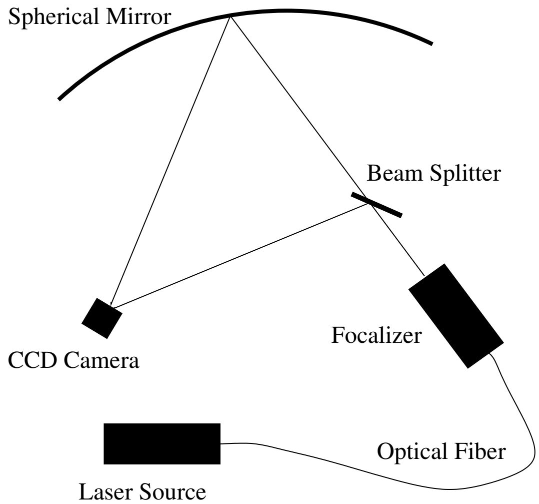

Figure 61. Sketch of the mirror alignment monitoring system.

nature of the test-beam set-up, the photoelectron yields observed were found to be in good agreement with those from the full LHCb Monte Carlo simulation, as shown in table 6. An example distribution of the number of pixel hits observed per event is shown in Fig. 62. Uncertainties in the corrections for the asynchronous test-beam set-up give rise to the residual differences between the data and simulated distributions.

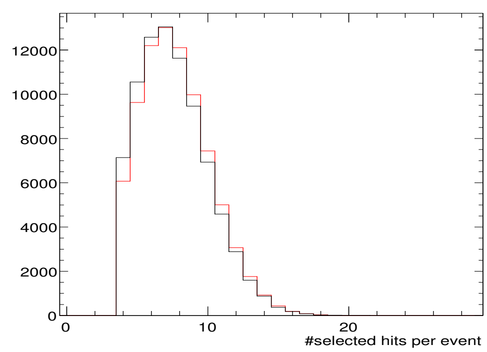

Figure 62. Number of pixel hits per event for data (solid) and simulation (dotted) on one of the HPDs used in the 2004 RICH test-beam.

The distribution of Cherenkov angles reconstructed was also found to be in good agreement with that from the simulation (Fig. 63).

The Cherenkov angle resolution was determined to be 1.66 ± 0.03 mrad, to be compared with the simulation expectation of 1.64 mrad. The dominant contributions to the Cherenkov angle resolution came from the uncertainty in the beam direction, the HPD point-spread function and the pixelisation. The latter two effects were significant owing to the test-beam geometry and are not expected to be dominant in the final LHCb experiment.

Detector Simulation and Offline Reconstruction The LHCb RICH system is modelled as part

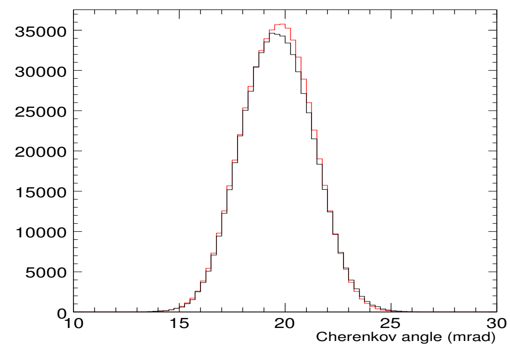

Figure 63. Reconstructed Cherenkov angle for data (solid) and simulation (dotted) on one of the HPDs used in the 2004 RICH test-beam.

HPD	μ_{fit}	μ_{exp}	$\mu_{\text{fit}}/\mu_{\text{exp}}$
L0	8.1 $\pm$0.10	11.4 $\pm$0.73	0.71 $\pm$0.07
L1	10.2 $\pm$0.16	10.0 $\pm$0.66	1.02 $\pm$0.07
C1	11.5 $\pm$0.33	11.2 $\pm$0.78	1.02 $\pm$0.07
R0	8.7 $\pm$0.24	8.9 $\pm$0.73	0.98 $\pm$0.09
R1	10.1 $\pm$0.03	10.7 $\pm$0.83	0.95 $\pm$0.08

Table 6. Comparison of the measured and expected photon yields for the individual HPDs in the 2004 RICH testbeam. The first column indicates the label of each HPD.

of the LHCb Simulation program named GAUSS which is based on the GEANT4 simulation toolkit (chapter 8). All important aspects of the geometry and the material description are fully consistent with what is known about the detectors themselves. A database is setup with all this information which is then input into the simulation and the reconstruction programs.

The Cherenkov generation is performed inside GEANT4, and the photons propagated with full knowledge of the expected reflectivities, transmissions and refraction effects at the various optical boundaries. The Rayleigh scattering in aerogel and absorption in the various media are also included. Inside the HPDs, using the measured quantum efficiencies and cross-focussing relations, the photoelectrons created from the photocathode are mapped down onto the silicon detectors. The location and the energy of the hits made by the photoelectrons are written out from GAUSS. From these hits, the charge sharing in the pixels and the response of front-end readout are modelled in a separate package named BOOLE (chapter 8), which creates the digitized hits in the same format which will be output from the operational detector. All known sources of background are simulated. The most important components come from Rayleigh scattered photons in the aerogel, rings from secondary particles without associated reconstructed tracks and Cherenkov generation in the HPD windows from traversing charged particles. More information on the LHCb RICH simulation may be found in [90]. Fig. 64 shows a simulated event display in RICH 1 displayed with the Panoramix event viewing package (chapter 8).

For a given pixel–track association, the apparent Cherenkov angle is reconstructed through knowledge of the track direction, the hit pixel location, and the geometry of the RICH optics. This reconstruction assumes that the Cherenkov photon was emitted at the midpoint of the radiator.

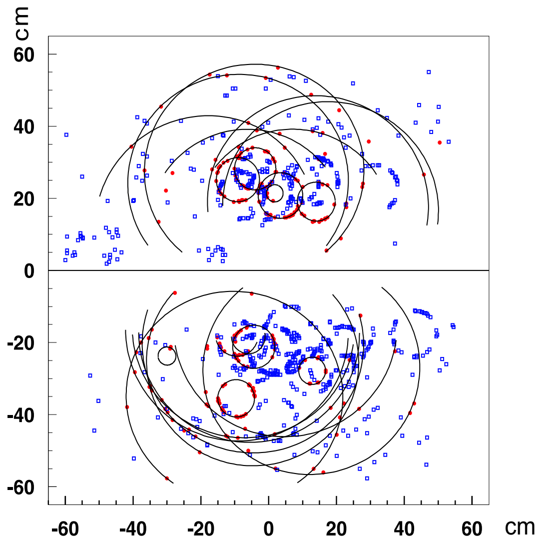

Figure 64. Display of a typical LHCb event in RICH 1.

In simulation the resolution on the Cherenkov angle per photoelectron, and the mean number of photoelectrons detected per $\beta \approx 1$ track can be determined by using truth information to ensure that the hit pixel–track association is correct. The mean number of photoelectrons is found to be $\sim$4 for aerogel, $\sim$16 for C_4F_{10} and $\sim$14 for CF_4. The resolution results are shown in table 7, where both the total resolutions and the individual contributions are listed. The single photoelectron resolution is largest for the aerogel, at 2.6 mrad, and smallest for the CF_4, at 0.74 mrad. These resolutions receive contributions from the uncertainty associated with the photon emission point, the chromatic dispersion of the radiator, the finite pixel size and the point spread function (PSF) arising from imperfections in the HPD optics. For the aerogel it is the chromatic dispersion error which dominates, whereas for the other two radiators the contributions are well matched. An additional uncertainty comes from the reconstruction of the track direction.

	Aerogel	C_4F_{10}	CF_4
Emission			
Chromatic			
Pixel			
PSF			
Track			
Total	2.6	1.59	0.74

Table 7. Single photoelectron resolutions for the three RICH radiators. All numbers in mrad. Individual contributions from each source are given, together with the total. *numbers to be filled in*

The baseline RICH reconstruction is based on a log-likelihood approach which matches the observed pattern of hit pixels in the HPDs to that expected from the reconstructed tracks under a given set of particle hypotheses [91]. In constructing this likelihood it is necessary to calculate the effective emission angle for all pixel–track combination which could physically be associated through Cherenkov radiation. The likelihood is maximised by varying the particle hypotheses of each track in turn, through electron, muon, pion, kaon and proton. As the likelihood considers all found tracks in the event, and all three radiators simultaneously, the method is referred to as the *global pattern-recognition*. A particular strength of the method is the manner in which it accounts for the background to a given ring coming from other rings (at least in the cases where tracking information exists). Background arising from the Cherenkov radiation of charged particles in the HPD windows is suppressed prior to the calculation of the log-likelihood by searching for and eliminating dense clusters of pixels. Background coming from rings emitted by secondary particles which do not have reconstructed tracks is accounted for by comparing the expected and found number of hits in given regions of the photodetectors, and performing a second iteration of the algorithm with thus difference assigned as a background term.

The global pattern-recognition has been optimised in order to minimise its CPU requirements. At present a similar amount of time is spent on the calculation of the pixel–track effective Cherenkov angles, and on the maximising of the likelihood. The mean CPU time per event which passes the LHCb trigger is around 100 ms, as estimated on those processors which will be available in 2007.

The output of the global pattern-recognition is a best hypothesis for each track, and the decrease in log-likelihood when changing from this solution to another hypothesis. The performance of the global pattern recognition can be seen in Fig 65 as a function of momentum between 2 and 100 GeV/c, evaluated on physics analysis quality tracks in events passing the LHCb trigger. Fig. 65(a) shows the probabilities for a pion to be identified as a light particle (best hypothesis = electron, muon or pion) and for a kaon to be misidentified as light. Fig. 65(b) gives the identification probabilities for a kaon to be identified as a heavy particle (best hypothesis = kaon or proton), and the probabilities for a pion to be misidentified as heavy. The performance is good over the entire momentum range. At around 30 GeV/c the identification probabilities are approximately 90 % and the misidentification probabilities less than 10 %.

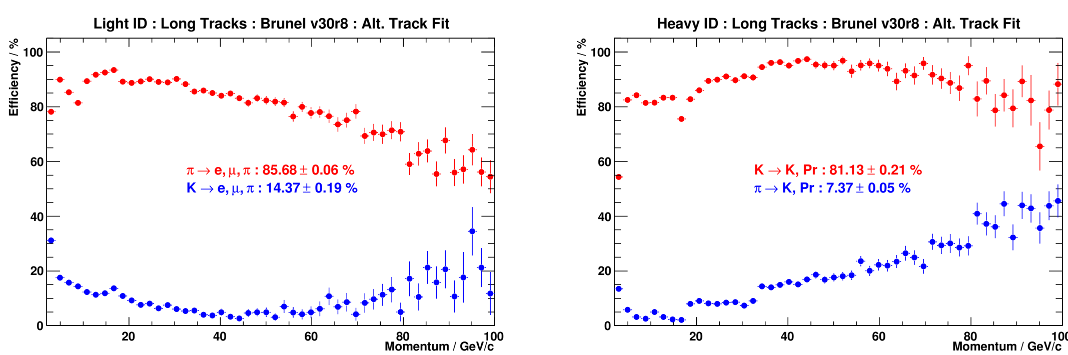

Figure 65. The particle identification performance of the RICH global pattern recognition as a function of momentum. (a) top, shows the identification and misidentification probabilities for the light particle hypothesis. (b) bottom, shows the same for the heavy particle hypothesis.

Other methods have been investigated to find the Cherenkov rings in the absence of tracking information. These include Markov chain Monte Carlo and Hough Transform techniques [92]. Such approaches may be particularly useful in the early stages of data taking, when it may be necessary to study the performance of the RICH detector in isolation from spectrometer information.

For physics analyses and detector diagnostics it will be very important to understand the performance of the RICH particle identification independently of simulation studies. As explained in section **??** the LHCb HLT has a high rate output stream dedicated to $D^{*\pm} \rightarrow D^0(h^+h^-)\pi^\pm$ events which will be selected without the use of any RICH information. The expected sample size on tape from this source is 10^8 events per 2 fb^{-1} of integrated luminosity. The dominant $D^{*\pm} \rightarrow D^0(K^-\pi^-)\pi^\pm$ decay mode provides a very high statistics unbiased sample of pions and kaons which may be used to measure the RICH performance directly. Studies indicate that with kinematic cuts alone purities of $> 95\%$ can be achieved.

RICH in the HLT Access to hadron identification information is highly desirable in the HLT, as this will enable the necessary rate reduction to be achieved without the loss of signal events which would otherwise pass the offline selection. The offline RICH global pattern recognition is an order of magnitude too slow for this purpose. For this reason alternative algorithms have been developed which can be accommodated in the allocated time budget.

Two approaches have been studied: a local algorithm, and an online implementation of the global pattern recognition. In both cases an increase of speed is achieved in the high multiplicity environment of RICH 1 by neglecting the sparsely populated aerogel rings and also by avoiding the full analytic reconstruction of the Cherenkov angle. Instead, the angle is calculated directly on the photodetector plane, after first having applied a local coordinate transformation to correct optical distortions. This method cannot be applied in RICH 2, because of the further difficulties introduced by the non-negligible curvature of the secondary mirrors. The global algorithm is further optimised in speed by only evaluating the pion and kaon hypotheses for each track, and by considering only those which lie within the Cherenkov angle resolution of the rings expected for these possibilities. This runs in a time equivalent to around 3 ms/event, estimated for 2007 processors. The local algorithm operates on a track-by-track basis. It is therefore more selective than the global approach and correspondingly quicker. It relies on histogramming nearby hit pixels in Cherenkov angle space and searching for peaks at the expected positions for each particle hypothesis of interest. Fig. 66 shows the performance of the local and global algorithms for identifying kaons and misidentifying pions as heavy particles. In both cases the 'best hypothesis' assignment has been used, although a 'change in likelihood' output is also available. This latter variable allows the efficiency of the offline particle identification to be retained at the expense of a higher misidentification rate.

Figure 66. *missing* The performance of the RICH online local and global particle identification as a function of momentum, showing the identification efficiencies for kaons and the misidentification efficiencies for pions as heavy particles.

5.2 Calorimeters

5.3 Muon Stations

5.3.1 Physics requirements

Muon triggering and offline muon identification are fundamental requirements of the LHCb experiment. Muons are present in the final states of many CP-sensitive B decays, in particular the two "gold-plated" decays, $B^0_d \rightarrow J/\psi(\mu^+\mu^-)K^0_S$ and $B^0_s \rightarrow J/\psi(\mu^+\mu^-)\phi$ [**?**]. In CP asymmetry and

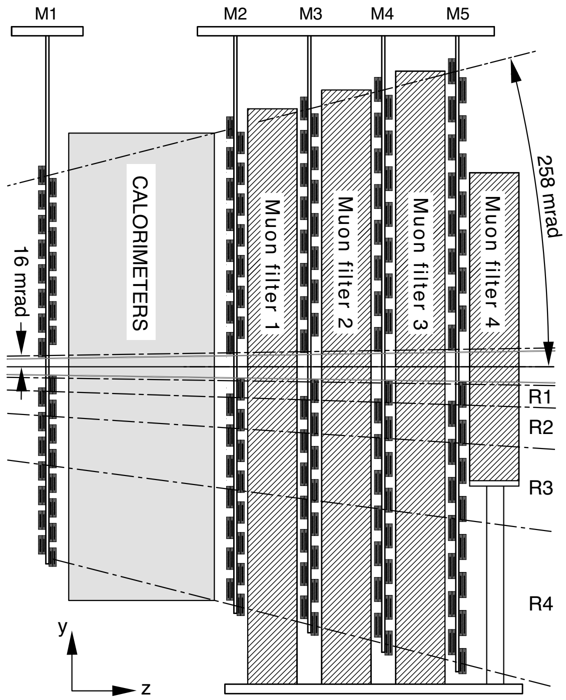

Figure 67. Side view of the muon system in the y,z plane.

oscillation studies, muons from semi-leptonic b decays provide a tag of the initial state flavor of the accompanying neutral B mesons. In addition, the study of rare B decays such as the flavor-changing neutral current decay, $B_s^0 \rightarrow \mu^+\mu^-$, may reveal new physics beyond the Standard Model [129].

The Muon System provides fast information for the high-p_T muon trigger at the earliest level (Level-0) and muon identification for the high-level trigger (HLT) and offline analysis.

5.3.2 Overview

Layout The Muon System [130, 131, 132] is composed of five Stations (M1-M5) of rectangular shape, placed along the beam axis (figure 93). The full system comprises 1392 chambers and has a total area of 435 m^2. The inner and outer angular acceptances of the muon system are 20 (16) mrad and 306 (258) mrad in the bending (non-bending) plane. This results in an acceptance of about 20 % for muons from inclusive b semileptonic decays.

Stations M2 to M5 are placed downstream the calorimeters and are interleaved with iron absorbers 80 cm thick to select penetrating muons. The total absorber (including the calorimeters) is approximately 20 interaction lengths. The minimum momentum of a muon to hit station M5 is about 6 GeV/c. Station M1 is placed in front of the calorimeters and is used to improve the p_T measurement in the trigger. The geometry of the five stations is projective, meaning that all their transverse dimensions scale with the distance from the IP.

The detectors provide space point measurements of the tracks, giving a binary (yes/no) information to the trigger processor [133] and to the DAQ. The information is obtained by partitioning the detector in rectangular *Logical Pads* whose dimensions define the x,y resolution.

The muon trigger is based on a stand-alone muon track reconstruction and p_T measurement and requires aligned hits in all the five stations. Stations M1–M3 have a high spatial resolution along the x coordinate (bending plane) and are used to define the track direction and to calculate the p_T of the candidate muon with a resolution of 20 %. Stations M4 and M5 have a limited spatial resolution, their main purpose being the selection of penetrating particles.

The layout of the muon stations is shown in figure 94. Each Muon Station is divided in four regions, R1 to R4, with increasing distance from the beam axis. The linear dimensions of the regions R1/R2/R3/R4 and of their segmentations scale in the ratio 1/2/4/8. With this geometry the particle flux and channel occupancy is expected to be roughly the same over the four regions of a given station and the (x, y) spatial resolution worsens far from the beam axis, where also multiple scattering increases. Figure 94 (right part) shows schematically the partitioning in Logical Pads and the (x, y) granularity. table 9 gives a detailed information on the geometry of the Muon Stations.

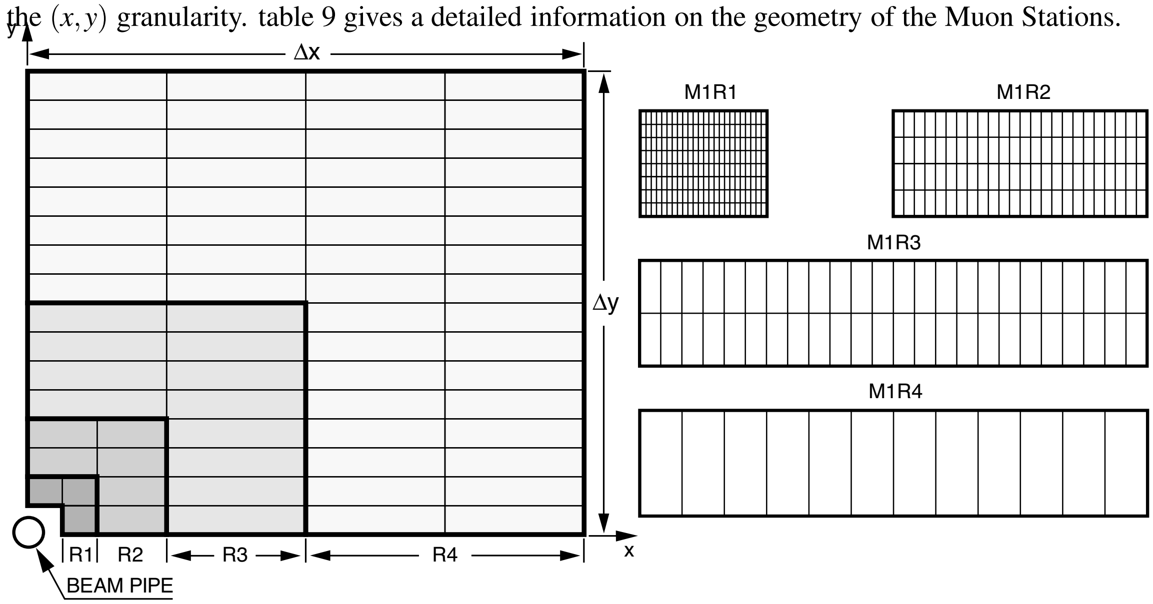

Figure 68. Left: front view of a quadrant of a muon station. Inside each region (R1−R4) the chambers are schematically shown. Right: division in Logical pads of four chambers belonging to the four regions of station M1. In all the regions of station M2-M3 (M4-M5) the number of pad columns per chamber is double (half) compared with the corresponding region of station M1, while the number of pad rows per chamber is the same (see table 9).

Detector technology The LHC bunch crossing rate of 40 MHz and the intense flux of particles in the Muon System [134] impose stringent requirements on the efficiency, time resolution, rate capability and ageing characteristics of the detectors, and on the speed and radiation resistance of the electronics.

Multi wire proportional chambers (MWPC) have been adopted for all regions except the inner part of M1 (M1R1) where the expected particle rate exceeds safety limits for ageing. In this region triple-GEM detectors have been adopted [132, 135, 136].

Since the trigger scheme uses a five-fold coincidence between all the stations, the efficiency of each station must be enough to obtain a trigger efficiency of at least 95 %, within a time window smaller than 25 ns in order to unambiguously identify the bunch crossing. This is ensured by the

Table 8. Basic information for the five stations M1–M5 and the four regions R1–R4. All dimensions in cm. z: distance of the stations from the IP; Δx and Δy: dimensions of a quadrant in each station (see figure 94). Rows R1-R4: granularity of the different regions of the muon detector as seen by trigger and DAQ. Number of Logical Pads per chamber (in brackets) and size of the Logical Pads, along x and y. In parentheses: size of the Logical pads projected onto station M1.

	M1	M2	M3	M4	M5
z	1210	1527	1647	1767	1887
Δx	384	480	518	556	594
Δy	320	400	432	464	495
R1	$[24\times 8]$ 1×2.5	$[48\times 8]$ 0.63×3.1 (0.5×2.5)	$[48\times 8]$ 0.67×3.4 (0.5×2.5)	$[12\times 8]$ 2.9×3.6 (2×2.5)	$[12\times 8]$ 3.1×3.9 (2×2.5)
R2	$[24\times 4]$ 2×5	$[48\times 4]$ 1.25×6.3 (1×5)	$[48\times 4]$ 1.35×6.8 (1×5)	$[12\times 4]$ 5.8×7.3 (4×5)	$[12\times 4]$ 6.2×7.7 (4×5)
R3	$[24\times 2]$ 4×10	$[48\times 2]$ 2.5×12.5 (2×10)	$[48\times 2]$ 2.7×13.5 (2×10)	$[12\times 2]$ 11.6×14.5 (8×10)	$[12\times 2]$ 12.4×15.5 (8×10)
R4	$[12\times 1]$ 8×20	$[24\times 1]$ 5×25 (4×20)	$[24\times 1]$ 5.4×27 (4×20)	$[6\times 1]$ 23.1×29 (16×20)	$[6\times 1]$ 24.8×30.9 (16×20)

use of a fast gas mixture and of an optimized charge-collection geometry both for the MWPC and GEM detectors. Moreover, the chambers are composed of two or four OR-ed gas gaps. In stations M2–M5 the MWPCs are composed of four gas gaps arranged in two sensitive layers with independent readout. In station M1 two gas gaps (two GEM chambers in R1) are used to minimize the material in front of the electromagnetic calorimeter.

In addition to the improved time resolution, the use of two layers with independent HV supplies and the flexibility of the readout provides a high degree of redundancy built into the system.

Readout To satisfy the requirements on spatial resolution and rate capability that are strongly varying on the detector, different technical solutions have been adopted for the MWPC in different stations and regions.

All the chambers are segmented into *Physical Pads*: anode wire pads or cathode pads in the MWPC and anode pads in the GEM chambers. Each Physical Pad is read out by one Front-End (FE) electronics channel. The FE electronics is essentially the same for all detectors and is based on custom radiation-hard chips specially developed for the Muon System. The input stage can be wired to handle both signal polarities: negative for anode pads, positive for cathode pads.

The electronics is provided with flexible logical units which can perform the OR (the AND in special cases) of a variable number of FE Electronics output channels following the requirements of the readout.

The readout methods used in different detector regions are summarized below:

Readout type	Region
MWPC	
Wire pads	R4
Mixed wire-cathode pads	R1-R2 in M2-M3
Cathode pads	everywhere else
GEM	
Anode pads	M1R1

In the regions R4, which comprise most of the chambers, the simplest and safest technology has been adopted: the Physical Pads are a group of adjacent wires connected together to the same FE channel. The length of the anode wires provides the y resolution, all the wires being aligned vertically. The requirement on y resolution then limits the vertical size of the chambers to $20-30$ cm, and results in a large number of chambers to equip the entire system.

Cathode pads (anode pads for GEM) are obtained by etching the desired pattern in the cathode (anode) planes. The pad structure of the chambers of station M1 is shown in figure 94 (right). All the pads in R3 chambers can be accessed from the borders. The chambers belonging to R1 and R2 have a chessboard pad structure so that a multilayer printed circuit board is used to bring the signals outside.

To keep the noise and the dead-time of the FE stage at an acceptable level, the pad electrical capacitance and the channel rate must be limited. This implies that in part of the detector the Physical Pads (wire pads or cathode pads) must be smaller than requested by the spatial resolution. In these cases up to four adjacent Physical Pads are OR-ed by the FE electronics to build-up a Logical Pad.

In regions R1-R2 of stations M2-M3 mixed readout has been adopted because the required spatial resolution in x imposes Logical Pads which would be too small to be practically built. In this case a narrow wire-strip defining the x resolution and a larger cathode pad defining the y resolution are the logical channels seen by the trigger and DAQ. Logical Pads are then obtained as an AND between wires and cathode pads (figure 95).

In M1 station, where the foreseen channel occupancy is high, the signals from the Logical Pads are directly sent to the trigger and DAQ. In most of the other regions, in order to reduce the number of output lines, several contiguous Logical Pads are further OR-ed to build larger *Logical Channels* in form of vertical and horizontal strips, whose signals are sent to the trigger and DAQ. The Logical Pads are then reconstructed by the coincidence of two crossing strips. Figure 96 shows the *Sectors* containing the crossing strips. The Sector size is adapted to the trigger processing elements that work on a fixed number of Logical Pads belonging to a projective tower over the five stations.

The full Muon System comprises ~ 122000 physical channels and ~ 26000 transmitted Logical Channels providing the ~ 55000 Logical Pads used for the muon tracking.

The specifications of MWPCs and GEMs and their performance are summarised in the following sections.

5.3.3 Wire chambers

Design The LHCb muon system comprises 1368 Multi-Wire Proportional Chambers. Prototype studies [137, 138, 139, 140, 141] showed that a time resolution of about 5 ns can be achieved in

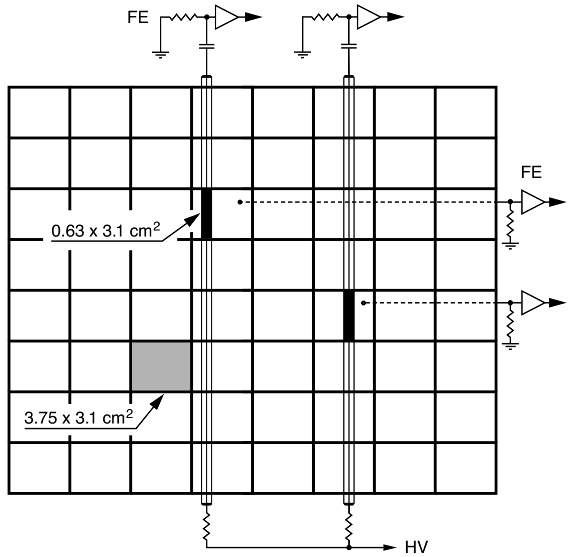

Figure 69. Scheme of the mixed wire-cathode pads readout in one M2R1 chamber. Two wire-pad and two cathode-pad readout channels are shown. The coincidence between crossing vertical wire-pads and the cathode Phyical Pads defines the Logical Pads (shown in black)

.

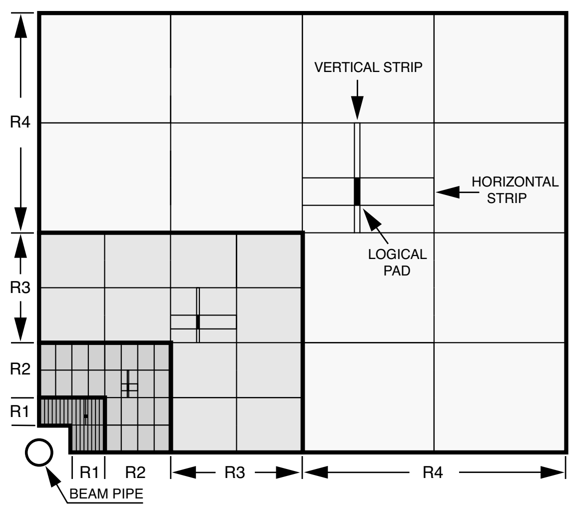

Figure 70. Front view of one quadrant of stations M2–M3 showing the partitioning into Sectors. In one sector of each region a horizontal and a vertical strip are shown. The intersection of a horizontal and a vertical strip defines a Logical Pad (see text). A Sector of region R1 (R2, R3, R4) contains 8 (4, 4, 4) horizontal strips and 6 (12, 24, 24) vertical strips.

a chamber with a wire plane of 2 mm spacing symmetrically placed in a 5 mm gas gap and by using fast, non-flammable gas mixtures of $Ar/CO_2/CF_4$ with 40% Ar and variable concentrations of CO_2/CF_4. Finally the mixture $Ar/CO_2/CF_4(40:55:5)$ was adopted. This elementary detector (*gas gap*) is the building block of the system. By OR-ing the signals from two adjacent gas gaps

Table 9. Main MWPC parameters

Parameter	Design value
No. of gaps	4 (2 in M1)
Gas gap	5 mm
Anode-cathode spacing	2.5 mm
Wire	Gold-plated Tungsten 30 μm dia.
Wire spacing	2.0 mm
Wire length	250 to 310 mm
Wire mech. tension	0.7 N
Total no. of wires	$\approx$ 3 Millions
Operating voltage	2.5-2.8 kV
Gas mixture	Ar / CO_2 / CF_4 (40:55:5)
Primary ionisation	$\simeq 100\,e^-/cm$
Gas Gain	$\simeq 10^5$ max
Gain uniformity	$\pm$30% typ
Charge/mip	$\simeq$ 0.8 pC @ 2.7 kV

the resulting *Double Gap* has an efficiency better than 95 % in a 20 ns window at a gas gain of $G < 10^5$. This gain is achieved [142] at a voltage of 2600-2700 V. Space charge effects due to the accumulation of ions are not expected for rates of up to $xx\,\mathrm{MHz/cm^2}$, which are much higher than the largest rates to be encountered in the Muon System.

The main parameters of the MWPC detectors are summarized in table 10. Detailed simulations [143] based on GARFIELD [144] have been used to optimize the design and to establish the required geometrical tolerances for the chamber construction. In the simulations a maximal tolerable variation in gas gain of $\pm$20% has been assumed, and the corresponding maximal deviation of the single construction parameters was established. The primary ionization has been simulated with HEED [145] and MAGBOLTZ [146] has been used to describe drift and diffusion.

PSpice has been used to optimize the cathode pad readout, in order to keep the cross-talk between pads below 5 %. Figure 97 shows the comparison of a Monte Carlo simulation for efficiency and time resolution as a function of the threshold with the values measured on a chamber prototype.

A chamber is made of four (two in station M1) equal gas gaps superimposed and rigidly stacked together, each constituting an elementary detector. The gas flows serially through all the gaps. As already mentioned, to ensure the required efficiency and redundancy, the chambers are composed by two detector layers with independent readout. In the case of four gap chambers two contiguous gas gaps are hard-wired together in OR to form a Double Gap layer. The readout electrodes of this Double Gap (either wires or cathode pads) are connected to separate FE channels and then OR-ed by the FE electronics. In the M1 chambers the two layers corresponds to the two gas gaps, and are readout in the same way. To have maximum operation flexibility each gap has its own HV line and can be powered or switched off independently of the others.

The structural elements of the chamber are the panels, each formed of an insulating core sandwiched between two conducting planes. The conducting planes inside a chamber form the

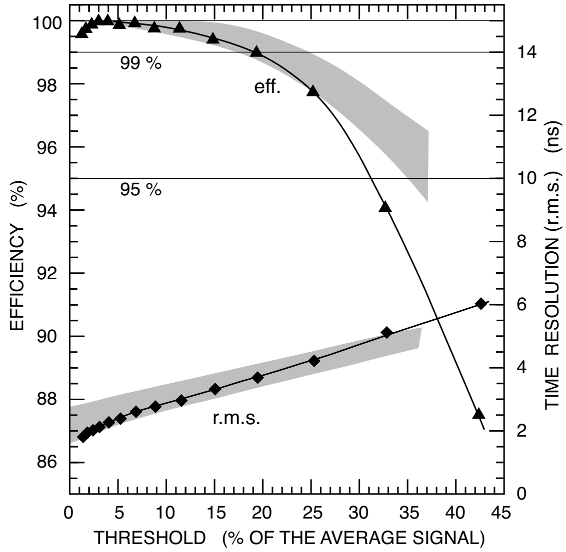

Figure 71. Efficiency and time resolution vs. discriminator threshold. Solid lines: experimental results. The Monte Carlo simulation is shown by the gray band.

cathodes while the two outer planes are grounded and act as an electrical shielding. The panels are stacked with two frames of 2.5 mm which support the wires (figure 98) while the 5 mm gas gap is ensured by several precise spacers placed on the panels (figure 99). An U-shaped brass

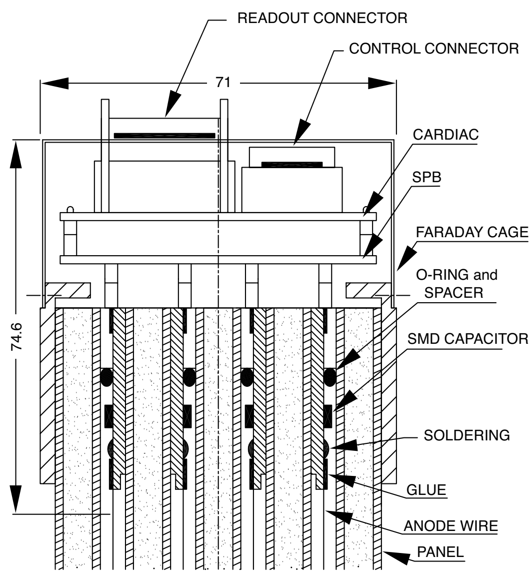

Figure 72. Cross section of a wire chamber showing the four gas gaps and the connection to the readout electronics. SPB: Spark Protection Board; CARDIAC: FE Electronics Board. In this case the hardwired OR forming the two Double Gaps (see text) is achieved in the SPB.

channel running around the chamber edges is soldered to the outer conductive planes to complete the chamber Faraday cage. The front-end boards, the LV voltage regulators and the HV filters are mounted inside the Faraday cage to minimize electrical pickup. The HV is brought in through a custom made multipin connector and multiconductor cable. LVDS shielded cables are used for signal transmission and control.

The general design and construction is the same for all chambers and is discussed in detail in Ref. [147].

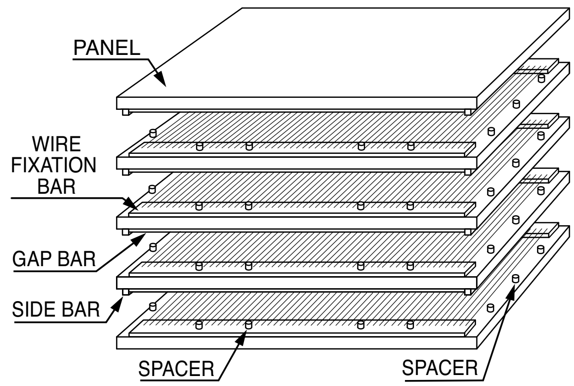

Figure 73. Exploded schematic view of a chamber showing the position of the various bars and spacers.

Chamber construction Given the large number of chambers, the production was distributed in several production sites. These were located in CERN, Ferrara, Firenze, LNF and PNPI. A great effort went in ensuring that all those sites had equivalent facilities and tooling, albeit with some flexibility. However, the same stringent quality criteria and test protocols have been adopted throughout to ensure a constant quality of the produced detectors.

Panels The panels are the basis of the chamber mechanical structure. A panel consists of two copper clad FR4 (Fire-resistant fibreglass epoxy) copper laminates, interleaved with a structural core. The copper foils constitute the chamber cathodes, which can be flat or etched with readout pads (see figure 100). The panels are individually wired and then are assembled to form the complete chamber.

The panel precision (thickness, planarity) defines the gap quality, therefore it is very important to achieve tight tolerances and stability in time, coupled with adequate robustness since they form the chamber's structure. The panels must be light and stiff and easily adapted to mass production. We have chosen a polyurethane foam core for stations M2–M5 and a light honeycomb core for station M1 where a lower material budget is mandatory.

The panels have been produced industrially. For M2–M5 Precision machined molds have been used to inject the polyurethane foam into the FR4 sandwich. Three molds of different sizes were used to produce all the panels. For M1 chambers, where a lighter construction is necessary, the core was made from a light NOMEX$^{(R)}$ sheet purchased with precise tolerance on thickness, and glued between the two FR4 foils. The pressure during gluing was assured by a vacuum bag. So the planarity of the panels was due to the quality of the honeycomb core.

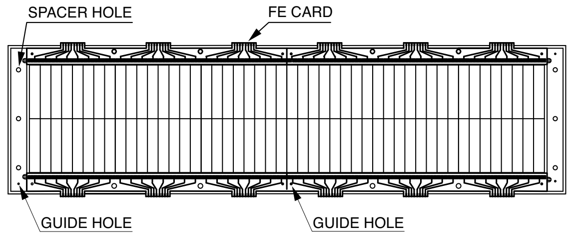

Figure 74. View of the 2×48 cathode pads of a chamber belonging to region R3. The pads are etched on the copper foil using the printed circuit technique.

All the panels were individully measured at production and then shipped to the production sites where they were checked once more before assembly. Despite the tight tolerances the yield of good panels was 90 %.

Wiring Since there are more than 3 Million wires in the system, specially built automatic winding machines have been used in all the production sites to perform panel wiring. The panel were fixed to a rigid frame, and grooved dowels (combs) were installed parallel to the panels long sides. To achieve the required precision on wire pitch, the wire spacing is determined by the combs while the wire height with respect to the cathode-plane is adjusted by precision bars mounted to the frame. On one side the bars are fixed, and on the other, they can be adjusted depending on the panel thickness to achieve a wire to cathode plane distance of 2.5 mm.

Figure 75. One of the automated winding machines ready to start panel wiring.

Once a frame was wired, it was be removed from the winding machine to have the wires glued and soldered. This procedure separated the three important steps of the chamber construction and allowed a parallel production.

The wires were glued to the wire fixation bars using epoxy adhesive (24 h polymerization time) before soldering. This procedure guarantees that the wires are kept in place with a fixed height with respect to the cathode plane. The gluing also keeps the wire tension to its nominal value.

Due to the large number of soldering points in the construction of the MWPCs ($\simeq$ 6 millions), the use of an automated and reliable method is desirable. Automatic soldering stations [147] both with conventional soldering heads and with laser heads have been employed. These stations performed well but in some cases also hand soldering has been used.

Final assembly In the final assembly of the chamber the panels were superimposed and glued together using cylindrical precision spacers for alignment (see figure 99). The glue ensures the chamber gas tightness. In the chambers produced at CERN no glue was used and gas tightness was ensured by O-rings seating in grooves machined in the panels. This allows to open the chambers easily if needed. Screws were used to keep together the panels. Then the brass channels forming the Faraday cage were soldered to the outer copper cladding of the panels.

Quality control and testing In order to have a uniform quality of the detectors produced in the various sites, stringent quality tests of the individual chamber components and of the assembled chamber have been carried out. All the panels were individully measured at production and then shipped to the production sites where they were checked once more before assembly. Despite the tight tolerances the yield of good panels was 90 %. On the production sites two automated measuring tables were used to measure the wire pitch and the wire mechanical tension. For measuring the wire pitch a system based on two digital cameras has been used [148]. In figure 102 the distribution of the pitch, measured on about 500,000 wires, is reported.

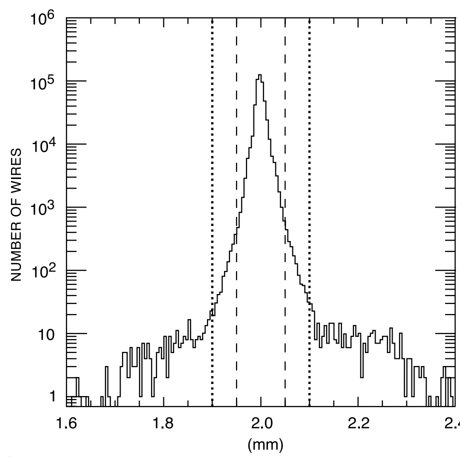

Figure 76. Distribution of wire pitch. Average value is 2.000 mm with an r.m.s. of 0.016 mm. The dashed (dotted) lines defines the interval which must comprise 95 % (100 %) of all the wires. Pitch values outside the dotted lines region have been visually checked. **figure to be replaced**

To determine the wire mechanical tension automated systems have been developed [149, 150] which measure the mechanical resonant frequency of each wire. In figure 103 the distribution of the wire tension measured on the abovementioned sample is reported.

The checks on the assembled chambers consisted in a gas leak test and a HV test at 2800 V using the standard gas mixture. The chamber gas leakage is measured by monitoring the decrease in time of an overpressure of 5 mbar. The method [148] has a sensitivity of about 0.01 mbar/hour well below the maximum gas leakage rate accepted (2 mbar/hour).

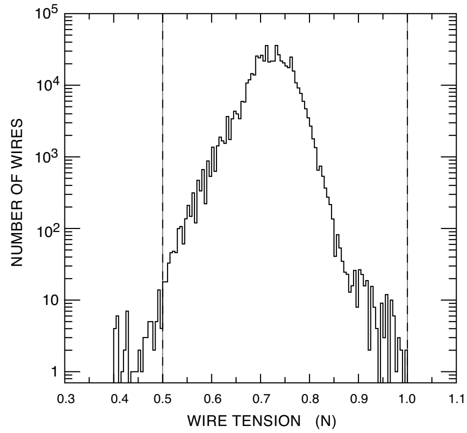

Figure 77. Distribution of wire mechanical tensions as deduced from a measure of the wire mechanical resonant frequency. The mean value of the tension is 0.73 N with an r.m.s. of 0.091 N. The wires having a tension outside the interval $0.5-1$ N (dashed lines) were replaced.

The uniformity of the gas gain in the gaps was systematically checked using radioactive sources mounted on automated tables which performed an (x,y) scan over the complete chamber surface (figure 104).

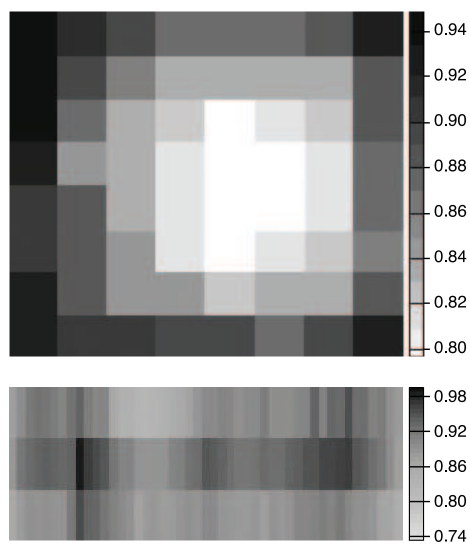

Figure 78. Measurement of gain uniformity of a gap of a chamber of M3R1 (top) and of M5R4 (bottom) performed with a movable radioactive source.

Chambers with local gain deviations of at least one Double Gap exceeding a factor $(1.7)^{\pm 1}$ with respect to the average were rejected.

All the above checks were performed without the readout electronics installed.

Finally all chambers, fully equipped with the front-end electronics, underwent a final test with cosmic rays prior to installation. The electronics noise was again checked after the chambers were mounted on the detector. The maximum accepted noise rate was 1 kHz, a value which has no measurable impact on the trigger system. Important information about individual components and the final chamber were stored in a database. This allows to retrieve at any time the results of all quality control measurements and will aid in understanding possible problems.

Aging properties An extensive aging test was performed on prototypes. The goal of the test was to prove that the performances of these detectors are not deteriorated by the large radiation dose expected in the experiment in several years of operation (10^8 s) at the nominal luminosity of $\mathcal{L} = 2 \times 10^{32}$ cm^{-2} s^{-1} .

The test was carried out with the ^{60}Co source of the Calliope facility (ENEA-Casaccia Research Center, Rome) [141, 151, 152]. During 30 days we monitored the currents in three four-gap MWPCs: two test chambers and a reference one. The tested chambers were operated at a voltage corresponding to a gas gain $\sim 8 \times 10^4$, about twice the value which will be adopted during the experiment and were exposed to doses in the ratio 1:4. The gain variation were measured by taking ratios of the currents with the reference chamber.

The largest integrated charge measured on one gap was 0.44 C/cm, corresponding to $\sim$ 4.6 years equivalent for region M1R2 and $\sim$ 8 years for region M2R1. (**these figures have to be checked against most recent rate evaluations**) No gain loss or other significant effect was detected.

5.3.4 GEM Chambers

In the innermost region R1 of the station M1, covering a total area of 0.6 m^2, the most stringent requirements for the detector is a rate capability of up to 500kHz/cm^2 of charged particles (with an average rate of 184kHz/cm^2). Due to the large particle flux this region has to be especially radiation hard such that no ageing effects will be visible in $\sim$ 10 years of LHCb operation. This is difficult to achieve in wire chambers unless the gain is decreased. The requirement on the cluster size ≤ 1.2 also was found to reduce the useful HV plateau length with MWPCs.

After an extensive R&D program triple-GEM detectors with active area of 20×24 cm^2 were selected for M1R1. Each of the 12 chambers consists of two triple-GEM detectors superimposed and forming the two sensitive layers, which are then digitally OR-ed.

Design The triple-GEM detector, which consists of three gas electron multiplier (GEM) [153] foils sandwiched between anode and cathode planes, can effectively be used as tracking detector, with good time and position resolution performances. A cross-section of the detector, together with the labeling of the different detector elements and the physical dimensions, is shown in figure 105.

The ionization electrons, produced in the gap between the cathode and the first GEM foil (drift gap) by the charged particles crossing the GEM, are attracted by electric fields through the three GEM foils where they get multiplied. Once they cross the last GEM foil they drift to the anode in the so called induction gap, giving rise to an induced current signal on the pads. With leading edge triggering, the discriminator crossing on the signal rising edge gives the time of the event.

Prototype tests have shown that the new and fast $Ar/CO_2/CF_4(45:15:40)$ gas mixture allowed to achieve a time resolution better than 3 ns (r.m.s.), to be compared with the time resolution of $\sim$ 10 ns obtained with the standard Ar/CO_2 (70:30) gas mixture[154].

Another improvement on the time performance has been obtained optimizing the detector geometry (figure 105). Mechanical considerations indicate that a minimum distance between GEM

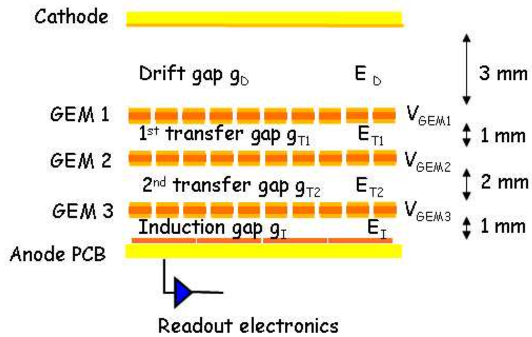

Figure 79. Schematic cross-section of a triple-GEM chamber showing the most relevant elements and dimensions.

foils of about 1 mm should be kept. The drift gap g_d size is large enough to guarantee full efficiency on charged tracks. The first transfer gap g_{t1} is kept as small as possible to avoid that primary electrons produced in the same gap give rise to a signal over threshold. The second transfer gap g_{t2} is larger than the first one to let the diffusion spread the charge over more holes and then lower the discharge probability. The induction gap g_i is small as possible to maximize the signal fraction integrated by the amplifier.

Measurements on prototype chambers have shown an efficiency in a 20 ns time window better than 96% with a gain of 6000. The pad cluster size measurement and the discharge studies suggested a maximum working gain of about 20000[136].

The exact values of the fields were found experimentally by optimizing time resolution versus discharge probability and are E_d=3.5 kV/cm, E_t=3.5 kV/cm and E_i=5 kV/cm. With these gap and field values at the entrance of the induction gap the diffusion gives a σ of 40 μm for the longitudinal position and the transverse of 70 μm.

In order to limit the damage in case of discharge, one side of the GEM foil is divided into six sectors, of about 33 mm$\times$240 mm. The separation between sectors is 200 μm. The pad printed circuit boards (PCB) is such that the pad to pad distance is 0.6 mm and the pads are interleaved by a ground grid of 0.2 mm thickness.

Chamber construction All the critical steps in chamber assembly have been carried out in class 10000 clean rooms or under class 1000 huts, in controlled temperature and humidity conditions. Class 1000 clean huts have been used for all the operations involving GEM foils.

The GEM foils (50 μm thick Kapton with two-sided Cu cladding of 5 μm) were manufactured by the CERN-EST-DEM workshop and have an active area of 202×242 mm^2. The holes have a bi-conical structure with external (internal) diameter of 70 (50) μm and a pitch of 140 μm. In order to limit the damage in case of discharge, one side of the GEM foil is divided into six sectors, of about 33 mm $\times$ 240 mm. The separation between sectors is 200 μm. Each individual foil was visually inspected for defects and for leakage current. Then the foils were stretched, with a mechanical tension of about 1 kg/cm with a special device visible in figure **??**

After the GEM stretching, a fiberglass frame was glued on the GEM foil using a Ciba 2012 epoxy. Both cathode and readout pad electrodes, realized on standard 1.0 mm thick printed circuit

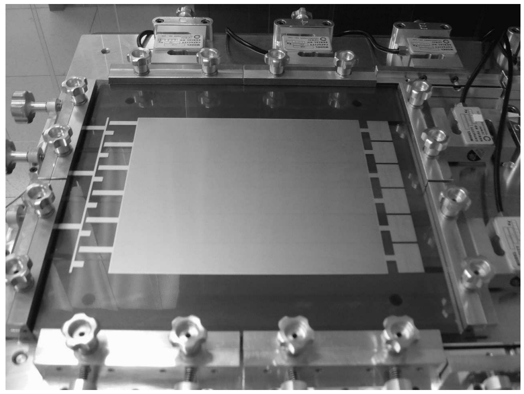

Figure 80. Stretching of a GEM foil.

board, are respectively coupled with a 1.0 mm fiberglass foil by means of a honeycomb structure, 8 mm thick. The back-panel with a 35 μm copper layer deposition on its external side is used as a

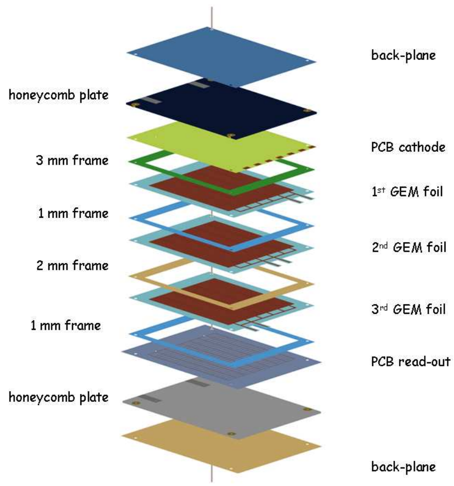

Figure 81. Exploded view of a GEM chamber

Faraday cage for the detector. The stiff cathode and pad panels act as support plates for the whole detector. These panels house two gas inlets and two gas outlets, made with machined Delrin inserts.

All fiberglass parts that are in contact with the sensitive volume of the detector, were visually inspected in order to find and eliminate any residual spikes or broken fibers, and then cleaned in a ultrasonic bath with de-mineralized water and dried in an oven at a temperature of 80 ^{o}C for one night.

The detector were built piling up and gluing together, on a reference plane machined to high precision, the single detector parts in the following order: cathode panel; the first GEM foil (GEM1) glued on a 3 mm thick fiberglass frame (see figure 107) ; the second GEM foil (GEM2) glued on a 1 mm thick frame; the third GEM foil (GEM3) glued on a 2 mm thick frame and then the last 1 mm thick frame that, followed by the pad panel, that closes the induction gap. All the gluing operation are performed using Araldite AY103 epoxy and HY991 hardener. A complete chamber is shown in figure **??**.

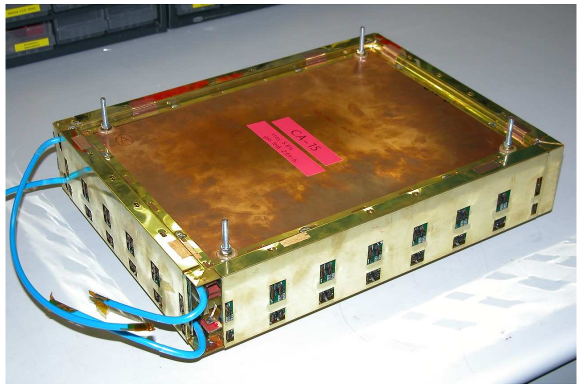

Figure 82. A triple-GEM detector fully assembled with the FE Electronics.

Quality control and testing Several quality check were performed, before the chamber assembly, on different detector components. Since the cathode and anode panels are the main mechanical structure of the detector, they were checked for planarity. Measurements on the whole panels showed that a displacement from an average plane was of the order of 50 μm (r.m.s).

The quality of GEM foils were checked performing various tests. A preliminary optical inspection was performed with a microscope to check for photolithographic imperfections. If the GEM foil passed the visual inspection, a high voltage test was performed. Such a test was done in a gas tight box, flushed with nitrogen in order to keep the relative humidity at $\sim$ 25% level. The acceptance requirement was a maximum leakage current less than 1 nA at 500 Volts. After the chamber construction the detector gas leak was measured. This measurement was performed by inflating to an overpressure of $\sim$ 10 mbar the chamber together with another chamber used as reference, and recording the overpressure decay. The typical gas leak rate of a chamber was of the order of few mbar/day, corresponding to a humidity level of the order of 100 ppm for a gas flow rate of 80 cm^3/min (as foreseen in the experiment). In order to check the uniformity response two kinds of tests were carried out. The first test was performed on a single chamber, while the second was executed on the two GEM detectors superimposed.

The uniformity gain test checked both the mechanical tolerance and the uniformity response of a single chamber. This test was performed with an X-ray gun. The current signal induced on each pad, 192 pads per chamber, was corrected for the temperature and the pressure variations. The current deviations from the average were all below 20 % in a detector surface scan, with a typical RMS $<$ 10%. Finally, a cosmic rays test was performed on all chambers fully equipped with the front-end electronics.

Discharge and aging properties The use of the new $\mathrm{Ar}/\mathrm{CO}_2/\mathrm{CF}_4(45:15:40)$ gas mixture required to demostrate the high robusteness of the GEM to discharges and ageing effect. In fact the choice of the electric field in the detector gaps and the unbalance configuration of the voltages applied to GEMs ($V_{GEM1} > V_{GEM2} > V_{GEM3}$) have been the result of an optimization of discharges produced by alpha particle. The measurements with alpha particle and a high intensity low momentum pion/proton beam, where more than 5000 discharges have been integrated without breakdown and performance losses, demostrated that GEMs in M1R1 can operate in safe conditions in the above mentioned working plateau. The large amount of CF_4 in the gas required a global irradiation test of the final chamber to check the compatibility between the construction materials (for both detector and gas system) and the gas mixture. For this reason we performed a test at the Calliope facility of the ENEA-Casaccia with the intense gamma ray flux from a ^{60} Co source. Performances of such chambers has been measured with beam test before and after the irradiation.

The results indicate that GEMs can tolerate without damages or performance losses an integrated charge of 2.2 C/ cm^2, i.e. the radiation dose foreseen in 10 years of operation[155, 156].

5.3.5 Mechanics

All the Muon System including the iron filters can be separated in two halves which can be moved on rails away from the beam pipe for maintenance and installation. The chambers, mounted on movable support walls, can move independently from the iron filters, which are normally located close to the beam pipe. Iron plugs minimize the free space among the beam pipe and the iron filters to reduce background.

In M2–M5 two rigid mechanical infrastructures built from iron beams (left and right) accomodate the suspensions for the four chamber support walls and have platforms for the racks of the electronics and gas systems. Each infrastructure can slide out of the filters if necessary. Cable chains are used to connect the DAQ cables, gas pipes and air cooling pipes with the outside. Station M1 structure is independent from the rest and the necessary racks are located on the floor of the cavern. Cable chains are used in M1 to rely the 1200 cables to the racks, whereas gas pipes are disconnected and reconnected in case of wall movement.

The walls are designed to support the weight of the half-stations with minimal thickess and are built of aluminumn sandwich plates. M1 wall is thinner to minimize material in front of the EM calorimeter. This is possible since the weight of M1 chambers is minimized thanks to the honeycomb panels.

Alignment Since the muon trigger relies on the projectivity of the stations an accurate space alignment has to be performed, in particular in the inner regions. Moreover, to minimize chamber overlaps since these would introduce ambiguities, the position of the chambers needs to be adjusted with an accuracy of about 1 mm.

The chambers are mounted on with screws on angle brackets fixed to the supporting walls. The walls are first precisely aligned by adjusting the overhead suspensions. Each chamber is then aligned relatively to the wall. The position can be adjusted vertically using spacers and horizontally via the slotted holes in the brackets. For horizontal (x-direction) and vertical (y-direction) alignment, the reference points are the support wall edge close to the beam pipe and the top edge, respectively. Finally, the equipped walls are precisely aligned together using as reference point of each half-station the center of the beam pipe. The reproducibility of the measurements is $O(1)$ mm and the reproducibility of the position after moving the walls is... (need final number).

5.3.6 Electronics

Figure 109 shows schematically the architecture of the Muon electronics. The task of the electron-

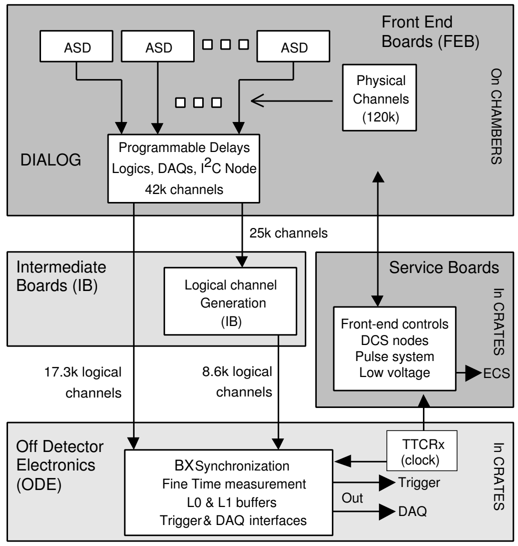

Figure 83. to be edited Simplified scheme of the Muon electronics architecture.

ics is twofold: prepare the information needed by the L0 muon trigger [133] and send the data to the DAQ system. The main steps are the following:

i. the front-end boards perform the amplification, shaping and discrimination of the chamber signals ($\approx 120,000$ physical channels). The time alignment within 1.5 ns of the different channels is also done in this step, to correct for different cable lenghts and different chamber bahavior. This is mandatory since the Muon Trigger is fully synchronous with the LHC cycle.

ii. The logical channel signals are generated by suitable logical ORs of the physical channels. This step is performed on the FE boards or on special Intermediate Boards (IB), when the logical channel span more than one FE board.

iii. The Off Detector Electronics (ODE) boards receive the logical channels. They are tagged with the number of the bunch crossing (BX identifier) and routed to the trigger processors via optical links without zero suppression.

iv. The fine time information, measured by the TDC on the ODE boards, is added and the data are transmitted to the TELL1 via optical links.

As far as possible, step ii. is performed on the chambers front-end boards in order to minimize the number of LVDS links exiting the detector. The other steps are performed on dedicated electronics boards mounted on racks on the left and right sides of the muon stations M2–M5. These racks also accomodate the HV power supplies and the gas mixers. Since the racks are fixed to the station infrastructure the cable do not have to move, thus simplifying the layout. For station M1 the racks are placed inside the bunker under the RICH and, as mentioned above, they are connected via cable chains to the chamber walls.

Radiation issues The Muon Electronics must operate reliably more than 10 years in a hostile environment. The FE boards in the regions close to the beam pipe are exposed to very high radiation dose: in M1 the maximum foreseen TID in 10 years is 5000 Gy. This decreases by only a factor 3 on M2. The maximum neutron and hadron fluences are in the range $10^{13}/\,\mathrm{cm}^2$. Therefore all the ASICs used in the FE Electronics have been produced using radiation-hard technology (see below).

The radiation decreases considerably on the detector sides where the electronics racks are located. The LHCb specifications require tolerance up to 22 Gy TID in 10 years in the electronics racks (including a safety factor of 2). This is reduced to 8.3 Gy for the HV system racks. These doses nevertheless dictate the use of radiation tolerant components or of a proper qualification of commercial components.

For complex logic design two kind of FPGA both of Actel family are used. The antifuse one time programmable a54SX falmily was chosen in the beginning for good radiation immunity and used in the IB. A more advanced type of FPGA the Flash based ProASICPLUS APA family from Actel was chosen in the design of ODE, PDM, SB (see below). Based on Actel Flash technology, ProASICPLUS devices offer reprogrammability and nonvolatility in a high density programmable logic product. The ProAsicplus is also good in radiation environment thanks to radiation hardiness of the the flash cells.

COTs (Commercial off-the-shelf) electronics was used for glue logic and signal conversion. All these chip were validated respect the radiation environment from our group or from others.

Dedicated ASICs The large number of channels in the Muon System, the flexibility necessary to adapt the readout scheme to the different regions, the necessity to synchronize all the channels in a 20 ns window, the high radiation expected in M1 and in the inner regions led us to develop three dedicated rad-hard ASICs (CARIOCA, DIALOG and SYNC) using IBM 0.25 micron technology. All chips provide/accept logical LVDS data.

Front-end board The chamber readout is performed via front-end boards (CARDIAC) [161] plugged directly on the chambers. Each CARDIAC has 16 inputs and 8 outputs, and is equipped with two CARIOCA chips and one DIALOG chip (figure **??**). A special diode circuit protects the front-end amplifiers from sparks and is mounted on a separated board.

The CARIOCA [157, 158, 159] is a front-end amplifier-shaper-discriminator chip with eight channels whose input polarity is selectable. The front-end current preamplifier can handle well the large spread in detector capacitances encountered in the Muon Chambers (from 20 pF for the GEMs to 220 pF for the M5R4 chambers). The input impedance is 50 Ohm, which is important to reduce internal pad-pad cross-talk. The chip includes tail cancellation and baseline restoring.

The CARIOCA has separate thresholds for all the channels in order to overcome the problem of channel-to-channel uniformity in the internal discriminators. ENC is about 2000 e at 0 capacitance, and increases as (42-45) e/pF. Power consumption is about 45 mW/ch.

The DIALOG [160] chip has 16 inputs to receive the outputs from two CARIOCA chips, and it performs the logical OR of the corresponding two layers of a chamber and forms the logical channels. The DIALOG is equipped with adjustable delays for every input line in such a way to allow to match in time the various input signals with a step of 0.78 ns. In addition the DIALOG also allows setting the CARIOCA thresholds and masking their inputs and contains features useful for system set-up, monitoring and debugging. Triple-voting and auto-corrected registers are used for better SEU immunity.

The CARDIAC boards are enclosed inside the chamber FC, together with radiation-tolerant voltage regulators (LH4913), which supply the necessary 2.5 V to the boards. The R4 chambers (48 pads) use only three boards and one regulator. On the mixed-readout chambers equipping R1 of M2-M3 there are 14 CARDIACs and three regulators.

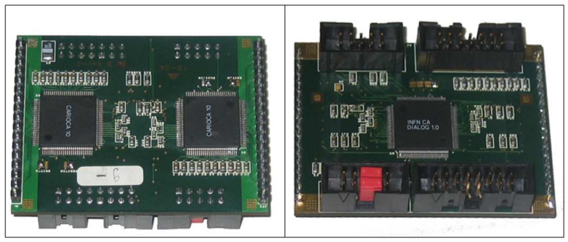

Figure 84. Top and bottom views of the CARDIAC board, showing the two CARIOCA chips and the DIALOG chip.

Given the large number of readout channels the Muon System comprises more than 8000 CARDIACs. All the boards are tested after assembly, and must pass successfully a thermal cycle to be accepted. Once mounted on the chamber the characteristics of each board, in particular with respect to noise, are measured again.

A special version of the CARIOCA chip, CARIOCAGEM [159], has been produced for use on the GEMs, with lower threshold and without tail cancellation circuit. A more compact CARDIAC card has also been designed, given the tight space available in the R1 region.

IB board Whenever the generation of the logical channels is not possible at the DIALOG level, an additional logical layer is needed. This happens in regions R2, R3 and R4 of stations M2 to M5. The layer is realized by the Intermediate Board System. The system comprises 168 IB boards. Each IB implements the logic on three Actel A54SX16A anti-fuse FPGA. Anti-fuse technology offers good tolerance against integrated radiation dose. The IB has been tested successfully up to 68.5 Gy without failures.

ODE board The ODE board contains the L0 pipelines, the L1 buffers and the DAQ interface. It synchronizes signals and dispatches them to the L0 trigger. Synchronous TFC signals are received by a TTCrx chip (one per board) [162]. On-board clock de-jittering and distribution is managed by a QPLL chip [163]. A total of 152 ODE boards are used, which are physically placed on the same corresponding IB racks.

Each board receives up to 192 logical channels (LVDS) and outputs data to the L0 Muon Trigger and the DAQ system. The signals are sent to the L0 Muon Trigger directly via 12 1.6 Gb/s optical links, each one served by one GOL chip [164]. Incoming signals are assigned the appropriate BX identifier and sent to the L0 pipelines. In parallel, the data and the four least significant bits of the associated BX Id are sent to the L0 trigger. Data resides on the L0 pipelines for 4 μs before receiving the `L0-accept` signal. Triggered data are then formatted and sent to the L1 buffer, placed on the TELL1 boards. The LVDS receivers, the L0 pipelines, the L0 derandomizer and a 4-bit TDC are integrated in a single component, the SYNC chip [160]. The chip incorporates a number of error-detection features, allowing remote control and diagnosis of possible malfunctioning in the boards. The other main board components (board controller, L0 and L1 interfaces) are based on one Actel Flash ProAsicPlus FPGA (APA450PQG208). Each ODE board also contains a CAN node, based on one Embedded Local Monitor Board (ELMB) [165], for board control via the ECS.

Channel mapping to the L0 trigger is organized by grouping the logical channels in three different ways. This is obtained placing the SYNC chips on daughter-boards of three different sizes, containing 2, 4 or 6 chips, while the ODE mother board is always the same. In the three cases 12, 8 or 6 links respectively are active on the board. The design of the ODE board has been very challenging because it copes with several, and sometime contradictory, requirements. Due to the small amount of space available, the board has a high channel density (192 differential input in a 6U card) maintaining at the same time a good level of flexibility to match the different Trigger Sector topologies. High signal integrity is also mandatory to guarantee a high quality of optical links (Bit Error Rate (BER) $< 10^{-16}$).

For radiation tolerance triple modular redundancy technique has been used whenever possible (both on the SYNC and the Board Controller) to increase SEU immunity.

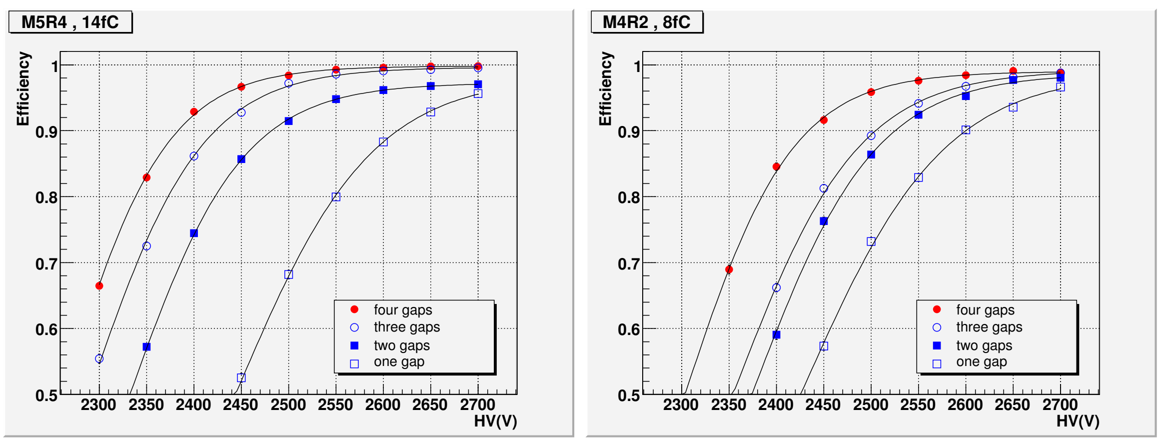

Figure 85. Plateau curves showing the efficiency in a 20 ns window for an anode-readout chamber (left) and a cathode-readout chamber (right). The four curves refer to different number of gaps being operated (see text). The threshold values in fC are also indicated, a lower value being used for cathode readout.

Experiment Control System The Muon System Experiment Control System (ECS) allows monitoring and control of all the Front-End boards. It consists of 10 Crates of equipment (6U height); each crate contains two types of modules: a Pulse Distribution Module (PDM) and up to 20 Service Boards (SB) connected via a custom backplane.

The ECS is a Distributed Control System where the processor capacity is distributed among all nodes in the system. We have chosen CANbus which has the advantages of being a multi-master protocol, with real-time capability, error correction, long communication distance and high noise immunity. The Service Board makes possible monitoring and control of up to 96 CARDIAC boards (1536 channels), as well as communication between front-end electronics and the external world. One SB hosts four ELMB [165] modules, based on an 8-bit microcontroller (Atmel ATMega128), running customized firmware. They allow communication via twelve serial links with Front-End cards and are able to communicate with each other using their on-board links.

The PDM module is used to generate and distribute a low-jitter synchronous test pulse in a chosen phase relation with the LHC machine clock, for timing alignment of the detector FE, a 40 MHz clock signal synchronous with LHC operation , and it allows distribution of 4 Canbus branches as well as routing of special messages to each SB using a custom data bus available on its backplane.

5.3.7 LV and HV systems

The Low Voltage and High Voltage systems are based on radiation-tolerant power supplies. They are installed in the pit on racks on both sides of the muon stations (M2-M5) and in the bunker (M1). For LV eight radiation tolerant power supplies are used [1], and for powering the Off Detector electronics ten. The LV is distributed to the FE-electronics on the various chambers through a LV-board with radiation-hard regulators providing the the 2.5 V for the CARDIAC boards.

The HV cabling for the wire chambers is designed to supply independently all the gaps, i.e. potentially about 5000 channels. Two distinct power supply systems are used. The first system, developed by PNPI and UF, is based on 36-channels modules and powers the R4 and R3 chambers in M2-M5. In the initial experiment running period the gaps of the R4 chambers are connected in parallel in groups of four via patch panels. In this way the number of independent HV channels is reduced from 3840 to 1536. The gaps in parallel belong always to different chambers to minimize the loss of efficiency in case, e.g., one gap should become shorted. The second system is a commercial one based on 32-channel modules [15] . It powers the chambers more exposed to radiation, whose operating conditions are more critical, and for this reason all the 1152 channels can be individually controlled and monitored.

Both HV systems have been tested for radiation resistance (hadrons, total dose and neutrons) and are apt to satisfy LHCb requirements for 10-year operation with a 10-fold safety factor.

The GEM detectors require several voltages to operate. To achieve maximum flexibility and safety we preferred not to use resistive voltage dividers and we designed a customized system [169] which is installed in the counting room.

The LV and HV systems are controlled via PVSS as part of the ECS.

5.3.8 Performances

MWPC Extensive tests have been performed on prototypes and final MWPC both on beam tests and with cosmic rays at the production sites, to measure the main chamber parameters.

Whereas MWPCs are expected to operate normally with four gaps in OR (two in case of M1 chambers) it is possible that a chamber will be operated with less gaps. Figure 111 shows the efficiency in a 20 ns window as a function of the applied voltage, for the largest chamber with

[1] Wiener Maraton
[15] CAEN Easy 3000 and A3535P

anode readout (M5R4) and one chamber with cathode readout (M4R2). The curves are given for the standard four-gap OR configuration, but also for three-gap configuration (one double-gap fully working and the second with one gap shut off), for the double-gap (the basic buiding block) and finally for the single-gap case. The curves show clearly the large improvement obtained by adding gaps. The plots also show that, in case one gap should be shut off during the data taking, the redundancy built into the chambers will make the loss in performance negligible.

Another important parameter is the cluster size, i.e. the average number of pads fired by an incoming particle. The cluster size affects directly the space resolution since the yes-no readout does not allow an interpolation between adiacent pads. A non-negligible contribution to the cluster size comes simply from the geometry given the fact that inclined tracks can cross two adjacent pads belonging to different gaps. Another (intrinsic) contribution comes from the chamber itself (inductive and capacitive cross-talk in the chamber and in the electronics). The design criteria required that the intrinsic cluster size should be less than 1.2. This is well satisfied as it can be seen in figure **??** which shows the cluster size distribution for the M5R4 and M2R4 chambers measured with perpendicular tracks [166]. Cosmic rays have also been used to check the gas gain uniformity

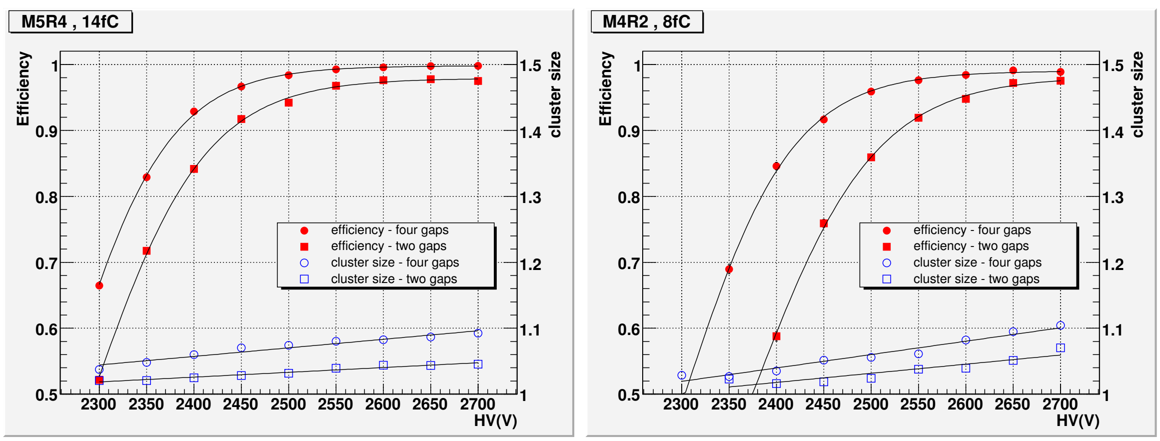

Figure 86. Efficiency and cluster size for the four-gap and double-gap configurations vs. HV. Anode readout-type chamber (left) and cathode readout-type chamber (right).

on final chambers [167].

Special tests have been performed also at the CERN Gamma Irradiation Facility (GIF). The chamber performance were studied for several detector gap configurations and different values of the background rate from the source. The details of the beam test studies can be found in [168].

The CARDIAC board did not show any anomalous behavior up to 1.2 MHz/channel, which is two times the highest rate per channel per chamber expected in the hottest region of the LHCb Muon System equipped with MWPC (M2R1). Up to a photon flux of 10 kHz/cm^2 (70% of the M2R1 rate/cm) there was no visible space-charge effect on the chamber efficiency and time performance.

GEM Chambers Triple-GEM detectors with the final FE electronics (CARIOCAGEM) have been tested both in a dedicated 40 MHz test beam a and with cosmics in the lab.

Efficiency and cluster size of the chambers in a 20 ns window as a function of the gain, for each single triple-GEM detector and the logical OR, are shown in figure 113.

Thanks to the longer shaping time of the CARIOCAGEM with respect to the CARIOCA chip the beginning of the efficiency plateau has moved from a detector gain of 6000 to 4000 with the choice of this electronics. This does not only decrease by 30% the integrated charge in the detectors, but it also allow the detectors to be operated in even safer conditions with respect to discharges.

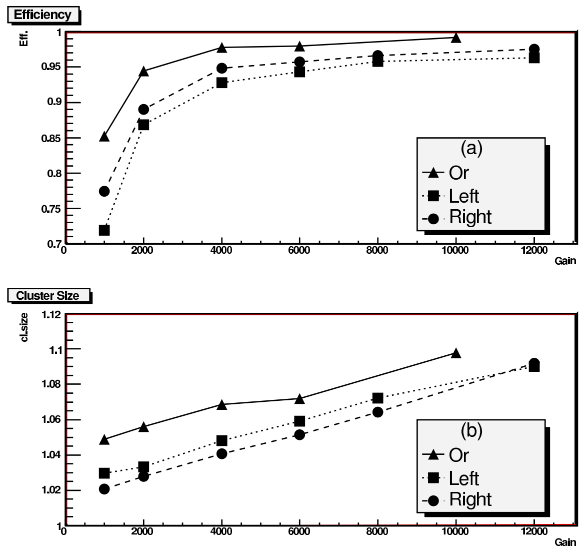

Figure 87. Efficiency and cluster size for one 3-GEM chamber as a function of the gain.

6. Trigger

6.1 Introduction

The LHCb experiment plans to operate at an average luminosity of 2×10^{32} cm^{-2}s^{-1}, much lower than the maximum design luminosity of the LHC, which makes the radiation damage more manageable. Futhermore, the number of interactions per crossing is dominated by single interactions, which facilitates the triggering and reconstruction by assuring low channel occupancy. Due to the LHC bunch structure and low luminosity, the crossing frequency with interactions visible[16] by the spectrometer is about 10 MHz, which has to be reduced by the trigger to about 2 kHz, at which rate the events are written to storage for further offline analysis. This reduction is achieved in two trigger levels as shown in Figure **??**: Level-0 and the High Level Trigger (HLT). Level-0 is implemented in custom electronics, while the HLT is executed on a farm of commodity processors.

At a luminosity of 2×10^{32} cm^{-2}s^{-1} the 10 MHz of crossings with visible pp interactions are expected to contain a rate of about 100 kHz of $b\bar{b}$-pairs. However, in only about 15% of the events will at least one B-hadron has all its decay products contained in the spectrometer acceptance.

[16]An interaction is defined to be visible if it produces at least two charged particles with sufficient hits in the VELO and T1–T3 to allow them to be reconstructible.

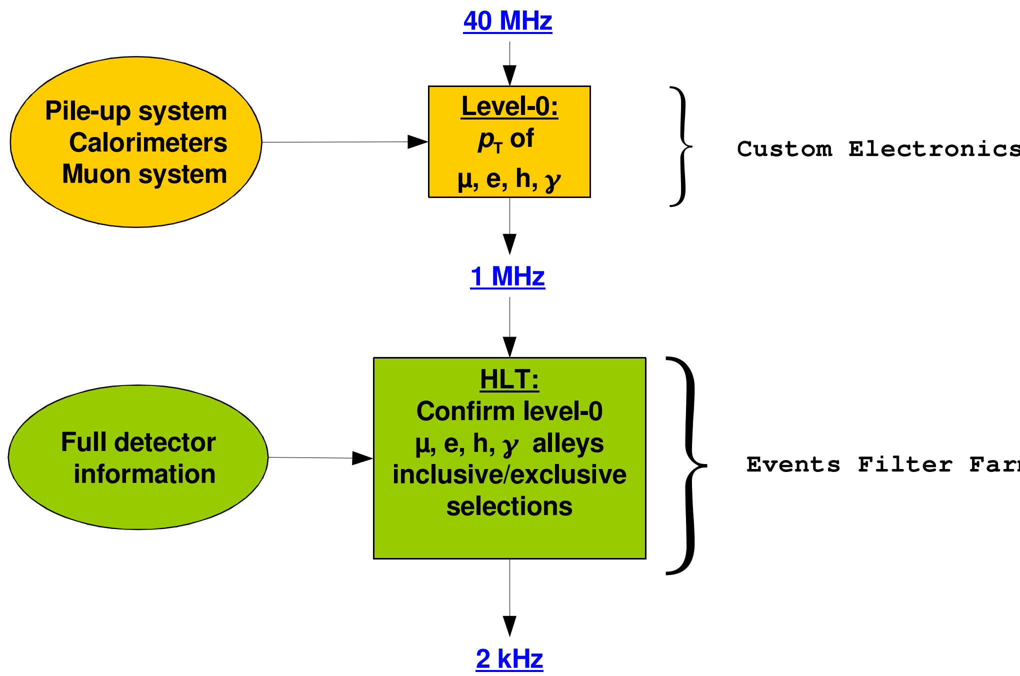

Figure 88. Scheme of the LHCb trigger.

Furthermore the branching ratios of B-particles used to study CP violation are typically less than 10^{-3}. The offline selections exploit the relatively large b mass and lifetime to select those b-hadrons, and stringent cuts have to be applied to enhance signal over background and thus increase the CP sensitivity of the analysis. Hence the requirement for the trigger is to achieve the highest efficiency for these offline selected events, and for a large variety of final states.

The trigger allows overlapping and pre-scaled triggers. It is possible to emulate the trigger from the data written to storage, which will give an additional handle on trigger efficiencies and possible systematics.

The purpose of Level-0 is to reduce the LHC beam crossing rate of 40 MHz to the rate of 1 MHz where the entire detector can be read out. Due to their large mass B-hadrons decay yielding a large E_T lepton, hadron or γ, hence Level-0 reconstructs:

- the highest E_T hadron, electron and photon clusters in the calorimeters,
- the two highest p_T muons in the muon chambers,

and the information is collected by the Level-0 Decision Unit to select events. Events can be rejected based on global event variables such as charged track multiplicities and the number of interactions, as reconstructed by the Pile-Up system, to assure that the selection is based on b-signatures rather than large combinatorics, and that these events will not occupy a disproportional fraction of the data-flow bandwidth or available processing power in subsequent trigger levels.

All Level-0 triggers are fully synchronous. Their latency does not depend upon occupancy nor on bunch crossing history. All Level-0 electronics is implemented in full custom boards.

The purpose of the HLT is to reduce the rate down to 2 kHz by using data of all sub-detectors. The measurements aimed at by LHCb requires a very high precision: hence systematic errors must be controlled to a very high degree. Amongst the 2 kHz of HLT accepted events, a large fraction is dedicated to precise calibration and monitoring of the detector and its performances.

The generic HLT algorithms refine candidates found by the Level-0 trigger. It is divided in four independent alleys: a) muons, b) muons and hadrons close to each other, c) hadrons and d) electrons, γ and π^0. The alleys to be followed are selected from the Level-0 decision. Each alley consists of four main steps:

- Level-0 confirmation,
- fast tracking using VELO and TT information to localize primary vertexes and to find tracks with a large p_{T} and large impact parameter,
- more precise measurement of the p_{T} using the T stations for tracks found previously and look for secondary vertexes,
- selection criteria specific to the alley.

After the generic algorithms interesting final states are selected using inclusive and exclusive criteria. The RICH information is not currently used by the HLT, and selection cuts have to be relaxed compared to the final selection to study the sensitivity of the selections and profit from refinements to the calibration constants.

6.2 Level 0 trigger

6.2.1 Overview

The Level-0 trigger shown in Figure 115 is subdivided in three components: the Pile-Up system, the Level-0 calorimeters and the Level-0 muon. Each components is connected to one detector and to the Level-0 Decision Unit collecting all candidates found to master the final decision.

The Pile-Up system aims at distinguishing between crossings with single and multiple visible interactions. It uses four silicon sensors of the same type as those used in the VELO to measure the radial position of tracks. The Pile-Up system provides the position of the primary vertices candidates along the beam-line and a measure of the total backward charged track multiplicity.

The Calorimeter Triggers looks for high E_{T} particles: electrons, γ, π^0 or hadrons. It forms clusters by adding the E_{T} of 2×2 cells and selecting the clusters with the largest E_{T}. Clusters are identified as electron, γ or hadron depending on the information from the SPD, PS, ECAL and HCAL Calorimeter. The E_{T} of all HCAL cells is summed to reject crossings without visible interactions and to reject triggers on muon from the halo. The total number of SPD cells with a hit are counted to provide a measure of the charged track multiplicity in the crossing.

The muon chambers allow stand-alone muon reconstruction with a p_{T} resolution of $\sim 20\%$. Track finding is performed by processing elements, which combine the strip and pad data from the five muon stations to form towers pointing towards the interaction region. The Level-0 Muon trigger outputs the two muons with the highest p_{T} for each quadrant of the muon detector.

The Level-0 Decision Unit collects all information from Level-0 components to form the Level-0 Trigger. The Level-0 Decision Unit is able to perform simple arithmetic to combine all signatures into one decision per crossing. This decision is passed to the Readout Supervisor which transmits it to the front-end electronics.

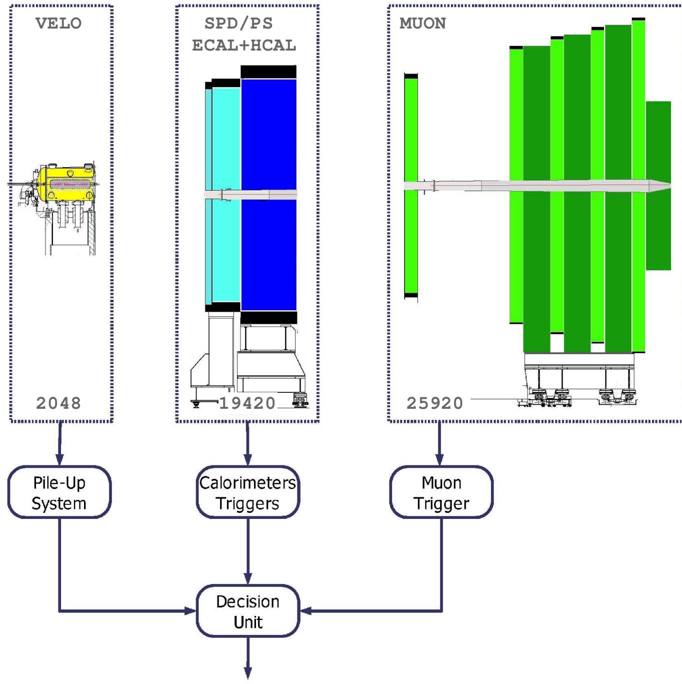

Figure 89. Overview of the Level-0 trigger. The Pile-Up system receives every 25 ns 2048 channels from the pile-up detector, the Level-0 calorimeters 19,420 channels from Scintillating Pad Detector, Preshower, Electromagnetic and Hadronic Calorimeters while the Level-0 muon handles 25,920 logical channels from the muon detector.

The latency of the Level-0, which is the time elapsed between a pp interaction and the arrival of the Level-0 trigger decision at the front-end electronics is fixed to 4 μs. This time includes the time-of-flight, cable length and all delays in the front-end electronics, leaving 2 μs for the processing of the data in the Level-0 trigger to derive a decision.

The Level-0 Trigger provides a decision for each bunch-crossing. Therefore, the architecture is pipelined and massively parallel.

6.2.2 Architecture

Calorimeter trigger A zone of 2 by 2 cells is used, large enough to contain most of the energy, and small enough to avoid overlap between various particles. Only the particle with the highest E_T is looked at. Therefore at each stage only the highest E_T candidate is kept, to minimize the number of candidates to process.

These candidates are provided by a three step selection system as shown in Figure **??**:

- A first selection of high E_T deposits is performed on the Front-End card, which is the same for ECAL and HCAL. Each card handles 32 cells, and the highest E_T sum over the 32 sums of 2×2 cells is selected. To compute these 32 sums, access to cells in other cards is an important issue.

- The Validation Card merges the ECAL with the PS and SPD information, prepared by the preshower front-end card. It identifies the type of electromagnetic candidate, electron, γ and π^0. Only one candidate per type is selected and sent to the next stage. The same card

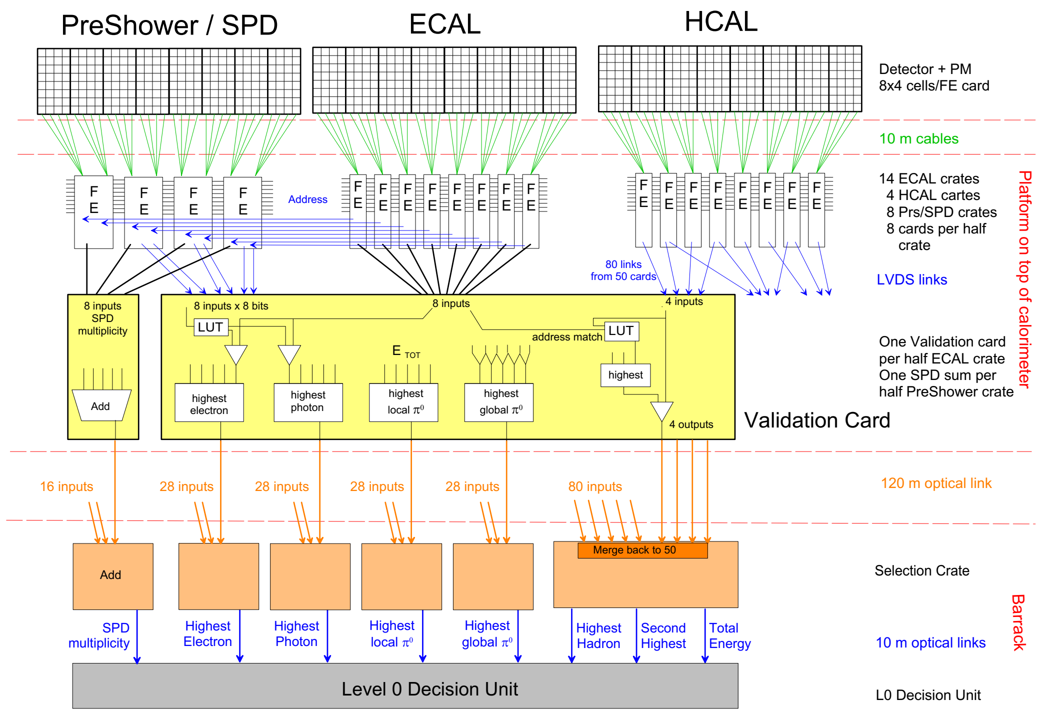

Figure 90. Overview of the Level-0 Calorimeters Architecture.

also adds the energy deposited in ECAL in front of the hadron candidates. A similar card computes the SPD multiplicity in the PreShower crates.

- The Selection Crate selects the candidate with the highest E_{T} for each type, it also produces a measure of the total E_{T} in HCAL and the total SPD multiplicity.

The first two steps are performed on the platform of the calorimeter, at a location with a radiation dose below 50 Gy during the whole lifetime of the experiment, and where single event upsets are expected to occur. Each component has been tested for radiation and single event upsets robustness. Anti-fuse PGAs are used, as well as "triple voting" techniques.

The trigger interface is housed in one Anti-fuse FPGA from Actel for ECAL/HCAL front-end cards and in one Flash EEPROM based FPGA for PS/SPD front-end boards. There is a large data flow between these components at a frequency of 40 MHZ through a dedicated backplane where interconnections are realized by point-to-point links running a multiplexed LVDS signals at 280 MHz. The same backplane is used for PreShower, ECAL and HCAL crates.

The Validation card is a 9U board with 16 layers. Clusters, PS and SPD hit maps arrive through the backplane via 20 LVDS links running at 280 MHz. The cluster identification is performed by two ProAsic FPGAs from Actel. Electrons, γ, hadrons and π^0 candidates are transmitted to the selection crates via an 8-channel optical mezzanine. The latter serialize data at 1.6 Gbit/s and drive a ribbon of 12 fibres. The control of the Validation and Calorimeter front-end cards are performed by a SPEC interface.

The Selection crate is located in the barrack behind the shielding wall in a radiation free place. It is a modular system which consits of eight Selection Boards. The design of the selection board

Boards	Number
ECAL/HCAL front-end	246
PS/SPD front-end	100
8-channels optical mezzanine	80
1-channels optical mezzanine	40
Validation card	28
SPD Control board	16
Selection Board	8

Table 10. Boards of the Level-0 Calorimeters.

is unique but adaptable to perform both the electromagnetic and the hadron clusters selection. The electromagnetic clusters selection is performed by using one board for each cluster type (electron, γ, π^0) while the hadron selection requires three boards. The results of the two first boards are transmitted to the third one where the final selection is performed. Finally, one board is also used to sum up the SPD multiplicity. The selection board is a 9U board with 16 layers. Input arrive via 28 optical links grouped in three ribbons of 12 fibres. High speed serial signal are deserialized by 28 TLK2501 chips from Texas Instrument. The selection of the highest candidate is performed by six FPGAs from Xilink Virtex II family. The selected candidates is sent to the Level-0 Decision Unit via 1-channel optical mezzanine driving high speed optical link. Inputs and output of the Selection Board are sent to the data acquisition system via two high speed optical link connected to the TELL1 board. The control of the Selection Board is performed by a credit card PC.

The total number of boards for the Level-0 Calorimeters Trigger is summarized in Table 11.

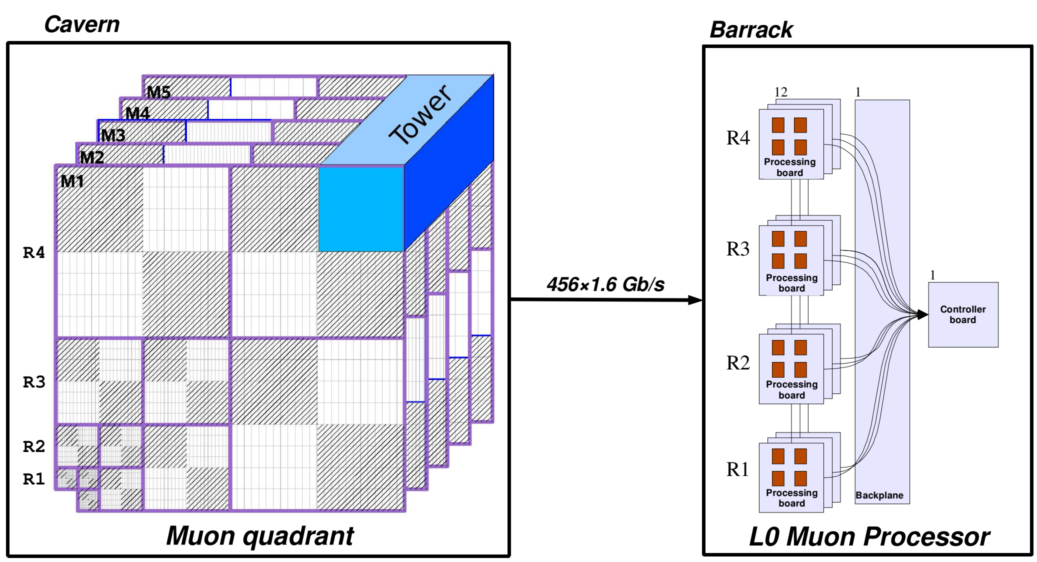

Figure 91. Overview of the Level-0 Muon Architecture. Level-0 Muon processor are located in the barrack behind the shielding wall in a place immune to radiation effects.

Muon trigger An overview of the Level-0 Muon architecture is given in Figure **??**. Each quadrant of the muon detector is connected to a Level-0 Muon processor via 456 optical links grouped in 38 ribbons containing 12 optical fibres each. An optical fibre transmits serialized data at 1.6 Gbits/s over a distance of approximately 100 meters.

A Level-0 Muon processor looks for the two muon tracks with the largest p_{T}. The track finding is performed on the logical pad layout. It searches for hits defining a straight line through the five muon stations and pointing towards the interaction point. The position of a track in the first two stations allows the determination of its p_{T}.

Seed of the track finding algorithm are hits in M3. For each logical-pad hit in M3, an extrapolated position is set in M2, M4 and M5 along a straight line passing through the hit and the interaction point. Hits are looked for in these stations in search windows termed Field Of Interest (FOI), approximately centred on the extrapolated positions. FOIs are open along the x-axis for all stations and along the y-axis only for stations M4 and M5. The size of the FOI depends on the station considered, the level of background and the minimum-bias retention required. When at least one hit is found inside the FOI for each station M2, M4 and M5, a muon track is flagged and the pad hit in M2 closest to the extrapolation from M3 is selected for a subsequent use.

The track position in station M1 is determined by making a straight-line extrapolation from M3 and M2, and identifying in the M1 FOI the pad hit closest to the extrapolation point.

Since the logical layout is projective, there is one-to-one mapping from pads in M3 to pads in M2, M4 and M5. There is also a one-to one mapping from pairs of pads in M2 and M3 to pads in M1. This allows the track finding algorithm to be implemented using only logical operations.

To simplify the processing and to hide the complex layout of stations, we subdivided the muon detector in 192 towers pointing toward the interaction point as shown in Figure **??**. A tower contains 288 logical pads[17] with the same layout. Therefore the same algorithm can be executed in each tower. Each tower is connected to a processing element, the key component of the Level-0 Muon processor.

To collect data coming from a tower spread over five stations and to send them to a unique processing element we use a fanout panel close to the muon processor.

Processing elements have to exchange a huge number[18] of logical channels to avoid inefficiency on the border of a tower. The topology of data exchange depends strongly on the location of the tower.

A processing element runs in parallel 96 tracking algorithms, one per M3 seed, on logical channels coming from a tower. It is implemented in a FPGA named Processing Unit (PU). A processing board contains four PUs and an additional FPGA to select the two muons with the highest transverse momentum within the board. A Level-0 Muon processor houses 12 Processing Boards, a custom backplane and a controller board. The custom backplane is mandatory to exchange logical channels between PUs. The controller board collects candidates found by the 12 Processing Boards and selects the two with the highest p_{T}. It also distributes signal coming from the TTC.

The Level-0 Muon implementation relies on the massive use of multigigabit serial links deserialized inside FPGAs. Processors are interfaced to the outside world via optical links while processing elements are interconnected with high speed copper serial links.

[17] 48 pads from M1, 2×96 pads from M2 and M3, 2×24 pads from M4 and M5.

[18] A processing element handles 288 logical pads. It sends a maximum of 224 and receives a maximum of 214 logical channels from neighbouring elements.

Boards	Number
Processing Board	48
Controller Board	4
Backplane	4

Table 11. Boards of the Level-0 Muon.

The Processing Board contains five FPGAs from the Stratix GX family and 92 high speed serial links with serialiazer/deserializer embedded in FPGAs. The board sends data to the data acquisition system via two high speed optical links. The processing board is remotely controlled via Ethernet by a credit card PC running Linux. The size of the printed circuit is 366.7×220 mm. It is composed of 18 layers where 1512 components are mounted. The power consumption is less than 60 W.

The Controller Board contains two FPGAs from the Stratix GX family. The board shares a lot of common functionalities with the Processing Board: same credit card PC, same mechanism to send information to the data acquisition system. The size of the printed circuit is 366.7×220 mm. It is composed of 14 layers where 948 components are mounted. The power consumption is less than 50 W.

The backplane contains 15 slots: 12 for the Processing Boards, one for the Controller Board and two for test. It distributes power supplies, signals coming from the TTC and assures the connectivity between the processing elements via 288 single-ended links (40 MHz) and 110 differential high speed serial links (1.6 Gb/s). The size of the 18 layers printed circuit is $395,4 \times 426,72$ mm.

The total number of boards for the Level-0 Muon Trigger is summarized in Table 12.

Pile-Up system The Pile-Up system consists of two planes (A and B) perpendicular to the beam-line and located upstream of the VELO. Each 300 μm thick silicon plane consists of two overlapping VELO R-sensors, which have strips at constant radii, and each strip covers 45°. In both planes the radii of track hits, $\rightarrow$ and r_b, are recorded. The hits belonging to tracks from the same origin have the simple relation $k = r_b/\rightarrow$, giving:

$$z_v = \frac{kz_a - z_b}{k-1} \tag{6.1}$$

where z_b , z_a are the detector positions and z_v is the position of the track origin on the beam axis, *i.e.* the vertex. The equation is exact for tracks originating from the beam-line. All hits in the same octant of both planes are combined according to equation 6.1 and the resulting values of z_v are entered into an appropriately binned histogram, in which a peak search is performed, as shown in Figure **??**. The resolution of z_v is limited to around 3 mm by multiple scattering and the hit resolution of the radial measurements. All hits contributing to the highest peak in this histogram are masked, after which a second peak is searched for. The height of this second peak is a measure of the number of tracks coming from a second vertex, and a cut is applied on this number to detect multiple interactions. If multiple interaction are found, the crossing is vetoed.

The architecture of the Pile-Up system is shown in Figure 119. It uses the signals of the integrated comparators of the Beetle chips on its four hybrids. The output of the four neighbouring

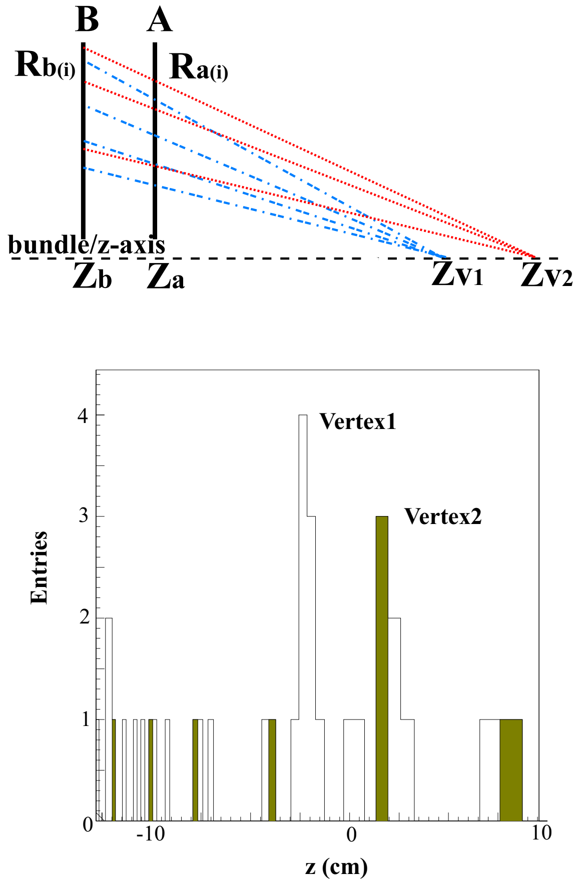

Figure 92. Basic principle of detecting vertexes in a crossing. The readout hits of plane A and B are combined in a coincidence matrix. All combinations are projected onto a z_v–histogram. The peaks indicated correspond to the two interaction vertexes in this particular Monte-Carlo event. After the first vertex finding the hits corresponding to the two highest bins are masked, resulting in the hatched histogram.

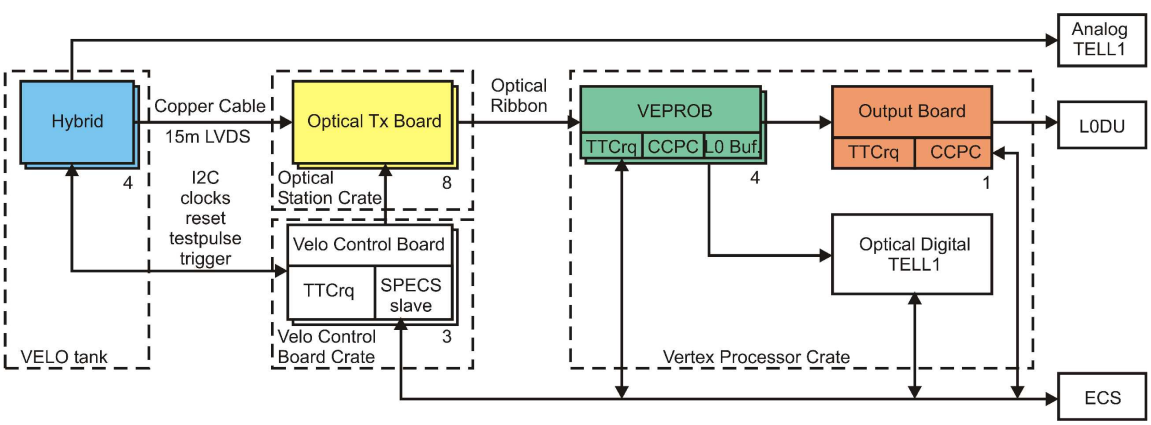

Figure 93. Overview of the Level-0 Pile-Up Architecture.

comparators is OR-ed, resulting in 256 LVDS links running at 80 Mbit/s per hybrid, which send the Level-0 signals to eight Optical Transmission Boards. Two Optical Transmission Boards cover a quadrant of the pile-up system. They time align and multiplex input hit maps to four Vertex

Boards	Number
Hybrids	4
Optical Transmition Board	8
Vertex Processing Board	4
Output Board	1

Table 12. Boards of the Level-0 Muon.

Processing Boards. Hit maps of the first bunch crossing are sent to the the first Vertex Processing Board in four consecutive clock cycles while hit maps of the second bunch crossing are going to the second Vertex Processing Board in four consecutive clock cycles. Bunch-crossing are distributed over the four Vertex Processing Board in a round-robin fashion. The Optical transmision board is a 9U board controlled by a SPEC interface via the VELO control board.

Vertex Processing Boards are 9U boards located in the radiation-free electronics barracks. They are connected to the Optical Transmissions Boards via 24 high speed optical links. The vertex Processing board is the key components of the pile-up system. It houses a large FPGA from Xilinx Virtex II family where the pile-up algorithm runs. A board handles one event over four and send its trigger decision which is the result of the algorithm to the output board via a high speed copper link (1.6 Gb/s). The board is controlled by a credit card PC and send its input/outputs to the data acquisition system via two high speed optical links after reception of a Level-0-yes from the TTC system.

The output board is a simple 9U board multiplexing the inputs from the Vertex Processing Board and outputting the data to the Level-0 Decision Unit. In addition, the output boards makes histograms of trigger decision made by the pile-up system. These histograms are accessible via the ECS interface.

The total number of boards for the Pile-Up system is summarized in Table **??**.

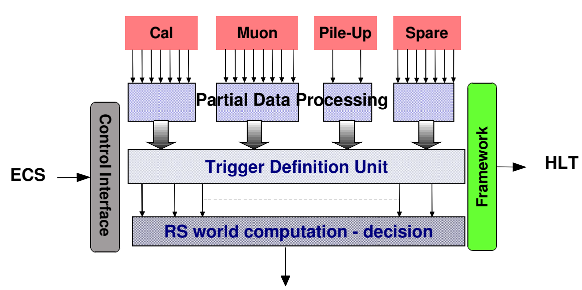

Figure 94. L0DU architecture

Decision Unit The Level-0 Decision Unit receives information from the calorimeter, muon and pile-up sub-triggers at 40 MHz, which arrive at different fixed times. The computation of the decision can start with a sub-set of information coming from a Level-0 sub-trigger, after which the sub-trigger information is time aligned. An algorithm is executed to determine the trigger decision, and a summary bank is constructed. The decision is sent to the Readout Supervisor, which has the ultimate decision about whether to accept an event or not. The Readout Supervisor is able to generate and time-in all types of self-triggers (random triggers, calibration, etc) and to control the trigger rate by taking into account the status of the different components in order to prevent buffer overflows and to enable/disable the triggers at appropriate times during resets.

The architecture of the Level-0 Decision Unit is shown in Figure 120. For each data source, a Partial Data Processing system performs a specific part of the algorithm and the synchronisation between the various data sources. Then a trigger definition unit combines the information from the above systems to form a set of trigger conditions based on multi-source information.

The trigger conditions are logically OR-ed to obtain the Level-0 decision after they have been individually downscaled if necessary.

The level-0 decision unit is based on the TELL1 board where optical cards are replaced by a single mezzanine where the Level-0 Decision Unit hardware is implemented. Inputs are received on two ribbons of 12 high speed optical links. Serial signals are deserialized by 24 TLK2501 from Texas Instruments and sent to two large FPGAs from the Stratix Family. Inputs and Outputs are sent to the data acquisition system via the TELL1 mother board.

6.2.3 Technology

The implementation of the Level-0 Trigger relies on the massive use of large Field Programmable Gate Arrays, high speed serial links and common techniques allowing an easy debug and commissioning.

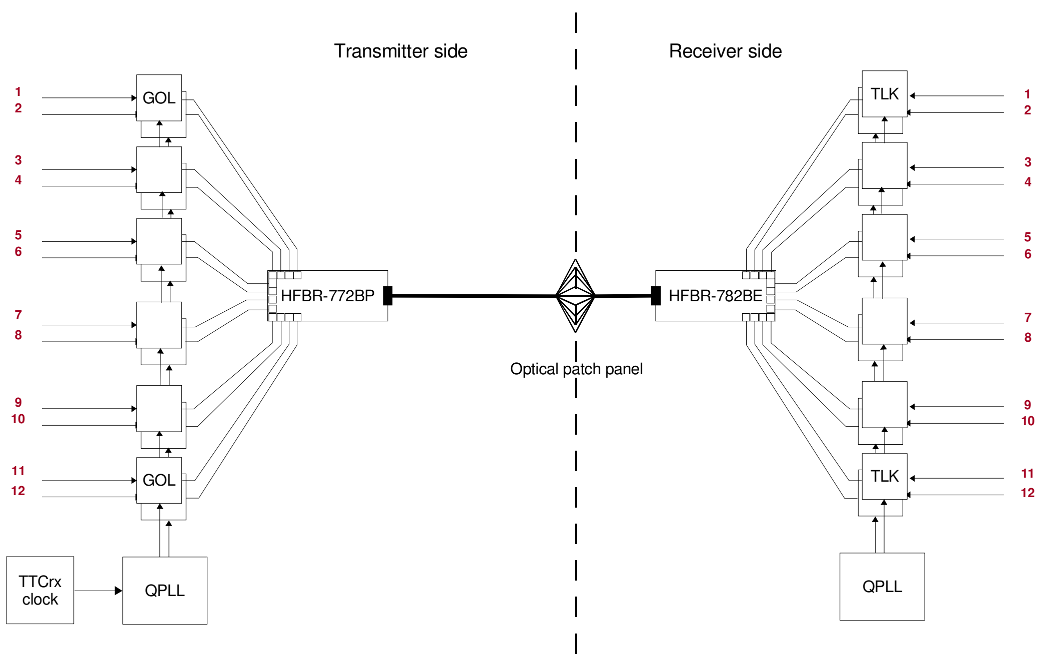

Figure 95. Overview of ribbon optical link.

High speed links The principle retained to transport information from the front-end electronics to Level-0 trigger boards located in the barrack is based on three concepts:

- serialization of the detector data;

- use of optical links as transport media;
- use of high density devices.

High speed serial transmission reduces the number of signal lines required to transmit data from one point to another. It also offers a high level of integration with many advantages: high reliability for data transfer over 100 meters; complete electrical isolation avoiding ground loops and common mode problems. In addition, the integration of several high speed optical links in a single device increases data rate while keeping the component count manageable at a reasonable cost.

Ribbon optical links integrated twelve optical transmitters (fibres, receivers) in one module. The important benefit of ribbon optical links is based on low-cost array integration of electronic and opto-electronic components. It also provides a low power consumption and a high level of integration.

An overview of ribbon optical link developed for the Level-0 trigger is shown in Figure **??**. The emitter stage relies on twelve serializer chips connected to one optical transmitter. The serializer is the GOL, a radiation hard chip designed by the CERN microelectronic group, transforming every 25 ns a 32-bit word into a serial signal with a frequency of 1.6 GHz using a 8B/10B encoding. High frequency signals are converted into an optical signal by the 12-channels optical transmitter from Agilent HFBR-772BE. The module is designed to operate multimode fibers at a nominal wavelength of 850 nm.

Initially the LHC clock distribution was not intended to be used for optical data transmission. Hence, it does not fulfill the severe jitter constraints required by high speed serializers. The GOL requires a maximum jitter of 100 ps peak to peak to operate correctly whereas the LHC clock jitter is as large as 400 or 500 ps. To reduce the jitter, a radiation hard chip, the QPLL, was designed by the CERN microelectronics group. It filters out the jitter up to an acceptable value with the help of a reference quartz crystal associated to a phase locked loop.

The emitter side is close to the detector in a place where the radiation dose is below 50 Gy over 10 years where single event upsets (SEU) are expected to occur. The GOL and QPLL chips are radiation hard chips immune to SEU. However, the optical transceiver is a commercial component designed to work in an environment free of radiation. An irradiation campaign took place at the Paul Scherrer Institute in December 2003. The component works within its specification up to a total dose of 150 Gy. The cross-section for single event upsets is equal to $(4.1 \pm 0.1) \times 10^{-10}$ cm^2 per single optical link.

The physical media between the front-end electronic boards and the Level-0 trigger board consists of ribbons of twelve fibers with MPO connectors on both side ($\sim$ 10 m.), MPO-MPO patch panels, long cables containing eight ribbons with MPO connectors ($\sim$ 80 m.), fanout panels (MPO-MPO or MPO-SC), short ribbons of twelve fibers ($\sim$ 3 m) with MPO connector on one side and a MPO or 12 SC connectors on the other side.

The receiving side is the mirror of the emitting side. Optical signals are converted into 1.6 Gb/s serial electrical signal by the 12-channels optical receiver HFBR-782BE. The twelve high frequency signals are deserialized into 16-bit words at 80 MHz by twelve TLK2501 chips from Texas Instrument. The receiving side is located in the barrack behind the shielding wall in a radiation free place. Therefore standard components can be used. In the muon processing board, where the density of input signal is high, TLK2501 chips are replaced by serializer/deserializer embedded in the Stratix GX FPGA.

The routing of the differential high speed trace between serializer/deserializer and the optical transceiver requires a lot of care since the geometry of the tracks must be totally controlled to guaranty a good impedance matching and to minimize electromagnetic emissions to the environment as well as sensitivity to electromagnetic perturbations from the environment.

Performance of the optical link have been measured with several setups in different ways. The bit error rate measured with Lecroy SDA11000 Serial Data Analyser is below 10^{-16} for a single fiber of 100 meter long.

Field Programmable Gate Arrays Three FPGAs technologies are used in the Level-0 trigger. They are characterized by the way they are configured:

- Anti-fuse based FPGAs (Actel AX family), that can be programmed only once;
- Flash-EEPROM based FPGAs (Actel pro-ASIC family), that can be erased and reprogrammed;
- Ram-based FPGAs (Altera Acex, Flex, Stratix and Stratix GX families or Xilinx Virtex family) that can be reprogrammed an unlimited number of times.

Anti-fuse and Flash FPGAs are used in the front-end boards close to the detector and therefore exposed to radiations. These components have been tested in heavy ion beams and have shown very low sensitivity to single event upsets andă single event latch-up. Special mechanisms like triple-voting or horizontal and vertical parity are implemented to increase the protection of registers containing critical data. Dose effects begin to appears in Flash based FPGAs for doses an order of magnitude above the total dose received during 10 years by the trigger front-end electronics.

Ram-basedă FPGAs are known to be very sensitive to single event upsets. For this reason their use is restricted to boards located in the barracks which is a radiation-free area.

All the FPGAs used in the trigger allow to give a good visibility on internal nodes behavior during the debug by providing embedded logic analyzer features (Silicon Explorer for Actel, SignalTap for the largest components of the Altera family and Chipscope for the Xilinx family).

Debugging and Monitoring tools Each Level-0 trigger board embeds a credit card PC or a SPEC components which is interfaced to FPGAs by a custom 16 bits bus. Therefore, we control the operation of any FPGAs and implement error detection mechanisms, error counters, spy and snooping mechanisms.

To test a complete sub-trigger in a stand alone mode, we implement data injection buffer to substitute input data. Results of the processing can be read back via the credit card PC at the output of dedicated SPY memories

The level-0 trigger is a very complicate system. Any malfunctions can therefore be difficult to understand and interpret. At each stage we log the input and results of the processing. In addition, we develop a software emulator reproducing the behaviour of the hardware on a bit to bit basis. By comparing results computed by the hardware with those of the emulator run on the same input data, we can point out any faulty components.

6.3 High Level Trigger

The High Level Trigger (HLT) consists of a C++ application which is running on every CPU of the Event Filter Farm (EFF), which contains between 1000 and 2000 computing nodes, and is described in Section 7. Each HLT application has access to all data in one event, and thus in principle could be executing the off-line selection algorithms. But given the 1 MHz output rate of the Level-0 trigger and the limited CPU power available, the HLT aims at rejecting the bulk of the events by using only part of the full information which is available.

The schematic of the different trigger sequences in the HLT is shown in Figure 122. The

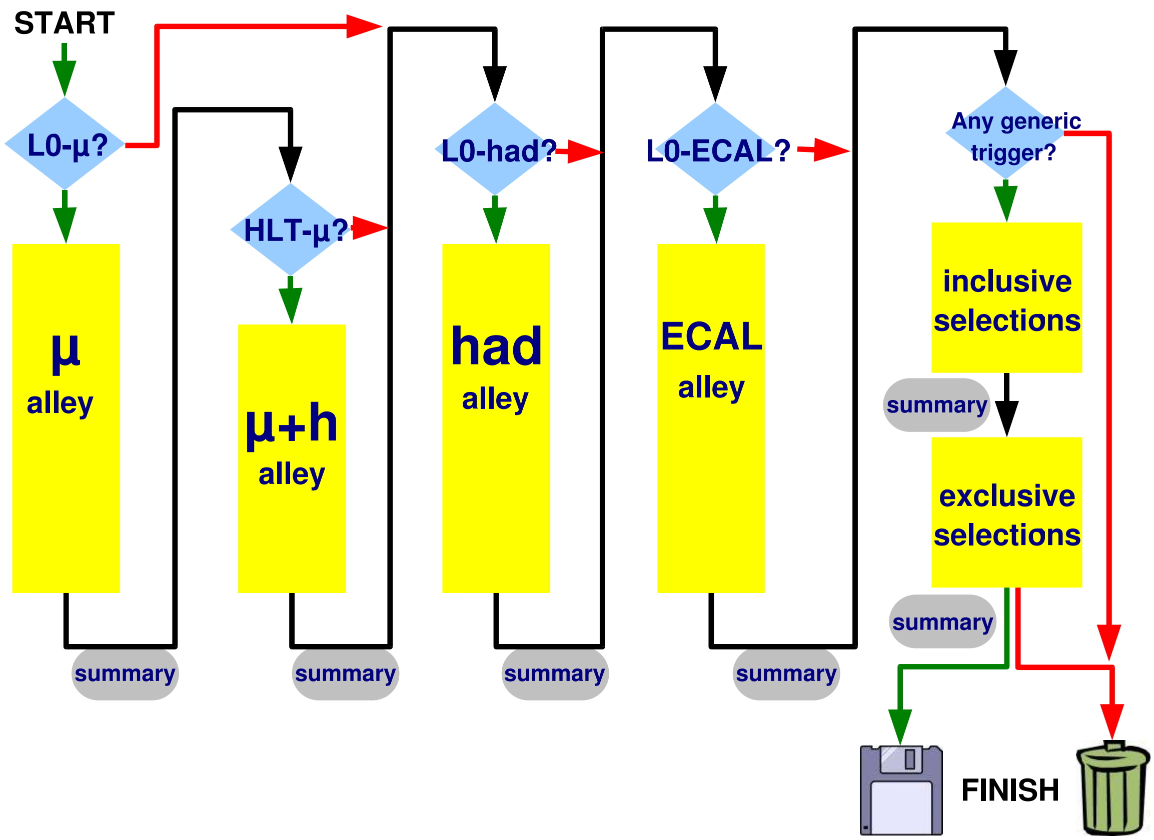

Figure 96. Flow-diagram of the different trigger sequences in the HLT. Squares represent reconstruction algorithms, while diamonds indicate where the trigger decisions are taken.

HLT starts with so-called alleys, where each alley addresses one of the trigger types of the Level-0 trigger, enriching the B-content of the events by refining the Level-0 objects, and adding impact parameter information. Most Level-0 triggers are only selected due to one Level-0 trigger type, and hence will only be processed by one alley. About ~ 15 % of the Level-0 events are selected by multiple triggers, and will consequently pass by more than one alley. If an event is selected by at least one alley, it is processed by the inclusive triggers, where specific resonances are reconstructed and selected, and the exclusive triggers, which aim to fully reconstruct B-hadron final states.

Figure **??** shows the flow-diagram within one alley. At the start of an alley it is checked if the Level-0 trigger was based on either the muon-system, the HCAL or the ECAL. The first aim of the alleys is to confirm the Level-0 objects with better resolution by matching them to

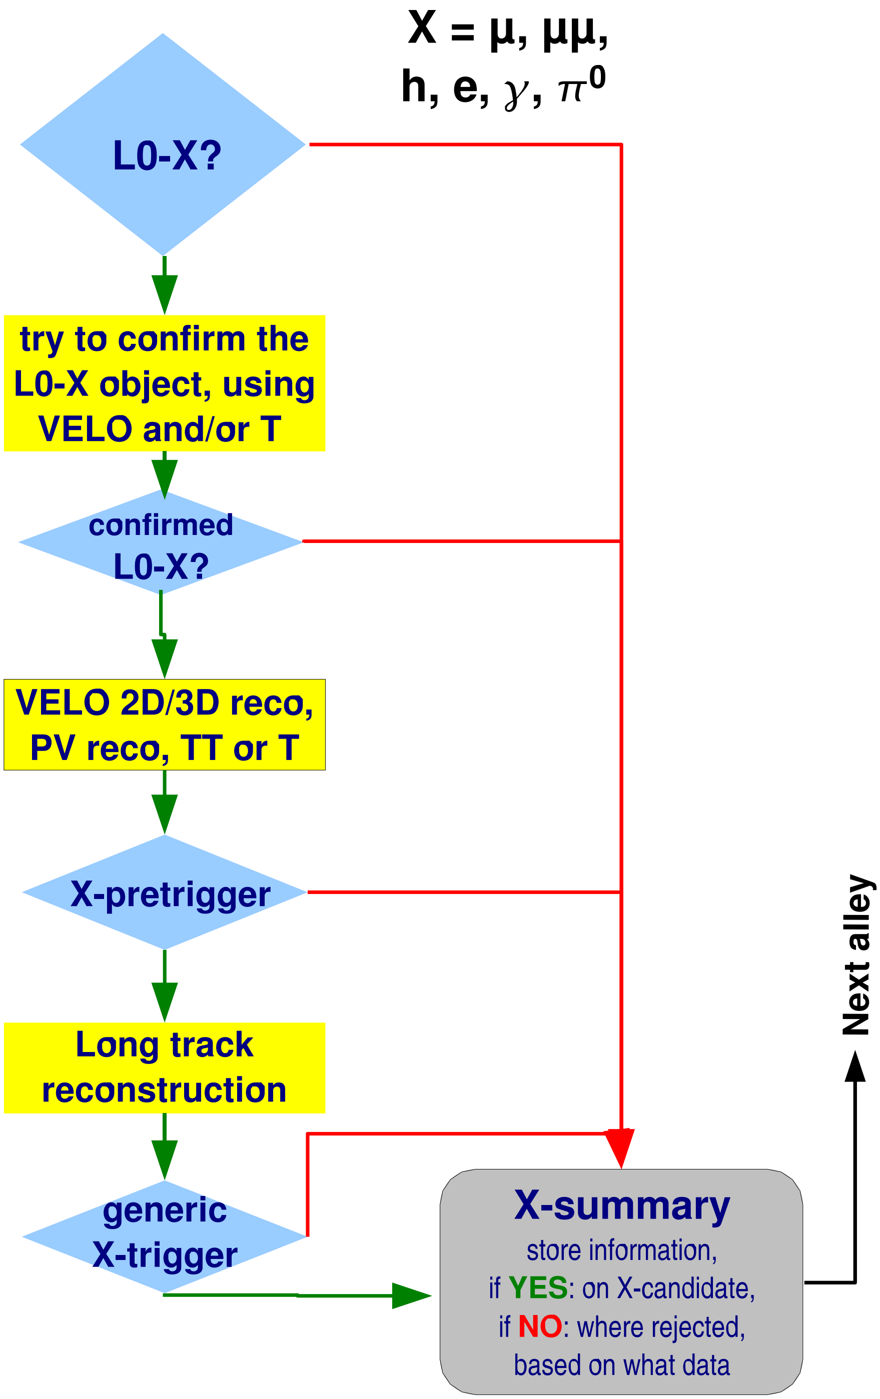

Figure 97. Flow-diagram within a HLT-alley. Squares represent reconstruction algorithms, while diamonds indicate where the trigger decisions are taken. T stands for the T1-T3 tracking stations, and TT is the Trigger Tracker

at least one tracking sub-detector, i.e. the VELO and/or T1-T3 tracking station. If the Level-0 object is confirmed, additional candidate B-decay tracks are reconstructed. B-mesons with their decay products in the LHCb acceptance move predominantly forward along the beam-line, which implies that the projection of the impact parameter in the plane defined by the beam-line and the track is large, while in the plane perpendicular to the beam it is almost indistinguishable with respect to primary tracks. This is exploited by reconstructing so-called 2D tracks using only the VELO R-sensors. These 2D tracks are sufficient to measure the position of the primary vertex with a precision of $\mathrm{RMS}^{\mathrm{PV}}_{\mathrm{x,y}} = 25\ \mu$m and $\mathrm{RMS}^{\mathrm{PV}}_{\mathrm{z}} = 60\ \mu$m, since the strips at constant radius are segmented in 45° ϕ-slices.

The Level-0 trigger is designed to only reconstruct muons with a relative large p_{T}, hence in the $\mu-$alley a search is made in stations M2-M5 to find additional muons, which are subsequently matched to their 2D tracks. These 2D muon candidates are combined with the ϕ-sensor clusters to form 3D tracks with a momentum resolution of $\mathrm{d}p/p \approx 6\%$ for correct matches, and their precise momentum is determined using the T1-T3 tracking stations.

In the other alleys the 2D tracks with a significant impact parameter are combined with the ϕ-sensor clusters to form 3D tracks. Unlike the muon-candidates no estimate is available for their momentum, and hence a search in the T1-T3 stations would be too time consuming. Their momentum is determined with a precision of $\mathrm{d}p/p = 20-40\%$ by combining 3D tracks with hits in TT, and using the fringe field of the magnet.

Based on these 3D tracks a pre-trigger decision is taken requiring a track, or a combination of tracks, with enough p_{T} and significant impact parameter. Each pre-trigger should contain at least one confirmed Level-0 object. The combined rate after the pre-trigger is expected to be around 30 kHz. This rate is sufficiently low to allow all candidate 3D-TT tracks to obtain a precision on their momentum of $\mathrm{d}p/p \approx 1\%$ using the T1-T3 stations. These so-called long tracks are then used to define the generic trigger of each alley, which like the pre-trigger is based on combining a confirmed Level-0 object with up to one long track with sufficient p_{T} and significant impact parameter.

Each alley produces summary information which is written to storage for accepted events. This summary contains the information of all tracks and vertexes which could have triggered the event. While the alleys are operating independently, care has been taken to avoid having to reconstruct the same track or primary vertex twice to avoid wasting precious CPU power.

The combined output rate of events accepted by the alleys is ~10 kHz, which is sufficiently low to allow to reconstruct long tracks out of the remaining 2D tracks. Prior to the final selection a set of tracks is selected with very loose cuts on their momentum and impact parameter. These tracks are used to form composite particles, like $\mathrm{K}^* \to \mathrm{K}^+\pi^-$, $\phi \to \mathrm{K}^+\mathrm{K}^-$, $\mathrm{D}^0 \to \mathrm{hh}$, $\mathrm{D_s} \to \mathrm{K}^+\mathrm{K}^-\pi^-$ and $\mathrm{J}/\psi \to \mu^+\mu^-$, which are subsequently used for all selections to avoid duplication in the creation of final states.

Previous trigger stages do not use cuts on invariant mass, nor precise pointing cuts to a primary vertex. The inclusive and exclusive selections aim to use these cuts to reduce the $\sim$ 10 kHz rate down to around 2 kHz, the rate at which the data is written to storage for further analysis. The exclusive triggers are more sensitive to tracking performance, while the inclusive triggers select partial B-decays to ϕX, J/ψX, D*X, $\mu^\pm$X, $\mu^\pm$hX and $\mu^+\mu^-$X to reduce the dependency on having to reconstruct all particles on-line. However, the exclusive selections of these channels produce a smaller rate, allowing for a more relaxed set of cuts. The final trigger is an or between the inclusive and exclusive selections.

7. Online System

The task of the Online system is to ensure the transfer of the data from the front-end electronics to permanent storage under known and controlled conditions. This includes the movement of the data themselves, but also the setting up of all the operational parameters of the experiment and the monitoring of these and the environmental parameters, such as temperatures or pressures. The online system also has to ensure that all detector channels are transferring their data at the same time, relative to the collisions of the particles in the accelerator. The LHCb Online system is described in detail in [170, 171, 172]

7.1 System Decomposition and Architecture

The LHCb Online system consists of three components

- Timing and Fast Control (TFC) system
- Experiment Control System (ECS)
- Data Acquisition (DAQ) system

Figure 124 shows the general architecture of the LHCb online system

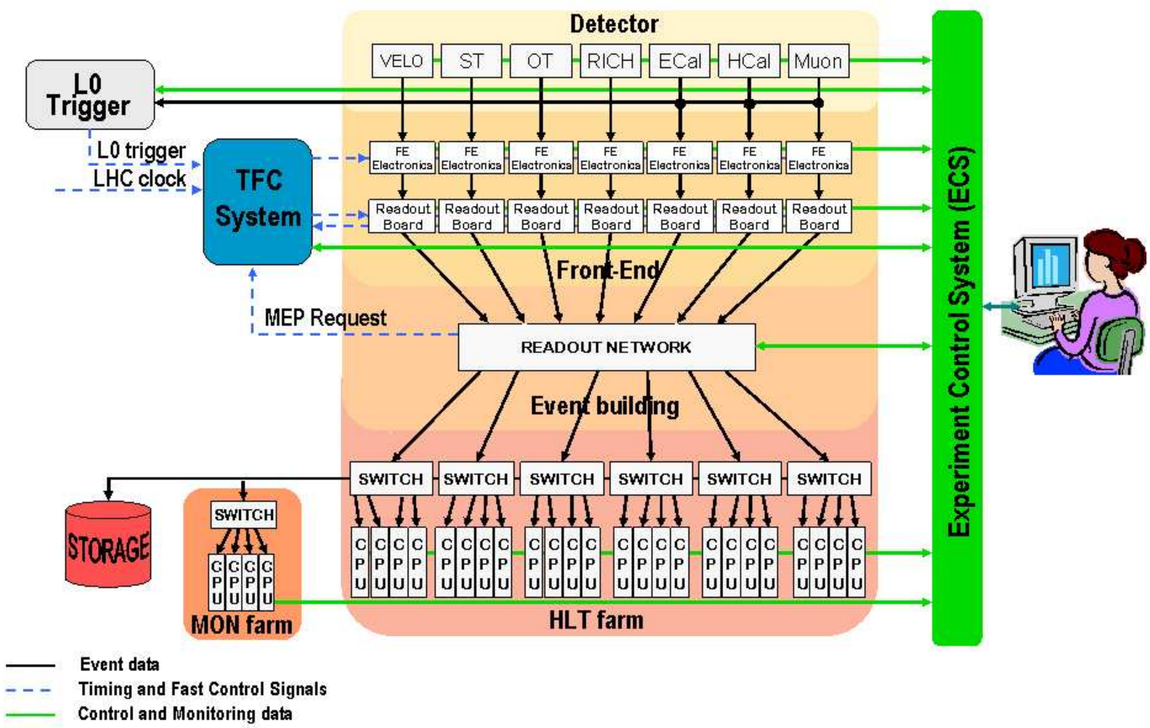

Figure 98. General architecture of the LHCb Online system with its three major components: Timing and Fast Controls, Data Acquisition and Experiment Control System

7.2 Data Acquisition System

The purpose of the Data Acquisition (DAQ) system is transporting the data belonging to one bunch crossing, one trigger, from the detector front-end electronics to permanent storage. The architecture of the DAQ system is depicted in Figure **??**. The design principles for the DAQ system were

- Simplicity
  simple protocols and small number of components with simple functionalities

- Scalability
  ability to react to changed parameters of the system, such as event sizes or trigger rates or CPU needs to execute the trigger algorithms.

- Only point-to-point links
  components are connected through point-to-point links only. No buses are used (outside monolithical boards). This leads to a more robust system.

- Use of COTS and wherever possible commodity components and protocols

With these principles it is assured that a reliable and robust system is built that also can cope with changed requirements without changes in architecture.

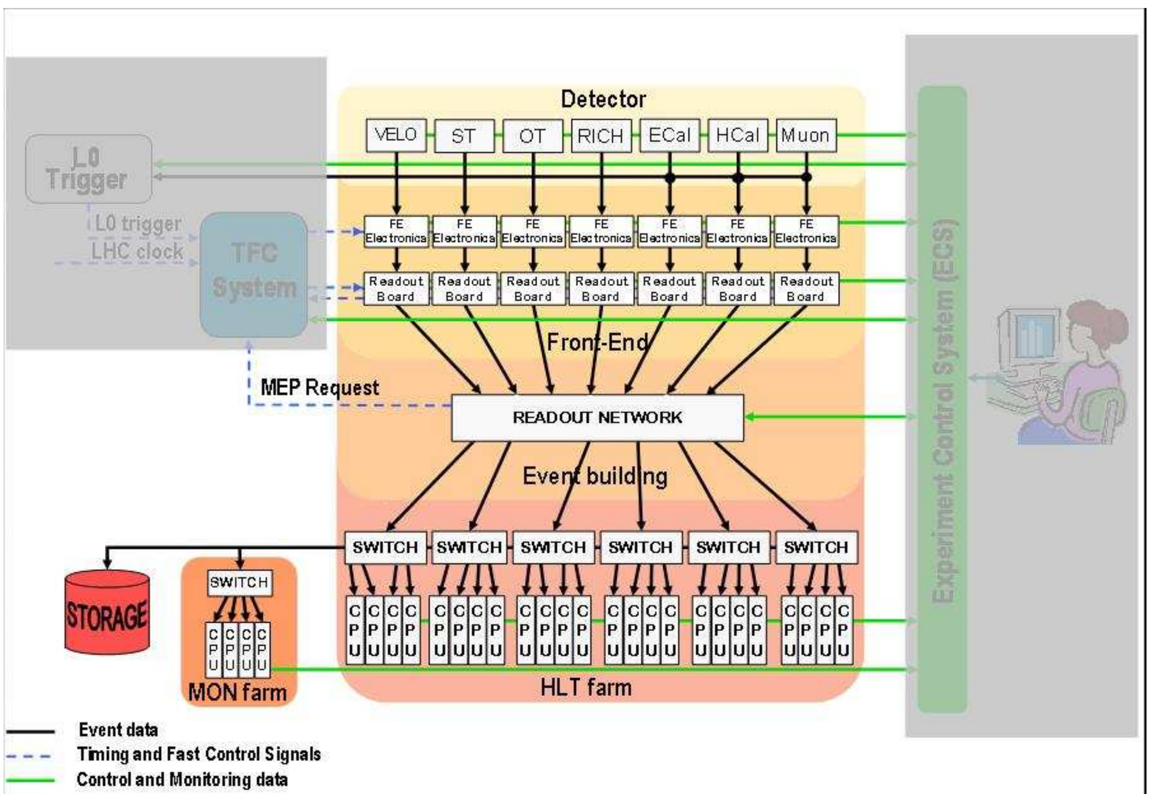

Figure 99. DAQ Architecture

Data from the on/near-detector electronics (front-end electronics) are collected in an LHCb-wide standardized electronics board (Tell11)[19]. Upon a positive trigger decision data are digitized in the Tell1 board if this is not already done in the front-end electronics. Subsequently, zero-suppression algorithms are executed in FPGAs and the zero-suppressed data are buffered in memory to be shipped further into the DAQ system towards a large CPU farm via a big Ethernet switch. The Tell1 board is described in detail in [173].

In the CPU farm an algorithm (High-Level Trigger, HLT) is executed that selects interesting interactions and subsequently, upon a positive decision, the data are sent to permanent storage. The HLT is expected to reduce the over-all rate from the original trigger rate of 1ăMHz to 2ăkHz, hence by a factor of 500. The storage system is expected to have a capacity of 40ăTB, which should offer sufficient buffer space to cope for situations where the transfer to permanent storage at Cern should be interrupted. As link technology for connecting the different components of the system Gigabit-Ethernet has been chosen. This mainly because of its wide, almost monopoly-like, acceptance in the LAN market and its low price. It also offers a very wide range of speed from 10 Mb/s to 10 Gb/s and the availability of very big switches (>1200 ports per chassis) are a big asset.

Consequently the Tell1-board offers 4 GbEthernet ports as output stage. (Some of) these will be fed into a large switching network providing the connectivity between the Tell1Šs and the individual Farm nodes. To overcome the significant overhead per frame of Ethernet the concept of Multi-Event Packets has been devised, in which the data of several triggers (10) are collected in one IP packet and transferred subsequently through the network. The size of the CPU farm running the HLT trigger algorithms is determined, of course, by the execution time it takes in average of the HLT algorithm per event but also the maximum bandwidth into an individual processing node, i.e. if the execution time should be very low, the input bandwidth might constitute the limiting factor and the number of ŚboxesŠ will have to be increased. The HLT algorithms are executed on a sizeable farm of CPUs. It's is expected to consist of 1000-2000 1U servers containing CPUs with multi-core technologies. The large number of CPUs are organized in 50 sub-farms of 20-40 CPUs each. One sub-farm is a functional unit and there is no cross-communication between sub-farms. This guarantees scalability since sowhere are large numbers exposed.

The quality of the acquired data is checked in a separate monitoring farm that will receive events accepted by the HLT and will house user-defined algorithms to determine e.g. the efficiencies of the detector channel or for example the mass resolution of the detector.

7.3 Timing and Fast Control

The TFC system drives all the stages of the data readout of the LHCb detector between the Front-End electronics and the online processing farm by distributing the beam-synchronous clock, the hardware trigger (L0), synchronous resets and fast control commands. The system is a combination of electronic components common to all LHC experiments and LHCb custom electronics. The TFC architecture shown in igure 126 can be described in terms of three main ingredients, the TFC distribution network, the trigger throttle network, and the TFC master (Readout Supervisor).

[19]The The RICH detectors use a local flavor of the Tell1 which is functionally identical to the Tell1.

The TFC optical distribution network with transmitters and receivers is based on the LHC-wide TTC (Trigger, Timing and Control) system developed at CERN. In addition to transmitting the beam synchronous clock, the protocol features a low-latency trigger channel and a second channel with framed user data used to encode the control commands. A switch has been developed and introduced in the distribution network to allow a dynamic partitioning of the LHCb detector to support independent and concurrent sub-detector activities such as commissioning, calibration and testing.

The optical throttle network is used to transmit back-pressure, that is a trigger inhibit, from the asynchronous parts of the readout system to the Readout Supervisor in case of data congestion. The network incorporates a Throttle Switch as a consequence of the requirement to allow partitioning of the readout system, and modules which perform locally an ŞorŤ of the throttle signals of each sub-system.

The heart of the system is the Readout Supervisor. It implements the interface between the LHCb trigger system and the readout chain. The Readout Supervisor synchronizes the trigger decisions and the beam-synchronous commands to the LHC clock and orbit signal, provided by the LHC accelerator. It is also capable of producing a variety of auto-triggers for sub-detector calibrations and tests, and performs the trigger control as a function of the load on the readout system. In order to load balance dynamically the online processing farm, the Readout Supervisor also selects and broadcasts the destination for the next set of events to the Readout Boards based on a request scheme in which the farm nodes send data requests directly to the Readout Supervisor.

For each trigger the Readout Supervisor transmits a data bank over the readout network which is appended to the event data and which contains the identifier of the event, the time and the source of the trigger.

7.4 Experiment Control System

The Experiment Controls System (ECS) ensures the control and monitoring of the operational state of the entire LHCb detector. This encompasses the traditional Detector Control, such as high and low voltages, temperatures, gas flows, or pressures, but also the control and monitoring of the Trigger system and the TFC and DAQ systems. The hardware components of the ECS are somewhat divers, mainly as a consequence of inhomogeneity of the equipment to be controlled. This ranges from standard crates and power supplies to individual electronics boards. In LHCb a large effort was made to minimize the number of different types of interfaces and connecting busses. The field busses have been restricted to

- SPECS (Serial Protocol for ECS)
- CAN (Controller Area Network)
- (fast)Ethernet

The first two (SPECS and CAN) are mainly used for equipment residing in the radiation area close to the detector. The associated interfaces tolerate modest levels of radioactivity but are not radiation hard. Ethernet is only used in the radiation free areas, such as the electronics barracks or on the surface. Ethernet is used to control individual PCs, but also the individual electronics

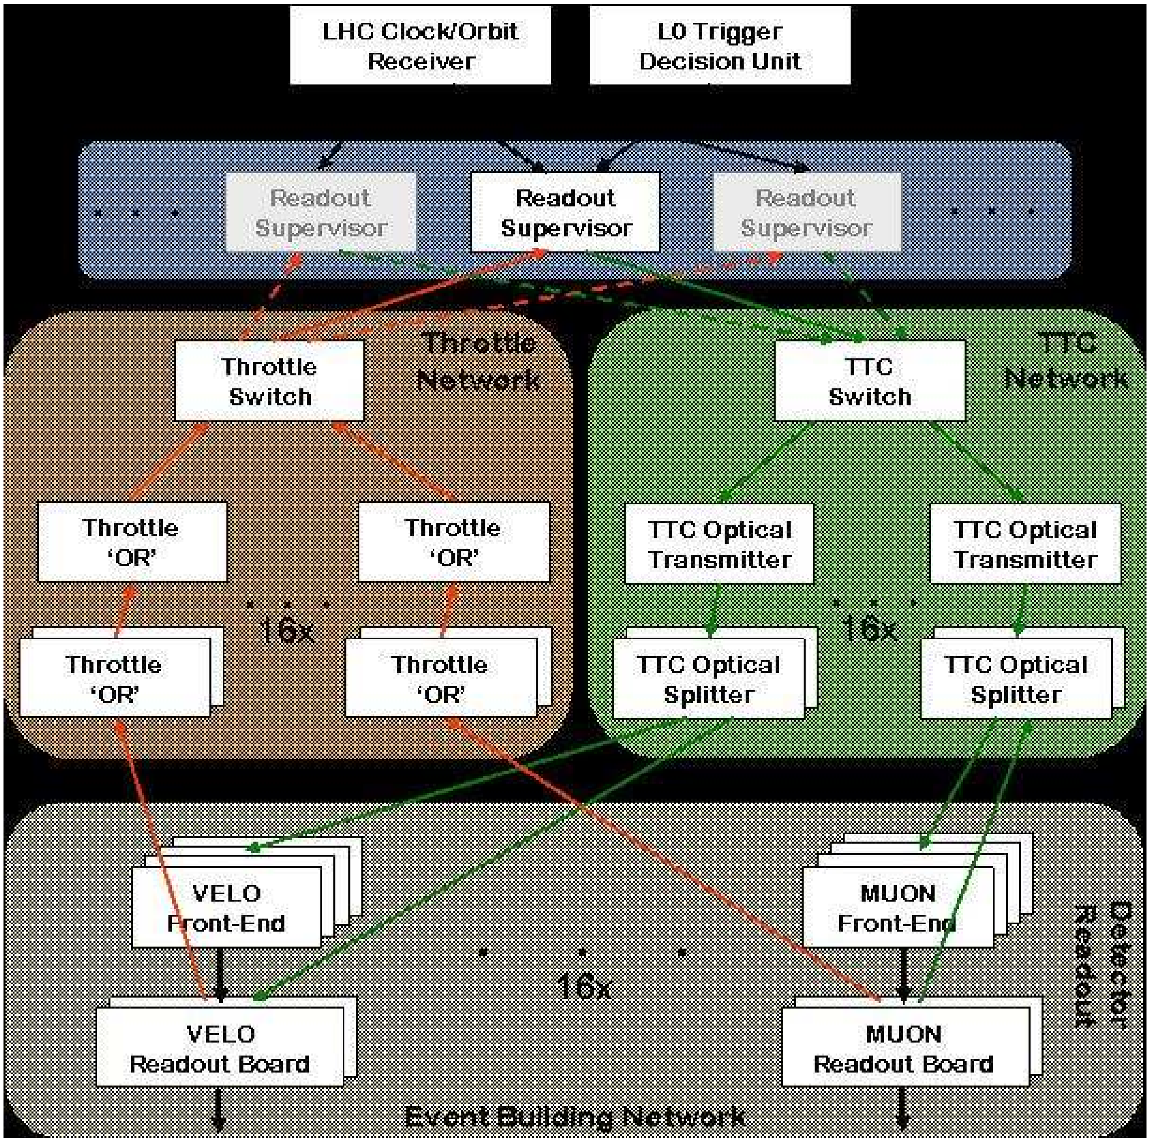

Figure 100. Schematic diagram of the TFC architecture

boards for the readout through Credit-Card sized PCs mounted directly on each board. This choice relieves us from having to use ŚCrate ControllersŠ. Thus we can use normal PCs over their standard Ethernet interfaces for controlling the readout electronics.

The ECS software is based on PVSS II, a commercial SCADA (Supervisory Control And Data Acquisition) system. This toolkit provides the necessary infrastructure for building the ECS system, such as a configuration database and communication between distributed components, graphical libraries to build operations panels, or an alarm system as well as components, such as OPC clients. Based on PVSS a hierarchical and distributed system was designed as depicted in Figure **??**.

Device Units, in Figure **??**, denote low-level access components that typically would commu-

Figure 101. ECS architecture

nicate directly to the hardware. They would, in general, not implement state machines or behavior. Examples of Device Units would be an OPC server for a power supply, but also a software process, such as the High-Level Trigger process would be controlled with a Device Unit.

Contrary to Device Units, Control Units would implement high-level states and transitions and also local logic to attempt recovering from errors of the subordinate Device Units. Typical examples of Control Units are a HV subsystem, or the component that would control the ensemble of the crates of a sub-detector or an entire sub-farm of the Event Filter Farm. Control Units are controlled by other Control Units, allowing building up a hierarchy of arbitrary depth. State sequencing in the ECS system is achieved using a Finite State Machine (FSM) package that allows creating complex logic such as implementing elaborate sequencing or automatic error recovery.

The (distributed) components of the ECS system are connected with a quite large Ethernet network consisting of several hundred Gigabit and Fast Ethernet links.

8. Performances

8.1 Introduction

The data used to study the performance of the LHCb detector had been obtained from detailed Monte Carlo simulations, which produces a raw data format identical to real data. The proton-proton collisions are simulated with the PYTHIA program [174]. The generated particles are tracked through the detector material and surrounding environment using the GEANT 4 package [175]. The geometry and material of the LHCb detector are described in detail. The detector responses, resolution, noise, cross-talk, etc. had been tuned to comply with testbeam results. Details about this can be found in the Reoptimization TDR and Computing TDR.

8.2 Track Reconstruction

In the track reconstruction software the registered hits of the VELO, the TT, the IT and the OT detectors are combined to form particle trajectories from the VELO to the calorimeters. The program aims to find all tracks in the event which leave sufficient detector hits, not only possible B-decay products. After fitting the reconstructed trajectory a track is represented by state vectors (x, y, dx/dz, dy/dz, Q/p) which are specified at given z-positions in the experiment.

Depending on their generated trajectories inside the spectrometer the following classes of tracks are defined, illustrated in Fig. 128:

1. **Long tracks:** traverse the full tracking set-up from the VELO to the T stations. They have the most precise momentum determination and therefore are the most important set of tracks for B-decay reconstruction.

2. **Upstream tracks:** traverse only the VELO and TT stations. They are in general lower momentum tracks that do not traverse the magnet. However, they pass through the RICH 1

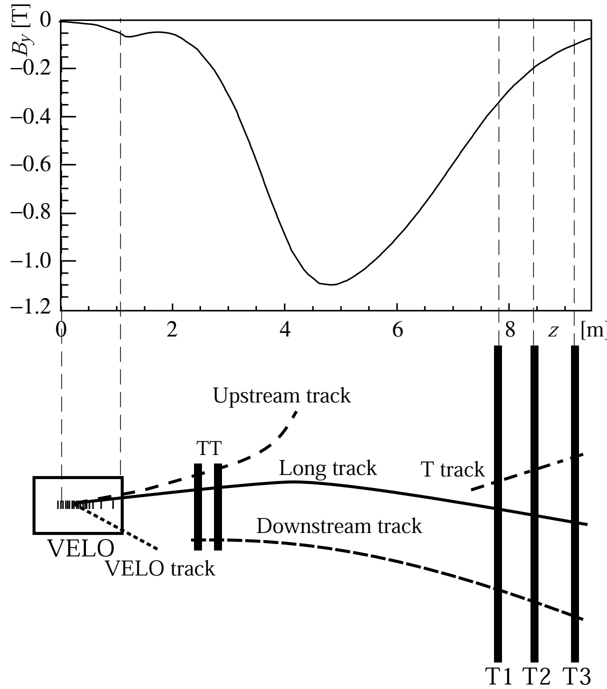

Figure 102. A schematic illustration of the various track types: long, upstream, downstream, VELO and T tracks. For reference the main B-field component (B_y) is plotted above as a function of the z coordinate.

detector and may generate Cherenkov photons. They are therefore used to understand backgrounds in the particle-identification algorithm of the RICH. They may also be used for B-decay reconstruction or tagging, although their momentum resolution is rather poor.

3. **Downstream tracks:** traverse only the TT and T stations. The most relevant cases are the decay products of K^0_S and Λ that decay outside the VELO acceptance.

4. **VELO tracks:** are measured in the VELO only and are typically large angle or backward tracks, useful for the primary vertex reconstruction.

5. **T tracks:** are only measured in the T stations. They are typically produced in secondary interactions, but are useful for the global pattern recognition in RICH 2.

The track reconstruction starts with a search for track “seeds”, the initial track candidates [176], in the VELO region and the T stations where the magnetic field is low After tracks have been found their trajectories are refitted with a Kalman filter fit [177] which takes into account multiple scattering and in addition corrects for dE/dx energy loss. The quality of the reconstructed tracks is monitored by the χ^2 of the fit and the “pull” distribution of the track parameters.

8.3 Performance

The pattern recognition performance is evaluated in terms of efficiencies and ghost rates. The efficiencies are normalized to the different reconstructible track samples. To be considered reconstructible, the requirements for each track type are as follows:

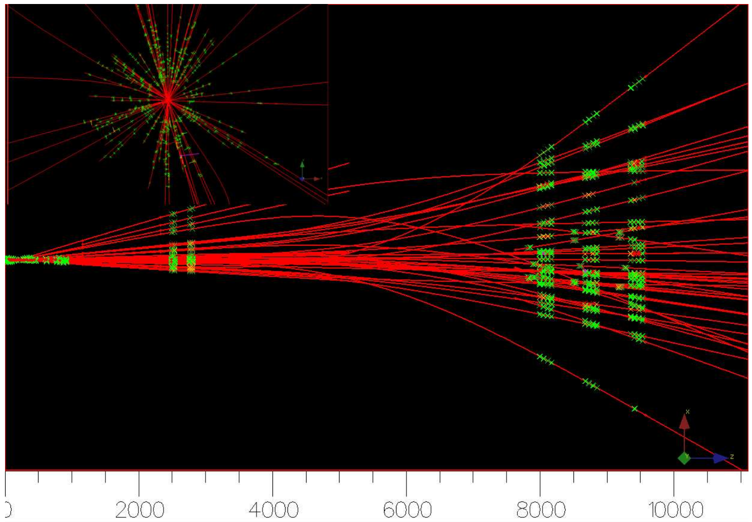

Figure 103. Display of the reconstructed tracks and assigned hits in an event. The insert shows a zoom into the VELO region.

- for VELO tracks the particle must give at least three r and three ϕ hits;
- for T tracks the particle must give at least 1 x and 1 stereo hit in each station T1–T3;
- for long tracks the particle must be reconstructible as a VELO and T track;
- for upstream tracks the particle must be reconstructible as a VELO track and give at least 3 hits in TT;
- for downstream tracks the particle must be reconstructible as a T track and give at least 3 hits in TT.

To be considered as "successfully reconstructed" a track must have at least 70% of its associated hits originating from a single Monte Carlo particle. The reconstruction efficiency is defined as the fraction of reconstructible particles that are successfully reconstructed, and the ghost rate is defined as the fraction of found tracks that are not matched to a true Monte Carlo particle.

The results quoted in this section are obtained from a sample of $B^0 \to J/\psi K^0_S$ events. Particle tracks originating from either decays of the signal B^0 or the other b-hadron in the event are referred to simply as B tracks in the text.

The average number of successfully reconstructed tracks in $b\bar{b}$ events is about 72, which are distributed among the track types as follows: 26 long tracks, 11 upstream tracks, 4 downstream tracks, 26 VELO tracks and 5 T tracks. The track finding performance is summarized for the most important cases: the long tracks, the low momentum (upstream) tracks and K^0_S decay (downstream) tracks. An example of a reconstructed event is displayed in Fig. **??**.

Long tracks The efficiency of long track reconstruction is plotted in Fig. 130 as function of the track momentum. For tracks with momentum higher than 10 GeV/c the average efficiency is 94%.

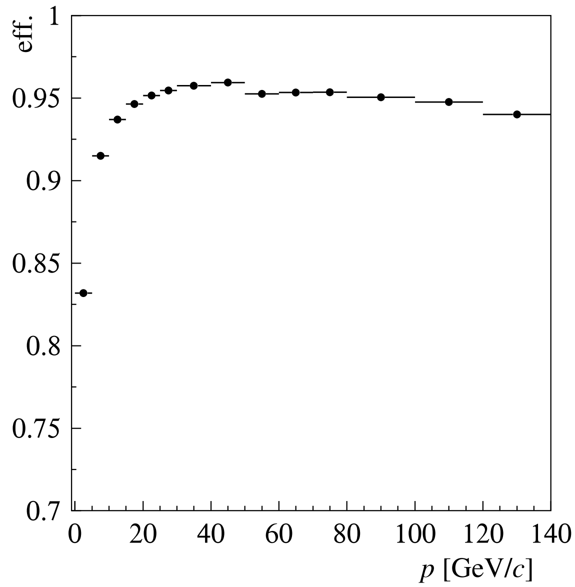

Figure 104. Efficiency as a function of the momentum of the generated particle for the long track finding.

For final state particles of specific B decays even higher efficiencies are observed (95–96%). The corresponding average ghost rate is about 9%.

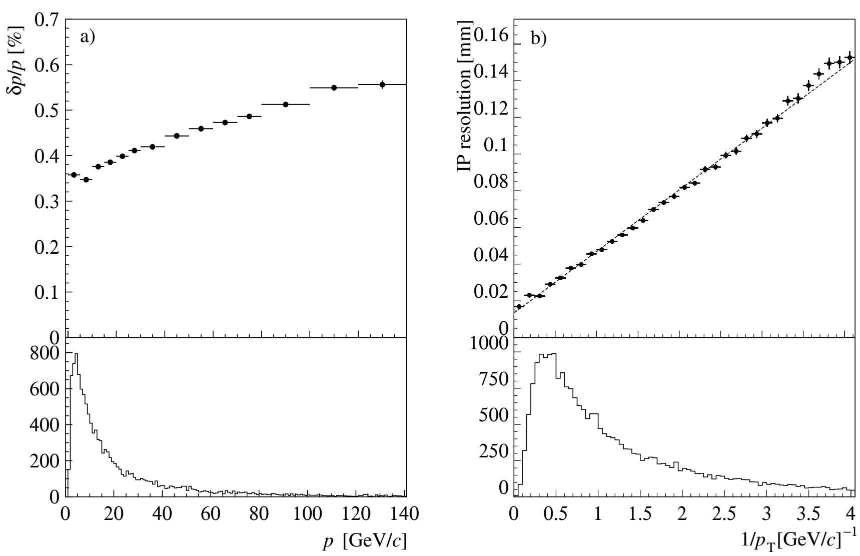

Figure 105. Resolution on the reconstructed track parameters at the production vertex of the track: (a) momentum resolution as a function of track momentum, (b) impact parameter resolution as a function of $1/p_{\mathrm{T}}$. For comparison, the momentum and transverse-momentum spectra of B-decay particles are shown in the lower part of the plots.

The momentum and impact parameter resolutions of the reconstructed track parameters are shown in Fig. **??**. The momentum resolution is plotted as a function of the track momentum and is seen to be increasing from $\delta p/p = 0.35\%$ for low momentum tracks to $\delta p/p = 0.55\%$ for tracks at the high end of the spectrum. In the same figure the momentum spectrum for B-decay tracks is also illustrated. The impact parameter resolution is plotted as function of $1/p_{\mathrm{T}}$ of the track. The linear dependence can be parametrized as $\sigma_{\mathrm{IP}} = 14\,\mu\mathrm{m} + 35\,\mu\mathrm{m}/p_{\mathrm{T}}$ with p_{T} in GeV/c. For comparison the momentum and $1/p_{\mathrm{T}}$ spectra of B-decay particles in the detector acceptance are plotted in the same figure.

Upstream and Downstream tracks Since most tracks with $p < 1$ GeV$/c$ are below threshold in RICH 1 the relevant efficiency for the upstream track finding is approximately 75% with a corresponding ghost rate of 15%. The momentum resolution due to the small fraction of the total B-field integral is only $\delta p/p \sim 15\%$.

The efficiency for finding downstream tracks above 5 GeV/c is about 80%. Since the downstream tracks traverse most of the magnetic field, the momentum resolution is relatively good with an average of $\delta p/p = 0.43\%$ for pions originating from $\mathrm{K_S^0}$ decays in $\mathrm{B^0} \to \mathrm{J}/\psi \mathrm{K_S^0}$ events. In order to maintain high efficiency the reconstruction allows for typically two or three track candidates in TT to be linked to a single track seed in the T stations. The ghost tracks introduced by this procedure are easily eliminated at a later stage when the pion track pairs are combined into $\mathrm{K_S^0}$ decays, as is shown in Chapter 8.4.

8.4 Particle Identification

The information from the various detectors, the two RICH detectors, the Calorimeter system and the Muon Detector, is combined for optimal identification of charged particle types $(\mathrm{e}, \mu, \pi, \mathrm{K}, \mathrm{p})$. Neutral electromagnetic particles (γ, π^0) are identified using the Calorimeter system, where the $\pi^0 \to \gamma\gamma$ may be resolved as two separate photons, or as a merged cluster. Finally $\mathrm{K_S^0}$ are reconstructed from their decay $\mathrm{K_S^0} \to \pi^+\pi^-$.

8.4.1 Kaon identification

Particle identification with the RICH system is performed as follows. The pattern of hit pixels observed in the RICH photodetectors is compared to the pattern that would be expected under a given set of mass hypotheses for the reconstructed tracks passing through the detectors, using the knowledge of the RICH optics. A likelihood is determined from this comparison, and then the track mass-hypotheses are varied so as to maximise the likelihood [178].

The kaon efficiency and misidentification rate is shown in Fig. 132. The effect of crossing the thresholds for Cherenkov light production in the three radiators is clearly seen at $p \sim 2$, 9 and 16 GeV$/c$. The average efficiency for kaon identification between 2 and 100 GeV$/c$ is 88%. The average pion misidentification rate, $\varepsilon(\pi \to \mathrm{K})$, between 2 and 100 GeV$/c$ is 3%.

8.4.2 Muon identification

Muons are identified by extrapolating well reconstructed tracks with $p > 3$ GeV$/c$ into the Muon stations.Muon Detector hits are searched within fields of interest (FOI) around the extrapolation point in each station, parameterized as function of momenta for each station and region. A track is considered a muon candidate when a minimum number of stations have hits in their corresponding FOI's. This number was optimized to maintain high efficiency [179].

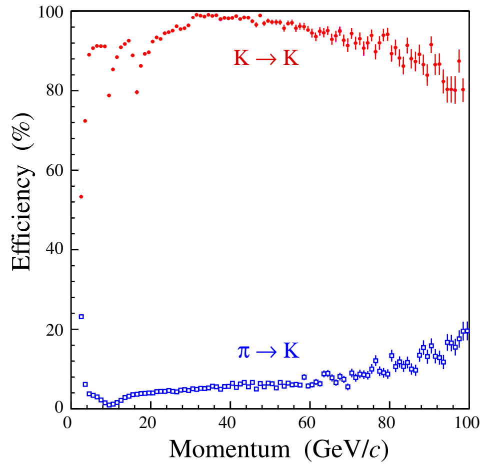

Figure 106. Kaon identification efficiency (solid points) and pion misidentification rate (open points) as a function of momentum. For this figure tracks are selected as kaons if their maximum-likelihood hypothesis kaon or heavier, and as pions if it is pion or lighter.

Using a sample of $\mathrm{B}^0 \to \mathrm{J}/\psi \mathrm{K}^0_\mathrm{S}$ the performance was measured to be $\varepsilon(\mu) = 94.3 \pm 0.3\%$ and $\varepsilon(\pi \to \mu) = 2.9 \pm 0.1\%$. The efficiency is a flat function of the momentum, above $10\,\mathrm{GeV}/c$, as shown in Fig. **??**.

Discriminating variables to help further improve the muon selection purity are constructed from the comparison of slopes in the muon system and the main tracker, and from the average track-hit distance of all hits in FOI's associated to the track. For each track the difference in log-likelihood between the muon hypothesis and pion hypothesis is determined, and summed with the values from the RICH and Calorimeter systems (if available). By doing this the pion misidentification rate can be reduced to 1%, whilst maintaining a muon efficiency of 93%, for muons above $3\,\mathrm{GeV}/c$.

The high purity that can be achieved with such cuts is illustrated in Fig. 134 (a), where the $\mu^+\mu^-$ mass plot is shown at the first step in the analysis of $\mathrm{B}^0_\mathrm{s} \to \mathrm{J}/\psi\phi$ events, taking all oppositely-charged pairs of tracks from signal events that pass the muon identification requirements. As can be seen, a clean J/ψ mass peak is reconstructed with a resolution of about $13\,\mathrm{MeV}/c^2$.

8.4.3 Electron identification

The electron identification [180] is mainly based on the balance of track momentum and energy

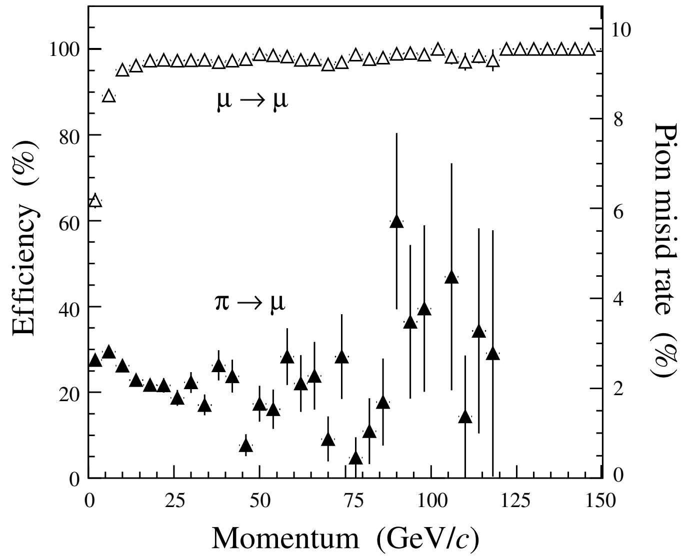

Figure 107. Muon efficiency (open triangles), and the pion misidentification rate (solid triangles, with scale on the right) as a function of track momentum.

of the charged cluster in the ECAL (Fig. **??**), and the matching between the corrected barycenter position of the cluster with the extrapolated track impact point.

A second estimator is related to the Bremsstrahlung photons emitted by electrons before the magnet. As there is little material within the magnet, such neutral clusters are expected in a well defined position given by the electron track extrapolation from before the magnet, as illustrated in Fig. 136.

Further improvement in electron identification is made by using the track energy deposition in the Preshower detector and the deposition of the energy along the extrapolated particle trajectory in the hadronic calorimeter, HCAL.

Finally, the Calorimeter information is combined with the RICH and Muon detectors.

To illustrate the performance of electron identification, the J/ψ mass plot is shown as the open points in Fig. 134 (b). The signal is fit with a function including a radiative tail, to account for the imperfect correction of Bremsstrahlung. The background tracks are dominantly of low p_T, and can be efficiently rejected by applying the requirement $p_T > 0.5\,\mathrm{GeV}/c$ for the electron candidates, as shown by the solid points in Fig. 134 (b).

The average efficiency to identify electrons in the calorimeter acceptance from $J/\psi \to e^+e^-$ decays in $B^0 \to J/\psi K^0_S$ events is 95%, for a pion misidentification rate of 0.7%, as shown in Fig. **??**.

8.4.4 Photon identification

Photons are reconstructed and identified with the electromagnetic calorimeter, ECAL, as neutral clusters [**?**, 181]. The reconstructed tracks are extrapolated to the ECAL face and a cluster-to-track position matching estimator, χ^2_γ, is calculated. The minimal value of the χ^2_γ estimator for each cluster is shown in Fig. 138. The clusters due to charged tracks are clearly identified as a peak at a small value of χ^2_γ. Clusters with $\chi^2_\gamma > 4$ are selected as photon candidates.

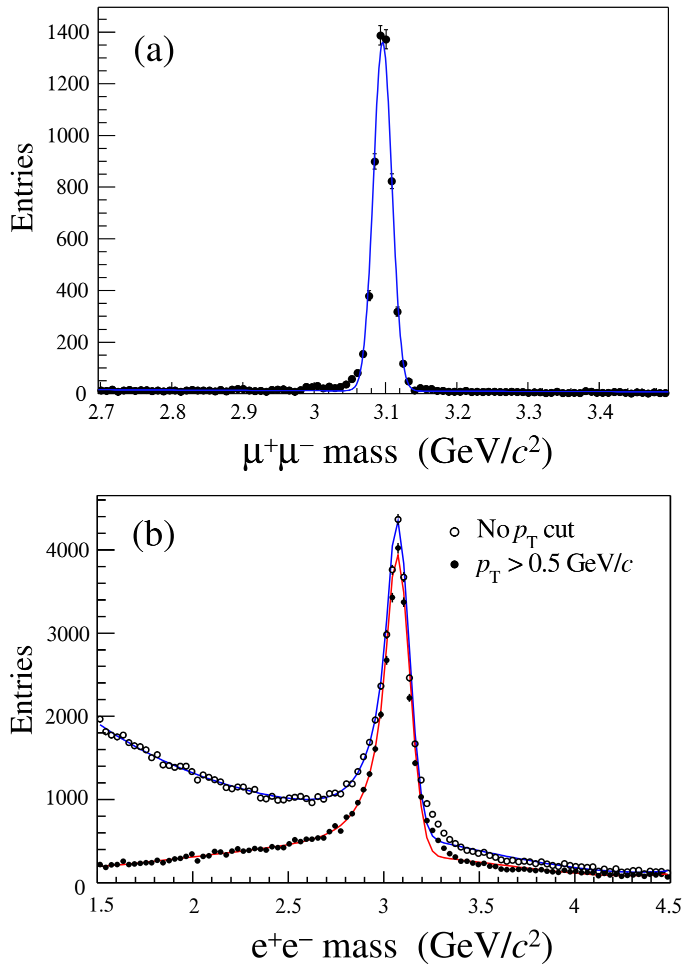

Figure 108. Invariant mass plots for the reconstruction of $J/\psi \to \ell^+\ell^-$ decays in $B_s^0 \to J/\psi\phi$ signal events: (a) for $\ell = \mu$, (b) for $\ell = e$, where the open points are before any p_T cut, and the solid points are after requiring $p_T > 0.5\,\mathrm{GeV}/c$ for the $e^{\pm}$ candidates.

The identification of photons converted in the passive material of the apparatus after the magnet, e.g. RICH 2 or M1, is based on whether there is a hit in the SPD cell that lies in front of central cell of the ECAL cluster. Seventy percent of reconstructed photons from $B^0 \to K^*\gamma$ decays are selected as non-converted photons, while the remaining photons are identified as converted photons, with correct assignment fractions of 90% and 79% respectively.

A cut on the energy deposition in the Preshower detector can improve the purity of selected samples both for converted and non-converted photons [182].

8.4.5 π^0 reconstruction

Neutral pion reconstruction has focussed on the $B^0 \to \pi^+\pi^-\pi^0$ decay channel, for which the mean transverse momentum of the π^0 is about $3\,\mathrm{GeV}/c$. Below this value the π^0 are mostly reconstructed as a resolved pair of well separated photons, while for higher p_T a large fraction of the pairs of photons coming from the decay of the π^0 cannot be resolved as a pair of clusters within the ECAL

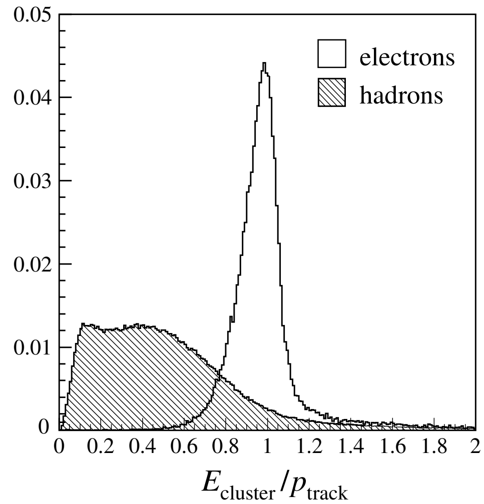

Figure 109. The ratio of uncorrected energy of the charged cluster in ECAL to the momentum of reconstructed tracks for electrons (open histogram) and hadrons (shaded histogram).

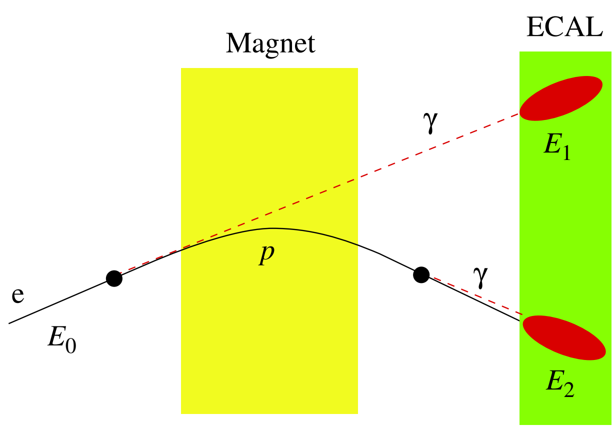

Figure 110. Schematic illustration of Bremsstrahlung correction. An electron may radiate photons when passing through material before or after the magnet: in the first case, a well defined cluster is seen in the ECAL, with energy E_1, whilst in the second case the Bremsstrahlung energy forms part of the electron cluster with energy E_2; for electron identification $E_2 = p$, the momentum measured in the spectrometer, while the energy of the electron at the origin, $E_0 = E_1 + E_2$.

granularity. About 30% of the reconstructible π^0 from $B^0 \to \pi^+\pi^-\pi^0$ lead to a single cluster, referred to as a merged π^0.

Figure **??** shows the mass distributions obtained in the cases where both photon candidates with $p_T > 200\,\mathrm{MeV}/c$ reached the calorimeter (a), or one converted before the calorimeter (b), according to the SPD signal. The distributions are fitted with the sum of a Gaussian and a polynomial function, giving a resolution for the π^0 mass of $\sigma = 10\,\mathrm{MeV}/c^2$.

An algorithm has been developed to disentangle a potential pair of photons merged into a single cluster. The energy of each cell of the cluster is shared between two virtual sub-clusters

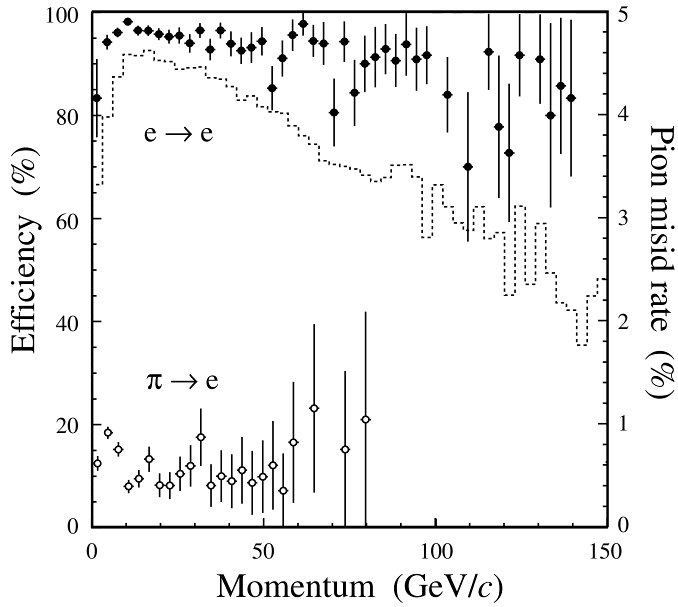

Figure 111. Efficiency for electron identification, after requiring $\Delta \ln \mathcal{L}_{e\pi} > 0$, as a function of momentum. The solid points are for electrons within the Calorimeter acceptance, the dotted histogram for all electrons from J/ψ decays. The pion misidentification rate is indicated by the open points, with scale on the right.

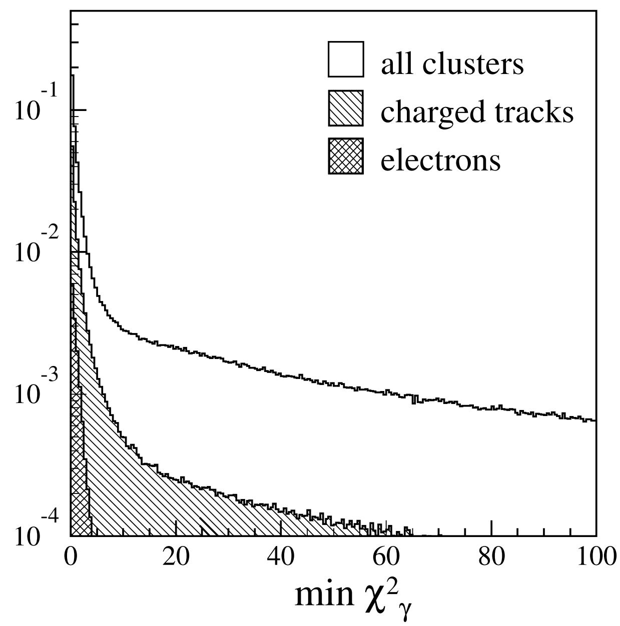

Figure 112. The minimal value of the χ^2_γ estimator from the track-cluster position matching procedure for ECAL clusters. The peak at small values of χ^2_γ corresponds to clusters due to charged particles, as indicated by the hatched distribution; the distribution for electrons is shown cross-hatched.

according to an iterative procedure based on the expected transverse shape of photon showers. Each of the two sub-clusters is then reconstructed as coming from a photon, as for isolated photons. The

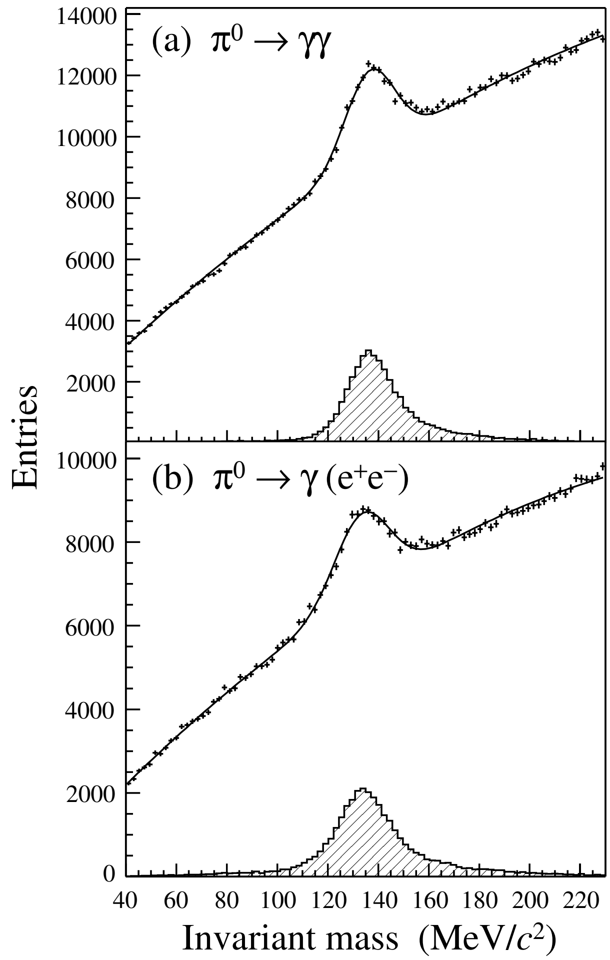

Figure 113. π^0 mass distributions, where (a) neither photon has converted, (b) one γ converted before the calorimeter. The separation converted/not-converted is obtained from the SPD information. The contributions of true π^0 are indicated by the shaded histograms.

reconstructed invariant mass obtained from all single clusters in $B^0 \to \pi^+\pi^-\pi^0$ events is shown in Fig. 140 (a). The same distribution for the clusters associated to the π^0 from B^0 decay is shown in (b). The distribution is fit with two Gaussians, to account for the broadening of the resolution due to photon conversion. A core resolution of about 15 MeV/c^2 is obtained. The reconstruction of the 4-momentum of merged π^0 is competitive with the resolved configuration, with an angular resolution of better than 1 mrad for merged π^0 above 20 GeV.

The reconstruction efficiency for π^0 that give photons inside the geometrical acceptance is summarized in Fig. **??** for the resolved and merged case.

8.4.6 K^0_S reconstruction

K^0_S are reconstructed through their decay to $\pi^+\pi^-$. For K^0_S from $B^0 \to J/\psi K^0_S$ decays, about 25% decay inside the active region of the VELO, 50% decay outside the active region of the VELO but upstream of TT, and the rest decay after TT (and will therefore be difficult to reconstruct).

The K^0_S that decay outside the active region of the VELO but before TT are reconstructed

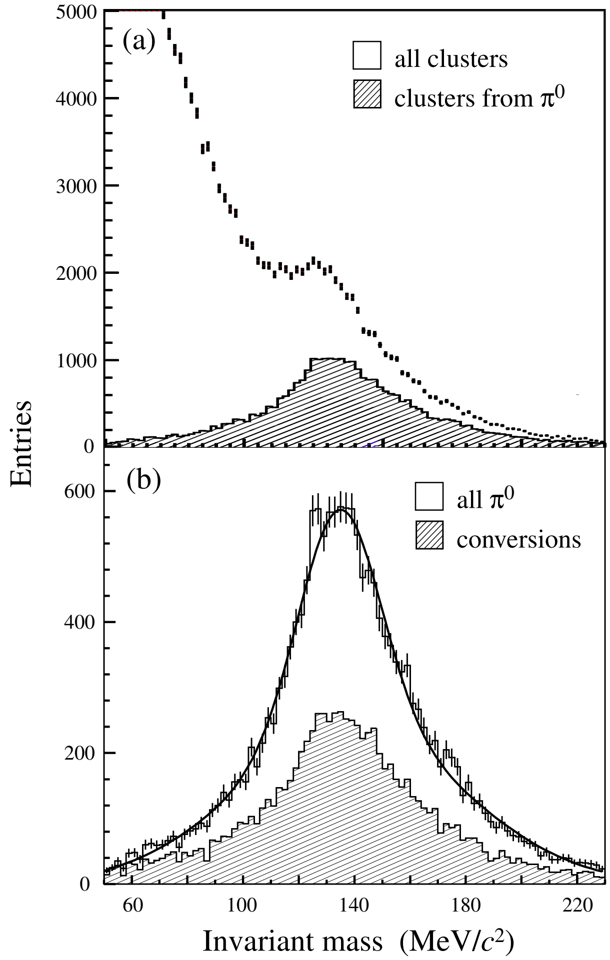

Figure 114. Invariant mass obtained with the merged π^0 algorithm (a) for all clusters in $B^0 \to \pi^+\pi^-\pi^0$ events, where the shaded histogram indicates the contribution from true π^0, and (b) for clusters associated to true π^0 from B^0 decays, where the shaded histogram indicates the contribution from pairs of photons with at least one conversion.

using pairs of oppositely-charged downstream tracks, which are found using the hits in the TT and T1–T3 tracking stations. The pions from K^0_S that decay within the VELO acceptance give either a long track or an upstream track, depending on whether they pass through the magnet to give hits in the downstream tracking stations.

The corresponding mass plots are shown in Fig. 142 (a), (b) and (c). As can be seen, there is some combinatorial background from other tracks in the signal events, particularly for the long-upstream category, but this background is removed by the additional requirements that are imposed when reconstructing a B meson.

8.5 Global performance

Given the reconstruction performances as described in the previous sections, the following global performances can be expected for the reconstruction of B-decays :

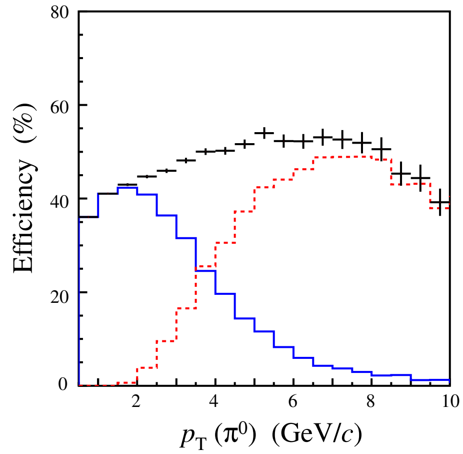

Figure 115. The reconstruction efficiency for π^0 that give photons inside the geometrical acceptance with $p_{\rm T} > 200\,\mathrm{MeV}/c$, versus the π^0 transverse momentum. The separate contributions from resolved and merged π^0 reconstruction are indicated by the solid and dashed histograms, respectively.

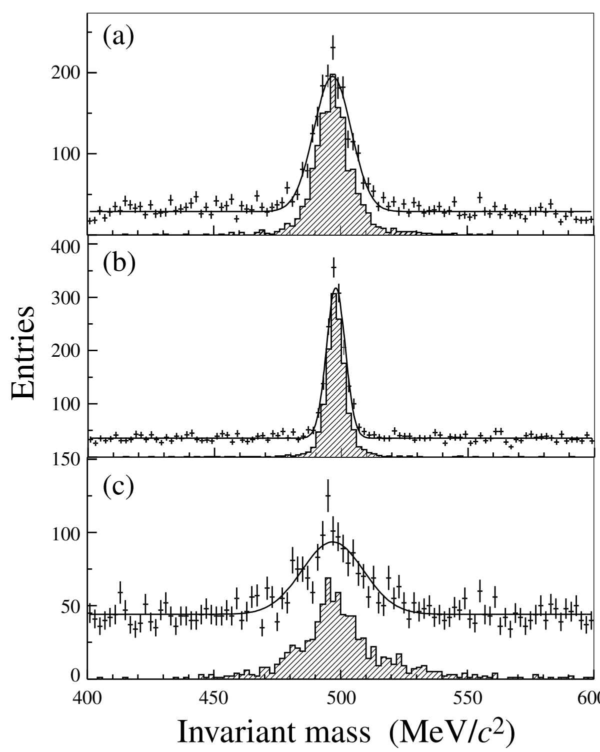

Figure 116. Reconstruction of $\mathrm{K_S^0} \to \pi^+\pi^-$. The $\pi^+\pi^-$ invariant mass is shown in $\mathrm{B^0} \to \mathrm{J}/\psi\mathrm{K_S^0}$ signal events, using different categories of tracks for the pion candidates: (a) downstream-downstream, (b) long-long, (c) long-upstream. Combinations coming from a $\mathrm{K_S^0}$ are indicated by the shaded histograms.

- primary vertex resolutions of 10 μm transverse to the beam axis and 60 μm along the beam axis;

- invariant mass resolutions typically in the range between 12 MeV/c and 25 MeV/c;
- proper lifetime resolutions of 40 fs.

8.6 Computing and Resources

This section describes the dataflow of the LHCb computing model for all stages in the processing of the real and simulated LHCb events. The roles of the various Tier centres are discussed and the distribution of the processing load and storage are outlined.

There are several phases in the processing of event data. The various stages normally follow each other in a sequential manner, but some stages may be repeated a number of times. The workflow reflects the present understanding of how to process the data. A schematic of the logical dataflow is show in Figure **??** and is described in more detail below.

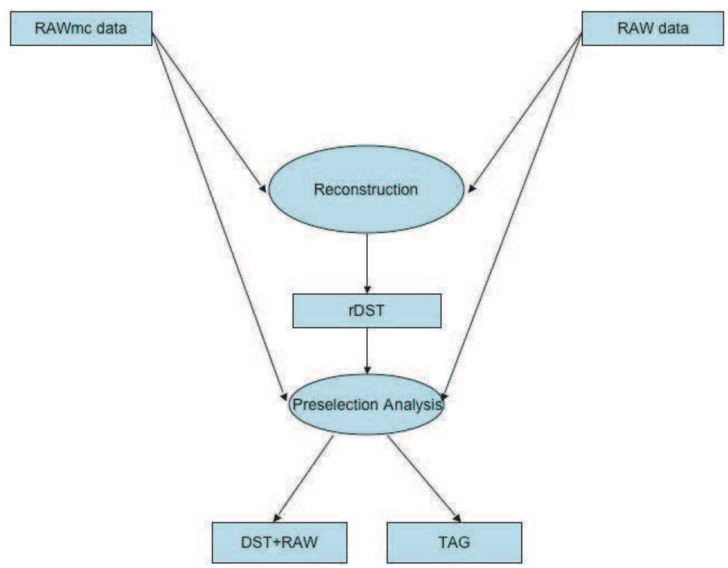

Figure 117. The LHCb computing logical dataflow model

The "real" raw data from the detector is produced via the Event Filter farm of the online system. The first step is to collect data, triggering on events of interest. The RAW data are transferred to the CERN Tier 0 centre for further processing and archiving. The RAW data, whether real or simulated, must then be reconstructed in order to provide physical quantities such as calorimeter clusters to provide the energy of electromagnetic and hadronic showers, trackers hits to be associated to tracks whose position and momentum are determined. Information about particle identification (electron, photon, π^0, hadron separation, muons) is also reconstructed from the appropriate sub-systems. The event reconstruction results in the generation of new data, the Data Summary "Tape" (DST). Only enough data will be stored in the DST that is written out during reconstruction to allow the physics pre-selection algorithms to be run at a later stage. This is known as a reduced DST (rDST.) The first pass of the reconstruction will happen in quasi-real time. It is planned to reprocess the data of a given year once, after the end of data taking for that year, and then periodically as required. This is to accommodate improvements in the algorithms and to make use of improved determinations of the calibration and alignment of the detector in order to regenerate new improved rDST information.

The rDST is analysed in a production-type mode in order to select event streams for individual further analysis. This activity is known as “stripping." The rDST information is used to determine the momentum four vectors corresponding to the measured particle tracks, to locate primary and secondary vertices and algorithms applied to identify candidates for composite particles whose four-momentum are reconstructed. Each particular channel of interest will provide such a pre-selection algorithm. The events that pass a physics working group's selection criteria are written out for further analysis. Since these algorithms use tools that are common to many different physics analyses they are run in production-mode as a first step in the analysis process. The events that pass the selection criteria will be fully re-reconstructed, recreating the full information associated with an event. The output of the stripping stage will be referred to as the (full) DST and contains more information than the rDST. Before being stored, the events that pass the selection criteria will have their RAW data added in order to have as detailed event information as needed for the analysis. An event tag collection will also be created for faster reference to selected events. It contains a brief summary of each event's characteristics as well as the results of the pre-selection algorithms and a reference to the actual DST record. The event tags are stored in files independent of the actual DST files. It is planned to run this production-analysis phase 4 times per year: once with the original data reconstruction; once with the re-processing of the RAW data, and twice more, as the selection cuts and analysis algorithms evolve.

The baseline LHCb computing model is based on a distributed multi-tier regional centre model. It attempts to build in flexibility that will allow effective analysis of the data whether the Grid middleware meets expectations or not. A schematic of the LHCb computing model is given in Figure 144.

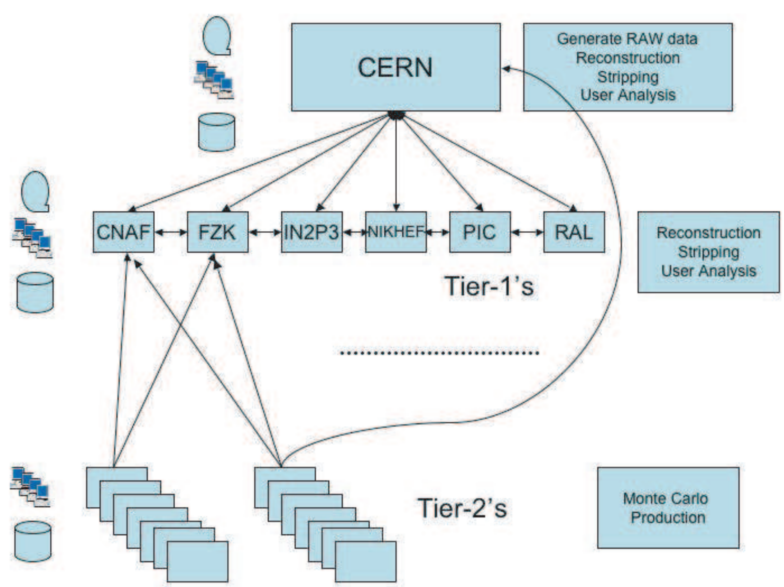

Figure 118. Schematic of the LHCb Computing Model.

CERN is the central production centre and will be responsible for distributing the RAW data in quasi-real time to the Tier-1 centres. CERN will also take on a role of a Tier-1 centre. An additional six Tier-1 centres have been identified: CNAF(Italy), FZK(Germany), IN2P3(France), NIKHEF(The Netherlands), PIC(Spain) and RAL(United Kingdom.) There are also a series of Tier-2 computing centres. CERN and the Tier-1 centres will be responsible for all the production-processing phases associated with the real data. The RAW data will be stored in its entirety at CERN, with another copy distributed across the other 6 Tier-1 centres. The 2nd pass of the full reconstruction of the RAW data will also use the resources of the LHCb online farm. As the production of the stripped DSTs will occur at these computing centres, it is envisaged that the majority of the distributed analysis of the physicists will be performed at CERN and at the Tier-1 centres. The current year's stripped DST will be distributed to all centres to ensure load balancing. The Tier-2 centres will be primarily Monte Carlo production centres, with both CERN and the Tier-1 centres acting as the central repositories for the simulated data. It should be noted that although we do not envisage any user analysis at the Tier-2's in the baseline model presented, it should not be proscribed, particularly for the larger Tier-2 centres.

It is expected that the reconstruction and the first stripping of the data at CERN and at the Tier-1 centres will follow the production in quasi real-time, with a maximum delay of a few days. The DST output of the stripping will remain on disk for analysis and be distributed to all other Tier-1 centres and CERN, whilst the RAW and rDST will be migrated to the mass storage system, MSS.

The re-processing of the data will occur over a 2-month period. During this process the RAW data will need to be accessed from the MSS both at CERN and the Tier-1 centres. The CPU resources available at the pit allow a significant fraction of the total re-processing and perhaps the subsequent stripping to be performed there. Hence at CERN there is an additional complication that the RAW data will also have to be transferred to the pit; similarly the produced rDST will have to be transferred back to the CERN computing centre. To enable later stripping it is necessary to distribute a fraction of the rDST produced at CERN during this re-processing to the Tier-1's; this is a consequence of the large contribution from the online farm.

The (two) stripping productions outside of the reconstruction or of the re-processing of the data will be performed over a one-month period. Both the RAW and the rDST will need to be accessed from the MSS to perform this production. The produced stripped DSTs will be distributed to all production centres.

The Monte Carlo production is expected to be an ongoing activity throughout the year and is the mainstay of the Tier-2 centres. The Tier-1 centres and CERN will act as the repository for the produced Monte Carlo data. The whole of the current year's Monte Carlo production DST will be available on disk at CERN and another 3 copies, on disk, distributed amongst the other 6 Tier-1 centres.

The 2008 resource requirements needed for the LHCb computing model at CERN and integrated across the Tier-1 centres and the Tier-2 centres are given in Table **??**.

9. Summary

References

[1] LHCb collaboration, A large Hadron Collider Beauty experiment, Technical Proposal, CERN document CERN/LHCC 98-4, 20 February 1998.

[2] LHCb Collaboration, “LHCb Reoptimized Detector Design and Performance”, CERN-LHCC-2003-030, 2003.

[3] G. Corti and G. von Holtey, "Study of beampipe induced background in the LHCb detector for the optimization of the vacuum chamber design", LHCb/2003-85.

[4] J.R. Knaster, "The vacuum chamber in the interaction region of particle colliders: a historical study and developments implemented in the LHCb experiment at CERN", Ph. D. Thesis, CERN-ETSII, May 2004.

[5] CERN Safety Instruction IS nr 25

[6] R. Veness et al., "Study of minimised UHV flanges for LHC experiments", CERN Vac. Tech. Note 2000-27.

[7] D. Ramos, "Beampipe supports in the high transparency region of the LHCb experiment", CERN Vaccum Tech. Note to be published.

[8] C. Benvenuti et al., "Vacuum properties of TiZrV non-evaporable getter films", Vacuum 60 (2001) 57.

[9] O. Grobner, Overview of the LHC vacuum system, Vacuum 60 (2001) 25-34

[10] LHCb Collaboration, LHCb Magnet Technical Design Report 1, CERN document CERN/LHCC 2000-007, 17 December 1999

[11] J.Andre’ *et al.,* Status of the LHCb Magnet System, IEEE Trans. Appl. Superconduct., Vol.12, No.1, pp.366-371, March 2002.

[12] F.Bersgma, Calibration of Hall sensors in three dimensions, 13th Int. Magnetic Measurement Workshop (IMMW13), SLAC Stanford, California, May 19-22, 2003.

Site	CPU (MSI2k.years)	Disk (TB)	Tape (TB)
CERN	0.36	350	631
Tier-1	1.77	1025	860
Tier-2	4.55	-	-

Table 13. Resource usage in 2008 at CERN, the Tier-1 and Tier-2 centres. A 1 GHz PIII processor is equivalent to 400 kSII2k.

[13] M.Patel,M.Losasso,T.Gys. Magnetic Shielding studies of the LHCb RICH Photon Detectors, RICH2004 Proceedings Nucl. Instr. and Methods A.

[14] J.Andre,*et al.,* IEEE Applied Superconductivity Vol. 14 (2004) 509.

[15] M.Losasso, *et al.,* Tests and Field Map of LHCb dipole magnet, IEEE Applied Superconductivity Vol. 16 (2005) 1700.

[16] The LHCb Technical Proposal. CERN/LHCC 98-4.

[17] M. Artuso and J. C. Wang, Nucl. Instrum. Meth. A **465**, 115 (2000) [arXiv:hep-ex/0007054].

[18] to be added

[19] V.Blobel and C.Kleinwort *A new method for the high-precision alignment of track detectors* hep-ex/0208021

[20] W. Baldini et al.,LHCb alignment strategy, LHCb note CERN-LHCb-2006-035

[21]

[22] EDMS ???

[23] EDMS 637676

[24] J. Gassner, M. Needham and O. Steinkamp, "Layout and Expected Performance of the LHCb TT Station", LHCb note 2003-140, 2003.

[25] LHCb Collaboration, "LHCb Inner Tracker Technical Design Report", CERN-LHCC-2002-029, 2002.

[26] J. Gassner, F. Lehner and S. Steiner, "The Mechanical Design of the LHCb Silicon Trigger Tracker", LHCb note 2004-110, 2004.

[27] A. Bay et al., "Hybrid Design, Procurement and Testing for the LHCb Silicon Tracker", Proceedings LECC 2005, Heidelberg, Germany, Sep 12-16, 2005, also available as: LHCb note 2005-061, 2005.

[28] S. Löchner and M. Schmelling, "The Beetle Reference Manual - chip version 1.3, 1.4 and 1.5", LHCb note 2005-105, 2005.

[29] A. Vollhardt, "A radiation tolerant fiber-optic readout system for the LHCb Silicon Tracker", Proceedings European Symposium on Semiconductor Detectors, Wildbad Kreuth, Jun 12-16, 2005, LHCb note 2005-032, 2005.

[30] **** reference for GOL chip (http://proj-gol.web.cern.ch/proj-gol/ does not work) ****

[31] D. Esperante, A. Vollhardt, "Design and Development of the Control Board for the LHCb Silicon Tracker", LHCb note 2007-xxx, 2007.

[32] C. Bauer et al., “Grounding, Shielding and Power Distribution for the LHCb Silicon Tracking”, LHCb note 2004-101, 2004.

[33] M. Agari et al., “Test Beam Results of Multi-Geometry Prototype Sensors for the LHCb Inner Tracker”, LHCb note 2002-058, 2002.

[34] R. Bernhard et al., “Measurements of Prototype Ladders for the Silicon Tracker with Laser”, LHCb note 2003-075, 2003.

[35] M. Agari et al., “Test-Beam Measurements on Prototype Ladders for the LHCb TT Station”, LHCb note 2003-082, 2003.

[36] J. Gassner et al., “Measurements of Prototype Ladders for the TT Station With a Laser”, LHCb note 2004-102, 2004.

[37] M. Agari et al., “Measurements of a Prototype Ladder for the TT Station in a 120 GeV/c π^- Beam”, LHCb note 2004-103, 2004.

[38] M. Needham, “Silicon Tracker Occupancies and Clustering”, LHCb note 2007-024, 2007.

[39] M. Needham, D. Volyanskyy, “Updated geometry description for the LHCb Trigger Tracker”, LHCb note 2006-032, 2006.

[40] A. Perrin, K. Vervink, “The Inner Tracker detector description and its implementation in the XML database”, LHCb note 2006-018, 2006.

[41] The LHCb Collaboration, P.R. Barbosa *et al.*,*Outer Tracker Technical Design Report*, CERN/LHCC 2001-024, Sept. 2001.

[42] S. Bachmann, *Specifications for the drift gas quality of the outer tracking system*, LHCb-2002-031, May 2002.
S. Bachmann, *Proposal for the gas distribution in the outer tracking system*, LHCb-2003-054, Jul. 2003.

[43] G.W. van Apeldoorn *et al.*, *Beam Tests of Final Modules and Electronics of the LHCb Outer Tracker in 2005*, LHCb-2005-076, Oct 2005.

[44] J. Nardulli and N. Tuning, *A Study of the Material in an Outer Tracker Module*, LHCb-2004-114, Aug 2004.

[45] Yu. Guz, G. van Apeldoorn, A. Pellegrino, T. Sluijk, N. Tuning, A. Zwart, *Study of the Global Performance of an LHCb OT Front-End Electronics Prototype*, LHCb-2004-120, Jan 2005.

[46] N. Dressnandt et al., *Implementation of the ASDBLR Straw Tube Readout ASIC in DMILL Technology*, IEEE (2000) Trans. on Nucl. Sci. **48** n4, 1239, 2000,
R. Bevensee et al., *An Amplifier-Shaper-Discriminator with Baseline Restoration for the ATLAS Transition Radiation Tracker*, IEEE (1996) Trans. on Nuc. Sci. **43**, 1725, 1996.
See: `http://www.hep.upenn.edu/atlas/asdblr`

[47] T.Sluijk and U.Uwer, *Specification of the OTIS-to-ASDBLR interface*.
See: `www.physi.uni-heidelberg.de/~uwer/lhcb/Electronics/Review/`

[48] H. Deppe, U. Stange, U. Trunk, U. Uwer, *The OTIS Reference Manual v1.1*, Feb 2004.
See: `www.physi.uni-heidelberg.de/~uwer/lhcb/Electronics/Review/`

[49] U. Stange, *Development and Characterization of a Rad Hard Readout Chip for the LHCb Outer Tracker Detector*, PhD Thesis, Uni Heidelberg, 2005.

[50] A.N.M. Zwart, *OTIS Board Version 1.0*, Aug 2003.
See: `www.nikhef.nl/pub/experiments/bfys/lhcb/outerTracker/Electronics/FE-Electronics/otisboard.html`

[51] P. Moreira, T. Toifl, A. Kluge, G. Cervelli, A. Marchioro, and J. Christiansen, *GOL Reference Manual, Gigabit Optical Link Transmitter manual*, Version 1.9, Oct 2005.
See: `http://proj-gol.web.cern.ch/proj-gol/`

[52] U. Uwer, D. Wiedner, A. Berkien, T. Sluijk, A. Zwart, *Specifications for the IF13-2 Prototype of the Auxiliary Board for the Outer Tracker*, LHCb-2005-039, Jul 2005.

[53] Th. Bauer, J. Nardulli and N. Tuning, *Flatness of an Outer Tracker Module*, LHCb-2005-009, Apr 2005.

[54] N. Tuning and A. Pellgrino, *Flatness of an Outer Tracker Layer in a Prototype C-frame*, LHCb-200, 2007, to be published.

[55] A. Pellegrino *et al.*, LHCb-200, 2007, to be published.

[56] H. Dekker, H. van der Graaf, et al., *The RASNIK/CCD 3-Dimensional alignment system*, 3rd International Workshop On Accelerator Alignment (IWAA 93), Sep 1993, Annecy.
See: `www.slac.stanford.edu/econf/C930928/papers/017.pdf`

[57] LHCb Collaboration, LHCb reoptimized detector design and performance, CERN/LHCC/2003-030, 2003.

[58] LHCb RICH group, LHCb RICH 1 Engineering Design Review Report, LHCb-2004-121; Geneva : CERN, 20 Oct 2005.

[59] R. Plackett: Photon detectors for the Ring Imaging Cherenkov detectors of the LHCb experiment, PhD Thesis, University of London, 2006.

[60] LHCb Collaboration, LHCb RICH Technical Design Report, CERN/LHCC 2000-0037, LHCb TDR 3, 7 September 2000.

[61] LHCb RICH group, LHCb RICH 2 engineering design review report, LHCb-2002-009; Geneva : CERN, 1 Mar 2002.

[62] T. Bellunato, Development of Ring Imaging Cherenkov Detectors for LHCb, PhD Thesis, Universita' degli Studi di Milano, Milano 2001.

[63] M. Laub, Development of opto-mechanical tools and procedures for a new generation of RICH-detectors at CERN, PhD Thesis, Ceské vysoké ucení technické v Praze, Prague 2001.

[64] L. Fernández Hernando, New Automatic Techniques to Test Optical Components of the Next Generation of RICH Detectors at CERN, PhD Thesis, Universitat Politècnica de Catalunya, Barcelona 2001.

[65] G. Aglieri Rinella, Development of the Photon Detection, Acquisition and Optical Systems of Modern Ring Imaging Cherenkov Detectors, PhD Thesis, Universita' degli Studi di Palermo, Palermo 2006

[66] A. Braem *et al.*, Nucl. Instrum. Methods Phys. Res., A 553 (2005) 182-186

[67] E. Albrecht *et al.*, Nucl. Instrum. Methods Phys. Res., A 433 (1999) 47-58

[68] E. Albrecht *et al.*, Nucl. Instrum. Methods Phys. Res., A 502 (2003) 236-240
P. Fauland, The COMPASS Experiment and the RICH-1 Detector, PhD Thesis, Universität Bielefeld, Mannheim 2004

[69] T. Gys, Magnetic field simulations for the LHCb-RICH 2 detector, LHCb-2002-029; Geneva : CERN, 06 May 2002
M. Patel et al., Magnetic Shielding Studies of the RICH Photon Detectors, LHCb-2005-055; CERN-LHCb-2005-055
M. Alemi, Passive magnetic shielding calculation for the photodetectors of RICH2, LHCb-98-017; Geneva : CERN, 02 Feb 1998

[70] A. Papanestis, Limits of software compensation for mirror misalignment of the RICH detectors, LHCb-2001-141; Geneva : CERN, 07 Feb 2002

[71] C. D'Ambrosio et al., The optical systems of LHCb RICHes : a study on the mirror walls and mirrors specifications, LHCb-2000-071; Geneva : CERN, 12 Apr 2001

[72] M.Y. Barnykov et al., *Nucl. Instr. and Meth.* A419 (1998) 584.

[73] T. Bellunato et al., *Nucl. Instr. and Meth.* A519 (2004) 493.

[74] T. Bellunato et al., *Nucl. Instr. and Meth.* A527 (2004) 319.

[75] T. Bellunato et al., *Nucl. Instr. and Meth.* A556 (2006) 140.

[76] O. Ullaland, *Nucl. Instr. and Meth.* A553 (2005) 107.

[77] D. Perego, Proceedings RICH2007- in preparation

[78] M. Alemi et al. , Nucl. Instrum. Meth. A449, 49-59 (2000)

[79] K. Wyllie et al., Nucl. Instr. and Meth. A 530 (2004) 82.

[80] P. Moreira et al., "A radiation tolerant gigabit serializer for LHC data transmission", Proceedings of the Seventh Workshop on Electronics for LHC Experiments, CERN/LHCC/2001-034 (2001) 145.

[81] J. Christiansen et al., IEEE Trans. Nucl. Sci, Vol 43, June 1996, pp 1773-1777.

[82] M.Adinolfi, Nucl. Inst. and Meth. A 553 (2005) 328.

[83] D. Breton et al., “SPECS: A Serial Protocol for the Experiment Control System of LHCb”,https://lhcb.lal.in2p3.fr/Specs/.

[84] C. Arnabolid et al., "The high voltage distribution system for the hybrid photodetector arrays of RICH1 and RICH2 at LHCb", submitted to IEEE Trans. Nucl. Sci.

[85] C. D’Ambrosio et al., The LHCb RICH Detector Control System, LHCb–2004–071; Geneva: CERN, 15 Sep 2004

[86] C. Gaspar et al, The Use of Credit Card-sized PCs for interfacing electronics boards to the LHCb ECS, LHCb–2001–147; Geneva: CERN, 5 Dec 2001

[87] C. Gaspar, M. Dönszelmann, A Distributed Information Management System for the Delphi experiment at CERN Presented at: IEEE Eight Conference REAL TIME ’93 on Computer Applications in Nuclear, Particle and Plasma Physics (Vancouver, 8–11 June 1993)

[88] B. Hallgren et al., The Embedded Local Monitor Board (ELMB) in the LHC Front-end I/O Control System, presented at the 7th Workshop on Electronics for LHC Experiments, Stockholm, Sweden, 10–14 Sep 2001

[89] 2004 RICH TB paper - in preparation

[90] S. Easo *et al.*, Simulation of LHCb RICH Detectors using GEANT4, IEEE Transactions on Nuclear Science, VOL. 52, NO. 5, October 2005, page 1665-1668

[91] R. Forty, O. Schneider, RICH Pattern Recognition, LHCb-98-040, April 1998.

[92] R. Forty, C. Jones, C. Lazzeroni, R. Muresan, M. Patel, G. Wilkinson, Reconstruction and calibration strategies for the LHCb RICH detector, to be published in Proceedings of *Computing in High Energy and Nuclear Physics*, 2006, Mumbai, India.

[93] E. Guschin, S. V. Laptev, Monte-carlo study of LHCb preshower, LHCb-2000-030.

[94] S. Barsuk *et al.*, Design and construction of the electromagnetic calorimeter for the LHCb experiment, LHCb-2000-043.

[95] R. Djeliadine,O. Iouchtchenko, V. F. Obraztsov, LHCb hadron trigger and HCAL cell size and length optimisation, LHCb-99-035.

[96] C. Beigbeder-Beau *et al.*, A joint proposal for the Level 0 calorimetric triggers, LHCb-99-017.

[97] C. Beigbeder-Beau *et al.*, The Front-End Electronics for LHCb calorimeters, LHCb-2000-028.

[98] S. Bota *et al.*, Scintillator Pad Detector Front-End Electronics, LHCb-2000-027.

[99] G. Böhner, R. Cornat, J. Lecoq, P. Perret, Front-End Electronics for the LHCb Preshower Detector, LHCb-2000-048.

[100] B. Delcourt, J. Lefrançois, Investigation of widening of the π^0 mass peak with electronic defects, LHCb-2000-029.

[101] S. N. Filippov *et al.*, Design and construction of the LHCb Scintillator-Pad/Preshower Detector, LHCb-2000-042.

[102] S. N. Filippov *et al.*, Experimental performance of PS/SPD prototypes, LHCb-2000-031.

[103] G. Böhner, A. Falvard, J. Lecoq, P. Perret, C. Trouilleau, Very Front-End Electronics for LHCb Preshower, LHCb-2000-047.

[104] Z. Ajaltouni et al., Nucl.Instr. Meth. A **504**, 9 (2003); Z. Ajaltouni et al., Nucl.Instr. Meth. A **504**, 250 (2003).

[105] E. Aguilo et al., Test of Multi-Anode Photomultiplier Tubes for the LHCb scintillator Pad Detector, Nucl.Instr. Meth. A **538**, 255 (2005).

[106] S. Monteil, on behalf of the LHCb collaboration, *Prepared for 11th Int. Conf. on Calorimetry in High-Energy Physics (Calor 2004), Perugia, Italy, 2004, 156-161.*

[107] L. Garrido, D. Gascón, R. Miquel; D. Peralta, Results of tagged photon test beam for the Scintillator Pad Detector, LHCb-2000-032.

[108] HERA-B Collaboration, HERA-B Proposal, DESY-PRC 94/02 May 1994.
HERA-B Collaboration, HERA-B Design Report, DESY-PRC 95/01 January 1995.

[109] J. C. Gregory *et al.* [Phenix Collaboration], "PHENIX: Preliminary conceptual design report”

[110] E.Tarkovsky, *Nucl. Instrum. Meth.* **A379**, 515 (1996).

[111] A.Bazilevsky *et al*, *IEEE Transactions on Nuclear Science* v.43, No.3 (1996).

[112] J.Badier *et al*, *Nucl.Instrum.Meth.* **A348** 74 (1994).

[113] Design, construction, quality control and performance study with cosmic rays of modules for the LHCb electromagnetic calorimeter. LHCb note in preparation.

[114] Study of the ECAL modules performance at the test beam. LHCb note in preparation.

[115] V. Brekhovskikh, R. I. Dzhelyadin, A. K. Konoplyannikov, V. I. Rykalin, The WLS Fiber Time Properties Study, LHCb-2000-39.

[116] R. Dzhelyadin, The LHCb hadron calorimeter, Proc. of the 8th Int. Conf. on Instrumentation for Colliding Beam Physics, Nucl. Instr. and Meth. A 494 (2002) 332

[117] L. G. Afanasieva *et al.*, The hadron calorimeter design and construction, LHCb-2000-045.

[118] M. Bonnet *et al.*, The hadron calorimeter prototype design and construction, LHCb-2000-035.

[119] C. Coca *et al.*, The hadron calorimeter prototype beam-test results, LHCb-2000-036.

[120] R. I. Dzhelyadin *et al.*, The hadron calorimeter module-0 construction, LHCb-2001-122.

[121] G. I. Britvich *et al.*, The HCAL optics radiation damage study, LHCb-2000-037.

[122] I. Koroko, J. Ocariz, A. Schopper, HCAL performance with irradiated sub-components, LHCb-2000-038.

[123] C. Beigbeder-Beau, D. Breton, O. Callot, J. Lefrançois, The readout of the LHCb calorimeters, LHCb-2000-046.

[124] G. Böhner, LHCb preshower signal characteristics, LHCb-2000-026.

[125] D. Gascón *et al.*, The Front End Electronics of the Scintillator Pad Detector of LHCb Calorimeter, p153 in report CERN 2007-001.

[126] D. Breton, D. Charlet, Using the SPECS in LHCb, LHCb-2003-005.

[127] A. Arefev *et al.*, Design of PMT base for the LHCb electromagnetic calorimeter, CERN-LHCb-2003-150.

[128] G. Avoni *et al.* [HERA-B ECAL Collaboration], The HERA-B ECAL electronics and monitoring *Prepared for 9th Conference on Calorimetry in High Energy Physics (CALOR 2000), Annecy, France, 9-14 Oct 2000.* Published in ”Annecy 2000, Calorimetry in high energy physics” 239-248.

[129] S. Amato *et al.*, “Analysis of the $B_s^0 \to \mu^+\mu^-$ Decay with the Reoptimized LHCb Detector”, LHCb-2003-165, 2003.

[130] LHCb Collaboration, “LHCB Muon System Technical Design Report", CERN-LHCC-2001-010, 2001.

[131] LHCb Collaboration, “LHCb Addendum to the Muon System Technical Design Report", CERN-LHCC-2003-002, 2003.

[132] LHCb Collaboration, “LHCb Second Addendum to the Muon System Technical Design Report”, CERN-LHCC-2005-012, 2005.

[133] LHCb Collaboration, “LHCb Trigger System Technical Design Report”, CERN-LHCC-2003-031, 2003.

[134] G. Martellotti *et al.*, “Particle Rates in the LHCb Muon Detector”, CERN-LHCb-2005-75, 2005.

[135] G. Bencivenni, et al., Nucl. Instr. and Meth. in Phys. Res. A 513 (2003) 264.

[136] M. Alfonsi et al., IEEE Trans. Nucl. Sci. NS-51 (2004) 2135.

[137] B. Bochin *et al.* Note LHCB-Muon 2000-102 (2000).

[138] D. Hutchcroft *et al.* Note LHCB-Muon 2001-024 (2001).

[139] B. Marechal *et al.* Note LHCb-Muon 2001-150 (2001).

[140] M. Anelli *et al.*, Note LHCb-Muon 2004-074 (2004).

[141] G. Lanfranchi, "Proceedings of the X International Vienna Conference on Instrumentation 2004", Nucl. Instr. and Meth. in Phys. Res. A 531 (2004) 221.

[142] E. Danè et al., Nucl. Instr. and Meth. in Phys. Res. A 572 (2007) 682.

[143] W. Riegler, Note LHCb-Muon 2000-060 (2000)

[144] R. Veenhof, "GARFIELD, a drift-chamber simulation program", W5050 (1999);`http://consult.cern.ch/ writeup/garfield/`

[145] I. Smirnov, `http://consult.cern.ch/writeup/heed/`

[146] S. Biagi, `http://consult.cern.ch/writeup/magboltz/`

[147] A. Kachtchouk *et al.*, Note LHCb-MUON 2001-026, (2001).

[148] M. Anelli *et al.*, IEEE Trans. Nucl. Sci. 53 (2006) 330.

[149] P. Ciambrone *et al.*, Nucl. Instr. and Meth. in Phys. Res. A 545 (2005) 156.

[150] W. Baldini *et al.*, "Wire Tension Meaasurement System with Laser Readout Technique", Note LHCb-MUON 2007-xxx, (2007).

[151] M. Anelli *et al.*, Note LHCb-Muon 2003-040 (2003)

[152] M. Anelli *et al.*, "Extensive aging test of two prototypes of four-gap MWPC for the LHCb Muon System", Note LHCb-Muon 2004-029 (2004)

[153] F. Sauli, Nucl. Instr. and Meth. in Phys. Res. A386 (1997) 531.

[154] A. Breessan *et al.*, Nucl. Instr. and Meth. A425 (1999) 262.

[155] M. Alfonsi *et al.*, IEEE Trans. Nucl. Sci. 52 (2005) 2872.

[156] M. Alfonsi *et al.*, Nucl. Phys. B, Proc. Suppl. 150 (2006) 159.

[157] D. Moraes *et al.*, Proc. of the 2001 IEEE International Symposium on Circuits and Systems, ISCAS 2001. Volume 1 (2001) 360.

[158] W. Bonivento *et al.*, Proceedings of the 7th workshop on electronics for LHC experiments Stockholm, Sweden (2001)
http://lhc-electronics-workshop.web.cern.ch/
LHC-electronics-workshop/ 2001/muon/moraesboniv.pdf

[159] W. Bonivento, Proc. of the 11th Workshop on Electronics for LHC and future Experiments Heidelberg, Germany (2005) http://www.lecc2005.uni-hd.de/

[160] S. Cadeddu *et al.*, IEEE Trans. Nucl. Sci. 51 (2004) 1961.

[161] W. Bonivento *et al.*, "Design and production of the CARDIAC boards", Note LHCb-Muon 2007-zzz (2007).

[162] J. Christiansen *et al.*, "TTCrx Reference Manual" V 3.10 (2005) http://ttc.web.cern.ch/TTC/TTCrx_manual3.10.pdf

[163] P. Moreira "QPLL Manual" V 1.1 (2005) http://proj-qpll.web.cern.ch/proj-qpll/images/qpllManual.pdf

[164] **Only broken links available on this item**

[165] Technical documentation on the ELMB can be found at http://elmb.web.cern.ch/ELMB/ELMBhome.html

[166] G. Sabatino *et al.*, Note LHCb-Muon 2006-011 (2006).

[167] S. de Capua *et al.*, Note LHCb-Muon 2006-010 (2006).

[168] M. Anelli *et al.* Note LHCb-Muon 2005-003 (2005).

[169] G. Corradi *et al.*, Nucl. Instr. and Meth. in Phys. Res. A572 (2007) 96.

[170] LHCb Collaboration, "LHCb Online System Technical Design Report",CERN-LHCC-2001-040, 2001.

[171] LHCb Collaboration, "LHCb Trigger System Technical Design Report",CERN-LHCC-2003-031, 2003.

[172] LHCb Collaboration, "Addendum to the LHCb Online System Technical Design Report",CERN-LHCC-2005-039, 2005.

[173] G. Haefeli, A. Bay, A. Gong, H. Gong, M. Muecke, N. Neufeld, O. Schneider, "The LHCb DAQ interface board TELL1", Nucl. Instrum. and Meth. A **560** (2006) 494–502.

[174] T. Sjöstrand *et al.*, Computer Physics Commun. **135** (2001) 238.

[175] GEANT detector description and simulation tool, CERN Program Library long writeup W5013 (1994).

[176] all references about pattern recognition algorithms, old and new. M. Benayoun and O. Callot, The forward tracking, an optical model method, LHCb/2002–8.
R. Forty, Track Seeding, LHCb/2001–109.
J. van Tilburg, Matching VELO tracks with seeding tracks, LHCb/2001–103.
Y. Xie, Short track reconstruction with VELO and TT, LHCb/2003–100.
R. Hierck, Track following in LHCb, LHCb/2001–112.

[177] R. Hierck *et al.*, Performance of the LHCb OO track-fitting software, LHCb/2000–86.
+ new notes

[178] R. Forty and O. Schneider, RICH pattern recognition, LHCb/1998–40.

[179] J.R.T. de Mello Neto and M. Gandelman, Muon ID performance with the reoptimized LHCb detector, LHCb/2003–89.

[180] H. Terrier and I. Belyaev, Particle identification with LHCb calorimeters, LHCb/2003–92.

[181] V. Breton, N. Brun and P. Perret, A clustering algorithm for the LHCb electromagnetic calorimeter using cellular automaton, LHCb–2001–123.

[182] O. Deschamps *et al.*, Photon and neutral pion reconstruction, LHCb/2003–91.