Chapter 5

Tracking

The LHCb tracking system consists of the vertex locator system (VELO) and four planar tracking stations: the Tracker Turicensis (TT) upstream of the dipole magnet and T1-T3 downstream of the magnet. VELO and TT use silicon microstrip detectors. In T1-T3, silicon microstrips are used in the region close to the beam pipe (Inner Tracker, IT) whereas straw-tubes are employed in the outer region of the stations (Outer Tracker, OT). The TT and the IT were developed in a common project called the Silicon Tracker (ST).

The VELO is described in section 5.1 the ST in section 5.2 and the OT in section 5.3.

5.1 Vertex locator

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 and c-hadron decays [28]. The VELO consists of a series of silicon modules, each providing a measure of the r and ϕ coordinates, arranged along the beam direction (figure 5.1). Two planes perpendicular to the beam line and located upstream of the VELO sensors are called the *pile-up veto system* and are described in section 7.1. The VELO sensors 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. The detectors are mounted in a vessel that maintains vacuum around the sensors and is separated from the machine vacuum by a thin walled corrugated aluminum sheet. This is done to minimize the material traversed by a charged particle before it crosses the sensors and the geometry is such that it allows the two halves of the VELO to overlap when in the closed position. Figure 5.2 shows a cross section of the VELO vessel, illustrating the separation between the primary (beam) vacuum and the secondary (detector) vacuum enclosed by the VELO boxes. Figure 5.3 shows an expanded 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 corrugated foils, hereafter referred to as *RF-foils*, form the inner faces of the boxes (*RF-boxes*) within which the modules are housed. They provide a number of functions which are discussed in the following sections.

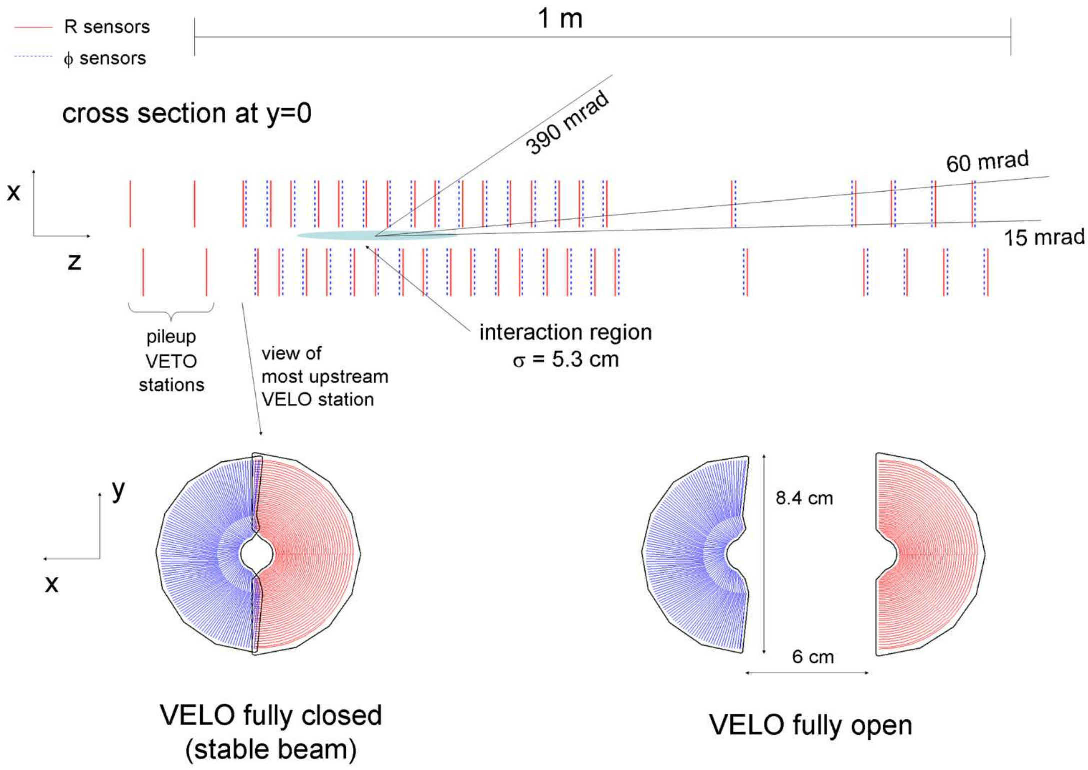

Figure 5.1: Cross section in the (x,z) plane of the VELO silicon sensors, at $y = 0$, with the detector in the fully closed position. The front face of the first modules is also illustrated in both the closed and open positions. The two pile-up veto stations are located upstream of the VELO sensors.

5.1.1 Requirements and constraints

The ability to reconstruct vertices is fundamental for the LHCb experiment. The track coordinates provided by the VELO are used to reconstruct production and decay vertices of beauty- and charm-hadrons, to provide an accurate measurement of their decay lifetimes and to measure the impact parameter of particles used to tag their flavour. Detached vertices play a vital role in the High Level Trigger (HLT, see section 7.2), and are used to enrich the b-hadron content of the data written to tape, as well as in the LHCb off-line analysis. The global performance requirements of the detector can be characterised with the following interrelated criteria:

- Signal to noise[1] ratio (S/N): in order to ensure efficient trigger performance, the VELO aimed for an initial signal to noise ratio of greater than 14 [29].
- Efficiency: the overall channel efficiency was required to be at least 99% for a signal to noise cut S/N> 5 (giving about 200 noise hits per event in the whole VELO detector).

[1]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.

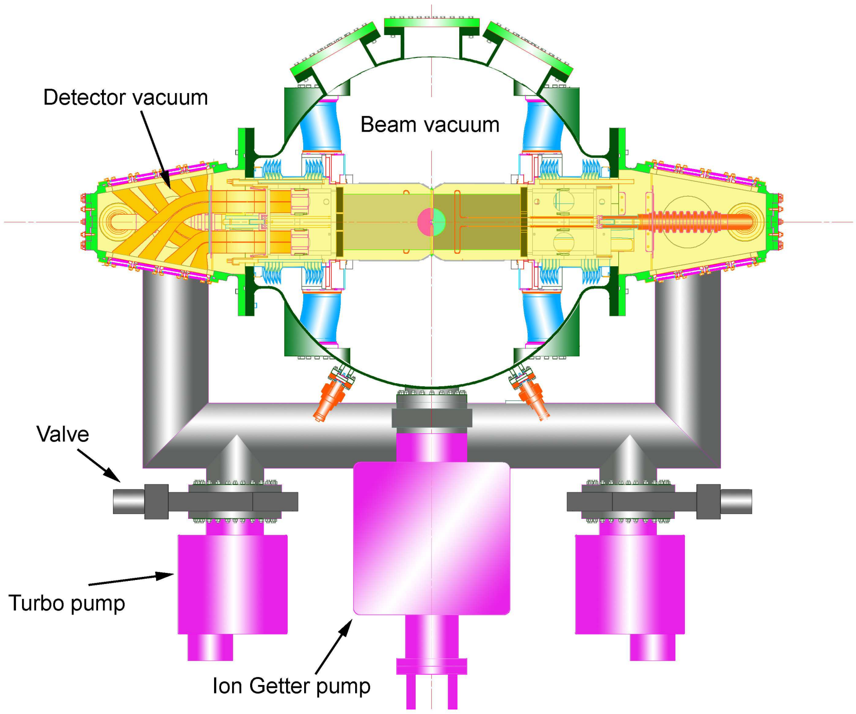

Figure 5.2: Cross section of the VELO vacuum vessel, with the detectors in the fully closed position. The routing of the signals via kapton cables to vacuum feedthroughs are illustrated. The separation between the beam and detector vacua is achieved with thin walled aluminium boxes enclosing each half.

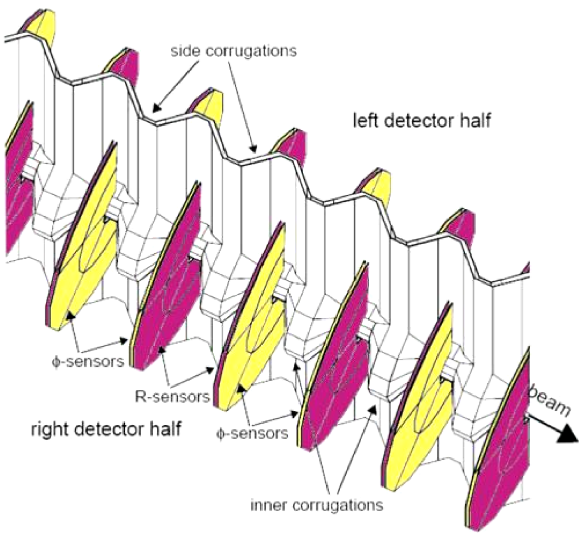

Figure 5.3: Zoom on the inside of an RF-foil, as modelled in GEANT, with the detector halves in the fully closed position. The edges of the box are cut away to show the overlap with the staggered opposing half. The R- and ϕ-sensors are illustrated with alternate shading.

- 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), in order to achieve the impact parameter resolution performance described in section 10. Furthermore, it was required that the resolution not be degraded by irradiation nor by any aspect of the sensor design.

Another important consideration is the *spillover probability*, which is defined as the fraction of the peak signal remaining after 25 ns. An additional requirement imposed on the system, affecting the readout electronics, is that the spillover probability be less than 0.3, in order to keep the number of remnant hits at a level acceptable for the HLT [30].

The construction of the VELO followed a number of requirements and constraints, which are briefly described in this section.

Geometrical

The VELO has to cover the angular acceptance of the downstream detectors, i.e. detect particles with a pseudorapidity in the range[2] $1.6 < \eta < 4.9$ and emerging from primary vertices in the range $|z| < 10.6$ cm. The detector setup was further constrained by the following considerations:

- Polar angle coverage down to 15 mrad for a track emerging at z=10.6 cm downstream from the nominal interaction point (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: $z_{N-2} \simeq 65$ cm.

- A track in the LHCb spectrometer angular acceptance of 300 mrad should cross at least three VELO stations. Given a maximum[3] 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. For this reason, an $r\phi$ geometry was selected for the design. Each VELO module was designed to provide the necessary 3D spatial information to reconstruct the tracks and vertices. One of the two sensors of the module, called the ϕ-measuring sensor, or ϕ-sensor, provides information on the azimuthal coordinate

[2] 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 Level-0 trigger (L0, see section 7.1).

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

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 rz tracking requirement imposes the additional constraint that the VELO circular strips should be centered as perfectly as possible around the beam axis. The result of simulation studies showing how the trigger performance would degrade as a function of various VELO R-sensor misalignments [31] indicate that the R-sensors should be mounted with a mechanical accuracy of better than 20 μm in x and y relative to each other within each half, and the two halves should be aligned to better than 100 μm relative to each other in these coordinates. The number of strips for both sensor types needed to satisfy the competing requirements of the LHCb environment, physics and a budgetary limit, is about 180000 channels.

Environmental

The VELO detector will be operated in an extreme radiation environment with strongly non-uniform fluences. The damage to silicon in the most irradiated area for one nominal year of running, i.e. an accumulated luminosity of 2 fb^{-1}, is equivalent to that of 1 MeV neutrons with a flux of 1.3×10^{14} $\text{n}_{\text{eq}}/\text{cm}^2$, whereas the irradiation in the outer regions does not exceed a flux of 5×10^{12} $\text{n}_{\text{eq}}/\text{cm}^2$. The detector is 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, the VELO cooling system was required to be 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 lifetime, continuous cooling after irradiation 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 ensure full depletion during the 3 years lifetime of the sensors was imposed as a further requirement.

Machine integration constraints

The required performance 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, imposing a number of special constraints which are briefly discussed here.

- The amount of material in front of the silicon detector is mainly determined by the necessity to shield against RF pickup and the mechanical constraint of building a sufficiently rigid foil. The detectors operate in a secondary vacuum and hence 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, at acceptable levels for efficient LHC and LHCb operation. The LHC beam vacuum chamber, and therefore also the VELO vacuum vessel, were required to be bakeable (to 160°C in the case of the VELO).

- 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 RMS spread 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 value 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 design of 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 into the shadow of the LHCb beampipe (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-foils, which represent a major fraction of the VELO material budget in the LHCb acceptance. In addition, the beam bunches passing through the VELO structures will generate wake fields which can affect the LHC beams. The RF foils, together with wake field suppressors which provide the connection to the rest of the beampipe, also provide the function of suppressing wake fields by providing continuous conductive surfaces which guide the mirror charges from one end of the VELO vessel to the other. These issues have been addressed in detail [32] and are further discussed in section 5.1.3.

5.1.2 Sensors and modules

Sensors

The severe radiation environment at 8mm from the LHC beam axis required the adoption of a radiation tolerant technology. The choice was n-implants in n-bulk technology with strip isolation achieved through the use of a p-spray. The minimum pitch achievable[4] 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 the vertex resolution.

The conceptual layout of the strips on the sensors is illustrated in figure 5.4. For the R-sensor the diode implants are concentric semi-circles with their centre at the nominal LHC beam position. In order to minimize the occupancy each strip is subdivided into four 45° regions. This also has the beneficial effect of reducing the strip capacitance. The minimum pitch at the innermost radius

[4] The company chosen to fabricate the LHCb sensors was Micron Semiconductor Ltd.

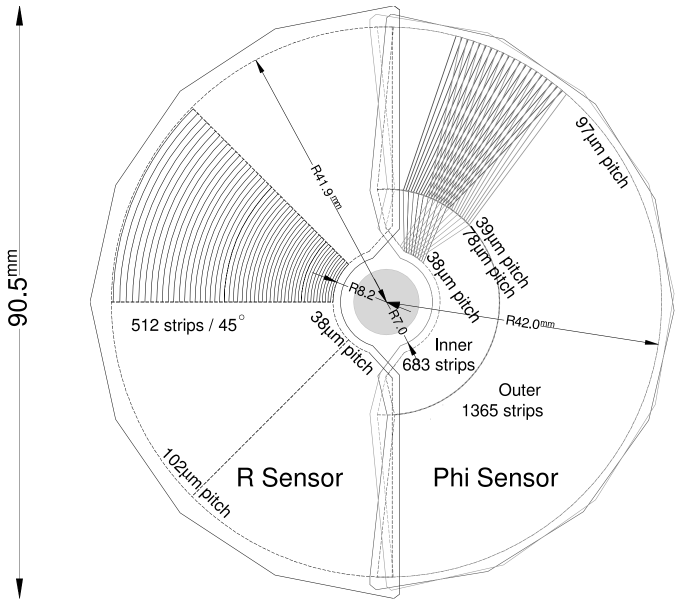

Figure 5.4: Sketch illustrating the $r\phi$ geometry of the VELO sensors. For clarity, only a portion of the strips are illustrated. In the ϕ-sensor, the strips on two adjacent modules are indicated, to highlight the stereo angle. The different arrangement of the bonding pads leads to the slightly larger radius of the R-sensor; the sensitive area is identical.

is 38 μm, increasing linearly to 101.6 μm at the outer radius of 41.9 mm. This ensures that measurements along the track contribute to the impact parameter precision with roughly equal weight.

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. However, this would result in unacceptably high strip occupancies and too large a strip pitch at the outer edge of the sensor. Hence, the ϕ-sensor 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 roughly 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 is complicated by the introduction of a skew to improve pattern recognition. At 8 mm from the beam the inner strips have an angle of approximately 20° to the radial whereas the outer strips make an angle of approximately 10° to the radial at 17 mm. The skew of inner and outer sections is reversed giving the strips a distinctive *dog-leg* design. The modules are placed so that adjacent ϕ-sensors have the opposite skew 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 principal characteristics of the VELO sensors are summarized in table 5.1.

The technology utilized in both the R- and ϕ-sensors is otherwise identical. Both sets of sensors are 300 μm thick. 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 μm of 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

Table 5.1: Principal characteristics of VELO sensors.

	R sensor	ϕ-sensor
number of sensors	42 + 4 (VETO)	42
readout channels per sensor	2048	2048
sensor thickness	300 μm	300 μm
smallest pitch	40 μm	38 μm
largest pitch	102 μm	97 μm
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

polysilicon 1 MΩ resistors and both detectors are protected by an implanted guard ring structure.

The pitch as a function of the radius r in μm increases linearly and is given by the following expressions:

$$\mathrm{R-sensor:}\qquad 40+(101.6-40)\times\frac{r-8190}{41949-8190}$$
$$\phi-\mathrm{sensor:}\qquad 37.7+(79.5-37.7)\times\frac{r-8170}{17250-8170}\quad (r<17250)$$
$$\phi-\mathrm{sensor:}\qquad 39.8+(96.9-39.8)\times\frac{r-17250}{42000-17250}\quad (r>17250)$$

The sensors were developed for high radiation tolerance. Early prototype detectors 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 was 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 inhomogeneous irradiation. It was shown that the transverse electric fields produce less than 2 μm effects on the cluster centroid.

A subset of the production sensors were exposed to a high neutron fluence ($1.3\times 10^{14}\,\mathrm{n_{eq}/cm^2}$) representing 1 year of operation at nominal luminosity. A strong suppression of surface breakdown effects was demonstrated. The evolution of the depletion voltage was found to correspond to the expectation over LHC operation: with an integrated luminosity of 2 fb^{-1} per year, the maximum deliverable full depletion voltage (500 V) is reached after approximately 3 years. During production the possibility arose of manufacturing full size n^+p sensors. These are expected to have similar long term radiation resistance characteristics to the n^+n technology, but feature some advantages, principally in cost of manufacture due to the fact that double sided processing is not needed. One full size module was produced in this technology and installed in one of the most upstream slots. It is forseen to replace all the VELO modules after damage due to accumulated radiation or beam accidents. The replacement modules will be constructed in the n^+p technology [33].

Modules

The module has three basic functions. Firstly it must hold the sensors in a fixed position relative to the module support. 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 is designed to hold the sensors in place to better than 50 μm in the plane perpendicular to the beam and within 800 μm along the direction of the beam. Sensor-to-sensor alignment (within a module) is designed to be better than 20 μm.

The module is comprised of a substrate, for thermal management and stability, onto which two circuits are laminated. This forms the hybrid. The substrate is fabricated *in-house* and is approximately $120\times170\times1$ mm. It has a core of 400 μm thick thermal pyrolytic graphite (TPG), and is encapsulated, on each side, with 250 μm of carbon fibre (CF). A CF frame of about 7 mm thickness surrounds the TPG and is bonded directly to the CF encapsulation to prevent delamination. The TPG is designed to carry a maximum load of 32 W away from the front-end chips. A semicircular hole is cut into the substrate under the region where the detectors are glued. Particular attention in the design and fabrication process is given to producing almost planar hybrids, to simplify the subsequent module production. Typical non-planarities of order 250 μm were achieved.

The circuits[5] were commercially hand populated to minimize exposure of the hybrid to high temperature and hence the possibility of delamination. The sensor front-end ASICs (Beetle 1.5 [34]) were then glued to the circuits. A total of 32 Beetle chips are used in each module. Kapton pitch adaptors,[6] 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 a sensor-to-sensor accuracy of better than 10 μm. The sensors were bonded to the electronics using a combination of H&K 710 and K&S 8090 bonding machines with 25 μm thick 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 pedestal and base. The pedestal is a low mass CF fibre construction designed to hold the hybrid stably. It is a hollow rigid structure approximately $140\times150\times10$ mm. Holes are drilled where appropriate to reduce the possibility of outgassing. One end was glued to the hybrid and the other to a CF base which contains two precisely manufactured invar feet. The design of the base allows repeatable mounting of the module to the module support with a precision a better than 10 μm. The accuracy of the final assembled modules satisfied the design criteria. The complete module, in schematic and after final assembly, is illustrated in figure 5.5.

The thermal performance of the modules was required to be such that the silicon could be operated below zero degrees, for a minimum cooling liquid temperature of -30°C. The performance of each module was monitored individually in a vacuum tank during construction with a thermal camera, as illustrated in figure 5.6. After the final cooling connections had been made on the module bases, all modules were rechecked for their thermal performance in vacuum. The temperature difference between the cooling blocks and the silicon is estimated to be about (21.4 ± 1)°C for individually cooled modules, and an improvement of up to 2°C is expected when the modules are

[5] The circuits were fabricated and laminated to the substrates by Stevenage Circuits, Ltd.

[6] The pitch adapters were manufactured at CERN.

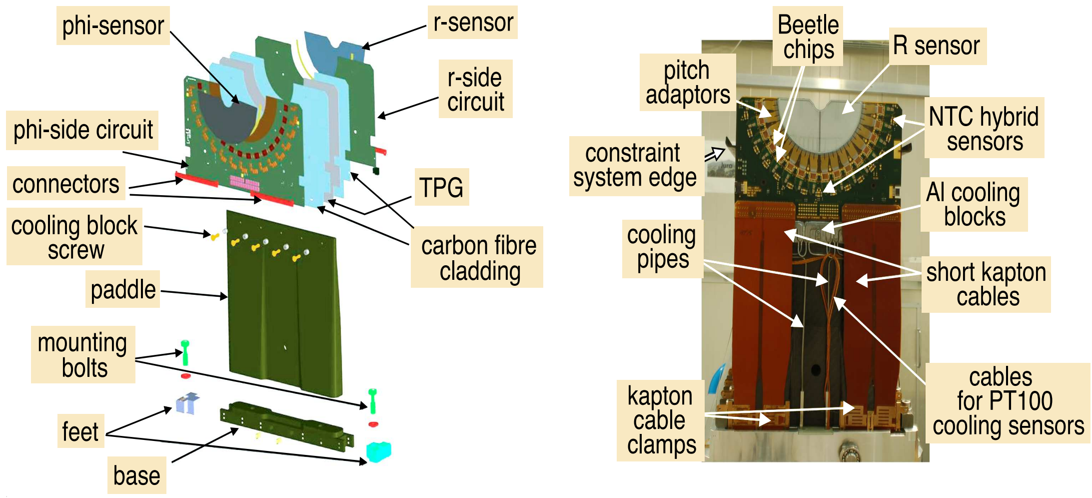

Figure 5.5: The left part of the figure illustrates the principal components of the VELO module. The right part of the figure is a photograph of the module as finally mounted. The short kapton cables are visible, along with the clamps at the base which prevent any long kapton cable movement coupling to the module itself. The aluminium blocks encasing the stainless steel CO_2 cooling tubes are also partially visible.

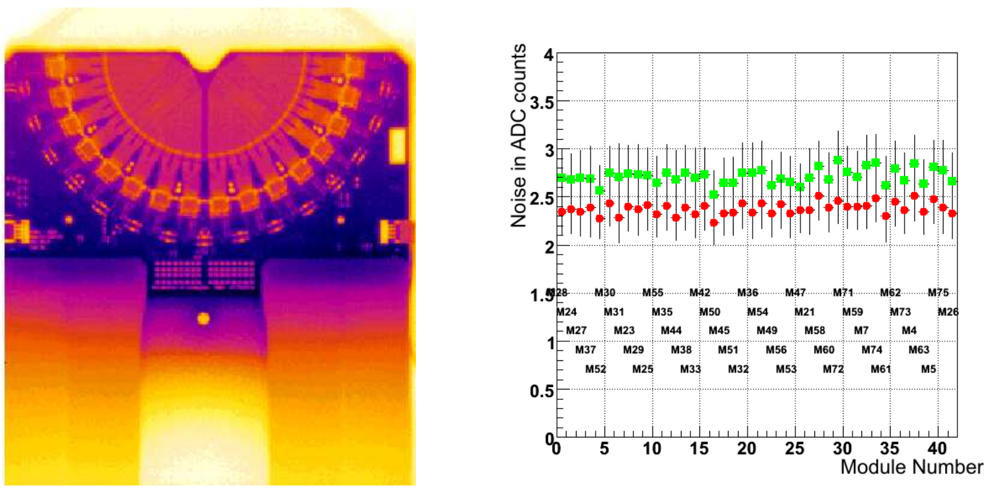

Figure 5.6: Quality control of the VELO modules. The left part of the figure shows an example of a thermograph which was routinely taken on both sides of the module during the construction. The right plot shows an example of a noise uniformity measurement for all the ϕ-sensors. The green squares are before common mode correction, and the red circles after the common mode correction. The error bars show the noise spread which is due to varying strip capacitances on the sensor.

cooled simultaneously. At the same time, the noise uniformity of the modules was checked, as shown in figure 5.6.

Each module was transported to CERN in a dedicated transport box [35] and submitted to a visual inspection with a high resolution microscope, including the survey of all bond wires, on

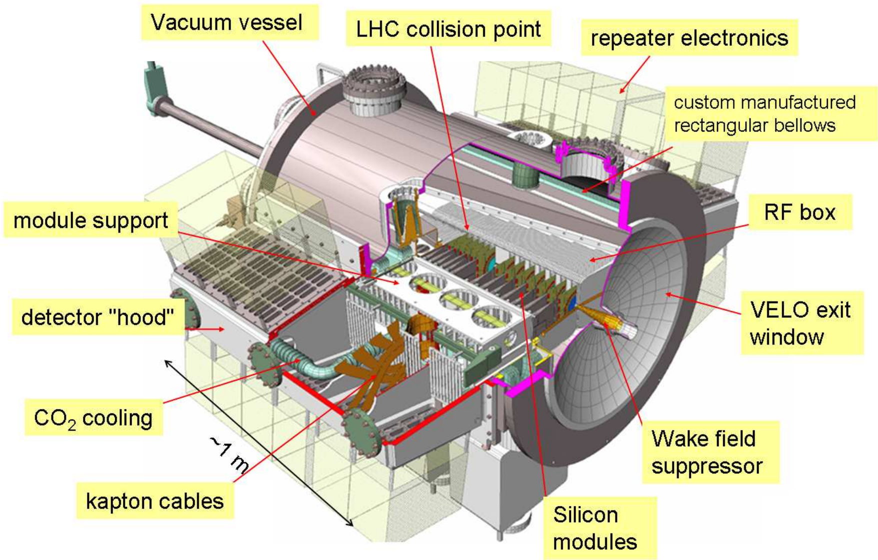

Figure 5.7: Overview of the VELO vacuum vessel.

reception. After this it underwent a 16 hour burn-in [36] consisting of a sequence of powering, cooling and data taking cycles, in vacuum. This was designed to uncover any inherent weaknesses introduced to the modules during manufacturing. Of all the modules transported, one was finally taken as a reserve as a result of the burn-in procedure, and a second was rejected after assembly due to a slightly worse thermal performance.

5.1.3 Mechanics

Introduction

The ultra high vacuum requirements of the LHC ring, the necessity for wake field suppression, the need to shield the detectors from electromagnetic effects induced by the high frequency 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 VELO mechanical design demanding. To meet all these constraints, a design with two detector halves was chosen, each placed inside a thin-walled aluminum box, as introduced in the previous sections. 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 5.7. 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 to

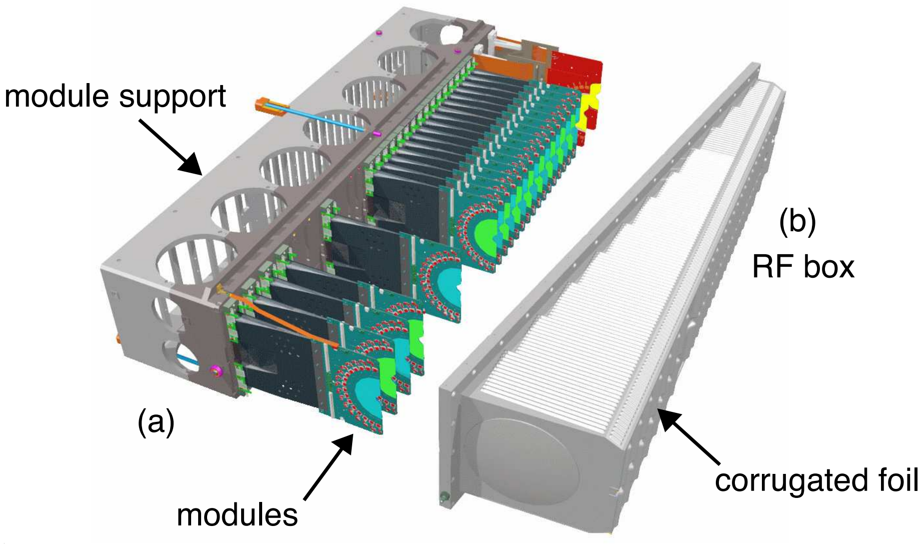

Figure 5.8: Exploded view of the module support and the modules (a), and the RF box (b). The corrugated foil on the front face of the box, which forms a beam passage can be seen. Its form allows the two halves to overlap when in the closed position.

the movement mechanism that is located outside the vacuum vessel. The module support, illustrated in figure 5.8, is mounted on bearings which lie in precisely machined slots inside the detector support and is bolted into position against three precision surfaces. The positioning of the modules on the module support is fully constrained with a combination of a slot and a dowel pin. The positioning of the modules relative to one another is determined by the precision of the machining of the module support. The positioning of the two halves relative to each other is determined by the precision of the module support, together with its positioning inside the detector support, which is adjustable. To suppress wake field effects, the dimension of the beampipe as seen by the proton beams has to vary very gradually. To match the beampipe upstream and downstream from the VELO, wake field suppressors made of 50 μm thick copper-beryllium have been constructed, so a good electrical match is provided in both the open and closed positions.

The durability of the wake field suppressors has been tested by performing an opening and closing movement 30,000 times, after which no damage was observed. A photograph of the downstream wake field suppressors is shown in figure 5.9. The exit foil of the vessel consists of a 2 mm thick aluminum window (see section 3).

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 into 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

Figure 5.9: Photograph of the downstream wakefield suppressor.

the optimal position. This is performed with a motion mechanism that can bring the detectors to their position with an accuracy of the order of 10 μm by means of a stepping motor with resolver readout. Additional potentiometers have been used to verify independently the proper functioning. The motion procedures are controlled by a Programmable Logic Controller (PLC). The boxes can move together in y and independently in x, and a series of stops prevents mechanical interference.

Vacuum system

In the LHC ring, Ultra High Vacuum conditions of better than 10^{-8} mbar are required, compatible with reduced beam-induced effects. To maintain these conditions, the beampipes are equipped with a Non-Evaporable Getter (NEG) layer. To maintain the low desorption yield of the NEG coating, the vacuum system will be vented during maintenance with ultra-pure neon. The thin walled detector boxes are expected to plastically deform at a pressure difference of 20 mbar (and to rupture above 50 mbar). 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 would saturate and contaminate the NEG material on the inside of the beampipe.

A 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 overpressure in the detector vacuum and 2 mbar overpressure in the beam vacuum. This asymmetry is designed to maximise the protection of the modules. This is achieved by using dedicated valves and restrictions, that are activated by membrane switches that react at the desired pressure difference. The complete vacuum system between the two sector valves around LHCb is controlled by a PLC.

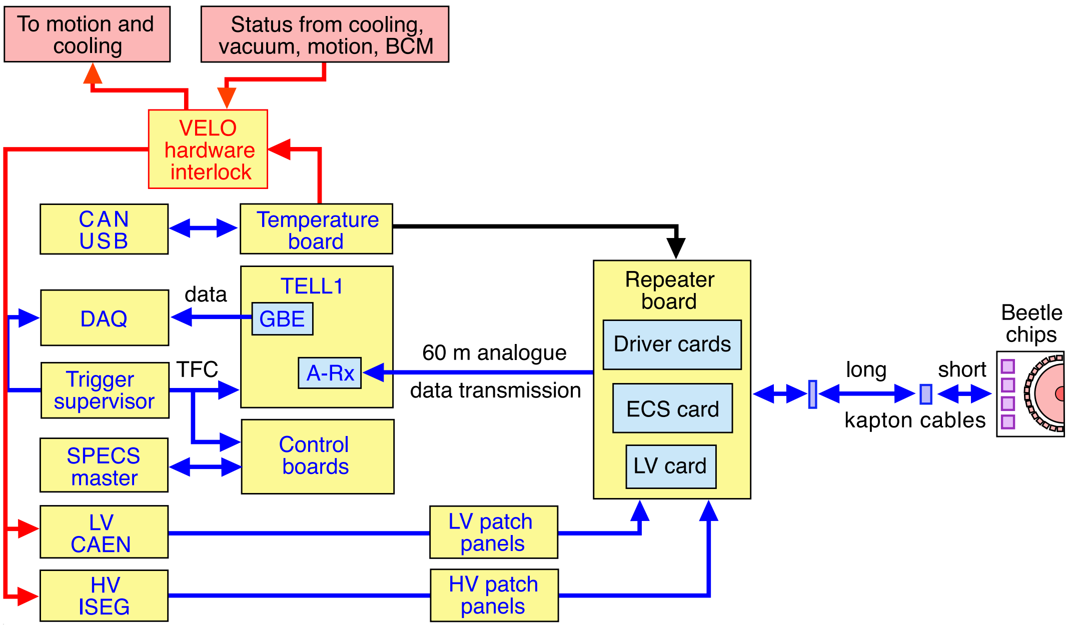

Figure 5.10: VELO readout chain for one side of one module.

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.

5.1.4 Electronics chain

The electronics layout is summarised in figure 5.10, which shows the readout chain for one side of one module. The individual components are discussed in this section.

Beetle chip

The VELO uses the Beetle, a custom designed radiation hard ASIC based on 0.12 μm CMOS technology[7] with an analog front-end, as the FE chip [34]. The chip was designed with a peaking time and sampling frequency to match the LHC bunch crossing rate of 40 MHz, and to be able to readout at a speed which can match the L0 accept rate of 1 MHz. For a full discussion of the design requirements see [28, 37]. Each of the 128 channels consists of a charge sensitive

[7] IBM, USA.

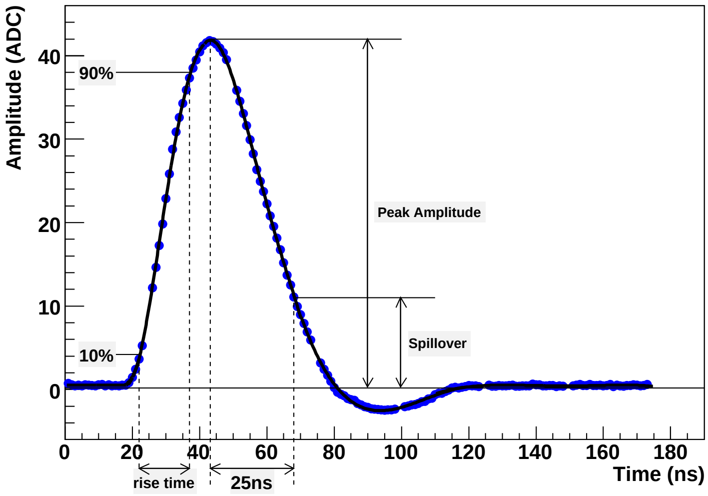

Figure 5.11: Beetle pulse shape as measured for a VELO sensor using test-pulses.

preamplifier/shaper and an analog pipeline of 160 stages designed to match the L0 4 μs latency. The data are brought off chip at a clock frequency of 40 MHz, with 32 channels multiplexed on each of 4 output lines. This allows the readout of one event to be achieved in 900 ns. The chip specifications required a survival more than 5 years of nominal operation at a dose of 2 MRad per year. In fact the chip is radiation hard against total dose effects of more than 100 MRad. Robustness against single event upset is achieved with redundant logic. It is programmable via a standard I2C interface, which controls the bias settings and various other parameters. The circuit can be tested via a charge injector (test-pulse circuit) with adjustable pulse height. In the case of the pile-up readout, where the data are used in the L0 trigger, a fast binary signal is required. This is achieved by routing the output of the front-end amplifier to a comparator stage with extra output pads on the chip.

For the VELO, the spillover value, or the remnant signal remaining after 25 ns, is of particular importance, and must be below 0.3 of its peak value for reliable HLT performance. Figure 5.11 shows the measured pulse shape on a Beetle fully mounted on a VELO module and bonded to the sensor, using a test-pulse corresponding to a minimum ionising particle in 300 μm of silicon, for the bias settings expected for LHCb operation. The rise time is measured to be 14.7 ± 0.5 ns and the spillover $(26.0 \pm 0.6)\%$. For more details on the measurement see [38].

Kapton cables

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 chosen as they are thin, flexible and radiation hard. The central part of the cable consists of a 17 μm thick copper layer with 150 μm wide strips. This layer is covered on each side with a 100 μm thick kapton foil, a 17 μm thick rolled annealed (RA) 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 the differential signals arriving at the board, including data signals, Time and Fast Control (TFC, see section 8.3) and FE chips configuration 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 and the L0 electronics 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 via a 60 m 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 (c.f. 8.4). Eight radiation hard voltage regulators are mounted on the board. Each voltage is monitored through an amplifier. The card is supplied by three power supplies located at 60 m cable distance. Sense-lines are used to compensate the voltage drop through the long supply cable.

- One ECS card: the ECS repeats the signals for the I2C configuration bus and controls and monitors the LV regulators. The applied voltage and the current limit signals are multiplexed and sent on differential lines. The ECS card communicates with the control boards located in intermediate crates 15 m away from the VELO.

Another part implemented on the board is the TFC functionality. These fast LVDS signals are signals required by the FE chips and are sent through a LVDS repeater mounted on the board. The boards and their components underwent radiation tests to confirm that they could tolerate more than ten times the expected level of 73 kRad during 10 years of operation [39].

Control board

The control board [40] is the heart of the control system of the L0 electronics. The timing signals and fast commands from the Readout Supervisor are distributed to the L0 electronics via the control board. Moreover the configuration data is received by the SPECS slave [41] on the control board and distributed to the configurable components via I2C. The monitored voltages are digitized on the control board and the values are propagated to the ECS system via the SPECS bus. Each control board supports six VELO electronics hybrids or two PU electronics hybrids and four PU optical boards.

Temperature board

The temperatures of the electronics hybrids and voltage regulators are monitored by the temperature boards. The temperature boards also generate signals to the VELO interlock board, which is the last safety mechanism in case of failure (see below). Each temperature board hosts one ELMB (Embedded Local Monitoring Board [9]) and four Interlock Boxes (IBox [42]) , which together serve 16 repeater boards. The ELMB has an embedded Atmel ATMega128 processor and communicates with the supervisory system via a CAN bus. It has digital I/O ports, 64 16-bit ADC channels and a voltage reference circuit. This provides a stand-alone system for temperature monitoring read out via a CAN bus. The IBox supplies a voltage across the NTC resistor and a precision resistor in series. The voltage drop across the NTC resistor is compared to a reference on the IBox and two bits per temperature channel are generated signalling the states *OK*, *too warm*, *too cold* or *error*. The ELMB is in this configuration connected in parallel with the IBox, passively monitoring the voltage drop and hence the temperature. The interlock states of the 64 temperature sensors are read out by one temperature board and combined into a common interlock state of the board. This interlock matrix is implemented in an FPGA and provides the facility to monitor the interlock status and mask faulty channels.

VELO specific TELL1 features

The TELL1 boards of LHCb are described in section 8.2. 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 as solutions for the VELO. As a consequence, the analog data are directly transmitted to the TELL1 board and digitized on *plug-in* cards. Each TELL1 board deals with the data from one sensor, i.e. 64 analog links, and features 4 A-Rx cards. One A-Rx card provides 16 channels of 10-bit ADC's to sample the analog data from 4 Beetle chips at 40 MHz. 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 digitization, the TELL1 performs data processing in the ppFPGAs (pre-processing FPGAs, of which there are four per TELL1) before sending the zero-suppressed data to the trigger farm [43]. The first stages are implemented in 10-bits and the precision is then reduced to 8 bits. The steps in this data processing include pedestal subtraction, cross-talk removal, channel re-ordering, common mode suppression, and clustering. Pedestal subtraction is implemented with a running average pedestal following algorithm available, if required, to calculate the value for each channel. Although most of the signal distortion in the long data cable is removed already through frequency compensation, a Finite Impulse Response Filter is used to correct for cross-talk and further improve the system performance. In both the R- and the ϕ-sensors, adjacent physical strips are scrambled in the readout chain. For the clustering it is essential to bring them back into order by implementing a channel reordering step. Common mode suppression algorithms may be performed in both the readout chip channel ordering and in the physical strip ordering. This allows correlated noise pickup or other baseline level shifts on the sensors or in the front-end chips and readout chain to be corrected. A mean common mode algorithm that corrects for a constant shift and a linear common mode algorithm that corrects for both a constant and a slope are implemented over appropriate groupings of channels. The last processing step is the clustering [44]. Strips are selected as seeding strips if the signal passes a certain seeding threshold. Strips next to the seeding strips are included in the cluster if their signals are above the inclusion threshold cut. A cluster can be formed by a maximum of four strips. The cluster centre is calculated with a 3-bit precision.

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 14-bit cluster position. For more refined calculations of the cluster centre and offline analyses, the ADC information of all strips is added in another data block [45].

The VELO hardware interlock system

The VELO is protected with a simple and failsafe hardware interlock system. An overview of its functionality is given in [46]. Switching on the low voltage when the VELO is not properly cooled is for instance prevented by this system. The status of the cooling, vacuum, motion and detector front-end systems and the BCM (cf. section 3) are combined in an interlock logic unit and fed back to high and low voltage systems and to the motion and cooling system. The interlock unit is based on a FPGA to provide flexibility and possible future interfacing to PVSS (see section 8.4) for monitoring. All input signals are continuously monitored and their status shown on the front panel LED's. The inputs are fail safe such that any disconnection or power loss will result in a bad status and the interlocks will fire. Any of the inputs may additionally be forced to a good status by internal switches for debugging or override purposes, and this is also indicated on the front panel. The interlock outputs are also fail safe such that cable removal or power failure will result in module power and cooling being removed. The status of the outputs is shown on the front panel and all individual outputs can be overidden by internal switches. Hardware signals are also exchanged between motion, cooling and vacuum PLCs to prevent, for example, cooling when there is no vacuum. As the VELO cannot be allowed to move in while the LHC beam is not stable, beam inhibit and beam status signals are exchanged with the BCM and the LHC beam interlock system [20]. Direct signals from the LHC sector valves and the neon injection system are given to the VELO vacuum system to prevent venting, evacuating or neon injection when one of these systems is not ready.

LV system

The VELO and pile-up low voltage system is based on a multi-channel power supply system.[8] All the power supplies are installed in the detector counting house. Each module has 12 fully floating channels and can supply 4 repeater boards. Each channel has its own voltage sense-line to permit a correction to be made for any voltage drop over the long distance cable. Each repeater board requires three voltages which are supplied over shielded twisted-pairs. The LV power supplies interface to the hardware interlock system.

HV system

The 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 biasing scheme ties the n^+ strips to ground, and applies a negative voltage to the backplane. The high voltage system utilises 6 power supply modules,[9] each controlling up to 16 sensors, which are housed in an

[8] Manufactured by CAEN.
[9] Manufactured by ISEG Spezialelektronik GmbH.

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. Three-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 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 interface to the hardware interlock system.

Grounding and power supply

The power distribution and grounding scheme partitioning of the VELO and the pile-up electronics follows the detector topology. Each silicon detector with its hybrid forms a group. The number of groups connected to the same power distribution connection is k ept to a minimum. There are in total 84 VELO and 4 pile-up hybrids with 16 FE chips each installed inside the tank. The tank is the central part of the system and is connected to the LHC machine. The LHC machine earth cables form part of the cavern network grounding. All metallic devices, such as electrical cabinets and cable trays, have ground connection to the main VELO tank support. All electrical devices forming part of the cooling and vacuum system must be grounded for safety reasons. All components connected to the main power network line are equipped with protection circuits and electrically floating devices made of conductive material are not permitted.

The analog FE electronics is potentially prone to pickup from external noise sources. The effects are minimised by keeping the signal current path as short as practicable. The reference voltage of the charge preamplifier in the FE chip is internally connected to the ground pins of the chip. The current loop for the signal generated in the sensor is closed through the sensor bias line. The bias line is AC-connected to the reference ground of the FE chip. The silicon detectors and hybrids are located close to the detector RF shield, which is connected to the ground. Potential differences between these two components give rise to the main source of noise in the detector. This is minimised by keeping the hybrid ground plane at the same potential as the RF shield. Each hybrid ground plane is tied to the module support with a silver-plated copper grounding strap.

5.1.5 Material budget

The material budget was investigated in the simulation by generating straight tracks originating from the interaction point and extrapolating them through the VELO. The radiation length seen by the track was calculated for each volume it crossed in the geometry description. The average radiation length is shown as a function of the azimuthal angle ϕ and the pseudorapidity η, with $1.6 < \eta < 4.9$ in figure 5.12. The contribution of the major components is also shown. The average radiation length of the VELO is 17.5% of a radiation length with the largest contribution coming from the RF-foil.

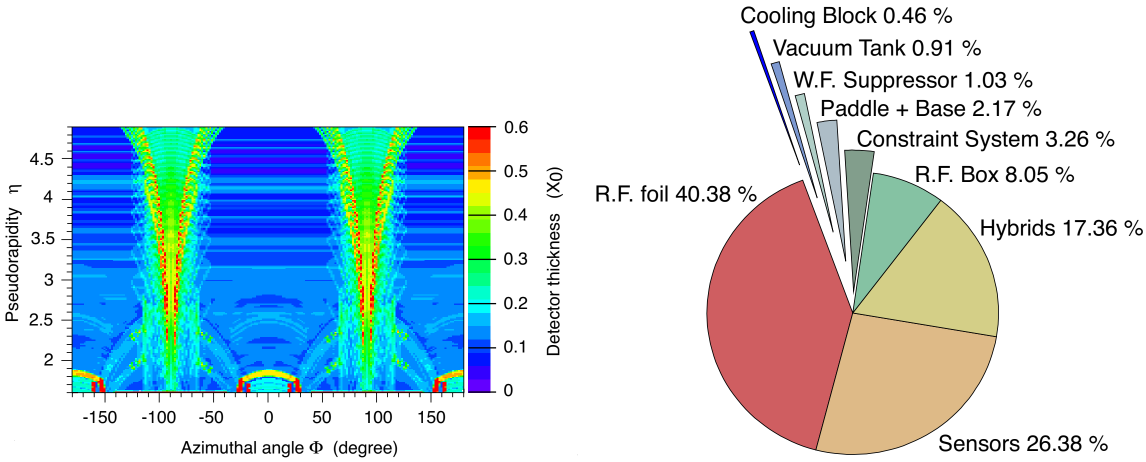

Figure 5.12: The average radiation length seen by particles passing through the VELO as a function of azimuthal angle, ϕ, and pseudorapidity, η, with 1.6< η <4.9. For a uniform sampling over these coordinates the average contribution to the radiation length is 0.175 X_0.

5.1.6 Test Beam detector commissioning

Test Beam 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 events produced by interaction in the detectors or on the entrance window of their enclosure, and, in the second data taking cycle, interactions in target 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°C. Only 1/4 of the electronics was read out, due to limitations in the available hardware. This data set provided considerable insight into the performance of the final detector modules. Figure 5.13 shows that the measured cluster multiplicity as a function of the strip number is in good agreement with the expectations based on a dedicated MonteCarlo simulation [47]. Based on Monte-Carlo simulation and the performance of the electronics a resolution of about 6 μm can be achieved at the inner radius of the sensors. This optimal resolution will be obtained by utilizing a non-linear charge-sharing algorithm that depends on the incident angle of the tracks (for the R-sensor) and the irradiation level of the sensors. All the parameters affecting the performance of the sensors and the readout electronics were explored, for example data was taken with the sensors biased with different high

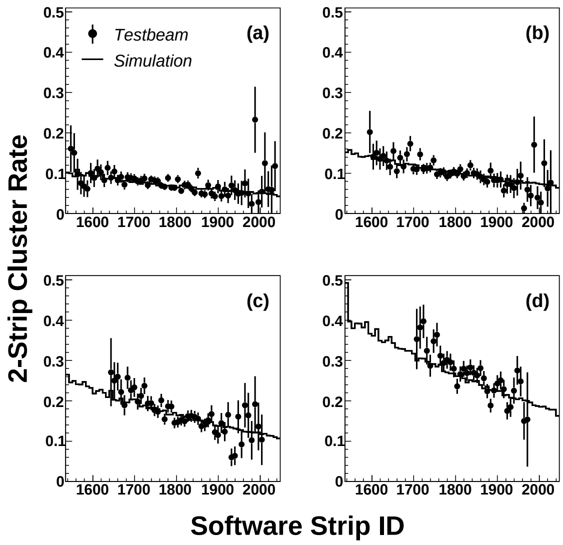

Figure 5.13: Fraction of two strip clusters in R-sensors, for different beam track angles with respect to the normal to the detector planes: a) 0°, b) 5°, c) 10° d) 15°. The horizontal axis represents the strip number, corresponding to a progressively coarser inter-strip pitch.

voltages, and the operating parameters of the Beetle chip were changed. In addition, the noise performance was studied 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. 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 taken 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 (MIP) and noise is the incoherent noise measured from calibration data after subtracting the coherent noise component. Their performance proved to be remarkably stable throughout the 14 days of data taking. Single track and interaction trigger data complement the first data set in assessing the hit resolution as a function of angle. In addition, target data produced reconstructed primary vertices from all the targets that help in tuning the vertex reconstruction algorithms.

2008 JINST 3 S08005

5.1.7 VELO software

Reconstruction software

The TELL1 data processing boards (see sections 5.1.4 and 8.2) perform a set of processing algorithms on the raw VELO data and identify clusters 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 [48]. 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 non zero-suppressed 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 algorithm 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, such as the signal thresholds used in the clustering algorithm.

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 analog signals are passed through the emulated TELL1 clustering algorithm and the output clusters are stored in the same format as for the real data [49].

Monitoring

The VELO data monitoring 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 on-line monitoring have already been performed in the VELO test beam.

A critical element of the 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 two halves and centre

2008 JINST 3 S08005

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 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 relies 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 silicon sensors are nominally located only 1 mm 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, used to identify high impact parameter tracks in the trigger, is performed for speed reasons initially in the rz projection. This requires that the strips on the R-sensors accurately describe circles around the beam position.

A survey of the individual modules and of the assembled halves on the module supports was performed. The relative positions of the modules and of the R- and ϕ-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 deviations of the modules from their nominal positions are shown in figure 5.14 for the two detector halves (A and C sides). The scatters give an impression of the total assembly precision. In the x coordinate the scatter is very small, and the overall offsets are fully compensated by the fact that the two halves move independently in x. In y there is a small slope seen in the C-side results, which was traced to the fact that the module support was not perfectly flat, but had a 40 μm corkscrew twist between the corners. The resulting slope was fully compensated by adjustments of the detector supports during the installation of the halves.

The precision survey is an important element of the 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 will not be possible to accurately measure with data. For example, the overall z positions of the sensors are 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 support will be maintained at a constant temperature of 20°C. The deviations of the system under vacuum have been determined from the software

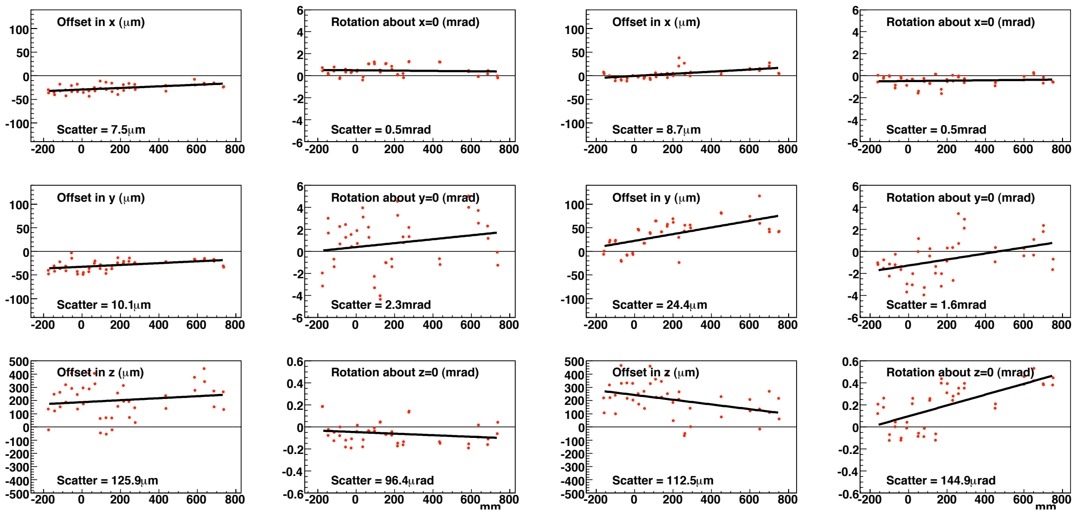

Figure 5.14: Results of the metrology of the fully mounted VELO halves vs z. The deviations from nominal are shown for every sensor. The six plots on the left show the result for the A side, and on the right for the C side. The rz tracking is sensitive to scatters and offsets in the x coordinate (between modules) and y coordinate (between modules, and between the detector halves).

alignment in a test beam using a partially assembled VELO and seen to be typically 10 μm or less, as shown in figure 5.15.

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 two 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 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 software alignment at the start of each fill remains, should this prove necessary.

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

- An alignment of the relative position of the R and Φ sensors on each module using the shape of the residual distributions across the sensors.
- An alignment of the modules within each VELO-half box using the residuals of hits on reconstructed tracks.
- A relative alignment of the two half-boxes with respect to each other principally relying on using tracks passing through the geometrical overlap between the modules in the two half-boxes.
- A global alignment of all sub-detectors relative to each other. This part is reviewed in [50].

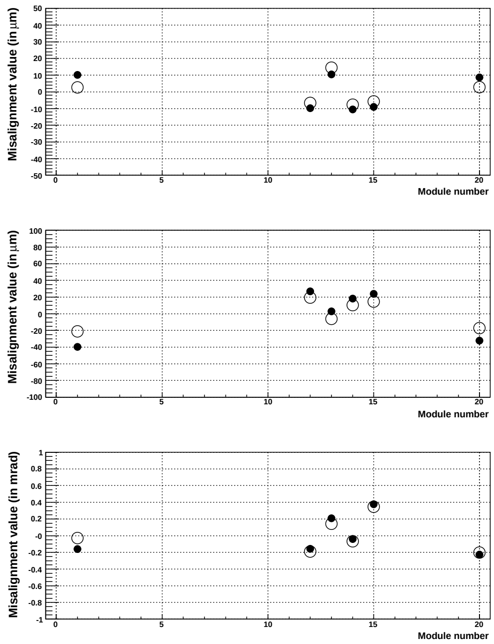

Figure 5.15: x-axis translation alignment constants of the VELO modules determined by software alignment from data recorded during the VELO test beam in air (open circles) and in vacuum (solid points). The modules are seen to be stable in the different conditions, including different thermal gradients across the modules.

Relative sensor alignment

The R- and ϕ-sensors bonded in an individual module have had their relative positions measured during the system metrology. To further improve and monitor their relative misalignment a method has been developed that exploits distortions observed in the residuals across the sensor. It can be shown that the relation between residuals ($\varepsilon_{R/\Phi}$) and misalignments in the sensor plane ($\Delta_x, \Delta_y, \Delta_\gamma$) is given by

$$\left\{ \begin{array}{lllllr} \varepsilon_R & = & -\Delta_x \cos\phi_{\text{track}} & +\Delta_y \sin\phi_{\text{track}} & & \text{(R sensor)} \\ \varepsilon_\Phi & = & +\Delta_x \sin\phi_{\text{track}} & +\Delta_y \cos\phi_{\text{track}} & +\Delta_\gamma r_{\text{track}} & (\Phi \text{ sensor}) \end{array} \right. , \tag{5.1}$$

where r_{track} and ϕ_{track} are the radius and the azimuthal angle of the extrapolated track position, respectively. The misalignments Δ_x, Δ_y are determined by a fit to the shape of the residual distribution as a function of the azimuthal angle, while Δ_γ, the rotation misalignment around the z-axis, is determined by a fit to the shape of the residual distribution as a function of the radius. The optimal alignment can be found after three iterations of this procedure, using the previously determined alignment constants as input to the next iteration.

The method has been tested using MonteCarlo simulations and VELO test beam data. It has been shown that an alignment precision at the μm level can be achieved for Δ_x and Δ_y, while Δ_γ will be determined to higher precision by the module alignment algorithm.

Module alignment and VELO half alignment

The module alignment and the VELO half 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. 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:

$$n_{\mathrm{total}} = n_{\mathrm{local}} \times n_{\mathrm{tracks}} + n_{\mathrm{global}} \tag{5.2}$$

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 $n_{\mathrm{global}} \times n_{\mathrm{global}}$ matrix inversion. Inversion by partitioning is handled by the Millepede program [51].

The number of tracks required for an effective alignment is relatively modest but the alignment is improved by using a mixture of tracks from primary vertex interactions and a complementary track set from a source such as beam-halo particles. The relative alignment of the two half-boxes is primarily constrained by tracks that pass through both halves. However, when the VELO is retracted an alternative technique is required which relies upon fitting primary vertices using tracks fitted in both halves of the VELO. The CPU requirements of the alignment are also low: of order minutes on a single PC.

Tests using data from the VELO test beam and on MonteCarlo 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. MonteCarlo 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.

5.1.8 VELO performance

The VELO layout has been optimised to minimise the amount of material in the acceptance 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 5.16.

The individual hit resolution of the sensors has been determined in a test beam and is a strong function of the sensor pitch and projected angle (the angle perpendicular to the strip direction), as shown in figure 5.17. The best raw resolution obtained is 7 μm. As shown in figure 5.18 perfor-

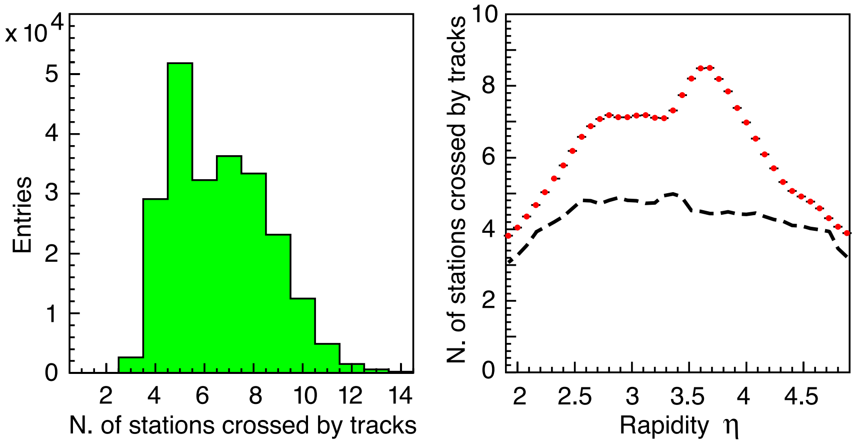

Figure 5.16: The left plot shows the number of stations hit per track in the VELO and the right plot shows the number of hits of a track in the VELO modules as a function of the pseudorapidity of the track. The dashed line indicates the limit above which 95% of the tracks lie.

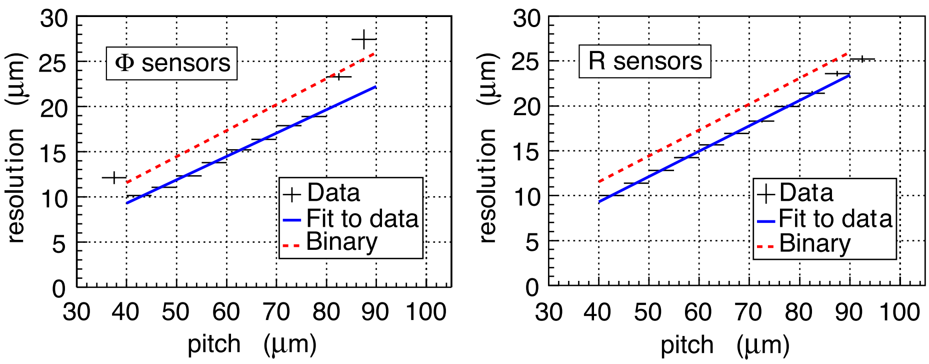

Figure 5.17: The raw hit resolution as a function of strip pitch as measured in the test beam for particles of normal incidence. The dashed line indicates the resolution expected for digital readout. The data points show the resolution as measured from the weighted centre of the charges on the strips.

mance improves for the low angle tracks when imperfections in the weighted charge distribution between two strips are taken into account [52].

In addition, crosstalk originating from inter-strip coupling, from coupling between electronic channels, and from signal feed-forward and backward in the analog transmission have not been taken into account. Once these have been fully parametrised further improvement in the resolution obtained from the system is anticipated.

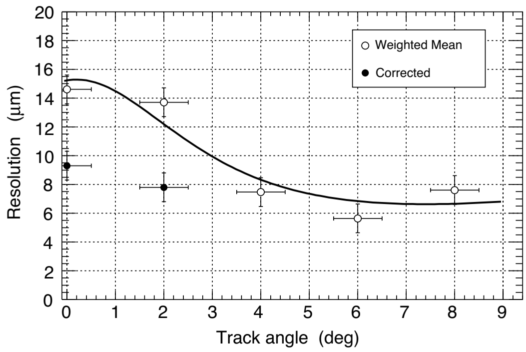

Figure 5.18: The individual hit resolution of a VELO R-sensor as a function of the projected angle for a pitch of 85 μm. The open circles show the resolution as derived from a weighted mean. The filled circles include the corrections for the imperfections in the weighted charge distribution between two strips.

5.2 Silicon Tracker

The Silicon Tracker (ST) comprises two detectors: the Tracker Turicensis[10] (TT) [2, 53] and the Inner Tracker (IT) [54]. Both TT and IT use silicon microstrip sensors with a strip pitch of about 200 μm. The TT is a 150 cm wide and 130 cm high planar tracking station that is located upstream of the LHCb dipole magnet and covers the full acceptance of the experiment. The IT covers a 120 cm wide and 40 cm high cross shaped region in the centre of the three tracking stations downstream of the 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 has an active area of about 8.4 m^2 with 143360 readout strips of up to 38 cm in length. The IT has an active area of 4.0 m^2 with 129024 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. Simulation studies have demonstrated that a single-hit resolution of about 50 μm is adequate for both the TT and the IT. 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. Charged particle densities of about 5×10^{-2} per cm^2 for minimum bias events are expected in the innermost regions of the TT. They fall off by two orders of magnitude to about 5×10^{-4} per cm^2 in the outermost regions of the detector. 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. For the IT, expected charged-particle densities for minimum-bias events range from about 1.5×10^{-2} per cm^2 close to the LHC beampipe to 2×10^{-3} per cm^2 in the outer regions of the detector. Also here, the chosen strip geometries result in maximum strip occupancies that do not exceed a few percent.

[10]The Tracker Turicensis was formerly known as the Trigger Tracker.

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. The benchmark parameter is the remaining fraction of the signal amplitude at the sampling time of the subsequent bunch crossing, 25 ns after the maximum of the pulse. Simulation studies have shown that signal remainders of 50% for the TT and 30% for the 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 radiation damage corresponding to ten years of operation at nominal luminosity.

Radiation damage. For ten years of operation at nominal luminosity, expected 1 MeV neutron equivalent fluences in the innermost regions of the detectors do not exceed 5×10^{14} per cm^2 for the TT and 9×10^{12} per cm^2 for the IT. Basic design rules for radiation hard silicon sensors were followed to ensure that the detectors will survive these fluences. The detector was designed and tested to work at bias voltages of up to 500 V. The sensors need to be operated at a temperature of 5°C or lower to suppress radiation damage induced leakage currents to a level where (a) shot noise does not significantly deteriorate the signal-to-noise performance of the detector and (b) the risk of thermal runaway due to the power dissipated by the leakage currents is avoided.

Material budget. As the momentum resolution of the LHCb spectrometer is dominated by multiple scattering, the material budget of the detectors had to be kept as small as possible. The TT was designed such that front-end readout electronics and mechanical supports are located outside of the LHCb acceptance. This was not possible in the case of the IT, which is located in front of the active region of the OT detectors. Here, a significant design effort was made to keep the amount of material for mechanical supports and for the cooling of front-end electronics as small as possible.

Number of readout channels. Readout electronics are 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 minimize the number of readout channels.

Different constraints on the detector geometries resulted in different designs for the detector modules and station mechanics of the TT and the IT. These are described in sections 5.2.1 and 5.2.2, 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 section 5.2.3. Finally, the expected detector performance, based on test beam measurements and simulation studies, is discussed in section 5.2.4.

5.2.1 Tracker Turicensis

All four detection layers of the TT are housed in one large light tight and thermally and electrically insulated detector volume, in which a temperature below 5°C is maintained [55]. The detector

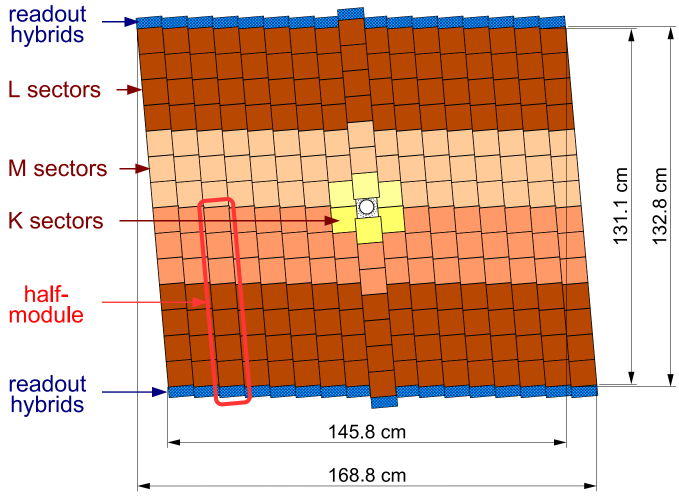

Figure 5.19: Layout of the third TT detection layer. Different readout sectors are indicated by different shadings.

volume is continuously flushed with nitrogen to avoid condensation on the cold surfaces. 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 5.19. Its basic building block is a half module that covers half the height of the LHCb acceptance. It consists of a row of seven silicon sensors organized into either two or three readout sectors. The readout hybrids for all readout sectors are mounted at one end of the module. The regions above and below the LHC beampipe are covered by one such half module each. The regions to the sides of the beampipe 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. These full modules cover the full height of the LHCb acceptance and are assembled from two half modules that are joined 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 u and v detection layers, each module is individually rotated by the respective stereo angle.

A main advantage of this detector design is that all front-end hybrids and the infrastructure for cooling and module supports are located above and below the active area of the detector, outside of the acceptance of the experiment.

TT detector modules

The layout of a half module is illustrated in figure 5.20. 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 beampipe, where the expected particle density is highest, the seven sensors are organized into three readout sectors (4-2-1 type half modules).

For the other half modules, the sensors are organized 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

2008 JINST 3 S08005

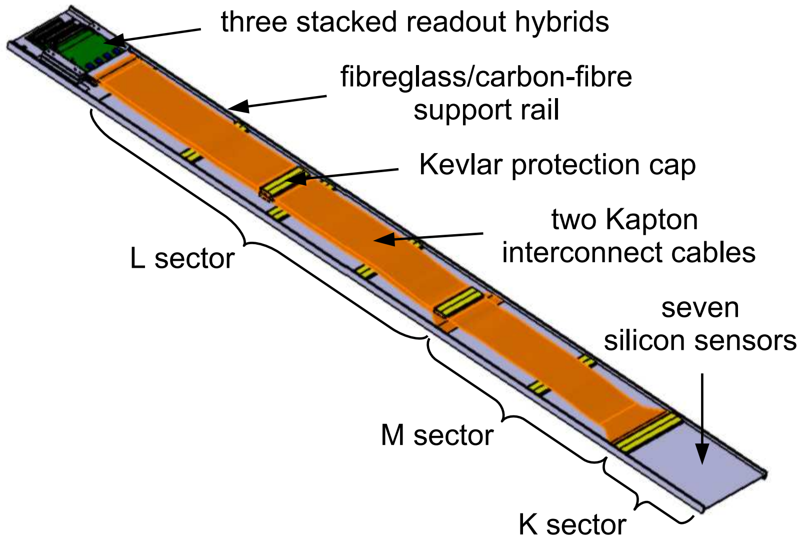

Figure 5.20: View of a 4-2-1 type TT detector module.

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 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 consisting 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. The half module is mechanically held together by two thin fibreglass/carbon fibre rails that are glued along the edges of the L hybrid and the seven silicon sensors.

Silicon sensors. The silicon sensors for the TT are 500 μm thick, single sided p^+-on-n sensors.[11] They are 9.64 cm wide and 9.44 cm long and carry 512 readout strips with a strip pitch of 183 μm.

Kapton interconnect cables. The Kapton interconnect cables for the M and K readout sectors were produced using standard plasma-etching technology.[12] 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. A short pitch-adapter 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. The small strip width was required to keep the strip capacitance of the cable small, but led to an unacceptably low production yield for fault

[11] The sensors are identical in design to the OB2 sensors used in the Outer Barrel of the CMS Silicon Tracker [56] and were produced by Hamamatsu Photonics K.K., Hamamatsu City, Japan.

[12] The Kapton interconnect cables were produced by Dyconex AG, Bassersdorf, Switzerland.

2008 JINST 3 S08005

free cables of the required length. The 39 cm long cables for the M sectors therefore had to be assembled from two shorter pieces, and the 58 cm long cables for the K sectors were assembled from three pieces. The pieces were joined together end-to-end by gluing them onto a thin strip of fibreglass reinforced epoxy. An electrically conductive adhesive tape was used to provide the electrical connection between the copper mesh backplanes of the two cable pieces. The signal, bias voltage and ground strips on the strip side of the cables were joined together by wire bonds.

Kevlar caps. Small Kevlar caps protect the wire bond rows on the strip side of the Kapton interconnect cables as well as those between silicon sensors. These caps are glued onto the surface of the cable or the sensors using standard two component epoxy glue.[13]

Readout hybrids. The front-end readout hybrids [57] consist of a carrier plate, a pitch adapter, and a four layer Kapton flex circuit that carries four front-end readout chips, some passive SMD components and an 80-pin board-to-board connector.[14] Through this connector, the multiplexed analog detector signals are read out and the control signals, low voltage and bias voltage are provided to the half module. Two of the four conductive layers of the flex circuit are used for digital and analog power and ground. The other two layers 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 upper hybrids have to be shorter than the lower ones to make the readout connectors on all hybrids accessible when these 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 front-end chips. The pitch adapter matches the input pitch of the front-end 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 these, the carrier plate is made from gold plated copper and the pitch adapter is a rectangular piece of alumina (Al_2O_3) substrate with strip lines produced using standard thin film technology. The pitch adapter is glued onto the carrier plate together with the Kapton flex circuit. The carrier plate for the L hybrid is a much more complicated piece that consists of aluminium nitride substrate and combines several functionalities.[15] First of all, it carries not only the strip lines of the pitch adapter for the L sector, produced using standard thick film technology, but also additional traces and vias that serve to connect the bias voltage for all readout sectors. Next, the support rails of the detector module are glued along the sides of this carrier plate. And finally, this carrier plate provides the mechanical and thermal interface of the half module to the detector station. All these functionalities will be described in more detail below. Aluminium nitride was chosen as the material for this carrier plate because of its high thermal conductivity and its small thermal expansion coefficient that is reasonably well matched to that of the silicon sensors.

The K and M hybrids are mounted on top of the carrier plate of the L hybrid using 2.5 mm thick spacers made of copper. These ensure good thermal contact between the carrier plates of the K and M hybrids and the one of the L hybrid.

Support rails. The two half module support rails are 5.5 mm high and consist of a 1 mm thick strip of carbon fibre glued to a 1 mm thick strip of fibreglass reinforced epoxy. A small groove

[13]Araldite 2011 (AW106/HV953U), by Huntsman Advanced Materials, Basel, Switzerland.

[14]The flex circuits were produced by Optiprint AG, Berneck, Switzerland.

[15]The production of the pitch adapters for the K and M hybrids and of the aluminium nitride carrier plate for the L hybrids as well as the assembly and bonding of all hybrids was done by RHe Microsystems, Radeberg, Germany.

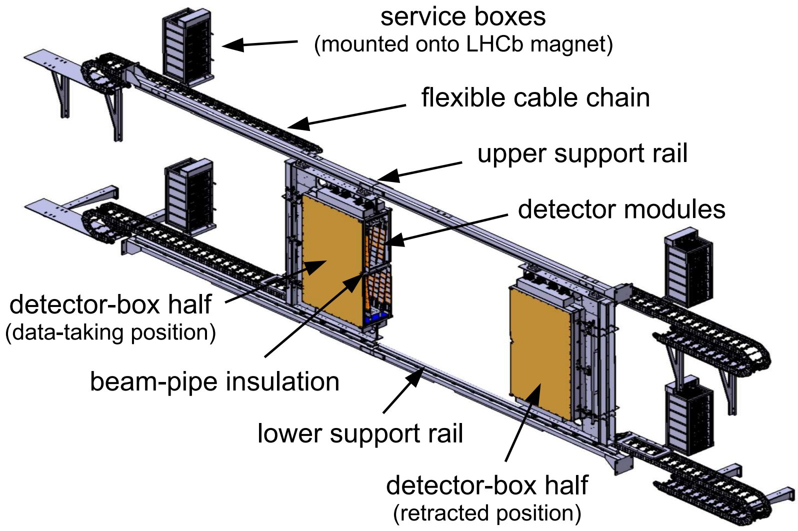

Figure 5.21: View of the TT station mechanics.

is milled into the flat side of the fibreglass strip which permits to slide the rail over the edges of the seven silicon sensors and the carrier plate of the L hybrid. The rail is glued to these using standard two component epoxy glue. The rôle of the fibreglass strip is to ensure the necessary electrical insulation between the strip side of the silicon sensors (which are at ground potential) and their backplanes (which carry the bias voltage of up to 500 V). The mechanical rigidity of the rail is defined mainly by the carbon fibre strip. The fibreglass strip spans the full length of the half module but the carbon fibre strip ends at the fourth silicon sensor from the hybrids. This permits to join two half modules to a 14-sensor long full module by gluing an additional carbon fibre strip to the free sections of fibreglass strip on both half modules.

Bias voltage. Bias voltage is supplied separately for each of the readout sectors on a half module via the cable that plugs into the corresponding front-end hybrid. From the hybrid it is connected to aluminium traces on the carrier plate of the L sector, using wire bonds in case of the L hybrid and thin copper wires in case of the M and K hybrids. From here, all bias voltages are brought to the back of the half module through aluminium vias that are embedded in the aluminium nitride 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, carries the bias voltage to the backplanes of the silicon sensors. The electrical connections between the Kapton flex cable and the sensor backplanes and between the cable and the bias voltage pads on the back of the aluminium nitride are made by wire bonding.

TT detector station

An isometric drawing of the detector station is shown in figure 5.21. It consists of two halves, one on each side of the LHC beampipe.

The half stations are mounted on rails and can be retracted horizontally for detector maintenance and bakeouts of the LHC beampipe. The main structural element of each half station is a C-shaped aluminium frame. It carries a detector box that is made of light weight aluminium clad foam. This box is open on the side facing the beampipe such that the two half stations in data taking position define one large volume that contains all detector modules. Mounted against the top and bottom walls of the detector box are horizontal cooling plates that provide the mechanical support for the detector modules as well as cooling of the front-end hybrids and the detector volume. The cooling plates 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 outer side walls of the detector volume. All electrical signals (detector signals, control signals and supply voltages) are transmitted on Kapton-flex cables through specially designed feedthroughs in the top and bottom walls of the detector box.

Detector box. The C-shaped support frames are assembled from 15 mm thick aluminium plates. They rest on a lower precision rail and are guided by an upper precision rail. These two rails are aligned parallel to each other to better than 100 μm in order to avoid possible distortions of the C-frame during insertion or retraction. Mounted flat against the inner surfaces of the C-frames are stiff sandwich plates consisting of a 30 mm thick aramid honeycomb structure with 1 mm thick aluminium cladding. These sandwich plates define the outer walls of the detector box. The front and rear walls of the box are made of 40 mm thick panels of polyetherimide foam[16] that are laminated on both sides with 25 μm of aluminium for electrical insulation and 275 μm of Kevlar for mechanical protection. These panels are screwed against the outer walls of the box and can be easily removed for the installation and maintenance of detector modules. Similar plates of 40 mm thick aluminium and Kevlar clad polyetherimide foam are also mounted flat against the inner surfaces of the outer walls of the box to improve thermal insulation here. Finally, the detector volume is closed off around the LHC beampipe by two specially machined semicylindrical pieces that again consist of the same polyetherimide foam clad with aluminium and Kevlar. The wall thickness of these beampipe insulation pieces is 30 mm, except for cutouts at the locations of the detection layers. Here, the wall thickness is reduced to about 5 mm to reduce the dead space in between the beampipe and the innermost detector modules. A clearance of 5 mm between the detector box and the beampipe had to be maintained to satisfy LHC safety demands.

Cooling plates. There are a total of four cooling plates, one for each quadrant of the detector. They are mounted horizontally onto pillars made of polyacetal (POM) that are fixed against the upper or the lower walls of the detector box. A drawing of a cooling plate is shown in figure 5.22. It is a precisely machined plate of 8 mm thick aluminium that measures 897 mm in x and 348 mm in z. Machined into its outer surface are semicircular 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 is machined to an overall flatness of better than 100 μm.

Cooling balconies. Mounted vertically against the flat inner surface of the cooling plate are the aluminium cooling balconies that provide the mechanical, thermal and electrical interface between cooling plate and detector modules. Precision holes and pins ensure the accurate positioning of the balconies on the cooling plate. The detector 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 aluminium nitride carrier plate of the half module. The correct positioning of the module is ensured by precision holes and pins in the balcony and laser cut holes in the aluminium

[16] Airex R82.60, by Gaugler & Lutz oHG, Aalen-Ebnat, Germany.

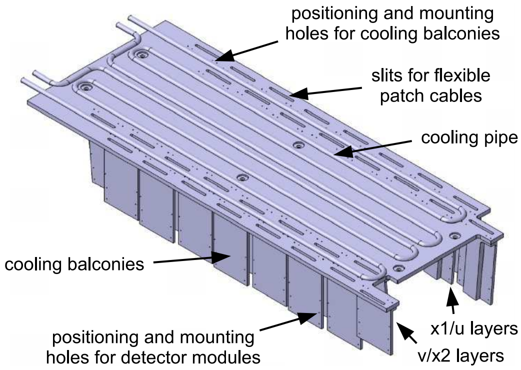

Figure 5.22: View of one TT cooling plate with mounted cooling balconies.

nitride carrier plate. There are two types of balconies, one for mounting modules vertically and one for mounting modules under an angle of 5°. 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.

Cooling elements. 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.

Electrical connections. Detector signals are read out and control signals, low voltage and bias voltage supplied to the detector modules via 50 cm long Kapton flex cables. These cables pass through dedicated slits in the cooling plate and through specially designed feedthroughs in the top and bottom walls of the detector box. There is a separate Kapton flex cable for each 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 a patch panel that is mounted on the outside of the detector box. From this patch panel, shielded twisted-pair cables lead through flexible cable chains to service boxes that are mounted against the front face of the LHCb dipole magnet. Here, the signals are prepared for digital optical transmission to the counting house as described in section 5.2.3. Each quadrant of the detector has a flexible cable chain and a stack of six service boxes.

5.2.2 Inner Tracker

Each of the three IT stations consists of four individual detector boxes that are arranged around the beampipe as shown in figure 5.23.

The detector boxes are light tight and electrically and thermally insulated, and a temperature below 5°C is maintained inside them. They are continuously flushed with nitrogen to avoid condensation on the cold surfaces. 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. Detector modules in the boxes above and below the beampipe (top and bottom boxes) consist of a single silicon sensor and a readout hybrid. Detector modules in the boxes to the left and right of the beampipe (side boxes) consist of two silicon sensors and a readout hybrid. The resulting layout and dimensions of one of the IT detection layers are illustrated in figure 5.24.

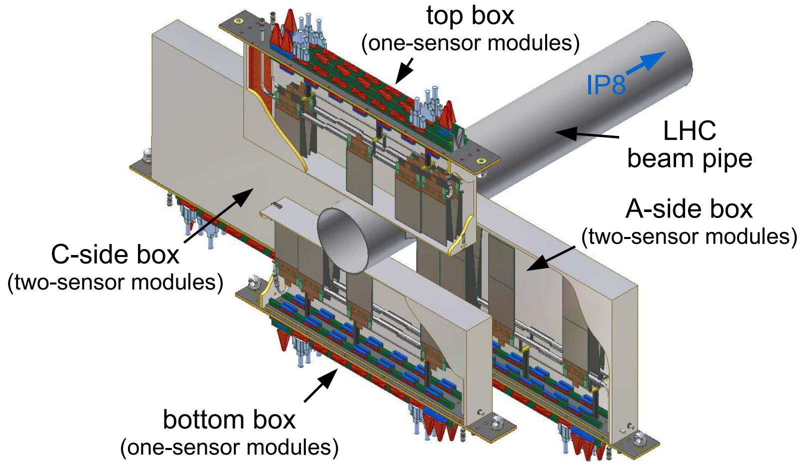

Figure 5.23: View of the four IT detector boxes arranged around the LHC beampipe.

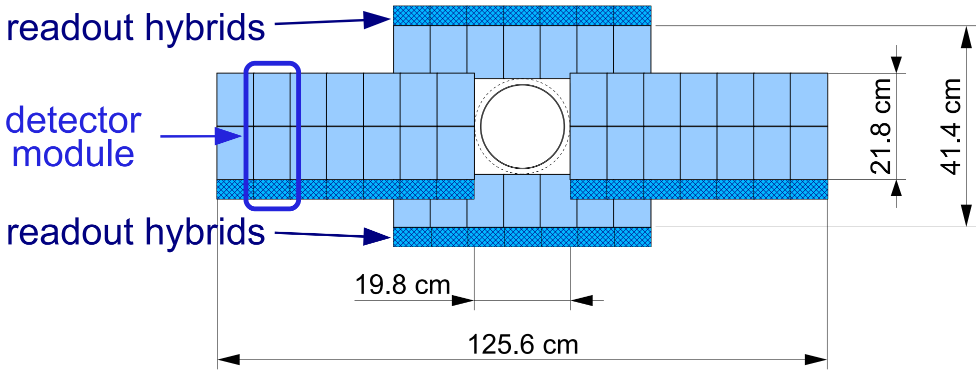

Figure 5.24: Layout of an *x* detection layer in the second IT station.

IT detector modules

An exploded view of a detector module is shown in figure 5.25. The 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 all glued onto a flat module support plate. Bias voltage is provided to the sensor backplane from the strip side through n^+ wells that are implanted in the n-type silicon bulk. A small aluminium insert (minibalcony) that is embedded into the support plate at the location of the readout hybrid provides the mechanical and thermal interface of the module to the detector box.

Silicon sensors. Two types of silicon sensors of different thickness, but otherwise identical in design, are used in the IT.[17] 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 thick, those for two-sensor modules are 410 μm thick. As explained in section 5.2.4 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.

[17]The sensors were designed and produced by Hamamatsu Photonics K.K., Hamamatsu City, Japan.

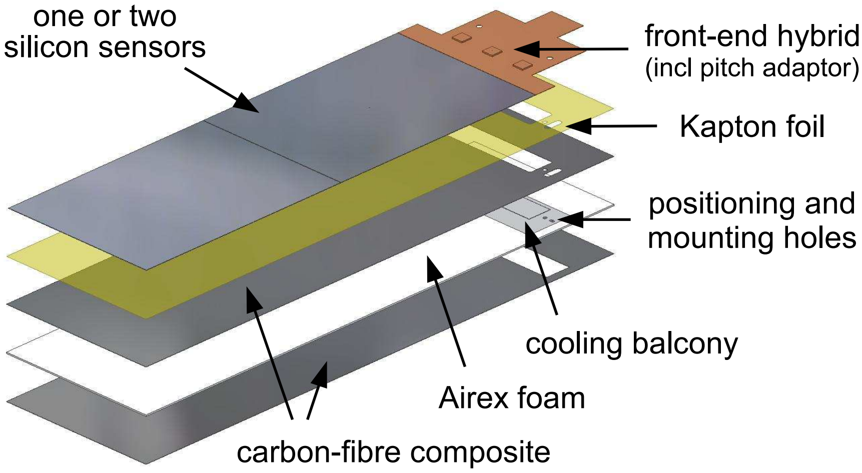

Figure 5.25: Exploded view of a two-sensor IT module. One-sensor modules are similar except that the support plate is shorter and carries only one sensor.

Readout hybrids. 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 adapter that is similar to that used for the M and K hybrids of the TT.[18] The only differences are (a) that the Kapton flex circuit for the IT carries only three front-end chips and (b) that it incorporates 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 adapter is glued onto the Kapton flex circuit and the Kapton flex circuit is glued directly onto the module support plate.

Module support plate. The module support plate[19] consists of a 1 mm thick sheet of polyetherimide foam[20] sandwiched in between two 200 μm thick layers of carbon fibre composite. The latter are produced from two layers of thermally highly conductive carbon fibres[21] 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 on all sides by 1 mm over the edges of the silicon sensors to protect these mechanically. The edges of the support plate are sealed with a non conductive epoxide resin to prevent loose fibres from sticking out and touching the sensor, where they might cause a short between the strip side and the backplane. At the location of the front-end readout chips, a 70 mm wide, 15 mm high and 1.5 mm thick aluminium piece called minibalcony is inserted in the support plate. It is glued to the support plate using thermally conductive glue.[22]

Minibalcony. The minibalcony defines the mechanical and thermal interface of the module to the detector box. It is machined to a flatness of better than 30 μm and contains precision holes for

[18]The production of flex prints and pitch adapters and the assembly of the readout hybrids took place at the same two companies as for the TT.

[19]The module support plates were produced by Composite Design, Echandens, Switzerland.

[20]Airex R82.60, by Gaugler & Lutz oHG, Aalen-Ebnat, Germany.

[21]K13D2U, by Mitsubishi Chemical Corp., Tokyo, Japan.

[22]TRA-DUCT 2902, by TRA-CON Inc., Bedford, Massachusetts, USA.

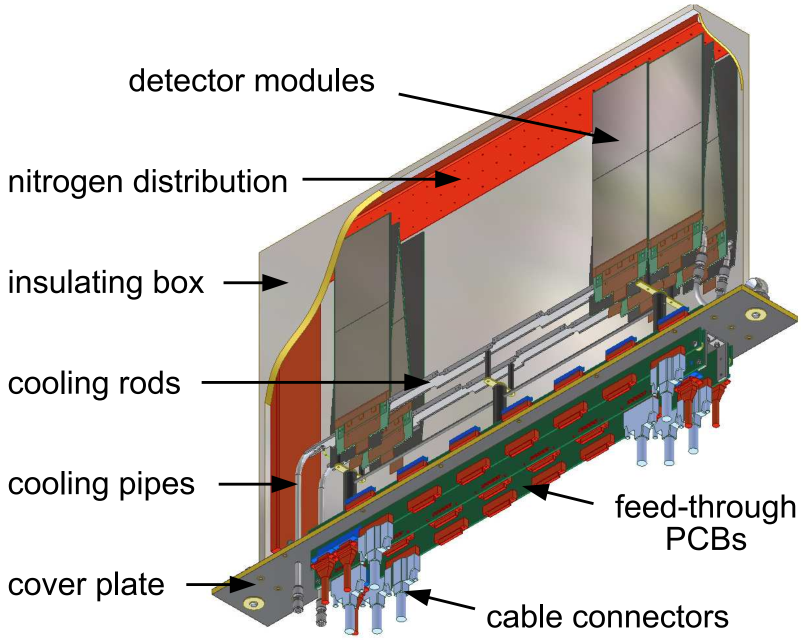

Figure 5.26: View of an IT side box. Top/bottom boxes are similar except that the box is shorter and contains one-sensor modules.

the mounting and exact positioning of the module. There are two types of minibalconies: one for modules that will be mounted vertically and another one for modules that will be mounted with the 5° stereo angle. The minibalcony provides a direct heat path from the front-end chips to the cooling rod, onto which the modules are mounted as described below. It also thermally connects the carbon fibre sheets of the module support plate to the cooling rod. Due to their good thermal conductivity along the module axis and their large surface area, the module support plates therefore contribute to the cooling of the silicon sensors and the detector volume.

The silicon sensors are glued onto the module support plate using strips of silicone glue.[23] The hybrids are glued onto the module support plate using standard two component epoxy glue. Small spots of thermally conductive glue[24] are applied at the location of the front-end readout chips in order to improve the thermal contact between the chips and the minibalcony.

IT detector boxes

An isometric view of a detector box is shown in figure 5.26. Its main structural element is a 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 that are inserted vertically through the cover plate serve to transmit supply voltage and detector and control signals 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.

[23] NEE-001-weiss by Dr. Neumann Peltier Technik, Utting, Germany; for a characterisation of this glue, see [58].
[24] EpoTek 129-4, by Epoxy Technology, Billerica, Massachusetts, USA.

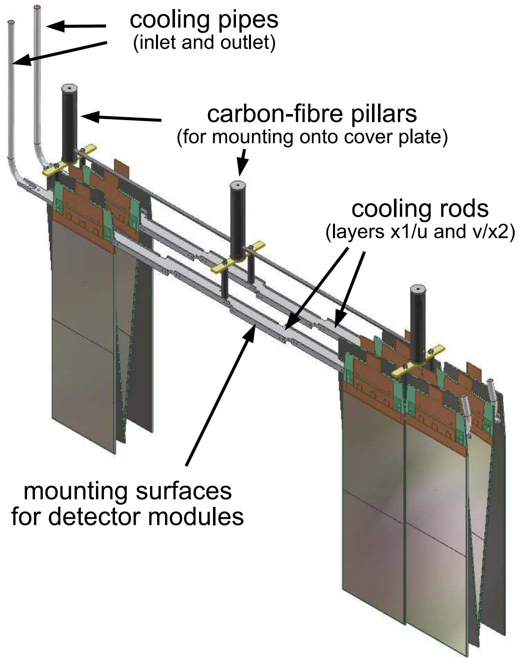

Figure 5.27: View of an IT cooling rod with a few detector modules.

Cover plate. The cover plate is made from a a 14 mm thick polymethacrylimide foam (Rohacell) sandwiched in between two layers of carbon fibre composite. Four printed-circuit boards, one for each detection layer, are inserted vertically through slits in the cover plate. They have four copper layers and serve to feed supply voltage and detector and control signals through the box. The outer layers carry the detector bias voltage and the analog and digital supply voltage, respectively. The two inners layers are used for the differential signals. Inside the detector box, the printed-circuit boards carry 60-pin board-to-board connectors into which the Kapton tails of the readout hybrids are plugged. Outside the box they are equipped with low mass connectors for signal cables and bias voltage cables. Cooling and nitrogen supply pipes also pass through the cover plate. Finally, the cover plate contains mounting holes for fixing the detector box on the IT support frames described below.

Cooling rods. An isometric view of the two cooling rods is shown in figure 5.27. They are mounted on the cover plate using pillars made out of a carbon fibre composite. Each cooling rod is machined out of a single piece of aluminium. It consists of a 3 mm thick central part and vertical mounting surfaces for each of the detector modules. An aluminium cooling pipe with an outer diameter of 6 mm and a wall thickness of 0.4 mm is glued into a semicircular groove that runs along the central part of the cooling rod. The mounting surfaces for the detector modules are 6 mm high and 70 mm wide and contain precision holes and pins to ensure the accurate positioning of the modules. Detector modules are mounted on both sides of the cooling rod, such that each cooling rod supports a pair of detection layers. The cooling pipes on the two cooling rods are connected in series using a short nitrile rubber hose.

Box enclosure. The box enclosure is assembled from 6 mm thick, flat sheets of polyisocyanurat (PIR) foam reinforced with a single, 200 μm thick carbon fibre skin and clad on both sides

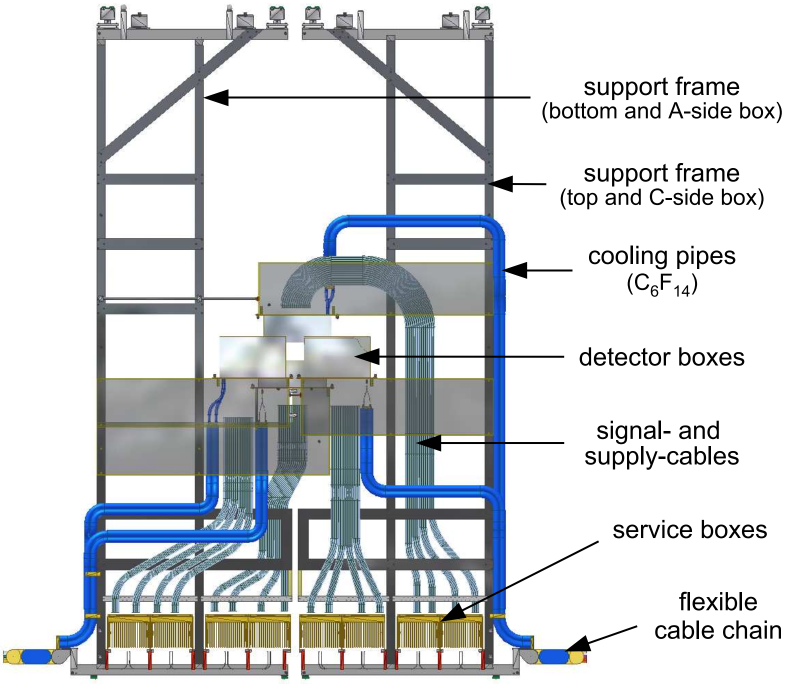

Figure 5.28: Front view of an IT detector station.

with 25 μm thick aluminium foil. For the side of the box facing the beampipe, the wall thickness is reduced to 2 mm to decrease 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 of fibreglass reinforced epoxy are embedded in the upper rim of the enclosure and permit to screw it onto the cover plate.

IT detector station

A front view of a detector station is shown in figure 5.28.

The detector boxes are mounted onto two support frames that are mounted on rails and can be retracted horizontally for detector maintenance and bakeouts of the LHC beampipe. The support frames are suspended from the upper support rail, which is mounted onto the Outer Tracker bridge (see section 5.3) and are guided by the lower rail that is mounted onto the LHCb bunker. The innermost sections of both support rails are precision machined to ensure an accurate positioning of the support frames in data taking position. The support frames are assembled from rectangular rods made of fibreglass reinforced epoxy and carbon fibre composite and from flat plates of aramide honeycomb clad with skins of a carbon fibre composite. Signal and bias voltage cables and flexible supply lines for C_6F_{14} and nitrogen are routed along the support frames from the detector boxes to the lower end of the support frame. Here, the service boxes are mounted in which the detector signals are prepared for digital optical transmission to the counting house as described in section 5.2.3. Optical fibres, cables and supply tubes 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.

5.2.3 Electronics

A central position in the ST electronics is occupied by the service boxes [59] that have already been mentioned in the descriptions of the TT and IT detector stations. The service boxes are located close to the detectors but outside of the acceptance of the experiment. They are custom-made crates that hold 12 (TT) or 16 (IT) digitizer cards, a backplane for the distribution of control signals and low voltage, and a control card that provides interfaces to the LHCb Timing and Fast Control system (TFC, see section 8.3) and Experiment Control System (ECS, see section 8.4). On the digitizer cards, the analog output signals from the front-end readout chips are digitized, multiplexed and converted to optical signals. They are further transmitted via 120 m long optical fibres to the counting house, where they are received on the TELL1 board (see section 8.2).

Bias voltage for the silicon sensors and operation voltage for the readout electronics are provided by commercial voltage supplies. The grounding scheme follows LHCb grounding rules layed out in [60]. Temperature and humidity sensors are installed at various locations inside the detector boxes and on the service box backplane to monitor environmental conditions.

Readout and data transmission

Front-end chips. Both the TT and the IT make use of the Beetle front-end readout chip.[25] Four (TT) or three (IT) Beetle chips are located on a 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 the analog data 32-fold multiplexed via four differential output ports.

Copper cables. All output signals from one front-end hybrid are transmitted from the detector boxes to the service boxes via a shielded 68-wire twisted-pair cable. Only 32 (TT) or 24 (IT) of the 68 wires are needed for the transmission of the detector signals. The remaining wires are used to provide timing and slow-control signals, low voltage and ground from the service box to the front-end hybrid. The twisted-pair cables are between 2.7 m (IT bottom boxes) and 8 m (TT) long. Since, in the case of the IT, they go through the acceptance of the experiment, custom-made cables with significantly reduced shielding braids are employed to minimize the material budget.

Digitizer cards. Each digitizer card processes the data from one front-end hybrid. There are two variants of the card: a TT version to process the data from the four chip hybrids of the TT and an IT version to process the data from the three chip IT hybrids. The basic functional block for the processing of the signals from a single Beetle chip is illustrated in figure 5.29. Four differential line receivers convert the signals from the four Beetle output ports from differential to single-ended and match the signal levels to the input range of the ADC chips. Four single channel 8-bit ADC chips are used to digitize the data. These ADC chips operate at 40 MHz and are phase locked to the sampling clock of the Beetle chip. The 4×8-bit wide output data from the four ADC chips is then fed into a single Gigabit Optical Link (GOL) chip [6]. The GOL multiplexes the data and encodes them to a single Gigabit Ethernet data stream with a data rate of 1.6 Gbit/s. The laser driver integrated into the GOL chip is used to drive a 850 nm wavelength VCSEL diode that feeds the digitized optical data into a single optical fibre. This functional block is repeated four times on

[25] The Beetle front-end chip is a common development for ST and VELO. It has already been described in section 5.1.4.

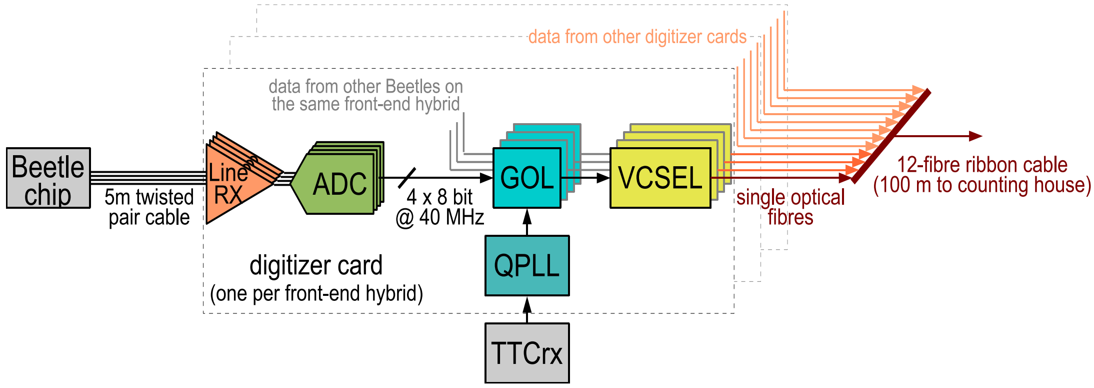

Figure 5.29: Functional block for the processing of the data from one Beetle chip. This functional block is repeated four times on a TT digitizer card and three times on an IT digitizer card.

a TT digitizer card and three times on an IT digitizer card. In addition, both variants of the card carry a central functional block for the distribution of the timing and control signals that the card receives from the service box control card.

Optical fibres. To transmit the data from the service boxes to the counting house, the outputs of twelve VCSEL diodes, corresponding to three TT digitizer cards or four IT digitizer cards, are connected to one twelve-fibre optical ribbon cable. This grouping defines a readout partitioning into groups of three readout sectors for the TT and four detector modules for the 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 twelve (four groups of three), an IT service box 16 (four groups of four) digitizer cards. Both types of service boxes, therefore, feed four twelve-fibre ribbon cables.

TELL1 boards. In the counting house, two fibre ribbon cables are connected to each TELL1 board. One TELL1 board therefore receives the data from six TT readout sectors or eight IT modules. A total of 84 TELL1 boards is required to read out the 280 readout sectors of the TT and the 336 detector modules of the IT. The functionality of the TELL1 board is described in section 8.2. Pedestals are calculated and subtracted for each readout channel, common-mode noise is calculated and subtracted for each event. A cluster-finding algorithm is applied and the positions and ADC values of the clusters are transmitted to the computer farm. Non-zero-suppressed data can be transmitted for monitoring and debugging purposes.

Detector control and monitoring

Detector control and monitoring is the main task of the service box control cards [61].

TFC. The control card holds a TTCrq mezzanine [62], which collects clock, trigger and timing information from the TFC network. The TFC signals are distributed to the digitizer cards via impedance controlled differential traces on the service box backplane. All GOL chips and Beetle chips associated with the same service box receive their clock signal from the same TTCrq mezzanine. The layouts of the backplane and the digitizer cards were optimized to equalize trace lengths, resulting in signal propagation time differences that do not exceed 3 ns.

ECS. The interface to the ECS is provided by two SPECS slave mezzanines that are mounted on each control card and that provide a total of eight I^2C busses. Four of these busses are used to control the GOL chips per group of three (TT) respectively four (IT) digitizer cards. The other four are used to control the Beetle chips per group of three respectively four readout hybrids. In addition, 36 I/O control lines permit to individually switch off the low-voltage regulators that provide the power for the readout hybrids and digitizer cards (see below).

Monitoring. 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 card. They are employed to monitor over-current conditions of the Beetle chips and to read out a PT1000 temperature sensor that is located on each readout hybrid.

Power distribution and grounding scheme

Low voltage. Low voltage levels of 2.5 V, 3.3 V and 5 V are required for the Beetle, GOL and ADC chips, the line receivers, and various LVDS drivers on the digitizer cards and the service box backplane. They are derived from voltage levels of about 5.5 V and 8 V that are generated by MARATON power supplies[26] in the LHCb cavern. The MARATON supplies are connected to the service box backplanes where they drive radiation tolerant programmable linear power regulators.[27] Two of these power regulators are used to provide each readout hybrid with 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. Another two power regulators provide each group of three (TT) respectively four (IT) digitizer cards with the required voltage levels, following the partitioning defined by the readout. These power regulators are located on the service box backplane and control card. They are cooled using the LHCb mixed water cooling system. The regulators for each readout hybrid and for each group of three respectively four digitizer cards can be individually switched off via the ECS.

Bias voltage. Detector bias voltage is provided by commercial high-voltage supplies[28] located in the counting house. It is connected to the detector boxes via 120 m long cables. The HV modules can provide up to 500 V and deliver a current of up to 10 mA per channel. Only the innermost readout sectors of the TT have individual HV channels to cope with the sensor leakage currents expected after several years of operation. Everywhere else, 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 from a HV patch panel in the counting house to the detector. The HV patch panel carries a jumper for each readout sector and detector module. This permits to manually disconnect individual readout sectors and detector modules from their HV supply.

Grounding scheme. The grounding scheme [63] of the detector boxes is illustrated in figure 5.30. The common ground for each detector box is defined by the cooling plates (TT) respectively cooling rods (IT) onto which the detector modules are mounted as described in sections 5.2.1 and 5.2.2. Thin copper wires are employed to connect ground pads on each of the Kapton flex prints

[26]MARATON low voltage power supply system for hazardous hostile environment, by W-IE-NE-R, Plein & Baus GmbH, Burscheid, Germany.

[27]LHC4913, developed by the CERN micro-electronics group and produced by ST Microelectronics, Geneva, Switzerland.

[28]CAEN SY1527 crates with A1511B modules, by CAEN S.p.A., Viareggio, Italy.

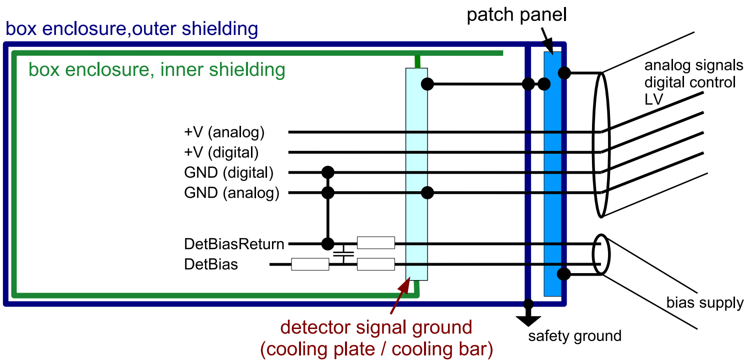

Figure 5.30: Grounding scheme of the detector boxes.

to the metallic screws with which the modules are fixed to the cooling balconies or cooling rods. The walls of the detector boxes are coated on both sides with 25 μm thick foils of aluminium. The inner shielding is connected to the common ground of the detector box and 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 shieldings of the signal and supply cables that connect to a detector box are connected to the outer box shielding. Both the LV and HV power supplies are kept floating and are connected to the LHCb general ground only in the detector boxes. In addition, the bias voltage for each module respectively readout sector is filtered using a passive low-pass RC filter that is implemented on the front-end readout hybrids.

5.2.4 Detector performance

An extensive R&D programme has been carried out to validate the detector concept, to optimize detector parameters and to estimate the expected performance of the detectors. It included simulation studies as well as various tests of prototype detectors in the laboratory and in test beams [64–68]. In view of the combination of the long readout strips and the fast pulse shapes employed, the signal-to-noise performance of the detectors was a major concern in these studies. The test beam measurements also confirmed that the expected spatial resolution of the detectors can be achieved. Other studies concerned the expected strip occupancies, which were estimated using events samples generated using the full GEANT 4 based simulation of the LHCb detector. A detailed analysis of the material budget of the detector was performed.

Signal-to-noise and efficiency

Various prototype detectors were built to establish the expected noise performance, charge collection efficiency and signal-to-noise performance of the final detectors. Effective readout strip lengths on these prototype detectors ranged from 108 mm up to 324 mm. Silicon sensors were employed that measured 320 μm to 500 μm in thickness, had strip pitches between 183 μm and 228 μm and ratios of strip implant width over strip pitch from 0.25 to 0.35. Some of the tested prototypes included Kapton interconnect cables of the same type and length as used in the M and K readout sectors of the TT. The performance of the prototype detectors was measured in an infrared laser

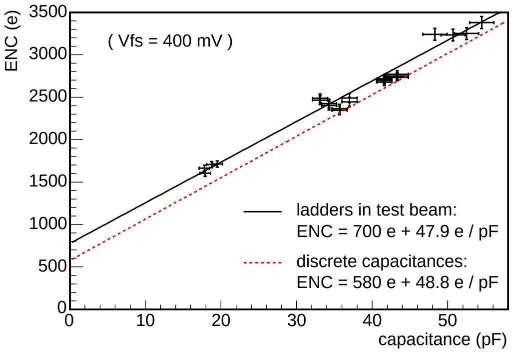

Figure 5.31: ENC obtained in test beam measurements as a function of the measured total strip capacitance of the tested prototype modules. The full line is a fit to the test beam results, the dashed line describes the results of laboratory measurements in which discrete capacitances were attached directly to the input of the Beetle chips.

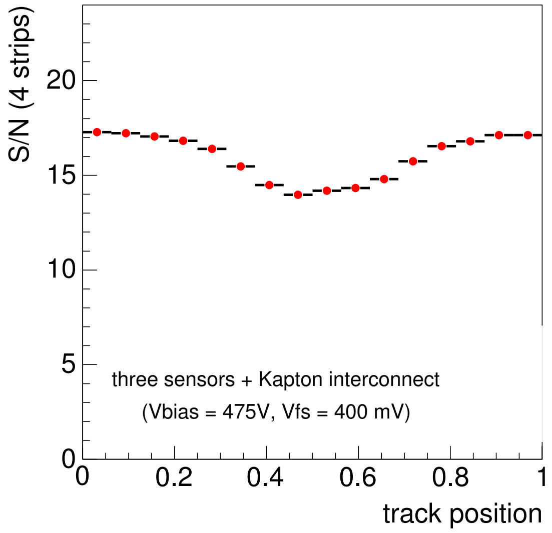

Figure 5.32: Most probable signal-to-noise ratio as a function of the inter-strip position (0 = centre of left strip, 1 = centre of right strip) measured in a CERN test beam. In order to rule out a potential bias due to clustering algorithms, the signal was calculated as the sum of the charges on the four strips closest to the particle impact point given by a beam telescope.

test stand and in several test beam periods in a 120 GeV π^- beam at CERN. The expected noise performance of the various detector configurations was also investigated in a SPICE simulation that 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 [69]. The results of this simulation agreed with the test beam measurements. The measured noise performance is summarized in figure 5.31. In this figure, the measured equivalent noise charge (ENC) for the different 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 of the Beetle chips, in which discrete load capacitances were attached to the Beetle inputs. Both the measurements and the SPICE simulations have demonstrated that the Kapton interconnect cables behave purely as an additional load capacitance and cause no deterioration of the quality and integrity of the detector signals.

Measurements of the charge collection efficiency were performed as a function of the position on the detector. No significant dependence was found on the position along the readout strips. However, in the direction orthogonal to the strips, a significant drop of the charge collection efficiency was observed in the central region between two readout strips. The effect is illustrated in figure 5.32 for a prototype module with the same detector geometry as a three-sensor readout sector of the TT. A similar charge loss was observed for all tested detector configurations. Its size did not depend on the strip length but was found to depend on the strip geometry. It decreased with increasing ratio of implant width to strip pitch (w/p) and with increasing sensor thickness (d) and depended roughly linearly on the ratio $(p-w)/d$. The charge loss could not be reduced by overbiasing 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 the silicon bulk and the silicon oxide in between the readout strips.

This charge loss does not affect particle detection efficiency as long as the signal-to-noise ratio in the central region in between the strips remains high enough. 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 thicknesses for the silicon sensors, most probable signal-to-noise ratios in excess of 12:1 are expected over the full surface of the detectors for both types of IT modules and all four types of TT readout sectors.

Spatial resolution and alignment

The spatial resolution of the detector modules was measured in test beams and was found to be about 50 μm, consistent with the expected resolution for the chosen strip pitches. Simulation studies have shown that, for two-strip clusters, it should be possible to further improve the spatial resolution by taking into account the position-dependent charge-sharing between the strips.

In order not to compromise the spatial resolution of the detectors, the positioning of each sensitive detector element in the x coordinate should be known to better than about 25 μm. The relative positioning accuracy of the individual silicon sensors on a detector module was monitored throughout the module production and was found to be better than 10 μm R.M.S. Various surveys of module positions inside the detector boxes and of detector boxes with respect to the LHCb reference frame have been performed. The results of these measurements are foreseen to be used to provide initial values for the software alignment of the detectors that will use reconstructed tracks from the LHCb spectrometer.

Strip occupancies

The strip occupancies presented here [70] were obtained from a sample of 5000 $B_d \rightarrow J/\psi(\mu^+\mu^-)$ $K_S(\pi^+\pi^-)$ events. The simulation software includes detailed descriptions of the detector geometry and of the signal collection, amplification and digitization. The simulation of the detector response was tuned to reproduce the results of the test beam measurements described before. For the TT, average strip occupancies of up to about 3.5% are found in the K sectors close to the beampipe. They 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 beampipe supports (see section 3) that are located at the exit of the LHCb magnet and are a prolific source of secondary particles.

2008 JINST 3 S08005

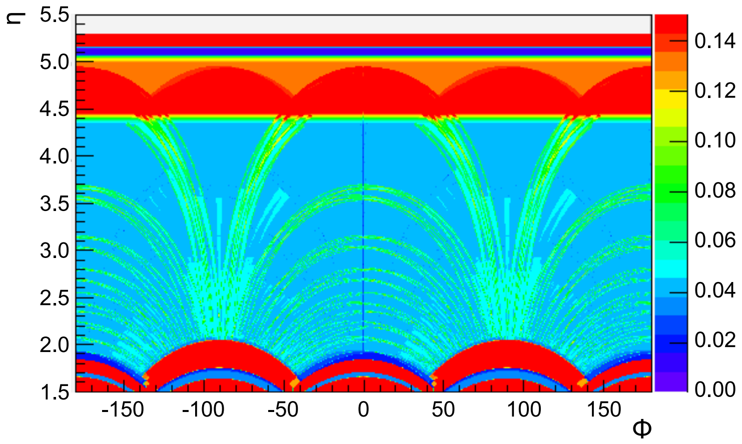

Figure 5.33: Radiation length of the TT as a function of the pseudorapidity η and the azimuthal angle ϕ.

Material budget

A careful analysis was performed of all materials that are located inside the acceptance of the experiment. Based on this analysis, a detailed description of all active detector elements and a simplified description of the passive components has been implemented [71, 72] in the XML based LHCb detector geometry description.

For the TT, where most of the dead material from detector supports, cooling etc. is located outside the acceptance, the material distribution is rather uniform. In total, it amounts to about 0.04 X_0, where more than 0.02 X_0 are due to the active material of the silicon sensors. An increase up to almost 0.13 X_0 is observed in the very forward region, due to the material of the insulation piece around the beampipe. The result of a material scan as a function of the azimuthal angle ϕ and the pseudorapidity η is shown in figure 5.33. The scan was performed by generating straight tracks originating from the nominal interaction point, extrapolating them through the TT station, and adding up the radiation lengths of all volumes in the geometry description that they crossed.

The material distribution for the 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 beampipe, the material budget adds up to about 0.035 X_0 per station, out of which more than 0.015 X_0 are due to the active material of the silicon sensors. The material budget peaks at almost 0.30 X_0 in the very narrow region of the cooling rods. The result of a material scan for one IT station is shown in figure 5.34.

5.3 Outer Tracker

The LHCb Outer Tracker (OT) is a drift-time detector [73], for the tracking of charged particles and the measurement of their momentum over a large acceptance area. Excellent momentum resolution is necessary for a precise determination of the invariant mass of the reconstructed b-hadrons: a mass resolution of 10 MeV/c^2 for the decay $B_s^0 \to D_s^- \pi^+$ translates into a required momentum resolution of $\delta p/p \approx 0.4\%$. The reconstruction of high multiplicity B decays demands a high tracking

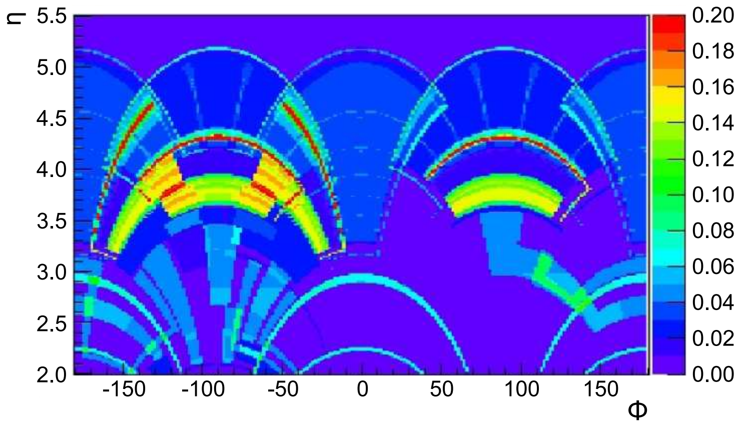

Figure 5.34: Radiation length of one IT station as a function of the pseudorapidity η and the azimuthal angle ϕ.

efficiency and at the same time a low fraction of wrongly reconstructed tracks: a track efficiency of 95% would result, for the decay $B_s^0 \to D_s^- \pi^+$, in an overall reconstruction efficiency of 80%.

5.3.1 Detector layout

The OT is designed as an array of individual, gas-tight straw-tube modules. Each module contains two staggered layers (monolayers) of drift-tubes with inner diameters of 4.9 mm. As a counting gas, a mixture of Argon (70%) and CO_2 (30%) is chosen in order to guarantee a fast drift time (below 50 ns), and a sufficient drift-coordinate resolution (200 μm). The gas purification, mixing and distribution system foresees the possibility of circulating a counting gas mixture of up to three components in a closed loop [74].

The detector modules are arranged in three stations (see figure 5.35). Each station consists of four layers, arranged in an *x*-*u*-*v*-*x* geometry: the modules in the *x*-layers are oriented vertically, whereas those in the *u* and *v* layers are tilted by $\pm\, 5^o$ with respect to the vertical, respectively. The total active area of a station is $5971 \times 4850\,\text{mm}^2$. The outer boundary corresponds to an acceptance of 300 mrad in the magnet bending plane (horizontal) and 250 mrad in the non-bending plane (vertical). The inner cross-shaped boundary of the OT acceptance was determined by the requirement that occupancies should not exceed 10% at a luminosity of $2 \times 10^{32}\ \text{cm}^{-2}\ \text{s}^{-1}$ (this area is covered by the IT, see section 5.2.2).

The OT assembly is shown in figure 5.35. The detector modules are supported by aluminium structures. Each station is split into two halves, retractable on both sides of the beam line. Each half consists of two independently movable units of two half layers (C-frames). The modules are positionned on the C-frames by means of precision dowel pins. The C-frames also provide routing for all detector services (gas, low and high voltage, water cooling, data fibres, slow and fast control). The OT C-frames are sustained by a stainless steel structure (OT bridge), equipped with rails allowing the independent movement of all twelve C-frames.

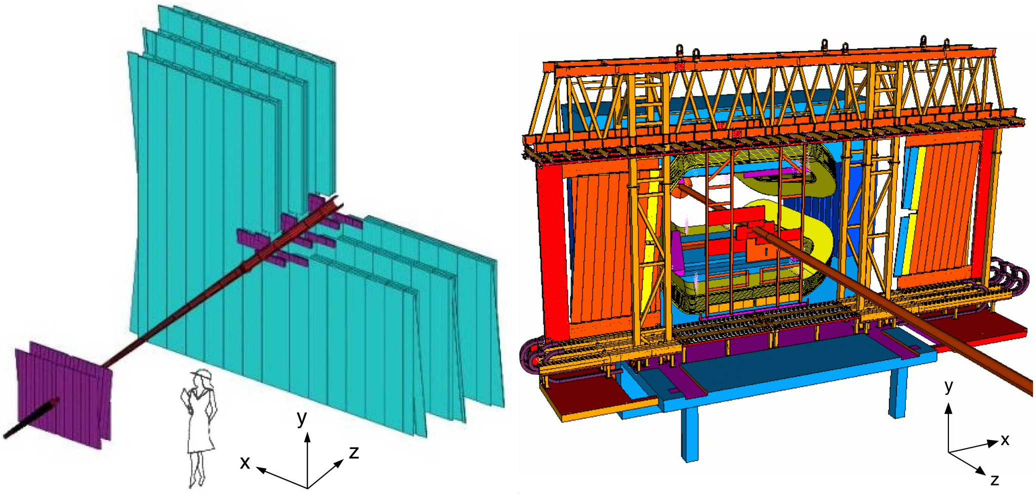

Figure 5.35: Arrangement of OT straw-tube modules in layers and stations (left) and overview of the OT bridge carrying the C-frames (right). The C-frames on both sides of the beam pipe are retracted.

5.3.2 Detector technology

Design

The design of the straw-tube module is based on the following requirements:

- Rigidity: the mechanical stability must guarantee the straw-tube position within a precision of 100 (500) μm in the x (z) direction; the anode wire has to be centered with respect to the straw tube within 50 μm over the entire straw length. The module box must be gas-tight and must withstand an overpressure of 10 mbar. The leak rate at this pressure has to be below 8×10^{-4} l/s.

- Material budget: to limit multiple scattering and the material in front of the calorimeters, the material introduced in the OT active area must not exceed few percent of a radiation length X_0 per station.

- Electrical shielding: the drift tubes must be properly shielded to avoid crosstalk and noise. Each straw must have a firm connection to the module ground. The module envelope itself must form a Faraday cage connected to the ground of the straw tubes and of the front-end electronics.

- Radiation hardness: the detector should withstand 10 years of operation at the nominal luminosity without a significant degradation of its performance. During that time the anode wires will accumulate a charge of up to 1 C/cm in the most irradiated area. As a consequence, all detector materials have to be radiation resistant and must have low outgassing.

The layout of the straw-tube modules is shown in figure 5.36. The modules are composed of two staggered layers (monolayers) of 64 drift tubes each. In the longest modules (type F) the monolayers are split longitudinally in the middle into two sections composed of individual straw

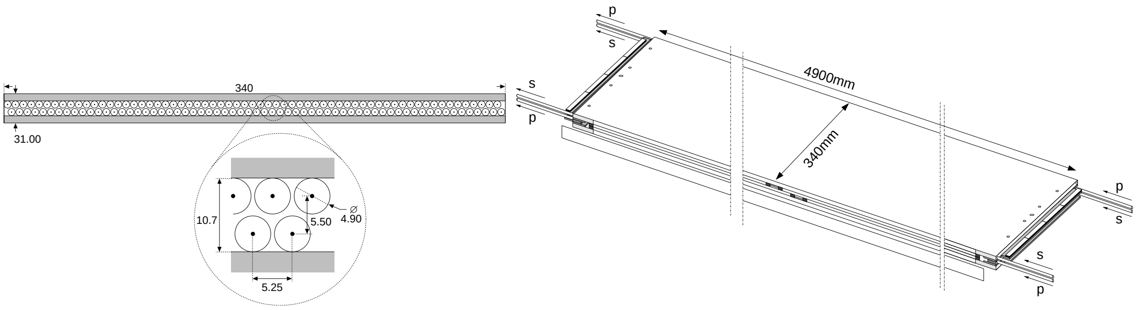

Figure 5.36: Cross section of a straw-tubes module (left) and overview of a straw-tubes module design (right).

2008 JINST 3 S08005

tubes. Both sections are read out from the outer end. The splitting in two sections is done at a different position for the two monolayers to avoid insensitive regions in the middle of the module. F-modules have an active length of 4850 mm and contain a total of 256 straws. In addition to the F-type modules there exist short modules (type S) which are located above and below the beam pipe. These modules have about half the length of F-type modules, contain 128 single drift tubes, and are read out only from the outer module end. A layer half is built from 7 long and 4 short modules. The complete OT detector consists of 168 long and 96 short modules and comprises about 55000 single straw-tube channels.

Construction

The straw tubes are produced by winding together two strips of thin foils,[29] as shown in figure 5.37: the inner (cathode) foil is made of 40 μm carbon doped polyimide (Kapton-XC[30]); the outer foil (Kapton-aluminium) is a laminate[31] made of 25 μm polyimide, to enhance the straws gas tightness, and 12.5 μm aluminium, crucial to ensure fast signal transmission and good shielding.

To build a monolayer the straw-tubes were glued to panels with a cored sandwich structure consisting of a 10 mm Rohacell core and two 120 μm carbon fibre skins. High precision aluminium templates (figure 5.37) were used during the glueing to position the straw-tubes to better than 50 μm over the entire module length. After the straw-tubes were glued to the panel the wiring was started. A gold-plated tungsten wire[32] with a diameter of 25.4 μm is used for the anodes. The wire was sucked through the straw-tube. At each end the wire is guided using injection-molded Noryl endpieces. To centre the wire also along the straw-tube Noryl wire locators had been placed every 80 cm inside the straws. The wires were strung with a tension of 0.7 N and were soldered to 5 mm long pads of a printed circuit board.

Special holding-devices, shown in figure 5.38, were used to keep the support panels flat to within 100 μm during the glueing of the straws and wiring. They were also used to assemble two monolayer panels into a detector module (figure 5.38). The sides of the modules were closed by 400 μm thick carbon fibre sidewalls. Spacers at the two module ends ensure the proper separation

[29] Lamina Dieletrics Ltd., UK.
[30] DuPont^{TM}.
[31] GTS Flexible Materials Ltd., USA.
[32] California Fine Wire, USA.

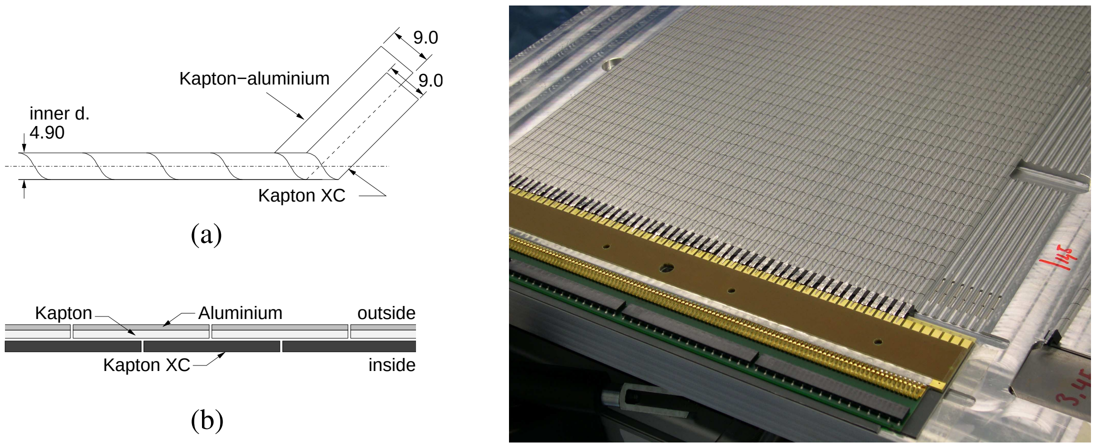

Figure 5.37: Left: (a) scheme of the straw winding using two foils; (b) Kapton-XC as inner foil and a Kapton-aluminium laminate as outer foil. Right: straws on the high precision aluminium template.

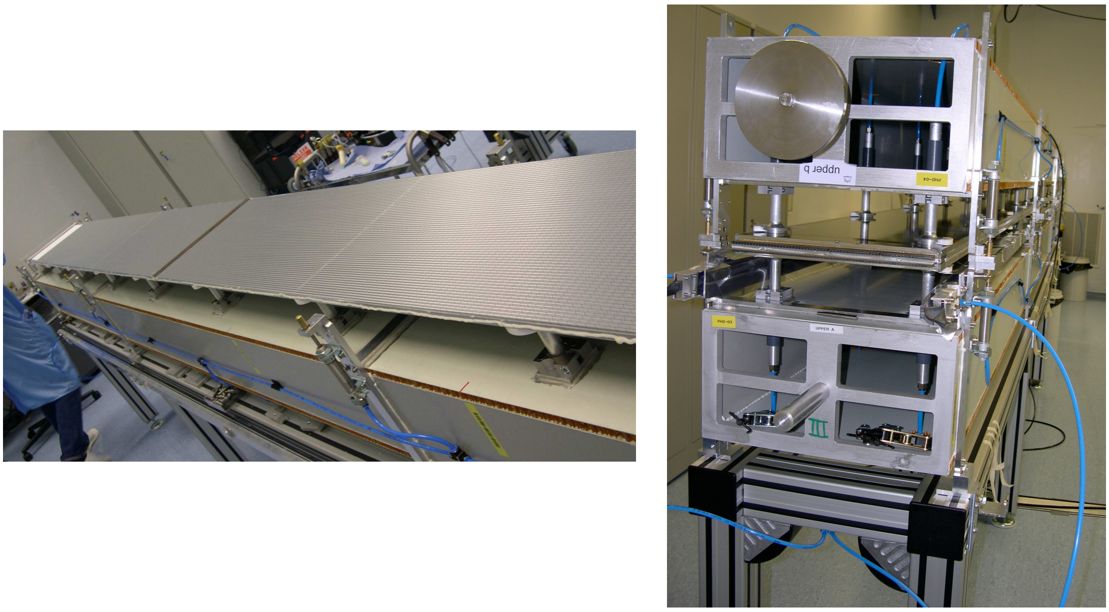

Figure 5.38: Left: straw-tube monolayer on a holding-device ready for gluing. Right: two straw-tube monolayers in a holding device while being glued to form a detector module.

of the two monolayers and provide an interface to the module gas pipes. All glueing steps were performed using Araldite AY103 in combination with the HY991 hardener. The glue viscosity was enhanced by adding silica gel.

The side walls and also the inner sides of the support panels are covered by the Kapton-aluminium laminate, which ensures the gas tightness of the module and provides a closed Faraday cage.

The construction of the detector modules was completed by three production sites in about two years. All detector parts were centrally checked and then distributed to the production sites. The production procedure and the quality monitoring steps and tools were the same in all production sites. Quality assurance included the check of the wire tension, pitch, and leakage current (in air) prior to the module sealing. Finished modules were tested for gas tightness. Using CO_2 as counting gas the modules were operated with a slowly increasing anode voltage to ensure that the leakage current per wire dropped below 1 nA for a voltage of 1700 V. Finally, the uniformity of the signal response of all drift cells was checked with a radioactive source.

Material budget

The forward geometry of the LHCb detector allows all detector services and supports of the OT to be located outside the detector acceptance. The OT material inside the LHCb acceptance (12 layers of modules) is estimated from the weights of the single components and the total amount of glue used (140 g per panel, 400 g to glue straw-tubes and to close the module) [75]. The total material of one station (four layers of modules, i.e. 8 monolayer panels) is on average equivalent to 3.2% of X_0. The total OT material sums up to 9.6% of X_0.

Ageing tests

In the region closest to the beam axis, the OT detector modules are foreseen to operate under particle rates per straw length of up 100 kHz/cm. Extensive ageing studies of test modules and of scaled-down prototypes showed that the detector technology is resistant to radiation doses corresponding to accumulated charges of up to 1.3 C/cm. No hints of a change of the gas amplificaton were seen in these studies which were performed with an average acceleration factor of a approximately 25.

However, irradiating mass production modules with low intensity β or γ sources with acceleration factors of O(1) produced a significant drop of the gas amplification. Most strikingly, this gain loss was not observed in the region of highest irradiation intensity (anode currents around 20 nA/cm) but at modest intensities and anode currents of 2–5 nA/cm [76]. The damage was only observed upstream (with respect to the gas flow) of the irradiating source. An analysis of the affected anode wires revealed a thin (less than 1 μm thick) insulating deposit. Although the exact mechanism is not yet understood, the glue (Araldite AY103) used to build the modules is identified as cause of the depositions. A number of preventive and remedial actions have been studied: long-term flushing of the detector modules with gas (CO_2) as well as warming-up the modules to 40^oC while being flushed, significantly reduce this effect. Large anode currents deliberately provoked by increasing the high voltage above the nominal operating voltage can be used to clean the anode wires, should the insulating deposits reduce the gain beyond an acceptable level.

5.3.3 Electronics

The front-end (FE) electronics measures the drift times of the ionization clusters produced by charged particles traversing the straw-tubes with respect to the beam crossing (BX) signal [77]. The drift times are digitized for every bunch crossing (25 ns) and stored in a digital pipeline to await the Level-0 decision. On a positive L0 decision, the digitized data of up to 3 bunch crossings (to cover a time range of up to 75 ns) is transmitted via optical links to the Level-1 buffer (TELL1)

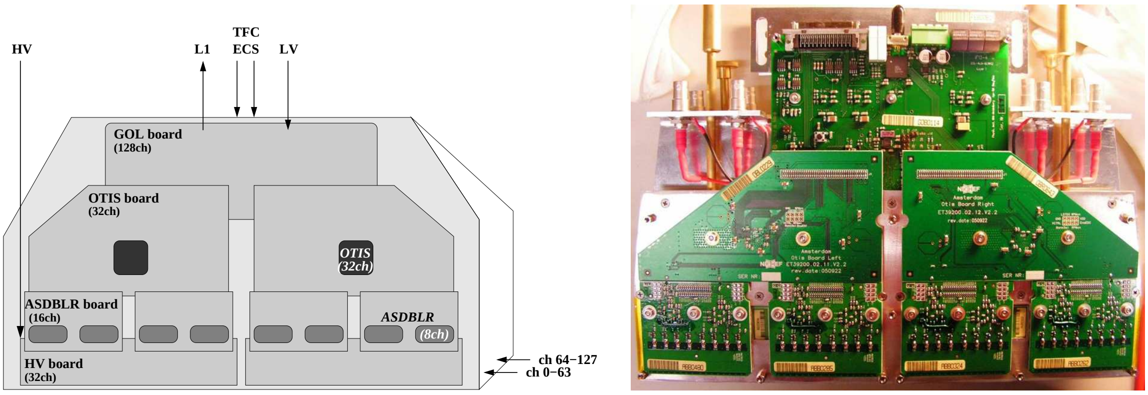

Figure 5.39: Design (left) and photograph (right) of the FE electronics mounted in a FE box. In the photograph the HV boards are not visible, because hidden by the ASDBLR boards.

boards. The radiation dose expected for the front-end electronics is only 10 kRad, well below the maximum tolerable dose of 2 Mrad.

As shown in figure 5.39, the FE electronics has a modular design, consisting of several interconnected boards housed inside a metallic box (FE box). These boxes are mounted at each end of the detector modules. A FE box is the smallest independent readout unit of the OT: the digitized data of the 128 channels of one module are sent via an optical link and received by the TELL1 board; high- and low-voltage, as well as fast- and slow-control signals are connected to each FE box individually. In total, 432 FE boxes are used to read out the OT detector.

The main components of the OT readout electronics are the High Voltage (HV) board, the ASDBLR board, the OTIS board, and the GOL auxiliary (GOL/AUX) board.

Four 32-channel HV boards plug directly into the signal feedthrough of each straw-tube module. The signal feedthrough is provided by a passive printed circuit board (PCB) built into the module end and defining the reference ground for the straw-tubes and readout electronics. The HV board has a single HV connection and is thus the smallest independent HV supply unit. Anode signals are decoupled using 300 pF capacitors[33] on the PCB.

Each ASDBLR board hosts two ASDBLR chips [78]. These are custom integrated circuits, providing the complete analog signal processing chain (amplification, shaping, baseline restoration, and discrimination) for straw sensors. The ASDBLR is implemented in bipolar technology and produced with the radiation hard DMILL process It includes eight identical channels per chip, with a peaking time of 7–8 ns, selectable shaping circuits and a three-state LVDS output. The Equivalent Noise Charge (ENC) at the ASDBLR input is about 2200 e^- + 140 e^-/pF, corresponding to a global ENC of about 0.9 fC when connected to the detector.

The hit outputs of two ASDBLR boards (32 channels) are connected to an OTIS board, which hosts one OTIS chip for drift time digitization. The OTIS ASIC is a radiation hard 32-channel Time-to-Digital Converter (TDC), developed specifically for the OT and manufactured in a standard 0.25 μm CMOS process [79, 80]. The OTIS block diagram is shown in figure 5.40. The TDC core measures the arrival time of the ASDBLR signals with respect to the LHC clock propagating through a 25 ns long Delay Locked Loop (DLL), a regulated chain of 32 double-staged

[33] JOHANSON 302R29W331KV4E.

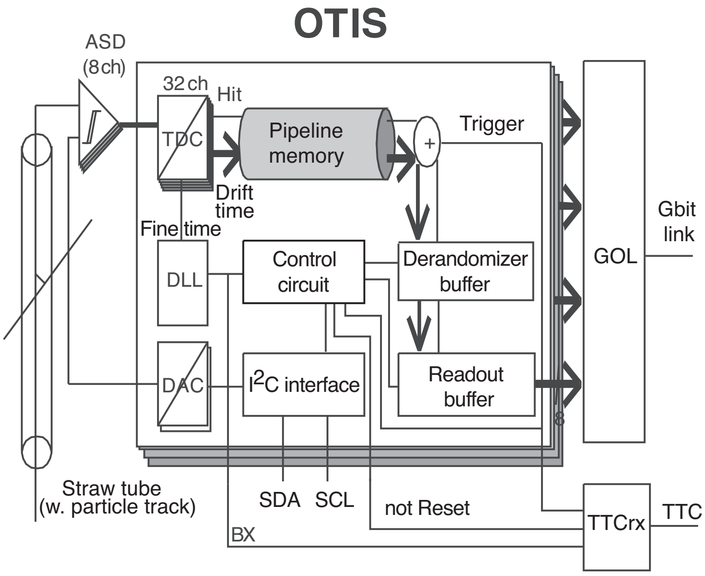

Figure 5.40: OTIS Block Diagram.

delay elements. The time digitization is done using the 64 delay-stages of the DLL (64 time bins) giving a theoretical bin size of 390 ps. The drift time data is stored in a pipeline memory with a depth of 164 events, allowing a L0 trigger latency of 160 clock cycles to compensate for trigger rate fluctuations. If a trigger occurs, the corresponding data words of up to 3 bunch crossings (the number of bunch crossings transferred is selectable) are transferred to a derandomizing buffer, able to store data from up to 16 consecutive triggers. Both, pipeline memory and derandomizing buffer are dual-ported SRAM memories. A control unit processes and reads out the data of each triggered event within 900 ns. The readout interface of the OTIS chip is 8-bit wide.

The data processing within the OTIS is clock driven. The chip operates synchronous to the 40 MHz LHC clock. A standard I^2C interface is used for setup and slow control.

The OTIS boards in a FE box are connected to one GOL/AUX board. This board [81] provides the outside connections to the FE box: the power connection, the interface to the fast-control (BX clock, triggers, resets) and the interface to the slow-control (I^2C). Three power regulators supply the different voltages (+2.5 V $\pm$3.0 V) needed by the OTIS and the ASDBLR chips. The GOL/AUX board hosts the GOL optical serializer chip [6] connected to an optical receiver mezzanine card on the TELL1 board.

As the data volume of the OT is large compared to other sub-detectors only the optical links of the nine FE boxes belonging to a layer quadrant are connected to a single TELL1 board. To read out the 432 FE boxes of the OT 48 TELL1 boards are used.

Cooling

As the FE box is entirely closed, cooling of the electronic components inside the boxes is necessary. All electronic boards are therefore mounted on a cooling frame (see figure 5.39) which is screwed to water-cooled plates on the C-frames.

2008 JINST 3 S08005

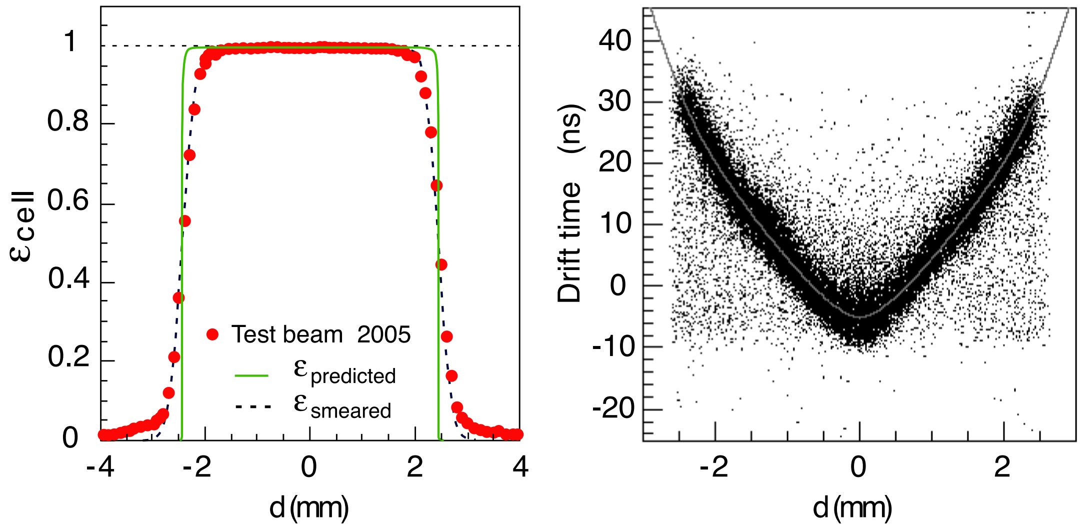

Figure 5.41: Left: hit efficiency as a function of the distance to the anode wire. Right: relation between the drift time and the calculated hit distance to the wire. Note the maximum drift time of approximately 45 ns.

High-voltage system

Each FE box has four independent HV connections, one for each 32-channels HV board. Two CAEN SY1527LC mainframes, each equipped with four 28-channels A1833BPLC supply boards, are used as HV supply. Using an 8-to-1 distribution scheme the 1680 HV connections of the detector are mapped on 210 CAEN HV channels. The distribution is realized using a patch panel wich offers the possibility to disconnect individual HV boards (32 channels) by means of an HV jumper. Both components, the HV supply as well as the patch panel are located in the counting house. Access to the HV system during data taking is therefore possible.

5.3.4 Test Beam results

To determine the performance of the detector modules in combination with the readout electronics, 4 short modules from the mass production equipped with FE boxes were tested at the 6 GeV electron beam of the DESY II facility in Hamburg [82]. Although the number of OT modules (4 modules, or 8 monolayers) was sufficient to allow full track reconstruction, a silicon strip telescope was used to provide redundant information for the determination of the coordinates at which the beam particle traverses the detector. The trigger signal (which was also used offline as a time reference) was generated by a coincidence of two scintillator counters installed downstream of the OT modules. The nominal counting gas mixture Ar/CO_2 was used.

The electron beam illuminated up to 7 straws per monolayer. In the offline analysis, the relation between the measured drift time and the distance to the wire (R-t relation, see figure 5.41) was established using the predicted distance of closest approach of the particle to the anode wire. The R-t relation was then used to convert measured drift times to coordinates. The hit finding efficiency, shown as a function of the distance to the wire in figure 5.41, is determined by verifying

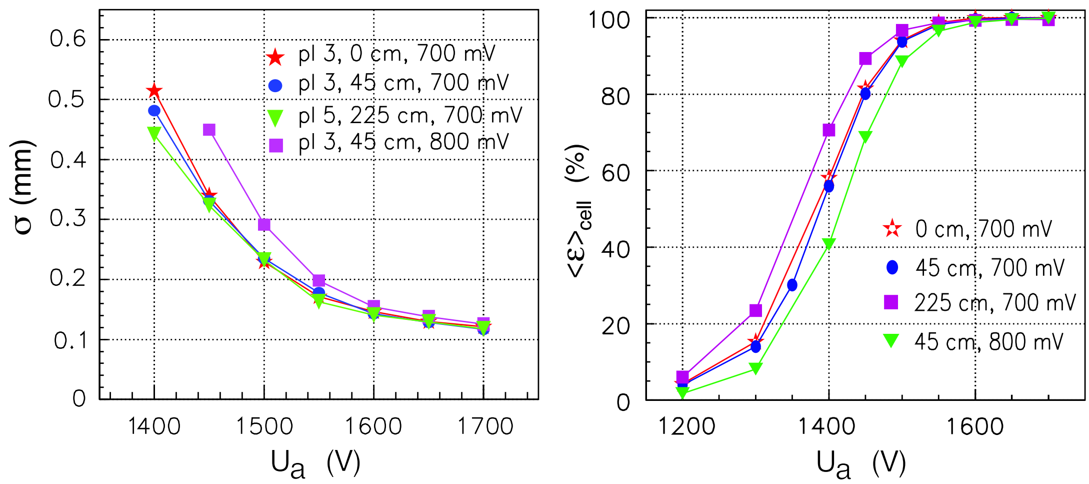

Figure 5.42: Performance of the straw-tube module: the position resolution (left) and the average efficiency (right) of a single cell are shown as a function of the HV value. The different curves correspond to different discriminator threshold voltages, and different distances between the electron beam spot and the front-end electronics.

whether the OT produced a hit at the predicted position. The position resolution is determined by comparing the measured hit coordinates to the predicted ones (residual) and by fitting a single gaussian to the distribution.

The performance of the OT straw-tube modules is determined for HV values ranging between 1200 and 1700 V, and for ASDBLR discriminator thresholds ranging from 1.5 fC to 5.5 fC. As shown in figure 5.42, efficiencies larger than 99% in the centre of the straw (dropping at the edge of the straw) and position resolutions of single cells below 200 μm can be attained for high voltage values above 1550 V. The noise level was found to be low: at a high voltage of 1550 V and a discriminator threshold voltage of 800 mV (corresponding to a signal threshold of 4 fC), the noise rate is below 1 kHz per channel. This corresponds to an average channel occupancy of (7.5×10^{-3})%. The probability to find a coherent hit in neighbouring channels (crosstalk) is found to be lower than 5% for high voltage values below 1600 V and discriminator threshold voltages above 800 mV (4 fC).

5.3.5 Alignment and monitoring

Errors in the mechanical alignment of the drift tubes can significantly degrade the track reconstruction. A single cell resolution of about 200 μm requires that the drift tubes be aligned within an accuracy of 100 μm (1 mm) in the x (z) coordinate. Therefore, care was taken during each step of the detector construction and installation to minimize alignment errors. The issue of the mechanical tolerances in the module production has been discussed in section 5.3.2. All detector C-frames which hold the modules were built with stringent requirements on the mechanical tolerances [83, 84]. During installation the positions of all modules were surveyed and the C-frames positions have been adjusted accordingly.

The reproducibility of the C-frames positioning after movement was checked to be better than the 200 μm precision of the optical survey.

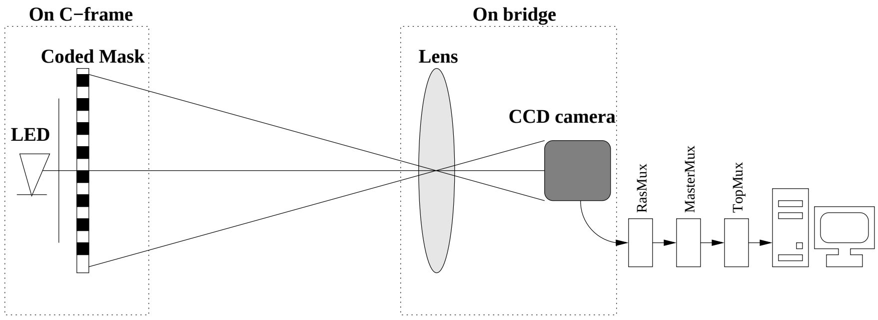

Figure 5.43: Schematic of the RASNIK alignment system.

The stability of the C-frames relative position during data taking is monitored by means of the RASNIK system [85], whose basic idea is to project the image of a detailed pattern through a lens onto a CCD camera (figure 5.43). Movements perpendicular to the optical axis are observed as change of the pattern position processed by the CCD camera, whereas movements along the axis are measured by the change in the image size. A total of 50 RASNIK lines are used for the OT monitoring. The intrinsic resolution of the system perpendicular (parallel) to the beam axis is better than 10 (150) μm respectively [86].