

Chapter 1

VELO

1.1 Theory

1.1.1 Silicon Particle Detectors

A semiconductor is a material that acts as part insulator, part conductor. Typical semiconductors, such as silicon, lie in group IV of the periodic table, and hence, have four electrons in their outer (valence) shells. When silicon crystallises, its atoms form covalent bonds. This means that the electrons are shared between neighbouring atoms in the crystal structure.

Due to the close proximity of adjacent atoms in a solid, the atomic wave-functions overlap and the individual energy levels form bands. The bands most relevant to this discussion are the valence band (the highest filled band) and the conduction band (the mostly empty band above the valence band). The bands are separated by a band gap and it is this band gap that distinguishes between conductor, insulator and semiconductor. Insulators have a very large band gap and the electrons are not energetic enough to make the leap to the conduction band. In conductors, the conduction band and the valence band overlap and there is an abundance of vacant energy states, allowing the electrons to move freely. In a semiconductor, the band gap is small enough that thermally excited electrons can make it across the energy gap at room temperature.

In a silicon crystal, when an electron jumps into the conduction band it leaves a vacant energy state, called a hole. A neighbouring electron can fill this hole, creating a new hole. In this manner the hole can move about the crystal, essentially acting as a positive charge carrier. In the presence of an E-field, electrons move in one direction and holes move in the opposite direction.

However, the number of electrons excited to the conduction band by thermal excitation is very small, and hence, density of charge carriers in silicon is also small. In order to increase the density of charge carriers, impurities can be introduced into the silicon. An element from group V in the periodic table can be added to the silicon and will fill one of the lattice crystal spaces. It has the four electrons necessary to make the covalent bond and one extra that is loosely bound to its nucleus and can easily be elevated to the conduction band. This process is called doping. Doping which adds extra conduction electrons is called n-type doping. Extra valence holes can also be added by doping with an element from group III in the periodic table - this is p-type doping. Note, however, that the silicon remains electrically neutral (since the numbers of electrons and protons are equal).

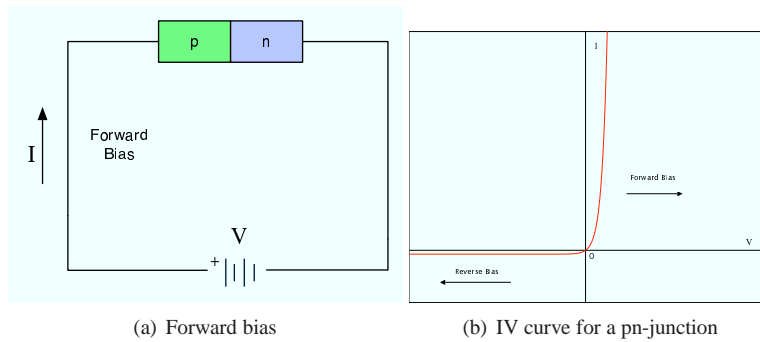


Figure 1.1: p-n junction bias

A p-n junction forms when n-type semiconductor comes into contact with p-type semiconductor. The zone where the two types meet is known as the depletion layer. This region contains no charge carriers - the conduction electrons from the n-type material travel over and fill the holes in the p-type material. The p-side of the junction becomes positively charged and the n-side negatively charged. This creates an electric field across the junction. This counteracts the continued diffusion of charge carriers. When the electric field is sufficient to repel incoming holes and electrons, the depletion region reaches its equilibrium width.

When a voltage is applied across the p-n junction, an external electric field forms (separate to the internal E-field at the pn-junction). For example, in Fig. 1.1 the positive terminal of a battery is connected to the p-side and the negative to the n-side of the junction. This creates a E-field pushing holes from the p-side where they are plentiful to the n-side, and vice-versa with the electrons. Thus, a current flows. This configuration (where p is at a higher electrical potential with respect to n) is known as a forward-biased p-n junction. Forward bias (P positive with respect to N) narrows the depletion zone and eventually reduces it to nothing, making the junction conductive and allowing free flow of charge carriers. A reverse-biased junction exists when the n-side is at a higher potential than the p-side. In this case, the field is in the opposite direction, but there are very few holes on the n-side to flow to the p-side and very few electrons on the p-side to flow to the n-side. Hence, very little current flows through a reverse-biased p-n junction. Under reverse bias (P negative with respect to N) the potential is increased, further widening the depletion zone. Fig. 1.1b shows an IV curve for a typical p-n junction. Above a certain voltage, a forward biased pn-junction has almost no resistance. With reverse biased, only a very small current flows due to thermally excited electrons and holes.

With a high enough reverse bias voltage the material can be fully depleted. A depleted pn-junction can be used as a particle detector. Due to the strong E-field in the depletion region, electron-hole pairs generated by a charged particle travelling through can be collected on the n or p side. Typically, the sensitive part consists of p strips in an n bulk. A cross section through a silicon detector is shown in figure 1.2. The deposited charge travels towards the nearest strip. This charge can be read out via aluminium strips.




Figure 1.2: Particle traveling through Silicon Sensor

1.1.2 The LHCb Vertex Locator

The Vertex Locator (VELO) is the primary vertexing and tracking detector for LHCb. The VELO surrounds the interaction point in LHCb, and has to provide precise measurements of track coordinates close to the interaction region. These are used to reconstruct production and decay vertices of beauty and charm hadrons, to provide accurate measurement of their decay lifetimes, and to measure the impact parameter of particles used to tag their flavour.

The VELO consists of 40 modules placed along the beam direction. Fig. 1.3 shows a cut-away view of the VELO. The modules can be seen in the centre of the tank. A single module can be seen in fig. 1.7.

Each module consists of 2 silicon sensors for measuring particle hits in R and ϕ . The modules are positioned to within 8mm of the beam and are retractable to protect the sensors, as the LHC beam is broader during initial beam injection.

Each sensor has 2048 aluminium strips. The ϕ strips are not perfectly symmetric, and are placed facing opposite directions relative to their closest neighbours. See fig. 1.4a. This means that strips on adjacent ϕ sensors, for example, are at a stereo angle relative to each other (fig. 1.4b), giving a better overall hit resolution.

The active material is the silicon sensor. A double-metal technology¹ was used in

¹The readout strips were placed on a second layer of metal, insulated from the metal that covered the detecting strips.




Figure 1.3: Velo Detector

the sensor design, enabling a strip close to the beam to be routed out over the top of other structures.

1.1.2.1 Radiation Hardness

In LHCb, the high levels of radiation are a concern. The silicon atoms can become displaced after irradiation, leaving vacancies and interstitials (extra energy levels in the band gap). Charges become trapped in the bulk of the silicon, leading to reduced charge carrier mobility and charge collection efficiency. As a result the sensor may become under-depleted. The sensor depletes from the n side, and this leads to a loss of efficiency and resolution in p-on-n sensors. For the VELO, $[n^+-on-n]^2$ was chosen, as the depleted layer is on the same side as the strips. This means that the sensors are radiation hard - they can retain 90% of their functionality after 3 years of LHCb running.

1.2 Experimental Setup At the University of Liverpool

I spent four months at the University of Liverpool to build and operate a DAQ system to test the Velo Modules. The DAQ system is a near replica of what will be used for data taking in the running experiment (the differences are detailed in section 1.2.2). The DAQ system was designed to read out data from a module, which was then analysed to characterise its electrical properties. Thermographs were taken of the module under

² n^+ - the + refers to high doping concentration ($\sim 1:10,000$) compared to the normal ($1:100,000,000$)

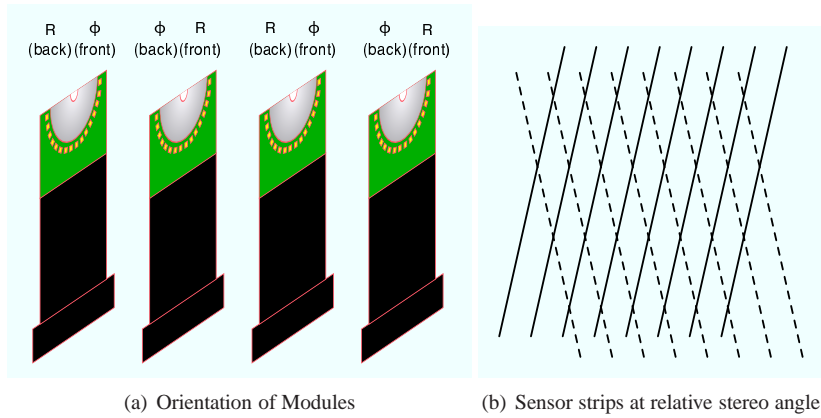


Figure 1.4:

vacuum and cooling. These were used to inspect the modules qualitatively for visually comparing modules and checking for abnormalities which can show up as “hot-spots” on the thermograph.

There are four main parts to the module testing setup - i. the module, ii. the readout electronics, iii. the thermographic apparatus, iv. and the system for creating a LHC-like environment, namely, the vacuum and cooling system.

Fig. 1.6 shows the connections in the DAQ system. The module, containing Sensors and Beetle Chips, is enclosed in the vacuum tank and connected to the repeater cards. The Beetle chips, repeater cards, and the Tell1 digitisation board form the readout network. Data flows through them as shown by the black lines. The data is finally stored on the PC and analysed at a later stage. The PC also connects to the Credit Card PC on the Tell1 which sends control signals (in blue) up the chain to the Beetle chips. The following sections describe these parts in more detail. Section 1.2.1 describes the module and its various parts. Section 1.2.2 describes the readout electronics. The Thermal camera and its use are described in section 1.2.3. The vacuum and cooling systems are outlined in section 1.2.4. The last section (1.2.5) gives the steps taken to run the DAQ system.

1.2.1 Module

A complete module is shown in fig. 1.7. A module consists of two halves; one with an R-sensor and one with a ϕ sensor. As can be seen from the figure, the sensors are connected via pitch adapters to Beetle chips. These chips are used for readout. All of these are glued and electrically bonded on a TPG³ Hybrid. This, in turn, is mounted onto a rigid carbon fibre pedestal.

1.2.1.1 Sensor

Each sensor has 2048 aluminum strips that can be read out. The ϕ sensors have radial strips, divided into inner and outer regions, of 683 and 1365 strips respectively. The R sensor’s strips are circular, divided into 4 quadrants of 512 strips each. The strips of the R sensor range from $40\mu\text{m}$ to $92\mu\text{m}$ in pitch, and the ϕ strips range from $37\mu\text{m}$

³Thermal Pyrolytic Graphite




Figure 1.5: Vacuum Tank, Cooling and Vacuum Pumps.

to $98\mu\text{m}$ in pitch. The .Fig. 1.7b shows the two sensors. Combined, the two sensors allow a position in 3D co-ordinate space to typically $5\mu\text{m}$ accuracy. The sensors are produced by Micron and are designed to be radiation hard for at least three years of LHCb running time, at an average luminosity of $2 \times 10^{32} \text{cm}^{-2} \text{s}^{-1}$. Each sensor spans 182 degrees, so there is a slight overlap between opposing sensors. n-on-n silicon was used as the active semi-conducting material.

1.2.1.2 Beetle

The beetle chips read out the charges on the sensor strips. For each sensor there are 16 Beetle chips; made using radiation hard CMOS technology. One chip processes a charge from 128 strips on the sensor at a rate up to 1.1MHz. Readings from the 128 strips are split among the four data output ports on the Beetle chip. The data is read as an analogue signal and arranged into frames of 32 strips. The data is output with some header information which is added to the data before it is propagated onwards. The header is in a binary format and includes information such as pipeline location, and some parity bits. Each chip has a pipeline (or buffer) with 187 locations. This starts to fill when the chip receives readout triggers.

1.2.1.3 Hybrid

The hybrid is the rigid structure on which the sensors, beetle chips and other circuitry are mounted. The bulk of the hybrid consists of TPG (Thermal Pyrolytic Graphite). TPG was chosen in order to cool the silicon and the chips. It has coefficient of thermal expansion (c.t.e.) close to zero and high thermal conductivity $\sim 1600 \text{W/mK}$ (Cu for instance is $\sim 400 \text{W/mK}$). TPG is a laminar substrate and it may delaminate at the edges when glued to another surface. TPG can be quite weak, so about 3mm Carbon Fibre

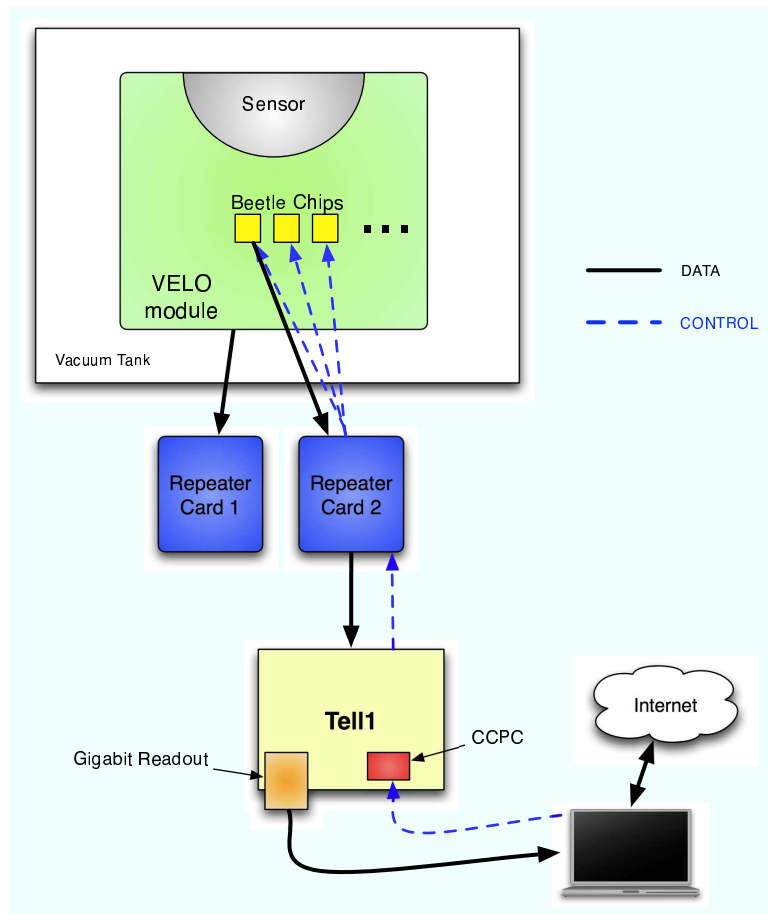



Figure 1.6: Data flow



(a) R and Phi silicon sensors

Figure 1.7: Velo Module and Sensor

(CF) is grafted on to the edges for structural integrity. A kapton circuit is laminated on top of the TPG in a vacuum press at 180°C . The kapton has a c.t.e. of $\sim 20\text{p.p.m}$. Thus, the hybrid is under tension after cooling. The CF also helps withstand this tension. On top of the Kapton circuit, discrete components (resistors, capacitors, etc) are added, then Pitch Adapters, Beetle chips and finally Silicon.

1.2.1.4 Thermistors

Thermistors on the hybrid are used to measure the temperature of the module and ensure that they are adequately cooled. There are two thermistors on each side of the module - one in the centre, roughly halfway between the cooling block and the sensor edge, and one close to the top right. These can be read out through a port on the repeater boards.

1.2.2 Readout Electronics

1.2.2.1 Repeater Boards

The primary function of the repeater boards is to shape and drive the signals from the Beetle chips over 60m of twisted-pair cables to the ADCs on the Tell1. For convenience, only 5m cables were used in Liverpool. The motherboard consists of a LV (low voltage) card and 4 driver cards. One repeater board will amplify one side of a module. Typically these boards also have an ECS (Experimental Control System) card. However in this instance, that has been removed and control signals are sent using I2C signals from the Tell1 board. The Repeater boards were designed and built at CERN.




Figure 1.8: Tell1 Board

1.2.2.2 Tell1

The Tell1 board is a motherboard with multiple functions. As can be seen in fig. 1.8 the Tell1 consists of 4 ADCs (Analogue to Digital Converters), for converting the Analogue Beetle signals to digital. There are also 4 FPGAs (Field-Programmable Gate Arrays) which can perform Pedestal and Common Mode subtraction (but lacking this functionality at time of testing). The Tell1 also assembles the digitised data into Ethernet packets. These packets are sent out the Gigabit Ethernet card to a computer for storage and analysis. Like the Repeater Boards, the Tell1 was also designed at CERN.

There's a Credit Card PC (CCPC) for controlling the Tell1 and sending signals to the Beetle chips. The CCPC is basically a 486 CPU. It boots from a combination of DHCP+NFS servers on the main PC. The CCPC runs Scientific Linux 4. Software on the CCPC is used to send control signals to the Beetle chips. The I²C control signals are as follows - power up chips, power down chips, configure chips, reset chips, trigger (to read out N events of data).

1.2.2.3 PC

The PC is used to both take data from the module and store it. There are three network ports on the PC - one for internet, one to connect to the Credit-Card PC on the Tell1 board and one for the Gigabit Ethernet output from the Tell1. The PC runs Scientific Linux CERN 3.




Figure 1.9: Top down view of the vacuum tank

1.2.3 Thermography

An FLIR (Forward Looking Infra-Red) Thermacam SC500 thermal camera was used to take thermographic images of the module under cooling and vacuum. The vacuum tank has two Germanium windows for the thermal camera to view the module. Thermographs were taken roughly every thirty minutes when testing the module. The module was in vacuum and under cooling with its chips powered up when the images were taken. The chips are the main source of heat on the module and along with the cooling, there is a heat gradient across the module.

1.2.4 Vacuum & Cooling Systems

The vacuum tank is a cross shape and the relevant points are marked in Fig. 1.9.

There are vacuum pipes entering at the top and bottom of the tank in the centre. The tank is designed so that the module's sensor will sit in the centre of the tank when mounted. The Germanium windows allow the Thermographic camera to take pictures of either side of the module.

The front tank door opens on rails and the module is mounted with the z-axis in line with the Germanium windows. (Assume z-axis is that of the experiment and the module is oriented similarly). The door consists of electronic feed-throughs to connect the Repeater Boards to the Kapton cables and also has feed-throughs for the cooling pipes (inlet and exhaust).

There are two vacuum pumps and three gauges. The Scroll pump works by the relative motion of two spiral shaped scrolls. Gas is trapped and compressed between the scrolls and pushed towards an exhaust in the centre. The scroll pump works from Atmospheric pressure and reduces the pressure to roughly $1\text{E}-2$ mbar. At this point the Turbo pump is used to further reduce the pressure to $1\text{E}-4$ mbar. The Turbo pump cannot be operated at atmospheric pressure. It operates by pushing out individual molecules rather than a fluid-like gas. The molecules are hit (batted) successively by interleaved fins inside the turbo pump. The fins turn clockwise and anti-clockwise relative to each other. See Fig. 1.10.

The vacuum system consists of four valves - two between the pumps and the tank and two to air. The vacuum gauges are in three different positions. One is connected to



Figure 1.10: Cut-away turbo-molecular pump

the scroll pump, one to the Turbo pump and one inside the tank. The one on the scroll pump is accurate from atmosphere to $5E-2$ mbar. The Penning gauge and the gauge inside the tank are designed to measure pressures below 10^{-1} mbar.

1.2.4.1 Cooling

The cooling is a bi-phase CO_2 system. Liquid CO_2 is kept under pressure and expands and changes to a gas in the cooling pads that are attached to the module. The phase change is isothermal - the module is cooled by a decrease in pressure and the latent heat of phase change. The pressure is kept at roughly 10bar during operation.

The cooling system consists of a 30kg bottle of liquid CO_2 (not including bottle), three valves, 2 over-pressure valves, a barometer, a flow meter, and a heater. There is also a CO_2 meter as a human precaution since CO_2 is highly poisonous at 10% air concentration. One of the valves can be used to regulate the CO_2 flow while monitoring the flow meter. The heater heats the gas as it exhausts so that it does not condense and solidify causing a blockage.

1.2.5 Procedure

Fig. 1.6 gives a reminder of how the data and control signals flow in the system.

All power supplies for the Tell1, PC and Repeater board were turned on and the various devices allowed time to power up. The Credit Card PC boots up by retrieving system files from the control PC. This takes about 2 minutes. Once booted, one can log into the CCPC via ssh. The DAQ program is started and this can be used to send I^2C signals to the Beetle chips on the module.

The VELO module is placed inside the vacuum tank, sealed and the vacuum pumps started. Once the vacuum had reached the desired pressure ($\sim 10^{-4}$ mbar), the CO_2 cooling system was turned on. At 0°C the module can be configured and powered up via I^2C signals from the CCPC. A data capture program is started on the control PC.

The events are saved to file in their binary format. This binary file can be processed at a later time using the Vetra analysis program.

The system is now considered to be in a “ready” state. The following sequence of events will capture N events from the sensor on the module: I^2C triggers are sent from the CCPC through the Repeater Motherboards to the Beetle chips on the module. The chips will send one event for each trigger. The data is passed from the Beetles to the driver cards on the Repeater board. These shape and drive the signals down 5m of twisted pair cables to the ADC cards on the Tell1 board. The events are digitized and the assembled into Ethernet packets along with the header information. These packets are sent out of the Gigabit Ethernet port and stored on the control PC.

1.3 Noise and Thermal Characterisation of Velo Modules

The overall goal is to characterise and reduce the noise in the data as much as possible in order to improve the signal-to-noise measurement for the experiment.

The following sections describe the analysis performed and the algorithms written in the Vetra analysis package. Vetra is based on the LHCb Gaudi software framework. It is simply a base program from which to write analysis routines. A systematic study was performed for all modules making up the Velo Detector. Results are presented here for Module 20⁴ whose response was typical.

1.3.1 Raw Data

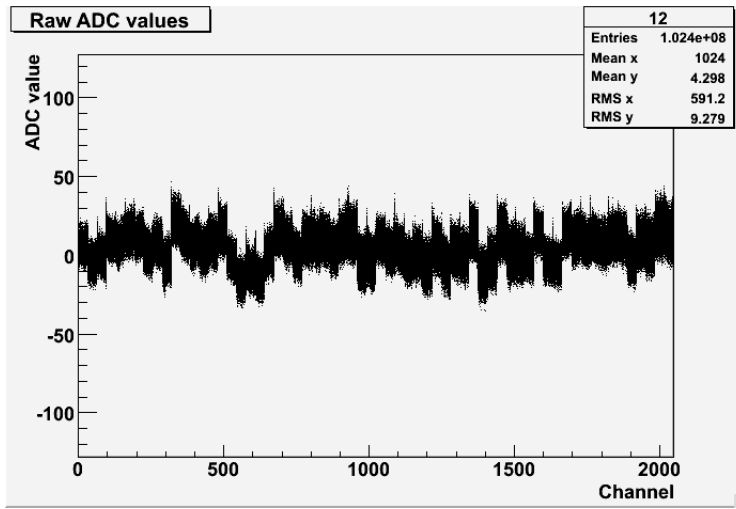
The first objective was simply to obtain a plot of ADC values by reading out data from the module. The raw noise was extracted from the data for the 50,000 events. Fig. 1.11a shows a plot of the raw data values against sensor channel. The overall r.m.s. of the values is 9.3 ± 0.1 , though significant structure can be seen. The total noise can be broken up into at least five constituents: fluctuations coming from (i) channel specific offsets, (ii) common mode noise, (iii), effects due to data header, (iv) beetle pipeline influence, and (v) the intrinsic noise. The following sections describe the procedure used to determine the intrinsic noise by subtracting the other constituents.

1.3.2 Pedestal Subtraction


It is clear from the plot (fig. 1.11a) that the data is shifted up or down relative to its neighbouring chip and channel. A pedestal subtraction was performed to normalise the values about zero. The pedestal is calculated as follows: An average of the raw values is calculated for a particular channel over all the events taken, i.e., for channel x , using N events the pedestal is

$$P_x = \frac{1}{N} \sum_{k=1}^N A_x^k$$

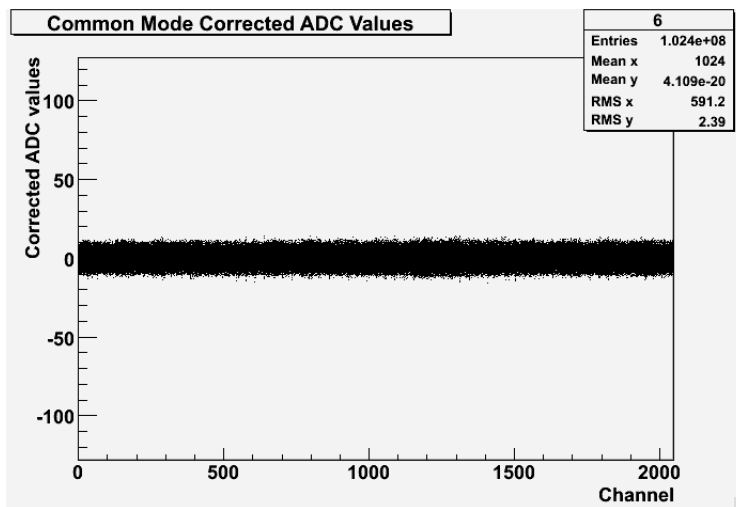
⁴Note to self: mostly using file - module20-Phi-230506_1435



(a) Raw Data



(b) Pedestal Subtracted data



(c) Common Mode corrected

Figure 1.11: Data Values

where A_x^k is the raw data value for channel x , and event k . P_x is the pedestal for channel x (i.e., a different pedestal is evaluated for each channel). A new set of pedestals is calculated for each module.

The pedestal is subtracted from each subsequent ADC value. The resultant pedestal corrected data is shown in fig. 1.11b. The new r.m.s. of the pedestal subtracted data values is 3.99 ± 0.06 .

1.3.3 Common Mode Noise

An analysis of common mode noise was performed to remove this contribution to the overall noise. Common-mode noise is the noise common to set of elements on an electrical pathway. For example, a slight electrical surge in the power grid might push all the data values up for one event and down for the next. Each Beetle chip processes 128 channels. The chip has 4 output ports, so each port outputs data for 32 channels. The port is a common electrical pathway to these 32 channels. The common mode noise was subtracted by grouping the channels into blocks of 32 and subtracting their average.

$$C = \frac{1}{n} \sum_{x=1}^n B_x$$

where C is the common mode noise for a block of n channels (32 in this case) and B_x is the pedestal subtracted data value for channel x . In contrast to the pedestal calculation, where all events were averaged, the common mode noise changes per event, so the average of all 32 channels was taken as the common mode noise for that particular event. Fig. 1.11c shows the data values after pedestal and common mode corrections. The r.m.s. has been reduced to 2.39 ± 0.03 .

1.3.4 Noise after pedestal and common mode subtraction

To obtain a plot of the noise in the data, the r.m.s. of corrected data values for each channel was calculated for the 50,000 events. Fig. 1.12 shows the pedestal and common mode corrected data values for one of the channels (771). The plot shows the data centred around zero, as would be expected after pedestal and common mode corrections. The noise is taken to be the r.m.s. of the data points on that channel. The plot shows that the noise is gaussian and an r.m.s. of 2.64 ± 0.03 was obtained for this channel.

This procedure was repeated for all channels. A plot of noise v. channel number is shown in fig. 1.13. Some structure is evident, in particular, one can see significantly higher noise on channels which are a multiple of 32, i.e., the first channel on every Beetle port. This was believed to be header spillover. On closer inspection, this can already be seen in the magnified view of fig. 1.11c (magnified in fig. 1.14).

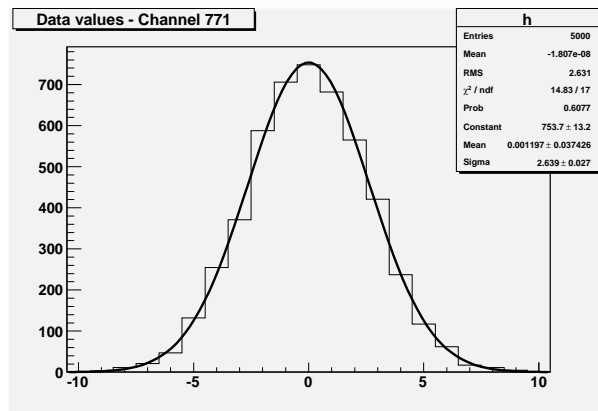


Figure 1.12: A corrected data channel





Figure 1.13: RMS Noise per channel

Table 1.1 shows how the Beetle data is organised. The cells represent frames of a data packet. The Beetle data stream is an analogue-digital hybrid. The data part is analogue, and the header is binary mixed into this stream. The header precedes the data as shown in the table. The “P” numbers refer to bits of the header that represent the Beetle pipeline number. It is these bits that affect the first data channels. The header is not a perfect binary square wave and the falling edge can spill over into the next cell modifying its data value. From the table it can be seen that the most significant bits are P0, P2, P4, and P6. So the effect of the header on the data depends on the relevant pipeline bit, which in turn depends on the originating Beetle Port. Fig. 1.15 shows four plots, each of a different channel. These channels were chosen because each is linked to a different Beetle port. The data values (pedestal and common mode corrected) for each channel were plotted against the pipeline location, and show how the header influence dependent on Beetle port.

Figure 1.14: High noise every 32nd channel

	Header			Data					
Beetle Port 0	...	P1	P0	0	1	...	29	30	31
Beetle Port 1	...	P3	P2	32	33	...	61	62	63
Beetle Port 2	...	P5	P4	64	65	...	93	94	95
Beetle Port 3	...	P7	P6	96	97	...	125	126	127

Table 1.1: Beetle Data Organisation

Due to the complicated nature of this header influence, the channels need to be corrected individually. As can be seen from the table the pipeline bits are spread over the four beetle output ports. A generic algorithm is not enough; a correction must be applied on a per pipeline bit per beetle port basis. A correction is applied by averaging the difference between the data values when a header bit is high and when it is low. The values are then normalised by subtracting this difference. Fig. 1.16 shows the RMS noise before and after this correction. The noise on every 32nd channels is reduced but not eliminated. The noise on every thirty-second channel was reduced from 2.92 ± 0.03 to 2.68 ± 0.03 . The average of the other channels, unaffected by the header influence remains unchanged at 2.37 ± 0.03 . Combined, this gives an average noise value for all channels of 2.38 ± 0.03 . The values in this type of plot are those that will be used to evaluate a signal to noise measurement for a module.

1.3.5 Pipeline Noise

The pipeline itself was analysed as a possible source of extra noise. Fig. 1.17 shows a plot of Pipeline location against channel and the colours show the noise in ADC counts. The vertical lines show that some channels have higher noise than others. Some of the strips on the silicon are longer than others and have a higher capacitance, causing higher

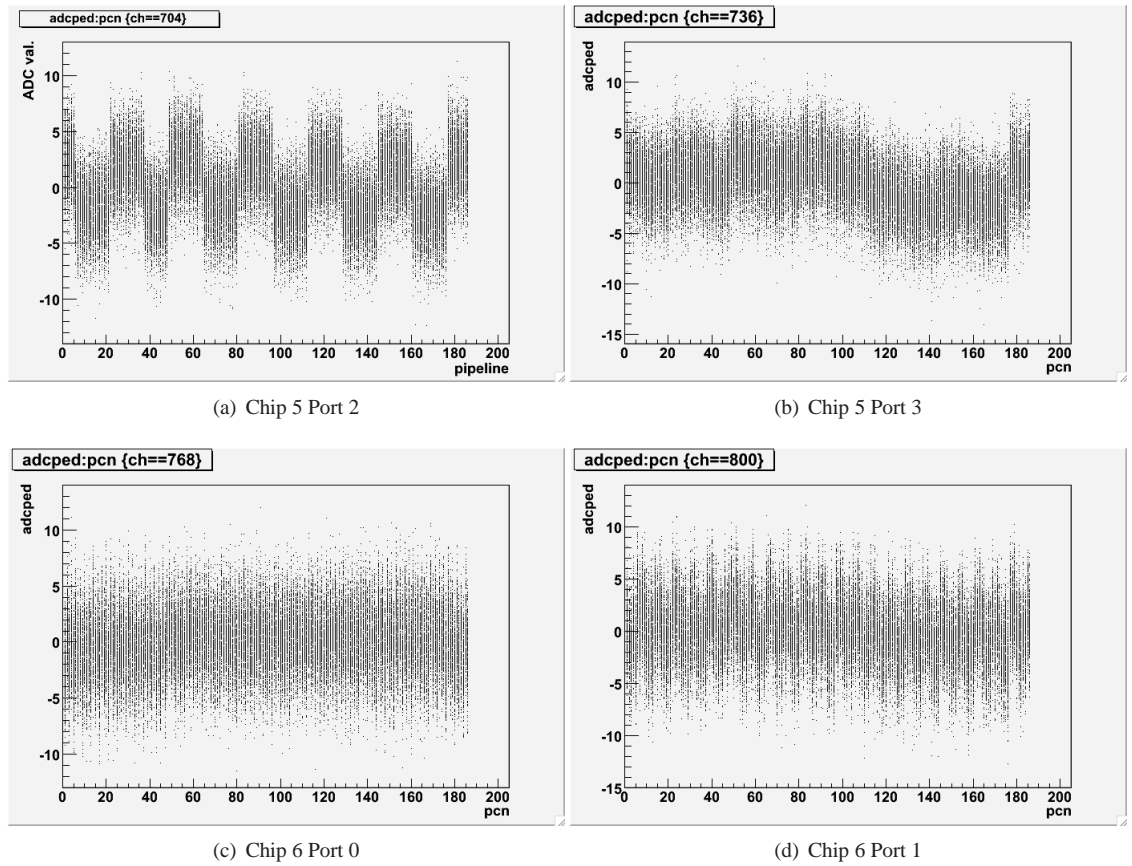



Figure 1.15: ADC values against pipeline location.

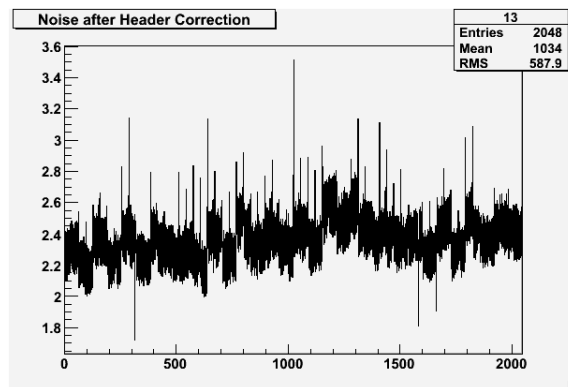
noise on these channels. Any horizontal lines would show particularly bad noise for a given pipeline location but none are obvious. Thus, we can conclude that for this particular module there are no bad pipelines.

1.3.6 Thermography

Thermographic images were taken to identify any obvious potential problem with the modules. Dead chips stand out against their neighbours, as do shorted wires, as cold and hot spots respectively. These are also used to qualitatively compare the modules. Fig. shows a thermograph of module 17. The cooling block is on the other side of this module in the area represented by the green box. It can clearly be seen that the area furthest from the cooling block is warmest. The sixteen Beetle chips can be seen “glowing” when fully powered. Any chip not powered up would look significantly darker than its neighbours. Due to reflectivity, the sensor and other components look warmer than they are relative to the hybrid.



(a) before



(b) after

Figure 1.16: RMS Noise before and after header correction

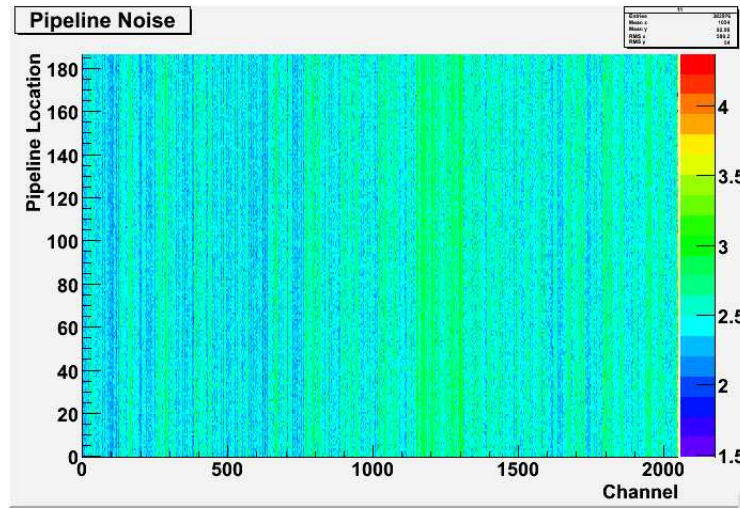



Figure 1.17: Pipeline Noise

Figure 1.18: Thermograph of module 17 (ϕ -side)

1.4 Conclusion

The DAQ system was successfully set up at the University of Liverpool and is currently being used to test VELO modules for the experiment. In the process of doing this, I have gained an intimate knowledge of the production and testing cycle of the modules at Liverpool and have become familiar with the operation of vacuum and cooling systems.

Control software for operating the Tell1 system and sending I2C signals to the Beetle chips was enhanced by writing a front-end for the Tell1 software control libraries. During my time at Liverpool, five modules were tested using the DAQ system. The data from the module is stored in a binary format when taken from the Tell1 digitisation board.

Algorithm	r.m.s.
Raw Data	9.3 ± 0.1
Pedestal Subtraction	3.99 ± 0.06
Common-Mode Correction	2.39 ± 0.03
Header Correction	2.38 ± 0.03
Channels unaffected by header	2.37 ± 0.03

Table 1.2: Systematic noise reduction

The binary data was then analysed using the Vetra program. Vetra is based on the LHCb Gaudi framework which will be used in the running experiment. Algorithms were added to Vetra to extract the raw data events, perform pedestal and common mode subtractions and produce noise plots of the corrected data. The goal of this was to reduce the systematic noise in the data. Table shows the reduction in r.m.s. values after each algorithm for Module 17 (ϕ -side).

Further corrections were made to the data to reduce “header spillover”. This is an effect of the header influencing the data when the signal is passed from the Beetle chips, effectively increasing the noise on every 32^{nd} channel. The average noise (r.m.s) on every 32^{nd} channel was reduced from 2.92 ± 0.03 to 2.68 ± 0.03 . This yields a final average noise value of 2.38 ± 0.03 .

An analysis of Beetle pipeline (or buffer) was investigated as another possible source of systematic noise. A plot was produced showing no obvious high noise from specific pipeline locations.

Thermographic images were taken of the module under vacuum and cooling to compare thermal properties of the modules. In some cases, shorted bond wires were found using this technique and sent for repair.

All of this information, both thermographic images and analysed data is stored in a database at the University of Liverpool. This database can be used as a reference for the electrical and thermal properties of the modules. All of the algorithms and software written are currently being used to test the modules that will eventually be used in the Vertex Locator for LHCb.

Bibliography

- [1] LHCb Collaboration, LHCb Technical Design Report, CERN-LHCC-98-004
- [2] Bowcock, T. et al., LHCb Velo Technical Design Report, CERN/LHCC 2001-0011
- [3] Collins, P. et al., LHCb CERN Velo Home page, <http://lhcb-vd.web.cern.ch/lhcb-vd/>
- [4] Bowcock T. et al., Velo Production Readiness Review Handbook, http://hep.ph.liv.ac.uk/lhcb/Documentation/Reviews/PRR_Home/prr_home.html
- [5] Turner, P., Design of two Silicon detectors for LHCb, Liverpool Velo Documentation webpage, 2003
- [6] T. Bowcock, G. Casse, Operating Conditions for the VELO Silicon, Liverpool Velo Documentation webpage, 2001
- [7] Young, Freedman, University Physics, Benjamin Cumming Pub., 2003
- [8] van Bakel et al., Beetle Reference Manual, LHCb 2005-105
- [9] van Bakel et al., Beetle Specification, HD-ASIC-2001