ATLAS SCT Readout and Barrel Macro-assembly Testing and Design of MAPS Test Structures

Bruce Joseph Gallop

Thesis submitted to
The University of Birmingham
for the degree of Doctor of Philosophy

Particle Physics Group
School of Physics and Astronomy
University of Birmingham

September 2005

Abstract

The semiconductor tracker sub-detector of the ATLAS experiment consists of 4088 individual silicon detector modules. A data acquisition system, SctRodDaq, required to configure, read out and test these modules both during large scale assembly and in the final experimental setting was designed, and is described here. SctRodDaq is shown to function well during assembly of the first stage of the SCT, Barrel 3.

The design and layout of pixel test structures to further the development of Monolithic Active Pixel Sensors for particle physics applications is described. Preliminary test results of the device show promise though they demonstrate the requirement of a relatively thin epitaxial layer for particle detection applications.

Author's Contribution

The ATLAS SCT, as for other modern high energy physics experiments, is the work of a large number of contributors.

A large part of my work was in the production of software for the SCT DAQ system using the final read-out hardware described in chapter 4 and known as SctRodDaq. This is the work of several people and is based upon the previous SCTDAQ software. The SCTDAQ software existed, though I contributed substantially to its stability and long term monitoring. For use during SCT module production at RAL I contributed to the development and maintenance of a web-based database system

Though the design of SctRodDaq was the work of the group, the implementation of each of the individual components was done by one or two individuals. My particular work on SctRodDaq was the development of the low-level SctApi software, described in section 4.8, and the configuration database, described in section 4.13, both essential parts. Low-level classes to access the ROD, BOC, TIM and VME drivers were written by others, to some of which I made minor contributions. Much effort was spent with many iterations of the ROD firmware and hardware, understanding its interactions with real SCT modules, and debugging problems.

The results described in chapter 5 demonstrate the first use of this software in a large scale setting. The analysis is the work mainly of others, though with much interaction with the central software. This is a clear demonstration of the viability of read-out for the full SCT system.

The final chapter describes a separate piece of work motivated by the author's background of a first degree in electronics: the results of part-time work carried out to develop new pixel system. My work is the layout of the 9 different pixels. The support electronics, submission and testing were the work of others.

Though not described in the following, I have spent some time at CERN working on the use of SCTDAQ for SEU tests in the irradiations programme and to a lesser extent for the system test and test beams. I was also present at CERN to help with the use of SctRodDaq in the ATLAS combined test beam.

Acknowledgements

First, I would like to acknowledge the supervision of Peter Phillips, Dave Charlton and Richard Apsimon.

For the SCT I would like to thank everyone involved in the irradiations, system tests and beam tests at CERN; barrel module assembly at RAL; and barrel macro-assembly at Oxford. Particular mention should go to Matt, Alan, Peter and John for discussions about SctRodDaq. Paul, Chris, Lars and Gareth for welcome to CERN and helping get to know the hardware. Gareth again for putting me up on a visit to CERN at short notice. Douglas for keeping unusual hours working on the ROD DSP code from across the Atlantic.

Work on MAPS was greatly helped by discussions with Giulio, Renato, Jaap, Arwell and Mike.

Both education and entertainment were gleaned from coffee breaks, lunches and dinners at Birmingham, CERN and RAL. Particularly everyone at Staff House and the Brass Band. Also, Sian, Kelly and John for suggestions of other things to do.

Finally, thanks to my family for all their support and I'm sorry Gran isn't around to see the finished article.

Contents

List of Figures

List of Tables

Chapter 1

Introduction

The ATLAS detector is designed to search for the Higgs boson which is predicted to be generated by the LHC collider. It is currently being built and commissioned at CERN on the Swiss-French border. The largest silicon part of the central tracker, the SCT, is being assembled in the UK and this requires testing for quality assurance at all points in the process.

The LHC, ATLAS and the SCT are introduced in chapter 2. This is followed by a more detailed description of the SCT hardware in chapter 3. Chapter 4 introduces the software used for testing and calibration and chapter 5 describes results from tests carried out on the first SCT barrel. Finally, chapter 6 describes some work done towards a next generation silicon detector.

The remainder of this chapter introduces the Standard Model, which describes the particles expected in the ATLAS detector, and how silicon detectors can be used in the detection of the products of collisions in the LHC.

1.1 The Standard Model

The Standard Model of particle physics describes a system of fundamental particles and their interactions. Many predictions of the Standard Model have been tested and found to agree very well with experiment.

The particles can be divided by their spin into three groups: the fermions, 6 quarks and 6 leptons, have spin $\frac{1}{2}$ † and are responsible for matter, they each have their own antiparticle; the force-carrying bosons, the photon (γ), $W^{\pm}$, Z, gluons, have spin 1 and are their own antiparticles; and the Higgs which is discussed further below, it has spin 0 and is unobserved. Figure 1.1 shows a diagrammatical representation of these particles.

The interactions between these particles form into distinct groups according to which bosons are involved. The electromagnetic force is responsible for the everyday forces, electricity and magnetism and is mediated by the photon. The W occurs in short-range weak decays, as in radioactivity. The short range implies a large mass for the W. The gluons bind the quarks into hadrons, pairs in mesons and triplets in baryons. The force is not felt outside the hadron as it is made neutral from the point of view of the colour charge. The most obvious force at scales larger than everyday, gravity, isn't described by the Standard Model but has a negligible effect on the particles and interactions at the smallest scales.

The electro-weak interaction is the result of the unification of electromagnetism and the weak nuclear force. This predicted the Z, which is another massive boson. Its couplings to the fermions are given in terms of the mixing angle θ_W. They have been studied in great detail at the LEP collider [1] and found to agree very well with the Standard Model.

These particles and interactions are described by a relativistic quantum field theory. In order that the theory can be used to make physical predictions, a useful feature is its

† $\hbar$ and c are taken to be 1 for the remainder of this thesis

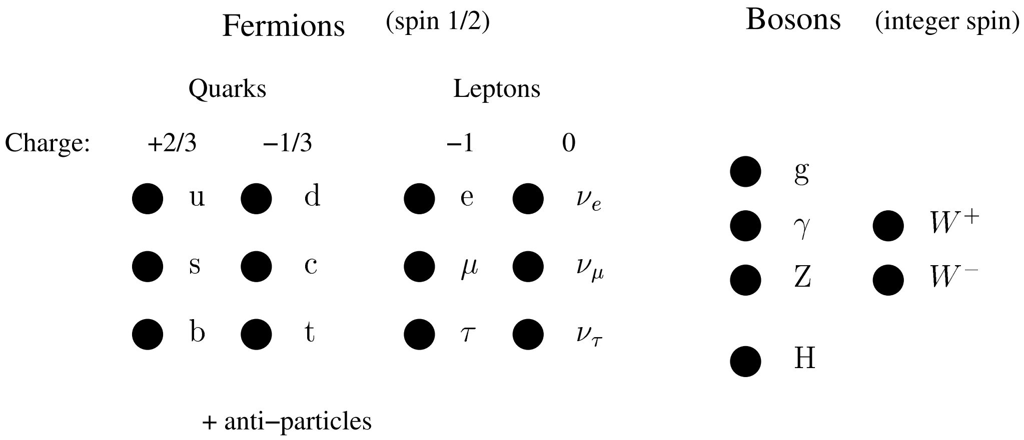

Figure 1.1: The standard model particles

renormalisability. It is therefore believed that the theory must also be gauge-invariant, which has been shown to imply renormalisability [2, 3].

The matter around us is made up of u and d quarks - bound into protons and neutrons - and electrons. This means that in order to learn anything about the other particles, we must observe the interactions between these particles and either the particle of interest or its decay products.

The properties of the Standard Model particles have been probed at particle colliders over many years, and all have been observed except for the Higgs. In order to probe for any particle, a machine that collides particles with a centre of mass energy higher than the rest mass of the new particle, is required. This is true of the Higgs which is described in more detail in the following section.

1.1.1 The Higgs Boson

The problem with the currently observed theory, is the difference in mass between the 4 electroweak bosons: the photon has zero mass, the W 80.4 and the Z 91.2 GeV. The zero mass photon arises naturally from the gauge invariant theory for the QED interactions, but the direct addition of a mass term for the massive bosons will break the symmetry and more seriously the renormalisability of the theory.

An alternative is to add mass to the particles using the mechanism of spontaneous symmetry breaking, where the mathematical formulation remains symmetric, but the physical system prefers a state in which the symmetry is broken.

The Standard Model specifies the Higgs mechanism for breaking the electro-weak symmetry. This introduces a complex doublet of spin 0 particles, three of which are subsumed as the W and Z masses and one which is physically present as the Higgs boson. This mechanism was used to predict the mass of Z from the mass of W:

$M_W^2 = \rho M_Z^2 cos^2\theta_W$

Where M_W and M_Z are the masses of the W and Z respectively, θ_W is the electroweak mixing angle and ρ is a correction factor which deviates slightly from 1 for higher order perturbations.

The mass of the Higgs boson is the remaining parameter that predicts all the other properties of the Higgs within the Standard Model. This means that though the mass itself is unknown, for a given mass its behaviour and hence how it may be observed can be fully determined. There are arguments [4] why it should be under 1 TeV and there are previous experiments which would have found a Higgs with a mass below 115 GeV [5]. Figure 1.2 shows the constraints from the LEP experiment: the shaded portion is excluded due to the Higgs not being found directly and the band shows the probable mass based on fits to the Higgs loop correction of the high statistics Z measurements.

In the case that the Standard Model Higgs is not found there are many other theories that have been proposed. A general test for new theories involves measuring interactions where the Higgs is required in the Standard Model. One example is in W-W scattering where, in the absence of the Higgs, the rate of production rises unphysically with increasing

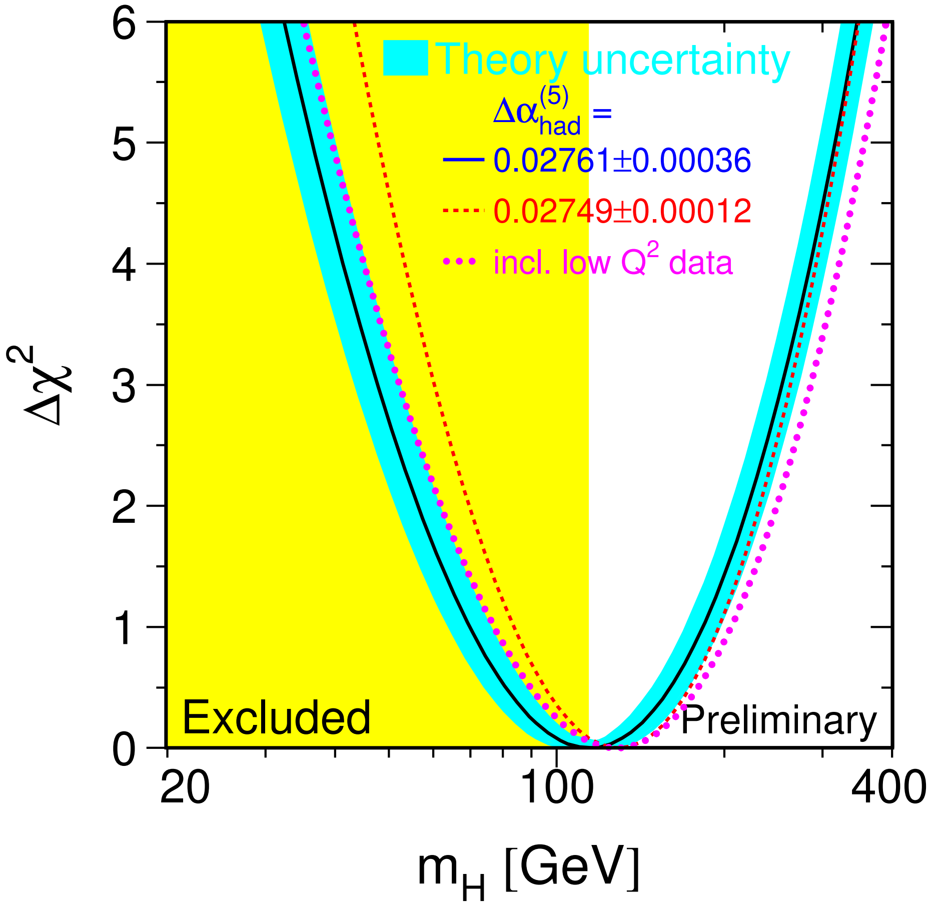

Figure 1.2: Constraints on the Higgs mass from LEP measurements [5]. $\Delta\chi^2$ represents the quality of a fit to the measurements for the different values of m_H

energy.

A popular extension of the Standard Model, which includes the Higgs, involves the introduction of super-symmetry, which adds an additional supersymmetric mirror particle for each currently known particle. The simplest form, MSSM (minimally super-symmetric standard model), has two doublets of complex fields which give rise to 5 Higgs particles, known as A, H, h and $H^{\pm}$.

1.1.2 Proton-proton Colliders

The LHC is a proton-proton collider. It has advantages in that the energy available in a collision is higher than in previous colliders, as described in section 2.3, but colliding hadrons also introduces unique problems. As the energies used increase, the proton is probed to very small distances. This means the proton can no longer be regarded as a point particle and must instead be regarded as a collection of partons, or the constituent parts of the proton. As a proportion of interactions at high energies, these are mostly gluons and some sea quarks.

The strongly interacting gluons and quarks produce a large number of QCD particles at a much higher rate than any interactions caused by new physics, see figure 1.3. As an example, the total proton-proton inelastic cross-section is estimated as 80mb, in comparison to 120pb for production of a nominal Higgs ($m_H = 120GeV$). This is a difference of 9 orders of magnitude. Due to this difference, a very large number of collisions are required in order to generate a statistically significant sample of events containing new physics.

Another problem caused by colliding protons is that a large number of hadrons are produced, including many neutrons. This large amount of radiation affects the material used to detect the particles adversely. This is described in section 1.2.4.

Once produced, the Standard Model Higgs will immediately decay into different par-

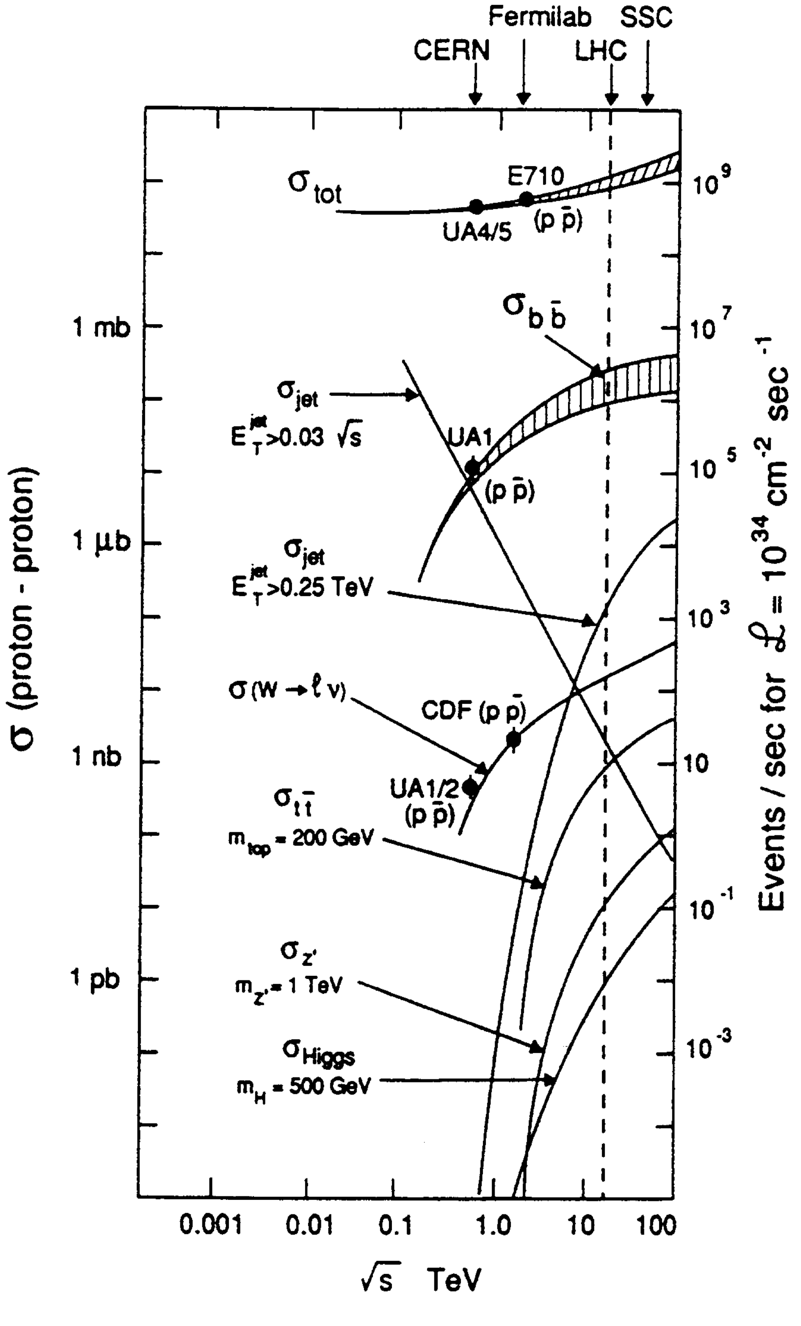

Figure 1.3: Cross-sections of p-p collisions as a function of cm energy, from [6]

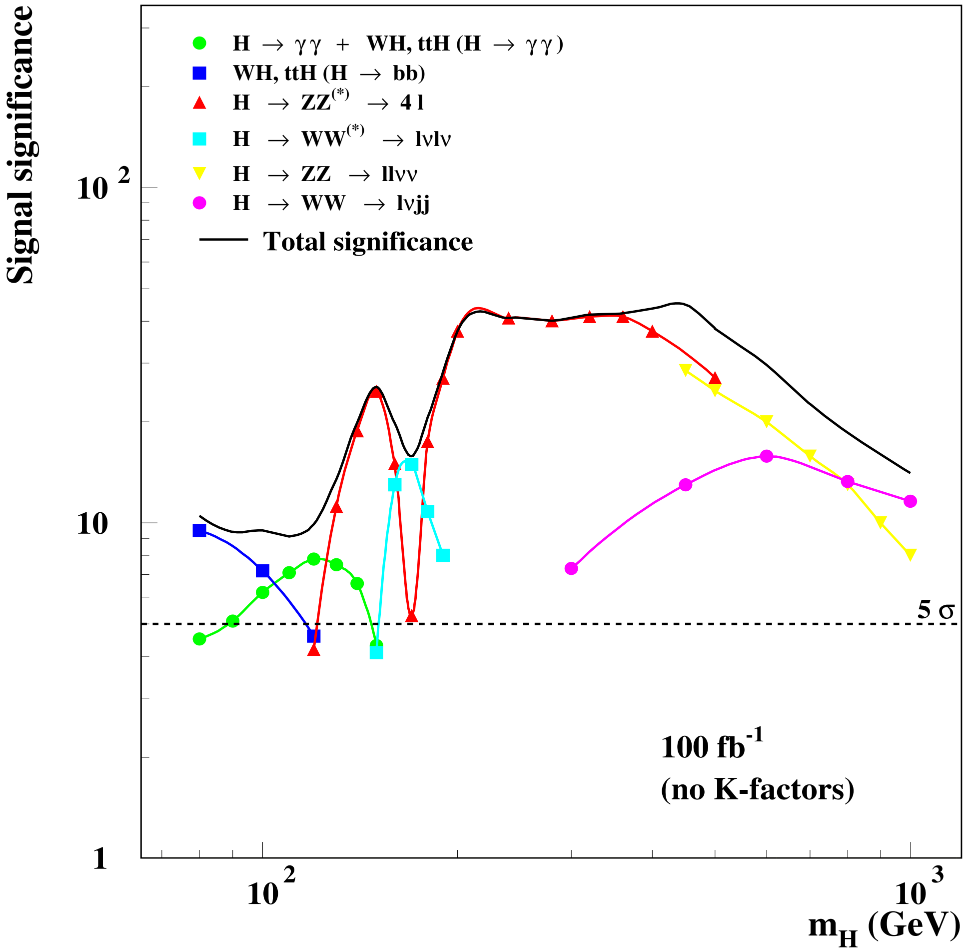

Figure 1.4: Potential for discovering the Higgs at ATLAS as a function of Higgs mass, from [7]

ticles as predicted by the Standard Model and dependent only on the unknown mass. Figure 1.4 shows the discovery potential of the Higgs at the ATLAS detector for the predicted range of masses. This is a function of the branching ratios to the different decay products, the background processes that generate similar decay products and the sensitivity of the ATLAS detector to the decay products. It can be seen that a Standard Model Higgs should be clearly discovered over the mass range 100-1000GeV.

Many of the decay modes include leptons, which must be reliably distinguished from the hadronic jets. Particularly useful is the ability to distinguish long lived muons. Also, in a decay to neutrinos, the only way to detect them is by measuring the absence of momentum needed to balance that seen in other particles. Decays of b quarks are recognisable by their relatively long life which allows them to travel several hundred microns before decaying further. More information on particle identification in the ATLAS detector in particular is provided in section 2.2.2.

1.2 Silicon Detectors

Particles interact with matter via interactions with either the electrons or nuclei of atoms. The electron field is most likely to be hit if the particle is electrically charged, resulting in ionisation of the material by knocking an electron away from its home atom. In semiconductors, such as silicon, this creates a so called electron-hole pair. The hole represents the absence of an electron and can migrate through the transfer of electrons between atoms.

1.2.1 Interaction of Particles with Matter

The mean energy loss due to ionisation by interactions of charged particles with a material is described by the Bethe-Bloch function [8, 9]. This is written in terms of the charge and speed of the incident particle and also properties of the material in question.

$$-\frac{dE}{\rho dx} = Kz^2\frac{Z}{A}\frac{1}{\beta^2}\left[\frac{1}{2}ln\frac{2m_ec^2\beta^2\gamma^2T_{max}}{I^2} - \beta^2 - \frac{\delta}{2}\right]$$

Where K is the constant $4\pi N_A r_e^2 m_e c^2$, and m_e, r_e, c, N_A and K are the standard constants: the mass of the electron, the classical radius of the electron, the speed of light and Avogadro's number respectively. Z, A and ρ are the atomic number, atomic mass, mass density of the detector material. β is the speed of the incident particle and γ is the Lorentz factor. Finally, z is the amount of charge on the incident particle; T_{max} is the maximum kinetic energy that can be imparted to a free electron in a single collision; I is the mean excitation energy; δ is a correction factor. There are additional correction factors at low energy that are not shown here.

In general, low velocity particles have more time to interact with the material and lose energy quickly. This rate falls as the velocity increases, with $\frac{1}{\beta^2}$. It then rises logarithmically, as relativistic effects become prominent, to the Fermi plateau. The minimum is at approximately $\beta\gamma = 3$ to 4.

The structure of this function gives rise to a relatively well defined point at which the particle minimally excites the material. The energy lost to the material at this point is $\frac{dE}{\rho dx} \approx 2MeV cm^2/g$ and a particle with the appropriate energy is referred to as a minimum ionising particle or MIP. This energy loss can be measured by the number of electrons which are ionised, as each electron is ionised by the same amount of energy. A detector can therefore be designed to detect a certain number of electrons, as ionised by the MIP, and be capable of detecting any other ionising particle which will, by definition, produce more electrons.

The amount of charge liberated by ionisation in a sheet of material for a given incident particle energy is described by the Landau function. This is a distortion of a Gaussian with a long tail towards higher charges, due to the possibility of a small number of particles depositing disproportionately large amounts of energy.

In order to detect the ionisation electrons produced by an incident particle, a detector allows them to migrate in an electric field created by a potential difference across the detector medium. This flow of electrons causes a build up of charge, which can then be measured by supporting electronics.

A common material used for detecting ionising particles is silicon. Its use relies on the semi-conductor properties of silicon and the depletion of free carriers from the silicon using a reverse biased p-n junction, as described in the following.

1.2.2 Silicon Detectors

Silicon is a semiconductor, meaning that the energy gap between the valence and conduction bands is small but not zero (as in a conductor). These have advantages over conductors, in that the small energy gap means an ionising particle can liberate a large number of electrons to be counted. The electrical characteristics of the material can be changed by doping the crystal with group III (eg boron) or V (eg phosphorous) elements. These are called acceptors and donors, respectively, because they lack an electron found in silicon or provide extra electrons to the valence band.

Free carriers are electrons or holes which can move independently of their home atom, in the conduction band. They are liberated by the thermal energy in the material. In doped silicon the majority free carriers are electrons in n-type silicon and they are holes in p-type silicon.

Once free, electrons and holes migrate due to drift and diffusion. Drift describes movement of charges under the influence of an electric field. Diffusion describes their net movement under random motion which is a function of the density gradient.

1.2.3 p-n Junctions

In order to detect the electron-hole pairs generated by an ionising particle, the signal must be easily detectable above the background, which in this case comes from the collection of the free carriers. Therefore a method of removing as many as possible is required.

A p-n junction is formed at the boundary between the silicon doped with acceptors (p) and that doped with donors (n). In an area either side of the junction, these carriers diffuse across the boundary and recombine with each other, creating a space charge layer in which there is a net positive/negative charge but no free carriers.

The potential difference this creates generates an electric field. This forms a potential barrier which stops most of the current flowing across the junction. This junction behaves as a standard diode; when forward biased, a current flows across the junction. The built-in potential can be increased by reverse biasing the junction, which has the effect of increasing the size of the space charge layer. This is also known as the depletion region due to the absence of free carriers. The remaining current, called the leakage current, is caused by thermally generated pairs which migrate, under the influence of the electric field, across the junction.

1.2.4 Effects of Radiation in Silicon

In addition to the useful ionising interactions with silicon, that enable the detection of particles, there are other interactions, with both temporary effects and long term damaging effects. Temporary upsets are called SEUs (single event upsets) and involve the generation of a large amount of charge in a place other than the detector. This is most likely to manifest itself in an error in the state of digital electronics.

Other effects are caused by the accumulation of defects caused by exposure to many particles. Interactions with the nucleus of an atom can result in displacement of the nucleus within the lattice, or a nuclear reaction which may change the charge of the nucleus.

Displacement of a nucleus results in a vacancy, where the nucleus was, and the interstitial nucleus itself. These can migrate through the crystal lattice and form relatively stable defect complexes, mainly through interaction with impurities in the silicon. These complexes may behave as donors or acceptors, which results in changes in doping type over time. For example, the main part of the SCT detector (see section 3.3.1) is initially n-type sandwiched between more heavily doped p-type silicon at the surface and n-type silicon at the base (p+ - n - n+). After irradiation the n-type behaves like p-type so looks like (p+ - p - n+). This means that the p-n junction has moved from the top of the wafer to the bottom.

These defects also lead to: increases in the leakage current, ie the current flowing through the junction when no particles are incident; loss of mobility, hence increased recombination and lower charge collection efficiency. Both of these effects lead to a reduction in the signal to noise ratio.

In addition to the detector, transistors in the supporting electronics are also damaged by the radiation. In silicon dioxide, used in the gate oxide of MOS transistors and which is an insulator, displaced electrons and holes, created by ionisation events, drift in the electric field and remain trapped near to the boundary of the material. This has the effect of introducing a fixed charge which affects the behaviour of the transistor. One effect is the reduction in the gate voltage required to turn on a transistor, which can be alleviated by using a thinner gate oxide. Another effect is the creation of parasitic transistors. These are formed around the edges of a transistor through interactions with the oxide that is used to insulate neighbouring transistors. The effect is to change the transistor parameters, including an increase in leakage current. As this is a function of the geometry of the transistor, using a different geometry can help reduce this effect. An example of this is shown in chapter 6.

1.3 Summary

Particle colliders are used to probe the properties of particles. Currently, one of the most interesting particles is the Higgs boson, but very little is known about it. The ATLAS experiment at the LHC collider, described in the next chapter, is designed to investigate this and other particles that might be found at high energies.

One of the materials used to detect the collision products is the semiconductor silicon. This is used both in the SCT strip detector described in chapter 3, and in the experimental HEPAPS3 pixel detector described in chapter 6.

Chapter 2

Detector Overview

The ATLAS detector is designed to detect the products of collisions in the LHC proton-proton collider.

This chapter describes the LHC and ATLAS. It starts with an overview of the LHC accelerator and then of the ATLAS detector. Next, it gives a description of the sub-detectors that make up the ATLAS detector. It ends with the flow of event data from the system.

2.1 LHC

The LHC (Large Hadron Collider) [10, 11] is a 14TeV centre of mass energy proton-proton collider, being installed in the 27km LEP (Large Electron Positron collider) [12] tunnel at CERN. It has almost an order of magnitude higher centre of mass energy than current hadron colliders, the highest energy collider is currently the Tevatron [13] which collides protons and anti-protons with a centre of mass energy of 1.96TeV.

In order to store particles accelerated to a higher energy, either the ring housing the accelerator needs to be bigger, or the magnets used to keep the particles in the ring must be stronger. Particles in a synchrotron suffer from energy losses due to synchrotron radiation as they are bent. This energy loss increases proportional to γ^4, where γ is the Lorentz factor for the particle. The energy of the LEP e^+-e^- collider was limited by these synchrotron radiation losses to a centre-of-mass energy of 210GeV. In the LHC, protons, which are about 2000 times heavier than electrons and therefore have a much lower γ for a given energy, are used instead.

As an alternative to colliding protons with oppositely charged anti-protons, as in previous experiments, the LHC collides two protons. Anti-protons are difficult to create, requiring large numbers of protons to be collided on a fixed target to produce each one, but have the advantage that the oppositely charged beams can be contained in the same magnetic channel. Colliding protons with protons requires the beams to follow two different channels through the bending magnets, with the field in opposite directions in each. In the LHC this is achieved using a complex magnetic field, see figure 2.1, allowing the two beams to be contained in the same magnet.

This enables the same tunnel as the LEP collider to be used with stronger magnets to produce higher energy collisions. The LHC accelerator consists of superconducting magnets, to bend the particles, running all the way around the tunnel, cooled using super-fluid helium to 1.9K.

Another consequence of the increased energy is that, in order to achieve the required event rate to see the low cross-section processes involving high energy particles, the luminosity must be much higher, of the order of $10^{34}\mathrm{cm}^{-2}\mathrm{s}^{-1}$.

Particles are injected into the LHC by the same injector chain as used for LEP, see figure 2.2. Protons come from the SPS (Super Proton Synchrotron) at 450GeV which come in turn from the PS (Proton Synchrotron) at 26GeV. The protons are initially accelerated by a small linear accelerator and subsequent booster before being injected

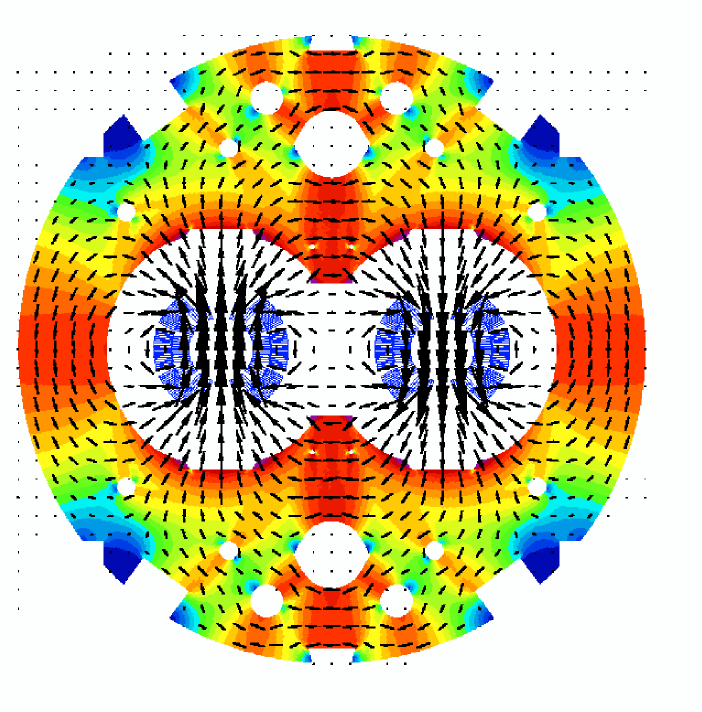

Figure 2.1: Magnetic flux in LHC dipole, showing twin cores with opposite field, from [14]

into the PS. The protons are accelerated in bunches which are spaced by 25ns intervals around the LHC ring. This results in the protons in the beams colliding at 40MHz.

Around the LHC ring are 8 straight sections, providing space for other parts of the machine, see figure 2.3. One section contains the RF cavities which accelerate the protons from 450GeV to 7TeV; another holds the beam-dump, where the particles are dumped at the end of a run. Two sections are used to clean the beam of any halo which can form due to unavoidable instabilities in the machine. The remaining four sections are used as interaction points, where the proton beams can be steered and focused on each other to generate collisions. Two general purpose experiments, ATLAS and CMS lie opposite each other. A B-physics experiment (LHCb) and a heavy-ions experiment (ALICE) lie at the interaction points on either side of the CERN main site.

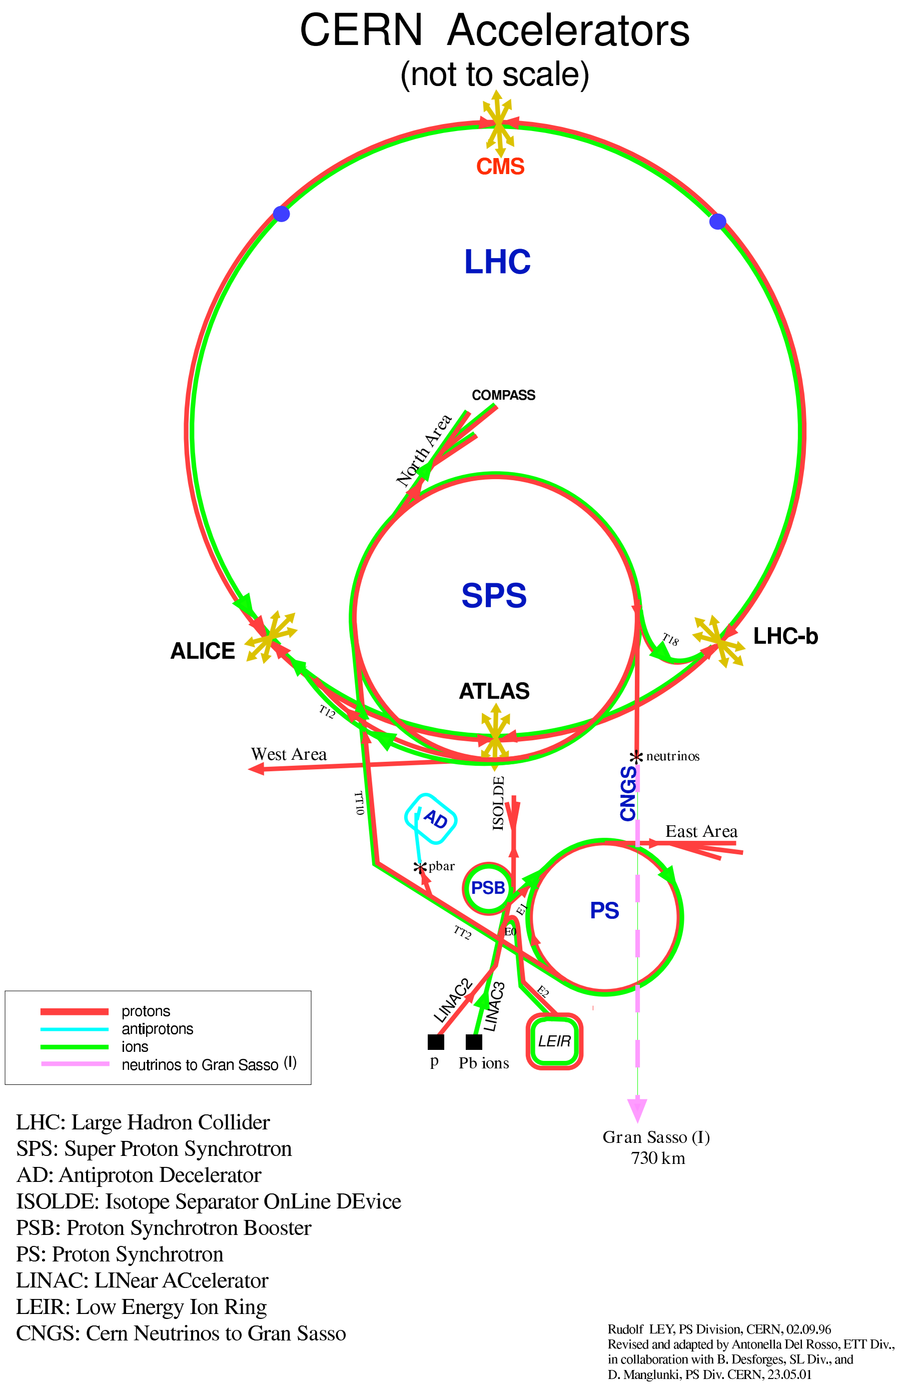

Figure 2.2: Injector chain for the LHC collider, from [15]

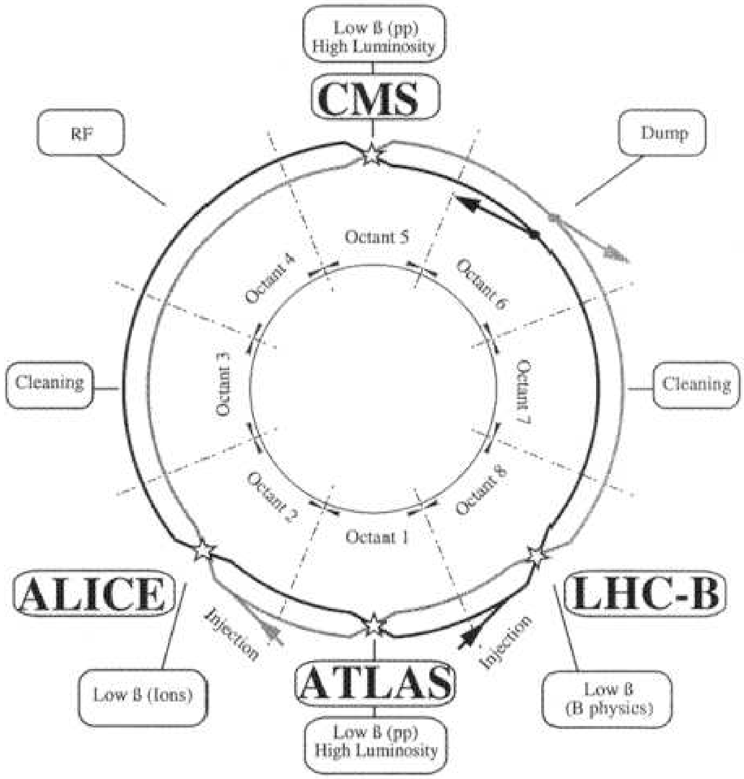

Figure 2.3: Octants of the LHC collider, from [11]

2.2 ATLAS

ATLAS (A Toroidal LHC ApparatuS) [16, 7] is a detector at interaction point 1 on the LHC ring, closest to the CERN main site. It is a general purpose detector, able to detect the particles expected in Higgs decays and other theories, beyond the Standard Model. One of the key features is its hermeticity, allowing estimates to be made of missing momentum as signatures of neutrinos and other invisible particles. A graphical representation is shown in figure 2.4. The cavern, known as UX-15, to house the detector has been completed and installation of the experiment is underway. Two smaller caverns (US-15 and USA-15) are also present on either side of the LHC ring, where the radiation environment is much less harsh. These are used for services, such as power supplies, cooling and data acquisition.

2.2.1 Co-ordinate system

There are several co-ordinate systems used within ATLAS, depending on context. At a global level, an origin is defined at the nominal interaction point: z describes distances along the beam-pipe; r measures the distance perpendicular to the beam-pipe and ϕ measures angular distance around the beam-pipe. Other co-ordinates are θ and η. Firstly, θ can be used to describe directly the angle a particle track makes to the beam-pipe, but more usually the pseudo-rapidity, η, is used.

When particles emerge from the collision point, they have components of momentum in both the longitudinal (z) and transverse (r) directions. Rapidity is a Lorentz invariant quantity that is used to combine the two; it is defined as:

$y = \frac{1}{2} ln \frac{E+p_L}{E-p_L}$

Where E is the particle energy and p_L is the longitudinal momentum of the particle. As this quantity varies for a given θ with the rest mass of a particle, the pseudo-rapidity

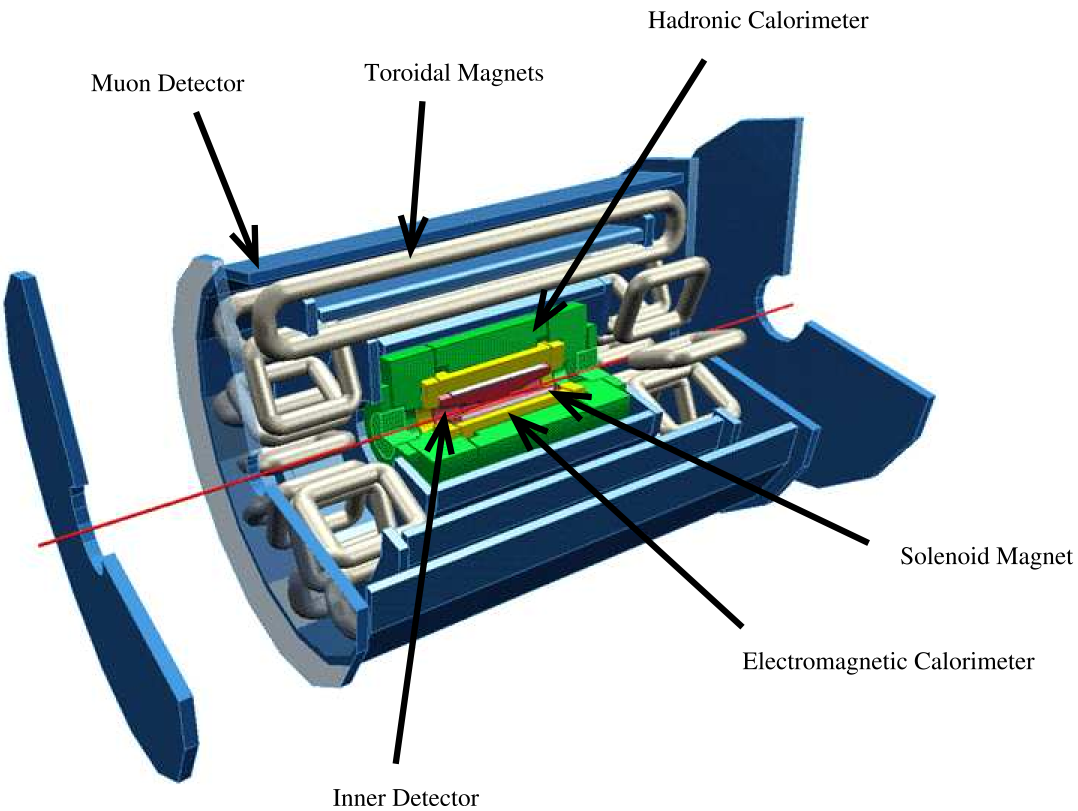

Figure 2.4: A cut-away schematic diagram of the ATLAS detector, from [17]

is used to describe the coverage of the detector. This is is defined as:

$\eta = -ln[tan(\frac{\theta}{2})]$

and is equivalent to rapidity for particles with zero rest mass.

2.2.2 Particle Detection

The aim of ATLAS is to detect as far as possible the final state of the proton-proton interaction. This consists of the types of particles that were produced, how much energy they have, and the direction they are travelling. The particles that may be produced consist of the currently known charged leptons, photons, hadrons and neutrinos, and possibly new sorts of stable particles. It is impractical to detect neutrinos, so as an alternative, the momentum balance in the detector is quantified as well as possible.

Some particles produced by the primary event decay very quickly, for example τ leptons and B hadrons. In the time taken to decay, a tau lepton may travel of the order of 100μm and a B meson as much as 1mm. These may be detected by their decay products as with other short lived particles, but they travel far enough that there is a detectable offset between the original vertex of the subsequent tracks compared to the original vertex of the other collision products, known as a displaced secondary vertex.

Particle tracks are used to discover the momentum of charged particles using a magnetic field. This bends the particles into a helical path where the radius of curvature of the track is proportional to the momentum of the particle. The track can also be used to find the sign of the charge. The energy of both charged and neutral particles are discovered by absorbing them in the calorimeters.

Particle types are mainly distinguished by their differing interactions with matter as shown in figure 2.5: charged particles interact with the material in the tracking regions, ionising the material, thus making them detectable; electrons and photons interact elec-

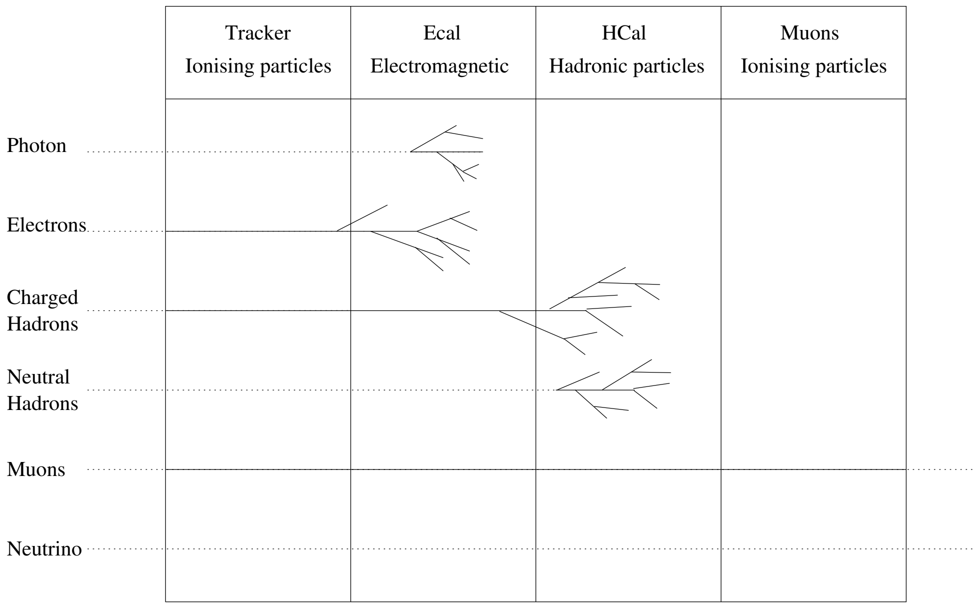

Figure 2.5: Particle identification by detection in different regions

tromagnetically with the electrons and nuclei in the denser electromagnetic calorimeter to produce showers of e's and γ's; hadrons interact strongly with the nuclei of the matter, but at a lower rate than the electromagnetic particles, and particularly in the hadronic calorimeter, finally, muons interact less than the other particles and remain to interact with the muon detectors on the outside of the experiment. Additionally, some particle types can be distinguished based on the velocity of the particle; this enables electrons (with low mass) to be distinguished from other charged particles with a similar momentum.

2.2.3 Triggers

The LHC has a bunch crossing rate of 40.08MHz and a readout of the whole detector contains ~1MB of information. Full readout of the detector at full speed would require a very high bandwidth and an inordinate amount of storage space.

As described in section 1.1.2, the rate of production of particles, such as the Higgs, which are of interest, is very low compared to the background produced by inelastic collisions. As there are identifiable signatures for the interesting particles, it is possible to trigger readout based on these signals and throw away the majority of the data before it needs either bandwidth or storage space. In ATLAS, these signals are such things as high or missing P_T (transverse momentum), the presence of a muon, or a displaced secondary vertex (signifying production of a b-quark).

ATLAS has three levels in the trigger system, each successive level taking more time to process more of the data in an event in order to cut the event rate.

Only a small part of the detector takes part in the first stage, called Level 1. The final decision to accept the event is made by the Central Trigger Processor (CTP), based on information from specialized parts of the muon chambers and calorimeters. Using this information the event rate can be reduced from 40MHz to a maximum of 75kHz. The decision is made by the trigger processors with a maximum latency of $1\mu s$, allowing time for the signal to be distributed among the other detectors in a total of $2.5\mu s$. In addition to the decision, the Level 1 trigger specifies Regions Of Interest (ROI), which are sectors of the detector that contain the large energies or muon particles on which the decision was based, and which may be of interest to higher level triggers.

The Level 2 trigger system further reduces the event rate. It has access to data from all the sub-detectors but restricts itself to the particular ROIs that Level 1 specified. Triggers are generated at an average rate of 1kHz with a maximum of 3kHz.

The data corresponding to the Level 2 events is sent to Level 3 (also known as the Event Builder). The Event Builder makes use of all the information from the detector for the specified bunch crossing to take a final on-line decision. Data is then written to disk for off-line analysis at about 200Hz or around 320MB/s.

2.2.4 ATLAS Overview

The detector itself is arranged in a cylindrical-onion like formation. It is divided into sub-detectors, which use different particle detection methods and most of which have a common barrel and end-cap structure. The barrel measures the positions of particles in ϕ and z at given radii, and the end-cap measures positions in ϕ and r at known z positions. The end-caps enable coverage of a wider range in η, where extending the barrels would provide little useful information.

The different sub-detectors are shown in figure 2.4. At the very centre of the detector is the beam-pipe, where the collisions occur. Next, going outwards, is the Inner Detector, which provides tracking for measuring particle momenta. This is itself divided into three: the silicon vertex detectors, called the Pixel detector; silicon strip detectors, the Semi-Conductor Tracker, and the Transition Radiation Tracker, which is a gas straw tracker.

Surrounding the Inner Detector are the calorimeters. The inner calorimeter measures the energy of electromagnetic particles and a fraction of the energy of the hadrons, the remaining hadronic energy is measured in the second 'hadronic' calorimeter layer.

Finally, muon detectors form an outer shell around the detector, which once again use a tracking system to determine the momentum of the muons, which have passed with only minimal interactions through all the previous detectors, before they escape the detector. This requires a large magnetic field in the muon volume which is provided by the barrel and end-cap toroids.

2.3 Inner Detector

The inner detector (ID) [18] is designed for the measurement of charged particles by tracking their positions as revealed by the ionisation of the material. This includes allowing an

Detector	Position	Area/m^2	Resolution $\sigma/\mu m$	Channels/10^6	η coverage
Pixels	B-layer	0.2	$r\phi = 12, z = 66$	16	± 2.5
	Barrel	1.4	$r\phi = 12, z = 66$	81	± 1.7
	End-cap	0.7	$r\phi = 12, r = 77$	43	$1.7 - 2.5$
SCT	Barrel	34.4	$r\phi = 16, z = 580$	3.2	± 1.4
	End-cap	26.7	$r\phi = 16, r = 580$	3.0	$1.4 - 2.5$
TRT	Barrel		170 per straw (ie $r\phi$)	0.1	± 0.7
	End-cap		170 per straw (ie $r\phi$)	0.32	$0.7 - 2.5$

Table 2.1: Resolution and Coverage of Inner Detector

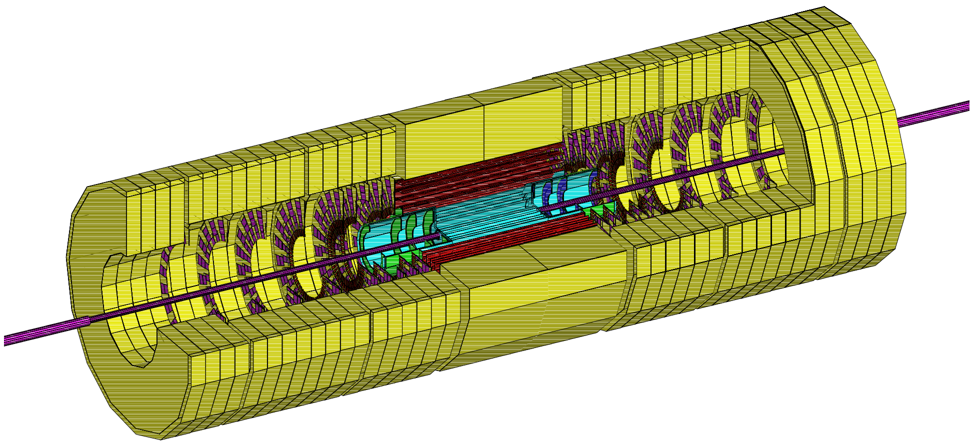

Figure 2.6: The inner detector, cut away to show SCT, Pixels and beam-pipe, from [18]

accurate reconstruction of the primary interaction point, and also of displaced secondary vertices. The ID is surrounded by a solenoid, allowing the curvature of the particles in the magnetic field to be used to work out the momentum. This provides a nominal 2T field that bends the particles into a helix. Three different tracking methods were chosen, suited to increasing distances from the beam pipe, they are represented in the three sub-detectors of the ID: the Pixels, a silicon vertex detector; the SCT, a silicon strip detector, and the TRT, a transition radiation gas wire tube tracker. The main features of these detectors are shown in table 2.1.

The inner detector is engineered as a whole, see figure 2.6, allowing it to be placed in the detector more easily and also allowing environmental controls to be shared between the three sub-detectors. Both the SCT and TRT barrel sections are contained in a volume of dry CO_2, which is kept at room temperature. This is the preferred environment for the TRT. Within this volume the SCT is hermetically sealed in a container which is filled with cold (around $-10°$C), dry nitrogen. The container for the pixel detector is separate so it can be removed independently, and is also filled with cold dry nitrogen.

2.3.1 Pixels

The Pixel vertex detector [19] is the innermost ATLAS sub-detector. It records positions of particles with a high spatial accuracy in at least three planes. In the barrel points are recorded in z and ϕ with a fixed r; the end-cap fixes z and records the r coordinate.

The detector is made up from many similar modules. The barrel and end-cap each have a single module design. The support structures are made up of three barrels and three disks at each end. The barrel is made up of two permanent barrels (at radii 9.9 and 12.3cm) with an extra layer of detectors at a very small distance (5.1cm). This layer is especially useful for b-tagging and is referred to as the B-layer. The three disks are placed at z's of between 11 and 20cm. The B-layer provides the same coverage ($|\eta| < 2.5$) as the remaining detector (both barrel and end-cap), but at a higher resolution.

Due to its close proximity to the beam, the radiation dose received by the pixel layers is extremely large, the inner permanent layer will receive 1×10^{14}n/cm^2 of 1MeV neutron equivalent and 50Mrad of ionising radiation over the 10 year LHC operation. The central B-layer receives 5 times this dose, which has consequences for the life time of this layer and it is designed to be replaced after a few years running.

There are 1456 barrel modules and 288×2 end-cap modules each with over 50,000 $50 \times 400\mu$m pixels. In the B-layer the pixels are slightly shorter (300μm). Each module covers an area of 16.4mm$\times$60.8mm. Each of 16 chips on each module reads out 18×160 pixels.

The modules are based on pairings of detector wafers and readout chips, which are bump-bonded to each other. The wafers are biased to a voltage of around 100V. The first level of readout is carried out in the 16 chips on each module. Each pixel on the detector corresponds to an area in the readout chip which contains an amplifier and discriminator. The information from each pixel is read out by a circuit per column which compresses the binary data (removing the zeros) and stores it while waiting for the level 1 trigger accept. On reception of the level 1 accept signal, a module level chip collects data for a single bunch crossing from the buffers and passes the event data to a control chip for transmission off the detector.

2.3.2 Semi-Conductor Tracker

The next division of the inner detector is the SCT silicon strip detector. This provides a balance between the higher resolution of the pixel detector and the need to put as little material as possible before the calorimeters.

The detector medium is a silicon crystal semi-conductor, with embedded strips separated by 80μm and biased with a voltage of around 150V. Each module is made up of two pairs of silicon detectors each read out by 6 readout chips, which each read out 128 strips.

The barrel portion of the SCT consists of 2112 barrel modules, covering the rapidity range $|\eta| < 1.4$. The modules are arranged on four concentric cylinders, at radii 30, 37.3, 44.7 and 52cm each with 32, 40, 48 and 56 rows of 12 modules respectively.

The end-cap structure is supported by two cylinders, and has 988 end-cap modules divided between 9 discs positioned at each end of the barrel. A fully populated disc contains 132 modules, with 52, 40 and 40 modules in the outer, middle and inner rings respectively. In order to make the most efficient use of silicon area, the modules are arranged such that a particle must travel through four layers of the detector (as with the barrel) and there is minimal coverage above $|\eta| = 2.5$. This means that some discs aren't fully populated, where this would introduce extra material for no gain in coverage.

This sub-detector is described in more detail in chapter 3.

2.3.3 TRT

The transition radiation tracker is the outermost portion of the inner detector. It is a gas tube tracker. Particles passing through a gas in an electric field generate currents in wires in the centre of the tube (or straw). The barrel straws are oriented axially and layered such that a particle is designed to pass through about 34 straws.

A special feature of the TRT is its ability to distinguish electrons due to their X-ray radiation on transitioning between two media of differing refractive index. This determines the structure of the plastic (polypropylene/polyethylene) surrounding the straws.

A straw is a hollow tube of diameter of 4mm with a 30μm W-Re wire through the centre between which is applied a high voltage (around 1.6kV). The tube is filled with a mixture of 3 gases: Xenon captures the X-rays of the transition radiation, CF_4 increases the gas speed and CO_2 is used to stabilise the mixture in the high electric field. The straws have a relatively small diameter compared to a standard gas detector. This is due to the very high hit occupancy that would result if the straws were a few centimetres across.

The straws are arranged with 50,000 axial straws in the barrel and 160,000 radial straws in each end-cap. A measurement of the drift-time across the tube enables space point measurements in ϕ with a resolution of 170μm per straw. In order to avoid excess occupancy where the coverage is already very good, the first few layers of the barrel aren't active in the central portion. All the barrel straws are divided at z=0 and read out at each end, the end-cap straws are read out at the outer end.

A large part of the read-out electronics for the TRT deals with the precise cancellation of the response from the Xenon gas. This takes microseconds to decay, compared to the several nanosecond timing precision and recovery needed. The generation of the artificial response is carried out using RC filters.

2.3.4 Track Reconstruction

Tracks are fitted to space-points which are gathered from the hit information in the Inner Detector, allowing reconstruction of the path the particle took. The track follows a decaying helical function defined by the momentum and pseudo-rapidity of the object particle and also the strength of the magnetic field used.

The particle momentum is proportional to the radius of curvature of the track: for a normally incident particle, the momentum of a particle is given by $p = 0.3BR$, where p is momentum in GeV, B is the magnetic field in T and R is the radius in m [9]. The error in the curvature $k = 1/R$ of the track due to the resolution of the measurement can be approximated by:

$$\delta k_{res} = \frac{\epsilon}{L'^2}\sqrt{\frac{720}{N+4}}$$

though this calculation is approximate for small N (< 10) and the accuracy is also affected by multiple scattering, particularly at low momenta. The track length in m is L', N is the number of measurements and ϵ is the measurement error in m.

The momentum resolution is therefore directly related to the accuracy and resolution of the space-point measurements and this argument motivates the very high resolution necessary in the inner detector. The precision of ID measurements allow measurements of particle momenta up to hundreds of GeV.

2.4 Calorimetry

Outside the tracking of the inner detector, the next layers are the calorimeters. These measure the energies of the incident particles. Two types of calorimeter are used, the electromagnetic calorimeter which wholly absorbs the lightest electromagnetically interacting particles, γ, e^-, e^+, and usually contains part of any hadronic showers. The hadronic calorimeter absorbs essentially all of the remaining strongly interacting particles, leaving only muons and neutrinos, of the known Standard Model particles, to escape the detector.

The principle of operation is similar in both sections of the detector. The incident particle interacts with a shower material or absorber, producing a shower of secondary particles, which are sampled in a detecting material. Two methods are used for detection: liquid argon in a high electric field produces cascades of electrons which are detected as charge and read out electrically; the tile calorimeters use a scintillating plastic which generates visible photons which are detected using photo-multiplier tubes. The size of the signal is approximately proportional to the incident particle energy, in both cases a calibration is made based on known particle energies.

Electrons, photons and hadrons interact with matter in different ways: the first two interact with electrons and the electric fields around the nuclei in the detector. The distance travelled before interaction is usually measured in radiation lengths. Hadrons interact via strong interactions with the nuclei in the detector. In this case, the distance travelled is usually measured in interaction lengths. The typical distance travelled in the detector before interaction for hadronic particles is longer than for electrons and photons, which enables the different shower types to be distinguished.

The ATLAS calorimetry is sub-divided into different systems, covering different pseudo-rapidity ranges and differing depending on the required resolution and particle type and also the prevailing radiation conditions. The subsystems are shown in figure 2.7 and described below.

2.4.1 Electromagnetic Calorimeter

The electromagnetic calorimeter is constructed from alternating layers of lead and liquid argon samplers, in an "accordion" geometry, such that gaps between detector modules are minimised.

A pre-sampler is placed in the barrel ($|\eta| < 1.8$), to allow corrections to be made for pre-showering within the inner detector volume. The remainder of the detector is built to sample showers created in the absorber layers.

The calorimeter is divided into cells which point towards the interaction region. This enables the measurements to be directly translated into η and ϕ.

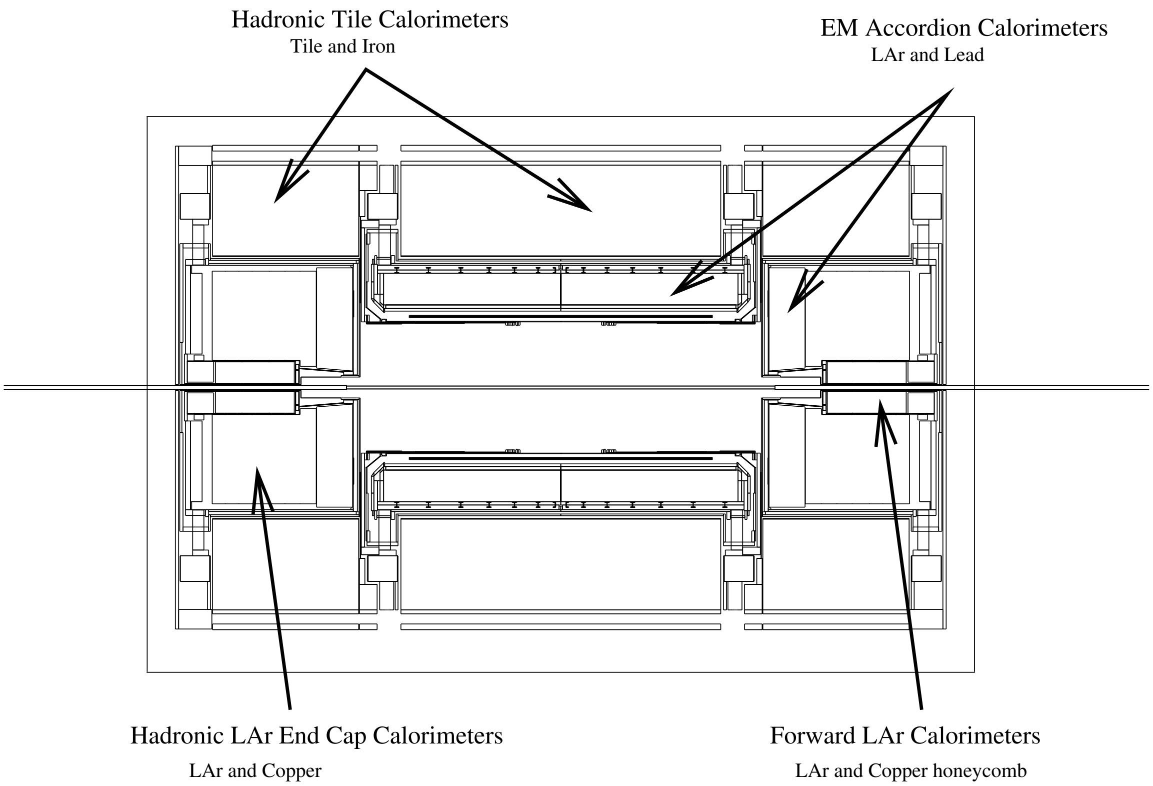

Figure 2.7: The calorimeters, adapted from [20]

2.4.2 Hadronic Calorimeter

The hadronic calorimeter is divided into portions that use LAr as the active medium, as in the EM calorimeter, and a scintillator tile based detector.

The barrel section ($|\eta| < 1.7$) is made up of an iron absorber, with plastic tile scintillator plates. This does not have a high radiation hardness and so a different technology is used in the end-cap.

The hadronic end-cap, covering $1.5 < |\eta| < 3.2$, works similarly to the electromagnetic system, detecting the shower using liquid argon as the active medium and uses copper as the passive shower material.

The forward calorimeter provides coverage in the extreme pseudo-rapidity region, $3.1 < |\eta| < 4.9$, where no other detectors are placed. This is a high radiation environment so there is a need for an intrinsically radiation hard system. Holes of diameter 5mm are formed in a block of copper or tungsten and rods of diameter 4.5mm are inserted. LAr is used as the active medium in the 250μm gap formed in-between.

Instead of pointing towards the interaction region, the hadronic calorimeter cells have a similar cross-section for all z in the barrel and in the end-cap. In order to provide data on jets in the barrel in η and ϕ, the cells are grouped as inputs to the PMTs. The end-cap cells already provide symmetry in ϕ and are formed into groups of similar η.

2.5 Muon Detector

The muon detector [21] measures the presence of, and finds the momentum of, the remaining detectable particles, muons. It is made up of several parts, the most visible being the external magnets which provide a toroidal field in order to bend the remaining particles to reveal momentum information.

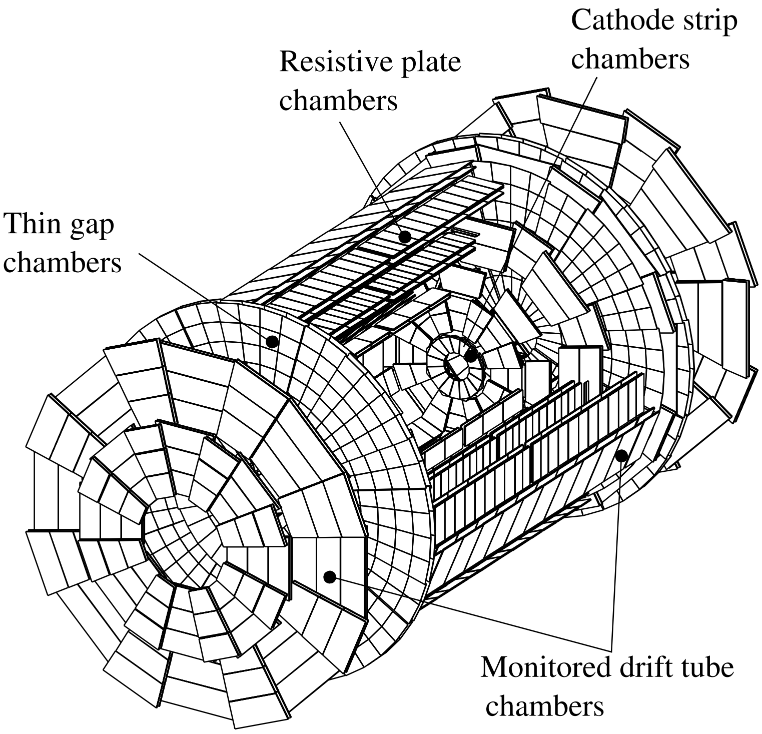

Figure 2.8: The muon detector, showing the coverage of the different sub-detectors (from [22])

Within the magnetic field the muon detectors are divided into different types, as shown in figure 2.8. The monitored drift tube chambers (MDTs) measure the precise co-ordinate, required for momentum measurements, in both the barrel and end-cap. They are complemented by the cathode strip chambers (CSCs) in the very high $|\eta|$ region. The muon detector also provides input to the level 1 trigger decision. This information is provided by the resistive plate chambers (RPCs) in the barrel and the thin gap chambers (TGCs) in the end-cap. They also provide positional information orthogonal to the momentum measurement. These detector types are described briefly below.

2.5.1 Toroids

Surrounding the whole experiment are the muon toroid magnets, which cause the tracks of the muon particles to bend in order to allow the determination of muon momenta. The magnets create a toroidal field running around the detector, therefore particles are bent such that the angle made with the beam-pipe changes. This is in contrast to the magnetic field of the inner detector and means that momentum resolution in the muon detectors is determined by the precision measurement of z in the barrel and r in the end-caps.

The toroidal field is created by 8 toroids in the barrel, which are air cooled superconducting magnets 26m long with an outer diameter of 20m. In the end-cap 8 toroids are rotated by 22.5 degrees so they fit in the end of the barrel toroids. They extend 5m from the calorimeter barrel and have inner and outer diameters of 1.65m and 10.7m respectively. These can be seen in figure 2.4.

2.5.2 Detector Chambers

The precision momentum measurements are supplied mainly by the MDT (monitored drift tube) detectors. These are drift tubes of 30mm diameter, with a central 50μm W-Re wire separated by a potential of 3kV. A timing resolution of 300ps over a maximum drift time of 500ns gives an average positional accuracy of 80μm per wire. The "monitored" in the name of the detector refers to optical imaging scans that measure the alignment of the detector a few times an hour. In order to provide the required accuracy, the position of the wires must be known to within 50μm. The drift tubes are supported by rigid cross-plates at either end, but the longest tubes are around 6m long and mechanical deformations of the order of 10μm are expected due to thermal effects, gravity and torsion.

The CSCs (cathode strip chambers) are located in the extreme end-cap ($2.0 < |\eta| < 2.7$), where the predicted hit occupancy would be too high for the MDT system. These are gas drift chambers with a 2.6kV bias. The cathodes are separated by 5.08mm, with anode wires separated by 2.54mm running through the centre. A precision measurement (r) is provided by measuring the charge recorded on neighbouring cathode strips, which are orthogonal to the wires and read-out at a pitch of 5mm. This system allows a track resolution of 60μm.

The barrel trigger functionality is provided by the RPCs (resistive plate chambers). These are formed from resistive plates surrounding gas chambers, with a high voltage (9kV) applied across the 2mm gap. Orthogonal strips with a pitch between 30 and 40mm on each side of the gap allow a capacitive readout. The strips are read out digitally giving a typical space-time resolution of 1cm$\times$1ns.

In the end-cap the trigger functionality is provided by the TGCs (thin gap chambers). A gas-filled chamber of width 2.8mm contains central anode wires separated by 1.8mm. A voltage of 3.2kV is applied across the gap. To match the required momentum resolution for the low precision trigger signals, groups of between 4 and 20 wires are connected together. Orthogonal read-out strips with a pitch between 15 and 50mm are capacitively coupled to the cathode and are read out to provide precision measurements of the ϕ component to complement the r component measured by the MDT.

2.6 The ATLAS Data Acquisition System

Data Acquisition, or DAQ, is the process of reading out the event data from the detector. In ATLAS it is strongly linked to the trigger system [23] described in section 2.2.3. Level 1 trigger information is gathered via a fast path for low resolution data from both the calorimeters and the muon trigger chambers. It is then processed rapidly to produce sufficient information for the decision.

The remainder of the detector must meanwhile buffer the data so it will still be available on reception of the Level 1 trigger accept decision, up to 2μs after the bunch crossing. In different sub-detectors this is implemented in different ways, but the data for the event corresponding to the Level 1 trigger bunch crossing must be available to the Level 2 trigger.

All the signals from the on-detector electronics chips leave the sub-detectors and are transmitted to off-detector electronics systems, mainly housed in two smaller caverns, US-15 and USA-15, either side of the main ATLAS cavern. Here the radiation levels are much lower and general purpose electronics can be used.

The TTC (Timing, Trigger and Command) system is used to distribute to the sub-detectors the L1 accept (L1A) trigger signal, the 40MHz bunch crossing clock and global resets of bunch crossing and event counters in order to keep the whole detector synchronized. This is done using a hierarchy of TTC crates which distribute the information to the sub-detector specific ROD (Read-out driver) crates in a standard format.

The first level of off-detector read-out is performed in the ROD crates. The RODs are sub-detector specific boards which convert signals directly from the sub-detectors into a common ATLAS event format. This data is then transmitted via a standard interface called the S-LINK [24]. The ROS (Read-Out Sub-system) receives the data from the sub-detectors, buffered by the ROB (Read-Out Buffer) and responds to requests for events from the Level 2 trigger systems and the event builder. The event builder is responsible for building complete events, from the portions provided by the sub-detectors, for processing by the event filter.

Once copied to persistent storage the events are available for off-line analysis. This takes place in a distributed manner, data being distributed to sites around the world in a hierarchical tier system. The provision of much of the infrastructure for this is the responsibility of the LCG (LHC computing grid). LCG also provides a database for “condition” data, recording information about the current behaviour of the system. This includes data on currents and voltages applied to the detectors, on-line settings of thresholds and current noise values.

2.7 Summary

The LHC is a new proton collider being built at CERN. ATLAS is one of the largest detectors that will examine the resulting collisions. A key part of the inner detector, the SemiConductor Tracker (SCT) is described in more detail in the next chapter.

Chapter 3

SCT Hardware

This chapter provides more details specific to the SCT after the summary in section 2.3.2.

The first section describes the infrastructure of the detector, together with the cooling and power systems. Next the detector modules are described, with the support electronics that are directly attached, within the detector cavern.

The final section describes, in some detail, the off-detector electronics used to convert the data coming from the detector to a format suitable for processing by the higher level analysis system.

3.1 SCT Infrastructure

The structure of the SCT was introduced in section 2.3.2. The 4088 SCT modules are attached to 4 barrels and 18 end-cap discs made of carbon fibre. The barrels are made up of unbroken sheets, which in themselves are quite flexible, but are supported by rigid flanges at either end. Connection of services for the barrels are made at either end. The end-cap structure is composed of the 9 discs, inserted into a slotted cylinder, such that the services leave the detector along the outside of the cylinder.

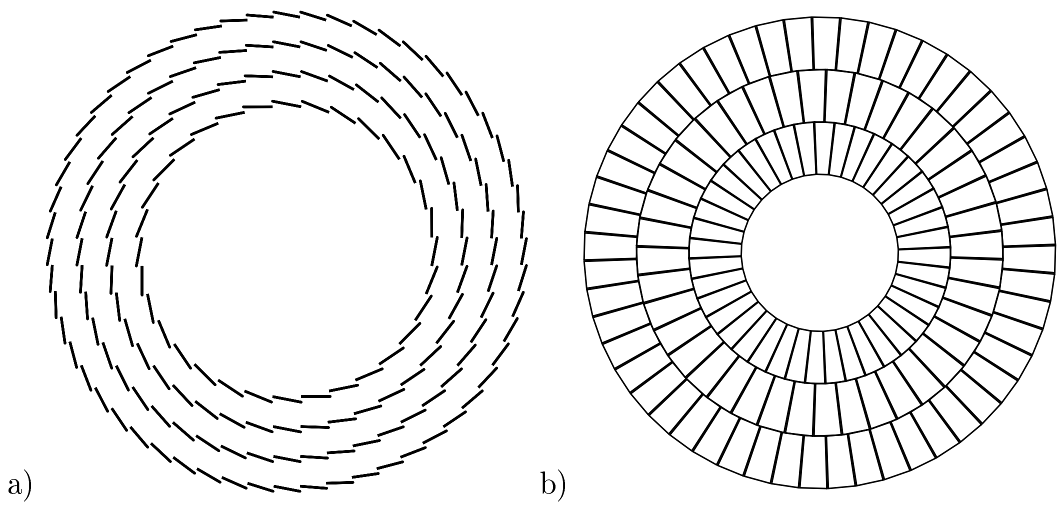

Figure 3.1: Module layout on a) barrel and b) endcap. Both are schematic views along the z axis of the experiment

Modules are attached to the barrels and discs so that they overlap slightly with their neighbours, ensuring there are no holes in the coverage. The positions are shown schematically in figure 3.1. In the barrel, this involves rotating each module around the axis parallel to the beam-line; and then, along each row of 12, alternate modules are mounted in a lower or upper position, allowing them to overlap. The endcap modules are placed in similar upper and lower positions around the disc; the middle ring of modules is then placed on the opposite side of the support to the inner and outer rings.

The services connected to each barrel and disc include: power supply cables, optical fibres for data and control, cooling pipes and cables for monitoring humidity and temperature. On the detector, the optical fibres and power supply cables are distributed to each module by a harness. The row of each barrel is serviced by two harnesses. Each disc of the endcap is divided into quadrants, each serviced by six harnesses.

The services are connected off the detector in the US-15 and USA-15 caverns. The data and control fibres are monitored by the DAQ (Data Acquisition system) and the remainder are controlled and monitored by the Detector Control System (DCS). The

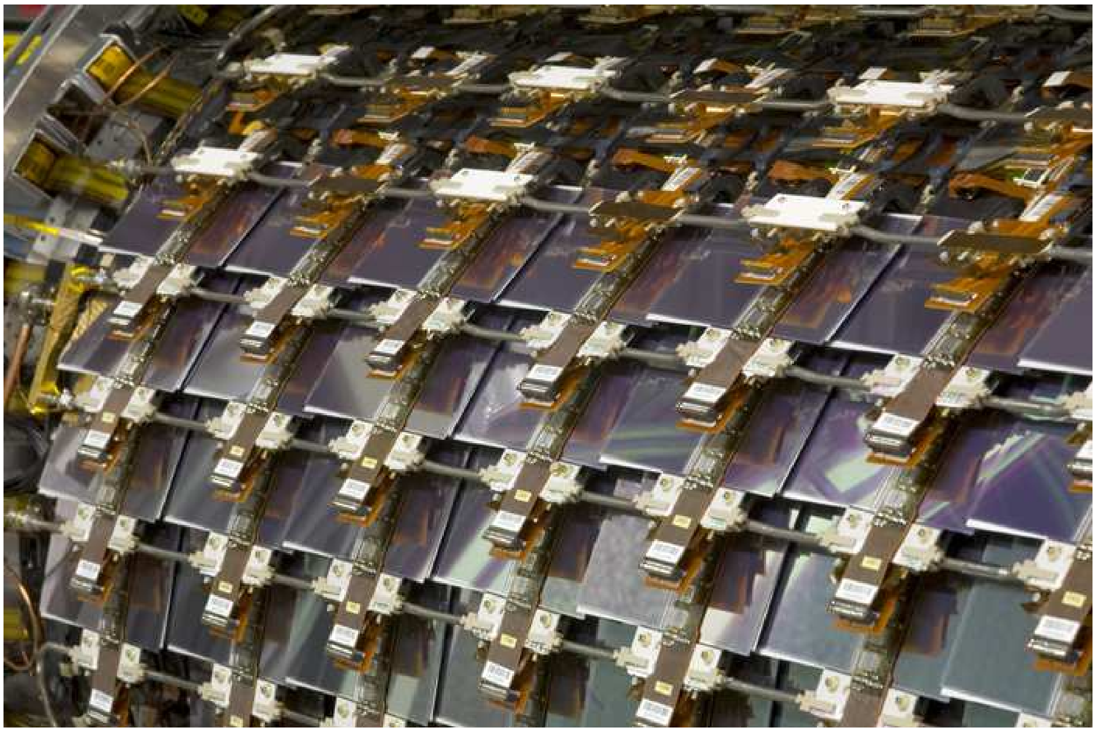

Figure 3.2: Photograph of modules on barrel 3. The upper two rows are empty, showing some of the services underneath

DCS is briefly described below and the DAQ system is described in section 3.4.

3.1.1 SCT Cooling System

Each SCT module produces about 8W of heat, about 6W from the read-out chips and the remainder due to the leakage current in the silicon detectors. This can't be dissipated passively from the small SCT volume so some active cooling is necessary. The modules are cooled using an evaporative cooling system based on C_3F_8, which is transported to and from the modules in tubes made from a very thin copper-nickel alloy.

The C_3F_8 is injected into the tubing at a low pressure using capillaries. The cold gas is then heated up by the detector as it passes through, absorbing heat from the modules. The gas is then compressed, cooled and returned to the system.

3.1.2 SCT Power Supplies

The power supply system consists of the low voltage and high voltage supplies. The low voltage supply powers the analogue and digital portions of the read-out chips. The high voltage supply supplies up to 500V for biasing the detector.

The power supply crates are housed in the US-15 and USA-15 caverns. Each crate contains 18 power supply cards [25]: 12 low voltage cards serving 4 modules each, and 6 high voltage cards with 8 channels each. Each crate therefore provides power for 48 modules. 88 crates (enough for 4224 modules) are used to power the whole of the SCT.

The low voltage cards [26] provide a variety of voltages to the modules. Due to losses in cables, the actual voltage which appears at the module end of the cable is measured and the source adjusted appropriately.

The module receives two main low voltages, 4V to the digital part of the read-out chip and 3.5V to the analogue part. The optical components receive their own power:

the transmitting lasers receive around 4V and the receiver diodes (PIN) get 6V. The low voltage cards also provide access to two thermistors mounted on the hybrid.

The high voltage cards [27] provide bias voltages up to 500V, and monitoring of currents over a wide range. An unirradiated module draws 3 orders of magnitude less current than an irradiated module due to an increase in leakage current. It may also need over 350V to bias the detector compared to the nominal 150V when unirradiated.

Each voltage to each module is supplied by an independent floating circuit. This enables control of grounding and shielding, all voltages are grounded to the SCT thermal shield.

3.2 SCT Module Specifications

The SCT goals require the module to be able to detect particles passing through with a high spatial accuracy and a temporal accuracy sufficient to identify which bunch crossing a hit is contained within, while contributing as few radiation lengths as possible to the inner detector.

3.2.1 Binary read-out

A major decision during the design period of the SCT was to use a binary read-out, as opposed to the more traditional analogue read-out. Instead of transmitting the analogue amplitude of the signal for a channel from the detector to be decoded by off-detector electronics, a single discriminator on the detector produces a single bit of data recording the hit/miss state of that channel.

The main advantages of using this system are: a simplification of the front-end electronics, as there are only two possible values that are allowed for each channel; the data can be easily compressed on the detector by removing the data that would otherwise need to be transmitted for the non-hit channels, and the off-detector read-out electronics can be simpler as they only have to deal with two values. Additionally, meta-data (such as event counters) can be transmitted in the same stream in a natural form.

The r.m.s. width of a strip is given by $d/\sqrt{12}$, where d is the strip pitch. A traditional analogue read-out system allows extra resolution to be interpolated from the signals from neighbouring strips, in cases where charge is shared between more than one strip. This was found not to produce significant improvements in the context of the SCT.

The binary threshold is set so that a channel can reliably distinguish between signal and noise. This means that the response to different signals must be known and also the noise must be low enough to be cut out.

3.2.2 Electrical Specifications

There are physics driven requirements for an efficiency of 99% per strip (low false negatives) and noise occupancy of 5×10^{-4} (low false positives) at the nominal threshold of 1fC [28]. The detector must also work after a radiation dose of 10Mrad and 2×10^{14} n/cm^2 of 1-MeV neutron-equivalent fluence. After irradiation, the detector noise increases and the charge collection efficiency reduces. The bias voltage needed to fully deplete the silicon increases from around 150V to about 400V after irradiation. Before irradiation the module should have a noise of ENC (Equivalent Noise Charge) 1500 electrons, allowing for an increase to over 2000 after irradiation. The signal-to-noise ratio is allowed to reduce from 15 to 10 over the ATLAS lifetime.

The loss of information from the binary read-out system implies that the threshold set on a chip must have a well known correspondence to the charge deposited in the detector. There is also a need for the threshold charge to be the same across the channels in a detector: if different channels responded differently to the deposition of the nominal threshold charge, the track finding algorithms would be biased by the potential extra hits. This leads to requirements that the channel-to-channel variation in threshold and noise be kept to a minimum.

Other requirements relate to the timing requirements of the module. A hit must be recorded in the correct time bin, as this is the only data that is recorded. The time-walk, the variation in response time to different amounts of charge, must be less than 16ns. Also the chip front-end must be able to discriminate two pulses separated by 50ns, allowing the edge detection circuitry to recover the the second pulse properly.

There is an allowance for bad channels during production. This is specified to be 1%, or 15 channels per module. As several bad channels in a row reduces the sensitivity to multiple hits, a limit of 7 consecutive bad channels is included.

The efficiency, signal to noise measurements and radiation requirements have been extensively tested in system tests [29], test-beams [30] and irradiations [31]. Other requirements have been found to vary more with module quality and are tested for each module during production. This motivates the need for calibration software to measure the variable parameters.

3.3 SCT Module Overview

The SCT is made of two distinct types of module, those located in the barrel region of the detector and those in the end-caps. Both types can be treated similarly by the data acquisition software, but apart from the read-out chips the barrel and end-cap use different components. The following description will focus mainly on the barrel.

Each module is made up of several distinct components: the silicon detectors, the read-out chips, the baseboard and the hybrid.

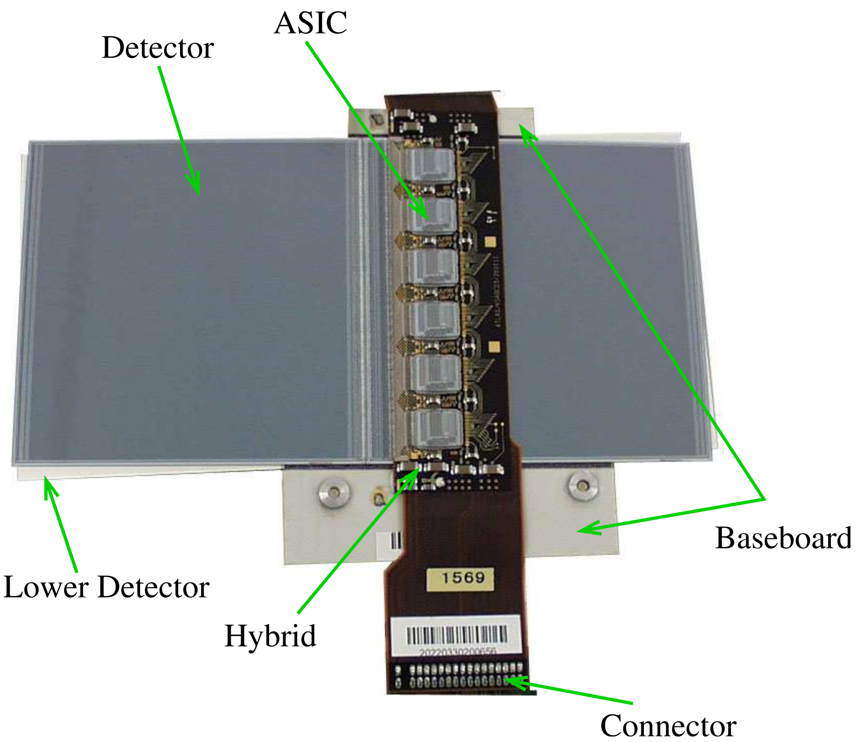

Figure 3.3: Picture of barrel module showing various components. Six ASICs are shown attached to the hybrid. The lower pair of detectors are visible below the upper pair due to the stereo angle

Barrel modules (figure 3.3) are rectangular and the hybrid is attached across the centre of the module. End-cap modules come in 3 different sizes, the hybrid is attached at one end and the detectors taper towards one end so that they tessellate on a disc.

3.3.1 Silicon Detectors

The active part of the SCT module is the silicon detector [32]. This reacts to ionizing particles that pass through, generating a charge that is discriminated by the chips on the hybrid.

The detector medium is a silicon crystal semi-conductor. The 285μm substrate is over-doped n^+, this is covered with a thick layer of lightly doped n-type silicon. Strips of p^+ silicon at the surface are covered with aluminium tracks with a pitch of 80μm, which conduct the charge to the electrical read-out. A voltage of around 150V is applied to the backplane, which fully depletes the n-type region and allows the collection of a minimum of about 2.2×10^4 electron-hole pairs, for a normally incident minimum-ionizing particle.

After running in ATLAS for 10 years, the radiation damage will result in "type inversion". This effectively converts the n-type silicon to p-type, which means that the depletion region is now formed from the p-n junction at the back of the crystal.

A barrel detector is 63.5 by 64mm with 768 active strips at a pitch of 80μm. A module consists of four identical detectors. The two pairs of detectors on each side of the module are skewed with respect to each other by 40mrad providing a stereo angle to give some resolution in the axial direction. Each detector pair is bonded together to create 768 120mm strips on each side.

3.3.2 Baseboard

The baseboard is used to hold all the other components in place, so is designed to be structurally sound while using a minimum of material. On a barrel module the cooling pipe is joined along a single strip on one side of the module. This means that the baseboard must also transport the heat generated (mainly by the read-out chips) from one end of a module to the other. A composite material called VHCPG (very high conductivity pyrolytic graphite) is used [33]. This has the useful feature of having a particularly high thermal conductivity in one plane, about $1700\mathrm{Wm^{-1}K^{-1}}$, compared to only 6 in the "transverse" direction (for comparison, silver is $400\mathrm{Wm^{-1}K^{-1}}$).

The baseboard is completed with electrically insulating beryllia plates to mechanically support the detectors.

The end-cap module has a similar support, known as the spine due to its position along the centre of the detectors. This also has the role of dissipating the heat from the detectors to the cooling pipe junction in the middle of the hybrid.

3.3.3 Hybrid

Around the baseboard and detector subassembly is wrapped a hybrid which supports the 12 readout chips. This provides data and power connections off the detector.

The barrel hybrid [34] is a multi layer PCB made of copper and Kapton*. This is shown in figure 3.4. The four layers are reduced to only two to provide additional flexibility at the connector end and also in the middle, where it is folded around the detectors. Two carbon-carbon bridges provide support over the detectors so that the hybrid does not come into contact with the sensors.

The pitch of the pads on the chips (48μm) is different to that of the strips on the

*Kapton ® is Dupont's trademark for polyimide

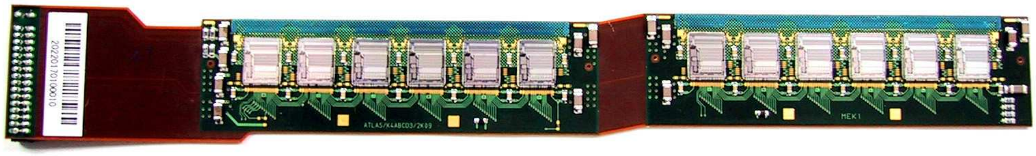

Figure 3.4: Barrel hybrid, showing the 12 read-out chips, the flexible centre part and the off-module connector at the left-hand end

detector (80μm). This discrepancy is resolved by using a pitch adaptor which is made from aluminium deposited onto a glass substrate.

Wire bonds are made from the pitch adaptor to the input pads of the front end chips and from the detector to the hybrid for each channel on the module. Additional bonds to the detector are made for the detector bias connections.

The end-cap hybrid is folded along its long edge. This is possible as the hybrid is attached to one end of the detector pair instead of to the centre.

3.3.4 SCT Read-out Chip Overview

Both the barrel and end-cap modules use the same front end chips and readout chain. This chip is designated the ABCD3TA [35] and is implemented in the DMILL process [36]. Each chip reads out 128 strips. The structure of the ABCD is shown in figure 3.5 and described in more detail below.

The whole chip runs synchronously with the 40MHz LHC bunch crossing clock. First, the charge generated in the detector is applied to the pad on the front-end of the chip.

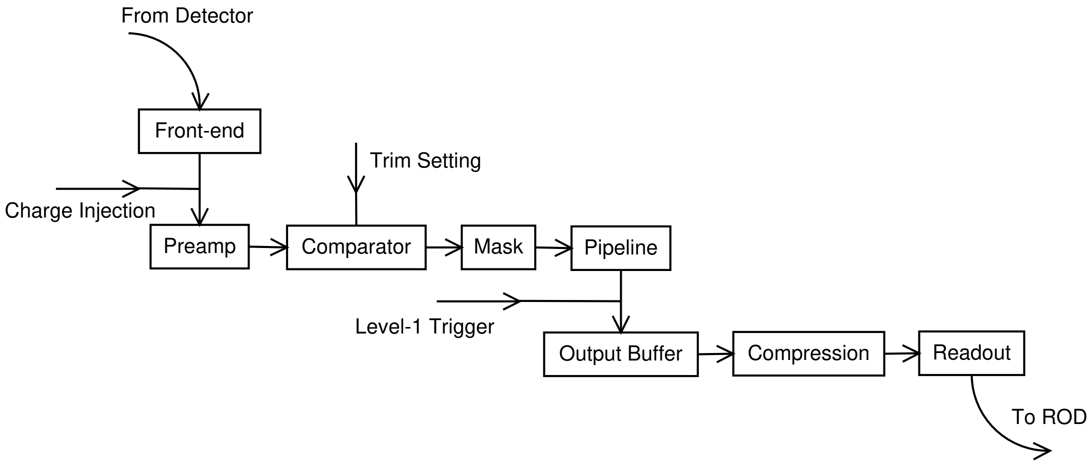

Figure 3.5: The components of the ABCD chip

The charge is amplified before a discriminator is used to generate the binary signal. On each clock the result passes through an optional edge detection circuit, which ensures that each hit on the detector generates only one pulse in the output. After passing through a mask register, the resulting value is stored in a pipeline storage buffer.

In addition to receiving signals from the detector, the chip also includes a calibration charge injection circuit which can inject a charge into the front-end of every 4th channel. All channels can be sent the charge in four stages by looping over 4 offsets. The first stage sends charge to channels 0, 4, 8..., the next stage will send a charge to 1, 5, 9... etc.

The pipeline is used to store the hit patterns in the chip while waiting for the L1A from the trigger system. It is built from 12 x 12bit shift registers per channel. Clock pulses are sent in sequence to 12 sets of cells; each set is a staggered chain between the shift registers such that each of the cells of a shift register receives a different clock pulse. The connections between them thus allow a bit to pass along one of the 12 shift registers in 132 clock cycles. Figure 3.6 shows a simplified 3x3 version where the input appears at one of the outputs 6 cycles later.

When a Level 1 trigger accept (L1A) signal arrives at the chip, the three bits from

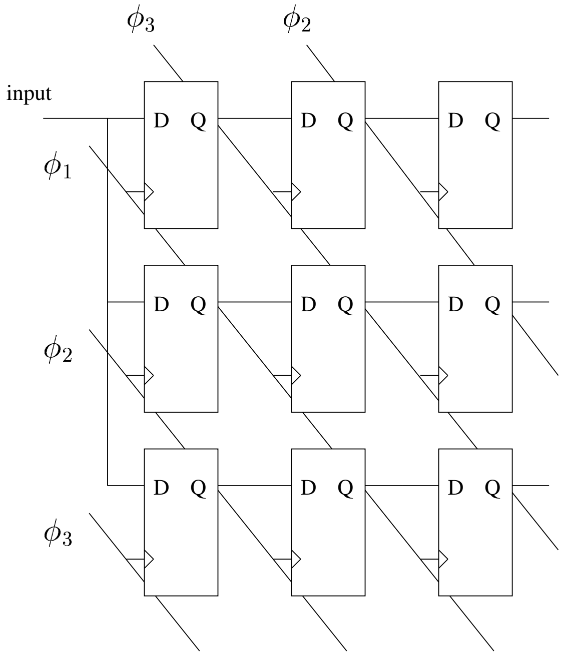

Figure 3.6: 3x3 illustration of the pipeline system. ϕ_1, ϕ_2 and ϕ_3 are clocks applied on subsequent bunch crossings to D-type buffers. A demultiplexer on the output is not shown

each channel currently at the end of the pipeline are moved into the readout buffer. These correspond to the hit/miss decision made at the bunch crossing (BC) selected by the L1A and also at the previous and subsequent BCs. This sequence of three bits is then filtered by one of four patterns: 'hit' which matches a hit in any of the three positions; 'level' which matches a hit in the central position; 'edge' which matches a central hit following a miss, and 'test' which will pass any pattern.

The data coming off the chip is compressed by reading out only those channels matching the specified pattern. The final bit-stream sent by the chip then includes the 4-bit chip number, followed by the number of the first channel of a string of consecutive hits, then the hits from all three time slots for all the consecutive channels. This is repeated for each set of consecutive hits. The bit-streams from individual chips are collected as described below and sent off the detector.

ABCD3TA read-out

The twelve chips on a module are arranged in two chains of six, one on each side of the module. Each chain is normally read out via one of two output streams off the module. The head of each chain is signified by a master chip (which is identified by a hardwired toggle) which passes data from the chips in the chain on to one of the output streams. The master starts the readout after the L1A has been received, first sending a header containing the BC counter and L1 counter for the event, then the hit data from the master chip itself. The master then passes a token to the next chip in the chain to tell it to send the bit-stream it has. This continues to the end of the chain where the trailer (a string of zeros that cannot otherwise appear in the output) is added.

In order to allow for the possibility of a dead chip in the read-out chain, the token passing may bypass a particular chip, so the token is passed to the next but one chip instead. This is carried out by specifying in the two working chips which of two pairs of token input and outputs to use to pass the token. The token can also be passed to the other half of the module, allowing chips on one side to be read out via the master chip on the other side. The end-cap hybrids allow an extra configuration (due to the geography of the hybrid), where the master on the other side may still take part in the read-out chain and all 12 chips can be read-out via a single stream. This can be used when the connection from the master chip to the off-detector read-out system has failed and is also used during module production to test hybrids with a simplified read-out system.

Configuration of the ABCD3TA

The ABCD provides several configuration parameters. These allow the chip to be configured for different behaviours, primarily during calibration so that the response of the detector can be precisely controlled. Some of the parameters affect the whole chip and a

Name	Values	Normal Value
Compression Mode	Hit, Level, Edge, Test	Edge
Calibration Index	0, 1, 2, 3	Any
Trim Range	0, 1, 2, 3	Low better
Accumulate Mode	On, Off	Off
Edge Detection	On, Off	On
Send Mask	On, Off	Off
Master Toggle	On, Off	On in master
End Toggle	On, Off	On at end
Input Bypass	Normal, Bypass	Normal
Output Bypass	Normal, Bypass	Normal
Clock Feed-through	On, Off	Off

Name	Level	Units	Minimum	Maximum	Step Size
Discriminator Threshold	Chip	mV	0	640	2.5
Calibration Amplitude	Chip	fC	0	16	0.0625
Strobe Delay	Chip	ns	0	$\sim$24	1 approx
Preamp Bias Current	Chip	μA	0	294.4	9.2
Shaper Current	Chip	μA	0	38.4	1.2
Trim DAC	Channel	mV	0	60-240	4-16
Mask	Channel	Boolean	-	-	-
Redundant Input Select	Module	Boolean	-	-	-

Table 3.1: Configuration values for a module. The upper table enumerates the values in the chip-level configuration register. The lower table describes other registers and toggles to configure a module

few are set at the channel level, a summary of all these is shown in table 3.1.

Five chip level configuration registers are used: one to set the discriminator threshold; two to set bias currents which modify the performance of the front-end amplifiers to compensate for radiation damage; and two registers to control the size and relative timing of the injection of the calibration charge.

Also at the chip level, a 16bit configuration register contains various flags which control the readout and bypass system:

- **Compression Mode** which compression pattern to use.

- **Calibration Index** which set of channels to send the calibration charge to.
- **Trim Range** specifies the size of the step used by the channel trim nibble.
- **Accumulate Mode** toggle the accumulation of hits at the end of the pipeline.
- **Edge Detection** a toggle for the edge detection circuits.
- **Send Mask** allows the channel mask to be sent as data through the pipeline.
- **Master Toggle** marks the chip as a master (this is ignored unless the hardware toggle is set).
- **End Toggle** marks the end of a read-out chain.
- **Input bypass** use the bypass connection for input tokens.
- **Output bypass** use the bypass connection for output tokens.
- **Clock Feed-through** the master chip sends the clock divided by 2 signal instead of data (set on startup).

For each channel there is a bit in the mask register and a four bit trim value. The mask enables the inputs to the pipeline to be switched off. The channel trim value is combined with the chip level trim range to enable the threshold used by the discriminator to be adjusted to compensate for the variation in the front-end amplifiers found between channels. The only module level parameter is the redundant input select which selects between possible inputs (see below).

The configuration is sent to the module over the same fibre as the trigger and reset signals and arrives at all chips at the same time. A chip identifies commands addressed to itself by comparison of an address encoded in the configuration stream with a hard-wired address encoding the position of the chip on the module. The redundant input select is controlled via the low voltage power supply and is also encoded in the chip address so that a single configuration stream may configure two modules.

Control of the ABCD3TA

In addition to commands to issue the level 1 accept (L1A) and to configure the chip registers, there are other signals which can be issued to the ABCD chip. Two reset signals are designed to be issued during physics runs to resynchronise the chip in the case that signals are lost. There are also two commands designed solely for testing purposes: one controls the injection of the calibration pulse into the front-end amplifiers, the second injects a single digital pulse into the pipeline. Both of these commands are used in conjunction with a delayed L1A signal to read out the affected pipeline bits once they have passed through to the end.

3.3.5 Module Optical Read-out and Control

The electrical LVDS (Low-Voltage Differential Signal [37]) signals generated by the two master ABCD chips on the module are converted to optical signals for transmission to the control room by two VCSELs [38] (Vertical Cavity Surface Emitting Laser) driven by the VDC (for VCSEL Driver Chip [39]).

The DORIC (Digital Optical Receiver IC [40]) decodes the encoded clock and command signals as received by the PIN diode and sends them to all twelve ABCD chips. A second pair of decoded clock and command signals is provided by the DORIC, which are sent to a second module. This provides a redundant command system to each module, as shown in figure 3.7, the pair of signals used by the module are chosen by the redundant input select.

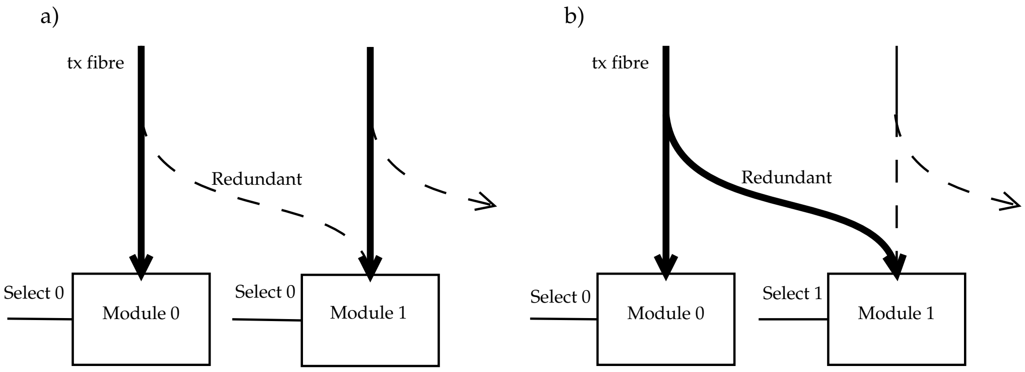

Figure 3.7: A pair of modules with redundant connections. a) shows two tx (transmit) fibres controlling two modules, b) shows the modules controlled by the same tx fibre, due to the change in the "select" signal

The VCSELs, PIN and converter chips constitute the opto-package. On the end-cap modules this package is plugged directly into the module. On the barrel modules, a short flexible circuit carries the LVDS signals a short distance to and from the opto-package.

3.4 Off-Detector SCT Read-out Hardware

The SCT ROD system is based on a VME [41] crate with a Single Board Computer (SBC). Up to 16 pairs of BOC (Back Of Crate) and ROD (Read Out Driver) cards are synchronized to the clock and trigger signals distributed by the TIming Module (TIM) in the centre slot, as in figure 3.8. The hardware design is shared between both the SCT and Pixel sub-detectors; firmware changes in the ROD and BOC allow for the differences in the event data.

The three optical fibres each from 48 detector modules are connected to the BOC card. This recovers the binary bit-stream from the two input fibres and simultaneously encodes the clock and command signal to be sent to the modules. The ROD is responsible for interpreting this bit-stream and generating the command signal to go with the clock from

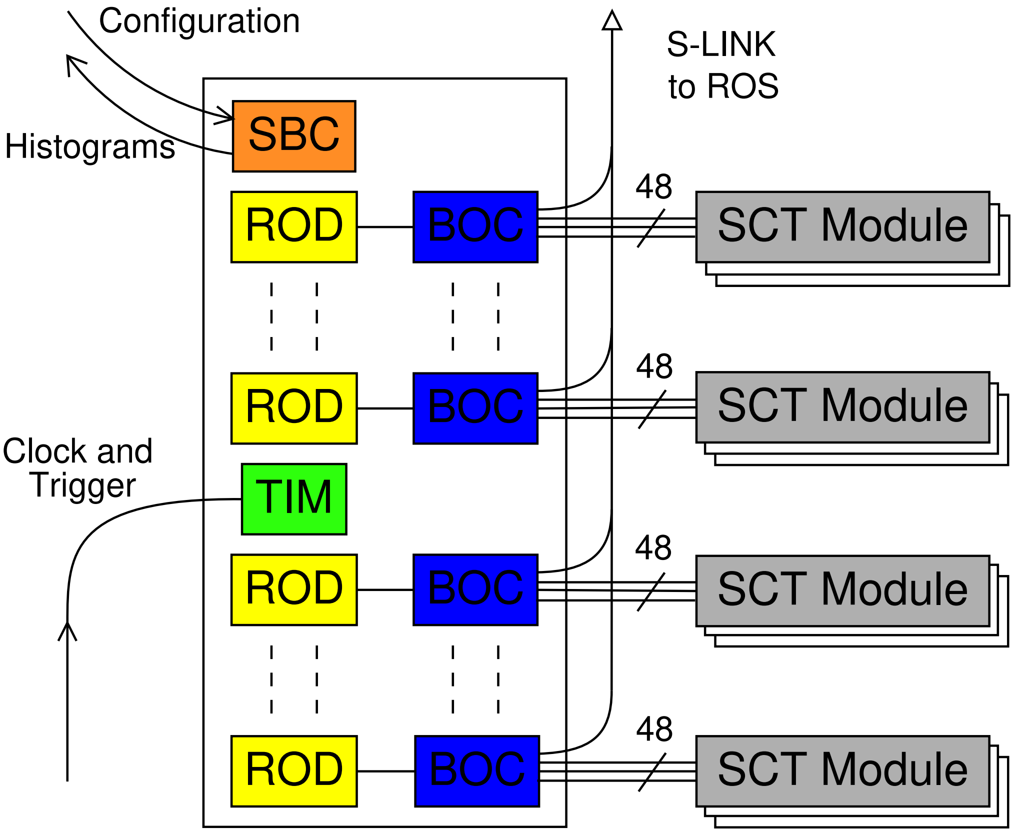

Figure 3.8: Schematic of the VME cards in a ROD crate, showing connections to the SCT modules and to the ROS which communicates with level 2 and the event builder

the TIM. The bit-stream is translated into the standard ATLAS event format which is transmitted via the BOC to the higher level ATLAS DAQ. The command signal encodes both the L1A signal and the configuration streams for the modules.

Each BOC/ROD pair is capable of controlling and reading out 48 SCT modules. In the final ATLAS configuration each one of 8 VME crates will contain on average 11 ROD/BOCs and read out half a partition (a partition is half the barrel or one complete end-cap). A crate of 14 RODs can read out the whole of barrel 6 (672 modules) which is the largest physical division of the sub-detector. This means that no more than one VME crate is needed during barrel or end-cap production. The SBC provides direct access to the ROD and TIM from the DAQ system, via ethernet; and is used to distribute configuration data and collect monitoring and diagnostic histograms. Each of these cards is now described in more detail.

3.4.1 TIM

The TIM [42] card is responsible for the timing within each ROD crate. It distributes signals from the TTC to the RODs and BOCs via a customized backplane in the VME crate. A common clock signal is distributed to the BOCs which then provide the clock signal to the connected RODs to synchronize transmissions between them. Also present on the backplane, and distributed to each ROD, are the L1 accept trigger and resets for the bunch crossing counter, the L1 counter and a full front-end reset. These enable the SCT system to synchronise itself with the rest of the ATLAS system and therefore send the correct trigger sequence numbers to the higher-level ATLAS read-out system.

There are two modes of operation: re-mastering and broadcasting the clock and trigger information received from the TTC, and a standalone mode where the TIM can generate any of the information internally. The standalone mode is used to test the functionality of a single ROD crate. In this mode, the TIM has the capability to generate single triggers, bursts of triggers or continuous streams. The rate at which the triggers are generated can be varied, and a pseudo-random jitter can also be applied so noise pickup from regular sources may be reduced.

3.4.2 BOC

The BOC is placed behind the ROD at the back of the VME crate (hence the name, Back Of Crate). This allows the distance between the two to be as short as possible, while enabling the independent development of functionality and construction.

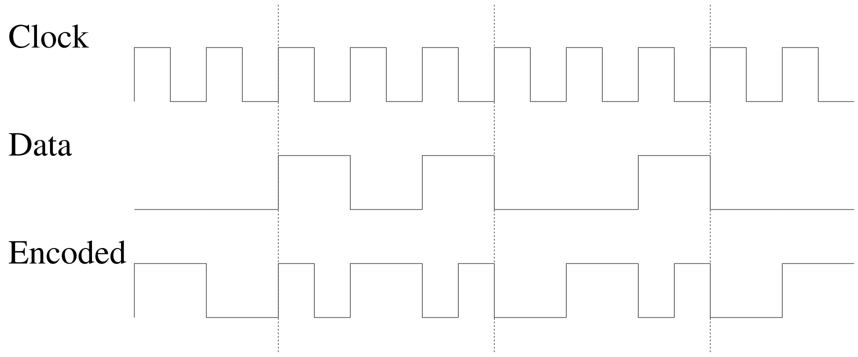

Figure 3.9: BPM encoding of the bunch crossing reset signal (1010010): for zero in the data a transition on the rising edge of the clock is present, for one in the data, an additional transition on the falling edge of the clock is introduced

The BOC can be divided into three independent systems: transmit (tx), receive (rx) and the S-LINK. For the S-LINK, data connections are provided to allow an ATLAS standard mezzanine card to be plugged in allowing transmission of the decoded event data from the ROD to the ATLAS DAQ. The tx and rx systems are described in more detail below.

Data Transmission

The data transmission part of the BOC provides the 48 modules each with one stream containing both trigger and configuration information, as well as the clock signal. The clock and command streams are encoded using the bi-phase mark (BPM) system (see figure 3.9), allowing only one fibre to carry both signals, and also guaranteeing that they arrive at the same time.

There are four programmable parameters controlling the transmission of data on each of 48 channels. These control: the laser current, in 256 steps; the BPM mark-space ratio, in 32 steps of approximately 0.25ns; the coarse delay of the BPM signal, in 32 steps of 25ns; and the fine delay of the BPM signal, in 128 steps of 0.3ns.

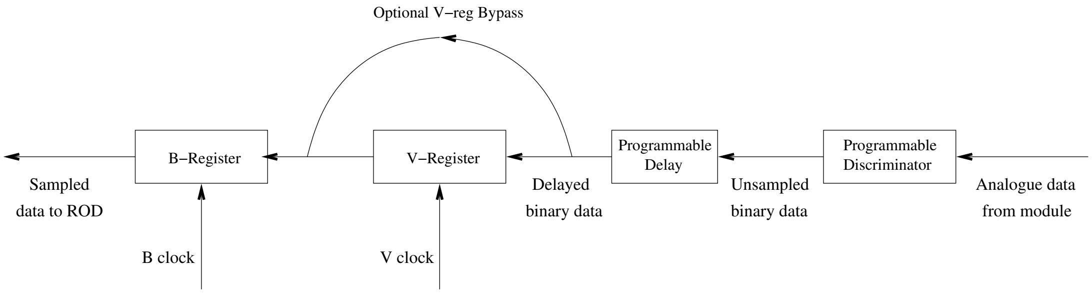

Figure 3.10: BOC buffers in the RX system

This ability to delay the signal by significant amounts allows the BOC to be configured such that each module receives the trigger signal at a precisely known time, regardless of the length of fibre between them.

Data Reception

The data receiver part of the BOC is used to receive the 2 bit-streams from each of the 48 modules, over 96 different fibres. The optical signal is sampled at the 40MHz clock rate to reproduce the bit-stream electrically. The flow of data through the receiver is shown in figure 3.10.

There are only two programmable rx parameters per channel. The first of these is the threshold which determines whether a 1 or 0 is produced when the data signal is fed into the discriminator, this is programmable in 256 steps. The second parameter is a delay which can be applied to the sample clock (0 - 24ns in 1ns steps). Though absolute timing is not critical at this stage as no information is contained in the timing, the delays allow corrections to phases on individual channels which are connected to different lengths of fibre and which may otherwise result in data corruption. Together these parameters allow compensation for deteriorations which may occur to the VDC components attached to the modules, which cannot be replaced once in the ATLAS pit.

There are other programmable parameters at the BOC level. The receiver can be configured to include an extra buffer which can be clocked at half the clock rate. This allows the clock-by-2 signal, from an unconfigured module, to be viewed either as a string of 0's or a string of 1's, allowing for measurement of the clock phase. This "v-clock" can be adjusted using two different "vernier" phase settings, which allow for comparison of timings that can then be adjusted in transmission. A coarse setting allow the phase to be varied over a range from 0 - 48ns in 1ns steps. The fine setting has a step of 40ps and a range of 0 - 10ns.

The result is then buffered in another register to synchronise the timing with the ROD clock, called the b-clock.

3.4.3 ROD

The ROD [43] is named identically to all the other cards which fulfil the same role in the other ATLAS sub-detectors. The design for the ROD hardware is in fact shared between the SCT and the Pixel detector, the only differences lie in the firmware which decodes the event bit-stream from the detector. The primary purpose of the ROD is to decode the ABCD bit-stream and transmit events via the S-LINK to the ROS (Read-Out Sub-system) and then the Level 2 trigger systems, and the event builder . It also enables access to the raw event stream for calibration and debugging of the system. For calibration purposes, it is able to generate internal triggers and histogram the events coming from the modules under its control.

ROD Firmware

The ROD is implemented using a series of FPGAs (Field Programmable Gate Arrays) and DSPs (Digital Signal Processors). The FPGAs implement the time critical part of

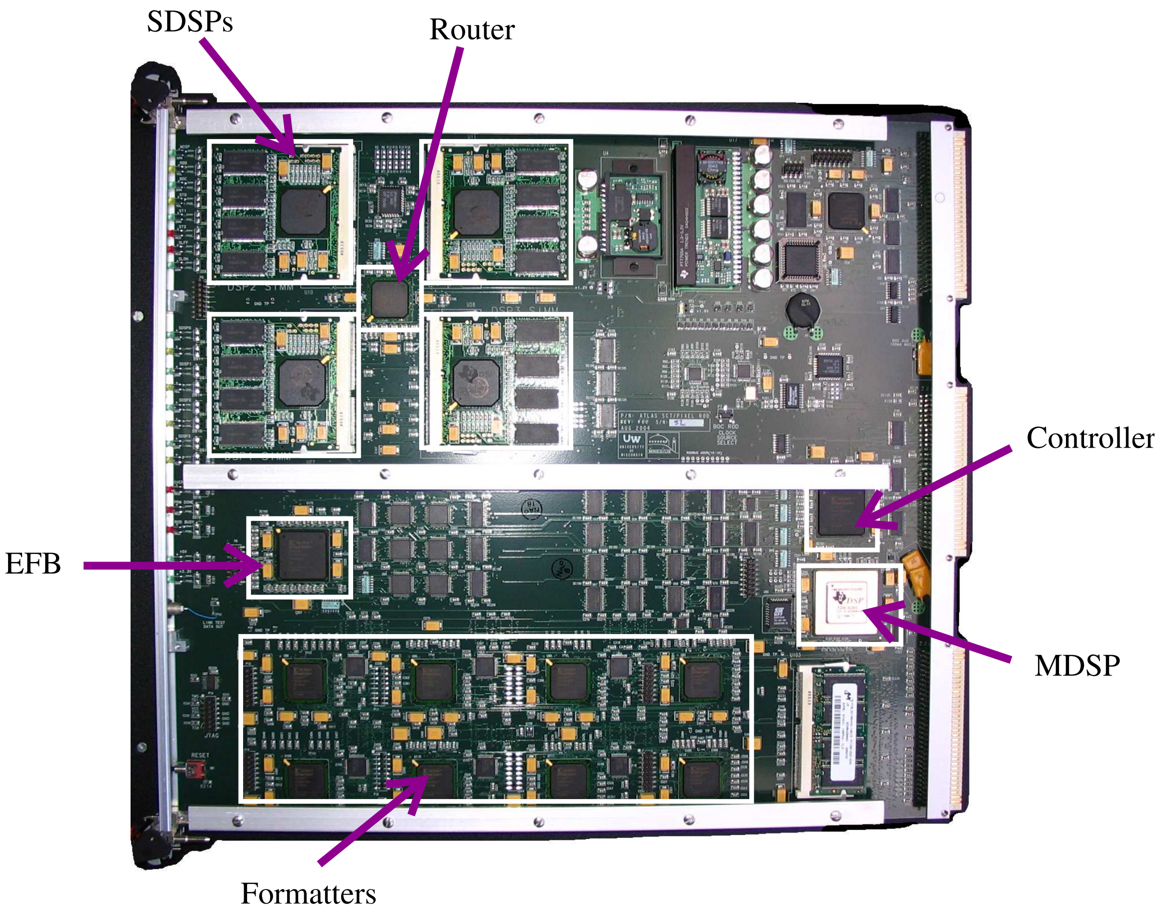

Figure 3.11: Picture of ROD module showing various components with FPGAs and DSPs labelled

the event data flow, while the DSPs are used for configuration, calibration and online-monitoring. Communication with the BOC takes place via a custom back-plane with (parallel) module data-links, S-LINK and a slower "setup bus" for changing the register settings on the BOC. The standard VME back-plane provides communication between the ROD and the SBC by exposing an interface to the DSP memories.

The FPGAs are loaded with one of four different programs, named the Formatter, Event Fragment Builder (EFB), Router and ROD Controller. These are shown in figure 3.11, also picked out are the four slave DSPs (SDSPs) and one master DSP (MDSP).

The ROD Controller is responsible for the high level control of the remainder of the system. This includes dealing with triggers from the TIM, keeping track of bunch crossing and event counters, and issuing arbitrary bit streams from two "serial ports" to arbitrary sets of modules. The serial ports are normally used for sending configuration data to the modules, but can also be used to send internally generated triggers and other command signals to the modules. Whatever the source of the trigger, when one occurs the controller signals the formatter FPGAs to look for an event header.

Eight formatter FPGAs each decode 12 "links", 6 modules worth of data. The primary purpose of the formatters is decoding the bit-stream from the ABCD chips into a stream of 16bit words for transmission to the EFB. At the front of the formatters are the input FIFOs, which enable arbitrary data to be played into the formatters, for testing, or the capture of data from the modules for independent processing. A second FIFO at the output of the formatter buffers the data ready for output.

When a token arrives from the ROD controller the formatter starts checking for the header pattern in the bit-stream. Once found, a header word is output and the remainder of the event is decoded. The data part of a normal event is decoded into a 16bit word which stores the chip and channel number and either one or a pair of hits. An optional expanded version encodes the data from all three time bins in the ABCD bit-stream. An event fragment for a link is completed with a link trailer word. Bit-stream errors, such as in the header, trailer or separator bits, and errors in the chip counter sequence are tagged for use by higher level processes.

The EFB (Event Fragment Builder) gathers the fragments from the formatters, building all 96 fragments into a single block to be sent to the router. The L1 and BC counters for each link are checked against the expected values and discrepancies are registered. These are then stored in the event header so that they can be discarded from the individual link headers. Links with no hits and no errors are removed to compress the data stream. The counter errors, together with those identified in the formatter, are counted and signalled in the event trailer.

The Router FPGA takes the event data and sends it over the S-LINK, and from the point of view of ATLAS data taking this is the end of the ROD's functionality. Additional features for calibration and monitoring are also implemented, though they are designed so as not to interfere with the main S-LINK data flow in data taking mode. A set of event traps in the Router monitor the event flow and can copy the data into slave DSP memory for further processing. An event trap can make a selection based on the trigger type of the current event (which may be specified by the ROD or passed by the TIM with the trigger), or on an internal count of the number of events that have passed through.

When it reaches a slave DSP the event data will start with a header which contains the expected L1 and BC counters, and the ATLAS trigger type. This is followed by the event fragments for each link, as decoded by the formatters. Each fragment is headed by a link header which specifies which link the data has come from, the channel and the formatter it was processed by.

The DSPs are of two types: one master DSP and four slave DSPs, which run at a slightly higher clock speed and have floating point math engines. The master DSP is responsible for non-timing critical operations, such as configuring the modules and controlling module resynchronization in case a reset or trigger is missed by a module. It also controls the generation of calibration histograms on the ROD, enabling triggers to be sent at a speed determined only by the ability for the events to be processed by the slave DSPs. The slaves are used to process events captured by the router event traps. This allows for generic on-line data monitoring and in particular for producing histograms of calibration data. The floating point engines allow for a fitting capability, so the histogram data can be fitted to reduce the bandwidth needed to transfer the resulting data off the ROD.

Primitives and Lists

The VME interface to the host SBC provides direct read and write access to the DSP memories via their HPI (Host Port Interface). This allows access to any single 32-bit word or block of words in the DSP memory. All the registers allowing access to the ROD FPGAs are mapped within the master DSP memory space. In order to allow a safer access and control method to the DSPs a system of primitive lists was developed which allows the DSP to remain in control of its memory while receiving commands from the host.

The structure of a primitive list is shown by Figure 3.12. An area of memory is reserved for the host (SBC) to load the primitive list into. A control bit is provided to let the DSP know that a primitive list has been loaded. The host may then monitor the status registers while the primitive list is executed. The master DSP monitors the primitive buffer in a loop, executing one primitive on each pass. The interpretation of the payload is chosen based on the primitive ID given in the header.

Primitives are provided for: reading and writing registers on the FPGAs; loading and copying portions of DSP memory; starting the slave DSPs; loading and modifying module configurations and also for starting more complex tasks, which run asynchronously with the primitive list loop. Primitives on the master DSP also allow for sending primitive lists to the slave DSPs and setting up the event traps in the Router.

An asynchronous message system allows the ROD to send messages back to the host. This works in a similar manner to the primitive lists. A bit in one of the status registers is set to let the host know a text buffer is present. The host then reads a structure from a

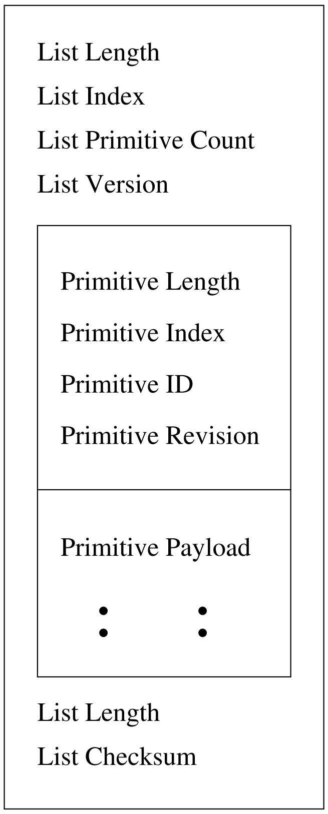

Figure 3.12: The ROD PrimList structure, each line represents a 32bit word

known address providing information about where the data is and how big it is. Separate text buffers are provided for error messages, warning messages and information messages. On the master a fourth buffer is filled with messages transferred from the slave DSPs.

Tasks

Tasks are processes that run in a cooperative multitasking mode and effectively change the character of a DSP. Tasks are started and stopped using primitives and are implemented as state machines which take turns on each iteration of the main loop. This allows primitive lists to run at the same time as tasks.

The Histogramming task runs on the slave DSPs and processes the event data which is trapped by the router into histograms. The Histogram Control Task, which runs on the master DSP during calibration, controls the sending of triggers to the modules and the changing of module configuration throughout the production of a histogram.

3.5 Summary

The ABCD chips on the detector modules produce event data which is sent to the readout system made up of crates of RODs, BOCs and TIMs. The ROD re-encodes the event data into a form used by the higher level ATLAS DAQ systems. It also provides access for monitoring purposes and enables the generation of histograms.

Software, running on the SBC, is required to make use of the ROD system to read out and produce histograms of event data from the ABCD chips. This software, and the use made of it, is described in the following chapter.

Chapter 4

A Data Acquisition and Characterisation System for ATLAS SCT Modules

This chapter describes the SCT module production process, the need for a characterisation system at all stages and what needs to be tested by the characterisation system. This continues with a description of the new SctRodDaq system designed for module characterisation with the ROD system, described in the previous chapter.

4.1 SCT Production Overview

The main effort in SCT production went into production of the detector modules. Each module was constructed from a few major components: twelve chips, a hybrid, a baseboard and four detectors. Each of these is described in section 3.3. The following describes the quality assurance tests carried out during the production of the SCT.

The first electrical tests are carried out on the chips before the silicon wafers are cut up.

Those with significant defects at this stage are discarded. This is done using a dedicated test system [44]. These tests include checking the effect of changes in clock speed and power supply voltages that are not repeated in further tests.

The individual components are then supplied to module construction clusters around the world. The barrel SCT modules are manufactured in four clusters: Japan, Scandinavia, UK and US. For end-cap production, there are 7 clusters: Freiburg, Geneva, Melbourne, Munich/Prague, NIKHEF, UK and Valencia. Each makes a number of modules consistent with the size of the cluster.

Each cluster breaks up the module assembly in different ways, but a standard testing procedure is followed. The procedure starts with two parallel steps: twelve chips are attached to each hybrid to create a "stuffed" hybrid, and four detectors are attached to each baseboard to create a baseboard sub-assembly.

The first step in the assembly of the hybrids is to attach the chips to the bare hybrids, and bond the hybrid-ASIC connections. At this point a confirmation test is carried out, which will find problems with bad channels or in the hybrid connections. This allows a problematic chip to be removed before the next stage. Wire bonds are then applied which provide connections from the chips to the pitch adaptor. A complete test is then carried out, including picking values for the channel trim DACs. The hybrid is run for 24 hours at a temperature of around 37°C. Triggers are sent continuously with a brief confirmation test performed at regular intervals. This provides a burn-in check and the status is monitored for any time dependent defects. Finally, the chips are tested cold, at a similar temperature to that in the final experiment and in a similar manner. This is designed to identify temperature dependent defects.

In parallel, a set of four detectors is mounted onto a baseboard. A leakage current check of the detectors is performed to check for breakdown problems. The alignment of the detectors on the baseboard is checked very carefully. The baseboard is the part of the module that is directly attached to the barrel structure. The position of the strips on the barrel is therefore directly related to the position of the detectors with respect to the mounting points on a baseboard. The position resolution of the detector can be reduced if the alignment is not well-known. Strict positional tolerances are imposed so that minimal corrections need be made to physics measurements. Therefore both the position of the detectors with respect to each other and also with respect to the mounting points on the baseboard is measured and controlled during production.

A hybrid may be transported to the module construction site where a further electrical confirmation test is run on it, to check for any deterioration in transit. The hybrid is then attached to a detector sub-assembly and wire bonds applied between the pitch adaptors and the detector silicon, to produce a finished module. A final electrical characterisation and metrology is carried out before a module has a 24 hour cold test. The metrology of the module also checks that the module fits within a well defined envelope. This ensures a module does not interfere physically with other modules, both while on the barrel and particularly during mounting where the module separation reaches a minimum.

All four SCT barrels are populated with modules at Oxford, a process called *macro-assembly*. Macro-assembly of the end-caps is split between NIKHEF and Liverpool. Each mounts the modules onto the discs in one end-cap and then mounts the discs into a cylinder. Before arriving at the macro-assembly sites, cooling services, electrical harnesses and module clips are fitted at RAL to both the barrel cylinders and the end-cap discs. Once the modules are mounted, a series of tests are carried out on the complete assemblies.

Following the individual macro-assembly of the SCT barrels, each is transported to CERN. The four barrels are fitted inside each other in a surface building near to the ATLAS pit. The barrel assembly is then tested together before being inserted into the

TRT barrel, then installed in the ATLAS pit. Due to time constraints, the end-caps have a shorter reception test before installation in the pit.

4.2 SCT Characterisation Overview

Testing of the modules during production is an integral part of the quality assurance procedure. The aim is to identify defects in the modules at as early a stage in production as possible so that effort is not expended on faulty modules and also to provide an opportunity to fix the defect. In the experiment, the ABCD chips need to be provided with various configuration parameters, so that the module operates as specified in section 3.2.2. These include module discriminator thresholds and the channel trims.

Once a module has been checked for basic power and readout functionality, the electrical performance can be tested. One of the first things to do in the characterisation procedure is make sure that an appropriate set of configuration values can be chosen to satisfy the specifications. This means that a channel must be tested to find an appropriate threshold along with a measurement of the noise.

The following describes how these calibrations are carried out and the tests that make up the characterisation procedure.

4.2.1 Calibration of Binary Readout

A key measurement to be made in the characterisation is the response of the analogue circuitry to specific charges arriving on the front-end. Charge generated within the detector is simulated by an injection of charge into the front-end circuits. The discriminator threshold that corresponds to this charge needs to be found using the readout available.

Using the binary readout, a scan could be made of the discriminator threshold: inject a known charge and read out the event corresponding to the charge injection at the end of the pipeline, then change the discriminator threshold and repeat. Looking for the change from an event recording a hit - where the threshold is low compared to the injected charge - to a miss - where the threshold is high - provides the appropriate threshold for the charge. Due to noise in the system however, mainly due to the detector leakage current, reading only one event doesn't give a reliable determination of the threshold.

Instead, the noise is explicitly allowed for in determining the threshold, as shown in figure 4.1. The noise is assumed to be Gaussian and a series of events is recorded for each threshold setting. This results in a measurement of occupancy at the different thresholds, with 100% at low threshold and 0% at high threshold. These measurements are accumulated into a histogram of occupancy vs (in this case) threshold, known as the occupancy histogram.

In between the high and low threshold regions, the occupancy curve is described by an error function, or S-curve, which can be fitted to the occupancy histogram for each channel, producing a mean value and standard deviation. These correspond to the discriminator threshold and the amount of noise observed on the channel.

This is used to link the threshold to the size of charge that was injected. For the experiment a threshold will be chosen that corresponds to an injected charge of 1fC.

4.2.2 Occupancy Histograms for Calibration

The occupancy histogram forms the standard unit for most of the tests in the characterisation sequence. The chip parameter which is varied across the histogram, along with the values scanned over, can be changed to analyse the performance in different ways. Also, the commands sent to the module with the trigger can be changed, enabling scans with charge injection and scans with a pure trigger and no charge injection.

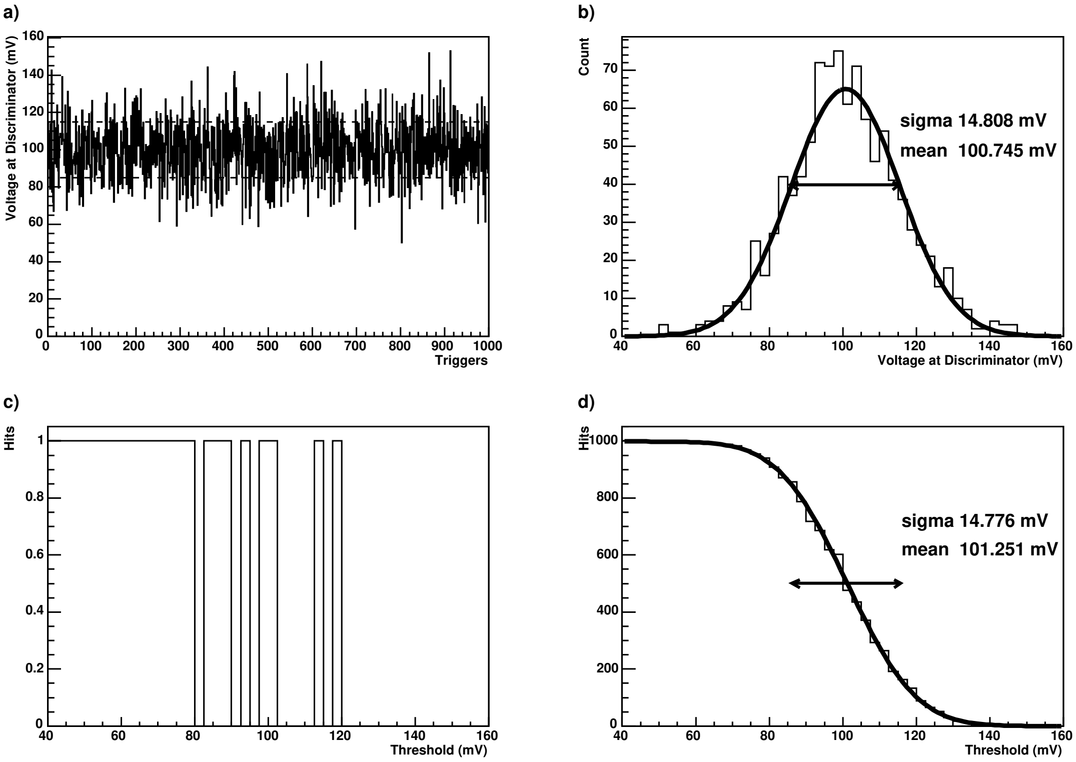

Figure 4.1: Simulation showing how detector noise is used for SCT module calibration. A Gaussian distribution centred at 100mV with a standard deviation of 15mV is taken as input to the discriminator. a) shows the raw voltage present at the discriminator over 1000 trigger events and b) shows the distribution of this voltage accumulated over 1000 events from which can be extracted the mean and sigma parameters. c) and d) use the information which can be gathered from the detector, showing scans over the discriminator threshold, a hit is recorded if the input voltage is above the threshold. c) records only one event per threshold setting and d) records 1000 events per threshold setting which allows the appropriate mean and sigma parameters to be extracted

For example, another measurement that is made during calibration is of the timing of the response of the electronics to the internally injected input charge. This involves a scan across the strobe delay parameter, which varies the time of the charge injection with respect to the read-out trigger. A relatively large charge is injected into the front end, while the discriminator threshold is set to a value which corresponds to a much lower charge. This results in a "top hat" function which represents the response to the calibration charge. With the strobe delay set to 0, the triggers will sample a point in the response after it has decayed. As the strobe delay is increased, the occupancy rises quickly to 100% showing the timing of the tail of the response. Further delay of the injection shows the rising edge of the response.

4.2.3 Characterisation Test Sequence

For testing of the module production, a series of tests called the Characterisation Sequence was devised. This sequence is designed to test as many features of the module system as possible and also to collect calibration parameters for the modules. Most tests consist of the generation of a series of occupancy histograms and a subsequent analysis to produce the results. The sequence is as follows:

- *IV Curve*: This tests the leakage current drawn by the module. This is measured by the high voltage bias system, as a function of applied bias voltage.
- *Hard Reset Test*: This tests the initialisation of the modules, which should provide a clock/2 signal after a hard reset.
- *Full Bypass Test*: Checks all of the bypass token links, in different combinations.
- *Redundancy Test*: This checks that a simple trigger burst is received correctly from the module for each setting of the select line. A simple check of the mask register

is also made.

- *Pipeline Test*: This tests the digital pipeline, channels can be identified as having stuck bits or dead bits in one of 12 positions.

- *Strobe Delay Test*: Sets the optimum Strobe Delay for each chip, based on the response to injection of a nominal charge.

- *Three Point Gain Test*: Provides a quick approximate measurement of the channel gains, by carrying out 3 threshold scans with different charge injected.

- *Trim Range Test*: Performs a threshold scan for each of the trim options and optimises the settings.

- *Response Curve*: Performs a 10 point gain scan with the channels trimmed. These are then used to generate a response function which maps injected charge to discriminator threshold.

- *Noise Occupancy Test*: One scan with no charge injection to find the noise value. This probes the tail of the noise distribution, which can show effects which are masked by the higher occupancy at low thresholds. It also provides a cross-check of the noise value obtained from the response curve measurement.

- *Timewalk Test*: Checks the variation of the optimum strobe delay value as different charge values are injected.

We can match the parts of the ABCD chip to the tests; referring to figure 4.2 and starting from the bottom. The IV curve checks that the detector leakage current is not too high. The response of the analogue system - detector, preamp, charge injection and comparator - is tested by the 3-point gain, response curve and timewalk tests. The noise

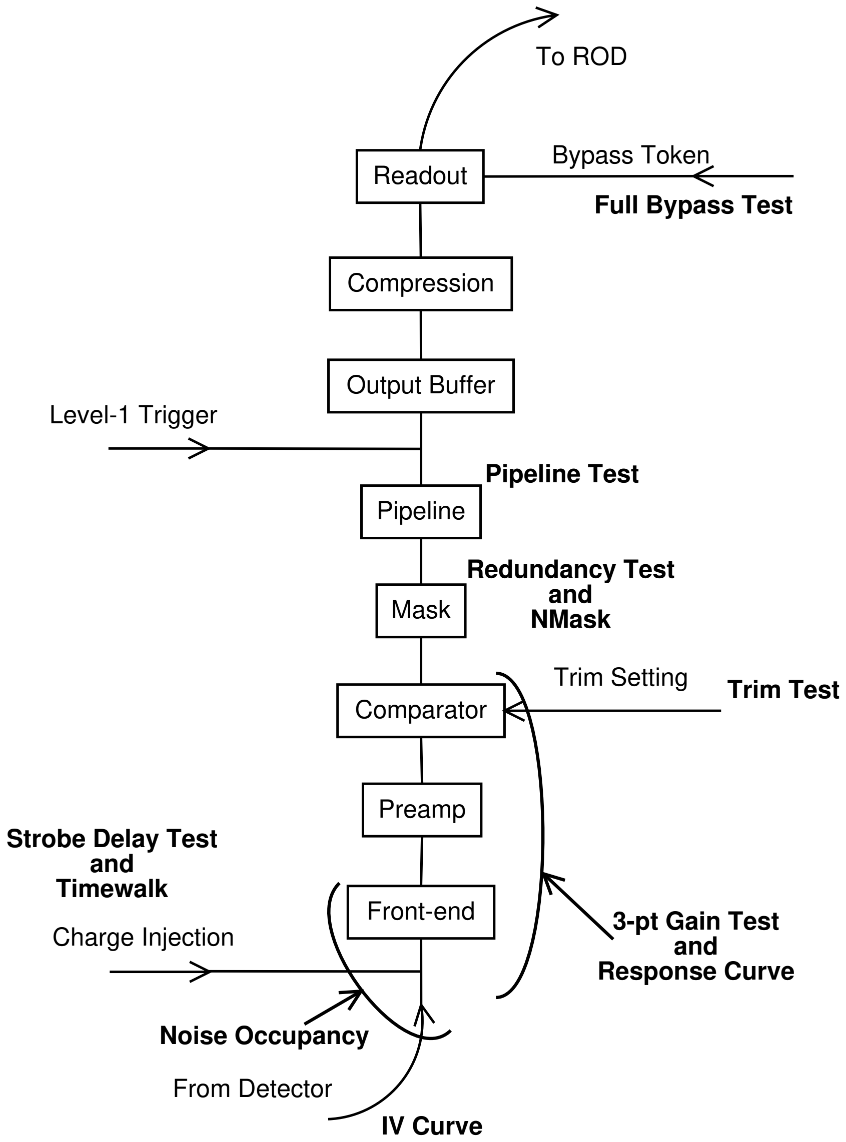

Figure 4.2: Overview of module data flow. Tests from the characterisation sequence are matched to the chip structure

occupancy test measures small effects in the tail of the noise, looking at both detector and electronics effects. The timing of the charge injection pulse is tested by the strobe delay test and also by the timewalk test. The trim settings are checked and adjusted by the trimming test. The mask register is tested by the redundancy test. The pipeline is examined in detail by the pipeline test. The trigger receiver, output buffer and readout systems are used by all of the tests. The chip bypass circuitry is tested by the full bypass test. At a module level, the redundancy test tests the use of the redundant clock and command signal.

4.3 SCTDAQ

During the production prior to macro-assembly, electrical testing was carried out using the SCTDAQ [45, 46] system. It was also used to take data in test-beams and during irradiation studies. It implements the characterisation sequence described above.

SCTDAQ controls the SCT modules using a hardware system based on the MuSTARD [47] because the full, final ROD readout system was still being developed at the time. This is based on a 6U VME crate and consists of the MuSTARD itself, the CLOAC [48] and SLOG [49]. The MuSTARD provides similar event decoding and histogramming functionality to the ROD, implemented using programmable logic. It implements only the receiver portion of the DAQ. The CLOAC provides timing information and can generate "fast" commands in a similar manner to the TIM. The SLOG or slow command generator forms the transmitter part of the system, which has been subsumed into the serial port functionality of the ROD. The SLOG is used to send long commands to the modules, including chip configuration to 12 modules, merging these with the single signal stream from the CLOAC. The output of the SLOG consists of 12 electrical signals which are sent to the modules.

In addition to the DAQ portion, SCTDAQ can also directly control the pre-production versions of the SCTLV and SCTHV, which can be placed in the same 6U VME crate as the DAQ system.

As regards the software implementation, each VME card has a library in SCTDAQ which represents the interface to the remainder of the software and uses VME drivers to talk to the card. On top of this is built a system for sending bursts of triggers, changing the module configuration appropriately and collecting the resulting histograms from the MuSTARD.

Each test in the characterisation sequence is controlled by a ROOT [50] script which is able to access the high level library. This results in an occupancy histogram called a *scan* corresponding to the scan variable and trigger selection requested. An appropriate function is fitted to each channel under control of a second ROOT script which then analyses the resulting data to produce the results in a form appropriate to each test. The results take the form of a set of summary plots and summary data which is uploaded to the production database so it may be reviewed at a later date.

This system works for small numbers of modules. For larger numbers, as in macro-assembly and in the final experiment, a large number of crates would be needed. This requires coordination between the computers attached to the VME bus and controlling the cards; this is particularly important during calibration scans. Another limitation of SCTDAQ is in the fitting of the histograms, which is one of the slowest parts of the analysis process and is carried out on the controlling computer. These limitations must be addressed in the final ROD DAQ system.

4.4 SctRodDaq

SctRodDaq is the name given to the DAQ system written to read out data from modules using the ROD system. This section describes the motivation behind the new system and gives an overview of the components. The remainder of the chapter describes the components in more detail.

4.4.1 SctRodDaq Specification

From the start of SCT macro-assembly onwards, the read-out of the SCT modules is carried out using the ROD system instead of that based on the MuSTARD. Due to the complete change in the VME cards the low level access must be rewritten. Much of the remaining code in SCTDAQ controls histogram creation and the building of the configuration streams to be sent to the modules. This functionality is now implemented in the ROD master DSP instead of in the controlling computer. Additionally, the fitting of histograms and the analysis of tests for only 12 modules has been found in SCTDAQ to consume a substantial amount of resources. In order to scale to larger numbers of modules, these tasks should be offloaded from the SBC. For use beyond macro-assembly, the system must be capable of coordinating between multiple ROD crates.

For these reasons, instead of retrofitting the previous SCTDAQ program, a new structure was developed, while making use of the knowledge gained from development of SCTDAQ.

In common with SCTDAQ, the new software must be able to:

- Configure modules according to an external specification
- Produce histograms of scans across module variables
- Fit different functions to the histograms

- Analyse data to produce new configuration parameters and defect lists
- Replicate the analysis carried out by SCTDAQ
- Provide access to low level functionality for expert users

Additionally, it must:

- Scale to at least 648 modules (for testing during barrel 6 macro-assembly)
- Make use of ATLAS provided software to facilitate future interaction with other ATLAS sub-detectors
- Communicate run state with the DCS system

4.4.2 SctRodDaq Overview

The SctRodDaq system was designed to implement these specifications. Fundamental to the design is a separation of the histogram production task from the fitting and analysis tasks. This allows the tasks to run concurrently where possible, to the point of running on different hosts. This means the SBC can spend most of its time on monitoring the ROD and moving histogram data. It also allows the different components to be implemented by different people.

An overview of the software elements of the system and the connections between them is given in figure 4.3. At start up, module configurations are read by the Configuration service from a database file and transferred to the ROD by the SctApi, which runs on the SBC. Calibration scans are carried out under the control of the Calibration Controller. First, histograms are requested from the SctApi, which passes the appropriate commands to the RODs. Histograms are returned to the system where they are fitted by the Fitting Service and analysed by the Analysis Service, as directed by the Calibration Controller.

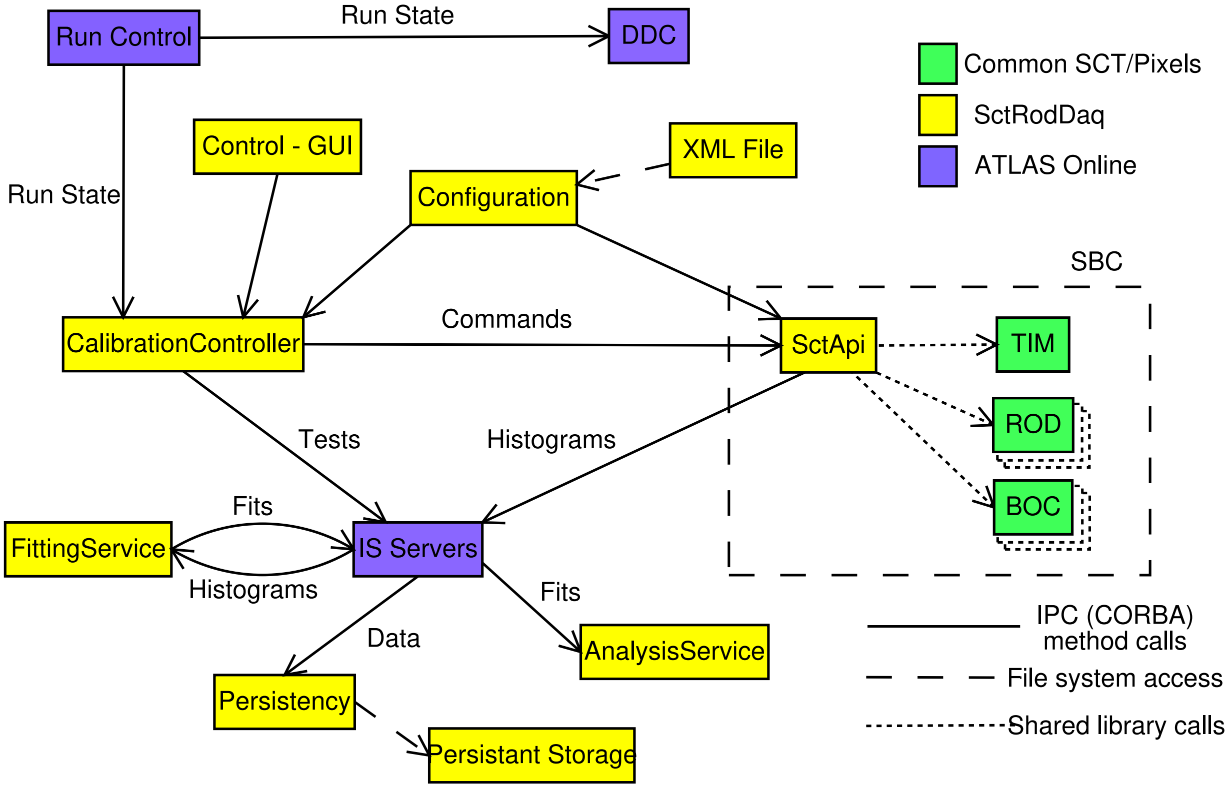

Figure 4.3: Overview of elements of SctRodDaq. For clarity, the connections between the GUI and various other components are not shown. The SctApi runs on the SBC in the VME crate. Using IPC (see section 4.5) allows the remainder of the system to run on different hosts

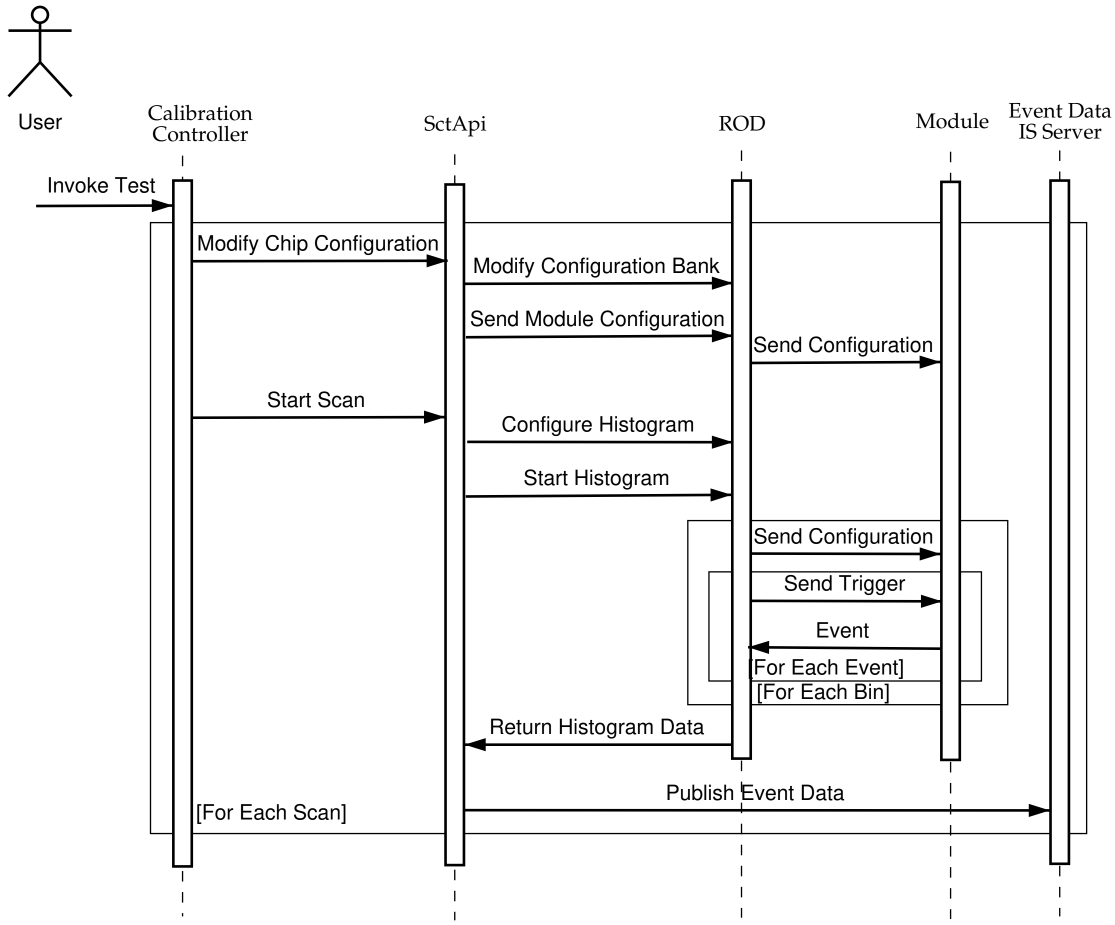

Figure 4.4: Sequence diagram showing the production of histogram data by the SctApi. The complete data is published to the Event Data IS Server. Further processing of this data is shown in the following figure

An SCT specific GUI is also provided within SctRodDaq. This allows the user to request characterisation sequences or individual tests and view the results of tests on the modules.

Figures 4.4 and 4.5 show interactions within the system in the process of carrying out a test.

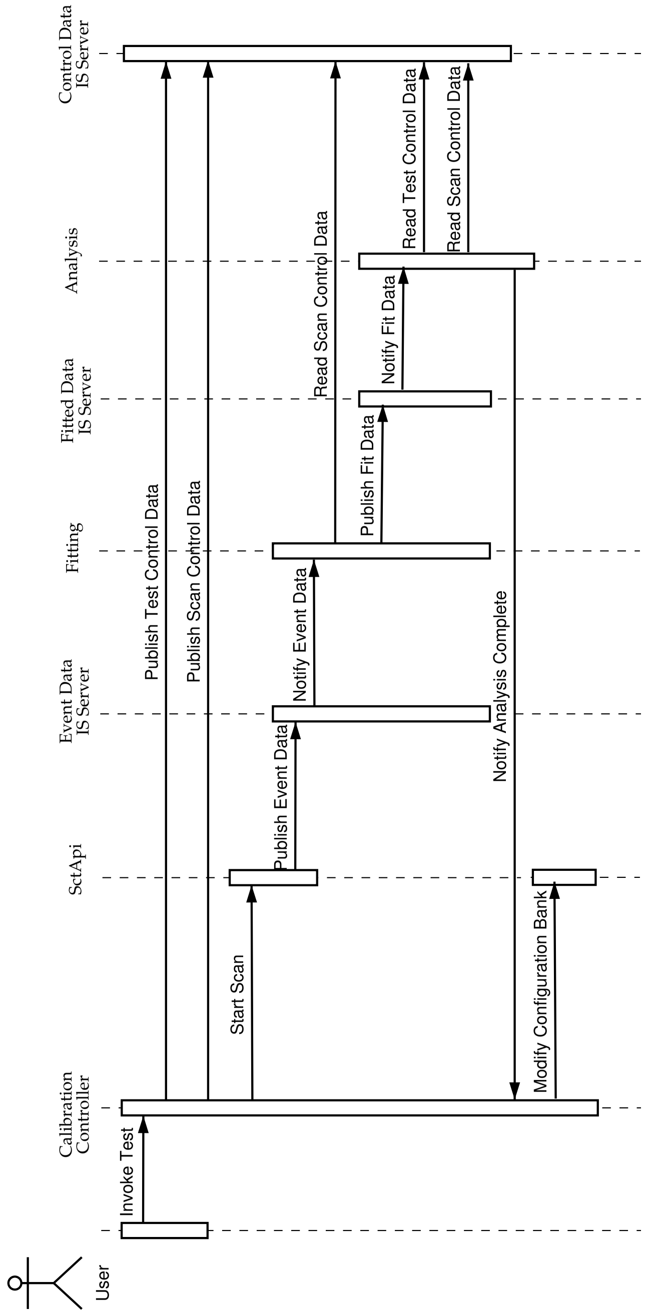

Figure 4.5: Sequence diagram showing processing of histogram data. The actions performed by SctApi between "Start Scan" and "Publish Event Data" are shown in the the previous figure

4.4.3 An Example: the Three Point Gain Test

The main use case for SctRodDaq is calibration, which is made up of a sequence of tests. These probe a particular part of the module and produce results which, for each module, consist of: raw data from which the remainder is derived; defects which apply to a range of channels in a module and are used to tag a module with a pass or failure tag; and a suggested change to the configuration of the module, used to set calibration parameters appropritately.

As an example, the Three Point Gain Test consists of three threshold scans with charges of 1.5, 2.0 and 2.5fC injected. It is used to measure the noise of a module and the similarity of response across the channels of a module.

The test is launched by selecting a menu item in the GUI. This invokes a test description class in the Calibration Controller. This first calls the SctApi on the SBC to set up the appropriate module configuration: compression mode to level; edge detect off; and send mask to normal. The size of the injection charge is set before each scan. For each scan a scan description class is invoked describing a Threshold Scan. This builds a description of the scan for the SctApi to execute. The variable over which to scan is set as the discriminator threshold, 500 events are sent for each bin and the range of threshold values is chosen according to the size of the injection charge. The sequence sent to the ABCD chip to trigger each event is set up: the charge injection command followed by a delay and an L1A command. This description object is then passed to the SctApi to generate the histogram. The Calibration Controller also places a description of each Test and Scan into the Control Data IS server, a distributed information store supplied by ATLAS online (section 4.5).

The SctApi starts by sending primitives to the ROD, checking that the configured modules are in the powered on state and returning valid events. Next, the scan is set up using more primitives: the router FPGA is set up to trap events into one of the slave DSPs; the slave DSPs are notified of the size and configuration of the histograms they are expected to create. Finally, the histogram control task is started on the master DSP, which produces the histogram. A thread is created in the SBC to monitor its progress, and the method call returns to the Calibration Controller.

The master DSP controls the sending of triggers to the modules and changes the configuration of the modules for each bin of the histogram. As the charge injection circuitry injects charge into a quarter of the channels at a time, the histogram code also includes a loop over the Calibration Index chip parameter.

The trigger sequence is sent to the modules. The calibration command triggers the injection of charge into the chip front-end for a single clock cycle. The level at the output of the amplifier is compared to the threshold and a 1 or 0 is clocked into the pipeline. The delay before the L1A command is sent is such that the value corresponding to the time of the charge injection is at the end of the pipeline when the L1A arrives. This triggers the hit pattern to be read out from the end of the pipeline and transmitted back to the ROD.

In the ROD, the event stream is processed by the Formatter, EFB and Router and is trapped into the slave DSP event buffer. This triggers the slave DSP to decode the event and accumulate the appropriate channels of the histogram. Channels where the charge was not injected are not incremented.

The SctApi monitoring thread polls the event counter and notices that the scan has completed. At this point, it reads the raw histogram data from the slave DSPs. Data corresponding to each module is separated and tagged with information about the scan, including the number of events which were received by the slaves and accumulated in the histogram. Each histogram is stored in the IS server corresponding to the Event Data, indexed by module serial number, run number and scan number.

The Fitting Service monitors the Event Data store. When a histogram appears, the scan type is examined and the complementary error function is chosen to fit against the data. The channel data is fitted producing mean and sigma parameters. Problem channels - for example where the fit fails - are tagged with defects and the generated data is placed in the Fitted Data IS server.

The Analysis Service monitors the Fitted Data IS server and choses the NPtGainTest analysis algorithm (the same algorithm is used for the 10-point Response Curve test). For each channel, a straight line is fitted to the mean and sigma parameters from the three scans in the test. The gradient of the slope represents the gain of the front end amplifier at this point. This is used to translate the noise recorded at the output of the amplifier to that seen on the input. The offset of the straight line fit is also recorded. Channels with too much noise are tagged as defective. A summary of this data is recorded which includes the mean values per chip of the gain, offset, output noise and input noise. Also recorded are the parameters for each chip of the straight line fit.

Finally, the Calibration Controller monitors the test results and when they are available, updates the configuration with the response curve parameters and masks channels that were recorded as defective.

The GUI monitors all of the information stores and displays a view of selected data, for instance the noise figures for all of the connected modules, in a colour coded diagram.

4.4.4 How it Fits Together

These subsystems are described in more detail in the following sections. Another description, focused on the analysis system and comparing it to the SCTDAQ system is given in [51].

4.5 Online Software

The ATLAS Online Software [52] is a collection of packages designed to form a common foundation for all ATLAS software running on-line. This provides an infrastructure which is used to create a distributed system where all the sub-detectors remain under central control while changes may still be made at a local level. Where possible these packages are employed by the SctRodDaq system. The main packages used are known as DDC, IGUI, IPC, IS, MRS and RC, and are explained below. In addition to explicit usage of the Online packages, the configuration and process management systems enable the different components to be run on different hosts.

The user interface is called IGUI and among other things enables the user to start up and shut-down the system using a series of run levels. This provides a uniform user interface for all sub-detectors, while allowing individual parts of the system to be temporarily removed.

The run level system used by the IGUI is implemented by the RC (Run Control) package. This allows a hierarchical system of Run Controllers to be specified, each of which corresponds to some functionality within the detector. Commands to move between run levels are distributed from the top to the bottom of the hierarchy. The Run Controller within SctRodDaq responds to the state changes and configures the modules and RODs appropriately. The run number is recorded by the Run Control system.

The fundamental method of communication between distributed components is implemented by the IPC (Inter-Process Communication) package. It is in turn based on CORBA [53]. IPC provides a standard name look-up system for the CORBA objects, based on a hierarchy of ipcservers. Within the SctRodDaq system it is used to provide access between different services. In addition to communication between components implemented on different hosts, IPC allows communication between implementation languages. This allows the majority of the system to be written in C++ while both the IGUI and the SctRodDaq GUI are written in Java.

The IS (Information System) package is built on top of IPC and provides an online repository for structured information. Repositories provide change notification and query services via their IPC methods. This is used in SctRodDaq to store and communicate information about the current calibration scan and test being performed, and also recorded and analysed data.

Messages can be sent to the user or any other component using the MRS (Message Reporting System) package. Messages have a simple structure, containing a text message, a severity, a set of named parameters and qualifiers. These can be used to filter which messages may be received. Various messages are used by almost all of the SctRodDaq components.

Finally, transfer of data between the DAQ and DCS (Detector Control System) systems is provided by the DDC (DAQ - DCS communication) package. This makes use of the IS package to store requests and data collected from the DCS system using the commercial PVSS (Prozessvisualisierungs- und Steuerungs-System) program [54].

4.6 SctRodDaq GUI

The GUI provides an overview of the module layout on a barrel or end-cap. It also allows scans and sequences of scans to be initiated, which are then analysed and the results displayed on top of the module layout using modifiable colour ranges.

The GUI is written in Java and most of the actions by the system are carried out using IPC calls to the Calibration Controller and the Configuration Service. Monitoring of results is done through the Online IS repositories, which are filled with data by the other components.

4.7 Calibration Controller

The Calibration Controller manages the actions of the SctApi and the Fitting and Analysis services. This is achieved by a system of sequences, tests and scans. At the lowest level, a scan request is sent to the SctApi which returns a histogram for analysis. Prior to this the module configuration is set up appropriately. A series of scans forms a test. All the scans in the same test are analysed together and result in a suggested change to the module configuration. A sequence is a series of tests, each of which performs updates of different module parameters, which results in an appropriately configured module.

Meanwhile, information about the requested scan is put into an IS server for use by the Fitting and Analysis services. The Calibration Controller watches the progress of the SctApi, Fitting and Analysis and prevents multiple sequences from running simultaneously.

4.8 SctApi

SctApi is the name given to the part of the system which interacts directly with the RODs. The remainder of the SctRodDaq system sends requests via an IPC interface. SctApi then translates these high level requests into commands to be sent to the ROD.

The main part of the SctApi is a process running on the SBC in the VME crate. This talks directly to the cards via the VME interface on the SBC. Commands sent to the SctApi are indexed with a crate number which enables an interface at a higher level to distribute these commands to the appropriate crate. This functionality is not required before more than one barrel is tested together at CERN and the following describes the SctApi implemented at the SBC level.

4.8.1 External APIs

The external interface to the SctApi is provided in two different ways, depending on the level of connectivity with the remainder of the system that is required. It is published either for use by the distributed system using the IPC package, or for local use through ROOT as the TApi class. Both interfaces provide identical functionality. The IPC interface is used by the remainder of the SctRodDaq system to perform any data acquisition tasks on the SCT modules. The alternative TApi class allows the user direct access to the hardware in a similar manner to SCTDAQ. This is of use mostly during development and debugging of the system and is particularly useful in allowing direct output to the same console as commands are issued from.

These interfaces provide access to the SctApi C++ class which implements the required functionality using a combination of the low level Crate class and the Configuration service. The methods provided can be divided into several types, as described below. The lowest level provides direct access to the hardware. Access to the ROD primitive list execution is provided, though is used mainly by the higher level SctApi methods. These latter methods allow initialisation, access to the module configurations, the production of occupancy histograms, and also access to raw data from the modules.

4.8.2 Low-level Hardware Access

The low level communication with the ROD is carried out via the RodModule class, which accesses the VME interface on the SBC. RodModule is a Pixel/SCT independent C++ interface to the ROD, which implements low-level DSP memory access methods and methods for writing the flash memories, for upgrading firmware. It also provides access to the ROD primitive list mechanism. This consists of: classes representing a primitive and a primitive list, which control generation of the primitive list buffer; and a simple state machine system for executing the primitive lists on the ROD. The state machine copies the buffer to the appropriate address in ROD memory and sets and monitors the appropriate registers to execute a primitive list. Both the BOC and TIM provide similar C++ interfaces, called BocCard and TimModule. BocCard provides methods to initialise the BOC and for protected access to the BOC registers, via methods in RodModule. TimModule provides access to the TIM registers and trigger sequencer.

Within the SctApi system, the instances of TimModule, RodModule and BocCard associated with the hardware in the crate are managed by the Crate class. The hardware configuration is read from the Configuration service and the Crate class initialises the hardware accordingly. The remaining methods are wrappers around the low-level access for use by the SctApi, providing proxies to one of multiple RodModule instances. The most important part of the Crate class is in the wrapping of the primitive list mechanism, which is used extensively at higher levels, and the text buffer readout, which is logged both to file and to MRS.

4.8.3 Primitive List Management

Primitive lists are the preferred means of accessing functionality on the ROD, as described in section 3.4.3. They are formed from command structures which are described by code common to both the DSP and the host. A PrimBuilder class allows primitive lists to be constructed by calling a series of methods representing the different primitive types. This creates a primitive list buffer which is then sent to a ROD for execution. Most primitives are executed in a short time but the functionality is provided to execute a primitive list asynchronously, where the list is sent and at a later time it is checked that it completed successfully. Another division is between primitives sent to only one ROD and those sent to all RODs in the crate together.

The primitive list thread polls the RODs for text buffers and activity in primitive list execution. Primitive lists are passed from higher levels to be sent to the ROD to be executed. A single thread, as opposed to a thread per ROD, is used so that the interleaving of the polling can be easily controlled. This also allows a separate queue for primitive lists sent to all RODs. A check is done on all the RODs to ensure that they are ready for a new primitive list, only then is the primitive list started on each ROD.

The primitive list mechanism is accessed by the higher level SctApi class through the `sendPrimList`, `awaitResponse` and `getResponse` methods. The `sendPrimList` method submits a primitive list for execution. The `awaitResponse` method waits for completion of the current primitive list. Finally, `getResponse` returns the result returned by the ROD of the primitive execution, if any.

4.8.4 Initialisation

The highest level method in SctApi is that which is called to initialise all the hardware. The `initialiseAll` method is called in response to a run state change in the IGUI. This initialises the connection to the VME backplane and then reads the hardware configuration from the configuration database. The selected RODs are reset and the firmware images for the slave DSPs are loaded and started. The associated BOCs are then loaded with appropriate parameters, followed by the TIM. Finally the module configurations are loaded from the configuration database and sent to the RODs.

In preparation for the beginning of a run, during which the characterisation tests may be performed, the run number is sent to the SctApi and the scan number reset. The module configurations are then sent from the RODs to the modules so they are ready to transmit events in response to triggers.

4.8.5 Module Configuration Methods

Module configurations are loaded from the Configuration service when the system starts up. They must then be sent to the modules and modified in preparation for calibration scans. The ROD takes care of storing the configurations and building the appropriate bit-stream to send to the modules. Configurations are stored in one of three banks which allows different settings to be stored for different circumstances, as described below.

The main functionality of the SctApi is transferring the initial configuration to the ROD, modifying the configuration appropriately and commanding the sending of the appropriate configuration to the modules. The methods `setABCDModule`, `sendABCDModule` and `getABCDModule` are used to move sets of module configurations around. These correspond directly to primitives that are sent to the ROD.

The module configurations can also be modified, both within the SctApi and in the ROD. The trim and mask registers within a chip can be modified in the SctApi then transferred to the ROD. The remainder of the configuration variables on the module can be modified using an identification number which is used to identify a scan type. These methods are used to set up the modules before a scan starts, to change the configuration on a module according to the analysis after the completion of a test, and finally to change the module configuration during a scan. In addition to parameters on the chip, parameters that can be set on the BOC are also identified for scans.

During a characterisation sequence the module configuration is modified many times. In some cases this modification should be stored but in others this is a temporary configuration, as is done to set up a threshold scan with a particular injection charge. In order to allow these temporary configurations to revert to the standard one the ROD stores banks of module configurations. The three banks are named *Scan*, *Calibration* and *Physics*. During a scan the *Scan* bank is overwritten with the changing configuration.

The *Calibration* bank is used by the CalibrationController to store the configuration as it is changed during a calibration sequence. The *Physics* bank is used for physics running and will be updated from the *Calibration* bank when it is known to contain a good configuration.

4.8.6 Scan Production

The main function of the SctApi during calibration is to produce occupancy histograms for analysis by the higher level SctRodDaq software. In SctRodDaq terminology, these are known as scans. The main parameters to describe a scan are:

- the configuration variable which is to change over the course of the scan
- the values the configuration variable should take
- the sequence of commands, known collectively as a trigger, sent to the modules for each event
- the number of triggers to send
- which modules will send events to the slave DSPs
- which slave DSP each module will send events to

A scan can also have some options which change the behaviour of the scan process, for instance whether the triggers are generated by the TIM or the ROD and whether the calibration line is in use, which requires a loop over the four calibration lines to generate a complete histogram. Once complete, a scan is identified to the SctRodDaq software by the run number, a scan number within the run and the serial number of the module.

One of the key reasons for generating histograms at the ROD level, is that many calibration histograms have an occupancy many times higher than that observed during a physics run. As the length of the bit-stream generated by the ABCD chip is in general proportional to the hit occupancy, then so is the time taken to read out the chip. This means that the module cannot be read out at full speed and the trigger speed must either be slow enough to allow the read out of the maximum bit-stream length, or it must vary according to the actual occupancy recorded by the module during the scan. The master DSP is placed in the appropriate place to control the trigger rate, where it can monitor both the incoming data in the formatters and the progress of the slave DSPs in accumulating the histograms.

Each module is assigned to one of eight groups. This allows the SctApi and the ROD to handle the modules differently. During a normal calibration scan, the master DSP sends a trigger to one of four groups (the low two bits of the group number) of modules in turn. The events from each group of modules is then trapped at the Router and distributed to a different slave DSP. This allows the time spent accumulating the histogram on one DSP to run in parallel to the time spent collecting and processing the events from the modules that are sent to the next DSP.

A second trigger stream can also be sent, during a scan, to another set of modules using the second ROD serial port. Modules that are sent this trigger are identified by the third bit of the group number, the first two bits of which identify which slave DSP events are sent to.

All these parameters are gathered into a scan description object which is passed to the SctApi by the Calibration Controller in order to request a scan. When this is received by the SctApi, it checks the current status of the hardware and modules and then starts the scan. The Formatters are set up to process the events from the modules and the Router is set up to trap the events into the correct slave DSPs. The histogram task is then started on the slave DSP and the histogram control task is started on the master. The histogram control task implements the feedback functionality which ensures that events arrive at the slave DSPs only as quickly as they can be processed.

Once the scan has started, SctApi passes control to a second thread which monitors the progress of the scan on all the RODs. This enables the method that initiated the scan to report in a timely manner to the Calibration Controller. The scan thread polls the progress of the histogram on the RODs. If no progress is made after 5 seconds then the histogram is aborted. When the histogram is complete the histogram buffers are read from the slave DSPs. Histograms from the individual modules are split up, a header is added and the result is published to the IS server for event data. From the IS server, the Fitting and Analysis services read the histogram and process it appropriately.

4.8.7 Raw Scan Production

The ROD provides different levels of interpretation of the data it receives. The lowest level is the access to the unprocessed input memories of the Formatters. This is useful for checking on the module status, and especially for adjusting the parameters of the BOC so that the bit-stream is reliably decoded.

The SctApi provides some different methods to read out these raw code bit-streams. The first is a modification of the histogram production procedure to produce raw scans. In each bin of a standard histogram is stored the hit count of all 768 channels on one side of a module. In a raw scan, the data from a link, from a number of bunch crossings, is recorded and accumulated in each bin. This allows the analysis software to examine the bit-stream as seen by the input to the ROD: after digitising by the BOC, but before being decoded by the formatters.

An additional option to a raw scan is to switch the BOC into its half-rate sampling mode. This is useful for reading the clock divided by two signal from a module as it means the BOC reads either a 1 or a 0 depending on the phase of the receiver strobe. A scan over all the settings of the BOC receiver delay parameter then results in a histogram with either 100% or 0% "occupancy". The step where the occupancy changes should be avoided as it is unreliable.

4.8.8 Simple Raw Event Decoding

The second method for reading out raw bit-streams is the probe. This records a brief section of the bit-stream in response to a trigger signal and decodes the result. It then assigns a single character code based on the type of response it gets. This allows a clock-divided-by-2 signal to be distinguished from a configured module, as defined by a match of the event header, a null response, a trigger, or a reset signal (these last two are found only with a dummy feedback module).

These codes allow a powered off module, an unconfigured module and a configured module to be easily distinguished. This procedure is carried out before each scan, before turning on the formatters so appropriate action can be taken before the scan fails.

The probe procedure can also be carried out in a scan over a configuration variable. This enables a clock signal to be recognised as such. The scan is set up over the BOC receiver delay and a change from 1 to 0 over the scan (or vice versa) is expected.

4.9 Special Scan Types

The standard scan provides a histogram of the occupancy on a module as the threshold is varied. This is normally carried out with a particular amount of charge injected. Many other scans are possible by modifying the various parameters on the scans. The following gives examples of some of these scans and how they make use of different features.

4.9.1 Configuration Check

A special scan, based on the probe scan, can be used to check the configuration of the connections between modules. Turning off the laser transmit current for all but one BOC channel will result in a signal returning from the module that it is directly connected to plus the redundant partner if it is set up. A scan is therefore implemented to turn on the tx channels on a ROD one at time and record the results. This can then be compared to the configuration found in the database.

4.9.2 Pipeline Test

The pipeline test is part of the standard characterisation sequence. It provides a systematic test of all the pipeline cells on a chip. Which set of cells is used in a chip is determined by the BC counter at the time. Therefore the pipeline test examines the data that results from varying the spacing between a soft reset and the L1A trigger. As a pulse travels through the pipeline, it is always stored in one of twelve sets of bits. The pulse is changed from all 1 to all 0 by modifying the mask register. This allows two fault conditions to be checked: a cell is stuck on, or a cell is stuck off.

The trigger sequence requested for this scan consists of the soft reset followed by a delay. Next is a command to the chip to inject a pulse of data into the pipeline, finally the L1A is sent a fixed time after the pulse, which corresponds to the length of the pipeline. This is the most complex trigger requested in the standard characterisation sequence. As a soft reset is issued before every trigger, this means that the L1 and BC counters must be handled differently on the ROD to other scans.

4.9.3 Full Bypass Test

The Full Bypass Test is another test from the characterisation sequence. This tests all the bypass token links between chips, using a predefined sequence. During part of this sequence, the module is configured so that one read-out link is not used. Instead, the master chip attached to the other link is used to read out a chip that would normally be read out by the first. This is treated specially by the ROD in that it must turn off the decoding by the formatters for the links that are turned off.

Another effect of changing the read-out chain is that the chip numbering changes. The EFB firmware checks the chip numbers to see if they are valid. This enables a bit swap in the bit-stream to be recognised. The chips are numbered 0, 1, 2, 3, 4, 5 on one side and 8, 9, 10, 11, 12, 13 on the other side. Bit 4 indicates the side of the module the chip is on. Instead of doing a straight test for monotonicity, the EFB must first toggle bit 4 depending on which side of the module the chips are being read out from.

4.9.4 Noise Occupancy

A final characterisation test, the Noise Occupancy test, scans thresholds around the nominal 1fC threshold, with no charge injection and measures the occupancy. Over the range of thresholds measured, the occupancy recorded can vary from around 0.1 to below 10^{-5}. In order to measure the (lowest) occupancy at high thresholds, of the order of a million events are sent to the modules. At low thresholds, and high occupancies, recording this many events takes a prohibitive amount of time, therefore the number is reduced.

The DSP histogram code is constrained to send a constant number of triggers for each threshold setting. In order to send different numbers of events to different bins in the histogram, the procedure is repeated starting at a different bin each time. The first round puts the lowest number of events in all the bins. Subsequent repetitions accumulate more

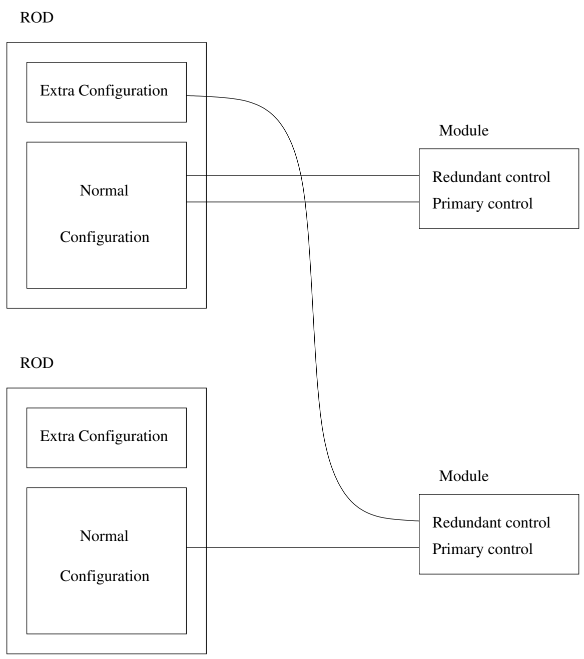

Figure 4.6: Two modules controlled by two RODs. The lower module is configured by the same ROD as the upper module, using configuration stored in an extra buffer

events in the later bins until all the requested events have been sent.

4.9.5 Redundancy

The redundant clock and command that can be provided to the modules provides additional complications for the system. The redundant input select signal is provided to the modules by DCS. The configuration data sent to a module contains an address which specifies which of two modules is configured. This means that the DAQ must agree with what the DCS sends to the modules.

Particularly within the end-cap structure, there are tx fibres that will control modules in the redundant configuration that are not in the same ROD or even crate as the primary module they control. For SctApi and for the ROD, this makes a difference when configuring modules and also when carrying out scans. As the select signal selects the redundant transmit only, and does nothing to the data signal, events from the modules are handled in the same way for both settings and only the master DSP knows the difference. The module configuration for the module being controlled from a different ROD is placed in a reserved area of this ROD, where the information is not transferred to the slave DSPs. The implementation of this is shown diagrammatically in figure 4.6. When in this mode, a ROD driven scan cannot be done. This would require coordination between the RODs about when to send the next trigger. Scans with redundant connections across RODs must therefore use triggers provided by the TIM.

4.9.6 Synchronous Triggers

Histograms can also be generated using the TIM as a trigger source. This exercises support for external triggers within the ROD and allows a single trigger to be simultaneously sent to all modules in a crate. These are both essential within the ATLAS experiment. Using this mode to generate calibration histograms introduces a significant problem. The feedback loop that is normally implemented in the master DSP to stop triggers being sent faster than they can be processed by the slave DSPs must instead include the SBC, all the RODs and the TIM.

A series of options are implemented in SctApi, which can be chosen as part of the scan request.

1. Ignore the problem. This forces the user to set the trigger frequency appropriately for the expected occupancy, or to trap into the slave DSPs only as many events as can be processed.

2. Full feedback. Send one trigger at a time and monitor the ROD for completion. This is very slow.

3. Semi feedback. Send a burst of triggers small enough to fit in the event FIFOs on the ROD, then monitor the ROD for completion.

4. Variable trigger rate. For each bin, measure how fast the events are processed using a small number of events. Then set the trigger rate appropriately.

All these options have problems, mostly that the trigger rate is very unpredictable. The first option is the simplest and is adequate for scans with a low enough occupancy. For noise measurements, such as to allow checking for any extra noise caused by running many modules at once, the occupancy is low enough to run at a nominal 100kHz.

The other options are provided so that it might be possible to carry out a full calibration sequence using synchronous triggers from the TIM. This is unlikely to provide additional useful information and requires additional development to work reliably.

4.10 Fitting Service

Once the histograms have been produced by the SctApi and published to the event data IS server, they are ready for analysis. The first stage is to make fits to the data for each channel. The Fitting Service monitors the event data IS server for new histograms. When a histogram appears, a selection is made based on the configuration variable being scanned over. This allows a choice to be made between the error function for threshold scans (S-curves) and a top-hat function for strobe delay scans. Once the fit has been made, the results are published to another IS server.

4.11 Analysis Service

Analyses are produced in a similar manner to fits, but data for all channels in a module are taken account of. The IS server with test control data from the Calibration Controller is monitored along with the IS servers for both event data and fitted data. Once the appropriate data is available the analysis can be performed.

Each test is analysed by a separate class which registers interest in the name of the test as reported by the Calibration Controller data. The result of the analysis of a test is a series of numbers describing the high level behaviour of a module. This can be displayed in various ways by the GUI. For instance after a 3 point gain test, the analysis produces information on the gain, offset and noise for each channel. These can be displayed in channel by channel graphs for each module. Alternatively the data can be summarized to produce a report in the same format as SCTDAQ uploads to the production database. This data is then used by the GUI to summarise the results for many modules, colour coding the graphical representation of the barrel or end-cap. Finally, data from the analysis is used by the Calibration Controller to update the configuration of each module.

4.12 Archiving Service

The IS servers store data that is produced online. The data is lost when the programs no longer run. The Archiving Service is responsible for saving the data to persistent storage. It monitors the IS servers that the other services place data on to. This is then serialized and stored permanently on disk.

4.13 Configuration Database

The final part of the SctRodDaq system is the configuration database. This stores the configurations of all the modules, RODs, BOCs, TIMs and power supplies in the system. This is used extensively by the SctApi to find out which RODs are in each crate and which modules are connected to which ROD. The GUI also uses the configuration to find out where each module is placed on the barrel or end-cap.

The power supply configuration is stored in the firmware of the power supply crate controllers. It is transferred to the DCS system outside of a run to save time when the parameters do not change and to avoid rewriting the flash memories more than necessary.

The bulk of the configuration needed by the SctApi are the parameters required to set up each module. In particular, the trim settings are different for every chip in the SCT. The configuration database also stores mappings of modules between different naming schemes for use by the user interface.

4.13.1 Module Indexing Schemes

During the module production phase, the primary scheme used for tests is the module serial number. This is used to index all the tests in the production database. For macro-assembly and for the final experiment this single address is not enough. The DAQ system needs to be able to: look up a configuration in the production database; control the correct module and receive data from it via the RODs; request changes in the power supply parameters from the DCS system, and also allow a geographical visualisation of the modules on the detector, the barrel or disc.

A particular ROD channel can be described by a set of four numbers: one of four "partitions"; a crate – of which there are 2 per partition; the slot a ROD is placed within the crate, and finally one of 48 channels in the ROD. Additionally the tx and 2 rx fibres

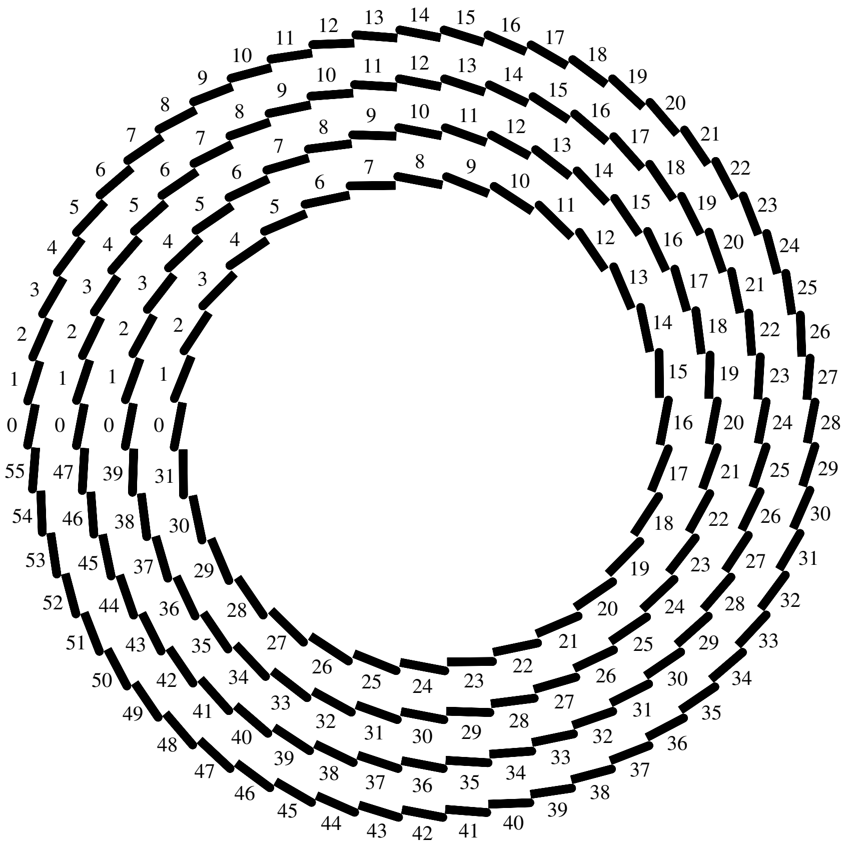

Figure 4.7: The row positions on the SCT Barrel

may be remapped from their defaults.

A module channel in the DCS system can be referenced by the crate (one of 88) and the channel in the crate (one of 48).

The barrel geography (the position of a module on the barrel) can be described by: the barrel number (3, 4, 5, 6), the row on the barrel (up to 55 as shown in figure 4.7) and a position along the row. 12 modules are distributed along the row on the z axis with no module at $z = 0$ so they are numbered -6, -5 ... -2, -1 and 1, 2 ... 5, 6.

The end-cap geography can be described by the end-cap number (-9 to -1 and 1 to 9), a quadrant number and a number which represents one of the 33 possible positions within an end-cap quadrant.

The configuration service therefore has to store and retrieve mappings between these 4 addressing schemes. In order to do this a canonical address is used within the configuration server which is based upon the grouping of optical fibres at the ROD: each bundle carries fibres for a maximum of 6 modules. This system is described by a number representing the bundle, or MUR (minimum unit of readout), and a module position (1-6) within the MUR.

The MUR scheme allows a simple mapping to both the ROD and the barrel geographical scheme: 8 harnesses are plugged into each BOC card in the ROD crate, and 2 harnesses support the modules on one row of each barrel, one for +z and one for -z.

For the end-cap a more sophisticated scheme is necessary. This is because the module scheme does not divide easily into the sets of 6 naturally supported by the ROD. On each quadrant of an end-cap disc there are up to 33 modules, which are connected to 6 harnesses as shown by the colours in figure 4.8. One harness is split between rows to service the 13th outer module.

The power supply mapping is also not tied directly to the MURs. One scheme for the barrel portion involves connecting the twelve modules on one half of two neighbouring rows. The configuration database stores the mappings directly from an MUR/module pair to a crate/channel pair.

4.13.2 Implementation

The Configuration service uses an IPC interface wrapper around a C++ class in a similar manner to the SctApi. The C++ interface is used directly by the TApi version of the SctApi interface to avoid unnecessary use of the IPC system. When run within SctRodDaq, the SctApi uses an interface that looks the same but wraps the IPC calls to the configuration service. This allows the GUI to make use of the same database and enables

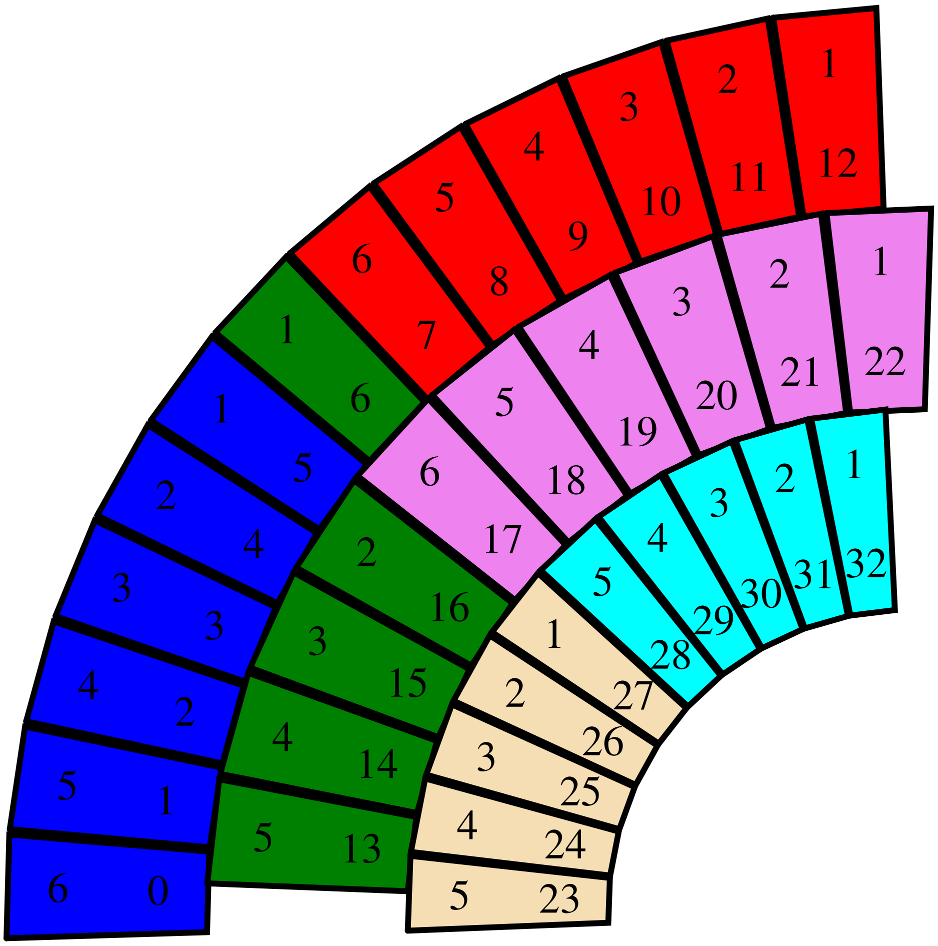

Figure 4.8: The positions on the SCT End-cap. The outer number represents the fibre number within an MUR. The inner number is the module number within the end-cap quadrant. Colours represent the six different harnesses

small modifications to be made.

All the information for the configuration service is stored in XML. This allows different types of structured information to be stored together, while support for XInclude [55] XML elements allows the data to be distributed among multiple files. The configuration can then be arranged on the file system in a more user friendly way, for instance by using one file for each module configuration. A utility program enables the XML file corresponding to a module to be downloaded from the production database.

The XML file is read using the libxml2 library [56]. This reads and parses the whole XML file into a tree structure in memory. Once there, both direct access to the tree and also XPath [57] queries are used to find information in the structure.

4.13.3 Optimizations

Optimization of the configuration service for speed didn't occur during development. The majority of functions used were found to take an acceptable time when compared to the start-up time due to resetting and loading the ROD hardware. Both the XPath queries and a manual search look for items in the tree by looking at element names and attributes in turn. This is a linear search which takes O(n) time to find an item in a list. For some queries, the time taken when larger numbers of modules are stored in the configuration was found to be prohibitive and the linear search was found to be the cause.

Finding the MUR number a module is connected to, given the module serial number, involves searching through all the MUR definition nodes for the serial number string. This is made worse by the fact that the serial numbers vary most in their last digits which means a comparison must always check most of the string even to find that they do not match. A second, similar, case is searching for a particular module configuration which is once again indexed by the module serial number.

Instead of the linear search, an index of the data is created based on the standard C++ `std::map` class. This allows the string to be mapped to a pointer into the tree using a heap. A match can then be found in O(ln(n)) time, where n is the number of modules. When an index was created in a `std::map` the time taken to do the look-up was reduced by a factor of 100 for a configuration file with a crate's worth of modules.

4.14 Summary

The SCT modules need monitoring at all stages of the experiment. During module production this was carried out by the SCTDAQ software, but a new system was needed for the numbers of modules in use during macro-assembly.

This new system is named SctRodDaq and is made up of several distributed components. The central part is the SctApi, which interacts with the ROD to configure modules and produce histograms for calibration.

The functionality of this system is demonstrated in the following chapter.

Chapter 5

SCT Barrel Macro-Assembly

Macro-assembly is the term used to describe the attachment of the SCT modules onto the barrel and end-cap support structures. SCT barrel macro-assembly is carried out at Oxford; end-cap macro-assembly is carried out at Liverpool and NIKHEF.

An important part of the macro-assembly procedure is the exercise of the electrical functionality of the module, by running appropriate tests to check continued operation. This is an important reason for the development of the SctRodDaq software described in chapter 4. This chapter describes the use of the SctRodDaq system at Oxford, including some configurations used during development of the software.

The following describes the hardware configurations on which the software has been tested, starting from small systems and extending to a complete barrel. Some results of tests run on the final Barrel 3 are also presented, which demonstrate the performance of the silicon detector modules and of the readout system.

5.1 Readout Hardware Configurations

All readout hardware configurations at Oxford are read out through one of two VME crates, which is enough to read out the whole of barrel 6. This contains the standard TIM, SBC and ROD/BOCs, as described in chapter 3. Two Linux based computer-systems run the calibration software: the disk-less SBC and a standard desktop PC as a server. The SBC is connected via ethernet to the server. The only executable (apart from ATLAS Online infrastructure) the SBC runs is SctApiServer, which implements the SctApi IPC interface. The remainder of the SctRodDaq system runs on the server.

The crate is an ATLAS standard VME crate which has 22 slots, the first 5 of which are 6U high, while the remainder are 9U. The SBC, as the VME bus master, must go in slot 0. The 9U slots contain the TIM and the RODs. This leaves 4 spare 6U slots which are available for other purposes. Figure 5.1 shows a picture of the cards in the crate.

Three module configurations were used at Oxford as part of the development for the full barrel tests: a two module system for small-scale tests, a larger test system based on part of a barrel, and the full barrel 3. A second system at Cambridge with four modules attached was also used, in a similar manner to the two module configuration described below.

5.1.1 Two Module Configuration

A simple module set-up in a box was used primarily during software development. This consists of two boxed SCT barrel modules on a desk, protected by a wooden box, see photo 5.2. The final version of the electrical harness was used to provide power and data links to the modules, but the modules were contained in the same boxes used for electrical testing during production. These allow chilled water to pass close to the module cooling plate to keep the modules from overheating, and they also provide support for

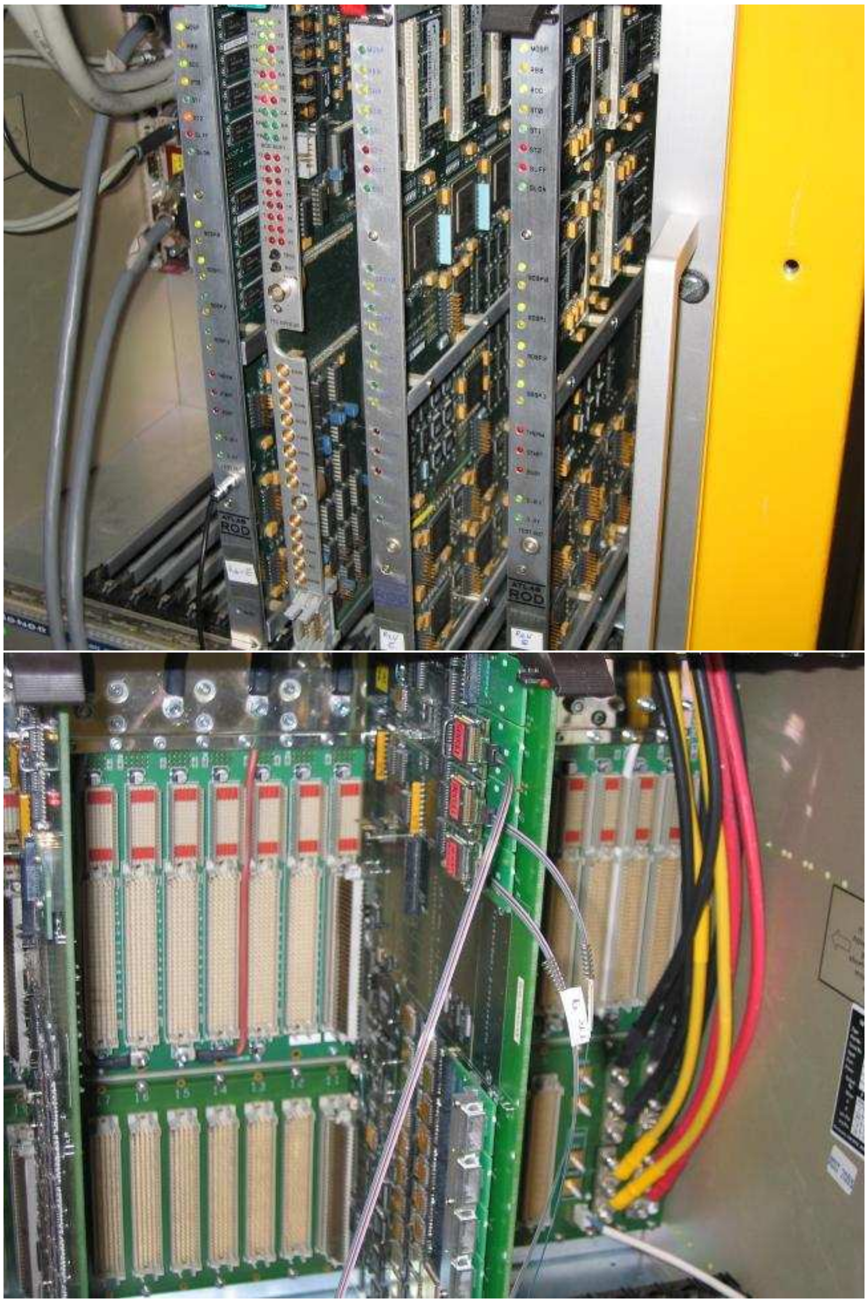

Figure 5.1: Photograph of cards in VME crate. The top picture shows three RODs with a TIM in the centre, to the left are power supplies and the SBC. The bottom picture shows a BOC in the rear of the crate with three tx fibre connections

the modules as the barrel is not available. Special patch cards are used to convert the electrical connections from the patch cards in the box to the production harness.

A pair of fibre ribbons is used to connect the two modules to the BOC, one for transmit (tx) and one for receive (rx). A standard MT-6 fibre is used for the tx connection, this is different from the final system where two 6-way fibres at the barrel end merge to one 12-way fibre connected to the BOC. The non-standard fibre connection involves a shift of the tx fibre channel relative to the nominal position. This requires the propagation of these mappings to the software configuration.

Patch cables from the harness allow the use of the same module power supplies as were used during module production, but which were not used in the remaining configurations reported here. The SCTHV and SCTLV were placed in slots 2 and 3, in the 6U area of the VME crate. The power supplies were driven from a Linux port of the old SCTDAQ software, also running on the SBC. As an alternative, used in some runs, the final ATLAS power supplies could be used, with a short cable to a dedicated crate. This required the use of the final, PVSS based, software system which is more complex and was also under development at the time.

5.1.2 Configuration of Test Sector

Development progressed using a barrel sector. This was designed for the earlier system tests and was also used for commissioning of the cooling system and other services. The full length carbon fibre sector of a barrel (the equivalent of number 5) was fitted with one cooling stave (a single loop of pipe). Three harnesses were attached to the sector. These were not of production quality, all having known bad connections and swapped fibres.

A total of 13 modules were mounted on the three harnesses. Fibres were connected as in the final detector: two MT-6 fibre bundles from the modules were joined into an

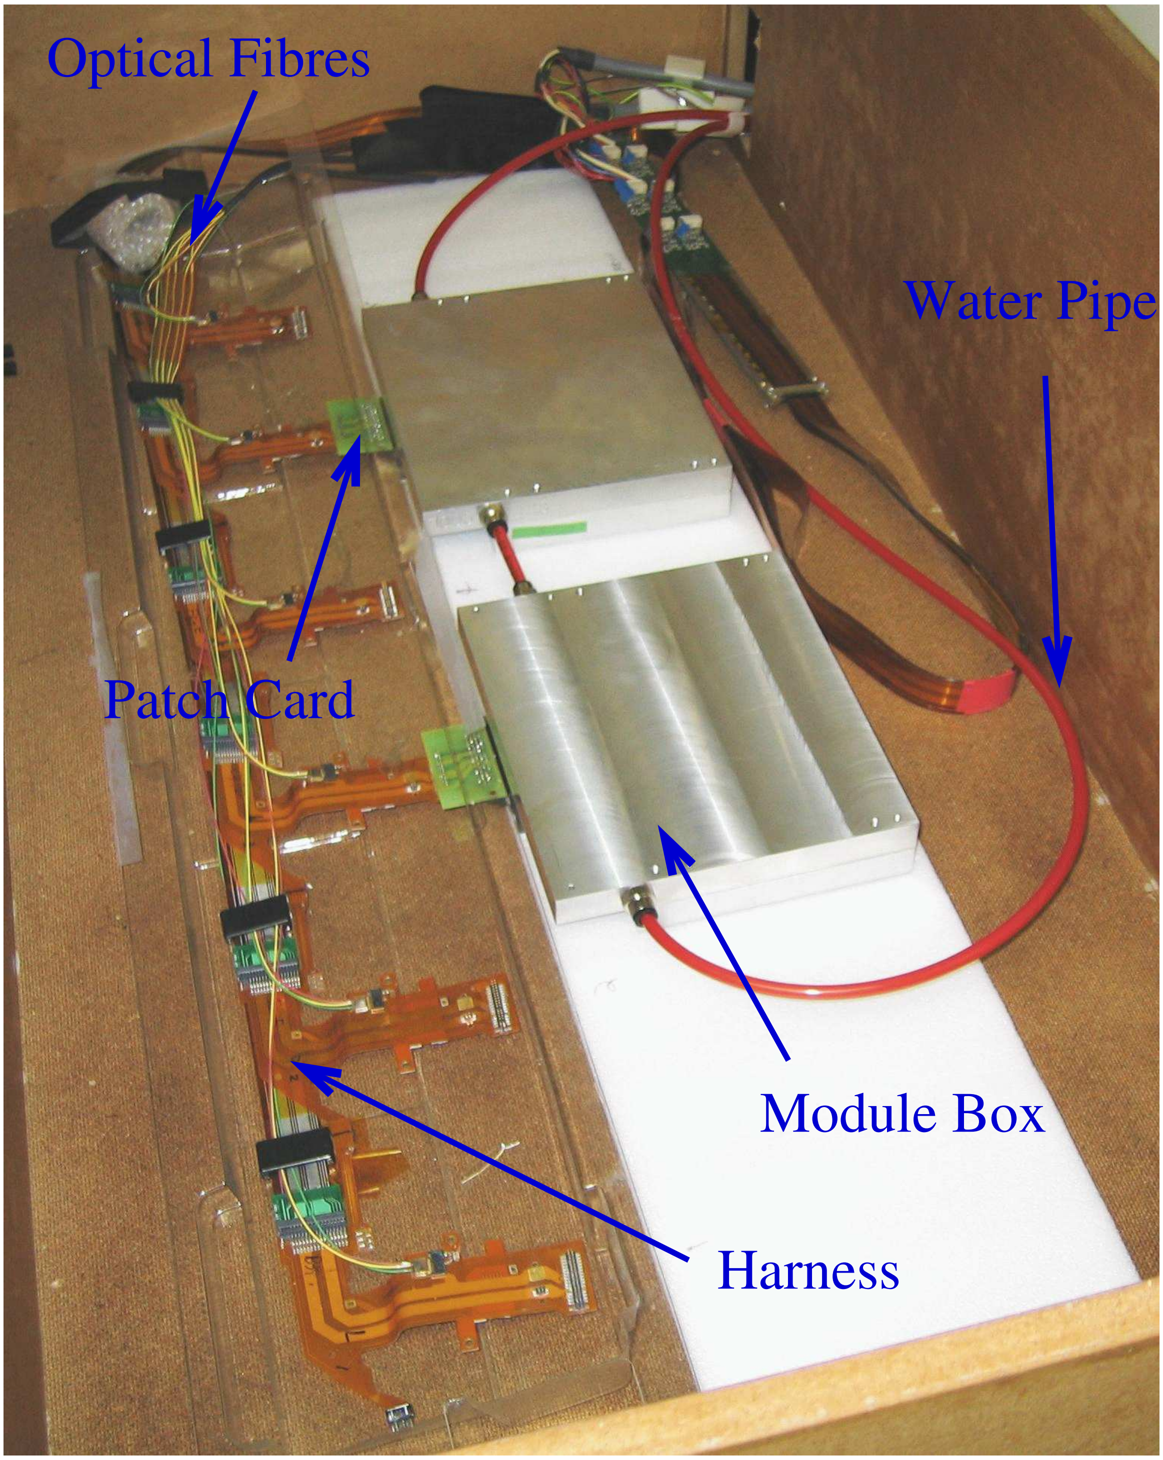

Figure 5.2: Photograph of modules in the box

MT-12 bundle plugged into the BOC.

The series power supplies were used due to the greater number of modules. These were housed in a dedicated power supply crate and controlled by a PVSS based software system running on a separate Microsoft Windows computer.

The modules were cooled using C_4F_{10} running through the cooling tubes with the environment at room temperature. The sector was placed in the same cold room as for the barrel macro-assembly, allowing module mounting procedures to be exercised, including the use of the robot [58]. This uses image recognition software to locate the module in its reception box and the correct barrel position before manipulating the module from the box to the barrel, avoiding obstacles created by modules and the harness on the barrel.

5.1.3 Configuration for Barrel 3

The final configuration was the full barrel 3 during macro-assembly. Modules to be mounted on the barrel arrived at Oxford from the four production clusters. They received a quick reception test using SCTDAQ before they were mounted onto the barrel using the robotic system. This verifies the configuration acquired at the production sites and that it has not changed during transportation. This configuration was then used as the initial state for the modules as they are tested on the barrel.

The barrel was housed in a specially constructed cold room. The cooling system made use of evaporative cooling with either C_4F_{10} or C_3F_8 as coolant. C_4F_{10} was used during room temperature operation, when the room required access. The cooling system for C_4F_{10} was less powerful, having enough power to cool only 24 to 36 modules. C_3F_8 was used to cool all of the modules to a detector temperature of around 0°C, as for running in ATLAS. This required careful control of the temperature and humidity within the room precluding casual access.

The operation of all 384 modules on barrel 3 requires 8 RODs and 8 crates of power supplies. These were connected to the barrel using cables and fibres to be used later during commissioning at CERN. A single fibre run connected the BOC to a short fibre extension from the harness on the barrel. The fibres were cut to length suitable for testing in the surface building at CERN. The excess fibre length formed a coil on the floor of the cold room.

The power cable connections were different for the macro-assembly from the final ATLAS system. Using the final power supplies requires the use of the appropriate connectors. The module ends of these cables from the power supplies were attached to temporary patch panels which connect directly to the patch panels on the barrel. This system bypasses two runs of cable that will be in the final system and also the chokes which are found in the patch panel on the cavern walls. The noise behaviour of the barrel during macro-assembly is therefore not directly comparable to the final system.

The absence of chokes will reduce the noise rejection of the system, but the cold-room environment at Oxford may be less electrically noisy than the final configuration inside ATLAS. Nevertheless noise figures are expected to provide a reasonable indication of performance in the final system.

5.2 Tests on Box Set-up

The ROOT based TApi interface (section 4.8.1) was used for running tests on the boxed modules. This allowed the state of the ROD to be examined in more detail for diagnostic purposes, using ROOT scripts to decode bit masks and toggles contained in the registers. Useful registers during histogramming include the status of the event trapping functionality and how the ROD formatters are set up. Additionally, options are present to output a list of all primitives sent to the ROD, in both binary and a human readable form. This is particularly useful during implementation of new ROD primitives and histogramming procedures.

Some of the first scans performed using the ROD system were the "NMask" scans. These provide a fixed pattern of data to the ROD for decoding; using the "send mask" bit in the ABCD configuration to copy the mask register into the pipeline. The mask is then varied across a scan, starting with all the channels on, and removing one channel at a time from the mask.

The resulting histograms are very regular and characteristic. The well known hit pattern allows a simple check that the ROD is histogramming the events from the modules correctly. Several potential problems with the histogramming and histogram read-out are diagnosable using the raw output from this scan, particularly with the help of a simple graphical display.

The test can also show up a variety of problems with setting the chip mask register. This test has therefore been included in the characterisation sequence for macro-assembly.

Figure 5.3 shows the result of an NMask test performed on a test module. This is in the standard histogram style for the ROD occupancy histograms. The two plots show the channels on the top and bottom sides of the module respectively. The second side is zoomed in to show a small region of the histogram. The x axis marks the channel within each side, the y axis shows the variable changed for each bin in the histogram and the value it took.

Of note is the extra occupancy visible in channel 116 of the second side, which is caused by a stuck cell defect in the pipeline for that channel. This fault is persistent and was also observed in previous test runs. During development this allows trivial identification of a particular module. This enables the physical module to be linked to a histogram buffer in the ROD. Otherwise, there is no information transferred from the module to provide

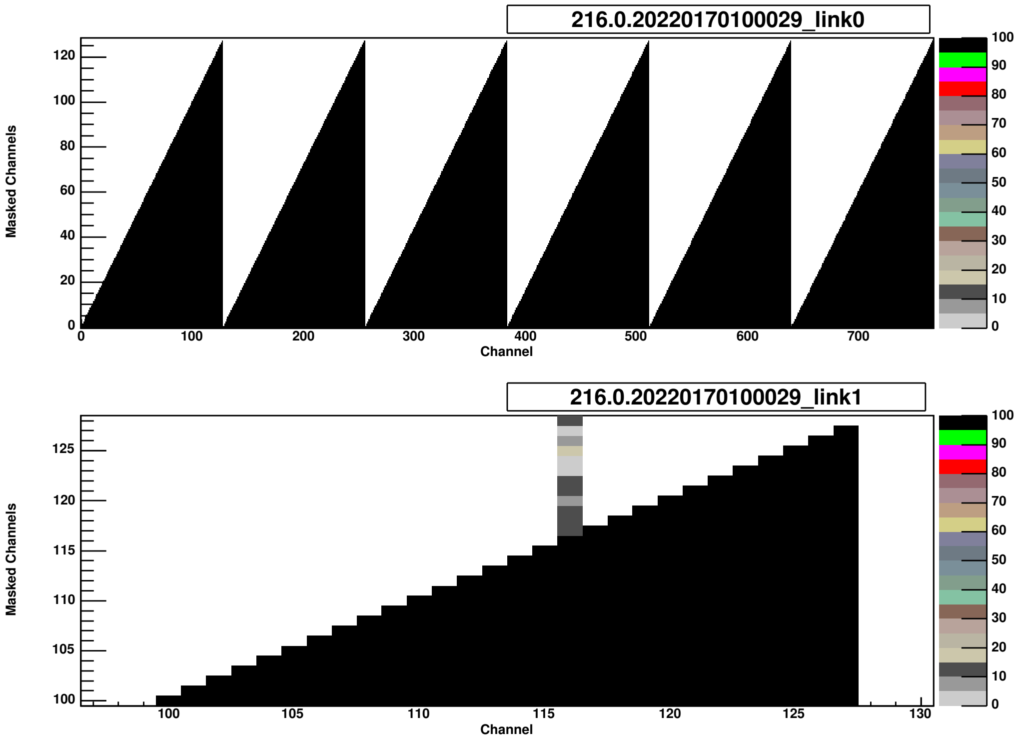

Figure 5.3: Plot of an NMask histogram. The top plot shows all of the channels on the first side. The second side is zoomed in to show a pipeline defect in channel 116

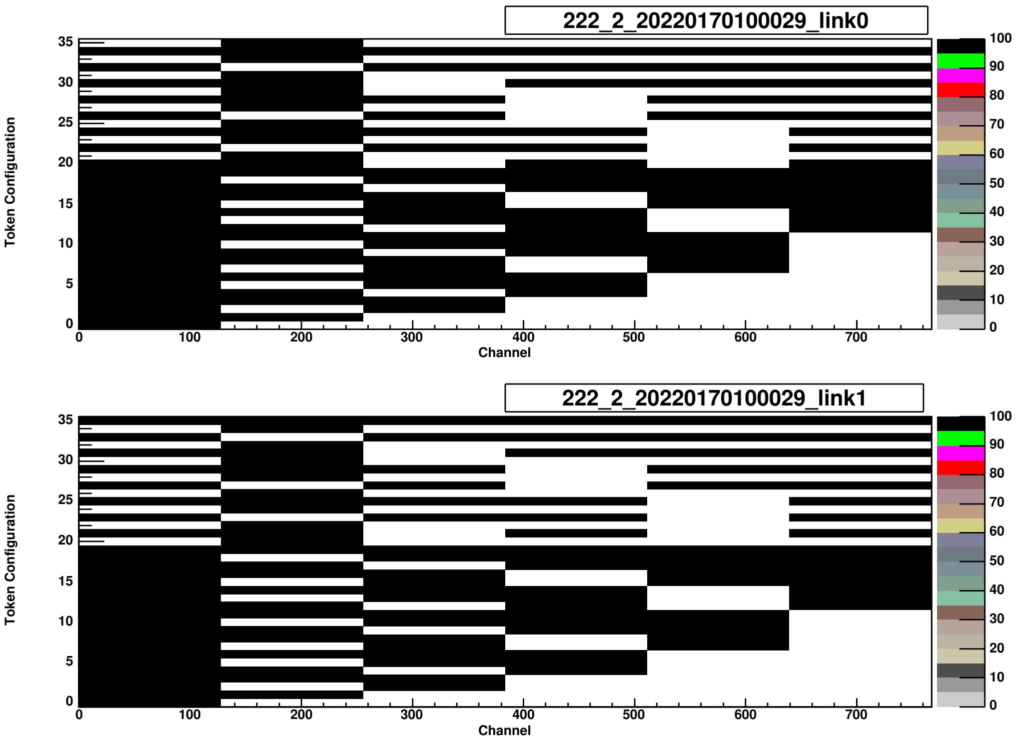

Figure 5.4: Histogram of a Full-Bypass test. For value 20 and above, the bins are read out by alternating sides

any identification to the DAQ system.

The Full-Bypass test exercises the module token passing. Example results are shown in figure 5.4. Each chip returns full occupancy when it is not bypassed. When bypassed no data is input in the histogram. The patterns below 20 bypass individual chips on both sides together. In pattern 20 and above, the master chip on one side is bypassed and up to 7 chips are read out through the master chip on the other side.

Within the ROD code, this test exercises the modification of formatter channel masks during a scan as described in section 4.9.3. Otherwise bad data is introduced when the formatter tries to find the event stream.

Between these two tests, the majority of basic ROD histogramming functionality can be verified. Remaining functionality, particularly the calibration line loop, is checked in a full characterisation sequence.

5.3 Sector tests

The barrel sector allowed testing and development of the software system for use with larger numbers of modules than possible with the box set-up. These modules were used to show up problems with the software at a higher level. Tests of both the SctApi and ROD firmware included: verifying histogramming with more modules than before, particularly the use of multiple harnesses; modules shared between groups in different combinations; and configuring modules in different patterns using both output serial ports on the ROD.

The analysis system and user interface were also used, enabling examination and comparison of the module characteristics. This also enabled the throughput of the system to be tested with real instead of simulated data.

The initial configuration for each module was taken from earlier test runs using SCT-DAQ. The chosen modules have known defects, including some that are unlikely to appear in the final experiment (the use of previous versions of the read-out chip, where the calibration charge is out of step compared to the current version). This allows the defect detection ability of the system to be verified using real modules.

It was found that the fibre mappings could be difficult to discover for multiple modules, where MT-12 fibre connections might be swapped or put upside down. This led to the development of the configuration checking scan described in section 4.9.1.

The sector provided the first use of the DDC (DCS - DAQ Communication) system to transfer DCS data from PVSS to the DAQ. This allows current information about the modules to be stored with the results from tests. It also allows for future integration, such that the DAQ may respond to events such as loss of bias on a module in a controlled manner. This enables a test of all the hardware and software reading out real modules to check for any adverse effects before being used on the final barrel system.

5.4 Barrel 3 Testing

Barrel 3 (the innermost one) was the first to arrive in Oxford after being fitted with the cooling system and harnesses at RAL.

The first tests were of the harnesses. These were checked for continuity using both the power supplies and the ROD DAQ system. Dummy modules with loop-back connections were attached to each position. Power supply currents and voltages were checked and the data path checked using an analysis based on the raw ROD data, as described in section 4.8.8, sending and receiving an L1A trigger signal.

Modules were mounted in stages and electrically tested before the next set is added. Lower modules (as described in section 3.1) must be mounted before upper modules so they can be replaced if necessary. During mounting, up to 36 modules were tested simultaneously, using two RODs at once. Figure 3.2 shows the barrel at a late stage of assembly.

Once all the modules were mounted they were retested under warm conditions (module temperature 30°C, cooling around 10°C) in groups of 24 with the final power supply connections, using the C_4F_{10} system. The whole system was then cooled down (module temperatures at 10°C using cooling to −8.5°C) for the cold tests to provide for more realistic running conditions.

During mounting, the modules were tested using a cut down version of the characterisation sequence: the NMask test, Strobe Delay, Pipeline, Full-Bypass, Three-Point Gain and Noise Occupancy. Once mounted, longer tests were run on the complete barrel. These include the addition of the new Double Trigger Noise and Synchronous Trigger

Noise tests in the cold run. The Synchronous Trigger Noise test is the only test carried out where all modules receive a trigger at the same time, as in the experiment. All the other tests use triggers generated independently on individual RODs. The results of these tests have been analysed in several ways, as described in [59, 60, 61, 62, 63]. Some of these results are described below.

5.4.1 Barrel 3 Noise Measurements

The noise is one of the critical parameters which will effect the physics performance of the SCT. Though a direct comparison to the specification is not possible, due to the differences in the noise suppression, the measurements can be used to verify uniformity and to pick out modules with particular problems. The noise can be measured in two ways: interpolated from the noise measurements in the threshold scans using the calculated gain, or from the noise occupancy measurements at different thresholds. A fundamental difference between these two methods is due to the presence of the injection charge in the threshold measurements, which introduces extra chip circuitry. Additionally, the two analysis methods make different assumptions about the structure of the noise and of the front-end amplifier response curve at different threshold levels .

Figures 5.5, 5.6, 5.7 and 5.8 show, in electrons, the ENC (equivalent noise charge) values for all the modules on barrel 3. The “noise” value is the input noise value calculated from the response curve test. The “noise occupancy” plots show the noise value calculated from the noise occupancy test. Figure 5.9 shows the distributions of the results.

The cold tests in general have lower noise due to the lower detector leakage current and therefore the associated noise. There is a structure visible in the noise maps, where the lower half (rows 16-31) has a lower noise value than the upper half. This is due to differences in the leakage current caused by the difference in temperature between the top

	-6	-5	-4	-3	-2	-1	+1	+2	+3	+4	+5	+6
Row 0	1593.02	1666.62	1708.53	1640.54	1688.94	1660.48	1602.81	1485.96	1718.72	1600.04	1611.14	1601.00
Row 1	1602.64	1704.39	1649.93	1633.06	1672.09	1616.61	1707.36	1687.41	1652.65	1652.79	1658.10	1669.69
Row 2	1605.88	1748.64	1718.70	1755.28	1752.37	1696.93	1647.20	1629.21	1684.85	1660.98	1722.73	1687.31
Row 3	1738.83	1693.04	1657.33	1642.12	1701.08	1623.05	1697.21	1644.85	1638.32	1664.70	1717.87	1746.11
Row 4	1709.62	1648.94	1675.10	1670.42	1695.59	1581.37	1662.44	1616.43	1633.39	1665.92	1650.87	1578.80
Row 5	1687.04	1664.58	1718.56	1691.82	1719.45	1675.88	1654.44	1518.57	1634.99	1587.52	1606.12	1684.97
Row 6	1681.29	1611.16	1697.16	1601.42	1734.37	1664.28	1661.91	1656.69	1618.31	1621.11	1689.36	1660.16
Row 7	1609.04	1657.87	1587.75	1655.31	1685.39	1686.70	1622.25	1609.86	1673.04	1653.98	1737.60	1669.27
Row 8	1749.62	1653.88	1708.21	1730.20	1759.72	1740.88	1655.32	1629.98	1595.94	1631.17	1580.23	1629.98
Row 9	1654.59	1676.15	1689.50	1708.53	1606.80	1549.06	1753.87	1695.09	1645.21	1624.71	1676.00	1676.42
Row 10	1603.93	1631.26	1636.95	1634.52	1603.22	1615.22	1638.93	1638.27	1565.44	1572.36	1732.99	1764.25
Row 11	1541.32	1575.39	1535.25	1546.91	1553.24	1563.26	1711.28	1653.95	1686.60	1622.66	1616.34	1667.76
Row 12	1554.26	1622.42	1580.30	1630.70	1595.23	1631.18	1737.97	1615.88	1632.12	1620.42	1720.29	1702.08
Row 13	1729.29	1767.67	1609.15	1651.22	1641.80	1664.90	1607.96	1631.61	1602.03	1703.99	1631.52	1624.70
Row 14	1511.04	1626.97	1736.59	1717.73	1728.47	1708.88	1652.56	1599.05	1623.25	1599.24	1651.68	1587.76
Row 15	1753.52	1784.45	1691.08	1784.01	1611.52	1701.99	1621.21	1639.54	1596.43	1637.40	1641.53	1613.09
Row 16	1616.36	1627.55	1563.10	1663.15	1665.05	1662.61	1573.35	1623.22	1607.86	1586.53	1663.77	1571.89
Row 17	1579.01	1513.87	1516.21	1509.37	1489.63	1532.26	1580.06	1614.84	1660.70	1670.50	1619.51	1522.90
Row 18	1551.91	1518.43	1548.34	1571.55	1585.18	1532.59	1597.57	1611.88	1633.22	1667.06	1567.13	1572.88
Row 19	1489.61	1537.12	1675.34	1610.21	1543.10	1670.61	1694.83	1570.34	1575.82	1554.42	1634.96	1618.61
Row 20	1657.68	1613.60	1521.12	1545.93	1424.47	1543.02	1515.04	1624.27	1692.77	1461.31	1665.57	1612.07
Row 21	1493.44	1572.18	1615.22	1522.62	1656.89	1462.68	1642.50	1628.58	1602.85	1582.27	1580.55	1506.64
Row 22	1558.73	1550.63	1595.44	1538.69	1549.83	1531.41	1474.93	1505.25	1430.60	1524.95	1589.04	1575.37
Row 23	1498.19	1538.27	1629.83	1598.67	1563.95	1540.91	1525.00	1561.94	1494.06	1562.98	1600.15	1487.59
Row 24	1578.00	1484.40	1561.30	1554.69	1569.80	1607.62	1534.63	1466.90	1542.99	1602.98	1645.61	1556.89
Row 25	1504.69	1578.98	1574.74	1563.22	1546.96	1506.88	1657.38	1589.51	1654.28	1659.98	1628.73	1655.52
Row 26	1579.46	1500.29	1623.66	1710.14	1674.52	1652.78	1626.14	1494.27	1560.91	1553.74	1643.15	1556.12
Row 27	1683.89	1700.85	1707.25	1660.57	1497.53	1690.05	1673.57	1767.54	1607.89	1662.30	1641.73	1562.25
Row 28	1666.87	1649.07	1718.47	1668.85	1669.08	1678.16	1695.61	1659.90	1641.10	1725.60	1623.30	1623.08
Row 29	1681.50	1714.36	1720.05	1760.74	1650.99	1731.09	1591.07	1568.18	1633.48	1548.31	1609.71	1503.89
Row 30	1653.38	1602.24	1602.93	1647.31	1585.20	1630.61	1543.68	1608.76	1601.69	1607.79	1481.60	1528.62
Row 31	1614.22	1601.79	1594.37	1663.10	1561.86	1626.59	1476.51	1546.06	1534.74	1488.59	1560.71	1532.12

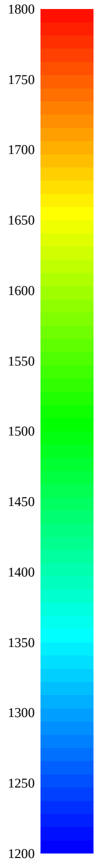

Figure 5.5: ENC noise figures from threshold scans for barrel 3 warm run, in electrons

	-6	-5	-4	-3	-2	-1	+1	+2	+3	+4	+5	+6
Row 0	1435.03	1458.79	1510.50	1507.32	1492.92	1532.11	1431.47	1312.29	1460.47	1385.49	1345.27	1425.89
Row 1	1536.56	1491.27	1430.09	1447.00	1471.62	1510.45	1397.95	1609.40	1565.95	1540.73	1441.58	1479.10
Row 2	1413.48	1508.04	1514.22	1499.26	1595.80	1459.86	1457.13	1440.43	1450.80	1477.50	1462.05	1702.47
Row 3	1593.78	1463.51	1558.83	1551.69	1664.82	1482.85	1617.15	1473.45	1465.72	1608.50	1582.96	1624.57
Row 4	1454.60	1429.89	1431.84	1450.85	1437.03	1446.04	1575.33	1518.42	1533.25	1542.44	1531.23	1474.44
Row 5	1548.25	1503.50	1542.84	1491.37	1511.07	1473.09	1459.49	1407.89	1411.33	1380.17	1391.67	1466.01
Row 6	1544.18	1442.27	1516.06	1477.67	1503.26	1568.65	1532.40	1540.53	1479.75	1473.13	1538.46	1568.87
Row 7	1476.49	1549.98	1520.69	1472.70	1500.13	1490.12	1399.86	1423.74	1341.51	1417.39	1616.26	1483.28
Row 8	1531.14	1550.14	1579.36	1611.43	1629.91	1605.68	1579.84	1539.34	1493.97	1428.94	1371.94	1416.60
Row 9	1469.08	1452.57	1490.98	1464.95	1532.99	1514.02	1529.47	1502.65	1566.13	1547.70	1564.62	1600.70
Row 10	1496.75	1491.31	1448.90	1411.86	1397.57	1410.46	1554.59	1535.46	1502.95	1574.36	1503.02	1515.54
Row 11	1491.27	1570.98	1550.46	1542.45	1360.22	1553.56	1599.84	1557.01	1500.49	1488.17	1531.79	1578.43
Row 12	1416.69	1692.80	1353.46	1432.17	1346.72	1404.68	1612.74	1533.94	1555.82	1485.49	1513.16	1446.26
Row 13	1516.04	1511.60	1467.11	1479.30	1567.40	1443.00	1495.18	1520.24	1474.72	1505.68	1533.00	1507.11
Row 14	1406.89	1413.76	1604.30	1566.27	1582.79	1598.26	1554.15	1396.18	1420.34	1414.95	1427.11	1391.42
Row 15	1525.89	1571.57	1477.48	1579.72	1516.35	1481.31	1535.41	1573.71	1488.22	1554.75	1556.40	1509.94
Row 16	1518.95	1554.50	1470.95	1595.41	1556.58	1568.44	1537.40	1508.14	1525.72	1504.64	1576.79	1480.11
Row 17	1497.04	1450.92	1408.89	1443.28	1375.03	1472.96	1476.89	1366.60	1480.87	1470.08	1416.80	1460.16
Row 18	1477.52	1471.79	1468.27	1625.34	1498.69	1463.63	1387.00	1386.97	1452.72	1461.38	1477.47	1474.63
Row 19	1402.11	1465.63	1435.94	1479.02	1329.09	1534.54	1534.17	1536.10	1356.79	1491.50	1541.34	1529.81
Row 20	1584.77	1415.80	1333.36	1310.02	1330.25	1415.41	1415.73	1394.41	1382.21	1360.73	1358.10	1400.93
Row 21	1410.33	1443.01	1511.20	1482.37	1514.27	1417.02	1546.45	1494.89	1491.74	1452.99	1481.09	1465.77
Row 22	1435.64	1364.10	1418.41	1346.68	1350.27	1348.12	1300.06	1321.32	1341.12	1357.51	1382.86	1378.59
Row 23	1393.22	1473.01	1448.93	1409.54	1419.91	1376.08	1378.48	1469.89	1336.92	1493.57	1558.83	1452.99
Row 24	1313.07	1333.00	1379.67	1400.59	1352.44	1338.55	1371.09	1295.84	1381.06	1328.42	1336.64	1375.40
Row 25	1376.79	1522.73	1494.19	1548.25	1541.94	1445.24	1587.79	1503.14	1432.77	1374.92	1349.68	1352.43
Row 26	1426.62	1331.56	1360.24	1439.86	1427.94	1385.52	1336.21	1328.25	1368.54	1396.20	1472.80	1400.20
Row 27	1604.48	1604.31	1598.03	1600.18	1424.27	1601.25	1373.85	1472.32	1497.74	1444.60	1399.25	1483.69
Row 28	1458.87	1439.29	1539.06	1478.32	1488.53	1470.38	1461.30	1410.25	1401.05	1486.69	1385.91	1398.63
Row 29	1364.06	1356.61	1368.46	1468.09	1441.80	1522.03	1406.38	1356.27	1433.70	1366.03	1472.35	1349.36
Row 30	1413.22	1400.29	1423.90	1488.75	1299.70	1462.44	1360.58	1559.01	1457.97	1433.12	1398.02	1362.52
Row 31	1396.36	1439.69	1364.06	1456.02	1450.11	1456.18	1310.99	1345.87	1439.17	1426.15	1353.98	1356.15

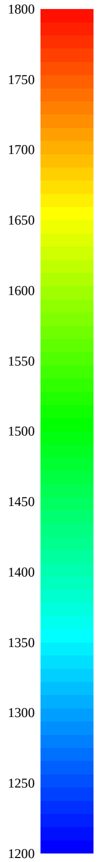

Figure 5.6: ENC noise from noise occupancy scan for barrel 3 warm run, in electrons

	-6	-5	-4	-3	-2	-1	+1	+2	+3	+4	+5	+6
Row 0	1412.49	1498.19	1513.98	1466.86	1505.47	1459.34	1488.87	1357.10	1592.27	1488.49	1479.43	1480.14
Row 1	1415.83	1519.92	1443.36	1434.83	1476.91	1427.71	1494.27	1492.98	1455.59	1461.95	1457.44	1484.36
Row 2	1418.95	1532.70	1496.97	1551.38	1533.38	1491.18	1487.85	1443.13	1495.36	1478.07	1525.94	1501.24
Row 3	1528.25	1491.48	1430.23	1451.13	1486.04	1418.72	1510.10	1444.90	1443.94	1469.59	1510.25	1534.24
Row 4	1519.79	1447.72	1476.65	1500.48	1504.43	1409.55	1464.68	1428.75	1441.42	1470.12	1477.34	1401.45
Row 5	1490.95	1435.48	1487.61	1462.54	1500.25	1453.80	1486.94	1344.43	1466.89	1444.20	1432.97	1498.63
Row 6	1468.93	1391.34	1483.97	1374.99	1517.42	1454.78	1504.27	1493.71	1438.91	1452.87	1526.82	1478.50
Row 7	1403.37	1469.12	1380.04	1449.43	1496.00	1482.81	1467.69	1445.45	1498.90	1504.53	1553.50	1515.19
Row 8	1538.40	1433.22	1501.30	1532.10	1537.65	1547.51	1490.92	1467.90	1453.60	1473.51	1417.74	1472.31
Row 9	1546.29	1552.30	1587.59	1608.19	1520.53	1449.03	1569.10	1541.05	1477.70	1472.93	1513.75	1531.75
Row 10	1504.46	1525.57	1524.19	1534.33	1498.51	1514.93	1485.70	1486.72	1401.44	1423.86	1577.31	1591.27
Row 11	1453.90	1458.38	1429.26	1457.43	1464.05	1476.51	1550.00	1486.08	1510.00	1464.31	1454.88	1513.03
Row 12	1461.57	1538.52	1484.47	1535.08	1496.14	1518.11	1548.30	1459.59	1467.91	1465.92	1552.79	1548.52
Row 13	1527.42	1553.87	1421.63	1461.14	1450.44	1468.21	1425.87	1441.43	1416.02	1520.05	1445.71	1457.65
Row 14	1355.38	1478.28	1509.09	1506.72	1528.32	1495.78	1482.68	1453.69	1452.49	1432.40	1492.82	1424.04
Row 15	1525.97	1542.27	1453.01	1564.50	1426.78	1465.18	1450.60	1451.85	1445.09	1457.48	1456.32	1452.86
Row 16	1403.70	1415.79	1378.75	1445.39	1478.14	1447.64	1393.80	1447.36	1441.44	1416.53	1490.77	1400.45
Row 17	1446.83	1400.30	1392.30	1429.75	1428.84	1439.19	1480.01	1486.43	1499.15	1525.63	1475.06	1377.60
Row 18	1459.33	1414.20	1440.12	1479.83	1484.76	1437.04	1486.92	1524.63	1533.29	1560.44	1432.48	1444.34
Row 19	1380.14	1451.24	1559.54	1495.72	1458.76	1554.98	1544.19	1430.37	1461.38	1414.42	1484.58	1468.41
Row 20	1529.77	1505.88	1404.10	1418.63	1308.46	1417.81	1385.24	1504.38	1564.08	1329.06	1519.56	1468.72
Row 21	1392.24	1440.97	1520.32	1403.14	1561.24	1386.12	1532.81	1504.60	1496.55	1506.86	1453.52	1400.22
Row 22	1468.09	1504.83	1522.62	1452.22	1477.38	1445.21	1383.31	1392.22	1316.24	1432.57	1489.28	1471.84
Row 23	1417.12	1443.79	1549.90	1500.75	1523.75	1459.40	1407.11	1455.29	1390.50	1445.50	1460.79	1382.84
Row 24	1485.64	1369.61	1467.35	1474.01	1509.27	1561.90	1416.87	1363.79	1434.98	1476.69	1517.97	1438.07
Row 25	1426.80	1471.25	1459.21	1442.21	1438.45	1396.51	1475.00	1435.70	1498.25	1486.04	1477.21	1490.19
Row 26	1457.27	1379.14	1501.85	1594.15	1550.42	1523.68	1474.79	1357.45	1424.80	1418.48	1490.45	1400.58
Row 27	1550.85	1565.04	1580.04	1523.41	1390.66	1568.50	1502.89	1595.85	1448.99	1501.01	1475.37	1395.00
Row 28	1572.01	1544.60	1583.63	1534.16	1532.70	1564.80	1542.82	1493.20	1483.16	1569.91	1460.43	1467.09
Row 29	1538.76	1557.84	1563.18	1581.45	1480.55	1546.29	1467.59	1456.15	1509.18	1431.53	1478.97	1394.64
Row 30	1498.48	1461.35	1435.60	1462.82	1442.28	1476.56	1428.68	1468.60	1477.54	1489.01	1361.33	1403.86
Row 31	1463.37	1445.64	1441.42	1499.54	1407.09	1474.24	1375.40	1446.81	1426.80	1361.43	1429.33	1409.05

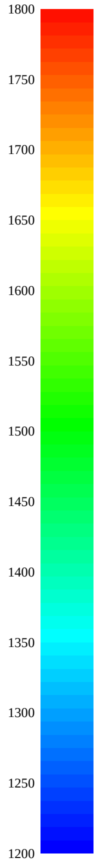

Figure 5.7: ENC noise figures from threshold scans for barrel 3 cold run, in electrons

	-6	-5	-4	-3	-2	-1	+1	+2	+3	+4	+5	+6
Row 0	1330.21	1392.32	1416.51	1422.15	1401.06	1421.71	1390.30	1250.49	1418.80	1326.18	1285.36	1390.96
Row 1	1433.39	1405.64	1328.94	1349.58	1378.65	1408.46	1286.36	1475.49	1415.56	1400.06	1309.20	1369.33
Row 2	1330.54	1408.86	1401.80	1413.89	1501.67	1367.08	1366.94	1343.88	1327.36	1373.42	1343.96	1595.24
Row 3	1502.98	1374.65	1454.62	1444.95	1492.02	1373.73	1488.71	1366.30	1345.42	1501.43	1457.64	1496.49
Row 4	1361.52	1345.10	1334.34	1368.76	1345.81	1360.75	1471.45	1413.36	1421.71	1422.45	1422.01	1360.89
Row 5	1417.15	1353.28	1399.48	1359.08	1375.98	1332.60	1374.92	1315.65	1320.79	1318.42	1288.88	1369.73
Row 6	1425.48	1311.89	1396.53	1332.69	1378.35	1424.28	1430.69	1448.98	1381.11	1379.27	1436.80	1458.15
Row 7	1321.76	1426.54	1395.95	1349.52	1381.16	1369.31	1321.03	1322.37	1255.64	1338.73	1512.13	1395.48
Row 8	1400.33	1411.08	1447.01	1470.95	1490.15	1472.35	1482.78	1448.79	1415.53	1333.57	1278.32	1336.00
Row 9	1437.18	1419.64	1451.81	1426.52	1504.08	1477.55	1423.22	1389.54	1454.07	1457.72	1448.09	1505.99
Row 10	1459.23	1450.62	1414.83	1378.01	1363.78	1376.35	1461.04	1449.37	1431.10	1431.02	1419.35	1414.58
Row 11	1454.96	1524.55	1517.44	1497.20	1334.33	1512.34	1496.51	1459.75	1414.58	1411.50	1434.01	1482.26
Row 12	1381.65	1699.46	1321.08	1395.20	1301.19	1370.63	1495.37	1424.82	1434.60	1388.02	1399.33	1359.81
Row 13	1393.41	1388.24	1341.68	1368.04	1435.85	1320.37	1390.42	1416.13	1367.14	1408.19	1433.13	1403.93
Row 14	1338.08	1332.58	1463.46	1436.07	1449.33	1465.61	1455.87	1314.60	1324.03	1336.80	1347.75	1310.03
Row 15	1379.38	1423.89	1329.78	1446.46	1401.97	1331.60	1424.46	1463.63	1390.56	1449.41	1430.95	1412.11
Row 16	1378.63	1400.38	1344.97	1445.69	1440.69	1406.05	1420.59	1375.66	1437.35	1395.83	1469.76	1384.65
Row 17	1444.14	1401.59	1375.67	1414.68	1373.51	1431.00	1432.20	1316.94	1399.78	1402.85	1347.36	1387.93
Row 18	1454.23	1442.85	1423.12	1631.76	1453.01	1418.92	1342.04	1368.11	1438.30	1413.21	1400.58	1411.80
Row 19	1358.72	1438.28	1425.07	1421.38	1309.66	1487.31	1460.81	1463.98	1311.28	1424.07	1468.51	1461.15
Row 20	1558.98	1380.64	1313.65	1282.67	1283.47	1369.75	1359.73	1347.45	1332.02	1313.17	1296.77	1337.19
Row 21	1374.29	1392.61	1484.80	1444.50	1495.97	1391.76	1500.01	1437.36	1438.57	1434.02	1406.31	1415.10
Row 22	1398.83	1354.36	1398.35	1338.12	1353.53	1334.22	1274.21	1291.72	1296.04	1321.30	1341.71	1329.32
Row 23	1361.87	1418.38	1436.24	1392.44	1451.23	1359.79	1320.11	1442.59	1294.20	1434.52	1503.16	1420.75
Row 24	1305.54	1300.55	1341.82	1380.26	1334.60	1327.59	1327.22	1260.62	1330.13	1281.58	1279.68	1326.69
Row 25	1331.43	1452.98	1418.08	1481.08	1478.17	1366.88	1505.19	1430.88	1374.55	1307.66	1291.02	1266.31
Row 26	1337.36	1253.07	1299.89	1372.76	1351.07	1311.40	1268.57	1274.43	1290.70	1320.99	1411.96	1337.17
Row 27	1559.18	1490.60	1496.10	1519.54	1370.54	1506.48	1291.28	1394.69	1544.57	1348.66	1306.49	1387.32
Row 28	1469.80	1438.61	1439.40	1436.33	1435.17	1407.75	1393.17	1317.07	1323.07	1421.07	1294.95	1320.99
Row 29	1310.14	1270.91	1299.18	1330.49	1331.45	1422.93	1358.39	1305.26	1375.17	1299.38	1422.27	1292.97
Row 30	1360.33	1336.97	1345.43	1404.77	1229.66	1376.30	1314.18	1355.81	1408.09	1385.61	1352.27	1308.70
Row 31	1314.61	1339.99	1272.34	1364.26	1343.63	1341.08	1257.28	1291.10	1389.26	1374.01	1296.90	1307.06

Figure 5.8: ENC noise from noise occupancy scan for barrel 3 cold run, in electrons

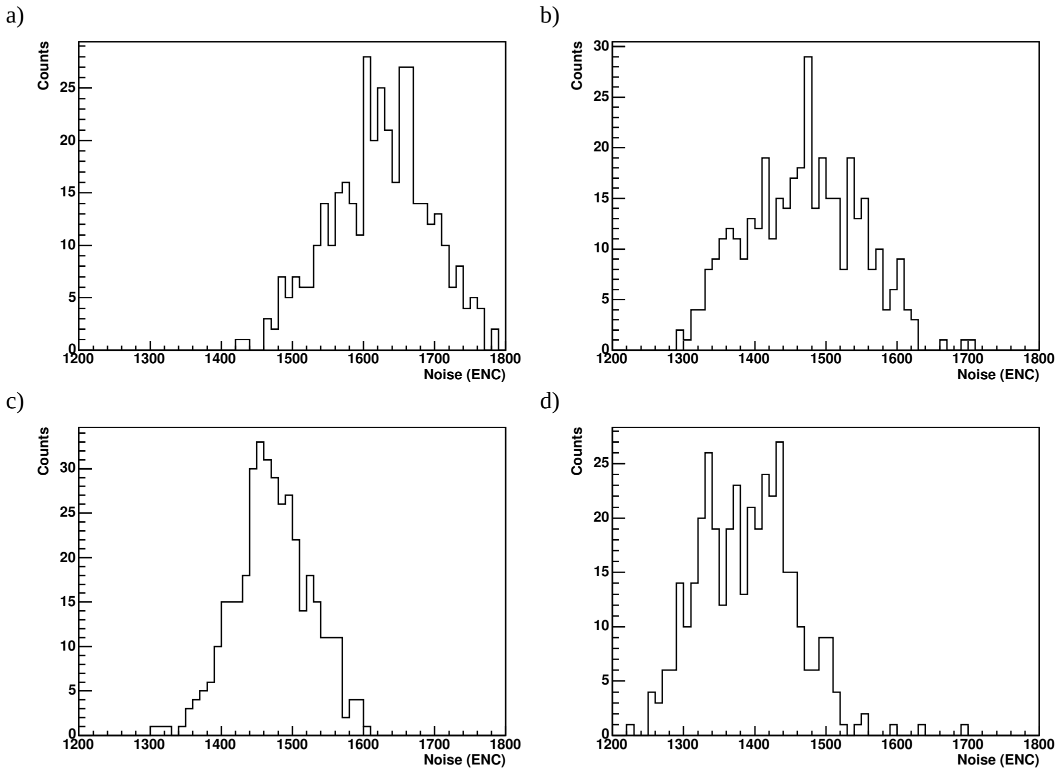

Figure 5.9: Histograms of ENC noise in electrons measured for barrel 3. The measurements are taked from the following scans: a) threshold scan warm conditions, b) noise occupancy scan warm conditions, c) threshold scan cold conditions, d) noise occupancy scan cold conditions

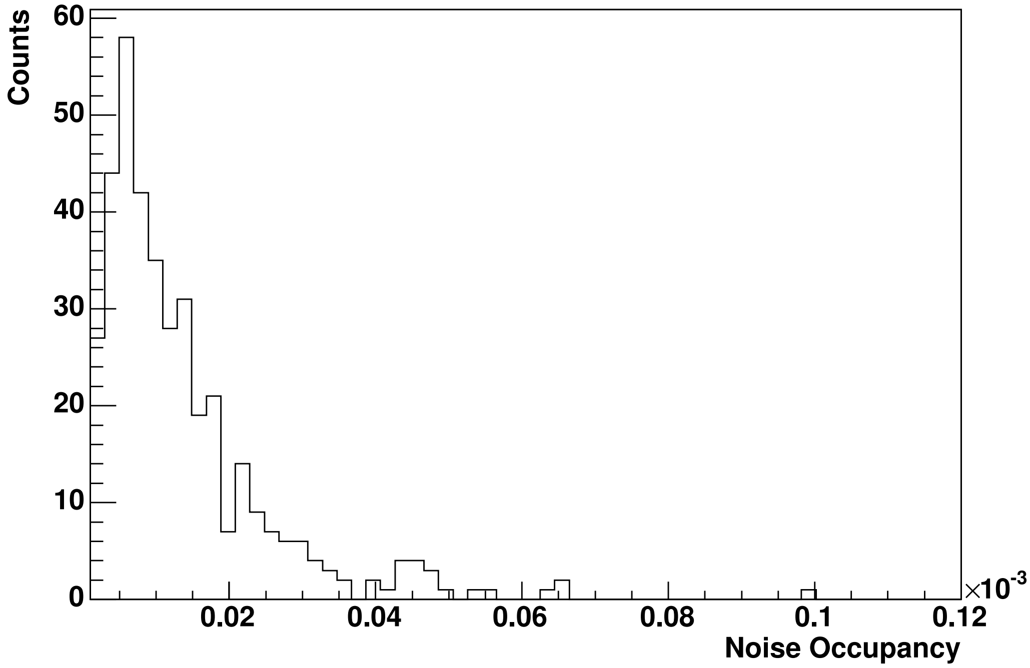

Figure 5.10: Distribution of mean noise occupancy in the cold run

(around 12°C) and the bottom (around 9°C) of the barrel.

The noise occupancy specification of 5×10^{-4} can be tested directly from the noise occupancy test. This is based on a measurement with the discriminator threshold set to 1fC. Figure 5.10 shows the distribution of the noise occupancy. All modules have a noise significantly below the requirement, allowing for the predicted increase in noise after irradiation.

5.4.2 Closely Separated Triggers

The Double Trigger Noise tests are designed particularly to look for effects of pick-up between closely spaced triggered events. This was observed at the system tests, where the noise was found to be three orders of magnitude higher for triggers separated by an interval of 132 clock cycles than for the other separations tested. A fix for this condition was found and applied to the production system. The test works by scanning the delay between two triggers while reading out and histogramming only the second event. A range of delays around 132 bunch crossings is scanned, corresponding to the pipeline depth. The second trigger is thus timed to read out the data being injected into the start of the pipeline while the data from the first trigger is being read out from the end of the pipeline.

Figure 5.11 shows the mean occupancy of each chip in a single module as a function of the delay between triggers. Most of the chips show slight deviations from the mean occupancy at particular trigger delay values, though all of these variations are much better than those observed in the system test. Chip 10 is the worst, showing a very large deviation from the mean occupancy for some trigger delays. Figure 5.12 shows the occupancies for each channel on the module over time, showing the localisation of the effect to a few channels on the module. According to an analysis of the results, involving

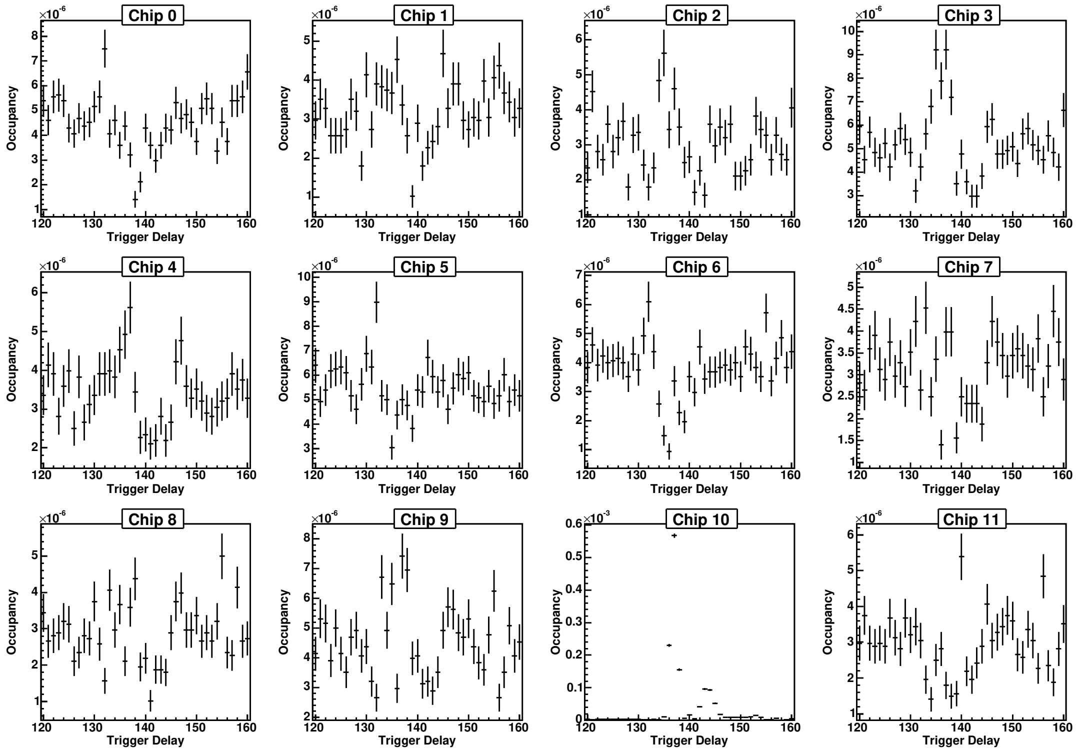

Figure 5.11: Occupancy histograms per chip scanned over trigger delay of the worst module in the Double Trigger Noise test

looking at deviations from the mean occupancy in a chip, this is the worst module tested. This is the only chip on the whole barrel with this feature, which is probably caused by a pick up of the read out of the first trigger. This is due to a quality assurance issue in the fix applied to barrel 3, which was modified for subsequent barrels.

5.4.3 Occupancy Per Event Plots

A new correlated noise detection method was introduced to the ROD histogramming code. This builds a histogram of the hit strip count for each chip, called the OPE (occupancy per event) histogram. This provides information on the structure of the events in an occupancy histogram. If the hits are evenly distributed among all the events, then each

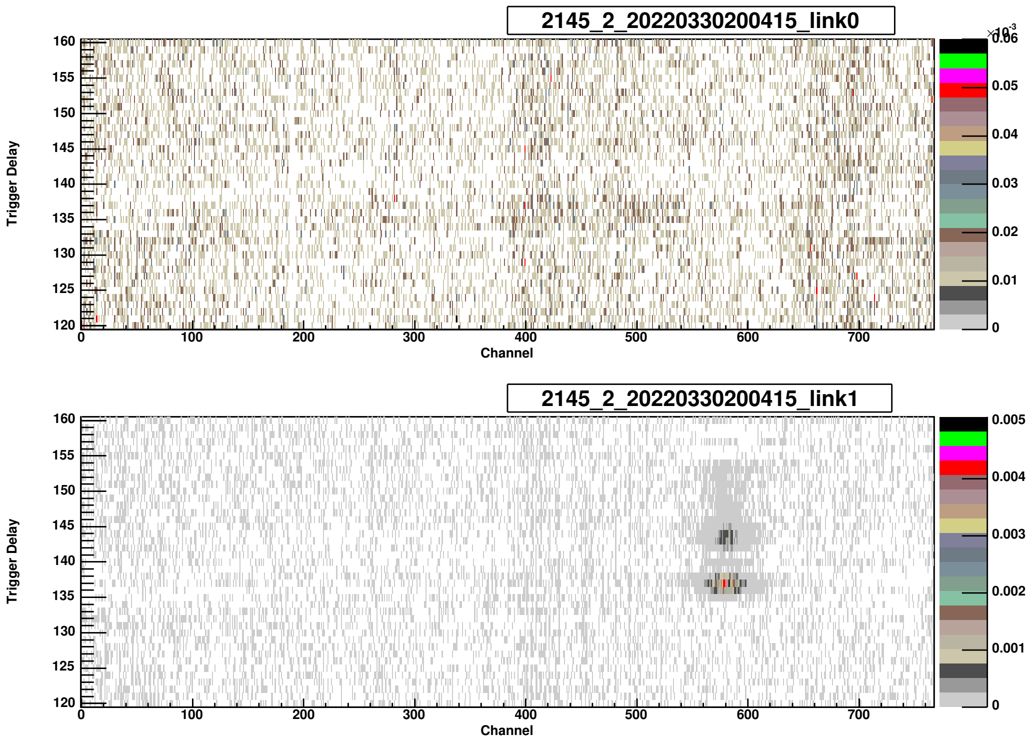

Figure 5.12: Occupancy plot of the worst module in the Double Trigger Noise test

event will have a small number of hits. Correlated noise manifests itself by inducing a large number of hits in a single event, which can otherwise be masked by the evenly distributed noise.

Figure 5.13 shows an OPE histogram recorded during a regular NMask scan, for illustration. When all the channels are masked, zero hits are recorded and this is counted in a special bin. The remaining histogram bins are four channel counts wide, so for instance when between zero and three channels are masked, a hit count of 125 - 128 is recorded, in the right hand bin.

Figure 5.14 shows a standard histogram from a noise occupancy scan with an artifact at high threshold. Figure 5.15 shows an OPE histogram revealing that the majority of these hits were in the same event.

An analysis of these histograms, based on looking for unexpected channel hit counts, was carried out on the noise occupancy data from the cold run. The previous figures are taken from the worst module found by this analysis. The fact that so few artifacts were found shows that the correlated noise is very low, even without the chokes in the power supply distribution.

5.4.4 Synchronous Trigger Test

The final new test run on barrel 3 was an occupancy measurement using triggers sent synchronously to the whole barrel by the TIM. A series of triggers were sent to the modules at a rate of 100kHz. Every 15th event was histogrammed. This is a restriction due to the event bandwidth in the router to slave DSP path in the ROD. Figures 5.16 and 5.17 show the noise values recorded from a run of synchronous triggers at a threshold of 1.0fC. All of the modules have a mean noise occupancy well below the specified 5×10^{-4}. Figure 5.18 shows one of the noisiest modules, that with the highest mean in figure 5.16,

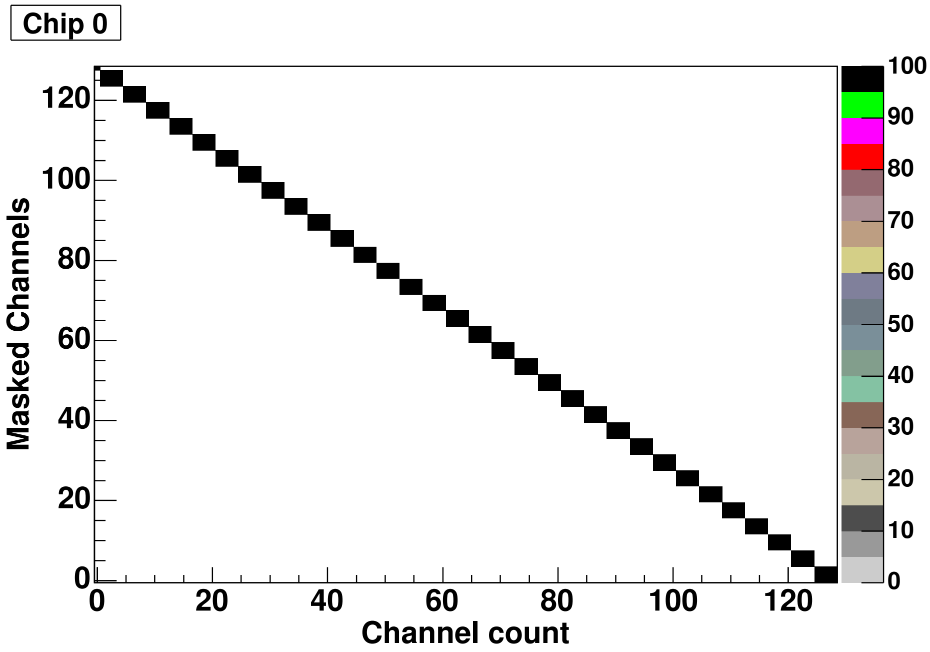

Figure 5.13: Occupancy per event histogram for a module during an NMask scan. The bins on the x axis are 4 channel counts wide except for 0

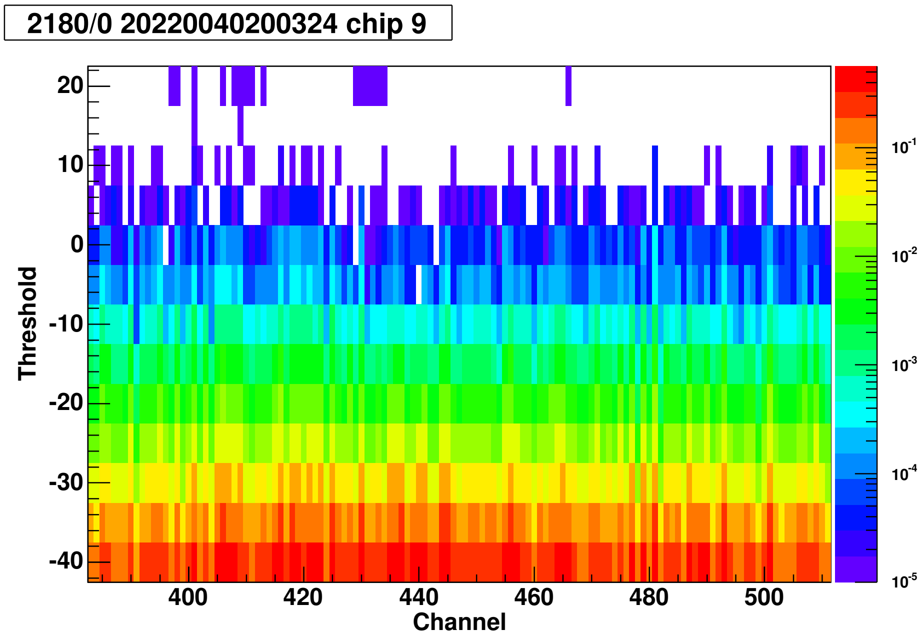

Figure 5.14: Occupancy histogram showing chip 9 of a module during a Noise Occupancy scan. This shows extra hits at threshold 20

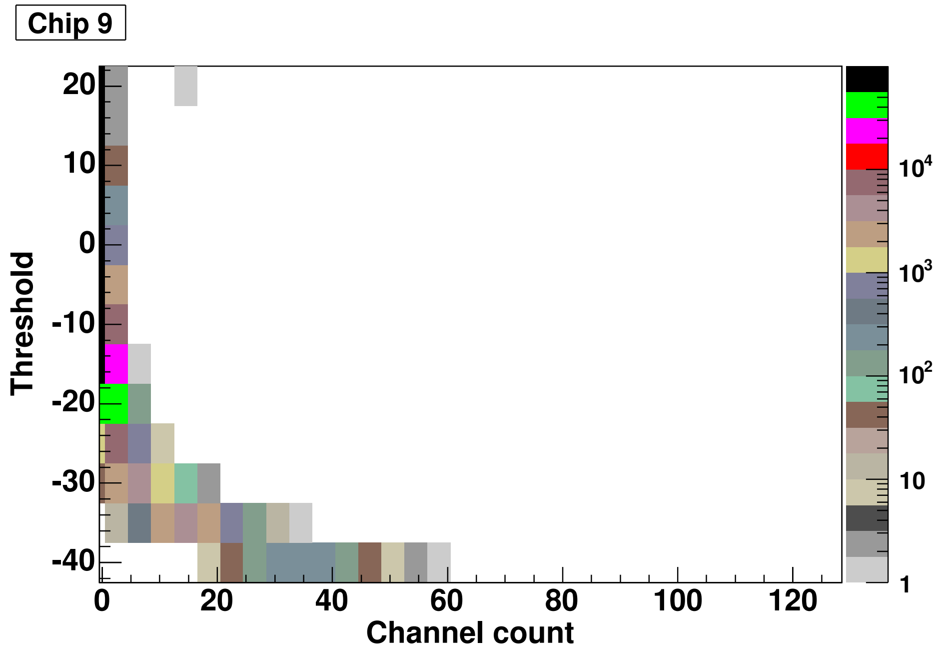

Figure 5.15: Occupancy per event histogram for a chip during a Noise Occupancy scan. At threshold 20 there are between 13 and 16 hits in one event

comparing the noise occupancy as recorded with TIM-generated synchronous triggers and that with triggers generated by individual RODs. This shows no significant difference in noise occupancy using synchronous triggers.

5.4.5 Barrel 3 Defect Summary

The remainder of the tests are substantially similar to the results obtained from individual modules tested during production. The noise results are substantially similar to these tests and are good even considering the differences in the noise environment.

Figure 5.19 shows the numbers of defective channels recorded in the barrel 3 cold run. These are as recorded after the three point gain test, which includes those channels masked off by the earlier digital tests.

A total of 5 chips are shown as missing (over 128 bad channels) due to problems

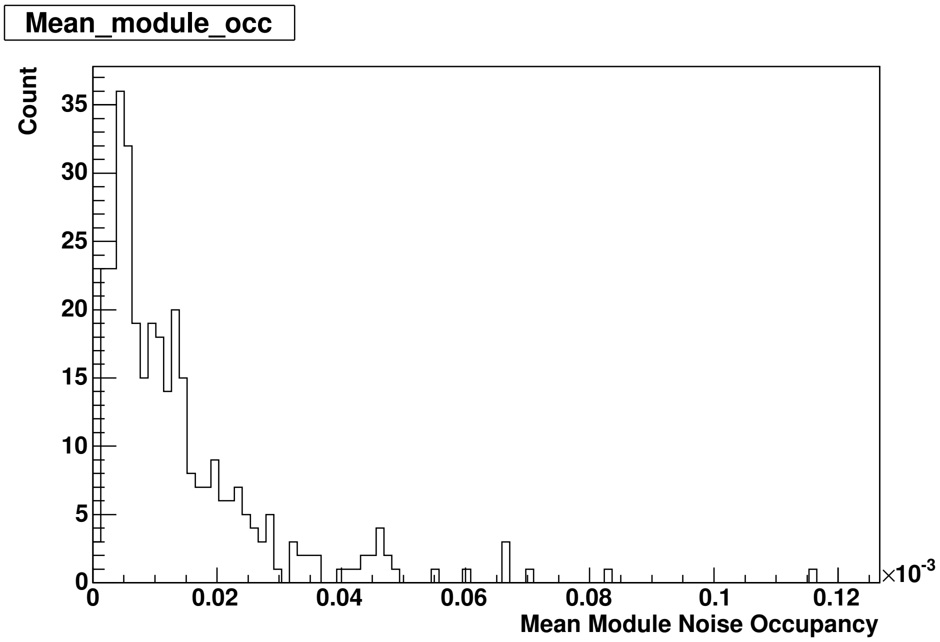

Figure 5.16: Histogram of noise values from synchronous triggers. Each entry is the mean noise occupancy across a module

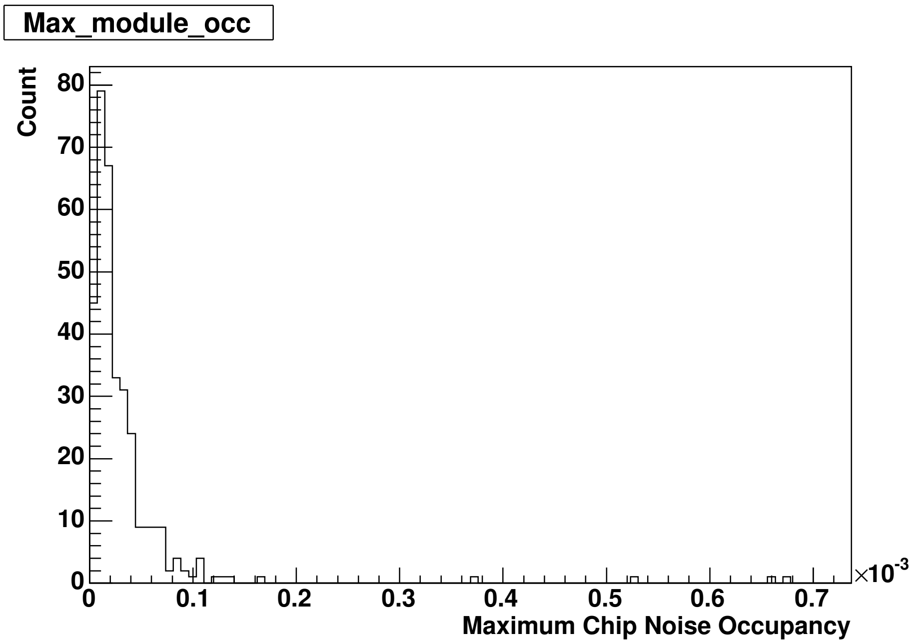

Figure 5.17: Histogram of noise values from synchronous triggers. Each entry is the maximum of the mean chip noise occupancies within a module

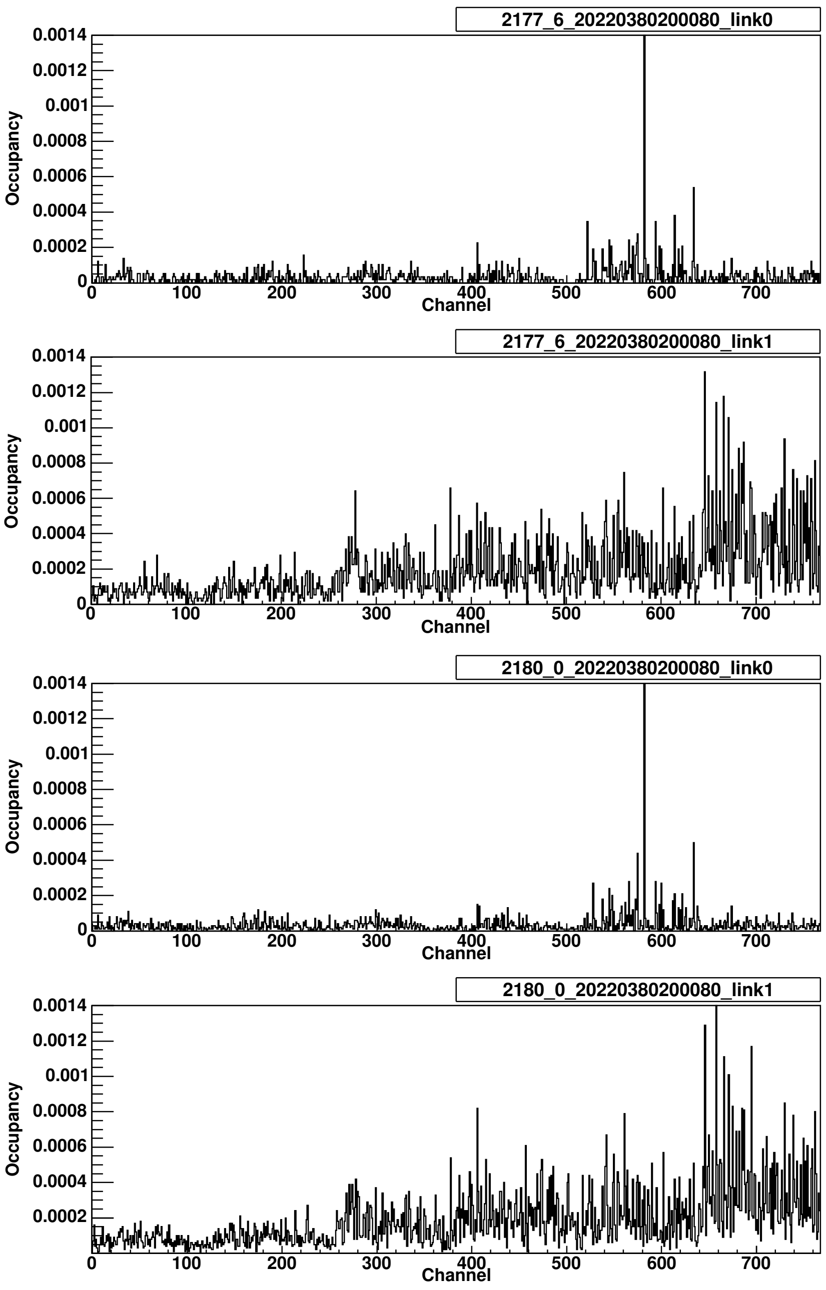

Figure 5.18: Histograms of a noisy module comparing noise occupancy with synchronous triggers (top pair) and ROD triggers (bottom pair)

	-6	-5	-4	-3	-2	-1	+1	+2	+3	+4	+5	+6
Row 0	0	2	2	2	1	4	1	1	1	0	0	1
Row 1	0	2	4	4	0	7	2	0	0	0	7	1
Row 2	9	1	11	1	0	2	11	0	12	0	10	4
Row 3	2	1	4	0	0	6	13	0	0	2	7	2
Row 4	1	4	1	0	2	0	2	0	0	1	0	1
Row 5	9	0	1	4	3	1	3	0	0	2	1	0
Row 6	0	5	10	11	9	0	3	3	3	1	7	5
Row 7	1	4	0	0	1	0	129	3	1	4	0	0
Row 8	6	7	4	2	1	2	1	1	1	0	2	3
Row 9	0	2	0	0	1	0	0	0	0	4	4	0
Row 10	3	1	1	2	1	0	1	0	0	1	0	0
Row 11	0	0	1	0	0	5	0	0	2	0	0	1
Row 12	2	6	9	1	11	10	4	0	0	2	7	2
Row 13	0	0	1	4	1	1	0	0	0	3	0	0
Row 14	4	1	1	3	0	0	0	1	0	0	0	0
Row 15	6	5	2	0	4	4	0	0	0	5	0	1
Row 16	0	0	0	4	1	0	1	0	0	0	1	2
Row 17	0	0	4	1	8	0	1	4	1	1	1	0
Row 18	11	2	1	3	2	3	0	0	1	0	0	0
Row 19	3	4	0	0	0	1	2	0	0	5	1	9
Row 20	0	2	2	1	0	1	1	5	3	1	5	1
Row 21	3	0	4	1	4	1	1	0	0	2	131	0
Row 22	2	1	0	0	3	1	2	1	2	0	1	0
Row 23	1	0	0	0	0	5	0	0	1	1	1	1
Row 24	0	3	0	0	1	0	0	1	0	5	4	2
Row 25	1	2	2	0	0	2	1	3	0	0	0	0
Row 26	3	0	0	0	1	0	1	0	0	0	7	6
Row 27	1	0	0	0	0	1	2	0	1	0	10	128
Row 28	2	0	0	128	0	0	0	1	0	0	2	1
Row 29	1	2	0	0	2	1	0	0	1	0	4	5
Row 30	1	0	1	0	1	1	0	129	2	2	1	0
Row 31	0	0	0	0	3	2	4	2	1	1	0	3

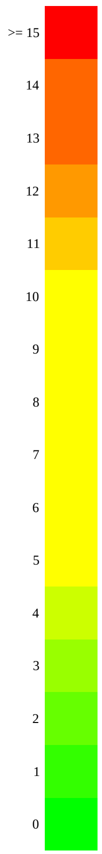

Figure 5.19: Defective channels for each module after NPt gain test of the barrel 3 cold run. See text for explanation of numbers over 128

reading out from one link of a module. Two of these have since been fixed by replacing a module in the first case, and by a better understanding of the problem in the second. The remainder are caused by mechanical problems with the optical link to one half of a module and cannot be fixed.

5.5 Conclusion

The SctRodDaq software has been used to read out SCT modules in a variety of configurations, including the whole of barrel 3, as reported here. More recently, it has also run on further barrels with similar results. Some new tests were run on all the modules together and the results were good. See news article [64].

The modules on barrel 3 passed the specified tests. In particular the noise values and the defective channel counts have been shown to be good. The cold tests showed no new defects appearing in the modules, only the expected reduction in noise due to the temperature dependent leakage current.

The software has been shown to scale from two modules to the size of a single barrel, needing only minor development in making this change. There are some improvements that can be made on the speed of histogramming. These include: more efficient histogramming code, more efficient transfer of similar module configurations and speed improvements in the transfer of histograms from the ROD to the SBC. These improvements involve modifications to the DSP firmware and were not carried out during barrel 3 assembly. Some improvements have been made for subsequent macro-assembly and all of these will be implemented for the final ATLAS configuration.

A major requirement for the full ATLAS system is to run the whole of the SCT under the same system. This requires more than one crate, and therefore SBC, and was unnecessary in the barrel 3 time frame. This has since been implemented, by Chris Lester, using a proxy IPC object to forward high level requests to the appropriate SctApi process running in each crate.

With these additions to the core software it can be used for configuration and calibration of the whole SCT in ATLAS, thus making a crucial contribution to the operation of the experiment.

Chapter 6

Pixels

MAPS or Monolithic Active Pixel Sensors is a relatively new technology that is being developed as an alternative to the more traditional CCD devices. They are currently being developed for use in physics applications, both for optical imaging in space science and particle detection in particle physics [65].

This chapter describes work carried out to forward their use in particle physics. A device testing new variations of the pixel layout was designed, laid out and tested.

6.1 Introduction

Taking the initials of MAPS in reverse order: the acronym describes a **S**ensing device made up of a 2D array of "**P**icture elements" (pixels); '**A**ctive' refers to the presence of active electronic components within each pixel; '**M**onolithic' means that the elements both for sensing and for readout are placed on the same substrate.

This can be compared to other pixel types such as hybrid pixels as found in the ATLAS pixel detector described in section 2.3.1, or CCD devices as used in the SLD experiment [66]. The key advantage of this technology over CCD is that it uses a standard

CMOS process with relatively small optimisations. This means a detector is cheaper than the equivalent CCD device. Also, standard CMOS logic can be put on the same substrate as the active detector. Another advantage over CCD techniques is that the read-out order is not predefined; the pixels can be read out so as to provide a lower resolution scan. Compared to the hybrid solution of the ATLAS pixels, the reduction in chip count makes the system cheaper and more reliable.

As with other silicon detectors (section 1.2), charged particles are detected through the collection of charges produced by ionisation. The silicon detector volume is ionised at a similar rate, producing a minimum of around 80 electron-hole pairs per μm of material traversed. This charge is collected by a potential gradient formed by a reverse biased diode (p-n junction) and read out by the electronics. In the case of MAPS the charge is collected within the pixel and the read-out electronics are on the same chip.

One of the main differences between MAPS and traditional silicon detectors is the depth of the depletion region. In the ATLAS SCT, the high resistivity silicon allows a large depletion region (285μm) to be formed, allowing the collection of a correspondingly large amount of charge. In contrast, MAPS uses a standard lower resistivity silicon and a lower bias voltage so that the depletion region is much shallower, of the order of a few microns. Charge is collected from both this region and some of the volume below it as described below.

As the active silicon is only partially depleted, only electrons ionised within the electric field are collected by drift. Electrons that are liberated from the undepleted region of the silicon diffuse into the depleted region and are also collected. This movement is bounded by the lifetime of the electron within the silicon; how long it takes for the electron to recombine with the lattice. This is therefore designed to be large by using a low doping density. In some processes, the surface silicon is deposited on top of a highly doped bulk

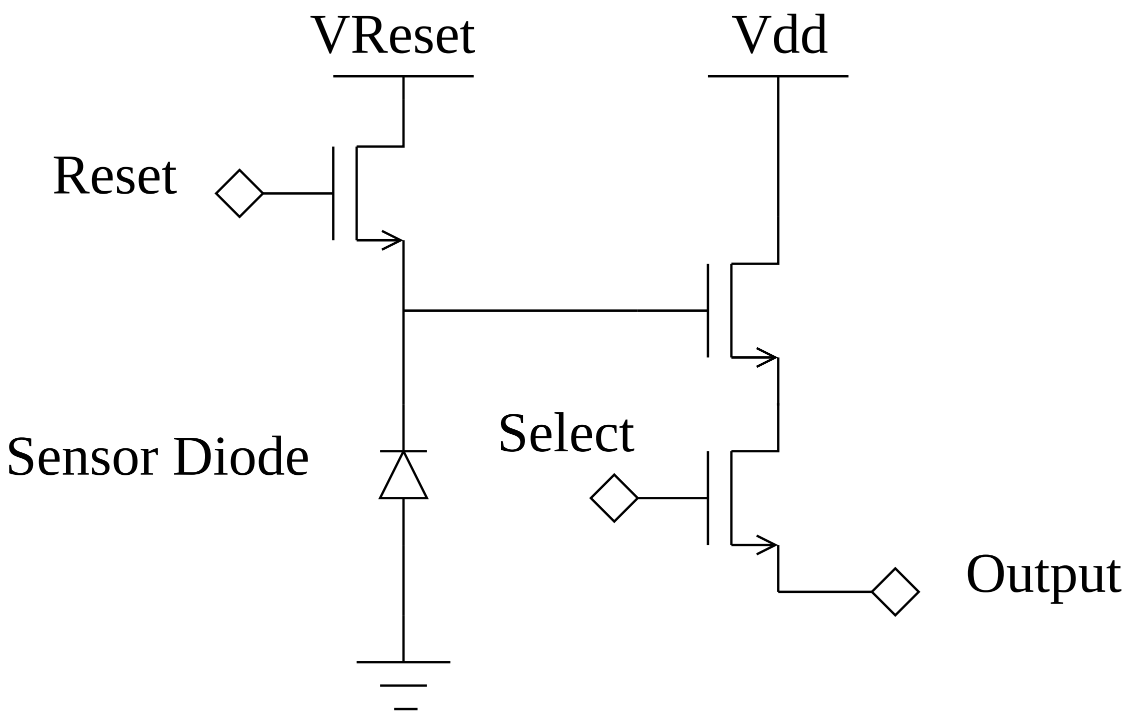

Figure 6.1: Schematic of 3MOS pixel

crystal forming an epitaxial layer. This has the effect of creating a potential barrier at the boundary between the epitaxial layer and the bulk. The barrier reflects electrons back into the epitaxial layer, containing the charge. The highly doped bulk silicon gives the electron a much shorter lifetime meaning fewer electrons are able to diffuse from the bulk across the boundary.

On the surface of each pixel in a standard MAPS chip is an arrangement of 3 NMOS transistors (figure 6.1) connected to the sensor diode. These allow the charge stored in the pixel to be read out first to the end of each column, and then transmitted off the chip. The reset transistor is used at the beginning of a readout cycle to clear the stored charge and return the potential to a known value. A second transistor provides a unity gain amplifier that mirrors the voltage of the stored charge. Finally the select transistor is used to select one pixel in a column to be read out by the end-of-column circuitry.

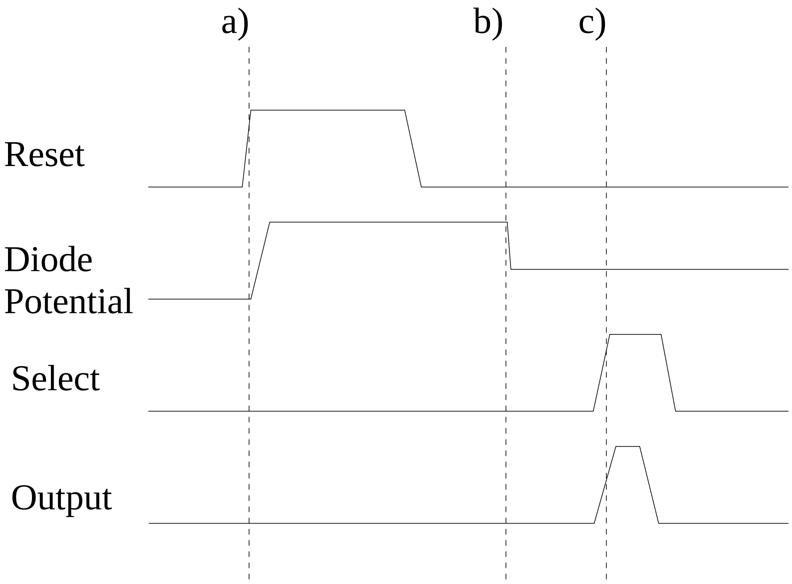

Figure 6.2: Timing diagram for 3MOS readout. At a) the reset signal is asserted. At b) a charge is injected into the anode, corresponding to the collection of electrons ionised by a particle. At c) the select signal is asserted to read out the charge from the output

Pixels are laid out in a grid of rows and columns with signals, metal tracks, common to all pixels along either a row or a column: the select and reset signals are raised in all pixels in a row simultaneously, and the output signal is common to all pixels in a single column.

At a chip level the read out proceeds as follows: a row is selected using the select line, allowing the output of all pixels in that row to mirror the stored value; the values are then read out column-by-column in sequence by the external logic. The voltage level is recorded using an analogue-digital converter. A simple picture of this is shown in figure 6.2 and the procedure is described in more detail in section 6.3.3.

6.2 Motivation for Design of HEPAPS3

The HEPAPS series of chips is part of an ongoing development of the MAPS technology at RAL. There are a series of design iterations testing differing architectures with the goal of providing sensors for both space and particle physics applications, for example see [67]. The first in the series, HEPAPS1, was a test of different diode topologies. This included testing of the correlated double sampling technique for reducing noise.

The HEPAPS2 compared other features, particularly the incorporation of a 10-cell pipeline within each pixel, called FAPS (F for flexible). This is used to enable a burst of rapid data taking so reading data off the chip can be carried out in dead time, which is particularly useful for particle physics. For example in the TESLA proposal [68] for a future linear collider, a pulse of electron bunches is contained in 950μs, and this is repeated 5 times a second. This allows many milliseconds to read out data collected in less than a millisecond.

For the current version, HEPAPS3, different variations on the basic pixel structure were tried. The primary variation was the deep n-well. This involves a special process which helped decide the production process used. This and other changes are described in more detail below.

6.2.1 Deep n-well Variation

In the standard APS device structure, the p-n junction is formed between a small n-well implant and the p-bulk. A thin p-type layer of silicon on the surface is used to implant transistors and other devices. The small p-n junction leads to a small volume of silicon containing an electric field and the majority of ionised electrons are collected by diffusion.

An alternative to this is to place a "deep" n-well underneath the surface p-layer. Both designs are shown in figure 6.3. The p-n junction is still formed between the n-well and

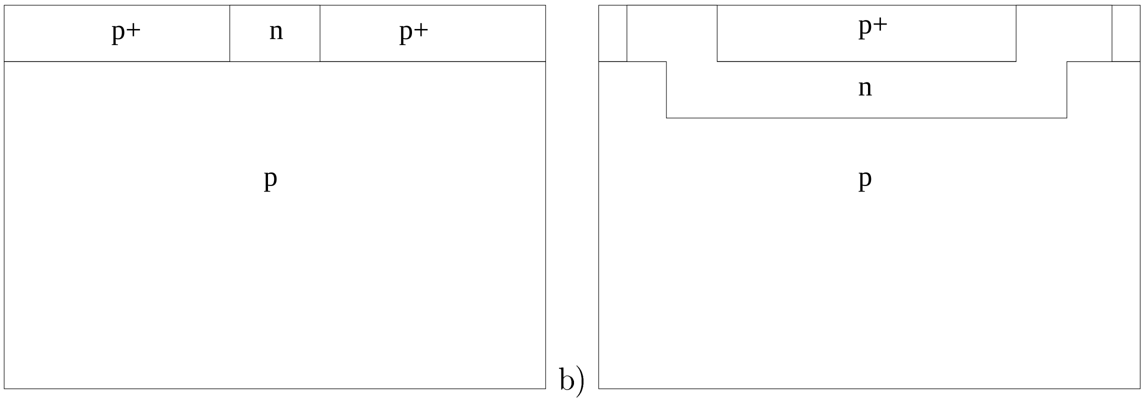

Figure 6.3: Comparison of sub-structure of pixel diodes: a) shows a normal diode, b) shows a deep n-well diode

the p-bulk, but is now formed over a larger area. The larger p-n junction can generate an electric field in a larger volume and more electrons can be collected by drift instead of diffusion. This should increase the speed of charge collection and also its efficiency. The main disadvantage is an increase in capacitance and leakage current which will add to the noise.

A simulation was carried out based on the proposed structure [69]. Some of the numerical results are shown in table 6.1. Figures 6.4 show the results graphically. This shows the expected increase in leakage current, capacitance and noise but the resultant signal-to-noise ratio is not decreased too much. The collection time is significantly reduced and the amount of collected charge is increased.

Another significant benefit of collecting electrons using drift instead of diffusion, is its radiation tolerance. The effects of radiation significantly reduce the diffusion length, so in a diffusion based detector the charge collection efficiency is significantly reduced. Simulations show the reduction is around 80% compared to around 5% for the deep n-well configuration.

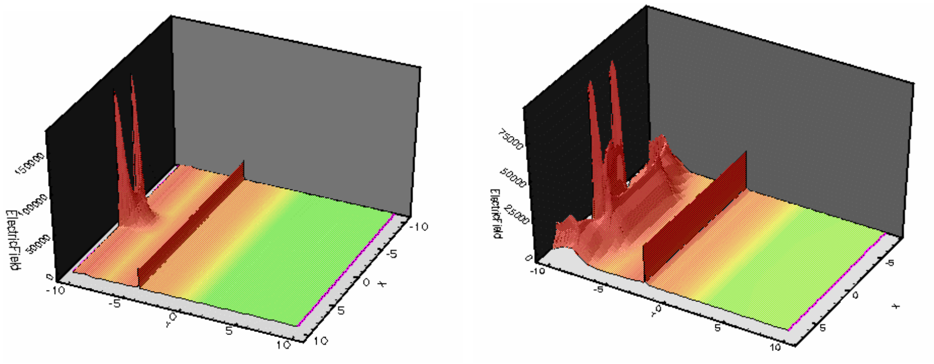

(a) Electric field

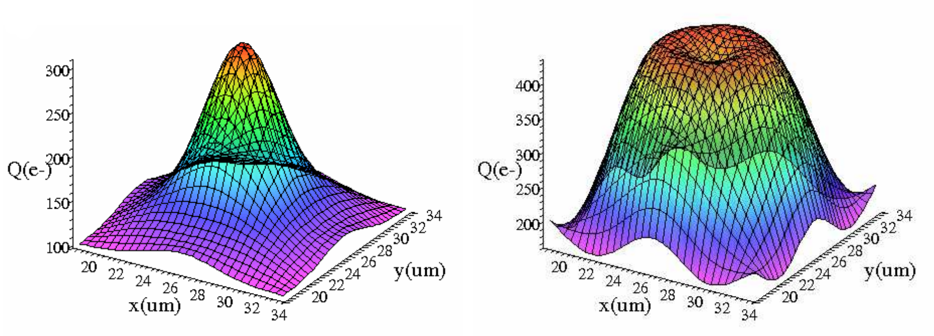

(b) Charge collected

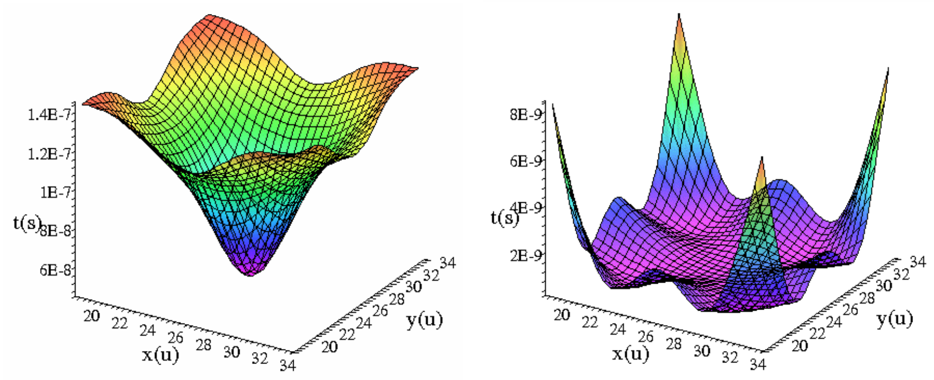

(c) Collection time

Figure 6.4: Simulation of shallow (left hand plots) and deep (right hand plots) n-well configurations

	Standard	Deep n-well
Pixel Size	15μm	15μm
Epitaxial Layer	8μm	8μm
Capacitance	1.5fF	22fF
Leakage Current	0.7fA	54fA
Average Collection Time	120ns	3ns
Average Collected charge	133e	261e
Min Collected charge	101e	168e
Max Collected charge	309e	422e
Noise RMS	25e	55e
Average S/N	5.3	4.7

Table 6.1: Simulation results for standard and deep n-well pixels

6.2.2 4MOS Variation

A 4th NMOS transistor can be added to the standard layout to provide a facility for Correlated Double Sampling (CDS). The two layouts are shown in figure 6.5, where the transistor has been added between the diode and the charge collection region and connected to the transmit signal. This allows for more control over how the pixel is read out.

Correlated double sampling combines two measurements in an attempt to lower the noise in the value read out. The first reading is made of the value following a reset. The second reading is made once the signal charge has been collected. The value of these two readings due to noise in the current flowing through the reset transistor should be correlated and thus can be cancelled.

It is thought that the reset noise is the largest noise constituent [65]. Thus subtracting the two readings will cancel the reset noise found in both, and should reduce the overall noise more than the increase in noise that comes from including the effect of reading out the uncorrelated noise twice.

Another type of noise in a pixel system is that known as fixed pattern noise. This is

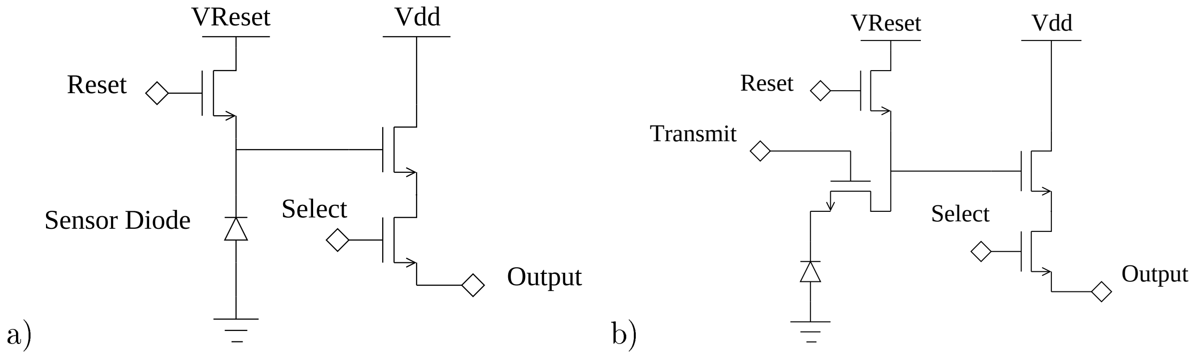

Figure 6.5: Schematic of 3 and 4 MOS pixels

caused by the intrinsic differences between each pixel. It can also be removed by cancelling out two readings as is done in the CDS procedure.

6.2.3 Enclosed Geometry Variation

Radiation tolerance is a significant design aim in physics applications, as discussed in section 1.2.4. While the deep n-well from the previous section can increase the tolerance of the detector, the MOS transistors are also affected by the high radiation dose. The CMOS process itself increases its tolerance to radiation damage at smaller feature sizes due to the reduction in the thickness of the oxide layer. The use of a thinner oxide layer means the change in the transistor threshold with a radiation dose is reduced.

Another effect of radiation on the transistors is the formation of parasitic leakage currents due to transistors which form at the edges of the intended transistors. Figure 6.6 shows the layout of both the standard NMOS transistor and the enclosed geometry, or Gate All Around (GAA), transistor which is designed to reduce the occurrence of these parasitic transistors.

An MOS transistor is formed by sandwiching a thin silicon oxide layer between a polysilicon layer which forms the gate and a p-well which forms the source and drain. The silicon is covered by a thick oxide layer which is designed to separate neighbouring

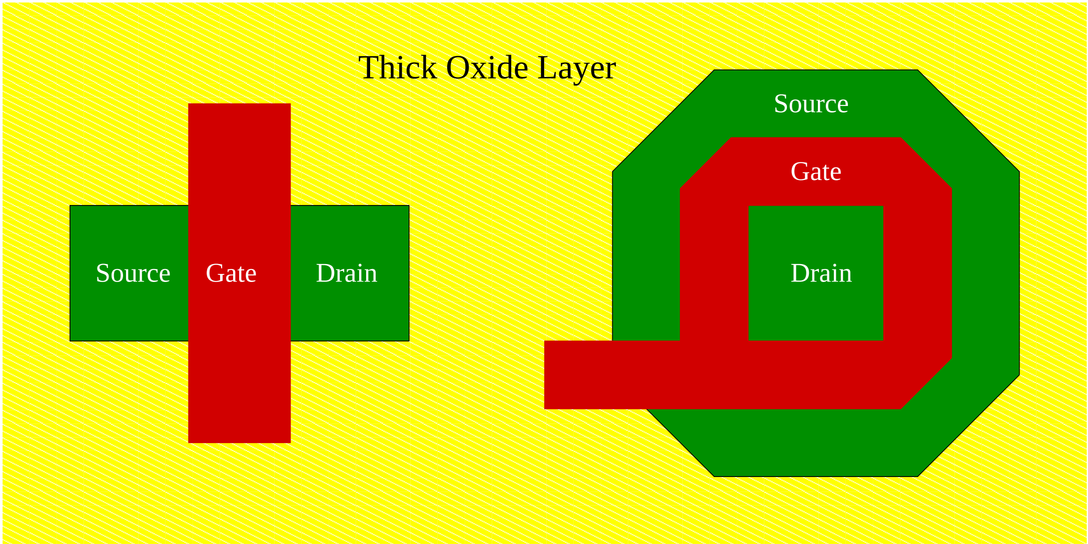

Figure 6.6: Layout of standard NMOS and GAA NMOS

transistors. After a large radiation dose, trapped holes make this oxide layer behave similarly to a much thinner one and transistors form at the junction between the polysilicon and p-doped silicon. The enclosed geometry reduces the boundaries at which these can form, but is much larger and calculation of the transistor parameters is more difficult. More information can be found in [70].

6.2.4 PMOS Variation

One of the goals of the MAPS system is to integrate more complex readout circuitry within the pixel, as has been done for instance in the ATLAS pixel system [71]. This can more efficiently be done using complementary MOS (CMOS), i.e. both NMOS and PMOS transistors. NMOS transistors are formed by embedding small n-type regions into a p-type region, which in this case is the highly doped surface layer. PMOS transistors require similar p-type regions embedded in an n-well, as shown in figure 6.7.

When using a pixel with a deep n-well, it may be possible to use this same n-well

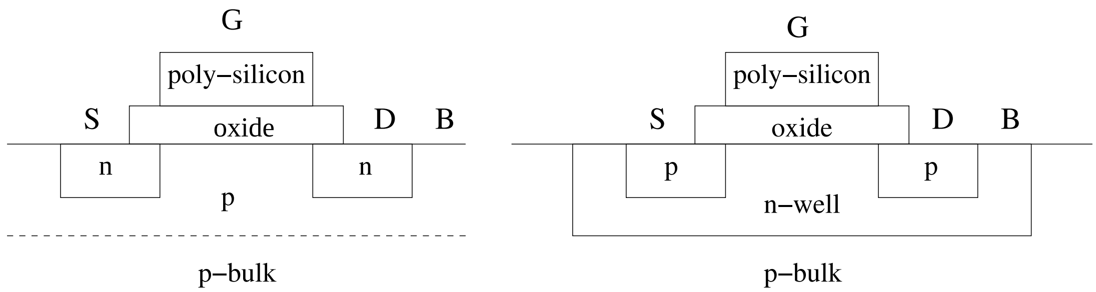

Figure 6.7: Cross-sections of standard NMOS and PMOS

to embed PMOS devices. This would reduce the surface space required, making more complex CMOS circuits possible. In order to test this, a PMOS transistor is added to the design. This directly replaces the NMOS reset transistor, reversing the polarity of the reset signal.

6.3 Design of HEPAPS3

6.3.1 Layout of Pixel Test Structures

A layout of nine pixel types in a 3×3 grid was decided upon. This enables the testing of a variety of different pixel structures. Starting from a standard 15μm square 3MOS pixel with no deep n-well, each test structure differs from another by only one variable, as shown in table 6.2. In addition to the variations described above, three of the test structures are implemented in a 30μm square pixel. This size is necessary to implement both the NWell and GAA variations in one pixel. It also provides a scaling comparison as larger sizes would allow more electronics in a pixel in future versions.

Schematics of the pixel circuits were drawn and simulated using tools from the Cadence [72] package. The next step was to lay out the components within the pixel. Tracks for the signals in both the 15 and 30 micron configurations were drawn in a separate

Size/μm	MOS	PMOS	NWell	GAA
15	3	-	-	-
15	3	-	-	yes
15	3	yes	-	-
15	3	-	yes	-
15	3	yes	yes	-
15	4	-	yes	-
30	3	-	yes	-
30	3	-	yes	yes
30	4	-	yes	yes

Table 6.2: Comparison of pixel types implemented

layer made common to all designs. The deep n-well was treated similarly. The required transistors and the diode contact were then drawn in the remaining space. The design was made to comply with the design rules for the process. This leaves little space for the transistors in the 15μm deep n-well design. Once laid out the design was successfully checked against the simulated schematic using an automatic extraction tool.

The finished pixels were then arranged into grids: the 30μm pixels arranged in 32×32 arrays and the 15μm pixels in 64×64 arrays. They were then incorporated into the chip level design.

6.3.2 On-chip Read-out Electronics

The pixels are surrounded by the electronics necessary to read out the voltage from each pixel. A schematic of this circuitry is shown in figure 6.8.

Not shown is the address decoding system which generates signals for the pixel columns and rows. The inputs to a decoding block are 8 address signals, a data signal and a reference signal, which are provided as inputs to the chip. The output is 192 signals representing the 192 rows/columns in the array. The data signal is applied to the output signal indexed by the address and the reference signal is applied to the remainder.

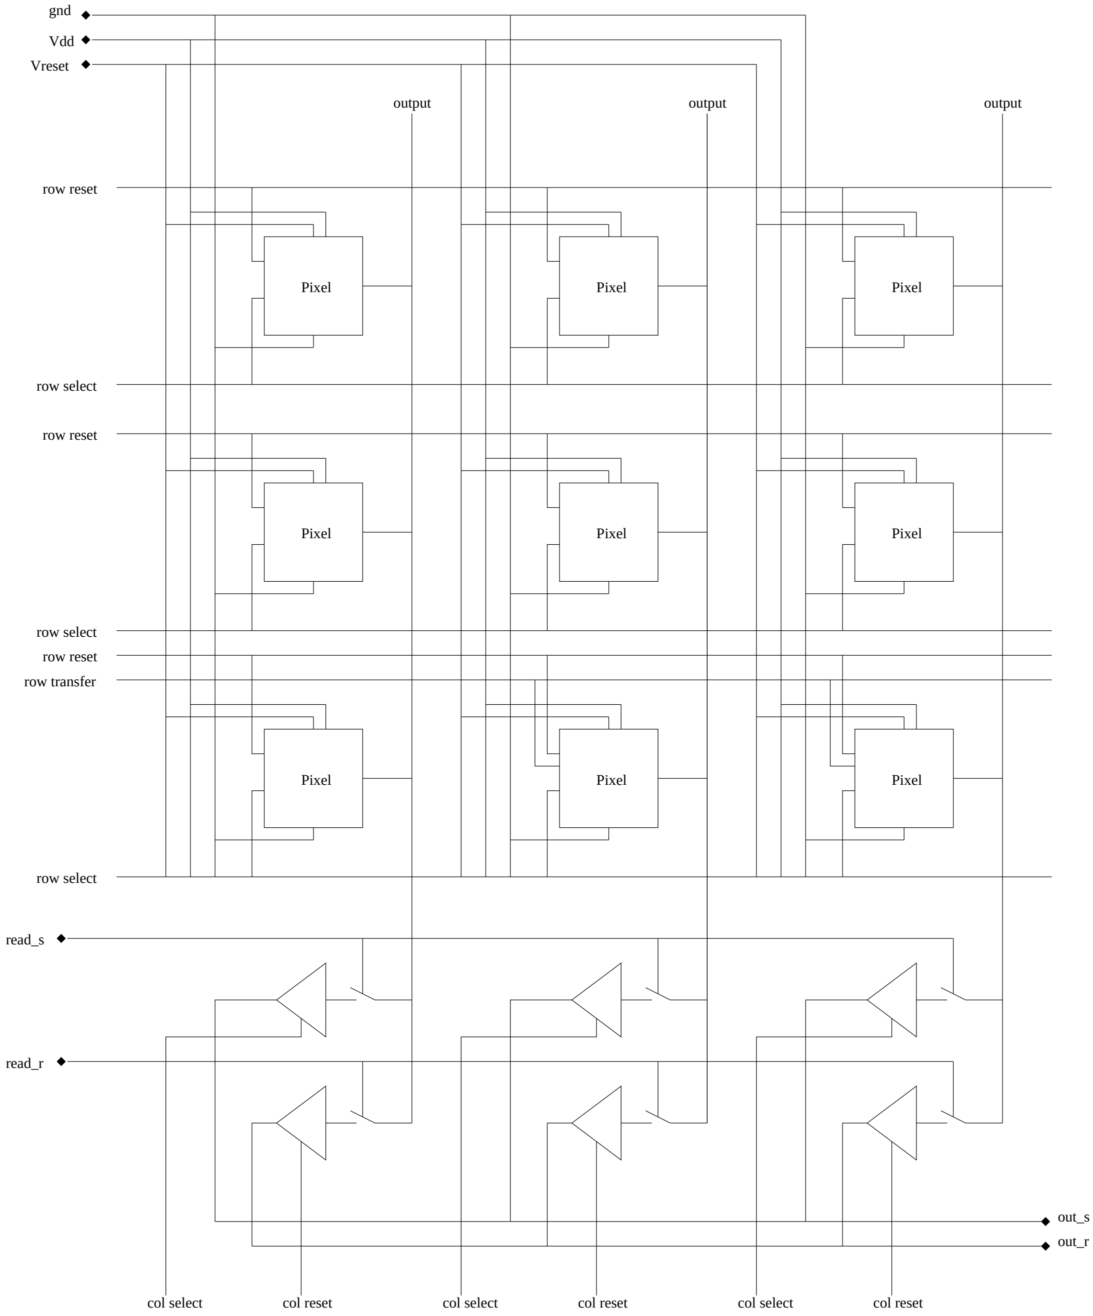

Figure 6.8: Schematic of readout electronics for HEPAPS3. Signals starting col and row are address decoded as described in the text. Two pixels have an additional "transfer" input (see section 6.2.2), based on [73]

The signals decoded to the rows are: select, transfer and reset. These are connected to the pixel inputs of the same name. They are used to pick a row of pixels and produce a voltage on the output signal. This voltage is stored in the end-of-column circuitry, in two buffers, on application of the read_s and read_r signals.

The end-of-column buffers are read out using the column decoded select and reset signals. These transfer the voltage from the buffer of the addressed column to the output pins.

The final layout is shown in figure 6.9.

6.3.3 Read-out Procedure

Figure 6.10 shows the timing of the different signals in the pixel. The diode potential is cleared by the assertion of the reset and read out by the select after some amount of charge has been collected.

At a chip level the read out proceeds as follows: a row is selected using the select line, allowing the output of all pixels in that row to mirror the stored value; the values are then read out column-by-column in sequence by the external logic. The voltage level is recorded using a analogue-digital converter.

The time between the de-assertion of the reset line and assertion of the select line, prior to readout, is called the integration time. As both these signals are asserted on a row by row basis, the integration periods of different rows do not coincide and this is called a “rolling shutter”. The reset signal is asserted soon after the select signal of the previous read-out cycle, allowing integration to occur while the other rows are read out. This system is not appropriate for particle physics applications, but greatly simplifies the on-chip electronics and read-out system during development.

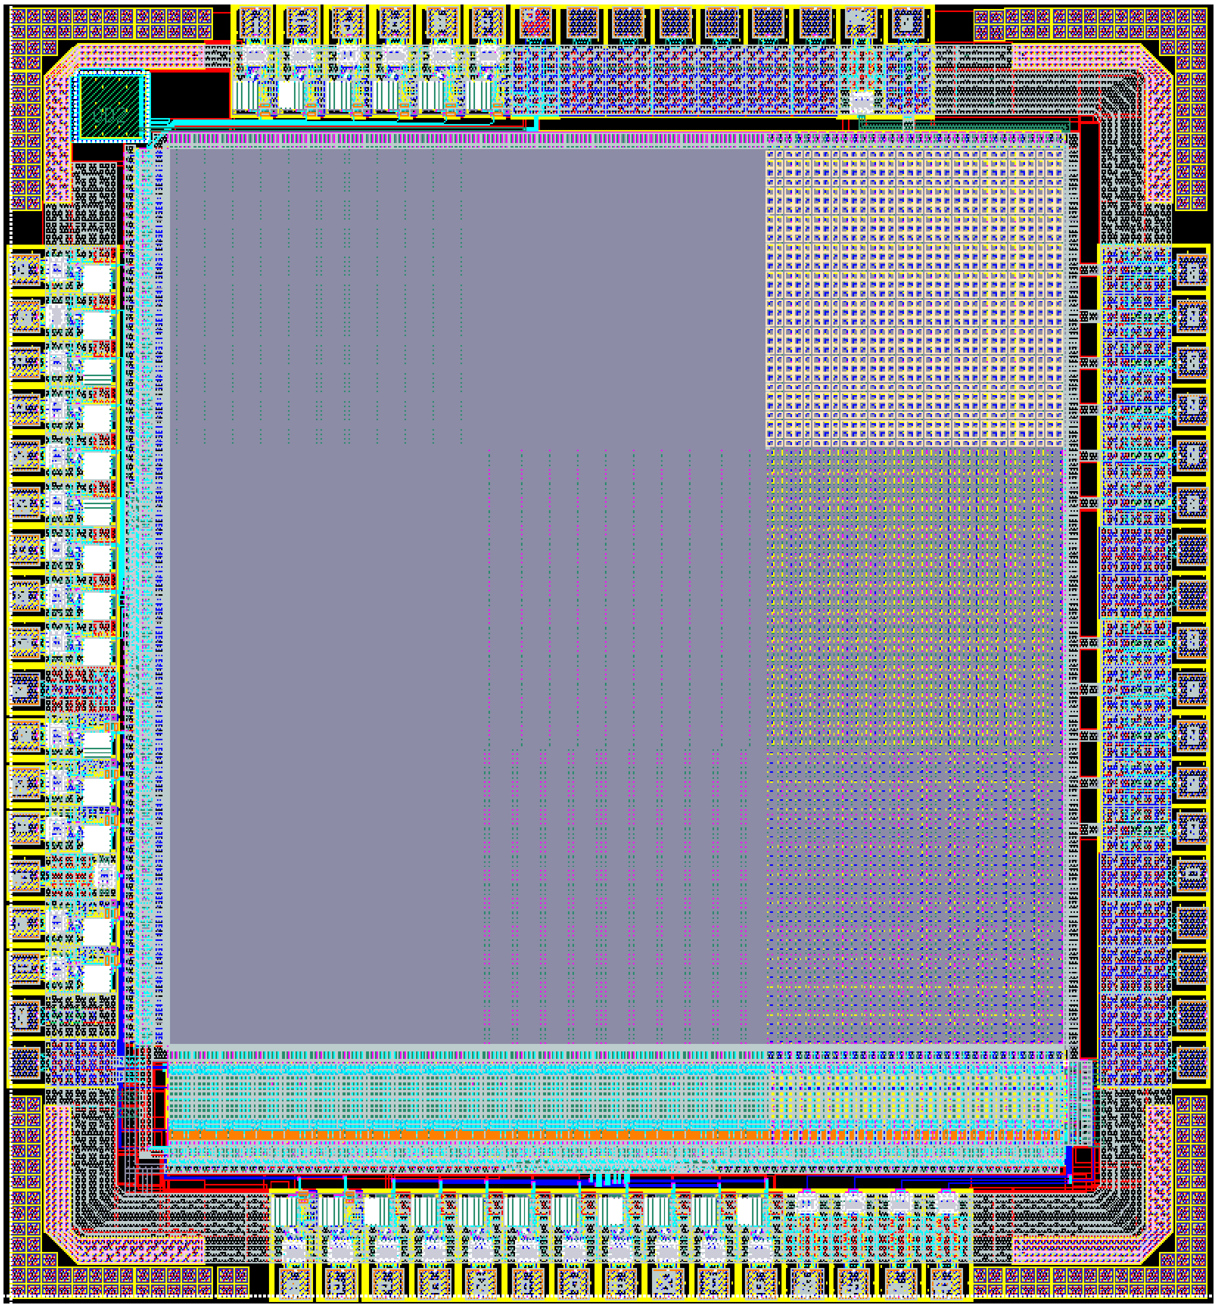

Figure 6.9: The HEPAPS3 chip

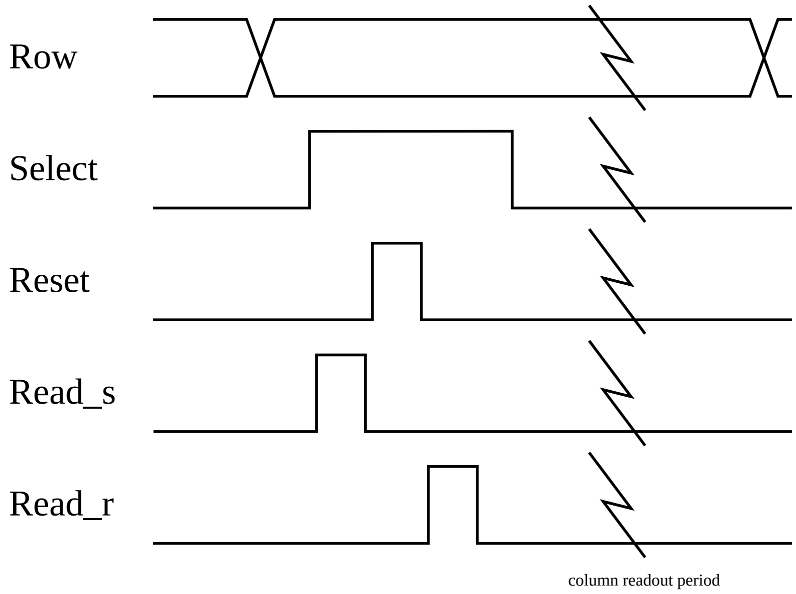

Figure 6.10: Read-out procedure for 3MOS pixels. For each row the read_s and read_r signals are asserted, loading the output value into the corresponding buffer. A similar procedure is used to read out the buffers by column

6.3.4 Silicon Process

The process used is a TSMC 0.25μm CMOS process. One of the options is the use of an epitaxial layer. This is a layer of minimally doped p-type silicon which is grown on top of the bulk (which is over-doped, p++) to around 10μm before the remainder of the device is implanted. It therefore ends up below the surface p-well implant (p+). The boundary between the epitaxial and bulk silicon forms a potential barrier due to the large difference in doping levels, which can prevent charge collection from deep within the silicon bulk, restricting it to electrons ionized in the well-defined epitaxial layer.

A process without an epitaxial layer was chosen as this was unavailable in combination with the deep n-well on a multi-project wafer. The bulk silicon has a doping equivalent to the epitaxial layer so the electron lifetime is similar. This means electrons can be collected from deep within the bulk of the silicon.

6.4 Device Testing

Once the pixels were fabricated, testing was carried out at Liverpool by Jaap Velthuis. A report on the results from the first three HEPAPS devices can be found at [74]. Results from HEPAPS3 are described below.

The first thing to test on a new chip is the signal-to-noise ratio. Once this is found to be sufficient, higher level analyses can take place.

In an analogue system, the value received and decoded by the ADCs is made up of different components. First is the signal, which is proportional to the charge deposited in the pixel; this is what needs to be extracted. Also present is a random background signal which is made up of a fixed offset for each channel and a varying background made up mainly of thermal noise. These can be characterised by a Gaussian distribution: the mean, or pedestal, represents the fixed pattern, and the standard deviation the variable noise. The signal to noise ratio is a measure of how visible the signal is when compared to the noise background. If it is too low then the signal cannot be distinguished from the noise.

6.4.1 Source Test Procedure

Four examples of APS3 devices were tested using an Ru-106 beta source. Each device is placed in a metal box with a scintillator mounted below the device, that is read out using a pair of photo-multiplier tubes. The device is read out using the procedure described in section 6.3.3. This results in a 14-bit sample for each pixel in the device. These frames are read out in sets of 100 and frames recorded on either side of a trigger signal (defined by a hit in both PMTs) are stored for further processing. The particle that caused the trigger should be found in one of these two frames, both must be read out due to the rolling shutter effect.

The first step in the analysis is to calculate and remove the pedestal. A sample of 500 non-signal frames is taken and the mean and standard deviation calculated for each channel. The sample is refined by dropping samples outside 3σ of the mean value. The resulting mean value is the pedestal and the standard deviation provides the noise value. The pedestal values are subtracted from the channel values for all subsequent events.

In order to reconstruct the position of a particle in the array, a cluster of pixels is needed. This enables charge shared with neighbouring pixels to be used in the position calculation. The first stage is to search a frame for seed pixels with values over a large threshold. The pixels immediately surrounding a seed are added to the cluster if they have values over a smaller threshold.

In this analysis, the sub-pixel particle position is not calculated, but the cluster sizes (number of pixels) and total values (summed over the pixels in the cluster) are histogrammed to provide information on how well the position could be found. For events with clusters in them, the values of pixels that are not near the clusters are histogrammed to provide a dynamic indication of the noise.

6.5 Test Results

Of the five pixel types tested (the 4MOS and PMOS pixels require modifications to the read-out system), clusters were reliably found in only three test structures, all 15μm pitch: 3MOS, 3MOS with GAA and Deep n-well 3MOS (DNW). The two 30μm devices tested didn't have any clusters. Four HEPAPS3 devices were tested, numbered 1, 4, 5 and 6 in the following text.

Table 6.3 shows counts of the clusters found in each of the three pixel types for the four devices. These were found using a seed threshold of 8 times the pixel noise and a low threshold of 2 times the noise. Device number 5 recorded substantially more clusters

Type	Device 1	Device 4	Device 5	Device 6
3MOS	83	36	594	231
DNW	21	303	574	306
GAA	28	106	757	94

Table 6.3: Cluster counts recorded for each device

than the other devices.

6.5.1 Background Noise

Figure 6.11 shows two kinds of histograms. In the left hand of each pair the pixel values from all events are histogrammed, excluding those that are close to designated clusters. This gives an indication of the noise in the device using a different procedure to the previous measurements. These values are well described by a Gaussian except for device 5 which shows an additional feature at a lower level. This suggests an additional source of noise. The extra clusters found in the analysis are likely to be caused by this extra noise source, so the results from device 5 are ignored in the following.

The right hand of each pair shows histograms of all the pixel values, each divided by the noise value for that pixel. This shows both the signal and the noise in the device. The tail to the right is the signal generated by the particles and this correlates well with cluster counts shown in the table above.

6.5.2 Cluster Size Histograms

Figure 6.12 shows the distribution of the cluster sizes in each test structure. This shows sharing of charge among a larger number of pixels than expected; to get good particle position information, the clusters should be spread over fewer than three pixels.

This increase in charge spreading is likely to be caused by the lack of an epitaxial

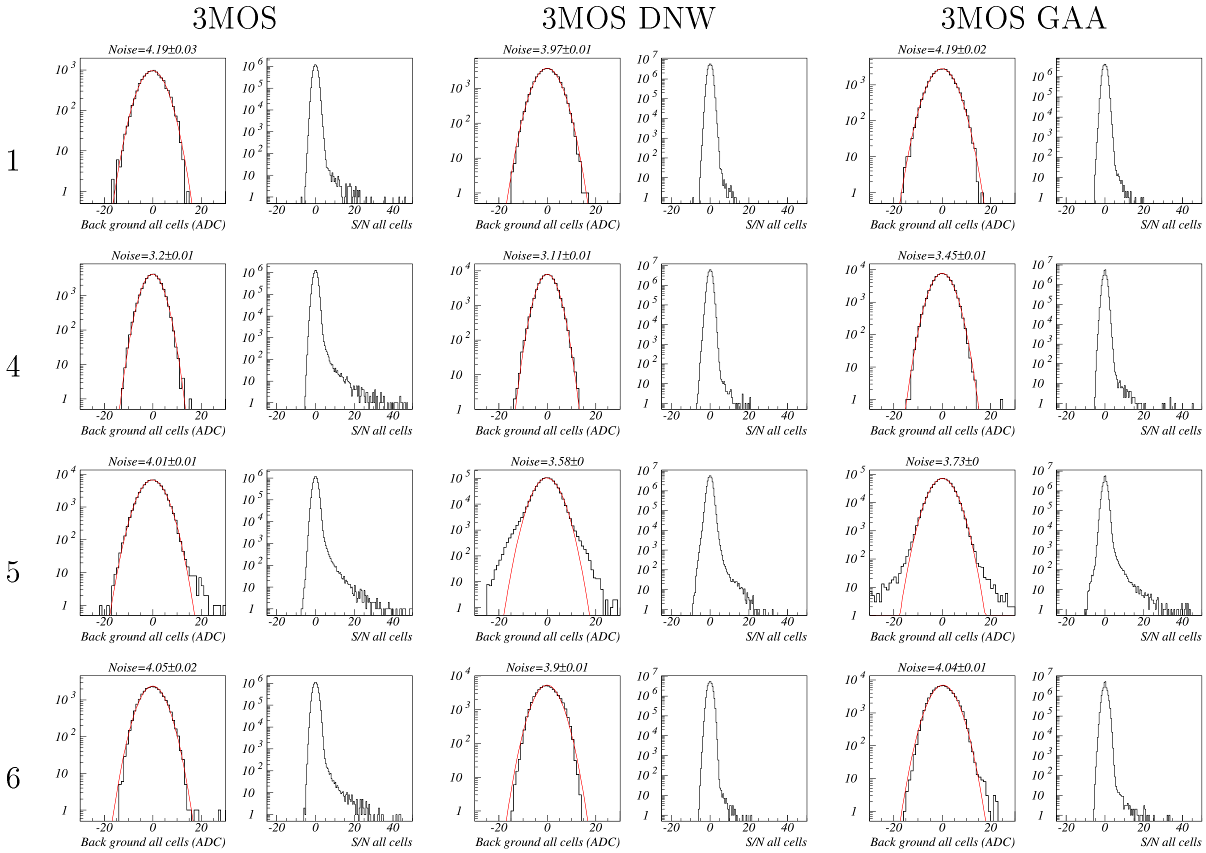

Figure 6.11: Distribution of background noise for three pixel types in the 4 devices. The left hand plot of each pair shows a histogram of pixel values in event frames excepting those pixels surrounding designated clusters. A Gaussian fit is shown in red. The right hand plot shows a histogram of all the values, each scaled by the cell noise. All plots are shown with logarithmic y axis

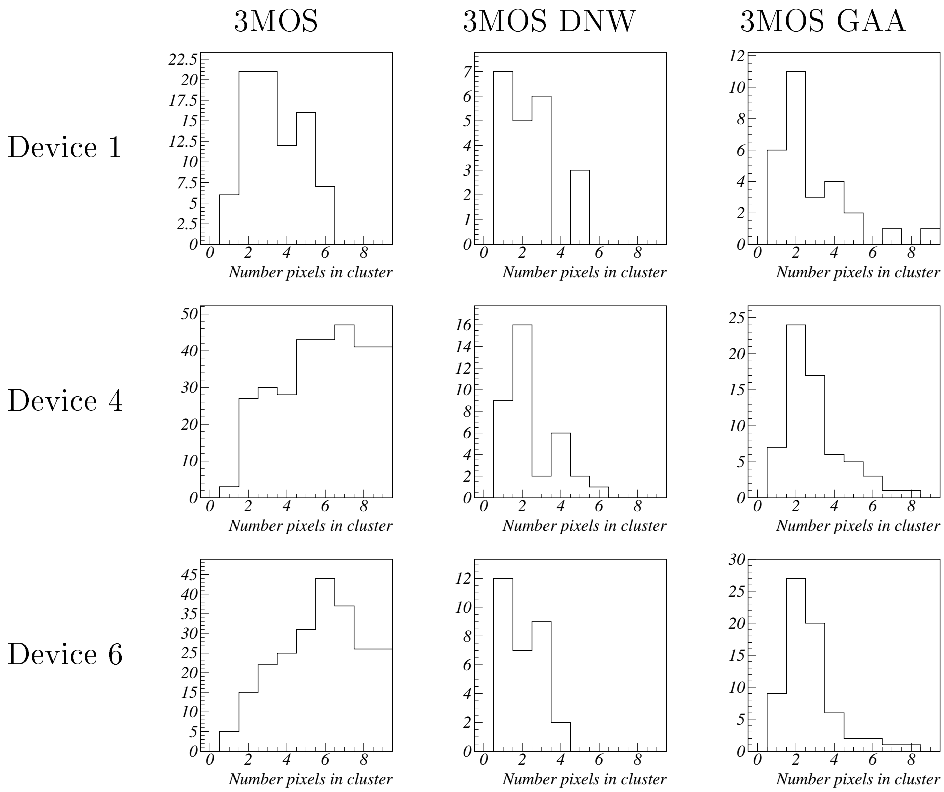

Figure 6.12: Distribution of cluster sizes for three pixel types in the 3 devices.

layer; it has not been observed in previous versions which did contain this layer. This is probably also the reason that no clusters are recorded in the 30μm devices.

There is significant variation in the cluster counts between the 3 devices, but within each device comparisons between test structures may be made. In general, the standard 3MOS pixel type shows larger and more numerous clusters, next is the GAA pixel and finally the deep n-well.

6.5.3 Signal Histograms and Signal-to-Noise

Figure 6.13 shows the distribution of the pixel values summed over each cluster. An approximated Landau function is fitted to the distributions of the 3x3 clusters and the mean value recorded as the signal for signal to noise calculations.

The signal to noise ratios for each test structure are shown in tables 6.4, 6.5 and 6.6. Across all the devices, the noise values are very similar, suggesting that the differences are caused by the geometry of the pixels. Comparing the signal to noise ratios among the pixel geometries, a general reduction from the standard 3MOS (around 25) to the GAA and deep n-well (10 to 15) is noticeable. The GAA transistors are significantly bigger than the standard ones. This has shown itself as a decrease in gain in previous revisions.

The quality of the clusters is not regarded as good enough to continue measurements using these devices. This is based on a combination of both the large clusters and the relatively low signal to noise values per pixel.

6.6 Conclusion

Only five test structures were tested due to time constraints. In three of these, a significant number of clusters were counted and quantitative results were acquired. The standard,

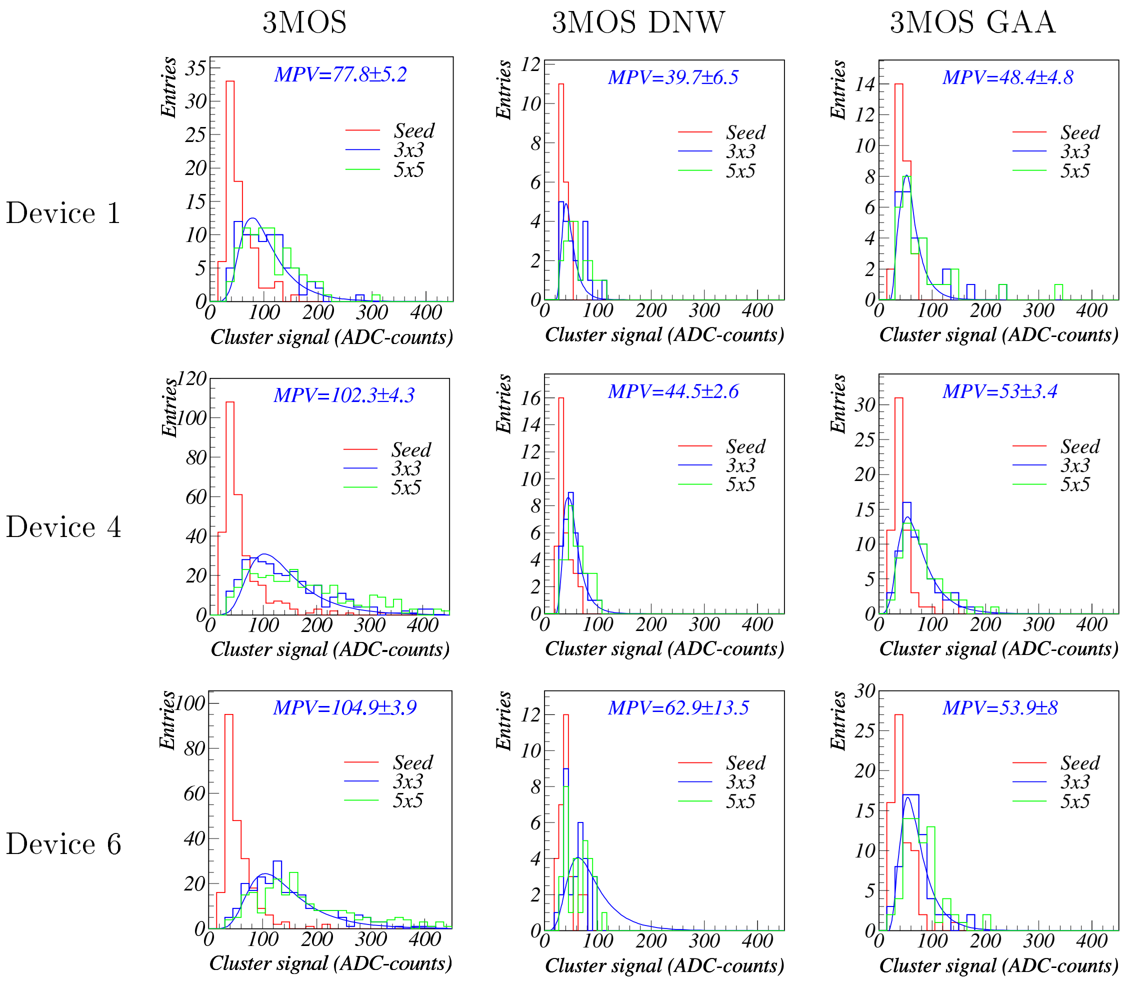

Figure 6.13: Distribution of cluster signals for the 3 devices. An approximated Landau function is fitted to the 3x3 cluster signals. The mean value is shown at the top of each figure

Device	Signal	Noise	S/N
1	78 ± 5	4.27 ± 0.02	18 ± 1
4	102 ± 4	3.53 ± 0.02	29 ± 1
6	105 ± 4	3.87 ± 0.01	27 ± 1

Table 6.4: Signal and noise results for the 15μm 3MOS devices

Device	Signal	Noise	S/N
1	39.7 ± 6.5	3.97 ± 0.01	10 ± 2
4	44.5 ± 2.6	3.11 ± 0.01	14 ± 1
6	62.9 ± 13.5	3.9 ± 0.01	16 ± 3

Table 6.5: Signal and noise results for the 15μm 3MOS Deep n-well devices

Device	Signal	Noise	S/N
1	48 ± 5	4.59 ± 0.02	11 ± 1
4	53 ± 3	3.84 ± 0.02	14 ± 1
6	54 ± 8	3.71 ± 0.01	15 ± 2

Table 6.6: Signal and noise results for the 15μm 3MOS GAA devices

reference, test structure was shown to have the highest signal to noise ratio, as expected. The GAA and deep n-well test structures showed a significant reduction. The recorded clusters are distributed over a larger number of pixels than is required to collect further data.

This is due to the lack of an epitaxial layer in the chosen process. A test simulation has shown that without the boundary at the epitaxial layer, electrons can be collected from as deep as 150μm. This option is more commonly used in optical image sensors, where photons penetrate only a few microns into the silicon.

The deep n-well test structure was successful in recording some clusters caused by interaction with a particle. Detailed measurements were not possible of the effect on collection speed, efficiency and radiation hardness. This may be possible using a process with an epitaxial layer.

Bibliography

[1] LEP Electroweak Working Group, SLD Electroweak Group, and SLD Heavy Flavour Group. A combination of preliminary electroweak measurements and constraints on the standard model. 2004. hep-ex/0412015.

[2] G. t'Hooft. Renormalization of massless Yang-Mills fields. *Nucl. Phys. B*, 33:173–199, 1971.

[3] M. Veltman. Limit on mass differences in the Weinberg model. *Nucl. Phys. B*, 123:89–99, 1977.

[4] C. H. Llewellyn Smith. High-energy behavior and gauge symmetry. *Phys. Lett.*, B46:233–236, 1973.

[5] LEP electroweak working group. http://lepewwg.web.cern.ch/LEPEWWG/.

[6] G. Flügge. Future research in high energy physics. In *1993 European School of High Energy Physics, Yellow reports. CERN 94-04*, 1994.

[7] ATLAS: Detector and physics performance technical design report. 1999. CERN-LHCC-99-14 and CERN-LHCC-99-15.

[8] R.K. Bock. *Particle Detector BriefBook.* Springer-Verlag Berlin and Heidelberg GmbH & Co. K, 1998.

[9] D. E. Groom et al. Review of particle physics. *European Physical Journal*, C15, 2000.

[10] The large hadron collider accelerator project. 1993. CERN-AC-93-03-LHC.

[11] LHC Study Group. The large hadron collider: Conceptual design. 1995. CERN-AC-95-05-LHC.

[12] S. Myers and E. Picasso. The design, construction and commissioning of the CERN large electron positron collider. *Contemp. Phys.*, 31:387–403, 1990.

[13] Design report Tevatron 1 project. FERMILAB-DESIGN-1983-01.

[14] LHC design report. `https://edms.cern.ch/nav/CERN-0000024512`.

[15] CERN accelerator chain. `http://ps-div.web.cern.ch/ps-div/PSComplex/accelerators.pdf`.

[16] W. W. Armstrong et al. ATLAS: Technical proposal for a general-purpose p p experiment at the large hadron collider at CERN. CERN-LHCC-94-43.

[17] ATLAS experiment tour. `http://atlasexperiment.org/etours_exper/etours_exper01.html`.

[18] ATLAS: Inner detector technical design report. 1997. CERN-LHCC-97-16 and CERN-LHCC-97-17.

[19] C. Gemme. The ATLAS pixel detector. *Nucl. Instrum. Meth. A*, 501:87–92, 2003.

[20] ATLAS: Calorimeter performance technical design report. 1996. CERN-LHCC-96-40.

[21] ATLAS: Muon spectrometer technical design report. 1997. CERN-LHCC-97-22.

[22] ATLAS: Detector and physics performance technical design report. Volume I. page 18, 1999. CERN-LHCC-99-14.

[23] ATLAS: High-level trigger, data acquisition and controls technical design report. ATLAS-TDR-016, CERN/LHCC/2003-022.

[24] S-LINK interface specification. https://edms.cern.ch/document/110828/4.

[25] J. Bohm, J. Stastny, et al. Power supply and power distribution system for the ATLAS silicon strip detectors. In *Stockholm 2001, Electronics for LHC Experiments*, pages 363–367, September 2001. CERN/LHCC/2001-034.

[26] SCT LV & HV power supply system project. http://www-hep2.fzu.cz/Atlas/WorkingGroups/Projects/MSGC.html.

[27] SCTHV VME model 203 user's manual. http://www.ifj.edu.pl/reports/2004/1947.pdf.

[28] ATLAS: Inner detector technical design report: The semiconductor tracker, readout electronics. pages 427–462, 1997. CERN-LHCC-97-17.

[29] Pamela Ferrari. The ATLAS semiconductor tracker system test results. *Nucl. Instrum. Meth. A*, 530:38–43, 2004.

[30] F. Campabadal et al. Beam tests of ATLAS SCT silicon strip detector modules. *Nucl. Instrum. Meth. A*, 538:384–407, 2005.

[31] Paul J. Dervan et al. The results of the irradiations of microstrip detectors for the ATLAS tracker (SCT). *Nucl. Instrum. Meth. A*, 514:163–166, 2003.

[32] D. Robinson et al. Silicon microstrip detectors for the ATLAS SCT. *Nucl. Instrum. Meth. A*, 485:84–88, 2002.

[33] G. A. Beck et al. Measurements of the thermal conductivity of pyrolitic graphite substrates for use in SCT modules. 1998. ATL-INDET-98-221.

[34] T. Kohriki et al. Development of the hybrid structure for the barrel module of the ATLAS silicon microstrip tracker. *Nuclear Science, IEEE Transactions on*, 49:3278–3283, 2002.

[35] ABCD3T specification ver 1.2. http://chipinfo.web.cern.ch/chipinfo/docs/abcd3t_spec_v1.2.pdf.

[36] M. Dentan et al. DMILL (durci mixte sur isolant logico-linaire) a mixed analog-digital radiation hard technology for high energy physics electronics. *Nuclear Physics B - Proceedings Supplements*, 32:530–534, May 1993.

[37] LVDS for SCI. Technical Report P1596.3, IEEE, 1995.

[38] P.K. Teng et al. Radiation hardness and lifetime studies of the VCSELs for the ATLAS semiconductor tracker. *Nucl. Instrum. Meth. A*, 497:294–304, 2003.

[39] VDC project specification. http://scipp.ucsc.edu/groups/atlas/elect-doc/VDC_spec.pdf.

[40] DORIC specification. http://scipp.ucsc.edu/groups/atlas/elect-doc/doric4_spec.pdf.

[41] VMEbus international trade association (VITA). http://www.vita.com.

[42] TIM manual. http://www.hep.ucl.ac.uk/∼jbl/DAQ/doc/TimModule.pdf.

[43] ROD manual. http://sctpixel.home.cern.ch/sctpixel/technical.html.

[44] V. Fadeyev et al. ASIC wafer test system for the ATLAS semiconductor tracker front-end chip. *Nuclear Science, IEEE Transactions on*, 49:1080–1085, 2002.

[45] L. Eklund , D. Ferrère, B. Gallop, and P.W. Phillips. Electrical tests of SCT hybrids and modules. ATLAS internal note ATL-COM-INDET-2003-004.

[46] SCTDAQ electrical test specifications. `http://hepwww.rl.ac.uk/atlas-sct/documents/Electrical_Tests.htm`.

[47] MuSTARD specification. `http://hepwww.rl.ac.uk/atlas-sct/mm/mustard/mustard.ps`.

[48] CLOAC specification. `http://www.hep.ucl.ac.uk/atlas/sct/cloac/Welcome.html`.

[49] SLOG specification. `http://hepwww.rl.ac.uk/atlas-sct/mm/slog/slog.pdf`.

[50] ROOT analysis framework. `http://root.cern.ch`.

[51] M. Palmer. PhD thesis, Cambridge, 2004.

[52] ATLAS online software. `http://atlas-onlsw.web.cern.ch/Atlas-onlsw/`.

[53] CORBA language. `http://www.opengroup.org`.

[54] PVSS SCADA system. `http://www.etm.at/index_e.asp?id=2&sb1=6`.

[55] XML include language. `http://www.w3.org/TR/xinclude`.

[56] libxml XML library. `http://xmlsoft.org/`.

[57] XML path language. `http://www.w3.org/TR/xpath`.

[58] S. Terada et al. Design and development of a work robot to place ATLAS SCT modules onto barrel cylinders. *Nucl. Instrum. Meth. A*, 541:144–149, 2005.

[59] H. Sandaker. The read-out and control system for the ATLAS semiconductor tracker. *Nucl. Instrum. Meth. A*, 541:96–103, 2005.

[60] A. J. Barr. Calibrating the ATLAS semiconductor tracker front end electronics. In *Italy 2004, IEEE Nuclear Science Symposium*.

[61] Kamil Sedlak. Macroassembly and performance of the ATLAS SCT barrels. In *Liverpool 2005, Position Sensitive Detectors*.

[62] B. M. Demirkoz. Construction and performance of the ATLAS semi-conductor tracker barrels. In *Hadron Collider Physics Symposium 2005*. submitted to Eur. Phys. J. C.

[63] P. Phillips, A. Barr, and B. Demirkoz. Talks given at ID week. `http://agenda.cern.ch/fullAgenda.php?ida=a042357#s17`.

[64] Device to probe limits of physics. `http://news.bbc.co.uk/1/hi/sci/tech/4035747.stm`.

[65] R. Turchetta et al. Monolithic active pixel sensors (MAPS) in a VLSI CMOS technology. *Nucl. Instrum. Meth. A*, 501:251–259, 2003.

[66] K. Abe et al. Design and performance of the SLD vertex detector, a 307 mpixel tracking system. *Nucl. Instrum. Meth. A*, 400:287, 1997.

[67] D. A. Evans et al. CMOS active pixel sensors for ionising radiation. *Nucl. Instrum. Meth. A*, 546:281–285, 2005.

[68] TESLA: Technical design report. `http://tesla.desy.de/new_pages/TDR_CD/start.html`.

[69] Enrico Giulio Villani et al. Simulation of a novel, radiation-resistant active pixel sensor in a standard 0.25 μm CMOS technology. *Nuclear Science, IEEE Transactions on*, 52, 2005.

[70] G. Anelli. PhD thesis, Grenoble, 2000.

[71] P. Fischer. Pixel electronics for the ATLAS experiment. *Nucl. Instrum. Meth. A*, 465, 2001.

[72] Cadence. `http://cadence.com`.

[73] Arwell Evans et al. Manual for HEPAPS3. Distributed internally.

[74] J. Velthius. CMOS sensor in 0.25μm technology. In *Vertex 2004*.