A setup for electrical quality assurance of ATLAS SCT barrel modules

Candidatus Scientiarum Thesis

by

Ola Kristoffer Øye

University of Bergen

Department of Physics

March 2003

Abstract

The ATLAS SemiConductor Tracker (SCT) is part of the inner detector in the ATLAS detector at the LHC complex at CERN. SCT will consist of a system of silicon strip detectors with digital readout, reading out space point information from high-energy proton-proton collisions. The detector will be assembled from over 4088 detector modules, with one type of module for the barrel part (2112 modules) and one for the end-caps (1976 modules). Before the assembly, the modules need to undergo a thorough quality assurance (QA) procedure, both in terms of physical properties and electrical performance. This thesis describes the setup of an electrical QA system for SCT barrel modules at the University of Bergen

The University of Bergen is part of the SCT group *Scandinavian Cluster,* together with the University of Oslo and the University of Uppsala. The group is responsible for building and testing 384 detector modules.

Acknowledgements

During the work on this thesis, I have had the opportunity to work with many helpful and resourceful people, and I would especially like to thank:

Professor Bjarne Stugu, supervisor, for letting me work on this project, and for advice both during the practical work and the writing process.

Bjørn Pommeresche, Lars G. Johansen and Arne Solberg in Bergen for always answering questions and helping me with practical issues.

Lars Eklund at Uppsala University/CERN for rewarding discussions and answers to numerous questions.

Werner Olsen for help with noise reduction in environmental chamber.

My fellow students at the physics department in Bergen.

Ola Kristoffer Øye
Bergen,
28/03-2003

Contents:

Chapter 1: Introduction to LHC and ATLAS

1.1 CERN

CERN (*Centre Européenne pour la Recherche Nucléaire/ European Organization for Nuclear Research*) is the world's largest particle physics centre, and is a collaboration between 20 European member states. It is situated approximately 10 km outside Geneva, on the border between France and Switzerland.

CERN was founded in 1953 by twelve founder states. A lot of questions in particle physics remained at the time unresolved, and the only way to get answers to these questions was to construct large particle accelerators for exploring the properties of matter at higher energies. CERN also had an important function as a unifying project for scientists from both sides of the recently ended World War 2.

Today, research at CERN is still mainly targeted at sub-nuclear physics, and the 6500 scientists working with CERN come from countries all over the world; USA, Japan, Canada, Russia, China, Israel and many more. This is over half of the particle physicists in the world, representing 500 universities and 80 nationalities.

CERN has housed a variety of detectors and accelerators since the start in 1953, and has always pushed the technology barrier forward. One of the first accelerators at CERN was the PS *(Proton Synchrotron)*. It accelerated and provided protons for fixed target experiments at a beam energy of 26 GeV. The PS was later used as a part of SPS (*Super Proton Synchrotron),* completed in 1981, with a centre of mass energy up to 900 GeV. This was sufficient to give experimental evidence of the postulated W and Z bosons of the weak interaction, and is considered to be one of CERN's major accomplishments.

The most recent major accelerator operated at CERN is the LEP (Large Electron Positron collider). The planning of LEP started in 1977, and used e^+e^- as collision particles. The reason for choosing these particles was to be able to do further studies of the weak interaction, especially the Z boson. Since the electron and positron are very light particles, 0.511 MeV/c^2, high beam energy was needed to reach the mass of the Z boson, 91.18 GeV/c^2. Since a particle traveling in a circle emits energy inversely proportional to the radius of the circle, a large circular orbit was needed not to lose too much energy. For this purpose, a circular tunnel, 3.8 m wide and 27 km long, was drilled 100 m underground, see figure 1.1. Like the PS was used as part of the SPS, SPS (and thus PS) was used in LEP as particle injectors.

During the years of operation, the centre of mass energy was increased from 91 to 209 GeV (referred to as LEP-II), which made it possible to search for additional particles with higher mass. No such particles were found, and this is part of the motivation for building even larger accelerators. [1]

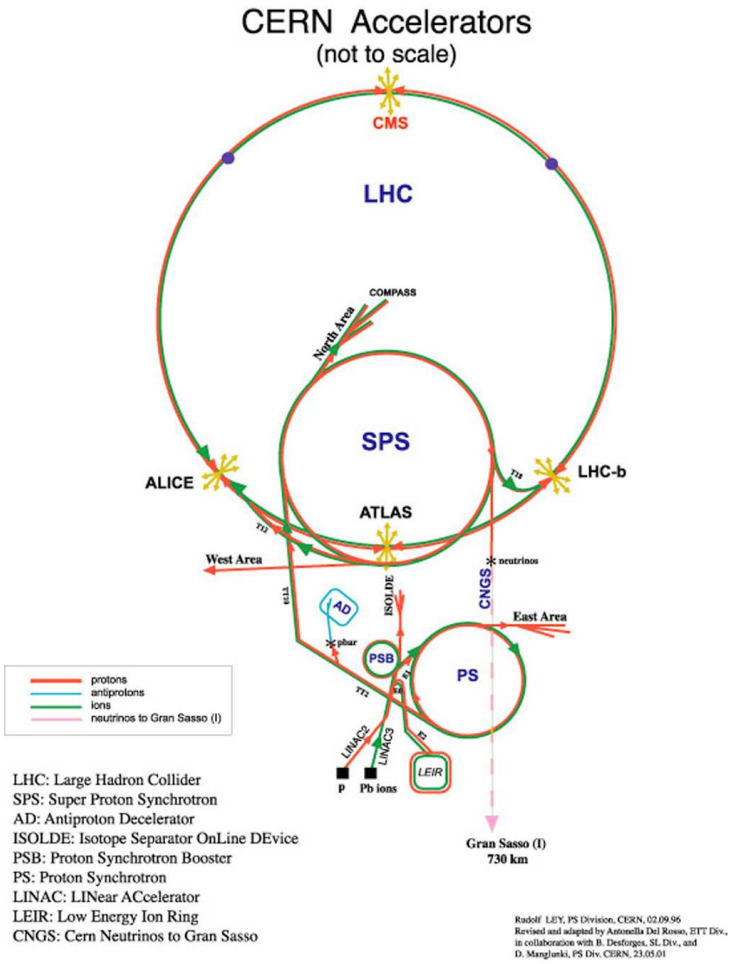

Figure 1.1: CERN accelerator- and detector complex [2]

1.2 LHC

While LEP and LEP-II delivered important data, the energy limit of 209 GeV was not sufficient to test all the predictions of the Standard Model (see chapter 2), and no new particles were found in the energy range.

The need for even higher centre of mass energies to search for new or theoretically predicted particles, is the motivation for building the Large Hadron Collider (LHC), the next large accelerator at CERN, planned for completion in 2007. It will be installed in the 27 km LEP ring, with SPS as a preaccelerator, and will be a proton-proton collider with

centre of mass energy 14 TeV, 7 TeV for each proton beam. The luminosity[1] will be 10^{34} $cm^{-2}s^{-1}$, with beam crossings 25 ns apart. Each beam will contain 2835 bunches, with $1x10^{11}$ protons in each bunch.

LHC can also operate as an ion collider, using Pb beams with a centre of mass energy up to 1250 TeV, about thirthy times higher than the highest energy reached previously in any accelerator.

The LHC project features four detectors at various points in the accelerator (see figure 1.1):

- **ATLAS**, ***A T****oroidal* ***L****arge hadron* ***A****pparatu****S***: High luminosity general purpose proton-proton detector.
- **CMS**, ***C****ompact* ***M****uon* ***S****olenoid*: High luminosity general purpose proton-proton detector, designed to provide particularly effective muon detection.
- **ALICE**, ***A L****arge* ***I****on* ***C****ollider* ***E****xperiment*: Ion-ion detector.
- **LHC-B**: Detector designed for studying B-physics.

To bend the proton beam around the ring, a strong magnetic field is needed. For this, 1296 superconducting dipole magnets operated at temperature 1.9 K, producing a field of 8.38 T will be used. A crossectional view of the ring, containing the two beampipes surrounded by the superconducting dipole magnets as well as cooling and shielding fascilities can be seen in figure 1.2.

For focusing the beam, 56 superconducting quadrupole magnets also operated at T=1.9 K is used. These are the magnetic equivalent of an optical lens, preventing the beam from spreading. The peak value of the field in the quadrupole magnets is 6.86 T.

Acceleration of the beam is done by 8 superconducting RF cavities placed in one of the straight sections of the ring. These create an accelerating field of 5 MV/m. The strong field will enable the protons to reach a speed of 9.71 km/h below that of light in vacuum, which means that a proton runs through the LHC ring in 88.924 μsec, also referred to as the LHC beam crossing rate of 11.2456 kHz.

The cooling system for the superconducting elements is based on super-fluid helium, which has very efficient heat transfer properties. [3]

[1] Luminosity is defined as $L = \frac{f \cdot n \cdot k^2}{A}$, where f = frequency of circulation , n = number of bunches, k = particles per bunch and A = cross-sectional area of bunch

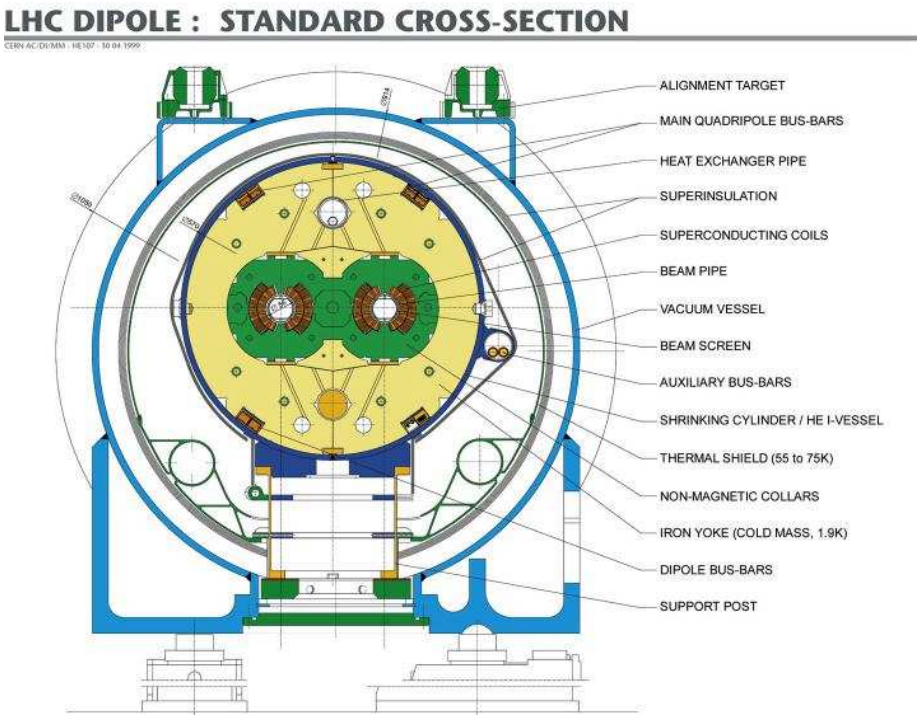

Figure 1.2: Cross-section of LHC beam pipes [2]

1.2.1. Proton-proton collisions at the LHC

The data from LEP gave a very detailed picture of the fundamental particles and their interactions. From the data it is possible to see traces of higher energy physics, and from this extrapolate to which energies is needed to reach the domain of "new physics". It is believed that this area lays around 1 TeV, so LHC must be able to deliver beams that match this energy. To achieve this, the LHC uses proton-proton collisions with a centre of mass energy 14 TeV.

In most cases, the collisions between protons will be *soft collisions;* large-distance collisions with small momentum transfers. These collisions will dominate, with an event rate of ~ 10^9 events/s, corresponding to 23 interactions creating 1000 particles at each bunch crossing, but will not be interesting. The interesting collisions will be those of single quark-quark/quark-gluon interactions, known as *hard scattering*, characterized by a large momentum transfer Q^2, giving rise to production of massive particles with high transverse momentum p_T. These will however have an event rate of only 1-10 events in 10^8 events, which means that selecting and extracting these events from all the uninteresting events (known as *pile-up*) will be a formidable challenge. [4]

Since the colliding quarks/gluons, also referred to as partons, only carry a fraction of the protons total momenta, this is the reason for the total centre of mass energy to be higher (14 TeV) than the energy needed for reaching the domain of 'new physics' (1 TeV).

The fraction of the momentum a parton carries is referred to as *x*, also called *Bjorken scaling variable*. The probability of finding a parton carrying this fraction *x* is given by a *Parton Distribution Function (PDF)*, $f(Q^2,x)$, depending both on *x* and the momentum exchange Q^2.

An example of such a function is shown in figure 1.3, at a large momentum transfer. We see that the probability of finding one parton carrying a significant part of the total momentum is small, and that the probability is shifted towards small *x*, which means that the sea-quarks carry a significant part of the momentum. For small Q^2, *x* is shifted towards ~1/3, meaning that the momentum is mainly carried by the three valence quarks.

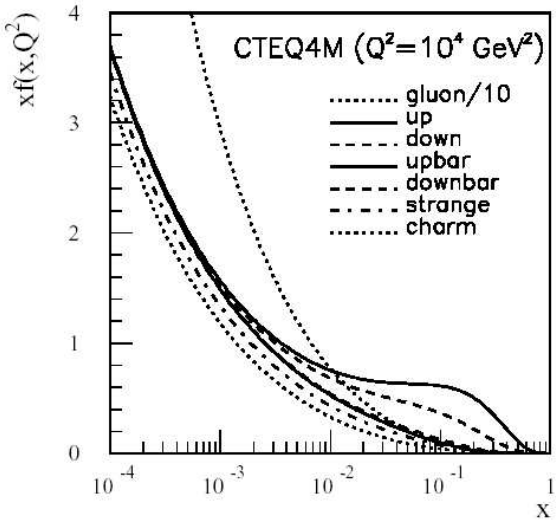

Figure 1.3: Parton distribution function of the proton for momentum transfer 100 GeV. [4]

1.3 ATLAS

ATLAS ***(A Toroidal Large hadron collider ApparatuS)*** is a general purpose proton-proton detector. It is the physically largest detector in the LHC, and is also the largest collaborative effort ever attempted in the physical sciences, with 2000 scientists from more than 150 universities and 34 countries working together on the project. Figure 1.4 shows the ATLAS detector.

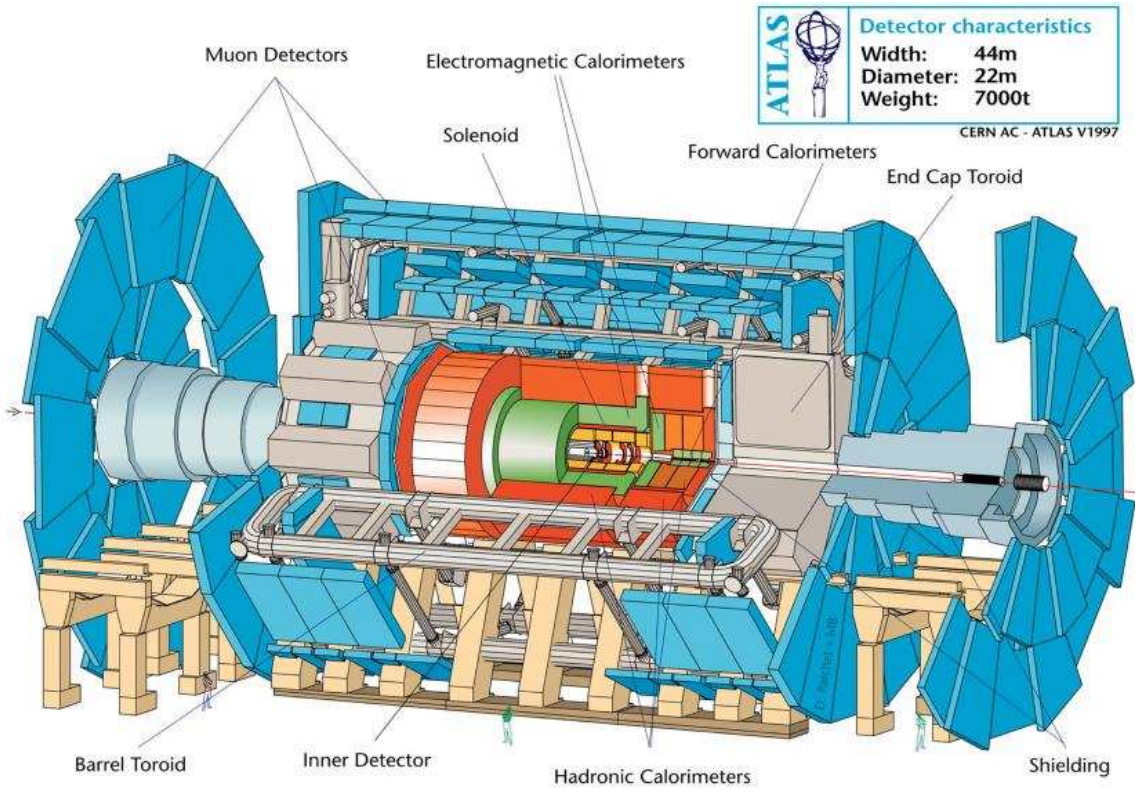

Figure 1.4: The ATLAS detector. [2]

Installation is expected to start in 2003, and the first run is scheduled to the start of 2007. Since ATLAS is a general purpose detector, it is designed to exploit the full potential of the LHC; it provides detection of leptons, hadrons, photons and neutrinos, featuring very precise tracking abilities and energy measurements.

ATLAS is divided into four major parts, each one described in the following sections:

- **Inner detector (ID)** – tracks and measures the momentum of charged particles.
- **Calorimeter** – measures the energy of electrons, photons and hadrons.
- **Magnet systems** – bending charged particles for momentum and mass measurement.
- **Muon spectrometer** – identifies and measures muons.

ATLAS will also feature a highly efficient trigger system to select the small number of interesting events from the large total number of events.

The coordinate system in ATLAS uses these axis and reference points:

- Origo placed at collision point (detector centre).
- X-axis horizontal, pointing towards the centre of the LHC-ring.
- Y-axis pointing upwards.

- Z-axis alligned with beam pipe so that it forms a right-handed coordinate system together with the X and Y axis.
- ϕ - azimuthal angle measured around the beam axis.
- θ - polar angle relative to the beam axis.
- η = - ln (tan θ/2). This is known as the *pseudo-rapidity angle*.
- r - distance from the collision point.

1.3.1 The Inner Detector (ID)

The inner detector is a tracking detector located closest to the interaction point in ATLAS. Its task is to reconstruct tracks and vertices of charged particles produced at the interaction point, as well as supplying the signature for short lived particles. The ID is situated in a solenoidal superconducting magnet system with a field of 2 T directed along the beam axis. This will cause the charged particles to bend, and based on the bending, the momenta of the particles can be determined with great precision.

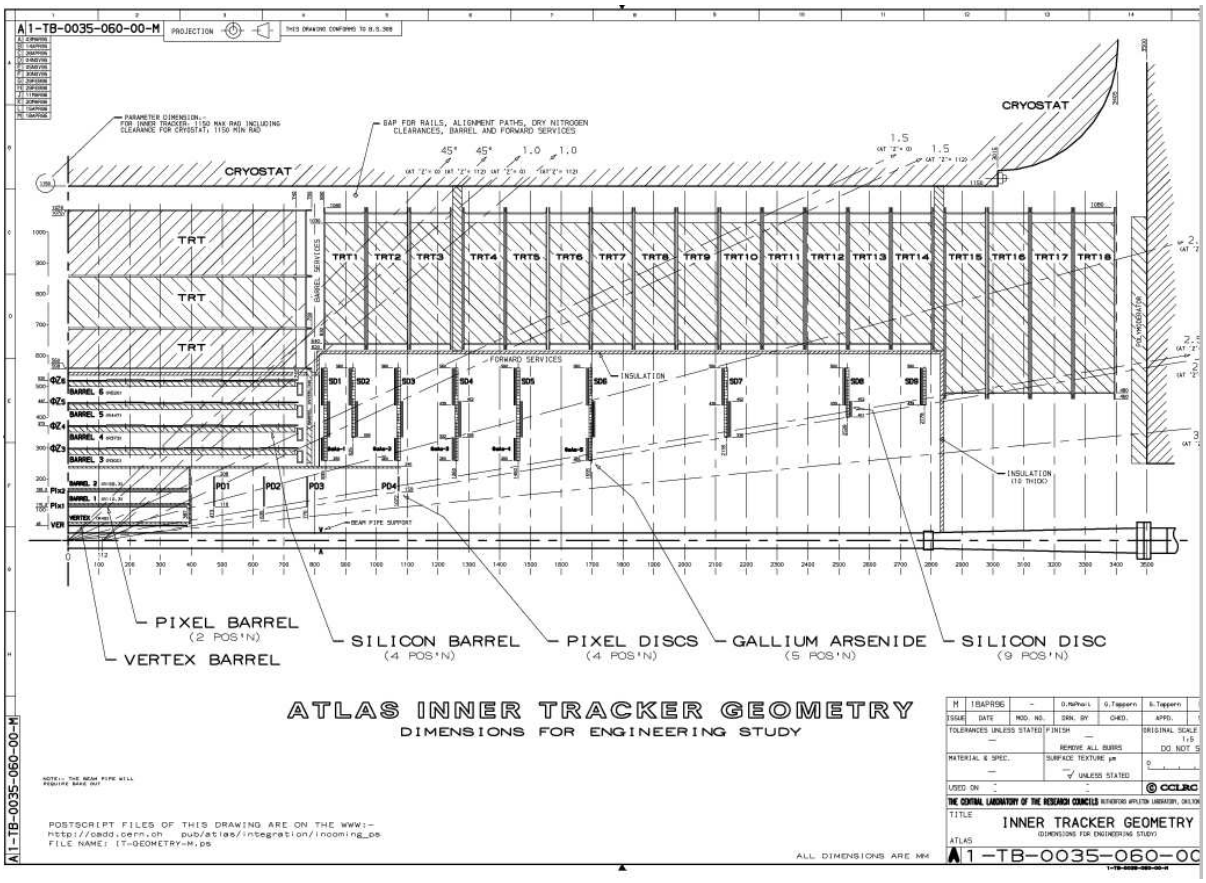

Figure 1.5: Cross-sectional view of inner detector along the beam (z)-axis with the interaction point down to the left.

The ID has a total length of 7 m and an outer radius of 115 cm, constrained by the cryostat containing the central solenoid manget (CS), and the liquid argon

electromagnetic calorimeter (LAr EM). It is divided into a barrel part covering $|\eta|<1$ and a forward/end-cap part covering $|\eta|<2.5$.

The inner detector consists of three detector subsystems:

- The pixel detector
- The silicon strip detector (SCT)
- The transition radiation tracker (TRT)

The track density and the radiation near the interaction point will be very high, so for the two innermost layers of the ID, semiconductor tracking devices have been chosen for detectors. They offer high resolution, and after an extensive development program they also offer high durability in the harsh radiation environment. The third layer of the ID is a straw detector, which offers tracking at lower granularity, but also at a much lower costs. Together, the three systems provides very robust tracking and pattern recognition abilities. Figure 1.6 shows a simulated event in the ID, demonstrating how the various parts of the inner detector will see the particles. The curving of the particle tracks is due to the 2T field in the inner detector, created by the central solenoid magnet.

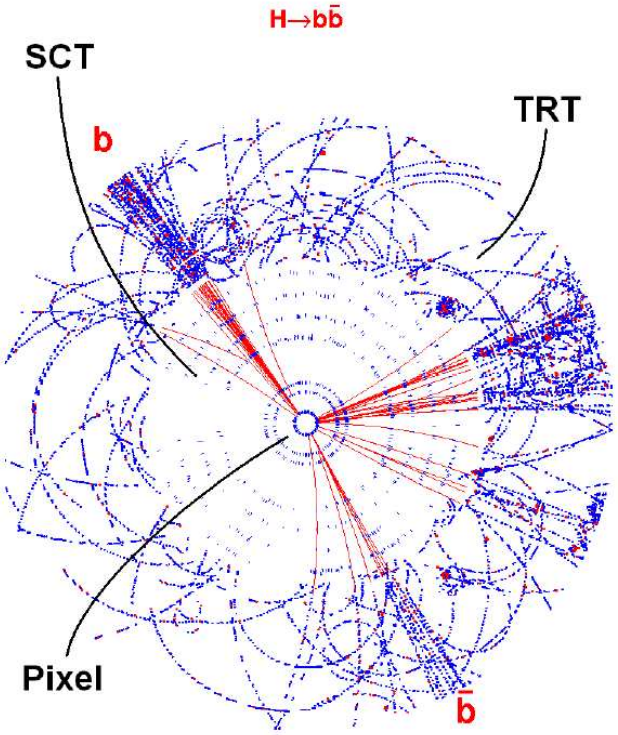

Figure 1.6: Simulated event in inner detector.

1.3.1.1 The Pixel detector

The pixel detector is the innermost layer in the detector, designed to provide a very high-granularity and high-precision set of measurements, as close to the interaction point as

possible. One of the important tasks of the inner detector is to discover signatures of short-lived particles such as B-mesons and tau leptons, important for finding traces of the Higgs boson. For this, very high resolution is needed.

The system consists of 140 million cells distributed over three barrel layers and eight forward discs, providing a minimum of three precision measurements over the full acceptance. It covers $|\eta|<2.5$, and the resolution is 12 µm in the Rϕ plane and 66 µm in the z plane, achieved by 50x300 µm pixel cells of n+ type technology. The pixels are arranged in modules, with the entire system consisting of 1500 barrel modules and 1000 endcap modules. Tracks can be distinguished if separated by more than ~200 µm.

Figure 1.7 shows the pixel detector, with the three barrels placed at approximately at r=4 cm, 10 cm and 13 cm, and the endcaps covering the radius 11 to 20 cm.

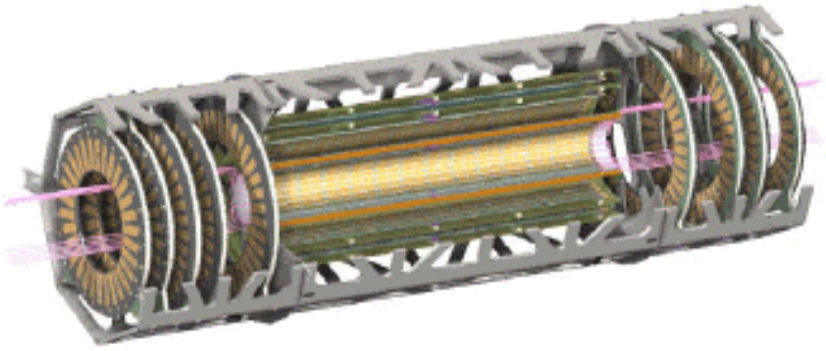

Figure 1.7: ATLAS Pixel detector

The radiation close to the interaction point is very high, estimated to 300 kGy of ionizing radiation, and over $5x10^4$ neutrons/cm^2 over the ten years of operation, and for this reason, the innermost barrel is designed for replacement. It is expected to have a lifetime of three years at reduced luminosity, and one year at full luminosity.

Chips reading 24x160 pixels each do the electrical readout of the pixels. These provide functions for buffering readout in case of an L1 trigger, with individual circuitry for each pixel element. [5]

1.3.1.2 SCT (*SemiConductor Tracker*)

The semiconductor tracker is the next layer in the detector. It operates at a lower resolution than the pixel detector, but was chosen because of the lower cost, and also because of the reduced need for resolution due to multiple scattering from the inner layers of the detector. The SCT is built from modules of siliconstrip detectors, see chapter 3 for a detailed description. The modules are assembled in 4 cylindrical central barrels at r=30.0, 37.3, 44.7 and 52.0 cm, in addition to 9 forward and 9 backwards disks, giving a

total sensitive area of 61 m^2 and 6.2 million readout channels. The SCT also covers $|\eta|<2.5$, and the spatial resolution is 16 µm in the Rϕ plane and 580 µm in *z*.

The SCT is built from two types of modules, endcap modules and barrel modules, with a total of 2856 modules required for the four barrel layers.

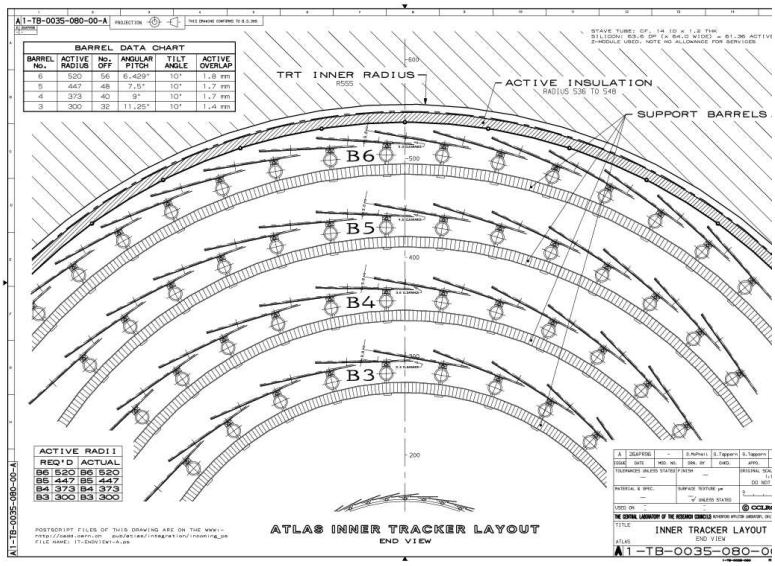

Figure 1.8: Crossectional view of SCT barrel down the z-axis, showing the alignment of the SCT barrel modules.

Both the electronics on the module and the detector strips will produce significant heat, so cooling of the modules is an important issue. Defects due to irradiation have also shown strong temperature dependence, so the SCT and pixel detector will share a nitrogen atmosphere of –7 °C. A cooling system based on cooling pipes containing perfluorinated propane (C_3F_8) kept at –14 °C will also be used. The pipes will be connected to each module to remove heat generated by the electronics and detecor leakage current. [6]

1.3.1.3 Transition Radiation Tracker (TRT)

The TRT surrounds the pixel and the strip detectors, and consists of 372 032 4 mm - diameter tubes of maximum length 150 cm, filled with gas (70 % Xe, 20% CF_4 and 10 % CO_2), with a 30 µm thick wire running through each tube. The TRT is, like the SCT and PIXEL detectors, divided in a barrel and endcap part, covering $|\eta|<2.5$ (see figure 1.5). The barrel covers the radius between r=56 and 107 cm, and the endcap r=64-103 cm. In the barrel, the tubes are placed parallel to the beam, and in the end-cap region perpendicular.

In each tube, a high voltage is maintained between the tube wall and the wire, so that when a particle traverses a tube, it produces a discharge, and a hit is detected. Through accurate timing of the discharge, it is also determined how far from the wire the particle passed. Each particle will on average hit 36 straws (see figure 1.6), and the points will together with the Pixel and SCT detector help reconstruct tracks at high resolution to determine momentum very precisely. The effective resolution of TRT will be 50 µm.

The detector features two different thresholds, one for particles hits (low threshold) and one for transition-radiation hits (high threshold). Transition radiation is radiation in the X-ray region that arises when ultra relativistic particles cross a boundary between two media with different dielectric constants. In the end caps, the transition radiation is created by placing stacks of 12 15 µm thick polypropylene foils (radiators) between each straw layer. In the barrel region, where space is a critical issue, the radiator will be put into all available space around the straws as sheets of loosely packed fibres with diameter 15 µm, made of polypropylene/polyethylene.[7]

1.3.2 Central Solenoid magnet (CS)

The next layer in the detector is the solenoid magnet. This is a superconducting magnet kept at T=4.5 K, sharing cryostat with the EM calorimeter. The magnet is 5.3 m long, creating a field of 2 T at the interation point, dropping to 0.5 T at the end of the ID. Since the CS is placed in front of the LAr EM, the material use is minimised to optimise calorimeter performance.

1.3.3 Calorimeters

ATLAS features two calorimeters, whose purpose is to absorb particles to measure their energy:

- Liquid Argon Electromagnetic (LAr EM) calorimeter
- Hadronic calorimeter

1.3.3.1 Liquid Argon Electromagnetic calorimeter (LAr EM)

The first calorimeter, the *LAr EM calorimeter*, is designed to absorb light electromagnetically interacting particles; $e^{\pm}$ and photons. It consists of an inner barrel cylinder and two end caps, covering $|\eta|<3.2$. On figure 1.9, the EM calorimeter is shown together with the other calorimeters of ATLAS. On the figure, the detector part of the EM cal. is marked as ‘accordion calorimeters’, reffering to the accordeon-like structure of the lead/liquid argon layers. When electrons and photons interact with the lead, which has a short radiation length, particle showers are created. These showers ionize the argon, in which the signal from the ionisation is picked up. The measurements are done in three samples, in addition to a presample done between the CS and the cryostat wall for energy loss adjustment due to the ID, solenoid and cryostat.

1.3.3.2 Hadronic calorimeters

The *Hadronic calorimeters* is designed to absorb hadrons (particles made from quarks), covering $|\eta|<4.9$. It consists of three parts; Hadronic Tile, Hadronic LAr and Forward LAr, see figure 1.9. The tile part is built from sandwiched layers of iron and scintillator, with iron as the passive (absorber) material in which the hadronic particles shower develops. The scintillators are the active sampling layers measuring ionization from the particle showers.

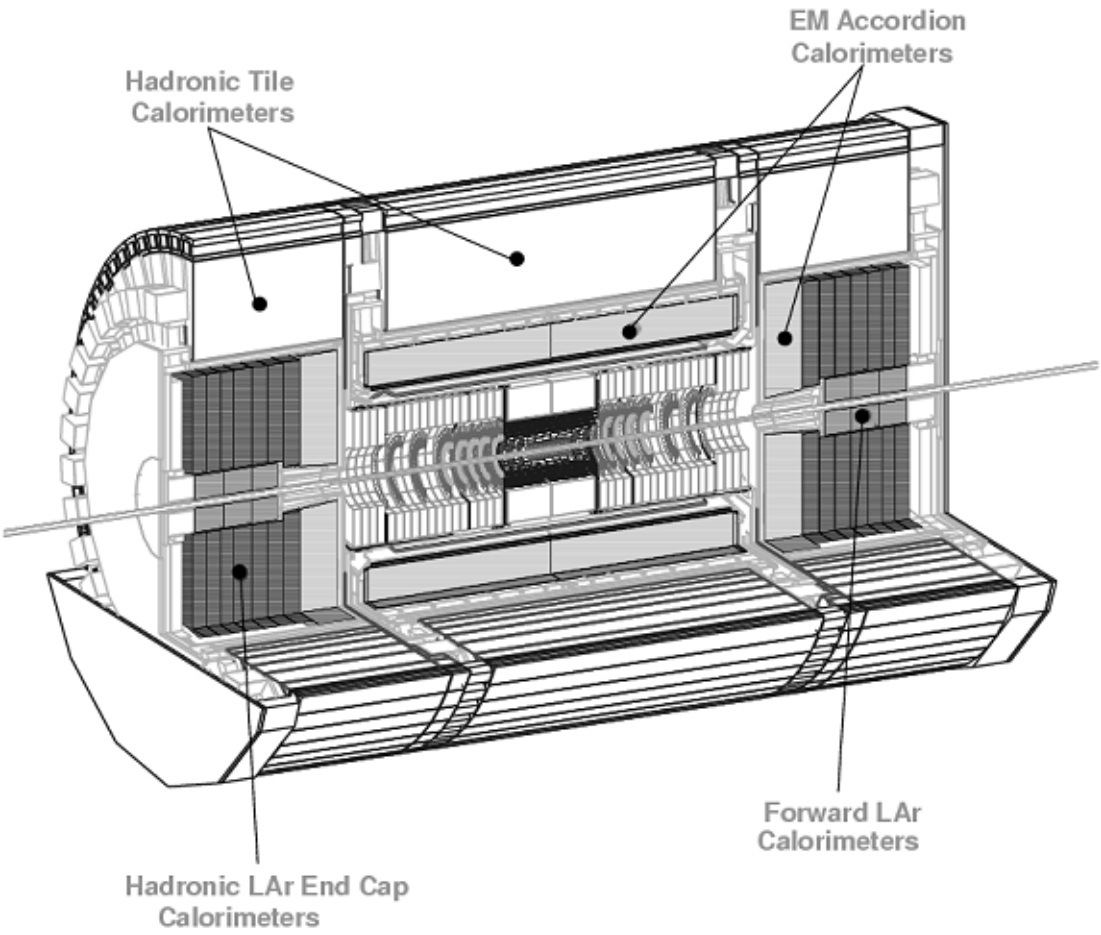

Figure 1.9: ATLAS calorimeters

The hadronic tile calorimeter covers r=228.0 to r=423.0 cm, and the region $0.8<|\eta|<1.7$. Its total length in z direction is 1266.0 cm.

The Hadronic LAr End Cap Calorimeters has an outer radius of 2.03 m, covers $1.5<|\eta|<3.2$ and uses much of the same technology as the EM LAr calorimeter. The forward LAr Calorimeters covers $3.1<|\eta|<4.9$, built of copper and tungsten. This part of

the calorimeter will operate under high radiation because of the small angle soft collisions from the majority of interactions.[5]

By comparing the total amount of energy measured with the centre of mass collision energy, one can find E_T^{miss}, the missing energy. This is due to neutrinos (see chapter 2), not directly observable in the detector. By determination of E_T^{miss}, the neutrinos in a process can be indirectly detected.

1.3.4 Muon system

Muons and neutrinos are the only particles that escape through the calorimeter. Since the muon has ~200 times the mass of the electron, the cross section for bremhsstrahlung is much smaller than for the electron, and muons mostly escape the calorimeter. They also have long enough lifetime to reach outside the calorimeter, in contrast to the heavier tau lepton, which has a much shorter lifetime. A separate muon detection system is thus needed.

1.3.4.1 Toroid muon magnet system

The barrel and endcap muon magnets create a toroidal magnetic field of approximately 4 T, consisting of 8 flat 25 x 5 m coils for the barrel part, and two 11 m high and 5 m thick endcaps, see figure 1.10. The magnets are superconducting, kept at a temperature 4.2-4.5 K. The toroids have air cores, to minimize scattering, and have also integrated cryogenics.

1.3.4.2 Muon spectrometer

The muon detector is the outermost layer in ATLAS, 26 m long with inner radius of 9.4 m and outer radius of 19.5 m. The deflection of the muon in the magnetic field is measured to identify the momenta and energy of the muons. The spectrometer covers $|\eta|<2.5$ and has a total area of 5300 m^2, and can be divided into two functional parts; trigger chambers and precision chambers [5] also see figure 1.10).

The trigger chambers are used by the triggering system to detect interesting events based upon high p_T, see the next section. They provide a low resolution, ~1 cm, but very fast readout, <25 ns, to uniquely identify bunch crossings. The technology used is Resistive Plate Chambers for the barrel part, and Thin Gap Chambers for the end cap parts.

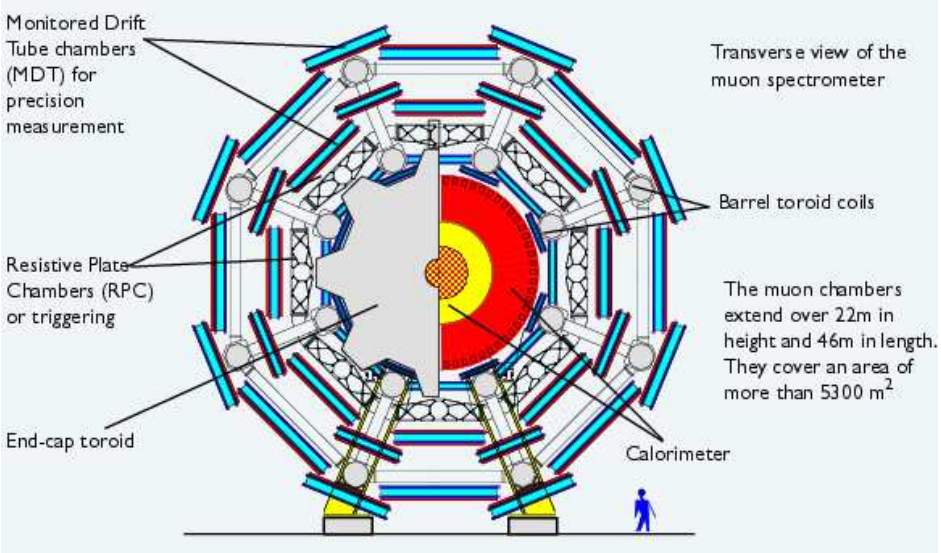

Figure 1.10: Crossectional view of muon spectrometer and air-core toroid magnet system.

The precision chambers provide a resolution of down to 60 µm, but with a drifttime up to 500 ns. The technology used is Monitored Drift-Tube chambers and Cathode Strip Chambers.

1.3.5 Triggers, Dataaqusition and Detector Control System (DCS)

With a beamcrossing rate of 40 MHz, one bunch each 25 ns, and 20-25 interactions per crossing, there will ~10^9 events per second at the interaction point in ATLAS. On average, 1000 particles traverses the pseudo-rapidity region of $|\eta|<2.5$ at each bunch crossing (see section 1.2.1.) A data amount of 7 Tb per second would be generated if all events were to be recorded, which is far too much to store. This means that ATLAS must have some kind of criteria for selecting the interesting events, which should occur 10-100 times per second. The large amount of uninteresting events is known as *pile-up*, or *minimum bias events*, characterised by a low p_T.

The interesting events will be selected through very rapid analysis, using various criteria for tagging an event as interesting. The trigger system is divided in three levels, with each level examining a larger dataset. The first trigger, the level 1 (L1) trigger, has a decision time of 2.0 µs, based on data from the calorimeters and dedicated trigger chambers in the muon spectrometer using high p_T muons, electrons, photons, hadrons and jets as well as large E_T^{miss}, as indications of intersting events. The L1 trigger is designed to be a conservative trigging system, i.e. it should select *all* potentially interesting events, which means that it also selects a lot of eventually uninteresting events. Because of the decision time of L1, all data from the detectors must be kept in pipelines long enough for a L1 decision to be taken. The L1 is very fast, analyzing several Tbits/s, and should create a maximal L1 trigger rate of ~75 kHz, which is the highest rate the ATLAS front-end

systems can accept. This is upgradable to 100 kHz. All data from all detectors are kept in pipeline memories long enough for a L1 trigger to arrive. [5]

After L1, there is L2 analyzing, which focuses on region of interests marked by the L1 trigger. L2 has a decision time of <10 ms, and is followed by the L3 trigger which reconstructs the detector data and sends it to storage.

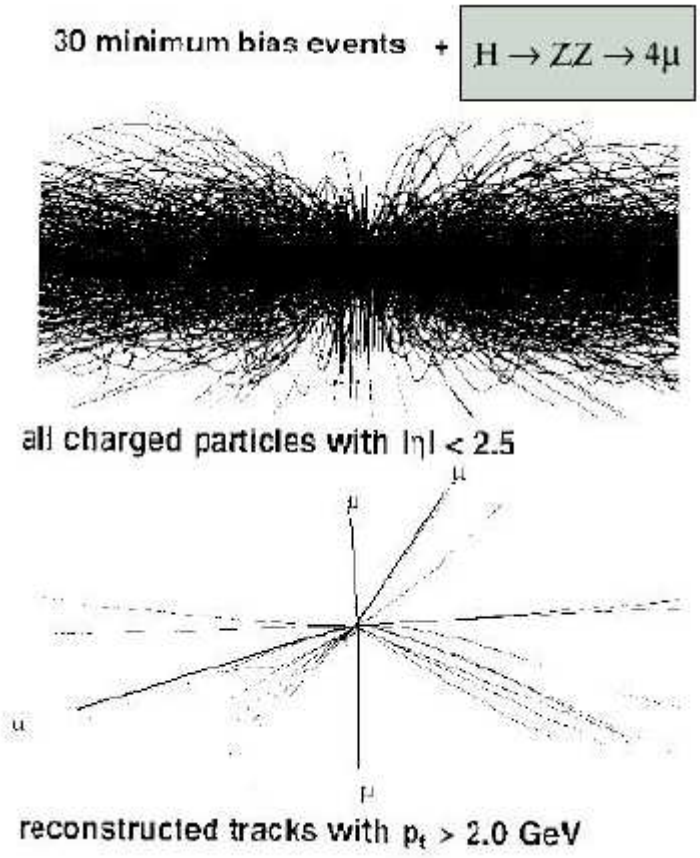

Figure 1.11: Event after L1 trigger cutoff due to high p_T **[]**

ATLAS will also use a sophisticated *Detector Control System* (DCS). The systems needs to monitor approximately 100 000 parameters of detector operation like currents, temperatures, humidity and such, and present these to operators in a practical way.

Each year, ATLAS will produce a data volume of 1 Pbyte (10^{15}) bytes, which needs to be analyzed. Distributed networking will become very important to enable physicists, independent of location, to access data and perform physics analysis. How this will be implemented, is still an issue under discussion, and will greatly depend on the evolution of networks in the years both before and during the run of ATLAS.

Chapter 2: Theory

This chapter gives an overview of the theoretical foundations of particle physics.

2.1 The Standard Model

Our present understanding of particle physics is contained in a theoretical framework called the Standard Model (SM). The theory has been gradually developed for over half a century, and is today in excellent agreement with experimental data, both from LEP and a variety of other experiments.

2.2 Fundamental particles

According to the theory, there are 12 fundamental particles, divided into three 'generations':

Family	**Leptons**		**Quarks**	
Spin	½		½	
Electric Charge	-1	0	2/3	-1/3
1.generation	**e** Electron m_e=0.511 MeV	ν_e Electr. neutrino m<3 eV	**u** m_u=1-5 MeV	**d** m_d=3-9 MeV
2.generation	**μ** Muon m_μ=105.6 Mev	ν_μ Muon neutrino m<0.19 MeV	**c** m_c=75-170 MeV	**s** m_s=1.15-1.35 GeV
3.generation	**τ** Tauon m_τ=1776.99 MeV	ν_τ Tauon neutrino m<18.2 MeV	**t** m_t=174.3±5.1 GeV	**b** m_b=4-4.4 GeV

Table 2.1: Fundamental particles of the Standard Model [8]

For each fundamental particle, there exists an antiparticle with identical mass and spin, but with opposite charge. We say that particles are fundamental in the sense that they seem to occupy a single point in space and hence are point-like, also called local entities.

2.2.1 Quarks

Quarks were first observed in 1968 at SLAC, as an indication of an inner structure in protons. Quarks exist in six *flavours*; *u*p, *d*own, *c*harm, *s*trange, *t*op and *b*ottom with corresponding anti-quarks. They are fermions, that is, they have an intrinsic spin of ½, and carry electric charge +(-)2/3 and –(+)1/3, thus interacting electromagnetically. The quarks also interact through the strong force, carrying a color charge. This charge exists in three different colors (anti-colors) red, green and blue, giving the quarks three additional degrees of freedom. The strong (color) force confines the quarks so that they create color-neutral objects, supported by the fact that a free quark has never been observed.

Particles composed of quarks are called hadrons, with particles made up from three quarks called baryons, and particles from one quark/anti-quark pair called mesons. The proton (u u d) and the neutron (u d d) are the only baryons found in nature, while the other baryons and the mesons are relatively short lived particles which eventually decays into stable particles. Whether the proton actually is stable or not, is still an unresolved issue, but it has been given a lower life time limit of 1.6×10^{25} years. [8]

2.2.2. Leptons

Leptons are the other class of elementary particles. They are also fermions, and are, like the quarks, divided into three generations. The charged leptons are the electron, muon and tauon, carrying an electric charge of –1. The three particles couple with equal strength in electromagnetic interactions, known as *lepton universality*. The only difference between them is the mass, ranging from 0.5 MeV to 1710 MeV. It is a challenge of particle physics to try to find out why these masses are so different.

For each charged lepton, there is a corresponding *neutrino*. Until recently, the neutrinos were believed to be massless, but are now believed to have a small mass. Since they do not carry any charge, they interact with matter only through the weak force, and this, combined with their low mass, makes them very hard to detect.

2.3 The fundamental forces

The three forces described in the standard model are presented in table 2.2. Each force has its mediating particle, called *gauge bosons*, of integer spin. Gravity is also listed, although it is not yet, despite serious efforts, incorporated in the Standard Model. However, gravity is very weak at small scales, so it can be safely neglected.

One of the primary goals of particle physics is to unite all the four forces into one fundamental force, but so far, only the electromagnetic and weak force have been unified successfully. This was done in Electroweak theory (Salaam-Weinberg 1969), and is not

yet fully experimentally verified, since the Higgs boson is yet to be found (see section 2.4).

To each of the interactions, certain conservations laws apply. Some are conserved in all interactions:

- Energy (E)
- Momentum (p)
- Angular Momentum (L)
- Electrical charge
- Colour charge
- Baryon number (B)
- Lepton number (L_e, L_μ, L_τ)

Mathematically, the Standard Model is based on local relativistic quantum field theory, describing interactions due to the three forces with the theories QED, Electroweak Unification and QCD.

Interaction	**Relative strength**	**Range**	**Mediating particle**	**Charge of particle**	**Felt by**
Strong	1	~10^{-15} m	**g (gluon)** Spin-1 m_g=0	1 of 8 color charge combinations	Quarks and gluons.
Electro-magnetic	~10^{-2}	$1/r^2$	**γ (photon)** Spin-1 m_γ=0	0	All particles carrying electric charge.
Weak	at 10^{-18} m: ~10^{-2} at 10^{-17} m: ~10^{-6}	~10^{-18} m	**$W^{\pm}$** Spin-1 m_W=80.4 GeV	Electric charge: ± 1	Left-handed particles, right handed anti-particles.
			Z^0 Spin-1 m_Z=91.18 GeV	0	Left-handed and right-handed particles, but with different strength.
Gravity	~10^{-43}	$1/r^2$	**Graviton** Spin-2 m_G=0	0	Massive particles.

Table 2.2: The three fundamental forces in the Standard Model + Gravity. [9]

2.3.1 The strong force

The strong force is the force responsible for holding the nucleus and hadrons together, and must thus be stronger than the repulsive force due to the electromagnetic interactions between the positively charged protons in the nucleus.

The mediating particle of the strong force is the massless *gluon*, which can carry one of eight different color combinations of the colors red, green and blue.

The strong force does not fall off with distance, but is rather believed to be proportional to *r*. This means that when two quarks move away from each other, the field between them, also called a *colour tube*, grows stronger. This is due to the fact that gluons carry colour charge and can thus couple to other gluons. Before the quarks are separated at an observable distance, the energy contained in the field between them reaches the pair-production energy, creating an additional quark/antiquark pair. This phenomenon is known as quark confinement, preventing single quarks from existing and always creating colour neutral objects, see figure 2.1.

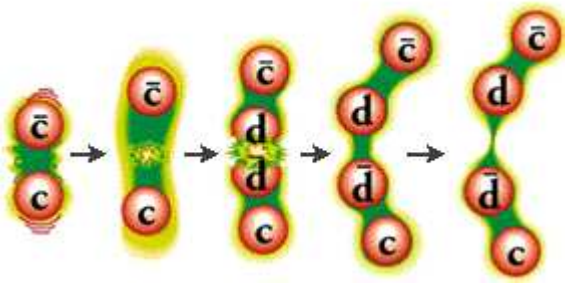

Figure 2.1: Confinement of quarks due to strong force.

The interactions due to the strong force are described in the gauge theory known as Quantum ChromoDynamics (QCD). The quarks are symmetric under exchange of the three colours, which results in an SU(3) symmetric theory. The theory also only applies at very high energies, where the quarks can be considered as free particles.

2.3.2 The electromagnetic force

The field quantum of the electromagnetic force is the photon, being massless and carrying no charge. Because of this massless and chargeless nature of the photon, it has infinite range, while the strength of the force goes like $1/r^2$.

Photons couple to particles carrying electric charge. The nature of the electromagnetic force is believed to be fully understood, and is described in the theory Quantum Electrodynamics (QED), which has been very successful in describing a variety of phenomena. Due to the electromagnetic force only having one type of charge, the theory is a U(1) symmetric gauge theory, making the photon its own antiparticle.

2.3.3 The weak force

While the EM force had infinite range due to the massless nature of the photon, the weak force has a very short range due to its heavy field quanta, the $W^{\pm}$ and Z^0 bosons. The Heisenberg uncertainty relation $\Delta E \Delta t \geq \hbar$ limits the effective range of the weak force to $\sim 10^{-18}$ m.

The charged currents of the weak force, $W^{\pm}$, only couples to left-handed particles or right-handed anti-particles. A left-handed particle has its spin in the opposite direction of its momentum (negative helicity). All massive particles have both a positive and negative helicity component, but since these can be of significantly different magnitude, we can say that for the weak force, P (spatial inversion)-parity is not conserved (since this inverts the helicity, making a left-handed particle a right-handed particle and opposite).

For the charged currents, this violation is maximal, but for Z^0 only partial, since Z^0 couples to both left- and right-handed particles, but with different strength. The weak force is also C (charge) parity violating, since changing the charge of a particle turns it into its antiparticle, but keeps the helicity and thus changes the coupling to weak interactions. However, the combined CP-parity is conserved for most weak processes. The fact that not all weak processes obey the combined CP-parity, is not satisfactorily explained in the standard model, and is an area that is currently being studied, and also will be studied in LHC, in particular through the studies of B-mesons.

The weak force is also responsible for the decays of the heavier 2. and 3. generation particles, being the only force that can 'change' the type of a particle. Two examples are shown in the Feynman diagrams below.

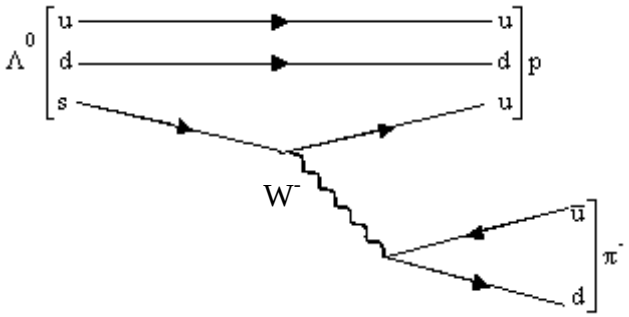

Figure 2.2: Feynman diagram showing $\Lambda^0 \rightarrow p + \pi^-$.

Figure 2.2 shows an example of an Λ^0 particle decaying into a proton and a π^- via a W^- boson. This is known as a *Flavour Changing Current*. In this case, an s quark changes to an u quark through emission of a W^-. This is a change across a generation, and is possible because the W bosons couple not to the mass eigenstates of the quarks, but to linear combinations of these, given by the *Cabibbo-Kobayashi-Maskawa (CKM)-matrix*.
The Standard model forbids Flavour Changing Neutral Currents (FCNC), meaning no interactions in which the Z^0 boson changes one quark to another.

For leptons, the question whether a lepton can change between different generations through charged currents, and thus violating the conservation of lepton number, is still an unresolved issue.

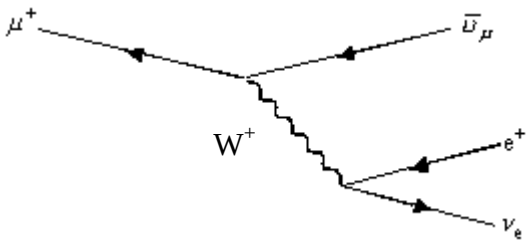

Figure 2.3: $\mu^+ \rightarrow e^+ + \nu_e + \bar{\nu}_\mu$

Figure 2.3 shows the decay of a μ^+ into a muon antineutrino, a positron and an electron neutrino through the emission of a W^+ boson, also known as muon decay.

2.3.4 Electroweak unification

In 1968, Salam and Weinberg showed that the electromagnetic and weak interactions could be understood as two aspects of the same interaction. They showed that one could, by introducing a new quantum number *weak isospin T*, group the left-handed (interacting through charged currents) fermions into isospin doublets with T=1/2, and right-handed fermions into isospin singlets with T=0. The doublet components could transform into each other via the weak interaction, obeying an SU(2) symmetry.

The gauge bosons have T=1, making it possible to group the $W^{\pm}$ bosons into a triplet, postulating a third W^0 boson with T_3=0. A singlet with T=0 was also needed, so a boson B^0 with T=0 was postulated. The B^0 gauge boson is connected with the quantum number hypercharge (a combination of isospin and electric charge), governed by a U(1) symmetry. It turns out that the real particles the photon and Z^0 can be described as linear combinations of the theoretical particles W^0 and B^0:

$$\begin{aligned} |\gamma\rangle &= \cos\theta_W |B^0\rangle + \sin\theta_W |W^0\rangle \\ |Z^0\rangle &= -\sin\theta_W |B^0\rangle + \cos\theta_W |W^0\rangle \end{aligned} \qquad (1)$$

This mixing is expressed as a rotation through the *electroweak mixing angle* θ_W, also known as the *Weinberg angle*. This showed the connection between the weak- and the electromagnetic force, with the resulting theory being a SU(2)xU(1) symmetric theory.

2.4 The Higgs model

For the electroweak theory to be gauge invariant, the particles involved need to have zero mass, which is obviously not the case. To solve this problem, the concept of spontaneous symmetry breaking is introduced in the theory, requiring an asymmetric ground state with a non-vanishing vacuum expectation value.

The symmetry breaking gives rise to four scalar (spin 0) fields, which at first give us expressions that have no physical meaningful interpretations. However, three of these fields can be removed by local gauge transformation, which in turn gives masses to the weak gauge bosons.

Since spontaneous symmetry breaking requires a non-zero vacuum expectation value, the value is interpreted as a field being present all over space, with the leptons acquiring mass by interacting with this field through the *Yukawa-coupling*.

This leaves us with one of the initial four scalar fields left, interpreted as the scalar (spin=0) particle the *Higgs boson*. If this particle exists, it should be detectable in ATLAS, verifying the theory of how the particles acquire their masses, one of the most central and fundamental questions in particle physics.

2.4.1 Higgs production/detection in ATLAS

As discussed in the previous section, the Higgs mechanism provides an explanation for the particle masses, so if signs of the Higgs boson are detected in ATLAS, it will be a major success of the Standard Model. If it is not found, the explanation of mass will need to be revised. It is clear that this is an important issue of the theory, and has to be tested.

The Standard Model does not predict the mass of the Higgs boson, but the present lower limit from LEP is around 108 GeV. If the Higgs boson is not found at the Tevatron at Fermilab, the lower limit could be extended up to ~120 GeV. There is an upper theoretical limit of 1 TeV to preserve unitarity at high energy.

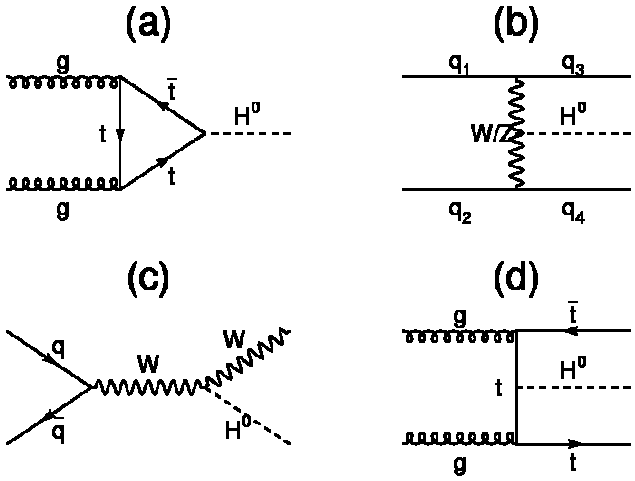

Figure 2.4: Feynman diagrams showing the most important Higgs production channels: a) Gluon fusion through top loop, b) Vector boson fusion, c) Associate production with W, d) Associate production with top pair.
Production of a Higgs boson will in LHC be dominated by gluon fusion, but other processes will also contribute. In figure 2.4, the four most important processes, valid for all realistic Higgs masses, are shown.

Although the production of Higgs does not depend on its mass m_H, the channels to which it can decay has a strong mass dependence. In figure 2.5, the expected total branching ratios of the Higgs boson is shown as a function of the Higgs mass.

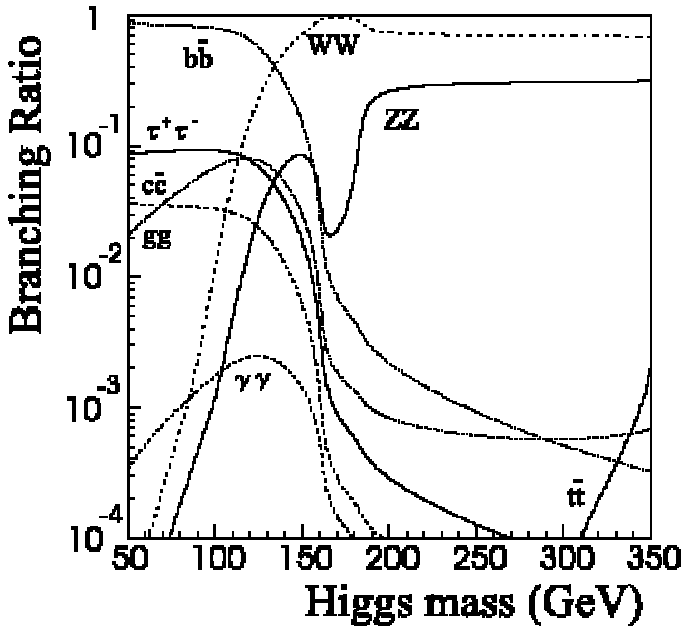

Figure 2.5: Total Higgs decay branching ratios as function of the Higgs mass.[4]

The mass range is usually divided in three, each with its own dominant decay channels:

- ***m_H < 120 GeV:*** *$H \rightarrow b\bar{b}$ dominates, since b-quarks are the most massive fermions cinematically accessible in this region. However, the QCD background will be 10^5 larger, meaning it will be impossible to extract this signal from the background, so other channels will be experimentally more useful in the detection of the Higgs boson.*
- ***80 < m_H < 150 GeV****: $H \rightarrow \gamma\gamma$. This process is very rare, but provides very clean signatures, so it can be used for identifying a Higgs boson.*
- ***130 GeV< m_H < 2 m_Z****: $H \rightarrow WW(*) \rightarrow l\nu l\nu$ and $H \rightarrow ZZ(*) \rightarrow 4\ l$ dominates, where one of the vector bosons can be virtual.*
- ***m_H > 2 m_Z****: WW and ZZ pairs dominate, both bosons real. This is the easiest region to discover the Higgs in, since the $ZZ \rightarrow 4\ l$ gives a very clear signature, almost background free.*

2.5 Open questions

The standard model is not a complete theory in the sense that it completely describes particle physics; there are still many open questions. 18 or more parameters, depending on the counting scheme, must be filled in the theory based on experimental values, and has no accepted or verified theoretical explanation. [10]

2.5.1 Physics beyond the standard model

There are several arguments that point in the direction that the standard model is not the ultimate theory of particle interactions. Among these is the lacking physical justification of the Higgs field, and that the Standard model fails to unify the coupling constants of the three forces at high energies (~10^{16} GeV). However, if the Standard Model is augmented with the concept of *SuperSYmmetry (SUSY),* it succeeds in combining the three coupling constants, as shown on figure 2.5. This is referred to as *MSSM (minimal supersymmetric standard model).*

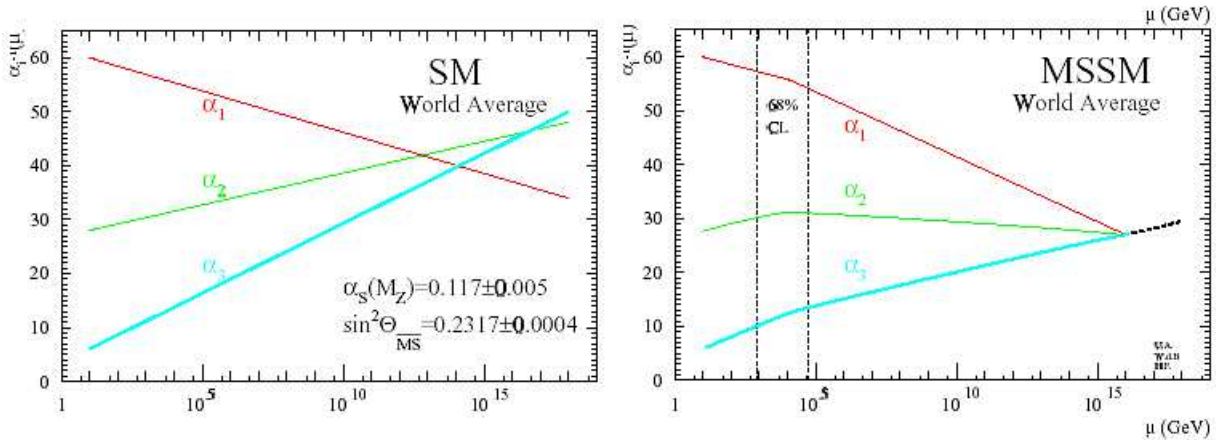

Figure 2.5: Coupling constants of the EM, weak and strong forces as function of energy in SM and MSSM [4].

Supersymmetry predicts that each particle should have its SUSY partner, resulting in quarks/squarks, leptons/sleptons, with the fermion partners being bosons, and opposite. The masses of these partners have a lower experimental limit of 100 GeV (250 GeV for squarks/gluinos), and the decay modes, i.e. the patterns by which the SUSY particles can be discovered, will like the Higgs, depend on the unknown mass of the SUSY partners.

2.5.2 Extra dimensions

There exists several theories using the concept extra dimensions. Some of these will be explored in ATLAS, for example the Randall-Sundrum scenario. The extra dimensions are introduced to solve the *hierarchy problem* of the Standard Model, the question why the forces are of so different scale, with a ratio 10^{-43} between the strength of the strong force and gravity. The answer to this could be that gravity propagates through extra dimensions, not making all its force 'visible' for us.

2.5.3 Other questions

There are several other unresolved issues that will be subject to further investigation in ATLAS/LHC:

- *Are quarks and leptons really elementary particles?*
- *Can gravity be incorporated in the Standard Model?*
- *Further insight in the nature of CP violation.*
- *Are there more particle generations, and if not, why three?*
- *How do the coupling constants evolve over the unprecedented range of energy?*

We see that there are many open questions, and hopefully, some of the answers can be found from the data from LHC/ATLAS.

Chapter 3: The SCT barrel module

The SCT barrel module will be used for constructing the SCT detector, part of the Inner Detector in ATLAS (see chapter 1 for an overview of the ID and ATLAS). The detector will be constructed from four barrels and 9 end caps, with the barrel and end caps being assembled from only two types of modules; barrel and end cap modules, making it a highly modular system. Figure 3.1 shows a barrel module.

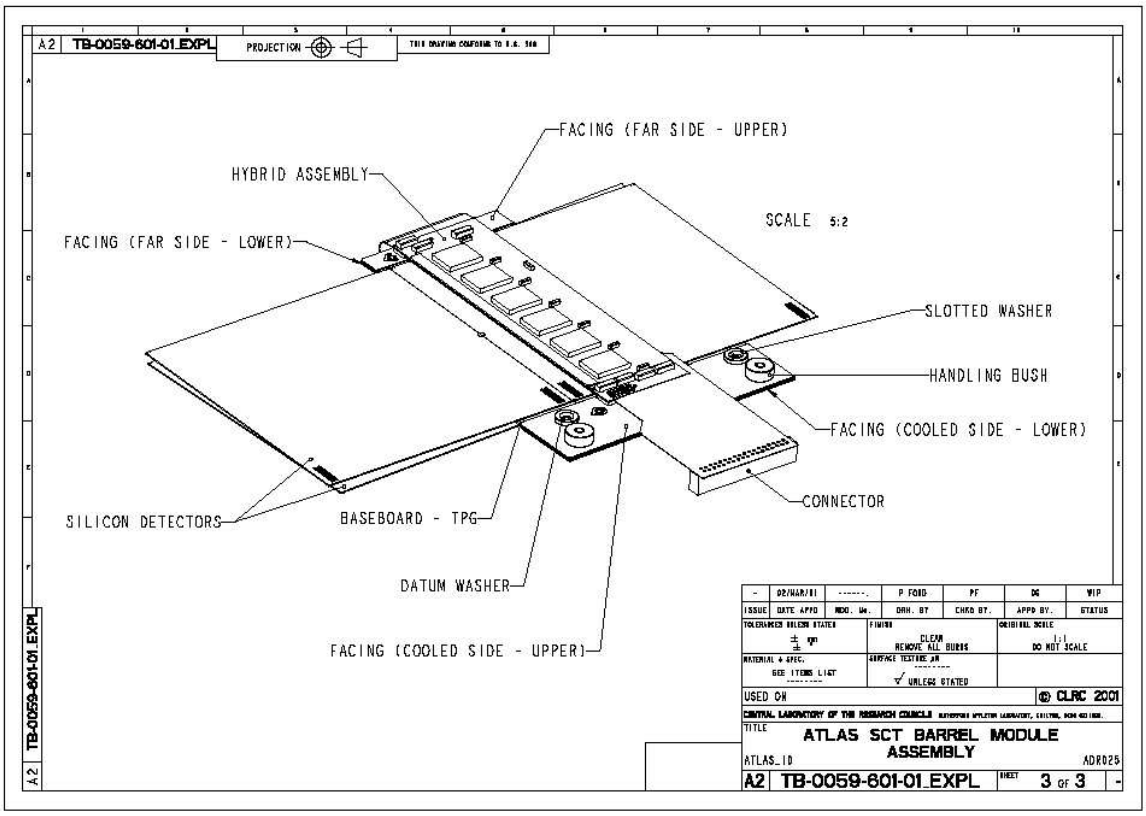

Figure 3.1: Schematic view of SCT barrel module [11]

Each SCT module consists of three major parts:

- ***Four silicon single sided reverse biased p-on-n strip detectors****. Two detectors are bonded strip-to-strip to form 126.09 mm long strips. Two such configurations are then glued together back to back with a 40 mrad stereoangle, separated by a VHCPG baseboard. The detectors have a strip spacing of 80 µm, so the resolution perpendicular to the strips is very high. However, because of the 40*

mrad stereo angle, a resolution along the strips is achieved as well, giving a resolution of 16 µm in Rϕ and 580 µm in z.

- ***VHCPG (Very High Thermal Conductivity Pyrolytic Graphite) Baseboard.*** *The baseboard is placed between the detectors, and works as a physical support, conductor and thermal transporter.*
- ***Hybrid with on-detector readout electronics.*** *The hybrid carries the front end electronics; 12 ABCD3T ASICS with one readout channel for each strip on the module.*

Each part can be seen on the figure, and in the following sections, each part will be given a more detailed description.

3.1 Silicon Strip Detectors

On the figure below, a single silicon strip detector is shown. It consists of 768 silicon strips (diodes), its dimensions are 63.56x65.6 mm, and the spacing between each strip is 80 µm. The detector is specified to be 285±15 µm thick.

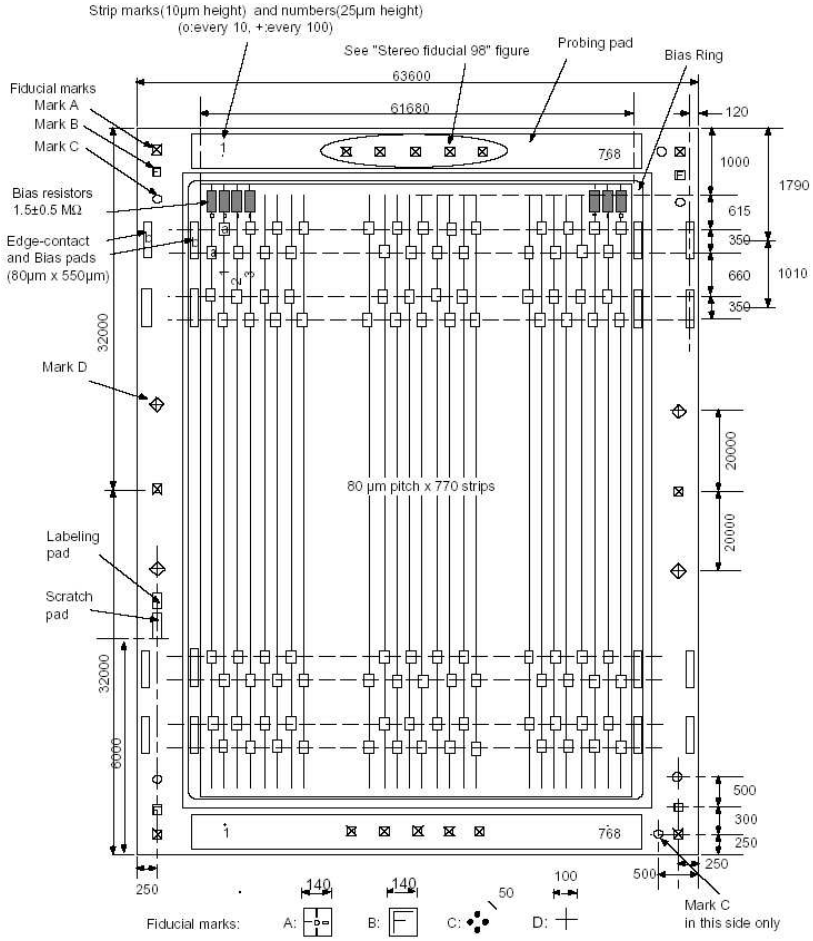

Figure 3.2: Geometrical layout of silicon strip detector (ATLAS98) [12]

In the next sections, an overview of semiconductors generally and the operation of the silicon detector specifically will be given.

3.1.1 Semiconductors

Semiconductors are a group of materials whose conductivity lies between that of conductors and insulators. The conductivity of semiconductors range from 10^{-8} to 10^{5} S/cm, and contrary to metals, the resistance decreases with temperature. Semiconductors have been known since early nineteenth century, but the fields in which they could be applied were not realized before 1947, when the first transistor was made at Bell Laboratories, clearly one of the most important technological discoveries of the 20th century.

3.1.2 Band theory of solids

The conductivity of different materials a can be explained in terms of energy bands. The figures below show the band structure of an insulator, a metal and a semiconductor.

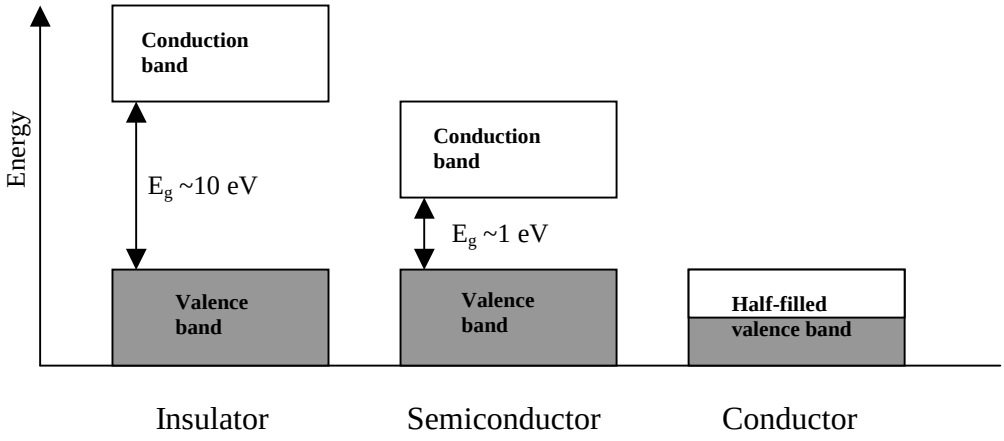

Figure 3.3. Band theory of solids

The bands represent the energy levels available to the electrons in the material, and are due to the splitting of energy levels that occur when atoms are brought close together. A material has a range of energy bands, and at low energy levels, there is an *energy gap* E_g between the levels. This means that the electrons in one band need a certain energy E_g to "jump" up to the next band.

In an insulator, the outermost band, also called the *valence band*, is completely filled, and the energy gap to the next band with unoccupied states, the *conduction band*, is large, typically E_g ~ 10 eV. This means that the energy needed for an electron to move from the valence band to the conduction band is very large, and there are thus very few electrons occupying these states. This means in turn that there are few "free" electrons to drift with an applied field, giving the material a high resistance.

Semiconductor materials are characterised by having a much smaller energy gap, $E_g \sim 1$ eV. This makes it possible for many electrons to be thermally excited to the conduction band. However, this depends on the thermal energy of the electrons, which means that semiconductors are poor conductors at low temperatures. However, as the thermal energy increases, the number of electrons excited to the conduction band increases, resulting in a strongly temperature dependent conductivity, in contrast to metals, where conductivity slowly decreases with temperature.

In a conductor, the outermost band in the material is not full. This means that the electrons have many unoccupied states very close to their own energy, and a small amount of thermal energy is all that is needed for the electron to move to a nearby state. So when an electric field is applied, the electrons move very easily with the field, due to the vacant states in the energy band.

3.1.3 p- and n-type semiconductors

There are both positive and negative charge carriers in a semiconductor. When an electron is excited from the valence band to the conduction band, it leaves a *hole* in the otherwise filled band. A valence electron from a nearby site can then move to the empty hole, and the hole thus appears as a travelling positive charge, *e+*.

To alter the properties of a semiconductor, a technique called doping is used. This involves adding specific impurities into the semiconductor to change its conductive properties. For example: A semiconductor from group 4 (four valence electrons) is doped with Arsenic, which has five valence electrons. Four of these electrons participate in the bonding, leaving one almost free electron in the arsenic atom. Its energy level is very close to the conduction band, and only a small amount of thermal energy is needed to excite the electron to the conduction band, see figure 3.4, making it practically free to move in the lattice of the semiconductor. The arsenic effectively *donates* an electron to the semiconductor, and is thus called a *donor*. The semiconductor then has an excess of free electrons, and is called a *n(egative)-type* semiconductor.

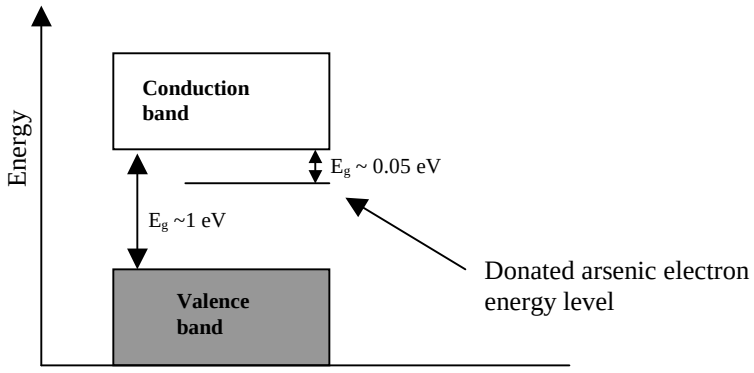

Figure 3.4: n-type semiconductor energy bands.

On the contrary, if a semiconductor is doped with atoms having *less* valence electrons than itself, the doping atom has a vacant state in its electron shells, referred to as a *hole*. This leaves a free state for other electrons in the valence band to move into, and the donor thus *accepts* an electron and is called an *acceptor*. The energy levels of these holes lie in the energy gap, but only right above the valence band, which makes it easy for an electron to move into a hole. This is called a *p(ositive)- type* semiconductor, since it have an excess of holes in its lattice.

3.1.4 p - n junction

If p- and n-type semiconductors are joined in a *p-n junction*, the holes and the electrons in the border between the two materials will combine to form a *depletion region*. The donor electrons of the n-type will occupy the holes of the p-type conductor, and the region will thus be *depleted* for moving charge carriers. When the electrons of the n-type material moves over to the p-type, they leave immobile positive ion cores, which will set up a field over the depletion region, of order $10^3 – 10^5$ V/cm, which keeps the depletion region free of mobile charges. This field creates a potential barrier, preventing further diffusion of holes and electrons across the border. The depletion region is typical a few µm wide. The p-n junction is also called a diode, and its most notable feature is its ability to pass current in only one direction. If a positive voltage is applied to the p-side of the junction, the potential barrier is decreased, and a current can flow through the p-n junction, known as *forward bias*. If a positive voltage is applied to the n-side, the depletion region increases, and only a very small leakage current can flow through the p-n junction. This is called *reverse biasing* of the p-n junction, and is the principle used in the silicon strip detectors.

3.1.5 Silicon as a semiconductor

Silicon is an elementary solid-state semiconductor, belonging to the IV group in the periodic system. It is the most used semiconductor material, mainly for two reasons; it has high availability, being a very usual material in the earth’s crust, and it yields a low leakage current.

3.1.6 Operational principle of silicon strip detectors

The reasons for using silicon strips as detectors are many; low cost, good spatial/energy resolution, fast response time, and good signal to noise ratio. The output signal is directly proportional to the energy deposited by the particle passing through the silicon, which makes it easy to measure this energy.

The ionization energy (the energy needed to produce an electron-hole pair) of silicon is 3.6 eV. The mean energy loss for a mip[2] particle traveling through the 285 µm silicon wafer is, according to the Bethe Bloch formula ~90 keV, which corresponds to about 25 000 electron-hole pairs. [13]

Below (figure 3.5) is shown a cross section of the detector. It consists of a p-n junction in the form of highly doped boron (Br) doped p+-type strips, 16-20 µm wide over a 285 µm thick n-type silicon bulk. The p+ strips are capacitively (AC) coupled to 1 µm thick Al strips, 16-22 µm wide, which collect the charge from the p+ strips. The Al strips are in turn bonded to the readout electronics. The material used between the Al and p+ strips are SiO_2, to avoid a direct (DC) coupling between the high voltage bias and the sensitive readout electronics. There is an n+ layer on the other side of the bulk to ensure good ohmic contact to the 1 µm layer of Al which is connected to the positive bias voltage (in the range 0-500 V).

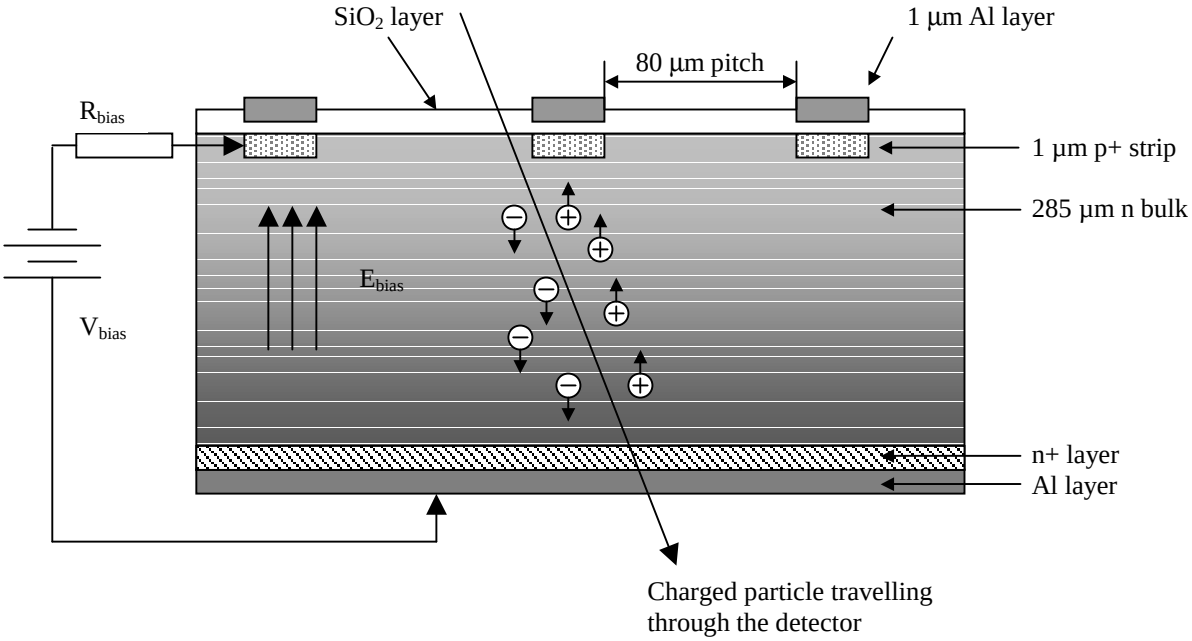

Figure 3.5: Cross-sectional view of a silicon strip detector

When a voltage is applied to the Al layer, this sets up a reverse bias over the p-n junction, creating a depletion zone in the bulk material. When a charged particle passes through the bulk, it ionizes atoms in its path, which causes electron-hole pairs to be created. Some of these recombine before they can be collected, but most of the charge should be picked up by the Al-strips of the detector. The overall charge collection efficiency of the detectors is specified to be at least 90 %.

[2] Minimum ionizing particle

To obtain a good signal to noise ratio, a low leakage current through the p-n junction is required. At the testing stage, a maximum leakage current of 6 µA at 150 V is specified. However, during operation at LHC, the high radiation doses will alter the properties of the silicon, giving rise to a higher leakage current, and thus higher noise. To minimize this effect, the system is cooled down, as the leakage current on irradiated modules has been shown to double for every 7 °C. Due to radiation, the n-bulk will over time change to p-type, a phenomenon known as type inversion. Since the strips are p+, it will still be possible to maintain a depletion region, but at a higher bias. After 10 years of LHC operation at planned luminosity, a bias of 350 V will be required.

3.2 Baseboard

The VHCPG (very high thermal conductivity pyrolytic graphite) baseboard is the central element in the SCT module, sandwiched between two pairs of detectors. Its task is to support the detectors mechanically, to transfer heat from the detectors and the readout electronics to the cooling pipes of the SCT. The latter will be gradually more important as the leakage currents increase due to radiation damage. The combination of these two capabilities is very important, as leakage currents have been shown to increase rapidly with temperature. It is also important for the baseboard to have the lowest possible mass, to minimize the total weight of the SCT detector. In figure 3.6 below, a baseboard is shown.

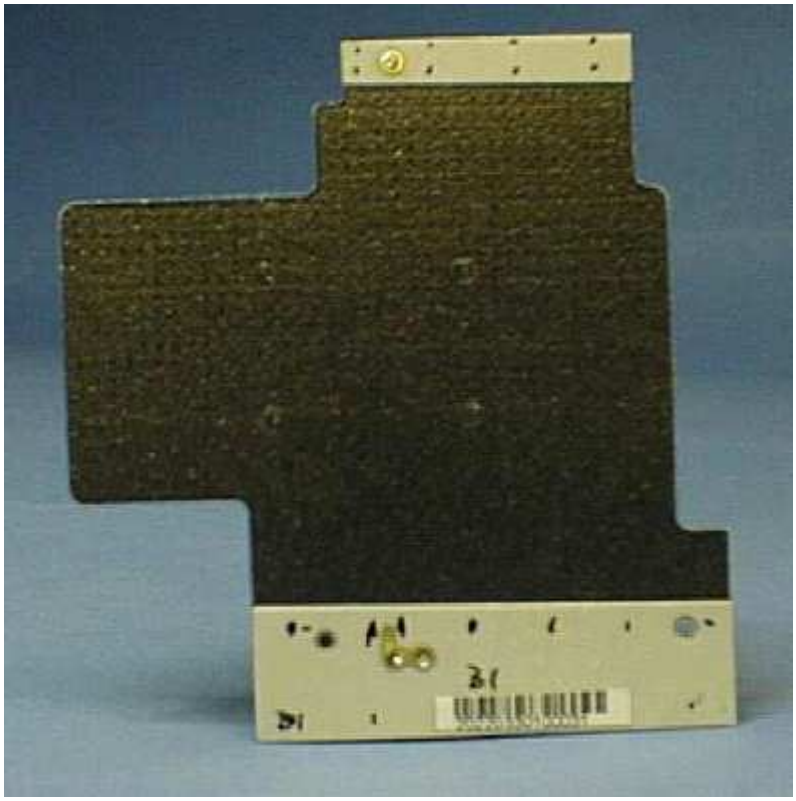

Figure 3.6: SCT barrel module baseboard [14]

The baseboard has a thickness of 380±15 µm, and has an in-plane thermal conductivity up to 1850 W/m/K at 20 °C, increasing 0.4 % for each degree lower in temperature it is operated. Orthogonal to the baseboard, the thermal conductivity is down to 6 W/m/K. This is due to the planar mosaic ordering of the carbon structures in the material, and ensures that heat gets transported efficiently. Attached to each side of the baseboard are BeO facings plates which connects the module to the environment. These have a thermal conductivity greater than 280 W/mK at 20 °C and are 250 ±10 µm thick. The BeO facings will be connected to the cooling pipes in the SCT detector. [14]

3.3 Hybrid

The hybrid is a four layer Cu/Polymide flexible circuit placed on a carbon-carbon bridge, carrying the front-end readout electronics for the silicon strip detectors. It consists of two main hybrid sections, each 74.6 x 21.3 mm connected with a 8.4 mm long interconnect (to wrap it around the module), and a 35.4 mm long connection to the pigtail cable. The thickness of the hybrid without components is 1.1 mm, with the highest components giving it a total height of 2.6 mm. Shown on figure 3.6 is a schematic drawing of a hybrid. 'Build-up' technology is used for constructing the hybrid, making it small and light and eliminating all cablings.

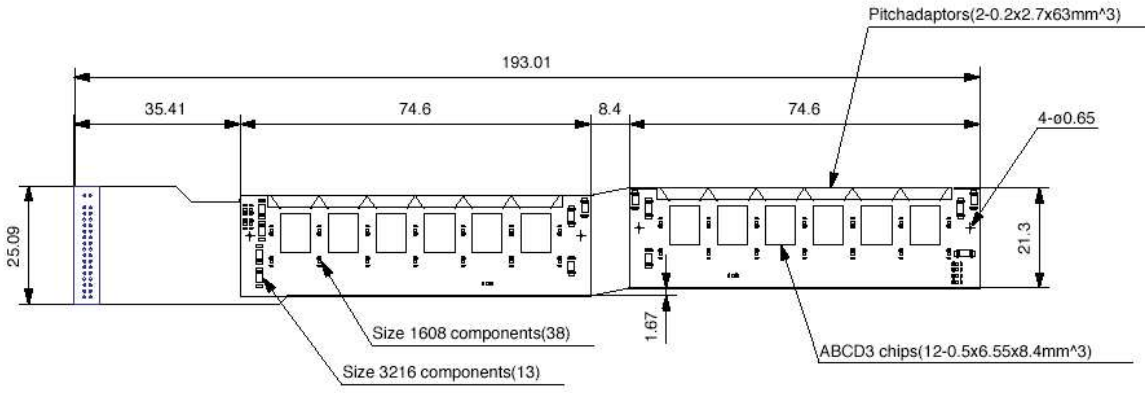

Figure 3.6: SCT barrel hybrid [15].

When mounted on the module, the hybrid forms a mechanical bridge over the detectors to avoid mechanical, electrical and thermal contact which could cause damage to the detectors.

The front-end electronics is placed on the hybrid, this means that there are 12 ABCD3T ASICS[3], six on each side, wire bonded to 128 strips each. The ABCD3T ASIC is described in section 3.4. The ASICs dissipate a power of 6 W nominally, which makes the thermal conductivity of the hybrid important. For this purpose, the Cu/Polymide flexible circuit is mounted on a carbon-carbon bridge with high thermal conductivity,

[3] ASIC – Application Specific Integrated Curcuit

700W/m/K. To connect the chips directly to this thermally conducting bridge, the flexible circuit has thermal “through-holes” with diameter 300 µm, vertically connecting the backside of the chips directly to the carbon-carbon bridge. When a circuit is glued to the carbon-carbon bridge, the through-holes are filled with electrically conductive adhesive, which gives each hole an estimated thermal conductivity of 40 W/mK. There are a total of 17 such holes per chip.

While the pitch between the silicon strips of the detector is 80 µm, the corresponding pitch between ASICS read-in channels is 48 µm. In order to make it possible to do a simple parallel wire-bonding when connecting the detector to the ASICS, a pitch adaptor is placed in front of the ASICS. It is made with thin-film technology, using 1-1.5 µm thick aluminium deposition on a glass substrate.

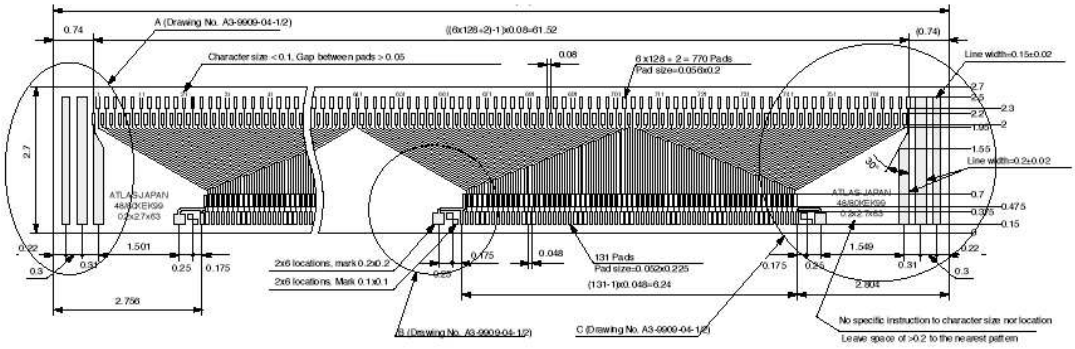

Figure 3.7: Pitch adapter connecting detector strips to input pads of ABCD3T, with the detector bonded to the upper side and ASCIS to the lower side [15].

The hybrid is also able to measure the temperature of the module, using thermistors with R_{25}=10 kΩ ± 1 % and B = 3435 K ± 1 %. The temperature can then be calculated from the measured resistance of the thermistors from the formula: [15]

$$R = R_{25}^{\left[B\left(\frac{1}{T}+\frac{1}{298}\right)\right]} \quad \text{(T is given in absolute temperature)} \qquad (2)$$

3.4 Front-end electronics

The front-end electronics on the hybrid consists of 12 ASICS of the type ABCD3T, interconnected via the circuitry of the hybrid. During development of the chip, low mass, low power consumption and high radiation tolerance have been important design goals.

The placing of the chips on the hybrid can be seen on figures 3.1 and 3.6. Each chip is wire-bonded to clock and control connections as well as the 128 outputs of the pitch adapter, which in turn is bonded to 128 silicon strips. A description of the chip and its operational principles will be given in the following sections.

3.4.1 The ABCD3T ASIC

The ABCD3T is an ASIC (*Aplication Specific Integrated Circuit)* designed for the purpose of digital readout of silicon strip detectors in the ATLAS SCT. Figure 3.8 shows the chip with dimensions and connections.

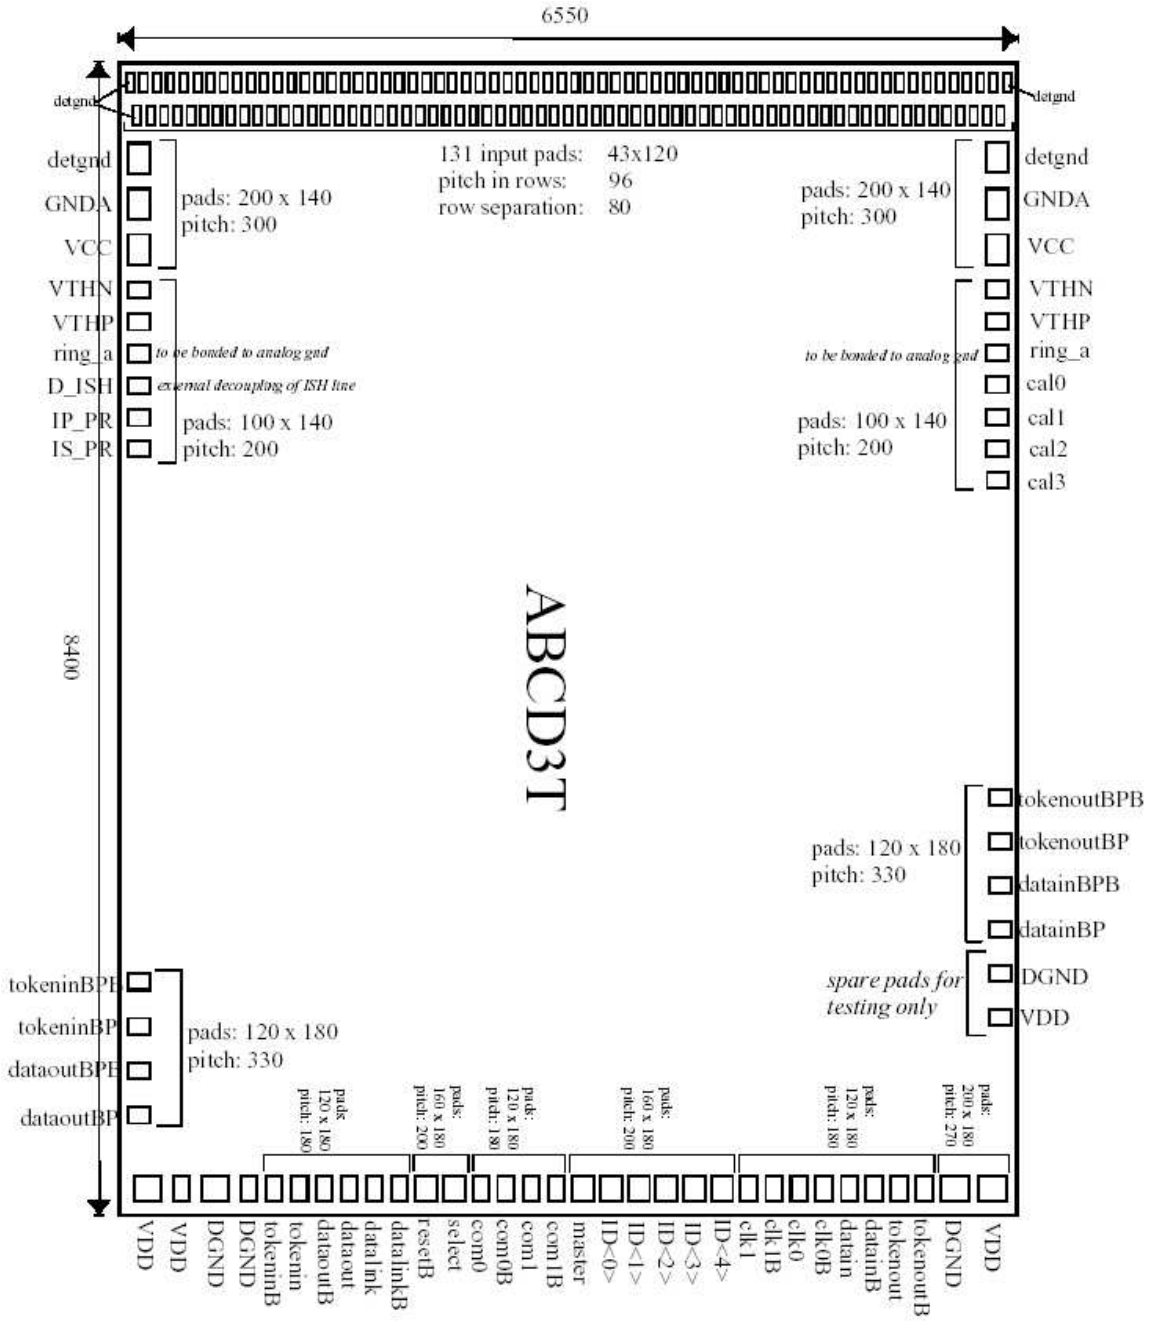

Figure 3.8: ABCD3T ASIC. Clock, and control connections can be seen on the lower side of the figure, power and tokens are on the left and right sides, while the connection to silicon strip are on the upper side of the figure. [16]

3.4.2 Command and clock reception for ASICS

Each chip has a dual set of clock and command inputs to increase the fault tolerance of the system. Which clock and command input to use, is set by the external SELECT signal.

3.4.3 Addressing scheme

To uniquely identify each ABCD3T on the hybrid, the chip has a 6 bit address. The MS bit is for historical reasons set to 1. ID(4) is tied to the SELECT input (see previous paragraph), selecting which clock/command input to use, while ID(3:0) is used to identify the twelve chips on the module, referred to as the geographical address. ID(6:0) is for each chip on the module wire bonded to a unique set of logic levels, to give them a unique address in case a chip needs to be individually addressed.

3.4.4 Signal flow

The ABCD3T has 128 inputs, one for each silicon strip. When a particle traverses the silicon, electron-hole pairs are made, and due to the field, the holes travel to the p+ strips, which are capacitively coupled to aluminium strips wire-bonded to the input of ABCD3T. The signal flow through the chip for the picked up charge from a silicon strip is shown in the diagram below, and each stage will be given a description in the following sections. In the chip, there are 128 such paths, each one to be connected to a silicon strip.

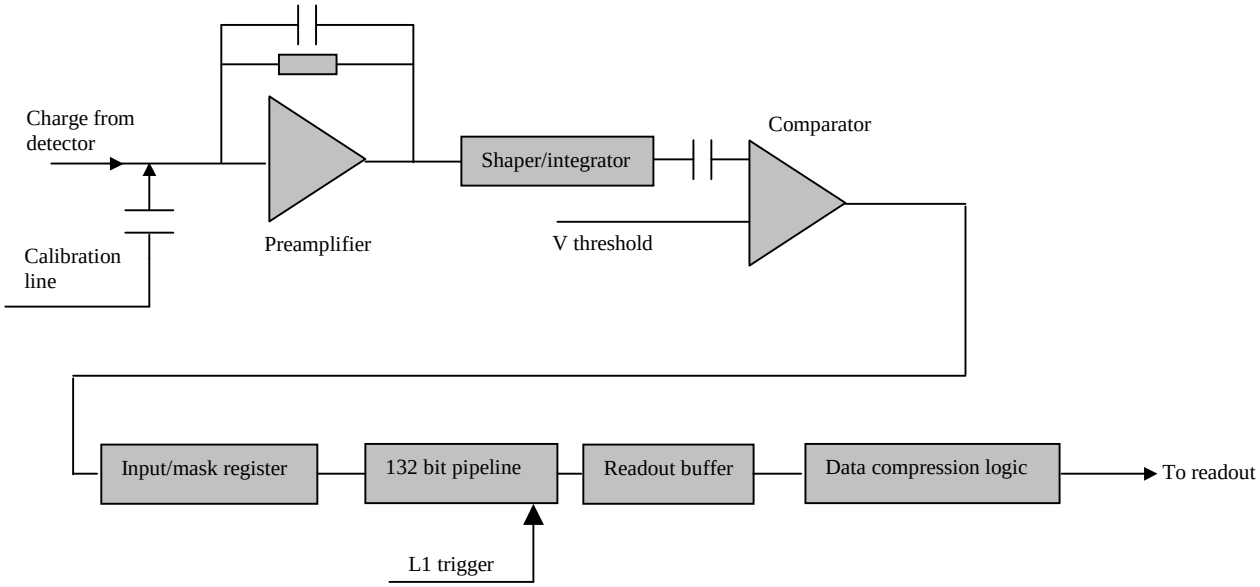

Figure 3.9: Signal flow for a single channel in the ABCD3T chip.

Calibration circuit

In the chip, each channel is connected to a 100 fF capacitor. At a certain control command, a pulse is sent from the capacitor, simulating a strip being hit. There are a total of four such calibration lines in the chip, connected cyclic to the channels. The calibration signal is selectable in the range 0-160 mV (0-16 fC) in steps of 0.625 mV (0.0625 fC). The ability to send in a known charge and see the chip's response to this is a very useful feature of the chip, used extensively during the electrical testing of the module.

The calibration circuit has a register associated with it, the strobe delay register, which determines the relative delay between the rising edge of the clock input and the rising edge of the calibration strobe.

Preamplifier/Shaper

When the charge is first picked up, it is sent to a preamplifier. The amplified signal is then integrated by a shaper with peak time 20 ns, converting the pulse to a voltage level with gain ~50 mV/fC for nominal shaper current of 20 µA. The maximum noise allowed is 1500 electrons rms for an unirradiated module, 1800 for irradiated. The linearity should be better than 5% in the range 0-4 fC. Before the pulse is sent to the comparator, it is sent through an AC coupling to remove DC offsets in the pulse level.

Comparator

The comparator compares a threshold voltage level to the input signal to determine if it should be recorded as a hit. If the input signal is higher than the threshold level, it outputs a logic 1 signal. The threshold voltage is applied from an internal DAC, set by the control software, and is common to all channels in the chip. The setting range from 0 fC (0 mV) to 12.8 fC (640 mV), with nominal setting 1 fC, adjustable in steps 0.05 fC (2.5 mV), with absolute accuracy 1 %. However, to compensate for channel-to-channel variations, there is for each channel a TrimDAC, allowing a threshold correction individual for each channel to be set. This is selected with a four bit register with four selectable ranges, ranging from 0-60 mV to 0-240 mV. This feature is crucial for compensating channel-to-channel variations, and will grow more important with time, as the offset spread in channels is expected to increase after radiation.

Input/Edge detection register

The input register latches the incoming data, delivering a well defined pulse width for further processing, and thus marks the beginning of the digital readout, see figure 3.11. The edge detection is able to detect a high to low transition, and if edge detection is enabled, only high to low transitions will be recorded as hits.

Mask register

The mask register enables the channel to be disabled in case of problems related to the channel. If a channel is masked, logic 0 will be sent independent of input. The mask register has a mode register, which can be used to select the output of masked channels to

be logic 1 instead of 0, convenient for sending test patterns through the rest of the digital readout chain.

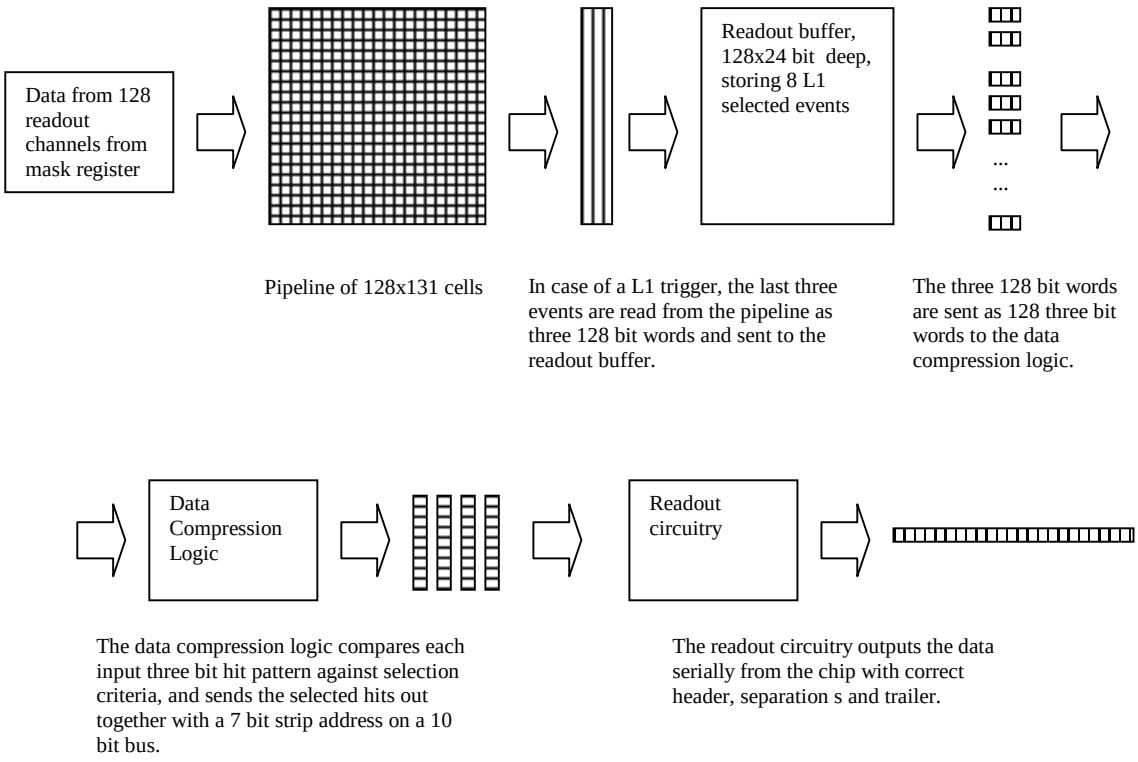

Figure 3.10: Digital readout stage of ABCD3T.

Pipeline

The next step is a binary pipeline, the first step in figure 3.10, realised as a multiplexed FIFO circuit. The reason for having a pipeline is to delay data long enough for the trigger system in ATLAS to determine whether an event is interesting or not, and thus sending a Level 1 trigger (see section 1.7 for a description of the trigger system in ATLAS).

An array of 12x12 dynamic memory cells is controlled by 12 non-overlapping clock signals, see figure. In each clock-cycle 12 cells are shifted, providing an effective delay of 12 x (12-1) = 132 clock cycles (3300 μs). When a L1 trigger arrives, the last three bits of the pipeline are read into the readout buffer, pending further readout. These three bits correspond to the three beam crossings centred on the L1 trigger time. All other data will continuously be shifted through the pipeline.

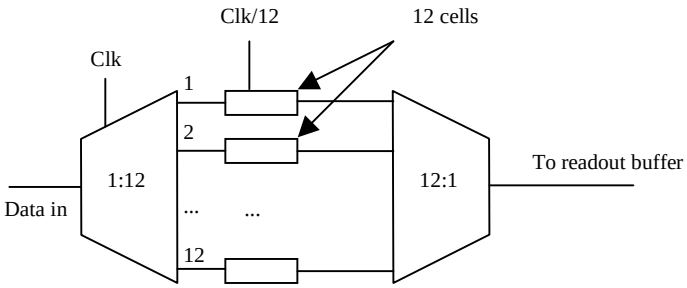

Figure 3.10: Pipeline architecture

Readout buffer

The readout buffer contains the events tagged as interesting by the L1 trigger. The buffer is 24 bits deep, which means that it can store data from up to 8 L1 triggers, providing a buffer for variations in L1 frequency. This satisfies the ATLAS specification of maintaining less than 1 % data loss at a L1 trigger rate of 100 kHz and a strip occupancy up to 1%. The register also contains an accumulator register, which tells whether the channel has been hit since last reset or not, and an overflow counter to track the number of events that have been overwritten in the buffer. A “data available” flag is set in the register when there are data available for readout. Two cyclic read and write pointers are used for addressing the register.

Data compression logic (DCL)

In an event, it is anticipated that very few of the 128 channels read out by the chip will contain hits. To exploit this fact, a fast and simple data compression logic is implemented in the chip. When the readout buffer signals that data is available, the data compression logic reads the L1 centered three bit data from all the strips as three 128 bit words, and then converts it to 128 three bit words, see figure 3.11. The data compression logic then compares the three bit words against one of four selectable criteria, determining for each channel if the three bit word should be recorded as a hit or not. The criteria are shown in table 3.1 below.

Criteria	**Hit Pattern**	**Usage**
Hit	1XX or X1X or XX1	Detector alignment
Level	X1X	Normal Data Taking
Edge	01X	Normal Data Taking
Test	XXX	Test mode

Table 3.1: Data compression criteria [16]

If the hit pattern of a channel fits the data compression criteria, the three bits are transmitted together with a seven bit address, identifying the strip hit, on a 10 bit bus. The data compression logic then waits for a ‘next’ signal from the readout logic, before sending data. If a strip adjacent to the previous strip also is hit, no address is sent, but the

signal line ‘adj’ is set to logic high. If no more hit data is present, the ‘end’ signal is set to logic high in the DCL register.

By using this technique, the amount data is greatly reduced, as well as the 128 bit readout bus being reduced to a 10 bit bus.

Readout circuitry

The readout circuitry of the chip converts the 10 bit bus data into a serial data stream with prescribed header, separators and trailers. It sends its data out when receiving a token, see next paragraph.

3.4.5 Data routing on hybrid

When hit data is sent from the hybrid, it is sent on two one bit lines, streams, one for each side of the detector. This is necessarily done with readout in serial fashion, with a technique called token passing. Shown below is a figure of the twelve chips, showing interconnections between each chip, used for token and data passing.

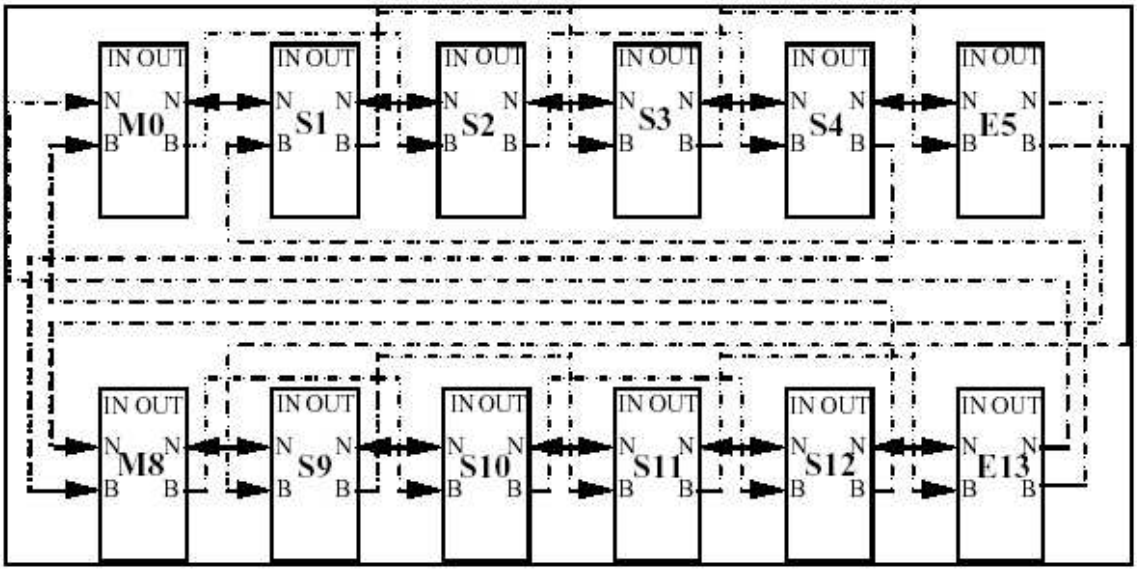

Figure 3.12: Normal token- and data flow between chips on the SCT barrel module, each chip is connected to its nearest neighbour. The dotted lines shows alternative token and data routings. M indicates Master, S slave and E end. [16]

A chip can be configured either as Master, Slave or End, which determines its role in the readout chain. When the chip configured as master receives a L1 trigger, its readout circuitry outputs a header containing the L1 trigger number and beam crossing number (see the next section for the actual bit format), followed by the hit data from the data

compression logic on the one-bit readout line. A few clock cycles before the last bit has been sent, it sends a token to the slave chip connected to it. The slave chips sends its data to the master chip, which appends the data to its own hit data. The slave chip then passes on the token to the next slave chip, which sends its data back through the other slave chip to the master chip. This token and data passing is repeated until the chip configured as End receives the token. This chip does not pass the token further on, but instead appends a trailer indicating end of readout, distinguishable from other data sent from the chips. When the master chip receives this trailer, it can start processing the next event.

Reading out the chips in serial fashion, makes the system vulnerable for single chip failures, since this could block the readout of working chips later in the readout chain. For this reason, alternative routings of token and data is available. If a chip failure should occur, for example in S2, the chips can be configured to route their token and data as shown on the figure, where the B connections marks the alternative paths.

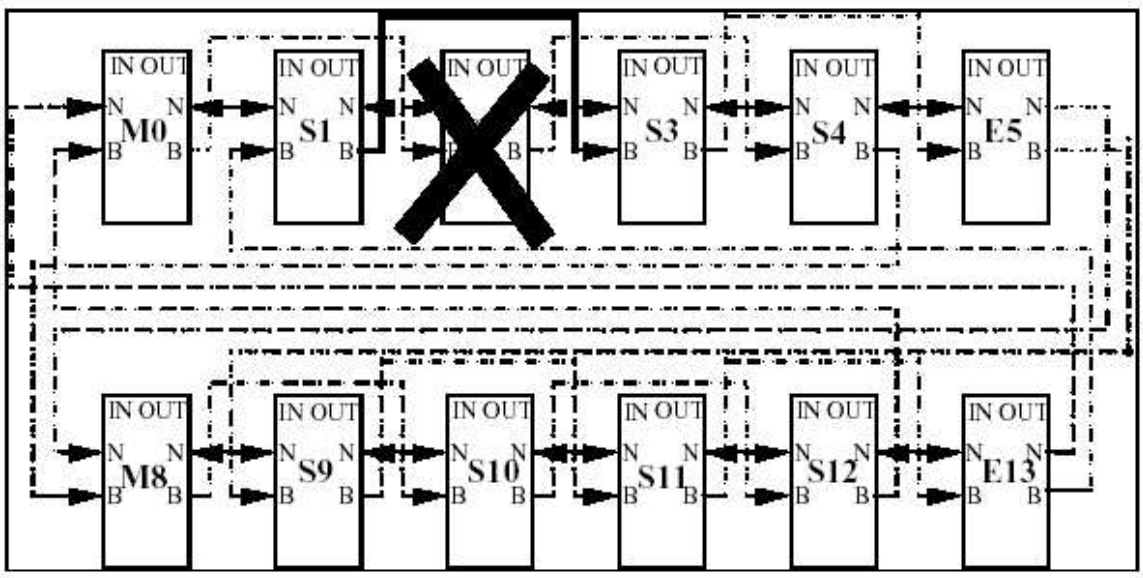

Figure 3.13: Data-/token passing in case of S2 chip failure [16]

In the case of two adjacent chips failing, the rest of the readout chain will be blocked from readout.

3.4.6 Readout protocol

The output to off-detector electronics for further processing is sent by the master chip only. The format of this bit stream is:

```
<11101><0><nnnn><bbbb,bbbb><1><data_block_1>...<data_block_n><1000,0000,0000,0000>
```

with <nnnn> being the number of L1 triggers recieved since last reset, needed to identify which L1 trigger the data should be associated with, and bbbb,bbbb the number of beam

crossings since last reset. There is one data block for each chip in the readout chain, sending one of five types of output data:

- **Physics data:** There are three types of physics data:
  1) **Isolated hit data-packet:** This type of packet is used to send the hit information from a hit channel on a chip when none of its neighbouring channels have been hit:

     `<01><aaaa><ccc,cccc><1><ddd>`
     Where aaaa is the chip address, ccc,cccc the strip address, and ddd the hit pattern.

  2) **Non-isolated hit packet:** This type of packet is used to send data from a group of 2 or more adjacent channels being hit. Only the channel address of the first channel is sent. This reduces the amount of data sent over the 1 bit dataout line, as the 7 bit addresses need not to be transmitted.

     `<01><aaaa><ccc,cccc><1><ddd><1><ddd>...<1><ddd>`
     where the fields have the same meaning as in 1), with <ddd> being the hit patterns from the adjacent channels.

  3) **No Hit data:** If the chip has recorded no hits corresponding to the received L1 trigger,

     `<001>`
     is sent.

- **Configuration/error data**

  4) **Configuration data:** A chip could be configured not to send data, a state called Send_ID mode. The chip will then output a bit pattern

     `<000><aaaa><111><cccc,cccc><1><cccc,cccc><1>`
     where aaaa is chip address, and c... is the contents of the 16 bit configuration register.

  5) **Error data:** This type of data is sent when the chip has detected an error, e. g. readout buffer overflow, meaning that L1 events are being overwritten. Format:

     `<000><aaaa><eee><1>`
     with <aaa> being chip address and <eee> one of three possible error codes.

Shown on figure 3.14 is an oscilloscope image of a L1 (`<110>`) trigger followed by a dataset sent from one stream on a module. The lower line is the L1 trigger (the signal is taken from the DATO output from CLOAC (see chapter 4), and shows 1 as low and 0 as high. The second line is taken from the MA output of the MuSTARD (also see chapter 4),

showing the header with L1 number followed by beam crossing number. Following this are six <001> signals, indicating that none of the chips have detected hits.

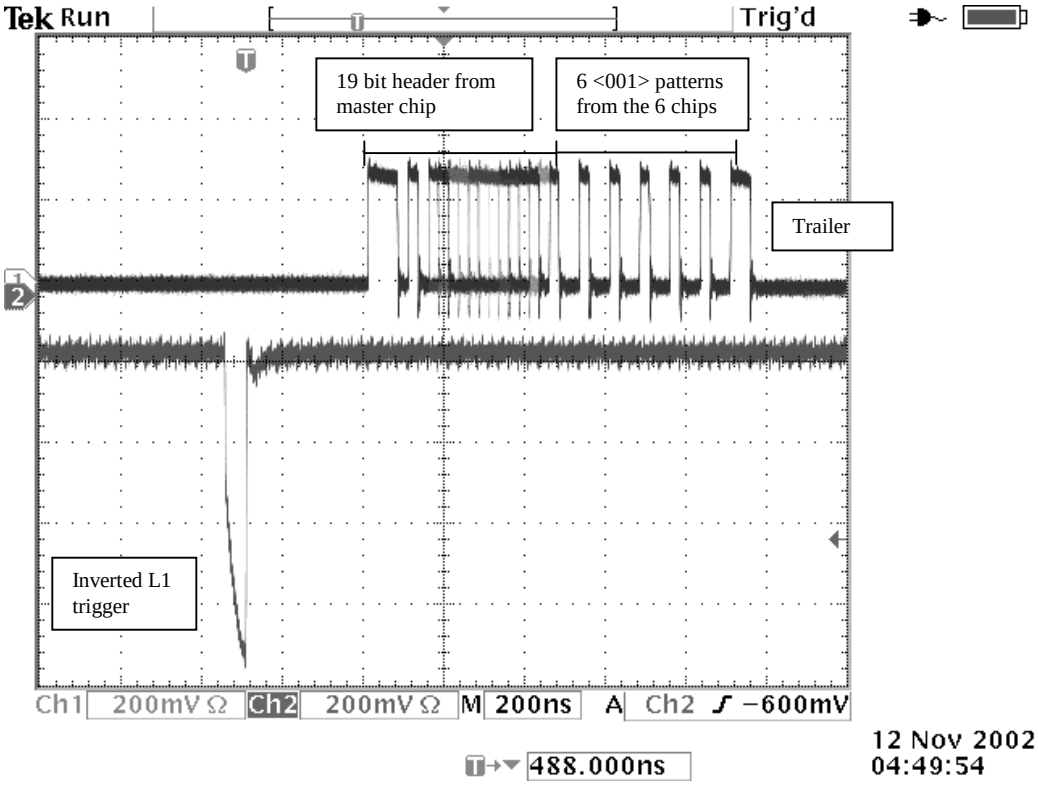

Figure 3.14: L1 trigger sent to module, followed by data output from module.

3.4.7 Control protocol

There are two main classes of commands: L1 trigger and control signals.
L1 trigger is, as explained earlier, the signal indicating that an interesting event has been identified by the trigger logic in ATLAS, and readout of this event is necessary. This is the shortest command, only three bits long, and will by far be the most issued command, with an upper limit of 100kHz.
Control signals are further divided into two classes; fast (short) and slow (long) commands, as shown in the table 3.2.

Type	**Field 1**	**Field 2**	**Field 3**	**Description**
Level 1	001			Level 1 trigger
Fast	101	0100/0010		Soft reset/BC reset
Slow	101	0111	21-156 bit long command string	Slow control command, 9 types

Table 3.2: ABCD3T commands

Fast commands

The fast commands are seven bits long, and are meant to be issued during data taking mode. There are only two fast control commands implemented, Soft Reset and BC reset. In LHC, these commands will be sent to the chips at regular intervals during periods of time when no L1 triggers are sent.
Soft reset clears all internal registers and tokens, and sets the chip in a no-data state while leaving the configuration of the chip unaffected.
BCReset, Beam Crossing counter Reset, resets the beam crossing counter, and otherwise leaves the chip unaffected. In ATLAS, this will be sent with a rate of 11.2456 kHz, the rate at which a single bunch crosses the LHC.

Slow commands

Slow commands are mainly commands for configuring the chip. They are longer bitstreams, ranging from 27 bits to 155 bits long commands. While a slow command is recieved, L1 triggers can not be sent to the chips. Slow commands includes all the commands necessary for configuring the chip, as well as commands for issuing calibration pulses.

3.4.8 Chip startup sequence

The startup of a session could typically look like this:

1) Send command to load the configuration register with appropriate settings
2) Send a command to load mask register
3) Send a series of commands to load the DAC registers and delay registers
4) Send a command to place the chip into data taking mode

The chip is then in a state to receive L1 trigger command and send data.

3.4.9 Power requirements

In table 3.3, the minimum, nominal and maximum voltage, current and power specifications for ABCD3T is given.

	Min	**Nominal**	**Max**	**Absolute max**
Analogue DC supply	3.3 V	3.5 V	3.7 V	5.5 V
Digital DC supply	3.8 V	4.0 V	4.2 V	5.5 V
Analogue DC current	600 mA	900 mA	1200 mA	1320 mA
Digital DC current	360 mA	400 mA	450 mA	470 mA
Power Consumption	3.54 W	4.75 W	6.0 W	7.0 W

Table 3.3: Voltage, current and power specifications for ABCD3T.[16]

Chapter 4: The setup

The setup for electrical testing of SCT modules is shown in figure 4.1. The setup consists of these main parts:

- An environmental chamber for temperature cycling and cooling of the module.
- CERN-standard VME crate for hosting VME modules.
- CLOAC VME module for clock and fast commands.
- SLOG VME module for slow commands.
- MuSTARD VME module for reading data from the module.
- SCT LV-3 VME module for powering the front-end electronics, some commands.
- PPR2 (Patch Panel Replacer) fanout passive VME card.
- SC2001 v. 2.0 support card for SCT module.
- HV supply for biasing the detector.
- A Windows NT computer with MXI interface, running SCTDAQ under ROOT.

In the following sections there will be a presentation of each of the parts. For a detailed description of the SCT module, see chapter 3. For a description of optimization of the system, see chapter 6.

The setup described in this chapter is not the final QA setup. The power supplies can only support one module at a time, making parallel testing impossible. To enable us to do parallel testing, LV and HV supplies which can support up to six modules (like SLOG, MuSTARD and CLOAC), is under development, and these will probably be used in the final setup.

4.1 Environmental chamber

An environmental chamber is needed for keeping the module at the right temperature, and for cycling temperature as required by the QA procedures (see chapter 7).

In the setup, we use a Design Environmental BS125-40 environmental chamber. It is operated with an FGH P1000 controller, programmed to do the required temperature cyclings. Also attached to the chamber is a nitrogen flow controller, type Vögtlin V100-140.05, providing nitrogen for low humidity.

The chamber had to be modified electrically to eliminate noise, see chapter 6.

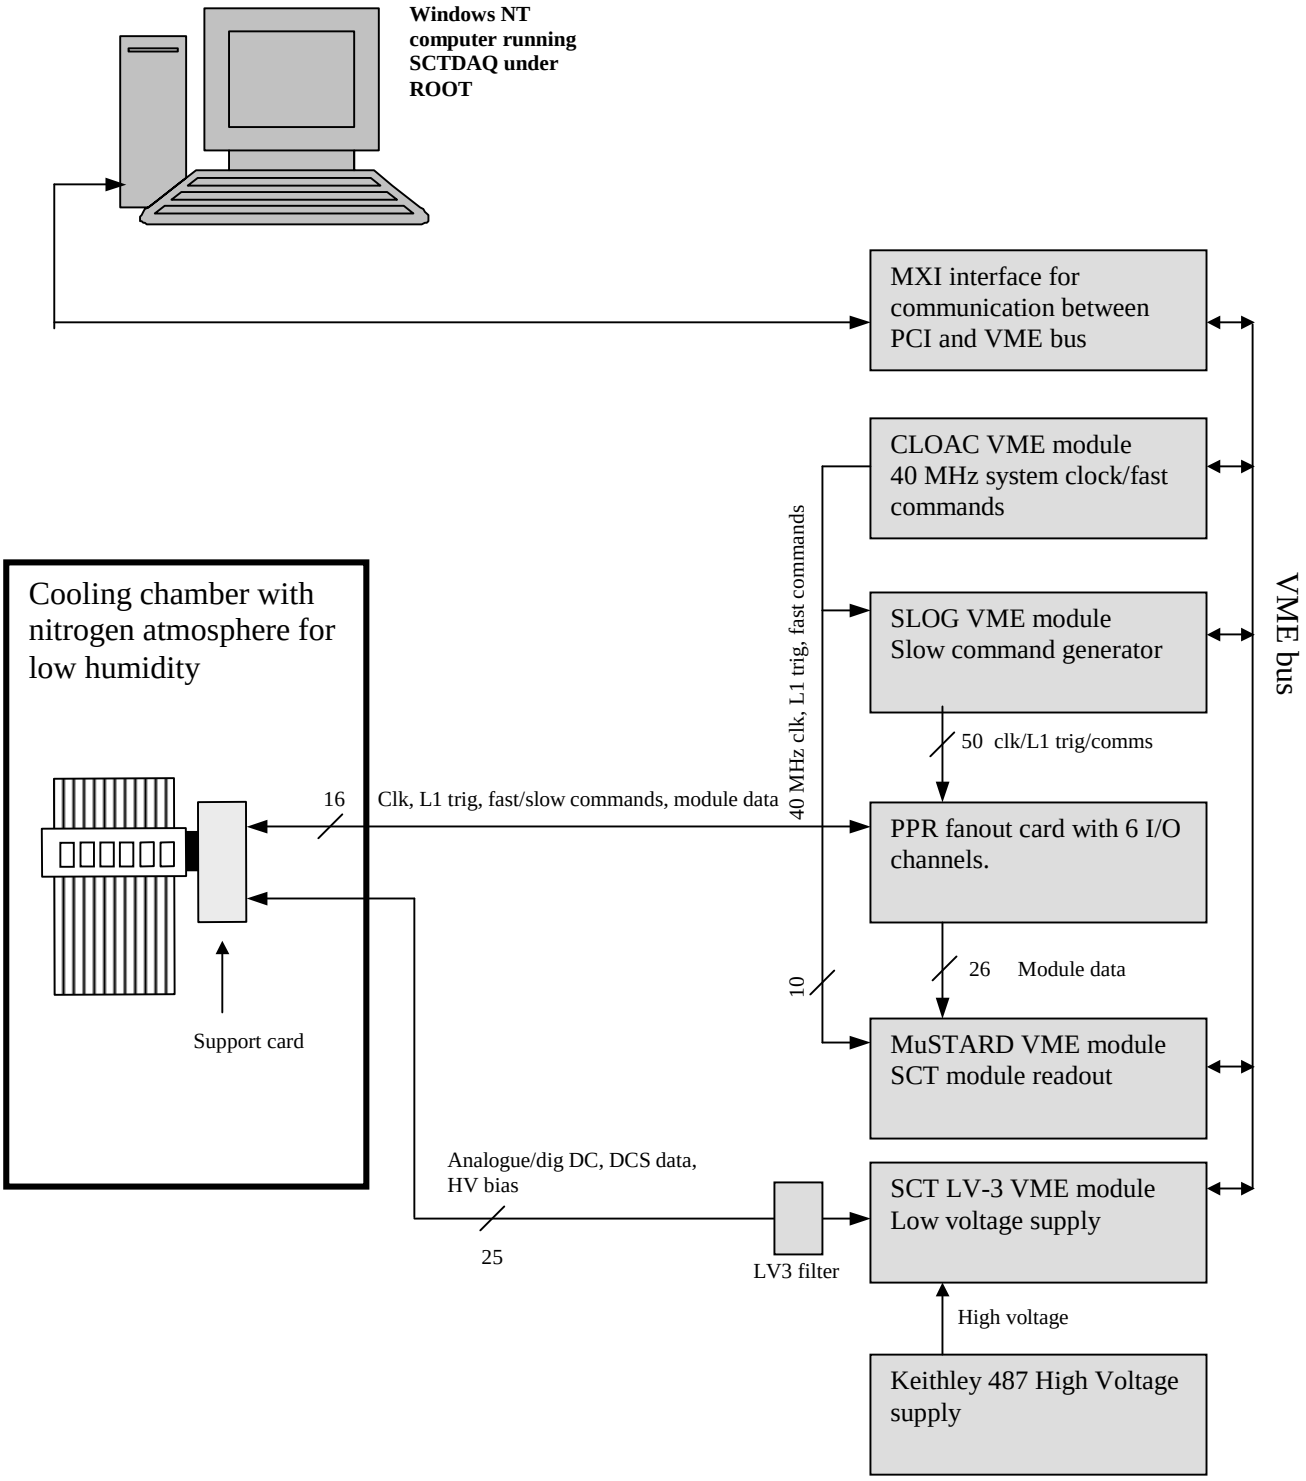

Figure 4.1: Schematic view of test setup

4.2 The VME system

The system that generates control commands, level 1 triggers and reads physics data from the module in the testsetup, is based on VME modules. This acts as an interface between the module and the computer.

VMEbus (*Versa Modular Eurocard bus*) is a general purpose bus standard, designed by Motorola, Signetics, Mostek and Thompson CSF in the eigthies, defined by the IEEE 1014-1987 standard. The VME modules used in the setup are placed in a VME crate, shown in figure 4.2.

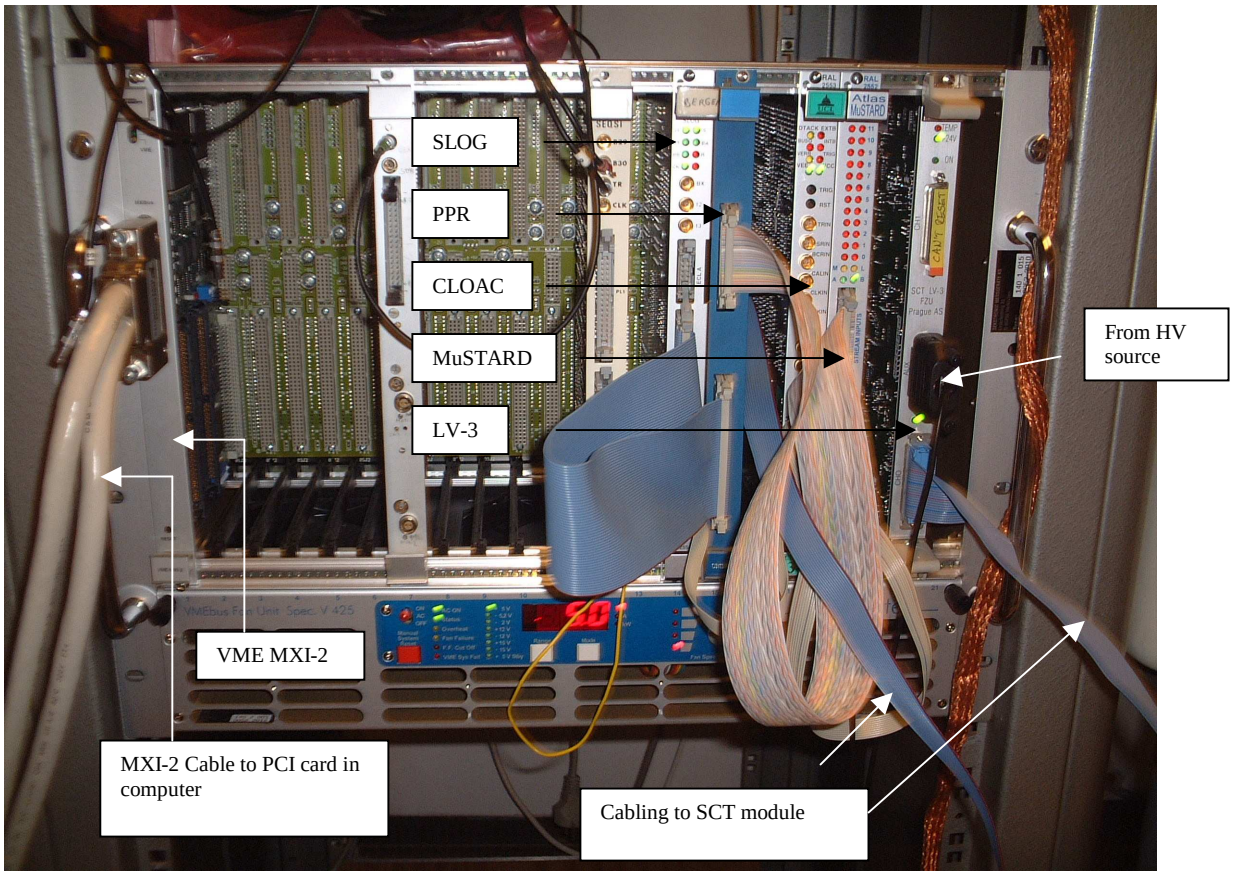

Figure 4.2: VME crate with VME modules for readout and control of SCT module.

The crate provides power, cooling and mechanical support for the modules. The modules connects to the VME bus with two 96 pin plugs (P1 and P2) and one 10 pin plug (JAUX) on the backplane of the crate. The JAUX is an enhancement of the standard VME, called CERN-standard VME. The JAUX provides +5.0 V, -5.2 V and can also be configured to supply ±12.0 V.

The VME modules used in this setup are all single width, 6U, using 24 and 32 bit transfers.

The bus system consists of four sub-buses; the data transfer bus, the arbitration bus, the priority interrupt bus and the utility bus. A typical VME bus data transfer consists of an arbitration cycle (to gain bus control), an address cycle (to select the correct register), and the actual data cycle. The bus is designed to use a completely memory mapped scheme, which means that each module can be seen as an address or a block of addresses.

To connect the VME bus to a personal computer, one has to use a bridge between the PCI bus of the computer and the VME bus. The bridge used is a standard called MXI-2, *Multisystem eXtention Interface* bus, and is actually a bus standard itself, designed to combine different buses into one, in this case VME and PCI. An MXI compliant PCI card, the PCI-MXI-2, is installed in a PCI slot in the computer, and the VME-MXI-2 VME module is installed in the VME crate. This is a single slot 6U VME bus device. The two devices connect via the MXI-VXI 2 cable. Both the MXI module and the cable can be seen on figure 4.2 on the previous page. Integrated in the cable is a dual-channel DMA (direct memory access) controller with standard interfaces for VXI, VME, MXI, PCI and PXI. By using this controller to transfer data and commands to and from devices, the controller frees up processor time in the computer, leaving the processor available for other tasks, such as data analysis and presentation. Each bus (VME and PCI) sees the other one as a virtual memory map, and can write directly to the addresses in this virtual memory space, creating a transparent connection between the two buses. This generates very little protocol overhead, and enables almost all of the MXI bus to be used for data. [1,17]

Configuring the VME/MXI bus

To communicate with the VME bus via the MXI interface, the PCI-MXI-2 card uses a driver, the NI-VXI system-level software. The software includes a Resource Manager, an interactive configuration and troubleshooting program, libraries of software routines for test and measurement (T&M) programming, as well as several debugging tools.

Shown in figure 4.3 is a screenshot from T&M explorer, displaying all installed VME modules. It also shows the control hierarchy, with the VXI Frame 1 being connected to the PCI-MXI-2 through the MXI-2 bus. There is an exclamation mark on VXI Frame 1 due to the fact that our crate is VME, not VXI. Each module is given a pseudo logical address to uniquely identify it in the system, as well as configured with the correct physical slot number and address space. Shown in table 4.1 are the settings for the four VME modules. These settings have to correspond to physical jumper/switch settings on the module, as well as settings in the compiled SCTDAQ software running on the computer (in the file *sct_hardware.h*). In this setup, the addresses were set to match the compiled SCTDAQ, to avoid the need for a recompilation of the software package.

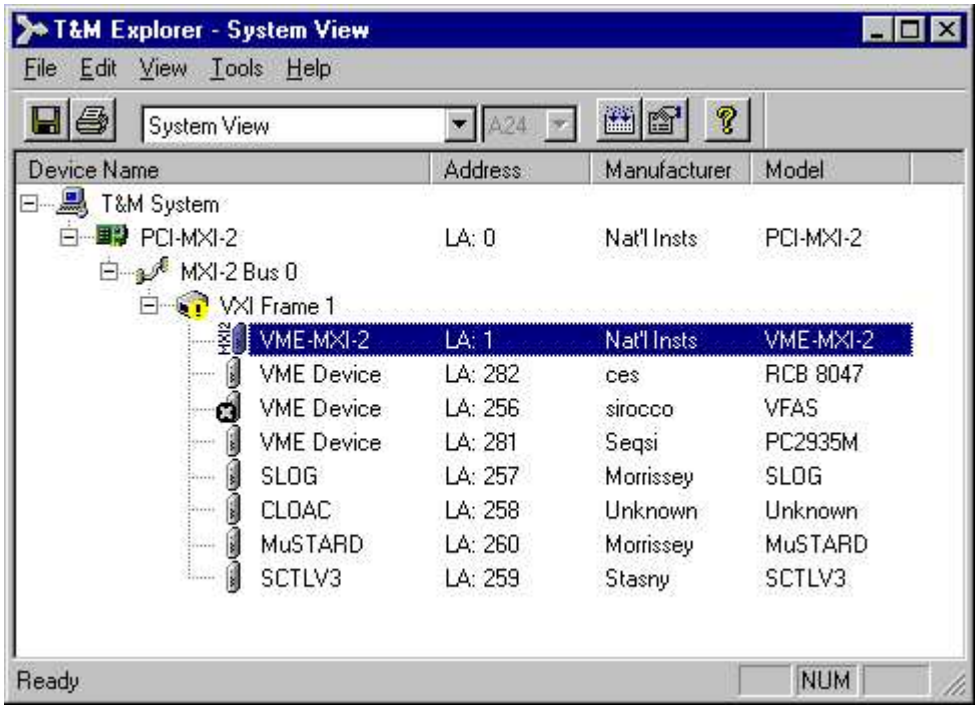

Figure 4.3: Screenshot from T&M explorer showing system configuration. Other VME modules can also be seen, not part of this setup.

VME module	Slot	Address space	Address range	Pseudo-logical address
SLOG	14	A32	0x1800000-0x181FFFF	257
CLOAC	17	A24	0x880000-0x8800FF	258
MuSTARD	18	A32	0x1900000-0x19FFFFF	260
SCTLV3	20	A24	0xFF0000-0xFFFFFF	259

Table 4.1: Addresses of the VME modules.

On system start up or after a crash, it is important to run the Resource manager, resman.exe. This correctly configures the system and maps all the addresses according to the configuration done in T&M explorer. If this is not done, the system may hang.

4.3 CLOAC

The CLOAC (*CLOck And Control*) is a 6U VME module designed for the SCT system- and beamtest. In the final LHC-ATLAS system, there will be a LHC-wide TTC (Timing, Trigger and Control system) for all detectors. In the test setup, it is the task of the CLOAC to provide some of the signals that in ATLAS will be generated by the TTC. This means that the CLOAC must provide clock, level 1 (L1) triggers, as well as control commands. However, although the CLOAC has slow command functionality, it is only used for fast commands in this setup, since the SLOG is a dedicated slow command generator (see section 4.4).

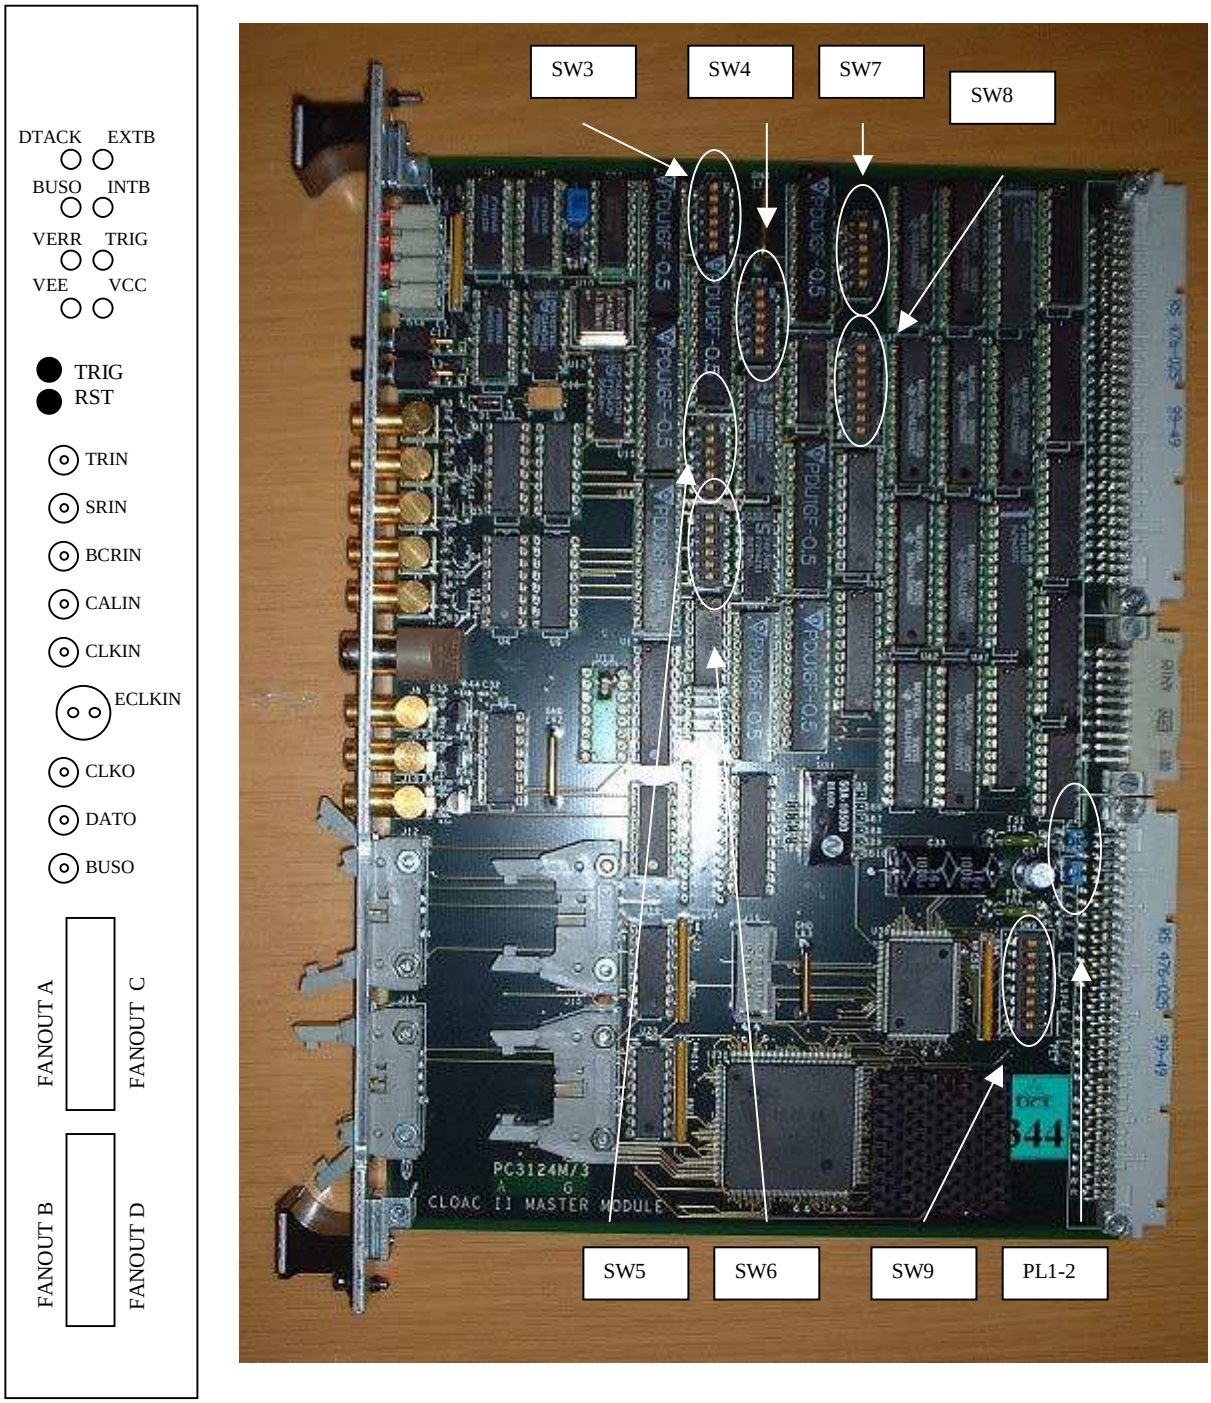

Figure 4.4: Front and side view of CLOAC VME module, showing relevant switches and jumpers, see table 4.2.

The CLOAC outputs a 40.08 MHz clock to the rest of the system, corresponding to the ATLAS BCO[4] clock. It is connected to the SLOG and MuSTARD (see figure 4.1.), sending clock, level 1 trigger and fast commands. [18]

The CLOAC can issue the following signals:

Hardware generated signals:

- **L1 trigger**: Single or repetitive with selectable frequency between 50 Hz and 600 kHz. It is also possible to send randomly spaced triggers with selectable average frequency between 12.5 Hz – 150 kHz.
- **Soft reset:** 7 bit fast command. Repetitive with selectable frequency between 0.005 Hz and 60 Hz.
- **BC (Beam Crossing) reset:** 7 bit fast command. Repetitive every 88.924 µsec. This corresponds to the time it takes for a single proton to traverse the entire LHC ring, known as the beam crossing rate, 11.2456 kHz.
- **INTBUSY –** internal busy signal for preventing conflict between commands.

Software generated signals:

- **Slow commands**: Limited to 64 bits.
- **Calibration sequence**: a calibration pulse followed by a L1 trigger after a programmable delay.

The CLOAC has the following inputs/outputs:

Outputs:

- **CLKO (J8):** LEMO NIM[5] connector outputting the 40.08 MHz clock
- **DATO (J9):** LEMO NIM connector outputting commands generated by CLOAC (mainly triggers and resets).
- **BUSO (J10):** LEMO NIM connector. Busy signal, this is an OR of the signals INTBUSY, EXTBUSY and VBUSY.
- **FANOUT A/B/C/D (J11-J15):** Clock and data out and EXTBUSY in on 10 bit IDC connectors. Busy is received from MuSTARD, indicating that its FIFO buffer is almost full, see section 4.5.

Inputs:

- **Five NIM LEMO + 1 ECL LEMO inputs for external clock, triggers and commands**. These are the uppermost inputs on the CLOAC. Since no external controllers are used in this setup, these inputs are not used.

Power requirements

The CLOAC uses the +5.0 V 10 A and –5.2 V 5 A power provided by the VME crate. The –5.2 V can be supplied by either J2 or JAUX. Upon recieval, the CLOAC had no

[4] BCO – *Beam Cross Over*.
[5] NIM – *Nuclear Instrument Module:* Logic level definition ‘low’ ~0.0V ‘high’ ~1.6 V

JAUX connector, so it had to be soldered on. To select JAUX as power supply, one had to close the jumpers PL1 and PL2, see table 4.2 and figure 4.4.

Address configuration

The CLOAC is an A24/D16 VME module, which means that it uses a 24-bit base address in the VME address space. The bits A16-A23 are selectable from the 8-bit DIL-7 switch SW9 on the CLOAC. The lever marked 1 is connected to A16. In the compiled version of SCTDAQ (the software controlling the VME modules), the CLOAC address is set to be 0x880000h. This could be changed, but a recompilation of the software is then needed. In our VME setup, this address was not in conflict with the existing setup, so SW9 were set as shown on figure 4.5, were switch ON corresponds to 0, giving and address of 10001000, corresponding to 0x880000h. It is also important that this address is the same as in the T&M explorer, so that the computer 'sees' the CLOAC at the right address.

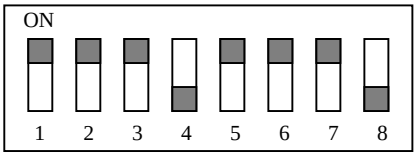

Figure 4.5: Example of setting of switch on CLOAC: SW9 Address A16-A24, with 8 as MSB and ON=0.

Most of the other DIP switches were left at their initial values. The switches are:

Name	**Value**	**Description**
SW1	-	Front panel, generates test trigger
SW2	-	Front panel, sends overall reset
SW3	111101	Sets CLK1 delay
SW4	110001	Sets CLK2 delay
SW5	110001	Sets CLK3 delay
SW6	110001	Sets CLK4 delay
SW7	111010	Sets CLK5 delay
SW8	00000000	Trigger window compensation delay
SW9	00010001	VME base address
PL1/PL2	Closed	Selects JAUX as –5.2 V supply
PL3-PL7	Open	Selects J2 as –5.2 V supply if closed

Table 4.2: Settings for switches and jumpers on CLOAC. DIP switches are written with first bit corresponding to switch lever marked 1.

Cabling

Cabling to and from CLOAC are as follows:

- 10 pin female-female 1-1 flat cable from Fanout A to MuSTARD. This could in principle be connected to any fanout connector, but due to compability issues with the VBURST routine in SCTDAQ, it has to be connected to Fanout A.
- 10 pin female-female 1-1 from Fanout B to SLOG IN.

The two lines carries the same signals, clock and command, but the external busy signal is only sent from MuSTARD to CLOAC.

It can be very useful monitoring the clock and command signals on one of the LEMO outputs J8/J9 of the CLOAC on an oscilloscope.

4.4 SLOG (SLOw command Generator)

The SLOG is designed to generate slow commands (see chapter 3) for the control and configuration of the ABCD3T chips. It receives clock and fast commands from the CLOAC and sends them, together with its own slow commands to the PPR2 (section 4.7). The PPR2 fans the signal out to 12 channels at LVDS[6] levels, which means that SLOG can support up to six detectormodules (one channel for each stream/master chip). A picture of SLOG is shown in figure 4.6, indicating inputs/outputs, switches and jumpers.
The SLOG also has an internal 40 MHz clock, but in this setup it gets clock from the CLOAC, so its internal clock is disabled. It can also generate L1 triggers and fast commands for standalone operation. Under normal conditions, the computer software assembles the slow commands, and loads them into SLOG RAM which outputs the slow commands synchronized to the CLOAC clock. The RAM is a 32K*16, with bits 0:11 corresponding to stream 1-12. [19]

The SLOG is, like the CLOAC, a 6U single width VME module with selectable A32 or A24 adressing using D16 programmed transfers.

[6] LVDS – Low Voltage Differential Signal (IEEE 1596.3),

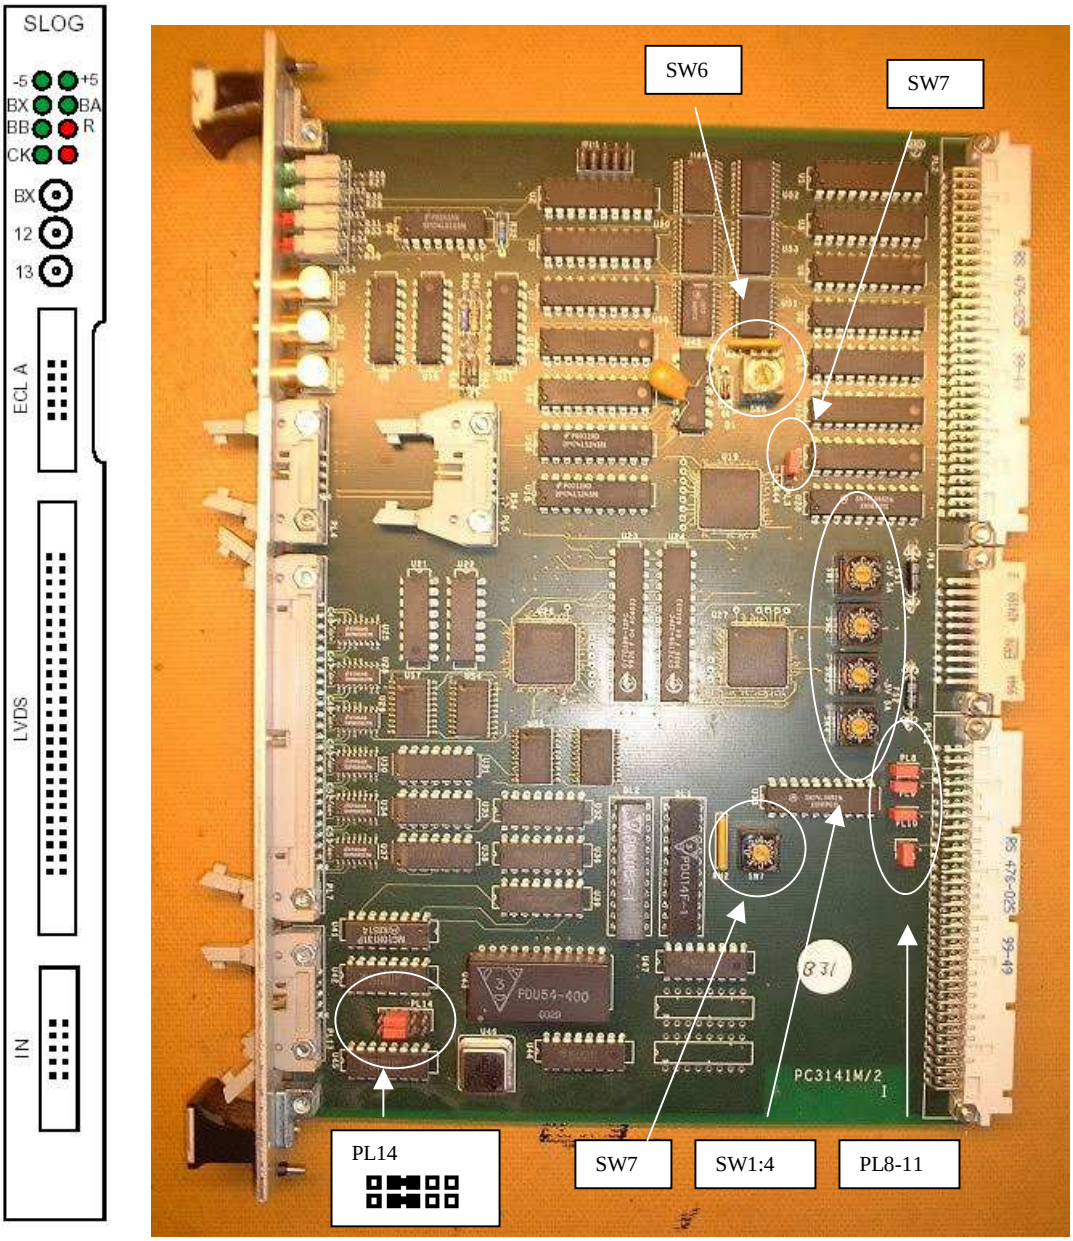

Figure 4.6: Front- and sideview of SLOG VME module showing relevant switches and jumpers, see table 4.3

Inputs:

- **IN:** This is a 10 pin IDC, receiving clock and commands from CLOAC fanout B. SLOG receives the differential ECL[7] signals clock and command from CLOAC.
- **BX:** This is a LEMO NIM input for the input of an external busy signal, not used in this setup.

[7] ECL – *Emitter-Coupled Logic.* A circuit/level standard for high speed communication.

Outputs:

- **LVDS:** A 50 pin header carries the 24 differential LVDS outputs, 12 pairs of clock and command for the 12 streams. This connects to PPR2 for further fanout to the six output channels.
- **ECL A/B:** Two 10 pin headers, ECL A and ECL B, carries copy of clock and command from the SLOG. This is primarily for controlling a MuSTARD in the absence of a CLOAC. These connectors can also receive the external busy signal from MuSTARD.
- **12/13:** These outputs carry the memory bits 12 and 13. These can be programmed and used as desired. In this setup, 12 monitors command channel 12.

LEDs

The LEDs gives information about voltage +-5, Run (L1 trigger is sent), BX BA BB are busy signals, CK indicates the presence of a 40 MHz clock.

Power requirements

SLOG requires +5.0 V and –5.2 V. Like the CLOAC, this is received via the JAUX which were soldered onto the card upon recieval. JAUX power supply must be selected with the PL8/9 jumper.

Addressing

In SCTDAQ, SLOG address is set to be A32 0x01800000h, with an address space of 0x00020000. The address of SLOG is set with the SW1 to SW4 corresponding to A19-A28 (see figure). This address corresponds to the address set in the T&M explorer.

The switches and jumpers of SLOG are summarized in the table below

Name	Value	Description
SW1	0	A19-A17
SW2	8	A23-A20
SW3	1	A27-A24
SW4	0	A31-A28
SW6	7	Default clock duty cycle
SW7	1	Default clock phase switch
PL3	Open	A32 selected
PL8-11	Open	JAUX –5.2 V supply selected
PL14	See figure 4.6	External clock and command selected

Table 4.3: Switch and jumper settings for SLOG

Cabling:

- 50 pin 1-1 female-female from LVDS to PPR2.
- 10-pin female-female 1-1 to IN from CLOAC FANOUT B.

The LEMO 12 NIM is useful for monitoring command, 13 for monitoring clock.

4.5 MuSTARD

The MuSTARD *(Multichannel Semiconductor Tracker ABC(D) Readout Device)* is the data receiver of the system. While the CLOAC and SLOG generates clock and fast/slow commands, the MuSTARDs task is to receive and decode data from the ABCD3T chips. The MuSTARD has, like the SLOG, 12 LVDS input streams, which means that it can receive and decode data from up to 6 SCT modules.

The Mustard is, like the CLOAC and SLOG, a 6U single width VME module using A32 addressing, occupying an address space of 1 Mbyte.

Inputs:

- **Stream inputs**: A 26 pin header connects the 12 LVDS input streams from the PPR2.
- **Controls:** A 10 pin header connects clock and command signals from the CLOAC, and carries a busy signal from the MuSTARD to the CLOAC. The MuSTARD reads only L1 trigger commands, since this is for the MuSTARD the only interesting command, indicating that readout is to take place.

Outputs:

- **MA/MB/MC**: These LEMO sockets output monitoring of the input streams, with four streams multiplexed on each connector. Which stream to look at is selected wit the SEL (SW4).
- **C:** System clock

Operational principle

When the CLOAC sends a L1 trigger to the module, the trigger is also sent together with the clock to MuSTARD. It does not directly use this signal in its data acquisition, but the number of L1 triggers can be used to keep track of the number of unread events in the front-end-buffers.

When the Mustard receives data on the streams, the different streams can have different phases relative to the system clock. This is adjusted for with three 4 channel delay chips with selectable delay from 0-24 ns. This is known as the Stream Delay, and is set by the control software (SCTDAQ, see chapter 5).

The signal is then sent through a FIFO buffer to real-time event histogramming in a 32 Kword memory. Each stream is assigned a 2Kword segment, with one 16 bit word for each strip in the stream. Normally, only 768 of the words are used, but in the case of redundant readout, more strips can be read through a stream. The hits are mapped in memory with the address `<stream number>:<chip number>:<strip number>`, and are accessible from the VME bus. [20]

If the FIFO should reach its 'almost full' limit, the external busy signal is sent to the CLOAC, telling it to stop sending L1 triggers, so that the MuSTARD is allowed to read in all pending events in the FIFO.

Power requirements

The Mustard requires +5V and –5.2 V, using JAUX for –5.2 V. This is selected by keeping PL7-10 open.

Addressing

The addressing of MuSTARD in the compiled SCTDAQ is A32 0x01900000h, with spacing 0x00100000h. This is selected with the hex switches SW1-3, corresponding to A19:A31, see table 4.4. The setting in T&M explorer must also correspond to this address.

Name	Value	Description
SW1	9	A23-A20
SW2	1	A27-A24
SW3	0	A28-A31
SW4	-	Selects monitored stream
PL3	Closed	Not documented
PL4	Closed	Not documented
PL5	Open	Not documented
PL6	Open	Not documented
PL7-10	Open	JAUX –5.2 V supply selected

Table 4.4: Switches and jumper settings of MuSTARD

Cabling

- 26- pin female-female 1-1 flat cable from PPR2 to STREAM INPUTS connector
- 10-pin female-female 1-1 flat cable from CLOAC FANOUT A to CONTROLS connector.

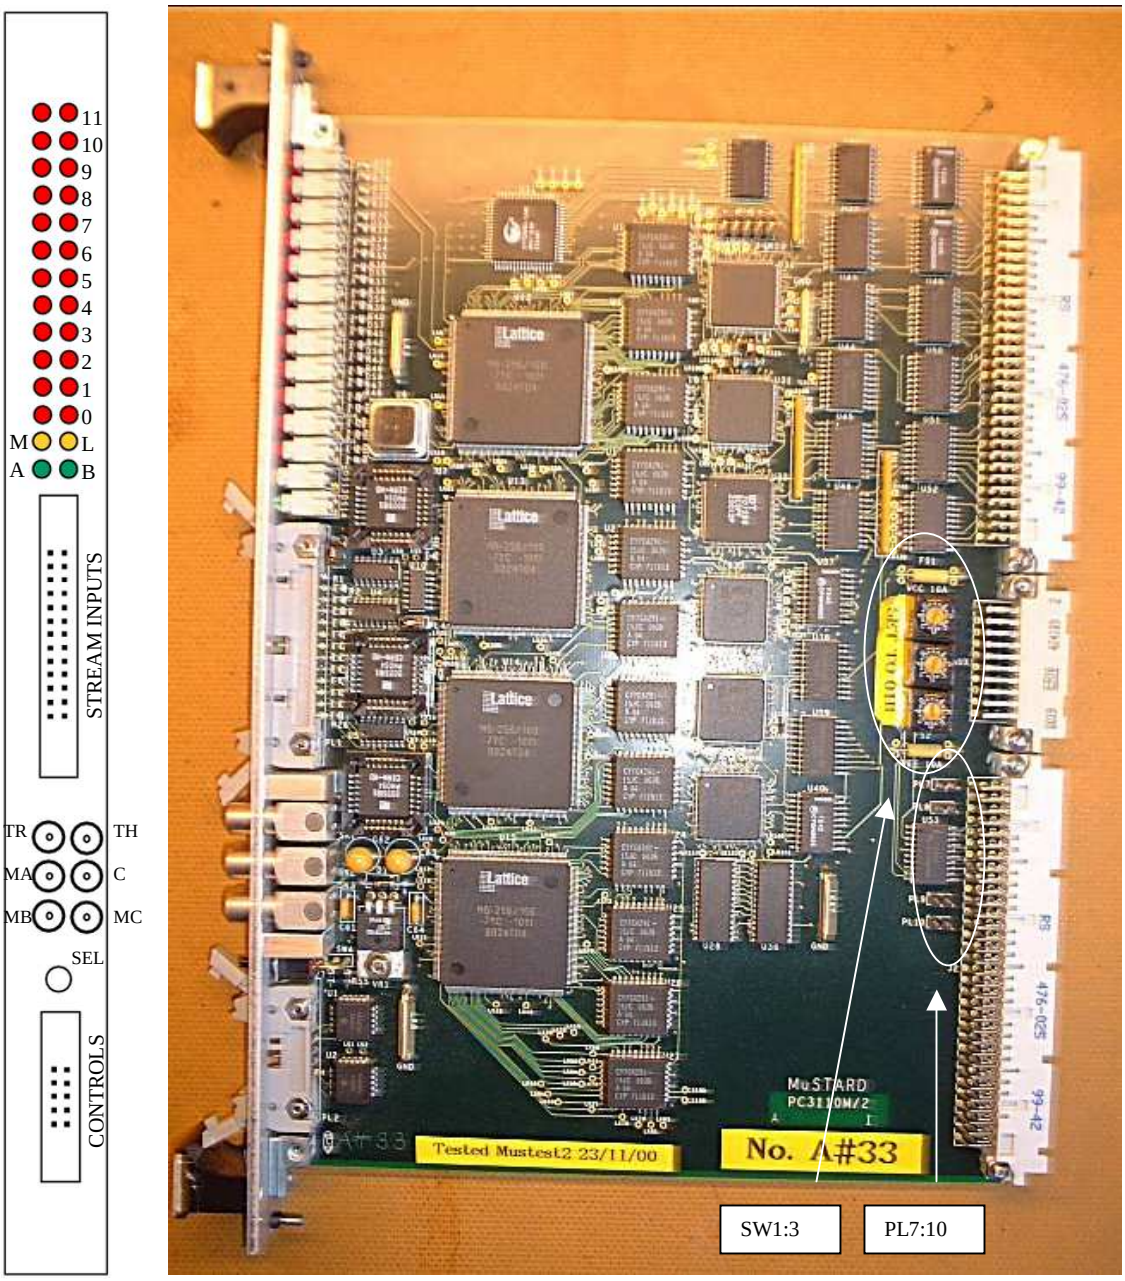

Figure 4.7: Front and sideview of MuSTARD.

4.6 SCT-LV3

SCT LV-3 is the low voltage power supply for the detector module testing. It can in principle support two modules at a time, but the LV-3 used in this setup has a faulty channel, so it can only support one module. It contains functions for resetting a module, measuring voltage and current levels, reading temperature, as well as detector bias

through-put. It is a standard 6U VME module with A24/A16 addressing, D16/D08 transfer.

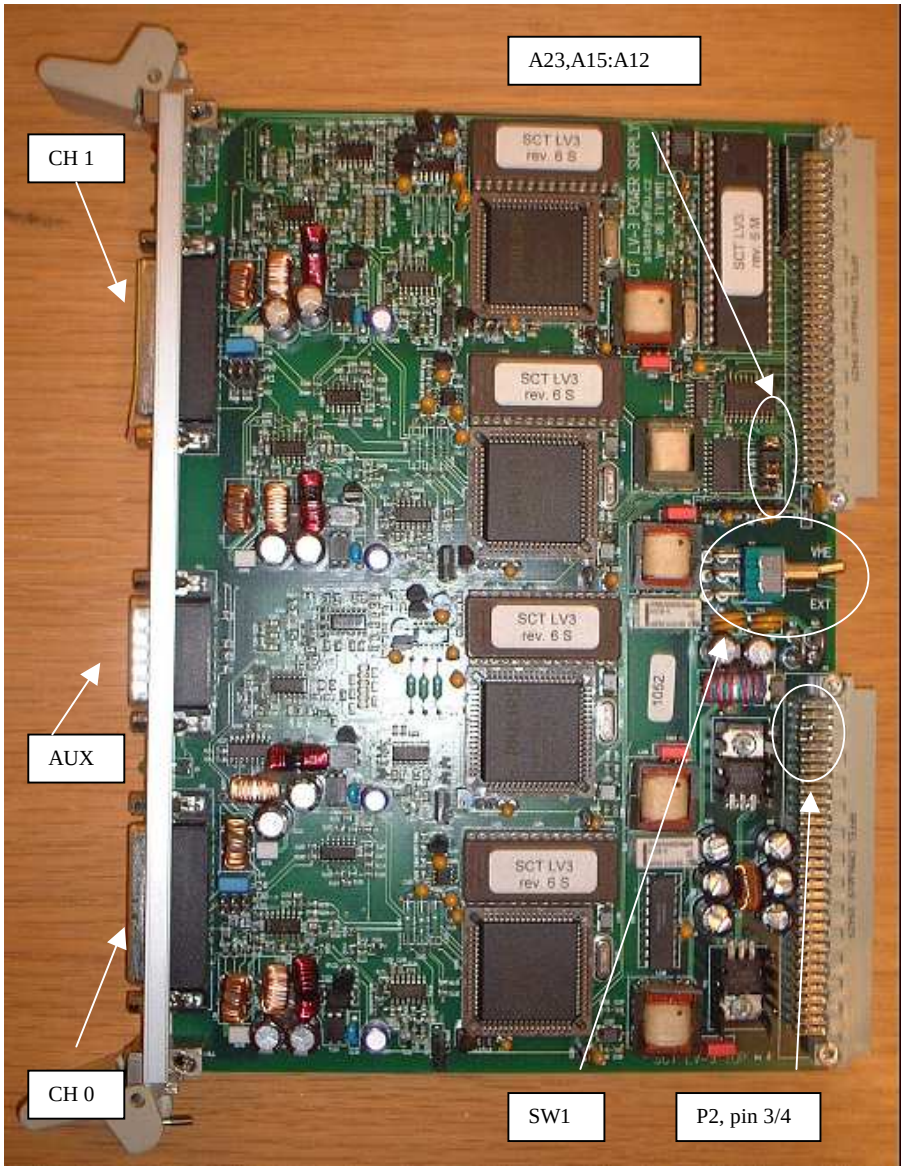

Figure 4.8: SCT LV-3 VME module.

Each of the two channels has two sepearate multi-voltage supplies, one for digital and one for analogue voltages. Both supplies are floating, i.e. insulated from ground. [21]

Connectors

CH0 and CH1 has the same pin layout, both are Female 25 pin CANNON D-SUB, sending the following signals:

- *Vdd/Vdd ground, digital power*

- *Vcc/Vcc ground, analogue power*
- *Logical signals: SELECT (for selecting which pair of clock/command input to use at ABCD3T) and RESET, which completely resets the chips.*
- *Sense of analogue and digital voltages on hybrid*
- *Sense of temperature on upper and lower side of hybrid.*
- *HV bias/ HV bias ground from AUX port*

The AUX male 15 pin CANNON D-SUB features the following signals:

- *HV Bias/HV return for both channels*
- *Signal ON for each channel*

Power requirements

SCT-LV-3 requires +12 V and –12 V. Since LV-3 does not have a JAUX connector, power is taken from the P2 connector of the VME bus. This is selected by SW 1 by setting this to VME (uppermost position). In our setup, the VME crate is configured to put –5.2 V on pins 3 and 4 on P2, this is required by MuSTARD, CLOAC and SLOG. It is however not compatible with SCT-LV-3, and would create a short circuit, so pin 3 and 4 on P2 was cut, as indicated on figure 4.8.

Addressing

The addressing in SCTDAQ for SCT-LV3 is A24, 0xFF0000h – 0xFFFFFFh. The jumpers A23/15/14/13/12 corresponds to the address bits with the same number, and is thus set to 10000, where open jumper equals 1. Address bits A22:A16 are hardwired to 1. This matches the settings in SCTDAQ and T&M measurement explorer.

Name	**Value**	**Description**
A12	Closed (0)	A12 bit
A13	Closed (0)	A13 bit
A14	Closed (0)	A14 bit
A15	Closed (0)	A15 bit
A23	Open (1)	A23 bit
SW1	VME	Power supply through VME bus
P2 pin 3-4	Cut	To prevent short of power supply
All other jumpers	Open	-

Table 4.5: Switch and jumper settings of SCT-LV3.

Cabling

- *15 pin female connector connected to AUX from floating Keithley 487 HV supply. HV+ to pin 13 and HV– to pin 14.*

- *25 pin female male connector from CH0 to male connector for supportcard. Cable is 1-1.*

4.7 Patch Panel Replacer (PPR2)

The PPR2's task is to fan out the signal from SLOG on six 10 pin IDC connections, and fan in data from the same six connections and pass it to the MuSTARD. The PPR2 fits in a 6U VME crate, but is not connected to the VME bus.

Each of the six fanouts are marked (P0-P5), and this should correspond to the value in the *st_system_config.txt* in SCTDAQ, see section 4.9.

Cabling

16 pin female connector from PPR2 1-1 to 15 pin male with the cable from pin 16 of PPR not connected to the supportcard connector.

26 pin female-female 1-1 to MuSTARD, 50 pin female-female 1-1 to SLOG.

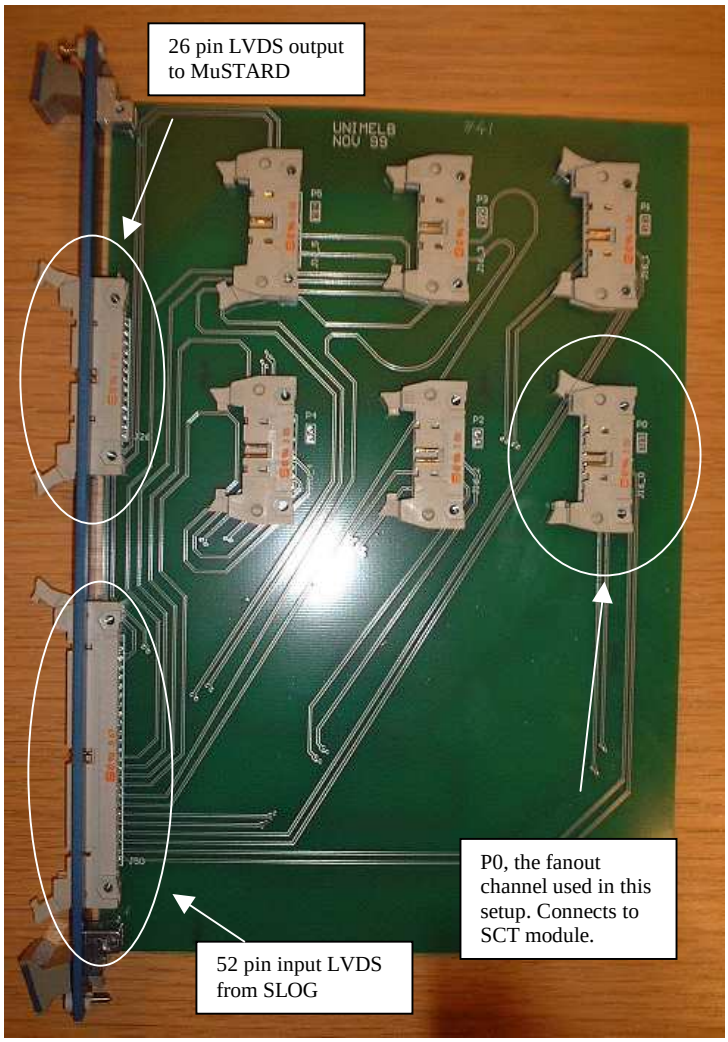

Figure 4.9: PPR2 (Patch Panel Replacer).

4.8 Support card SC2001

To connect the two connectors from the SCT LV-3 and PPR2 to the SCT module, a support card is needed. On figure 4.10, a picture of the support card is shown. The support card has circuitry to adjust the LVDS voltages to the right level before entering

the module, thus acting as a simple repeater. It is also worth to notice that on the support card, digital ground, DGND, is connected to signal ground from the SLOG, which in turn is connected to VME crate ground, making this the ground connection from VME crate ground to the floating high and low voltage supplies. See chapter 6 for a more detailed discussion of grounding of the system.

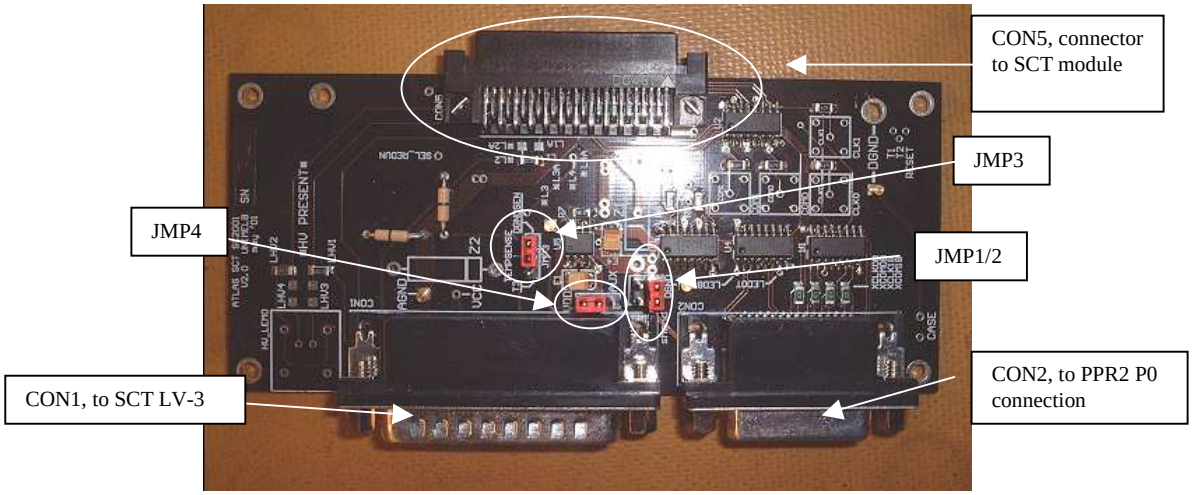

Figure 4.10: SC2001 supportcard

Name	Setting	Description
JMP1	Open	Connects shield ground to AGND if closed
JMP2	Closed	Connects shield ground to DGND if closed
JMP3	Closed 1-2	Connects Tempsense and DGND sense
JMP4	Closed 2-3	Connects Lvpower and Vdd

Table 4.5: Jumper settings on SC2001 support card.

4.9 Windows NT computer running SCTDAQ

SCTDAQ is a software package running under ROOT, developed specifically for controlling the VME modules, which again control the SCT barrel modules. It provides a graphical user interface for configuring the module and a wide range of tools for testing both digital and analogue functionality of the module. In chapter 5, a more detailed description of SCTDAQ will be given. In the next sections, the installation and configuration procedures of the software are covered. During the work on the thesis, SCTDAQ has been under development, so new updates have frequently been available. The versions used in this setup range from SCTDAQ v.3.20 to v.3.38.

4.9.1 ROOT

ROOT is an object oriented C++ interpreter containing many powerful objects for data-analysis. It is the successor of PAW and PIAF, needed because these systems cannot handle the challenges of the data amount from LHC. It contains a framework with all the tools needed to analyse large amounts of data, and is continuously developed. The heart of ROOT is the C++ interpreter CINT, allowing uncompiled C++ scripts to work together with compiled code, a feature SCTDAQ relies heavily on. Because of its C++ framework, ROOT can easily be extended to other domains than particle physics analysis.

The ROOT version used for the test setup is v3.02/07

Installation of ROOT under Windows NT is fairly straightforward, and detailed instructions can be found on the root homepage [22].

4.9.2 Download/configuration of SCTDAQ

The software can be downloaded from the SCTDAQ homepage [23] and is installed in the folder *c:\sctdaq* (if not, a recompilation is needed). On the homepage of SCTDAQ there is an excellent description of how to install SCTDAQ on a machine, so only an outline of what was concretely done in this setup will be given.

A presupposion of running SCTDAQ is that ROOT is run from the SCTDAQ catalogue, allowing the use of SCTDAQ specific ROOT configuration files.

The path in the SCTDAQ specific ROOT configuration file *(c:\sctdaq\.rootrc)* was modified from

```
WinNT.*.Root.DynamicPath:    .;$(ROOTSYS)/bin;$(PATH)
```

to

```
WinNT.*.Root.MacroPath:
.\;$(ROOTSYS)\macros;$(ROOTSYS)\tutorials;c:\sctdaq\stdll;c:\sctdaq\anal
ysis;c:\sctdaq\macros;c:\sctdaq\tests;d:\sctvar\macros
```

to get ROOT to search in the SCTDAQ paths when the user executes a macro.

The following changes were also done

```
Rint.History:            $(HOME)/.root_hist
Rint.Load:               rootalias.C
Rint.Logon:              rootlogon.C
Rint.Logoff:             rootlogoff.C
```

exchanged with

```
Rint.Load:               c:\sctdaq\rootalias.C
```

```
Rint.Logon:             c:\sctdaq\rootlogon.C
Rint.Logoff:            c:\sctdaq\rootlogoff.C
Rint.History:           c:\sctdaq\root_history.txt
```

This causes ROOT to run the logon/off files in the *c:\sctdaq* directory. These macros load the correct .dll, stdll.dll, needed for the macros in SCTDAQ to run properly. Depending on wether or not one is using a CLOAC, one has to choose the correct .dll. This is done in the *\sctdaq\bin* folder, where there are two .dll files, stdll_cloac.dll and stdll_nocloac.dll. The correct file, in this case stdll_cloac.dll is renamed to stdll.dll.

4.9.3 The SCTDAQ configuration files

One must also create a directory *d:\sctvar\config* for the configuration and results files from SCTDAQ. This can be changed, but a recompilation is then needed, since the directory is hardcoded in stdll.dll. A different drive is used to ensure that the configuration files do not get overwritten if a reinstallation of SCTDAQ is performed. The configuration files required by SCTDAQ are:

- **st_config_system.dat**
- **[modulename].det**

st_config_system.dat:

This file determines parameters for the configuration of the entire system. The file used in this setup is shown below:

```
/*
       ST Configuration file
       =====================
*/
       DETECTOR  LV     HV     SLOG      RSLOG     OPTO_TX        ROPTO_TX   OPTO_RX
id    pr ac  cr ch  id ch  id ch pg  id ch pg  typ id ch iset typ id ch  id c0 c1 th0 th1
Module  0  1  1   0  0   0  0   0  0  0   0  1  0  -1 -1  1  0    -1  -1  1  -1  0  1 100

MuSTARD         Module
id s0 s1 d0 d1  Filename
------------------------------
100   0  0  1 5 5  20220380200002
/*END*/
```

The parameter *detector id* referes to which plug on the PPR the module is connected to (see section 4.7.). This enables the software to correctly map the physical connections of the modules in the system. This is set to 0 here, corresponding to P0 on the PPR.

The file contains various parameters for the VME modules, and MuSTARD has four additional parameters, numbering of streams (s0/s1) being 0 and 1, and a stream delay for each stream (d0/d1). This stream delay, possible to set via a macro in the SCTDAQ package, is the delay between the clock from the CLOAC and the data from the module (sent in response to a L1A trigger from the CLOAC).

The last column gives each module in the system a name. This serves two purposes, it identifies the module in plots and output data, but the name is also used for module specific configuration files on the form:

***[modulename]*.det**

The file contains all the necessary information for correctly configuring the module. One can set parameters global to the module, and specific for each chip, like compression mode, threshold, strobe delay and a more.

Additional configurations files are (not required for operation):

[*modulename*].mask

This is the mask file, containing a complete listing over the channels that should be masked. It is just a listing of the channels number, with possible comments marked by #.

[*modulename*].trim

This is the trim file, it describing the trimming of each channel. The file starts with a trim range for each chip (0-3, see chapter 5) where the format Chip 0 1 means that chip 0 should be set to trim range 1. It then lists the trim values for each channel as a number between 0 and 1 from historical reasons, instead of a number between 0 and 15, which would be the most logical representation since the TrimDAC is a four-bit register.

Chapter 5: The SCTDAQ software

5.1 Overview

The electrical tests of the SCT barrel modules are controlled and analyzed with the SCTDAQ package under ROOT, which also provides a graphical user-interface and various plots of the measurements. In this section, the basic functionality of SCTDAQ will be overviewed. A short presentation of the software and how to install and configure it were given in chapter 4.

Since ROOT is based on the CINT C++ interpreter, uncompiled C++ scripts can be run directly in ROOT. SCTDAQ rely on of such uncompiled scripts (macros) for data analysis and graphical interface, and a compiled dynamic linked library, stdll.dll, for communication with the VME modules (See chapter 4 for an introduction to these). A schematic view of this is presented in figure 5.1.

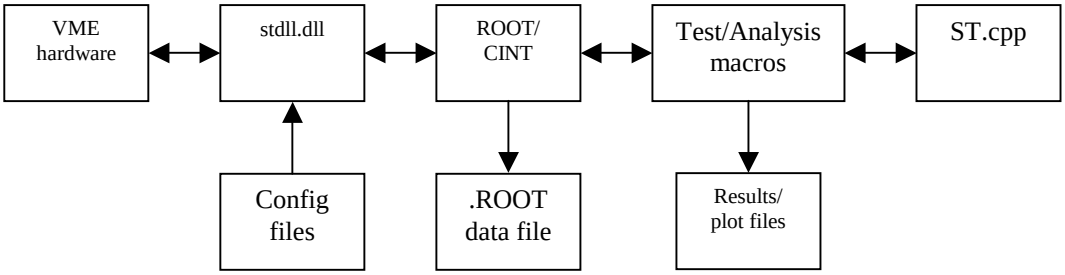

Figure 5.1: Structure of SCTDAQ

SCTDAQ is started with the top level macro ST.cpp, with the command *.x ST.cpp* in ROOT. *ST.cpp* handles the interaction with the user through menus and various displays. Most tests involved in the electrical testing of the module are implemented in separate macros, represented by the “test/analysis macro” box in figure 5.1.

These macros communicate with the VME hardware through the compiled dynamic linked library *stdll.dll*, represented by an object *e* of the class TST, which provides the high level methods for passing on hardware calls to lower level functions, see figure 5.2 for an overview of the structure and dataflow of *stdll.dll*. All hardware tasks go through this object, which again calls lower level functions. The source code of the library consists of a set of files located in the *stdll* and *stlib* directories of SCTDAQ.

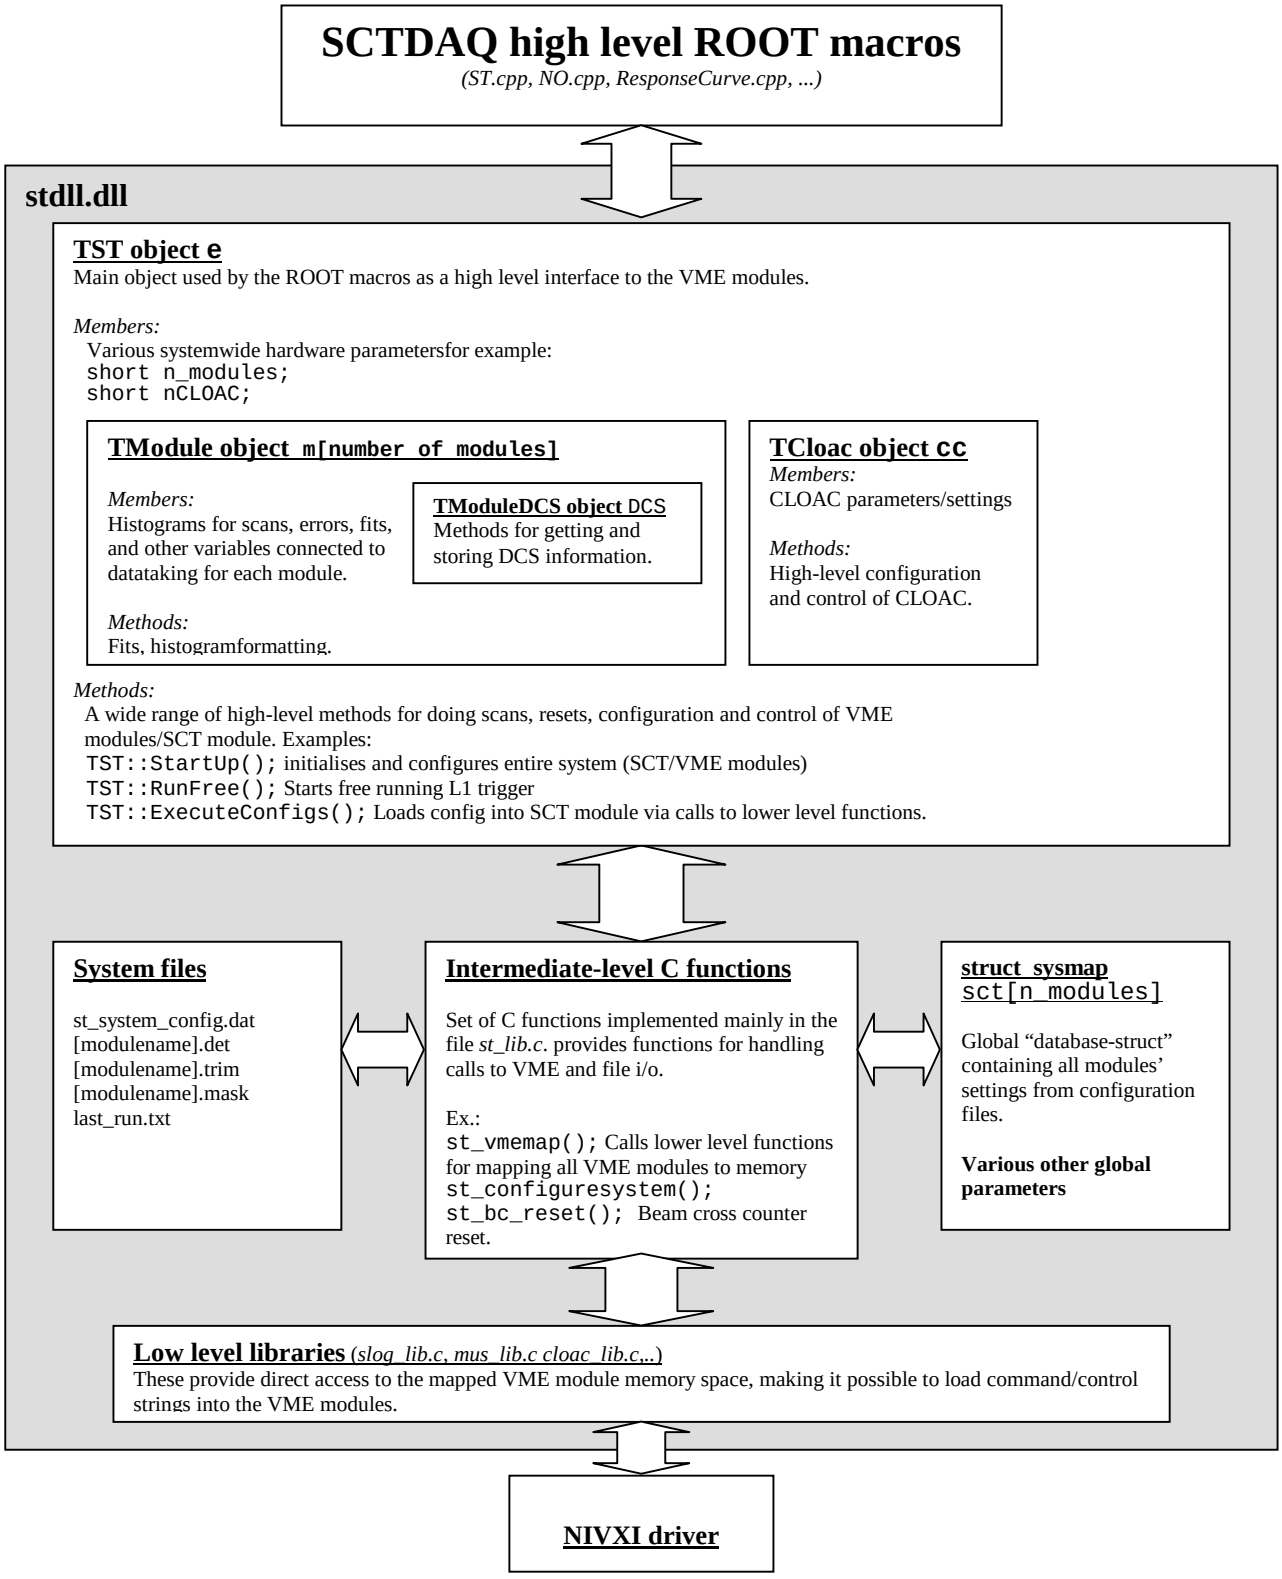

Figure 5.2: Internal structure of stdll.dll showing communication between high level macros and VME hardware.

Each time SCTDAQ is run, it stores all recorded data in a .ROOT file with the name *strun[run number]_[scan number].ROOT* in the directory *d:\sctvar\data*, so that it can be accessed at a later stage. The macros also create PostScript plots from most of the standard tests and place them in *d:\sctvar\ps*. All test results are written to an ASCII file on the form *d:\sctvar\results\[module name]_[date].txt*. This is very convenient for easy access to the test data and various numerical values not available from the .ps plots. The plots and result files are described in section 5.4 together with the main SCTDAQ macros.

5.2 Example of operational principle : Changing the threshold of the SCT module

To get a practical example of the structure of hardware calls shown on figure 5.2, we will see how the process of changing the threshold value in the ABCD3T chips on the SCT module is handled internally by SCTDAQ. The operation is performed by choosing *ChangeVariable* → *SetVThr* from the main menu. This calls the function

```
st_dialog_change_variable(short typ)
```

in the top level *ST.cpp* macro. The function promts the user through the console for a new value for the threshold (passed to the function as typ).

The reconfiguration of the SCT module has to be issued as slow commands from the SLOG (see section 4.4.), so communication must go through the TST object `e`, figure 5.2. Once the user has given a value, the TST method

```
  e->ConfigureVariable(typ,var);
```

is run with `typ` being threshold, defined to numerical value 1, and `var` the value input by the user, say 200.0 mV. The implementation of the method is shown below:

```
void TST::ConfigureVariable(int typ, float var){
  float newvar;
  short n;

  st_coerce_variable(typ,var,&newvar,0);
  st_configure_variable(-1,typ,newvar);
  ...(if-sentences which are not used for threshold)

}
```

The method first calls the intermediate-level C function `st_coerce_variable`. This is implemented in *st_lib.c*, see figure 5.2, and ensures that the value passed as `var` is valid for the variable type `typ`. In the case of threshold, the value is checked for being between 0.0 mV and 637.5 mV, rounded to nearest 2.5 mV and returned through the pointer `&newvar`. It also displays the message

```
      st_coerce_variable: parameter 1 unchanged from 200.0
```

in the console, to tell the user that the passed value, 200.0 mV, is a valid value for threshold (referred to as parameter 1).

After this, the function

```
st_configure_variable(-1,typ,newvar)
```

is called, with arguments `-1` , meaning configure all modules in the system, `typ` for variabletype (threshold), and `newvar` as the valid rounded value passed from `st_coerce_variable`.
This function, which is implemented in the file *st_lib.c*, changes the threshold parameter in the struct `sct[n_modules]`, which contains all SCT module settings, see figure 5.2.

`e->ConfigureVariable` is then finished running, and returns to the function where we started; `st_dialog_change_variable`.

Notice that no configs has been sent to the SCT module yet. `e->ExecuteConfigs();` is then called, the method which will put the updated contenst of the `sct[]` database to the SLOG, which in turn will send it to the SCT module:

```
void TST::ExecuteConfigs(){
  st_execute_configs(0,1);
}
```

This simply calls a new C function from *st_lib.c* with arguments 0 (meaning no dump to screen) and 1 (update every parameter in module).

It starts by calling

```
slog_disable_external(slog_id)
```

This is an example of a function from the low level libraries. It does what it says, it disables external inputs on the SLOG, this is because we do not want external interrupts while writing to SLOG memory. The function is implemented in *slog_lib.c*;

```
int slog_disable_external(short id){
  unsigned short *ip;

  if ((id<0) || (id>=slog_total)) return -1;

  ip = (unsigned short *)(slog_info[id].address + SLOG_EXTERNAL);
  *ip = 0;

  return 0;
}
```

After checking that the SLOG id is valid, it accesses the parameter `slog_info[id].address`. This contains the memory mapped address of the SLOG, so when writing to this address, the NIVXI drivers ensures that it is written to the correct VME address through the MXI bus. The address is provided by the function

`MapVXIAddress(...)` during the initialisation of the system (see section 4.3 for an overview of the MXI bus used for communication between VME and PCI). `SLOG_EXTERNAL` is a parameter hardcoded in the headerfile *slog_private.h*, and gives the offset from the base address to the external enable bit in the SLOG.

The pointer `*ip` now points to the correct address, and the register is set to 0 to disable external signals in the SLOG.

Various other configurations of SLOG is also necessary before sending a command to the module, but we shall not go into the details of this here.

The next step is to create the binary commandstring for the SLOG, which it should send to the ABCD3T chips on the SCT module. There are functions written for this in the file *abc.c*, called *control block generators.* These functions take the `sct` struct database containing all module parameters and converts it to binary control blocks which the ABCD3T chips can read. The control blocks are then loaded into SLOG memory with the same principle as for `slog_disable_external` with the function

```
slog_load(slog_id, cbr, cbr_size);
```

where `cbr` are the control block. The sending of the control block, which is a *slow* command, is initiated with the function

```
slog_start(slog_id);
```

This sets the *RUN* bit in the SLOG control register, which starts the sending of the slow command string to the SCT module. This updates the configuration registers in the ABCD3T chips, and thus changes the threshold of the module to 200.0 mV.

5.3 Scans

Most of the analogue tests used in the QA of the module are based on *scans*. A scan is characterised by gradually running a SCT module parameter from one value to another, for example threshold from 0.0 mV to 400.0 mV in steps of 5.0 mV. For each scan point, a package of signals is sent, usually involving a calibration signal and a L1 trigger for readout of the module. This combination of signals and triggers is usually sent a given number of times for each scanpoint, in order to get a good statistical basis. A set of such signals are called a *burst*, typically containing 1000 triggers.

SCTDAQ has a number of predefined scans and burst, some of these are listed in table 5.1 and 5.2.

Scan type	**Name**	**Description/Max range/Min stepsize**
1	ST_VTHR	Threshold scan / 0-637.5 / 2.5 mV
2	ST_VCAL	Calibration pulse scan / 0-159.375 / 0.625 mV

3	ST_STROBE_DELAY	Strobe delay scan / 0-63 / 1 bit
4	ST_PREAMP	Preamplifier scan / 0-285.2 / 9.2 µA
5	ST_SHAPER	Shaper scan / 0-37.2 / 1.2 µA
6	ST_TRIM	Trimrange scan / 0-15 / 1 bit

Table 5.1: Some common SCTDAQ scan types

Burst type	**Name**	**Description**
1	ST_TBURST	Triggerburst, uses VCAL from CLOAC, which consists of a calibration pulse followed by pipeline and system delay and finally a L1 trigger to read out the signal, histogramming in MuSTARD (default)
5	ST_RAWBURST	Same as 1, but histogramming in software
6	ST_SENDIDBURST	Module in SENDID mode
7	ST_NSEBURST	For Noise occupancy scan, only L1 trigger without calibration signal, MuSTARD histogramming
11	ST_TBURST2	Same as 1, but with errordecoding

Table 5.2: Some common SCTDAQ Burst types.

5.3.1 Operational principle of scans

While running a scan, SCTDAQ presents a graphical view of the hit-data from the scan, showing a plot of hits vs. channel for each module and link, figure 5.3.

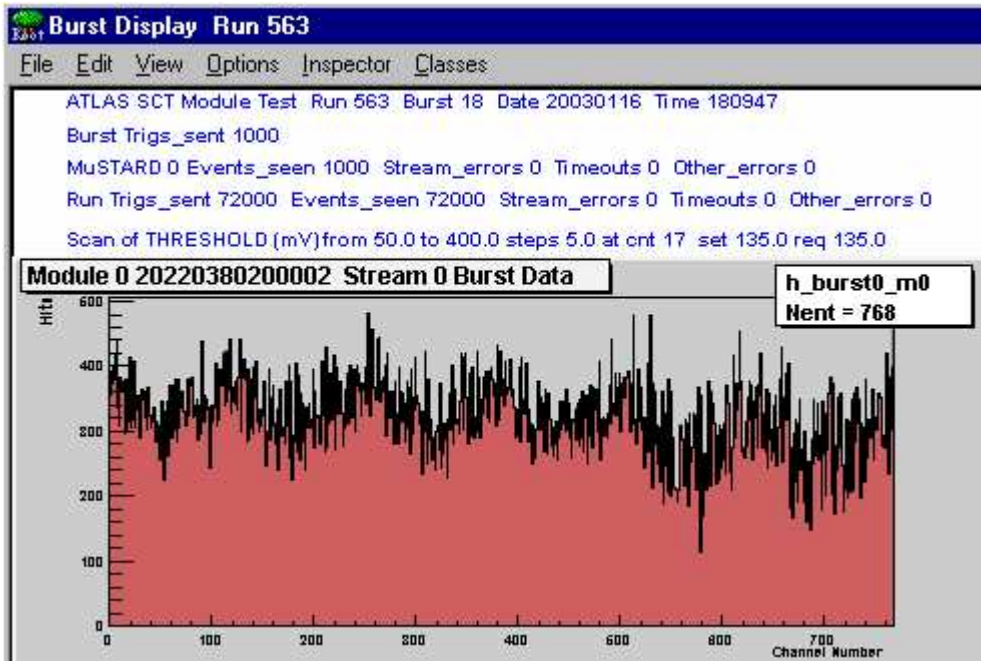

Figure 5.3: The burst display, showing number of hits for the current scan point.

Once the burst and scan type are selected, the scan is performed as a loop over all the points in the scan, and for each scan point, the system runs the method:

```
e->ScanPoint();
```

This method takes care of executing the bursts and increments the scan point until the end point is reached. The scan point, for example a threshold value, is loaded into SLOG with `st_execute_configs(..)` in the same manner as when changing a variable in the section 5.2, and sent to the SCT module.

Central in the method `e->ScanPoint()` is:

```
e->ExecuteBurst();
```

which resets the MuSTARD histograms, and starts the sending of the selected burst through various lower level functions, depending on the burst type selected. The burst is sent a set number of times, typically 1000, followed by a module soft reset for clearing possible hit data in the module.
When the burst is completed, the results are read out from the MuSTARD (if we use MuSTARD histogramming, which is usual). This is done by using the low level functions for getting a pointer to the MuSTARD histogram memory and later using the C method `st_give_data(...)` for copying its contents into a ROOT histogram. The MuSTARD histogram is then cleared for the next burst and scanpoint.

The procedure for performing a single scan point can be summarised as follows:

- Assemble slow command control block in SLOG from struct `sct` database for sending correct scan point to SCT module.
- Send control block from SLOG to SCT module.
- Reset MuSTARD histograms (if they are used).
- Send predefined trigger type from CLOAC to SCT module and MuSTARD a specified number of times (For standard threshold scan: 1000 VCAL (Calibration signal + delay + L1A) triggers).
- Read and copy hits from MuSTARD histogram to memory.

This loop is performed for each scan point until the last scan point is reached.

When the scan is completed, the hit data is stored in two two-dimensional ROOT histograms (type `TF2H`) in the *TModule* object `m`, (figure 5.2), one for each link/stream:

```
e->m[n_modules]->h_scan0
e->m[n_modules]->h_scan1
```

in the *TST* object `e`. The histograms are displayed in a window called the *Scan Display* once the scan is completed. In figure 5.4, one of the streams in the Scanb Display is shown. The histogram are plotted as scanpoints vs. channels.

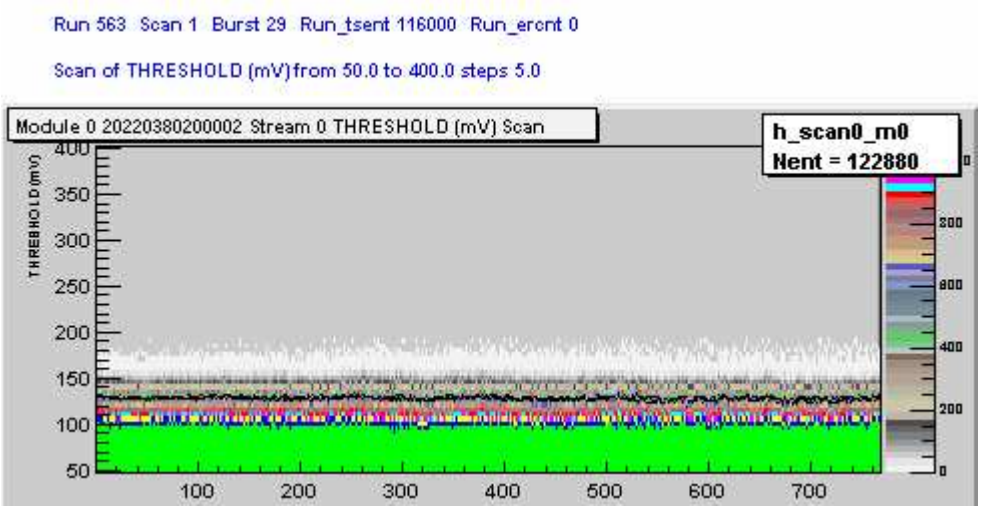

Figure 5.4: The scan display. This shows the result of a threshold scan from 50.0 mV to 400.0 mV in 5.0 mV steps, with 1000 triggers sent for each scan point. For each threshold value (y-axis), the number of hits in each channel (x-axis) is indicated by a color coding. At approximately 130.0 mV a dotted black line of points indicates the VT50 points, see section 5.3.3. One could also say that the scan display shows the *occupancy*[8] for the channels, however not normalized to 1.

5.3.2 S-curves

While the scan display gives a rough impression of the occupancy of each channel of the SCT module, SCTDAQ can also plot the scan data for each channel as number of hits vs. scanpoint, for example threshold. These curves give a more precise view of the hits for the channels, and are called *s-curves.* They are selected from the *ShowScurves* button in the SCTDAQ main menu, which plots s-curves in a window and writes them to a .ps file.

On figure 5.5 on the next page there is an example of a set of s-curves from a Threshold scan from 30.00 mV to 270.00 mV with an injected charge of 2.0 fC, with 1000 triggers sent per point. The x-axis shows threshold in mV, while the y-axis represents number of hits. The plot shows one side of the module, Link 0, with its 6 readout chips. Each chip has its own column, while the four calibration lines are shown as four rows. Each window contains 32 channels, so the plot shows all the 768 readout channels from the link.

Ideally, the S-curves should have been pure step-functions, but due to the intrinsic noise of the system, they get their characteristic s-shape. The steepness of the slope is a very good indication of the noise in the system. If the module is trimmed (see section 4.9 and 5.4.4), there should be little relative spread between the curves.

[8] The occupancy is the fraction *hits recorded/triggers sent.*

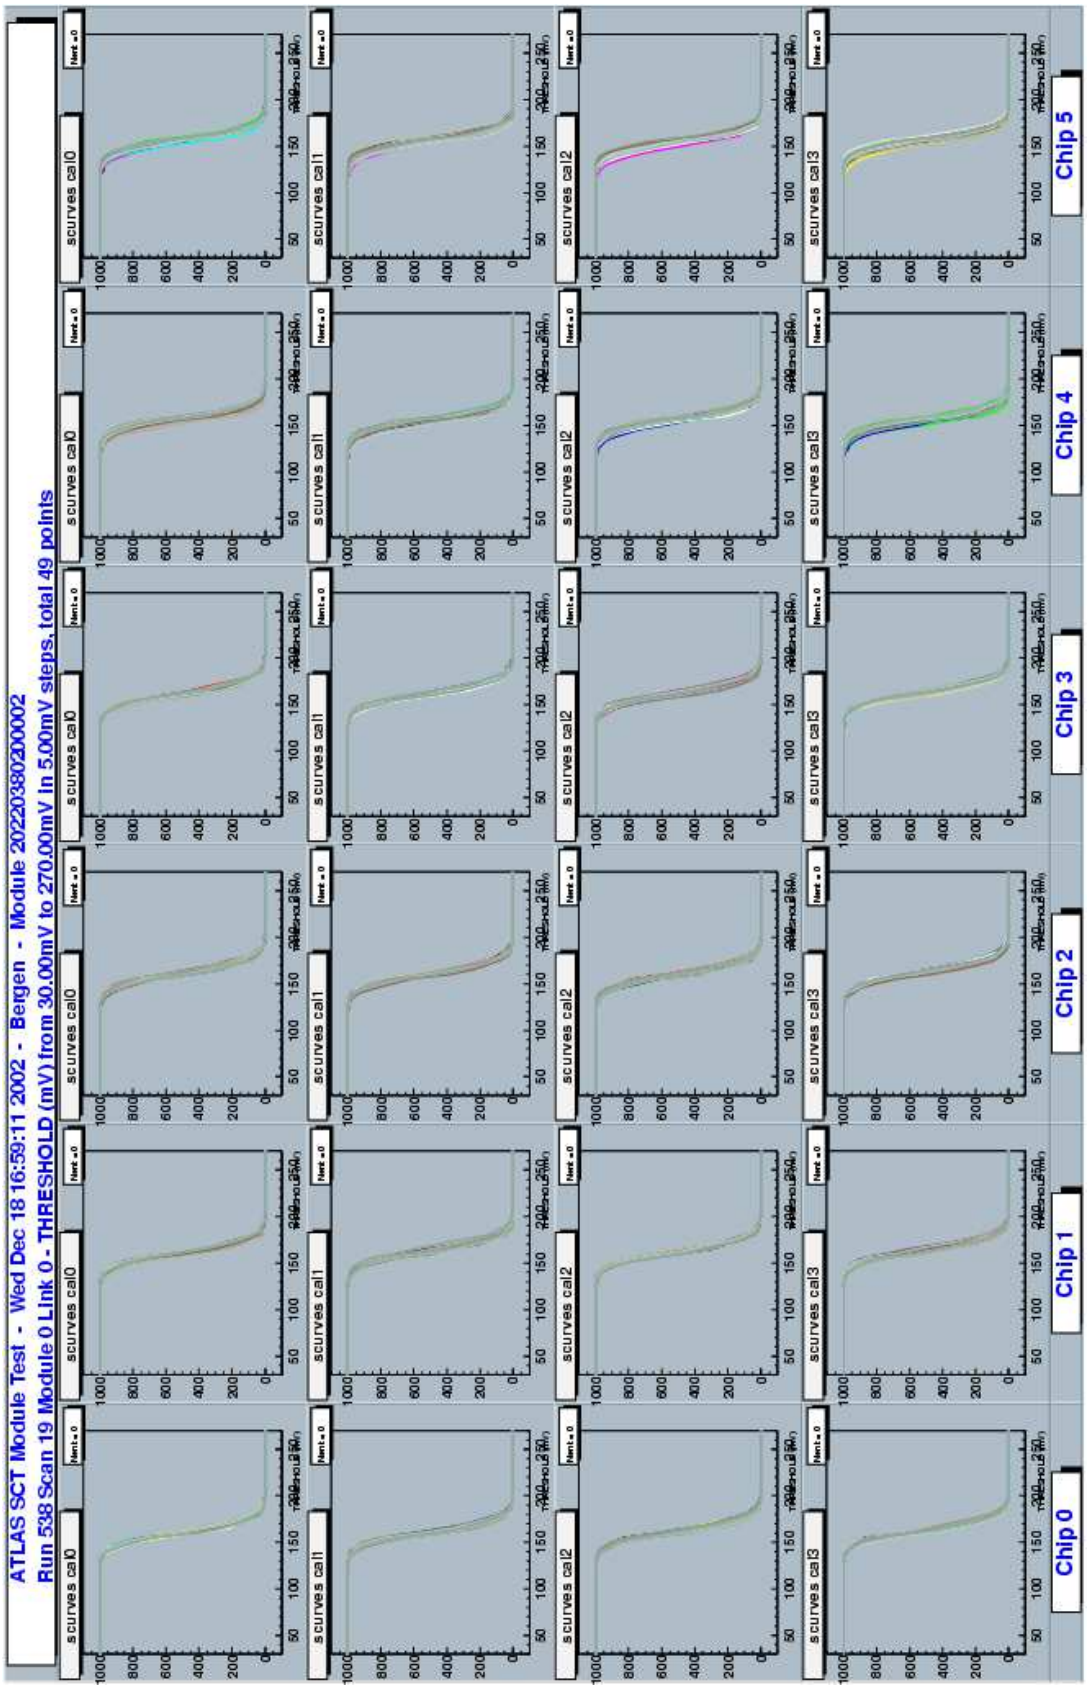

Figure 5.5: Threshold scan s-curves. Explanations are given in text.

5.3.3 Fitting of scans

To get numerical values from the s-curves, functions is fitted to the scan data. This is done by the method

```
void TModule::FitScan(short fit_type, short cal_mode)
```

There are three fit functions used for analyzing, each with three parameteres ***C****onstant*, ***M****ean*, and σ *(sigma)*:

- **Error function:**

$$\text{Erf(x)} = \frac{C}{2}\left(1 - \frac{2}{\sqrt{\pi}}\int_0^G e^{-t^2}dt\right) \text{ where } G = \frac{M-x}{\sqrt{2}\sigma} \qquad (2)$$

- **Complementary error function**

$$\text{Erfc(x)} = \frac{C}{2}\left(1 + \frac{2}{\sqrt{\pi}}\int_0^G e^{-t^2}dt\right) \text{ where } G = \frac{M-x}{\sqrt{2}\sigma} \qquad (3)$$

- **Gauss function**

$$\text{Gauss(x)} = C \cdot e^{-\frac{1}{2}\left(\frac{x-M}{\sigma}\right)^2} \qquad (4)$$

ROOT has custom function fitting as a built-in feature, so it takes care of determining the parameters Constant, Mean and Sigma. On figure 5.6, a ROOT plot of the three functions defined in SCTDAQ is shown.

In the threshold scan case, a complementary error function is fitted to each s-curve. The *Mean* parameter is interpreted as the VT50[9] point, while σ is interpreted as the output noise of the channel in mV. We see from figure 5.6 that the complementary error function strongly resembles the s-curves from figure 5.5.

After fitting the appropriate function to the scan data, $|\chi|^2$ is computed for determining the goodness of the fit. All computed values are dumped to the console after the fit is performed. The value $|\chi|^2$*/(Number of degrees of freedom)* typically gives a value around

[9] The VT50 point is defined as the point at which the chip has 50 % efficiency, for example, if 1000 L1 triggers have been sent, the channel has 500 hits.

1.0 for a threshold scan with input charge >1.0 fC, indicating that the complementary error function is a good approximation to the channel occupancy.

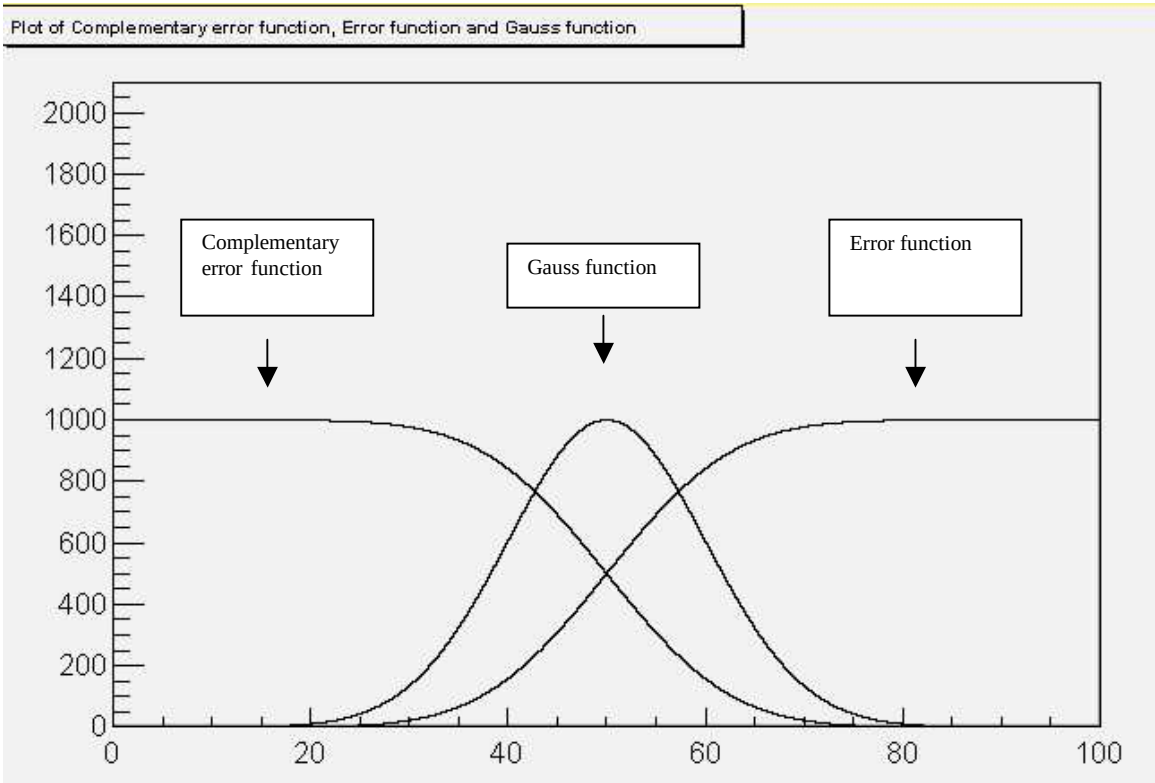

Figure 5.6: Complementary error function (3), gauss function (4) and error function (2) with parameters Constant = 1000, Mean = 50 and Sigma = 10.

5.4 The macros used for electrical testing

In the electrical testing, two sets of macros are used, the *Characterisation Sequence* and the *Confirmation Sequence*. The tests are started from the *ABCD test* menu in SCTDAQ.

5.4.1 Characterisation Sequence and Confirmation Sequence

The Characterisation Sequence aims to give a complete characterisation of both the digital and analogue functionality of the module, while the confirmation sequence only tests the aspects that is crucial for correct operation of the module, and is for this reason much quicker to perform. Table 5.3 shows the submacros run in each of the tests.

Submacro	**Characterisation**	**Confirmation**
Clock and command reception *(HardReset.cpp)*	X	
Bypass functionality *(FullBypassTest.cpp)*	X	X

Clock and command redundancy *(RedundancyTest.cpp)*	X	X
Pipeline Efficiency test *(PipelineTeset.cpp)*	X	
Strobe delay scan *(StrobeDelay.cpp)*	X	X
Three point estimation of Gain, Noise, Offset *(ThreePointGain.cpp)*	X	X
TrimRange scan *(TrimRange.cpp)*	X	
Determination of the Response Curve *(ResponseCurve.cpp)*	X	
Noise Occupancy scan *(NO.cpp)*	X	
Timewalk scan *(TimeWalk.cpp)*	X	

Table 5.3: Characterisation and confirmation sequence

5.4.2 Results- and plot files

Each test macro creates a set of data, which is written to the file *d:\sctvar\results\[serial_number]_YYYYMMDD.txt*, one for each day. Each time a macro is run, this is indicated in the file by a header containing date, time, serial number and DCS information (see section 3.6), followed by macro specific data. An example of such a header is shown below:

```
#
%NewTest
#
SERIAL NUMBER  : 20220380200002
TEST MADE BY   :
LOCATION NAME  : Bergen
Run number     : 548-1
TEST_DATE      : 06/01/2003
PASSED         : YES
PROBLEM        : NO
#
%DAQ_INFO
#
#HOST
"ATLAS1"
#VERSION
"3.20"
#DUT
"."
#TIME
"20:48:26"
#
%DCS_INFO
#
#T0     T1
25.0    27.0
#VDET   IDET
0.0     0.00
#VCC    ICC
3.48    970
#VDD    IDD
4.00    530
#TIME_POWERED
.
#
```

The reason that the DCS variables VDET and IDET (bias and leakage current) are both 0.00, is that we are using an external HV supply, not integrated in SCTDAQ.

Some of the tests generate additional textfiles or PostScript (.ps) files as indicated on figure 5.1. These are placed in the directory *d:\sctvar\ps*. Examples of such plots are shown later in this chapter.

The characterisation and confirmation sequences can be divided into digital and analogue tests. In the next sections, each macro will be given an overview, and its output data explained. The tests are specified in the documents [24] and [26].

5.4.3 Digital tests macros

The following macros test the digital functionality of the SCT module.

Clock and command reception test: *HardReset.cpp*:
The power is turned off and on, to put the chips in start-up mode. The user is then prompted to use an oscilloscope connected to the MuSTARD MA output to check that the master chip on each stream outputs 20 MHz clock (system clock divided by 2). This is the default start-up mode. The chips are then configured via SLOW commands issued from the SLOG, and the clk/2 signal should stop. The user is prompted to verify this. The HardReset signal is then sent, and the user is once more prompted to check for clk/2 signal. The analogue and digital currents Icc and Idd are recorded and written to the results file. If the user does not see the expected signals, the module is classified as faulty.

Bypass functionality Test: *FullBypassTest.cpp*:
The FullBypassTest is testing the internal routing of signals and tokens on the hybrid, see section 3.4.5. The module is programmed to each possible configuration with regard to the routing of token and data between each chip. The test runs through the 36 different token/data configurations, and checks each link, using SLOG to program the module to the different token routing configurations. The test is done for Vdd from 4.0 mV to 3.5 mV in 0.1 steps. If the links are not working at the minimum voltage 3.5, the module does not pass the test.

Clock and Command redundancy test: *RedundancyTest.cpp*:
This tests the redundancy functionality of the module, see section 3.4.2. A mask pattern 010101... is applied to the mask register of each chip. Calibration signals and triggers are sent, and a total occupancy of 50% should then be recorded due to the masking of 50% of the channels. The inverted mask pattern is then applied, 101010..., and the occupancy should still be 50 %. Select is then set to 1, to select the other pair of clock and command inputs, and the test is then repeated. If both pair of clock and command gives the same result, the test is passed.

Pipelinetest :*PipelineTest.cpp*:
The readout chain in the module contains a 132 cell pipeline, as discussed in section 3.4. First, no channels are masked, and hits should be detected. If there are channels without hits, this may come from dead cells in the pipelines, always outputting 0.
Then all channels are masked, and if there still are hits, this comes from stuck cells in the pipeline, always outputting 1.

5.4.4 Analogue tests

All the analogue tests are based on scans of threshold, strobe delay and trim values for different input charges and other conditions.

Strobe Delay Scan: *StrobeDelay.cpp***:**
The strobe delay of each chip determines the relative delay between the rising edge of the calibration strobe output and the rising edge of the clock input. To determine the optimal setting for this, a StrobeDelay scan is performed. The delay varies with temperature, so it is important that once a strobe delay has been set, the temperature must be constant, or a new strobe delay scan has to be performed.

The macro first sets a threshold based on VT50 values (see section 5.3.3) from a 2.0 fC threshold scan. It then scans the strobe delay register on each chip from 0 to 63, and fits an error/complementary error function to the rising/falling edge of the strobe delay peak to determine the timing of the edges of the signal. It then selects a strobe delay setting which is 25 % of the distance between the rising edge and the falling edge with respect to the rising edge (see figure 5.7). For acceptance, the rising edge must lie between 0 and 28, and the falling edge between 35 and 63.

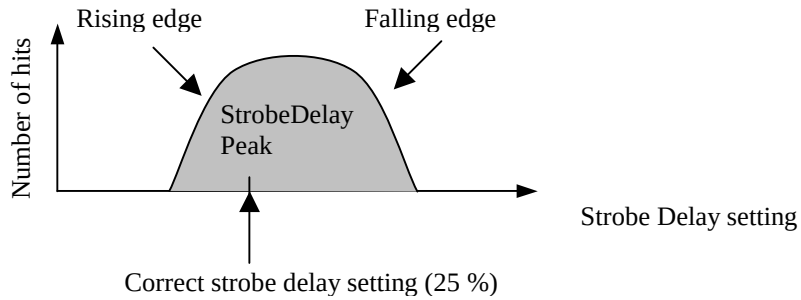

Figure 5.7: Setting of strobe delay.

Three Point Gain (Three point estimation of Gain, Noise and Offset)
ThreePointGain.cpp

This test measures quickly the gain and noise of a module. Three charges, 1.5, 2.0 and 2.5 fC are injected on the calibration lines, and a threshold scan is done for each of the charges. A complementary error function is fitted to each threshold scan to yield values for VT50 (mean) and output noise (sigma) for each channel.

For charges <1.5 fC, the lower ends of the s-curves get distorted due to noise, and the complementary error function underestimates the noise.

When the three scans are completed, a straight line

$y = ax + b$

is fitted to each set of VT50 points for each channel, where

y = VT50 point (from complementary error function fit) in mV
x = Injected charge in fC

and the parameters to be determined are interpreted as:

a = Approximation of gain in mV/fC
b = Offset in mV

This resulting linear fit is known as a *Response Curve*, giving a relationship between an injected charge in fC and the threshold level in mV. From this, the input noise can be calculated.
Since the output noise in mV for each channel is known from the sigma parameter of the fit, the input noise in ENC[10] can be calculated from the formula:

$$Input\_noise[ENC] = \frac{6250 \cdot Output\_noise[mV]}{Calculated\_gain[mV/fC]} \quad (5)$$

where 6250 is the ratio 1 fC/1.6e-4 fC (the elementary charge).

The response curve fit is plotted for each chip, and for the 2fC input charge, values for each channel is plotted. The plot itself uses the same plotting macro as ResponseCurve, and since this is a more interesting plot, an explanation of the various plots will be given there when discussing the Response Curve test, see figure 5.11 and 5.12.

Trim Range test *(TrimRange.cpp)*

This test sets out to determine the optimal trim settings (see section 3.4.4. for an introduction to the concept of trimming). The test uses threshold scans with 1.0 fC injected charge, since this is the targeted operating threshold in ATLAS. All four trim ranges are scanned, starting with range 0, where scans are performed for all 16 steps, and then range 1, 2 and 3, with only step 3, 7,11 and 15 scanned. An example of a TrimDAC plot is shown on figure 5.8.

[10] ENC – *E*quivalent *N*oise *C*harge, is the noise in numbers of electrons.

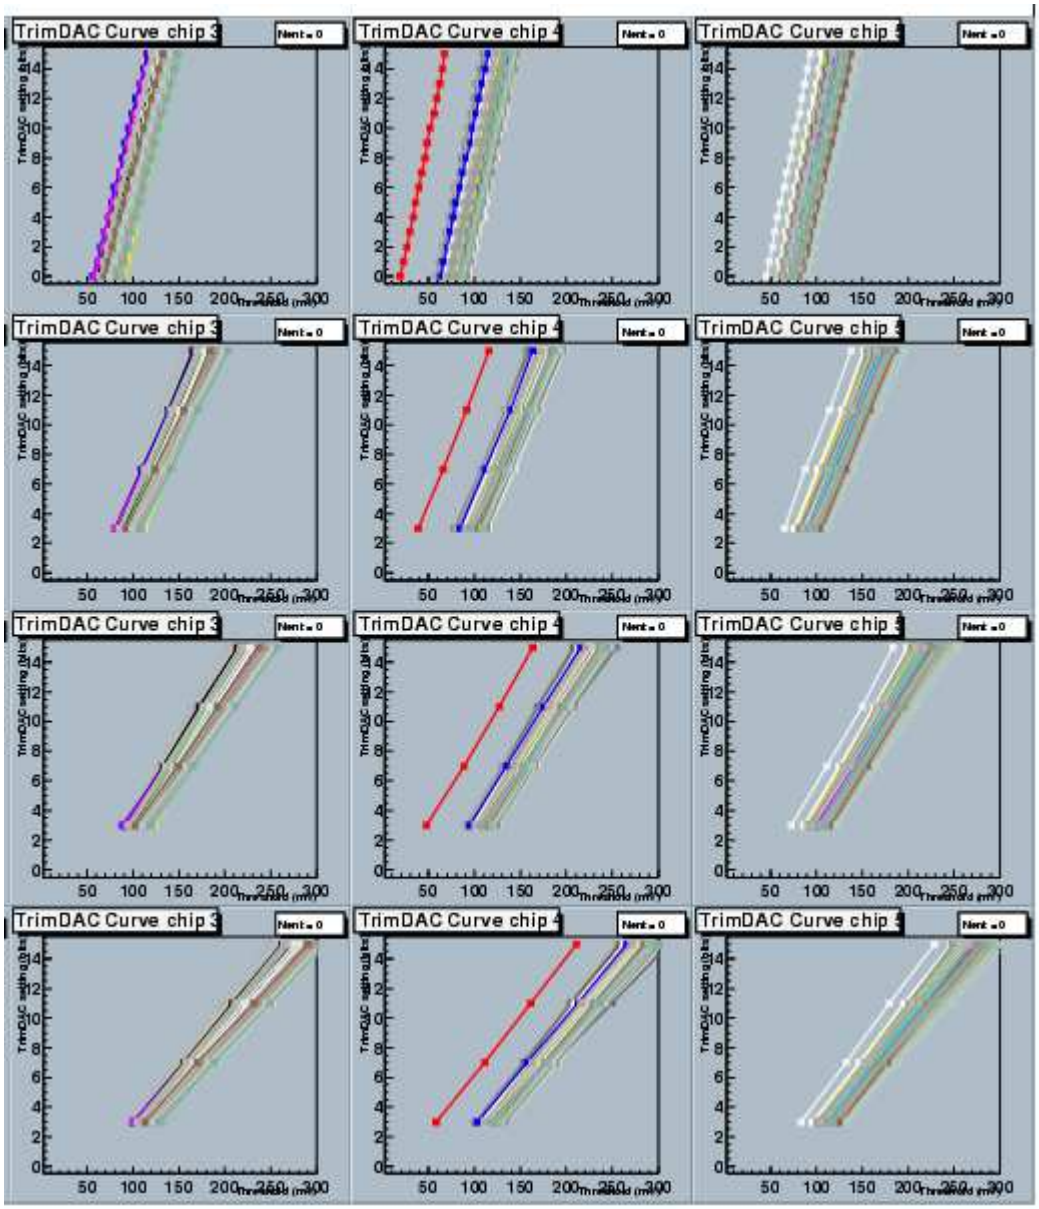

Figure 5.8: TrimDAC plots, showing linearity of trim settings.

The plot shows the results of the threshold scans for all four ranges (the four rows) with one chip in each column. Each plot shows trim setting vs. VT50 point for the 1.0 fC threshold scan. The 16 step points of range 0 are plotted for each channel in the uppermost row, and a straight line is fitted to them to determine the step sizes and offsets. This also works as an important test to verify that there is a linear relationship between trim setting and VT50 point. The three next rows show the scans for range 1, 2 and 3 repsectively.

We see that for chip 4, there is a channel that has a notable offset compared to the other channels.

For each range, the maximum number of channels that can be brought into line with trimming is calculated, and this, together with the parameters of the fitted line and distribution of trimmings of channels, is shown on the plot on figure 5.9.

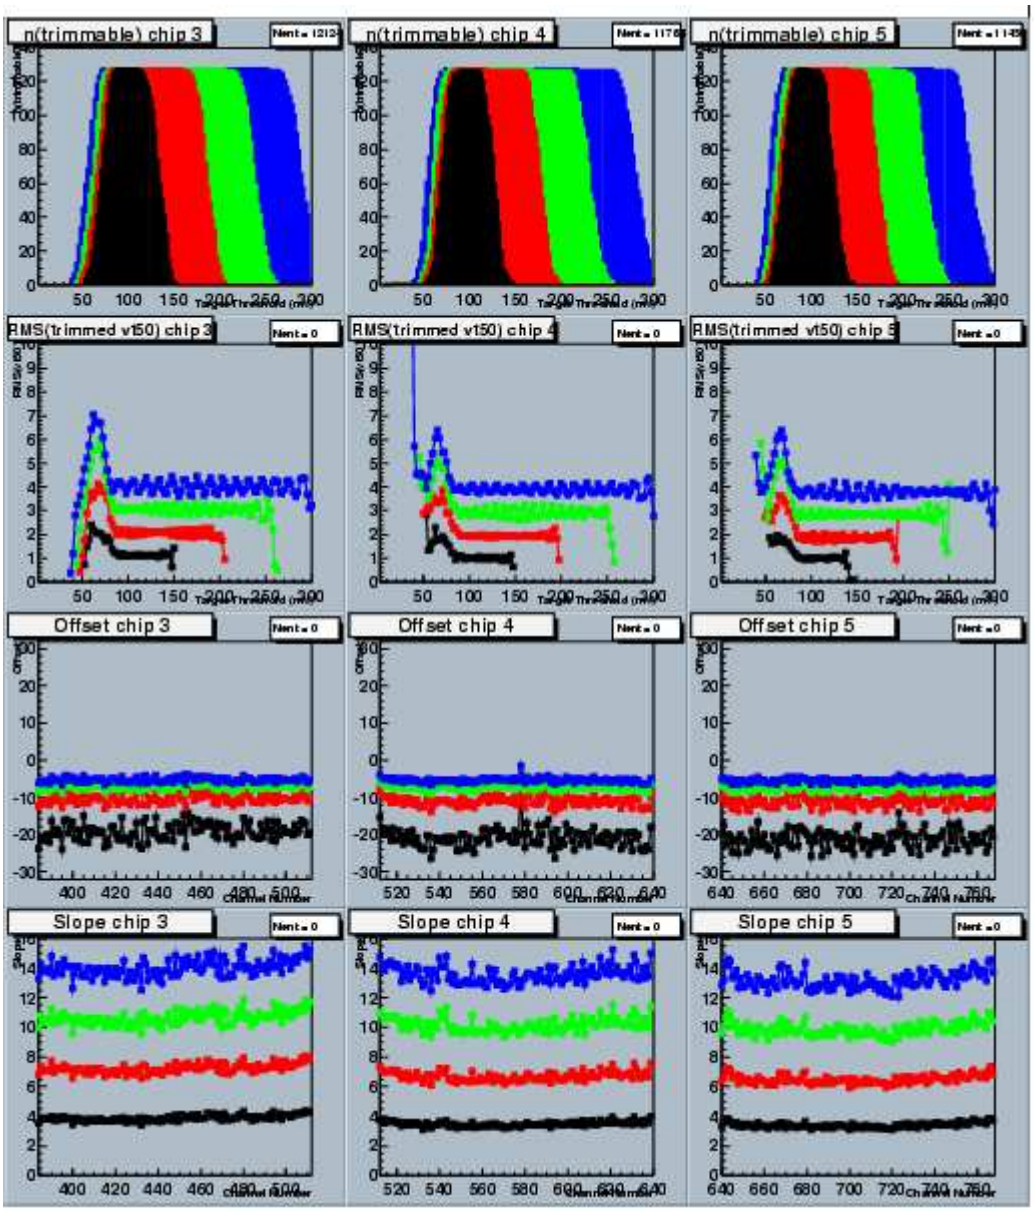

Figure 5.9: TrimDAC plots, showing number of trimmable channels for each range, RMS of trimmed VT50 points, offset from linear fit in figure 5.8 and slope from the same fit.

The uppermost row in figure 5.9 plots trimmable channels (y axis) vs. target threshold (VT50, x-axis) for each range, with range 0 as the black curve. This tells us how many channels can be brought into line for a certain threshold value for a certain range.

For range 0 on chips 3 and 5, we see that for VT50 = 100.0 mV, all 128 channels can be brought into line. If we for example wish to trim chip 3 to have a VT50 of 200.0 mV for a

1.0 fC injected charge, we see that for range 0, no channels are trimmable. For range 1, the curve 'behind' the range 0 curve, approximately 10 channels are trimmable, while for range 2 and 3, all 128 channels are trimmable.

For chip 4, the chip having the channel with the large offset in figure 5.8, only 127 channels are trimmable for range 0. This is not very easy to see from the plot in figure 5.9, but the range 0 curve lies a little lower than the range 1/2/3 curves. This means that for trimming chip 4, we have to use range 1.

We see that the available trim interval is much larger than the interval actually used, this is due to the expected increase in channel-to-channel offset after irradiation.

The next row in figure 5.9 shows the RMS distribution of the trim values for each target threshold. The lowest line is range 0 and the highest range 3.

The two next rows in figure 5.9 shows the offsets [mV] and slope [mV/trim step] of the linear fit vs. channel for each of the ranges in figure 5.8. We can again see that for chip 4, the single channel with the large offset is easy to spot in the plot in row three, while it has a slope similar to the other channels, see the plot in row 4.

The TrimRange macro writes its results to the *results* file, but also creates additional ASCII files containing recommended .trim and .mask files for each range, and finally a .trim and .mask file containing recommended ranges and trims for each chip. This is marked as range –1, and is placed in the *\results* directory, so it has to be manually copied over to the *\config* directory by the user, and renamed to the appropriate filename (see section 4.9) to become the permanent trim setting. However, the new trim settings are loaded into to the SCT module at the end of the macro.

On figure 5.10, a plot of the final trim settings for each channel of the module is shown for the range –1 fit. The two uppermost large plots to the left show trimsettings vs. channels, while the twelve uppermost plots to the right shows the distribution of trim settings for each chip with gauss functions fitted. The two lower plots show the resulting VT50 distribution for 1.0 fC. We see that the target level chosen is 100.0 mV, chosen from the center of the uppermost plots of figure 5.9.

The criteria TrimRange.cpp uses for masking a channel, is based on whether a channel is trimmable or not, that is, if it can be brought into line with the other chips in some range. If a channel has too low or high range, or if the step size or offset differ with more than 3 rms from the mean value for the chip, the channel is recorded in the results file as defective.

The module will not pass the test if it has trim step sizes outside a specified interval, or if some channels are untrimmable.

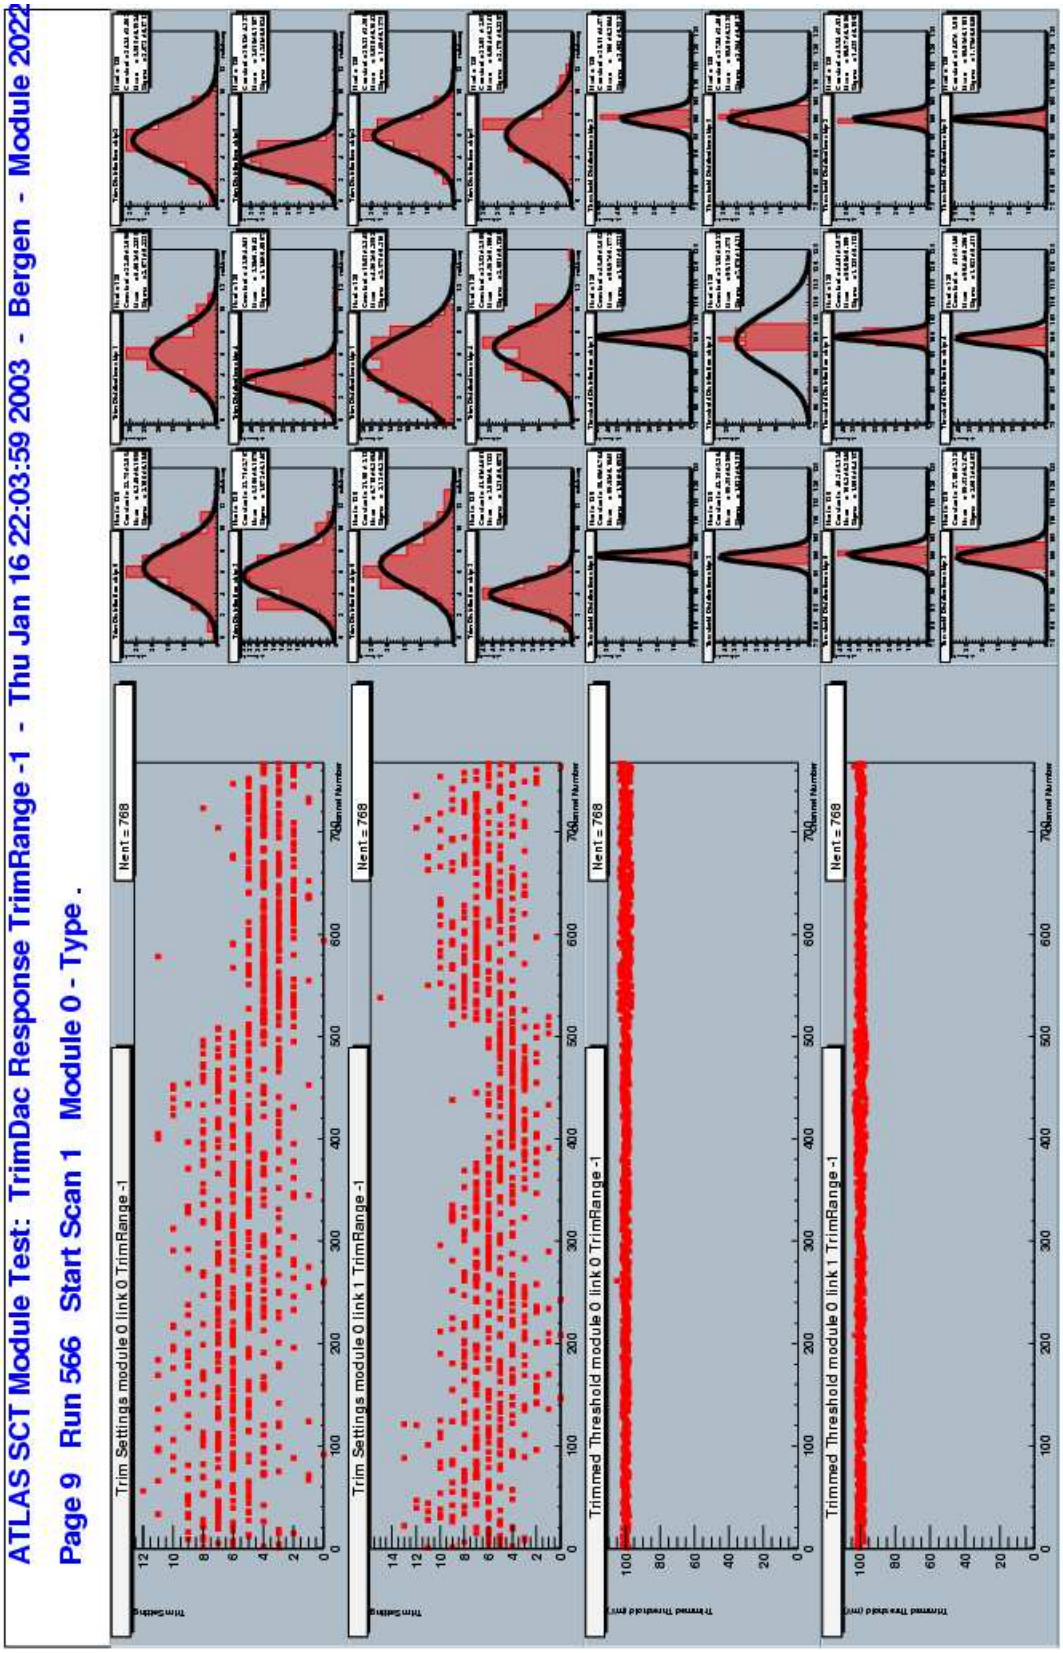

Figure 5.10: Distributions of trim settings. See text for explanation.

Response Curve *(ResponseCurve.cpp)*

This is in principle the same test as ThreePointGain, but this time run with the new trim settings loaded, and with 10 scans for 10 injected charges instead of 3, giving a more precise response curve.

The injected charges are 0.50, 0.75, 1.00, 1.25, 1.50, 2.00, 3.00, 4.00, 6.00, 8.00 fC. The response curve fit used is not linear, but a type of exponential function defined in the function `st_exp(double *x. double *par)` in SCTDAQ. This function has the form:

$$y = p2 + \frac{p0}{1 + e^{-\left(\frac{x}{p1}\right)}} \qquad (6)$$

with y = threshold [mV] and x = input charge [fC]. The three parameters p0, p1 and p2 no longer directly corresponds to a physical quantity, but the slope of the function is still interpreted as the gain:

$$gain = \frac{dy}{dx} = \frac{p0 \cdot \left(e^{-\frac{x}{p1}} \right)}{p1 \cdot \left(1 + e^{-\frac{x}{p1}} \right)^2} \qquad (7)$$

and the offset as the value of the function taken in x = 0.

The fitted function can be seen in the upper row of figure 5.11 as a fitted response curve to the ten VT50 points, having a gradually decreasing slope, corresponding to the fact that the gain (mV/fC) will decrease for larger input charges as the amplifier becomes saturated. That the gain becomes lower with higher input charge can also be seen in the second row in figure 5.11, plotting the calculated gain for each input charge.

Input noise is still computed from the formula (5), and is plotted in the lower rows of figure 5.11. Also notice that for the smallest input charges (<1.0 fC), the noise is noticeably lower. This is due to the fact that for small input charges, the error functions are no longer good approximations to the scans, since the s-curves gets distorted at low thresholds due to noise occupancy. The error function fits thus underestimates the noise, creating the characteristic noise drop for small input charges.

The macro identifies channels which are dead (always 0 hits), stuck (always full occupancy), lo/hi gain, lo/high offset, un/partbonded channels, noisy and inefficient channels, and writes this to the results file, together with the computed values. ResponseCurve.cpp also generates a file [module_name]_RC_[run number].txt, containing the input noise, gain and offset for each channel at 2 fC input charge.

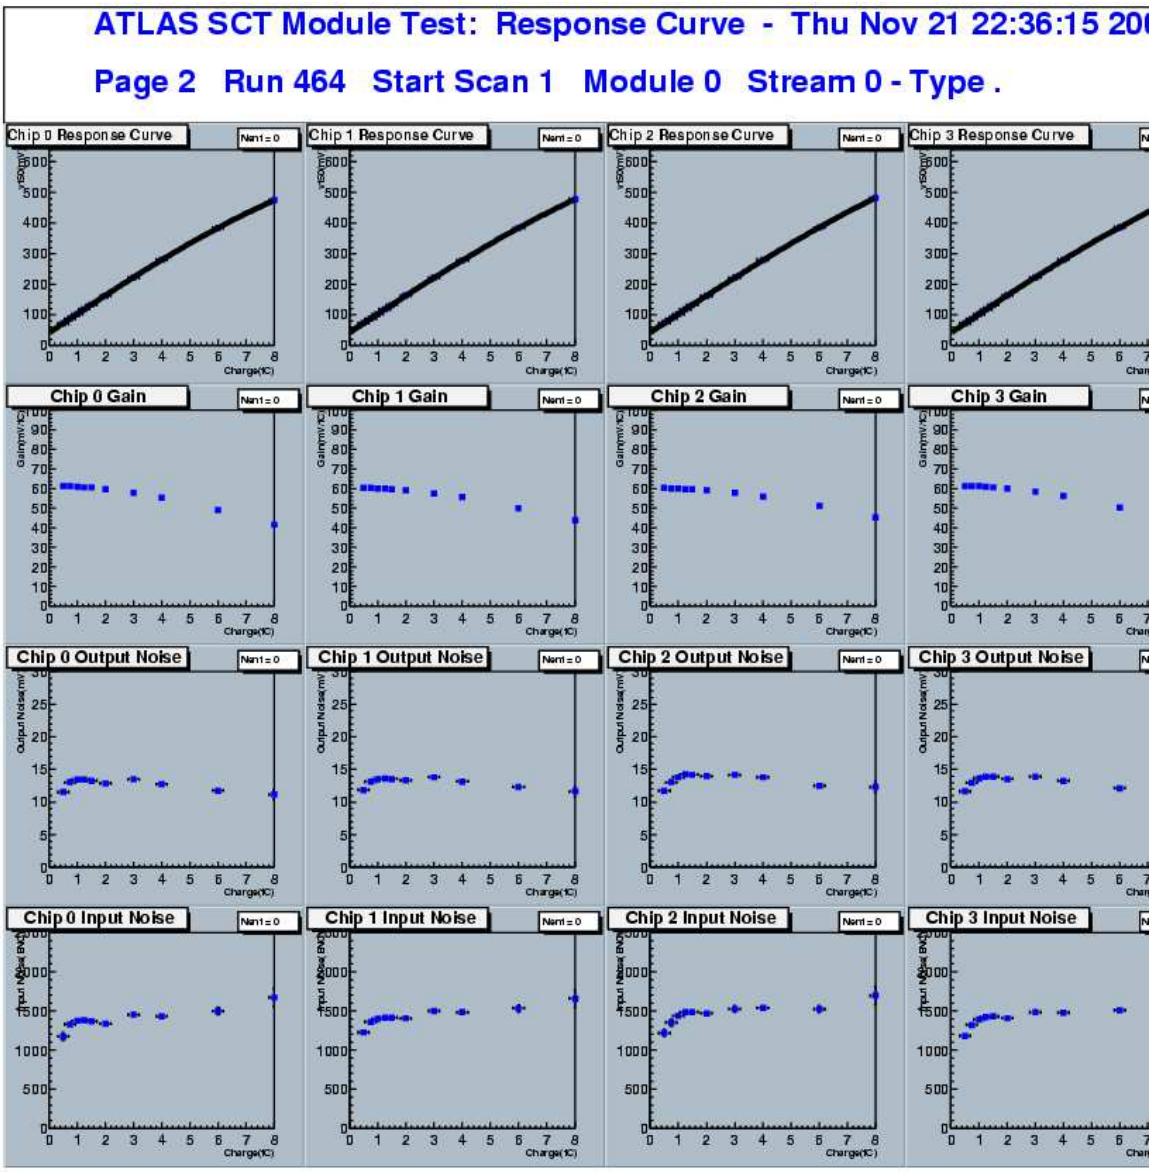

Figure 5.11: Plotted response curve, gain, output- and input noise for each chip for each injected charge.

Also generated is a plot specifically for the 2.0 fC scan, showing VT50 points, gain, offset and input noise vs. channel, shown in figure 5.12.

To the right of each large plot, there are six smaller plots. These show the distribution of values for each chip for the larger plot. Gauss fits are done to the distributions for each chip, with fitted values of Constant, Mean and Sigma. The plotting macros used for plotting figure 5.11 and 5.12 is also used for plotting the results from the ThreePointGain macro.

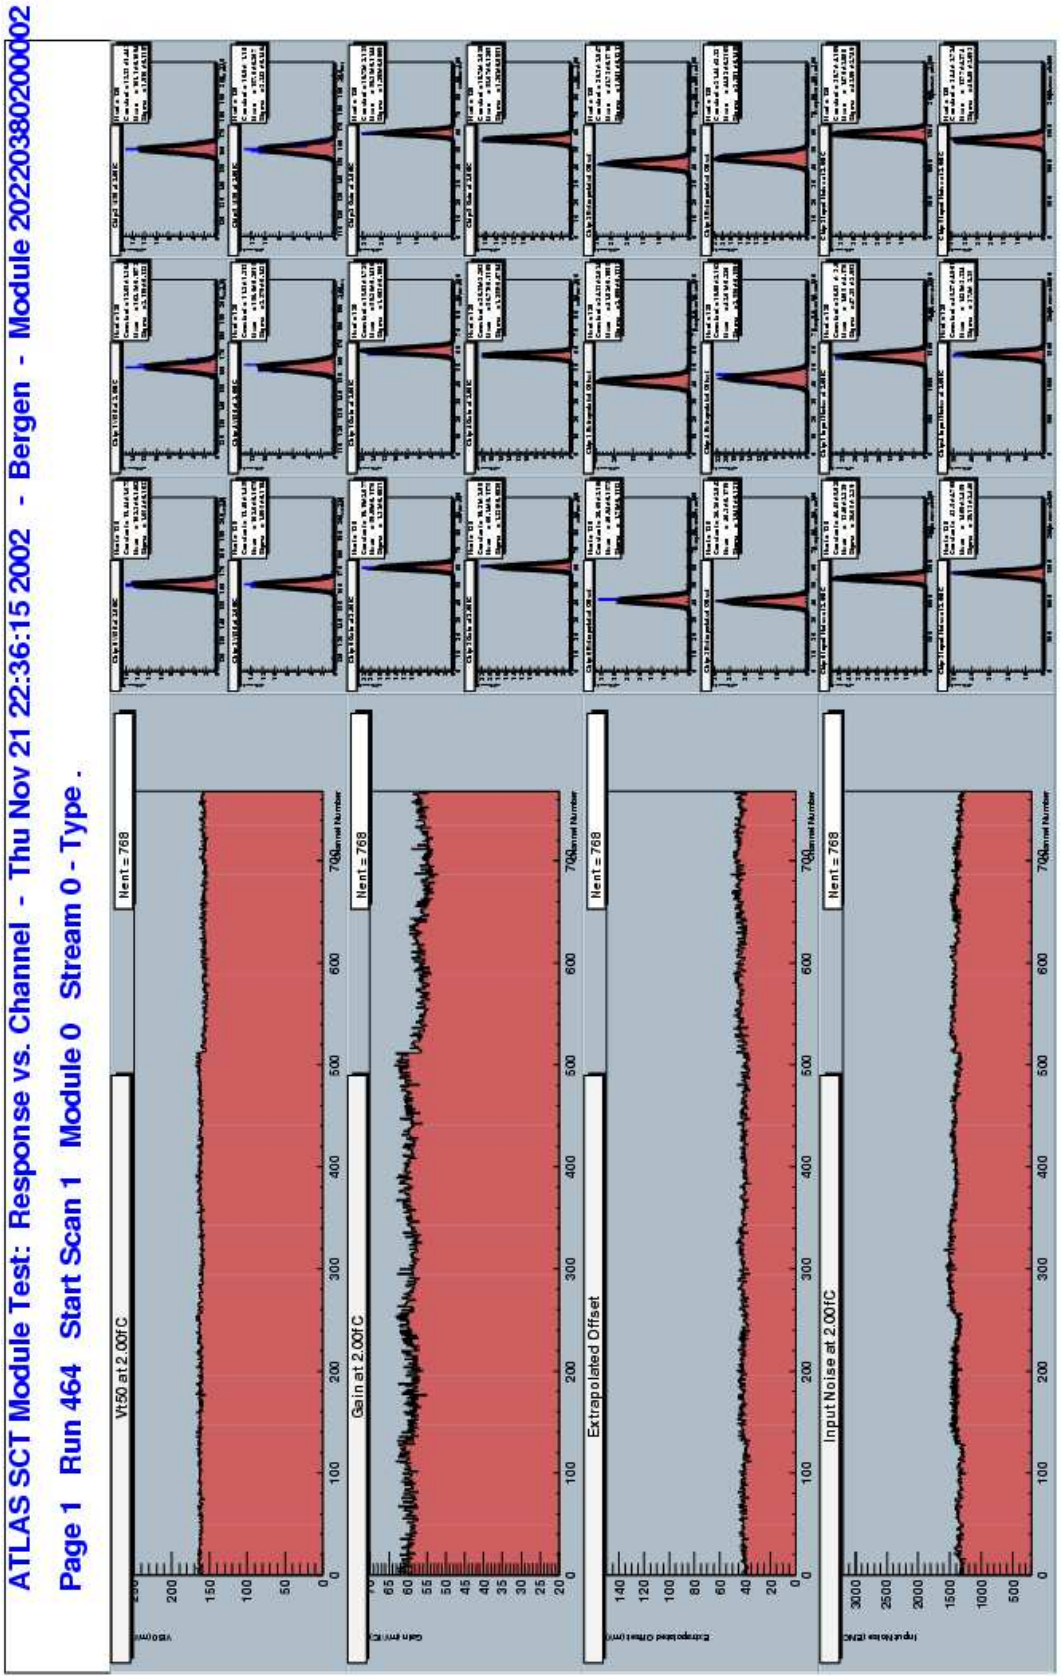

Figure 5.12: Plot for 2.00 fC injected charge generated by the response curve test.

Noise Occupancy *(NO.cpp)*

Noise occupancy is a test that measures the noise in the system at the 1.0 fC threshold. For this, a threshold scan is performed with trigger type ST_NSEBURST, type 7, meaning that the CLOAC only sends single L1 triggers, using no calibration charges. This means that the triggers will only read out hits due to noise being higher than the threshold value.

The threshold scan is performed from 0.0 mV to 150.0 mV in 2.5 mV steps. The number of triggers sent for each scan points varies through the scan, as the occupancy decreases with higher threshold. It starts out by sending 2000 triggers, which should create 2000 hits, because the noise is always higher than 0.0 mV. As the threshold rises, the occupancy decreases, and the number of triggers is increased so that there is a minimum of 50 hits in at least 50 % of the active readout channels. When the number of triggers has reached 1000000, it is no further increased, and when there are only a few hits over the whole detector, the scan is stopped.

On the next page, the plot from a Noise Occupancy scan is shown (figure 5.13). This shows the scan histograms for each stream to the left, and the occupancy vs. threshold to the right, plotted on a logarithmic scale. A numerical value of the occupancy is plotted for the 1.0 fC point. This point is extracted from the response curve of fig 5.11, read from the latest results file, since the response curve gives a relationship between threshold level and input charge.

The 1.0 fC point is again selected as a reference since this is the point at which the modules will operate in ATLAS. The design goal is to have a noise occupancy of ~1e-5 at 1.0 fC, so if a channel has an occupancy higher than 5e-4, it is marked as noisy in the results file. The test also calculates mean and RMS noise occupancy for each chip at 1.0 fC and writes this to the results files. The noise occupancy offset is calculated from the mean of an Erfc fit.

An estimate of input noise in ENC, shown on figure 5.14, is also done with the following technique:

The upper part (threshold ~60 mV and higher) of the NO plot is expected to have the form of a falling gauss function:

$$Occ(x) = C * e^{-\frac{1}{2}\left(\frac{x-M}{\sigma}\right)^2} \qquad (8)$$

where x is the threshold, M mean and Occ the recorded noise occupancy.

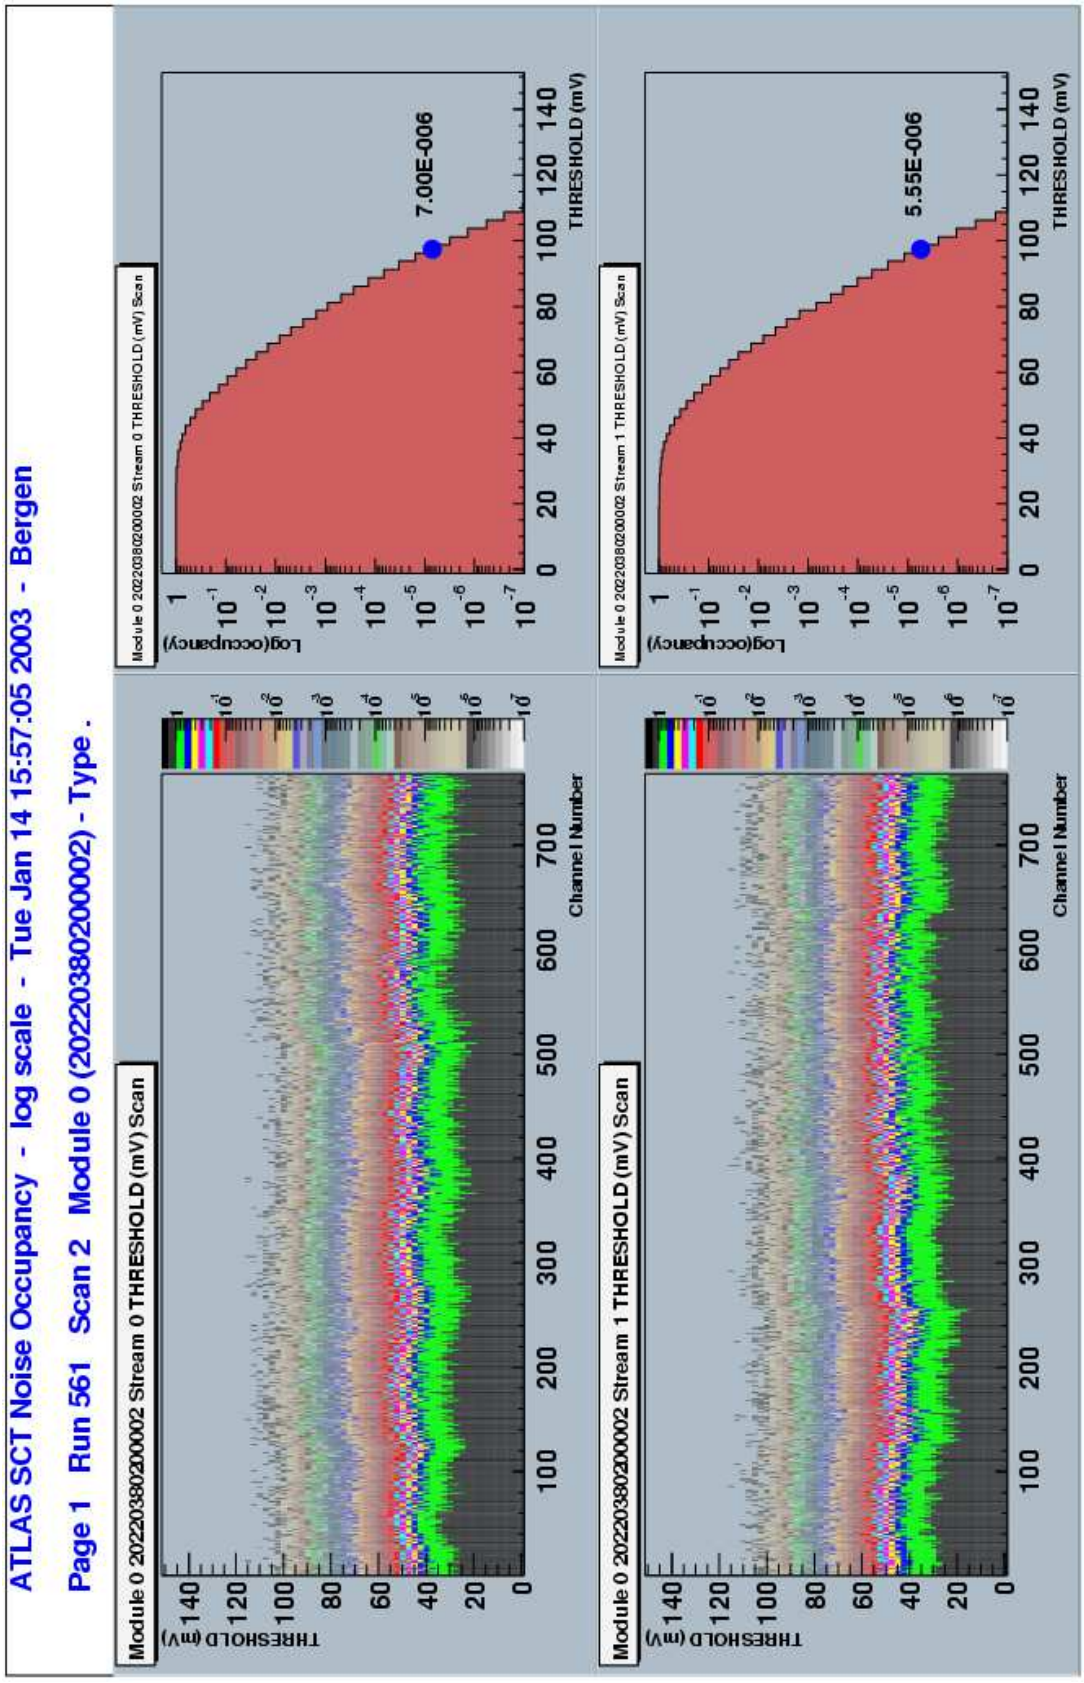

Figure 5.13: Noise occupancy scan plot

By taking the natural logarithm on each side of the expression, one gets:

$$\ln(Occ) = -\frac{1}{2}\left(\frac{x-M}{\sigma}\right)^2 + \ln(C) \tag{9}$$

This could be expressed as a linear function of the form

$$y = az + b \tag{10}$$

where

$$\begin{aligned} y &= \ln(Occ) \\ a &= -\frac{1}{2\sigma^2} \\ z &= (x-M)^2 \\ b &= \ln(C) \end{aligned} \tag{11}$$

The fitted function can be seen in figure 5.14, fitted to a plot of ln(Occupancy) vs. Threshold^2 [fC^2]. The estimate of the input noise is 1367 ENC, calculated from the σ of the linear fit, multiplied with the 1fC/e ratio 6250. The gauss fit (the line) are fitted from 0.3 fC^2, corresponding to ~0.54 fC, which again can be converted to ~60 mV from the Response Curve. This is inside the interval where the noise occupancy should fit the falling edge of gauss a curve.

The estimated input noise should correspond to the input noise measured by the ResponseCurve test, acting as a cross-check of the measurement. By measuring the same parameter by two independent methods, it is also an important check of the reliability of the setup.

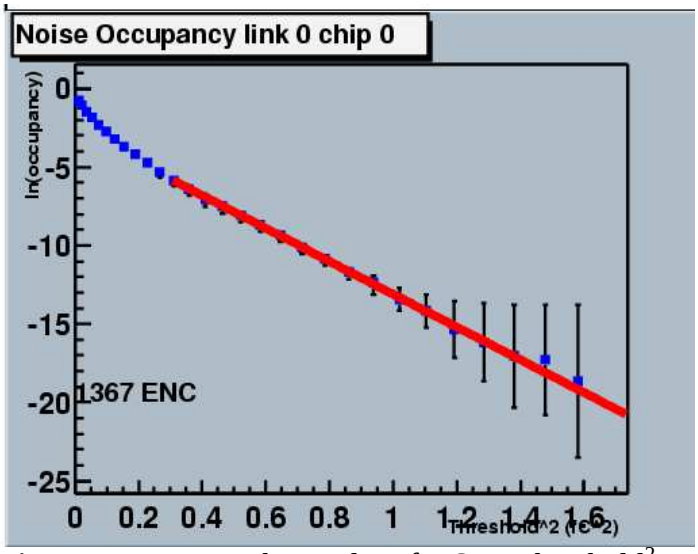

Figure 5.14: Logarithmic plot of NO vs. threshold^2 with linear fit giving ENC estimation.

Timewalk (*TimeWalk.cpp)*

Timewalk was defined in chapter 3 as "*The maximum time variation in the crossing of the time stamp threshold over a signal range of 1.25 fC to 10.0 fC, with the comparator set to 1 fC*". Figure 5.15 shows a graphical view of this definition:

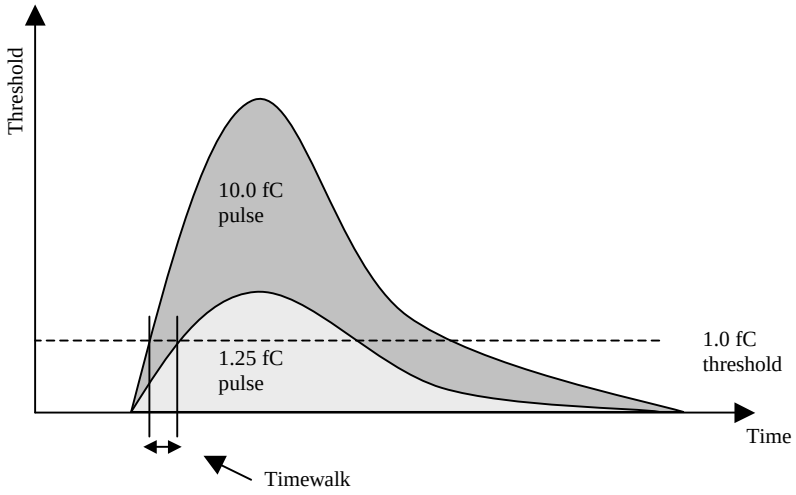

Figure 5.15: Timewalk definition

To determine this difference between timestamp crossings, a strobe delay scan is performed for a set of 10 charges ranging from 1.25 fC and 10.0 fC with the threshold set to 1.0 fC. The strobe delay is scanned from 0 to 63, and the result of this scan for the 10.0 fC charge can be seen in the second row of figure 5.16, as a plot of strobe delay vs efficiency (how large part of the calibration charges the chips recorded). Two error functions are fitted to the rising and falling edge of the plots to determine the 50 % efficiency point, their values are also plotted on the figure. Since the range in which the module records hits should be within a clock unit, the width of the plot corresponds to 25 ns (from 40 MHz clock frequency), giving a relationship between strobe delay step and time.

The uppermost row of figure 5.16 shows the Timewalk curve. This is a plot of Strobe Delay units vs. input charge, and shows for each channel the strobe delay setting at which the VT50 point of the falling edge of the strobe delay scan is reached (as in the second row). The difference between the 1.25 fC point and 10.0 fC is then calculated to get the correct timewalk for each channel. The distributions of timewalk values is plotted vs. Strobe Delay in steps and nS respectively for each chip in the two lower rows.

If a channel has a timewalk less than 5 steps, it is classified as TW_LO, too low timewalk, and correspondingly TW_HI if it has timewalk above 16. Both of these conditions is classified as a defect, and the module will not pass the test if any of these occur for a channel.

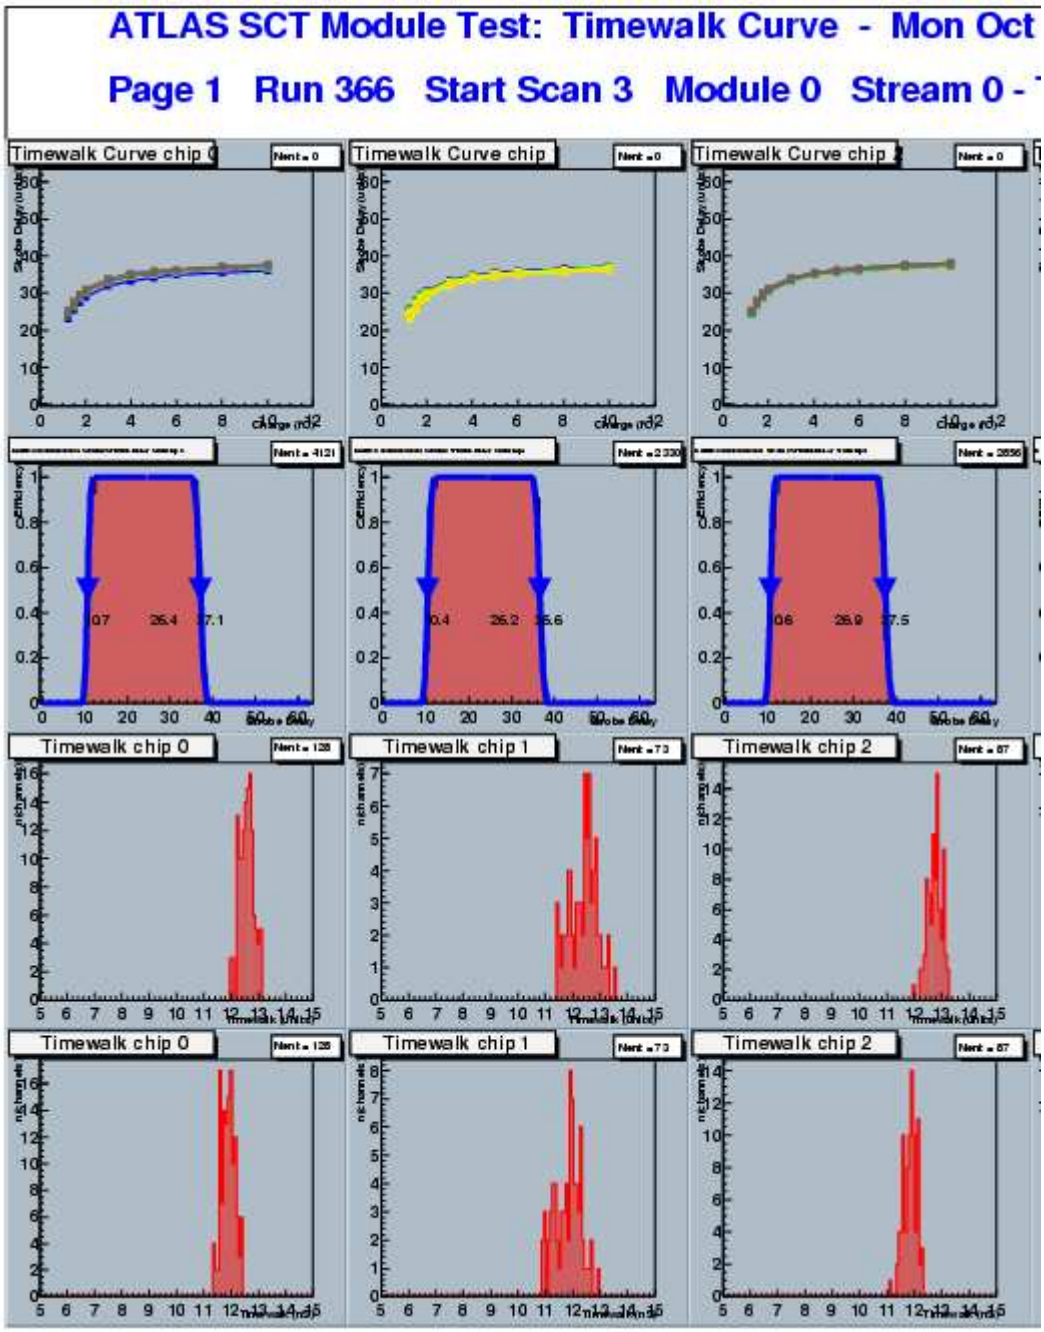

Figure 5.15 : Timewalk plots

This covers the basic functionality of SCTDAQ and the tests used in the QA of the SCT modules.

5.4.5 Pass/Fail criteria for electrical tests

For each test, there are a set of criteria for the module to pass the test. These are summarized in table 5.4.

Test	**PASS criteria**
Hardreset.cpp	No defects.
RedundancyTest.cpp	No defects.
FullBypassTest.cpp	No defects.
PipelineTest.cpp	Max 2 defective channels.
StrobeDelay.cpp	No defects.
ThreePointGain.cpp	- No chip failures. - No chip with gain >100 or <0. - Not more than 8 consecutive bad* strips. - Not more than 15 bad strips in total.
TrimRange.cpp	No TR_RANGE** defects.
ResponseCurve.cpp	Same as ThreePointGain.cpp.
NO.cpp	Criteria not yet set.***
TimeWalk.cpp	No defects.

Table 5.4: Pass criteria for electrical tests.

*) bad strips are strips that are faulty or having bad performance, classified as follows:

- DEAD – No output
- STUCK – Continuous output
- LO_GAIN – gain < (0.75*mean_chip_gain)
- HI_GAIN – gain > (1.25*mean_chip_gain)
- LO_OFFSET – Offset < -100 mV
- HI_OFFSET – Offset > 120 mV
- UNBONDED – Noise < 800 ENC
- PART_BONDED – Noise <1100 ENC
- NOISY – Noise > (1.15*mean_chip_noise)
- INEFFICIENT – Max efficiency < 100 %

**) TR_RANGE- A channel is marked with TR_RANGE defect if the step size is not inside a specified expected interval. Channels could also have the defects TR_STEP and TR_OFFSET, which is deviations from the mean step and offset values, but these are not critical.

***) For the Noise Occupancy test, no PASS criteria is yet set in SCTDAQ. There is however the design goal of 1e-5 noise occupancy, so if a channel has noise occupancy higher than 5e-4, it is marked as NO_HI in the results file.

For the tests allowing some defects on the module, it is possible that additional criteria will be implemented at a later stage, when there are better statistics on overall module quality.

Chapter 6: Noise reduction in the setup

As the SCT barrel module is very sensitive to noise, minimation of noise in the setup through grounding, shielding and filtering has proved crucial to the precision of the system.

Noise can come from many sources, both through signal wires, ground and air. In the following sections, the steps that have been taken for minimizing noise will be discussed.

6.1 Noise from low voltage power supply (SCT-LV3 VME module)

The SCT LV3 VME module delivers low voltage DC to both the analogue and digital part of the SCT module. On the DC lines, there were large spikes showing up every 10 µs. An oscilloscope image of these, measured between the analogue DC line and its ground line on the output of CH0 of SCT-LV3, is shown in figure 6.1.

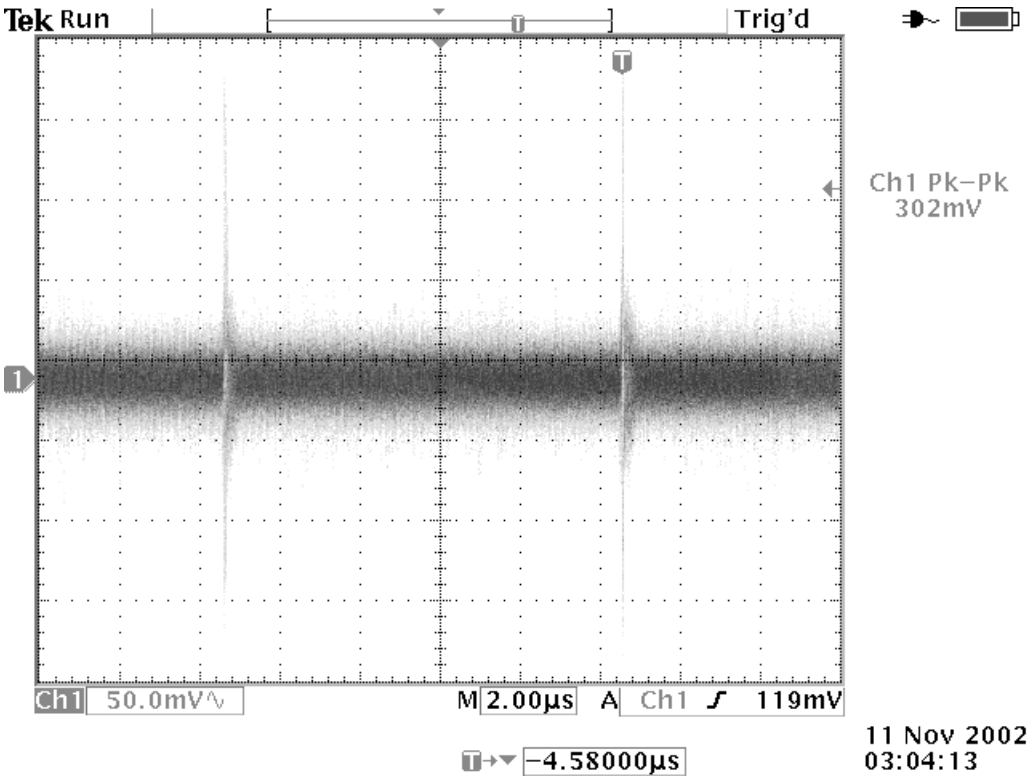

Figure 6.1: Measurement between analogue DC and analogue ground on SCT LV3 CH0 ouput (pins 4 and 2), showing spikes of amplitude 302 mV and period 10 µs. Resolution is 2.00µs/div in x and 50.0 mV/div in y.

By looking at the transients at a higher time resolution, it can be seen that they are composed of high frequency signals, figure 6.2.

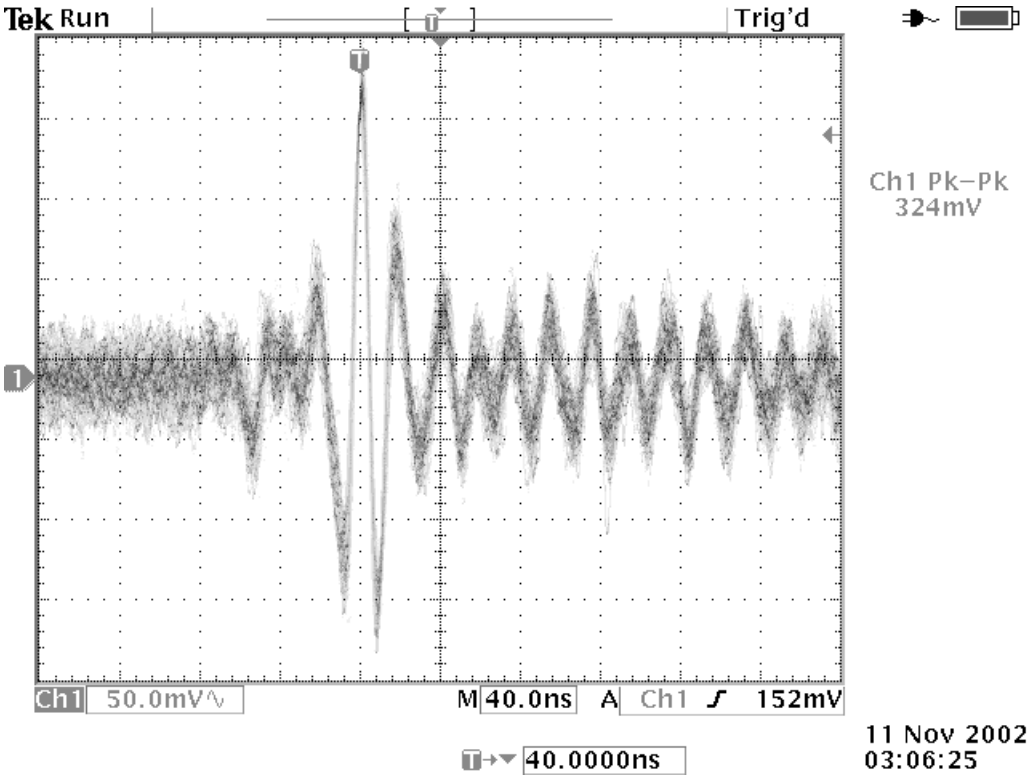

Figure 6.2: LV3 noise at higher time resolution, 40.0 ns/div and 50.0 mV/div.

We see that the signal has a period of approximately 20-30 ns, resulting in a frequency around 30-50 MHz. Since 40 MHz is the frequency at which the SCT module works, it is sensitive to noise in this frequency domain. These spikes seemed to create a 'tail' on the noise occupancy plots, see figure 6.7.

To diminish the amplitude of these spikes, a filter was created, using standard components ELFA 65-641-40 EMI (Electro Magnetic Interference) - guard 22 nF 12 V. This is a simple filter, which consists of two solenoids and a varistor which directs the high frequency components of the signal to ground. It has effective damping in the 1-1000 MHz area, with highest damping around 30 MHz, which is the region in which the SCT module is sensitive to noise. A schematic view of the filter is shown on figure 6.3.

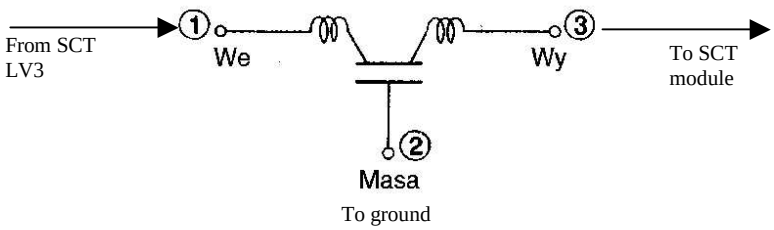

Figure 6.3: 65-641-40 EMI-guard 22 nF filter with connections.

The LV3 filter was realized as a PCB board with two 25 pin connectors. Each channel of the flat cable was laid out on the board, but only the analogue and digital DC lines, their respective grounds and sense wires were filtered (pins 1-4, 11, 12, 14, 17, 23 and 24). The rest of the lines in use; HV bias, temperature sense and digital signals (pins 6, 7, 10, 11, 19, 20, 22) runs directly through the PCB. The lines not in use from LV3 were not connected. The middle feet of all components (point 2 on figure 6.3) are connected on the middle of the PCB, which in turn is grounded to the VME crate. A picture of the filter is shown below in figure 6.4:

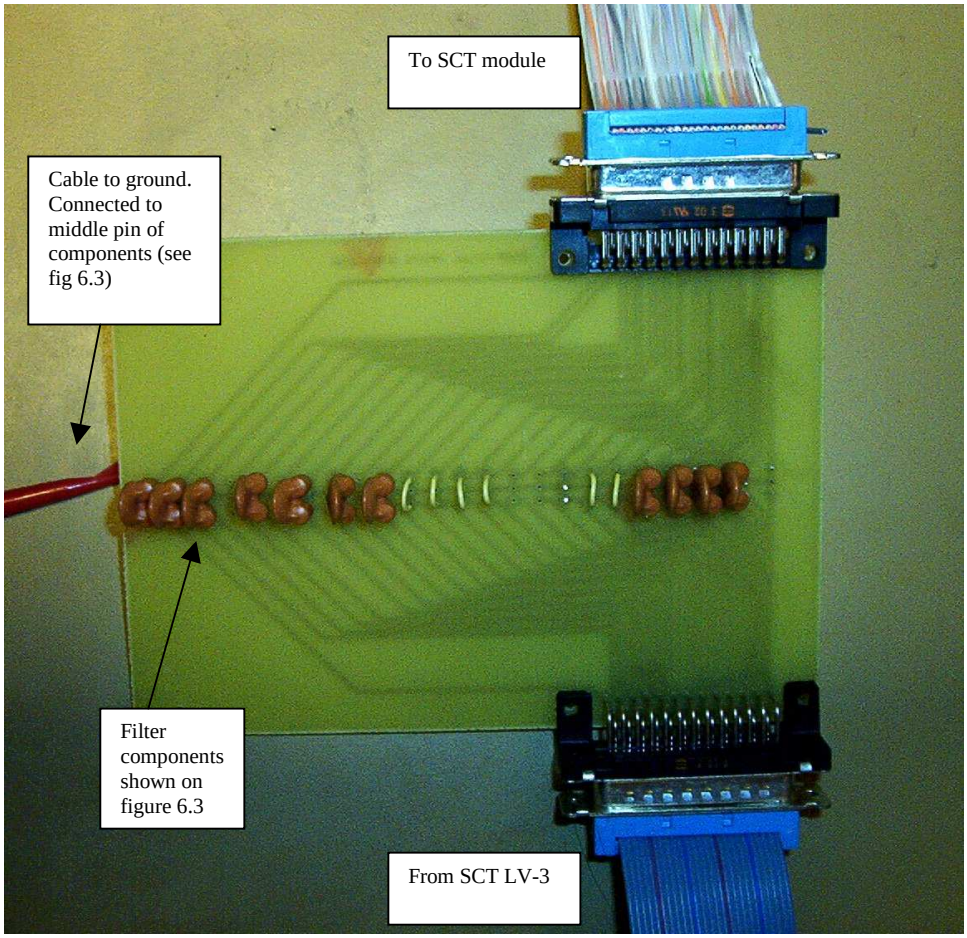

Figure 6.4: LV-3 filter.

When first testing the filter, a new problem showed up. During start up of the system, LV3 reported the error *LV trip*. LV3 has an internal limit for how high current it can send out, to protect the system against being short circuited. If this level is crossed, the DC channels are turned off, and an LV trip is reported. When connecting the filter to the system, too high current was drawn during the power-up of the module. The LV trip also occured when powering up at low temperatures (thermistor temperature< 0°C), independent of the filter.

This proved to be a known problem from some other testing sites. When testing the filter with another module, no LV trips were reported, so the tendency to do LV trips can vary from module to module.

An upgraded version of the LV3 firmware, supposed to solve the problem, was available. The upgrade was intended for the four 28 pin EPROMs marked SCT LV3 rev. 6S, seen on figure 4.8, of type AMD AM27C512, which is a 64-bit x 8k CMOS EPROM. Four new chips were programmed with the new firmware, returning checksums[11] 00F7D99F. The old chips were kept as a backup.

This solved the LV trip problem, and the filter proved very efficient in reducing the noise. Figure 6.5 and 6.6 shows an oscilloscope measurement between analogue DC and analogue ground after the filter was applied:

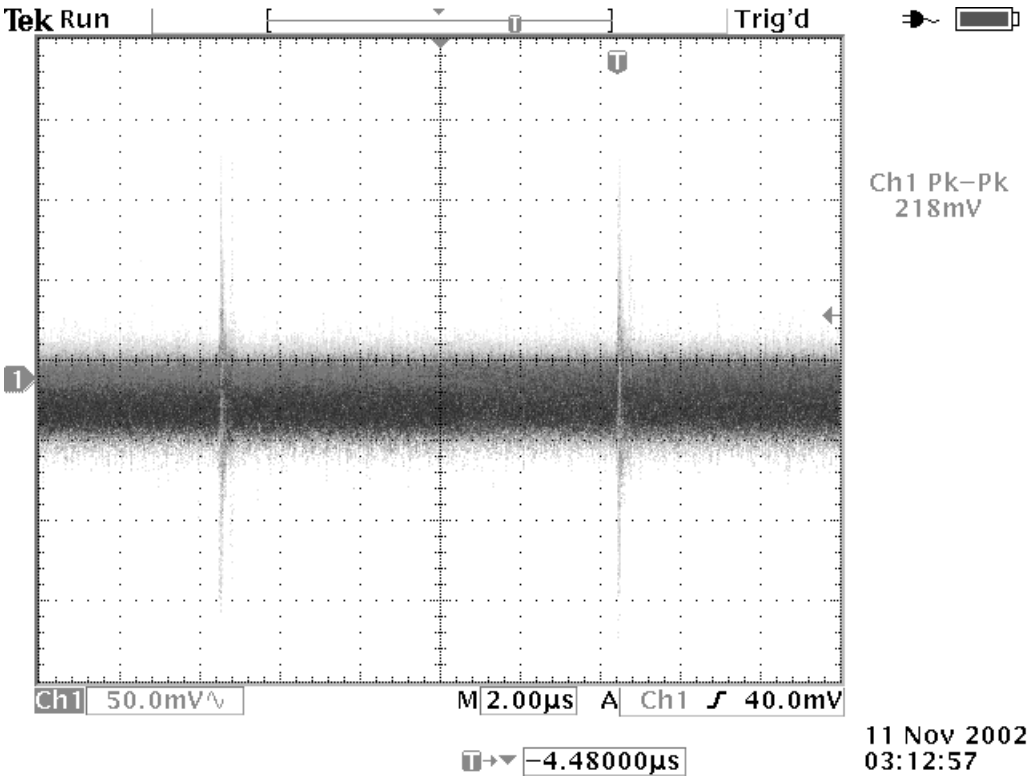

Figure 6.5: Analogue DC from LV3, measured on the output of LV3 filter. We see that the amplitude of the spikes are reduced from ~300 mV to ~200 mV.

[11] The checksum is the 64-bit sum of all the 64 bit locations in the EPROM. All chips programmed with the same software should return the same checksum, making it a simple check for programming errors.

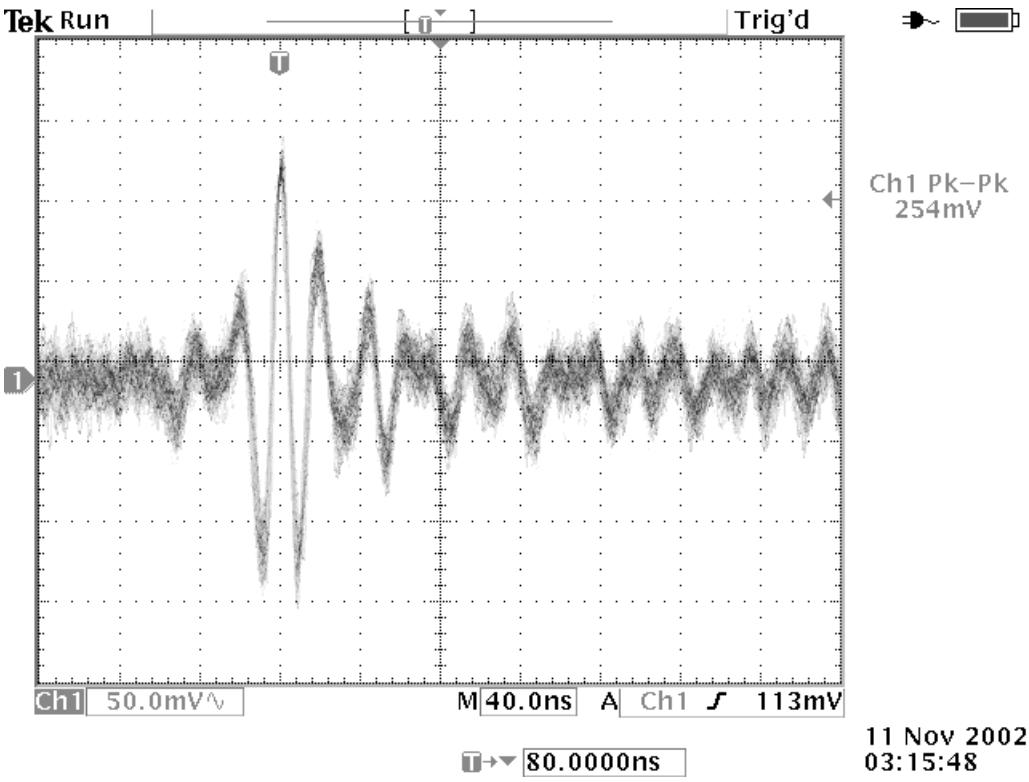

Figure 6.6: Measurement of analogue DC with LV3 filter at higher time resolution (40.0 ns/div).

6.1.1 LV3 filter measurements

As already mentioned, the spikes seemed to create a 'tail' on the noise occupancy plots, as can be seen on figure 6.7.

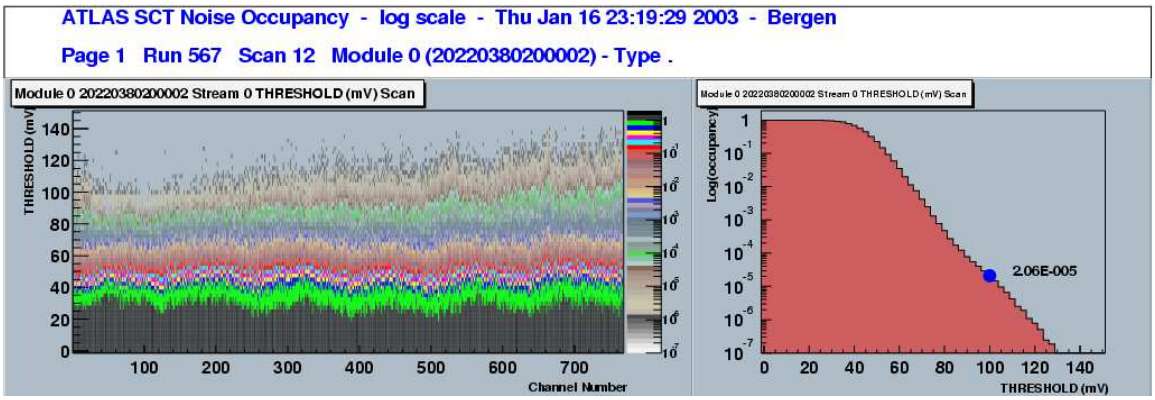

Figure 6.7: Noise occupancy scan without LV3 filter connected, performed at hybrid temperature 0°C and 200 V detector bias.

The noise occupancy on figure 6.7 is quite high, 2.06 x 10^{-5}, and we see that the falling plot to the right do not fall off constantly, but starts to fall slower around 80 mV

threshold, creating a 'tail' on the plot. We also see from the scan histogram to the left that the channels to the right (far away from the master chip), has a higher noise than the channels to the left (near or on the master chip). The plot suggests that we have two gauss curves superpositioned on to each other, with different mean and sigma.

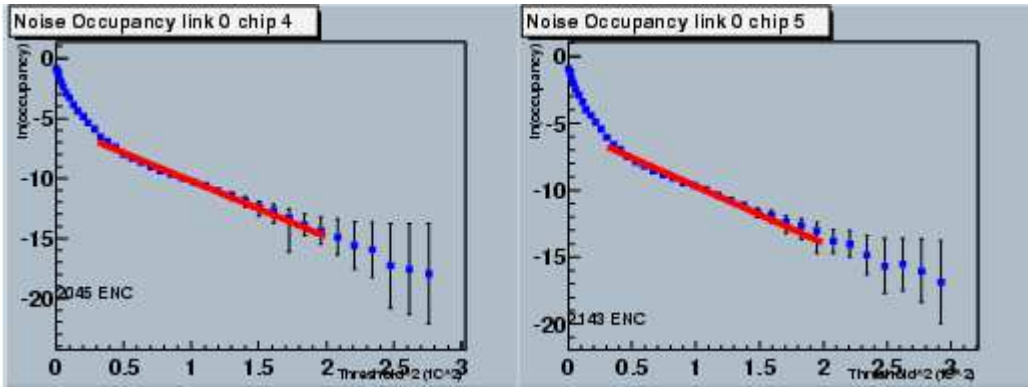

Figure 6.8: Linear fit of noise occupancy without filter, same scan as figure 6.7. Estimated noise is ENC 2045-2143.

Figure 6.8 shows the linear noise occupancy fit without filter (see chapter 5 for a discussion of the fitting methods used). The plot itself is clearly not linear, but supports the assumption that we have two superpositioned gauss curves, giving us two 'linear' components, each dominating in its interval.

We also see a quite high estimated input noise, in the order of 2000 ENC. The input noise measured under the same conditions by the response curve was 1248 and 1256 respectively, so we have a difference of ~800 ENC, with the difference being largest for the chips farthest away from the master chip, see the plots in figure 6.11 and 6.12.

Since the linear fit from which sigma is extracted is done in the interval where the 'extra' gauss curve resides, we see that we get large deviations between the input noise measured by the two different methods. This comes from the fact that the noise occupancy scans are very sensitive to external noise around the 1.0 fC threshold, since this is where the intrinsic noise of the module vanishes, making any external noise visible.

In contrast, the response curve does its input noise calulations based on scans with injected charge ranging from 0.5 fC to 8.0 fC (see chapter 5). Except for the two first charges, the charges are of the same magnitude, or larger than the noise generated from LV3, which can be seen to vanish at a threshold value of approximately 130 mV. In addition, the values output by SCTDAQ are the input noise calculated at the 2.0 fC point. The input charge in this case overshadows the noise generated by LV3, resulting in input noise values unaffected by the external noise.

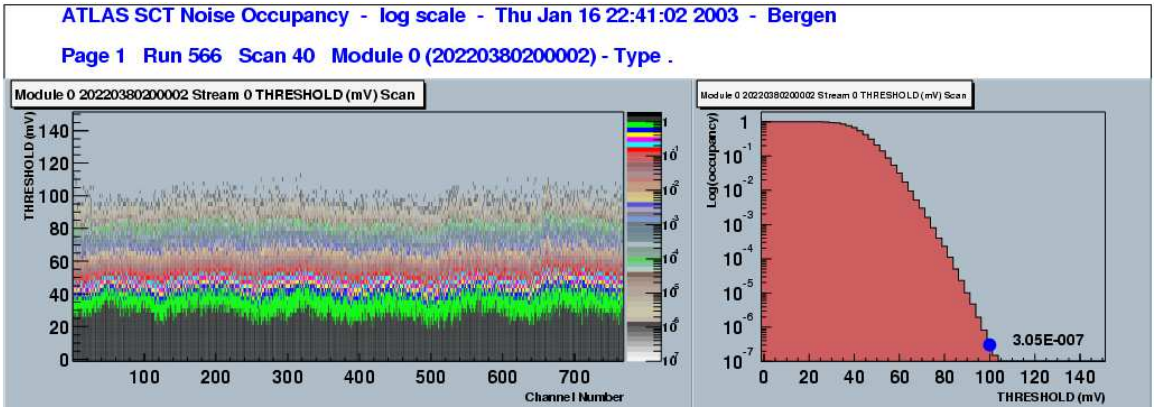

Figure 6.9: Noise occupancy scan at 0°C with LV3 filter.

Figure 6.9 shows the plot for a scan run with the LV3 filter, also at 0°C. This shows a very low noise occupancy at 1.0 fC; 3.05 x 10^{-7}, which is far below the critical limit of 5 x 10^{-4}. We see that the plot falls off nicely, and the scan histogram shows an even distribution of noise over all channels.

On figure 6.10 the corresponding linear fit is shown. We see that the plot is far more linear this time, and the fit is much better in line with the entire range of occupancy measurements. The noise estimated is 1294 and 1311 ENC, which is in agreement with the input noise measured by the Response Curve test under the same conditions; 1234 and 1239.

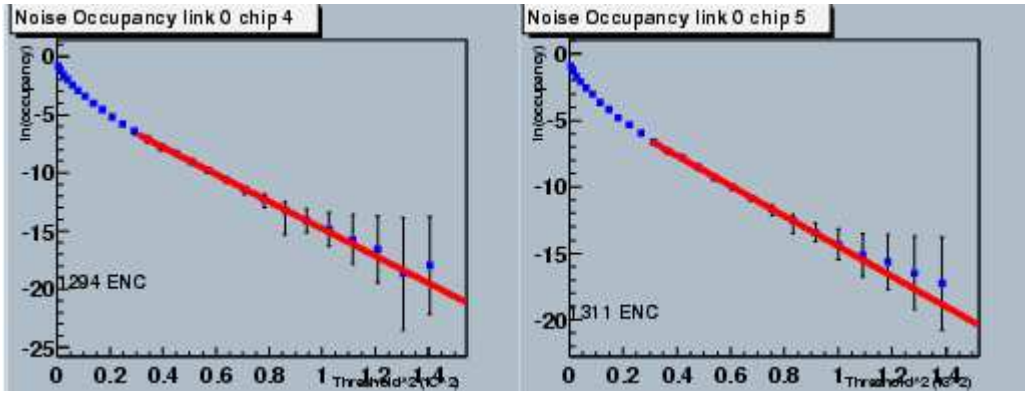

Figure 6.10: Linear fit of noise occupancy with filter, same scan as on figure 6.6. Estimated noise is now ENC 1294-1311.

Tables and plots showing data from Response Curves and noise occupancy scans performed with and without filter for hybrid temperatures +27°C and 0°C is presented in figures 6.11 and 6.12, with corresponding data in tables 6.1, 6.2 and 6.3. Errorbars are not plotted for reasons of clarity, but are listed in the tables.

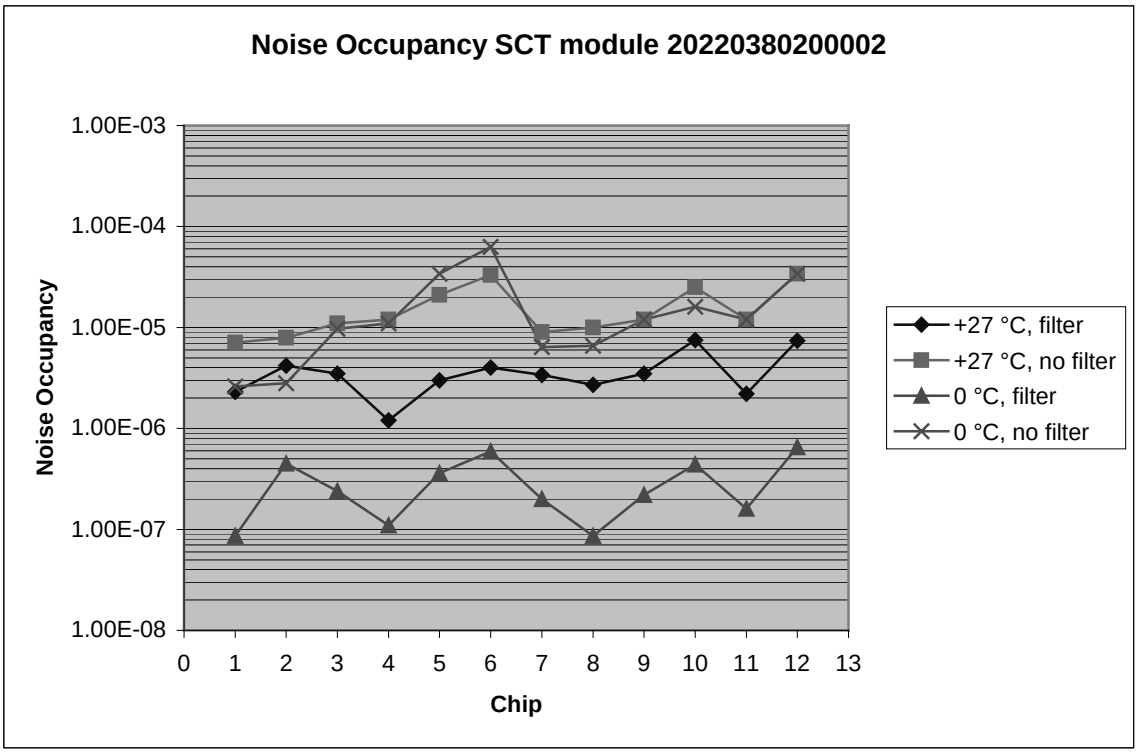

Figure 6.11: Noise occupancies at 1.0 fC threshold. Measurements done with hybrid temperature +27°C and 0°C, with and without filter (data in table 6.1).

Test	Noise Occupancy							
Runnum-scan	564-14		563-42		566-40		567-12	
Condition	+27 °C, filter	Error	+27 °C, no filter	Error	0 °C, filter	Error	0 °C, no filter	Error
Chip 0	2.30E-06	1.52E-06	7.10E-06	2.66E-06	8.60E-08	2.93E-07	2.60E-06	1.61E-06
Chip 1	4.20E-06	2.05E-06	7.90E-06	2.81E-06	4.50E-07	6.71E-07	2.80E-06	1.67E-06
Chip 2	3.50E-06	1.87E-06	1.10E-05	3.32E-06	2.40E-07	4.90E-07	9.70E-06	3.11E-06
Chip 3	1.20E-06	1.10E-06	1.20E-05	3.46E-06	1.10E-07	3.32E-07	1.10E-05	3.32E-06
Chip 4	3.00E-06	1.73E-06	2.10E-05	4.58E-06	3.60E-07	6.00E-07	3.40E-05	5.83E-06
Chip 5	4.00E-06	2.00E-06	3.30E-05	5.74E-06	5.90E-07	7.68E-07	6.30E-05	7.94E-06
Chip 6	3.40E-06	1.84E-06	9.00E-06	3.00E-06	2.00E-07	4.47E-07	6.40E-06	2.53E-06
Chip 7	2.70E-06	1.64E-06	1.00E-05	3.16E-06	8.60E-08	2.93E-07	6.60E-06	2.57E-06
Chip 8	3.50E-06	1.87E-06	1.20E-05	3.46E-06	2.20E-07	4.69E-07	1.20E-05	3.46E-06
Chip 9	7.50E-06	2.74E-06	2.50E-05	5.00E-06	4.40E-07	6.63E-07	1.60E-05	4.00E-06
Chip 10	2.20E-06	1.48E-06	1.20E-05	3.46E-06	1.60E-07	4.00E-07	1.20E-05	3.46E-06
Chip 11	7.40E-06	2.72E-06	3.40E-05	5.83E-06	6.50E-07	8.06E-07	3.40E-05	5.83E-06
Mean	3.74E-06	1.88E-06	1.62E-05	3.88E-06	2.99E-07	5.19E-07	1.75E-05	3.78E-06

Table 6.1: Data for plot in figure 6.11. Errors calculated from binomial statistics.

We clearly see that the scans done with the LV3 filter applied have the lowest noise occupancy. The two scans done without filter show approximately the same noise occupancy values, with a mean value of order 1×10^{-5}. This is a clear indication that the LV3 noise almost completely dominates the temperature dependent noise occupancy. For

the scans done with filter, the one at 0°C has, as should be expected, the lowest noise occupancy.

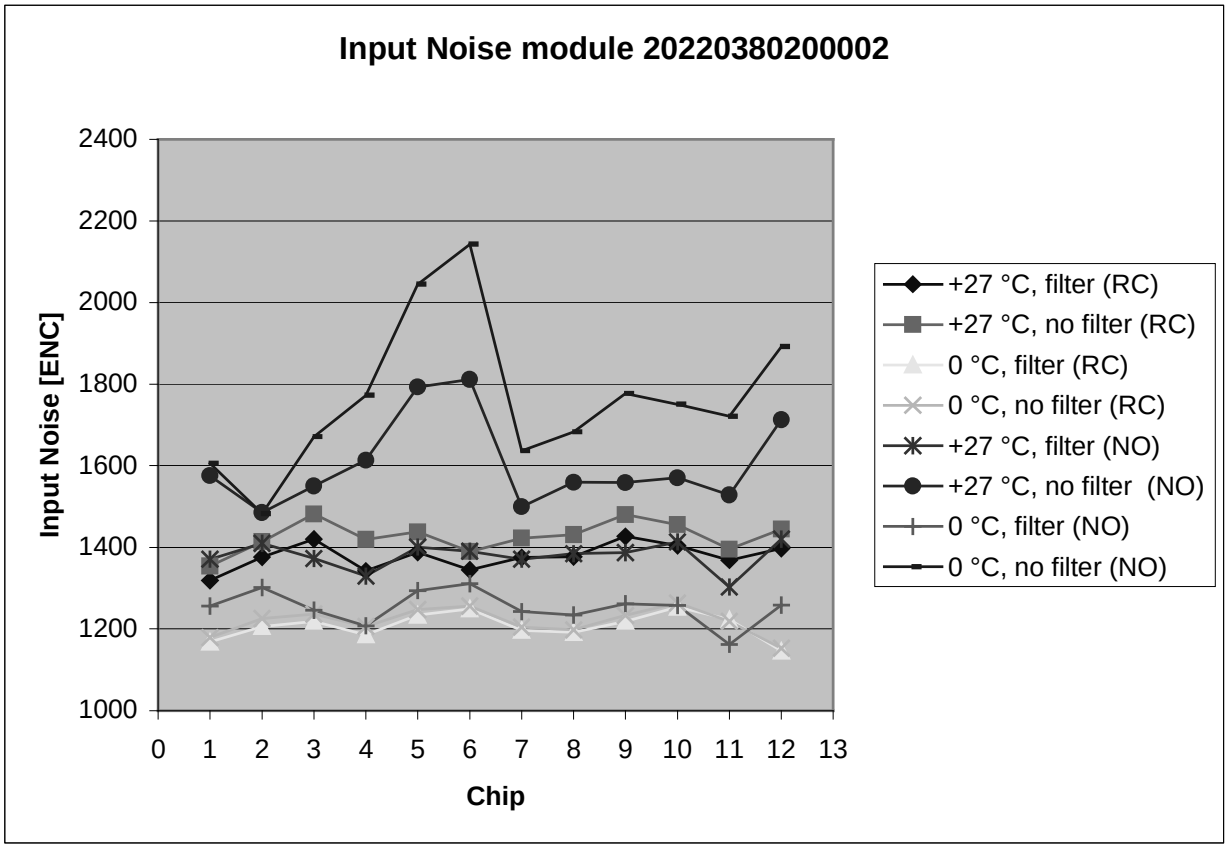

Figure 6.12: Input noise in ENC under various conditions. The lines marked with RC are input noise calculated from the response curve, while the ones marked with NO are the estimated input noise from noise occupancy. (Data in table 6.2 and 6.3)

Test	Response Curve, ENC							
Temp	+27 °C		+27 °C		0 °C		0 °C	
Runnum-scan	564-3		563-32		566-29		567-1	
Condition	Filter (RC)	Error	No filter (RC)	Error	Filter (RC)	Error	No filter (RC)	Error
Chip 0	1319	34	1354	36	1168	46	1179	44
Chip 1	1376	49	1414	51	1206	55	1226	56
Chip 2	1420	40	1482	45	1219	52	1237	49
Chip 3	1342	35	1419	42	1186	42	1206	42
Chip 4	1388	34	1438	34	1234	64	1248	61
Chip 5	1345	50	1390	53	1249	103	1256	106
Chip 6	1375	28	1423	33	1197	43	1205	43
Chip 7	1377	31	1431	35	1192	48	1198	45
Chip 8	1427	26	1480	37	1219	58	1233	52
Chip 9	1404	51	1456	58	1254	63	1264	60
Chip 10	1369	41	1396	44	1223	44	1219	44
Chip 11	1397	38	1445	37	1146	60	1153	60
Mean	1378.25		1427.33		1207.75		1218.67	

Table 6.2 : Data for plot in figure 6.12 (errors is RMS distribuition of chip noise)

Test	Noise Occupancy, estimated ENC							
Temp	+27 °C		+27 °C		0 °C		0 °C	
Runnum-scan	564-14		563-42		566-40		567-12	
Condition	Filter (NO)	Error	No filter	(NO) Error	Filter (NO)	Error	No filter (NO)	Error
Chip 0	1371	20	1576	20	1256	20	1606	20
Chip 1	1410	‘	1485	‘	1302	‘	1483	‘
Chip 2	1373	‘	1550	‘	1246	‘	1671	‘
Chip 3	1330	‘	1614	‘	1208	‘	1773	‘
Chip 4	1401	‘	1793	‘	1294	‘	2045	‘
Chip 5	1391	‘	1812	‘	1311	‘	2143	‘
Chip 6	1371	‘	1500	‘	1243	‘	1637	‘
Chip 7	1385	‘	1560	‘	1234	‘	1683	‘
Chip 8	1387	‘	1559	‘	1262	’	1777	‘
Chip 9	1413	‘	1571	‘	1258	‘	1751	‘
Chip 10	1303	‘	1528	‘	1162	‘	1721	‘
Chip 11	1420	‘	1713	‘	1259	‘	1892	‘
Mean	1379.58		1605.08		1252.92		1765.17	

Table 6.3 : Data for plot in figure 6.12, errors estimated from uncertanty of linear fit.

We see from the plot in figure 6.12 that the filter does not have a large influence on the response curve measurements; the differences with or without filter are less than 50 ENC for both temperatures. For the ENC estimated from the linear fit of the logarithmic noise occupancy plots, there are however large differences. We see that for both temperatures, the use of filter lowers the estimated values with several hundred ENC, most clear for 0°C where there is a mean of 500 ENC difference between running with filter and without.

The estimated ENC of the NO-scans with filter are now in nice agreement by the input noise measurements from the response curve, as should be expected.

6.1.2 Using two filters in serial connection

While the LV3 filter removes much of the noise, we see that it does not eliminate the high frequency spikes altogether. Could another filter applied in series further reduce the amplitude of the spikes? To check this, another filter was made, to see if it would have any effect on the measurements. Below is a table of measurements done between VME ground and the filter on the analogue power line:

Measurement	**Max peak**
VCC-VME ground on LV3 output	918 mV
VCC-VME ground after first filtering	815 mV
VCC-VME ground after second filtering	749 mV

Table 6.4: Measurements done on two filters connected in serial. We see that the second filter reduces the peak value further.

The measurements of noise occupancy and input noise does however not show any drastic improvement, see figure 6.13. We see that the measurements done with two filters are slightly lower than the ones done with one, but this might as well come from the fact that the 2-filter measurements were done after the one filter measurement, with succeeding scans often resulting in slightly lower values for the later scans.

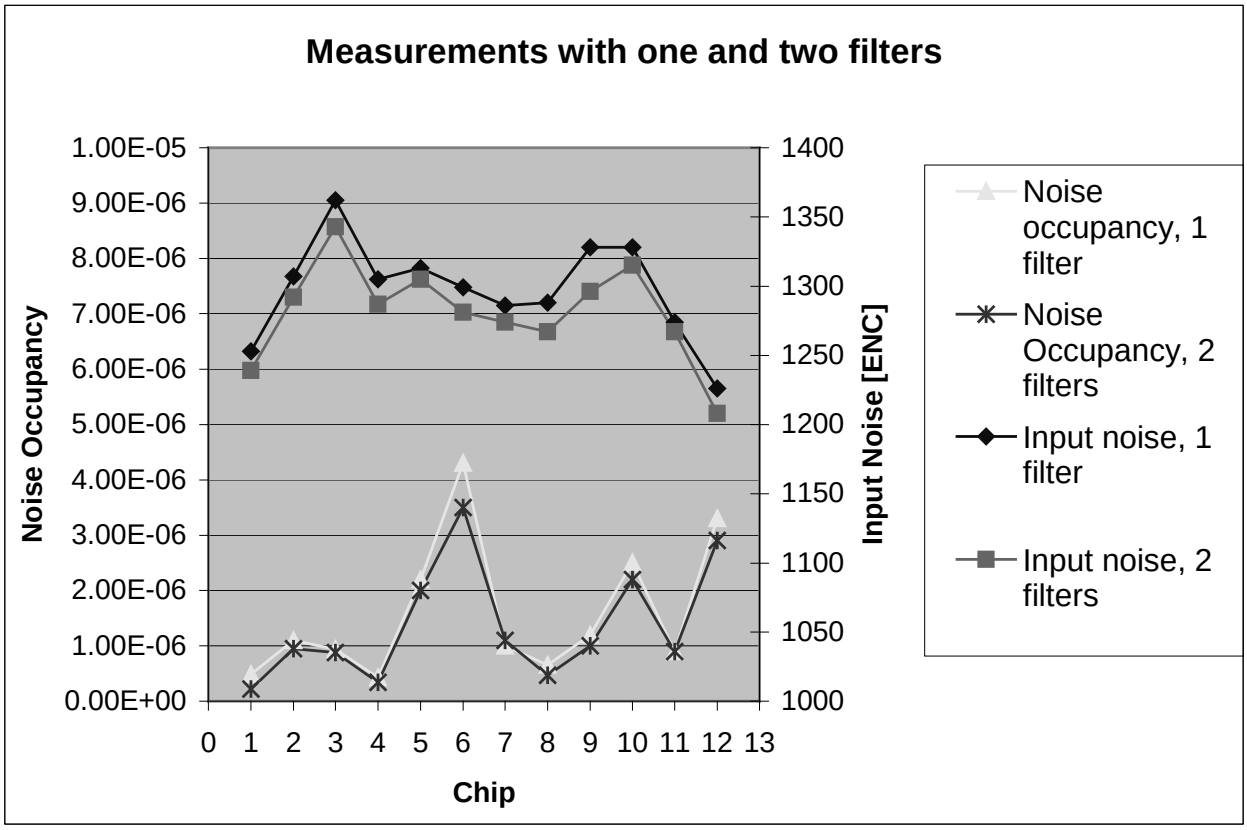

Figure 6.13: Measurements of noise occupancy and noise done with one and two LV3 filters.

It seems like the first filter made the spikes from the power supply sufficiently small, so that they do not influence the measurements. Further reducing their amplitude does not seem to improve noticeably on the already good results.

6.1.3 LV3 filter conclusions

We can draw some conclusions from the experience gained when using the filter.

- Applying a filter to the LV3 improves Noise Occupancy and ENC measured from Noise Occupancy scans considerably.

- Filter has small effect on input noise measured by Response Curve test.
- ENC estimates from noise occupancy scans with filter are in agreement with ENC measured by response curve, while data from unfiltered scans show large differences.
- Temperature dependence of input noise as measured by Response Curve and filtered Noise Occupancy measurements is of order ~ 5 - 10 ENC/°C.
- Temperature dependence of noise occupancy is of order ~ 1 x 10^{-7}/°C.
- Applying a second filter to the SCT-LV3 does not improve noticeably on the results.

6.2 Noise from environmental chamber

As described in section 4.1, the SCT module is placed in an environmental chamber for cooling during the tests. The chamber has been a serious noise source, and the work related to the removal of this noise is described in the following section.

6.2.1 Loss of communications due to relay switching

Initially, the chamber used electro-mechanical relays for controlling the cooling and heating. The switching of these relays generated sufficient electrical noise to 'blackout' the SCT module, resulting in loss of communications and test data during the runs. This was a serious problem with the reliability of the setup, which should be able to perform continuous 24 hour testing.

On the next page, on figure 6.13, the circuit diagram for the chamber control system is shown, with the noisy relays indicated.

By contacting the manufacturer, we received instructions on how to change from electro-mechanical relays to more 'silent' zero-crossing solid state relays, denoted by SSR in the following. These relays have no mechanical parts, and are supposed to do the switching only at the zero crossing of the AC current, thus generating a minimum of noise.

Figure 6.14 shows the changes done to the original setup, showing the old relay disconnected and the new connections to the SSRs.

The new relays were of type RS 352-519 with 85-280 V AC control, SSR 10A 250V ac load, 600V block, and the resistors of type MEGGITT SBCHE6, CGS 10K J DMHG

The three switched lines of RL1 on figure 6.13 were replaced by only two by the SSRs, so there were some concern in disabling the lines 27/28, thus disabling the FASCIA Heater. This seems to have something to do with heating of the tubes in the cooling system, preventing these from freezing if run at low temperatures over long periods of time, which will not be done in this setup. Disabling this functionality did not seem to have any direct influence on the operation of the chamber.

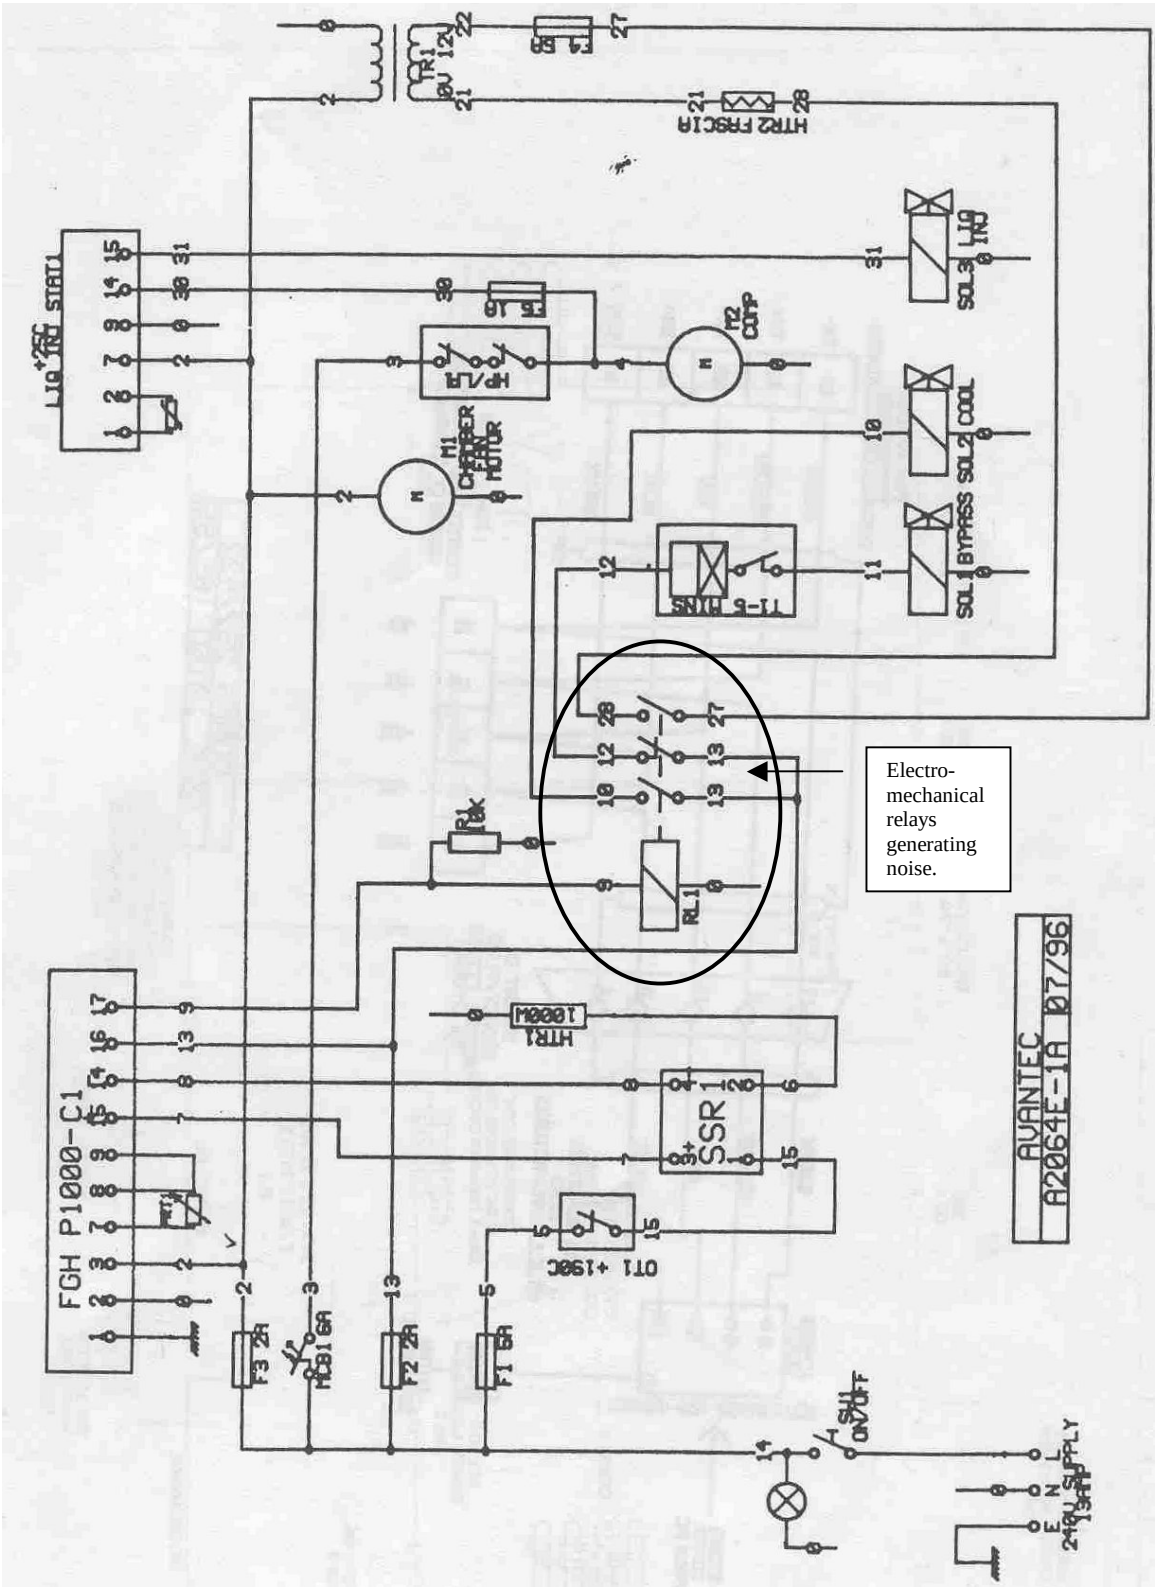

Figure 6.13: Original control circuit of cooling chamber. [25]

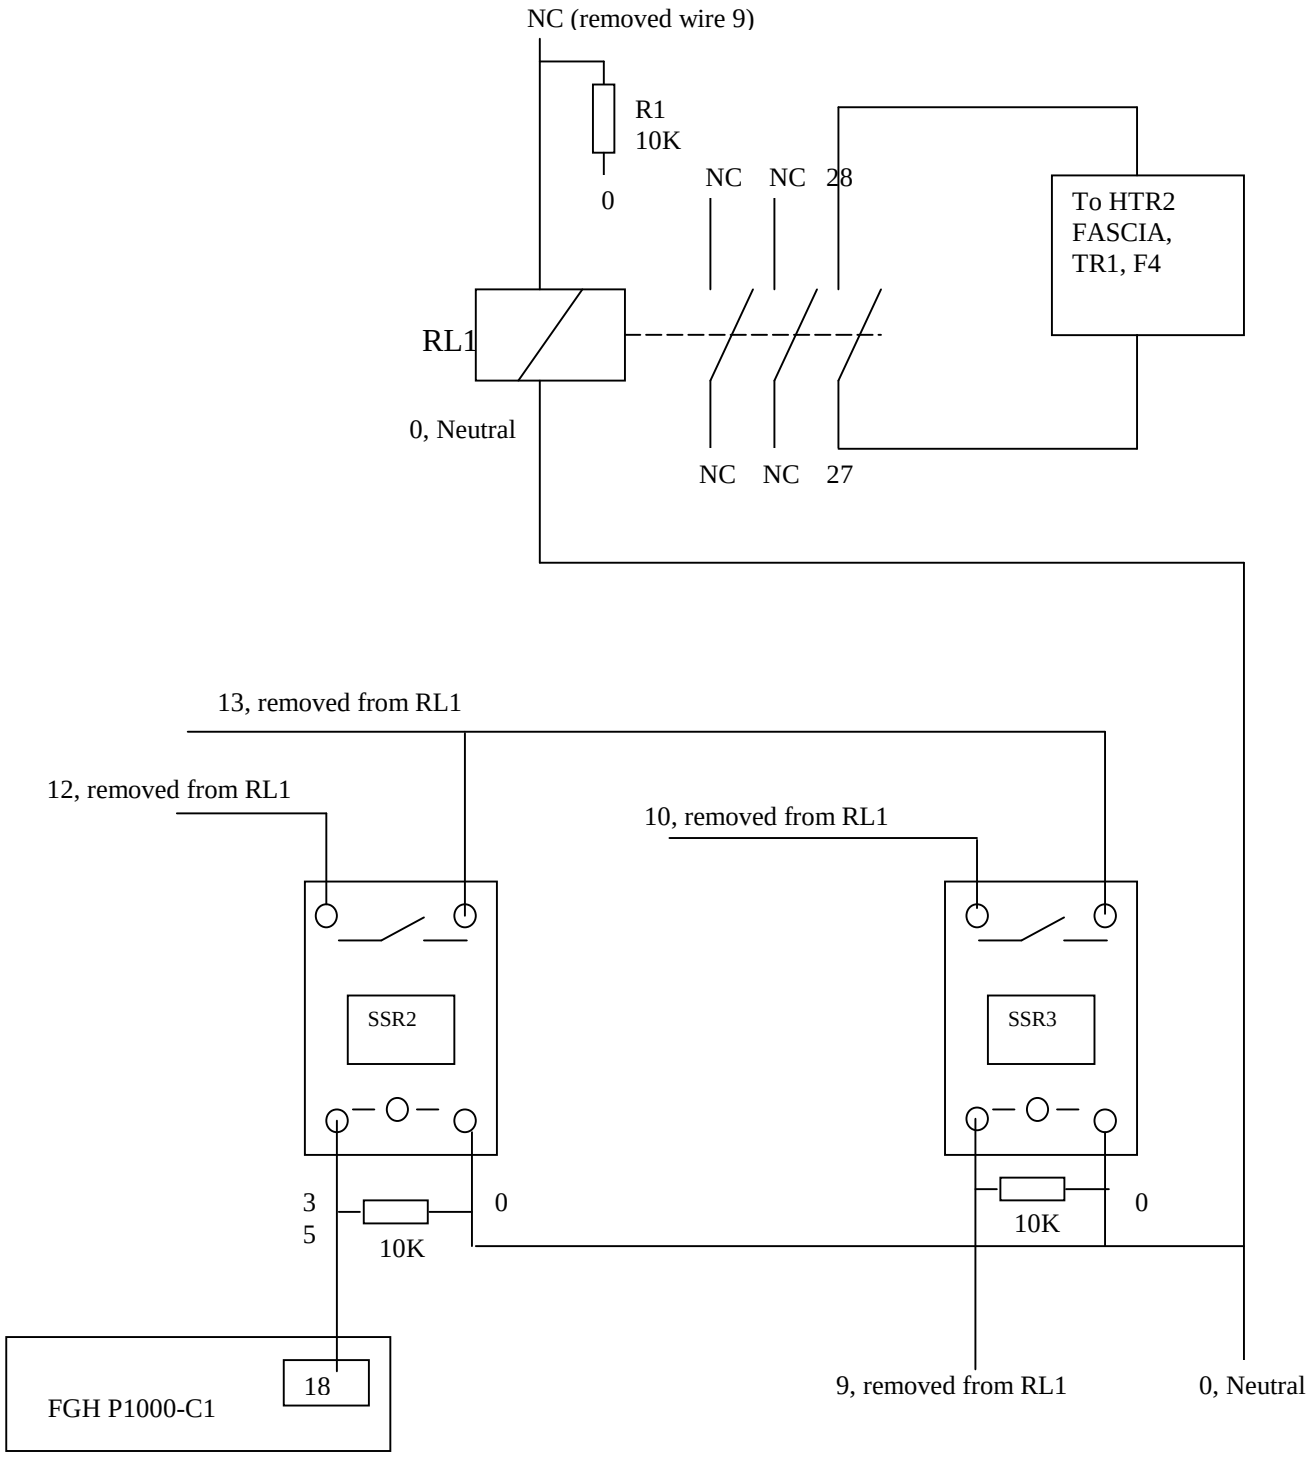

Figure 6.14: Modifications done when installing solid state relays in cooling chamber.

The change to SSRs lowered the switching noise enough to solve the 'blackout' problem, preventing the critical communication losses, but did not remove the switching noise entirely.

6.2.2 Noise occupancy spike problem

Instead of loosing communications, there often showed up a 'spike' in the noise occupancy plots due to the switching of relays, see figure 6.15.

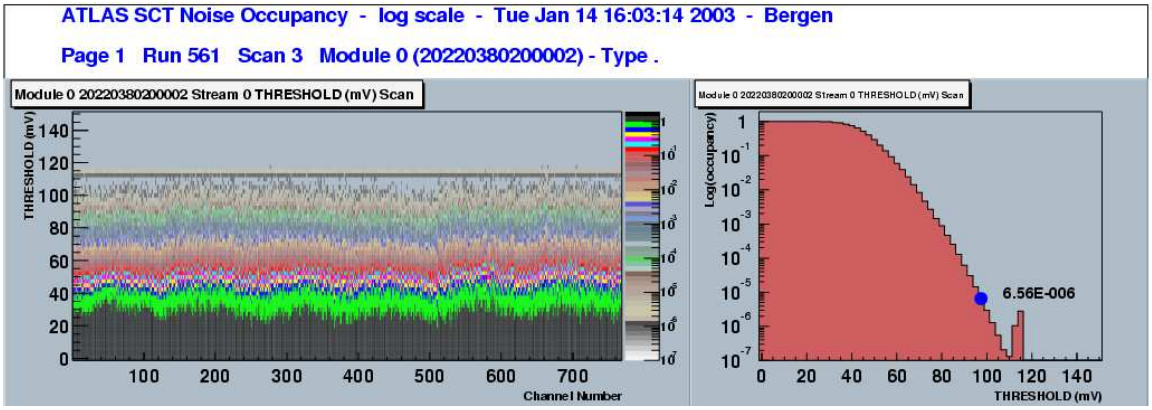

Figure 6.15: 'Spikes' on noise occupancy plots due to relay switching in cooling chamber. We clearly see the spike, also in the scan histogram, as a uniform line of hits in all channels.

In the plot of figure 6.15, there are actually two spikes, one of occupancy 1 x 10^{-6} and one of 3 x 10^{-6}. Each scan point in the upper threshold range of the scan consists of 10^6 triggers, taking roughly half a minute to send. For the first spike, there is approximately 1 hit for each channel (1 hit/10^6 triggers = Occupancy of 1x10^{-6}), meaning the chamber has done one "noisy switching" during the half minute interval. The next spike should likewise mean three switchings in the interval.

The rate at which these spikes appeared varied strongly, but seemed independent of temperature and filters. They would show up in approximately one of three noise occupancy plots, but did not have any direct effect on the noise occupancy measurement unless they hit the 1.0 fC point precisely. They did, however have a certain influence on the linear fit/ENC estimation, see figure 6.16. Since the fit tries also to fit the points from the spikes to the line, we get a somewhat higher ENC than is the case.

To locate the source of the noise, measurements were done with an oscilloscope at various points in the system. Transients were measured both on the VME modules and in the chamber, but of different magnitude. The largest amplitudes were found inside the chamber, ranging from 3 up to 16 V. The corresponding transients measured on for example the MA output of the MuSTARD were of the order 100 mV.

It was observed that the largest spikes occurred when the cooling was switched on and off by the cooling chamber controller. To do this, the system opens the valve with the solenoid marked SOL2 COOL in figure 6.13. In the modified chamber, this solenoid is controlled by SSR3. We disconnected the solenoid from SSR3 (but with SSR3 still operational), to see if this had any effect on the noise measured on the inside of the chamber, but the transients did not disappear.

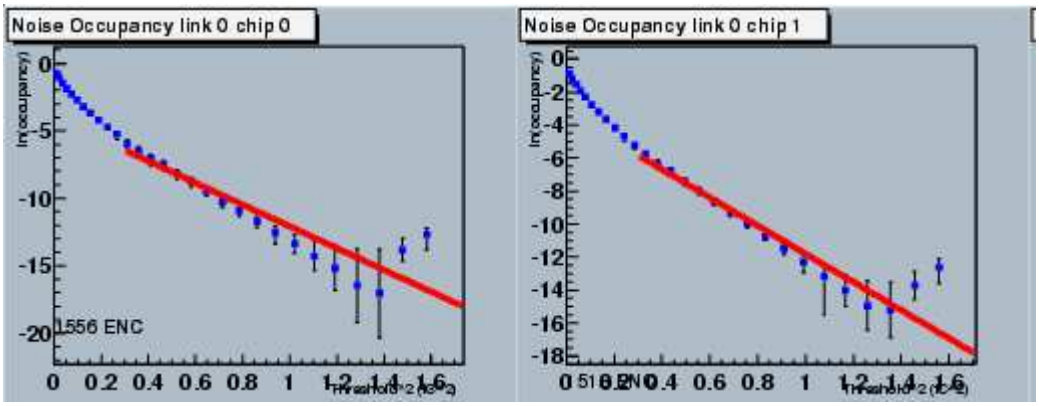

Figure 6.16: Deviation from linear fits due to switching noise (from the same scan as figure 6.15).

SSR3 is controlled from line 9 of the P1000 controller. This line is in turn controlled from a relay inside the controller, so to check if this generated any noise, all cablings from the controller unit were removed. The switching of the relay inside the controller did not create any transients inside the chamber, so it could be eliminated.

The large transients therefore seemed to come from the SSR itself or from cabling and other circuitry, possibly because of fast switching generating rapid voltage drops. In an attempt to dampen the transients, 0.1 µF capacitors were placed in parallel with the SSRs. Some of the transients could then be used for charging the capacitor, but this did not seem to have any effect on the problem.

Since the problem could come from the fast switching of the SSRs, a slow switching solid state relay was built. The final design of the relay is shown on figure 6.17.

The requirements for replacing SSR3, was to keep wire 9 as the control signal, and turn SOL2 COOL on when 9 was active. Wire 9 comes from the controller unit of the chamber, see figure 6.13. It should close wire 13 and 10 when it is active, so that a current flows through the solenoid, generating a magnetic field which exerts a pull on the iron bolt connected to the cooling valve. This should start the cooling of the chamber.

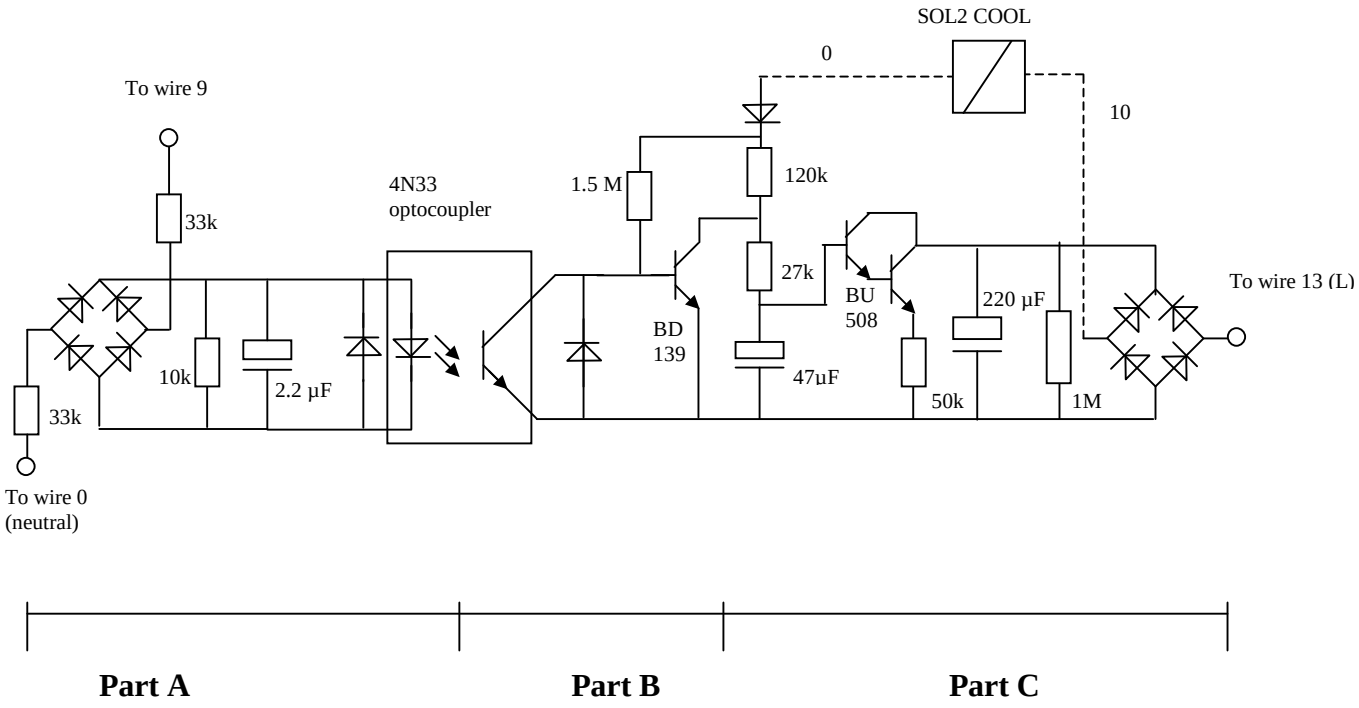

Figure 6.17: Slow-switching solid state-relay.

The relay can be divided into three main parts, as shown on the figure. A description of each part and its principle of operation follows:

Part A:

Wire 9, which is phase 1 from the net, and 0, which is phase 2 from the net (with 180° relative phase), is input to a rectifier. If wire 9 is inactive, we see that there can go no current through part A, because the current from 0 has nowhere to go. If wire 9 is active, current will flow through the optocoupler. The 2.2 µF capacitor is used to make the rectified current smooth.

The 4N33 optocoupler was introduced in the circuit as a result of voltages in the circuit influencing the level of the input signal. This caused troubles with switching at the correct voltage levels. When using an optocoupler, the diode and transistor are galvanically separated, eliminating problems related to voltage differences between the two parts of the circuit.

The diode of the 4N33 works as a base for the transistor, making C-E conductive when it is active.

Part B/C:

If wire 9 in Part A is inactive, no current flows through the diode of the optocoupler, creating no conductivity between collector and emitter of the transistor in 4N33. Then the voltage difference between B and E of the transistor BD139 becomes sufficiently large (>0.7 V) to make its C-E connection conductive. This means that there can be no voltage across the 27 k and 47 µF components, which in turn means that there is no voltage between B and E of the transistors BU508. All current must then flow through part B, and out through wire 13 through the rectifier to the right on the figure.

If wire 9 is active, the C-E of the optocoupler is conducting, and there can is no voltage difference between B and E of the BD139 transistor, resulting in no current through the transistor. Very little current will go through the 1.5 M resistor, so most of the current will flow through the two 120 k and 27 k resistors. In the start, most of the current will go to the 47 µF capacitor, but as this charges, voltage also builds up at the BU508 transistor base. The current from wire 13 is then free to run through the rectifier, the BU508 transistors, and back through the rectifier to the SOL2 COOL solenoid, which should open the valve and start the cooling.

The 47 µF capacitor plays an essential role in the circuit, as it is the component that creates the slow switching effect. As wire 9 becomes active, instead of putting current instantaneously on wire 10, as was done in SSR3, the current is increased gradually while the 47 µF capacitor charges. This also causes the magnetic field in the solenoid to rise gradually, opening the valve at a slower rate, and thus creating a less rapid voltage drop. The capacitor has a time constant of 47µF * (120 k +27 k) = 6.9 s, but in practice, the switching was done in ~1 s, because the BU508 transistors starts to conduct before the 47 µF capacitor is fully charged.

Due to the diode placed before the 120k resistance, only half of the 0 phase will flow through, which could result in fast on/off switching of the BU508 transistors. This is prevented from happening because of the 47µF capacitor charging out in each such 10 ms interval.

The 220µF capacitor is placed in the circuit for smoothing out the voltage after it has passed the rectifier. Across it, a 1M resistor is placed, to let the capacitor discharge when not powered. This is to protect against electrical shock through touching.

The coupling of the two BU508 transistors is called a Darlington coupling, giving a amplification equivalent to the product of amplification in each separate transistor.

Although the relay did work, it did not solve the noise occupancy problem, and the spikes continued to show up in the plots, at roughly the same rate. Trying to disconnect the relays one by one, it was seen that although the other relays created smaller transients, these were also sometimes, although not as often, sufficiently large to give a spike in the

NO plot. This did not happen for every switching, but seemed to depend on the amplitude of the transient, as expected.

We could have tried to further rebuild the control circuits of the chamber, but this would require quite a lot of work, so the strategy was changed towards trying to isolate the system from the noise generated by the chamber.

The transient could sneak into the system either through the power net, or through the air as electromagnetic radiation, eventually as both. Both cases had to be taken into account.

Various strategies were tried:

- **Noise filter on the power inputs,** supposed to remove transients from the 220 V power, were applied to both the VME crate and the Keithley high voltage source It could not be applied to the chamber since the filter had a 5 A fuse, while the chamber draws up to 13 A. This did not remove the spikes.
- **Different groundings schemes of the VME crate and cooling chamber** were tried (three different courses in various combinations). The inside of the chamber were also disconnected from chamber ground and grounded to the VME crate. This seemed to have no effect on the problem. All available grounds are however connected at some point in the power net in the building, so two truly separate grounds were not available.
- **Shielding of module.** The module was placed in a metal box, electrically isolated from the box, which again was electrically isolated from the chamber, see figure 6.19. The metal box was then grounded outside the chamber. This was done in an attempt to make the metal box pick up possible electromagnetic radiation from the inside of the chamber, but the spikes were still visible in the plots.
- **Shielding of cables.** The two flat cables providing power and communications were covered with two layers of conducting metal tape. The tape was electrically isolated from the cooling chamber, and grounded outside the chamber. This alone did not solve the problem, but by also shielding the module with the metal box as described in the previous point, and grounding the cable shield to the metal box, the spike problem seemed to vanish. The configuration is shown in figure 6.18. During previous testing, the maximum number of consecutive noise occupancy plots without spikes were 6. After properly shielding the setup, the spikes have not shown up again, despite extensive testing (100+ noise occupancy scans).

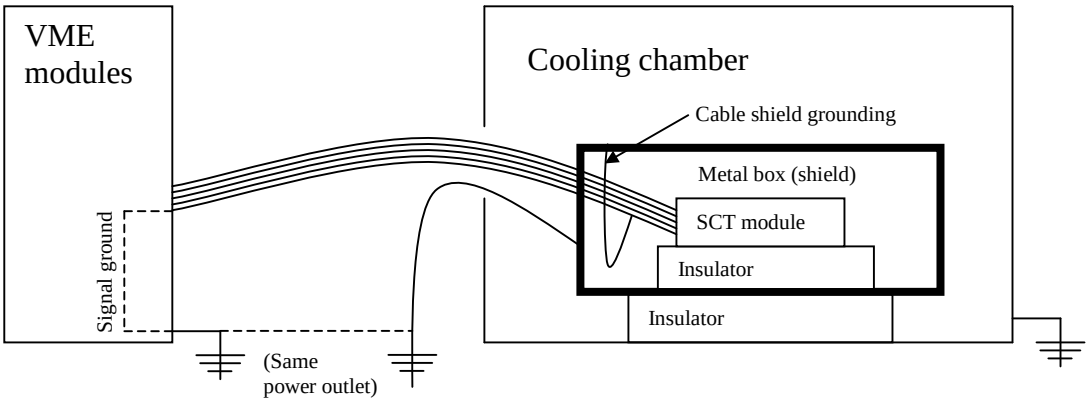

Figure 6.18: Final shielding of the setup.

We conclude that the spikes on the Noise Occupancy plots came from disturbances due to electromagnetic radiation picked up by the cables and SCT module electronics.

6.3 Grounding of system

To get a proper grounding of a system is always important. It is important to avoid ground loops, i.e. grounding of the system at different points in the setup, which can cause a current to flow between two ground points, generating noise.

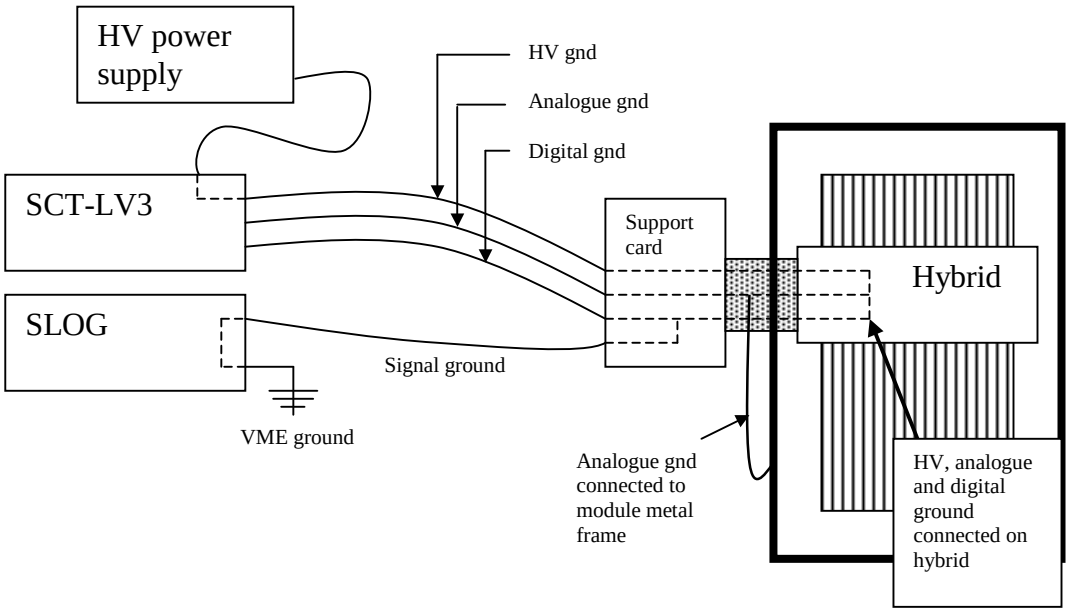

Figure 6.19: Final grounding of setup.

The grounding of the components in the setup is shown in figure 6.18 and 6.19. The entire readout system is grounded through SLOG, which defines signal ground, the level which digital communication signals are related to, from VME ground. The power supplies, both LV and HV, are floating, meaning that the grounds of bias, analog and digital voltage are separated from the ground. They are, however connected to signal ground on the support card and the hybrid, as shown on figure 6.19. The grounds are also connected to the module frame, to pick up eventual induced currents.

As can be seen on figure 6.18, the system ground (VME ground) and cooling chamber ground is taken from two different power outlets. This to prevent the already mentioned transients from influencing the system through the ground. When the cooling chamber is connected to the same outlet as VME, the transients are clearly visible through the output of VME modules on an oscilloscope. The effect is greatly reduced when using other outlets, but it has not been seen to influence measurements.

Different grounding schemes were tried; grounding digital ground on support card to power outlet ground, using separate cable from module frame to AGND on LV3 filter, disconnecting ground from inside of environmental chamber, keeping the module frame floating and so on. The setup in figure 6.19 has at the time of writing proved the most efficient and simple grounding scheme.

6.4 Noise elimination conclusion

The two main sources of noise in the system have been eliminated; the SCT-LV3 power supply and the cooling chamber. Grounding has also been investigated, but as long as the system is grounded only through the SLOG module in the VME crate, the various power outlets used was found to play little role in the overall setup.

By reducing the noise, the system is better suited to do precise characterizations of the SCT modules.

Chapter 7: QA procedures for the SCT barrel module

Before an SCT module can be used in ATLAS, it needs to go through a QA (Quality Assurance) procedure. In addition, each part of the module (detectors, hybrids, baseboards and readout electronics) undergoes its own QA before being assembled. This chapter describes the steps of the module QA, but first gives an overview of the QA for each module part, also specified in the document [26].

7.1 QA of module parts

QA of ASICS

The ASICS have undergone their own QA before being mounted on the hybrid. The procedure consists of both analogue and digital tests of the ASICS, to ensure that both the digital and analogue readout works properly. The power consumption and currents in the ASICs are also tested to be within spec.

QA of hybrid

The hybrid undergoes a QA before it is mounted and wirebonded to the rest of the detector sandwich. The QA consists of a test of the passive components of the hybrid, without ASICS mounted. Visual checks, and a verification that the hybrids have the correct dimensions in all planes are also performed. The ASICS are then mounted on the hybrid with electrically conductive epoxy. The hybrid is tested by running a Characterisation Sequence (see chapter 5).

After this, a long term test is performed. This is a 100-hour test run with the hybrid thermistors at 37°C, with the ASICs temperature expected to be 50°C. During the test the hybrids are powered, clocked, configured and triggered at the nominal L1 trigger frequency of 100 kHz. The currents and temperatures are measured every few minutes. The ASICS are then bonded to the pitchadapter, and a conformation sequence followed by a visual check is finally performed to ensure that the hybrid operates as expected.

QA of detectors

Each detector is checked visually for dust or cracks or other damages. An IV-curve (See section 7.2.2.) up to 500 V is recorded for each detector, allowing a maximum leakage current of 6 µA at 150 V and 20 µA at 350 V.

QA of Sensor-Baseboard Sandwich

The sensor baseboard sandwich consists of four silicon strip detectors. The assembly starts with the detectors being glued glued to the VHCPG (*Very High Thermal Conductivity Pyrolytic Graphite*) baseboard. The assembly is checked visually. An IV curve is recorded up to 500 V for each detector individually, and if the 500 V bias differ by more that 1 μA from that last recorded in the database for the detector, the assembly is put aside for further visual checks and currents stability measurements.

A metrology[12] survey is also performed. This measures the relative positions of the four sensors, the stereoangle and the holes in the baseboard. If the sandwich passes this test, a hybrid is mounted on the sandwich. Leakage current is checked for up to 500 V, and the strips are wire-bonded to the pitch-adaptor of the hybrid.

7.2 QA of completed module

The various stages of testing for a completed module can be summarized as follows:

- Metrology (not part of this setup)
- IV-curve and visual inspection
- Initial electrical test at room temperature
- Thermal cycling
- Cold long-term electrical-/leakage current stability test

In the following sections, each of the testing stages will be described.

7.2.1 Metrology

The completed module is surveyed for mechanical precision. A well-defined set of fiducial marks on the sensors is used (in-plane survey) as well as a 5x5 matrix for the hybrids and mounting (out-of plane survey).

7.2.2 IV-curve/visual inspection

Upon receival, the module is checked visually for damage. It is important to look for cracks in the baseboard, scratches or other damages to the detectors and damages to the wire-bondings.

The module is then placed in the environmental chamber, connected to the setup, and the chamber temperature is set to 15°C (this is the temperature at which the electrical test in room temperature is to take place).

[12] Metrology is a precise measurement of the physical dimensions of a module.

Since the leakage current is temperature dependent, it is important to wait for the temperature of the chamber to stabilize. The ASICS are left unpowered, and the bias is incremented from 0 to 500 V in 5 V steps, waiting 20 s between each step to give the leakage current time to stabilize. This is plotted as an *IV*-curve, leakage current I vs. applied bias V, as shown in figure 7.1. All measurements should be normalised to 20 °C for comparison with earlier measurements of the detectors. This is done by the formula:

$$I(20[C]) = I(T[C]) \cdot \left(\frac{293}{273 + T[C]} \right)^2 \cdot e^{\left(-7019 * \left(\frac{1}{293} - \frac{1}{273 + T[C]} \right) \right)} \qquad (12)$$

The measured leakage current should not exceed the sum of the leakage current recorded in the database for each of the four detectors. An example of an IV-curve for the module 2022038020002 measured in this setup is shown in figure 7.1, with both the measured curve at +15° C and the corresponding normalized curve at +20 °C.

If the leakage current suddenly shoots up at a certain voltage (a phenomenon known as *micro discharge*) the module should be subject to long term (24 hour) leakage current stability testing at 350 V. If the current decays to a normal current within the 24 hours and stays stable for at least 3 hours, the test is passed.

The IV- curve in figure 7.1 was taken with a Keithley 487 Picoamperemeter/Voltage source, the HV source of the current QA setup. In the final setup, the IV curves should be taken by the new HV source under development, which will be better integrated in the setup (see chapter 3).

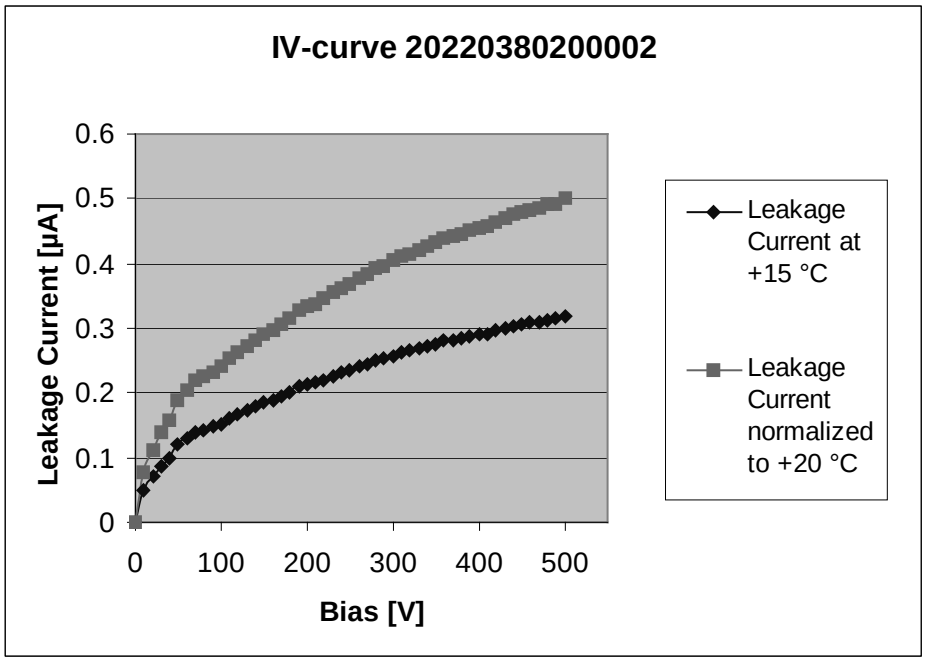

Figure 7.1: IV-curve of module 20220380200002.

7.2.3 Electrical testing

The module is now ready for the first electrical test. This is performed in room temperature, defined to be 27 ± 3 °C measured by the hybrid. This temperature is achieved by setting the cooling chamber to 15.0 °C, already done under the IV-curve scan. If the humidity gets to high, this can cause too high leakage current and eventually short circuit of the electronics, so the humidity should stay below 60% to have a safety margin. The leakage current has been seen to shoot up at humidities near 90 %.

The detector bias is set to 200 V, and a characterisation sequence is started from the *ABCD tests* menu in SCTDAQ.

7.2.4 Thermal cycling

The module is, unpowered, brought through ten thermal cycles from –25°C to +40°C. The ramp up/down time is 30 minutes, and the soak time (time before starting a new ramp up/down), is also 30 minutes, with the first cycle going up to +50°C. This results in a total testing time of approx. 20 hours. This is a purely mechanical test to see if the module can handle large and relatively rapid changes in temperature (2-3°C/min.). The cycling is controlled by the programmer unit of the cooling chamber.

To prevent dewing during the cycles, the chamber is required to go through two volume changes of nitrogen during the 20 hours. The inside of the chamber is approximately 150 l, so this requires a flow of 15 l N_2/h. To achieve this, the flow controller is set to 100 mm, with a pressure of 3 bar. The chamber is also flushed with nitrogen before starting the cycling to ensure an initial dry atmosphere. During the cycling, the humidity is monitored by a hygrometer.

7.2.5 Cold long-term stability test

This test is performed to ensure that the module shows long-term stability at the ATLAS operating temperature, 0°C. The ASICs are powered, clocked and triggered, and the detector bias is maintained at 200 V, with the current and temperature monitored every 15 minutes over the period. The cooling is set to –13.5 °C so that the hybrid thermistors measures 0 ± 3 °C. To check the electrical properties of the module over the long term test, a confirmation sequence is performed every second hour. Every five minutes, a noise occupancy is recorded with a burst of 32 000 triggers with threshold at 1.0 fC. At the end of the test, a Characterisation sequence is performed.

A Long Term Test plot is shown in figure 7.2. In the plot, the last NO point (After the Characterisation sequence) is noticeably higher than the rest. This is due to a bug in SCTDAQ, and does not show up in newer versions of the software. We see in the plot temperature (upper row), Icc (analogue current) and Idd (digital current) in the second row, detector bias in the third row (always zero due to our HV supply not being integrated in the setup), and the NO measured by the 32 000 triggers in the lower row.

The test is started by choosing Module LTT (Long Term Test) from the *ABCD test* menu in SCTDAQ.

7.2.6 Data examination

After running the tests, the data from all tests must be examined. For most of the tests there exists a set of criteria for a module to PASS or FAIL (see section 5.4.5), but some of the tests are still under revision.

7.2.7 SCT Database/Web publishing

When a complete electrical QA is performed on a module, the results should be uploaded to the SCT database. This will be done with a separate Java interface, which still is under development. There is however a set of Perl macros that makes it easy to publish the results from a module test on the web. This is run by the command:

```
Perl ArchiveAndPublish.pl <modulename> <label>
```

from the command prompt in Windows, where label is a comment given by the user to identify the test. The macro takes all tests found in the results file for the specified module and makes a hierarchy of webpages with plots and data for each module. The results can be found on the address [27].

7.3 Estimated time for QA

In table 7.1, an estimate over the time needed to do a complete QA for a set of modules is presented. With the current setup, only one module can be tested at a time, but the plan is to test six modules in parallel once the required equipment is available.

Test		**Duration**	**Note**
IV-curve scan		35 min	
Eventual long term stability leakage current test		24 h	This test is performed if there is a *micro discharge* during the IV-curve scan.
Electrical test at room temperature		1.5 t	Characterisation sequence. Time estimated is for testing one module, might take several hours for a set of six modules
Thermal cycling		~20 t.	
Cold long-term electrical/leakage current test		24 t	
Handling, cooling, and other tasks		1-2 t	
Total	No micro discharge	~50 h	
	Micro discharge	~80 h	

Table 7.2: Time estimation for module QA.

Test	**Time**
Hard reset	1 min
Set stream delay	1 min
Full bypass test	2 min
Redundancy test	1 min
Pipeline test	1 min
Set strobe delay	3 min
Three Point Gain	4 min
Trim range	31 min
Response curve	17 min
Noise occupancy (dependent of number of triggers that needs to be sent)	15 min
Timewalk	17 min
Total	**~90 min**

Table 7.3: Characterisation sequence duration. Times are for one module, expected to increase when several modules are tested in parallell.

We see that for a set of modules with no micro discharge, the test takes ~50 hours. The thermal cycling and LTT test runs overnight, and the minimum time taken for testing a set of modules is two days.

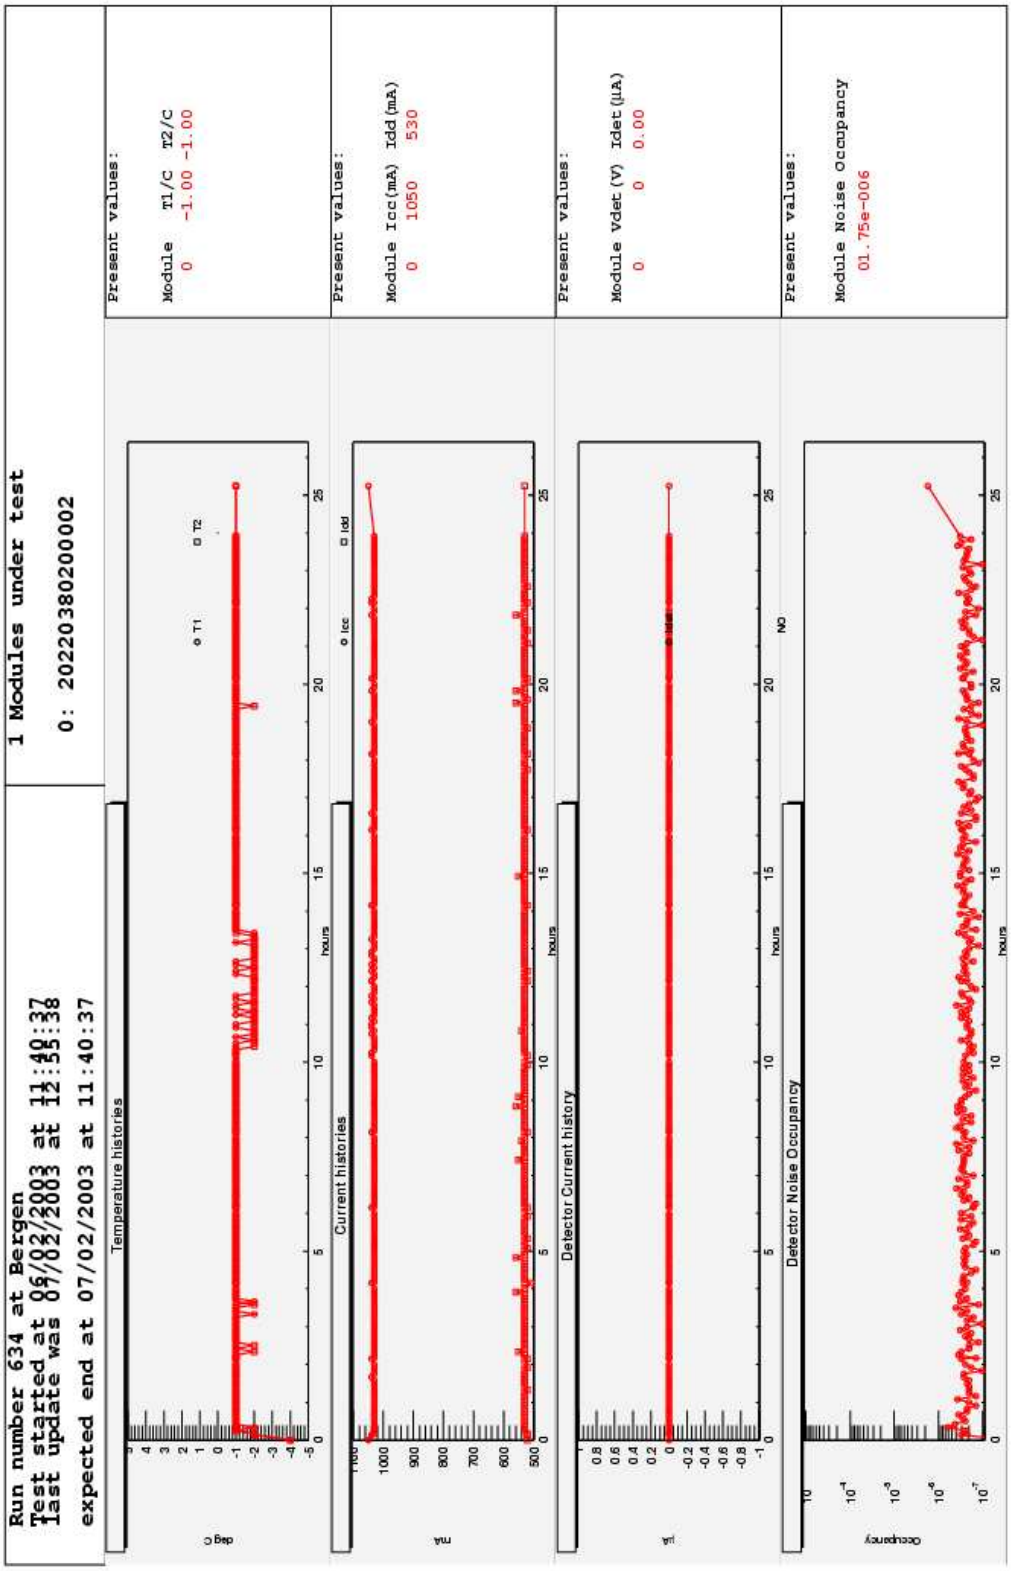

Figure 7.2: Cold long term test. The last high NO point is due to a bug in SCTDAQ.

Chapter 8: Oscillations in s-curves

As discussed in chapter 5, an s-curve is a plot of occupancy vs. threshold, and should appear as a smooth curve. However, on approximately 50 % of the modules produced in the Scandinavian cluster (detector-baseboard sandwich glued in Oslo, hybrid mounted and bonded in Uppsala), the s-curves have shown a behavior known as oscillations. An example of oscillating s-curves can be seen in figure 8.1. The effect only appears at low calibration charges (0.0-1.0 fC), for example when performing a noise occupancy scan (no input charge). The effect also only shows up at the bottom side (link 1) of the module, and usually increases towards the end of the hybrid. The effect was first noticed in Uppsala, as the module used as test module for the setup in Bergen, module 002, did not show any oscillations. To determine whether the oscillations were due to intrinsic problems with the modules or with the test setup, production modules 6-13 were shipped to Bergen for testing, to see if the oscillations seen in Uppsala also could be seen here.

The test specifications do not put any requirement on the shape of the s-curves, but it was nevertheless an effect that needed to be understood. In the following sections, the steps that have been taken to gain insight in this problem will be discussed.

8.1 Characterisations in Bergen and Uppsala

Initially, all the 8 modules were run through a characterisation sequence in Bergen, and the results from this can be seen in table 8.1 and figure 8.2. It was found that the s-curves showed more or less the same oscillating behavior at both sites, as shown in figure 8.1.

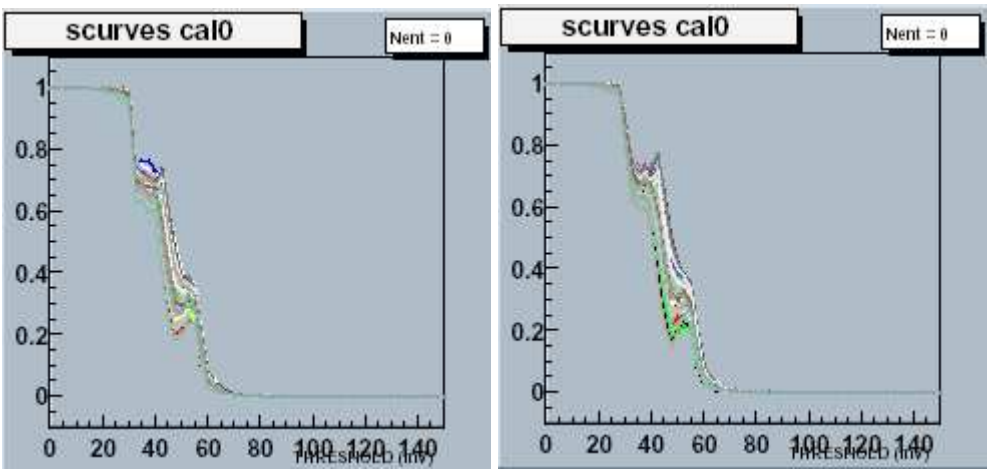

Figure 8.1: Oscillating s-curves from module 20220380200006 at 0 °C. The s-curve to the left is taken with the setup in Bergen, while the one to the right is from Uppsala. We see that the curves show the same oscillating behavior.

As can be seen from table 8.1 and figure 8.2, the results from the characterisations are in agreement, apart from the noise occupancy, which is somewhat higher in Uppsala, probably due to higher external noise. The characterisations also acted as a good cross-check for the two test setups, which have some fundamental differences. In Uppsala, the modules are connected to a large aluminium block which are cooled with cooling pipes running through the block. In Bergen, the modules are kept in their plastic boxes and cooled with air in the cooling chamber.

Since the measurements were in agreement, the question whether the two testing setups had a common problem had to be taken into account.

8.2 Oscillations with respect to different grounding schemes

The grounding schemes used in Bergen and Uppsala was quite similar, so to check if the grounding of the setup had any influence on the oscillations, four different grounding schemes were tried, also listed in table 8.1, named 1, 2, 3 and 4.

From a noise point of view, grounding 2 had already been found to give good values with respect to noise, see chapter 6. It was found, as can be seen from table 8.1, that the groundings 1, 2 and 4 gave approximately the same results, that is; oscillations on some modules, but good values from characterisations. Grounding 3, using a separate cable from the module frame to AGND on SCT-LV3, in some cases removed the oscillations altogether, but resulted in too high noise occupancy, >5e-4. The grounding obviously introduced heavy noise in the system, which probably acted as a ‘screen’ for the oscillations.

Summary of characterisations in Bergen and Uppsala

	Bergen					Uppsala					
Module	**Input Noise**	**Gain**	**NoiseOcc**	**T**	**Gnd***	**Input Noise**	**Gain**	**NoiseOcc**	**Gnd***	**T**	**Comments**
#02	1100-1219	58.8-68.1	3.9e-8 - 6.8e-7	0°C	2	1200-1302	63.2-67.1	5.8e-7 - 6.2e-6		0°C	No oscillations.
#06	1277-1455	52.7-58.0	3.3e-6 - 1.5e-5	0°C	1	1312-1478	53.1-58.3	3.9e-6 - 9.0e-5		0°C	Strong oscillations.
	1347-1510	55.8-61.2	4.1e-6 - 1.3e-5	0°C	2						Same osc as gnd 1.
	1288-1566	52.8-66.3	9.1e-6 - 5.8e-4	0°C	3						Small oscillations, but high noise occupancy
	1485-1648	50.3-54.3	9.7e-6 - 2.5e-5	27°C	1						Same oscillations as for 0°C
#07	1254-1476	51.3-56.9	1.5e-6 - 1.2e-5	0°C	1	1217-1426	51.4-57.0	--		0°C	Some oscillations
	1235-1461	54.2-60.4	1.8e-6 - 8.8e-6	0°C	2						
	1328-1543	49.5-53.4	1.3e-6 - 4.5e-6	27°C	2						
#08	1271-1433	49.4-57.3	8.1e-7 - 3.1e-6	0°C	1	1306-1407	49.7-57.3	1.4e-6 - 5.2e-5		0°C	Very small oscillations
	1328-1482	49.3-57.0	1.5e-6 - 4.4e-6	0°C	2						Same osc as gnd 1.
	1342-1517	49.9-57.9	2.1e-5 - 4.4e-4	0°C	3						No osc, high noiseocc.
#09	1380-1512	52.4-56.5	3.3e-5 - 9.9e-4	0°C	3						No osc, high noiseocc.
	1390-1487	52.5-56.3	1.5e-6 - 7.1e-6	0°C	2	1379-1486	53.1-57.1	--		0°C	Strong oscillations.
	1546-1665	50.3-52.9	3.5e-6 - 3.1e-5	27°C	2						Strong oscillations.
	--	--	3.4e-6 - 2.6e-5	27°C	4						Strong oscillations.
#10	1426-1511	48.3-53.5	5.2e-6 - 1.8e-5	0°C	1	1555-1661	46.7-52.3	--		20°C	Small oscillations.
#11	1380-1480	51.3-54.5	8.9e-6 - 2.8e-4	0°C	3	1440-1508	51.1-54.2	5.6e-6 - 2.2e-5	4	5°C	No oscillations.
								5.3e-6 - 2.4e-5	2	5°C	
#12	1248-1436	56.0-61.5	4.0e-7 - 1.1e-5	0°C	2	1499-1666	54.5-58.7	--		27°C	No oscillations.
#13	1367-1507	51.7-58.5	2.6e-6 - 2.0e-5	0°C	2	1564-1751	47.3-56.0	--		27°C	Medium oscillations

***Gnd**

1) Floating module frame
2) Frame connected to AGND on patchcard
3) Frame connected to AGND on LV3 via separate cable.
4) Frame connected to AGND both through patchcard and separate cable to LV3.

Table 8.1: Summary of characterisations in Bergen and Uppsala [27, 28].

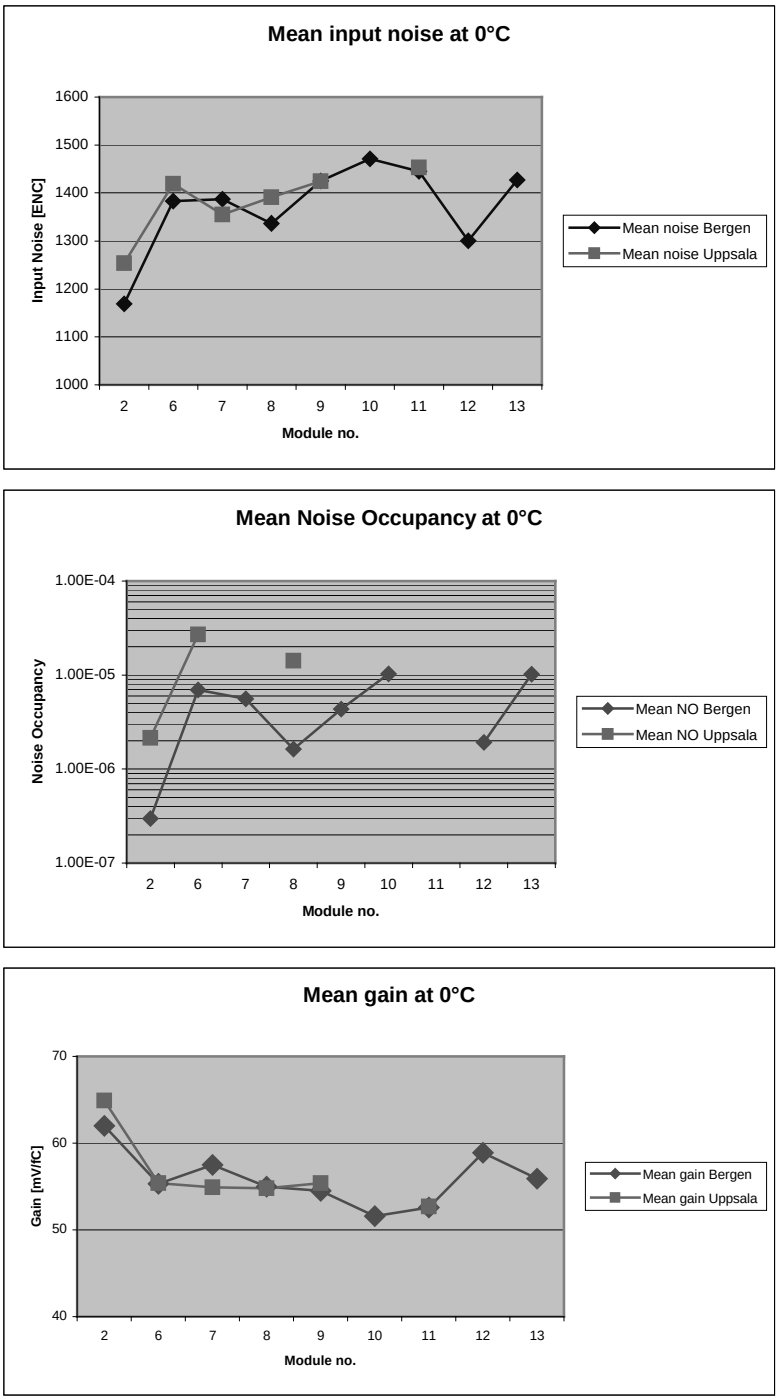

Figure 8.2: Mean values from characterisations in Bergen and Uppsala

8.3 Temperature dependence

To see if the oscillations depended on temperature, scans were done at the two specified operating temperatures, 0°C and 27 °C, and the oscillations compared channel to channel. Figure 8.3 shows s-curves for the last 32 channels on link 1, module 9 at the two

temperatures. There are minor differences, but the s-curves follow approximately the same oscillating pattern at both temperatures.

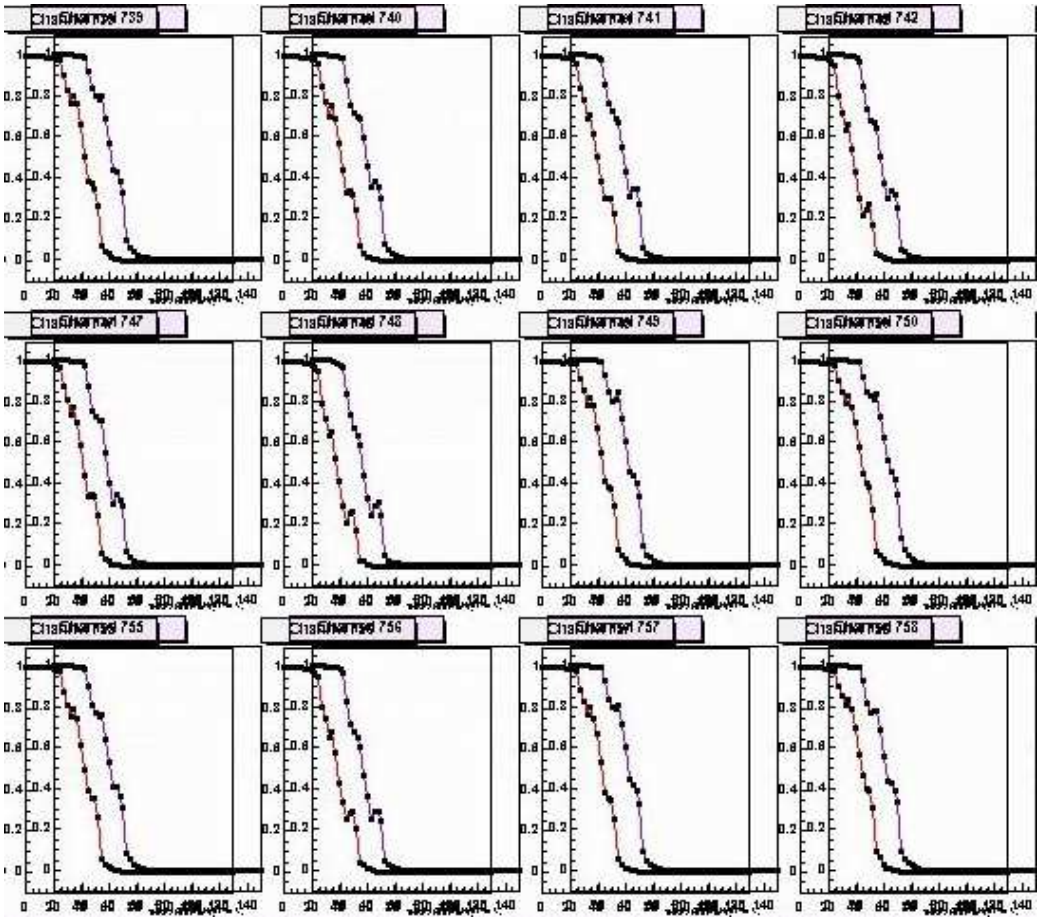

Figure 8.3: S-curves for 12 channels on link 1, module 20220380200009. The s-curves for a threshold scan at +27 °C (left curve) and 0°C (right curve) is shown, with a 20 mV offset for readability. We see that the oscillations are very similar at the two temperatures.

8.4 Oscillations versus input charge

As already mentioned, the oscillations only appear for small input charges. To do a more systematic study of this, 0-200 mV threshold scans were performed with input charges from 0.0 to 1.6 in 0.1 steps for three modules, one with large oscillations (006), small oscillations (008) and no oscillations (002).

$|\chi|^2$(Chi square)/NDF (Number of Degrees of Freedom) gives an indication of the goodness of the fit, and will thus be higher for larger oscillations. This can be seen on figure 8.4, which is a plot of $|\chi|^2$/NDF vs. readout channel for an oscillating module. We see that the values for $|\chi|^2$/NDF are much higher on stream 1 (bottom side of module), and increases towards the end of the hybrid. This agrees with the observed oscillations, the oscillations increases towards the end of the hybrid.

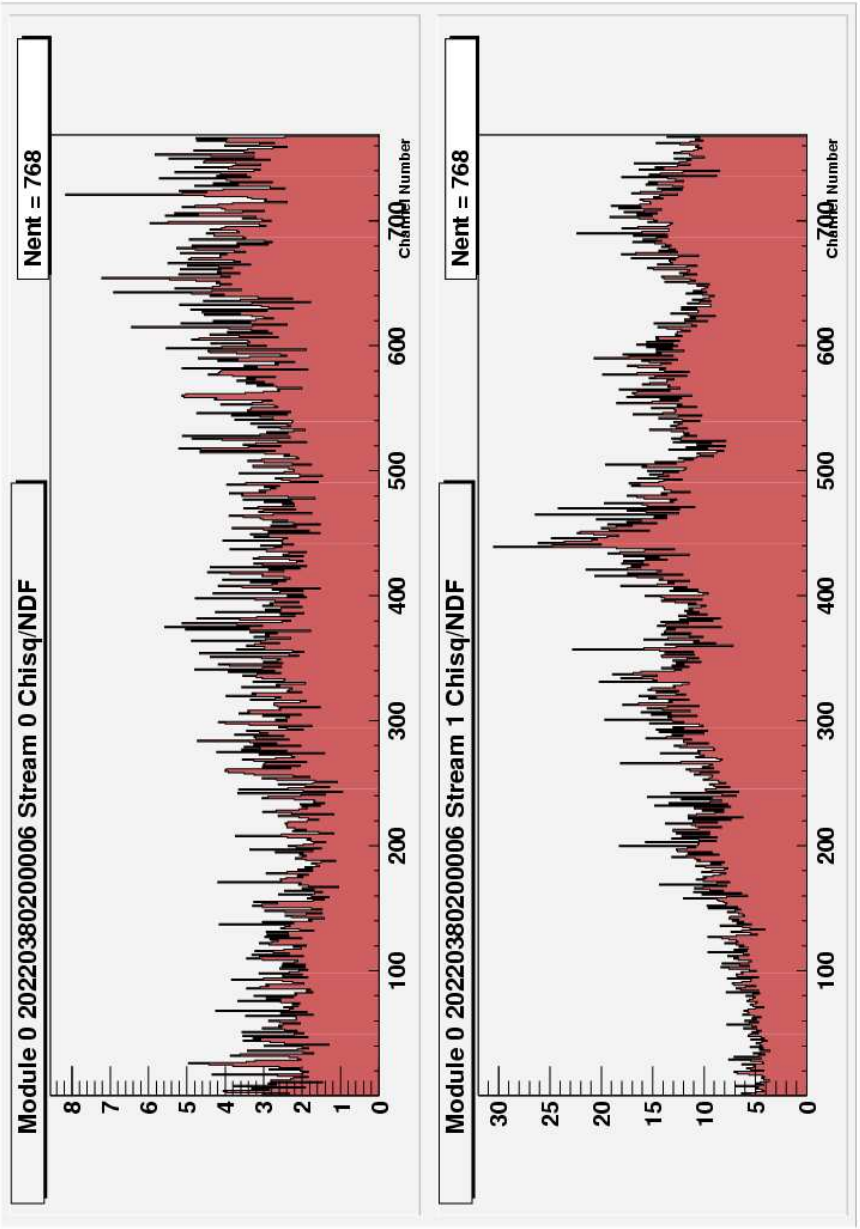

Figure 8.4: $|\chi|^2$/NDF vs. readout channel for 0.7 fC input charge. We see that $|\chi|^2$/NDF is higher for stream 1 and increases further towards the end of the chip. This corresponds to the behavior of the oscillations.

The mean values of $|\chi|^2$/NDF for the different input charges are shown in figure 8.5. We already now that the s-curves are not good approximations for input charges <1.0 fC, but it is however interesting to notice that the complementary error function gives worse fits at 0.4 fC than at 0.0 fC. Also notice that the values for link 0 are very consistent on the three modules.

For link 1, we see that the oscillations have a large effect at low thresholds, but that the $|\chi|^2$/NDF values join at a value of approximately 1 for input charge 1.0-1.3 fC for all

modules, which is the lower limit at which the modules are expected to operate at in ATLAS.

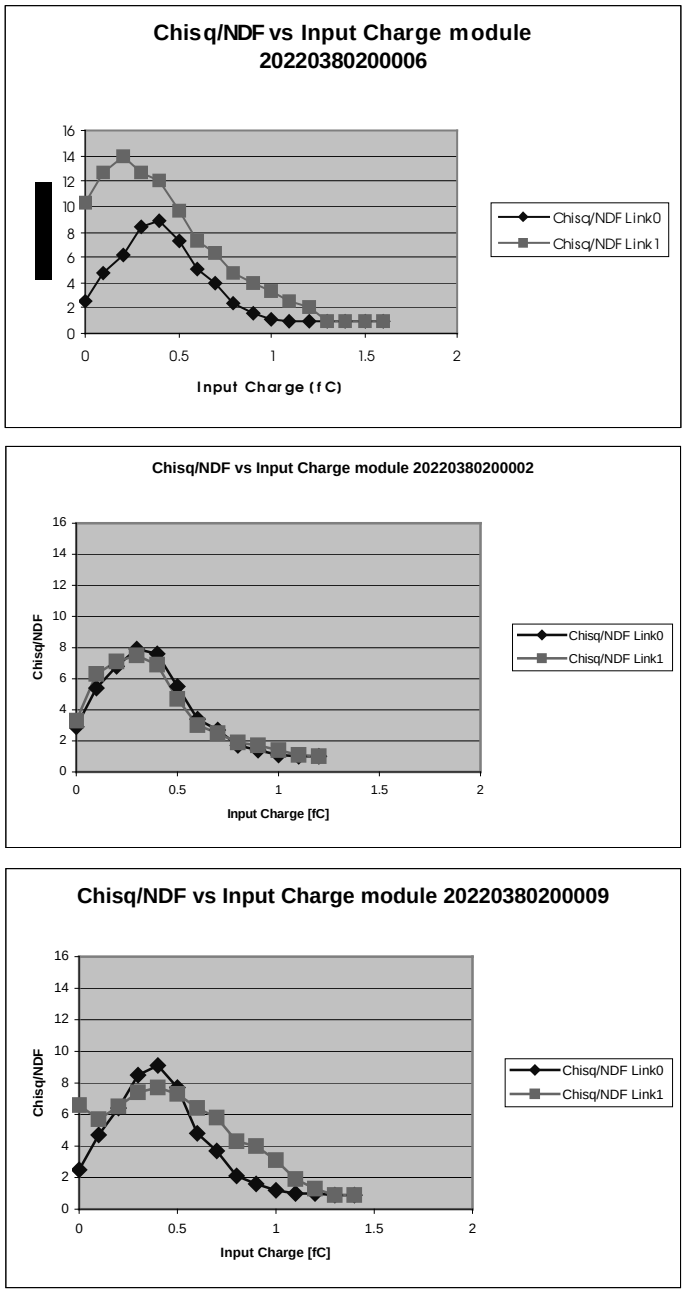

Figure 8.5: Oscillations represented by Chisq/NDF vs. input charge from threshold scans for three modules; 006 with large oscillations, 009 with medium oscillations, and 002 with no oscillations. Oscillations represented by Chisq/NDF vs. input charge from threshold scans for three modules; 006 with large oscillations, 009 with medium oscillations, and 002 with no oscillations.

8.5 Power

The tendency to oscillate seemed to increase towards the end of the hybrid. The end of the hybrid is also the end of the power line, and the phenomenon could thus be due to too low voltage at the end of the hybrid. The voltage, both analogue and digital, was therefore increased to see if this would have any effect on the oscillations. It was seen that when increasing the voltage, the oscillations also increased, and when reducing the voltage, the oscillations diminished. This was linked to the fact that when the power increases, so do the gain, and as the oscillations probably are a feedback effect, the oscillations would increase with the gain. The results are shown in table 8.2.

Scan/S-curves	Vdd	Vcc	Chisq/NDF L0	Chisq/NDF L1	Sigma L0	Sigma L1	Mean L0	Mean L1
Run 655 scan 1	4.0	3.5	3.1	6.3	9.59	10.56	44.042	43.669
Run 655 scan 2	4.3	3.5	2.8	8.3	9.958	11.15	44.279	43.772
Run 655 scan 3	4.3	3.8	2.9	10.6	10.66	11.61	45.737	45.577
Run 655 scan 4	3.8	3.2	4.0	4.6	7.557	8.389	41.636	41.615
Run 655 scan 5	4.0	3.5	3.2	6.1	9.262	10.31	44.16	43.916
Run 655 scan 6	3.5	3.0	25.3	39.9	3.457	3.647	31.121	29.843
Run 655 scan 7	4.3	3.9	3.0	11.6	10.68	11.58	46.073	46.026
Run 655 scan 8	3.8	3.2	4.0	4.7	7.508	8.493	41.627	41.657

Table 8.2: Oscillations at different analogue (Vcc) and digital (Vdd) voltages. Chisq/NDF. Sigma and Mean from the complementary error function fits are also plotted to give an indication on noise and gain. We see that mean (gain) increases with analogue (Vcc) voltage.

8.6 Crosschecks at Cambridge, RAL and CERN

Since the measurements gave very similar results in Uppsala and Bergen, the oscillations seemed like a module problem rather than a setup problem. To further check this hypothesis, module 6 was tested in a setup at CERN, and showed the same oscillating behavior here. It was then sent, together with module 8 and 9, to two testing sites in the UK cluster; Cambridge and RAL (Rutherford Appelton Laboratory), where the modules produced in the UK have been tested. These modules have not shown any oscillations in the UK setups, so testing our modules here would give us a good indication on whether the oscillations came from the module or the setup. The results from this testing is shown in table 8.4.

It was found that the modules oscillating most (006), showed similar oscillations at all sites. Module 9 oscillated in Bergen, but was completely free of oscillations when tested at RAL. For module 8, the results were opposite; the module showed very small oscillations in Bergen and Uppsala, but large oscillations at Cambridge and RAL. Apart from this, the results from the characterisations in UK were in agreement with those from Bergen and Uppsala.

Summary of characterisations at Cambridge, RAL, Bergen and Uppsala (module mean values)

		Cambridge	RAL	Bergen	Uppsala
Mod 009	**Input noise**		1570.1	1592	
	Gain		51.7	51.7	
	Offset		35.7	35.6	
	Noise Occ		1.97E-05	1.24E-05	
	Est. in noise		1497	1521	
	Temperature		27°C	27°C	
	Oscillations		No	Yes	
	Comment	No oscillations at RAL, medium+ oscillations in Bergen. Oscillations could be removed by 'noisy' grounding scheme.			
Mod 006	**Input noise**	1513	1520.3	1573.5	1419.5
	Gain	52.8	52	52	55.4
	Offset	39.1	39.9	42.3	40.8
	Noise Occ	1.58E-05	1.43E-05	1.45E-05	2.71E-05
	Est. in noise	1491.4	1484.4	1589	1614.1
	Temperature	27°C	27°C	27°C	0°C
	Oscillations	Yes	Yes	Yes	Yes
	Comment	Similar oscillations at all four sites.			
Mod 008	**Input noise**	1515.9	1560.7	1385.3	1391.9
	Gain	51.7	51.9	54.6	54.8
	Offset	37.9	37.8	37.1	39.4
	Noise Occ	1.44E-05	1.63E-05	2.48E-06	1.42E-05
	Est. in noise	1487.3	1480.3	1358.5	1537.1
	Temperature	27°C	27°C	0°C	0°C
	Oscillations	Yes	Yes	Very small	Small
	Comment	Medium oscillations at Cambridge and RAL. Small oscillations in Bergen and Uppsala. The s-curves in Scand/UK are taken at different temperatures, but oscillations on other modules have shown little dependence on temperature.			

Table 8.3: Summary of characterisations at Cambridge, RAL, Bergen and Uppsala [26,27]. We see that the characterisations from the UK Cluster are in agreement with those from the Scandinavian Cluster. There are however differences in the oscillations seen by UK and Scand.

8.7 Oscillations conclusion

We can make some conclusions regarding oscillations based on the tests performed:

- Oscillations on Scandinavian modules are seen also at test sites at CERN and in the UK. The UK has previously not seen any serious oscillations on its own modules. There are however not complete agreement between the oscillations seen in Scand and UK. For modules 006, 008 and 009, the oscillations in the UK

were respectively the same, larger and smaller than the oscillations seen in Bergen and Uppsala for the same modules. This suggests that there are differences in our test setups, but the UK does nevertheless see larger oscillations than they have seen on any of their own modules, indicating that the oscillations is a module problem.

- We have not found a grounding scheme for the setup that can be said to improve the oscillations without increasing the noise.
- Oscillations show little or no dependence on temperature.
- The oscillations are probably a feedback effect, and increase when we apply a higher voltage to the chips, because of the higher gain.
- The oscillations disappear at ~1.0 fC, which is the ATLAS operating threshold, making the oscillations irrelevant under normal ATLAS operation. It is however an effect that it is important to understand and explain.

We can conclude from this that the oscillations seem to be a problem intrinsic to the SCT module. The problem either lies in the components or in the assembly of the module; this is at the time of writing still an unresolved issue, and is subject to investigation.

However, although UK sees oscillations, they are not in agreement with the oscillations seen in Uppsala and Bergen for all modules. Why these differences show up, is a question that should be further looked into.

Conclusion

The setup for SCT barrel module testing at the University of Bergen can now perform the electrical QA tests according to the specifications given in [26].

External noise has been eliminated from the system to a satisfactorily degree, and the results from the system have been cross-checked with results from other testing sites for 8 modules tested in the setup. Infrastructure for publishing results on the web is available and working, and the published results can be found at [27]. There is however some tasks that remain to be done:

- **We need to be able to test more modules than 1 at a time.** This would greatly improve testing time, since the time needed to test a module is quite long. To do this, we need both a HV and LV supply that supports more than 1 module. Supplies for supporting six modules are under development in Oslo and Uppsala. This will also enable us to do IV-curves as an integrated part of the electrical testing. We also need to create a reliable system for physically placing the six modules in the environmental chamber.

- **A stringent testing procedure must be developed.** For making the testing procedures as simple and effective as possible, detailed instructions on how to test SCT modules, how and where to save data, what to do if problems occur and so on, needs to be created.

- **The setup must be moved to a clean-room.** During the work with the setup, the setup has been placed in an ordinary room. This is not satisfactorily clean for serial module testing, and the setup must thus be moved.

- **The test results need to be uploaded to the CERN SCT database.** An early version of the SCTDB interface is available, but this has not yet been tested for the data from Bergen. First, data from module construction and module components must be registered in the database. This is at the time of writing not yet completed, but the system will probably be fully operational when the series testing is ready to start. Until that, the QA results from the cluster are published on the web [27,28].

- **Further insight into oscillating s-curves.** The origin of the oscillating s-curves, and why the same modules oscillate differently at different sites, needs to be understood. The oscillations seems to be an effect intrinsic to the modules, but the difference in oscillations between Scand and UK also suggests that we have differences in our test setups that we can not yet account for.

The Scandinvian Cluster is still not qualified for module series production by the SCT project management, but once it is, it is important that these issues are resolved and the system is ready for series testing.

References

[1] Bjørn Pommeresche, *A setup for electronic tests of ATLAS silicon strip detectors*, Cand. scient thesis, University of Bergen 2001

[2] CERN Photo database, *http://info.web.cern.ch/info/Press/PhotoDatabase/*

[3] LHC-General information, *http://lhc.web.cern.ch/lhc/general/gen_info.htm*

[4] Fabiola Gionotti, *Collider physics: LHC,* ATL-CONF-2000-01.

[5] ATLAS Collaboration, *Detector and Physics Performance, Technical Design Report,* CERN/LHCC/99-14, 22. May 1999.

[6] ATLAS Collaboration, *Inner detector Technical Design Report,* CERN/LHCC/97-16, 30. April 1997

[7] Ulrik Egede, *The search for a Standard Model Higgs at the LHC and electron identification using transition radiation in the ATLAS tracker*, Ph. D. thesis, Lund University, ISBN 91-628-2804-5

[8] K. Hagiwara et al. , *The Review of Particle Physics,* Physical review D66, 010001 (2002)

[9] Serway, Moses, Moyer, *Modern Physics,* Saunders College Publishing 1997, ISBN 0-03-001547-2

[10] Povh, Rith, Scholz, Zetsche, *Particles and Nuclei 2nd edition,* Springer Verlag 1999, ISBN 3-540-66115-8

[11] A. Carter, J. Carter, J. Morin, Y. Unno, *SCT Barrel Module: Requirements and Specification of Barrel Modules,* SCT-BM-FDR-4, May 2001
(http://atlas.web.cern.ch/Atlas/GROUPS/INNER_DETECTOR/SCT/module/SCTbarrelmod.html#FDR)

[12] Y. Unno, *Atlas 98 silicon microstrip detector layout,* Sep. 1998

[13] Lars G. Johansen, *Probing p+-type and n+-type silicon microstrip detectors for the ATLAS Semiconductor Tracker,* Cand. Scient Thesis University of Bergen March 1999.

[14] R. Apsimon, A. Carter, J. Carter, *SCT Barrel Module: Module Components section 5.2: Baseboards,* SCT-BM-FDR-5.2, May 2001

[15] Y. Unno, T. Kohriki, T. Kondo, KEK, *Cu/Polyimidebarrel hybrids*, SCT-BM-FDR5.3, May 2001.

[16] ABCD3T Project Specification V1.2, July 24, 2000, *http://scipp.ucsc.edu/groups/atlas/elect-doc/abcd3t_spec.pdf*

[17] An introduction to VME, *http://www-esd.fnal.gov/esd/catalog/intro/introvme.htm*

[18] D. A. Hayes, J. B. Lane, M. Postranecky, *Some notes on design and functionality of the "CLOAC" (C+C Master) VME module for the ATLAS –SCT System &Beam Tests,* University College London, 15. november 2000

[19] M. Morrissey, *SLOG Documentation,* August 1999

[20] M. Morrisey, M. Goodrick, *MuSTARD Documentation,* September 1998

[21] J. Stastny, I. Polak, *VME SCT LV Power Supply,* Institute of Physics, Academy of Sciences of the Czech Republic.

[22] ROOT homepage: *http://root.cern.ch/*

[23] SCTDAQ homepage: *http://s.home.cern.ch/s/sct/public/sctdaq/sctdaq.html*

[24] P.W. Phillips, L. Eklund, *Electrical Tests of SCT Hybrids and modules, rev. 3.38,* http://hepwww.rl.ac.uk/atlas-sct/documents/Electrical_Tests.htm

[25] Design Environmental Limited, *Instruction Manual for Model Ref: BS125-40*

[26] R. Apsimon, J. Carter, D. Charlton, L. Eklund, P, Phillips, Y. Unno, *SCT Barrel Module: Module QA,* SCT-BM-FDR-7, May 2002

[27] SCT Barrel module electrical QA Bergen, *http://www.fi.uib.no/Atlas/*

[28] ATLAS at Uppsala University, *http://www4.tsl.uu.se/~atlas/*