The ATLAS Semiconductor Tracker Barrel System Test

P.J. Bell, P. Brückman, P. Ferrari, S. Gadomski, J. Grosse-Knetter,
C.A. Magrath, M. Morrissey, K. Nagai, J. Pater, P.W. Phillips, B. Skubic

Version 5

Abstract
Results are presented of tests with a prototype of a partial section of the ATLAS SCT barrel. Up to 16 non-irradiated and 6 irradiated modules have been tested on this system and showed a performance well within specifications. Different grounding and shielding configurations have been tried and sufficient protection from external noise has been achieved.

1 Introduction

ATLAS is one of the two multi-purpose detectors for new physics in the TeV-mass range which are being constructed for operation at the LHC [1]. ATLAS comprises several layers, extending from the Inner Detector (ID), to the super-conducting magnets, calorimeters, and muon chambers. The ID is designed for precision tracking of charged particles with 40 MHz bunch crossing identification [2]. It combines tracking straw tubes in the outer transition-radiation tracker (TRT) with the semiconductor microstrip detectors of the semiconductor tracker (SCT) in the middle and semiconductor pixel detectors in the innermost PIXEL component.

The SCT [2, 3, 4] consists of four nested cylindrical barrels in the centre and nine disks in each of the two end-caps. The barrels cover a region from 300 to 520 mm in radius and have an active length of 1530 mm, centred around the interaction point. The respective barrel layers are fully covered by 32, 40, 48, and 56 rows of twelve identical modules, making a total of 2112 modules.

The SCT modules [4, 5] are made of four sensors, glued in pairs on either side of a thermally highly conductive baseboard. The sensors are approximately $6{\times}6$ mm^2 p-in-n silicon wafers of approximately 300 μm thickness. Each sensor has 768 readout strips with 80 μm pitch. The modules are equipped with hybrids holding twelve readout chips, six on the top and six on the bottom side of the module. Each chip provides high-gain (50mV/fC) fast-shaping (20ns peaking time) binary readout of 128 channels using radiation-hard technology. All channels have a common threshold for discrimination, which can be adjusted individually for each channel with a 4-bit trim DAC to cope with the increased offset spread after irradiation. The heat produced by the chips and the detectors is extracted via the baseboard at one end of the module. Tests with individual prototype modules showed a performance well within the specifications [5, 6].

The main goal of the system test is to run as many modules as possible in a physical configuration which is as close as possible to that of the planned ATLAS SCT, thereby testing the performance of the modules in such a system and comparing to their stand-alone performance [3, 5, 6]. The system test complements beam tests and irradiation tests [5, 6, 7, 8, 9, 10, 11] in assuring full functionality of the final SCT.

2 Setup of the system test

The system test simulates a section of one of the SCT barrels in all its details of mounting, cooling, supply, and shielding. Figure 1 shows a schematic drawing of an SCT barrel in the design as used for the system test setup[1] and figure 2 shows a photograph of the system test setup. Specifications of the location of objects on the sector (left, front, etc) refer to the orientation as shown in figure 2.

2.1 Mechanical support

Modules are mounted on a sector of a carbon-fibre-Corex-sandwich cylinder with dimensions very near to those of the innermost ATLAS SCT barrel. The mechanical support of modules is realised with carbon-fibre brackets which are mounted on the Corex sector. The sector can hold up to 48 modules, in four rows of twelve. The rear row is not supplied with power or optical-readout fibres, nor are the right halves of the two front rows.

The Corex sector is mounted on three aluminium feet which are supported by a carbon-fibre rod. The rod itself is fixed on either end by vertical aluminium plates. The whole system sits on an aluminium base.

A plastic cover with square cross-section, complementing the aluminium support, is used to provide a gas- and light-tight enclosure.

2.2 Cabling

Modules are powered and read out via prototype *opto-harnesses* [13], which transmit low and high voltages to the modules. Data and command transfer employ optical signals (see section 2.4 below) which are processed by dedicated circuits on the harnesses [14]. These circuits decode the incoming optical clock and commands, transmit the signals electrically to the modules, and transform the data from the modules into optical signals for transmission to the readout electronics. Each opto-harness serves up to six modules. Each harness consists of six pairs of approx. 1.5m-long 50 μm-aluminium-on-kapton tapes and six bunches of three optical fibres. Six *doglegs* interface the power lines on the main tapes to each module and carry the optical circuits. Thus two harnesses are needed per row of twelve modules. These two harnesses are separated from each other and connections are made only at their respective ends outside the sector enclosure.

To guard against data loss due to defects on the optical circuits, redundancy options exist both for the incoming clock and commands and for the outgoing data. Each optical circuit can optionally produce a copy of its (electrical) clock and command signals and send it to one of its neighbouring doglegs to replace defective signals there. Small copper-kapton tapes connect neighbouring doglegs for the transport of these redundancy signals. The data from each module comes in two streams, one for each side, which are separately transformed into optical signals. In case of problems with one of the data streams both module sides can be read out via one stream.

All patch panels and power tapes used are close to the final ATLAS design, except for extra provisions on the patch panels to allow testing of various coupling schemes. The opto-harness ends are connected to a first patch-panel (PPB1) that provides filtering on the power lines (see section 2.6 below). It interfaces to 4m of 100μm aluminium-on-kapton power tapes which connect to the next patch panel (PPB2). The following section of the cable run is not very close to the ATLAS design since a simplified version of the foreseen third patch panel (PPB3) is directly connected to PPB2. The 30m-long conventional cables connecting PPB3 and the power supplies are shorter than the required length for ATLAS although the cable resistance is similar. Dedicated tests with full cable length for a subset of the modules are described in section 5.7. The optical signals are transmitted through 25m-long fibres.

2.3 Cooling

The hybrids and detectors of each module produce between 5 and 10 W of heat which must

[1]Several details in the design have changed since construction of the system test setup, such as the shape of the cooling pipes.

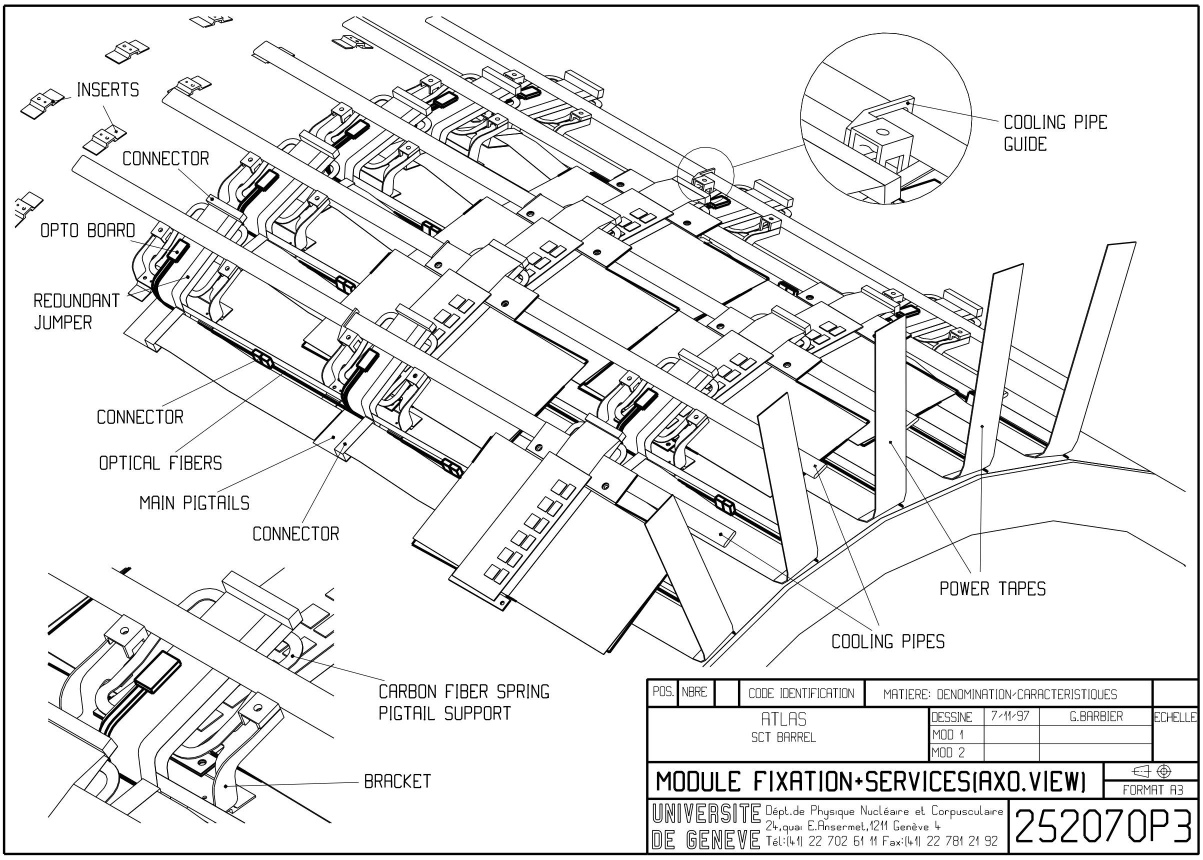

Figure 1: Drawing of a section of a SCT barrel in 3D-view in the design as used for the system test setup [12].

be dissipated. For this purpose the three front rows on the sector are equipped with aluminium cooling pipes. They have a flat-oval cross section onto which the module baseboards are clamped. The pipes are mounted into an aluminium support bracket attached to the sector on either side. An ethanol-water mixture is pumped through the pipes; the temperature of this fluid is controlled by a commercial chiller as described in section 2.7.

2.4 Readout

The modules are read out using a VME-based prototype system [15] that consists of four components as shown in figure 3. Two components, *Cloac* and *Slog*, generate the 40 MHz-clock, command and trigger signals, while a third, *Mustard*, decodes the data received from the modules. All three are connected to an opto-transceiver, *Optif*, which performs the conversion of the clock and command from electrical into optical signals, and the reverse for the data from the optical fibres from the harnesses. These DAQ components are used in place of the final-design readout drivers (ROD) [2]. The low-voltage power supplies (see section 2.5 below) are also used to read module temperatures for monitoring purposes. A ROOT-based DAQ-software package is used [16], running on a Windows-NT PC which is connected to the VME crates via a National Instruments interface card.

2.5 Power supplies

Power for the chips on the modules and the optical circuits on the doglegs is provided by prototype VME supplies (*SCTLV*); detectors are currently biased with the companion *SCTHV* units [17]. These units provide two separate low-voltage channels (one for the analogue circuits and one

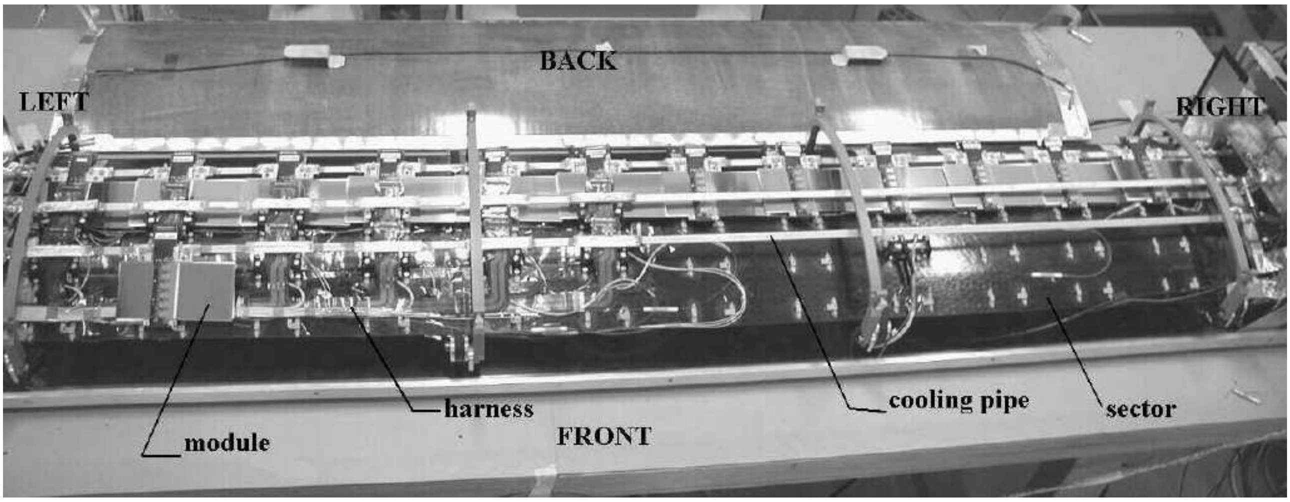

Figure 2: Photograph of the barrel system test showing a configuration with one row completely equipped with modules.

for the digital circuits on the modules, later referred to as analogue/digital voltage/ground) and one separate high-voltage channel for each module. The individual channels are floating with respect to each other. It should be noted, however, that digital, analogue, and high-voltage grounds are connected on the modules.

2.6 Grounding and shielding

The system test considered several grounding and shielding schemes, the two main configurations being based on the two proposals for the final detector [18, 19, 20]. The key parts and their connections, with possible options to modify the shielding scheme, are outlined in section 2.6.1. The following two sub-sections 2.6.2 and 2.6.3 list the actual options used to implement the two proposals.

2.6.1 General layout

The cooling pipe ends can be connected to the aluminium parts of the support structure which provide the reference point of the shield. Each pipe is split electrically in its mid-point, with the use of a plastic joint to provide mechanical continuity. Each pair of pipe-halves can be re-connected electrically with a short copper cable.

To control stray capacitance between the cooling pipe and the SCT detector backplane, shunt shields are placed between the modules and the cooling pipe. These shunt shields are 36 μm-copper-on-kapton pads 13$\times$75 mm^2 in size, the copper being soldered to analogue ground on the dogleg. Optionally, analogue ground can be connected to the cooling pipes with a short copper cable, effectively shorting the shunt shields and referencing the module ground to the pipes.

The sector is also connected electrically to the metalwork which supports it. A cover made of fibreglass with a 100-μm-thick aluminium foil on its underside is placed approximately 10 cm above the carbon fibre structure and simulates the thermal shield, an aluminium barrel surrounding the four SCT barrels. The cover is connected to the metal support structure at either end. The brackets that hold the modules are connected to digital ground on the dogleg, individually for each module.

The power tapes for each harness are wrapped together with 25μm-aluminium-on-25μm-kapton shielding from the end of the thermal shield to PPB1. The wrap is directly connected to the metal surrounding the sector. At PPB1 digital and analogue grounds are capacitively connected (with 2.7 μF) to the cable shielding. As an option, digital ground can be directly connected to the shield. PPB1 provides similar capacitive filtering of all other power lines to the cable shielding. The shield on the PPB1s on either side of the sector are connected to two 300$\times$900 mm^2-sheets of 100μm-thick copper which are connected

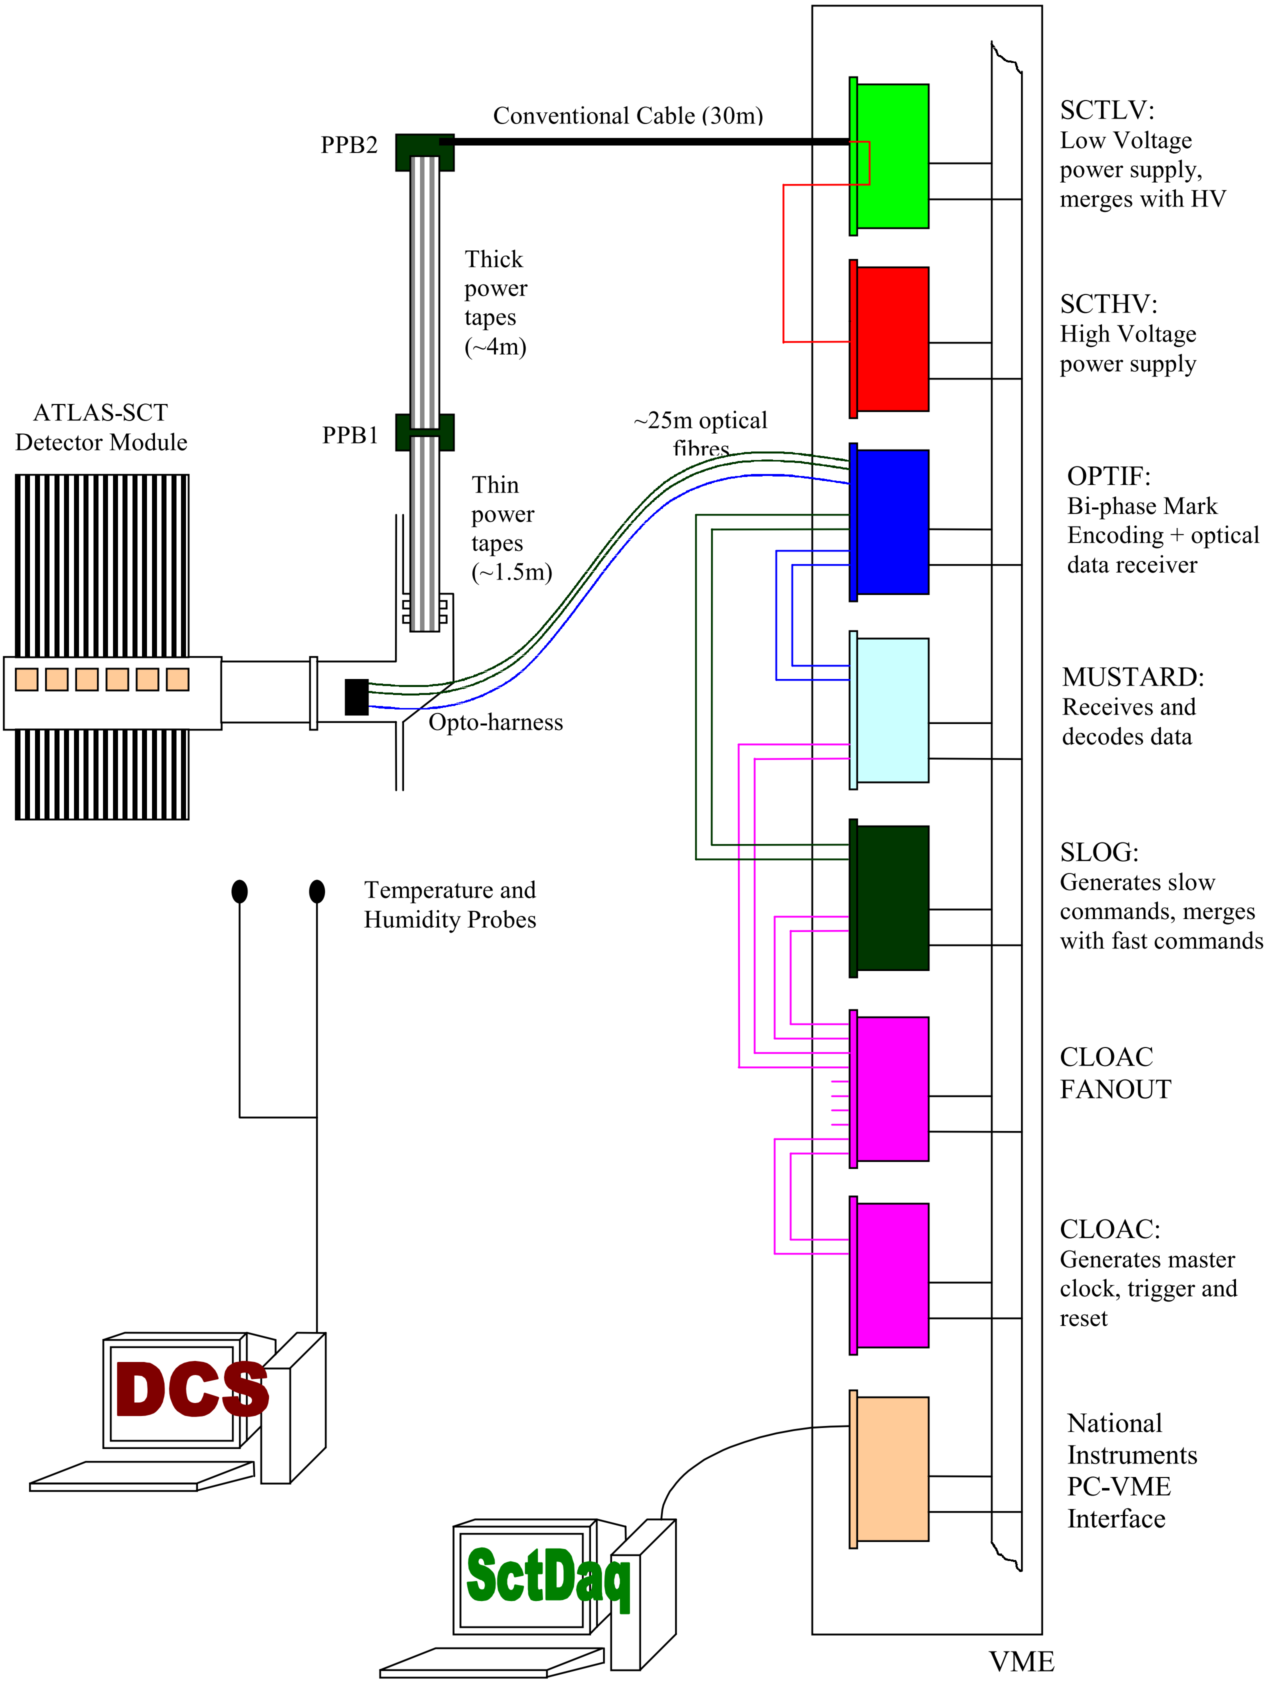

Figure 3: Schematic diagram of the system test DAQ-system.

to the main metal structure at their other end. The sheets are used to simulate the extra shield connection provided by the radial heat spreader plates running radially outward from the thermal shield to the PPB1 location in the SCT detector.

From PPB2 to PPB1 the system test shield differs from the ATLAS SCT; there is no surrounding metal (e.g. heat spreader plates, cooling supply tubes) to connect to as is intended [18]. However, the power tapes are wrapped in bundles of six in the same aluminium-on-kapton foil as the harness ends. The wrap is connected to the copper sheet at PPB1.

Capacitive filtering to the shield and common-mode chokes are used at PPB3 but as discussed above there is no cable between PPB2 and PPB3. The conventional cable shield is connected through at PPB3, optionally as a simple cable or connected through one of the coils, and connects to the tape wrap at PPB2. On the low-voltage supplies the cable shield connects capacitively either to the respective digital grounds or to the crate metal.

All shielding connections are checked regularly with an LCR meter to ensure constant quality. All connections typically show an AC resistance at 100 kHz of a few tens of mΩs or better.

2.6.2 Scheme A

This scheme involves direct connections between all metal in the system and direct connection of the module grounds to the cooling pipe. The main features are as follows [20]:

- Shield of conventional cables connected to crate metal.
- Shield routed through chokes at PPB3.
- Cooling pipes electrically connected in the middle.
- Cooling pipe halves connected to the respective end plate.
- Shunt shields between modules and cooling pipes *shorted*.

2.6.3 Scheme B

This scheme attempts to minimise the flow of externally induced currents involving the detectors. The key feature is the shunt shield reducing the stray capacitance between module and cooling pipe [18]. Other features are:

- Shield of conventional cables connected to digital grounds at LV supplies.
- Shield routed through chokes at PPB3.
- Digital ground connected to shield at PPB1 for each power tape.
- Cooling pipes electrically split in the middle.
- Cooling pipe halves connected to the respective end plate.
- Shunt shields between modules and cooling pipes are *used*.

2.7 Operating conditions

The system test includes a prototype of the *detector control system* (DCS), which monitors ambient temperature and humidity as well as the cooling pipe and module temperatures. In addition it is used to control the chiller. One important safety concern is the possibility of condensation forming on sensitive components on the sector (modules, opto-chips), therefore during operation the dew-point inside the sector enclosure is measured by the DCS system and used as a safety limit. Similarly, the module and pipe temperatures are checked for overheating.

Under standard operating conditions, the enclosed volume is flushed with dry nitrogen while running, thus providing an atmosphere with a relative humidity of 40% or less. The only cooling is provided by the cooling pipes operating at 10°C, resulting in temperatures measured on the module hybrids between 25 and 30°C. The ambient temperature typically ranges from 15 to 20°C.

In order to simulate the situation of running at low temperatures in ATLAS, when a detector temperature of -7°C will be required [2], a *cryophasing* procedure has been established to safely cool down the whole sector to temperatures below zero, take data, and subsequently warm it up again. The actual cool-down or warm-up is performed using the chiller to cool the pipes, with the sector situated inside a commercial freezer which is flushed with dry nitrogen to keep the humidity low. As for the warm operation, DCS is used to check for overheating or risk of condensation. After cool-down the pipe temperature is -16°C, resulting in temperatures measured on the module hybrids between 0 and 5°C. The ambient temperature is set to either -16 or -10°C.

Cold operation is a prerequisite for the operation of irradiated modules on the sector, due to the possibility of thermal runaway caused by the increased leakage currents drawn by irradiated detectors.

The chips on the modules are operated at their nominal voltages of 3.5 V for the analogue and 4 V for the digital circuits, respectively. Unless noted otherwise, the detectors of unirradiated modules are biased with 100 V, and those of irradiated modules with 350 V, so that the detectors are fully depleted [8].

3 Analysis methods

This section describes different methods of testing the noise performance of the system.

3.1 Module readout

All analysis methods use the binary hit information constantly being read out by the modules every 25 ns. These events are written into 132-slot-deep FIFOs on the chips such that upon issuing a trigger command the last event in a FIFO is written to the output data stream. This means that the event sent back corresponds to a time bin from $25\,\mathrm{ns} \times 132 = 3.3\,\mu\mathrm{s}$ before the arrival of the trigger command at the module. The data acquisition is handled sequentially, i.e. one trigger is sent, the data is read and decoded, after which another trigger is issued etc.

The following trigger modes are used:

- regular triggers, typically at a rate of a few tens of kHz;
- regular triggers with calibration pulses: as above, using the calibration circuits on the module hybrids to inject a well defined charge[2] into the pre-amplifiers, individually for each detector strip, adjusted to coincide with the trigger command.
- two successive triggers: two triggers sent within a short interval (between 75 ns and 10 μs is used), ignoring the data from the first trigger.
- external triggers: initiated from an outside source and adjustable for timing between the external trigger and the trigger sent to the module.

3.2 Calibration scans

Calibration scans are used to determine the amplifier gain and the equivalent noise charge Q_{noise} at the inputs to the pre-amplifiers, for every channel of a module. Regular triggers in conjunction with calibration pulses are used, and the occupancy is measured as a function of increasing threshold.

Repeating such scans for several values of injected charge yields the gain from a set of discriminator threshold values at 50% occupancy, taken from fits with an error function, against charge. The gain determined in this way is used to convert the noise as seen after the amplifiers, derived from the error function, into equivalent noise charge (ENC), as seen on the amplifier input.

The standard approach is to use injected charges of 1.5, 2.0 and 2.5 fC to determine the gain, with the actual noise value being taken from the result of the 2 fC scan. The gain and the ENC value are usually quoted per chip, quantified by averaging over all its 128 channels.

[2] The magnitude of the injected charge is defined by the amplitude of a voltage pulse and the nominal value of a capacitor. The actual capacity of the latter can deviate from the nominal value by up to 20%. Any difference from the nominal capacity is measured before module assembly and taken into account during operation.

3.3 Noise occupancy scans

The noise occupancy is measured as a function of threshold charge, Q_{thr}, with the number of regular triggers being varied between 2000 and 10^6 to compensate for decreasing number of hits as the threshold increases. The calibration of the Q_{thr}-scale is as described in section 3.2.

To a good approximation the logarithm of the occupancy O depends linearly on the square of the threshold for $Q_{\mathrm{thr}} \gg Q_{\mathrm{noise}}$ [21]:

$$\ln(O) \approx -\frac{1}{2}\frac{Q_{\mathrm{thr}}^2}{Q_{\mathrm{noise}}^2}\,. \tag{1}$$

Deviations from the linear behaviour indicates (non-Gaussian) noise pick-up. If a linear behaviour is observed an appropriate fit is used to determine Q_{noise}. The fit range is restricted to $Q_{\mathrm{thr}} > 0.5$ fC to be well above Q_{noise} which is typically around 0.2 fC.

This method offers a way to measure Q_{noise} alternatively to the calibration-scan method discussed above. It should be noted, however, that such a calibration is needed in order to determine Q_{thr} so that this method and the calibration scan cannot be considered as totally independent of each other.

3.4 Repeated noise measurements

The calibration measurement described in section 3.2 is repeated for several hours. Both the variation of the gain and of the noise is analysed. With many modules in operation this measurement is very slow – it typically takes about a half-hour per individual measurement.

In a different mode, the noise occupancy is measured repeatedly in a regular trigger mode. Repeated bursts of 10^6 triggers at a threshold of $Q_{\mathrm{thr}} = 1$ fC, the threshold at which the SCT will be operated, are used. Both the occupancy as a function of time and the average occupancy are analysed. This method is much faster than the above; it typically takes a few minutes for an individual measurement, which makes it very sensitive to fluctuating noise pick-up.

3.5 Common mode noise

Common mode noise pick-up is a particularly worrying problem for binary readout systems as used for the SCT modules. Thus special analysis methods are dedicated to this problem [21]. Data is taken in a regular triggering mode, storing the full hit information for each event. The number of hits H per event on a set of strips, typically the 128 strips on one chip, is determined. The mean and variance of H over a large sample of events are related assuming a random noise source (the strips). Any deviation from this relation indicates common mode noise pick-up, and the amount of pick-up is determined subsequently without need of comparison to "bench mark" measurements taken outside the sector [22].

Ideally the event samples should be taken at the SCT-operation threshold of 1 fC. However, to compromise between the need to take a large event sample for sufficient number of total hits (which decreases with threshold) and the aim of being close to the 1 fC threshold, data are normally taken at $Q_{\mathrm{thr}} = 0.8$ fC. No difference in the result is expected [22].

3.6 Double-triggers

In this mode two successive triggers are sent so that the second trigger probes what happened when the data corresponding to the first trigger was sent out. This tests for interference effects or voltage drops caused by a module writing out its data. Such effects should be seen when the delay between the triggers is approx. 3.3 μs (see section 3.1). This corresponds to a trigger rate of 300 kHz which is well above the expected maximum average trigger rate in ATLAS of 100 kHz. However, according to the specifications, triggers may be succeeding each other within two bunch crossings, ie 50 ns [2].

The rest of the analysis is very similar to an individual measurement of the repeated noise occupancy analysis detailed in section 3.4. The information taken with the first trigger is ignored since this trigger is only used to initiate activity in the readout, while the data from the second trigger is used to calculate the occupancy. The result is compared to the occupancy measured with a burst of single regular triggers.

3.7 Crosstalk

To measure crosstalk between modules, a real signal must be induced in the detector part of the modules. To first order, this is checked in the test beam on a single-module level. However, since multi-module geometry is only correctly simulated in the system test, crosstalk between two or more modules can only be tested here.

For the signal production, a ^{90}Sr-source is used, emitting β^- radiation with an energy spectrum up to approximately 2 MeV. A scintillator is put under the module into which the signal is to be injected, such that the main axis is parallel to the module strips. The scintillator serves as an external trigger source, see section 3.1. It is connected to Cloac via a discriminator and the trigger timing is adjusted for maximum efficiency.

The data-acquisition mode combines this external triggering with writing the full hit information for each event, similar to the common mode noise analysis discussed above. Events with an isolated hit strip on the top and the bottom side of the target module are selected offline. To determine the beam profile the 2D-coordinates of each hit are calculated, event by event. For the coordinate perpendicular to the strips, X, the average of the position of the hit strip from top and bottom side is used. The coordinate parallel to the strips, Z, is determined using the 40 mrad relative rotation between top and bottom side. The crosstalk is estimated from the probability $O_{\rm src}$ of getting a measured hit in conjunction with a real hit from the source, comparing to a set of data without source yielding the probability $O_{\rm ref}$. Both the difference in $O_{\rm src}$ and $O_{\rm ref}$ as well as the quantity

$$\chi = (O_{\rm src} - O_{\rm ref})/\sqrt{2}\sigma\,, \qquad (2)$$

where σ is the spread in probability which is estimated from the reference measurement probability, are analysed.

3.8 Noise injection

The noise level in the laboratory in which the measurements were made is considered relatively low compared to what is expected inside ATLAS. To enforce a situation of large noise, a sinusoidal noise current is injected into various components of the system by a magnetic field created with a ferrite and a coil, the latter being connected to a signal generator. The ferrite is placed around three different components:

1. an individual cable between power supplies and PPB3;
2. a bundle of six tape pairs between PPB2 and PPB1 (corresponding to one harness), including their shielding wrap;
3. a cable connected to the copper-shields at PPB1 on either side via a 1 Ω resistor each.

The first two items simply simulate noise pick-up on the cables, the last item simulates noise current running through the central shield parts (thermal shield, cooling pipes) which is expected since the central shield forms an electrical short between cable wraps coming in from either side of the barrel.

The noise is measured with a single calibration scan (see section 3.2) with 2 fC injected charge. To convert the measured noise to ENC, archived values of the gain from calibration scans are used. The amount of noise pick-up is determined as the difference in noise measured with an injection scan and a scan with the signal generator still connected but set to zero output.

4 Individual tests

When a module arrives at the system test, it is accompanied by the results from a standard characterisation as performed at the manufacturing institute [16]. The first step in integrating the module into the system test is to repeat this standard characterisation on an *electrical-readout test bench* to verify that the module has not suffered in transit. The electrical-readout test bench is considerably simpler than the full system test as it bypasses the optical communication, employs only very short power and signal cables, and has well controlled grounding and shielding since the module is mounted in a metallic box. Therefore a module running stand-alone on the test bench tends to give its best performance. The noise figures from the calibration scans are used as bench marks when analysing similar data from multi-module runs on the sector.

Once the module is verified to be in good working order, it is mounted on the system test sector with all grounding and shielding connections made, and the standard characterisation sequence is repeated, powering only that one module, and this performance is compared to its performance on the bench. Any differences are noted, and investigated if possible. When this comparison is complete, the module is considered ready to be included in multi-module tests. A list of all modules used is given in table 1. The first four modules listed in the table were made with an older version of readout chips and module 0009 would not qualify as a good module due to three noisy chips. Thus those five modules are only used to complete a certain geometry of modules but not quoted for noise measurements despite being operated like the other modules.

Part of the characterisation sequence applied during these individual tests are full trimming (adjustment of the 4-bit trim DACs) and calibration scans [16]. The results of these are used to initialise and calibrate all modules in all subsequent multi-module runs.

5 Tests with several modules

This section contains a summary of the most important results produced in the system test. The full information is available elsewhere [23].

5.1 Twelve modules in a row

A row of twelve modules is the most natural subset of a completely equipped barrel. In particular, the two harnesses concerned come in from the left or right, respectively, and are floating with respect to each other. It must thus be checked for noise pick-up and crosstalk between the two halves of the row. The measurements described in this section use twelve unirradiated modules mounted on the right harness and the rear left harness, similar to the arrangement shown in figure 2.

The measurements described in this section use grounding and shielding scheme A (see section 2.6.2) since no difference to scheme B is observed in the measurements without noise being injected into the system. (Measurements with injected noise and detailed shielding studies are discussed in section 5.3.)

The most basic measurement is the calibration scan (see section 3.2), which returns the noise of each module using an injected charge of 2 fC. The result is comparable to noise results from individual measurements (see section 4) as shown in figure 4. The tendency of the sector results to be lower for some modules might come from advance biasing to avoid the effect of initial reduction in noise (see discussion below), a precaution not taken for the individual measurements. The noise ranges from 1400 to 1800 ENC.

Scanning the noise occupancy O versus the threshold as described in section 3.3 shows the linear behaviour of $\ln(O)$ against Q^2_{thr} expected from eq. (1). An example of $\ln(O)$ against Q^2_{thr} with a linear fit is shown in figure 5. The noise values derived from such a fit as shown in figure 6 look similar to the results from the calibration scans discussed above, thus giving no hint of significant noise pick-up. The tendency of the noise occupancy scan to return smaller noise values than the calibration scans is in agreement with the observation of a dependency of the measured noise on the injected charge for calibration scans: results for an injected charge below 1 fC can yield noise values up to 20% lower than those seen for higher values of injected charge.

The noise occupancy at a threshold of 1 fC shown in figure 7, obtained from more than 100 repeated individual measurements (see section 3.4), is well below the quality limit of $5 \cdot 10^{-4}$ required for the readout speed to not be dominated by noise hits [2].

Name	Database S/N	chip	irrad?	gain (mV/fC)	noise (ENC)
k3103	none	2T	NO	57.5	1462
k3104	none	2T	NO	57.3	1507
scand1	none	2T	NO	58.9	1426
k3112	none	2T	NO	57.9	1197
0008	20220170100008	3T	NO	48.0	1598
0009	20220170100009	3T	NO	47.9	1909
0011	20220170100011	3T	NO	45.3	1599
0018	20220170100018	3T-A	NO	50.7	1579
0019	20220170100019	3T-A	NO	46.6	1724
0020	20220170100020	3T-A	YES	34.3	1750
0022	20220170100022	3Tm	NO	48.1	1668
0026	20220170100026	3T	NO	46.3	1725
0029	20220170100029	3T-A	NO	49.4	1611
0030	20220170100030	3T-A	NO	49.6	1672
0035	20220170100035	3T-A	NO	46.3	1544
0036	20220170100036	3T-A	NO	46.8	1574
0037	20220170100037	3T-A	YES	34.1	1779
0044	20220170100044	3T-Aepi	YES	29.2	1714
0046	20220170100046	3T-Aepi	NO	49.9	1703
0047	20220170100047	3T-Aepi	YES	32.0	1649
0050	20220170100050	3T-A	YES	34.0	2053
0052	20220170100052	3T-Aepi	NO	50.5	1631
0007	20220170200007	3T-Aepi	YES	41.2	1820
0034	20220170200034	3T-Aepi	YES	34.6	1790

Table 1: List of all modules used in the system test. The names listed are used in the text to identify the modules, the serial number is that used for the SCT data base, the “chip”-column gives the version of the ABCD chip and the “irrad?” column specifies if the module was irradiated before being put onto the sector. Gain and noise (first four modules uncorrected for calibration capacitor deviations) are the average across the entire module, taken from the calibration scans run during the initial tests on the electrical-readout test bench at room temperature (hybrid temperature approx. 25°C) for unirradiated and at -30°C inside a freezer (hybrid temperature approx. -16°C) for irradiated modules.

The same set of measurements allows checking for the time behaviour of the occupancy. For the measurements shown here, the modules have been biased for several hours before the measurements were started, yielding a constant noise occupancy in time, see figure 8. If this precaution concerning the bias voltage is not taken, a strong drop in noise occupancy can be observed, the size of which varies from module to module. An extreme example is shown in figure 9. This behaviour is presumeably related to a similar change that has been observed in the silicon's inter-strip capacitance [24]. The noise measured with repeated calibration scans shows the same time dependence. It should be noted however that the gain measured in such scans is observed to be very stable over several hours, see figure 10.

Since none of the above tests show any sign of noise pick-up, the more sensitive technique of measuring common mode noise as described in section 3.5 is applied. A sample of 500,000 events is taken at a threshold of 0.8 fC. The measurement yields an upper limit in noise pick-up of less than 300 ENC. No distinct geometrical pattern in the noise pick-up is observed, see figure 11.

An occupancy measurement with bursts of two successive triggers as discussed in section 3.6 yields largely the same occupancy as a measurement with a burst of single triggers, except for a delay $\Delta t_{\mathrm{Trigger}}$ between the triggers in the range 3 to 4 μs. Figure 12 shows the occupancy for a single chip on one of the twelve modules as a function of $\Delta t_{\mathrm{Trigger}}$ in that range. A short increase followed

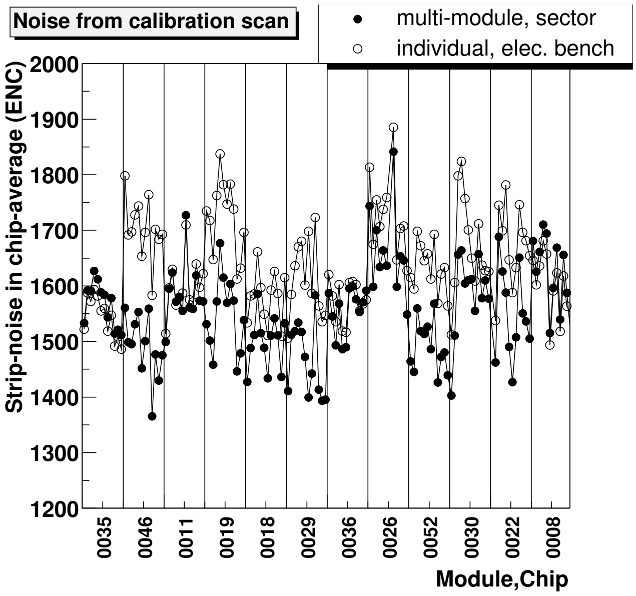

Figure 4: Noise as measured with calibration scans, comparing between a measurement of a row of twelve modules on the sector and individual measurement on the electrical-readout test bench. Each slot on the graph corresponds to one module as labelled on the X-axis, in order as they are mounted on the sector, and each of the symbols represent the average noise per chip. The noise is quoted in units of the electron charge.

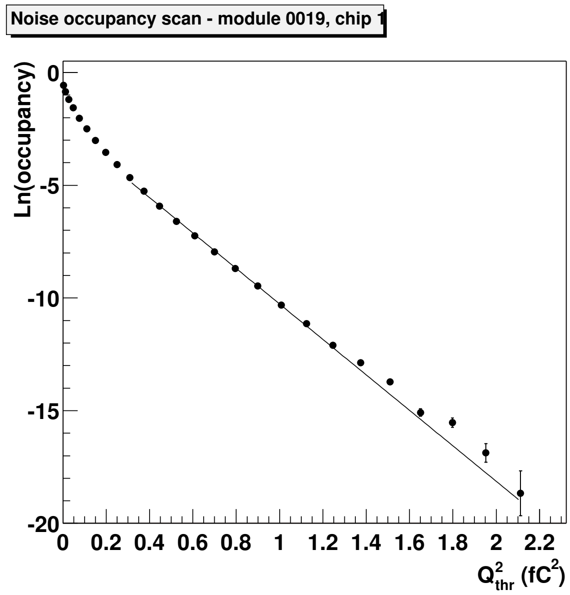

Figure 5: Noise occupancy scanned versus Q^2_{thr}, measured with a row of twelve modules on the sector. The plot shows the occupancy averaged over one chip of a selected module. The line shows the result of a linear fit according to eq. (1).

by a dip is observed. This effect is seen on all twelve modules and is essentially the same for all chips on the modules. Such a dip might be caused by a drop in voltage on the module hybrid, causing an effective change in the threshold. Threshold, occupancy, and noise are related by eq. (1), the approximation being valid for the threshold used here. A shift in threshold, ΔQ_{thr}, changes the occupancy by a factor R_O as follows:

$$R_O \approx \exp\left(\frac{Q^2_{\mathrm{thr}}}{Q^2_{\mathrm{noise}}}\,\Delta Q_{\mathrm{thr}}\right), \qquad (3)$$

in approximation for small ΔQ_{thr}. The observed change in occupancy of less than a factor of 10 thus corresponds to a change in threshold of less than 15%, resulting in a negligible loss in efficiency.

Crosstalk is measured according to the description in section 3.7, with the source placed over module 0029. The beam profile shown in figure 13 is reasonably well placed on the target-module and shows a Gaussian shape. Figure 14 shows the quantity χ as per equation (2) for each strip which is very low for all modules apart from that

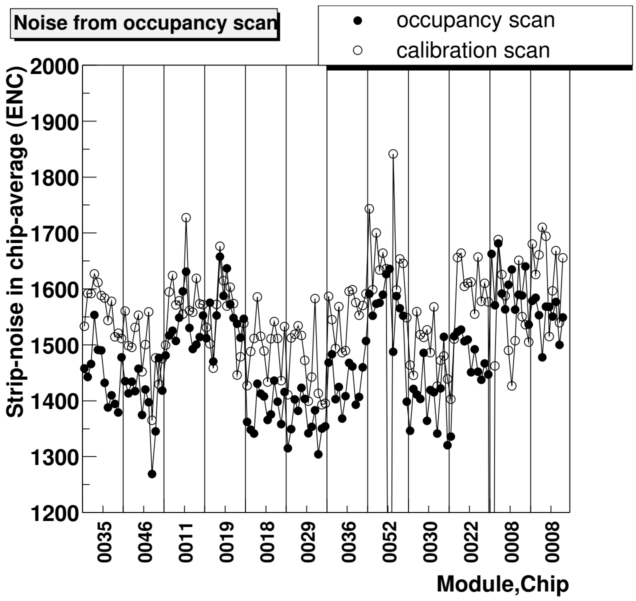

Figure 6: Noise as measured with a noise occupancy scan, comparing to the results from calibration scans. Both measurements use a row of twelve modules on the sector. Each slot on the graph corresponds to one module as labelled on the X-axis, in order as they are mounted on the sector, and each of the symbols represent the average noise per chip. The noise is quoted in units of the electron charge.

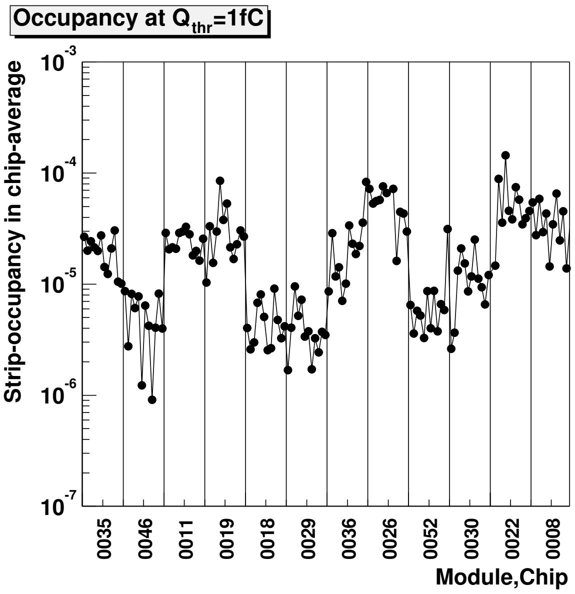

Figure 7: Noise occupancy measured at a threshold of 1 fC with a row of twelve modules on the sector, averaged over more than 100 individual measurements. Each slot on the graph corresponds to one module as labelled on the X-axis, in order as they are mounted on the sector, and each of the symbols represent the average occupancy per chip.

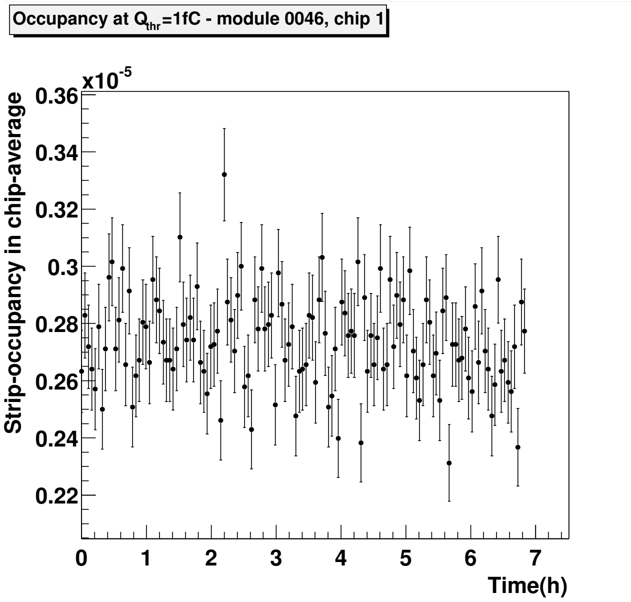

Figure 8: Noise occupancy repeatedly measured at a threshold of 1 fC with a row of twelve modules on the sector over several hours. Shown is the average occupancy for one chip on a selected module, plotted against time.

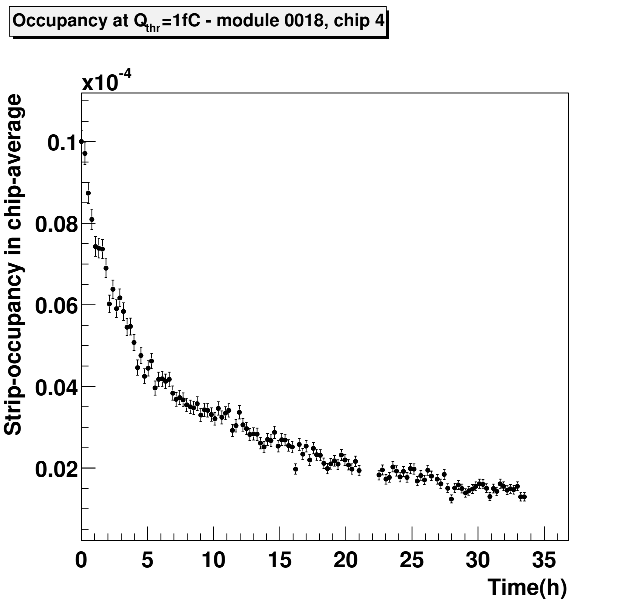

Figure 9: Similar measurement as shown in figure 8 *without* having had the modules biased for several hours before starting the measurements.

into which the particles are sent, thus not showing a significant amount of crosstalk. The narrow peak on the back side of module 0036 is due to an unmasked noisy channel.

5.2 Modules overlapping azimuthally

Although the module geometry used for the measurements described above deals with the most worrying overlap (across the electrical split between harnesses in the middle), different effects that only show up when modules overlap azimuthally are not checked there. To investigate such a scenario, all 16 unirradiated modules were mounted on the three left harnesses such that the nine best-performing modules formed a block of 3×3 in the middle of this configuration. Thus the module in the centre of this block overlaps with a module on each edge. The measurements run for the configurations of a row of twelve modules as discussed above were repeated in this configuration, using grounding and shielding scheme A (see section 2.6.2) since no difference to scheme B was observed in the measurements with a row of

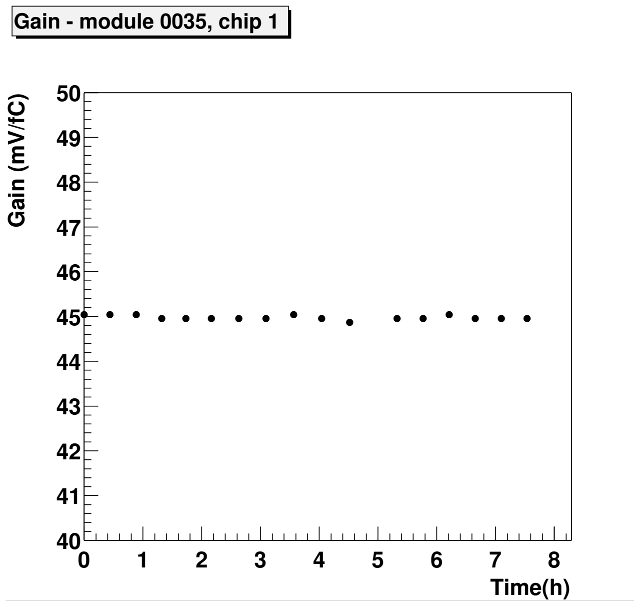

Figure 10: Gain repeatedly measured with calibration scans with a row of twelve modules on the sector over several hours. Shown is the gain of one chip on a selected module.

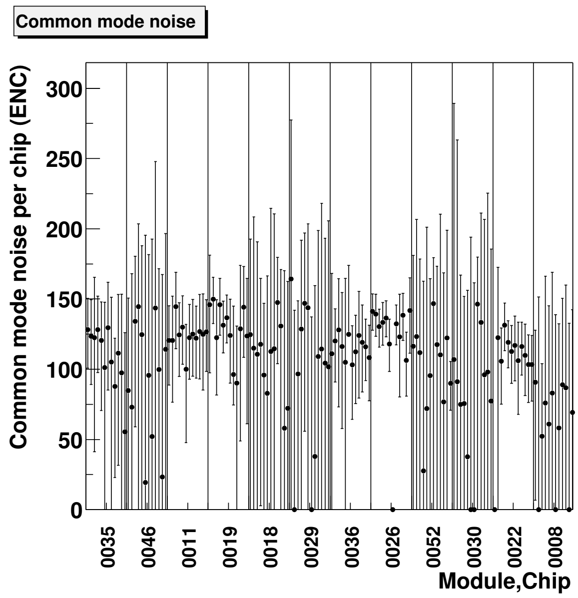

Figure 11: Results from a common mode noise measurement with twelve modules on the sector. The area between two vertical lines on the graph corresponds to one module as labelled on the X-axis, shown in the same order as they are mounted on the sector, and each of the symbols represents the amount of noise pick-up on a chip, quoted in units of the electron charge.

twelve modules. Again, no indication of significant noise pick-up is observed. As an example, the result from calibration scans, the average noise occupancy and the result from a common mode noise measurement are shown in figures 15, 16 and 17. The order of the modules is as follows (from left to right): front row: 0052, 0030, 0029; mid-row: 0046, 0035, 0036; rear row: 0018, 0011, 0019.

5.3 Shielding studies

To demonstrate the necessity of good shielding connections, the noise performance of the system was checked after having loosened the connection of the aluminium wrap around the end of the right harness to the main shield. Repeated measurements of the noise occupancy at a threshold of 1 fC (see section 3.4) are shown in figure 18. A periodic pick-up is observed that is related to the air conditioning in the laboratory switching on after the first three hours of data taking. Fixing the degraded shielding connection yields an occupancy vs time plot similar to figure 8.

As demonstrated in the two sections above, all measurements with the system undisturbed show no significant sign of noise pick-up. To get a

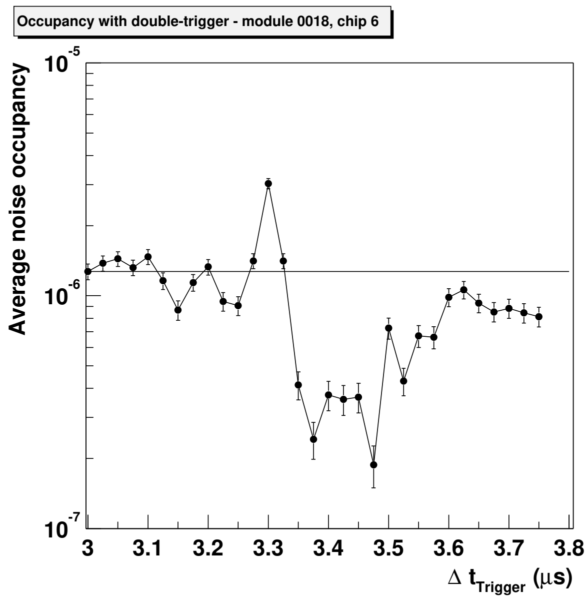

Figure 12: Noise occupancy measured with the second of a pair of successive triggers as a function of the time $\Delta t_{\mathrm{Trigger}}$ between the triggers. The line indicates the occupancy measured with a burst of single triggers.

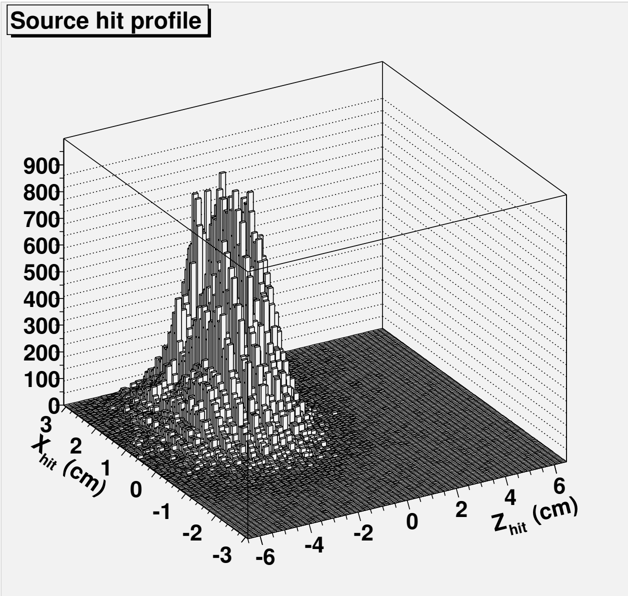

Figure 13: Profile of hits from a β^- source placed above the module shown here.

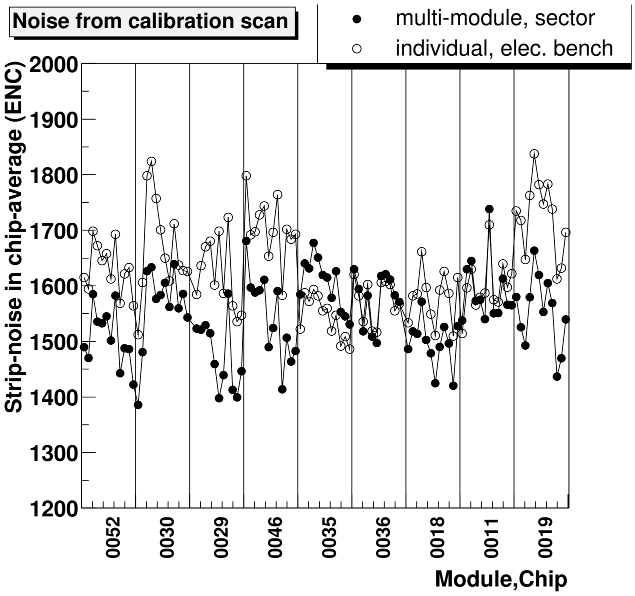

Figure 15: Noise as measured with calibration scans, comparing between a measurement of a block of 16 modules (only the central nine are shown) on the sector and individual measurement on the electrical-readout test bench. Each slot on the graph corresponds to one module as labelled on the X-axis, and each of the symbols represent the average noise per chip, quoted in units of the electron charge.

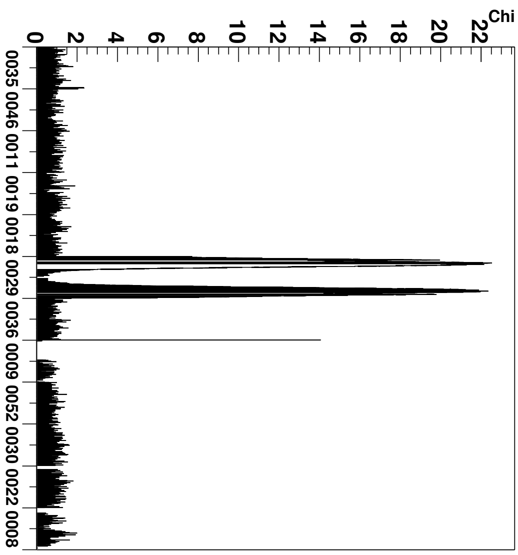

Figure 14: Probability of crosstalk between the module as labelled on the X-axis and module 0029 into which the signal was injected, represented by χ as defined in eq. (2).

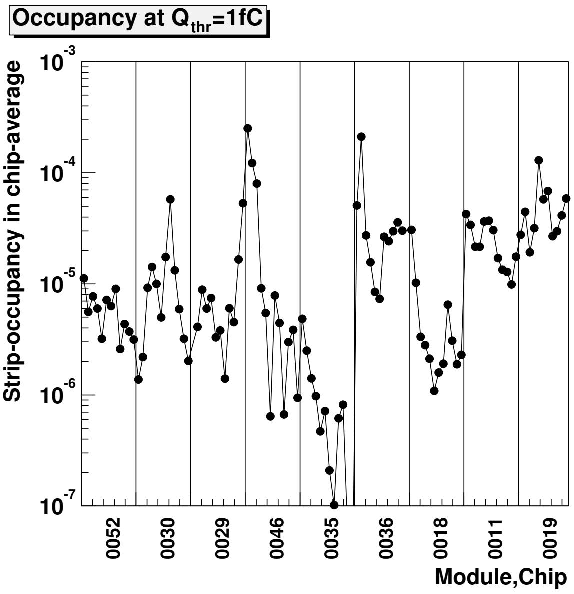

Figure 16: Noise occupancy measured at a threshold of 1 fC with a block of 16 modules (only the central nine are shown) on the sector, averaged over more than 100 individual measurements. Each slot on the graph corresponds to one module as labelled on the X-axis and each of the symbols represent the average occupancy per chip.

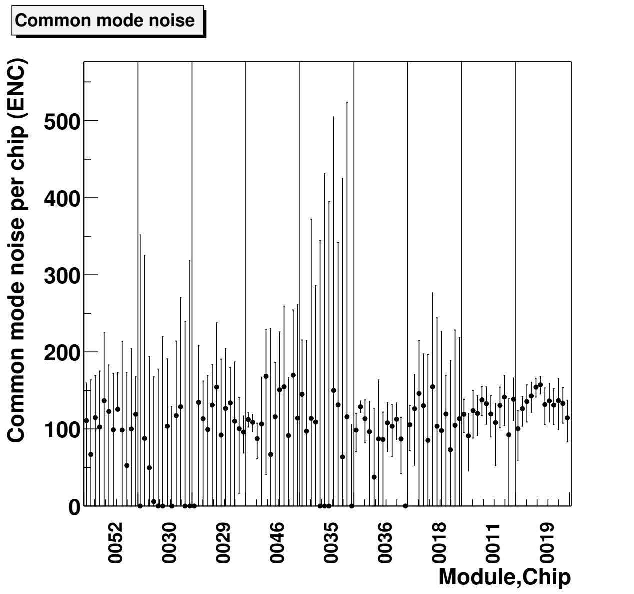

Figure 17: Results from a common mode noise measurement with a block of 16 modules on the sector (only the central nine are shown). Each slot on the graph corresponds to one module as labelled on the X-axis and each of the symbols represent the amount of noise pick-up on a chip.

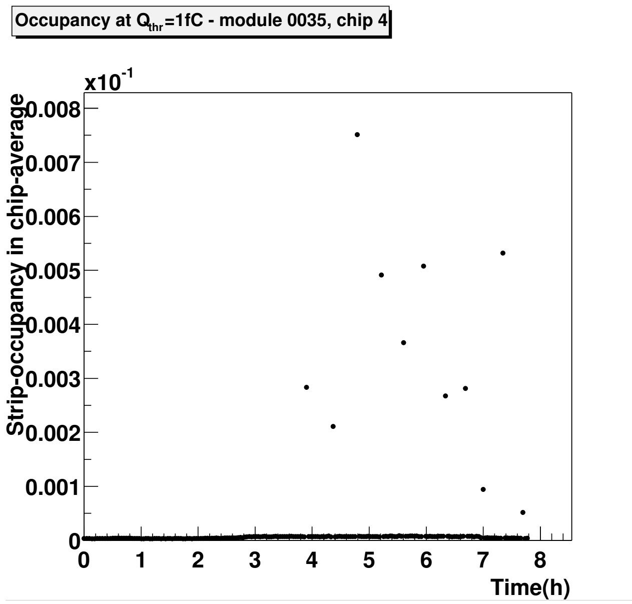

Figure 18: Noise occupancy repeatedly measured at a threshold of 1 fC with a row of twelve modules on the sector over several hours, with one of the shielding connections being deliberately poor (see text). Shown is the average occupancy per one chip on a selected module plotted against time.

clearer picture of the pattern of possible noise pick-up and to study various shielding options (see section 2.6), extra noise is forced into the system as described in section 3.8, using a setup with twelve modules in a row, similar to that described in section 5.1. Figure 19 shows the amount of pick-up as a function of the frequency of the injected current for one chip on one of the twelve modules in shielding scheme A, injecting with method 3 (see section 3.8). A clear peak between 10-20 MHz is seen, the exact behaviour varying slightly with module and chip. A similar frequency behaviour is observed for shielding scheme B so that the following discussion is restricted to injection at one frequency only. Injection modes 1 and 2 show a more complicated structure versus frequency, but the maximum is still located in the same frequency region. Figure 20 shows the location of noise pick-up for all three injection modes at 10 MHz, comparing shielding schemes A and B. For injection mode 1 scheme B performs slightly better, for mode 2 there is barely any difference between the two schemes. In mode 3, i.e. when injecting directly into the main shield, scheme A performs significantly better. Other shielding schemes have been tried, which still leave scheme A best for mode 3 and scheme B best for mode 1; no improvement for injection mode 2 has been achieved.

5.4 Cold running

The configuration of a row of twelve modules (see section 5.1) is run cold following the procedures outlined in section 2.7, requesting an ambient temperature of -10°C. This proved the method of cooling to sub-zero temperatures to be safe, an important check needed before operating irradiated modules (see section 5.5 below). In addition, the noise measured at temperatures approx. 20°C lower than under standard warm conditions allows a check of the temperature dependency of the performance of the system. Some of the modules show a change in gain measured with calibration scans (see section 3.2) of up to 11% and were excluded from the following analysis. The effect must be investigated further. Figure 21 shows the difference, warm to cold, in noise normalised to the difference in temperature measured on the hybrids of the modules. Most modules show a

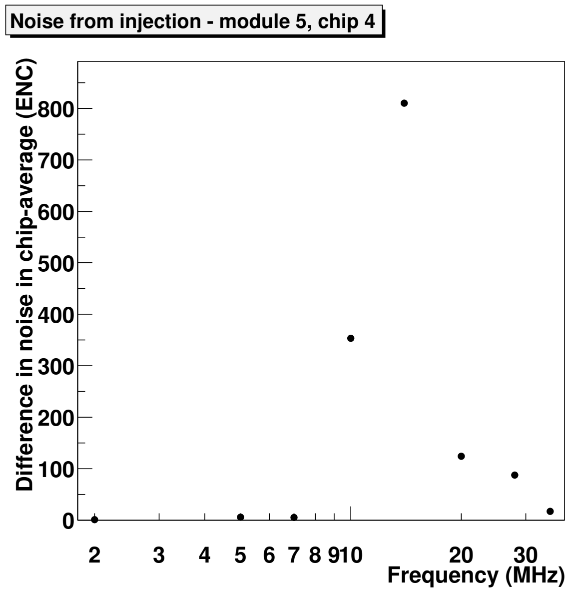

Figure 19: Noise pick-up as difference in noise quoted in units of the electron charge when comparing a measurement with and without a sinusoidal current being injected into the shield (mode 3 as explained in section 3.8). Shown is the pick-up as function of the frequency of the injected signal for one chip on one module in a configuration of twelve modules.

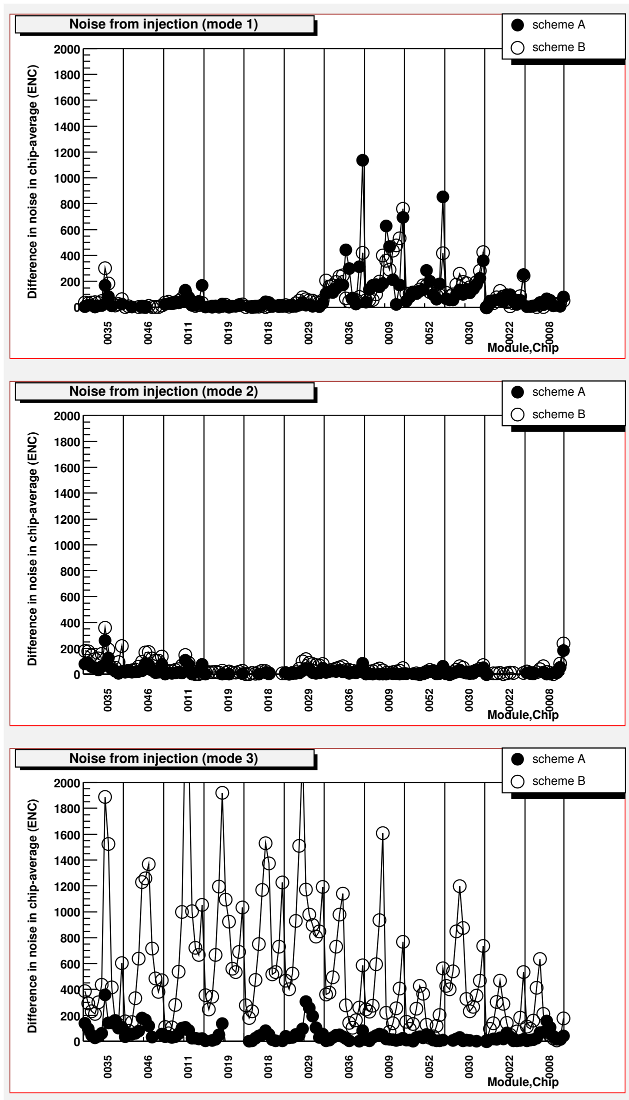

Figure 20: Noise pick-up as difference in noise when comparing a measurement with and without a sine current being injected into different parts of the system (modes as explained in section 3.8). Each slot on the graph corresponds to one module as labelled on the X-axis, in order as they are mounted on the sector in a configuration of twelve modules, and each of the symbols represent the pick-up quoted in units of the electron charge.

noise reduced by approx. 6 ENC/°C. The noise occupancy at a threshold of 1 fC is essentially stable when measured repeatedly (see section 3.4). The average occupancy is lower than that measured warm by a factor of two to three, compare figure 7 with 22.

5.5 Irradiated system test

The configuration of a row of twelve modules (see section 5.1) was modified to have the modules on the right harness replaced by six irradiated modules. The sector was operated cold following the procedures outlined in section 2.7, requesting an ambient temperature of -16°C.

Due to limitations in the cooling system, only three of the six irradiated modules could be fully biased with detector currents below 4 mA. Of the three modules that could not be fully biased, module 0020 could only be biased to 250 V but still showed reasonable performance. The other two modules (0007 and 0050) could only be biased to voltages slightly above 100 V; they were biased

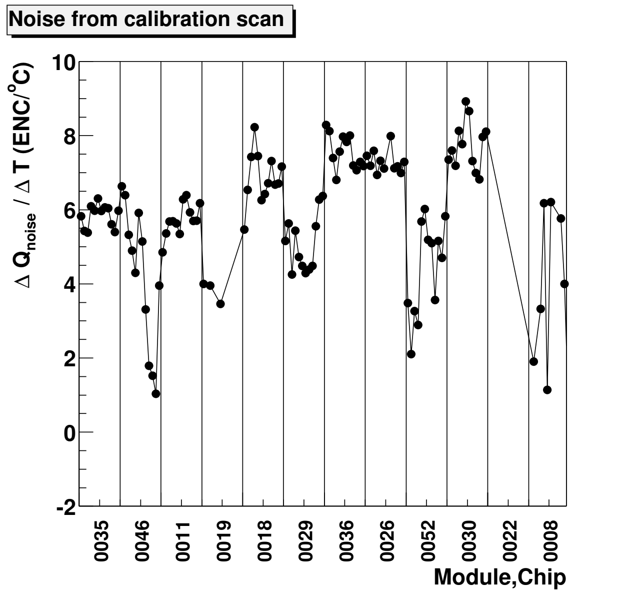

Figure 21: Results from a noise measurement with a row of twelve modules on the sector, comparing warm and cold operation. Shown is the difference in noise measured with calibration scans, normalised to the difference in temperature measured on the hybrids of the modules.

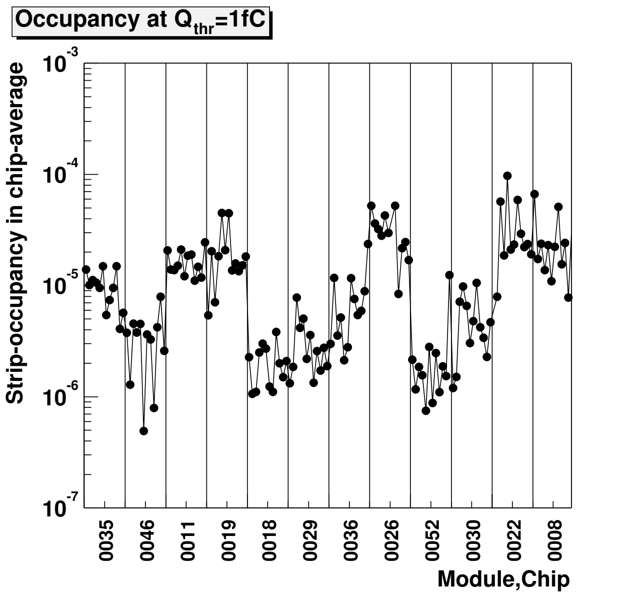

Figure 22: Noise occupancy measured at a threshold of 1 fC with a row of twelve modules on the sector operated cold, averaged over more than 100 individual measurements. Each slot on the graph corresponds to one module as labelled on the X-axis, in order as they are mounted on the sector, and each of the symbols represent the average occupancy per chip.

and read out but their performance will be ignored in the following discussion. Despite the cooling problems no indication of thermal runaway is observed. Figure 23 shows the bias current of module 0047 and the ambient temperature. On average, the former is stable, in detail it follows the variation in ambient temperature caused by the freezer switching on and off. A similar behaviour is observed in the occupancy at a threshold of 1 fC. Assuming that the bias current doubles when the detector temperature increases by 7°C, the above fluctuations in current correspond to a change in detector temperature of 0.6°C, much less than the 1.8°C variation in ambient temperature.

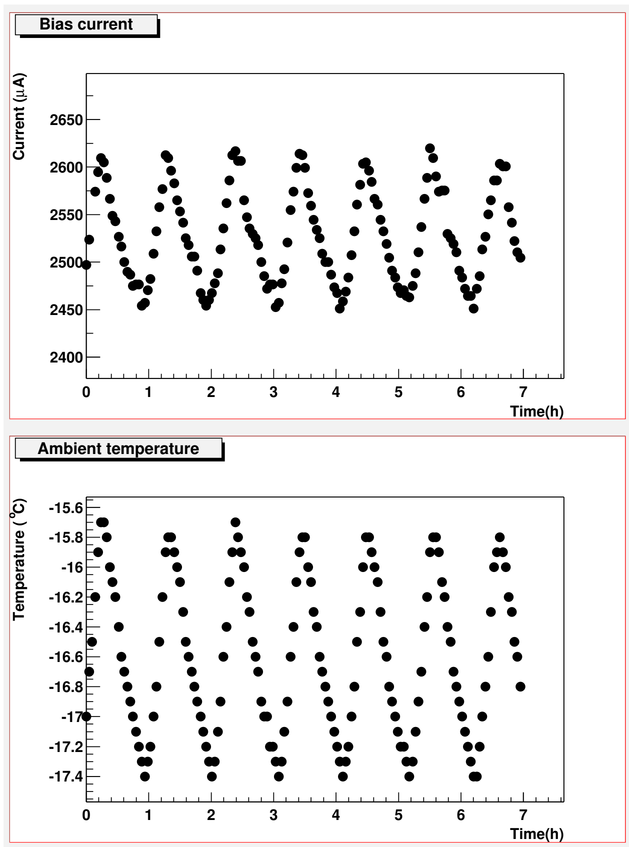

Figure 23: Bias current (top) and ambient temperature (bottom) against time during repeated noise occupancy measurements for the irradiated module 0047 in a row of six unirradiated and six irradiated modules.

The noise resulting from calibration scans is shown in figure 24. The values are corrected for changes in gain observed in the irradiated modules, similar to those mentioned in section 5.4 above. The reference values from similar measurements on the electrical-readout test bench are higher for the unirradiated modules since those measurements were taken at temperatures approx. 20°C higher than the sector measurements. Similarly, the irradiated modules were tested individually at temperatures approx. 20°C *lower* than on the sector, resulting in the reference noise figures being lower than the results from the sector measurement.

The noise occupancy at 1 fC threshold is found to be high for the irradiated modules, see figure 25, which was already observed during the individual tests of those modules. The unirradiated modules are not strongly affected. It should be noted that a common mode noise measurement as per section 3.5 yields pick-up compatible with zero for the irradiated modules.

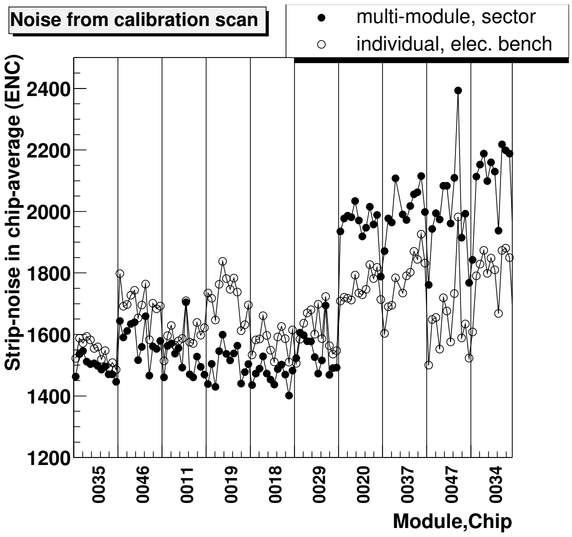

Figure 24: Noise as measured with calibration scans, comparing between a measurement of a row of twelve modules on the sector and individual measurement on the electrical-readout test bench. Each slot on the graph corresponds to one module as labelled on the X-axis, in order as they are mounted on the sector, the right six having been irradiated (only four of them are shown, see text). Each of the symbols represent the average noise per chip. The noise is quoted in units of the electron charge.

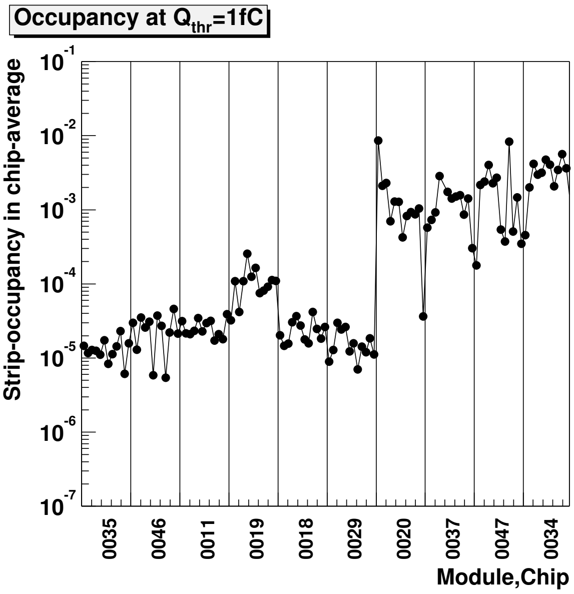

Figure 25: Noise occupancy measured at a threshold of 1 fC with a row of twelve modules on the sector, the right six (only four of them are shown, see text) having been irradiated, averaged over more than 100 individual measurements. Each slot on the graph corresponds to one module as labelled on the X-axis, in order as they are mounted on the sector, and each of the symbols represent the average occupancy per chip.

5.6 Redundancy

As described in section 2.2, the clock and command signals for each module can be obtained electrically from a neighbouring dogleg as a redundancy option. This means that high-frequency signals are transported close to the detectors, causing concern about the detectors picking up such a signal. This mode is thus tested with two left harnesses, completely connected to form a "ring" of redundancy links. They are equipped with six modules each. Two sets of calibration scans with modules getting normal or redundancy command[3], respectively, yield the module noise. Both operation modes show similar noise, see figure 26, indicating negligible pick-up from the redundancy tapes.

5.7 Power supply and cable tests

As already mentioned in section 2.2, the cabling from PPB2 to the power supplies in the system test is not very close to the ATLAS design. Special tests with the 30 m cable normally used between PPB3 and the power supplies ("type 3") connected between PPB2 and PPB3 and 60 to 120 m of the design-type cable ("type 4") run-

[3]It should be noted that due to technical problems only nine out of twelve modules could be operated in redundancy mode, leaving k3112, k3104 and scand1 in normal mode for both measurements.

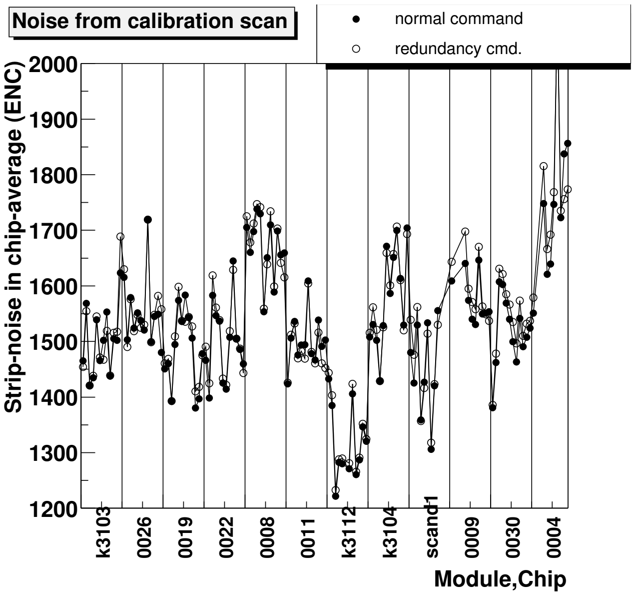

Figure 26: Noise as measured with calibration scans with two half-rows of six modules each, comparing between a measurement sending the clock and command signals to the modules in the normal way and one using the redundancy links. Each slot on the graph corresponds to one module as labelled on the X-axis, in order as they are mounted on the sector (the left half corresponding to the front row, the right to the rear row), and each of the symbols represent the average noise per chip.

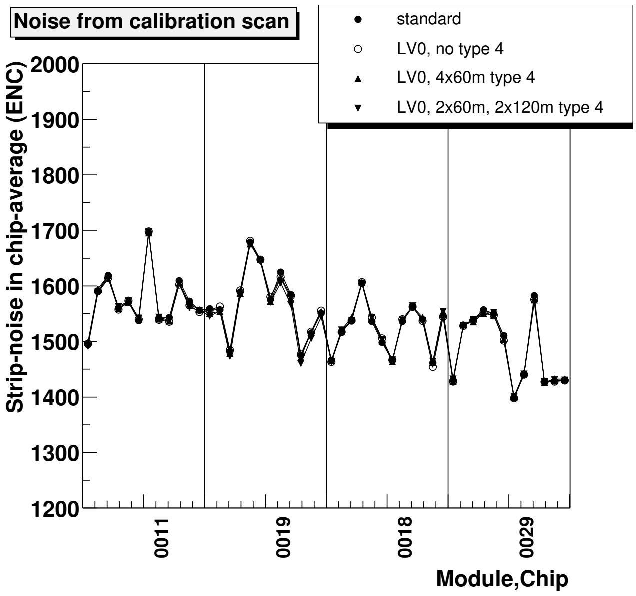

Figure 27: Noise as measured with calibration scans with four modules on one harness, comparing between setups using different cable between PPB2, PPB3 and a new low-voltage supply prototype (see text). Each slot on the graph corresponds to one module as labelled on the X-axis, and each of the symbols represent the average noise per chip.

ning between PPB3 and the power supplies were performed. At the same time, the tape between PPB1 and PPB2 was replaced by conventional cable ("type 2"). These tests were run with a newer version of the low-voltage supplies ("LV0") which is close to the final version. Figure 27 shows the noise measured with calibration scans for four modules, comparing to measurements with the standard setup. No difference is observed.

6 Summary

Up to 16 modules have been operated in two configurations on a barrel SCT system test. Noise performance close to that measured with individual module tests and well within specifications is observed. The amount of common-mode noise pick-up is measured to be well below the intrinsic module noise.

The necessity of good shielding has been demonstrated. Differences in noise pick-up are observed between two proposed grounding and shielding configurations when injecting extra noise into the system, the conclusion depending on the details of the injection. More information on this subject is expected from a revised system test with a grounding and shielding system much closer to that of the SCT in its parts near the sector. This second system test will as well have all on-sector components according to the latest design so that possible changes in the noise performance can be checked.

Some open questions, like the drop in occupancy observed with two triggers succeeding each other within about 3.5 μs or the temperature dependence of the (measured) gain of the amplifiers of the module readout, can be studied further on a single-module basis.

Acknowledgements

The efforts of R. Fortin and B. Herskind as well as G. Barbier and E. Perrin during the production of components for the system test and their maintenance are highly appreciated. The helpful and competent assistance of B. Gallop with all our software and DAQ-related problems was invaluable. We acknowledge the advice of N. Spencer and T. Smith in setting up and improving our grounding and shielding schemes. Last not least, the system test could not have been run without the generous donations of equipment from the various SCT institutes and of modules from the production sites.

References

[1] ATLAS Technical Proposal for a General-Purpose pp Experiment at the LHC at CERN, CERN/LHCC/94-43 (1994).

[2] ATLAS Collaboration, ATLAS Technical Design Report - Inner Detector, CERN/LHCC/97-16 and CERN/LHCC/97-17 (1997).

[3] M. Turala for the ATLAS SCT collaboration, Nuclear Instr. Meth. A 466 (2001) 243.

[4] Y. Unno for the ATLAS SCT collaboration Nuclear Instr. Meth. A 453 (2000) 109.

[5] T. Kondo for the ATLAS SCT collaboration, Nuclear Instr. Meth. A 485 (2002) 27.

[6] H. Pernegger for the ATLAS SCT collaboration, VERTEX2001, Brunnen, Switzerland (2001).

[7] Y. Unno for the ATLAS SCT collaboration, 2001 IEEE NSS/MIC, San Diego, USA (2001).

[8] M. Vos for the ATLAS SCT collaboration, Nuclear Instr. Meth. A 485 (2002) 67.

[9] A. Barr et al., ATLAS note ATL-INDET-2002-005 (2002).

[10] R.S. Harper et al., Nuclear Instr. Meth. A 479 (2002) 548.

[11] P. Dervan for the ATLAS SCT collaboration, 4th International Conference on Radiation Effects on Semiconductor Materials Detectors and Devices, Firenze, Italy (2002).

[12] G. Barbier et al., SCT barrel technical drawings, available at:
http://dpnc.unige.ch/atlas/atlaspage/mecan-act/webmec1.html

[13] T. Weidberg, ATLAS-EDMS technical note ATL-IS-ES-0049 (2001).

[14] D.G. Charlton et al., Nuclear Instr. Meth. A 443 (2000) 430.

[15] M. Morrissey, M. Postranecky, M. Warren, et al., documentation of the Slog, Cloac, Mustard DAQ system, available at:
http://hepwww.rl.ac.uk/atlas-sct/mm/Slog/slog.pdf
http://www.hep.ucl.ac.uk/atlas/sct/cloac/Welcome.html
http://hepwww.rl.ac.uk/atlas-sct/mm/Mustard/Mustard.ps

[16] P.W. Phillips, *Operation of Barrel Modules*, to be submitted as ATLAS-EDMS technical note.

[17] J. Bohm et al., Power Supply and Power Distribution System for the ATLAS Silicon Strip Detectors, CERN-2001-005 (2001).

[18] N. Spencer, ATLAS-EDMS technical note ATL-IC-EN-0004 (1999).

[19] M. Morrissey and T. Weidberg, ATLAS-EDMS technical note ATL-IS-ES-0056 (2001).

[20] D. Greenfield and T. Smith, ATLAS-EDMS technical note ATL-IS-ER-0019 (2001).

[21] L. Feld, S. Roe, A. Ahmad, and S. Snow, Nuclear Instr. Meth. A 487 (2002) 557.

[22] P.J. Bell and J. Grosse-Knetter, *Theory and measurement of common mode noise in binary read out systems*, to be submitted as ATLAS note.

[23] The SCT system test WWW page:
http://home.cern.ch/asct186/systemtest.html

[24] A. Chilingarov, *Time dependence of the interstrip capacitance*, talk at the SCT week June 2002.