The Data Acquisition and Calibration System for the ATLAS Semiconductor Tracker

Tom Barber, University of Cambridge, on behalf of the ATLAS SCT Collaboration.

Abstract—The Semiconductor Tracker (SCT) data acquisition (DAQ) system has been built to calibrate, configure, and control the approximately six million front-end channels of the ATLAS silicon strip detector.

It provides a synchronized bunch crossing clock to the front-end modules, communicates first-level triggers to the front-end chips, and transfers information about hit strips to the ATLAS high-level trigger system.

The system has been used extensively for calibration and quality assurance during SCT barrel and end-cap assembly and for performance confirmation tests after transport of the barrels and end-caps to CERN.

Operating in data-taking mode, the DAQ has recorded millions of synchronously-triggered events during commissioning tests including cosmic ray triggered events both in standalone and ATLAS-wide mode.

In this paper we describe the components of the data acquisition system, discuss its operation in calibration and data-taking modes and present some detector performance results from these tests.

Index Terms—Silicon detectors, Data Acquisition

I. INTRODUCTION

THE ATLAS experiment [1] is a general-purpose detector at the Larger Hadron Collider (LHC), a proton-proton collider with a centre of mass energy of 14 TeV and a design luminosity of 10^{34} cm^{-2}s^{-1}.

High precision spacepoint measurement is essential in ATLAS for track reconstruction and particle momentum measurements. In ATLAS this task is performed by the triumvirate of the Inner Detector: the Pixels, the Semiconductor Tracker (SCT) and the Transition Radiation Tracker (TRT). The Pixels lie closest to the interaction point and allow secondary vertex reconstruction. The TRT is furthest from the primary vertex and contributes around 36 track measurements based on straw tubes. The SCT is a silicon strip tracking detector housed between the Pixels and TRT. The entire Inner Detector is enclosed inside a 2 T solenoidal magnetic field. Figure 1 shows how the three subsystems are combined to make up the Inner Detector.

II. DETECTOR LAYOUT

The SCT provides an average of four high-precision spacepoints for ATLAS tracking, within a pseudorapidity of $|\eta| < 2.5$. The central region, covering $|\eta| < 1.4$ is instrumented by the barrel subsystem, consisting of four concentric cylinders, with radii from 299 mm to 514 mm. Mounted on the barrel are a total of 2112 identical modules [2].

The regions either side of the barrel, down to $|\eta| < 2.5$ are covered by two end-cap subsystems, each of which consists of
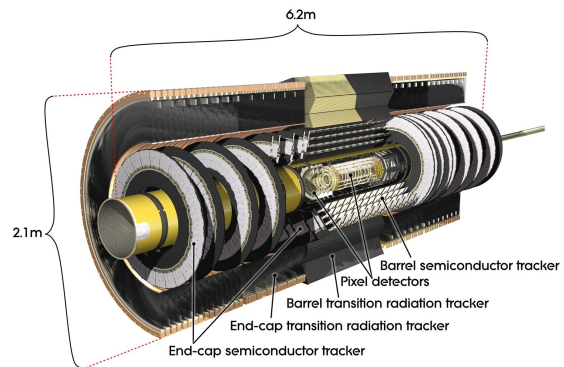

Fig. 1. A schematic of the ATLAS Inner Detector, including the barrel and two end-caps that make up the SCT.
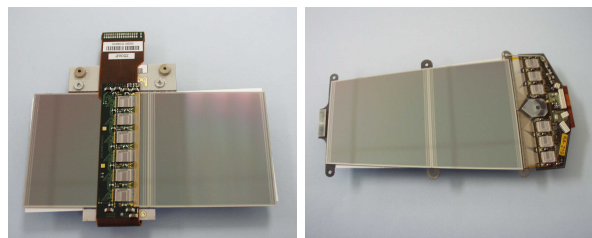

Fig. 2. An SCT barrel (left) and end-cap module (right). The barrel modules are all identical, whereas the end-cap modules come in three different types: inner, middle (shown here) and outer. The six ABCD chips are clearly visible mounted on the hybrid.

nine discs. Each end-cap is constructed from 988 modules of three different designs [3]. Photographs of a barrel and end-cap module are shown in figure 2.

In total the SCT comprises 4088 modules, with 61m^2 of silicon sensors and 6.2 million read-out channels.

III. THE SCT DATA ACQUISITION SYSTEM

The SCT DAQ system [4] is comprised of the hardware component of the silicon detector modules, optical links and off-detector hardware and the DAQ software package which controls and configures these components for data-taking.

A. Hardware

The hardware required to control, configure and read out information from the SCT can be broadly divided into two parts: The sensors and front-end electronics that make up the on-detector components, and the off-detector read-out system that resides in the counting rooms. These two sections are connected by infra-red optical links.
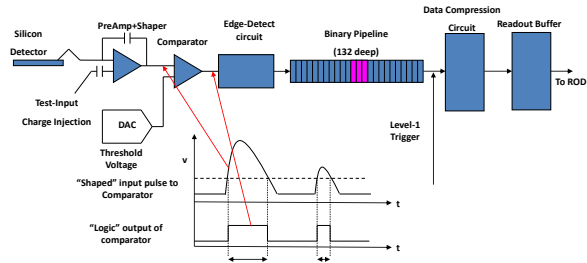
Fig. 3. The components of the ABCD chip.

1) SCT Modules: The SCT module is the basic unit of the detector and is comprised of the silicon strips, front-end read-out chips and support structure.

a) Silicon Sensors: Each module has two planes of silicon with 768 strips of p^+ implant on n-type bulk, which are offset by a small stereo angle of 40 mrad to provide space-point resolution. There is a nominal bias voltage of 150 V applied to the strips, a value that may be amended in the future. The module is supported by a baseboard which provides structural stability as well as dissipating heat generated by the module.

b) ABCD Chips: The implant strips are read out by the ABCD chip, which are radiation hard DMILL technology application-specific integrated circuits (ASICs) each with 128 channels. The chips are mounted on a flex circuit known as the hybrid mounted on the module. A schematic of the ABCD chip is shown in figure 3.

The charge generated in each strip on the sensor is first amplified, shaped and then forms the input to a discriminator. The discriminator output is a binary signal that is passed through a mask register and stored in a pipeline buffer each time a clock signal is received by the chip.

An optional edge-detect circuit can be applied which ensures that each detector hit only generates one output pulse. On receipt of a Level 1 Accept (L1A), the result at the end of the pipeline is moved into an output buffer and compressed before being sent to the readout.

To aid synchronisation throughout the detector, and to diagnose different hit patterns, hits are read-out in three different 25 ns time bins, centred on the L1A trigger.

For calibration purposes, known charges can be injected into every fourth channel by applying voltages across a calibration capacitor.

2) Optical Links: Each SCT module is connected to the off-detector read-out system by three infra-red optical links [5].

A module receives information over a single fibre, with the clock and command signals combined using a bi-phase mark (BPM) encoding. A PIN diode converts the incoming data from optical into electrical signals, which are then decoded by the DORIC (Digital Optical Receiver IC) and distributed to all chips on the module. A second set of clock and command signals is also generated which can be sent to the adjacent module on the detector, providing redundancy in the system should a link fail.

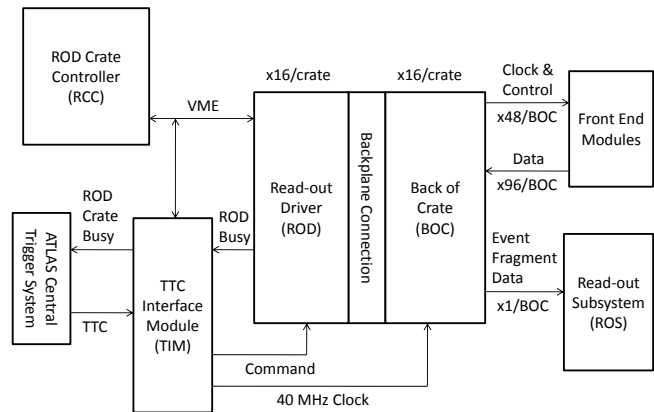Two designated master ABCD chips on the module are

Fig. 4. Off-detector electronics of the SCT.

each connected to a VCSEL (Vertical Cavity Surface Emitting Laser), driven by the VDC (VCSEL Driver Chip), allowing transmission of hit data off-detector via two optical links. Under nominal conditions, the hit data from six ABCD chips are transmitted per link, corresponding to one side of a module. Should one of the links fail, the modules can be configured to read out information for all chips through a single link.

The entire optical package containing the PIN, DORIC, VCSEL and VDC and fibre connections are housed on the detector adjacent to the module.

3) Off-Detector Read-out Hardware: The off-detector DAQ system is installed in the ATLAS large service cavern, around 100m from the SCT itself, which resides in the experimental hall. The off-detector hardware consists of eight 9-U VME crates and nine rack mounted servers. The crates contain various VME modules that communicate via a custom backplane. An overview of the most important components in the system is shown in figure 4.

An SCT crate will typcially contain a single Timing Interface Module (TIM), a number of Readout Driver (ROD) and Back-of-Crate (BOC) pairs and a single ROD Crate Controller (RCC).

a) ROD Crate Controller (RCC): The RCC is a commercial 6U Single Board Computer (SBC) running Linux, which acts as the VME host, and as a result occupies the first slot in the crate. Dedicated software runs on the RCC providing an overall interface for control, configuration and communication of all the hardware present in the crate.

During data-taking physics mode runs, the RCC is used to configure modules, after which the data flow is handled solely by the FPGAs on the ROD. After a run has started, the SCT DAQ software takes on the role of monitoring.

b) Timing Interface Module (TIM): There is a single TIM in each crate, which is responsible for distributing the trigger, timing and control (TTC) signals from the the ATLAS Central Trigger Processor (CTP) system to every ROD over the backplane. For every L1A trigger sent, the TIM also distributes level-1 trigger and bunch clock counters (L1ID and BCID respectively) to assist event synchronisation. If a ROD cannot

keep up with the read-out rate, it will assert a BUSY signal that is sent to the TIM. This signal is propagated back to the ATLAS trigger system, halting triggers until the BUSY is cleared. The TIM can also generate a BUSY internally to veto fixed frequency triggers, which may induce harmful wire-bond resonances [6].

c) Back-of-Crate Card (BOC): The BOC is responsible for transmitting commands and data between the optical fibre connections and the ROD with which it is paired.

Each command designated for the front-end modules is routed via a TX plug-in, which converts the clock and command into a single BPM signal. This electrical signal is converted to optical form by a 12-way VCSEL array before being transmitted to the modules along a 12-way fibre ribbon. The intensity of the optical signal can be configured to individual fibre connections, using a digital-to-analogue converter (DAC) on the BOC. Timing of the outgoing signals can be adjusted to ensure that clock received by the modules has the correct phase relative to the particles from LHC collisions. This is set on a module-by-module basis to allow for differences in fibre lengths and time-of-flight variations from different module locations.

Incoming data from modules are received in optical form from a total of 96 input links, and converted into electrical signals by eight 12-way PIN diode arrays. The signals are then discriminated and sampled at a variable phase to ensure reliable reconstruction of the binary data stream. The electrical data are finally forwarded to the ROD module that is paired with the BOC for event processing.

In addition, the BOC is also responsible for transmitting formatted event fragments from the ROD to the first level of the ATLAS high-level trigger system, the Read-out Subsystem (ROS). The ROD generates a single, formatted data stream which is forwarded to the ROS via an S-Link connection.

The 40 MHz ATLAS clock is normally distributed directly to the front-end modules from the TIM via the BOC. However, in the absence of this clock, a phase-locked loop on the BOC can generate a local replacement. This ensures that the modules always received a clock signal, without which they generate less heat which could effect detector alignment due to thermal distortions.

d) Readout Driver (ROD): One basic function of the ROD [7] is to transmit the control commands and configuration data through the BOC to the front-end modules. These data can be Level 1 triggers, Event Counter or Bunch Crossing resets, calibration commands or module register data. The second function of the ROD is to receive and pass on data streams from the front-end modules, which can either be event data or module register data. The ROD boards themselves are a hybrid architecture of Field-Programmable Gate-Arrays (FPGAs), which implement the data path, and Digital Signal Processors (DSPs) for control and calibration. Each ROD can control and process data from 48 SCT modules, with up to 16 RODs per crate.

The Master DSP (MDSP) oversees the operation of the entire ROD, although it does not explicitly take part in data taking during physics runs. It provides access to the FPGAs and can run tasks such as histogram control and link masking,
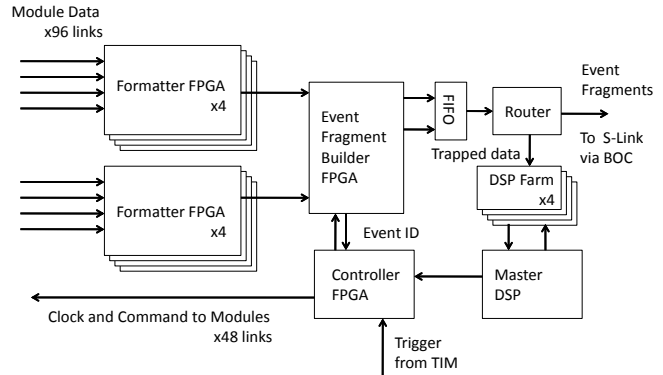

Fig. 5. SCT ROD data path implemented by a hybrid architecture of FPGA and DSP components.

communicating between the host via transfer of so-called primitives.

The ROD controller FPGA (RCF) coordinates all of the control path operations required for data-taking, module calibration and on-board diagnostics, provides connections from the other FPGAs, slave DSPs and BOC to the MDSP and interfaces with the TIM for access to clock and trigger data. In normal data-taking one of the key functions of the RCF is to distribute clock and triggers to the modules via the BOC.

The remaining FPGA components are used to implement the data path on the ROD, as shown in figure 5. The formatters receive serial data from the module input links, convert them into a 32 bit data word format and provide derandomising buffering of the event fragments for each link in parallel. The formatters also detect module errors and can send a ROD BUSY to the TIM if one of the buffers is close to maximum occupancy. There are eight formatter chips on each ROD, each capable of reading out data from 12 links. The formatters are arranged into two banks of four devices, an architecture that allows the processing bandwidth to approach 80 MHz.

The next stage in the data path is the Event Fragment Builder (EFB) FPGA, which is responsible for creating ATLAS standard event fragments [8]. The EFB contains two processing engines that can each collect the output from a bank of four formatters. The data stream is also monitored and flagged for errors, including checks of L1ID and BCID synchronisation. As the event fragment is constructed, it is stored in a FIFO until a complete fragment is ready to be sent on to the router.

The router FPGA is the final stage of data processing on the ROD, with the primary function of transmitting event fragments to the S-Link, via the BOC. As data flow through the router, error flags are added to each link header, and link information replaces the L1ID and BCID counts. The router also contains four event traps which can be configured to siphon events at variable frequency and, if desired, filter them for specific trigger source (either ATLAS, TIM or ROD). Each trap is connected independently to a slave DSP (SDSP),

where event fragments can be transferred for counting and monitoring.

The slave DSPs have a similar program structure to the MDSP, sharing the same primitive and task functionality. During normal operation, the master DSP is continuously communicating with the four slaves, and so direct access to them is not possible from the host. Instead, primitives intended for the slaves are sent using the master DSP as a middleman.

The main function of the slaves is to histogram event data from the router during calibration scans. The DSP retrieves data frames of 256 words using a direct memory access (DMA) transfer with an interrupt service routine (ISR) on completion. Once a complete event has arrived, the DSP ISR places it onto an event queue. On each pass through the DSP polling loop, the next event in the queue is transferred for processing by dedicated histogramming tasks, after which the event is removed and memory released. These tasks can either be in the C programming language, or in the faster assembly language, and can be configured to process hits from all three time bins together, or in separate histograms. At the end of a calibration scan the histogram data can then be read out to the host through the MDSP using primitives.

The ROD and TIM module designs are common to both the SCT and Pixels, with differing firmware tailored to suit the needs of the two sub-detectors.

B. Software

In addition to software running on the RCC, other distributed calibration and analysis tasks are run on the rack-mounted servers, written in a combination of C++, Java and Python. These machines are connected to each other and the user in the ATLAS control room via ethernet.

The DAQ software is written using the ATLAS Trigger and DAQ (TDAQ) framework, which allows it to be smoothly integrated into central ATLAS data-taking. In particular, the SCT DAQ makes use of the inter-process communication (IPC) tools provided by ATLAS online software.

To run the SCT detector at maximum efficiency, the optimum values must be found for various crate and module parameters, such as optical delay and threshold, and response of silicon strips to injected charge. Due to a number of factors, such as variations in fibre lengths and manufacturing tolerance of chips, optimisation needs to be done on a module-by-module basis. Channel-by-channel calibration may also be necessary due to degradation in detector response caused by irradiation.

The customary procedure to find the optimum value of some parameter is to perform a scan over all possible settings of that parameter, using algorithms to fit and analyse the results. In some cases, a test is performed which involves running a series of scans in succession.

During an SCT calibration run, the SCT DAQ package is used to control and implement the starting of scans. An overview of the SCT DAQ software is shown in figure 6.

The SctApi crate server, running on each of the eight crates, loads the individual crate and module settings from the Configuration Service at the start of a run and applies them to
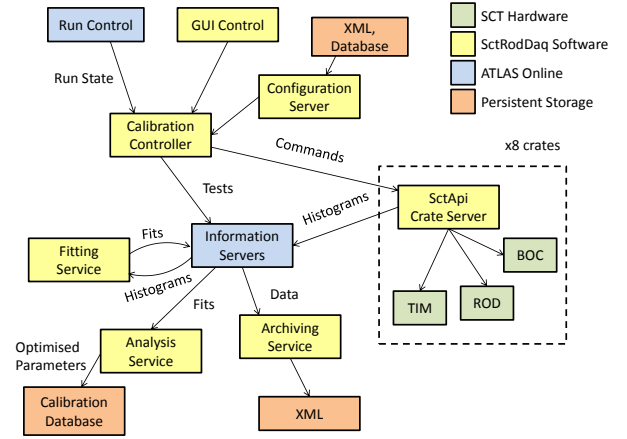

Fig. 6.   The structure of the SCT DAQ software.

the relevant hardware over VME. During calibration, the crate servers control setting up and varying scan parameters and collecting data into occupancy histograms. For each setting of a scan parameter, a sequence of internally generated triggers is sent to each module.

The resulting hit data are trapped on the router, which are then histogrammed on the SDSP and finally read out by the host at the end of the scan.

The Calibration Controller is responsible for overseeing the process of setting up tests, starting scans on the crate servers, and choosing the relevant fitting and analysis algorithms.

Once a test is completed on the crate servers, the histogrammed results are published to an Information Server (IS). The Fitting Service contains a listener thread which takes these histograms and adds them to a queue. Worker threads then undertake the processor-intensive job of performing fits, while the listener threads are free to respond to further data. This listener and worker implementation allows the computing to be done in close to real-time.

In a similar way, fit data published by the fitting service are picked up by the Analysis Service which then extracts the optimum operating parameters from the scan, and decides if the test has been successful. The Analysis Service may also run on raw histograms in the case where a fitting procedure is not needed.

Finally, the Calibration Controller monitors the test results when they are available and updates the configuration with the updated optimum values.

IV. CALIBRATION OF THE SCT

A number of different tests are necessary to fully calibrate the SCT detector. These tests fall into the three main categories of optical tuning, digital tests, and analogue characterisation.

A. Optical Tuning

To ensure reliable communication between off-detector electronics and the front-end modules, it is essential to ensure that the optical links are well configured.

The most significant variables in the optical tuning are the current supplied to the VCSEL chips on the BOC (Tx current),

the threshold at which a received signal is discriminated, and the phase that the received signal is sampled (Rx threshold and delay respectively). As the Rx threshold and delay are correlated, an important optical scan varies the two parameters simultaneously to find a region of parameter space where reliable communication occurs.

For the optical tests, the modules are put into 'clock/2' mode, in which they return a clock signal which is half the frequency of the 40 MHz input signal. For each setting of threshold and delay, counters on the MDSP record the number of 1's within a set time window. As no triggers are sent, this is one of the fastest tests performed.

B. Digital Tests

Once communication is established, the next stage is to confirm that the ABCD chips are functioning correctly. This is done with a variety of tests which exercise the channel mask registers (NMask), trigger and bunch crossing counters pipeline cells and chip token-passing logic.

C. Analogue Calibration

An essential part of calibration is to characterise the detector's response to injected charge. Each front-end ABCD chip has an 8-bit DAC which sets the threshold globally across that chip. Channel-by-channel variations can be compensated for by using a 4-bit DAC known as the TrimDAC.

By injecting various known charges and measuring the occupancy at different thresholds, the analogue properties of each channel can be determined. For each value of charge injected the threshold is scanned over the entire range and a complementary error function is fitted to the average occupancy. The threshold at which occupancy is 50% corresponds to the median of the injected charge, with the sigma giving the noise. By repeating this scan at different injected charges, the per-chip response can be built up. A full response curve contains data points from ten different charges. The values used to fit the response curve are stored in the module configuration and can be used to set the threshold at an arbitrary charge. For nominal conditions, the threshold is set to a value of 1.0 fC.

Similar scans are used to optimise the TrimDAC thresholds, and threshold scans without injected charge are used to find the noise occupancy of the detector.

V. DETECTOR PERFORMANCE RESULTS

In 2006, the SCT barrel and end-caps were installed in their final position in the ATLAS experimental hall at point 1. Since then, efforts have focused on calibrating and preparing the detector for data taking.

A. DAQ Experience

During SCT commissioning, a number of developments have been made to the DAQ allowing it to evolve into a faster and more reliable system.

1) Errors during scans: During a calibration scan, errors in the data stream may occur due to badly tuned optical links or noisy module strips. When calibrating large numbers of modules, it is important that scans are robust against these errors.

Events which contain errors are not immediately histogrammed by the SDSP, but sent to a separate routine that keeps track of which links might be causing the error. If an event is marked as containing errors it is picked up by the histogram control task running on the master, which attempts to reset modules before resending triggers. If there are still error events being sent after a set number of resets, the error handling routine will try to switch off links until an error free event is received. This has allowed calibration scans to run much more reliably in the presence of large numbers of modules.

2) ABCD Simulation: A simulation block has recently been added to the formatter, which replaces the input data from the BOC with internally generated events. By injecting the ROD with data at the earliest possible stage, it has allowed the full readout chain to be exercised in both physics and calibration modes when real modules were unavailable.

The size and type of events generated has been made fully configurable using a spare register on the formatter. Events can vary in size from one hit per module to full occupancy of 128 hits in each chip. Occupancy can be further varied by changing the fraction of empty events that occur and events can be configured to generate hits in different time bins. In addition to generating hits, module chip errors can be inserted into the data stream, allowing the flagging and monitoring of errors to be tested upstream. Finally, bit-flips can also be introduced, simulating optical noise in the system. Pseudo-random numbers are generated using a 19-bit linear feedback shift-register (LFSR). Constraints on formatter size mean that there is only space for one simulation engine in each formatter, so an identical data stream is produced for each pair of links.

Allowing the SCT ROD to produce events with specific size and error properties has been very useful in testing the robustness of the system to large, noisy and erroneous events.

3) Masking Noisy Links: During physics mode data taking, the high data volume produced by noisy modules, chips and strips can cause the ROD to exert a BUSY signal which propagates to the ATLAS central DAQ via the TIM. During combined ATLAS runs, the only way to clear these BUSY signals is to mask off an entire ROD from the readout.

As an alternative to this, a more efficient masking system was developed whereby a task running on the MDSP monitors the status of all of the input links to the ROD. If a formatter buffer fills beyond a certain fraction of its maximum size, the link will enter a mode whereby only the header and trailer are stored (header-trailer limit, or HTL), and the hit information is discarded. If the buffer continues to fill, then the formatter will assert a BUSY and which will halt triggers. To prevent this happening, the MDSP will mask any links that enter the HTL so that no data are received from the noisy link. Instead, a warning flag is raised in the data stream, marking that link as having been masked. This way both online and offline software has a record of which links are present in the detector for
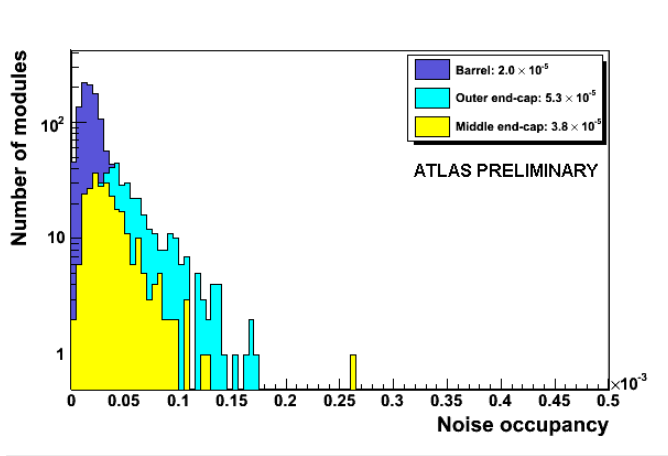
Fig. 7. Average strip noise per module, for all modules in the SCT, after calibration.

efficiency measurements and track reconstruction. In addition, a record of masked links is kept on the MDSP and transferred to the crate server, as a notification to the DAQ operator.

A natural development for the future would be to try to reset and reconfigure noisy modules before unmasking them and reintroducing them into the read-out.

B. Calibration Results

Extensive calibration of the SCT modules was undertaken over the summer of 2008 in preparation for physics running in September. During this period, typically 99% of the barrel modules were operational, and 97% of end-cap modules. A number of modules in the end-cap were not operational due to problems with the cooling plant, some of which have been repaired for 2009.

Characterisation tests during this period resulted in a factor of two reduction in the noise occupancy of modules in the SCT, with smaller tails in the distributions. Figure 7 shows the average strip noise per module above a 1 fC threshold. Variations in noise related to strip length in the different types of modules are clearly seen, with the majority of values between $2 - 5 \times 10^{-5}$. This is around an order of magnitude better than the specification of 5×10^{-4}, giving some headroom for noise occupancy increase after irradiation.

C. Data-taking

In parallel with the calibration runs, the SCT has also collected triggered events in data-taking mode, with a variety of different configurations.

1) Milestone Runs: During 2007 and 2008 a number of week-long milestone runs took place, with the aim of integrating different ATLAS sub-detectors for combined read-out using either random or cosmic triggered events.

Participation of the SCT was limited to the latter of these runs, where the DAQ was exercised with the rest of ATLAS. In the first of these, a small number of modules in a test box were read-out, testing the performance of the trigger system and synchronisation issues. In the next run almost the whole of the SCT barrel took cosmic data, after timing in with the
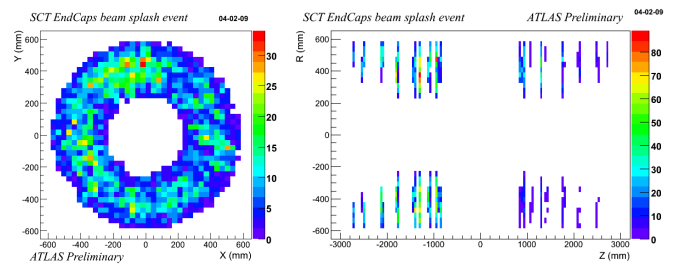

Fig. 8. An example of an LHC beam splash event as seen by the SCT, showing the number of spacepoints, with the average noise from empty events subtracted. The left-hand plot shows the spacepoints in all 18 disks, projected into the plane perpendicular to the beam axis. The right-hand plot shows the same event from the side, with the beam originating from the negative z direction.
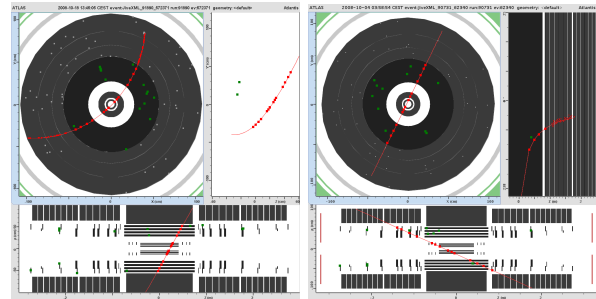

Fig. 9. Two cosmic muon tracks as recorded by the ATLAS inner detector with (left) and without (right) magnetic field. In both cases, the SCT shows hits in all eight barrel layers. The right-hand plot is of additional interest as it shows a cosmic ray traversing one of the SCT end-caps.

rest of ATLAS. For the final two runs the formatter simulator was used, enabling validation of software changes and the successful testing of 70 kHz high rate triggers.

2) First Circulating Beam in the LHC: On the 10th September 2008, the LHC began injecting protons into the accelerator, which were stopped on collimators close to ATLAS. Splashes of particles originating from the collimators reached the ATLAS cavern and were detected by a number of subsystems, including the SCT. Due to concerns about module safety, the bias voltage was reduced to 20 V with a threshold of 1.2 fC, and only the end-caps were powered. An example of a so-called splash event is shown in figure 8.

3) Cosmic Muons: Between September and December 2008, the SCT was mainly concerned with data-taking using cosmic muons, with either the whole of ATLAS, or as part of dedicated Inner Detector cosmic runs. Examples of two cosmic-ray events are shown in figure 9. In total, the SCT DAQ collected around 1.15 million cosmic tracks without the magnetic field switched on, and 0.88 million with the magnetic field at 2 T. During this period, the efficiency of the barrel modules was measured to be well above the 99% nominal value. By examining the distance between a fitted track and the measured SCT spacepoint, detector alignment studies were also performed.

VI. CONCLUSIONS AND OUTLOOK

The SCT data acquisition system is responsible for the configuration, calibration and read-out of 4088 modules of

the SCT detector. The hardware consists of front-end modules, optical links and off-detector read-out crates. Controlling this hardware is the job of the SCT DAQ software, which also oversees the starting and analysis of calibration scans.

During SCT commissioning and operation between 2007 and 2009, a number of DAQ improvements have been made, including masking noisy modules during calibration and physics runs and the addition of a hardware simulator to the ROD.

The commissioning of the SCT was completed in September 2008 in time for the first LHC circulating beams, with detector performance well within design specifications. The following months of cosmic data-taking collected 2 million tracks in the SCT, providing measurements of hit efficiency and detector alignment.

As the LHC prepares for restart towards the end of 2009, the SCT DAQ community will strive for continued improvements of the system and further fine tuning of the detector calibration. It is expected that the SCT will be operational and ready for first LHC collisions with over 99% working modules.

REFERENCES

[1] G. Aad *et al.*, "The ATLAS Experiment at the CERN Large Hadron Collider," *Journal of Instrumentation*, vol. 3, no. 08, p. S08003, 2008. [Online]. Available: http://stacks.iop.org/1748-0221/3/S08003

[2] A. Abdesselam *et al.*, "The Barrel Modules of the ATLAS Semiconductor Tracker," *Nucl. Instrum. Meth.*, vol. A568, pp. 642–671, 2006.

[3] Abdesselam, A. and others, "The ATLAS Semiconductor Tracker End-cap Module," *Nucl. Instrum. Meth.*, vol. A575, no. 3, pp. 353 – 389, 2007.

[4] ——, "The Data Acquisition and Calibration Systems for the ATLAS Semiconductor Tracker," *JINST*, vol. 3, p. P01003, 2008.

[5] ——, "The Optical Links of the ATLAS Semiconductor Tracker," *JINST*, vol. 2, p. P09003, 2007.

[6] T. J. Barber *et al.*, "Resonant Bond Wire Vibrations in the ATLAS Semiconductor Tracker," *Nucl. Instrum. Meth.*, vol. A538, pp. 442–457, 2005.

[7] J. Joseph *et al.*, *ATLAS Silicon ReadOut Driver (ROD) Users Manual*. [Online]. Available: http://www-eng.lbl.gov/~jmjoseph/Atlas-SiROD/Manuals/usersManual-v164.pdf

[8] C. Bee *et al.*, "The Raw Event Format in the ATLAS Trigger and DAQ," *ATLAS Communication*, vol. ATL-DAQ-98-129, 2004.

Tom Barber is a PhD student at the University of Cambridge, but has been based at CERN since he started working on the SCT DAQ in October 2007. At some point he will have to go back to Cambridge and write an analysis on neutral diboson production in ATLAS, and is hoping to get some real data to look at before having to write up. In his spare time he enjoys running, cycling, swimming, hiking and writing conference proceedings.