

# The data acquisition and calibration system for the ATLAS Semiconductor Tracker

A. Abdesselam,<sup>k</sup> T. Barber,<sup>a</sup> A.J. Barr,<sup>k\*</sup> P. Bell,<sup>c†</sup> J. Bernabeu,<sup>q</sup> J.M. Butterworth,<sup>p</sup> J.R. Carter,<sup>a</sup> A.A. Carter,<sup>m</sup> E. Charles,<sup>r</sup> A. Clark,<sup>e</sup> A.-P. Colijn,<sup>i</sup> M.J. Costa,<sup>q</sup> J.M. Dalmou,<sup>m</sup> B. Demirköz,<sup>k‡</sup> P.J. Dervan,<sup>g</sup> M. Donega,<sup>e</sup> M. D'Onofrio,<sup>e</sup> C. Escobar,<sup>q</sup> D. Fasching,<sup>r</sup> D.P.S. Ferguson,<sup>r</sup> P. Ferrari,<sup>c</sup> D. Ferrere,<sup>e</sup> J. Fuster,<sup>q</sup> B. Gallop,<sup>bl</sup> C. García,<sup>q</sup> S. Gonzalez,<sup>r</sup> S. Gonzalez-Sevilla,<sup>q</sup> M.J. Goodrick,<sup>a</sup> A. Gorisek,<sup>c§</sup> A. Greenall,<sup>g</sup> A.A. Grillo,<sup>n</sup> N.P. Hessey,<sup>i</sup> J.C. Hill,<sup>a</sup> J.N. Jackson,<sup>g</sup> R.C. Jared,<sup>r</sup> P.D.C. Johansson,<sup>o</sup> P. de Jong,<sup>i</sup> J. Joseph,<sup>r</sup> C. Lacasta,<sup>q</sup> J.B. Lane,<sup>p</sup> C.G. Lester,<sup>a</sup> M. Limper,<sup>i</sup> S.W. Lindsay,<sup>g</sup> R.L. McKay,<sup>f</sup> C.A. Magrath,<sup>i</sup> M. Mangin-Brinet,<sup>e</sup> S. Martí I García,<sup>q</sup> B. Mellado,<sup>r</sup> W.T. Meyer,<sup>j</sup> B. Mikulec,<sup>e</sup> M. Miñano,<sup>q</sup> V.A. Mitsou,<sup>g</sup> G. Moorhead,<sup>h</sup> M. Morrissey,<sup>l</sup> E. Paganis,<sup>o</sup> M.J. Palmer,<sup>a</sup> M.A. Parker,<sup>a</sup> H. Pernegger,<sup>c</sup> A. Phillips,<sup>a</sup> P.W. Phillips,<sup>m</sup> M. Postranecky,<sup>p</sup> A. Robichaud-Véronneau,<sup>e</sup> D. Robinson,<sup>a</sup> S. Roe,<sup>c</sup> H. Sandaker,<sup>j</sup> F. Sciacca,<sup>p</sup> A. Sfyrla,<sup>e</sup> E. Stanecka,<sup>cd</sup> S. Stapnes,<sup>j</sup> A. Stradling,<sup>r</sup> M. Tyndel,<sup>l</sup> A. Tricoli,<sup>k\*\*</sup> T. Vickey,<sup>r</sup> J.H. Vossebeld,<sup>g</sup> M.R.M. Warren,<sup>p</sup> A.R. Weidberg,<sup>k</sup> P.S. Wells<sup>c</sup> and S.L. Wu<sup>r</sup>

<sup>a</sup> Cavendish Laboratory, University of Cambridge, J.J. Thomson Avenue, Cambridge CB3 0HE, U.K.

<sup>b</sup> School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, U.K.

<sup>c</sup> European Laboratory for Particle Physics (CERN), 1211 Geneva 23, Switzerland

<sup>d</sup> Institute of Nuclear Physics PAN, Cracow, Poland

<sup>e</sup> DPNC, University of Geneva, CH 1211 Geneva 4, Switzerland

<sup>f</sup> Department of Physics and Astronomy, 12 Physics Hall, Ames, IA 50011, U.S.A.

<sup>g</sup> Oliver Lodge Laboratory, University of Liverpool, Liverpool, U.K.

<sup>h</sup> University of Melbourne, Parkville, Vic. 3052, Australia

<sup>i</sup> NIKHEF, Amsterdam, The Netherlands

<sup>j</sup> Department of Physics, P.O. Box 1048, Blindern, N-0316 Oslo, Norway

<sup>k</sup> Physics Department, University of Oxford, Keble Road, Oxford OX1 3RH, U.K.

<sup>l</sup> Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire OX11 0QX, U.K.

<sup>m</sup> Department of Physics, Queen Mary University of London, Mile End Road, London E1 4NS, U.K.

<sup>n</sup> Santa Cruz Institute for Particle Physics, University of California, Santa Cruz, CA, U.S.A.

<sup>o</sup> Department of Physics and Astronomy, University of Sheffield, Sheffield, U.K.

<sup>p</sup> Department of Physics and Astronomy, UCL, Gower Street, London, WC1E 6BT, U.K.

<sup>q</sup> Instituto de Física Corpuscular (IFIC), Universidad de Valencia-CSIC, Valencia, Spain

<sup>r</sup> University of Wisconsin-Madison, Wisconsin, U.S.A.

E-mail: [a.barr@physics.ox.ac.uk](mailto:a.barr@physics.ox.ac.uk)

\* Corresponding author.

† Now at the School of Physics and Astronomy, University of Manchester, Manchester M13 9PL, U.K.

‡ Now at the European Laboratory for Particle Physics (CERN), 1211 Geneva 23, Switzerland

§ Now at the Jožef Stefan Institute and Department of Physics, University of Ljubljana, Ljubljana, Slovenia

\*\* Now at the Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire OX11 0QX, U.K.

**ABSTRACT:** The SemiConductor Tracker (SCT) data acquisition (DAQ) system will calibrate, configure, and control the approximately six million front-end channels of the ATLAS silicon strip detector. It will provide a synchronized bunch-crossing clock to the front-end modules, communicate first-level triggers to the front-end chips, and transfer 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 endcap assembly and for performance confirmation tests after transport of the barrels and endcaps to CERN. Operating in data-taking mode, the DAQ has recorded nearly twenty million synchronously-triggered events during commissioning tests including almost a million cosmic ray triggered events. 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.

**KEYWORDS:** Particle tracking detectors; Data acquisition concepts; Software architectures (event data models, frameworks and databases); Control and monitor systems online.

---

## Contents

<b>1. Introduction</b>	<b>1</b>
<b>2. Off-detector hardware overview</b>	<b>3</b>
2.1 The Read-out Driver (ROD)	4
2.1.1 Operating modes	5
2.1.2 Physics data-taking	5
2.1.3 Calibration	6
2.1.4 ROD communication	6
2.2 Back of Crate card (BOC)	7
2.2.1 BOC Hardware Implementation	8
2.3 TTC Interface Module (TIM)	8
2.3.1 TIM Hardware Implementation	10
2.4 Data rates	10
<b>3. Readout software</b>	<b>11</b>
<b>4. Detector setup and calibration</b>	<b>13</b>
4.1 Module communication optimisation	13
4.2 Front-end calibration	13
<b>5. Application and results</b>	<b>16</b>
5.1 Barrel and endcap macro-assembly	16
5.2 Commissioning and cosmic ray tests	17
<b>6. Conclusions</b>	<b>19</b>
<b>A. Abbreviations</b>	<b>20</b>


---

## 1. Introduction

The ATLAS experiment is one of two general-purpose detectors at CERN's Large Hadron Collider (LHC). The SemiConductor Tracker (SCT) is a silicon strip detector and forms the intermediate tracking layers of the ATLAS inner detector. The SCT has been designed to measure four precision three-dimensional space-points for charged particle tracks with pseudorapidity<sup>1</sup>  $|\eta| < 2.5$  (figure 1).

---

<sup>1</sup> The pseudorapidity,  $\eta$ , is given by  $\eta = -\ln \tan (\theta/2)$ , where  $\theta$  is the polar angle relative to the beam axis.



**Figure 1.** Cross section of the ATLAS Inner Detector showing a quarter of the barrel and half of one of the two endcap regions. The SCT is within a Transition Radiation detector (TRT) and surrounds a Pixel detector [1]. The dimensions are in mm.

**Table 1.** Radius and number of modules on each of the four SCT barrel layers.

Barrel	B3	B4	B5	B6	Total
Radius / mm	299	371	443	514	
Modules	384	480	576	672	2112

**Table 2.** Longitudinal position and number of modules for the nine disks on each SCT endcap.

Disk	1	2	3	4	5	6	7	8	9	Total
z  / mm	847	934	1084	1262	1377	1747	2072	2462	2727	
Modules	92	132	132	132	132	132	92	92	52	988

The complete SCT consists of 4088 front-end modules [2],[3]. Each module has two planes of silicon each with 768 active strips of  $p^+$  implant on n-type bulk [4]. The planes are offset by a small stereo angle (40 mrad), so that each module provides space-point resolutions of 17  $\mu\text{m}$  perpendicular to and 580  $\mu\text{m}$  parallel to its strips. The implant strips are capacitively coupled to aluminium metalisation, and are read out by application-specific integrated circuits (ASICs) known as ABCD3TA [5]. Each of these chips is responsible for reading out 128 channels, so twelve are required for each SCT module.

The SCT is geometrically divided into a central barrel region and two endcaps (known as ‘A’ and ‘C’). The barrel region consists of four concentric cylindrical layers (barrels). Each endcap consists of nine disks. The number of modules on each barrel layer and endcap disk is given in table 1 and table 2. The complete SCT has 49,056 front-end ASICs and more than six million individual read-out channels.

For physics data-taking the data acquisition (DAQ) system must configure the front-end ASICs, communicate first-level trigger information, and transfer data from the front-end chips to the ATLAS high-level trigger system.

The role of the DAQ in calibrating the detector is equally important. The SCT uses a “binary” readout architecture in which the only pulse-height information transmitted by the front-end chips is one bit per channel which denotes whether the pulse was above a preset threshold. Further information about the size of the pulse cannot be recovered later, so the correct calibration of these thresholds is central to the successful operation of the detector.

The discriminator threshold must be set at a level that guarantees uniform, good efficiency while maintaining the noise occupancy at a low level. Furthermore the detector must maintain good performance even after a total ionizing dose of 100 kGy(Si) and a non-ionising fluence of  $2 \times 10^{14}$  neutrons/cm<sup>2</sup> of 1-MeV neutrons, corresponding to 10 years of operation of the LHC at its design luminosity. The performance requirements, based on track-finding and pattern-recognition considerations, are that channel hit efficiency should be greater than 99% and noise occupancy less than  $5 \times 10^{-4}$  per channel even after irradiation.

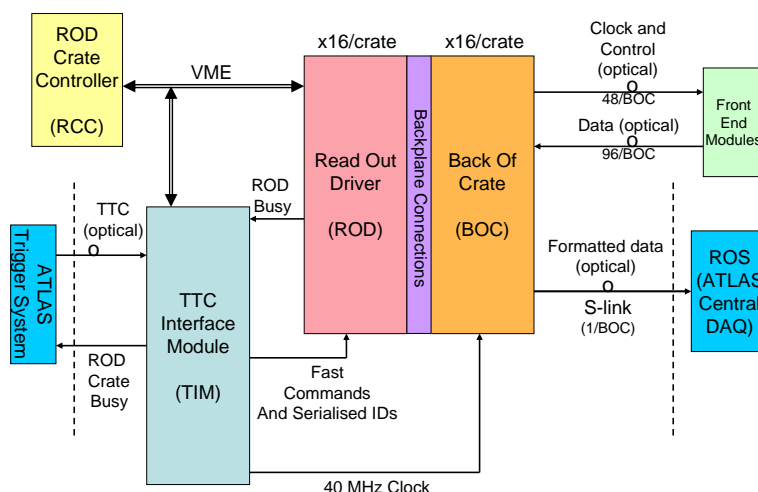
During calibration, internal circuits on the front-end chips can be used to inject test charges. Information about the pulse sizes is reconstructed by measuring occupancy (the mean number of hits above threshold per channel per event) as a function of the front-end discriminator threshold (threshold “scans”). The calibration system must initiate the appropriate scans, interpret the large volume of data obtained, and find an improved configuration based on the results.

This paper is organized as follows. In section 2 there is a description of the readout hardware. The software and control system are discussed in section 3. In section 4 there is a description of the calibration procedure. A review of the operation of the data acquisition system is given in section 5 together with some of the main results, covering both the confirmation tests performed during the mounting of SCT modules to their carbon-fibre support structures (“macro-assembly”) and more recent tests examining the performance of the completed barrel and endcaps at CERN (“detector commissioning”). We conclude in section 6. A list of some of the common abbreviations used may be found in the appendix.

## 2. Off-detector hardware overview

The off-detector readout hardware of the SCT DAQ links the SCT front-end modules with the ATLAS central trigger and DAQ system [6], and provides the mechanism for their control. The principal connections to the front-end modules, to the ATLAS central DAQ and between SCT-specific components are shown in figure 2.

The SCT DAQ consists of several different components. The Read Out Driver (ROD) board performs the main control and data handling. A complementary Back Of Crate (BOC) board handles the ROD’s I/O requirements to and from the front-end, and to the central DAQ. Each ROD/BOC pair deals with the control and data for up to 48 front-end modules. There can be up to 16 RODs and BOCs housed in a standard LHC-specification 9U VME64x crate with a custom backplane [7], occupying slots 5-12, 14-21. In slot 13 of the crate is a TTC Interface Module (TIM) which accepts the Timing, Trigger and Control (TTC) signals from ATLAS and distributes them to the RODs and BOCs. The ROD Crate Controller (RCC) is a commercial 6U Single Board Computer running Linux which acts as the VME master, and hence it usually occupies the first slot in the crate. The RCC configures the other components and provides overall control of the data acquisition functions within a crate. The VME bus is used by the RCC to communicate with the RODs and with the TIM. Communication between each ROD and its



**Figure 2.** Block diagram of the SCT data acquisition hardware showing the main connections between components.

partner BOC and between the TIM and the BOCs is via other dedicated lines on the backplane. The highly modular design was motivated by considerations of ease of construction and testing.


In physics data-taking mode, triggers pass from the ATLAS TTC [8] to the TIM and are distributed to the RODs. Each ROD fans out the triggers via its BOC to the front-end modules. The resultant hit data from the front-end modules are received on the BOC, formatted on the ROD and then returned to the BOC to be passed on to the first module of the ATLAS central DAQ, known as the Read-Out Subsystem (ROS) [9]. The RODs can also be set up to sample and histogram events and errors from the data stream for monitoring.

For calibration purposes, the SCT DAQ can operate separately from the central ATLAS DAQ. In this mode the ATLAS global central trigger processor (CTP) is not used. The TIM generates the clock and SCT-specific triggers are taken from other sources. For most tests they are generated internally on the RODs, but for tests which require synchronisation they can be sourced from the SCT's local trigger processor (LTP) [10] or from the TIM. The resultant data are not passed on to the ROS, but the ROD monitoring functions still sample and histogram the events. The resultant occupancy histograms are transferred over VME to the ROD Crate Controller and then over the LAN to PC servers for analysis.

In both modes, the data sent from the front end modules must be identified with a particular LHC bunch crossing and first-level trigger. To achieve this, each front-end ASIC keeps a count of the number of triggers (4 bits) and the number of clocks (8 bits) it has received. The values of the counters form part of each ASIC's event data header. Periodic counter resets can be sent to the front end ASICs through the TTC system.

## 2.1 The Read-out Driver (ROD)

The Silicon Read-out Driver (ROD) [11] is a 9U 400mm deep VME64x electronics board. The primary functions of the ROD are the front-end module configuration, trigger propagation and event data formatting. The secondary functions of the ROD are detector calibration and monitoring. Control commands are sent from the ROD to the front-end modules as serial data streams. These commands can be first-level triggers, bunch-crossing (clock counter) resets,



**Figure 3.** An overview of the ATLAS Silicon Read-out Driver (ROD) data and communication paths.

event (trigger) counter resets, calibration commands or module register data. Each ROD board is capable of controlling the configuration and processing the data readout of up to 48 SCT front-end modules. After formatting the data collected from the modules into 16-bit words, the ROD builds event fragments which are transmitted to the ROS via a high speed serial optical link known as the S-Link [12].

A hybrid architecture of Field Programmable Gate Arrays (FPGAs) and Digital Signal Processors (DSPs) allows the ROD the versatility to perform various roles during physics data-taking and calibrations. Four FPGA designs are used for all of the real-time operations required for data processing at the ATLAS trigger rate. The Formatter, Event Fragment Builder and Router FPGAs are dedicated to performing time-critical operations, in particular the formatting, building and routing of event data. The Controller FPGA controls operations such as ROD setup, module configuration distribution and trigger distribution. A single “Master” (MDSP) and four “Slave” (SDSP) DSPs on the board are used to control and coordinate on-ROD operations, as well as for performing high-level tasks such as data monitoring and module calibration. Once configured, the ROD FPGAs handle the event data-path to the ATLAS high-level trigger system without further assistance from the DSPs. The major data and communication paths on the ROD are shown in figure 3.

### 2.1.1 Operating modes

The ROD supports the two main modes of operation: physics data-taking and detector calibrations. The data-path through the Formatter and the Event Fragment Builder FPGAs is the same in both modes of operation. In data-taking mode the Router FPGA transmits event fragments to the ROS via the S-Link and optionally also to the SDSPs for monitoring. In calibration mode the S-Link is disabled and the Router FPGA sends events to the farm of Slave DSPs for histogramming.

### 2.1.2 Physics data-taking

After the data-path on the ROD has been set up, the event data processing is performed by the FPGAs without any intervention from the DSPs. Triggers issued from the LTP are relayed to the ROD via the TIM. If the S-Link is receiving data from the ROD faster than they can be

transferred to the ROS, back-pressure will be applied to the ROD, thereby halting the transmission of events and causing the internal ROD FIFOs to begin to fill. Once back-pressure has been relieved, the flow of events through the S-Link resumes. In the rare case where the internal FIFOs fill beyond a critical limit, a ROD busy signal is raised on the TIM to stop triggers.

The Router FPGA can be set up to capture events with a user-defined pre-scale on a non-interfering basis and transmit them to the farm of SDSPs. Histogramming these captured events and comparing them against a set of reference histograms can serve as an indicator of channels with unusually high or low occupancies and the captured data can be monitored for errors.

### 2.1.3 Calibration

When running calibrations, the MDSP serial ports can be used to issue triggers to the modules. In calibration mode the transmission of data through the S-Link is inhibited. Instead, frames of data (256 32-bit word blocks) are passed from the Router FPGA to the SDSPs using a direct memory access transfer. Tasks running on the SDSPs flag these transferred events for processing and subsequent histogramming. A monitoring task can be run on the SDSPs that is capable of parsing the event errors flagged by the FPGAs and reporting these errors back to the RCC. More details on the use of the ROD histogramming tasks for calibration can be found in section 4.

### 2.1.4 ROD communication


The ROD contains many components, and is required to perform many different operations in real time. For smooth operation it is important that the different components have a well-defined communication protocol. A system of communication registers, “primitives”, “tasks” and text-buffers is used for RCC-to-ROD and Master-to-Slave inter DSP communication and control.

The communication registers are blocks of 32-bit words at the start of the DSP’s internal memory which are regularly checked by the Master DSP (MDSP) inside the main thread of execution running on the processor. The MDSP polls these registers, watching for requests from the RCC. These registers are also polled by the RCC and so can be used by it to monitor the status of the DSPs. Such registers are used, for example, to indicate whether the event trapping is engaged, to report calibration test statistics, and for communicating between the RCC and the ROD the status of “primitive” operations. The ROD FPGA registers are mapped in the MDSP memory space.

The “primitives” are software entities which allow the MDSP to remain in control of its memory while receiving commands from the RCC. Each primitive is an encoding in a block of memory which indicates a particular command to the receiving DSP. These are copied to a known block of memory in groups called “primitive lists”. It is through primitives that the ROD is configured and initialized. Generally each primitive is executed once by the receiving DSP. Primitives exist for reading and writing FPGA registers, reading and writing regions of SDSP memory, loading or modifying front-end module configurations, starting the SDSPs, and to start and stop “tasks”. The MDSP can send lists of primitives to the SDSPs, for example to start calibration histogramming. The DSP software is versatile enough to easily allow the addition of new primitives representing extra commands when required.

“Tasks” are DSP functions which execute over an extended period of time. These are started and stopped by sending primitives from RCC to MDSP, or from MDSP to SDSP and continue to execute in cooperation with the primitive list thread. They run until completion or until they are halted by other primitives. Examples of tasks are the histogramming and the histogram control tasks. The former runs on the SDSPs handling histogramming of events while the latter runs on the MDSP and manages the sending of triggers, as well as changes in chip configuration and histogram bin changes.





**Figure 4.** Block diagram showing the layout and main communication paths on the BOC card.

## 2.2 Back of Crate card (BOC)

The BOC transmits commands and data between the ROD and the optical fibre connections which service the front-end modules, and is also responsible for sending formatted data to the ROS. It also distributes the 40 MHz bunch-crossing clock from the TIM to the front-end modules and to its paired ROD. A block diagram of the function of the BOC is shown in figure 4.

The front-end modules are controlled and read out through digital optical fibre ribbons. One fibre per module provides trigger, timing and control information. There are also two data fibres per module which are used to transfer the digital signal from the modules back to the off-detector electronics. A more detailed description of the optical system is given in [13].

On the BOC, each command for the front-end modules is routed via one of the four TX plug-ins as shown in figure 4. Here the command is combined with the 40 MHz clock to generate a single Bi-Phase Mark (BPM) encoded signal which allows both clock and commands to occupy the same stream. Twelve streams are handled by each of four BPM12 chips [14]. The encoded commands are then converted from electrical to optical form on a 12-way VCSEL array [14] before being transmitted to the front-end modules via a 12-way fibre ribbon. The intensity of the laser light can be tuned in individual channels by controlling the current supplied to the laser using a digital to analogue converter (DAC) on the BOC. This is to cater for variations in the individual lasers, fibres and receivers and to allow for loss of sensitivity in the receiver due to radiation damage.

The timing of each of the outgoing signals from the TX plug-in can be adjusted so that the clock transmitted to the front-end modules has the correct phase relative to the passage of the particles from the collisions in LHC. This phase has to be set on a module by module basis to allow for different optical fibre lengths and time-of-flight variations through the detector. It is also

necessary to ensure that the first-level trigger is received in the correct 25 ns time bin, so that the data from the different ATLAS detectors are merged into the correct events. For this reason, there are two timing adjustments available – a coarse one in 25 ns steps, a fine one in 280 ps steps.

Incoming data from the front-end modules are accepted by the BOC in optical form, converted into electrical form and forwarded to the ROD. As each front-end module has two data streams and each ROD can process data for up to 48 modules, there are 96 input streams on a BOC. The incoming data are initially converted from optical to electrical signals at a 12-way PIN diode array on the RX plug-in. These signals are then discriminated by a DRX12 chip [14]. The data for each stream are sampled at 40 MHz, with the sampling phase and threshold adjusted so that a reliable ‘1’ or ‘0’ is selected. The binary stream is synchronized with the clock supplied to the ROD so that it receives the data at the correct phase to ensure reliable decoding.

After the data are checked and formatted in the ROD, they are returned to the BOC for transmitting to the first element of the ATLAS higher-level trigger system (the ROS) via the S-Link connection. There is a single S-Link connection on each BOC.

The 40 MHz clock is usually distributed from the TIM, via the backplane and the BOC, to the front-end modules. However, in the absence of this backplane clock, a phase-locked loop on the BOC will detect this state and generate a replacement local clock. This is important not only because the ROD relies on this clock to operate, but also because the front-end modules dissipate much less heat when the clock is not present, and thermal changes could negatively affect the precision alignment of the detector.

### 2.2.1 BOC Hardware Implementation

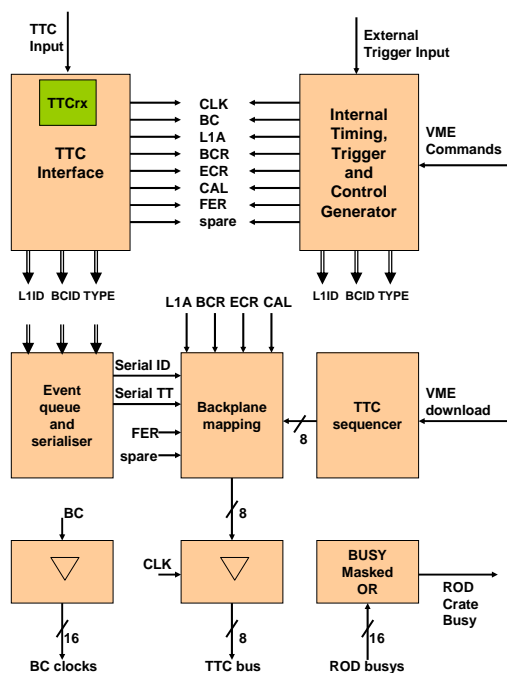
The BOC is a 9U, 220mm deep board and is located in the rear of the DAQ crate. It is not directly addressable via VME as it only connects to the J2 and J3 connectors on the backplane and so all configuration is done over a set-up bus via the associated ROD.

A complex programmable logic device (CPLD) is used for overall control of the BOC. Further CPLDs handle the incoming data – these have been used rather than non-programmable devices as the BOC was designed to be also usable by the ATLAS Pixel Detector, which has different requirements. As can be seen from the previous section, there is a significant amount of clock-timing manipulation on the BOC. Many of these functions are implemented using the PHOS4 chip [15], a quad delay ASIC which provides a delay of up to 25 ns, in 1 ns units. The functions of the BOC (delays, receiver thresholds, laser currents etc.) are made available via a set of registers. These registers are mapped to a region of ROD MDSP address space via the setup bus, so that they are available via VME to the DAQ. The S-Link interface is implemented by a HOLA [16] daughter card.

### 2.3 TTC Interface Module (TIM)

The TIM [17] interfaces the ATLAS first-level trigger system signals to the RODs. In normal operation it receives clock and trigger signals from the ATLAS TTC system [18] and distributes these signals to a maximum of 16 RODs and their associated BOCs within a crate. Figure 5 illustrates the principal functions of the TIM – transmitting fast commands and event identifiers from the ATLAS TTC system to the RODs, and sending the clock to the BOCs (from where it is passed on to the RODs).

The TIM has various programmable timing adjustments and control functions. It has a VME slave interface to give the local processor read and write access to its registers, allowing it to be configured by the RCC. Several registers are regularly inspected by the RCC for trigger counting and monitoring purposes.



**Figure 5.** Block diagram showing a functional model of the TIM hardware. Abbreviations are used for the bunch crossing clock (BC/CLK), first-level trigger (L1A), event counter reset (ECR), bunch counter reset (BCR), calibrate signal (CAL), first-level trigger number (L1ID), bunch crossing number (BCID), trigger type (TYPE) and front end reset (FER).

The incoming optical TTC signals are received on the TIM using an ATLAS standard TTCrx receiver chip [19], which decodes the TTC information into electrical form. In the physics mode the priority is given to passing the bunch-crossing clock and commands to the RODs in their correct timing relationship, with the absolute minimum of delay to reduce the latency. The TTC information is passed onto the backplane of a ROD crate with the appropriate timing. The event identifier is transmitted with a serial protocol and so a FIFO buffer is used in case of rapid triggers.

For tests and calibrations the TIM can, at the request of the local processor (RCC), generate all the required TTC information itself. It can also be connected to another TIM for stand-alone SCT multi-crate operation. In this stand-alone mode, both the clock and the commands can be generated from a variety of sources. The 40 MHz clock can be generated on-board, derived from an 80.16 MHz crystal oscillator, or transferred from external sources in either NIM or differential ECL standards. Similarly, the fast commands can be generated on the command of the RCC, or automatically by the TIM under RCC control. Fast commands can also be input from external sources in either NIM or differential ECL. These internally or externally generated commands are synchronised to whichever clock is being used at the time, to provide the correctly timed outputs. All the backplane signals are also mirrored as differential ECL outputs on the front panel to allow TIM interconnection.

A sequencer, using  $8 \times 32\text{k}$  RAM, allows long sequences of commands and identifiers to be written in by the local processor and used for testing the front-end and off-detector electronics. A ‘sink’ (receiver RAM) of the same size is also provided to allow later comparisons of commands and data sent to the RODs.

The TIM also controls the crate's busy logic, which tells the ATLAS CTP when it must suspend sending triggers. Each ROD returns an individual busy signal to the TIM, which then produces a masked OR of the ROD busy signals in each crate. The overall crate busy is output to the ATLAS TTC system. ROD busy signals can be monitored using TIM registers.

The CDF experiment at Fermilab found that bond wires could break on front-end modules when forces from time-varying currents in the experiment's magnetic field excited resonant vibrations [20]. The risk to the ATLAS SCT modules is considered to be small [21], even on the higher-current bond wires which serve the front-end optical packages. These bonds have mechanical resonances at frequencies above 15 kHz so, as a precaution, the TIM will prevent fixed-frequency triggers from being sent to the front-end modules. If ten successive triggers are found at fixed frequencies above 15 kHz, a period-matching algorithm on the TIM will stop internal triggers. It will also assert a BUSY signal which should stop triggers from being sent by the ATLAS CTP. If incoming triggers continue to be sent, the TIM will enter an emergency mode and independently veto further triggers. The algorithm has been demonstrated to have a negligible effect on data-taking efficiency [22].

### 2.3.1 TIM Hardware Implementation

The TIM is a 9U, 400 mm deep board. The TTCrx receiver chip and the associated PIN diode and preamplifier developed by the RD12 collaboration at CERN [19] provide the bunch-crossing clock and the trigger identification signals. On the TIM, a mezzanine board (the TTCrq [23]) allows an easy replacement if required.


Communication with the BOCs is via a custom J3 backplane. The bunch-crossing clock destined for the BOCs and RODs, with the timing adjusted on the TTCrx, is passed via differential PECL drivers directly onto the point-to-point parallel impedance-matched backplane tracks. These are designed to be of identical length for all the slots in each crate to provide a synchronised timing marker. All the fast commands are clocked directly, without any local delay, onto the backplane to minimise the TIM latency-budget.

On the TIM module, a combination of FastTTL, LVTTL, ECL, PECL and LV BiCMOS devices is used. The Xilinx Spartan IIE FPGA series were chosen as the programmable logic devices. Each TIM uses two of these FPGAs. These devices contain enough RAM resources to allow the RAMs and FIFOs to be incorporated into the FPGA.

The TIM switches between different clock sources without glitches and, in the case of a clock failure, does so automatically. To achieve this, dedicated clock-multiplexer devices have been used. These devices switch automatically to a back-up clock if the selected clock is absent. Using clock detection circuits, errors can be flagged and transmitted to all the RODs in the crate via a dedicated backplane line, allowing RODs to tag events accordingly.

### 2.4 Data rates

The system has been designed to operate at the expected ATLAS first-level trigger rate of 75 kHz and up to a maximum rate of 100 kHz [24]. At 100 kHz, the front-end module to BOC data-links will on average require 40% of the available bandwidth at 1% average front-end hit occupancy and 70% of the available bandwidth at 2% average hit occupancy (when both data-links on that module are operational and are equally loaded.) An eight-deep readout buffer in the front-end ASICs ensures that the fraction of data which can be lost due to the buffer overflowing will remain less than 1%, even for mean hit occupancy up to 2% and an average trigger rate of 100 kHz [5]. This includes a large safety factor as the expected worst case strip occupancy (for the highest occupancy front-end modules) averaged over strips and time is less than 1%.



**Figure 6.** Schematic control and data flow diagram for the SCT calibration and control system.

The S-Link interface has been tested with ROD-generated test data at rates of up to 158 MBytes per second. This data rate corresponds to 1% time-averaged and module-averaged occupancy at first-level trigger rates of 100 kHz. The actual occupancy, averaged over modules on each ROD, is expected to be smaller than 1%; simulations suggest hit occupancies should range from 0.33% for the outermost barrel up to 0.64% for the innermost barrel [1].

Further system tests with large numbers of real detector modules are described in section 5.


### 3. Readout software

The complete ATLAS SCT DAQ hardware comprises many different elements: nine rack-mounted Linux PCs and eight crates containing eight TIMs, eight Linux RCCs and ninety ROD/BOC pairs. The SctRodDaq software [25],[26],[27] controls this hardware and provides the operator with an interface for monitoring the status of the front-end modules as well as initiating and reviewing calibrations. The software can optimise the optical communication registers as well as testing and calibrating the front-end ASICs.

It is important that the calibration can proceed rapidly, so that the entire detector can be characterized within a reasonable time. To achieve this, an iterative procedure is generally used, fixing parameters in turn. The results of each step of the calibration are analysed, and the relevant optimisation performed before the subsequent step is started. Both the data-taking and the data-analysis of each step must therefore be performed as quickly as possible, and to satisfy the time constraints parallel processes must run for both the data-taking and the analysis.

A diagram of the main software components is shown in figure 6. The readout software comprises approximately 250 thousand lines of code written largely in C++ and Java. The hardware-communication parts of the software (the SctApi crate controllers) run on the RCCs and control the RODs, BOCs and TIMs over VME. They are responsible for loading configuration data, setting up the on- and off-detector hardware, performing the actions required during run state transitions, and retrieving monitoring histograms from the RODs. During calibration, they initiate calibration scans and retrieve calibration histograms from the RODs.

The analysis subsystem and user interface run on dedicated rack-mounted Linux PCs. The calibration controller is responsible for synchronizing control during calibration operation. The fitting and analysis services perform data-reduction and calculate the optimal values for



**Figure 7.** A view of the SCT graphical user interface. This display is showing the BOC receiver thresholds for the RX data links on all BOCs on a particular DAQ crate. The colour scale indicates the current value of the threshold. The display is set in a crate view where vertical lines of modules represent complete ROD/BOC pairs, each of which services up to 48 modules.

calibration parameters. The archiving services read data from transient objects and write them to persistent storage on disk. Inter-process communication is based on a number of ATLAS online software tools [28] one of which (IPC) provides a partition-based naming service for CORBA-compliant interfaces.

Since many operations need to be done in near real-time, most of the processes have concurrent threads of execution. For example the service which fits the occupancy histograms implements a worker/listener pattern. As new data arrive, they are added to a queue by a listener thread which is then immediately free to respond to further data. Meanwhile one or more worker threads undertake the processor-intensive job of performing the fits. The fit algorithms have been optimised for high performance since for most tests several fits are required for every read-out channel (see section 4).

The front-end and DAQ system configuration can be stored either in an XML file or in a relational database. It is populated with data from previous calibrations, including quality assurance tests taken during front-end module assembly, and confirmation tests performed during macro-assembly and detector commissioning. A Java-based graphical user interface (figure 7) allows the operator to launch calibration tests, to display the results of tests, to display calibration information and to compare results to reference data.

## 4. Detector setup and calibration

Good front-end calibration is central to the correct operation of the detector because a large amount of pre-processing and data reduction occurs on the SCT's front-end ASICs. Each front-end chip has an 8-bit DAC which allows the threshold to be set globally across that chip. Before irradiation it is generally straightforward to find a threshold for which both the noise occupancy ( $< 5 \times 10^{-4}$ ) and efficiency ( $> 99\%$ ) specifications can be satisfied. After the module is irradiated setting the threshold at a suitable level becomes even more important. Irradiation decreases the signal collection, and increases the noise seen by the front-end. This means that after 10-years LHC equivalent radiation the working region narrows, and to satisfy the performance requirements the channel thresholds need to be set within a more limited range [29]. To assure uniformity of threshold, every channel has its own 4-bit DAC (TrimDAC) which is used to compensate for channel-to-channel threshold variations. The TrimDAC steps can themselves be set to one of four different values, allowing uniformity of thresholds to be maintained even as uncorrected channel-to-channel variations increase during irradiation.

### 4.1 Module communication optimisation

Before the front-end modules can be calibrated, the system must be configured to ensure reliable communication between the modules and the BOCs.

When first powered, the SCT modules return a clock signal on each of their two optical links which is half the frequency of the 40.08 MHz input clock signal. The first round of optimisation uses counters implemented in ROD firmware to determine the number of logical ones received in a fixed time period for each point of a matrix of two variables: the optical receiver threshold and the sampling phase used to latch the incoming data. An operating point is chosen within the range of values for which the 20.04 MHz signal is received correctly.

For the second round of optimisation, the front-end modules are set up to return the contents of their configuration registers, so that a known bit-pattern can be expected. Triggers are sent and the value of the receiver threshold varied in order to locate the region in which the binary stream is faithfully transmitted. This technique is slower than that used in the first round of optimisation, but is necessary because of its greater sensitivity to slow turn-on effects exhibited by a small number of VCSELs used in the detector. The optical tuning process is described in more detail in [13][27].

### 4.2 Front-end calibration

Most of the calibration procedures are designed to set registers on the front-end ASICs. The most important features of these ABCD3TA chips [5] can be seen in figure 8. The analogue front-end carries out charge integration, pulse shaping, amplitude discrimination, and latching of data. The digital pipeline stores the resultant binary 1-bit per channel information for 132 clock cycles pending the global ATLAS first-level trigger decision. If such a trigger is received, the data are passed into an eight event deep buffer and read out serially with token-passing between the chips daisy-chained on the module. The data are compressed using an algorithm which only transmits information about hit channels.

For calibration purposes, known charges are injected into the front end of each readout channel. Every fourth channel is tested simultaneously, the set of active channels being determined by two bits in each chip's configuration register. The calibration charges are generated by applying voltage pulses of known amplitude, set by a dedicated 8-bit DAC, across the calibration capacitors. To compensate for wafer-to-wafer variations in capacitance which

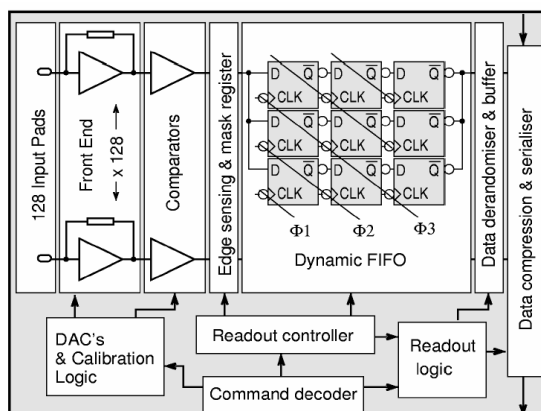


Figure 8. Simplified Block Diagram of a SCT ABCD3TA ASIC [5].





Figure 9. Occupancy as a function of front-end discriminator threshold. (a) The shading scale shows the fraction of hits, as a function of the channel number (x-axis), and comparator threshold in mV (y-axis) for all channels on one side of a barrel module (6 ASICs; 768 channels). The front-end parameters were already optimised before this scan, so the channel-to-channel and chip-to-chip occupancy variations are small. (b) Mean occupancy and complementary error function fits for each of the six ASICs.


can occur during ASIC manufacture, correction factors were obtained from measurements on a number of test structures on each wafer. The applied voltage step used during charge injection is adjusted in accordance with these factors on a chip-to-chip basis.

For each channel, a histogram of occupancy as a function of discriminator threshold is created, and a complementary error function fitted. The threshold at which the occupancy is 50% corresponds to the median of the injected charge, while the sigma gives the noise after amplification. An example threshold scan is shown in figure 9. During this calibration scan 500 triggers were sent per threshold point, and the charge injected was 1.5 fC.





**Figure 10.** Response curves, showing the value of the discriminator threshold at which the mean chip occupancy is 50% as a function of the charge injected, for each of the 12 chips on one module.



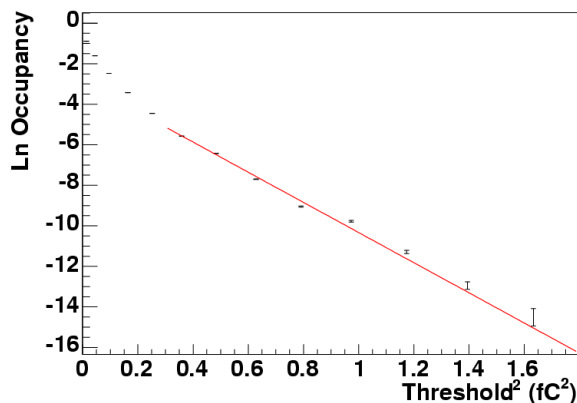
**Figure 11.** Histograms of the threshold DAC value at which the occupancy is 50% ( $V_{th}$  50%) for each channel, relative to the mean for the module. (a) Before the TrimDAC registers were optimised, and (b) after optimising the TrimDAC registers.

To calibrate the discriminator threshold, the DAQ system initiates threshold scans for several different values of injected charge. Example ten-point response curves for a particular module are shown in figure 10. The points are fitted with curves of the functional form:  $y = c_2 + c_0(1 + \exp(-x/c_1))^{-1}$ , with the parameters  $c_{0,1,2}$  allowed to vary during the fit. From the data and the fitted curves the front-end gain and noise are calculated. The gain is the gradient of the response curve. The noise before amplification can be calculated by dividing the noise after amplification by the gain [30]. The gain and the noise are usually quoted at 2 fC input charge.

A similar technique is used to optimise the TrimDAC registers. For this test the injected charge is held constant and threshold scans are performed for different values of the TrimDAC registers. Using the results, an algorithm chooses optimal trim values, which reduce the channel-to-channel variations in the threshold (figure 11).

Threshold scans with no injected charge (figure 12) are used to find the noise occupancy. The response curve allows the chip threshold to be calibrated in units of front-end input charge. The parameter of interest is the noise occupancy near the 1 fC nominal working point.

A variety of different noise scans have been used to search for any signs of cross-talk or noise pick-up. One example is a test designed to be sensitive to any electrical or optical activity associated with the ASIC readout. For that test, pairs of triggers are generated with a variable separation, close to the duration of the pipeline delay, so that the second trigger is received when the data associated with the first trigger is at different stages of being read out. The noise



**Figure 12.** The (natural) logarithm of one front-end chip’s average noise occupancy as a function of the square of the discriminator threshold, measured after calibration.

occupancy associated with the second trigger is examined for any dependence on the trigger separation time.

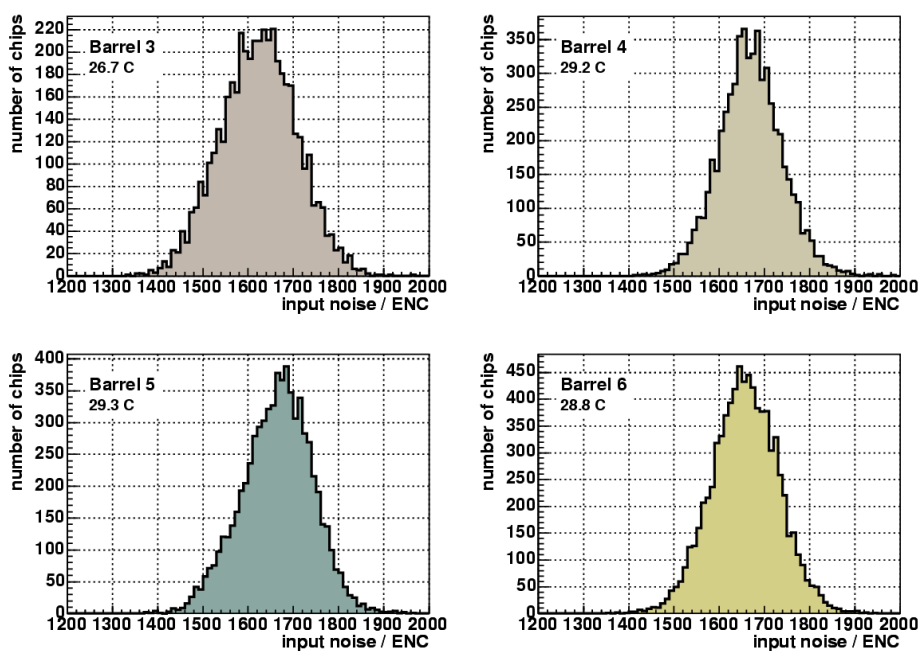
A full test sequence contains other procedures which verify the digital performance of the ASICs. These exercise and test the front-end trigger and bunch-crossing counter registers, the channel mask registers, pipeline cells and chip token-passing logic as described in [30],[31],[32]. The readout system can also initiate specialised scans, for example for timing-in the detector to the LHC bunch crossing, for fine tuning the relative timing of the front-end modules, and for modifying the TX optical duty-cycle to minimise the clock jitter seen by the front end ASICs.

## 5. Application and results

### 5.1 Barrel and endcap macro-assembly

The SctRodDaq software was used extensively to test the performance of large numbers of modules after mounting onto their support structures at the assembly sites [2],[3],[30],[32],[33],[34]. Groups of up to 672 modules (the complete outermost barrel, B6 in table 1) were tested simultaneously with single crate DAQs. The ATLAS central elements (CTP, LTP, ROS etc) were not present, so the DAQ was operated in calibration mode, with triggers generated either on the RODs or on the TIM. The hit data were histogrammed on the RODs with S-Link transmission inhibited. Tests were performed to measure the noise performance, confirm known problem channels, and check that no new electrical defects had been introduced during the assembly process.

A typical time to run, analyse and feedback the results of a calibration test consisting of three front-end threshold scans is about 20 minutes, for a test in which 500 triggers are sent for each of 100 different ASIC configurations. This time includes the period required to transfer about two hundred megabytes of histogram data from the RODs to the analysis system, as well as to fit occupancy histograms for all 1536 channels on each of the several hundred modules. The parallel nature of the system means that the time required is not strongly dependant on number of modules. It is expected that when the base-line performance of the detector is well understood, the duration of tests can be shortened, for example by decreasing the number of triggers or the number of configurations tested, or by reducing the amount of information exported from the RODs.



**Figure 13.** Average input noise per chip for each of the four SCT barrels. The units on the x-axis are equivalent noise charge in electrons. These tests were performed “warm” with the mean of the modules’ temperature sensors (located near the ASICs) as indicated for each barrel.


Histograms of the chip-averaged input noise values found during barrel assembly are shown in figure 13. The noise values are consistent with single-module tests performed during module production. The modules have been designed to operate at colder temperatures (sensors at  $-7$  C) at which the input noise will be about 150 ENC lower because the noise decreases by about 5 ENC per degree [5]. The noise levels for endcap modules were also found to be consistent with expectations [35].

Performance confirmation tests during assembly were also used to identify any problematic channels such as those which were dead, had unacceptably high noise, or with other defects such as missing wire bonds which made them unusable in practice. For the barrel and for both the endcaps the fraction of fully functional channels was found to be greater than 99.7% – much better than the build specification of 99% good channels.

## 5.2 Commissioning and cosmic ray tests

At CERN, the SCT barrel and endcaps were each integrated [36] with the corresponding sections of the gaseous polypropylene-foil Transition Radiation Tracker (TRT) [37]. Further tests, including combined SCT/TRT cosmic ray studies, were then performed [38],[39]. These were the first large-scale tests of the SCT DAQ in physics mode.

For the barrel test, 468 modules, representing 22% of all modules on the four barrels, were cabled to make “top” and “bottom” sectors in azimuthal angle,  $\phi$ . There was no applied magnetic field and care was taken to reproduce, as far as possible, the service routing and grounding of the final setup in the ATLAS experimental cavern. All data were taken with the modules running “warm”, that is with their temperature sensors, located adjacent to the ASICs, at approximately 28 C.



**Figure 14.** Number of coincident hits as a function of the TIM trigger delay (in units of bunch crossing clocks).

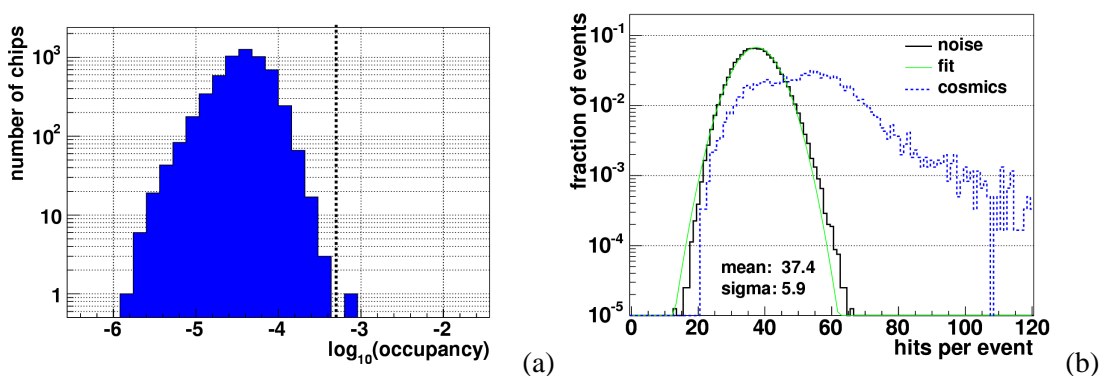
Cosmic rays were triggered using coincident signals from scintillators located above and below the barrel. Unlike during the assembly tests, the clock and trigger identifier information was distributed to the SCT TIM and to the TRT DAQ using ATLAS LTP modules. The resultant hit data were transferred from the SCT DAQ via the S-Link to a ROS and then written to disk. As well as using the cosmic trigger, noise data were recorded in physics mode under a variety of test conditions, using fixed frequency or random triggers sent from a master LTP.

To time-in the SCT with the cosmic ray and the TRT, the modules' relative timings were calculated from known differences in optical fibre lengths. The global delay was optimised using dedicated ROD monitoring histograms which recorded, as a function of the global delay, the number of coincident hits on opposite sides of each module (figure 14) with the front-end discriminator threshold set to its physics value of 1 fC. The front-end chips were configured to read out three consecutive time bins, and the TIM trigger delay was changed in steps of 75 ns (three bunch crossing clocks) to efficiently produce the effect of 25 ns steps in the TIM trigger delay. A 'hit' was defined to be coincident if there was a matching hit on any of the three chips on the opposing side of the module which have sensitivity in the physically overlapping region. The delay was then fine-tuned using the 270 ps TX fine delay on the BOCs to centre the peak of the coincidence signal on the middle of the centre of one of the time bins. After timing-in, hits from cosmic rays traversing the SCT and the TRT could be observed [38] on the event display.

In the noise tests, the occupancies obtained were not significantly different from those found for tests made on the individual barrels before integration. No significant change in the noise occupancy was observed when running concurrently with the TRT, when running at trigger rates from 5 Hz to 50 kHz, or for synchronous versus asynchronous triggering.

More than 450 thousand cosmic ray events and over 1.5 million synchronous noise events were recorded in the barrel tests. Figure 15a shows that the average noise occupancy was about an order of magnitude below the  $5 \times 10^{-4}$  specification even though the modules were not cooled to their design temperature. The distribution of the number of hits in noise runs is very well described by a Gaussian curve (figure 15b), so there is no evidence of correlated noise. By contrast, events which are triggered by cosmic rays have a long tail showing the expected correlated hits.

A further nine million physics-mode events were recorded in the synchronous operation of 246 modules during the commissioning of Endcap C. Again no significant change in noise occupancy was found for the endcap when integrated with the TRT compared to assembly tests, for synchronous versus asynchronous triggers, or for different trigger rates in the range 1 kHz to 100 kHz.



**Figure 15.** (a) Histogram of the logarithm of the average noise occupancy for each chip on the 468 barrel modules under test. The thick vertical dotted line indicates the performance specification of  $5 \times 10^{-4}$ . (b) Histogram showing the number of hits per event for a noise run (solid histogram) and for a cosmic-ray triggered run (dashed histogram). The curve is a Gaussian fit to the noise data. The cosmic-ray-triggered events frequently have multiple tracks so can contain a large number of hits per event.

Further information about the setup and results of these commissioning tests can be found in [38],[39]. In particular the hit-finding efficiency for cosmic-triggered tracks was found to be greater than 99% for all barrel layers after alignment.

## 6. Conclusions

The ATLAS SCT data acquisition system has been used extensively since the autumn of 2004 for performance testing and quality assurance during assembly and commissioning of the detector. Quality assurance tests in calibration mode, made simultaneously on groups of up to 672 modules (11% of the complete SCT), have helped ensure that the barrel and both endcaps were each ready for installation with more than 99.7% of channels performing to specification.

Commissioning tests in physics data-taking mode have demonstrated the continuing good performance of the SCT barrel and endcaps after integration with the TRT. Over ten million events have been successfully taken with synchronous triggers, demonstrating successful operation of both the DAQ system and the SCT detector with the final ATLAS trigger and data chain.

The complete DAQ system required for readout of the full SCT has been installed, integrated and tested. The system works well, and further development is expected as the system's performance, efficiency, and robustness are optimized in preparation for routine ATLAS data taking. The DAQ system will continue to monitor the SCT and will fine-tune the calibration as the detector's properties change during irradiation.

## Acknowledgements

We acknowledge the support of the funding authorities of the collaborating institutes, including the Netherlands Organisation for Scientific Research (NWO), the Spanish National Programme for Particle Physics; the Research Council of Norway; the Science and Technology Facilities Council of the United Kingdom; the Office of High Energy Physics of the United States Department of Energy; the United States National Science Foundation.

## A. Abbreviations

Alphabetical list of selected abbreviations:

ABCD3TA	ATLAS Binary Chip in DMILL, version 3 with TrimDacs, revision A
BOC	Back of crate card
CTP	The ATLAS central trigger processor
DAQ	Data acquisition
ENC	Equivalent noise charge
FE	Front-end
HOLA	High speed optical link for ATLAS
LHC	The CERN Large Hadron Collider
LTP	Local trigger processor module
MDSP	Master digital signal processor
RCC	Rod crate controller
ROD	Read-out driver
ROS	Read-out subsystem
RX	Receiver
S-Link	CERN serial optical link
SCT	Semiconductor tracker
SDSP	Slave digital signal processor
TTC	Timing, trigger and control
TIM	TTC interface module
TrimDAC	DACs which allow correction of individual channel thresholds
TRT	Transition Radiation Tracker
TX	Transmitter

## References

- [1] ATLAS collaboration, *ATLAS Inner Detector Technical Design Report*, CERN/LHCC/97-16, CERN/LHCC/97-17, [http://atlas.web.cern.ch/Atlas/GROUPS/INNER\\_DETECTOR/TDR/tdr.html](http://atlas.web.cern.ch/Atlas/GROUPS/INNER_DETECTOR/TDR/tdr.html) (1997).
- [2] A. Abdesselam et al., *The barrel modules of the ATLAS semiconductor tracker*, *Nucl. Instrum. Meth. A* **568** (2006) 642.
- [3] A. Abdesselam et al., *The ATLAS SemiConductor tracker end-cap module*, *Nucl. Instrum. Meth. A* **575** (2007) 353.
- [4] A. Ahmad et al., *The silicon microstrip sensors of the ATLAS semiconductor tracker*, *Nucl. Instrum. Meth. A* **578** (2007) 98.
- [5] F. Campabadal et al., *Design and performance of the ABCD3TA ASIC for readout of silicon strip detectors in the ATLAS semiconductor tracker*, *Nucl. Instrum. Meth. A* **552** (2005) 292.
- [6] ATLAS collaboration, *ATLAS high-level trigger, data-acquisition and controls : Technical Design Report*, CERN-LHCC-2003-022 (2003) <http://cdsweb.cern.ch/record/616089>.
- [7] [http://atlas.web.cern.ch/Atlas/GROUPS/FRONTEND/documents/Crate\\_Technical\\_Specification\\_final.pdf](http://atlas.web.cern.ch/Atlas/GROUPS/FRONTEND/documents/Crate_Technical_Specification_final.pdf).
- [8] B.G. Taylor, *Timing distribution at the LHC*, in *8<sup>th</sup> Workshop on Electronics for LHC Experiments* Electronics for LHC Experiments, Colmar, France , 9-13 Sep 2002, pp 63-74.

- [9] <http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/ReadOut/>.
- [10] <https://edms.cern.ch/document/551992/2>.
- [11] T. Vickey, *A Read-out Driver for Silicon Detectors in ATLAS*, Proceedings 12th Workshop on Electronics for LHC and Future Experiments, Valencia, Spain (2006).
- [12] <http://hsi.web.cern.ch/HSI/S-Link/>.
- [13] A. Abdesselam et al., *The Optical Links for the ATLAS SemiConductor Tracker*, 2007 JINST 2 P09003.
- [14] M-L. Chu et al., *The off-detector opto-electronics for the optical links of the ATLAS SemiConductor Tracker and Pixel detector*, *Nucl. Instrum. Meth. A* **530** (2004) 293.
- [15] T. Toifl and R. Vari, *A 4-Channels Rad-Hard Delay Generator ASIC with 1 ns Minimum Time Step for LHC Experiments*, Workshop on Electronics for LHC Experiments, CERN/LHCC/98-36 (1998) <http://sunset.roma1.infn.it:16080/prode/leb98.pdf>.
- [16] <http://hsi.web.cern.ch/HSI/S-Link/devices/hola/>.
- [17] J.M. Butterworth et al., *Timing, Trigger and Control Interface Module for ATLAS SCT Read Out Electronics*, ATL-INDET-99-018.
- [18] <http://ttc.web.cern.ch/TTC/intro.html>.
- [19] [http://www.cern.ch/TTC/TTCrx\\_manual3.9.pdf](http://www.cern.ch/TTC/TTCrx_manual3.9.pdf).
- [20] G. Bolla et al., *Wire-bond failures induced by resonant vibrations in the CDF silicon tracker*, *Nucl. Instrum. Meth. A* **518** (2004) 277.
- [21] T.J. Barber et al., *Resonant bond wire vibrations in the ATLAS semiconductor tracker*, *Nucl. Instrum. Meth. A* **538** (2005) 442.
- [22] A.J. Barr et al., *A Fixed Frequency Trigger Veto for the ATLAS SCT*, ATL-INDET-PUB-2007-010.
- [23] <http://proj-qpll.web.cern.ch/proj-qpll/images/manualTTCrq.pdf>.
- [24] ATLAS collaboration, *ATLAS level-1 trigger: Technical Design Report*, ATLAS-TDR-012, CERN-LHCC-98-014.
- [25] M.J. Palmer, *Studies of Extra-Dimensional Models with the ATLAS Detector*, University of Cambridge Thesis (2005).
- [26] B.J. Gallop, *ATLAS SCT readout and barrel macro-assembly testing and design of MAPS test structures*, University of Birmingham Thesis (2005).
- [27] B.M. Demirköz, *Construction and Performance of the ATLAS SCT Barrels*, University of Oxford Thesis (2007).
- [28] S. Kolos et al., *Experience with CORBA communication middleware in the ATLAS DAQ*, ATL-DAQ-2005-001.
- [29] F. Campabadal et al., *Beam Tests Of ATLAS SCT Silicon Strip Detector Modules*, *Nucl. Instrum. Meth. A* **538** (2005) 384.
- [30] P. Phillips, *Functional tests of the ATLAS SCT barrels*, *Nucl. Instrum. Meth. A* **570** (2007) 230.

- [31] L. Eklund et al., *Electrical tests of SCT hybrids and modules*, ATLAS note ATL-INDET-PUB-2007-006.
- [32] A.J. Barr, *Calibrating the ATLAS Semiconductor Tracker Front-end Electronics*, ATL-INDET-CONF-2006-001, *IEEE Nucl. Sci. Symp. Conf. Rec.* **2** (2004) 1192.
- [33] G. Viehhauser, *Macro-assembly of the ATLAS barrel SCT*, *IEEE Nucl. Sci. Symp. Conf. Rec.* **2** (2004) 1188.
- [34] D. Ferrère, *The ATLAS SCT endcap: From module production to commissioning*, *Nucl. Instrum. Meth.* **A 570** (2007) 225.
- [35] T.J. Jones, *The construction of the ATLAS semi-conductor tracker*, *Nucl. Instrum. Meth.* **A 569** (2006) 16.
- [36] H. Pernegger, *Integration and test of the ATLAS Semiconductor Tracker*, *Nucl. Instrum. Meth.* **A 572** (2007) 108.
- [37] T. Akesson et al., *Status of design and construction of the Transition Radiation Tracker (TRT) for the ATLAS experiment at the LHC*, *Nucl. Instrum. Meth.* **A 522** (2004) 131.
- [38] ATLAS SCT and TRT collaborations, *Combined SCT and TRT performance tests before installation*, in preparation.
- [39] B.M. Demirköz, *Cosmic tests and performance of the ATLAS SemiConductor Tracker Barrels*, *Nucl. Instrum. Meth.* **A 572** (2007) 43.