

2 **The Data Acquisition and Calibration System for the** **ATLAS Semiconductor Tracker**

4 A. Abdesselam^k, T. Barber^a, A.J. Barr^{k,†}, P. Bell^{c,1}, J. Bernabeu^q, J.M. Butterworth^p, J.R.
 Carter^a, A.A. Carter^m, E. Charles^r, A. Clark^e, A.-P. Colijnⁱ, M.J. Costa^q, J.M. Dalmau^m, B.
 6 Demirköz^{k,2}, P.J. Dervan^g, M. Donega^e, M. D'Onofrio^e, C. Escobar^q, D. Fasching^r, D.P.S.
 Ferguson^r, P. Ferrari^c, D. Ferrere^e, J. Fuster^q, B. Gallop^{b,l}, C. Garcia^q, S. Gonzalez^r, S.
 8 Gonzalez-Sevilla^q, M.J. Goodrick^a, A. Gorisek^{c,3}, A. Greenall^g, A.A. Grilloⁿ, N.P. Hesseyⁱ,
 J.C. Hill^a, J.N. Jackson^g, R.C. Jared^r, P.D.C. Johansson^o, P. de Jongⁱ, J. Joseph^r, C.
 10 Lacasta^q, J.B. Lane^p, C.G. Lester^a, M. Limper^l, S.W. Lindsay^g, R.L. McKay^l, C.A.
 Magrathⁱ, M. Mangin-Brinet^e, S. Martí i García^q, B. Mellado^r, W.T. Meyer^f, B. Mikulec^e,
 12 M. Miñano^q, V.A. Mitsou^q, G. Moorhead^h, M. Morrissey^l, E. Paganis^o, M.J. Palmer^a, M.A.
 Parker^a, H. Pernegger^c, A. Phillips^a, P.W. Phillips^l, M. Postranecky^p, A. Robichaud-
 14 Véronneau^e, D. Robinson^a, S. Roe^c, H. Sandaker^j, F. Sciacca^p, A. Sfyrla^e, E. Stanecka^{c,d}, S.
 Stapnes^j, A. Stradling^r, M. Tyndel^l, A. Tricoli^{k,4}, T. Vickey^r, J.H. Vossebeld^g, M.R.M.
 16 Warren^p, A.R. Weidberg^k, P.S. Wells^c, S.L. Wu^r

^a *Cavendish Laboratory, University of Cambridge, J.J. Thomson Avenue, Cambridge CB3 0HE, UK*

^b *School of Physics and Astronomy, University of Birmingham, Birmingham B15 2TT, UK*

^c *European Laboratory for Particle Physics (CERN), 1211 Geneva 23, Switzerland*

^d *Institute of Nuclear Physics PAN, Cracow, Poland*

^e *DPNC, University of Geneva, CH 1211 Geneva 4, Switzerland*

^f *Department of Physics and Astronomy, 12 Physics Hall, Ames, IA 50011*

^g *Oliver Lodge Laboratory, University of Liverpool, Liverpool, UK*

^h *University of Melbourne, Parkville, Vic. 3052, Australia*

ⁱ *NIKHEF, Amsterdam, The Netherlands*

^j *Department of Physics, P.O. Box 1048, Blindern, N-0316 Oslo, Norway*

^k *Physics Department, University of Oxford, Keble Road, Oxford OX1 3RH, UK*

^l *Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire OX11 0QX, UK*

^m *Department of Physics, Queen Mary University of London, Mile End Road, London E1 4NS, UK*

ⁿ *Santa Cruz Institute for Particle Physics, University of California, Santa Cruz, CA, USA*

^o *Department of Physics and Astronomy, University of Sheffield, Sheffield, UK*

^p *Department of Physics and Astronomy, UCL, Gower Street, London, WC1E 6BT*

^q *Instituto de Física Corpuscular (IFIC), Universidad de Valencia-CSIC, Valencia, Spain*

^r *University of Wisconsin-Madison, Wisconsin, USA*

[†] *Corresponding author, email: a.barr@physics.ox.ac.uk*

¹ *now at the School of Physics and Astronomy, University of Manchester, Manchester M13 9PL, UK*

² *now at the European Laboratory for Particle Physics (CERN), 1211 Geneva 23, Switzerland*

³ *now at the Jožef Stefan Institute and Department of Physics, University of Ljubljana, Ljubljana, Slovenia*

⁴ *now at the Rutherford Appleton Laboratory, Chilton, Didcot, Oxfordshire OX11 0QX, UK*



40 **The SemiConductor Tracker (SCT) data acquisition (DAQ) system will calibrate,**
42 **configure, and control the approximately six million front-end channels of the**
44 **ATLAS silicon strip detector. It will provide a synchronized bunch-crossing clock**
46 **to the front-end modules, communicate first-level triggers to the front-end chips,**
48 **and transfer information about hit strips to the ATLAS high-level trigger system.**
50 **The system has been used extensively for calibration and quality assurance during**
52 **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.

1 Introduction

54 The ATLAS experiment is one of two general-purpose detectors at CERN's Large Hadron
56 Collider (LHC). The SemiConductor Tracker (SCT) is a silicon strip detector and forms the
58 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 pseudo-rapidityⁱ $|\eta|$
 < 2.5 (Figure 1).

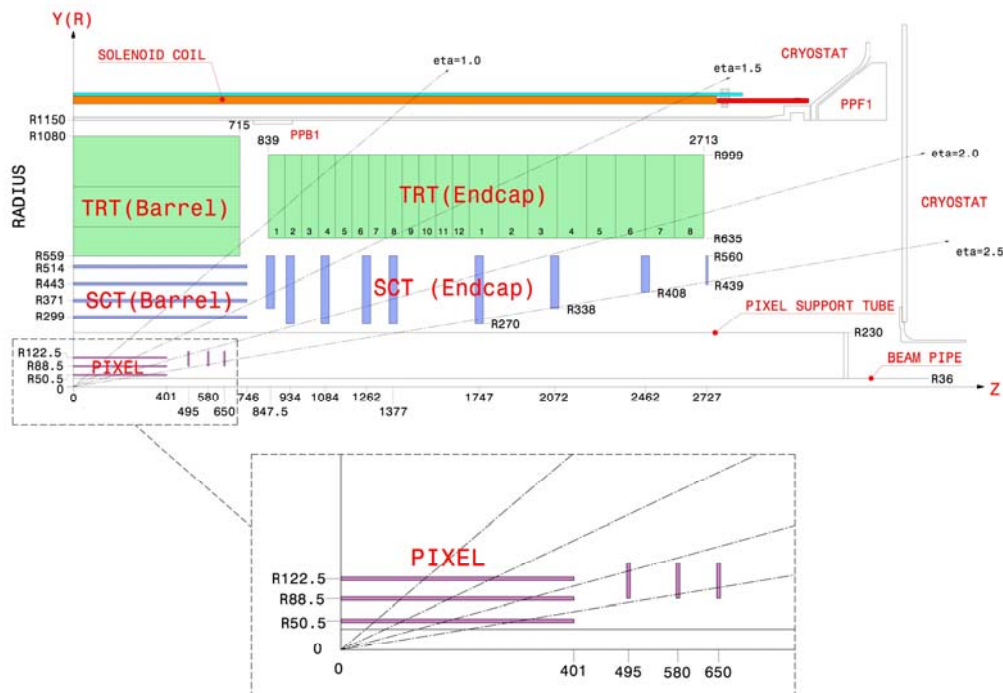
60 The complete SCT consists of 4088 front-end modules [1,2]. Each module has two planes of
62 silicon each with 768 active strips of p⁺ implant on n-type bulk [3]. The planes are offset by a
64 small stereo angle (40 mrad), so that each module provides space-point resolutions of 17 μm
 perpendicular to and 580 μm 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 [4]. Each of these chips is responsible for reading out 128 channels, so
 twelve are required for each SCT module.

66 The SCT is geometrically divided into a central barrel region and two endcaps (known as 'A' and
68 'C'). The barrel region consists of four concentric cylindrical layers (barrels). Each endcap
70 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.

72 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.

74 The role of the DAQ in calibrating the detector is equally important. The SCT uses a "binary"
76 readout architecture in which the only pulse-height information transmitted by the front-end chips
78 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.

ⁱ The pseudorapidity, η , is given by $\eta = -\ln \tan(\theta/2)$, where θ is the polar angle relative to the beam axis.



80 **Figure 1. Cross section of the ATLAS Inner Detector showing a quarter of the barrel and half of one**
 82 **of the two endcap regions. The SCT is within a Transition Radiation detector (TRT) and surrounds a**
Pixel detector[5].The dimensions are in mm.

84 The discriminator threshold must be set at a level that guarantees uniform, good efficiency while
 86 maintaining the noise occupancy at a low level. Furthermore the detector must maintain good
 88 performance even after a total ionizing dose of 100 kGy(Si) and a non-ionising fluence of 2×10^{14}
 neutrons/cm² 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 5×10^{-4} per channel even after irradiation.

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

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

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

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

92 During calibration, internal circuits on the front-end chips can be used to inject test charges.
 94 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
 96 volume of data obtained, and find an improved configuration based on the results.

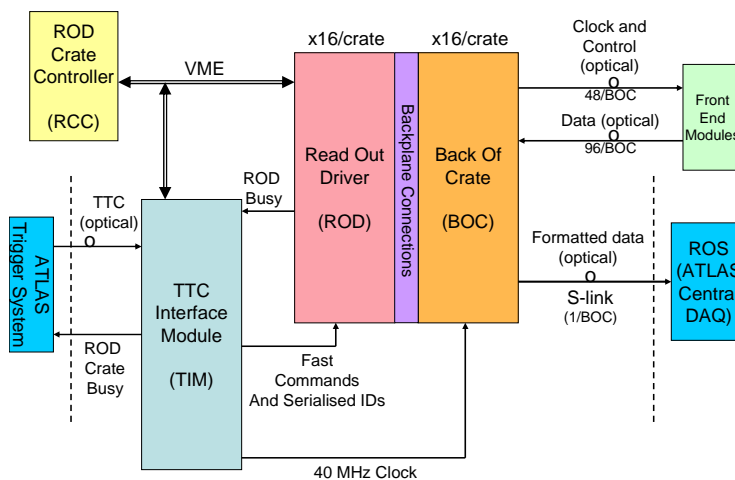
98 This paper is organized as follows. In Section 2 there is a description of the readout hardware.
 The software and control system are discussed in Section 3. In Section 4 there is a description of
 the calibration procedure. A review of the operation of the data acquisition system is given in

100 Section 5 together with some of the main results, covering both the confirmation tests performed
 102 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
 104 abbreviations used may be found in the appendix.

2 Off-detector hardware overview

106 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
 108 principal connections to the front-end modules, to the ATLAS central DAQ and between SCT-
 specific components are shown in Figure 2.

110 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
 112 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
 114 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
 116 (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
 118 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
 120 data acquisition functions within a crate.



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

124 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.
 126 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
 128 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.

130 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

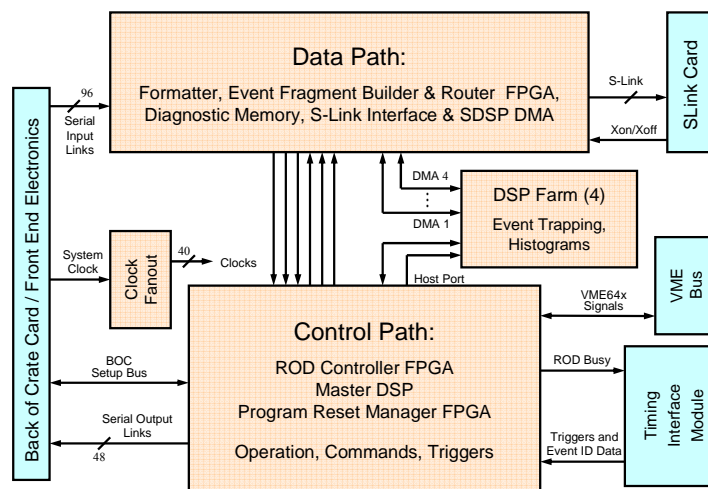
132 clock and SCT-specific triggers are taken from other sources. For most tests they are generated
internally on the RODs, but for tests which require synchronisation they can be sourced from the
134 SCT's local trigger processor (LTP) or from the TIM. The resultant data are not passed on to the
ROS, but the ROD monitoring functions still sample and histogram the events. The resultant
136 occupancy histograms are transferred over VME to the ROD Crate Controller and then over the
LAN to PC servers for analysis.

138 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
140 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
142 front end ASICs through the TTC system.

2.1 The Read-out Driver (ROD)

144 The Silicon Read-out Driver (ROD) [10] is a 9U 400mm deep VME64x electronics board. The
primary functions of the ROD are the front-end module configuration, trigger propagation and
146 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
148 streams. These commands can be first-level triggers, bunch-crossing (clock counter) resets, event
(trigger) counter resets, calibration commands or module register data. Each ROD board is
150 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
152 builds event fragments which are transmitted to the ROS via a high speed serial optical link
known as the S-Link [11].

154 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
156 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
158 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
160 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
162 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
164 assistance from the DSPs. The major data and communication paths on the ROD are shown in
Figure 3.



166

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

168 2.1.1 Operating Modes

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

176 2.1.2 Physics data-taking

177 After the data-path on the ROD has been set up, the event data processing is performed by the
 178 FPGAs without any intervention from the DSPs. Triggers issued from the LTP are relayed to the
 179 ROD via the TIM. If the S-Link is receiving data from the ROD faster than they can be
 180 transferred to the ROS, back-pressure will be applied to the ROD, thereby halting the
 181 transmission of events and causing the internal ROD FIFOs to begin to fill. Once back-pressure
 182 has been relieved, the flow of events through the S-Link resumes. In the rare case where the
 183 internal FIFOs fill beyond a critical limit, a ROD busy signal is raised on the TIM to stop triggers.
 184 The Router FPGA can be set up to capture events with a user-defined pre-scale on a non-
 185 interfering basis and transmit them to the farm of SDSPs. Histogramming these captured events
 186 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.

188 2.1.3 Calibration

189 When running calibrations, the MDSP serial ports can be used to issue triggers to the modules. In
 190 calibration mode the transmission of data through the S-Link is inhibited. Instead, frames of data
 191 (256 32-bit word blocks) are passed from the Router FPGA to the SDSPs using a direct memory
 192 access transfer. Tasks running on the SDSPs flag these transferred events for processing and
 193 subsequent histogramming. A monitoring task can be run on the SDSPs that is capable of parsing
 194 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

196 The ROD contains many components, and is required to perform many different operations in
198 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.

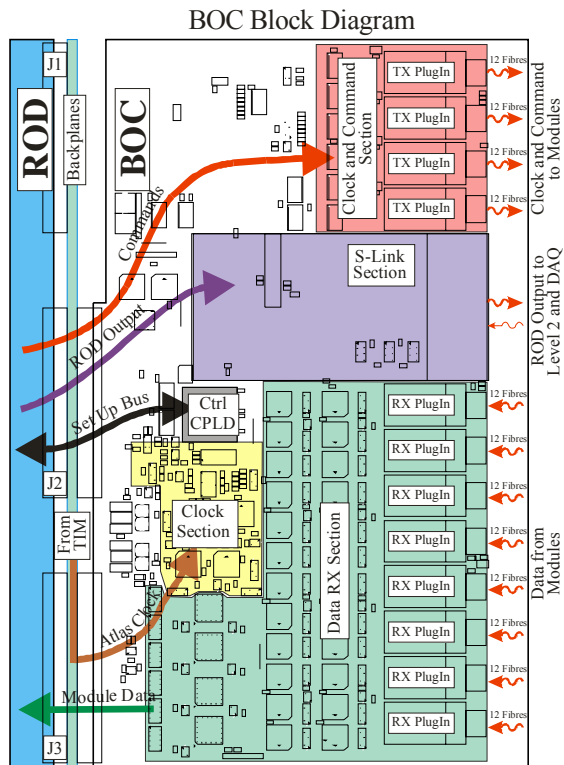
200 The communication registers are blocks of 32-bit words at the start of the DSP’s internal memory
202 which are regularly checked by the Master DSP (MDSP) inside the main thread of execution
running on the processor. The MDSP polls these registers, watching for requests from the RCC.
204 These registers 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
206 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
208 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
210 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
212 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
214 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
216 extra commands when required.

“Tasks” are DSP functions which execute over an extended period of time. These are started and
218 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
220 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
222 the MDSP and manages the sending of triggers, as well as changes in chip configuration and
histogram bin changes.

224 **2.2 Back of Crate card (BOC)**

The BOC transmits commands and data between the ROD and the optical fibre connections
226 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
228 and to its paired ROD. A block diagram of the function of the BOC is shown in Figure 4.



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

232 The front-end modules are controlled and read out through digital optical fibre ribbons. One fibre
 233 per module provides trigger, timing and control information. There are also two data fibres per
 234 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 [12].

236 On the BOC, each command for the front-end modules is routed via one of the four TX plug-ins
 237 as shown in Figure 4. Here the command is combined with the 40 MHz clock to generate a single
 238 Bi-Phase Mark (BPM) encoded signal which allows both clock and commands to occupy the
 239 same stream. Twelve streams are handled by each of four BPM12 chips [13]. The encoded
 240 commands are then converted from electrical to optical form on a 12-way VCSEL array [13]
 241 before being transmitted to the front-end modules via a 12-way fibre ribbon. The intensity of the
 242 laser light can be tuned in individual channels by controlling the current supplied to the laser
 243 using a digital to analogue converter (DAC) on the BOC. This is to cater for variations in the
 244 individual lasers, fibres and receivers and to allow for loss of sensitivity in the receiver due to
 radiation damage.

246 The timing of each of the outgoing signals from the TX plug-in can be adjusted so that the clock
 247 transmitted to the front-end modules has the correct phase relative to the passage of the particles
 248 from the collisions in LHC. This phase has to be set on a module by module basis to allow for
 249 different optical fibre lengths and time-of-flight variations through the detector. It is also
 250 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.

252 Incoming data from the front-end modules are accepted by the BOC in optical form, converted
 253 into electrical form and forwarded to the ROD. As each front-end module has two data streams
 254 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

256 array on the RX plug-in. These signals are then discriminated by a DRX12 chip [13]. The data for
each stream are sampled at 40 MHz, with the sampling phase and threshold adjusted so that a
258 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.

260 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-
262 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
264 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
266 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
268 precision alignment of the detector.

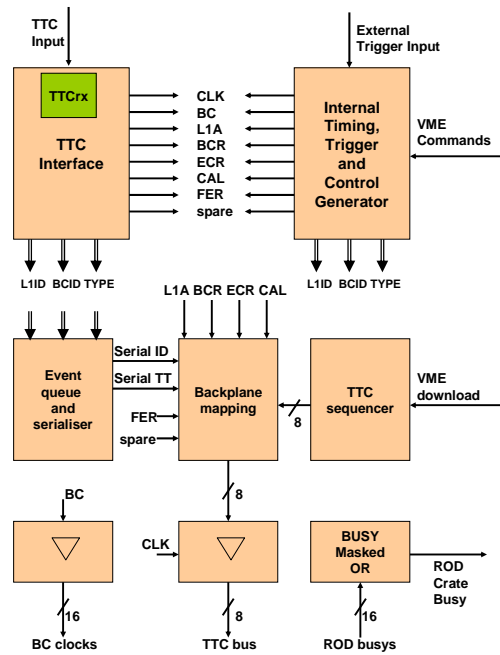
2.2.1 BOC Hardware Implementation

270 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
272 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
274 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
276 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
278 chip [14], 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
280 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
282 [15] daughter card.

2.3 TTC Interface Module (TIM)

284 The TIM [16] 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 [17] and distributes
286 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
288 from the ATLAS TTC system to the RODs, and sending the clock to the BOCs (from where it is
passed on to the RODs).



290

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

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

300 The incoming optical TTC signals are received on the TIM using an ATLAS standard TTCrx
 302 receiver chip [18], which decodes the TTC information into electrical form. In the physics mode
 304 the priority is given to passing the bunch-crossing clock and commands to the RODs in their
 306 correct timing relationship, with the absolute minimum of delay to reduce the latency. The TTC
 308 information is passed onto the backplane of a ROD crate with the appropriate timing. The event
 310 identifier is transmitted with a serial protocol and so a FIFO buffer is used in case of rapid
 312 triggers.

306 For tests and calibrations the TIM can, at the request of the local processor (RCC), generate all
 308 the required TTC information itself. It can also be connected to another TIM for stand-alone SCT
 310 multi-crate operation. In this stand-alone mode, both the clock and the commands can be
 312 generated from a variety of sources. The 40 MHz clock can be generated on-board, derived from
 314 an 80.16 MHz crystal oscillator, or transferred from external sources in either NIM or differential
 316 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.

318 A sequencer, using 8×32k 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

320 ‘sink’ (receiver RAM) of the same size is also provided to allow later comparisons of commands
and data sent to the RODs.

322 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
324 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.

326 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
[19]. The risk to the ATLAS SCT modules is considered to be small [20], even on the higher-
328 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
330 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
332 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
334 veto further triggers. The algorithm has been demonstrated to have a negligible effect on data-
taking efficiency [21].

336 **2.3.1 TIM Hardware Implementation**

338 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 [18] provide the bunch-crossing
340 clock and the trigger identification signals. On the TIM, a mezzanine board (the TTCrq [22])
allows an easy replacement if required.

342 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
344 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
346 marker. All the fast commands are clocked directly, without any local delay, onto the backplane
to minimise the TIM latency-budget.

348 On the TIM module, a combination of FastTTL, LVTTTL, ECL, PECL and LV BiCMOS devices
is used. The Xilinx Spartan IIE FPGA series were chosen as the programmable logic devices.
350 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.

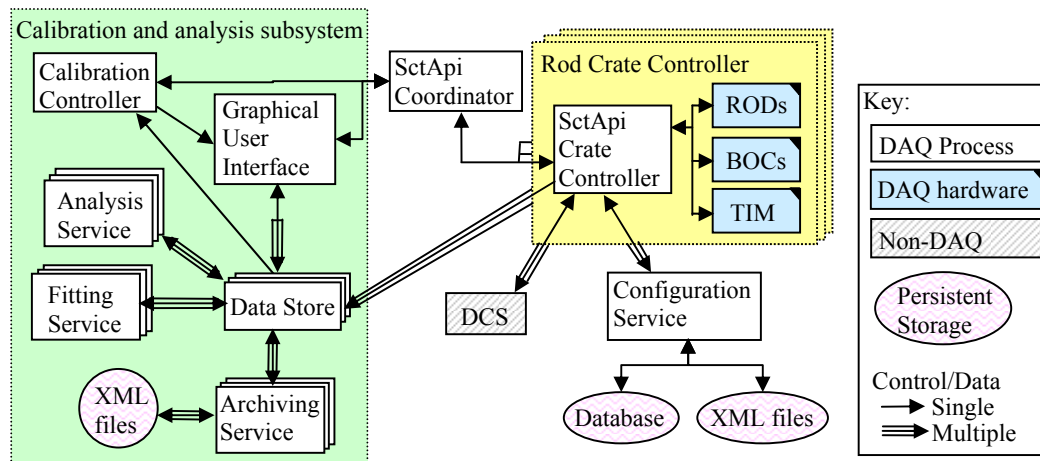
352 The TIM switches between different clock sources without glitches and, in the case of a clock
failure, does so automatically. To achieve this, dedicated clock-multiplexer devices have been
354 used. These devices switch automatically to a back-up clock if the selected clock is absent. Using
clock 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.

356 **3 Readout Software**

358 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
360 ROD/BOC pairs. The SctRodDaq software [23,24,25] controls this hardware and provides the
operator with an interface for monitoring the status of the front-end modules as well as initiating
362 and reviewing calibrations. The software can optimise the optical communication registers as well
as testing and calibrating the front-end ASICs.

364 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,

366 fixing parameters in turn. The results of each step of the calibration are analysed, and the relevant
 368 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.



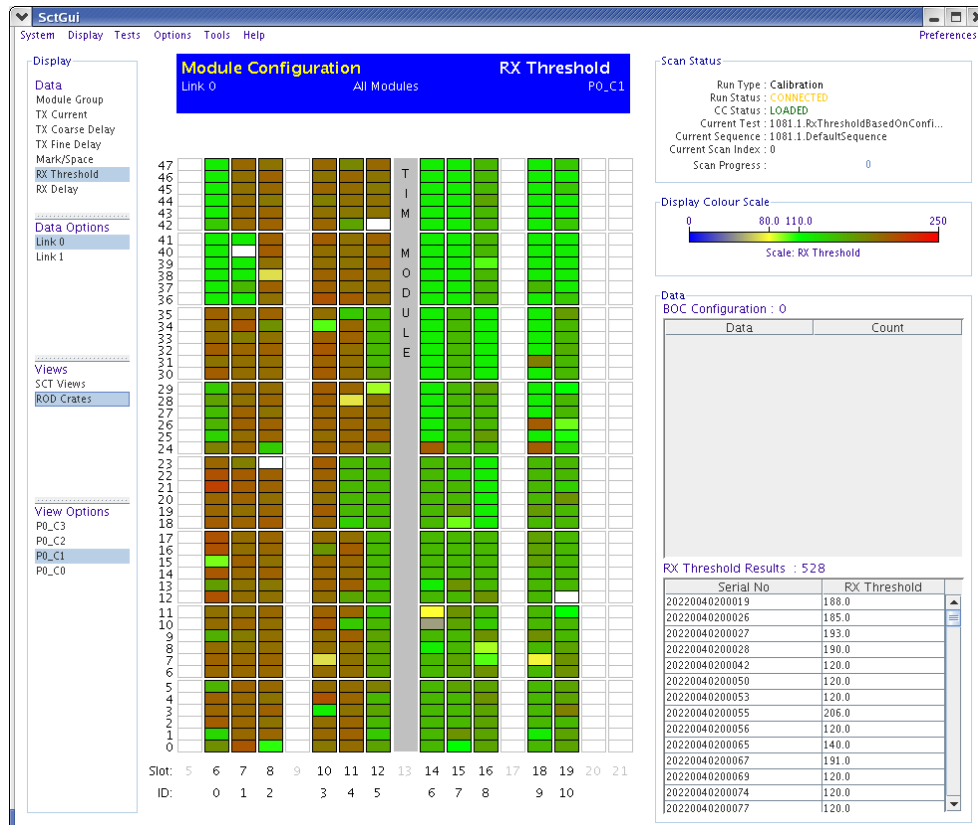
370 **Figure 6. Schematic control and data flow diagram for the SCT calibration and control system.**

372 A diagram of the main software components is shown in Figure 6. The readout software
 374 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
 376 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.

378 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
 380 fitting and analysis services perform data-reduction and calculate the optimal values for
 calibration parameters. The archiving services read data from transient objects and write them to
 382 persistent storage on disk. Inter-process communication is based on a number of ATLAS online
 software tools [26] one of which (IPC) provides a partition-based naming service for CORBA-
 384 compliant interfaces.

386 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
 388 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
 390 optimised for high performance since for most tests several fits are required for every read-out
 channel (see Section 4).

392 The front-end and DAQ system configuration can be stored either in an XML file or in a
 394 relational database. It is populated with data from previous calibrations, including quality
 396 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.



398
 400 **Figure 7. A view of the SCT graphical user interface. This display is showing the BOC receiver**
 402 **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.

4 Detector setup and calibration

404 Good front-end calibration is central to the correct operation of the detector because a large
 406 amount of pre-processing and data reduction occurs on the SCT's front-end ASICs. Each front-
 408 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
 (< 5×10^{-4}) and efficiency (> 99%) specifications can be satisfied. After the module is irradiated
 410 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
 412 requirements the channel thresholds need to be set within a more limited range [27]. To assure
 uniformity of threshold, every channel has its own 4-bit DAC (TrimDAC) which is used to
 414 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
 416 uncorrected channel-to-channel variations increase during irradiation.

4.1 Module communication optimisation

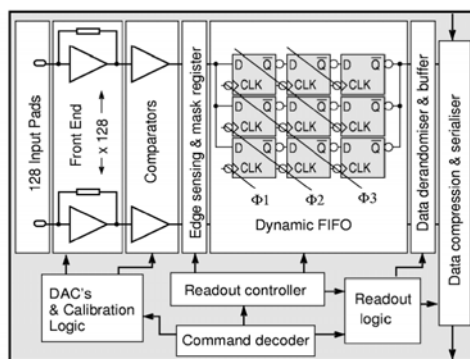
418 Before the front-end modules can be calibrated, the system must be configured to ensure reliable
communication between the modules and the BOCs.

420 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
422 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
424 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.

426 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
428 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,
430 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
432 [12,25].

4.2 Front-end calibration

434 Most of the calibration procedures are designed to set registers on the front-end ASICs. The most
important features of these ABCD3TA chips [4] can be seen in Figure 8. The analogue front-end
436 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
438 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
440 daisy-chained on the module. The data are compressed using an algorithm which only transmits
information about hit channels.




442 **Figure 8. Simplified Block Diagram of a SCT ABCD3TA ASIC [4].**

444 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
446 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.
448 To compensate for wafer-to-wafer variations in capacitance which can occur during ASIC
manufacture, correction factors were obtained from measurements on a number of test structures
450 on each wafer. The applied voltage step used during charge injection is adjusted in accordance
with these factors on a chip-to-chip basis.

452 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%


454 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
 456 triggers were sent per threshold point, and the charge injected was 1.5 fC.



458

Figure 9. Occupancy as a function of front-end discriminator threshold. (a) The shading scale shows
 460 the fraction of hits, as a function of the channel number (x-axis), and comparator threshold in mV (y-
 462 axis) for all channels on one side of a barrel module (6 ASICs; 768 channels). The front-end
 464 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.


To calibrate the discriminator threshold, the DAQ system initiates threshold scans for several
 466 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:
 468 $y = c_2 + c_0(1 + \exp(-x/c_1))^{-1}$, 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
 470 the response curve. The noise before amplification can be calculated by dividing the noise after
 amplification by the gain [28]. The gain and the noise are usually quoted at 2 fC input charge.



472


474 **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.**

476 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.
478 Using the results, an algorithm chooses optimal trim values, which reduce the channel-to-channel variations in the threshold (Figure 11).



480 **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**
482 **(b) after optimising the TrimDAC registers.**

484 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.



486

488

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.

490

492

494

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

496

498

500

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 [28,29,30]. 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.

502

5 Application and results

5.1 Barrel and endcap macro-assembly

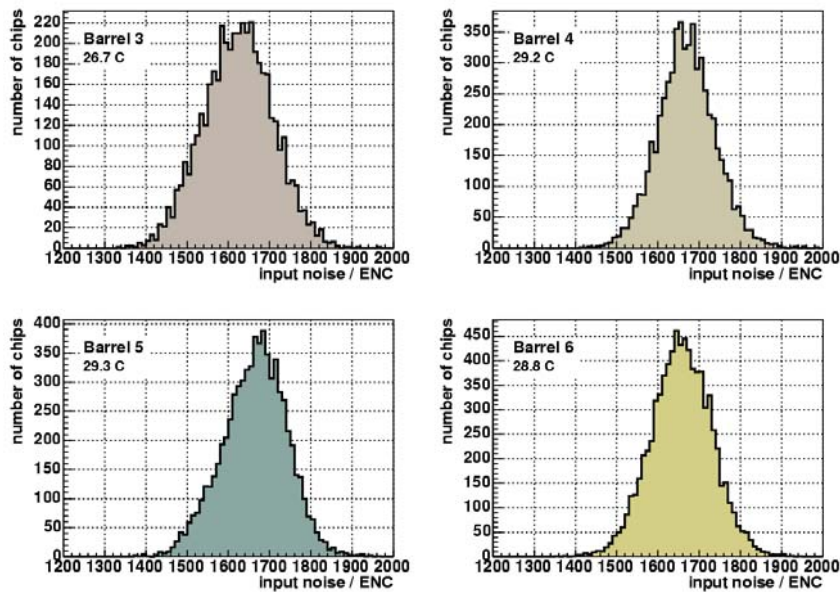
504

506

508

510

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 [1,2,28,30,31,32]. Groups of up to 672 modules (complete Barrel 6) 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.



512

514 **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.**

516 A typical time to run, analyse and feedback the results of a calibration test consisting of three
 518 front-end threshold scans is about 20 minutes, for a test in which 500 triggers are sent for each of
 520 100 different ASIC configurations. This time includes the period required to transfer about two
 522 hundred megabytes of histogram data from the RODs to the analysis system, as well as to fit
 524 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.

526 Histograms of the chip-averaged input noise values found during barrel assembly are shown in
 528 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
 530 which the input noise will be about 150 ENC lower because the noise decreases by about 5 ENC
 per degree [4]. The noise levels for endcap modules were also found to be consistent with
 expectations [33].


532 Performance confirmation tests during assembly were also used to identify any problematic
 534 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
 536 better than the build specification of 99% good channels.

5.2 Commissioning and cosmic ray tests

538 At CERN, the SCT barrel and endcaps were each integrated [34] with the corresponding sections
 of the gaseous polypropylene-foil Transition Radiation Tracker (TRT)35. Further tests, including
 540 combined SCT/TRT cosmic ray studies, were then performed [36,37]. These were the first large-
 scale tests of the SCT DAQ in physics mode.

542 For the barrel test, 468 modules, representing 22% of all modules on the four barrels, were cabled
 544 to make “top” and “bottom” sectors in azimuthal angle, ϕ . There was no applied magnetic field
 546 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 warm with module temperature
 sensors, located adjacent to the ASICs, at approximately 28 C.

548 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
 550 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
 552 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.



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

556 To time-in the SCT with the cosmic ray and the TRT, the modules’ relative timings were
 558 calculated from known differences in optical fibre lengths. The global delay was optimised using
 560 dedicated ROD monitoring histograms which recorded, as a function of the global delay, the
 562 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
 564 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
 566 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 [36] on the event display.

568 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
 570 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.

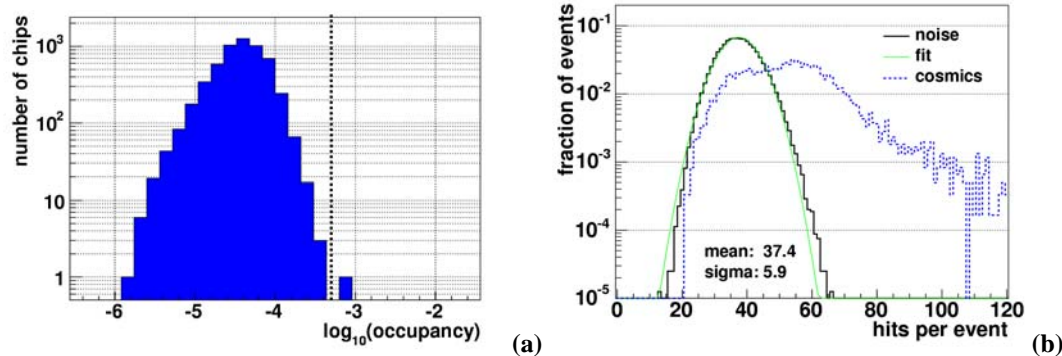


Figure 15. (a) Histogram of the logarithm of the average noise occupancy for each chip on the 468 barrel modules under test. The thick vertical dotted line indicates the performance specification of 5×10^{-4} . (b) Histogram showing the number of hits per event for a noise run (solid histogram) and for 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.

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 5×10^{-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 comic 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.

Further information about the setup and results of these commissioning tests can be found in [36-37]. 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

608 system will continue to monitor the SCT and will fine-tune the calibration as the detector's
 609 properties change during irradiation.

610 Acknowledgements

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

616 Appendix

Alphabetical list of selected abbreviations:

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

[1] A. Abdesselam et al, "The barrel modules of the ATLAS semiconductor tracker",
 Nucl.Instrum.Meth.**A568**:642-671 (2006)

[2] A. Abdesselam et al, "The ATLAS SemiConductor Tracker End-cap Module", Nucl. Instrum.
 Meth. **A575**:353-389 (2007)

[3] A. Ahmad et al, 'The silicon microstrip sensors of the ATLAS semiconductor tracker', Nucl.
 Instrum.Meth. **A578**:98-118 (2007)

[4] 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.**A552**:292-328 (2005)

[5] ATLAS Inner Detector Technical Design Report, CERN/LHCC/97-16, 1997

[6] "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

-
- [8] B.G. Taylor “Timing distribution at the LHC” in 8th 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] “A Read-out Driver for ATLAS Silicon Detector Modules”, in preparation
- [11] <http://hsi.web.cern.ch/HSI/S-Link/>
- [12] A. Abdesselam et al. “The Optical Links for the ATLAS SemiConductor Tracker”, JINST **2** P09003 (2007)
- [13] 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.**A530**:293 (2004)
- [14] T. Toifl, 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>
- [15] <http://hsi.web.cern.ch/HSI/S-Link/devices/hola/>
- [16] “Timing, Trigger and Control Interface Module for ATLAS SCT Read Out Electronics”, J. M. Butterworth et al. ATL-INDET-99-018
- [17] <http://ttc.web.cern.ch/TTC/intro.html>
- [18] http://www.cern.ch/TTC/TTCrx_manual3.9.pdf
- [19] G. Bolla et al, Wire-bond failures induced by resonant vibrations in the CDF silicon tracker, Nucl.Instrum.Meth. **A518**, 227 (2004)
- [20] T. J. Barber et al, “Resonant bond wire vibrations in the ATLAS semiconductor tracker”, Nucl.Instrum.Meth. **A538**, 442-457 (2005)
- [21] A. J. Barr et al., “A Fixed Frequency Trigger Veto for the ATLAS SCT”, ATL-INDET-PUB-2007-010
- [22] <http://proj-qpll.web.cern.ch/proj-qpll/images/manualTTCrq.pdf>
- [23] M. J. Palmer, “Studies of Extra-Dimensional Models with the ATLAS Detector”, University of Cambridge thesis, Jan 2005
- [24] B. J. Gallop, “ATLAS SCT readout and barrel macro-assembly testing and design of MAPS test structures” University of Birmingham thesis, Sep 2005
- [25] B.M. Demirköz, "Construction and Performance of the ATLAS SCT Barrels" University of Oxford thesis, April 2007
- [26] S. Kolos et al, “Experience with CORBA communication middleware in the ATLAS DAQ” ATL-DAQ-2005-001
- [27] F. Campabadal et al., “Beam Tests Of ATLAS SCT Silicon Strip Detector Modules” Nucl. Instrum.Meth. **A538** (2005) 384.
- [28] P. Phillips, “Functional tests of the ATLAS SCT barrels”, Nucl.Instrum.Meth.**A570** 230-235 (2007)
- [29] L. Eklund et al, “Electrical tests of SCT hybrids and modules”, ATLAS note ATL-INDET-PUB-2007-006

-
- [30] A.J. Barr, "Calibrating the ATLAS Semiconductor Tracker Front-end Electronics", ATL-INDET-CONF-2006-001, IEEE NSS Conference Records, 16-22 Oct. 2004, 1192-1195
- [31] G. Viehhauser, Conference Record of the 2004 IEEE Nuclear Science Symposium, Rome, Italy, 16-22 October 2004 (2004), pp. 1188-1191.
- [32] D. Ferrère, "The ATLAS SCT endcap: From module production to commissioning" Nucl.Instrum.Meth. **A570** 225-229 (2007).
- [33] T.J. Jones, "The construction of the ATLAS semi-conductor tracker", Nucl.Instrum.Meth. **A569** 16-20 (2006)
- [34] H. Pernegger, "Integration and test of the ATLAS Semiconductor Tracker", Nucl.Instrum.Meth. **A572** 108-112 (2007)
- [35] 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. **A522** (2004) 131
- [36] "Combined SCT and TRT performance tests before installation", in preparation.
- [37] B.M. Demirköz, "Cosmic tests and performance of the ATLAS SemiConductor Tracker Barrels", Nucl.Instrum.Meth. **A572** 43-47 (2007)