

Construction and Performance of the ATLAS SCT Barrels and Cosmic Tests

by

Bilge Melahat Demirköz

Balliol College

B.S., Massachusetts Institute of Technology (2001)

S.M., Massachusetts Institute of Technology (2004)

Submitted to the Department of Physics for
the degree of


Doctor of Philosophy

at

Balliol College

OXFORD UNIVERSITY

April 2007



© Bilge Melahat Demirköz, MMVII. All rights reserved.

The author hereby grants to Oxford University permission to reproduce and distribute publicly paper and electronic copies of this thesis document in whole or in part.

Construction and Performance of the ATLAS SCT Barrels and Cosmic Tests

by

Bilge Melahat Demirköz

Submitted to the Department of Physics
on April 17, 2007, in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

Abstract

ATLAS is a multi-purpose detector for the LHC and will detect proton-proton collisions with center of mass energy of 14TeV. Part of the central inner detector, the Semi-Conductor Tracker (SCT) barrels, were assembled and tested at Oxford University and later integrated at CERN with the TRT (Transition Radiation Tracker) barrel. The barrel SCT is composed of 4 layers of silicon strip modules with two sensor layers with $80\mu\text{m}$ channel width. The design of the modules and the barrels has been optimized for low radiation length while maintaining mechanical stability, bringing services to the detector, and ensuring a cold and dry environment. The high granularity, high detector efficiency and low noise occupancy ($< 5 \times 10^{-4}$) of the SCT will enable ATLAS to have an efficient pattern recognition capability.

Due to the binary nature of the SCT read-out, a stable read-out system and the calibration system is of critical importance. SctRodDaq is the online software framework for the calibration and also the physics running of the SCT and has been developed and tested during construction and commissioning of the detector with cosmics. It reliably measures the SCT performance parameters for each of the 6.3×10^6 channels in the SCT, identifies defects and problematic modules and writes them to an offline database for access from Athena, the ATLAS offline software framework. This dataflow chain has been exercised during the cosmics run at CERN, where a 5×10^5 cosmics sample for the combined SCT and TRT detectors was collected with a scintillator based trigger. It is now being commissioned in the ATLAS pit.

Acknowledgments

I would like to thank the Particle Physics and Astronomy Council (PPARC) and Scottish Power for awarding me the Dorothy Hodgkin Award which funded me for the duration of my graduate studies and again to PPARC for awarding me a Long Term Attachment (LTA) position at CERN. I would like to thank the Oxford Physics department and Balliol College for providing funding which enabled me to go to conferences. I would like to thank Profs. Ulrich Becker and Susan Cooper for inspiring me and helping me start my studies at Oxford University.

In a collaboration like ATLAS, it is simply impossible to thank each and everyone for their ideas. Here simply, I would like to thank everyone who took the time to discuss things with me. Irregardless, there will be unspoken heroes.

I would like to thank the team that I am part of, for it has been an honor to work with them: all members of the SctRodDaq and readout hardware development team and everyone involved in barrel macro-assembly and integration. Especially, Drs. Alan Barr, Chris Lester, Bruce Gallop, Dave Robinson, Matt Warren and John Hill have provided insight into the inner workings of the SCT and have made a creative and constructive group to work with. I would especially like to thank Peter Phillips whose help and strong support made this thesis possible. I would also like to thank *Bugs Bunny*, now a necessary figure in daily SCT operations, for surviving all attempts on his life.

I would like to thank Drs. Tony Weidberg and Georg Viehhauser who have given me wise advice and helped me navigate through a few rough seas. I would also like to thank Drs. Heinz Pernegger, Janet Carter, Mike Tyndel and Pippa Wells for their enthusiasm and leadership has been an inspiration. Drs. Nigel Hessey and Todd Huffman have both been wonderful as thesis committee members. I thank them both very much for their efforts in finalizing this thesis.

Last but not least, I would like to thank my friends and my family.

This thesis is dedicated to the memory of Prof. Gian-Carlo Rota, friend and mentor.

Author's Contribution

The work presented in this thesis is necessarily the work of a large number collaborations. Here I will try to briefly describe my specific contributions.

My work began with the start of macro-assembly of modules onto barrels at Oxford in October 2004. Since I was exempted from taking classes, I started immediately running cabling, fibring and testing shifts on the barrels on a daily basis, until September 2005. Testing involved learning how to operate the DCS (Detector Control System), monitor the cooling system and interlocks as well as using the testing software, called SctRodDaq. Initially, I wrote the XML-configurations for the barrels and all MURs (Minimum Unit of Readout) on the barrels. Daily testing required new configurations everyday. I also wrote the testing guidelines and updated them regularly, documenting software and hardware problems and solutions on the SctWiki. After becoming an “expert” user of SctRodDaq, I helped improve the testing software. I wrote several tests and analyses for pick-up, common-mode noise diagnosis, such as CounterErrorTest, DoubleTriggerNoiseTest and OccupancyPerEvent. I found that there were a few cases of lightleaks from opto-packages onto some modules on Barrel 3 which increased the noise. A fix was found for other barrels that were assembled later. I refined the analysis of some tests such as NMaskTest, RxThresholdTest and NPtGainTest. I improved some tools to define defects and modified the software so that pipeline defects can be updated after a PipelineTest. I identified a few modules which were swapped during the macro-assembly process from their trimming “fingerprints.”

As I moved to CERN in September 2005, with the last assembled barrel, I worked on the reception testing of barrels at CERN and started working on preparing the barrels for a combined cosmics test with the TRT, while the barrels were being integrated. I wrote a new package (ISSummaryWriter) to put NPtGain and Noise Occupancy test results into a form that can be easily inserted into a database and for the database, wrote a translation method in the configuration service to be able to code the offline identifier for a given module. I worked on modifying CDI (Conditions

Database Interface) for SCT requirements and used CDI later to put two data sets into COOL during cosmics commissioning.

In view of improving SctRodDaq for the final system, I ran the test system, H8Testbox, to debug several versions of DSP code, for faster and on-ROD histogramming. I wrote a new package (ISTestResult) which makes all the test summaries available in IS (InformationService) under the general ATLAS software framework. The SctGui now reads this information instead of text summaries and this improved the speed of the calibration procedure significantly.

Towards the test with the TRT, I worked on combining the readout of the SCT with the general ATLAS-wide software framework and I helped setup the scintillator system for the trigger of this combined test. I wrote a dedicated DAQ to read out TDC/ADC information from the scintillators (over VME) and insert this information into the ATLAS bytestream during physics datataking, also in the same readout cycle with the SCT and the TRT. Then I setup the timing and trigger chain for the combined run, paying special attention to the distribution of the signals for the relative and absolute timing of the detectors. After working on understanding the trigger latency, I had the idea to count coincidences while scanning through trigger delay to time-in the SCT to within 5nsec. For diagnosing problems quickly during the cosmics running, I improved a light-weight bytestream converter for the online system, which decodes the information from the ATLAS readout chain so that it can highlight problematic modules, or desynchronisation problems. This new version of the bytestream converter also could plug in to the TRTviewer, which is a light-weight online event displayer for the InnerDetector, originally intended for the TRT, but was expanded later to cover the SCT. After running commissioning shifts of the SCT sectors and running cosmics, I helped debug problems experienced during the cosmics running, such as ROS (ReadOut Sub-system) memory-limitations, which is part of the ATLAS event-building hardware.

In addition to my duties with the SCT barrel commissioning, I helped debug problems during the reception testing of both of the SCT endcaps at CERN and during the cosmics running of a sector of Endcap-C. I discovered that Endcap-C had

been running in the edge-mode while taking cosmics data instead of level-mode which increased the noise occupancy by ten-fold.

After the barrel cosmics, I have been working on moving my online expertise to the offline by helping the monitoring experts implement the coincidence timing histograms offline. With this method, Endcap-C was timed in. I have worked on decoding the configuration information from the online databases in the offline software to use it in the module efficiency calculations.

Contents

1	Introduction	25
1.1	LHC: The Large Hadron Collider	26
1.2	ATLAS: A Toroidal LHC Apparatus	27
1.2.1	ATLAS Coordinate Systems	27
1.2.2	Physics Goals	29
1.2.3	Energy loss and particle identification	30
1.2.4	Magnet System	32
1.2.5	Inner Detector	33
1.2.6	Calorimeter System	40
1.2.7	Muon System	43
1.2.8	Trigger and Data Acquisition System	45
1.2.9	Offline Processing	48
2	SCT: Semi-Conductor Tracker	49
2.1	Silicon Detectors	49
2.1.1	Silicon Detectors in a radiation environment	51
2.1.2	Design requirements of the SCT module	51
2.2	SCT barrel modules	53
2.2.1	SCT wafers	54
2.2.2	Radiation damage to an SCT module	58
2.2.3	Principle of binary read-out	59
2.2.4	Front-end electronics	61
2.2.5	Read-Out of modules	63

2.3	Detector control and readout	65
2.3.1	Optical Communication with Modules	65
2.3.2	Detector Control System	67
2.3.3	Off-detector readout	70
2.3.4	Configuration of the SCT system	78
3	SCT Barrels: Production to Commissioning	85
3.1	Calibration of binary readout	85
3.2	Module Production	87
3.3	Macro-Assembly and testing at Oxford University	90
3.3.1	SCT Hardware	92
3.4	SCT Online Software	93
3.4.1	TDAQ	94
3.4.2	SctRodDaq	98
3.5	Data Acquisition Runs	105
3.5.1	Calibration Runs	105
3.5.2	Physics Mode Runs	106
3.5.3	ATLAS Event Building and Event Format	107
3.6	SCT Calibration	108
3.6.1	SCT Scans and Tests	110
3.6.2	ModuleProbe	111
3.6.3	TxCurrentTest	114
3.6.4	RxDelayTest	115
3.6.5	RxThresholdTest	116
3.6.6	2D BOC Scan	119
3.6.7	MarkSpaceRatioTest	120
3.6.8	NMaskTest	121
3.6.9	PipelineTest	122
3.6.10	FullByPassTest	123
3.6.11	Counter Error Test	124

3.6.12	Strobe Delay Test	124
3.7	Analog Tests	127
3.7.1	Threshold scans	127
3.7.2	Occupancy per Event	128
3.7.3	Three Point Gain Test	131
3.7.4	Response Curve	135
3.7.5	TrimRange	142
3.7.6	Noise Occupancy Test and Occupancy Per Event (OPE)	146
3.7.7	Double Trigger Noise Test	149
3.7.8	Synchronous Triggers Noise Test	151
3.8	Barrel Test Results	157
4	SCT Conditions Database	163
4.1	Implementation of a realistic detector	163
4.1.1	SCT Conditions Databases	164
4.1.2	SCT DAQ Conditions Database	165
4.1.3	The SCT IS servers for the conditions database	166
5	Cosmic Tests of the ATLAS Inner Detector Barrels	171
5.1	Description of the Cosmic Setup	171
5.1.1	Measurement of the phase of the trigger with respect to the system	178
5.1.2	Timing Results	181
5.2	Trigger Distribution and Timing	184
5.3	Physics Data Taking	192
5.4	Results	201
6	Conclusions	211
A	Appendix	213

List of Figures

1-1	A schematic drawing of the ATLAS detector	28
1-2	A 3-dimensional cut-away view of the ATLAS inner detector.	34
1-3	Position and η -coverage of the inner detector components, [2].	40
1-4	A cutaway of one half of the calorimeter system	42
1-5	A cutaway view of the 3 barrel layers and 3 wheels of each end-cap.	44
1-6	Timing and trigger distribution to the SCT modules and ATLAS read-out schema.	47
2-1	A barrel module shown in a test-box with a test connector.	53
2-2	A schematic cross-section of an SCT sensor.	55
2-3	The reconstructed pulse shape due to beam particles from the test beam for a non-irradiated and an irradiated module.	57
2-4	Noise distribution and Landau charge deposit distribution for a MIP at normal incidence.	60
2-5	Signal processing chain within the ABCD chip	62
2-6	Different chip by-passing schemes for normal and modified modules	66
2-7	The principle of bi-phase mark(BPM) encoding.	67
2-8	A photo of two Low Mass Tapes and doglegs.	68
2-9	Schematic drawing of a clockwise TX redundancy loop.	80
3-1	Occupancy as a function of front-end discriminator threshold.	86
3-2	The distribution for calibration factors for all chips on all barrels.	88
3-3	A photo of completed Barrel 6.	91
3-4	The data-flow in SetRodDaq while in calibration mode.	95

3-5	SCT GUI example snapshot, showing RxThreshold values for one ROD crate.	100
3-6	The schematic for the control and monitoring of the histogramming task up to the stage where the EventData is published.	102
3-7	The schematic for the control and monitoring of the histogramming task and the processing of the histogram data, [71].	103
3-8	A TxCurrent scan test result for one Rx link.	115
3-9	A RxDelay scan test result for one Rx link.	116
3-10	An oscilloscope picture of the optical signal from a VCSEL.	117
3-11	An RxThreshold scan test result for one Rx link.	118
3-12	An RxThreshold scan test result for a slow turn-on VCSEL.	118
3-13	A 2-dimensional BOC scan for one Rx link.	120
3-14	A MarkSpaceRatio test result for one Rx link.	121
3-15	An NMask Test result for one side of the module.	122
3-16	A FullBypass Test result for one side of the module.	124
3-17	A Counter Error Test result for both Rx links.	125
3-18	A Strobe Delay Test result for one link.	126
3-19	A Strobe Delay fit result for the average of one chip.	127
3-20	The overlaid threshold scan fits for all channels on two chips with different trim ranges.	130
3-21	Occupancy per event histogram for one chip for an NMask scan. . .	131
3-22	The slopes in offset, dependant on module clock and command source.	134
3-23	A quadratic fit to the 10 point gain scan for one chip	136
3-24	The analysis plots for the gain and noise for both sides of one module from a 10 point gain test.	137
3-25	The measured V_{t50} values for all chips for different injected charges in the sector plotted against their offsets.	138
3-26	The gain measured at 2 fC for all chips of the cosmics sector.	140
3-27	The noise occupancy measured at 1 fC before and after the configuration is updated with a new response curve.	141

3-28	Histogram of trim range settings for all chips in the SCT barrels.	143
3-29	The effect of trimming on offset.	144
3-30	The effect of trimming on a raw scan result.	145
3-31	An OPE histogram for a chip.	147
3-32	Noise Occupancy test result plot and an OPE badness plot for a typical chip.	152
3-33	A chip with a high OPE badness at low threshold.	152
3-34	A chip with a high OPE badness at high threshold.	153
3-35	Histogram of the worst badness parameter for each chip from Barrel 3.	153
3-36	Evidence for light leaks on the barrels from opto-packages.	154
3-37	A barrel schematic in the r - ϕ plane.	154
3-38	Evidence of light leaks from a plot of occupancy versus time.	155
3-39	Electrical pickup candidates.	155
3-40	The double trigger noise occupancy plot for an electrical pickup candidate.	156
3-41	The histogram of the number of bad channels per module on all barrels.	159
3-42	The comparison of average input noise for every chip for Barrel 3 during Oxford and CERN tests.	160
3-43	Histogram of average input noise for every chip for all barrels.	161
4-1	A Pipeline test summary object as seen in the IS GUI.	169
5-1	The insertion of Barrel 3 into the other barrels and thermal enclosure.	172
5-2	A photo just as the TRT barrel was rolled over the SCT barrel.	174
5-3	Drawing showing the positions of the scintillators with respect to the TRT in the SR1 cosmics setup at CERN.	175
5-4	Schematic of the cosmics NIM trigger logic.	176
5-5	A drawing of the power supply crate mapping for the sector of the SCT that was read out during cosmics data taking.	177

5-6	Efficiency versus trigger time from the test beam for a non-irradiated (left) module biased with 250V and an irradiated (right) module biased with 450V.	179
5-7	Distribution of the time measured in the TDC for the meantimer signals of the three scintillators.	182
5-8	Distributions of the difference of the TDC measurements in the three scintillators	183
5-9	The trigger chain for cosmics, with the SCT LTP running as the master LTP.	185
5-10	A schematic drawing of a module showing which chips overlap and so might have hits in coincidence due to the stereo-angle between the two sides of the module.	190
5-11	The first coincidence histogram produced of coincidences of hits from two sides of each module, from the top sector, with the initial timing.	191
5-12	The histogram of single coincidence of hits from all modules versus bunch-crossings, made by scanning through the coarse trigger delay.	192
5-13	The physicsTxDelayScan progress in physics mode as shown on the SetGui.	193
5-14	The track of a cosmic muon traversing through the top sectors of the TRT and the SCT.	194
5-15	The track of a cosmic muon traversing through the TRT and the SCT sectors.	195
5-16	Run number versus the total number of events taken in one run.	197
5-17	The unbiased residual sigmas measured for the different layers and sides of the SCT modules for simulation and data using global χ^2 alignment.	202
5-18	The unbiased hit efficiency measured for the different barrel layers and sides from cosmic data, before and after alignment.	203
5-19	The unbiased SCT residuals for fast and slow particles.	205

5-20	The histogram of number of hits in an event for a cosmic run for all 3 time-bins.	207
5-21	The histogram of number of hits in an event for a noise run for all 3 time-bins.	208
5-22	The histogram of the noise occupancy difference between physics and calibration mode in standard deviations for each module in the cosmics sector.	209

List of Tables

1.1	Some basic parameters of the LHC at design luminosity, [106]	27
1.2	Main parameters of the Inner Detector.	33
1.3	The SCT barrel geometry.	36
1.4	The module positions for one LMT.	37
2.1	Some default DCS settings and allowed monitored values for the “ON” state of an SCT module.	71
2.2	The error codes issued by the ROD event formatter.	74
2.3	The module output format as sent by the router to the DSP and up the S-link.	75
2.4	The keys used in the module output format in Table 2.3.	75
2.5	The definition of the barrel MUR (Minimum Unit of Read-out) number: bsXX.	79
2.6	The attributes of the MUR (Minimum Unit of Read-out) element in the XML-based configuration files.	79
2.7	The attributes of the ROD element in the XML-based configuration database.	80
2.8	The attributes of the module element, in the XML-based configuration database.	82
2.9	The chip configuration register contents and default values.	83
2.10	The last five bits of the ABCD chip configuration register need to be changed if reading through one link.	84

3.1	Mapping of groups of LMTs to power-supply crates for Barrel 3 testing at macro-assembly.	92
3.2	Mapping from LMT to power-supply crate channels for part of Barrel 3.	93
3.3	Power-supply crate slot to card-mapping	93
3.4	Contents of the ROD Event header	108
3.5	List of all defects that can be issued by SctRodDaq with the meaning and action taken.	112
3.6	The values and their meaning of the results returned by SctApi after a probescan and the actions that the DAQ shifter might need to take to get all the modules returning events.	113
3.7	The table of known TTC links which do not work. All these modules when operated in redundant TTC mode return valid data.	114
3.8	List of all the known optical Rx links that do not work on the barrels.	129
3.9	A list of some known slope problems as identified by an N-point gain scan test and solutions.	135
3.10	Defect list from the latest full barrel tests: reception testing of Barrels 3, 5 and 6 at CERN and cold testing of Barrel 4 at Oxford.	158
4.1	Encoding of the Athena offline identifier	165
4.2	The SummaryData object for NPtGain test as defined in the OKS database.	168
4.3	The SummaryData object for NoiseOccupancy test as defined in the OKS database.	168
4.4	The generic ISSummaryData class as defined in the OKS database.	170
5.1	TDC and ADC mapping information for the scintillator measurements	180
5.2	Results of the fits to the meantimer time difference distributions.	181
5.3	The latency in the trigger chain up to the SCT front-end for physics running in the pit.	188
5.4	The latency in the cosmics trigger chain up to the SCT front-end for SR1 setup.	189

5.5	List of physics data taken during cosmics commissioning.	196
5.6	List of modules with edge detect turned on for the cosmics sector.	208
A.1	The abbreviated list of ATLAS source identifiers as defined in the eformat library of the TDAQ-common package.	213
A.2	A list of commonly used acronyms	214

Chapter 1

Introduction

The Large Hadron Collider (LHC) has a high discovery potential for new physics with a high center of mass energy of collisions ($\sqrt{s} = 14$ TeV) and a high design luminosity of 10^{34} cm⁻²s⁻¹. ATLAS, A Toroidal LHC Apparatus, will be the world's largest particle detector when the LHC starts operating in mid-2007 at CERN, in Geneva, Switzerland.

The semiconductor tracker (SCT) is of vital importance to ATLAS since it provides good tracking and momentum resolution up to a pseudo-rapidity η of 2.5 extending from a radius of 0.3 m to 0.52 m. The SCT has a central barrel system and two end-cap systems. The barrel system was assembled at CERN by inserting four barrels inside each other. Each barrel was macro-assembled at Oxford University, by putting modules onto a barrel shaped support structure. The end-caps have been macro-assembled at Liverpool and at NIKHEF and consist of 9 disks each. In total, there are 4088 silicon detector modules in the SCT, with 6 million channels, each providing a 1-bit binary signal at each bunch-crossing every 25 nanoseconds.

This thesis first briefly discusses some of the design and performance requirements of the ATLAS SCT in the LHC environment. Then it details my involvement in the SCT barrel construction and testing as well as the development of analysis tools to study its performance. In the next chapter, the cosmic commissioning of the detector and results are detailed. Then the implementation of a realistic SCT description and simulation in the ATLAS offline software framework is discussed.

1.1 LHC: The Large Hadron Collider

The LHC, [95], will be the highest energy and largest operational particle accelerator and collider in the world with the start of 14 TeV collisions data taking in 2008. The initial running period of 3 years will be at luminosities up to $10^{33} \text{ cm}^{-2}\text{s}^{-1}$. After a major upgrade to the collimation system, the design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ will be reached. The machine will be commissioned initially with 450 GeV beam and then with 7 TeV beam at low luminosity before reaching full luminosity.

At the beginning of the accelerator chain, a small linear accelerator and a subsequent booster injects protons into the PS (Proton Synchrotron), which accelerates them to 26 GeV. At this energy, the protons are injected into the SPS (Super Proton Synchrotron) which accelerates and injects them into the LHC at 450 GeV. In fully operational mode, LHC accelerates these particles to 14 TeV. In the fall of 2007, collisions at a center of mass energy of 900 GeV, as injected from the SPS, are foreseen for the commissioning phase of the LHC.

The LHC is being installed in the LEP tunnel which has a circumference of 27 km. Around the ring, there are 8 straight sections, each followed by a section which bends the particles. One straight section contains the RF cavities which accelerates the particles and another contains the beam dump. Two sections clean the beam halo and the other four straight sections each have an interaction point for the beams, where detectors are being installed: ATLAS, CMS (Compact Muon Solenoid), LHCb(LHC beauty experiment) and ALICE(A Large Ion Collider Experiment). ATLAS and CMS are large multi-purpose detectors optimized for discovery of new high energy particles, whereas LHCb will mainly focus on b-physics and ALICE will elucidate the physics of the quark-gluon plasma.

Some basic parameters of the LHC are listed in table 1.1. It is worthwhile to note that the biggest engineering challenge for the LHC is the design and production of the high field superconducting magnets which bend, squeeze and focus the beam and the necessary cryogenic system which uses super-fluid liquid helium.

Parameter	Value
Energy at collision	7.0 TeV
Energy at injection	0.45 TeV
Machine circumference	26658.833 m
Time between bunches	25 ns
Frequency of bunches	40.08MHz
Number of particles per bunch	1.15×10^{11}
Number of bunches per beam	2808
RMS bunch length	7.55 cm
RMS beam size at ATLAS and CMS	$16.7 \mu\text{m}$
Circulating beam current	0.582 A
Magnetic field strength	8.36 T
Dipole field at 7 TeV	8.33 T
Dipole magnet temperature	1.9K
Number of dipole magnets	1232
Number of quadrupole magnets	392

Table 1.1: Some basic parameters of the LHC at design luminosity, [106]

1.2 ATLAS: A Toroidal LHC Apparatus

ATLAS is a multi-purpose detector which can directly detect charged and most neutral particles that will be decay products of unstable standard model particles such as the Z and the W^\pm , the predicted Higgs boson and other particles from theories beyond the standard model. Hermeticity of the detector is an important design requirement because it allows for a measurement of missing momentum for neutral and invisible particles where a direct detection is virtually impossible, such as the neutrino.

1.2.1 ATLAS Coordinate Systems

The ATLAS detector, shown in Fig. 1-1, is housed at interaction point 1 on the LHC ring, close to the CERN main site. The vector that points from the interaction point to the centre of the LHC ring defines the “+x” axis and the “+y” axis points upwards. The “+z” direction is along the beam axis and the “-z” side is known as the C-side and the “+z” side is known as the A-side for the naming conventions in the pit. Besides the standard Cartesian coordinate system, especially for physics analyses, a coordinate system with (r, ϕ, θ) is useful. Here, “r” is the transverse radius from the




Figure 1-1: A schematic drawing of the ATLAS detector showing the main detector components. Notice the scale from the physicists standing on the shielding.

beam-pipe and ϕ the azimuthal angle, measured from the x-axis. θ can be used to directly measure the angle away from the beam-pipe.

For describing tracks of particles in a detector, rapidity, y , is especially useful because it is invariant under longitudinal (in z) Lorentz boosts and in hadron collisions, the original longitudinal momentum of the interacting partons is not known. Rapidity is defined as

$$y = \frac{1}{2} \ln \frac{E + p_L}{E - p_L} \quad (1.1)$$

where E is the energy of the particle and p_L is the longitudinal component of the momentum of the particle. For a particle with zero rest mass, this equation is reduced to

$$\eta = -\ln \left(\tan \left(\frac{\theta}{2} \right) \right) \quad (1.2)$$

where η is the pseudo-rapidity. η is a good approximation for y in the relativistic limit. This parameter is convenient for describing the coverage of a detector. A high η coverage, meaning $\eta \gg 1$, means that a detector has good coverage in the forward regions.

1.2.2 Physics Goals

The Standard Model has been incredibly successful in explaining most particle physics measurements. There is, however, one missing ingredient to this recipe that has not yet been discovered and that is the Higgs boson, responsible for giving mass to all other particles by means of a symmetry breaking mechanism.

Despite the success of the Standard Model, it is by no means a complete description of observed physical phenomena. For example, it contains no treatment of gravity or general relativity nor does it include any mathematical mechanism to solve the hierarchy problem or to generate neutrino masses.

Furthermore, there are few – yet significant – unexplained cosmological observations, such as the apparent matter-antimatter asymmetry of the universe as well as the measured dark matter and dark energy content of the universe. The cosmological constant prediction is divergent due to vacuum fluctuations. These phenomena do not merely suggest but rather require more fundamental laws at work in the universe around us.

ATLAS can elucidate these topics by testing the Standard model by allowing for precision measurements and an exhaustive search for the Higgs boson and the search of physics beyond it. Here is a brief list of the main physics interests:

1. The Higgs boson: A main goal of ATLAS is to find this necessary, yet still missing piece in the Standard Model. The low mass limit from LEP searches is 114.4 GeV at 95% confidence level [137], with a theoretical upper bound of 1TeV. ATLAS and CMS can detect a Standard Model Higgs boson with high significance ($> 5\sigma$) within this entire mass range.
2. Precision measurements of the top quark and heavy bosons: Since the top quark

was found at Fermilab, there has been a yet unsuccessful pursuit to measure its couplings with as high precision as other particles. At LHC, the cross-section for $t\bar{t}$ production is ≈ 1 nb. A high statistics sample collected by ATLAS will yield a top mass measurement which is dominated by theoretical calculation uncertainties rather than statistical errors.

3. CP violation: A necessary ingredient in trying to explain the apparent matter anti-matter asymmetry of the universe is CP violation. ATLAS make a competitive measurement of CP violation in the B_q^0 system in the initial low luminosity running when b-quark identification will not be hindered by pile-up.
4. Search for a dark matter candidate and physics beyond the standard model: There are several different theories of what may be beyond the realm of the standard model. Theories such as SUSY (Super-Symmetry), Kaluza-Klein or extra-dimensions would produce several interesting particles and decay channels among them at least one particle which is a dark matter candidate. Such candidates can be discovered in ATLAS by missing E_T owing to the hermeticity of the detector.

1.2.3 Energy loss and particle identification

Before a discussion of the ATLAS detector, I would like to discuss how tracking detectors work on the principle of energy loss of a particle. Charged particles moving through matter lose energy due to basic physics processes such as bremsstrahlung and ionization, while photons lose energy through Rayleigh and Compton scattering and pair production. Each of these processes depends on the energy and mass of the particle as well as the properties of the material that the particle is traversing. In general, the mean rate of energy loss due to ionization (does not include bremsstrahlung) by a charged particle is given by the Bethe-Bloch formula [137]:

$$-\frac{dE}{dx} = Kz^2 \frac{Z}{A} \frac{1}{B^2} \left(\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 T_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right) \quad (1.3)$$

where T_{max} is the maximum kinetic energy which can be transferred to a free electron in a single collision, K is the constant $4\pi N_A r_e^2 m_e c^2$ and N_A , r_e , m_e and c are respectively, Avogadro's number, the classical radius of the electron, the mass of the electron and the speed of light. Z is the atomic number and A is the atomic mass of the material. β is the speed of the incident particle and γ is the Lorentz factor. z is charge of the incident particle and I is the mean excitation energy and $\delta(\beta\gamma)$ is the density effect correction to ionization energy loss.

A low energy particle loses energy fast compared to a relativistic particle, as it has time to interact with the material it is traversing. The rate at which the particle loses energy decreases as β^{-2} until it reaches a minimum at $\beta\gamma \approx 3$, where relativistic energy losses start to become appreciable and the energy loss rate increases logarithmically until the Fermi plateau. This minimum is important as a particle spends quite some time as a “minimum ionizing particle” or MIP before stopping in the material. Furthermore, if a detector can detect a MIP, then it can detect any other particle as surely, other particles will deposit more energy in the detector material. The number of electron ion pairs produced in an SCT module is discussed in chapter 2. The distribution of energy loss in a thin sheet of material is given by the Landau function, which can be approximated as a Gaussian at the peak, but has a long tail towards higher energies.

Most tracking detectors aim to detect the resulting ionization from energy loss of a particle in a particular medium. The exceptions are detectors based on the detection of scintillation light. All tracking detectors, except those, employ electric fields in order to drift, sometimes avalanche and multiply and eventually detect these ionization electrons. Detectors employ different media – gaseous, liquid and solid – to balance the critical detection efficiency versus radiation length that they place in front of subsequent detectors.

Radiation length, X_0 , is the distance in which an electron or positron with energy > 1 GeV, loses $1 - e^{-1}$ or 63.3% of its original energy by bremsstrahlung. Ignoring the material-dependent effects to first order, the dimensions of an absorber are quite often expressed in this unit.

It is important for detectors closest to the interaction point to have good efficiency while maintaining a low radiation length so that the momentum measurement does not degrade due to bremsstrahlung, and particles live long enough to reach the outer sections of the detector to obtain full tracking information. Furthermore, support structures and services as well as the necessary magnet system to bend the particles to provide momentum information, are under the same requirement of maintaining low radiation length.

1.2.4 Magnet System

The magnet and its support structure visually dominates ATLAS. There are two magnet systems: a central solenoid that surrounds the inner detector and an outer toroid system for the momentum measurement in the muon spectrometer. The magnet coils are made from aluminum stabilized NbTi superconductor, which will be cooled down to 4 K by a cryogenic liquid helium system.

The central solenoid is designed to be as thin as possible to minimize radiation length as the particles are on their way to the calorimeter. The solenoid is a 5.3 m long cylinder with a radius of 1.15 m, with a thickness of only 45 mm and provides a 2 T field for inner detector tracking

Placed on the outside of the calorimeter system, there are 8 superconducting coils with an outer diameter of 20 m and 26 m length in the barrel area, providing a toroidal magnetic field. These huge toroid magnets and their necessary support structure were the first components to be installed in the pit, as the rest of ATLAS fits within them. The end-caps are smaller and will be installed last to “close off” the detector. In each end-cap, there are 8 coils. Each end-cap is rotated by $\pi/4$ with respect to the barrel so that they fit in the cracks between the ends of the barrel coils. The end-caps are 5 m long.

Given its size and its 1300 ton weight, the magnet system of ATLAS is a considerable engineering challenge. The system stores 1600 MJ of energy when fully operational and protection against quenches when the system needs to dissipate this energy is of critical importance.

1.2.5 Inner Detector

The ATLAS inner detector has three components: an inner-most silicon pixel detector, followed by a silicon strip detector – the SCT – which is then surrounded by a transition radiation tracker – the TRT – [15] as seen in Figure 1-2.

Detector	Section	η coverage	Resolution per measurement (μm)	Layers	Area (m^2)	Channels ($\times 10^6$)
Pixel	B-layer	± 2.5	$r\phi = 12, z = 66$	1	0.2	16
	Barrel	± 1.7	$r\phi = 12, z = 66$	2	1.4	81
	End-cap	1.7 - 2.5	$r\phi = 12, r = 77$	3	0.7	43
SCT	Barrel	± 1.4	$r\phi = 16, z = 580$	4	34.4	3.2
	End-cap	1 - 2.5	$r\phi = 16, r = 580$	9	26.7	3.0
TRT	Barrel	± 0.7	170			0.1
	End-cap	0.7 - 2.0	170			0.32

Table 1.2: Main parameters of the Inner Detector.

The system requirements for the tracker for the LHC environment are various such as the cooling, noise performance, lifetime considerations, radiation hardness of all components, [37] redundancy in readout, and mechanical stability over time. Also, since the inner detector is placed inside the calorimeters, which measure particle energy, the tracker must be light-weight and optimized for minimum radiation length.

A back of the envelope calculation of some ATLAS silicon tracker design requirements

We can understand the order of magnitude of the required resolution for the ATLAS tracker by the following naive calculation. Lets take a singly charged particle with momentum of 1 TeV in our detector and say that we would like to determine the charge. The outer radius of the silicon tracker is 0.52 m away from the interaction point and to clearly resolve the sagitta of such a particle curving in the 2 Tesla ATLAS solenoid field, we need to have spatial resolution smaller than this sagitta. Taking the particle to be at $\eta = 0$, and ignoring the effects of multiple scattering, the radius of curvature is given by the following approximate relation:




Figure 1-2: A 3-dimensional cut-away view of the ATLAS inner detector, showing the various sub-detectors. The TRT C-wheels have been staged and are not shown.

$$r \simeq \frac{p_T(\text{GeV})}{0.3 \times B(\text{T})} = \frac{1000}{0.3 \times 2} = 1660\text{m} \quad (1.4)$$

Then the sagitta of this particle in the detector with dimensions, $L = 0.52$ m as before, is given by

$$s \simeq \frac{L^2(\text{m}^2)}{8 \times r} = 20\mu\text{m} \quad (1.5)$$

It follows then that the resolution of the silicon trackers per space-point should be less than $20\mu\text{m}$ in the bending direction to be able to do charge recognition in the volume of the silicon tracker for a 1 TeV particle. With the addition of the TRT, the momentum resolution for a 1 TeV particle is 36%, [17].

Pixels

The first layers of the tracker, being closest to the interaction point, have to provide precise tracking information, crucial for B-physics and also pattern recognition. A pixel tracker is ideal for this, with its inherently small pixel size and low noise and will provide excellent tracking and vertex resolution, [76]. The high cost per covered area prohibits the use of this system at intermediate to outer radii.

The pixel tracker consists of one “B-layer,” two barrel layers and two end-caps, each with three disks. The pixel B-layer is located at a radius of 5.1 cm and the two barrel layers are at 9.9 cm and 12.3 cm from the interaction point as seen in 1.2. The pixel modules, which are identical in all regions of the detector, are rectangular active devices of 6 cm by 2 cm with 46,080 pixels, with a pitch of 50 μm in the bending plane and 400 μm along the beam. Each pixel module comprises of a silicon pixel array bump-bonded to 16 readout chips which each serve an array of 18×160 pixel diodes. For each pixel channel, the electrical signal is amplified, shaped and compared to a threshold to provide a binary output to the readout system. Some analogue information is available through a time over threshold (TOT) measurement. In total there are about 67 million readout channels in the barrel and 13 million in the disks.

Being situated closest to the beam pipe, the pixel tracker has to cope with a high particle flux density as well as the extreme radiation environment. Pixel modules have been designed with this requirement in mind and tested extensively in such an environment, [11, 74]. The data transmission from the module is zero-suppressed with a maximum rate of 160 Mbits/sec, which can accommodate up to an occupancy of about 1×10^{-3} for a Level-1 trigger rate of 75 kHz with no dead-time while the expected physics occupancy is 1×10^{-4} . The pixel detector components will see the equivalent of 5×10^{14} neutrons/cm² over a 10 year lifetime. The radiation dose the B-layer will receive is four times higher than the first barrel layer. The B-layer can not withstand this dose and is designed to be replaced every few years.

As both the SCT and the pixels read out binary tracking information, some of the

readout development has been shared between the two detector communities, such as the optical readout system [69] and the read-out drivers [82].

SCT

The SCT (Semi-Conductor Tracker), is a silicon tracker with strips. The SCT can afford to have fewer readout channels since the occupancy from physics will be less than the pixels due to the larger radius. The coarser segmentation in the SCT as compared to the pixels is considerably cheaper and a larger area detector can be built.

The SCT barrel has 4 layers and each SCT end-cap consists of 9 disks, supported by a cylinder. The four SCT barrels, which will be discussed in detail later, are numbered from 3 to 6, since the 0th, 1st and 2nd layers of the tracker are pixel barrels. On each barrel, the modules are placed in rows with a tilt angle of $\approx 11^\circ$ from the exactly tangential orientation to provide coverage in $r\phi$ and to reduce the effects of the ionization electrons drifting in the magnetic field (Lorentz angle) on the module's spatial resolution. The number of rows on each barrel is given in Table 1.3. There are 12 modules in each row, which also corresponds to an LMT (Low Mass Tape) since each row is serviced by an LMT from each end. To ensure hermetic coverage in $r\phi$ and z , the modules are staggered in upper (U) and lower (L) positions in order to provide overlap as seen in Table 1.4. Between the upper and lower modules, there is an air-gap of $\approx 1\text{mm}$, limited by high-voltage breakdown issues.

Barrel	Radius	LMTs	Modules	sensor z-aligned
Barrel 3	299mm	32	384	bottom-side
Barrel 4	371mm	40	480	top-side
Barrel 5	443mm	48	576	bottom-side
Barrel 6	514mm	56	672	top-side
Total		176	2112	

Table 1.3: The SCT barrel geometry: the radius is to the centre of a module; the last column gives the u-v orientation. All the barrels have a length of 1492mm.

The 9 disks are numbered from 1, closest to the interaction point, to 9, the furthest. There are four different kinds of end-cap modules. They are all trapezoidal in shape and different in size, but otherwise similar to the SCT barrel modules in electronics

z+						z-					
-6	-5	-4	-3	-2	-1	1	2	3	4	5	6
1	2	3	4	5	6	7	8	9	10	11	12
L	U	L	U	L	U	L	U	L	U	L	U

Table 1.4: The module positions for one LMT.

and readout, [3]. They are placed on 3 rings on disks with the requirement that modules must overlap. To provide overlap in $r\phi$, inner and outer modules are placed on one side of each end-cap disk, the middle modules are placed on the other side. There are 52, 40 and 40 modules respectively in the outer, middle and inner rings. As the coverage of the SCT end-caps must extend to a $\eta = 2.5$, not all rings are used on the outer disks. For example, disk 9 only has its outer ring populated with modules. There are 988 modules in each end-cap.

An SCT barrel module comprises four single-sided p-in-n silicon detectors. Each silicon detector is $6.36 \times 6.40 \text{ cm}^2$ and has 768 readout strips, each of $80 \text{ }\mu\text{m}$ pitch, [2]. For each side of an SCT barrel module, two of these silicon detectors are wire-bonded together to give an active strip length of 123.2mm. The two sides of a module are glued together with a small (40 mrad) stereo angle to provide positional information along the z-direction. To provide a slight advantage in tracking, for barrels 3 and 5, the strips on the bottom side of the modules are aligned parallel with z while for barrels 4 and 6, the strips of the top side are aligned with z. This so called “u-v” orientation is also listed in Table 1.3. Overall, the SCT provides a spacial resolution of $16 \text{ }\mu\text{m}$ in $r\phi$, providing excellent tracking up to a radius of 51.4cm. The end-cap modules are arranged with their tapered strips aligned radially.

TRT

The Transition Radiation Tracker (TRT) is a light-weight, low-cost detector which provides a large number of tracking measurements as well as particle identification information. The TRT consists of gaseous proportional counters embedded in radiator material. The transition radiation measurement works on the principle that an ultra-relativistic particle emits X-rays when crossing the boundary between two materials

with different dielectric constants, [61, 83]. The transition radiation is dependent on the Lorentz γ factor, which is different for particles with the same momentum but different mass. Significant amounts of transition radiation is only produced if the γ factor is above 1000. Specifically, detecting this X-ray radiation from electrons with momentum higher than 1 GeV/c allows the TRT to distinguish them from other particles, such as pions which only radiate transition radiation if their momentum is close to 100 GeV/c. The TRT gaseous counters detect all charged particles from their energy loss as in Eq. 1.3, but also detect the energy deposition from the transition radiation X-rays. For this reason, the counters have two thresholds: the lower threshold detects the usual energy loss hits in the gas, and the higher threshold detects the energy loss plus the energy deposition from the X-ray. With the ATLAS TRT, the pion rejection is expected to be ≈ 150 for an electron efficiency of 80% at $p_T = 2$ GeV and ≈ 60 at $p_T = 20$ GeV, [22].

The TRT geometry is optimized to maximize the production and detection of the transition radiation X-rays in the polypropylene radiator material and the gaseous counters. The counters are 1.4 m long, 4 mm diameter straw tube, each fitted with a 30 μm diameter gold-plated tungsten wire, [8]. The wire is placed at high negative voltage and the aluminum coated inner wall of the straw tube acts as a cathode. When a particle passes through the straw, it ionizes the gas and the electrons drift towards the wire, forming an avalanche close to the wire where the electric field is high. The multiplication process is in the proportional regime, providing a signal amplification of $\approx 2.5 \times 10^4$. The wires are split in half at the center of the TRT and read out at each end, every 3.125 ns or in 24 time bins in total of 75 ns. For each channel, this provides a drift time measurement allowing for an estimation of where the particle crossed the straw. Using $Xe : CO_2 : O_2 = 70 : 27 : 3$ in the straws, the TRT can measure the position with an accuracy $\sim 140 \mu\text{m}$ [6, 7].

In the barrel TRT, there are 73 layers of straws oriented axially. Due to the geometry of the placement of the TRT straws, an average of 36 hits along a track is expected. In total, there are 52,224 straws in the TRT barrel and $\approx 160,000$ radial straws in each end-cap. The barrel straws are read out on both ends as they

are divided in the middle except for the first 9 layers. Due to the high occupancy during physics running, these layers are divided into three segments and the outer two are read out. The noise occupancy of the TRT is 2% if integrated over the 3 bunch-crossing read-out cycle of 75 ns. The physics occupancy of the TRT will be high, but with the aid of track seeding in the silicon detectors, the TRT will provide valuable tracking and particle identification information, [5].

BCM

The beam conditions monitor (BCM) is a late addition to the ATLAS inner detector. It will be situated very close to the beam-pipe. It consists of diamond detectors which are radiation hard and have fast signal properties, [77, 102]. It will provide valuable information on the instantaneous luminosity that ATLAS receives which is important for measuring the total cross-sections of physics channels as well as measuring background events, for example from beam-gas interactions. It can also provide a trigger for diffractive physics studies and physics of pomeron-exchange.

LUCID

LUCID is envisaged to be a dedicated luminosity monitor for ATLAS. It consists of 400 Cerenkov tubes surrounding the beam pipe at $z = \pm 18$ m, [105] and of Roman pots at $z = \pm 240$ m, which detect protons deflected through a very small-angle. Absolute calibration of the detector is possible by a measurement of the Coulomb scattering and the fast timing characteristics of this detector allow a bunch by bunch online luminosity measurement. LUCID should provide a 2% determination of luminosity as well as provide a trigger for diffractive physics studies.

Tracking in the Inner Detector

Nominally, a track will be identified from seven space-points in the silicon tracker, three from pixels and four from the SCT. “Coincident” hits on both sides of an SCT module define a space-point in the SCT. The very low occupancy of the pixels and the SCT space-points allow for seeding of the track finding algorithms.




Figure 1-3: Position and η -coverage of the inner detector components, [2].

Track finding in the inner detector involves finding and fitting a helical function to the particle's path as a function of the particle's momentum and the magnetic field strength. However, the magnetic field is not totally uniform and the particle does lose energy along its path and so the helix radius will decrease. For a large number N of measurements ($N > 10$) along a track, the resolution of measurement of the track curvature, k , is approximated by

$$\delta k = \frac{\epsilon}{L'^2} \sqrt{\frac{720}{N+4}} \quad (1.6)$$

where L' is the track length and ϵ is the measurement error. However, this formula is valid only when the tracking accuracy is not limited by multiple scattering at low momentum. For ATLAS, multiple scattering effects dominate tracking accuracy below 40 GeV.

1.2.6 Calorimeter System

The goal of the calorimeter system which surrounds the central solenoid is to stop and precisely measure directly the energy of the photons, electrons and jets and distinguish them. Fine granularity segmentation of the detector is necessary for good shower shape pattern recognition and robust particle identification. A reliable measurement

of missing E_t requires detector hermeticity, which has been a driving factor for the design of the calorimeter system.

The calorimeter system consists of a liquid-argon electromagnetic calorimeter and a surrounding hadronic calorimeter in the barrel region, [13]. The liquid-argon “hadronic” end-caps and the forward calorimeters cover up to a η of 4.9. A drawing of the different sub-systems of the calorimeter system is shown in Fig. 1-4.

The showers produced by light electromagnetically interacting particles such as the γ and e^\pm are wholly contained in the electromagnetic calorimeter, as they can penetrate much less than hadrons and produce narrower showers. Often a hadronic shower will start in the electromagnetic calorimeter, most of which will be absorbed in the hadronic calorimeter. Only the standard model neutrinos and muons with enough energy, and maybe some yet-undiscovered particles, can escape through the calorimeter system.

It is worth noting that the ATLAS calorimeter system is non-compensating, which means that the energy deposition detection response from a hadronic shower and an electromagnetic shower is not the same and has to be corrected in the offline software.

Electromagnetic Calorimeter

The electromagnetic calorimeter is a lead and liquid-argon sampling calorimeter with innovative accordion geometry, which ensures hermetic coverage as it minimizes the gaps between detector modules. It is inherently radiation hard due to the materials used and provides long-term stability. The liquid-argon calorimeter is contained in a cryostat with an outer radius of 2.25 m and extends 6.65 m along the beam, [20].

Electromagnetic showers may start to develop in the inner-detector region since the total material before the calorimeter is about $2.8X_0$ radiation lengths at $\eta = 0$. Hence, a pre-sampler of an active layer of liquid-argon with a thickness of 1.1 cm is placed in the barrel region, $\eta < 1.8$ to allow for corrections in the energy measurement.

The thickness of the electromagnetic calorimeter is above $24X_0$ in the barrel region and is about $26X_0$ in the end-caps. The high-granularity of the detector elements allows for a precise position measurement, and the information obtained by the energy




Figure 1-4: A cutaway view of one half of the calorimeter system, showing the positions of the different components.

deposition and shower development shape in the calorimeter and the pre-sampler allows for an excellent γ/π and e/π separation, [43].

Hadronic Calorimeter

The hadronic calorimeter employs different calorimetry techniques for different pseudo-rapidity regions, suited to the different needs and to different radiation environments, [14]. Good containment of the hadronic showers is needed to keep punch-through to the muon system to a minimum. A thickness of ≈ 10 interaction lengths is enough to achieve this containment. The barrel part is a scintillating tile and iron based sampling calorimeter, [60]. The outer radius of the scintillator-iron calorimeter is 4.25 m and covers $\eta < 1.7$. At higher pseudo-rapidities, $1.5 < \eta < 3.2$, a copper and liquid argon detector with parallel plates, similar to the electro-magnetic calorimeter, is employed due to its better radiation tolerance, [63]. The forward calorimeter, is composed of rod-shaped electrodes in a tungsten/copper honeycomb matrix, filled with liquid-argon and covers $3.1 < \eta < 4.9$.

1.2.7 Muon System

The muon spectrometer is a very large tracker which determines the size of ATLAS, as seen in Fig. 1-5. The outermost chambers of the barrel muon system are at a radius of 11 m. The furthest position measurement from the interaction point is in the 3rd layer of the muon end-caps, which is located about 23 m from the interaction point. The alignment of the muon system is critical since its dimensions are so large and prone to displacement and deformations. An optical alignment system will monitor the chambers and offline analysis can correct for their mis-alignments to reach the design precision of $30\mu\text{m}$.

The muon system employs different technologies depending on the η and the purpose, [16]. The muon system has to track muons with high-precision over large distance as well as provide a muon trigger for the rest of ATLAS. For tracking measurements, Monitored Drift Tube chambers (MDTs) are employed in the low η barrel




Figure 1-5: A cutaway view of the 3 barrel layers and 3 wheels of each end-cap. The positions of the different technologies are indicated.

and end-cap regions while Cathode Strip Chambers (CSCs) with higher granularity and higher rate and radiation tolerance are used in the $2 < \eta < 2.7$ region. The muon trigger is provided by Resistive Place Chambers (RPC)s in the barrel region and by Thin Gap Chambers (TGC)s in the end-cap region.

In contrast to the other tracking detectors in ATLAS, the tracking resolution of the muons chambers have to be precise along the z direction in the barrel and along r in the end-caps. The reason is that the toroidal field bends the particles in η , effectively changing the angle they make with the beam-pipe. In the muon system, the average magnetic field is 0.6 T, allowing for a measurement of muon momentum up to 6 TeV.

1.2.8 Trigger and Data Acquisition System

The amount of data detected by ATLAS is rather large for today's technology to record fully. ATLAS has ≈ 200 million readout channels in total, some digital and some analog. The multiplicity of events at 14 TeV can not be predicted exactly but the total inelastic cross section at 14 TeV is predicted to be ≈ 80 mb. With the rate of collisions at 40 MHz, it is impractical and not necessary to store information for each channel for each event. For each event, there will be hits only on some channels, and so zero-suppression can be employed to reduce the amount of information that needs to be stored. Nor is it necessary to store every event because most of the events will not reveal any new physics.

In essence, the interactions in pp collisions can be divided into two categories: “minimum-bias events” and “hard scattering events.” The minimum bias events are “soft” collisions between two incoming protons interacting from a long distance. They represent the bulk of collisions at the LHC with a cross section of ≈ 70 mb. While most of the particles from these collisions escape down the beam pipe as their final states have low p_T , a few particles with high enough p_T may leave tracks in the central detectors, causing “pile-up.” The hard-scattering events are more interesting due to their discovery potential for new physics. In these interactions, the momentum transfer can be very large as they result from short range parton interactions from incoming protons. A clear signature of such events are high p_T particles and jets. Since about 100 events per second can be written to disk, the ATLAS trigger is designed to identify these hard-scattering events for specific physics searches and precision measurements, [18], with a rejection factor of 10^7 from minimum-bias events. As avoiding pile-up is impossible, a sample of minimum-bias events for QCD-background studies is needed as well.

The trigger system has three levels of online event selection, each of which reduce the accepted event rate from the previous level. The first level, Level-1 trigger uses reduced-granularity information from a subset of detectors – specifically the muon and calorimeter systems – to identify events which contain high p_T particles or jets

and large missing E_T . The Level-1 decision is made by the CTP (Central Trigger Processor) and distributed to all detectors in ATLAS at a maximum rate of 75 kHz by the TTC (Timing, Trigger and Command) system. The TTC system distributes the LHC-synchronized 40 MHz bunch-crossing clock, the global bunch-crossing and event counter resets as well as the Level-1 accept (L1A) trigger signal to all TTC “partitions” or sub-detector readout crates.

The Level-1 trigger requires signals from various detectors to reach the CTP and the Level-1 trigger signal to be generated and distributed to all the front-ends of all sub-detectors. The Level-1 decision needs to be made quickly, as all detector front-ends have to pipeline and store information until the Level-1 decision arrives. Allowing for the distribution of the signals to the front-ends, this total Level-1 trigger “latency” is defined to be $\approx 2.5 \mu\text{s}$ for ATLAS which leaves about $1 \mu\text{s}$ for the CTP to make the Level-1 decision. With Level-1 trigger decision, all information from the front-ends of all sub-detectors is read out through the RODs (ReadOut Drivers) into the ROBs (ReadOut Buffers) by the ROS (ReadOut Sub-system), as shown in Fig. 1-6. The RODs and ROSeS are housed in USA-15, which is a cavern on the side of the ATLAS cavern where radiation levels are low. Zero-suppression of channels with no information occurs on the front-ends and the RODs.

The Level-2 trigger uses the Region of Interest (ROI) information provided by the Level-1 trigger to access full-granularity detector information from every sub-detector but only for those sectors which triggered Level-1 and now are stored in the ROBs. It is worth noting here that the inner detector tracking information is available for the Level-2 trigger decision. Level-2 triggers are generated at an average rate of 1 kHz with a maximum of 3 kHz.

Level-3, also known as the Event Filter (EF) uses the fully reconstructed event information to select which events should be written to disk for offline analysis. Data is written to disk at about 200 Hz with one event containing ≈ 1.5 MB of information.




Figure 1-6: Timing and trigger distribution to the SCT is through the TIM (Timing Interface Module) which gets the TTC information from the CTP (Central Timing Processor). Data is collected at the ROD level and then sent up to the Event Builder through the ROSes.

1.2.9 Offline Processing

The offline framework for ATLAS is called Athena. Its main task is to reconstruct and analyze the immense amount of data that is written to disk as well as simulate an equivalent sample of data for back-ground physics studies. Although it is called “offline,” some portion of it works “online,” specifically for high-level event selection (Levels 2 and 3) and to monitor the conditions of various detectors and components.

For simulation, Athena depends on other packages such as PYTHIA and HERWIG to generate events. Specific particle filter algorithms can simulate the online-trigger and generate samples of physics studies.

Both for reconstruction and simulation, an incredible amount of information is needed beyond what is provided as hits in the data files. An extremely detailed detector description as well as conditions and configurations of the beam, the magnets, all sub-detectors and triggers is necessary for the long and excruciating journey of correlating hits with real particles and vice versa. The conditions and configuration information needs to be available from online databases, which get updated regularly. Detailed magnetic field maps, material budget, muon system alignment monitoring information and temperatures of modules in the SCT are just a few examples of the ingredients provided to the offline in this long process. The offline also writes information to the databases for use in reconstruction. Such information comes from monitoring of detectors as well as alignment studies.

The offline seeds tracks from hits in the silicon tracker and tries different combinations of hits on subsequent layers as track segments are allowed to grow into the TRT and outer detectors. However, it can not reconstruct each and every track in the under-lying event, but needs to estimate pile-up from them. The tracks from the hard-scattering event are reconstructed through the magnetic-field lines with updated energy and momentum information at every point along it. This calls for a good understanding of energy loss as well as elaborate tools for bremsstrahlung recovery.

The high statistics obtained by the ATLAS detector and the detailed offline analysis tools have the potential to achieve the goals anticipated by the physics community.

Chapter 2

SCT: Semi-Conductor Tracker

The SCT is a silicon tracker designed for the LHC environment. Before an in-depth discussion of the SCT and its performance, a brief description of silicon technology and the design decisions for the SCT given its environment is necessary.

2.1 Silicon Detectors

The working principle behind silicon detectors is that silicon is a semi-conductor whose charge carrier density can easily be modified and finely controlled in well-defined regions, statically by introducing impurities into its crystal structure, and dynamically by applying an electric field. In pure silicon, an ionizing particle may excite an electron out of its energy level if the energy transferred exceeds the necessary energy needed for an electron-hole pair production, which is on average 3.62 eV at 0 K [94]. Although the band gap is only 1.1 eV between the valence and conduction bands in silicon, the energy needed to produce an electron-hole pair is higher as the rest of the energy goes into exciting lattice vibrations. Electron-hole pairs are constantly formed in the material from the thermal heat, but subsequently recombine to leave a neutral medium.

By introducing impurities, known as doping, the intrinsic electrical properties of silicon bulk can be modified. Dopants are either electron donors or acceptors, forming respectively *n-type* and *p-type* silicon. A donor atom, often a group V atom from the

periodic table with one more electron than silicon such as phosphorus, donates its weakly-bound valence electrons, creating an excess of negative charge carriers, while an acceptor dopant, is often a group III atom, [120]. Junctions between regions of n-type and p-type semiconductors have built-in electric fields and the electrons from n-type and holes from p-type regions recombine and eliminate each other. This creates a non-conducting depletion zone in the middle where charge carriers are almost totally absent. Reverse-biasing such a junction, by putting the p-type region to a higher electrical potential than the n-type region, increases the width of this depletion zone. A sufficiently large applied high voltage can increase this to the full thickness of the wafer. When an ionizing particle passes through it, it will liberate a large number of electron-hole pairs, as seen in Fig. 2-2. These electrons and holes will drift respectively toward the anode and the cathode and can be collected on implant strips either directly or by capacitive AC coupling to a metalization layer, [53]. There is no amplification process required in a semi-conductor detector as the primary charge collected is large (on the order of 10^4 electrons) compared to gaseous detectors. As the deposited charge is easily detectable above the thermal background noise, such a reverse-biased junction can be made into a particle detector. Nonetheless, pre-amplification of the collected signal is necessary.

While depletion allows for all primary charge to be detected, over-depletion is essential if the detector is desired to have fast charge collection. The higher electric field when over-depleted speeds up the charge carrier flow which reduces losses due to charge trapping.

Although the junction is reverse-biased, a small fluctuating current flows through it, resulting in output noise and setting the limit on the smallest signal that can be observed. This current, known as the leakage current is caused by the movement of minority carriers in each of the p and n region. Thermally generated pairs in the depletion region will also drift under the influence of the electric field. A large contribution to the leakage current can come from surface currents, which depend on the existence of contaminants and non-uniformities on the surface, the minimization of which requires a clean environment for detector production.

2.1.1 Silicon Detectors in a radiation environment

The primary radiation damage to silicon detectors is a result of the changes in the crystal structure of the semiconductor. An inelastic scattering of a high energy hadron on a nucleus in this crystal structure can cause it to recoil, removing it away from its original lattice site. Such displaced nuclei form stable defects in the bulk material, effectively increasing the number of acceptors, making the material more p-type. Heavy irradiation of a lightly doped n-type material can even cause it to become p-type. If the radiation damage is not too great, higher electric fields can still fully deplete the detector, maintaining the full charge collection.

The damage results in an increase in leakage current as well as a loss of mobility of drifting electrons and ions through recombination. Lower charge collection efficiency and increased noise leads to a reduction in the signal to noise ratio.

In addition, detector read-out components can also be affected by the irradiation. As they are specific to the design of the read-out system, this will be discussed later specifically for the SCT.

2.1.2 Design requirements of the SCT module

For the requirements on the design of the SCT detectors and the front-ends, the long-term performance of the SCT and its impact to ATLAS can be described by the following parameters:

1. Resolution: The standard deviation of a measurement from a rectangular strip is its width divided by $\sqrt{12}$ since

$$\sigma^2 = \int_{-1/2}^{1/2} x^2 dx = \frac{1}{12} \quad (2.1)$$

Then the spatial resolution of two independent ideal binary silicon strip measurements is given by

$$resolution = \frac{pitch}{\sqrt{12} \times \sqrt{2}} \quad (2.2)$$

The factor of $\sqrt{2}$ comes from there being two independent measurements. At

high energies, the resolution of the silicon tracker per space-point is required to be $< 20\mu\text{m}$, the strip pitch of the detector should be less than $100\mu\text{m}$. At low energies, the energy resolution is dominated by multiple-scattering and there is no significant benefit of having a smaller strip width.

2. Low physics occupancy: Physics occupancy should be $< 1\%$ at the highest luminosities, considered the limit for offline pattern recognition capabilities, [15].
3. Efficiency: The efficiency of SCT modules should be more than 99% for good pattern recognition, which also requires fast signal shaping and small time-walk to make sure that the hit is registered in the correct time-bin. (Time-walk is the difference in arrival time of signals of different charge deposits.)
4. Low noise: At the design luminosity of the LHC, the detector occupancy from collisions is estimated to be between 0.2% and 0.5%. It is important that the noise occupancy be negligible compared to the collision occupancy so that it does not cause inefficiency through decreasing the bandwidth for real hit data transmission or obscuring the real hits. The design requirement is that silicon modules have noise occupancy less than 5×10^{-4} per strip at the operating sub-zero temperatures.
5. Radiation length: Minimal material in the inner tracker is ideal, to keep the degradation of the momentum and impact parameter resolution to a minimum by keeping bremsstrahlung and multiple-scattering to a minimum. Photon conversions in the inner tracking volume should also be kept to a minimum.
6. Radiation hardness: SCT is expected to withstand fluences of up to $\approx 2 \times 10^{14} n_{eq}/\text{cm}^2$, normalized using the non-ionizing energy loss cross-sections from expected damage of 1 MeV neutrons and also an ionizing dose of 100 kGy(Si), [15].




Figure 2-1: A barrel module shown in a test-box with a test connector. The six readout chips of the hybrid are clearly visible. The lower wafers can be seen protruding due to the stereo angle between the two sides.

2.2 SCT barrel modules

The SCT barrel modules are made of 4 wafers, as discussed in Chapter 1 and shown in Fig. 2-1. They are of type: (p+, n, n+), which means that a lightly-doped n-type silicon is sandwiched between a more heavily doped n-type silicon at the base and a heavily doped p-type at the surface. The four sensors are all glued onto a stiff, high thermal conductivity graphite (TPG) baseboard with BeO facings on both surfaces for thermal management, [26]. Strict tolerances are imposed on the positioning of the wafers on the baseboard as the resolution of a module depends on it. The read-out electronics are mounted on a wrap-around hybrid which supports 12 read-out chips. Here we will first discuss the wafers and then the read-out electronics of the SCT barrel modules. End-cap modules are similar in many ways and their design is documented in [100].

2.2.1 SCT wafers

The above requirements were taken into consideration while making design decisions and optimizing SCT sensor parameter. *P-in-n* and *n-in-n* silicon were considered for the design and *p-in-n* was chosen, [64]. Some of the important values are shown on the schematic of the sensor design in Fig. 2-2. The reasons for these design decisions are as follows: [4]

- Thickness: The goal of maximizing charge collection efficiency while maintaining low radiation length sets the scale for the thickness of the wafers. Also, a small collected charge requires more power from the front-end pre-amplifier to achieve a high signal-to-noise. A silicon thickness of $285\ \mu\text{m}$ was chosen, which presents 0.61% of a radiation length to particles crossing it at a normal incidence. Given that the average density of silicon is $2.3\ \text{g/cm}^3$, a Minimum Ionizing Particle (MIP) is expected to deposit an average of 120 keV. As the most probably energy lost by a MIP is 90 keV, the most probable number of electron-hole pairs created is 25,000.
- Inter-strip distance: The inter-strip distance is $80\ \mu\text{m}$ which provides a spatial resolution in the bending direction of $16\ \mu\text{m}$, fulfilling the earlier criteria and providing a 2σ measurement of the sagitta for a particle with $p_T \approx 1\ \text{TeV}$. For such a particle, at $\eta = 0$, using the result from full simulation, the momentum resolution, $\sigma(1/p_T)$, will be $1/360\ \text{GeV}$, [17]. A smaller inter-strip distance would cost more and may suffer from high-voltage breakdowns.
- Strip width: While minimizing the strip width helps to improve the resolution of the tracker, this increases charge sharing between the strips which reduces efficiency. The width of the strip also affects the inter-strip capacitance, which if too thin, is a source of cross-talk. The strip width for the SCT has been chosen as $22\ \mu\text{m}$ as the optimal value.
- Inter-strip capacitance and strip length: Inter-strip capacitive noise dominates the noise of an SCT channel. Inter-strip capacitance is $\approx 1.3\ \text{pF/cm}$ and scales




Figure 2-2: A schematic cross-section of an SCT sensor, a p-in-n silicon micro-strip detector, [24]. A track passes through the depletion region and liberates electrons and holes which are swept to the electrodes by the electric field.

with strip length. While longer strip-length means that the module will cover larger area with the same number of read-out channels, it is limited by the requirement of having a low physics occupancy at high luminosities. A strip length of 12 cm was chosen, which keeps noise occupancy less than the specification of 5×10^{-4} per channel, corresponding to a noise requirement of $< 1500e^-$ equivalent noise charge (ENC).

- AC vs. DC coupled detector read-out: AC-coupled detectors simplify the complexity of the read-out chip design. The signal formed in the silicon, is capacitively coupled to the aluminum electrode through a SiO_2 insulation layer. The large capacitance through the SiO_2 layer from the strip to the readout electrode means that the impedance is low and that the signal collection is fast, < 10 ns.

The signal induced on one aluminum strip is the result of the holes drifting towards it. While the same number of electrons drift away from it, the weighting function, [118] for electrons is smaller than for holes as they are, on average further away from the electrode. The drifting of electrons induce a signal on several neighboring strips whereas the signal generated from the drifting of the holes is localized to one or two

strips, as illustrated in Fig. 2-2. Charge collection in the electrodes couples into the design of the SCT front-ends with the following important considerations:

- **Charge collection:** Minimizing noise requires minimizing the input capacitance to the pre-amplifier, [107]. For this reason, the SCT front-ends are mounted on a hybrid, directly bonded to the wafer, minimizing distance and hence capacitance between charge collection and the pre-amplifier.
- **Power:** Power dissipation of the front-ends and removal of the heat is crucial as the front-ends are inside the detector volume. Also, the leakage current drops exponentially with temperature, so even a small reduction in temperature brings a significant improvement in reducing leakage currents and hence noise. Keeping the SCT at low temperatures, $< 0^\circ \text{C}$, is essential to minimize the detrimental effects of irradiation as well.
- **Read-Out:** A critical design decision for the SCT is the binary read-out. Often the charge deposited in the wafer is shared between multiple-strips. Having analogue information from the cluster of hits slightly improves the resolution of the detector, while the advantage of binary read-out compared to analog or digital read-out is its lower cost and that it requires less data transfer as well as consumes less power, which then requires less material for cooling. However, the main disadvantage of a binary read-out is that it is harder to calibrate and harder to monitor its performance, which in turn makes the read-out and calibration system a critical issue for a successful and working SCT.
- **Noise:** The front-end amplification and digitization of the signal by use of a threshold needs to be efficient ($> 99\%$) in signal detection and should not introduce any coherent noise into the system.
- **Shaping time:** There is an LHC bunch-crossing every 25 ns, which means that the signal from the resulting particles must be collected in a window shorter than 25 ns. While it takes a reverse-bias of about 70 V to fully deplete the wafer, the SCT modules will have 150 V applied to over-deplete the detector and to achieve




Figure 2-3: The reconstructed pulse shape due to beam particles from the test beam for a non-irradiated (left) module biased at 150 V and an irradiated (right) module, biased at 500 V, from [37]. The peaking time increases from 20 ns to 24ns with radiation.

fast charge collection, [4]. This will ensure that the signal will be identified with the correct bunch-crossing and not spread over to the next bunch-crossing. The SCT reads out 3 bunch-crossings of data, known as the “hit pattern,” centered around the Level-1 trigger time-bin to reduce mis-identification of hits. Fig. 2-3 shows the pulse shapes at the discriminator input. These are reconstructed from test-beam data, [37].

- Speed: The average Level-1 trigger rate is 75 kHz. Correct timing and identification of bunch-crossing cycles, pipelining signals until the arrival of the trigger decision and fast-response to a trigger is essential for Level-2 input.
- Robustness: Once the SCT detectors are installed inside the ATLAS detector, they will be inaccessible. Hence it is essential that the sensors and the read-out operate reliably over long periods under irradiation and possibly, repeated shutdown. Redundancy in the control and read-out circuitry of the front-end chips have been implemented with this in mind.

2.2.2 Radiation damage to an SCT module

The characteristics of the SCT modules will change under irradiation. Nonetheless, it will work reliably and have a high resolution while maintaining a good signal to noise ratio. To initially select technologies to achieve this and later to demonstrate reliable performance, several irradiation studies have been performed, [98]. Type inversion with radiation damage makes full depletion more important. A higher leakage current requires higher high-voltage for full depletion as well as charge-trapping which also requires a higher high-voltage for better signal collection.

Initially, the SCT wafers are of type: (p+, n, n+), but after irradiation for a couple of years of ATLAS running, the n-type silicon will become p-type, known as “type-inversion.” The SCT structure will look like (p+, p, n+) which will amount to the p-n junction moving from the top of the wafer to the bottom.

The depletion region starts at the bottom. If the module is not fully depleted, the depletion region will not reach the p-strips and the signal is lost. The generated holes will be lost due to recombination in the bulk. So the detector must be fully depleted after irradiation, requiring increase high voltage. In the SCT, the high-voltage bias can be increased up to 500 V from the original 150 V, [84].

Furthermore, leakage currents and noise will increase with radiation, [58] as the dominant noise source is the inter-strip capacitance, which increases with irradiation due to trapped charges at the Si-SiO₂ boundary. While higher leakage currents lead to higher heat dissipation, in turn higher temperatures lead to an increase in thermally generated electron-ion pairs, resulting in higher leakage currents. This process places tougher requirements on the cooling system to stop what is potentially a catastrophic thermal run-away. The binary threshold of the read-out chip can be increased as well to compensate for the increase in noise, [124]. The shaping time of the signal will get longer as well, causing the signal to leak over to the next time-bin. While the median signal-to-noise is $\approx 13 : 1$ for unirradiated modules, it will be $\approx 9 : 1$ for irradiated modules. The tracking performance of irradiated modules has been verified in a test-beam in 2002, [128].

Radiation damage is strongly dependent on the temperature of the silicon. If the SCT is kept constantly at low temperatures ($< 0^\circ \text{C}$), the reverse-annealing can be minimized. Radiation hardness and functionality after irradiation is not only essential for the wafers but also for the front-end chips and read-out electronics, such as VCSELS, [121].

For a silicon detector with binary read-out, another hazard of being in a radiation environment is the possibility of “Single Event Upsets” (SEU), which has been studied in detail, [62]. As a particle crosses the read-out chip, it can corrupt the event content, which consists of hit patterns, or its identifiers. An event is identified by the bunch-crossing and the L1 counters as read-out from the first chip on each side of the module. Especially, due to the high rate of bunch-crossing counter errors that would arise from SEU, both counters will be reset regularly to remedy the situation, [73]. How they are reset will be described in section 2.3.3.

2.2.3 Principle of binary read-out

Binary read-out provides a simplification for front-end electronics and read-out system as there are only two possible values for any channel. Binary electronics requires a smaller area chip than for analog or digital electronics which would employ ADCs. Such binary chips are cheaper and consume less power than the alternatives, putting less stringent requirements on the cooling system, which in turn reduces the material in the SCT cooling system.

Furthermore, the hit data can be compressed easily by suppressing information from channels with no hits. This zero-suppression allows for a high Level-1 trigger rate to be accepted. Also, a reasonable number of events can be buffered by the on-detector electronics until the data can be transferred to off-detector electronics with minimal dead-time.

Naturally, the challenge of a binary read-out is to set the binary threshold to its optimal value, which minimizes noise and maximizes efficiency for each single channel. An optimal threshold can be set only if the noise generation and the signal deposition is well understood and if the calibrating charge-injecting mechanism is accurate. A




Figure 2-4: Noise distribution with arbitrary normalization (dashed line) and Landau charge deposit distribution for a MIP at normal incidence (thin line). The 1 fC threshold gives good separation. The thick solid line shows the average expected occupancy as a function of threshold, [99].

threshold of 1 fC was found to be the optimal working point for the SCT, [37]. Figure 2-4 shows that an RMS equivalent noise charge (ENC) of $\approx 1500e^-$ or 0.2 fC,¹ is below the 1 fC threshold. While the most probably value deposited by a MIP is 3.6 fC, the 1 fC working point is $> 99\%$ efficient with Landau fluctuations even with $\approx 50\%$ of the charge collected on one strip due to charge sharing between strips. This nominal operating threshold satisfies the earlier criteria of a noise occupancy of $< 5 \times 10^{-4}$ and an efficiency of $< 99\%$.

Setting the 1 fC threshold uniformly across all channels is one of the most important tasks of the calibration electronics and software. A mechanism is needed to calibrate out the channel to channel variations in response to the same deposited charge as to avoid the offline tracking algorithm from being biased by the potential extra or missing hits. Module to module as well as channel to channel variations need to be minimized to provide a well-calibrated SCT. A charge-injection calibration mechanism is in place in the read-out chips and will be described in the next

¹1fC is equivalent to $\approx 6250e^-$.

chapter.

2.2.4 Front-end electronics

The front-end electronics of the module is mounted on a flexible multi-layer PCB copper-kapton hybrid[90] above the detectors. This hybrid provides data and power connections to and from the detector, folding in the middle to service the lower side of the module, as seen in Fig. 2-1. Wire bonds are made from a pitch adaptor to the input pads of the front-end chips and also to each channel on the module. There are 12 front-end chips called ABCD3TA ASICs [38], 6 for each side of the module, which provide the binary read-out of 128 detector channels each. The ABCD technology is based on DMILL, [57].

The ABCD read-out chain, shown in Fig. 2-5, consists of a bi-polar analogue circuitry, [65] followed by a series of digital CMOS devices, [39]. The signal collected on each channel is first amplified and then shaped. The calibration circuit can also inject charge into this analogue circuitry. The circuit injects charge into every 4th channel and looping over 4 offsets allows for all channels to be calibrated.

After the amplifier, a programmable threshold discriminator turns the signal into binary information. This comparator has a programmable 8-bit DAC for the threshold adjustments common to all strips in a chip. A 4-bit *trim* DAC adjusts for strip to strip variations. A common *trim range* for the chip sets the step size for a trim DAC bit.

The discriminator compares the level of the signal with what is the equivalent signal level from a 1fC charge deposition. There are two modes of latching this signal to the clock, respectively called “edge-detect” mode and “level” mode. Both modes discriminate the signal at 1fC, initially producing binary pulses with a duration a few ns longer than the time over threshold and not latched to the clock cycle. The level mode samples at the rising edge of the clock, producing a signal latched to the clock with as many clock-cycles as the length of the initial signal. On the other hand, the edge-detect mode is built to detect a high to low transition in the data, generating a pulse of duration of 1 clock-cycle irrespective of the length of the incoming pulse.

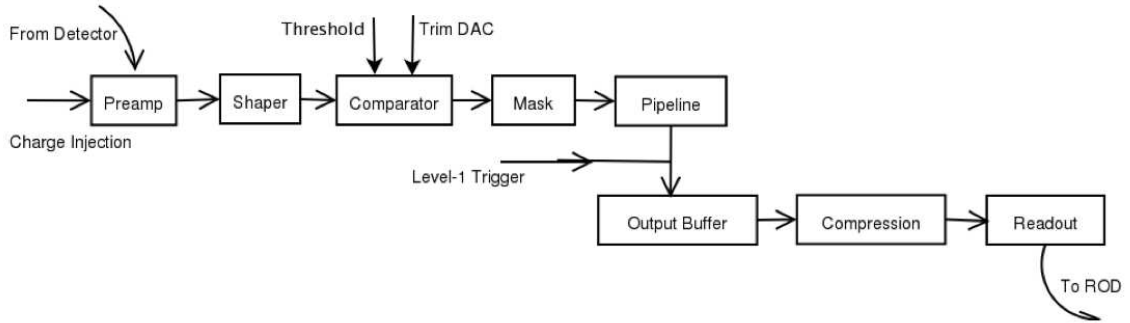


Figure 2-5: Signal processing chain within the ABCD chip, [71]. The analogue processing takes place in the pre-amp, shaper and comparator while the rest of the chain is digital. In normal running, the final output is zero-suppressed compressed data.

Running in edge-detect mode results in higher noise occupancy due to the detection of short duration pulses, while it reduces the “shadow” hits for the Level-1 triggered time-bin from the time-bin before it, effectively cleaning signal for the triggered time-bin.

The rest of the circuitry is in CMOS. After passing through a mask register, which is for example used to mask noisy channels, the signal is fed into a binary pipeline, [51]. For each channel, there are 12 *pipeline cells* which store the information in a staggered chain of 12 bits, and shifts the registers at every clock-cycle. Passing of the clock in different clock phases allows for the bit pattern to appear at the end of the pipeline after 132 bunch-crossings corresponding to the time it takes for a Level-1 trigger for a specific bunch-crossing.

The Level-1 accept signal is 3-bits long as sent by the off-detector components and if there was a trigger, the pipeline output for 3 bunch-crossings centered around the Level-1 trigger bunch-crossing is transferred to a de-randomizing buffer of 8 *events* deep for read-out. An event here is defined as 3 bunch-crossings worth of information for every single channel. While the expected worst case occupancy for the inner most layer of the SCT is 0.6%, a safety margin was built in so that a 2% detector occupancy with a buffer depth of 8 events leads to a dead-time of only 1%, [133].

When the event is dequeued out of the read-out buffer, it is compressed to reduce the event size in the mode specified by the configuration. The compression

logic is built to suppress channels with specific hits patterns, [119]. There are four compression patterns available:

- *Hit-mode*, also known as *xxx*, returns a hit in any of the 3 time-bins and suppresses only those channels with no hits in any of the 3 time-bins.
- *Level-mode*, also known as *x1x*, returns only those channels with a hit in the central time-bin.
- *Edge-mode*, not to be confused with *edge-detect* mode, also known as *01x*, matches a central hit following a no-hit time-bin.
- *Test-mode* returns all hit patterns and provides no compression.

Hit-mode provides a good measurement of noise in patterns such as *100* or *001*, if there is no hit on the other side of the module and so the hit is not due to a particle. However, the read-out of the hit-mode takes longer as the information to be transferred off-detector is larger. For commissioning and low occupancy running, hit mode is ideal to debug the detector, but for high luminosity running, level-mode and edge-mode will be used. Edge-mode ensures that the hit in the central bin is not due to a hit that spilled over from the previous bin and further reduces the event size.

Further compression is provided by “clustering” of hits. Quite often, the deposited charge from a particle will be spread over several channels. The information from such a cluster of such hits is read out by identifying the first hit channel and then giving the number of consecutive channels and then the hit patterns for each of them. While 17 bits are needed for an isolated hit in the data stream, only 4 additional bits are needed for a neighboring hit.

2.2.5 Read-Out of modules

When a chip reads out the event content, it adds a 4-bit chip number to identify the bit-stream. All chips are the same, but their position in the read-out sequence and hence their configuration gives them different roles. The “first” chip on each

side of the detector is designated to be the *master* chip as it is the closest one to the module connector and hence the read-out electronics. The master chips are known as *M0* on the top side and *M8* on the bottom side of the module. The rest of the chips are known as *slaves* and although the last chip to read-out is also a slave, since its configuration has to be different, it is called an *end* chip. On each side of the module, the read-out starts with the master after an L1A has been received. The master sends its bit-stream with an event header, containing its bunch-crossing and Level-1 identifiers to the read-out circuitry, connected to the module. The master passes a *token* to the next chip to tell it to send its bit-stream, when there is only one bit left for it to read out. This ensures that there is no break in data streaming. When that slave chip has completed reading out, it passes the read-out token to the next chip and so on, until the end chip is reached, which adds a string of zeros that can not otherwise appear in the output. If there are no hits for a particular chip, the circuitry sends a special “no hit data” string and passes the token to the next chip in the chain.

The token-passing is allowed through special electrical lines on the hybrid that connect the chips. Each chip’s token passing circuitry is connected to its neighbor as well as the chip that comes after its immediate neighbor. This provides immunity to single point failure so that in case of a single chip failure that chip can be by-passed and the data from the rest of the chips can be read out. Fig. 2-6 shows various chip-bypassing possibilities.

If a master chip or opto-communications fail for one side of the module, the data can be read out for all chips except the master chip on that side, since the first slave on one side is connected to the end slave on the other side. For barrel modules, the master chip’s token passing line is not connected to the end-chip by default whereas for end-cap modules, it is. In special locations on the barrels, where opto-communications are known not to work, special *modified* modules were produced which allows the master chip to receive a token from the end-chip on that side, enabling all chips to read out, but out of sequence.

Modified modules became available slightly after the completion of Barrel 3. They

were placed on those harness locations on Barrels 4, 5 and 6, where it was known from barrel reception harness testing that the optical communications were malfunctioning.

The token passing of the read-out chips is controlled by configuration of a register in the chips, which will be discussed in the next section.

2.3 Detector control and readout


The control of the modules is provided by the on-detector opto-electronic readout hardware which is in communication with the off-detector hardware, known as a ROD (ReadOut Driver) crate. We will now detail how the opto-communications, the DCS (Detector Control System) and the ROD crate works. Then we will discuss how the read-out system is configured for specific module read-out modes.

2.3.1 Optical Communication with Modules


The BOC (Back of Crate) card provides optical communication with the modules. The bi-phase mark (BPM) encoding is used to send commands to the module in one optical signal from the BOC as shown in as seen in Fig. 2-7. BPM combines both clock and command in one signal, and is “balanced” – ie. it is 50% high and 50% low, independent of the bit pattern being transferred, [41]. This is important for high-speed data transmission to diodes. In return, the BOC receives one data stream from each side of the module. The DORIC4A (Digital Optical Receiver IC) and the VDC (VCSEL² Driver Chip) ASICs are used in the conversion of these signals between electrical and optical form at the module end, [135]. The master chip on each side of the module is responsible for the electrical LVDS (Low-Voltage Differential Signal) transmission of data to a VCSEL driven by a VDC. The DORIC4A decodes the encoded clock and command signals as received by the $p-i-n$ ³ diode and passes them on to all 12 ABCD chips. The fibres which send commands to the module though the $p-i-n$ diode and the DORIC4A are called TX for “transmitter” fibres while the

²Vertical Cavity Surface Emitting Laser


³In the $p-i-n$ diode, the i stands for intrinsic silicon.




(a) The standard module read-out scheme.



(b) If any chip fails, it can be by-passed. Here, slave chip S1 is by-passed by changing the configurations for chips M0, S1 and S2.



(c) If a master chips or opto-communications fail for one side of the module, the master chip has to be by-passed and therefore, can not be read out. Here, the configurations for E5, M8 and S9 are changed.



(d) If a modified module is in place, the end chip can pass the read-out token to its original master chip, enabling it to be read out. Here, the configurations for E5, M8, S9 and E13 are changed.

Figure 2-6: Different chip by-passing schemes for normal and modified modules




Figure 2-7: The principle of bi-phase mark(BPM) encoding,[41]. BPM combines clock and command into one stream. A chain of 0's gives a 20 MHz clock. As illustrated, a 1 gives a flip of the clock phase. The BPM chip delays the signal by three bunch-crossings.

fibres which receive back data from the modules are called RX for “receiver” fibres.

The system can be configured to send a copy of the decoded clock and command signals to a designated neighboring module through a special connector, which provides TTC redundancy, as seen in Fig. 2-8. The configuration for the module receiving the redundant TTC needs to be changed to enable the physical connection called the “select” line, which will be discussed later. The *p-i-n* diode, VCSELs, DORIC4A and the VDCs are housed in what is known as an opto-package on the *dogleg*, a copper/kapton flex circuit. The *dogleg* provides a connection between the opto-package and the module and provides power and temperature monitoring for the module.

2.3.2 Detector Control System

The DCS (Detector Control System) provides temperature monitoring of both sides of the hybrid mounted on a module, supply voltages to the opto-package, as well as the high voltage (HV) for the module and low voltage (LV) for the ABCD3T chips, [32]. Most of the 6 W dissipated by a module is due to the power consumption of the read-out chips although with radiation the current drawn by the silicon detector will increase. A maximum of 10 W is allowed for after radiation and this heat is removed




Figure 2-8: The two Low Mass Tapes shown here are laid on top of one other and service two modules. The doglegs which attach to the LMTs provide the electrical connectivity and the opto-package and are responsible for converting electrical signal to optical signals and vice versa.

by an evaporative cooling system, [127].

To minimize the impact of a single point failure and possible noise pickup from ground loops, each module is connected to its own low voltage and high voltage power supply channel through a dogleg and an aluminum/kapton LMT (Low Mass Tape) which runs under the modules, along the barrel. The power and the optical devices are housed on the dogleg which is a copper/kapton flex circuit that is attached to the module hybrid, [1], as seen in Fig. 2-8. For LMTs an aluminum conductor thickness of $50\ \mu\text{m}$ was chosen to minimize its effect on radiation length. These tapes require very careful handling to avoid line breaks near solder joints and also frequent continuity checks during assembly and installation. Also, the VCSELs housed in the opto-packages have no protection against electro-static discharge (ESD) damage, which requires all operations on module services to have good ESD precautions.

High voltage and low voltage cards are housed in a power supply crate. The power supply to the CMOS part of the ABCD circuitry (digital) is independent of the bi-polar part of the circuitry (analogue) which minimizes the pick-up from the 40 MHz clock and readout cycle. The Low Voltage cards which provide the modules

with V_{cc} , which supplies power to the bi-polar part, and V_{dd} , for the CMOS part, have 4 channels. The High Voltage cards provide HV from 150 V before irradiation and 500 V after irradiation to the modules. The high voltage cards have 8 channels.

The DCS system has to supply low voltage and high voltage to the modules as well as monitor and adjust them according to the voltage drop in the cables, [96]. In ATLAS, the distance between the power supply crates and the detector is ≈ 90 m in cable length and there is a non-negligible potential drop in the cables. Hence, there are “sense” lines from the module back to the power supply cards which provide a measurement of this potential drop and allow the card to set the correct voltage independent of the voltage along the cables, [104].

The low voltage cards also supply voltages to the components in the opto-package and allow for their monitoring. The toggle of the “reset” line allows for a hard-reset of the module without powering it off while a high setting for the “select” line selects the redundant TTC provided by a designated neighboring module. After a hard-reset, the module goes in to an unconfigured state, a mode known as *clock-by-2* as links return a 20 MHz clock.

In each crate, there are 12 LV cards and 6 HV cards, supplying 48 modules in total. Each power crate has three special cards:

1. CC: The Crate Controller. This card controls all the HV and LV cards in the crate and talks natively in CAN to a DCS Computer through the CAN bus. Each of these cards has a dip-switch which lets the user set the Crate number.
2. SIC: Interlock Card. This card monitors the states of the HV and LV cards and talks through a connection on the back of the crate to a safety interlock, called the Matrix.
3. Shorting card, shorts the cable shield ground to the VME ground for ESD protection while connecting cables.

The safety interlock, called the Matrix, provides the mapping between the inputs, the cooling pipe temperatures and the BOC door interlock and the power supply

cards. Its purpose is to send “enable” signal to the SIC cards in the Power Supply crates if the corresponding pipe temperatures are cold enough to be able to operate the SCT safely and if the BOC door is closed. The BOC card provides 3 mW fibre-coupled power. This laser light is harmful to the eye and the lasers are required to be disabled if the BOC door is open.

The matrix sends signals to the SIC card on 13 lines, 1 of which is the VCSEL current and 12 of which enable power supply channels in groups of 4. For enabling the VCSELs, it checks that the BOC door is closed. For enabling HV/LV cards, it verifies and routes the TTL signals from the IBOX (Interlock BOX) which indicate low temperatures, [113].

FSM (Finite State Machine) is a software structure for monitoring and controlling the power supplies. It can issue warnings and alarms when certain values are exceeded. While the HV and LV cards have “hardware” trips which turn off power supply to a module immediately, the FSM has software controlled “soft” trips which turn off power. Such warnings and alarms are listed in Table 2.1. The FSM has the job of communicating with the data-acquisition system and providing a “go-ahead” for the start of a run as well as determining and providing information on what sections of the detector are operating nominally. Information on modules which are turned off or have tripped will be inserted into the header of a luminosity block that identifies ≈ 1 minute of ATLAS running.

2.3.3 Off-detector readout

The ROD crates are responsible for the control and the read-out of the modules and sending the data to the ROB for ATLAS event building.

A ROD crate is a VME crate containing up to 16 RODs (Read-out Drivers) [87] and BOC (Back of Crate) cards [75] and one SBC (Single-Board Computer) and one TIM (Timing Interface Module) [130, 34]. One ROD crate can control up to 672 modules, or all of Barrel 6, which is the largest tested natural sub-detector of the SCT.

The crate has a front and a back side. The RODs are placed on the front of the

Parameter	Set Value	Low Alarm	Low Warn	High Warn	High Alarm
LV V_{cc} [V]	3.5	0.0	3.4	3.6	3.7
LV V_{cc} sense [V]		0.0	3.4	5.6	5.7
LV I_{cc} [mA]		0	750	1100	1250
LV V_{dd} [V]	4.0	0.0	3.9	4.1	4.2
LV V_{dd} sense [V]		0.0	3.9	6.1	6.2
LV I_{dd} [mA]		0	400	700	1000
LV V_{VCSEL} [V]	4.0	0.0	3.9	4.1	4.2
LV I_{VCSEL} [mA]		0.0	1.8	2.4	4.0
LV V_{P-I-N} [V]	6.0	0.0	5.9	6.1	6.2
LV I_{P-I-N} [mA]		0.0	0.1	0.5	0.6
LV Temp [$^{\circ}$ C]		-273	10	33	35
HV V [V]	150.0	140.0	145.0	155.0	157.5
HV I [μ A]		0.00	0.01	3	5

Table 2.1: Default DCS settings and allowed monitored values for the “ON” state of an SCT module. Monitored values do not have “set” values by definition.

crate while a corresponding BOC, which it controls, is placed on the back side of the crate. Slot #13 in the front side of this VME crate is dedicated for the TIM for a uniform distribution of timing signals to the RODs as it is the middle slot. A custom J3 back-plane is used to distribute signals between the TIM and the RODs.

The RODs and BOC cards which are responsible for the control and the read-out of the modules. The ROD is responsible for generating the command signals for the BOC communication with the modules as well as interpreting the data as received by the BOC from the modules.

The TIM is responsible for relaying TTC (Timing and Trigger Control) information to the SCT RODs from the ATLAS TTC system.

The SBC acts as a RCC (ROD Crate Controller) which also allows the user to interface with the system. The RCC has an ethernet connection, which lets the user control and configure the RODs and the BOC cards. The SBC runs the SctApi application which acts as the interface between the RODs and the SctRodDaq application, which we will discuss in detail in Section 3.4.2.

The control and readout fibres are attached to the BOC cards. Each BOC card has connections for 4 TX fibre ribbons and 8 RX fibre ribbons, which means that a

card can control and readout 48 modules. A fibre ribbon has 12 individual fibres in it, so for example, could carry RX signals from half a row, or send TX signals to two half-rows.

Now we will describe the functionality of each card in more detail.

Timing Interface Module

The Timing Interface Module (TIM) distributes the TTC signals to the RODs and the BOCs through the VME back-plane. The TTC signals have two sources. In standalone mode, the TIM can generate these signals internally or it can be under the control of ATLAS TTC system. In standalone mode, it can be configured to broadcast single triggers, trigger bursts or continuous triggers, with an adjustable rate or with a pseudo-random jitter. The standalone mode is useful for synchronously testing a crate's worth of modules in the SCT in calibration mode, especially during macro-assembly or commissioning. Under the control of the ATLAS TTC, the TIM can veto triggers if the TIM or any of the RODs are busy using a special BUSY line.

TIM has a Fixed Frequency Trigger Veto (FFTV) functionality, [21], which issues a RodBUSY signal if the triggers are in a certain frequency range that would be resonant with the wire-bond mechanical resonant frequency. Triggers initiate read out activity in the chips which lead to large variations of current flow through the wire-bonds between the readout cycle and the idle time. Since the currents flow through a magnetic field, the wire-bonds experience a force, which if the trigger rates are close to the mechanical resonant frequency, could result in the breaking of these wire-bonds. Under the random physics triggers condition, the probability of getting such fixed frequency triggers is low, but to avoid this low risk damage to this wire-bonds, a programmable TTFV is in place to veto such triggers. It is designed to veto triggers at any fixed rate in a programmable frequency range from $\approx 15 - 500$ kHz.

The common distribution of the system clock from the TIM ensures the synchronization between the RODs and the BOCs and stability of their communication. Also, the TIM distributes other TTC commands such as L1-Accept, bunch-crossing counter reset, L1 counter reset and a full front-end reset, known as a *soft-reset*. Such

resets are essential in avoiding communication failures that would arise due to single event upsets (SEU). They also enable the SCT sub-system to synchronise itself with the rest of ATLAS and identify the events with the correct bunch-crossing and L1 counters, essential for event building algorithms.

Read-Out Driver

The Read-Out Driver (ROD) is the main work-horse of the SCT calibration procedure as well as being the bank for configuring of modules in different modes.

In physics mode, it decodes the bit-stream as received by the BOC and encodes them with the appropriate identifiers, headers and trailers and transmits them back to the BOC for transfer to the ROB up the S-LINK, for further event-building. In calibration mode, the ROD has several tasks, one of which is being able to generate internal triggers and sending them at the fastest rate compatible with other running tasks. The power of the ROD comes from its histogramming task which allows the user to scan through different variables, meshing the returned millions of triggers-worth of binary information from the front-ends into a simple picture like a “scope-photo” of the full scan range.

The ROD achieves these complicated tasks by using 4 *slave* DSPs (Digital Signal Processors) for the configuration and histogramming tasks under the control of 1 MDSP (Master DSP) while 4 FPGAs (Field Programmable Gate Arrays) control the event data flow. The 4 FPGAs are:

- ROD Controller: It is responsible for interpreting and executing TTC signals from the TIM and controlling the event data flow at the top level, by signalling the formatter when to look for the received data stream from the module.
- Formatter: The formatter consists of 8 formatter FPGAs, each of which can decode 6 modules worth of data. It decodes the data into 16-bit words for transmission to the EFB(Event Fragment Builder). The formatter has two modes: expanded and compressed. The data received from the module always contains 3 time-bins worth of information. This information is kept in the

Value	Meaning
0x1	Error in header
0x2	Bit error in trailer
0x4	Error flagged by ABCD (eg. output FIFO full)
0x8	Bit error in sync
0x10	Unexpected hit pattern in the three time bins
0x20	L1 ID mis-match
0x40	BC ID mis-match
0x80	Timeout error (no data after trigger sent)
0x100	ROD buffer almost full
0x200	Data overflow
0x400	Error in chip sequence
0x800	Invalid chip number (eg. 6, 7, 14, 15)

Table 2.2: The error codes issued by the ROD event formatter.

expanded mode, while in the compressed mode, only the $01x$ pattern is tagged as a hit and the time-bin information is discarded, reducing data size significantly. The formatter also checks for errors in the header, the trailer and chip sequence and passes this information to the EFB as well. The list of errors and codes are given in Table 2.2.

- EFB: Event Fragment Builder receives the data for each link from the formatter and builds what is known as a ROD fragment. It first suppresses information from those links with no hits and no errors. Data for each module with a hit is represented in the format shown in Table 2.3 with keys explained in Table 2.4. The EFB also tags Level-1 and Bunch-Crossing (BC) errors by checking them against the ROD's internal counter and inserts this information into the data stream. Then it wraps up these module data strings in the standard ATLAS data format, called a ROD fragment, which will be explained in the next chapter by adding a ROD header and trailer to package the data.
- Router: The ROD fragment produced by the EFB, if running in physics mode, is sent by the router over the BOC and the S-LINK to the higher levels of the ATLAS event building mechanism. The router also allows for an optional configurable event-trapping mechanism to function in any run mode, copying

data to a designated slave DSP's memory for further processing.

Name	Bits [15:0] or [31:16]
Header	001pt1bKxMMMMMM
Trailer	010zhvxxxxxxxxxxx
1 hit condensed	1FFFFCCCCCxfx0
2 hits condensed	1FFFFCCCCCsfx1
1st hit cluster expanded	1FFFFCCCCCODDD
1 hit cluster expanded	1xxxxxxxx0xxx1DDD
2 hit cluster expanded	1xxxxxxxx1DDD1DDD
Flagged error	000xxxxxxxxFFFFEEE
Raw data	011nnnxwWWWWWWW

Table 2.3: The module output format as sent by the router to the DSP and up the S-link. The keys are given in Table 2.4.

x	do not care (The ROD fills these with 0's)
W	raw data
C	cluster base address
D	3 bit hit data
E	ABCD error code
F	Front-end or chip number
h	header trailer limit error
n	count of raw data
f	error in condensed mode data, 1st hit
s	error in condensed mode data, 2nd hit
p	preamble error
z	trailer bit error
v	data overflow error
t	time out error
K	condensed mode
l	L1 error
b	BCID error
M	link number

Table 2.4: The keys used in the module output format in Table 2.3.

A histogramming task runs on each slave DSP, on the events that have been trapped in the router, generating histograms for each link. The histogramming task can run in physics mode as well as in calibration mode. In the calibration mode, the MDSP is responsible for the configuration of modules, for varying scan variables, controlling the rate of triggers and providing module synchronization with soft resets

if a trigger is missed. The DSPs histogram in a new bin with the varying scan variable, keeping track of the number of events and errors in each bin. On a ROD, there are 4 *groups* of modules, defined in the configuration files. These 4 groups correspond to the 4 DSPs on a ROD. Each DSP deals with all the links in a certain group. This way, each group gets its triggers synchronously and the data from all links belonging to that group is histogrammed on one DSP.

While the histogramming is running, an asynchronous message system allows the ROD to send messages back to the host SBC, allowing for the monitoring of the progress and the errors. When the histogramming task is finished, a VME interface allows for the user to transfer the histograms from the DSP memories via the host, as well as buffers for errors and warning messages.

A software framework called SctRodDaq, which will be described in detail later, is responsible for further processing of the data, such as fitting and analyzing it. Still, a fitting task can also run on the DSPs, which further reduces the data that needs to be transferred to the host. Such a task is used for some tests.

Histogramming *on-line*, or histogramming while running physics triggers is more complicated as, for now, it requires the stopping of triggers during the start of the histogramming task as well as at points when the bins are changed and during the transfer of the histogram buffers. Triggers can be inhibited by issuing RodBusy but such inhibitions during ATLAS physics running is prohibited. Since on-line histogramming has proved to be extremely useful for commissioning and cosmic runs, there is on-going work to make on-line histogramming run without inhibiting triggers.

BOC

The BOC provides communication with the modules through the TX and RX fibres and also sends ROD-processed events via the S-LINK to the ATLAS event building mechanism.

The BOC can sample at a 20MHz rate as well as the standard 40MHz rate. In the 20MHz rate, the return from an unconfigured module is viewed as a string of alternating 0's and 1's. The receipt of a soft reset can be confirmed by seeing that

the module returns this string.

Stable opto-communications with the modules require programmable parameters in the transmission and the receiving of the data for each link, which we shall list here. For the TX links, it is essential that the clock and commands are received by the module on time and with enough optical power to convert them correctly to electrical signals. The adjustable parameters for TX links are:

- Laser current: The laser currents are adjustable in 256 steps, with a typical value around 160, which is the equivalent DAC for 10 mA laser current.
- BPM mark-space ratio: The bi-phase mark to space ratio adjusts the ratio of the two halves of the clock cycle to be 50%. It is in 32 steps of approximately 0.25 ns. An optimized setting of this ratio is important as it minimizes the timing jitter of the digitizing circuitry of the read out chip.
- Coarse delay: The coarse delay delays the clock and command sent to the module by 32 steps of 1 bunch-crossing each. The coarse delay is used to ensure that the Level-1 trigger arrives on the module 132 bunch-crossings after the triggered bunch-crossing.
- Fine delay: There are 7 bits fine delay adjustment from 0 to 35 ns, used to adjust for the fibre-length differences between different modules to synchronize them all with the system-clock and also adjust for the time of flight of particles from the interaction point, which will be discussed in Chapter 3.

For the RX links, it is important that the data received back from the module is interpreted correctly. Its adjustable parameters are:

- RxDelay: This delay adjusts the phase of RX data relative to system clock in nanosec, in 25 steps of 1 ns each, and is used to set the BOC sampling as far away from the clock-edge as possible. The absolute timing of the signal from the module is not important as it contains no additional information, but individual adjustment for each module is necessary due to different fibre lengths.

- RxThreshold: This threshold is used to discriminate the received signal, is adjustable in 256 steps and corresponds to the $p-i-n$ current in microAmps.

Such opto-communication parameters need to be adjusted for each and every single channel before a full characterisation of the SCT to avoid data errors during long test scans. How they are adjusted will be discussed in detail in the next chapter. Such optimized values need to be stored in the configuration database so that the user does not have to re-optimize the parameters every time.

2.3.4 Configuration of the SCT system

The configuration of the SCT system is possible through the loading of the configuration through the SBC onto the RODs. The RODs have 4 separate *banks* which hold the configuration. Each is used for a different purpose, such as one is for physics mode while another one is the starting configuration for calibration mode. Another bank is used for modifications to the configuration while scanning through a variable. One bank has not been used yet.

The user defines the configuration in XML (Extensible Markup Language) [44], through a well-defined schema in a database which resides in a file tree. Tables 2.5 to 2.8, give some details of the schema. This database will soon be migrated to a CORAL database by a dedicated SctRodDaq tool, [72].

While part of the configuration deals with multi-crate capability, here we will mostly discuss the configuration for a smaller number of modules, or an MUR (Minimum Unit of Readout), which forms the building block for the rest of the configuration. A ROD configuration “includes” up to 8 MUR files and will be described later. In the barrel, an MUR consists of half a row or LMT, containing 6 modules. In the end-cap, it contains 4 to 6 modules depending on the location. Each MUR is identified by its location on the barrels, given in Table 2.5. A special *geography* section in the tree specifies how the MURs relate to physical location in the SCT. Although it is trivial for the barrel, this mapping is rather complicated for the end-caps.

An MUR corresponds to an RX fibre ribbon, with 12 single fibres, which allow for

Field	Meaning	Offline definition
b	barrel number: 3, 4, 5 or 6	“layer” + 3
s	side of the barrel, 0 for Z- or 1 for Z+	none
XX	lmt row number: 01, 02,...	“phi” + 1

Table 2.5: The definition of the barrel MUR (Minimum Unit of Read-out) number: bsXX.

Element	Attribute	Contains
MUR	order	order of fibre ribbon at BOC. start counting at 0 from bottom of BOC.
	id	bsXX as in Table 2.5
module	id	Position 1..6 and references corresponding module files.
rmodule	group	Trigger group from 0..7. Actually, only 0..3 is used to route data to corresponding DSPs in calibration mode.
	id	Position 1..6 and references the redundant module. Example: < rmodule id="2" > tag means, that the dogleg at position 2 can send a signal to the module that is defined in this field.
channel	id	same as module id: 1..6
	output current	the bias current for the VCSEL
	output delay	adjusts phase of TX to synchronize modules and accommodate different fibre lengths + TOF from IP
	output markSpace	adjusts the markSpace(duty cycle) of BPM chip
	output fibre	optional definition in case of a fibring-mapping error
	stream{0,1} threshold	threshold (0..255) used to discriminate received optical signal.
	stream{0,1} delay	adjust phase of RX data relative to system clock in nanosec. (0..24)

Table 2.6: The attributes and sub-elements of the MUR (Minimum Unit of Read-out) element, and description of their contents in the XML-based configuration files.

the readout of the full MUR, hence the name minimum unit of readout. Each MUR corresponds to a BOC *order*, which is the plug-in for its RX fibre at the BOC. It is natural that the 0 th and the 1 st RX plug-in correspond to the 0 th TX plug-in at the BOC. This is “assumed” to be true, unless there has been a fibring-mapping error, which can be corrected in the configuration, by an optional *output fibre* definition, as shown in Table 2.6. Most of the parameters in the MUR definition have already been introduced in the BOC section. Here, we will clarify a couple of points.

TX redundancy flows in the barrel, in loops of 12 modules. The loops never cross $\eta = 0$ and flow on two adjacent MURs, on consecutive barrel half-rows. Such half-




Figure 2-9: Schematic drawing of a clockwise TX redundancy loop. The orientation of the redundancy loop is determined with the lower LMT number on top.

Element	Attribute	Contains
rod	id	From 0..13 in each ROD crate.
	slot	The RodCrate slot where the ROD card is.
	RODId	The unique ATLAS “source identifier” for use by the EFB.
	ROBId	Corresponding ROB id, as defined on the ROS.
	“xi:include href=...”	Includes the MUR files

Table 2.7: The attributes of the ROD element, and description of their contents, in the XML-based configuration database.

rows are required to be serviced by the same TX fibre, allowing for the continuity of the TX-redundancy loops. The loops may flow in the clockwise or anti-clockwise direction depending on their location on the barrels. How the orientation can be determined from the redundancy flow is shown in Fig. 2-9. The 1st and the 6th module on each half-row either gets or provides TX redundancy to the module in the same position on the next half-row. It is worth noting here that the connectivity of TX redundancy loops have only been tested on Barrels 3 and 6 due to time constraints on the macro-assembly process.

The TX fibres are identified in the schemas as *output* having the current, delay, MSR adjustments as well as the optional fibre definition. The configuration file has two 8 bit words to store the information for the TxDelays: *0b 000cccc 0tttttt* where *c* = coarse and *t* = fine delay in binary.

RX fibres used to transmit data from the top and the bottom faces of the module are identified in the MUR schema as *streams*, with *stream0* referring to the top side of the module and *stream1* referring to the bottom side. Both RX streams have individual threshold and delay adjustments as described already in the BOC section.

Up to 8 such MUR definitions are included in a ROD definition, whose schema is given in 2.7. The ROD definition forms the ultimate building block for a simple physical *crate* definition which includes up to 14 RODs. Each ROD in ATLAS is uniquely identified through the 32-bit *RODId* defined in this schema. Such information is inserted by the EFB into the header of the ROD fragment to identify it. The first 2 bytes of the *RODId* identify the SCT and which ROD crate the ROD is in. The last byte is the same as the *id* defined in the first attribute. The *ROBId* contains the same SCT identifier as the *RODId* but in the last byte identifies the S-link plug-in on the ROS, and hence the ROB, that the ROD is connected to. The SCT identifiers are listed in Table A.1 in Appendix A. For the SCT, there is a one-to-one correspondence between ROD and ROB, which is not necessarily true for other detectors. The *ROBId* is used in the building of the ROS fragment, which combines ROD fragments from the same crate. ATLAS event building will be described in the next chapter.

At the top level of the XML-tree, the DCS mapping of the modules and the default values, alarms and trips are defined. The DAQ system can communicate with the DCS. Control and synchronization of values is especially important for special cases such as modules reading out from one stream or modules working under redundant TX connection, which require changes not only in the DAQ configuration but also the DCS configuration.

Another important part of the XML-tree is the configuration of the ROD itself. The code that runs on the ROD DSPs has gone through several iterations and these files are also specified in the XML file to be loaded onto the RODs at DAQ start-up.

SCT Module Configuration

All the module configuration files are included in the top level XML-tree and referred to in the MUR definitions. Generally, the module configuration files mentioned above are stored under `BarrelB/lmtXX/module/20220ss0200nnn.xml` in the XML directory tree. The module serial numbers are 14 characters long where *ss* represents production site with 04 defined as the U.S., 17 as Japan, 33 as U.K. and 38 as Scandinavia. *nnn* is the module number.

Element	Attribute	Contains
module	sn	module serial number
	active	whether the whole module is enabled active or not. is 1 by default unless something is wrong.
	select	is 0 by default. select = 1 means that the modules will use the redundant TX signal chain. Note that power supplies in DCS need be configured for this redundancy separately.
chip	id	ABCD chip id, defined as 0..11
	address	hardware address
	active	whether the chip is enabled active or not. is 1 by default unless something is wrong.
	config	Chip configuration register, explained below in Table 2.9.
	mask	is all f's unless some channels are masked off before the input to the pipeline. 4× 8-nibbles corresponding to 128 channels
	c_factor	Calibration factor
	rc_function	Response curve function

Table 2.8: The attributes and sub-elements of the module element, and the description of their contents, in the XML-based configuration database.

A module configuration is sent to each module by a TX fibre using special command header. Each chip is responsible for identifying the part which addresses itself by a comparison of the hardware addresses and changes its configuration accordingly.

Part of the schema of the module definition is given in 2.8. Besides the parameters given here, there are five more chip configuration parameters. One sets the threshold in the mV equivalent of 1 fC collected charge. Two settings control the size and the timing of the injected calibration charge. Another two settings optimize the performance of the front-end amplifier to compensate for radiation damage. There is also a 4-bit trim setting for every single channel as well as chip-wide trim-range setting to compensate for the variations in response from the front-end amplifiers. Two attributes require further discussion as they directly effect how well the ABCD circuitry can be calibrated: calibration factor and the response curve. They will be discussed in depth in the next chapter.

The chip configuration register sets up the chip readout and token passing schema. Calibration line sets which group of channels receives the calibration charge. During the calibration, this line is changed to inject charges to all channels.

Field	Name	Value	Default
1	Read-out Compression Mode	Hit(0), Level(1), Edge(2) or Test(3)	2
2	Calibration Line	0, 1, 2 or 3	Any
3	Trim Range Setting	0, 1, 2 or 3	Ideally low
4	Edge detect	Not level mode	0
5	Send mask	Copy the mask register into the pipeline	0
6	Accumulate	Specify accumulate mode	0
7	Input bypass	Uses the bypass connection for input tokens if 1	0
8	Output bypass	Uses the bypass connection for output tokens if 1	0
9	Master	0 to specify master chip ^a	0 for master only
10	End	1 if end chip	1 for end only
11	Feed-through	Normally 1, set to 0 to get clk/2	1

Table 2.9: Content and default of the chip config attributes from Table 2.8. The ROD sets these bits in the ABCD configuration register.

^a1 enables the master chip to be used as part of a read out chain from the other side

RX Redundancy scheme

From the chip configuration register description, one can implement the chip-bypassing scheme described in Fig. 2-6. The last five bits of the chip configuration register need to be changed to reflect that one side of the module will be read out from the other side. If there is a modified module available in place, then the information from the master chip can also be included in the readout by asking that side's end chip to pass its token to the master chip. If there is no modified module, the end chip stays as an end-chip and the master chip is not read out. Here, we give the changes needed in the configuration registers if either link is dead.

Reading through Link 0			
	Normally	Modified Module	No Modified Module
Chip 5 or E5	0 0 1 1 1	1 0 1 0 1	1 0 1 0 1
Chip 6 or M8	0 0 0 0 1	0 0 1 1 1	0 0 0 1 1
Chip 7 or S9	0 0 1 0 1	0 1 1 0 1	0 1 1 0 1
Chip 11 or E13	0 0 1 1 1	0 0 1 0 1	no change
Reading through Link 1			
	Normally	Modified Module	No Modified Module
Chip 0 or M0	0 0 0 0 1	0 0 1 1 1	0 0 0 1 1
Chip 1 or S1	0 0 1 0 1	0 1 1 0 1	0 1 1 0 1
Chip 5 or E5	0 0 1 1 1	0 0 1 0 1	no change
Chip 11 or E13	0 0 1 1 1	1 0 1 0 1	1 0 1 0 1

Table 2.10: The last five bits of the ABCD chip configuration register need to be changed if reading through one link.

Chapter 3

SCT Barrels: Production to Commissioning

The SCT modules have been tested several times from production, [110] to installation in the pit. Here, we will describe in detail the testing starting with macro-assembly and ending with the commissioning in the pit. Besides the production and testing of the modules placed on the barrels, there were system tests of both the barrel and the end-cap modules where complete sectors similar to the final detector were assembled. In such system tests, different grounding schemes[68], long term stability [101] were tested as well as the noise and standard performance figures. We will describe how to quantify the performance of modules in detail in this chapter. Before that however, it is worth mentioning a few words about how the modules were produced and qualified for barrel assembly.

3.1 Calibration of binary readout

The testing of SCT modules is a check of their calibration and performance and also of the test system itself. The key to the characterization of the modules lies in reconstructing the analogue response of the modules from the binary readout by first setting the optimal timing for the charge injection, then injecting a set of charges into the front-end and then scanning through the threshold to map out the response




Figure 3-1: Occupancy as a function of front-end discriminator threshold. The circles represent data points while an S-curve fit is overlaid.

curve.

The optimal time for the internal charge injection is determined by scanning through the strobe delay parameter which delays the time of the charge injection with respect to the read-out trigger. For this scan, a relatively large charge is injected into the front-end with a low discriminator threshold, so when the correct timing is set, there is always a hit. Scanning through strobe delay results in a “top hat” function. If the strobe delay is too low, the charge detection efficiency is 0% and as the strobe delay is increased, the efficiency quickly rises to 100% showing the timing of the tail of the response, while increasing it further shows the rising edge of the response as the efficiency sharply falls again to 0%. It is important to inject the charge at the top of this “top hat” distribution, while keeping away from the edges.

In the characterization, the range of injected charges covers the full operational regime of the SCT front-ends, from 0.5 fC to 8 fC, accounting for charge sharing and Landau fluctuations. For each injected charge, an error-function fit, also known as an *S-curve*, is performed for all channels, as shown in Fig. 3-1. At low thresholds

compared to the injected charge, the probability of recording a hit is 100%, while at higher thresholds the probability of getting a hit is 0%. The point at which the probability of getting a hit or a miss is equal provides the appropriate threshold corresponding to the injected charge. This point is known as $V_{t_{50}}$. Determination of this point is achieved by sending different number of triggers to the modules by the ROD for each threshold. A low number of triggers are issued at very low thresholds as the occupancy is high. With higher thresholds, the number of triggers are increased to the level allowed by the bandwidth as occupancy starts dropping. At these thresholds, up to millions of triggers may be sent since deviation from the expected low occupancy might indicate pick-up or common mode noise issues which need to be flagged and resolved for a trustable read-out system. Besides the determination of the $V_{t_{50}}$ point the standard deviation of the S-curve is also calculated, and is a measure of the output Gaussian noise on each channel.

The $V_{t_{50}}$ point for each injected charge is plotted against the known injected charge and is fitted with a quadratic response curve function. Although the response curve is generally nearly linear, the fit allows for the observed slight deviations from a linear curve at low thresholds. The response curve for each channel is summarized by several parameters such as the offset defined as $V(0)$, the gain defined as dV/dQ at 2 fC, the output noise and the input noise in electrons also quoted for 2 fC. Quoting the summarized values at a higher threshold of 2 fC ensures that the information is extracted from the linear region of the curve. By looking at the distribution of these extracted parameters over all the barrel modules, we can make a coherent picture of the performance of the barrel system and this will be presented at the end of the chapter. The rest of this chapter now will be dedicated to the production, macro-assembly and testing of the modules and will lead to the full characterization results.

3.2 Module Production

Before module assembly, each ASIC underwent an extensive testing and qualification process, which checked first digital and then analog performance, [12]. One important







Figure 3-2: The distribution for calibration factors for all chips on all barrels. Overall, the mean calibration factor is 1.036 with spread of 0.034 and a systematic uncertainty of 0.020.

constant measured during these tests is the *calibration factor*. The analog circuitry of the chips can be calibrated for the 1 fC working point only if there is an absolute measure of the charge injected by the calibrating capacitor. The capacitance is inversely proportional to the thickness of the gate oxide layer, called *TOXP1N_PLUS* in [31] and is measured for each wafer. The ratio of this thickness to the nominal design specification of 420 Å is called the calibration factor [123], and needs to be stored and used for a uniform calibration over all chips. The calibration factor and other values of these tests were transferred to the SCT “Geneva” database for later use in calibration and qualification tests, [116]. The calibration factors included in the module configuration represents these values as they have been downloaded from this database. The measurement error in the calibration factors directly affect how accurately the ABCD circuitry can be calibrated.

Figure 3-2 gives the distribution of the calibration factors for each barrel. From the general distribution, we find that the average calibration factor is 1.036 with a statistical variation of 0.034. However, the systematic error on the measurement of the calibration factors itself is substantial. The calibration factor is measured for each wafer at the 5 test structures on the wafer. The deviation on the measurement of these 5 test structures is 0.020, showing that the thickness of a wafer is not very uniform. Only the average value is recorded for each wafer in the database. There are 256 chips manufactured from a wafer and normally about 20% of the chips from one wafer qualified for module assembly. Although the calibration process accounts for the c_factors by correcting the injected charge by this factor, this variation of 2% arising from the inter-wafer thickness variation can not be taken out and is the largest source of chip-to-chip variation in the calibration process.

The SCT modules were selected carefully for each barrel according to their electrical and noise performance, measured at their production sites and reception at Oxford, [33]. The end-cap modules were similarly tested and qualified for each disk, [97].

As irradiation will require modules to be biased at a higher voltage than the starting value of 150 V, all modules were tested to check leakage currents up to

500 V. Those modules with the least leakage currents and with the lowest noise were placed on Barrel 3, the barrel which will be irradiated to the highest dose. On the other hand, Barrel 6 contains some modules which can not be biased at 500V but this will not be necessary since this barrel will not receive enough radiation to warrant such a voltage.

3.3 Macro-Assembly and testing at Oxford University

All modules were initially tested in groups of 4 modules inside a test-box at Oxford reception, using a system called SctDaq, [66]. SctDaq forms the basis of the final testing software called SctRodDaq, but read out modules electrically instead of optically. Differences in algorithms and performance of SctDaq and SctRodDaq are documented in [99]. SctDaq was for small scale testing of modules, during the test-beams, and for the irradiation studies, as well as Oxford reception testing. Some results from this reception testing will be shown later in the results section in this chapter. The aim of this testing was to identify modules whose characteristics may have changed since their initial testing at their production sites and to re-calibrate all modules uniformly under the same conditions. This qualification ensured that the modules had not been damaged during shipment and had not developed further defects.

The accepted modules were placed on the barrel by two purpose built robots in a clean room, [122]. First the lower modules and then the upper modules were assembled for each row. Barrel 3 was macro-assembled first, followed by Barrel 6 and Barrel 5 and finally by Barrel 4. Figure 3-3 shows the completed Barrel 6.

The testing of the SCT modules, starting with the macro-assembly phase, was performed using an online software, called SctRodDaq. It has the tasks of configuring, calibrating and controlling the modules and analyzing the data. Here we shall describe the setup, procedure and software for the testing of the barrels. The testing of the SCT end-caps was done in a similar fashion as the software system is shared. Testing




Figure 3-3: A photo of Barrel 6 completed and tested, ready to be sent to CERN. The services extend horizontally out at each end. The support spindle is visible at the right hand end. The modules and doglegs are visible in the centre.

of the modules generally took place after a whole cooling loop worth of modules had been placed on a barrel, which is 4 rows.

The regular warm testing of modules that were newly assembled on a cooling loop began after the lights in the testing room were off and the humidity was lowered with a dry air supply to avoid condensation. Then a C_4F_{10} cooling system was started. After checking the temperature of each module was below 20°C , a point comfortably low to avoid thermal run-away, the modules were powered. After being powered and configured, the modules would warm up to about 28°C .

After a barrel had been fully assembled, a whole barrel “cold” test was performed on all barrels except for Barrel 5. Barrel 5 cold test, which due to timing constraints could not be performed at Oxford, was performed later at CERN. During cold testing, a different cooling system, running on C_3F_8 , similar to the final ATLAS cooling system, was used. When the modules were powered up during this period, they are colder at around 10°C as the coolant is different and the whole testing room is also cooled down. Due to the high number of modules being tested at the same time, the

rather large data output also tested the capabilities of the SctRodDaq and the DCS systems. A long testing sequence was employed to ensure that the modules perform according to specifications during warm and cold testing.

3.3.1 SCT Hardware

The readout for testing of barrels at Oxford was performed using one of two 9U VME 64 crates which is enough to read out the whole barrel 6. The spare VME crate was a test-bed for new software with the test-box of two modules and later for the simultaneous testing of Barrel 5 and Barrel 6 during construction.

In the BOC, the Z- cards and fibre ribbons were placed in the corresponding Z- side of the crates and the Z+ were placed accordingly. The fibre mapping was sequential in RX and TX fibres from Z- to Z+ with care taken not to break the TX redundancy loops.

For the DCS side, a maximum of 14 power-supply crates were used. The crate number determined the power-supply mapping, as crates with even numbers (starting from 0) belong to the Z- side and crates with odd numbers belong to the Z+ side.

A configuration file was generated by the DAQ and was copied to the DCS for its configuration. This file called *Cnf_Map.txt* contains the mapping of modules to HV and LV cards, as well as default settings. In Table 3.1, an example of Barrel 3 “cold” run mapping is given.

Rows	Crate ID z-	Crate ID z+
LMT 15-22	0	1
LMT 23-30	2	3
LMT 31-6	4	5
LMT 7-14	6	7

Table 3.1: Mapping of groups of LMTs to power-supply crates for Barrel 3 testing at macro-assembly.

Table 3.2 gives the mapping of power channels for LMTs 15 – 22 Z+ as an example. The rest of Z+ side followed the same pattern. The channels listed here also correspond to the number on the connections (minus 1) on the back of the power

crates. The Z- side was the mirror image of this mapping. For the front side of the PS crate, the mapping of the cards is shown in Table 3.3.

LMT	+6	+5	+4	+3	+2	+1
22	47	46	45	44	43	42
21	36	37	38	39	40	41
20	35	34	33	32	31	30
19	24	25	26	27	28	29
18	23	22	21	20	19	18
17	12	13	14	15	16	17
16	11	10	9	8	7	6
15	0	1	2	3	4	5

Table 3.2: Mapping from LMT to power-supply crate channels for part of Barrel 3. Other LMTs follow the same pattern.

Slot	Card	Slot	Card
0	SIC	11	LV
1	CC	12	HV
2	LV	13	LV
3	HV	14	LV
4	LV	15	HV
5	LV	16	LV
6	HV	17	LV
7	LV	18	HV
8	LV	19	LV
9	HV	20	Empty
10	LV		

Table 3.3: Power-supply crate slot to card-mapping

3.4 SCT Online Software

The SCT online software, called SctRodDaq, is the software interface between the SCT readout hardware and the rest of the ATLAS DAQ. It controls and configures the hardware, and collects the data coming from the modules. It can analyse the data to monitor the detector or to determine calibrations. It also enables the data path from the RODs to the ROB. It depends heavily on the general ATLAS TDAQ package, [134], described in the next section. Section 3.4.2 then describes SctRodDaq

in more detail. Fig. 3-4 illustrates the main applications in SctRodDaq and their communication lines.

3.4.1 TDAQ

The ATLAS Trigger and Data Acquisition software framework is known as TDAQ. The data acquisition part provides controls, configuration, data flow and monitoring systems. A TDAQ configuration is a collection of hardware and software which allows readout of detectors. The configuration defines the control tree and endpoint applications.

The control tree consists of “segments” which are readout sub-systems, such as an SCT end-cap, the event builder or a monitoring application. Applications can have “resources” which can be enabled or disabled, such as a ROB channel. This structure lets the user have control over detector, sub-detectors, and so on down to modules.

The concept of partitions in TDAQ allows physicists to work in parallel without interference. A partition manager allocates hardware and software elements to partitions, ensuring no element is in more than one partition. Running a TDAQ configuration is often referred to as “running a partition.” For macro-assembly and in calibration mode, the SCT readout is run as its own partition. For cosmics testing, in physics mode, the SCT and the TRT read out under the same partition. During ATLAS-wide physics data-taking, the SCT could drop out of the ATLAS partition, and set up an SCT-only partition to run calibration triggers.

TDAQ is implemented as a finite state machine, defining different states that the machine can be in, what steps are to be taken during transitions from one state to another and what each sub-system can do in each state, [93]. This has advantages for both physics mode and calibration mode. The IGUI allows the user to make transitions between allowed states as follows:

1. Boot: At the boot command, TDAQ creates the control tree and then launches the endpoint applications. Once booted, the DAQ has switched from the controller’s internal BOOTED state to the INITIAL state.

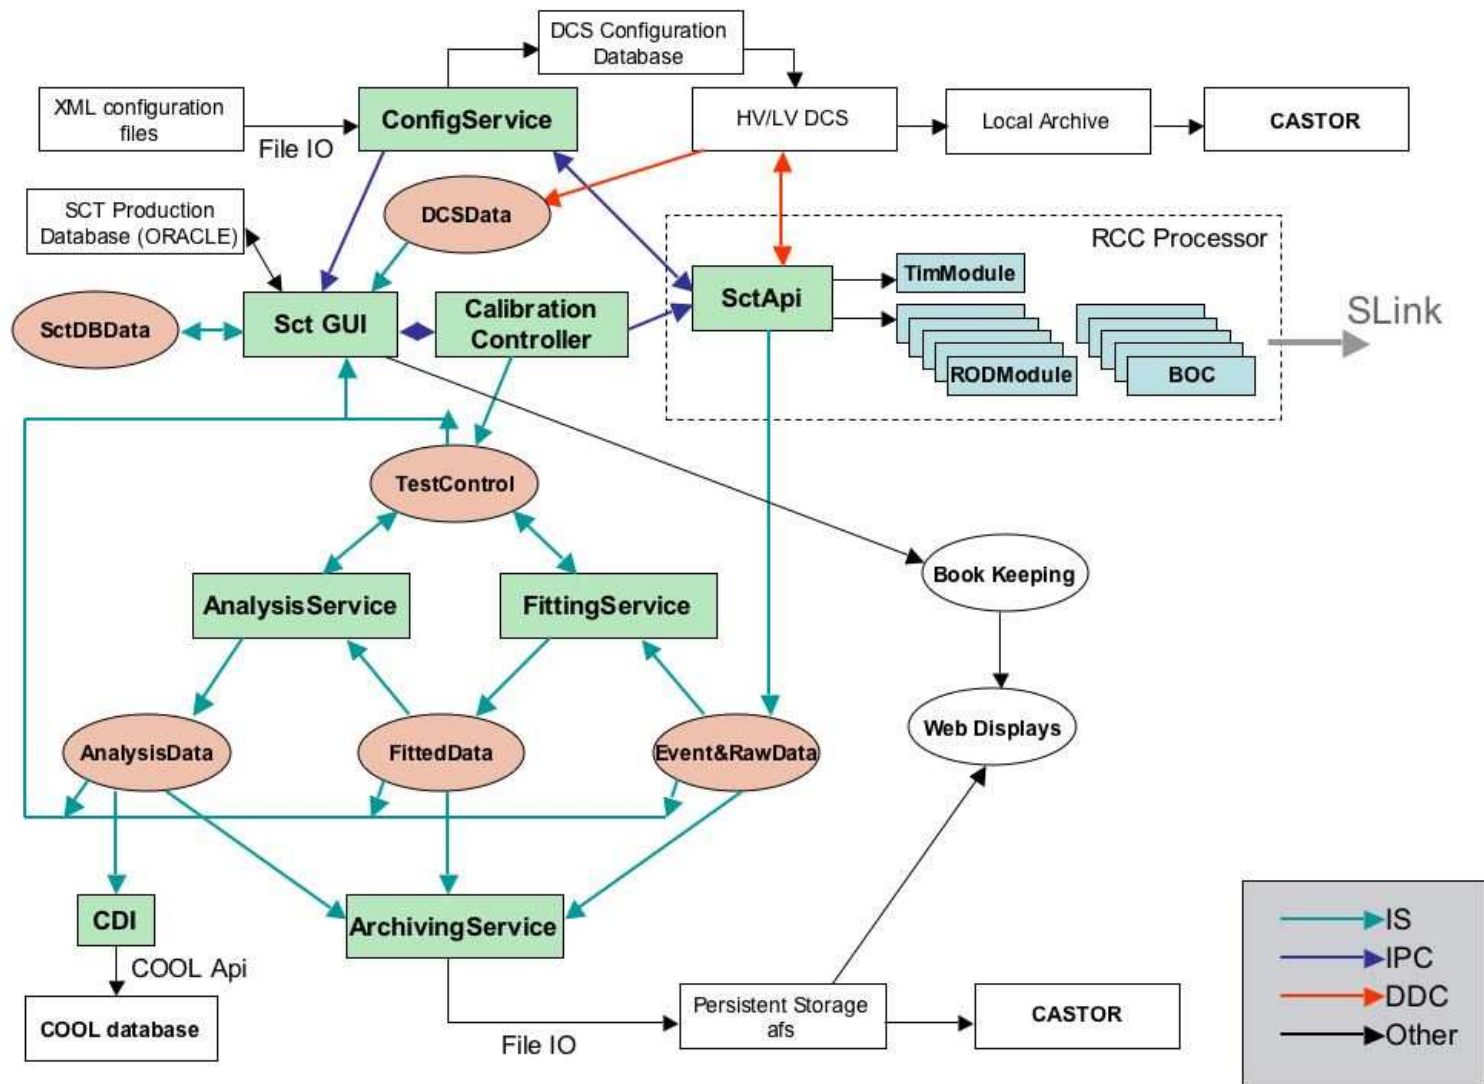


Figure 3-4: The data-flow in SctRodDaq while in calibration mode. Applications are represented by green rectangles while data objects are represented ellipses.

2. **Configure:** In this state transition, the endpoint applications read relevant configuration information from a database and configure their respective segments. On completion of this state transition, the DAQ is in the CONFIGURED state and for example, the RODs and the SCT front-end chips are configured.
3. **Start:** This state transition is foreseen to be fast as the computing intensive tasks for preparing to run are in the configure transition. This state transition will be used to reload only those parameters that have been changed since configure or those that are only valid for a single run. The endpoint applications are readied for accepting triggers and receiving data. The trigger path is enabled if in physics mode. All busy signals except the CTP (Central Trigger Processor) busy signal are switched off. As the last step, the busy of the CTP is removed to allow triggers to flow. Once started, the DAQ is in the RUNNING state and the data acquisition commences.
4. **Stop:** In the first step of this transition, the CTP is set to BUSY so that the triggers stop. TDAQ stops the rest of the applications asynchronously to allow for queued event fragments to be sent to the event builder. The monitoring applications are stopped last. The state returns to CONFIGURED.
5. **Unload:** This transition un-configures the endpoint applications and unloads the database. The state returns to INITIAL.
6. **Shutdown:** This transition should be used sparingly and only if there is no plan to run the same partition soon. Ideally, the operator should leave the partition infrastructure running and change the partition if necessary. The state returns to BOOTED.

The TDAQ framework has over 100 packages, some used for higher level online control and processing. Some packages, such as the event format library, are also used in Athena and so are called TDAQ-common packages. Here we give a brief description and functionality of packages that are relevant to the scope of this thesis. More detailed information about TDAQ components can be found in [59].

- IGUI: The Integrated Graphical User Interface is started by the command 'play_daq' and lets the user control and monitor the DAQ.
- IPC: The Inter-Process Communication package provides communication between processes that run under the same "IPC partition," regardless of which computer they are running on. If any detector needs debugging during ATLAS physics running, IPC allows it to be taken out of the general ATLAS partition and run in calibration mode on its own in a different partition. IPC is implemented by the use of CORBA, [10].
- IS: The Information Service facilitates the sharing of data between different endpoint applications. Just where the data is is transparent to an application, allowing distributed computing. It is built on top of IPC, [91].
- CDI: Conditions Database Interface is used to store and retrieve IS objects in the ATLAS conditions database called COOL.
- DDC: DCS-DAQ Communication is used to pass data and commands between the DCS system and the DAQ applications. For example, conditions data from the DCS control computers can be published into IS. The required power-supply voltages can be passed from the IGUI to the power-supply project.
- OH: The Online Histogrammer lets the user specify and make histograms online and publishes these histograms into IS.
- MRS, Message Reporting Service allows applications to report errors to other applications. Most commonly, they are sent to the IGUI application which presents them to the user. TDAQ is migrating to a new system called Error Reporting Service (ERS) which is more flexible.
- OKS: OKS (Object Kernel Support) [86] is a library to support in-memory objects, in a way which can be persistent (long lasting, e.g. between runs) and active (fast access). It is used for example to store and retrieve the TDAQ configuration.

- EMON, GNAM: These are online monitoring packages that can be attached to different HLT levels. GNAM is a light-weight package that runs with its own byte-stream converter, while EMON uses the Athena-monitoring package and allows the use of the full Athena cabling, byte-stream converter and database packages.
- RCD: ROD Crate DAQ is the ATLAS standard for communicating with a VME ROD crate. SctRodDaq does not use RCD. However, it does use the RodDaq and VMEInterface packages that RCD depends on, [49].
- DVS: The Diagnostics Verification System is the diagnostics and monitoring tool for TDAQ setup and run phases [88] and is very useful for debugging a running TDAQ partition. It has access to MRS application logs, files and IS objects.

3.4.2 SctRodDaq

The software to read out the SCT is called SctRodDaq. It is built on TDAQ. It is designed to be scalable and runs in a very distributed environment. It was written to be usable starting with the macro-assembly of the barrels to the end of LHC running.

SctRodDaq is a large system with about 250,000 lines of code in about 2000 files excluding auto-generated code and the TDAQ software it relies on. It is written in C, C++, Java and IDL (Interface Definition Language), [114]. Maintenance is expected to be carried out by several generations of graduate students and so modularity and good documentation of the code is essential. The code was designed to be scalable and runs in a very distributed environment. Flexibility is required to cope with developments in TDAQ as well. SctRodDaq has evolved with the TDAQ releases which are documented at [89]. The evolution of SctRodDaq is documented on the Wiki, [23].

There are two main running modes: calibration and physics. In calibration mode, the TIM or ROD controls the triggers sent to the modules. The ROD histograms the data and sends the histograms to the SctRodDaq for analysis to extract calibrations

such as required threshold-settings. In physics mode, the TIM listens for triggers from the LTP (Local Timing Processor) and SctRodDaq controls and monitors the flow of data from the RODs to the ROB. It is also possible to histogram the data in the RODs in physics mode.

SctRodDaq separates the task of filling the histograms from the analysis and storage tasks. This is crucial for processing the large amount of data (17 GB of histograms for Barrel 6 only) in a short time. The histogramming is massively parallel, running concurrently in all RODs. The analysis can be shared between a large number of PCs. The multi-crate functionality has only recently been tested on a large scale during the commissioning of the barrels in the pit.

Fig. 3-4 shows the main applications and their links. The following sections describe the functionality of each of the main applications:

SctGUI

This is the interface between the user and SctRodDaq. It is started from the command line for interacting with the SCT. The user can select what type of calibration runs to do and set parameters for different scans. It displays views of which modules are active in the current configuration in a rolled out view for the barrel and disc-by-disc view for the end-cap geometry. It can also display information by ROD crate and ROD-id.

Results of scans can be displayed on top of the module layout with a user-defined color range. It also displays DCS conditions, which it accesses via DDC. Fig. 3-5 shows an example SctGUI view of the results of an RxThreshold scan for all link 0's in Crate 3 in the pit. The box in the bottom-left corner shows which crate view is shown. In the middle, each column corresponds to one ROD and the order of modules corresponds to the actual cabling of fibres in the BOC (back of the crate) corresponding to that ROD. At the same time, an RxThreshold test is running and its progress can be seen in the upper right hand corner. Soon the data will arrive and be listed in the right center box and the RxThreshold values for both Rx links will be automatically updated.

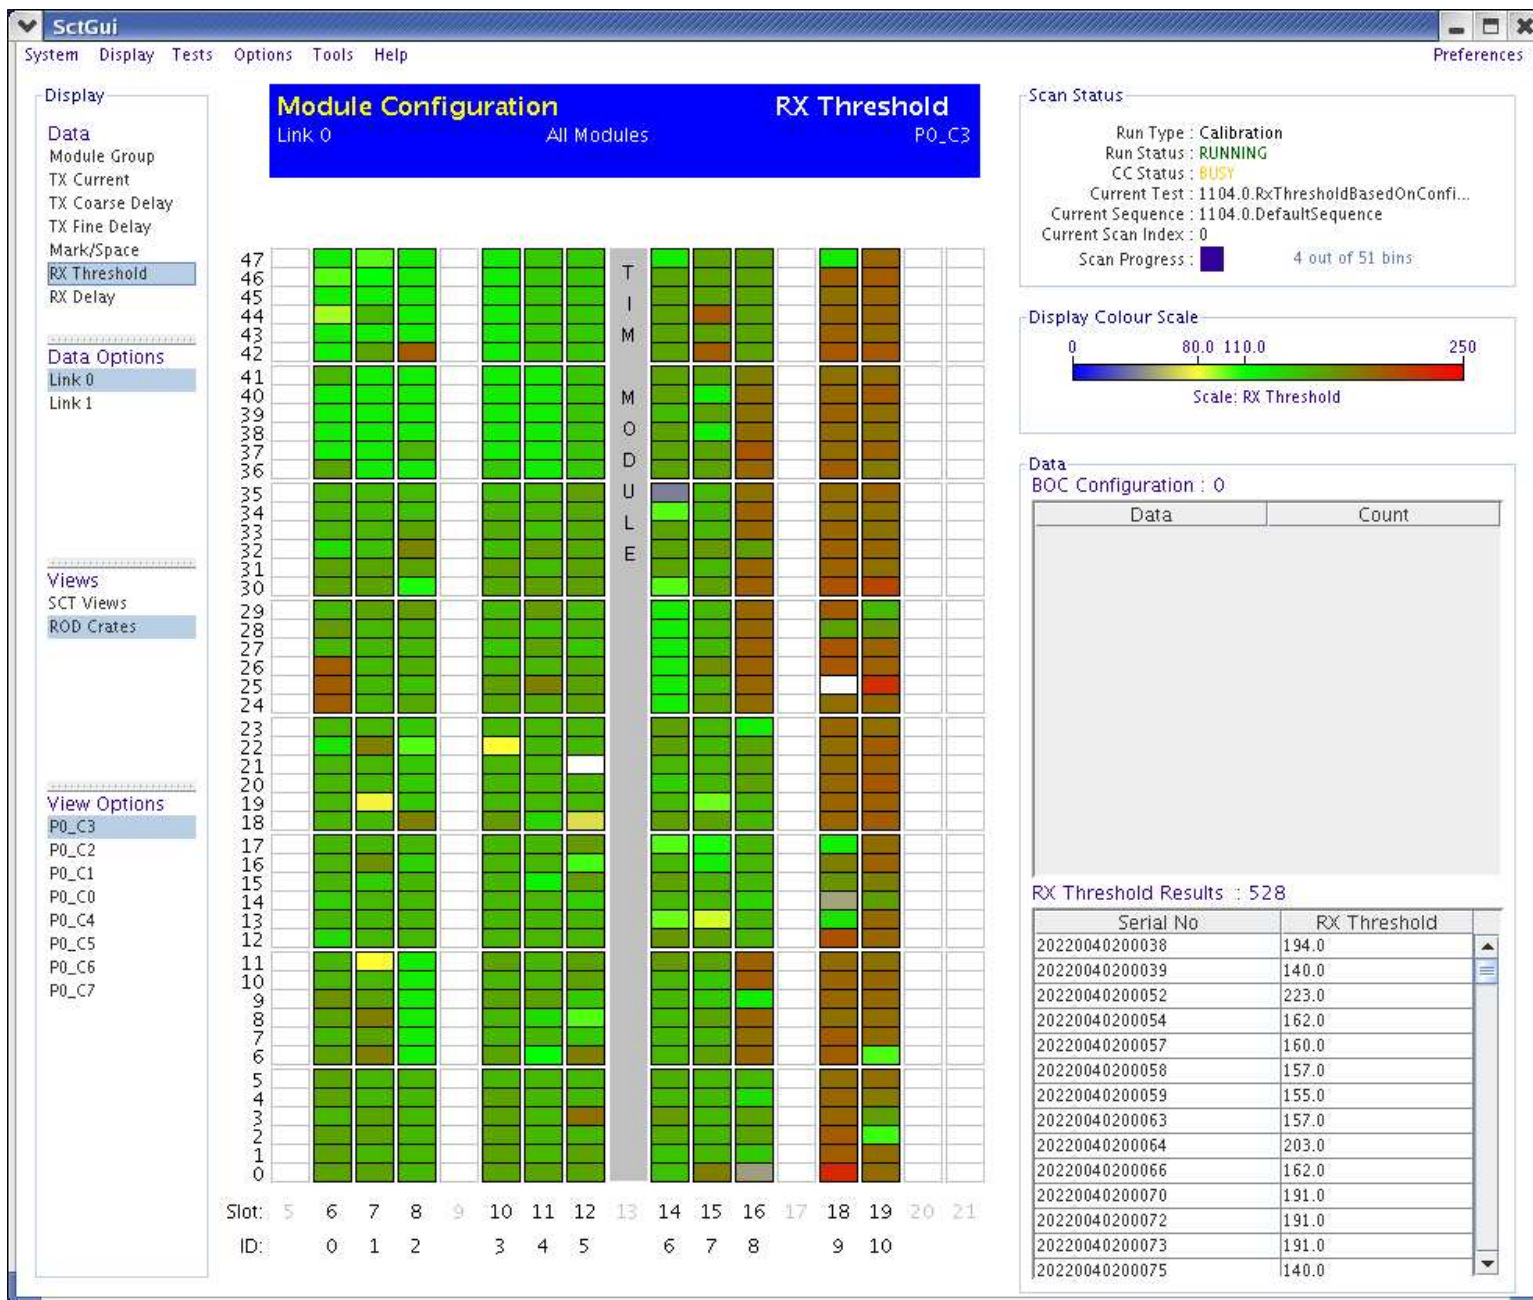


Figure 3-5: The SCT GUI here is showing the RxThreshold values from the original configuration for one ROD crate, while at the same time, showing the progress of a running RxThreshold scan on the right hand side. Once the test is finished, the RxThreshold values will be automatically updated.

ConfigService

The Configuration Service reads the configuration files and distributes the information to the SctApi applications. The configuration files are xml-based and contain information such as which modules to include in the readout and all the module parameters to program the ABCD chips.

Calibration Controller

The Calibration Controller receives commands for what scan to perform from the SctGUI. It prepares a description of the scan and distributes this to the SctApi.

At the end of a calibration run, the calibration controller checks the test results, and updates the configuration parameters of modules where appropriate. For example, at the end of a 10-point gain test, the response curve function is updated and the defective channels are masked off.

During a calibration run, the calibration controller controls the DAQ. The SctGUI changes its state from “InControl” to “Running,” once the run starts. When the run ends and all results have been published in IS, the calibration controller hands the control back to the SctGUI which changes its state back to InControl.

SctApi

SctApi is the interface between the ConfigService plus Calibration Controller and the TIM plus RODs. Each ROD Crate runs its own instance of SctApi. These instances are indexed with the ROD Crate number, allowing the ConfigService and Calibration Controller to send commands to the appropriate crate.

SctApi configures the TIM either to generate triggers in calibration mode or to listen for triggers in physics mode. It configures the RODs according to the information received from the configuration controller. Most communication with the RODs is via “primitives” - a list of commands which are downloaded to the ROD to be executed, [71]. SctApi prepares the primitives according to the scan type. Before a run, a primitive is downloaded to make the RODs do a “pre-scan hardware check,” to

make sure basic communication with the module is working. After that, primitives are downloaded to carry out the scan and fill histograms.

During histogramming, SctApi can read and write to the ROD DSP memory. This allows a thread of SctApi to monitor the run progress. At the end of histogramming, SctApi reports back to the Calibration Controller and reads the ROD histograms out. These are wrapped with information such as scan details, number of triggers, a time stamp and then published into the Event and RawData IS servers.

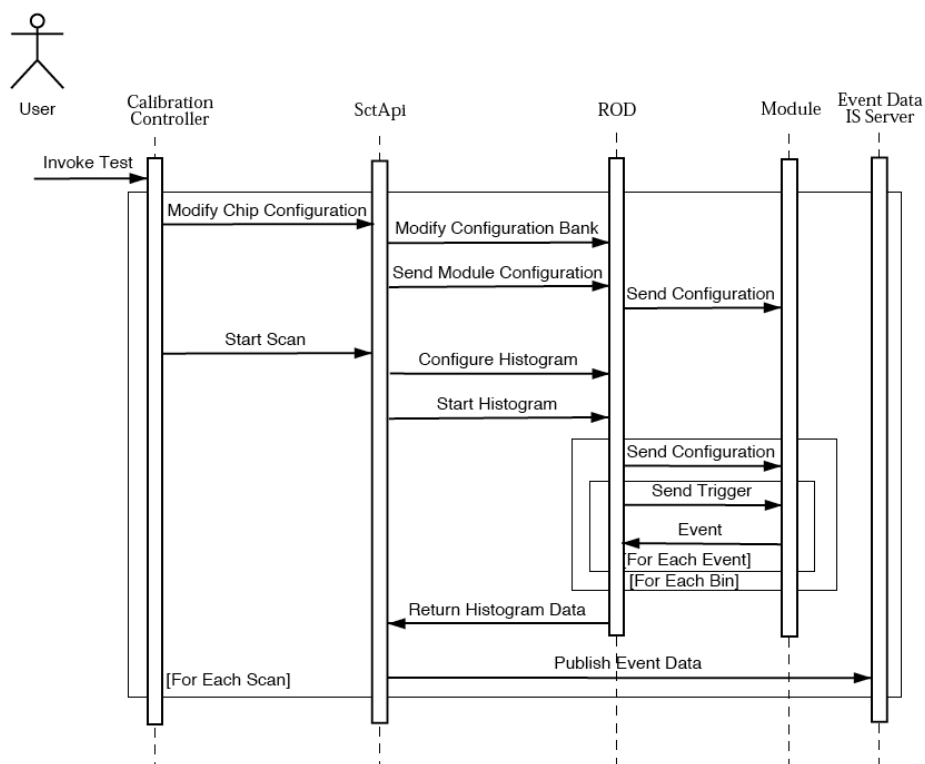


Figure 3-6: The schematic for the control and monitoring of the histogramming task up to the stage where the EventData is published, [71]. Tall rectangles represent hardware and software components. Arrows indicate data and command flow. Large rectangles indicate loops. The rest of the data flow is shown in 3-7.

FittingService

The Fitting Service makes fits to histograms – for example, S-curve fits to threshold scans. When SctApi publishes histograms to IS, IS notifies the Fitting Service. The

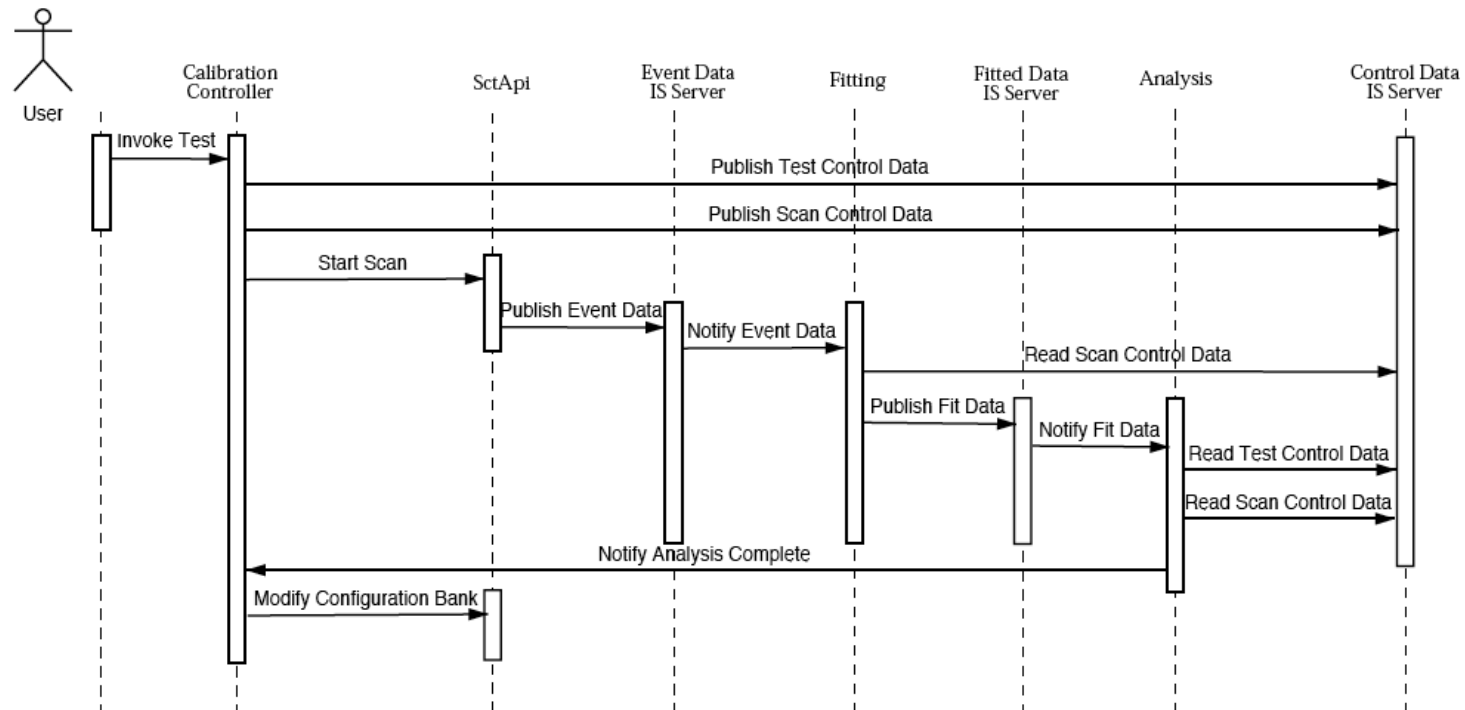


Figure 3-7: The schematic for the control and monitoring of the histogramming task and the processing of the histogram data, [71].

Fitting Service adds it to its queue. It processes this queue, popping histograms from the list, reading the scan type and making the appropriate fits. The results of the fits are then published via IS to the FittedData IS service.

AnalysisService

The Analysis Service processes the fit results and extracts the important parameters and identifies defects. It makes summaries for each module, such as chip average and RMS deviations for parameters. It also makes a defect list, with details of the defective channel or chip.

Analysis Service monitors the FittedData IS service and adds items to its queue when they arrive. If a test has multiple scans, it waits until all scans have been published. It then processes this queue, producing TestResult data. This is published in the TestData IS server, from where the Calibration Controller can pick it up and for example modify the configuration parameters for the next run. Summaries of TestResult data are published to the ISTestResult IS server. The summaries typically contain the mean and the RMS values of the fit parameters for each chip.

When a strip or chip parameter is outside some range, it is considered a defect. Table 3.5 lists the possible defects, the test which can determine it, what the limits are and what action is taken. Analysis Service makes a list of these defects, which is used by the Calibration Controller to mask off channels or by experts to take further action. For some minor defects, no action is taken and the strip remains unmasked. But the defect may need to be accounted for in the offline analysis.

ArchivingService

The Archiving Service monitors several of the IS servers and writes out those objects required for permanent storage – “persistifies” them. They are stored as text files, in a custom data format, and compressed to save disc space. These files can be read back into SctRodDaq or by a stand-alone tool called DataDisplayer for further analysis.

3.5 Data Acquisition Runs

The SctRodDaq software is started and initialized in the “Boot” step of the TDAQ run-control. All the applications defined in the OKS database for the partition are started. The SCT has two run modes, namely the Calibration and the Physics mode. The run modes can only be changed before the “Configure” step of TDAQ run control. The user chooses which mode to run with from the IGUI before configuring TDAQ.

In the “Configure” step of TDAQ run-control, the ConfigService reads the configuration files and sends the appropriate data to the SctApis. The configuration of hardware and software data-flow is slightly different in these two modes. In a calibration run, only the SCT is defined as a TDAQ resource. In physics mode, the configuration also contains the source of physics triggers, e.g. ATLAS CTP, or just one LTP for physics test running, or cosmics running. We will now describe the two modes in more detail.

3.5.1 Calibration Runs

For a calibration run, at the “Configure” step, the RODs capture and histogram all the data but do not send any data up the S-Link. In the “Start” step of TDAQ run-control, the trigger control is handed to the SctGUI. The user can start various scans and tests from the SctGUI. Fig. 3-7 summarises what happens in a typical scan. Fig. 3-6 illustrates in more detail the histogram production during a calibration run.

When SctGUI invokes a test, the Calibration Controller assumes control. It downloads any chip parameters that need modifying to the SctApis and hence to the modules via the RODs. Then it starts the scan: SctApi assumes control and first configures the RODs to do a pre-scan hardware check, then configures the histograms in the RODs and finally letting triggers flow and starts histogramming. The RODs send any configuration changes needed for each step of the scan, e.g. the next threshold, to the modules. They then send the appropriate number of triggers, and histogram the resulting data from the modules.

When histogramming is finished, SctApi reads the histograms from the RODs

and publishes them to IS, and finally reports to Calibration Controller that the run is finished. The Calibration Controller gains control. The analysis chain then starts with the Fitting Service fitting the histograms and publishing the results. The Analysis Service extracts the important fit results, giving new configuration parameters and making a list of faulty strips and chips. The new configuration parameters are stored, ready for the next scan. Analysis Service informs the Calibration Controller when the analysis is complete.

Finally, the scan and analysis is complete and the Calibration Controller returns the control to the SctGUI. The Archiving Service stores the results on disc for future reference.

3.5.2 Physics Mode Runs

For a physics run, at the “Configure” step, the RODs are configured to send any data up the S-Link, while maintaining the ability to capture and histogram the data on the RODs as well. Some resource, such as the CTP or an LTP is configured to be the source of the trigger chain while the rest is configured to receive the TTC from upstream. In the “Start” step of TDAQ run-control, the triggers are started.

In this mode, the SctGUI has no control over the triggers but serves mainly as a monitoring tool. It can, however, still start histogramming tasks on the RODs from the events captured while running physics triggers. Since the ROD no longer controls the triggers, the number of triggers sent and the number of events received constantly has to be monitored and the histogram binning changed only when the criteria specified in the SctApi is fulfilled. The data flow once the histogram is produced is the same as in the calibration mode. Such ROD monitoring histograms could promptly identify problematic modules and report them due to the high rate of L1A. Compared to higher level monitoring tools, which histogram events after L2, such ROD histograms are quicker in identifying problems. There is on-going work towards implementing such monitoring ROD histograms in physics mode.

There can be no updates to the module configurations during one physics run. However, strips could be masked off dynamically in the ROD. This could be important

to avoid noisy strips clogging up the bandwidth for sending data higher up the chain.

In the very special case of physics running to time-in the SCT to the rest of ATLAS, it is possible to scan through the BOC RxDelay. The ROD can histogram the occupancy versus RxDelay, using a bean-shell Java script. This is further discussed in Section 5.2.

3.5.3 ATLAS Event Building and Event Format

The ROD assembles the data into ROD event fragments and passes these via the BOC and an S-link to the ROBs. The ROBs are PCI modules housed in a VME-board PC known as a Readout System (ROS) PC. The PC with its ROBs is known as a ROS.

When Level-2 accepts an event, it requests the ROS to send its ROB data to the Event Builder (EB). The ROS assembles the ROD fragments into ROB fragments, and the ROB fragments into ROS fragments and passes this to the EB. For the SCT, one ROD fragment goes into one ROS fragment while in other systems, a ROB fragment can contain many ROD fragments.

The EB uses the Sub-Farm Input to assemble the ROS fragments into sub-system fragments. For example, the SCT data from the event is collected into four sub-system fragments (Barrel side A, Barrel side C, Endcap A and Endcap C) and these sub-system fragments are assembled into the event.

Thus there are five organizational layers in a finally-built ATLAS event, [27], which are known as Event, Sub-detector, ROS, ROB and ROD fragments. Each type of fragment is identified by the first word of its header: *0xAA1234AA*, *0xBB1234BB*, *0xCC1234CC*, *0xDD1234DD* and *0xEE1234EE* respectively. Once assembled, the event is passed to the Event Filter (EF), the final trigger level in ATLAS. If accepted by the EF, the event is sent to permanent storage.

ATLAS TDAQ provides packages with configurable components for event building. These can process the data throughout the chain from the RODs to ROBs to Level-1 and Level-2 and the EB and EF, to final file writing. As already mentioned in Section 3.4.1, the SCT uses a custom Rod Crate DAQ package, but from then on the

SCT data follows the same path and processing as all other detectors.

The SCT ROSs run an application built on TDAQ readout packages. The TDAQ packages provide example configuration for simple tasks. SCT ROS readout applications have been written based on these. The packages are configured via OKS.

The SCT ROS assembles the data it receives from the modules into a ROD fragment. This starts with the ROD fragment header given in Table 3.4. The header contains a source identifier; these are defined ATLAS-wide and are given in Table A.1. Headers also contain information about the run-type and run sequence, as provided by the run-control; and the Level-1 trigger type, L1ID and BCID as broadcast by the TTC system to the RODs via the TIM.

It is worth noting here that the SCT front-ends have 8 bits of BC ID information, while ATLAS counts up to 12 bits of BC ID information. The SCT RODs compare only the last 8 bits of BC ID and the last 4 bits of L1 ID to check synchronization of the front-ends.

Word	Contents	Comment
0	0xB0F00000 + UCTRL	Beginning of fragment marker
1	0xEE1234EE	Start of header
2	0x9	Header size
3	0x30000000	Format Version Number (Ver 3.0)
4	0x002XMMM	SCT Source Identifier: M = ROD ID, X = LS Nibble of Sub-detector ID
5	0xTTSSSSSS	Run Type(T) and Sequence within Run Type (S): T = 0x00 if Physics, 0x01 if Calibration 0x02 if Cosmics and 0x0f if Test
6	0xEELLLLLL	Extended Level 1 ID: E = ECR ID, L = L1ID
7	0x00000BBB	Bunch crossing ID
8	0x000000AA	ATLAS Level 1 trigger type
9	0x00RR000T	Detector event type R = ROD or T = TIM

Table 3.4: Contents of the ROD Event header

3.6 SCT Calibration

Due to the binary nature of the SCT readout, the calibration software is of critical importance to a working and efficient SCT. The analog behaviour of the SCT modules

can only be extracted by scanning through chip thresholds. But extracting this analog information requires optimal and stable optical communication with the modules. The optimization of this optical communication is the first step of a SCT test in the calibration mode. After the optical communication is established by making sure that all modules are returning events, digital tests are used to check that the ABCD chips are performing to specifications. Only after this, can one go ahead with the analog tests, which really characterize the performance of the silicon modules.

As well as the measurement of variables that quantify the performance of the modules, a good understanding of the defects is important as they can influence tracking performance. For example, the number of dead channels allowed per module is 15 channels to achieve the 99% efficiency requirement while the number of consecutive bad channels allowed is 7 – more would reduce the sensitivity to multiple hits.

The testing of the SCT modules is not necessarily straight forward. Even if a module returns events when probed, this does not mean that the optical communication is stable enough to read out complicated and long data streams back from that module. Often, the DAQ shifter is required to optimize the optical communication once more after a test has failed.

As presently coded, the histogram control task running on the ROD's master DSP will abort the scan if any of the ROD's four slave DSPs receives a single error event from any module. The histogram task running on that slave will decode the error event and place diagnostic information into its text buffer, from where the information can be read out and written to a log file. After identifying which modules are causing errors, the modules need to be debugged. Quite often, the problem is with the optical-communications, which may need to be re-optimized for longer data packets. If the optical communication for one link is particularly unstable, the data for that link can be readout using the RX redundancy from the other side at the cost of losing the data from the master chip. Rarely, one single chip or channel might cause event formatting errors and needs to be found by doing a "binary-search" through different module configurations. In the future, it is foreseen that while running calibrations between physics runs, the histogramming task will be coded to be more error tolerant

to let scans continue to completion.

Another complication in the calibration procedure during warm module testing is the power trips due to higher temperatures at lower thresholds. As the threshold is lowered, the occupancy per chip increases and the ABCD chip has more information to process, resulting in higher currents to be drawn and temperature of the modules to rise. If there is insufficient headroom between the mean operating temperature of the modules and the software trip limit of 33°C then modules with increased temperature differences between the two sides (“ ΔT ”) are likely to trip off as thresholds are reduced.

Under such circumstances the ROD will see a time-out error as there will be no reply from the module to a trigger. This will cause the the ROD to abort the scan, making the completion of a full calibration difficult. Also, at lower thresholds, the packets sent back from the modules grow in size, as the data packets are longer, and there is a higher chance that a small instability in the optical communication causes data errors in the stream, again causing the ROD histogramming task to abort.

Scan histograms will read out up to the point at which the scan aborted. As the optimum running 1 fC threshold point and the 2 fC point where the gain is most linear are better measures of the calibration accuracy than the lower threshold points, the scans are done starting at higher thresholds and then going to lower thresholds to minimize the chances of it aborting before it reaches the determining 1 fC point. With this method, most of the scan is done without warming up the detector by running at low thresholds.

3.6.1 SCT Scans and Tests

All scans and histograms are 2D projections of 3D histograms and are of a common form unless stated otherwise. For each module, there are two histograms, one for each link, but in this chapter, only one histogram is shown as an example for each test for brevity. For BOC parameter scans, the *x-axis* is time (in bunch-crossings) while for “module scans”, it is channel number. The y-axis is the parameter being varied through the scan. The colour in the histogram is a projection of the 3rd axis,

which is the number of triggers which returned ‘1’ for that channel and parameter. The colour scale on the right side generically is normalized so that full-occupancy is ‘1’ as default. Quite often the chip structure of the module is apparent in the scans as a repetitive pattern of 128 channels. For example, Fig. 3-15 shows a “module scan” of the masking register with the channel along the *x-axis*, and the number of channels masked along the *y-axis*. In this case, the normalized contents are either 1 (unmasked) or 0 (masked). We now will go through the tests illustrating results with example plots.

3.6.2 ModuleProbe

ModuleProbe sends a trigger and then traps the raw data that is returned. It is a check that all modules are returning correct readout patterns. Table 3.6 is a list of patterns that the module could be returning and what actions that the DAQ shifter might need to take to get all modules returning events.

After the modules are powered on, they go into CLK/2 mode, which is a mode where the modules return the input clock divided by 2 in frequency. If they are seeing light from the BOC, then they return a clock signal at 20 MHz. If a module is not in CLK/2 at this stage, the I_{p-i-n} should be checked and if low, the TxCurrent can be increased. Another possibility is that the Tx to Rx fibre mapping is wrong or the fibre connections are not mated optimally. Yet another possibility is that the RxThreshold settings are too high while the VCSEL has a low output level.

When the modules receive their configuration, they should start returning events. If not, checks are performed on the DCS side to see what state the module is in. For example, if the select line is set high on either the DAQ or the power-supply configuration, but not both, the module will remain in CLK/2 mode. After this check, if the module is still not configured then a HardReset can be issued to the module to try to reset the module and the configuration sent again. If the module is not configured after a couple of tries, then Tx-settings may need to be optimized by a TxCurrent test and fibre connections may also need to be checked.

If after all these checks, the module is not responding, it is possible that the TTC

Defect	Severity	Criteria	From test	Action if any
DEAD	UNUSABLE	Output always < 1%	Pipeline/NO	now updated
STUCKON	UNUSABLE	Output always > 98% occupancy	Pipeline	now updated
UNDER	DODGY	Occupancy always < 95%	NMask	masked
OVER	DODGY	Occupancy > 100%	NMask	masked
BADFIT	DODGY	Fit not good, chi2/dof > 5000	FittingServ	online
LO_GAIN	UNUSABLE	Gain < 0.75* chip average	NPT	masked
HLGAIN	UNUSABLE	Gain > 1.25* chip average	NPT	masked
LO_OFFSET	UNUSABLE	Offset < -100	NPT	masked
HLOFFSET	UNUSABLE	Offset > 120	NPT	masked
UNBONDED	SERIOUS	Noise < 800	NPT	Digi/Reco
PARTBONDED	SERIOUS	Noise < 1100	NPT	Digi/Reco
NOISY	SERIOUS	Noise > 1.15* chip average	NPT	Digi/Reco
MEAN_ERROR	SERIOUS	Error in the mean	FittingServ	online
SIG_ERROR	SERIOUS	Error in the sigma	FittingServ	online
NOINIT	SERIOUS	Couldn't initialize fit	FittingServ	online
NO_HI	SERIOUS	High noise occupancy of > $5 * 10^{-4}$	NoiseOcc	Digi
DEAD_CELL	SERIOUS	Dead cell in the pipeline	Pipeline	Digi/Reco
STUCK_CELL	SERIOUS	Stuck cell in the pipeline	Pipeline	ROD/Digi/Reco
TR_RANGE	UNUSABLE	Unusual chip trim step size	TrimRange	masked
TR_STEP	DODGY	Channel trim step size not 4	TrimRange	online
TR_OFFSET	DODGY	Channel trim range offset not 4	TrimRange	online
TR_NOTRIM	UNUSABLE	Untrimmable channel	TrimRange	masked
TOKEN	DODGY	Direct token fails	FBT	
RTOKEN	DODGY	Bypass token fails	FBT	
TW_HI	DODGY	Time walk too big, > 16.	TW/RX/TC	online
TW_LO	DODGY	Time walk too small, < 5.	TW/RX/TC	online
SD_LO	SERIOUS	Strobe delay rise < 0 or fall < 28	StrobeDelay	online
SD_HI	SERIOUS	Strobe delay rise > 35 or fall > 63	StrobeDelay	online
VLO_GAIN	UNUSABLE	Gain < 0.3*chip average	NPT	masked
V_NOISY	UNUSABLE	Noise is > 1.25*chip average	NPT	masked
NOISE_SLOPE	DODGY	Noise slope/chan > 1.0	NPT	online
OFFSET_SLOPE	DODGY	Offset Slope/chan > 0.07	NPT	online
GAIN_SLOPE	DODGY	Gain slope/chan > 0.04	NPT	online
BAD_OPE	DODGY	OPE/binomial variance > 2.0	NoiseOcc	
DOUBTR_HI	DODGY	High double trigger noise > 5.0	DoubleTrigger	
L1_COUNTER	SERIOUS	Level 1 counter defect	ChipCounter	Replaced
BC_COUNTER	SERIOUS	Bunch crossing counter defect	ChipCounter	Replaced

Table 3.5: List of all defects that can be issued by SctRodDaq with the meaning and action taken.

Value	Meaning	Action
0	stuck OFF (no response)	check DCS and fibre mapping and RxThreshold setting
1	stuck ON	hard reset
2	clk/2: unconfigured	send configuration or check that select line setting is same on DAQ and power-supply mapping
4	clk/4	power cycle with TTC enabled
A	Valid ABCD error code (1, 2 or 4)	issue hard reset
a	Another error code (corrupt data)	try RxThreshold scan
C	Module in Config Readback mode	reconfigure module
E	Returning events	No action

Table 3.6: The values and their meaning of the results returned by SctApi after a probescan and the actions that the DAQ shifter might need to take to get all the modules returning events.

link is dead, either due to a fibre break or an electrical connection problem. The list of known TTC links that do not work are given in Table. 3.7. Links which pass the diode test, but see no *p-i-n* current, have their data link fibres broken. Those for which there is an electrical fault due to a break in the LMT, which is now inaccessible, the *p-i-n* diode can not be biased and so the link does not work. All these modules listed here can be operated with redundant TTC, getting their clock and command from their neighbors. With this setting, they all returns events.

Location	Serial Number	Problem
Barrel 4 LMT38 Z-6	20220170200136	No <i>p-i-n</i> current. Passes diode test.
Barrel 4 LMT27 Z+1	20220170200378	No <i>p-i-n</i> current. Passes diode test.
Barrel 4 LMT15 Z+5	20220170200695	Electrical fault from cable testing.
Barrel 5 LMT31 z+1	20220170200079	No <i>p-i-n</i> current. Passes diode test.
Barrel 5 LMT31 z+6	20220170200179	No <i>p-i-n</i> current. Passes diode test.
Barrel 6 LMT44 z+1	20220170200575	Electrical fault from cable testing.

Table 3.7: The table of known TTC links which do not work. All these modules when operated in redundant TTC mode return valid data.

After a probe, generally a HardReset is issued by the user to check that all modules can be reset to the state of initial power-on, which is CLK/2. If HardReset works as well then modules are configured and checked one last time before starting the test sequence. There are a handful of modules where the toggle of the reset line does not HardReset those modules. In such rare cases, a module may need to be powered on and off to achieve the same effect.

A module probe is run before the start of tests (other than BOC scans) by the SctApi, and if a module is not returning events under this “pre-scan hardware check” the scan is aborted.

3.6.3 TxCurrentTest

The TxCurrent scans through the TxCurrent register on the BOC and histograms the reply back from the module. For each TxCurrent bin, a hard-reset is sent to the modules to put them in clk/2 and then the module is configured to send back the contents of the configuration register.




Figure 3-8: A TxCurrent scan test result for one Rx link. The optimum setting is indicated by the thick line at a DAC setting of 166. Below 110 and above 220, more and more bit errors occur.

As shown in Fig. 3-8, at low TxCurrent, the module does not even see the clock and does not reply with any light. At slightly higher TxCurrent, the module will see the clock, and so will return $\text{clk}/2$, seen as a 50% occupancy or green band in the histogram. At higher TxCurrent, the module will configure and will reply with the contents of the configuration register stably for the width of the working regime, indicated here by the analysis by the red lines. The optimal TxCurrent setting, indicated by the blue line, is defined to be the middle of the working regime. At very high TxCurrent settings, the DORIC4A/PIN system saturates and loses clock and command and the module either returns no data or erroneous data.

3.6.4 RxDelayTest

RxDelay delays the local copy of the 40 MHz bunch-crossing clock to sample the incoming data link at the optimal phase.

The RxDelay test first sets all the modules in $\text{clk}/2$ by issuing a hard-reset and then scans through the RxDelay register at the BOC, from 0 to 24 ns, while setting up the BOC to sample at half frequency. This means the transition will be from all 1's to all 0's. The analysis sets the optimal RxDelay to be equidistant from the clock edges.

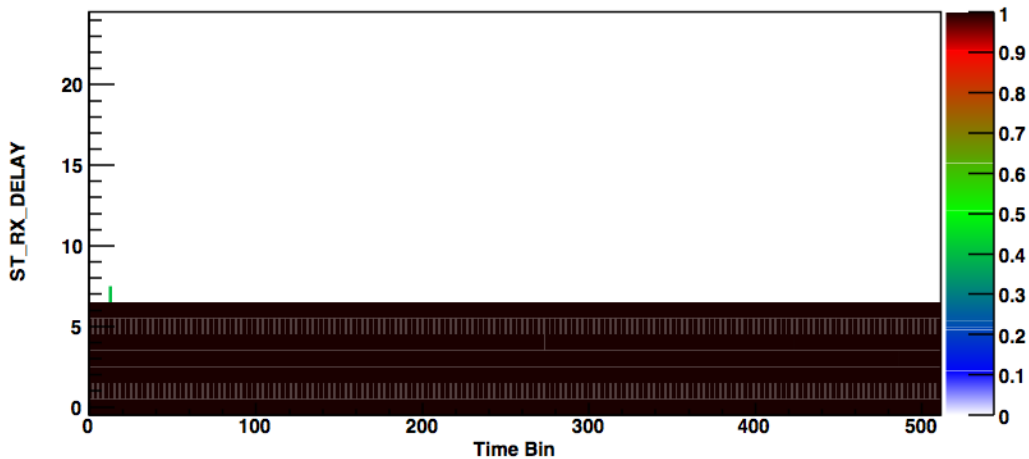


Figure 3-9: A RxDelay scan test result for one Rx link. At small delays (up to 6 ns) the clock is high and a bit is always seen. After that, no bits are seen. The delay is set to 18 ns, equidistant from the clock edges, for this case.

3.6.5 RxThresholdTest

The RxThreshold test measures the amount of light being received back from the module at the BOC by scanning through the threshold setting at the BOC. It sets the RxThreshold to its optimal value at the end of the analysis to allow for a stable communication with the module. This test asks each chip to read back the contents of the configuration register so that a constant data packet is returned. The headers are deliberately suppressed since the BC and L1 counters vary from event to event, which would complicate the analysis. The calibration controller sets up the test to return a slightly different pattern on odd and even datalinks, to give a better indication of optical crosstalk, which may occur if the infineon connectors at the BOC are not seated correctly. The chips on link 0 are configured to have the calmode, compression and trim range set to 2, while chips on the other side have these registers set to 1.

Fig. 3-11 shows a typical histogram of this scan. The *x-axis* is time in bunch-crossings while the *y-axis* is the RxThreshold setting. 10 triggers are sent to the module as the light output can vary for some VCSELs. The analysis calculates the minimum and the maximum threshold values for which all data-packets are received correctly, and sets the optimal threshold to the 70% point between the minimum and




Figure 3-10: An oscilloscope picture of the optical signal from a VCSEL using a fast optical probe with an optical attenuator to ensure that the signal is well below the saturation level for the probe, [41]. The light level drops very quickly to zero, implying that the rise in noise in Figs. 3-11 and 3-12 is due to the BOC.

the maximum. The light output of the VCSELs quite often increases with time; also the slow fall-off of the *p-i-n* diode in the BOC produces a higher background noise with longer data packets. The light output from the fibre does not have this slow-fall off as shown in Fig. 3-10 where the light level drops off very quickly to zero. It has been found that by setting the threshold at 70% rather than the more obvious choice of 50%, the optical communication is less prone to errors.

If the minimum working RxThreshold value is above 100, it is advisable to turn down the VCSEL voltage to reduce the likelihood of errors with longer data packets. If the maximum RxThreshold is below 150, it is advisable to turn up the VCSEL voltage a little bit to increase the error-free zone of the optical communication. After any applied changes to the VCSEL voltage, a new RxThreshold scan is needed to set the RxThreshold values and to check that no other changes are necessary.

Rarely, a link will return no light. In such a case, either the fibre is broken or the VCSEL is dead. A suspected cause of damage to the VCSELs is electrostatic discharge, [132]. Extra care during construction and macro-assembly has been

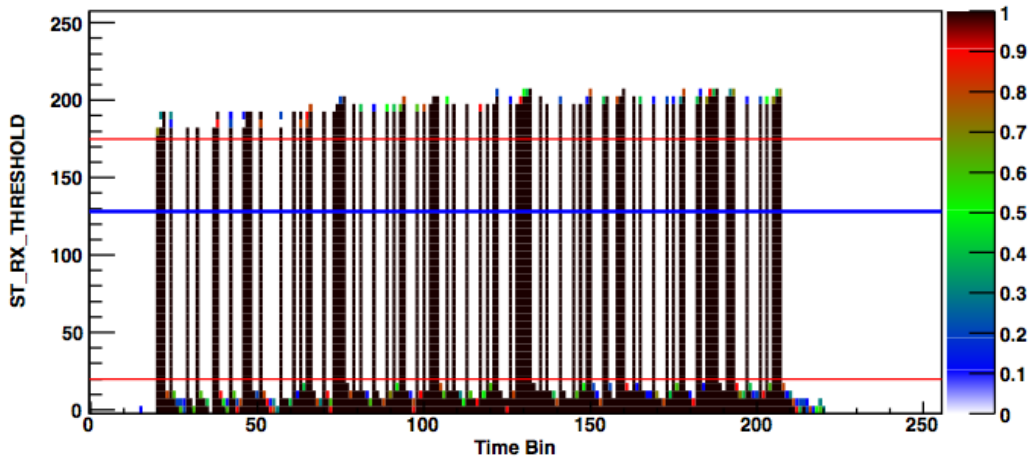


Figure 3-11: An RxThreshold scan test result for one Rxlink. For each trigger, the module sends the same reply. So each time bin should either be black (always 1) or white (always 0). At high thresholds, the 1's are not seen by the BOC; at low thresholds, noise converts 0's to 1's. The minimum and the maximum RxThreshold value for which the opto-communication is stable is shown with horizontal red lines. The best RxThreshold value, which is set, is shown with a horizontal blue line.




Figure 3-12: An RxThreshold scan test result for a link with a slow turn-on VCSEL. There is only one value of RxThreshold for which the communication might work for this particular module, serial number 20220380200139 at location Barrel3 LMT28 z+6. This link has a high error rate and this side of the module must be read out from the other link.

practiced to avoid such electro-static discharge on the tools, cables, robots and shift-takers. In the rare case that a link will return no light, 11 out of the 12 chips of the module can be read out through a single link. A list of all known optical Rx links which do not work are given in Table 3.8.

Some VCSELs have a slow turn on time which is seen as increasing light output in time. As seen in Fig. 3-12, there is only one value of RxThreshold for which the data would be decoded correctly. However, with longer data packets, communication will be unstable. Sometimes increasing the VCSEL voltage helps to increase the working regime of the opto-communications. Occasionally, the slow fall-off of the *p-i-n* diode signal in the BOC forces the RxThreshold values to be very high. A way to fix this is to set the VCSEL voltage lower. Quite often, a stable communication can not be established for these slow turn-on VCSELs, and the module has to be read out from the other side's link.

3.6.6 2D BOC Scan

The 2-dimensional BOC scan is not “a standard test” which can be executed directly from the SctGui yet but is soon foreseen to be one. It can be run from the Bean-Shell scripting window which is available from the SctGui after putting all the modules in $\text{clk}/2$ by hard-resetting them. This test scans through RxDelay and RxThreshold for each link and histograms the number of events returned by that link.

It finds the optimal values for a link of both the RxDelay and the RxThreshold by finding the worst RxDelay value and setting the delay to be 9 ns away from the falling edge and the RxThreshold value to be 2/3 s of the way in the working regime.

It is not sensitive to “slow turn-on” VCSELs and so may set the RxThreshold lower than its ideal for such links, but it is essential in de-coupling the dependency of RxThreshold on the RxDelay. Running an RxThresholdTest after a 2D-BOC scan to fine tune the RxThreshold, has been demonstrated to be the best way to adjust optical Rx settings.




Figure 3-13: A 2-dimensional BOC scan for one Rx link. The RxDelay and Rx-Threshold settings are dependent on one another and this test sets the optimal value of both at the same time.

3.6.7 MarkSpaceRatioTest

The DORIC4A recovers the 40MHz clock from the BPM signal. If the BOC duty cycle is far from 50%, the recovered clock has significant time jitter, [41]. This test finds the correct setting of the Mark to Space ratio (MSR) in the BPM chip for 50% duty cycle of the BOC output.

It is difficult to find this point because the Rx signal in a BOC also has effects from the DORIC4A and RxThreshold settings. This difficulty has been overcome by measuring the duty cycle of the clock/2 from the module as a function of MSR settings. This is measured for two opposite phases of clock/2. When the duty cycle from the BOC is correct (50%), the duty cycles from the module for the two phases are equal.

The module is set in clock/2 and an RxDelay scan performed for each MSR register setting. The duty cycle is calculated and plotted as in Fig. 3-14. Then the phase of the clock is flipped by sending a “soft reset” to the module, and the process is repeated. The two sets of data have opposite slopes. Where they meet indicates the MSR register setting which gives 50% BOC duty cycle.

This test is very time consuming and so has only been applied to channels which




Figure 3-14: MarkSpaceRatio test result for one Rx link. Each point is the result of a scan at the given MSR register value.

still have communication problems after optimizing TxCurrent, RxThreshold and RxDelay. Otherwise, a default value of $0x14$ is assumed. A fast 2-D histogrammer, similar to the fast-RxThreshold-RxDelay scan, that will run on the ROD directly is being developed.

3.6.8 NMaskTest

The NMask test is the first digital test that establishes that optical communications are working and checks functionality of the ABCD chip. It checks the mask register for each channel of the module and reports a NMASK defect if the channel can not be masked off.

For each bin in this test, a different masking register pattern is set. The DAQ sends a pulse output command followed by an L1 trigger which causes all the channels




Figure 3-15: An NMask Test result for one side of the module. When zero channels are masked off, all channels fire (bottom of histogram). When all 128 channels are masked off, no channels fire (top of the histogram).

to record a hit. As the masking register is applied to the data, data is returned only for those channels which are not masked. The DAQ decrements the mask registers of the chips, producing the triangular form for each chip as seen in Fig. 3-15, which is easy to analyze.

Channels with very high noise should be masked off during the running of the experiment since they complicate track finding. Such high noise channels also have a detrimental effect on the bandwidth. If a channel needs to be masked due to high noise but also has a NMASK defect, it will have to be masked offline.

This test validates that the DAQ can read out various-length signals from the modules and generally fails if the optical communication parameters are not optimal. The DAQ user has to re-optimize these parameters if an NMASK scan fails for a module.

3.6.9 PipelineTest

This is a test of the pipeline of the cells of the ABCD3T in order to identify cells which may permanently output 0 (dead) or 1 (stuck). The pipeline for a channel is 132 deep, but ordered in 12 blocks of 11 cells deep. We test for a dead or stuck cell in each block. If for a given channel, all 12 blocks are found to contain dead or

stuck cells, then a defect type of DEAD or STUCK is recorded. If for any channel, the number of blocks containing dead cells or stuck cells is greater than zero but less than 12, each individual defect is noted as either type DEADCELL or STUCKCELL and a mask of defective cells is listed in the analysis summary. In this test, the DAQ first sends a soft reset, which resets the BC counter. Which set of pipeline cells is used is determined from the BC counter at any time. The soft reset is followed by a variable delay before sending the pulse output command and a Level-1 trigger. Scanning over the delay tests each of the cells in the pipeline buffer.

Since the pipeline is downstream of the masking registers in the data flow, these defects can not be masked off and in the case of STUCKCELL defects contribute to the noise occupancy of a module. It is necessary to mask defective cells off either on the ROD or offline in Sct_Digitization because one STUCKCELL causes an occupancy of $1/12 = 8\%$ on that channel. A masking of such channels on the ROD is foreseen in the long-term to reduce the occupancy transferred by the S-link and seen by the Level-2 trigger. However, this is not trivial and so for now, stuckcell information is written into the COOL database for masking offline.

3.6.10 FullByPassTest

In this test, the module is programmed to each possible configuration of routing token and data between the chips while the readout chain is checked, as seen in Figure 3-16. Irradiation tests have shown that different chips can have different susceptibilities to radiation damage effect with regard to the token passing. When a chip is bypassed, it returns no events and if it is not bypassed, it should return full occupancy. Here, first, the token passing on one side is practiced for the patterns up to 20 on the *y-axis* and for those above 20, the master chip on one side is bypassed and combinations up to 7 chips are read out from one of the links, [71]. Two defects, namely TOKEN and RTOKEN are issued if respectively direct and bypass token passing fails.




Figure 3-16: A FullBypass Test result for one side of the module. All combinations of chip by-passing on one link are exercised in this test.

3.6.11 Counter Error Test

The ABCD3T ASIC chip has 4 bits of Level1 trigger counter and 8 bits of Bunch Crossing Counter. The purpose of this test is to check that these bits in the data headers read from the master chips on a module are correct. If a counter was faulty, the module was replaced during macro-assembly, as it would not be possible to verify that the module is synchronized properly with the rest of the experiment during physics running. The test sends 128 consecutive triggers to the module and reads them back. Only the first 20 bits of the reply are plotted. The analysis compares the counters from the two links of a module, as seen in Fig. 3-17. It also checks to make sure that the counters return non-zero counts. Two defects, namely L1_COUNTER and/or BC_COUNTER and the erroneous header bit are reported. One module on Barrel 6 was replaced at the macro-assembly stage because one of its master chips had a faulty BC counter.

3.6.12 Strobe Delay Test

The strobe delay varies the phase of the charge injection relative to the L1A command. An optimal setting of the strobe delay for each chip is important for the accuracy of the threshold calibration. A 4 fC charge is injected with the threshold set to 2 fC:




Figure 3-17: A Counter Error Test result for both Rx links. The headers from each link of a module are compared to identify an error in the L1 or the BC counters of the master chips. The significance of the bit in each time bin is indicated between the two plots: *n* is the L1 counter and *b* is the BC counter. An error is highlighted in the BC counter.




Figure 3-18: A Strobe Delay Test result for one link. At short delays, the signal has not arrived and no hits are registered. At long delays, the signal has decayed below threshold. In between, the efficiency reaches 100%. Fig. 3-19 shows a typical fit for a chip.

hence for the correct timing there should be 100% hit efficiency. Fig. 3-18 shows the results of a typical scan. Ideally the value is chosen such that the peak of the shaper response is sampled.

When looked edge on for a single chip, the strobe delay has a smeared top-hat shape, as shown in Fig. 3-19. The strobe delay is set to a point 40% of the way from the mid-rise to the mid-fall of the shape. This point was chosen to minimize the effects of “Large Gain Spread (LGS) cold” chips.

If a module has been trimmed and calibrated at some temperature, at lower temperatures it may be affected by LGS cold resulting in threshold variations between channels. A possible cause is extra leakage current in the shaper DAC making the actual shaper current higher than expected from the DAC settings, [78]. This moves the shaper operating point close to the edge of the linear range (as the linear range becomes limited and dependent on the actual channel offset), affecting responses to large input charges, [50]. Although the gain around the operating point of 1 fC is still uniform, fits to the response curve are distorted and an apparent large gain spread is observed at 2 fC. This situation can generally be remedied by reducing the shaper current to avoid saturation of the pre-amplifier, [42]. In the rare case that it can not be, setting the strobe delay at the 40% mark between the rise and the fall of the strobe delay curve instead of closer to the rising edge has been shown to reduce




Figure 3-19: A Strobe Delay data and fit result for the average of one chip. The optimum delay setting is indicated at the 40% position.

the effect of timing spread and hence the large gain spread that is observed at lower temperatures. This 40% point is now used as the default for all modules.

3.7 Analog Tests

Analog tests measure the performance of the SCT modules and also provide the configuration for the critical parameters of the calibration of the SCT front-ends. Here we will first discuss the threshold scan which forms the basis of most analog tests and then discuss occupancy per event histograms which actually can be produced for any kind of scan, but are particularly useful for analog scans. After that, we will describe the particular uses of these two methods in measuring responses under different conditions from the detector.

3.7.1 Threshold scans

A threshold scan forms the basis of all analog tests, with or without charge injection. The occupancy versus threshold data (measured in mV) for each channel is fit with

a complementary error function, also known as an S-curve.

The threshold at which occupancy is 50% corresponds to the median of the injected or noise charge distribution, known as Vt_{50} . The variance of the distribution is a measure of the output noise (in mV). For a threshold scan, the fitting service reports the Vt_{50} and the width for each channel to the analysis server, which then combines one threshold scan with others to extract other parameters.

Figure 3-20 shows threshold scans at 1.5 fC for two chips. The responses from all the channels in each chip have been fit with the complementary error function. The difference between these plots is a result of their different trim ranges. The chip with the trim range of 0 has a smaller spread in Vt_{50} than the chip with trim range 1. There are four trim ranges 0 to 3, increasing in range and coarseness. A chip with high average gain tends to have a higher spread and so needs a higher trim range. Trim range will be discussed more in detail in the next sections.

3.7.2 Occupancy per Event


The occupancy per event (OPE) is a histogram type which can be requested for any type of scan, but is particularly useful for analog tests. When the CalibrationController requests this kind of a histogram, the ROD will fill a histogram with the number of hit strips for each event for every chip in addition to the normal histogram and return it at the end of the scan. As the OPE request is also included in the test description object, the fitting and the analysis servers are aware of it and the analysis of such a histogram can be attached to the TestResult easily by the inclusion of the “OPETool” analysis tool in the analysis algorithm of the test.

OPE histograms provide information about the structure of hit occupancies for events. If hits in strips are uncorrelated, then the number of strips hit in an event will have a binomial distribution. But if there is correlated noise, then events with a large number of hits will occur more often than the binomial expectation. This is the basis of the OPE analysis which we will discuss for a NoiseOccupancy scan where it is most useful.


Here, in Fig. 3-21, a OPE histogram for one chip during an NMask scan is shown.

Location	Serial number	Problem	Discovered
Barrel 3			
LMT28 Z+6	20220380200139	link1 VCSEL - the light output has a slow risetime	Oxford
LMT24 Z-6	20220330200131	link1 dead - standard module in place	CERN SR1
LMT29 Z-3	20220380200111	link1 dead - standard module in place	Oxford
LMT16 Z-3	20220330200301	broken VCSEL bias line	CERN pit
LMT22 Z+6	20220040200322	link0 dead	CERN pit
LMT25 Z-4	20220330200090	broken VCSEL bias line	CERN pit
Barrel 4			
LMT02 Z+4	20220170200653	link0 dead - modified module in place	Oxford
LMT29 Z+1	20220170200113	link0 dead - modified module in place	Oxford
LMT35 Z-6	20220170200941	link1 flakey - modified module in place (precautionary)	Oxford
LMT24 Z-2	20220330200461	link0 dead - standard module in place	CERN pit
LMT29 Z+6	20220330200559	link1 dead - standard module in place	CERN pit
Barrel 5			
LMT19 z-6	20220330200701	link0 dead - modified module in place	Oxford
LMT21 z-6	20220330200637	link0 dead - modified module in place	Oxford
LMT40 z+6	20220170200183	link1 dead - modified module in place	Oxford
LMT09 Z+3	20220170200609	link0 low output: OK with RX threshold of 50	Oxford
LMT42 Z+3	20220170200107	broken VCSEL bias line	CERN pit
Barrel 6			
LMT08 z-2	20220330200606	link1 dead - modified module in place	Oxford
LMT09 z+2	20220330200209	link0 dead - modified module in place	Oxford
LMT17 z+5	20220330200505	link0 dead - modified module in place	Oxford
LMT19 z+3	20220170200932	link1 dead - standard module in place	Oxford
LMT28 z+1	20220330200117	link0 dead - modified module in place	Oxford
LMT30 z-5	20220330200693	link1 dead - modified module in place	Oxford
LMT39 z-6	20220380200010	link0 dead - standard module in place	CERN SR1
LMT16 Z-4	20220380200012	link1 dead - standard module in place	CERN pit
LMT24 Z+3	20220170200281	link1 dead - no light	CERN pit
LMT22 Z+6	20220380200225	link0 dead - no light	CERN pit

Table 3.8: List of all the known optical Rx links that do not work on the barrels.



(a) A threshold scan at 1.5 fC injected charge for a chip with trim range 0.



(b) A threshold scan at 1.5 fC injected charge for a chip with trim range 1.

Figure 3-20: The overlaid threshold scan fits for all channels on two chips with different trim ranges. Chips in different trim ranges are evenly distributed through the individual barrels. The data points lie close to the fit and are not shown.




Figure 3-21: Occupancy per event histogram for one chip for an NMask scan. The NMask test result is shown in Fig. 3-15. In the upper left hand corner, when all chips are masked off, there are no hits. With decreasing number of channels being masked, the number of hit channels increases.

The histogram has bins that are four channel counts wide except for the initial special bin only for 0 hits. When all the channels are masked as seen in Fig. 3-15, the OPE histogram records zero hits for all triggers. When the NMask test starts to unmask some channels, the occupancy per event changes accordingly, reporting the number of channels which have hits. The OPE pattern could also be used to analyze the data and report defects, however the original NMask plot is easy to analyze and so no OPE analysis exists for the NMask test.

3.7.3 Three Point Gain Test

This test performs a threshold scan at injected charges of 1.5, 2 and 2.5 fC, which is the most linear region of the response curve. Its purpose is to verify the analog performance of the modules but not to calibrate it. The 3 point gain scan does not update the configuration; the 10 point gain scan, which we shall discuss next, not only verifies the performance more accurately, but also updates the configuration with new parameters.

The parameters important in quantifying the performance are extracted from the fits. The gain of each channel is calculated from a linear fit to Vt_{50} of the three scan points. The output noise at 2 fC is divided by the gain to determine the input noise in fC. The equivalent noise charge (ENC) is reported by converting noise in fC to electrons by multiplying it by the number of electrons in 1 fC, which is 6250. The offset is extracted from the linear fit and its spread is a good indicator of the trimming uniformity.

In module testing using SctDaq, which were all tests up to and including reception testing at Oxford, reported parameters for the 2 fC point. However, a recently discovered bug in SctRodDaq revealed that Vt_{50} and its RMS were reported for 1.5 fC for the 3 point gain test, while all the other parameters were reported again for 2.0 fC. This bug in the TestResult streamer has now been fixed and all parameters will be reported for 2.0 fC starting with the second phase of barrel commissioning in the pit. Extracting the summaries for the 2.0 fC point for tests previously done is now under consideration to allow for direct comparisons in the future.

The 3 point gain test summary reports the offset, gain, input and output noise as well as their RMS values for each chip, as well as the 3 point fit values, the defects and the slopes for offset, gain and noise for each chip. As can be seen from Fig. 3-24, there is a pattern which is repeated for each chip. The noise is higher in the center of the chip. The digital and analogue power source for each chip is delivered at the side of the chip and there are decoupling capacitors near the chip. Therefore the resistance in the power bus causes a voltage drop across the chip and increases the noise in the middle of the chip. The lower side of the module is always slightly noisier than the upper side of the module, attributed to the fact that the lower side is further away from the module ground.

Besides these unavoidable effects also apparent in single module tests, it was noticed during macro-assembly that some modules would display a “slope” in noise, offset or gain which was indicative of pickup. Such features had a strong correlation with either that module’s or its neighbor’s select line settings. While in the design of the module attention was paid to reducing the noise generated by the operation


of the module, in this case, activity correlated with the clock generates a switching current on the power buses, increasing the noise. However, why some modules show a marked effect and why others do not is not understood.

An example offset slope is shown in Fig. 3-22. The module shows an offset slope if its neighbor is getting its own clock and command. However, if the neighbor gets its clock and command from this module instead, the module no longer exhibits such a slope. There are other examples where a module shows some slope feature only if running with its own clock and command and this effect goes away if it is running in redundant clock and command. In general, slope features have been demonstrated to be solvable by changing a combination of the module's and its neighbor's clock and command streams. Table 3.9 lists such known problems and solutions. It is thought that the slopes are due either to pickup from the neighbor module's clock and command or to the effect of the DORIC on that module's power busses. It is interesting to note that such slopes have only been observed on the lower side of the module and often toward the end chip as it is the furthest away from the ground. This observation backs the theory that the slopes are related to the module ground connectivity.


It is worth mentioning that the module on Barrel 4 LMT40 z-5 was replaced at the end of barrel macro-assembly to try to cure the offset slope. The replaced module showed the same features confirming that this is not a problem with the module but rather with the services. Running this module with it receiving redundant TX signal from its neighbor does not fully cure the offset but reduces it. Turning off its neighbor on z-6 cures the problem, but is not an option.

As well as the slope defects, there are several other defects the three point gain test finds. Table 3.5 lists these with the tag NPT as the test name. Since 3 point gain test is relatively quick compared to a 10 point scan, it is often used as a check that the total number of channels that are functional have not changed during macro-assembly and commissioning. During ATLAS physics running, 3 point gain tests can be used to quickly indicate when a recalibration is needed with a 10 point gain scan.

The summary and the defects issued by the 3 point gain scan is the same as a 10



(a) A plot of offset versus channel number if the module's neighbor has select set to 0 whereby it gets its own clock and command.



(b) A plot of offset versus channel number if the module's neighbor has select set to 1 whereby it gets its clock and command from this module.

Figure 3-22: The slopes in offset, dependant on module clock and command. These plots are for module serial number 20220170200008 on Barrel3 LMT30 at z-6 and if its neighbor module on the same row at z-5 with serial number 20220170200009 has its own clock and command, then the module at z-6 develops an offset slope. The offset goes away when the module at z-5 gets its clock and command from the module at z-6.

point gain scan or any other future N point gain scan. Some of the defects are not masked, as can be seen in Table 3.5, as these channels are still considered usable. An example as such are the *noisy* channels, which are 15% to 25% noisier than the rest of the chip but not noisy enough to warrant masking off. Such information can be used in the offline digitization simulation and event reconstruction. Another identified defect that is not masked off is for the *partbonded* channels. On such channels, one bond has failed but the other bond is still intact. The geometry of the module bonding dictates that such a partbonded channel has to be functioning only for the wafer which is under the ABCD chip. Such channels may only be 50% efficient, but their hits contain more accurate position information, as only half the strip length is active.

Location	Serial number	Problem	Solution
Barrel3 LMT24 z+3	20220330200446	Offset slope	Use select=1 on this module
Barrel3 LMT30 z-6	20220170200008	Offset and noise slope	Use select=1 on its neighbor
Barrel4 LMT40 z-5	20220170200095	Offset slope	Use select=1 on its neighbor
Barrel5 LMT36 z+6	20220170200690	Offset and noise slopes on several of its neighbors	Use select=1 on this module

Table 3.9: A list of some known slope problems as identified by an N-point gain scan test and solutions.

3.7.4 Response Curve

This test extends the three point gain test to 10 scan points where the injected charge is varied over a larger range to 0.5, 0.75, 1.0, 1.25, 1.5, 2, 3, 4, 6, and 8 fC points. This test is also known as a 10-Point Gain scan. It gives a precise measurement of gains and offsets with which to update the configuration. This test is crucial for the long term stability and performance of the SCT as with irradiation, the characteristics of the modules are expected to change.

For 10 point gain response curves, the fitted Vt_{50} point was reported at 2 fC for SctDaq and 0.5 fC for SctRodDaq. This bug is now fixed in SctRodDaq in favor of the 2 fC reference point.

The response curve may be fit with different functions and parameters. As of now in SctRodDaq, a quadratic fit is used, as seen in Fig. 3-23. Some of the higher






Figure 3-23: A quadratic fit to the 10 point gain scan for one chip. The pre-amplifier non-linearity at high signals is visible.

charge injection points in the curve may need to be excluded from the fit because it is possible that the threshold DAQ saturates above 5 fC. Saturated threshold DAQ leads to a threshold scan where the occupancy does not drop to zero even at high thresholds, leading to fitting problems. Since this effect is most apparent for the 8 fC point, it is known as the “8 fC effect.” TestResult and summary reports this if any points are excluded from the response curve fit. The fit parameters for the chips are stored in the updated module configurations, so that they will be used to accurately set the discriminator threshold corresponding to the required charge for future tests.

The results from this test are often quoted to quantify the performance of the barrels, for they provide the most accurate measure of the response of the detector for most of its operational range. Unfortunately, it takes a longer time to complete than other tests and hence was not employed during the macro-assembly nor the reception testing at CERN stages for the barrels. However, it was used for a sector of the integrated SCT barrel, before the cosmics testing, to update the configurations. The cosmics sector is described in detail in the next chapter. For our purposes here, it is a sample of 461 modules, for whom all tests passed, selected evenly from all



(a) A plot of gain versus channel number for one module



(b) A plot of input noise in ENC versus channel number for one module

Figure 3-24: The analysis plots for the gain and noise for both sides of one module from a 10 point gain test.




Figure 3-25: The measured V_{t50} values for all chips for different injected charges in the sector plotted against their offsets.

barrels.

It is important to set the threshold accurately to 1 fC for all channels. The trim setting to achieve 1 fC varies due to the spread in offsets and the spread in gains. The former dominates, as can be seen in Fig. 3-25: the offsets have a spread of about 80 mV, while the spread in Vt_{50} due to the gain and the error in the calibration factors at 1 fC is about 10 mV (the width of the 1.0 fC band). The offsets are very well determined in a 10-point response curve measurement and do not contribute significantly to the error in threshold setting. The spread in gains can also be corrected for. What remains is the uncertainty in the charge-injection calibration factors (Fig. 3-2) which dominate the error in the threshold setting.

As explained in Section 3.2, the charge-injection capacitor calibration-factors have about 2% uncertainty. This leads to a spread in the thresholds during running, affecting efficiencies. It is useful to have a cross-check of this spread. If we represent the error in the i th chip calibration factor by f_i then the actual charge injected for a nominal charge Q is

$$Q_i = f_i Q \quad (3.1)$$

so that the Vt_{50} points are at

$$Vt_{50,i} = V_{0,i} + f_i g_i Q \quad (3.2)$$

Here for simplicity, we assume a linear response, which is a good approximation at 1 fC. So the slope of the response curve measures $f_i g_i$ and not g_i . Fig. 3-26 is a histogram of the $f_i g_i$, showing a mean and sigma of 53.6 ± 2.2 mV/fC.

We can unfold the true spread in gains using

$$\frac{\sigma_g}{\langle g \rangle} = \sqrt{\frac{\sigma_{fg}^2}{\langle fg \rangle^2} - \frac{\sigma_f^2}{\langle f \rangle^2}} = 3.6\% \quad (3.3)$$

The gain spread is much bigger than the uncertainty of the calibration factors, which is 2%. If it had been much smaller, we could have assumed all gains to be the

same, and then the response curve becomes a measurement of the calibration factors. This would have allowed the thresholds to be set more accurately.

The default configuration files for the modules in the barrel system were created during reception testing of modules at Oxford, using SctDaq. At that time, the calibration factors were not available and all calibration factors were assumed to be “1.” The calibration were downloaded into the configuration files during barrel macro-assembly at Oxford. During these tests, the charge injected into the front end was corrected by the calibration factor, but the 1 fC threshold DAC was not recalculated since no new response curve had been obtained. When a 10 point gain test was run on this cosmics sector, the configurations were updated with the correct response curve reflecting a more accurate calculation of the 1 fC working point.






Figure 3-26: The gain measured at 2 fC for all chips of the cosmics sector. The width is due to both the variations in gain and the errors in the calibration factors.



(a) The noise occupancy measured at 1 fC for all chips from a Noise Occupancy test before a response curve test.



(b) The output noise measured at 1 fC for all chips from a Noise Occupancy test after a response curve test.

Figure 3-27: The noise occupancy measured at 1 fC before and after the configuration is updated with a new response curve, accounting for the calibration factor differences.

3.7.5 TrimRange

As mentioned earlier, the channel trims must be set to minimise the channel to channel variations in response. An untrimmed module is easily recognizable from the large variation in offset, as shown in Fig. 3-29 and Fig. 3-30. The main aim of trimming is to minimize the variations in the optimal 1 fC threshold setting for physics running so that the efficiency for all channels is the same. This test is also crucial for the long term stability and performance of the SCT as the channel-to-channel variations will increase with irradiation and stable operation of the SCT will require the chips to be set in higher trim ranges.

The trim range scan injects a charge of 1 fC for all events and does threshold scans for different trim DAC settings. Usually it is performed after a 3-point gain scan has checked that charge injection works and before a 10-point gain scan which sets the threshold for 1 fC. 28 threshold scans are made with the first 16 corresponding to each of the 4-bits of trim DAC allowed by the 0th trim range. As most unirradiated module can be trimmed with this smallest trim range, the other trim ranges have only 4 threshold scans done with trim DACs set to 3, 7, 11 and 15 to cover the whole range. The spread of response grows with radiation and larger trim-ranges will have to be used to get a similar response.

For each channel and for each of the 4 trim range settings, the calculated V_{t50} points for the threshold scans are plotted against the trim DAC values. The resulting graphs are approximately linear and are fitted with straight lines, allowing for the estimation of the required trim DAC setting for a particular threshold, [99].

The next step of the analysis determines how many channels are trimmable for each trim range setting. It chooses the trim range setting which maximizes the number of trimmable channels. If several trim ranges can trim the same number of channels, the smallest trim range setting is chosen. A channel is called trimmable in a certain trim range if for each discriminator threshold setting in the range 2.5 to 302.5 mV, called a trim target, there is a trim DAC setting which brings the channel into line with the chosen discriminator setting. A common *optimum* trim target for each chip





Figure 3-28: Histogram of trim range settings for all chips in the SCT barrels.

is centre of the range of possible settings of trim targets for which the number of trimmable channels is maximized. The optimum trim target is related to the offset distribution of the whole chip before it is trimmed, however the offset is increased as a result of the trimming exercise.


The number of channels that are trimmable on each chip for each trim range setting is stored. The lowest trim range that maximizes the number of channels that are trimmable is used as the main criteria for setting the trim range for a chip. Untrimmable channels are masked off at the end of the analysis.

In the TestResult and the summary of the trim range test, the trim range, target, the number of trimmable channels as well as the V_{t50} for the trim target, the offset and the step size derived from the trim DAC is reported. The configuration is updated with the new trim target and trim range for each chip and the new trim DAC settings for each channel.

The trim range settings for all chips in a sector of the SCT barrels are shown in 3-28. While most chips are now in trim range 1, with irradiation, chips will migrate to higher trim ranges. The high trim range chips correspond to those chips with high offset as from Fig. 3-25.




(a) A plot of offset versus channel number for the module after trimming.




(b) A plot of offset versus channel number for the module before trimming.

Figure 3-29: The effect of trimming on offset. This particular module with serial number 20220120001135 on Disk 2 Top-Right Inner 7 of endcap-A was noticed not to be trimmed during CERN reception testing.



(a) A raw scan result at 1.5 fC injected charge for the module after trimming.



(b) A raw scan result at 1.5 fC injected charge for the module before trimming.

Figure 3-30: The effect of trimming on a raw scan result at 1.5 fC for the same mode as in Fig. 3-29. Channel to channel variations in threshold are much smaller after trimming.

3.7.6 Noise Occupancy Test and Occupancy Per Event (OPE)

This test measures the noise occupancy as a function of threshold. The noise occupancy test sends a number of triggers between 2×10^3 and 10^6 depending on the threshold. As the threshold is increased, the occupancy of noise hits decrease from 0.1 to below 10^{-5} and a higher number of triggers are required to measure the noise accurately. At low thresholds and hence high occupancies, recording the same number of triggers as at high thresholds would take a very long time. The scan is started with bins at high thresholds and then the threshold is decreased to avoid heating up the detector and also to take data at high thresholds before possibly running into event errors with long data strings at low thresholds, as mentioned before.

A linear fit to a plot of $\log(\text{noise occupancy})$ vs threshold^2 , as shown in Figure 3-32(a) allows for the estimation of Gaussian noise of each module in ENC. A deviation from this linear behavior, particularly at high thresholds is indicative of non-Gaussian behavior such as presence of common mode noise. The occupancy at a nominal 1 fC threshold and the input noise in ENC are recorded in the test summary and offer a less accurate but different method of measuring the noise from the NPtGain scans. The two methods are not completely independent as the NPtGain scans determine the equivalent charge to be found from the threshold DAC value in mV, that is used in the noise occupancy scan. The input noise values extracted from the noise occupancy method are slightly lower than those from an NPtGain scan, consistent with the observation that the noise measured from a 0.5 fC injected charge is lower by 20% than that from a 2.0 fC injected charge, [28].

An interesting feature of the noise occupancy is that it is heavily dependant on the time the high voltage was turned on, if the noise occupancy is measured within the first few hours. During this time, the noise occupancy decreases and eventually settles to a level, which will continue to be stable for a long time, [29]. A corresponding decrease is also apparent in inter-strip capacitance, [40]. For this reason, the NPtGain results and noise occupancy results from the macro-assembly or the reception testing of barrels, are not easily comparable.

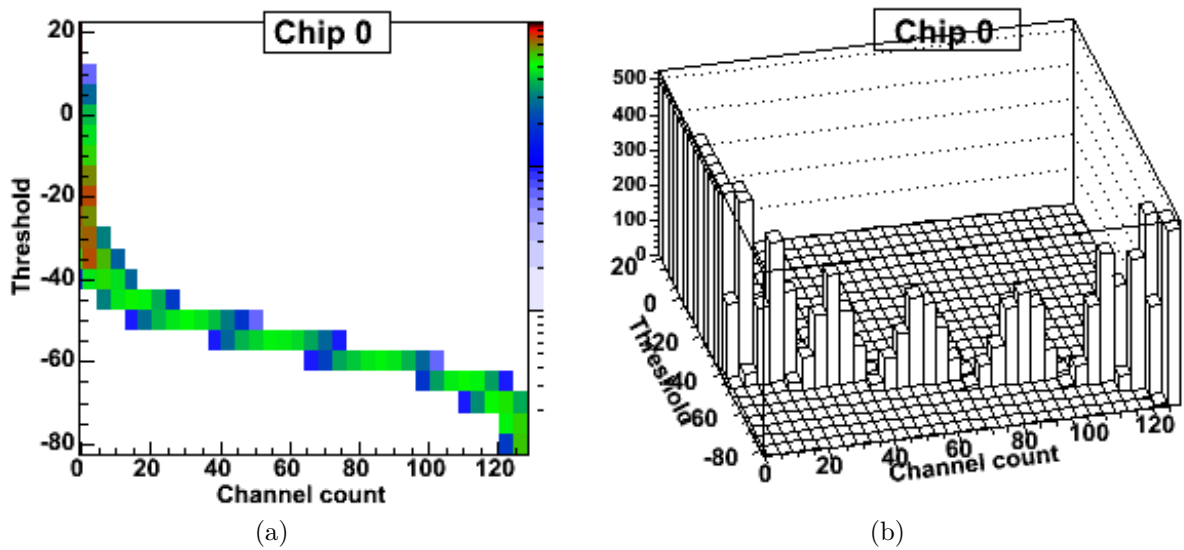


Figure 3-31: A histogram of number of channel fired (occupancy per event) at different thresholds for a chip. (a) is the standard view and (b) is a lego plot of the same data. At a given threshold, the number of channels fired has a Binomial distribution.

The noise occupancy results for the cosmics sector before and after the response curve is compared in Fig. 3-27. The mean noise is lower after the response curve and so is its standard deviation. This is an effect of being able to set the 1 fC point more accurately. Setting a true value of 1.01 fC threshold instead of the nominal 1 fC has a smaller effect on noise occupancy than setting a 0.99 fC true value. The difference in noise occupancy between the 0.99 fC to 1 fC is larger as noise occupancy drops exponentially. Therefore, correcting the response curve has the effect of bringing the noise occupancy down as well as its variance.

An efficient estimate of common mode noise is the occupancy per event (OPE) for each chip. At each threshold, some “n” number of channels out of the 128 channels on a chip are expected to fire on average, as can be seen in Figure 3-31(a). The same plot is shown in a 3-dimensional lego plot in Figure 3-31(b). The distribution of the number of channels firing in an event is given by the Binomial distribution if there is no correlation between the channels. The analysis compares the observed variance of this distribution with the calculated variance, from a simple formula that relates the mean to the variance of the Binomial distribution. The OPE “badness” parameter is defined as the ratio of these two variances. A plot of the badness parameter for

the same chip is shown in Figure 3-32(b). In the summary of an OPE analysis, only the highest OPE badness parameter and the threshold bin number for that badness parameter is reported.

OPE analysis plots are generally displayed with threshold minus the trim target on the y-axis. This is natural as the scan is actually performed with respect to trim target. Modules have different trim targets corresponding to different offsets after trimming. To minimize the number of scan points in the test while scanning through the operational region for all chips, the RODs are instructed to scan through the variable “threshold with respect to the trim target.” Subtracting out the trim target allows for a direct comparison of behaviour around the threshold corresponding to 1 fC. The RODs scan “threshold with respect to trim target” from -80 mV to 20 mV instead of scanning the full range of threshold.

There are two kinds of observed high OPE badness cases. At low thresholds, effects of s-curve wiggles, a known feature of some SCT modules, can produce a high badness parameter. An example is shown in Fig. 3-33, where the scan was terminated before reaching full occupancy at very low threshold. At the lowest threshold setting in this plot, the badness is high as the ratio of the variance of the distribution is very high compared with the average. This is because the occupancy does not decrease monotonically as the threshold is raised, causing the effect known as s-curve wiggles.

However, such cases of s-curve wiggles are harmless for the operating of the detector at 1 fC. For barrel 3, the histogram of all OPE badnesses are shown in Fig. 3-35. As can be seen almost all chips in the barrel have a badness parameter below 2. Only two cases shown in this plot, are at thresholds higher than -20 mV over trim target. As only OPE badnesses at thresholds close or higher than the 1 fC working point will affect the common noise pickup at the operating regime, after the Barrel 3 whole barrel cold test, the analysis was changed to flag only those badnesses greater than 2 for thresholds higher than -20 mV over trim target. For Barrel 3 and other barrels, only a few such cases of badness parameter were observed at high threshold. More importantly, these were shown to be non-reproducible. The worst case of a high badness parameter at high thresholds is given in Fig. 3-34. Here at a threshold

higher than 1 fC, there is one event which has 13 to 16 hits on one chip. This test was repeated several times but this effect was not reproducible. Moreover to put such a cluster of hits in perspective, there are millions of triggers issued for each bin of the noise occupancy test and over 40,000 chips in the barrel system alone. As such an event is non-reproducible, it is negligible as a source of common mode noise and suggests that the pickup in the SCT is minimal.

There are other techniques on how to diagnose common mode noise issues and some of them are mentioned in [30].

3.7.7 Double Trigger Noise Test

The purpose of this test is to identify problematic electrical and optical pickup during the readout of the module. The Level1 buffer depth on an ABCD chip is 132 deep so that the readout of an event on the module happens always 132 bunch-crossings after the event was taken. The test is performed in the following sequence: sending one trigger, waiting for “n” bunch-crossings, sending another trigger. The number “n” is chosen to be close to the Level1 buffer depth as to identify pick-up from the readout of the first trigger, in the readout by the second trigger.

In this test, “n” is varied between 120 and 160, while noise occupancy data is taken at 0.9 fC, down from the nominal 1.0 fC to increase the statistics. A DOUBTR_HI defect is reported if the peak occupancy is 5 sigma away from the baseline or if the peak occupancy is higher than 1×10^{-4} . Two modules on Barrel 3 were discovered to show light leak related problems as seen in Figure 3-36(a). The header bits of the data package can be read off in time. The header is *111010* followed by 4 bits of Level-1 counter and then 8 bits of bunch crossing counter, which are changing throughout the test.

The VCSELS operate at 850 nm and silicon has a good quantum efficiency at this wavelength. The VCSELS operate synchronously with the readout by definition. Also they output 1 mW of optical power and even a small percentage of leak is detectable. A careful visual inspection of Barrels 5 and 6 at Oxford and of Barrel 4 at RAL showed that there were some opto-packages which were not fully covered, as seen in

Figure 3-36(b). These opto-packages are now fully sealed by either gluing the original aluminum foil back or gluing a new aluminum foil on the original one with Araldite 2011. After this fix, no new light leaks were found.

A unique case of high double trigger noise on Barrel 5 deserves to be mentioned. Module 20220170200522 on Barrel 5, LMT10 at Z-4 was tagged as having a very high double trigger ratio of highest peak to background noise for every single chip on the module, as seen in Fig. 3-39(b). Both sides of the module showed stripes in time and the occupancy plot for the lower side is shown in Fig. 3-40. However, the time structure is not the "11101" expected from a light leak, but rather resembles "1100110011".

To debug this problem, the module was set to operate in clock/2 mode but a difference in the high voltage leakage current, which would indicate a light leak, was not observed. Reducing the VCSEL voltage to 3 V to reduce the light output did not have any effect either. Switching off the power to all other modules did not change the behaviour which indicated that this extra noise must be generated on the module or its services. Putting it in select=1, the effect went away, confirming this hypothesis. Going back to select=0 and reducing the TX current to reduce the amplitude of the optical BPM signal did not change the amount of this extra noise either. We therefore believe this effect to be electrical in origin and in some way connected to the operation of the local DORIC, although the details have not been understood. This module was operated in select=0 during the cosmics data-taking but will be operated in select=1 mode in the future.

There are also minor electrical pickup candidates on the barrels which are hardly discernible by eye on the noise occupancy plots. Such chips have on average 80% more noise occupancy for a few bunch-crossing around the time when the read-out cycle starts. About half of the chips on the barrel show some higher noise during these few bunch-crossings. A typical case is shown in Fig. 3-39(a). High-frequency switching noise from the readout is thought to couple into the chip power lines, producing a few extra hits. Such small deviations in noise for a few time-bins is negligible when averaged out over the ≈ 530 time-bins, the average distance between two triggers


running at 75 kHz.

3.7.8 Synchronous Triggers Noise Test

The goal of this test is to diagnose common mode noise problems that might arise only when the trigger is sent at the same time to all modules, as is the case for physics running. The Synchronous triggers analysis is the same as the noise occupancy analysis and so an OPE analysis is also performed. The main question is whether there is any change in occupancy when triggers are synchronous as opposed to when they are not. For all other tests, each four trigger group on a ROD is issued triggers by the MDSP, which only allows for up to 12 modules getting their triggers synchronously. In this test, all modules should get their triggers synchronously. To ensure this, the TIM is asked to take control of sending the triggers to all the RODs at the same time and the triggers are generated at 100 kHz (this is the default value, if not set in the configuration). One event out of 15 is transferred to the SlaveDSPs and histogrammed.


For truly synchronous triggers, the trigger signal needs to reach all the modules at the same time. To achieve this, there are several effects which need to be taken into account:

- Light takes 5 ns/m to travel down a fibre.
- Fibre-lengths on the barrels: A very detailed list of all the fibre lengths of all the different harness types on the barrels is available. The lengths are known to cm-accuracy, [131].
- Fibre-lengths from BOC to barrel: Not well known. Assumed to be equal for all at the Oxford setup. For carrying the signals from the ATLAS pit to the BOC, the fibres were ordered from the factory with a length of 90 m and have a production tolerance of about 1 m. At fibre reception testing, the fibre lengths were measured using a time-of-flight method and the achieved accuracy is better than 1ns, [115].





(a) A typical plot of $\ln(\text{occupancy})$ versus threshold for one chip. The noise occupancy at 1 fC is of observed to predicted variance) versus threshold quoted on the graph. The fitted straight line is a for each threshold bin of the noise occupancy scan fit in the linear region, from which an estimate of for one chip. The badness is usually close to 1 as expected. A large deviation (> 2.0) would indicate coherent noise.

Figure 3-32: Noise Occupancy test result plot and an OPE badness plot for a typical chip.



(a) The plot of the badness parameter showing event with high number of hits at low threshold. (b) The OPE plot showing an event with high number of hits at low threshold.

Figure 3-33: A chip with a high OPE badness at low threshold. Such effects are due to well-known s-curve features at low thresholds, which are not a problem.



(a) The plot of the badness parameter showing event with high number of hits at high threshold.

(b) The OPE plot showing an event with high number of hits at high threshold.

Figure 3-34: A chip with a high OPE badness at high threshold. This occurred in one event and was not reproducible.

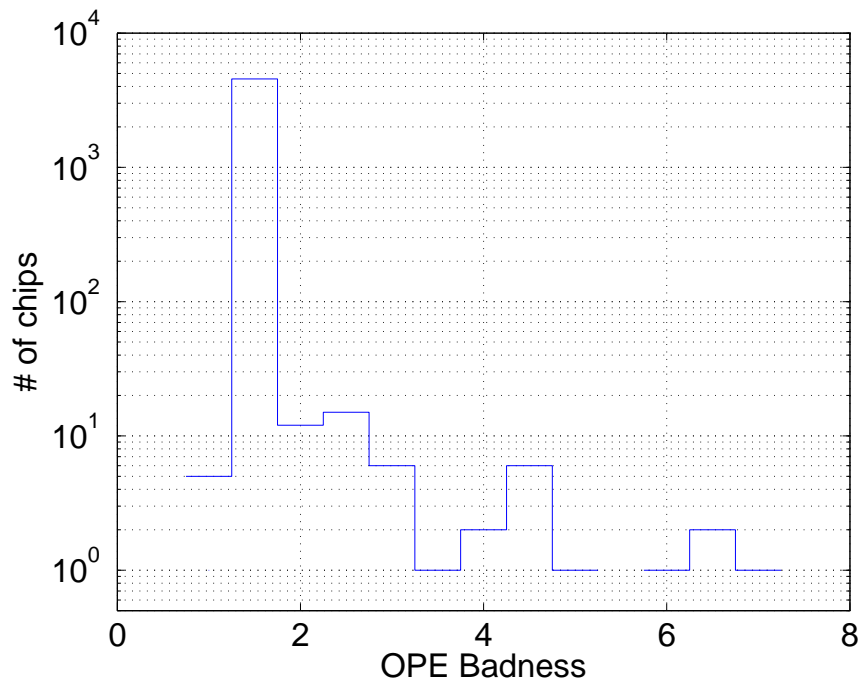




Figure 3-35: Histogram of the worst badness parameter for each chip from Barrel 3. Most chips have worst values below 2; most of those above 2 are due to s-curve features at low thresholds and the others were due to some rare event and are not reproducible. Significant coherent noise would give far more chips above 2.



(a) Occupancy versus time for channels in chip S11 of Barrel 3 module LMT09 z-4. Channels are more likely to fire when readout of a previous event is giving a 1; the header bit pattern of an event has been superimposed.



(b) A potential light leak on an opto-package on Barrel 6, where the cover foil has been damaged.

Figure 3-36: Evidence for light leaks on the barrels from opto-packages.

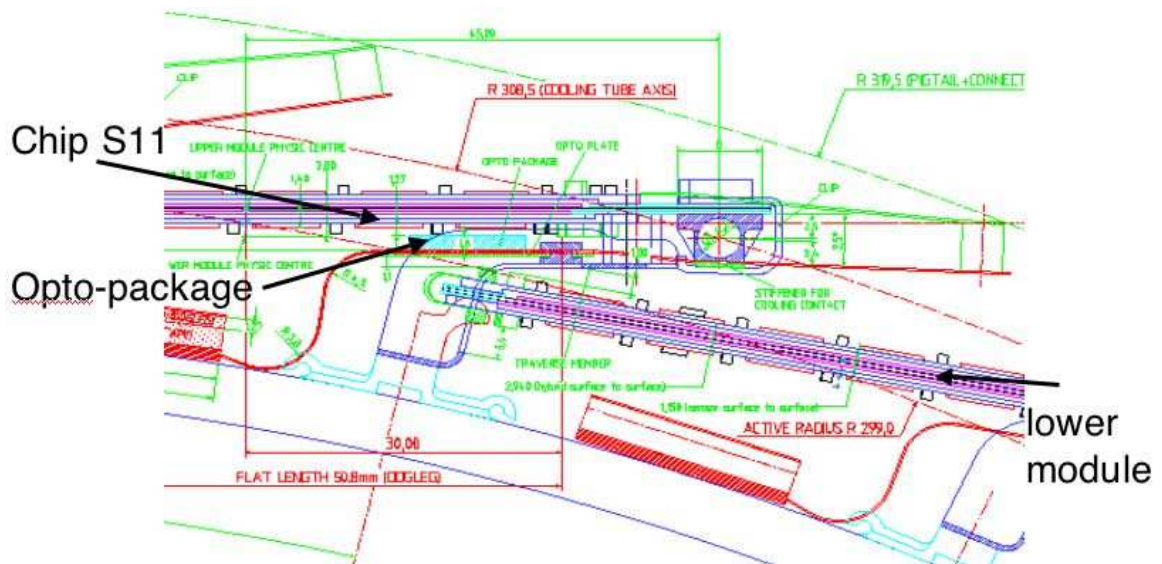


Figure 3-37: A barrel schematic in the $r\text{-}\phi$ plane showing that the silicon strips corresponding to chip S11 of a module are right over that module's opto-package, explaining why chip S11 is most susceptible to light leaks.





Figure 3-38: Evidence of light leaks from a plot of occupancy versus time: approximately 136 bunch-crossings after an event, that event is being optically read out. Then light leaks give far more hits than normal.



(a) A typical electrical noise pickup candidate. The occupancy increase at 126 and 130 bunch-crossings precede any optical activity.

(b) The exceptional electrical pickup candidate on Barrel5, shown in more detail in Fig. 3-40 and discussed in the text.

Figure 3-39: Electrical pickup candidates.




Figure 3-40: The double trigger noise occupancy plot for Link1 of the electrical pickup candidate on Barrel5 LMT10 at z-4, module 20220170200522. The bands of high occupancy around 138, 133 and 137 are very clear. The pattern is not fully understood, but the tests showed that this is self-induced electrical pickup.

The configuration files contain the TX delay variable which by default is set to 0 for all links. As there are 46 different types of “fibre stacks” on the barrels, a script was written to calculate the length of the fibre-stack associated with an MUR (Minimum Unit of Readout) and to replace the delays in the configuration file with the appropriate values. This was done starting with Barrel 5 cold testing. Ever since, no significant change in noise has been found between running synchronously and asynchronously. A handful of high OPE (Occupancy Per Event) badness values have been found for a few chips for one trigger each.

In actual physics running, the goal is not to have the triggers run perfectly synchronously, but to adjust for the time-of-flight of a particle from interaction point to module so that the trigger arrives at the module when the response of the detector will be optimum. Another delay will be added to the TX links to achieve this. This delay will be different for each barrel and module position, but otherwise same in phi.

3.8 Barrel Test Results

A good indicator of how well the SCT works is the total number of working channels. Table 3.10 shows the distribution of defects through the barrels. 99.7% of the channels on the SCT barrel can be read-out. A large contribution to the number of defects is from the dead Rx links, which due to previous ESD damage are foreseen to increase with time. 10 out of these 12 Rx links shown in the table will be recovered as there are modified modules in place and all the data should be available once the software is in place to allow for the offline to handle chip data that is not ordered linearly. However, since this table was compiled after full barrel tests, 14 optical Rx links have been discovered to be dead from electrical tests as shown in Table 3.8. As there are no modified modules in place, these chips will not be recovered and assuming that the master chips from the modified modules are read out, bring the total number of channels that can not be read out due to optical link problems to 2048.

A histogram of individual channel defects on modules is shown in Fig. 3-41. It shows that most modules do not have any defects on them. The bi-modal structure of

Barrel	Channels	Not bonded	Dead	Dead Rx link	Part bonded	Noisy	Other	Total Defects
3	589824	180	357	384	91	460	11	1483
4	737280	55	245	256	16	242	27	841
5	884736	173	770	256	97	492	30	1818
6	1032192	385	2513	640	197	1936	49	5720
Total	3244032	793	3885	1536	401	3130	117	9862

Table 3.10: Defect list from the latest full barrel tests: reception testing of Barrels 3, 5 and 6 at CERN and cold testing of Barrel 4 at Oxford.

this distribution is due to historic reasons. Initially, only the modules which had no dead channels were used for module production. However, by the end of the module production, insufficient number of ABCD chips with 128 working channels forced the SCT to use ABCDs with one bad channel in the module production. Such chips were grouped together on modules, forming a small peak around 12 bad channels per module.

A comparison of the input noise from 3-point gain tests at Oxford module reception, Oxford macro-assembly and CERN barrel reception testing is shown in Fig. 3-42. Although the plots are similar in shape, the input noise between the Oxford and CERN barrel testing is different although there is very good agreement between the Oxford module reception and CERN barrel reception testing. About half of the 140 electron difference in the average noise figures could be explained by the 5 ± 1 ENC decrease in noise per 1° C decrease in temperature, [111, 38] and the other half of the difference is thought to be due to differences in the grounding and screening of the barrels.

The input noise for all 4 barrels are shown in 3-43, where again the CERN test for Barrel 5 is slightly higher than the other barrels. The ongoing commissioning of the barrel in the pit is of high importance as the final cooling and grounding systems are now in place. The test results from that commissioning phase will give the definitive answer on the barrel performance values.

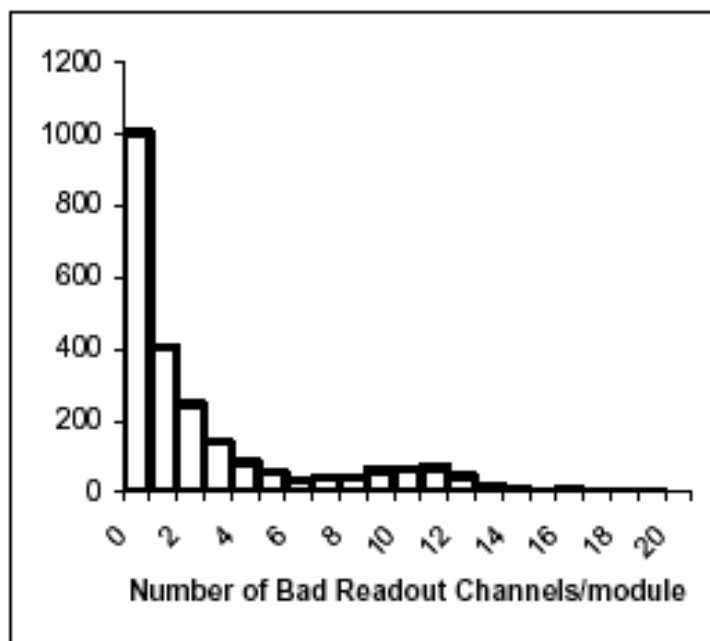


Figure 3-41: The histogram of the number of bad channels per module on all barrels, [2]. Most have zero defects; the second peak around 10 is due to the use of one dead-channel chips.

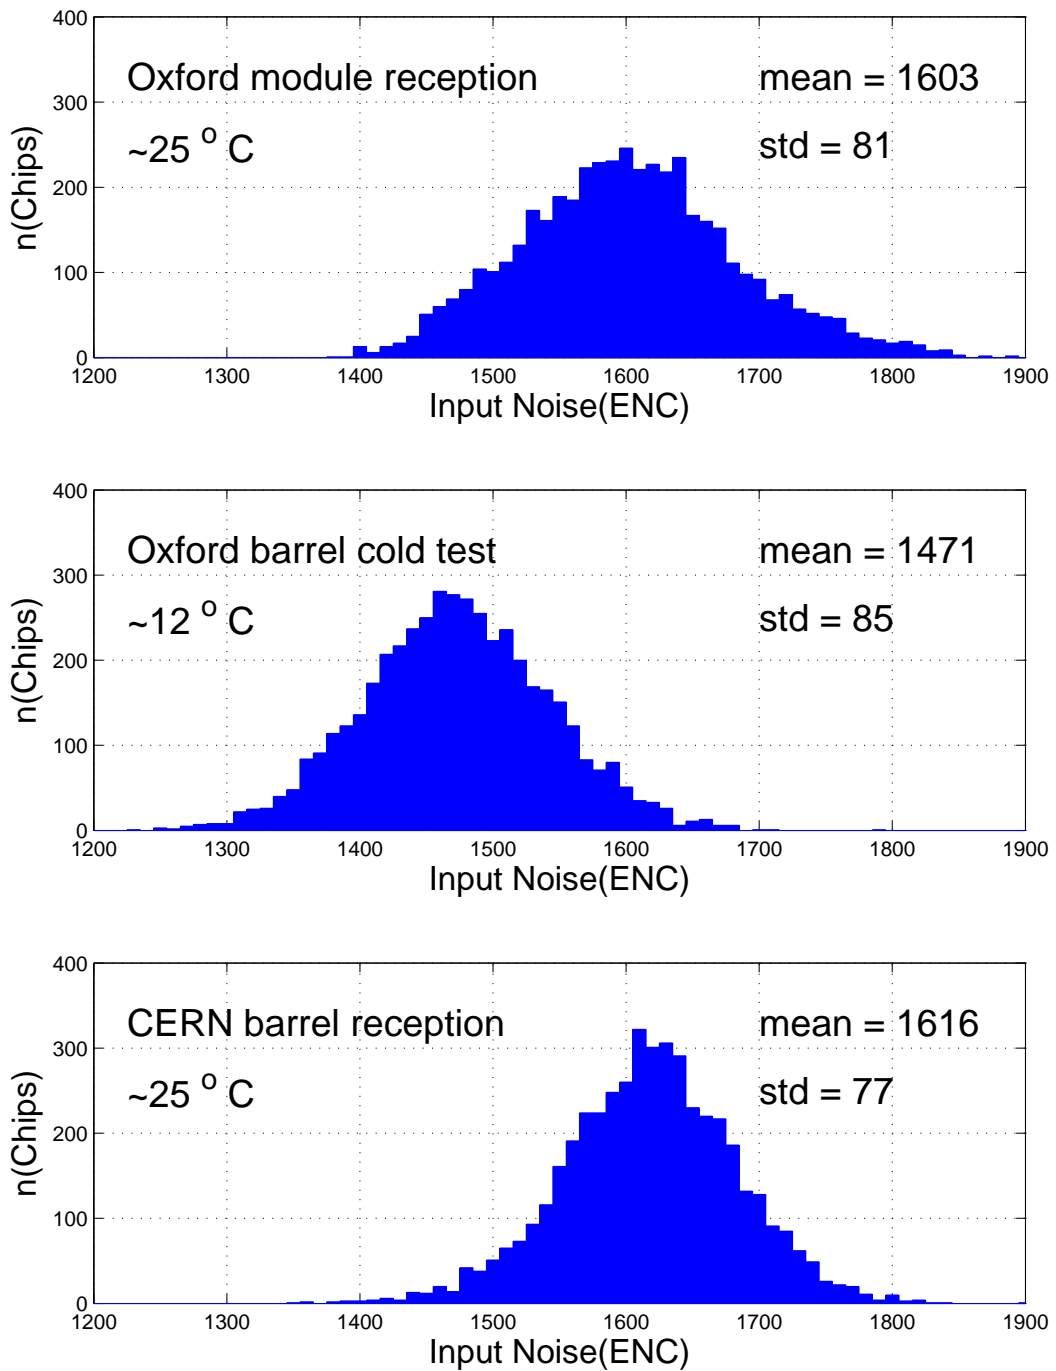


Figure 3-42: The comparison of a histogram of average input noise for every chip for Barrel 3 from Oxford and CERN tests. The difference in input noise can partially be explained by the different temperatures. The assembly of modules onto barrels did not alter the noise characteristics.




Figure 3-43: Histogram of average input noise for every chip for all barrels. Oxford data is from barrel cold test at $\approx 12^\circ$ C. Barrel 5 data is from CERN barrel reception test at $\approx 25^\circ$ C. There is no evidence for barrel to barrel variations.

Chapter 4

SCT Conditions Database

4.1 Implementation of a realistic detector

A good implementation of the detector is necessary in the ATLAS simulation to understand the true physics resolution. To implement a realistic detector, each component of the detector must be modeled first as it is constructed initially and in a way that allows for performance and alignment variations over time. Such an understanding is also important for event reconstruction from real data. For example, missing sectors of the detector change event quality, and alignment is crucial for good particle momentum reconstruction.

There are numerous parameters which quantify the performance of the complicated detectors in ATLAS and most of these parameters will change on a daily basis, and if not, then certainly during the lifetime of the experiment. Such information needs to be stored in a database and be updated regularly by the online tools, monitoring programs and detailed offline analyses. Both the simulation and reconstruction of ATLAS events will need information from these databases and so the access needs to be fast and reliable.

For ATLAS, COOL, a CERN JCOP (Joint Controls Project) solution [125], has been designated to be the database where conditions parameters for simulation and reconstruction will be stored.

Cool is a database with hierarchical structure, [79]. Objects of the same type are

stored or referenced as *channels* in the same *folder*. Each channel has its own interval of validity (IOV), which means they have an associated start and an end “time” between which they are valid. This “time” can either be a run and event number or an absolute time-stamp. The *meta-data* associated with each folder contains the choice which “time” type was used in that folder.

4.1.1 SCT Conditions Databases

For the SCT, there are several parameters that need to be monitored and recorded in a database for such a simulation. Alignment information and some performance variables are also necessary for the reconstruction of actual physics data. It is foreseen that the offline alignment tools such as the FSI (Frequency Scanning Interferometer), the robust alignment, and the global Chi-squared track alignment methods will all write to this database.

The offline needs information from the online software as well, specifically from the DAQ and the DCS. The SCT DCS conditions database will use the general ATLAS-wide tool, [45], the DCS conditions will be stored in COOL [54]. This database will contain crucial information about which modules were turned on and what temperatures the modules were running. The SCT DAQ configurations database stored in CORAL will also contain a sub-set of this information, such as which modules are active.

The SCT DAQ conditions database is also known as the “Calibrations Database.” It will provide information on the noise, defects and other parameters of the SCT which are crucial to a good modeling of the SCT response.

For the SCT conditions database, a decision was taken to use the ATLAS offline identifier as the “key” for identifying channels in COOL. This identifier is defined in Table 4.1. An absolute time-stamp was chosen to index the data in “time” as opposed to “run number.”

Field	mask	bit shift	values	encoded values
indet	0x7	29	2	0 (for indet)
sct	0x3	27	2	1 (for SCT)
bec	0x3	25	-2,0,2	0,1,2
lay_disk	0xf	21	0:8	0:8
phi_mod	0x3f	15	0:55	0:55
eta_mod	0xf	11	-6:6	0:12
side	0x1	10	0:1	0:1
strip	0x3ff	0	0:767	0:767

Table 4.1: Encoding of the Athena offline identifier

4.1.2 SCT DAQ Conditions Database

The SCT DAQ conditions were stored for the SCT barrel cosmics run using a specially modified version of CDI (Conditions Database Interface). CDI is a tool which allows users to archive all IS (Information Service) objects published under a particular IS server in a table format in COOL.

Under tdaq-01-04-01 which was used for SCT barrel commissioning, cdi-00-03-04 was standard. However, there were several features that were implemented only for the SCT and tagged in a release as cdi-00-05-07, [55]. Here is the list of those changes:

- CDI (Conditions Database Interface) was changed to allow for custom folder and channel names in COOL, using the object name in IS as basis. The objects in IS are of the form: *ISSummary.ISNPtGainTestResult.offlineID*. ISNPtGainTestResult or some other name in that field became the folder name and the offline identifier became the channel name.
- CDI stored arrays in COOL version 1.2.x by encoding the arrays as strings, separated by some characters.
- If the string that was trying to be stored in COOL was larger than the maximum that COOL version 1.2.x would allow (≈ 4000 characters), CDI was changed to truncate the string. If a module's HV trips during a test then the defect list grows very large, exceeding this limit. Without this truncation, the CDI would crash.

- CDI was designed to be started before any IS objects were published in the IS server that it is configured to subscribe to and then to archive the IS objects as they were published and then updated. It did not store objects which were already there where CDI was started. However due to the high number of trials needed to get all module communications stable enough to have a test pass on all modules, the SCT required functionality to be able to start CDI after running a test and after deciding that the results were indeed all good enough to be stored. CDI was changed to store all objects in the server upon subscribing to it and then to update those objects if new data was available later.

ONASIC, [9], became the official conditions database interface starting with TDAQ version 01-06-00. SCT decided to move to TDAQ version 01-06-00 to test the new software, especially the new ROS software for the end-cap C calibration and cosmics tests in November 2006. The ONASIC/CDI developers agreed collaborate with the SCT to implement an SCT-specific CDI for TDAQ version 01-06-00 which successfully uploaded DAQ-conditions data for the duration of the endcap-C tests, [85]. With the release of TDAQ version 01-07-00 in February 2007, the SCT is now planning on moving to ONASIC. As the COOL versions have changed between different TDAQ releases, the data previously taken also needs to be migrated to be accessible to the offline software packages.

4.1.3 The SCT IS servers for the conditions database

A decision was made to store only the most basic performance data for the barrel cosmics for proof of principle and to limit the amount of data stored in COOL which needs to be accessed by the offline. For this purpose, the summaries from the NPtGain test and the NoiseOccupancy test were published with a well-defined schema under an IS Server called *ISSummary*. Both objects whose schemas are provided in Tables 4.2 and 4.3 inherit from a common *ISSummaryData* object whose schema is provided here in Table 4.4. The version of these schemas must be incremented if the definition class changes. The *TestData* objects that are produced everytime a test is performed

online are also stored in the database to provide information about the conditions under which the test was performed. It also helps to map run and scan number to the time stamp and provides a lookup table between the offline and online representations.

The data for the barrel sector commissioning was stored in the database before and after the recalibration of the front-ends by a response curve test. Both of these data sets are now available offline through the *SCT_ConditionsData* and *SCT_ConditionsAlgs* packages in Athena, [52].

After the proof of principle of the data-flow from the SCT AnalysisServer to COOL and then to the offline Athena framework, a new SCT package, *ISTestResult*, was written to publish TestResults produced by the online testing procedure to an IS server. This package is now used by the SCT GUI instead of the previous text summaries and has been instrumental in speeding up the GUI display and access to data. An example of an *ISTestResult* object for the Pipeline test is shown in Fig. 4-1. In the future, an implementation of ONASIC subscribing to this new *ISTestResult* server could store all necessary information from the SCT database into COOL.

There is ongoing work for using the SCT DAQ conditions database in the ATLAS simulation. As it stands right now, the *SctDigitization* package which simulates hits on the SCT modules, has a constant value for the average of noise, gain, offset and other parameters and their RMSes for every module. While all modules are implemented in the same way with the same exact set of parameters, each module has 1% of its strips that are dead in the general ATLAS simulation. However for the cosmics simulation and reconstruction, the masked channel list is obtained from a temporary SCT configurations database in COOL, and these channels are masked in the digitization package.

The information provided in the SCT conditions database can improve this simulation significantly. As average and RMS of parameters are stored for each chip, a Gaussian distribution of parameters can be implemented for each strip. Those strips for which a parameter does not fit into the Gaussian distribution are listed in the database under defects as in Table 3.5, so defect masks can be applied to the channel list. Hence a representation of the SCT behaviour on a channel-by-channel basis can

be achieved. The digitization package is also foreseen to get DCS and configuration parameters from the respective databases and uniquely re-create each module's characteristics in the offline software. There is on-going work to implement these in the SctDigitization package.

The reconstruction software can also benefit from the information stored in the databases. For example, knowledge of part-bonded channels as listed on Table 3.5, which are only 50% efficient, would improve their space-point resolution. For channels with stuck-cells, a hit should be masked only if the bunch-crossing counter is equal to the stuck-cell's identifier modulo 12, as opposed to the whole channel being masked off. The information from noisy but still usable channels could be weighed by some factor in the reconstruction which might improve the quality of hit pattern recognition.

With the commissioning in the pit and migration to ONASIC, DAQ conditions data for all modules should soon be available in the COOL database, ready for the offline to use it in simulation and reconstruction of the data.

class = ISNPtGainSummaryData superclass = ISSummaryData	Summary of the N-point gain test Defined in Table 4.4	Type
gainByChip	The mean gain number for each chip.	double array
gainRMSByChip	The gain RMS number for each chip.	double array
offsetByChip	The mean offset number for each chip.	double array
offsetRMSByChip	The offset RMS number for each chip.	double array
noiseByChip	The mean noise number for each chip.	double array
noiseRMSByChip	The noise RMS number for each chip.	double array
version	Version of this class - used in streaming.	u32 init-value1

Table 4.2: The SummaryData object for NPtGain test as defined in the OKS database.

class = ISNoiseOccupancySummaryData superclass = ISSummaryData	Summary of the noise occupancy test Defined in Table 4.4	Type
offsetByChip	The mean offset number for each chip.	double array
occupancyByChip	The mean occupancy for each chip.	double array
occupancyRMSByChip	The occupancy RMS for each chip.	double array
noiseByChip	The mean noise number for each chip.	double array
version	Version of this class - used in streaming.	u32 init-value1

Table 4.3: The SummaryData object for NoiseOccupancy test as defined in the OKS database.

Partition "SCT", server "ISSctData"

Name	Type	Modified	Description
ISSctData.ISTestResult.9567.3.	ISPipelineTestResult	21/2/06 21:21:19	Describes a pipeline test result

Value	Type	Name	Description
9567	U32	runNumber	The run number in which this Summary was extracted
3	U32	scanNumber	The scan number in which this Summary was extracted
20220170200734	String	serialNumber	Serial number of the module.
11/8/05 16:15:06	Time	startTime	Time at the start of the scan.
21/2/06 21:21:19	Time	endTime	Time at the end of the scan.
ST_TRIG_DELAY1	String	configVariable	The variable that was modified in this scan if any.
0	Boolean	passed	Whether the test passed
0	Boolean	problem	Whether the test had problems
3, 4	U32[2]	allScanNumbers	The scan numbers used in this test
0, 1	Double[2]	testPoints	The value of the configVariable used for each scan
	String	user	Who issued the test.
CERN	String	location	Location of test.
	String	host	The host on which the test was done.
SctRodDaq_0_00_Exp	String	daqVersion	The version of SctRodDaq used.
SCAN_OPTIONS:	String[1]	comments	Comments about this test.
	String[0]	dcsDataType	Type of this DCS Data.
	Double[0]	dcsDataValue	Value of this DCS Data.
STUCK_CELL	String[1]	defectType	Type of the defect.
265	U16[1]	defectBeginChannel	The channel that this defect begins on.
265	U16[1]	defectEndChannel	The channel that this defect begins on.
1	Boolean[1]	hasDefectParameter	Whether or not this defect has a parameter that quantifies
2048	Double[1]	defectParameter	The parameter that quantifies the defect if any.
	Double[0]	opeScanVariable	The OPE scan variable.
	Double[0]	opeBadness	The OPE badness parameter for the scan variable.
Pipeline	String	analysisAlgorithmName	The name of AnalysisAlgorithm which created this result
0	U16	analysisAlgorithmVersion	The version of AnalysisAlgorithm which created this result
1	U32	headerVersion	Version of this summary header class - used in streaming.

27 attributes 1 object

Figure 4-1: A Pipeline test summary object as seen in the IS GUI after the re-analysis of the data. Note that the startTime reflects the start of the scan but the endTime reflects the time when the latest analysis was finished.

class = ISSummaryData superclass = Info	Describes a summary data object Defined by InformationService	Type
serialNumber	Serial number of the module.	string
location	Location of test. SR1 or Pit	string
runNumber	The run number in which this Summary was extracted	u32
scanNumber	The scan number in which this Summary was extracted	u32
startTime	Time at the start of the scan.	time
endTime	Time at the end of the scan.	time
defect0Type	Type of the defect on link0.	u16 array
defect0BeginChannel	The channel that this defect begins on link0.	u16 array
defect0EndChannel	The channel that this defect ends on link0.	u16 array
hasDefect0Parameter	Whether or not this defect has a parameter that quantifies it on link 0.	bool array
defect0Parameter	The parameter that quantifies the defect if any on link0.	double array
defect1Type	Type of the defect on link1.	u16 array
defect1BeginChannel	The channel that this defect begins on link1.	u16 array
defect1EndChannel	The channel that this defect ends on link1.	u16 array
hasDefect1Parameter	Whether or not this defect has a parameter that quantifies it on link 1.	bool array
defect1Parameter	The parameter that quantifies the defect if any on link1.	double array
headerVersion	Version of this summary header class - used in streaming.	u32 init-value1

Table 4.4: The generic ISSummaryData class as defined in the OKS database.

Chapter 5

Cosmic Tests of the ATLAS Inner Detector Barrels

5.1 Description of the Cosmic Setup

The four SCT barrels were integrated at CERN with the last barrel being inserted in September 2005. The SCT and the TRT barrels were integrated in February 2006 and a sector of both detectors was prepared for cosmics testing.

The SCT barrel system was integrated from the four individual barrels by first mounting the thermal enclosure of the final barrel into a support cradle and then sliding it over Barrel 6 while it was suspended on an axial spindle. The weight of the barrel is transferred to the thermal enclosure. The services of Barrel 6 were folded out to allow for the insertion of the next barrel inside it. The procedure is repeated for every barrel. Fig. 5-1 shows the insertion of the last barrel, Barrel 3. The services of the three outer barrels have been folded out. There is an inner thermal enclosure fitted inside Barrel 3 and the services for the barrels are sealed into the slots of the thermal enclosure feed-throughs on the sides of the SCT barrel. Then the services are folded on a temporary support structure for the services and the barrel system is checked for gas tightness.

For the SCT and TRT integration, the SCT was inserted into the TRT by rolling the TRT over the SCT, as seen in Fig. 5-2. Since the clearance between the outside of




Figure 5-1: The insertion of Barrel 3 into the other barrels and the thermal enclosure. At the top of the picture, Barrel 3 services are folded in axially; the modules of Barrel 3 are also visible, entering the circular flange of Barrel 4. The services of the outer three barrels are folded out and cover the rest of the picture.

the SCT and the inside of the TRT is about 1 mm, this procedure was very slow and the TRT was guided with rails on the ground. During the movement, the alignment and concentricity of the two detectors were repeatedly verified. After the insertion, the SCT was positioned on rails inside the TRT with a precision of $\approx 250 \mu\text{m}$, [103].

Soon after the insertion of the TRT and SCT, sectors of both detectors were cabled up via the temporary patch panels (TPP). The TPPs for those LMTs (Low Mass Tapes) which were not in the sector were removed to prepare the LMTs for pit cabling. In Fig. 5-2, it is possible to see that for the sector, the SCT TPPs are still attached and the top sector of the TRT had already been cabled before the insertion.

The sector of SCT prepared for cosmic tests consists of 468 SCT modules which can all be read out using one ROD crate. The TRT sector consisted of 12 TRT modules and so a total of 3284 straws. The TRT used $Ar : CO_2 = 70 : 30$ gas for cosmic running. In addition, there were three scintillators as shown in Fig. 5-3. The trigger used the top two scintillators while the bottom scintillator, which was below 20 cm of concrete, was used offline to select higher momentum tracks (above about 300 MeV). These scintillators were read out by one photo-multiplier at each end, [56]. All scintillators were $144 \text{ cm} \times 40 \text{ cm} \times 2.5 \text{ cm}$ and positioned with an accuracy of 1 cm.

The signals from the 6 photo-multiplier tubes are processed to provide the trigger and timing information as sketched in Fig. 5-4. Only the signals from the top and middle scintillator were used to provide the trigger to maintain a high trigger rate. The pulse height and arrival time information from all 6 photo-multipliers was read out and stored.

Each photo-multiplier output signal is fed through an ORTEC935 constant fraction discriminator (CFD) such that the timing signal is always triggered at the same fraction of the PM output signal. The width of the output NIM signal was set to 40 ns, longer than one bunch-crossing (bco) to avoid a double trigger in one time-bin. The threshold level for the input signal was set to 45 mV, below the typical input signal of 200 mV, to be fully efficient even to small output signals from the photo-multiplier.




Figure 5-2: A photo just as the TRT barrel was rolled over the SCT barrel. The TRT is supported on rails on each side of the orange structure; the SCT is supported by the cantilever square beam in the center. The SCT services are folded in axially and arranged in four rings. The top sector of the TRT has already been cabled for the cosmics run. Soon after, the SCT weight was transferred onto the TRT and the cantilever removed.




Figure 5-3: Drawing showing the positions of the scintillators with respect to the TRT in the SR1 cosmics setup at CERN. The black lines in blue boxes represent scintillators in their frames. Shaded regions are concrete.

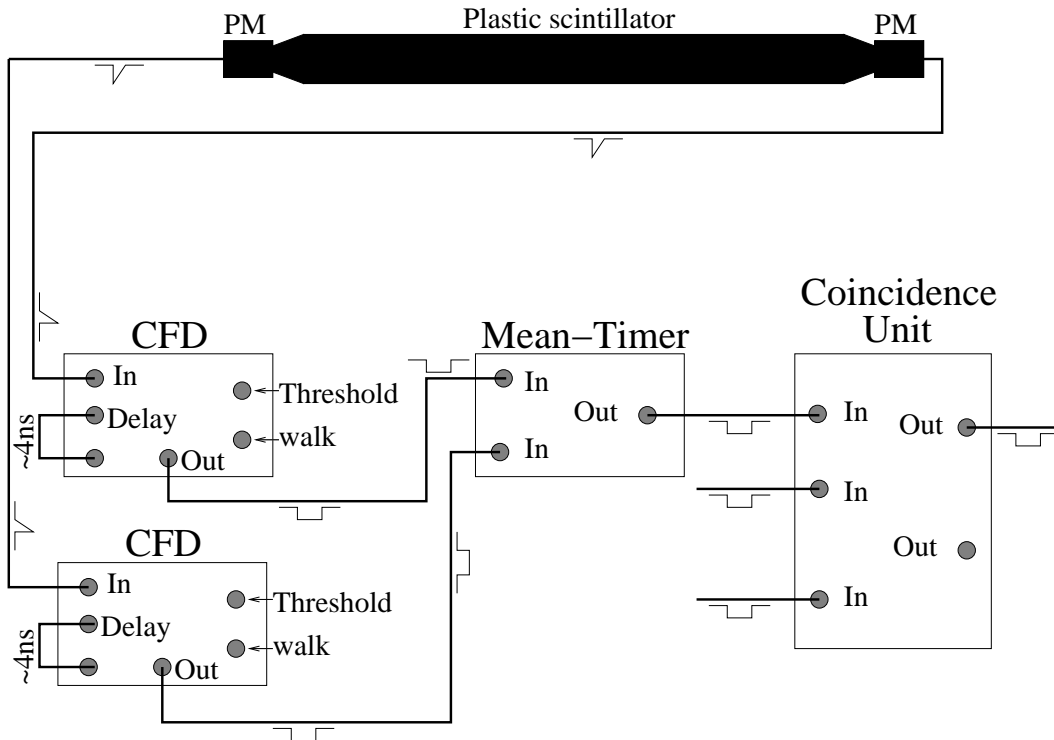


Figure 5-4: Schematic of the cosmic NIM trigger logic. Only one scintillator is shown. The trigger required a coincidence between the mean-timer outputs of the upper two scintillators shown in Fig. 5-3.

The output of two CFD signals originating from one scintillator were fed into a LECROY 624 meantimer channel, which delivers a NIM output signal after a delay with the meantime of the start of the two input signals with an accuracy of 0.5 ns, [46]. The signal duration of such a unit is approximately equal to the total overlap time of its two input signals. Finally, the signals of the top and middle scintillator are fed into a LeCroy LRS466 coincidence unit. The trigger rate of the two-scintillator coincidence was 2.4 Hz, which is compatible with expectations from simulation.

The signal from the middle scintillator was delayed to make it always the last to arrive. Then the muon trigger time is always determined by the time the cosmic ray crosses the middle scintillator. This reduces the trigger timing jitter and allows the TRT to be able to make timing measurements of its own. This muon trigger is fed into the LTP (Local Timing Processor), which waits until the next 40 MHz clock-edge arrives before issuing an “L1A out” (see Fig. 5-9 and section 5.2.)

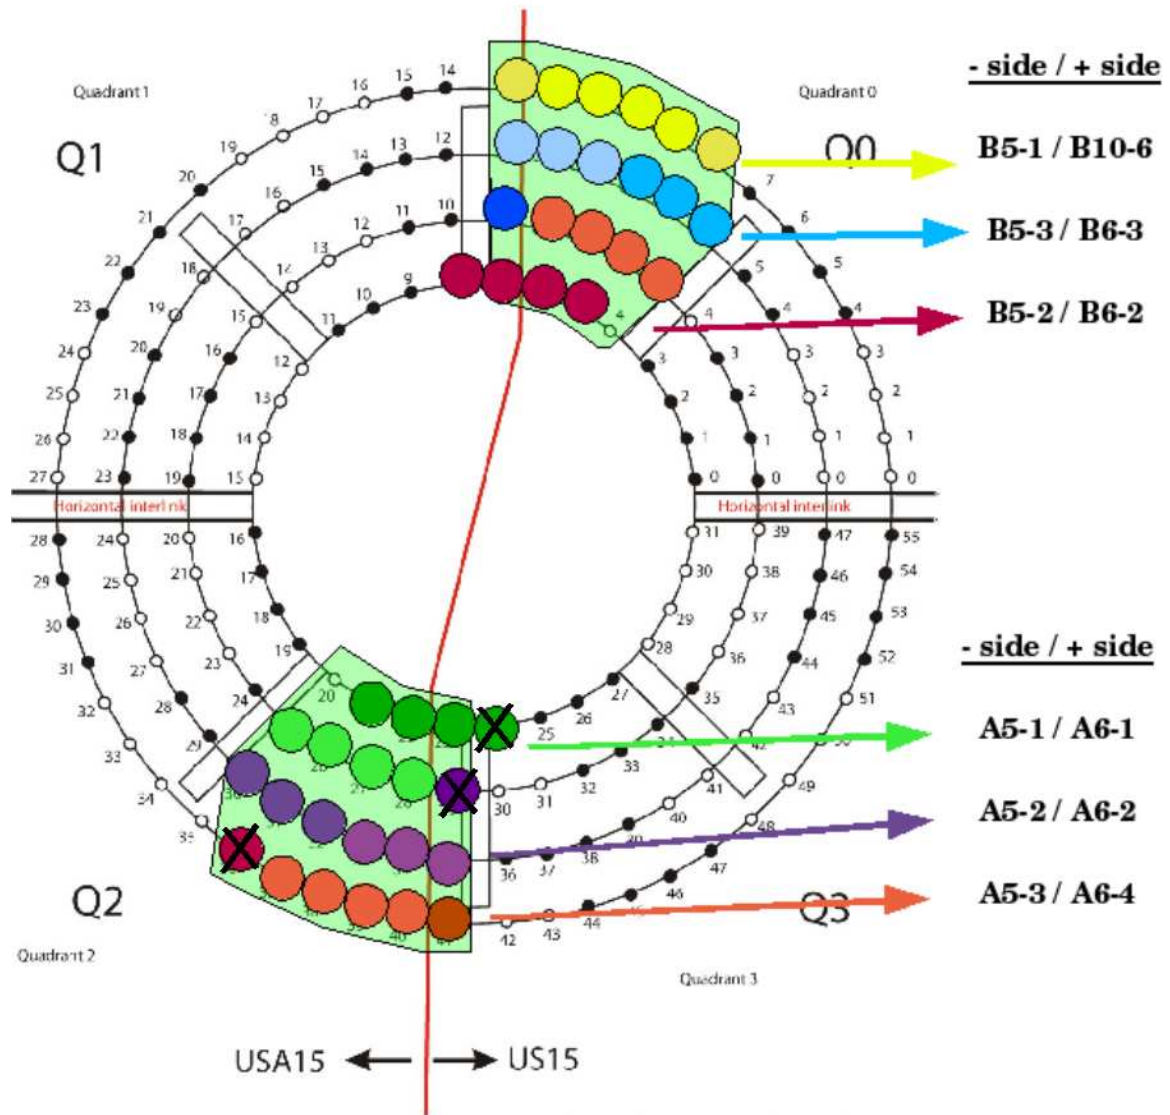


Figure 5-5: A drawing of the power supply crate mapping for the sector of the SCT that was cooled down, powered and read out during the cosmics data taking. The colored pipes represent the sector that was run. Different colors are used for different cooling pipes. The view is from outside the barrel, looking towards Z-. Filled dots represent cooling pipes with manifolds at the Z+ end while open dots are those at Z-. The access to the cooling pipes and read-out fibres was limited and those pipes which are crossed out were not used although originally planned. USA15 and US15 are diff

Stable trigger timing is necessary for the TRT's drift time measurement and for the SCT to be timed-in for optimal charge collection. An estimate of the trigger signal time resolution is presented in the next sections.

5.1.1 Measurement of the phase of the trigger with respect to the system

The efficiency depends critically on the timing, as explained in Section 2.2.1. During ATLAS running, the hits arrive at a fixed time relative to the clock. During cosmics running, the hit times are independent of the clock, and so timing needs special attention. For this, the muon trigger time relative to the clock was measured and recorded.

During test-beam running, the efficiency was measured as a function of beam particle arrival time with respect to the clock, [37]. Fig. 5-6 shows the results. The efficiency reaches a plateau of almost 100% for a time window spanning about 25ns. The timing needs tuning to make all these arrive in one time-bin. Storing the muon-trigger time allows this analysis to be repeated, giving a check that the barrel SCT achieves the required timing precision.

For the cosmics running, a dedicated DAQ chain was commissioned for the measurement of the time difference between the muon trigger and L1Aout, called "trigger time". Charge and the time measurement of the signal on each of the three scintillators was made using a a 6U VME crate with a single board computer with the following read-out cards:

1. CORBO, RCB 8047, counts the number of triggers (L1A) and provides the trigger chain with a busy out, [48].
2. TDC, Time to Digital Converter, CAEN V775, has 12 bit resolution and was used in the mode with 140 ns Full Scale Range [35], giving a 35 psec time resolution.
3. QDC, Charge to Digital Converter, CAEN V792, has 12 bit resolution and the




Figure 5-6: Efficiency versus trigger time from the test beam for a non-irradiated (left) module biased with 250 V and an irradiated (right) module biased with 450 V, [37]. For both the discriminator threshold setting is 1 fC. For a non-irradiated module, the efficiency is shown if the discriminator is operated in level sensing (filled circles) and edge sensing mode (open circles).

input range is from 0 to 400 pC with 50Ohm impedance, [36].

Each of the six signals from the CFD 935 as well as the three signals from the meantimer, were fed into the TDC module. The TDC is used in common start mode using the LTP clocked trigger for the start. It measures the time difference between the LTP clocked trigger and the scintillator stop signals. A delay loop of about 180 ns has been added to ensure that the stop signals arrive later than the common start signal. Each of the analog signals from the six PMs was fed into the QDC to make a measurement of the charge deposited in the scintillators. The mapping of these signals to channels is shown in Table 5.1.

These TDC and ADC measurements were wrapped in the ATLAS data format, [27] and inserted into the ATLAS byte-stream at the event-builder level for the information to be available offline, allowing for charge collection efficiency studies and momentum selection. The VME drivers for the TDC, ADC and the CORBO modules were written by the ATLAS Tile Calorimeter collaboration [117]. The data acquisition system ran under the TDAQ framework (version 01-04-01) and used RCD (ROD

Name	TDC channel	ADC channel
HSC1-B	0	7
HSC1-A	1	6
HSC2-B	2	5
HSC2-A	3	4
HSC3-B	4	3
HSC3-A	5	2
meantime HSC-1	6	none
meantime HSC-2	7	none
meantime HSC-3	8	none

Table 5.1: TDC and ADC mapping information for the scintillator measurements

Crate DAQ) [49].

While a data-driven software trigger-in was used for the read-out of these modules, special attention was paid to the Level1A count that the CORBO module generated from the number of triggers it received. This was done to ensure that this “read-out segment” would not get out of synchronization with the SCT and the TRT segments. If there was a jump in the L1A count from the previous event, an error was reported. Furthermore, a L1A veto was implemented in the LTP (Local Timing Processor) using the BUSY signals from the TDC, ADC and the CORBO.

An ATLAS ROD fragment was formed from the ADC and TDC measurements using the L1A ID provided by the CORBO module. This fragment was then passed through RCD into a ROS application and then to the Event Builder. The source ID of the SubDetector fragment was defined to be “0x70” or “TDAQ Beam Crate,” as listed in Table A.1.

The DAQ system also has a mode, which dumped the raw data from a single VME module for a given number of triggers. This “test” mode was instrumental in optimizing the different setup parameters.

To summarize, trigger timing information is recorded with the data. This allows for charge collection efficiency studies and momentum cuts. A simple online decoder for this information was written and is now included in Athena.

5.1.2 Timing Results

The trigger time resolution can be determined from the information provided by the TDC measurements. Fig. 5-7 shows the distributions of the time measurements for the meantimer signals in the three scintillators in Run 3099, a combined SCT-TRT cosmics run. Since timing is measured with respect to the bunch-crossing cycle of 25 ns, the distribution of the meantime of the middle scintillator, which provides the trigger, is approximately a 'block' distribution with a width of 25 ns.

Events outside the 25 ns window are anomalous. The small tail on the low end of the distribution is due to events in which the time of the signal from the top scintillator arrived later than the signal from the middle scintillator. Such events can occur if the signals in the top and middle scintillators are not generated by the passage of a single particle, but rather by different particles in the cosmic shower.

Fig. 5-8 shows the difference in the measured arrival times of the three scintillators for events with at least one reconstructed track. Once corrected for a possible relative offset due to small differences in signal shape or cable length, this TDC time difference is a measure for the time-of-flight (TOF) of the cosmic particle between the scintillator planes. This TOF measurement can be used for momentum selection.

	mean [ns]	sigma [ns]	expected TOF [ns]
mid - top	14.42 ± 0.01	0.34 ± 0.01	12.85
bot - top	18.43 ± 0.01	0.40 ± 0.01	18.48
bot - mid	4.04 ± 0.01	0.32 ± 0.01	5.62

Table 5.2: Results of the fits to the meantimer time difference distributions. The last column shows the expected time-of-flight for a particle with infinite momentum at perpendicular incidence.

The measured TOF distribution is a convolution of the true TOF distribution and the TDC resolution function. The true TOF distribution is determined by the momentum and angular distribution of the cosmics particles in the trigger. The simple case of a particle with infinite momentum ($\beta = 1$) that traverses the setup perpendicularly can be used to put an upper limit on the time resolution of the setup. Gaussian function fits to an interval of 2 ns around the mean of the distributions are

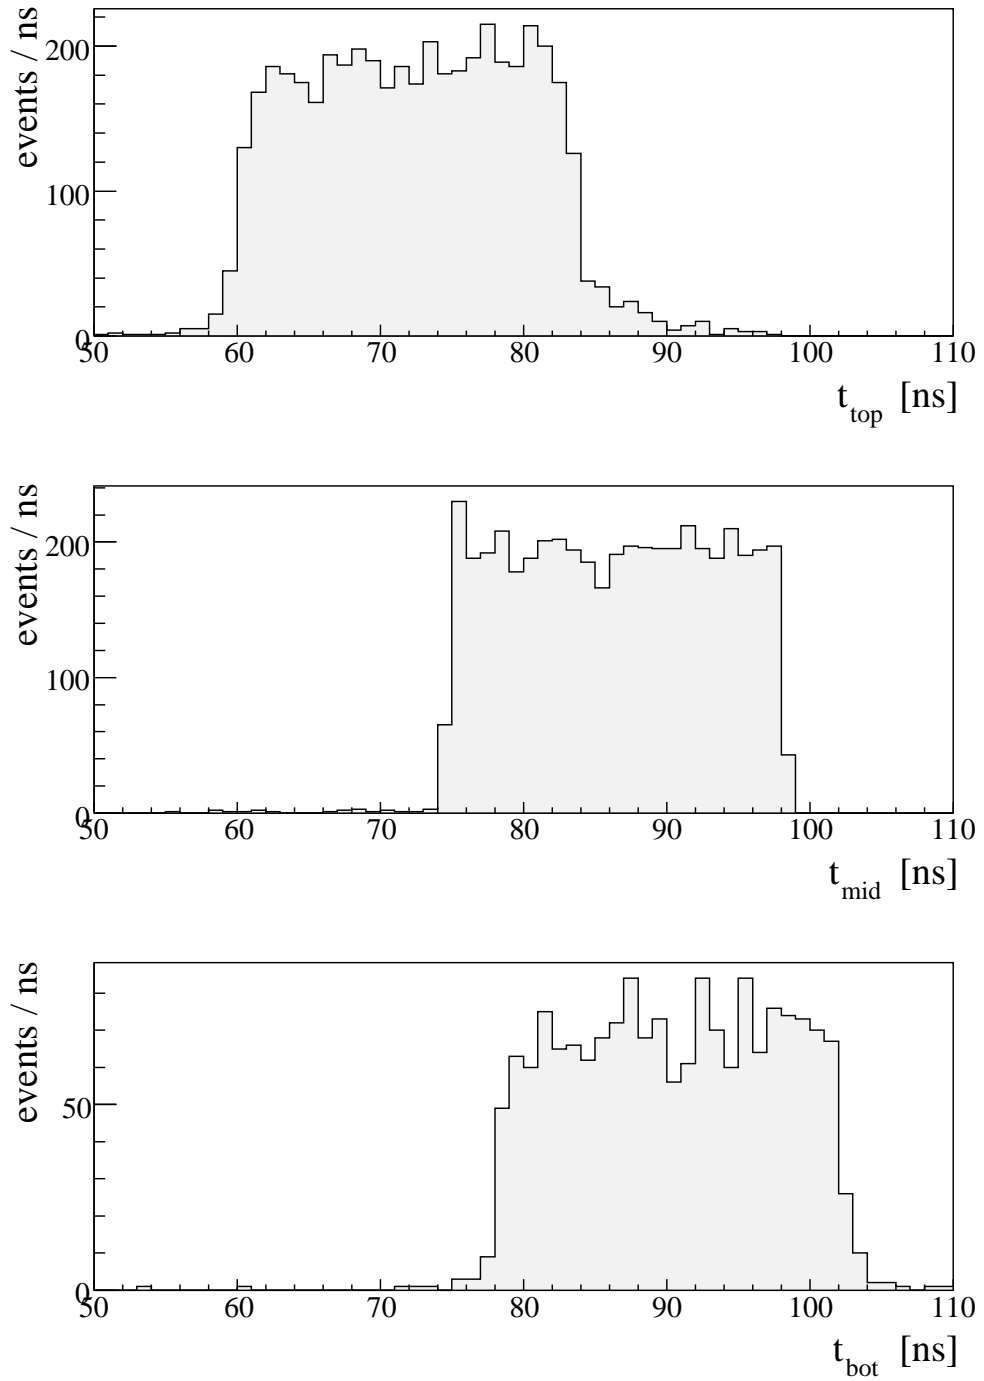


Figure 5-7: Distribution of the time measured in the TDC for the meantimer signals of the three scintillators. For the middle scintillator, which sets the trigger time, all hits are within the 25 ns clock cycle apart from a very small number of accidentals.

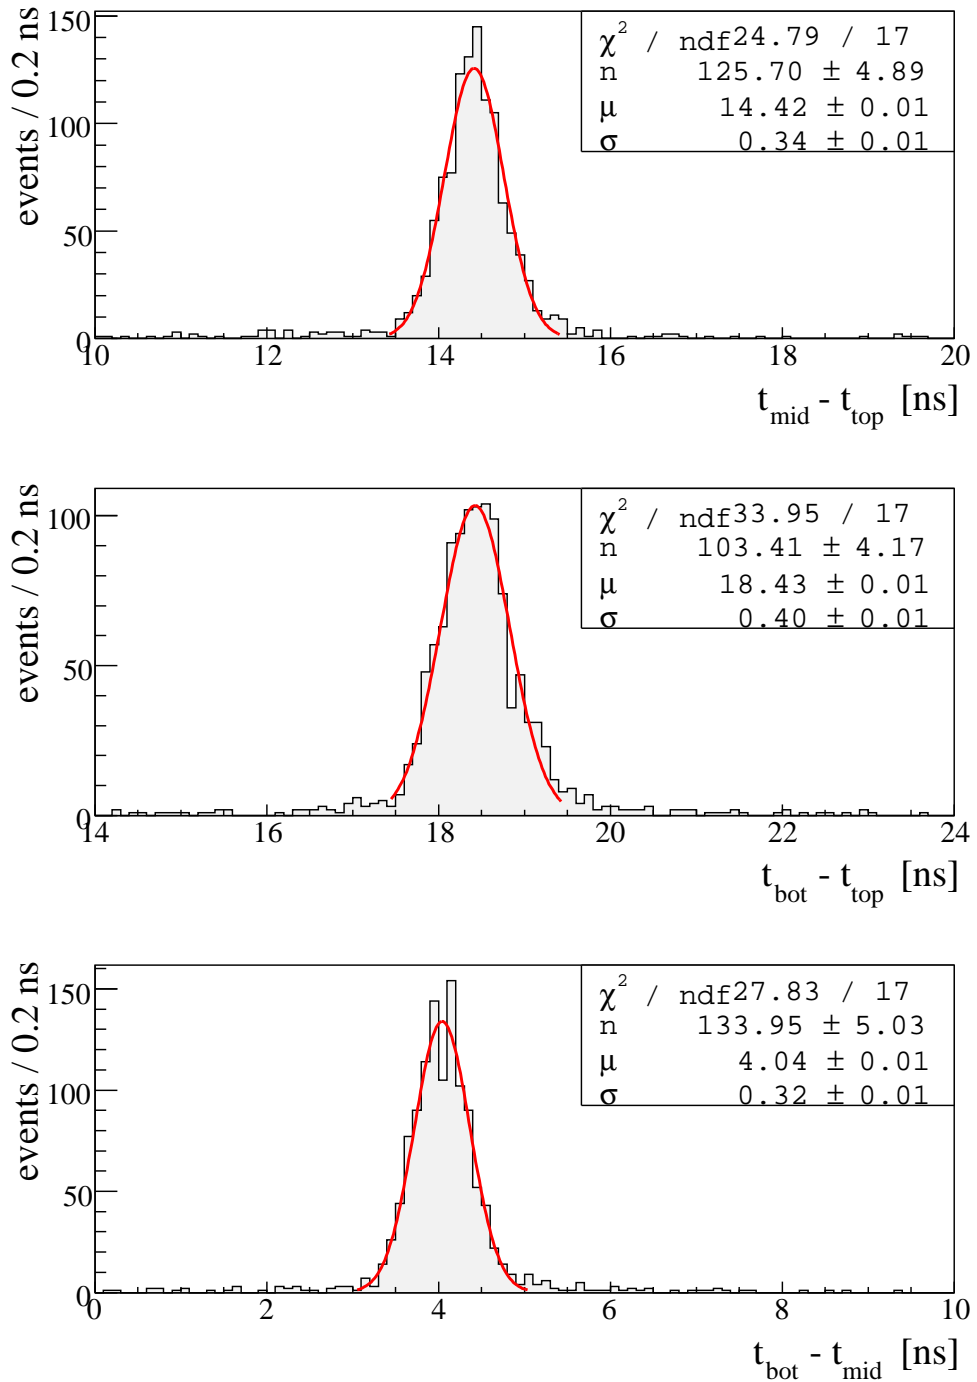


Figure 5-8: Distributions of the difference of the TDC measurements in the three scintillators. The spread is Gaussian with a width of under 400 ps. This is due to time-of-flight differences as well as timing resolution. The means correspond well with the expected time-of-flight.

shown in Fig. 5-8.

The fit results along with the expected TOF from a particle with infinite momentum and perpendicular incidence are summarized in Table 5.2. The fact that the width of the distribution is larger with the larger distance between the scintillators suggest that the spread in momentum or pathlength contributes to the observed width. As the timing resolution is dominated by these considerations, we can conclude that mean time resolution is better than about $0.3 \text{ ns}/\sqrt{2}$.

Care was taken to keep cable lengths from the photomultipliers to the TDC's equal. Still, slight transit time differences are possible. The good agreement of the measured and expected TOF for top and bottom scintillators suggest equal transit times. An extra delay of 1.6 ns for the middle scintillator would explain the discrepancies for middle to top and bottom to middle times.

5.2 Trigger Distribution and Timing

The cosmics trigger was distributed to the SCT, TRT and the TDC/ADC read-out systems by the use of standard ATLAS trigger logic cards:

1. LTP: Local Timing Processor, [70]
2. TTCvi: Timing, Trigger and Control vme interface module [67]
3. TTCvx: Timing, Trigger and Control vme transmitter module
4. TTCrx: Timing, Trigger and Control vme receiver, which is part of the TIM card (Timing Interface Module)

The LTP received the scintillator-trigger NIM signal and generated a Level-1 Accept (L1A) from it at the next rising clock edge. The clock was also generated in the LTP.

The LTP trigger was vetoed by the BUSY signal. The BUSY signal has to be the "OR" of several signals, notably the SCT BUSY signal provided by the SCT TIM, the TDC, ADC and CORBO BUSY signals, as shown in Fig. 5-9. The ORBIT signal

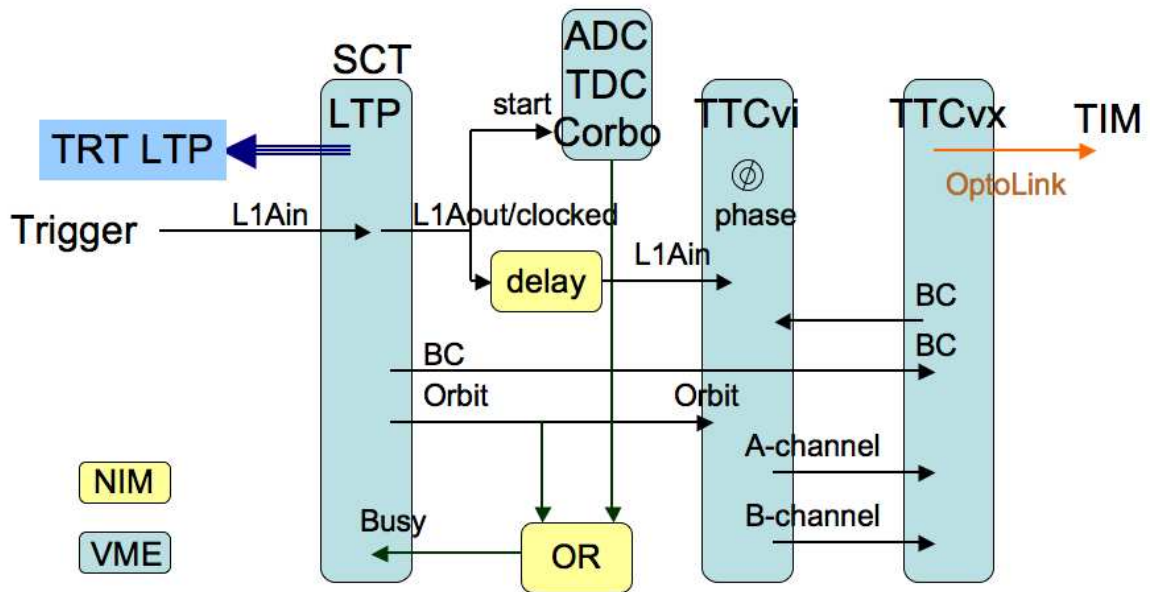


Figure 5-9: The trigger chain for cosmics, with the SCT LTP running as the master LTP. The SCT LTP internally generates the clock (BC) and the orbit signals. The scintillator muon trigger (L1Ain) is fed into the LTP, which waits until the next 40 MHz clock edge arrives and then issues an “L1A out”.

was also added to the BUSY signal. Every $89 \mu\text{s}$, the LHC beam has a gap of 128 bunches, known as the orbit gap. A Bunch-Crossing Reset (BCR) is issued each orbit gap to reset front end and ROD counters. Since the BC counter is not well defined during the BCR, the cosmic triggers have to be suppressed for the duration of the ORBIT signal.

The triggers are suppressed for 6 bunch-crossings after the L1A is issued in the LTP, following the ATLAS requirement. While the SCT read-out chip only requires that L1A be suppressed for 2 bunch-crossings after an L1A because the trigger command is 3 bits long (110), [119], the calorimeters require at least 6 bunch-crossings between two L1As.

With the LHC beam, the SCT will need Event Counter Resets (ECRs) approximately at 1Hz due to Single Event Upsets (SEU), [62]. An ECR sets the ABCD event counter to zero and changes the “extended L1 counter,” which is the most-significant byte of the L1ID. This counter is known as the “ECR Counter” and is stored stored

on the RODs, as shown in Table 3.4. Special handling is needed at the start of the run; the ABCD event counters are zeroed, but then can change before the actual start of a run either due to an SEU or probing of modules. Hence an ECR is required also just before the triggers are enabled.

The SCT needs an ECR just before enabling triggers and during cosmics running. So the SCT sub-system would start with ECR counter at 1. This could be simulated in the ADC/TDC system, however it was not possible for the TRT. The TRT ran with prototype RODs which had no ECR counter and so TRT ROD fragments always gave zero for this counter. To solve this problem, an ECR from the SCT TIM was issued at the “Prepare for Run” stage of the TDAQ “Start” transition and then the ECR counter in the SCT RODs were cleared.

The SCT LTP was configured to be the master and distributed the L1A, BUSY and BC signals to a VME crate with a TTCvi and TTCvx. The signals were carried by a special “CTP Link Out” cable from the SCT LTP to the slave TRT LTP. The BUSY signal from the master is designed to be carried in this cable to the slave, however this was noticed not to be the case. This LTP firmware bug was fixed in a later release, but for the duration of the cosmics running, the BUSY signal had to be sent to the TRT LTP over a separate LEMO cable.

The SCT and the TRT timed-in their read-out to the cosmics separately. Each detector needs to delay the LTP-issued L1A by an appropriate amount to match the pipeline depth of 132 bunch-crossings. The latency of the SCT trigger timing has been carefully investigated before the design of the SCT read-out system. The final latency budget for the SCT operation with the LHC beam is given in Table 5.3.

For the cosmics read-out, the latency budget had to be recalculated for various differences, shown in Table 5.4. The latency in components such as the trigger, TIM, ROD, BOC, cable and fibre lengths is well understood and is about ~ 30 bunch-crossings (bunch-crossings), [92]. The extra delay needed for the L1A to reach the front-ends at the end of the pipeline were calculated to be $\approx 2.350 \mu\text{s}$. Note that is a calculation to guide the timing efforts in the right direction. Exact timing needed was found by scanning through time delays at the BOC, which we shall discuss later.

It is also worthwhile to note here that since clock and trigger should be synchronous for all modules during the cosmics test, each individual clock and command signal is delayed to compensate for differences in fibre length between the modules. For data taking with the LHC beam, the time of flight to the modules will also be taken into account in this adjustment.

For the SCT read-out sub-system, the maximal delay that can be issued in the BOC is 32 bunch-crossings, which is not enough to reach the pipeline depth during the cosmics run. Therefore, such a delay had to be issued before the TTCvi. The TTCvi generates the ECL A-channel and B-channel signals which are sent to the TTCvx for optical transmission to the TTCrx. The A-channel carries the L1A information while the B-channel carries the BCR and clock related signals. The A-channel and B-channel signals have to be synchronized and such a synchronization is provided by the adjustment of the phase on the TTCvi module. For the final operation of the SCT, such a long delay before the TTCvi is not necessary as seen in Table 5.3.

Several delay units were tried to introduce the timing delay needed, but the significant time jitter due to these delay units caused the loss of synchronization of the phase and therefore of the read-out subsystems. In the end, the needed delay for the SCT timing was introduced by a very long cable, which introduced a delay of $2.350 \mu\text{secs}$. The phase of the TTCvi was adjusted to resynchronize the L1A and the clock signals. The trigger signal is distributed to the TDC/ADC/CORBO read-out system before this delay so as to minimize the time that the measured signals need to be delayed by to arrive after the trigger.

Due to the difficulties with having a stable delay without jitter, after the cosmics run, a new method on the TIM has been implemented which will let the SCT set an accurate delay into the timing if needed. This new method introduces 7 bunch-crossings of constant delay plus a variable programmable delay into the trigger chain, [129].

The method above was used for coarsely timing the SCT into the right bunch-crossing for the trigger. The fine timing of the SCT used the histogramming capabilities of the ROD in the physics mode. While in physics mode, the ROD decodes

BC	ns	Latency Item	Responsible
77.9	1947	Central Trigger Processor output	ATLAS Level1 Trigger
0.4	10	Fan-out module	ATLAS Timing, Trigger & Control
1.6	40	8m cable: CTP to TTCvi	ATLAS Technical Coordination
0.1	3	TTC vme interface module	ATLAS Timing, Trigger & Control
0.1	3	0.6m cable: TTCvi to TTCvx	ATLAS Timing, Trigger & Control
0.9	22	TTC vme transmitter module	ATLAS Timing, Trigger & Control
6.4	160	32m fibre: TTCvx to TTCrx	ATLAS Technical Coordination
3.0	75	TTC receiver chip	ATLAS Timing, Trigger & Control
2.0	50	Timing Interface Module	SCT Off-Detector Electronics
0.2	5	Back-plane of ROD crate	SCT Off-Detector Electronics
3.0	75	Read-Out Driver module	SCT Off-Detector Electronics
0.5	13	Back Of Crate card	SCT Off-Detector Electronics
2.0	50	Bi-Phase Mark chip	SCT Links
19.4	485	97m fibre: ROD to Detector	ATLAS Technical Coordination
1.0	25	DORIC decoder chip	SCT Links
7.0	175	ABC(D) read-out chip	SCT Front-End Electronics
6.5	162	SCT Contingency	SCT Electronics Coordinator
132.0	3300	TOTAL = Pipeline Depth	

Table 5.3: The latency in the trigger chain up to the SCT front-end for physics running in the pit.

and formats events and sends them for the event-building up the S-link, [126]. But at the same time, it can sample events and histogram them. For timing-in the SCT, the coincidences on each module were histogrammed and the histograms for all the modules in the cosmics sector were summed, as seen in Fig. 5-12.

One coincidence is defined as a hit in a chip on one side of a module and a hit in a chip in a module on the other side, for which the corresponding strips of the two chips physically overlap, as shown in Fig. 5-10. As the average noise occupancy for the SCT is 5×10^{-5} . As there are 7×10^5 channels in the cosmics sector, we expect 35 single hits from noise. The probability of getting a coincidence from noise hits in the whole cosmic sector is $35 \times 3\text{chips} \times 128\text{channels}/\text{chip} \times 5 \times 10^{-5}$ which is about 70%. As cosmic particles produce coincidences from single muons that traverse the whole detector, we expect ≈ 8 coincidences in each event.

The histogramming task on the ROD, looks at the chip hits for each time-bin for each chip. For each hit, it checks to see if any of the 2 or 3 chips on the other side of the module have hits. As each check between one chip and a chip on the other side of

BC	ns	Latency Item
5.0	125	25m cable: Scintillator to NIM trigger crate
3.2	80	Trigger logic
2	50	Local Timing Processor
0.1	3	TTC vme interface module
0.9	22	TTC vme transmitter module
3	75	15m fibre: TTCvx to TTCrx
3.0	75	TTC receiver chip
2.0	50	Timing Interface Module
0.2	5	Back-plane of ROD crate
3.0	75	Read-Out Driver module
0.5	13	Back Of Crate card
2.0	50	Bi-Phase Mark chip
6.0	150	30m fibre: ROD to Detector
1.0	25	DORIC decoder chip
7.0	175	ABC(D) read-out chip
94	2350	EXTRA DELAY
132.9	3320	TOTAL

Table 5.4: The latency in the cosmics trigger chain up to the SCT front-end for SR1 setup. Note that the table contents do not exactly add up to the 132 bunch-crossings but a higher number. It is thought that this difference is due to the phase adjustment in the TTCvi.

the module is defined as a possible coincidence, the maximum number of coincidences on a module is 16.

The first coincidence histogram produced during the commissioning of the top sector of the cosmics is shown in Fig. 5-11. To produce this histogram, the timing was adjusted with the calculated delays using the trigger budget, and already, the observed coincidences indicated that the SCT was coarsely timed into the cosmics. The Tx coarse delay in the BOC is scanned, while the number of coincidences only in the middle time-bin of the 3 time-bins is plotted against this delay. The data compression mode selected for cosmics reads out if there is a hit in any of the 3 time-bins. In this histogram, coincidence bin 0 shows the noise hits for which there was no coincidence on the other side of the module. The coincidence bin 1 is where there is a clear peak at TxDelay of 0 and 1.

In the first version of the on-ROD histogram, the time of accumulation for each bin was a constant, while in the final version, the number of triggers was monitored



Figure 5-10: A schematic drawing of a module showing which chips overlap and so might have hits in coincidence due to the stereo-angle between the two sides of the module.

and the bin was changed after it had at least a certain number of triggers in it, as can be seen in Fig. 5-13. It is known as `PhysicsTxDelayScan` and can be started from the SCT GUI. The most crucial part of running a histogram in physics mode is the need for issuing of the `TIM BUSY` before the ROD can change bins. This histogramming task on the ROD can not handle triggers while changing bins so the triggers need to be suppressed during this time. As this is inconvenient for ATLAS timing-in, a work-around is currently under study.

In the final and faster version of the on-ROD histogramming code, the coincidences in each of the 3 time bins are read out in expanded mode while scanning through the trigger delay in steps of 3 bunch-crossing in the BOC card, as seen in Fig. 5-13. Fig. 5-12, where only the coincidence bin 1 is plotted for simplicity, shows an enhancement of coincidences in the 0th and 1st time bins.

Once the coarse delay for the timing was set, the fine delay adjustments began. Several `physicsTxDelayScans` were performed at different Tx fine delay settings, and a global 8 ns fine delay for all SCT BOC was chosen since it evenly distributed the signal over two time-bins. As the phase of the cosmics trigger is random, the signal can not be confined to one time-bin.

After the SCT and the TRT were timed into the cosmics trigger, both systems took synchronous data along with the TDC and ADC system. A single muon track is shown in Fig. 5-14. One coincidence hit, or space-point, is seen on each SCT barrel. The TRT track has 33 hits along the track. There are no hits in the first 9 layers since the track is close to $\eta = 0$ where these layers are not read out.




Figure 5-11: The first coincidence histogram produced of coincidences of hits from two sides of each module, from the top sector, with the initial timing. The frequency of one coincidence is high for coarse delay of 0 and 1, showing the SCT was well-timed using latency calculations. The excess of zero-coincidence module for TxDelay of 0 or 1 bunch-crossing compared to the flat noise level at other delays is not fully understood and should be investigated offline.




Figure 5-12: The histogram of single coincidence of hits from all modules versus bunch-crossings, made by scanning through the coarse trigger delay. This plot is for an 8 ns fine time delay, which minimizes the coincidences in the side-bins (-1 and +2) and was selected for the final running.

5.3 Physics Data Taking

The hardest challenge for the combined SCT and TRT DAQ teams was the synchronization of the readouts. Once both systems were timed-in to the trigger, sometimes, either the TRT, the SCT or the TDC/ADC sub-system would lose one L1A. Such a loss was tracked down first to the BUSY signals that were not propagated properly and once that was solved to the critical phase adjustment in the TTCvi.

Initially, it was hard to diagnose a loss of synchronization, as the Read-Out Sub-systems (ROSeS) and the EventBuilder buffers the events and runs asynchronously, building and processing events whenever they arrive. During this debugging phase, the offline byte-stream converter developers were trying to solve their configuration problems so the light-weight online byte-stream converter and the TRTviewer both proved to be very useful. The online byte-stream decoder written by the SCT DAQ team can recognize the ROD headers from the TRT, SCT and the TDC/ADC segments, and can display the important L1ID and BCIDs in the headers as well as decode SCT and TDC/ADC fragments. The TRTviewer is a GUI which was written by

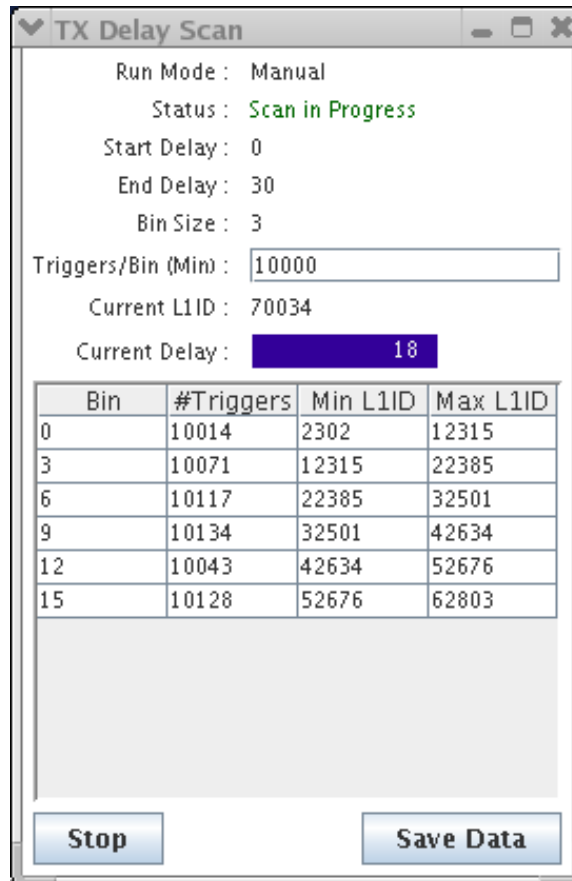


Figure 5-13: The physicsTxDelayScan progress in physics mode as shown on the SctGui. The SctGui monitors the number of L1As histogrammed for each bin. The bin is changed as soon as the number of L1As exceeds 10,000. The slight differences in the number of triggers in each bin are due to the slow and relatively infrequent nature of this monitoring process.




Figure 5-14: The track of a cosmic muon traversing through the top sectors of the TRT and the SCT. The points used in the display to illustrate the hits are exaggerated in size. Here, a TRT hit is enlarged by a factor of ~ 10 and an SCT space point is enlarged by a factor of ~ 200 . The noise hits in the TRT sector can be seen and the noise occupancy is around 2%, [47].




Figure 5-15: The track of a cosmic muon traversing through the TRT and the SCT sectors.

the TRT DAQ team as an online TRT display and has a light-weight byte-stream converter for the TRT. When the SCT byte-stream converter was modified and plugged into the TRTviewer, it became a very powerful tool for viewing events and diagnosing synchronization problems, eventually leading to a stable read-out system.

During the eventual continuous cosmics data taking, the up-time was 57%. The continuous running mode lasted for 92 hours and a total of 450,000 cosmics were recorded in 52 hours during this time (at 2.4Hz trigger rate). The other 40 hours were spent trying to solve problems that arose. The stop and start nature of this week of data taking can be observed in Fig. 5-16. A problematic run would either crash or be diagnosed and stopped generally before the number of triggers reached 1000. Most cosmics physics runs reached 20,000 events with the biggest one having 56,473 events, while the noise runs towards the end of the running could run for a larger number of events, but in a much shorter time. The dataset collected during this combined physics data taking is given in Table 5.5.

Run type	Number of events
Cosmics:	
cosmics at trim target	280,000
cosmics at 1 fC with new-RC	170,000
Noise Runs:	
events at 0.9 fC RC to look for correlated noise	900,000
events [0.9 – 1.2 fC RC] threshold scan	490,000
events in trigger rate noise scan	220,000
events with different grounding schemes	14,000
events looking for pickup from TRT	90,000
TRT heater pickup tests (on/off)	66,000

Table 5.5: List of physics data taken during cosmics commissioning.

The problems encountered during this combined data taking phase could not all be debugged at the time and were debugged later. Two of the most important and time consuming problems were the ROS crashing due to a large ROD fragment and the ROD boot-up problems. We now will discuss these in some detail.

Occasionally, a run would stop with the SCT ROS issuing an error “ROD fragment too large for allocated page size.” Page size is the memory that can be allocated for




Figure 5-16: Run number versus the total number of events taken in one run. Noise runs with high trigger rate have 10^5 events in, while the cosmic runs have $\approx 10^4$ events corresponding to a few hours of running with 2.4 Hz cosmic trigger rate. The runs with problems were diagnosed usually within a few minutes and have less than 10^3 events in them.

one event on the FILAR. The FILAR is the card that receives the data from the S-link in the ROS machine. The page size for cosmics running was set to 4kB, significantly larger than the expected average ROD fragment size of 100 bytes from noise occupancy and cosmics. If the incoming ROD fragment is larger than this allocated page size of 4 kB, then the ROS gives an error and the DAQ goes into an error state, stopping the run. This particular mode of error was specifically very hard to debug since the ROS machine would crash at this point and would not allow for log-ins to try and debug this problem. It is likely that the ROS crashes while trying to dump the contents of its memory to file. When the ROS crashes, it needs to be rebooted, which in return requires the ROD crate to be rebooted, since the S-links are turned off when the ROS is rebooted and will not be turned on until the ROD crate is rebooted. Hence, this error mode caused a significant downtime for the cosmics running.

To try to understand the problem, the page size was increased from 4kB to 8kB. However, the runs which stopped due to this error were not less frequent than before. During the course of the cosmics running, there was no further attempt made to debug this problem directly.

During the running of the cosmics, it was noticed that if a module went into clock-by-two mode, the ROD could still recognise this as a valid event header and attempt to send a necessarily infinitely long fragment up the S-link. In this case, the ROD fragment certainly would exceed any allocated page size and would crash the ROS. A new formatter firmware was developed to reject ROD fragments without the valid header. However, the use of this new firmware did not improve the situation; the ROS errors still happened at a similar rate as before.

In the weeks following the cosmics running, there was an effort to reproduce the same exact problem with the H8Testbox. However, this effort was unsuccessful. Next, simulated single-event fragments produced on the ROD were pushed up the S-link to understand at which point the ROS gives this error. The ROS did not crash in these tries. When the page size was set to 8 kB and a 4 kB ROD fragment was sent up the S-link, the ROS still issued the error that the “ROD fragment was too large for

allocated page size.” At this point, the ROS experts were contacted and the following valuable information was revealed: The page size for the ROS is restricted to certain discrete values and any attempt to set a value except these will default to the closest lower value. The page size can be increased to 4 MB, however this functionality could not be tested since the ROD can simulate single-event fragments of 12 kB maximum. In the next running of the SCT, it is foreseen to increase the page size to its maximum size to try and capture these events which must have either extremely high occupancy or a corrupt very long event which the ROD is not filtering out.

The other big problem during combined running was the ROD boot-up problems. When the RODs are first booted up, after a hard ROD crate reboot, the communication between the SBC and a ROD will occasionally time out. Once a ROD does boot, in the next attempts, it will always boot on the first try. The reason for this ROD booting problem is not well understood but the buffer transfer between the ROD and the SBC is under investigation as a possible source of this error.

Unfortunately, Athena monitoring could not be utilized to its full capabilities during this running period. If the Athena monitoring task was in the Run configuration, then the start-up would take a very long time as Athena had to reference several offline repositories and databases. Quite often, the EMON service would crash and even if everything else was stable, the number of histograms that the monitoring service would issue would bring the DAQ to a grinding halt. Also, the prerequisite for being able to start Athena monitoring is that the configuration for the run that is about to be started is in the configurations database before the system is configured. But, this is unnatural for the detectors, as they have not been configured yet and the configuration for that particular run might need to be changed to get all the modules to return events. Quite often, the SCT DAQ team would upload the configurations in advance into the database, only to find out later that the configuration in the database had been outdated by the time the run could be started. Development of workarounds to these problems are now under way, as the input from Athena monitoring will be very valuable for the pit commissioning of the detector.

During physics running, there were also some data errors in the byte-stream.

As stopping and re-starting a run was time consuming and contributed to dead-time, some errors were fixed by changing DCS and BOC parameters during physics running. Here we will give a list of such problems and if possible, fixes:

- A High Voltage trip could be diagnosed by a sudden rise in occupancy from a module or by the DCS monitoring system. A HV trip could be fixed by masking the HV channel off and then turning the HV on again. Some of these HV trips were later found to be related to a firmware bug on the HV card, which has now been fixed.
- A Low Voltage trip could be diagnosed by a “Timeout error” from a module or by the DCS monitoring system. Although LV can be turned back on, the module can not be configured while the run is in progress. Only if there were two or more LV trips would the run be stopped and the modules re-configured.
- Level1 and/or BC errors as seen in the byte-stream. As all the counters on the modules have been tested and there is no significant irradiation that might lead to SEU, such errors must result from the instabilities of the optical communication, e.g. from the TX signal not being received correctly, resulting in the missing of a clock cycle or a trigger; or, if the errors are only in a handful of events during a run, they must be due to the RX signal not being correctly discriminated by the BOC. In such a case, the RX values can be adjusted on the BOC. RxThreshold values with 10 more and 10 less than the current setting were tried with the link that was returning errors, and this was sometimes successful in avoiding errors from that link. If RxThreshold adjustments do not improve the situation, it was assumed that RxDelay tuning was needed, but this was not attempted by hand. If the errors are in several consecutive events, then it means that the module has lost synchronization with the rest of the SCT. BC errors were often seen for only one event, as after a BC-reset, the BC counter is re-synchronized with the rest of the system. A persistent BC or L1 error implies that the TX communication is unstable. In such rare cases, the *p-i-n* current was checked and if low, then the TxCurrent setting was increased

by a few DACs. Rare cases of persistent BC errors were fixed by this method, however there was no trivial way to fix persistent L1 errors. The run could be “paused” and a L1 reset and then L1-counter reset sent to the modules to synchronize them back with the TRT. However, this only worked if the TDC/ADC segment was disabled in the running, as this segment always gave fatal errors if paused. For this reason, persistent L1 errors could only be fixed by stopping and re-starting the run. A few L1 errors were not considered to be important enough to go through the tedious stopping and re-starting procedure.

- Zeroth bin anomaly. There was an interesting feature which was first noticed in the byte-stream by use of the online decoder and later confirmed by the offline, known as the “zeroth time bin anomaly.” When the SCT is run in any-hit mode, there can be a hit any of the 3 time-bins, but there should be at least one hit. However, it was noticed that a handful of modules were returning channel-hit information but with no hit in any of the 3 time-bins. This anomaly is currently not understood and is documented in detail in [25].

5.4 Results

The cosmic runs were the first test of the whole dataflow from the online to the offline, of reading out the detector in time, checking synchronizations, building and then writing events, monitoring tools, byte-stream converters, data displays, offline detector-configuration tables called “cabling” information and offline analysis tools. As it is a substantial sample of data with 5×10^5 muons, it was used for noise and efficiency measurements but also was instrumental in establishing the first alignment results. The present residuals from tracks are better than the specified build tolerances, [112].

The efficiency of a given SCT wafer using the cosmic data can be calculated using a sample of good tracks and after a set of fiducial cuts, [19]. The efficiency algorithm removes any hits located on the barrel layer under investigation, or the i^{th} layer. A track refit is then performed excluding these hits. From the perigee parameters of this




Figure 5-17: The unbiased residuals measured for the different layers and sides of the SCT modules for simulation (red triangle) and data using global χ^2 alignment (blue squares). A road width of 2 mm is used.




Figure 5-18: The unbiased hit efficiency measured for the different barrel layers and sides from cosmic data, before and after alignment. The 99% hit efficiency design criterion has been met.

new track, an extrapolation to the i^{th} layer is performed to obtain the intersection point or “predicted position” of a hit on a module. If this intersection point is within the fiducial area of the module it is included in the denominator of the efficiency calculation. If an SCT cluster is found to be located within a set distance from the predicted position, an entry is made for the numerator. Only the following tracks are considered for the efficiency calculation:

- The track must have 10 or more SCT hits.
- The track fit must have a χ^2 per degree of freedom ≤ 24 : This is an extremely loose cut, and only removes a small number of very strange tracks, [80].
- the track must have an incidence angle on the module of less than 20 degrees in the azimuthal direction.
- the track intersection point on the module must be at least 2 mm from the bond gap and 1.5 mm from the edges of the module
- the track intersection point must be at least 1.5 mm from any masked strips or chips.

In order for a hit to be entered into the numerator of the efficiency calculation, in addition to the above requirements, a hit must be found within a certain roadwidth around the predicted hit position. By default, and unless specified otherwise, this roadwidth is 2 mm. In order to choose a proper value for this cut, simulation was used. Since the momentum of the particle is unknown (no magnetic field), the effect of multiple scattering is not known.

Efficiencies have also been computed in data making use of the global χ^2 alignment method, Fig. 5-18. After alignment the unbiased hit efficiency in all barrel layers is measured to be within specifications, i.e. greater than 99%.

The resolution of the SCT is also as predicted. Fig. 5-17 compares the unbiased residuals from simulation and cosmic data. The data is very close to the simulation. It is greater than the 22 μm hit precision because it includes errors in the track fit prediction and multiple scattering.




Figure 5-19: The unbiased SCT residuals for fast (left) and slow (right) particles, [81]. The time-of-flight cut for producing these plots is placed at 18.5 ns. The RMS of residuals for slow particles without any alignment corrections is 229 μm while it is 200 μm for fast particles, showing an improvement due to a smaller multiple-scattering effect for higher momentum muons.

It is important for the alignment efforts to have a clean single-track muon sample devoid of low energy muons which multiple scatter in the detector volume and produce secondary hits. The use of the timing information from the TDCs was demonstrated to improve the resolutions. Using the RobustAlignment method, a time-of-flight cut was used to have a cleaner sample of particles. As shown in Fig. 5-19, the RMS of residuals of the track fits (in the direction perpendicular to the strips in the local frame) without any previous alignment is 229 μm , if only particles with time-of-flight higher than 18.5 ns are considered, [81]. It is 200 μm for particles with time-of-flight lower than this cut, showing the improvement due to suppression of low momentum muons.

The cosmic data set was the first real alignment challenge for a substantial part of the ATLAS tracker. The results of the χ^2 track alignment method from cosmic hits on tracks can be quantified by fitting Gaussians to the peak of the residual distribution.

The width before alignment was 65 μm , while after alignment, it was reduced to 32 μm . This agrees well with the Monte-Carlo prediction of 31 μm for a cosmic sample with energies higher than 200 MeV for a perfectly aligned detector, [109].

The noise runs taken under different conditions have shown no effect on the stability of the system. For example, analyses have looked for dependence of noise on trigger rate. Runs with 5, 50, 500 Hz, 5 kHz and 50 kHz trigger rate have been compared and no change in noise occupancy was seen. Different grounding schemes with and without shorting the thermal enclosure to ground, with and without the TRT reading, or with the heaters switching on and off, produced very little change in noise. The biggest effect comes with the application of the response curve from the calibration to the physics configuration, which reduces the noise, as documented in Chapter 3.

A good example that there is no self-pickup comes from two plots shown in Figs. 5-20 and 5-21. Here, the multiplicity of hits in the event is histogrammed for a noise run and a cosmic run. For the noise run, a Gaussian fits well to the observed noise with no events with high multiplicity, which would indicate pickup. The addition of cosmic adds a correlation tail to the plot, which is lacking in the noise plot, showing that indeed there are no correlated hits in the noise run.

An interesting exercise is to compare the noise from the online and offline analyses. For this comparison, noise run number 2900 and calibration run number 2981 were used. Both runs are after the cosmic sector configurations had been updated with the results from the response curve. The physics run 2900 was done in any hit mode with a random trigger, but only the X1X hits were selected to compare the noise for only for time-bin, as in the calibration run. The difference in standard deviations of the noise occupancy for these two runs is shown in Fig. 5-22. Here, the difference is calculated in terms of the standard deviation of the noise from the calibration run as

$$difference = \frac{NO(calibration) - NO(physics)}{std(NO(calibration))} \quad (5.1)$$

where NO stands for noise occupancy. The mean is slightly above 0, so the calibration




Figure 5-20: The occupancy of events histogram for a cosmic run for all 3 time-bins, [108], with a cosmics trigger.




Figure 5-21: The occupancy of events histogram for a noise run for all 3 time-bins, [108], with a random trigger.

run is seen to be more noisy than the physics run. This effect is expected since the offline imposed a tighter cut on the noisy channels than the online software did. The average noise occupancy calculated using the offline software is 4.5×10^{-5} , [136].

Location	Serial number	File origin	Reason	Difference	Calc. eff.
Barrel3 LMT08 z+2	20220330200475	Cambridge file	unknown	1.94	0.98
Barrel3 LMT09 z-4	20220330200415	QML file	unknown	3.17	0.97
Barrel4 LMT08 z+1	20220170200887	KEK	LGS candidate	3.07	0.99
Barrel6 LMT10 z+5	20220040200377	Berkeley	LGS candidate	4.22	0.97

Table 5.6: List of modules with edge detect turned on for the cosmics sector and the noise difference between the offline and the online software in standard deviation of noise from the online. Here the measurement efficiencies are also given from the online Athena monitoring framework.

The tail of this distribution requires more discussion. Four modules which were in the barrel sector cosmics run were recently discovered to be in the wrong read-out mode, namely the edge-detect mode, while running in physics mode. The DAQ should re-configure modules to be in level-mode for a physics run, just as it does for




Figure 5-22: The histogram of the noise occupancy difference between physics and calibration mode in standard deviations for each module in the cosmics sector.

calibration mode. Instead, the configuration registers were set as in the database files. This bug was discovered during the end-cap C cosmics, which was after the barrel cosmics run.

The latest configuration files in the database for end-cap C modules were from module production institutes, which wrote configuration files with edge-detect turned on. As the noise of all modules on end-cap C in physics mode was consistently higher than in calibration mode, the problem was noticed sooner and a fix was found. On the other hand, almost all of the barrel configuration files that were downloaded were from Oxford module reception testing, where the DAQ had been altered to write the edge-detect register off in the configuration file, leading to similar average noise figures for calibration and physics modes. For the barrel cosmics run, four modules had their configuration file downloaded from module production sites instead of Oxford, and hence edge-detect mode turned on. These modules are listed in Table 5.6.

Two of these modules showed large gain spread (LGS) features when tested at cold-temperatures, so instead of the warmer Oxford module reception testing configuration files, the configuration files produced at colder temperatures at the module reception sites were downloaded and used as a cure to these features. It is unclear why the other two modules have files which originated from elsewhere but it might be because the Oxford reception configurations had not been uploaded to the database at the time when the barrel was assembled and the configuration files were downloaded to test these modules as assembled onto the barrels.

It is interesting to note that the efficiencies for these 4 modules running in edge-detect mode are also lower than most modules. Almost all modules had better than 99% efficiency during the cosmics run, [80]. All of these 4 modules have lower efficiencies than the average, but the sample size is insufficient for this effect to be statistically significant. It should be noted that the efficiency for the modules have been measured with no timing considerations. As seen from Fig. 5-6, the efficiency of the modules depends significantly on the trigger timing. It is foreseen that the timing information from the TDC setup could be used to further refine the efficiency measurements in the future.

Chapter 6

Conclusions

The ATLAS SCT Barrel integration with the TRT barrel is now complete. 99.7% of the 3.3 million SCT channels are working.

The final data-acquisition system, which is of critical importance for a stable and working SCT due to the binary nature of the read-out, has been developed and tested extensively. Both the physics and the calibration mode of the SCT have been exercised. Multi-crate read-out allowing for systems larger than 600 modules to be read out at the same time is under progress. SctRodDaq has been built to minimize the downtime of the detector and to be fast enough to calibrate the SCT during the LHC beam-fills. The method of timing into triggers has been demonstrated during the cosmics run and will be used to time into cosmics triggers in the pit and later to the LHC beam.

The methods for quantifying the performance of the SCT have been established by the SctRodDaq software and refined over the macro-assembly stage of the SCT integration. The input noise measured at 2fC on the barrels at $\approx 25^\circ\text{C}$ is ≈ 1600 ENC, which is very close to the results from module reception testing, and shows that there is no large contribution from coherent noise. The noise occupancy at 1fC is $\approx 5 \times 10^{-5}$, an order of magnitude lower than the design specification of 5×10^{-4} . There is a very low level of electrical pickup related to read-out activity which will not contribute significantly to the overall noise level.

The cosmic tests of the SCT have provided a good test of the operation of the

SCT in physics mode. The tools to time the detector in were developed and used successfully. The offline analysis showed that the SCT modules had the expected high hit efficiency, while maintaining very low noise occupancy. The cosmic tests also allowed for the development of the offline tools and an initial understanding of the alignment of the barrels system. The residuals from fitted tracks indicate that the detector was built to a much higher precision than that specified by the build tolerances, and the hit position resolution is as expected.

The SCT and TRT barrels are now installed in the ATLAS detector and are in the process of being commissioned; they will soon be followed by the end-caps.

Appendix A

Appendix

TTC Partition	Source ID
FULL_SD_EVENT	0x00
PIXEL_BARREL	0x11
PIXEL_FORWARD_A_SIDE	0x12
PIXEL_FORWARD_C_SIDE	0x13
PIXEL_B_LAYER	0x14
SCT_BARREL_A_SIDE	0x21
SCT_BARREL_C_SIDE	0x22
SCT_ENDCAP_A_SIDE	0x23
SCT_ENDCAP_C_SIDE	0x24
TRT Ancillary Crate	0x30
TRT_BARREL_A_SIDE	0x31
TRT_BARREL_C_SIDE	0x32
TRT_ENDCAP_A_SIDE	0x33
TRT_ENDCAP_C_SIDE	0x34
LAR crates	0x4*
TILECAL crates	0x5*
MUON crates	0x6*
TDAQ_BEAM_CRATE	0x70
TDAQ crates (other)	0x7*
OTHER	0x81

Table A.1: The abbreviated list of ATLAS source identifiers as defined in the eformat library of the TDAQ-common package. Here, * is used as a wild-card.

API	Application Programming Interface
ASIC	Application Specific Integrated Circuit
ATLAS	A Toroidal LHC ApparatuS
BOC	Back Of Crate card
CDI	Conditions Database Interface
CERN	The European Organization for Nuclear Research
CMOS	Complementary Metal-Oxide Semi-conductor
CORBA	Common Object Request Broker Architecture
CTP	Central Trigger Processor
DAC	Digital to Analogue Converter
DAQ	Data AcQuisition
DCS	Detector Control System
DORIC	Digital Optical Receiver Integrated Circuit
DSP	Digital Signal Processor
EFB	Event Fragment Builder
ENC	Equivalent Noise Charge
FFTV	Fixed Frequency Trigger Veto
FPGA	Field Programmable Gate Array
FSM	Finite State Machine
GUI	Graphical User Interface
IDL	Interface Definition Language
IPC	Inter-Process Communication
IS	Information Service
LHC	Large Hadron Collider
LMT	Low Mass Tape
LTP	Local Timing Processor card
MIP	Minimum Ionising Particle
MRS	Message Reporting Service
MUR	Minimum Unit of Readout
OPE	Occupancy Per Event
RCC	ROD Crate Controller
RCD	ROD Crate DAQ
ROD	ReadOut Driver
ROS	ReadOut System
SBC	Single Board Computer
SCT	Semi-Conductor Tracker
SEU	Single Event Upset
TDR	Technical Design Report
TIM	TTC Interface Module
TRT	Transition Radiation Tracker
TTC	Trigger, Timing and Control system
VCSEL	Vertical Cavity Surface Emitting Laser
VDC	VCSEL Driver Chip
XML	eXtensible Markup Language

Table A.2: A list of commonly used acronyms

Bibliography

- [1] A. Abdesselam et al. The ATLAS SCT Optoelectronics and the Associated Electrical Services. 2006. ATLAS note, ATL-COM-INDET-2006-011, submitted to NIM.
- [2] A. Abdesselam et al. The Barrel Modules of the ATLAS SemiConductor Tracker. *Nucl. Instrum. Meth.*, A568:642–671, 2006.
- [3] A. Abdesselam et al. The ATLAS SemiConductor tracker end-cap module. *Nucl. Instrum. Meth.*, A575:353–389, 2007.
- [4] A. Ahmad et al. The Silicon Microstrip Sensors of the ATLAS SemiConductor Tracker. *Nucl. Instrum. Meth.*, A578:98–118, 2007.
- [5] T. Akesson et al. Tracking performance of the transition radiation tracker prototype for the ATLAS experiment. *Nucl. Instrum. Meth.*, A485:298–310, 2002.
- [6] T. Akesson et al. ATLAS Transition Radiation Tracker test-beam results. *Nucl. Instrum. Meth.*, A522:50–55, 2004.
- [7] T. Akesson et al. Operation of the ATLAS Transition Radiation Tracker under very high irradiation at the CERN LHC. *Nucl. Instrum. Meth.*, A522:25–32, 2004.
- [8] 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.*, A522:131–145, 2004.

- [9] A. Amorim et al. ATLAS ONASIC User Guide. http://sim.fc.ul.pt/sim_en/ATLAS:ONASICUserGuide.
- [10] A. Amorim et al. Use of CORBA in the ATLAS prototype DAQ. *IEEE Trans. Nucl. Sci.*, 45:1978–1982, 1998.
- [11] A. Andreazza. Test beam performance of the ATLAS Pixel Detector modules. *Nucl. Instrum. Meth.*, A565:23–29, 2006.
- [12] F. Anghinolfi et al. Asic wafer test system for the atlas semiconductor tracker front-end chip. *IEEE Trans. Nucl. Sci.*, 49:1080–1085, 2002.
- [13] ATLAS Collaboration. *Technical Design Report*, volume ATLAS Calorimeter Performance. CERN, 1996. CERN/LHCC/96-40.
- [14] ATLAS Collaboration. *Technical Design Report*, volume ATLAS Tile Calorimeter Performance. CERN, 1996. CERN/LHCC/96-42.
- [15] ATLAS Collaboration. *Technical Design Report*, volume ATLAS Inner Detector. CERN, 1997. CERN/LHCC/97-16 and CERN/LHCC/97-17.
- [16] ATLAS Collaboration. *Technical Design Report*, volume ATLAS Muon Spectrometer. CERN, 1997. CERN/LHCC/97-22.
- [17] ATLAS Collaboration. *Technical Design Report*, volume ATLAS Detector and Physics Performance. CERN, 1999. CERN/LHCC/99-14 and CERN/LHCC/99-15.
- [18] ATLAS Collaboration. *Technical Design Report*, volume ATLAS High-level trigger, data acquisition and controls. CERN, 2003. CERN/LHCC/2003-022.
- [19] ATLAS SCT and TRT Collaborations. Combined SCT and TRT performance tests in the SR1 assembly area. In preparation as an ATLAS note.
- [20] B. Aubert et al. Construction, assembly and tests of the ATLAS electromagnetic barrel calorimeter. *Nucl. Instrum. Meth.*, A558:388–418, 2006.

- [21] T. J. Barber et al. Resonant bond wire vibrations in the atlas semiconductor tracker. *Nucl. Instrum. Meth.*, A538:442–457, 2005.
- [22] D. Barberis. ATLAS Inner Detector developments. *Nucl. Instrum. Meth.*, A446:331–337, 2000.
- [23] A. Barr et al. The SctRodDaq Wiki. <http://www.hep.phy.cam.ac.uk/~atlasdaq>.
- [24] A. J. Barr. Studies of supersymmetry models for the ATLAS experiment at the Large Hadron Collider. Ph.D. Thesis, Churchill College, Cambridge University, 2003.
- [25] A. J. Barr and T. Byatt. Investigation into the zeroth time bin anomaly in the SCT barrel modules in the ATLAS detector. 2006. ATLAS note, ATL-COM-INDET-2006-013.
- [26] G. A. Beck et al. Measurements of the thermal conductivity of pyrolytic graphite substrates for use in SCT modules. 1998. ATL-INDET-98-221.
- [27] C. Bee et al. The raw event format in the ATLAS Trigger & DAQ. 2005. ATLAS communication, ATL-DAQ-98-129.
- [28] P. J. Bell. Constraints on Anomalous Quartic Gauge Boson Couplings at OPAL and the System Test of the ATLAS Barrel SCT. Ph.D. Thesis, Birmingham University, 2003.
- [29] P. J. Bell and J. Grosse-Knetter. Investigation of the Decay of Noise Occupancy with Time Using Module 0009 on the Electrical Stand. http://asct186.home.cern.ch/asct186/barrel_web/es_investigation_2001-10.html.
- [30] P. J. Bell and J. Grosse-Knetter. Techniques for the measurement of common mode noise in binary read-out systems. *Nucl. Instrum. Meth.*, A516:153–159, 2004.

- [31] W. Bialas et al. Testing specification for the wafer screening. Project Name: ABCD3T ASIC, Version: 2.3, 21 August 2001, <http://chipinfo.web.cern.ch/chipinfo/docs/TestSpecABCDV23.pdf>.
- [32] J. Bohm et al. SCT LV & HV Power Supply System Project. <http://www-hep2.fzu.cz/Atlas/WorkingGroups/Projects/MSGC.html>.
- [33] P. Bruckman et al. Oxford set barrel module reception test. <http://www.physics.ox.ac.uk/atlas/ModuleProduction/index.shtml>.
- [34] J. Butterworth et al. TIM (TTC interface module) for ATLAS SCT and pixel read out electronics. Prepared for 7th Workshop on Electronics for LHC Experiments, Stockholm, Sweden, 10-14 Sep 2001.
- [35] CAEN. Technical Information Manual, Revision n.11 for Mod. V775 series, 32/16 channel Multievent TDCs. 2005. <http://www.caen.it/nuclear/product.php?mod=V775>.
- [36] CAEN. Technical Information Manual, Revision n.11 for Mod. V792 series, 32/16 channel QDCs. 2005. <http://www.caen.it/nuclear/product.php?mod=V792>.
- [37] F. Campabadal et al. Beam tests of ATLAS SCT silicon strip detector modules. *Nucl. Instrum. Meth.*, A538:384–407, 2005.
- [38] 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.
- [39] D. Campbell et al. The development of the rad-hard CMOS chip for the binary readout of the ATLAS Semiconductor Tracker. Proc. of the Fifth Workshop on Electronics for LHC Experiments, Snowmass, Colorado, USA, 20-24 September 1999, CERN/LHCC/99-33, 29 October 1999, p.152.

- [40] A. Chilingarov. Time Dependence of Interstrip Capacitance. Talk presented at ATLAS SCT week, Valencia, Spain, June, 2002.
- [41] 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.*, A530:293–310, 2004.
- [42] A. Ciocio. Diagnostic analysis of silicon strips detector readout in the ATLAS Semi-Conductor Tracker module production. *Nucl. Instrum. Meth.*, A541:259–266, 2005.
- [43] J. Colas et al. Position resolution and particle identification with the ATLAS EM calorimeter. *Nucl. Instrum. Meth.*, A550:96–115, 2005.
- [44] W3 Consortium. Extensible Markup Language (XML). <http://www.w3.org/XML/>.
- [45] J. Cook et al. COOL Conditions Database for PVSS. <http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/DCS/CONDDDB/CondDBhome.html>.
- [46] LeCroy Corporation. Le Croy 624 Octal Meantimer Manual. <http://www-esd.fnal.gov/esd/catalog/main/lcrynim/624.htm>.
- [47] M. J. Costa and C. Schmitt. ATLAS Inner Detector commissioning with cosmic rays. XI International Workshop on Advanced Computing and Analysis Techniques in Physics Research, Amsterdam, the Netherlands, April 23-27 2007 .
- [48] Creative Electronics Systems S.A. RCB 8047 CORBO, VME Read-out Control Board Manual. 1995. Version 1.0.
- [49] G. Crone et al. ROD Crate DAQ User’s Guide. 2006. ATLAS note, ATLAS-DQ-EN-0020.
- [50] W. Dabrowski. Personal communication.

- [51] W. Dabrowski et al. Design and performance of the ABCD chip for the binary readout of silicon strip detectors in the ATLAS semiconductor tracker. *IEEE Trans. Nucl. Sci.*, 47:1843–1850, 2000.
- [52] J. Dalmau. Reading SCT calibration data from COOL to Athena. Talk given at ID week, Jul 03, 2006, <http://indico.cern.ch/conferenceDisplay.py?confId=a0625>.
- [53] C. J. S. Damerell. Vertex detectors: The state of the art and future prospects. Lectures given at 23rd Annual SLAC Summer Institute on Particle Physics: The Top Quark and the Electroweak Interaction, Stanford, CA, 10-21 Jul 1995.
- [54] S. D’Auria. SCT DCS online databases. Talk given at ID week, July 05, 2006, <http://indico.cern.ch/conferenceDisplay.py?confId=a0625>.
- [55] N. de Barros. Conditions Database Interface: Implementation updates and user requirements. ATLAS Online Database Meeting, CERN, Dec 12th, 2005.
- [56] B. Demirkoz et al. Atmospheric Muon based Trigger Setup for the Integration of the ATLAS Inner Detector. ATLAS note, in preparation.
- [57] M. Dentan et al. DMILL (Durci Mixte sur Isolant Logico-Lineaire): A mixed analog-digital radiation hard technology for high energy physics electronics. *Nucl. Phys. Proc. Suppl.*, 32:530–534, 1998.
- [58] P. J. Dervan. The results of the irradiations of microstrip detectors for the ATLAS tracker (SCT). *Nucl. Instrum. Meth.*, A514:163–166, 2003.
- [59] TDAQ developers. ATLAS Online Software Components. <http://atlas-onlsw.web.cern.ch/Atlas-onlsw/components/components.htm>.
- [60] B. Di Girolamo. The ATLAS Tilecal hadronic calorimeter. *Nucl. Instrum. Meth.*, A453:233–236, 2000.
- [61] B. Dolgoshein. Transition radiation detectors. *Nucl. Instrum. Meth.*, A326:434–469, 1993.

- [62] J. D. Dowell et al. Single event upset studies with the optical links of the ATLAS semiconductor tracker. *Nucl. Instrum. Meth.*, A481:575–584, 2002.
- [63] B. Dowler et al. Performance of the ATLAS hadronic end-cap calorimeter in beam tests. *Nucl. Instrum. Meth.*, A482:94–124, 2002.
- [64] L. M. Drage. The Evaluation of Silicon Microstrip Detectors for the ATLAS Semiconductor Tracker and Supersymmetry Studies at the Large Hadron Collider. Ph.D. Thesis, Trinity College, Cambridge University, 1999.
- [65] T. Dubbs et al. The development of the CAFE-P/CAFE-M bipolar chips for the ATLAS Semiconductor Tracker. Proc. of the Fifth Workshop on Electronics for LHC Experiments, Snowmass, Colorado, USA, 20-24 September 1999, CERN/LHCC/99-33, 29 October 1999, p.123.
- [66] L. Eklund et al. ATLAS SCT Test DAQ Online Documentation. <http://s.web.cern.ch/s/sct/public/sctdaq/sctdaq.html>.
- [67] Ph. Farthouat and P. Gällnö. TTC-VMEbus Interface, TTCvi-MKII. 2000. Version 1.6.
- [68] P. Ferrari. The ATLAS Semi-conductor Tracker System Test Results. *Nucl. Instrum. Meth.*, A530:38–43, 2004.
- [69] T. Flick et al. Optical readout in a multi-module system test for the ATLAS pixel detector. *Nucl. Instrum. Meth.*, A565:85–89, 2006.
- [70] P. Gällnö. ATLAS Local Trigger Processor (LTP) Technical description and users manual. 2005. Version 0.4.
- [71] B. Gallop. ATLAS SCT Readout and Barrel Macro-assembly Testing and Design of MAPS Test Structures. Ph.D. Thesis, Birmingham University, 2005.
- [72] B. Gallop. SCT DAQ Conditions Database. Talk given at ID week, Feb 28, 2007, <http://indico.cern.ch/conferenceDisplay.py?confId=9661>.

- [73] B. Gallop and A. R. Weidberg. SEU Induced Bunch Crossing Counter Errors in the SCT Readout. 2002.
- [74] M. Garcia-Sciveres. The ATLAS Pixel Detector. Prepared for the International Workshop on Semiconductor Pixel Detectors for Particles and X-Rays, Carmel, CA, 9-12 Sep 2002.
- [75] M. Goodrick. Back-Of-Crate Documentation. <http://www.hep.phy.cam.ac.uk/~goodrick/AtBOC>.
- [76] G. Gorfine. Tracking performance of the ATLAS pixel detector in the 2004 Combined Test Beam. *Nucl. Instrum. Meth.*, A565:43–49, 2006.
- [77] A. Gorisek. The design and test of the ATLAS DIAMOND beam conditions monitor. Prepared for 9th ICATPP Conference on Astroparticle, Particle, Space Physics, Detectors and Medical Physics Applications, Villa Erba, Como, Italy, 17-21 Oct 2005.
- [78] A. Grillo. Personal communication.
- [79] R. Hawkings et al. COOL conditions database deployment in ATLAS. <https://uimon.cern.ch/twiki/bin/view/Atlas/CoolATLAS>.
- [80] H. Hayward. Personal communication.
- [81] F. Heinemann. Personal communication.
- [82] J. Hill. ATLAS SCT/Pixel Software Workshop Homepage. <http://sctpixel.home.cern.ch/sctpixel/WorkshopFeb2003/welcome.html>, 6 - 7 Feb 2003, Cavendish Laboratory, Cambridge, U.K.
- [83] J. D. Jackson. *Classical Electrodynamics*. John Wiley & Sons, Inc., New York, third edition, 1998. p. 646-654.
- [84] L. G. Johansen et al. Radiation studies of silicon-microstrip detectors for use in ATLAS and SCT. *IEEE Trans. Nucl. Sci.*, 49:2919–2922, 2002.

- [85] P. Johansson. Personal communication.
- [86] R. Jones et al. The OKS persistent in-memory object manager. *IEEE Trans. Nucl. Sci.*, 45:1958–1964, 1998.
- [87] J. Joseph et al. ATLAS Silicon ROD Documentation. <https://uimon.cern.ch/twiki/bin/view/Main/AtlasSiliconRodGroup>.
- [88] A. Kazarov. Implementation and User’s Guide of DVS: Diagnostics Verification System. <http://rd13doc.cern.ch/Atlas/DaqSoft/components/diagnostics/dvsug/dvsug.html>.
- [89] A. Kazarov et al. ATLAS DAQ/HLT-I Software releases. <http://atddoc.cern.ch/cmt/releases/>.
- [90] T. Kohriki et al. Development of the hybrid structure for the Barrel module of the ATLAS silicon microstrip tracker. *IEEE Trans. Nucl. Sci.*, 49:3278–3283, 2002.
- [91] S. Kolos. Information service user’s guide. <http://atlas-onlsw.web.cern.ch/Atlas-onlsw/components/is/Welcome.htm>.
- [92] J. Lane. Survey of SCT latency budget. http://www.hep.ucl.ac.uk/~jbl/SCT/SCT_latency.html.
- [93] G. Lehmann Miotto. Operations at different activity stages. 2005. ATLAS-TDAQ-CONTROLS-2005-1.
- [94] W. R. Leo. *Techniques for Nuclear and Particle Physics Experiments*. Springer-Verlag, Berlin, First edition, 1987.
- [95] LHC Study Group. *The Large Hadron Collider: Conceptual design*. 1995. CERN-AC-95-05-LHC.
- [96] P. Malecki et al. Multiprocessor system controlling power supply distribution for the ATLAS SCT. 10th Workshop on Electronics for LHC and Future Experiments LECC 2004, Boston, MA, 13-17 Sep 2004, p.210-213.

- [97] M. Mangin-Brinet et al. Electrical test results from ATLAS-SCT end-cap modules. *ATLAS internal note*, ATL-COM-INDET-2003-004.
- [98] G. Moorhead et al. ATLAS SCT Testbeam. http://atlas.web.cern.ch/Atlas/GROUPS/INNER_DETECTOR/SCT/testbeam/.
- [99] M. J. Palmer. Studies of Extra-Dimensional Models with the ATLAS Detector. Ph.D. Thesis, Sidney Sussex College, Cambridge University, 2005.
- [100] S. Peeters. ATLAS SemiConductor Tracker Endcap. Ph.D. Thesis, Amsterdam University, 2003.
- [101] H. Pernegger. Recent test results on the ATLAS SCT detector. *Nucl. Instrum. Meth.*, A501:80–86, 2003.
- [102] H. Pernegger. First test results of a high-speed beam conditions monitor for the ATLAS experiment. *IEEE Trans. Nucl. Sci.*, 52:1590–1594, 2005.
- [103] H. Pernegger. Integration and Test of the ATLAS Semiconductor Tracker. *Nucl. Instrum. Meth.*, A572:108–112, 2007.
- [104] H. Pernegger et al. Baseline SCT Detector Supply Distribution. ATLAS Engineering Specification, ATL-IS-ES-0042, ver. 3.3.
- [105] J. Pinfeld. Plans for the very forward region of atlas: The lucid luminosity monitor. Prepared for 9th ICATPP Conference on Astroparticle, Particle, Space Physics, Detectors and Medical Physics Applications, Villa Erba, Como, Italy, 17-21 Oct 2005.
- [106] J. Poole. LHC Design Report: Beam Parameters and Definitions. 2004. EDMS document 445830, v. 5.
- [107] V. Radeka. Low-Noise Techniques in Detectors. *Ann. Rev. Nucl. Part. Sci.*, 38:217–277, 1988.
- [108] T. Reisaburo et al. SCT Noise Occupancy - Offline summary. Talk given at ID week, Jul 06, 2006, <http://indico.cern.ch/conferenceDisplay.py?confId=a0625>.

- [109] P. Bruckman De Renstrom. Alignment strategy for the ATLAS tracker. Presented at VERTEX 06, Perugia, Italy, Sept. 25-29, 2006, proceedings to appear in NIM A.
- [110] D. Robinson et al. Silicon microstrip detectors for the ATLAS SCT. *Nucl. Instrum. Meth.*, A485:84–88, 2002.
- [111] A. Robson and P. W. Phillips. The Temperature Dependence of the Electrical Performance of ABCD2T readout chips in Prototype ATLAS SCT Barrel Modules. Sept 15, 2000, http://hepwww.rl.ac.uk/atlas-sct/documents/Temperature_Dependencies.pdf.
- [112] H. Sandaker. SCT and TRT Performance from Cosmic Ray Runs. Proceedings of Topical Workshop on Electronics for Particle Physics, Prague, Czech Republic, 3-7 Sept 2007.
- [113] H. Sandaker. The read-out and control system for the ATLAS Semi-conductor tracker. *Nucl. Instrum. Meth.*, A541:96–103, 2005.
- [114] SCT Collaboration. The Data Acquisition and Calibration System for the ATLAS SemiConductor Tracker. In preparation, to be submitted to NIM.
- [115] K. Sedlak. Personal communication.
- [116] R. Shaun et al. Correction factors for ABCD internal calibration capacitors. Correspondence with W. Bialas, http://atlas.web.cern.ch/Atlas/GROUPS/INNER_DETECTOR/SCT/ecmod/abcd_internal_calibration_correc.htm.
- [117] O. Solovianov. Personal communication.
- [118] H. Spieler. *Semiconductor Detector Systems*. Oxford University Press, 2005.
- [119] R. Szczygiel et al. ABCD3T Chip Specification. Version: 1.2, 24 July 2000, http://chipinfo.web.cern.ch/chipinfo/docs/abcd3t_spec_v1.2.pdf.
- [120] S. M. Sze. *Physics of semiconductor devices*. John Wiley and Sons, Second edition, 1981.

- [121] P. K. Teng et al. Radiation hardness and lifetime studies of the VCSELs for the ATLAS Semi-conductor tracker. *Nucl. Instrum. Meth.*, A497:294–304, 2003.
- [122] S. Terada et al. Design and development of a work robot to place ATLAS SCT modules onto barrel cylinders. *Nucl. Instrum. Meth.*, A541:144–149, 2005.
- [123] N. Unno. Personal communication.
- [124] Y. Unno et al. Beamtest of nonirradiated and irradiated ATLAS SCT microstrip modules at KEK. *IEEE Trans. Nucl. Sci.*, 49:1868–1875, 2002.
- [125] A. Valassi et al. Conditions Database Project (Subproject of the POOL Persistence Framework). <http://lcgapp.cern.ch/project/CondDB/>.
- [126] H.C. van der Bij et al. S-LINK: A Prototype of the ATLAS Read-out Link. Prepared for the 4th Workshop on Electronics for LHC Experiments, Rome, September 2005.
- [127] G. Viehhauser et al. Temperature Behaviour and Uniformity of SCT Barrels during Assembly and Reception Testing. 2006. ATLAS note, ATL-COM-INDET-2006-014.
- [128] M. A. Vos et al. A study of the tracking performance of irradiated SCT prototypes using test beam data. 2002. ATLAS note, ATL-INDET-2003-003.
- [129] M. Warren et al. Personal communication.
- [130] M. Warren et al. TTC Interface Module for ATLAS Read-Out Electronics: Final production version based on Xilinx FPGA devices. Prepared for 10th Workshop on Electronics for LHC and future Experiments, Boston, 13 - 17 Sep 2004.
- [131] A. R. Weidberg. Personal communication.
- [132] A. R. Weidberg et al. The Production of ATLAS SCT Optical Links. 11th Workshop on Electronics for LHC and Future Experiments, Heidelberg 12-16 September 2005.

- [133] A. R. Weidberg et al. Updated Dead Time Calculations for the SCT Binary Readout Architecture. ATLAS Internal Note, INDET-NO-124, 1996.
- [134] T. Wengler et al. The ATLAS TDAQ Project: Trigger, DAQ and DCS. <http://atlas.web.cern.ch/Atlas/GROUPS/DAQTRIG/daqtrig.html>.
- [135] D. J. White et al. Radiation hardness studies of the front end ASICs for the optical links of the ATLAS SemiConductor Tracker. *Nucl. Instrum. Meth.*, A457:369–377, 2001.
- [136] M. White. Searching for Supersymmetry with the ATLAS Detector. Ph.D. Thesis, Churchill College, Cambridge University, 2006.
- [137] W-M. Yao et al. Review of Particle Physics. *J. Phys. G: Nucl. Part. Phys.*, 33:1+, 2006.