

Design, Development, Characterisation and Operation of ATLAS ITk Strip Staves



Peter William Phillips

Linacre College, Oxford

Thesis submitted in fulfilment of the requirements for the degree of Doctor
of Philosophy at the University of Oxford

Hilary Term, 2020

Abstract

ATLAS (A Toroidal LHC ApparatuS) is a general purpose particle physics experiment designed to exploit the physics potential of the Large Hadron Collider. Two further upgrades are planned to increase the luminosity of the LHC, the latter of which (“Phase 2” scheduled for 2023/2026) will give rise to the High Luminosity Large Hadron Collider (HL-LHC). With peak luminosity increased by a factor of 3.75 beyond today’s values, a new Inner Tracker (ITk) will be required for ATLAS. An all-silicon design is chosen with a mixture of Pixel and Strip layers.

This thesis documents the electronic system design of the ITk Strip Staves to be used in the barrel region, and presents results from tests of front-end chips, modules and staves built using prototype chipsets made in the same 130 nm technology as will be used in the experiment. Most of the results come from systems using the first of these chipsets, HCC130 / ABC130, but the final chapter documents wafer screening and single chip results using a later prototype chipset, HCCStar / ABCStar, which has a revised front-end amplifier design and an updated data transfer topology which makes greater use of the available bandwidth. Although beyond the scope of this thesis, the first staves built with the latter chipset have recently been completed. The HCCStar and ABCStar designs are presently being reworked to give improved tolerance to Single Event Effects (SEE), leading toward the assembly of the first pre-production staves later this year. The production phase is due to start next year with commissioning of the completed system, in the ATLAS cavern, scheduled to take place in 2025.

Acknowledgements

I would like to thank the Particle Physics Department of the STFC Rutherford Appleton Laboratory for agreeing to fund my studies, especially the ATLAS ITk Strip group comprising Stephen Haywood, Steve McMahon, Jens Dopke, Craig Sawyer and Bruce Gallop for their encouragement, support and proof reading!

Having been active in the field for over 25 years in principle I should thank rather a lot of people, but there isn't space to list everybody. I must first acknowledge my long-term colleagues in the ITk Strip Data Acquisition team, Bruce Gallop (again) and Matt Warren (UCL): without their continued collaboration this thesis (and many others) would be very short indeed. My forays into ASIC testing have drawn heavily on the experiences of Francis Anghinolfi (formerly of CERN, now retired), Jan Kaplon (CERN), Paul Keener, Jeff Nandoy and Mitch Newcomer (Penn) and Jaya John John (Oxford). In recent years I have visited Brookhaven National Laboratory several times to work with the US stave test team: I would like to thank Dave Lynn, Stefania Stucci and colleagues for their hospitality and collaboration, especially Stefania! I must also thank my fellow Oxford student, Zhiying Li, for her help obtaining many of the more recent ASIC and stave results at RAL.

My thesis advisor Tony Weidberg has been a friend and colleague for many years: his guidance has been very much appreciated. I also wish to record my thanks to those who helped guide the early stages of my career: Richard Apsimon, Martin Morrissey and Mike Tyndel (formerly of RAL, now retired), and Gareth Moorhead (formerly of The University of Melbourne, now at CSIRO). Without their help and guidance, I very much doubt that this thesis would ever have been written.

With thanks to my family, Pauline and Eleanor, for their understanding.

Contents

Contents	vi
List of Figures	xii
List of Tables	xiv
Contribution	xv
Glossary	xviii
1 Introduction	1
1.1 The Large Hadron Collider	1
1.2 The ATLAS Experiment	4
1.3 The High Luminosity Large Hadron Collider, HL-LHC	7
2 Silicon Strip Detector Systems	10
2.1 An introduction to Silicon Strip Detector Systems	10
2.1.1 A Generic Detector System	10
2.1.2 Silicon Detector Systems	12
2.1.3 Radiation Damage in Silicon	15
2.1.4 Noise in Silicon Detector Systems	17
2.2 Chapter Summary and Conclusions	20
3 System Architecture, Test Methodology and System Metrics	22
3.1 System Architecture	22

3.1.1	Layout of the ITk Strip Detector	22
3.1.2	Electronics Architecture	25
3.1.3	Powering Strategies	26
3.2	Module Architecture	31
3.2.1	ATLAS Trigger Architecture	32
3.2.2	The Front End Chip ABC130	34
3.2.3	The Hybrid Controller Chip HCC130	36
3.2.4	The Power Board	37
3.3	Test Methodology and System Metrics	38
3.3.1	Threshold Scans and the Complementary Error Function	38
3.3.2	Calibration Timing	39
3.3.3	Three Point Gain	41
3.3.4	Trimming	42
3.3.5	The Response Curve	44
3.3.6	Noise Occupancy as a function of Threshold	45
3.3.7	Timewalk	46
3.3.8	Double Trigger Noise Occupancy	47
3.3.9	Pedestal Trim	49
3.4	Chapter Summary and Conclusions	50
4	Evaluation of the ABC130/HCC130 chipset	51
4.1	Test System Design for ABC130/HCC130 chipset	51
4.2	First Batch of ABC130 wafers	54
4.3	Second Batch of ABC130 wafers	55
4.4	Third Batch of ABC130 wafers	57
4.4.1	First 14 wafers (Batch 3A)	57
4.4.2	Remaining 10 wafers (Batch 3B)	58
4.5	Fourth Batch of ABC130 wafers	63
4.6	Yield across all batches	64
4.7	Evaluation of HCC130	66

4.7.1	HCC130 Irradiation Programme (UK)	67
4.8	Chapter Summary and Conclusions	68
5	ABC130 Noise Performance	70
5.1	Devices Tested	70
5.1.1	Prototype ASIC	70
5.1.2	Single Chip Boards	71
5.1.3	Hybrids and DAQloads	72
5.1.4	First Short-Strip Module	74
5.2	Initial DAC tuning with nStrobed = 256	74
5.3	Revised DAC tuning and reduction of nStrobed	76
5.4	Chapter Summary and Conclusions	78
6	Non-Segmented Stave with ABC130/HCC130 Chipset	79
6.1	Stave Specifications and Modifications	79
6.1.1	Bus Tapes	79
6.1.2	End of Substructure Cards	81
6.1.3	Power Supplies and DAQ	83
6.2	Modules and Modifications	85
6.3	Results	91
6.3.1	Comparison of noise on and off the stave	91
6.3.2	Noise as a function of Sensor Bias	91
6.3.3	Gain and Noise as a function of nStrobed	95
6.3.4	Noise as a function of Time	97
6.3.5	Estimation of Input Noise by Noise Occupancy method	98
6.3.6	Use of module SS2 as an Aggressor	99
6.3.7	Powering the stave with a prototype "PP2" DC-DC converter	100
6.3.8	Power disturbances at startup / shutdown	102
6.4	Chapter Summary and Conclusions	103

7	Segmented Staves with ABC130/HCC130 Chipset	104
7.1	Stave Specifications	104
7.1.1	Bus Tapes	104
7.1.2	End of Substructure Cards	106
7.2	UK Single Sided Stave	109
7.2.1	Build and Test Information	109
7.2.2	Results	111
7.3	UK Double Sided Stave	116
7.3.1	Build and Test Information	116
7.3.2	Results: Master Side	118
7.3.3	Results: Slave Side	121
7.4	US Single Sided Stave	129
7.4.1	Build and Test Information	129
7.4.2	Results	130
7.5	Chapter Summary and Conclusions	132
8	The Three Stave Test	134
8.1	Introduction	134
8.1.1	The Support Structure and Stave Installation	135
8.1.2	Services	137
8.1.3	Data Acquisition and Debugging	141
8.1.4	Operating Conditions and Standard Tests	143
8.2	Position 1: Separation of supports 237 mm, 12°	145
8.2.1	“Like for Like” comparison with Original Single Stave Results	145
8.2.2	Effect of Optimised Burst Conditions	147
8.2.3	Effect of running all staves together	147
8.2.4	Investigation of additional noise of modules near EoS	152
8.3	Position 2: Separation of supports 90 mm, 12°	154
8.3.1	Comparison against 12°, Separation of support centres 237 mm	157
8.4	Position 3: Separation of supports 80 mm, 20°	157

8.4.1	Comparison against 12°, Separation of support centres 237 mm	160
8.5	Comparison of Noise Occupancy results for selected modules	163
8.6	Chapter Summary and Conclusions	168
9	Evaluation of the Star Chipset	170
9.1	Test System Requirements for the Star Chipset	170
9.1.1	The FMC-1701 Driver Board	173
9.1.2	StarChip Single Chip Boards	176
9.1.3	StarChip Probe Cards	178
9.2	ABCStar Single Chip Board Noise Results	180
9.3	ABCStar Wafer Probing	181
9.3.1	Development of the ABCStar Test Sequence	181
9.3.2	Offline Analysis	188
9.3.3	Results and Discussion	199
9.4	ABCStar Single Chip Board Supplemental Results	202
9.4.1	Optimisation of Measurement Conditions	203
9.4.2	Operation across a range of bias conditions	205
9.4.3	A more detailed characterisation of ABCStar’s internal ADC	207
9.4.4	Shmoo Plots	209
9.4.5	Calibration of ABCStar’s internal Temperature Diode	211
9.5	HCCStar Wafer Probing	212
9.5.1	Development of the HCCStar Test Sequence	212
9.5.2	Results and Discussion	214
9.6	Chapter Summary and Conclusions	217
10	Conclusions and Ongoing Work	219
A	Cuts for Automated Analysis of ABC130 Wafer Probing Results	222
B	Reduction of noise generated by the SemiProbe Probe Station	232
	References	234

List of Figures

1.1	The CERN Accelerator Complex	2
1.2	The ATLAS Detector	4
1.3	Particle Detection in ATLAS	6
1.4	The LHC Roadmap	7
2.1	Generic Detector System	11
2.2	Electrode Segmentation	11
2.3	A simple p-n junction	12
2.4	A Detector Front-End Circuit and its Equivalent Noise Schematic	18
3.1	ITk Layout	23
3.2	ITk Strip Stave and Petal Concepts	24
3.3	ITk Strip Stave Cross Section	25
3.4	ITk Strip Electronics Overview	25
3.5	Four Possible LV Powering Schemes	27
3.6	Topology of the DC-DC Synchronous Buck Converter	29
3.7	The Short Strip Module Concept	31
3.8	Short Strip Module SS1	32
3.9	Block Diagram of ABC130 ASIC	34
3.10	Block Diagram of HCC130 ASIC	35
3.11	Block Diagram of AMAC ASIC	37
3.12	Simulation of noise and signal distributions	38
3.13	An example of a Strobe Delay Test	40

3.14	Typical Three Point Gain	42
3.15	Trim Plots	43
3.16	Typical Response Curve	45
3.17	Typical Noise Occupancy	46
3.18	Typical ABCStar Timewalk Plots	47
3.19	An example of a Double Trigger Noise Test	48
3.20	An example of a Pedestal Trim Test	50
4.1	ABC130 Single Chip Board and Driver	52
4.2	ABC130 Probe Card and Probed Die	53
4.3	3D Electron Micrograph showing correction of ABC130 by Focussed Ion Beam	54
4.4	The distributions of the offset of the IR8B DAC before (left) and after (right) modifications to sense the true ground potential of the ABC130 chip.	59
4.5	Typical Batch 3B LDO Characterisation	60
4.6	Batch 3B: The distributions of tuned DAC values and the resulting monitor voltages for parameter VREF.	61
4.7	Batch 3B: Comparison of mean input noise values under nominal (left) and tuned (right) conditions.	62
4.8	HCC130 Single Chip Board	66
4.9	HCC130 Current Increase with TID	68
5.1	ABC130 Barrel DAQload	73
5.2	Noise as $f(\text{capacitance})$ with initial ABC130 DAC tunings. For all ABC130 measurements, nStrobed, the number of channels into which charge is simul- taneously injected, is 256.	75
5.3	Gain Increase as $f(\text{nStrobed})$	77
5.4	Noise as $f(\text{capacitance})$ with revised ABC130 DAC tunings. For the ABC130 measurements made using discrete capacitors, nStrobed, the number of chan- nels into which charge is simultaneously injected, is 1; for all other ABC130 measurements nStrobed is 16.	77

6.1	EoS Board with modifications	82
6.2	Non-Segmented ABC130 Stave	86
6.3	Schematic of the HV-MUX circuit and its oscillator waveform	87
6.4	Schematic of the modified HV-MUX circuit and its oscillator waveform	89
6.5	Image of a modified HV-MUX circuit	89
6.6	Noise occupancy of SS5, Hybrid 1, Column 0.	90
6.7	Input Noise by Three Point Gain method: comparison of results on and off the stave	92
6.8	Sensor Capacitance and Input Noise Predictions	94
6.9	H5 Breakdown	95
6.10	Gain and Noise as a function of nStrobed	96
6.11	RC vs time	97
6.12	RC vs NO	98
6.13	Insertion of aggressor module SS2 beneath the stave	100
6.14	Input Noise by Response Curve method: comparison of results with aggressor SS2 on and off	100
6.15	Input Noise by Three Point Gain method: comparison of results with bench LV supply and DC-DC board DC1073A	101
6.16	Current and Voltage delivered to the stave during power on and power off . .	102
7.1	The Segmented Bus Tape	104
7.2	UK Single Sided Stave: Master EoS	106
7.3	CCR Bus Termination	109
7.4	UK Single Sided Stave: IV Curve	111
7.5	UK Single Sided Stave: Gain on and off Stave	112
7.6	UK Single Sided Stave: Input Noise on and off Stave	113
7.7	UK Single Sided Stave: Gain Ratio on/off the Stave	114
7.8	UK Single Sided Stave: Input Noise Ratio on/off the Stave	115
7.9	Data Transfer Architecture within the ABC130 Barrel Hybrid	119
7.10	UK Double Sided Stave, Master Side: Gain on and off Stave	120

7.11 UK Double Sided Stave, Master Side: Input Noise on and off Stave	121
7.12 UK Double Sided Stave, Master Side: Input Noise ratio (on-stave / off-stave) for inner columns	122
7.13 Threshold Scans for the Cambridge Modules	124
7.14 UK Double Sided Stave, Slave Side: Gain on and off Stave	125
7.15 UK Double Sided Stave, Slave Side: Input Noise on and off Stave	126
7.16 UK Double Sided Stave, Slave Side: Gain on and off Stave	128
7.17 UK Double Sided Stave, Slave Side: Input Noise on and off Stave	129
7.18 US Single Sided Stave: Gain on and off Stave	131
7.19 US Single Sided Stave: Input Noise on and off Stave	132
8.1 Barrel Structures	135
8.2 Stave Clearances	136
8.3 Three Stave Support Structure	137
8.4 Services for Two Staves	140
8.5 Position 1: Images	145
8.6 Position 1: Comparison of Gain against Single Stave Results (Standard Burst Conditions)	148
8.7 Position 1: Comparison of Input Noise against Single Stave Results (Standard Burst Conditions)	149
8.8 Position 1: Gain (Optimised Burst Conditions) divided by Gain (Standard Burst Conditions)	150
8.9 Position 1: Input Noise (Optimised Burst Conditions) divided by Input Noise (Standard Burst Conditions)	151
8.10 Position 1: Gain (sides together) divided by Gain (sides separately)	152
8.11 Position 1: Gain (sides together) divided by Gain (sides separately)	153
8.12 Position 1: Gain running one module on each of the three stave sides	155
8.13 Position 1: Input Noise running one module on each of the three stave sides	156
8.14 Position 2: Images	157
8.15 Position 2: Gain divided by reference data measured in Position 1	158

8.16	Position 2: Input Noise divided by reference data measured in Position 1 . . .	159
8.17	Position 3: Images	160
8.18	Position 3: Input Noise divided by reference data measured in Position 1 . . .	161
8.19	Position 3: Input Noise divided by reference data measured in Position 1 . . .	162
8.20	Noise Occupancy for Module 12 of the Slave Side of the UK Double Sided Stave	164
8.21	Noise Occupancy for Module 12 of the US Single Sided Stave	165
8.22	Noise Occupancy for Module 11 of the Slave Side of the UK Double Sided Stave	166
8.23	Noise Occupancy for Module 11 of the US Single Sided Stave	167
9.1	Data Transfer Architecture of the Barrel Hybrid	171
9.2	Outline of the Test System for Star Chips	173
9.3	FMC-1701	174
9.4	Single Chip Boards for the Star Chips	177
9.5	Probe Cards for the Star Chips	179
9.6	ABCStar Single Chip Noise Measurements	180
9.7	Good and Bad LDO Fits	189
9.8	Issues with the ABCStar's ADC data which needed to be addressed offline .	190
9.9	Distribution of optimised bias DAC settings	193
9.10	Distributions of optimised LDO settings	193
9.11	The distribution and spatial variation of Mean Channel Gain	195
9.12	The distribution and spatial variation of Mean Channel Input Noise	195
9.13	Gain vs Channel and the accompanying distribution for Chip VPCVRRH_046	196
9.14	Gain vs Channel and the accompanying distribution for Chip VXCQV1H_005	196
9.15	The distribution of χ^2 of a Gaussian fit to each chip's Channel Gain values .	197
9.16	Gain vs Channel and the accompanying distribution for Chip VICVPYH_231	197
9.17	Gain vs Channel and the accompanying distribution for Chip VICVPYH_174	198
9.18	The distribution and spatial variation of "Gain Slope"	198
9.19	Wafer Map showing location of rejected ABCStar Die	202
9.20	Gain and Input Noise as a function of Strobe Delay	203
9.21	Gain and Input Noise as a function of nStrobed	204

9.22	Gain and Input Noise as a function of nTrigs	205
9.23	Gain and Noise as f(VDDA, BVREF, BIREF)	206
9.24	Characterisation of ADC on all ranges, chip VJCVPXH_010	207
9.25	Characterisation of ADC on “best” range, chip VJCVPXH_010	208
9.26	Shmoo plots for irradiated and unirradiated ABCStar Chips	210
9.27	Wafer Map showing location of rejected HCCStar Die	216
A.1	Powering ABC130 during ASIC testing	223
A.2	Example Batch 3A LDO Characteristic (single die)	226
A.3	Example Batch 3A Calibration and Threshold DAC Characteristic (single die)	227
A.4	Slope of Threshold and Calibration DACs for all die of wafer AKQ8E9H . .	227
A.5	Example Batch 3A SLVS DAC Characteristic	228
A.6	Typical ABC130 Gain and Noise Plots	229
A.7	Mean Gain and Mean Input Noise for all die of wafer AKQ8E9H	230
A.8	Typical ABC130 Wafer Pickmap	231
B.1	Voltage Noise measured on the SemiProbe PS4L chuck	233

List of Tables

3.1	Main layout parameters for the ITk Strip Barrel	24
4.1	ABC130 Probing Yield, Batch 2	56
4.2	ABC130 Probing Yield, Batch 3A	58
4.3	Batch 3B: The distributions of tuned DAC values and the corresponding mon- itor voltages.	61
4.4	Batch 3B: Cuts on optimised bias DAC and Strobe Delay settings.	62
4.5	Batch 3B: Additional cuts on DAC settings.	63
4.6	ABC130 Probing Yield, Batch 3B	63
4.7	ABC130 Probing Yield, Batch 4B	65
4.8	ABC130 Probing Yield across All Batches	66
6.1	Non-Segmented ABC130 Stave Module Selections	85
7.1	UK SS Stave Module Selections	110
7.2	UK DS Stave Module Selections	117
7.3	UK DS Stave Slave Side Backplane Contact Resistance	127
7.4	US SS Stave Module Selections	130
8.1	Barrel Layer Parameters	136
8.2	Stave Low Voltages and Currents	143
8.3	Excluded ABC130 Chips	143
8.4	Test Sequence	144
9.1	StarChip Power Test Configurations	182

9.2	StarChip Fuse Registers	183
9.3	ABCStar DAC Targets	184
9.4	ABCStar Summary File Content	189
9.5	ABCStar Power Test Cuts	191
9.6	ABCStar DAC Cuts	191
9.7	ABCStar ADC Cuts	192
9.8	ABCStar Tuned DAC Cuts	193
9.9	ABCStar Analogue Cuts	194
9.10	ABCStar Probing Yield, Batch 1	200
9.11	ABCStar Probing Yield, Batch 2	200
9.12	ABCStar Defect Statistics	200
9.13	ABCStar Analogue Defect Statistics	201
9.14	ABCStar Low Yield Sites	201
9.15	ABCStar Temperature Diode	212
9.16	HCCStar Probing Yield, Batch 1	214
9.17	HCCStar Defect Statistics	215
9.18	HCCStar Low Yield Sites	215
A.1	Table of Cuts on ABC130 Constant Voltage Power Measurements	225
A.2	Table of Cuts on ABC130 Constant Current Power Measurements	225
A.3	Table of Cuts on ABC130 MUX (DAC) Parameters	228
A.4	Table of Cuts on ABC130 SLVS driver measurements	229
A.5	Table of Cuts on ABC130 Analogue (Gain) Parameters	230

Author's Contribution

As a member of the Particle Physics Department at the Rutherford Appleton Laboratory I have been working with Silicon Microstrip Detectors since 1994, however the time span of this thesis begins with the delivery of the ABC130 chip in 2014, the first ASIC in the 130 nm CMOS technology which will be used in the ITk Strip detector. I was responsible for the conceptual design of the driver board and single chip PCB used in the early testing programme, subsequently completed by a professional engineer, and for the design and layout of the probe card PCB. The first version of the chip gave no output due to a bug regarding the polarity of the transceiver's output enable line, and I worked with the ASIC designers and UK industry to have a small number of chips partially corrected by means of Focussed Ion Beam (FIB) surgery. The very first tests of functional ABC130 chips and the readout system could then be made by the core ITSDAQ team, including myself.

Additional wafers were ordered to a corrected design. Once we had established that both the chip and readout system were fully functional, I assembled the wafer test sequence, writing software to accommodate the non-standard mapping of our multi-project wafer (MPW). I proceeded to work with a summer student to test the five remaining wafers of this batch using the MicroManipulator probe station at RAL. Before the next batch of wafers arrived some months later the control electronics of this machine had failed: I successfully re-engineered the system using Commercial Off The Shelf (COTS) components and a custom PCB of my own design. As wafer testing proceeded I wrote analysis tools within the ROOT environment to apply offline cuts to the test data in order to generate pickmaps for the dicing vendor. In all 72 wafers were delivered, each of which was tested by me using the MicroManipulator probe station.

I was also responsible for the conceptual design of the HCC single chip card, developed for use with the same readout system. After initial tests of the chip and software had been completed by the ITSDAQ team, I developed a test sequence which was subsequently used during proton irradiations at Birmingham and X-ray irradiations at RAL. Bringing the ABC130 and HCC chips together, I tested the first readout hybrid to be built with the complete HCC/ABC130 chipset, and presented these results to the community in February 2015. As part of this work I developed the "chipID" algorithm, which compares the pedestals measured on a hybrid to wafer test data, such that the identities of the ASICs fitted to each hybrid may be confirmed. Early modules and "DAQ loads" (hybrids bonded to one or more miniature strip sensors) were taken to test beams at DESY, read out using low capacitance network cables by means of adaptors which I made, their events being assembled for transfer to EUDAQ by an event builder which I co-authored within the ITSDAQ framework.

Once the community was confident that the noise of the ABC130 front end had been understood we proceeded to populate a stave, built with non-segmented bus tapes, with "second generation" ABC130 modules using a powerboard (DC-DC converter and HV GaN-FET switch) controlled by a commercial one-wire chip. As I was responsible for the stave testing, I needed to ensure that we had a system in place which could operate it. To that end I designed three small PCBs to facilitate the tests. I made further modifications to the EoS board itself and, for the modules fitted with HV GaN-FET switches, I made a modification to each completed module such that the switch could be operated by the on-board one wire device. I later made modifications to apply a RC filter to the multiplexer drive oscillator on each board to eliminate fast edges of this signal which had been seen to feedback into the sensor. I was also the sole author of the software used to control and monitor the stave LV and HV supplies and temperatures such that it may operate unattended.

Population of the first single-side short-strip stave built with segmented bus tapes began at RAL in March 2018, and population of a double-sided short-strip stave at RAL began in November 2018. I was responsible for testing both of these staves. Together with a third, single sided stave built by our US collaborators at Brookhaven National Laboratory, these staves were used to form a "three stave test" operated through autumn 2019. I was

responsible for providing the services (cables, pipes, PCB adaptors) and control software needed to operate these staves, and for the definition and execution of the test schedule. I also visited BNL twice during the course of 2018 to ensure the compatibility of test procedures and results between the two stave assembly sites.

Meanwhile the Preliminary Design Review (PDR) of the production chipset, comprising ABCStar and HCCStar, was held in January 2018 following on from the PDR of the AMAC2 chip which was held in November 2017. I was on the review committee for all three chips. My main contribution here is towards the testing of the chips, notably the ABCStar. I designed a new FPGA Mezzanine Card (FMC) to provide the interface between a commercial FPGA board and the ABCStar/HCCStar chips together with a new ABCStar single chip PCB and probe card PCB, taking all three designs through to manufacture. I also collaborated with engineers at Penn as they completed designs for a single chip board and probe card for HCCStar. Results from these chips form the final chapter of this thesis, including data from 14 wafers of ABCStar die tested at RAL by me using software to which I contributed the top level sequence, wafer map, DAC and ADC characterisation code and the FMC driver class. I also led the offline analysis of this data to determine the final pickmaps for use by the dicing vendor, and transformed the data into a format suitable to be uploaded to the ITk Production Database. The first 10 of these wafers also had their HCCStar die tested at RAL largely by myself, however the software for this activity was led by Penn.

Glossary

ABC130 ATLAS Binary Chip in 130 nm technology

ABCN25 ATLAS Binary Chip Next generation in 250 nm technology

ABCStar ATLAS Binary Chip with Star topology, 130 nm technology

AMAC Autonomous Monitor And Control ASIC, 130 nm technology

AMAC1 Version of AMAC which communicates over the single ended I2C bus

AMAC2 Version of AMAC which communicates over a differential bus using a custom “morse code” protocol named “Endeavour”

ASIC Application Specific Integrated Circuit

ATLYS a commercial FPGA module supplied by Digilent, using a Xilinx Spartan-6 device

BCC Buffer Control Chip, a minimal hybrid controller ASIC in 250 nm technology

BCR Bunch Counter Reset, a signal which resets the Bunch Crossing Counter in the ABC130 or ABCStar chips.

BCO Bunch Crossing clock, a 40.08 MHz clock signal representative of the LHC beam structure to which all detector elements much be synchronised.

BER Bit Error Ratio, the number of bit errors divided by the total number of transferred bits.

bPOL12V Buck Point of Load converter with 12 V input, the ASIC used in the on-module DC-DC converter

CCR Clock Control Reject, an ITk Strip specific acronym for the trio of control signals needed by a hybrid, which carry the TTC information.

CLPS CERN Low Power Signalling, a CERN standard for differential signalling between low voltage ASICs

DESY Deutsches Elektronen-Synchrotron, a Research Centre of the Helmholtz Association

EoS End of Substructure card, the PCB at the end of a stave or petal upon which the VTRx transceiver and lpGBT chips are located

ENC Equivalent Noise Charge: the number of electrons which would need to be injected into the sensor to create a signal equal to the noise.

FELIX Front-End Link eXchange, the data acquisition system to be used in future upgrades of the ATLAS Experiment.

FDR Final Design Review

FEC Forward Error Correction

FMC FPGA Mezzanine Card, a VITA 57.1 standard which defines an I/O mezzanine module used with FPGA “carrier” cards

FPGA Field-Programmable Gate Array, an integrated circuit containing logic blocks which may be interconnected by the application developer to implement the desired functionality.

GBTx GigaBit Transceiver chip

GBT-SCA GigBit Transceiver Slow Control Adaptor (for use with GBTx)

GENESYS 2 a commercial FPGA carrier card supplied by Digilent, using a Xilinx Kintex-7 device

GLIB Gigabit Link Interface Board, a custom FPGA module made by CERN, using a Xilinx Virtex-6 device

HCC130 Hybrid Controller Chip in 130 nm technology

HCCStar Hybrid Controller Chip with Star topology, 130 nm technology

HSIO High Speed Input Output board, a custom FPGA module made by SLAC, using a Xilinx Virtex-4 device

ITSDAQ ITk Strip Data Acquisition firmware and software, the system used to test our chips, hybrids, modules and staves. Delivery is largely a UK responsibility.

lpGBTx Low Power GigaBit Transceiver chip

LCB An acronym formed of the acronyms L0A/CMD/BCR: a protocol which transfers the L0A trigger, fast commands, slow commands and bunch counter reset signals to the HCCStar and ABCStar chips.

LDO Low Drop Out regulator, a linear voltage regulator device.

linPOL12V Linear Point of Load regulator 12 V input, a linear regulator used on the powerboard to power AMAC independent of the DC-DC converter.

L0A The zeroth level trigger signal in the ITk Strip detector which selects an event from transfer from the input FIFO to the event buffer. The data does not leave the ABC130 or ABCStar chip.

L1A The first level trigger signal in ATLAS which selects an event for transfer off-detector.

LVDS Low Voltage Differential Signalling, a well-defined commercial standard for differential signalling (TIA-644). The common mode level is 1.2 V.

LS Long Strip. Refers to an ITk Strip Barrel module with strip length 5.0 cm approx.

Mbit/S Megabit per second, unit of data transfer rate equal to 1,000,000 bits per second

M-LVDS Multipoint Low Voltage Differential Signalling, a high current variant of the LVDS driver tuned to drive multiple receivers (TIA-899).

NEXYS Video a commercial FPGA carrier card supplied by Digilent, using a Xilinx Kintex-7 device

NTC Negative Temperature Coefficient thermistors (temperature sensors)

PDR Preliminary Design Review

PRR Production Readiness Review

TID Total Ionising Dose

TTC Trigger, Timing and Control.

SFP Small Form-factor Pluggable transceiver, a compact hot-pluggable network interface module.

SFP+ Small Form-factor Pluggable transceiver Plus, a compact hot-pluggable network interface module rated at 10 Gbit/s.

SMD Surface Mount Device

SLVS Scalable Low Voltage Signalling, a commercial standard for differential signalling in low voltage devices (JESD8-13). Different variants set the common mode level as 200 mV, or half the voltage rail of the devices.

SS Short Strip. Refers to an ITk Strip Barrel module with strip length 2.5 cm approx.

UDP User Datagram Protocol. One of the simpler members of the Internet Protocol Suite which omits much of the error checking and retransmission associated with other protocols. Often used in real time systems such as data acquisition.

VL Versatile Link - optical links for GBTx

VL+ Versatile Link plus - optical links for lpGBT

VLDB Versatile Link Demonstration Board, a demonstration module with GBTx on board, made by CERN.

Chapter 1

Introduction

This thesis documents the electronic system design of the ITk Strip Staves, with a special focus on the performance of module and stave prototypes built using a prototype chipset in 130 nm CMOS technology. This chapter introduces the Large Hadron Collider (LHC), the ATLAS experiment, and the path towards operation at higher luminosity.

1.1 The Large Hadron Collider

The Large Hadron Collider (LHC) [1] is the largest and highest energy particle collider ever built. It occupies a tunnel of 27 km circumference, running under the Swiss-French border, which was originally built for the Large Electron Positron collider (LEP) [2]. Preparatory work for the LHC began in 1998, two years before LEP was decommissioned, with installation complete by 2008.

An image of the CERN accelerator complex, of which LHC is the biggest part, is shown in Figure 1.1. Most of the time the LHC collides two beams of protons. These protons come from a bottle of hydrogen gas located in the Linear Accelerator 2 (Linac 2) [4] facility¹: an electric field is used to remove the electrons from the hydrogen atoms and the resulting protons are accelerated through a series of radio frequency cavities until they reach an energy of 50 MeV. The proton beam then passes to the Proton Synchrotron Booster (PSB) [5] where

¹From run 3 onwards, Linac 2 will be replaced as the source of the protons by Linac 4 which will accelerate them to 160MeV before transfer to the Proton Synchrotron Booster.




Figure 1.1: A diagram of the CERN Accelerator Complex, to illustrate how the proton beam passes from one machine to the next before reaching the Large Hadron Collider (LHC). The locations of the experiments are also shown. (Reproduced from Reference [3].)

they are accelerated to 1.4 GeV. Next the protons are transferred to the Proton Synchrotron (PS) [6] for acceleration to 25 GeV, before the Super Proton Synchrotron (SPS) [7] accelerates them to the LHC injection energy of 450 GeV. The LHC itself comprises two separate proton beams running in opposing directions: each has its own beam pipe within the common ring. Once each beam has been filled with protons, the beams are accelerated to the desired energy, presently 6.5 TeV per beam (or 13 TeV centre of mass energy) which takes approximately 20 minutes. Within the experimental regions, where the beams share a common beam pipe, the beams may now be brought into collision by means of additional steering magnets. Typically beams would then be maintained in the LHC for 20 hours during which period the experiments may record physics data.

An alternative mode of operation of the LHC collides heavy ions, usually lead. In this case the ions begin their journey in the Linear Accelerator 3 (Linac 3) [8] facility, being accelerated to 4.2 MeV before transfer to the Low Energy Ion Ring (LEIR) [9] where they are accelerated to 72 MeV. After this they are transferred to the Proton Synchrotron. From

this point their route into the LHC is the same as for the protons.

A total of seven experiments are installed at the intersection points of the LHC beams, although only four of these are shown in Figure 1.1. The experiments are:

- A Large Ion Collider Experiment (ALICE) [10]. This is a specialist detector located at Point 2 which studies p-Pb and Pb-Pb collisions to determine the physics of strongly interacting matter at extreme energy densities (the quark-gluon plasma).
- A Toroidal LHC ApparatuS (ATLAS) [11], a general purpose experiment located at Point 1.
- Compact Muon Solenoid (CMS) [12], a general purpose experiment located at Point 5.
- Large Hadron Collider beauty (LHCb) [13]. Located in the Point 8 cavern this specialised b-physics experiment measured the parameters of CP violation in the interactions of b-hadrons (heavy particles containing a bottom quark).
- Large Hadron Collider forward (LHCf) [14]. This detector shares the cavern at point 1 with ATLAS. It is a special purpose experiment built for astroparticle physics which uses particles thrown into the far forward region by LHC collisions as a means to simulate cosmic rays under laboratory conditions.
- MOнопole and Exotics Detector At the LHC (MoEDAL) [15]. A specialist detector which seeks to find the magnetic monopole or highly ionizing stable or pseduo stable massive particles, this detector shares the cavern at Point 8 with LHCb.
- TOTal Elastic and diffractive cross section Measurement (TOTEM) [16]. This detector shares the cavern at point 5 with CMS and targets accurate measurement of the proton-proton interaction cross-section and detailed studies of proton structure.

Design objectives for the LHC include, in the case of proton-proton running, a maximum beam energy of 7 TeV (14 TeV centre of mass energy) and a peak luminosity of

$10^{34} \text{cm}^{-2} \text{s}^{-1}$. Whilst the maximum energy has yet to be reached, as will be shown later the design luminosity has already been surpassed.

1.2 The ATLAS Experiment




Figure 1.2: The ATLAS Detector, labelled to show the various sub-detectors. The ATLAS Detector weighs about 7000 tonnes, is 46 metres long and 25 metres in diameter. (Reproduced from Reference [17].)

ATLAS (A Toroidal LHC Apparatus) is a general purpose particle physics experiment designed to exploit the physics potential of the Large Hadron Collider. Born of a merger of the EAGLE (Experiment for Accurate Gamma, Lepton and Energy measurements) [18] and ASCOT (Apparatus with Super Conducting Toroids) [19] collaborations, ATLAS was originally proposed in 1994. The completed detector is shown in Figure 1.2. The detector is 44 m long, 25 m high and weighs approximately 7000 tonnes. Construction was completed in 2008.

The Inner Detector (ID) [20][21] is located within a 2 Tesla solenoidal magnet and is designed to measure the direction, momentum and charge of electrically charged particles.

The ID comprises the Pixel Detector [22], the SemiConductor (strip) Tracker (SCT), and the Transition Radiation Tracker (TRT). The pixel and strip systems are both silicon technologies whereas the TRT tracks particles by means of the ionisation of gas in a series of straw tubes. The tracking system provides the information necessary to identify the location of the primary vertex of each proton-proton interaction and the location of secondary vertices, points where particles generated in the proton-proton interaction decay into series of less energetic particles and photons.

Outside the solenoid lies the calorimetry system. Surrounding the Inner Detector (ID) is the electromagnetic calorimeter, also known as the Liquid Argon (LAr) calorimeter [23], designed to absorb and record the energy of electrons and photons leaving the tracker volume. The electrons and photons lose energy as they pass through a series of lead and stainless steel plates, generating showers of secondary electrons which produce ionisation in the Liquid Argon between the plates. Summing these signals across many layers, the energy of the showers and thus the incident electrons or photons may be determined. Beyond the LAr calorimeter in the central region is the hadronic calorimeter, also known as the Tile Calorimeter [24], designed to absorb and record the energy of hadrons. This comprises stacks of iron plates alternating with plastic scintillating tiles. Hadrons interact with the iron plates, producing showers which generate light in the scintillator medium. The light is transmitted to a series of photomultiplier tubes (PMTs) by means of wavelength shifting fibres which shift the blue scintillator light down in frequency to better match the sensitive region of the PMTs.

At the outermost radii of ATLAS lies the Muon Spectrometer (MS) [25]. Muons are not stopped in the calorimeter system, indeed they are not stopped in the muon spectrometer. The spectrometer measures their momentum by tracking the trajectory of the (charged) muons as they bend in a magnetic field generated by a series of toroidal magnets.

The detection of different particle types is summarised in Figure 1.3. Event data from ATLAS and CMS together led to the discovery of the Higgs Boson in July 2012 [27],[28]. The tracking systems of both experiments had an important role to play in this discovery, for example in the identification of muons and electrons resulting from the Higgs decay chan-




Figure 1.3: Schematic representation of particle detection in ATLAS. The tracking system determines the charge and momentum of charged particles; electrons and photons are stopped in the electromagnetic calorimeter; hadrons are stopped in the hadronic calorimeter; muon energies are determined by the muon spectrometer. Neutrinos have a very low probability to interact with matter, and are hence not detected. (Reproduced from Reference [26].)

nel $H \rightarrow ZZ^{(*)} \rightarrow 4\ell$. In the case of ATLAS, muon candidates are formed by matching reconstructed ID tracks with track segments reconstructed in the Muon Spectrometer [29]. Electron candidates require a reconstructed ID track aligned to a cluster in the electromagnetic calorimeter which satisfies a number of selection criteria [30]. Further detail is given in reference [27].




Figure 1.4: The Large Hadron Collider (LHC) Roadmap shows that two further luminosity upgrades are planned. This thesis relates to the preparation of an upgraded Strip Tracker for ATLAS, a phase 2 upgrade to be installed during Long Shutdown 3 (LS3). (Reproduced from Reference [31].)

1.3 The High Luminosity Large Hadron Collider, HL-LHC

The Large Hadron Collider (LHC) roadmap is shown in Figure 1.4, which documents a number of actual and planned energy and luminosity increases. Further work to the magnet protection system will be needed before the machine can reach its design energy of 14 TeV, hence this will not now happen until Run 3. However the machine successfully reached its design luminosity of $1.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in June 2016, which was further increased to $2.0 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in 2017.

Two further upgrades are planned to increase the luminosity of the LHC, the latter of which (“Phase 2” scheduled for 2023/2026) will give rise to the High Luminosity Large Hadron Collider (HL-LHC). In order to maximise the data taking efficiency of the experiments, luminosity levelling techniques are used to manage the instantaneous luminosity at each of the interaction points [32], with the ultimate design goal for HL-LHC being a levelled luminosity of $7.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. This is a factor of 3.75 beyond today’s values: the TRT would saturate at this multiplicity, so cannot be used; similarly the readout bandwidth of

the present strip and pixel designs would be insufficient to handle the increased detector occupancy. In addition the existing pixel and strip sensors cannot survive the integrated luminosity of $4,000 \text{ fb}^{-1}$, which HL-LHC will target due to radiation damage. For these reasons a new Inner Tracker (ITk) will be required for ATLAS. An all-silicon design is chosen, with the initial design objective that the new Inner Tracker (ITk) must perform at least as well as the present Inner Detector (ID) despite the harsh HL-LHC environment in which many more tracks will be present in each event.

An important design goal is therefore that the total radiation length (material) of the new tracker should be less than the old one. Low power design techniques and deep sub-micron ASIC technology will be used to minimise the power consumption of ITk. As a result, chip voltages will be lower and currents higher than in the present tracker. Traditional individual or parallel powering methods are no longer practical due to the power losses in off-detector cables which such methods would bring or, to look at this another way, due to the high mass which would need to be added to those cables in order to deliver the power with acceptable efficiency. (These issues will be explored further in the next chapter.)

This thesis documents the electronic system design of the ITk Strip Staves to be used in the barrel region, including performance studies of prototypes built using chipsets in the selected 130 nm CMOS technology. A stave functions both as a thermo-mechanical core to support and cool the modules, and as a cable to provide electrical services. The module itself uses a silicon sensor of approximately 10 cm by 10 cm (the largest square which may be cut from a 6" wafer) upon which the polyimide flex hybrids which support the readout ASICs are glued. Each long-strip module has one hybrid serving two banks of 4.8 cm strips whilst each short-strip module has two hybrids each serving two banks of 2.4 cm strips. The strip tracker will use DC-DC Point Of Load (PoL) conversion to develop the required operating voltage from a higher input level. This offers efficiency gains by increasing the voltage at which power is delivered to the detector and by reducing the current, hence the cost and mass (radiation length) of cable plant may be reduced. However this also gives rise to additional considerations when designing the system, for instance the need for the PoL

converters to be adequately shielded such that switching signals are not picked up in the sensors.

Further details of the stave concept, module implementation and an introduction to silicon strip detectors will be given in later chapters.

Chapter 2

Silicon Strip Detector Systems

This chapter introduces the reader to Silicon Strip Detector Systems and the mechanisms behind radiation damage and noise. Whilst a more complete treatment is beyond the scope of this thesis, additional information is available from references [33, 34, 35, 36]

2.1 An introduction to Silicon Strip Detector Systems

2.1.1 A Generic Detector System

A generic detector system, based upon an ionisation chamber, is shown in Figure 2.1. Radiation passing through the chamber liberates charge pairs, which are pulled towards the collection electrodes by means of an electric field applied across the chamber. A current pulse is induced in the circuit which is then amplified, shaped and digitised. Shaping of the signal in the analogue domain is typically required to maximise the signal to noise ratio of the detector system and, if required, to control the duration of each pulse in order to minimise system deadtime. The resultant data may be further processed, or buffered pending receipt of a readout request, before onward transmission to the Data Acquisition system (DAQ).

In order to get position information out a generic ionisation chamber, one may segment the electrodes as shown in Figure 2.2. The energy lost by the particle is proportional to the




Figure 2.1: A Generic Detector System based upon an Ionisation Chamber. Radiation passing through the chamber liberates charge pairs, which are pulled towards the collection electrodes by means of an electric field applied across the chamber. A current pulse is induced in the circuit which is then amplified, shaped and digitised. The resultant data may be further processed or buffered, before onward transmission to the Data Acquisition system (DAQ).




Figure 2.2: The collection electrodes may be segmented in order to obtain position information. A particle passing perpendicular to the plane of the sensor may develop a signal on (a) one electrode or (b) a pair of adjacent electrodes. Inclined particle tracks may produce signals on multiple segments as shown in case (c).

distance it travels through the sensing material: a greater charge is generated for inclined tracks than for perpendicular ones, although at higher angles the signal upon each electrode may in fact be reduced as the charge spreads over multiple electrodes. If the objective of the detector is to record the energy lost by the particle, then the signal is usually sampled by a multi-bit Analogue to Digital Converter (ADC): in some systems the signal may even be sent off detector in analogue form for digitisation in the counting room. However for tracking applications, where one requires only position information, it has become common to record only signals which exceed a fixed, programmable threshold. Sometimes referred to as “binary readout”, this is akin to an ADC with one bit resolution. This helps to minimise




Figure 2.3: Thermal diffusion of holes and electrons across the p-n boundary produces a region free of charge carriers known as the Depletion Region (a). By applying an electric field to reverse bias the junction the width of the Depletion Region is further increased (b).

the amount of data which needs to be transmitted off-detector, which can be very useful if one is designing a low mass system with a large channel count.

2.1.2 Silicon Detector Systems

A semiconductor detector system may be thought of as a solid-state ionisation chamber. Whilst a variety of semiconductor materials such as CVD Diamond, Gallium Nitride (GaN), Germanium or Silicon Carbide may be used, silicon is a popular choice due to its abundant usage in consumer electronics: the fabrication of silicon devices is well understood, which helps to keep cost under control, and the consistency of the bulk material is generally excellent.

In Figure 2.3 the generic ionisation chamber is replaced by a p-n junction in silicon. Thermal diffusion of holes and electrons across the p-n boundary produces a region free of charge carriers known as the Depletion Region with a corresponding potential difference between its two sides, the “built-in” potential V_{bi} . A particle which passes through the junction gives rise to electron-hole pairs throughout the silicon, but outside the Depletion Region these may only move by diffusion which is too slow for our purposes. Within the Depletion Region the electron-hole pairs are pulled by the intrinsic electric field to form a faster, drift signal. It can be shown that [33] application of reverse bias across the junction increases the depletion depth w_d according to the following equation:

$$w_d = \sqrt{\frac{2\epsilon(V_b + V_{bi})}{Ne}} \quad (2.1)$$

where:

- ϵ is the dielectric constant (11.9 ϵ_0 for silicon)
- V_b is the applied bias
- V_{bi} is built-in junction potential (of order 0.5 V for a typical silicon sensor)
- N is the bulk dopant concentration
- e is the electronic charge

Considering the diode as a parallel plate capacitor of area A it therefore follows that its capacitance C is given by

$$C = \epsilon \frac{A}{w_d} = A \sqrt{\frac{\epsilon e N}{2(V_b + V_{bi})}} \quad (2.2)$$

Hence for the common case where $V_b \gg V_{bi}$:

$$C \propto \frac{1}{\sqrt{V_b}} \quad (2.3)$$

The diode is said to be “partially depleted” until the depletion depth w_d extends across the full thickness of the silicon, which occurs at a bias potential known as the Depletion Voltage: the diode is now said to be “fully depleted”, the sensor capacitance is minimised and the charge collected by drift when a particle traverses the sensor is maximised.

A particle passing through the depletion region of a p-n junction gives rise to a number of electron-hole pairs which are pulled to the diode terminals by the action of the electric field. Not all the energy is converted to ionisation: the average amount of energy lost per

electron–hole pair is 3.6 eV [33]. The charge Q which passes between the diode terminals in response to a particle which traverses it is therefore proportional to the energy deposited:

$$Q = \frac{E}{E_i} e \quad (2.4)$$

Where E is the energy deposited by the particle and E_i is the ionisation energy. (For silicon $E_i = 3.6$ eV.)

In turn the rate of energy loss as a particle traverses any medium is given by the Bethe-Bloch formula[37]:

$$-\frac{dE}{dx} = K z^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right] \quad (2.5)$$

where:

- K is the constant $\frac{4\pi e^4}{c^2 m_e} N_A$
- e is the electronic charge
- c is the speed of light
- m_e is the mass of the electron
- N_A is Avagadro's number
- Z is the atomic number
- A is the atomic mass of the material
- β is the speed of the incident particle
- γ is the Lorentz factor
- W_{max} is the maximum kinetic energy which can be transferred in a single collision
- I is the mean excitation energy

- $\delta(\beta\gamma)$ is the density effect correction to ionization energy loss

A Minimum Ionising Particle (MIP) would be any particle for which the rate of energy loss is close to the minimum predicted by this formula, which occurs when $\beta\gamma \approx 3$. An important point here is that if a detector system is designed to deliver good performance for Minimum Ionising Particles, it is able to detect any other interacting particle as its rate of energy loss will be higher.

In 300 μm thick silicon, a MIP will deposit approximately 84 keV which, dividing by the ionisation energy of 3.6 eV, corresponds to a liberated charge of 23k electrons or 3.7 fC. In practice radiation damage effects may reduce the charge which can be collected, even in a fully depleted sensor, for reasons which are outlined in the following section.

2.1.3 Radiation Damage in Silicon

Whilst radiation damage has numerous effects upon the operation of silicon sensors and readout electronics, there are essentially only two damage mechanisms. The first of these is Displacement Damage [38], in which an incident particle displaces an atom from the crystal lattice. Modern, deep sub-micron CMOS electronics is essentially immune to this form of damage as the electronic circuit is formed on the surface of the wafer: the main body (or 'bulk') of the silicon does not contribute to the circuit's function. However in the sensor the signal is formed within the silicon bulk, so displacement damage is a major concern. During the lifetime of the experiment, the depletion voltage of the sensor increases due to displacement damage. Ultimately we may reach the limit of the power supply system, meaning that we are no longer able to reach full depletion, with a consequent reduction in signal levels as the active thickness is reduced. A further degradation of the available signal comes about due to a reduction of carrier lifetimes: some of the carriers may become trapped by the damaged lattice sites. In tandem with these effects there will be an increase in the sensor leakage current which, as will be explained later, can increase the intrinsic noise. Hence it is important that detector systems are designed to deliver the required performance

for the duration of the experiment, taking into account sensor characteristics across the full lifetime.

The second of these damage mechanisms is that of Ionization Damage. In this case, energy deposition in insulating (oxide) layers generates electron-hole pairs. Some of the holes become trapped, leading to a buildup of positive charge in the oxide layer and a consequent increase in leakage between adjacent structures. Although this will affect sensors at some level, it is more of a concern with respect to the readout electronics. Any leakage between the source and drain of a MOS transistor will affect its operating threshold: analogue circuits may have degraded function and digital circuits may ultimately stop working. Indeed increased leakage current may in itself be a problem in a tracking detector application, as the additional power needed to bias the circuit must first be supplied, requiring extra cross section in the power cabling, and the additional heat must be removed, requiring additional cooling capacity.

A first study of the effects of ionising radiation upon transistors made in the 130 nm process which is used for the ITk strip chipset was published in 2005[39]. In this process, the standard gate oxide¹ is only 2.2 nm thick which limits the extent to which charge trapping may occur, however transistors are three dimensional structures. In particular the Shallow Trench Isolation (STI) oxide found at the edge of an open layout transistor is thicker: positive charge trapped here changes the characteristic of the parasitic transistor formed by the source, STI (in lieu of the gate) and the drain. However one must also consider the effects of ionising radiation upon the depletion region of the transistor. In the case of a PMOS transistor, the interface states which form at the boundary of the STI and depletion region are also positive. Both types of defect will increase the threshold of the parasitic transistor: as such there is no increase in leakage, and operation of the main transistor is not affected. In the case of an NMOS transistor, the accumulation of positive charge in the STI oxide will decrease the threshold of the parasitic transistor: there will be additional leakage and the main device will have degraded function. However in an NMOS transistor the interface

¹Transistors using the standard gate oxide are rated for operation at 1.5 V. A thicker oxide is available for transistors which need to be operated at up to 3.3 V.

states have negative charge, and will ultimately compensate for the trapped STI charge, restoring the original performance of the device.

Whilst it is possible to use special “Radiation Hardening by Design” (RHBD) techniques to eliminate radiation induced leakage within a transistor by the use of Enclosed Layout Transistors (ELT) [40], in which the source is surrounded by the gate such that possible parasitic transistors are eliminated, use of this geometry increases the size of each transistor and therefore imposes limits on circuit density. In the ITk Strip chipset, this approach is used only for the most critical (analogue) parts of each chip². The remainder of the circuits are built using the standard library with the use of additional constraints regarding the minimum transistor size.

Hence whilst the ITk Strip chipset will continue to function well beyond the expected worst case lifetime dose of 53.3 Mrad [41], the leakage current will peak at approximately 0.6 Mrad due to the (fast) accumulation of trapped positive charge in the STI oxide of NMOS parts, after which it will start to fall due to the (slower) accumulation of compensatory interface states. This phenomenon is known in ATLAS as the “TID peak” [42]: its magnitude varies as a function of temperature and dose rate, with the highest currents being observed for high dose rate and low temperature. Preliminary studies involving the “pre-irradiation” of prototype ITk Strip front end chips to 8 Mrad have given very positive results. In production, it is likely that all ABCStar front-end chips will be pre-irradiated in order to reduce the susceptibility of the detector system to this effect.

2.1.4 Noise in Silicon Detector Systems

A Detector Front-End Circuit and its Equivalent Noise Schematic are shown in Figure 2.4. Capacitance C_d represents the sensor, biased through resistor R_b . For high frequencies, bypass capacitor C_b has a low impedance, hence bias resistor R_b may be considered as being

²In fact the NMOS input transistor of the ABC130 front end amplifier was not an enclosed layout device, with the result that a significant noise increase was observed during radiation tests. The input transistor of the later chip, ABCStar, is a fully enclosed design with the result that this problem has been solved.




Figure 2.4: A Detector Front-End Circuit and its Equivalent Noise Schematic. Reproduced from H. Spieler, *Semiconductor Detector Systems* [33]. © Oxford University Press 2005. Reproduced with permission of the Licensor through PLSclear. The figure is further explained in the text.

connected in parallel with the sensor. Coupling capacitance C_c represents the readout strip in an AC-coupled strip sensor: this also has low impedance at high frequencies and hence may be neglected from noise calculations. Series resistance R_s represents the summation of all resistive elements between the sensor and the amplifier input, including the resistance of the readout strip on the sensor; the resistance of the wirebonded connection between the readout strip and the amplifier; the resistance of the input fan-in routing within the readout chip itself.

In the equivalent noise schematic, parallel noise due to statistical fluctuations in the sensor bias current (“shot” noise) is represented as current source i_{nd} in parallel with detector capacitance C_d . Parallel noise due to thermal noise in the bias resistor is represented as current source i_{nb} in parallel with resistance R_b . Series noise due to thermal variations in the strip resistance is represented as voltage source e_{ns} in series with resistance R_b . Finally the noise sources within the amplifier itself are represented by series noise voltage e_{na} and parallel noise current i_{na} .

Noise voltage e_{nd} arising from statistical fluctuations in the sensor bias current I_d may be represented as:

$$e_{nd}^2 = i_{nd}^2 \frac{1}{(\omega C_d)^2} = 2eI_d \frac{1}{(\omega C_d)^2} \quad (2.6)$$

Noise voltage e_{np} arising from thermal variations in parallel resistance R_b may be repre-

sented as:

$$|e_{np}|^2 = \frac{4kTR_p}{1 + (\omega R_p C_d)^2} \quad (2.7)$$

Noise voltage e_{nr} arising from thermal variations (thermal noise) in series resistance R_s may be represented as:

$$e_{nr}^2 = 4kTR_s \quad (2.8)$$

Noise voltage e_{na} arising from (white) thermal noise and (frequency dependent) “flicker” (“1/f”) noise sources within the amplifier may be represented as:

$$e_{na}^2 = e_{nw}^2 + \frac{A_f}{f} \quad (2.9)$$

where e_{nw} represents the white noise component and A_f represents the gain of the amplifier as a function of frequency.

Noise voltage e_{nia} arising from statistical fluctuations in transistor bias currents, i_{na} , within the amplifier, may be represented as:

$$e_{nia}^2 = i_{na}^2 \frac{1}{(\omega C_d)^2} \quad (2.10)$$

As there is no correlation between any of these noise sources, the noise voltage spectral density at the amplifier input, $e_{ni}(f)$, is given by their sum in quadrature:

$$e_{ni}^2(f) = \frac{2eI_d}{(\omega C_d)^2} + \frac{4kTR_p}{1 + (\omega R_p C_d)^2} + 4kTR_s + e_{na} + \frac{i_{na}^2}{(\omega C_d)^2} \quad (2.11)$$

It can further be shown [33] that, for amplifiers with a simple, first order $CR - RC$ shaper, the Equivalent Noise Charge Q_n is given by:

$$Q_n^2 = \left(\frac{\epsilon^2}{8}\right) \left[\left(2eI_d + \frac{4kT}{R_p} + i_{na}^2\right) \cdot \tau + (4ktR_s + e_{na}^2) \cdot \frac{C^2}{\tau} + 4A_f C^2 \right] \quad (2.12)$$

where $\frac{\epsilon^2}{8}$ normalises the noise to the gain and C represents the total input capacitance, dominated by the sensor but including, for example, additional capacitances coming from wirebonds and/or any PCBs glued on top of the sensor. The first term represents the total parallel (current) noise and increases with the amplifier shaping time. The second term represents the total series (voltage) thermal noise and increases with the total input capacitance, but decreases with shaping time. The third term represents the flicker (1/f) noise, which is independent of shaping time but also increases with the total input capacitance.

The input amplifiers of both front end chips discussed in this thesis, ABC130 and ABCStar, are Regulated Cascode amplifiers made in CMOS technology, and in reality they are closer to a second order CR-RC filter. It can however be shown[43] that the dominant series noise contribution comes from thermal noise in the MOS input transistor: thermal noise in the second and third MOS transistors can generally be neglected. In the case of an irradiated module, where the sensor leakage current can approach 1 mA, the dominant contribution to the equivalent parallel noise current will come from shot variations in the sensor leakage current itself³. In the case of ABCStar there will be an additional series noise contribution from the feedback resistance within the front end amplifier itself⁴.

2.2 Chapter Summary and Conclusions

This chapter has introduced a generic detector system and the specific case of a silicon microstrip detector. In any HL-LHC experiment, the long-term effects of radiation are a major concern: in the case of the ATLAS ITk Strip Detector the major concerns are

³This is one of the reasons why the operational temperature of the ATLAS ITk Strip detector must be carefully controlled: the sensor leakage current is also temperature dependent.

⁴The ABC130 amplifier uses active feedback in which a MOS transistor is placed in parallel with the feedback capacitance C_f , however as this leads to signal compression for the negative charges collected in the ATLAS ITk's n^+onp sensors, in ABCStar this has been replaced by a simple fixed resistor R_f .

displacement damage in the sensors and ionization damage in the electronics. Whilst a robust treatment of the electronic noise in a silicon microstrip detector system must consider a variety of noise sources, these may be characterised as contributions toward the series (voltage) or parallel (current) noise. The series noise contribution, dominated by thermal noise in the MOS input transistor, scales directly with the detector capacitance. As the detector system is subjected to radiation, the parallel noise contribution quickly becomes dominated by shot noise due to increased leakage currents in the silicon sensor itself.

Chapter 3

System Architecture, Test Methodology and System Metrics

This chapter documents the system architecture chosen for the ITk Strip barrel and the methods used to calibrate the detector front end are introduced.

3.1 System Architecture

3.1.1 Layout of the ITk Strip Detector

To deliver the tracking performance needed at HL-LHC, it is necessary to replace the tracking system at the heart of the ATLAS experiment. The design of the present Transition Radiation Tracker (TRT) is not compatible with such high particle fluences, hence an all silicon design is chosen.

The layout concept of one quadrant of the ATLAS Inner Tracker (ITk) is shown in Figure 3.1. Pixel layers are shown in blue and strip layers in red. The design of the ITk Strip Detector is further detailed in the Technical Design Report [41], but it can be seen from the figure that the system comprises four barrels, with one endcap region to each end. The barrel region is 2.8 m long, with radii and segmentation as shown in Table 3.1. Each barrel is fitted with a number of “staves”, each 1.4 m long, which are tilted so as to minimise gaps in coverage. Similarly each endcap disk is tiled with a number of castellated “petals”.




Figure 3.1: Schematic layout of the ITk Layout for the Phase-II upgrade of ATLAS [44]. Only one quadrant and only active detector elements are shown. The active elements of the barrel and end-cap Strip Detector are shown in blue, for the Pixel Detector the sensors are shown in red for the barrel layers and in dark red for the end-cap rings. The horizontal axis is along the beam line with zero being the interaction point. The vertical axis is the radius measured from the interaction region.

These concepts are introduced in Figure 3.2.

Each stave or petal comprises a support structure made of carbon fibre facings, honeycomb and low density carbon foam, into which a titanium cooling pipe has been integrated. The carbon fibre skins have polyimide flex circuits known as “bus tapes” co-cured upon them to route electrical services between the End of Substructure (EoS) card, which provides the interface to the off-detector systems, and the detector modules. Modules are glued to each face of the stave (petal) to form two back-to-back sensor planes. On each side the modules are glued at fixed, non-zero angles to give a small angle stereo geometry. In this way each layer provides a space point in three dimensions. In the case of the barrel region, the two innermost layers, where the track density will be highest, are populated with short strip modules with 2.4 cm strip length. The two outermost layers are populated with long strip modules having 4.8 cm strip length: in this manner the power dissipation of the system and indeed its cost are reduced whilst the required performance is achieved.

Layer	Radius	Segments in ϕ	Strip Pitch	Strip Length	Stereo Angle
0	405	28x1280	75.5	24.1	52
1	562	40x1280	75.5	24.1	52
2	762	56x1280	75.5	48.2	52
3	1000	72x1280	75.5	48.2	52
	mm		mm	mm	mrاد

Table 3.1: Main layout parameters for the Strip Barrel (from the ATLAS ITk Strip Technical Design Report (TDR) [41])




Figure 3.2: ITk Strip Barrel and Petal Concepts (from the ATLAS ITk Strip Technical Design Report (TDR) [41])

The cross-sectional structure of the stave (petal) core is further explored in Figure 3.3. The low density carbon foam has high thermal conductivity [45] and is confined to the regions around the cooling pipes, to which it is attached with carbon loaded epoxy adhesive, with the remainder of the core comprising carbon fibre honeycomb due to its better mechanical properties. The edges of the stave cores are completed by carbon fibre closeouts, also known as “C-channels”, into which the mounting brackets may later be glued.




Figure 3.3: Cross Sectional view of an ITk Strip Barrel Stave (from the ATLAS ITk Strip Technical Design Report (TDR) [41])

3.1.2 Electronics Architecture




Figure 3.4: ITk Electronics Overview (from the ATLAS ITk Strip Technical Design Report (TDR) [41])

An architectural overview of the ITk Strip On-detector Electronics is given in Figure 3.4. Whilst drawn in a stave geometry, it applies equally to the forward region. The power and communications interface between each stave or petal and the outside world is provided by the End Of Substructure (EoS) card. The Versatile Link+ (VL+) [46] provides radiation hardened optical transport for the control signals and data between the on- and off-detector systems. For each short stave side, one downlink and two uplink fibres are used. Each long strip stave side will utilise one downlink and one uplink fibre only as the required uplink bandwidth is within the capability of a single optical link.

Each downlink fibre’s raw bandwidth of 2.56 Gbps is decoded by the low power GigaBit Transceiver (lpGBT) chip [47] to provide up to 1.28 Gbps user bandwidth. Forward Error Correction is used to provide resilience against error bursts of up to 12 bits duration (FEC12) which are expected during operation in ATLAS due to Single Event Effects within the PIN diode that receives the optical signal. In our application the 1.28 Gbps user bandwidth is demultiplexed by the lpGBT to give the 8 streams of Clock, Control Reject (CCR) data at 160 Mbit/s which are needed to configure the front end modules. A further 80 Mbit/s Internal Control (IC) bandwidth is reserved to configure the lpGBT chip itself. An additional 80 Mbit/s External Control bandwidth, whilst originally intended to control an additional lpGBT, will instead be used by us to provide a control bus to configure the Autonomous Monitor And Control (AMAC) [48] chip on each module’s power board. As the lpGBT chip may service no more than one uplink fibre, it can be seen that each short strip stave side requires two lpGBT chips: only the “primary” lpGBT is connected to the downlink, with the “secondary” lpGBT being configured by the primary lpGBT over an I2C link. Other on-detector components of the high bandwidth Versatile (optical) Link (VL+) are the GigaBit TransImpedance Amplifier (GBTIA) [49] which amplifies the downlink signal received by the PIN diode and the low power GigaBit Laser Driver (lpGBLD) [50] used to modulate the output of a Vertical Cavity Surface Emitting Laser (VCSEL) diode to drive the uplink back to the counting room.

Low voltage power is supplied to the stave in the form of a nominal 11 V power rail which must be locally converted to the (lower) voltages required by the active components, both at EoS and on the modules themselves. The rationale behind this design choice is outlined below.

3.1.3 Powering Strategies

Four possible low voltage powering strategies are outlined in Figure 3.5. The first two schemes are used by the silicon trackers of running LHC experiments. Individual powering is used in the ATLAS SCT [51]. Each module has its own floating power supply channel and its own





(a) Individual Powering, $P = nI^2R$ (b) Parallel Powering, $P = n^2I^2R$ (c) Serial Powering, $P = I^2R$ (d) DC-DC Powering, $P = n^2I^2R/r^2$

Figure 3.5: Four Possible LV Powering Schemes, showing the expected power losses in off-detector cabling of total resistance R for n loads each drawing current I . (In the case of DC-DC powering, r is the conversion ratio V_{in}/V_{out} .)

cable: power losses in the cables scale with the number of modules. Parallel powering is used in the present CMS strip detector [52], with a small number of modules being connected in parallel to a common cable and supply. Power losses in the cables now scale with the square of the number of modules, hence whilst the number of cables is reduced their resistance must be kept low in order to maintain high efficiency for the power distribution system. In both cases the cables make a significant contribution to the cost of the detector, and its material budget.

The ATLAS ITk will occupy a much larger volume than the present SCT [53], and its segmentation will be greatly increased: channel counts are increased by a factor of 10, making the efficiency of the detector electronics and the power delivery system extremely important. The traditional individual or parallel powering methods outlined above are no longer practical due to the amount of material (copper) which would be required to implement it with reasonable power efficiency, and its cost.

Ignoring leakage currents - which, as discussed in the previous chapter, can actually be quite important in irradiated circuits - CMOS logic dissipates power only when transistors are switching. The work being done is primarily charging the gate capacitance of the device. The power P required to switch a gate is therefore proportional to the gate capacitance C , the switch frequency f , and the square of the voltage V : $P \propto CV^2f$. Whilst reducing the operating voltage of the digital circuitry will bring the power consumption down, this competes against the need to maintain sufficient margin for the front end amplifier to operate without excessive compression. The chosen technology node for ATLAS ITk strips is 130 nm, which offers the possibility to run digital logic at 1.0 V whilst running the analogue front end circuitry at 1.2 V.

The remaining two powering schemes are realistic candidates for use to power the silicon trackers of HL-LHC experiments. In Serial Powering a constant current is passed through a chain of serially connected modules: each module uses one or more shunt transistors, which may be embedded within the existing ASICs, to dissipate any surplus current in order to maintain the required operating voltage. The current flowing through the chain must

therefore be equal to that of the module in the chain which has the highest demand. This is extremely efficient if all modules are the same, however in scenarios such as the ITk endcap where several module designs are used, additional power must be dissipated in the shunts of the smaller modules such that the current through the chain remains sufficient to power the largest modules. Power losses in the cabling are now independent of the number of modules in the chain, hence the cost and material of the cables may be minimised.




Figure 3.6: Topology of the DC-DC Synchronous Buck Converter

In DC-DC powering, a Point Of Load (PoL) switching regulator is used to develop the required voltage from a higher input level. This could be done using a switched capacitor technique, however the community has been developing a PoL converter based upon the more common “Synchronous Buck” topology illustrated by Figure 3.6. The desired output voltage is achieved by adjustment of the duty cycle of the two switch transistors, alternately charging and discharging inductor L . Whilst commonly used in industry, this technique requires some customisation for use in tracking detectors. Firstly, air core inductors must be used as any ferromagnetic materials would saturate in the magnetic field present in the tracking volume. Secondly, the controller chip must be radiation tolerant, which has led to

the development of the “FEAST” series of buck converter ICs [54]. Thirdly, for use near sensitive front end circuits, the converters must be carefully screened to control both electric and magnetic field emissions. Finally, for use in tracking detectors, the form factor of the converters must be carefully controlled to meet space constraints and keep material to the absolute minimum. As was the case with serial powering, DC-DC conversion offers efficiency gains by increasing the voltage at which power is delivered to the detector and by reducing the current, hence the mass (and cost) of the cables may be reduced. However there will be some power dissipation within the converter itself: dependent upon detailed design choices, a figure of 75% efficiency is not unrealistic.

A first generation of stave prototypes was made to evaluate the performance of the Serial Powering and DC-DC schemes using a preliminary chipset made in 250 nm technology, comprising the ABCN25 [55] front end chip and the Buffer Control Chip [56] (hybrid controller). It was determined that as more modules were added to the stave it became increasingly difficult to maintain good performance for all modules of a serially powered chain, with the modules showing excess occupancy correlated with readout activity¹. However the solution based upon DC-DC converters was found to deliver good performance through to a complete (12 module) stave, with no additional noise burden correlated with readout activity. Considering also the difficulty of making serial powering work with good efficiency in the endcap region, which covers a petal geometry with several module designs of different sizes and hence different power requirements, the decision was taken to baseline the use of DC-DC technology. This has led to the development of miniature, low mass, radiation and magnetic field tolerant DC-DC converters suitable to be built into each strip detector module as described in the following sections.

Sensor bias (high voltage) cables also contribute to the mass of the system, material which may be a concern for tracking applications. As the sensor bias current of a typical silicon strip module is perhaps three orders of magnitude down from the current needed to power the front end chips, parallel powering is a possibility, however if one sensor were to

¹The method used for this determination, the Double Trigger Noise Occupancy test, is explained in Section 3.3.8

fail as a short circuit it would become impossible to bias all sensors connected in parallel to it. To mitigate against this possibility, ATLAS ITk have been investigating the possible use of radiation hard high voltage switch transistors as a means to disconnect any failed sensor from a common power bus. Known as High Voltage MULTipeXing (HV-MUX), this is detailed below.

3.2 Module Architecture




Figure 3.7: The Short Strip Module Concept, exploded view (from the ITk Strip Technical Design Report (TDR) [41], updated to reflect final sensor dimensions.)

The Short Strip module concept is introduced in Figure 3.7, and an annotated photograph of the completed Short Strip Module SS1 is shown in Figure 3.8. Although Figure 3.7 refers to the later “Star Chips” chipset and Short Strip Module SS1 was built using the ABC130/HCC130 chipset as used for most work reported in this thesis, the module layout is essentially the same.

The base of the ABC130 Short Strip Barrel Module [57] is the silicon sensor, approxi-




Figure 3.8: Short Strip Module SS1

mately 98 mm x 98 mm, which is divided into four “columns” of strips, each approximately 24 mm long. Two hybrids are glued to the sensor: it can be seen that there are two mirror symmetric hybrid layouts, designated as “left-handed” (LH) and “right-handed” (RH) respectively. Each hybrid comprises a row of 10 ABC130 (or ABCStar) Front End Chips in communication with a single HCC130 (or HCCStar) Hybrid Controller Chip, and its Front End chips are wirebonded to two columns of strips on the sensor, one which runs underneath the hybrid (inner column), and one which does not (outer column). Also glued to the sensor is the module’s Power Board, located between the two hybrids, with the DC-DC converter block located within the metal shield can.

The following sections describe in detail the architecture of the ABC130/HCC130 chipset in 130 nm technology. (Later chapters which refer to work performed with the “Star Chips” chipset will begin by summarising significant differences in functionality.)

3.2.1 ATLAS Trigger Architecture

Whilst the trigger architecture which will be used for ATLAS operations at HL-LHC [58] is largely beyond the scope of this thesis, a brief overview is necessary in order to understand

the operation of the ITk Strip chipset. Two hardware trigger architectures are foreseen: the “baseline” scheme, a single-level trigger architecture, and the “evolved” scheme, a two-level trigger architecture. Both schemes start by making a preliminary event selection based upon data from the Calorimeter and Muon systems, to form the Level 0 (L0) trigger (command) stream. In the baseline scheme each event selected by the L0 trigger will be transmitted off-detector to be further filtered in software. In the evolved scheme the trigger hardware may identify, for certain events which have already been selected as part the L0 trigger stream, one or more regions of the ITk which may be of potential physics interest. A Regional Readout Request (R3) command is sent to the selected part(s) of the detector, which must prioritise transmission of the selected event. The resulting data is then processed by the Level 1 (L1) trigger hardware and events which remain of interest are then read out for the entire detector by transmission of a L1 readout command.

This is handled by the ITk Strip chipset as follows. Upon receipt of a L0 trigger, the hit pattern corresponding to the selected event is transferred to the L1 buffer. Here it will remain until a Level 1 (L1) or Regional Readout Request (R3) command is received, which requests an event specified by the corresponding L0 identifier (id). Hence in the baseline scheme, L0 and L1 readout commands are essentially sent as matched pairs: the strip detector is designed to support transmission of events at a mean rate of 1 MHz (one event every 25nS). In the evolved scheme, part of the available readout bandwidth is reserved for prioritised Regional Readout data, however the system must also support transmission of events selected by the Level 1 trigger hardware at up to 800 kHz. In order to simplify the off-detector data flow, no attempt is made to merge data which has already been sent in response to receipt of an R3 command back into the L1 data stream: rather the data remains persistent in the front end and can also be read in response to receipt of a L1 readout command, until overwritten by another event with the same (seven bit) L0 identifier.

Note that whilst the ABCStar/HCCStar chipset which will be used in production is designed to comply with the above specifications, with L0/L1 rates at 4 (2) MHz/600 (800) kHz respectively, the earlier ABC130/HCC130 chipset which forms the basis for most of this the-




Figure 3.9: Block Diagram of ABC130 ASIC (from the ITk Strip Technical Design Report (TDR) [41])

sis was only designed for L0/L1 rates of 500 kHz/200 kHz.

3.2.2 The Front End Chip ABC130

A block diagram of the ABC130 chip is shown in Figure 3.9. In order to reduce the part count in the finished detector, each ABC130 die serves 256 channels, twice that of the previous ABCN25 chip [55], with 128 channels serving each of two columns of strips on the sensor. Input signals from the strip sensors are amplified, shaped, discriminated and then filtered according to the content of a mask register to suppress noisy channels. The resultant hit pattern is then fed into the L0 buffer. Upon receipt of a Level 0 (L0) trigger the event corresponding to the programmed latency is transferred from the L0 buffer into the L1 buffer. Here it will remain until a Level 1 (L1) or Regional Readout Request (R3) command is received which requests an event specified by its L0 identifier (id). Whilst events which are requested by the L1 command are queued to be transferred off detector in sequence, events requested by the R3 command are prioritised and will be transmitted immediately.

Two versions of the ABC130 chip were submitted, one containing an additional Fast Cluster Finder (FCF) block. The chip requires two input clocks, the bunch crossing (BC) clock and the data recovery clock (DRC). Two 80 Mbit/s control signals are required, each comprising a pair of time multiplexed streams at 40 Mbit/s, namely the L0 and command




Figure 3.10: Block Diagram of HCC130 ASIC

signals (L0/COM) and the L1 trigger and R3 signals (L1/R3). Output data is transmitted as a series of packets of 60 bits in length. Within a barrel hybrid, the ABC130 chips are arranged into two loops of 5 chips each. Each chip is able to pass data to the “left” or to the “right”: although a complete loop would typically be configured to pass its data in the same direction, as the chips at both ends of the loop are able to pass data to HCC130 the loop may be split if required, for example in order to route data around a failing chip. To this end the ABC130 design provides a pair of bidirectional SLVS transceivers named DATA and XOFF at opposing corners which are configured according to the desired data flow direction such that at one corner DATA is received and XOFF is transmitted, whilst at the other corner XOFF is received and DATA is transmitted. The XOFF signal is used for flow control purposes, ensuring that ABC130 may only transmit data when the next link in the chain, ABC130 or HCC130, is able to receive it.

3.2.3 The Hybrid Controller Chip HCC130

A block diagram of the HCC130 chip is shown in Figure 3.10. The chip forms the interface between the ABC130 chips and the End of Substructure card (EoS) which, in production staves, would in turn interface the stave to the off-detector systems by means of the GigaBit Transceiver (GBTx) chip and Versatile (optical) Links (VL). The HCC130 receives three signals from the stave bus: the BC clock, the L0/COM multiplex and the L1/R3 multiplex. The design incorporates a Phase Locked Loop, the ePLL taken from the GBTx chip, which generates a series of duty and phase controlled 160 MHz and 320 MHz clocks from the input BC clock. Some of these are used internally within HCC130, others are used to generate the 40 MHz BC clock and 80 MHz or 160 MHz data clock (DCLK) as sent to the ABC130 chips. In the other direction, data packets from the front end chips are received and phase aligned at 80 Mbit/s or 160 Mbit/s, then fed into the Data Concentrator block. This buffers them and interleaves the two streams for transmission along the stave at 160 Mbit/s or 320 Mbit/s. One novel feature of the HCC130 chip is that it incorporates two sets of pads with the same functions along each of the long edges of the die. This was done to facilitate the use of the circuit on mirror symmetric left- and right-handed hybrids as introduced in Figure f:SS1.

In addition HCC130 incorporates an Autonomous Monitor (AM) block, foreseen as part of the detector control and interlock system. The analogue and digital regulators within the ABC130 chips are enabled by control buses driven by the AM, reverting to a minimal current state when these are disabled. AM has an ADC with seven input channels to digitise key parameters: (unregulated) voltage, (bandgap) voltage reference, HCC130 (regulated) voltage, internal (HCC130) temperature, hybrid temperature, sensor bias (return) current and one external (spare) voltage. For each of these parameters, the AM may be programmed such that, in the event that programmable upper or lower bounds are exceeded, the ABC130 voltage regulators are disabled. In this way a module may be protected, for example, against events such as temperature excursions. (Having been prototyped in the HCC130, the AM block was subsequently reused in the standalone Autonomous Monitor and Control (AMAC) chip to be used in later revisions of the power board.)




Figure 3.11: Block Diagram of the Autonomous Monitor And Control (AMAC) ASIC (from the ITk Strip Technical Design Report (TDR) [41])

3.2.4 The Power Board

The final component of the module, the Power Board, comprises a low pass filter for the sensor bias supply, a miniature, low-mass DC-DC converter and space for an optional High Voltage Switch block (to be integrated in later versions). The present DC-DC converter utilises the FEAST2.1 buck converter ASIC together with a novel solenoidal coil, and operates with a voltage gain of around 6.6. Later versions will use an upgraded chip developed from FEAST2.1, named bPOL12V, to achieve the required radiation tolerance. The High Voltage switch block will use radiation hard GaNFET devices to ensure that the full sensor bias potential may be switched throughout the lifetime of the experiment. The Autonomous Monitor And Control ASIC shall form part of later versions of this board, enabling the DC-DC output, operating the HV switch and monitoring parameters such as hybrid temperatures, voltages and the sensor bias (return) current.




Figure 3.12: Simulation of noise and signal distributions

3.3 Test Methodology and System Metrics

3.3.1 Threshold Scans and the Complementary Error Function

The ATLAS ITk strip system, like the SCT which it will replace, is a binary detector system. Input signals are compared against a programmable threshold, with hits being fed into the readout buffers only when the threshold is exceeded. Consider the simulations shown in Figure 3.12. The top left plot represents the noise of the system, a Gaussian distribution with sigma 14 mV and mean 0 mV. The step function shown in the top right plot represents the signal: in a system without noise, a signal of 150 mV would result in hits for thresholds below 150 mV only. The bottom left plot shows the complement of the integral of the Gaussian noise distribution shown in the top left plot, added to the signal distribution shown in the top right. This represents the expected response of the complete system: for injected signals of constant magnitude 150 mV in a system with a Gaussian noise distribution of

sigma 14 mV, a threshold scan would recover this distribution. The bottom right plot fits the complementary error function to this distribution, thereby recovering the mean signal and the width (sigma) of the noise distribution.

This illustrates the most common way by which the noise of a binary detector system is measured: the threshold scan. Whilst injecting pulses of known magnitude, by means of charge injection circuitry built into the front end chip, the threshold is scanned from low to high. The ITk Strip front end chips give us the possibility to inject charges of amplitude defined by an integrated calibration DAC into any number of their channels, from 1 to 256, in a single event. The resultant histogram of occupancy as a function of threshold is fitted to the complementary error function on a channel by channel basis to determine VT50, the threshold for which the probability that a hit is recorded is 50%.

One point to note here is that, at very low thresholds, channels into which charge is not being deliberately injected are likely to have significant occupancy due to noise. This may affect the analogue functionality of the chip due to increased digital activity, as the discriminators will, on average, fire more often. In principle the susceptibility of measurements to this effect can be reduced if, instead of scanning the threshold as one injects pulses of a fixed charge, one fixes the threshold and scans the injected charge. By this method one may determine QT50, the charge for which the probability that a hit is recorded is 50%. However the resolution of the ABC130's eight bit calibration DAC, compared to the noise of an unbonded chip, gives insufficient points on the resulting plot of the error function to yield an accurate determination of the noise. This has been addressed by the extension of the calibration DAC to nine bits in the ABCStar design.

3.3.2 Calibration Timing

In order to make accurate measurements using the internal calibration signal, the timing of the signal must first be adjusted for good efficiency. In response to the calibration command, ABC130 generates a digital pulse which is passed through a chain of 63 current starved inverters. The point in the chain at which the signal is picked up, from 0 to 63, is set by

programming the calibration delay register. The digital pulse is then routed to the calibration circuit, a fast analogue switch which moves between two fixed levels: V_{dd} and (V_{dd} - calibration DAC setting). This pulse is made available for external monitoring through the chip's internal analogue multiplexer and is further distributed within the chip. At each channel it is buffered and may be suppressed by means of the calibration mask register, before injecting charge into the front end amplifier by means of a 60 fF capacitor. Charge of opposite polarities will be injected into the front end on the falling and rising edges of the pulse, corresponding to the signal polarity expected from $n^{+}onp$ or $p^{+}onn$ sensors respectively. The width of the pulse, 200 ns^2 , is chosen to be a factor of ten larger than the shaping time of the amplifier to ensure that the charges injected by the two edges of the pulse do not interfere with one another. The polarity of the pulse is programmable such that the first edge can inject charge of either polarity.




Figure 3.13: An example of a Strobe Delay Test for 128 channels of a single ABCStar chip. The plot represents, for fixed Threshold and Charge Injection values, the occupancy of each channel as a function of its Calibration Delay (Strobe Delay) value. It can be seen that between the values of 10 and 30, all channels have 100% occupancy. The analysis fits these two edges and sets the calibration delay to be a specified fraction of the distance between them.

The method used to adjust the calibration timing first records a threshold scan for a 2 fC

²In the ABCStar design the calibration pulse duration is increased to 400 ns.

charge, using the default timing, in order to approximate the threshold setting equivalent to 2 fC. The second stage of the process changes the compression mode of the chip to report hits occurring only in the central time bin (bunch crossing) of the three which are stored in the pipeline, then using a 4 fC charge at a 2 fC threshold the calibration delay is varied across the full range of 64 steps. The expected result is a “top hat” function, giving the extent of the calibration delay for which the injected signal falls into the central time bin. The software fits the two edges of the distribution, and sets the calibration delay to be a specified fraction of the distance between the edges. The optimal fraction varies according to the characteristic of the front end circuit, but may be expected to remain constant for different chips of the same design. Detailed studies of gain and noise of the both the ABC130 and ABCStar chips, such as those shown in Section 9.4.1, have confirmed the optimal “Strobe Delay Fraction” to be 0.57.

An additional register allows adjustment of the delay range by changing the current through each inverter stage: three options are provided, designed such that the full 63 steps will cover approximately 50, 60, or 80 ns. The actual size of each delay step is process and temperature dependent: if necessary this may be calibrated by measuring the width of the strobe delay “top hat”, which corresponds to one 25 ns time bin. Hence in Figure 3.13, it can be seen that one Strobe Delay count is approximately 1.25 ns. Note also that due to the temperature dependence of each strobe delay step, it is good practice to adjust the calibration timing at the start of each test session, having first waited for the system to reach thermal equilibrium.

3.3.3 Three Point Gain

Once the timing of the system has been optimised, one typically performs a “quick” measurement of the gain and noise in order to identify any channels with excess noise such that they may be masked off (excluded from readout) in order to avoid interference with subsequent procedures. A series of threshold scans are performed for injected charges of 0.5, 1.0 and 1.5 fC. A linear fit is made to the plot of VT50 as a function of charge to determine the




Figure 3.14: Typical Three Point Gain Plot for a single ABC130 chip

correspondence between charge and threshold. The derivative of this function gives us the gain (dV/dQ) as a function of input charge. This may then be used to convert the measured output noise in mV, from the width of the complementary error function, to input noise in fC, the noise at the amplifier input. A typical Three Point Gain result for a single ABC130 chip is shown in Figure 3.14.

3.3.4 Trimming

Within each front end chip, the amplifier offset (pedestal) and gain will vary from channel to channel due to matching limitations within the ASIC fabrication process. Each ABC130 (and ABCStar) front end channel has a dedicated 5 bit trim DAC which may be used to adjust its offset to compensate for pedestal and gain variation. The operational range of these trim DACs is in turn set by a chip wide 5 bit trim range DAC: the minimum range covers 36 mV and the maximum range 150 mV. The predicted operating threshold in ATLAS is 0.7 fC, therefore the trim DACs would ideally be adjusted in order to best align the channel thresholds at or near this point, using the smallest possible setting of the trim range DAC. The DACs show very little temperature dependence, hence the trimming only




Figure 3.15: A selection of trim plots for one column of strips (128 channels) of a single ABC130 chip. Figures a) and b) give, for Trim Ranges 2 and 8 respectively, the correspondence between the equivalent threshold, VT50, and the Trim DAC setting. Each line represents data for a single channel for an injected charge of 1.0 fC. Figure c) shows, again for an injected charge of 1.0 fC, the number of channels which may be trimmed to target thresholds in the range 0 to 100 DAC counts for trim ranges 2 (black), 4 (red), 6 (green) and 8 (yellow) respectively. Figure d) illustrates, using the same colour scheme, the predicted RMS of the trimmed threshold distribution as a function of possible trim targets for each of ranges 2, 4, 6 and 8. For larger Trim Range settings the range covered by the Trim DACs is bigger, hence all 128 channels of the chip may be aligned across a greater range. However, as the Trim DAC step size has been increased, the RMS of the trimmed threshold distribution also increases.

requires infrequent adjustment to compensate for any (deliberate) adjustment of other front end parameters or for the effects of radiation upon the channel offsets.

In practice the trim DACs are studied as illustrated in Figure 3.15. For a fixed injected charge of 0.7 fC or 1.0 fC, a series of threshold scans are recorded for a series of values to cover the 5 bit range of the channel trim DAC with VT50, the point of 50% efficiency, being calculated for each channel. This is repeated with the chip wide trim range being set to cover increasing voltage ranges. For each combination of channel and trim range, a graph of VT50 as a function of the trim DAC setting is formed: examples are shown in Figures 3.15a and 3.15b. The effect of the trim DAC upon the channel threshold is not linear, hence a second order polynomial is fitted to each channel's data. Across each chip, for each combination of channel and trim range the fitted data is solved to determine the number of channels which may be brought into line for a range of possible target threshold values. The analysis is repeated for each of the studied trim ranges, and the software selects the minimum trim range for which the number of successfully trimmed channels may be maximised. Figure 3.15c shows, for our example chip, the number of channels within a column of 128 strips which may be trimmed to target values in the range 0 to 100 DAC counts for trim ranges 2, 4, 6 and 8. The RMS of the predicted threshold distribution for each of the four ranges, were this chip to be trimmed to trim target x , is shown in Figure 3.15d.

3.3.5 The Response Curve

Once the timing and trimming of the system have been optimised, one may determine the noise of the system by recording each module's response curve. A series of threshold scans are performed for injected charges of 0.21, 0.50, 0.805, 1.01, 1.24, 1.51, 2.00, 3.01, 3.99, 6.02 fC: these values are chosen with a focus on the region below 2 fC, where the amplifier response is approximately linear, with the addition of a few points at higher values. A fit is made to the plot of VT50 as a function of charge to determine the correspondence between charge and threshold: typically an exponential function is used to fit the full response curve as this fits the data better for low charges than, for example, a second order polynomial.




Figure 3.16: Typical Response Curve for a single ABC130 chip

The derivative of the fitted function gives us the gain (dV/dQ) as a function of input charge. This may then be used to convert the measured output noise in mV, from the width of the complementary error function, to input noise in fC, the noise at the amplifier input. A typical Response Curve for a single ABC130 chip is shown in Figure 3.16.

3.3.6 Noise Occupancy as a function of Threshold

It can be shown[33] that, in a binary detector system, the logarithm of occupancy due to noise is inversely proportional to the square of the threshold:

$$\log P_n = \log\left(\frac{\Delta t}{4\sqrt{3}\tau}\right) - \frac{1}{2}\left(\frac{Q_t}{Q_n}\right)^2 \quad (3.1)$$

where:

- P_n is the noise power
- Δt is the time interval
- τ is the time constant of an RC low pass filter




Figure 3.17: Typical Noise Occupancy result for two ABC130 chips

- Q_t is the threshold level (in fC)
- Q_n is the equivalent noise charge (in fC)

Hence the slope of a plot of $\log(\text{noise occupancy})$ as a function of threshold squared (in fC) may be used to determine the noise. A typical result for two ABC130 chips is shown in Figure 3.17. The error bars are bigger at higher thresholds due to the lower occupancy.

3.3.7 Timewalk

In the ABC130 and ABCStar specification documents, timewalk for modules with unirradiated detectors is defined as “the maximum time variation in the crossing of the time stamp threshold over a signal range of -0.75 fC to -10 fC, with the comparator threshold set to -0.5 fC”. Hence the timewalk measurement first runs a threshold scan with an injected charge of 0.5 fC in order to set the threshold, then scans of the calibration delay are made for injected charges between 0.75 fC and the upper limit of the calibration circuit which is 9.0 fC for ABC130 and 9.8 fC for ABCStar. For each of these scans the rising edge of the injected signal, which corresponds to the edge of the peak at higher delay values, is then




Figure 3.18: Typical ABCStar Timewalk Plots (Single Chip Result).

fitted to determine how the timing of the discriminated signal has varied across this range. This is shown in Figure 3.18a. In the case of the final scan, taken with the maximum possible injected charge, both edges are fitted: the width of the peak will correspond to one Bunch Crossing (BCO) clock period of 25 ns duration and hence may be used to calibrate the scale of the strobe delay axis. This is shown in Figure 3.18b.

3.3.8 Double Trigger Noise Occupancy

In a detector system with binary readout, it is very important that the readout signals should not couple back into the front end amplifiers, thereby increasing the noise rate. Whilst the output signals themselves are usually transmitted differentially to reduce the probability that they may influence the sensor and amplifiers, it remains possible that current pulses due to the normal operation of the chip, for example when a trigger is received, may also cause problems. Whilst assembling the ATLAS SCT, which uses an optical data transmission scheme in which logic 0 is represented by “light off”, there were even cases in which the optical

signal was picked up directly in the sensor due to unforeseen breaks in the hermeticity of the optical transceiver packaging [59]: thankfully these problems were identified and corrected before the detector was installed.

The Double Trigger Noise Occupancy Test was developed to look for time dependent changes in the noise rate in a binary detector system which may be correlated with readout activity. Two triggers are sent to the system separated by a controlled, variable time interval. By varying the time of arrival of the second trigger with respect to the first one may explore changes in the occupancy of the front end which may occur at various points in the readout cycle. The sequence is repeated across a range of thresholds to determine the magnitude of any correlated noise signals which may be present.




Figure 3.19: An example of a Double Trigger Noise Test which finds occupancy increases correlated with readout activity. This plot is for a module built with ABCN25 chips mounted on a serially powered stave STAVE250-SP, and was recorded in February 2014 at a threshold of 0.5 fC. An increase in occupancy is observed for trigger separations of 131 and 136 Bunch Crossing (BCO) clock periods.

An example of a Double Trigger Noise result is shown in Figure 3.19. This comes from a serially powered stave STAVE250-SP, at this time loaded with four short strip modules, using the earlier ABCN25 chipset made in 250 nm CMOS technology. The plot shown was recorded for a fixed threshold of 0.5 fC and is for the module closest to the End of Substructure (EoS) board, as introduced in Figure 3.2. The spacing between the two L1A (single level) triggers was varied from 120 to 150 Bunch Crossing (BCO) clock periods with the latency of the

ABCN25 chips set to 131. For this serially powered stave, at this threshold one starts to see occupancy increases for L1A spacing of 132, close to the assigned latency, and 136 BCO periods. STAVE250-DC, a companion stave built using DC-DC converters, did not show this issue: these tests were a key input to the subsequent decision to use DC-DC Point of Load power conversion instead of Serial Powering.

The situation is a little more complicated for ABC130 (and ABCStar) due to their use of a two level trigger scheme: in principle there may be digital activity correlated with both reception of the Level 0 (L0) trigger, which transfers hit data between buffers within the chip, and reception of the Level 1 (L1) or Regional Readout Request (R3) trigger, used to request transmission of data off the chip. Hence we must run tests to look for possible changes in occupancy which may occur as a result of all possible trigger signals, L0, L1 or R3: we must run three separate variants of the DTN test.

3.3.9 Pedestal Trim

The ABC130 has significant gain variation within each die meaning that, for a chip which has been trimmed to maximise threshold uniformity at the nominal operating threshold of 0.7 fC, threshold uniformity is significantly degraded as one moves away from that threshold. The pedestal trim was originally developed as a means to align channel offsets (pedestals) before making noise occupancy or double trigger noise studies, where it may be important to maximise threshold uniformity at 0 fC as this represents the most challenging operating condition for the detector system.

An example of a Pedestal Trim data set is shown in Figure 3.20. For this test, the threshold and chip wide trim range are set to pre-determined values and the channel trim DACs are scanned across the full 5 bit range. No charge injection takes place: for each channel, the point at which the channel records 50% occupancy is the trim DAC value to be loaded to bring the channel's offset into line with the target threshold across the chip. With ABCStar the Pedestal Trim can achieve good threshold uniformity across the operating range, as the gain variation with a chip has been much reduced: this is an attractive option




Figure 3.20: An example of a Pedestal Trim Test for 128 channels of a single ABCStar chip corresponding to one column of strips (the two columns being handled separately by the DAQ system). The plot represents, for fixed Threshold and Trim Range conditions, the occupancy of each channel as a function of its TrimDAC value. It can be seen that each channel reaches 50% occupancy at a different Trim DAC value. Analysis of this data yields a set of Trim DAC values to be written to the chip to bring these “pedestals” into line.

due to its speed.

3.4 Chapter Summary and Conclusions

The system design of the ITk Strip Tracker and in particular the stave concept used in the Barrel region have been outlined; the Module and its key components have been introduced; test procedures which will occur in the following chapters have been described.


Chapter 4

Evaluation of the ABC130/HCC130 chipset


This chapter documents the initial evaluation of ABC130 and HCC130 on single chip test boards, and the wafer screening of ABC130 as required to deliver tested good die to the module construction programme. It can be seen that the initial screening criterion continued to evolve as further experience was gained with the chip.

4.1 Test System Design for ABC130/HCC130 chipset

Drawing upon the model used to test the ABCN25, a custom driver board was made to interface the High Speed Input Output (HSIO) [60] readout board to a single wire bonded or probed ABC130 ASIC, as shown in 4.1b. The driver board is designed around a Xilinx Spartan 3A Field-Programmable Gate Array (FPGA) [61] which is programmed to respond to a number of ABC130 style commands to change the mapping, polarity and direction of fast control signals between the device under test and the data acquisition system. This firmware also drives several slow control signals needed by the chip. The driver accommodates common mode differences between the Low Voltage Differential Signalling (LVDS) outputs offered by the Spartan 3A FPGA, which have a common mode level of 1.2 V, and the Scalable Low Voltage Signalling (SLVS) expected by ABC130, with a common mode level of 600 mV, by lifting the ground of the ABC130 by 600 mV with respect to the ground of the driver



(a) ABC130 Single Chip Board



(b) ABC130 Driver

Figure 4.1: ABC130 Single Chip Board and Driver

such that the mean level is correct for each chip. Finally a number of commercial ADCs are fitted to the driver to monitor several aspects of the device under test, such as the output of its threshold and bias DACs, and the magnitude of signals coming from its current programmable SLVS outputs.

The single chip board shown in Figure 4.1a was designed such that a miniature strip sensor could be connected to the ABC130 to facilitate early studies of signal to noise performance

with particles of known energy from laboratory sources. A later version of this PCB added a hole under the sensor to eliminate unnecessary support material. In addition the board allows known capacitive loads to be added to two front end channels in order to measure the dependence of the noise upon load capacitance. Design and manufacture of both the driver board and single chip PCB took place at Cambridge, with additional batches of single chip PCBs being made at Oxford. The single chip board has also been used to facilitate the irradiation of the chip using x-rays and proton beams.




Figure 4.2: ABC130 Probe Card and Probed Die

Replacing a faulty ASIC on a hybrid or module is a time consuming and potentially risky operation. It is therefore necessary to screen each wafer of chips such that defective die may be excluded from the production chain. The ABC130 probe card is shown in Figure 4.2. This is a custom printed circuit board fitted with an epoxy probe ring by Probe Test Solutions Limited (PTSL) [62]. The probe card is pin compatible with the ABC130 single chip board such that the driver firmware and test software may be developed and debugged with wire-bonded chips.

In addition to the needles required to service the chip, the ring includes a probe contact switch. This comprises a pair of tungsten probes of which only one, the “mobile” needle, can touch the wafer. The second, “static” needle is not a probe in the usual sense in that it maintains a fixed, rotated aspect in which it is normally in contact with the mobile needle to make a closed circuit. The mobile needle is designed to be the first point of contact with the wafer, touching down on the scribe line between die around $100\text{ }\mu\text{m}$ earlier than the

other probes. As it touches down the mobile probe moves away from the static probe and the circuit is broken. In reality there will be some variation in the resistance between the two probes before they become fully open, hence the driver actually measures the resistance of this contact. Although the reproducibility of the measurement can be poor, the probe contact switch still provides useful information whilst setting up the system¹.

4.2 First Batch of ABC130 wafers




Figure 4.3: 3D Electron Micrograph showing correction of ABC130 by Focussed Ion Beam. The hole is approximately 10 μm by 10 μm . (Credit: Nanoscope Services Ltd. [63])

A first batch of four ABC130 wafers was received in November 2013. Having diced the first wafer and mounted samples onto single chip PCBs, it was quickly established that the chips gave no output. The designers were able to identify the cause of this problem: the polarity of an enable line in the bidirectional SLVS transceiver was incorrect. A small number of die were corrected by NanoScope Services Ltd. [63], specialists in the repair of integrated circuits by Focussed Ion Beam correction. To effect this repair, several layers of oxide and metal were removed from the chip to allow access to the erroneous trace. The relevant track was then cut, and a tungsten strap was placed from the newly disconnected direction node to the power or ground plane to permanently enable or disable transceiver output. Each

¹During production, it is planned to monitor the contact switch, but not to rely upon it.

correction was completed within a region approximately 10 microns by 10 microns, as shown in Figure 4.3. In this manner we were able to obtain first data from the chip and verify other aspects of the design.

With the screening of future wafers in mind, the corrected chips were used to develop a sequence to test all key features of the chip. Firstly the chip is powered, its analogue and digital currents measured, and the function of its internal current shunt (for serial powering) checked² Secondly the function of the chip’s integrated Digital to Analogue Converters (DACs) is studied, using commercial ADCs on the driver board to measure DAC output over a selection of possible values. The third stage executes a series of test vector blocks designed to test the digital functionality of the chip: reading/writing registers, passing data through in both directions, flow control etc. Finally a brief analogue characterisation is performed comprising a three point gain measurement as described in subsection 3.3.3 followed by a quick check of the trim DAC functionality, used to compensate for pedestal offsets in the front end amplifiers as explained in subsection 3.3.4.

4.3 Second Batch of ABC130 wafers

A second batch of five ABC130 wafers, in which the known bug had been corrected, was received in June 2014. A Micromanipulator Model 4460 [64] probe station was used to test these wafers. Two operational issues were not understood at the time. Firstly for some wafers, die with the Fast Cluster Finder (FCF) block as described in subsection 3.2.2 had high failure rates for digital test vectors involving the passing of simulated event data from HSIO through the chip and back to HSIO. Although suggestive of marginal timing differences between FCF and non-FCF die, at the time this was worked around by simply repeating the test until it was passed. Secondly the characterisation of the chip’s variable current SLVS outputs would yield an approximately equal mix of positive and negative slopes. This issue

²Although the DC-DC powering scheme will now be used, the shunt functionality shall be retained for production modules such that the load may be adjusted to partially compensate changes in currents due to radiation effects.

Wafer	nA die	nB die	n Rejects	A Yield	A+B Yield
A5Q2XLH	122	12	1	90.4%	99.3%
A8Q2YOH	124	10	1	91.9%	99.3%
A6Q2XKH	114	15	6	84.4%	95.6%
AYQ2XTH	105	23	8	77.0%	94.1%
AQQ3OGH	121	11	1	89.6%	97.8%
Total	586	71	17	86.8%	97.3%

Table 4.1: ABC130 Probing Yield, Batch 2

was accommodated by accepting both positive and negative slopes as valid results. (Note that both issues were later fully resolved as described in subsection 4.4.1.)

Whilst the analysis and classification procedures proposed for production are described later in this chapter, a different methodology was used for these early wafers. Die were automatically classified as rejects if they failed power or digital tests, or had significant analogue defects such as low gain or high noise. For each wafer the RMS of the distribution of each DAC parameter was calculated, and die with one or more parameters more than 4 x RMS away from the mean were highlighted for manual classification based upon plot inspection. In addition the SLVS output plots were visually inspected for every die. A-grade die passed all tests successfully; B-grade die had up to 2 bad or noisy channels, or exhibited slightly anomalous DAC behaviour, but may still be useful during the R&D phase of the project. Die were additionally classified as type “0” or “1”, according to whether or not they contain an FCF block. This information was incorporated into the pick map such that the dicing vendor may respect the four useful streams. The yield of each wafer is shown in Table 4.1. Category “A” yield was found to vary between 77.0% and 91.9%.

4.4 Third Batch of ABC130 wafers

4.4.1 First 14 wafers (Batch 3A)

A third batch of 24 ABC130 wafers was received in August 2015, and again tested using the Micromanipulator probe station³. Whilst readout of ABC130 chips using HSIO remains possible, in 2015 the board of choice within the community was the cheaper, commercially available ATLYS [67] board, which was therefore used during the wafer test campaign. The ATLYS board and firmware build were found to have different test vector timings compared to the previous HSIO version, however this was easily accommodated.

The causes of the two readout issues left unresolved during earlier tests were identified and corrected whilst running the first few wafers of the second batch. The second wafer tested, A7Q8IKH, had a high number of “FCF” die failing the digital tests in which data packets were passed to the chip from ATLYS. To resolve this issue the ATLYS firmware was revised to permit the adjustment of the phase of output signals in addition to the phase of input signals. The output phases were duly scanned to determine the operating window common to both ABC130 types, and set to the middle of the window. Whilst from the third wafer AEQ8FXH onwards this issue has been resolved, this does highlight subtle timing differences between FCF and non-FCF types.

The fifth wafer tested, AIQ8I9H, initially exhibited a high number of SLVS outputs seemingly switching between the usual states having positive and negative slopes. Upon investigation the causal factor was that the direction of the bidirectional ports on ATLYS was not being set correctly. This was corrected in software and the affected die retested. From this point onwards, all die tested had SLVS DAC measurements with consistent slope polarities, making automated offline analysis a realistic prospect.

The development of the offline analysis and the cuts which it uses are documented in

³By this time the control electronics of the Micromanipulator probe station had failed, so it was refurbished using a selection of Commercial Off The Shelf (COTS) parts including a motion controller [65] and stepper drives [66] from Arcus Technology.

Wafer	nA die	nB die	n Rejects	A Yield	A+B Yield
A7Q8IKH	112	11	12	83.0%	91.1%
AGQ8DWH	108	16	11	80.0%	91.9%
AEQ8FXH	110	18	7	81.5%	94.8%
APQ8E5H	114	14	7	84.4%	94.8%
AIQ8I9H	119	6	10	88.1%	92.6%
A4Q8INH	110	9	16	81.5%	88.1%
AZQ8FCH	119	7	9	88.1%	93.3%
ANQ8HMH	117	9	9	86.7%	93.3%
ACQ8EHH	113	12	10	83.7%	92.6%
A8Q8IJH	119	13	3	88.1%	97.8%
ALQ8I6H	111	14	10	82.2%	92.6%
ACQ8IFH	115	11	9	85.2%	93.3%
AUQ8HGH	112	10	13	83.0%	90.4%
AKQ8E9H	110	17	8	81.5%	94.1%
Total	1477	156	122	84.2%	93.0%

Table 4.2: ABC130 Probing Yield, Batch 3A. Note that wafer A7Q8IKH was retested as part of batch 3B and is hence excluded from the totals above.

Appendix A. Running this analysis on all wafers gives the yield figures shown in Table 4.2. The average category “A” yield across all wafers is 84.5%, and the average category “A+B” yield is 93.2%. This compares favourably to earlier designs made at other technology nodes, yet sufficient die are still rejected so as to justify the screening process.

4.4.2 Remaining 10 wafers (Batch 3B)

Probing of the remaining wafers of this batch began in October 2016 in anticipation of the go-ahead being given for the community to begin module assembly for the stave and petal projects. By this time the module programme had given us an indication that the default bias settings for the ABC130 chip may not be correct. When this first came into question it became obvious that the DAC fit equations stored with the wafer test data often made little or no sense. To follow this up a study was made which compared DAC characterisation data for a few chips tested first on the probe station but later retested on PCBs, and found that the agreement was very poor. This was attributed to the variability of contact resistance when testing the wafers, combined with the lack of a sense wire to make




Figure 4.4: The distributions of the offset of the IR8B DAC before (left) and after (right) modifications to sense the true ground potential of the ABC130 chip.

an accurate measurement of local ground within the chip under test.

The issue was discussed with the ASIC designers who were able to suggest a way forwards. One pad on the ABC130 is connected to the chip’s internal ground plane through a 90 k Ω test resistor. Whilst this had been intended as a process monitor, by coupling this node through a voltage follower into a spare ADC channel it was possible to obtain an accurate measurement of the chip’s internal reference plane. This worked well for bonded chips but moving to the probe station it was noted that for some chips this potential was not constant for the duration of the tests, again due to poor contact. It was experimentally determined that such chips typically gave better contact when tested for a second time, so an automated loop was added to cycle the probe contact in cases where the ground potential moved by more than 10 mV during the test. This largely cured the problem. A further contribution to these measurement errors came from a lack of settling time between the chip being powered and the measurements being taken which was eliminated by means of a short delay.

The effect of these improvements may be illustrated by Figure 4.4. The left hand plot shows the offset of the IR8B DAC as measured for the 14 wafers of batch 3A using the original procedures and the right hand plot shows the same parameter as measured for the 10 wafers of batch 3B, once the identified causes of systematic errors had been eliminated. It can be seen that the mean value has shifted by approximately 12 mV, this corresponds to the mean difference between the true “ground” potential within the chip and the potential of the previous reference on the driver board. In addition the RMS of the distribution is

reduced from 2.8 to 2.1 mV. A software bug regarding the measurement of the comparator bias DAC was also corrected, which was buffered on the driver board by an op-amp set to have a gain of two as opposed to the unity gain of all other channels.




Figure 4.5: Typical Batch 3B LDO Characterisation. For low settings the output voltage has a constant value which is largely defined by the voltage delivered to the chip. When the programmed target first falls below the delivered voltage, a “knee” is observed, beyond which the characteristic falls linearly.

In order to increase the speed of the LDO Characterisation sequences, the software was updated to program the LDO control DACs using unary coding, as the designers had intended. A typical result is shown in Figure 4.5, in which the x-axis now shows the number of approximately equal currents being summed by the unary DAC as explained in Section 4.4.1. With only 17 possible values to test the sequence is much shorter: the downside is that the accuracy of adjustment is not as good, but with a worst case error of 20 mV it is good enough.

Now that the DAC characterisations were known to be reliable, an extra section was added to the test sequence to tune the DACs online to set the various bias voltages within the chip to give the correct nominal values, as specified by the ASIC design team, on the relevant monitor nodes. The strobe delay (timing of the charge injection pulse) is then optimised for the die and its gain and noise is measured using the tuned bias parameters. The new procedure was used to test the wafers of this batch using the refurbished MicroManipulator probe station. One wafer held over from batch 3A was also retested with the updated procedure, with screening being completed in January 2017.

There were further subtle developments as the run progressed. Each wafer has its gain




Figure 4.6: Batch 3B: The distributions of tuned DAC values and the resulting monitor voltages for parameter VREF.

Parameter	Nominal DAC	Nominal V	Mean DAC	RMS (DAC)	Mean V	RMS(V)
VREF	13	50	13.36	0.72	50.06	0.52
IREF	13	50	13.27	0.69	50.04	0.58
IR8B	13	50	13.22	0.68	50.04	0.56
VCOM	16	36.5	16.59	1.16	36.57	0.44
IPRE	16	123	15.58	1.08	123.0	2.9
IFEED	8	30	9.40	0.57	29.99	1.01
	-	mV	-	-	mV	mV

Table 4.3: Batch 3B: The distributions of tuned DAC values and the corresponding monitor voltages, over a population of 1295 chips.

and noise measured twice: once for “default” DAC tunings and once for the optimised values determined online. Upon completion of the first wafer it was obvious that the distributions of the tuned values were centred upon the designer’s predictions and not the values used previously, hence for subsequent wafers the default settings were updated to match the centre of the measured distributions. From the fourth wafer onwards an additional test procedure was added to make a direct measurement of the low and high DC levels on the calibration bus for a nominal 1fC charge.

The distribution of tuned DAC values and the corresponding voltage measurements for bias parameter VREF are shown in Figure 4.6 for all wafers of batch 3B⁴. Together with the other bias parameters, the mean and RMS values of these distributions are summarised in Table 4.3. It can be seen that there is good agreement between the measured distributions

⁴The number of samples in Figure 4.6 is lower than the number of samples in Figure 4.4 as, by this point in the sequence, testing has been discontinued for die which failed the digital test vectors.




Figure 4.7: Batch 3B: Comparison of mean input noise values under nominal (left) and tuned (right) conditions.

Parameter	Mean	RMS	Low Cut	High Cut	Fail Code	Category
VREF	13.36	0.72	11	16	0x0001	X
IREF	13.27	0.69	11	16	0x0002	X
IR8B	13.22	0.68	11	16	0x0004	X
IPRE	15.58	1.08	13	18	0x0010	X
IFEED	9.40	0.57	8	11	0x0020	X
Strobe Delay	21.0	2.1	16	27	0x0040	X

Table 4.4: Batch 3B: Cuts on optimised bias DAC and Strobe Delay settings.

and the expected nominal values for all parameters with IFEED being the only parameter for which the mean DAC setting is not within half a count of the expected nominal value.

The combined benefits of DAC and Strobe Delay tuning is illustrated by Figure 4.7. The left hand figure shows the measured noise distribution for 1167 functional die operated under nominal conditions, and the right hand figure repeats this for the same sample under tuned conditions. Not only is the mean input noise value reduced from 468 to 452 electrons, but the RMS of the distribution falls from 33 to 27. Hence it can be seen that the optimisation is having a beneficial effect. (Note that wafer ADQ8IEH has been excluded from these histograms as the nominal conditions were redefined after this had been tested.)

The revised operating conditions for the chip mean that some of the cuts used previously are no longer valid, notably those on the slope of the calibration and threshold DACs. Acceptance cuts of these two DAC slopes together with the bandgap voltage and the tuned LDO output voltages are shown in table 4.5 together with the summary of the distributions from which they have been derived. In addition, the cut upon the range of the trim DAC

Signal	Parameter	Mean	RMS	Low Cut	High Cut	Fail Code	Category
Bandgap	Voltage	593.2	33.1	550	630	0x0001	X
VDDA (tuned)	Voltage	1200	8.5	1180	1220	0x0002	X
VDDD (tuned)	Voltage	1200	8.6	1180	1220	0x0004	X
Thresh. DAC	Slope	0.520	0.037	0.46	0.59	0x0020	X
Calib. DAC	Slope	0.567	0.047	0.47	0.65	0x0040	X
		mV	mV	mV	mV		

Table 4.5: Batch 3B: Additional cuts on DAC settings.

Wafer	nA die	nB die	nRejects	A Yield	A+B Yield
A7Q8IKH_retest	119	8	8	88.1%	94.1%
ADQ8IEH	117	13	8	86.7%	94.1%
AEQ8IDH	114	13	8	88.4%	94.1%
A5Q8IMH	119	12	4	88.1%	97.0%
A9Q8IIH	117	10	8	86.7%	94.1%
AFQ8FWH	105	15	15	77.8%	88.9%
ABQ8EIH	122	7	6	90.4%	95.6%
AJQ8HRH	122	5	8	90.4%	94.1%
A5Q8F6H	114	14	7	84.4%	94.8%
ARQ8E3H	120	7	8	88.9%	94.1%
ASQ8E2H	120	11	4	88.9%	97.0%
Total	1289	115	84	86.8%	94.5%

Table 4.6: ABC130 Probing Yield, Batch 3B. Wafer A7Q8IKH, initially tested as part of batch 3A, is included in these totals.

over which all 128 channels may be trimmed has been reduced from the previous limit of 5 counts to 4 counts, and the value for input noise above which single channels are considered bad has been reduced from 750 to 700 electrons. Applying the cuts as described above to the 10 wafers of batch 3B the mean category “A” yield is 87.0% and the mean category “A+B” yield is 92.6% as shown in Table 4.6.

4.5 Fourth Batch of ABC130 wafers

A further batch of 48 ABC130 wafers arrived at The Rutherford Appleton Laboratory in September 2017. All wafers were tested using the MicroManipulator machine. Some small procedural changes were necessary: firstly the power supply voltages for VDDA and VDDD

were increased from 1.50 V to 1.55 V in order to compensate for increased voltage drops in newly extended cabling. Secondly it was noted that contact resistance was in many cases reduced by breaking and re-making contact: this was duly automated. Finally it was noted that after the testing of each ABC130's integrated shunt transistor, which forces a high current through the VDDD_RAW and GND needles, the delivered voltage under normal conditions often increased, indicating further reduced contact resistance. Hence the test sequence was again modified to perform this shunt test earlier in the sequence in cases where the delivered voltage was deemed to be marginal (VDDD_REG less than 1.2 V before the chip is configured).

Applying the usual cuts, the yield for each of these wafers is presented in Table ???. The setup wafer VECR3ZH has a lower yield than usual: in part this is due to a system malfunction which led to the needle array being scraped across 11 die⁵. Wafer VVCQ22H also has a slightly lower yield than is typical, however the worst wafer of the batch is clearly wafer VVCQ5IH which has more than 40% bad die, generally failing the initial power tests or the digital test vectors. Whilst this might be indicative of a contact problem, it seems to have been confined to this wafer, as when the next wafer was loaded the yield returned to acceptable levels. Across this series of wafers, the mean category "A" yield is 89.0% and the mean category "A+B" yield is 93.9%.

4.6 Yield across all batches

The yield across all batches is summarised in Table 4.8. A total of 77 wafers were probed, yielding a total of 9121 category A (good) die. A further 658 die were placed into category B (minor defect) whilst 618 die were rejected. As regards the yield across batches, less stringent cuts were in place for batch 2 than for batch 3, which explains the initial fall in yield. Thereafter the category "A + B" yield remains more or less constant. The fraction

⁵The causes of this incident - a combination of a faulty USB hub and a software bug that failed to trap the resultant error - were eliminated before work resumed, and the damaged probe card was replaced by a spare.

Wafer	nA die	nB die	nRejects	A Yield	A+B Yield
VECR3ZH	110	5	20	81.5%	85.2%
VLCR68H	129	4	2	95.6%	98.5%
VYCQ5FH	125	3	7	92.6%	94.8%
VTCQ62H	128	2	5	94.8%	96.3%
VFCQ8EH	124	7	4	91.9%	97.0%
VTCQ24H	125	8	2	92.6%	98.5%
VXCQ20H	125	5	5	92.6%	96.3%
V5CQ58H	117	9	9	86.7%	93.3%
VJCQ7TH	124	6	5	91.9%	96.3%
VPCQ66H	125	6	4	92.6%	97.0%
VSCQ63H	126	4	5	93.3%	96.3%
VNCQ67H	122	9	4	90.4%	97.0%
VRCQ26H	125	6	4	92.6%	97.0%
VOCQ5DH	120	4	11	88.9%	91.9%
VZCQ5EH	125	6	4	92.6%	97.0%
VJCQ5UH	122	5	8	90.4%	94.1%
VICQ1XH	126	6	3	93.3%	97.8%
VVCQ22H	109	6	20	80.7%	85.2%
VVCQ5IH	71	8	56	52.6%	58.5%
VKCQ7SH	118	7	10	87.4%	92.6%
VPCQ7NH	122	6	7	90.4%	94.8%
V7CQ2QH	125	7	3	92.6%	97.8%
VCCQ70H	126	4	5	93.4%	96.3%
VDCQM9H	114	5	16	84.4%	88.1%
VLCQM1H	120	5	10	88.9%	92.6%
V0CQMMH	119	9	7	88.1%	94.8%
VZCQMNH	119	3	13	88.1%	90.4%
VICQLMH	126	3	6	93.3%	95.6%
VGCQM6H	117	6	12	86.7%	91.1%
VXCQQ6H	120	7	8	88.9%	94.1%
VLCQSGH	124	5	6	91.9%	95.6%
V1CQMLH	122	9	4	90.4%	97.0%
VCCQMAH	121	8	6	89.6%	95.6%
V7CQLYH	120	7	8	88.9%	94.1%
V2CQMKH	111	7	17	82.2%	87.4%
VFCQLQH	119	11	5	88.1%	96.3%
VDCQLSH	117	4	14	86.7%	89.6%
V6CQMGH	114	14	7	84.4%	94.8%
VACQLVH	122	7	6	90.4%	95.6%
V8CQLXH	118	12	5	87.4%	96.3%
VTCQQAHA	122	7	6	90.4%	95.6%
V6CQLZH	125	4	6	92.6%	95.6%
VBCQLUH	122	10	3	90.4%	97.8%
V9CQQUH	121	10	4	89.6%	97.7%
V5CQMHH	123	7	5	91.1%	96.3%
VCCQLTH	122	9	4	90.4%	97.0%
VYCQQ5H	119	8	8	88.1%	94.1%
VMCQM0H	123	6	6	91.1%	95.6%
Total	3000	183	192	89.0%	93.9%

Table 4.7: ABC130 Probing Yield, Batch 4

Batch	nWafers	nA die	nB die	nRejects	A Yield	A+B Yield
2	5	586	71	17	86.8%	97.3%
3A	13	1477	156	122	84.2%	93.0%
3B	11	1289	115	84	86.8%	94.5%
4	48	5769	316	395	89.0%	93.9%
Total	77	9121	658	618	87.7%	94.1 %

Table 4.8: ABC130 Probing Yield across All Batches

of good die (category A) is at its lowest for batch 3A, after which it recovers due to the introduction of tuning procedures to optimise the bias DAC and strobe delay settings.

4.7 Evaluation of HCC130




Figure 4.8: HCC130 Single Chip Board

The ABC130 driver board was made with the additional requirement to test HCC130 in mind. In terms of the digital functionality of the device, the requirement is broadly the same, a number of differential IO lines at SLVS levels, however sufficient lines are provided to connect only the left or right handed IO bank at any one time. The driver board's commercial ADCs are used to monitor adjustment of the HCC130's internal regulator, and a dual channel DAC was included to facilitate calibration of the autonomous monitor's internal ADC. A single chip board, shown in Figure 4.8, was designed and made at Oxford

in order to support initial tests. Whilst the option to make an HCC130 probe card for use with the driver board was considered, HCC130 probing activities were later taken up by our US collaborators, who chose a different route.

Early tests concentrated upon the digital functionality of the chip, with a couple of significant errors being found with the HCC130 documentation, notably the polarity of the ePLL's clock enable register. Very basic tests were made to confirm operation of the autonomous monitor block, but once HCC130 was known to be a functional design we quickly moved on to tests of the first HCC130/ABC130 hybrid. However the HCC130 single chip board was used to support x-ray and proton irradiation of this chip.

4.7.1 HCC130 Irradiation Programme (UK)

As outlined in section 2.1.3, the increase in leakage current within our front-end chips as a function of Total Ionising Dose (TID) is a major concern. Hence the focus of the HCC130 irradiation programme was to record the evolution of the current drawn by the chip as a function of dose. In the UK one chip was irradiated using the Siefert X-ray set at RAL, and one chip was irradiated with protons at the Birmingham Cyclotron facility. A complementary study at low dose rates was later initiated by US collaborators using a ^{60}Co source at Brookhaven National Laboratory, but this uses completely different hardware and software from that described here.

During the UK study, the beam was stopped at regular intervals in order to run the test sequence under constant conditions. The sequence starts by (re)configuring the chip to default settings, and recording the current. The chip's single LDO is then scanned in unary fashion, recording the current drawn by means of the power supply, three voltages by means of the ADCs on the driver board (VDDREG, VHALF, VBANDGAP) and two of the same voltages by means of the HCC130's integrated Autonomous Monitor (AM) ADC (VDDREG and VHALF). The LDO is then set to deliver the desired 1.2 V supply, before the drive current of each SLVS output is scanned. Next a commercial DAC on the driver board is used to recalibrate two of the AM's ADC channels under both the original and newly

optimised LDO settings. Finally the current drawn by the chip is recorded once more, again under both the original and newly optimised LDO settings. One difference in methodology between the x-ray and proton studies is that the shutter of the x-ray set may be controlled directly by the test software, making the procedure largely automated. During irradiation at Birmingham, the beam must be stopped manually each time it is desired to run a test.




Figure 4.9: Current increase factor for HCC130 as a function of Total Ionising Dose (TID) during x-ray irradiation (from the ATLAS ITk Strip TDR). The chip has a single supply rail hence it is not possible to distinguish between current drawn by the analogue (Autonomous Monitor) and digital parts.

The development of the current drawn by the HCC130 during X-ray irradiation is shown in Figure 4.9. It can be seen that the current increase is much smaller than is the case for ABC130, which is as expected as the larger ABC130 chip contains more digital functionality

⁶ Note that the star architecture of the future chipset requires that HCCstar contains much more memory, so this may be expected to exhibit a greater increase.

4.8 Chapter Summary and Conclusions

The ABC130 Driver Board and its companion Single Chip Boards for the ABC130 and HCC chips facilitated a wide range of studies of the two chips. Whilst the first submission of ABC130 gave no output due to a hardware bug, this was corrected in a second submission.

⁶In the case of ABC130 another significant contribution to TID induced leakage current comes from an NMOS transistor, part of the integrated current shunt, which uses a thicker oxide rated to higher voltages. Whilst exactly the same block was later used in the ABCStar prototype, for production chips the shunt has been redesigned to omit this part.

A total of 77 ABC130 wafers were tested, yielding a total of 9121 good die.

Chapter 5

ABC130 Noise Performance

This chapter documents noise results from measurements of ABC130 chips built into a variety of devices: Single Chip Boards (SCBs) with capacitor loads or mini sensors; hybrids with mini sensors (also known as DAQloads), and early barrel modules. In particular the beneficial effect of the revised DAC tuning derived during wafer probing measurements and the variation of the number of strobed channels are explored.

5.1 Devices Tested

5.1.1 Prototype ASIC

Before the ABC130 was submitted to the foundry, a test chip was made in the same 130 nm technology in order to evaluate the performance of the proposed front end design [68]. This “prototype” chip has 32 front end channels which are identical to those in ABC130, however the complex back end and integrated DACs of the latter design are omitted. Instead bias voltages and currents are derived externally and fed into the prototype by means of wire bonds, and a simple CMOS output is provided for each front end channel. The test system generates the calibration pulse by means of a commercial pulse generator, however the 60 fF calibration capacitors through which charge is injected remain on-chip.

It is usual to characterise the performance of a Strip Detector front end amplifier by measuring its equivalent noise charge as a function of load capacitance, as the series contri-

bution to the noise scales directly with the added capacitative load, and the test system for the prototype was made with this requirement in mind. The test board permits the addition of each of a series of commercial SMD capacitors to be soldered into place, connected between one front end channel and a ground pad on the chip. Results from the prototype design set the benchmark against which the noise of the completed 256 channel ABC130 chip, with or without loads, are to be compared.

5.1.2 Single Chip Boards

The Single Chip Board (SCB) for ABC130 and its companion driver board were introduced in section 4.1. These boards were widely used during ABC130 radiation studies, with more than a hundred boards being made, although the additional complexity of the driver board led to the use of these boards being confined to a limited number of institutes. However this section outlines the ways in which the SCB was used to deliver the noise measurements needed to understand the performance of the ABC130 front end amplifier.

In order to measure noise as a function of load capacitance, traces are provided to connect discrete surface mount capacitors directly between each of two of the chip's input channels and the reference ground (on the chip). In addition the layout of the board includes keep out zones in the power and ground layers in the vicinity of these capacitors in order to minimize stray capacitance. The SCB design also includes the option to mount a single miniature strip sensor which may be bonded to 100 of ABC130's input channels. Indeed one board was fitted with two mini-sensors in series, with 50 of the 100 channels bonded to the first sensor being bonded to the second one. As the bulk and interstrip capacitance of fully depleted mini-sensors is well known, such measurements can provide additional input to the "Noise Slope" graph using more realistic loads.

5.1.3 Hybrids and DAQloads

The first Short Strip Barrel hybrids were ordered in June 2014: 12 hybrids of each of two “handed” designs were manufactured. For reasons of economy these were made as individual circuits to a thin FR4 build rather than the polyimide flex circuits envisaged for future use. Population of the first such hybrid was delayed until ABC130, HCC130 and the DAQ system had been fully debugged at RAL using single chip boards. To first order the hybrid worked as expected, with some uncertainties regarding the measurement of gain and noise as discussed later in Section 5.3.

In order to further probe the characteristics of the ABC130 front end and to distribute knowledge of test procedures amongst the community without building complete strip modules, the concept of “DAQloads” was developed. A DAQload is a partially loaded hybrid attached to a “Hybrid Carrier” PCB upon which up to two miniature strip sensors may be fitted. An example built with a single ABC130 / sensor pair is shown in Figure 5.1. In this way detailed studies may be made using the HCC130/ABC130 chipset without the need of a full size sensor, an expensive component with restricted availability. Furthermore as the external connectivity of the Hybrid Carrier PCB is the same as the Module Frame PCB, groups which have successfully read out a DAQload are in principle equipped and ready to read out a complete module.

A major difficulty faced when trying to get early hybrids, DAQloads and later modules to run was related to their use of two and four channel capacitor arrays to AC-couple the three differential control signals. It transpired that such arrays have an undocumented parasitic capacitance between elements, leading to crosstalk between the control signals to the point at which the working margin was too small when using standard LVDS drivers. To be able to operate hybrids and modules of this design reliably an alternate buffer circuit was made using drivers to a related standard, Multidrop LVDS (M-LVDS): designed for bus applications these have a slower rising edge and increased drive current as compared to LVDS devices. Signal integrity over longer cable runs was further improved by the manufacture of adaptors allowing the lossy ribbon cables to be replaced by low capacitance network cables.




Figure 5.1: ABC130 Barrel “DAQload” DLRH5, comprising a Right-Hand hybrid fitted with one HCC130 chip and one ABC130 chip, attached to a “Hybrid Carrier” PCB upon which one ATLAS12A miniature sensor has been mounted. The carrier PCB has the same external connectivity as the Module Frame PCB used to test Barrel Modules.

These components were successfully used to operate DAQloads and modules with 5 m cables in testbeams at CERN and at DESY, and with 20 m cables during proton irradiations at CERN.

As a long term fix the possibility to DC-couple the differential control signals upon entry to the hybrid was considered, however the receivers on an unpowered HCC130 chip were found to pull down the inputs, affecting the operation of other hybrids which may be connected to the same bus. Instead the hybrids were redesigned to make use of individual coupling capacitors instead of the previous arrays. In this way the parasitic capacitance between array elements was eliminated and the control signals could still be AC-coupled to give protection against abnormal DC loading of the buses caused by unpowered or faulty HCC130 die. In addition to this the ribbon cable pinout was revised to reduce crosstalk between signals by grounding every second pair, with adaptors being used to maintain compatibility with existing parts. Finally a new buffer board was designed according to the

updated pinouts, based upon a custom SLVS buffer chip [69].

5.1.4 First Short-Strip Module

Short-Strip module SS1 was assembled and first tested in April 2016 with the power board being added a couple of weeks later. Once the power board was added, additional noise was seen in the strips in its immediate vicinity. As the DC-DC converter on the powerboard operates at a switching frequency of 2 MHz, which is within the bandpass of the ABC130 front end amplifier, it is possible for the electric and magnetic fields to inject signals into the sensor if not properly shielded. The shielding can is intended to eliminate this effect: the magnetic field induces eddy currents within the metallic box which cancel out the field, and the metal also functions as an electrostatic shield. However in this instance the shield had been soldered to the PCB only at its corners, leaving a small gap through which the magnetic field may escape. To validate this theory we took an identical power board and placed it, powered up, on the sensor at one corner of the module. Again additional noise was seen. Next the seams between the shield and the PCB were fully soldered and the excess noise went away. This soldering was then carefully repeated on the power board which had already been glued to the module. Whilst due to access restrictions it was not possible to fully seal one side of the box, the other side was sealed and the extra noise observed to that edge of the shield was seen to go away. This important result confirms that it is possible to glue a DC-DC converter onto a silicon microstrip module without changing the front end noise ¹.

5.2 Initial DAC tuning with nStrobed = 256

Figure 5.2 plots noise as a function of capacitance using, for ABC130, the initial DAC tunings, and injecting charge into all 256 channels simultaneously. (These DAC tunings

¹In the production phase, the field emissions of each powerboard will be measured before use to ensure that leaky converters are never mounted onto sensors.




Figure 5.2: Noise as $f(\text{capacitance})$ with initial ABC130 DAC tunings. For all ABC130 measurements, nStrobed, the number of channels into which charge is simultaneously injected, is 256.

were later found to be incorrect due to a lack of compensation for ground rise occurring within the test system, as explained in Section 4.4.2.) The blue line with triangular markers and the light green line with square markers represent measurements of the ABC130 front end prototype chip using discrete load capacitors, injecting positive and negative charges respectively. The difference between the two lines is due to the design's use of active feedback which, when used with negative signals as expected from n on p sensors, leads to signal compression². The red line with circular markers represents similar measurements of ABC130 using discrete load capacitors and negative charge injection, which progressively deviates from the prototype measurements as the capacitive load is increased. The yellow plus signs denote measurements made using discrete capacitors added to a hybrid and fall further above the prototype data. The remaining markers represent data from a variety of Short Strip modules, Long Strip modules, and ABC130 Single Chip Boards: again, the deviation from expectation (the prototype results) is significant. This includes Short Strip Module SS1:

²The later ABCStar front end design has resolved this issue by using simple, resistive feedback.

its noise with or without the converter, whilst consistent between the two configurations, was higher than the ASIC designers had anticipated. Long strip modules with 4.8 cm strip length deviated from the expected curve even more.

5.3 Revised DAC tuning and reduction of nStrobed

Once the systematic improvements in DAC characterisation at wafer test had been made, as described in Section 4.4.2, noise measurements made under the revised bias conditions showed a small yet significant improvement. Attention then turned to something which had been attempted before, seemingly without benefit, namely to vary the number of channels being strobed simultaneously during the measurement cycle. To speed up the measurement process, the baseline had been to strobe all 256 channels at once, and to send triggers to the module in bursts of 16. A reduction in the recorded noise was seen whilst strobing reduced numbers of channels when running with single trigger cycles, which is painfully slow, but not with bursts of 16 triggers. Software bugs were found and corrected: both trigger options now yield similar results. The software was also updated such that 1 in $\{1,2,3,4,5,8,16,32\}$ channels may now be strobed during any of the common calibration tests.

A variety of devices were now retested to see how the measured noise varied as a function of the number of strobed channels. Whilst unloaded channels show a small effect, channels bonded to one or two mini sensors (8 or 16 mm respectively), short or long strips (24 or 48 mm) show an increasingly large effect. The designer of the front end amplifier has suggested that this effect may be due to a reduction in dV/dt of the calibration signal as the load capacitance is increased: the full charge does not arrive within the 20 ns shaping time of the amplifier, hence the effective magnitude of the calibration pulse is reduced by ballistic deficit. This leads to an underestimation of the true gain, and hence an overestimation of the input noise. The effect is summarised in Figure 5.3.

Figure 5.4 plots noise as a function of capacitance using the corrected DAC tunings and with a reduction in nStrobed, the number of channels into which charge is simultaneously




Figure 5.3: Fractional Gain Increase as $f(\text{number of strobed channels})$, normalised to the 256 channel case. Data are shown for a variety of loads.




Figure 5.4: Noise as $f(\text{capacitance})$ with revised ABC130 DAC tunings. For the ABC130 measurements made using discrete capacitors, nStrobed, the number of channels into which charge is simultaneously injected, is 1; for all other ABC130 measurements nStrobed is 16.

injected. The blue line with triangular markers and the light green line with square markers represent measurements of the ABC130 front end prototype chip using discrete load capacitors, injecting positive and negative charges respectively: these are unchanged from Figure 5.2 and are shown for reference only. The dark green line with triangular markers represent similar measurements of ABC130 using discrete load capacitors and negative charge injection: it can be seen that this line is now very similar to the performance of the prototype chip, which was not the case previously. The remaining markers represent data from Short Strip module SS1 (diamonds), Long Strip Module LS4 (plus signs) and an ABC130 Single Chip Board loaded with two mini-sensors (blue asterisks). Whilst the Long Strip module in particular still has excess noise in comparison to a simple capacitive load, at 14% this is much reduced from the previous 35% excess, and close to the 8.5% excess expected by the designer when taking into account the additional series noise due to the strip resistance of the sensor as outlined in Section 2.1.4.

5.4 Chapter Summary and Conclusions

For many months, ABC130 chips loaded with silicon strip sensors were found to have excess noise. Key contributions to this were identified from two effects, which have now been corrected: the use of incorrect DAC settings for the front end, and the unfortunate decision to strobe all the channels at once which, whilst having little or no effect upon an unbonded chip, has a significant effect upon the performance of the charge injection circuitry in short strip and long strip modules. In conclusion the performance of the ABC130 chip was now sufficiently well understood for us to proceed to build the modules for the first electrical stave. Assembly of the hybrids for these modules began in January 2017.

Chapter 6

Non-Segmented Stave with ABC130/HCC130 Chipset

This chapter documents the specification of the partially populated, non-segmented stave, and its performance.

6.1 Stave Specifications and Modifications

6.1.1 Bus Tapes

Whilst the stave concept was introduced in section 3.1.1, it is again summarised here for reasons of clarity.

Each stave comprises a support structure made of carbon fibre skins glued to a core of carbon fibre honeycomb and foam, into which a titanium cooling pipe has been integrated. Polyimide flex circuits known as “bus tapes” are co-cured onto the carbon fibre skins in order to route electrical services between the End of Substructure (EoS) card, which provides the interface to the off-detector systems, and the detector modules. Modules are glued to each face of the stave at different angles to give a small angle stereo geometry. The first electrical stave replicates one of the two innermost layers, where the track density will be highest, which are populated with short strip modules having four banks of 2.4 cm strip length. The design of the bus tape for this stave is referred to as being “non-segmented” as all 26 hybrid

locations are served by a single multidrop Clock, Control, Reject (CCR) bus. This approach had been adopted in order to minimise the amount of material in the tape, however this was subsequently determined to be sub-optimal with regard to signal propagation as will be discussed below.

Studies of point-to-point transmission on the non-segmented tapes were made at RAL using a high bandwidth Time Domain Reflectometer (TDR) with companion simulations being made at Oxford using Ansys High Frequency Structure Simulator (HFSS) [70]. The TDR measures the Scattering Parameters across a range of frequencies for each differential microstrip on the bus tape. Commonly abbreviated to “S-parameters”, the scattering parameter matrix records the relationship between reflected and incident signals at the beginning and end of the transmission line, from which the attenuation, reflection, transmission and eye diagram¹ for the expected signals may be computed. (For this study, such computations were made using the commercial software package Signal Integrity Studio, from Teledyne LeCroy [71], which is no longer available.)

This work showed that the signal quality of the 160 Mbit/s point-to-point data signals would be very poor if the tape was laminated directly to the carbon fibre surface of the stave as had been planned. There were two main reasons for this: firstly carbon fibre sheeting makes a poor reference plane for our transmission lines as it is lossy at high frequencies (in other words the AC resistance is high). Secondly the characteristic differential impedance of the point-to-point transmission lines on these tapes was determined to be of order 80 Ω , lower than the 100 Ω characteristic of the EoS board. These issues were resolved by adding one extra polyimide/copper layer under the bus tape, with the copper adjacent to the carbon fibre skin, such that the data traces were an increased distance from a reference plane of higher conductivity. The increased dielectric thickness results in the nominal differential impedance of the point-to-point lines being increased to a nominal 110 Ω , closer to the 100 Ω target. In addition, as the copper thickness of 17.5 microns exceeds the skin depth for

¹An eye diagram is a graphic often used to illustrate the margins present in digital communication systems. The received signal is repetitively sampled such that many (0 to 1) and (1 to 0) transitions are overlaid. The resulting image represents the opening of an eye.

frequencies of interest, the carbon fibre no longer transports any of the return current so its lossiness is not a concern. (For production staves a new build with revised trace geometry is proposed which will get close to the desired $100\ \Omega$ characteristic without the additional material arising from an extra layer.)

Next the effect of loading the multi-drop CCR bus with differing numbers of hybrids was approximated by adding a $3.7\ \text{pF}$ load at each hybrid location. For these tests representative pseudo-random $160\ \text{Mbit/s}$ data was driven onto one CCR pair, with eye diagrams being recorded on an oscilloscope with $1\ \text{GHz}$ bandwidth. It was found that locations near the driven end of the bus were affected by reflection; locations near the centre were affected by both reflection and attenuation and had the worst signal quality; locations at the far end had relatively clean signals affected only by attenuation [72]. For this first stave it was decided to build a partially loaded object, using only the four module locations (8 hybrids) with the best signal quality².

6.1.2 End of Substructure Cards

End of Substructure (EoS) boards for this first stave were made by DESY using commercial parts. The EoS board used for the first, Left Hand (LH) stave side is shown in Figure 6.1: a number of modifications may be seen which are further described in the following paragraphs.

The origin of this EoS design predates the decision to make DC-DC Point of Load power conversion the baseline powering scheme, hence the board was made to support the Serial Powering scheme. This has two consequences. Firstly the CCR signals sent to the modules are driven by a commercial LVDS driver which sets the common mode level at $1.2\ \text{V}$. For a serially powered stave, this does not matter as the signals must be AC coupled at the hybrids anyway: as one goes along the stave, each module will be at a higher potential than

²For future staves this issue has been resolved by a new, segmented design in which a series of CCR buses serve no more than 10 hybrids each. The layout, with two copper layers, is outlined in Chapter 7.1.1.




Figure 6.1: End of Substructure (EoS) board with modifications.

the last by 1.5 V. In the case where a DC-DC powered stave is to be built, in principle the AC coupling of the CCR signals at the hybrids may be omitted as long as the common mode of the driver falls within the limits of the receiver, but this is not the case for the standard LVDS driver. Furthermore, for the non-segmented bus tape, the impedance of the CCR bus loaded with eight hybrid drops remains of order $68\ \Omega$ even after the addition of the extra polyimide/copper layer. A standard LVDS driver delivers 3.5 mA current, resulting in a 238 mV swing across a $68\ \Omega$ termination resistor as opposed to the 350 mV expected across a $100\ \Omega$ termination resistor. To increase the drive current in order to achieve the required signal amplitude at the end of the CCR bus, and to set the common mode of the CCR signals within the common mode of HCC130's receivers in case AC coupling were to be dropped from the hybrids, a small PCB was made which hosts a custom SLVS buffer IC [69]. This was mounted on the stave as shown toward the bottom of Figure 6.1 and bonded between the EoS CCR outputs and the corresponding signal lines on the bus tape to ensure that the common mode levels are correct and to boost the drive current by a factor of two.

Secondly, in a serially powered stave the data sent from the hybrids to EoS must be AC coupled at the receiver: the common mode of the data signals across a serially powered stave with 14 modules would differ by 19.5 V. In principle an AC-coupled data connection should

also work in the DC-DC powered case, however in the EoS design the AC coupling was implemented using compact four capacitor arrays. Although not known at the time of EoS manufacture, these arrays have significant parasitic capacitance between elements, leading to crosstalk which can degrade signal quality. Such capacitor arrays are best avoided hence, together with their associated hysteresis feedback resistors (used to maintain DC bias), they were removed from the EoS boards used on the stave. As can be seen toward the bottom left hand corner of Figure 6.1, fine, enamelled copper wire was instead used to complete the DC circuit path into the EoS receiver chips.

A few smaller issues apply to this EoS design. Due to a design error, a number of pads intended to be wire bonded to the stave were left covered by solder resist, including LV power and its return. The offending resist was removed and the pads were soldered to the bus tape using copper braid. In addition the provision to connect to the stave's dual embedded Negative Temperature Coefficient (NTC) thermistors, used to determine the temperature of the cooling pipe for purposes of a safety interlock, was incorrect, having both the wrong form of connector (pins up vs pins down) and the wrong pin mapping (two wires instead of four). This problem was worked around by making a separate cable to connect the NTC polyimide flex directly to the interlock system, bypassing EoS completely, as can be seen at the top right of the picture. Finally the LV and sense connections were provided by means of a 0.1" header which is not well suited to onward connection using high cross section cables: a small PCB was made to route each node to a connector from the Molex Mini-Fit JR series, providing a more manageable means to cable up the system.

6.1.3 Power Supplies and DAQ

With the initial EoS design, each stave side requires two low voltage power supplies, one for the EoS card at 4.5 V and one for the on-module DC-DC converters at 11 V. To meet this need a pair of TTI model CPX400DP [73] low voltage power supplies are used, connected to the control PC by means of their built-in USB interfaces. These supplies receive mains power which routes through the Cambridge Interlock box [74]. Powered by Arduino [75], a popular

open-source microcontroller based electronics platform, this box monitors the temperature of the dual NTC thermistors mounted within the core on the stave cooling pipe and will turn off the mains power if either sensor exceeds a programmable limit. This limit is presently set to 22 °C.

Similarly each stave side has two high voltage buses. Previous studies of stave prototypes using a chipset in 250 nm technology had used SCTHV VME modules [76] originally made for SCT module production tests, with each of four channels being capable of delivering 500 V. This system which was brought back into use for these tests, which required the installation of updated, commercial driver software and the modification of the ITSDAQ SCTHV control software to access the hardware by means of a more modern API library. For connection to the stave a pair of D7W2 connectors are provided on the EoS pigtail, each with two HV rated pins. During the very first stave tests the HV cabling seemed to be a conduit for common mode noise to enter the system. This was remedied by the addition of a simple RC filter at the point of power entry: a 200 Ω resistor is fitted at the stave end of each HV cable with a 47 nF capacitor between each HV line and stave GND being attached to the rear of the EoS pigtail. To choose the value of such a resistor one must bear in mind that when the ITk strip detector reaches the end of its life, sensor bias currents will have increased significantly due to radiation damage³. Hence the voltage drop across the resistive element of any filter in the bias chain, at end of life sensor currents, must to be carefully considered. Each bias segment will serve no more than four modules, and each module's powerboard already contains a two stage RC filter with a total series resistance of 11 k Ω . Hence for the RC filter added upon entry to the stave a filter resistance 200 Ω was selected to ensure that any additional voltage drop would be one magnitude lower than that already expected, due to the presence of four 11 k Ω "resistors" in parallel.

In order to interface both sides of the stave to a single Nexys Video [77] FPGA board, a pair of adaptor PCBs were made to reduce the connectivity required for a single, partially loaded stave side from two 50-way ribbons down to one, to match the two 50 pin headers

³For the worst affected regions of the strip barrel, end of life sensor leakage currents are expected to be of order 1.3 mA.

Side	Position	HV	Name	Breakdown	Tested	Sensor
LH	0	0	RAL_SS4	400	Part	ATLAS12 W721
LH	1	0	RAL_SS7	>500	Part	ATLAS12 W659
LH	2	1	RAL_SS6	260	Part	ATLAS12 W660
LH	3	1	LIV_W682	280	Part	ATLAS12 W682
RH	0	2	RAL_SS5	>500	Part	ATLAS12 W606
RH	1	2	LIV_W683	>500	Part	ATLAS12 W683
RH	2	3	LIV_W680	155	Part	ATLAS12 W680
RH	3	3	LIV_W668	165	Part	ATLAS12 W668

Table 6.1: A list of modules which became part of the Non-Segmented ABC130 Stave. Modules were grouped according to the sensor breakdown voltage observed during module reception testing, such that sensors with similar breakdown potentials were placed on a common HV line. All HCC130 IDs are unique within the stave.

available from the DAQ side. It should also be noted that communications between the FPGA board and the host PC take place over a gigabit ethernet link, which has insufficient bandwidth available to transfer the ABC130 data for so many modules to the host in realtime: the packet buffers within the firmware are quickly exceeded leading to data loss. Whilst the development of histogramming firmware for the ABC130/HCC130 protocol was started in response to this limitation, for the majority of work with this stave, a slower methodology was used in which, for a given event, every ABC130 chip in the system must report that its data has been fully transmitted before the next trigger is sent.

6.2 Modules and Modifications

Table 6.1 shows the eight modules which were made for use with this stave. Four were built at RAL and are identified by the order of assembly (SS5, SS6, SS7 and SS8), and four were built at Liverpool which are identified by the sensor serial numbers (W668, W680, W682 and W683). A photograph of the completed “LH” side is shown in Figure 6.2.

For reasons which remain under investigation with the vendor, this batch of sensors seems particularly susceptible to localised breakdown leading to increased leakage current, especially during standalone module tests. The fixture used for single module reception tests, before they are mounted on the stave, must provide both electrical (sensor backplane






Figure 6.2: The Non-Segmented ABC130 Stave, “LH” side. Note that the four modules were mounted at the four locations furthest away from the End of Substructure (EoS) card.

contact) and thermal (cooling) connectivity to the module. In order to isolate the coolant from the HV potential, the present design comprises two aluminium plates separated by a polyimide isolation sheet. The lower plate has both cooling channels through which water may flow and vacuum channels, routed through to the top of the assembly by a matrix of drilled holes. Vacuum must be applied such that the sensor makes a good electrical and thermal connection with the top plate, however this may also deform the sensor. There is evidence to suggest that this may lead to localised increases in sensor leakage current due to increases in mechanical stress [78].

As each stave side has only two HV buses to be shared between four sensors, an attempt was made to group modules according to the maximum voltage which their sensors were able to withstand during single module tests. The first stave side (“LH” geometry) was



(a) Schematic



(b) Oscillator Waveform

Figure 6.3: Schematic of the prototype HV-MUX circuit (a) and its oscillator waveform (b). For figure (b) the horizontal scale is 5 ns per box and the vertical scale is 1 V per division.

fitted with the modules which were able to withstand higher potentials (SS4, SS7, SS6 and W682) leaving the other modules (SS5, W668, W680 and W683) to be fitted with HV-MUX circuitry for use on the second stave side (“RH” geometry). (It would later be seen that most modules behaved rather better, in terms of breakdown voltage, on the stave than during the single module reception tests: this will be discussed in section 6.3.2). The first side (“LH” geometry) was made available for testing in March 2017, followed by completion of the second side (“RH” geometry) in May 2017.


The first module to have the HV-MUX circuitry fitted was RAL SS5, chosen due to its low leakage current. The schematic of this circuit, which was fitted before the module was mounted upon the stave, is shown in Figure 6.3a. The NC7S14 Schmitt Trigger IC (U1) together with resistor R1 and capacitor C1 forms a square wave oscillator which drives the left hand side of capacitor C3 with the waveform shown in Figure 6.3b. This drives

the voltage multiplier formed by diode array U2 and capacitors C4, C5 and C6 in order to generate the gate voltage from GaNFET U3, which must be held at least 2 V above the negative HV supply in order to connect the sensor to the high voltage supply.


Standalone testing of module SS5 before and after the addition of the HV-MUX circuit comprised Three Point Gain (3.3.3) and Response Curve (3.3.5) tests which revealed there to be additional, localised noise correlated with the position of capacitor C3. Considering the oscillator waveform shown in Figure 6.3b this is not surprising: the fast edges of the oscillator waveform are within the bandpass of the ABC130 amplifier and may be expected to inject signals. For this module a piece of copper foil was fitted in the shadow of C3 to shield the silicon sensor from its field emissions, and at the level of Three Point Gain and Response Curve tests this was found to remedy the problem. Three further modules then had the HV-MUX circuit added: Liverpool modules W680, W683 and W668. Together with RAL module SS5, these modules were added to the second (“RH”) side of the stave. In each case the copper shield was fitted before module attachment, so localised noise increases correlated with C3 were not expected.

Once the four HV-MUX modules had been added to the stave and digital communication issues had been debugged, Strobe Delay, Three Point Gain and Response Curve tests were made followed by the Noise Occupancy test. It was unfortunate that the noise occupancy test had not been performed earlier as this revealed that for three modules, SS5, W668 and W680, the copper foil treatment was not sufficient to shield the sensor from the HV-MUX oscillator signal. An example of this is shown in Figure 6.6a which clearly shows excess noise centred around channel 280. (Note that module W683 does not exhibit this problem at all: the reasons for this remain to be understood as the oscillator waveform looks much the same).

To remedy this problem a different approach was taken, namely to change the oscillator waveform from a square wave to a sawtooth. The modified schematic is shown in Figure 6.4a, with resistor Rx and capacitor Cx being added to form an RC filter. The signal which drives this filter is still a square wave, so it is important to minimise the copper area upon which



(a) Schematic



(b) Oscillator Waveform

Figure 6.4: Schematic of the modified HV-MUX circuit (a) and its oscillator waveform (b). The modification adds the RC filter formed by components Rx and Cx. For figure (b) the horizontal scale is $2 \mu\text{s}$ per box and the vertical scale is 1 V per division.




Figure 6.5: Image of a modified HV-MUX circuit.




Figure 6.6: Noise occupancy of SS5, Hybrid 1, Column 0.

this “aggressor” signal remains. Bearing this in mind, the modification was implemented as shown in Figure 6.5. Firstly resistor R1 is removed such that the trace carrying the original “aggressor” signal can be cut. A dual resistor array wrapped in shrink tubing is then mounted above U1 to provide resistors Rx and R1 and capacitor Cx is mounted vertically on the ground pad intended for capacitor C8, which is not fitted.

As a result of these changes capacitor C3 is now driven with a “clean” signal as shown in Figure 6.4b: the fast edges are no longer present, hence there is no significant frequency content within the bandpass of the ABC130 front end and there should be no additional noise burden in our modules. Indeed this is demonstrated by Figure 6.6, which shows the occupancy due to noise (pedestal scan) as a function of channel number as first obtained on-stave, with copper shielding tape adjacent to HV-MUX capacitor C3, and in the modified condition in which the shield has been removed, but the drive waveform is now a sawtooth. The first plot clearly shows a noise excess centred upon channel 280, which has been eliminated by the modification. As a consequence of this study, the RC filter formed by the added components Rx and Cx was built into the next powerboard design. For production modules, the AMAC2 chip will additionally include a series resistor akin to Rx on each HV-MUX

oscillator output pin.

6.3 Results

6.3.1 Comparison of noise on and off the stave

A comparison of noise estimates made by the Three Point Gain method as described in Section 3.3.3 for the eight modules of this stave, before and after their attachment to the stave core, is shown in Figure 6.7. It can be seen that there are good correlations between the noise values recorded both on and off the stave, with the on-stave results typically being higher by only 15 electrons. At this point the stave was operating with conservative bias settings consistent with those used for standalone tests: no attempt was made to push toward higher bias voltages until the basic operation of the stave had been established.

6.3.2 Noise as a function of Sensor Bias

It is well known that the capacitance of a diode varies as a function of the applied bias voltage, and silicon strip sensors are no exception to the rule. Furthermore as the series noise of the front end amplifier varies according to the load capacitance it can be seen that the noise in our detector system will vary as a function of the applied sensor bias voltage until full depletion is reached.

The load capacitance C_l which a silicon strip sensor presents to the front end amplifier is typically calculated as:

$$C_l = \frac{C_b}{n} + C_i + C_{i2} \quad (6.1)$$

Where C_b is the backplane capacitance, n is the total number of strips, C_i is the inter-strip capacitance to nearest neighbours and C_{i2} is the interstrip capacitance to next nearest neighbours.




Figure 6.7: Input Noise by Three Point Gain method: comparison of results on and off the stave. In both plots, each point represents the arithmetic mean of the input noise for 128 channels, either “outer” strips which do not run under the hybrid or “inner” strips which do. These must be treated separately as the noise level is different. Hence each module is represented by 40 points. In the case of the upper plot, within a module these points are ordered as follows to represent the order of the columns of strips along the stave: 10 points for outer strips, 10 points for inner strips, 10 points for inner strips, 10 points for outer strips. Aside from the last module, the noise measured on the stave is marginally higher than that measured off-stave.

The backplane and interstrip capacitance of the sensors used to build the modules fitted to the non-segmented stave were measured at The Cavendish Laboratory, The University of Cambridge, as part of sensor reception testing. Plots of $\frac{C_b}{n}$ as a function of the applied Bias Voltage, for the eight sensors used in this stave, are given in Figure 6.8a. It can be seen that full depletion, the point beyond which the measured capacitance is constant, occurs at between 360 and 390 V, which supports the decision to test the modules at 400 V. The total number of strips on a short strip sensor does not in fact match the 5120 instrumented channels of the short strip module as an additional “field shaping” strip is placed at each end of the four columns, therefore the total number of strips given that there are four columns is 5128.

Whilst the backplane capacitance characteristic may vary from sensor to sensor, the interstrip capacitance is extremely well controlled by the vendor, meaning that this part of





the load may be considered identical for all sensors. This measurement, C_i , is shown in Figure 6.8b, which illustrates that a certain voltage must be applied to the sensor in order to achieve separation of the strips. The voltage at which separation occurs can however vary according to the historical treatment of a sensor with respect to bias voltage and humidity [79]. Note also that the figure shown includes only the nearest neighbours. It is usual to estimate the additional capacitance to the next nearest neighbours as 10 to 15 % of the interstrip capacitance to the nearest neighbours, and to include this in the load calculations. Figure 6.8c therefore estimates the total load capacitance of a single front-end amplifier, C_l , as $\frac{C_b}{n} + 1.1C_i$, scaling the interstrip capacitance by a factor of 1.10 to accommodate the additional capacitance to the next nearest neighbours.

The “Noise Slope” of the ABC130 front end amplifier as a function of capacitive load, measured using discrete capacitors, was presented as the dark green line in Figure 5.4. This fit is described by the following equation:

$$N = \sqrt{219981 + (16999.4C_l^2)} \quad (6.2)$$

Where N is the input noise in electrons and C_l is the load capacitance.

This fit may be used together with the data shown in Figure 6.8c to estimate the input noise expected from the ABC130 short strip modules built with these sensors. These predictions are shown together with measurements of input noise as a function of sensor bias, made by the Response Curve method, in Figure 6.8d. The predictions are the lines and the measurements are the markers: each marker represents the average input noise of all known good channels of the outer columns of that module, such that we do not need to consider the additional capacitance which the overlaid hybrids and powerboard will present to the inner columns. Although the noise measured on the stave is in general lower than the prediction, with the exception of RAL_SS4, at the usual 400 V bias all modules are within 30 electrons of the prediction. In the case of RAL_SS4 the measurement is around 80 electrons lower than the prediction. Whilst the reasons for this have not been understood, the higher serial

(a) Bulk Capacitance, C_b (b) Interstrip Capacitance, C_i (c) Load Capacitance, C_l 

(d) Estimated and Measured Input Noise

Figure 6.8: Sensor Capacitance and Input Noise Predictions. The figure shows (a) Bulk Capacitance, (b) Interstrip Capacitance and (c) Load Capacitance as a function of sensor Bias Voltage. The measured variation in the ABC130 front end amplifier noise as a function of load capacitance may be used together with the measured load capacitance to predict the variation in Input Noise as a function of Bias Voltage for short strip modules built with these sensors. Figure (d) shows the predicted Input Noise values for each of these eight sensors (lines) together with measurements (markers) for the corresponding modules, made on the stave using the Response Curve Method. Each set of markers corresponds to the line in the same colour. In the case of figure (b) data is presented for sensor W604: the sharp fall in capacitance at around 110 V occurs as the depletion zone extends beyond the point at which the strip implants become isolated from one another. All plots are shown using a logarithmic scale for the Y-axis. (Data from the Cavendish Laboratory, University of Cambridge [80].)

number of this sensor might imply it was manufactured later than the others.

Given that during standalone module tests it was frequently impossible to reach the desired 400 V bias target, the ability to run this test for so many sensors was an encouraging result. Indeed only one module now breaks down, namely module 2 on the “LH” stave side (RAL_SS6) upon which HV-MUX is not fitted, so the sensor may not be isolated. Whilst inspection of the occupancy data suggests that breakdown first starts at 350 V, at which point the HCC130’s (uncalibrated) measurement of the sensor bias current is 5 counts, by 400 V this measurement has increased to 56 counts and the 15 channels are affected as shown in Figure 6.9. As a result modules 2 (RAL_SS6) and 3 (LIV_W682) which share a common sensor bias line, could not be tested above 400 V, which can be seen as the omission of point markers in Figure 6.8d.






Figure 6.9: Channels of module 2 affected by sensor breakdown

6.3.3 Gain and Noise as a function of nStrobed


The ABC130 has a single calibration bus, with each front end amplifier’s calibration capacitance connected to this bus through a transistor. Registers are provided to operate these transistors in any desired pattern. Hence a pulse on the calibration bus can be sent simul-




(a) Gain of Outer Columns




(b) Gain of Inner Columns




(c) Output Noise of Outer Columns



(d) Output Noise of Inner Column



(e) Input Noise of Outer Columns



(f) Input Noise of Inner Column

Figure 6.10: Gain and Noise as a function of nStrobed. Figures a) and b) shown the measured Gain, Figures c) and d) show the measured Output Noise and Figures e) and f) show the Input Noise calculated by dividing the Output Noise by the Gain, for Outer and Inner columns respectively. In each case, one line represents data for a single hybrid. The points and error bars represent the mean and RMS of the distribution of the selected parameter across the ten chips of each hybrid: the spread is dominated by gain differences between the chips. It can be seen that higher Gain values are recorded as nStrobed is reduced, which has a direct effect upon the calculated Input Noise. The Output noise is nearly static throughout. All measurements are Response Curves recorded with nTriggers = 512.

taneously to any number of channels in the chip, from 0 to 256. Whilst 256 is generally favoured for reasons of speed as charge injection measurements would otherwise need to be repeated several times until all channels have been tested, the ABC130 is known to exhibit some deficit in the quantity of injected charge as the number of strobed channels is increased, leading to errors in the calculation of the gain (and hence the input noise). This effect was introduced in 5.3 and is shown again for the modules on the stave in Figure 6.10. There is a difference of 40 electrons between the Input Noise values recorded when injecting charge into 256/256 channels (1 in 1) at once and those recorded when injecting charge into 32/256 channels (1 in 8) at once. This is due solely to differences in the measured gain upon which the calculated Input Noise depends: the noise at the discriminator output (Output Noise in mV), remains flat throughout. Hence this is confirmed as an effect of the charge injection circuit, not the amplifier.

Note that despite this effect the default for the measurements of the modules on this stave remains to strobe all channels at once for reasons of speed. Hence unless otherwise stated, input noise values are expected to be over estimates by approximately 40 electrons.

6.3.4 Noise as a function of Time




Figure 6.11: Input Noise by Response Curve method: Noise as a function of Time

The stave was operated under constant conditions (400 V bias) for 60 hours, repeatedly executing the Response Curve test. The results are shown in figure 6.11. Noise is generally very stable as a function of time, the notable exception being hybrid 5 which exhibits a small glitch in the very first test. This is the module which has localised breakdown effects illustrated in Figure 6.9, so the glitch is consistent with this anomaly having taken some tens of minutes to settle into a stable state.

6.3.5 Estimation of Input Noise by Noise Occupancy method




Figure 6.12: A Comparison of Input Noise calculated by the Response Curve and Noise Occupancy Methods, presented together with the ratio between the two measurements.)

Figure 6.12 compares the Input Noise of the stave measured by the Response Curve method against an estimation of the Input Noise made by the Noise Occupancy method as described in Section 3.3.6. It can be seen that Input Noise measured by the Response Curve method with Optimised Burst Conditions ($n\text{Strobed} = 32$, $n\text{Trigs} = 256$) is on average 30 electrons (5

6.3.6 Use of module SS2 as an Aggressor

Previous work [81] has shown that electric field emissions from DC-DC converters may capacitively couple into both the front end wire bonds and the sensor strips, hence it is very important that the DC-DC part of the module powerboard should be carefully shielded. When assembled into barrels, there will be overlap between the staves, hence it is possible that electric field emissions from the on-module DC-DC converters or HV-MUX circuitry may be picked up in opposing modules. It was decided to make a first study of this possibility by placing an “Aggressor” module beneath the stave.

Module SS2 was originally built without HCC130 chips as a test piece in 2016. Following completion of the modules for the stave, SS2 was completed by the addition of HCC130 die, a powerboard and by completion of the wire bonding. It had originally been planned to add the HV-MUX circuitry, but this would not have been straightforward as the module uses the previous generation of hybrids using FR4 substrates. The low voltage connections between the hybrids and powerboards would be covered by the HV-MUX PCB, hence a spacer would need to be used to ensure that the bonds would not be squashed. This operation was considered relatively high risk, hence the HV-MUX circuit was omitted.

The module was placed onto the standard module test fixture and vacuum used to hold it down. This assembly was then placed under the stave, with SS2 facing module SS4, their respective sensors being approximately 25 mm apart. (In the installed barrels the minimum distance between the sensors of adjacent staves is much smaller than this, of order 10 mm, however the area of overlap is much smaller than was the case during this test.) The aggressor module was connected to the same Nexys Video board as the stave to ensure that all modules receive synchronised triggers; its sensor bias was connected to a Keithley 2410 Source Measure Unit (SMU); DC-DC input power was taken from a spare TTi power supply channel and the whole assembly was placed under a series of opaque, static dissipative black plastic bags. Side views of the module, underneath the stave, may be seen in Figure 6.13.

A comparison of Input Noise estimates made by the Response Curve method for the eight




Figure 6.13: Insertion of aggressor module SS2 beneath the stave




Figure 6.14: Input Noise by Response Curve method: comparison of results with aggressor SS2 on and off. The noise is not affected by operation of the aggressor.

modules of this stave, with the aggressor module SS2 running and not running, is shown in Figure 6.14. No significant differences are seen.

6.3.7 Powering the stave with a prototype "PP2" DC-DC converter

A comparison of input noise estimates made by the Response Curve method for the eight modules of this stave, with the "RH" side powered by a TTI CPX400DP [73] bench supply or by a Linear Technology DC1073A [82] DC-DC converter evaluation board is shown in Figure




Figure 6.15: Input Noise by Three Point Gain method: comparison of results with bench LV supply and DC-DC board DC1073A, which uses the Buck converter IC proposed for the DC-DC converter at patch panel “PP2”. The noise is unchanged. (One module is missing from the reference run due to the presence of digital errors.)

6.15. The DC1073A uses the LT3845 chipset [83] which, at the time, was under evaluation for possible use in the ITk strip detector to provide a 48 V to 12 V DC-DC converter to be located at Patch Panel 2 (PP2)⁴. PP2 is a set of active patch panels to be located between the cryostat which surrounds the ITk detector and the ATLAS Toroid which will map the cables arriving from power supplies located in the service caverns to the cabling which routes down to the petals and staves. As such this is an interesting test to verify that a typical circuit which uses this chip does not cause excess noise when used to power a stave. Admittedly at this stage the cable setup is not representative of the impedances one may expect in the final installation, however it is encouraging that no significant difference is observed.

⁴The LT3845 is no longer recommended for new designs. The PP2 DC-DC converter design will now use its replacement, the LT3845A [84].




Figure 6.16: Current and voltage delivered to the stave turning the power supply on turning the power supply off. The horizontal scale is $500 \mu\text{s}$ per box, the vertical scale is 500 mA per box for current and 2 V per box for voltage. (For the purposes of this test, the on-module DC-DC converters were set to switch on immediately power is applied, which would not normally be the case.)

6.3.8 Power disturbances at startup / shutdown

The Buck converter ASIC used in the stave’s DC-DC converters is known to fail if more than 14 V is applied to its input. Keeping the input voltage below this limit is therefore a key objective of the power distribution scheme to be used in the installed detector. Although the cables used for the present stave test stand are much shorter than the final cable chain, and the bench supplies typically used to provide power would be expected to perform well, we can already look at the transients observed when turning the stave on or off to see if one might expect any problems.

A Keysight N2893A current probe [85] (uncalibrated) was connected around the “RH” stave side positive LV feed and a standard single-ended voltage probe was connected between the positive LV feed and return at the point of entry to the stave. The traces shown in Figure 6.16 show what happened when the LV power, supplied by a TTI CPX400DP power supply, was turned on and off. For the purposes of this test the one wire controllers which enable the DC-DC converters, operating from a separate power supply, had already been programmed such that when the stave low voltage bus was powered, the DC-DC converters would power up in the “on” state.

Turning the power on (Figure 6.16a) there is an inrush current (yellow) of more than 3 A, which falls slowly until the delivered voltage (green) crosses 5.5 V, at which point the Buck converter ICs commence operation. The current now falls sharply, then begins to rise until all converters are fully operational. Beyond this point current falls smoothly as voltage rises, the efficiency of the power conversion increasing, until the steady state is reached with 11 V delivered to the stave. Turning the power off (Figure 6.16b) current rises as the voltage falls, until one by one each DC-DC falls below the minimum operational voltage, resulting in four downward steps in the current drawn. There is also a small glitch in the current at the end, as one converter briefly resumes operation.

The behaviour shown by these plots is acceptable: the feature we are keen to avoid, namely sharp glitches in the voltage rail which may push towards the 14 V limit, is seen not to occur. Naturally we must repeat this test with a more representative cable and power supply, once such units become available.

6.4 Chapter Summary and Conclusions

A partially loaded, double sided stave was built using a non-segmented bus tape. Although limited to four modules each side and built without the GBTx optical transceiver ASIC, this stave was a useful test vehicle. The HV-MUX circuitry, initially found to be a source of noise, was successfully modified on-stave, the added filter components being taken forward in future powerboard designs. Thereafter module performance was in agreement with expectation, input noise being of order 15 electrons higher for all modules running together on the stave than for each module running individually on the test bench. A study of input noise as a function of sensor bias confirmed the expected form, and stave performance was found to be stable as a function of time. Power disturbances observed powering up and powering down the stave were found to be within acceptable limits.

Chapter 7

Segmented Staves with ABC130/HCC130 Chipset

This chapter documents the specification of prototype segmented staves, both single-sided and double-sided, and their performance.

7.1 Stave Specifications

7.1.1 Bus Tapes




Figure 7.1: The copper layers of the Master Side Segmented bus tape. The top layer carries the Clock, Control, Reject (CCR) traces and LV power whilst the bottom layer carries signal ground and LV power. The CCR traces are split into four segments serving, from left to right, 1, 5, 4 and 3 modules respectively.

Section 6.1.1 makes reference to the effects which attenuation and reflection have on

signal quality for the multidrop Clock, Control, Reject (CCR) signals on the stave bus tape and concludes that in order to maintain good margins one should limit the number of hybrid loads on each bus to ten. This is achieved by the layout shown in Figure 7.1, which is for the outer (Master) side of the stave. Although not shown, the layout for the inner (Slave) side is similar but mirrored. All of the fast signals to and from HCC130 route in the top copper layer together with 11 V power, control lines for AMAC and two HV lines. The bottom layer carries the reference ground plane for each of the fast signals, 11 V return (0 V) and the other two HV lines.

Although of similar construction to standard kapton flex circuit technology, the length of the bus tape is beyond the capability of many vendors. The build is centred upon a polyimide sheet upon which copper has been deposited on both sides. This is patterned and etched to form both copper layers. As the fast signal traces and their reference plane are formed on opposing sides of the same sheet the separation between them is extremely well controlled which gives better impedance control than any build in which a compressible glue layer may come between the active copper layers. Additional polyimide layers are cured onto each face of the double sided core to provide insulation from the sensor backplanes and the carbon fibre stave skin. One novel feature of this build is that traditional through vias are not used to provide connectivity between the layers, instead a laser drilling process is used to open apertures through the polyimide base layer such that copper on the lower layer is made accessible from the top. The same process is used to open apertures in the two cover layers, to selectively expose the fast signal ground trace in order for it to connect with the carbon fibre stave skin to ensure that the stave structure is grounded and to allow connections to be made to the silicon strip modules, bus terminations and End of Substructure (EoS) cards on the top. Once a bus tape has been completed, all exposed pads are given an Electroless Nickel Immersion Gold (ENIG) plating to make them suitable for aluminium ultrasonic wedge (wire) bonding.

With reference to the CCR traces, it can be seen from Figure 7.1 that the design provides four segments serving a total of 13 modules: from left to right, 1, 5, 4 and 3 modules

respectively¹. One feature worthy of note is that the CCR buses for the 5 module segment start at the EoS at the top left and route almost halfway along the stave, under the sensors of the modules they serve, before routing to the stave edge and then doubling back. In this manner a trace length of order 50 cm is inserted between the line driver at EoS and the hybrid loads which helps to attenuate the high frequency content of the signals, thereby reducing the severity of the reflections which occur at each hybrid drop. (For the four and three module segments similarly long unloaded trace lengths fall naturally out of the layout.)

7.1.2 End of Substructure Cards




Figure 7.2: The Master End Of Substructure (EoS) card as fitted to the UK Single Sided Stave. Using a mixture of commercial and custom ICs, this board provides power and control signals to the stave. Fast signals arrive and depart the stave in the optical domain, being converted from/to electrical form by the SFP+ transceiver, and are demultiplexed/multiplexed by the GBTx chip.

End of Substructure (EoS) boards for these staves were again provided by DESY, this time using a mixture of custom and commercial parts. The EoS Master board fitted to the outer side of the UK Single Sided Stave, the first to be assembled, is shown in Figure 7.2. The key component is the GigaBit Transceiver (GBTx) [86] chip, which is partially

¹Production staves will host 14 modules on each side, henceforth the rightmost CCR segment will be extended to serve one additional module.

hidden under the red plastic cover towards the left of the image. The precursor of the Low Power GigaBit Transceiver (lpGBT) [47] chip which will be used in production staves, GBTx is a radiation tolerant ASIC designed to provide an interface from/to high bandwidth serial streams, transmitted in the optical domain, to/from the detector front end hardware. In this application it provides four sets of CCR signals each comprising one 40 MHz Bunch Crossing clock and two control streams at 80 Mbit/s, and it receives up to 26 data streams at 160 Mbit/s, one from each hybrid fitted to the stave side. Conversion between the electrical and optical domains is provided by a commercial transceiver mounted within the SFP+ cage at the bottom right of the image. At the centre of the image there is the GBT Slow Control Adaptor ASIC (GBT-SCA): it had been hoped to use this to provide a bridge between the differential serial streams provided by GBTx and the low voltage I2C streams required by AMAC1A, but the specific implementation of the I2C protocol within GBT-SCA is not fully compatible with AMAC1A so this feature was not used. Instead an external control path was provided for AMAC1A. Each EoS contains a further three custom ASICs: located under each of the black plastic covers located near the upper edge of the board is a custom SLVS buffer IC [69] used to shift the 200 mV common mode level of the GBTx's SLVS outputs to the level of 600 mV, better suited to the HCC130².

In terms of electrical connectivity, the wires which carry the (high current) 11 V power are shown and labelled in Figure 7.2. Underneath this a brown flex circuit may be seen which routes to a second PCB fitted with commercial connectors. This flex carries the High Voltage lines needed by a complete double sided stave; sense wires for the 11 V low voltage power; the I2C lines needed to configure AMAC1A; and four NTC lines used to monitor the temperature of the stave cooling pipes. Whilst in principle the GBTx contains a series of eFuse registers which may be programmed to bring up the chip in the correct operating mode, for all work reported in this thesis the GBTx was configured by means of an external I2C bus. The GBTx, GBT-SCA and buffer ICs require 1.5 V power: whilst a space is left to fit an optional DC-DC converter to generate this from the 11 V supply, for the tests reported

²Whilst a common mode level of 600 mV would have been necessary for DC coupled connections to HCC130, in fact the clock and control inputs to the chip are AC coupled on each barrel hybrid, meaning that this additional buffer stage could have been omitted.

in this chapter this voltage was supplied directly. Finally, along the top edge of the EoS PCB, wirebonded power and signal connections to the stave bus tape may be seen.

For the double sided stave, the inner side of the stave is fitted with an EoS Slave board. This is largely the same as the EoS Master although without the flex circuit pigtail or the space for the optional DC-DC converter. Instead the Slave EoS plugs into the back of the Master EoS by which means it receives its four HV lines and 1.5 V power. Wires soldered between the Slave EoS and the Master's connector PCB are used to deliver 11 V power. The Slave EoS also requires its own I2C connection in order to configure the GBTx chip.

In order to minimise reflection, each differential pair must be terminated by a resistor which matches the differential impedance of the transmission line. Whilst the GBTx chip provides an integrated termination for each data input, it is necessary to provide a line termination at the end of each of the multidrop CCR buses. The schematic of the termination is shown in Figure 7.3a. At the end of each differential pair, represented as nodes BUS_P and BUS_N, is a pair of $45.3\ \Omega$ resistors connected in series. The chosen resistors are the stock values which come in just under the measured differential impedance of the bus tapes used, confirmed by TDR measurements to be $92\ \Omega$. At the centre of the resistor pair there is a 10 nF capacitor in order to keep the common mode aligned to the signal reference ground, although this part was omitted from the boards used with the US stave. The actual board represented by Figure 7.3b implements this circuit three times, once for each signal of the CCR triplet. As each stave side has four CCR segments, four termination boards are required.





Figure 7.3: In order to minimise reflection, each differential pair must be terminated by a resistor which matches the differential impedance of the transmission line: (a) shows the circuit used to terminate the multidrop CCR buses on the stave and (b) shows its physical implementation.

7.2 UK Single Sided Stave

7.2.1 Build and Test Information

The first fully populated, single sided stave with segmented bus tapes and GBTx was assembled at RAL using UK Stave core number 15. Assembly was completed in July 2018 using the modules listed in Table 7.1. Only two of these modules, fitted at the first and last locations, have real ATLAS12 Short Strip (SS) Sensors. The remainder were built using Mechanical Sensors (MS), a facsimile of a real sensor comprising the correct top metal layer printed on silicon without any of the other processing steps. For the modules built with mechanical sensors the ABC130's front end channels were not bonded as the load capacitance would be unrealistically high, hence these modules should return noise values consistent with unbonded chips. The HCC130 IDs (addresses) assigned to the modules were carefully chosen to avoid repetition, such that the four CCR segments could be driven with duplicate signals whilst retaining the ability to configure each hybrid independently.

In terms of low voltage power, 1.5 V was supplied to the EoS board by a TTi MX180TP[87] power supply, with 11 V being supplied to the stave from a TTi CPX400DP[73] model which delivers higher current. Mains power to the CPX400DP unit, which delivers the majority of power to the stave, was routed through the Cambridge Interlock box[74]. This monitors Neg-

Position	HV	Name	AMAC ID	HCC0 ID	HCC1 ID	Tested	Sensor
0	0	RAL_SS10	0	2	5	Yes	ATLAS12
1	0	IHEP_MS01	1	6	7	No	Mechanical
2	0	BHAM_MS01	2	21	24	Part	Mechanical
3	1	BHAM_MS02	3	22	23	Part	Mechanical
4	1	SHEF_MS01	4	13	16	Part	Mechanical
5	1	LIV_MS01	5	9	12	Yes	Mechanical
6	1	LIV_MS02	6	11	10	Yes	Mechanical
7	2	CAM_MS01	7	17	18	Yes	Mechanical
8	2	CAM_MS02	8	19	20	Yes	Mechanical
9	2	GLA_MS01	9	29	26	Yes	Mechanical
10	3	SHEF_MS03	10	14	15	Yes	Mechanical
11	3	SHEF_MS02	11	4	8	Yes	Mechanical
12	3	BHAM_SS01	12	25	27	Yes	ATLAS12

Table 7.1: A list of modules which became part of the UK Single Sided Stave. The table is separated into sections to match the CCR segmentation of the bus tape. All HCC130 IDs are unique within the stave. Only two of the modules have real ATLAS12 sensors, the remainder were built using mechanical samples.

ative Temperature Coefficient (NTC) thermistors attached to the stave’s embedded cooling pipes and will turn off the mains supply if a temperature of 22 °C is exceeded. The two modules with real sensors received high voltage bias from a VME crate filled with SCTHV[76] VME modules, rated to deliver up to 500 V at 5 mA. The coolant was water at 17 °C.

As mentioned previously, the EoS cards used for these staves provide optical control and readout by means of the GBTx chip. The ITSDAQ firmware was ported to the Gigabit Link Interface Board (GLIB) [88], provided by CERN as a supported hardware platform for GBTx testing. Using the SFP+ cages built onto the GLIB main board, the system was able to support up to two GBTx chips. Unfortunately as the network link between GLIB and the host computer runs at 1 Gbit/s, test methods which involve the full transfer of event data between the GLIB and the host do run slowly. Firmware histogramming was implemented as a means to speed up the testing, however unresolved flow control issues meant that runs using the firmware histogramming were limited to only a few tens of events: this was not widely used.

In addition to the HCC130 and ABC130 front end chips, which are configured by the DAQ system through the GBTx chip, testing a stave involves two other control and monitoring

paths. Firstly the GBTx chip must itself be configured. This is done using a custom 1.2 V USB-I2C adaptor[89] and accompanying Java software[90]. Secondly the AMAC1A ASICs on each module's powerboard must be initialised over a dedicated I2C bus running at 1.5 V, such that each module's low voltage and high voltage power may be enabled. For this first stave this was handled by means of an existing low voltage I2C interface, PMOD-AMAC, which had been made in support of the module test programme. For stave test purposes, this was connected to the GLIB by means of a custom patch cable.

7.2.2 Results

A partially completed stave, fitted with modules 0 to 4, was first tested in June 2018. It took a few days to implement and debug the mappings between the various modules and streams (e-links) within the lpGBT, and to tune the various delays and sample phases built into the system to the point at which a Three Point Gain test could be completed. Everything seemed to be in order: module mounting then resumed, with the completed stave being made available for testing in July.




Figure 7.4: An IV curve, up to 400 V, for the two SS modules mounted on the UK Single Sided Stave. The current drawn reflects the presence of a 10 M Ω drain resistor on the module powerboard, in parallel with the sensor. There is no evidence of sensor breakdown.

During testing, the SS modules with real sensors were biased at 400 V. The corresponding

IV curves, with currents recorded at the power supply, are shown in Figure 7.4. The IV characteristic is dominated by that of a $10\text{ M}\Omega$ drain resistor fitted to the powerboard in parallel with each sensor: the actual current drawn by the sensors is much smaller and hence cannot be seen. In principle the sensor bias return current may also be monitored by the AMAC1A chip, however in practice the current monitoring resolution of this version of AMAC is too low for the current drawn by a well behaved, unirradiated sensor to be seen. However we should be able to see currents of order $1\text{ }\mu\text{A}$ and above, such as may occur if a sensor were to breakdown.




Figure 7.5: A comparison of Gain for the modules of the UK Single Sided Stave measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are full Response Curves taken at RAL under Standard Burst Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 192$). In both plots, each point represents the arithmetic mean of the input noise for 128 channels, either “outer” strips which do not run under the hybrid or “inner” strips which do. These must be treated separately as the noise level is different. Hence each module is represented by 40 points, ordered as follows to represent the order of the columns of strips along the stave: 10 points for outer strips, 10 points for inner strips, 10 points for inner strips, 10 points for outer strips. The gain does not change when the modules are mounted on the stave. (Note that off-stave data is not available for Module 1.)

Several days were spent adjusting the configuration of each hybrid, flipping the direction of data loops, adjusting the phase at which the data is sampled upon entry to HCC130 and adjusting the sampling phase of data upon entry to the GBTx chip in order to minimise




Figure 7.6: A comparison of Input Noise for the modules of the UK Single Sided Stave measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are full Response Curves taken at RAL under Standard Burst Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 192$). The Input Noise of the MS modules with mechanical sensors is not changed, however the Input Noise of the SS modules with real sensors, fitted at the first and last locations, is increased. (Note that off-stave data is not available for Module 1.)

error rates. Once the stave was working sufficiently well, a Strobe Delay Test as described in Section 3.3.2 was performed in order to optimise the timing of the charge injection pulse, followed by a Response Curve Test as described in Section 3.3.5. The results of the Response Curve test are shown in Figures 7.5 (Gain) and 7.6 (Input Noise). The first plot shows that the mean gain of each ABC130 chip on the stave is little changed with respect to the off-stave measurements. The second plot shows that the noise of the MS modules built with unbonded mechanical sensors is also little changed, however the noise of the SS modules built with real sensors shows a clear increase.

Figure 7.7 plots, for the two SS modules only, the ratio of the Gain of each channel measured on-stave to the Gain of the same channel measured off-stave. The top four plots refer to Module 0, furthest away from the EoS card, and the bottom four plots refer to Module 12, adjacent to the EoS card. Each plot represents data for one column of 1280 strips. The plots are displayed in a similar orientation to the geometry within the module,




Figure 7.7: Gain Ratio on/off the Stave. For each channel of the two SS modules with real sensors, the ratio of Gain measured on-stave to Gain measured off-stave is plotted. Hybrids 0 and 1 refer to Module 0, hybrids 24 and 25 to Module 12. In both cases Column 1 refers to the bank of strips which runs away from the hybrid whilst Column 0 refers to the bank of strips which runs underneath the hybrid and, in part, the powerboard.

with those labelled as “Col 1” (Column 1) running away from the module centre whilst those labelled as “Col 0” (Column 0) run towards the centre of the module, under their respective hybrids and, in part, the powerboard. Each mark on the vertical axis represents the boundary between ABC130 chips. Excepting a few bad channels, for which an off-scale value of zero is shown, it can be seen that the highest Gain Ratio of any channel is 1.08, and for most channels the Gain Ratio is much closer to one. About one quarter of the chips stand out as having a greater variation in Gain Ratio channel to channel than their neighbours, for example channels 768 to 896 of Hybrid 0, but this is still a variation of less than 8%. The reasons for this variation have not been understood.

Figure 7.8 shows, in a similar format, the ratio of the Input Noise of each channel measured on-stave to the Input Noise of the same channel measured off-stave. Firstly it can be seen that the baseline noise increase of Module 12 is larger than that of Module 0. Secondly there are spikes in the Input Noise Ratio occurring in the inner columns only, between channels 768 and 896 for Hybrids 1 and 25 and between channels 640 and 768 for Hybrids 0 and




Figure 7.8: Input Noise Ratio on/off the Stave. For each channel of the two SS modules with real sensors, the ratio of Input Noise measured on-stave to Input Noise measured off-stave is plotted. Hybrids 0 and 1 refer to Module 0, hybrids 24 and 25 to Module 12. In both cases Column 1 refers to the bank of strips which runs away from the hybrid whilst Column 0 refers to the bank of strips which runs underneath the hybrid and, in part, the powerboard.

24. Whilst these spikes, correlated with known weak points in the powerboard shield design, are sometimes visible in the off-stave results, they are much more obvious when the modules are tested on the stave.

It was hypothesised that the noise spikes may have been suppressed in the off-stave tests due to the presence of a large metal cooling block under the sensor, providing additional cancellation of the DC-DC converter’s magnetic field by means of Eddy currents. This hypothesis was later checked at Brookhaven National Laboratory (BNL)[91]. A sheet of double sided, copper clad PCB material was cut to fit between the sensor backplane and the cooling block. On one side the metal was removed from the region under the DC-DC converter (“perforated” side) whilst on the other side complete coverage was retained (“complete” side). It was confirmed that the attenuation of the noise spikes was better when the “complete” side was in contact with the sensor backplane than when the “perforated” side was in contact with the sensor. This is because the missing metal in the “perforated” side reduces the amount of electrically conductive material in which the cancellation of the

magnetic field by Eddy currents may occur. Hence we may expect any effects related to localised deficiency of the power board shield to be more visible on the stave.

As this stave had only two “real” modules, no further tests were performed at this stage. It had fulfilled its major objectives in so far as the segmented multi-drop bus concept had been shown to work with HCC130/ABC130 modules and the GBTx chip for the first time. Construction of additional modules for a fully populated, Double Sided UK stave now began.

7.3 UK Double Sided Stave

7.3.1 Build and Test Information

The first fully populated, double sided stave with segmented bus tapes and GBTx was assembled at RAL using UK Stave core number 16. Assembly was completed in April 2019 using the modules listed in Table 7.2, with assembly of the first side having been completed in February 2019. Every module was built using a real ATLAS12 Strip Sensor (SS). The minimum number of bond wires were removed³ from the modules as built in order to make the HCC130 IDs unique within each Clock Control Reject (CCR) segment but not within the stave, or even within a stave side. Hence it was necessary to redesign the ITSDAQ stave test firmware and software to support the routing of discrete chip configuration commands to each of eight separate CCR segments, four on each side of the stave, rather than sending multiple copies of the same configuration stream as had been done previously.

Low voltage power was again supplied to the common voltage rail which feeds both EoS boards at 1.5 V from a TTI MX180TP[87] power supply. Each stave side drew 11 V power supply from its own output on a dual channel TTI CPX400DP[73] model. High voltage was delivered by an ISEG[92] high voltage power supply crate housing three 16 channel modules rated to deliver up to 2 kV at 8 mA but limited to 500 V output, with the stave under test connected to a single card. Data acquisition was performed by ITSDAQ as described

³HCC130 IDs are set by adding wire bonds to pull the desired address lines down to zero.

Side	Position	HV	Name	AMAC ID	HCC0 ID	HCC1 ID	Tested
Master	0	0	RAL_SS11	0	5	8	Yes
Master	1	0	BHAM_SS02	1	26	28	Yes
Master	2	0	RAL_SS13	2	6	12	Yes
Master	3	1	SHEF_SS01	3	18	19	Yes
Master	4	1	LIV_SS04	4	14	15	Yes
Master	5	1	LIV_SS03	7	12	16	Yes
Master	6	1	RAL_SS12	11	9	17	Yes
Master	7	2	RAL_SS14	5	18	16	Yes
Master	8	2	SHEF_SS02	8	17	20	Yes
Master	9	2	BHAM_SS03	9	9	10	Yes
Master	10	3	BHAM_SS04	10	11	12	Yes
Master	11	3	BHAM_SS05	6	2	3	Yes
Master	12	3	RAL_SS15	12	20	19	Yes
Slave	0	0	CAM_SS04	0	19	20	Yes
Slave	1	0	LIV_SS07	1	25	27	Yes
Slave	2	0	CAM_SS03	2	17	18	Yes
Slave	3	1	GLA_SS04	3	15	12	Yes
Slave	4	1	GLA_SS03	4	13	14	Yes
Slave	5	1	CAM_SS02	5	17	18	Yes
Slave	6	1	CAM_SS01	6	19	20	Yes
Slave	7	2	GLA_SS01	7	6	8	No
Slave	8	2	GLA_SS02	8	11	10	Yes
Slave	9	2	SHEF_SS03	9	21	23	Yes
Slave	10	3	SHEF_SS04	10	22	24	Yes
Slave	11	3	LIV_SS06	11	18	20	Yes
Slave	12	3	BHAM_SS06	12	29	30	Yes

Table 7.2: A list of modules which became part of the UK Double Sided Stave. The table is separated into sections to match the CCR segmentation of the bus tape. All HCC130 IDs are unique within a Clock Control Reject (CCR) segment, but not between segments: addresses which occur more than once within a stave side are shown in bold.

previously, however operation of the AMAC1A chips was expanded to cover multiple stave sides by the addition of a custom FPGA Mezzanine Card (FMC), FMC-ITk-I2C[93]. This implements eight separate, multiplexed I2C buses, and is plugged directly into the GLIB. Configuration of the GBTx chip(s) still used a single USB-I2C adaptor[89], the presence of multiple devices being handled by physically swapping cables. The coolant was water at 17 °C.

7.3.2 Results: Master Side

The completed Master Side of the stave was made available for test in February 2019. Although data mappings were the same as for the electro-mechanical stave which had already been understood, it took a few days to fully debug the functionality which had been added to the data acquisition system to allow the reuse of HCC130 IDs across multiple CCR segments, whilst still being able to provide each hybrid with unique configuration data. It was also necessary to fine tune the configuration of some of the hybrids, adjusting clock and sample phases, and changing the direction of data transfer within the hybrid in order to achieve reliable data transmission.

Figure 7.9a shows the default data transfer topology within an ABC130 Barrel Hybrid, with every ABC130 chip being configured to send its data to its left hand neighbour: only the leftmost chip of each group of five transmits data to HCC130. Figure 7.9b illustrates the case where chip 23 has failed: in order to recover data from chips 24 and 25, it is necessary to change the direction in which they transfer their data. In practice the (software) event decoder for a given hybrid will stall if it fails to receive a complete event from every ABC130 chip, meaning that it is necessary to investigate any such failing links by means of a binary search, progressively excluding chips from readout until a working configuration is established. One may then proceed to attempt to recover the remaining chips by adding them back, changing their direction of data transfer. If the problem is a failure of data transfer at a particular point in the chain, for example a bond between chips, then it should be possible to recover all chips. If however the problem is a faulty chip, or perhaps a missing




Figure 7.9: Data Transfer Architecture within the ABC130 Barrel Hybrid. In order to minimise the impact of any failed chip or interconnection, each ABC130 may transmit its data to the left or to the right. Figure (a) shows the default case where the ABC130 chips are all configured to transmit data to the left, sending data to HCC130 as two loops of five chips each. Figure (b) shows the potential impact of the failure of chip 23: by changing the direction of readout of the right most two chips, a total of nine ABC130 chips may still be read out.

or degraded clock or control signal, then this may not be possible. (Although the detailed circumstances of each case where a chip had to be excluded from readout were not fully investigated, this is not a particular concern for production modules as the readout topology will be changed to a “Star” topology, in which each ABCStar front end chip sends data directly to the HCCStar Hybrid Controller Chip.)

Another issue with this stave side is that the silicon sensor of Module RALSS12 was initially subject to breakdown at 50 V, increasing to 350 V after flushing the test enclosure with dry nitrogen for some days. For the duration of the initial testing of this stave side the voltage applied to HV segment 1 was therefore limited to 350 V⁴.

The results of a Response Curve test for the Master side of the UK Double Sided Stave are shown in Figures 7.10 (Gain) and 7.11 (Input Noise), together with comparison data for the same modules taken off-stave. The first plot shows that, as expected, the mean gain of each ABC130 chip on the stave is little changed with respect to the off-stave measurements. The second plot shows that the noise of the modules shows a clear increase, with input noise measured on-stave being, in most cases, a factor of 1.1 higher than when measured off-stave. For the three modules nearest EoS (chip number > 400) the increase is even higher, for some

⁴Later, as part of the Three Stave Test, this segment was successfully taken to the 400 V target.




Figure 7.10: A comparison of Gain for the modules on the Master Side of the UK Double Sided Stave, measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are full Response Curve tests taken at RAL under Standard Burst Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 192$). The gain does not change when the modules are mounted on the stave.

chips up to a factor of 1.3 increase is observed.

Figure 7.12 presents the Input Noise ratio (x) as a function of Strip Number (y) for the inner columns of the modules on the master side of the UK Double Sided stave, the strips which run under the hybrids and powerboard. From this plot we learn three things. Firstly every “Hybrid 1” plot shows a localised increase in noise somewhere between strips 768 and 896. This is due to deficiencies of the shielding around the DC-DC converter on the powerboard as discussed in Section 7.2.2. Secondly many “Hybrid 0” plots show greater noise increases for the five chips above channel 640 than they do for the five chips below channel 640. This is especially true for the three modules nearest the EoS card. Thirdly it becomes clear that four chips of module 1 Hybrid 0 could not be read reliably on the stave and had to be disabled. The increased noise of the modules nearest EoS will be further explored in the next chapter.




Figure 7.11: A comparison of Input Noise for the modules on the Master Side of the UK Double Sided Stave, measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are full Response Curve tests taken at RAL under Standard Burst Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 192$).

7.3.3 Results: Slave Side

Assembly of the Slave Side of the stave was completed, and the stave made available for test, in April 2019. The setup was modified to use the FMC-I2K-I2C board to provide dual 1.2 V I2C interfaces, one to configure the AMAC1A chips of each stave side, but was otherwise unchanged.

The stave is fitted with four modules assembled at Cambridge and it transpires that the three of them have large numbers of dead (zero gain or zero noise) channels. The electrical isolation of the test fixture used with the first three modules was inadequate: when breakdown occurs between the sensor backplane and ground, the resultant charge injected into the front end amplifiers can overwhelm ABC130's input protection resulting in damage. Examples of threshold scans for these modules, in each case for the strips running underneath the hybrids, are given in Figure 7.13. It can be seen that channel loss does not affect all chips equally and occurs mostly at the centre of each chip. The latter effect is an artefact of the way power and ground routing is done within the chip: the channels towards the centre





Figure 7.12: Input Noise ratio (on-stave / off-stave) is plotted against strip number for the inner columns of the modules on the Master Side of the UK Double Sided Stave. Both measurements are full Response Curve tests taken at RAL under Standard Burst Conditions (nStrobed = 256, nTrigs = 192).

of the chip exceed the damage threshold before those towards the edges. Furthermore the damage threshold may vary from die to die, hence some die are affected more than others. Recurrence of this problem during the production phase will be avoided by redesign of the module test tooling to provide better insulation: in addition the next generation of the front end chip, ABCStar, has been redesigned to offer more resilient input protection. Meanwhile additional insulation was applied to the test fixture such that the problem did not recur.


The initial results of a Response Curve test for the Slave side of the UK Double Sided Stave are shown in Figures 7.14 (Gain) and 7.15 (Input Noise), together with comparison data for the same modules taken off-stave. The first plot shows that, as expected, the mean gain of each ABC130 chip on the stave is little changed with respect to the off-stave measurements, however the second plot shows a number of interesting noise effects which need to be discussed. Note that none of these effects may be directly attributed to the damaged channels of the affected modules built at Cambridge, as known bad channels have been excluded from the calculations.

It has been shown in Section 6.3.2 that the noise of a Silicon Microstrip module varies according to the bias voltage, so one possible reason for the additional noise may of course have been that sensors did not receive bias as intended. Furthermore problems with the resistance of sensor backplane connections made using silver loaded epoxy adhesives were seen during construction of the SCT [53], hence the resistance of these connections on the Slave Side of the UK Double Sided stave was investigated further.


Each silicon sensor is of course a large area photo diode and hence will generate a voltage when exposed to light. All of the sensors on the stave were made to the same design and should therefore deliver similar photo voltages in response to similar levels of illumination: modules which deliver especially low photo voltages most likely have failed or failing backplane connections. The photo voltage generated by each sensor was measured using a Digital Volt Meter (DVM). This was done by using the DVM probes to contact the bias filter capacitor on each module's powerboard, with the cover having been removed from the test frame and with the stave unpowered. The results of this study are shown in the second




(a) CAM_SS01 Hybrid 0




(b) CAM_SS01 Hybrid 1




(c) CAM_SS02 Hybrid 0




(d) CAM_SS02 Hybrid 1




(e) CAM_SS03 Hybrid 0



(f) CAM_SS03 Hybrid 1



(g) CAM_SS04 Hybrid 0



(h) CAM_SS04 Hybrid 1

Figure 7.13: Threshold Scans with 1 fC charge injection for the four Cambridge Modules mounted on the Slave Side of the Double Sided Stave. Data is presented for the column of strips going underneath each hybrid. Three modules have large numbers of dead channels caused by breakdown between the sensor backplane and the module test fixture.




Figure 7.14: A comparison of Gain for the modules on the Slave Side of the UK Double Sided Stave, measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are full Response Curve tests taken at RAL under Standard Burst Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 192$) before repair of the module HV connections. The gain does not change when the modules are mounted on the stave. (Note that off-stave data is not available for Module 7.)

column of Table 7.3, “Initial Photo V”. It can be seen that modules 11 and 12, found to have anomalous high noise during the first Response Curve Test of this side, deliver much lower photo voltages than the other modules.

To further investigate this effect, for each module where the photo voltage was less than 350 mV the DVM, set to diode test mode, was used to force current through the conductive glue to see if this may reduce its (assumed) resistance. One meter lead was used to contact a piece of foil which had been inserted under the sensor under test, thereby making a direct connection to the backplane. The current would then flow through the backplane metallisation and through the conductive glue dot onto the HV trace on the bus tape. The circuit is then completed by a wirebond up to the power board and through copper trace to the filter capacitance, the second meter lead was applied to its HV terminal. It can be seen in the third column of Table 7.3, “Photo V after DVM Treatment”, that a significant increase in the measured photo voltage was observed in each case, hence the resistance had been decreased. Unfortunately the effect was found to be temporary, lasting no longer than




Figure 7.15: A comparison of Input Noise for the modules on the Slave Side of the UK Double Sided Stave, measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are full Response Curve tests taken at RAL under Standard Burst Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 192$) before repair of the module HV connections. Several modules have increased noise in excess of the factor of 1.1 experienced on previous stave sides: this is addressed in the text below. (Note that off-stave data is not available for Module 7.)

ten minutes.

Some hours later the resistance of each module's backplane connection was measured, connecting the DVM between the inserted backplane foil and the HV terminal of the filter capacitor. These results are shown in column four of Table 7.3, being found to vary between $67\ \Omega$ and $375\ \text{k}\Omega$. For each module found to have a backplane resistance in excess of $200\ \Omega$ a more aggressive treatment was applied, using a bench power supply to push $100\ \text{mA}$ through the backplane connection for 4 seconds. During the treatment the positive supply was connected to the filter capacitor and the negative supply to the sensor backplane and the voltage output of the supply was limited to $1.0\ \text{V}$, generally falling to $80\ \text{mV}$ once current began to flow. A few minutes later the backplane contact resistance of the treated modules was remeasured, results which are shown in column five of Table 7.3. It can be seen that there was a substantial improvement in each case, in some cases a reduction of five orders of magnitude and with all treated connections now having less than $20\ \Omega$ resistance.

Module	Initial Photo V	Photo V after DVM Treatment	Initial Backplane Contact Resistance	Resistance after Application of 100 mA	Final Photo V
0	375	-	18k	2.5	370
1	397	-	150k	3	410
2	260	375	300k	10	370
3	414	-	350	3	410
4	417	-	110	-	410
5	370	-	69	-	370
6	420	-	100	-	410
7	345	-	67	-	350
8	395	-	300	3	400
9	410	-	675	15	410
10	331	410	375k	5	410
11	90	420	375k	7	420
12	120	400	375k	8	400
	mV	mV	Ω	Ω	mV

Table 7.3: A survey of the photo voltage produced by each sensor of the Slave Side of the Double Sided stave was carried out, with three modules delivering particularly low voltages. A temporary improvement was made by pushing a small current through the connection, made by conductive glue, using a Digital Volt Meter (DVM) set to diode test mode. A bench power supply was then used to push 100 mA through each connection, which resulted in a (more) permanent reduction in the resistance of the conductive glue. The final column shows the photo voltages achieved after the treatment had been completed.

Finally the photo voltage of each sensor was remeasured. These values, presented in the final column of Table 7.3, show that the photo voltage delivered by those sensors which initially gave anomalous low values has now increased to acceptable levels. Furthermore these values were retained after several days: at this stage the effect of the treatment was expected to be permanent⁵.

With the sensor backplane connections having been repaired, the Response Curve test was repeated upon the Slave side of the UK Double Sided Stave. Results are shown in Figures 7.16 (Gain) and 7.17 (Input Noise), together with comparison data for the same modules taken off-stave. As seen previously the mean gain of each ABC130 chip on the

⁵It is important to note that conductive epoxy will not be used to make the sensor backplane connections on production staves. This technology will be replaced by an aluminium foil or “tab” which will be ultrasonically bonded to the sensor backplane during module production. Once the module has been glued to the stave core, the tab will also be ultrasonically bonded to the HV pad on the bus tape.




Figure 7.16: A comparison of Gain for the modules on the Slave Side of the UK Double Sided Stave, measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are full Response Curve tests taken at RAL under Standard Burst Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 192$) after repair of the module HV connections. The gain does not change when the modules are mounted on the stave.

stave is little changed with respect to the off-stave measurements. The noise performance has been improved as a result of the currents passed through the conductive glue dots, with the ratio between Input Noise measured on- and off- stave now being generally similar to the factor 1.1 observed for the Master side of the stave. There remains an increase in noise for the last few modules, but a similar pattern was observed for the Master side. Module 6 (CAM_SS01) remains an outlier with high noise: as this is one of the modules damaged during production tests, with a large number of anomalous channels which may not always be correctly identified and excluded by the analysis code, this is not a major concern.




Figure 7.17: A comparison of Input Noise for the modules on the Slave Side of the UK Double Sided Stave, measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are full Response Curve tests taken at RAL under Standard Burst Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 192$) after repair of the module HV connections. With the exception of chips 240 to 270 corresponding to Module 6 (CAM_SS01) the input noise has been greatly reduced.

7.4 US Single Sided Stave

7.4.1 Build and Test Information

This stave was assembled at Brookhaven National Laboratory (BNL) using US Stave core number 7. One feature of this stave core was that the “closeout” at the end of the stave away from the EoS board was incorrectly handed, meaning that if a module were to have been mounted at position zero its sensor would have partially overhung the end of the stave. For this reason it was decided not to mount a module at that location.

During tests of this Stave at Brookhaven a simpler powering configuration was used, with all four HV buses connected to a single Keithley 2410 Source Measure Unit capable of delivering 1.1 kV. A Sorensen XPF60-20DP power supply [94] - a rebranded version of the TTI CPX400DP used at RAL - was used to deliver 11 V to the stave with 1.5 V being supplied by an older Instek unit. Data acquisition was again provided by ITSDAQ running

Position	HV	Name	AMAC ID	HCC0 ID	HCC1 ID
0	0	N/A	-	-	-
1	0	SCIPP-EL-07	1	3	4
2	0	SCIPP-EL-09	2	5	16
3	1	SCIPP-EL-10	3	17	20
4	1	SCIPP-EL-11	4	18	19
5	1	SCIPP-EL-12	5	21	22
6	1	SCIPP-EL-06	6	2	9
7	2	LBL-05	7	11	7
8	2	LBL-10	8	13	6
9	2	LBL-11	9	14	15
10	3	LBL-07	10	23	27
11	3	LBL-06	11	30	29
12	3	LBL-04	12	38	8

Table 7.4: A list of modules which became part of the US Single Sided Stave. The table is separated into sections to match the CCR segmentation of the bus tape. All HCC130 IDs are unique within the stave. Module 0 was not loaded.

on GLIB, the same setup used at RAL. Configuration of AMAC1A was possible either direct from ITSDAQ or from a Raspberry Pi. The coolant was water at 15 °C.

7.4.2 Results

The stave was first made available for testing at BNL in July 2018, with six modules loaded, then again in September 2018 with 12 modules loaded. Figure 7.18 plots the Gain of the 12 modules running together on the stave against their individual module reception test data, recorded at BNL before module mounting took place. Two modules stand out, Module 3 (chips 80-119) and Module 9 (chips 320-360)⁶. For the duration of the initial stave tests the identities of the two hybrids of module 9 were swapped in the system configuration file. This was first realised during the preparation of Figure 7.18, in which the order of these two hybrids has been corrected. A channel by channel analysis looking for gain correlations confirms that the revised hybrid mapping is correct and the similarity of the channel mean gain values has been improved, however it is not as good as usual as a further consequence of the mapping error may have been the failure of the strobe delay adjustment. At the time

⁶Recall that module zero was not fitted to the stave, hence for this stave only chip 0 in the plots corresponds to the first chip of Module 1.




Figure 7.18: A comparison of Gain for the modules of the US Single Sided Stave, measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are Three Point Gain tests taken at BNL under Standard Burst Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 192$). It can be seen that, unlike the UK stave results, the Gain measured on this stave is typically 2.5% lower than that recorded during the module reception testing.

of writing the anomalous gain of module 3 has not been understood, however this hybrid is missing from several alternate datasets it was a frequent cause of data transmission errors. Hence it is likely that the on-stave data is erroneous.

For all modules mounted on the UK single sided and double sided staves the ratio of the gain measured on- and off- stave is essentially unity, however Figure 7.18 shows that for the modules mounted on the US stave this ratio is approximately 0.95 for all except Module 3 (chips 80-119). For all other modules the gain recorded on-stave is systematically lower than that recorded off-stave. As the on-stave gain values for the US stave are similar to both the on-stave and off-stave values for the UK staves, the gain values from the US module reception test may be considered anomalous. It has been confirmed that the burst conditions (n_{Strobed} , n_{trigs}) were indeed the same for the two data sets: the reasons for the higher gain of the off-stave results recorded at BNL has not yet been understood.

Figure 7.19 plots the Input Noise of the 12 modules running together on the stave against




Figure 7.19: A comparison of Input Noise for the modules of the US Single Sided Stave, measured on and off the stave, presented together with the ratio between the two measurements. Both measurements are Three Point Gain tests taken at BNL under Standard Burst Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 192$). With the exception of Module 2 (chips 40-79), Input Noise measured on this stave is higher than that recorded during the module reception testing: for many modules the increase is consistent with the decreased gain shown in Figure 7.18, although the modules nearest the EoS card show noise increases in excess of this, consistent with the other staves.

their individual module reception test data. Module 2 (chips 40 to 79 SCIPP-EL-09) has anomalously high off-stave noise values, perhaps due to a lower bias voltage being used during the test. The second hybrid of Module 3 (chips 100 to 119 SCIPP-EL-10) also has a higher scatter than usual but as stated previously this hybrid was associated with a high probability of data transmission errors. Otherwise the noise results conform to the same pattern as the other staves, with the ratio between on-stave and off-stave Input Noise values being of order 1.1, with some increase in the module(s) nearest the EoS card.

7.5 Chapter Summary and Conclusions

For the purposes of this study, a total of three staves with non-segmented bus tapes were built and populated with ABC130 short strip modules. Of these staves the first used mostly

“Electro-mechanical” modules fitted with dummy sensors although two modules were fully functional silicon microstrip detectors. Of the remaining staves, one has modules on both sides but the other is single sided.

With the exception of the stave assembled in the USA, a comparison of gain measurements made on- and off- stave shows they are unchanged. (In the case of the US stave measurements, it is the off-stave results which are anomalous.) In general terms the noise of the modules measured on-stave is a factor of 1.1 higher than when measured off-stave, with additional noise increases in the modules nearest the EoS board: this will be further investigated in the following chapter. There are also a couple of outliers which correlate with other problems of various kinds: these effects are not directly related to the stave.

With three functional staves in hand, the next stage of the programme was to try to operate them together.

Chapter 8

The Three Stave Test

This chapter documents the specification, execution and performance of the Three Stave Test.

8.1 Introduction

The ATLAS ITk Strip Barrel comprises four layers. An end view of this arrangement is shown in Figure 8.1, although not all staves are shown.

Figure 8.2 gives a more detailed view of three staves of the outermost layer, Layer 3, although electrical services and cooling are not shown. Each stave is supported by five mounting points, shown in green, each comprising a bracket attached to the stave and a housing which is glued onto the carbon fibre support cylinder. Using custom tooling, each stave is pushed into position from the cylinder end until its integrated mounting brackets, shown in dark green, are aligned with the support housings. At this point a 1.4 m long hex key will be used to rotate the locking cam which forms part of each housing, shown in purple, such that the stave is locked into position. The insertion tooling may now be removed, leaving the stave at a tilt angle of (for layer 3) 11 degrees. Looking at Figure 8.2 one may also discern that there is overlap between opposing staves, hence the sensors and hybrids of the modules of the lower face of each stave directly face the modules on the upper face of its neighbour. In order to maintain the hermeticity of the detector for for all tracks with transverse momenta above 1 GeV [41], the region of overlap for L3 covers approximately

1 cm and at the point of closest approach, the sensors are separated by 11.87 mm. Numeric parameters for this and other layers are given in Table 8.1.




Figure 8.1: The drawing gives an end view of the proposed four barrel structure, with services omitted. During stave insertion the barrel interlinks (shown in red) must be removed from the affected area. The diameter of the innermost and outermost barrels are 701.4 mm and 1920.8 mm respectively.

8.1.1 The Support Structure and Stave Installation

In order to facilitate investigations of possible cross talk between adjacent staves a space frame structure was constructed as shown in Figure 8.3. This allows three staves to be oriented across a range of angles and separations representative of the final structure. Rather than the edge insertion method which will be used in the experiment, for the purposes of this test the staves were installed by hand. This is practical only because the Three Stave Test Structure permits the separation between the staves to be increased to give easy access.

Each stave has five mounting brackets glued along one edge. In turn these are supported by five housings, made to the final design, which are glued to aluminium plates. The aluminium plates are free to slide along an extruded aluminium support beam until their bolts




Figure 8.2: The drawing gives an end view of three staves of the outermost layer, with services omitted. There is no direct overlap between any DC-DC converter and the sensor of an opposing stave. A small overlap between sensors is necessary to provide hermetic tracking coverage: the minimum separation between opposing sensors for the featured layer (L3) is 11.87 mm, however for the innermost layer (L0) it is 9.25 mm. An explanation of the stave support mechanism comprising Bracket, Locking Cam and Housing, is given in the text.

are tightened. Prior to mounting a stave, the support housing nearest the $z = 0$ end (away from the EoS card and services) is fixed at the desired location. The remaining support housings are moved a short distance from their target positions and left free to slide. The stave being inserted is then placed into the position defined by the fixed housing, after which the remaining housings are carefully slid into position. Next each mounting point is locked by inserting a hex key into its bracket's locking cam and rotating it to the locked position. Finally the remaining housings are made secure by tightening the corresponding bolts.

Layer	Active Stave Radius	Tilt Angle	Sensor Separation
L0	399	13	11.87
L1	562	12	10.49
L2	762	12	10.05
L3	1000	11	9.25
	mm	degrees	mm

Table 8.1: Barrel Layer Parameters. For each of the four barrel layers L0 to L3, three parameters are shown: Active Stave Radius, the radius to the geometric centre of the stave; Tilt Angle, the angle at which the staves are tilted; Sensor Separation, the closest approach of a silicon sensor of one stave to a silicon sensor of the neighbouring stave.




Figure 8.3: CAD model of the Three Stave Support Structure. A space frame approach is taken in which the angles of the staves and the separations between them may be adjusted to orient the staves in a variety of realistic configurations.

8.1.2 Services

The service arrangement for two staves is shown in Figure 8.4. Power is delivered to each prototype stave through a pair of commercial D17W6 connectors situated on the Master EoS pigtail. The connector comprises 4 high voltage pins, 2 high current pins and 7 standard pins. Each connector serves one side of the stave with: 11 V power and return (2 high current pins); 11 V voltage sensing (2 standard pins); four HV bias lines (4 high voltage pins); a common HV return (1 standard pin); one NTC connection (2 standard pins) and two unused pins. In the original single stave test stand, it was possible to locate the low voltage supply immediately adjacent to the stave such that cables of only 60 cm length could be used. For the Three Stave Test stand this was impractical: power supplies were placed in

an adjacent rack and connected to each stave's Master EoS pigtail by 3 m long cables. Each stave connector's sense pins were routed back to its power supply's sense inputs to maintain the correct voltage upon entry to the bus tape.

Final staves will generate low voltage power for their lpGBT chips at 1.2 V from the 11 V supply by means of a pair of additional DC-DC converter boards mounted within a cut-out on the Master EoS, however these were not available for these tests. Instead the Master EoS provides a two pin jumper where 1.50 V may be fed in to power the GBTx chips: this is used for two of the staves but for the UK electromechanical stave wires were soldered directly to the PCB. In all cases one feed and one return wire leave each stave and run for approximately 60 cm before connecting to two pairs of wires (one pair for the current and one pair for sensing) which run back to the power supplies. Under load, voltage drops will occur in the last 60 cm of these cables, notably in the case of the double sided stave where the current is doubled. In the latter case the power supplies were set to deliver 1.56 V to compensate for this effect. A pair of TTi CPX400DP [73] dual channel power supplies were used to provide each of the four instrumented stave sides with 11 V; a single TTi MX180TP [87] triple channel supply was used to provide each stave with 1.5 V.

Although the low voltage cabling is not screened, the high voltage cabling is. A high voltage patch box with D50 connector is placed within the three stave test enclosure: a multicore high voltage cable with overall braided screen runs from this point back to the ISEG High Voltage Crate, and from this point up to the D17W6 connectors which mate with the EoS pigtails individual screened twisted pairs are used. The patch box becomes in effect the Star point with respect to the grounding of the staves, the point at which a (ground) reference is provided by the screened HV cables to which the otherwise floating low voltage returns are referenced upon entry to EoS. The low voltage power supplies and the high voltage crate are connected to a local network, as is the data acquisition computer by means of a dedicated Network Interface Card (NIC): in this manner the power supplies may be operated from the ROOT (ITSDAQ) console if required.

Additional electrical services were required to provide the I2C buses needed to configure

the GBTx chips and the AMAC chips on the module powerboards. The source of the configuration signal for the GBTx chips was the custom USB-I2C adaptor[89] used previously however its output, now raised to 3.3 V, was connected to an FMC-ITk-I2C [93] multiplexer board located underneath the stave test enclosure which provides four multiplexed outputs at the correct voltage of 1.5 V. Custom cables were used to route the signals from here to each EoS board directly, and the GBTx configuration utility provided by CERN was modified to support operation of the FMC's I2C multiplexer chips. A second FMC-ITk-I2C board was used to provide four I2C buses to configure the AMAC chips: once more this was placed under the Three Stave Test enclosure although in this case the signal to drive them was sourced from the FPGA board used for data acquisition. Ribbon cables were used to route these signals to the Master EoS of each stave.

A pair of manifolds located under the Three Stave Test enclosure was used to ensure that each stave receives equal coolant flow, water at 17 °C. Although the humidity dependence of the breakdown of the ATLAS12 strip sensors is well known [95], for the duration of the work reported here the test enclosure was not fitted with any humidity monitoring system. However whilst running, the enclosure was flushed with dry nitrogen at up to 10 l/min to keep the sensors sufficiently dry to reach 400 V bias without breaking down. Without the dry gas flow it was not possible to run two of the sensors on the UK Double Sided Stave over depletion.




Figure 8.4: Services for Two Staves. Labelled services include: Cooling Water (6 mm diameter red pipes); Low Voltage power at 1.5 V for GBTx (red/black twisted pair); Low Voltage power at 11 V for the modules (thick grey cables); High Voltage sensor bias (white screened twisted pair cables); Optical Fibres with aqua or purple jackets; I2C signals to program AMAC (ribbon cables); and the High Voltage Patch box which serves as the grounded point to which each stave is referenced. In addition thin, grey “phone” cables (unlabelled) carry sense wires and NTC connections for the staves, and I2C signals to program GBTx.

8.1.3 Data Acquisition and Debugging

The GLIB board used for previous stave tests did not have sufficient resources to accommodate a firmware build for four stave sides, for this reason the system was migrated to the commercial Genesys 2 [96] FPGA development board. As this board does not support on-board optical transceivers, the commercial HTG-FMC-X4SFP+ FPGA Mezzanine Card (FMC) [97] was used to provide four SFP+ cages. Commercial transceivers, AVAGO (now Broadcom) part AFBR-709SMZ [98], were used to physically implement both ends of each link. The opportunity was taken to move to the most recent version of the “GBT core” firmware rather than the early version which had been used previously.

Scaling the system to support so many channels was not without difficulty. The first issue encountered was the difficulty of bringing up four GBT links at the same time. The initial system design had foreseen that the four firmware instances of the GBT core would be reset simultaneously, both for transmit (TX) and receive (RX), but this was seldom sufficient to bring up all four links. Instead the firmware was revised to make accessible the individual TX and RX resets for each GBT core. Software was written to continually reset the transmit part of a link until TX lock was achieved, before proceeding to continually reset the receive part of a link until RX lock is achieved. Repeating this process over all four links (stave sides) from a cold start can take of order 40 seconds¹. Once lock has been achieved, three of the four links will happily keep running for several days at a time and perhaps longer, however one stave side tends to drop the link every few hours which makes it very difficult to run any prolonged test. This would appear to be a property of the stave rather than the link, in so far as fibre and channel swaps were unable to resolve the matter. Similarly the voltages delivered to the stave were found to be correct, and in any event increasing them within safe limits had no effect. Thankfully the affected stave is the single-sided electromechanical stave which only has two modules fitted with real silicon sensors, and in any event its instrumented side faces away from the adjacent stave so we would not expect any crosstalk between this stave and any other. Hence the decision was taken not to debug this problem further, and

¹Although beyond the scope of this thesis, forming links with staves using the later lpGBT chip, as will be used for production staves, is a much more reliable procedure.

to bring the system up running three stave sides: the UK Double Sided stave and the US Single Sided Stave.

A second problem arose from the rate at which the ITSDAQ software is able to process packets which arrive across the network. When the system is started, the software sends a command to the firmware to ask which streams it supports, where two streams are required to receive data from each hybrid². The firmware responds by sending one “stream config” packet across the network for each supported stream: in the case of firmware built for four stave sides, 224 packets are sent in quick succession. The initial observation was that the software would start up with a random selection of missing streams, meaning that some number of hybrids could not be read out. After some investigation it was established that all 224 stream config packets were being received by the linux kernel, however they were not reaching ITSDAQ because it could not process the packets sufficiently quickly to prevent overflow of the operating system’s receive buffer. This was resolved by increasing the size of the buffer to 10 MB, although in reality a lower value may have been sufficient.

The third problem was one of data loss when running scans. This time diagnostics suggested that packet loss may now be occurring within the kernel, which by default is configured to buffer up to 1000 incoming User Datagram Protocol (UDP) packets: if this is exceeded packet loss will occur³. However the issue was not resolved by adjustment of the corresponding system parameter, so the cause was elsewhere. The problem was finally traced to the firmware design which had assumed that the open source UDP implementation which it uses would be able to receive data from the modules and transmit it to the network as fast as the modules could deliver it. With so many modules feeding data into a single 1 Gbit/s network link this was no longer the case: corrupt UDP packets were being transmitted which the kernel was unable to process. Within the firmware, an additional flow control mechanism was added such that data was only passed to the UDP core when it was able to accept it. Hence the issue was resolved in December 2019, and data taking could finally begin.

²Recall that the HCC130 multiplexes the data streams from two loops of five chips. Upon receipt, ITSDAQ demultiplexes this into one pair of streams.

³A different data acquisition system, FELIX [99], will be used in the experiment. Rather than UDP this will use the reliable NetIO [100] protocol for the network transmission of physics data.

Stave Side	GBTx V	GBTx I	Stave V	Stave I	Modules
UK SS Master	1.51	657	11.00	0.75	Unpowered
UK DS Master	1.56	1831	11.00	6.51	Powered and Configured
UK DS Slave	-	-	11.05	6.29	Powered and Configured
US SS Master	1.50	910	11.00	5.66	Powered and Configured
	V	mA	V	A	

Table 8.2: Stave Low Voltages and Currents: the voltages delivered to each stave side and the currents drawn during normal operation. Note that the GBTx chips of the double sided stave share a common 1.5 V power rail which is running 60 mV over voltage to compensate for voltage drops in the cabling. As the GBTx of the UK Single Sided Stave would frequently lose its lock its GBTx chip was generally left powered but unconfigured and its modules unpowered, hence the currents drawn by this stave are lower than the others.

Stave Side	Position	Module	Hybrid	nExcluded Chips
UK DS Master	1	BHAM_SS02	0	4
UK DS Slave	7	GLA_SS01	1	1
UK DS Slave	8	GLA_SS02	1	1

Table 8.3: The table lists the ABC130 chips which had to be excluded from readout in order to achieve error free operation.

8.1.4 Operating Conditions and Standard Tests

The voltages supplied and currents drawn by the staves are shown in Table 8.2. Although it had by now been decided to concentrate upon the UK Double Sided and US Single Sided staves due to the high frequency with which the UK Single Sided Stave would drop its lock with the incumbent GBT signal, it was still necessary to provide power to the UK Single Sided Stave as an unpowered GBTx chip could adversely affect operation of the I2C multiplexer used to program the other three GBTx chips. Hence the currents drawn by this stave are much lower. The GBTx chips on both sides of the Double Sided Stave share a common 1.5 V supply which draws almost exactly twice the current of that drawn by the US Single Sided Stave: the output of its power supply had to be increased to 1.56 V to compensate for voltage drops in the cabling.

In addition to the exclusion of the UK Single Sided stave, a further six ABC130 chips had to be excluded from readout in order to achieve reliable operation. These are shown in Table 8.3. With this set of exclusions in place, data acquisition was robust and error free.

Scan Index	Test	nStrobed	nTrigs	Trimmed at
1	Strobe Delay	256	32	1 fC
3	Three Point Gain	256	192	1 fC
6	Three Point Gain	32	512	1 fC
9	Response Curve	256	192	1 fC
19	Response Curve	32	512	1 fC
29	Noise Occupancy	N/A	10000	1 fC
30	Pedestal Trim	N/A	100	N/A
31	Response Curve	32	512	0 fC
41	Noise Occupancy	N/A	10000	0 fC

Table 8.4: The Three Stave Test Sequence. (The parameter nStrobed has no meaning for the Pedestal Trim or Noise Occupancy tests, as no charge injection takes place.)

For each mechanical configuration of the staves, the same set of tests was executed, chosen to permit a wide range of performance comparisons. This test sequence is shown in Table 8.4 and begins with a Strobe Delay Test as described in Section 3.3.2 such that the timing of the charge injection pulse is adjusted to compensate for temperature variations. Next the Three Point Gain Test as described in Section 3.3.3 is executed twice, first under Standard Burst Conditions (nStrobed = 256 and nTrigs = 192) and then again under Optimised Burst Conditions (nStrobed = 32 and nTrigs = 512). Then the Response Curve Test described in Section 3.3.5 is executed twice, once with Standard Burst Conditions and then once with Optimised Burst Conditions, followed by the Noise Occupancy test as described in Section 3.3.6. Up to this point the trim DACs of the ABC130 chips have been configured with the result of a Trim Range test as described in Section 3.3.4, designed to maximise threshold uniformity in response to 1 fC injected charge. In order to better explore stave performance at lower thresholds, the next test in the sequence executes a Pedestal Trim as described in Section 3.3.9 which maximises threshold uniformity at a threshold equivalent to 0 fC injected charge. Next a Response Curve Test is taken with using Optimised Burst Conditions in order to re-establish the calibration of each chip with the new trim settings. Finally the Noise Occupancy scan is repeated in order to be better able to spot anomalies which may become visible at low thresholds.




Figure 8.5: Images of the First Three Stave configuration. Figure 8.5a shows the UK single sided electro-mechanical stave at the left, the UK Double Sided stave in the middle, and the US Single Sided stave at the right. Figure 8.5b gives an end-on view of the two staves of interest, the UK Double Sided stave at the left and the US Single Sided stave at the right. The separation between the centres of the stave supports is 237 mm.

8.2 Position 1: Separation of supports 237 mm, 12°

For the first tests the three staves were set to a tilt angle of 12 degrees and spaced approximately equally, as far apart as the services would permit, as shown in Figure 8.5a. The spacing between the support centres of the UK Double Sided and US Single Sided staves, shown end-on in Figure 8.5b, was measured as 237 mm. At such separation the effect any stave may have upon its neighbour should be minimised.

8.2.1 “Like for Like” comparison with Original Single Stave Results

The single stave results described in the previous chapter were all taken under Standard Burst Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 192$). For the US Single Sided stave the only available results for which all modules were operational are for a single Three Point Gain test, however for the other staves we have a choice of Response Curve or Three Point Gain results.

Figures 8.6 and 8.7 present the Gain and Input Noise ratios for the three operational stave sides running individually in the Three Stave test stand (numerator) and on the single

stave test bench (denominator). A like for like comparison is made, using Response Curve results for the UK Double Sided stave and Three Point Gain data for the US Single Sided stave. The tests chosen here use Standard Burst Conditions.

Certainly the gain of the modules measured under the same operating conditions should be identical. This is largely confirmed by Figure 8.6, the only module for which the gain has changed by more than 2.5% being Module 9 (chips 320 to 359) of the US Single Sided Stave. It is believed that in this case that it is the single stave result which is anomalous, certainly the gain measurements for this module during the initial on-stave measurement were unusually low as shown in Figure 7.18. Where data is missing from the plot, this is because a chip had been disabled or, as is the case with modules 2 (CAM_SS03), 5 (CAM_SS02) and 6 (CAM_SS01) of the UK Double Sided Stave Slave side, because less than half of its channels were working.

Figure 8.7 tells a slightly different story as regards the Input Noise of the staves. One may disregard Modules 2 (CAM_SS03), 5 (CAM_SS02) and 6 (CAM_SS01) of the UK Double Sided Slave side and Module 9 of the US Single Sided Stave for the reasons outlined above, however this still leaves seven modules for which the noise has increased by 5% or more and one hybrid for which the noise has gone down by 5%. Of the modules with increased noise one is on the Master side of the UK Double Sided Stave (Module 11); five are on the Slave side of the UK Double Sided Stave (Modules 2, 7, 9, 11 and 12); one are on the US Single Sided Stave (Modules 10). As regards the Slave side of the UK Double Sided Stave, four of the affected modules (2, 9, 11 and 12) had previously benefited from remedial treatment of the conductive glue used to make their sensor backplane connections as described in Section 7.3.3. It is possible that this treatment was not permanent after all, however the probing necessary to confirm this cannot be safely accomplished whilst the stave is mounted in the Three Stave Support structure. A less invasive possibility would be to measure module noise as a function of sensor bias as explained in Section 6.3.2, to determine if each sensor reaches full depletion at the expected voltage: unfortunately this was not possible for reasons of time. Although different conductive epoxies were used during the construction of the UK

(EPO-TEK H20E) and US (Loctite TRA-DUCT 2902) staves, an increase in the resistance of the backplane connection would be a possible explanation for each module for which a noise increase has been observed⁴.


8.2.2 Effect of Optimised Burst Conditions

Figures 8.8 and 8.9 present the Gain and Input Noise ratios for the three operational stave sides running individually with Optimised Burst Conditions (numerator, (nStrobed = 32 and nTrigs = 512) and with Standard Burst Conditions (denominator, (nStrobed = 256 and nTrigs = 192)). All data is taken from Response Curve Tests. Aside from chips 327, 337, 347 and 357 of the slave side of the UK Double Sided Stave, which are not bonded to the sensor, the Gain is uniformly increased and the Input Noise is uniformly decreased: this is to be expected when nStrobed is reduced from 256 to 32, as explained in Section 6.3.3. As the Optimised Burst Conditions give a more accurate result than the Standard Burst conditions, the results shown in the numerator become our new baseline, representative of initial stave performance in the Three Stave Test Stand, to which all other configurations shall be compared.


8.2.3 Effect of running all staves together

Figures 8.10 and 8.11 present the Gain and Input Noise ratios for the three operational stave sides running individually (numerator) and running together (denominator). All data is taken from Response Curve Tests taken under Optimised Burst Conditions (nStrobed = 32, nTrigs = 512). The Gain is unchanged. The only significant change in the Input Noise plot occurs for module 3 of the US Single Sided Stave which shows an increase of order 2.5%. Note however that this increase does not occur in any subsequent measurements: furthermore the results shown in the denominator were recorded three days before the results shown as the numerator, so it is possible that some kind of sensor training effect came into play here.


⁴Recall that it is planned to replace the electrically conductive epoxy used to make the sensor backplane connection by an ultrasonically bonded aluminium foil in production staves.



(a) UK Double Sided Master Side




(b) UK Double Sided Slave Side




(c) US Single Sided Master Side


Figure 8.6: The ratio of the Gain, measured in Position 1 running each stave side separately, divided by the corresponding results for each stave as presented in the previous chapter. For the UK Double Sided stave, both results are Response Curves and for the US Single Sided stave both results are Three Point Gain tests. All tests were recorded under Standard Burst Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 192$). Apart from chips 320 to 359 of the US Single Sided Stave (red box), for which the Single Stave result is believed to have been anomalously low, the gain remains essentially constant.



(a) UK Double Sided Master Side




(b) UK Double Sided Slave Side




(c) US Single Sided Master Side


Figure 8.7: The ratio of Input Noise, measured in Position 1 running each stave side separately, divided by the corresponding results for each stave as presented in the previous chapter. For the UK Double Sided stave, both results are Response Curves and for the US Single Sided stave both results are Three Point Gain tests. All tests were recorded under Standard Burst Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 192$). In the majority of cases the noise is little changed. Outliers are discussed in the text.



(a) UK Double Sided Master Side




(b) UK Double Sided Slave Side




(c) US Single Sided Master Side


Figure 8.8: The Gain measured running the three stave sides separately with Optimised Burst Conditions ($n_{\text{Strobed}} = 32$, $n_{\text{Trigs}} = 512$) divided by the Gain measured running the three stave sides separately with Standard Burst Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 192$), both in Position 1. The reduction in the number of channels which are strobed simultaneously gives rise to an increase in the measured gain as explained in 6.3.3. The gain of the channels represented by points 327, 337, 347 and 357 of the slave side of the UK Double Sided Stave does not change as they are not bonded to the sensor. Both sets of results are from Response Curve tests.



(a) UK Double Sided Master Side




(b) UK Double Sided Slave Side




(c) US Single Sided Master Side


Figure 8.9: The Input Noise measured running the three stave sides separately with Optimised Burst Conditions ($n_{\text{Strobed}} = 32$, $n_{\text{Trigs}} = 512$) divided by the Input Noise measured running the three stave sides separately with Standard Burst Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 192$), both in Position 1. The reduction in the number of channels which are strobed simultaneously gives rise to a decrease in the measured noise as explained in 6.3.3. The noise of the channels represented by points 327, 337, 347 and 357 of the slave side of the UK Double Sided Stave does not change as they are not bonded to the sensor. Both sets of results are from Response Curve tests.



(a) UK Double Sided Master Side



(b) UK Double Sided Slave Side




(c) US Single Sided Master Side

Figure 8.10: The Gain measured running the three stave sides together divided by the Gain measured running the three stave sides separately, both in Position 1. Both sets of results are from Response Curve tests under Optimised Burst Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 512$). The gain is unchanged.


(It is for this reason that the data set shown as the numerator is taken as the reference data set against which the performance in other stave orientations will be compared.)

8.2.4 Investigation of additional noise of modules near EoS


To date, all tests of fully or near fully populated stave sides have shown increased noise for the last three or four modules, adjacent to the EoS region. In order to explore this more fully a series of runs were taken in which, for each of the three active stave sides, only one module was powered at a time. The tests covered the three modules 9, 10, 11



(a) UK Double Sided Master Side



(b) UK Double Sided Slave Side



(c) US Single Sided Master Side


Figure 8.11: The Input Noise measured running the three stave sides together divided by the Gain measured running the three stave sides separately, both in Position 1. Both sets of results are from Response Curve tests under Optimised Burst Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 512$). One module exhibits higher noise when the three staves are operated together (red box), however as this does not occur in later data sets this is assumed to have been a temporary effect.

and 12. Figure 8.12 presents the ratio of the Gain measured during these tests divided by the reference Gain measured running each complete side separately: the gain is unchanged. Figure 8.13 similarly presents the Input Noise, which has increased for all modules in the range 1% to 5%. This immediately discounts the possibility that some kind of noise signal may be developing along the stave due to the operation of multiple modules.


Other factors which may vary according to the number of powered modules include the voltage delivered to the input of each DC-DC converter, due to the resistance of the wirebonds between EoS and the bus tape and indeed the resistance of the bus tape itself, and the module temperatures. In principle AMAC1A is able to monitor both of these parameters however for the duration of the Three Stave Tests the data acquisition system was not configured to log them. Hence further study is required to be able to eliminate the possibility that the increased noise of the modules when running only one module on each of the three stave sides may be caused by either of these effects. Certainly it has previously been shown [101] that the input noise of a strip module powered by a local DC-DC converter increases as the input voltage is raised, which would be consistent with these observations.

8.3 Position 2: Separation of supports 90 mm, 12°


For the second position the three staves were left at a tilt angle of 12° but the two “active” staves, the UK Double Sided stave and US Single Sided stave, were placed as closely together as could be safely achieved. The limiting factor is the proximity of the SFP+ cage on the Slave Side of the UK Double Sided stave and the wirebonds of the adjacent module of the US Single Sided Stave, as shown in Figure 8.14b. The spacing between the support centres of the UK Double Sided and US Single Sided staves was measured as 90 mm which results in a minimum sensor to sensor clearance of order 13 mm as shown in Figure 8.14a. Whilst this is greater than the 10 mm minimum sensor to sensor clearance expected in the experiment, this is the closest we may get in the three stave setup unless one of the staves is offset by 10 cm along the Z axis. (Note that the SFP+ cages are not used in the experiment: they will




(a) Module 9 of the Three Stave Sides



(b) Module 10 of the Three Stave Sides




(c) Module 11 of the Three Stave Sides




(d) Module 12 of the Three Stave Sides


Figure 8.12: The Gain measured running one module on each of the three stave sides together, divided by the reference Gain measured running each of the three sides separately. The procedure was repeated for the three modules 9 (Figure 8.12a), 10 (Figure 8.12b), 11 (Figure 8.12c) and 12 (Figure 8.12d), leaving all other modules unpowered. Each plot shows the result for the corresponding module of the Master side of the UK Double Sided Stave (chips 0-39), the Slave side of the UK Double Sided Stave (chips 40-79) and the US Single Sided Stave (chips 80-119). All results are from Response Curve tests taken under Optimised Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 512$) in Position 1. The gain remains constant.




(a) Module 9 of the Three Stave Sides



(b) Module 10 of the Three Stave Sides



(c) Module 11 of the Three Stave Sides



(d) Module 12 of the Three Stave Sides

Figure 8.13: The Input Noise measured running one module on each of the three stave sides together, divided by the reference Input Noise measured running the three sides separately. The procedure was repeated for the three modules 9 (Figure 8.12a), 10 (Figure 8.12b), 11 (Figure 8.12c) and 12 (Figure 8.12d), leaving all other modules unpowered. Each plot shows the result for the corresponding module of the Master side of the UK Double Sided Stave (chips 0-39), the Slave side of the UK Double Sided Stave (chips 40-79) and the US Single Sided Stave (chips 80-119). All results are from Response Curve tests taken under Optimised Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 512$) in Position 1. It can be seen that the noise increases when a smaller number of modules is operated.




Figure 8.14: Images of the Second Three Stave configuration. Figure 8.14a shows an end-on view of the two staves of interest, the UK Double Sided stave at the left and the US Single Sided stave at the right. The separation between the centres of the stave supports is 90 mm and the minimum distance between the two sensor planes is of order 13 mm. Figure 8.14b shows the point of minimum clearance between these two staves: the limit comes from the optical transceiver mounted on the lower face of the double sided stave, which will be replaced by a low profile VTRx+ transceiver in production staves. Clearance at this point is of order 1 mm.


be replaced by the custom VTRx+ optical transceivers which have a much lower profile.)

8.3.1 Comparison against 12°, Separation of support centres 237 mm


Figure 8.15 clearly shows that the Gain has not been affected by the reduction in the separation of the support centres from 237 mm to 90 mm. Similarly Figure 8.16 shows that the Input Noise is little changed with the possible exception of one hybrid on Module 10 on the US Single Sided Stave, for which bringing the staves closer together seems to have decreased the noise of some chips by 2.5%.

8.4 Position 3: Separation of supports 80 mm, 20°


For the third position the two “active” staves were pushed together such that the mounts of their supports were touching each other. In order to achieve this the tilt angle was increased to 20°. This arrangement is shown in Figure 8.17a. The sensor to sensor overlap is increased as shown in Figure 8.14b, however the minimum separation between opposing sensors is



(a) UK Double Sided Master Side




(b) UK Double Sided Slave Side




(c) US Single Sided Master Side


Figure 8.15: The Gain measured in Position 2, running the three stave sides together, divided by the reference Gain measured in Position 1, running the three sides separately. Both sets of results are Response Curve tests taken under Optimised Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 512$). The gain remains constant.



(a) UK Double Sided Master Side




(b) UK Double Sided Slave Side




(c) US Single Sided Master Side

Figure 8.16: The Input Noise measured in Position 2, running the three stave sides together, divided by the reference Input Noise measured in Position 1, running the three sides separately. Both sets of results are Response Curve tests taken under Optimised Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 512$). Some chips of Module 10 of the US Single Sided Stave show decreased noise by 2.5% (red box).



(a) End View




(b) Minimum Clearance

Figure 8.17: Images of the Third Three Stave configuration. Figure 8.17a shows an end-on view of the three staves, with the UK Double Sided stave at the centre in close proximity to the US Single Sided stave at the right. Their metal support structures are touching giving a separation between the centres of the stave supports of 80 mm. The minimum separation between the two sensor planes is of order 20 mm. Figure 8.17b shows the point of minimum clearance between these two staves, the limit coming from the optical transceiver mounted on the lower face of the double sided stave, which will be replaced by a low profile VTRx+ transceiver in production staves. Clearance is now of order 8 mm.


increased to 20 mm.

8.4.1 Comparison against 12°, Separation of support centres 237 mm


Figure 8.18 shows that the Gain measured in Position 3 remains consistent with earlier results. Figure 8.19 shows the same 2.5% decrease in Input Noise for some chips of Module 10 on the US Single Sided Stave as was recorded in Position 2, however there is one new effect: an increase of noise by 2.5% for Module 9 of the Master side of the UK Double Sided Stave.



(a) UK Double Sided Master Side




(b) UK Double Sided Slave Side




(c) US Single Sided Master Side


Figure 8.18: The Input Noise measured in Position 3, running the three stave sides together, divided by the reference Input Noise measured in Position 1, running the three sides separately. Both sets of results are Response Curve tests taken under Optimised Conditions ($n\text{Strobed} = 256$, $n\text{Trigs} = 512$). The gain remains constant.



(a) UK Double Sided Master Side



(b) UK Double Sided Slave Side



(c) US Single Sided Master Side

Figure 8.19: The Input Noise measured in Position 3, running the three stave sides together, divided by the reference Input Noise measured in Position 1, running the three sides separately. Both sets of results are Response Curve tests taken under Optimised Conditions ($n_{\text{Strobed}} = 256$, $n_{\text{Trigs}} = 512$). Some chips of Module 10 of the US Single Sided Stave show the same 2.5% noise decrease that was observed in Position 2 (red box), and Module 9 of the Master side of the UK Double Sided Stave shows a 2.5% noise increase (cyan box).

8.5 Comparison of Noise Occupancy results for selected modules

It was demonstrated in Section 6.2 that Noise Occupancy scans can be a sensitive tool to detect localised noise phenomena. Figures 8.20 and 8.21 present Noise Occupancy scans at the three stave positions for Module 12 of the Slave Side of the UK Double Sided Stave and the US Single Sided Stave respectively. Each scan was prefaced by an online pedestal trim as described in Section 3.3.9 in order to maximise the uniformity of thresholds at 0 fC equivalent threshold. As the Slave EoS of the UK Double Sided Stave directly faces the sensor of the latter module, this is one place where cross talk between staves may be expected to occur. Although the plots are not absolutely identical, in part because of the online nature of the pedestal trim, no significant effect is seen.

Figures 8.22 and 8.23 present similar Noise Occupancy plots for Module 11 of the Slave Side of the UK Double Sided Stave and the US Single Sided Stave respectively. Again no significant effect is seen.




Figure 8.20: Noise Occupancy for Module 12 of the Slave Side of the UK Double Sided Stave. Data is presented for Positions 1 (a-d), 2 (e-h) and 3 (i-l). The Z axis represents the fractional occupancy of each Channel Number (X axis) as a function of Threshold (Y axis) from a sequence of (up to) 64000 events. No significant difference is observed.




Figure 8.21: Noise Occupancy for Module 12 of the US Single Sided Stave. Data is presented for Positions 1 (a-d), 2 (e-h) and 3 (i-l). The Z axis represents the fractional occupancy of each Channel Number (X axis) as a function of Threshold (Y axis) from a sequence of (up to) 64000 events. No significant difference is observed.




Figure 8.22: Noise Occupancy for Module 11 of the Slave Side of the UK Double Sided Stave. Data is presented for Positions 1 (a-d), 2 (e-h) and 3 (i-l). The Z axis represents the fractional occupancy of each Channel Number (X axis) as a function of Threshold (Y axis) from a sequence of (up to) 64000 events. No significant difference is observed.




Figure 8.23: Noise Occupancy for Module 11 of the US Single Sided Stave. Data is presented for Positions 1 (a-d), 2 (e-h) and 3 (i-l). The Z axis represents the fractional occupancy of each Channel Number (X axis) as a function of Threshold (Y axis) from a sequence of (up to) 64000 events. No significant difference is observed.

8.6 Chapter Summary and Conclusions

It took much longer than expected to debug the data acquisition system to the point at which Three Stave Studies could begin, however with the exception of the UK Single Sided (electro-mechanical) stave stable operation was achieved. Three positions were studied and whilst no evidence of cross talk between opposing staves has been found, the closest approach between sensors in the experiment is almost 4 mm less than the 13 mm separation studied here. The noise level of several modules was higher throughout these tests than that achieved during previous, single stave tests, which may indicate continuing issues with the resistance of the electrically conductive glue used to make the backplane connections⁵.

For technical reasons, the scope of this thesis is restricted to studies performed in December 2019. Had more time been available, it would have been prudent to conduct further work as follows:

- The voltages delivered to the DC-DC converters of each module along a stave should be measured under normal load conditions, and the voltage delivered to a stave should be varied in order to measure the noise of a subset of modules across a range of input voltages. This is important to be able to discount the possibility that the increased noise of the modules closest to the EoS region may be related to their operation at a slightly higher input voltage than the modules further down the stave.
- Throughout the Three Stave study, several modules have shown increased noise compared to the corresponding single stave results. The sensor backplane bias connections of some of these modules had previously been treated by pushing large currents through the conductive epoxy in order to reduce its resistance, with consequent noise benefits. One hypothesis for the observed noise increases would be that the benefit of such treatment may not be permanent. Therefore it would be important to re-measure the backplane contact resistance of each of the affected modules, or to determine the

⁵As has been stated previously, it is not planned to use conductive glue to make the sensor backplane connections for production staves.

voltage at which full depletion is reached by recording module noise as a function of sensor bias, in order to prove (or disprove) this hypothesis.

- Although the earlier non-segmented stave was shown to have stable noise performance over a period of 60 hours as shown in Section 6.3.4, a similar study should be made using the UK Double Sided and US Single Sided staves. Some of the low level variations which occurred between previous test runs may easily be explained by training effects, if such effects were to be observed during long term stability tests.
- The closest approach between sensors which has been studied to date is 13 mm (Position 2), limited by the proximity of the SFP+ optical transceiver of the upper stave's EoS card to the hybrid of the opposing module as shown in Figure 8.14b. In the experiment a lower profile VTRx+ transceiver is used, which permits the closest approach between sensors to be reduced to 9.25 mm. It should be possible to get closer to this figure if the staves were translated relative to each other in Z, however the limitation would then come from the proximity of support structures rather than the staves themselves.

Furthermore the temperature range over which the performance of the Three Staves could be studied was limited both by the lack of humidity monitoring within the test enclosure and by the operating range of the cooling system. For future studies of staves using the Star Chipset, both of these points are being addressed. Nevertheless, the results presented in this chapter indicate that crosstalk between adjacent, correctly referenced staves should be of negligible concern in the experiment.

Chapter 9

Evaluation of the Star Chipset

This chapter documents the initial evaluation of ABCStar and HCCStar on single chip test boards and the wafer screening of the first 14 Star Chip MPW wafers as performed at RAL.

9.1 Test System Requirements for the Star Chipset

Although the requirements of a driver board to test ABCStar and HCCStar, the final prototype front-end and hybrid controller chips made for the ITk Strip detector, are similar to those of a driver board to test the earlier ABC130 and HCC130 chips, they are not the same. One of the most significant differences is a topological change in the way in which data is transmitted within the hybrid. Figure 9.1a shows the situation for a barrel hybrid using ABC130/HCC130. The ABC130 chips are formed into loops of five chips, the chip at each end of a loop is able to transmit its data to HCC130 directly, although normally the chips would be configured such that only the chip at one end of the loop has this function. All other chips must pass their data along the chain, in response to transitions of the XOFF flow control signal generated by their neighbour, such that the chip at the end of the chain is able to send the data to HCC130. In the case of ABCStar/HCCStar a radically different scheme is adopted as shown in Figure 9.1b. Now each ABCStar chip has a direct data path back to the HCCStar, which allows a reduction in the complexity of the digital logic within the ABCStar whilst reducing the probability that excess data generation by a single chip may result in data loss from other chips. However it does mean that additional point-to-




Figure 9.1: Data Transfer Architecture of the Barrel Hybrid: comparison of the “loop” architecture used by ABC130/HCC130 (left) with the “Star” architecture used by ABCStar/HCCStar (right). In both cases consecutive ABC addresses are chosen such that the four least significant bits of the chip address may be considered as a direct extension of the channel number, channels being numbered from left to right within the chip, however for ABC130 the addresses start at 16 as address 0 must be avoided for technical reasons.

point differential pairs must be incorporated in the hybrid layout, which a consequent small increase in mass and radiation length. In addition the ASIC test system must now be able to provide HCCStar with stimuli equivalent to 11¹ ABCStar chips, which the ABC130 driver is not able to do.

Within a hybrid, the HCCStar will provide the four “hybrid side” control signals required by the ABCStar chip. These are:

- BCO - Beam CrOssing clock at 40 MHz, the primary clock to which analogue sampling of the front end inputs is aligned.
- DRC (aka RCLK) - Data Recovery Clock at 160 MHz, used to clock data out of the chip.
- LCB - An acronym formed of the acronyms L0A (trigger), CMD (command) and BCR (Bunch Crossing Reset). A command signal at 160 Mbit/s which transmits these three key functions over four BCO clock periods by means of a 16 bit “symbol” which is decoded by the chip.
- LPPR - an 80 Mbit/s time multiplexed signal comprising the Low Priority request (LP) and Priority Request (PR) signals used to enact the transfer of data stored on-chip as

¹Although each barrel hybrid has 10 ABCStar chips, some of the endcap hybrids have the full complement of 11.

the result of a previous L0A trigger².

However the HCCStar itself has receives only three “stave side” control signals:

- BTC - Bus Tape Clock at 160 MHz, the clock used to receive and decode the command signals.
- LCB - An acronym formed of the acronyms L0A (trigger), CMD (command) and BCR (Bunch Crossing Reset), a command signal at 160 Mbit/s.
- R3L1 - a 160 Mbit/s stream in which the Regional Readout Request and Level 1 trigger signals are encoded. This signal is processed by HCCStar to form the LPPR signal needed by ABCStar.

Hence the StarChip test system must be able to generate two clocks at up to 160 MHz and two command signals at up to 160 Mbit/s. Furthermore whilst the data rate output by ABCStar remains 160 Mbit/s, HCCStar’s data output may operate at 320 or 640 Mbit/s: the ABC130 driver was not designed to operate at the higher rate.

A further requirement for a system to test the StarChips is the ability to program each die’s built-in “fuse register”, which gives us the ability to program a unique, persistent serial number into the chip. In order to do this one must iterate, writing a single bit which one wants to end up in the fused pattern into a control register, then pulsing a dedicated voltage bus (VFUSE) from the usual 1.2 V up to 3.3 V for a few hundred μ s in order to blow the corresponding “fuse”. Again, the ABC130 driver is not able to do this.

The conceptual design of the test system for the StarChips, ABCStar and HCCStar, is outlined in Figure 9.2. At the heart of the system is a commercial FPGA board, the Nexys Video[77], which runs the ITSDAQ firmware. Upon this is mounted a custom FPGA Mezzanine Card (FMC) which provides an interface between the FPGA and the chip(s) to be tested, including any custom or commercial ICs which may be needed to test all aspects

²The LP and PR signals correspond to the Level 1 (L1) and Regional Readout Request (R3) signals respectively, as outlined in Section s:subsection321B.




Figure 9.2: Outline of the Test System for Star Chips.

of the device(s) under test. Commercial high density twisted pair ribbon cables (SAMTEC FFTP series) are used to connect the FMC to either a Single Chip Card (SCB) fitted with a wirebonded StarChip, or a custom Probe Card which may be used to contact a StarChip on the wafer. Separate SCBs and Probe Cards were developed for both ABCStar and HCCStar. The design of the FMC, SCBs and Probe Cards are further discussed in the following section.

9.1.1 The FMC-1701 Driver Board

A specification was drawn up for a custom FPGA Mezzanine Card (FMC) to meet the requirements of a system to test both the ABCStar and HCCStar chips. The resulting board, known as “FMC-1701”[102], is shown in Figure 9.3. In order to fully test HCCStar, a total of 14 differential outputs are provided corresponding to the 160 MHz clock, two control signals and 11 simulated ABCStar data streams. Although not all are required to test the StarChips, a matching number of differential inputs is provided such that the differential IO of each FMC-1701 may be tested by addition of a loopback cable. The differential signals are distributed between headers J1 and J2 shown in Figure 9.3: the pinout is chosen such that the third ribbon cable of three is only required when testing HCCStar; in addition header J1 retains compatibility with the buffer made to test HCC130/ABC130 hybrids and modules meaning that partial tests of Star Chip functionality may be made using previous hardware designs. Differential outputs are driven by the SLVS ASIC [69], also used in the HCC130/ABC130 module test buffer, however on FMC-1701 the common mode and




Figure 9.3: FMC-1701, a driver board for ABCStar and HCCStar testing. The annotations give the function of each connector.

differential mode voltage levels of the SLVS chip are driven from commercial DACs such that they may be varied during testing to facilitate exploration of signal margins. The bandwidth of this chip is 320 Mbit/s. Differential inputs are buffered by a commercial chip with 800 Mbit/s bandwidth, such that HCCStar can transmit data at its maximum rate of 640 Mbit/s. In addition sufficient buffered ADC channels are provided to be able to measure the DC level of (one side of) each differential input.

As the SLVS buffers are wire bonded ASICs, a protective cover made from static dissipative acrylic is fitted to protect them from damage. Again with reference to Figure 9.3, header J3 provides power, monitoring and control capabilities to the device under test. Two analogue control voltages are provided which may vary between 0 V and 1.5 V, generated by a commercial dual-channel DAC on FMC-1701. These are used to drive the power load balancing shunt device integrated into the ABCStar and to provide a calibration stimulus to the ADC which forms part of each Star Chip design. A total of 15 single ended control

lines are provided. Nine of these, which are fully bidirectional, route to a commercial port expander with open collector drive capability. The remaining six lines pass through bidirectional level translating buffers³ directly to the FMC connector and hence to FPGA pins, meaning that the state of these lines can be registered or driven more quickly. This latter functionality had in mind the possible need to test the Joint Test Action Group (JTAG) Boundary Scan [103] functionality built into ABCStar: an industry standard interface used for the verification of PCBs and ASIC designs, this was included in the initial ABCStar prototype for debug purposes but will be dropped from the production version in order to further reduce power consumption and Single Event Effect (SEE) susceptibility.

Monitoring is provided in part by a 12 channel ADC which may be configured to operate as a six channel differential device. In the case of ABCStar the four most important monitoring nodes comprise measurements of the analogue and digital ground planes within the chip, the comparator threshold test point and the output of the ASIC's internal multiplexer which allows one to characterise each of the internal bias DACs. These four nodes are buffered by an operational amplifier chosen for its low offset error and high linearity before passing to adjacent ADC channels as matched signal/ground pairs. In this way the ADC may operate in its differential mode such that any ground movement of the chip does not contribute to measurement errors on these critical nodes. In the case of HCCStar three of these buffered ADC inputs are re-purposed to be able to monitor the DC point of three AC-coupled SLVS control signals, as maintained by HCCStar, after the coupling capacitors on the Single Chip test Board (SCB) or probe card. In all these cases the high input impedance of the buffer amplifier ensures minimal disturbance to the voltages being output by the chip. Additional ADC channels measure the power supply voltages delivered to the SCB or probe card; the output of each chip's internal Low Drop Out (LDO) voltage regulators; a Negative Temperature Coefficient (NTC) thermistor on the SCB or probe card and, in the case of a probe card, the resistance of a wafer contact probe which may be incorporated into the needle ring.


In terms of powering the device under test, two options are supported. With reference to

³One of these buffers is permanently wired as a driver, however the direction of the remaining five buffers may be varied under software control.


Figure 9.3 one may connect a pair of external 1.5 V power supplies to connector J6, with sense wires available at J10 if required. This option is recommended for wafer testing as commercial supplies generally have programmable current limits, which helps ensure that the current drawn by any faulty device remains within safe limits. Alternatively a 1.5 V Low Drop Out (LDO) regulator is provided on the FMC-1701 main board which may be connected to the device under test by means of jumper bank J5. Each voltage rail is monitored by an INA230 current and power monitor IC [104], which also measures the voltage drop across a 50 m Ω shunt resistor in order to determine the current with 50 μ A precision. The FMC-1701 itself takes most of its power directly from the FMC header, however an external -1.5 V supply at header J11 is required for optimum performance of the various buffer amps which form part of its design. Additional features include access to the FMC-1701's internal I2C bus at header J7, which enables the connection of external temperature and humidity sensors such as may be required for low temperature ASIC tests, and two general purpose LEMO sockets J8 and J9 which may be used to provide monitor outputs for connection to an oscilloscope or as external trigger inputs. Connector J4 provides an additional option for the HCCStar's 640 Mbit/s data output (and diagnostic clock signal) to enter FMC-1701: screened, low capacitance mini DisplayPort cables maintain better signal integrity than ribbon cables of similar lengths which can give a better indication of driver performance and hence Bit Error Ratios (BER).

9.1.2 StarChip Single Chip Boards

Initial studies of the ABCStar and HCCStar designs were made using wire bonded die on custom Single Chip Boards (SCBs). These are shown in Figure 9.4. Connected to FMC-1701 by means of two twisted pair ribbon cables, the ABCStar SCB[105] shown in Figure 9.4a allows the complete functionality of the ASIC to be tested. The front-end chip may be operated on its own or loaded by up to two “baby” strip sensors. Other than this the circuit is entirely passive. It is also possible for two of the front-end channels to be loaded by capacitors of known values: this is an important feature which allows one to determine



(a) ABCStar Single Chip Board



(b) HCCStar Single Chip Board


Figure 9.4: Single Chip Boards (SCB) for the Star Chips: ABCStar (top) and HCCStar (bottom). These were used widely during the initial verification of each ASIC design and shall be used throughout production to facilitate the X-ray irradiation of sample ASICs for purposes of Quality Control.

the noise of the amplifier across a range of capacitive loads representative of the sensors to be used in the experiment.


The HCCStar SCB is shown in Figure 9.4b. Again, apart from the ASIC, the board is entirely passive. It supports a variety of possible connection options which are selected by soldered jumpers, making it a little more complicated than the ABCStar board. The most usual configuration routes all HCCStar’s inputs and outputs to FMC-1701 by means of three twisted pair ribbon cables, however secondary options allow the board to be wire bonded to a stave bus tape to facilitate tests of the quality of signal transmission along the tape, or for an ABCStar SCB to be connected to the HCCStar SCB in order to simulate a basic hybrid. Although the HCCStar SCB is bigger than the ABCStar SCB, the mounting holes have the same spacing to ensure compatibility with pre-existing cooling fixtures used to control ASIC temperatures during irradiation tests.

9.1.3 StarChip Probe Cards

Probe card designs for the two chips were also prepared as shown in Figure 9.5. Each is designed to be as similar as possible to the corresponding SCB circuit, although a probe contact switch is added in each case. This comprises a pair of tungsten probes which are normally in contact with one another but come apart when one of the probes contacts the wafer. This is further explained in section 4.1. In reality there will be some variation in the resistance between the two probes before they become fully open, hence the FMC-1701 driver measures the resistance of the contact switch. Features omitted from the ABCStar probe card[106] include the ability to load the front-end with capacitors or sensors, and the HCCStar probe card omits the ability to connect an ABCStar SCB directly to the chip under test.



(a) ABCStar Probe Card



(b) HCCStar Probe Card

Figure 9.5: Probe cards for the Star Chips: ABCStar (top) and HCCStar (bottom).

9.2 ABCStar Single Chip Board Noise Results




Figure 9.6: ABCStar SCB Noise Measurements. The red and green data show input noise as a function of discrete load capacitance before and after irradiation of the chip to 80 Mrad. The diamonds refer to noise measurements made using miniature n^+ on p strip sensors. Also shown, in cyan, are measurements of noise as a function of capacitance made of the ABCStar front-end prototype circuit. (Negative charge injection is used throughout, consistent with the signals developed in n^+ on p sensors.)

Once the readout system had been debugged and the core functionality of the ABCStar had been demonstrated, noise studies were performed using ABCStar SCBs to verify that, in the final strip module application, tracking performance will not be impaired by excess occupancy due to noise in the front-end system.

In the first noise study, the load capacitors of one ABCStar SCB were varied across the range 0 pF to 5.6 pF, sufficient to cover the expected range of sensor loads. The design allows load capacitors to be fitted to each of two channels: for each combination of channel and capacitive load a Three Point Gain was recorded, as described in Section 3.3.3, in order to measure the noise. The average input noise of the two channels is plotted as the red line in Figure 9.6. The green line in Figure 9.6 repeats this procedure on a second ABCStar

SCB, the chip having previously been irradiated to the lifetime dose of 80 Mrad⁴ using the X-ray source at RAL. This demonstrates that the excessive noise increase shown in irradiated ABC130 die, due to leakage in the Shallow Trench Isolation (STI) layer, has been largely eliminated in the ABCStar design by the use of Enclosed Layout input transistors. However a smaller effect remains due to increased leakage currents within the gate oxide itself.

The second noise study used additional ABCStar SCBs fitted with varying combinations of miniature strip sensors of varying lengths: “mini” (1 cm), “SS_mini” (2.5 cm) and “LS_mini” (5 cm). Each of these sensors were made on the same wafer as one of the full size ATLAS ITk strip sensor designs and therefore had similar n^+ on p processing. These were connected in differing combinations to produce SCBs fitted with strip sensors equivalent to four capacitance points: 1.1 pF (one mini), 2.7 pF (one SS_mini), 3.8 pF (one mini plus one SS_mini) and 5.6 pF (one LS_mini). In each case, 100 channels were bonded to the sensors. The average input noise for the bonded channels for each of these SCBs, again resulting from a Three Point Gain test, is shown by a Diamond plotted in Figure 9.6. Although the mini sensor noise values are higher than those determined for single channels with equivalent capacitor loads, this is in part because the additional load capacitance to the next nearest neighbour strips has not been taken into account. In addition there is a small noise contribution in the sensor case due to the series resistance of the metal readout strips as outlined in Section 2.1.4. Hence the agreement is satisfactory and the noise values are considered to meet the required specification.

9.3 ABCStar Wafer Probing

9.3.1 Development of the ABCStar Test Sequence

Before wafer tests could start, a test sequence was established built from functional blocks written and debugged using wire bonded chips on Single Chip Boards.

⁴As stated in Section 2.1.3 the expected worst case lifetime dose is actually 53.3 Mrad. The 80 Mrad end point used during TID studies includes a safety factor of 1.5.

Config	RESETB	PWR_LOW	Clocks	LCB PRLP	LDO_LOW	ASIC Registers
0	low	high	off	off	clear	N/A
1	low	high	on	on	clear	N/A
2	low	low	on	on	clear	N/A
3	high	low	off	off	clear	Initial
4	high	low	on	off	clear	Initial
5	high	low	on	on	clear	Initial
6	high	low	on	on	clear	Default
7	high	low	on	on	set	Tuned
8	high	low	on	on	clear	Tuned

Table 9.1: The configuration of control lines and registers during ABCStar Power Tests. Configurations 0-6 are tested at the beginning of the test sequence whilst configurations 7 and 8 are recorded later, once the LDO and DAC tunings have been made.

The first part of the test sequence measures the currents drawn by the chip under a series of pre-determined configurations as shown in Table 9.1. The power dissipated by ABCStar is expected to vary in response to a number of factors including register configurations; activation of POWER_LOW mode which disables many parts of the chip, a feature which is necessary in order to be able to briefly perform a limited test of an uncooled stave upon its insertion into the barrel structure; lowering the digital LDO output voltage by asserting LDO_LOW; holding the logic core in the reset state by asserting the RESETB pin; disabling the BCO and DRC clock signals; disabling the LCB and PRLP command and trigger signals. First the chip is held in reset by pulling the RESETB node high and low power mode by pulling the POWER_LOW node high whilst the effects of disabling then enabling the clock and control signals are studied, before low power mode is released. The reset signal is then released and the effect of disabling then separately enabling the clock and control signals is determined. Only configurations 0 to 6 are tested at this point in the sequence: the final one of these tests sets the chip's registers to operational defaults. For each of these configurations the recorded analogue and digital currents and voltages are compared against expectation online and any device for which a measurement is outside the permitted range will fail the test. During wafer probing any unusually low reading is likely to be due to a failure of contact: in such cases the needles are automatically re-seated and a second attempt is made.

fuseID	Meaning	Comment
fuseID[23:22]	WaferSeries[1:0]	prototype = 0, pre-production = 1, production = 3 up to 2047 wafers in each series set true if AMAC2 up to 2047 die per wafer
fuseID[21:11]	WaferID[10:0]	
fuseID[10]	isAMAC?	
fuseID[9:0]	DieID[9:0]	

Table 9.2: The allocation of bits within the 24 bit Fuse Register of an ATLAS ITk “Star Chip”.

On a hybrid, each ABCStar will be assigned a permanent, static chipID (address) by means of wirebonding. It is however important to verify during wafer tests that each chip will function at any possible address. This is done by the second part of the test sequence which loops over all possible addresses, setting the address on the four chipID control inputs before writing a number to the chip’s latency register using the newly set address. This number is then read back. If the correct number is received, the chip is judged to be able to operate at that address. The test sequence is discontinued for any chip which fails to respond at one or more of the possible sixteen addresses.

At this point in the sequence we have established bi-directional communication with the chip and it is drawing the expected current. The next step writes the desired serial number of the chip to its fuse register: this is done as early as possible in the test sequence - as soon as we know we can talk to it - to ensure full traceability of as many die as possible.

All three StarChip designs - ABCStar, HCCStar and AMAC2 - feature the same 24 bit fuse register, however the scheme used to assign a unique identifier to each device is complex and must take into account a number of factors. The subdivision of these 24 bits is outlined in Table 9.2. Firstly the scheme takes into account the need for three generations of these designs: prototype, pre-production and production, so two bits are reserved for this. Secondly, whilst the present prototype wafers are a Multi-Project (MPW) design comprising 378 ABCStar die, 96 HCCStar die and 96 AMAC2 die, pre-production and production ASICs will be processed as two wafer designs, one for HCCStar and AMAC2 and one for ABCStar alone. This is the cheapest way to obtain the chips in production quantities, as the numbers of die of each type can be better tuned to meet our needs. The combined

Parameter	Description	Monitor Voltage	Tolerance
LDOA	Analogue Regulator	1.20 V	50 mV
LDOD	Digital Regulator	1.20 V	50 mV
VREF	Internal Bias Voltage	80 mV	2 mV
IREF	Internal Bias Current	86 mV	2 mV
R8B	Bias for 8-bit DACs	80 mV	2 mV
VCOM	Discriminator Bias	35 mV	2 mV

Table 9.3: Tuning Targets for the two adjustable voltage regulators and four bias DACs of the ABCStar chip.

HCCStar/AMAC2 wafer will contain approximately 1050 die with an unequal split between HCCStar and AMAC2, hence 11 bits of the serial number are allocated to record the die number within the wafer with the top bit set only for AMAC2 die. This leaves 11 bits to record the identity of each wafer which, as the total number of production wafers across both designs is estimated as 750, gives a factor of two safety margin.

Each wafer bears a seven character alphanumeric code, etched upon it by the vendor. Upon delivery, these codes are added to a master Look Up Table (LUT) to define a unique integer corresponding to each wafer “name”. These integers are then used to give the wafer identity, WaferID, in the serial number scheme as described above. The identity of the chip within the wafer, chipID, corresponds to the index of that die within the wafer map. It is only necessary to fuse those bits of the register which are to be set at logic “1”. In order to minimise the fusing current, these bits are written into the control register one at a time and the fuse voltage VFUSE is pulled from 1.2 V up to 3.3 V in order to blow the fuse. This is repeated until all ones have been written, at which point the fuse register is read back. Any die for which the expected serial number is not read back is deemed to have failed this test and the test sequence is suspended.

The next part of the test sequence aims to characterise the DACs and ADCs built into ABCStar. As part of this process, and in order to compensate for process variations, four of the critical bias DACs are tuned to maximise the performance of each chip. The list of parameters to be tuned, and the targets and required tolerances, are shown in Table 9.3. Any chip which may not be tuned in accordance with these values will be rejected at the

end of the test or, in the case of LDOA and LDOD, immediately.

Firstly the chip’s internal analogue and digital Low Drop Out (LDO) regulators, controlled by 16 bit unary DACs, are scanned across the entire range and set to deliver the nominal operating voltage of 1.2 V. Next the software cycles through all possible outputs of the chip’s internal multiplexer, which monitor a number of bias nodes and DACs including three of the remaining parameters to be tuned. An additional node used to monitor the bias of the comparator block, VCOM, is available elsewhere on the chip pad ring, the last of the parameters to be tuned. The final DAC characterised by this test adjusts the current of the ABCStar’s output SLVS driver. In all except the last case, the signals are available to the ABCStar’s internal ADC as well as our external ones, so measurements are made by both methods.

Once the DACs have been fully characterised, the next step aims to calibrate the function of the ABCStar’s internal ADC. For four of the possible ADC gain settings, a known voltage in the range 0 V to 1.2 V is provided to the ADC calibration input. Similarly the ABCStar’s internal shunt device is characterised by sweeping its (external) control voltage between 0 V and 1.2 V and recording IDDD, the current drawn on the digital power rail. Finally for this stage of the sequence, the current drawn on the analogue power rail, IDDA, is studied as a function of one of the front-end bias parameters, IREF. Measurements of the analogue and digital ground of the chip are also made at the start and end of the test. The temperature of an NTC thermistor located on the probe card and the resistance of a probe contact switch, as explained in Section 9.1.3, are also recorded.

In addition to the checks upon the LDO output voltages mentioned earlier, four additional cuts upon ADC measurements are applied online. Two of these ensure that the chip’s analogue and digital ground potentials have moved less than 5 mV during the test. The remainder check that the “knees” of the analogue and digital LDO characteristics, as introduced in Figure 4.5, occur no later than bin 10, to ensure that the analogue and digital voltages delivered to the chip exceed 1.3 V. These tests are designed to detect problems which may occur due to poor control of probe contact resistance: in such cases the test

sequence will be aborted⁵.

At this point in the sequence the bias parameters of the chip have been successfully optimised. Additional power measurements are now made, shown as configurations 7 and 8 in Table 9.1, in order to record the currents which the tuned chip may be expected to draw in the experiment when clocked and correctly configured, but otherwise idle. Also a quick check is made of the function of additional bit of the digital LDO control DAC, LDOD_LOW, which aims to drop the output voltage of the LDO by a nominal 200 mV, as may be required in order to limit the digital power of the circuit during the phenomenon of the “TID peak” which was introduced in Section 2.1.3. Online cuts are applied to these measurements and the test sequence will be stopped for any die which do not conform to expectation.

Next a series of digital tests are performed. Known regular patterns are written to the chip’s mask register in order to establish a background pattern of channels with occupancy in all time bins. The facility to pulse the chip’s input register is then used to generate a subset of time bins which have occupancy in all channels. Each of the resulting two dimensional patterns is compared against expectation: should any of the tests yield the incorrect result, the test is repeated before marking a chip as having failed the test, in which case the sequence is stopped.

The final portion of the test sequence aims to verify the performance of the ABCStar’s 256 front-end inputs. Firstly a Pedestal Trim is performed as outlined in Section 3.3.9. This scans the channel trim DACs in the absence of charge injection, in order to measure the amplifier offset (pedestal) of each channel. This data is then analysed and the trim values adjusted to bring all channel offsets into line. Die with more than two untrimmable channels are immediately rejected and the test sequence aborted. Next a the strobe delay setting is scanned and tuned in order to optimise the charge injection timing as outlined in Section 3.3.2: die for which the optimised strobe delay setting does not fall within expected, generous limits are again immediately rejected. Finally a Three Point Gain (3PG) test as described in Section 3.3.3 is executed to determine the gain and noise of each channel.

⁵Such chips typically pass when retested after adjustment of the needle contact heights.

It should be noted that additional, more rigorous cuts are applied as part of an offline analysis which is used to generate final wafer pick maps to be used by the dicing vendor. This is further explained in Section 9.3.2.

The initial batch of 12 StarChip MPW wafers (“Batch 1”) was delivered in October 2018. One wafer was sent immediately to a dicing vendor and a second wafer was sent untested to Carleton University⁶, Ottawa, Canada. This left 10 StarChip MPW wafers to be tested at RAL. Testing began in February 2019, once the basic performance of the chip had been successfully verified using Single Chip Boards. The 12” SemiProbe PS4L Probe Station [107] at RAL, having been subjected to remedial noise treatment as described in Appendix B, was used for these tests.

An additional four StarChip MPW wafers (“Batch 2”), available for purchase as a result of over-production by the vendor, were delivered to RAL in May 2019. Minimal changes were made to the software to address specific deficiencies which had been noted during offline analysis of data from the original 10 wafers.

The first of these changes was to eliminate the spurious values which were occasionally obtained when reading the ABCStar’s internal ADC. By design, ABCStar sends uncommanded High Priority Register (HPR) data in response to a variety of error conditions, or may be programmed to do so periodically. Running tests with a bonded chip, it was determined that the bad ADC readings were due to the misinterpretation of these packets. The code was updated to identify and reject HPR packets and this has resolved the problem.

Secondly during the original run of ten wafers it was noted that the ASCII log file for the internal MUX data was filled with zero content, meaning that for the first 10 wafers this data, where possible, had to be recovered from the ROOT output. This was traced to a coding error and duly corrected. As a result of this fix the measurement of the bandgap voltage using the internal ADC is available for the second batch of wafers: for the first batch only a measurement using the external ADC was recorded. Finally some changes were made to

⁶In the production phase Carleton University and their industrial partner, DA Integrated, will be responsible for the testing half of all ABCStar die. RAL will be responsible for testing the other half.

the ADC/DAC characterisation scans to speed up the test sequence by only writing register values which change during each scan, rather than sending the complete set each time.

9.3.2 Offline Analysis

An offline analysis is used to determine the final yield of each wafer and to produce wafer pick maps for transmission to the dicing vendor. The methodology is similar to that used with the ABC130 chip as documented in Appendix A. The first stage of this procedure, “Compilation”, pulls together data from a variety of ROOT and ASCII files generated during the tests to produce a single, summary ROOT file for each wafer. This file contains a number of TTree and histogram objects as shown in Table 9.4. During this procedure some of the data is reprocessed to address limitations of the fits made online. The first items to be replaced are the fits to the LDO characteristics. In many cases the online fit fails to converge, as shown in Figure 9.7. The raw data is recovered from the stored ROOT files and processed to determine the “knee” of the graph and the slope of the linear region. To aid further study, the raw LDO data and the values determined from the new analysis are stored to a dedicated TTree object, `ldoTree`.

The second and third of these issues relate to measurements made using the ABCStar’s internal ADC and are shown in Figure 9.8. Figure 9.8a shows a typical result in which a known voltage is applied to the ADC input in order to calibrate it. This is repeated for four of the possible 16 gain settings, 0x4, 0x8, 0xC and 0xF. In the figure presented the ADC reading saturates on the highest range, an effect which had not been anticipated as it did not occur on the first bonded chips. This phenomenon was discussed with the analogue designer, who confirmed that the ADC was designed to have a linear response up to 1 V. To accommodate the effect, the offline code initially fits a straight line to the data for voltage stimuli up to 0.90 V. A second pass includes data points beyond this as long as they deviate from the previously fitted line by less than 20 counts: the straight line fit is then repeated for this expanded data set, from which the saturation region has been excluded. Figure 9.8b illustrates another effect whereby attempts to read the internal ADC may occasionally

Object Name	Type	Content	Scope
pwrTree	TTree	Voltages and Currents from Power Tests	one wafer
emuxTree	TTree	DAC characterisation using ext. ADC	one wafer
imuxTree	TTree	DAC characterisation using int. ADC	one wafer
dacTree	TTree	DAC tuning information	one wafer
tvTree	TTree	Results of Digital Test Vectors	one wafer
gnTree	TTree	Summaries of Analogue Front End Tests	one wafer
ldoTree	TTree	Reprocessed LDO tuning data	one wafer
TRIM_XXX	TH1F	Trim data for die XXX	one die
GAIN_XXX	TH1F	Gain vs channel for die XXX	one die
OFFSET_XXX	TH1F	Offset vs channel for die XXX	one die
INNSE_XXX	TH1F	Input Noise vs channel for die XXX	one die

Table 9.4: The content of the ROOT Summary File resulting from testing an ABCStar Wafer.




Figure 9.7: Typical LDO tuning characteristics for ABCStar, implemented as 16 bit unary DACs. In the right hand example the online fit has failed, but this is addressed offline. All that matters online is that the desired 1.2 V output may be achieved: this is illustrated in each figure by the asterisk with the achieved voltage being written on the plot.




Figure 9.8: Two issues with the ABCStar’s ADC data were addressed offline. The left hand plot illustrates saturation of the ADC on its highest gain range, which the online fit algorithm did not anticipate. The right hand plot shows that the software occasionally records a “bad” ADC reading due to packet loss. Both effects were accommodated by re-analysis of the data offline; in addition the second of these effects was fixed in the version of the online software used for the second batch of four wafers.

return a bad reading. The online code made no attempt to exclude these from the data before making a linear fit, hence the offline code must identify and exclude such points before repeating each fit.

The second stage of this procedure, “Analysis”, processes the content of the summary objects against a series of carefully chosen cuts in order to separate die into three streams: categories A (good), B (minor defect) and X (reject). The cuts applied to the power test measurements are shown in Table 9.5 and mirror exactly the cuts used in the online software. The cuts upon current measurements are loose, being based upon early measurements made using bonded chips but having been adjusted based upon experience with the first wafer. Cuts upon voltage measurements are similarly loose, with the exception of checks of the regulated voltages once the chip’s LDO regulators have been tuned. The cuts are sufficient to detect that each chip under test has responded correctly to the test stimuli.

The cuts applied to the characterisation of ABCStar’s internal DACs are shown in Table 9.6: the first four entries in the table represent 4 sigma cuts on the distribution of the relevant parameter across a sample of the first 10 wafers. The cuts on the knees of the LDO graphs are there to reject chips for which insufficient voltage was delivered to the chip due

Parameter	Configuration	Lower Limit	Upper Limit	Comment
IDDD	0,1	0 mA	12 mA	Digital Current
IDDD	2,4,5,6	20 mA	40 mA	
IDDD	3	10 mA	40 mA	
IDDD	7	15 mA	25 mA	
IDDD	8	20 mA	40 mA	
IDDA	0,1	4 mA	12 mA	Analogue Current
IDDA	2,3,4,5	30 mA	50 mA	
IDDD	6,7,8	60 mA	80 mA	
VDDD	all	1.40 V	1.60 V	V _{digital} on FMC-1701
VDDA	all	1.40 V	1.60 V	V _{analog} on FMC-1701
VDRAW	all	1.40 V	1.60 V	V _{digital} on Probe Card
VARAW	all	1.40 V	1.60 V	V _{analog} on Probe Card
VDREG	0-5	1.40 V	1.50 V	Regulated V _{digital}
VDREG	6	1.10 V	1.35 V	” ”
VDREG	7	0.90 V	1.10 V	” ”
VDREG	8	1.15 V	1.25 V	” ”
VAREG	0-1	0.20 V	0.70 V	Regulated V _{analog}
VAREG	2-5	1.34 V	1.50 V	” ”
VAREG	6	1.10 V	1.35 V	” ”
VAREG	7	1.15 V	1.25 V	” ”
TESTRES	all	900 Ω	1100 Ω	Test Resistor

Table 9.5: Cuts applied to the “Power Test” measurements. Any chip falling outside the cuts is rejected (category “X”). Note that for some parameters, different cuts are applied to each tested configuration: the configurations are outlined in Table 9.1.

Parameter	Comment	Lower Limit	Upper Limit
TRDAC slope	Slope of Trim DAC	2.50e-3 mV/count	3.74e-3 mV/count
Vbandgap	Bandgap Voltage	571 mV	635 mV
VCAL slope	Slope of FE Calibration DAC	3.19e-4 mV/count	3.45e-4 mV/count
VTHR slope	Slope of Threshold DAC	4.95e-4 mV/count	5.40e-4 mV/count
LDOA Knee	Knee of LDOA Characteristic	N/A	10 counts
LDOD Knee	Knee of LDOD Characteristic	N/A	10 counts
TESTRES	Integrated Test Resistor	900 Ω	1100 Ω

Table 9.6: Cuts applied to the characterisation of ABCStar’s internal DACs. Any chip falling outside the cuts is rejected (category “X”).

Parameter	Lower Limit	Upper Limit	Category
ADC Slope, range 0x4	9000 counts/V	15000 counts/V	X
ADC Slope, range 0x8	4246 counts/V	6245 counts/V	X
ADC Slope, range 0xC	3395 counts/V	4699 counts/V	X
ADC Slope, range 0xF	2929 counts/V	4041 counts/V	X
Maximum count, range 0x4	3000 counts	N/A	X
Maximum count, range 0x8	3000 counts	N/A	X
Maximum count, range 0xC	3000 counts	N/A	X
Saturation voltage, range 0xC	0.8 V	N/A	X

Table 9.7: Cuts applied to the characterisation of ABCStar’s internal ADC.

to contact resistance. The TESTRES cut is redundant, as the same test and cut are made during the power tests as previously shown in Table 9.5.


Whilst no cuts are applied to measurements of the ABCStar’s internal voltages and monitor nodes as measured using its internal ADC, cuts are applied to the calibration of the ADC as shown in Table 9.6. Recall that the ABCStar’s integrated ADC has 16 possible gain settings, four of which are tested on the wafer: 0x4, 0x8, 0xC and 0xF (4, 8, 12 and 15)⁷. The most important cut in terms of operation of an ABCStar chip in ATLAS, is that on the nominal gain setting of 0xC (12) the ADC should not exhibit saturation until beyond 0.8 V, which is sufficient to cover normal use. In addition it is verified that the maximum count on all tested ranges exceeds 3000, again this is to eliminate chips with ADCs subject to early saturation. The cuts applied to the recorded ADC slopes are motivated to eliminate extreme outliers only.

The cuts applied to the ABCStar’s tuned DAC parameters are shown in Table 9.8. Die which may not be tuned to the correct operational targets will have already been rejected online, so the intent here is to exclude chips which fall well into the tail of the distributions. The distributions of the tuned values of two parameters VREF_DAC and IR8B_DAC, as determined across all fourteen wafers, are shown in Figure 9.9. The mean value of the VREF_DAC distribution is 12.21, which is in close agreement with the mean value of 13

⁷For production ABCStar wafers, this procedure will be revised. It is planned to make a quick assessment of all 16 ADC gain settings to select the one which best covers the desired 0.9 V operating range, before making a detailed calibration for that setting only.


Parameter	Lower Limit	Upper Limit	Category	Comment
VREF_DAC	10	14	X	Internal Bias Voltages
IREF_DAC	12	17	X	Internal Bias Currents
R8B_DAC	10	14	X	Bias for 8 bit DACs
VCOM_DAC	8	13	X	Bias for Discriminator
LDOA_DAC	9	14	B	Analogue Regulator
LDOD_DAC	9	14	B	Digital Regulator
LDOA_Output	1.18 V	1.22 V	X	Analogue Voltage
LDOD_Output	1.18 V	1.22 V	X	Digital Voltage

Table 9.8: Cuts applied to the tuning of ABCStar’s internal DACs.



(a) Distribution of VREF_DAC settings (b) Distribution of IR8B_DAC settings

Figure 9.9: The distributions of the optimised DAC settings for two parameters, VREF_DAC and IR8B_DAC, across all 14 wafers. Negative entries correspond to cases where the desired output value could not be achieved.



(a) Distribution of LDOA_DAC settings (b) Distribution of LDOD_DAC settings

Figure 9.10: The distributions of the optimised DAC settings for two parameters, LDOA_DAC and LDOD_DAC, across all 14 wafers. Negative entries correspond to cases where the desired output value could not be achieved.

Parameter	Lower Limit	Upper Limit	Category
UnTrimmable Channels	-	0	B
UnTrimmable Channels	-	2	X
Tuned Strobe Delay	15	30	X
Mean Gain	70 mV/fC	89 mV/fC	X
Mean Input Noise	360 e-	460 e-	X
Bad Channels	-	0	B
Bad Channels	-	2	X
Chisq(Gain)	-	27.5	X
Gain Slope	-0.023	0.023	X

Table 9.9: Cuts applied to the measurements of ABCStar’s Front End Performance.

predicted by simulation. The value predicted by simulation for IR8B_DAC is also 13 and the mean value of the measured distribution, which is 11.85, is similarly close. For these parameters, the tuned values for all chips fall within the applied cuts. The distributions of the tuned values of two LDO DACs, LDOA_DAC and LDOD_DAC, are shown in Figure 9.10. The mean values of the two distributions are 12.41 and 12.24 respectively. In both cases the cuts reject a few chips in the tail for which the optimised setting is 15: whilst this does not leave sufficient headroom to accommodate bandgap variations which may be expected to occur during operation of the ITk detector due to radiation effects these chips remain perfectly functional, hence they are categorised as having a minor defect (Category “B”). In addition a tighter offline tolerance of 20 mV is applied to the tuned outputs of the analogue and digital regulators than the 50 mV used online.

Cuts applied to measurements of the ABCStar’s front-end performance are shown in Table 9.9. Referring back to the test sequence, the first stage of the analogue test sub-sequence is a Pedestal Trim as described in 3.3.9 which aims to align the amplifier offsets (pedestals) for all channels by adjustment of the trim DACs. Whilst several trim ranges (gain of the trim DAC) are available, a mid range setting is used for the test. Die having more than two untrimmable channels - those which may not be brought into line with the majority - are rejected outright (category “X”) and die having one untrimmable channel are marked as having minor defects (category “B”). The next stage of the sub-sequence performs a Strobe Delay Test as described in 3.3.2, which optimises the timing of the charge




Figure 9.11: The distribution of the Mean Channel Gain and the variation of Mean Channel Gain as a function of Die Number. Data for all 14 wafers are shown.




Figure 9.12: The distribution of the Mean Channel Input Noise (INNSE) and the variation of Mean Channel Input Noise as a function of Die Number. Data for all 14 wafers are shown.

injection pulse. Chips for which the optimised strobe delay value does not fall into the range 15 to 30 counts will be rejected.

The final part of the analogue test sub-sequence performs a Three Point Gain Test as outlined in Section 3.3.3 in order to determine the gain and noise of the chip. The measured distribution of the Mean Channel Gain is shown in Figure 9.11 together with a plot showing the variation in the Mean Channel Gain as a function of Die Number. On the basis of this distribution the cuts on the Mean Channel Gain were set at 70 mV/fC to 89 mV/fC, a three sigma cut. It is also apparent that the measured gain varies across the wafer, being at its lowest toward the edges and at its highest near the wafer centre. Similarly the distribution of the Mean Channel Input Noise (INNSE) is shown in Figure 9.12 together with a plot showing the variation in this parameter as a function of Die Number. As a result of this (non-Gaussian) distribution the cuts on the Mean Channel Input Noise were set at




Figure 9.13: The plot of gain vs channel for Wafer VPCVRRH Chip 046 and the accompanying distribution for which χ^2 is 60. A pronounced variation in gain is seen across the chip: such chips must be rejected.




Figure 9.14: The plot of gain vs channel for Wafer VXCQ1H Chip 005 and the accompanying distribution for which χ^2 is 29. A region of abnormally high gain is observed: such chips must be rejected.

360 electrons (-2 sigma approx) to 460 electrons (+4 sigma approx). The measured input noise is in general more constant across each wafer than is the case for gain. The motivation of these cuts is the identification of gross outliers, and any such chips are rejected.

Figures 9.13 and 9.14 show cases in which chips exhibit localised variations in channel gain, yet do not fail any of the existing selection criteria based upon bad channel counts, mean gain or mean noise. In order to reject chips such as these, an additional cut was developed based upon the χ^2 of a Gaussian fit to the distribution of Channel Gain values within each chip. As the test case shown in Figure 9.14 has a χ^2 of 29, any cut made on the χ^2 of the Gain Distribution must be set below this. The distribution of such χ^2 values across a sample of all 14 wafers is shown in Figure 9.15. By manual inspection of additional chips highlighted as potentially bad as the cut level was reduced, it was decided to reject any




Figure 9.15: The distribution of χ^2 of a Gaussian fit to each chip's channel gain values, across a sample of 14 wafers.




Figure 9.16: The plot of gain vs channel for Wafer VICVPYH Chip 231, showing a number of randomly distributed low gain channels, and the accompanying distribution for which χ^2 is 27.5.

chip for which χ^2 of the distribution of channel gain values is greater than 27.5. Cutting at this level was sufficient to identify and reject the two cases shown in Figures 9.13 and 9.14 and others with obvious non-uniformity of gain across the chip, without picking up any false positives.

A subsequent study looked in more detail at every chip failing the cut at 27.5 on the χ^2 of the distribution of the channel gain values. A total of 14 chips failed the cut. Eight of these fall into the category illustrated by Figure 9.16, in which the distribution of the channel gain values has a non-Gaussian tail which may be attributed to number of randomly distributed low gain channels. The remaining six chips exhibit a pronounced variation (or slope) in the gain measured across the chip as shown in Figure 9.17. The offline analysis was expanded to calculate the “Gain Slope” for each chip: the resultant distribution and a plot of the Gain Slope as a function of die number are shown in Figure 9.18. It can be seen in Figure




Figure 9.17: The plot of gain vs channel for Wafer VICVPYH Chip 174, showing a pronounced slope across the chip, and the accompanying distribution for which χ^2 is 60.0.




Figure 9.18: “Gain Slope” is, for each ABCStar die, the slope of a straight line fitted to a plot of gain vs channel. Here the distribution of this parameter is shown together with a plot showing how the parameter varies against die number (and hence a chip’s location on the wafer). The sample includes all 14 tested wafers.

9.18b that the gain slope varies as a function of location on the wafer, with a tendency for its magnitude to be greater at the wafer periphery than at the centre. A cut was placed on Gain Slope at ± 0.023 , with chips falling outside this limit being rejected (Category “X”). It is important to reject die with excessive gain slopes as this clearly indicates a problem with the chip.

9.3.3 Results and Discussion

The yield across the 10 tested wafers of “Batch 1” is shown in Table 9.10. It can be seen that a total of 3453 good die were found, corresponding to a category “A” yield of 91.3%. The yield across this set of 4 wafers of “Batch 2” is shown in Table 9.11. It can be seen that a further 1384 good die were found, corresponding to a category “A” yield of 91.5 %.

In total 369 ABCStar die were rejected as a result of these tests, 261 from the first batch of ten wafers and 108 from the second batch of four wafers. Table 9.12 shows at which point in the test sequence each of the rejected die was deemed to have failed. It can be seen that the 33.1% of rejected die, 122 die in total, fail the initial power tests. The second most common reason for a chip to be rejected is due to the offline cuts on the analogue test results, accounting for 27.4% of rejected die or 101 die in total. The reasons for these rejections are further explored in Table 9.13. This is in addition to the 30 chips which failed online cuts on the number of trimmable channels and / or the strobe delay setting, for which test sequence was not completed.

Another useful way to look at the results across the 14 wafers is to identify the locations at which faulty die are frequently found. The resulting map of the number of die rejected at each location is shown in Figure 9.19. It can be seen that there are seven locations for which a yield of 50% or less was found. The reasons for these rejections are further explored in Table 9.14. Locations 1, 5, 46 and 128 are towards the edge of the 200 mm diameter wafer and clip the “known good” diameter of 194 mm provided by the vendor, so a poor yield at these locations is not unexpected. Note however that there are four additional sites which clip the known good radius, locations 286, 358, 359 and 378, and these have a good yield. This may

Wafer	nA die	nB die	nRejects	A Yield	A+B Yield
VTCVRMH	334	8	36	88.4%	90.5%
VPCVRRH	327	8	43	86.5%	88.6%
V0CVRFH	353	9	16	93.4%	95.8%
V2CVRDH	343	4	31	90.7%	91.8%
V7CVR8H	342	8	28	90.5%	92.6%
V1CVREH	359	4	15	95.0%	96.0%
VKCV PWH	346	11	21	91.5%	94.4%
VXCVQ1H	352	5	21	93.1%	94.4%
VUCVRLH	347	7	24	91.8%	93.7%
VSCVRNH	350	4	24	92.6%	93.7%
Total	3453	68	259	91.3%	93.1%

Table 9.10: ABCStar Probing Yield, Batch 1

Wafer	nA die	nB die	nRejects	A Yield	A+B Yield
VMCVPUH	352	8	18	93.1%	95.2%
VXCVRIH	342	3	33	90.5%	91.3%
VHCVPZH	344	5	29	91.0%	92.3%
VICVPYH	346	5	27	91.5%	92.9%
Total	1384	21	107	91.5%	92.9%

Table 9.11: ABCStar Probing Yield, Batch 2

Sequence	Reason Rejected	Number Rejected	Fraction of Rejected Die
1	Initial Power Tests Failed	122	33.3%
2	ID Test Failed	41	11.2%
3	fuseID Not Written	0	0.0%
4	DAC Cuts Failed	49	13.4%
5	Internal ADC Fails Cuts	4	1.1 %
6	DAC Tunings Fail Cuts	5	1.4 %
7	LDO Tunings Fail Cuts	0	0.0%
8	Additional Power Tests Failed	3	0.8%
9	Digital Tests Failed	9	2.5%
10	Online Analogue Cuts Failed	30	8.2%
11	Offline Analogue Cuts Failed	103	28.1%

Table 9.12: ABCStar Defect Statistics. For the 366 rejected (category “X”) die, the point in the test sequence at which the die was first identified as bad is tabulated. There is only one entry for each rejected chip.

Cut	Number of Die Rejected
Mean Channel Gain < 70 mV/fC	30
Mean Channel Gain > 89 mV/fC	0
Mean Channel Input Noise < 360 e-	0
Mean Channel Input Noise > 460 e-	39
More than 2 Bad Channels	18
RMS of Channel Gain distribution > 3.0	0
χ^2 of Channel Gain distribution > 27.5	16
Gain Slope < -0.0023	12
Gain Slope > +0.0023	3

Table 9.13: ABCStar Analogue Defect Statistics: for the 96 rejected (category “X”) die for which the analogue test sequence was completed, only four of the offline cuts were triggered. These were low Gain (30 chips); high Input Noise (39 chips); more than 2 bad channels (18 chips); excessive χ^2 of the Channel Gain distribution (16 chips); chips with excessive Gain Slopes (12+3 chips). There may be more than one entry for each chip.

Site	1	5	46	128	185
Initial Power Tests Failed	5	2	1	10	4
ID Test Failed	9	2	0	4	0
DAC Cuts Failed	0	0	3	0	0
Digital Tests Failed	0	0	1	0	0
Online Analogue Cuts Failed	0	2	0	0	0
Offline Analogue Cuts Failed	0	3	2	0	7
Number of Bad Die at location	14	9	7	14	11

Table 9.14: ABCStar Low Yield Sites. Across a population of 14 wafers, this table shows the five sites at which a yield of 50% or less is observed, and explores the reasons why these die were rejected.




Figure 9.19: Map showing the location of all rejected (category “X”) ABCStar die across a population of 14 wafers. Each cell is labelled with the corresponding die number. Cells 1, 5, 46, 128 and 185 stand out as being the sites with the highest numbers of rejected chips.

suggest that some parts of the ABCStar circuit are more easily affected when they clip the known good boundary than others. The fifth location at which poor yield is experienced is site 185, right at the centre of the wafer: although similar observations have been reported for other designs in this process, the reasons for this have not been understood.

9.4 ABCStar Single Chip Board Supplemental Results

After wafer probing of ABCStar had been completed, several additional tests were performed using Single Chip Boards to provide input to the Final Design Review (FDR) of the ITk Strip ASICs. Four samples were available for testing, all from the initial diced wafer VJCVPXH: chips 004 and 010 (unirradiated) and chips 8 and 9 (irradiated with x-rays to 80 Mrad at RAL).




Figure 9.20: (a) Gain, (b) Output Noise and (c) Input Noise as a function of Strobe Delay for chip VJCVPXH.010. The red lines denote the chosen optimum setting of 22 corresponding to the maximum Gain and minimum Input Noise.

9.4.1 Optimisation of Measurement Conditions

Additional gain and noise measurements of the ABCStar front-end, over a range of bias conditions, were requested for the review. As a precursor to this we reviewed the basic conditions used during a Three Point Gain Test to ensure that no significant systematic error was being introduced by this test procedure. In all cases the chip under test is first tuned to optimum bias conditions.

It remains important to optimise the Strobe Delay setting, which governs the timing of the internal charge injection pulse, prior to any Three Point Gain Test. The speed of the delay gates inside ABCStar is known to vary as a function of temperature, hence it is usual to repeat the Strobe Delay Test at the start of any testing session or when the operating temperature is changed. Figure 9.20 shows how Gain, Output Noise and Input Noise vary as a function of Strobe Delay, Input Noise being calculated by dividing the Output Noise by the Gain and scaling by the electronic charge. During this scan the data compression logic within ABCStar is set to return hits occurring in any one of the three available time bins, hence the plots show a recurring pattern. The edges of the central 25 ns time bin




Figure 9.21: (a) Gain and (b) Input Noise as a function of nStrobed, the number of channels receiving the charge injection pulse at once. Test sample is chip VJCVPXH.004. Error bars represent the standard deviation of the measurement across the channels of the chip.

correspond to the Gain minima occurring at Strobe Delay values of 10 and 30. Hence it can be seen that, for the operating conditions at the time of the test, one Strobe Delay count was approximately equal to 1.25 ns. The Strobe Delay Test as described in Section 3.3.2 looks at the central time bin on its own, and sets the delay to a point 57% of the way between the edges of that bin. For the chip shown above, VJCVPXH.010, the delay chosen by the Strobe Delay Test is 22, illustrated in Figure 9.20 by the vertical red lines. It can be seen that this value gives the maximum Gain and minimum Input Noise: in effect this plot illustrates why the Strobe Delay fraction was chosen to be 0.57.

The ABCStar ASIC, in common with the ABC130 design, allows one to inject charge pulses into any number of its channels at once, from 1 to 256. For ABC130 modules the difference in the noise recorded by a Three Point Gain Test is almost 10% higher when charge is injected into all 256 channels simultaneously than the figure recorded when charge is injected into only 16 channels at once, as shown in Section 6.3.3. ABCStar exhibits the opposite effect as shown in Figure 9.21, with the noise actually falling slightly as increasing numbers of channels receive the charge injection pulse, however the magnitude of this effect is negligible (less than 1%). It would be interesting to repeat this test on a greater number of samples, for example the first ABCStar stave, in order to confirm or deny the statistical significance of this observation.

It is also possible to vary the number of events recorded at each point of each scan of the Three Point Gain Test, parameter nTrigs. The results of this study are shown in Figure





Figure 9.22: (a) Gain and (b) Input Noise as a function of nTrigs, the number of events recorded at each point of each threshold scan. Test sample is chip VJCVPXH_004. Error bars represent the standard deviation of the measurement across the channels of the chip.

9.22. The number of events taken at each point, nTrigs, is varied from 128 to 8192 in powers of 2. When the number of events recorded is small, 128, the measured Gain, Input Noise and Output Noise (not shown) are all underestimated, however for numbers of events per point equal to or greater than 1024 events the gain and noise remain constant.


Whilst at times during the production process it remains necessary to record a smaller number of events - for example to maximise throughput during wafer screening - one should ideally record 1024 events at each point for the most accurate results.

9.4.2 Operation across a range of bias conditions


Measurements of the Gain and Noise of the ABCStar front-end were made across a range of supply voltages and bias conditions in order to verify that the chip would operate stably across the full range of process variations one may expect to encounter during production. The analogue supply voltage VDDA was varied by adjustment of the LDOA regulator, selecting the closest possible voltage to each of three targets: 1.1 V, 1.2 V and 1.3 V. At each of the selected voltages the gain was measured for a matrix of 9 bias points to cover BVREF from 9 to 17, equivalent to the variation of voltage bias in the range of 15% to 25%, and to cover BIREF from 8 to 22, equivalent to the variation of the input transistor bias current from 100 to 155 μA . At each point of the current matrix the chip is first trimmed - a pedestal trim - before the threshold DAC is re-characterised using the ADC on FMC-1701. Next the Strobe Delay test is performed to compensate for any variation in response times




(a) Gain at VDDA = 1.127 V




(b) Noise at VDDA = 1.127 V




(c) Gain at VDDA = 1.187 V



(d) Noise at VDDA = 1.187 V



(e) Gain at VDDA = 1.306 V



(f) Noise at VDDA = 1.306 V

Figure 9.23: Gain and Noise as a function of bias parameters BVREF and BIREF, for supply voltage VDDA in the range 1.127 to 1.306 V. Test sample is chip VJCVPXH_009, which has been irradiated to 80 Mrad.

which may occur under different bias conditions before finally a Three Point Gain Test is made using the measured threshold calibration to ensure that the calculated gain and noise will be as accurate as possible. For the Three Point Gain Test, the number of Strobed Channels is 256 and the number of events taken at each point is 1024 in accordance with the data presented in Section 9.4.1. Four samples were studied: data for chip VJCVPXH_009 is presented in Figure 9.23 however the results for all chips were similar. The front-end was found to perform well across the full range that was studied, in agreement with simulation.

9.4.3 A more detailed characterisation of ABCStar’s internal ADC

Although not apparent from the initial specification document, the four register bits available to adjust the bias current of the ABCStar’s internal ADC were not given equal weight. The most significant bit (MSB) was intended to select a “high gain” or “low gain” option targeting a full scale reading of 250 mV or 1.0 V respectively with the remaining three bits being used to accommodate process variations within the chip, notably variation of a (fixed) bandgap voltage which regulates the ADC’s reference current.




Figure 9.24: Characterisation of ADC on all ranges, chip VJCVPXH_010.

Figure 9.24 shows the response of the ADC of chip VJCVPXH_010 for each possible bias setting (left) together with the calculated Integral Non-Linearity (INL, right). In the left hand plot each point represents the arithmetic mean of 10 ADC readings with the RMS of




Figure 9.25: Characterisation of ADC on “best” range, chip VJCVPXH_010.

the readings shown as the error bar. For each range, a straight line is fitted to all points below 1.0 V such that the Integral Non Linearity (INL) may be calculated. It can be seen that measurements taken on the lowest ranges have very high error bars in the left hand plot and large INL values in the right hand plot, for example the black circles and red squares. This effect was subsequently reproduced in simulations by the analogue design team: for such low currents the ADC is not properly biased. As a result of this discovery the next version of the chip the “high gain” option will be dropped and the ADC bias adjustment will be refocused to provide 16 options to tune the chip to a full scale reading of 1.0 V across a wider range of process variations, with good linearity across the range 0 to 0.9 V.

Under such circumstances it seems fair to concentrate on the “low gain” region, to choose the bias current which most closely matches the full scale reading to 1.0 V, and to study the INL using a fit made to the expected linear region of 0 to 0.9 V. The results of such a study for chip VJCVPXH_010 are shown in Figure 9.25. It can be seen that between the limits of 0 and 0.9 V the maximum deviation from the fitted line is 2 counts, or 0.5 mV. In this region the ADC performs in accordance with specification and will meet our requirements in ATLAS.


9.4.4 Shmoo Plots

It is common industrial practice for the operation of digital circuits to be tested across a wide range of operating conditions including bias voltage, clock frequency, clock phase and temperature, to ensure that operation of the circuit is in no way marginal when operated under the conditions for which it was designed. Execution of such tests may also be used to help identify any weak spots in the design, which may be improved in subsequent versions of the product. The results of tests such as these are typically presented in the form of a two dimensional histogram showing the performance of the system as a function of two variables, known as a “Shmoo Plot” [108] in honour of a similarly-shaped cartoon character.


With the exception of the work reported in this section, all tests of ABCStar documented in this thesis were performed using fixed frequency Bunch Crossing (BCO) and Data Recovery (DRC) clocks of 40 and 160 MHz respectively, which represents the normal operating condition of the chip in ATLAS. These clocks are internally generated in the Nexys Video’s FPGA. It is also possible to supply the FPGA with an external clock signal, in which case one can run at any (BCO) frequency up to the 52 MHz limit supported by early firmware builds or the 60 MHz limit supported by current ones. In all cases the DRC clock is generated by multiplication of the input BCO clock within the FPGA and hence the required constant phase relationship is maintained.

The source of the external clock signal is a PCB based upon the FS7140 Programmable Phase Locked Loop (PLL) Clock Generator IC [109], which may be programmed from the Nexys Video by means of its I2C interface. The vendor provides a Windows application which may be used to calculate the register values to be written to the chip in order to obtain the desired frequency output: this was used to derive a Look Up Table (LUT) to cover frequencies from 20 MHz to 65 MHz in 1 MHz steps. In order to operate the chip at voltages below the limit of the internal LDO regulator, the regulator is set to provide the minimum drop and the required voltage is fed directly to the chip.


Initial tests to explore writing and reading ABCStar register content across a range of




(a) VJCVPH.004 (0 Mrad) Shmoo at 20 °C



(b) VJCVPH.004 (0 Mrad) Shmoo at -40 °C



(c) VJCVPH.009 (80 Mrad) Shmoo at 20 °C



(d) VJCVPH.009 (80 Mrad) Shmoo at -40 °C

Figure 9.26: Shmoo plots for irradiated and unirradiated ABCStar chips at differing temperatures, showing how performance varies as a function of frequency and digital supply voltage VDDD. The metric used is (one plus) the number of channels for which pedestal trim data fails basic integrity checks, dark blue (one) being the desired result. Test sample VJCVPH.009 has been irradiated to 80 Mrad.

voltages and frequencies found that this part of the chip was fully functional from 20 to 52 MHz (limited by the test system) right down to VDDD = 0.8 V. In order to identify any part of the chip which may be failing, the next step was to run the full wafer test sequence for one of the SCB test chips across a range of VDDD settings, at a fixed BCO frequency of 40 MHz. It was found that during digital pattern tests and threshold scans data loss began to occur at around 1.0 V. The software did not report any decoding errors, hence the effect represented a real loss of data within the chip.

The pedestal trim scan was then turned into a standalone test vector to allow exploration of this effect by counting the number of channels for which the returned data failed basic

integrity checks, such as not reaching 100% occupancy, or by not maintaining 100% occupancy through to the end of the scan once that level had been reached. The “Shmoo” plot shown in Figure 9.26a is for unirradiated chip VJCVPXH_004 operated in an environmental chamber at 20 ° Celsius. It shows that at the usual digital operating voltage of 1.2 V the chip operates correctly up to the 52 MHz limit imposed by the test system, and at the usual operating frequency of 40 MHz the chip operates down to 1.05 V. Reducing the temperature of the environmental chamber to -40 °C, margins are further increased as shown in Figure 9.26b. “Shmoo” plots at the same temperatures were also recorded for irradiated chip VJCVPXH_009 as shown in Figures 9.26c and 9.26d: margins are slightly degraded with respect to the unirradiated chip, but still good.

In all cases the apparent recovery for certain frequencies whilst running at 0.8 V is real: it is likely that for some part of the circuit the phase relationship between clocks is now such that we have crossed over to the next clock period, and operation is recovered. This suggests that the performance may be recoverable at other voltage/frequency combinations subject to precise tuning of the relative phase of the clocks and command signals: this was attempted but the delay steps available from the firmware, at 1/4 of the phase of the output DRC signal, were too coarse for this to be of any real benefit. However the results are very encouraging, indicating that we should have 200 mV margin at the intended operating point of 40 MHz and 1.2 V at -25 °C.

9.4.5 Calibration of ABCStar’s internal Temperature Diode

The ABCStar incorporates a diode structure intended to serve as a temperature monitor. The voltage across the diode is accessible to the chip’s integrated ADC, but is not available elsewhere. Hence in order to measure the voltage vs temperature characteristic of the temperature diode it was necessary not only to vary the temperature of the chip across the expected operating range, but also to recalibrate the internal ADC against a known external voltage source (DAC) at each temperature studied.

A study was performed for the usual four samples on Single Chip Boards (SCBs), two

Chip	TID	Diode Slope	Diode Offset
VJCVPXH_004	0	-0.002069	0.6813
VJCVPXH_010	0	-0.002076	0.6827
VJCVPXH_008	80	-0.002189	0.6414
VJCVPXH_009	80	-0.002192	0.6422
	Mrad	mV/°C	mV

Table 9.15: Calibration of ABCStar’s integrated Temperature Diode. Results are presented for four die, two irradiated and two unirradiated, calibrated against an NTC thermistor across the temperature range -40 °C to +20 °C. Whilst there is a subtle difference between the irradiated and unirradiated chips, within each group the results are similar.

irradiated chips and two unirradiated chips, in an environmental chamber across the temperature range -40 °C to +20 °C. The SCB design places a Negative Temperature Coefficient (NTC) Thermistor adjacent to the chip and this was used to provide a more accurate temperature reference than that reported by the environmental chamber. Results are shown in Table 9.15: it can be seen that whilst the temperature sensitivity of irradiated and unirradiated samples differ slightly, within each category consistent performance is observed.

9.5 HCCStar Wafer Probing

9.5.1 Development of the HCCStar Test Sequence

The probe test sequence for HCCStar comprises a number of individual tests each focussed upon a subset of the chip’s functionality. The first stage, the Configuration Test, configures the chip and tunes the internal bandgap reference such that the LDO delivers 1.2 V. If this may not be achieved, the chip is assumed dead and the system moves to the next chip. The second stage, the Power Test, makes measurements of analogue parameters under a number of different conditions such as low power mode, chip clocked and chip unlocked. The parameters measured include the current drawn by the chip; the common mode voltages of the HCCStar’s SLVS outputs; the common mode voltages of the (AC coupled) control inputs to the HCCStar, R3L1 (trigger), BTC (clock) and LCB (command). Cuts are applied online and bad chips are immediately rejected.

The third stage, the Communication Test, confirms that the HCCStar’s configuration registers have been written correctly, and that the chip only responds to commands which are sent to the correct address. The fourth stage of the sequence writes a unique identifier, the FuseID, to the fuse register in accordance with the rules given in Table 9.3.

The fifth stage, the Digital Test, sends simulated ABCStar data into each of ABCStar’s 11 data inputs, individually and in several combinations, and verifies that HCCStar passes the data through correctly. In addition a number of possible error conditions are forced to ensure that the HCCStar will respond correctly. The sixth stage, the Memory Test, completely fills HCCStar’s L0 (trigger number) SRAM and event FIFO buffers with known patterns and verifies that the data is read back correctly.

The seventh stage, the Hybrid Passthrough Test, looks at the signals which HCCStar will send to the ABCStar chips on the hybrid and verifies that they are correct under a variety of conditions. For instance, configuration commands should only be passed to ABCStar if the incoming command is labelled with the correct HCCStar address, or marked as a broadcast command for the attention of all hybrids. The eighth and final stage, the Analogue Test, repeats the measurements made as part of the initial Configuration Test before the HCCStar’s internal ADC, used for monitoring purposes, is calibrated. This uses a different method from that used with ABCStar. The ADC gain is measured for each and every one of eight possible gain settings and the software selects the setting which most closely matches unity gain (1000 counts = 1000 mV).

Chips with ADCs for which unity gain may not be achieved to within 5%, assuming they have passed all the other tests, are placed into category “B”⁸. Any chip which fails any other criteria, or any other test, is rejected. Only chips which pass all the tests are placed into category “A”. Any die which is rejected was retested to reduce the probability that an erroneous result may be recorded due to probe contact failures.

⁸During screening of the prototype HCCStar chips it was not considered necessary to add an additional cut to reject completely die for which the ADC ramp gain was grossly incorrect, as in any event “B” grade chips shall not be used on strip modules.

Wafer	nA die	nB die	nRejects	A Yield	A+B Yield
V1CVREH	85	0	11	88.5%	88.5%
VKCVPPWH	91	1	4	94.8%	95.8%
V0CVRFH	88	1	7	91.7%	92.7%
V2CVRDH	91	2	3	94.8%	96.9%
V7CVR8H	92	1	3	95.8%	96.9%
VPCVRRH	88	2	6	91.7%	93.8%
VSCVRNH	90	1	6	93.8%	94.8%
VTCVRMH	91	0	5	94.8%	94.8%
VUCVRLH	93	1	2	96.9%	97.9%
VXCVQ1H	87	3	6	90.6%	93.8%
Total	896	12	53	93.3%	94.6%

Table 9.16: HCCStar Probing Yield, Batch 1

9.5.2 Results and Discussion

The ten Multi Project Wafers (MPWs) of batch 1 not only had their ABCStar die tested at RAL, but their HCCStar die also. HCCStar die were tested on the older 8” MicroManipulator probe station so as not to interfere with ongoing tests of ABCStar die on the newer 12” SemiProbe system. The yield across this set of 10 wafers is shown in Table 9.16. It can be seen that 896 good die were found, corresponding to a category “A” yield of 94.6%.

For each wafer, the number of die which failed the test sequence and the point in the sequence at which the die was rejected is shown in Table 9.17. It can be seen that the two most common reasons why a chip is rejected are failure of the Digital Test (22 die) followed by failure of the initial Configuration Test (18 die). It is possible that some of the chips identified as bad by the digital tests, which verify the reception of ABCStar data by HCCStar, may be false positives due to contact resistance: the ABCStar data connections are clustered at one end of the chip which isn’t really used for anything else, hence this is the only test which requires good contact on all four sides of the die. All other rejections are believed to be correct. Note also that failure of the HCCStar Analogue test does not result in complete rejection of the die, rather it is categorised as having a minor defect (category “B”).

A map showing the number of HCCStar die rejected at each location is shown in Figure

Wafer	Con	Pwr	Com	Fuse	Dig	Mem	Pas	Ana
V1CVREH	2	1	0	0	8	0	0	0
VKCVPTH	3	0	0	0	1	0	0	1
V0CVRFH	1	0	0	0	4	2	0	1
V2CVRDH	1	2	0	0	0	0	0	2
V7CVR8H	1	0	1	0	1	0	0	1
VPCVRRH	2	1	1	0	1	1	0	2
VSCVRNH	2	0	1	0	2	0	0	1
VTCVRMH	3	0	0	0	1	1	0	0
VUCVRLH	1	0	0	0	1	0	0	1
VXCVQ1H	2	1	0	0	3	0	0	3
Total	18	5	3	0	22	4	0	12

Table 9.17: HCCStar Defect Statistics. For each chip which did not complete the test sequence successfully, the point at which the die was first identified as bad is tabulated. There is only one entry for each chip. Note that die which failed the Analogue Tests are not categorised as rejects but rather as having minor defects (category “B”).

Site	4	5	6	50
Configuration Test Failed	0	0	8	0
Power Test Failed	0	0	0	3
CommunicationTest Failed	0	0	0	0
FuseID not written Correctly	0	0	0	0
Digital Test Failed	2	8	0	0
Memory Test Failed	0	0	0	0
Hybrid Passthrough Test Failed	0	0	0	0
Analogue Test Failed	1	1	0	0
Number of Bad Die at location	3	9	10	3

Table 9.18: HCCStar Low Yield Sites. Across a population of 10 wafers, this table shows the three sites at which a yield of 70% or less is observed, and explores the reasons why these die were rejected.




Figure 9.27: Map showing the location of all rejected HCCStar die across a population of 10 wafers. Each cell is labelled with the corresponding die number. Cells 4, 5, 6 and 50 stand out as being the sites with the highest numbers of rejected chips.

9.27 and the reasons for these rejections are further explored in Table 9.18. It can be seen that there are four locations at which the yield is 70% or less. At site 6 the yield is 0%: every chip fails the bandgap test. After the wafers had been diced it was discovered that this site had been added to the wafer map by mistake: not only does the chip cross the 194 mm “known good” radius but it physically crosses the 200 mm wafer boundary. During HCCStar wafer testing some difficulty was encountered getting good contact with die in the top right hand corner of the wafer, so the lower yield of sites 4 and 5 may be down in part to contact resistance. The final site at which poor yield was experienced is site 50 at which the yield is 70%. This is at the middle of the wafer, adjacent to ABCStar site 185 at which a

poor yield was also experienced. As the two ASIC designs were tested using different probe cards and on different probe stations, this seems to support the hypothesis that the central region of the wafer may be subject to processing difficulties.

9.6 Chapter Summary and Conclusions

A custom driver board was made in the form of an FPGA Mezzanine Card (FMC) in order to test the ABCStar and HCCStar chips. Initial investigations including software and firmware debugging were performed using wire bonded die on custom Single Chip Boards (SCBs). The input noise of the ABCStar front-end was measured as a function of load capacitance and found to be in good agreement with the performance of the front-end prototype, both before and after irradiation. Later tests focused on the exploration of margins within the ABCStar design, both Analogue and Digital, confirmed the design to be robust.

Wafer screening was performed for both ABCStar and HCCStar. The yield for ABCStar across a sample of 14 wafers was found to be 91.4% and the yield for HCCStar, across a sample of 10 wafers, was found to be 93.3%. Hence despite the different areas of the two chips, with ABCStar being approximately 4 times larger than HCCStar, the yield of the two designs is similar.

Work remains to be completed before these screening processes may be used in production. Rate estimates for the production phase are based around running one ABCStar wafer of 466 chips each working day, yet the 378 ABCStar die of each prototype MPW wafer were typically tested over two days. The main reason for this low throughput is mechanical: the wafer contact lift frequently needed adjustment during the tests. In principle the need for this can be reduced by further adjustment of the probe station hardware to improve chuck and needle planarity, and the need for any residual adjustment may be automated. Once this issue is resolved, a test time of 180 s per chip would be sufficient to test one production wafer of 466 ABCStar chips in each 24 hour period. However, in order to completely test one wafer and to start a second one within a typical eight hour shift, such as may be required to

get back on schedule after unforeseen interruptions, it would be necessary to bring the chip test time below 60 s.

Subsequent to the tests reported here, two key optimisations have brought the time taken to test a single ABCStar chip down from 150 to 90 seconds. The first change was to increase the clock frequency used to read the commercial I2C devices on the FMC-1701 driver board from 10 kHz to 500 kHz, and the second was to suppress the online graphical display. Further reductions should be possible by reduction of the amount of network traffic between the host and the FPGA, notably by re-coding the digital tests such that the test vector sequences are stored in the Nexys Video's on-board memory for execution and analysis by the FPGA directly. If the 60 s target test time may not be achieved by optimisation of the software and firmware alone, the next step would be to reduce the number of points recorded for certain DAC characterisation measurements. It would not be advisable to drop any part of the sequence completely as this would leave the required functionality partially untested.

Chapter 10

Conclusions and Ongoing Work

This thesis has presented results from two generations of prototype front end ASICs designed for the ATLAS ITk Strip detector together with the test methodologies used to obtain them.

Although the initial version of the ABC130 chip could not be read out due to a design error, wafers made after this had been corrected were found to have a high yield. Several months were spent testing early modules built with this chipset, trying to obtain sufficient agreement with the designer’s expectations: once a maladjustment of the front end’s bias settings had been corrected, and the significant dependence of the measured gain upon the number of channels simultaneously receiving charge injection stimulus was worked around by strobing a small subset of each chip’s 256 channels, this was finally achieved in November 2016. The ABC130 front end design was however found to exhibit a significant noise increase in response to ionising radiation, which would need to be corrected in a future design.

A partially loaded, double sided stave built using electrical readout and non-segmented bus tapes was made available for testing in May 2017. The HV-MUX circuitry, initially found to be a source of noise, was successfully modified on-stave. Thereafter module performance on this stave was in agreement with expectation. A further three ABC130 short strip staves were built with non-segmented bus tapes and GBTx (optical) readout, although the first of these staves used mostly “Electro-mechanical” modules fitted with dummy sensors. In general terms the noise of the modules measured on-stave is a factor of 1.1 higher than when

measured off-stave, with additional noise increases in the modules nearest the End of Stave (EoS) card.

It took much longer than expected to debug the data acquisition system to the point at which Three Staves could be read simultaneously, however no evidence of cross talk between opposing staves has been found at any of the three studied positions. The noise level of several modules was however higher than that achieved during previous, single stave tests, which requires further investigation. Operation of the Three Stave Test system was halted in January 2020 by failure of the dry gas systems, meaning that the sensors may not be fully biased, which limits the work reported in this thesis. It would be prudent to conduct additional studies once the dry gas supply is restored, in particular once a cooling system capable of taking staves to lower temperatures becomes available, which could provide useful information about the performance of ABC130 staves at operating temperatures approaching those expected in ATLAS (coolant at $-35\text{ }^{\circ}\text{C}$).

The second generation prototype chipset, HCCStar / ABCStar, was also found to have a high yield. The revised front end amplifier design in the ABCStar delivers low noise values across the range of expected sensor loads and no significant increase in noise is observed when the chip is irradiated. Unlike the earlier ABC130 chip, gain measurements made using the internal charge injection circuitry are almost independent of the number of strobed channels, so accurate gain measurements can be quickly recorded. The ABCStar front end design fully meets our experimental requirements.

Although beyond the scope of this thesis, the main limitation of the HCCStar / ABCStar chipset is their susceptibility to Single Event Effects, which was found to be unacceptably high. Although all critical registers were triplicated in the first version of these chips the control logic was not, making the operation of these chips susceptible to radiation-induced errors. The digital part of all three designs (HCCStar, ABCStar and AMAC) is now being redesigned to protect all critical functions by means of fully triplicated, spatially separated control logic running from three separate clock trees. At the time of submission the revised ABCStar was being manufactured with delivery due in May 2020. The front end ampli-

fier design is unchanged, although there are changes in some of the peripheral analogue blocks: the tunable ranges of the linear regulator blocks have been re-centred in response to measurements shown in this thesis, and the shunt regulator block has been revisited to remove an unnecessary transistor with thick gate oxide due to its inherent susceptibility to increased leakage current in response to Total Ionising Dose. Hence the analogue results currently being obtained from modules and staves built with the initial HCCStar / ABCStar chipset, including a partially loaded Short Strip staffe built at the Rutherford Appleton Laboratory (RAL) assembled in August 2019 and a fully loaded Long Strip staffe completed at Brookhaven National Laboratory (BNL) in February 2020, should give a very good indication of the analogue performance of the final detector.

The work documented in this thesis underpins many aspects of the ATLAS ITk Strip programme and has been instrumental to the success of several areas of activity. Single chip test cards have contributed directly to the basic understanding of the front end chipset and wafer probing activities have provided known good parts to the hybrid and module construction programmes. In turn the module construction programme has built modules for more detailed studies, such as testbeam activities, and of course for construction of the staves. Getting each staffe to work has been an important vindication of the electrical concepts behind the staffe, and getting the three staves to work together has indicated that there is no measurable crosstalk between staves. Many of these results have been presented in internal collaboration meetings, and referenced repeatedly as part of ATLAS Upgrade review processes. This work has brought the ATLAS ITk closer to reality: the completed Strip detector will help deliver the tracking performance needed for a new decade of physics at the energy frontier.

Appendix A

Cuts for Automated Analysis of ABC130 Wafer Probing Results

The online software stores results in a mixture of wafer and chip specific text files and chip specific ROOT [110] files. It was clear that the practice adopted for earlier wafers, which was largely manual, was not scalable to production quantities. A series of additional offline tools were developed such that the analysis of wafer data may largely be automated. The first tool builds one ROOT “reference” file per wafer comprising: four TTree objects summarising the results of power tests, DAC characterisations, digital tests and analogue (gain) tests respectively; three TH1 objects per chip storing the channel by channel gain, noise and offset (vt50) numbers. The second tool applies a series of cuts to each reference file in turn, generating a wafer pickmap in ASCII form. The third tool turns the ASCII pickmaps into graphical form, and the fourth tool enables a variety of summary plots to be made for each wafer or across a batch.

The ABC130 has separate power domains for its analogue and digital blocks, VDDA and VDDD, each of which has its own Low Drop Out (LDO) regulator to provide some isolation from movement of the input voltage planes. The power distribution schema for a single ABC130 power domain during wafer test is outlined in Figure A.1, for the purposes of this discussion this is equally applicable to either domain. Measurements of the supply




Figure A.1: Powering ABC130 during ASIC testing. The power supply is represented by voltage source V_{PSU} on the left. Power is lost in the cabling from the power supply to the driver board, of resistance $2R_a$, and again in the cabling from the driver board to the ABC130 Single chip board, of resistance $2R_b$. The input voltage to the single chip test PCB or probe card chip is therefore lower than V_{PSU} : this voltage is represented by voltmeter V_{RAW} . Additional power losses will occur in the PCB, wirebonds and, in the case of a probed chip, the contact resistance of the probe needles, of total resistance $2R_c$. Hence the voltage which actually reaches the chip will be lower than V_{RAW} . The chip also contains a linear regulator which is set to deliver voltage $V_{REG} = 1.2$ V by adjustment of digitally programmable resistance R_2 . The diagram also makes reference to V_{MG} (“Module Ground”), a measurement of the chip’s 0 V rail with respect to system ground. (The schematic is equally applicable to the chip’s analogue or digital power domain.)

voltage V_{PSU} VDDA, VDDD and current i_{ps} {IDDA, IDDD} are made at the power supply, connected to the host PC by means of its integral Universal Serial Bus (USB) interface. The resistance of the cable between the power supply and the driver board, R_a , is low and can be generally be neglected as can the resistance of the connection through the driver board itself. Between the driver board and the probe card 0.05” twisted pair ribbon is used: although the supply current i_{ps} passes through multiple wires, the equivalent resistance of this cable $R_b > R_a$. The voltage delivered to the probe card, V_{RAW} {VDDA.RAW, VDDD.RAW}, is monitored by means of sense wires and a commercial ADC chip (AD7998) on the driver

board. Further resistance due to the probe card PCB and probe needles, contact resistance and indeed the power planes within the chip itself exists before the LDO is reached. The sum of these factors is represented by R_c in the diagram. The LDO output voltage V_{REG} {VDDD_REG, VDDA_REG}, routed off-chip to permit the addition of external decoupling capacitance, is also measured. As there is essentially no DC current flowing through these connections, represented by R_d , we can be confident that once the LDO is tuned to set V_{RAW} to the nominal operating voltage of 1.2 V - by adjustment of a DAC within ABC130 which sets the value of R_2 - the chip is powered correctly.

Figure A.1 also shows V_{MG} (VMOD_GND), a measurement of “Module Ground”, the potential of ABC130’s 0 V rail with respect to system ground. Recall that the driver board offsets the ground of the chip by approximately 600 mV with respect to system ground such that standard LVDS drivers, which have a common mode voltage of 1.2 V, may communicate with ABC130 which expects to receive SLVS signals with a common mode voltage of 600 mV with respect to its 0 V rail. As V_{MG} is measured at the driver board it can be seen that ground rise across resistance $R_b + R_c$ will introduce a small error to the measurement.

The first portion of the test supplies constant voltages VDDA and VDDD to the chip, and attempts to configure it. Die which draw abnormal currents at power on, or die for which the analogue current does not increase and/or the locally regulated voltages are abnormally low upon configuration are flagged as rejects and the test sequence is discontinued. Most frequently such anomalies are due to high contact resistance such that insufficient voltage reaches the chip, and the affected die will pass once the contact height has been adjusted. The cuts used during this portion of the test sequence are summarised in Table A.1.

A second part of this test verifies the function of the ABC130’s built-in shunt transistor, which is connected to the digital power domain. This circuit block, not shown in the schematic above, was originally envisioned as part of the possible serial powering scheme. The test verifies that an external op-amp circuit is able to control the gate of the shunt in order to dissipate surplus current and regulate VDD_RAW correctly. The cuts used during this portion of the test sequence are summarised in Table A.2.

Parameter	Condition	Measurement	Cut	Code	Classification
VDDA	Power up	PSU voltmeter	<1.45 V	0x0001	X
IDDA	Power up	PSU ammeter	<25 mA	0x0002	X
VDDD	Power up	PSU voltmeter	<1.45 V	0x0004	X
IDDD	Power up	PSU ammeter	<28 mA	0x0008	X
IDDA	Configured	PSU ammeter	Does not increase	0x0010	X
VDDA_RAW	Configured	Driver ADC	<1400 mV	0x0020	X
VDDA_REG	Configured	Driver ADC	<1100 mV	0x0040	X
VDDD_RAW	Configured	Driver ADC	<1400 mV	0x0080	X
VDDD_REG	Configured	Driver ADC	<1100 mV	0x0100	X
VMOD_GND	Configured	Driver ADC	>700 mV	0x0200	X

Table A.1: Cuts on ABC130 Constant Voltage Power Measurements resulting in Classification “X” (reject)

Parameter	Condition	Measurement	Cut	Code	Classification
VDDD	Shunt Active	PSU voltmeter	<1.50 V	0x0001	X
IDDD	Shunt Active	PSU ammeter	<80 mA	0x0002	X
VDDA_RAW	Shunt Active	Driver ADC	<1400 mV	0x0020	X
VDDA_REG	Shunt Active	Driver ADC	<1100 mV	0x0040	X
VDDD_RAW	Shunt Active	Driver ADC	<1400 mV	0x0080	X
VDDD_REG	Shunt Active	Driver ADC	<1100 mV	0x0100	X
VMOD_GND	Shunt Active	Driver ADC	>700 mV	0x0200	X

Table A.2: Cuts on ABC130 Constant Current Power Measurements resulting in Classification “X” (reject)

The next portion of the test sequence aims to characterise each DAC which forms part of the ABC130 design. As most of these are made available by means of an analogue multiplexer built into the chip, this is referred to as the “MUX” test. A total of 20 nodes are characterised however only a subset are described here, starting with the characterisation of the chip’s dual (analogue and digital) Low Drop Out (LDO) regulators.




Figure A.2: Example Batch 3A LDO Characteristic (single die)

Within each LDO block, a control DAC is used to tune the feedback resistance in order to vary the set point, in effect adjusting resistor R2 in the schematic shown in Figure A.1. Whilst all other DACs in the ABC130 are binary weighted designs, in which a series of currents each defined by a different power of two are electronically switched in and out of circuit to give the desired output, the LDO block uses a thermometer-coded or unary-coded¹ design in which all of the switchable currents are the same. Hence a unary DAC with 16 switches is intended to have, including zero, 17 possible values, whereas a binary weighted design with 16 switches has $2^{16} = 65536$ possible values. In reality each of the 16 switchable currents of the LDO control DAC will be slightly different due to parasitics: the ABC130 design gives us access to switch each bit of the unary DAC in and out of circuit separately, meaning that one can take advantage of these small differences in order to achieve a better result. (There are several possible combinations of the 16 switches which may be closed to achieve summation of, say, 10 currents, and each will give a subtly different output value.) A representative LDO tuning result for a single chip of this batch is shown in Figure A.2. The software programs each regulator to develop the closest match to the desired 1.2 V output,

¹Unary coding, also known as thermometer coding, represents a natural number n as a sequence of n ones followed by zero.

values which will remain persistent for the duration of the test sequence. Die for which the voltage of 1.2 V may not be attained are rejected online, and the test is discontinued.




Figure A.3: Example Batch 3A Calibration and Threshold DAC Characteristic (single die)

Figure A.3 shows the characteristic of the 8-bit calibration and threshold DACs of a typical ABC130 chip, the only 8-bit DACs in the design ². Cuts on the slope of these DACs are applied offline to ensure that any non-monotonic behaviour is excluded. The distribution of these values for a typical wafer, AKQ8E9H, is shown in Figure A.4 to show that the cuts have deliberately been set outside the usual distribution of values. A further basic check is made by measuring the (static) bandgap voltage. If this lies outside a certain range, the calibration of most ADCs shall be affected, hence this is generally seen together with other anomalies. The cuts on the bandgap voltage and the slope of the calibration and threshold DACs are described in Table A.3.

²As these voltages are used to understand the performance of the front end amplifiers a greater resolution is required here than for the other DACs: those used to configure the operation of the amplifier utilise a common 5-bit design, and the DACs used to set the driver output current for differential signalling are a 3-bit design.




Figure A.4: Slope of Threshold and Calibration DACs for all die of wafer AKQ8E9H




Figure A.5: Example Batch 3A SLVS DAC Characteristic. In the left plot the XOFF_R transceiver (blue) is active and varied; in the right plot the DATA_L transceiver is active and varied; other transceivers remain at static levels throughout.

Signal	Parameter	Low Cut	High Cut	Fail Code	Classification
Bandgap Voltage	-	540.0	650	0x0001	X
Threshold DAC	Slope	0.42	0.51	0x0020	X
Calibration DAC	Slope	0.43	0.62	0x0040	X
		mV	mV		

Table A.3: Cuts on ABC130 MUX (DAC) Parameters resulting in Classification “X” (reject)

ABC130 incorporates four bidirectional SLVS transceivers as described in Section 3.2.2. The plots shown in Figure A.5 detail the effect of adjustment of the output current for the XOFF_R (blue) and DATA_L (black) transceivers respectively. Similar characterisations are made for the output current setting of the DATA_R (green) and XOFF_L (red) transceivers, although these are not shown. It will be noted that we monitor the inverted side of the differential pair for the XOFF signals and the non-inverted side for the DATA signals. A fit is made to the slope of the plot for each active transceiver and cuts re-applied to the fit parameters, as described in Table A.4. The main effect of the cuts is to exclude any drivers which exhibit non-monotonic behaviour.

Other DACs which form part of the ABC130 design are tested, notably those used to set the various biasing nodes and the trim DACs. In general any gross anomaly will result in the chip being rejected due to abnormal gain and/or noise measurements, hence cuts were not applied to these parameters during this phase of wafer screening activity.

The third part of the test sequence verifies that each front end input channel is fully functional and that the measured noise and gain fall within acceptable limits. The trim

Signal	Parameter	Low Cut	High Cut	Fail Code	Classification
SLVS XOFR	Offset	400	620	0x0100	X
SLVS DATAL	Offset	400	620	0x0200	X
SLVS DATAR	Offset	400	620	0x0400	X
SLVS XOFL	Offset	400	620	0x0800	X
		mV	mV		
Signal	Parameter	Low Cut	High Cut	Fail Code	Classification
SLVS XOFR	Slope	-42.5	-15	0x1000	X
SLVS DATAL	Slope	38	65	0x2000	X
SLVS DATAR	Slope	38	65	0x4000	X
SLVS XOFL	Slope	-42.5	-15	0x8000	X
		mV/count	mV/count		

Table A.4: Cuts on ABC130 SLVS driver measurements resulting in Classification “X” (reject)




Figure A.6: Typical ABC130 Gain and Noise Plots (Wafer Probe result)

DAC for each channel is set to a mid-range value (16 counts) before threshold scans are recorded for injected charges of 0.5, 1.0 and 1.5 fC. This is the three point gain measurement as described in subsection 3.3.3. Finally the trim DAC is raised to the maximum value (31 counts) and a threshold scan for an injected charge of 1.5 fC is repeated. From these four scans it is possible to determine the gain and noise of each channel, and to verify that each channel may be trimmed.

A typical gain and noise result for a single chip is shown in Figure A.6. The left hand plot shows the gain of each channel of the chip, which has a mean value of 31.0 DAC counts/fC (σ 1.38)³. The right hand plot shows the input noise (Equivalent Noise Charge) of each

³These figures are presented in DAC counts/fC rather than mV/fC as, at the time of this study, the calibration of the threshold DAC had not been fully understood, however one count is approximately 2.5 mV.




Figure A.7: Mean Gain and Mean Input Noise for all die of wafer AKQ8E9H

Parameter	Cut	Unit	Fail Code	Classification
No. Trimmable Channels	< 256	-	0x0002	B
Trim Width	< 5	mV	0x0004	B
No. Bad Channels	>0	-	0x0020	B
Mean Gain	<25	DAC counts/fC	0x0100	B
	>35	DAC counts/fC	0x0200	B
No. Bad Channels	>2	-	0x0040	X
No. Trimmable Channels	<254	-	0x0400	X
Mean Gain	<24	DAC counts/fC	0x1000	X
	>36	DAC counts/fC	0x2000	X
Mean Input Noise	<350	ENC	0x0008	X
Mean Input Noise	>550	ENC	0x0010	X

Table A.5: Cuts on ABC130 Analogue (Gain) Parameters resulting in Classification “B” (minor defect) or “X” (reject)

channel, which has a mean value of 458.2 electrons (sigma 29.3). Individual channels are considered as bad if the measured noise is less than 100 electrons or more than 750 electrons. If a chip has more than 2 bad channels it is rejected, alternatively die with 1 or 2 bad channels are categorised as grade “B”. In addition die are downgraded to category “B” if any channel(s) may not be trimmed, or if the portion of threshold space over which all 256 channels can be correctly trimmed is less than 5 mV (two counts on the threshold DAC).

Figure A.7 shows histograms of the mean gain and mean input noise for all die of wafer AKQ8E9H. Die having a mean gain value outside the range of 25 to 35 are graded as category “B”, those outside the range 24 to 36 are rejected outright (category “X”). Die having a mean input noise outside the range of 350 to 550 electrons are also rejected. The cuts applied to the gain and trim tests are summarised in Table A.5: it should be noted that these cuts

are applied equally to all wafers.

As stated earlier, the offline analysis also generates the pick maps which are sent to the dicing vendor, taking all cuts on all parameters into account. By way of an example, the pick map for wafer AKQ8E9H is shown in Figure A.8. This particular wafer has a total of 110 good (category “A”) die, comprising 52 A0 (non-FCF) and 58 A1 (FCF) die.




Figure A.8: Typical ABC130 Wafer Pickmap


Appendix B

Reduction of noise generated by the SemiProbe Probe Station

During early attempts to commission the SemiProbe PS4L [107] Probe Stations systems installed at Oxford and RAL, it was found that the noise measured for an ABC130 chip on either system was higher than that measured on the Micromanipulator machine by almost a factor of two, making sensible characterisation of ABC130 die impossible. Further investigations of this noise problem were made in October 2017.

Working with the machine at RAL, the noise voltage on the wafer chuck and metal structure of the system, measured with respect to mains earth, was found to be 32.25 V peak to peak as shown in Figure B.1a. By disconnecting the cables from the control unit one at a time it was determined that the noise signal was mostly coming from the stepper drive cables for the chuck axes. These cables were made of two shielded sections plugged together, however the screens were not connected across the break. This meant that the shield of the upper section of cable nearest the chuck was left floating with respect to the controller unit, leaving it free to pick up the stepper drive signals and hence to radiate these to nearby metalwork. Once this error, made by the vendor, had been corrected, the noise voltage on the chuck fell to 7 V.

Investigations continued, disconnecting cables from the system controller one by one, and



(a) Before Modifications

(b) After Modifications

Figure B.1: Voltage Noise measured on the SemiProbe PS4L chuck before and after a number of modifications were made, including screening of the stepper motor drive cables and the fitting of a number of ferrite chokes inside the control housing. The noise voltage was reduced from 32.25 V peak to peak to 1.5 V peak to peak. The vertical scale for Figures (a) and (b) are 5 V and 500 mV per division respectively. The horizontal scale is 5 μ s per division in both cases. A complete list of remedial measures is given in the text.

it was noted that the remaining noise signal was now dominated by the stepper drive cables for the microscope. Again the cables were made in two sections, but in this case the upper section was not screened at all. New upper microscope drive cables were made from stock parts and upon fitting them to the system the chuck noise voltage was further reduced to 3.5 V peak to peak. Additional minor improvements in the chuck noise voltage came from the fitting of ferrite chokes around the stepper wiring within the controller housing; connection of the screen of the system's signal (limit switch and encoder) cables; replacing the emergency stop wiring, which had been a coaxial cable, by a screened twisted pair type. At this point the noise voltage was down to 2.5 V: with the addition of heavy earth strapping between the system's base plate (an optical bench) and mains earth this was further reduced to 1.5 V as shown in Figure B.1b. This compares favourably to the 2 V peak to peak magnitude of the noise voltage observed on the chuck of the MicroManipulator machine (not shown).

Although the ABC130 wafer screening programme was completed using the MicroManipulator machine as documented in Chapter 4, the corrected SemiProbe system at RAL was successfully used to screen ABCStar wafers as documented in Section 9.3. The remedial treatment developed at RAL has also been applied to the Oxford machine.

References

- [1] CERN. The Large Hadron Collider (LHC). <https://home.cern/science/accelerators/large-hadron-collider>. [Online; accessed 19-June-2020].
- [2] CERN. The Large Electron-Positron (LEP) collider. <https://home.cern/science/accelerators/large-electron-positron-collider>. [Online; accessed 19-June-2020].
- [3] C. Lefèvre. The CERN accelerator complex. Complexe des accélérateurs du CERN. <http://cds.cern.ch/record/1260465>, Dec 2008.
- [4] CERN. Linear Accelerator 2 (Linac 2). <https://home.cern/science/accelerators/linear-accelerator-2>. [Online; accessed 19-June-2020].
- [5] CERN. The Proton Synchrotron Booster (PSB). <https://home.cern/science/accelerators/proton-synchrotron-booster>. [Online; accessed 19-June-2020].
- [6] CERN. The Proton Synchrotron (PS). <https://home.cern/science/accelerators/proton-synchrotron>. [Online; accessed 19-June-2020].
- [7] CERN. The Super Proton Synchrotron (SPS). <https://home.cern/science/accelerators/super-proton-synchrotron>. [Online; accessed 19-June-2020].
- [8] CERN. Linear Accelerator 3 (Linac 3). <https://home.cern/science/accelerators/linear-accelerator-3>. [Online; accessed 19-June-2020].
- [9] CERN. The Low Energy Ion Ring (LEIR). <https://home.cern/science/accelerators/low-energy-ion-ring>. [Online; accessed 19-June-2020].

- [10] The ALICE Collaboration. A Large Ion Collider Experiment (ALICE). <https://home.cern/science/experiments/alice>. [Online; accessed 19-June-2020].
- [11] The ATLAS Collaboration. The ATLAS Experiment. <https://home.cern/science/experiments/atlas>. [Online; accessed 19-June-2020].
- [12] The CMS Collaboration. The Compact Muon Solenoid (CMS) Experiment. <https://home.cern/science/experiments/cms>. [Online; accessed 19-June-2020].
- [13] The LHCb Collaboration. The Large Hadron Collider beauty (LHCb) experiment. <https://home.cern/science/experiments/lhcb>. [Online; accessed 19-June-2020].
- [14] The LHCf Collaboration. The Large Hadron Collider forward (LHCf) experiment. <https://home.cern/science/experiments/lhcf>. [Online; accessed 19-June-2020].
- [15] The MoEDAL Collaboration. The Monopole and Exotics Detector at the LHC (MoEDAL). <https://home.cern/science/experiments/moedal>. [Online; accessed 19-June-2020].
- [16] The TOTEM Collaboration. The 'TOTAL cross section, Elastic scattering and diffraction dissociation Measurement' (TOTEM) experiment. <https://home.cern/science/experiments/totem>. [Online; accessed 19-June-2020].
- [17] J. Pequenaio. Computer generated image of the whole ATLAS detector. <http://cds.cern.ch/record/1095924>, Mar 2008.
- [18] P. Jenni. EAGLE : Experiment for Accurate Gamma, Lepton and Energy measurements: expression of interest. 1992. Conference : Towards the LHC experimental programme : General Meeting on LHC Physics and Detectors, Evian-les-Bains, France, 5 - 8 Mar 1992.
- [19] P.R. Norton. The ASCOT detector at the LHC: expression of interest. 1992. Conference : Towards the LHC experimental programme : General Meeting on LHC Physics and Detectors, Evian-les-Bains, France, 5 - 8 Mar 1992.

- [20] The ATLAS Collaboration. ATLAS inner detector: Technical Design Report, 1. Technical Report CERN-LHCC-97-016, CERN, Geneva, 1997.
- [21] The ATLAS Collaboration. ATLAS inner detector: Technical Design Report, 2. Technical Report CERN-LHCC-97-017, CERN, Geneva, 1997.
- [22] The ATLAS Collaboration. ATLAS pixel detector: Technical Design Report. Technical Report CERN-LHCC-98-013, CERN, Geneva, 1998.
- [23] The ATLAS Collaboration. ATLAS liquid-argon calorimeter: Technical Design Report. Technical Report CERN-LHCC-96-041, CERN, Geneva, 1996.
- [24] The ATLAS Collaboration. ATLAS tile calorimeter: Technical Design Report. Technical Report CERN-LHCC-96-042, CERN, Geneva, 1996.
- [25] The ATLAS Collaboration. ATLAS muon spectrometer: Technical Design Report. Technical Report CERN-LHCC-97-022, CERN, Geneva, 1997.
- [26] J Pequeno and P Schaffner. How ATLAS detects particles: diagram of particle paths in the detector. <http://cds.cern.ch/record/1505342>, Jan 2013.
- [27] The ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716(1):1 – 29, 2012.
- [28] The CMS Collaboration. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, 716(1):30 – 61, 2012.
- [29] The ATLAS Collaboration. Muon reconstruction efficiency in reprocessed 2010 LHC proton-proton collision data recorded with the ATLAS detector. Technical Report ATLAS-CONF-2011-063, CERN, Geneva, Apr 2011.
- [30] The ATLAS Collaboration. Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data. Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data. *Eur. Phys. J. C*, 72(arXiv:1110.3174. CERN-PH-EP-2011-117):1909. 45 p, Oct 2011.

- [31] HL-LHC Project Schedule. <https://project-hl-lhc-industry.web.cern.ch/content/project-schedule>. [Online; accessed 19-June-2020].
- [32] O. Brüning. Machine Plans through HL-LHC: Goals & Status. 2020. ACES 2020, CERN, 26 - 28 May 2020.
- [33] H. Spieler. *Semiconductor Detector Systems*. Series on Semiconductor Science and Technology. OUP Oxford, 2005.
- [34] S.M. Sze and K.K. Ng. *Physics of Semiconductor Devices*. Wiley, 2006.
- [35] G. Lutz. *Semiconductor Radiation Detectors: Device Physics*. Springer Berlin Heidelberg, 2007.
- [36] F. Hartmann. *Evolution of Silicon Sensor Technology in Particle Physics*. Springer Tracts in Modern Physics. Springer International Publishing, 2017.
- [37] M. Tanabashi et al. Review of Particle Physics. *Phys. Rev. D*, 98:030001, Aug 2018.
- [38] Michael Moll. Displacement damage in silicon detectors for high energy physics. *IEEE Trans. Nucl. Sci.*, 65(8):1561–1582, 2018.
- [39] F. Faccio and G. Cervelli. Radiation-induced edge effects in deep submicron CMOS transistors. In *IEEE Transactions on Nuclear Science, Volume 52 Issue 6*, pages 2413–2420, December 2005.
- [40] G. Anelli et al. Radiation tolerant VLSI circuits in standard deep submicron CMOS technologies for the LHC experiments: practical design aspects. *IEEE Trans. Nucl. Sci.*, 46(6):1690–1696, 1999.
- [41] The ATLAS Collaboration. Technical Design Report for the ATLAS Inner Tracker Strip Detector. Technical Report CERN-LHCC-2017-005, ATLAS-TDR-025, CERN, Geneva, Apr 2017.
- [42] K. Dette. Total Ionising Dose effects in the FE-I4 front-end chip of the ATLAS Pixel IBL detector. *Journal of Instrumentation*, 11(11):C11028–C11028, Nov 2016.

- [43] J. Kaplon and S. Kulis. Review of input stages used in front end electronics for particle detectors. Technical Report PH-EP-Tech-Note-2015-001, CERN, Geneva, Mar 2015.
- [44] The ATLAS Collaboration. Expected Tracking Performance of the ATLAS Inner Tracker at the HL-LHC. Technical Report ATL-PHYS-PUB-2019-014, CERN, Geneva, Mar 2019.
- [45] I. Bonad. *Thermo-mechanical characterisation of low density carbon foams and composite materials for the ATLAS upgrade*. Doctoral Thesis. University of Glasgow, 2011.
- [46] F. Vasey et al. The Versatile Link PLUS Project. <https://espace.cern.ch/project-Versatile-Link-Plus/Shared%20Documents/Versatile%20Link%20PLUS%20Project%20V2.3.pdf/>. [Online; accessed 04-January-2018].
- [47] P. Moreira. <https://espace.cern.ch/GBT-Project/LpGBT/Specifications/LpGbtSpecifications.pdf/>. [Online; accessed 04-January-2018].
- [48] P. Keener and M. Newcomer. ATLAS ITk Electronics Specification: Autonomous Monitor And Control (AMACv2) ASIC. Technical Report AT2-IS-CD-0004, CERN, Geneva, Oct 2017.
- [49] M Menouni, P Gui, and P Moreira. The GBTIA, a 5 Gbit/s Radiation-Hard Optical Receiver for the SLHC Upgrades. In *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP09): Paris, France, 2009*.
- [50] G. Mazza, F. Tavernier, P. Moreira, D. Calvo, P. De Remigis, L. Olantera, C. Soos, J. Troska, and K. Wyllie. High-Speed, Radiation-Tolerant Laser Drivers in 0.13 μm CMOS Technology for HEP Applications. *IEEE Transactions on Nuclear Science*, 61(6):3653–3659, 2014.
- [51] P.W. Phillips. ATLAS SCT power supply system. In *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP07): Prague, Czech Republic, 2007*.
- [52] S. Lusin. The CMS low voltage system. In *Proceedings, Topical Workshop on Electronics for Particle Physics (TWEPP07): Prague, Czech Republic, 2007*.

- [53] A. Clark, U. Parzefall, and M. Tyndel. Status of the ATLAS SCT. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 579(2):580 – 591, 2007. Proceedings of the 6th International “Hiroshima” Symposium on the Development and Application of Semiconductor Tracking Detectors.
- [54] F. Faccio, G. Blanchot, C. Fuentes, S. Michelis, S. Orlandi, S. Saggini, and I. Troyano. FEAST2: A Radiation and Magnetic Field Tolerant Point-of-Load Buck DC/DC Converter. In *2014 IEEE Radiation Effects Data Workshop (REDW)*, pages 1–7, July 2014.
- [55] W. Dabrowski et al. Design and performance of the ABCN-25 readout chip for ATLAS inner detector upgrade. In *2009 IEEE Nuclear Science Symposium and Medical Imaging Conference*, pages 373–380, 2009.
- [56] M. Warren et al. The Buffer Control Chip. <http://www.hep.ucl.ac.uk/atlas/upgrade/BufferControlChip.pdf>, 2012. [Online; accessed 05-August-2020].
- [57] L. Poley et al. The ABC130 barrel module prototyping programme for the ATLAS strip tracker. *Journal of Instrumentation*, 15(09):P09004–P09004, Sep 2020.
- [58] The ATLAS Collaboration. Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System. Technical Report CERN-LHCC-2017-020. ATLAS-TDR-029, CERN, Geneva, Sep 2017.
- [59] B. M. Demirköz. *Construction and Performance of the ATLAS SCT Barrels and Cosmic Tests*. Doctoral Thesis. University of Oxford, 2007.
- [60] D. Nelson. HSIO Development Platform Users Guide. http://www.slac.stanford.edu/~djn/Atlas/hsio/HSIO%20Development%20Platform%20Users%20Guide_c02.pdf. [Online; accessed 28-May-2019].
- [61] Xilinx. Spartan-3A FPGA Family:Data Sheet. https://www.xilinx.com/support/documentation/data_sheets/ds529.pdf. [Online; accessed 24-June-2020].

- [62] Probe Test Solutions Limited (PTSL). Company web site. <http://www.probetestsolutions.com/>. [Online; accessed 28-June-2020].
- [63] NanoScope Services Ltd. FIB Circuit Edit. <http://www.nanoscpeservices.co.uk/fib-circuit-edit/>. [Online; accessed 29-January-2020].
- [64] The MicroManipulator Company. Company web site. <https://www.micromanipulator.com/>. [Online; accessed 28-June-2020].
- [65] Arcus Technology. 4 Axis USB Controllers: PMX-4EX-SA. <https://www.arcus-technology.com/products/multi-axis-motion-controller/4-axis-usb-controller/>. [Online; accessed 16-October-2019].
- [66] Arcus Technology. 3 Amp Driver: ACE-SDX-V2. <https://www.arcus-technology.com/products/stepper-motor-drivers/stepper-driver/>. [Online; accessed 16-October-2019].
- [67] Digilent. Atlys Spartan-6 FPGA Trainer Board. <https://store.digilentinc.com/atlys-spartan-6-fpga-trainer-board-limited-time-see-nexys-video/>. [Online; accessed 16-October-2019].
- [68] J. Kaplon and M. Noy. Front end electronics for SLHC semiconductor trackers in CMOS 90 nm and 130 nm processes. In *2011 IEEE Nuclear Science Symposium Conference Record*, pages 1955–1960, Oct 2011.
- [69] M. Karagounis. 8 Channel LVDS Receiver. Technical Report ATL-IP-ER-0025, CERN, Geneva, Dec 2010.
- [70] Ansys. High Frequency Structure Simulator (HFSS). <https://www.ansys.com/products/electronics/ansys-hfss>. [Online; accessed 05-August-2020].
- [71] Teledyne LeCroy. Signal Integrity Studio. <https://teledynelecroy.com/search/default.aspx#?cludoquery=SI%20Studio&cludopage=1>. [Online; accessed 5-August-2020].

- [72] J. Dopke et al. Lessons learned in high frequency data transmission design: ATLAS strips bus tape. *JINST*, 12(01):C01019, 2017.
- [73] Aim-TTi. CPX Series. <https://www.aimtti.com/product-category/dc-power-supplies/aim-cpxseries>. [Online; accessed 26-June-2019].
- [74] L.B.A. Hommels and S. Sigurdsson. Temperature Interlock. <http://www.hep.phy.cam.ac.uk/~hommels/interlock/TemperatureInterlock.htm/>. [Online; accessed 04-January-2018].
- [75] Arudino AG. Arduino home page. <https://www.arduino.cc>. [Online; accessed 29-May-2019].
- [76] E. Gornicki, S. Koperny, and P. Malecki. SCTHV VME Model 203 User's Manual. Technical Report 1947, IFJ, Krakow, 2004.
- [77] Digilent. Nexys Video Artix-7 FPGA Trainer Board. <https://store.digilentinc.com/nexys-video-artix-7-fpga-trainer-board-for-multimedia-applications/>. [Online; accessed 16-October-2019].
- [78] M. Stegler, I. Bloch, L. Poley, A. A. Affolder, L. B. A. Hommels, C. T. Klein, and Y. Unno. Investigation of the impact of mechanical stress on the properties of silicon strip sensors. *Nucl. Instrum. Meth.*, A924:133–136, 2019.
- [79] C. T. Klein. *Investigation of performance and the influence of environmental conditions on strip detectors for the ATLAS Inner Tracker Upgrade*. Doctoral Thesis. University of Cambridge, 2020.
- [80] L. B. A. Hommels. Private Communication. [Email; dated 04-January-2018].
- [81] S. Kilani. *Electrical Tests of the ATLAS Phase-II Strip Tracker Upgrade*. Doctoral Thesis. University College London, 2014.
- [82] Analog Devices. DC1073A, LT3845EFE Demo Board. <https://www.analog.com/en/design-center/evaluation-hardware-and-software/evaluation-boards-kits/dc1073a.html>. [Online; accessed 27-June-2019].

- [83] Analog Devices. LT3845 High Voltage Synchronous Current Mode Step-Down Controller with Adjustable Operating Frequency. <https://www.analog.com/en/products/lt3845.html>. [Online; accessed 22-June-2020].
- [84] Analog Devices. LT3845A High Voltage Synchronous Current Mode Step-Down Controller with Adjustable Operating Frequency. <https://www.analog.com/en/products/lt3845a.html>. [Online; accessed 22-June-2020].
- [85] Keysight. N2893A 100 MHz/15A AC/DC Current Probe. <https://www.keysight.com/en/pd-1912832-pn-N2893A/100-mhz-15a-ac-dc-current-probe>. [Online; accessed 22-June-2020].
- [86] P. Moreira, J. Christiansen, and K. Wyllie. GBTx Manual. <https://espace.cern.ch/GBT-Project/GBTX/Manuals/gbtxManual.pdf/>. [Online; accessed 20-December-2019].
- [87] Aim-TTi. MX Series. <https://www.aimtti.com/product-category/dc-power-supplies/aim-mxseries>. [Online; accessed 26-June-2019].
- [88] CERN. The Gigabit Link Interface Board (GLIB). <https://espace.cern.ch/project-GBLIB/public/default.aspx>. [Online; accessed 22-June-2020].
- [89] D. Porret. USB-I2C Adaptor Schematics. https://espace.cern.ch/GBT-Project/VLDB/Schematics/usb_i2c_schematics.pdf. [Online; accessed 26-June-2019].
- [90] P. Leitao. GBTx Programmer. <https://gitlab.cern.ch/gbtproj/gbtxprogrammer>. [Online; accessed 26-June-2019].
- [91] D. Lynn, R. Burns, P. Kuczewski, G. van Nieuwenhuizen, R. Roisin, Z. Schillaci, and S. Stucci. Module/Stave noise performance. <https://indico.cern.ch/event/776671/contributions/3241136/attachments/1765341/2865895/modulemeeting4Dec.pdf/>. [ATLAS internal; accessed 14-January-2020].
- [92] ISEG. EHS Series High Voltage Module. <https://iseg-hv.com/en/products/detail/EHS>. [Online; accessed 26-June-2019].

- [93] P.W. Phillips. FMC-ITk-I2C, I2C Multiplexer Schematic. <https://cernbox.cern.ch/index.php/s/DQCrRrMYdoAUP5Z>. [Online; accessed 26-January-2020].
- [94] Sorensen. XPF Series. <https://www.powerandtest.com/power/dc-bench-power-supplies/xpf-series>. [Online; accessed 22-June-2020].
- [95] L.B.A. Hommels et al. Detailed studies of full-size ATLAS12 sensors. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 831:167 – 173, 2016. Proceedings of the 10th International “Hiroshima” Symposium on the Development and Application of Semiconductor Tracking Detectors.
- [96] Digilent. Genesys 2 Kintex-7 FPGA Development Board. <https://store.digilentinc.com/genesys-2-kintex-7-fpga-development-board/>. [Online; accessed 16-October-2019].
- [97] Hitech Global. HTG-FMC-X4SFP+, FMC module with Quad SFP+ ports. http://www.hitechglobal.com/FMCModules/FMC_4SFP+_Module.htm. [Online; accessed 24-June-2020].
- [98] Broadcom (formerly AVAGO). AFBR-709SMZ 10Gb Ethernet, 850 nm, 10GBASE-SR/SW,SFP+ Transceiver. <https://www.broadcom.com/products/fiber-optic-modules-components/networking/optical-transceivers/sfpplus/afbr-709smz>. [Online; accessed 24-June-2020].
- [99] S. Ryu on behalf of the ATLAS TDAQ Collaboration. FELIX: The new detector readout system for the ATLAS experiment. *Journal of Physics: Conference Series*, 898:032057, Oct 2017.
- [100] J. Schumacher, C. Plessl, and W. Vandelli. High-Throughput Network Communication with NetIO. Technical Report ATL-DAQ-SLIDE-2016-842, CERN, Geneva, Nov 2016.
- [101] A. Greenall. Private Communication. [Email; dated 05-February-2020].

- [102] P.W. Phillips. FMC-1701, Star Chip Driver Schematic. <https://cernbox.cern.ch/index.php/s/y9hiQytDws7dW1P>. [Online; accessed 22-April-2020].
- [103] IEEE Standard for Test Access Port and Boundary-Scan Architecture. *IEEE Std 1149.1-2013 (Revision of IEEE Std 1149.1-2001)*, pages 1–444, 2013.
- [104] Texas Instruments. INA230 28V, 16-bit, I2C output current/voltage/power monitor w/alert. <https://www.ti.com/product/INA230>. [Online; accessed 28-June-2020].
- [105] P.W. Phillips. StarChipsSCB, ABCStar Single Chip Board Schematic. <https://cernbox.cern.ch/index.php/s/xT3xoMW8pAHJqpR>. [Online; accessed 22-April-2020].
- [106] P.W. Phillips. StarChipsProbe, ABCStar Probe Card Schematic. <https://cernbox.cern.ch/index.php/s/hUnouQ4050dFuRA>. [Online; accessed 22-April-2020].
- [107] SemiProbe. Probe System For Life (PS4L). <https://www.semiprobe.com/probe-station-product-lines/probe-system/>. [Online; accessed 28-June-2020].
- [108] K. Baker and J. von Beers. Shmoo plotting: the black art of IC testing. In *Proceedings International Test Conference 1996. Test and Design Validity*, pages 932–933, 1996.
- [109] ON Semiconductor. FS7140: I2C Programmable 1-PLL Clock. <https://www.onsemi.com/products/clock-timing/clock-generation/pll-clock-generators/fs7140>. [Online; accessed 27-June-2020].
- [110] CERN. ROOT Data Analysis Framework. <https://root.cern.ch/>. [Online; accessed 16-October-2019].