



FIRMWARE DEVELOPMENT FOR THE ATLAS LEVEL-1 CALORIMETER TRIGGER CLUSTER PROCESSOR MODULE

Tamsin Helen Moye

*Thesis submitted for the degree of
Doctor of Philosophy*



Particle Physics Group,
School of Physics and Astronomy,
University of Birmingham.

September, 2006

Acknowledgements

My viva examiners pulled me up on the absence of an acknowledgements section in this thesis, and it should not be underestimated how many people to whom I owe a huge amount of gratitude for their assistance in completing this, so I found a few minutes to identify at least those to whom I owe the most:

Firstly I must thank my supervisor, Professor Peter Watkins, who must have found me one of the more difficult students to guide, and yet his patience and encouragement never waned. Short of actually writing this thesis for me, he couldn't have done more.

I would like to extend much gratitude to the many people who I worked alongside on the ATLAS Level-1 Calorimeter Trigger project, especially those in Birmingham. I can't name everyone here, but I must at least mention Gilles Mahout for all his guidance in firmware development, Bruce Barnett for his encouragement during all testing at RAL, Kambiz Mahboubi for all the patience he found for us students during the long shifts on the CERN testbeam, and to Ethan Woehrling (especially for picking me off those German trainlines), Pavel Weber, Frederik Rühr, Adam Davis and Stefan Rieke, to name but a few of the people who brightened some of the longer meetings.

Socially, for my years in Birmingham, I owe huge thanks to all of the students I met both in the department and outside (mostly Staff House or poker games). Special

mention must go to my first year housemates and peers: Paul (Bacon), (Little)Matt and Alistair (poker), to the guys I shared offices with, especially Stevie (for his tea and inspiration), Daniel (especially for the printing, binding and Mexican-Spanish lessons), and (chatty) Richard (for his crossword assistance). For their ongoing inspirational inputs, even after they (or I) left the Uni, thanks hugely to Yves (for canal runs and housemate adoption), Oliver (for tea, crosswords, and at least one good link a week), Bruce (for all the last crossword answers), Rory (for the strangest and funniest stories ever), Devans (for skiing, drinking and diverse debates), Richard (Penny - for the half-empty glass, and questionable websites) and cold-hands Ian (for drinks, flumps, and gum).

Finally I must extend my gratitude to those whose encouragement has enabled me to bring this all together. Thanks to all my buddies outside Uni, especially those who've taken me walking (especially Marky and Roz), and listened to my rants over irc (especially the #craq population). Thanks hugely to Simon for all of his encouragement, support, motivation and patience over the stupidly late nights. For the final year of inspiration and perspective I'd also like to thank my new colleagues, especially Rachel for the gentle pestering and encouragement, and David and Andy for helping me to find spare days to finish things off.

More than anyone though, I must thank my family, especially Adie, for everything.

Synopsis

The first chapter of this thesis introduces the motivation behind the Large Hadron Collider (LHC) at CERN - the Standard Model of Particle Physics and possible physics beyond the Standard Model. The second chapter gives an overview of the LHC, and the ATLAS detector in particular, with a focus on calorimetry which provides the input for the Level-1 Calorimeter Trigger. Chapter three explains the ATLAS Trigger system and focuses on the Level-1 Calorimeter Trigger and Chapter four provides more detail about the Level-1 Calorimeter Trigger Cluster Processor Module, and discusses the logic within the FPGAs on the board. Chapter five describes the testing environment used for the CPM, including the software tools for firmware development and the hardware modules designed specifically for aiding tests. Chapter six goes into detail about testing of the CPM and specifically firmware-based testing using the chips on the board, and Chapter seven describes integration tests with the full Level-1 Trigger, focusing on the H8 CERN testbeam during the autumn of 2004.

Author's Contribution

The development of firmware versions, as described throughout the various tests in Chapter 6, was conducted by the author both during and after familiarisation with the tools, environment and programming language described in Chapter 5. The development of bit-error rate tester firmware versions for the Serialiser chips, and subsequently for the CP chips, was carried out by the author. The firmware versions to correct the bit-order of data arriving on the CPM, and to handle data received from the old and new MCM ASICs, were also the work of the author. The author also conducted extensive latency measurements across the CPM, which led to optimisations for later versions both of the CPM PCB and the CP firmware, as described in Chapter 6.

The ensuing tests described in Chapter 6 were also carried out by the author, with assistance from Gilles Mahout in Birmingham, and in integration tests guided by Bruce Barnett at RAL.

The author also provided general assistance during the running of the H8 CERN testbeam described in Chapter 7. Work with both the old and new MCMs at the CERN testbeam, including establishing pedestals for the new MCMs, BCID of calibration pulses, and energy measurements comparing data received in the CPM and JEM were also conducted by the author alongside PhD student Pavel Weber from the University of Heidelberg, and the wider Level-1 Trigger group.

Contents

1	Physics Context	1
1.1	The Standard Model of Particle Physics	1
1.1.1	Elementary Fermions	1
1.1.2	Gauge Bosons	2
1.1.3	The Higgs Mechanism	4
1.2	Beyond the Standard Model	4
1.2.1	Grand Unified Theories	6
1.2.2	Supersymmetry	7
1.3	Physics at the Large Hadron Collider	9
1.3.1	Accelerator Luminosity	12
1.3.2	Proton-Proton Physics	13
2	Experiment Overview	17

2.1	The Large Hadron Collider	17
2.2	The ATLAS Experiment	21
2.2.1	Particle Detection	22
2.2.2	Coordinate Systems	24
2.2.3	Detector Geometry	24
2.2.4	Inner Detector	25
2.2.5	Calorimeters	29
2.2.6	Muon Spectrometer	35
3	The ATLAS Trigger System	39
3.1	The Trigger Concept	40
3.1.1	Latency	41
3.1.2	The Level-1 Trigger	41
3.1.3	The Level-2 Trigger	44
3.1.4	The Level-3 Trigger	45
3.2	Calorimeter Trigger Architecture	46
3.2.1	PreProcessor	46
3.2.2	Jet/Energy-Sum Processor	51
3.2.3	Cluster Processor	53

3.2.4	Common Merger Module	53
3.2.5	ReadOut Driver	55
3.2.6	Processor Backplane	56
4	The Level-1 Calorimeter Trigger Cluster Processor	59
4.1	Overview of the Cluster Processor	60
4.1.1	The Cluster Processor Module	60
4.2	FPGAs	63
4.3	Serialiser FPGA	66
4.3.1	Real Time Data Path	67
4.3.2	Readout Logic	68
4.4	Cluster Processor FPGA	69
4.4.1	Trigger Algorithms	70
5	Development Environment	73
5.1	Software Tools	73
5.1.1	HDL Designer	74
5.1.2	ModelSim	75
5.1.3	Leonardo Spectrum and Xilinx	78

5.2	Hardware	79
5.2.1	DSS Module	80
5.2.2	LSM Module	81
6	CPM Testing	83
6.1	Overview of CPM Tests	83
6.2	Serialiser Chip Tests	84
6.2.1	Real Time data tests	84
6.2.2	BER Tests	87
6.2.3	Bit-order modifications	92
6.2.4	All-0s PPrASIC parity error	93
6.3	CP Chip Tests	94
6.3.1	Real Time Testing	94
6.3.2	Latency Measurements	95
6.4	Summary	103
7	H8 CERN Testbeam	105
7.1	Introduction	105
7.2	Setup	106

7.2.1	Beam Characteristics	107
7.2.2	Trigger Towers and Cables	108
7.2.3	Timing and Trigger	109
7.3	Operation	109
7.4	Timing, Trigger, Control and Readout	111
7.5	Results	113
7.5.1	Establishing pedestals in the MCMs	113
7.5.2	BCID of TileCal calibration pulse	116
7.5.3	Beam Energy Measurements with the CPM and JEM	118
7.6	Summary	120
7.7	Conclusions	121

List of Figures

1.1	Mass Dependence of Common Higgs Decay Channels [7]	5
1.2	Typical behaviour of the strong and electroweak couplings as functions of the four-momentum transfer squared, Q^2 in a typical GUT as per the current SM (left) and the MSSM (right) [10]	7
1.3	Statistical significance of the discovery of a SM Higgs at ATLAS [9]	11
1.4	Cross-section of p-p collisions as a function of cm energy for some key processes [13]	15
2.1	LHC dipole magnet flux plot - twin cores with opposite fields [17]	18
2.2	The accelerator chain at CERN [18]	20
2.3	The LHC Layout [19]	21
2.4	The ATLAS Detector [24]	22
2.5	Particle Penetration of Detector [25]	23
2.6	The Inner Detector [26]	26
2.7	ATLAS Calorimetry [27]	29

2.8	Muon Chambers [28]	35
2.9	Muon Spectrometer Subdetectors [29]	37
3.1	ATLAS Trigger System [30]	40
3.2	ATLAS Level-1 Trigger [30]	42
3.3	ATLAS Level-1 Calorimeter Trigger [32]	47
3.4	The PreProcessor Module [33]	48
3.5	The PreProcessor Multi-Chip Module [34]	49
3.6	The Jet/Energy-Sum Processor Module [34]	52
3.7	Photo of the CPM [34]	54
3.8	The Common Merger Module [34]	55
3.9	The ReadOut Driver [34]	56
3.10	Photo of a system crate containing (from left to right) CPU, CMM, JEM, 3 CPMs, another CMM and a TCM [41]	57
3.11	The Processor Backplane - front view where the modules plug in inside crate (left) and the back where cables plug in (right) [34] . . .	58
4.1	Organisation of Level-1 Calorimeter Trigger into CPMs and FPGAs [44]	61
4.2	Schematic layout of the CPM [45]	62

4.3	Data Chain through the CPM [46]	66
4.4	Block diagram of Serialiser functionality [46]	67
4.5	Serialiser data capture and synchronisation logic [46]	68
4.6	CP chip processing window [47]	69
4.7	Arrangement of trigger towers used for CP algorithm [47]	70
4.8	CP chip block diagram [47]	71
5.1	The HDL Designer control window and two design block-diagram windows	75
5.2	Screenshot of the ModelSim command window (top left), log window (top right) and the wave window (foreground)	76
5.3	ModelSim Wave Window [48]	77
5.4	ModelSim List Window [48]	78
5.5	Five DSS modules feeding LVDS data to one CPM [52]	81
5.6	LVDS Source Module [53]	82
6.1	Schematic Layout of Test Procedure [52]	84
6.2	Error profile of data received inside the Serialiser as a function of the delayed 40 MHz clock [55]	85
6.3	Errors by Serialiser of data received as a function of the delayed 40 MHz clock [55]	86

6.4	Schematic Layout of the Serialiser	88
6.5	Wave diagram of bit error rate tester as viewed in ModelSim	90
6.6	Expected order of data arriving on CPMs [59]	92
6.7	Actual order of data arriving on CPMs [59]	92
6.8	Swapping pin allocations in Serialiser .ucf file [59]	93
6.9	Expected parity in all-0s case [59]	93
6.10	PPrASIC parity error in all-0s case [59]	93
6.11	Diagram of delayed data tests between CP and Serialiser chips [52]	95
6.12	Error profile of data received inside the CP chip as a function of the delayed 40 MHz clock	96
6.13	Latency measurements across the CP chip using ModelSim	97
6.14	HDL Designer block diagram illustrating the top level of the CP chip, cp_top	98
6.15	Diagram of how the trigger towers relate to the CP inputs	98
6.16	HDL Designer block diagram illustrating the main functional part of the CP chip, cp_core	99
6.17	HDL Designer block diagram illustrating the constituents of the front end block within the CP chip	100
6.18	HDL Designer block diagram illustrating the constituents of the al- gorithm block within the CP chip	100

6.19	HDL Designer block diagram illustrating the composition of the dmux_shell block within the CP front_end logic	101
6.20	HDL Designer block diagram illustrating the composition of the alg_4x4 block within the main algorithm block	101
6.21	Latency across the components of CPM 1.8	102
7.1	Schematic of the Testbeam setup [61]	106
7.2	The ATLAS Trigger System at the Testbeam [62]	110
7.3	A liquid argon pulse digitised by a PPM [62]	112
7.4	Timing slices illustrating hits, and the pedestal set by the MCM is visible in the empty (5th) pulse [64]	114
7.5	Timing slices illustrating saturated pulses (4th and 5th) in the MCMs [64]	115
7.6	Time slices illustrating the pedestals with BCID implemented [64] . .	116
7.7	Time slices readout from the CPM, bypassing BCID (left) and with BCID (right) [64]	117
7.8	Readout of the total energy sum over all channels in the CPM and JEM [65]	118
7.9	250 GeV beam events read out from the CPM and JEM [62]	119
7.10	Correlation between energies of events read out from the CPM and JEM [62]	120

List of Tables

1.1	Standard Model - Leptons [2]	2
1.2	Standard Model - Quarks [3]	2
1.3	Standard Model - Bosons [4]	3
1.4	Standard Model particles and their superpartners [11]	8
1.5	Dominant Decay Modes of the Higgs [11]	12
3.1	Level-1 Trigger criteria for some physics processes	43
3.2	Level-1 Trigger rates at full luminosity	44

Chapter 1

Physics Context

1.1 The Standard Model of Particle Physics

The last fifty years has seen the development of one of the most comprehensive theories of modern physics: the Standard Model (SM) of Particle Physics. It is a local gauge theory describing the behaviour of, and interactions between, fundamental constituents of matter. It has been tested over the past decade to unprecedented accuracy, but parameters such as the particle masses are still phenomenologically introduced into the theory.

1.1.1 Elementary Fermions

The SM identifies twelve fundamental particles which constitute all matter: six leptons (see Table 1.1) and six quarks (see Table 1.2), which make up all hadrons. These twelve particles are all fermions and have $\frac{1}{2}$ -integer spins. Additionally, each of these particles has its own antiparticle. Furthermore, leptons can be classified into three flavours which differ mainly by the different masses of the charged leptons. Table

1.1 summarises the most important properties of the leptons. The corresponding values for the antiparticles can be obtained by adjusting quantum numbers accordingly, e.g. by multiplying the electric charge by -1. Recent discoveries [1] indicate that neutrinos flavour oscillate over time, i.e. depending on the distance over which they have travelled, a fraction of electron neutrinos will appear to have oscillated into muon neutrinos. The possibility of this phenomenon, while not yet completely understood theoretically, is the reason for the undetermined entries in the third column.

Particle	Mass (MeV)	Mean Life(s)	Electric Charge
electron, e^-	0.511	∞	-e
electron neutrino, ν_e	$< 3 \times 10^{-6}$	∞	0
muon, μ^-	105.7	2.197×10^{-6}	-e
muon neutrino, ν_μ	< 0.19	∞	0
tau, τ^-	1777	290.6×10^{-15}	-e
tau neutrino, ν_τ	< 18.2	∞	0

Table 1.1: Standard Model - Leptons [2]

The same categorisation is done for quarks (see Table 1.2). To obtain the correct values for the corresponding antiquarks, some quantum numbers have to be adjusted, i.e. the colour charge must be inverted and the electric charge multiplied by -1.

Quark	Electric charge (e)	Mass
Up, u	$\frac{2}{3}$	1 - 5 MeV
Down, d	$-\frac{1}{3}$	2 - 9 MeV
Charmed, c	$\frac{2}{3}$	1.15 - 1.35 GeV
Strange, s	$-\frac{1}{3}$	75 - 170 MeV
Top, t	$\frac{2}{3}$	170 - 179 GeV
Bottom, b	$-\frac{1}{3}$	4 - 4.4 GeV

Table 1.2: Standard Model - Quarks [3]

1.1.2 Gauge Bosons

Additionally, the SM identifies particles with integer spins (bosons), which act as mediators of the four fundamental forces (electromagnetic, weak, strong and gravi-

tational). All interactions between particles are the result of these forces.

Force	Mediator	Acts on	Range/m	Relative Strength
Strong	gluon, g	colour charge	10^{-15}	1
Electromagnetic	photon, γ	electric charge	∞	10^{-2}
Weak	$W^\pm, Z^0$	weak charge	10^{-18}	10^{-5}
Gravitational	(graviton, G)	mass	∞	$\sim 10^{-40}$

Table 1.3: Standard Model - Bosons [4]

The strong force is mediated by the massless gluon which couples to the colour charge (red, green or blue) carried by all quarks. Since gluons are colour charged themselves (they carry a colour and an anticolour), they can interact with each other, resulting in a finite range of the strong force even though gluons are massless. The theory describing the strong interaction is called Quantum Chromodynamics (QCD) [5].

The massless photon is the mediator of the electromagnetic force which couples to the electric charge. The electromagnetic interaction is described by the theory of Quantum Electrodynamics (QED) first developed by R. Feynman and others in 1948 [5].

The weak interactions are mediated by the exchange of either a neutral vector boson Z^0 ($m_Z = 91.19$ GeV, [3]), or one of the charged vector bosons W^+ or W^- ($m_{W^\pm} = 80.43$ GeV, [3]). These massive bosons couple to the weak charge which is carried by all leptons and quarks.

In 1968, Glashow, Weinberg and Salam were able to unify the electromagnetic and weak interactions using a single mathematical framework: the theory of electroweak interactions. The relative strengths of the fundamental forces and their respective ranges are shown in Table 1.3. The graviton has not yet been discovered, but the SM has helped lead to the discovery of bosons that transmit the other three of the four known forces, so the graviton is postulated as the particle which transmits the gravitational force.

1.1.3 The Higgs Mechanism

The Electroweak theory was developed in 1968 and unifies the electromagnetic and weak interactions. It involves a spontaneous symmetry breaking, which gives rise to the Z^0 and $W^\pm$ boson masses and the zero mass photon.

If a fourth ‘Higgs’ field is introduced to the SM, it will affect all particles, and its lowest energy state will be non-zero. The Higgs field permeates all space, and is a scalar field, holding a single value at all points in space.

The symmetry-breaking mechanism of the electroweak interaction predicts a neutral spin-0 scalar particle, which could be the Higgs boson, H . It has not yet been observed, but searches have eliminated masses < 114 GeV [6]. The number of Higgs bosons could be higher, and even a mechanism without a boson is conceivable. The mass of the Higgs boson cannot be theoretically predicted, but unitarity arguments, requirements of the stability of the electroweak vacuum, and the validity range of the SM imply certain bounds. These bounds are highly dependent on specific assumptions, though, and to include most sensible possibilities, the Higgs mass could be anywhere between 50 and 800 GeV. In the SM, once the mass of the Higgs is known, all decay properties of the Higgs are predicted, see Figure 1.1

The most likely decay channels for the Higgs, for masses < 140 GeV are $H \rightarrow f\bar{f}$ with $H \rightarrow b\bar{b}$ with a branching ratio of 0.85. For larger masses, $H \rightarrow W^+W^-$ or Z^0Z^0 are predicted to have the largest branching ratios.

1.2 Beyond the Standard Model

Although the SM is a very powerful predictive theory and at the present time all experimental observations are consistent with it, much remains unexplained includ-




Figure 1.1: Mass Dependence of Common Higgs Decay Channels [7]

ing the differing strengths of coupling constants, and why there are exactly three generations of quarks and leptons. There may also remain phenomena awaiting discovery, especially at higher energies, which are not described by the SM.

One of these possible areas is identified when attempts are made to carry unification further, by combining the electroweak and strong interactions in a higher, unified symmetry, which could be manifest at extremely high energy. The scale of these grand unified theories (GUTs) is believed to be $E_{GUT} \sim 10^{16}$ GeV. These are discussed in more detail in the following section.

Once an exploration of higher mass scales begins, other aspects begin to play a role, for example, the ‘hierarchy problem’, where radiative corrections to the electroweak

parameters, like corrections to the boson masses, come from loops with circulating virtual fermions and bosons. If there were more massive particles, $M \sim M_{GUT}$, these would occur in virtual processes at lower energy scales. If the mass of the Higgs particle is driven by these more massive Higgs objects of the GUT scale, its value would become unstable unless quite precise cancellations occurred at a level of $M_W M_{GUT} \sim 10^{-14}$ [8]. Supersymmetry (SUSY) models are designed to overcome this problem, and are discussed in Section 1.2.2.

Other shortcomings of the SM are that it is both incomplete and ugly. Gravity is not included, whereas a ‘theory of everything’ should encompass that too, as well as account for the many arbitrary parameters in the SM - some 17 or 18 empirical masses, couplings, mixing angles, etc. that have to be inserted ‘by hand’. Ambitious attempts to incorporate gravity with the other fundamental interactions, the so-called ‘supergravity theories’ will probably not be investigated by the ATLAS detector and so are not discussed here.

1.2.1 Grand Unified Theories

Following the successful unified theory of electroweak interactions, it naturally led to the question as to whether the strong interaction could also be unified with the weak and electromagnetic interactions into a single GUT. The strong interaction has a much greater strength than the electroweak at currently accessible energies, making this seem unlikely, but the strength of an interaction depends on the distance over which it acts, or equivalently, the four-momentum transfer squared (Q^2). The strong interaction coupling decreases with Q^2 while the electroweak interaction increases with Q^2 but varies much more slowly. A naive extrapolation from the low-energy values suggests that their various couplings might become equal at an enormous value $Q^2 = M_X^2 c^4$, where the unification mass, M_X is of the order 10^{16} GeV [9], see Figure 1.2. The Standard Model as it stands doesn’t enable these couplings

to cross (left graph), whereas with the introduction of the simplest SUSY theory, the Minimal Supersymmetric Standard Model (MSSM), they can become unified (right).




Figure 1.2: Typical behaviour of the strong and electroweak couplings as functions of the four-momentum transfer squared, Q^2 in a typical GUT as per the current SM (left) and the MSSM (right) [10]

In GUTs, all three interactions are united into a single interaction, characterised by a single coupling constant, at the unification mass. Differences between them emerge as one interpolates downwards to currently available energies.

1.2.2 Supersymmetry

The most popular GUTs incorporate a new proposed symmetry of nature, SUSY. According to this theory, every known particle has a SUSY partner, or superpartner, which is like it in all respects except its spin differs by $\frac{1}{2}$. Fermions gain boson superpartners, named by prefixing the original name with an ‘s’, producing sleptons and squarks, and bosons gain fermion superpartners, which gain the ending ‘ino’, see Table 1.4.

Particle	Symbol	Spin	Superparticle	Symbol	Spin
quark	q	$\frac{1}{2}$	squark	$\tilde{q}$	0
electron	e	$\frac{1}{2}$	selectron	$\tilde{e}$	0
muon	μ	$\frac{1}{2}$	smuon	$\tilde{\mu}$	0
tau	τ	$\frac{1}{2}$	stau	$\tilde{\tau}$	0
W	W	1	Wino	$\tilde{W}$	$\frac{1}{2}$
Z	Z	1	Zino	$\tilde{Z}$	$\frac{1}{2}$
photon	γ	1	photino	$\tilde{\gamma}$	$\frac{1}{2}$
gluon	g	1	gluino	$\tilde{g}$	$\frac{1}{2}$
Higgs	H	0	Higgsino	$\tilde{H}$	$\frac{1}{2}$

Table 1.4: Standard Model particles and their superpartners [11]

If SUSY was an exact symmetry, a particle and its superpartner would have exactly the same mass. This would trivially solve the Higgs mass problem described previously, by cancelling the virtual loops because of the opposite signs of the loops caused by the superpartners. This is evidently not the case as no SUSY particles have yet been detected, so SUSY is only an approximate symmetry. This is assumed to be the case in the SUSY version of GUTs in which even the lightest supersymmetric particles have masses which are of the same order of magnitude as the W and Z masses. When these supersymmetric particles are taken into account, it can be shown that the extrapolation of Figure 1.2 is modified in such a way that the grand unification mass, M_X is increased by an order of magnitude, while the value of the grand unified coupling constant, g_U remains relatively constant. It does enabled these three lines to meet, unifying the strong force with the electroweak, at this level.

To verify SUSY, it is of course necessary to detect the superparticles. According to many versions of the theory, these can only be created or destroyed in pairs, so that the decay of a superparticle must yield at least one superparticle in the final state, and the lightest such particle must be stable. There are several candidates for the lightest superparticle, but many models assume it is a neutralino, $\tilde{\chi}_0$, which

is a mixture of the photino, Higgsino and Zino. A simple reaction would then be:

$$q + \bar{q} \rightarrow \tilde{q} + \tilde{\bar{q}} \quad (1.1)$$

followed by decays of the squark and antisquark:

$$\tilde{q} \rightarrow q + \tilde{\chi}_0, \text{ and} \quad (1.2)$$

$$\tilde{\bar{q}} \rightarrow \bar{q} + \tilde{\chi}_0 \quad (1.3)$$

giving an overall reaction:

$$q + \bar{q} \rightarrow q + \bar{q} + \tilde{\chi}_0 + \tilde{\chi}_0. \quad (1.4)$$

The squark and antisquark decay very quickly before they can reach the detector. The neutralinos interact extremely weakly with matter, and so escape detection; so the final state quark and antiquark (seen as hadronic jets) are all that would be detected, and the events would be characterised by the fact that they will only account for a fraction of the energy of the initial state. They are also unlikely to be emitted back-to-back in the centre of mass frame. Hence there are clearly defined criteria for detecting these events, so long as there is sufficient energy to produce them. The Large Hadron Collider (LHC) will facilitate searches for SUSY particles, providing the masses of the SUSY partners are less than ~ 2 TeV. Another telltale sign of SUSY will be the manifestation of five SUSY Higgs bosons ($H^\pm$, h , H and A), rather than just the single Higgs predicted by the SM.

1.3 Physics at the Large Hadron Collider

Bunches of protons counter-rotate in the LHC with an energy of 7 TeV each. As they collide head-on the centre of mass (CM) energy $\sqrt{s}$ of a proton-proton collision

is 14 TeV. As the partons in the protons can be considered free particles at these energies and over the timescales of the interaction, typically only about one sixth of this energy is available for production of new particles through a parton-parton interaction. Hence the LHC will be able to observe hard interactions with typical $\sqrt{s} \approx 2$ TeV.

The total cross-section of an inelastic proton-proton collision is expected to be large at LHC energies, and when combined with the high luminosity, more than twenty collisions can occur in a single bunch-crossing. The huge amounts of data to be processed, and the difficulties in identifying a certain event in the pileup of background events, makes triggering for interesting physics events in the ATLAS detector very challenging.

It is quite possible that the SM will be valid for the whole or most of the LHC energy scale, making it a sound theoretical background for LHC events. The SM originated in the 1960s and is unsurpassed in precise predictions for the outcome of a wide spectrum of experiments since. Most aspects of the SM are well established, the exception being the mechanism which gives particles their masses. Several theories exist which attempt to provide explanation for this, the most prominent being the Higgs mechanism as described previously in this Chapter. The ATLAS experiment will put such theories to the test.

The LHC will allow the study of particle interactions at an energy scale never before observed at an accelerator. At the LHC, the focus of activities will include:

- Spontaneous symmetry breaking at the electroweak scale.
- The search for leptoquarks.
- The search for the lightest stable sparticle.
- The search for evidence of quark compositeness.

Spontaneous symmetry breaking could manifest itself as the single Higgs boson predicted by the Standard Model, or the quintet of Higgs bosons predicted by minimal supersymmetric extensions to the SM. The search strategies for both of these are well known. To cover the full potential mass range, $80 \text{ GeV} < m_H < 1 \text{ TeV}$, sensitivity is required for a variety of final states, see Figure 1.3 [11].




Figure 1.3: Statistical significance of the discovery of a SM Higgs at ATLAS [9]

Furthermore the search for the five SUSY Higgs bosons will also require sensitivity to other decays such as:

Decay	Mass Range / GeV/c^2
$H \rightarrow \gamma\gamma$	$90 < m_H < 150$
$H \rightarrow ZZ^* \rightarrow 4l^\pm$	$130 < m_H < 180$
$H \rightarrow WW \rightarrow l^\pm \nu 2jets$	$180 < m_H < 800$
$H \rightarrow WW, ZZ \rightarrow 2l^\pm 2\nu$	
$H \rightarrow ZZ \rightarrow 4l^\pm$	

Table 1.5: Dominant Decay Modes of the Higgs [11]

$$\begin{aligned}
A &\rightarrow \tau^+ \tau^- \rightarrow e + \mu + \nu's \\
&\rightarrow l + hadrons + \nu's \\
H^\pm &\rightarrow \tau^\pm + \nu
\end{aligned}$$

It is thought that the Higgs decay channels will be visible above the dominant hadronic background at the LHC. However, the expected cross-sections for these processes are small for both low Higgs mass (< 180 GeV) and the very high mass ranges, hence the requirement of very high luminosity.

There are further models of electroweak symmetry breaking, predicting the existence of currently undiscovered particles, with masses in the TeV range. The LHC will cover this range, providing insight into the mechanism of symmetry breaking, whether in the form of new particles, or in their absence.

Furthermore, it aims to investigate certain important parameters of the Standard Model, including high-precision measurements of the masses of the W boson and the top quark. Running at design luminosity, $10^7 t\bar{t}$ pairs and 3×10^8 W bosons are expected to be produced each year [12].

1.3.1 Accelerator Luminosity

The target for the initial physics run luminosity is $10^{33} \text{cm}^{-2} \text{s}^{-1}$. This will be gradually increased to the design luminosity of $10^{34} \text{cm}^{-2} \text{s}^{-1}$. The high luminosity is

required as the cross-sections for heavy particle production (eg. the Higgs boson) are extremely small [11]. The luminosity, $\mathcal{L}$ of the LHC is given by

$$\mathcal{L} = \frac{N_a \cdot N_b \cdot j \cdot v / C}{A} \quad (1.5)$$

where N_a and N_b are the number of protons in bunches a and b , j is the number of bunches in each beam, v is the velocity of the protons, C is the circumference of the accelerator ring, and A the cross-sectional area of the beam. The reaction rate, R , at a given interaction point, is related to the cross-section, σ , of a particular reaction and the luminosity by:

$$R = \sigma \cdot \mathcal{L} \quad (1.6)$$

It is evident that the LHC will be an ideal environment for the high energy search for new physics, surpassing the collision energy and luminosity of the Tevatron at Fermilab (the current highest) by factors of ~ 7 and ~ 10 respectively.

1.3.2 Proton-Proton Physics

The LHC is a proton-proton collider. It has advantages in that the energy available in a collision is higher than in previous colliders, but colliding hadrons also introduces unique problems. As the energy used is increased, the proton is probed to very small distances. This means the proton can no longer be regarded as a point particle and must instead be regarded as a collection of partons, or the constituent parts of the proton which are identified as three valence quarks (up, up, down), a sea of virtual quarks, antiquarks, and gluons. Those partons interact via the strong force and therefore QCD is the appropriate theory to describe the behavior of partons inside the proton itself as well as for interactions between partons of different protons. Unfortunately, QCD predictions can only be calculated perturbatively in the short-distance region since the coupling grows much stronger at large distances.

Proton-Proton Scattering

In many cases, two protons will scatter off each other elastically, corresponding to a low momentum transfer, which means that the internal structure of both protons remains stable throughout the interaction. The signature of such events are marginally deflected protons, i.e. the scattered proton lies in proximity of the original beam axis. Hard scattering, i.e. two partons colliding or even annihilating each other, will only take place in a tiny fraction of all collisions. Characteristic signatures of such events are particle jets in or near the plane perpendicular to the beam axis, corresponding to a high momentum transfer between two partons. The centre of mass energy available for the production of new particles, for example Higgs bosons, is not equal to the total centre of mass energy $\sqrt{s}$ of the overall process since the involved partons carry just a fraction, x_1 and x_2 of the total momentum of their respective proton. These fractions are determined by the experimentally measured structure function $F_2(x_i, Q^2)$, with $i = 1, 2$ and Q being the four-momentum transfer between the two partons. Therefore the amount of energy available for final state particles is reduced to [4]:

$$\sqrt{x_1 x_2 s} \approx \frac{1}{6} \sqrt{s} \quad (1.7)$$

Figure 1.4 shows the partial cross-sections for heavy quark and Higgs boson production as well as the total cross-section for proton-proton scattering as a function of the total centre of mass energy $\sqrt{s}$. Some physically interesting processes, like the direct production of the Higgs boson, only contribute an extremely small part to the overall cross-section. The fraction of the cross-sections in this case is [14]:

$$\sigma_{Higgs}/\sigma_{Total} \approx 10^{-11} \quad (1.8)$$

The strongly interacting gluons and quarks produce a large number of QCD particles at a much higher rate than any interactions likely to be caused by new physics, see Figure 1.4. There is a difference of nine orders of magnitude, and consequently, a




Figure 1.4: Cross-section of p-p collisions as a function of cm energy for some key processes [13]

large number of collisions are required to generate a statistically significant sample of events containing new physics.

Another problem caused by colliding protons is that a large number of hadrons are produced, including neutrons. This provides around twenty extra hadrons at each collision in addition to the high E_T particles, causing a complex underlying event.

Once produced, the SM Higgs will immediately decay into different particles as predicted by the SM, and dependent only on the Higgs mass. The likelihood of

discovery of the Higgs by the ATLAS detector is a function of the branching ratios to the different decay products, the background processes that generate similar decay products, and the sensitivity of the ATLAS detector to the decay products.

Many of the decay modes include leptons, which must be reliably distinguished from the hadronic jets. In a decay to neutrinos, the only way to detect them is by measuring the absence of momentum needed to balance that carried by other particles. Decays of b quarks are recognisable by the relatively long life of hadrons containing b quarks, which allows them to travel several hundred microns before decaying. The ATLAS detector is designed to specifically detect these various decay modes and details of how this is done are discussed in the next chapter.

Chapter 2

Experiment Overview

To study the Standard Model to a high accuracy, and to search for physics beyond it, a high-energy, high-luminosity collider is required. The LHC is expected to be operational at CERN in Geneva, Switzerland in 2007. This chapter briefly describes the LHC and outlines the main components of the ATLAS experiment.

2.1 The Large Hadron Collider

The LHC [15] [16], is installed in the 27km long LEP (Large Electron-Positron Collider) tunnel and will make use of CERN's existing accelerators as injectors. It will collide counter-rotating bunches of protons at a frequency of 40.08MHz, with a centre of mass energy, $\sqrt{s}= 14\text{TeV}$. There will also be a programme of running with lead ions with $\sqrt{s}= 1148 \text{ TeV}$.

The LEP collider was limited to a centre of mass energy of 210 GeV by synchrotron radiation losses: particles radiate as they are bent around a ring. In order to achieve even higher energies, either a larger ring is required to house the accelerator, making

the curvature more gentle for the particles, or heavier particles must be used. In using the same sized ring to achieve higher energies, stronger magnets are required to bend the more energetic particles around the ring. The LHC will use protons instead of the electrons/positrons that LEP used. Since protons are approximately 2000 times heavier than electrons, synchrotron radiation losses are much reduced.

By using protons rather than antiproton-proton collisions, the LHC avoids the fact that antiprotons are difficult to create because they require a large number of protons to be fired onto a fixed target. However, it also means we cannot exploit the fact that, as was the case for LEP, two beams of oppositely charged particles could be contained in the same magnetic channel. The LHC is required to have two different channels for its beams, with opposite fields in each. This is achieved using 8.3 T superconducting dipole magnets, allowing the two beams to be contained in the same magnet (see Figure 2.1).




Figure 2.1: LHC dipole magnet flux plot - twin cores with opposite fields [17]

There will be 1296 large magnets made from superconducting Niobium-Titanium wire to generate the required 8.3 T field to bend the particle trajectories around the

ring. The magnets are cooled by liquid helium at 1.8 K. At this temperature the refrigerant is a superfluid with minimal viscosity and maximal thermal conductivity.

The protons are accelerated by a 50 MeV Linac, by the 1.4 GeV Proton Synchrotron (PS) booster, the 30 GeV PS, and finally the 450 GeV Super Proton Synchrotron (SPS), before being injected into the LHC accelerator ring (see Figure 2.2). The protons are accelerated in bunches which are spaced at 24.95 ns intervals around the LHC ring, producing the bunch crossing frequency of 40.08 MHz.

There are eight straight sections around the LHC ring, providing space for the other parts of the machine (see Figure 2.3). There are two sections (labelled 3 and 7) for cleaning the beam from any halo that might have formed from instabilities in the machine. One section (4) holds the RF cavitites which accelerate the protons from 450 GeV to 7 TeV. One section (6) holds the beam-dump, where the particles are dumped at the end of a run. The other four sections (1, 2, 5 and 8) are the interaction points, where the beams are steered and focussed to cause them to intersect and produce collisions.

Four experiments are to be situated at the four interaction points round the ring, where the beams cross (see Figure 2.3):

- ATLAS (A Toroidal LHC ApparatuS), a general purpose detector, [20].
- ALICE (A Large Ion Collider Experiment) will examine Pb-Pb interactions, [21].
- LHCb (LHC beauty experiment) will focus on CP violating decays of bottom quarks, [22].
- CMS (Compact Muon Solenoid), another general purpose detector, [23].




Figure 2.2: The accelerator chain at CERN [18]




Figure 2.3: The LHC Layout [19]

2.2 The ATLAS Experiment

More than 2000 physicists from 150 universities and laboratories in 34 countries all over the world participate in the ATLAS collaboration. The detector itself (Figure 2.4) measures 45 m long by 11 m radius, and weighs around 7000 tons.

ATLAS is a general-purpose experiment, and is thus designed for optimum performance over a broad range of physics tasks. The main aims of ATLAS, as detailed in the previous chapter, are to search for the Higgs boson and physics beyond the Standard Model as well as making many precision measurements.




Figure 2.4: The ATLAS Detector [24]

2.2.1 Particle Detection

The detector should record sufficient information to allow reconstruction of the final state of the proton-proton interaction. In order to do this, it is necessary to identify the particles, their direction, and their energy. Charged leptons, photons, hadrons and neutrinos will all be produced, as well as possibly some other stable new particles. It is not practical to detect neutrinos directly, so the vector sum of momentum deposited in the detector is carefully determined in order to reveal any that has been carried off by such particles.

A magnetic field is used to help identify the momentum and charge of particles as positive and negative charges curve in opposite directions with their tracks following helical paths where the momentum of the particle is proportional to the radius of curvature of the track. This measurement is made in the inner detector (Section 2.2.4). Charged particles are detected in the tracking detectors by ionising the material (see Figure 2.5).




Figure 2.5: Particle Penetration of Detector [25]

The energy of particles is then measured by absorbing them in calorimeters (Section 2.2.5). Electrons and photons interact electromagnetically with the dense material in the electromagnetic calorimeter, producing showers of electrons, positrons and photons.

Hadrons interact strongly with the nuclei of the dense materials in the calorimeters, but shower deeper into the calorimeter than the electromagnetic particles so leave little deposit of energy in the electromagnetic calorimeter but continue through to deposit most of their energy in the hadronic calorimeter.

Finally, muons interact little and pass through most of the detector, depositing some energy through ionisation in each layer of the detector, including the muon chamber at the external limit of the detector (Section 2.2.6).

2.2.2 Coordinate Systems

The origin is taken to be at the centre of the detector and the beamline lies along the z-axis. The y-axis is then vertical and the x-axis is horizontal, perpendicular to the beamline.

Spherical coordinates are used for the detector, with the polar angle, θ , the angle the particle makes with the beamline measured from the beam axis and ϕ is the angle around the beam axis.

Particles emerging from the interaction point have longitudinal momentum p_L in the z-direction and transverse momentum p_T in the x and y directions (r) directions. The Lorentz invariant quantity, rapidity (y) is expressed as:

$$y = \frac{1}{2} \log \frac{E + p_L}{E - p_L} \quad (2.1)$$

where E is the particle energy and p_L is the longitudinal momentum of the particle.

As rapidity varies for a given θ with the rest mass of the particle, pseudorapidity (η) is often used, where:

$$\eta = -\ln \tan \frac{\theta}{2}. \quad (2.2)$$

2.2.3 Detector Geometry

The ATLAS detector is a cylindrical, multi-layered structure, with sub-detectors assembled in concentric layers. Each sub-detector comprises a barrel section and two end-caps. The barrel measures the ϕ and z positions of particles at given radii whilst the end-caps measure ϕ and r at known z's. Effectively, the end-caps cover the region of high- η when extending the barrel would provide little additional information.

At the full operating luminosity of ATLAS, ~ 25 proton-proton collisions are expected for every bunch crossing which produce up to ~ 500 charged particles over $|\eta| < 2.5$. This spray of charged particles will mask any interesting physics channels we are looking for. The detector must have a fast response time and fine granularity to cope with pile-up of multiple interactions and each detector component must be radiation-hardened to survive the large radiation flux.

At the centre of the detector is the beampipe, along which the protons pass and where the collisions occur. Next is the inner detector (Section 2.2.4) comprising trackers to measure particle momenta and charge. Surrounding the inner detector are the calorimeters (Section 2.2.5). Innermost is the Electromagnetic Calorimeter (ECAL) and beyond that the Hadronic Calorimeter (HCAL) layer. Outermost is the Muon Spectrometer (Section 2.2.6) for detecting and tracking muons which have passed through the rest of the detector.

The following sections in this chapter give a brief summary of the ATLAS subdetectors: the inner detector, the calorimeters and the muon spectrometer. Emphasis will be put on the calorimeters, as they provide the signals for the Level-1 Calorimeter Trigger, on which the remainder of this thesis will focus. Further information can be found in [12].

2.2.4 Inner Detector

The Inner Detector is 7 m long, 3.3 m in diameter, and is divided into three sections: 2 forward sections and 1 barrel section. In the forward section the detector elements are placed radially on disks. The barrel section has detector elements arranged in concentric layers about the beam line. The transition from barrel to disk geometry occurs at $|\eta| \sim 1$. The Inner Detector provides precision vertex reconstruction and momentum measurements. It uses a combination of high spatial resolution discrete




Figure 2.6: The Inner Detector [26]

detectors and lower resolution continuous tracking for robust track reconstruction.

At full LHC luminosity, the basic functions of the inner detector are:

- The efficient reconstruction of charged particle tracks within the range $|\eta| < 2$
- Efficient momentum measurements for charged particles over the range $|\eta| < 2$
- The tagging of b-jets
- The measurement of the z-coordinate of the primary vertex of interactions

At lower luminosities, the inner detector performs the additional task of reconstructing secondary vertices from b and τ decays.

The sub-detector consists of three units: the silicon vertex detector (the Pixel Detector), silicon strip detectors (the Semi-Conductor Tracker (SCT)) and a gas straw tracker (the Transition Radiation Tracker (TRT)). These are explored in more detail in the following sections.

The Pixel Detector

The pixels are closest to the p-p interaction point and are the most precise, providing spatial resolution of $14\text{ }\mu\text{m}$ in three layers. This enables the separation of hundreds of tracks from high-energy p-p collisions and to identify whether particles originate from the primary interaction point or from a displaced secondary vertex. In the barrel, points are recorded in z and ϕ at a fixed r . In the end-caps, z is fixed and r and ϕ are recorded.

The information from each pixel is read out by a circuit which compresses the binary data and stores it whilst waiting for the Level-1 Accept. On receipt of such a signal, a module-level chip collects data for a single bunch crossing from the buffers and passes the event data to a control chip for transmission off the detector.

Due to its proximity to the beam, the pixels receive high doses of radiation: 10^{14} n/cm^2 of 1 MeV neutron equivalent and 50 Mrad of ionising radiation over the 10 years of estimated LHC operation. The closest layer receives approximately 5 times this dose and will need replacing after a few years.

The Semi-Conductor Tracker

Next out from the beamline are the silicon microstrip detectors (barrel and forward SCT) with poorer resolution, which is acceptable as track density decreases with increased distance from the interaction point. The detector is comprised of a silicon crystal semi-conductor, with embedded strips separated by $80\text{ }\mu\text{m}$ and biased with a voltage of $\sim 150\text{ V}$. Each module contains two pairs of silicon detectors read out by 6 readout chips - each reading out 128 strips. Each end-cap is supported by two cylinders and has 988 end-cap modules over 9 disks. A full disk comprises 132 modules with 40, 40 and 52 modules in three concentric rings. While the Pixel detector measures a single space-point (both z and r - ϕ with a given interaction),

the strip geometry of the SCT means that only the r - ϕ -coordinate is measured one detector layer, and then the next layer of strips, being at a stereo angle of 40 mrad, enables a lower-resolution z -coordinate to be measured.

The Transition Radiation Tracker

The number of pixel and strip layers that can be employed by the inner detector is limited due to their need for on-board electronics. As well as increasing the power dissipation and overall cost of the detector, these electronics significantly increase the thickness of the inner detector in terms of radiation lengths. For this reason a TRT is used at the outer radius of the inner detector.

The TRT is comprised of drift tubes (straws). A straw is a hollow tube of 4mm diameter. It contains a 30 μm W-Re wire through the centre across which is applied a high voltage (~ 1.6 kV) and a mixture of three gases: CF_4 increases the speed of propagation of the x-rays through the gas, Xenon captures the x-rays of the transition radiation, and CO_2 is used to stabilise the mixture in the electric field. Transition radiation occurs as charged particles traverse the interface between substances with different dielectric properties. The intensity of the emitted radiation is sensitive to the particle's energy rather than its velocity. When a particle passes through a gas in a high electric field, it will ionise the gas and generate current in a wire at the centre of the tube (or straw). In the barrel section these run parallel to the beam pipe in layers so a particle will typically pass through 34 straws. In the forward regions the straws are oriented radially. These straws have a small diameter compared to a standard gas detector due to the very high hit occupancy that would result if the straws were bigger. This detector technique is a non-destructive way of detecting electrons.

The Solenoidal Magnet

Surrounding the inner detector is the barrel cryostat containing the solenoidal coil for the 2 T inner magnet. This solenoid lies before the calorimeter and is integrated with the inner wall of the cryostat to minimise the thickness of material in front of the calorimeter. A presampler detector is installed immediately behind the cryostat cold wall to compensate for this lost energy.

2.2.5 Calorimeters




Figure 2.7: ATLAS Calorimetry [27]

The importance of the calorimeter lies in the fact that many interesting processes, including decay channels of the Higgs boson, are calorimetrically detectable in the form of isolated leptons, photons and hadronic jets, and by high missing transverse energy (E_T^{MISS}). Calorimetry is an ideal element of a detector for the extreme environment of the LHC because the energy resolution improves with higher energy. Figure 2.7 shows the ATLAS calorimetry. It comprises a cylinder of outer radius 4.25

m and total length 12.20 m. The inner barrel and end-cap sections of the calorimeter all use liquid argon (LAr) technology and are housed inside three cryostats. These are then surrounded by a scintillator tile calorimeter. Hadronic calorimetry is performed mainly within the tile calorimeter, with some being performed in the LAr endcaps.

The inner calorimeter is the ECAL which fully absorbs the lightest, electromagnetically interacting particles, γ , e^- , e^+ , and sometimes part of hadronic showers. The outer calorimeter layer is the HCAL which absorbs all remaining strongly interacting particles. Muons cause ionisation so leave a small signal in the calorimeters but are not completely absorbed. Of the known SM particles, only muons and neutrinos will escape the detector.

Electrons and photons interact with electric fields of the atoms in the detector material. Hadrons, however, interact with nuclei via the strong force, generally travelling further and making it possible to distinguish between the different shower types.

The calorimeters are made of alternating sensitive and absorbing layers, so the detected particles lose energy in the absorbing material by creating secondary particle showers which are then detected in the sensitive layers. Two methods are used for detection. LAr in a high electric field produces cascades of electrons which are detected as charge and read out electrically. The tile calorimeters use a scintillating plastic. Particles cause ionisation and a small fraction of the excitation energy emerges as visible light during de-excitation and these visible photons are detected using photomultiplier tubes. The signal size is proportional to the energy of the incident particle.

Electromagnetic Calorimeter

The ECAL provides energy measurements for electrons and photons. The calorimeter is split into three sections: one barrel $|\eta| < 1.475$ and two end-cap sections $1.375 < |\eta| < 3.2$. The end-caps are further divided into two coaxial wheels: the inner covering $1.375 < |\eta| < 2.5$ while the outer covers $2.5 < |\eta| < 3.2$. In the pseudorapidity range $|\eta| < 1.8$ the calorimeter is preceded by a presampler detector, installed immediately behind the cryostat cold wall, and used to correct for the energy lost in the material (inner detector, cryostats, coil) upstream of the calorimeter.

The total thickness is $\approx 24\chi_0$ in the barrel and $\approx 26\chi_0$ in the end caps. (The radiation length, χ_0 , is the thickness of absorber over which the electron or photon energy is reduced by a factor of $\frac{1}{e}$).

In order to be sensitive to all of the processes of interest at the LHC, good electron and photon identification and reconstruction are required over a large range of energies. For the efficient tagging of b-jets via the search for leptonic b decays, it is necessary to reconstruct electrons with transverse energies down to 2 GeV. At the other end of the scale, sensitivity to processes such as $Z' \rightarrow ee$ and $W' \rightarrow e\nu$ requires electrons with energies of up to 5 TeV to be reconstructed. A 16-bit dynamic range is required to adequately reconstruct photons and electrons within these limits. An energy resolution of $\Delta E/E \leq 10\%/\sqrt{E} + 1\%$ is thought to provide adequate sensitivity to all of the physics processes of interest at the LHC. At design luminosity and $|\eta| = 0$, the total noise contribution to the resolution of an electromagnetic shower has a root mean squared value of 400 MeV.

The barrel ECAL uses lead absorber plates in a liquid argon ionisation medium. There are 1024 absorber plates, which are of an accordion geometry and are arranged on radii around the ϕ axis. Copper readout electrodes are centred between each pair of lead plates. The barrel ECAL extends to $|\eta| = 1.4$, with electromagnetic (EM)

calorimetry at high pseudo-rapidity ranges being provided by the end-cap calorimeters. LAr is used as the detecting medium because of its good resistance against high radiation, good insulating properties and its low probability of capturing free electrons. Lead is a good absorber due to its high cross-section for bremsstrahlung and pair-production which are the two main processes through which electrons and photons lose their energy in an absorbing medium. These produce secondary particle showers which excite/ionise the LAr atoms, leading to an ionisation cascade. This ionisation is then amplified and transformed, by a dedicated pulse shaper, into an electrical signal proportional to the energy deposit.

The end-cap ECALs also use lead absorber plates with an accordion geometry. Each end-cap ECAL consists of two coaxial wheels of plates, which radiate out from the beam axis like spokes. There are 768 plates in the outer wheel, with half this number in the inner wheel.

Optimising the segmentation of the calorimeter involves balancing performance issues such as the efficiency of electron and photon identification, spatial resolution and the level of pile-up, against cost and technical constraints such as the routing of the signals from inside the cryostats. The dominant background to electrons and photons are jets. These produce showers in the ECAL which are longer and broader than those of isolated electrons and photons. The longitudinal and transverse segmentation of the calorimeter gives a measure of the shape of a shower and is therefore an essential tool for jet rejection. The segmentation of the ECAL is $\Delta\eta \times \Delta\phi \approx 0.025 \times 0.025$ in the second layer and $\Delta\eta \times \Delta\phi \approx 0.025 \times 0.05$ in the third layer, for $|\eta| < 2.4$. The remaining background is mainly due to high- P_T isolated π^0 s from jet fragmentation. To reduce this, a fine-grained, position-sensitive device is required to identify the two merged photons produced by a π^0 decay. This is achieved by segmenting the first layer of the barrel ECAL into narrow strips of $\Delta\eta \times \Delta\phi \approx 0.003 \times 0.01$, providing an integrated pre-shower detector.

Stringent requirements are placed on the ECAL by such decays as the Higgs to two photons. In the search for such rare processes, it is essential to maintain excellent acceptance. A large rapidity coverage is therefore required. Simulation studies have shown that the coverage provided by the electromagnetic calorimeter, of $|\eta| < 3.2$, allows high-efficiency reconstruction of essentially all the interesting high-transverse momentum physics channels.

To be sensitive to all processes of interest, good electron and photon identification and reconstruction are required over a large range of energies. For the efficient tagging of b-jets via the search for leptonic b-decays, it is necessary to reconstruct electrons with transverse energies down to 2 GeV.

Hadronic Calorimeter

The HCAL is situated around the ECAL since hadrons travel through the ECAL without significant loss of energy. Again, it is divided up into one central barrel section and two end-caps, but has an additional integrated forward calorimeter.

The main creation method of secondary particles in the HCAL is the strong interactions between incident particles and absorber nuclei. This calorimeter uses different technologies in the barrel and end-cap regions. The barrel ($|\eta| < 1.7$) uses iron as the absorbing medium and scintillating plates of plastic as its sensitive material, and is therefore called the tile calorimeter. The end-caps ($1.5 < |\eta| < 3.2$) use copper and LAr, similar to the electromagnetic calorimeter. There are approximately 21,000 channels in this part of the HCAL.

The major tasks of the HCAL are to identify jets and measure their energy and direction, to contribute towards the measurement of the total missing- E_T , and to enhance the particle identification capability of the ECAL by measuring such quantities as isolation and leakage. The reconstruction of both jet energy and missing- E_T

are limited by intrinsic effects. Uncertainty is introduced into the measurement of jet energy via such effects as jet fragmentation and the loss of energy into non-interacting particles (neutrinos and muons), whilst pile-up noise contributes a significant uncertainty to the measurement of missing- E_T . It is therefore acceptable for the HCAL to possess a lower energy resolution than the ECAL. It is believed that a resolution of $\Delta E/E = 50\%/\sqrt{E} + 3\%$ for the range $|\eta| < 2.4$, and $\Delta E/E = 100\%/\sqrt{E} + 10\%$ for $|\eta| > 2.4$, provides an adequate sensitivity to all of the physics processes of interest.

The search for processes such as $H \rightarrow WW \rightarrow l\nu + 2\text{jets}$ requires the reconstruction of the decay of a high- P_T W to a jet pair. This can only be performed with a well segmented HCAL. A segmentation of $\Delta\eta \times \Delta\phi \approx 0.1 \times 0.1$ is used for the range $|\eta| < 2.4$, with a coarser segmentation of $\Delta\eta \times \Delta\phi \approx 0.2 \times 0.2$ for the range $|\eta| > 2.4$. The required pseudo-rapidity coverage of the HCAL is determined by the measurement of missing- E_T . This measurement obviously requires good acceptance, so while the HCAL end-caps extend up to $|\eta| < 3.2$, an integrated ‘forward calorimeter’ increases this coverage up to $|\eta| < 4.9$.

The forward calorimeter ($3 < |\eta| < 4.9$) provides coverage in the extreme pseudo-rapidity region. This is closest to the beamline and is hence subject to an extreme particle flux, so is built using a metallic tube and rod electrode structure embedded in copper (inner section) or tungsten with a very small LAr gap. Usually in such conditions, LAr calorimeters do not work properly due to the accumulation of positive ions in the LAr gap which distort the electric field and degrade signals from the calorimeter. This is why such a narrow LAr gap is used between the electrode and absorber - 250 μm compared to the more normal gap of 1.94 mm used in the LAr barrel of the ECAL. The forward calorimeter covers the full ϕ range and $|\eta|$ up to 4.9. The forward calorimeter is used both for e/γ and isolated hadron/jet detection and measurements, and has roughly 11,000 readout channels.

2.2.6 Muon Spectrometer

The muon spectrometer surrounds the calorimeters and identifies muons which have passed through virtually unhindered. A strong toroidal magnetic field, generated by superconducting magnets, causes the charged muons to leave curved tracks, allowing calculation of their charge and momentum. Both the barrel and end-caps contain Monitored Drift Tubes for high precision coordinate measurements, complemented by Cathode Strip Chambers at high $|\eta|$. Resistive Plate Chambers in the barrel and Thin Gap Chambers in the end-caps provide the rapid response necessary to contribute to the Level-1 Trigger decision. Quick identification of muons is crucial for the trigger system in order to identify interesting processes like the Higgs decay $H \rightarrow ZZ^* \rightarrow 4\mu$.




Figure 2.8: Muon Chambers [28]

Toroidal Magnets

The size of the ATLAS detector is dictated by the toroidal magnets in the muon detector, required to bend the tracks of the muons so their momenta can be determined. They create a toroidal field running round the detector to cause the particles to bend in the direction parallel to the beam pipe. This means that momentum resolution in the muon detectors is determined by the precision measurement of z in the barrel and r in the end-caps.

There are 8 toroidal magnets in the barrel, comprising superconducting magnets 26 m long and are situated at a 10 m radius from the beamline. 8 smaller toroids are rotated in ϕ by 22.5° for the end-caps (see Figure 2.8), extending 5m from the calorimeter barrel and interleaving with the barrel toroids.

Monitored Drift Tubes

The precision momentum data are provided predominantly by the monitored drift tube detectors. The tubes are 30mm diameter, with a central $50\text{ }\mu\text{m}$ W-Re wire separated by a potential of 3 kV. A timing resolution of 300 ps over a maximum drift time of 500 ns gives an average positional accuracy in r of $80\text{ }\mu\text{m}$ per wire. Optical imaging scans monitor the alignment of the detector a few times an hour since the position of the wires must be known to within $50\text{ }\mu\text{m}$. Rigid cross plates at either end support the drift tubes, but since they are up to ~ 6 m long, mechanical deformations of the order $10\text{ }\mu\text{m}$ are expected to result from thermal, gravitational and torsion effects.

Cathode Strip Chambers

The extreme end-caps contain cathode strip chambers ($2.0 < |\eta| < 2.7$), where the predicted hit occupancy would be too high for monitored drift tubes. These are gas chambers with a 2.6 kV bias. The cathodes are separated by 5.08 mm and the anode wires run through the centre, separated by 2.54 mm. Precision measurements are obtained by measuring the charge on adjacent cathode strips which are orthogonal to the wires, producing a track resolution of $60 \mu\text{m}$.




Figure 2.9: Muon Spectrometer Subdetectors [29]

Resistive Plate Chambers

The Level-1 Trigger functionality in the barrel section is provided by resistive plate chambers. These are formed from resistive plates within gas chambers and a high (9 kV) voltage applied across a 2 mm gap. Orthogonal strips on each side of the gap, with 30-40 mm pitch, allow a capacitive readout. The typical space-time resolution

attained is $1\text{ cm} \times 1\text{ ns}$.

Thin Gap Chambers

In the end-caps, the Level-1 Trigger functionality is provided by Thin Gap Chambers. A 2.8 mm wide gas-filled chamber houses central anode wires separated by 1.8 mm. A 3.2 kV voltage is applied across the gap. At the trigger level, lower resolution is required so groups of 4-20 wires are connected together. Orthogonal readout strips, with a 15-50 mm pitch, are capacitively coupled to the cathode. They provide precision ϕ measurements to complement the r measurement from the monitored drift tubes.

Chapter 3

The ATLAS Trigger System

The ATLAS detector has approximately 147 million readout channels, and at the design luminosity of $10^{34} \text{ cm}^{-2}\text{s}^{-1}$ and interaction rate of 10^9 Hz , 1.28 MB of data are produced in every bunch crossing (BC), i.e. every 25 ns. This produces a data rate of 51.2 TB/s. It is hugely impractical to write to disk, store or later analyse this amount of data.

As described in Chapter 1, the rate of production of particles of interest, such as the Higgs Boson, is very low compared to the background produced by inelastic collisions, but such processes have identifiable signatures and so it is possible to trigger readout based on these signals and discard the majority of data before bandwidth or storage space is required.

A multi-level trigger system has been designed for ATLAS; each successive level taking more time to process more of the data in an event to reduce the event rate by a factor of 10^6 , from the LHC 40 MHz bunch-crossing frequency, to a data rate of $<100 \text{ Hz}$ ($\sim 100 \text{ MB/s}$); a more manageable amount to be written for permanent storage.

3.1 The Trigger Concept

The design goal of the ATLAS trigger is to remove inelastic collision background events whilst retaining the interesting physics. The ATLAS detector uses three layers of triggering; each one refining the selected events from the previous layer. A diagram of the trigger system is shown in Figure 3.1.




Figure 3.1: ATLAS Trigger System [30]

Only a small part of the detector is used for the Level-1 Trigger and the final decision is made by the Central Trigger Processor (CTP), based on information from reduced calorimeter granularity and the fast resistive plate and thin gap muon chambers. Using this information, the Level-1 Trigger makes its decision within $1 \mu\text{s}$, allowing a total round-trip latency (taking the cable lengths into account) of $2 \mu\text{s}$ (~ 80 bunch-crossings) and reducing the event rate from 40 MHz to ~ 75 kHz. The Level-1 Trigger also provides Region of Interest (RoI) information to Level-2, identifying areas of the detector that contain muons or high energy deposits, which may be of interest to higher level triggers.

Following a Level-1 Accept (L1A), the detector data from the selected event is stored in ReadOut Buffers (ROBs) until the Level-2 Trigger reaches a decision. If Level-2 rejects an event, the data are discarded; otherwise it is passed to the Level-3 Trigger (Event Filter). Due to a latency limit of 10 ms, Level-2 cannot investigate all of the data in the ROBs, but starts with the RoIs directly provided by Level-1. Level-2 reduces the rate from 75 kHz to ~ 1 kHz.

Following a Level-2 Accept, the Event Filter reconstructs the whole event in full granularity and resolution. It employs algorithms identical to those used offline later with improved calibration, such as cluster and jet reconstructions. This reduces the data rate to ~ 100 Hz which is then written to disk for offline analysis.

3.1.1 Latency

The time taken for a decision to be made as to whether or not an event should be accepted is known as the trigger latency. Ideally this would be smaller than the time between two bunch-crossings, but with the LHC this time is far too short since it takes ~ 1 s to fully reconstruct an event. Therefore, all data are temporarily stored (buffered) by the front-end electronics in pipeline memories. The number of channels varies for the different sub-detectors: $\sim 10^5$ for the calorimeters and $\sim 10^7$ for the inner tracker. As these numbers are extremely high, the total cost increases very rapidly with pipeline length, and so it is desirable to keep the latency to a minimum.

3.1.2 The Level-1 Trigger

The Level-1 Trigger consists of several trigger processors - each associated with a specific sub-detector, and a global processor - the CTP, to combine results and pro-

vide an overall trigger decision (see Figure 3.2). The processors are all pipelined so many events can be processed concurrently. The Level-1 Trigger only uses data from the calorimeters and muon detectors. It consists of a digital processor, taking information at reduced granularity from the LAr ECALs and the tile HCAL. Events are kept which contain high- E_T electrons, photons, taus, or jets, and events containing a large missing E_T . The Level-1 Calorimeter Trigger is explained in more detail in Section 3.2.




Figure 3.2: ATLAS Level-1 Trigger [30]

The Level-1 Muon Trigger processes information solely from the resistive plate chambers and thin gap chambers of the muon detectors, which provide digital hit information. Events containing high- E_T muon tracks which point back to the interaction point are accepted. The input comes from over 800,000 trigger channels so RoIs can be sent to Level-2 with every positive trigger decision. Both the Muon Trigger and Calorimeter Trigger have several programmable thresholds for E_T . They also have several lower thresholds, not usually used in the Level-1 Trigger decision, but used to define RoIs for Level-2.

The Level-1 CTP combines results from the Calorimeter and Muon Triggers and sends the final trigger decision to Level-2 and back to the front-end electronics.

The Level-1 Trigger receives events at the full LHC bunch-crossing rate of 40.08MHz and is required to reduce this rate to less than 100kHz for the Level-2 Trigger. This requires a high rejection rate with only one event in 10^3 being accepted.

The Level-1 Trigger applies various algorithms that search for isolated electrons and photons, hadrons, taus, jets and muons as well as summing energy deposited to calculate missing transverse energy, required to identify non-interacting particles. Since the latency has to be $< 2 \mu\text{s}$, all trigger algorithms are performed in dedicated hardware, but it is also important for the ATLAS trigger to be able to react to unexpected physics, so some flexibility is maintained by the use of Field Programmable Gate Array (FPGA) chips, and by having programmable parameters. Table 3.1 shows some signatures by which the Level-1 Trigger will identify events of interest. The thresholds are examples only, as these are programmable.

Process	Level-1 Trigger
Higgs $\rightarrow \gamma\gamma$ (80 GeV $< m_H < 130$ GeV)	2 γ 's with $P_T > 20$ GeV
Higgs $\rightarrow 4$ leptons (120 GeV $< m_H < 800$ GeV)	ee, $e\mu$ or $\mu\mu$ with $P_T > 20$ GeV for both leptons
Higgs $\rightarrow 2$ leptons + 2ν (very large m_H)	ee, $e\mu$ or $\mu\mu$ with $P_T > 20$ GeV for both leptons
$t\bar{t} \rightarrow 3$ jets + lepton	e or μ with $P_T > 40$ GeV
$t\bar{t} \rightarrow 2$ leptons	ee, $e\mu$ or $\mu\mu$ with $P_T > 20$ GeV for both leptons
W-Z pairs	ee, $e\mu$ or $\mu\mu$ with $P_T > 20$ GeV for both leptons
SUSY $\rightarrow$ jets + missing- E_T	≥ 3 jets with $P_T > 200$ GeV + missing- $E_T > 200$ GeV
SUSY cascade to leptons	ee, $e\mu$ or $\mu\mu$ with $P_T > 20$ GeV for both leptons
$Z', W' \rightarrow$ leptons	e or μ with $P_T > 40$ GeV
$Z', W' \rightarrow$ jets	2 jets with $P_T > 200$ GeV

Table 3.1: Level-1 Trigger criteria for some physics processes

These processes are expected to be visible above the large (mainly hadronic) background of the LHC but efficiencies will depend on details of the new physics. The estimated acceptance rates for each of the Level-1 Triggers are listed in Table 3.2.

Trigger	Rate
≥ 1 isolated em cluster with $P_T > 40$ GeV (isolation not required for clusters with $P_T > 65$ GeV)	31 kHz
≥ 2 isolation em clusters, each with $P_T > 20$ GeV	16 kHz
≥ 1 μ with $P_T > 20$ GeV	8 kHz
≥ 2 μ 's, each with $P_T > 20$ GeV	67 Hz
≥ 2 jets, each with $P_T > 200$ GeV	5 kHz
Total:	60 kHz

Table 3.2: Level-1 Trigger rates at full luminosity

3.1.3 The Level-2 Trigger

If an event is accepted by the Level-1 Trigger, the data for that event are read out from the pipeline memories in the front-end electronics into a derandomising buffer. The data are compressed, multiplexed (data from two BC are combined into one signal, with a flag to indicate which BC it corresponds to) and then sent to Level-2 buffer memories. In the event that either the derandomising buffer or the Level-2 memories become full, an inhibit signal is sent to the Level-1 CTP to prevent Level-1 triggers being sent while this condition persists.

The Level-2 Trigger consists of 2 main parts: feature extraction, in which candidate trigger features are parameterised using information from a localised area in one or more detectors as highlighted by the RoI information from Level-1; and global processing in which information from such features is combined and an overall trigger decision is made.

The Level-2 trigger needs to refine the Level-1 decision and reduce the event rate by a further factor of ~ 100 . It uses full calorimeter granularity to sharpen the electron, photon, jet and missing E_T cuts, and the electron trigger is refined using information

from the inner tracker so calorimeter deposits can be matched with high E_T tracks. The muon cuts can be enhanced using the precision muon chambers and the inner tracker. The Level-2 Trigger can also require events to satisfy additional signatures, such as extra leptons at low E_T or large $|\eta|$.

Following a positive Level-2 Trigger decision, the data from all the detectors are read out to the Level-3 Trigger.

3.1.4 The Level-3 Trigger

The Level-3 Trigger, or Event Filter is the last stage of the selection process in the ATLAS experiment. Level-2 and the Event Filter are globally known as High Level Triggers (HLT). While Level-2 only receives a subset of the available physical information (RoIs), the Event Filter is a set of event processing sub-farms and receives fully assembled events from the Data Collection subsystem. It uses offline physics and event reconstruction algorithms accessing the full event data and has a latency of a few seconds. The data processing and local control will be independent between sub-farms, but they will be controlled by a single global Event Filter supervisor which itself will be under the control of the experiment's online run control system [31]. The Event Filter makes a final event selection before writing data permanently to disk at a rate of ~ 100 Hz (~ 100 MB/s). The data are saved to another farm of processors of sufficient power to allow similar algorithms to be run during offline analysis with improved calibration.

Given that it has access to complete event data, the Event Filter is also able to perform global monitoring, calibration and alignment functions online, which are not possible at the detector readout level. This will be a vital element in the overall quality control of the experiment both for the physics quality and the detector optimisation and performance. The quality of the detector calibration will have a

direct bearing on the quality of the Event Filter decision itself.

3.2 Calorimeter Trigger Architecture

The ATLAS Level-1 Calorimeter Trigger is a pipelined processor designed to receive and analyse 7200 summed analogue trigger tower signals from the ECAL and HCAL. Each trigger tower corresponds to a region of the detector measuring 0.1×0.1 in $\eta \times \phi$. The analogue signals are amplified, then sent down 60m cables to the PreProcessor (PPr); see Figure 3.3. The PPr digitises and processes the signals, assigns them transverse energy values and identifies them with a specific bunch crossing. The PPr then passes the digitised data to the Cluster Processor (CP) and the Jet/Energy-sum Processor (JEP). The CP subsystem identifies isolated electron/photon and tau/hadron candidates in the data. The JEP receives data in 0.2×0.2 sums of trigger towers (jet elements), then identifies jets and sums transverse and missing-transverse energy. The CP and JEP send information to the CTP for every BC within $2 \mu\text{s}$ of the event taking place. The CTP then combines the Calorimeter and Muon Trigger results and produces an L1A trigger decision. On receipt of an L1A from the CTP, the stored input, intermediate, and output data are sent to the data acquisition system (DAQ) to be included in the data stream. The types and locations of cluster candidates and jets are passed to the RoI builder for the Level-2 Trigger [32].

3.2.1 PreProcessor

Differential signals from the ATLAS ECAL and HCAL are brought to analogue receiver stations where they are calibrated to an E_T scale before being distributed to the front panels of the PPr crates.




Figure 3.3: ATLAS Level-1 Calorimeter Trigger [32]

There are 8 standard VME64 crates to the PPr subsystem - 4 process EM trigger towers and 4 process hadronic (HAD) towers. Each crate is populated by 16 9U PreProcessor Modules (PPMs), see Figure 3.4, each of which can process 64 analogue

inputs. Initially the signals enter one of four analogue input stage daughter modules where they are converted to single-ended signals for digitisation.




Figure 3.4: The PreProcessor Module [33]

The actual signal processing is performed by 16 Multi-Chip Modules (MCM), see Figure 3.5, each of which processes 4 trigger towers using 4 10-bit FADCs with a sampling frequency of 40 MHz. A special ASIC (PHOS4) (to the left of the ADCs in Figure 3.5) allows fine adjustments to be made, providing programmable delays in steps of 1 ns across the 25 ns LHC clock period. The digitised values are then sent to a custom PreProcessor ASIC (PPrASIC).

The PPrASIC synchronises the four inputs, assigns the signals to the correct BC,




Figure 3.5: The PreProcessor Multi-Chip Module [34]

and uses a lookup table to produce calibrated 8-bit transverse energy values. It also sums the four values into 2×2 jet elements, performs ‘bunch-crossing multiplexing’ (see below) and pipelines the input data for eventual readout. Finally, two 10-bit LVDS serialisers operating at 400 Mb/s transmit the processed trigger tower data to the CP and a third serialiser sends summed 9-bit jet elements to the JEP.

A daughter card on the PPM distributes data to the CP and JEP subsystems via LVDS serial links operating at 400 Mb/s over at least 10m of shielded ‘parallel-pair’ cables. To provide the trigger algorithms with the needed overlapping data between detector quadrants, some links must be fanned out $1 \rightarrow 2$. Measures are also taken to improve signal driving capabilities, since observed signal attenuation and distortion from such long cables can compromise data integrity.

Bunch Crossing Identification

The electronic pulse caused by a particle’s interaction with the detector actually extends over several BC periods, so it is necessary to assign the signal to a BC (bunch-crossing identification, BCID).

There are three different methods whereby the PPrASIC identifies the BC associated with a signal, allowing sufficient redundancy for consistency checks.

- Normal, non-saturated signals (below ~ 250 GeV) are processed by a pipelined Finite Impulse Response (FIR) filter. Five consecutive FADC samples are taken, multiplied by pre-defined coefficients, and the resulting values are summed. A subsequent ‘peak-finder’ attributes the maximum value of the sum to the corresponding BC. The working range for this method extends from small signals of a few GeV up to the near-saturation level of 250 GeV.
- For saturated signals, two consecutive samples are compared to a ‘low’ and a ‘high’ threshold, making use of the finite ‘peaking-time’ (50 ns) of an analogue input signal. Detection of a ‘leading edge’ allows a virtual peak to be attributed to a specific BC. This method is valid from ~ 200 GeV through to the maximum energy extent of the calorimeter.
- Thirdly, and mainly for consistency checks, comparators are used with programmable thresholds on the analogue input daughter modules to present a ‘rising edge’ signal to the MCM. Given the known peaking time, BCID can be performed using an appropriate programmed delay in the PPrASIC. This method is valid for much of the range of the previous two, allowing plenty of overlap for consistency checks during setup.

The output from the FIR filter is presented to a lookup table to obtain a calibrated E_T value for the tower. If the BCID criteria are met, this value is sent to the CP and JEP outputs. Saturated signals are assigned the maximum 8-bit value (255 GeV). A 2×2 jet element with a sum that overflows, or which contains a saturated tower, is also assigned the maximum 9-bit value (511 GeV). All data words are accompanied by an odd-parity bit for error detection. Odd-parity is where the 8-bits of data, plus the BCmux flag, plus the parity bit (10-bits in total) contain an odd number of set bits.

Bunch-Crossing Multiplexing

The MCM outputs are sent to a daughter card on the PPM via 10-bit LVDS serial links for electrical fanout to the CP and JEP. The four trigger towers from each MCM are sent to the CP, and a 9-bit sum of these towers (jet element) is sent to the JEP. The ‘peak-finding’ nature of the BCID algorithm means it will always follow an occupied BC with an empty (zero) one, allowing two concurrent trigger towers to share a single serial link (bunch-crossing multiplexing, BC-mux).

Towers are paired in the PPrASIC output stage. If a tower has a non-zero value, this value is sent to the link along with a flag bit to indicate which of the two towers is being transmitted first. In the next LHC clock cycle, an E_T value (or zero) for the other tower is sent to the link, again with a flag bit. For this clock cycle, the flag bit is used to indicate whether the second tower’s value belongs to the same BC as that of the first tower, or to the following BC.

By using BC-mux, data for the CP subsystem can be achieved with only two links per MCM instead of four. This is not possible for the the summed 4-tower jet elements as there is no assurance that such a sum will be followed by a zero in the next BC.

3.2.2 Jet/Energy-Sum Processor

The JEP identifies jets and sums the total transverse energy, E_T^{tot} . Each Jet/Energy-Sum Module (JEM) receives and deserialises data from 88 LVDS links, corresponding to 4×11 jet elements in $\eta \times \phi$, for both electromagnetic and hadronic, see Figure 3.6. Four input FPGAs (XC2V1500) receive data at 40 MHz and sum the electromagnetic and hadronic parts of each jet element to 10-bit values, then multiplex these sums to 80 MHz for fan-out across the board and over the backplane. Backplane and on-

board transmission of data to the main processor units are carried out at 80 MHz. Pipelines in each input FPGA save input data for readout to the DAQ system upon an L1A [35].




Figure 3.6: The Jet/Energy-Sum Processor Module [34]

The jet algorithm uses coarser granularity, identifying transverse energy sums within overlapping windows consisting of 2×2 , 3×3 or 4×4 jet elements (corresponding to sizes 0.4, 0.6 or 0.8 in $\eta \times \phi$). The size of blocks is programmable to provide flexibility. If a jet is detected whose energy exceeds one of eight programmable thresholds, it is recorded. Multiple counting of jet candidates is avoided by requiring the window to surround a 2×2 cluster whose sum is a local maximum. The location of this 2×2 cluster also defines the coordinates of the jet RoI. All three jet window sizes are calculated by default, and eight independent combinations of energy threshold are available for trigger menus.

The energy-summation algorithm produces unsigned sums of E_T , E_x and E_y and uses the system-level sums of these to report total and vector E_T values to the CTP. This is possible because the location of quadrants in the detector and which signals correspond to where, is known.

The jet and energy-summation algorithms are implemented in one large ‘main processor’ FPGA (XCV2000) per JEM [36]. The main processor is also responsible for reporting results to the crate-level mergers (see Section 3.2.4), as well as pipelining of DAQ and RoI information for readout. The jet and energy outputs of each JEM are both 25-bit data streams; the former consisting of 8 3-bit jet multiplicities and one odd-parity bit, and the latter containing the unsigned values of E_T , E_x and E_y , each compressed from 12 bits to an 8-bit ‘quad-linear’ scale (6-bit mantissa plus two multiplier bits). The energy output is also accompanied by an odd-parity bit.

A single ReadOut Controller (ROC) FPGA collects input data from the input FPGAs and output and RoI data from the main processor FPGA for readout to DAQ and Level-2.

3.2.3 Cluster Processor

The CP is described in detail in the next chapter. It is designed to identify transverse energy clusters associated with electron/photon and isolated hadron candidates. Algorithms based on the trigger tower energies have been designed to identify such clusters. The CPM is shown in Figure 3.7.

3.2.4 Common Merger Module

Two modules in each CP and JEP crate carry out crate-level merging of results received from the crate’s processor modules. In the CP crates, each of the two




Figure 3.7: Photo of the CPM [34]

merger modules is responsible for calculating 3-bit cluster multiplicities for 8 of the 16 e/γ and hadron/ τ cluster definitions. In the JEP crates, one merger module produces 3-bit multiplicities for the 8 jet definitions, while the other produces sums of E_T , E_x and E_y . A Common Merger Module (CMM) has been developed for all merging functions, using a large FPGA with multiple firmware configurations to do crate-level merging on up to 400 bits of data per BC from the crate's processor modules [37], see Figure 3.8.

The CMM also performs system-level merging. Parallel LVDS cable links between the subsystem crates bring the crate-level results to one CMM of each type, which is designated the system-merger. A second large FPGA on the CMM carries out the system level merging and reports final results to the CTP.

The FPGAs on the CMM also store input and output data into pipelines for readout to DAQ upon an L1A.




Figure 3.8: The Common Merger Module [34]

3.2.5 ReadOut Driver

There are 2 separate readout systems for the Level-1 Calorimeter Trigger:

- Input, output and some intermediate data from each module are read out to the DAQ system;
- The CP and JEP subsystems report feature types and coordinates of RoI data to the Level-2 Trigger.

The readout system is designed to report 1 BC of RoI data, and up to 5 BCs of DAQ data per event at an L1A rate of up to 75 kHz. A common, FPGA-based ReadOut Driver (ROD) is used by all of the subsystems, see Figure 3.9.

Readout FIFOs on each processor FPGA or ASIC are read out as 40 MHz serial streams to a ROC for timing alignment, which is then passed in parallel to inputs on




Figure 3.9: The ReadOut Driver [34]

a 20-bit G-link transmitter (Agilent HDMP 1022) [38]. The G-link then transmits them at 800 Mb/s to a ROD.

The ROD is a 9U module residing in a standard VME64 crate. It has 18 G-link receivers (Agilent HDMP 1024) which pass the parallel outputs to the FPGA for data compression, zero suppression, and some data monitoring. A common ROD module is used by both the DAQ and RoI readout subsystems to gather and report data from the PPr, CP and JEP subsystems, using a different firmware configuration in a large FPGA for each readout task [39]. The ROD also contains four S-link transmitters [40] for passing compressed event data to the DAQ and RoI readout buffers. The number of G-link and S-link outputs used, and the mapping between them, depends on the type and source of data being read out.

3.2.6 Processor Backplane

A common, custom processor backplane is used by the CP and JEP subsystems. It is 9U high, and can host up to 21 single-width modules. Each position in the backplane is intended for a single kind of module. There are 16 CPM/JEM positions,

flanked by 2 CMM positions and a further 2 for a VME CPU (mounted in a 9U adaptor module) and a Timing and Control Module (TCM), see Figure 3.10.




Figure 3.10: Photo of a system crate containing (from left to right) CPU, CMM, JEM, 3 CPMs, another CMM and a TCM [41]

The backplane is almost entirely populated with 2mm HM connectors with 1148 signal and ground pins in each JEM/CPM and CMM position. DC power is provided via 3 high-current DIN pins at the bottom of each module. The backplane has 8 signal layers sandwiched between 10 ground planes, providing good impedance control and low cross-talk [42]. Two signal layers provide point-to-point links between neighbouring processor modules for input data fan in/out. A Controller Area Network bus (CANbus) on one of these layers provides an interface to the detector control system (DCS). The CANbus is a high-integrity serial data communications bus used for crate and module monitoring. Another four layers provide diagonal connections from outputs at the top and bottom of each CPM/JEM to the two merger modules at the right and left of the processor modules. The final two layers distribute TTC signals from the TCM.



Figure 3.11: The Processor Backplane - front view where the modules plug in inside crate (left) and the back where cables plug in (right) [34]

The LVDS serial link and merger interconnect cables are connected to the rear side of the backplane and passed through it to the front panels of the modules. The crate also provides extended geographic addressing. Upon power-up, address pins tell each module its location in the crate, and the crate's location in the system. The modules use this information to set unique VME, CAN and TTC addresses. Modules with multiple FPGA configurations can also use this information to automatically load the appropriate configuration.

Chapter 4

The Level-1 Calorimeter Trigger Cluster Processor

The ATLAS Level-1 Calorimeter Trigger consists of three main data-processing components, as detailed in the previous chapter. This chapter examines the CPM which has been developed jointly by the University of Birmingham and Rutherford Appleton Laboratory (RAL).

The CP system looks for e/γ and τ candidates within the coverage of the ATLAS tracking system ($|\eta| < 2.5$). Its input signals are formed by analogue summation of calorimeter cells to form ‘trigger towers’ of granularity $\Delta\eta \times \Delta\phi \sim 0.1 \times 0.1$, separately in the ATLAS electromagnetic and hadronic calorimeters in this η range. The Preprocessor subsystem receives 64×50 ($\phi \times \eta$) trigger tower signals from the LAr electromagnetic calorimeters, and a matching array of projective trigger towers of the same granularity from the hadronic calorimeters (TileCal and HEC). The signals are transformed from energy to transverse energy, and digitised to 8-bit precision before being input to the trigger hardware. All trigger tower signals from every bunch-crossing are supplied by the PPr synchronously to the CP at the 40

MHz LHC bunch-crossing rate via dedicated trigger cables [43].

4.1 Overview of the Cluster Processor

The Cluster Processor divides the trigger space into four quadrants in ϕ , each of which is processed by a single Cluster Processor crate containing 14 CPMs. A CPM processes 64 windows arranged logically in a 16×4 ($\phi \times \eta$) array. A crate of 14 CPMs is therefore sufficient to process 56 trigger windows in η , while data will only be available from the 50 trigger towers spanning the pseudo-rapidity range of $|\eta| < 2.5$. Most CPMs process 16×4 trigger towers in $\phi \times \eta$, so to fully process the full array of trigger towers, the CPM requires an array of 20×7 , as shown in Figure 4.1, to allow isolation rings to be included for every trigger tower. The CPMs at both ends of each crate will have a number of unused windows and those that are used will have only some of their input towers populated.

4.1.1 The Cluster Processor Module

The CPM is the module on which the electron/photon and tau/hadron algorithms are executed. It is a complex 16-layer (400 mm deep) 9U (366 mm high) board, with fine 0.1 mm printed tracks, and with 2000 resistors and capacitors. It comprises (see Figure 4.2):

- 80 400 Mbit/s LVDS links, to collect data from Preprocessor Modules
- 80 LVDS deserialisers, to convert data to 40 MHz 10-bit parallel words
- 20 Serialiser chips, to distribute data at 160 MHz: on-board to perform cluster finding algorithm, via backplane to adjacent modules.




Figure 4.1: Organisation of Level-1 Calorimeter Trigger into CPMs and FPGAs [44]

- Receive data at 160 MHz from neighbouring modules
- 8 CP chips, to perform the Cluster Finding Algorithm
- 2 Hit Merger chips, to calculate and transmit multiplicities to Level-1 Trigger decision via CMM
- 2 ROC chips, to pipeline Trigger Tower Data, multiplicities and RoI coordinates
- On Level-1 request, to send readout information to ROD module to help build Level-2 decision and read out to DAQ

The main requirements of the the CPM are to:

- identify possible isolated electrons, photons and hadronic τ decays.




Figure 4.2: Schematic layout of the CPM [45]

- calculate multiplicities of e/γ candidates and τ candidates for different E_T thresholds.
- transmit these to the Common Merger Module (CMM) for Level-1 trigger decision.
- transmit trigger tower data, multiplicities, and RoI coordinates to RODs.

Each CPM receives and deserialises data from 80 LVDS links, corresponding to 4×20 trigger towers in η and ϕ , both electromagnetic and hadronic. Twenty Serialiser FPGAs (Xilinx XCV100E) receive the data and serialise them to 160 MHz. The data are then shared between neighbouring modules via the backplane, and finally fanned out to eight CP FPGAs (Xilinx XCV1000E) which perform the cluster algorithms. The Serialiser FPGAs also store input in pipelines for eventual readout to DAQ upon receipt of a L1A.

The eight CP FPGAs are implemented in large devices, each of which services 8 overlapping 4×4 windows. Pipelines implemented in each CP FPGA save output data for readout to the DAQ system, and cluster types and coordinates for readout as RoIs to Level-2.

Two ‘hit multiplicity’ FPGAs collect and sum the cluster multiplicities from the CP FPGAs for reporting to the crate-level merging of Level-1 results. The multiplicities are reported on two 25-bit data streams, each containing 3-bit multiplicities for 8 of the 16 cluster types, plus one parity bit. If more than 7 instances of a cluster type are identified, the multiplicity is reported as 7.

Two additional ROC FPGAs collect input data from the Serialiser FPGAs, RoI data from the CP FPGAs and output data from the hit multiplicity FPGAs upon a L1A, and transmit them to DAQ and Level-2.

4.2 FPGAs

The Trigger Algorithms are carried out on FPGAs. An FPGA is a user-programmable logic device which can be configured for a specific application by uploading configuration data to an on-chip static RAM.

The main blocks of an FPGA are:

- Configurable Logic Blocks (CLBs) which perform the various logical operations.
- Input and Output Blocks (IOBs) which form the interface between the CLBs and the external pins.
- The routing network which processes global signals like the clock and reset, and dynamically connects the IOBs and CLBs.

The configurable logic blocks provide the functional elements from which the user-specified logic is constructed and are arranged in a matrix at the centre of the chip. They contain Look-Up Tables (LUTs) to perform combinatorial logic functions (adders, comparators etc.) and storage elements in the form of flip-flops. The data path through a logic block is controlled by multiplexers established during configuration.

Like ASIC chips, FPGAs are designed to carry out numerous tasks, in parallel, during one clock cycle. The advantage FPGAs have over ASICs is the programmability. Once an ASIC is produced, its functionality cannot be altered, and the chip can only be used to perform its original tasks. The functionality of an FPGA can change with time.

Programming an FPGA links the blocks in such a way that the intended task can be completed using the wired logic. Once a program is loaded, the FPGA retains this information until the power is cut or a reset signal is sent.

Signals are transmitted between logic blocks within the FPGA via three types of configurable interconnections:

- General purpose connections form a grid surrounding the logic blocks. Intersection points on the grid route signals according to the downloaded configuration.
- Long line connections also form a grid around the logic blocks, but they bypass the switching matrices. They are intended for use with high fan-out signals which must travel long distances with a minimum propagation delay. Additional long lines driven by a global buffer are provided for the distribution of a clock signal.
- Direct connections provide a high-speed link between adjacent logic blocks.

The input and output connections of each logic block to the direct connection network, the long lines, and the local segments of the general purpose network are controlled by switches on that logic block which are established during configuration.

Signals enter and exit the FPGA via the I/O blocks which provide an interface between the internal logic and the device package pins. Each I/O block contains both input and output buffers and the downloaded configuration determines not only whether a pin is to act as an input or output, but also whether the voltage thresholds will be TTL or CMOS compatible. Before configuration, all I/O levels are TTL compatible. The FPGAs on the CPM are set to be CMOS compatible.

The FPGA can be configured either at power-up, or on command. A variety of modes exist for the transfer of the configuration data, and establishing a 3-bit pattern at the ‘mode pins’ of the FPGA informs the device of which mode pattern is to be used. In ‘Master Parallel’ mode, for example, the FPGA retrieves the data from an external memory, supplying both the memory addresses and the clock signal itself (using an internal crystal oscillator). Alternatively, in ‘Peripheral’ mode, a microprocessor can be used to supply both the configuration data and a clock signal to the FPGA, which provides a READY/BUSY signal as a handshake.

The bitstream itself is generated by proprietary software, and is the end result of producing a logic design for an FPGA (see Chapter 5). The length of the configuration bit-stream is independent of the complexity of the logic design and depends only upon the size of the device to be configured.

4.3 Serialiser FPGA

The Serialiser FPGA converts the parallel data received from a link de-serialiser chip to 160 Mbit/s data streams for the CP FPGA and for module-to-module data transfer via the backplane. It also captures the received data and transfers it to the DAQ via the ROC.

The CP receives serial data from the PPM using LVDS deserialiser chips (10 bits/chip every 25 ns), as shown in Figure 4.3. The Serialiser FPGA receives data from four BC-multiplexed channels, a total of 40 bits every 25 ns. The 40 bits are split into two 20-bit fields, each of which is serialised on to five 160 Mbit/s serial links in four-bit nibbles. The 20 bit data field consists of two 8-bit trigger towers and their corresponding flag bits and error checking bits. The serialiser FPGA converts the two 8-bit words on to four 160 Mbit/s serial lines and the other four bits (two error bits and two flag bits) on to another 160 Mbit/s serial line.




Figure 4.3: Data Chain through the CPM [46]

The functionality of the Serialiser FPGA is illustrated in Figure 4.4.




Figure 4.4: Block diagram of Serialiser functionality [46]

4.3.1 Real Time Data Path

The Serialiser FPGA receives data from four LVDS deserialiser chips. Each chip provides 10 bits of parallel data and a recovered 40 MHz clock, which is used by the Serialiser FPGA to clock in the data. Since two 10-bit words are serialised on to five 160 Mbit/s serial lines, the data from a pair of channels must arrive within an

acceptable time window. Due to the timing spread between the serialiser-deserialiser chip-sets, additional timing adjustment may be required, so a flipflop (* in Figure 4.5) provides this adjustment by allowing early signals to be delayed by one clock tick.

After the data have been synchronised, the parallel-to-serial converter takes a 4-bit nibble of data and serialises it to 160 Mbit/s, and provides a 4-way fan-out. The data which pass through the real-time path are also copied to a dual port RAM which can be readout to the DAQ.




Figure 4.5: Serialiser data capture and synchronisation logic [46]

4.3.2 Readout Logic

The input data are extracted from the real-time data path after synchronisation, as shown in Figure 4.4. These data are BC-demultiplexed and written to a dual-port RAM (DPR), which also captures the status of the link and parity error flags associated with the data. The data for one bunch crossing, after BC-demultiplexing, is referred to as a ‘slice’. The DPR holds a history of data for 128 bunch crossings before being over-written. This allows for the latency of the trigger system and for a L1A decision to be made.

4.4 Cluster Processor FPGA

The most complex part of the logic required by the CP is implemented on the CP FPGA. The CP chip will process a 2×4 region from EM and HAD calorimeters, see Figure 4.6.




Figure 4.6: CP chip processing window [47]

The lateral dimensions of electromagnetic showers in the calorimeters are of the order of a few centimetres. Electrons and photons entering a trigger tower deposit most of their energy in one trigger tower unless they hit the join between two towers, in which case the energy will be spread over them both. The cluster energy is based on the sum of two adjacent trigger towers. It would be even more accurate to sum four trigger towers, but this decreases jet rejection.

The algorithm uses a 4×4 sliding window to search for isolated EM energy clusters and tau candidates, and provide triggers and RoIs, see Figure 4.7. The CP chip applies this sliding window in steps of one trigger tower in both η and ϕ directions to trigger towers 11, 12, 13, 14, 21, 22, 23 and 24, as illustrated in Figure 4.6.

The functional specification of the CP chip is illustrated in a block diagram in Figure 4.8.




Figure 4.7: Arrangement of trigger towers used for CP algorithm [47]

4.4.1 Trigger Algorithms

The cluster processor is designed to find electromagnetic clusters, from electrons or photons, and isolated hadrons, many of which will come from taus. Using the transverse energies of the trigger towers within the trigger window, we obtain [44]:

- EM clusters: Sum of 2 adjacent of the 4 central EM trigger towers.
- EM isolation sum: Sum of the 12 ECAL isolation ring trigger towers.
- HAD clusters: Sum as EM cluster, plus the sum of the 4 corresponding central HAD trigger towers.
- HAD isolation sum: Sum of the 12 HCAL isolation ring trigger towers.
- Central HAD isolation sum: Sum of the 4 HCAL central trigger towers.
- Central RoI: Sum of the 8 central trigger towers (4 EM, 4 HAD).
- Peripheral RoI's: Sums of the additional 2x2 arrays of EM and HAD trigger towers within the trigger window (8 combinations).




Figure 4.8: CP chip block diagram [47]

The EM cluster trigger and isolated HAD trigger are formed from these trigger elements. The CPM identifies and reports clusters satisfying 8-16 e/γ and 0-8 $hadron/\tau$ criteria, for a total of 16 independent cluster types.

Electromagnetic cluster trigger:

The electron/photon cluster algorithm identifies 2×2 clusters of trigger towers with at least one two-tower sum (1×2 or 2×1) of nearest neighbouring EM towers ex-

ceeding a predefined threshold. Isolation thresholds are set for the 4×4 surrounding ring in the ECAL, as well as the 2×2 hadronic towers sum behind the cluster, and the 12-tower hadronic ring behind it.

- At least one EM cluster must exceed a threshold.
- The EM, HAD and central HAD isolation sums must all be below specific thresholds.
- Central RoI must be a local transverse energy maximum, ie. at least as energetic as the peripheral RoIs.

Hadronic cluster trigger:

The hadron/tau algorithm is similar to the electron/photon algorithm; nearest-neighbour sums of EM plus HAD trigger towers are compared with an energy threshold, and separate isolation thresholds are set for the surrounding 12-tower rings in the ECAL and HCAL. Multiple counting of clusters is avoided by requiring the central 2×2 region of the cluster to be a local maximum.

- At least one HAD cluster must exceed a threshold.
- EM and HAD isolation sums must be below specific thresholds.
- Central RoI must be a local transverse energy maximum.

Chapter 5

Development Environment

There are two aspects to the development environment used; firstly, there are the software tools used to modify firmware for the FPGAs, and secondly there is the hardware environment in which tests on the chips could be carried out. Both of these areas are briefly described in this chapter.

5.1 Software Tools

Developing a logic design for an FPGA is a process with two distinct stages. The first is to specify the functionality of the logic design, which may be done using either schematic capture, a hardware design language, or a combination of these two methods. The design at this stage may be kept independent of the technology on which it is to be implemented, or such device-specific features as I/O pins may be allocated via the schematic or language file.

Simulation may be performed before progressing further to ensure that the logic specified is functionally correct.

Once this has been completed, the next stage is to target the logic design towards FPGA technology. For most designs, this is a largely automated process. Software supplied by the FPGA manufacturer can target the design towards a specific FPGA, or the smallest FPGA in a given family which will accommodate the specified logic. The routing of designs is optimised to maximise the speed of the logic and a configuration bit-stream is produced.

At this stage the design can be simulated once again, to verify the functionality and the timing of the logic.

5.1.1 HDL Designer

The HDL Designer Series is a family of tools for electronic system design. Designs can be imported, and the tool automatically creates editable graphic diagrams. The graphics tools also allow block diagrams, state diagrams, flow charts and symbol views to be used.

It enables a design to be created schematically, using block diagrams and flow charts. The design is also coded at a lower level within these blocks, down to detailed functionality, using a hardware design language, in this case VHDL. The graphical design enables complicated designs to be created and edited, see Figure 5.1.

The individual component can be opened up, and the corresponding block of code is revealed, built by the software. It included all inputs that had been put on the graphical diagram, and main structure required. The functional code can then added within this component.

An additional testbench can also written, to provide inputs for the main code in order to test it. The testbench is not implemented on the chip, and therefore the code is not constrained in the same way. It can be given wait periods which cannot be expressed




Figure 5.1: The HDL Designer control window and two design block-diagram windows

in hardware. HDL Designer includes a simulation analyser interface supporting error cross-referencing and animation facilities to assist with design debug operations. It interfaces with ModelSim.

5.1.2 ModelSim

Once the main code and testbench have both been written, they can be loaded into ModelSim. ModelSim is a VHDL simulator with full debugging capabilities. It enables functional and timing models of the design to be verified, see 5.2. ModelSim has a graphical user interface (GUI) but commands can also be passed by the command line.

You can override ModelSim's default resolution by specifying the `-t` option on the command line or by selecting a different Simulator Resolution in the Simulate dialog box. Available resolutions are: $1\times$, $10\times$, or $100\times$ of fs, ps, ns, μ s, ms, or sec. Clearly you need to be careful when doing this type of operation. If the resolution set by `-t` is larger than a delay value in your design, the delay values in that design unit are rounded to the closest multiple of the resolution.

Waveform Analysis

When the simulation finishes, it is possible to analyse waveforms to assess and debug the design using the Wave window, see Figure 5.3. Numerous commands give advanced functionality, including setting pointers, zooming the resolution, searching along particular signals for its next event, and formatting between binary, decimal and hexadecimal displays.




Figure 5.3: ModelSim Wave Window [48]

It is also possible to look at waveform data in a textual format in the List window,

see Figure 5.4. The List window displays simulation results in tabular format. It is useful for gating expressions and trigger settings to focus in on particular signals or events and for debugging delay issues. The window is divided into two adjustable panes, which allows you to scroll horizontally through the listing on the right, while keeping time and delta visible on the left.




Figure 5.4: ModelSim List Window [48]

5.1.3 Leonardo Spectrum and Xilinx

At this stage, however, ModelSim does not tell a full story, but merely gives an ideal view of how the code runs. Leonardo Spectrum allows the design to be synthesised for a specific chip, taking into account timing delays due to routing within the chip, and the speed of the chip itself. The code is compiled into a netlist corresponding to the design using logical combinations. It is then mapped onto specific CLBs of a particular FPGA type using a Place and Route (PAR) tool.

Further simulations can be subsequently run in ModelSim which will now reveal the code's behaviour in a manner specific to the chip. Once the code has been fully tested, the output file from the PAR tool can be compiled into a bit stream, using Xilinx software supplied by the FPGA manufacturer. The routing of designs is optimised to maximise the speed of the logic and a configuration bit-stream is

produced, which can be loaded into the FPGA.

5.2 Hardware

Numerous system tests have been carried out on the CPM, including those at the H8 CERN Testbeam, which are later discussed in Chapter 7 and integration tests carried out at RAL. The majority of CPM-specific tests were carried out at The University of Birmingham. Tests make use of the Data Sink and Source (DSS) signal board, which is a carrier of daughter boards. Daughter boards can be chosen to generate CMOS signals, to be a Glink receiver, or to emulate Low-Voltage Differential Signalling (LVDS) Tx signals from the PPM [49]. Later the LVDS Source Modules (LSM) were designed and built which emulate a full PPM, see Section 5.2.2. A typical test rig at Birmingham consisted of:

One 6U crate with:

- TTCvi and TTCex to run at 40.08MHz
- Electrical TTC fan out for DSS and LSM
- 1 DSS with GIO source card, to generate L1A signals
- 1 DSS with Glink Rx
- 2 LSMs

One 9U crate with:

- 3 CPMs
- 1 CMM

- 1 TCM

The system used at RAL was more complex, integrating the CPM with downstream modules such as RODs and running at high L1A rates.

5.2.1 DSS Module

DSS modules are required to generate the test patterns for testing links (G-links, LVDS, and others) and to test a range of trigger modules. Also the DSS module can be also used for testing inputs, such as the ASICs and MCMs. It is a general-purpose motherboard designed to meet all the data source and sink functions required for testing the PPMs, CPMs and JEMs, with the requirement of only having to design specific daughter boards for various applications.

The DSS motherboard is a 6U VME and implements common functions such as generate, receive and compare test patterns in real time up to 40MHz; VME interface, control and status registers, etc. The daughter boards implement the physical layer to transmit/receive data up to 80 bits wide and can configure the mother board to perform source (transmit) or sink (receive) functions.

The CMC card enables translation between CMOS and LVDS signals and can be used as an LVDS source to drive data into the CPM test system [50]. They can be used to send a known stream of data, or can be set up with a Pseudo-random Data Generator, to send random data for testing [51]. Each DSS can provide 16 LVDS sources.

In order to feed a full CPM, as developed with later testing, it was necessary to use five DSS modules, see Figure 5.5. The custom backplane receives the LVDS data from the rear, providing flexibility as the modules can be removed without uncabling.




Figure 5.5: Five DSS modules feeding LVDS data to one CPM [52]

This limits testing of CPMs as they require high numbers of DSS modules to drive them. The LSM was developed to provide this capability and alleviate the pressures on DSS requirements.

5.2.2 LSM Module

The LSM provides 480Mb/s LVDS serial data signals for testing CPM and JEM processor modules. It is similar in function to the DSS/LVDS combination, but drives far more LVDS links by using FPGAs for the data source and serialisation. The serial format matches that of existing LVDS links [54].

The LSM uses the TTCdec daughter card to provide the master clocks for the module and is shown in Figure 5.6.




Figure 5.6: LVDS Source Module [53]

Chapter 6

CPM Testing

A common approach has been used in testing different areas of the board or to validate data sent downstream to the Level-1 Trigger system chain. A detailed simulation package is able to generate expected data at different points inside the trigger system, down to the level of each individual chip component. It is necessary to select the type of data to process, and then test vectors are generated and loaded inside available playback memory on the system. The simulation uses the same test vectors, and the computed output is compared to the hardware data uploaded from spy memories to check for errors, see Figure 6.1.

6.1 Overview of CPM Tests

Tests have been run to test all of the chips on the CPM, including LVDS inputs to the Serialiser chips, Real Time Data tests of the CP chips, Hit Multiplicity outputs of the hit merger chips read out to the CMM, and Readout Frame Definition from the ROC chips read out to the ROD. Many tests require modifications to the FPGA firmware, either to resolve issues that have been identified within the chip, to compensate for




Figure 6.1: Schematic Layout of Test Procedure [52]

issues identified elsewhere (e.g. in other modules), or just to include extra temporary logic to capture data for specific tests. Modifications made to the Serialiser and CP chip firmware fulfilling each of these requirements are further detailed in this chapter.

6.2 Serialiser Chip Tests

6.2.1 Real Time data tests

The CPM receives its data from the PPM in LVDS format. Two LSMs were used to source two CPMs. Playback memories of the LSM are synchronously started on a TTC command, allowing the CP algorithm to be tested between two adjacent CPMs. Cables of 12 m were connected to the backplane. Once the data were recovered correctly inside the spy memory of the serialiser chip, a timing investigation of the input data was performed. It consisted of delaying the 40 MHz clock strobing the incoming data in steps of 104 ps, the minimal value provided by the TTCrx chip. The results are shown in Figure 6.2.




Figure 6.2: Error profile of data received inside the Serialiser as a function of the delayed 40 MHz clock [55]

Each line plots the cumulative errors incurred per chip as a function of the timing offset. It can be seen that errors begin to occur when the offset is approximately 11 ns, and continue to occur across a window of about 5 ns. When the offset is more than approximately 16 ns, the data return to being error free. At the level of each individual chip, the transition is around 2 ns, as can be seen more clearly in Figure 6.3, but the overall spread of the tracks length expand it up to 5 ns.

The large spread (5 ns) compared to the individual error of each chip (2 ns) is largely due to the group of Serialisers closest to the backplane having their errors occurring earlier than those in the next row, as can be seen in Figure 6.3. Each patch on the graph indicates the spread of errors on one Serialiser chip. Those towards the left, around TTC-value 170 (every alternate patch) represent the errors on the first row of chips on the board, i.e. those closest to the backplane, while those around




Figure 6.3: Errors by Serialiser of data received as a function of the delayed 40 MHz clock [55]

TTC-value 190 are in the second row.

This plot is taken from the latest version of the CPM (Version 1.9). The previous version showed a much larger spread between when the errors occurred for the first and second rows of Serialiser chips, which significantly reduced the error-free window. This led to the identification of the impact of track lengths on timing windows and the new version of the PCB made efforts to keep track lengths similar for parallel data.

6.2.2 BER Tests

As detailed in the previous chapter, much testing has been undertaken to ensure the CPM will perform up to required specifications. The signals arriving via the LVDS cables require a Bit-Error Rate of $< 10^{-13}$.

The serial links from the PPr to the CP use LVDS over shielded (un-twisted) pair cables. LVDS is a high-speed (~ 500 Mb/s), low-power, general-purpose interface standard. It features a low voltage swing of ± 400 mV with power-supply voltages ranging from 2.7 V to 5.0 V [56].

Each pair of trigger towers is described by a 10-bit word for each bunch-crossing, of which 8 bits are the calibrated E_T value for one of the towers, 1 bit is a flag indicating which tower, and the final bit gives odd parity encoded over those 9. Each link carries these 10 bits every bunch-crossing: an effective bandwidth of 400 Mbit/s per link.

Each CPM hosts 80 links from the PPr, 40 carrying EM and 40 carrying HAD trigger tower data. The serial links are combined in cable assemblies of 4 links, then the CPM receives 4 such cables from each of 2 PPMs (one EM, one HAD) covering the same quadrant. Another cable brings shared data from a further 4 PPMs in the 2 adjacent quadrants. An error rate of $< 10^{-13}$ is required over these cables. A test was devised to carry out this objective.

Source code was written in Very High Speed Integrated Circuit (VHSIC) Hardware Description Language (VHDL). Software tools were used to design, test, debug and synthesise the design, as detailed in the previous chapter.

The Serialiser Code was opened up in HDL Designer where the graphical form could be modified to include an error_comparator (blue) block, see Figure 6.4. Within this block, the source code was written to compare data being received over one of the

channels with data that were stored in the DPR. The memory allocated within the Serialiser was 128 words long, so a pattern of this length could be tested.




Figure 6.4: Schematic Layout of the Serialiser

Pseudo-random Data Generator

Since a pattern length of 128 words does not cover all possible 10-bit words, it is unsatisfactory to just loop these data for a run as long as 8 hours to demonstrate a high bit-error rate. Instead, the code was modified to include a Linear Feedback Shift Register (LFSR) to generate pseudo-random data [57].

A piece of code would continually generate a new pseudo-random word, without relying on the 128 word RAM. The code was triggered by the clock, so a new word was produced on every clock tick to compare with the data being received from the

DSS. This was exactly the same way that data were produced in the DSS, so by starting with the same seed the exact same data could be generated.

Synchronisation

One of the main problems was that of synchronisation. Separate resets had to be sent to the serialiser chips and the DSS which results in the data from each not necessarily starting to flow at the same time. A pointer was held at the top of the memory, waiting for a matching value from the DSS. Once a match was found, a ‘lock’ went high and the data from the RAM were allowed to flow, and each pair of values was compared. In the event they did not match, an error counter was incremented.

This was not a complete solution because the LFSR could generate the same number more than once, and so a false lock may be established, after which none of the data would actually match. As this was clearly not a satisfactory solution, the `error_comparator` code was modified to check for two consecutive matches before locking and counting errors. This produced a far more robust solution.

Figure 6.5 illustrates the signals in ModelSim. When a reset is sent (`rst_ld_cntrs`, 7th line down, goes high), some time passes while data feeding from the RAM (`unused_data_dpr`, 8th line down) is held whilst waiting for a match. When two consecutive events match, the offset (`dpr_addr_offset`, 11th down) can be set. Both sets of data can now flow, and are compared, and the `reset_pointer` (10th line down) drops so errors can be counted. As can be seen from the figure, the error counter (`error`, 12th line down) counts no errors.




Figure 6.5: Wave diagram of bit error rate tester as viewed in ModelSim

LVDS input BER results

An overnight run of 8 hours at 40 MHz on a full CPM of 20 Serialiser chips, each with 4 input channels of 10 bits, accounts to 4.6×10^{13} bits tested. Such a test yielding zero errors implies a bit-error rate of $< 2 \times 10^{-14}$, which is 20 times lower than the target. The Bit-Error Rate testing of the LVDS cables proved successful.

A bit-error rate test was also performed on each input of the CP chips. The same error_comparator block of code was copied into the CP firmware and the firmware regenerated. An upper limit of 10^{-13} per channel was measured.

Study of frequency variation of BER results

A frequency generator was used to alter the clock frequency on the CPM to establish the impact this would have on this result. While the firmware is designed to run at 40.08 MHz, there is some flexibility in this window. The aim of this test was

to establish the extent of the tolerance we had, how future possible modifications to the LHC bunch-crossing rate would be handled by the board, and to establish that we are not operating on the edge of the frequency window which could begin to introduce bit-errors.

The BER tester firmware was loaded, and the frequency increased. Primarily in coarse granularity (steps of 5 MHz) but then in finer granularity (1 MHz steps) to better identify the edge at which errors began.

The system demonstrated a good deal of flexibility at low frequency, with a window of as much as 20 MHz. In the upper range, this was of the order of 10 MHz. While this means that significant increases to the BC rate of the LHC would require new hardware in the Level-1 Trigger, it provides ample contingency for the first several years' running at least.

Temperature variations of BER results

An environmental chamber was used to vary the operating temperature of the CPM between 5°C and 50°C [58]. This is considered to sufficiently cover the temperature range to which the CPM can be exposed, and was used to establish if extreme temperatures would incur additional bit-errors.

The clock jitter showed an increase with temperature but was found to be less than 10 ps over the 45°C range and was therefore considered to be a negligible effect.

The increase in the temperature caused the timing of the data errors (and the corresponding error-free windows) to shift. This is expected and is a result of the electronics behaving as would be expected with temperature variations. The error free window remained approximately constant during the entire temperature range, and so it was concluded that the CPM operates across the full anticipated

temperature range without incurring bit-errors.

6.2.3 Bit-order modifications

Data arrives at the CPM in 10-bit nibbles from the PPM, and should arrive in the format shown in Figure 6.6. Through testing, it emerged that data were being processed in the CPM in the wrong bit order, see Figure 6.7.

Parity	BCMux	8 Data-Bits								LSB
1	0	0	0	1	0	1	1	0	1	

Figure 6.6: Expected order of data arriving on CPMs [59]

8 Data-Bits								LSB	BCMux	Parity
0	0	1	0	1	1	0	1	0	0	1

Figure 6.7: Actual order of data arriving on CPMs [59]

This meant that most data were being read incorrectly, and hence incurring parity errors.

A user constraints file (.ucf) is produced during the synthesis process for a firmware design, which allocates signals and outputs to specific pins. It was possible to edit the .ucf file associated with the serialiser chips, identify the input signals, and swap the bit order around, as shown in Figure 6.8. The serialiser chips then receive the information in the expected order, and hence it is handled correctly, eliminating parity errors.

<pre> <snip> NET "dr(3)" LOC = "n2"; NET "dr(4)" LOC = "p1"; NET "din_xl(0)" LOC = "b16"; NET "din_xl(1)" LOC = "c16"; NET "din_xl(2)" LOC = "d16"; NET "din_xl(3)" LOC = "e13"; NET "din_xl(4)" LOC = "e14"; NET "din_xl(5)" LOC = "e15"; NET "din_xl(6)" LOC = "f12"; NET "din_xl(7)" LOC = "f13"; NET "din_xl(8)" LOC = "f14"; NET "din_xl(9)" LOC = "f16"; NET "din_xm(0)" LOC = "j14"; NET "din_xm(1)" LOC = "j13"; NET "din_xm(2)" LOC = "k16"; NET "din_xm(3)" LOC = "j15"; NET "din_xm(4)" LOC = "g14"; NET "din_xm(5)" LOC = "g15"; NET "din_xm(6)" LOC = "h13"; NET "din_xm(7)" LOC = "h14"; NET "din_xm(8)" LOC = "h15"; <snip> serialiser_oldbitorder.ucf </pre>	<pre> <snip> NET "dr(3)" LOC = "n2"; NET "dr(4)" LOC = "p1"; NET "din_xl(0)" LOC = "d16"; NET "din_xl(1)" LOC = "e13"; NET "din_xl(2)" LOC = "e14"; NET "din_xl(3)" LOC = "e15"; NET "din_xl(4)" LOC = "f12"; NET "din_xl(5)" LOC = "f13"; NET "din_xl(6)" LOC = "f14"; NET "din_xl(7)" LOC = "f16"; NET "din_xl(8)" LOC = "c16"; NET "din_xl(9)" LOC = "b16"; NET "din_xm(0)" LOC = "k16"; NET "din_xm(1)" LOC = "j15"; NET "din_xm(2)" LOC = "g14"; NET "din_xm(3)" LOC = "g15"; NET "din_xm(4)" LOC = "h13"; NET "din_xm(5)" LOC = "h14"; NET "din_xm(6)" LOC = "h15"; NET "din_xm(7)" LOC = "h16"; NET "din_xm(8)" LOC = "j13"; <snip> serialiser_newbitorder.ucf </pre>	BCmux flag Parity bit
--	--	-------------------------------------

Figure 6.8: Swapping pin allocations in Serialiser .ucf file [59]

6.2.4 All-0s PPrASIC parity error

An additional error arose regarding the parity bits; namely that all-0s data from the PPM were arriving with the incorrect parity bit due to an error on the PPrASIC, see Figure 6.10. While the detected error was to be corrected in new versions of the PPrASIC, there was still an issue for testing and integration as all-0s data, which is quite common, would regularly incur parity errors in the CPM.

8 Data-Bits								LSB	BCMux	Parity
0	0	0	0	0	0	0	0	0	0	1

Figure 6.9: Expected parity in all-0s case [59]

8 Data-Bits								LSB	BCMux	Parity
0	0	0	0	0	0	0	0	0	0	0

Figure 6.10: PPrASIC parity error in all-0s case [59]

A new version of the Serialiser firmware was developed which intentionally checked for the incorrect parity in the all-0s case.

Once the new PPrASIC was produced, it was used in testing, but old versions continued to be used as they had already been fully tested and calibrated. So integration tests required multiple versions of the Serialiser firmware in order that the CPM could handle data received from the PPM. In all, three versions of the Serialiser firmware were developed, all with the correct handling of the bit-order:

- Version 9 - which checked correctly for parity and would work with the new PPrASIC
- Version 10 - which checked for wrong parity and would work with the old PPrASIC
- Version 11 - which ignored the parity bit and so could be used with PPMs with old and new MCMs on them

6.3 CP Chip Tests

6.3.1 Real Time Testing

One of the challenges the CPM has to overcome is to correctly receive the 160 MHz data, on board and from the backplane. As detailed in the Serialiser section, the Serialiser playback memories were used to drive data to the CP chips at 160 MHz. The delay between the CP chip clock (deskew2) and the Serialiser chip clock (deskew1) was varied. Delays were added using the available 104 ps step sizes, see Figure 6.11.

Once data were being correctly received inside the CP chip, a timing investigation was also performed. Error checks were performed on each individual pin input of the CP chip and results are shown in Figure 6.12.




Figure 6.11: Diagram of delayed data tests between CP and Serialiser chips [52]

A pattern with an expected period of 6.25 ns is observed, the output being serialised at 160 MHz.

On board input data exhibit an error-free timing window of 2.5 ns, which is about 1 ns wider than for the backplane data. The PCB was then optimised to keep the delay between tracks as short as possible, hence widening this window. The size of this error peak is due to the spread of the components themselves, rather than noise on the tracks. It was also evident from Figure 6.12 that two separate clocks are required in order to overlap the error-free windows of the on-board data and the backplane data.

6.3.2 Latency Measurements

Measurements were made across the CPM to test the latency and see where time could be saved.





Figure 6.12: Error profile of data received inside the CP chip as a function of the delayed 40 MHz clock

Using ModelSim, and a testbench containing simulated hits in order to track signals through the full CP chip, the latency across various block of the chips was measured, see Figure 6.13.

The top level of the CP chip, as visible in HDL Designer, is shown in Figure 6.14. It shows the signals received by each individual CP chip (LPE down to NSE), as well as all other inputs including resets, VME access, and the clocks (separate block bottom right in the figure). The way these inputs relate to trigger towers is illustrated in Figure 6.15.



As can be seen in Figure 6.15, each chip covers a 4×3 array of trigger towers. Each produces five 4-bit nibbles of data except for column 'p' which is along the edge, so only produces three - the rest of these data come from an adjacent chip across the board. This array is duplicated for both the EM and HAD trigger towers. Each trigger tower is received as a 20-bit word (or 12-bit in the case of column p), and these are the inputs seen in Figure 6.14, labelled by the row letter (L, M or N) then the column letter (P, Q, R or S) and then E for EM or H for HAD, producing the signals LPE down to NSE.

The contents of the main block in Figure 6.14, the `cp_core`, is shown in Figure 6.16. The functional blocks within the `cp_core` are the `front_end` (left block), the algorithm (centre block), and the `readout_logic` (smaller block on the right). The relationship between these blocks, and signals which pass between them can be seen in the figure.




Figure 6.14: HDL Designer block diagram illustrating the top level of the CP chip, cp_top



Figure 6.15: Diagram of how the trigger towers relate to the CP inputs

The front_end and algorithm blocks are in the realtime data path, and so the latency across these blocks is important and impact the latency across the whole CPM,




Figure 6.16: HDL Designer block diagram illustrating the main functional part of the CP chip, cp_core

whereas the readout_logic does not. These blocks could be broken down into their constituent logic blocks, see Figures 6.17 and 6.18. The timing across individual components of the front_end block are annotated on the diagram in Figure 6.17. These were measured by tracking the signals entering, and the corresponding signals leaving, each of these blocks, using ModelSim, once the individual signals had been identified at this level. As can be seen in Figure 6.17, the ob_correction (on-board correction) block, second from the left takes 10 ns to process a signal. In parallel with this, the bp_correction block, fourth from the left, would incur the same latency (not marked). Places where most time was being lost could be identified and the impact of using the two available clocks was investigated.

The largest blocks in the front_end block, which have a latency of 47.9 ns, as annotated on Figure 6.17, are the dmux_shell blocks. The dmux_shell was investigated further, and the latency across the components is annotated on Figure 6.19. These clock data in on the rising edge of the clock, but it is possible to clock data on the




Figure 6.17: HDL Designer block diagram illustrating the constituents of the front end block within the CP chip




Figure 6.18: HDL Designer block diagram illustrating the constituents of the algorithm block within the CP chip

falling edge too, which offered the potential of saving half a tick in each case. ModelSim simulations were carried out to ensure that this did not have any unforeseen impacts on the rest of the design. Bit files were then synthesised and uploaded to the chips and tests were run to measure the full impact on latency, as well as to ensure that the chips were handling data correctly.




Figure 6.19: HDL Designer block diagram illustrating the composition of the dmux_shell block within the CP front.end logic

The algorithm block contains the alg_4x4 block which is duplicated several times to carry out the processing of each trigger tower, and each threshold, see Figure 6.20. This block is also in the realtime data path, and so this block was investigated to see if any latency could be saved.




Figure 6.20: HDL Designer block diagram illustrating the composition of the alg_4x4 block within the main algorithm block

In total the Level-1 Trigger is limited to 2 μ s of latency, which includes the cables to and from the detector. This equates to the CPM part of the chain requiring a latency in the order of 20 BC, with a target of 13.8 (345 ns), but this includes the

CMM collecting the data from the CPM crate, latest measurements of which show it to have a latency of 4.9 ticks. LHC ‘clock ticks’ are equivalent to the BC rate, and are the same as the clock ticks across the CPM which also runs at 40MHz, i.e. a tick is 25ns.

Latency measurements across Version 1.8 of the CPM are shown in Figure 6.21. It can be seen that the total latency across the board at 13.4 ticks. Adding this to the 4.9 ticks of the CMM exceeds the target latency, but is just within tolerances. It doesn’t give any contingency though, and so it was preferable to reduce this latency.

A new version (Version 1.9) of the CPM was produced in late 2005. It contained modifications identified in the previous sections, and meant that 12 ns were saved across the Serialiser chips. However, 6.25 ns were added to cope with the clock alignment scheme, so only a net gain of 6.25 BC.




Figure 6.21: Latency across the components of CPM 1.8

In addition, an obsolete calibration scheme was identified in the CP chip firmware, using extra flipflops in the real time data path. This was a section of the logic that

was duplicated, and first created, in the Serialiser chips, but in that case the logic was in the readout logic, and so wasn't on the critical path, so optimisation was not such a high priority. This was revisited in the CP firmware, and the additional flipflops could be removed, saving 2 clock ticks.

6.4 Summary

FPGAs were used on the CPM to provide the flexibility to change the processing ability of the module without having to remove the module when it is in a crate in the ATLAS pit, without incurring the severe time penalty and expense of producing new ASICs and desoldering and reattaching them to the board, and with less effort and time expended.

They have proved to be very versatile, and even prior to operation have permitted many modifications to the design, to work constantly on optimisation, right into final PCB production, to add extra logic temporarily to the CPM to assist other tests, to make minor adjustments to compensate for the changes during the development of adjacent modules in the trigger chain, and to fix any errors in the firmware design which emerge during testing. The main firmware development of the CP and Serialiser chips in Birmingham has been for each of these reasons.

Latency measurements were made across the CPM, and in detail across the CP chip. The CP code was scrutinised to identify if there was anywhere where latency savings could be made. Subsequent modifications were made to both the CP firmware and the PCB design to implement these identified improvements.

Temporary modifications were made to measure bit-error rates when the CPM was exposed to different environments, and to test the LVDS inputs, on both the operation of the Serialiser chips and the CP chips. Modifications were also made to

compensate for the differing parity received in specific cases, from the PPMs, whilst the new MCMs were being developed. This enabled continued testing and integration of the modules whilst awaiting the new MCMs to be developed and produced.

Issues were identified with the order that the CPM was receiving data, and it was established that the data were not being handled correctly. This was quite a fundamental problem, making all data received meaningless as the bits were being dealt with in the wrong order. A permanent modification was made to the Serialiser firmware to rectify the way in which the bits were handled.

Many integration tests were also carried out with the CPM, in small-scale tests in the laboratory either in Birmingham or at RAL, with the PPMs (upstream) or the CMMs and RODs (downstream). The major integration test, though, was the H8 CERN Testbeam in autumn 2004, which is described in the next chapter.

Chapter 7

H8 CERN Testbeam

7.1 Introduction

The feasibility of a trigger design as described in the Chapter 4 was demonstrated following a long series of test beam and laboratory runs using small-scale versions of the trigger during the ‘calorimeter trigger demonstrator program’ [60]. This program consisted of several generations of hardware modules to test algorithms and technologies for the final trigger system and ultimately proved the workability of the trigger concept. The program ended in 1998 with a system consisting of conceptual modules performing the tasks of the final system on a small scale to prove the trigger concept, and the design for the Level-1 Calorimeter Trigger was defined in the ATLAS Technical Design Report [56].

Since then, the three main subsystems of the Level-1 Calorimeter Trigger have been developed individually. Although this has been done in suitable laboratory environments or dedicated test benches, through several iterations and design changes, and with close collaboration between groups, no functional series of modules in the real-time data path had been assembled since. Nor had a full combined readout

been tested. In the autumn of 2004, each of the modules were sufficiently developed that it was now feasible to bring them together and attempt integration. It was also an appropriate time to attempt this, giving enough time for any arising issues to be resolved prior to modules entering final production runs.

The concept of the H8 test beam setup was to simulate a section of the ATLAS barrel, with modules from all parts of the detector being integrated into a full slice test. It also provided an opportunity to exercise the DAQ, monitoring and reconstruction software.

7.2 Setup

The ATLAS H8 Test Beam is located at the SPS, North Area (Prevessin, France) at CERN, where the SPS beam is extracted for various beam tests. The H8 test beam setup uses barrel modules and components of ATLAS detectors to simulate a small section of the ATLAS detector barrel region, see Figure 7.1.




Figure 7.1: Schematic of the Testbeam setup [61]

7.2.1 Beam Characteristics

The SPS beam consisted of protons with an energy of typically 400 GeV (with a maximum of 450 GeV). In normal operation the SPS did not supply a beam structure comparable to the LHCs tight bunches in 25 ns intervals, but instead long spills of ~ 4.8 s in length, every ~ 16.8 s. As the SPS will be part of the accelerator chain for the LHC, operation in a mode with 25 ns bunches was possible, and was tested for a short period.

For the H8 test beam, the SPS proton beam was mainly shot at a target where a beam of secondary particles was created. With the H8 beam optics, a combination of particle type and energy could be selected for use, typically pions or electrons with energies up to 250 GeV. A tertiary beam of low energy particles (2 - 10 GeV) and the attenuated primary proton beam were also available.

The beam line is considered to go through a virtual point of interaction and the η/ϕ -planes of detector modules were positioned in its path. The final beam position in the detectors depended on a set of magnets at the virtual point of interaction, and could be moved in both the η and ϕ directions.

One LAr electromagnetic barrel module and three hadronic TileCal barrel modules were installed. The LAr module covered an area of 0.4 in ϕ and 0.0 to 1.5 in η , while the TileCal modules covered a total of 0.3 in ϕ and -1.0 to +1.0 in η . All modules shared the ϕ direction and the $\eta = 0$ plane, meaning only one half of the TileCal Modules had the LAr module in front of them.

Both calorimeters were centered around the beamline in ϕ and the calorimeters shifted downwards, so the LAr module covered $-0.2 < \phi < +0.2$ and the TileCal covered $-0.15 < \phi < +0.15$. Calibration systems were in place for both calorimeters, allowing precise amounts of energy to be injected directly to obtain pulses for calibration, similar to those resulting from energy deposition from beam particles but

of precisely known energies. These systems could also be used to provide signals for the Level-1 Calorimeter Trigger.

7.2.2 Trigger Towers and Cables

The front-end electronics for summing calorimeter cells into trigger towers, Tower Builder Boards (TBB) for LAr and 3-in-1 boards for the TileCal, were set up and cabled as they will be in the final ATLAS setup. This resulted in the typical trigger tower size of 0.1×0.1 in $\eta \times \phi$.

Three cables had been laid from the detector to the trigger control room for each calorimeter type. The TileCal cables were 60 m long and the LAr ones 70 m. As the TileCal signals went through a patch panel and another 10 m of cable before arriving at the receiver system, the length of cable through which signals passed was roughly equal.

Each TileCal cable carried signals from one half of one module in η , so from an area one ϕ bin high and 10 trigger towers wide from $0 < \eta < 1$. Similarly, each LAr cable had 15 trigger tower channels from one ϕ bin and all η values of the module, $0 < \eta < 1.5$, similar to the final ATLAS setup.

Ultimately, all trigger tower signals of the TileCal arrived at the trigger control room while the LAr calorimeter had four ϕ bins of which only three could be selected simultaneously due to a limited number of cables. Furthermore, only one TBB board was installed for most of the test beam while two are required to cover the LAr ϕ range, so only the upper or lower ϕ half of the module could be accessed at once.

7.2.3 Timing and Trigger

Timing and the transmission of L1A signals was distributed to all detector and Level-1 Trigger components. Generation and distribution of these was done by the CTP group.

During normal SPS operation, the clock was a TTC generated 40 MHz clock, not linked to the beam structure. This meant that it was not possible to calibrate the timing between the arrival of a signal and the clock, as the signals floated freely with respect to the digitisation points. This hindered optimisation of BCID coefficients and prevented fine calibration of resulting energies.

During the 25 ns run, the clock was provided by the SPS and signals arrived in a fixed relationship to the clock. Moreover a clock synchronised with the calibration signals generated by the calibration system could be provided in all modes.

The L1A signals were generated by 2 methods. The first was a mock-up CTP consisting of a simple beam trigger, originating from scintillating counters in the beamline and delayed by the expected Level-1 latency. The second method, planned for the end of the full Level-1 Trigger run, was a real CTP prototype supplied with information from the Level-1 Trigger subsystem.

7.3 Operation

The goal of the test beam setup at CERN was to encompass all detector and Level-1 Trigger systems of ATLAS, and to emulate ATLAS on a small scale. The desired maximally realistic environment was not only important for tests, but sticking to the design of ATLAS also saved preparation work as the setup had already been specified in full detail.

During the test beam, both electromagnetic and hadronic calorimeters were available. On the calorimeter detectors, cells are combined to form trigger towers, with a reduced granularity of 0.1×0.1 in $\eta \times \phi$. The calorimeter trigger towers covered a total region of 0.2×0.8 in $\eta \times \phi$. A picture of the overall system during the test beam is shown in Figure 7.2.




Figure 7.2: The ATLAS Trigger System at the Testbeam [62]

A total of 5 trigger crates were used:

- 1 Preprocessor crate containing 1 PPM
- 1 receiver crate
- 1 processor crate containing 2 CMMs, 1 JEM and 1 CPM
- 1 control crate for the clock control (TTCex) and local control of the Level-1 Accept rate

- 1 readout crate containing 4 RODs.

A single 9U crate with a VME 64xP backplane was used to house the PPr. Two crate controller CPUs were installed, a ‘homebrew’ one based on a standard PC motherboard and a Concurrent Technologies crate controller. In addition, one TCM and one PPM were installed in the crate. A second PPM and two rear-mounted G-link modules were also tested, but not used during data-taking.

The PPM digitises calorimeter pulses and prepares data for the CP and JEP. The board carries 16 MCMs [63], each holding 4 12-bit FADCs, 1 ASIC and 1 timer chip (Phos4). Only half of them were used in the test beam corresponding to a total of 32 channels. Analog calorimeter data are received from the front, and digital data are output to the CP and JEP systems using LVDS 400 MB/s serial links through the back of the board. Energy calibration and BCID are performed on digital data, such as that shown in Figure 7.3.

One CPM and one JEM were installed in a shared crate with one crate controller, one TCM and two CMMs. The LVDS cables coming from the PPM were connected based on the ATLAS cabling scheme, as well as all outputs of the crate: the G-links from the processors to the RODs in the readout data path and connections from the CMMs to the CTP in the real-time data path.

7.4 Timing, Trigger, Control and Readout

Timing and trigger signals were provided by the CTP group and fanned out to all other groups. The Level-1 Calorimeter Trigger used a pair of 6U TTC modules to encode those signals and distribute them to the TCMs via optical links, from where they were further distributed over the backplanes of the crates. For most of the test beam, a mock-up CTP, basically a delayed beam trigger from scintillating counters,




Figure 7.3: A liquid argon pulse digitised by a PPM [62]

provided the L1A signals to initiate readout, but at the end a real CTP was also fed with trigger objects from the CPM and JEM to produce L1As. For the readout of data from the Level-1 Calorimeter Trigger, a 6U crate with 4 RODs was installed and connected to the system along with a CCT crate controller, a Busy module (to issue and manage ‘pause’ commands during a run) and a Local Trigger Processor (LTP) to substitute for the CTP.

All control of the system during the data taking runs was done using the Online Software environment, a component of the ATLAS DAQ system. Data was written out and histogrammed offline as well as directly histogrammed online using a Monitoring package.

7.5 Results

During the test beam, the trigger system was successfully integrated with both LAr and Tile Calorimeters. The CMMs sent results to the CTP and thereby successfully produced triggers from the full trigger chain. The trigger menu was used to select electromagnetic cluster, jet and total energy triggers, and RoIs were successfully sent to the RoI Builder. The trigger system was run routinely under the main ATLAS Run Control system.

A root package was available to produce histograms in realtime of the data being received from the modules within the Level-1 Calorimeter Trigger, and could be used to troubleshoot, calibrate, and check the results as they emerged. The package utilised data from the readout path of each of the modules as received by the TriggerDAQ software, and allowed flexibility over timing integrations, axes scales, and simultaneously produced multiple plots from corresponding data to allow full comparisons to be carried out. Some results from work carried out at the testbeam are presented in this section.

7.5.1 Establishing pedestals in the MCMs

Shortly prior to the testbeam, a new set of MCMs had been produced containing the new ASIC. These had not yet been tested, and so a combination of a limited number of tested and approved new MCMs were used alongside well-established, calibrated old ones.

A pipeline system stores five consecutive FADC-samples, applies weights by multiplying the samples with pre-defined coefficients and sums up the resulting values. The FADC-samples are consecutive in time, so the pulse shape illustrates the electronics response. The coefficients can be set via VME access to the chip firmware,

see Figure 7.4. Thus, the integral over a calorimeter signal is derived, improving transverse energy-resolution and minimising noise-contributions. A subsequent peak-finder attributes the maximum value to the corresponding time-slice (LHC clock cycle).




Figure 7.4: Timing slices illustrating hits, and the pedestal set by the MCM is visible in the empty (5th) pulse [64]

The working range of the method spans from small signals (few GeV equivalent) over the linear signal range (up to ~ 250 GeV) into near saturation. The limit is reached when several consecutive samples are clipped to the maximum FADC-value, see Figure 7.5. It is evident in the fourth and fifth samples that there are two consecutive slices with the same (maximum value). These samples cannot be trivially BCID as there is no individual peak value, but two which share the same value. These pulses are saturated. Although not as clearly visible, it is likely that the third pulse is also saturated as its peak is cut off at 1024 which is the maximum 10-bit value.




Figure 7.5: Timing slices illustrating saturated pulses (4th and 5th) in the MCMs [64]

During the testbeam, these pedestals had to be established and calibrated with the data being received by the AnIn cards on the PPM so that data being passed on to the CPM and JEM were reliable. Measurements of noise were made, identifying the conversion factor for ADC counts to voltage, both for the data being received by the PPM, and also data calibrated for what is being sent from the detector (compensating for cable attenuation, Receiver Board gain, etc.), and allowing for calorimeter and pre-amp noise. These pedestals assume uniform noise from the calorimeter and pre-amp, because Receiver Board gain and read-out electronics noise are comparably negligible, so are adjusted such that their nominal values are the same, providing a flat baseline above which signals can be detected.

Communication with the PPM was possible via the TriggerDAQ software, and also via a specific PPM DAQ package. Used in conjunction with the SimpleTBA outputs

for certain values, it was possible to calibrate the pedestals (four per MCM; but only 8 MCMs were placed on the board (16 slots)). The pedestals were established using 10-bit data for accuracy, and then the 8-bit data was checked. The PPM DAQ also allowed the established pedestal thresholds to be loaded at each restart, so once calibrated they continued to produce reliable data.

The digitised pulse is converted to the single 8-bit E_T value, phased to enter the trigger pipeline at the correct time (correct bunch number) for all calorimeter towers, see Figure 7.6. Most of the calibration is ultimately incorporated into this process.




Figure 7.6: Time slices illustrating the pedestals with BCID implemented [64]

7.5.2 BCID of TileCal calibration pulse

The standard SPS beam structure is quite different from that of the LHC, with there being no correlation between the clock used for the systems and the occurrence of beam particles. Compounded with issues with non-functioning channels in the PPM,

and parity errors at the JEM input, data taking was difficult to interpret. A period of a lower energy tertiary particle beam (~ 2 GeV) was used to check and verify time offsets of the different subsystems and the functioning of channels making use of the TileCal calibration system. Uniform signals with an amplitude of 1.5 V at the PPr input were created with a fixed correlation to the clock.




Figure 7.7: Time slices readout from the CPM, bypassing BCID (left) and with BCID (right) [64]

At first the DAQ readout offsets were set to catch the raw 10-bit ADC results of the pulses, bypassing BCID. The results for one channel are shown on the left of Figure 7.7. The abscissa is the time in 25 ns ticks, but it is not continuous as empty events are now being suppressed and only data actually read out are placed in the histogram. Five slices are taken each time, to allow the peak value to be identified by using the shape of the pulse. If the system is properly calibrated, this will be in the central slice. Adjacent pulses are separated by an empty bin for clarity.

The plot on the right of Figure 7.7 shows the situation after the BCID logic has been enabled and the readout time offset adjusted back to normal operation. The peak-finding algorithms permit the highest of the five values to be identified as the peak, and hence the signal corresponding to the peak energy of the BC, to be identified.

A single 8-bit value remains from each pulse, representing the corresponding energy determined by the BCID.

The total energy sum readout from both the CPM and JEM are shown in Figure 7.8. A distinct peak is visible corresponding to the uniform calibration signals, but the energy sum over all channels in the JEM is much lower. This strongly suggests that some channels have constant parity errors at the JEM input, as data with parity errors is zeroed. Some CPM data also had parity errors arriving at the CMM, but this could be fixed with minor timing adjustments.




Figure 7.8: Readout of the total energy sum over all channels in the CPM and JEM [65]

7.5.3 Beam Energy Measurements with the CPM and JEM

A data taking run was conducted during a period with a 250 GeV beam at $\eta = 0.55$ consisting primarily of pions with about 10% electrons. Any non-functioning channels had been masked out on both the JEM and CPM. Figure 7.9 compares the event histograms of the two processors, while the energies read out from the JEM and CPM are compared on a per-event basis in Figure 7.10. Data shown is from the electromagnetic LAr calorimeter barrel module.




Figure 7.9: 250 GeV beam events read out from the CPM and JEM [62]

The first point to note is that the system was not calibrated. The attenuation of the signal on the cables from the front-end electronics alone reduces the amplitude by about 35%. The sharp peaks at the left edge of the plots in Figure 7.9 are false BCIDs caused by noise. These were read out when a trigger was produced by a particle which then hit the calorimeter outside the region visible to the Level-1 Calorimeter Trigger setup. The thick tail of the distribution towards low energies is a combination of effects including the composite beam being mainly pions, and some electrons occasionally hitting the edge of, and losing energy in, multiple trigger towers.

There is a clear correlation between the energy read out from the JEM and CPM as shown in Figure 7.10. Although this would be naively expected as both systems receive the same parity protected digital data, it demonstrates that there are no problems in the realtime path that affect data taking. For example, the absence of points with significant energy in the CPM but none in the JEM provide proof that the channels used did not suffer from parity errors upon arrival at the JEM. The line of points directly below the diagonal is to be expected as in some cases the




Figure 7.10: Correlation between energies of events read out from the CPM and JEM [62]

reduced granularity and resolution of the jet elements lead to slightly lower energy sums in the JEM.

7.6 Summary

The results presented above are only a small sample, and a much larger number of achievements were made, and milestones reached, at the testbeam. All Level-1 Calorimeter Trigger subsystems were successfully combined, supplying reasonable data and results. The Online Software was used to control the modules, and the

Level-1 Calorimeter Trigger was integrated into the combined test beam, allowing common runs with the CTP, the Muon System and the detectors. Significant progress was made in connecting the Level-1 Calorimeter Trigger to the DAQ read-out system and the CTP, where many problems were identified and overcome.

Furthermore, a prototype CTP was successfully supplied with trigger objects from both Level-1 Calorimeter processors; multiplicities of clusters passing energy thresholds from the CPM, jets from the JEM, and the thresholds passed by the global energy sums of the JEM. An L1A was generated from this data and used to initiate readout, thus the entire functionality of the real-time data path has been verified.

To conclude, it was demonstrated that the Level-1 Calorimeter Trigger is able to meet its requirements, and is fully functional. Some problems were identified in all modules, however, and work identified which it is necessary to complete before commissioning of the system.

7.7 Conclusions

The ATLAS detector at CERN will be at the frontier of particle physics development, and hopes to find the Higgs boson, and possible physics beyond the SM. At full operating capability, it will exceed the energy and luminosity of current particle accelerators by factors of 7 and 10 respectively, and will explore new energy regions for new physics.

The extreme conditions of the ATLAS experiment place stringent requirements on the detector and the Level-1 trigger, and the high data-rate is processed in dedicated hardware which has been developed at a modular level over the past ten years.

Over the past few years, the CPM has developed from conception to a final production design. Many tests have been done to aid the evolution of the board and the

firmware in the chips, and to test the module's ability to perform its tasks in the extreme environment in which it will be placed.

The use of FPGAs on the CPM provides ongoing flexibility as new, emerging physics can be properly investigated using the flexible thresholds and even algorithm modifications if this should be necessary. The software environment provides a flexible platform on which to develop the firmware still further and minimal disruption is caused to upgrade the trigger in this case.

Extensive integration tests, including the H8 CERN testbeam, have been carried out between the CPM and modules both up and downstream, so confidence can be placed in the full Level-1 Trigger's integration and full data processing capability. Early running at the LHC will not be at full design luminosity or energy, so a period will be available for calibration and any minor modifications required in-situ. The trigger is currently being installed in the ATLAS pit at CERN, and hopes to see its first data by the end of 2007.

Glossary

ALICE	A Large Ion Collider Experiment
ASIC	Application Specific Integrated Circuit
ATLAS	A Toroidal LHC Apparatus
BC-mux	Bunch-Crossing Multiplexing
BCID	Bunch-Crossing Identification
CANbus	Controller Area Network bus
CLB	Configurable Logic Block
CM	Centre of Mass
CMC	Common Mezzanine Card
CMM	Common Merger Module
CMOS	Complimentary Metal Oxide Semiconductor
CMS	Compact Muon Solenoid
CP	Cluster Processor
CPM	Cluster Processor Module
CSC	Cathode Strip Chambers
CTP	Central Trigger Processor
DAQ	Data Acquisition
DPR	Dual Port RAM
ECAL	Electromagnetic Calorimeter
EM	Electromagnetic
FADC	Flash Analogue-to-Digital Converter

FIFO	First In First Out
FIR	Finite Impulse Response
FPGA	Field Programmable Gate Array
GUT	Grand Unified Theory
HAD	Hadronic
HCAL	Hadronic Calorimeter
IOB	Input and Output Block
JEM	Jet/Energy-Sum Module
JEP	Jet/Energy-Sum Processor
L1A	Level-1 Accept
LAr	Liquid Argon
LEP	Large Electron-Positron Collider
LFSR	Linear Feedback Shift Register
LHC	Large Hadron Collider
LHCb	LHC Beauty Experiment
LTP	Local Trigger Processor
LUT	Look-Up Table
LVDS	Low Voltage Differential Signal
MCM	Multi-Chip Module
MDT	Monitor Drift Tube
PAR	Place And Route
PPM	PreProcessor Module
PPr	PreProcessor
PS	Proton Synchrotron
QCD	Quantum Chromodynamics
QED	Quantum Electrodynamics
RAL	Rutherford Appleton Laboratory
RAM	Random Access Memory
RF	Radio Frequency

ROB	Readout Buffer
ROC	Readout Controller
ROD	Readout Driver
RoI	Region of Interest
RPC	Resistive Plate Chamber
SCT	Semi-Conductor Tracker
SM	Standard Model
SPS	Super Proton Synchrotron
SRL	Serialiser
SUSY	Supersymmetry
TBB	Tower Builder Boards
TCM	Timing and Control Module
TGC	Thin Gap Chamber
TRT	Transition Radiation Tracker
TTC	Timing, Trigger and Control
TTL	Transition Time Length
UCF	User Constraints File
VHDL	VHSIC Hardware Description Language
VHSIC	Very High Speed Integrated Circuit
VME	Virtual Machine Environment

Bibliography

- [1] Y.Fukuda *et al.* ‘Evidence for oscillation of atmospheric neutrinos’, Phys. Rev. Lett. B 433, **9** (1998).
- [2] Particle Physics Booklet (2004 edition), extracted from Review of Particle Physics, Eidelman *et al.*, Phys. Rev. Lett. B 592, **1** (2004).
- [3] D. E. Groom *et al.*, ‘Review of Particle Physics’, Springer Verlag (2001).
- [4] D. Perkins, ‘Introduction to High Energy Physics’, Cambridge University Press (2000).
- [5] D. Griffiths, ‘Introduction to Elementary Particles’, John Wiley and Sons (1987).
- [6] The LEP Collaborations ALEPH, DELPHI, L3, OPAL, the LEP Electroweak Working Group and the SLD Heavy Flavour and Electroweak Groups. A Combination of Preliminary Electroweak Measurements and Constraints on the Standard Model. CERN-EP preprints, CERN-EP-2001-098 (2001). <http://lepewwg.web.cern.ch/LEPEWWG/stanmod/>.
- [7] The ATLAS Detector and Physics Performance. Technical Design Report ATLASTDR-14, CERN/LHCC 99-14, CERN (1999).
- [8] B.R. Martin & G. Shaw, ‘Particle Physics’, Second Edition, The Manchester Physics Series, John Wiley & Sons (1999).

- [9] T. Wengler, ‘Frontiers of Particle Physics 1 - The Search for the Higgs’, University Lectures, CERN / PH-Dep.CH-1211 (2005). <http://twengler.web.cern.ch/twengler>.
- [10] W. de Boer, ‘Grand Unified Theories and Supersymmetry in Particle Physics and Cosmology’, IEKP-KA/94-01 hep-ph/9402266 v5, 7 February 2001.
- [11] F. Halzen & A. D. Martin, ‘Quarks and Leptons. An Introductory Course in Modern Particle Physics’, John Wiley and Sons (1984).
- [12] The ATLAS Collaboration. The ATLAS Detector and Physics Performance. Technical Design Report ATLASTDR-14, CERN/LHCC 99-14, CERN (1999). <http://atlasinfo.cern.ch/Atlas/GROUPS/PHYSICS/TDR/access.html>.
- [13] I. P. Brawn, ‘Bunch-Crossing Identification for the ATLAS Level-1 Calorimeter Trigger’, PhD Thesis, University of Birmingham (1996).
- [14] European Committee for Future Accelerators. Large Hadron Collider Workshop, Proceedings. CERN 90-10 (2001).
- [15] The Large Hadron Collider Accelerator Project (1993), <http://www-td.fnal.gov/LHC/USLHC.html>.
- [16] LHC Study Group, The Large Hadron Collider Conceptual Design, CERN-AC-95-05-LHC (1995).
- [17] LHC Design report, CERN-0000024512 (2005) <https://edms.cern.ch/nav/CERN-0000024512>.
- [18] Rudolf LEY *et al.*, PS Division, CERN accelerator chain. <http://ps-div.web.cern.ch/ps-div/PSComplex/>.
- [19] LHC Layout from the CERN Webpages: <http://st-div.web.cern.ch/st-div/Projects/>.

- [20] The ATLAS Collaboration. ATLAS Technical Proposal CERN/LHCC 94-43, LHCC/P2, CERN (1994).
- [21] The ALICE Collaboration. ALICE Technical Proposal Cern/LHCC 95-71, LHCC/P3, CERN (1995).
- [22] The LHCb Collaboration. LHCb Technical Proposal CERN/LHCC 98-4, LHCC/P4, CERN (1998).
- [23] The CMS Collaboration. CMS Technical Proposal CERN/LHCC 94-38, LHCC/P1, CERN (1994).
- [24] CERN ATLAS website, <http://atlas.ch/>.
- [25] B. J. Gallop, ‘ATLAS SCT Barrel Macro-assembly testing and design of MAPS test structures’. PhD Thesis, University of Birmingham (2005).
- [26] ATLAS Experiment Webpages: http://atlasexperiment.org/atlas_photos/innerdet/innerdet_general/.
- [27] Carlton University Webpages: <http://www.physics.carleton.ca/research/experimental/atlas/>.
- [28] Boston University ATLAS homepage: <http://physics.bu.edu/ATLAS/ATLAS-info/>.
- [29] ATLAS: Detector and Physics Performance Technical Design Report. Volume I. page 18, CERN-LHCC-99-14 (1999).
- [30] CERN ATLAS Webpages, DAQ and Trigger Control Group, http://atlas.web.cern.ch/Atlas/SUB_DETECTORS/DetStatus/daqtriggercontrol/.
- [31] F. Touchard *et al.*, ATLAS TDAQ/DCS High Level Trigger, HLT Infrastructure Operational Analysis and Requirements to other (sub-)systems, v1.0, November 2002.

- [32] J. Garvey *et al.*, ‘The ATLAS Level-1 Calorimeter Trigger Architecture’, IEEE Trans. Nucl. Sci. **51**, 356-360 (2004).
- [33] E. Eisenhandler, ‘Level-1 Calorimeter Trigger Status’, *presentation at ATLAS Overview Week, Stockholm, 12 July 2006*.
- [34] S. Silverstein, ‘ATLAS Level-1 Calorimeter Trigger Architecture’, presentation on behalf of ATLAS Level-1 Calorimeter Trigger Collaboration *Proceedings of the 13th Conference on Real Time (RT2003), Montreal, Canada, 18-23 May 2003*.
- [35] J. Garvey *et al.*, ‘Atlas Level-1 Calorimeter Trigger: Subsystem Tests Of A Jet/Energy-Sum Processor Module’, IEEE Trans. Nucl. Sci. **51** (2004) 2356.
- [36] J. Garvey *et al.*, ‘Test Results for the Jet/Energy Processor of the ATLAS Level-1 Calorimeter Trigger’, 235-239; Atlas Internal Note ATL-DAQ-2003-049, 5 October 2003, *Proceeding of the 9th Workshop On Electronics For LHC Experiments (LECC 2003), Amsterdam, The Netherlands, 29 September - 3 October 2003*.
- [37] G. Anagnostou *et al.*, ‘One size fits all: Multiple uses of Common Modules in the ATLAS Level-1 Calorimeter Trigger’, CERN/LHCC/2001-034, p253. *Prepared for 7th Workshop on Electronics for LHC Experiments, Stockholm, Sweden, 10-14 September 2001*.
- [38] Agilent G-link information available at: <http://www.seiconductor.agilent.com/>.
- [39] G. Anagnostou *et al.*, ‘Prototype Readout Module for the ATLAS level-1 Calorimeter Trigger Processors’, CERN/LHCC/2001-034, *Prepared for 7th Workshop on Electronics for LHC Experiments, Stockholm, Sweden, 10-14 September 2001*
- [40] CERN S-link Specification: <http://www.cern.ch/HIS/s-link>.

- [41] N. Ellis, ‘Trigger, DAQ and DCS’ presentation, TDMT on behalf of the Trigger/DAQ system, 20040227-umwlap002-06-ellis, April 2004.
- [42] M. K. Jayananda, ‘ATLAS Level-1 Calorimeter Trigger - Simulation of the backplane for Common Merger Module’, ATLAS ATL-DAQ-2004-004, CERN, Geneva (2000).
- [43] G. Anagnostou *et al.*, ‘Prototype Cluster Processor Module for the ATLAS Level-1 Calorimeter Trigger’, CERN Yellow Report 2002-003, 256-260, *Prepared for the 8th Workshop on Electronics for LHC Experiments (LEB 2002), Colmar, France, 9-13 September 2002*
- [44] J. Garvey *et al.*, ‘Use of an FPGA to identify electromagnetic clusters and isolated hadrons in the ATLAS Level-1 Calorimeter Trigger’, Nucl. Instrum. Meth. A **512**, 506 (2003).
- [45] S. Hillier *et al.*, ATLAS Level-1 Calorimeter Trigger Prototype Cluster Processor Module, Project Specification (FDR), Version 2.00, 15 March 2005.
- [46] V. Perera *et al.*, Project Specification: ATLAS Calorimeter First Level Trigger - Serialiser FPGA, Version 1.3, 11 March 2005.
- [47] V. Perera *et al.*, Project Specification: ATLAS Calorimeter First Level Trigger - Cluster Processor Chip, Version 2.0, March 2005.
- [48] ModelSim SE User’s Manual, Mentor Graphics, Version 6.0, 20 July 2004.
- [49] G. Anagnostou *et al.*, ‘Study of LVDS Serial Links for the ATLAS Level-1 Calorimeter Trigger’, CERN 2000-010, CERN/LHCC/2000-041 333-337 (2000), *Proceedings of the 6th Workshop On Electronic For LHC Experiments (LEB 2000), Crakow, Poland, 11-15 September 2000.*
- [50] R. Staley, University of Birmingham, ATLAS First Level Trigger - LVDS source CMC daughter boards for Data Source and Sink (DSS) Module, Design Specification, Version 3.0.x, 16 September 2002.

- [51] V. Perera *et al.*, Data Source and Sink (DSS) Specification, Version 2d, March 2000.
- [52] G. Mahout, ‘ATLAS Level-1 Calorimeter Trigger: Subsystem Tests of a Prototype Cluster Processor Module’, *Prepared for the 9th Workshop on Electronics for LHC Experiments, Amsterdam, 29 September - 3 October 2003.*
- [53] G. Mahout, ‘Performance of the new Cluster Processor Modules - Version 1.9’, *Prepare for The Level-1 Calorimeter Trigger System Joint Meeting, Mainz, April 2005.*
- [54] R. Staley, LVDS Source Module, Design Specification v2.1, 11 March 2005.
- [55] J. Garvey *et al.*, ‘ATLAS Level-1 Calorimeter Trigger: Subsystem Tests of a Prototype Cluster Processor Module’, Atlas Internal Note ATL-DAQ-2003-051, 12 December 2003.
- [56] ATLAS First Level Trigger Technical Design Report, CERN/LHCC/98-14 and ATLAS TDR-12, 30 Jun 1998.
- [57] P. Alfke, ‘Efficient Shift Registers, LFSR Counters, and Long Pseudo-Random Sequence Generators’, Version 1.1, Xilinx Application Note, XAPP 052, 7 July 1996.
- [58] J. Booth *et al.*, ‘Environmental Chamber Test of the Cluster Processing Module’, 18 March 2005.
- [59] T. Moye, ‘Parity-error Handling Firmware’, *Prepared for the Level-1 Trigger Meeting, RAL, 2 June 2004.*
- [60] C. Bohm *et al.*, ‘Test-Beam Results from the ATLAS Level-1 Calorimeter Trigger Demonstrator’, IEEE Trans. Nucl. Sci. **46**, 904-909 (1999), *Proceedings of the Nuclear Science Symposium & Medical Imaging Conference (NSS/MIC), Toronto, Canada, 8-14 November 1998.*

- [61] J. Haller, ‘Searching for New Physics at High Energies’, *presentation in Hamburg, Germany, 19 September 2005*.
- [62] J. Garvey *et al.*, ‘Beam Test of the ATLAS Level-1 Calorimeter Trigger System’, Atlas Internal Note ATL-COM-DAQ-2005-010; Atlas Note ATL-DAQ-CONF-2005-006, 2 March 2005.
- [63] G. Anagnostou *et al.*, ‘The final Multi-Chip Module of the ATLAS Level-1 Calorimeter Trigger Preprocessor’, *Prepared for 7th Workshop on Electronics for LHC Experiments, Stockholm, Sweden, 10-14 Sep 2001*.
- [64] D. Sankey, RAL, SimpleTBA testbeam analysis software used to create root graphs during the testbeam at CERN during Autumn 2004.
- [65] F. Ruhr, ‘Initial Tests of the ATLAS Level-1 Trigger Pre-Processor’, Diploma Thesis, Kirchhoff-Institut für Physik, University of Heidelberg, HD-KIP-04-14 (2005).