

DEVELOPMENT OF THE DETECTOR
CONTROL SYSTEM FOR THE ATLAS
LEVEL-1 TRIGGER AND
MEASUREMENT OF THE SINGLE TOP
PRODUCTION CROSS SECTION

Christopher James Curtis

*Thesis submitted for the degree of
Doctor of Philosophy*



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

October 2009

UNIVERSITY OF
BIRMINGHAM

University of Birmingham Research Archive

e-theses repository

This unpublished thesis/dissertation is copyright of the author and/or third parties. The intellectual property rights of the author or third parties in respect of this work are as defined by The Copyright Designs and Patents Act 1988 or as modified by any successor legislation.

Any use made of information contained in this thesis/dissertation must be in accordance with that legislation and must be properly acknowledged. Further distribution or reproduction in any format is prohibited without the permission of the copyright holder.

Abstract

This thesis discusses the development of the Detector Control System (DCS) for the ATLAS Level-1 Trigger. Microcontroller code has been developed to read out slow controls data from the Level-1 Calorimeter Trigger modules into the wider DCS. Back-end software has been developed for archiving this data. A Finite State Machine (FSM) has also been developed to offer remote access to the L1 Trigger hardware from the ATLAS Control Room.

This Thesis also discusses the discovery potential for electroweak single top production during early running. Using Monte Carlo data some of the major systematics are discussed. A potential upper limit on the production cross section is calculated to be 45.2 pb. If the Standard Model prediction is assumed, a measured signal could potentially have a significance of up to 2.23σ using 200 pb^{-1} of data.

Author's Contribution

The ATLAS collaboration is a large organisation, and although I have received a lot of support from both the central ATLAS DCS group, L1Calo collaboration and the Single Top Working Group, I would like to make clear my own contribution to the work presented in this thesis.

In terms of the L1 Trigger DCS, my contribution consisted of work at the LCS level. I designed and wrote the monitoring code for the L1Calo Module Microcontrollers, enabling the modules to communicate with the DCS back-end. At the level of the LCS, I designed and wrote the first implementation of the L1 Trigger PVSS project, using some software components made available by ATLAS DCS and JCOP. This initial design consisted of datapoint configuration enabling communication with both Wiener crates and the L1Calo modules, data archiving and configuration via XML files. I also designed and implemented both the FSM and the graphical user interface. This project is still used as the base for the current L1 Trigger DCS.

In terms of the Single Top Analysis, I made use of centrally produced AOD files (forming part of the reprocessed MC08 datasets). I developed my own analysis code for use in Athena (release 14.5.0.5) and ROOT, the results of which are presented in this thesis.

This work is dedicated to Alex Rendall

Acknowledgements

First thanks must go to Pete Watkins. His guidance more than anyone else over the last five years has shaped my work and development.

To Steve Hillier and Gilles Mahout, who introduced me to the joys of L1Calo, taught me how to code (and debug) and helped me stick at it when it all failed (again and again and again).

To Murrough Landon and Bruce Barnett, whose helpful comments regarding the DCS helped shape it into the slick machine it is today (“I’m not sure that 30 minutes to initialise the DCS will be popular...”).

To Paul Thompson, who stepped in and took the reigns when time was short. Chris Hawkes, who helped correct my errors. Miriam Watson for asking the right questions. To Juergen Thomas, for his unending wisdom on so many topics (especially the impenetrable Athena trigger code).

To Mark Slater and Yves Coppen, who sat and patiently fixed my grid and ganga problems after I had resorted to mild acts of violence towards my PC. To Lawrie Lowe, who taught me everything I know about linux.

To the many and varied residents of 316 and beyond - Joseph, Neil, Dave, Owen, Sparsh, Zoe, Tim, Dan, Jody - and your welcome (and many) distractions.

And to Mark, who constantly explained, challenged and stole my code from CVS.

To Richard, Ivan and Ethan, who taught me everything I needed to know about going to CERN. To Tamsin, who set the standards to which we all aspire. To Dan, who taught me random Spanish phrases and how to organise. And to Liam, without

whom the first year and summer school would have been a total disaster.

To Desperate!, Team Easy and The Medical Harem, who all reminded me that there are still normal people outside of physics.

To Matt, Al, Sam and Emmsie for food and frolics. To Dom for chocolate flavoured Go.

To Ria and Ben, who were there for good times, bad times and schnitzel. To you both I owe so much.

To Mum and Dad, without whom I most definitely could not and would not have got to this stage. To Zoe and Caroline, who reminded me that physics is not the centre of everything.

Most importantly, I would like to thank Alex, who has been a constant source of support, inspiration and motivation throughout this project.

Finally, I would like to thank the Science and Technology Research Council (STFC, formerly PPARC) for their financial contribution, as this project would not have been possible without their support.

“Virtual particle? Is this going to be like your pretend numbers?”

Zoe Curtis, my sister, keeping me real.

Contents

1	Overview	1
2	Motivation	3
2.1	The Top Quark	4
2.2	Top Quark Production	6
2.3	Electroweak Production	8
2.4	Top Quark Decay	11
2.5	Conclusion	13
3	The LHC and the ATLAS Detector	14
3.1	The ATLAS Detector	16
3.2	Tracking	19
3.3	Calorimetry	21
3.3.1	Electromagnetic Calorimetry	22

3.3.2	Hadronic Calorimetry	22
3.4	Muon Spectrometry	23
3.5	Triggering and Data Acquisition	25
3.6	Luminosity Measurement	27
3.7	Conclusion	28
4	The Detector Control System	29
4.1	Controlling ATLAS	29
4.1.1	Detector Safety System	30
4.1.2	JCOP	31
4.2	The ATLAS DCS Structure	32
4.3	The DCS Front-End	34
4.4	The DCS Back-End	34
4.4.1	Local Control Stations	34
4.4.2	Subdetector Control Stations	35
4.4.3	Global Control Stations	36
4.4.4	Finite State Machine	36
4.4.5	Alarms	39
4.4.6	Archiving	41

4.4.7	Security	42
4.5	Level-1 Trigger DCS	43
5	Level-1 Trigger DCS Front-End	47
5.1	Project Requirements and Scope	47
5.2	Monitored VME Crates	48
5.3	L1Calo Trigger Modules	49
5.4	CAN Microcontroller Code	51
5.4.1	Node ID	52
5.4.2	Data Transfer	54
5.4.3	Network Management	54
5.4.4	Module Initialisation	55
5.4.5	Generic Event Cycle	56
5.4.6	TCM Event Cycle	57
5.4.7	Error Messages	61
5.4.8	Other Asynchronous Events	63
5.5	Conclusion	64
6	Level-1 Trigger DCS Back-End	66

6.1	Project Requirements and Scope	66
6.2	Hardware Connection	67
6.3	PVSS	69
6.3.1	Wiener Crates	69
6.3.2	L1Calo Modules	70
6.3.3	TCM	71
6.4	FSM	73
6.4.1	Crate-only Partitions	74
6.4.2	L1Calo Partition	77
6.4.3	Module Device Units	78
6.4.4	TCM Device Units	80
6.4.5	Higher Level Control Units	82
6.5	FSM Operation	82
6.6	Conclusion	84
7	Single Top Production	86
7.1	Monte Carlo Model	86
7.1.1	Single Top	87
7.1.2	$t\bar{t}$	88

7.1.3	W/Z+jets	89
7.1.4	Other backgrounds	90
7.2	Object Reconstruction	92
7.2.1	Electrons	93
7.2.2	Muons	94
7.2.3	Jets	95
7.2.4	Overlaps	96
7.2.5	b-tagging	96
7.2.6	Missing Transverse Energy	99
7.3	Event Selection	99
7.3.1	Trigger and Lepton Identification	100
7.3.2	Jet Identification	103
7.3.3	Further Selection	103
7.4	Cross Section Measurement	106
7.4.1	Monte Carlo Statistics	109
7.4.2	Luminosity	110
7.4.3	Background Cross section	111
7.4.4	Parton Density Function	111

7.4.5	b-tagging	113
7.4.6	Jet Energy Scale	113
7.4.7	Cross Section Upper Limit	117
7.5	Conclusion	118
A	Glossary	120
B	Cut Flow	123

List of Figures

2.1	W mass corrections	5
2.2	Higgs Mass Constraint.	5
2.3	$t\bar{t}$ Pair Production.	6
2.4	Parton Density Function.	7
2.5	Single Top Production.	8
2.6	BSM Single Top Cross Section Variations.	10
2.7	Alignment of single top spin.	12
2.8	Spin Correlations in Single Top Decay Products	13
3.1	The LHC and associated accelerators.	15
3.2	The ATLAS Detector.	18
3.3	The Inner Tracking Detectors.	19
3.4	The Calorimeters.	21
3.5	The Muon Systems.	24

3.6	The ATLAS Trigger.	26
4.1	Wiener VME Crate Display.	32
4.2	ATLAS DCS Overview.	33
4.3	ATLAS DCS FSM Structure.	37
4.4	FSM Operational Modes.	38
4.5	DCS Alarm Screen.	40
4.6	DCS Databases.	42
4.7	Level-1 Trigger	45
5.1	Subdetector Wiener Crates.	48
5.2	L1Calo Module CANBus.	53
5.3	CANopen Data Frame.	55
5.4	L1Calo Module Event Cycle.	58
5.5	TCM Sync Cycle.	59
5.6	TCM Event Cycle.	60
5.7	L1Calo Module Error Levels.	62
5.8	L1Calo ROC Crate.	64
5.9	Updating L1Calo Operating Limits.	65

6.1	DCS OPC Servers.	68
6.2	L1Calo Module Synchronisation.	72
6.3	Wiener Crate FSM.	75
6.4	Wiener Crate FSM Logic.	76
6.5	L1Calo Crate FSM.	77
6.6	L1Calo FSM Overview.	78
6.7	L1Calo Module FSM Logic.	79
6.8	L1Calo TCM FSM Logic.	81
6.9	L1Calo Crate FSM Logic.	83
6.10	L1Calo FSM GUI.	85
7.1	Single Top t-channel Production.	87
7.2	$t\bar{t}$ production.	88
7.3	W/Z+jets production.	89
7.4	Wt Associated Production.	90
7.5	Diboson Production.	91
7.6	W $b\bar{b}$ +jets Production.	91
7.7	Electron ID cuts.	94
7.8	(Truth - Reconstructed) Muon P_T	95

7.9	Electron/Jet Overlap Removal	97
7.10	b-tagging algorithm output	98
7.11	Primary lepton P_T	101
7.12	Lepton Invariant Masses.	102
7.13	Primary b-jet P_T	103
7.14	N Light Jets	104
7.15	Event geometry.	105
7.16	Centrality of b-jet and light jet.	106
7.17	$\Delta\eta$ between the primary b-jet and primary light jet.	107
7.18	$\Delta\eta$ between the primary light jet and lepton.	107
7.19	b-tagging Systematics.	114
7.20	JES Systematics.	115
7.21	Upper Limit on t-channel cross section as function of integrated lu- minosity.	118
7.22	Standard Model cross section prediction for t-channel production. . .	119

List of Tables

5.1	L1Calo Module Types.	50
5.2	L1Calo Module NMT States.	56
7.1	Number of expected and simulated events for the single top t-channel and selected backgrounds channels.	92
7.2	Number of selected events in 200 pb ⁻¹ data.	108
7.3	Systematic Effect of Limited MC Statistics	110
7.4	Systematic Effect of PDF	113
7.5	Systematic Effect of b-tagging	116
7.6	Systematic Effect of Jet Energy Scale	116
7.7	Systematic Scenarios	117

Chapter 1

Overview

At the time of writing, the Large Hadron Collider (LHC) is once again entering the final stages of commissioning before the proton beams are injected. Collisions are expected before the end of 2009, with the accelerator expected to run throughout 2010. During this time, the ATLAS detector will be recording data with the intention of exploring a wide physics programme.

This thesis represents activity in two areas of the ATLAS collaboration - the Detector Control System and Electroweak Top Production.

Chapter 2 gives a short overview of top quark physics at the LHC and ATLAS, discussing in part why electroweak top production should be studied. Chapter 3 discusses some of the major components of the ATLAS detector and how they relate to object reconstruction.

Chapters 4 - 6 are concerned with the Detector Control System (DCS) for the Level-1 Trigger. The need for the DCS is highlighted, and the integration of controls for the Level-1 Trigger discussed. Some of the technical achievements in reading monitoring

data from the hardware into the wider ATLAS DCS are also explored.

Chapter 7 discusses electroweak top production during early data running at ATLAS, and how it might be observed. An upper limit on the production cross section is calculated, making some conservative estimates on some of the dominant systematic errors. Finally, the significance of a measurement at the standard model prediction is calculated.

Chapter 2

Motivation

The Large Hadron Collider (LHC) has been designed to collide protons with a centre of mass energy of 14 TeV. This order of magnitude increase in the collision energy, when compared to the Tevatron, will allow physics studies to work in a completely new energy regime. A lot of effort has already been made into understanding how physics Beyond the Standard Model (BSM) might be explored, but it is important to remember that the LHC experiments will also be improving on existing measurements. Some of these measurements have been limited by very low statistics at previous experiments which is where, with its high luminosity, the LHC will be able to make the largest impact. This is especially true in the field of the top quark.

This chapter explores some of the aspects of top quark production and decay, and how studying these areas might lead to the discovery of new physics at the LHC.

2.1 The Top Quark

The top quark is the heaviest of the known quarks, with a mass of 171.3 ± 1.6 GeV [1]. The first direct observation was at the Tevatron [2, 3]. Since its discovery, the DØ and CDF experiments have continued to investigate properties of the top quark, putting upper limits on its lifetime [4] and measuring its electric charge [5]. These studies have been limited by the low statistics achievable at the Tevatron. The LHC experiments will not suffer the same problem.

The top quark in the standard model is the electroweak partner of the b quark, and is predicted to share similar characteristics with the other u-type quarks. For example, ATLAS intends to measure the electric charge of the top quark, which is predicted to be $+2/3$ [6]. ATLAS studies also show that the top mass may be measured with a precision of a few GeV with only 1 fb^{-1} of data [6].

One of the important consequences of the top quark's high mass is that it contributes radiative loop corrections to the W and Z boson masses (figure 2.1(a)). By accurately measuring the Z mass, the LEP experiments placed limits on the top mass before it was discovered at the Tevatron [7].

If it exists, the Higgs boson is also expected to provide loop corrections to the W mass. By measuring the W and top masses with a high accuracy, constraints may be placed on the Standard Model Higgs mass. Figure 2.2 shows the 68% confidence contour for direct measurements of the W and top masses at LEP and the Tevatron. Also shown are lines corresponding to different Higgs masses. [8].




Figure 2.1: Corrections to the W and Z masses coming from (a) top quarks and (b) Higgs boson.




Figure 2.2: Constraints on the Higgs mass due to measurements on the masses of the W boson and top quark.

2.2 Top Quark Production

Because of their high mass, top quarks are only produced at high energies. Pair production through strong interactions is the most prolific source of top quarks at the Tevatron, and is also expected at the LHC. The leading order diagrams are shown in figure 2.3.

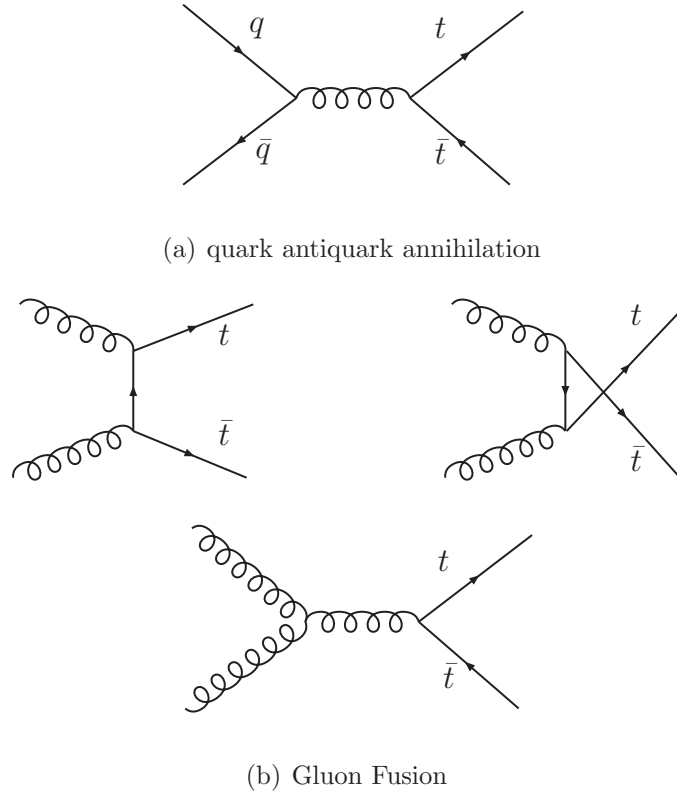


Figure 2.3: Leading order pair production of top quarks via (a) quark annihilation and (b) gluon fusion.

At the Tevatron, the quark annihilation process is dominant. As production is most likely to occur around the threshold energy, $2m_t$, the typical parton momentum fraction, x is given by

$$x \approx \frac{2m_t}{\sqrt{s}} \quad (2.1)$$

assuming the energies of the two interacting partons are roughly equal [9]. For the Tevatron, $\sqrt{s} = 1.96$ TeV which corresponds to $x \approx 0.18$. The Parton Density Function (PDF) shown in figure 2.4 shows that this value of x lies in a region dominated by the valence quarks [10].




Figure 2.4: Parton Density Function for $Q^2 = 175 \text{ GeV}^2$. At high x , the valence quarks dominate.

At the LHC $x \approx 0.03$, so gluon fusion is expected to make a 90% contribution to $t\bar{t}$ production [11].

2.3 Electroweak Production

Although production via the strong interaction is predicted to be the largest source of top quarks at the LHC, electroweak production is also possible [12, 13]. Three channels are predicted to be observable at the LHC, shown in figure 2.5.




Figure 2.5: Single Top production.

The highest contribution to the production cross section comes from the t-channel process, which is characterised by the top quark and a high momentum spectator quark in the final state. Even though it is an electroweak process, the cross section is expected to be as high as 256 pb at 14 TeV, which is roughly a third of the strong production cross section [14]. There are a number of reasons as to why the single top t-channel and $t\bar{t}$ cross sections are of the same order:

- Single top production is kinematically favoured because only one top mass must be produced, allowing partons with much lower x to take part in the production.

- Strong production is colour suppressed as the initial partons can only have specific colour combinations.
- The t-channel cross section scales as $(1/M_W^2)^2$, whereas the gluon fusion channels scale as $1/\hat{s}$.

The Wt associated production channel is also predicted to have a significant cross section at the LHC, with a predicted cross section of 62 pb at 14 TeV [15]. This channel is not observable at the Tevatron, due to the low b quark density. The smallest contribution will come from the s-channel process, with a predicted cross section of 10 pb [14].

The first direct observation of these modes of production was published by the $D\emptyset$ collaboration in 2009 [16]. The observation is statistically limited, so more data will be required before the different production channels can be distinguished.

The ability to measure the individual production cross sections is important because each channel is sensitive to different predictions of BSM processes. For example, if a fourth generation of quarks were to exist, the t-channel cross section would be expected to be enhanced. Equally, Flavour Changing Neutral Currents, which are forbidden at tree level in the Standard Model, would increase both the s and t-channel cross sections. Predicted cross sections for different BSM theories are shown in figure 2.6 [17].

In addition to BSM physics, single top production offers the only method for making a direct measurement of the $|V_{tb}|$ CKM matrix element. Attempts have been made to measure the ratio




Figure 2.6: Single top t and s -channel cross sections at $\sqrt{s} = 14$ TeV, taking into account different BSM processes. The solid black circle shows the Standard Model prediction. The pink and blue crosses show how top-flavour and top-pion models respectively change σ_s . The red asterisk and green circle show how a fourth generation and FCNC would affect σ_t [17].

$$R = \frac{Br(t \rightarrow Wb)}{Br(t \rightarrow Wq)} = \frac{|V_{tb}|^2}{|V_{td}|^2 + |V_{ts}|^2 + |V_{tb}|^2} \quad (2.2)$$

However, these analyses either assume the unitarity of the CKM matrix and become insensitive to new physics (such as a fourth generation of quarks) or they make no assumptions at the cost of not being able to measure the absolute magnitude of $|V_{tb}|$. These studies have shown that $|V_{tb}| \gg |V_{ts}|$ and $|V_{td}|$ [1].

Alternatively, by measuring the single top production cross section, the CDF and DØ collaborations have both published direct measurements on $|V_{tb}|$ [18, 16].

2.4 Top Quark Decay

The top quark has a lifetime of approximately 5×10^{-25} s [9]. This is an order of magnitude smaller than the characteristic hadronisation time scale ($\Lambda_{QCD}^{-1} \approx 3 \times 10^{-24}$ s).

An interesting consequence of its short lifetime is that the probability of gluon radiation is very small, which means that there is very little chance of spin flip before the quark decays [19]. In single top production this can lead to highly polarised top quark samples, where the top spin is predicted to be aligned with the direction of the d quark in the top rest frame (figure 2.7). Note that for the single top t-channel, the direction of the d quark corresponds to the spectator quark in roughly 75% of cases [20].

This spin polarisation gives rise to calculable angular distributions of the top decay products. For a spin up t quark, the angle between a decay product, i , and the d quark as measured in the top rest frame, is denoted χ_i^t , and is given by




Figure 2.7: Single top spin correlations. (a) Single top quarks are polarised in the direction of motion of the d -type quark, which may either be one of the initial state quarks or (in the case of the t -channel) the spectator quark. (b) The angle between this axis and a decay product in the top rest frame is denoted χ_i^t .

$$\frac{1}{\Gamma_T} \frac{d\Gamma}{d(\cos\chi_i^t)} = \frac{1}{2} (1 + \alpha_i \cos\chi_i^t) \quad (2.3)$$

where the correlation coefficient α_i is equal to 1 for antilepton coming from the W . There are weaker anticorrelation coefficients of ≈ -0.324 for the neutrino and ≈ -0.403 for the b -quark. When the W decays hadronically, the value of α_l is also predicted for $\alpha_{\bar{d}}$ and $\alpha_{\bar{s}}$. Equally, α_u and α_c are predicted to have the same value as α_ν [19]. Figure 2.8 shows these correlations.

Measuring these correlations will be a good test of the $V - A$ theory, which is the source of the predictions for α_i . They may also be a powerful discriminant between single top and other background processes.




Figure 2.8: Spin Correlations in Single Top Decay Products.

2.5 Conclusion

Although the top quark was first observed 14 years ago, some of the studies remain statistically limited. As the LHC is expected to produce 8 million top pairs a year at a luminosity of $10^{33} \text{ cm}^{-2}\text{s}^{-1}$, precision measurements are expected to become limited by systematic errors very quickly.

One interesting mode recently observed at the Tevatron experiments is electroweak production of the single top. Studying this channel may offer insight into BSM physics channels. An important milestone for the ATLAS experiment will be to re-establish evidence of this signal at the LHC.

Chapter 3

The LHC and the ATLAS

Detector

The Large Hadron Collider (LHC) [22] is an accelerator with a 27 km circumference based at CERN on the Swiss border. It has been designed to collide bunches of protons together with a centre of mass energy of 14 TeV. The accelerator will provide a nominal luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, delivering roughly 100 fb^{-1} of data per year.

Protons emerge from the CERN linear accelerator at 50 MeV (figure 3.1). They then pass through a number of synchrotrons designed to boost their energy by at least an order of magnitude per machine before being injected into the LHC at 450 GeV. The LHC then uses a number of RF cavities to accelerate two beams of protons to energies of 7 TeV in opposite directions. The bunches of protons are separated in time by 25 ns and are guided around the LHC ring by more than 1000 superconducting dipole magnets.

There are 4 interaction points around the LHC ring. The proton beams will be

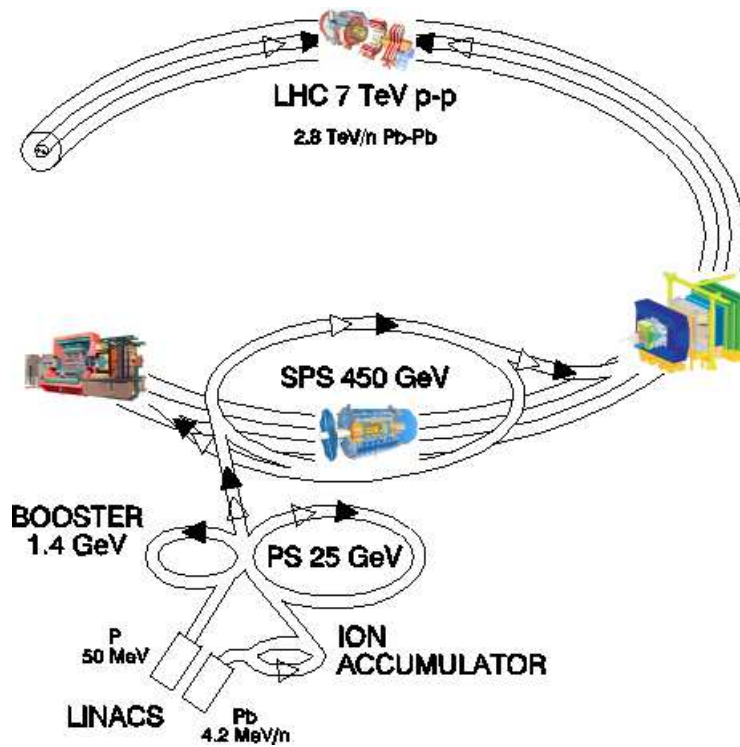


Figure 3.1: The LHC, and associated accelerators at CERN. Also shown are the relative positions of the four main detectors (starting from the bottom and working clockwise) ATLAS, ALICE, CMS and LHCb.

crossed at these points, colliding bunches at a rate of 40 MHz. The LHCb detector is a low luminosity experiment located at one of these interaction points and has been designed to investigate B-physics. The CMS and ATLAS detectors are designed to operate at peak luminosity. Both experiments have wide physics programmes.

The LHC will replace proton-proton collisions for one month a year with heavy ion (specifically lead) collisions. These will be studied in detail by the ALICE detector, located at the final interaction point. Heavy ion collisions allow for a different physics programme to be pursued, studying the very high particle multiplicities expected and searching for evidence of quark-gluon plasmas.

3.1 The ATLAS Detector

The ATLAS detector [23] has been designed to fulfil the needs of a wide physics programme. It will perform measurements of known physics in the new energy regime whilst looking for evidence of new physics beyond the standard model. Because of the wide programme, ATLAS is required to:

- identify muons and record their momenta accurately over several orders of magnitude
- identify electrons and photons with a low misidentification rate
- measure both electromagnetic and hadronic energy deposits with a high accuracy and maintain a high acceptance in the calorimetry so that missing transverse energy may be recorded
- record with high precision track information so that secondary vertices may be observed

These aims are realised through a number of subdetectors which make up ATLAS, as shown in figure 3.2.




Figure 3.2: The ATLAS detector. The inner trackers are shown in red. the electromagnetic calorimeter in orange, the hadronic calorimeter in green and the muon spectrometer in blue. The detector is 44m long and 25m in diameter.

3.2 Tracking

High resolution tracking is an important requirement for the ATLAS detector because this determines the detector's ability to measure secondary vertices. As the tracking detectors also sit within a 2T magnetic field produced by the central solenoid, momentum measurements for charged particles are also possible, based on the curvature of the reconstructed tracks. It is possible to reconstruct tracks of particles with $P_T > 0.5$ GeV in this manner.

The tracking detectors consist of three main subdetectors (figure 3.3). The pixel detectors sit closest to the beam pipe. They consist of 3 layers of silicon detectors, covering the pseudo-rapidity¹ region of $|\eta| < 2.5$. Each pixel is $50 \times 400 \mu\text{m}^2$, but offers a resolution of approximately $10 \mu\text{m}$ in ϕ and $115 \mu\text{m}$ in z . Because the inner layer is so close to the beam pipe (5 cm), it is expected that it will be replaced after 3 years of running.




Figure 3.3: The Inner Tracking detectors.

The pixel layers will also play an important role in identifying secondary vertices.

¹ ϕ is defined as the angle between the particle track and the y-axis in the x/y plane. η is defined as $-\ln(\tan\frac{\theta}{2})$, where θ is the angle between the track and the z-axis in the y/z plane.

These displaced vertices may be indicative of particle decay, such as B hadrons, and accurately identifying them will help discriminate between different types of objects.

The middle layers of the Inner Detector are made up of the Silicon Microstrip Trackers (SCT). The SCT modules are strip detectors, arranged into 4 layers in the barrel and end-cap, extending over the same η region as the pixel detectors. Each module consists of two wafers of micro strip sensors, glued back to back with a relative angle of 40 mrad. Each wafer contains more than 750 strips of active semiconducting material, where each strip is separated by a distance of 80 μm . The stereo angle between the wafers allows the SCT modules to make a track measurement in z as well as $R - \phi$ (though the resolution is over 30 times worse in the z direction).

Finally, the outer most sub-system of the Inner Detector is the Transition Radiation Tracker (TRT). This consists of large number of straw detectors, aligned with the beam pipe in the barrel and radially in the end-caps. Each straw is 4 mm in diameter and 1440 mm in length and covers the region $|\eta| < 2.0$. Due to the high volume of straws used, it will provide 36 hits per track on average, with a resolution of 130 μm per straw in ϕ .

In addition, the TRT will also aid in the identification of electrons. When a charged particle crosses a boundary between two dielectrics it radiates. Photons emitted by particles as they pass through the TRT are absorbed by the xenon gas, resulting in a larger signal amplitude than in the case where no transition radiation is produced. Because the number of photons emitted is sensitive to the Lorentz boost, the strength of the signal may be used to discriminate between different masses for a particle of known energy.

The whole Inner Detector resides inside a super conducting solenoid, which provides

a magnetic field with a strength of 2T. The solenoid operates at a temperature of 4.5K. Because it sits between the interaction point and the calorimeter, the solenoid has been carefully designed so as to minimise the amount of material particles must pass through.

3.3 Calorimetry

The ATLAS physics programme requires very good electromagnetic (EM) calorimetry to aid in the identification of photons and electrons. In order to measure jet energies and make an estimate of the missing transverse energy (\cancel{E}_T), the hadronic calorimeters are required to have as wide an acceptance as possible. The different calorimeter subdetectors are shown in figure 3.4.




Figure 3.4: The Electromagnetic and Hadronic Calorimeters

3.3.1 Electromagnetic Calorimetry

The EM calorimeter is a sampling calorimeter with Liquid Argon (LAr) active region interspersed with lead absorbers. In the region $|\eta| < 2.5$, high granularity strip cells are available, to complement information from the trackers. Cells in the inner layer have the finest divisions in η , enabling a precision measurement. This can help separate electrons from other objects, such as neutral pions decaying to photons. These cells are combined with cells from another two layers, of typical size 0.025×0.025 in $\phi \times \eta$ to form modules which are at least 22 radiation lengths (X_0) deep.

The EM Calorimeter extends into the end-caps to $|\eta| = 3.2$, with the active thickness increasing to $36 X_0$. Calorimetry in the region $3.1 < |\eta| < 4.9$ is performed by the Forward Calorimeters (FCal). The FCal consists of three modules, of which the first is designed for electromagnetic calorimetry. It contains a LAr active region, but the lead absorbers are replaced by copper, which is better suited to working in the high radiation forward region.

In addition, there is an extra layer of liquid argon between the central solenoid and the start of the EM calorimeter over the region $|\eta| < 1.8$. This presampler is designed to detect EM showers emerging from the solenoid, in an attempt to correct for material energy losses before particles enter the calorimeter proper.

3.3.2 Hadronic Calorimetry

The hadronic calorimeter surrounds the EM calorimeter and is split into three pseudorapidity regions. The tile calorimeter extends over the region $|\eta| < 1.7$ in the barrel and end-caps. It consists of plastic scintillating tiles distributed between steel

absorbers. The tiles are arranged in wedge-shaped blocks, representing a single calorimeter module. Photons produced in the scintillators are read out by fibres into photomultiplier tubes.

The Hadronic End-cap Calorimeters (HEC) cover the region $1.5 < |\eta| < 3.2$. These consist of copper and liquid argon sampling calorimeters. Finally, the FCal contains two more modules to aid in calorimetry between $3.1 < |\eta| < 4.9$. These modules consist of a liquid argon active region with tungsten absorber, which is better at containing hadronic showers.

3.4 Muon Spectrometry

ATLAS is required to detect charged particles which penetrate the calorimeter with a high degree of accuracy. The muon spectrometers make an estimate of the momentum of charged particles based on a number of high precision position measurements as they move through an 4T magnetic field. This field is provided by superconducting air core toroidal magnets - 8 coils in the barrel region and 8 in each of the end-caps.

Two types of technology have been used to track particles through the muon spectrometer (figure 3.5). The Monitored Drift Tubes (MDTs) are 30 mm in diameter and between 1 and 6 meters in length, depending on their location within the detector. They give tracking information in the barrel and outer layers of the end-cap, providing a z measurement in the pseudorapidity range $|\eta| < 2.7$. The drift tubes are arranged in layers, typically 3 tubes deep. An MDT chamber consists of a pair of tube layers, separated by a spacer of varying height, again depending on the location of the chamber within ATLAS. More than 1000 chambers have been used in




Figure 3.5: The Muon Detectors and Toroidal Magnets.

the detector, providing typically 20 measurements per muon.

In the forward regions, a different technology has been used. The Cathode-Strip Chambers (CSCs) are proportional chambers, which provide both radial and azimuthal measurements in the pseudorapidity region of $2 < |\eta| < 2.7$. In this region the muon track density is predicted to be particularly high, so the CSC's ability to measure R and ϕ simultaneously will greatly benefit the track reconstruction algorithms. Accuracy is expected to be limited by multiple scattering.

In order to make a spatial measurement, the MDTs and CSCs must be timed in to dedicated triggers which signal the arrival of a particle. In the barrel region, $|\eta| < 1.05$, Resistive Plate Chambers (RPCs) are used. In the end-cap, $1.05 < |\eta| < 2.4$, Thin Gap Chambers (TGCs) are used. Both technologies provide a ϕ measurement and either z (RPC) or R (TGC) measurement. These spatial measurements have a lower resolution than those performed by the MDTs and CSCs. However, they have

a timing resolution better than the LHC clock frequency, making them suitable for muon triggering and beam crossing tagging.

3.5 Triggering and Data Acquisition

The LHC will collide protons at a rate of 40 MHz in the ATLAS detector. As the offline systems will be able to cope with an event rate of 200 Hz, ATLAS is equipped with a trigger for filtering off events suitable for permanent storage. These events will share some common characteristics, such as isolated leptons or high energy jets, which may be indicators of a significant physics event.

The ATLAS trigger consists of three levels (figure 3.6). The Level-1 trigger [24] is responsible for making an initial decision on whether to process further an event. It takes information from the calorimeters and the muon chambers, at a reduced granularity, and looks for key signatures such as isolated energy clusters. If an event is accepted, it is passed on to the Level-2 trigger. Until Level-1 has made its decision, data must be stored in the memory of each subdetector. In order to keep the pipeline memories short, Level-1 must make a decision as quickly as possible, and so has been implemented in custom hardware and designed with a latency of less than $2.5 \mu\text{s}$.

Events are expected to pass through the first level trigger at a rate of up to 75 kHz. Once accepted, data from potentially interesting regions is read out at full granularity and passed to the level-2 trigger. This processes the data, along with information from the trackers, and accepts events at a rate of 3.5 kHz. The final layer

²Although Level-1 is expected to take $2 \mu\text{s}$ to make a decision, it processes collision data in parallel, so it will accept or reject events every 25 ns

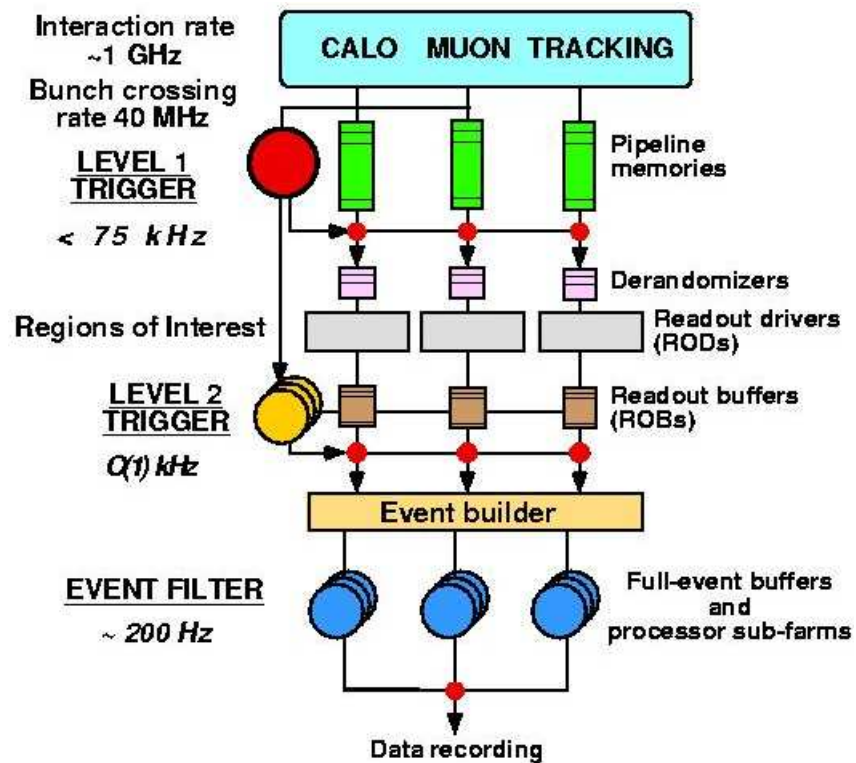


Figure 3.6: The ATLAS Trigger. Each of the three layers is responsible for reducing the event rate down to a manageable rate for the next level. Events passing the Event Filter are sent for permanent storage and further processing on the Grid.

of the trigger is the event filter, which consists of a dedicated CPU farm running the offline analysis software (Athena). The event filter reduces the event rate to 200 Hz, taking approximately 4 seconds to fully process each event. From here events are sent for permanent storage on the Grid, where they will be processed further.

The trigger belongs to the wider Data Acquisition system (TDAQ), which is responsible not only for triggering, but data readout and distribution of timing signals. TDAQ is also responsible for the Detector Control System (DCS) which offers an interface for operating the ATLAS detector.

3.6 Luminosity Measurement

Measuring the luminosity of interactions for the purposes of physics analysis is the responsibility of the experiments. ATLAS intends to use a number of methods to make this measurement. The LUCID detectors operate in the very forward regions of ATLAS (± 17 m). They consist of a number of Cerenkov tubes surrounding the beam pipe, and have been designed to detect the number of inelastic proton scatters in this region. This will allow for a relative luminosity measurement to be made.

Absolute luminosity will be measured by the ALFA detectors, located at ± 240 m on either side of the interaction point. Scintillation fibres mounted on Roman Pots will be used to measure the elastic scattering rate at very small angles ($11 < |\eta| < 13.5$) via the Optical Theorem, which can be used to extract the luminosity. Because the ALFA detectors must be so close to the beam line, they can only be used during low luminosity ($10^{28} \text{ cm}^{-2} \text{ s}^{-1}$) runs.

During the early collisions the ALFA detectors will not be available. Instead, LUCID

will be calibrated from beam parameters provided by the LHC. This is expected to provide an initial luminosity measurement with an accuracy of 20%. During later running, it is expected that the production rate of W and Z bosons may be used to estimate the luminosity. This will require good knowledge of the production cross sections.

3.7 Conclusion

The ATLAS detector is a large machine dedicated to the search for new phenomena and precision measurements at high luminosities. It has been designed with the aim of studying a wide range of physics topics.

At the time of writing, the LHC is expected to be ready for beam again by late 2009. Initially, beam energy will be limited to 3.5 TeV. It is expected that the beam energy will be increased to 5 TeV and a peak luminosity of $\approx 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. At the end of the first year, approximately 200 pb^{-1} of data may have been recorded. At this point, ATLAS will be able to make competitive measurements in some channels when compared to other accelerator experiments, such as CDF and DØ at the Tevatron [25].

Chapter 4

The Detector Control System

4.1 Controlling ATLAS

The ATLAS experiment consists of a large volume of hardware which must be operated in a coherent and safe manner in order to record data from the LHC collisions. Because of the physical size of the detector, and because of the harsh working environment, the capability to operate and monitor the detector hardware remotely is essential. This requirement is met by the Detector Control System (DCS). The DCS is responsible for monitoring the detector and ancillary systems, such as power supply and cooling.

The DCS is responsible for providing and maintaining a homogeneous interface to the detector. Once commissioned, the detector will be operated by shifters who may have limited experience with some subsystems. Providing a clear and consistent method of control will aid the smooth operation of the detector during data taking runs.

The DCS is also responsible for recording and storing hardware monitoring data. This data may be useful for diagnosing problems, but it may also be relevant to the offline analysis.

Finally, the DCS can also flag potential problems with the detector hardware so that action, whether it be manual or automatically triggered, may be taken to avert further difficulties [28].

4.1.1 Detector Safety System

It is important to stress the limit of the DCS remit. The DCS is not responsible for human or machine safety. These responsibilities belong to the Detector Safety System (DSS) and the CERN Safety and Alarm Monitoring (CSAM).

Abnormal events in ATLAS are classified according to their severity. Level 1 Alarms have the lowest severity and might consist, for example, of a high temperature reading or a fault with a fan. These cases, where neither human nor machine safety is at risk, are the concern of the DCS. The DCS may at this point take action in an attempt to stop the problem from becoming more severe.

Level 2 Alarms represent a more serious risk, such as a cooling failure or a water leak. These are the concern of the DSS which defines protocols and actions to deal with these events. Level 3 Alarms are the most severe, and include events such as fire or lack of oxygen. The DSS is responsible for these events, but it will also transmit the alarm to the CERN Fire Brigade and trigger an evacuation.

Because of its importance, the DSS remains independent from the DCS. Although information is shared between the two systems where relevant, the DCS is by design

forbidden from interfering with actions performed by the DSS.[29]

4.1.2 JCOP

The Joint Controls Project (JCOP) was set up at CERN to address common problems in controlling the LHC experiments. Although the experiments are very different, they must all employ some method of controlling and monitoring their hardware. The JCOP group sought to unify and coordinate efforts on common areas.[30]

One of the most important tasks of the JCOP group was to search for a suitable software package with which the experiments could develop their control projects. A large study was undertaken, and the PVSS-II package developed by ETM[31], was found to be the most suitable candidate.[32] PVSS has since been adopted by all of the LHC experiments.

PVSS works by defining a number of variables called datapoints, which are managed through an internal database. Datapoints are simple software constructs and can take any number of types (eg integer, floating point number, string). Their online values can be set, retrieved and manipulated by a scripting layer, allowing users to monitor values and display them in Graphical User Interfaces (GUI). Crucially, a datapoint may be configured to read its online value from a hardware source, making PVSS a suitable environment for developing monitoring applications.

The JCOP group are also responsible for developing some software components for use within PVSS. These include support for common items of hardware, such as the Wiener VME crate (figure 4.1) and libraries for configuring alarms and archiving. These components have been built on by other groups. For example, using tools produced by the JCOP group, the Central ATLAS DCS group developed a framework

for defining a Finite State Machine (FSM) to represent the ATLAS hardware.[33]

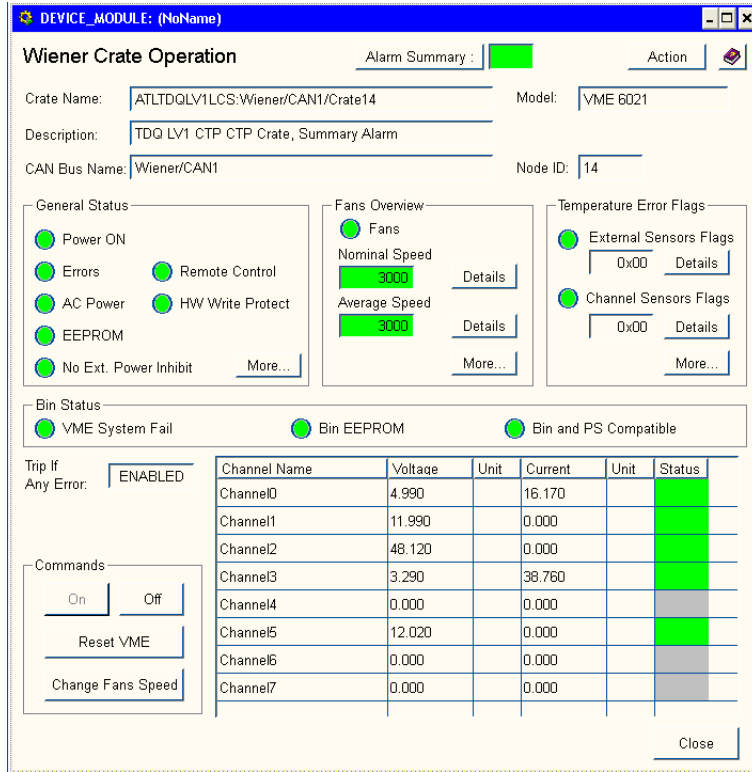


Figure 4.1: The JCOP panel used to display information related to a standard Wiener VME crate. This panel, along with other JCOP tools, may be built into a wider, experiment specific interface.

4.2 The ATLAS DCS Structure

The ATLAS DCS follows a tree-like structure (figure 4.2). The lowest layer consists of the detector hardware and is referred to as the front-end (FE). Higher levels consist of networked PCs which are responsible for processing data from the front-end. These layers are collectively referred to as the back-end (BE).

It is important to note that monitoring data and commands are transferred only in the vertical direction on the DCS tree. There can be no direct transverse communi-

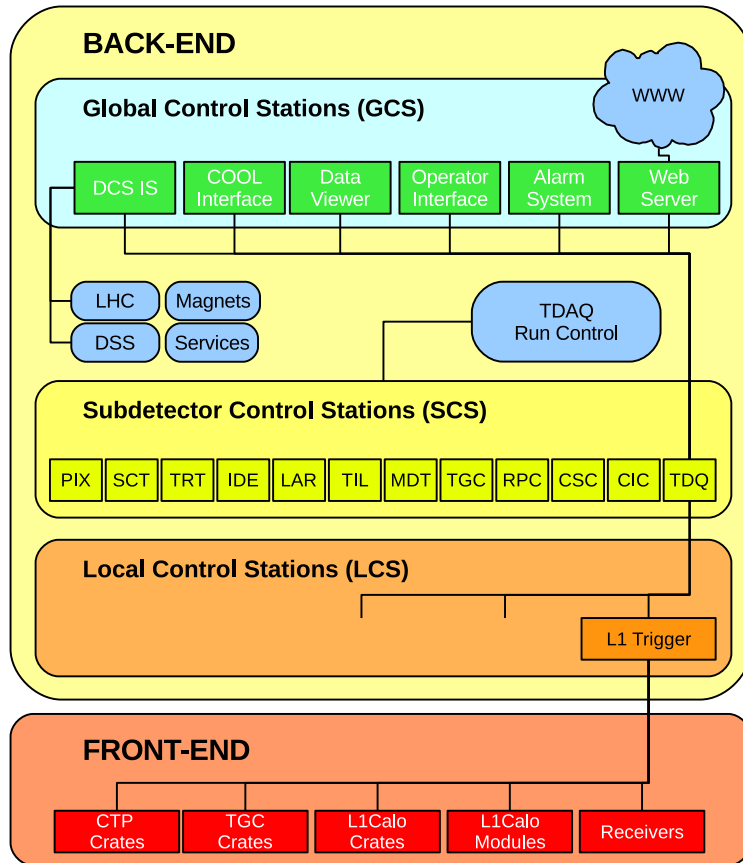


Figure 4.2: Overview of the ATLAS DCS. The system is divided into two sections - the front-end, which consists of the monitored hardware, and the back-end, which consists of a number of networked PCs. The entire system is modular, enabling sub-systems to be partitioned and placed under the control of expert users.

cation between nodes in the same layer. Also, every node is limited to have exactly one parent. Taken together, these constraints make it possible to partition control of ATLAS. For example, it may be desirable for expert users in the Trigger group to have exclusive control over the trigger hardware for the purposes of debugging whilst leaving the rest of the detector under the control of the central ATLAS DCS. The partitioning mechanism makes it impossible for the two groups of controllers to issue commands unintentionally to the same hardware.[28]

4.3 The DCS Front-End

The DCS front-end refers to the system monitoring hardware distributed on and around the ATLAS detector. These systems are designed to provide co-ordinated remote control over the detector hardware. They are also responsible for reading out monitoring data, making it available to the back-end. The DCS front-end can also refer to the control systems associated with ATLAS infrastructure, such as the monitoring of electronics racks, cooling and environment.[34]

Most front-end systems are connected to the DCS back-end by a number of fieldbus types, CANbus (Controller Area Network) being the most commonly used. This is a general purpose, industry standard for communicating with microprocessors responsible for system control and automation.[35]

4.4 The DCS Back-End

The ATLAS DCS back-end consists of a large number of networked, rack-mounted PCs running PVSS. The exact specification of the PC depends on its role, but all machines are required to be stable and robust. The PCs are organised into a tree structure split into three levels (figure 4.2).

4.4.1 Local Control Stations

The Local Control Stations (LCS) represent the interface between the detector hardware and the DCS. These PCs are responsible for receiving monitoring data and making it available to PVSS. As such, they are required to have good I/O capa-

bility. They have 3 PCI slots, which are capable of housing up to 12 CANbus interfaces.

The LCS are also responsible for the lowest levels of the FSM (section 4.4.4). These PCs interpret monitoring data and define an operational state. State information from many hardware sources are collated and summarised by leaf nodes in the FSM (ie the nodes at the bottom of the FSM tree) and then passed further up the tree. Requests coming from higher FSM nodes will ultimately arrive at the LCS, and be converted into low-level commands to be sent to the hardware.[33]

4.4.2 Subdetector Control Stations

The Subdetector Control Stations (SCS) represent entire subdetectors, such as the SCT or the LAr Calorimeter. Because of this, the SCS represent the highest level at which the ATLAS detector controls may be partitioned¹. They are responsible for summarising the hardware operational states reported by connected LCS.[28]

The SCS also represents the main interface between the DCS and the TDAQ. During a data taking run it may be desirable for TDAQ to request certain commands be sent to the hardware, or to receive a limited amount of DCS information related to the status of the hardware. This DAQ DCS Communication (DDC) takes place at the SCS level, and is implemented through tools provided by the Central ATLAS DCS team.[36]

¹The next level up requires taking control of the whole ATLAS detector, and this operation is exclusively reserved for the DCS desk in the ATLAS Control Room.

4.4.3 Global Control Stations

The Global Control Stations (GCS) perform a number of functions. They offer the main human interfaces to the DCS, through the ATLAS Control Room desk and the read-only web interface [37]. They are also responsible for a number of global services, such as managing the main Alarm System and interacting with the Information Server, through which the DCS communicates with external systems, such as the LHC and the DSS.[28]

4.4.4 Finite State Machine

Each PC in the back-end tree is responsible for running a Finite State Machine (FSM). The role of the FSM depends on its placement within the back-end (figure 4.3).

FSM leaf nodes on an LCS are called Device Units (DU) and represent an item of monitored hardware. DUs interpret monitoring data from the hardware and derive an FSM state. They are also responsible for converting user commands propagated from the FSM into a form which may be sent to the hardware.

Subsequent layers within the FSM consist of either Logical Units (LU) or Control Units(CU). Nodes of these types are responsible for summarising the states of their children. For example, in figure 4.3, Device Units representing hardware modules are connected to a Logical Unit. This LU derives a summary state based on the states of the module DUs. Logical Units represent abstract logical groupings within the detector, so in this example the children of the LU are all modules housed within the same crate of electronics.

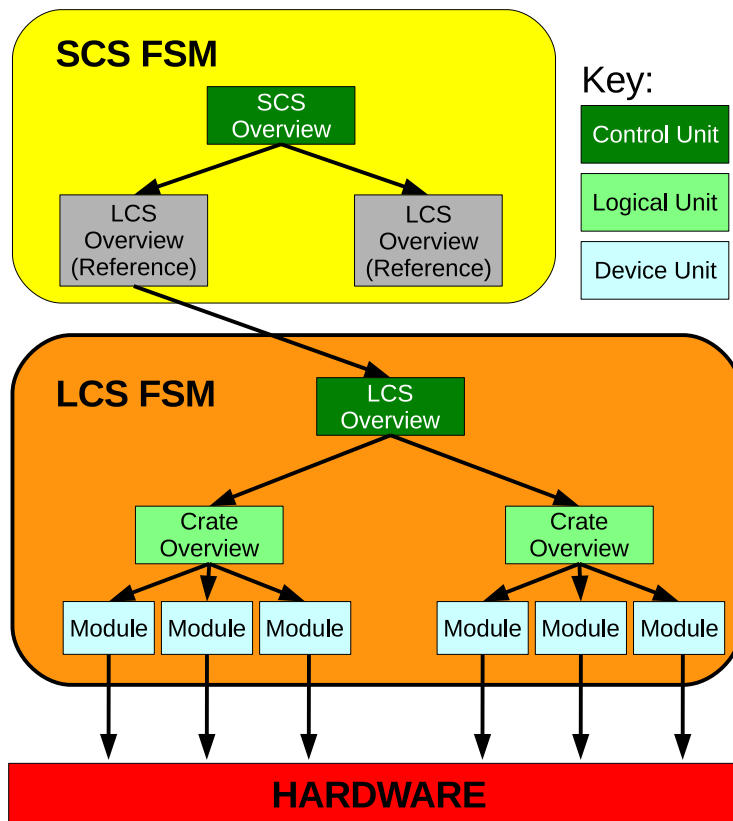


Figure 4.3: ATLAS FSM Structure. Each PC runs an FSM, which links into other remote FSMs running on other PCs. In this way data can flow up from the hardware to the top and commands can flow down. Users can take full and exclusive control of an FSM sub-tree only where a Control Unit has been instantiated.

Logical Units and Control Units are generally interchangeable, as they both perform the same summary roles. Control Units have added functionality for implementing the FSM partitioning system. A user may take control of a specific CU and its children, specifying how state and commands are propagated between partitions. The main use cases are shown below (figure 4.4). In (a), User 1 has exclusive control over the whole FSM - no other user may issue commands. This is the normal method by which the FSM is used, with ownership being held by the DCS operator in the ATLAS control room.[33]

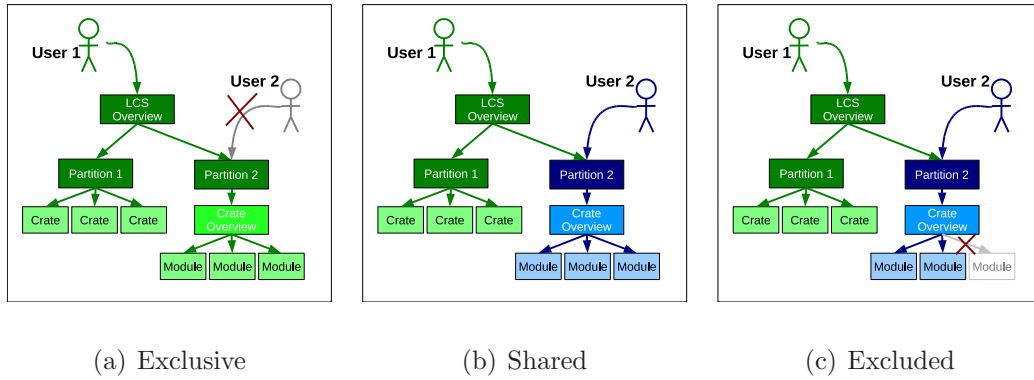


Figure 4.4: FSM Operational modes. In (a), User 1 has exclusive control over the entire tree, and no other users may issue commands. In (b), User 1 has shared a partition with the User 2, who may also issue commands to the shared partition. This could represent an intervention by an expert user to fix a known problem. In (c), a problem module has been excluded from the tree, no longer receiving commands or propagating state information.

Control of a partition may be shared, as shown in figure (b). In this case an expert user (User 2) may also issue commands to the shared partition, shown in blue. Finally, if a single node or partition is the cause of a persistent problem, it may be excluded as shown in figure (c). In this case, no state information for the excluded node is propagated up the tree and no commands will be issued. Note that in both (b) and (c) User 2 cannot interfere with the running of partition 1.

The ATLAS DCS FSM is unique amongst the LHC experiments in that each FSM node consists of two pieces of information - an FSM state and a status. The state represents the physical condition of the represented hardware. As a simple example, a node representing a module might define the states ON and OFF. The status represents the well-being of the hardware. In the case of the simple module, it might have a status OK if all data is within the expected limits. However, the status might change to WARNING if a channel moves beyond those limits (for example if a temperature reading got too high).

Four status levels are defined for the ATLAS FSM. In order of increasing severity they are:

- OK - The monitored object lies within expected parameters
- WARNING - The least severe status. Minor problem, to be dealt with within normal working hours.
- ERROR - A more serious fault. Could affect the performance of the detector
- FATAL - The most serious category. Problem seriously affects detector and should be dealt with immediately.

The state and the status elements for a given FSM node are in general not correlated. For example, it is possible for a module to be in the state ON with an OK status, or ON with a WARNING status, OFF with ERROR status etc. However, in practice, a particular state may restrict the available status options.

The FSMs running on different back-end PCs are linked by defining references to nodes in other trees. For example, the leaf nodes on an SCS FSM consist of references to the root nodes on LCS. Note that the vertical communication constraint is enforced again when linking FSMs, ensuring that the detector can be partitioned effectively.[33]

4.4.5 Alarms

PVSS provides a monitoring facility by which individual datapoints may be monitored and flagged in the case of a problem. A range of acceptable values can be set, and if the online value of the datapoint falls outside that range an alarm is raised.

The alarms from a specific system are then summarised in a GUI, as shown in figure 4.5.

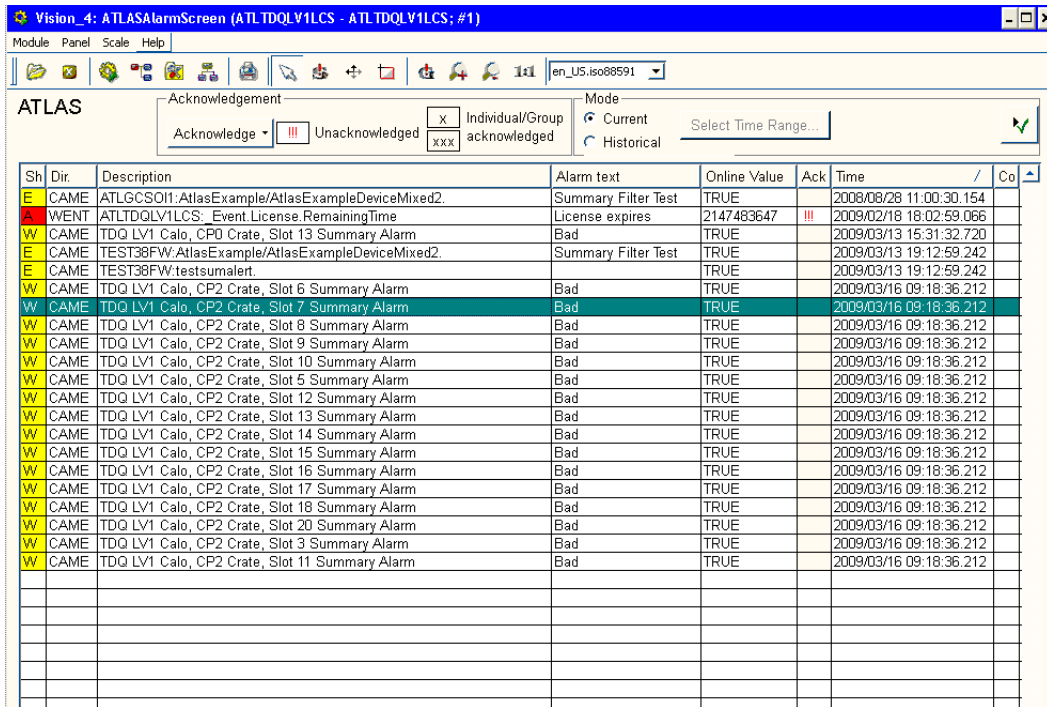


Figure 4.5: The ATLAS Alarm Screen. Alarms attached to specific datapoints are displayed on this screen. In this particular example, summary alarms for modules within the L1Calo Trigger system are displayed. The user may obtain more detailed information about an alarm by clicking on the panel.

ATLAS Alarms are categorised using the same levels as are used for the FSM status nodes (ie OK - FATAL). In general the FSM and the Alarm System are independent. In practice, status nodes and alarms usually derive their values from the same datapoints, so they will often display the same level (ie a status node and a particular alarm might both be in the WARNING state). However, the alarm system offers more detailed information about problems within the system.

PVSS also allows for the definition of Summary Alarms. These alarms are sensitive to the state of other alarms, taking on the value of the most severe. In this way,

alarms may be filtered so as not to flood the Alarm Screen. This does not invalidate the requirement that the Alarm Screen report detailed information about problems in the system, as the original alarms are recoverable from the summary alarm.[28]

4.4.6 Archiving

Persistent storage of datapoint values is achieved in PVSS by exporting data to a backend database. In the case of ATLAS, recorded datapoint values are stored in the ORACLE Online Database (figure 4.6). In order to minimise the amount of data transported to the database, value and time dependent smoothing is used to configure the conditions under which an online value is archived. For example, after a datapoint is archived, a dead-band is defined, requiring that the value of the datapoint change by more than the dead-band before it is archived again. However, this condition is only valid for a fixed time interval after the initial archive. After the time interval, any change in the datapoint value will trigger it to be archived.

Due to security considerations, this database is not accessible from outside of the ATLAS Control Network. However, it is often desirable that external institutes have access to archived data for the purposes of detailed monitoring and debugging. To this end, data is replicated to the ORACLE Offline Database, which is externally accessible, at regular intervals. Any DCS information which is required for offline calibration, reconstruction and analysis is copied to the Conditions Database.[28]

Finally, there is also a configuration database, which feeds different settings into PVSS.

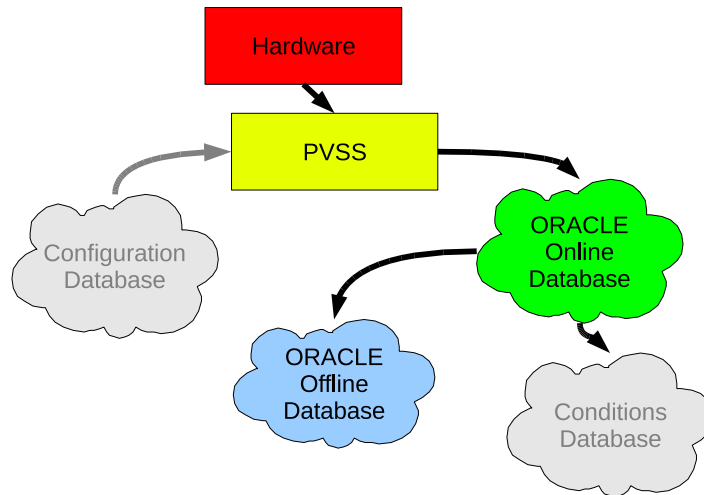


Figure 4.6: Overview of the ATLAS DCS Databases. Hardware delivers monitoring data to PVSS. Appropriate datapoints are marked for persistent storage, causing their values to be stored in the ORACLE Online Database. Values are replicated to the Offline Database to enable off-site access. A subset of the data, relevant to offline analysis, is also copied to the conditions database.

4.4.7 Security

Security is a very important consideration of the DCS project. It is essential that the integrity of the detector be maintained and kept secure from both external threats and inexperienced users.

The ATLAS controls network is isolated from the CERN general purpose network (GPN), so a breach of the GPN does not necessarily expose the ATLAS hardware. A gateway between the two networks does exist, but access is granted only to approved CERN users².

²This is a compromise; it allows some users access the controls network, but in doing so it relies on the strength of their CERN Computing Account credentials. A compromised user account may still pose a real risk to the machine.

In terms of DCS machines (LCS etc), specific users must be granted permission to access specific machines. Permission is usually limited only to users who absolutely need it, such as system experts. Shift operators will not normally have access, or require access, to these machines.

The DCS software maintains another level of Access Control (AC) through the use of the JCOP AC Framework. Different subsystems are defined to exist within specific domains. For example, the Level-1 Trigger forms part of the TDQ³ AC Domain. Different actions within a domain require different privileges. For example, operator privileges might be required to power on a VME crate, whereas expert user privileges would be required to deactivate alarms.

A set of privileges for a particular domain is referred to as a Role. Individual users are assigned to groups, each group having a different set of privileges corresponding to different roles. Users will only be allowed to execute actions within the domains for which they have the appropriate permissions. For example, a user with expert privileges in the TDQ domain, may only have observer privileges in the Pixel Detector domain.[28]

4.5 Level-1 Trigger DCS

The main focus of the subsequent chapters is on the DCS for the Level-1 Trigger[24]. The role of the Level-1 Trigger is to identify signatures in the detector which might be an indication of rare physics processes. It has been designed to select events at a maximum rate of 75 kHz, a factor of 500 smaller than the bunch crossing rate.

³Historical convention: In ATLAS the Trigger and Data Acquisition is generally referred to as TDAQ. In the DCS, it is referred to as TDQ

In order to make a decision within the fixed latency, the Level-1 Trigger analyses reduced granularity data from the detector. Figure 4.7 shows the main components of the Level-1 Trigger.

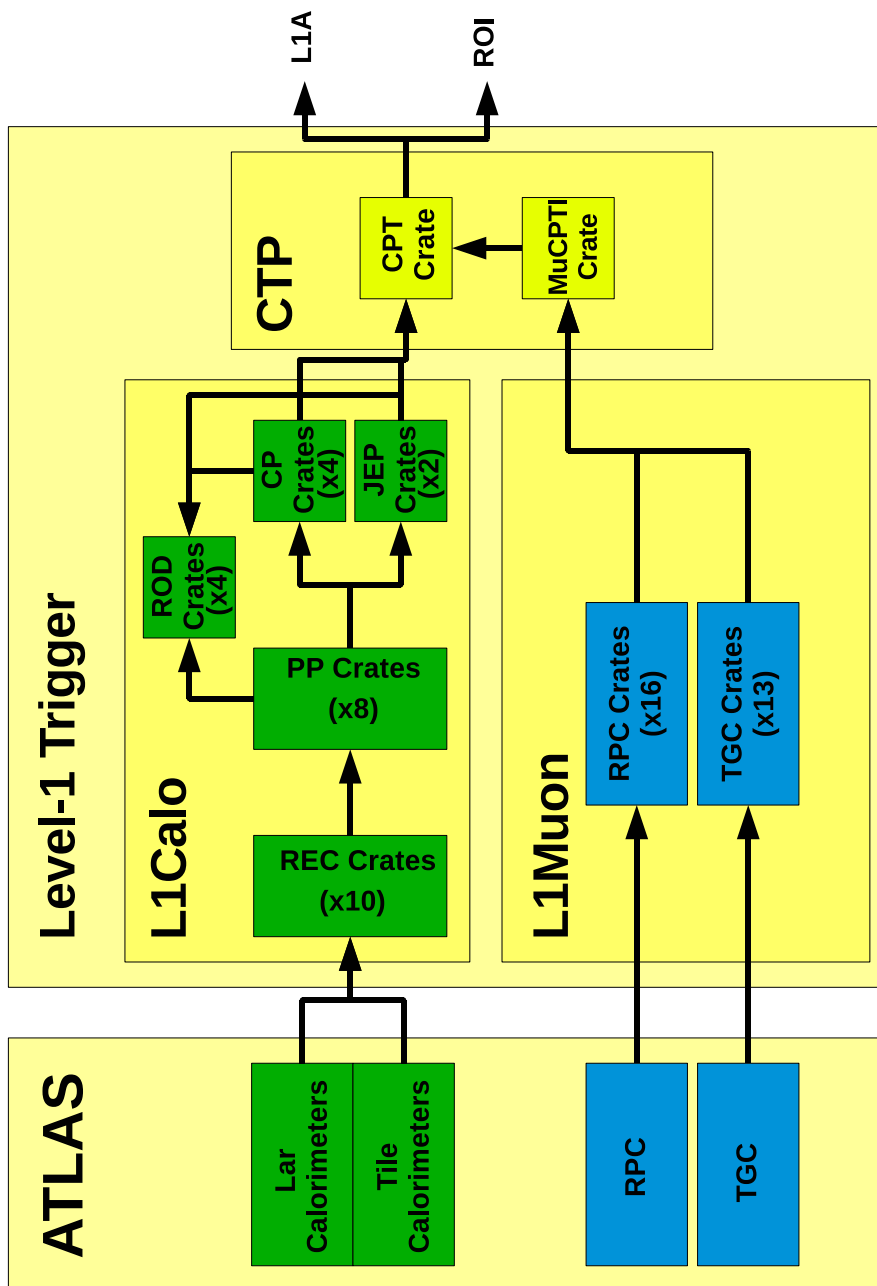


Figure 4.7: Level-1 Trigger

Analogue calorimeter signals are received by the receiver crates, which then pass them into the Level-1 Calorimeter Trigger (L1Calo). The signals are first digitised by the Preprocessors (PP), which then distribute the signals into the Cluster Processor (CP) and Jet/Energy Processor (JEP) crates. The CP crates search for signatures corresponding to isolated electrons, photons and taus. The JEP crates search for high energy jets, and also perform a missing energy calculation. Data from all three types of crate is recorded by the Read Out Drivers (RODs), which are housed in separate crates.

The number of objects passing the defined trigger thresholds in the CP and JEP systems is passed to the Central Trigger Processor (CTP), which receives similar information from the TGC and RPC triggers. The CTP then broadcasts the Level-1 Trigger decision to the rest of the ATLAS readout.

In terms of the DCS, the Level-1 Trigger consists of a larger number of VME crates, which must be modelled in the FSM. In addition, a large number of modules in the L1Calo crates must also be monitored and modelled. The focus of the subsequent chapters is how the DCS has been implemented, mainly at the LCS level, for the L1 Trigger. Chapter 5 considers the front-end and how monitoring data is collected and broadcast by the hardware via a CANbus. Chapter 6 is mainly concerned with how the back-end deals with the monitoring data through the execution of an FSM.

Chapter 5

Level-1 Trigger DCS Front-End

The Level-1 Trigger DCS is responsible for monitoring and controlling the trigger hardware. For the most part, this consists of offering remote access to the VME crates housing the trigger electronics. In the case of the L1Calo Trigger though, individual modules are also monitored.

5.1 Project Requirements and Scope

The Trigger DCS is expected to:

- Offer remote access to crates, allowing them to be powered on and off
- Report the state of the hardware
- Take automatic and appropriate action in the case of problems

These requirements must be met in a robust and reliable manner as it is important that the trigger hardware remains stable during a run.[24]

5.2 Monitored VME Crates

The TGC, CTP and L1Calo subdetector trigger crates are all monitored by the L1 Trigger DCS¹. Although not part of the Level-1 Trigger, the Receiver Crates are also monitored by the L1 Trigger DCS.

The trigger crates are standard VME crates, manufactured by the German company W-IE-NE-R Plein & Baus GmbH.[40] Each subdetector has a CANbus into which its crates are connected (see figure 5.1). As far as possible, each subdetector's CANbus is kept separate and independent of the other CANbuses. In this way, problems on one bus will not interfere with the running of another.




Figure 5.1: The L1 Trigger Subdetector Wiener crates, spread over 4 CAN buses. Note that the other CAN buses (CAN 0, 3, 5 and 7) are reserved for future additions to the system.

The VME crates monitor various channels, of which approximately 20 are actively used by the L1 Trigger DCS. The most important channels are:

¹At the time of writing, the RPC crates had not been connected.

- Power status
- Current and Voltage information on active channels
- Temperature
- Fan status

The L1 Trigger Wiener crates are polled every 30 seconds for this data. Commands are also sent over the CANbus to individual crates, requesting that they power up or down etc. Operating intervals for the temperature, current and voltage channels are stored on-board, and are set via the crate front panel. If a channel moves outside these limits, or if the fans fail, the crate will automatically power off. This happens instantly and requires no external action from the DCS.

5.3 L1Calo Trigger Modules

In addition to the VME crates, the L1Calo Trigger DCS also monitors individual modules within the crates. With the exception of the L1Calo Timing Trigger Control (TTC) crate², the L1Calo crates house between 10 and 19 monitored modules, depending on crate function. Each crate also contains a CPU module in the first slot and a Timing Control Module (TCM) in the last slot. The TCM is connected to the DCS, but does not monitor any data channels directly. The CPU is not connected to the DCS at all.

Modules within a crate communicate DCS information via a CANbus on the crate backplane (Figure 5.2). This bus connects all modules with the TCM. Although the

²The TTC crate houses no monitored modules.

TCM does not monitor any datapoints directly, it does act as a bridge between the crate CANbus and the external DCS. The 16 TCMs connect to an external CANbus, which then interfaces with the DCS back-end.

In total, 6 different types of module are monitored (see table 5.1), and each module type monitors a different number of channels (5950 in total across the whole system). Every connected module contains a Fujitsu MB90F594 microcontroller [41], which is responsible for reading monitoring data and broadcasting it on the crate CANbus.

<i>Module</i>	Type Code	No. Mods	No. Chnls/Mod	Info Monitored
PPM	1	144	24	8 Voltages 16 Temperatures
CPM	2	64	16	3 Voltages 10 Temperatures 3 Currents
JEM	3	32	42	28 Voltages 14 Temperatures
CMM	4	12	8	3 Voltages 5 Temperatures
ROD	6	20	32	6 Voltages 26 Temperatures
CAM	9	6	5	5 Voltages

Table 5.1: The 6 module types monitored by the L1Calo DCS. The total number and type of channels varies between module types. The type code is used in CAN messages to identify the type of module present.

In terms of requirements, the DCS for the L1Calo modules should:

- Offer a uniform implementation for all module types
- Readout up to 42 datapoints per module, with 8 bit precision
- Contain a fast feedback loop, such that modules can request crates be powered down

The uniformity requirement is important - the microcontroller code should not make assumptions about module type or position. This is necessary because it is feasible that modules may be moved or swapped between runs. The DCS for replaced modules should work “straight out of the box”, without having to rely on expert users to load specific code versions. In addition, crate layouts differ between the production system installed at CERN and the various testbenches. Uniformity means that the same DCS software can be used on production and test bench systems.

The modules are passive data sources; they cannot interact directly with the parent VME crate CANbus (to request the crate power off, for example). Because of this, an important requirement of the front-end is that it be able to notify the back-end of a problem very quickly, so that the back-end may take appropriate action. The exact details of this feedback mechanism are discussed in the next chapter.

5.4 CAN Microcontroller Code

A custom software solution was developed to run on the CAN microcontrollers. This software is responsible for inter-module communication on the CANbus and also communication with the DCS back-end. In order to interface successfully with centrally provided back-end software, the microcontroller code was based loosely on the CANopen message protocol. The CANopen protocol defines a number of

message types and dictates how they are identified. A subset of these message formats has been implemented for the L1Calo project.

It is the use of CANopen which dictates why the TCM must act as a bridge when communicating with crate modules (figure 5.2). The CANopen protocol only allows 127 nodes to exist on the same CANbus³. As there is the potential for up to 320 (ie 16×20) monitored modules in total, they must be grouped into smaller CANbuses. A logical way of doing this is to define one sub-bus per crate (ie the internal CAN bus in figure 5.2). Messages from the DCS back-end are addressed to individual TCMs, but contain a module ID in the first byte of data. This ID is interpreted by the receiving TCM, so that it can forward the message on to the appropriate module.

5.4.1 Node ID

On an internal crate CANbus, each module is assigned a node ID based on its position within the crate. Numbering starts from 1 and extends to 21, labelling modules from left to right when facing the front of the crate. As the TCM is always located in the slot on the far right, a TCM always communicates on the internal CANbus with a node ID of 21 (figure 5.2). The crate CPU is always found in slot 1, but this is not monitored by the DCS⁴.

TCMs also require another Node ID to communicate on the external CANbus. In this case, the node ID is based on the crate number, which is set by a variable resistor on the backplane. Each backplane provides a different number, ensuring

³7 bits are used to define the node ID, which must range from 1 - 127 inclusive. Messages containing the node ID 0 are considered to be addressed to all modules.

⁴Though provision has been left in the front-end software should this change in the future.

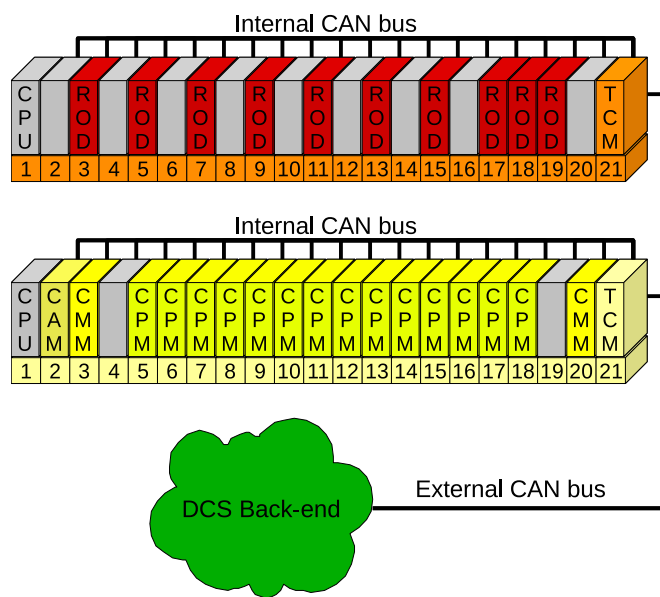


Figure 5.2: Two of the sixteen crates which make up the L1Calo module CANbuses. Each VME crate defines an internal CANbus on the backplane, which is used by modules to communicate with the TCM. Each TCM can also communicate with the DCS back-end via the external CANbus.

that TCMs may be uniquely identified on the external CANbus.

5.4.2 Data Transfer

A number of message formats have been defined for communicating with the L1Calo modules. The most important of these are:

- **Error messages** for reporting errors on individual channels.
- **Module type messages** for reporting the module type present in a crate slot and CAN code version it is running
- **Module data messages** for reporting the online values of the channels monitored by a module
- **Configuration messages** for configuring individual modules

Message types are identified using 4 bits close to the beginning of the CAN frame. Data from each monitored channel is digitised with a resolution of 8 bits. Each CAN message contains a data payload of 8 bytes, which is not enough to broadcast data from all channels for the majority of modules. For this reason, most messages are multiplexed - the first byte is used as an extra identifier (see figure 5.3).

5.4.3 Network Management

In addition to data-carrying messages, there are also a number of Network Management (NMT) messages defined. Messages of this type are sent to move a node into a new state or to notify the master node of a change in a slave node state. The




Figure 5.3: Part of a typical CANopen data frame. The first 11 bits are used to identify the message - 4 bits define the message type and 7 the node ID. In this example, the message type is a Module Data message being broadcast by the module in slot 9. The data packets consists of 8 bytes. As module data messages are multiplexed, the first byte is used to define the multiplex ID. In this example the Mux ID is equal to 1, which means this is the second data packet. The remaining 7 bytes are used to broadcast data from monitored channels.

most important NMT states are Operational, Preoperational and Stopped, which are summarised in table 5.4.3. Modules move between states on receiving NMT commands to do so, usually from the master node.

The NMT commands also define a Reset command, which forces a module's CAN microcontroller to reinitialise.

5.4.4 Module Initialisation

When the module is powered on, the CAN code initialises. The microcontroller starts by determining the module type and its position within the crate. This information is recorded and used to define CAN message IDs. Execution then splits depending on whether the module is a TCM or not. I/O ports are initialised in the case of generic (ie non-TCM) modules, so that monitoring data may be read. CAN ports and buffers are then initialised.

NMT State	Integer Code	Allowed Actions
Operational	5	Receive and broadcast NMT, error, data and config messages
Preoperational	127	Receive and broadcast NMT and config messages
Stopped	1	Receive NMT messages only

Table 5.2: NMT States supported by L1Calo modules. Each module broadcasts a 7 bit integer, representing its NMT state. This integer is broadcast by each module on a regular basis (typically every 16 seconds). This heartbeat message can then be used by the back-end to determine if the module is still present.

Default operating limits are configured depending on module type. Note that these default intervals are hardcoded into the CAN microcontroller code, allowing for much faster initialisation times. They may be changed, on a module-by-module basis, by the back-end once the crate has finished initialising.

5.4.5 Generic Event Cycle

Once the default limits have been loaded, the module has finished initialising and it is considered to be in the Operational NMT state. It then enters its normal event cycle.

Generic modules in the Operational state read data from various sensors located on the module and record them in the Object Dictionary. This consists of a large block of ordered, reserved memory. I/O operations on this block are atomic - they cannot be interrupted by any other process - which ensures that data cannot be corrupted.

Once all sensors have been read, the data are compared to the on-board operating intervals. If a channel lies beyond the acceptable limits, an error message may be broadcast to the TCM, depending on priority and severity. After this, the event cycle then restarts.


The length of the event cycle depends to first order on the number of channels a module monitors. The CAM has the shortest event cycle at approximately 1.5 ms (figure 5.4). The JEM, which monitors the highest number of channels, takes 18.5 ms to complete a cycle.

5.4.6 TCM Event Cycle


The TCM does not monitor any data directly, so its event cycle is entirely driven by asynchronous events. The back-end is responsible for broadcasting a NMT Sync message to all TCMs every 17 seconds. On receiving this message, each TCM waits for a number of seconds equal to its node ID (see figure 5.5). It then broadcasts the cached data about the modules detected within the crate, followed by their monitored data.

One second after broadcasting module data, the TCM clears its cache and rescans the crate. It sends CAN messages to each slot in turn, asking any module present to respond with a module type code. The TCM then requests each module present in the crate broadcast the data values currently stored in their Object Dictionary. This is written to the Object Dictionary on-board the TCM, ready to be broadcast on receiving the next NMT Sync message⁵. This sequence is shown in figure 5.6.

⁵Because the TCM broadcasts data first and then probes the crate, the first data messages from the TCM after initialisation are explicitly zero. Meaningful data will instead be broadcast on receiving the second Sync message



(a) CAM Event Cycle



(b) JEM Event Cycle

Figure 5.4: Generic module event cycle time. For the purposes of this test, the beginning and end of the event cycle were marked by the broadcast of a heartbeat message (shown in green). The length of the event cycle time varies between 1.5 and 18.5 ms, depending on module type. Note that the time scale is 500 $\mu\text{s}/\text{unit}$ in the CAM event cycle trace and 5 ms/unit for the JEM trace.




Figure 5.5: TCM Sync cycle. The DCS back-end broadcasts a sync message on the L1Calo Module External CANbus, labelled A, every 17 seconds (the time scale is 5 s/unit). Crate TCMs then wait for a number of seconds equal to their external node ID before replying with cached data. In this example, only crates 11 (labelled B) and 15 (labelled C) are connected to the CANbus.

The replies from the four monitored modules present in the crate can be seen.

Communication between the TCM and the back-end, and between the TCM and the crate modules takes up approximately 1 second of the 17 second event cycle. For the rest of the time, the TCM waits to receive asynchronous messages - either commands from the back-end or error messages from the modules.

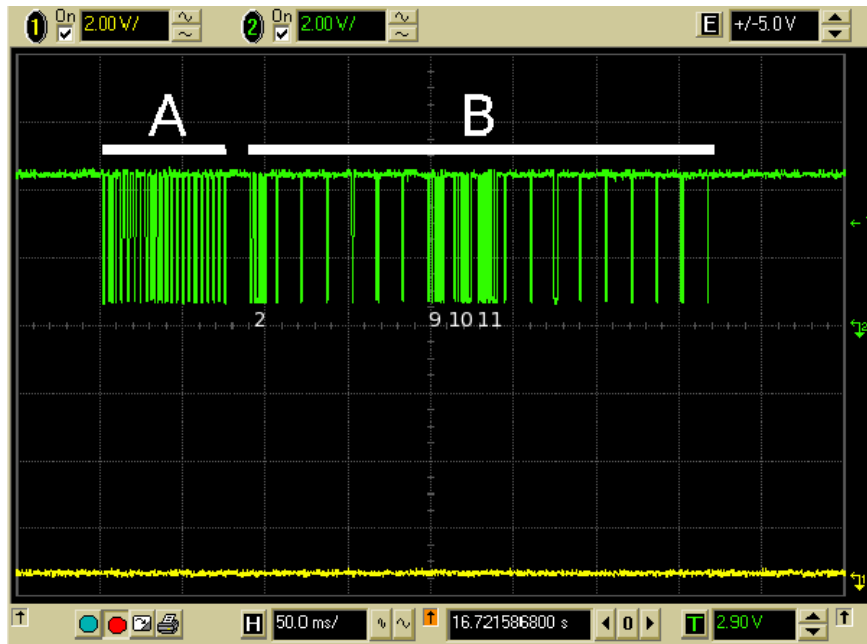


Figure 5.6: TCM event cycle. The TCM probes each slot in the crate for information regarding modules present (labelled A). The TCM then probes each slot again, requesting each modules cache of monitoring data (labelled B). In this example, there is a CAM in slot 2 and 3 CPMs in slots 9, 10 and 11. The thin spikes in section B represent requests from the TCM. Wider spikes represent replies from modules. The CAM and the CPMs can clearly be seen to reply.

5.4.7 Error Messages

Monitored channels are expected to operate within well defined limits. If one or more channels move beyond these limits, the module is determined to be in a state of error and will broadcast an error message on the internal CANbus. This message has a higher priority than any other CAN message, and contains information about the location (ie the module and crate ID) and the severity of the error⁶.

Module error states fall into two categories, Warning and Fatal (see figure 5.7). Warning states represent the situation where a channel has moved beyond acceptable operating limits (typically $\pm 10\%$ of the normal expected value). In this case, the role of the Warning message is to notify of a potential problem - no further action is required. Fatal messages represent a more serious problem (deviations of typically $\pm 15\%$ or more) and are considered to be a request for external action.

Error messages are broadcast only on a change of module error state, with Fatal message taking precedence over Warnings. For example, if a module has already broadcast a Warning message and another channel moves into the Warning operating interval, a further warning message will not be sent. If a channel moves into the Fatal operating regions though, a Fatal message will be sent.

There is also an accompanying Clear message, which is used to represent transitions from an error state to the normal working state. This message is also broadcast when a module (re)initialises, explicitly clearing any cached errors in the DCS back-end.

One important note is that the state transitions have been designed with built in hysteresis to avoid rapid bursts of error messages. All of the monitored channels read in by the CAN microprocessor are digitised. As such, it is possible for a

⁶The channel number and value are also read out for debugging purposes




Figure 5.7: L1Calo module error levels for a chip temperature. Error messages are broadcast only on a state change. Note hysteresis in moving between states. This avoids saturating the bus with error messages if a particular channel lies on a state boundary.

particular channel to oscillate between two values over a number of event cycles. If this happens at the value of an operating limit it's possible that an error message will be broadcast every event cycle, which is undesirable.

In order to avoid this, channels must pass lower thresholds when returning to a less severe error state. Consider the example shown in figure 5.7. In this case, the CPM will enter the Fatal state if the chip temperature is equal to or greater than 70 °C. However, it will only move to the Warning state when the same chip temperature is equal to or less than 68 °C.

5.4.8 Other Asynchronous Events

Once a module has finished initialising, it is possible to check the operating limits and change them on an module-by-module basis. Relaxing these intervals is useful, for example, if a particular slot is known to suffer from reduced cooling performance as is the case of the L1Calo ROCs (figure 5.8).

When reading an operating limit from a module, the back-end starts by sending a request to the crate TCM for a single limit on a specific channel on a particular module. The TCM forwards this request to the appropriate module, which then responds immediately. The TCM relays the module response to the DCS back-end. This process can take up to 6 ms per limit. A similar pattern is followed when changing a limit to a new value, with the new operating interval being specified in the request from the back-end (figure 5.9).

Finally, the TCM also supports a reinitialise command. On receiving this command, the TCM will instruct all modules within a crate to revert to their default operating limits. This can be useful for synchronising the hardware with the back-end, as

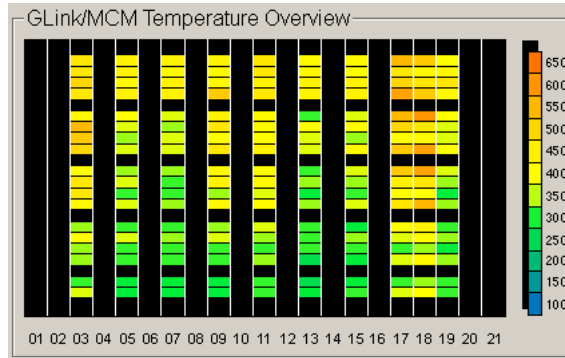


Figure 5.8: L1Calo ROC crate. Each column represents a ROD module within a crate and each block a chip on the module. The colour of each block represents the chip temperature. Due to insufficient air flow, the chips on modules 17 and 18 record higher temperatures on average. It may be desirable to relax the operating limits for the chips on these particular modules.

discussed in chapter 6. Again, cached data values are compared against the default limits immediately after reinitialising, which may trigger fresh error messages.

5.5 Conclusion

The L1 Trigger DCS front-end is responsible for monitoring over 7000 data channels, the majority belonging to the L1Calo Trigger project. These channels are probed every 17 seconds, although higher priority error messages give an effective response time of 6 ms. The next chapter will discuss how this data is used by the DCS back-end to provide an overview of the system, and how the back-end is used to control the hardware.

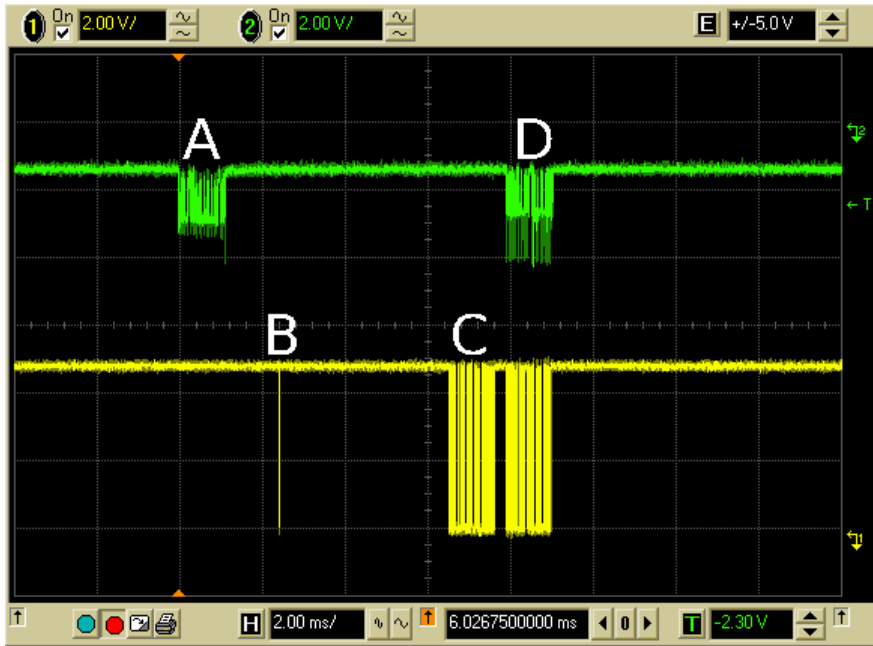


Figure 5.9: Updating and checking operating limits. At point A, the TCM (green) transmits a new operating limit to a CPM (yellow). This limit is received by the CPM at point B, 1.3 ms later. In this case, the change in operating limit triggers an error message. Error message is emitted from the CPM at point C, 4.1 ms after the limit is changed. The error is received by the TCM at point D, approximately 5.5 ms after the initial update is received by the CPM.

Chapter 6

Level-1 Trigger DCS Back-End

6.1 Project Requirements and Scope

The DCS back-end is responsible for processing slow controls data originating from the hardware. In terms of the Level-1 Trigger DCS, the back-end is responsible for monitoring Wiener crates and a large number of modules in the L1Calo system. The scope of this project is limited to the level of the Local Control Station (LCS). The main requirements of the project are to:

- Define the connection between hardware and PVSS, enabling data readback
- Archive relevant monitoring data for further analysis
- Provide a mechanism for remote control of the hardware
- Flag visually any faults in the system so that they be identified easily

It was desirable to develop the project in a way that it may be deployed on systems isolated from the main DCS network, specifically test benches at various institutes.

This adds the extra requirement that software be configurable and generic. Tools have been developed to allow users to tailor the software to a local system, allowing them to add or remove crates for example. It was especially important that this generality requirement be met when modelling the L1Calo modules, as they are removed and replaced on a test bench regularly and may even be found in locations which differ from the production system.

6.2 Hardware Connection

PVSS is unable to communicate directly with the hardware on the CANbus. Instead, the gap between PVSS and the CANbus is bridged by the use of OPC servers. OPC¹ is an industry standard which allows control data from different devices to be shared.

The standard requires that supported devices be accompanied by an OPC server. This software is responsible for the direct communication with the hardware. It provides a standardised interface for I/O operations, allowing other software programs (OPC Clients) to connect, read back information from the server and write commands.

In the case of the L1 Trigger DCS, two OPC servers are used (figure 6.1). The first is provided by Wiener and is used to communicate with the VME crates. The second server, provided by the ATLAS Central DCS group, enables communication with devices supporting the CANOpen protocol. This OPC server is used to communicate with the L1Calo modules. PVSS has a built in OPC client, and so is capable of connecting to both of these OPC servers directly to read back hardware data.

¹Formerly OLE for Process Control, but marketed as an acronym without meaning since 2006.

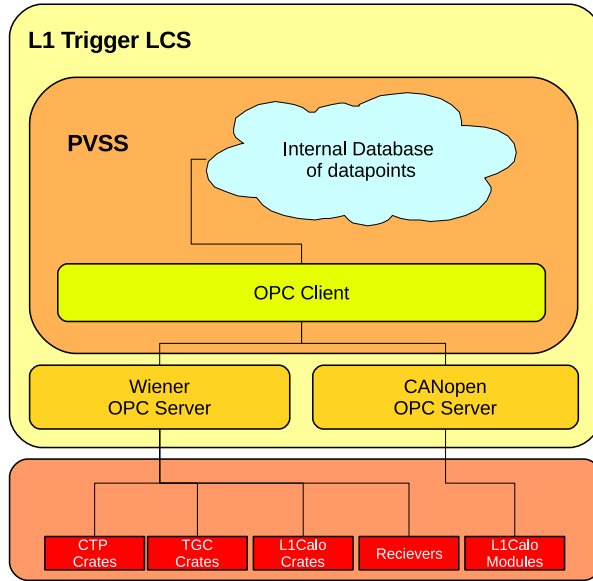


Figure 6.1: OPC Servers. Two OPC servers interpret data coming from the CANbus and make it available to PVSS.

OPC clients may subscribe to all data made available by the OPC server, or just a subset. When a client connects to an OPC server, it creates a number of internal groups into which subscribed items are assigned. Different groups can be updated with different refresh rates. This is important because it means that fast-changing data (such as a voltage) can be monitored separately from slowly changing data (module barcode), reducing CPU load. In the L1 Trigger project, Wiener information is refreshed at either 30 second or 15 minute intervals, depending on importance.

L1Calo module information is refreshed at a higher rate (100 ms), as the modules are more reliant on the back-end for intervention when a problem occurs. Commands sent to both the Wiener crates and the L1Calo modules are also refreshed at 100 ms, which is the fastest time the OPC server can execute them. It should be noted that it is this delay in reading back OPC items which dictates how fast the DCS can react to error messages coming from the L1Calo hardware.

The majority of OPC items will only be updated if there is a difference between the currently held value and the value last polled. The exception to this is the NMT state coming from the L1Calo TCMs. As only the lowest seven bits of an NMT message are used to specify a particular state, the OPC server toggles the eighth bit after every heartbeat (so for example, the OPERATIONAL state can be represented by 5 or 133). Toggling this bit means that the NMT OPC item will change everytime a new message is received, enabling PVSS to use the item to confirm a module's continued presence on the CANbus.

In keeping with the generality requirement for the L1Calo system, the CANopen OPC is configured to provide address space for all 21 slots in a crate. This means that the OPC server does not have to be reconfigured every time a module is added or removed.

6.3 PVSS

6.3.1 Wiener Crates

The Wiener crates are represented within PVSS by a number of datapoint types made available by the JCOP team. The datapoint types are generic in that they support the monitoring of up to eight voltage channels per crate. The majority of Level-1 Trigger crates only require 2-5 input channels (5 in the case of the CTP crate in figure 4.1), so surplus datapoints are deactivated.

6.3.2 L1Calo Modules

Each L1Calo module is represented within PVSS by a number of custom datapoints. The design of the datapoints is generic so that they can be used to represent all types of module, much like a template. One module datapoint exists for every slot in every Wiener crate in the L1Calo system². Each module can potentially monitor up to 42 channels, but as with the Wiener crates, surplus channels are deactivated.

The module datapoints are currently configured by a number of XML files which are parsed by scripts in PVSS. This system will eventually be replaced by the use of the Configuration Database. Using these tools, modules may be quickly, and individually, configured to represent the available hardware, whether it be the production system at Point-1 or a local test bench. Configuration consists of filling in expected values to a module datapoint. For example, a module datapoint corresponding to a particular slot in a crate may be configured to expect a CPM. The datapoint will be told what the default operating limits for that module are, how many channels it should be monitoring and what version of the CAN code it should be running.

Depending upon its expected type, each module datapoint will be configured to transform data received from the CANbus into a human readable form before it is displayed or archived. The modules broadcast all channel data as an 8 bit integer, which may require scaling and the use of an offset before it is meaningful. This transformation system is flexible in that, although each module starts with an appropriate default transformation, it can be varied from module to module. This allows scope for modules to be individually calibrated in the future.

²Excluding the TTC crate, which does not monitor any modules, and slots containing TCM's as these are handled separately. This corresponds to 20×16 module datapoints in total.

It is also possible to adjust the operating limits on a module-by-module basis. Rather than continuously transfer and confirm every operating limit in the system, PVSS keeps a local cache of the limits stored on each module. However, this requires that the limits stored in PVSS and those stored on a module are carefully synchronised when either list is updated. This is especially true when a crate is power cycled, as modules will reload default limits. This synchronisation process is demonstrated in figure 6.2.

An alarm is set on the error flag of active modules, ensuring that if a module does report a fault it will be logged and displayed in the DCS Alarm Screen. In addition though, WARNING alarms are set on the module type and CAN code version data received from the CANbus. Different CAN code versions may contain different default operating limits, so it is important that these items match the expected items, otherwise the wrong set of operating limits may be listed in PVSS. Currently, inconsistencies are flagged for the attention of an expert user, who can reconfigure the module datapoint if necessary.

6.3.3 TCM

The TCM represents a break with the requirement that L1Calo module representations be generic and configurable within PVSS. This is necessary because when sending information to the CANbus via the CANopen OPC server, there can only be a one-to-one mapping between PVSS datapoints and OPC items.

The TCM datapoint type inherits³ from the L1Calo module datapoint type. It defines the same alarms and hardware addresses as the generic module. The TCM

³PVSS does not directly support the programming concept of inheritance, so in this context inheritance refers to a straight copy of the L1Calo module type, with extra elements added on.

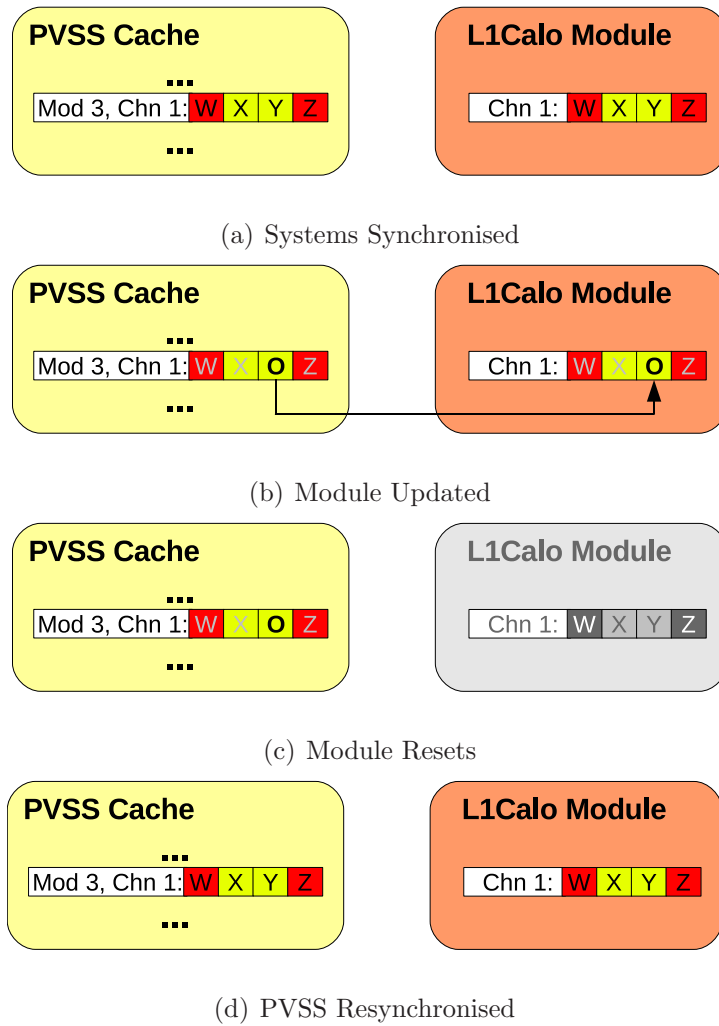


Figure 6.2: Synchronisation of Operating Limits between PVSS and L1Calo Modules. PVSS and an L1Calo module start with synchronised copies of the same operating limits, (a). Limits may be changed in PVSS and then sent to the hardware, (b). If a module is powered off, (c), it will restart with its default limits. This means that PVSS must reset any changes, (d), when the module comes back online.

does not monitor any data channels directly, but it does broadcast a module type code and its CANcode version numbers, so in this way the TCM can be represented in PVSS as a generic L1Calo module with zero monitored channels. However, extra datapoint elements are included for writing commands to the CANbus. These commands include:

- Changing the NMT state of modules within the crate
- Sending new operating limits to specific modules within the crate
- Resetting the CAN microcontrollers on board specific modules.

Datapoint elements are also provided for reading back operating limits on specified modules, for the purposes of confirmation and debugging.

This failure to meet the generality requirement is justified due to the special role of the TCM. Firstly, due to the crate mechanics, it is impossible to find the TCM in any slot other than 21 - it has been designed not to fit into any other slot. Secondly, as the back-end cannot communicate with crate modules directly, the TCM must be present in order for the DCS to function. These facts allow the additional specialisations outlined above to be made.

6.4 FSM

The FSM is one of the most important parts of the Level-1 DCS back-end, as it represents the main interface for controlling and monitoring the hardware. It is required to:

- Display the status of the Wiener crates and monitored L1Calo modules

- Allow users to securely issue commands to Wiener crates and TCMs
- Power off a Wiener crate if a problem occurs with an L1Calo module
- Synchronise PVSS and the L1Calo modules, to ensure that both systems have the same operating limits

The LCS FSM has a single root node which is designed to give summary information of the status of the Level-1 trigger. This node is referenced in the parent SCS FSM (as shown in figure 4.3). Below this, the FSM is divided into four partitions, representing the four monitored subsystems. Three of these systems share a common structure, as they only monitor Wiener crates. The fourth partition, representing the L1Calo Trigger, contains extra complexity due to the monitored modules.

6.4.1 Crate-only Partitions

The structure of partitions which only monitor Wiener crates is shown in figure 6.3. The top node is a Control Unit (CU) which gives an overview of the entire partition. Below this, there is one Device Unit (DU) for every monitored Wiener crate.

It is important to note that the entire Wiener crate FSM is completely passive. Commands are propagated down and state information up, but no automatic action is taken by this part of the FSM. This is because the Wiener hardware will decide what action to take in the case of a fault. This independence from the FSM makes the system robust as a Wiener crate is capable of tripping off even if communication with the DCS back-end is lost.

The Wiener crate DU is relatively simple, supporting 3 (+2) states⁴. The DU

⁴In addition to any normal operation states, all FSM nodes support the UNKNOWN and

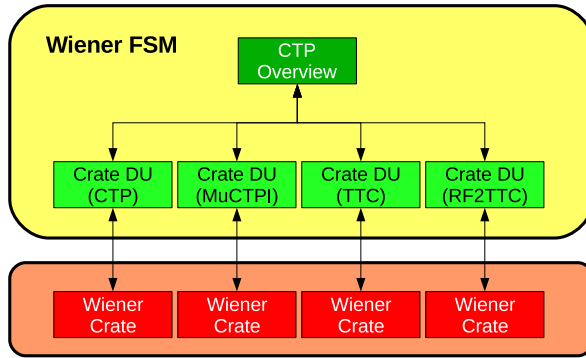


Figure 6.3: Wiener crate FSM. One FSM partition like this exists for each Wiener CANbus connected to the LCS (excluding L1Calo). The root node is a Control Unit giving users an overview of the entire partition, the CTP in this case. Below this, there is one Device Unit representing each Wiener crate connected to the CANbus (4 crates in the case of the CTP).

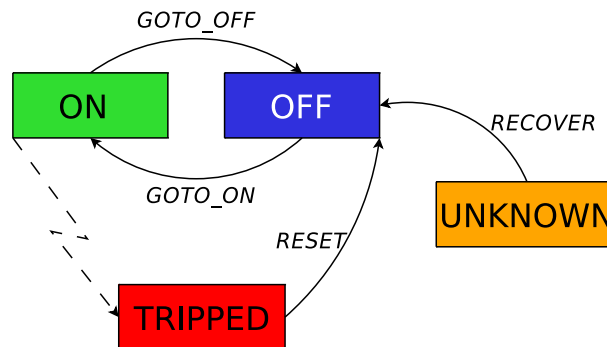
monitors two datapoints in order to determine the state of any given Wiener crate – the power status of the crate (whether it is switched on or not) and the error flag. In the case of a fault, such as a fan failure, the Wiener hardware will set the error flag and decide for itself whether or not to trip off, based on the conditions set via the crate front panel. Figure 6.4(a) shows how the FSM state is derived, and Figure 6.4(b) shows the valid state transitions.

The Wiener DU status is derived from the crate summary alarm, which checks the error flag along with temperature, fan status flags and flags representing the active voltage channels. If the summary alarm has been set by the crate, but the crate is still powered, the crate is determined to have an ERROR status. If the alarm has been set and the crate has tripped off, the crate is determined to have a FATAL status.

DEAD states. These two additional states are used to reflect problems with the FSM rather than the hardware.

	Power	
	On	Off
No Error	ON	OFF
Error	ON	TRIPPED

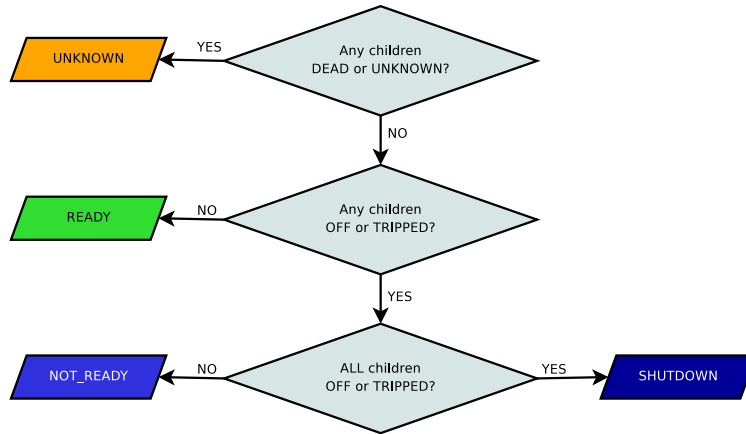
(a) Derivation of Wiener crate state



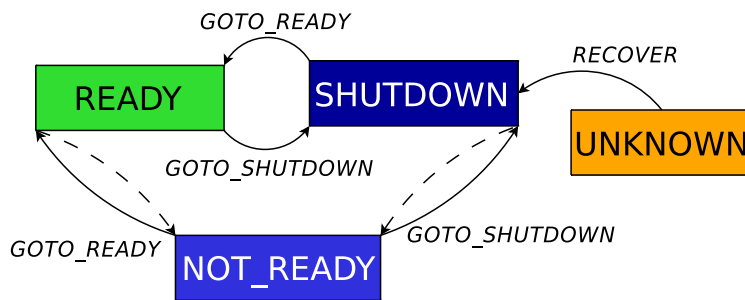
(b) Wiener crate state machine

Figure 6.4: Wiener VME crate FSM. Figure 6.4(a) shows how the crate state is derived. Figure 6.4(b) shows the allowed transitions between states and the available commands which may be used to trigger a transition. Note that dashed lines represent allowed transitions for which no user command is available.

The partition overview CU derives its state from its children (ie the Wiener crate DUs in this case), as shown in figure 6.5(a). If any crates are powered off, the partition will move to the NOT_READY state. When all crates are powered off, the partition moves to the SHUTDOWN state. The status of the partition CU is derived from the most serious child status. For example, if one Wiener crate reports a FATAL status, and the rest are OK, the partition overview will also report a FATAL status.



(a) Derivation of the LV1 Crate Overview state



(b) LV1 Crate Overview state machine

Figure 6.5: LV1_CRATES Overview FSM node. Nodes of this type provide a summary of child Wiener crate states. They are used to represent the TGC, CTP and Receiver FSM trees.

6.4.2 L1Calo Partition

The structure of the L1Calo FSM is shown in figure 6.6. DUs are grouped to represent the 17 crates that make up the L1Calo trigger. Each group consists of 22 nodes – 20 representing the 20 potential modules per crate, 1 node for the TCM and a final node for the Wiener crate itself. Each group has a parent CU, which summarises the state of the crate. Finally the 17 crate CUs are summarised in a root node, which represents the L1Calo partition.

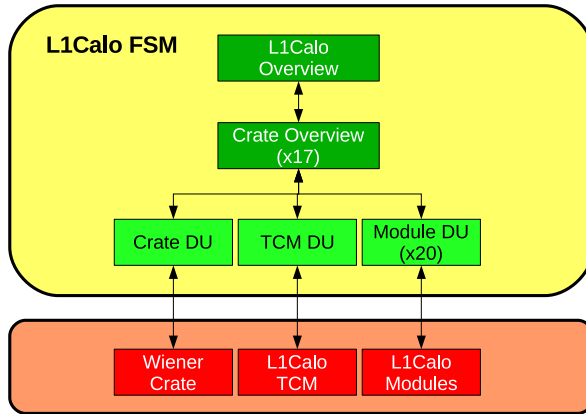


Figure 6.6: L1Calo FSM. The root node references 17 children, each of which gives an overview of a specific L1Calo crate. Each crate overview gives summary information from 20 module DUs, 1 Wiener and 1 TCM DU.

6.4.3 Module Device Units

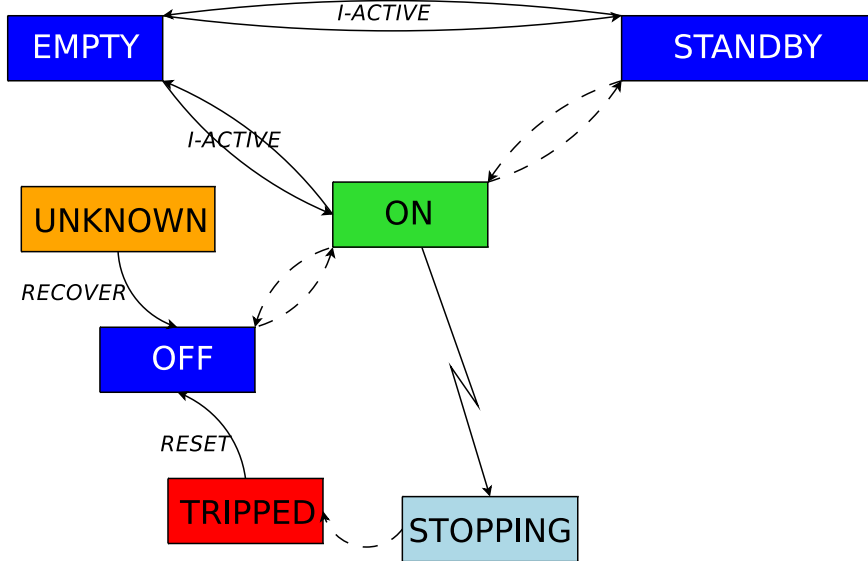
DUs representing modules are generic, and make no assumptions about specialisations for specific module types. The DU state is derived from a number of indicators (figure 6.7(a)). Assuming that the module is powered, and has not reported a fatal fault, the DU state depends on whether or not a module has been detected in that slot and the value of the NMT state⁵. If a module is detected, and is broadcasting monitoring data, the DU will move to the ON state. If no module is detected, the DU will move to the EMPTY state.

When deriving the module DU status, the FSM first checks the module fault flag. If a fault has been reported by a module, the DU status will match the severity of the fault (either WARNING or FATAL in the case of L1Calo modules).

⁵Strictly, NMT state here refers to that of the crate TCM. Because all communication goes through the TCM, only the NMT state of the TCM is important. If a specific module is in an NMT state other than operational, the TCM will not detect that module and so will report an empty slot to PVSS.

		Crate Power				
		No	Yes			
Error	2	TRIPPED	STOPPING	{0, 1}	NMT	
	!2	OFF	OFF			5
			EMPTY			ON
			STANDBY	{6, 127}		
		0	!0			
		Module Type				

(a) Derivation of the L1Calo Module state



(b) L1Calo Module state machine

Figure 6.7: L1Calo Module FSM. Figure (a) shows how the module state is derived. Figure (b) shows the allowed transitions between states.

If a module does report a fatal fault, the FSM will take action to automatically power off the Wiener crate. The process starts with the module DU moving into the STOPPING state. This state is detected by the parent CU, which then issues the GOTO_OFF command to all child nodes, including the Wiener crate DU. This causes the Wiener crate to power off. Once the crate is off, the problem module is

moved to the TRIPPED state, allowing operators to clearly distinguish the source of the trip.

Assuming no fault is detected, the FSM will then check that the module detected in a specific slot is the module that was expected. If the reported module type does not match the expected module type, or if the CANopen code version does not match, the module DU status will be set to WARNING. This gives a clear indication to a user that there may be a configuration problem. It is important that these problems be checked and fixed, because other information originating from that slot (ie monitoring data) may be archived and displayed incorrectly.

Because PVSS does not communicate with L1Calo modules directly (routing messages instead through the TCM), very few FSM commands are available for use with the module DU. Other than commands associated with error recovery, the only other command available is I-ACTIVE. This command allows expert users to deactivate the modules OPC addresses in PVSS, effectively cutting the link between the module and PVSS. This can be useful when trying to debug and isolate problems.

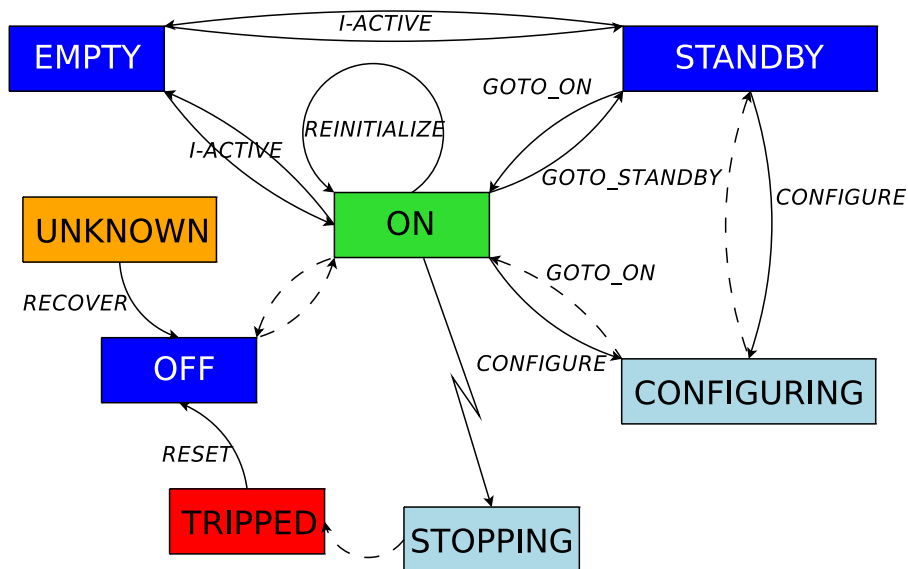
6.4.4 TCM Device Units

As with the datapoint representation, the TCM FSM DU inherits from the generic module DU, but adds functionality to support its role in configuring the other modules. The FSM state derivation is shown in figure 6.8(a).

One of the most important roles of the TCM DU is to synchronise the list of module operating limits PVSS holds with the limits cached by the hardware. If new limits have been loaded into PVSS (either from the Configuration Database or via some other mechanism), they may be sent to the hardware by issuing the CONFIGURE

		Crate Power					
		No	Yes				
Error	2	TRIPPED	STOPPING				
	!2	OFF	OFF	ON	CONFIGURING	{0, 1}	NMT
			EMPTY			5	
		STANDBY				{6, 127}	
		0	!0				
		Module Type					
		0					
				> 0			
				Configure Flag			

(a) Derivation of the TCM state



(b) TCM state machine

Figure 6.8: L1Calo TCM FSM. Figure 6.8(a) shows how the TCM state is derived. Figure 6.8(b) shows the allowed transitions between states and the user commands which will trigger them. Although it is based on the L1Calo Module DU, the TCM DU supports extra states and commands for configuring modules.

command from the TCM DU. This command will first move all modules in the crate to the STANDBY state, clearing the CANbus of traffic. The TCM will then enter the CONFIGURING state and proceed to broadcast the updated operating limits

to the hardware, returning to its original state once the transfer is complete.

If a crate is powered off the modules will revert to their default operating limits when powered back on. It is important that this reset is applied to PVSS too. As such, when an L1Calo Wiener crate is powered back on (either via the DCS or manually), the crate overview CU issues the REINITIALIZE command to the TCM. This triggers a process which resets the operating limits held in the PVSS cache, as the power cycled modules will have reverted to their default values. It is possible to change this behaviour so that in the future any non-standard configuration held in PVSS is sent to the modules instead of being reset.

6.4.5 Higher Level Control Units

The L1Calo crate CU logic may be seen in figure 6.9. CUs of this type play an important role in reacting to faults with modules (section 6.4.3). These CUs also provide an overview of the crate status, and follow a similar logic to the Wiener crate Partition overview CUs. The L1Calo crate CUs then feed into a generic CU, to give an overview of the L1Calo partition. Finally all of the partition overview CUs (L1Calo, Receivers, CTP and TGC) feed into the FSM root node, which provides an overview of the entire LCS FSM. It is this node which is referenced by the SCS FSM.

6.5 FSM Operation

The normal use case for controlling the Level-1 Trigger hardware is through the DCS FSM in the ATLAS control room. During a run the FSM is included in the

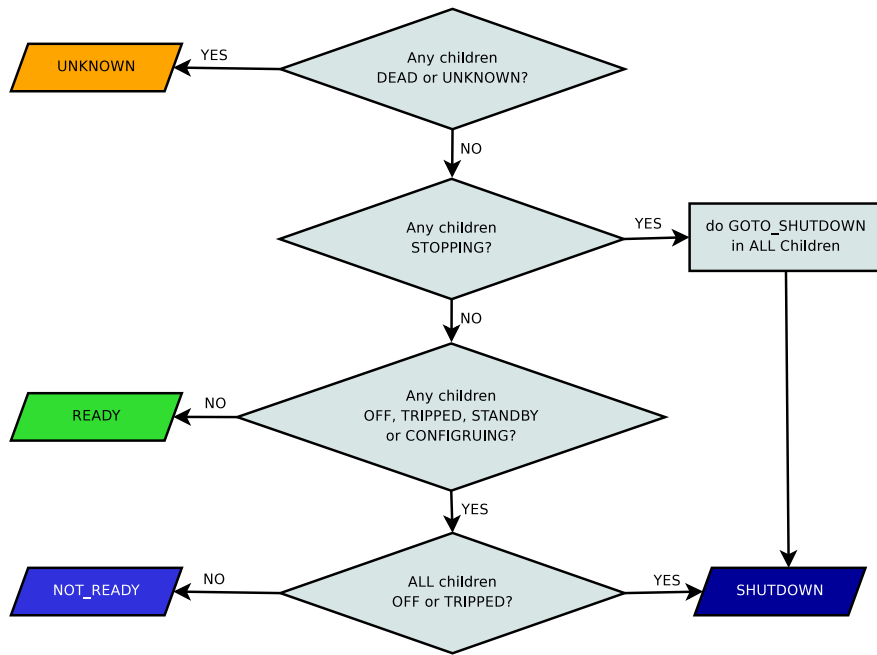


Figure 6.9: LV1_CALO_CRATE Overview CU. They are used to represent an overview of a single L1Calo crate, summarising information from 1 Wiener crate, 1 TCM and 20 module children. The allowed transitions between states are the same as those shown in figure 6.5(b).

main ATLAS partition, which is controlled by the operator on the DCS desk. Other users may view the status of the system, either at the Level-1 Trigger desk, or via the DCS terminal server.

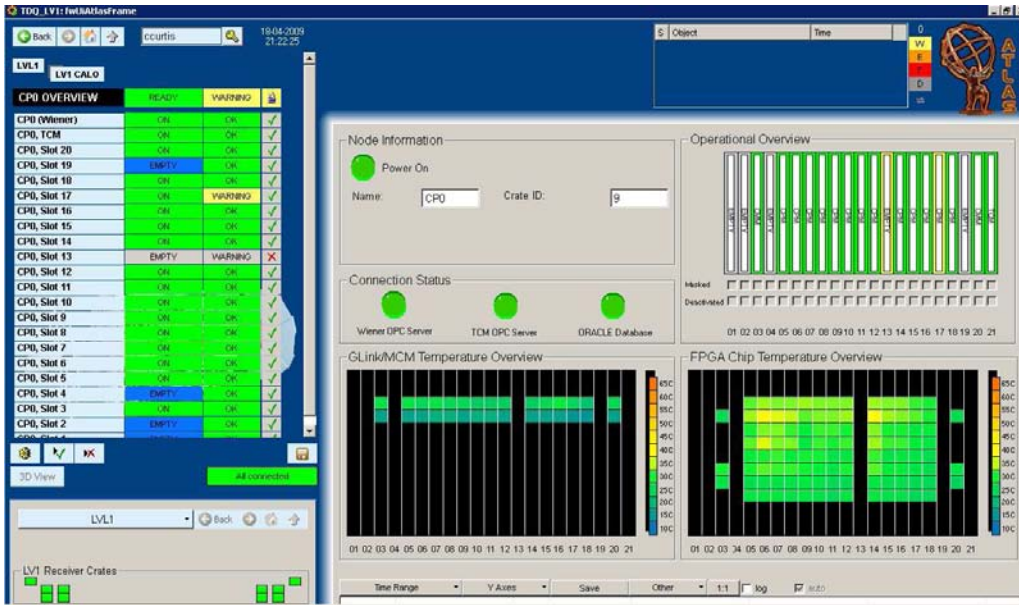
Every node type in the LCS FSM is accompanied by a graphical interface, designed to display information relevant to that particular level of the FSM . Figure 6.10 shows an example of screens relating to an L1Calo CP crate. In 6.10(a), an overview of the crate is shown. A temperature map of the monitored chips is available, and users may see at a glance what type of modules have been detected in each slot. In this example, there are problems with two of the modules. The user may click through using the navigation panel on the left to find more information about each Module DU . In this case, the screen in 6.10(b) shows that Warning status on module 17

is associated with a mismatch between the expected CAN code and the detected versions (as represented by the orange warning light). This screen also shows a table detailing online values for the monitored channels (chip temperatures in this example). Clicking on this table will bring up archived data for each channel.

Casual observers may also view a limited subset of ATLAS DCS information on the Point-1 website. This is a read-only facility which requires no special privileges.

6.6 Conclusion

The Level-1 Trigger DCS has been available for use at Point-1 since 2007 and has been tested during Milestone runs and other commissioning activities. This culminated in the project being included in the ATLAS wide DCS during the September 2008 switch on event, giving shift operators a clear indication of the hardware status during real data taking.



(a)



(b)

Figure 6.10: L1Calo FSM GUI. These screens are used to display summary information about a specific L1Calo crate. For example, a complete temperature map of all of the monitored module processors can be seen in (a), with more detailed information about a module in (b).

Chapter 7

Single Top Production

The analysis in this chapter aims to investigate the possibility of measuring the Standard Model single top t-channel production cross section using early LHC data at $\sqrt{s} = 10$ TeV. The important sources of background are discussed and a series of cuts are then defined which improve the signal to background ratio whilst also reducing the sensitivity to some of the major systematic uncertainties. Finally, the possibility of excluding the background only hypothesis is explored.

7.1 Monte Carlo Model

As the LHC has yet to start colliding protons at the time of writing, this study focuses only on simulated data. A number of Monte Carlo generators have been employed - AcerMC, MC@NLO and Alpgen to model the hard interactions and Pythia and Herwig to model the parton showers. All Monte Carlo data have passed through the full ATLAS detector simulation and reconstruction.

As the search strategy for the single top is based on isolated, high P_T leptons accompanied by at least one b-jet, the background model was restricted to processes fitting these criteria. The following sections describe some of the details of how these processes were simulated.

7.1.1 Single Top

The production of single top quarks was modelled by the AcerMC generator [42]. Production via the t-channel consists of two major Feynman Diagrams (figure 7.1). The leading order (LO) contribution consists of a b type sea quark combining with a radiated W boson. The most significant next to leading order (NLO) diagram consists of a gluon splitting into a virtual b quark which then combines with a W boson. Both diagrams are characterised by the presence of a high energy quark (q') in the final state, in addition to the top quark decay products.




Figure 7.1: Single Top t-channel production.

AcerMC simulates the process to leading order, but includes the leading log terms, and hence the dominant NLO diagram [43]. The cross section has been scaled to NLO by a k-factor of 1.05. The NLO Standard Model cross section¹ is 43.2pb, which

¹This is the cross section for the case where the W decays leptonically. The full SM production

will result in approximately 8600 events for an integrated luminosity of 200pb^{-1} . [44].

7.1.2 $t\bar{t}$

The $t\bar{t}$ process (figure 7.2), has been simulated at NLO using the MC@NLO generator [45]. As this analysis requires a high P_T lepton trigger, the generated sample was filtered for at least 1 leptonically decaying W boson. The sample cross section is 402pb , which corresponds to approximately 80 000 events at 200pb^{-1} . [44]

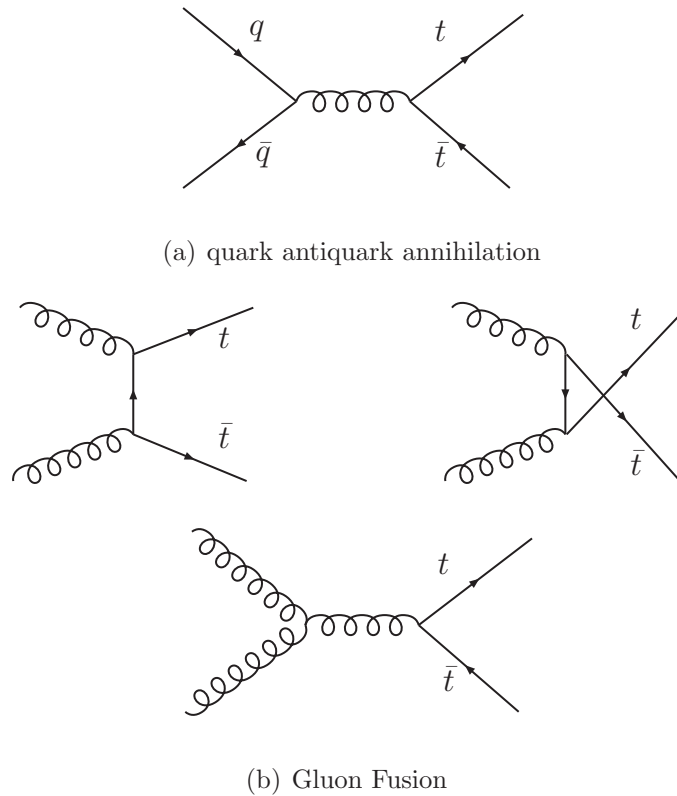


Figure 7.2: LO $t\bar{t}$ production diagrams. At the LHC, 10% of $t\bar{t}$ production will come from quark antiquark annihilation, (a). The remaining 90% will come from gluon fusion, (b). [46]

cross section is 122pb .

7.1.3 W/Z+jets

Along with $t\bar{t}$, W+jets is expected to be one of the major backgrounds in real data (figure 7.3). A large sample of W+(0-5) hard partons was simulated using the Alpgen generator [47]. A filter was applied at the generator level requiring the W to decay leptonically. A k-factor of 1.22 scales the samples to a NLO cross section of 48nb, which corresponds to nearly 10 million events in 200pb^{-1} of data [44].




Figure 7.3: W/Z + jets production

The Z+jets channels are also expected to be an important background, as one of the leptons from the Z decay may not be reconstructed properly. The processes Z+(0-5) partons were simulated and have been scaled to a NLO cross section of 4 nb, resulting in a further 800 000 background events in 200pb^{-1} of data. Again, a generator level filter was applied, requiring the Z to decay to either an electron, muon or tau.[44]

7.1.4 Other backgrounds

A number of other background sources have also been investigated in this study. Although their cross sections are in general small, they are expected to be selected with a high efficiency.

The Wt single top associated production channel (figure 7.4) has also been generated using AcerMC, though it is considered to be a background channel in this study. There is required to be at least 1 leptonically decaying W , resulting in a cross section of 14 pb [44].




Figure 7.4: Wt Associated Production.

The WW and WZ diboson channels (figure 7.5) have also been investigated. Both processes were generated at LO using Herwig, with a generator filter requiring at least one leptonic decay. The lepton was also required to have $P_T > 10$ GeV and $|\eta| < 2.8$. The combined cross section, after filter, is 20 pb. This corresponds to 4 000 events at 200pb^{-1} [44].

The $Wb\bar{b}$ +jets sample (figure 7.6) is expected to be selected with a high efficiency due to the presence of at least one good b-tagged jet. A 0-3 parton sample was generated using Alpgen with a cross section of 18 pb [44].

There is a complication when combining the W +jets and $Wb\bar{b}$ +jets samples in that an overlap exists between the two samples, mainly occurring in the region where the




Figure 7.5: Diboson Production.




Figure 7.6: $Wb\bar{b}+\text{jets}$ production.

$b\bar{b}$ pair are generated with low P_T . In order to reduce this overlap, the $Wb\bar{b}+\text{jets}$ events were filtered at generator level, requiring that the hard b quarks have $P_T > 20$ GeV and that ΔR between the b and \bar{b} be greater than 0.7, where $\Delta R^2 = \Delta\eta^2 + \Delta\phi^2$. This method of overlap removal is not perfect and the amount of double counting is estimated to be roughly 4% of the total $Wb\bar{b}+\text{jets}$ cross section.[44]

Although the QCD multijet signals in general lack the high momentum lepton required to trigger the event, the production cross section (~ 9 mb) is high enough to make the lepton misidentification probability significant. It is expected that the QCD processes will form a substantial background, before cuts are applied, to single top production.

Due to the lack of fully simulated events, the QCD background has not been studied

in this analysis; it will require the use of data-driven methods in order to accurately model the processes. Based on the experience of the Tevatron experiments, it is expected that a cut on the angle between the primary lepton and the missing transverse energy, E_T , parametrised as a function of the lepton P_T , will be a highly efficient method of reducing this background.[16]

A summary of the number of expected events and the number of simulated events can be seen in table 7.1.

Process	Expected Number of Events	Number Simulated
<i>t-channel</i>	8600	17 000
$t\bar{t}$	43 000	1 100 000
$W+jets$	10 000 000	7 600 000
$Z+jets$	800 000	1 900 000
<i>Other</i>	2 000	9 000

Table 7.1: Number of expected and simulated events for the single top t-channel and selected backgrounds channels.

7.2 Object Reconstruction

The main strategy employed by this analysis is to separate signal and background events by observing differences in the distributions of reconstructed leptons and jets. Information from the missing energy calculation is also used. Some of the considerations related to reconstructing these objects are discussed below.

7.2.1 Electrons

Reconstructed electrons start off as electromagnetic (EM) calorimeter clusters. These clusters have a fixed size, either 3×7 EM cells in $\eta \times \phi$ or 5×5 , depending on whether the cluster is in the barrel calorimeter or endcap. The candidate clusters are then matched to reconstructed tracks in the tracking detectors. If a track is suitably aligned with the cluster, and if the cluster energy is comparable to the track momentum, then the cluster and track become a candidate for an electron [48].

Electron candidates are then subjected to a series of cuts designed to provide good identification efficiency whilst rejecting background objects such as misidentified jets. The loosest cuts rely only on calorimeter information, taking into account the shower shape and hadronic leakage. A medium set of cuts improves the background rejection by requiring at least 9 hits in the SCT and at least 1 hit in the pixel layers, in addition to the loose cuts. They also use high granularity data from the strip cells in the first layer of the EM calorimeter, rejecting candidates with two local maxima which may instead be evidence for two photons from neutral pion decay within a jet. Finally the tightest cuts also require a minimum number of hits in the TRT and improved track to cluster matching in both the E/p ratio and $\eta \times \phi$ co-ordinates [6].

Figure 7.7 shows that the majority of electron candidates reconstructed in a single top t-channel process only pass the loosest cuts (black line). However, if the reconstructed electrons are restricted to those which are best matched to the truth electron coming from the top decay (red line), it can be seen that the majority of these electrons pass the tight cuts. In this study only electrons passing the medium or tight criteria are kept for further analysis. In addition, electrons reconstructed in the calorimeter crack region ($1.37 < |\eta| < 1.52$) are excluded from further analysis.




Figure 7.7: The number of reconstructed electrons passing the identification cuts. The black line shows the proportion of electrons passing the cuts before any further selection is applied. The red line shows only the reconstructed electrons which have been matched to the MC truth electron coming from the single top t -channel decay.

7.2.2 Muons

There are several muon reconstruction algorithms used in ATLAS, most of which fall into two families - STACO and MOORE [6]. Muons reconstructed by STACO/MOORE algorithms start by reconstructing tracks in the muon spectrometer which are then extrapolated to the beam line. Reconstructed muons may then be combined with tracks in the inner detector, to improve the momentum measurement and to offer a veto on muons created in the calorimeters from charged pion decay. There also exist muon tagging algorithms which search for matching hits in the muon spectrometer for a given inner detector track.

In this analysis the default STACO algorithm was used. Figure 7.8 shows the difference in P_T between the matched truth and reconstructed muon for the two algorithm types.




Figure 7.8: Difference in truth and reconstructed muon P_T for truth matched STACO and MOORE family muons.

7.2.3 Jets

The jet reconstruction algorithms are seeded by combinations of calorimeter cells. Two types of seed are available - either towers or topological clusters. Calorimeter towers are formed by summing the energy of cells in projections of η and ϕ from the point $z = 0$. Alternatively, topological clusters are formed by summing neighbouring calorimeter cells, provided the cell energy is higher than the expected noise value. This results in a three-dimensional, irregular cluster of calorimeter cells.

The calorimeter seeds are then fed into several jet finding algorithms. In this study, a cone algorithm has been used to reconstruct jets. This builds jets iteratively; starting with the seed, the four momenta of all objects (either towers or clusters) within a radius ΔR are combined to form a candidate jet. Another cone of similar ΔR is then centred on the candidate jet and the objects within the cone are recombined. This process continues until the cone position remains stable. The algorithm then moves on to reconstruct the next jet [6].

In this analysis the default tower algorithm is used, with a cone size of 0.4.

7.2.4 Overlaps

Because both the jet and electron reconstruction algorithms take EM Calorimeter clusters as inputs, there is a chance that the same clusters will be reconstructed as both an electron and a jet. These overlaps must be removed before further selection criteria are imposed. Any jet which lies within $\Delta R < 0.3$ of a reconstructed electron is subsequently ignored by the analysis.

Figure 7.9 shows the ΔR between all reconstructed jet candidates in a single top event and the closest reconstructed electron, plotted against the ratio of the transverse energies. A large number of jet candidates lie very close to the electron ($\Delta R < 0.1$). Given that the electron has already passed stringent identification cuts and that the ratio of transverse energies are similar, it is reasonable to assume that these jet candidates are electrons which have been reconstructed by the jet algorithm.

7.2.5 b-tagging

The ability to identify jets which include the decay of a B hadron is an important tool when studying top quark events; b-jets are much less likely to be produced in a large proportion of the background processes. A number of different b-tagging algorithms have been developed for use in ATLAS reconstruction. The tagging algorithms typically assign a weight to each reconstructed jet, which is a measure of how likely it is that the jet is a b-jet.

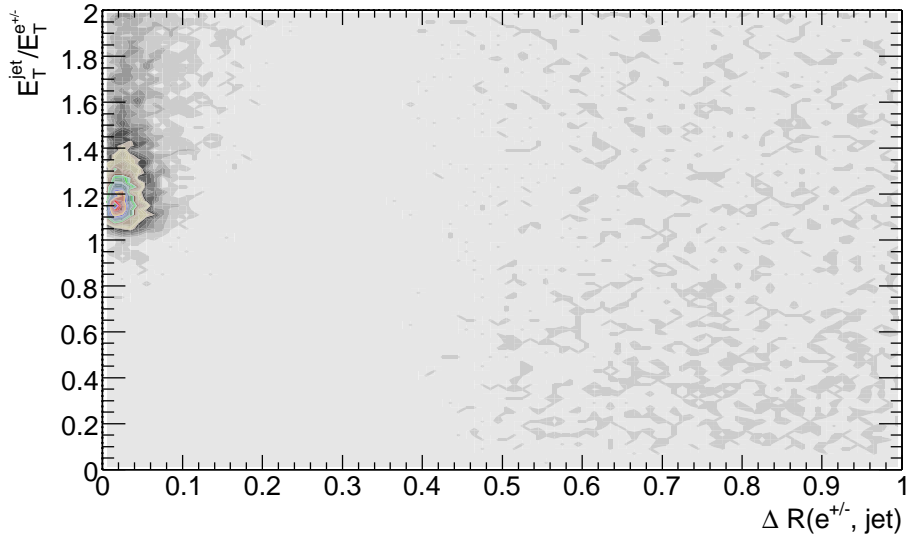




Figure 7.9: Electron/Jet overlap removal. Jet candidates are matched to well reconstructed electrons.

The ATLAS tagging algorithms fall into a number of broad categories. The largest class relies on likelihood functions based on spatial variables, such as the transverse impact parameter (d_0) or the number of reconstructed secondary vertices. Because the likelihood functions require a priori knowledge about the distribution of these variables, they will require varying amounts of real data to commission. Taggers which rely only on parameters in the transverse plane (such as IP2D, which is based solely on the transverse impact parameter) are expected to require 100pb^{-1} to commission. More complicated taggers (such as IP3D+SV1, which takes the 3D impact parameter as an input, as well as information from secondary vertices), will require much more data to commission [49]. The outputs for both of these taggers are shown in figures 7.10(a) and 7.10(b) for jets in the single top t-channel.


During early data taking, less complicated algorithms will be used [6]. The JetProb algorithm is based on the b-tagging algorithm developed at LEP and used at the Tevatron [50]. For all tracks associated with a jet, the significance of the transverse



(a) IP3D+SV1



(b) IP2D



(c) JetProb

Figure 7.10: Output from three different b -tagging algorithms used in the ATLAS reconstruction. The blue lines represent the weights from reconstructed jets that have been matched to truth b -jets. The red lines represent the weights from all other jets. See main text for details about the algorithms.

impact parameter (ie $d_0/\sigma(d_0)$) is compared to a resolution function, estimating the probability that the track emerged from the primary vertex. The probabilities for all tracks are then combined for each jet. The resulting probability can be used as a figure of merit as to how likely it is that the jet contained a relatively long lived particle (such as a B hadron). The output of this tagger can be seen in figure 7.10(c).

Although it does not discriminate between b and other jets as well as the likelihood taggers, it can be commissioned with any prompt tracks. Current estimates suggest that it might be commissioned with only 50 pb^{-1} of data, making it available for use during the initial data taking period [51].

7.2.6 Missing Transverse Energy

Leptonically decaying single top quarks will produce a neutrino, and this is expected to result in a significant missing transverse energy, \cancel{E}_T . The x and y components of energy deposits in calorimeter cells are summed to form the basis of the \cancel{E}_T calculation. Corrections are then added in for muons, which leave little energy in the calorimeters, and for energy losses in the Liquid Argon cryostat. Finally, cells associated with reconstructed objects (such as jets) are replaced in the energy sum, as they have a more accurate energy calibration. This final step results in a linearity better than 1% [6].

7.3 Event Selection

The single top t-channel cross section

$$\sigma_t = \frac{N_{Data}}{\epsilon_t \mathcal{L}} - \sum \frac{\epsilon_b \sigma_b}{\epsilon_t} \quad (7.1)$$

may be calculated from data by a simple counting analysis, where ϵ_t is the t-channel selection efficiency and ϵ_b is the efficiency of selecting events from the background process with cross section σ_b . At the truth level, the signal to background ratio (S/B) is 8×10^{-4} . Selection cuts which maximise the S/B ratio whilst maintaining a suitably high number of signal events have been developed.

7.3.1 Trigger and Lepton Identification

The lowest unrescaled isolated lepton triggers are applied to the simulated data. In the muon channel, the `EF_mu10` trigger menu item is required to be passed, which corresponds to an isolated muon with $P_T > 10$ GeV. The signal trigger and reconstruction efficiency is 0.30. In the electron channel, the `EF_e20_loose` trigger menu item is required to be passed. This corresponds to a $P_T > 20$ GeV electron candidate and results in an efficiency of 0.26. In both cases the efficiency is defined as the ratio between the number of events passing the trigger and the total number of events in the dataset.

A series of selection criteria are then applied to the primary (ie highest P_T) lepton (figure 7.11)². The lepton is required to have a $30 < P_T < 100$ GeV. An isolation condition is then applied, requiring that there be less than 6 GeV deposited within $\Delta R < 0.2$ in the calorimeters. A veto on secondary leptons is also applied, requiring that there be no further reconstructed electrons or muons, irrespective of channel, with $P_T > 10$ GeV.

²Note that the histograms in figures 7.11 - 7.18 have been normalised to unit area.




Figure 7.11: Primary lepton P_T . Note that each histogram has been normalised to unit area.

This secondary lepton cut is important not only because of its ability to reduce the Z+jets and dilepton $t\bar{t}$ channels, but also because it ensures that searches in the electron and muon streams are orthogonal - an event from the muon stream passing this selection is, by definition, not going to be found in the egamma stream, despite the inclusive streaming strategy employed by ATLAS.

Finally, two cuts using invariant mass are applied (figure 7.12). The transverse invariant mass of the primary lepton and E_T is required to be greater than 30 GeV whereas the invariant mass of the primary and remaining secondary leptons is required to be less than 65 GeV. These vetoes reduce the number of Z+jets events by 78% whilst reducing the number of single top events by 20%.

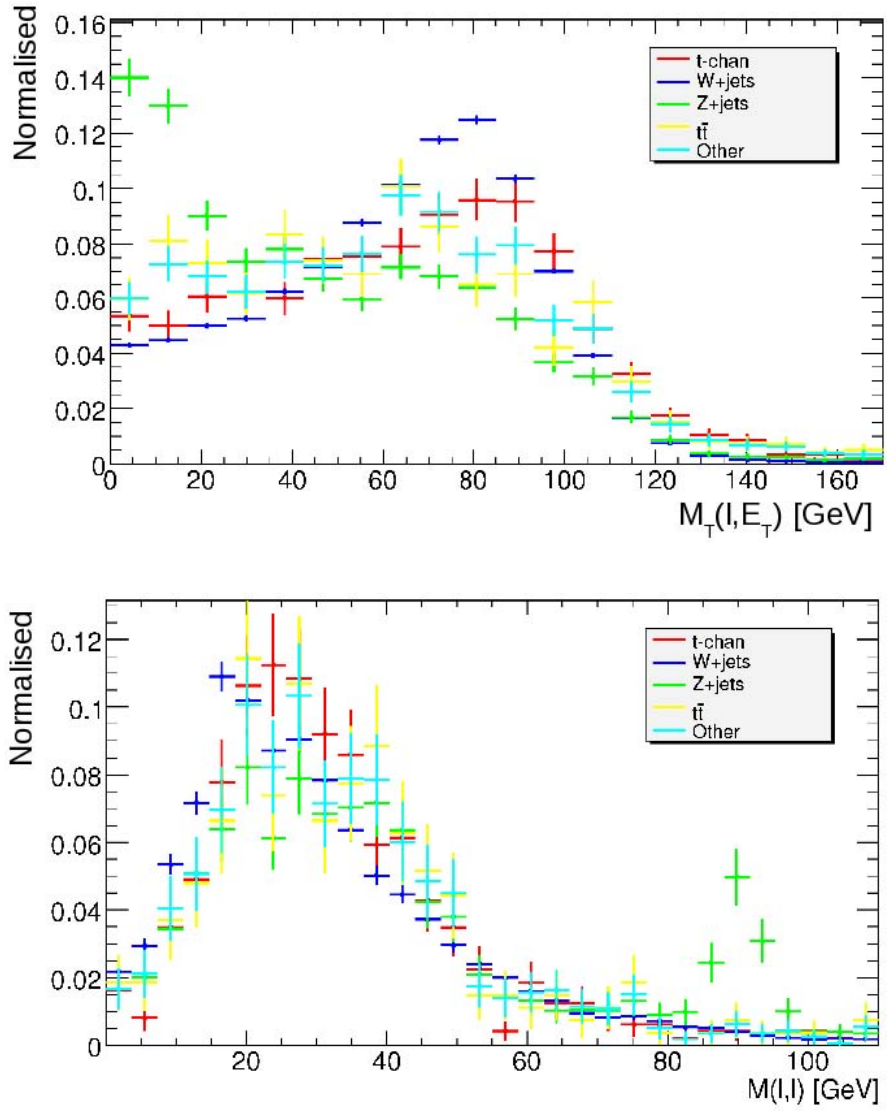


Figure 7.12: (a) Transverse invariant mass of the primary lepton and the E_T , and (b) invariant mass of the primary and secondary leptons. These plots show the distributions after the secondary lepton P_T cut, but before the mass cuts were applied.

7.3.2 Jet Identification

Reconstructed jets are first separated depending on whether they pass the b-tagging requirements. All taggable jets (ie jets with $P_T > 15$ GeV and in the range $|\eta| < 2.5$) passing the JetProb weight cut of 2.5 (figure 7.10(c)) are considered to be b-jets. This weight cut corresponds to a tagging efficiency in the t-channel simulation of 27%. The light jet rejection, defined as the inverse of the probability of mis-tagging a light jet, is 60. The b-jet is also required to have $P_T > 30$ GeV at this point (figure 7.13). Finally, the number of non-tagged jets with $P_T > 15$ GeV is limited to be between 2 and 4 inclusive (figure 7.14). This targets the W+jets channels at low multiplicity and the $t\bar{t}$ channels at high multiplicity.




Figure 7.13: Primary b-jet P_T .

7.3.3 Further Selection

After the initial preselection, the signal-to-background ratio is approximately 7%. This may be improved by exploiting the single top geometry. Figure 7.15, shows a typical t-channel truth event. The event is characterised by a light jet (yellow),




Figure 7.14: Number of untagged jets above 15 GeV.

recoiling against the highly boosted single top (blue). The decay products of the top move off in the same direction of motion as the top. This means that there is usually a large opening angle between the b-jet (red) and the light jet, and between the lepton (green) and the light jet. The angle between the b-jet and the lepton tends to be smaller. Cutting on these angles can improve the signal-to-background ratio and reduce sensitivity to systematic errors.

A powerful classifier is given by the centrality, c

$$c = \frac{E_T^{j_b} + E_T^{j_l}}{|E^{j_b}| + |E^{j_l}|}. \quad (7.2)$$

which is a function of energy of the primary b-tagged and the light jet. Events in the single top t-channel tend towards lower ratios, due to the high momentum of the forward light jet (figure 7.16). A hard cut on this value ($c < 0.15$) reduces the signal by a factor of 2, whereas the background is reduced by a factor of 6.

A further increase in the signal to background ratio can be achieved by exploiting the angular distribution of the single top decay products. The differences in η




Figure 7.15: Typical single top t -channel event geometry. The length of each line represents the parton momentum. The truth top quark (dark blue) recoils against a high energy d quark (yellow). The top quark is highly boosted, so its decay products (in this case a muon in green and a b -jet in red) move off in the same direction. Note that the length of a track is proportional to the particles momentum.




Figure 7.16: Centrality of b-jet and primary light jet.

between the primary b-tagged and non-tagged jet and the primary lepton and the non-tagged jet may be used as classifiers. In the single top, the primary non-tagged jet will typically recoil away from the top quark, resulting in larger angles between the non-tagged jets and the other primary objects (figures 7.17 and 7.18).

Requiring that $1.4 < \Delta\eta(j_b, j_l) < 5.0$ and that $0.8 < \Delta\eta(j_l, l) < 4.0$ reduces the background by 46% and the signal by 24%.

The numbers of remaining events after these cuts have been applied to 200 pb^{-1} of data are shown in table 7.2. The final signal-to-background ratio is 0.28 ± 0.02 (statistical error only).

7.4 Cross Section Measurement

An upper limit on the cross section, as a function of integrated luminosity, may be calculated by means of a Bayesian technique [1]. For a given confidence level, $1 - \alpha$,




Figure 7.17: $\Delta\eta$ between the primary light jet and the b-jet.




Figure 7.18: $\Delta\eta$ between the primary light jet and the lepton.

Process	Number Selected
<i>t-channel</i>	105 ± 6
<i>t\bar{t}</i>	171 ± 3
<i>W+jets</i>	169 ± 11
<i>Z+jets</i>	11 ± 2
<i>Other</i>	18 ± 2
<i>Total Background</i>	369 ± 11

Table 7.2: Number of selected events in 200 pb⁻¹ data.

the maximum cross section, $\sigma^{1-\alpha}$, is given by

$$1 - \alpha = \frac{\int_0^{\sigma^{1-\alpha}} \pi(\sigma) p(N_{Data}|\sigma) d\sigma}{\int_0^\infty \pi(\sigma) p(N_{Data}|\sigma) d\sigma} \quad (7.3)$$

where $1 - \alpha$ is taken to be equal to 0.95. The prior[52], $\pi(\sigma)$ is assumed to be 1 for $\sigma \geq 0$, and 0 for all other values. $p(N_{Data}|\sigma)$ is the probability density function for measuring N_{Data} for a given value of σ , and assuming a known, fixed background.

In the first instance, for a known background, N_B (where N_B is the sum over all background channels), and a signal efficiency ϵ_t , the probability density function may be assumed to be the Poisson likelihood function

$$P(\sigma) = \frac{(S + N_B)^{N_{Data}} e^{-(S+N_B)}}{N_{Data}!} \quad (7.4)$$

where $S = \epsilon_t \mathcal{L} \sigma$. However, there are uncertainties in ϵ_t and B due to the systematic errors, so that a nuisance parameter per systematic must be introduced such that

$$\epsilon_t \rightarrow \epsilon_t + \sum_j e_j$$

$$N_B \rightarrow N_B + \sum_j b_j$$

for j systematic errors. These nuisance parameters are assumed to have a Gaussian distribution, g , with mean 0 and width (either $\Delta\epsilon_t$ or ΔN_B) equal to the systematic error. Introducing the nuisance parameters requires that the probability density function given by equation 7.4 be transformed [53] such that

$$p(\sigma) = \int \int_{-\infty}^{\infty} P(\sigma) \cdot g(e, \Delta\epsilon_t) \cdot g(b, \Delta N_B) de db \quad (7.5)$$

In order to establish single top production at the LHC the null hypothesis, ie the background only hypothesis, must be first excluded to a reasonable level. Taking $N_{Data} = N_B$, equation 7.3 may be solved numerically [54] if the nuisance parameters can be quantified. While it is impossible to accurately predict exactly how these errors will affect the measurement from simulation alone, some initial estimates may be made on the more dominant effects. The following sections make initial estimates on how significant the nuisance parameters might be.³

7.4.1 Monte Carlo Statistics

There is a systematic error in the cross section measurement due to the limited Monte Carlo statistics available. For ΔN_B , the error is derived from the uncertainty in the various background efficiencies, ϵ_{B_i}

$$\frac{\Delta N_B}{N_B} = \sqrt{\sum_i \left(\frac{s(\epsilon_{B_i})}{\epsilon_{B_i}} \right)^2} \quad (7.6)$$

³The convention in the following passages is to denote the standard error on a quantity as $s(x)$, to avoid confusion with cross sections denoted σ .

A Binomial error is assumed on the efficiency, such that

$$\frac{\Delta\epsilon_{B_i}}{\epsilon_{B_i}} = \sqrt{\frac{1 - \epsilon_{B_i}}{\epsilon_{B_i} B_i}} \quad (7.7)$$

where B_i is the total number of events in each background dataset. A similar equation exists for the error on the signal efficiency. The nuisance parameters are shown in table 7.3. Note that although this form for ΔN_B and $\Delta\epsilon_t$ does not have a luminosity dependence (ie the scale of the error remains fixed for all integrated luminosities), the error may be reduced by using a larger simulated dataset.

	Δ
$\Delta N_B/N_B$	3.4 %
$\Delta\epsilon_t/\epsilon_t$	6.2 %

Table 7.3: Systematic effect of limited MC statistics on the expected number of background events and signal efficiency.

7.4.2 Luminosity

During very early running the luminosity will be estimated from machine parameters, leading to a fractional uncertainty as high as 20%. The error on the luminosity measurement is expected to drop to 5% later on [48]. This fractional error on the luminosity equates to the same fractional error in N_B , such that

$$\frac{\Delta N_B}{N_B} = \frac{s(\mathcal{L})}{\mathcal{L}} \quad (7.8)$$

Because of its dependence on luminosity, there is a similar fractional error on the

quantity S in equation 7.4.

7.4.3 Background Cross section

The cross section measurement relies on an accurate estimation of the expected number of background events, making it sensitive to uncertainties in the background cross sections. This may be expressed as

$$\frac{\Delta N_B}{N_B} = \frac{\sqrt{\sum_i \epsilon_{B_i}^2 S(\sigma_{B_i})^2}}{\sum_i \epsilon_{B_i} \sigma_{B_i}} \quad (7.9)$$

Taking the error on the theoretical calculation of $\sigma_{t\bar{t}}$ to be 6%, σ_{Wt} to be 3% and all other cross sections (W/Z+jets etc) to be 20% [44], the resulting $\Delta N_B/N_B$ is $\pm 4.3\%$.

This error is dominated by the W+jets cross section uncertainty. Assuming that the Monte Carlo accurately describes the shapes of the distributions, data-driven methods may be employed to derive normalisation factors, reducing the sensitivity of the analysis to the background cross section uncertainties. One study [6], suggests that the W+jets normalisation factor could be known to an accuracy of 5%, and the Z+jets to 3%. This improvement would reduce $\Delta N_B/N_B$ to $\pm 3.1\%$. Again, there is no associated nuisance parameter for the signal efficiency.

7.4.4 Parton Density Function

The majority of Monte Carlo samples used in this study took the CTEQ6L Parton Density Function [55] set as an input (the exception was the MC@NLO $t\bar{t}$ dataset,

which used CTEQ6M). The CTEQ PDF sets consist of fits to real data and therefore will have some degree of uncertainty which will in turn propagate into the cross section measurement. One method of estimating this uncertainty is to vary each of the free parameters in the fit (20 in total) by $\pm\sigma$ and judge the effect on the t-channel selection efficiency and the number of expected background events. CTEQ provide 40 auxiliary PDF sets representing these 2×20 variations.

Rather than regenerate 40 Monte Carlo datasets for each error PDF (which would be prohibitively CPU intensive), the original datasets may be weighted, on an event-by-event basis, so that the effect of the new PDF is taken into account [56]. The event weight is given by

$$w_i^\pm = \frac{f_1(x_1, Q, S_i^\pm) f_2(x_2, Q, S_i^\pm)}{f_1(x_1, Q, S_0) f_2(x_2, Q, S_0)} \quad (7.10)$$

where f_1 and f_2 are the PDF values for initial incoming partons with momentum fractions x_1 and x_2 , Q is the event scale and S_0 and S_i^\pm are the central (original) and error PDFs.

Having rescaled an event by w_i , a measurable (X) such as signal efficiency or number of background events, may be recorded. The expected difference in the measurable for the complete ensemble of PDF error sets is given by

$$\Delta^+ X = \sqrt{\sum_i^N \max(X_i^{\max} - X_0, 0)^2}$$

$$\Delta^- X = \sqrt{\sum_i^N \max(X_0 - X_i^{\min}, 0)^2}$$

where N is the number of free parameters in the PDF parametrisation (ie 20 in the

case of CTEQ6).

Table 7.4 shows the errors in ϵ_t and N_B derived from the CTEQ6M error sets.

	Δ
$\Delta N_B/N_B$	5.4 %
$\Delta\epsilon_t/\epsilon_t$	1.8 %

Table 7.4: Error in signal efficiency and number of expected background events due to uncertainty in PDF.

7.4.5 b-tagging

It has been shown that it may be possible to estimate the b-tagging efficiency from data with an absolute accuracy of 5%. In order to estimate the effect of this uncertainty on the cross section, the tag weight cut was shifted by an amount δ_i such that the efficiency of tagging jets labelled as truth b-jets shifted by $\pm 5\%$ (figure 7.19) [51]. Note that δ_i is specific to each channel. It has also been suggested that the relative error in the light jet rejection is $\pm 10\%$. In a similar manner to the b-tagging efficiency variation, the tag weight cut was adjusted for all non-truth b-jets (ie reconstructed jets that are better matched to a light truth quark or gluon) such that the rejection for each channel varied by $\pm 10\%$ [51].

The resulting variations in N_B and ϵ_t are displayed in table 7.5.

7.4.6 Jet Energy Scale

Previous studies have shown that the scaling between truth and reconstructed jet energies may have an uncertainty of between 5 and 20% depending on reconstruc-




Figure 7.19: *b*-tagging systematics variations. The weight cut is varied so that the efficiency for each channel varies by $\pm 5\%$. For example, in the single top *t*-channel shown here, a weight cut of 2.50 corresponds to an efficiency of 27%. Cuts of 1.88 and 3.30 correspond to efficiencies of 32% and 22% respectively.

tion location [57]. The selection criteria used in this analysis have been chosen to minimise the effect of this uncertainty where possible.

In order to estimate how the Jet Energy Scale uncertainty affects the cross section measurement two new data ensembles were produced - scaling all jet energies up or down by some proportion depending on their η co-ordinate (either $\pm 5\%$ for $|\eta| < 2.5$ and $\pm 10\%$ otherwise or a more conservative estimate of $\pm 10\%$ for $|\eta| < 2.5$ and $\pm 20\%$). The smaller scale energy distributions can be seen in figure 7.20.

In addition, because reconstructed jets are used in the final step of the calculation, the x and y components of the \cancel{E}_T must also be adjusted such that

$$\cancel{E}_{x,y} = \cancel{E}_{x,y} - \sum_i E_{x,y}^{j_i} \frac{s(E^{j_i})}{E^{j_i}} \quad (7.11)$$




Figure 7.20: Uncertainties due to the variations in the Jet Energy Scale. The black line represents the energy spectrum of all jets in the simulated t -channel events. The red line shows how the spectrum changes when each jet energy is scaled up by either a factor of 5 or 10% (dependent on η position). The blue line shows how the spectrum changes when each jet energy is scaled down.

	b-tagging efficiency $\pm 5\%$	Light jet rejection $\pm 10\%$
$\Delta N_B/N_B$	4.3%	3%
$\Delta \epsilon_t/\epsilon_t$	7.4%	$< 1\%$

Table 7.5: Uncertainties due to errors in the b-tagging efficiency and rejection. The b-tagging efficiency error dominates. Note that the required change in the b-tagging efficiency is absolute, whereas the change in the rejection is relative.

where E^{j_i} is the energy of each jet. Note that all jets, including those removed in the overlap procedure, are used in this calculation.

The selection criteria were applied to the scaled data ensembles, and an expected difference in ΔN_B and ϵ_t calculated. The uncertainties are summarised in table 7.6.

	$\pm 5/10\%$	$\pm 10/20\%$
	+	-
$\Delta N_B/N_B$	5.4%	5.7%
$\Delta \epsilon_t/\epsilon_t$	3.8%	6.0%

Table 7.6: Uncertainty due to the Jet Energy Scale.

It is interesting to note that the P_T cut on the b-tagged jet actually reduces the sensitivity to the JES. Figure 7.20 shows a large variation in the number of scaled jets below and around 25 GeV, whereas the number of jets in the region of 30-80 GeV shows very little variation after scaling. By imposing the P_T cut, and assuming a 10% uncertainty, only jets with a $P_T > 27\text{GeV}$ after scaling are eligible to be selected. As this P_T region is less sensitive to JES variations, the cut reduces the effect of the systematic error.

7.4.7 Cross Section Upper Limit

Assuming the background only hypothesis, H_0 , the 95% Confidence Level limit given by equation 7.3 may be taken as a figure of merit as to how well H_0 may be resolved from other hypotheses. Figure 7.21 shows how the H_0 upper limit may scale with integrated luminosity. Two different systematic scenarios have been plotted, which are detailed in table 7.7. Scenario 1 (the blue line) uses a more conservative set of systematic errors. In this case, the upper limit is completely dominated by the systematic effects, and lies just above the Standard Model prediction (the black dashed line) at 45.2 pb.

Error Source	Scenario 1	Scenario 2
Luminosity	10%	5%
JES	10/20%	5/10%
W+Jets Cross section	20%	5%

Table 7.7: Uncertainty scenarios for the luminosity and JES systematic errors. All other sources of error are consistent in both scenarios.

Scenario 2 uses a more optimistic set of systematic errors. In this case, the statistical effects appear to dominate for the first 100 pb⁻¹, before settling on a value 35.2 pb. This is below the standard model prediction, which suggests that this analysis, under the assumptions made for the systematics, would be able to distinguish between the Standard Model and background only hypotheses. Note that the upper limit has only been calculated for integrated luminosities above 50 pb⁻¹.

Using equation 7.5, it is possible to explore the significance of any measurement made from data for a fixed integrated luminosity. For example, figure 7.22 shows the probability distributions for the different systematic scenarios at 200 pb⁻¹ in the




Figure 7.21: Upper limit, calculated at the 95% confidence level, for the t -channel cross section as a function of integrated luminosity. The blue line shows a conservative estimate of the systematic errors. The red line depicts a more optimistic estimate. The black dashed line shows the Standard Model prediction of the cross section.

case where the Standard Model cross section is assumed. In this case, the probability of the measurement still being consistent with H_0 is between 1.27 - 3.93%, which corresponds to a significance of between $1.76 - 2.23\sigma$.

7.5 Conclusion

Electroweak single top production will be a challenging channel to observe at the LHC. The detection of this channel quickly becomes limited by systematic errors, and success depends on how well these systematics are understood and controlled. The methods explored in this study have been shown that with 200 pb^{-1} of data a t -channel production cross section greater than 35.2 pb could be excluded to the 95% confidence level. Alternatively, a measurement using this much data could




Figure 7.22: Probability distributions for different production cross sections assuming the Standard Model at 200 pb^{-1} .

correspond to a significance of between $1.76 - 2.23\sigma$ assuming the Standard Model prediction. However, it should be noted that this significance is likely to decrease when other factors, such as the QCD background, are taken into account.

Nevertheless, current estimates for the LHC physics programme suggest that in the first year approximately 200 pb^{-1} might be recorded at a collision energy of 10 TeV. If this estimate proves to be true, then it may be feasible to gain evidence of Standard Model single top t-channel production within one year of start up.

Appendix A

Glossary

CAN - *Controller Area Network, an industry standard field bus.*

CANopen - *An open message protocol for use on a CAN bus.*

CAM - *Clock Alignment Module. Used by the L1Calo trigger and monitored by the L1Calo DCS.*

CPM - *Cluster Processor Module. Used by the L1Calo trigger and monitored by the L1Calo DCS.*

CU - *Control Unit. High level node in the DCS FSM giving summary information of partitions. Differs from a Logical Unit (LU) in that users may take control of the partition from a Control Unit.*

DCS - *Detector Control System.*

DSS - *Detector Safety System.*

DU - *Device Unit. Lowest node in the FSM, represents a physical item of hardware.*

FSM - *Finite State Machine. Software construct used to model and control detector hardware.*

GCS - *Global Control Stations. Main interface to the DCS.*

JCOP - *Joint Controls Project. Collaboration between the LHC experiments.*

JEM - *Jet Energy Module. Used by the L1Calo trigger and monitored by the L1Calo DCS.*

LCS - *Local Control Station. Computers at the lowest level of the DCS backend used to slow controls data from the hardware.*

L1Calo - *Level-1 Calorimeter Trigger.*

LU - *Logical Unit. High level node in the DCS FSM giving summary information of a partition. Users may not take control of the partition at the level of a Logical Unit.*

OPC - *Ole (Object Linking and Embedding) for Process Control. A server/client model for making data available.*

PPM - *Preprocessor Module. Used by the L1Calo trigger and monitored by the L1Calo DCS.*

PVSS - *Integrated Design Environment used by all LHC experiments to create and run their DCS.*

ROD - *Read Out Driver. Used by the L1Calo trigger and monitored by the L1Calo DCS.*

SCS - *Subdetector Control Station. Organises and summarises groups of related*

Local Control Stations.

TCM - *Timing Control Module. Used by the L1Calo trigger and forms an integral part of the L1Calo DCS readout chain.*

TDQ - *Timing and Data Quality. In the context of the DCS, TDQ is a subdetector.*

Appendix B

Cut Flow

The number of events left in 200 pb^{-1} of data after each selection cut was applied. Note that the line labelled preselection gives the number events after the trigger condition has been applied, along with a few very loose preselection and overlap criteria used to reduce the analysis time. These criteria require:

- At least one lepton be reconstructed in a non-excluded region (for example, at least one electron reconstructed outside of the region $1.37 < |\eta| < 1.52$).
- At least 2 jets be reconstructed, one of which must have a tag weight of > 0.85 .

Cut	t-chan	$t\bar{t}$	W+jets	Z+jets	Other
-	8635	43412	9.7×10^6	864902	18082
Preselection	2674	20473	134674	28937	3548
$30 \text{ GeV} < l_p P_T < 100 \text{ GeV}$	1625	12566	82066	19292	2147
$l_s P_T < 10 \text{ GeV}$	1403	8729	80912	10295	1803
$M_T(l_p, \cancel{E}_T) > 30 \text{ GeV}$	1138	6986	67339	2296	1407
$M(l_p, l_s) < 65 \text{ GeV}$	1124	6878	66840	2260	1385
b-tag weight > 2.5	588	3574	9688	401	504
$jet_b P_T > 30 \text{ GeV}$	535	3410	4578	250	411
$2 \leq N_{jets} \leq 4$	327	1651	2466	134	240
$Centrality > 0.15$	146	292	371	26	35
$1.4 < \Delta\eta(j_b, j) < 5.0$	122	230	278	17	25
$0.8 < \Delta\eta(j_b, l) < 4.0$	105	171	169	11	18

Bibliography

- [1] C. Amsler et al. Review of particle physics. *Physics Letters B*, 667, Pages 1-1340, 2008.
- [2] F Abe et al. Observation of top quark production in p anti-p collisions with the collider detector at fermilab. *Phys. Rev. Lett.*, 74, 1995.
- [3] S Abachi et al. Observation of the top quark. *Phys. Rev. Lett.*, 74(14), 1995.
- [4] CDF Collaboration. First direct limit on the top quark lifetime. *CDF Note*, 8104, 2006.
- [5] V M Abazov et al. Experimental discrimination between charge $2e/3$ top quark and charge $4e/3$ exotic quark production scenarios. *Physical Review Letters*, 98, 2007.
- [6] The ATLAS Collaboration. Expected performance of the atlas experiment - detector, trigger and physics. (CERN-OPEN-2008-020), 2008.
- [7] Combined preliminary data on Z parameters from the LEP experiments and constraints on the Standard Model. (CERN/PPE/94-187), 1994.
- [8] Lep electroweak working group. 2009.
- [9] Marc-Andr Pleier. Review of top quark properties measurements at the tevatron. Technical Report arXiv:0810.5226. FERMILAB-PUB-08-495-E, Oct 2008.

Comments: 190 pages, 55 figures, 22 tables. To be published in Int. J. Mod. Phys. A.

- [10] Durham Online Parton Distributions. <http://durpdg.dur.ac.uk/hepdata/pdf3.html>.
- [11] A Quadt. Top quark physics at hadron colliders. *European Physical Journal C*, 48, 2006.
- [12] Z Sullivan. Understanding single-top-quark production and jets at hadron colliders. *Phys. Rev. D*, 70, 2004.
- [13] B W Harris et al. Fully differential single-top-quark cross section in next-to-leading order qcd. *Phys. Rev. D*, 66, 2002.
- [14] J Campbell and R K Ellis. Single top-quark production and decay at next-to-leading order. *Phys. Rev. D*, 70, Nov 2004.
- [15] A. Belyaev and E. Boos. Single top quark $t\bar{w}+x$ production at the cern lhc: A closer look. *Phys. Rev. D*, 63(3), 2001.
- [16] V M Abazov et al. Observation of single top-quark production. *Physical Review Letters*, 103, 2009.
- [17] Tim M. P. Tait and C.-P. Yuan. Single top quark production as a window to physics beyond the standard model. *Phys. Rev. D*, 63(1):014018, Dec 2000.
- [18] T. Aaltonen et al. Observation of electroweak single top-quark production. *Physical Review Letters*, 103, 2009.
- [19] G Mahon. Observing spin correlations in single top production and decay. *arXiv:hep-ph/0011349v1*, 2000.

- [20] Z Sullivan. What does single-top-quark production teach us about lhc physics. *Talk given at CTEQ Workshop, Physics at the LHC: Early Challenges*, May 2007.
- [21] W Wagner. Top quark physics in hadron collisions. *Reports on Progress in Physics*, 68, 2005.
- [22] L Evans et al. The lhc machine. *Journal of Instrumentation*, 3, 2008.
- [23] G Aad et al. The atlas experiment at the cern large hadron collider. *Journal of Instrumentation*, 3, 2008.
- [24] R Achenbach et al. The atlas level-1 calorimeter trigger. *Journal of Instrumentation*, 3, 2008.
- [25] Summary of lhc performance workshop - chamonix 2009. Feb 2009.
- [26] F Ragusa and L Rolandi. Tracking at lhc. *New Journal of Physics*, 9, 2007.
- [27] H Burkhardt and P Grafstrom. Absolute luminosity from machine parameters. Technical Report LHC-PROJECT-Report-1019. CERN-LHC-PROJECT-Report-1019, 2007.
- [28] A Barriuso Poy, H Boterenbrood, H J Burckhart, J Cook, V Filimonov, S Franz, O Gutzwiller, B Hallgren, V Khomutnikov, S Schlenker, and F Varela. The detector control system of the atlas experiment. *Journal of Instrumentation*, 3, 2008.
- [29] H J Burckhart. Recommendations for the usage of the detector safety system in atlas. Technical Report EDMS ID 754656, CERN, 2009.
- [30] Oliver Holme, Manuel Gonzalez-Berges, Piotr Golonka, and Sascha Schmeling. The jcop framework. Technical Report CERN-OPEN-2005-027, CERN, 2005.

- [31] ETM. <http://www.etm.at>.
- [32] P C Burkimsher. Jcop experience with a commercial scada product, pvss. <http://accelconf.web.cern.ch/accelconf/ica03/PAPERS/TU201.PDF>, 2003.
- [33] S Schlenker A Barriuso Poy. Atlas dcs fsm integration guide. Technical Report ATL-DQ-ON-0010, CERN, 2007.
- [34] O Gutzwiller. Atlas dcs common infrastructure control. Technical Report EDMS ID 964256, CERN, 2009.
- [35] CAN Bus. ISO 11898.
- [36] V Khmoutnikov. Atlas daq-dcs communication software. Technical Report ATL-DQ-ON-0008v.5.5, CERN, 2008.
- [37] DCS Online Status. <http://pcatdwww.cern.ch/atlas-point1/dcs/>.
- [38] R Barillere et al. Results of the OPC evaluation done within JCOP for the control of the LHC experiments. Prepared for 7th International Conference on Accelerator and Large Experimental Physics Control Systems (ICALEPCS 99), Trieste, Italy, 4-8 Oct 1999.
- [39] H Boterenbrood. Canope high-level protocol for can-bus. 2000.
- [40] Wiener Plein and Baus Elektronik. <http://www.wiener-d.com/index2.php>.
- [41] Fujitsu. F2mc-16lx mb90590/590g series data sheet. Technical Report DS07-13704-5E.
- [42] B Kersevan and E Richter-Was. The monte carlo event generator acermc and package acerdet (version 3.5). 2008.

- [43] Z Sullivan T Stelzer and S Willenbrock. Single top-quark production via w-gluon fusion at next to leading order. (hep-ph/9705398v2), 1997.
- [44] M Bosman et al. Understanding monte carlo generators for top physics. *ATLAS Note*, ATL-COM-PHYS-XXX, 2009.
- [45] S Frixione and B Webber. The mcnl0 3.1 event generator. (hep-ph/0506182v1), 2005.
- [46] V M Abazov et al. Measurement of the ttbar production cross section in ppbar collisions at $\sqrt{s} = 1.96$ tev using kinematic characteristics of lepton + jets events. *Physical Review D*, 75, 2007.
- [47] M L Mangano et al. Alpgen, a generator for hard multiparton processes in hadronic collisions. *JHEP*, 07, 2003.
- [48] The ATLAS Collaboration. The atlas experiment at the cern large hadron collider. *JINST* 3, (S08003), 2008.
- [49] G Piacquadio on behalf of the ATLAS Collaboration. B-tagging at lhc: Expected performance and its calibration in data. (ATL-PHYS-PROC-2008-070), 2008.
- [50] The CDF Collaboration. Measurement of the ttbar production cross section in ppbar collisions at $\sqrt{s} = 1.96$ tev using lepton + jets events with jet probability b-tagging. *Physical Review D (Particles and Fields)*, 74(7), 2006.
- [51] The ATLAS Collaboration. Expected performance of the atlas experiment - detector, trigger and physics: b-tagging performance chapter. (CERN-OPEN-2008-020), 2008.
- [52] E Gross. Lhc statistics for pedestrians. CERN-2008-001.

- [53] T Hebbeker. Calculating upper limits with poisson statistics. L3 Note 2633.
- [54] J Heinrich. Bayesian limit software: multi-channel with correlated backgrounds and efficiencies. *CDF Statistics Memo*, (7587), 2005.
- [55] J Pumplin et al. New generation of parton distributions with uncertainties from global QCD analysis. *JHEP*, 07, 2002.
- [56] G Khorauli and M Cristinziani. Review of pdf re-weighting technique. *Talk given at ATLAS MC Generator Meeting, 26 August 2008*, May 2008.
- [57] M Hurwitz. Understanding the jet energy scale using photon-jet balancing. *ATLAS Internal Note*, ATL-COM-PHYS-2009-046, 2009.
- [58] R Brock et al. Handbook of perturbative qcd. *Rev. Mod. Phys.*, 67, 1995.
- [59] Z Sullivan. Angular correlations in single-top-quark and wjj production at next-to-leading order. *Phys. Rev. D*, 72, 2005.
- [60] L Vacavant. b-tagging algorithms and performance in atlas. Technical Report ATL-PHYS-PROC-2009-023. ATL-COM-PHYS-2009-016, CERN, Geneva, 2009.
- [61] B Clement et al. Search for single-top events at 10 tev with the atlas detector. Technical Report ATL-COM-PHYS-2009-303, CERN, Geneva, 2009.
- [62] P Francavilla. Jet calibration in the atlas experiment at lhc. Technical Report ATL-PHYS-PROC-2009-108. ATL-COM-PHYS-2009-514, CERN, Geneva, 2009.