Speckle Interferometry and a Study of Γ_W for the ATLAS Experiment

Daniel Carlos Buira Clark

Magdalen College

University of Oxford

Thesis submitted in partial fulfilment
of the requirements for the degree of
Doctor of Philosophy

University of Oxford, Hilary Term 2006.

Speckle Interferometry and a Study of Γ_W for the ATLAS Experiment

Daniel Carlos Buira Clark
Magdalen College, University of Oxford

Hilary Term 2006

Thesis submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy

Abstract

The ATLAS experiment is being assembeled as a general purpose particle detector to exploit the physics discovery potential of the Large Hadron Collider. The SemiConductor Tracker (SCT), one of the central subsystems of ATLAS, requires alignment of $\mathcal{O}(10\mu\text{m})$ for good performance. The alignment program consists of a set of proceedures to establish the location of detector elements. This includes optimised design and construction prior to operation.

This thesis describes Electronic Speckle Pattern Interferometry (ESPI), a technique used to measure deformation at the micron level which has been applied to SCT components during design and prototype development. ESPI creates contour maps of objects under study, helping a qualitative understanding of object behaviour as well as providing measurements. ESPI measurements of the thermal deformation of SCT detector modules are presented.

A modified ESPI interferometer was developed to create contours of vibration amplitude. Tests were performed on well-understood objects to ensure consistency between the interpretation of Vibration ESPI results and known object behaviour. The technique was then used to measure the behaviour under acoustic driving of prototype SCT support structures made of Carbon-Fibre Reinforced Plastic (CFRP) materials. Several low modes of oscillation were found, as well as the oscillator quality factor, for each of the CFRP structures studied. These results provide design input for SCT construction by specifying resonant frequencies to avoid, and corroborate the validity of Finte Element Analysis studies on these inhomogeneous materials used in the ATLAS SCT.

High event statistics will permit ATLAS perform precision Electroweak measurements. This thesis presents a study of the statistical resolution with which the decay width of the W boson, Γ_W will be measured.

This thesis was typeset with LaTeX 2_ε.

Published at the University of Oxford, Oxford, United Kingdom.

To my family, with love

Acknowledgements

It is a pleasure to thank those who have helped me complete this work. Richard Nickerson, my supervisor, with his combination of patience and focus, has brought to bear a clarity that has proven very valuable. Tony Weidberg has also been a great help, always encouraging and willing to think about things. David Howell, ever generous with his time, made significant contributions to my understanding of optics, interferometry, and many other areas of physics applied to the real world, for which I am very grateful.

The fellow graduate students with whom I shared room 613 deserve a round of applause: Adrian Fox-Murphy, Stephen Hunt, Anish Grewal, Paul Coe, and Ankush Mitra. The companionship and the love of physics would each on their own be enough to earn my gratitude. To have both has been wonderful. Thanks guys.

The Oxford HEP group has been a stimulating environment in which to develop as a physicist. Many other graduate students have provided advice, enthusiasm, clear thinking, and fun times. The academics posses an enormous passion for their subject, and their drive is what makes Particle Physics happen. I am grateful, both for specific guidance and overall inspiration, to Todd Huffman, Armin Reichold, and Farrukh Azfar.

Others within the field have spent time helping me come to grips with the tools of High Energy Physics. I'd like to thank Michel Lefebvre, Elzbieta Richter-Was, and Törbjorn Sjöstrand.

Much of this work has required the participation of several people within the different groups of the Department. First mention goes to Paul Pattinson, without whom it truly would not have been possible. His dedication, common sense, and good cheer proved essential when stumbling upon the unforseen as we worked towards inflexible deadlines. Roy Wastie and Pete Shield helped with several aspects of electronics, including the design and construction of the VESPI control circuit. Brian Hawes was a great help when taking many of the ESPI measurements. Alan Holmes, Wing Lau, and Stephanie Yang of the Design Office helped a great deal with the FEA modelling. My thanks to all of you.

Another important thank-you goes to computing support. Pete Grönbech, Ian McArthur, Chris Hunter have all helped in many ways, as has the whole Oxford team. Andrew Samsun and Jonathan Wheeler were always helpful and responsive in their management of the RAL Central Simulation Facility where the simulation work of Chapter 6 was carried

out.

Another essential contribution came from the generous financial support of the Consejo Nacional de Ciencia y Tecnologia (CONACyT), México (no. reg. 92865), which is gratefully acknowledged.

There are also many to thank outside the world of Physics. At Magdalen I met Marnix, Jeff, Hannah, Gian Paolo†, and Andy. Through the Mexican Society I got to know Gabriel, Lorena, Andres, Jose Luis & Fabiola & Griselda, Raul & Rocio, and Cesar. Without these friends my Oxford story would have been very different.

Friends from back home also did their bit, just by being friends. Looking at an old photo, I should single out Carlos, since we somehow ended up doing similar things for most of the past 18 years, and he introduced me to a very cute friend of his who is now my wife. I'd also like to mention the other Carlos† of the group - even though he didn't make it to the photo, we're not going to forget him.

The deepest gratitude goes to my family. My parents have been there, unconditionally, from before the start with love, encouragement, and strength, as well as wise words. Sebastian, Lau, Fina, Jorge, and Ann have also each played their role. Thank you.

Regardless of the doctoral degree, I've already gained the best possible reward from pursuing it: the World's Greatest Wife. She has supported, encouraged, and guided me, giving me confidence, strength, security, joy, and love. Every day she inspires me to live life to the full, and that's not the half of it. Gracias $A\gamma\gamma\epsilon\lambda\iota\kappa\acute{\eta}$, $\sigma'\alpha\gamma\alpha\pi\acute{\alpha}\omega$ $\gamma\iota\alpha$ $\pi\acute{\alpha}\nu\tau\alpha$.

Contents

1 Introduction **1**

1.1 The Standard Model and How to Test It 1

1.2 Electronic Speckle Pattern Interferometry and ATLAS SCT Alignment . 16

1.3 Monte Carlo Studies of LHC Electroweak Physics and Γ_W 17

2 Speckle Interferometry **19**

2.1 Principles of Speckle Interferometry 19

2.2 Principles of Time Averaged ESPI 29

2.3 Noise and Visibility in ESPI 34

3 The Oxford Static Speckle Interferometer **38**

3.1 Physical Requirements for ESPI 38

3.2 ESPI Hardware 40

3.3 Information Processing 48

3.4 Measurements 49

4 The Oxford Vibration Speckle Interferometer **55**

4.1 Vibration of ATLAS SCT Structures 55

4.2 Assembly of the VESPI interferometer 60

4.3 Tests of VESPI 62

4.4 Measurements of ATLAS SCT Prototypes 70

5 The W width and the transverse mass method at the LHC **92**

5.1 The W boson at the LHC 92

5.2 Measurement techniques of the W decay width 94

5.3 The Transverse Mass spectrum 96

5.4 Γ_W from the M_T tail 98
5.5 Data rates and triggering 102
5.6 Sources of systematic error 103

6 The ATLAS statistical resolution on Γ_W using the M_T method 105
6.1 Scope of this study 105
6.2 W Event Generation 106
6.3 Data Analysis and Results 111
6.4 Summary of Results 120

A Speckle Statistics and Autocorrelation 122
A.1 Objective Speckle and Speckle Correlation 122
A.2 Subjective Speckle Autocorrelation 127

B Binned Log Likelihood Fit 129
B.1 The Maximum Likelihood Method 129
B.2 The Binned Log Likelihood Method 130
B.3 Finding the likelihood or log likelihood as a function of the parameter . . 130

List of Figures

1.1 The LHC 7
1.2 The ATLAS Experiment 9
1.3 The ATLAS Inner Detector 10

2.1 A Simple Speckle Pattern 20
2.2 Airy pattern 20
2.3 Speckle formation 20
2.4 Diagram of surface displacement in ESPI 22
2.5 ESPI circular deformation pattern due to pressure change 24
2.6 Out-of-plane sensitive interferometer 27
2.7 In-plane sensitive interferometer 28
2.8 Out-of-plane interferometry of an oscillating surface 32
2.9 Modulus of the zeroth-order Bessel function 33

3.1 Optics layout of the Oxford Speckle Interferometer 41
3.2 Camera, beam splitter, and reference fibre 43
3.3 Experimental setup for Out-of-plane ESPI 47
3.4 Experimental setup for In-plane ESPI 48
3.5 Diagram of a KEK End Tap module 51
3.6 Thermal deformation of a KEK End Tap module 51
3.7 Diagram of RAL End Tap module 52
3.8 Thermal deformation of RAL End Tap module 53
3.9 Thermal deformation of RAL module support structure 54

4.1 Flügge theory plots of vibrations in cylindrical shells 57
4.2 Discrepancies between different analytical theories of thin shell vibration 58

4.3 Diagram of camera timing circuit for VESPI 63
4.4 Circular membrane under tension, lowest vibrating modes 63
4.5 Drum test piece 64
4.6 Tunable drum test piece driven at mode (1,1) 65
4.7 Circular capacitor and guard ring 66
4.8 Diagram of bracket used to hold metal test plate in place 67
4.9 Metal test plate VESPI amplitude measurement 68
4.10 Metal test plate covered in polymer film, VESPI amplitude measurement 70
4.11 Cutaway view of a Honeycomb Sandwich structure 71
4.12 VESPI experimental setup for CFRP square plate measurements 74
4.13 VESPI and FEA modes for CFRP square plate, part one 76
4.14 VESPI and FEA modes for CFRP square plate, part two 77
4.15 CFRP square plate Quality Factor 78
4.16 Photograph of CFRP cylinder in VESPI lab 80
4.17 Experimental setup for CFRP cylinder measurements 82
4.18 Modal classification of a vibrating cylinder clamped at both ends 83
4.19 Cylinder VESPI and FEA results: (2,1) mode 84
4.20 Cylinder VESPI and FEA results: (4,2) mode 85
4.21 Cylinder VESPI and FEA results: (4,3) mode 85
4.22 CFRP cylinder mode splitting amplitude, (2,1) 87
4.23 Quality factor of CFRP cylinder 90

5.1 A typical M_T spectrum 97
5.2 Jacobian Peak 98
5.3 Jacobian peak after detector smearing effects 99
5.4 Normalised gaussian and lorentzian 99
5.5 Effect of varying Γ_W on M_T spectrum 100

6.1 Effect of varying the WGAM parameter 108
6.2 Jacobian migration high-tail events into low M_T region 109
6.3 Generation cuts and CPU time savings 110
6.4 Likelihood as function of Γ_W 113
6.5 Variation of Γ_W and σ_{Γ_W} due to finite template statistics 116

A.1 Spatial intensity variation: the Objective Speckle Pattern 123

A.2 Effect of an apperture on an image plane: the Subjective Speckle Pattern 128

List of Tables

1.1 The Elementary Particles of the Standard Model 3
1.2 LHC Operating Parameters 7
1.3 ATLAS Readout Channels 8

2.1 Zeros of the zeroth-order Bessel function 34

4.1 Square plate VESPI and FEA results comparison 75
4.2 CFRP Cylinder VESPI and FEA results 86

6.1 Effect of PDFs on Γ_W 114
6.2 Estimate of template statistics required for LHC analysis 118
6.3 Estimate of possible PDF systematic effect 119

Chapter 1

Introduction

1.1 The Standard Model and How to Test It

1.1.1 Why build the World's Biggest Machine?

The Large Hadron Collider (LHC) is the next-generation particle physics accelerator, which is scheduled to begin operation in 2007 in CERN, the European Laboratory for Particle Physics near Geneva, Switzerland. This 27 km circular machine will collide protons into each other with a centre-of-mass energy of $\sqrt{s}$=14 TeV. The energy provided *to each proton* is about seven times the kinetic energy of a flying mosquito [CER05]. Because a mosquito is so much more massive than a proton, this is proportional to taking all the kinetic energy necessary to make twenty *billion* Cheops Pyramids travel at the speed of a commercial airliner, placing it on a 1 kg object, and colliding it head on with another 1 kg object travelling with the same energy in the opposite direction.

The ATLAS (A Toroidal LHC Apparatus) experiment is being assembled around one of the interaction points, in order to collect the resulting debris from the proton collisions, which will be the most energetic to have been studied to date in this manner. ATLAS, which will be five storeys high, will surround the interaction point with an array of detectors feeding 150 million electronic readout channels. It will register copious amounts of raw data from the collisions: about one terabyte per second, which is equivalent to the total information exchange across all telecommunications networks of Western Europe at peak usage. This data will be filtered over three real-time stages, and eventually "only" 2.5 millionths of what is generated will be stored and analysed.

It is anticipated that the majority of the data will be highly consistent the current

theoretical framework of particle physics, the Standard Model (SM), permitting improved precision measurements of several known SM parameters. Furthermore, the SM itself predicts certain - as yet undiscovered - phenomena to exist within the LHC energy regime, whose confirmation or disproof would both be of great interest. Finally, since the accelerator will probe an energy regime never before explored in this manner, many physicists hope it will yield phenomena that are beyond the Standard Model. This would enrich scientific knowledge by stretching - or breaking - the theoretical framework currently used to explain the nature of matter, and the underlying structure of the universe itself.

1.1.2 The Standard Model of Particle Physics

The Standard Model, the context in which the search for the fundamental constituents of matter is currently pursued, has emerged by piecing together concepts over the past 50 years. It is described in several references, including [HM84], [AH89], and [Gri87], and will thus only be outlined briefly.

The Standard Model uses relativistic quantum field theory to describe the known particles of matter and the forces that govern their behaviour by means a Lagrangian which displays local gauge invariance under certain groups of transformations. The SM Lagrangian starts with terms representing the known spin $\frac{1}{2}$ fermions - quarks, leptons, and neutrinos. Gauge invariance under the group symmetry U(1)$\otimes$SU(2)$\otimes$SU(3) is then required, accounting for the electromagnetic, weak, and strong forces, and their corresponding exchange particles, the spin 1 gauge bosons: the photon (γ), the W and Z, and the gluons. An additional field is then added to account for the observed masses of the W and Z by means of spontaneous symmetry breaking - the Higgs field. The masses of all other particles, not previously explained by the model, can then be associated with each particle's coupling to the Higgs field. The particles of the Standard Model are summarised in table 1.1.

The Standard Model is solidly based on experimental evidence. Many of the key ideas in the model were first suggested by experimental observation. To date no significant discrepancy between the theory and experimental observation has been found. Most compellingly, the SM predicted the existence of the W, Z, and top quark before they were discovered. The only remaining SM particle yet to be found is the Higgs boson. In spite of these successes, there are numerous reasons for which physicists believe the SM does not tell the whole story. Unresolved questions include:

	Fermion Generation 1st	2nd	3rd	Q	I	I_3	Y	$SU(3)_C$
Quarks	$\begin{pmatrix} \mathrm{u}_{\text{1-5}} \\ \mathrm{d}'_{\text{3-9}} \end{pmatrix}_L$	$\begin{pmatrix} \mathrm{c}_{\text{1.15-1.35 GeV}} \\ \mathrm{s}'_{\text{75-170}} \end{pmatrix}_L$	$\begin{pmatrix} \mathrm{t}_{\text{174.3 GeV}} \\ \mathrm{b}'_{\text{4.0-4.4 GeV}} \end{pmatrix}_L$	$\frac{2}{3}$ $-\frac{1}{3}$	$\frac{1}{2}$ $\frac{1}{2}$	$\frac{1}{2}$ $-\frac{1}{2}$	$\frac{1}{3}$ $\frac{1}{3}$	3 3
	u_R	c_R	t_R	$\frac{2}{3}$	0	0	$\frac{4}{3}$	3
	d_R	s_R	b_R	$-\frac{1}{3}$	0	0	$-\frac{2}{3}$	3
Leptons	$\begin{pmatrix} \nu_{e\ \text{< 3eV}} \\ \mathrm{e}_{\text{0.511 MeV}} \end{pmatrix}_L$	$\begin{pmatrix} \nu_{\mu\ \text{< 0.19(95\% CL)}} \\ \mu_{\text{105.7}} \end{pmatrix}_L$	$\begin{pmatrix} \nu_{\tau\ \text{< 18.2(95\% CL)}} \\ \tau_{\text{1776}} \end{pmatrix}_L$	0 -1	$\frac{1}{2}$ $\frac{1}{2}$	$\frac{1}{2}$ $-\frac{1}{2}$	-1 -1	1 1
	e_R	μ_R	τ_R	-1	0	0	-2	1

	Gauge and Higgs Bosons			Q	I	I_3	Y	$SU(3)_C$
Spin 1	γ	QED	zero mass	0	0	0	0	1
	Z^0	EW	91.19 GeV	0	1	0	0	1
	W^+		80.42 GeV	1	1	1	0	1
	W^-		(same as W^+)	-1	1	-1	0	1
	$g_a (a = 1 \ldots 8)$	QCD	zero mass	0	0	0	0	8
Spin 0	h^0	Higgs	> 107.7GeV(95% CL)	0	0	0	0	1

Table 1.1: The Elementary Particles of the Standard Model: All left-handed fermions are grouped into doublets with masses given underneath. The right-handed fermions have identical masses to their left-handed partners. All masses are in MeV unless stated. The down-type quarks are the weak eigenstates. They are related to the mass eigenstates via the CKM mixing matrix. The quantum numbers are given in the right hand columns with electric charge defined as $Q = I_3 + \frac{1}{2}Y$ and normalised to the electron charge; where I is the weak isospin and I_3 is its third component and Y is the hypercharge. In addition the quarks carry colour charge. The gauge bosons are given in the bottom table with their mass and quantum numbers. A complementary table exists for antiparticles. [Gro00]

Massive Neutrinos [Ahm01] [Fuk01] The combined data from SNO and Super-Kamiokande neutrino observatories provides strong evidence that neutrinos of one flavour oscillate into another. The favoured theory is that neutrinos are massive and cross-coupling is related to a mixing matrix. This is not described by the SM, although it could be incorporated through an ad-hoc extension.

19 free parameters The SM does not predict particle masses or coupling constants, they must all be inferred from experiments. In a manner analogous to other physical theories for which underlying explanations were later discovered, such as the Periodic Table of Elements, it is hoped several of these currently independent parameters may be related through physical mechanisms involving other, more "fundamental" objects and theories.

Origin of Three generations There is no theoretical constraint on the the number of generations. Ordinary matter just requires the first generation. The extra generations could point to a higher group structure.

Dark Matter Astronomical observations have shown the universe contains a significant amount of matter which does not participate Electromagnetic interactions. This is expected to be massive non-baryonic weakly interacting matter. No such particle type exists within the SM.

Origins of CP Violation This phenomenon has been observed and incorporated into the SM by the introduction of the CKM matrix and a non zero phase. However the physical origins of such a phase and the values of the CKM matrix have not been explained.

Possibly Absent Higgs All the elementary particles of the SM have been discovered except the Higgs Boson. Global fits to measured electroweak data suggest a Higgs mass of around 100 GeV [CGHFG]. If the Higgs is not discovered, some other mechanism for EW symmetry breaking must be found.

The Hierarchy Problem As it stands, the SM Higgs boson, which is a scalar field, is subject to loop corrections to its self-energy which are quadratically divergent. This means that without extreme fine tuning, the Higgs mass would be expected to run to

a scale 30 orders of magnitude greater than the $\mathcal{O}$ (TeV) scale needed for compatibility with existing Electroweak physics measurements [DM73]. A theoretical mechanism to keep the Higgs mass at the same "hierarchy" as the rest of EW physics would extend and improve the model.

Inclusion of Gravity A theory which aims to account for all elementary interactions must describe gravity. The SM does not address the gravitational force at all.

In an attempt to address some of these questions, new theories have been proposed which aim to include SM physics but also explain phenomena outside its regime, hopefully explaining the SM as a partial theory applicable to the sub-TeV energy regime. These theories are collectively referred to as "Physics Beyond the Standard Model". Three of these ideas are mentioned briefly.

SuperSymmetry (SUSY) [ATLb] This theory assumes a symmetry exists between fermions and bosons, thus allowing for the existence of spin 1 partners to the known leptons and quarks (referred to as *sfermions*) as well as spin $\frac{1}{2}$ partners to the gauge particles (called *gauginos*). It is a maximal extension of the Lorentz group and provides a solution to the hierarchy problem by offering a new particle with different spin for every existing particle, which automatically cancels the diverging loop corrections. Precision data from LEP and SLC show the coupling constants are currently not expected to unify at a higher energy scale but the inclusion of SUSY could allow for this. Thus SUSY may lead to Grand Unified Theory which encompasses the forces into a common group such as SU(5) or SO(10).

Technicolour [ATLb] Technicolour provides a dynamical mechanism for breaking electroweak symmetry and so is a replacement for the Higgs mechanism. It assumes the existence of technifermions which carry technicolour charge and interact strongly. Technipions are created when techniquark condensates break EW symmetry.

Large Extra Dimensions [AHDD98] Recent theoretical work has explored the possible existence of large extra dimensions of $\mathcal{O}$ (mm). It proposes the SM is confined to a four-dimensional plane, but gravity is allowed to propagate in all dimensions. This

reduces the Planck scale to the electroweak scale which allows quantum gravity to be observed at $\mathcal{O}$ TeV and solves the hierarchy problem.

1.1.3 The Large Hadron Collider [LP]

The LHC is a high-energy, high-luminosity proton-proton synchrotron collider being assembled within the 27 km-long circular tunnel which previously housed the Large Electron Positron (LEP) Collider at CERN. It will be fed by the SPS and its pre-accelerators with 450 Gev protons, further accelerating them up to 7 TeV [LP]. The centre of mass energy of 14 TeV means the collisions will recreate the energy density of the universe 10^{-12} s after the Big Bang. The full design luminosity of $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ will lead to about 20 proton-proton collisions per bunch crossing, and the bunch-crossing frequency of 40 MHz will lead to very high collision statistics [1]. The accelerator is expected to operate for the first three years at lower luminosity, $10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Significant technical challenges have been overcome in order to construct the LHC. The target centre of mass energy combined with the fixed tunnel size has meant including superconducting (Niobium-Titanium) bending magnets to generate the required 8.4 T field. The magnets will operate at a temperature of 1.8 K, cooled by superfluid liquid Helium. Space and cost considerations have lead to an innovative "two-in-one" magnet design which allows both rings to be contained in the same cryostat.

Four experiments are under construction to study collisions at the LHC. Given the high energy and collision rate, they also face significant technical challenges. ATLAS and CMS will both be general purpose experiments, needing to fully surround and contain events with layers of high-resolution detectors, and managing extremely high volumes of raw data. LHC-b will be a forward experiment dedicated to bottom-quark physics, aiming to learn more about CP violation. ALICE is the fourth experiment, dedicated to the study of heavy ion (Pb-Pb) collisions which are also programmed for the LHC, in order to learn about the quark-gluon plasma.

Some of the key parameters of the LHC are presented in Table 1.2.

[1] The number of events generated per year in all of the relevant physics channels will be a factor of 10 to 10 000 times greater than the total number accumulated in previous experiments [ATLa].

Ring Circumference	26.7 km
Number of Main Bends	1232
Field at Main Bends	8.386T
Number of Bunches	2835
Bunch Spacing	7.48 m
	25 ns
Proton Energy	7 eV
Collision Energy ($\sqrt{s}$)	14 TeV
Total Cross Section	100.0 mb
Inelastic Cross Section	60.0 mb
Luminosity	$10^{34}\text{cm}^{-2}\text{s}^{-1}$
Number of Events per Collision	≈20
RMS Beam Size is x,y at IP	16.7 μm
Lead-Ion Beam Data	
Total Centre-of-Mass Energy	1148 TeV
Energy per Nucleon	2.76 TeV
RMS beam radius at Collision	15.9 μm

Table 1.2: LHC Specifications [?]

.

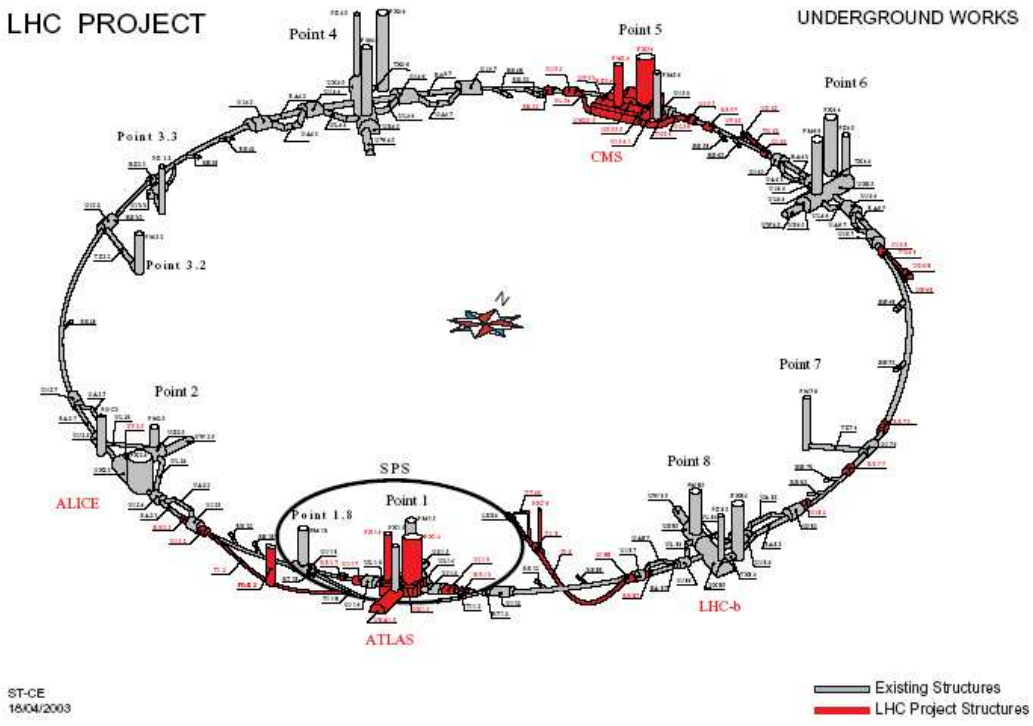

Figure 1.1: The LHC.

Detector System	Number of Active Elements
Pixels	140 000 000
SCT	6 280 000
TRT	420 000
LAr accordion calorimeters (em)	170 000
LAr & end-cap calorimeters (hadronic)	9 000
Tile calorimeters (hadronic)	10 000
Muon System	1 230 000

Table 1.3: ATLAS Read-out Channels per System [ATLa]

.

1.1.4 The ATLAS detector [ATLa]

ATLAS is a general purpose detector designed to exploit the full physics potential of the LHC. At the time of writing it is being assembled, and is expected to become operational in 2007. The detector must be able to make accurate measurements over a broad range of channels while operating at full LHC luminosity ($10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$). This detector has been planned with built in redundancies, and aims to use multiple signatures for internal cross-checking. This is necessary to face the harsh LHC environment.

In order to meet its goals, the detector must have excellent calorimetry, being able to identify and measure the energy of electrons and photons to high precision, even at full luminosity. The hadronic calorimetry must also have good precision and full coverage for jet and missing energy reconstruction. Stand-alone precision muon tracking is also necessary. The detector must have large η acceptance as well as full ϕ coverage. At high luminosity, ATLAS must be capable of efficient high p_T lepton momentum measurements, electron and photon identification, and τ and heavy flavour identification. At low luminosity, full event reconstruction is desired. In addition, suitable triggering is needed to ensure efficient selection of rare physics events.

Figure 1.2 gives a cut-away view of the ATLAS detector. The configuration is typical for a colliding beam detector with tracking chamber at the centre surrounded by calorimeters. The outermost layer is the muon spectrometer. Table 1.3 shows the number of active elements, and therefore readout channels, present in ATLAS.

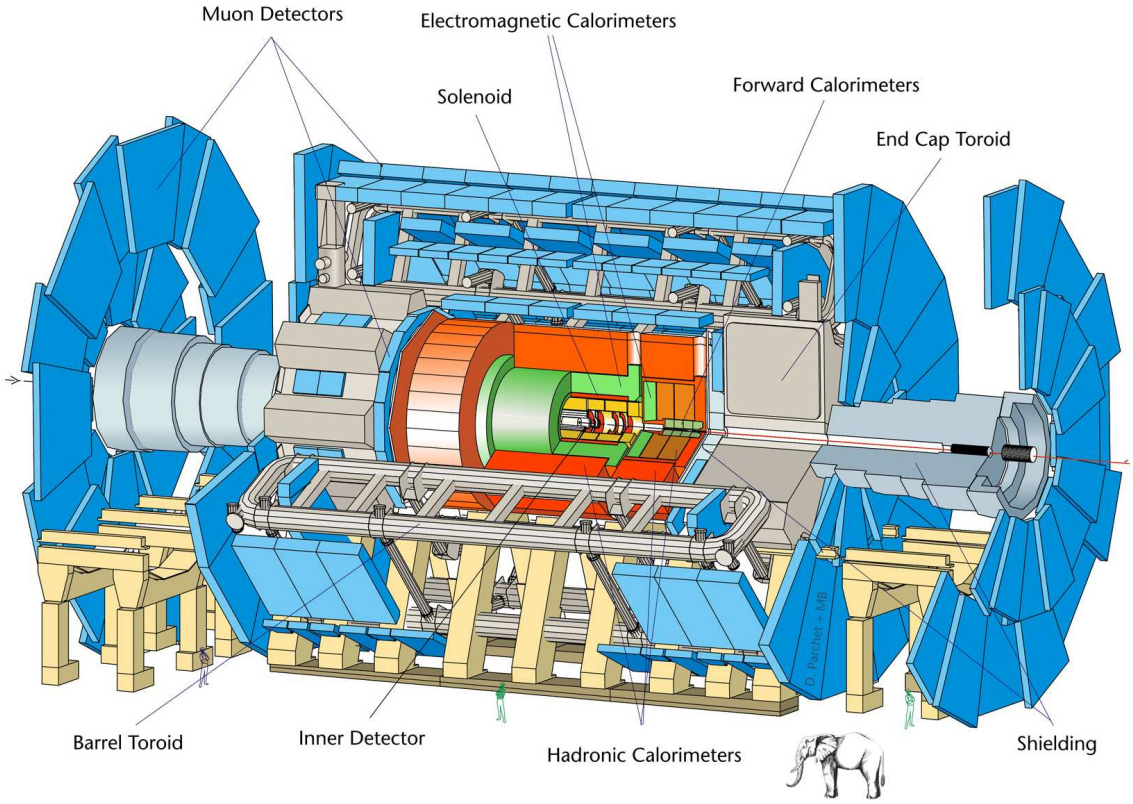

Figure 1.2: Drawing of the ATLAS Detector. The concentric array of detectors is shown to scale, with an African elephant for comparison. The subsystems and their functions are described in Section 1.1.4. The overall dimensions of the ATLAS detector are length of 44 m, diameter of 22 m, and weight of 7000 tons. ATLAS contains approximately 150 million active detector elements, see Table 1.3. The experiment will be contained within an underground cavern 100 m beneath the surface. During LHC operation, radioactivity within the ATLAS cavern will be too high to permit human entry. This high number of components within a hostile environment requires each system to be extremely robust: the Silicon Central Tracker is designed to be immune to single point failures.

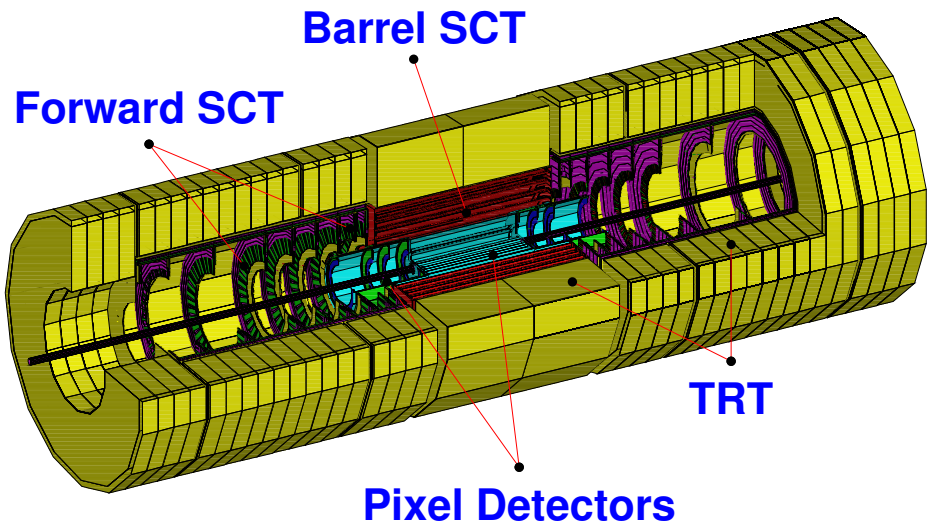

Figure 1.3: ATLAS Inner Detector, with Pixels, SCT (Silicon strips) and the TRT (straw tube wire chambers), see Section 1.1.4.1

1.1.4.1 Inner Detector

The Inner Detector (ID) provides tracking, momentum measurement, and precision vertex reconstruction. It uses a combination of high precision/high granularity discrete detectors and lower resolution continuous tracking for robust track reconstruction.

The ID is placed within the Central Solenoid (CS). This superconducting solenoid is 5.3 m long and 1.22 m in (internal) radius, generating a nominal magnetic field of 2.0 T which permeates the ID, aiding momentum and electric charge determination. The field return of the solenoid is through the steel tiles of the hadron calorimeter.

The ID itself is 7 m long with an outer radius of 1.15 m. It is divided into three sections: 2 forward and 1 (central) barrel section extending ±80 cm. In the forward section the detector elements are placed radially on disks. The barrel section has detector elements arranged in concentric cylindrical layers about the beam line, see fig. 1.3

The ID subsystems are identified by the type of active detector used: Pixels in the innermost layers (barrel radius from 4 cm to 13 cm), Silicon strips in the Semi-Conductor Tracker (SCT) (barrel radius from 30 cm to 52 cm) and Straw Tubes (single-wire chambers) of the Transition Radiation Tracker (TRT) (barrel radius from 56 cm to 107 cm). The Pixels and SCT from an inner section of discrete measurements while the outer section TRT gives continuous lower precision measurements. A typical track might pass through 3 pixels, 8 SCT elements and 36 TRT tubes.

The ID provides full tracking for $|\eta| < 2.5$ including impact parameter and vertexing for τ and heavy-flavour tagging.

Pixels The Pixel layers give high precision, high granularity measurements which contribute towards impact parameter resolution and tagging of τ and heavy flavours. Each pixel is 50μm in $R\phi$ and 300μm in z and arranged, in the barrel section, in three layers at radii of 4 cm, 10 cm and 13 cm. In the forward section there are five pixel disks from 11cm to 20cm from the centre. The total number of pixel channels is 140 million. The fine granularity of the pixels require the readout chips to be bump-bonded onto the detector substrate.

While secondary vertex reconstruction is enhanced by placing the innermost pixel layer at 4cm from the beam pipe, the severe radiation flux will damage the detectors, requiring replacement during the experiment's lifetime.

Semi-Conductor Tracker (SCT) The SCT detector elements are silicon microstrips at a pitch of 80 μm. Wafers containing these strips are used to form SCT modules, in which two wafer layers are glued back-to-back, with the strips offset by a stereo angle of 40mrads. This allows precision measurements in $R\phi$ and Z, while minimising "ghost" hits. In the barrel section, the modules are mounted on 4 barrels at radii of 30.0 cm, 37.3 cm, 44.7 cm and 52.0 cm, In the forward section, the modules are placed on nine wheels. There are 62 million SCT channels in total, and the detector system is designed to be immune to single-point failures.

The SCT contributes to tracking, vertex reconstruction, and impact parameter measurements with its high granularity.

Transition Radiation Tracker (TRT) The TRT is an array of straw tubes which provide continuous tracking hits. Radiation foils are added between straw tubes to generate transition radiation, which is detected by ionising the Xenon gas. This adds discrimination between electrons and hadrons to the ID. Each straw is a single (anode) wire chamber within a cylindrical cathode sleeve, 4mm in diameter and 144cm long in the barrel with a 30μm gold plated W-Re wire. In the barrel section there are 50,000 straws while the forward section has 320,000. Each channel provides a drift time measurement, giving a spatial resolution of 170μm, and two thresholds which are used to discriminate between tracking hits (low threshold) and transition radiation (high threshold) hits.

1.1.4.2 Electromagnetic Calorimeter (E-Cal)

The E-Cal provides energy measurements for non-hadronic particles, principally electrons and photons, by means of causing particle showers and sampling these showers within a medium. It is structured in three sections: one barrel $|\eta| < 1.475$ and two end-cap sections $1.375 < |\eta| < 3.2$. The end-caps are further divided into two coaxial wheels: the outer covering $1.375 < |\eta| < 2.5$ while the inner covers $2.5 < |\eta|3.2$.

The EM-Cal uses liquid Argon (LAr) as a sampling medium, with lead absorber plates to produce the showers and kapton bonded electrodes to collect the signal. An "accordion" geometry is used for the plates and electrodes which gives complete ϕ symmetry with no azimuthal cracks. The thickness of the lead plates has been optimised as a function of rapidity to maintain the required energy resolution. In the end-cap section the amplitude of the "accordion waves" increases with radius. The total thickness is > 24 radiation lengths in the barrel and > 26 in the end cap.

The outer end-cap and barrel ($|\eta| < 2.5$) are given over to precision physics; so have a finer segmentation compared to the inner end-cap. The barrel section has a pre-sampler to compensate for energy loss through the Inner Detector and cryostat.

The overall expected energy resolution is $10\%/\sqrt{E} \oplus 0.7\%$. This will be important for many physics processes. For example the $H \rightarrow \gamma\gamma$ channel will rely heavily on the E-Cal to separate the signal from the background prompt photon continuum.

1.1.4.3 Hadronic Calorimeters (H-Cal)

The H-Cal provides measurements of jet momentum and is used to calculate missing transverse energy. It has total acceptance of $|\eta| < 4.9$ and is physically divided into three sections: two end cap and one central barrel region. The calorimeter uses different technologies to suit the varying radiation environment, as the scintillator used in the central region cannot withstand the radiation in the end-cap regions. The three H-Cal sub-units are:

- Tile Calorimeter ($|\eta| < 1.7$)
- Hadronic End Cap (HEC) ($1.5 < |\eta| < 3.2$)
- Forward Calorimeter (FCAL) ($3.1 < |\eta| < 4.9$)

The Tile Calorimeter is sampling calorimeter using steel absorber plates and 3mm thick plastic scintillating tiles. They are arranged radially, staggered in depth, and periodic along z. Two sides of the scintillating tiles are readout by optical fibres[2] into separate photo-multiplier tubes.

Both the HEC and FCAL used liquid Argon for the harsher radiation at high $|\eta|$. The HEC uses copper plates with three parallel electrodes between each one. The two outer electrodes are set to 2 kV while the central rod is read out the charge.

The FCAL is split into three sections: the first one made from copper and two made from tungsten. All the sections have a metal matrix with concentric tubes and rods in the longitudinal gaps. The tubes and matrix are grounded while the rod is at high potential. The LAr filling the tube forms the active material.

The overall resolution of the H-CAL is $50\%/\sqrt{E} \oplus 3\%$. for $|\eta| < 3.0$. An accurate measure of jet momenta and missing energy is important for physics channels such as SUSY which has a signature of multiple jets and missing energy.

1.1.4.4 Muon Spectrometer

The Muon Spectrometer provides independent precision measurements of the muon momentum by placing detector "stations" around and within the large, air-filled volumes permeated by magnetic fields created by superconducting toroids. The air minimises multiple scattering, while the large separation between stations provides large lever arms for measuring track momenta.

The muon magnet system is the most distinctive feature of ATLAS, and it determines the overall dimensions of the experiment. It is composed of three large air-core toroids, one Barrel Toroid (BT) and two End-Cap toroids (ECT). The overall dimensions of the BT system are 26 m in length and 20 m in diameter, with the smaller end-cap toroids inserted in the BT and aligned with the CS. The magnetic field is therefore generated in three regions: barrel ($|\eta| < 1.0$), end-caps ($1.4 < |\eta| < 2.7$), and transition region ($1.0 < |\eta| < 1.4$) with a mixture of barrel and end-cap fields. Each toroid is made of eight coils, which for the BT are each contained in individual cryostats. For each ECT, all eight coils are contained within one cryostat. The peak field for the BT is 3.9 T, and 4.1 T for the ECT. The stored energy within the muon toroidal system exceeds 1200 MJ.

[2]These fibres are wavelength-shifting in order to ensure that photons entering from the side of the fibre, which lies along the scintillator face, propagate down its length.

The muon spectrometer, like the toroidal magnets, is divided into three sections: two end-caps and a central barrel. Detector "stations" are placed at the beginning, in the middle, and at the end of the field-permeated volume in the barrel, and on either side of the field in the end-caps. Each station is comprised of triggering and precision detector types. As in other ATLAS detector systems, the detectors are placed along the beam direction in the barrel and in wheels perpendicular to the beam in the end-caps.

The detector technologies used in the barrel (up to $|\eta| = 2.0$) are Monitored Drift Tubes (MDT) for precision and Resistive Plate Chambers (RPC) for triggering. In the end-caps ($2.0 < |\eta| < 2.7$), Cathode Strip Chambers (CSC) make precision measurements, while Thin Gap Chambers (TGC) act as triggers. Each of these is described briefly below.

MDT are aluminium tubes with a diameter of 30mm and a central 50μm W-Re wire. When the gas pressure is at 3 bar the MDTs have single wire precision of 80μm. Several MDT layers, and shape-sensitive readout electronics contribute to the overall accuracy of the MDTs.

The **RPC**s are a pair of parallel insulating plates, with a narrow gap filled by a specific gas mixture. They are read out by a pair of orthogonal pick-up strips placed on either side of the detector. Primary ionisation electrons are multiplied into avalanches by high (4.5kV/mm) electric fields. Each one has a space-time resolution of 1cm×1ns. This fast timing makes the RPCs convenient for triggering

CSCs are multi-wire proportional chambers in which the anode-cathode spacing equals the anode wire pitch, and in which the cathodes are segmented strips running orthogonally to the anode wires. The anode wires cause ionised electron avalanches, of which capacitive readout of the cathodes yields a resolution of 60μm. The position is determined by the centre of gravity of induced charge.

TGCs are multi-wire proportional chambers where the anode wire pitch is greater than the anode-cathode distance, and the overall chamber thickness is 3 mm. These small distances, as well as the 3kV potential drop, account for the TGC's fast response time. The anode wires are laid out parallel to the MDTs and readout-strips are placed orthogonally. The signal along the anode wires provides the trigger information and gives a measure of the second-coordinate of the muon trajectory.

The expected resolution of the muon spectrometer requires knowledge of detector element locations to an accuracy of 30μm for a given projective tower. An optical alignment

system will be used to monitor detector positions and to correct for any displacements, since it will not be feasible to ensure mechanical stability to such a small scale.

1.1.4.5 Trigger and DAQ

The ATLAS trigger and data-acquisition (DAQ) system is based on three levels of online event selection. Each of these triggers further discriminates the data which has passed the previous level. The aim is to remove minimum bias events while keeping interesting physics and processing large data rates.

Over the three trigger levels, the event rate is reduced from the bunch crossing rate of 40MHz to 100Hz to allow data to be written for permanent storage. This corresponds to reducing the data rate from 40 TB/s to 100 MB/s for 1MB/event.

Level 1 Level 1 makes the initial cut, reducing an input rate of 40MHz to an output of 75 KHz (upgradable to 100KHz). It uses reduced granularity data from calorimetry and the muon spectrometer to search for high p_{T} leptons or photons, jets, $\tau \rightarrow$ hadrons decay and missing energy. The trigger is implemented on custom hardware processors which store data in "pipeline" memory buffers mounted on or close to the detector. The pipeline length is 2.5μs: 2.0μs trigger latency (equivalent to 80 bunch crossings) plus 500ns buffer.

Level 2 Level 2 uses Regions-of-Interest (RoI) information identified by level 1, with access to full granularity and precision, including tracking data, where necessary. The data rate may be reduced to ≈ 1 KHz with a variable trigger latency of $\approx 1 - 10$.

Level 3 Level 3 is the Event Filter (EF) and the last stage of online selection. It uses offline method and algorithms adapted for online analysis and will include latest alignment and calibration data. The output data is ≈ 100 Hz or 100 MB/s for a full event.

1.1.4.6 This Thesis within the ATLAS Collaboration

This thesis presents two research projects carried out within the context of the development and assembly of the ATLAS experiment. The first is a hardware project, which developed an interferometer to measure vibrational properties of ATLAS SCT support

structure prototypes before SCT assembly. The second project, a Monte Carlo analysis, estimates the statistical limitation on the accuracy with which ATLAS will be able to measure the width of the W boson, a well known SM parameter.

1.2 Electronic Speckle Pattern Interferometry and ATLAS SCT Alignment

Electronic Speckle Pattern Interferometry (ESPI) is a technique that creates images of objects covered in fringes which represent contours of either deformation or vibration amplitude, depending on how the technique is used. ESPI has been used during SCT design and prototyping to measure thermal distortions of silicon strip modules and modes of vibration of support structures.

SCT performance requires the location of the modules to be known to a precision of of $\approx 10\,\mu$m. This alignment requirement is challenging for several reasons. SCT components are designed to a tight mass budget to minimise scattering within the tracker. This leads to low mass / high rigidity structures, which may vulnerable to relatively high amplitude oscillations, as the elastic moduli will tend to be high and the mass density low.

While running, the SCT modules will dissipate 15 - 30 kW of power in total, leading to severe local temperature gradients, which may cause local deformations. A cooling system will run by all the modules to remove the excess heat. Stresses and vibrations transfered through the cooling system may also cause displacement, deformation, and vibration.

The relatively large scale of the SCT increases its vulnerability; the SCT volume is a factor of ten larger than the largest current silicon system used in particle physics - that of CDF II [CDF].

Tracks will be used for alignment, as has been done in previous high-energy physics experiments. In addition, the SCT will have a grid of Frequency Scanning Interferometers (FSI) measuring absolute distances between fiducial points on the barrels and wheels to determine overall deformation of the system.

Chapter 2 of this thesis explains the theory behind both static ESPI, and vibrational ESPI (VESPI), along with some comments on visibility and noise in this type of interferometry.

Chapter 3 describes a previously assembled interferometer, as well as some thermal deformation measurements of prototype SCT modules made with it.

Chapter 4 describes the development of an interferometer which uses the VESPI theory from Chapter 2 to successfully measure vibration contour maps. The chapter then presents vibrational mode measurements made with this new interferometer on prototype Carbon Fibre SCT support structures. The resulting measurements are compared with Finite Element simulations, which are found to match, validating the use of the modelling approach on sophisticated, non-isotropic materials. VESPI was also used to find the quality factor of the prototypes studied, showing the Carbon Fibre structures have a high Q, which means that excitement of a normal mode during ATLAS operation could lead to motion over a period of $\mathcal{O}\,1$ second. This could have negative impact on both tracker alignment directly and on FSI's capacity to measure absolute length.

1.3 Monte Carlo Studies of LHC Electroweak Physics and Γ_W

The LHC will produce a very large number of gauge bosons and gauge boson pairs. This combined with the high centre-of-mass energy will allow precision EW measurements to be performed. Preliminary studies have previously been carried out within the ATLAS collaboration on the precision which can be expected when measuring W mass and on Triple Gauge Couplings (TGC) [ATLb]. Chapters 5 and 6 describe an additional precision measurement which can be carried out by ATLAS, that of the width of the W, Γ_W.

The W mass study referred to above anticipates that the existing W mass accuracy before the LHC will be near 30 MeV, while a precision of 15 MeV is necessary for it not to be the dominant source of error on the estimation of the Higgs mass, provided the expected accuracy of top mass, 2 Gev, is attained. The study concludes that, if systematic sources of error are understood and corrected for, ATLAS may achieve a resolution of about 20 MeV. By combining this measurement with that obtained by CMS, (the other LHC general-purpose detector), the total error should be close to 15 MeV.

Triple Gauge Couplings are of interest because the couplings between the EW gauge bosons have very specific SM predictions, making measurements of such couplings a powerful test of the SM. Any deviations would make for interesting physics in themselves, and provide complementary information to that obtained by direct searches for new

physics.

Furthermore, measurements involving inclusive gauge boson production / gauge-boson pair production / gauge boson plus jet(s) production will all be important to understand events and backgrounds for new particle searches, and Z boson production will play a key role in the calibration of the detector mass scale.

Within this context, chapters 5 and 6 present a study of the Γ_W statistical resolution expected from ATLAS:

Chapter 5 describes the transverse mass (M_T) technique, with which Γ_W has been measured successfully by the CDF (Collider Detector Facility) collaboration at the Tevatron accelerator in Fermilab. It explains how this technique could be used at ATLAS, and what steps will be followed to simulate the measurement.

Chapter 6 reports on a set of Monte Carlo simulations and subsequent analyses carried out following the methodology set out in Chapter 5 in order to estimate the statistical accuracy ATLAS can expect to achieve on a Γ_W measurement. Possible systematic sources of error are also briefly discussed.

Chapter 2

Speckle Interferometry

This chapter describes Electronic Speckle Pattern Interferometry (ESPI). A speckle pattern, which is created by imaging an optically rough surface illuminated by coherent light, is combined with another wave front (plane or speckled) to create an interferogram which is then stored. If the imaged surface undergoes a small distortion, displacement, or vibration, and another interferogram is taken, the correlation between speckle intensities in the two interferograms yields information about the change in shape undergone by the imaged surface.

2.1 Principles of Speckle Interferometry

2.1.1 Diffraction and the Subjective Speckle Pattern

It can be readily observed that an optically rough surface illuminated by coherent light presents a speckled appearance (fig 2.1).

Optically rough means that the surface is non-specular. More precisely, it means that variations in surface height are present within a distance smaller than the resolution limit of the imaging system[1]. Because the speckled image depends on the optical system used, it is called the *subjective speckle pattern.*

To undestand the origin of the subjective speckle, it is recalled that the diffraction due to an aperture can be derived from the Huygens principle, which states that every point on a wave front is itself a source of spherical waves. For the case of a circular

[1]These height variations must be of a magnitude comparable to (or greater than) the wavelength of the illuminating light.

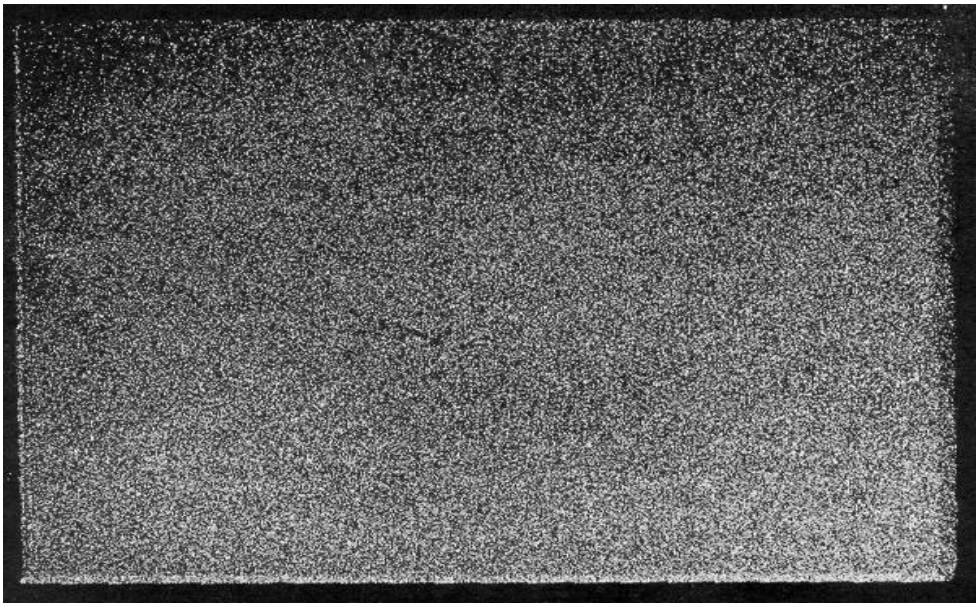

Figure 2.1: A speckled image of an optically rough rectangular plate.

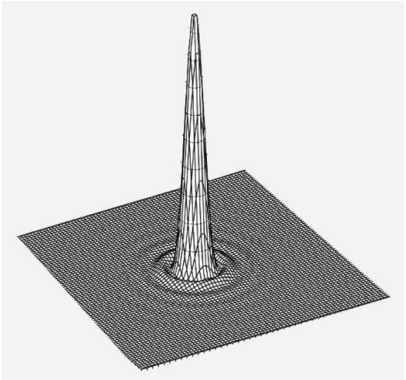

Figure 2.2: Airy pattern.

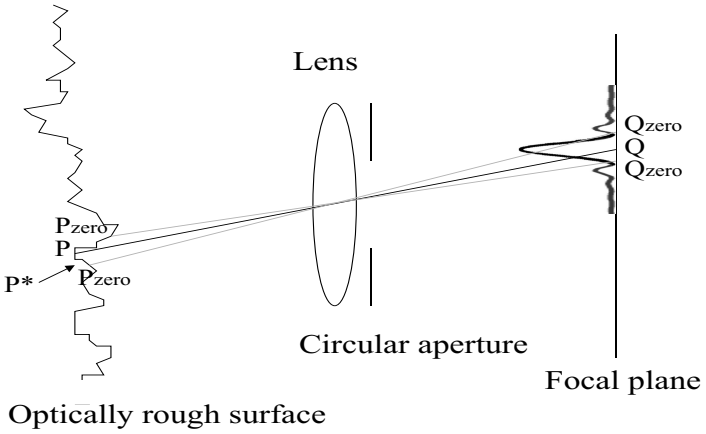

Figure 2.3: Speckle formation.

aperture, it can be shown [Hec87] that the diffraction pattern is described by the Airy formula:

$$I(\theta) = I(0) \cdot \left[\frac{2J_1(\frac{2\pi}{\lambda} a \sin\theta)}{\frac{2\pi}{\lambda} a \sin\theta} \right]^2 \tag{2.1}$$

where I is the intensity of light on the image plane, J_1 is the first order Bessel function of the first kind, θ is the angular deviation from the normal of the incident wave front, λ is the wavelength of the light, and a is the radius of the aperture (fig 2.2).

Consider an optically rough surface illuminated by coherent light and focused onto a screen by a lens of a circular aperture. Light reflected off a point P is focused by the lens onto point Q and diffracted by the aperture into an Airy pattern around Q. Since most of the light is received inside the first zero of the Airy function, it can be said that the area of influence of a given surface point on the image plane is the circle formed by this first zero (Q_{zero} in fig 2.2).

Light reflected from another point P^* will fall within the area of influence of point P if P^* is very close to P. Given that the surface height is uneven, the optical path between P and Q is different to that between P^* and Q. This path difference is random since it depends on the exact surface profile and the ratio of the surface height variation to the wavelength of the illuminating light.

Thus the intensity at any image point Q is determined by the interference of the light reflected from all the points within the corresponding P_{zero} circle on the object surface, where P_{zero} is the locus of points closest to P whose reflected light does not contribute to the intensity at Q. For some points on the image this interference will be bright, for others less so. The resulting image will be a random distribution of dark and bright spots: the speckle pattern.

2.1.2 Interferograms and Contour Maps

Let $A_o(x,y)e^{i(\omega t - \phi_o(x,y))}$ be the electric field of arbitrary amplitude, A_0, and phase, ϕ_0, at a given point (x,y) on a screen illuminated by the speckled image of an object[2]. To use this speckle pattern for interferometry, it is necessary to combine it with another wave

[2]Strictly speaking the electric field is described by the real part of this expression, i.e. Electric Field $= \Re[A_o(x,y)e^{i(\omega t - \phi_o(x,y))}]$. However, it is customary to omit the $\Re$ symbol in the calculations and only take the real part explicitly when describing a quantity that is measured. This does not alter the results and reduces clutter within the equations.

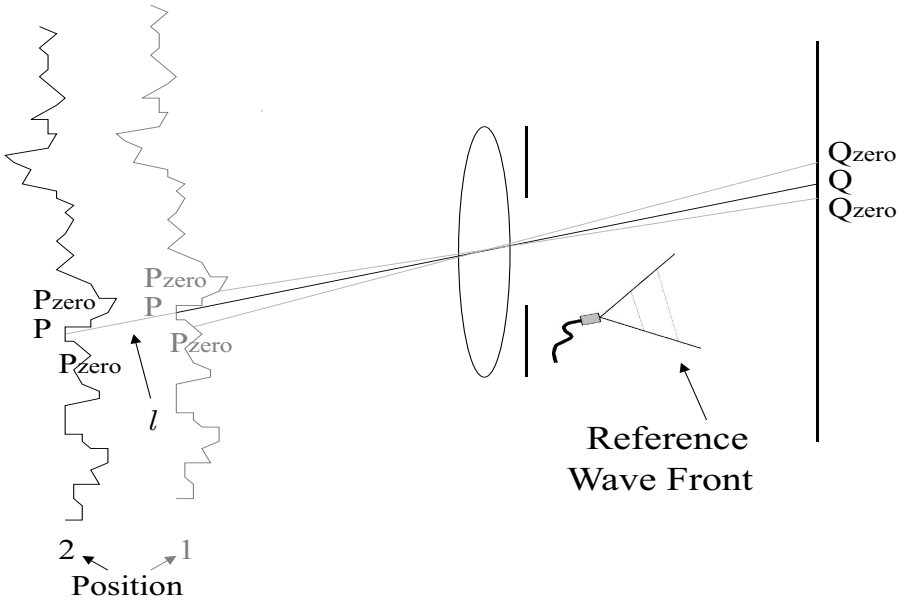

Figure 2.4: Surface displacement from position 1 to position 2.

front on the image plane[3]. Let this second (reference) wave front be $A_r(x,y)e^{i(\omega t-\phi_r(x,y))}$. The sum of these two will be a new speckle pattern with the same general appearance as before but a different distribution of bright and dark speckles. Such an image will be called an *interferogram* for the remainder of this thesis, in order to draw a distinction between speckle patterns generated by the combination of two wave fronts and those caused solely by the object wave front("raw" speckle patterns).

To find the intensity of the electromagnetic field at a point, the amplitude of the field must be multiplied by its complex conjugate. Thus the intensity of the interferogram is described by the modulus squared of the sum of the amplitudes:

$$I(x,y) = A_o^2(x,y) + A_r^2(x,y) + 2A_o(x,y)A_r(x,y) \cdot \cos\left(\phi_o(x,y) - \phi_r(x,y)\right). \qquad (2.2)$$

where I is the intensity of the interferogram.

Now consider what happens if the object surface is slightly displaced as shown in figure 2.4. The difference in optical path between the Q_{zero} region and the P_{zero} region in positions 1 and 2, Δl, introduces a phase difference $\Delta\phi = \frac{2\pi}{\lambda}\Delta l$ between the speckle patterns formed. This leads to a change in intensity in the interferogram. Given that the

[3]Only the combination of two or more wave fronts can produce a phase-dependent intensity. See eq. 2.2.

reference wave front remains unchanged and the P_{zero} region still maps onto the Q_{zero} region, the only term that differs in eq. 2.2 between interferograms 1 and 2 is ϕ_o.

If the interferograms are subtracted and the modulus taken, the result is

$$I_{|sub|} = |I_1 - I_2| = 2A_oA_r \cdot [\cos(\phi_{o_1} - \phi_r) - \cos(\phi_{o_2} - \phi_r)] \tag{2.3}$$

which can be rewritten as

$$I_{|sub|} = 2A_oA_r \cdot |\cos\Phi - \cos(\Phi + \Delta\phi)| \tag{2.4}$$

where $\Phi = \phi_{o_1} - \phi_r$ and $\Delta\phi = \phi_{o_2} - \phi_{o_1} = \frac{2\pi}{\lambda}\Delta l$ as before.

This quantity, $I_{|sub|}$, is useful because it is zero in the regions where $\Delta\phi = 0, 2\pi, 4\pi, ...$ Furthermore, because $I_{|sub|} \not< 0$, it can be displayed as an intensity[4]. Thus when $I_{|sub|}$ is viewed simultaneously for all x and y, forming a two dimensional image of the surface, the image is dark in the regions that have undergone displacements of an integer number of wavelengths. If it is assumed that the surface is continuous, contours of equal intensity on the image can be interpreted as contours of equal displacement on the object. This is best seen with an example. Figure 2.5 shows a subtraction of two interferograms. Theses interferograms were taken under nearly identical conditions. The object imaged is a circular silicon wafer, the edge of which is glued to the end of a cylindrical chamber of the same diameter, forming a lid. Between the two interferograms, the pressure inside the cylindrical chamber was increased, which caused the wafer to bulge. The subtraction has a clear structure of concentric intensity contours which, by the above reasoning, represent contours of equal surface displacement due to the pressure change.

Such patterns are present in all interferogram subtractions where surface continuity ensures contour continuity, provided that the interferograms are correlated, (see below). These patterns of bright and dark lines across an image are often called 'fringes' because they are reminiscent of interference fringes, Moire fringes, and many other interference phenomena. Strictly speaking, however, the are contours of equal intensity correlation which are interpreted as contours of equal object displacement, deformation, or vibration amplitude, depending on the nature of the experiment.

[4]Intensities must be ≥ 0 because they are defined as $A \times$ Complex Conjugate. Subtracting two intensities can yield values < 0, which cannot be displayed as an intensity

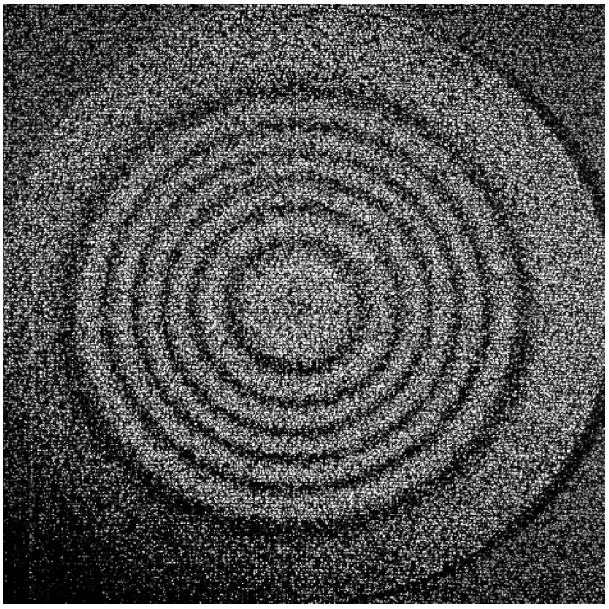

Figure 2.5: ESPI image of a circular Si diaphragm bulging due to a pressure change of 50 Pa. Each fringe represents a displacement of $\sim$260 nm, implying a total peak movement of about 1.6 microns.

2.1.3 Speckle Size and Correlation

In order to understand when two interferograms can be combined in such a way that they yield useful information, it is necessary to study the speckle itself in more detail.

A thorough statistical analysis of speckle size and speckle pattern autocorrelation has been carried out by Goodman [Goo75]. The autocorrelation function is defined as as the average of the product of a speckle pattern at a certain position times itself at a different position, $R(x_1, y_1; x_2, y_2) = \langle I(x_1, y_1)I(x_2, y_2)\rangle$, where the average is carried out over an ensemble of macroscopically similar but microscopically different surfaces (for example different samples of the same material). It is useful because it is a measure of how similar the patterns are on a microscopic scale: large R implies a high degree of similarity between them, small R implies little similarity [5].

[5]To see why this is true, consider two special cases. First consider the average over the product of two different speckle distributions that have no intensity correlation whatsoever. Some of the terms contributing to the average will be due to bright speckles of the first pattern being multiplied by bright speckles of the second, some terms will be due to bright speckles on dark ones, some will be due to dark on dark. Bright multiplied by bright will be relatively rare. The average will be composed of many small contributions and very few larger ones. Now consider the average over the product when $I_1 \equiv I_2, x_1 \equiv y_1$, and $x_2 \equiv y_2$. Here the speckle patterns are completely correlated. In this case bright will always be multiplied by bright, and dark speckles from one pattern will not reduce the contributions

Goodman's study begins with the assumption that the electric field at a given point on a plane at a distance z from a rough, coherently illuminated surface is a circular complex Gaussian random variable (a reasonable assumption given the Central Limit Theorem). A summary of this work is presented in Appendix A. Only the result is stated here, namely that the autocorrelation function for the speckle pattern imaged through a circular aperture on the image plane (x, y) is given by

$$R(r) = \langle I \rangle^2 \left[1 + \left| \frac{2J_1\left[\frac{\pi D r}{\lambda v}\right]}{\frac{\pi D r}{\lambda v}} \right|^2 \right]. \tag{2.5}$$

where r is $(\Delta x^2 + \Delta y^2)^{1/2}$ on the image plane, D is the diameter of the aperture, v is the distance between the lens and the image plane, and J_1 is the first order Bessel function as in equation 2.1.

It should be noted that the second term in $R(r)$ is an Airy function (see eqn 2.1 and fig 2.2). Thus it is sensible to define speckle diameter to be twice the value of r for which the autocorrelation function reaches its first minimum.

$$d_{speckle} = 2.44 \frac{\lambda v}{D} \tag{2.6}$$

where as before, λ is the wavelength of light, v is the distance between the lens and the image plane, and D is the diameter of the circular aperture.

It should be noted that the simple assumption that the speckle diameter is the diameter of the area of influence of point P on the image plane (i.e. twice the distance between Q and Q_{zero} in fig 2.3) yields the same result. Thus the naïve estimate and the complete analysis agree on equation 2.6.

When two interferograms are subtracted to obtain a fringe pattern, the fringes appear because of *intensity* correlations between the interferograms: speckles that have the same intensity in both interferometers appear dark in the subtraction. However, if the speckle pattern moves across the image plane by more more than $d_{speckle}$, *spatial* correlation is lost, so comparing the intensities on point (x, y) on the image plane before and after some change is no longer equivalent to comparing the speckle caused by the P_{zero} circle before and after it.

of the bright speckles from the other. The average will be maximal, and so greater than for the previous case.

Therefore spatial correlation must be preserved so that intensity correlation can be measured. Loss of spatial correlation between interferograms is one of the major obstacles to ESPI measurements. It is necessary to ensure that the object under study, the illumination sources, and the imaging device do not move enough to cause the speckle pattern to be displaced by more than a speckle diameter between interferogram captures.

2.1.4 Interferometer Geometries

It has been seen that subtracting two interferograms yields contours of equal displacement. The exact dependence of phase on displacement magnitude and direction depends on the geometry of the interferometer setup. Two geometries have been used for measurements in Oxford: out-of-plane sensitive and in-plane sensitive.

2.1.4.1 Out-of-plane sensitive

In the out-of-plane arrangement, the object is illuminated with one wave front and imaged on an image plane which is illuminated by another (reference) wave front. This reference wave front maintains the same phase throughout the measurement[6]. The phase difference introduced is thus due only to the change in path length from the source S to the lens L via a point P on the object when P moves from position 1 to position 2. (See figure 2.6.)

The important quantity is $\Delta l = [SP_2 + P_2L] - [SP_1 + P_1L]$. In the figure the displacement $D = P_1P_2$ is greatly exaggerated, in reality $D \ll SP_1, SP_2$ so SP_1 is practically parallel to SP_2. Thus

$$SP_2 - SP_1 = \frac{\vec{D} \cdot \vec{SP_1}}{SP_1} \quad \text{and} \quad P_2L - P_1L = \frac{\vec{D} \cdot \vec{P_1L}}{P_1L} \quad . \tag{2.7}$$

For small angles from the normal of the viewing plane (dotted lines in fig. 2.6), the projection of D along this angle is very similar to the z component of D. Thus for small illumination and viewing angles, θ and $\psi \leq 15°$, it is true to a very good approximation that

$$\Delta l = (\cos\theta + \cos\psi) D_z \implies \Delta\phi = 2\pi(\cos\theta + \cos\psi) D_z/\lambda. \tag{2.8}$$

[6]In specific cases the phase of the reference wave front may be intentionally varied. See sections 2.2.3 and 4.2.2.

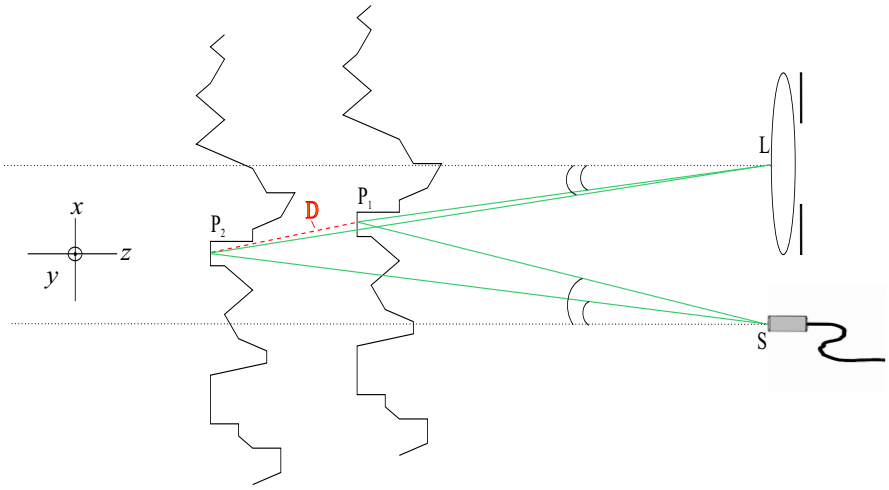

Figure 2.6: Out-of-plane sensitive intererferometer.

This interferometer geometry is sensitive to motion that is normal to the viewing plane, i.e. it is out-of-plane sensitive, with adjacent contours of equal intensity corresponding to a relative displacement slightly smaller than $\lambda/2$. Using a wavelength of $514\,nm$, as is the case in figure 2.5 and in all other interferograms shown in this thesis, yields out-of-plane sensitivity of $\sim 260\,nm$ per contour of equal intensity.

Motion of the object along the plane will not register a signal in this geometry, but will lead to spatial decorrelation which will reduce or destroy contour visibility.

2.1.4.2 In-plane sensitive

In-plane sensitivity is achieved by illuminating the object from two sources, S and S', at different angles (fig 2.7). In this case there is no 'reference' beam and the phase of both beams is altered by object movement. Thus the change of path differences between the two wave fronts before and after a displacement, Δl, can be seen from figure 2.7 to be $\Delta l = [S'P_2 - SP_2] - [S'P_1 - SP_1]$, which can also be written $\Delta l = \Delta S' - \Delta S$ where $\Delta S' = [S'P_2 - S'P_1]$ and $\Delta S = [SP_2 - SP_1]$. Note that Δl does not depend on changes in PL.

Now consider $\Delta S'$ in detail. It is convenient to draw an auxiliary line from $S'P_2$ to P_1 (or from $S'P_1$ to P_2 if $S'P_1 > S'P_2$) to complete a right angle triangle with $\Delta S'$ and the displacement D. Since $S'P_1$ and $S'P_2$ are practically parallel,

$$\Delta S' = D \sin \alpha. \tag{2.9}$$

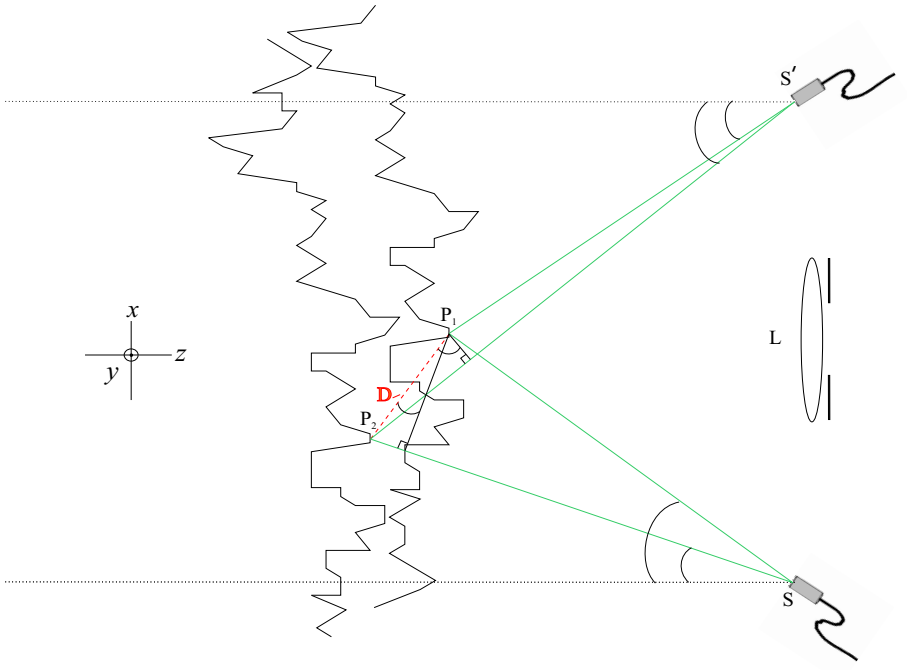

Figure 2.7: In-plane sensitive interferometer.

Here α is the angle between the auxiliary line and D. It can be seen from the diagram (fig. 2.7) that $\alpha = \delta + \psi$ where δ is the angle between D and the xy plane. Thus

$$\Delta S' = D \sin(\delta + \psi) = D[\sin\delta\cos\psi + \sin\psi\cos\delta] \tag{2.10}$$

and, using a similar construction for ΔS,

$$\Delta S = D \sin(\delta + \theta) = D[\sin\delta\cos\theta + \sin\theta\cos\delta]. \tag{2.11}$$

The change in path differences can now be expressed as

$$\Delta l = \Delta S' - \Delta S = D[\sin\delta(\cos\psi - \cos\theta) + \cos\delta(\sin\psi - \sin\theta)] \tag{2.12}$$

which, if $\theta = -\psi$ simplifies to

$$\Delta l = D\cos\delta \cdot 2\sin\theta. \tag{2.13}$$

Because δ is the angle between the displacement and the xy plane, this interferometer is

only sensitive to the component of the motion along it. In a typical arrangement with $\theta = 60^\circ$ and $\lambda = 514\,nm$, each contour represents a displacement of $297\,nm$ relative to an adjacent contour.

Surface motion perpendicular to the plane will not produce a signal, but will lead to speckle decorrelation, reducing or destroying contour visibility.

2.2 Principles of Time Averaged ESPI

2.2.1 Interferometry and Dynamic Measurements

The interferometry discussed in section 2.1 involves objects that remain static while the interferograms are created. The subtraction of two interferograms taken at different times yields intensity correlation fringes which contain information about the deformation undergone by the object between these times. However, interferometry can also be used on objects while they are in motion. The dynamical measuring techniques used to do this divide into two categories: multiple imaging and time averaging techniques.

Multiple imaging is essentially a sped-up version of the static measurements discussed in section 2.1. The idea is to take pictures of a very short duration so the object appears static in the interferograms. (Technically, the requirement is that the distance travelled by the object during the shutter or integration time of the imaging system be much smaller than a wavelength of light.) Several of these short-duration interferograms are taken as the object moves and are later subtracted from each other, two by two. The motion of the object between different times is then deduced from these subtractions exactly as if they were static measurements.

The advantages of this technique are the following. The interpretation of the images is identical to that of static displacements and thus no further understanding of speckle interferometry is required. Furthermore, it is applicable to all sorts of dynamic situations: oscillations, displacements, and even transients, provided that the interferograms can be taken closely enough in time. The main disadvantage of this technique is its power requirement since, for an imaging system of a given sensitivity, a shorter exposure time demands a corresponding increase in illumination. Thus the higher the frequency or the faster the transition under study, the more powerful a source of light is needed.

The second technique, time averaging, takes single interferograms of an exposure time long enough for the motion of the object to be apparent. The exact nature of the motion

is deduced from the effect it has on the image. Thus, if multiple imaging can somehow be compared to cinematography in that motion is reconstructed from a sequence of sharp images, time averaging is rather like studying motion by using a blurred photograph.

Time averaging comes into its own when studying regular oscillations, as will be explained in the following subsection. Long integration time gives more oscillation cycles to integrate over, and thus, in general, better results (see discussion immediately prior to equation 2.22). Time averaging is particularly useful for studying large areas, because the disadvantage incurred by spreading the available light over a large area is reduced by increasing the camera integration time.

2.2.2 Time Averaged Intensity

Consider a time-dependent intensity pattern, $I = I(x, y, t)$, such as the interferogram created by combining a stationary reference wave front with a wave front reflected off a moving object in an out-of-plane interferometer. When recording this with a camera, the resulting intensity registered by the photographic film or CCD is the value of the intensity at each point averaged over the camera integration (shutter) time. More precisely, if the camera generates the image over a time T, the registered interferogram is

$$I(x, y) = \frac{1}{T} \int_0^T I(x, y, t)\, dt. \tag{2.14}$$

Recall that the intensity in an out-of-plane interferometer has a $\cos(\phi_o(x, y) - \phi_r(x, y))$ term as stated in equation 2.2. Motion of the surface will lead to $\phi_o(x, y)$ being replaced by

$$\phi_o(x, y) \rightarrow \phi_o(x, y, t) = \phi_{const}(x, y) + \frac{2\pi}{\lambda} k a(x, y, t), \tag{2.15}$$

where $\phi_{const}(x, y)$ is a an arbitrary offset phase, k is a numerical factor which accounts for interferometer geometry (*e.g.* the effect of the illumination angle), and $a(x, y, t)$ is the displacement of the object from its initial position. For a surface in simple harmonic motion (figure 2.8),

$$a(x, y, t) = a_0(x, y) \sin(\omega t + \psi) \tag{2.16}$$

where a_0, ω, and ψ are, respectively, the amplitude, frequency and phase of the object vibration. Substituting equations 2.16 and 2.15 in equation 2.2 gives

$$I_{int}(x,y,t) = DC_{terms} + C \cdot \cos\left(\frac{2\pi}{\lambda} k a_0 \sin(\omega t + \psi) + \phi_{const}(x,y) - \phi_r(x,y)\right) \quad (2.17)$$

where $DC_{terms} = A_o^2(x,y) + A_r^2(x,y)$ and $C = 2A_o(x,y)A_r(x,y)$.

Once the form of $I(x,y,t)$ is known, equation 2.14 can be used to calculate the intensity registered by the camera:

$$I(x,y) = DC_{terms} + \frac{1}{T} C \int_0^T \cos\left(\frac{2\pi}{\lambda} k a_0 \sin(\omega t + \psi) + \phi_{const}(x,y) - \phi_r(x,y)\right) dt. \quad (2.18)$$

Setting $\frac{2\pi}{\lambda} k a_0 \sin(\omega t + \psi) = A$ and $\phi_{const}(x,y) - \phi_r(x,y) = \Phi$, this can be rewritten as[7]

$$I(x,y\ ) = DC_{terms} + \frac{1}{T} C \left[\cos\Phi(x,y) \int_0^T \cos\left(\frac{2\pi}{\lambda} k a_0 \sin(\omega t + \psi)\right) dt - \sin\Phi(x,y) \int_0^T \sin\left(\frac{2\pi}{\lambda} k a_0 \sin(\omega t + \psi)\right) dt\right]. \quad (2.19)$$

Equation 2.19 is the exact expression for the image of a sinusoidally oscillating object. However, it is of interest to consider the special case when the camera integration time is a multiple of the period of oscillation of the object; $T = n\tau$ where n is integer and $\tau = 2\pi/\omega$, the period of object oscillation. Since sine is an odd function the second integral of equation 2.19 vanishes, and since the zeroth order Bessel function of the first kind is [Arf85]

$$J_0[x] = \frac{1}{2\pi} \int_0^{2\pi} \cos(x \sin t)\, dt. \quad (2.20)$$

The intensity after time integration can be written as

$$I(x,y) = A_o^2 + A_r^2 + 2A_o A_r \cos\Phi(x,y) \cdot J_0\left[\frac{2\pi}{\lambda} k a_0(x,y)\right] \quad (2.21)$$

where the (x,y) dependence of A_o and A_r is omitted for brevity.

Equation 2.21 is applicable in all cases where $T \gg \tau$, even if the time of integration

[7] Use trigonometric identity $\cos(A \pm \Phi) = \cos A \cos\Phi \mp \sin A \sin\Phi$.

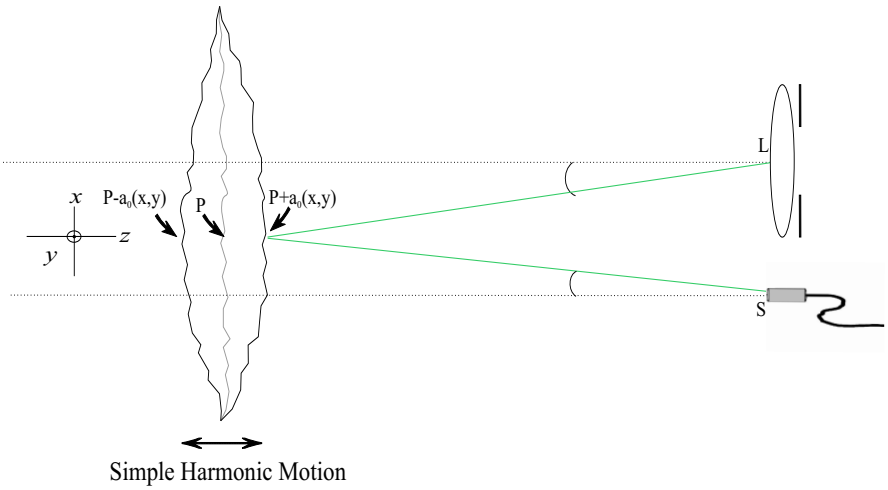

Figure 2.8: Out-of-plane sensitive interferometer imaging surface in simple harmonic motion. The surface is shown in the extremum positions $\pm a_0$ with the equilibrium position shown in grey This motion can be studied with time-averaging techniques, see equation 2.19.

is not a whole number of object oscillation periods. This is because, for periodic motion,

$$\int_0^T = \int_0^{(n+\varepsilon)\tau} = \int_0^{n\tau} + \int_0^{\varepsilon\tau} = n\int_0^{\tau} + \int_0^{\varepsilon\tau} \tag{2.22}$$

where $T = (n+\varepsilon)\tau$, n is an integer and $\varepsilon < 1$. Therefore, for $n \gtrsim 10$,

$$n\int_0^{\tau} + \int_0^{\varepsilon\tau} \approx n\int_0^{\tau} \text{ provided } \int_0^{\tau} \geq \int_0^{\varepsilon\tau} . \tag{2.23}$$

Equation 2.19 meets condition 2.23 because all functions inside the integrands are harmonic.

2.2.3 Fringes in Time Averaged ESPI

Equation 2.21 states that the intensity registered by a camera on which the interferogram of a vibrating surface is formed has a Bessel function dependence on the amplitude of the vibration, and it is applicable provided certain conditions are met (see subsection 2.2.2). These fringes form a contour map of the surface, with regions of equal values of J_0 appearing as contours of equal amplitude.

As it stands, the interferogram described by equation 2.21 cannot be interpreted

visually. The J_0 signal sits atop a $A_o^2 + A_r^2$ background which is generally greater than $2A_oA_r$ ($A_o \approx A_r$ is rarely true on a pixel-by-pixel basis). Whether the interference term adds to or subtracts from $A_o^2 + A_r^2$ depends not only on the Bessel function, but also on $\cos\Phi(x, y)$, which varies randomly from point to point because the surface is optically rough. The interferogram displays a uniform average intensity, with a wider than average spatial variation of intensities in the regions that have a high value of J_0. These fringes of intensity variation are not distinguishable by eye.

The J_0 signal is made visible by taking two time averaged interferograms in rapid succession, keeping all conditions identical except for Φ (in practice ϕ_r), and then subtracting them. As in equation 2.4, it is the modulus of the subtraction which is displayed. This allows cancellation of the background terms without destroying the signal. The result is

$$I_{|sub|}(x, y) = \left| \Big(\cos\Phi_1(x, y) - \cos\Phi_2(x, y) \Big) \cdot J_0 \left[\frac{2\pi}{\lambda} k a_0(x, y) \right] \right| . \tag{2.24}$$

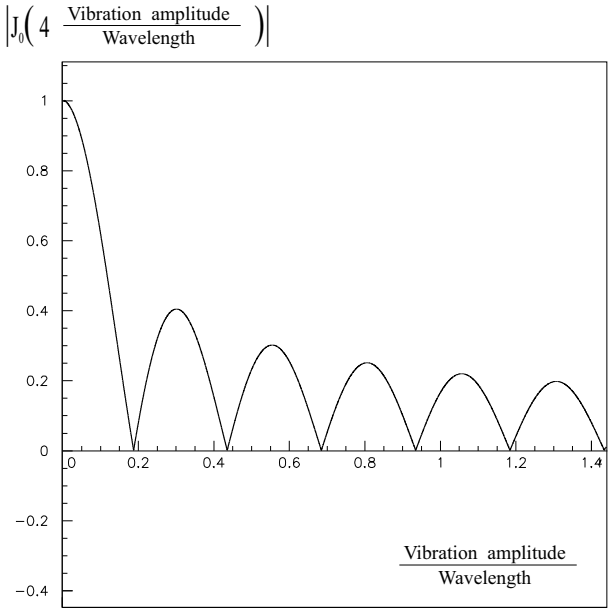

Figure 2.9: Modulus of the zeroth-order Bessel function, see equation 2.24. The abscissa is the ratio of the vibration amplitude a_0 to the wavelength of light λ. Thus, for a fixed wavelength, intensity contours show regions of equal amplitude.

The regions where the Bessel function is zero appear as dark contours in this subtraction, separated by brighter fringes. The $|J_0|$ dependence of brightness on the ratio

[h]

Zero No.	a_0/λ	a_0 for $\lambda = 514\,nm$
1	.1914	$98.4\,nm$
2	.4393	$225.8\,nm$
3	.6887	$354.0\,nm$
4	.9383	$482.3\,nm$
5	1.1881	$610.7\,nm$
6	1.4379	$739.1\,nm$
7	1.6877	$867.5\,nm$
8	1.9377	$996.0\,nm$
9	2.1883	$1124.8\,nm$
10	2.4350	$1251.6\,nm$

Table 2.1: Zeros of the zeroth-order Bessel function of the first kind. After the third zero, zeros occur for values of $a_0/\lambda \sim 1/4$. Thus sensitivity of Time Averaged ESPI for measuring vibration amplitude is $\sim \lambda/4$.

of amplitude to wavelength is shown in figure 2.9 and the values for which the first 10 dark fringes occur are presented in table 2.1, both for a_0/λ and for a_0 in the case where $\lambda = 514\,nm$. The brightest contour in the interferogram, which shows the stationary region, is quite distinctive. It will be referred to as the *stationary fringe*.

2.3 Noise and Visibility in ESPI

In any subtraction of interferograms used to make a measurement, the noise is defined as the difference between the intensity a pixel should have due to the behaviour that is being measured and the intensity displayed in the final image. This value varies from pixel to pixel. The possible causes of noise are speckle decorrelation between interferograms and instrumental effects such as fluctuations in laser intensity, non-linearities in the video camera response, and video camera noise. In time averaged ESPI, a low number of integration cycles can also reduce or destroy visibility[8].

Any displacement of a magnitude comparable to the speckle size (eqn 2.6) of either one of the interfering wave fronts on the plane of the CCD leads to decorrelation between interferograms taken before and after the movement. Once the movement has occurred,

[8]This is not the case if the integration time of the imaging device is an integer number of cycles, *e.g.* section 4.2.2 and figure 4.3.

a given camera pixel is illuminated by contributions from the speckle that illuminated it beforehand and from the neighbouring speckle which has partially moved onto the pixel. Thus the average pixel intensity after decorrelation can be described as

$$I_2 = RI_1 + (1 - R)I_N, \tag{2.25}$$

where R is the correlation as defined in equations A.10 and 2.5, I_1 and I_N are the intensities registered on a pixel and the neighbouring pixel respectively before the movement, and I_2 is the intensity of the pixel after the movement. Explicitly, equation 2.25 states that given

$$I_1 = A_0^2 + A_r^2 + 2A_0A_r\cos(\theta)J_0\left[\frac{2\pi}{\lambda}ka_0(x,y)\right] \ ,$$

$$\begin{aligned} I_2 \ = \ & R\left[A_0^2 + A_r^2 + 2A_0A_r\cos(\theta + \Delta\theta)J_0\left[\frac{2\pi}{\lambda}ka_0(x,y)\right]\right] \\ + \ & (1 - R)\left[A_{0_N}^2 + A_{r_N}^2 + 2A_{0_N}A_{r_N}\cos(\theta_N + \Delta\theta)J_0\left[\frac{2\pi}{\lambda}ka_0(x,y)\right]\right] \ , \end{aligned} \tag{2.26}$$

where the subscript $_N$ indicates a variable corresponds to the neighbouring speckle, $\Delta\theta$ is the phase change due to the deformation of the object which leads to intensity variation as a function of object distortion. It is thus the signal in Deformation ESPI. J_0 is the Bessel function due to object vibration, and thus the signal in VESPI; $J_0 \equiv 1$ if there is no vibration.

If the illuminating wave fronts are close to uniform, it is reasonable to expect that A_{0_N,r_N} are in general similar to the values of neighbouring speckles. The value of θ_N however, has no correlation with that of θ. In general, therefore, unless J_0 is small, the dominant contribution to noise comes from θ_N.

The camera and the laser can both contribute to the noise. Fluctuations in laser intensity lead to different values between $A_{0,r}$ in I_1 and I_2. Non-linear camera response causes the subtraction of registered values for two intensities to yield a value that is not the difference between the true intensities.

Camera saturation by excessive brightness creates another difficulty. All saturated pixels register identical values and so lead to black subtractions. Therefore regions of saturation appear as dark fringes in the ESPI picture.

Artifacts in the illuminating wave fronts can also create structure in the ESPI images which is indistinguishable from the signal fringes. This can only be corrected during interpretation: the operator recognises that some feature of the image is a product of the illumination and disregards it.

Noise in ESPI is very noticeable on the dark fringes of an image. This is because the ESPI image is the absolute value of the subtraction of the interferograms. Regardless of whether the noise makes an interferogram pixel brighter or darker than the true value, the subtraction will be brighter. Thus noise looks like 'snow' on the dark regions of the ESPI image. Noisier images have more 'snow', until the noise is such that the dark regions are indistinguishable from the bright ones and information can no longer be obtained from ESPI.

It should be noted that the greater the deformation, the more ESPI is vulnerable to noise. This is because greater fringe density implies smaller fringes, which are more easily disrupted by noise.

ESPI visibility depends on the overall brightness of the ESPI image, and on the level of noise. For each image, the range of the display must be adjusted to give optimal visibility of the fringe structure. This is done by setting an upper (max) and a lower (min) threshold on the intensities, and stretching the range in between to fill the whole range of intensities from black to white on the display. Thus if the brightness range of the computer display is represented by a number between 0 (black) and 1 (white), the displayed intensity is obtained using

$$I_{display} = \begin{cases} 1 & ; \quad I_{subtraction} \geq max \\ \frac{I_{subtraction} - min}{max - min} & ; \quad min < I_{subtraction} < max \\ 0 & ; \quad I_{subtraction} \leq min. \end{cases}$$

The optimal value of max depends on the image brightness and the noise level; there is no benefit in reserving a section of the display range for intensities that do not occur, or occur only due to noise. The optimal value of min depends solely on the noise, because its function is to suppress the 'snow'. Because there is no easy way to determine the noisiness of one interferogram with respect to another, good values of max and min must be found by trial and error for every image. Clean images show their structure very clearly over a wide range of thresholds, while noisy images require more thorough

searching.

This discussion of noise and visibility is qualitative because the true test of an image is evaluated by eye: whether the fringe pattern can be understood and its fringes counted. During data taking the ESPI operator must continuously monitor fringe visibility by carrying out subtractions in real time. As long as the fringe patterns are comprehensible, the data collected is fully reliable and will need no further correction. When difficulties emerge, it is far more effective to tackle the possible causes of noise than to continue data talking in the hope of later correction through statistical methods. In this context, the lack of a full mathematical description of the noise, the signal to noise ratio, and the fringe visibility, presents no obstacle to the technique.

Chapter 3

The Oxford Static Speckle Interferometer

The Oxford ATLAS group has built an Electronic Speckle Pattern Interferometer based on an Argon ion laser to provide coherent light, optical fibres for flexibility of illumination, a video camera for image acquisition, and a computer for image display, manipulation, and storage. This chapter describes the Interferometer in detail, first explaining the general requirements and then presenting the specific hardware, geometry, and software used for measurements. The functionality of the Interferometer is then illustrated with measurements.

3.1 Physical Requirements for ESPI

3.1.1 Illumination

ESPI requires two polarised, monochromatic wave fronts of roughly uniform intensities to illuminate the object and/or image plane. The theoretical description of ESPI presented in Chapter 2 assumes that both wave fronts used to create the interference pattern are monochromatic and polarised in the same plane. If the polarisation angles differ, the interference effect diminishes with the cosine of this difference, leading to reduced visibility. If the difference in polarisation angles changes across the wave front, or with the passage of time, this leads to intensity variations in the final fringe pattern. These intensity variations are not due to object motion and are therefore noise.

Uniform intensity is necessary for uniform fringe visibility. Fringe pattern continuity

across the whole surface is necessary for fringe pattern interpretation. An interruption of fringe visibility in a fringe pattern not only destroys information about the obscured region, but also hinders efforts to relate the motion of segments of the object that are on either side of the region. Thus good fringe visibility across the whole fringe pattern is important.

Objects must be optically rough in order to produce the speckle. Because of this roughness, the illuminating light is dispersed over a wide solid angle after reflecting off the object. In general, the imaging camera only covers a small solid angle around the object. Therefore most of the illuminating light does not enter the imaging system and is wasted. This intrinsic inefficiency requires the light source be powerful.

3.1.2 Image Formation and Acquisition

Speckle interferometry is at its most useful when many images can be taken quickly and subtracted, filtered, and analysed in real time. For this the images are taken by an electronic camera connected to a computer with a dedicated 'Frame Store' or 'Frame Grabber' card. They are then processed as a a two dimensional array of numbers which represent intensities of the picture elements or *pixels* that make up the monochrome image. This array is easily stored and manipulated with image-handling software.

The speckled image of the object is formed by the optics and projected onto the photosensitive surface of the camera. The speckle size depends on the imaging optics, and the optimal speckle size for ESPI visibility and effectiveness depends on the spatial resolution of the photosensitive surface. In the case of the electronic video cameras used throughout the work presented in this thesis, the photosensitive surface is a Charge Coupled Device (CCD), with a spatial resolution (pixel size) of between 6 and 10 microns.

As the CCD cannot register structure that is smaller than its pixel size, it is useful to define the *registered speckle* as the area of uniform intensity registered and stored by the detector. In the case of a CCD camera, the registered speckle is the size of the CCD pixel. For optimal ESPI *visibility*, the optical speckle size on the CCD must be larger than or equal to the pixel size. If the optical speckle size is smaller than the pixel size, there will be more than one optical speckle per pixel, and the registered speckle will be an average of several optical speckles. This averaging reduces the effective dynamic range of intensities that the interferogram can display. If a pixel contains only one speckle, it will display the full range of intensities the speckles display, within the dynamic range of

the camera.

Optimal *spatial resolution* of ESPI fringes for a given pixel size requires that the optical speckle size be equal to or smaller than the pixel size. If the speckles are as large as several pixels, then the interferogram will register areas of equal intensity of this size, within which no fringe structure will be visible. Thus for fringe resolution to be limited by the CCD and not by the optics, the speckle size should not be bigger than the pixel size.

The ideal speckle size is the pixel size, or, more generally, the size of the spatial resolution of the photosensitive surface used to register the interferograms.

Speckle size is described by equation 2.6, namely

$$d_{speckle} = 2.44\frac{\lambda v}{D}$$

where λ is the wavelength, v the distance between the aperture and the imaging surface, and D is the diameter of the aperture.

3.1.3 Temperature Control and Thermometry

One of the main reasons for constructing the Oxford Speckle Interferometer is to study the deformation ATLAS SCT components undergo when subjected to temperature changes between room temperature and ATLAS ID running temperature (-20 C). This measurement requires a temperature controlled environment in which prototype modules can be cooled while being illuminated and imaged. Humidity control is also required, otherwise chilling the environment leads to condensation.

To help deduce the exact cause of any deformation registered, accurate thermometry is required.

3.2 ESPI Hardware

3.2.1 Laser, Fibres, and Light-Scattering Powder

Laser An Argon ion (Ar^+) laser was chosen for the ESPI system, based on a combination of power and cost. The $4p \rightarrow 4s$ transition used emits emerald-green light with a $514\,nm$ wavelength [Sve89]. The high plasma temperature in the tube causes the laser to have a Doppler line width of about 3500 MHz, less than one part in a thousand of the

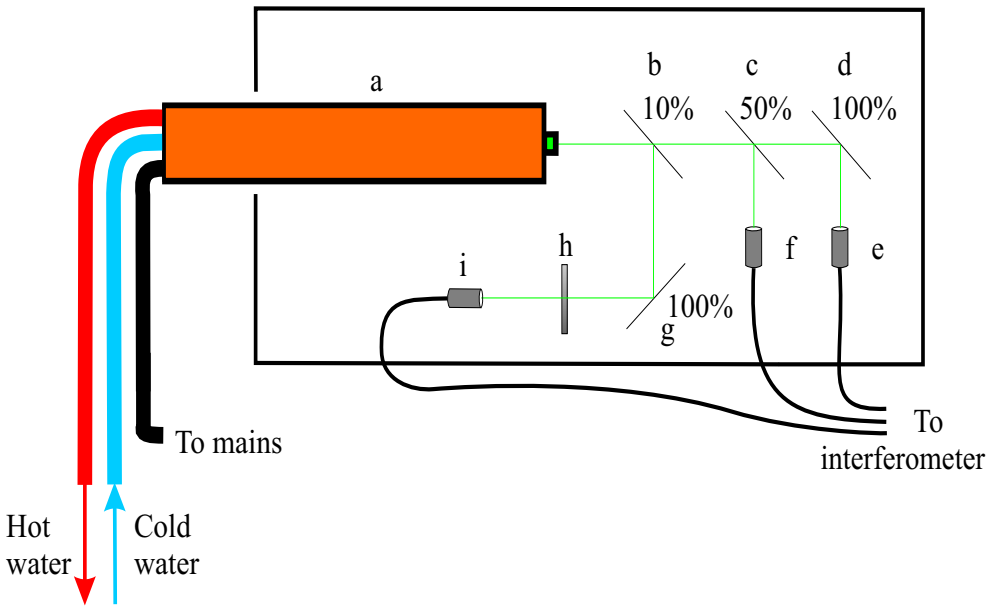

Figure 3.1: Optical layout of Oxford Speckle Interferometer. *a)* Argon ion laser. *b)* 10% beam splitter. *c)* 50% beam splitter. *d)*, *g)* Mirrors. *e)*, *f)* Optical fibre ends for object illumination: these ends have lenses which focus incoming light into the single-mode fibres. *h)* Variable neutral density filter.

laser frequency. It thus plays no role in ESPI, a technique that registers changes of order a quarter or half wavelength as it's highest resolution.

The laser used is a COHERENT Innova 90 Argon ion laser, with a 4 Watt maximum power output, and single-mode output of up to 900 mW. It has a power supply capable of delivering 40 Ampères at 300 Volts. The cooling circuit used to dissipate up to 15 kW has its own pump which circulates water through the laser tube and on to a heat exchanger which is coupled to a chilled water circuit in that runs through the laboratory.

Fibres and Coupling Optics Flexibility is gained by the use of optical fibres for light delivery, since the delivery end can be placed freely positioned. However, the requirements of wave front polarisation and uniformity necessitate the use of polarisation-preserving, single-mode fibres. The fibres also need a coupling lens at one end to direct the beam into the core, and a dispersing lens at the other end to open the beam into a broad cone of light. The fibres used, which meet these specifications, are manufactured by Point Source. They are 10 meters long and surrounded by protective metal cladding. These fibres are described in [Sou] and [Rei96].

Two fibres should be able to illuminate the object so that in-plane ESPI can be performed. In addition one fibre should also be able to illuminate the CCD directly so that out-of-plane ESPI can be done. To avoid frequent realignment, three fibres are used. The beam is first directed through a 90/10 beam splitter. The 10% goes to a variable neutral density filter and is then directed into a fibre to deliver the reference wavefront in the out-of-plane geometry. The 90% goes through a 50/50 splitter, and each of these is sent into its own fibre. The result is to gain flexibility of use at the cost of a factor of two in power for out-of-plane ESPI. A diagram of this arrangement is shown in figure 3.1.

White Powder To ensure the object has uniform optical roughness and albedo, it is covered with Developer, an aerosol that contains an extremely fine suspension of talc powder (Magnesium Silicate Hydroxide $Mg_3Si_4O_{10}(OH)_2$) in trichloroethane[1]. After thorough shaking, the developer is sprayed onto the surface of the object, depositing a uniform layer of white powder. When the solvent dries, only the powdered talc remains. Repeated sprayings result in an opaque layer which has very little mechanical integrity and thus faithfully follows the motion of the object surface.

3.2.2 Imaging Optics and CCD

A video camera connected to a computer is used to acquire the ESPI interferograms. Monochrome Charge Coupled Device (CCD) video cameras are readily available; they are compact, inexpensive, and much used in surveillance and industry. Two different CCD cameras have been used extensively for ESPI measurements. The first is a Hitachi KP-M1 with a 8.8 mm by 6.6 mm (useful area) CCD (756 by 581 effective pixels) and analogue output. The second is a Hitachi KP-F100 with 8.71 mm by 6.9 mm (useful), 1300 by 1030 (effective) pixel CCD (pixel pitch 6.7*6.7um) and 10 bit digital output.

For out of plane ESPI it is necessary to shine a reference beam directly onto the CCD of the camera. The easiest way to do this is by positioning a beam-splitter between the lens and the CCD and illuminating the beam-splitter from the side. Unfortunately, the standard casing of commercial CCD video cameras offers no access to the space between lens and CCD, and the space itself is quite small, about 1 cm. This problem has been

[1]Developer ARDROX 9D6F, sold by Brent International PLC,65 Denbigh Rd, Bletchley MK1 1PB, United Kingdom

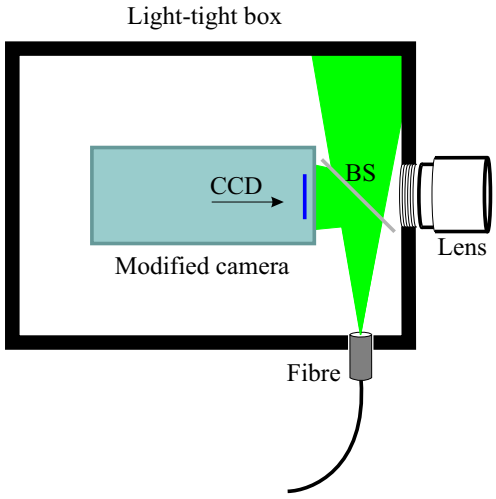

Figure 3.2: Camera, beam splitter and reference fibre for out-of-plane ESPI.

solved by mounting the camera in a light-tight box with a lens mounting cut into one side of it, and removing the front plate and original lens from the camera itself. The camera is placed inside the box leaving enough space between the CCD and the lens mounting in which to insert the necessary beam splitter. See figure 3.2.

The lens mounting allows the use of 35mm reflex-camera lenses. These lenses are designed to have $\gtrsim 2\,\text{cm}$ clearance between the last surface of the lens and the image plane, since they are meant to be used with cameras that have a movable mirror between the lens and the film. This is a convenient space for inserting the beam splitter. Thus the arrangement allows the formation of a focused image on a CCD which is also illuminated by a reference beam.

The most important lens parameter for ESPI is the diameter of the aperture, not the focal length. This is because, as shown in equation 2.6, the speckle size on the image plane depends on the size of the camera aperture, the wavelength used, and the distance between the aperture and the image plane. Given the wavelength ($514\,\text{nm}$), pupil to CCD distance ($\approx 3\,\text{cm}$), and pixel size ($\approx 8\,\mu\text{m}$) involved, small lens apertures are used. This seriously reduces the lens' light gathering capacity. To ensure enough light incides on the CCD, the camera must be brought as close as possible to the object under study, provided the whole object can be imaged.

The focal length used is thus deduced from the magnification required at a given object distance. This object distance must be previously determined by the light requirement of

the aperture-CCD combination, where the aperture itself is chosen to obtain the required speckle size.

Standard photographic lenses have variable apertures which are marked with the f-number. Once a lens of a suitable focal length has been selected, the aperture must still be fine tuned to optimise signal visibility. Note speckle size decreases linearly with aperture, while light gathering increases quadratically.

3.2.3 Chamber, Chillers, and Thermometry

Chamber and Chillers Environmental chambers have been constructed to keep the components at controlled temperatures while measurements are carried out. There are two chambers made of 10 cm thick polyurethane supported by a steel frame. The interior clearances are 30 cm height, 45 cm depth, and 45cm length for one chamber, 200 cm length for the other.

The back plate of the chamber is a 2 cm thick aluminium plate. Samples can be bolted to the back plate through the polyurethane using thermally insulating bolts. These bolts are made of tubes and thus have a low thermal cross-section. Secure mechanical contact between the sample and the backing plate is important, since the sample must be firmly fixed. This is especially true for measurements that involve subtracting intereferograms taken before and after a period of several hours. The front of the chamber is sealed off with a double-glazed window unit. The chamber is not fully air-tight, but it can be run with positive pressure of dry air (Nitrogen) to prevent condensation during cool down.

The chamber is cooled by a Neslab RTE-140 refrigeration unit, using Lotoxane®, a paraffin, as a refrigerant. This fluid is chilled and circulated through a heat exchanger on the inside of the chamber. The minimum chamber temperature that can be reached with this arrangement is close to -15 C.

Another RTE-140, identical to the one used to cool the chamber, is available and can be used to cool the object directly. The cooling pipe of prototype modules studied typically have a diameter of $\sim$ 3 mm. This small bore requires a cooling fluid that maintains a low viscosity at -15 C. Lotoxane, which has a viscosity of 1.2 mPa$\cdot$sec at minus 20 C, was chosen because of this.

Thermometry In most static displacement measurements, deformation is caused by temperature change. Object temperature is determined by gluing platinum film RTDs

(Resistance Temperature Detectors) to the object. These detectors are small ceramic tabs onto which a thin layer of platinum with a nominal resistance of 100 Ω is deposited. The variation of this resistance with temperature is well known: 0.0039327 Ω/Ω/C. The value of the resistance is accurately measured with the four terminal method, which is not affected by the resistance of the leads.

The four leads of each thermometer are connected into one input channel of a Keithley 7001 Switch System. The 7001 can switch between its twenty input channels, and is connected to a Keithley 2000 Digital Volt Meter, which performs the 4 lead resistance measurement on the selected channel.

3.2.4 Mechanical Support

Optical Support The interferometer is built upon a large aluminium optical bench. The bench is 3.5 m long by 1.5 m wide. Strengthened aluminium rails and posts can be bolted to this bench using a regular arrangement of threaded holes. These posts are used to support the video camera and the illuminating fibre ends. Thus there is great spatial freedom over the location and direction of the different interferometer elements. Schematics of the support elements can be seen in figures 3.3 and 3.4.

Object Support During a typical thermal deformation study, a prototype SCT component is fixed to a support structure which is in turn bolted to the back plate of the environmental chamber. As the temperature changes, the module and the support structure both undergo deformation. The manner in which module is fixed to the chamber must be carefully designed to minimise the driving effect of the support structure deformation on the module deformation.

Several different support jigs have been constructed, some with sliding rails or rotating axes, aimed at mechanically decoupling jig motion from the module. However, it has been observed that the static friction which must be overcome to initiate motion along such rails or pivots, even if very small, can be sufficient to drive object deformations on the scale of ESPI sensitivity. This is shown in the measurements of section 3.4.2

The most useful support structure has proved to be that with the simplest design. It is a thick quartz glass square plate with a Thermal Pyrolitic Graphite (TPG) block manufactured by T. Kohriki in KEK, Japan. It is shown in figure **??**. The monolithic plate deforms as a single object. Over constraining the plate can transform simple expan-

sion into complex twisting and buckling, so constraints must be minimal. An important feature of this support structure is the low Coefficient of Thermal Expansion of fused silica, $< 1 \times 10^{-6}$ m $\cdot$ m$^{-0.5}$ $^{\circ}$C^{-1}. This is a factor of five lower than that of silicon.

The plate has four holes, one in each corner, and they are at the right spacing to fit four thermal bolts on the back plate of the environmental chamber. It has been found through trial and error that the smallest number of fringes on this plate and the least elaborate (i. e. easier to interpret) fringes are seen when it is bolted down tightly at one of the top corners, lightly at the other upper corner to stop rotation, and just tight enough to push the plate flat against the back surface of the thermal bolt on the lower corner under the tight top bolt. This prevents flapping along the axis made by the line that joins the two top bolts.

Onto the plate can be bolted a small block of any material through which the cooling pipe can be run and to which the sample module is itself bolted down. Berillya, Thermal Pyrolitic Graphite, and Quartz Glass have all been tried. This bolting can be done with spring-loaded bolts to ensure pressure is applied regardless of thermal contraction or expansion of the nuts and bolts.

3.2.5 Interferometry Layout

The basic interferometer layout for deformation ESPI is shown in figures 3.3 and 3.4. The object is prepared for measurement and then fixed to the back of the environmental chamber. The chamber is usually sealed off with the double-glazed window unit.

The camera is placed at a distance from the object such that, with the appropriate aperture, the light entering the camera is enough for good ESPI visibility. The lens used must be of a focal length[2] that allows the image to fill as much of the CCD as possible without being truncated. Determination of aperture, object distance, and focal length are described in detail in section **??**. The direction of viewing is as close to the surface normal as possible. This makes object imaging most efficient, and also makes fringe pattern interpretation easier.

[2]Strictly speaking, it is the magnification of the lens that determines the size of the image, but in the regime of values in which ESPI is used, it is true to a very good approximation that

$$\text{Magnification} \approx \frac{\text{Focal Length}}{\text{Object Distance}}$$

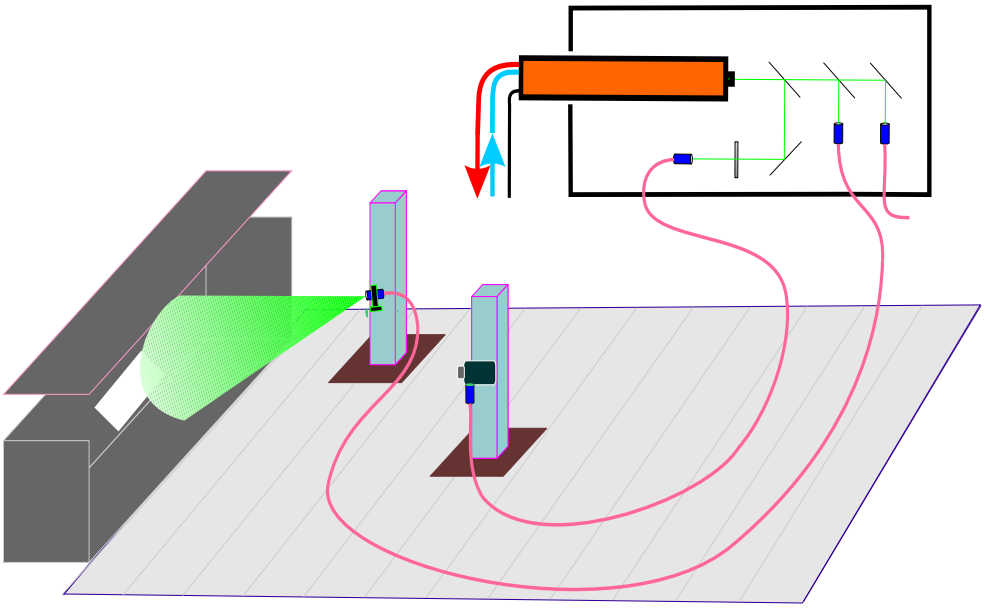

Figure 3.3: Experimental setup of out-of-plane ESPI

The distance at which the illuminating fibre ends are placed from the object is adjusted to optimise signal visibility. The ends emit a Gaussian cone of light. Therefore a large distance between fibre end and object yields greater uniformity of illumination than a smaller distance. a smaller distance, however, results in more light entering the camera from the object. There is a trade off between these two considerations. Illumination angles from the direction of viewing vary according to whether in-plane or out-of-plane sensitivity is in use. The illumination angles are usually 30° for in-plane ESPI, simply because this value makes equation 2.13 easy to use. The sensitivity becomes $\Delta l = D \cos \delta$.

For out-of-plane sensitivity, equation 2.8 shows that normal incidence is ideal. However, if the chamber is closed off by the window, some of the illuminating light is specularly reflected from it. This reflection must not enter the camera since the intense brightness of the specular reflection will saturate pixels on which it is imaged, making them useless for ESPI. The illuminating angle in out-of-plane ESPI is typically between 5° and 15° to keep this specular reflection out of the camera.

In general, the laser is run at the highest power practically achievable. Signal visibility improves with power on two fronts: aperture can go down, and fibre-object distance can go up. The intrinsic inefficiency of using light scattered from a rough surface means ESPI is usually limited by available laser power.

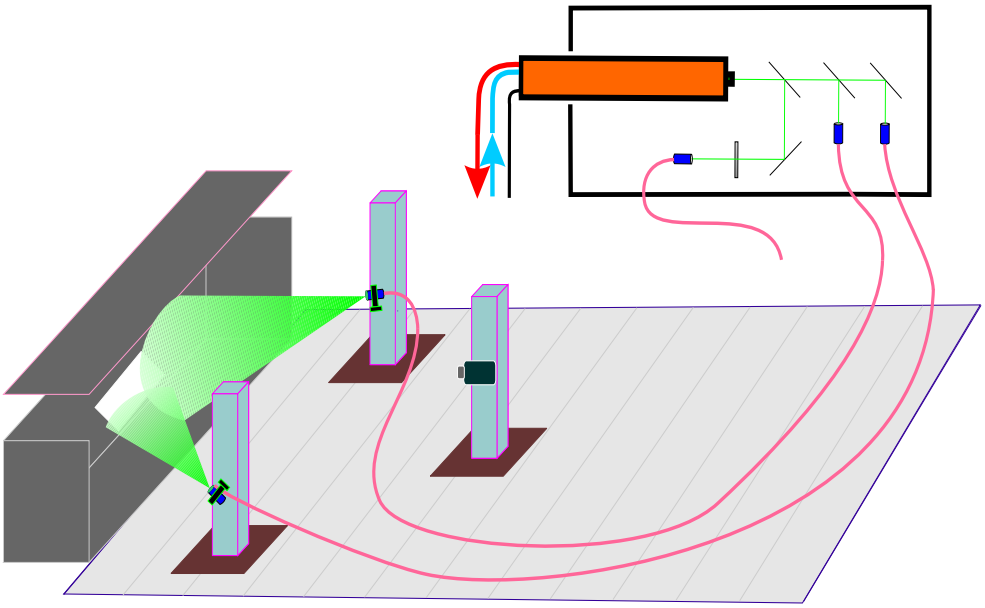

Figure 3.4: Experimental setup of in-plane ESPI

3.3 Information Processing

Two image acquisition and handling arrangements have been used, involving a CCD camera, and a desktop computer with a frame-grabber card, and specific software. In both cases, specific code was written to facilitate all ESPI operations.

3.3.1 Analogue Camera and 8 bit ADC

Earlier work was done with a Hitachi KP-M1 and a Pentium PC. The analogue signal of the camera was sent into the frame grabber card which contains an 8 bit Analogue Digital Converter (ADC). The dynamic range of intensities in the image file was thus from 0 to 255. Camera pixel size was 11 μm.

The PC ran SEMPER® by Synoptics, a programming language designed for image manipulation which contains many image-handling commands. In SEMPER, programs are written in lines of text, much like in FORTRAN or C. Whole programs can be run, but it can also be used line by line to carry out specific commands. Thus it is extremely versatile since the operator can input one or a few lines of text and execute them in seconds.

The measurements were taken in this manner, with existing code being modified if

necessary and a few lines being selected and executed when required. To prepare for a new set of measurements meant at most to alter some parameters in existing lines of code and to write two or three new loops. The array of mathematical tools available in SEMPER is excellent.

3.3.2 10 bit Digital Camera

Later measurements, including most of the vibrational work, were done with a Hitachi KP F-100. This camera has a 10 bit digital output, giving 4 times greater dynamic range than the KP-M1. The pixel size is 6.7μm, which is a noticeable improvement in spatial resolution. As discussed in section 2.1.3, smaller pixel size can allow greater pupil aperture and thus more light acceptance.

A 10 bit frame grabber card capable of handling images of up to 4 Mbytes was installed in a 233 MHz Pentium II PC. It was decided to use National Instruments' LabVIEW® as the image handling software for the KP F-100. Unlike SEMPER, LabVIEW is not custom built for image manipulation. It was selected to enable integration of image acquisition software with other ESPI software (such as thermometry readout and timing control) and, and the desire to have software compatibility with other interferometry measurement programs under development at Oxford.

LabVIEW is in general easy to use for controlling hardware since specific control modules are for many different devices. These modules (or *VI*s[3] have been written) have control panels that are equivalent to the front of the instrument, and are fully linkable to any other VIs, which makes them easy to integrate into larger programs.

A special set of VIs for image manipulation (IMAQ) has been developed by National Instruments. These image VIs together with camera control VIs for the KP-F 100 and VIs developed to control the Keithley 2000 and 7001 were used to construct VIs to control the ESPI system as a whole.

3.4 Measurements

One of the main applications of the Oxford ESPI facility has been to measure the deformation of prototype modules caused by changes of temperature. Measurements of this

[3]The VI or Virtual Instrument is the unit of functional code in LabVIEW, analogous to a subroutine in FORTRAN

kind are described in detail in [Rei96]. Two such measurements, done on prototype ATLAS ID silicon strip modules, are included here to illustrate the technique. One has been chosen to highlight the strength of ESPI, and one to illustrate the difficulty of making deformation measurements at the micron scale. SCT silicon strip modules are described briefly in 1.1.4.1.

3.4.1 KEK End-Tap Module

Module A prototype End-Tap SCT module was designed and constructed at KEK, Japan. A diagram of the module is shown in figure 3.5. The 'End Tap' design places the readout electronics at the end of the the module, rather than across its centre. This design offers the advantage of separating the heat generated by the readout electronics from the detectors themselves, but it is electrically less favourable for readout and was eventually abandoned in favour of a 'Centre Tap' design. However, this measurement is presented because the unobstructed surface of the silicon of an End Tap module presents an ideal sample for ESPI.

The module was constructed by gluing the two layers of silicon on either side of a 1 mm thick plate of Thermal Pyrolitic Graphite. This directional thermal conductor is placed between the silicon wafers in order to conduct the heat generated by the silicon into the module cooling pipe. Four cutouts were made in the TPG to reduce the total mass of the module. The diagram clearly shows the cutouts. A broad tab of TPG sticks out 1 cm beyond the side of the silicon, to conduct heat out of the module across a wide front. This tab is also used to hold the module in place.

Mounting The module was mounted by bolting the TPG tab on a TPG block which had a cooling pipe running through it. This block is in bolted down on a massive glass plate, which is itself affixed to the back of the environmental chamber.

Measurement The measurement shown in figure 3.6 was taken by placing the KEK module in the environmental chamber, sealing the chamber, and cooling it. No cooling fluid was circulated through the module cooling pipe. The whole module was given several hours to achieve thermal stability. Thermometers on the module registered a temperature difference of 4C. The interferometer geometry used was out-of-plane sensitive, with a relative displacement of 0.24μm per fringe. The overall out-of-plane motion of the far

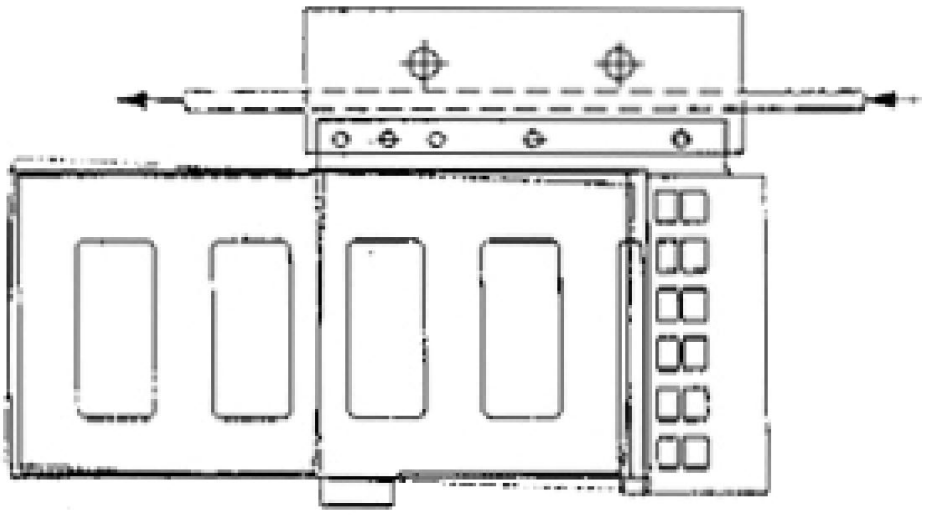

Figure 3.5: Diagram of KEK End Tap Module Note TPG cutouts to reduce mass.

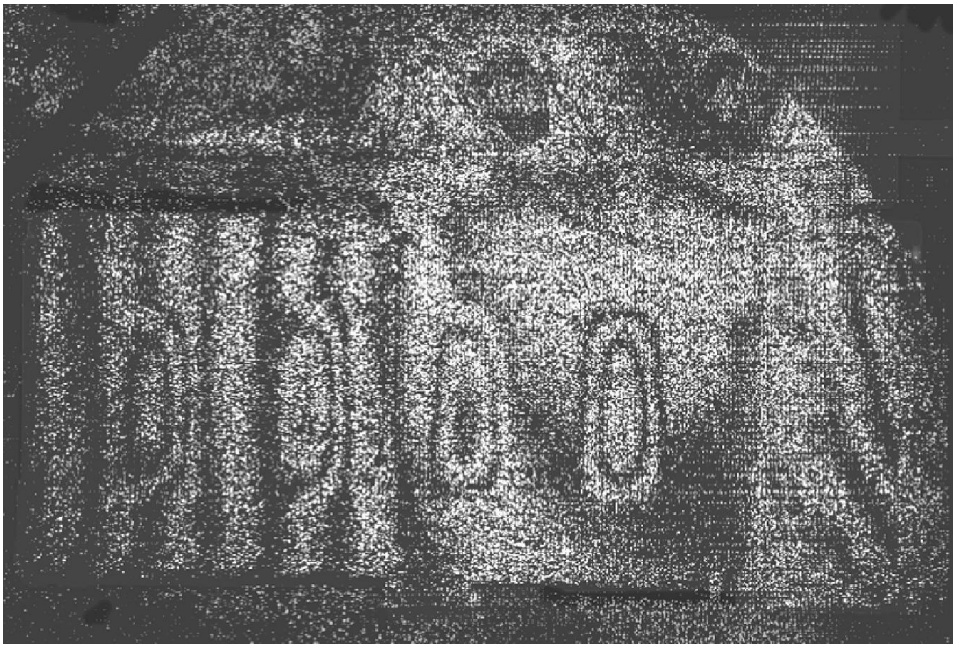

Figure 3.6: Thermal deformation of KEK End Tap module Out-of-plane ESPI fringe pattern due to a temperature change. The silicon bulges where it is not glued down to the TPG, revealing the location of the cutouts.

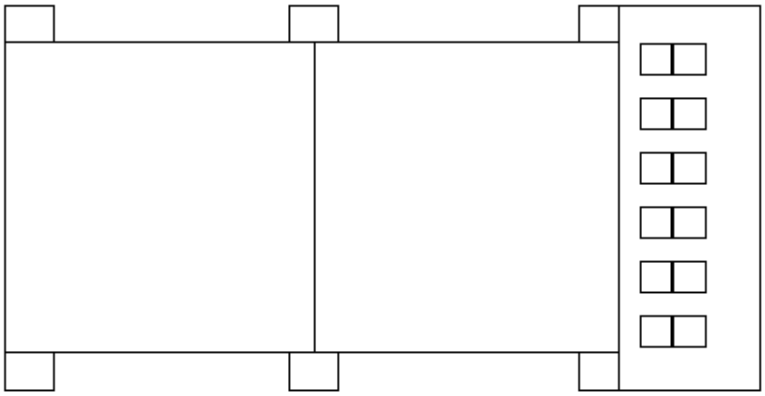

Figure 3.7: Diagram of RAL End Tap module

(left) end of the module is 6.5 fringes, or 1.6 μm, relative to the TPG mounting block above the module

The most distinctive features of figure 3.6 are the silicon bulges over the TPG cutouts. They show that a local variation in height over a small area of module surface may be of up to 1/3 of that over the whole module length. More importantly they illustrate the true strength of ESPI, which presents an image that helps visualise the displacements, as well as measuring them. This gives the module designers insight into what is driving module behaviour, and how to fix it.

3.4.2 RAL End-Tap Module

Module The other measurement presented in this chapter is of an 'End Tap' prototype developed by RAL (Rutherford Appleton Laboratory), UK. This module has three thin slabs of Berillya between the silicon layers, one in the middle where the two 6 cm tiles are joined, and one at each end. These stick out sideways from the silicon on both sides, providing three fixing points per side.

Mounting The support structure is a long aluminium bar with the cooling pipe running through it. This bar is free to rotate, and one side of the module is bolted to it at the three tabs. The module is then allowed to hang from the cooling bar. A second, flexible, cooling pipe is connected to the three bottom tabs of the module. This pipe hangs from the module and is not considered part of the support of the module.

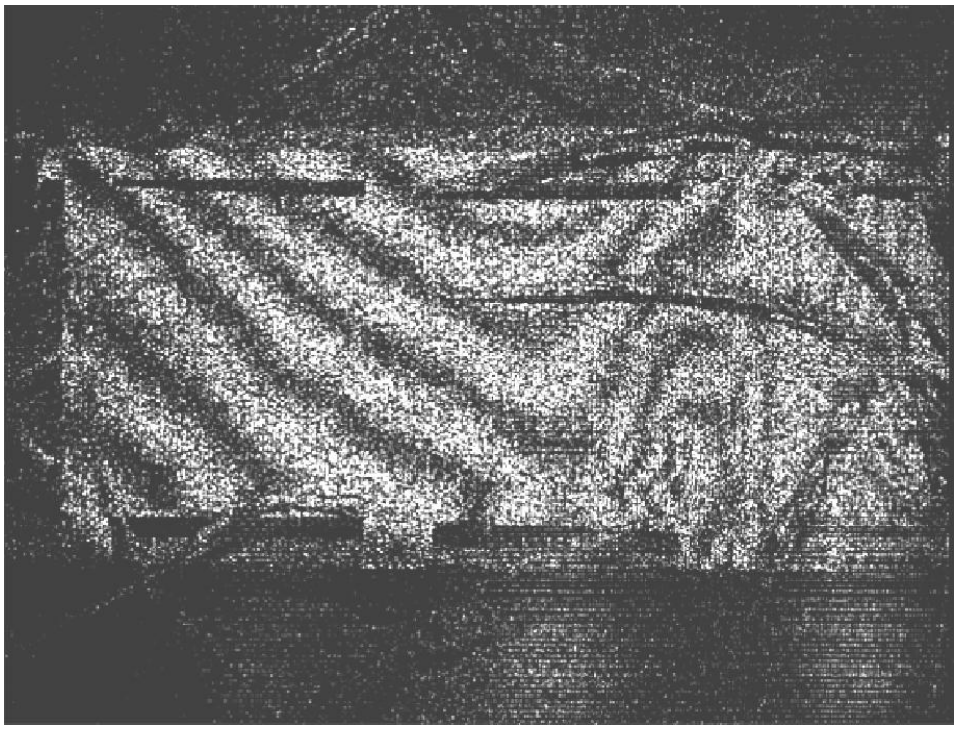

Figure 3.8: Thermal deformation of RAL End Tap module
Out-of-plane ESPI of RAL module subjected to a power change.

Measurement The module was placed in the environmental chamber which was cooled down to -10C. After the module reached thermal equilibrium, the first interferogram was taken. The module was then subjected to thermal variation by dissipating 1 W of power on heaters stuck to the hybrid area to simulate the thermal effect of the readout chips. No fluid was circulated through the cooling pipes. After 10 minutes, once the module had achieved thermal equilibrium, the second interferogram was taken. The resulting out-of-plane measurement of the deformation is shown in figure 3.8.

The salient feature of this fringe pattern is the clear fringes along the length of the support bar running along the top of the module. It is significant that the support bar fringes seem to align with fringes along the silicon, suggesting that the bar plus silicon system is deforming as one continuous object. The next observation to note is the curve of the fringes along the centre of the module. This can be interpreted either as a buckling of the module under compression, or as a reduction of an existing buckle under the reduction of compression. Given that the difference is caused by an increase in temperature and power dissipated, which should lead to expansions, the second option

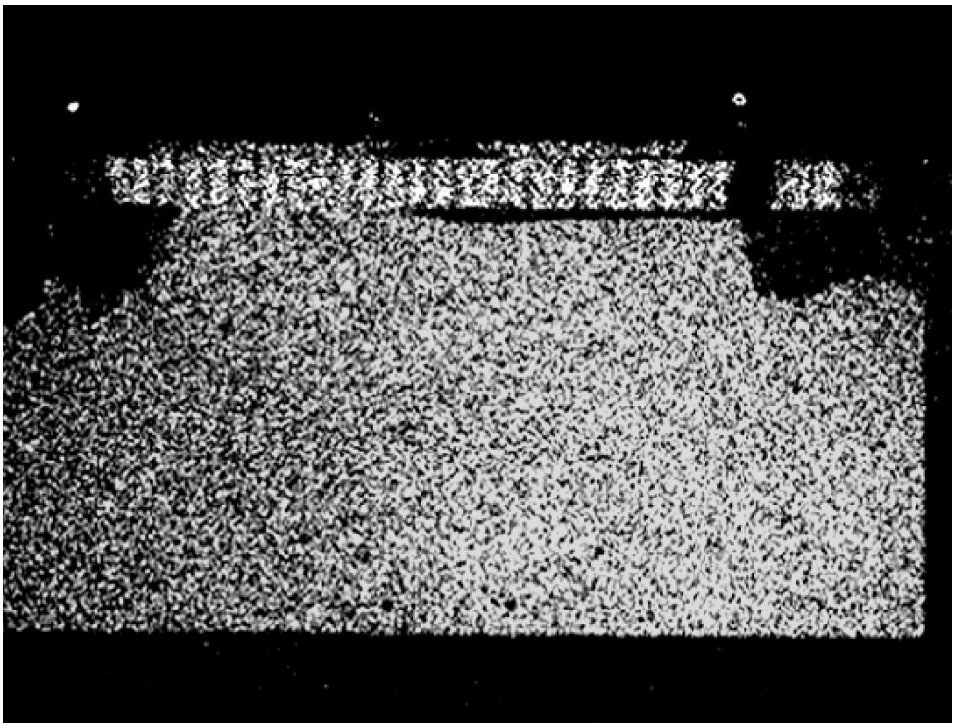

Figure 3.9: Thermal deformation of RAL support structure In-plane ESPI of RAL support bar.

seems more likely. It implies that the initial cooling of the module and environmental chamber created a buckling. This in turn points to the possibility of the aluminium bar contracting faster than the silicon, during the cooling phase, causing the buckling. The CTE of aluminium of aluminium is $23\,ppm/K$ and that of silicon is $2.6\,ppm/K$.

3.4.2.0.1 Module Support This analysis lead to a measurement being taken of the support bar on its own. Figure 3.9 was created using the in-plane sensitive geometry, and thus shows a clear, uniform contraction due to cooling. Overall, the RAL module measurements show that it is difficult to study a module in isolation, and furthermore, that for any ESPI module measurement to be a credible reflection of how modules will behave inside ATLAS, it is essential to perform measurements on them mounted in the support structure that will be used in the detector.

Chapter 4

The Oxford Vibration Speckle Interferometer

This chapter describes the Time Averaged Vibrational Electronic Speckle Pattern Interferometry (VESPI) work done at Oxford, which, in combination with Finite Element modelling, is used to validate the understanding of certain prototype Carbon Fibre support structures used in the ATLAS SCT.

4.1 Vibration of ATLAS SCT Structures

4.1.1 Possible Causes and Effects of SCT Vibration

The ATLAS Inner Detector Semiconductor Tracker is being assembled from thousands of components, ranging from support brackets and modules, centimetres in size, to the Carbon Fibre Reinforced Plastic (CFRP) cylindrical shells, 1.6 metres long, to support the modules. Most of these components must meet the two criteria of low mass and high stiffness. Physics performance specifications require that the location of all active elements be know to a precision of $\approx 10\,\mu$m, see section 1.2.

Given this low mass / high stiffness requirement, the SCT is both susceptible and sensitive to vibration. Furthermore, vibrations in the support structures could interfere with the Frequency Scanning Interferometry (FSI) alignment system that has been developed for the ATLAS SCT. Thus it is of interest to anticipate the vibrational spectrum of the SCT components where possible.

4.1.2 Predicting Fundamental Frequencies of Plates and Shells

4.1.2.1 Analytical Description of Oscillations in 3D Plates and Shells

Although some of the basic results of vibrational physics are well known, the general problem of oscillations in three-dimensional objects is a complex one. Historically, the accepted practice was to introduce simplifying assumptions based on the characteristics of the object, in order to reduce the number of terms in the equations of motion. Specific frameworks were developed to address the motion of certain forms, and in certain parameter regimes, such as *plates* (3D objects where extension in one dimension, named *thickness*, is much smaller than the other two) and *shells* (plates with curvature normal to the thickness dimension). Both circular plates and closed cylindrical shells will be used in the ATLAS ID. [ATLa]

A consequence of this complexity is that estimating even the lowest frequencies of oscillation for simple shapes is a challenging problem. As an example, figure 4.1 shows four plots displaying the lowest n (circular nodes) modes of a cylindrical shell as predicted by the Flügge theory for four values of R/h where R is the radius of the cylinder and h the thickness of the shell. It should be noted that in order to address a question as simple as 'what modal shape has the lowest frequency' it is necessary to know Ω (function of frequency, density, Youngs modulus and Poissons ratio), the length-to-radius ratio, and the the radius-to-thickness ratio of the shell. The four plots, as well as their individual captions, were taken from a 1973 monograph on the vibrations of shells [Lei73].

Figure 4.2, also taken from [Lei73], shows the discrepancies between two of the several accepted theoretical approaches used to address the problem of vibrations in shells, the Flügge and the Donnell theories.

Within this context, FEA was seen as a valid alternative to predict the motion of plates and shells, although of limited usefulness because accuracy can only be achieved with a large number of elements, and thus many computations. The present-day abundance of powerful computers, however, has made FEA the method of choice.

4.1.2.2 Finite Element Analysis of Vibrating Bodies

Finite Element Analysis (FEA) is a numerical technique which models a continuous body as a number N of discrete, mutually connected components of a finite size, called *elements*. These elements are linked together to form a system of equations which de-

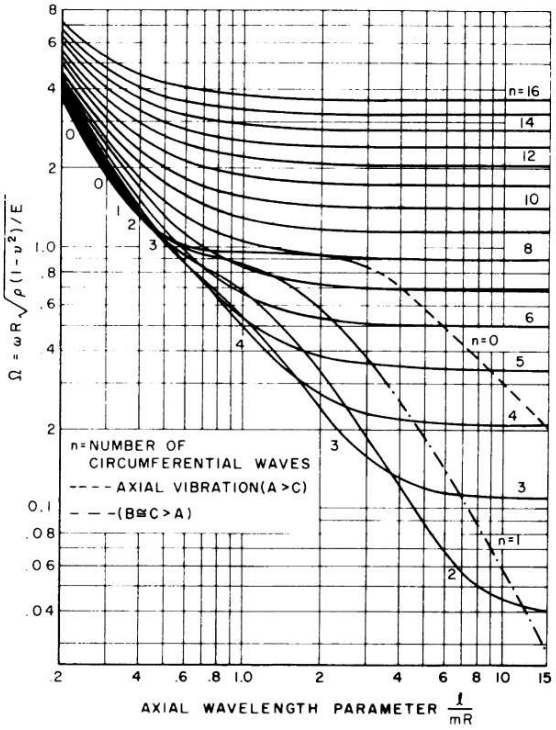

FIGURE 2.14.—Variation of the fundamental frequency parameter Ω with l/mR according to the Flügge theory; $\nu = 0.3$, $R/h = 20$. (After ref. 2.69)

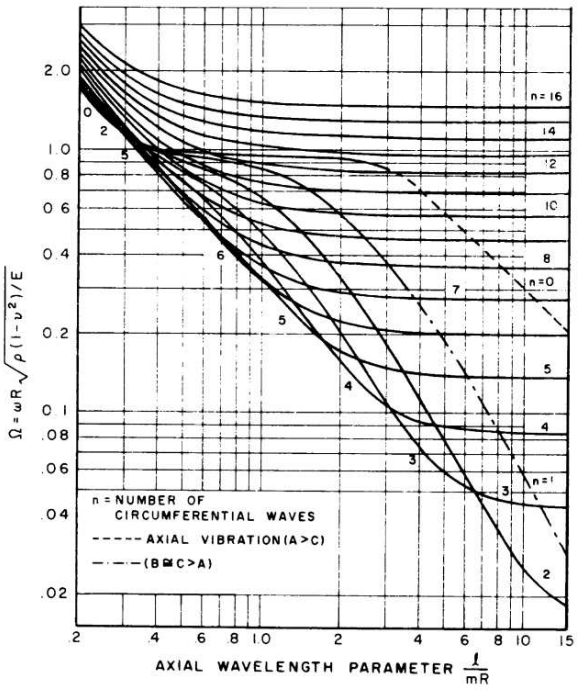

FIGURE 2.15.—Variation of the fundamental frequency parameter Ω with l/mR according to the Flügge theory; $\nu = 0.3$, $R/h = 50$. (After ref. 2.69)

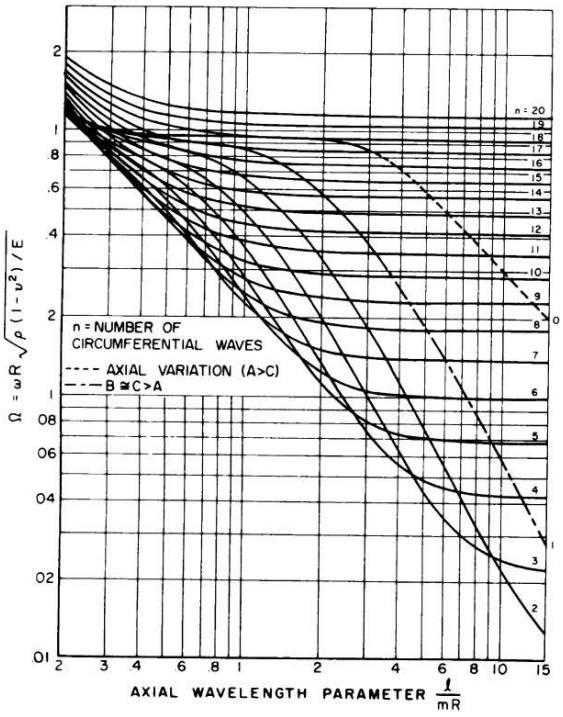

FIGURE 2.16.—Variation of the fundamental frequency parameter Ω with l/mR according to the Flügge theory; $\nu = 0.3$, $R/h = 100$. (After ref. 2.69)

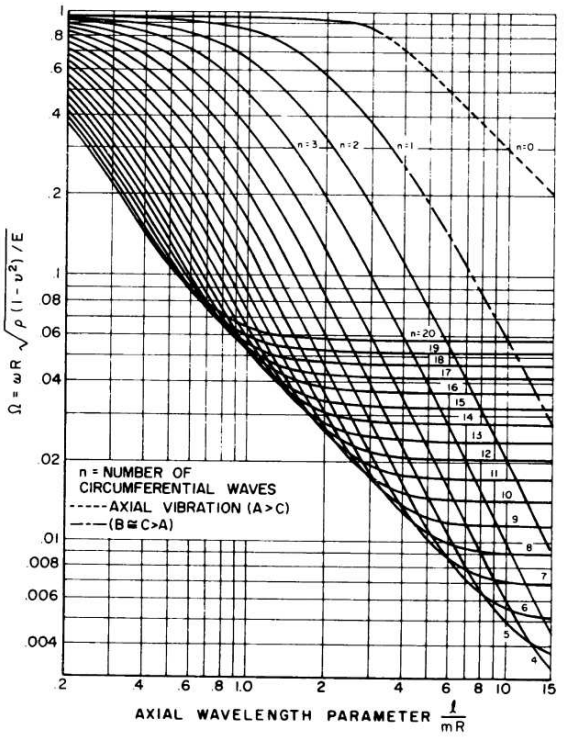

FIGURE 2.17.—Variation of the fundamental frequency parameter Ω with l/mR according to the Flügge theory; $\nu = 0.3$, $R/h = 2000$. (After ref. 2.69)

Figure 4.1: Flügge theory plots. Figure is facsimile of [Lei73] p64. See caption to figure 4.2.

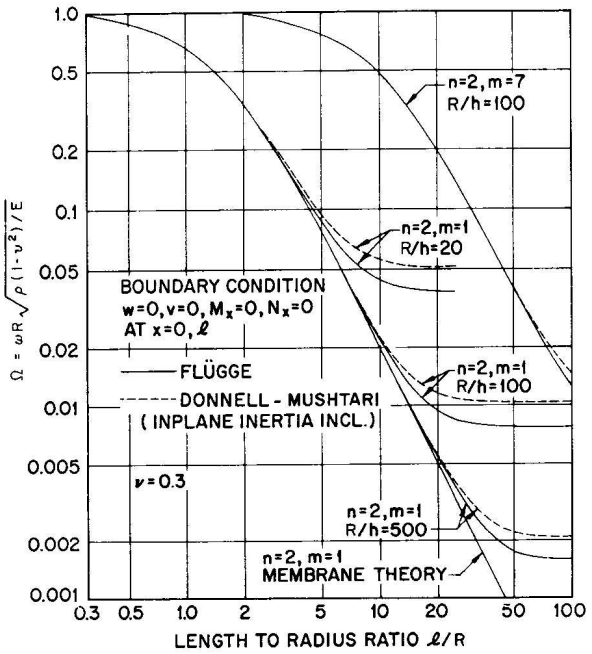

FIGURE 2.13.—Comparison of Flügge and Donnell frequency spectra for $n = 2$. (After ref. 2.35)

Figure 4.2: This figure, taken from [Lei73], shows the discrepancies between two accepted theories of vibration of cylindrical shells, the Flügge theory and the Donnell theory. **Previous figure, figure 4.1**, presents four Flügge theory plots showing circular modes, indicated by the letter n, which vibrate at the lowest Ω (the "frequency parameter", a function of frequency, density, Youngs modulus, and Poissons ratio) for a given m (the number of axial nodes), as a function of the ratio of the cylinder's length to radius, l/R. The four plots show different values of R/h, the cylindrical shell showing circular modes, indicated by the letter n, which vibrate at the lowest Ω (the "frequency parameter", a function of frequency, density, Youngs modulus, and Poissons ratio) for a given m (the number of axial nodes), as a function of the ratio of the cylinder's length to radius, l/R. The four plots show different values of R/h, the cylindrical shell radius-to-thickness ratio.

termines the motion for all the elements, ie of the whole body. FEA thus replaces a calculation for a single object, which could be intractable, with many, more manageable calculations [Spy94].

The first step in FEA is to divide the body into a finite number of elements, *i.e.* discretise it. Choosing the right element shapes, sizes, types, and number is key to a successful analysis. Structures and inhomogeneities within the material must be considered, in order to make each element as simple as possible. The number of elements adds accuracy at the cost of greater computation, so tradeoffs must be made, assigning many elements to critical or complex regions and fewer to simpler parts.

The next step, which is entirely done by the computer, is to work out the stiffness and damping matrices that will determine the motion of the system of coupled oscillators:

$$[M]\ \ddot{x} + [C]\ \dot{x} + [K]x = F(t). \tag{4.1}$$

The vibrational motion of the object, is then approximated by calculating the motion of the resulting coupled oscillator of N components.

Here $\ddot{x}, \dot{x}$, and x represent the acceleration, velocity, and location vectors, and $[M], [C]$, and $[K]$ are the mass, damping, and stiffness matrices, and F(t) is the vector of applied forces [Spy94]. The specific coefficients of the matrices are derived from the material properties as entered into the model. The material properties needed to specify vibrational motion are the density and the elastic moduli (Youngs modulus, Poissons ratio, bulk modulus, shear modulus). As a final step, the constraints of motion applied to the body must be entered.

The eigenfrequencies are found by solving the characteristic equation of the system. Direct substitution of these eigenvalues yields the normal modes, of which there are as many as there are degrees of freedom. The lowest modes of an FEA calculation have shapes and frequencies that are very similar to those of the corresponding continuous object.

Professional FEA software packages exist for a wide range of applications, including thermal, electromagnetic, and mechanical. These packages contain pre-defined element types and segmentation algorithms that make the analysis easier to perform.

The Oxford Physics Department Design Office uses ALGOR, a powerful Engineering package which allows the whole process to be carried out within one application, from drawing the geometry of the object to viewing the deformation in a coloured, contoured

graphical output. This is useful for comparison with ESPI work since the deformation contours are recognisable and easily matched with ESPI fringe patterns.

Specific FEA analyses carried out for this thesis are described in sections 4.4.2.1 and 4.4.3.1.

4.1.3 Measurement Program Aims

The objective of the present VESPI measurement program is to analyse ATLAS SCT Carbon Fibre prototype structural components to find a set of resonant modes and frequencies that allow validation and possibly tuning of FEA simulations of these prototypes.

This objective requires the assembly and testing of a working VESPI system, capable of finding and identifying fundamental modes on samples of over a meter in scale. Prototype samples must be prepared and measured, and the results analysed to form as complete a spectrum as possible. Separately, detailed FEA simulations of these prototypes must be carried out using manufacturers' specifications where known, and educated estimates where not known, as inputs for the material properties of the samples.

Comparison of the VESPI and FEA modal frequencies may yield consistent or inconsistent results. Gross inconsistencies will most likely be a sign of error, either in the experimental procedure or in the computer simulation work. Moderate inconsistencies may be ameliorated through FEA parameter tuning, which may be justified within some range for parameters which are not well known. Fully consistent results should give confidence that the support structures are well understood and that the FEA modelling is accurately simulating the object under study.

4.2 Assembly of the VESPI interferometer

4.2.1 Physical Requirements

VESPI has the same illumination requirements as deformation ESPI, section 3.1.1. Section 2.2.2 shows that the integration time over which the interferograms are made must be either an integral number of periods of oscillation of the object under study, or greater than or about ten periods. This is to ensure all parts of the oscillation cycle influence the average with equal weight. Equation 2.24 shows that subtracting two interferograms of the same object vibrating in the same way increases fringe visibility, provided there

is a difference in phase differences between the object and reference beams in the two interferograms.

VESPI can be used to search for the modes of vibration of an object at different frequencies, particularly its natural frequencies. This is done by holding the object in a support structure that resists vibration - (typically by being massive), and stimulating it to vibrate at a specific frequency. The driving signal must be close to monochromatic and measured accurately in order to know the frequency at which the object is stimulated. The support structure should hold the object in such a way that its influence on the vibrational behaviour of the object is minimised, and preferably in a way that results in simple boundary conditions.

Thus VESPI requires camera timing control, interference beam relative phase control, a driving signal generator with mechanical (*e.g.* acoustic) coupling to the sample, a frequency counter, and adequate mechanical support for the sample.

4.2.2 Hardware and Software

Piezo controlled mirror The relative phase of the two beams of the interferometer is controlled by a mirror mounted on a piezo-electric crystal. Due to the piezoelectric effect, the size of the crystal is controlled by a voltage applied across it. The mirror is placed in the path of the reference beam and the voltage is varied between image acquisitions. Thus the optical path travelled by the beam can be varied by varying a voltage.

A voltage controller circuit has been designed and built for the piezo. The voltage difference needed to change the path of the laser by one wavelength is applied between two chains of resistors. One chain has four equal resistors, and the other ten. The operator selects the phase difference between two images by switching the terminals of the piezo voltage between different resistors. Phase differences of an integer number of quarters or tenths of 2π can thus be chosen.

Acoustic driving A sinusoidal signal generator, capable of generating an AC voltage with variable frequency and amplitude, is connected to a frequency counter and a loudspeaker of 4 or 8 ohms. This is used to acoustically drive all samples (except the Carbon Fibre cylinder of section 4.4.3) at known frequencies. The signal generator used is an ADVANCE LF Type S. G. 66, and the frequency counter is an Aplab Digital 15MHz Time and Frequency Counter Type IC12. Typical output powers range from 1 to 20

watts.

Timing electronics Control electronics have been built to ensure that the camera integration time matches an integral number of object oscillation periods. The circuit contains a Maxim 038 High-Precision Waveform Generator, a Maxim 442 Power Amplifier, a Mach 210 Programmable Logic Device (PLD), a hexadecimal selector switch, and a potentiometer. A diagram of the control system can be seen in figure 4.3. The frequency at which the Waveform Generator oscillates is determined by the potentiometer. The Generator has one sinusoidal output and one square wave output. The sinusoidal output is sent through the amplifier to drive the loudspeaker. The square output is one of the three inputs of the PLD. The other two are the hex selector switch and a TTL signal from the computer. The PLD, which is programmed in Palsam, starts operation when it receives an enable pulse from the computer. It waits for the next edge in the square wave and sends a trigger pulse to the camera to begin integrating. It then counts the number of square wave cycles, and when this number equals the number selected on the hex switch, it sends the 'stop' trigger pulse to the camera.

This control system allows the operator to determine the integrating time, modulo the oscillation period. The amount of light registered on the CCD can thus be varied while ensuring the camera integrates over whole cycles of motion. Fewer cycles are selected at lower frequencies, and more at high frequencies, to place the registered intensities within the dynamic range of the camera by controlling the total integration time.

4.3 Tests of VESPI

4.3.1 Rubber Drum and Preliminary Amplitude Contours

To establish that vibrational contour fringes could be obtained, a tunable circular drum was constructed. The circular membrane was chosen because it is easily approximated by a mathematically tractable problem, it is easy to construct and tune, and it can be made to resonate to large amplitudes. The fundamental modes of oscillation of the circular membrane are available from an Acoustics text such as [Hal92]. The lowest modes, labelled (d, c) for the number of nodal diameters and circles, are shown in figure 4.4.

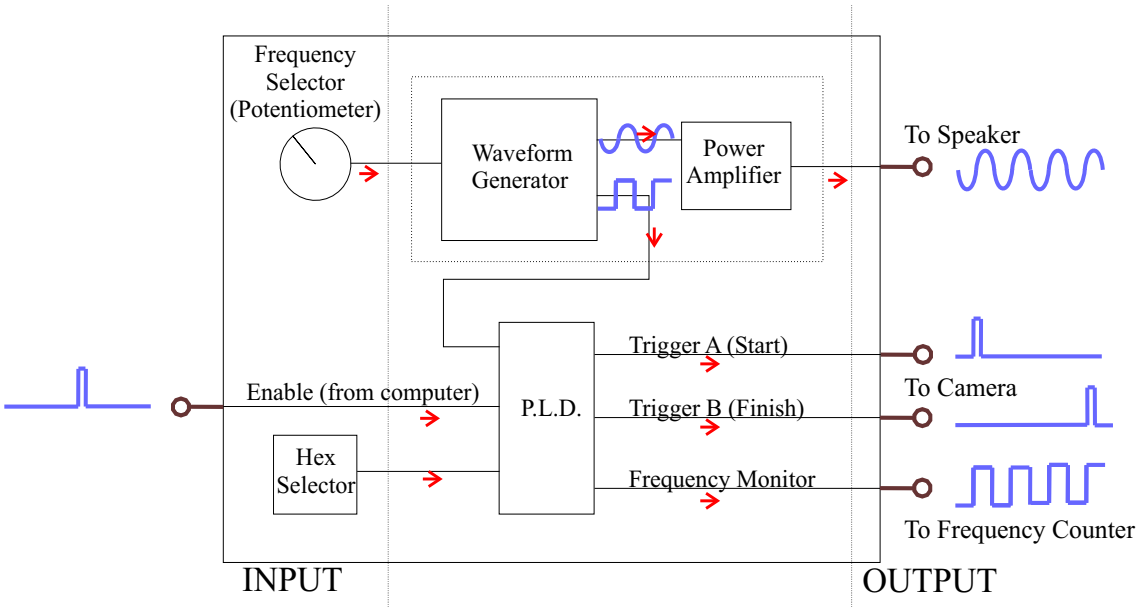

Figure 4.3: Schematic of camera timing control circuit for low-frequency VESPI. The system ensures the camera integration time is an integer multiple of the period of the driving oscillation. The frequency and the number of cycles is selected by the user. Acquisition is triggered by a TTL pulse sent by the image acquisition computer running VESPI software on LabVIEW.

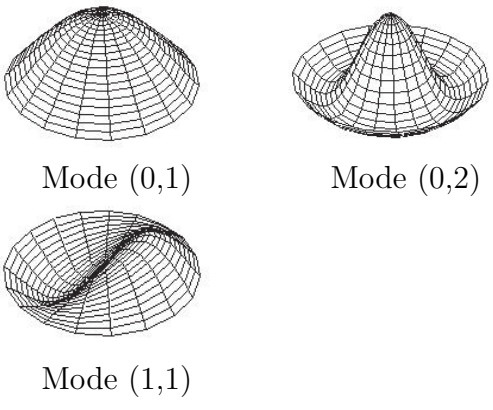

Figure 4.4: Flexible Circular Membrane under Tension. Modes are labelled (d,c) to indicate the number of nodal diameters and circles. The modes corresponding to the the three lowest resonance frequencies are shown.

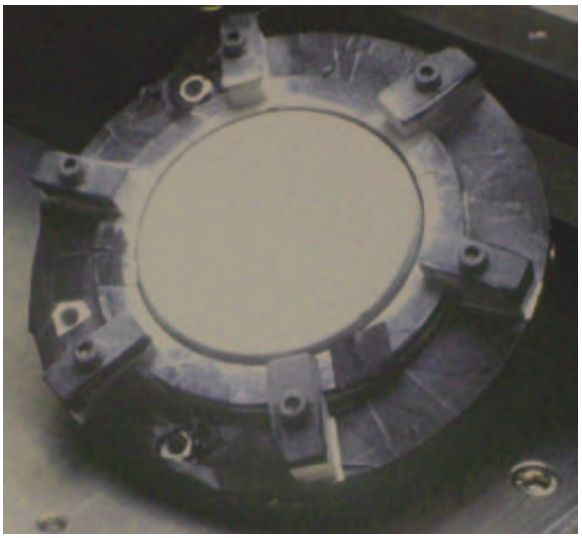

Figure 4.5: Drum Test piece. 10 cm diameter drum frame made of aluminium. Rubber membrane is tensioned and covered in white powder (developer) to make the surface reflective and rough. Aluminium parts covered in black tape to reduce unwanted reflections.

A thick circular frame capable of holding and tensing a circular membrane was built out of aluminium, and circular sheet of rubber was stretched in it, see figure 4.5.

A 20-watt car loudspeaker, driven by the signal generator, was held near the membrane and the driving frequency was varied. When the driving frequency coincided with the fundamental frequency of the membrane, there was a notable increase in the sound heard, and the vibration of the membrane could be seen.

Vibration patterns of a membrane can be visualised by sprinkling sand on the vibrating surface. If the surface is held horizontally, sand settles into the regions that are static (the nodal regions) and jumps away from the regions of high-amplitude oscillation. Using this method, the (1,1) mode of oscillation of the test membrane was found while driving it with the speaker at 396 Hz. The black sand formed the contour seen in figure 4.6 *a*). The surface was then illuminated with laser and driven at a much lower amplitude. Several interferograms at different phases were taken and subtracted. The resulting pattern is shown in figure 4.6 *b*).

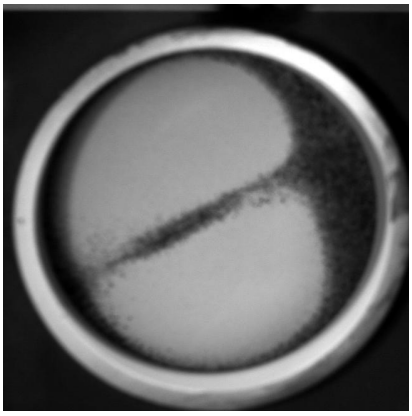

a)

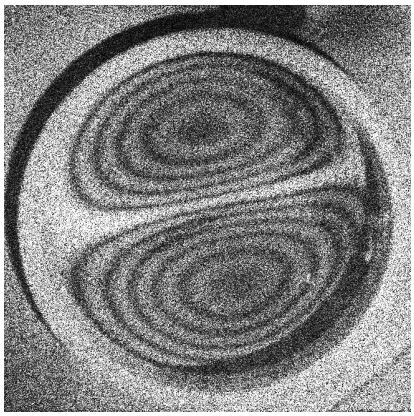

b)

Figure 4.6: Tunable drum skin driven at Mode (1,1)
a) Photograph of membrane several seconds after being uniformly sprinkled with black sand. The sand travelled away from the high amplitude regions and into the nodal regions.
b) Vibrational ESPI fringe pattern created immediately afterwards at a much lower amplitude. The six fringes between the node and the anti node indicate a maximum amplitude of 740 ± 70 nm according to table 2.1.

4.3.2 Capacitive Sensor and Amplitude of Vibration

Figure 4.6 *b)* presents a clear fringe pattern showing the contours of the (1,1) mode of oscillation. The shape of the mode shown by these contours is seen independently in the sand pattern of figure 4.6 *a)*. However, it does not establish that the value for the amplitude of object vibration obtained from VESPI using table 2.1 is in agreement with the true amplitude. An independent measurement of amplitude was required to validate the use of VESPI for amplitude measurement.

A capacitance distance sensor was used to independently measure the amplitude of vibration of a metallic oscillating surface. Such sensors can measure the distance between two plates of a capacitor to great accuracy. Neglecting edge effects, the capacitance of a two-plate capacitor is given by

$$C = \epsilon \cdot \epsilon_0 \frac{A}{D}$$

where ϵ_0 is the permittivity of the vacuum and ϵ is the relative permittivity, A is the plate area and D is the plate separation. Measurement of z_C, the impedance of the capacitor, is done by passing an AC current of known frequency ω through it, measuring the AC voltage V, and using $V = I/j\omega C$. From this the capacitance is deduced using $z_C = 1/j\omega C$.

The device used is a Micro-Epsilon capaNCDT which consists of a cylindrical head

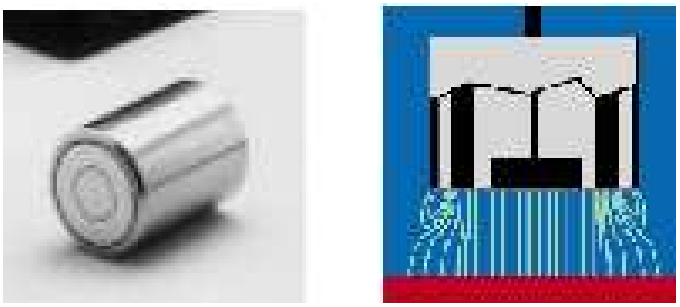

Figure 4.7: Micro-Epsilon circular capacitor plate and guard ring

with a flat circular surface, divided into two concentric circles, and electronics for readout. The inner circle is a capacitor plate, and the outer ring is a guard ring that protects the inner plate from edge effects. This is illustrated in figure 4.7.

To use this device, the capacitive plate surface must be held very close (less than 1mm) to a flat metal surface, which acts as the second plate of the capacitor. The electronics send a high-frequency signal down a coaxial cable to the capacitor head. The electronics output signal is a voltage, V_{out} which is proportional to the capacitance measured. The V_{out} must be calibrated such that a 10 V difference corresponds to a difference of 1 mm in the distance between the capacitive sensor head and the metallic surface under study. Once done, the output voltage resolution is millivolts on a voltmeter or an oscilloscope. If the surface of the object oscillates relative to the fixed capaNCDT sensor, the voltage output will oscillate at the same frequency and with an amplitude ratio of 10 mV/μm[1].

In order to simultaneously measure, both with the capacitive probe and with VESPI, the amplitude of vibration of a driven object, a suitable test piece was needed. A flat metal 10cm^2 plate was used. A support structure which could pin down two corners of the square plate while holding the capacitor head fixed near its surface was constructed out of a thick aluminium bar. A loudspeaker was held behind the support bar and capacitor, and the other surface of the plate was covered in developer powder. The capacitor measured the movement of the back surface of the plate, and the front surface was simultaneously analysed using VESPI.

Figure 4.8 shows a schematic drawing of the setup, seen from the side. Figure 4.9 shows 2 pictures of the setup seen from the front, and the resulting measurement. Pic-

[1]This assumes the frequency of mechanical oscillation is much lower than the AC frequency used by the Micro-Epsilon electronics. The work done in this thesis is in the low audio regime, where this condition is always true.

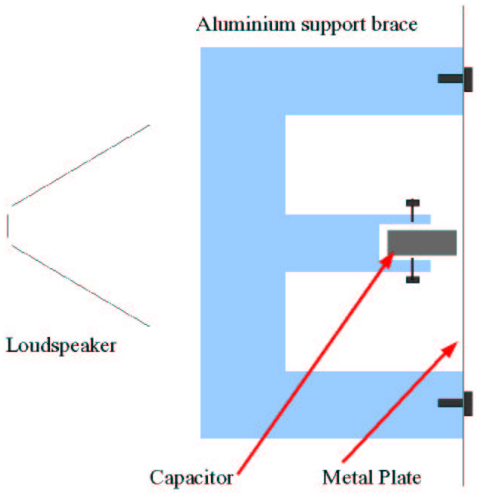

Figure 4.8: Side view diagram of the experimental setup used to hold the metal test plate and capacitive sensor in place.

ture *(a)* shows the aluminium support structure holding the capacitor sensor, with the loudspeaker behind it. In picture *(b)* plate is shown covered in white powder and screwed down to the support structure in the top and bottom corners. Picture *(c)* shows the time averaged fringe pattern produced while driving the speaker at a frequency of 406 Hz. The small white disk near the centre of this picture shows the location of the capacitor behind the plate (c.f. picture *(a)*).

The interpretation of the fringe pattern displayed in figure 4.9 *(c)* is straightforward. From equation 2.24 and figure 2.9 it is clear that a stationary region of a vibrating object displays the brightest fringe on a fringe pattern. The brightest contour in figure 4.9 *(c)* connects the two points where the plate is bolted down to the 'massive' support structure. Not only is it clearly much brighter than all other fringes, but there are darker fringes on both the inside and the outside of this contour, suggesting it is a nodal contour with movement on either side.

The central region of the plate shows a series of concentric fringes, which indicates that this part of the lid is moving in and out of the surface in an 'umbrella' mode, rather like a circular membrane in the (0,1) mode (fig. 4.4-(0,1)). The VESPI result for the amplitude at the centre of each dark fringe can be read off table 2.1. From this table, the centre of the plate is the point of greatest amplitude and, being beyond the eighth dark fringe, has a maximum displacement of between 996 and 1125 nanometres. The white

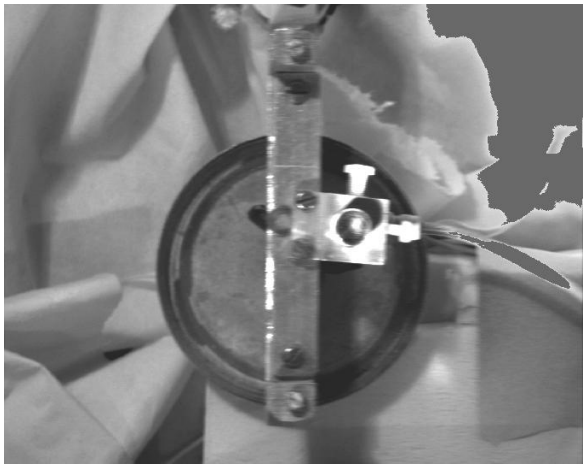

(a)

(b)

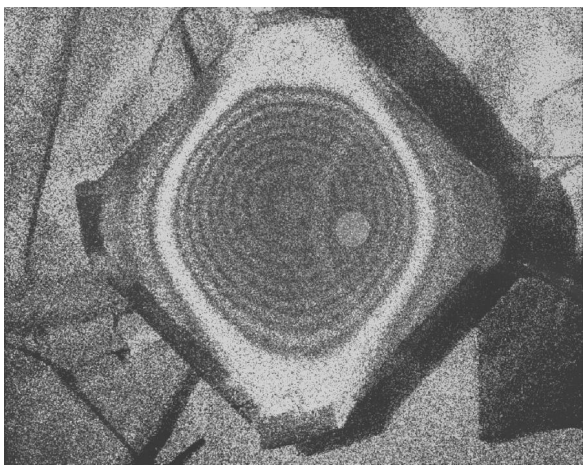

(c)

Figure 4.9: Capacitive and VESPI measurements of amplitude. *(a)* Support structure, capacitor, and loudspeaker. The aluminium bar holds the capacitive head near the centre of the image. The points where the plate is fixed are visible at the ends of the bar. *(b)* Metal plate covered with white powder and fixed in front of capacitor and loudspeaker. *(c)* VESPI image of plate driven at $406\,Hz$. A white disk has been added to show the location of the capacitor. VESPI measurement at that point, $675 \pm 65\,nm$ is compatible with capacitor measurement, $700 \pm 50\,nm$. A secondary reflection, most likely due to an internal reflection within the camera, can also be seen in figure *(c)*. The repetition of the bright stationary contour can be observed at lower intensity shifted a few centimetres to the right.

disk is behind fringes five and six, which results in a displacement of between 610 and 740 nanometres.

The voltage signal from the capaNCDT readout was measured on an oscilloscope. The signal was a sinusoidal waveform of the same frequency as the signal driving the loudspeaker, and the amplitude was $14 \pm 1mV$ peak-to-peak, which, by the calibration mentioned earlier in this section, leads to 700 ± 50 nm as the capacitor amplitude reading.

The two techniques used in the simultaneous analysis were thus seen to generate consistent results.

4.3.3 Amplitude Test with a Plastic Film

All ESPI measurements mentioned until now were carried out on objects covered in white powder. The reason for this is discussed in section 3.2.1. There are disadvantages in using this powder, however, because it can only be totally removed from certain surfaces, such as silicon wafers. It cannot be removed from a module, for example, since the spray deposits powder into small gaps between components. Thus, although ESPI is a non-destructive measuring technique, it cannot be applied to objects which must be free of powder afterwards.

A possible solution to this is to cover the object in a thin, opaque, optically rough plastic film[2] instead of the powder. There are two mechanical difficulties which must be overcome for this to work. The first is that the object shape might not allow film to be tensioned in a manner that allows it to faithfully follow the surface of the body, particularly where cavities are present. The second problem is more insidious: since the film has a mechanical integrity of its own, it is difficult to ensure that it will always follow the movement and deformation of the surface it is stretched across. Should the film motion be different to that of the object, ESPI studies the behaviour of the film rather than that of the sample.

For out-of-plane vibrational studies of flat or convex objects with simple surfaces, however, a film might be viable. To test this, metal plate of the previous section was covered in a tensioned film of white polymer. The film was held down to the sides of the lid with adhesive tape. The procedure of the previous experiment was repeated, and the capacitor reading was the same: 700 ± 50 nm. The VESPI result is shown in figure 4.10,

[2]In fact the film itself need not be opaque or rough; it can be sprayed with Developer after being affixed to the object.

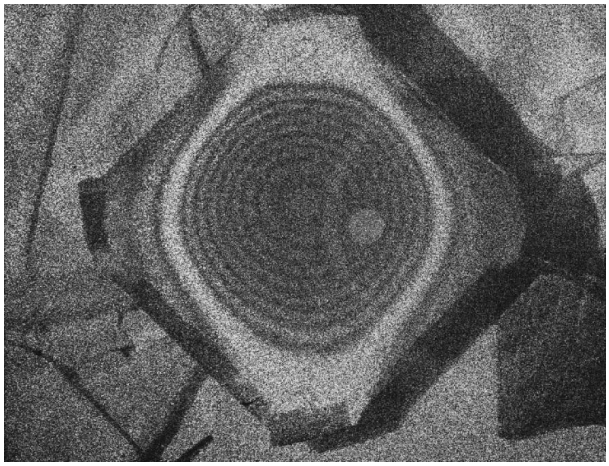

Figure 4.10: As figure 4.9 *(c)* with a white plastic film covering the plate instead of the white powder. As in figure 4.9 *(c)*, a secondary reflection of the main fringe can clearly be seen shifted to the right of the main pattern.

which is virtually indistinguishable from figure 4.9*(c)*.

4.4 Measurements of ATLAS SCT Prototypes

4.4.1 Remarks on Carbon-Fibre Reinforced Plastics

In order to meet the rigidity and low mass specifications necessary for good physics performance, the ATLAS SCT support structures must be made of materials with a high ratio of mechanical strength to radiation length. Some modern materials, in particular Carbon Fibre Reinforced Plastics (CFRP) have a far higher strength to mass ratio than most traditional lightweight materials. CFRPs can also be constructed with a Coefficient of Thermal Expansion of practically zero by combining resins of negative CTE with fibres of positive CTE.

CFRPs are made by arranging long thin filaments of graphite in a specific geometry and fixing them together with epoxy resins. The resulting fibres or tapes can then be stuck together to create layers or bulk materials. The properties of graphite determine the material's strength along the direction of the fibres, while the resin's properties determine it in the perpendicular direction. Therefore the strength of a certain CFRP material can be highly directional. The particular arrangement of the filaments in fibres, fibres in layers, and layers in a material must be tailored to the application for which the material

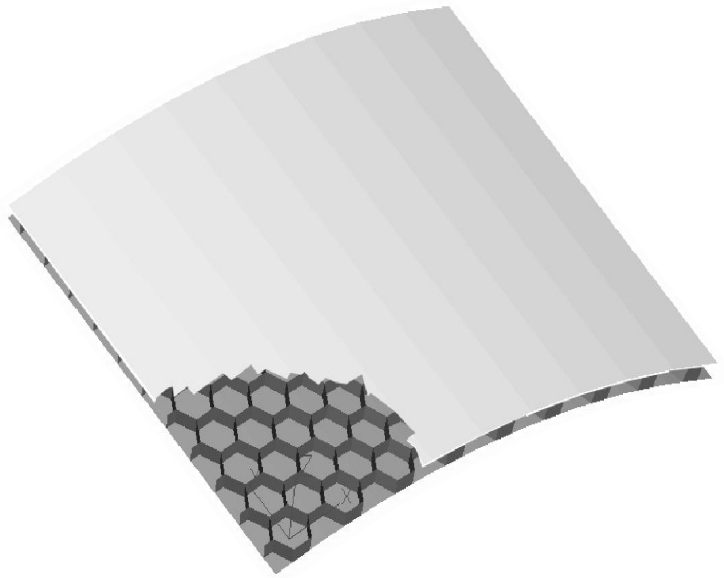

Figure 4.11: Cutaway view of a cylinder segment made of a sandwich of honeycomb placed between two layers of laminate.

is being produced. It is worth emphasising that the mechanical properties of CFRPs are, in general, highly non-isotropic.

The ATLAS SCT requires support structures to hold the modules in position: cylinders for the central region and disks for the forward regions. In addition, support struts are needed to build cages for the TRT. All of these will be constructed of CFRP. The cylinders and disks are to be built of a 'sandwich' rather than a bulk. This means that the object is made of two thin CFRP skins kept a fixed thickness apart by some lower density material. For the SCT, this 'filling' is an array of hexagonal cells (honeycomb) which are also made interwoven CFRP tapes. This is illustrated in figure 4.11. The advantage of the honeycomb sandwich is that it gives thickness and thus mechanical strength to the structure while adding very little material.

The laminates above and below the honeycomb, as well as the honeycomb itself, are constructed from several layers of fibres aligned at different angles. The overall mechanical behaviour of such a sandwich is not easy to predict even if the properties of the fibres and resins are well known. It is thus of interest to make dynamic measurements which give an insight into the properties of the structure as a whole.

4.4.2 VESPI measurement of a CFRP honeycomb square plate

A square plate of honeycomb sandwich, cut out of a prototype Front-End wheel, was studied with VESPI. The objective was to find several of the normal modes, and compare the results (shapes and frequencies) with FEA predictions, in order to gain an insight into the validity of the modelling of such materials.

The square plate was constructed with 0.2 mm thick laminates encasing a 1 cm high honeycomb with a hexagon size of 1 cm. The square has a side length of 33 cm. There are four 8 mm diameter holes in it, one near each corner, 1.5 cm away from the two closest sides. It also has a central hole of 4 mm diameter. The honeycomb is 10 mm thick, and the two facings are each 0.2 mm thick, giving a total sample thickness of 10.4 mm. The facings are made of three layers of parallel fibres, and these layers are placed at 0, +60, and -60 degrees. The honeycomb walls are not airtight. The honeycomb density is 32 kg/m^3, while the density of the whole plate is 125 kg/m^3.

4.4.2.1 FEA Modelling Procedure

A finite element computer model of the square plate was constructed using pre-defined element types available in ALGOR. The honeycomb sandwich was simulated by modelling a sandwich of two different types of element: *laminate* elements were used to build up the skins, while *solid* elements were used to simulate the honeycomb.

The laminate element type exists to simulate materials composed of aligned fibres, such as the CFRP skins used in the ATLAS SCT. Therefore it allows for a specification of laminate layers and fibres. The laminate was modelled as composed of three layers, each of a one-fibre thickness, and each made of parallel fibres lying adjacent to each other, held together by a resin. The layers were placed one on top of the other at relative angles of 0, +60, and -60 degrees. Each layer was made of fibres of 0.07 mm diameter, and a mass density of [2450 kg/m^3]. The Youngs modulus, E, for each layer was specified along the fibre direction as $E_{xx} = 280$ MPa, and across the fibre direction as $E_{yy} = 7.8$ MPa[3]. These values are based on standard Carbon Fibre laminate properties, as can be found in [Jon99].

The density used for the solid element was that supplied by the sample manufacturer

[3]Note the factor of nearly 40 between E_{xx} and E_{yy}. This shows the far greater strength in the direction where the fibres themselves determine E than in the perpendicular direction where the resin determines it.

for the honeycomb. Values for the elastic moduli of the honeycomb were not available. However, it was assumed that the primary role of the honeycomb in the dynamics of the sandwich was to constrain the relative motion of the two laminates, as well as to contribute slightly to the overall mass, rather than to affect the motion significantly due to its own elastic moduli or anisotropies. Therefore, a solid element type, which is homogeneous and isotropic, was used as the filler of the sandwich structure.

This assumed low sensitivity to the elastic moduli of the honeycomb was expressed by giving the solid element a value of $E = 10\,\mathrm{MPa}$, *i.e.* significantly lower than that of the laminates along the preferred direction. Furthermore, it was verified that the resulting analysis had a low sensitivity to variations in the solid element E.

Two squares were created out of the laminate, and one of the solid, each of the size of the CFRP sample, $33\,\mathrm{cm} \times 33\,\mathrm{cm}$. These were aligned and affixed as a sandwich. The whole sample was then broken down into a square grid of 20×20 elements.

A "mode search" was carried out on this model, using the constraint that the centre of the object should remain static. During this search, ALGOR scanned a range of frequencies and presented the resonance frequencies with a graphical output of the corresponding mode shapes as contour maps. The mode shapes and frequencies found, as well as how they compare with VESPI measurements, are discussed in section 4.4.2.3.

4.4.2.2 Experimental Description

A diagram of the experimental arrangement is shown in figure 4.12.

Mounting The plate was held in place from the central hole by a 4 cm long screw which pressed flexible washers against the plate. The plastic material was necessary to protect the fragile edge of the holes from the metal of the nut and screw head. The screw was then threaded directly into a tapped hole in a solid aluminium block which was in turn bolted onto a massive sand-filled alloy pillar.

Mechanical driving An 8 Ohm 20 Watt loudspeaker was placed 40 cm away from the plate. It was necessary to place the speaker in front of the plate rather than behind it because of the sound shadow produced by the massive support pillar. It had to be placed off the normal to the plane to not obstruct the line of sight of the camera. The speaker was connected to the ADVANCE LF signal generator (section 4.2.2).

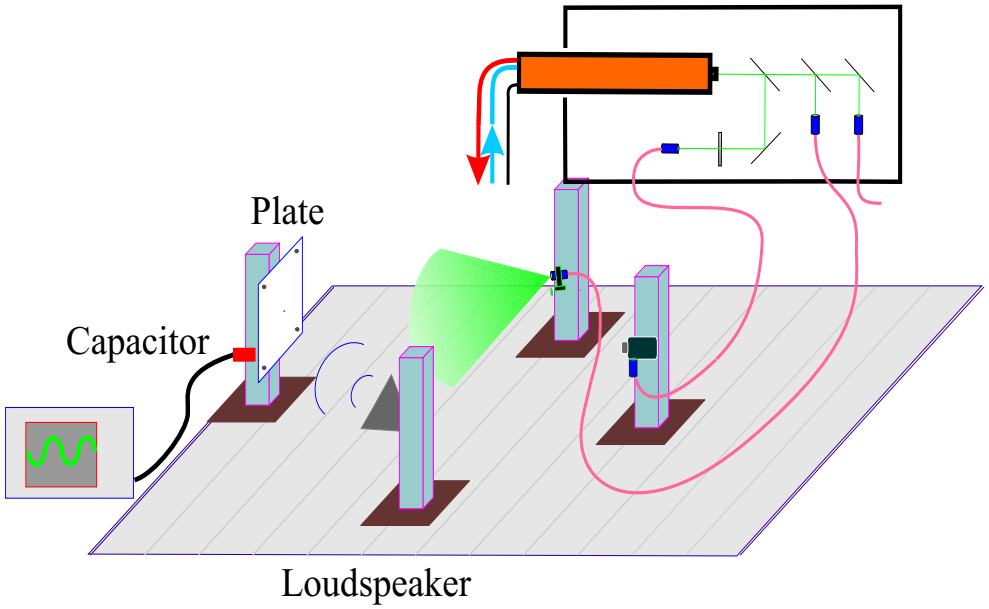

Figure 4.12: VESPI experimental setup for CFRP square plate measurements

Interferometer Geometry and Illumination Since the samples measured by VESPI in this study are either planar or thin shell objects, only out-of-plane VESPI measurements are presented in this thesis. This is because in such objects, oscillations along the plane meet the resistance of much more material than those out of it, and are thus correspondingly smaller in amplitude [Soe81]. Out-of-plane modes are more visible and easier to interpret, and are sufficient for the measurements required. The illumination was thus set up in the out-of-plane configuration (section 2.1.4). One fibre end was placed at a small angle from the surface normal, illuminating the surface, while the CCD was illuminated with a second fibre, as shown in figure 3.2. The CCD fibre delivers much less power than the object fibre, see section 3.2.2. In order to ensure good reflectivity and speckles, the plate was covered in developer powder before being measured.

Procedure The capaNCDT sensor was placed directly behind the plate, less than 1 mm away from the back, powder-free surface. The output of the NCDT readout electronics was monitored with an oscilloscope. The driving frequency was scanned from 20 Hz to 2000 Hz, once with the capacitor placed near the centre of a side, and again with the

Mode Number	5	6	7	12
FEA frequency (Hz)	429	678	1136	2235
VESPI frequency (Hz)	443	662	1063	2026

Table 4.1: Square plate frequency comparison: FEA and VESPI. The resonant frequencies for four distinctive modes which were clearly resolved by VESPI are shown. The FEA frequency was calculated based on the assumptions described in section **??**. The VESPI frequency is the driving frequency at which the resonant mode was found and imaged. Note agreement to within 10%.

capacitor near a corner[4], to ensure that a mode was not missed by placing the capacitor very close to a nodal line. These frequency scans revealed clear eigenmodes of vibration at definite frequencies. The most distinctive eigenmodes, *i.e.* those of greatest amplitude for a given driving amplitude, are shown in figure 4.14.

4.4.2.3 Results

The modes found with VESPI and their natural frequencies are compared with the corresponding FEA modes in table 4.1 and figures 4.13 and 4.14.

The lowest four modes found by ALGOR are in-plane and thus cannot be detected by the VESPI arrangement used. The fifth mode was found at 443 Hz. (figure 4.14 *a*), which should be compared with 429 Hz for the FEA model. The sixth was found at 662.2 Hz, (figure 4.14 *b*) which is close to the FEA prediction of 678 Hz (figure 4.13 *b*). Both these patterns are distinct, and the FEA and VESPI mode shapes practically identical.

A VESPI pattern very similar to the seventh mode is shown in in figure 4.14 *c)*.

Mode 12 was found at 2026.5 Hz, see picture *d)* of figure 4.14. This agrees to within 10% with the FEA frequency, 2235 Hz (*d)* of figure 4.13)

The close match between the FEA and VESPI frequencies validates the FE modelling technique applied to the non-homogeneous Carbon Fibre material. It further confirms the accuracy of the assumed input parameters.

Quality Factor The quality factor of an oscillator is given by the ratio of ω_0, the peak frequency of a resonance to γ, the frequency width of the peak (equation 4.2) [KK73].

[4]Because graphite is a reasonably good electrical conductor the motion of a CFRP surface can be accurately measured with the capaNCDT system described in section 4.3.2.

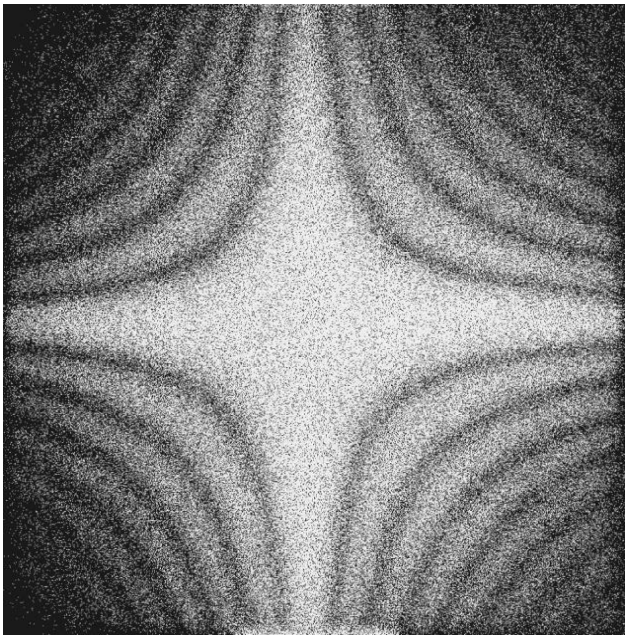

a) VESPI Mode 5, 433 Hz

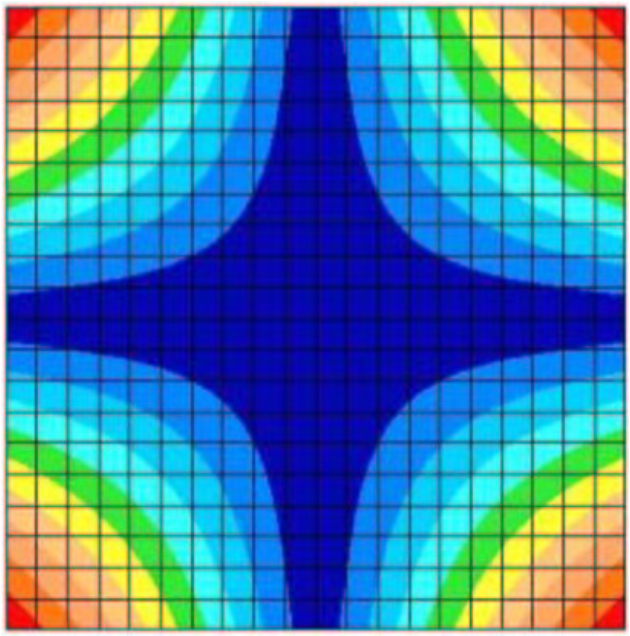

b) FEA Mode 5, 429 Hz

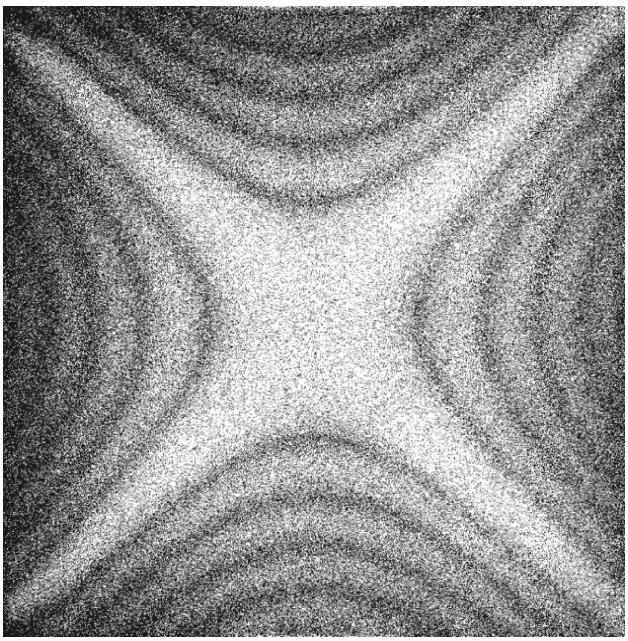

c) VESPI Mode 6, 662 Hz

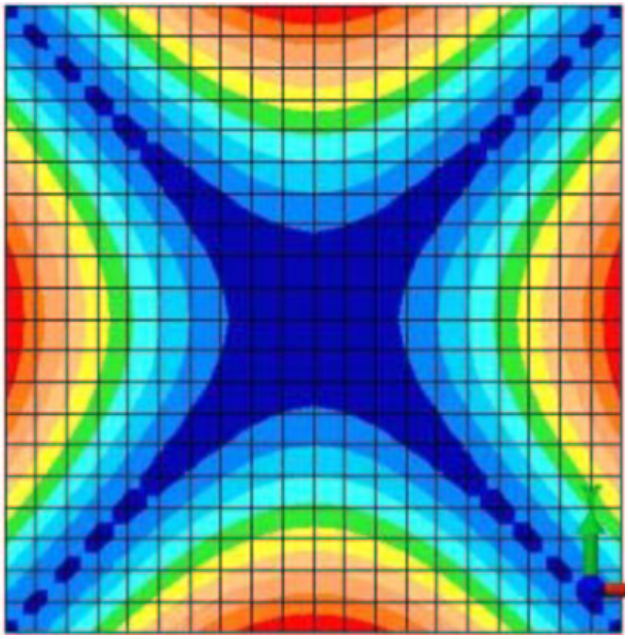

d) FEA Mode 6, 678 Hz

Figure 4.13: Comparative VESPI and FEA normal modes and frequencies for CFRP Square Plate, page 1 of 2.

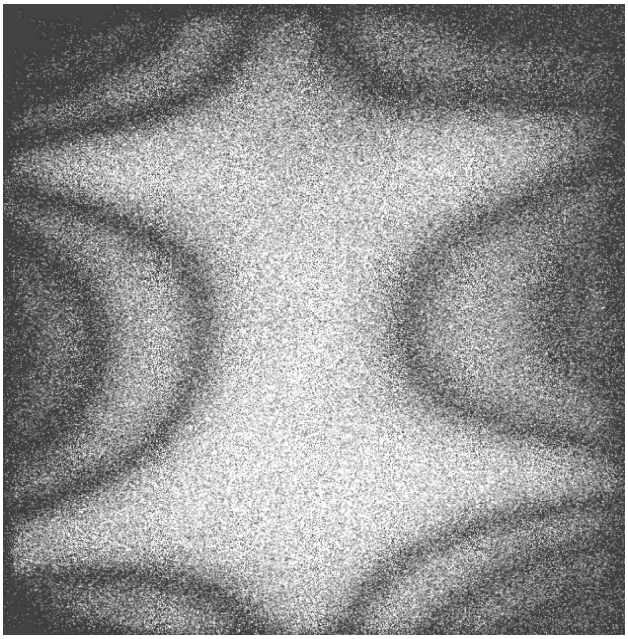

a) VESPI Mode 7, 1063 Hz

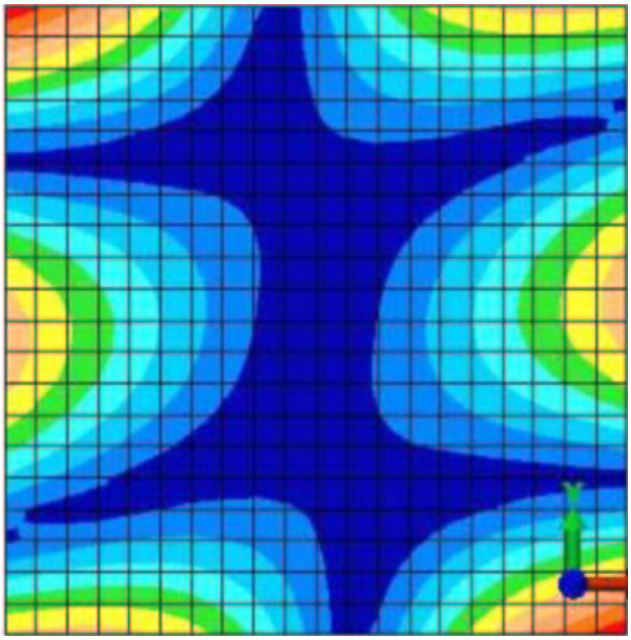

b) FEA Mode 7, 1136 Hz

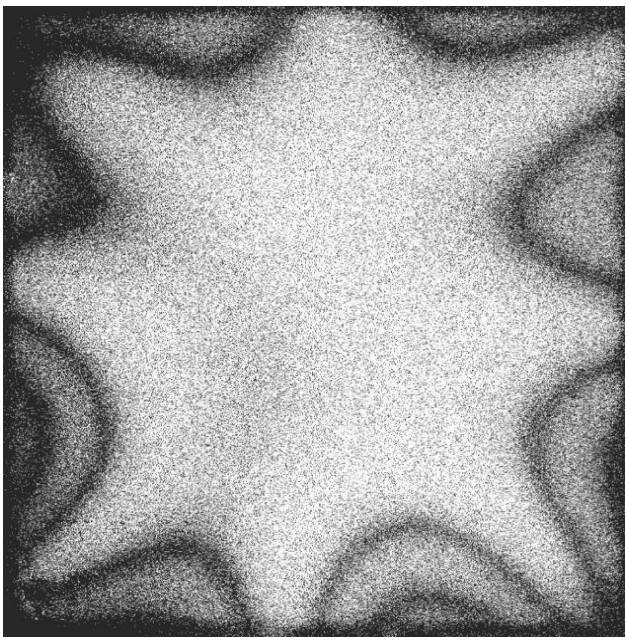

c) VESPI Mode 12, 2026 Hz

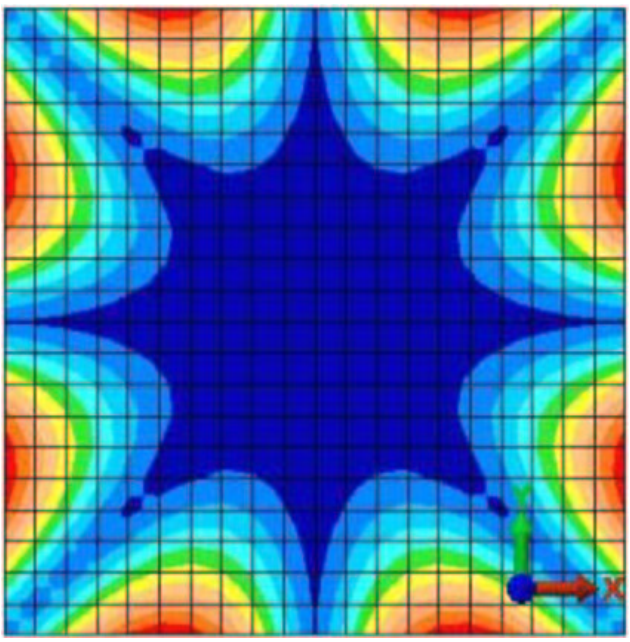

d) FEA Mode 12, 2235 Hz

Figure 4.14: Comparative VESPI and normal modes and frequencies for CFRP Square Plate, page 2 of 2.

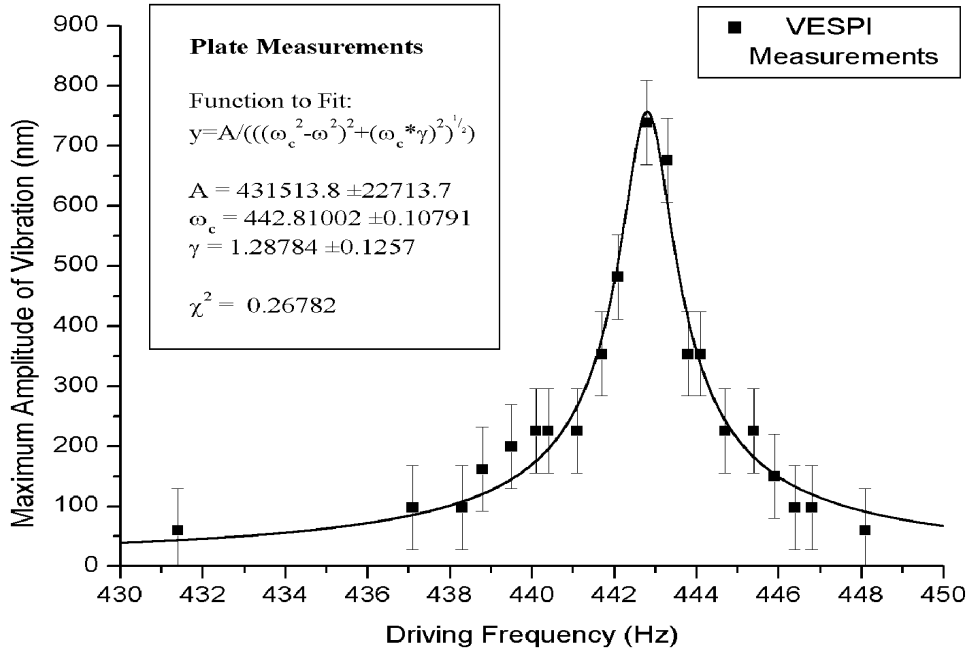

Figure 4.15: Plate Q Measurements

To determine the value of Q for the CFRP Plate, a scan of object vibration amplitudes was done around the 443 Hz resonant frequency.

$$\frac{1}{\left[\left(\omega_0^2 - \omega^2\right)^2 + \left(\gamma\omega\right)^2\right]^{1/2}} \tag{4.2}$$

The plate was driven with a constant driving voltage amplitude while the driving frequency was varied in small steps. Since the region of frequency covered is small, the coupling efficiency of the speaker to the air was assumed to be constant. The resonance curve (equation 4.2) was fitted to the data, and the resonance frequency ω_0 and the width γ were obtained from the fit.

The result of this procedure is shown in figure 4.15. The data points are amplitudes measured by counting fringes on VESPI images. Since each successive fringe shows a difference of ≈ 130 nm in amplitude, the half the minimum scale error is $\approx \pm 70$ nm. The values of the fitted parameters, as well as the χ^2 of the fit, are shown in the box inside the plot. Using $Q = \omega_0/\gamma$,

$$Q = \frac{442.8 \pm 0.1\,\mathrm{Hz}}{1.29 \pm 0.13\,\mathrm{Hz}} = 344 \pm 34 \tag{4.3}$$

Recalling the interpretation of Q as the fraction of energy stored in the oscillator dissipated per cycle, this number means that a disturbance on a CFRP structure could last for about one second, and thus must be avoided for optimal ATLAS SCT physics performance.

4.4.3 VESPI measurement of a CFRP honeycomb cylindrical shell

VESPI measurements were performed on an ATLAS Inner Detector prototype CFRP cylindrical shell. This cylinder was manufactured as a test piece for ATLAS SCT assembly work, although its radius is different from the final specified barrel radii. As before, the objectives were to find the normal modes of oscillation, and to compare measured behaviour with FEA predictions, in order to establish the validity of the method and possibly tune material property parameters. A measurement of the quality factor Q of the cylinder was also obtained.

The prototype cylinder is 1580 mm long with an outer diameter of 560 mm[5]. The overall shell thickness is 5.8 mm, with the honeycomb sandwiched between 0.2 mm skins. A solid CFRP circular flange was fixed to both ends of the cylinder in order to make it more rigid. These circular flanges have an outer radius of 560 mm, an inner radius of 500 mm, and a thickness of 2 mm. There are 4 holes of 6 mm diameter going through each flange at a radius of 520 mm, separated by 90 degrees. These holes were used to hold the cylinder and fix it to a support structure. A photograph of the cylinder (covered in white plastic film) is shown in figure 4.16.

4.4.3.1 FEA Modelling Procedure

The finite element modelling for the CFRP cylinder was similar to that of the square plate, section 4.4.2.1. As before, the sandwich was made using laminate and solid element types, although for the cylinder the solid element had smaller thickness and somewhat higher density. The initial inputs for the elastic moduli were those used for the plate, and

[5]The closest final design barrel diameter is 520 mm

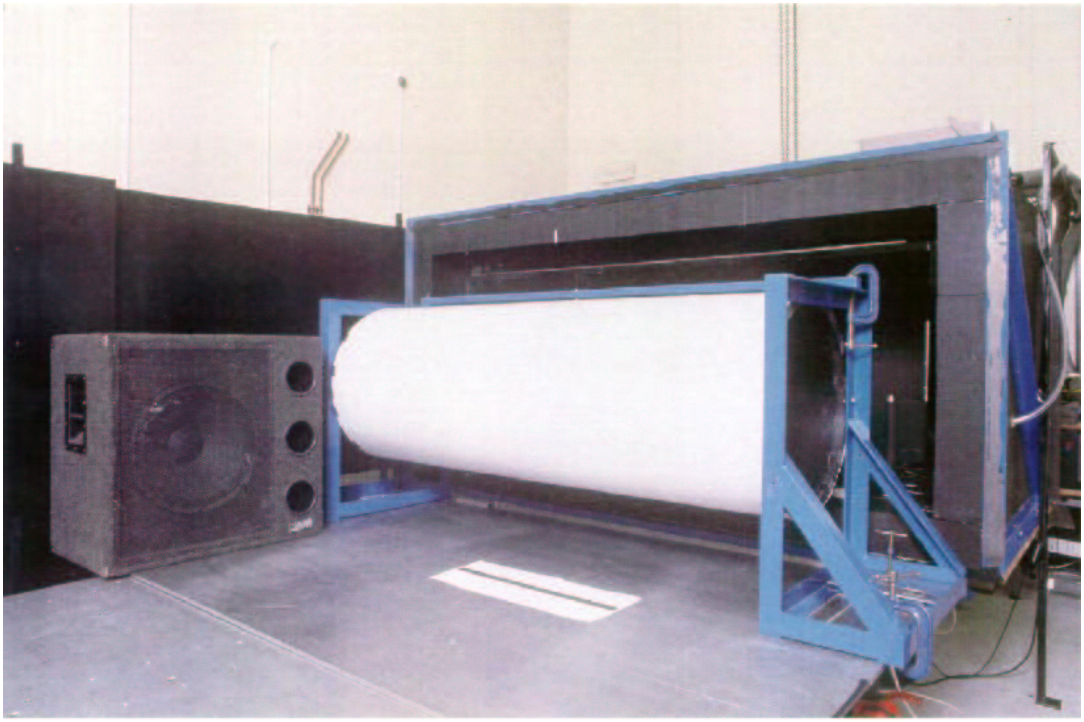

Figure 4.16: CFRP prototype cylinder for ATLAS SCT covered in white polymer film. Note the large loudspeaker to the left of the cylinder. This was used to acoustically drive the natural modes of oscillation. For comparison, the metal ruler beneath the cylinder is 50 cm long

were later tuned to match the lowest measured VESPI frequency with its corresponding FEA mode.

The cylinder was broken down into a grid of 40 azimuthal×16 axial small rectangular plates. The boundary conditions of the cylinder were simplified by assuming the stiffness of the flanges to be infinite, *i.e.* the edges of the cylinder were assumed static. This approximation seems reasonable given the much greater stiffness of the flanges compared to the honeycomb and skins.

Once more a "mode search" was carried out on this model. Some of the modes found, together with the corresponding VESPI interferograms, are shown in figures 4.19 to 4.21.

4.4.3.2 Experimental Description

Mounting A steel frame was designed and built to hold the cylinder, made of two 90 degree angle brackets kept the length of the cylinder apart by steel bars. The structure had 6mm holes which aligned with those in the flanges. The cylinder was placed within the holding structure and fixed to it with 4 mm diameter nylon screws. The screws were capped by nuts which were barely tightened, and spaces were filled with flexible rubber washers. Although this left the cylinder loose to move in the frame, it minimised transmission of vibrations from the frame to the cylinder. This vibrational decoupling was necessary to stop vibrations from being transmitted to the cylinder through the support structure, and to ensure that the motion registered by VESPI was that of the cylinder on its own, not of a combined object composed of the cylinder and the support frame.

Brightness The cylinder was covered with the same thin white plastic film tested in section 4.3.3. The film was cut to a length slightly greater than the cylinder circumference. The length was wrapped around the circumference, a little tension was applied to pull the film taut, and the two overlapping ends were stuck together with adhesive tape. The width of the cylinder not covered by the first sheet was covered by a second sheet which was cut down to size. The overlapping ring of the two sheets is visible in some of the VESPI images because the film is not completely opaque, so the region with the overlap reflects more light.

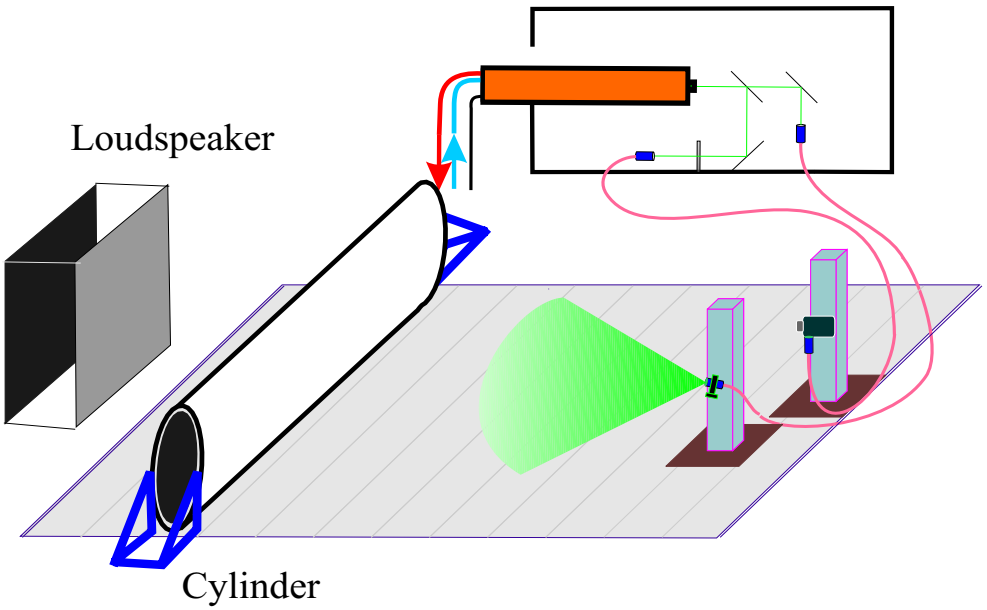

Figure 4.17: VESPI Setup for Cylinder Measurements

Interferometry Geometry and Illumination VESPI measurements on the cylinder were made more difficult than previous measurements by its large size. In order to ensure the image of such a large object fell within the CCD of the video camera, a 17 mm wide angle lens was used, and the camera was placed as far from the cylinder as the optical bench allowed (3.5 m).

The intensity of the illumination was a serious challenge given the large size of the object. The optical arrangement was modified to send the full output of the laser (minus that required for the reference beam) into a single fibre, which was then used to illuminate the object. This was still insufficient to illuminate the whole cylinder with enough intensity and uniformity to carry out a measurement.

This problem was circumvented by illuminating only half of the cylinder for each interferogram captured. The illuminating fibre end was held near the camera, pointing at one half of the cylinder, and a VESPI image was generated. The fibre holder was then rotated to illuminate the other half of the cylinder, and a second VESPI image was generated, while all other things remained equal.

The two separate VESPI images were merged after the measurements using Adobe Photoshop, an image manipulation package. The results of these manipulations show clear fringe patterns, which were used to make the measurements described in this section.

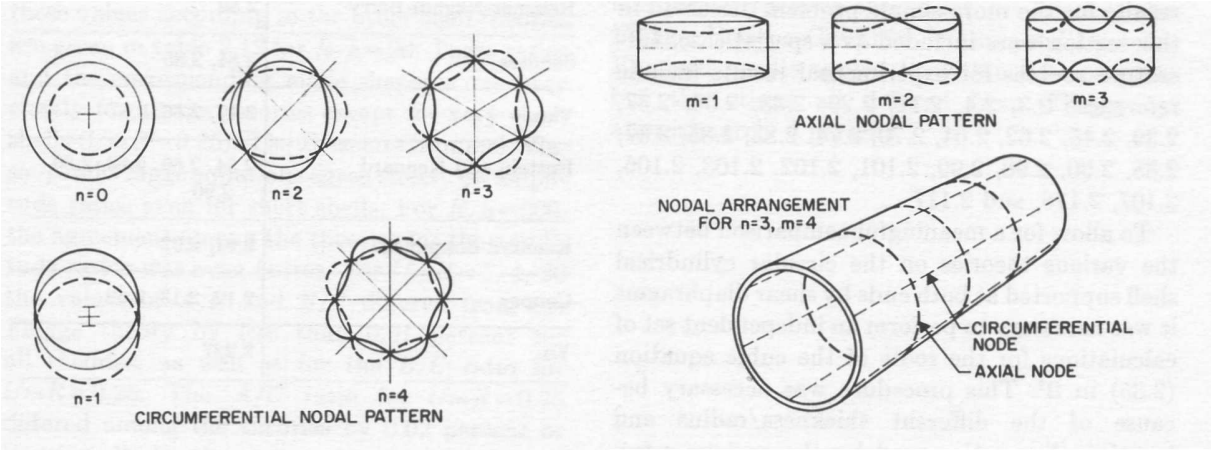

Figure 4.18: Illustration of the (n,m) modal classification of a thin cylinder based on equation 4.4: n is the azimuthal or circular number and m is the axial number. From [Lei73].

Procedure As before, the capaNCDT sensor was placed close to the cylinder, less than 1 mm away from the surface. The output of the NCDT readout electronics was monitored with an oscilloscope. The driving frequency was scanned from 20 Hz to 2000 Hz several times, with the capacitor placed at different locations, to ensure that modes were not missed by placing the capacitor very close to a nodal line.

4.4.3.3 Results and Analysis

The frequencies at which modes were found by the FEA calculation and by VESPI are shown in table 4.2: values found by VESPI are shown in red, FEA values are shown in blue. Modes can classified by the number pairs (n,m), as these describe all shapes permitted by the boundary conditions though equation 4.4 below [Soe81]. An illustration of this classification, taken from [Lei73], is in figure 4.18.

$$u = A \sin \frac{m\pi z}{L} \cos n(\theta - \phi) \tag{4.4}$$

Here u is the displacement normal to the cylinder surface, z is the axial coordinate, θ is the azimuthal coordinate, ϕ is an arbitrary phase, A is an arbitrary constant, L is the length of the cylinder, and n and m are positive integers. The interpretation of this result is aided by noting it is the product of the normal modes of a clamped bar $(B \sin m\pi z/L)$ and those of a circular ring $(C \cos n(\theta - \phi))$.

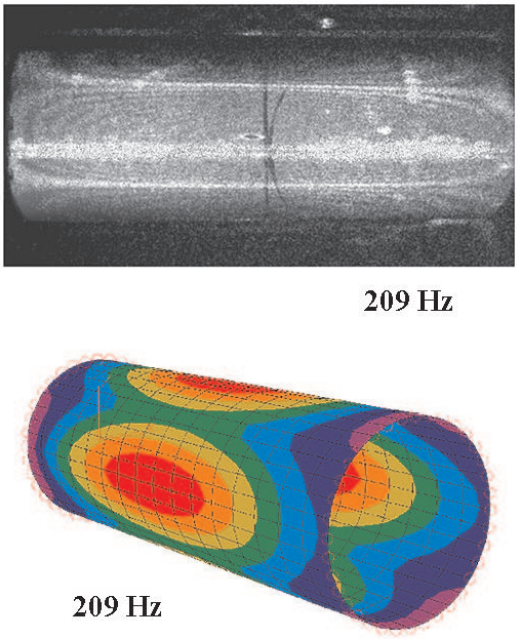

Figure 4.19: Lowest measured mode, (2,1), with four nodal lines along cylinder length and two nodal circles (one at each end).

Three of the VESPI images obtained are matched with their corresponding FEA prediction and shown in figures 4.19 to 4.21.

The FEA predictions for most modes display mode splitting into two frequencies only a few Hertz apart. This splitting is discussed later in the chapter. VESPI mode shapes were recognised and matched to the corresponding FEA modes. Figure 4.22 shows that, for the same driving amplitude, the vibration amplitude of the lower resonant frequency of the (1,1) mode is greater than that of the higher frequency. Thus, in cases where FEA predicted two frequencies to a given (n,m) and only one was measured, it was assumed the measured value corresponds to the lowest frequency predicted value.

Noise The contours of equal amplitude are clear in all three pictures. However, there are also other bright regions in the pictures. These are due do non-uniformities in the object reflectivity. The white film used is not totally opaque, letting a little light through. Therefore there are two lines on the barrel which reflect extra light: one is the vertical line to the right of centre where the two pieces of film overlap. The second is the horizontal line along which a specular barrel would reflect into the camera. Since the barrel, without

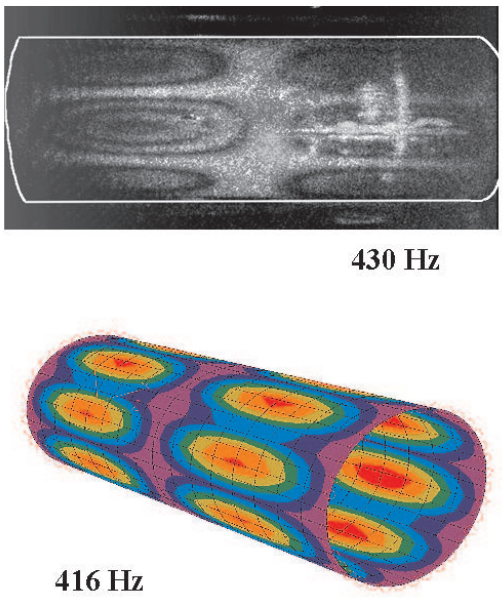

Figure 4.20: Mode (4,2), with eight nodal lines along cylinder length and three nodal circles (one at each end and one in the middle).

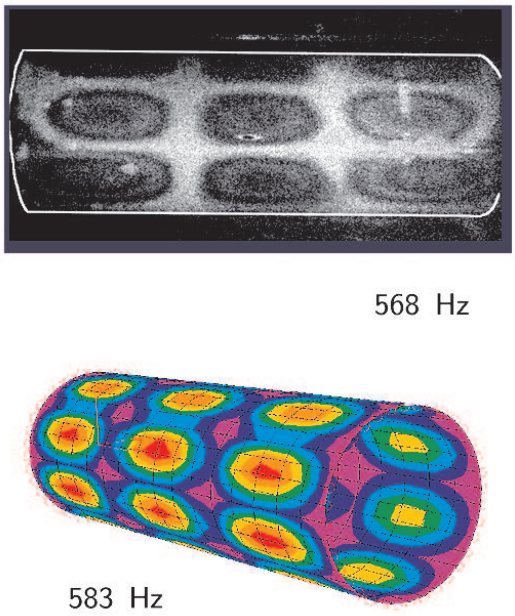

Figure 4.21: Mode (4,3), with eight nodal lines along cylinder length and four nodal circles (one at each end and two in the middle).

		Azimuthal			
	Modes	n = 2	n = 3	n = 4	n = 5
Axial	m = 1	209 217 209 215	239 232 236	371 352 354	603 628
	m = 2		419	430 405 416	
	m = 3			568 583	

Table 4.2: Measured and modelled modes of vibration for the CFRP cylinder. The numbers (n, m) refer to the number of azimuthal and axial nodes as given by the model of a thin-shelled cylinder constrained at the ends through equation 4.4 and figure 4.18. FEA calculated frequencies shown in blue and VESPI measured frequencies shown in red. Note mode splitting in both FEA and VESPI, see text.

the film, is nearly specular, most of the light that gets through the film along this line is indeed reflected into the camera. In these regions, A_0 and A_r of equation 2.26 become considerably greater than those of the rest of the image. As a consequence, slight relative differences between A_{0_N,r_N} and $A_{0,r}$ lead to absolute differences which are greater than the ESPI signal across the rest of the object. When the display thresholds are set to optimise the visibility over the bulk of the object, these over-reflective regions have above threshold brightness and thus appear white, destroying any structure ESPI might have displayed in them.

Mode splitting Both the FEA analyses and the VESPI measurements display some mode splitting, *i.e.* a separation in energy (frequency) between modes of the same modal numbers (n,m). Several characteristics of a vibrating real (*i.e.* of finite thickness) shell can give rise to this splitting, including material inhomogeneities, initial internal stress conditions, and variations from cylindricity. This in turn can be due to the action of an external cause that breaks the cylindrical symmetry, such as the action of gravity. It is also possible to generate false splitting in the FEA due to the finite size of the elements.

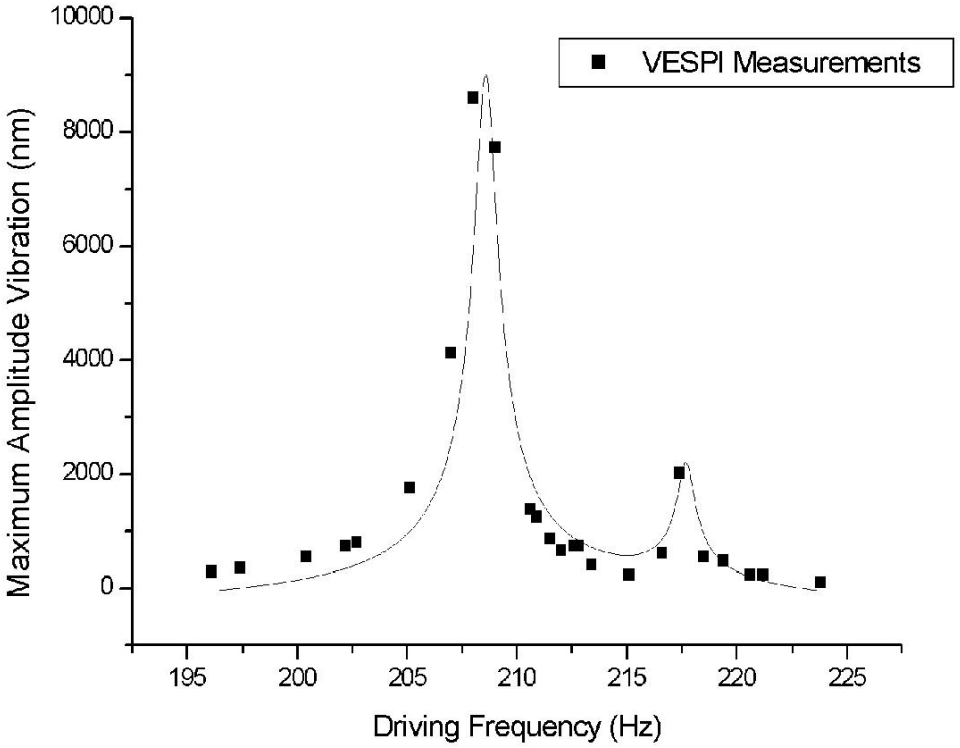

Figure 4.22: VESPI amplitude measured for lowest (2,1) mode of CFRP cylinder. Two peaks correspond to mode splitting into 209 Hz and 217 Hz, see table 4.2. Note second peak amplitude has significantly lower amplitude than first peak (approx 1/4) for equal driving amplitude.

Furthermore, the complex, three-dimensional motion possible in a cylindrical shell may result in *apparent* mode splitting, *i.e.* *different* modes of same (n,m), which are different due to axial or torsional (in-plane) motion and relative phases between the (n,m) radial motion and these other types of motion [Lei73]. Given the complexity of finding an analytical description of the full vibration spectrum of a real cylindrical shell, no attempt has been made to systematically predict the specific mode splitting shown in table 4.2.

It should be noted that both predicted and measured splitting result in relatively small frequency differences. Furthermore, splitting was only detected for the strongest (209 Hz) mode, and the relative amplitude of the higher frequency state was noticeably weaker than that of the lower frequency state. It is thus possible that other split modes predicted by the FEA were not detected due to low relative amplitude. However, the fact that several of the modes were seen at frequencies close to those predicted by the FEA calculations should still give us confidence that the model is an accurate approximation of the sample under study.

Youngs modulus and other elastic moduli Tuning the elastic moduli to match the lowest measured mode with the VESPI frequency, 209 Hz, was sufficient to ensure an agreement of within 10% between FEA and VESPI frequencies in the measured range. The need for this tuning is not surprising, since it is believed the wheel from which the square plate was cut out and the cylinder were not manufactured from identical materials. The match between FEA and VESPI without further tuning suggests that the tuned values were close to the true properties of the materials in the cylinder. Thus the match is seen as an overall validation of the combined VESPI and FEA modelling procedure put forward in 4.1.3.

Discussion of Driving Amplitude. VESPI has been used to measure the amplitude of object vibration in response to a driving force. The amplitude of this driving force, and the relationship between the driving amplitude and resulting object amplitude, are more difficult to establish. It impractical to directly measure the mechanical power driving the sample: while the electrical power dissipated by the amplifier could be measured, the distribution of this power into the acoustic spectrum is a complex one, with an unknown frequency-dependent speaker response convoluted with an unknown frequency-dependent coupling between the speaker and the surrounding air. Even if the acoustic

power spectrum delivered into the laboratory were known, the propagation of the sound through the room, and the coupling with the object under study present more unknown frequency and directional dependencies.

The best way to gain a clear understanding of the acoustic power driving the sample would be to put in place a dedicated apparatus to measure the spectrum of incoming sound close to the location of the object, both with the object present and with it absent. This in itself is a metrological project of some scale, which was not embarked upon. Fortunately, exact knowledge of the driving amplitude is not essential to analyse the data presented in this thesis. Direct comparison of object vibration amplitudes as measured by VESPI is only done in two instances. In both cases a comparison of vibration amplitude was based on the assumption that driving amplitude remains constant over small frequency range.

The first case is the comparison of the (2,1) modes at 209 Hz and the 217 Hz under identical amplifier settings and experimental geometry (figure 4.22). It seems reasonable to believe that the variation in amplitude caused by this small frequency shift is due primarily to the relative strength of the modes in the cylinder rather than to any frequency-dependent changes in the amplifier-to-speaker, speaker-to-air, or air-to-cylinder coupling efficiencies.

VESPI amplitudes are also compared when stating that this lowest 209 Hz mode is the "strongest" mode observed, *i.e.* it requires significantly less driving force to be excited to a given amplitude, not only than the degenerate 217 Hz mode, but than any other mode measured. This comparison was done with crude measures of driving power, such as the amplifier volume setting, or the sound loudness as perceived by the experimenter, however the effect is noticeable enough to be reported with some confidence[6].

Quality Factor The lowest measured (2,1) mode was studied to obtain the quality factor of the cylinder. The frequency was scanned at a set driving amplitude. The vibration amplitude obtained at each frequency is shown in figure 4.23. The figure also shows the best fit of a resonance curve 4.2 to the data. All data points save the three

[6]The 209 Hz mode was driven to oscillate at an amplitude greater than that measurable by the interferometer while the sound system was playing the monotone 209 Hz signal at a low volume. To reduce the number of fringes to a value $\leq$ 16 fringes ($\approx 2\,\mu$m) which could be readily counted, the volume of the sound system amplifier was turned down to a setting of 1 in the range of 0 to 14. The 217 Hz mode showed similar amplitude at a driving volume of nearly 2, and none of the other modes showed more than 5 fringes for a volume of 2. Many modes still showed a countable number of fringes when the room was shaking from the sound levels, at a volume setting of 3 or 4.

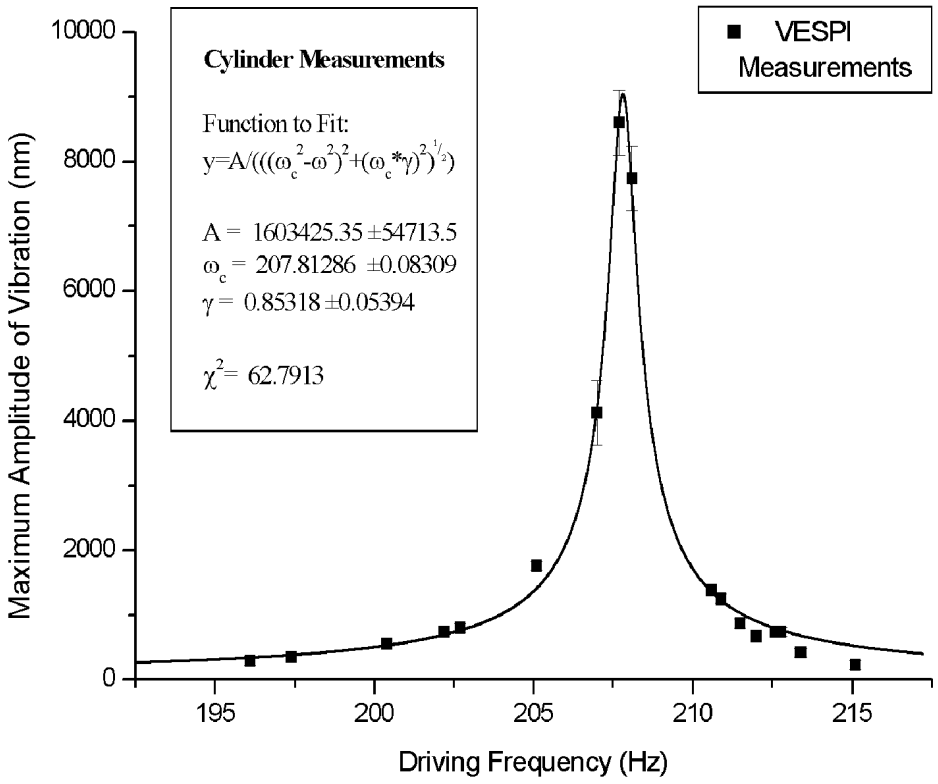

Figure 4.23: Cylinder Q Measurements

highest have errors of about the same size as the squares. The three highest points in the plot have greater error bars because they were not measured directly; only ≈16 fringes per node (≈200 nm) were resolvable with the apparatus used. To deduce the amplitude of cylinder vibrations at these frequencies, a similar frequency scan was done at a lower driving amplitude. In this second scan the central values of the peak were measured clearly, although the lower values were difficult to measure since they were of order one fringe. In the range where both were measured, the ratio between the high amplitude and low amplitude was found and used to scale the highest amplitude measurements up to the plot of the initial scan. Because this procedure introduced a greater uncertainty, these three data points have greater error bars.

The Q value deduced from this fit is

$$Q = \frac{207.8 \pm 0.08\,\mathrm{Hz}}{0.85 \pm 0.05\,\mathrm{Hz}} = 244 \pm 15 \tag{4.5}$$

As in the case of the square plate, this measurement shows that the fractional energy lost per cycle is low, allowing a perturbation to last up to ≈ 1 s.

4.4.4 Conclusions

An Electronic Speckle Pattern Interferometer has been developed to measure fundamental modes of vibration. An ATLAS SCT prototype CFRP cylinder has been studied and several natural modes of oscillations found, of with the most susceptible to excitation occurring at 209 Hz. The frequency and mode shape information is compatible with Finite Element Analysis results. The quality factor of the oscillating cylinder has been found to be 244±15, implying any excitation of a normal mode could last for a second before decaying.

The implications for the ATLAS SCT are that excitation of the normal modes of the cylinders during LHC operation must be avoided. The frequencies of the SCT barrels will be lower than those measured due to the loading of the modules. However, the value of Q is likely to remain high.

The Oxford Vibration Interferometer can thus help assess the mechanical properties of large (≈ 1 m) objects and, combined with FEA modelling, it can be used to validate assumptions on material properties of the object under study.

Chapter 5

The W width and the transverse mass method at the LHC

The LHC will produce a very large number of W bosons[1], giving ATLAS a unique opportunity to make high statistics measurements of these particles. This chapter presents some of the W's properties, as well as a technique to measure its width, Γ_W, using the transverse mass spectrum. This material is background for the following chapter, Chapter 6.

5.1 The W boson at the LHC

The W^+, W^- particles are massive, charged vector bosons whose properties were accurately predicted by the Glashow-Salam-Weinberg 'Standard' $SU(2) \times U(1)$ model of electroweak interactions. They were first observed by the UA1 and UA2 experiments at the S$p\bar{p}$S accelerator at CERN in 1983. They have been subject to precision measurements by all four LEP experiments and by the Tevatron experiments.

First published in 1968 and 1969, the GSW model incorporates the previously known properties of the electromagnetic and weak interactions - such as the long range of EM, the weak Fermi theory, and the parity-violating V minus A coupling of weak interactions - while using one of the most fruitful ideas of modern theoretical physics: invariance

[1] A total of 20 billion may be generated at the ATLAS interaction point alone during the scheduled 10 years of operation. However, high luminosity running is likely to result in extremely high event rates for many phenomena of interest, so there may be insufficient bandwidth to store Ws. Low luminosity running will generate approximately 900 million Ws, of which the more interesting events will most likely be recorded for analysis.

under symmetry transformations or gauge invariance.

It assumes a world of massless fermions and requires that interactions between them obey the SU(2) and U(1) gauge invariances of the weak and electromagnetic interactions. Four massless bosons are proposed to mediate the interactions: one each for the 3 generators of SU(2), and the single generator of U(1), labelled $W^{\pm}, W^0$ and B^0 respectively.

Mass terms are not included at the gauge requirement stage, because they would break the symmetry. Instead, mass is introduced through the interactions of an additional particle, the Higgs boson. The result is a mixing of the B^0 and the W^0 into a massless and a massive state, the photon A^0 and the Z^0. The angle that describes this mixing, θ_W, also relates the masses and the couplings of the $W^{\pm}$ and Z^0. The Fermi theory is explained as very short range by a suitably high value of M_W, and M_Z is then deduced by using $M_W/M_Z = \cos\theta_W$.

The $W^{\pm}$, which carries a unit of both isospin and charge, emerges from this as the only particle capable of changing the flavour of fermions. It does so by coupling to the left-handed fermion doublets, which are $\binom{\nu_e}{e}, \binom{\nu_\mu}{\mu}, \binom{\nu_\tau}{\tau}, \binom{u}{d'}, \binom{c}{s'}, \binom{t}{b'}$. The d' type quarks are related to the mass eigenstate d type quarks by the CKM matrix V_{ij}.

The SM prediction for the partial width of the W to the $\binom{\nu_e}{e}$ doublet can be calculated directly from GSW model and is found to be [Gro00]:

$$\Gamma(W \to e^+\nu_e) = \frac{G_F M_W^3}{6\sqrt{2\pi}} \approx 226.5 \pm 0.3 \text{ MeV} \tag{5.1}$$

The coupling of the W to all fermion doublets is the same. Thus, considering 3 leptonic channels plus two permitted quark channels ($\binom{t}{b'}$ being too massive) multiplied by a colour factor of three, the total width could be expected to be 9 times the $\binom{\nu_e}{e}$ width. In fact, after QCD corrections, the quark channel widths look like

$$\Gamma(W \to u_i d_j) = \frac{C G_F M_W^3}{6\sqrt{2\pi}} |V_{ij}|^2 \approx (707 \pm 1)|V_{ij}|^2 \text{ MeV} \tag{5.2}$$

where G_F is the Fermi constant, V_{ij} is the Cabibo-Kobayashi-Maskawa matrix element, and $C = 3\times$(QCD corrections) where the 3 is due to colour. This leads to the following Standard Model prediction of the total W decay width [Gro00]:

$$\Gamma_W^{SM}(Total) = 2.0927 \pm 0.0025\text{GeV} \tag{5.3}$$

Thus the width of the W is very precisely predicted by the Standard Model. A precise measurement is an interesting test of the model, and any discrepancy between the measured and predicted values would be a significant result. For example, a greater width could be evidence for a previously unknown decay channel.

W boson production At the LHC single Ws will be produced in collisions of u type quarks (anti quarks) with d type anti quarks (quarks). The production rate thus depends on the parton luminosities inside the protons as well as on the fermion-fermion cross section. These parton luminosities can be obtained from the Parton Distribution Functions (PDFs). This can be expressed as follows:

$$\sigma_{proton-proton} = \sum_{a,b} \int dx_1 \; dx_2 \; f_{a/p}(x_1, Q^2) \, f_{b/p}(x_2, Q^2) \; \hat{\sigma}_{ab}(\hat{s}) \tag{5.4}$$

where $f_{i/p}(x, Q^2)$ is the probability density of finding parton type i carrying momentum fraction x in a proton probed by a momentum exchange Q^2, and $\hat{\sigma}_{ab}(\hat{s})$ is the parton-parton cross section, at parton centre-of-mass energy $\hat{s}$, to produce a W of mass $\hat{s}$. The $\hat{\sigma}(\hat{s})$ contains the physics, and is of the form [BP87]

$$\hat{\sigma}_{ij}(\hat{s}) = \frac{1}{3} \frac{|V_{ij}|^2}{3\pi} \left(\frac{G_F M_W^2}{\sqrt{2}} \right)^2 \frac{\hat{s}}{(\hat{s} - M_W^2)^2 + (\hat{s}\Gamma_W / M_W)^2} \tag{5.5}$$

The cross section for $pp \to W \to e\nu, \mu\nu$ at the LHC is 30 nb, so 900 million events should be expected in the ATLAS detector during three years of running at the initial luminosity of $10^{33} cm^{-2} s^{-1}$ [ATLb]. Although precision measurements are intrinsically more difficult at hadron colliders than at electron colliders due to the lack of knowledge of the longitudinal momentum of the colliding particles, such high statistics present a unique opportunity to study the W properties. Studies have already been carried out to estimate the statistical error on the mass of the W, M_W. However, no such study has been carried out to date on the expected statistical error of the total W width, Γ_W.

5.2 Measurement techniques of the W decay width

The ratio of branching ratios Two methods exist to measure Γ_W. The method by which the LEP experiments have obtained the current best value is to extract it from the ratio of the cross sections $\sigma(e^+e^- \to W \to e\nu)$ to $\sigma(e^+e^- \to Z \to e^+e^-)$. This can

be expressed in terms of the total cross sections and the branching ratios (BR) as

$$R = \frac{\sigma_W(Total)}{\sigma_Z(Total)} \cdot \frac{BR(W \to e\nu)}{BR(Z \to e^+e^-)} = \frac{\sigma_W(Total)}{\sigma_Z(Total)} \cdot \frac{\Gamma_W(l\nu)}{\Gamma_W} \cdot \frac{\Gamma_Z}{\Gamma_Z(ll)} \tag{5.6}$$

where the quantity $\Gamma_Z/\Gamma_Z(ll)$ has been measured to high precision by the LEP collaborations. The ratio of total cross sections, σ_W/σ_Z, and $\Gamma_W(l\nu)$, the W partial width to $l\nu$, can be calculated to high precision in the Standard Model. Although this method has yielded the most precise result to date, it clearly has a crucial dependence on these Standard Model calculations being a correct description of reality.

The transverse mass method This is a direct measurement of the W width. It has been used by the CDF collaboration to make a competitive measurement of Γ_W, as is described in references [Abe95] and [Ash98]. This last work is a PhD thesis which describes the method in detail.

A $pp \to W \to e\nu, \mu\nu$ event results in two very high momentum particles, the lepton and the neutrino, being emitted back-to-back in the rest frame of the W. The neutrino propagates undetected, while the lepton is almost certainly tracked to high accuracy, provided that it travels within the acceptance range of the detector. The *transverse momentum* of a particle, defined as the component of the momentum perpendicular to the beam axis, has a magnitude given by

$$|\vec{p}_T| = |\vec{p}|\sin(\theta) \tag{5.7}$$

where θ, the polar angle, is the angle between the momentum vector and the beam axis. Since the total momentum of a proton-proton collision perpendicular to the proton beam direction is zero[2], conservation of momentum requires that the final p_T also be zero. Thus if the vector sum of the final state transverse momenta is different from zero, it is deduced that undetected particles have carried away the difference. In the case of an event where there is one lepton with very large $|\vec{p}_T^{\,l}|$, it is reasonable to associate the missing transverse momentum, $\vec{p}_T^{\,miss}$, to a single neutrino. Thus large $|\vec{p}_T^{\,l}|$ and $|\vec{p}_T^{\,miss}|$ make a clear experimental signature for a W boson decay.

The transverse mass, which is defined in section 5.3, is obtained by reconstructing a "mass" of the decaying W using the *transverse* momenta rather than the full momenta of

[2]The transverse momenta of the partons inside the proton is negligible with respect to the momenta involved in the proton-proton collision, as is any transverse motion of the proton inside the beam.

the decay products. This is an important distinction, since the longitudinal momentum - hence the full momentum - of the neutrino cannot be known in a hadron collider. The resulting quantity is useful because it has a clear endpoint determined by the W mass. Assuming there is no detector error, an event can only register an $M_T > M_W$ if the decaying W had a mass $> M_W$. Thus a study of the M_T spectrum beyond M_W reveals information on Γ_W. Furthermore, M_T is independent (to first order) of the W transverse momentum, as well as entirely independent of its longitudinal momentum.

5.3 The Transverse Mass spectrum

The *transverse mass* of a W boson which decays into a lepton and neutrino system is defined as follows:

$$M_T^2 = (|\vec{p}_T^{\,l}| + |\vec{p}_T^{\,\nu}|)^2 - (\vec{p}_T^{\,l} + \vec{p}_T^{\,\nu})^2 = 2|\vec{p}_T^{\,l}||\vec{p}_T^{\,\nu}|(1 - \cos(\Delta\phi)) \tag{5.8}$$

where $\Delta\phi$, the difference in azimuthal angle[3] between the two tracks, is the angle between $\vec{p}_T^{\,l}$ and $\vec{p}_T^{\,miss}$. The physical significance of M_T is illustrated by considering the simple example of a W created at rest in the lab frame and decaying to a lepton (e or μ) and a neutrino. In such a case, since the lepton and the neutrino are emitted back-to-back, $\cos(\Delta\phi) = -1$, and $M_T = 2|\vec{p}_T^{\,l}|$. Taking the leptons to be massless implies $|\vec{p}^{\,l}| = M_W/2$, from which it can be seen that

$$M_T = M_W \sin\theta \qquad \text{(if } W \text{ is produced at rest).} \tag{5.9}$$

If a W produced in motion is now considered, it can be shown that to first order the transverse momentum of the W has no effect on M_T [BP87]. The longitudinal W momentum cannot affect M_T, a function of purely transverse quantities.

Thus M_T is a quantity which can be measured using the lepton and missing transverse momenta, and whose maximal value is highly sensitive to M_W. A value of M_T greater than M_W depends primarily on Γ_W or detector error. A typical M_T spectrum is shown in fig 5.1.

[3] An azimuthal angel is an angle on the plane perpendicular to the beam axis.

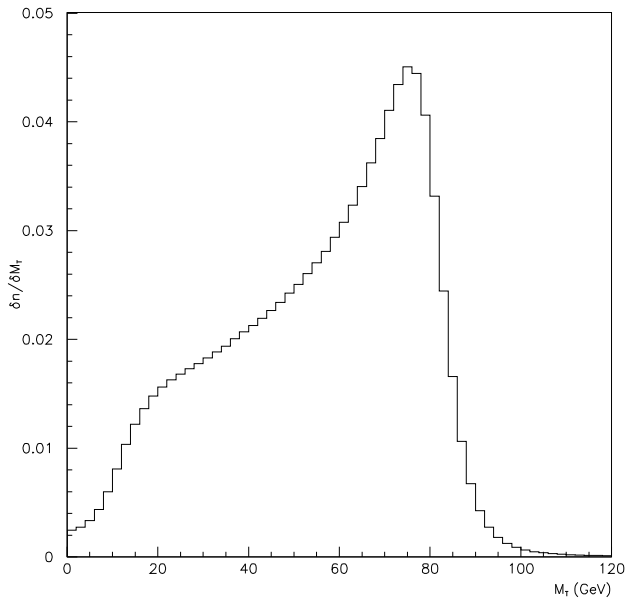

Figure 5.1: A typical M_T spectrum.

Jacobian peak The angular distribution of the decay products of the W clearly determines the M_T distribution. However, to convert a distribution in θ to a distribution in M_T, a change of variable from θ to $\sin\theta$ must be performed. The Jacobian of this transformation is of the form $1/\sqrt{1-\sin^2\theta}$. The singularity at $\theta = 90°$, which is the most distinctive feature of the M_T distribution (see fig 5.2), is called the *Jacobian peak*.

Far Tail As has been mentioned before, the M_T tail ($M_T > M_W = 81\,\text{GeV}$) depends on the value of Γ_W and detector effects. The effect of the detector can be seen in figure 5.3, namely that the measured peak is less sharp than the real peak, and the measured M_T tail, particularly close to the peak, has a high content of hits with a real M_T of, or slightly below, the W mass.

The detector error generally displays a gaussian distribution, smearing a sample of events with a true value of X into a detected gaussian distribution of values centred on X with some standard deviation σ. The natural tail of the M_T spectrum, on the other hand, which is due to the width of the decaying particle, follows a lorentzian curve, as described in section 5.1. Figure 5.4 shows a lorentzian and gaussian, both normalised. It can be seen that the lorentzian has higher far tails. Therefore, while the low end of the

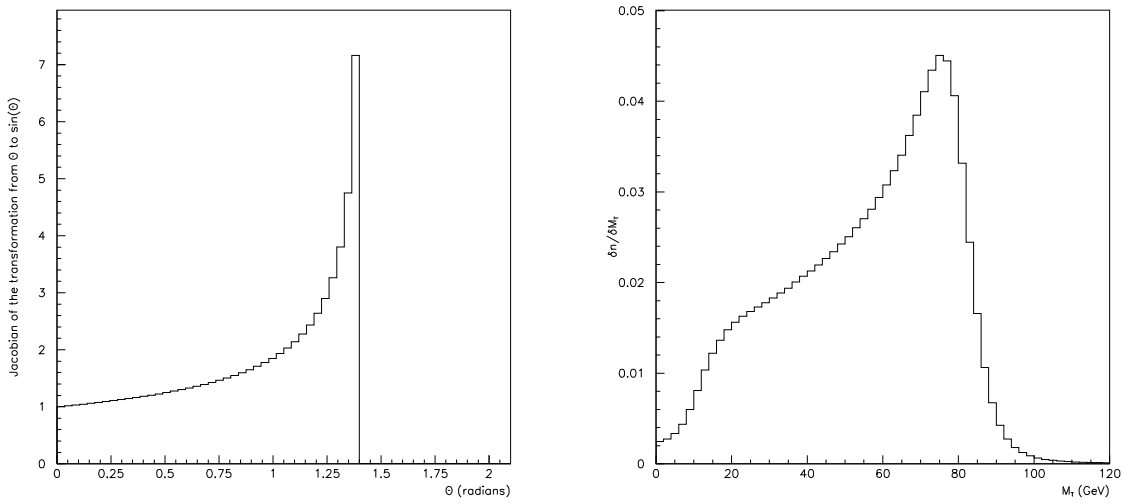

Figure 5.2: Jacobian peak. Jacobian of transformation from θ to $\sin\theta$ and the M_T spectrum of $W \to e\nu$ in ATLAS, normalised to one, with no cuts. Low M_T behaviour differs mainly due to parton luminosity, high M_T tail due to detector smearing and W width.

M_T tail may be swamped by the detector effects, it should be possible to find a section of the detected M_T spectrum populated primarily by Ws with genuinely high mass by looking at high enough values of M_T.

This focus on the high tail unfortunately implies low efficiency: less than one in every 300 $W \to l\nu$ events detected is usable for this measurement.

5.4 Γ_W from the M_T tail

5.4.1 Introduction

The far M_T tail contains information regarding the value of Γ_W primarily in its height relative to the M_T peak, as well as in its shape. This is can be seen in figure 5.5, which shows two M_T spectra, generated with different widths, and normalised to unit area under the curve. The relative height of the tail to the peak is in effect a relative W cross-section measurement: the high Γ_W samples will display a higher overall tail compared to the peak.

The value of the Γ_W parameter can be extracted from the M_T tail through a bin-by-

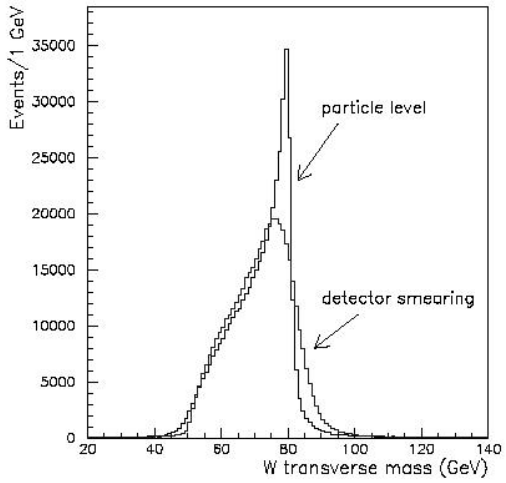

Figure 5.3: Detector smearing of the particle-level Jacobian peak results.

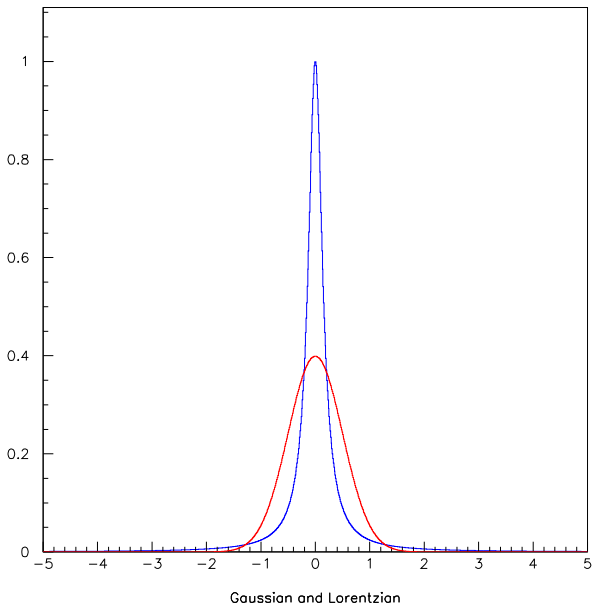

Figure 5.4: Normalised gaussian and lorentzian. Note that in the tails, Gaussian (red) falls off far more rapidly, and Lorentzian dominates.

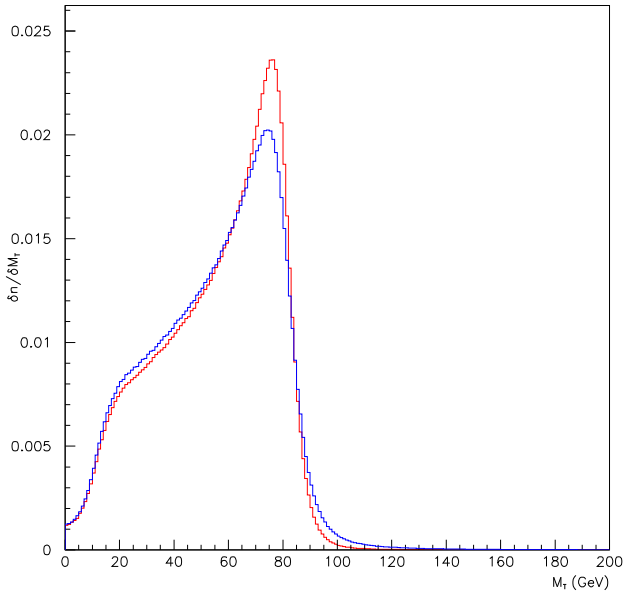

Figure 5.5: Variation of Γ_W: note how the greater Γ_W has a lower peak but higher tail

bin comparison with a set of adequately normalised spectra or templates generated with a range of known values of Γ_W. This bin-by-bin approach has the advantage that it can also use any information contained in the shape of the tail.

5.4.2 The binned log likelihood fit

The binned log likelihood method can be used to find the value of the parameter Γ_W from the high M_T tail obtained by the ATLAS detector. In order to extract the most likely value of Γ_W to account for the measured data, this method requires a set of pre-existing M_T spectra, or *templates*. Each of these templates is generated using a slightly different value of Γ_W , with the whole set spanning the expected value. The collected data set is compared to each of these templates, one by one, and the *likelihood*[4] that the collected M_T tail is a daughter distribution of that template is calculated. Once this has been

[4]The likelihood of a set of data points x_i with probability $y(x)$ is the product of the probabilities: $\mathcal{L} = \Pi\, y_i$. In the case of histograms, the *binned* likelihood is computed in a similar manner, taking the x_i to be the number of entries in bin i, and the probability y to be the probability that bin i contains x entries. For further explanation, see Appendix B.

done for all the templates, a simple polynomial can be fit to the data points of likelihood $\mathcal{L}$ *vs* Γ_W and the Γ_W for which this *likelihood function* is maximal will be the result:

$$\mathcal{L}({\Gamma_W}^{result}) = \mathcal{L}^{max}. \tag{5.10}$$

The likelihood method is described in more detail in Appendix B. There it is shown that the error on this measurement is such that

$$ln(\mathcal{L}({\Gamma_W}^{result} \pm \sigma)) = ln(\mathcal{L}^{max}) - 1/2. \tag{5.11}$$

Because of this, it is the log-likelihood function $ln(\mathcal{L}(\Gamma_W))$ that is fitted to the data points. Γ_W is determined by finding the minimum of the -2$ln(\mathcal{L})$ function. The errors on the measurement are given by adding 1 (one) to the value of -2$ln(\mathcal{L}^{min})$.

5.4.3 Template requirements

Ideally, any spectra used as templates for a likelihood fit should be true distributions, *i.e.* they should be calculated analytically or, if generated by Monte Carlo simulations, they should have infinitely high statistics. In practice they can be obtained through MC generation provided the statistics in the templates are significantly higher than the data set to ensure that, in each bin, the statistical fluctuations of the templates around the true Poisson expected value are small compared to the difference between this expected value and the value measured by the experiment. The effect of limited template statistics on the accuracy of the fit can be quantified, and it is assumed that for the analysis performed on actual ATLAS data after 3 years of low-luminosity LHC running, enough CPU time will be available to generate templates which make this effect significantly smaller than the accuracy limit imposed by the event statistics collected.

For the study in Chapter 6, templates have been generated with 10 times higher statistics than the simulated data set. This level of statistics was selected as a compromise between CPU time needed, and the accuracy required to make the study meaningful. A more detailed discussion, as well as an estimation of the template statistics required to reduce the template statistics error to an acceptable level, is presented in sections 6.3.4 and 6.3.5.

Once they are generated, the templates must be normalised and then scaled to the data set under study. The normalisation uses the total area under the M_T curve, which

requires knowledge of the peak region and low M_T region of the spectra as well as the high tail. The normalisation process is discussed in more detail in section 6.3.2.

The greatest challenge in generating templates to analyse real data collected by ATLAS will be the accuracy of the detector simulation. Whether the understanding of the detector will be good enough to bring systematic error down to a level comparable with the statistical uncertainties that can be achieved is difficult to determine before the detector is operational. However, for the study presented in Chapter 6, both the templates and the data have been generated with the same particle and detector level simulation packages, so this has not been a problem. A further discussion of possible sources of systematic error is presented in section 5.6.

5.5 Data rates and triggering

The LHC experiments will face a serious technical challenge in handling all the information registered by the detectors. In spite of a proton-proton event rate[5] of about 800 MHz (40 MHz bunch crossing rate $\times \sim 20$ collisions per bunch crossing), ATLAS will only save events at a rate of 100 Hz. This reduction factor of 8 million implies that "only" $\sim 10^9$ events will be stored per year.

The current study assumes that all W events generated during the low luminosity stage will be available for offline analysis. However, as is mentioned in section 5.1, $\sim 3{\times}10^8$ Ws are expected during one year of low luminosity running. It is not realistic to assign 1/3 of the ATLAS bandwidth to a precision study of Electroweak physics.

As is stated in section 5.3, only 1 in 300 Ws will have a high enough M_T to contribute to the tail of the spectrum. Thus the requirement is reduced to $\sim 1/1000$ of the available bandwidth.

A drawback of not registering events below the far tail is that the height of the M_T peak will be lost. This is necessary for normalisation and scaling, as is discussed in section 6.3.2. Comparable statistics in the peak and far tail regions are achieved by registering similar numbers of Ws across the whole spectrum, thus doubling the bandwidth requirement to $\sim 2/1000$ of the total.

[5]This is for high luminosity, low luminosity event rate will be a factor of 10 smaller.

5.6 Sources of systematic error

The purpose of the study presented in Chapter 6 is to estimate the *statistical* accuracy with which ATLAS will be able to measure Γ_W. Therefore, the systematic errors will not be discussed in detail. What follows is simply a list of the main sources of systematic error with a few comments.

5.6.1 Backgrounds

$\mathbf{Z \to e^+e^-, \mu^+\mu^-}$ The Z signature is usually very clear, making it unlikely for a Z to be mistaken for a W. However, if the calorimetry has regions of poor performance, some fraction of the produced Z will "lose" a lepton into a region where the energy may be undermeasured. This could result in an apparent event with only one high p_T lepton and missing p_T. In particular, the CDF detector suffered from such a problem in previous measurements, see [Abe95] and [Ash98]. Once ATLAS is operational, the calorimetry will be thoroughly studied. The response to Z events should eventually be well understood. Thus, any effect of this type should be simple to model and include in analyses. As the fraction of Zs involved is expected to be small, this should result a manageable source of error.

QCD This category includes many types of events that could result in a high p_T lepton and missing p_T. These include neutral pions, charged pions, and b and c mesons with a plethora of decay channels. An effective cut can be carried out by requiring W candidates to present isolated high p_T tracks, since QCD events generally result in multiple tracks. However, some QCD decay channels will mimic the W signature a small fraction of the time. It is believed, based on the CDF experience, that these will be manageable, but this remains to be seen in the high interaction rate of the LHC.

$\mathbf{W \to \tau\nu \to e\nu\nu\nu, \mu\nu\nu\nu}$ This irreducible background can in principle be modelled as accurately as the e and μ channels. Its effect should therefore be accurately accounted for in the production of templates and subsequent analysis, leading to a manageable systematic error. It is worth pointing out that the τ's decay products (which include the e, μ and νs) will be collimated in the τ's direction, distributing its momentum among them. Therefore, the e, μ will have a reduced p_T, and the event will thus be recorded with a reduced missing p_T, making it unlikely for a given τ event to register in the far

M_T tail, even if the initial $W \to \tau\nu$ was a high M_T event.

A general point on the backgrounds to this analysis is that they all diminish considerably with M_T in the far tail. Therefore, when a full analysis is done with ATLAS data, the cutoff at which the far M_T tail starts will be determined based on the trade off between systematics and statistics.

5.6.2 Parton Distribution Functions

The PDFs describe the momentum distributions of the partons in the hadrons. Therefore, they affect the shape of the W cross-section, and thus influence the M_T spectrum and the measured value of Γ_W.

A simple way to gauge this effect of the PDFs is to generate different sets of templates using a range of PDFs. In the study that follows, an equivalent but opposite test is carried out by generating templates with one PDF set and data sets with several PDFs, including that used for the templates, and obtaining a value of Γ_W with each of these data sets. The variation of Γ_W with a sample of PDFs is presented in section 6.3.3.

5.6.3 ATLAS detector resolution

The ATLAS detector subsystems that will primarily be used to determine the M_T spectrum are the Inner Detector, the Electromagnetic Calorimetry, and the Muon Chambers. The Hadronic Calorimetry will also play a role, since its input will be necessary to determine the missing energy and momentum, as well as the W recoil.

Of these, the greatest source of error is likely to be the EM calorimetry. The energy resolution is expected to be $\sim 10\%/\sqrt{E} \oplus 0.7\%$ [ATLa], which implies that a 50 GeV electron may be mismeasured by 1 GeV. This could certainly affect an M_T measurement by a GeV or more. Given that it is the relative height of the far tail to the peak that is used to determine the fit of the data to the template, this discrepancy does not seem too large. However, the real challenge will be in modelling the detector behaviour correctly to incorporate this effect and convert it into a manageable systematic error.

Chapter 6

The ATLAS statistical resolution on Γ_W using the M_T method

This chapter presents a study of the statistical error with which the ATLAS detector can be expected to measure the width of the W boson, Γ_W, by means of a binned likelihood fit of Γ_W to the high tail of the W transverse mass distribution (M_T). The effects of using templates based on finite statistics, as well as the possible systematic effects of the Parton Distribution Functions on the error on Γ_W, are also studied.

6.1 Scope of this study

This study aims to measure the statistical accuracy with which ATLAS can measure the width of the W boson, Γ_W, during the first three years of LHC operation by means of a binned log-likelihood fit of the W boson M_T spectrum (see section 5.4.2) . About 900 million $pp \rightarrow W \rightarrow e\nu, \mu\nu, \tau\nu$ events are expected in the ATLAS detector during this period with LHC running at the initial luminosity of $10^{33} cm^{-2} s^{-1}$ [ATLa]. Such high statistics present a unique opportunity to study the W properties. Studies have already been carried out to estimate the statistical error on the mass of the W, M_W [ATLb]. However, no such study has been carried out to date on the expected statistical error of the total W width, Γ_W.

The log-likelihood ($\ln(\mathcal{L})$) method is used to find the value of Γ_W that best fits the generated data of the transverse mass M_T. This method requires a set of high-statistics templates which must be created with a range of values of the parameter to fit, see section

5.4. Thus the study entails the generation of several (≈ 10) data sets with at least 10 times higher statistics than the data set to be studied. These high-statistics event sets are used as the templates against which the data samples are to be compared, bin by bin, to obtain the overall likelihood that each template is a true distribution of which the data sample is a subset.

The main steps carried out in this study are presented in the following list:

- Identification of necessary hardware and software tools to simulate the LHC W events and record the M_T spectra
- Determination of necessary generation statistics and analysis cuts
- Creation and/or modification of code to ensure the different packages implement the necessary physics criteria and cuts, as well as to ensure they have compatible inputs and outputs
- Verification that code is working and producing sensible physics results (testing and debugging)
- Generation of templates with a representative PDF over a range of likely Γ_W values, followed by normalisation of each template generated
- Generation of the primary data set with the same PDF as templates and a typical low-lumi number of events
- Generation of data sets with a range of different PDFs and similar number of events as the primary data set
- Execution of binned log-likelihood fits of the datasets on the normalised templates
- Implementation of any additional simulations and fits deemed necessary to study the effects of finite template statistics

6.2 W Event Generation

6.2.1 Computer Hardware and Software Tools

Computer Hardware: The bulk of the generation work was carried out on the RAL (Rutherford Appleton Laboratory) CSF (Central Simulation Facility). This facility is

a farm of over 150 fast (>1 GHz) Intel processors running Linux on which users can submit batch jobs. The actual coding and compilation to generate scripts, data cards, and executables was carried out on the Oxford Particle Physics Linux servers, as was the post-generation analysis. For a discussion of the total computing power used, see section 6.3.5.

Software: The main elements of the software used are the event generator and the detector simulator. The event generator used also calls an external library of PDFs. Finally, a histogram analysis package was used to view and analyse the data.

The $pp \rightarrow W$ events were generated using PYTHIA 6.136 [Sjö94]. This is a large, well structured Monte Carlo event generator that simulates whole particle physics events, from colliding beams to final state, at the particle level. PYTHIA has built-in options for initial conditions, so LHC conditions can be easily selected. The program then models the incoming beams, with default or selected PDFs, selects initial partons, builds up initial showers, models the hard parton-parton process, models all outgoing partons, fragmentations, and decays from the hard process, models any final state showers and beam remnant effects, and writes all final state particles in a common block, using a standard event format which can be read by many other particle physics programs.

PYTHIA has full capability to model the W events required for this study. However, it does not envisage the possibility of a direct modification of Γ_W, since it is calculated by from other parameters, namely the Fermi constant G_F, The W mass M_W, and the CKM matrix elements, see section 5.1. Thus, the PYWIDT routine was modified to introduce a variable to multiply with the calculated Γ_W to yield the new Γ_W. Figure 6.1 shows the effect of the change in the new parameter WGAM. Note that WGAM=1 does indeed yield the same curve as the unmodified PYTHIA.

ATLFAST 2.50 [RWFP98], a fast detector simulator, was used to introduce the detector effects on the reconstructed particles. The ATLFAST input is the output of PYTHIA, the final-state particles of each event. It passes these particles through simulated ATLAS detector systems to obtain the signal that ATLAS would reconstruct for each event. This program is not a full detector simulation, but rather a tool written for quick implementation of high statistics ATLAS studies. Full detector simulation is too time consuming for many of the feasibility, background, and other high-statistics studies carried out by

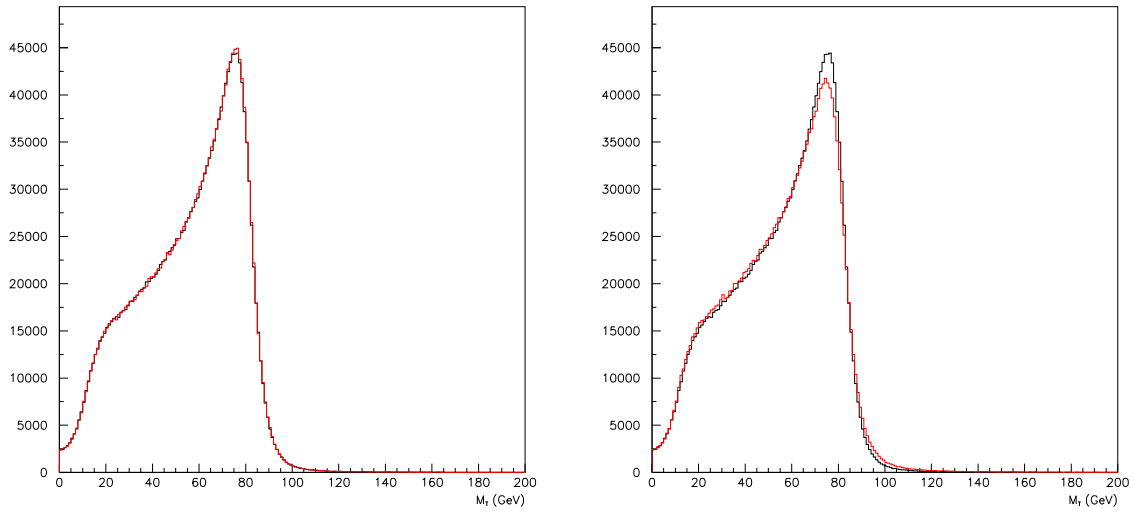

Figure 6.1: The first plot superimposes an M_T spectrum generated with the original PYWIDT routine over a similar spectrum generated with a modified PYWIDT routine with the value of the W width parameter = 1. The second plot superimposes the original PYWIDT sample with an M_T spectrum generated with the W width set to 2. Note how the modification only has effect if the parameter is $\neq 1$.

the collaboration, while particle-level studies have often proved too optimistic. ATLFAST does not contain full detector geometry information, but only key features such as the η coverage for precision physics and calorimetry, size of barrel/end cap transition region for the electromagnetic calorimeter, and the granularity of the calorimeters. It attempts to reproduce as well as possible the expected ATLAS mass resolution for important physics signals, as obtained from full simulation. It parametrises electron energy resolution, photon energy resolution, and muon energy resolution. In particular, the electron energy is parametrised as

$$\frac{\delta E_e}{E_e} = \frac{0.12}{\sqrt{E_e}} \oplus \frac{0.245}{E_e} \oplus 0.007 \tag{6.1}$$

which is consistent with the detector resolution stated by the ATLAS TDR and quoted in section 5.6.3. The output of ATLFAST is a PAW Ntuple, which contains the events as detected by ATLAS.

PDFLIB 8.04 [Plo00] is a package which compiles all the available parton density func-

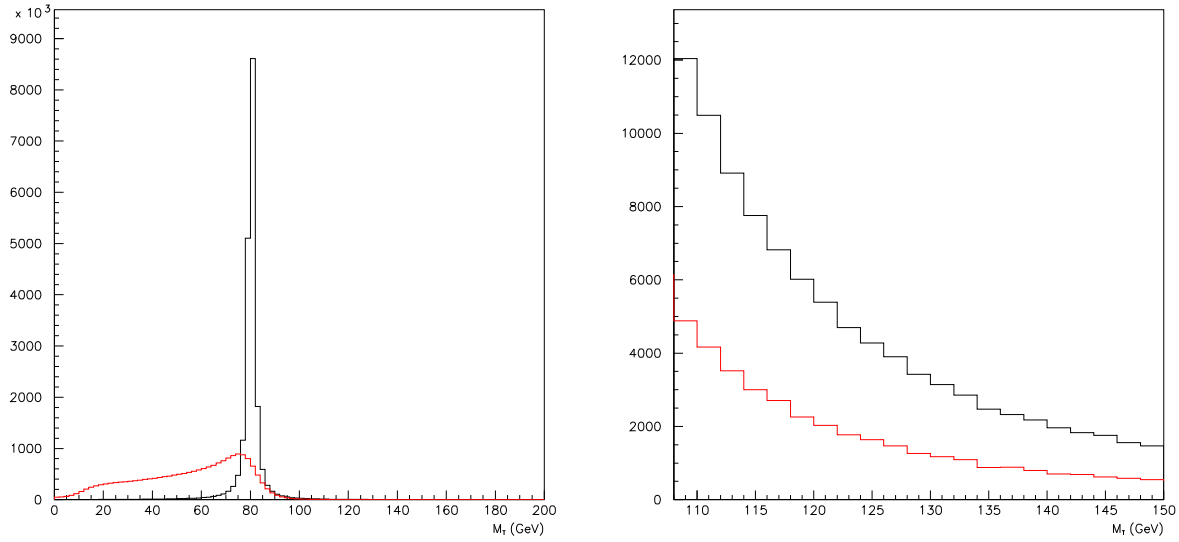

Figure 6.2: The Jacobian shape of the M_T spectrum migrates most W events with high W mass to lower values of M_T. The black line shows the M_W spectrum for a given distribution of Ws, the red line is the corresponding M_T spectrum for the same sample.

tions of the proton, pion, and photon in a consistent manner, so they can be used by other packages. This package was called by PYTHIA to make use of the selected PDFs.

PAW 2.12 [Bru00], a histogram analysis and visualisation program used extensively in particle physics, was used to analyse the histograms, extract the measurements, and produce the plots.

6.2.2 Cuts for W Event Generation

The majority of the Ws produced in the LHC will not be of use for a study which targets the tail of the M_T spectrum. The W must decay semi-leptonically for the M_T to be defined in the first place. Furthermore, most particles will fall to the left of the M_T peak due to the Jacobian shape of the spectrum. Fig. 6.2 compares a typical M_T spectrum of $W \to e\nu$ events generated in PYTHIA and 'smeared' with ATLFAST with the W mass M_W of the same data set.

In order to save time by not generating W events that would not pass the analysis cuts, two generation-level cuts were imposed. The first was a cut by decay channel,

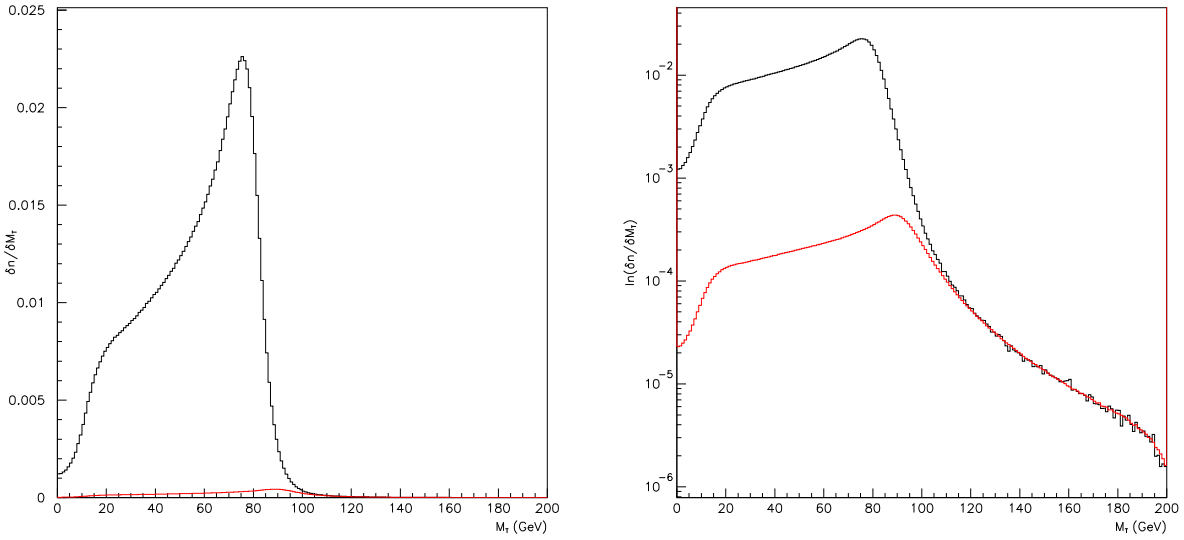

Figure 6.3: A generation cut on W mass saves CPU time, leaving high M_T tail unaltered (black line generated without the cut, red line generated with the cut, shown on linear and logarithmic scales.)

which only generated Ws decaying to e, μ, and τ. This avoided generating the more time-consuming final states of the hadronic decays.

The second cut was done by W mass. Only Ws with a mass greater than 90 GeV were generated, in order for the generated events to have a much greater likelihood of being in the high M_T tail. Fig. 6.3 compares two M_T spectra, one created without this 90 GeV mass cutoff and one with the cutoff. The plot shows how the events generated with the cut can produce comparable high M_T tails for a fraction of the particles generated.

6.2.3 Data Generated

The data sets generated for this study can be classified as follows:

- **Templates:** 9 sets of 7.5×10^8 $W \rightarrow l\nu$ ($e\nu$, $\mu\nu$, $\tau\nu$) events with a W mass $> 90\,\text{GeV}$, each set generated with a different value of Γ_W. The high M_T tails ($120\,\text{GeV} > M_T > 180\,\text{GeV}$) of these sets contain 3.4×10^7 W events each, and are used as the templates in the $\ln(\mathcal{L})$ fit. The exact values of Γ_W used for the nine templates are 1700, 1900, 2000, 2050, 2100, 2150, 2200, 2300, 2500 MeV. These values were selected to emphasise the quality of the fit near the minimum of the $\ln(\mathcal{L})$ curve.

- **Template normalisation spectra:** 9 sets of 2.5×10^8 $W \to l\nu$ ($e\nu$, $\mu\nu$, $\tau\nu$) events without a W mass cut, each one generated at one of the 9 values of Γ_W listed in the previous point. These are used to normalise the templates, as is explained in subsection 6.3.2.

- **Primary data:** 1 set of 7.5×10^7 $W \to l\nu$ ($e\nu$, $\mu\nu$, $\tau\nu$) events with a W mass $> 90\,\text{GeV}$, generated with a value of $\Gamma_W = 2120\,\text{MeV}$, the PDG value of July 2000 [Gro00].

- **PDF data:** All events listed so far were generated with PDF MRST($c-g$). A further two sets of 7.5×10^7 $W \to l\nu$ ($e\nu$, $\mu\nu$, $\tau\nu$) "primary" events (i.e. with $\Gamma_W = 2120\,\text{MeV}$) with a W mass $> 90\,\text{GeV}$ were generated with the following PDFs: MRST($l-\alpha_s$) and MRST($h-\alpha_s$).

- **Data normalisation spectra:** Three sets of 4.5×10^7 $W \to l\nu$ ($e\nu$, $\mu\nu$, $\tau\nu$) events without a W mass cut, all generated with $\Gamma_W = 2120\,\text{MeV}$, the PDG value of July 2000. Each of these three spectra was generated with one of each of the PDFs mentioned in the Primary data and PDF data entries above.

6.3 Data Analysis and Results

6.3.1 Cuts for Data Analysis

To ensure well-behaved Ws, following cuts were used

- $E_T^{miss} > 25$ GeV
- No jets with $p_T > 30$ GeV
- W recoil $|u| < 20$ GeV

The first cut significantly reduces the QCD background events, which tend to have a lower E_T^{miss}. The second and third cuts reject some useful events in order to maintain a clean signal by avoiding both QCD jets which may smear the E_T^{miss} and p_T^{miss} signals, and high p_T Ws. The losses due to these cuts are $\sim 25\%$. However, in this study the data generation cuts of section 6.2.2 ensure the QCD background is not generated in the first place, so the loss factor is lower.

6.3.2 Template Normalisation and Scaling

It is necessary to normalise the templates before they can be compared to the primary dataset (or datasets generated with different PDFs). The object of the normalisation is to ensure that the resulting curve is an M_T probability density function for a single W event, which can then be scaled up to the number of events detected.

Normalisation is carried out by using the template normalisation spectra described in section 6.2.3. For each of these spectra, the content of each bin in the range 20 to 180 GeV is divided by the total number of events in the range, ensuring the area under the normalised histogram equals one. The higher statistics templates are then scaled down such that their high M_T tails match the corresponding, now normalised, normalisation spectra. The normalisation range of 20 to 180 GeV was chosen as it is broad enough to satisfy a key requirement of the normalisation procedure, namely that the total W cross section integrated over the range does not change due to a variation of Γ_W between 1.7 to 2.5 GeV.

In order to perform the likelihood calculations, the normalised high statistics M_T tails were scaled up to match the simulated data set generated with $\Gamma_W = 2120$ and 7.5×10^7 $W \to l\nu$ ($e\nu$, $\mu\nu$, $\tau\nu$) events (equivalent to three years LHC low luminosity running). This was done by comparing the data set with the corresponding data normalisation spectrum, and scaling up this spectrum until it matched the data tail. The area under this scaled-up data normalisation spectrum was multiplied by the normalised templates to provide the templates used in the analysis.

6.3.3 Results of the ln($\mathcal{L}$) fits

As is described in section 5.4, the binned log-likelihood fit is carried out by comparing the data set with each template, bin by bin, to obtain the likelihood that the data set value of each bin could be a measurement of a Poisson distribution whose expected value is the template content of the bin. These bin-by-bin probabilities are summed up for each template to obtain the corresponding likelihood that the template is the parent distribution of the data set. Finally, a fit is performed to find the likelihood $\mathcal{L}$ as a function of Γ_W. The value at which -2ln($\mathcal{L}$) is minimised is the measured value of Γ_W. The statistical error on this value is determined by finding the values of Γ_W for which the value of -2ln($\mathcal{L}$) is 1 greater than the minimum, see section 5.4.2.

See curve 6.4.

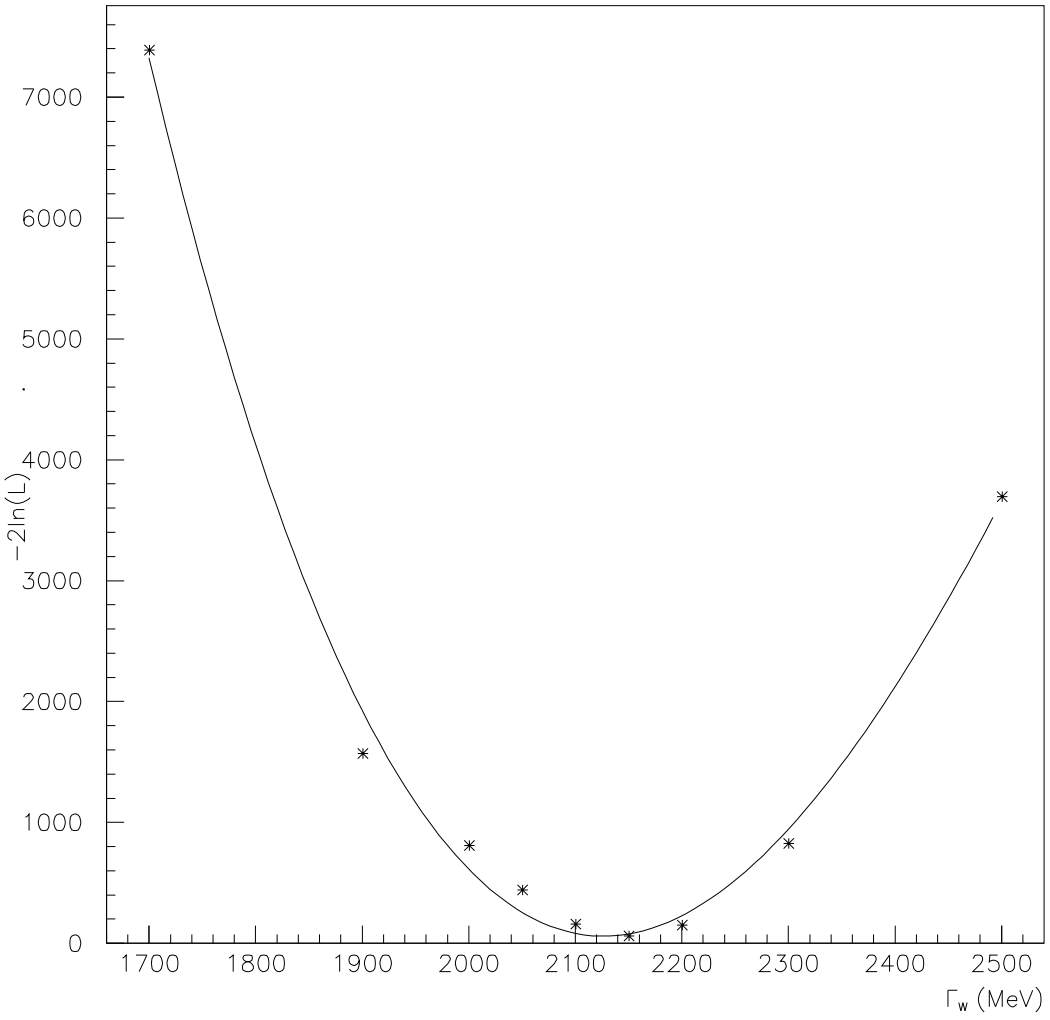

Figure 6.4: Likelihood as function of Γ_W. Data points (asterisks) show -2lnl, where $\mathcal{L}$ is the likelihood that the Primary Data Set, generated with $\Gamma_W = 2.12\,\text{GeV}$, is a daughter of each of the templates in turn: Γ_W = 1.7, 1.9, 2.0 2.05, 2.1, 2.2, 2.3, and 2.5 GeV. Primary data and templates produced with PDF = MRST(c-g). Curve is 3rd degree polynomial fit to the points, used to approximate a continuous -2ln($\mathcal{L}$) function. Resulting fit yields a minimum -2ln($\mathcal{L}$) (maximum likelihood) for $\Gamma_W = 2.1263\,\text{GeV}$, with an implied accuracy due to data set statistics (3 years LHC low luminosity) of $\sigma_{\Gamma_W} = 0.0056\,\text{GeV}$. For a discussion on the accuracy of the likelihoods and the resultant fit, see section 6.3.4.

Primary Data The fit of the primary data set to the templates is presented in figure 6.4. A 3rd degree polynomial is fitted to the points. The minimum of the fitted function occurs at $\Gamma_W = 2.1263\,\text{GeV}$. The corresponding value of -2ln($\mathcal{L}$) is 59.224. The values of Γ_W corresponding to a -2ln($\mathcal{L}$) of 60.224 are 2.121 and 2.132 GeV, implying a statistical error of ± 0.0056 GeV.

PDF Data Data sets of the same statistics as the primary data set were generated calling different Parton Distribution Functions from PDFLIB. Table 6.1 shows the results of the ln($\mathcal{L}$) fit for the datasets generated with the different Parton Distribution Functions used.

	"Primary" MRST $(c-g)$	MRST $(l-\alpha_s)$	MRST $(h-\alpha_s)$
Γ_W [MeV]	2126	2147	2105
σ_+ [MeV]	5.57	5.59	5.48
σ_- [MeV]	5.54	5.64	5.43
Min[ln($\mathcal{L}$)]	59.224	84.073	35.884

Table 6.1: Variation of best fir Γ_W and σ_{Γ_W} for selected PDFs Possible systematic effects of modelling templates with a PDF different from that governing LHC physics are explored by using two additional PDFs, each different from that used in template generation, to create secondary data sets. The Γ_W variation between the different PDFs should be compared with the total uncertainty in this study, 12 MeV, (see section 6.3.4). For a discussion of these results, see section 6.3.6. Note how the σ_+, σ_- across this table are consistent with $5.6 \pm 0.2\,\text{MeV}$, the value of $\sigma_{\Gamma_W(DATA)}$ presented in section 6.3.4.

Discussion of the range of Γ_W obtained with the PDFs is postponed until section 6.3.6, in order to include the effects of template statistics on the measured uncertainty on Γ_W (sections 6.3.4 and 6.3.5).

6.3.4 Effects of Template Statistics on Γ_W and σ_{Γ_W} Accuracy

In principle, the role played by a template in the ln($\mathcal{L}$) calculation is that of the true distribution, of which the observed data is assumed to be a sampling. In reality, the ln($\mathcal{L}$) and subsequent Γ_W fit reflect the match between the data and the template, rather than the match between the data and the true distribution the template aims to represent.

As a result, any deviation of the template from the true distribution affects the result of the Γ_W fit.

It is necessary to measure the effect of this on the accuracy of the results obtained with the templates used in this study, in order to ensure that the reported value of σ_{Γ_W} is meaningful. Furthermore, an error due to the finite template statistics could contribute to any discrepancy between the input value of Γ_W and the reconstructed value, particularly if this difference is substantially greater than σ_{Γ_W}.

The effect of finite template statistics can be observed by studying the variation in the Γ_W fit caused by a gaussian fluctuation of the value of every bin in every template *prior* to normalisation and scaling. In practice, 200 different randomised template sets were generated by replacing the content of each bin with a gaussian random variable of mean μ equal to its original value and standard deviation σ equal to the square root of that original value.

The rationale for this procedure follows from the Poisson distribution that governs the content of the templates. If a certain number of events is used to generate a given template, the resulting spectrum should be, bin by bin, within $\sqrt{N}$ (where N is the content of the bin) of the true underlying distribution 68% of the time[1]. Therefore, by allowing a template to fluctuate, bin by bin, with a gaussian distribution of $\mu = N$ and $\sigma = \sqrt{N}$, the actual range of variation due to finite template statistics is explored. The effects of this variation on functions of the template (*i.e.* Γ_W, σ_{Γ_W}) are found by obtaining each one from each of the randomised template sets, and studying the resulting distributions.

The variation presented by Γ_W and σ_{Γ_W} due to the 200 randomised templates is presented in figure 6.5. Both distributions are approximately gaussian; the σ of the best-fit gaussian are $\sigma_{\Gamma_W(TEMPLATE)} = 10.4$ MeV and $\sigma_{\sigma_{\Gamma_W}(TEMPLATE)} = 0.17$ MeV.

Thus the error on σ_{Γ_W}, $\sigma_{\sigma_{\Gamma_W}(TEMPLATE)}$ is found to be about 0.2 MeV. The overall error on the Γ_W values reported in this analysis is calculated by adding $\sigma_{\Gamma_W(DATA)}$ of 5.6 MeV and $\sigma_{\Gamma_W(TEMPLATE)}$ of 10.4 MeV in quadrature, yielding $\sigma_{\Gamma_W(TOTAL)}$ of 11.8 MeV.

As a consistency check, the reconstructed value of Γ_W, 2126 MeV$\pm$ 5.6 MeV$_{(DATA)}$ $\pm$ 10.4 MeV$_{(TEMPLATES)}$, is consistent with a "true" value of 2120 MeV, the value with which the data simulation was generated.

[1]Assuming a gaussian uncertainty, which is valid for Poisson distributions of high statistics

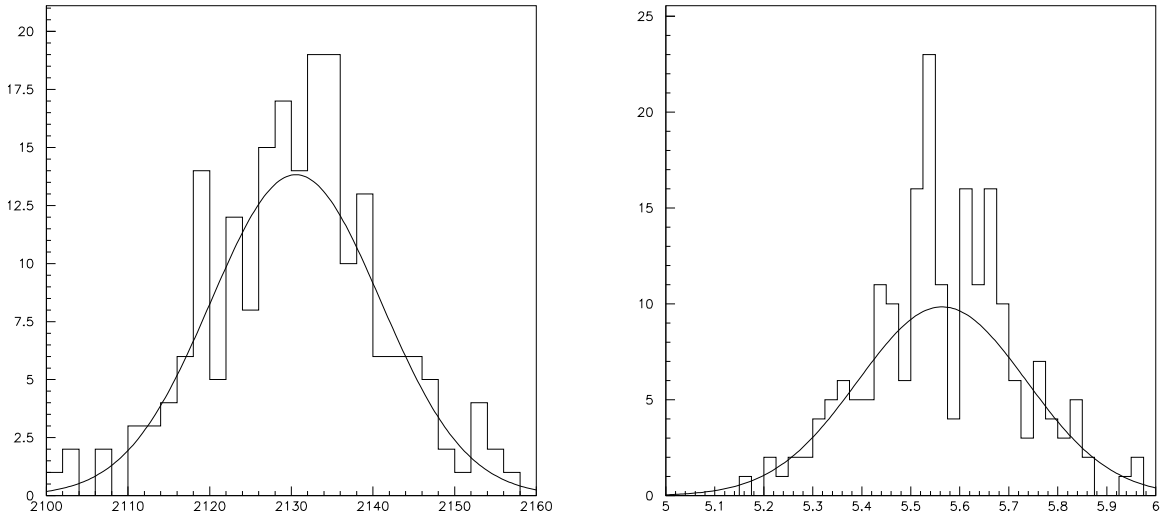

Figure 6.5: Variations of Γ_W and σ_{Γ_W} due to finite template statistics. Both plots show histograms analysing 200 randomised template sets. Left histogram is of Γ_W: gaussian fit has $\mu = 2130\,\text{MeV}$ and $\sigma = 10.39\,\text{MeV}$. Right histogram is of σ_{Γ_W}: gaussian fit has $\mu = 5.56\,\text{MeV}$ and $\sigma = 0.172\,\text{MeV}$

6.3.5 Template Statistics Required for an LHC Analysis

The template statistics used for this study would be not be sufficient to perform an adequate analysis of data generated at the LHC, since the uncertainty on Γ_W due to template statistics, $\sigma_{\Gamma_W(TEMPLATE)}$, is greater than the uncertainty due to the data set statistics [2].

In order to estimate the template statistics necessary for a meaningful LHC measurement of Γ_W, and to better understand the behaviour of the template error, the dependence of $\sigma_{\Gamma_W(TEMPLATE)}$ on the number of events used to generate the templates is now studied in more detail.

For a total number of events $\mathcal{N}$ generated to create a template, assume a specific bin in the generated spectrum contains N events. Given that N is determined by a Poisson probability, then (provided that N is larger than a few tens) it can be assumed that N will be within $1/\sqrt{N}$ of $\langle N \rangle$, the expectation value of the bin content, 68% of the time.

The template value of each bin, λ, which is used for the $\ln(\mathcal{L})$ calculation and the

[2]However, the template statistics *are* enough to allow an estimate of σ_{Γ_W}, as is discussed at the beginning of section 6.3.4

Γ_W, σ_{Γ_W} fits, is determined by normalising the template and scaling it to the number of events in the observed data set. Therefore, for each bin $\lambda \propto N$ and $\Delta\lambda \propto \Delta N$ with the same constant of proportionality. This implies that λ is expected to be within $\Delta\lambda$ of $\langle\lambda\rangle$ 68% of the time, and also that $\Delta\lambda/\lambda = \Delta N/N = 1/\sqrt{N}$.

The Poisson probability, calculated for each bin, that r instances of independent events occur (*i.e.* the data) for an expected value of λ (*i.e.* the template) is described by

$$P(\lambda, r) = \frac{\lambda^r}{r!} e^{-\lambda} \tag{6.2}$$

so the rate of change of the probability with respect to λ is given by

$$\frac{\delta P(\lambda, r)}{\delta\lambda} = P(\lambda, r) \cdot (\frac{r}{\lambda} - 1)$$

hence

$$\Delta P(\lambda, r) = P(\lambda, r) \cdot (\frac{r}{\lambda} - 1) \cdot \Delta\lambda. \tag{6.3}$$

Thus, the dependence of the Poisson probability on the template statistics can be written as

$$\Delta P(\lambda, r) \propto \Delta\lambda = \lambda \cdot \frac{1}{\sqrt{N}} \tag{6.4}$$

This result permits a relative change in $\Delta P/\ P$ to be estimated for separate instances in which a given template was generated with different statistics. Namely, if a template generated with an overall number of events $\mathcal{N}_1$, which through a bin content of N_1 results in a Poisson uncertainty in a given bin of $\Delta P_1/P_1$, a subsequent generation with a new number of events $\mathcal{N}_2$ will lead to a new Poisson uncertainty of

$$\frac{\Delta P_2}{P_2} = \frac{\Delta P_1}{P_1} \cdot \sqrt{\frac{N_1}{N_2}}. \tag{6.5}$$

The calculation of the $\ln(\mathcal{L})$ is carried out by summing the logarithms of the Poisson probabilities (see Appendix B), which is equivalent to multiplying the probabilities themselves. Thus, if every probability has an error which scales with template statistics as equation 6.5, the overall error on $\ln(\mathcal{L})$ should scale accordingly. Therefore it is expected that the errors on Γ_W and σ_{Γ_W} both scale with $1/\sqrt{N}$.

This conclusion has been tested by performing a simulation of the effect of different template statistics on the error of the extracted quantities Γ_W and σ_{Γ_W}. These simula-

tions are basically a repetition of the previous simulation used to find the error on Γ_W and σ_{Γ_W}, performed for a range of assumed template statistics. The effect of different template statistics can be produced by introducing a scaling factor $\sqrt{F}$ next to the $\sqrt{N}$ term in the randomising procedure described earlier in this section. With this, the *relative* error which corresponds to scaling the template statistics by F, the square of the factor, is found. As has been previously argued, this relative error is sufficient to simulate the effect of higher statistics on Γ_W and σ_{Γ_W} since the templates are normalised before use.

No. events per template	*'F'*, no. events as multiple of present study	Γ_W uncertainty (MeV)	σ_{Γ_W} uncertainty (MeV)
7.5×10^{10}	100	1.08	0.016
3.8×10^{10}	40	1.74	0.028
7.5×10^{9}	10	3.18	0.047
7.5×10^{8}	**1 (current study)**	**10.4**	**0.17**
1.8×10^{8}	0.25	23.2	0.35

Table 6.2: Result of randomising templates with $\sqrt{N}$ chosen to simulate a range of template statistics. Note the scaling of $\sigma_{\Gamma_W(TEMPLATE)}$ is consistent with $1/\sqrt{N}$ to within 10%

Table 6.2 shows how the accuracy of Γ_W and σ_{Γ_W} scales with N; the results are consistent with a $1/\sqrt{N}$ relationship to within 10%. This is also consistent with what could be expected from the Central Limit Theorem.

Table 6.2 shows that in order to have a total error determined primarily by LHC statistics, templates generated through this technique would require over one order of magnitude higher statistics than those used in this study. As an example, if 100 times higher statistics were used, the total error would be:

$$\sigma_{\Gamma_W(TOTAL)} = 5.6\,\mathrm{MeV}_{(DATA)} \oplus 1.1\,\mathrm{MeV}_{(TEMPLATES)} = 5.7\,\mathrm{MeV}_{(TOTAL)}.$$

For the current study, approximately 10^{10} template events were generated, using approximately 10^{17} floating point operations, and taking about 30 days of Central Simulation Facility runtime with an average allocation of 45 Gflops/s (30 processors $\times$ 1.5 Gflops/s rating per processor). While a similar resource allocation would require 3000

days for a full LHC study (≈100 times more events), a significant increase in computing power per researcher by 2010 is likely. However, it must be noted that generating templates of sufficient statistics to achieve LHC accuracy through Monte Carlo simulation will be a CPU intensive undertaking.

6.3.6 Discussion of PDF results

The discussion of PDF results is now resumed from section 6.3.3. Table 6.1 presents the results of several Γ_W fits, each performed for a data set generated with a different PDF. The object of this is to determine what effect the choice of PDF has on the value of Γ_W found.

The MRST collaboration has argued [MRST99] that the uncertainty in α_s is likely to be one of the more important sources of uncertainty in PDFs at the LHC in the W and Z production regime. This is partly because, given $\sqrt{s} = 14\,\text{TeV}$, production of a particle at $\approx 100\,\text{GeV}$ is in the low x regime, *i.e.* the partons that actually collide to create the gauge bosons carry a small fraction of the beam proton momentum. In this low x regime, the speed at which PDFs evolve with x has a high α_s dependence.

In order to explore this effect, the collaboration has put forward a low α_s PDF and a high α_s PDF, which describe the range of uncertainty. These are the test PDFs used for comparison with the default"primary" PDF in this study.

Section 6.3.4 shows that the overall resolution of the current study, including the effects of both finite data statistics and finite template statistics, is $\sigma_{\Gamma_W(TOTAL)} = 12\,\text{MeV}$. The best fit values of Γ_W obtained for the PDFs used are expressed in terms of this in Table 6.3.

	"Primary" MRST $(c-g)$	MRST $(l-\alpha_s)$	MRST $(h-\alpha_s)$
Γ_W [MeV]	2126	2147	2105
Difference from 2120, in σ	0.5	2.3	-1.3

Table 6.3: Effect of PDF variation as a function of $\sigma_{\Gamma_W(TOTAL)}$. The low and high α_{strong} PDFs present a variation which may be greater that that attributable to the statistical error on the current study, 12 MeV, although current template statistics are insufficient to resolve this unambiguously. The possible effect would be resolved by templates with higher statistics.

The primary data set reconstructed Γ_W differs from the PDG 2000 value (=2120 MeV),

with which all events were generated, by 0.5 $\sigma_{\Gamma_W(TOTAL)}$. The first test PDF, MRST(l-α_s) (low α_{strong}), is within 2.3 $\sigma_{\Gamma_W(TOTAL)}$ of the generating value. This suggests there may be a PDF effect since it implies a 2% probability of the difference being only due to the resolution of this study. However, a 2.3σ effect is not indicative of an unambiguous discrepancy. The second test PDF, MRST(h-α_s) (high α_{strong}), shows a deviation of -1.3$\sigma_{\Gamma_W(TOTAL)}$, therefore its effect, if any, is beyond the resolution of this study.

Overall, this analysis suggests that PDF variations may have an effect, but if so it is on the edge of the resolution of the current template statistics. This finding validates the present study as being worth pursuing in spite of the systematic limitations from PDF accuracy, although the accuracy of a full LHC analysis may be limited by PDF uncertainties. Such a result could in turn provide new information to constrain future PDFs.

6.3.7 Consistency with CDF

The current study uses a simulated data set of 208,270 electrons in the far M_T tail to achieve the fit shown in fig. 6.4, yielding a statistical uncertainty due to the data statistics of 5.6 MeV. A similar study carried out at CDF, using a real data set of 456 electrons in the far tail, yielded $\sigma_{stat} = 130$ MeV. If the CDF error were to scale with $1/\sqrt{N}$, the expected error for the current study would be approximately 6 MeV, which is consistent with this result.

6.4 Summary of Results

Estimate of statistical accuracy on Γ_W expected from M_T analysis of LHC 3 years low lumi data: Section 6.3.3 shows the expected statistical uncertainty is 5.6 MeV. Section 6.3.4 shows the error on this uncertainty is approximately $\pm$ 0.2 MeV

Template statistics needed in order to achieve LHC accuracy through Monte Carlo templates: Section 6.3.5 argues that in order for real data gathered over the first three years of LHC operation to be analysed with MC templates, about 100 times the events generated for the present study would be required ($\approx 10^{12}$ events), requiring a significant amount of processor operations ($\approx 10^{19}$ flops).

Estimate of possible systematic effect due to PDF uncertainty: The current study suggests that PDF variations may have an effect, but lacks the resolution to detect it unambiguously. A future LHC study with higher resolution templates may have a systematic error comparable with its statistical error due to PDF uncertainties. Such a result could in turn provide a possible constraint on future PDFs.

Present results consistent with comparable CDF results: The current study obtains a statistical uncertainty of 5.6 MeV with three years LHC data (208,270 electrons in the far M_T tail). A similar study carried out at CDF (456 electrons), yielded $\sigma_{stat} = 130$ MeV. The statistical error scales with $1/\sqrt{N}$, which shows the measurements are consistent.

Appendix A

Speckle Statistics and Autocorrelation

In this appendix, which closely follows the work of Goodman [Goo75] , the autocorrelation function of the subjective speckle pattern is defined, derived, and used to obtain equation 2.5. This is a more rigorous treatment than that presented in subsection 2.1.3 of the 'speckle diameter' or, more precisely, the distance which a speckle pattern can be displaced along the image plane before the intensity correlation necessary for ESPI is lost.

In order to achieve this aim the *objective* speckle pattern is introduced. The statistics of its amplitude and phase are studied, and the autocorrelation of the intensity is derived. Finally, the effect of a circular apperture is considered and the *subjective* speckle autocorrelation on the image plane is obtained.

A.1 Objective Speckle and Speckle Correlation

A.1.1 Speckle Intensity Probability Distribution

When an optically rough surface on plane (ξ, η) is illuminated with coherent light, the intensity on plane (x, y) (see figure A.1) is a random distribution of bright and dark speckles. This is called the *objective speckle pattern*. It can be is explained by considering the electric field A at (x, y) as a sum of N different phasors, with amplitude $(1/\sqrt{N})|a_k|$ and phase ϕ_k, which are reflected off different points of the (ξ, η) plane:

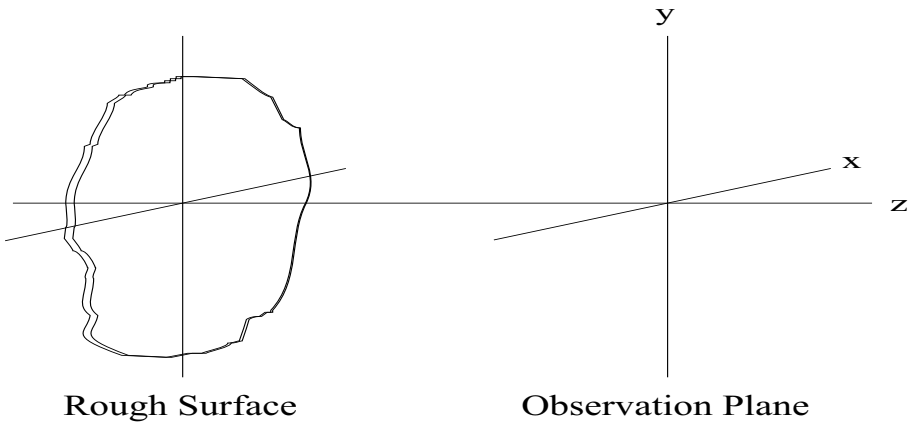

Figure A.1: The objective speckle pattern. The intensity on the (x, y) plane is a random variation of bright and dark speckles due to the coherent, dephased light reflected from a finite rough surface at (ξ, η).

$$A(x,y) = \frac{1}{\sqrt{N}} \sum_{k=1}^{N} |a_k| e^{i\phi_k} \tag{A.1}$$

If both amplitude and phase of each phasor are independent of each other and every other amplitude and phase, the sum is the result of a two dimensional 'drunkard's walk'. Given that the surface is large with respect to the roughness of the surface, it can safely be assumed that the number N of phasors contributing to the total electric field $A(x, y)$ at any point is very large. The Central Limit Theorem of statistics **??** ensures that both the real and the imaginary parts of $A(x, y)$, $\Re A(x,y) = \sum \Re a_k$ and $\Im A(x,y) = \sum \Im a_k$, have a Gaussian distribution because they are both sums of many independent random variables of the same distribution. Therefore the joint probability density function of the real and imaginary parts is

$$p_{\Re A, \Im A}(\Re A, \Im A) = \frac{1}{2\pi\sigma^2} \exp\left(-\frac{[\Re A]^2 + [\Im A]^2}{2\sigma^2} \right) \tag{A.2}$$

where the variance is

$$\sigma^2 = \lim_{N\to\infty} \frac{1}{\sqrt{N}} \sum_{k=1}^{N} \frac{\left\langle |a_k|^2 \right\rangle}{2} .$$

This is a circular Gaussian distribution, i.e. a Gaussian distribution with circular contours of equal probability in the complex plane.

It is now convenient to express equation A.2 in terms of the intensity I and phase θ of the amplitude instead of its real and imaginary parts. Once this is done, the joint probability for I and θ can be integrated over all θ to yield the probability distribution for I. Using the transformation

$$\begin{aligned} \Re A &= \sqrt{I}\cos\theta \\ \Im A &= \sqrt{I}\sin\theta \end{aligned} \tag{A.3}$$

the joint probability is expressed in terms of I and θ as

$$p_{I,\theta}(I,\theta) = p_{\Re A,\Im A}(\sqrt{I}\cos\theta, \sqrt{I}\sin\theta)||J||. \tag{A.4}$$

Here $||J||$ is the modulus of the Jacobian of the transformation

$$||J|| = \left|\left| \begin{array}{cc} \frac{\partial \Re A}{\partial I} & \frac{\partial \Re A}{\partial \theta} \\ \frac{\partial \Im A}{\partial I} & \frac{\partial \Im A}{\partial \theta} \end{array} \right|\right|. \tag{A.5}$$

For transformation A.3 it is easily shown that $||J|| = \frac{1}{2}$. Substituting this value in equation A.4 and using equation A.2 gives

$$p_{I,\theta}(I,\theta) = \frac{1}{4\pi\sigma^2}\exp\left(-\frac{I}{2\sigma^2}\right). \tag{A.6}$$

To obtain the probability distribution of the intensity regardless of phase, equation A.6 must be integrated over all θ. Thus

$$p_I(I) = \int_0^{2\pi} \frac{1}{4\pi\sigma^2}\exp\left(-\frac{I}{2\sigma^2}\right) d\theta = \frac{1}{2\sigma^2}\exp\left(-\frac{I}{2\sigma^2}\right). \tag{A.7}$$

Because the probability distribution of I has a negative exponential form, it can be shown that

$$\langle I^n \rangle = n\,!\,(2\sigma^2)^n = n\,!\,\langle I \rangle^n \tag{A.8}$$

where $\langle I \rangle = \int_0^\infty I\, p_I(I)\, dI$ is the expectation value or mean. For $n = 2$,

$$\left\langle I^2 \right\rangle = 2\,\langle I \rangle^2. \tag{A.9}$$

A.1.2 Intensity Autocorrelation

The autocorrelation function is defined as

$$R(x_1, y_1; x_2, y_2) = \langle I(x_1, y_1) I(x_2, y_2) \rangle \tag{A.10}$$

where the average is carried out over an ensemble of macroscopically similar but microscopically different surfaces. Using equation A.9 this can be rewritten as

$$R(x_1, y_1; x_2, y_2) = \langle I(x_1, y_1) \rangle \langle I(x_2, y_2) \rangle + |J_A(x_1, y_1; x_2, y_2)|^2 \tag{A.11}$$

where the *mutual intensity* J_A is defined as

$$J_A = \langle A(x, y) A^*(x, y) \rangle \,. \tag{A.12}$$

(From equation A.12 it is clear that $J_A = \langle I(x, y) \rangle$.)

The problem of finding the autocorrelation function for the objective speckle pattern is thus reduced to finding the mutual intensity. Consider the field at a point (x, y) due to the field (ξ, η) (see figure A.1). Using the Fresnel approximation **??** it can be written as follows.

$$\begin{aligned} A(x, y) &= \frac{1}{\lambda z} \exp\left[-i\frac{\pi}{\lambda z}(x^2 + y^2)\right] \iint_{-\infty}^{\infty} \alpha(\xi, \eta) \\ &\quad \exp\left[-i\frac{\pi}{\lambda z}(\xi^2 + \eta^2)\right] \exp\left[-i\frac{\pi}{\lambda z}(x\xi + y\eta)\right] d\xi \, d\eta \end{aligned} \tag{A.13}$$

where $\alpha(\xi, \eta)$ is the field at the scattering surface, z is the distance between the scaterring plane and the plane of observation (see figure A.1), and λ is the wavelength of the scattered light.

Let $A(x_1, y_1)$ be expressed as such an integral over variables of integration (ξ_1, η_1) and let $A(x_2, y_2)$ be similarly expressed in terms of (ξ_2, η_2). Now substitute these expressions in equation A.12, and invert the order of integration and averaging. The mutual intensity in the (x, y) plane becomes

$$
\begin{aligned}
J_A(x_1, y_1; x_2, y_2) &= \frac{1}{\lambda^2 z^2} \exp\left[-i\frac{\pi}{\lambda z}(x_1^2 - x_2^2 + y_1^2 - y_2^2)\right] \\
&\iiiint_{-\infty}^{\infty} J_\alpha(\xi_1, \eta_1; \xi_2, \eta_2) \exp\left[-i\frac{\pi}{\lambda z}(\xi_1^2 - \xi_2^2 + \eta_1^2 - \eta_2^2)\right] \\
&\exp\left[-i\frac{\pi}{\lambda z}(x_1\xi_1 + y_1\eta_1 - x_2\xi_2 - y_2\eta_2)\right] \, d\xi_1 \, d\eta_1 \, d\xi_2 \, d\eta_2 \, .
\end{aligned} \tag{A.14}
$$

Two simplifications can now be made. Since it is the modulus of the mutual intensity which contributes to the autocorrelation of the speckle, the initial exponential factor in equation A.14 can be ignored. Secondly, it is assumed that the roughness of the scattering surface is too fine to be resolved by a lens of the size of the (x, y) observing region. Thus, the mutual intensity at the scattering plane

$$
J_\alpha(\xi_1, \eta_1; \xi_2, \eta_2) \cong \kappa \, P(\xi_1, \eta_1) P^*(\xi_2, \eta_2) \, \delta(\xi_1 - \xi_2, \eta_1 - \eta_2) \tag{A.15}
$$

where κ is a proportionality constant, the function $P(\xi, \eta)$ represents the amplitude of the field incident on the scattering spot (ξ, η), and $\delta(\xi, \eta)$ is a two-dimensional delta function. The result of these simplifications is that

$$
\begin{aligned}
J_A(x_1, y_1; x_2, y_2) &= \frac{\kappa}{\lambda^2 z^2} \iint_{-\infty}^{\infty} |P(\xi_1, \eta_1)|^2 \\
&\exp\left[i\frac{2\pi}{\lambda z}\left(\xi_1(x_1 - x_2) + \eta_1(y_1 - y_2)\right)\right] \, d\xi_1 \, d\eta_1 \, .
\end{aligned} \tag{A.16}
$$

Equation A.16 shows that the mutual intensity of the observed field is proportional to the Fourier transform of the incident intensity distribution $|P(\xi, \eta)|^2$ and depends only on the differences in observation plane coordinates.

It should be noted that

$$
J_A(x_1, y_1; x_2, y_2) = \frac{\kappa}{\lambda^2 z^2} \iint_{-\infty}^{\infty} |P(\xi, \eta)|^2 \, d\xi \, d\eta = \langle I \rangle^2 \, .
$$

Thus the autocorrelation function for the objective speckle, as described in equation A.11, can be written as

$$
\begin{aligned}
R_I &= \langle I \rangle^2 \left[1 + \left| \frac{\frac{\kappa}{\lambda^2 z^2} \iint_{-\infty}^{\infty} |P(\xi,\eta)|^2 \exp\left[i \frac{2\pi}{\lambda z} \left(\xi \Delta x + \eta \Delta y \right) \right] \, d\xi \, d\eta}{\langle I \rangle^2} \right|^2 \right] \\
&= \langle I \rangle^2 \left[1 + \left| \frac{\iint_{-\infty}^{\infty} |P(\xi,\eta)|^2 \exp\left[i \frac{2\pi}{\lambda z} (\xi \Delta x + \eta \Delta y) \right] d\xi \, d\eta}{\iint_{-\infty}^{\infty} |P(\xi,\eta)|^2 \, d\xi \, d\eta} \right|^2 \right] . \qquad \text{(A.17)}
\end{aligned}
$$

Equation A.17 should be studied in detail. It states that the dependence of the autocorrelation on $(\Delta x, \Delta y)$ is governed by the size, shape and illumination of the scattering surface. The explicit dependence is found by integrating this formula over the exact shape under study. Consider as an example a uniformly lit square surface of dimensions $L \times L$, where $|P(\xi,\eta)| = 1$ within the square and zero otherwise. The autocorrelation function of the intensity is

$$
R_I(\Delta x, \Delta y) = \langle I \rangle^2 \left[1 + \mathrm{sinc}^2 \frac{L \Delta x}{\lambda z} \mathrm{sinc}^2 \frac{L \Delta y}{\lambda z} \right] . \qquad \text{(A.18)}
$$

A.2 Subjective Speckle Autocorrelation

The results obtained in the previous section are modified if the speckle pattern occurs in an imaging geometry such as that shown in figure A.2. Since the speckle pattern in this configuration depends on the size and shape of the aperture, it is called the *subjective* speckle pattern.

Let the object area be uniformly reflective (or transmissive), uniformly illuminated, and broad relative to the resolution of the lens pupil. If these conditions are met, the size of the speckles incident on the lens pupil is extremely small compared to its diameter. Thus, to a good approximation, the mutual intensity of the fields *in the pupil* is given by equation A.15 where the coordinate system (ξ, η) now lies in the pupil plane. In effect, the pupil plane can be treated as a uniformly bright rough surface, since the distribution of the electric field along its surface is has the same characteristics as the distribution of

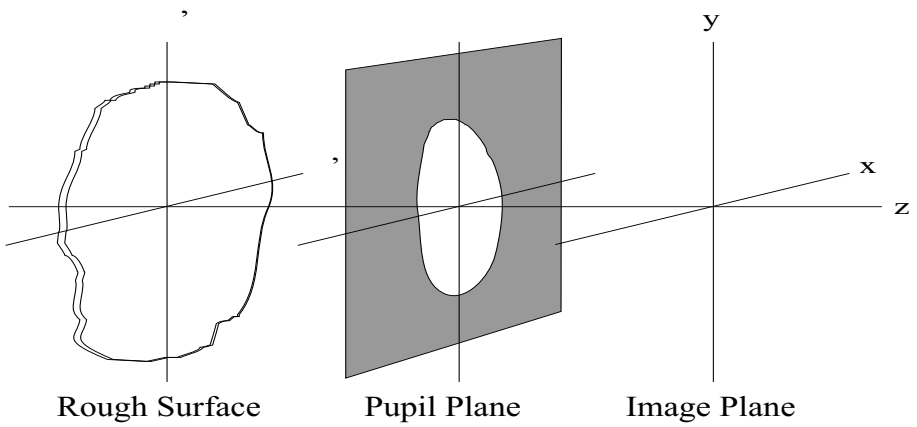

Figure A.2: The subjective speckle pattern.The imaging geometry at (ξ, η) determines the speckle on the image plane (x, y).

the field across an optically rough surface.

What remains is to calculate the autocorelation function of the image plane speckle. This is done using the results of the previous section together with the new interpretation of $P(\xi, \eta)$. As is stated in equation A.17, the autocorrelation function of the intensity pattern consists of a constant term plus the square of the modulus of the normalised Fourier transform of the intensity transmisittance $|P(\xi, \eta)|^2$ of the lens pupil. Note that the autocorrelation function is thus independent of any abberations of the imaging system, since these abberations only affect the phase of $P(\xi, \eta)$.

For the case of a circular pupil of diameter D, the autocorrelation function of the image-plane speckle is

$$R_I(r) = \langle I \rangle^2 \left[1 + \left| \frac{2J_1\left[\frac{\pi D r}{\lambda z}\right]}{\frac{\pi D r}{\lambda z}} \right|^2 \right] \tag{A.19}$$

where $r = ((\Delta x)^2 + (\Delta y)^2)^{1/2}$. This is identical to equation 2.5.

Appendix B

Binned Log Likelihood Fit

B.1 The Maximum Likelihood Method

The maximum likelihood method is a parameter fitting technique. It requires a measured data sample, and a hypothesis expressed as a probability distribution which is function of a parameter P. Thus every observed data point x_i has a probability given by [Lyo93]

$$y_i = y(x_i; \mathrm{P}). \tag{B.1}$$

The *likelihood* is defined as the product of the probabilities,

$$\mathcal{L} = \prod_i y_i. \tag{B.2}$$

The value of P for which the likelihood is maximal is taken to be the best estimate of P. Thus it is necessary to find the likelihood as a function of the parameter, and obtain the maximum. This can be done by numerical methods, see section B.3.

In order to estimate the error on P, the logarithm of the likelihood function is considered.

$$l = ln(\mathcal{L}) = \sum_i y_i \tag{B.3}$$

It is convenient to assume that $\mathcal{L}$ distribution tends to a gaussinan,

$$\mathcal{L} \sim e^{\frac{(\mathrm{P}-\mathrm{P}_0)^2}{2c}} \quad , \tag{B.4}$$

at least near its maximum. If this is true, the l distribution will tend to a parabola. This is shown by expanding l into a Taylor series about the maximum and using the assumed gaussian shape of $\mathcal{L}$ by replacing l'' with 1/c.

$$\begin{aligned} l &= l_{max} + \frac{1}{2!} l'' \delta \mathrm{p}^2 + ... \\ &= l_{max} - \frac{1}{2c} \delta \mathrm{p}^2 + ... \end{aligned} \tag{B.5}$$

It can be seen that a $\delta \mathrm{P} = (\mathrm{P} - \mathrm{P}_0) = \sigma$ corresponds to a difference of 1/2 in l. Therefore, if a parabola-like curve can be reasonably fitted to the l function (*i.e.* to the l *vs* P points), the $\mathcal{L}$ shape can be taken as approximately gaussian. It is then reasonable to estimate the error on P by taking the values of P for which

$$l = l_{max} - \frac{1}{2} \qquad \text{or} \qquad ln(\mathcal{L}) = ln(\mathcal{L})_{max} - \sigma \ . \tag{B.6}$$

B.2 The Binned Log Likelihood Method

The method can also be applied to histograms. If each bin in a histogram has an expected content given by the hypothesis and a value of the parameter P, (in this thesis such a histogram of expectation values is called a template, see chapters 5 and 6), a measured content (the data), and a known probability distribution relating an expected value and an observed value to a probability (*e.g.* a Poisson distribution), then a probability can be calculated for every bin. Once more, the likelihood is defined as the product of these probabilities, see equation B.2.

The value of P for which the expected histogram yields the highest $\mathcal{L}$ or l is the result of the fit. The error is estimated by finding the values of P for which l yields $l_{max} - 1/2$.

B.3 Finding the likelihood or log likelihood as a function of the parameter

Since the expected value histograms may be obtained through numerical calculations, there may not be an analytical form of $\mathcal{L}$ as a function of P. In this case, the likelihood function itself may be approximated numerically, for example by generating a number of

expectation value histograms with a range of values for the parameter, and then fitting some function to the $\mathcal{L}$ *vs* P points.

This can also be done for l rather than $\mathcal{L}$, since the result is the same but the function and the errors are easier to work with. In particular, fitting a 2^{nd} or 3^{rd} degree polynomial to the l *vs* P points provides some confidence that the assumed gaussian distribution of the $\mathcal{L}(\mathrm{P})$ function is not too far from reality.

Bibliography

[Abe95] F. Abe *et. al.* Direct measurement of the W Boson width. *Phys. Rev. Lett.*, 72(342), 1995. CDF Collaboration.

[AH89] I. J. R. Aitchison and A. J. G. Hey. *Gauge Theories in Particle Physics.* Institute of Physics Publishing, second edition, 1989.

[AHDD98] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali. The hierarchy problem and new dimensios at a milimeter. *Phys. Lett.*, B429:263-272, 1998. hep-ph/9803315.

[Ahm01] Q.R. Ahmad *et. al.* Measurement of the rate of $\nu_e + d \rightarrow p + p + e^-$ interactions produced by 8b solar neutrinos at the sudbury neutrino observatory. *Phys. Rev. Lett.*, 87(7:1301), 2001.

[Arf85] George Arfken. *Mathematical Methods for Physicists.* Academic Press, third edition, 1985.

[Ash98] William Joseph Ashmanskas. *A Direct Measurement of the W Boson Decay Width in Proton-Antiproton Collisions at 1.8 TeV.* PhD thesis, University of California, Berkley, 1998.

[ATLa] ATLAS Collaboration. *ATLAS Detector and Physics Performance, Technical Design Report Vol I.* CERN/LHCC/99-14, 1999.

[ATLb] ATLAS Collaboration. *ATLAS Detector and Physics Performance, Technical Design Report Vol II.* CERN/LHCC/99-15, 1999.

[BP87] V. D. Barger and R. J. H. Phillips. *Collider Physics.* Addison Wesley, third edition, 1987.

[Bru00] Rene Brun. *Phyisics Analysis Workstation (PAW).* CERN software, 2.14 edition, 2000.

[CDF] CDF Collaboration. *The CDF II Detector Technical Design Report Vol.* FERMILAB-Pub-96/390-E, 1996.

[CER05] CERN. *What's next at CERN?* http://public.web.cern.ch/Public/Content/Chapters/Ab en.html, 2005.

[CGHFG] LEP Collaborations, LEP Electroweak Working Group, SLD Heavy Flavour, and Electroweak Working Group. *A combination of preliminary electroweak measurements and constraints on the Standard Model.* CERN-EP 2001-098.

[DM73] D.A. Dicus and V. S. Mathur. Upper bounds on the values of masses in unified gauge theories. *Phys. Rev.*, D7(3111:3114), 1973.

[Fuk01] S. Fukuda *et. al.* Constraints on neutrino oscillarions using 1258 days of super-kamiokande solar neutrino data. *Phys Rev Lett.*, 86(5656-5660), 2001.

[Goo75] J.W. Goodman. Statistical properties of laser speckle patterns. In J.W. Dainty, editor, *Laser Speckle and Related Phenomena.* Springer-Verlag, 1975.

[Gri87] David Griffiths. *Introduction to Elementary Particles.* Wiley, 1987.

[Gro00] D.E. Groom *et. al.* Review of Particle Physics. *The European Physical Journal*, C15:1, 2000.

[Hal92] Donald Hall. *Basic Acoustics.* Krieger Publishing Company, reprint edition, 1992.

[Hec87] Eugene Hecht. *Optics.* Addison Wesley, third edition, 1987.

[HM84] Francis Halzen and Alan D. Martin. *Quarks and Leptons.* Wiley, 1984.

[Jon99] Robert M. Jones. *Mechanics of Composite Materials.* Taylor & Franics, second edition, 1999.

[KK73] Daniel Kleppner and Robert Kolenkow. *An Introduction to Mechanics.* McGraw-Hill, 1973.

[Lei73] Arthur Leissa. *Vibrations of Shells*. Acoustical Society of America / AIP, reprint edition, 1993 (1973).

[LP] P. Lefèvre(Ed) and T. Petterson(Ed). *The Large Hadron Collider: Conceptual Design*. CERN/AC/95-05, 1995.

[Lyo93] Louis Lyons. *Statistics for Nuclear and Particle Physicists*. Cambridge University Press, reprinted edition, 1993.

[MRST99] A. D Martin, R. G. Roberts, W. J. Stirling, and R. S. Thorne. *Parton Distributions and the LHC: W and Z Production*, July 1999. hep-ph/9907231.

[Plo00] Hans Plothow. PDFLIB *8.04: Proton, Pion and Photon Parton Density Functions, Parton Density Functions of the Nucleus, and α_s Calculations*, 17 April 2000. CERN W5051 PDFLIB.

[Rei96] Armin Reichold. *Development and Performance Measurements of a Silicon Tracker Prototype for ATLAS and a Study of Gluino Cascade Decays into Same Sign Dileptons using the ATLAS Inner Detector*. PhD thesis, Dortmund University, 1996.

[RWFP98] Elzbieta Richter-Was, Daniel Froidevaux, and Luc Poggioli. ATLFAST *2.0 a fast simulation package for ATLAS*, 10 November 1998. ATL-PHYS-098-131.

[Sjö94] Törbjorn Sjöstrand. *Computer Physics Communications*, 82(74), 1994. Also see http://www.thep.lu.se/torbjorn/Pythia.html.

[Soe81] Werner Soedel. *Vibrations of Shells and Plates*. Marcel Dekker, Inc., 1981.

[Sou] Point Source.

[Spy94] Constantine Spyrakos. *Finite Element Modeling in Engineering Practice*. West Virginia University Press, 1994.

[Sve89] Orazio Svelto. *Principles of Lasers*. Plenum Press, third edition, 1989.