Narrow Linewidth Lasers for use with Neutral Strontium as a Frequency Standard

by

Steven Johnson

A thesis submitted to
The University of Birmingham
for the degree of
DOCTOR OF PHILOSOPHY

Ultracold Atoms Group
School of Physics and Astronomy
College of Engineering and Physical Sciences
The University of Birmingham

June 2013

Abstract

The continual development of precise timing has always pushed technology forward, and the latest generation of clocks uses electronic transitions within cooled atoms as their frequency reference. For atoms where the clock transition is at an optical frequency an ultra stable laser is used to probe the transition; the construction and development of this laser for use with strontium is the subject of this thesis. The system consists of a laser, an optical cavity and the electronics to lock the laser to the optical cavity. The theoretical component of this thesis looked into the possible materials used to make the optical cavity from and how the optical cavity should be supported to minimize the frequency noise of the laser. A compact external cavity diode laser was constructed and shown to have a linewidth of 26 ± 15 kHz. To lock and narrow the linewidth of this laser, two ultra stable optical cavity assemblies were constructed with vibration and thermal stabilisation enclosures. To characterise the laser stability a beat measurement was performed between the two stabilised lasers and a beat note linewidth of 1.37 ± 0.28 Hz was measured. The instability of the beat note was measured; an Allan deviation of 2×10^{-14} was obtained between averaging times of 0.2 s and 1 s.

ACKNOWLEDGEMENTS

Completing a PhD is like climbing a mountain; you work hard and climb up the mountain until you reach the summit, and you can tell you have reached the summit because the world looks different. You are no longer looking at sheer rock but you can look around you and see what is around without having to hang on to the rock face. Previously you were pressing forward, trying to head upwards but sometimes heading back, going and trying different approaches or even having to talk to someone who had tried another path up the mountain. But the most significant part of being at the top of the mountain is you look down and you see where you could have gone a different route; you may see an easier route that would have been quicker but you never knew it was there. Sometimes you realise you could have taken a cable car, so you would have reached the summit but it would not have been the same; you would not have learnt how to scale the mountain. The main thing is you are not just seeing what is below you but realising that you are not on your own and you are surrounded by mountains, all with their own summit and other people trying to climb them.

This PhD is the story of how the mountain was climbed and the crags reached, the sheer rock I had to find my way round and the cliff edges I avoided. But it could not have been done without the people who supplied me with the ropes and crampons, the people who showed me the way round the cliffs and drops, and occasionally suggesting I try a new route that no one thought of taking before.

First and foremost I would like to acknowledge my lab partner Ole Kock who has been invaluable for discussions to develop ideas, assistance in the lab and for pointing out when I

am being less than clear. My supervisor Prof. Kai Bongs has been a huge support throughout my PhD and I would like to thank him for his time spent with me and in particular his time spent in my laboratory. I have had a number of undergraduate students who have assisted with taking results for this thesis and I must acknowledge Lyndsie Smith for her construction and testing of one of the ECDL laser systems and her assistance with the optical cavity locking for the measurement of the laser linewidth. I also thank Harry Proud for the investigation into EOMs which led to the realisation of the piezo-electric resonances effect. Jochen Kronäger I must thank for his patience with me when assisting with electronics and for designing a significant quantity of the electronics I used. I would also like to thank the other members of the University of Birmingham Cold Atoms group for making my PhD time so enjoyable. From other research groups at Birmingham I wish to thank the members of the Gravitational wave group who have given their time to assist me in understanding all aspects of optical cavities.

I had the great pleasure of spending a week with Liam Cunningham and Ronny Nawrodt at the University of Glasgow Institute for Gravitational Research and thank them for assistance with modelling of thermal noise in optical cavities. I have also had many useful and fruitful discussions with the people at the National Physical Laboratory in London, with very special thanks and huge gratitude to Liz Bridge who has been of great assistance and always willing to answer my questions. For technical help I thank Steve Brookes and the rest of the Physics machine workshop team who have produced a vast number of mechanical components for me, often in very short times, and for discussion on how things should best be built. I also thank Jon Fellows for the LaTeXfiles to produce this thesis; even if he believes all experimentalists are just Monkeys who have been trained to use spanners.

Finally thanks to my friends within the Archery Club, Guild of Students, house mates and my family. A very special thank you to Miss Rancins, my A-level Physics teacher, who convinced me to come to Birmingham to study Physics.

CONTENTS

LIST OF FIGURES

LIST OF TABLES

CHAPTER 1

INTRODUCTION AND MOTIVATION

1.1 Why do we need better clocks?

Why have scientists always strived for better measurements of time? Countries around the world have metrology institutes that make standards for length, voltage, mass, time and many other units. The reason so many countries spend money on these organisations is that having standardised units is fundamental to allowing technology to develop. Having standardised length units ensures manufacturers can ensure their products fit with those of other manufacturers. Developments in precision time keeping began with nautical chronometers used to determine longitude on long sea voyages using celestial navigation; John Harrison in 1759 was the first man able to build a precise clock that would maintain its accuracy on long sea voyages so that the ship's position could be calculated accurately. Clocks have developed a long way since then and precision clocks no longer use mechanical components to keep time. High precision clocks are still important to navigation today; every Global Positioning System (GPS) satellite contains an atomic clock. The time signal broadcast by these satellites allows us to use SatNavs and mobile phones for navigation. High precision clocks have become fundamental to communication technology; the rate of data sending for very fast systems is only possible with very precise clocks so that the computer systems are

correctly synchronized.

Improvement in the instability of frequency standards is enabling the testing of fundamental physics such as symmetry breaking and changes in the fine structure constant over time [1]. Not only can we use these clocks to test fundamental physics but we can also use them to measure the gravitational field itself. The theory of general relativity predicts that space-time is distorted by mass [2]; the greater the mass the greater the distortion. The further away you are from the distortion the less the effect it will have. In a gravitational field clocks will run slower compared to a clock in a lower gravitational field; this is a key prediction of general relativity that can be tested with high precision clocks [3].

Clocks in different gravitational potentials will vary in frequency. As the Earth orbits the Sun in an elliptical path, the distance between the Earth and Sun varies and therefore the gravitational potential will change during the year. By comparing measurements made at different times of the year the annual variation has been demonstrated [4].

Tests of the theory of special relativity have also been done with lithium ions which have been accelerated in a storage ring and, by examining the Doppler shift of the light, the relativistic frequency shift has been observed in frequencies emitted by these atoms [5].

A current planned space mission is to build the Atomic Clock Ensemble in Space (ACES) [6]. On this mission a laser cooled atom clock known as PHARAO is to be included, which is designed to be accurate to 1 part in 10^{16} . The objectives of ACES are to test the time and space invariance of fundamental constants, theories of relativity and to provide a proof of principle for putting a cold atom clock in space. Another mission based on using the ACES clock is the Einstein Gravity Explorer [7], which aims to measure gravitational time dilation at different points of the elliptical orbit of the clock around the earth due to the different strength of gravitational field.

1.2 How Clocks Work

All types of clock work in a similar way; they are based on some oscillator that serves as a frequency reference. The simplest of these is a pendulum that takes the same time to swing from one side to the other; the period of a pendulum swing is only dependent upon the length of the pendulum and the acceleration due to gravity at the location of the pendulum. By knowing a pendulum takes say one second to swing from one side to the other and back again we can use this to keep time; if we count the number of times a pendulum swings back and forth we know how many seconds have passed. But if our pendulum swing is not exactly one second but one second plus a millisecond, which seems fairly insignificant, over extended periods of time our clock will become increasingly wrong. Over a whole day that is 86400 seconds our clock will be almost a minute and a half slow, and over a whole year we will lose nine hours. This is why pendulum clocks need some natural reference such as the rotation of the earth to correct for these shifts. We could correct for the inaccuracy in our pendulum oscillator if every day we waited for the sun to be exactly at its highest point in the sky and reset the clock; then we would only ever be a maximum of a minute and a half wrong. If we were to do this we would need to correct for the variation in the position of the Sun, due to the Earth taking an elliptical path around the Sun, which causes a sundial to be out by up to 16 minutes compared to the time measured by a precision clock [8]. This difference can be calculated by what is known as the 'equation of time'. Without a natural reference of time passing to correspond to our man-made oscillator it is useless at keeping accuracy for long periods of time. If we imagine that we have a perfectly adjustable pendulum in a vacuum that has a period of exactly one second we would have a perfect oscillator. If we could measure the position of the Sun with absolute precision and take the equation of time into account, by measuring with our perfect oscillator we would find that the length of the day actually changes due to the rotation of the Earth slowing down; therefore for a time standard maybe the rotation of the Earth is not the best thing to use as our reference.

Fortunately we have been given something else that is more precise: atomic transitions. Within an atom electrons can move between energy levels and the energy difference in these levels relates to a particular frequency. The electron transitions found in the isotope of an element are the same in the UK as they are anywhere in the world; a laboratory in any country that measures the frequency of a particular transition will be able to measure exactly the same value anywhere in the world. However there are frequency shifts that must be taken into account. These shifts are predominantly cause by black body radiation, gravitational potential differences, Zeeman shift and the AC Stark effect. The frequency shift due to each process must be calculated to find the total frequency shift for each individual experiment if direct comparisons are to be made.

In 1955 Louis Essen at the National Physical Laboratory in Teddington, London, built the world's first caesium atomic clock using a microwave frequency transition in caesium as the natural reference. In this clock a microwave source that was adjusted to be resonant with the microwave frequency acted as the pendulum.

The pendulum in a mechanical clock can have a frequency of 1 period per second or 1 Hz, whereas a microwave frequency will have tens of billions of periods per second. The great benefit of using microwaves is the increased precision of the measurement we can make; with a pendulum we could measure to half or maybe a quarter of an oscillation with a fairly high certainty but any smaller fraction would be very difficult to define, so the smallest measurement of time we can make with a pendulum clock would be a quarter of a second. With microwaves even if we can only measure to one period we can still measure to a time of one tenth of a billionth of a second, which is significantly better than our pendulum clock.

However atoms also have optical frequency atomic transitions that have even more oscillations per second; by optical we mean that the radiation absorbed or emitted is in or near the frequency of visible light of the electromagnetic spectrum. These optical frequencies have several hundred trillions of periods per second; a suitable source of this radiation to use

as our oscillator would be a laser. However the light emitted by a laser is not sufficiently stable to directly measure the desired atomic transition. We need to stabilise the frequency of the laser and this is achieved by locking the frequency to a Fabry-Pérot optical cavity. The stability of the laser will depend on how stable the length of the optical cavity is so significant efforts are made to minimise the changes in length of the optical cavity.

1.3 This Work

The aim of this thesis is to demonstrate how a 'pendulum' for an optical frequency standard can be built. Our pendulum is a narrow linewidth laser system that has been built to be compact and rugged in design. The majority of work has been on the design, development and characterisation of the laser systems and the vacuum systems for the ultra stable optical cavities that are required to stabilise the laser. Alongside the work presented in this thesis, an apparatus has been built to prepare strontium atoms in order to be interrogated so that a measurement of an optical atomic transition can be made. Our experiment uses neutral strontium atoms that are laser cooled; the subject of the apparatus built and procedure used for cooling and trapping of the strontium atoms are to be published in the thesis of my lab partner Ole Kock [9].

This thesis is organised as follows: In chapter 2 the concept of atomic clocks will be discussed and how the atoms are prepared by laser cooling in order to be interrogated will be covered. In chapter 3 the design of a new external cavity diode laser will be discussed and an experimental test of its linewidth will be presented. We will show that it is possible to build an external cavity diode laser with a significantly improved intrinsic stability by extending the length of the external cavity. In chapter 4 the development of ultra stable Fabry-Pérot optical cavities and the various noise sources associated with them will be discussed. This includes computational work on the thermal noise of optical cavities showing what materials

should be selected for the optical cavities. Also presented is the modelling used to determine the best position to support our optical cavity in a way that is vibrationally insensitive. In chapter 5 the bulk of the experimental work will be presented including the design of the systems enclosing the ultra stable optical cavities, the temperature control and locking electronics used. The key results presented will be the beat measurements of the locked laser systems where a linewidth of 1.37 ± 0.28 Hz is measured; also the system is characterised by measuring an Allan variance, displaying a short term stability of 2×10^{-14} between 0.2 and 1 s averaging time. How these measurements are made and the justification of why this is done will be discussed.

CHAPTER 2

OPTICAL CLOCKS

2.1 Introduction

This chapter introduces the electronic structure of strontium atoms, highlighting the relevant energy levels for the clock transition and laser cooling. Laser cooling is used to prepare the atoms to reduce Doppler shifts in the individual atoms. The fundamental requirement of an atomic clock is the clock transition on which the frequency reference is based, and it is demonstrated why strontium is a suitable candidate to be a frequency standard. This optical frequency can be compared to the current caesium standard using an optical frequency comb to enable direct comparisons to the primary standard so that the frequency may be determined. Finally how an optical lattice is used to strongly trap the atoms is discussed.

2.2 Atomic Structure

Within an atom electrons occupy energy levels; these energy levels are defined by the physical laws of how the electrons interact with each other and the atom's nucleus [10]. By 'excite an atom' we mean that an electron from that atom is moved from its current energy state to another unoccupied state at a higher energy. An excitation is caused by a photon interacting

with the atom and being absorbed; this interaction can only occur if the energy of the photon is the same as the energy needed to excite the electron from one energy level to the next. As each electron state within an atom has a defined energy, to make an electron transition occur the incident photon must have a specific energy. The energy the photon requires is the difference in energy of the two states; if this is the case the photon will be absorbed and the electron moves to the higher energy state. The energy E of a photon is proportional to its frequency ν multiplied by Planck's constant h; $E = h\nu$.

The strontium atom has two valence electrons that can be excited by optical frequency transitions. In figure 2.1 the notation used on the diagram of the strontium energy levels describes which energy level these two valence electrons occupy [10]. The ground state orbital configuration of these electrons is $5s^2$ where 5 is the principle quantum number, s^2 indicates that both electrons occupy the zero orbital angular momentum state.

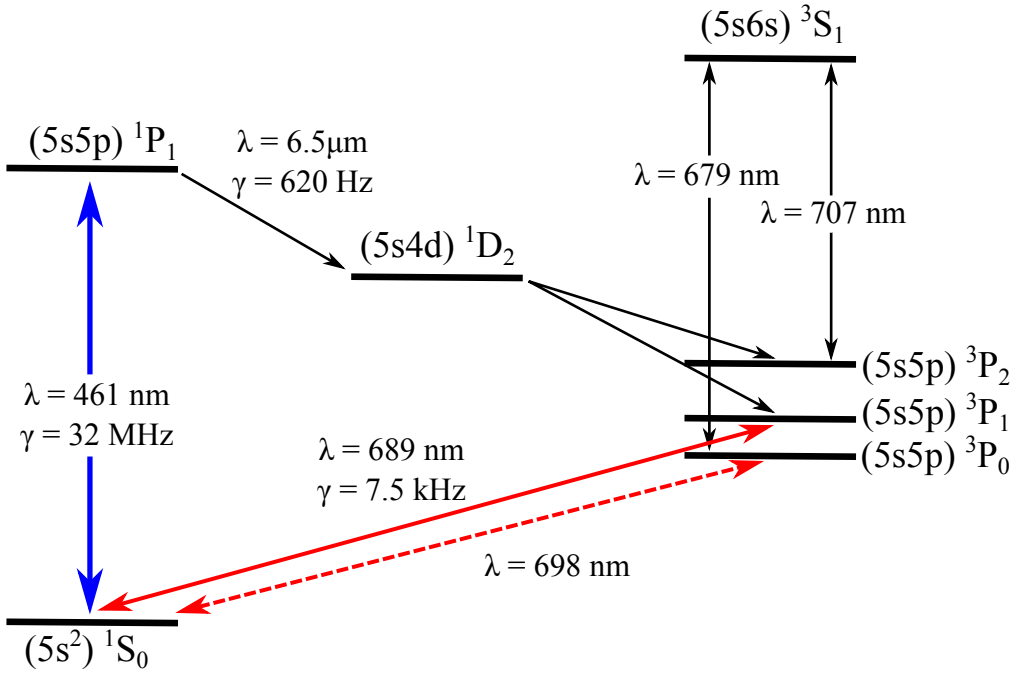

Figure 2.1: Partial diagram of the strontium energy levels.

We write the state the valence electrons occupy as $^{2S+1}L_J$ where S is the electron spin state that can be either singlet or triplet state, L is the orbital angular momentum and J

is the total angular momentum. The total orbital angular momentum J is the sum of the quantised vectors **L** and **S**, therefore the minimum value of J = |L| − |S| and the maximum value is J = |L| + |S|. The electrons have a total magnetic dipole moment due to the orbital and spin dipole moments; under the influence of a magnetic field this dipole moment will have a potential energy due to the magnetic dipole moment orientation with respect to an external magnetic field. For energy levels with J ≠ 0 there will be a splitting of the energy levels known as Zeeman splitting with an energy depending on the projection of the total angular momentum on the direction of the magnetic field. This projection is classified by the magnetic quantum number m_j, the minimum and maximum values of the m_j projection are -J and +J respectively.

Quantum mechanics dictates that an emitted or absorbed photon can only transfer one unit of orbital angular momentum, therefore we can say that ΔL = ±1. An electron couple transition cannot flip the spin of an electron, so ΔS = 0. From these statements we can conclude that our total angular momentum must have the relation ΔJ = 0 or ±1 but not J = 0 to J = 0. These limitations are known as selection rules and will give *allowed* transitions. If a transition does not obey these rules it is known as a *forbidden* transition. It is not impossible for a forbidden transition to occur; the transition can happen due to higher order electromagnetic transitions but it is extremely unlikely.

As discussed the photon emitted or absorbed during an atomic transition has a frequency given by the energy difference between the two states; the natural linewidth of this transition is given by the inverse lifetime of the excited state. According to the uncertainty principle the uncertainty in energy ΔE is inversely proportional to the lifetime of the excited state Δt:

$$\Delta E \Delta t \approx \frac{\hbar}{2}. \tag{2.1}$$

As a forbidden transition is very unlikely to occur this is a very long lived transition and

therefore the uncertainty principle tells us the energy is very well defined. As frequency is proportional to energy, the frequency variation will also be small and the natural linewidth will be narrower than an allowed transition.

The laser cooling for neutral strontium is achieved by utilising the 461 nm transition and the 689 nm transition. The 461 nm transition ($^1S_0$ to $^1P_1$) is an allowed transition giving a linewidth of 32 MHz enabling large capture velocities. The 689 transition ($^3S_0$ to $^1P_1$) changes from a singlet state to a triplet state and is forbidden giving a linewidth of 7.5 kHz, which allows deep cooling of atoms previously trapped using the 461 nm transition.

Our clock transition ($^1S_0$ to $^3P_0$) is a J = 0 to J = 0 transition and a change from singlet to triplet, this is known as a doubly forbidden transition and has a much narrower natural linewidth than the forbidden transition. This is precisely why we use the clock state because it is so well defined in energy. For our clock the chosen atomic species, neutral strontium, the clock transition has a linewidth of ≤ 10 mHz [11].

2.3 Optical Clocks

To better understand how we can improve the instability of our clock transition measurement we need to understand what determines the instability of our measurement. If the measurement is not limited by any part of our measurement apparatus, the instability $\sigma(\tau)$ of an atomic clock is given by the formula

$$\sigma(\tau) = \frac{\Delta\nu}{\pi\nu_0}\sqrt{\frac{T}{N\tau}} \tag{2.2}$$

where $\Delta\nu$ is the linewidth of the clock transition, ν_0 is the frequency of the clock transition, T is the cycle time required to take a measurement, N is the number of atoms being interrogated and τ is the averaging time [12].

We wish to use an optical transition that has a higher frequency than a microwave tran-

sition because the higher the frequency we use the smaller instability we will have. Optical atomic clocks use either trapped ions or laser cooled clouds of neutral atoms as a frequency reference. Ion traps have shown to be very good frequency standards but as can be shown from equation 2.2 a neutral atom clock with several thousand atoms will have a smaller instability than a clock with a single ion or very few ions. In principle better stability can be gained from using large clouds of neutral atoms where many atoms can be averaged over. However ions can be confined to much smaller volumes such that their environment can be more accurately controlled; this leads ultimately to better accuracy but at the cost of larger averaging times.

When a photon is emitted from an atom the frequency is dependent on the velocity of the atom due to the Doppler effect. If an atom is moving towards the detector when a photon is emitted the frequency measured would be higher that if the atom were at rest; if the atom was moving away the frequency measured would be lower. Because thermal atoms have a range of velocities the precise frequency measured is blurred out because the Doppler shift makes the measured frequency much broader than the natural atomic transition. If we reduce the velocities of the atoms the frequency emitted will not have such a wide frequency broadening due to the Doppler effect. The velocities of particles in a gas are described by the Maxwell-Boltzmann distribution; the mean velocity is proportional to the square root of the temperature of the gas, by cooling the gas down the particles travel slower. This reduction in temperature can be achieved using laser cooling.

Our experiment uses neutral strontium atoms that are laser cooled. By preparing the atoms into a particular state the clock transition frequency of the atoms can be probed using fluorescence detection [13]. As the linewidth of the energy transition we will use in strontium is very narrow we find that the precision of the frequency measurement is limited by the frequency stability of the laser. The laser also acts as a "flywheel" to maintain the frequency in between interrogations of the atoms.

To make the system into a clock in the future a frequency comb will be required to make a comparison between the microwave frequency standard and our optical frequencies being measured [14]. A frequency comb uses a femtosecond mode-locked laser to produce a continuous stream of very short pulses of light; the spectrum is a comb of narrow peaks in frequency space with spacing that is equal to the repetition rate f_{rep} at which pulses are emitted from the mode-locked laser. By frequency broadening the input pulses in an air-silica microstructured optical fibre, more modes are produced with frequency spacing of f_{rep} until an octave-spanning comb is reached. The frequency of the n^{th} mode of the comb is given by

$$f_n = nf_{\text{rep}} + f_0 \tag{2.3}$$

where f_0 is a frequency offset from zero. To accurately know the frequency of each mode we need an accurate knowledge of both the repetition rate and the offset frequency. If the frequency comb spans a whole optical octave, for example from 1064 to 532 nm, we can determine f_0. This is achieved by using a frequency doubling crystal to produce a second harmonic of the 1064 nm mode ($2nf_{\text{rep}} + 2f_0$) and beating it with the 532 nm mode ($2nf_{\text{rep}} + f_0$). Beat measurements are discussed at length in section 3.4. The beat frequency produced will be the difference in frequency between the two modes and from this f_0 can be calculated. The RF beat frequency can be compared with a stable microwave oscillator thereby making a connection between the optical frequency reference and a microwave frequency standard.

2.4 The Definition of the Second

As the energy of a clock transition is fixed and has a specific frequency associated with it, this frequency can be used to define how long a second is by counting the number of oscillation in one second.

The International System of Units (SI) has the second as an integral part of the system

as many other definitions of units include the second within their definition [15]. The first atomic frequency standard was based on rubidium using a microwave frequency that had an instability around 10^{-10}; equivalent to being correct to one second in 300 years. The early atomic clocks were already more stable than mechanical clocks at the time that had an accuracy close to one second a year [16].

The second is currently defined as "the duration of 9 192 631 770 periods of the radiation corresponding to the transition between the two hyperfine levels of the ground state of the caesium-133 atom" [17]. In 1967 when the second was defined the measurement was not as precise as today, primarily because the caesium atoms were not cooled. The precise microwave transition can be found in caesium using state selection and observing the Ramsey fringes; by detecting alternately on either side of the central fringe with a microwave oscillator the frequency can be fixed to the atomic reference and the oscillator will maintain its stability.

The caesium clock has been constantly improved however and is reaching the limits of what precision can be reached. If we wish to have more accurate clocks a new standard for the second needs to be developed. The main consideration in choosing an atomic system for a clock reference is the characteristics of the energy levels we use; we wish to use an element that has energy levels that maintain the energy difference between them as a change in energy will change the associated frequency. Energy levels can be shifted by magnetic or electric fields therefore we require an element that is minimally perturbed by these. Broad atomic transitions will have linewidths of several megahertz; however the atomic transition we wish to use can have linewidths of less than 10 mHz.

There are a number of other frequency standards that have been defined as secondary representations of the second [18]. These standards include ionic species such as mercury [19], ytterbium [20] and neutral atom clocks such as strontium [21]. There are also highly precise frequency standards such as ionic aluminium that have very high reproducibility [22]. All these elements have an atomic transition that is extremely well defined with minimal

change in frequency due to external perturbations and that is used as the clock transition.

An aluminium ion optical clock with a 1.121 PHz clock transition has been shown to have a fractional frequency inaccuracy of 8.6×10^{-18} [23]. However, the current stability record for a frequency standard is held by NIST; using spin-polarized, ultra cold atomic ytterbium two clocks with an instability of 1.6×10^{-18} after only 7 hours of averaging have been demonstrated [24]. Our atomic clock will use strontium-87 that has a clock frequency of 429 228 004 229 873.65 (37) Hz according to the best current measurements of the standard [25]. A comparison of two strontium clocks have shown fractional instabilities of 1×10^{-17} [26].

2.5 Laser Cooling

Laser cooling has given us the ability to cool neutral atoms to temperatures near to absolute zero by using a magneto-optical trap (MOT). Since the time when laser cooling was developed in the 1980s [27] the field has grown to be a major area of research; the 1997 Nobel Prize was awarded to the creators of laser cooling [28].

The basis of laser cooling is that photons of light carry momentum; if a photon hits an atom and the photon is absorbed then the atom changes its velocity due to the change in momentum. If the atom and photon are travelling in opposite directions when they collide the atom will slow down. This is analogous to a bowling ball being slowed down by shooting ping-pong balls at it, each ping-pong ball will slow the bowling ball slightly but with enough ping-pong balls the bowling ball will eventually stop. This analogy is wrong in two ways, firstly to be completely analogous the bowling ball would have to be 300 million times heavier. Secondly, as already discussed, the only way an atom can absorb the photon is if the energy of the photon is equal to the difference in energy of two energy levels within the atom. We must use the correct frequency laser to perform the laser cooling on a specific transition

within an atom. The frequency of light used must account for the Doppler effect and be at a slightly lower frequency than required for a photon to be absorbed if the atom were at rest; this is known as red-shifted. This ensures that atoms moving towards the laser will absorb photons with a higher probability than atoms moving away from it. If an atom moves towards the laser the frequency of the light will appear to be higher therefore it will absorb the photon and receive a momentum kick to slow it down. Once the photon is absorbed it must then be emitted. The release of this photon is by spontaneous emission, causing the photon to be emitted in a random direction; if we average over lots of photon emissions the momentum change has a net value of zero. To cool the atoms six laser beams are used from three orthogonal directions [29]. This reduces the individual velocities of the atoms and as the temperature of the system is determined by the average velocity spread of the entire ensemble of atoms as they are individually slowed down we can consider the temperature of the system to become lower.

Strontium was chosen for our experiment because the cooling methods and techniques are very well known and relatively easy to achieve [30]. A partial diagram of the strontium energy levels and the associated atomic transitions is shown in figure 2.1. The first stage cooling is achieved using a 461 nm laser for Doppler cooling and this is followed by a secondary cooling laser at 689 nm. The atoms are then stored in a far off-resonant optical lattice at 813 nm. The clock transition is at wavelength of 698 nm. For the cooling transitions the laser linewidths must be similar to the linewidth of the transition used; this will require stabilised lasers. For our detection laser more work is required to produce an extremely narrow linewidth laser.

To trap the atoms spatially we require a magnetic field; by using two coaxial parallel coils carrying current in opposite directions a quadrupole field is produced. The field is zero in the middle and increases as you move away from the centre. This magnetic field causes Zeeman splitting of the J = 1 state into three levels $m_j = -1, 0$ and $+1$. The transition energy will change as we move further away from the middle of the trap; the light used to

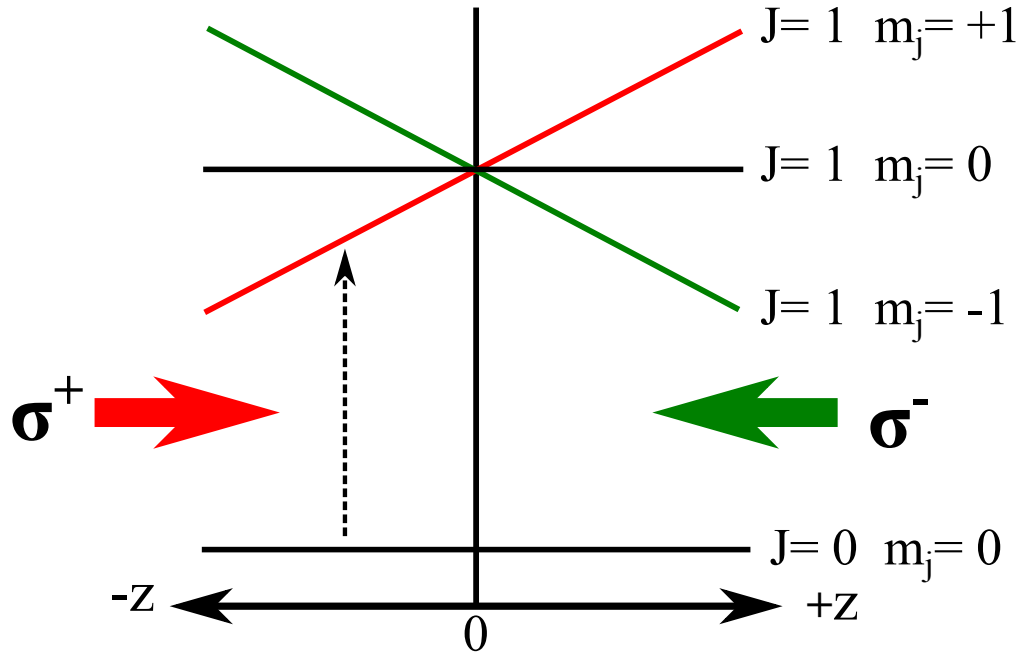

Figure 2.2: Arrangement of a MOT in one dimension.

trap is circularly polarised with light coming from one direction having σ^+ and from the other direction having σ^-. The arrangement for this set up is shown in figure 2.2 for one dimension. The quantity $\sigma^\pm$ is defined as causing a transition with a change in the magnetic quantum number $\Delta m_j = \pm 1$. The frequency of the light used will be red-detuned, when an atom moves from the centre of the trap the m_j energy levels of the excited state will be shifted such that the atoms interact more with the light that corresponds to the polarisation of light to drive the atom towards the centre of the trap.

Ultra High Vacuum (UHV) technology must be used to produce a vacuum with extremely low pressures at the level of $< 10^{-10}$ mbar to produce an environment free from any other atoms which could interact with the atoms and make it impossible to cool them. A vacuum chamber with good optical access is required as several lasers are used during the experiment.

2.6 Optical Lattice

The atoms in a MOT will still move but far less than the thermal atoms. To further reduce the movement of the atoms we make an optical trap using a high power laser beam reflected back on itself to interfere and make a standing wave of light. Atoms will be trapped in areas of high intensity in the beam due to a dipole force. If the lattice potential is sufficiently high we find that when an atom emits a photon the atom cannot gain enough energy to move around in the trap and hence the atom does not recoil. Therefore there will be no Doppler shift of the emitted photon; this is known as the Lamb-Dicke regime.

The most accurate neutral atom measurements have been achieved with the atoms stored in an optical lattice with a magic frequency where the frequency of the lattice is chosen to not perturb the energy difference between the levels of the clock transition [31].

CHAPTER 3

STABILISED DIODE LASER SYSTEMS

3.1 Introduction

For the optical clock several lasers will be required at different wavelengths that are needed to cool, trap and interrogate the atoms. These lasers require narrow linewidths to probe the transitions of the atoms. For the laser interrogating the clock transition we need a laser with a very narrow linewidth, which can be achieved by locking it to an ultra stable high finesse optical cavity. However to lock to this optical cavity we need a sufficiently stable laser to probe a resonant mode of the optical cavity. The development of such a stable laser for our experiment will be discussed in this chapter. Since we wish to build a portable optical clock we require the laser system to be portable and that it does not require vast amounts of cooling or other paraphernalia that would be required for a laser such as a Titanium-Sapphire laser or gas laser. The clear choice is to use a stabilised diode laser system that is compact, requires electronics that are portable and has low power consumption.

Throughout this thesis the terms laser frequency ν and laser wavelength λ will be used depending on the particular situation. The wavelength is expressed as $\lambda = c/\nu$ so the conversion is very simple and one can be considered to be equivalent to the other for the purpose of this thesis. To give some feeling for numbers the frequency of the clock lasers built

for this experiment is 429 THz; this corresponds to a wavelength of 698 nm. The variation in the stability will be expressed in frequency, such as a hertz level change of $\Delta\nu = 1$ Hz. The conversion of change is given by $\Delta\lambda = (-c/\nu^2)\Delta\nu$. Therefore the change in wavelength equivalent to 1 Hz will be 1.6×10^{-21} m, otherwise written as 1.6 zeptometres. To describe a laser its wavelength is given in nanometres and then its stability will be quoted in units of frequency, as is standard by most optical manufacturers.

3.2 External Cavity Diode Lasers

A diode laser contains two types of semiconductor, one area that is rich in electrons and one area lacking electrons [32]. When an electron and a 'missing electron' known as a hole combine a photon of light is released. To use a diode laser we require a current to be applied across the two areas of semiconductor. Diode lasers can have the frequency of the light they emit changed by altering the current supplied to the diode or by changing the temperature of the diode. A free running diode laser will emit a large range of frequencies with a spectral width of around 1 GHz. Lasers are often described as being monochromatic, that is only producing one colour of light; even though the range of frequencies emitted by a free running diode laser is very small it will have too wide a frequency range to probe an optical cavity resonance or many of the atomic transitions used for cooling of atoms.

We therefore need to stabilise a diode laser by a technique of feeding light back into the laser diode by using additional optics. A diode laser has a very wide linewidth because of the small area the light is emitted from; the length of the semiconductor is several 100 μm so the coherence length is very short, and this gives a wide laser linewidth. By making an external cavity outside the semiconductor the spatial coherence increases, producing a narrower linewidth laser; this is achieved by reflecting the beam back into the laser diode and this optical feedback will cause stimulated emission in the laser cavity. Stimulated emission

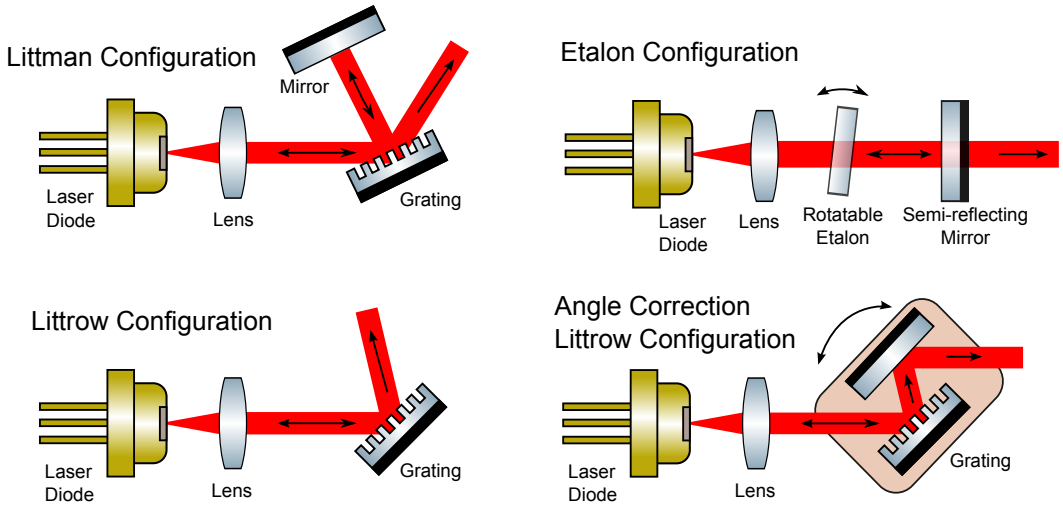

Figure 3.1: Several arrangements that are used for External Cavity Diode Lasers. A conventional casing of laser diode package is shown, inside which would be a small semiconductor that is the emitting light active medium. The beam from a diode laser is divergent and requires a lens to focus the light to a collimated beam. The other optics are used to form the external cavity by directing light back into the diode laser.

is when a photon interacts with an atom and causes a photon of the same wavelength to be emitted.

Possible ways to make a laser diode with an external cavity are by using an etalon or grating to select the frequency wanted; some possibilities are shown in figure 3.1. The design using an etalon configuration uses an intra-cavity etalon between the laser diode and a partially reflecting mirror directing the light back into the cavity; the frequency of the laser can be chosen by rotating the etalon as this will choose the frequency that is resonant in the cavity. This method has been used to produce a laser with a 14 kHz linewidth for a range of laser frequencies [33]. There are several ways to use a reflection grating to stabilise the diode laser. The angle of the grating selects the wavelength and the diffraction of the cavity increases the stimulated emission. The grating will spread a beam of light into its component frequencies by reflecting the light out at some angle depending on its wavelength. For a laser beam which is mostly monochromatic the beam will be split slightly for each reflected order

and the frequency reflected back into the diode can be carefully selected. For a reflection grating which has a line spacing of d the reflected beam is given by

$$m\lambda = d(\sin\alpha + \sin\beta) \tag{3.1}$$

where α is the input angle, β is the output angle of the beam and m is the order of the diffracted beam.

There are two arrangements for grating stabilised diodes: the Littrow configuration and the Littman-Metcalf configuration [32][29]. These configurations are designed so that the frequency of the laser can be carefully tuned by changing the length of the cavity by moving the grating. We can use a piezoelectric transducer to change the position of the grating by very small distances. The Littrow configuration is the most simple as the beam is reflected off the grating straight into the diode; this requires a grating with the required number of lines per millimetre to be selected so that the input angle α is the same as the output angle β. The wavelength of the laser can be chosen by rotating the grating but this also changes the output angle. In the Angle Correction Littrow configuration the output angle change can be corrected for by using a mirror that rotates through the same angle as the grating when it is rotated ensuring the output beam angle does not change. The Littman-Metcalf configuration has the first order beam reflected off a grating onto a mirror which reflects the beam back onto the grating and into the laser diode. The advantage of this design is that the output beam is at a fixed angle if the mirror is rotated; the disadvantage is the increased complexity as an additional mirror is used.

It was decided that the Littrow configuration would be used for our set up due to the simplicity of the design; the additional mirror in the Littman-Metcalf configuration could have produced further vibrational modes that may have increased the laser linewidth. The angle of the beam would change when the grating was turned but the tuning of the laser

frequency would not require major movement in the grating once the central frequency has been set. From equation 3.1 it can be shown that

$$\delta\alpha = \frac{m}{2d\cos\alpha}\delta\lambda \tag{3.2}$$

$$= -\frac{m}{2d\cos\alpha}\frac{c}{\nu^2}\delta\nu \tag{3.3}$$

from which it can be calculated that for a frequency change $\delta\nu$ of 100 GHz the grating will rotate by 1.7×10^{-5} ° for the design used. Later it will be shown that a frequency change of 7 GHz is all that should be required and therefore minimal grating rotation will be needed. The chromatic resolving power of a grating is given by

$$R = \frac{\lambda}{\Delta\lambda} = mN \tag{3.4}$$

where N is the total number of illuminated slits and m is the order of diffraction. Reflection gratings can have around 1000 lines per millimetre. The beam size will usually be less than 10 mm so the maximum chromatic resolving power will be 50 GHz, significantly larger than the linewidth of the free running laser diode therefore having no effect on narrowing the laser frequency.

The theoretical linewidth of an ECDL formed by a grating reflecting light back into the diode is given by the Schawlow-Townes linewidth [34][35][36]. First we introduce the theoretical cavity linewidth $\delta\nu_c$:

$$\delta\nu_c = \frac{c}{2\pi(nL_d + L_{\text{ext}})}(\kappa_L L_d - \ln\sqrt{R_{\text{eff}}}). \tag{3.5}$$

The internal diode transmission is κ_L but we can assume $\kappa_L = 1$, L_{ext} is the external length of the cavity given by the distance from the diode to the grating, L_d is the length of the semiconductor with refractive index n, R_{eff} is the total reflectivity of the cavity, the reflectivity

of the grating back into the diode is R_1, the reflectivity of the back facet of the laser diode is R_2 and $R_{\text{eff}} = R_1 R_2$. From the cavity linewidth we can calculate the Schawlow-Townes laser linewidth Γ_m of the cavity mode m

$$2\Gamma_m = \frac{\pi h\nu(\delta\nu_c)^2}{P_m} n_{\text{sp}} \tag{3.6}$$

where P_m is the power in mode m and n_{sp} is the number of spontaneously emitted photons in the mode; n_{sp} tends to 1 when the laser diode is operated above the threshold current for the laser diode. From equations 3.5 and 3.6 we can tell that by increasing the length of the ECDL cavity the laser linewidth decreases proportionally to the square of the length of the extended cavity. However if the mechanical stability of the longer cavity is significantly worse when we make the cavity longer this causes the linewidth to broaden due to vibrations in the laser system changing the length of the external cavity. Therefore it is believed a rigid and compact design is required for the linewidth to be minimal [37].

3.3 Mechanical Design of the ECDL

We endeavoured to design a mechanically stable laser that could be made intrinsically more stable by making a longer cavity. The chosen design uses a monolithic block of aluminium to mount the optical components, similar to a design presented in [38]. The design developed has support for the laser diode, precision collimation lens positioning and a commercially available optical mount built into the side of the monolithic block [39]. The monolithic design used is shown in figure 3.2. The ECDL is required to be as rigid and isolated from acoustic noise as possible, this is achieved when a lid is sealed down onto the monolithic block. To stop air flow all holes are glued over and a window is glued onto the box. To position the lens solidly a fine thread lens tube is used to position the lens with the focus of the lens on the laser diode to produce a collimated beam; once a locking ring is tightened

the lens mount becomes fixed and very stable.

The diode laser used was a custom wavelength-selected diode specifically for single mode operation at 698 nm made by Sacher Lasertechnik [40]. The semiconductor was anti-reflection coated to ensure that the laser facets did not act as an optical cavity; this should stop multi-mode operation and improve the wavelength tuning range. No multi-mode operation was observed when this laser diode was used. Previously non anti-reflection coated laser diodes have been used, which often displayed multi-mode operation. The semiconductor was supplied mounted in a standard 9 mm diameter laser diode package; the maximum current that could be supplied to the diode laser was 100 mA.

A reflection grating with 1800 lines per millimetre was glued to the face of the optical mount. The mechanical actuators of the optical mount were used to angle the beam back into the diode. To allow fine horizontal movement, a low voltage piezo stack actuator [41] was mounted between one of the actuators and the face of the optical mount. The piezo used was 5 mm $\times$ 5 mm and 2 mm deep and it has a range of 3 μm when 50 V is applied, therefore we are able to change the laser frequency over 5 GHz without having to use high voltage electronics as required by many piezo actuators.

The free running diode laser frequency depends on the current supplied and the temperature of the laser. To have a diode laser with long term stability we need to keep the temperature of the semiconductor constant; either the laser diode could be temperature controlled or the entire ECDL monolithic block could be temperature controlled. It was thought that by stabilising the temperature of the entire ECDL the external cavity would have minimal changes in length due to thermal expansion so the long term stability of the ECDL would be better. The temperature is stabilised using a peltier element between the ECDL and the aluminium base of the laser system. The base is part of a containing box housing the ECDL and the other optics required; the base is a 20 mm thick aluminium plate making it more than sufficient to act as a heat sink for the peltier element. The temperature was detected with

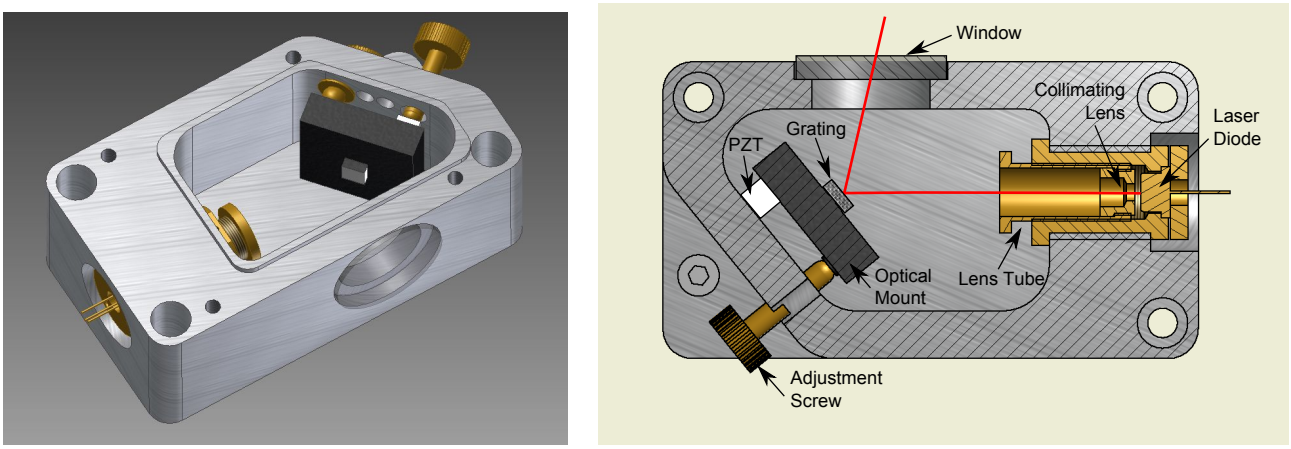

Figure 3.2: Mechanical design of the External Cavity Diode Laser; to give a scale a laser diode has an outer diameter of 9 mm and the distance between the diode and the grating is 50 mm. The right hand image is a section cut through of the middle of the monolithic block of aluminium.

an AD590 temperature sensor glued to the side of the ECDL at a position above the peltier element. A home built analogue PI controller was used to feedback to the peltier; the controller was based on the WTC3243 Temperature Controller built by Wavelength Electronics [42].

3.4 Measuring the Laser Linewidth

To characterise the linewidth of lasers a beat measurement needs to be made by heterodyning the beams onto a fast photodiode. The interference of the two laser frequencies ν_1 and ν_2 causes a beat note to be produced. Assuming we have two independently running lasers with equal intensity ${E_0}^2$; the electric field incident on the photodiode will be

$$E = E_0 \cos \nu_1 t + E_0 \cos \nu_2 t \tag{3.7}$$

and we can write the measured intensity I as the square of the electric field

$$\begin{aligned} I &\propto {E_0}^2 (\cos^2 \nu_1 t + \cos^2 \nu_2 t + 2 \cos \nu_1 t \cos \nu_2 t) \\ &\propto {E_0}^2 (\cos^2 \nu_1 t + \cos^2 \nu_2 t + \cos(\nu_1 + \nu_2) t + \cos(\nu_1 - \nu_2) t) \end{aligned} \tag{3.8}$$

The high frequency signal, anything except $\nu_1 - \nu_2$, will not be detected because it is not possible to build electronics with a fast enough response to measure terahertz frequencies. The high frequency signals will only be measured as an offset on the detector. The measured signal will be an offset plus the oscillating intensity of the hetrodyned electric field. The measured signal will be of the form

$$I \propto 1 + \cos(\nu_1 - \nu_2) t. \tag{3.9}$$

The beat note is defined as $\Delta\nu = |\nu_1 - \nu_2|$; the two laser wavelengths need to be similar enough such that the beat note frequency can be measured by our electronics. By measuring the width of the beat signal we can measure the variation in frequency between the two lasers. The linewidth of the beat note will be wider than the linewidth of each laser separately. If we assume the lasers are separate and independent, such that there is no feedback from one to the other, and we assume the laser linewidths are the same size we can deduce the individual laser linewidth. By error propagation the individual laser linewidths can be calculated from the beat note linewidth; $\delta\nu_1 = \delta\nu_2 = \frac{1}{\sqrt{2}}\delta\nu$.

3.5 Experimental Linewidth Determination

One of the ECDLs built has a cavity length between the diode and the grating of 50 mm and the other has a length of 100 mm. If we consider equation 3.6, by doubling the length of the cavity the linewidth should be reduced to 25% of the linewidth of the shorter laser. The

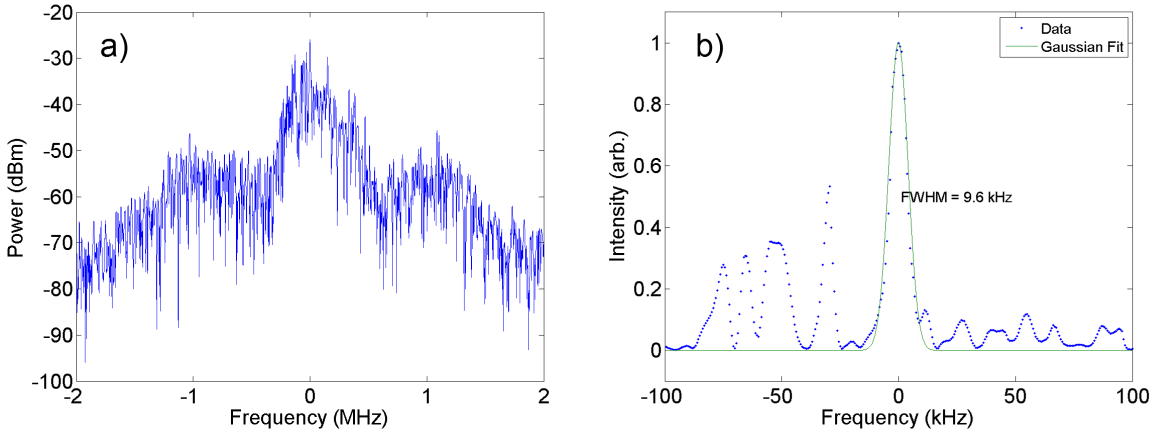

Figure 3.3: A single set of spectrum analyser data of the free running long (100 mm) ECDL laser system shown on two different scales. a) is the full width of the frequency scan in log scale. b) is focused in on the central peak in linear scale and the fitted gaussian measuring the linewidth is shown on the plot.

lasers were aligned onto a fast photodiode; the photodiode had a bias voltage applied using a bias-tee with the output being connected to a Rohde and Schwartz FPS spectrum analyser with a maximum frequency span of 7 GHz. The lasers could be brought sufficiently close in frequency using a wavemeter to measure the beat frequency of the lasers; when the frequency difference was less than 7 GHz apart the spectrum analyser could be used to examine the linewidth of the beat note. The beat note can be compared when the lasers are unlocked or when one of the lasers is locked to a high finesse optical cavity; as will be shown later the linewidth of the locked laser is *significantly* less than the unlocked laser, therefore we can assume that the measured beat note is almost entirely due to frequency fluctuations in the unlocked laser. By comparing the measured linewidth of both the lasers we can determine whether the longer laser has a significantly better intrinsic stability. To reduce the effect of vibration on the laser systems they were both mounted on separate minus-K passive vibration isolation platforms [43].

Full details of these results have been published elsewhere [44]. The linewidth of the beat frequency was calculated from the spectrum analyser data by fitting a gaussian to the

highest peak in the linear data and measuring the width of this gaussian at half the maximum height. Figure 3.3 shows an individual set of data from the free running long laser. On 3.3 b subsidiary peaks can be seen, these peaks will be caused by vibrations shifting the frequency of the laser. If the measurement time were much longer the random spread in the peak frequency would produce be a wider linewidth as the spread in peaks broaden the central narrow peak. All of the results were taken with a frequency span of 5 MHz and a resolution of 10 kHz. The final result showed that when the long laser system was locked to the optical cavity and the short laser system (50 mm) was free running the measurement had a linewidth of 39 ± 23 kHz. When the short laser system was locked to the optical cavity and the long laser system (100 mm) was free running the measurement had a linewidth of 26 ± 15 kHz. Therefore we can determine that the longer cavity is more stable, but not by the amount expected. We would expect that by doubling the length of the cavity the stability would improve by the square of the length change; a fourfold improvement would be expected. However the longer design is clearly not mechanically stable enough for this to happen but does give us an improved instability. Our result is in reasonable agreement for linewidth reduction with similar laser system designs built for mechanical stability [45].

CHAPTER 4

ULTRA STABLE FABRY-PÉROT CAVITIES

4.1 Introduction

Fabry-Pérot optical cavities are used in atomic physics to stabilise lasers to make the linewidth of a laser narrower so that it can be used to probe narrow atomic transitions. By using an ultra stable optical cavity, with extremely stable length, we can produce an ultra narrow linewidth laser that can probe the very narrow clock transition required for an atomic clock. Ultra stable optical cavities have become common use in research groups building optical clocks. How we lock our laser to the optical cavity is very important; the method used is the Pound-Drever-Hall method and how we use this technique is explained. Presented here are simulations of intrinsic thermal noise of the optical cavities and how the magnitude of the noise changes depending on the materials the optical cavities are made from. The sensitivity to vibration of a particular design of optical cavity is also modelled and the most effective theoretical supporting position is found.

4.2 Fabry-Pérot Optical Cavity Theory

The first people to use two highly reflective surfaces as an etalon were Charles Fabry and Alfred Pérot in the late 1800s [46]. A conventional Fabry-Pérot etalon consists of two parallel reflective mirrors set a certain distance apart from each other; this is also known as an optical cavity. An optical cavity will usually comprise of a spacer made out of a low thermal expansion material that has a mirror fixed to each end of the spacer. One or both of the mirrors will have some concave curvature so that the cavity is stable; this allows a Gaussian beam to propagate in the optical cavity. The optical cavity acts as an interferometer, light enters through one of the mirrors, is reflected back off the other mirror, reflected again of the first mirror and so on. If the reflected light after one round trip is in phase with the incoming light they will interfere constructively, if they are π out of phase they will interfere destructively. The condition to have constructive interference is that there be an integer or half-integer number of wavelengths in the length of the cavity L; the n^{th} longitudinal mode will have a wavelength λ where $\lambda = 2L/n$. This gives us our free spectral range (FSR), which is the frequency separation between the frequencies of consecutive modes that will be transmitted through the optical cavity; an example is shown in figure 4.1. The frequency separation between resonant modes is $\nu_{\text{FSR}} = c/2L$ where c is the speed of light. The range of frequencies around the resonant mode that can be transmitted through the cavity is determined by the reflectivity of the mirrors. For a low reflectivity each photon will spend less time in the optical cavity, therefore the coherence of the light will be less and a wide range of wavelengths will be transmitted. However for very high reflectivity each individual photon will spend longer in the cavity so the coherence will be greater; the range of allowed frequencies will therefore be narrower giving very narrow resonant modes. The range of frequencies transmitted is often measured as the linewidth $\Delta\nu$ and is defined as the width of the transmitted frequency peak at half of the maximum transmitted intensity or full width half-maximum (FWHM). The finesse $\mathcal{F}$ is the ratio between the free spectral range and the

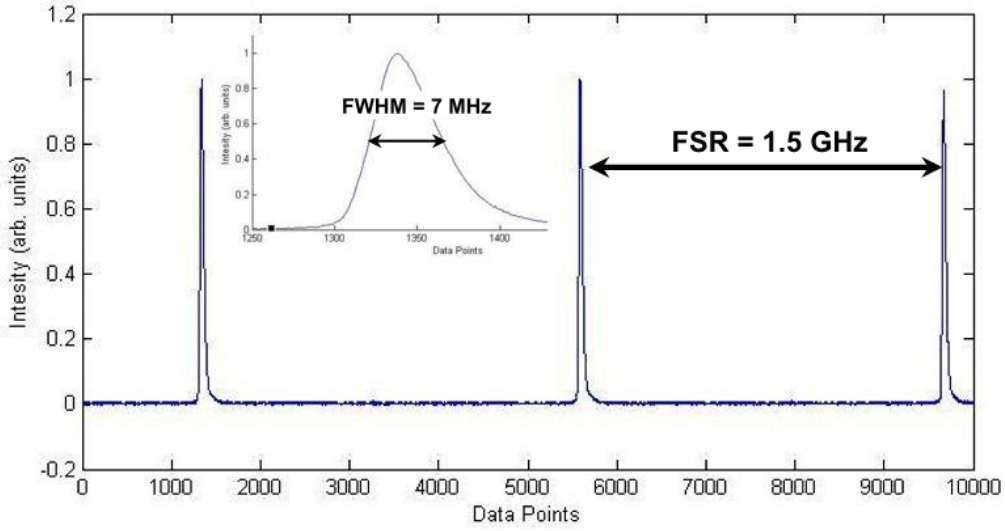

Figure 4.1: The transmission spectrum of a Fabry Pérot Cavity with a finesse of 200. The measurement was done by changing the length of the Thorlabs optical cavity described in table 4.1. The cavity length was changed by having one of the mirrors mounted on a piezo-electric transducer and changing the voltage supplied. The 'data points' are an arbitrary measure of the linear change in length of the cavity. The insert is a magnified image of the transmission peak. The asymmetry observed in the shape of the linewidth is believed to be due to higher order modes being transmitted by the optical cavity.

linewidth. The higher the finesse we have for our cavity the narrower the resonant modes and the higher resolution frequency discrimination we have for light passing through the optical cavity. For mirrors with a high reflectivity R the finesse is

$$\mathcal{F} = \frac{\nu_{\mathrm{FSR}}}{\Delta\nu} \approx \frac{\pi\sqrt{R}}{1-R} \tag{4.1}$$

We wish to use an optical cavity to reduce the frequency variations and therefore reduce the linewidth of our laser's frequency by electronic feedback to the laser frequency control. An optical cavity will change in length due to vibrations, thermal expansion and fluctuations in air pressure so all of these things must be controlled.

To measure the stability of an optical cavity we require two full set ups. If we wanted to measure the frequency range of a radio frequency source we could use a spectrum analyser

	Thorlabs Cavity	ATFilms Cavity
Arrangement (Curvatures)	Confocal (50 mm,50 mm)	Plano-concave (1 m,∞)
Optical Cavity Length	50 mm	100 mm
Free Spectral Range ν_{FSR}	1.5 GHz *	1.5 GHz
Finesse $\mathcal{F}$	200	400 000
Linewidth $\Delta\nu = \frac{\nu_{\mathrm{FSR}}}{\mathcal{F}}$	7.5 MHz	3.75 kHz
Mirror Reflectivity	94%	99.998%

Table 4.1: A comparison of parameters for the optical cavities used. The Thorlabs cavity is a SA200-Series Scanning Fabry Perot Interferometer, the ATFilms cavity is Notched Spacer Cavity Assembly ATF 6020-01. All data has come from the data sheets for the relevant cavities or values calculated from those numbers. *The 50 mm Thorlabs cavity has a FSR of 1.5 GHz because the arrangement is confocal and a closed light path takes two round trips.

directly but this cannot be done with light as the best electronics are not fast enough to analyse frequencies higher than several gigahertz, and visible light frequencies of hundreds of terahertz are completely inaccessible. By having two lasers we can look at the difference in frequency between them and measure the width of the frequency comparison; this is known as a beat measurement.

4.3 Laser Locking to an Optical Cavity

Our aim is to improve the stability of our laser by locking to an optical cavity. We use the resonant modes of the optical cavity to fix the frequency of the laser. We can use one of these resonances to give electronic feedback to the laser frequency control; we can electronically reduce the linewidth using a closed-loop servo system. We must also consider the bandwidth of our locking electronics; the bandwidth is linked to the range of frequencies we are able to correct for.

To lock to the optical cavity one method we could use would be to set the laser frequency to maintain maximum transmitted intensity of the light through the optical cavity. However when the laser frequency moves off resonance we would only measure a reduction in transmitted intensity and be unsure in which direction we needed to correct the frequency.

If we locked at a point at half of the maximum transmitted intensity we could tell in which direction correction is required. If the intensity were to increase we could feedback to change the frequency and to reduce the intensity; if the intensity were to decrease we could change the frequency in the opposite direction to increase the intensity. However any changes in amplitude of the input light would also cause the frequency to be adjusted. Instead of these methods we used the Pound-Drever-Hall locking technique.

4.4 Pound-Drever-Hall Locking

To lock to an optical cavity we use a method known as Pound-Drever-Hall locking (PDH) [47]. There are very good reviews on this subject [48][49] that describe the method in a more coherent and detailed way. The generic set up for PDH is shown in figure 4.2. We require frequency sidebands on our light, these sidebands can be produced by either frequency or phase modulation of the laser beam. The method shown in figure 4.2 is phase modulation where the laser beam passes through a phase modulating electro-optic modulator (EOM); by modulating the light we produce sidebands on the light at the frequency of the oscillations applied to the EOM. The size of these sidebands is controlled by the amplitude applied to the modulator. Sidebands can also be produced when using a diode laser by modulating the current supplied to the laser diode but that also modulates the laser's intensity. When the laser beam with its carrier and sidebands are incident on an optical cavity the reflected sidebands and the carrier will experience a phase change, depending on their relative position with respect to the resonance. By looking at the reflected light from the cavity and mixing this signal with the original modulation frequency we can see the phase change of the reflected light; this is known as the error signal. The reflected light is used because the light transmitted through the optical cavity has very little of the sidebands when a high finesse cavity is used and a very small error signal would be observed.

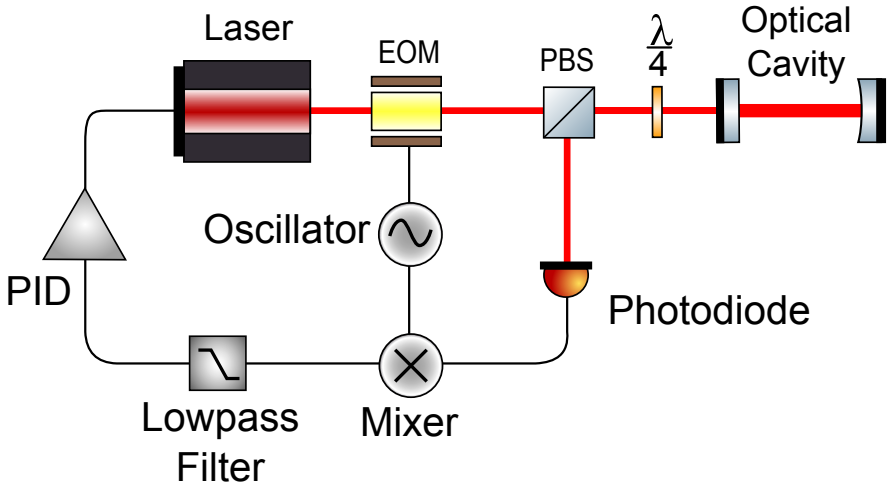

Figure 4.2: Generic set up for Pound-Drever-Hall Technique.

So that we can use the reflected light from the optical cavity, we use polarisation to change the path the reflected light will take. After the light has passed through the EOM it is then incident on a polarising beam splitter (PBS) followed by a quarter wave plate ($\lambda/4$) that turns the linearly polarised light into circularly polarised light. The light is reflected off the cavity and the sidebands interfere with the carrier; the reflection also reverses the circular polarisation of the light. As the circularly polarised light is travelling with the opposite rotation when it passes through the quarter wave plate it will be linearly polarised orthogonal to the incoming beam, the light will not be transmitted but will now be reflected by the polarising beam splitter and be incident on the photodiode. The signal from the photodiode is mixed with the phase modulation signal so differences in phase can be detected; this is the error signal. The error signal is fed into a PID controller linked to the laser current controller that modifies the current to bring the laser back to the resonant frequency for the cavity.

The modulation of the laser frequency or carrier $\omega = 2\pi\nu$ with an RF frequency Ω causes sidebands to be created at $\omega - \Omega$ and $\omega + \Omega$. The electric field of the light after the EOM

with a modulation strength of β will be

$$E = E_0 \left(J_0(\beta) e^{i\omega t} + J_1(\beta) e^{i(\omega+\Omega)t} - J_1(\beta) e^{i(\omega-\Omega)t} \right). \tag{4.2}$$

The size of the sidebands are given by Bessel functions denoted by J_α where α is an integer denoting which order of Bessel function. $J_0(\beta)$ is the proportion of the power in the carrier frequency and $J_1(\beta)$ is the power in the first order sidebands. The power in the carrier is $P_c = J_0^2(\beta) P_0$ and the power in the side band is $P_s = J_1^2(\beta) P_0$, where P_0 is the total power of the laser beam. In this equation higher order Bessel functions have been disregarded as they are usually small and are not required for the production of the error signal.

We wish to measure the phase of the reflected light however this cannot be measured directly. By the interference of the carrier and sidebands in the reflected light there will be a beat pattern. The phase of the beat pattern will tell us the phase of the reflected beam; this is known as the error signal. The error signal is given by

$$E_{\text{ref}} = E_0[F(\omega) J_0(\beta) e^{i\omega t} + F(\omega+\Omega) J_1(\beta) e^{i(\omega+\Omega)t} - F(\omega-\Omega) J_1(\beta) e^{i(\omega-\Omega)t}] \tag{4.3}$$

where the reflection coefficient is the ratio of the incident and reflected light $F(\omega) = E_{\text{ref}}/E_{\text{inc}}$, which is dependent on the phase of the light. When close to resonance there will be a phase shift of the carrier, the sidebands would be far of resonance and therefore experience no phase shift. When the frequency of incident light is on resonance there will be no reflection and no phase shift of the carrier; in this case there will be an error signal of zero because only the two sidebands are reflected and as they have opposite phase they will cancel out. When the frequency is slightly off resonance there will be a non-zero signal and the sign of the error signal will indicate in which direction a correction to the laser frequency is needed. The transmitted intensity and the error signal from the experiment when the laser frequency is scanned over a cavity resonance is shown in figure 4.3.

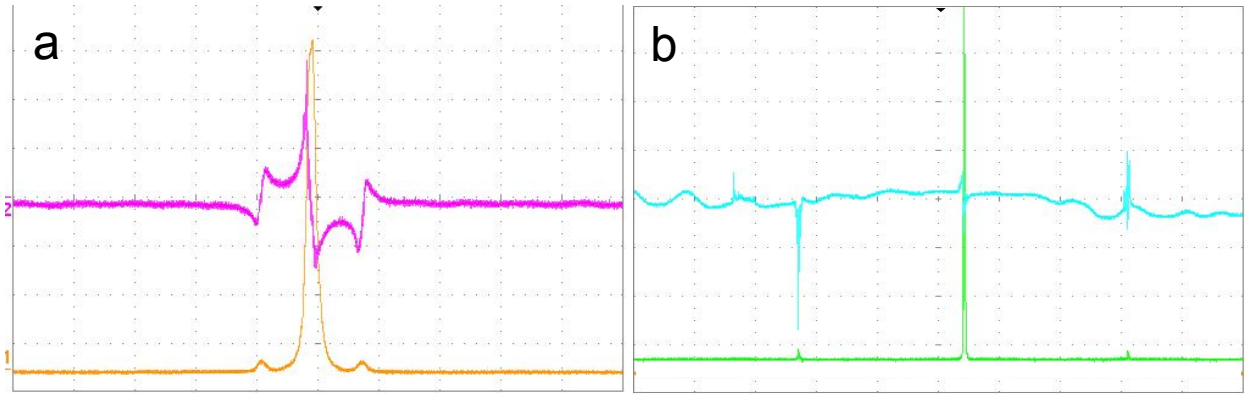

Figure 4.3: Oscilloscope traces of the optical cavity resonance. a) Using the Thorlabs optical cavity with a finesse of 200 we see the transmission signal from the cavity (yellow) and the error signal (pink). The sidebands at a modulation frequency 10 MHz on are shown either side of the main carrier. The error signal is the change in phase of the light so as we pass through a peak the phase changes; we lock to the point where the the error signal crosses zero. b) Showing the transmission signal and the error signal for the ATFilms high finesse optical cavity; the higher finesse gives a significantly narrower linewidth of the cavity and a significantly steeper error signal. Details of the optical cavities used are shown in table 4.1

.

Given fast enough electronics with a wide bandwidth and low noise it is possible to lock to 1 part in 10,000 of the linewidth of an optical cavity. Electronics for PDH will easily reach a bandwidth of several hundred kilohertz. The slope of the error signal is determined by the linewidth of the cavity and the linewidth is determined by the finesse, therefore we can increase the finesse to increase the resolution of the cavity. Modulation strength has an optimal value to give a maximum magnitude of error signal (figure 4.4) and operating at this point will give the electronics the maximum sensitivity to a frequency change. However modulation of the phase can also cause a change in amplitude therefore it may not be best to use the maximum error signal. It may be beneficial to use lower modulation strength and reduce the amplitude modulation. PDH does not overcome the problem that if the optical cavity changes in length the frequency of the laser will also change.

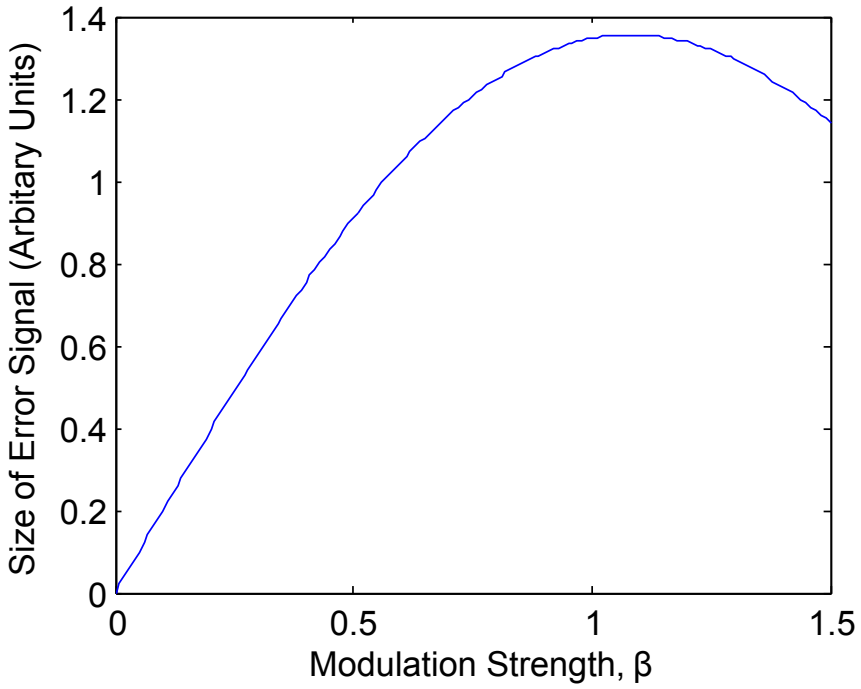

Figure 4.4: The graph shows the modulation strength applied against the magnitude of the error signal. It shows a maximum at $\beta = 1.08$ that corresponds to a side band peak height of 40% of the height of the main carrier. This number is independent of the finesse of the cavity or the modulation frequency, assuming the modulation frequency is much greater than the linewidth of the optical cavity. These figures were calculated using the formulae found in [48].

4.5 Ultra-stable High Finesse Cavities

We wish to produce a laser that not only has a narrow linewidth, such that the laser does not change frequency on short time scales, but also has a well-defined frequency that does not drift over long time scales. From [50] we find the relation for the change in frequency $\Delta\nu$ given a change in optical cavity length of ΔL

$$\frac{\Delta\nu}{\nu} = \frac{\Delta L}{L}. \tag{4.4}$$

For an optical cavity where $L = 100$ mm, a change in length of 1 fm corresponds to a change in the laser frequency of several hertz for an optical frequency; ensuring a constant length of the optical cavity is therefore key to frequency stability. The solution is to use low expansion materials such as Ultra Low Expansion (ULE) glass [51] that has a thermal expansion coefficient of $0 \pm 0.10 \times 10^{-9}$ K^{-1} when the material is at a temperature where there is a turning point of thermal expansion.

To have very high optical frequency differentiation we require a very high finesse optical cavity. The mirrors are made from dielectric coatings specifically designed for the wavelength of light being used to reflect as much light as possible. The high reflectivity coating is applied to the mirror substrate that is also made from low expansion materials. By using alternating layers of different refractive index dielectric a coating reflectivity of 99.999% can be produced, this gives a finesse of 200 000 or more. The coating run is a very expensive part of the optical cavity production due to the expertise needed to make these coatings with such high reflectivity. The other side of the mirror substrate is coated with an anti-reflection coating to minimise etalon effects caused by the back surface of the mirror. The mirrors of an optical cavity must make a stable resonator for a gaussian beam, one or both mirrors will have a concave curvature; this is explained in section 5.2. The mirrors are attached to the spacer using a process known as optical contact bonding; this requires both surfaces to be

polished atomically flat and the surfaces pushed together such that the material will bond via intermolecular forces such as the van der Waal force making them a solid mass with no detectable bond.

If there is gas inside an optical cavity there can be pressure changes due to fluctuations in the temperature of the gas. A change in the pressure of the gas changes its refractive index; the light will travel at a different speed in the optical cavity and this will cause a frequency drift of the locked laser. By having the optical cavity in vacuum this alleviates the problem of fluctuations in the refractive index of the gas. A vacuum stability of 10^{-8} mbar is sufficient to limit these fluctuations in order to reach a 1 Hz frequency linewidth [50].

Vibrations affecting the optical cavity span from audio frequencies at 10 kHz to very low frequencies of 0.1 Hz. Efforts must be made to reduce the effect these have on the length of the optical cavity. The very small changes in length due to vibration will affect the laser linewidth by adding its frequency to the laser noise and therefore random fluctuations causing the linewidth to broaden. To reduce the effect of the vibrations an optical table with isolation supports is used; in addition to this a passive or active vibration isolation platform can be used. Optical tables are very good at suppressing the transmission of vibration at frequencies above 10 Hz but often have a resonance around 1 Hz. The acoustic vibrations are suppressed by using an acoustic box to house the experiment [52].

An important aspect of limiting vibration effects is the shape and mounting method of the optical cavity [53]. The apparatus needs to be designed such that any vibrations transmitted to the cavity have a minimal effect on the optical cavity length. If the cavity's axis lies in the horizontal direction the cavity should be supported in a way that only allows the mirrors to move parallel to each other. An alternative method is to mount the cavity with the optical axis in the vertical direction supporting the cavity about its mid-plane.

4.6 Development of the Ultra-Stable Optical Cavity

Many research groups have developed narrow linewidth lasers using ultra-stable optical cavities. The initial development of this technology was for the gravitational wave community; to detect changes in space-time due to gravitational waves very large Michelson-Morley interferometers are used. The interferometer splits a laser beam which is sent down two long arms of the interferometer and reflected back to interfere with each other. By measuring the intensity of the combined beams any change in phase can be detected, the phase change will be due to a change in arm length due to gravitational waves. The gravitational wave detectors require a high stability in both the frequency and phase of the laser beam [54]. To reach the level of sensitivity required to detect gravitational waves using very large interferometers a laser with extremely stable phase is required. For ground-based gravitational wave detectors such as LIGO and GEO the detectors' frequency range is currently in the region of 100 Hz to 10 kHz, whereas for the planned space-based detector LISA the frequency range will reach as low as millihertz sensitivity. To develop a laser for our purposes with a linewidth of several hertz requires technology drawn from both the ground and space-based detectors as the frequency range is between the spectral ranges of both sets of projects. Drawing on experience from these groups is an important step to producing the next level of stable cavities.

Quantum optics groups have used Fabry-Pérot cavities to reduce the linewidth of lasers so that cooling on long-lived transitions with a narrow linewidth could be achieved. However the development of ultra stable cavities for the purpose of optical clocks has been a relatively recent development. The first sub-hertz linewidth laser used cryogenically cooled optical resonators; a linewidth of 0.7 Hz was achieved with a 1064 nm laser for integration times of 20 s corresponding to a fractional stability of 2.3×10^{-15} [55]. In these experiments a sapphire optical cavity was cooled to 1.9 K; at this temperature it was observed that the 30 mm long optical cavity had a thermal expansion coefficient of 5×10^{-12} K^{-1}.

Scientists at NIST in the United States were the first to report a sub-hertz linewidth

for a stabilised laser at optical frequencies [56]. This was done using a dye laser, stabilised to an optical cavity with a finesse of 150 000. The Pound-Drever-Hall technique was used to stabilise the laser to an optical cavity was made from ULE. The cavity was supported by a v-block mounted inside a vacuum system. To damp high frequency vibrational noise the vacuum system was mounted on an optical table suspended from the ceiling on surgical tubing. The table suspended on surgical tubing would act like a pendulum allowing lower frequencies to be transmitted to the optical cavity enclosure but higher frequencies would be damped.

Stable cavity designs for optical frequencies were then developed by the JILA group, also in the United States, and they were able to produce a laser system with a typical linewidth of 1 Hz [57]. They tried various combinations of low expansion materials, such as ULE, Zerodur and fused silica, for the mirrors and spacer of the cavity itself [58].

The JILA group tested a variety of cavity designs and shapes available to minimize acceleration and vibration within the cavity [59]. They used computer modelling to perform finite element analysis on the range of cavity shapes to establish the optimum design. They determined that the best structure was a vertically mounted tapered cavity. This design has been used widely in many of the following developments.

Meanwhile in the UK the National Physical Laboratory (NPL) developed a specialised support system for a horizontally mounted cavity with 'cut outs' [60][61]. This was to minimise vertical low-frequency vibrations. Once again finite element analysis was used to determine an optimal supporting position for the cavity and the design was proven to work experimentally.

The JILA group built a vertically-mounted tapered cavity with cooperation from other research laboratories around the world. This cavity limited noise such that there is an appropriate compromise between acceleration sensitivity and thermal noise. This was developed specifically to probe a transition in atomic strontium [21].

The group at the Max-Planck Institute for Quantum Optics in Germany have developed a 0.5 Hz linewidth laser using a vertically-mounted tapered cavity made from ULE, with a finesse of 400 000 [50]. The cavity was kept in a vacuum system, which was designed to be thermally and acoustically isolated from the surrounding environment. The thermal fluctuations were minimised by using large amounts of thermal insulation and feedback systems incorporating peltier elements to actively control the temperature. The optical cavity was cooled down to the thermal expansion turning point of ULE; for their optical cavity this was 12 °C. It has been shown that simpler systems using a horizontally oriented optical cavity have been used to produce comparable results to the more sophisticated vibration isolation systems [62]. However it is believed that these cavities are limited by the thermal fluctuations of the mirror coatings as predicted by Numata [63].

Ultra stable cavities are now commercially available with full vacuum system and thermal shields. They can be purchased from companies such as Stable Laser Systems [64], which is a company originating from the NIST research group.

The record for the narrowest laser linewidth is a laser system where the optical cavity is a comparatively long spacer (210 mm) made from silicon [65]. The system built by PTB uses a silicon optical cavity because it has significantly lower thermal noise than ULE, but the turning point of its thermal expansion occurs at 120 K. The silicon optical cavity requires cooling with cryogenic techniques using liquid nitrogen. The laser was measured to have a linewidth of less than 40 mHz and an instability of 1×10^{-16} for short time scales. Infrared light must be used with this optical cavity because silicon is opaque for wavelengths below 1000 nm. To use these silicon cavities with an optical clock using strontium, either a frequency doubler or frequency comb is required.

4.7 Thermal Noise in Optical Cavities

As mentioned the research into thermal noise in optical cavities began in the field of gravitational wave detectors where kilometre-sized Michelson-Morley interferometers are used in an attempt to detect fluctuations in space-time. These gravitational waves are due to phenomena such as supernova and collapsing binary stars [66]. One of the limitations of the detector's sensitivity was due to thermal fluctuations of the mirrors changing the length of the cavity [67]. The magnitude of these fluctuations is given by the material properties of the mirror substrate, the mirror spacer and the dielectric mirror coatings. We consider the light hitting the mirror to be an oscillating field and the response of the system will depend on the frequency at which we are trying to measure. The perturbation can be calculated either analytically [68] or numerically using finite element analysis [63].

The first thermal noise source to consider is the Brownian thermal noise. We use the fluctuation dissipation theorem [69][70] to calculate the spectral density $S_x(f)$ of the thermal fluctuations that can be calculated from the force of the oscillations imposed on the mass by a light beam,

$$S_x(f) = \frac{2k_B T}{\pi^2 f^2} \frac{W_{\text{diss}}}{F_0^2} \tag{4.5}$$

where W_{diss} is the dissipated power from the laser beam into the system, F_0 is the total force the laser beam has on the mirror, T is the temperature of the system, k_B is the Boltzmann constant and f is the frequency of the noise.

The analytical solution [68] considers a mirror with an applied beam pressure; taking the case of a material which has a mechanical loss $\phi_{\text{substrate}}(f)$ the spectral density of the Brownian noise is given by

$$S_x(f) = \frac{2k_B T}{\pi f} \frac{1-\sigma^2}{\sqrt{2\pi} E r_0} \phi_{\text{substrate}}(f), \tag{4.6}$$

where E and σ are the Young's modulus and Poison's ratio of the mirror substrate material respectively and r_0 is the radius of the laser beam at the point where the light intensity has fallen to $1/e$ of the maximum. There is also thermal noise due to the strain energy in the dielectric mirror coatings [71]. This is proportional to the substrate thermal noise and is calculated using a multiplication factor of the mirror substrate noise to give the total Brownian noise for the system; this is given by

$$\sim \left(1 + \frac{2}{\sqrt{\pi}} \frac{1-2\sigma}{1-\sigma} \frac{\phi_{\text{coating}}(f)}{\phi_{\text{substrate}}(f)} \frac{l}{\sqrt{2} r_0}\right) \tag{4.7}$$

where l is the thickness of the coating and $\phi_{\text{coating}}(f)$ is the mechanical loss of the optical coating on the substrate.

However this analytical solution does not take into account effects due to the shape of the mirrors used or the effect of a spacer as this changes the energy dissipation in the mirrors and coatings. The gravitational wave community has developed further theory to numerically calculate the thermal noise [72]. The separate energy dissipation in each component can be calculated using finite element analysis [63]. By using a computer model of the optical cavity and applying a gaussian force profile to each of the two internal faces of the mirrors, the separate strain energy in each component can be extracted from the model. The strain energy U_{max} can be used to calculate the dissipated power $W_{\text{diss}} = 2\pi f U_{\text{max}}$.

A further source of noise is the thermoelastic fluctuations due to heating of the mirror material at random positions and causing temperature gradients [73]. This noise is only dominant in the mirror substrate and we disregarded the noise due to mirror coatings. The noise in the mirror substrate is given by

$$S_{\text{te}}(f) = \frac{\sqrt{2}}{\pi^{5/2}} \frac{k_B T^2 \alpha^2 (1+\sigma)^2 \kappa}{\rho^2 C^2 r_0^3 f^2} \tag{4.8}$$

where α is the coefficient of thermal expansion, C the specific heat, κ the thermal conductivity

and ρ the density of the material. Further examination of thermoelastic fluctuations shows that the noise is not linear in frequency and there is a frequency cut-off at frequencies lower than 1 Hz due to adiabatic behaviour of the mirror substrate [74]. At this frequency the thermal diffusion length is greater than the beam radius and the thermoelastic fluctuations do not affect the optical cavity length.

A final source of noise that may be significant in some systems is the photothermal noise, due to the fluctuations in the number of photons absorbed by the mirror [75]. The photothermal noise is given by

$$S_{\text{ph}}(f) = 2\alpha^2(1+\sigma)^2 \frac{\hbar f_0 W_0}{\rho^2 C^2 r_0^3 f^2} \tag{4.9}$$

where f_0 is the frequency of the photons hitting the mirror and W_0 is the peak absorption power. Again it can be shown that the photothermal noise has a frequency cut-off at frequencies lower than 1 Hz due to adiabatic behaviour of the mirror substrate [74].

4.8 Thermal Noise Simulations

A commercial finite element analysis program ANSYS [76] was used to simulate the deformation of the optical cavity due to a gaussian shaped beam incident on the mirrors. The methods used for these simulations were similar to the methods shown in work by Numata [63]. We are able to measure the deformation of the optical cavity due to the photon pressure deforming the surface of the mirror. An exaggerated deformation of an optical cavity due to the photon pressure is shown in figure 4.5. The program calculates the strain energy W_{diss} in various parts of the optical cavity due to the deformation; from which we can calculate the Brownian thermal noise in the mirror substrate, the mirror dielectric coating and the spacer. The optical cavity used for the model is a simple cylindrical design that has a spacer 100 mm in length and has a diameter of 50 mm with mirrors 25.4 mm in diameter optically

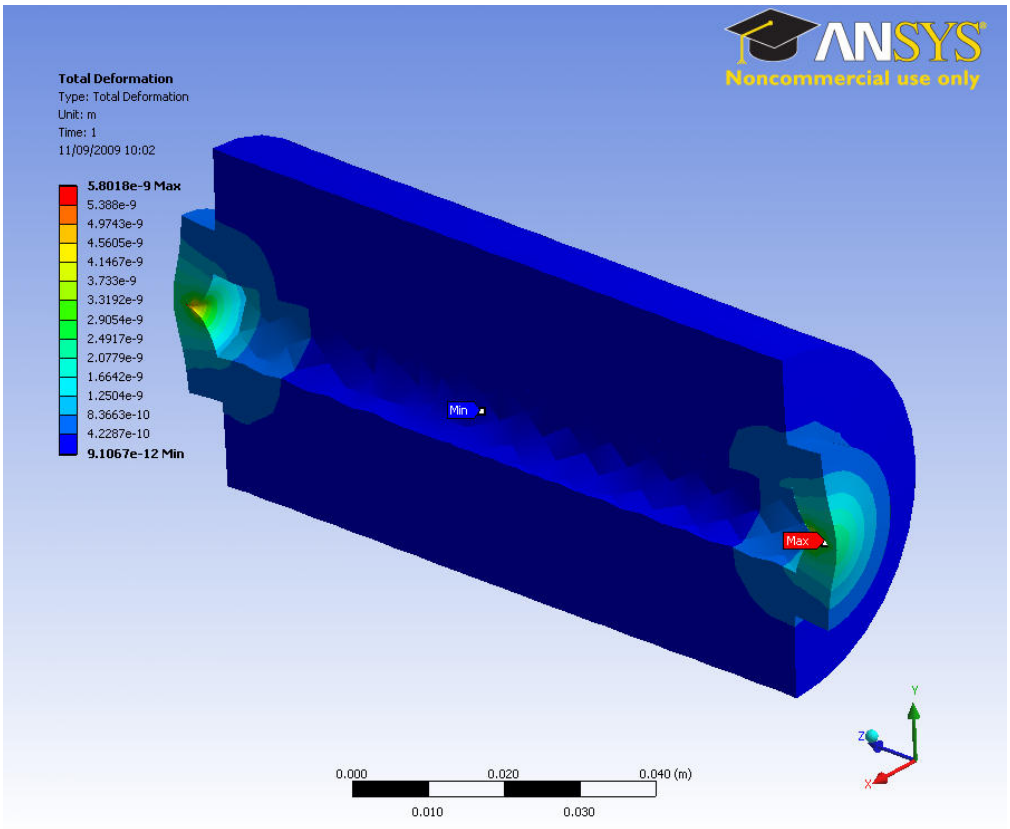

Figure 4.5: This is the image produced by ANSYS finite element analysis program to calculate the thermal noise. The response of the entire cavity is modelled using a laser beam with a gaussian beam shape incident on both mirrors. For clarity only half of the optical cavity is shown and deformation due to the laser beam has been exaggerated significantly so that the change in shape is obvious in this illustration of the optical cavity.

contacted to each end of the spacer; the optical coatings cover the internal mirror surface and are assumed to be a homogeneous 140 μm thick dielectric coating. In reality the dielectric coatings are made from many alternating layers of dielectric but this is not important for the modelling.

In figure 4.6 we validate the calculated values of the Brownian noise thermal obtained from the model against the analytical values of the Brownian thermal noise calculated from the equations in [67]. The total of the analytically calculated values is smaller than the numerical values produced by the finite element analysis program for the Brownian thermal noise by a factor of 0.65. However for the optical cavity spacer the analytical values are 0.33 of numerical

values; this significant difference between the modelled values and analytical values is possibly due to the strain of the mirror also deforming the spacer not being taken into account for the analytical values. In figure 4.7 all noise sources are shown: the Brownian noise for different components, the thermoelastic noise and the photothermal noise. The photothermal noise and the thermoelastic noise are shown to have a frequency cut-off due to the adiabatic behaviour of the mirrors because of how the energy is dissipated. The thermoelastic and photothermal noise were not calculated using the strain energy as calculated by the ANSYS program; the values for these noise sources were only calculated analytically. As we can see from figure 4.7 we can clearly disregard the photothermal noise for the ULE optical cavity; being more than 10^{10} times smaller than any other noise source, it has an insignificant effect.

In an attempt to reduce the total noise it would be advisable to change the mirror substrate as this is the most significant noise source. In figures 4.8 and 4.9 the differences in noise levels are shown for different materials used for the mirror substrate. For the optical cavity with ULE mirror and a ULE spacer the Brownian noise of the mirror substrate will dominate. If we replace the ULE of the mirror substrate with fused silica the thermoelastic noise becomes far more significant but the overall noise is lower. By comparing other materials we can find the best materials to minimise the thermal noise.

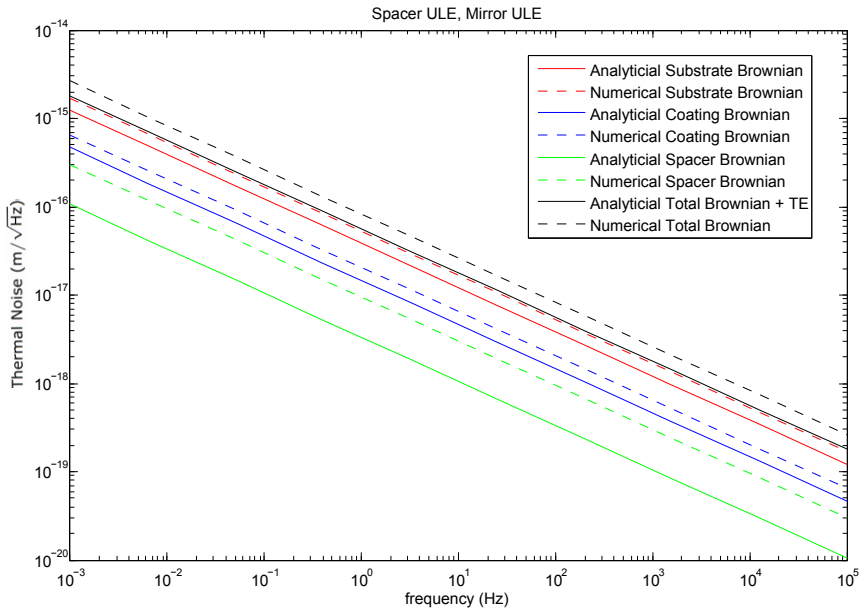

Figure 4.6: A comparison of noise calculated from the numerical simulations from finite element analysis and the analytically calculated noise. This was for an optical cavity with a spacer made from ULE and the mirrors substrates made from ULE.

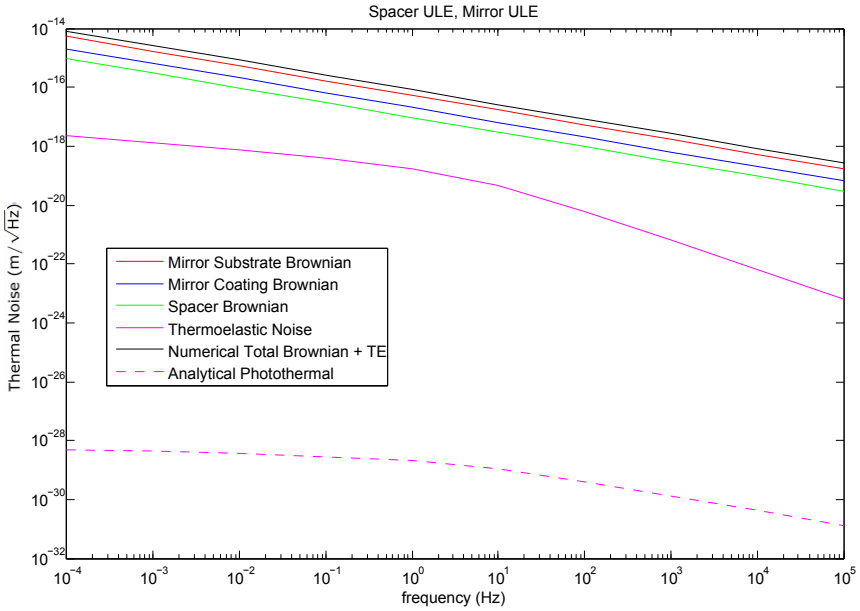

Figure 4.7: All noise sources are shown for an optical cavity with a spacer made from ULE and the mirrors substrates made from ULE.

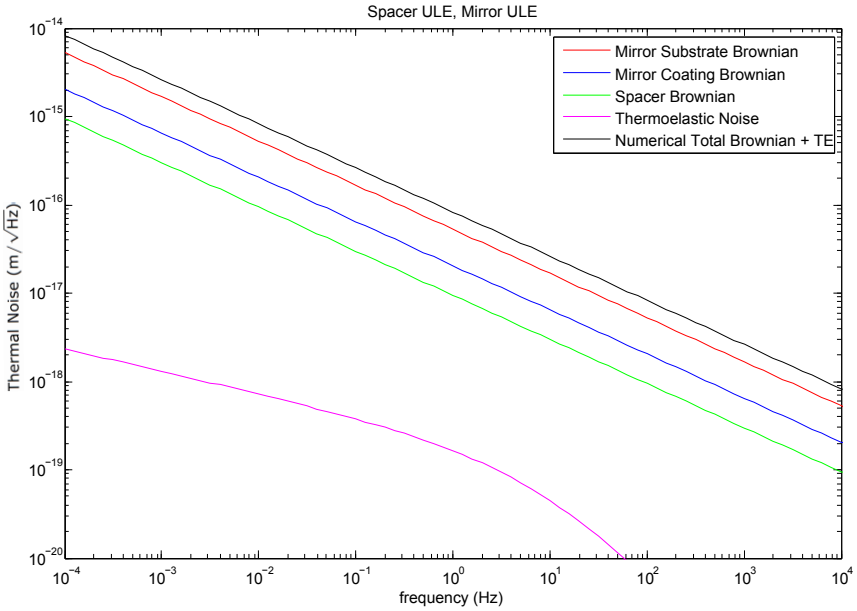

Figure 4.8: The thermal noise of an optical cavity made with a ULE mirror substrate and a ULE spacer.

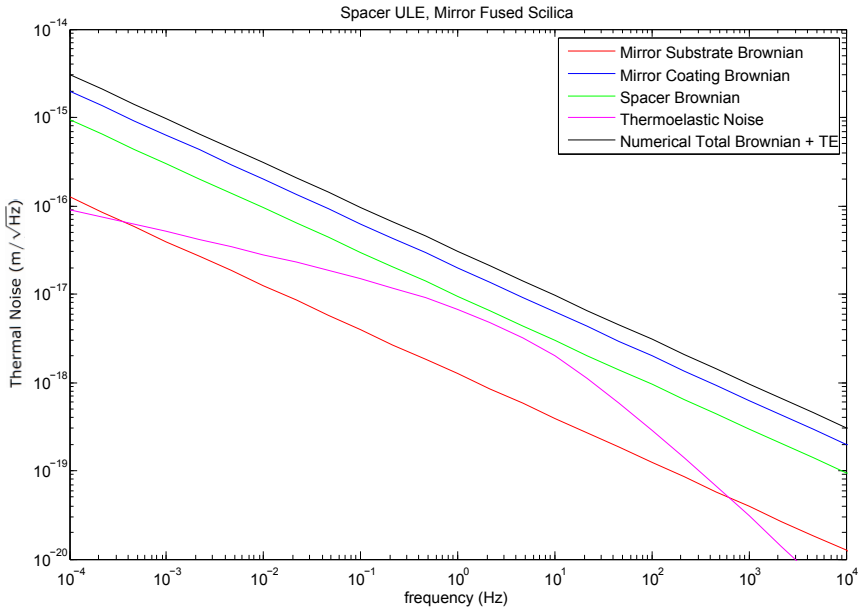

Figure 4.9: The thermal noise of an optical cavity made with a Fused Silica mirror substrate and a ULE spacer.

Cavity Spacer	Zerodur	Zerodur	ULE	ULE	Fused Silica	Silicon	Sapphire	Sapphire
Mirror Substrate	ULE	Fused Silica	ULE	Fused Silica	Fused Silica	Silicon	Sapphire	Fused Silica
Mirror Brownian Noise ($m/\sqrt{Hz}$)	5.27×10^{-17}	1.24×10^{-18}	5.29×10^{-17}	1.24×10^{-18}	1.24×10^{-18}	8.45×10^{-19}	2.94×10^{-18}	1.24×10^{-18}
Coatings Brownian Noise ($m/\sqrt{Hz}$)	2.05×10^{-17}	1.99×10^{-17}	2.05×10^{-17}	1.99×10^{-17}	1.99×10^{-17}	1.45×10^{-17}	1.38×10^{-17}	1.99×10^{-17}
Spacer Brownian Noise ($m/\sqrt{Hz}$)	7.52×10^{-17}	7.47×10^{-17}	9.47×10^{-18}	9.44×10^{-18}	2.24×10^{-19}	1.55×10^{-19}	5.21×10^{-19}	6.05×10^{-19}
Thermoelastic Noise ($m/\sqrt{Hz}$)	1.65×10^{-19}	6.71×10^{-18}	1.65×10^{-19}	6.71×10^{-18}	6.71×10^{-18}	1.91×10^{-17}	6.40×10^{-17}	6.71×10^{-18}
Photothermal Noise ($m/\sqrt{Hz}$)	2.03×10^{-29}	9.34×10^{-28}	2.03×10^{-29}	9.34×10^{-28}	9.34×10^{-28}	1.02×10^{-26}	1.16×10^{-26}	9.34×10^{-28}
Total Displacement Noise ($m/\sqrt{Hz}$)	1.49×10^{-16}	1.03×10^{-16}	8.30×10^{-17}	3.73×10^{-17}	2.81×10^{-17}	3.46×10^{-17}	8.12×10^{-17}	2.85×10^{-17}
Total Frequency Noise ($Hz/\sqrt{Hz}$) *	0.64	0.44	0.36	0.16	0.12	0.15	0.35	0.12

Table 4.2: The spectral density of thermal displacement noise for materials at a frequency of 1 Hz. Values are given for a range of possible materials used to make ultra-stable optical cavities. All of the calculations were performed at a temperature of $T = 300$ K. *Total frequency noise for a 100 mm optical cavity when using a 698 nm wavelength laser.

Table 4.2 shows the thermal displacement noise for each of a range of material combinations at a frequency of 1 Hz. As can be seen from the table photothermal noise is insignificant in every material tested. An alternative representation of the data is shown in figure 4.10. Also included in the table is the total frequency noise which is calculated from the displacement noise. As the frequency of a mode is given by

$$\nu = \frac{nc}{2L} \tag{4.10}$$

this can be differentiated to give the change in frequency $\Delta\nu$ given a displacement in the optical cavity length ΔL

$$\Delta\nu = -\frac{nc}{2L^2}\Delta L = -\frac{\nu}{L}\Delta L. \tag{4.11}$$

For example, a 100 mm optical cavity with a laser wavelength of 698 nm will have a frequency change of 5 Hz due to a length variation of only 1 fm. The thermal displacement noise for fused silica is 2.8×10^{-17} m/$\sqrt{\text{Hz}}$, which corresponds to a frequency noise of 0.15 Hz/$\sqrt{\text{Hz}}$. From [63] we find that

$$\sqrt{S_x(f)} = \sigma_y \nu / \sqrt{2(\ln 2) f} \tag{4.12}$$

where σ_y is the instability of the laser. Assuming a laser locked to the optical cavity is only limited by this frequency noise, for a measurement time of 1 second the fractional frequency instability would be 1.1×10^{-15} for a 698 nm laser.

From the modelled results we can conclude that the best material for an optical cavity to have minimal thermal noise would be either fused silica or silicon. Using silicon for our experiment would cause significant issues as cryogenic cooling would have to be used because the turning points in thermal expansion for silicon are 4.2 K and 124 K [65]. Also silicon is opaque to light with a wavelength shorter than 1000 nm and therefore not suitable for the wavelength of light we require for the strontium clock transition. If we wished to use silicon a

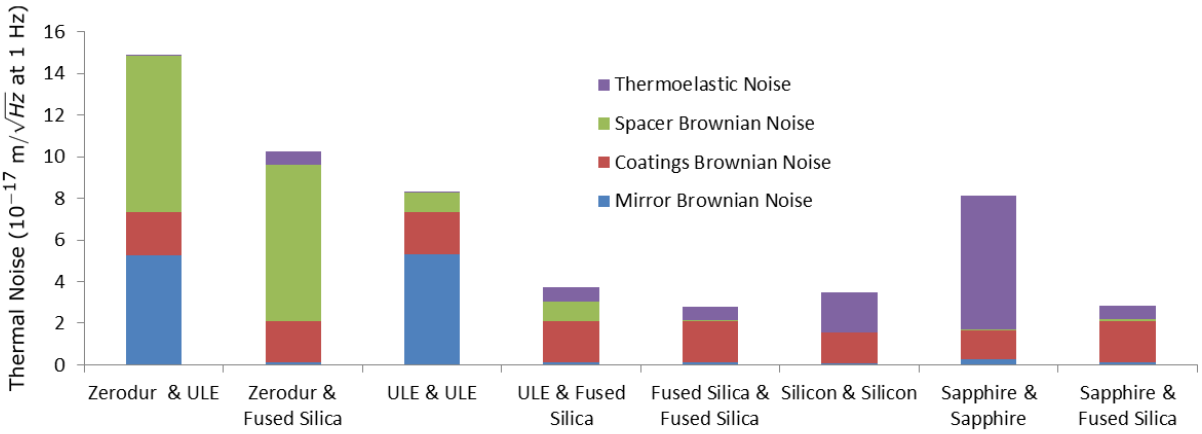

Figure 4.10: A graphical representation of the thermal noise values given in table 4.2.

long wavelength laser would need to be locked to the optical cavity and a frequency doubling arrangement would be required to gain the correct frequency or a frequency comb would be required. For our use, according to the results found, an optical cavity made completely from fused silica would be the best choice but it would have significantly higher thermal expansion at room temperature compared to ULE. Even though the fused silica is better for thermal noise the length stability of the cavity would be much lower. To improve the length stability a ULE spacer with fused silica mirror substrates could be used with a ring of ULE on the back of the mirror to correct for the difference of thermal expansion between the two materials [77].

In addition to changing materials there may still be other developments to reduce the thermal noise using techniques developed by the gravitational wave community, such as using higher order Laguerre-Gauss modes that increase the beam size on the mirrors thereby reducing the thermal noise [78]. Another option would be mirror coatings constructed in a different way to the usual dielectric coatings used for our highly reflective mirrors, by structuring the surface to have resonant coupling of the light modes with the mirror coating, it may be possible to have high reflectivity mirrors without the thermal noise associated with our current dielectric coatings [79].

4.9 Mechanical Stability of Optical Cavities

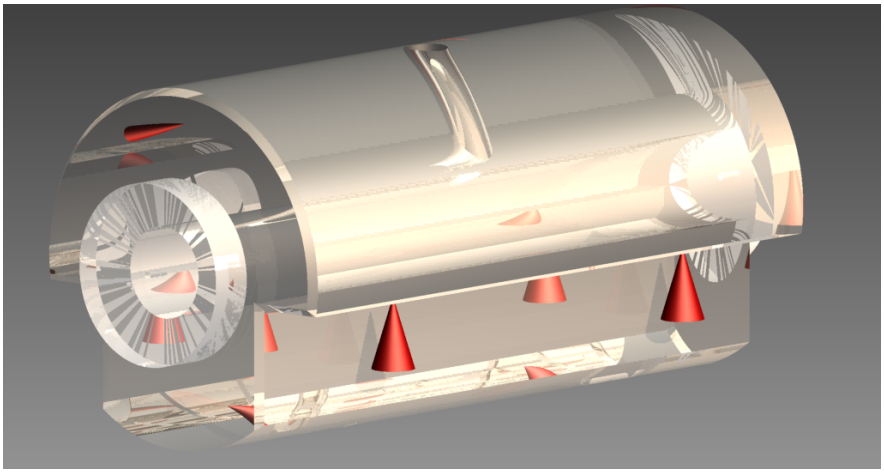

Figure 4.11: A computer image of how the cavity supports are modelled. Four points, two on each side symmetrically about the centre, support the cavity. The red cones shown illustrate how the supporting points are modelled; as four defined points on the surface of the horizontal cavity.

To minimise the problem of vibrations affecting the length of the optical axis we can design a cavity that is mechanically stable; we can also design the supporting structure such that the distance between the mirrors does not change if vibrations cause the optical cavity to deform. We are able to use finite element analysis to produce a design that is passively stable to vibrations applied to the cavity. Previous work has shown there are supporting points where the distance between the mirrors will not change when the cavity is affected by accelerations [59]. By using the same modelling techniques as Webster *et al* [60], using a horizontal notched cavity with a mushroom shaped cross section, four supporting points can be tuned to a particular position on the optical cavity making it vibrationally insensitive; these supporting points have been calculated for our stable cavity. Our optical cavity was slightly different to the ones discussed in the literature so modelling was used to find our optimal support points positions. The modelling program COMSOL [80] was used to perform the finite element analysis of a stress-strain simulation for the optical cavity with a range of

supporting point positions. Figure 4.11 is an image of what was being modelled; the ULE cavity is supported on four points that may be moved in or out and left or right so the ideal supporting position can be found. Due to the symmetry of the cavity the modelling was done with only one quarter of the cavity to reduce the computation time required. The faces that would join to another quarter were constrained such that they were only able to move parallel to the constrained surface, simulating the constraint of the symmetry on a full cavity. As only a quarter of the cavity was modelled only one supporting point was required, a single point was constrained on the surface, this is not a perfect model as the supporting points will have a physical size but it was felt this was a sufficient approximation. The optical cavity was drawn using CAD, fitted with a triangular mesh, the symmetry and supporting point had the constraints applied and finally a force equal to $1g$ was applied in the vertical direction.

Intuitively we can imagine that with the supports at the far edges away from the middle, the mirrors would tip towards each other due to the drooping in the middle of the cavity and the cavity becomes shorter, whereas if the points are near the middle the mirrors would tip away from each other due to the ends bending away and the cavity becomes longer. This is what the model predicts and is demonstrated in figure 4.12; the red region is where the edge is moving away from the centre, the blue region is where the edge is moving towards the centre. The two extremes are shown; the best position is where the top and bottom edges do not move when gravity is applied. By adjusting the supporting points on the cavity the stable positions can systematically be found.

We wish to find a point at which if a force is applied the length of the optical axis does not change, therefore if acceleration affects the optical cavity caused by a vibration in the vertical direction, this does not affect the optical cavity length. The stable positions are where if a vertical acceleration is applied there is zero axial displacement of the mirrors. The simulations were run for various positions of the supporting point along the cavity and

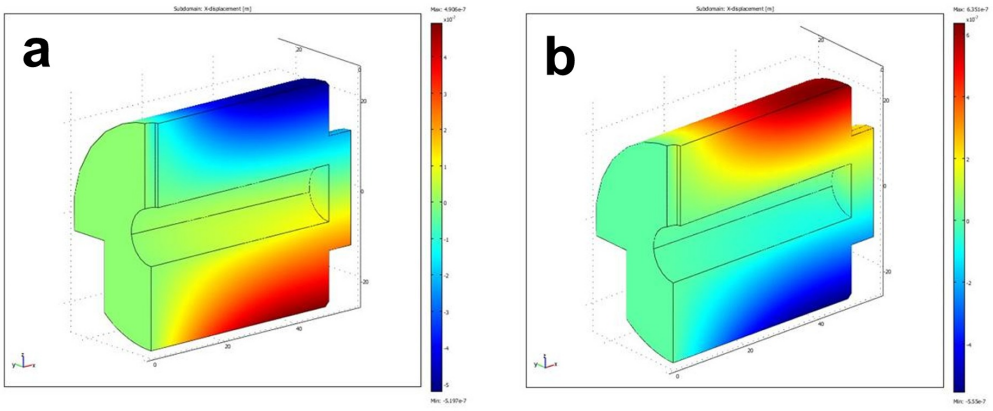

Figure 4.12: Images from COMSOL of the mechanical stability of a supported optical cavity. a) Support Position 45 mm from the cavity axis. b) Support position 10 mm from the cavity axis. Red indicates movement away from the cavity axis; the blue indicates movement towards the cavity axis.

at different distances from the axis centre across the optical cavity. The distance change between the centre of the mirrors when a force is applied compared to when the optical cavity is under no force is shown in figure 4.13. The measurement was only in the horizontal direction along the length of the optical axis. There are a number of possible points that could be used; any position which crosses the zero of the axial displacement axis could be used as a possible position for vibration isolation. The same simulation was completed with a higher force of $10g$ applied and the results of the best position to use were consistent with the $1g$ case. The response values found here are similar to results in [60] where a similarly shaped optical cavity was modelled; the dimensions of our optical cavity were different from the one used in the previous results so we should not expect them to be the same. In comparing the modelled response to the experimental results measured in [60] we found they differed from each other and the vibrationally insensitive supporting position was found to be different. This difference between our model and experimental values is still to be experimentally tested; this difference may be due to the physical parameters used in the model such as Young's Modulus not being the same for the batch of ULE the cavity was constructed from. It was also assumed the optical cavity is symmetric and the supporting

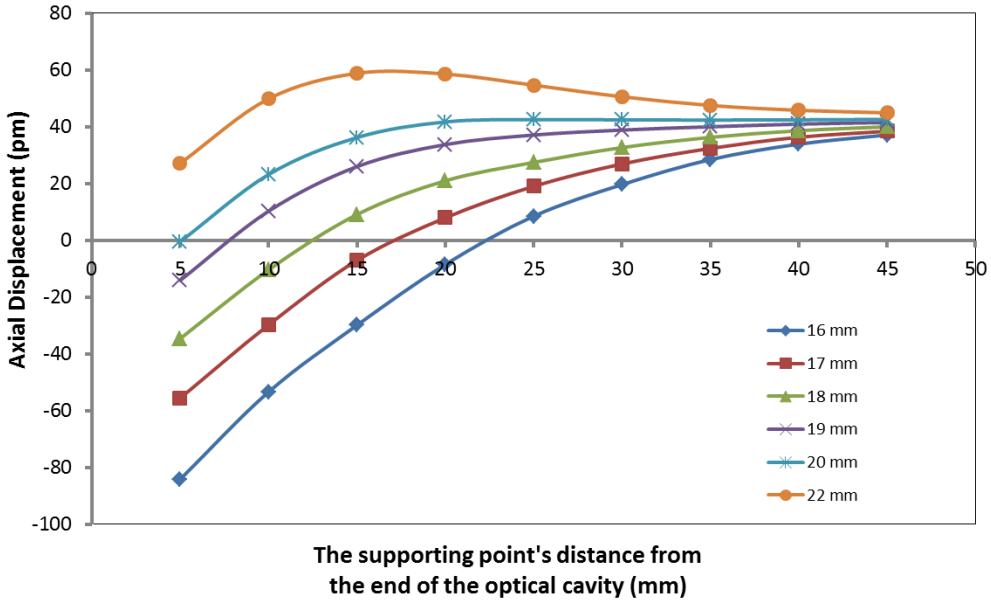

Figure 4.13: Plot of the difference in horizontal axial displacement between the centre of the mirrors when $1g$ is applied to the optical cavity. Each plot represents a different supporting position from the centre of the spacer at various points along the length of the cavity. Positive displacement is movement away from the centre and negative displacement is movement towards the centre of the cavity.

points in contact with the cavity were modelled as a point instead of a small area that occurs in the real system. The optimisation will need to be done with our cavity experientially to have the best performance for vibrationally insensitive response.

If the precision when positioning the optical cavity on the supporting points is 1 mm the response to a $1g$ acceleration will be of the order of 10 pm. Our fractional length change will therefore be $1 \times 10^{-10}/g$. From equation 4.4 the frequency response is calculated to be 40 kHz/g. To reach a frequency stability of 1 Hz we require vibrations of no more than $10^{-5}g$ transmitted to the optical cavity.

Alternative shapes of the optical cavity and the respective modelled sensitivity to vibration is presented in a work by Chen *et al* [59]. For a cylindrical optical cavity mounted horizontally, similar to the cavity modelled here but the without the cut-out sections, when supported at the optimum position with a U-bracket the modelling showed a fractional length

change of $2 \times 10^{-9}/g$. For an optical cavity mounted vertically that is supported by a central ring with the spacer tapered towards the mirrors the sensitivity of the fractional length change was shown to be $7 \times 10^{-11}/g$. This shape of optical cavity is less sensitive to accelerations due to the symmetry of the spacer; under gravity the same force causing the top mirror to move down will cause the bottom mirror to also move down by a very similar amount. The vertical cavity modelled was the same design as another optical cavity that was used for this experiment, as discussed in section 5.3, so a direct comparison can be made to the results from this modelling.

Work has been undertaken to improve designs for the vertical cavity to increase the passive mechanical stability of this system [81]. Also at NPL, modelling and testing a new design developed for space flight has been undertaken with promising results for a rigidly held cavity that is insensitive to inertial forces acting in any direction [82].

CHAPTER 5

EXPERIMENTAL METHODOLOGY

5.1 Introduction

This chapter will present experimental designs and developments of the vacuum enclosures built to house the ultra stable optical cavities to ensure thermal and vibrational control of the systems. It is fundamentally important to use the optical cavities without environmental perturbations to enable us to produce a laser with a significantly reduced linewidth. The optical set up shown in figure 5.7 is the basis of the experiment and should be referred to so that the experimental explanation can be followed. The bandwidth requirements needed for the optical cavity locking electronics will be discussed. As an aside the issues encountered with home built electro-optic modulators will be discussed and the solutions to the problems encountered will be explained. The fundamental result of this thesis is the linewidth measurement and the stability measurement of the locked lasers. The linewidth was measured by a heterodyne beat measurement using a spectrum analyser; the instability will be presented as a plot of Allan deviation measured using a frequency counter.

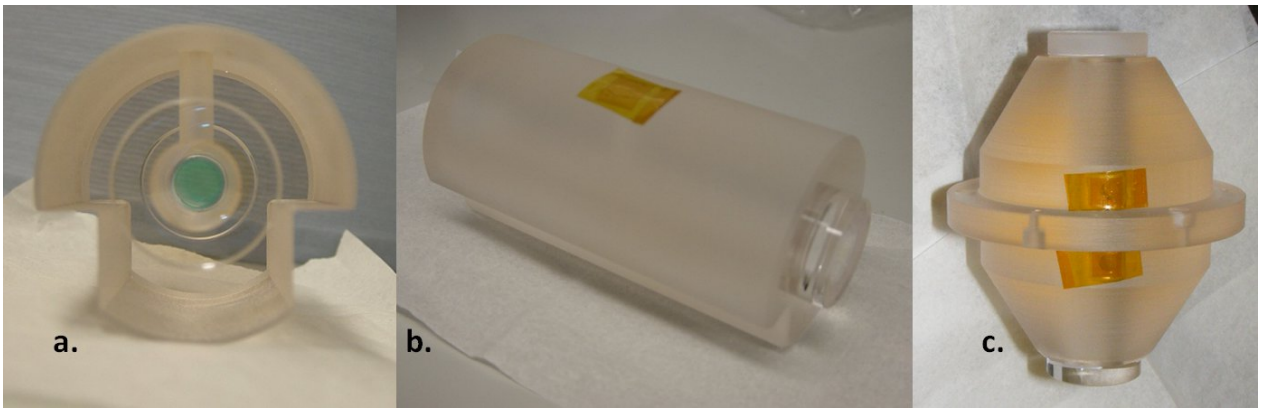

Figure 5.1: Photographs of the ultra stable optical cavities used. a) End view of the H-Cav, with mushroom shaped cross section; the reflective dielectric coating that is transparent at many visible frequencies can be seen on the mirror substrate. b) Side view of the H-Cavity. c) Side view of the V-Cav; the mark on the side of the optical cavity is kapton tape used to seal in the argon filling the cavity during transport, the tape was removed before installation into the vacuum chamber.

5.2 Design of the Ultra Stable Optical Cavities

Two high finesse optical cavities constructed by ATFilms were used [83]. One of the optical cavities was based on a JILA design that has also been used by the Max-Planck Institute for Quantum Optics [50] and will be referred to as the vertical cavity (V-Cav) as the optical axis is mounted in the vertical direction; the other optical cavity was designed based on the NPL design [60] and will be referred to as the horizontal cavity (H-Cav). Both the optical cavities are shown in figure 5.1. The cavities are similar in that they are made from the same ultra-low expansion material (ULE) for both the spacer and mirror substrates. All the high reflectivity dielectric mirror coatings are of the same coating run specifically designed for light with a wavelength of 698 nm. The two optical cavities only differ in the shape and length of the cavity spacer. As these optical cavities where purchased before the modelling in section 4.8 was performed the ULE material recommended by ATFilms was used, future experiment would benefit by considering the materials we have analysed such as fused silica.

Choosing to use two different shapes of cavity allowed a direct comparison to be made. If two identical cavities were used they may share similar vibrational modes, these modes would be common to both cavities so the comparison would possibly not show the true characteristics of the individual laser systems.

The standard mirror choice for commercially available ATFilms cavities is to have one planar mirror and one curved mirror with a radius of curvature of 0.5 m; this arrangement gives a spot size on the mirrors of 0.210 mm and 0.235 mm respectively [84]. To increase the spot size and reduce the average power per unit area a radius of curvature of 1 m was used giving spot sizes of 0.258 mm and 0.272 mm, which equates to an increase in mean spot area of $\sim 40\%$ which will reduce the thermal Brownian noise as predicted by equation 4.6.

5.3 Vertical Cavity Experiment

The vertical cavity experiment was our research group's first experience with ultra stable optical cavities. It was considered that our design should be kept close to the one developed by Alnis *et at* [50] so that a working system could be constructed. The V-Cav spacer is wide in the middle tapering off towards the mirrors (see figure 5.1 c); it is 77.5 mm long and constructed out of ULE. There is a supporting ring around the middle of the spacer with six holes located uniformly around the ring and three of these holes are used to symmetrically support the cavity, the other three holes are in place so the cavity may be mounted the other way up. The mirrors are 25.4 mm in diameter and 6.35 mm thick; these are optically bonded to each end of the cavity spacer. One side of the mirror is coated with a high reflectivity dielectric coating and the other side was coated with an anti-reflection coating. The high reflectivity dielectric coating used was made from alternating layers of Tantalum Oxide (TA_2O_5) and Silicon Oxide (SiO_2); each layer was 3.8 μm thick and 37 layers of each were used.

The design of the vacuum system for the V-Cav mostly uses commercially available CF components from Vacom [85]. The vacuum system is shown in figure 5.2; the outer can was an unequal CF 4-way cross made from stainless steel: a DN150 tube with two short DN35 tubes welded on to the sides. Vacuum components are measured by their internal diameter in millimetres, therefore a DN35 tube has an internal diameter of 35 mm and the outer diameter of the flange is much larger. The top and bottom flanges were custom designed so that a laser window could be mounted on each flange. Through-holes with a diameter of 12.7 mm were drilled in the middle of each flange, the surface around the hole was machined flat to a tolerance of 1 μm and surrounding the flat area were 6 tapped holes. An in-house vacuum sealing method was used to attach the laser windows [86]; a lead-indium wire (Indalloy 165) was used to make a soft gasket to seal the windows on to the flange. The lead-indium wire was soldered together and pre-flattened using a press. The excess solder around the joint was cut off to make a uniform ring; this was to ensure an even pressure was applied across the whole window when sealing. One lead-indium ring was placed between the machined flange and the window and another lead-indium ring between the window and the window holder. This sealing method worked well with the flat laser windows used; however ideally an angled window would be used to ensure no etalon effects could occur in the window or between the window and other optics, but the sealing method used has not been tested for angled windows.

The vacuum system contained two heat shields that were constructed from aluminium because it has a high thermal conductance. The temperature of these shields was controlled using two peltier elements, one between the inner and outer shell and one between the outer shell and the steel vacuum can. The aim was to lower the temperature of the vacuum system to reach the zero crossing of the coefficient of thermal expansion of the ULE, usually around 12 °C. However the ability to cool was limited by the conductivity of the stainless steel flange; the steel was not an effective heat sink therefore the peltier elements were not as efficient as

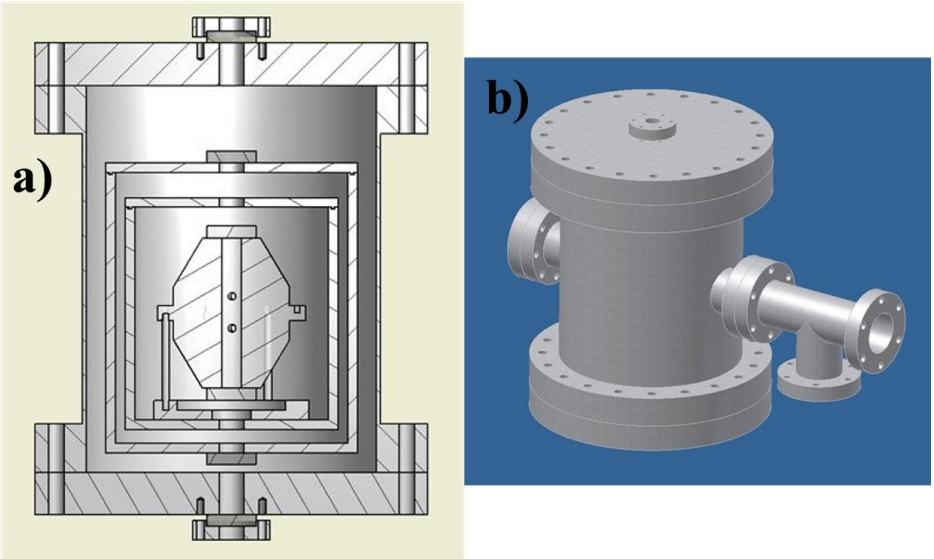

Figure 5.2: a) A drawing of the V-Cav vacuum and thermal shielding assembly. The vertically mounted cavity is placed on Teflon posts that sit on a Zerodur ring, which is enclosed within 2 aluminium heat shields and a Stainless Steel CF150 camber. Not shown are the peltier elements between the boxes to temperature stabilise the system. b) An image of the finished system with electrical feed through and connection to vacuum pumps.

calculated and so the lowest temperature reachable was 18 °C. This was improved by gluing a heat sink to the bottom of the steel flange directly below the peltier element but still only a temperature of 16 °C could be reached; this is with respect to a room temperature of around 20 °C. The temperature was measured with an AD590 temperature sensor; the output of this temperature sensor was a linear current output of 1 μA/K. To measure this current a 1 kΩ resistor was used and the voltage was measured across the resistor. Four AD590 temperature sensors were used with two on each aluminium heat shield, one next to the peltier element to control the temperature of each individual thermal shield and one higher up to take an average temperature to monitor the system. The electrical connection for the temperature control was done using a 19-pin 'Mill' connector, used because it is small enough to fit into a DN35 flange but still has a sufficient number of pins to connect two peltier elements and four temperature sensors as required for our set up.

To ensure the system was able to reach a vacuum pressure of around 10^{-8} mbar all

the components had to be checked to ensure they would not 'outgas'. Outgassing is where gases are released from materials in the vacuum; this outgassing increases the number of particles and therefore increases the pressure within the vacuum system. A major concern was hydrocarbons as these chemicals can interact with the high reflectivity dielectric coatings of the optical cavity mirrors and reduce the finesse. All items put into vacuum were meticulously cleaned using an ultrasonic water bath. All the wires used were kapton coated instead of normal PVC coated wires; for the peltier elements new wire needed to be soldered on [87]. The vacuum system was pumped down using a turbo-pump and a 20 litre per second ion pump was used to maintain the vacuum pressure; a stable pressure of 1.3×10^{-8} mbar was achieved.

The major item known to be an issue for outgassing was the peltier elements. When a current is passed through an element it heats up and releases material into the vacuum, so every effort was made to reduce the outgassing from the peltier elements. Sealed and unsealed peltier elements are commercially available. Naively it could be assumed that a sealed peltier element will be better as the material will be contained in the peltier element but this is not the case as the sealing material is usually something that will outgas anyway; also this seal may possibly leak after some time and if this happens a good vacuum will be ruined by the release of material out of the peltier. We used unsealed peltier elements that were prepared before use; firstly they were cleaned in an ultrasonic bath, washed in high concentration hydrochloric acid, then washed in distilled water to remove the hydrochloric acid and finally washed in acetone as this evaporates easily and leaves no residue in the element. Even with these preparations, when the peltier elements were installed in the vacuum and had current passed through them, the vacuum pressure increased. The elements were left on for some time in vacuum to outgas before the vacuum was broken and the optical cavity placed into the system.

The aluminium heat shields were glued to the peltier elements with Torr Seal epoxy glue

that has a good thermal conductivity and very low outgassing. To support the aluminium cans spacers made from the plastic Vespel were used; this material was chosen because it is a low outgassing material and has sufficient strength to hold the cans in place. Inside the inner aluminium can a low thermal expansion ceramic disk made from SMAHT ceramic was used [88]; three holes were made in the disk to fit Teflon posts for the optical cavity to sit on. These soft posts should minimize the effect of any vibration being transmitted to the optical cavity.

The entire vacuum system was supported by aluminium posts bolted to a minus-K passive vibration isolation platform. Being passive, the weight of the apparatus on the platform had to be carefully calculated so that the platform had optimum performance. The tension in the spring supporting the load can be adjusted to tune the response of the platform to maximise the damping if different load masses were used. The passive system requires no power supply or control electronics therefore would be easier to use than an active vibration isolation platform if the system were ever taken out of the lab; they have similar performance to active vibration isolation platforms.

The entire set up was enclosed in an acoustic box; the box was made from fibre-board lined with rubber and acoustic foam. The rubber ensures the box is sealed and the acoustic foam is used so that any acoustic noise is not transmitted to the air inside the box. The box was designed so that any one side could be removed and the box would still stand, this ensured that equipment could be easily accessed from all sides; the full set up is shown in figure 5.3.

5.4 Horizontal Cavity Experiment

The main consideration for the horizontal vacuum system design was to move towards having a portable system by reducing the size of the vacuum system; the shape and size of the

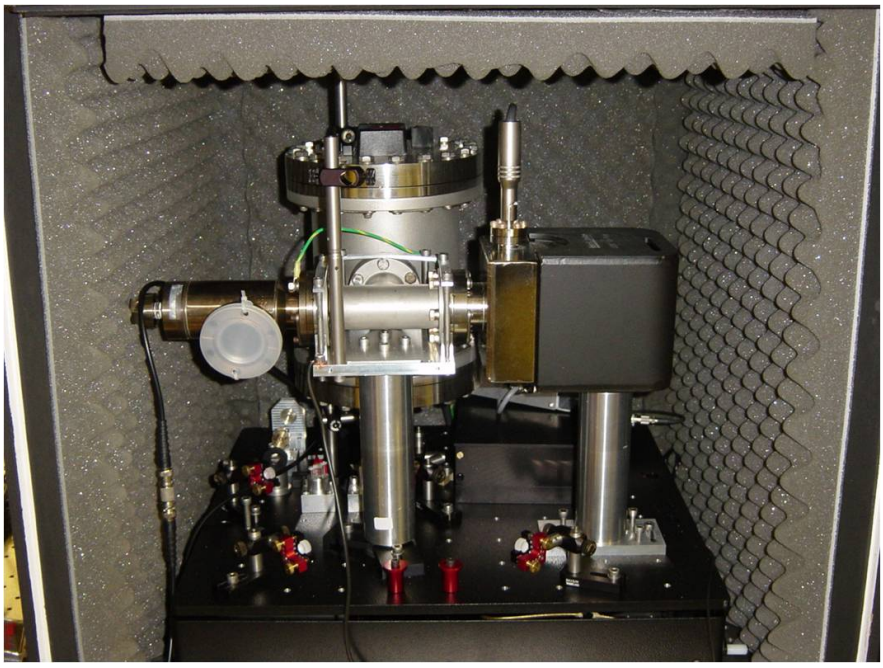

Figure 5.3: Photograph of the full V-Cav assembly. The vacuum system and ion pump can be seen mounted on the minus-K vibration isolation platform. The black box on the right hand side is the laser system with stable ECDL enclosed. Also on the platform are the optics for modulating the light and coupling to the cavity. This is encased in the acoustic box lined with acoustic foam.

enclosure was carefully considered. It was decided that the mushroom-shaped cross section design of cavity spacer by NPL was the best choice if we wished to reduce the size of the vacuum system. The horizontal cavity spacer used was 100 mm long and 50 mm in diameter with cut outs on either side (see figure 5.1 a and b), with 1 inch mirrors optically contacted to the spacer.

The system was designed to be large enough for the cavity to fit inside and minimise the foot-print of the system; a major limitation was the size of the vacuum valve and ion pump. The vacuum valve is the smallest valve commercially available but with a size of 140 mm $\times$ 57 mm is a large component compared to the vacuum system (200 mm $\times$ 100 mm); the ion pump is a 3 l/s pump from Gamma Vacuum. The other components for the vacuum system were custom built in house. For thermal shielding a vacuum-tight outer box made from aluminium was used, inside which was an inner box that was made from copper. Copper has better thermal conductivity than aluminium and, as the copper block was larger than the aluminium shells used for the V-Cav system, the inner box should be a better thermal shield ensuring a uniform heating of the cavity. The copper box was also gold coated to reduce the resistivity of the box; this reduced the black-body radiation between the inner and the outer box. The construction of the vacuum system is shown in figure 5.4. As the outer thermal shield is exposed to the air, the temperature gradient over the box will be a lot higher than the V-Cav so it is possible that the temperature fluctuations will be higher. Windows were fixed to the aluminium box using Torr Seal vacuum epoxy. The windows used were flat and AR coated; again it would have been preferable to use angled windows to reduce the etalon effect but the flat windows have presented no problem. The optical cavity may have been at a slight angle, which would also alleviate etalon effects.

As discussed in section 4.9, the position of the support points is extremely important; the copper box was built with two brass sliders inside so the supporting points of the optical cavity could be moved to the desired positions. The slider has a dove-tail shaped grove that

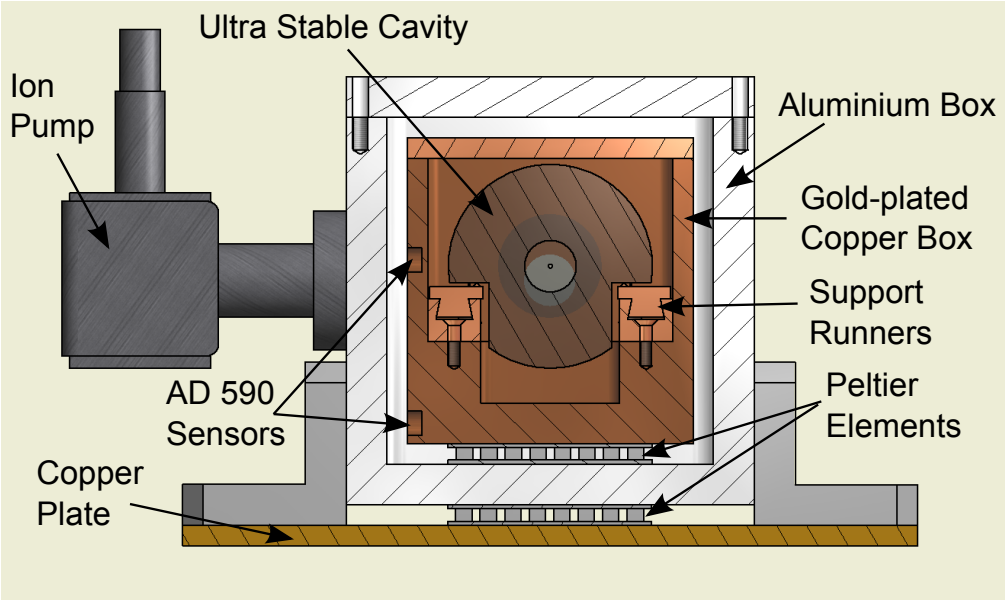

Figure 5.4: A cut through image of the H-Cav vacuum system and thermal shields. The internal copper box (gold coating not shown) sits on a peltier element inside an aluminium outer box that also acts as a vacuum can and a second larger peltier is underneath the outer aluminium box. A copper plate is used as a heat sink so the system can be cooled to low temperatures. The thermal insulation between the aluminium right-angle brackets between the aluminium outer box and the copper heat sink is not shown.

the support sits in as shown in figure 5.5; to ensure the experiment matched the conditions of the modelling the supporting points are made from hardened stainless steel so they do not deform and remain small defined points when supporting the optical cavity. The slider can be easily moved within the system and a grub screw can be tightened down to secure the supporting point in place. As the slider position will need to be changed to tune the supporting points of the cavity the vacuum system must be able to be opened many times.

The inner box is simply sealed with screws; since the outer box is required to be vacuum tight, the seal between the top of the box shell and the lid was made from lead-indium alloy wire. Both the top edge of the box and the lid were machined flat to a high tolerance (100 μm flatness) so that when the lead-indium wire was pressed between the two smooth surfaces the seal was vacuum tight. The lead-indium seal was made by shaping the wire into a rectangle of the size required and twisting the ends together. The round wire was squashed by sealing

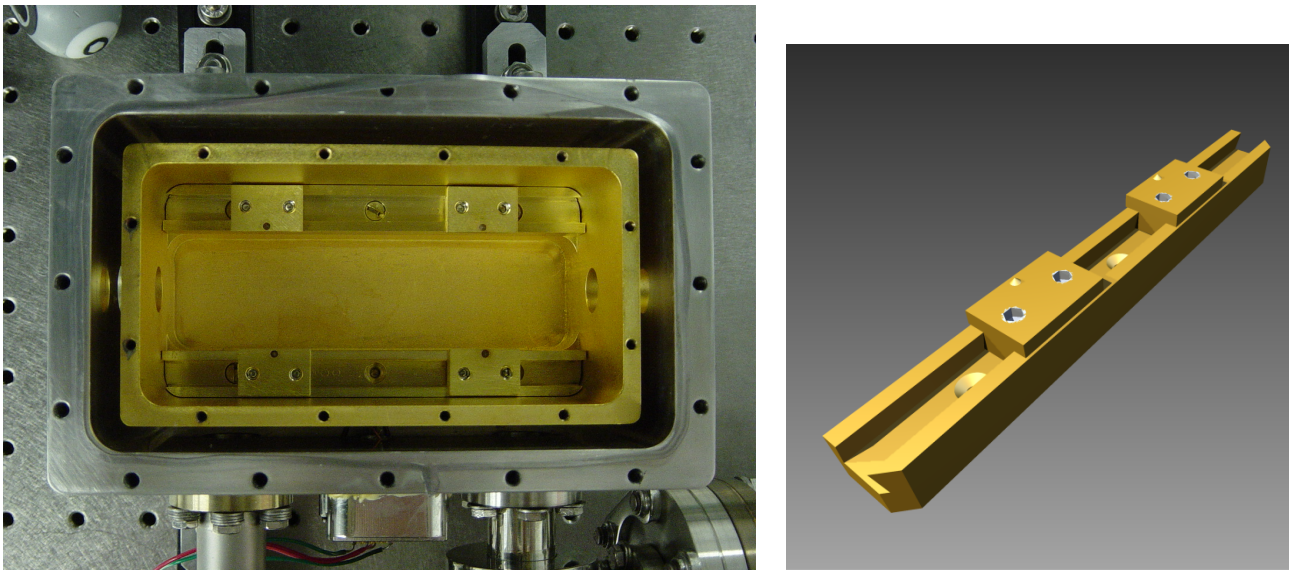

Figure 5.5: Photograph of the open horizontal vacuum system. The outer aluminium can encloses the inner gold-plated copper box. At the bottom of the image are the vacuum connections to the ion pump and the valve to the turbo pump. The electrical feedthrough supplies connections to the peltier elements and AD590 temperature sensors. On the rim of the outer box a lead-indium seal is used; on the bottom edge is the flattened twist in the lead-indium seal to make a vacuum tight ring. The brass dove-tail runners and movable support points for the cavity support are shown inside the inner box. The right hand image is a slider used in the horizontal vacuum system. The dove tail groove in the slider allows the supporting position to easily be moved and fixed.

the box; the lid has 16 screws to give even pressure across the whole box, uniformly pressing the wire flat. It was found that each seal can be used five or six times and the same vacuum pressure can be reached. The vacuum connections for the ion pump and the valve to the home built aluminium box was also sealed using a lead-indium seal; the seal was positioned between the aluminium box and a copper gasket, and the gasket then goes directly onto the knife edge of the CF vacuum component. The alternative would be to have a knife edge machined into the aluminium, but this would be very soft and may not have sealed; a way to make the aluminium more resilient would be to coat the knife edge with some material to harden it but this would be costly.

Inside the vacuum system are two AD590 temperature sensors, one near the inner peltier and one higher up to monitor the temperature of the system. These sensors and the peltier are connected via two vacuum feedthrough 9-pin sub-D connectors; the wires are soldered on and the vacuum feedthrough was glued in place using Torr Seal. On the outside of the box is a two stage peltier between the aluminium outer vacuum chamber and a copper plate. The copper plate acts as an excellent heat sink and we were able to reach temperatures as low as 14 °C. A further AD590 sensor is attached to the outside of the aluminium box next to the peltier element to control the temperature of the outer box. An acoustic box was also constructed for the H-Cav set up; similar in construction to the one shown in figure 5.3. However there were issues when the acoustic box was used; as the air was contained in the acoustic box convection was not able to cool the copper plate and the temperature of the vacuum system rose. For example, when attempting to cool to the system to 14 °C it instead heated up to 22 °C. It was decided a temperature closer to room temperature would be used so less power to cool was required; at room temperature of 20 °C the H-Cav vacuum system could maintain a steady temperature when it was enclosed by the acoustic box. This system also had significant outgassing from the peltier element; however only having one peltier within the vacuum system reduced this problem. The smaller ion pump was believed

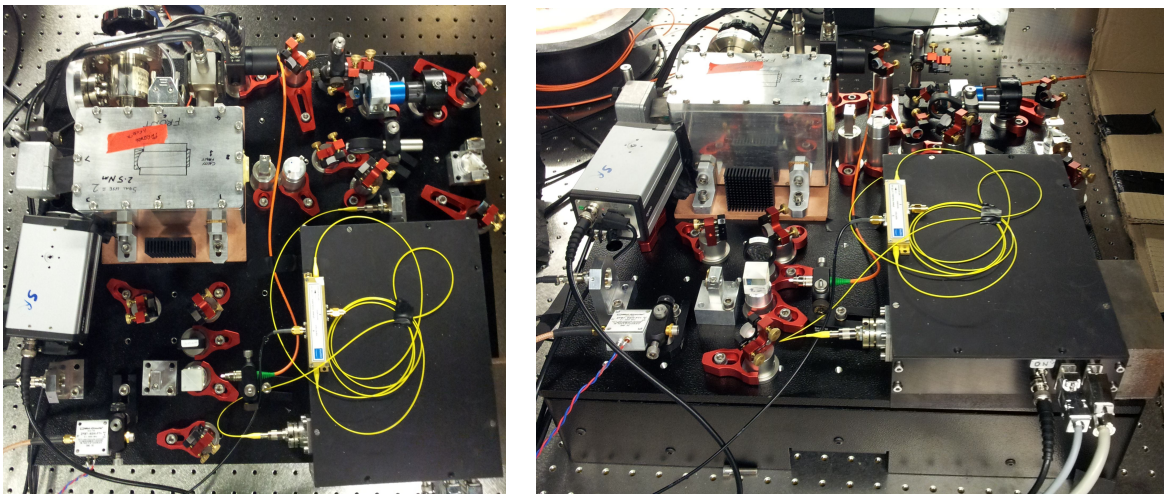

Figure 5.6: Photographs of horizontal cavity system. The vibration isolation platform base plate is 510 mm by 460 mm. The compact optics allows the cavity, laser system and other optics to be placed entirely on one vibration isolation platform. In the lower right is the laser system with the fibre coupled EOM (yellow) mounted on top, in the upper left is the horizontal cavity vacuum system, in the lower left is the optics for beat measurement between light picked off from the horizontal cavity and the light coming in via fibre from the V-Cav system to the vibration isolation platform.

to limit how low the vacuum could reach; the stable pressure measured on the ion pump was 1.2×10^{-7} mbar.

5.5 Optical Set Up

The full optical set up is shown in figure 5.7. For each optical cavity a laser system was built that comprised of an ECDL with two -40 dB optical isolators and a fibre coupler; the isolators ensure that there is no optical feedback to the laser diode that may cause future problems. The ECDL design used has been fully discussed in section 3.3. The remaining optics in the box were required to couple light into the fibre. The fibre between the laser system and the table made repositioning of the systems simple and allowed the laser systems to be swapped without requiring realignment of the rest of the optics. A major advantage of using the optical fibre instead of the beam directly from the ECDL was spatial mode

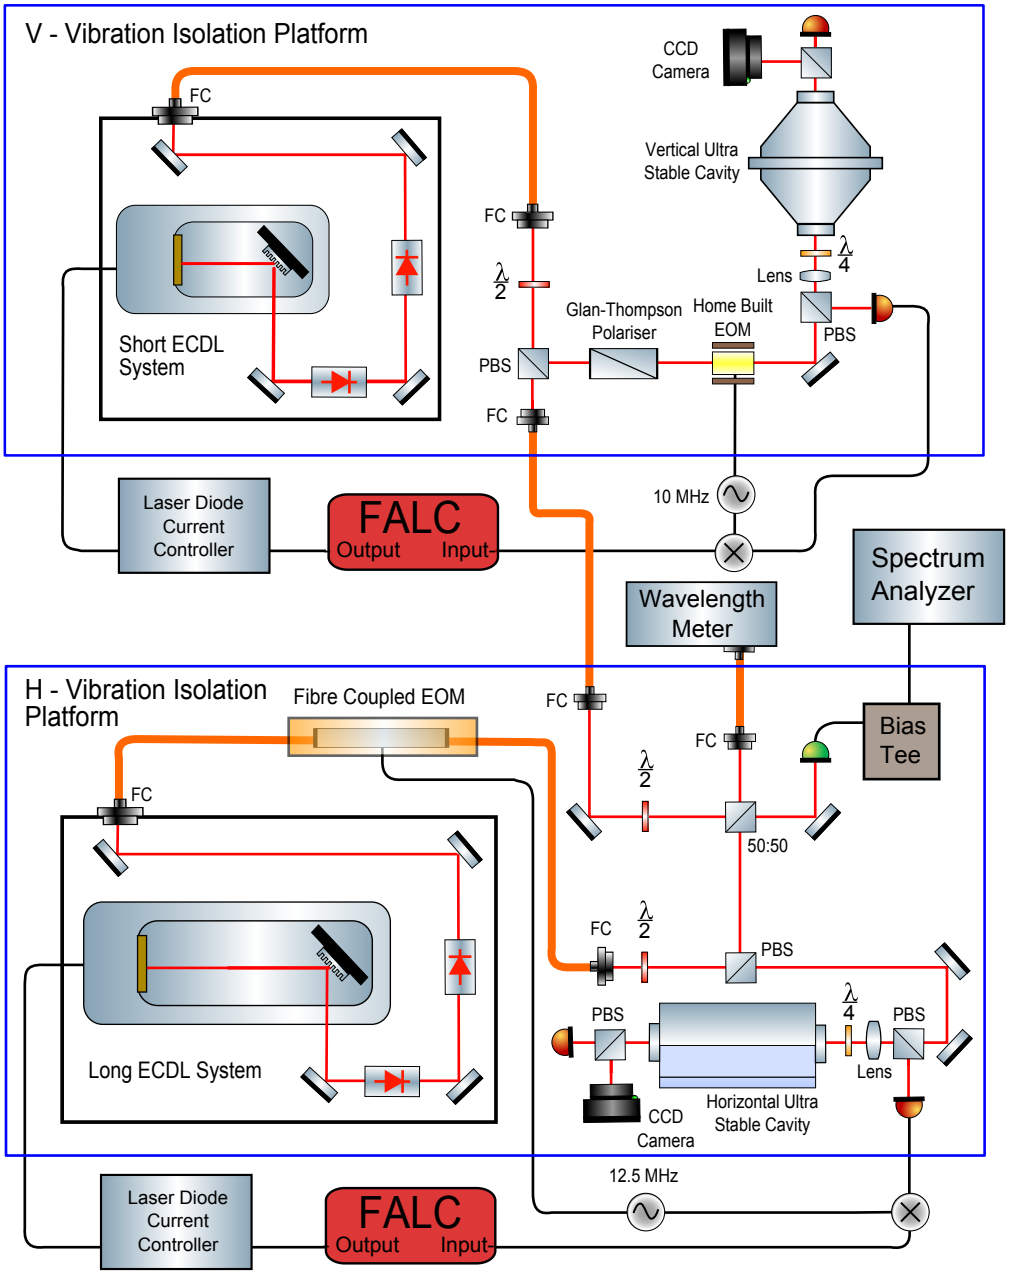

Figure 5.7: Complete experimental arrangement of the beat measurement system. The blue boxes indicate a vibration isolation platform; the vertical cavity system and laser are on one and the horizontal cavity, laser and beat measurement set up are on the other. FC: Fibre Couple, PBS: Polarising Beam Splitter, $\lambda/2$: half-wave plate, $\lambda/4$: quarter wave plate.

cleaning; the beam from the ECDL was very poor so there was significant optical power loss coupling to the fibre, but after the fibre a clean gaussian beam shape was obtained. Having a clean gaussian beam is important for coupling light into the optical cavity, as a gaussian mode will be propagated in the optical cavity.

Being able to couple light to the optical cavity for the first time is a difficult task that requires experience. We require at least two mirrors to align the light into the cavity; a rough alignment can be achieved by ensuring the beam reflected from the cavity is in close alignment with the light beam entering the cavity. A post-it note with a 1 mm diameter hole is highly recommended to align the two beams to be on top of each other as the beam can pass through the hole and the reflected beam can be seen on the post-it. To study the transmission from the cavity a camera and a photodiode are both positioned after the cavity; using a 50:50 beam splitter to split the light between the two sensors they may be used at the same time. The camera used was a colour CCTV camera with the lens removed. The photodiode was a large area photodiode making any alignment easier because there is greater area to hit; this is very useful when the very weak transmitted light cannot be seen with the naked eye. A commercially available large area photodiode from Thorlabs and a home built photodiode were used on different systems; both worked effectively but the ease of use and reduced noise on the signal from the Thorlabs photodiode made it a clear choice for future experiments. The camera is used to determine the shape of the resonant modes in the cavity, whereas the photodiode will give the relative intensity transmitted; both are important to get the desired mode dominant inside the optical cavity. We wish to use lowest order mode, which is the TEM_{00} mode. The modes that can be present within a high finesse cavity can be easily calculated [84] and examples of the modes seen are shown in figure 5.8.

To have maximum transmission we require the laser beam to be 'mode matched' to the optical cavity in size, shape and position. The beam is input on to the flat mirror of the optical cavity where the beam waist will be positioned. The mode matching is used to ensure

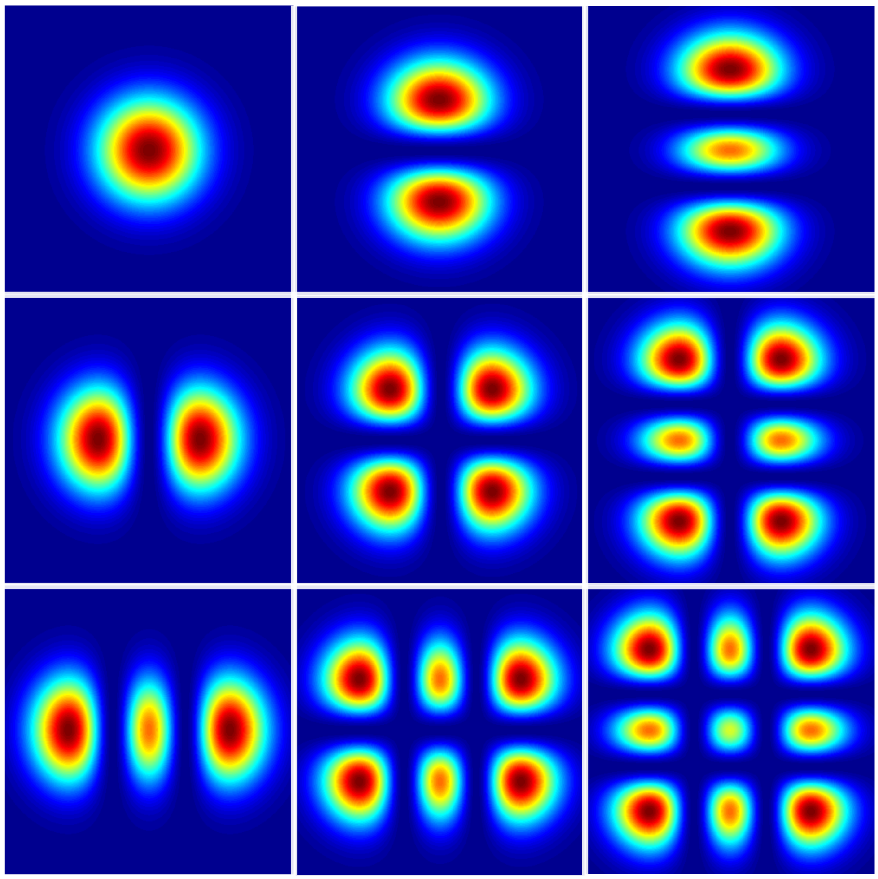

Figure 5.8: Hermite Gaussian modes of an optical cavity calculated from [84]. Modes from TEM_{00} to TEM_{22} are shown. These modes, and higher order modes, were observed on the transmitted signal of the optical cavity. The desired mode we wish to use in the optical cavity is the TEM_{00} mode on the top left.

the input beam matches the optical modes that can propagate inside the cavity. A lens with a carefully calculated focal length should be used to achieve roughly the right beam size to couple to the cavity; the beam size and hence the lens required needed to be calculated from the literature [89]. A series of lenses are used to adjust the spot-size and the beam waist to couple the most light into the TEM_{00} mode. The position of the lenses will need to be adjusted to optimise the coupling. For our set up the lens in the fibre coupler and a secondary lens nearer the optical cavity were used to focus the beam. Fine adjustments could be made by moving the lens in the fibre coupler to change the beam size. Moving the lens would not change the direction in which the beam emitted from the fibre coupler and therefore other optics did not have to be adjusted to compensate.

If we adjust the laser frequency by changing the voltage supplied to the piezo of the ECDL we will see various Gaussian modes within the optical cavity; the frequency should be adjusted until a lower order mode can be found. It was not always possible to find a TEM_{00} mode using the ECDL piezo voltage adjustment so the current to the laser often had to be changed to find the desired mode. If there is not a dominant TEM_{00} mode in the cavity the steering mirrors must be used to correctly align the input beam. On scanning across several modes it may be found that a TEM_{00} mode is not the brightest mode in the optical cavity; when locked to a TEM_{00} mode it is possible that another lower order mode will dominate and the laser will jump to the other mode instead. It was found that when locked the mirrors could be adjusted to increase the light transmitted by the mode observed. This made the TEM_{00} mode dominant; however a steady hand must be used to ensure the mirror is not knocked as the laser will no longer be locked to the desired mode. The power measured in the beam incident on the optical cavity was 1.2 mW and the power of output light transmitted though optical cavity when locked was 5.8 μW; therefore 0.5% of the input light was coupled to the optical cavity. In comparison to [50] who were able to couple 20% into the optical cavity it is clear the beam shaping optics will need to be improved to increase the coupling

efficiency.

The beat note detection apparatus was mounted on the vibration isolation platform with the horizontal cavity set up. A polarising beam splitter was used before the EOM on the V-Cav set up to pick off light and a fibre optic was used to transfer the light from the V-Cav vibration isolation stage to the H-Cav vibration isolation stage. There was concern that noise in the fibre would limit the measurement of the linewidth but this was believed to not be an issue due to the close proximity of the two laser systems and the fibre optic being only 1 m long. If there was an issue with noise in the fibre there are schemes that can be used to correct for the noise using an acousto-optic modulator to correct the frequency [90]. This fibre optic will be used in the future to take light to the optical clock experiment to interrogate the strontium atoms. As a fibre-coupled EOM was used for the H-Cav system to transport the light from the laser system to the optics on the vibration isolation platform the light used for the beat measurement had been modulated. A half-wave plate and polarising beam splitter was used on the H-Cav system to pick off the light needed to be used for the beat measurement. The beams from the H-Cav and V-Cav laser systems were incident on a 50:50 beam splitter; for the beams to interfere they must have the same polarisation. A half waveplate was used to rotate the polarisation to match. As the 50:50 beam splitter is not polarisation dependent this could be done easily; if we were to use a polarising beam splitter additional optics would be required to give the beams the same polarisation.

To measure the beat signal a fast photodiode was used with a bias-T so that a bias voltage could be applied to the photodiode to increase the speed and sensitivity of the detector. The heterodyned beams were aligned and incident on the photodiode. The signal from the photodiode was measured using a Rhode and Schwartz spectrum analyser with a frequency range of 9 kHz to 7 GHz. The fast photodiode and any amplifying electronics may limit the bandwidth of detection; the fast photodiode used had a bandwidth of 10 GHz and for the spectrum analyser no amplification was required. The current drawn by the bias-T is

proportional to the intensity of light on the sensor. By observing the current drawn by the bias-T each beam could be aligned onto the sensor; by observing the signal on the spectrum analyser fine adjustment of the beam alignment could be made to increase the amplitude of the beat signal. It is essential to note when searching for a small signal it is important that no beam blocks are positioned in the system so that both beams are in fact incident on the photodiode.

Results were also taken using an EIP 575 Source Locking Microwave Frequency counter that could measure frequencies up to 40 GHz. The frequency counter was used in place of the spectrum analyser so that the frequency variation over time could be measured. The data could be collected using a computer via the GPIB port on the frequency counter. The signal directly from the fast photodiode was too small to be detected by the frequency counter so two Mini-Circuits amplifiers (ZX 60-3018G-S+) were used to boost the signal with a gain of 20 dB per amplifier; however the bandwidth of these amplifiers was limited to 3 GHz therefore limiting the bandwidth that could be measured by the frequency counter. The frequency change was measured so that the Allan deviation for the lasers could be calculated.

5.6 Laser Locking Electronics Set Up

To control the current supplied to the laser diode, home built laser current controllers were used. Two different current controller versions were used with the same electronics for the current supply but with different ways of changing the current supplied to the laser. One controller had a digital input and the other had an analogue input. The analogue system controlled the current via a potentiometer that was adjusted by hand, this required the controller to be accessible and therefore sit a long distance from the laser system such that the cables supplying current to the laser system passed by many cables carrying RF signals. The RF may have interfered with the laser current supply cable causing fluctuations in

the current supplied to the laser and therefore causing frequency fluctuations. The digital controller was in two parts: the laser current supply and a digital control box; the two devices were connected via an Ethernet cable. This allowed the current supply to sit on top of the laser systems and the cable current supply to the laser could be very short, minimising the possibility of RF interference with the cable. Both the current drivers had an input for fast frequency corrections so that fast electronics could be used to alter the current supplied to the laser and correct any changes in the laser frequency.

A key component in the set up is a commercially available Fast Analogue Linewidth Controller (FALC) built by Toptica [91]. The device is a fast tuneable proportional-integral-derivative (PID) controller specifically built for high finesse cavity locking; the major benefit is the ability to tune the controller locking parameters with ease using an array of dip switches to control the differential and integral frequency response of the PID controller. The FALC has a fast and slow branch; the correcting signal is filtered and the fast part of the signal is sent to the laser current controller, while a slow correction is sent to the ECDL piezo so that we can correct for changes in the length of the ECDL. So that the FALC could be used to lock without unplugging cables, a switch was added in the output BNC cable; this allowed the locking signal to the laser current driver to be turned on easily when required. The limitation of the locking bandwidth is the speed of the electronic devices; the bandwidth of the devices must be significantly greater than the linewidth of the laser used to have minimal optical power outside of the central resonant peak. As discussed in section 3.5 the linewidth of the laser diode is approximately 40 kHz; having a locking bandwidth of 500 kHz would be sufficient to keep the vast majority of light in the central peak.

To change the laser frequency a summing amplifier was built. By moving the grating in the ECDL the laser frequency could be changed; this is done by changing the voltage supplied to the piezo stack. So that the frequency could be moved, a voltage offset was set using a potentiometer acting as a voltage divider with the amplifier supply voltage. To scan the laser

frequency the ECDL was supplied with a saw-tooth voltage signal. The slow branch signal from the FALC used to correct the ECDL piezo position can be fed back to the piezo control voltage. These three signals were summed together with unity gain using a TL072 op-amp in an inverting summing amplifier arrangement. However the ECDL piezo requires a positive voltage, a unity gain inverting amplifier was used after the summing amplifier to invert the signal giving a positive voltage that could be supplied to the ECDL. The maximum output voltage of the summing amplifier was 10 V due to this being the limit of the TL072. To ensure the frequency did not drift, L7812 and L7912 voltage regulators were used to give a regulated ± 12 V supply; this ensured the offset voltage would not drift due to fluctuations in the voltage supply.

5.7 Production of the PDH Error Signal

The PDH locking technique is discussed in section 4.4. For the production of the PDH error signal we added frequency sidebands on the lasers using an Electro Optic Modulator (EOM). For the V-Cav system a home built EOM, which is to be discussed in section 5.10, was used driven by a 10 MHz signal from an amplified crystal oscillator. For the H-Cav a commercially available fibre coupled EOM from Jenoptik [92] was used driven by a 12.5 MHz signal from an arbitrary waveform generator made by Keithley.

A frequency generator will output the frequency required and may also produce harmonics of that frequency in small amplitudes; this can cause issues when being used to drive an EOM. A precaution to ensure that a frequency generator gives only the desired frequency is to use a low pass filter; for the 10 MHz oscillator a Mini-circuits BPL 10.7+ low pass filter was used to remove the higher frequencies from the crystal oscillator.

Using a polarising beam splitter directly before the optical cavity, we were able to separate the incident and reflected beams using a quarter-waveplate after the beam splitter to make

circularly polarised light. The light reflected from the optical cavity was circularly polarised in the opposite direction; the oppositely polarised light would pass back through the quarter-waveplate and become linearly polarised orthogonally to the incoming beam, therefore the beam reflected from the optical cavity would be reflected by the polarising beam splitter instead of passing through it. The signal reflected from the cavity was detected with a fast photodiode with a 150 MHz bandwidth. The signal from the photodiode needs to be mixed with the oscillating signal supplied to the EOM. The mixer used was a Mini-circuits ZAD6+ mixer with a frequency range between DC and 100 MHz. To produce a maximum signal, a phase shifter is required to phase match the reflected signal and the oscillator; lengths of cable were used as a phase shifter so that the bandwidth of the error signal would not be limited by phase shifting electronics. The signal from the mixer is the error signal; this error signal is then filtered by the FALC to give the desired frequency response to achieve the most stable lock to the optical cavity. The filtered signal is input to the current driver for the diode laser so that the frequency can be adjusted to shift the laser back to the resonance of the optical cavity. Usually in PDH locking systems, a low pass filter is used after the mixer to remove the high frequency $\nu_1 + \nu_2$ signal; the FALC has a built in low pass filter so this was not required. However it is often useful to view the error signal directly from the mixer, in which case a low pass filter with a cut off around 100 kHz can be used remove the high frequency components to see a clear signal.

To lock to the signal we ideally want the error signal to have a zero offset; however issues with amplitude modulation and etalon effects were identified in the system that caused the error signal offset to drift over time. The problem appeared to be due to an issue with maintaining the polarisation of the light incident on the optical cavity. The original optical set up had mirrors between the quarter-waveplate and the optical cavity; the circularly polarised light was scrambled by these mirrors. To resolve the issue the beam cube and quarter waveplate were moved to be directly in front of the optical cavity and this drastically

reduced any drift in the error signal and thus was no longer an issue.

5.8 Calculation of the Optical Cavity Finesse

The quoted value for the finesse of the optical cavities produced by ATFilms was in the region of 400 000. This can be tested by seeing how long it takes for the light to be emitted from the optical cavity; this is known as a ring-down measurement. We measure the time required for the light to be dissipated through the highly reflective mirrors; this is similar to how long it would take a bell to become silent by losses of the vibrational energy due to friction. The ring-down will measure the loss of the optical cavity as this will determine the reflectivity of the optical cavity mirrors. The measurement was achieved by turning off the locking to the optical cavity by turning off the PID controller input to the current driver; this caused the laser frequency to jump and no longer be on resonance with the optical cavity mode therefore the light intensity transmitted I will decay exponentially with some time constant τ.

$$I = I_0 e^{-\tau t} \tag{5.1}$$

We can calculate the finesse by calculating the reflectivity of a mirror R; after one round trip time $(2L/c)$ the light is reflected off each mirror once so the proportion of the light remaining in the cavity will be $\rho = R^2$. We must calculate the number of round trips n in a given time t and calculate the power remaining in the cavity ρ^n.

$$I = I_0 \rho^{\frac{ct}{2L}} \tag{5.2}$$

By comparing equations 5.1 and 5.2 we can see that the mirror reflectivity can be calculated.

$$R = \sqrt{\rho} = \sqrt{\exp\left(\frac{-2L\tau}{c}\right)} \tag{5.3}$$

For the high finesse optical cavities we took a ring down measurement using a photodiode to measure the light power transmitted by the optical cavity and logging the decay on a LeCroy oscilloscope. The oscilloscope trigger was set for when the intensity of the transmitted light dropped below 95% of the maximum intensity of transmitted light when the laser is locked to the optical cavity. An example trace of a ring-down measurement for the high finesse optical cavity is shown in figure 5.9. For the ring-down four measurements were taken for each optical cavity and the time constant was calculated by fitting an exponential to the data. Using a MATLAB script the time constant of the decay could be measured and the reflectivity of the mirrors were calculated using equation 5.3. It was found that the start of the decay curve does not match with an exponential function; this may be due to the electronics not immediately shifting the light off resonance when the PID was switched off. As the electronic lock was turned off to move the beam off resonance, this may not have been quick enough initially and light was still entering the cavity, so this part of the data has been disregarded. It would have been better to use an acousto-optic modulator to quickly shift the position of the incident light off the optical cavity; however one was not available. To calculate the finesse from the mirror reflectivity equation 4.1 was used. The finesse of the H-Cav was calculated to be 285000 $\pm$ 2000 and the finesse of the V-Cav was calculated to be 341000 $\pm$ 2000.

5.9 Digital Temperature Controller

To control the temperature of the optical cavities it was decided that a digital PI (Proportional Integral) temperature controller would have significant advantages over analogue electronics. The optical cavities are both housed in vacuum systems with large heat shields surrounding the cavities controlled by peltier elements. The large metal heat shields have very long time constants; to control the temperature the integration time constant for the

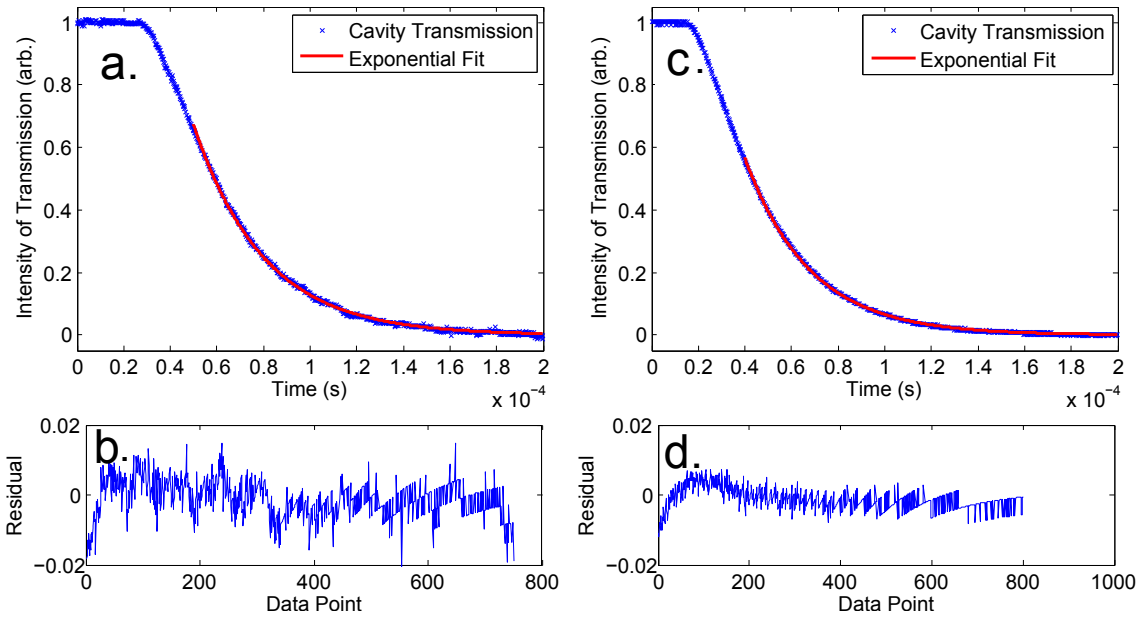

Figure 5.9: Ring down measurement of the high finesse optical cavities. When the laser is moved off resonance there is an exponential decay of light power within the cavity. The red line is a exponential fit from which a decay rate can be found. One of the H-Cav measurements is shown in **a** and the residuals of that measurement is shown in **b**. One of the V-Cav measurements is shown in **c** and the residuals of that measurement is shown in **d**.

controller must be very long. Other ultra stable optical cavity experiments using analogue electronics have achieved long integration times with banks of 10 capacitors in parallel [50], but in a digital system the data can be stored and the calculations for the power required can be made by the computer with an infinite integration time if desired. The temperature controller that was used was built by the HiROS group at the University of Birmingham for the purpose of temperature control of ground based telescopes. The controller is known as "Iguana" because it regulates the temperature of a body; a Linux computer was required to communicate with the controller. It was advantageous to use a digital system because temperature readings from the sensors used for temperature control and the monitoring sensors could be logged on the computer so they may be referred to easily. The ability to change the parameters for the PI control is simple via the computer interface; compare this to analogue electronics where changing parameters would require adjusting a potentiometer or possibly

re-soldering electrical components. The temperature is measured using an AD590 temperature sensor that gives a current of 1 μA/K, which is read by measuring the voltage across a 1 kΩ resistor inside the temperature controller. The temperature measured could change if the temperature of the 1 kΩ resistor changed but the temperature controller box itself is thermally controlled with a cooling fan mounted on the box. The controller readings will fluctuate when it is initially turned on as the box changes temperature and a current flows in the resistor causing it to heat up; this will become stable after several minutes. The Iguana controller takes a reading every second to recalculate the power to be supplied and logs the temperature in the computer's memory every five seconds; readings are collected via a 24-bit ADC (analogue to digital converter), which gives a temperature measurement precision of approximately 0.05 mK. The temperature from one of the AD590 sensors is used to calculate the current supplied to the peltier elements using a PIC micro controller; the computer need not be connected to the Iguana to control the temperature, but the computer is required to change variables and log the data. The current output to the peltier elements is via a 12-bit DAC (digital to analogue converter) and has a range of ± 1 A; the output is bipolar to allow positive and negative outputs to be generated to heat or cool.

The output power $P(t)$ of the Iguana is calculated by comparison of the sensor temperature $T(t)$ with the set-point temperature S. There are two parameters that can be changed; the 'full scale temperature' F and the 'power per temperature' C. The full scale temperature is used to determine the power required given the temperature difference from the set-point temperature; this acts like a proportional control. The power per temperature is used to calculate the power required given the temperature difference prior to the time of the calculation; this acts like an integral control. The power supplied to the peltier element can go between -100% and $+100\%$ and is defined as

$$P(t) = \frac{1}{F}\left(T(t) - S\right) + \frac{1}{100C}\int_0^t \left(T(t') - S\right) dt'. \tag{5.4}$$

In our case the integration is a summation of the discrete temperature measurements. Like any other system, if the wrong values for the PI parameters are selected the system will oscillate. Tuning the parameters to optimise the system took significant time. Periodic variation in the temperature could have an oscillation period of up to two hours and several cycles are needed to identify the fluctuations when the amplitude is small, so a reasonable time scale to examine the temperature over was twelve hours. The Iguana program calculates the temperature measured by the AD590 temperature sensor using an offset value and a scaling factor; these values were adjustable so that any offset in the AD590 could be corrected for. By taking the temperatures of the system at two different temperatures, all the sensors could be calibrated by a method described in section A.4.

To cool the vacuum systems down to the temperatures required for the ULE turning point in thermal expansion around 12 °C more current was needed, a lot more power than the 1 A that the Iguana could supply. For the H-Cav system to reach 12 °C, we required 2 A of current. This current output was above the specifications of the Iguana temperature controller so a power supply was used to pull the system down to the temperature required and the internal peltier element was then used to regulate the temperature. Another issue found was that when the temperature controller switched polarity, a spike would appear in the temperature measured by the controller; these spikes would cause the integrator part of the PI control to increase the power supplied and the system would no longer have good temperature stability. Using a set-point temperature for the inner shell one degree hotter or cooler than the temperature of the outer shell solved this issue, but is not the optimum way to control temperature. The two shields being at different temperatures causes temperature gradients and this may increase the temperature fluctuations in the system. To avoid the spikes it was decided that the inner shell peltier elements on both systems would not be used as a further temperature stabilisation stage, but instead the inner can would only be used as a passive thermal shield.

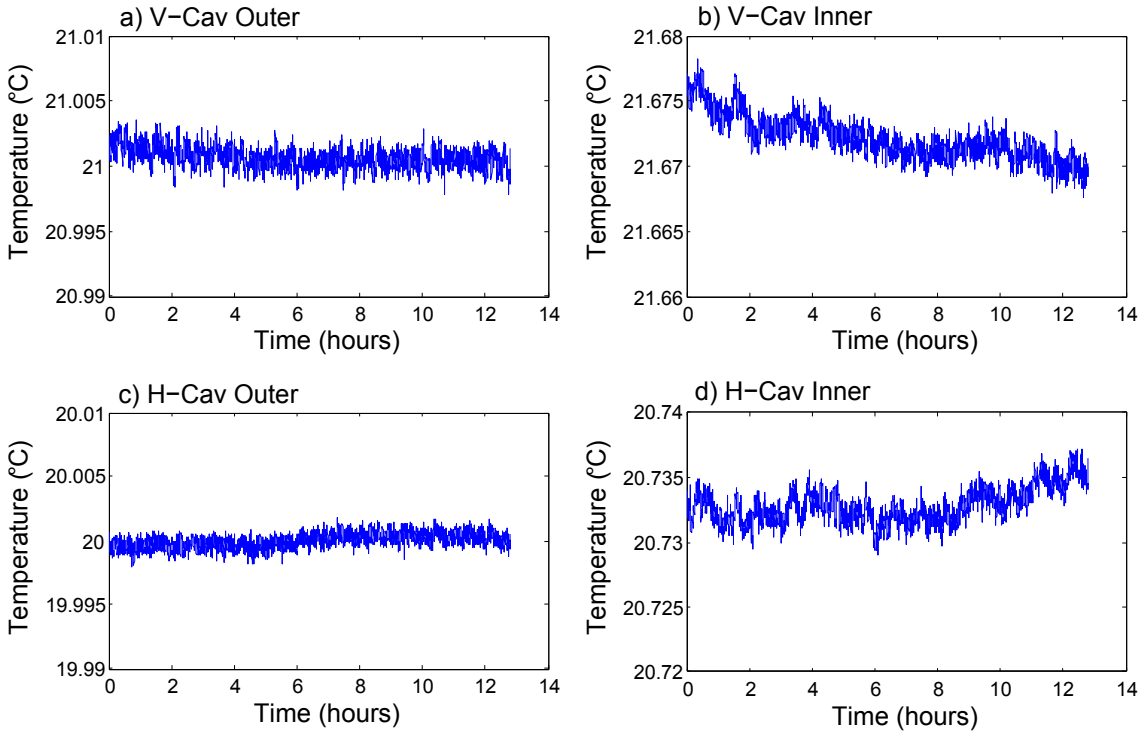

Figure 5.10: Temperature measurements of optical cavity vacuum systems over a period of 12 hours. The *Outer* readings are from the control sensors positioned next to the active peltier elements on the outer heat shield. The *Inner* readings are from a sensor on the inner shield which will detect the passive temperature of the system. The Inner temperature appears to fluctuate more because the sensor is further from the temperature controller so there is more chance for fluctuations. The difference in temperature between two sensors is either due to a temperature gradient or an error in the calibration between the two sensors.

When the acoustic boxes were used, the additional thermal insulation caused by the acoustic foam and wooden box did not allow the vacuum systems to cool as effectively. If cooling to a low temperature the air inside the box would heat up dramatically and cooling further would not be possible due to the increased air temperature. To ensure that the optical cavities had a constant temperature both systems were operated at room temperature; the stability over 12 hours is shown in figure 5.10. The *Outer* sensor used for feedback clearly stays close to the temperature wanted. The measurement that will give a more realistic reading of the thermal stability is the *Inner* measurement, as this is further from the peltier element giving temperature feedback. It would be expected that the inner and outer sensors would give the same temperature, but due to minor differences in the calibration slightly different temperatures are measured. The Inner measurements show that both systems have a thermal stability of approximately 5 mK over 12 hours. The frequency shift due to a change in temperature will depend on how far the system is from the thermal expansion turning point temperature. If the optical cavity is within 10 mK of the thermal expansion turning point temperature the sensitivity will be 50 Hz/mK [50], therefore a long term stability of 250 Hz could be achieved over 12 hours for our systems if operated close to the turning point of thermal expansion.

5.10 Electro Optic Modulators

As discussed, locking to an optical cavity by PDH locking requires that we modulate the light at a megahertz frequency; this can be done either by modulating the current supplied to the laser diode or by using an electro-optic modulator (EOM). If the current to the laser diode is modulated the intensity will also be modulated; the amplitude modulation causes an offset in the error signal. We therefore use an EOM otherwise known as a Pockels cell [93]. A Pockels cell is a crystal with a high electro-optic coefficient; when a voltage is applied across the

crystal the refractive index changes and this causes a phase change of light passing through the crystal; by modulating the voltage applied across the crystal the phase of the light will therefore be modulated. This will cause sidebands on the light to be formed on either side of the main carrier at a frequency offset from the carrier equal to the modulation frequency. The difficulty with EOMs is choosing the correct crystal, as different crystals have different electro-optic coefficients and even within the same crystal there will be different electro-optic coefficients on different axes. With careful selection of the crystal and the direction it is cut, a crystal with the desired characteristics can be found; crystals normally used to make EOMs are Lithium Niobate ($LiNbO_3$) and Lithium Tantalate ($LiTaO_3$).

Light entering the crystal must have its polarisation in the plane of the electric field applied to the crystal for the light to be modulated; we require alignment of the polarisation to the crystal to be precise. Passing the light through a Glan-Thompson polariser before the EOM cleans the polarisation ensuring only linearly polarised light at the correct orientation is used. A Glan-Thompson polariser is used because it has a very high extinction ratio removing the unwanted polarisation. Birefringence can cause amplitude changes of the transmitted light if two different light polarisations are used; birefringence is where different polarisations of light travel at different speeds through the crystal and interference can occur between these beams. Amplitude modulation may arise from etalon effects in the EOM crystal; light reflecting from the output facet of the crystal will interact with the incoming light causing interference. To avoid any etalon effects the crystal can either be anti-reflection coated or Brewster cut at the faces; having incident light on the crystal facet at the Brewster angle will help ensure the desired polarisation is transmitted. The modulation index β of an EOM is calculated from the electric field E applied to the crystal. The capacitance of the crystal is given by $C = \epsilon A/d$, where d is the distance between the conducting plates, A is the area of the conducting plates and ϵ is the permittivity of the crystal. The electric field is $E = \frac{1}{2}CV^2$

where V is the voltage applied across the crystal. The modulation index is given by

$$\beta = \frac{\pi n_o^3 r E l}{\lambda} \tag{5.5}$$

The refractive index of the crystal is n_o, the electro-optic coefficient is r and l is the length of the crystal. From this we determine that a small but long crystal will give the largest phase modulation for a given applied voltage.

For our experiment we tried using several EOMs. Initially we used a home built EOM made with a Lithium Niobate crystal of size 2 mm × 2 mm × 20 mm, with two of the long sides gold coated. To give a high enough modulation the crystal was built into a resonant circuit to increase the voltage potential across the crystal. The resonant circuit was constructed to supply a maximum voltage at the desired modulation frequency; a variable capacitor was used to finely tune the circuit's resonant frequency. The EOM would occasionally produce large amounts of amplitude modulation as well as phase modulation. To try and resolve the problem we studied the frequency response of the crystal using a network analyser. A network analyser outputs a series of frequencies over some range and detects the response of the system to each of these frequencies. The crystal was connected between the signal contacts of the output and the input of the network analyser while the ground wires were shorted together. The data from the network analyser is shown in figure 5.11. As the crystal is constructed with two parallel conducting plates separated by the insulating crystal, it is unsurprising that the EOM acts like a large capacitor; the transmitted signal increases as the input frequency increases however small peaks can also be seen on this linear trace.

The presence of these peaks at particular frequencies is believed to be due to piezoelectric resonances of the crystal. Lithium Niobate as well as being an electro-optic material is also a piezo-electric material so when a voltage is applied the crystal will change in size. As we are applying an oscillating signal we find mechanical resonance modes in the crystal. The

Figure 5.11: The frequency response of a straight cut EOM crystal. This is the network analyser scan of frequencies from 0 MHz to 20 MHz; the y-axis is in log scale. The EOM has a increase in response as the frequency increases due to the system acting like a capacitor; peaks of piezoelectric resonances can be seen at particular frequencies.

speed of sound in a solid is given by $v = \sqrt{C/\rho}$ where C is the bulk modulus and ρ is the density. For Lithium Niobate [94] we find that the speed of sound is around 7000 m/s. Our modulation frequency $\Omega/2\pi$ gives the wavelength of the modulating signal in the material

$$\lambda_{\text{mod}} = \frac{v}{\Omega/2\pi} \tag{5.6}$$

and hence we can calculate the number of wavelengths in the system for a given crystal dimension l

$$\text{Number of Longitudinal Modes} = \frac{l}{\lambda_{\text{mod}}} \tag{5.7}$$

For our crystal of 2 mm $\times$ 2 mm $\times$ 20 mm, we can be certain that the modes were in the direction perpendicular to the conducting plates with a distance of $l = 2$ mm as the same frequencies of resonance are found if a crystal of 2 mm $\times$ 2 mm $\times$ 10 mm is used; this would

Frequency (MHz)	Number of Wavelengths
1.85	1/2
5.24	3/2
8.86	5/2
12.28	7/2
15.90	9/2
19.39	11/2

Table 5.1: The resonant longitudinal modes of a $LiNbO_3$ crystal seen in figure 5.11 and the corresponding number of wavelengths within the crystal.

not have been true if the modes were longitudinal. We measure that the modes occur at specific frequencies and that these relate to a number of wavelengths as shown in table 5.1. Why the resonances only occur when the wavelength is at a half mode is unclear but a link can certainly be shown.

When a voltage is applied to a piezo-electric crystal it will deform slightly; if this voltage is oscillating and is at a frequency that is at a mechanical resonant mode of the crystal, the crystal will change shape significantly. When a crystal changes shape the birefringence can also change; the modulated signal will cause a constantly changing birefringence to occur and cause amplitude modulation at frequencies where the crystal resonates.

A Brewster cut crystal was constructed to see if an etalon effect was causing the amplitude modulation; an image of this EOM is shown in figure 5.12. The facets of the crystal are cut so that linear polarised light in the plane of the crystal light hitting the crystal surface at Brewster's angle will be totally transmitted to the crystal with no reflection; any other polarisation will be partially reflected. Having the EOM crystal cut in this way stops any etalon effects because if light is reflected from the facet of the crystal, it will not be reflected back down the crystal and therefore unable to be interact with the incident light. The Brewster cut crystal was more difficult to align than the straight cut EOM as the angle of the incident light had to be carefully set to have light transmitted through to the other side. When the network analyser was used to see the frequency response of the Brewster cut

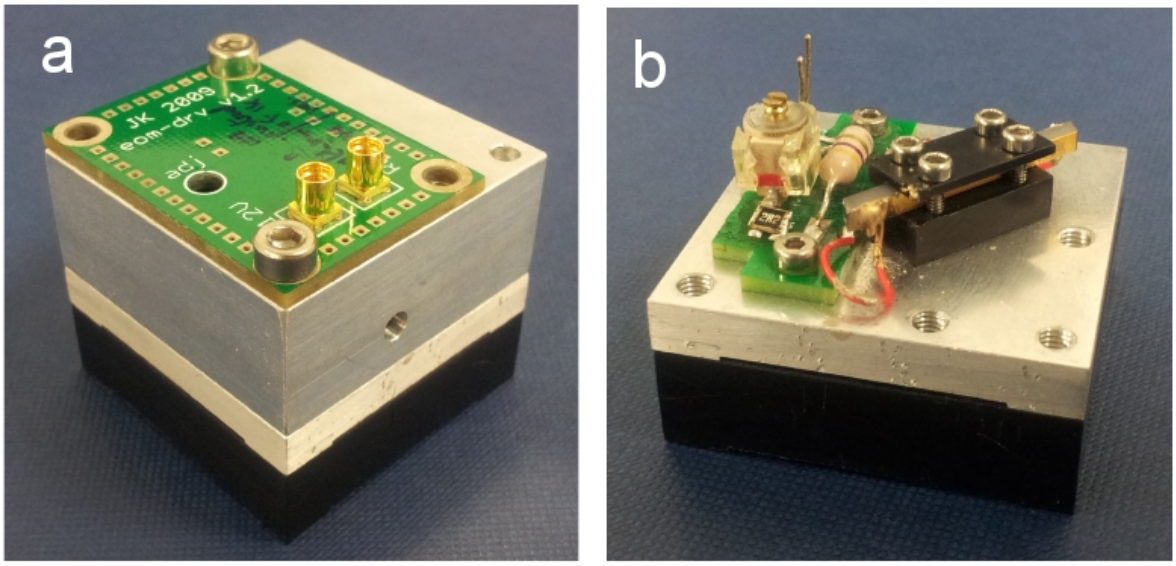

Figure 5.12: a) The EOM crystal in its home built aluminium housing; on top is the electronics for the amplifying signal for the driving voltage. b) The open EOM; visible are the Brewster cut EOM with gold coated sides and the resonant circuit components. A variable capacitor is used so the resonant frequency can be varied; by measuring the current drawn by the electronics the resonant peak of the electronics can be adjusted; this is also useful to alter the amount of modulation needed to achieve the optimum error signal.

crystal, a greater number of peaks can be seen than in the straight cut crystal as shown in figure 5.13 a; the extra peaks will be associated with other longitudinal modes in the crystal due to its odd shape. To check whether these peaks in the frequency response were also giving amplitude modulation we aligned light from a laser through the crystal and detected the intensity of the transmitted light with a fast photodiode. The output signal from the network analyser was used to drive the EOM modulation and the input to the network analyser came from a fast photodiode positioned in the beam after the crystal. To ensure a high voltage was supplied to the crystal the resonant circuit on the EOM was used. The resonant circuit had a response peak at 15 MHz and the amplitude modulation can be seen to increase around the resonant frequency. A comparison of figures 5.13 a and b show that many of the peaks in the frequency response of the EOM crystal correspond in frequency to peaks in the transmission amplitude of the light from the crystal.

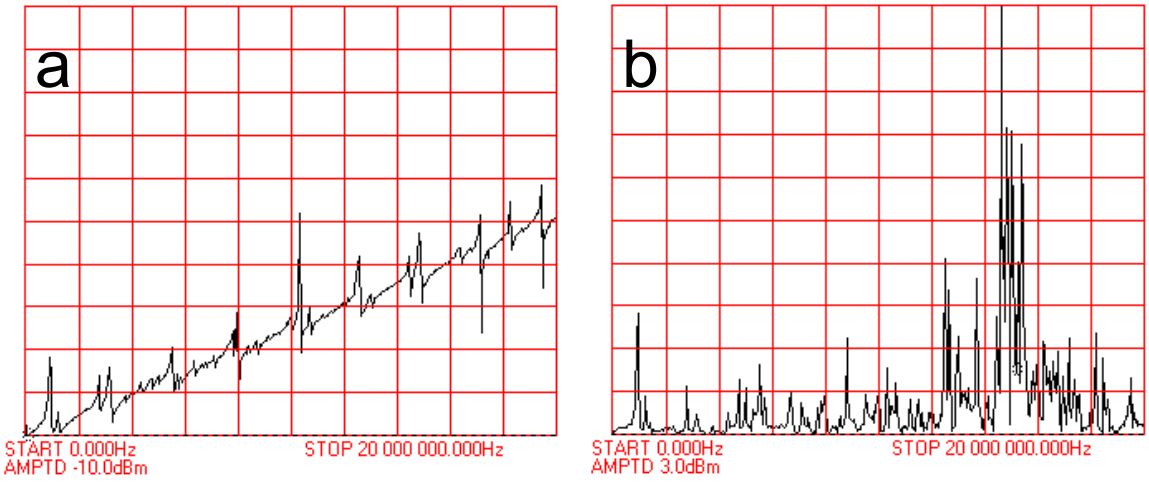

Figure 5.13: a) The frequency response of the Brewster cut EOM crystal. b) The amplitude modulation measured in the light due to the EOM. The signal can be seen to be without amplitude modulation unless there is a resonant mode in the crystal's capacitance. A correlation can be seen between these two sets of results.

A fibre coupled EOM built by Jenoptik [92] was used in the final set up; this device used a Lithium Niobate wave guide as the phase modulator. The system is built with two optical fibres connected to the crystal; the fibres had FC/APC (fibre connector, angled physical contact) compatible with the fibre coupler on the laser systems, so the fibre EOM could be swapped to replace a fibre optic without requiring the optics to be realigned. APC connectors were used to ensure there was no back reflection down the fibre. This EOM was very versatile to use, as the modulation could be set to almost any frequency or modulation voltage. The crystal is small so only low voltages were required, alleviating the need for a resonant circuit; the maximum intensity of light that could be input to the EOM was 30 mW. A major disadvantage of the fibre EOM was the significant loss of light power transmitted through the EOM, even when taking into account the usual drop in total light power due to the loss when coupled into an optical fibre, the EOM had significant losses only transmitting a maximum 20% of the light power before the fibre. However there was no measurable amplitude modulation from the fibre coupled EOM. One EOM that was not used in the final set up but was tested was a commercial EOM constructed by Photonics Solutions. This

was similar in construction to the home built EOMs; the unit was an enclosed module that included a crystal and a resonant circuit set to a frequency of 12.5 MHz. The amplitude modulation was small compared to the home built EOM but the reason for this is unclear. This EOM was not used in the final set up because the resonant frequency could not be changed, this flexibility may have been required depending on the frequency filters available.

5.11 Beat Measurement Experimental Results

The method used to measure the beat signal when both the lasers were locked to an optical cavity was the same method used to measure the free running laser linewidth as discussed in section 3.4. To characterise the stability of the optical cavity systems both lasers had to be locked to a TEM_{00} mode at similar frequencies. A wavemeter with a precision of 0.1 GHz made by Advantest was used to get the two laser frequencies close enough that the beat frequency was less than 7 GHz, the maximum span of the spectrum analyser that was used. The wavemeter was known to drift in frequency over time but this was not an issue as only the frequency comparison was required, an accurate absolute value was not required. When the beat frequency is within the bandwidth of the spectrum analyser, the alignment of the beams becomes significantly easier as the size of the beat signal is related to how well aligned the laser beams are; by better alignment of the beams with each other and alignment onto the fast photodiode the magnitude of the signal could be increased.

One issue when attempting to take measurements was the vibrations due to the air conditioning unit in the laboratory. The significant effect due to the acoustic noise shaking the experimental set up can be seen in figure 5.14. When measurements were being made the air conditioning was turned off so that the optical cavities were not perturbed.

When both lasers were locked to their respective optical cavities, different characteristics could be seen when the frequency span of the spectrum analyser was changed. A range of

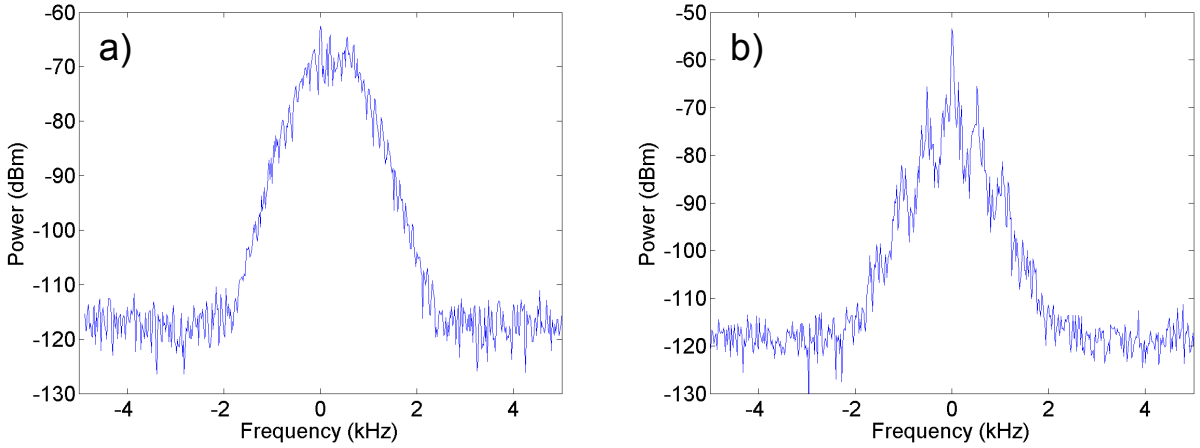

Figure 5.14: a) A beat measurement where the linewidth is broadened due to the acoustic vibration. b) Once the air conditioning is switched off the peak can clearly be seen.

these frequency spans are shown in figure 5.15. In figure 5.15 a) the modulation peaks can be seen above and below the central peak; second order harmonics of the modulation frequency can also be seen. The beat signal shows sidebands at 12.5 MHz because the beat is between the pure frequency of the vertical laser system and the modulated signal of the horizontal system as the light is picked off after the fibre coupled EOM. There is a beat frequency with the primary frequency but also beats with the sidebands. In figure 5.15 b) the lasers are locked and we can see what the bandwidth of our locking electronics is by seeing where the signal is suppressed around the central locked carrier; this bandwidth is shown to exceed 500 kHz. As we reduce the span of the spectrum analyser further features can be seen; figures 5.15 c) and d) show the central peak having features at 500 Hz from the central peak, possibly due to some electronics noise. Figure 5.15 e) is shown on a linear scale where the central peak stands out strongly against the other peaks. In figure 5.15 f) a gaussian is fitted to the central peak to measure the linewidth of the beat frequency.

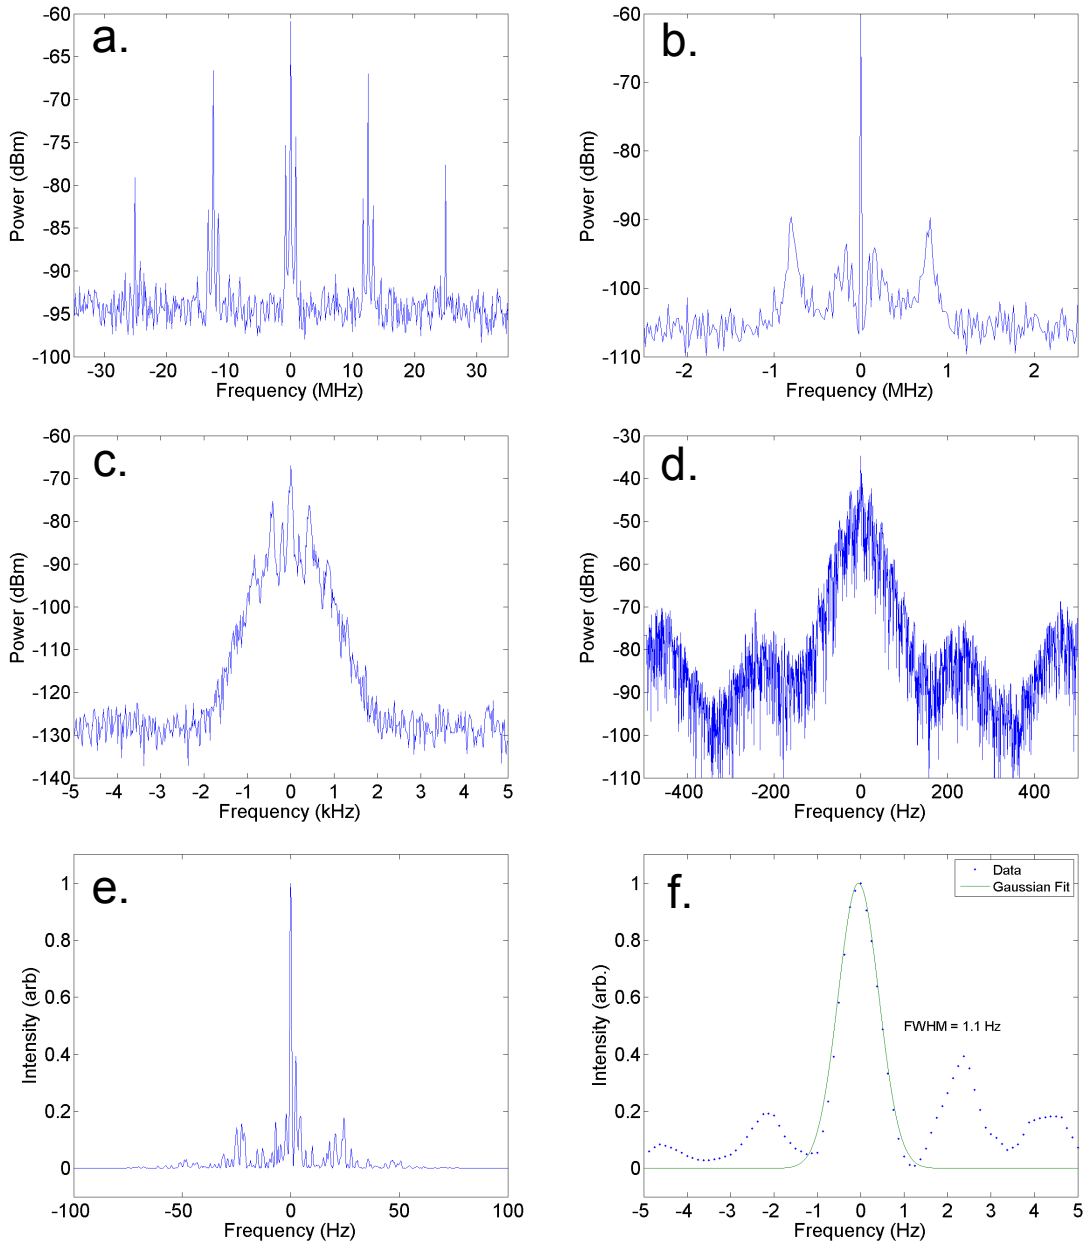

Figure 5.15: Beat measurement of the lasers for varying frequency spans; on the plots an offset of 4.41 GHz has been subtracted from the beat frequencies. a) Widest frequency span: the additional beat notes due to the harmonics of the modulation frequency are shown. b) Span around the main beat frequency: the side peaks show the bandwidth limitation of the locking electronics. c) Span in the kilohertz range showing the inner structure and the increase of noise due to the acoustic noise in the system. d) The main peak width shown in dBm scale. e) The main peak width shown in normalised linear scale. f) Smaller span image of the main peak shown in normalised linear scale with a gaussian fit to the data points.

5.12 Calculation of the Beat Frequency Linewidth

To measure such a narrow frequency the highest resolution of the spectrum analyser was required. The parameters used to take measurements of the beat linewidth were a span of 1 kHz using 8001 points and a resolution bandwidth of 1 Hz; for these settings we had an acquisition time of 4 seconds. Various filters could be used on the spectrum analyser to acquire the data; it was found that the best choice for our measurements was an FFT filter as all the channels would be sampled at the same time. To convert a power measured by the spectrum analyser in dBm to a power measured in milliwatts we use the following conversion.

$$P_{(\mathrm{mW})} = 10^{\left(P_{(\mathrm{dBm})}/10\right)} \tag{5.8}$$

Once we have the measurement in linear units of milliwatts the data can be normalised; then a gaussian is fitted to the peak in the data and the FWHM of the gaussian is calculated.

Over a period of two days linewidth measurements were taken; fifteen sets of data were used to calculate the final narrow laser linewidth measurement. The average beat note linewidth for both lasers locked to their respective optical cavities was found to be 1.37 Hz with a standard deviation of 0.28 Hz.

5.13 Calculation of Allan Deviation

The Allan deviation describes the stability of an oscillator over a range of averaging times [16][95]. We compare the fractional frequency average over a specific time interval. The fractional frequency is the measured frequency divided by the frequency of the oscillator; in our case the oscillator frequency is the frequency of the laser (429.2 THz). The Allan deviation is defined as

$$\sigma_y(\tau) = \sqrt{\frac{1}{2}\langle(\bar{y}_{n+1} - \bar{y}_n)^2\rangle} \tag{5.9}$$

where $\bar{y}_n$ is the fractional frequency given by the measured frequency at a specific time and $\bar{y}_{n+1}$ is the fractional frequency a time τ later. The angle brackets $\langle x \rangle$ denote the calculation of the expectation value of x; for discrete values with the same probability density such as our data the expectation value is simply the mean of the values.

The beat frequency was measured using a frequency counter that takes a frequency reading for a certain time period; the readings are stored by a computer using the GPIB interface on the frequency counter. From this data a comparison of frequency readings can be made. The calculation of the Allan deviation and the plotting was completed using a freely available MATLAB program available from the MATLAB file exchange [96]. For the measurements both the H-Cav and V-Cav lasers were locked to a TEM_{00} mode of their respective high finesse optical cavities. The measurements were taken for long periods but this was difficult to achieve due to various external perturbations causing the laser to cease being locked or to jump to another TEM mode; instead of the intended TEM_{00} mode of the optical cavity.

The results measured are presented in figures 5.16 and 5.17; it can be seen that 2×10^{-14} is the limit of the laser instability; this would indicate a short term linewidth of around 10 Hz, which is significantly greater than the direct measurement of the beat note linewidth but this may be due to the way in which the heterodyne beat frequency width was measured. The short term frequency stability of the frequency counter is $< |1 \times 10^{-9}|$ rms for one second of averaging time [97]. For a beat frequency of 1 GHz this is a stability of 1 Hz which corresponds to an Allan deviation of 2×10^{-15} for our frequency of light. As the frequency counter is of significant age, and has not been calibrated, the stability could be worse than quoted on the data sheet; there is a good chance this was limiting our measurement. If a lower frequency beat note were to be used the measurement of Allan deviation would still be limited by the frequency counter as we are unable to measure to a resolution of less than 0.1 Hz; if we were to use a beat frequency of less than 100 MHz the limitation will be the resolution of the frequency counter. An alternative frequency counter would need to be used

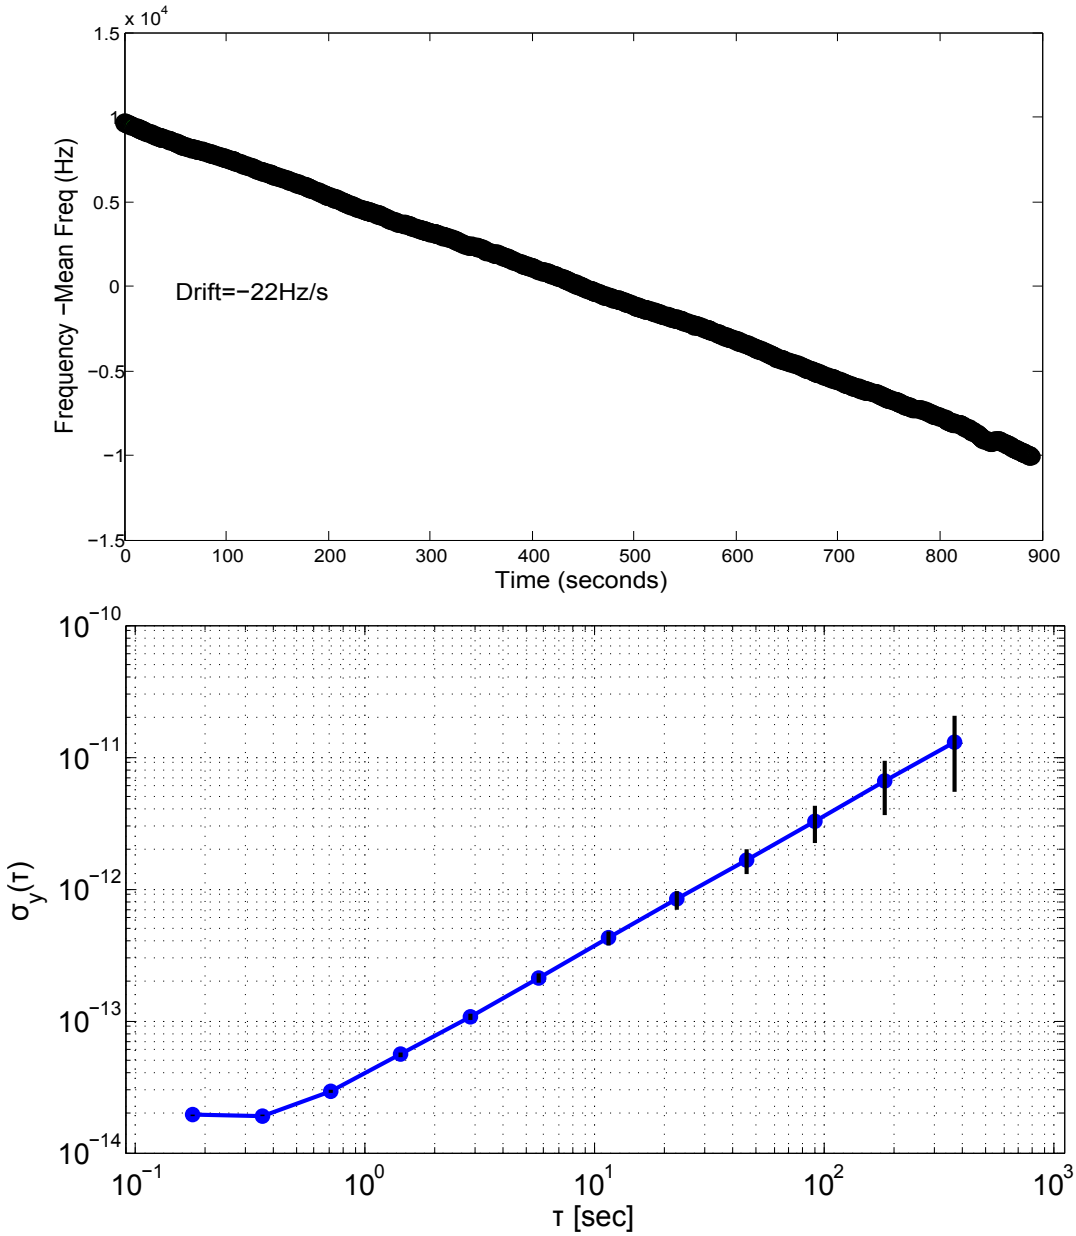

Figure 5.16: Allan deviation measurement for both lasers locked to their optical cavities. The top graph shows the frequency measurement over 15 minutes of recording the beat frequency between two lasers both locked to a TEM_{00} mode of a high finesse cavity. The frequency measurement was taken every 0.17 seconds (5.6 Hz) and a frequency reading resolution of 10 Hz was used. The beat frequency was at 3.36 GHz, a drift of -22 Hz/s was seen on the data. The bottom graph is the Allan deviation calculated from the frequency counter data. The noise floor appears at 2×10^{-14}; however this was only observed at less than 0.4 seconds due to the high rate of drift of the frequency.

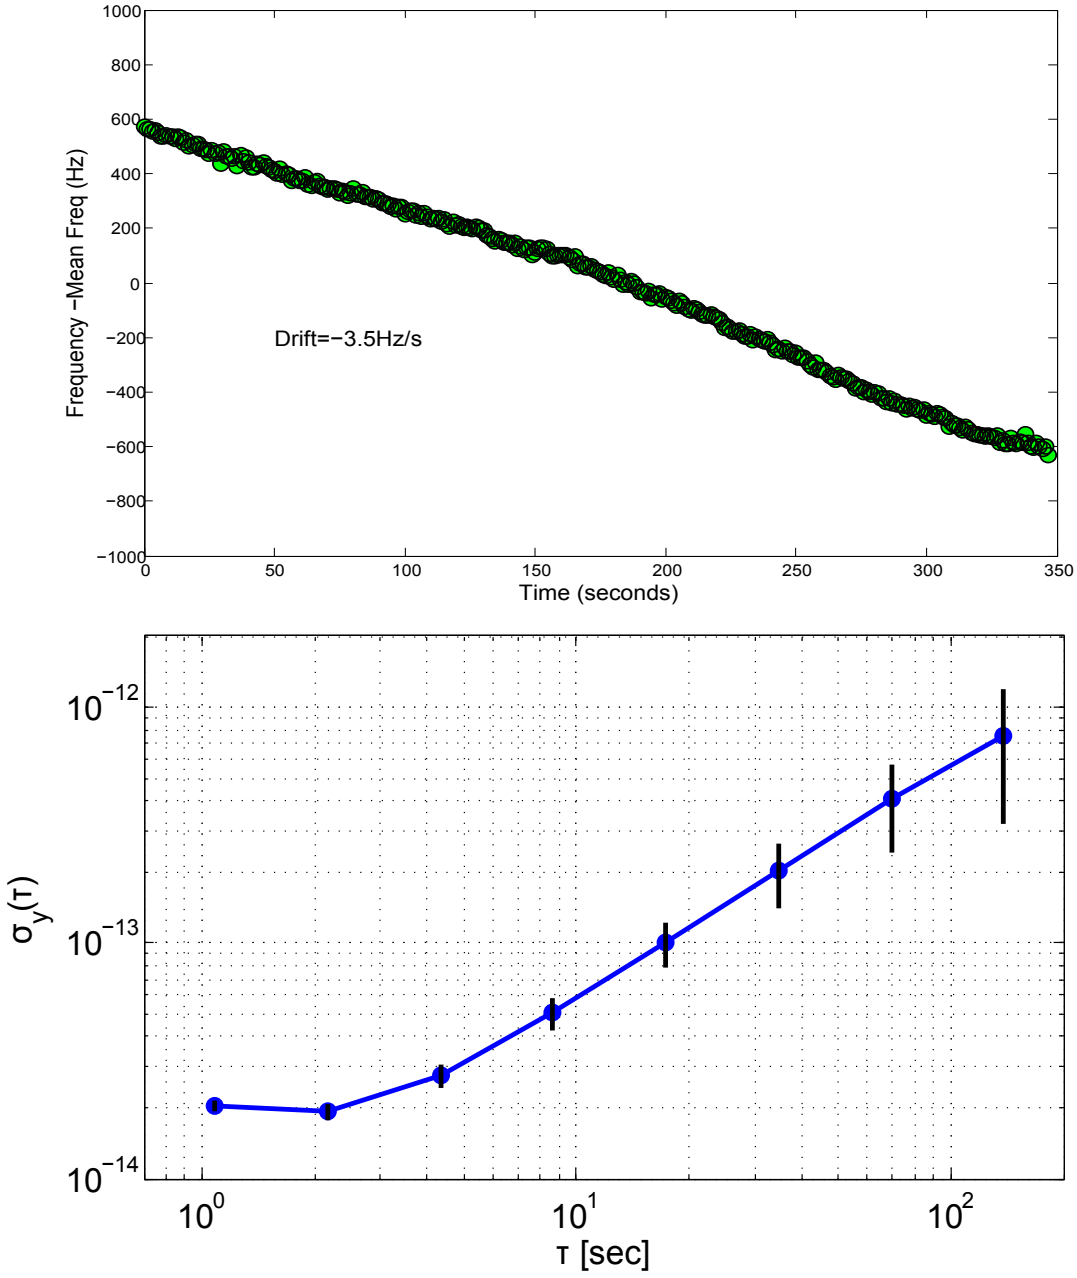

Figure 5.17: Allan deviation measurement for both lasers locked to their optical cavities. The top graph shows the frequency measurement over 6 minutes of recording the beat frequency between two lasers both locked to a TEM_{00} mode of a high finesse cavity. The frequency measurement was taken every 1.018 seconds (0.92 Hz) and a frequency reading resolution of 1 Hz was used. The beat frequency was at 5.54 GHz, a drift of -3.5 Hz/s can be calculated from the data. The bottom graph is the Allan deviation calculated from the frequency counter data. Once again the noise floor can be seen at 2×10^{-14} but the lasers remain at that stability for times less than 2 seconds.

if a better precision is required when the stability of the optical cavity locking is improved. Also it was unknown what the dead time of the measurement was so the possible bias could not be calculated which may have shifted the Allan deviation measurement [98]. A further reason for the stability measured by the frequency counter being much greater than the linewidth measurement would be the bandwidth of the respective measurements. As the frequency counter will measure all frequencies transmitted by the filters, between 0 Hz and 3 GHz, the noise from these frequencies will be include in the measurement. The spectrum analyser had a resolution bandwidth of 1 Hz therefore the noise is lower.

The smallest frequency drift measured was -3.5 Hz/s. A significantly smaller frequency drift due to optical cavity length changes has been published [50] where the frequency drift was two orders of magnitude better than our measurement; they were able to achieve this because of a better temperature stability used for the control of their optical cavities and one of the optical cavities was held at the thermal expansion turning point.

To improve the characterisation of the stability measurements a faster frequency counter could have been used to take frequency measurements with a τ of less than 0.1 seconds; however a suitable device was not available. As stated beat measurement and Allan deviation showed some disagreement. A 4 second acquisition time was used and, as shown by the frequency counter, it would be expected that the drift of several hertz over this time scale would cause a wider linewidth to be measured but, this was not the case. It is also possible that as the Allan deviation measurements were taken over a short time period, the linewidth measurements were taken at a time when there was less drift than during the beat frequency measurements.

5.14 Comparison of Results

If our laser were only limited by thermal noise, for an optical cavity made with a ULE spacer and ULE mirror substrates, the frequency noise would be 0.36 Hz/$\sqrt{\text{Hz}}$ as calculated in section 4.8. Using equation 4.12 this corresponds to a fractional frequency instability of 8.2×10^{-16} for the 698 nm laser system measured with a 4 second acquisition time, as used by the spectrum analyser. However as has been shown the beat frequency linewidth measurement between the two locked lasers is currently 1.37 $\pm$ 0.28 Hz for a 4 second measurement time. This directly measured linewidth would be equivalent to an instability of 3.2×10^{-15}.

The measurement of the Allan deviation showed that for times of less than 2 seconds there is a fractional frequency instability of 2×10^{-14} (compared to a theoretical limit of 1.2×10^{-15} at 2 seconds). The variation in the results from the linewidth measurement and the Allan deviation show the frequency must be varying in a way that is not shown by the linewidth measurement, probably due to the short time period over which the measurement was taken or that the stability of the frequency counter limited the measurement. Either result shows that our laser is not at the theoretical limit of instability and noise sources will need to be characterised to reduce this instability down further.

In comparison, a similar system with a ULE optical cavity [50] achieved a 0.5 Hz optical beat note linewidth with Allan instability of 2×10^{-15} between 0.1 and 10 s averaging time. The very best laser instability using a cooled silicon optical cavity [65] was measured to have a linewidth of less than 40 mHz and an instability of 1×10^{-16}.

CHAPTER 6

CONCLUSIONS

This thesis has presented the difficulties and limitations of building a stable narrow-linewidth laser using an ultra stable Fabry-Pérot cavity. The progress of developing a neutral strontium optical lattice clock at the University of Birmingham is limited by the ability to produce a narrow linewidth laser and I have gone some way to addressing this problem.

I have designed, built and characterised a laser system suitable for use with an ultra high finesse optical cavity. The design of the extended cavity diode laser system is compact and rugged; linewidths as low as 26 ± 15 kHz have been measured for the free running linewidth of the laser system. I made a comparison of a *long* and a *short* system and it was shown that increasing the length of the ECDL cavity did improve the linewidth; however the mechanical stability clearly limited the laser stability as the reduction in linewidth for the laser system was not as much as predicted by the Schawlow Townes linewidth.

It was previously known that the limiting factor for the instability of an optical cavity is the mirror coating thermal noise. I have performed extensive numerical analysis to show that by using alternative materials to make the optical cavity, the noise due to thermal fluctuations can be reduced. I have calculated that an all silicon optical cavity would perform extremely well and have a frequency noise of 0.15 Hz as limited by the thermal fluctuation of the optical coatings. Silicon can transmit infra-red light but is opaque to our frequency of light. For

our purpose, to produce a laser to probe the strontium clock transition at 698 nm, either a complex frequency doubling system or frequency comb would be needed or a different optical cavity would be required. The best option would be to use a fused silica optical cavity for which I have calculated the theoretical fractional frequency instability to be 3.5×10^{-16}.

I have also performed finite element analysis for various supporting points for one of our optical cavity designs, and have found the optimum position to support the optical cavity such that vibrations transmitted to the optical cavity have a minimum effect on the length of the optical axis.

I have built a temperature controlled vacuum system from conventional vacuum components for one of our ultra stable optical cavities, and have also designed and built a miniature vacuum system from custom made components to house the other ultra stable optical cavity. I tuned the temperature control parameters and was able measure a temperature stabilised to within 5 mK over a 12 hour period for both systems.

During the investigation into the electro optic modulators I was able to clarify many of the issues that the group encountered when using these phase modulators. I have showed that mechanical resonances in the $LiNbO_3$ occurs at the same frequencies where light passing through the crystal will undergo amplitude modulation. This is assumed to be due to the piezo-electric effect that changes the shape of the crystal and causes an amplitude change in the light, but the exact mechanism by which this happens is not completely clear.

The key result of this thesis is that by using the home built diode laser systems and commercial locking electronics, I was able to lock to two stabilised Fabry-Pérot optical cavities and make a heterodyne beat measurement between the two lasers. The linewidth of the heterodyne beat measurement was shown to be 1.37 ± 0.28 Hz. An Allan deviation measurement showed that the two laser systems could reach an instability of 2×10^{-14} for times less than 2 seconds. The best drift rate measured for the system when the temperature was stabilised was shown to be -3.5 Hz/s. To have confidence in these results we need to

fully characterise the spectrum analyser to investigate if there is any dead-time during the measurement. Currently we have sufficient results to demonstrate that an extremely narrow linewidth laser has been developed and built.

There is still much to do to finish the optical strontium clock. The outlook for the future of the project will be to improve the stability of the laser system and be able to use it to make a frequency measurement of the clock transition within strontium.

To increase the stability of the laser systems further, work needs to be done to improve the thermal stability of the optical cavities. Cooling the systems down and taking drift rate measurements at different temperatures will enable us to find the turning point of the thermal expansion for each of the optical cavities; operating at this point will ensure the optical cavities have a smaller thermal drift and will therefore be more stable. However, the systems were limited in how low a temperature they could be cooled to by insufficient heat removal from the systems; future designs will need to be developed so that the zero-crossing temperature of the ULE spacer can be reached. Much work needs to be done to reduce the temperature of the acoustic box; conventional methods such as water cooling or fans will produce vibration. Care must be taken that these vibrations are not transferred to the apparatus within the box. To narrow the laser linewidth further work needs to be done to improve the locking stability; this may involve examining the electronics, to change the equipment if needs be to remove any limiting factors.

The supporting position for the H-Cav system will need to be adjusted and have vibrations applied to the optical cavity to measure the response to shaking. The model shown in section 4.9 will be a good guide; however it is expected that the measured point of stability will be different from the supporting point found in the model.

The optical properties of the light input to the optical cavity needs to be investigated. The mode matching to the cavity could be improved to increase the amount of light coupled into the optical cavity. We also need to stabilise and reduce the light input power to the optical

cavity as the resonant frequency is dependent on the intensity of the incident light [50].

To be able to investigate the Allan deviation at shorter time periods a frequency counter with a higher frequency sampling rate will be required. To improve the linewidth measurement a spectrum analyser with a lower resolution bandwidth would also be required as the current apparatus is limiting the precision of our measurement to a 1 Hz bandwidth.

The next stages of this experiment will be to use the constructed laser system with a strontium clock set up with cooled atoms so that a measurement of the clock transition frequency can be carried out. At this stage a cold atom cloud of strontium has been produced in the laboratory but no optical lattice has yet been used. Once this is achieved and a clock measurement made the portability of the system can be tested. A long term aim would be to transport the clock to another laboratory with another strontium clock so a direct comparison of the two clocks can be made.

APPENDIX A

HOW TO USE THE IGUANA TEMPERATURE CONTROLLER

A.1 Introduction

The Iguana temperature controller is used to control the temperature of the horizontal and vertical optical cavity vacuum systems. It has 5 analogue to digital converters (ADC) for temperature input and 5 digital to analogue converters (DAC) for current outputs. The system is based on a PIC micro controller within the Iguana box and all control is via a computer interface. Parts of the controller operation are also discussed in section 5.9 of the main thesis.

A.2 Set up of the Iguana

The Iguana temperature controller inputs are for an AD590 temperature sensor using a female sub-D 9 pin connector; the positive side is connected to pin 1 and the negative to pin 5. The output is a bipolar current output with a range of ± 1 A; there is a connection on pins 1 and 5 of a male sub-D connector. The polarity will depend on which way the peltier is orientated so the easiest way to find the polarity is by trial and error. It is possible to

change the polarity in the control program but I rarely did this. There are outputs 0,1,2,3,4 and inputs 0,2,4,5,6; the missing inputs are used for the temperature sensors on the circuit board to ensure the DACs do not get too hot.

A dedicated computer with a Linux installation is used for the program; control is via the command line. The connection to the Iguana is via a RS-232 cable. To change control parameters we edit the configuration file. The default text editor for the computer is `ed`, the control file is `zoo.conf`. To find the file use the following commands:

```
cd /home/zoo/etc/
ed zoo.conf
```

To open a previously opened file at the same place it was closed use the command `ed` without the file name; this will automatically open the last edited document. Once changes have been made to save a file use **Ctrl+X, Ctrl+S**; to close the file use **Ctrl+X, Ctrl+C**.

Within the `zoo.conf` file are many animals; this is where it gets interesting. The system was developed to control a ground based telescope for solar observation. There are animals that control the dome movement, the telescope position and the temperature of components; in addition to animals that control other devices. We use the Iguana because it controls the temperature of a body. Fairly low down in the file `zoo.conf` is the Iguana section with the inputs and outputs.

To run the program to enable the Iguana to control the temperatures we need to be in the directory `/home/zoo/etc` and use the command `iguana`. To view the temperatures we use a 'paddock' to see the animal characteristics; use the command `squid` to view the paddock.

A.3 Administrator Rights

There are three levels of Administrator: `visitor, keeper` and `vet`. The visitor can view all data from the system but cannot make any changes. The keeper is required to make changes

via an `echo` command when changing the temperature of the system. The vet is required to change any parameters in the `zoo.conf` file. To change what level of rights are used type the level into the command line and the admin level will change.

A.4 Changing Parameters

The input ADC channel parameters in the file `zoo.conf` are shown below. This is shown for channel **0** and is of the form

```
Iguana.adcChannel0.name:    V-Cav Out
Iguana.adcChannel0.scale:   0.0000602262
Iguana.adcChannel0.offset:  4535402
```

Also in this block are warnings but these are not used. The temperature reading T is calibrated by using the ADC value a,

$$T = (a - \text{offset}) \times \text{scale} \tag{A.1}$$

The AD590 sensor is made such that it has a good linearity but can have a manufacturing offset. If we heat the system and cool the system to specific temperatures we can calibrate the temperature sensors. If left for a long time we can assume that the inner and outer sensors will be close in temperature and a calibration can be done. Using two temperatures T_1 and T_2 and the temperature measurements from another sensor T_1' and T_2' we can calibrate the second sensor to match the first sensor by changing the parameters to the new scale$'$ and offset$'$ parameters.

$$\text{scale}' = \frac{T_1 - T_2}{T_1' - T_2'} \tag{A.2}$$

$$\text{offset}' = \frac{T_1' T_2 - T_2' T_1}{T_2 - T_1} \tag{A.3}$$

The DAC channels have many aspects that can be changed. This is shown for channel **0**.

```
Iguana.dacChannel0.engine:         detector
!Iguana.dacChannel0.engine:        fixed
Iguana.dacChannel0.dacValue:       0x800
Iguana.dacChannel0.adcChannel:     2
Iguana.dacChannel0.setpoint:       20
Iguana.dacChannel0.startPower:     0
Iguana.dacChannel0.time1K1P:       3
Iguana.dacChannel0.fullScaleTemp: 0.1
Iguana.dacChannel0.minPower:       -100
Iguana.dacChannel0.maxPower:       100
Iguana.dacChannel0.minTemp:        5
Iguana.dacChannel0.maxTemp:        40
```

The output can either be based on the input from a *detector* or be a *fixed* value, both these options are shown but the fixed option has been commented out with an `"!"`. If a fixed value is used the `dacValue` dictates the power supplied, this value is in hexadecimal between `0x000` giving maximum cooling and `0xfff` giving maximum heating (depending on having the correct polarity); `0x800` corresponds to no power being supplied. The parameters shown are for a `setpoint` temperature of 20 °C. The `startPower` is used if we wish the controller to begin supplying a power before the PI control has started to calculate the power to supply. The output power $P(t)$ of the Iguana is calculated by comparison of the control sensor temperature $T(t)$ with the set-point temperature S. There are two parameters that can be changed: the `time1K1P` C and the `fullScaleTemperature` F. The power supplied is given by this equation

$$P(t) = \frac{1}{F}\left(T(t) - S\right) + \frac{1}{100C}\int_0^t \left(T(t') - S\right) dt' \tag{A.4}$$

To use the whole range of the current supply (± 1 A) the `minPower` and `maxPower` are set to -100% and $+100\%$ respectively. Changing values in `zoo.conf` requires re-running `iguana` to implement the changes of the values. The `minTemp` and `maxTemp` are safety features, such

that if the temperature of the control channel goes beyond these limits the system power will stop being supplied to the peltier element.

To change parameters 'on-the-fly' it is necessary to be in the administrator level `keeper`. To view all the channels and their output use the command

```
ed /home/zoo/var
remora "cat iguana"
```

To change a value such as the set point temperature use the `echo` command with the correct channel name.

```
echo chan0_setpoint=18
```

The temperature should be changed via this method; changing other values in the `remora` list should be done like this. Making these changes in `zoo.conf` may not work without a reset of the system.

A.5 Viewing Results

All the results are saved to the directory

```
/home/zoo/Results/
```

The file name will be of the form *'location' 'day/night' 'year' 'month' 'day'*`.dat`. The location will always be `b` for Birmingham; for day use `w` and for night use `x`. What classes as day and night is defined by sunrise and sunset as the controller was originally used for solar observation; this time will obviously change from day to day. Also the time is always UTC so be aware of this during the summer. The year, month and day are expressed in two digit format. To look at the results for the day of 14th August 2012 we use

```
ed bw120814.dat
```

There will be columns with the time and the temperatures read by all the system's temperature inputs, all temperatures will be expressed in °C.

The time is expressed in hours; when viewing overnight negative time should be used to see times before midnight. To see the whole graph times `-100 100` can be used as arbitrarily large numbers. We need to have a column number between 1 and 6 to see the desired value. For example to view the data from the 14th August 2012 between times 10:30 and 14:00 of column 4 use the command

```
tail -f -n 100000/home/zoo/Results/bw120814.dat|colselect 1 10.5 14 codf -p=c=4
```

The `tail` command allows us to view a graph of the results that will update as more results are written into the file. To close the graph use the **Esc** key and use **Ctrl+C** to stop the program running. A different set of results cannot be viewed without stopping the previous program.

A.6 MATLAB Script

A short MATLAB script has been written so that the temperature files can be viewed on other computers.

```
% Iguana_Read.m
%
% File written to read in files directly from the Iguana recording of temperatures.
% The four graphs of the four temperature channels will be produced on one plot.

filename='bw120720.dat';

[fid,message]=fopen(filename,'r');
disp(message); %Will tell us if the file path is wrong.

%Read in first line and the file with many columns
textscan(fid,'%s',1,'delimiter','\n');
```

```
data=textscan(fid,'%n %n %n %n %n %n %n %n %n %n');

%Name of each file for each column
names={'time','a)V-Cav Outer','b)V-Cav Inner','c)H-Cav Outer','d)H-Cav Inner'};

time=time-time(1);
limit=1;

for n=2:5
subplot(2,2,n-1);
plot(time,data{n});
min=mean(data{n})-limit;
max=mean(data{n})+limit;

title(names(n),'fontsize',13);
ylabel('Temperature (\circC)','fontsize',13);
xlabel('Time (hours)','fontsize',13);
end
```

APPENDIX B

HOW TO USE THE FREQUENCY COUNTER

B.1 Introduction

Using an EIP 575 Source Locking Microwave Counter it is possible to produce graphs of the Allan deviation. The EIP has a frequency range of up to 18 GHz; however amplifiers were required to raise the amplitude of the signal to a minimum of -30 dBm so the signal could be measured. The amplifiers had a maximum bandwidth of 3 GHz. On the microwave counter is a rotational knob; turning this all the way anti-clockwise causes the gate time (i.e. measurement time) to be as short as possible. To connect to a computer a GPIB to USB interface is used, a "Prologix GPIB-USB Controller" was the one that was able to operate with the Linux system. The instructions here are for a Linux command line interface.

When connected, the USB-GPIB should appear in the list of devices, in our case called `ttyUSB0`. The address for the frequency counter to use is pre-set 16. The frequency counter is made to listen to the controller device using the `da` command. This should cause a **RMT** light on the frequency counter to light up and the front panel buttons to become unresponsive, pressing **Reset** on the frequency counter will make the counter respond to panel buttons but it will no longer listen to commands sent for the computer. The resolution can be changed by the computer by sending a command to say how high a resolution is wanted; sending `R0`

gives 1 Hz resolution, sending `R2` gives 100 Hz resolution. A frequency counter measuring for 1 second can theoretically only measure to 1 Hz precision, measuring for 0.1 seconds can only measure to 10 Hz precision. On this particular counter a resolution of 1 Hz has a data rate of 0.9 Hz, having a resolution of 10 Hz has a data rate of 5 Hz. The readings can be displayed on the command line and also saved to a file. The file only contains the frequencies measured; it is required to measure the time taken with a stopwatch so that the data rate can be calculated. Once the device is connected to the computer via the GPIB-USB the sequence of commands to use is shown below:

```
echo "addr 16" > /dev/ttyUSB0
echo da > /dev/ttyUSB0
echo R2 > /dev/ttyUSB0
cat /dev/ttyUSB0 | tee filename.txt
```

B.2 MATLAB Script

The analysis is done with a MATLAB program `allan.m` available from MATLAB Central by searching for Allan Deviation. However to get the data read in the correct format to be used with `allan.m` a short script has been written to read in the data file and call the function.

```
% function [DATA_OUT,TAU,rate,drift]=allan_read(time,filename,savename)
%
% A file to take data from the EIP 575 Source Locking Microwave Counter and
% convert this to a graph of Allan deviation.
%
% --Outputs--
% DATA_OUT - Structure file will required data for allan.m
% TAU -      An array of times at which the ADEV was calculated
% rate -     The frequency in Hertz of at which the measurement was taken
% drift -    The linear drift in the system taken away before ADEV calculated
%            expressed in Hz/s
```

```
%
% --Inputs--
% time -     The period of time taken to do the measurement in SECONDS only
% filename - Name of the file where the data is held for a measurement
% savename - Name to appear on the figures produced
%
% Steven Johnson 22/06/12 (Re written on 05/07/12)

function [DATA_OUT,TAU,rate,drift]=allan_read(time,filename,savename)

laser_freq=429.2e12;

[fid,message]=fopen(filename,'r');
disp(message); %Will tell us if the file path is wrong.

readindata=textscan(fid,'%n'); % Read strings delimited by a carriage return

DATA=readindata{1};

%Transpose data to have colomb vector
DATA=DATA';

%Time in Hertz per reading
rate=(size(DATA,2)-1)/time;

%TAU is the values at which to calculate the ADEV, Make time intervals of t*(2^n)
n=1;
TAU(n)=1/rate;

while TAU(n) < time
    n=n+1;
    TAU(n)=(1/rate)*2^(n-1);
end
TAU(n)=[]; %Remove last two values
TAU(n-1)=[];

disp(['Data Readings = ',num2str(size(DATA,2))]);
disp(['Data Rate = ',num2str(rate),'Hz']);

%Calculate the drift rate of the data
X=1:size(DATA,2);
P=polyfit(X,DATA,1);
drift=(polyval(P,size(X,2))-polyval(P,1))/time;
```

```
%Make the data a ratio of change by dividing by the laser freq
DATA=DATA./laser_freq;

%Make data structure for allan.m
DATA_OUT=struct('freq',DATA,'rate',rate);

%Call the allan.m function to produce the graphs of drift and Allan deviation
allan(DATA_OUT,TAU,savename,3);

%Save the plot of Allan deviation as a png file
filename=strcat(filename(1:7),'.png');
is=strcat(filename);
saveas(gcf,is,'png');
```

APPENDIX C

HOW TO MAKE A BEAT NOTE BETWEEN LOCKED LASERS

C.1 Introduction

Making a beat note between two locked lasers is a relatively complex task so I hope by writing a guide this will not only clarify any of the issues encountered when using the set up constructed for this thesis, but that the solutions found could also be used for other experimental set ups.

I will discuss the following: the use of the EOMs and how to gain the best error signal, how the other electronics are used to scan the laser frequency and find resonant frequency in the optical cavity, how to use the FALC and change the parameters used for laser locking. Also the alignment to the optical cavity will be discussed and how to make the lasers close enough in frequency to measure a beat frequency between them. Be aware that much of this, especially the lock to the cavity, took months to achieve so much of this will take time and dedication is required.

C.2 Modulation of the Laser Light

To lock to the optical cavity we require modulated laser light. On this experiment it has been possible to lock to a high finesse cavity by directly using modulation of the laser diode current, applying a 10 MHz signal to the current driver. However it will be better to use an EOM as the light from the laser diode will be spectrally clean if not modulated.

We need to know how large the modulation strength is; the size of the sidebands can be seen using a low finesse cavity. A scanning cavity with a finesse of 200, available from Thorlabs, was used for measuring the size of the sidebands. This optical cavity was used as it takes very little time to align (easily done by correctly aligning the input beam with the output beam) and it has a built in photodiode to see the transmitted signal. One of the mirrors is mounted on a ring piezo, scanning the cavity length with a triangular wave at 10 Hz with a voltage around 1 V_{pp} makes the modes easy to see; once a mode is observed it can be increased in size by adjusting the mirrors coupling the light to the optical cavity. When coupling a laser beam to any optical cavity adjustments will need to be made in angle and position, to do this we require *two* mirrors directly before the cavity that are *easily accessible*. If a high finesse optical cavity is used, the heights of the central and sideband peaks will fluctuate because not all the light will be transmitted through the optical cavity. Due to these fluctuations making a measurement of the amplitude of the central peak or the amplitude of the sidebands is very difficult; taking a measurement at one time will produce different results to a measurement taken moments later.

According to the simulations (figure 4.4) to get the biggest error signal possible a modulation strength of $\beta = 1$ should be used; this relates to a sideband height compared to the central peak height of 40%. Speaking to other research groups a sideband height of 25% is the recommended number and this is sufficient to get a good error signal; this seems to be an arbitrary number but shows most values will work.

If using modulation of a voltage supplied to a current driver I strongly recommend using

a very low voltage (10 mV) and increasing from there until a suitable sideband height is reached, just in case high modulation damages the laser diode.

If using an EOM there is chance of damage as the crystals may be damaged by very high voltages applied to them. The home built EOMs have a resonant circuit with a variable capacitor to adjust the resonant frequency. By measuring the current drawn by the circuit the resonance of the EOM driving circuit can also be found without having to set up the low finesse cavity. The sideband height can be changed by adjusting the frequency of the driving voltage or the resonant frequency of the circuit. The modulation signal voltage should not be changed and should always have a magnitude of 1 dBm (0.71 V_{pp}).

C.3 Locking to a High Finesse Cavity

Once the correct modulation strength for the sidebands has been established the correct alignment of the optical cavity needs to be achieved. It is far easier to see a mode if we are able to scan the ECDL frequency. An arbitrary frequency generator can be used to give a triangular wave voltage to constantly change the piezo position and therefore the laser frequency. So that the offset could be easily changed with a switch to turn the scanning signal on and off a summing amplifier with a range of 0 V to +10 V was used. It was felt this voltage range was not wide enough, so an amplifier with a gain of 2.5 was built to be connected after the summing amplifier.

To get light into the high finesse optical cavity, align the input beam and reflected beam to be in-line with each other using a Post-it note with a 1 mm hole for the light to pass through. Once a rough alignment has been made, scan the laser frequency and use a photodiode to see the size of the transmitted signal. If no modes can be detected with the photodiode it may be required to use the eye to see modes and to increase the brightness by adjusting the coupling mirrors; often the eye can detect lower levels than a photodiode. The signal from

the oscillator for modulating the light and the signal from the photodiode for the reflected light are both fed into a mixer that outputs the error signal. To give a maximum error signal we must match the phase of the signals; we maximise the size of the error signal using the extra BNC cable as a phase shifter. The speed of transmission in a BNC is around 66% of the speed of light c. At our frequencies of around 10 MHz the wavelength will be

$$\lambda = \frac{c}{f} \tag{C.1}$$

$$= \frac{c \times 0.66}{10 \times 10^6 \text{ Hz}} \tag{C.2}$$

$$= 19.8\text{m} \tag{C.3}$$

To change the phase of the signal by 90 °C will require about 5 m of BNC cable. The cable length should be changed for one of the arms of the signal until the error signal is maximal. The direction of the slope is unimportant as this can be corrected for on the FALC by using the inverting or non-inverting output.

C.4 Using the Fast Analogue Linewidth Controller (FALC)

The FALC is extremely well designed for its use and has full specifications in the FALC user manual. To understand how the FALC works read pages 150 and 169 of the user manual before attempting to use the equipment. The FALC has three locking parameters for the fast branch: Slow Limited Integrator (SLI), Fast Limited Differentiator (FLD) and Fast Limited Integrator (FLI). The parameters are changed by adjusting the dip switches on the front of the FALC, table C.4 gives the typical parameters used.

To lock using the FALC the output from the mixer should be connected to either the inverting `invert` or non-inverting (noninv) input of the FALC. The choice is down to the locking direction required and is best chosen through trial and error. The FALC output is

	SLI	FLD	FLI
Straight Through Connection	10	1	1
Recommended Starting Parameters	5	1	6
Parameters used for H-Cav Locking	7	7	6
Parameters used for V-Cav Locking	9	5+6	7

Table C.1: Parameters for various stages of the locking to the high finesse optical cavity. The Straight Through Connection does not alter the signal at all and will output what is input to the system. The Recommended Starting Parameters are the numbers recommended by the FALC User Manual. The Parameters used for Locking are the parameters currently used on the experimental set ups.

connected to the current driver for the laser. Ensure the `run-reset` switch is set to reset. Looking at the monitor output from the FALC on a oscilloscope is very useful to see the filtered error signal. On the FALC there is a Trimpot for Offset Adjust; this can be used to move the offset of the error signal to have zero offset. This can also be achieved by slightly adjusting the orientation of the quarter-wave plate before the optical cavity but it is not clear why. The gain Trimpot was never used; it is recommended this is set to minimum when possible.

A simple BNC switch is very useful for locking and unlocking the laser, this is a switch between the FALC output and the input to the current driver. Using a small amplitude scan (500 mV_{pp}) the offset of the summing amplifier can be changed to find a mode. Images of what to expect during the various locking stages is shown in figure C.1. Once a mode can be seen on the scan the FALC can be turned on; the signs of the laser being roughly locked are that the transmission signal will increase and the error signal will become wider. A gain on the FALC of around 0.1 should be used, higher than this and the system will be unstable, but experiment to see what gives the biggest signal. If it does not work try using the other input to the FALC. The likely outcome will be a *grass* looking lock, i.e. lots of small spikes coming up. Once some sort of lock is found the parameters can be adjusted to reduce the grass and have the transmission signal be as high as possible; this can be done with the lock

running. There is no magic method for the parameter selection and trial and error is the best approach. It is expected that the locked signal transmission amplitude will be a lot higher than the peaks seen scanning the laser without any lock. Once the signal sticks at a level, the error signal can be looked at and the parameters changed further to reduce the width of the error signal. As is shown in figure C.1 the error signal in e) is much narrower than in d) where the locking parameters are not set correctly.

The scan can now be turned off on the summing amplifier and the laser will stay in lock. If need be the unlimited integrator can be used if the stability of the ECDL is not good. This is explained well on page 170 of the User Manual. The unlimited output from the FALC `Unlim` should be connected to one of the inputs on the summing amplifier.

C.5 Locking to the Correct Gaussian Mode

Now the laser locks it should lock to many different modes. Having the summing amplifier with the scan off and the FALC locking running it is possible to adjust the offset of the summing amplifier and the laser will lock to various TEM modes. The TEM_{00} mode can be found by adjusting the ECDL piezo; if one cannot be seen within the scanning range adjust the current supplied to the laser diode and try again. Once found it is possible to adjust the coupling mirrors while the laser is locked to increase the transmission signal size for the TEM_{00} mode, but most likely just touching the mirror actuator will kick it out of lock. It is also possible to increase the amount of light in the TEM_{00} mode by adjusting the focus of the fibre output collimator lens to change the mode matching into the optical cavity.

C.6 Making a Beat Signal

The beams from each laser system must be aligned onto the fast photodiode so the heterodyne beat measurement can be made; this alignment fortunately means both beams can be aligned

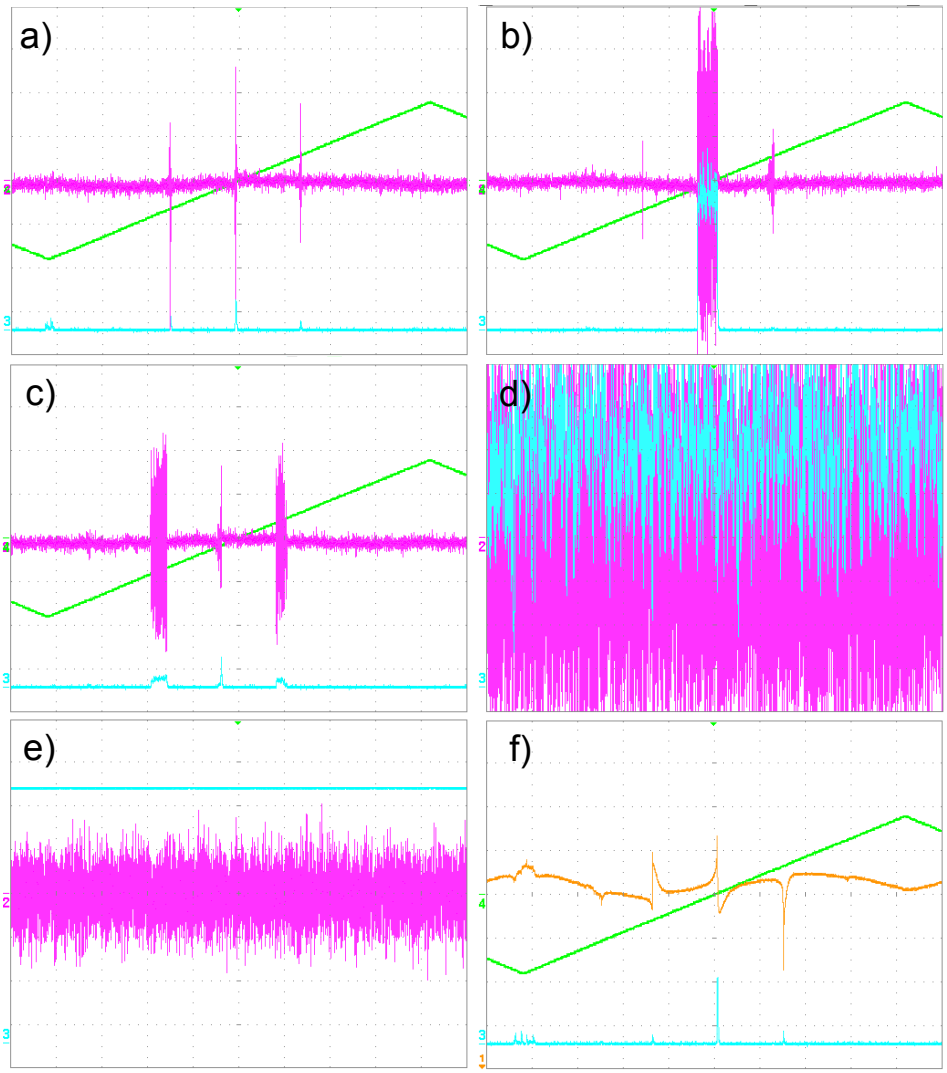

Figure C.1: The oscilloscope trace of the various stages of locking a laser to the high finesse optical cavity. **Blue** is the transmitted intensity through the optical cavity. **Purple** is the error signal from the mixer filtered with a 5 MHz low-pass filter. **Green** is the scanning voltage to the ECDL piezo. **a)** The laser locking is off, the transmitted signal can be seen through the cavity and the error is signal produced. The scanning voltage for the piezo was 600 mV at 3 Hz. **b)** The lock is turned on and the laser begins to lock to the central peak as desired. **c)** This is where the locking is set to the wrong direction and the locking is to the sidebands not the central peak; to resolve this swap to the other FALC input. **d)** The scan is turned off and the laser is locked but the parameters are not correct. **e)** The parameters are correctly set (note the scale for the transmitted signal has been increased $\times 2$ to still be visible). **f)** The lock is off and the scanning is on. **Orange** is the FALC monitor output with SLI=10, FLD=1 & FLI=8; now the slope direction of the error signal can be clearly seen.

at the same time to the fibre optic for the wavemeter. The wavemeter can be used to compare the laser frequencies directly. By covering one beam before the 50:50 beam splitter cube the frequency of the other beam can be measured. The current should be adjusted until one laser matches the other frequency to better than 7 GHz (0.007 THz on the wavemeter). Once this is achieved the signal can be viewed on the spectrum analysers; the high frequencies mean a suitable fast photodiode with a bias-tee must be used. The current drawn by the DC side of the bias-tee indicates the light power incident on the active area, this can be used to increase the power from both beams on the photodiode. Now the signal should be visible on the spectrum analyser; the mirrors can be adjusted further to increase the signal size.

To have two *locked* lasers beat together is much more difficult; by changing the offset to one of the lasers the TEM_{00} mode can be found but make sure the beat signal is still within the frequency range of the spectrum analyser. Once one is locked the other can be adjusted to find a TEM_{00} mode and now a measurement of the beat note for the locked laser can be achieved by reducing the frequency range around the peak. Modes in the optical cavities were not always seen every 1.5 GHz, as they should be given by the free spectral range of the cavities, but the reason a mode may not appear is unclear. It may be useful to use an auto-tracking function on the spectrum analyser that keeps the peak in the centre of the frequency scan. This will all take time so be patient as the mode being used will disappear for apparently no reason.

LIST OF REFERENCES

[1] S Bize, P Laurent, M Abgrall, H Marion, I Maksimovic, L Cacciapuoti, J Grunert, C Vian, FP dos Santos, P Rosenbusch, P Lemonde, G Santarelli, P Wolf, A Clairon, A Luiten, M Tobar, and C Salomon. Cold Atom Clocks and Applications. *Journal of Physics B*, 38(9):S449–S468, 2005. 2

[2] W Rindler. *Relativity: Special, General and Cosmological*. Oxford University Press, 2001. 2

[3] R Vessot, M Levine, E Mattison, E Blomberg, T Hoffman, G Nystrom, B Farrel, R Decher, P Eby, C Baugher, J Watts, D Teuber, and F Wills. Test of Relativistic Gravitation with a Space-Borne Hydrogen Maser. *Physical Review Letters*, 45:2081–2084, 1980. 2

[4] S Blatt, A D Ludlow, G K Campbell, J W Thomsen, T Zelevinsky, M M. Boyd, J Ye, X Baillard, M Fouché, R Le Targat, A Brusch, P Lemonde, M Takamoto, F-L Hong, H Katori, and V V Flambaum. New Limits on Coupling of Fundamental Constants to Gravity Using ^{87}Sr Optical Lattice Clocks. *Physical Review Letters*, 100(14):140801, 2008. 2

[5] S Reinhardt, G Saathoff, H Buhr, La Carlson, A Wolf, D Schwalm, S Karpuk, C Novotny, G Huber, M Zimmermann, R Holzwarth, T Udem, T Hänsch, and G Gwinner. Test of Relativistic Time Dilation with Fast Optical Atomic Clocks at Different Velocities. *Nature Physics*, 3(12):861–864, 2007. 2

[6] L Cacciapuoti and Ch Salomon. Space Clocks and Fundamental Tests: The ACES Experiment. *European Physical Journal-Special Topics*, 172:57–68, 2009. 2

[7] S Schiller, G Tino, P Gill, C Salomon, U Sterr, E Peik, A Nevsky, A Goerlitz, D Svehla, G Ferrari, N Poli, L Lusanna, H Klein, H Margolis, P Lemonde, P Laurent, G Santarelli, A Clairon, W Ertmer, E Rasel, J Mueller, L Iorio, C Laemmerzahl, H Dittus, E Gill,

M Rothacher, F Flechner, U Schreiber, V Flambaum, W Ni, L Liu, X Chen, J Chen, K Gao, L Cacciapuoti, R Holzwarth, M Hess, and W Schaefer. Einstein Gravity Explorer - A Medium-Class Fundamental Physics Mission. *Experimental Astronomy*, 23(2):573–610, 2009. 2

[8] P Duffett-Smith. *Practical Astronomy with your Calculator*. Cambridge University Press, 1988. 3

[9] O Kock. *Magneto-Optical Trapping of Strontium for use as a Mobile Frequency Reference*. PhD Thesis, University of Birmingham, 2013. 5

[10] R Eisberg and R Resnick. *Quantum Physics of Atoms, Molecules, Solids, Nuclei and Particles*. Wiley, 1985. 7, 8

[11] A Ludlow, M Boyd, T Zelevinsky, S Foreman, S Blatt, M Notcutt, T Ido, and J Ye. Systematic Study of the ^{87}Sr Clock Transition in an Optical Lattice. *Physical Review Letters*, 96(3):033003, 2006. 10

[12] P Gill. Optical frequency standards. *Metrologia*, 42:S125, 2005. 10

[13] S Falke, H Schnatz, J Vellore Winfred, T Middelmann, S Vogt, S Weyers, B Lipphardt, G Grosche, F Riehle, U Sterr, and Ch Lisdat. The ^{87}Sr Optical Frequency Standard at PTB. *Metrologia*, 48(5):399–407, 2011. 11

[14] S Diddams, D Jones, J Ye, S Cundiff, J Hall, J Ranka, R Windeler, R Holzwarth, T Udem, and T Hänsch. Direct Link between Microwave and Optical Frequencies with a 300 THz Femtosecond Laser Comb. *Physical Review Letters*, 84(22):5102–5105, 2000. 12

[15] Bureau International des Poids et Mesures. http://www.bipm.org/en/si/base_units/. 13

[16] C Audoin and B Guinot. *The Measurement of Time: Time, Frequency and the Atomic Clock*. Cambridge University Press, 2001. 13, 96

[17] National Institute of Standards and Technology. http://physics.nist.gov/cuu/Units/. 13

[18] BIPM Recommended values of standard frequencies. http://www.bipm.org/en/publications/mep.html. 13

[19] W Oskay, S Diddams, E Donley, T Fortier, T Heavner, L Hollberg, W Itano, S Jefferts, M Delaney, K Kim, F Levi, T Parker, and J Bergquist. Single-Atom Optical Clock with High Accuracy. *Physical Review Letters*, 97(2), 2006. 13

[20] K Hosaka, S Webster, A Stannard, B Walton, H Margolis, and P Gill. Frequency Measurement of the $^2S_{1/2} -^2 F_{7/2}$ Electric Octupole Transition in a Single $^{171}Yb^+$ Ion. *Physical Review A*, 79(3):033403, 2009. 13

[21] A Ludlow, T Zelevinsky, G Campbell, S Blatt, M Boyd, M de Miranda, M Martin, J Thomsen, S Foreman, J Ye, T Fortier, J Stalnaker, S Diddams, Y Le Coq, Z Barber, N Poli, N Lemke, K Beck, and C Oates. Sr Lattice Clock at 1×10^{-16} Fractional Uncertainty by Remote Optical Evaluation with a Ca Clock. *Science*, 319(5871):1805–1808, 2008. 13, 41

[22] T Rosenband, D Hume, P Schmidt, C Chou, A Brusch, L Lorini, W Oskay, R Drullinger, T Fortier, J Stalnaker, S Diddams, W Swann, N Newbury, W Itano, D Wineland, and J Bergquist. Frequency Ratio of Al^+ and Hg^+ Single-ion Optical Clocks; Metrology at the 17th Decimal Place. *Science*, 319:1808–1812, 2008. 13

[23] C Chou, D Hume, J Koelemeij, D Wineland, and T Rosenband. Frequency Comparison of Two High-Accuracy Al^+ Optical Clocks. *Physical Review Letters*, 104:070802, 2010. 14

[24] N Hinkley, J Sherman, N Phillips, M Schioppo, N Lemke, K Beloy, M Pizzocaro, C Oates, and A Ludlow. An Atomic Clock with 10^{-18} Instability. *arXiv:1305.5869*, 2013. 14

[25] G Campbell, A Ludlow, S Blatt, J Thomsen, M Martin, M de Miranda, T Zelevinsky, M Boyd, J Ye, S Diddams, T Heavner, T Parker, and S Jefferts. The Absolute Frequency of the ^{87}Sr Optical Clock Transition. *Metrologia*, 45(5):539–548, 2008. 14

[26] T Nicholson, M Martin, J Williams, B Bloom, M Bishof, M Swallows, S Campbell, and J Ye. Comparison of Two Independent Sr Optical Clocks with 1×10^{-17} Stability at 10^3 s. *Physical Review Letters*, 109:230801, 2012. 14

[27] E Raab, M Prentiss, A Cable, S Chu, and D Pritchard. Trapping of Neutral Sodium Atoms with Radiation Pressure. *Physical Review Letters*, 59(23):2631–2634, Dec 1987. 14

[28] Nobel Prize Foundation. http://nobelprize.org/nobel_prizes/physics/laureates/1997/. 14

[29] H Metcalf and P van der Straten. *Laser Cooling and Trapping*. Springer, 1999. 15, 21

[30] T Legero, S Winfred, F Riehle, and U Sterr. Ultracold Sr-88 Atoms for an Optical Lattice Clock. In *Proceedings of the 2007 IEEE International Frequency Control Symposium-Jointly with the 21st European Frequency and Time Forum*, pages 119–122. IEEE, 2007. 15

[31] M Takamoto, F Hong, R Higashi, and H Katori. An Optical Lattice Clock. *Nature*, 435(7040):321–324, 2005. 17

[32] C Ye. *Tunable External Cavity Diode Lasers*. World Scientific, 2004. 19, 21

[33] X Baillard, A Gauguet, S Bize, P Lemonde, Ph Laurent, A Clairon, and P Rosenbusch. Interference-Filter-Stabilized External-Cavity Diode Lasers. *Optics Communications*, 266(2):609–613, 2006. 20

[34] S Saliba and R Scholten. Linewidths Below 100 kHz with External Cavity Diode Lasers. *Applied Optics*, 48(36):6961–6966, 2009. 22

[35] CH Henry. Theory of the Linewidth of Semiconductor-Lasers. *IEEE Journal of Quantum Electronics*, 18(2):259–264, 1982. 22

[36] MW Fleming and A Mooradian. Spectral Characteristics of External-Cavity Controlled Semiconductor-Lasers. *IEEE Journal of Quantum Electronics*, 17(1):44–59, 1981. 22

[37] L Ricci, M Weidemuller, T Esslinger, A Hemmerich, C Zimmermann, V Vuletic, W Konig, and TW Hänsch. A Compact Grating-Stabilized Diode-Laser System for Atomic Physics. *Optics Communications*, 117(5-6):541–549, 1995. 23

[38] S Papp. *PhD Thesis, Experiments with a Two-Species Bose-Einstein Condensate Utilizing Widely Tunable Interparticle Interactions*. University of Colorado, 2001. 23

[39] Radiant Dyes Lasers. http://www.radiant-dyes.com/. 23

[40] Sacher Lasertechnik. http://www.sacher-laser.com/. 24

[41] A low voltage piezo stack PCh 50/5x5/2 produced by Piezomechanik Gmbh. http://www.piezomechanik.com/en. 24

[42] Wavelength Electronics. http://www.teamwavelength.com/. 25

[43] Los Angeles Minus-k Technologies Inc. http://www.minusk.com/. 27

[44] L Smith. Lasers for Atomic Clocks. *4th Year Masters Project, University of Birmingham*, 2012. 27

[45] E Cook, P Martin, T Brown-Heft, J Garman, and D Steck. High Passive-Stability Diode-Laser Design for use in Atomic-Physics Experiments. *Review of Scientific Instruments*, 83(4):043101, 2012. 28

[46] C Fabry and A Perot. On a New Form of Interferometer. *Astrophysical Journal*, 13(4):265–272, 1901. 30

[47] R Drever, J Hall, F Kowalski, J Hough, G Ford, A Munley, and H Ward. Laser Phase and Frequency Stabilization Using an Optical Resonator. *Applied Physics B*, 31(2):97–105, 1983. 33

[48] E Black. An Introduction to Pound-Drever-Hall Laser Frequency Stabilization. *American Journal of Physics*, 69(1):79–87, 2001. 33, 37

[49] G Bjorklund, M Levenson, W Lenth, and C Ortiz. Frequency-Modulation (FM) Spectroscopy - Theory of Lineshapes and Signal-to-Noise Analysis. *Applied Physics B*, 32(3):145–152, 1983. 33

[50] J Alnis, A Matveev, N Kolachevsky, Th Udem, and T Hänsch. Subhertz Linewidth Diode Lasers by Stabilization to Vibrationally and Thermally Compensated Ultralow-expansion Glass Fabry-Perot Cavities. *Physical Review A*, 77(5):053809, 2008. 38, 39, 42, 59, 60, 74, 82, 86, 100, 101, 105

[51] Corning Speciality Glass and Ceramic. http://www.corning.com/specialtymaterials/products_capabilities/ULE.aspx. 38

[52] D Hils, J Faller, and J Hall. Practical Sound Reducing Enclosure for Laboratory Use. *Review of Scientific Instruments*, 57(10):2532–2534, 1986. 39

[53] T Nazarova, F Riehle, and U Sterr. Vibration-Insensitive Reference Cavity for an Ultra-Narrow-Linewidth Laser. *Applied Physics B*, 83(4):531–536, 2006. 39

[54] E Flanagan and S Hughes. Measuring Gravitational Waves from Binary Black Hole Coalescences. I. Signal to Noise for Inspiral, Merger, and Ringdown. *Physical Review D*, 57(8):4535–4565, 1998. 40

[55] S Seel, R Storz, G Ruoso, J Mlynek, and S Schiller. Cryogenic Optical Resonators: A New Tool for Laser Frequency Stabilization at the 1 Hz Level. *Physical Review Letters*, 78(25):4741–4744, 1997. 40

[56] B Young, F Cruz, W Itano, and J Bergquist. Visible lasers with subhertz linewidths. *Physical Review Letters*, 82(19):3799–3802, 1999. 41

[57] A Ludlow, X Huang, M Notcutt, T Zanon-Willette, S Foreman, M Boyd, S Blatt, and J Ye. Compact, Thermal-Noise-Limited Optical Cavity for Diode Laser Stabilization at 1×10^{-15}. *Optics Letters*, 32(6):641–643, 2007. 41

[58] M Notcutt, L Ma, A Ludlow, S Foreman, J Ye, and J Hall. Contribution of Thermal Noise to Frequency Stability of Rigid Optical Cavity via Hertz-Linewidth Lasers. *Physical Review A*, 73(3):031804, 2006. 41

[59] L Chen, J Hall, J Ye, T Yang, E Zang, and T Li. Vibration-Induced Elastic Deformation of Fabry-Perot Cavities. *Physical Review A*, 74(5):053801, 2006. 41, 53, 56

[60] S Webster, M Oxborrow, and P Gill. Vibration Insensitive Optical Cavity. *Physical Review A*, 75(1):011801, 2007. 41, 53, 55, 59

[61] S Webster, M Oxborrow, S Pugla, J Millo, and P Gill. Thermal-Noise-Limited Optical Cavity. *Physical Review A*, 77(3):033847, 2008. 41

[62] Y Zhao, J Zhang, A Stejskal, T Liu, V Elman, Z Lu, and L Wang. A Vibration-Insensitive Optical Cavity and Absolute Determination of its Ultra High Stability. *Optics Express*, 17(11):8970–8982, 2009. 42

[63] K Numata, A Kemery, and J Camp. Thermal-Noise Limit in the Frequency Stabilization of Lasers with Rigid Cavities. *Physical Review Letters*, 93(25):250602, 2004. 42, 43, 44, 45, 51

[64] Stable Laser Systems. http://www.stablelasers.com/. 42

[65] T Kessler, C Hagemann, C Grebing, T Legero, U Sterr, and F Riehle. A Sub-40 mHz Linewidth Laser Based on a Silicon Single-Crystal Optical Cavity. *Nature Photonics*, 6:687, 2012. 42, 51, 101

[66] A Abramovici, WE Althouse, RWP Drever, Y Gurse, S Kawamura, FJ Raab, D Shoemaker, L Sievers, RE Spero, KS Thorne, RE Vogt, R Weiss, SE Whitcomb, and ME Zucker. LIGO - The Laser-Interferometer-Gravitational-Wave-Observarory. *Science*, 256(5055):325–333, 1992. 43

[67] S Rowan, J Hough, and D Crooks. Thermal Noise and Material Issues for Gravitational Wave Detectors. *Physics Letters A*, 347(1-3):25–32, 2005. 43, 46

[68] Y Levin. Internal Thermal Noise in the LIGO Test Masses: A Direct Approach. *Physical Review D*, 57(2):659–663, 1998. 43

[69] T Kessler, T Legero, and U Sterr. Thermal Noise in Optical Cavities Revisited. *Journal of the Optical Society of America B*, 29(1):178–184, 2012. 43

[70] F Bondu, P Hello, and JY Vinet. Thermal Noise in Mirrors of Interferometric Gravitational Wave Antennas. *Physics Letters A*, 246(3-4):227–236, SEP 14 1998. 43

[71] G Harry, A Gretarsson, P Saulson, S Kittelberger, S Penn, W Startin, S Rowan, M Fejer, D Crooks, G Cagnoli, J Hough, and N Nakagawa. Thermal Noise in Interferometric Gravitational Wave Detectors due to Dielectric Optical Coatings. *Classical and Quantum Gravity*, 19(5):897–917, 2002. 44

[72] N Nakagawa, A Gretarsson, E Gustafson, and M Fejer. Thermal Noise in Half-Infinite Mirrors with Nonuniform Loss: A Slab of Excess Loss in a Half-Infinite Mirror. *Physical Review D*, 65(10):102001, 2002. 44

[73] M Fejer, S Rowan, G Cagnoli, D Crooks, A Gretarsson, G Harry, J Hough, S Penn, P Sneddon, and S Vyatchanin. Thermoelastic Dissipation in Inhomogeneous Media: Loss Measurements and Displacement Noise in Coated Test Masses for Interferometric Gravitational Wave Detectors. *Physical Review D*, 70(8):082003, 2004. 44

[74] M Cerdonio, L Conti, A Heidmann, and M Pinard. Thermoelastic Effects at Low Temperatures and Quantum Limits in Displacement Measurements. *Physical Review D*, 63(8):082003, 2001. 45

[75] VB Braginsky, ML Gorodetsky, and SP Vyatchanin. Thermodynamical Fluctuations and Photo-thermal Shot Noise in Gravitational Wave Antennae. *Physics Letters A*, 264(1):1–10, 1999. 45

[76] ANSYS Inc. Simulation Driven Product Development. http://www.ansys.com/. 45

[77] T Legero, T Kessler, and U Sterr. Tuning the Thermal Expansion Properties of Optical Reference Cavities with Fused Silica Mirrors. *Journal of the Optical Society of America B*, 27(5):914–919, 2010. 52

[78] C Bond, P Fulda, L Carbone, K Kokeyama, and A Freise. Higher Order Laguerre-Gauss Mode Degeneracy in Realistic, High Finesse Cavities. *Physical Review D*, 84:102002, 2011. 52

[79] F Brueckner, T Clausnitzer, O Burmeister, D Friedrich, E Kley, K Danzmann, A Tuennermann, and R Schnabel. 100% Reflectivity from a Monolithic Dielectric, Micro-Structured Surface. *Proceedings of the Society of Photo-Optical Instrumentation Engineers (SPIE)*, 6883:X8830, 2008. 52

[80] COMSOL Multiphysics Engineering Simulation Software Environment. http://www.comsol.com/. 53

[81] D Guyomarc'h, G Hagel, C Zumsteg, and M Knoop. Some Aspects of Simulation and Realization of an Optical Reference Cavity. *Physical Review A*, 80(6):063802, 2009. 57

[82] S Webster and P Gill. Force-Insensitive Optical Cavity. *Optics Letters*, 36(18):3572–3574, 2011. 57

[83] Advanced Thin Films. http://www.atfilms.com/. 59

[84] H Kogelnik and T Li. Laser Beams and Resonators. *Applied Optics*, 5(10):1550, 1966. 60, 72, 73

[85] VACOM Vacuum Components and Mechanics GmbH. http://www.vacom-uk.com. 61

[86] S Cox, P Griffin, C Adams, D DeMille, and E Riis. Reusable Ultrahigh Vacuum Viewport Bakeable to 240 °C. *Review of Scientific Instruments*, 74(6):3185–3187, 2003. 61

[87] NASA Outgassing Data for Selecting Spacecraft Materials. http://outgassing.nasa.gov/. 63

[88] SMAHT Ceramics Inc. http://www.smahtceramics.com/. 64

[89] S Self. Focusing of Spherical Gaussian Beams. *Applied Optics*, 22(5):658–661, 1983. 74

[90] L Ma, P Jungner, J Ye, and J Hall. Delivering the Same Optical Frequency at 2 Places - Accurate Cancellation of Phase Noise Introduced by an Optical-fibre or Other Time-varying Path. *Optics Letters*, 19(21):1777–1779, 1994. 75

[91] Toptica Photonics. http://www.toptica.com/. 77

[92] Jenoptik AG. http://www.jenoptik.com/. 78, 92

[93] A Yariv. *Optical Electronics*. Saunders College Publishing, Fourth edition, 1991. 86

[94] Data Sheet for Lithium Niobate from Boston Piezo Optics. http://www.bostonpiezooptics.com/. 89

[95] CA Greenhall. Frequency Stability Review. *Telecommunications and Data Acquisition Progress*, 42:200, 1986. 96

[96] M Hopcroft. Matlab program for calculation and plotting of an allan deviation. http://www.mathworks.com/matlabcentral/fileexchange/13246-allan, October 2010. 97

[97] EIP 575 Source Locking Frequency Counter data sheet. http://wa6vhs.com/Test20equipment/EIP/578.pdf. 97

[98] J Barnes and D Allan. Variances Based on Data with Dead Time Between Measurements. *NIST Technical Note*, page 1318, 1990. 100