Sub-hertz optical frequency metrology using a single ion of $^{171}Yb^{+}$

Steven A. King

Linacre College
University of Oxford

A thesis submitted in partial fulfilment of
the requirements for the degree of
Doctor of Philosophy at the University of Oxford

Trinity term, 2012

Experimental work carried out at the
National Physical Laboratory
Teddington
Middlesex

Sub-hertz optical frequency metrology using a single ion of $^{171}\mathrm{Yb}^{+}$

A thesis submitted for the degree of Doctor of Philosophy

Trinity Term 2012

Steven A. King

Linacre College, Oxford

Abstract

Optical frequency standards offer the possibility of a step improvement of up to two orders of magnitude in the accuracy with which the SI second can be realised. $^{171}\mathrm{Yb}^{+}$ possesses two dipole-forbidden optical transitions that are promising candidates for a redefinition of the second. In this thesis, absolute frequency measurements of these two transitions are presented.

A number of experimental upgrades have been implemented, which have resulted in a large reduction in both the statistical and systematic uncertainties associated with the measurements and have improved both the reliability and simplicity of the experimental setup. In particular, the replacement of two frequency-doubled Ar^{+}-pumped Ti:sapphire lasers with extended cavity diode lasers has eliminated the downtime associated with their maintenance. Additionally, the introduction of polarisation modulation on the cooling light has allowed the residual bias magnetic field required for laser cooling to be reduced by a factor of thirty.

The first measurement at the National Physical Laboratory (NPL) of the frequency of the $^{2}\mathrm{S}_{1/2}$ $(F = 0) \rightarrow$ $^{2}\mathrm{D}_{3/2}$ $(F' = 2)$ electric quadrupole (E2) transition at 436 nm is presented. The transition frequency was measured against a hydrogen maser using a femtosecond optical frequency comb, and was determined with a relative standard uncertainty of 1.3×10^{-14}.

A commercial diode-based laser system was then implemented in order to drive the $^{2}\mathrm{S}_{1/2}$ $(F = 0) \rightarrow$ $^{2}\mathrm{F}_{7/2}$ $(F' = 3)$ electric octupole (E3) transition at 467 nm. The laser frequency was actively stabilised to the ultra-narrow atomic absorption with a resolved linewidth of 11 Hz, allowing the acquisition of ninety hours of frequency data measured relative to the NPL's primary frequency standard CsF2. Combined with a thorough evaluation of the systematic perturbations, the total fractional uncertainty in the absolute frequency of the transition has been reduced by a factor of twenty to 1×10^{-15}. Recent complementary results from Physikalisch-Technische Bundesanstalt (PTB) show that the E3 transition in $^{171}\mathrm{Yb}^{+}$ has the potential to be a highly accurate and reproducible optical frequency standard, and to date these measurements demonstrate the best international agreement between trapped ion optical frequency standards.

Acknowledgements

The work presented in this thesis would not have been possible without the considerable input of a large number of scientists. Huge thanks are due to the people with whom I have worked most closely during the last three years: Stephen Webster, Rachel Godun, Patrick Baird and Patrick Gill, not just for imparting their tremendous experience and knowledge from day one but also for making my D.Phil an incredibly enjoyable experience. They have always made time to answer my questions, and for that I am extremely grateful.

The other ion trappers at NPL have brought a lot to this project via their considerable experience and expertise: Hugh Klein, Geoffrey Barwood and Guilong Huang who have always been available to discuss science over a cup of coffee; and Alastair Sinclair and Guido Wilpers who provided valuable advice on technical issues that arose along the way.

Thanks also to the occupants of the frequency comb lab: Giuseppe Marra for many, many long conversations about rf electronics; Helen Margolis for putting up with my excess enthusiam; Luke Johnson for analysing the vast quantities of data we sent to him down the optical fibre; Barney Walton and Veronika Tsatourian for the late-night frequency measurements; and Stephen Lea for his contributions to the fundamental physics side of the experiment.

I was fortunate to start my doctorate at NPL at the same time as three other graduate students: Ian Hill, Liz Bridge and Ben Parker. They have been great people to work with, and I have had many good times with them outside the lab.

Special thanks are saved for my family, who have supported me throughout my now rather lengthy progression through the education system and have not chased me to go out and get a proper job! I hope that this thesis proves that it has been time well spent.

Finally, thanks to NPL for providing the funding to allow me to work on this exciting project. I have learned a great many things and had a lot of fun doing so.

Contents

Chapter 1

Introduction

The International System of Units (SI) consists of seven base units: the metre, kilogram, second, ampere, kelvin, candela and mole [1]. Each of these units is rigorously defined, and is then realised experimentally by measurement institutes worldwide. The measurement institutes are then responsible for the maintenance and dissemination of the standards. Of the seven base units, it is the second that has been realised most accurately. This is largely due to the advent of microwave atomic clocks in the 1950s.

Proposed changes to the SI will redefine all but one of the base units in terms of fixed values of the fundamental constants, which are then experimentally realised in the laboratory [2]. The definitions of these quantities (except the mole, which is simply a number) will all make use of the definition of the second.

It has not taken long for highly precise and accurate clocks to find applications that extend beyond the laboratory environment, ranging from high-speed communications to global navigation. Perhaps the most well-known example is the global positioning system (GPS). To determine its location, a receiver relies on signals from at least four space-based clocks with timing precision that would have been impossible to realise without the development of atomic clocks.

In addition to their increasing role in everyday life, precise clocks are gaining importance in tests of fundamental physics. Experiments are ongoing aimed at detecting any time-variation of the fundamental constants; this could have profound implications for modern physical theories that attempt to unify gravitation and quantum mechanics. Future space-based projects such as the Atomic Clock Ensemble in Space (ACES) plan to test theories of gravitation by observing the behaviour of clocks in low gravity environments, and ground-based geodesy experiments will benefit from such precise timekeeping.

Despite the exceptional performance of microwave atomic frequency standards, work is ongoing to improve the accuracy of the realisation of the SI second even further. The development and continued improvement of optical frequency standards has led to the proposition of a re-definition of the second in terms of an optical atomic transition rather than the present-day definition involving the caesium-133 ground state hyperfine transition (see subsection 1.1.2). In preparation, the International Committee for Weights and Measures (CIPM) has proposed some of the leading optical frequency standards as secondary representations of the second [3, 4].

This chapter introduces the main components of any clock, and how they are characterised. It also introduces the concept of the atomic clock, and why research is ongoing into the development of clocks that operate in the optical and ultraviolet regions.

1.1 Clock fundamentals

Every clock consists of three major components: an oscillator that provides the 'ticks', a reference to define the duration of one tick, and a counter which reads out the ticks. This is true whether the clock is based on the oscillation of a pendulum, piezoelectric vibrations of a quartz crystal or the oscillations of a microwave signal.

1.1.1 The oscillator

The choice of oscillator and its physical characteristics are important considerations. Taking the pendulum as an example, the stability of its oscillation period can be improved by aerodynamic design and by reducing losses at the pivot point. But even if the pendulum is very stable in the short term it is likely that, due to manufacturing imprecision or temperature drifts causing its length to change, a 'seconds' pendulum (which is designed to swing once every two seconds) may not swing exactly 43200 times in one twenty-four hour day. A typical pendulum clock might gain or lose a few seconds in a week. However, by comparing the time as measured by a clock against a more stable reference (for example, the rotation of the earth) it can be adjusted to bring it back into line. An appropriate choice of reference can compensate for even a very noisy oscillator if frequent comparisons are made.

1.1.2 The reference

Early attempts to quantify the passage of time involved astronomical observations. In particular, the second was defined in terms of the period of rotation of the Earth about its own axis. As timekeeping methods improved, it became clear that the mean solar day was lengthening, but at a variable rate. In 1960, the second was re-defined in terms of the Earth's orbital period about the Sun in the year 1900 [1]. This was still an unsatisfactory definition, but was designed to bring the definition of the second into line with the ephemeris timescale that had been adopted by the International Astronomical Union eight years previously.

The 1950s and 60s heralded the arrival of the atomic clock. Work by Louis Essen and John Parry at the National Physical Laboratory (NPL) using two separated microwave cavities to interrogate beams of caesium atoms [5] demonstrated that atomic transition frequencies could be measured extremely precisely, and perhaps more importantly that when the systematic conditions under which they are measured are well controlled then it is possible to extract *unperturbed* transition frequencies. This is the key to their reproducibility: an atomic clock built anywhere in the universe should tick at the same rate when corrected to its unperturbed frequency. As frequency is simply the reciprocal of time, a clock had been developed that ticked at a rate of over nine billion times every second. Since 1967, the SI second [6] has been 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.’

The limiting systematic uncertainties of clocks based on atomic beams are related to the velocities of the thermal atoms in the beam. Slowing the atomic beam allows a reduction in these uncertainties, but also limits the maximum possible physical separation between the microwave cavities and hence the maximum time between the interrogations* as the atoms simply fall out of the apparatus under gravity. The apparatus was refined to allow a cloud of cold atoms to be launched vertically, producing the so-called caesium fountain. This has two advantages: a reduction in motional systematic effects, and since the atoms experience the same microwave cavity twice (once on the way up and again on the way down) the potential phase offset between the two microwave cavities becomes smaller (as long as the atoms pass through exactly the same point in the cavity twice. If the fountain is tilted there is a potential systematic offset known as distributed cavity phase shift, which is related to any microwave phase gradients across the cavity). In 1997, the definition of the second was clarified to state that the transition frequency was to be corrected to an environmental temperature of 0 K, as shifts due to ambient blackbody radiation were now becoming significant.

As the fountains continue to improve, it is becoming necessary to evaluate an increasing number of systematic shifts. Recoil from the absorption of microwaves will cause the two wavefunctions corresponding to the two hyperfine levels to become slightly spatially separated (this effect is used in atom interferometry). As the two sub-clouds pass through apertures in the microwave cavity, it is possible that they could be ‘clipped’ in a non-common way. As the operation of the fountain relies on the excitation of a precise fraction of the atoms, a loss that is not common to the two states could introduce a bias

*As will be elaborated upon in chapter 5, increasing the interrogation time allows narrower absorption features to be resolved and thereby reduces the level of statistical uncertainty on the transition frequency.

on the measured transition frequency. After a thorough evaluation of this effect and that of the distributed cavity phase shift, recent measurements at the NPL indicate that the fountain CsF2 has reached a total fractional systematic uncertainty of 2.3×10^{-16} [7, 8].

1.1.3 The counter

The final component of any clock is a counter, which reads out the ticks. In a mechanical watch or pendulum clock, the counting is performed by clockwork mechanisms which then relay the information to a dial or screen. In a microwave frequency standard, a sample of the oscillating signal is sent to a microwave counter, which can then store this data electronically. This approach can not be directly followed for an optical frequency standard, as the oscillations of an optical wave are much too rapid to be counted electronically. Instead, an optical frequency comb (see section 8.1) is used to down-convert the signal to radio (rf) frequencies, which can then easily be counted using conventional rf electronics.

1.2 The Allan variance

The performance of an atomic clock is quantified by its stability and accuracy. When characterising frequency fluctuations, the simplistic approach would be to quote a standard deviation about some mean value. This approach is not suitable for oscillators, due to the presence of correlated noise. Instead, it is usual to quote the Allan variance [9], which corresponds to the level of frequency fluctuations when the dataset is averaged into bins of period τ:

$$\sigma^2(\tau) = \frac{1}{2} \left\langle (\bar{y}_{n+1} - \bar{y}_n)^2 \right\rangle \tag{1.1}$$

where $\langle\rangle$ denotes an average over all bins. When the Allan variance is plotted as a function of averaging time on a log-log plot, the gradient of the curve can be used to extract information about the particular noise process on the relevant timescale (see figure 1.1). A summary of the main noise processes is given in table 1.1.

Noise Process	Frequency dependence	Allan variance
White phase	f^2	$1/\tau^2$
Flicker phase	f	$1/\tau^2$
White frequency	f^0	$1/\tau$
Flicker frequency	$1/f$	τ^0
Random walk	$1/f^2$	τ
Drift	N/A	τ^2

Table 1.1: Frequency dependence of the common noise processes at work in oscillators, along with the characteristic functional dependence of the Allan variance on averaging time.

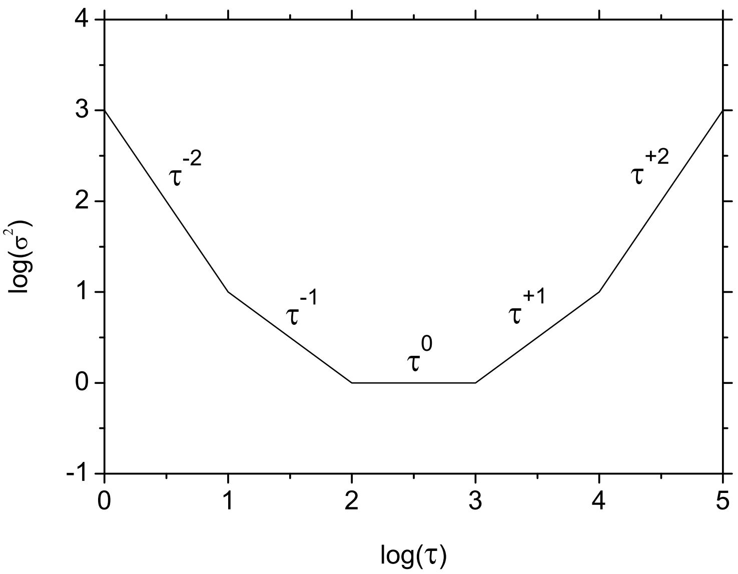

Figure 1.1: Log-log plot of the Allan variance σ^2 against averaging time τ. The gradient of the plot provides information on the primary noise contribution on the timescale considered.

In the present work, it is the fractional Allan deviation $\sigma(\tau)/\nu = \sqrt{\sigma^2(\tau)}/\nu$ that is quoted to allow the stabilities to be quantified in terms of 'parts in 10^{x}' at a specified averaging time.

1.3 Oscillator statistics

The statistical fluctuations of an oscillator can be described by the simple expression:

$$\sigma(\tau) \propto \left(\frac{1}{\mathrm{SNR}}\right)\left(\frac{1}{Q}\right)\sqrt{\frac{t_P + t_D}{\tau}} \tag{1.2}$$

where Q is the quality factor of the oscillator resonance defined in equation 1.4, SNR is the signal-to-noise ratio that can be extracted from the system, t_P is the time spent probing the system, t_D is the dead time in the probe cycle and τ is the averaging period. Each of these parameters must be optimised in order to realise the true potential of an oscillator.

1.3.1 Quantum projection noise

The signal-to-noise from an atomic clock is determined by the number of atoms n, and the probability p with which the resonance is driven. In the simplest case of a single atom, a pulse of near-resonant radiation will excite the atom into a coherent superposition of the initial and final states. When the state is read out, the wavefunction collapses into one of the two states. As the readout of the excitation rate of any individual atom is binomial, the statistical noise on the excitation is given by the standard deviation of the binomial distribution and referred to as quantum projection noise (QPN) [10]. This forms the ultimate limit to the statistical uncertainty of any atomic standard. The signal-to-noise is given by:

$$\mathrm{SNR} = \frac{np}{\sqrt{np(1-p)}} \tag{1.3}$$

where n is the number of atoms in the sample (or the total number of times the transition is probed) and p is the probability of driving the transition at the relevant detuning.

1.3.2 Transition quality factor

The quality factor of a resonance is given by the ratio of the frequency ν at which it occurs to the full-width at half maximum (FWHM) $\Delta\nu$ with which the resonance is resolved:

$$Q = \frac{\nu}{\Delta\nu} \tag{1.4}$$

In the absence of any other line-broadening effects, the resolved FWHM is roughly given by the inverse of the time spent probing the system (this will be discussed further in chapter 5). Thus the quality factor can be improved for a given probe time by choosing a transition with a high resonance frequency. Optical frequencies are around 10^5 times higher than microwave frequencies, and this is one of the main advantages of optical frequency standards over their microwave counterparts.

1.3.3 Dead time

Dead time in the clock cycle degrades the stability at a given averaging time by slowing down the rate at which the oscillator can be corrected. In addition, if the dead time becomes a significant fraction of the cycle time then the system is liable to the Dick effect, where high frequency noise from the oscillator can be aliased due to the low sample rate (see section 1.4.2).

1.4 Optical frequency standards

Typical optical transitions have natural linewidths on the order of megahertz, with resulting quality factors of around 10^9. However, some atomic species have low-lying metastable states. For instance, the lowest-lying excited state in Yb^+ is ${}^2F_{7/2}$. Decay out of this state

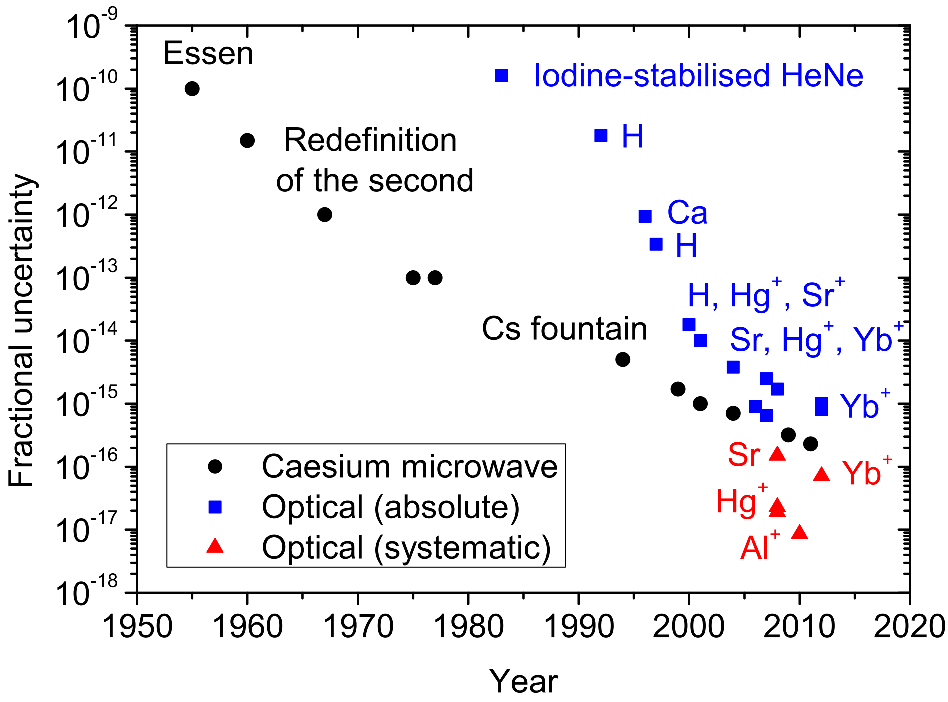

Figure 1.2: Improvement of frequency standards over time with some key results highlighted [12]. It can be seen that some optical frequency standards have now reported systematic uncertainties below that of the primary standard.

to the $^2S_{1/2}$ ground state is forbidden by dipole selection rules as it requires a change of three units of angular momentum. Hence the S→F transition has a very narrow natural linewidth. Alternatively, in singly-ionised group III metals such as Al^+ the lowest lying excited state is $^3P_0$ which is totally forbidden from decaying to the $^1S_0$ ground state. The transition can only be induced by the presence of mixing between states, either by the hyperfine interaction or the presence of a magnetic field.

Candidates for optical frequency standards are being developed worldwide [11] and are continually improving, with some now reporting systematic uncertainties below that of the caesium primary standard (see figure 1.2). These systems fall into two main categories: those based on trapped ions and those based on neutral atoms confined in optical lattice potentials.

Species	Transition	Systematic uncertainty
^{115}In$^+$	^{1}S$_0 \rightarrow {}^3$P$_0$	1.4×10^{-14} [24]
^{88}Sr$^+$	^{2}S$_{1/2} \rightarrow {}^2$D$_{5/2}$	2.5×10^{-15} [19]
^{40}Ca$^+$	^{2}S$_{1/2} \rightarrow {}^2$D$_{5/2}$	5.0×10^{-16} [16]
^{171}Yb$^+$	^{2}S$_{1/2} \rightarrow {}^2$D$_{3/2}$	4.5×10^{-16} [28]
^{171}Yb$^+$	^{2}S$_{1/2} \rightarrow {}^2$F$_{7/2}$	7.0×10^{-17} [17]
^{199}Hg$^+$	^{2}S$_{1/2} \rightarrow {}^2$D$_{5/2}$	1.9×10^{-17} [29]
^{27}Al$^+$	^{1}S$_0 \rightarrow {}^3$P$_0$	8.6×10^{-18} [27]
^{229}Th$^+$	nuclear	N/A

Table 1.2: Comparison of the lowest published systematic uncertainties for the current state-of-the-art ion-based optical frequency standards, correct as of June 2012.

1.4.1 Ions

Systems based on single ions in radio frequency (rf) potentials have the advantage of tight confinement. This results in low systematic uncertainties associated with motion, and narrow linewidths from extended interrogation times. However they suffer from a low signal-to-noise ratio which results in high levels of quantum projection noise, and if the ions are not confined to the centre of the trapping potential then they become liable to a Stark shift from the residual trapping field.

Systems based on Ca$^+$ [13, 14, 15, 16], Yb$^+$ [17, 18], Sr$^+$ [19, 20, 21], Hg$^+$ [22], In$^+$ [23, 24], Th$^+$ [25, 26] and Al$^+$ [27] are being developed, with systematic uncertainties of below 1×10^{-17} being reported in the case of Al$^+$. The leading systems are summarised in table 1.2.

1.4.2 Neutral atoms

Neutral atom systems have the benefit of a much lower statistical uncertainty at a given averaging time as more than 10^5 atoms can be interrogated simultaneously, resulting in a much higher signal-to-noise. However, confinement to allow long interrogation times requires the atoms to be loaded into an optical lattice, greatly reducing the duty cycle of the standard as the atoms must first be cooled in two stages. This makes them liable to the Dick effect: aliasing of laser noise can occur due to the low sampling rate, thus

Species	Transition	Systematic uncertainty
^{40}Ca	$^1S_0 \rightarrow {}^3P_1$	6.6×10^{-15} [33]
^{88}Sr	$^1S_0 \rightarrow {}^3P_0$	2.9×10^{-15} [44]
^{174}Yb	$^1S_0 \rightarrow {}^3P_0$	1.5×10^{-15} [45]
^{171}Yb	$^1S_0 \rightarrow {}^3P_0$	3.4×10^{-16} [35]
^{87}Sr	$^1S_0 \rightarrow {}^3P_0$	1.5×10^{-16} [36]
Mg	$^1S_0 \rightarrow {}^3P_0$	N/A
Hg	$^1S_0 \rightarrow {}^3P_0$	N/A

Table 1.3: Comparison between the current state-of-the-art neutral atom-based optical frequency standards, correct as of June 2012.

degrading the short-term instability of the standard. This effect is not applicable in ion-based standards as the instability due to projection noise generally exceeds the short-term oscillator instability. In order to reduce this effect, and to realise the true QPN limit of an optical lattice clock, a reduction in the cycle dead time is required in conjunction with a reduction in the laser noise.

Care must be taken with the choice of wavelength for the optical lattice as this can introduce a large perturbation onto the clock transition through an ac Stark shift [30, 31]. There is, however, the possibility of a 'magic' wavelength at which the upper and lower states are equally perturbed and hence there is no net effect on the clock transition*.

Unlike a single-ion system, more than one atom can occupy a site in the lattice and therefore density-dependent systematic shifts due to collisions between the atoms must be considered [32]. Some groups have even reported collisional perturbations with fermionic isotopes, where cold s-wave collisions are normally forbidden. This is a subject of an ongoing debate.

Systems based on Ca [33]†, Yb [34, 35], Sr [36, 37, 38, 39, 40], Mg [41] and Hg [42, 43] are being developed, with focus on both fermionic and bosonic isotopes. The leading systems are summarised in table 1.3.

*This is only true for a given polarisation, in reality the light shift due to the lattice laser has scalar, vector and tensor components which must all be considered.

†Thermal beam, not confined in a lattice.

1.5 Tests of fundamental physics

In addition to their potential to realise a more accurate SI second, the extreme precision offered by optical frequency standards makes them sensitive tools for tests of fundamental physics. Some atomic states have large relativistic corrections to their energy, and therefore have a strong dependence on the fine structure constant, α. By making repeated measurements of a frequency ratio $R = \nu_1/\nu_2$ between two optical transitions to or from states with energies that depend strongly on α, it is possible to amplify any time-variation in this constant [46, 47]:

$$\frac{\dot{R}}{R} = (A_1 - A_2)\,\frac{\dot{\alpha}}{\alpha} \tag{1.5}$$

where A_1 and A_2 are sensitivity coefficients for the two transitions. This forms a stringent test of the validity of local position invariance. The present limit to this time-variation comes from high-accuracy comparisons of Al^+ and Hg^+ optical clocks at NIST, and currently stands at $\dot{\alpha}/\alpha = (-1.6 \pm 2.3) \times 10^{-17}$ year^{-1} [29].

$^{171}Yb^+$ in particular lends itself to such a test as it possesses two clock transitions with different sensitivity to $\dot{\alpha}$, and the sensitivity coefficients are of opposite sign thus further amplifying any change. The ratio between the frequency of the $^2S_{1/2} \rightarrow {}^2F_{7/2}$ electric octupole (E3) transition at 467 nm and that of the $^2S_{1/2} \rightarrow {}^2D_{3/2}$ electric quadrupole (E2) transition at 436 nm has an overall sensitivity coefficient $(A_{E3} - A_{E2})$ of -6.83 [48, 49], making it more than twice as sensitive as the Al^+/Hg^+ comparison. The benefit of measuring an optical frequency ratio in a single ion rather than between different systems is that there will be a cancellation of any common-mode frequency shifts such as the gravitational redshift and the second-order Doppler effect. There is also no need to correct for the blackbody radiation shift, as long as the conditions in the trap are repeatable between measurements.

1.6 Thesis outline

This thesis details the $^{171}Yb^+$ optical frequency standard at the NPL, focusing on recent improvements to the experimental setup and concluding with absolute frequency measurements of the two optical 'clock' transitions. In chapter 2, the fundamentals of the trapping and cooling of single ions are introduced. The lasers required to cool and manipulate ytterbium-171 are described in chapter 3. Chapter 4 contains a discussion of the two ultrastable 'clock' lasers that form the oscillators for the clock. The coherent interaction of these lasers with their respective transitions is explained in chapter 5. The lasers are actively stabilised to these transitions, with the techniques detailed in chapter 6. The various systematic effects that perturb the transition frequencies are covered in chapter 7. These systematic shifts are then evaluated in order to measure the absolute frequencies of the two transitions with the results outlined in chapter 8. Finally, conclusions are drawn in chapter 9 along with suggestions for future work in order to improve the performance of the standard even further.

This thesis aims to provide the reader with an overview of the experiment as a whole, but the large-scale and complex nature of the apparatus means that many other people have contributed greatly to the measurements undertaken in this work. Helen Margolis, Veronika Tsatourian, Barney Walton and Luke Johnson have operated the optical frequency combs in order to make absolute frequency measurements. The majority of the experimental control software was written by Barney Walton and Rachel Godun. The ultrastable optical cavities used to stabilise the two probe lasers were designed and constructed by Stephen Webster. The two ion traps were built by Matthew Roberts and Kazu Hosaka. Finally, Guilong Huang and Stephen Webster played a large part in the running of the experiment in the build-up to the absolute frequency measurement of the E2 transition performed in 2009.

Chapter 2

Single ion trapping and cooling

A system consisting of a single ion 'at rest' has many advantages for use as a frequency standard. In particular the fact that the ion is tightly confined means that systematic shifts associated with the motion of the ion are low. The long lifetime of the trap allows for very long interrogation times and thus extremely narrow resonances can be resolved*. In addition, environmental conditions need only be known at the one point where the ion resides.

This chapter will cover the main techniques involved in realising such a system. Section 2.1 contains an explanation of the basic operation principles of the ion trap, and the concept of laser cooling of the trapped ion is introduced in section 2.2. The detection of the ion is described in section 2.3. Section 2.4 describes the loading of the trap, and measurements of its depth are detailed in 2.5. The importance of a low ion heating rate is explained in section 2.6, and section 2.7 covers how the residual motion of the ion can be minimised through a technique known as rf-photon correlation.

*As discussed in chapter 5, in the absence of any other line-broading mechanisms the resolveable linewidth is inversely proportional to the time spent probing the system.

2.1 Confining a single ion

2.1.1 Creating an ion trap

Owing to their charge, ions can be trapped by electric fields. However, it is impossible to create an electrostatic confining field in three dimensions as it would violate the Laplace condition (this is also referred to as Earnshaw's theorem):

$$\nabla^2\Phi = 0 \tag{2.1}$$

Hence any electrostatic potential Φ that is trapping in two dimensions must be anti-trapping in the other dimension ('saddle' potential). In a static potential of this geometry, an ion would be quickly lost in the direction of anti-trapping. However, the ion can be retained by adding a time-dependence to the potential ($\Phi(t)$) so that the ion sees a time-averaged pseudopotential that is trapping in all three dimensions. This revelation led to Wolfgang Paul receiving the Nobel Prize in 1989.

2.1.2 The Paul trap

The most common design of single-ion Paul trap is a set of two dc 'endcaps' with a ring carrying the ac voltage as shown in figure 2.1a. This design has one inherent problem - the ring limits access for laser beams and restricts the line of sight for optics used to image the ion.

A modified design is shown in figures 2.1b and 2.2. The ring has been removed and replaced with two 'inner' endcaps that sit inside the dc 'outer' endcaps [50], with the two separated by a sleeve of dielectric material.

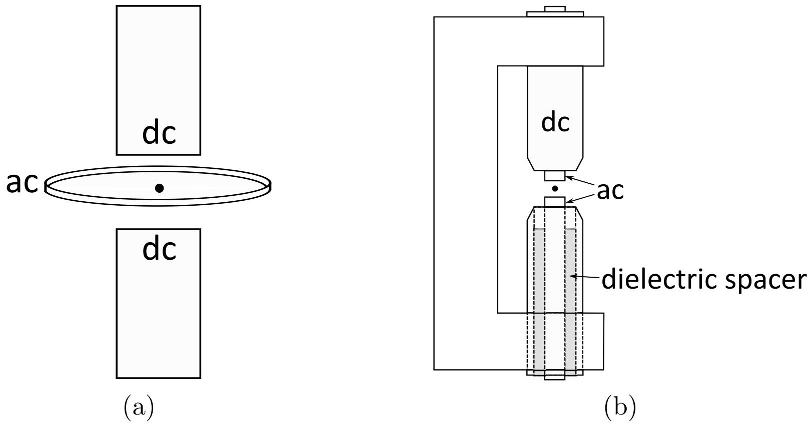

Figure 2.1: The two main designs of single ion Paul traps. (a) Ring trap, with the ac voltage applied to the outer ring. (b) Endcap trap, with the ac voltage applied to the inner endcaps.

Figure 2.2: Photograph of the NPL endcap trap design [51] with the inner and outer endcaps visible. Ceramic ovens containing isotopically enriched metal (bottom) and the two compensation electrodes (centre) surround the trap in the horizontal plane. The inner endcaps have a radius of 250 μm and are separated by 560 μm.

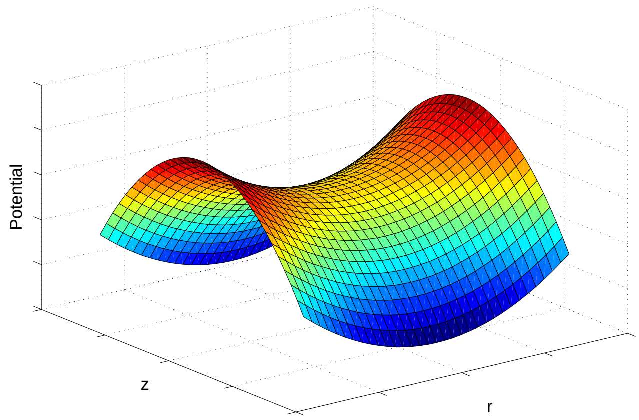

(a)

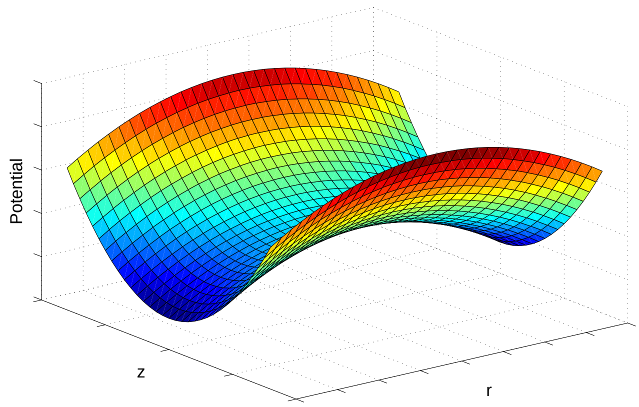

(b)

Figure 2.3: Saddle potential for the phase of the ac drive voltage equal to (a) 0 and (b) π.

2.1.3 The electrodynamic potential

The quadrupolar potential produced by the trap electrodes is shown in figure 2.3 and takes the form:

$$\Phi = \frac{\xi(V_{\mathrm{dc}} + V_{\mathrm{ac}})}{2r_0^2}(r^2 - 2z^2) \tag{2.2}$$

where $V_{\mathrm{ac}} = V_0 \cos(\Omega t)$ is the potential applied to the inner endcaps at frequency $\Omega/2\pi$, V_{dc} is a static voltage applied to the two outer endcaps, and r_0 is the distance between the rf electrode and the trapped ion. It can be seen that this potential satisfies the Laplace condition $\nabla^2\Phi = 0$. ξ is referred to as the efficiency of the trap and is a measure of how well the applied voltage is converted into the quadrupolar potential. For an ideal endcap trap, $\xi \approx 0.6$ [50].

The motion of an ion in such a potential is given by solutions of the Mathieu equation:

$$\frac{d^2x_i}{dt^2} + \frac{\Omega^2}{4}(a_i - 2q_i \cos\Omega t)x_i = 0 \tag{2.3}$$

where i represents the degrees of freedom. As the trap is cylindrically symmetric, $i = z, r$. The trap is characterised by the stability parameters a and q:

$$-a_z = 2a_r = \frac{8e\xi V_{\mathrm{dc}}}{mr_0^2\Omega^2} \tag{2.4}$$

$$q_z = -2q_r = \frac{4e\xi V_{\mathrm{ac}}}{mr_0^2\Omega^2} \tag{2.5}$$

A dc confining potential is required for axial confinement in linear ion traps as the ac potential is purely radial, but this is not required in the case of the endcap trap as the rf field also provides the axial confinement and so $a_i = 0$. Bound motion only exists in the trap for certain values of $\beta_i \approx \sqrt{a_i + q_i^2/2}$, as shown on the stability diagram (figure 2.4). There are also lines of instability within the 'stable' region that must be avoided. These lines generally are absent for $q_i < 0.4$, and so historically the traps have been operated with $q_z \approx 0.4$.

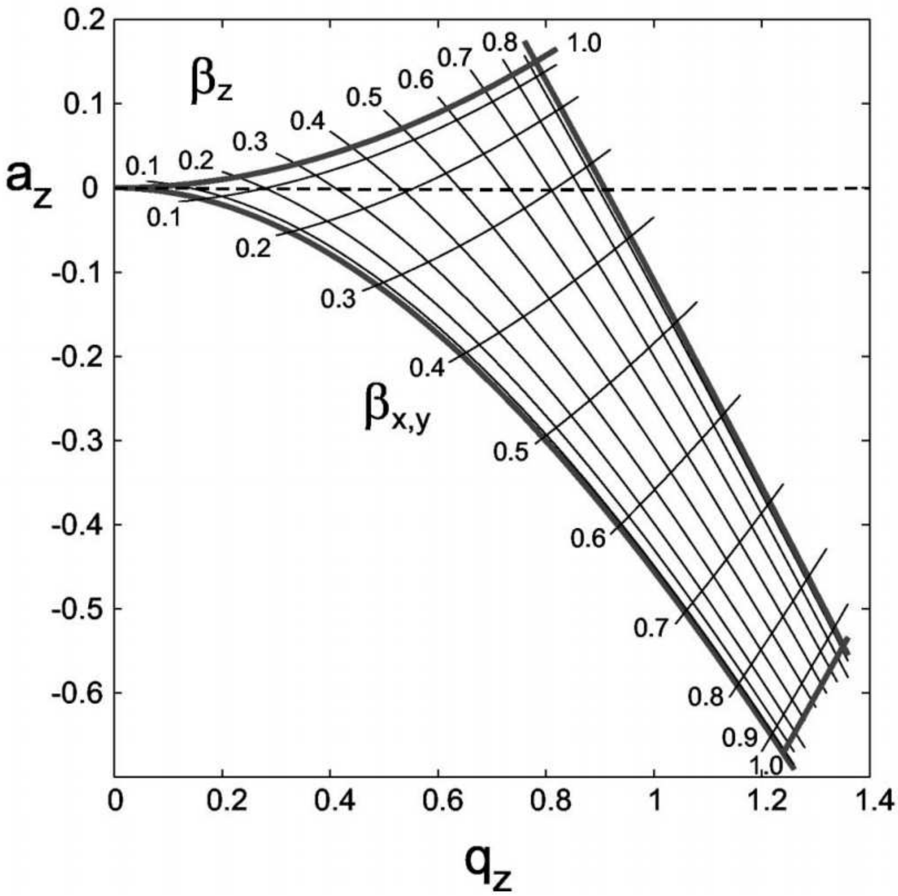

Figure 2.4: Stability diagram for a cylindrically symmetric ion trap, reproduced from [52]. The trap is stable within the region bounded by $0 < \beta_i \approx \sqrt{a_i + q_i^2/2} < 1$.

The full derivation of the solution to the Mathieu equation in the case of $q < 0.4$ can be found in reference [53] and is summarised well in the thesis of Matthew Roberts [54]. The overall motion of the ion is described by:

$$z(t) = \bar{z} + z_a \cos\omega_z t - \frac{z_a q_z}{4}[\cos(\Omega + \omega_z)t + \cos(\Omega - \omega_z)t] - \frac{\bar{z} q_z}{2}\cos\Omega t \tag{2.6}$$

where $\bar{z}$ is the mean position of the ion in the trap. z_a is the amplitude of the slow 'secular' oscillations of the ion in the trap that occur at frequency ω_z. The third term is referred to as intrinsic micromotion and is always present. The final term represents excess micromotion which occurs at the trap drive frequency and is only present if $\bar{z} \neq 0$, i.e. the ion has been displaced from trap centre. This is usually the result of stray electric fields or inaccuracies in the machining of the trap. By applying dc voltages to an array of 'compensation' electrodes, the ion can be returned to trap centre. This is important to ensure the ion does not see any residual rf field, and also has only a small residual velocity. The importance of this is discussed in chapter 7.

2.1.4 Choice of trap drive parameters

In general, it is desirable for the ion to have high secular frequencies. This improves the efficiency of Doppler cooling (i.e. the ion is cooled to a smaller vibrational quantum number $\langle n \rangle$, see subsection 2.2.2), and reduces the rate of heating by reducing the ion's sensitivity to low frequency electric field noise (as discussed in subsection 2.6). The secular frequency is given by:

$$\omega_i = \frac{\beta_i \Omega}{2} \approx \frac{q_i \Omega}{2\sqrt{2}} \tag{2.7}$$

and so the secular frequency can be increased by increasing q or by increasing Ω. These quantities are not independent, as discussed earlier. q_i is generally fixed in order to create a stable trap. In order to increase Ω by a factor of two whilst maintaining a constant q_i requires the ac voltage to be increased by a factor of four. As the trap drive signal is generally produced by an amplified synthesiser or by direct digital synthesis (DDS), an increase of a factor of four in voltage becomes an increase of a factor of 16 in applied power, which is often not feasible (e.g. the NPL Yb^+ traps operate at around 13 MHz, and the trapping voltage is produced by stepping up the output of an amplifier working at 3 W. Increasing the drive frequency to 26 MHz would require a forward power of 48 W).

2.1.5 Generating the ac drive signal

Ion trap drive signals have typical peak-to-peak amplitudes of several hundred volts. There are several ways of generating such signals, and until recently the two ytterbium traps were driven by self-oscillating transistor circuits. This circuit suffered from high harmonic levels, as shown in figure 2.5, and required a high dc voltage source. A better option is a resonant transformer, which can amplify a small ac supply voltage to the levels required to trap an ion. A simplified circuit diagram for such a circuit is shown in figure 2.6. Two designs of resonant transformer were tried, the first being a large helical resonator (see figure 2.7). This design was found to be quite temperamental, and did not reliably store ions overnight.

The second design tried is shown in figure 2.8. The circuit design is a simple resonant step-up transformer based on the mutual inductance between two coupled coils wrapped around an iron powder core (Micrometals T50-6). The circuit resonates at a frequency given by:

$$\omega^2 = \frac{1}{L_1 C_1} = \frac{1}{L_2 C_2} \tag{2.8}$$

Maximum coupling [55] occurs when:

$$\omega M = \sqrt{R_1 R_2} \tag{2.9}$$

where $M = k\sqrt{L_1 L_2}$ is the mutual inductance of the two coils (k is a dimensionless coupling constant between 0 and 1). The inductance of each coil is given by:

$$L = N^2 A_L \tag{2.10}$$

Parameter	Helical resonator	Toroidal transformer
Primary turns	2	1
Secondary turns	12	27
Unloaded Q	1210	180
Loaded Q	91	58
Loaded resonance	12.5 MHz	13.3 MHz
VSWR on resonance	< 1.1	< 1.2

Table 2.1: Summary of the two trap drives trialled to replace the self-oscillating transistor circuit.

where N is the number of turns and A_L is the inductance rating of the toroid in H/turn^2. Finally, the resistance of each coil is determined through the quality factor Q_{ind} of the toroid:

$$R = \frac{\omega L}{Q_{\text{ind}}} \tag{2.11}$$

The output voltage from the finished transformer (assuming perfect impedance matching) is given by:

$$\frac{V_{\text{out}}}{V_{\text{in}}} = \frac{\omega L_2}{2\sqrt{R_1 R_2}} \tag{2.12}$$

Impedance matching is generally measured in terms of the voltage standing-wave ratio (VSWR) at resonance:

$$\text{VSWR} = \frac{1 + V_{\text{r}}/V_{\text{f}}}{1 - V_{\text{r}}/V_{\text{f}}} \tag{2.13}$$

where V_{f} and V_{r} are the forward and reflected voltages respectively. This was measured to be less than 1.2 at the resonance frequency of 13.3 MHz using an in-line VSWR meter, indicating a coupling of better than 91%. The quality factor of the loaded resonance was measured to be 58. The transformer is driven by an amplified synthesiser, which results in a clean drive signal with low harmonic levels as shown in figure 2.5. To reach $q_r = 0.2$, $q_z = 0.4$ requires around 3 W to be applied to the primary coil of the transformer. The parameters for the toroid and helical resonator are summarised in table 2.1.

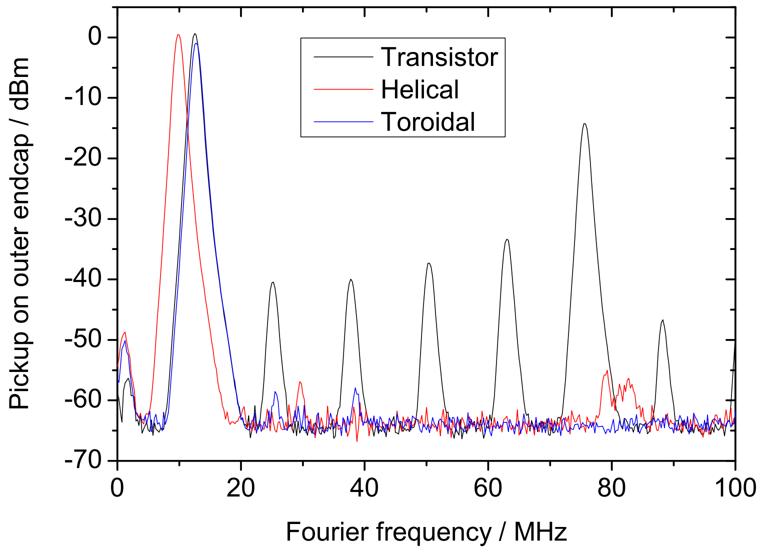

Figure 2.5: Comparison of the harmonics of the drive signal for a transistor oscillator circuit, helical resonator and toroidal transformer. These were monitored by observing the pickup on the outer endcaps when the inner endcaps were driven with each of the voltage sources. The feature near dc is the window function of the spectrum analyser and should be ignored. The resonant frequency of the loaded helical resonator was later increased from 9 MHz to 12 MHz by careful adjustment of the connection to the vacuum feedthrough.

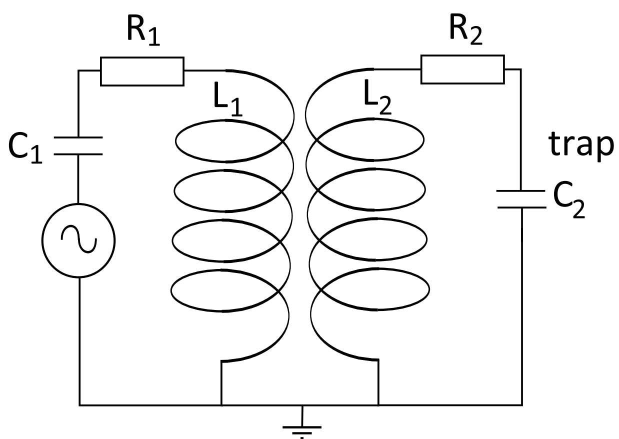

Figure 2.6: Simplified circuit diagram for the two types of step-up transformer used to drive the traps.

Figure 2.7: Photograph of the helical resonator. The secondary coil is wound around a Rexolite core and enclosed within a cylindrical copper shield. The primary coil is attached to one of the endcaps, and its length iterated to optimise coupling.

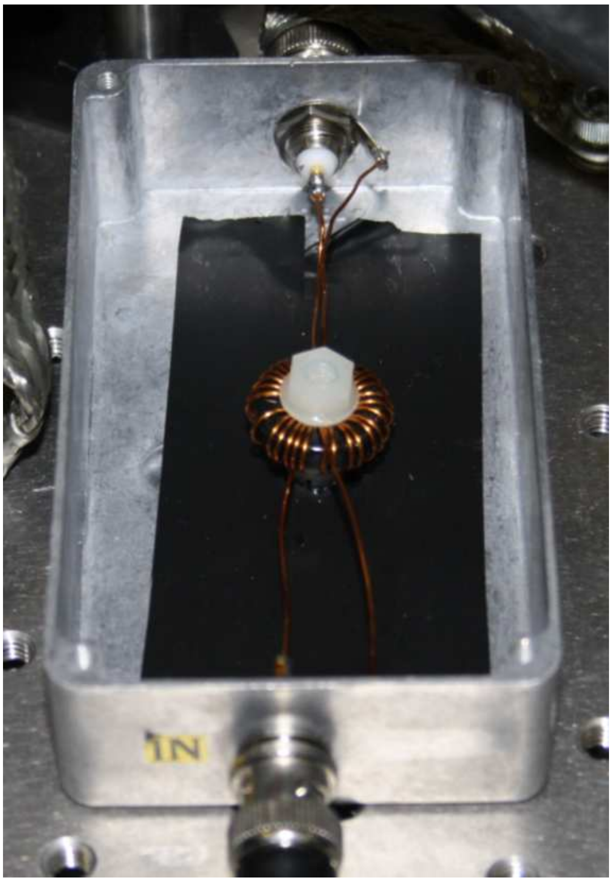

Figure 2.8: *In-situ* photograph of the toroidal resonator. The core is mounted on a nylon post for mechanical stability.

2.2 Laser cooling

2.2.1 The Lamb-Dicke regime

Motion in the z-direction is used as an example in this section, but the equations and conclusions are equally valid for motion in the radial plane.

Doppler shifts associated with the secular motion of the ion in the trap ($z = z_a \cos \omega_z t$) lead to a modulation of an interrogating laser ($E = E_0 \exp[i(kz - \omega t)]$) in the reference frame of the ion. The resulting Lorentzian spectral profile of the transition is described classically by:

$$I(\omega) = \sum_{-\infty}^{\infty} \frac{|J_n(\beta)|^2}{(\omega - [\omega_0 - n\omega_z])^2 + (\gamma/2)^2} \tag{2.14}$$

where ω is the laser frequency, J_n is the Bessel function of order n, $\beta = kz_a = 2\pi z_a/\lambda$ is the modulation index, ω_0 is the resonance frequency of the transition and γ is its natural linewidth. This equation describes a central carrier with an infinite number of sidebands spaced by ω_z.

Confining the ion to a region smaller than the wavelength of the light used to interrogate the transition ($\beta \ll 1$) leads to $J_0(\beta) \approx 1 - (\beta/2)^2$ and $J_1(\beta) \approx \beta/2$. Hence:

$$\frac{I_1}{I_0} \approx \left(\frac{\beta}{2}\right)^2 \tag{2.15}$$

in which case the first order sidebands are small compared to the carrier and higher order sidebands are completely absent. This is referred to as the Lamb-Dicke regime. Under this condition, first-order Doppler broadening is eliminated and the natural linewidth of the transition can be resolved. If the motional sidebands can be resolved, we can consider the ion spending most of its time 'at rest' and hence the interaction time with a laser tuned to the carrier frequency is maximised (see figure 2.9).

As the motion of the ion is quantised, this model requires a small modification. It is useful to think about the energy levels of the trap as a vibrational 'ladder'. The ion can

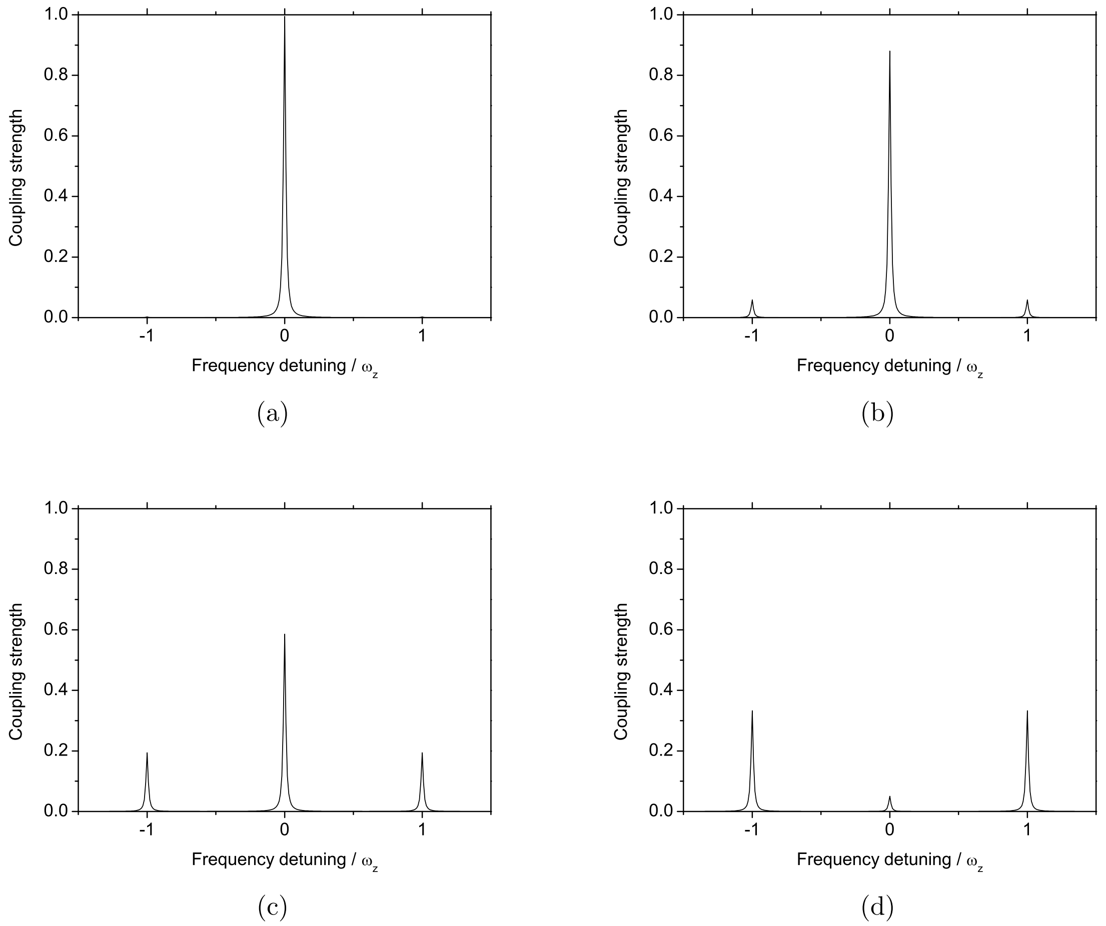

Figure 2.9: Profile of a transition with resolved motional sidebands, focusing on the carrier and first order sidebands in the cases of (a) $\beta = 0.1$, (b) $\beta = 0.5$, (c) $\beta = 1$ and (d) $\beta = 2$.

occupy a range of vibrational levels with a probability distribution [56] well described by:

$$p(n) = \frac{1}{1 + \langle n \rangle} \left(\frac{\langle n \rangle}{1 + \langle n \rangle} \right)^n \tag{2.16}$$

where n is an individual vibrational level and $\langle n \rangle$ is the expectation value of the vibrational levels. For the case of low $\langle n \rangle$, the motional sidebands become asymmetric [57]:

$$I_{-1} = \langle n \rangle \eta^2 \tag{2.17}$$

$$I_0 \approx 1 \tag{2.18}$$

$$I_{+1} = (\langle n \rangle + 1)\eta^2 \tag{2.19}$$

where $\eta = (k \cos\theta)\, z_0$ is known as the Lamb-Dicke parameter (θ is the angle between the direction of motion and the propagation vector of the interrogating laser). This represents the modulation index for the ground state of the harmonic well:

$$z_0 = \sqrt{\frac{\hbar}{2m\omega_z}} \tag{2.20}$$

2.2.2 Doppler cooling

Once confined in the trap, the ion can be cooled by driving a strong electronic transition (in the case of Yb^+, the $^2S_{1/2} \rightarrow {}^2P_{1/2}$ transition at 370 nm is used). If the laser is red detuned from the centre of the transition, when the ion is moving towards the laser it will be blueshifted into resonance in the frame of the ion. On absorption, the ion receives the momentum of the photon and hence is slowed. The subsequent re-emission of the photon is in a random direction, and so the net result of many absorption/re-emission cycles is a cooling of the ion. Conversely, if the laser is blue detuned from the centre of the transition then the ion will preferentially scatter photons when it is moving away from the laser, leading to a net heating. Optimum cooling occurs for the laser being

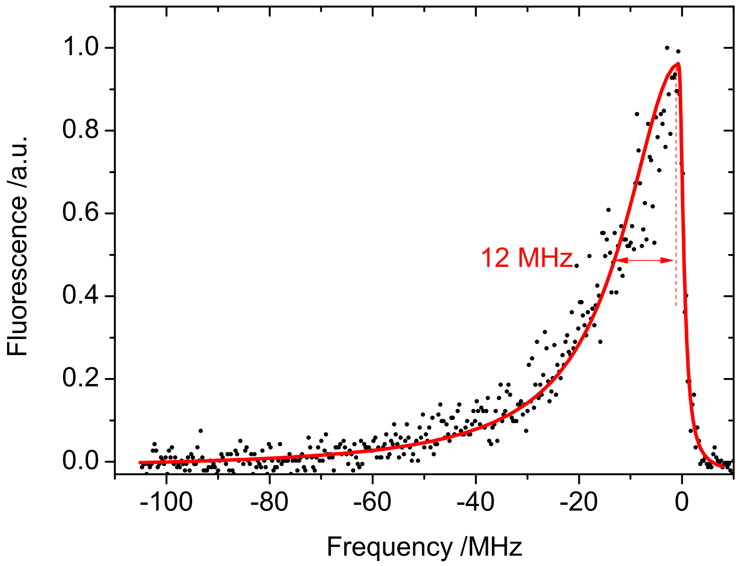

Figure 2.10: Lineshape of a broad cooling transition in a cold ion. Note the sharp decay to zero fluorescence at line centre as the laser becomes blue-detuned and begins to heat the ion.

red detuned to the half-maximum of the line profile. A typical lineshape produced by scanning the cooling laser over the strong transition is shown in figure 2.10.

The lowest temperature attainable by Doppler cooling is determined by the linewidth of the transition used to cool the ion. If this is greater than the secular frequency ($\omega_z < \gamma$) then the secular sidebands are unresolved ('weak binding limit'). The energy full-width of the excited state limits the energy of the ion after cooling to $\hbar\gamma/4$. This is referred to as the Doppler limit. The minimum temperature attainable by Doppler cooling is given by equating the thermal energy to the minimum kinetic energy:

$$\frac{kT_{min}}{2} = \frac{\hbar\gamma}{4} = \left(\langle n\rangle + \frac{1}{2}\right)\hbar\omega_z \tag{2.21}$$

The full-width of the 370 nm transition in Yb^+ is approximately $2\pi \times 20$ MHz, which leads to a Doppler temperature of about 0.5 mK. This corresponds to a mean vibrational level $\langle n\rangle = 9$ for typical radial secular frequencies of $\omega_r = 2\pi \times 1$ MHz and $\langle n\rangle = 5$ for the axial frequency of $\omega_z = 2\pi \times 2$ MHz.

2.3 Detecting the ion

The ion scatters on the order of 10^7 photons per second at 370 nm into the full 4π solid angle. Some of the scattered photons are collected by a lens with a numerical aperture of 0.35 and are imaged with a magnification of 10 onto a photomultiplier tube (PMT). The pre-amplified photo-pulses from the photomultiplier tube are sent to an Ortec amplifier-discriminator to remove spurious noise from the signal and generate TTL rising edges for each detected photon, which are then counted using a National Instruments counter card connected to the laboratory control PC. With the cooling transition well saturated ($\approx 5I_{\rm sat}$) and a bias field of 300 μT applied to destabilise the ground state of the transition (see section 3.4.1), 2.5×10^4 photons are detected per second. The main reasons for this low detection efficiency are the limited solid angle encompassed by the imaging lens ($\approx$ 3% of the full 4π steradians), losses in the material of the trap window and lens ($\approx$ 50%), and the quantum efficiency of the photomultiplier tube ($\approx$10-20%).

Several precautions are taken to ensure a low background rate on the fluorescence signal. A PMT with a low dark count rate of below 5 counts per second is used, and ambient light is excluded by a black cylinder between the trap vacuum system and the PMT. A colour (UG1) filter is also placed on the input to the PMT to ensure that only 370 nm photons are detected. The background rate attributable to scatter from the electrodes is reduced by placing a pinhole with 200 μm diameter in the image plane of the lens, resulting in a background detection rate of only 100 photons per second. Unfortunately this comes at the expense of some signal as the imaging system suffers from some spherical aberration, and the peak detection rate is reduced to 1.6×10^4 photons per second.

2.4 Loading the trap

A thermal beam of neutral ^{171}Yb is produced by a small ceramic oven containing isotopically-enriched metal foil, through which a current of 6 A is passed for up to ninety seconds.

Around 0.5 mW from a home-built, free-running extended cavity diode laser at 399 nm is focused to a waist of 5-10 μm at the centre of the trapping region in order to drive the strong $^1S_0 \rightarrow {}^1P_1$ transition in the neutral atoms. No attempt is made to load in an istopically selective manner, but it is only on very rare occasions that we observe loading of an incorrect isotope.

The oven flux is sufficiently low that atomic fluorescence is not observed by the PMT monitoring the trapping region. The absorption of a photon at 370 nm from the cooling laser (10 μW focused to a 30 μm waist) then excites the electron to the continuum. In around 50% of cases, a single ion is loaded. The presence of additional ions is obvious since: a) higher peak fluorescence rates are observed, and b) the cooling transition profile is modified [54]. If multiple ions are loaded, the cooling laser is blocked for periods of up to fifteen minutes. During this time there is a lot of anomalous heating due to interaction between the ions, and ions are ejected from the trap until only one remains. At this point the heating rate slows to a level dependent on the electric field noise seen by the ion. Re-application of the cooling laser then resumes Doppler cooling.

2.5 Measuring the potential depth of the trap

The trap depth can be determined by measurement of the secular frequencies of the ion. A small oscillating voltage is applied to a nearby electrode, thus 'tickling' the ion at this frequency. If the drive frequency is matched to one of the ion's motional frequencies, the ion couples strongly to the drive voltage and is heated. This leads to an increased Doppler shift on the cooling transition. If the cooling laser is initally far red detuned, an increase in the scatter rate can be detected when the resonance is found. The radial secular frequency is shown as a function of rf drive power in figure 2.11. The secular frequency scales linearly with the trapping voltage, and hence should scale as the square root of the drive power P. However, the secular frequency was found to scale roughly as $P^{0.3}$. It was found that the VSWR for the toroid rises from below 1.2 at 1 W of

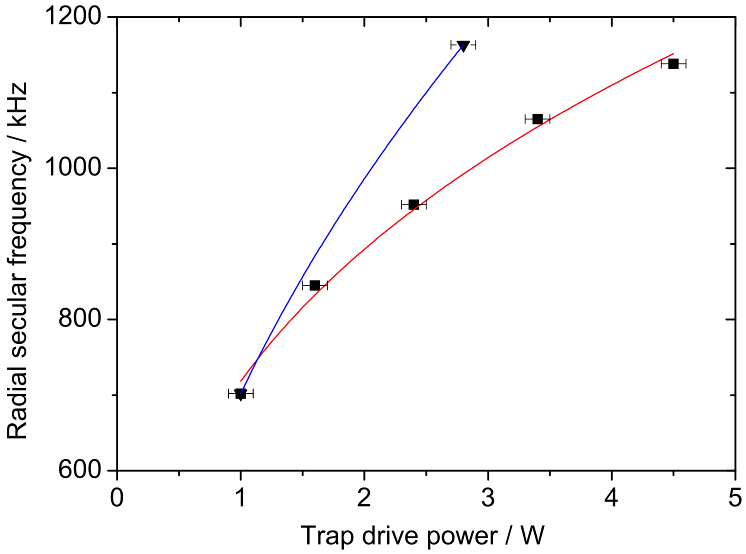

Figure 2.11: Radial secular frequency of a single ion against forward rf power into the toroidal transformer. The square points were taken with a single half-turn on the primary coil, and the triangular point at 2.8 W was taken with one full turn on the primary coil.

drive power to over 1.4 at 4 W of drive power, which was attributed to resistive loss causing heating and hence a change in the impedance matching of the toroid. Once the impedance mismatch was corrected by adjusting the number of turns on the primary coil, the data point at 2.8 W was re-taken and is consistent with a scaling of $P^{0.5}$ as expected. At this secular frequency, the ion experiences a peak quadrupolar trapping voltage of $\xi V_{ac} = 92$ V. Assuming $\xi = 0.6$, a peak of 154 V is applied to the inner endcaps.

2.6 Heating rate and ion storage

To prevent heating of the ion due to collisions with background room-temperature gas molecules, the trap is mounted within a cubic vacuum chamber maintained at a pressure of 10^{-9} Pa (10^{-11} Torr). This also acts to prevent collisional quenching of the upper states of the clock transitions.

The presence of noise in the electric field used to confine the ion can excite its motional degrees of freedom, leading to heating. The effect of this noise scales favourably with the electrode-ion separation d (comparison of heating rates in various traps have led to an empirical scaling of d^{-4}, which matches some theoretical models [58]), and hence heating rates in macroscopic traps tend to be smaller than those in miniature traps. The primary cause of this noise has been theorised to be 'patch' potentials arising from the trap loading process. It is therefore desirable to retain an ion for as long as possible to minimise the disruption caused by the deposition of ytterbium metal on the trap electrodes. For similar reasons it is also desirable to load the trap using photoionisation rather than electron beam bombardment as it is more efficient (meaning a lower oven flux) and it does not produce as many stray charges.

In some systems, the primary ion loss mechanism is chemistry with the residual gas present in the vacuum chamber, usually hydrogen. After the reaction, the Doppler cooling laser is no longer resonant with the charged molecule and hence the ion is lost from the trap. By a fortunate coincidence, the dissociation wavelength of YbH^+ is close to 370 nm [59, 60]. Hence the molecule is dissociated almost as soon as it forms, leading to very long ion storage times. Ions have been observed to remain in the trap for over 8 hours without laser cooling. With continuous cooling, ions have been retained for well over a month (and only lost after a prolonged absence of cooling light after a laser mode-hop over a weekend).

2.6.1 Ground state cooling

After cooling, the ion occupies a distribution of vibrational levels determined by the Doppler limit of the cooling transition (as described in section 2.2.1). The ion can be further cooled to the ground state of the trap ($\langle n \rangle = 0$) using transitions with linewidths that are sufficiently narrow to allow the secular sidebands to be resolved ('strong binding limit'). Often it is a dipole-forbidden transition (e.g. a clock transition) that is used for this purpose, but ground-state cooling has been demonstrated using coherent effects such as electromagnetically-induced transparency (EIT) to produce narrow absorption features on the strong cooling transition. In order to cool the ion, the first red motional sideband of the transition is excited (see figure 2.12), leading to the loss of one motional quantum. The ion is then cleared out from this state in a manner that involves a spontaneous decay, which (on average) does not change the vibrational quantum number. The ion is thus pumped to the vibrational ground state of the trap by repeated trips around this cycle.

The main advantages of ground-state cooling in a single-ion frequency standard experiment are an increase in excitation rate (see section 5.3) and a minimisation of systematic shifts associated with the ion's motion. In this experiment and other single-ion frequency standards, no such cooling is required as the ion is in the Lamb-Dicke regime and the loss of contrast on the first Rabi flop is small. In addition, as the cooling light is not present during the probe phase, the ion will heat and will not remain in the ground state for very long. This is not a problem in quantum information processing (QIP) experiments where probe pulse durations are typically on the order of microseconds, but in an optical clock where the probe time approaches one second any minor gains in contrast or systematic shifts are almost immediately lost.

Clocks based on quantum logic spectroscopy do however require ground state cooling, but for a different purpose. A clock ion with no accessible cooling transition (e.g. Al^+) is sympathetically cooled by another 'logic' ion that does possess such a transition (e.g. Be^+), resulting in the formation of a two-ion crystal. In order to read out the state of the Al^+ ion after a probe pulse, it is necessary to perform a logic operation using its vibrational quantum number as a register [61] and hence the ion crystal must first be initialised in the ground state of one of its motional degrees of freedom.

2.7 Micromotion minimisation

As mentioned earlier, if the ion is displaced from trap centre it undergoes motion at the trap drive frequency, Ω. If not corrected for, this leads to: a distorted cooling lineshape that affects the efficiency of laser cooling; a large second-order Doppler shift on the clock transitions; reduced coupling of the probe laser to the clock transition; and a Stark shift

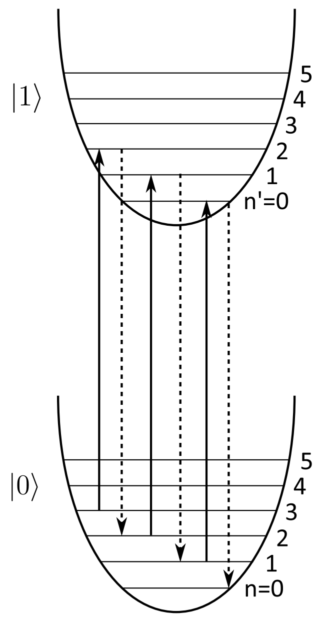

Figure 2.12: Simplified depiction of cooling to the zero point energy of the harmonic well using a transition with resolved motional sidebands. A laser excites the $\Delta n = -1$ component of the $|0\rangle \rightarrow |1\rangle$ transition, and a spontaneous decay follows on the $\Delta n = 0$ component. Repeated trips around this cycle result in the ion being pumped to the ground state. In a real system, the state $|1\rangle$ can be metastable and so the state is quenched to a higher-lying, shorter-lived state from which the spontaneous decay occurs.

on the clock transitions from the residual rf field where the ion resides. It is possible to determine the ion's motion in three dimensions using a technique referred to as rf-photon correlation.

The motion of the ion leads to a modulation of the fluorescence rate. If the ion is moving towards the beam, it observes the laser to be blueshifted towards line centre and more photons are scattered. If the ion is moving away, the laser is redshifted and fewer photons are scattered. Crucially, this modulation occurs at the frequency of the ion's motion and hence can be extracted.

A sample of the trap drive signal is converted to a sequence of pulses by an amplifier-discriminator. A time-to-amplitude converter (TAC) is armed when it receives a rising edge from the discriminator. Once armed, the TAC is stopped when it receives a rising edge from the photomultiplier tube. A PC interface card is used to record the time delays between the rf edge and the detected photon. A plot of detected photons against delay time shows modulation at the trap drive frequency above a flat background corresponding

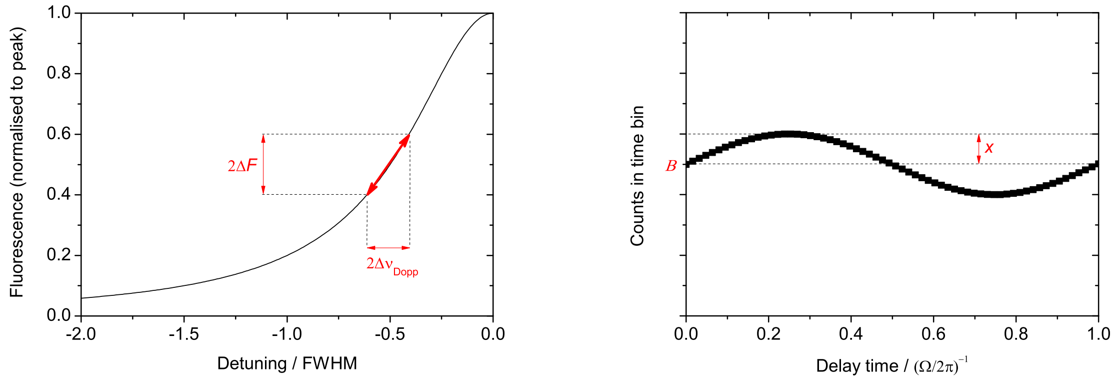

(a) (b)

Figure 2.13: (a) Illustration of the modulation of the fluorescence level by the ion's motion. (b) Corresponding rf-photon correlation signal showing modulation at the trap drive frequency over a flat background.

to the mean photon scatter rate as shown in figure 2.13b. The depth of the modulation corresponds to a Doppler shift, with the cooling lineshape providing the frequency discriminant as shown in figure 2.13a. If the cooling laser is tuned to the half-maximum of the cooling transition, this discriminant is maximised and can be approximated to be linear for low levels of modulation:

$$\frac{2\Delta F}{2\Delta\nu_{\mathrm{Dopp}}} = \frac{F_{\mathrm{max}}}{\delta\nu} \tag{2.22}$$

where ΔF is the peak modulation of the fluorescence rate, $\Delta\nu_{\mathrm{Dopp}}$ is the peak Doppler shift caused by the motion, F_{max} is the maximum fluorescence rate and $\delta\nu$ is the full-width at half-maximum of the cooling lineshape. For a peak modulation of x counts on the rf-photon correlation signal over a background of B counts, the corresponding fluorescence modulation is given by:

$$\frac{x}{B} = \frac{\Delta F}{F_{\mathrm{max}}/2} \tag{2.23}$$

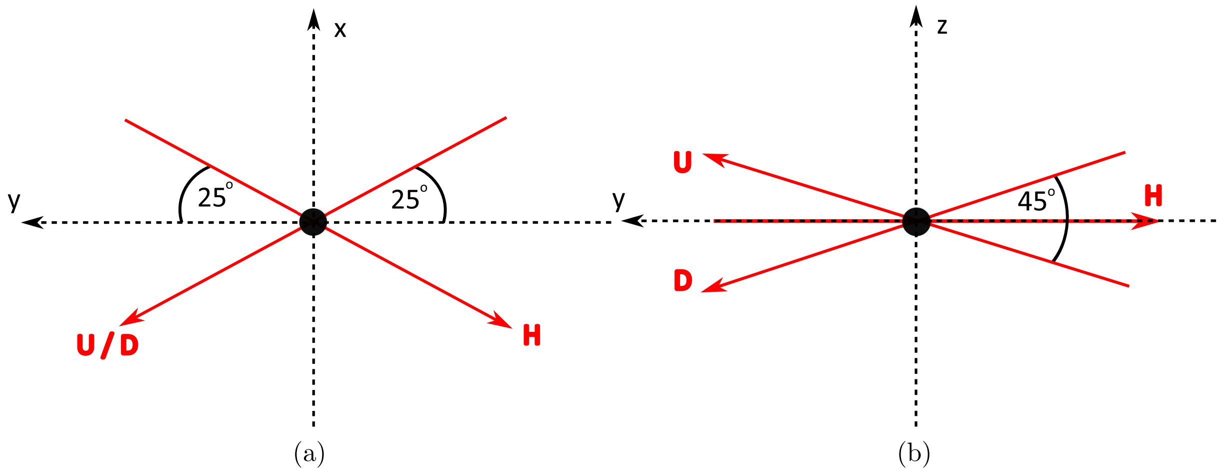

Figure 2.14: Geometry of the three cooling beams used to determine the ion's excess micromotion. The 'horizontal' (H) beam lies in the plane perpendicular to the rf electrodes. The 'up' (U) and 'down' (D) beams are tipped oppositely out of the horizontal plane by 22.5°.

This leads to a peak Doppler shift of:

$$\Delta\nu_{\text{Dopp}} = \frac{x}{2B}\delta\nu \tag{2.24}$$

from which the peak velocity v_{max} of the ion can easily be determined:

$$\frac{v_{\text{max}}}{c} = \frac{\Delta\nu_{\text{Dopp}}}{\nu_{370}} \tag{2.25}$$

where ν_{370} is the frequency of the 370 nm transition.

The depth of modulation is minimised by applying dc voltages to compensation electrodes around the centre of the trap. In order to allow 3D minimisation, three non-planar cooling beams are used to ensure there is no hidden direction of motion. One beam is in the horizontal plane of the trap and the other two are oppositely tipped out of the horizontal plane by 22.5° (this is the maximum angle allowed by the vacuum chamber windows). The micromotion is minimised along the line-of-sight of each of the beams in turn, and the process iterated until no further reduction in modulation depth is observed. A typical correlation signal after rough minimisation is shown in figure 2.15.

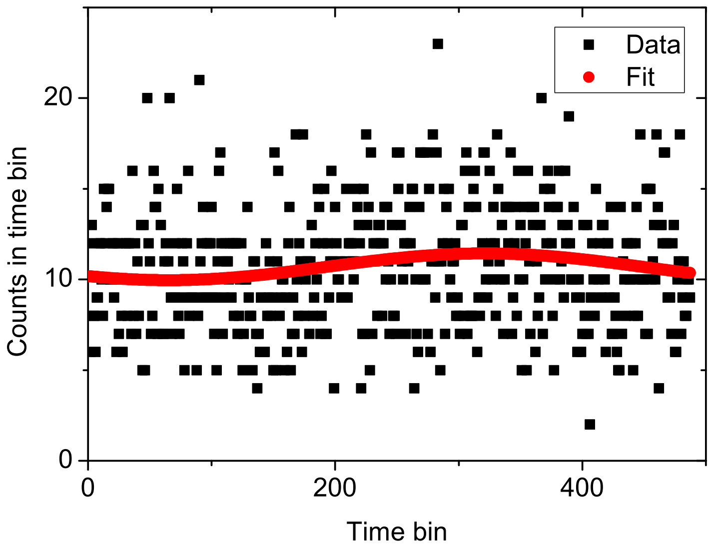

Figure 2.15: Typical rf-photon correlation signal for the horizontal cooling beam, with the depth of modulation corresponding to a peak Doppler shift of 1 MHz. This gives a peak residual velocity of 38 cm/s in this direction.

2.7.1 Automated micromotion compensation

Recently, software has been developed to allow the excess micromotion to be compensated automatically. A fast Fourier transform (FFT) of the rf-photon correlation signal is performed, and the amplitude and phase of the peak corresponding to the excess micromotion are extracted. The phase of the modulation takes values close to $\pm\pi$ (i.e. the modulation is either in phase or in antiphase with the drive voltage), and changes sign as the optimum compensation voltage is passed. By calibrating the sign of the phase and the modulation depth as a function of voltage (as shown in figure 2.16), an error signal can be generated and a digital servo designed in order to monitor and compensate this excess motion continuously. These servos have not yet been implemented as the remotely tuneable high voltage supplies are still under construction.

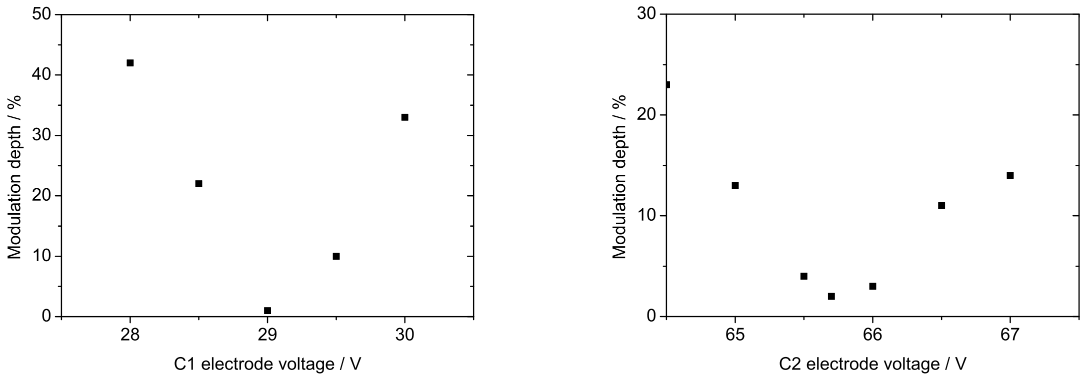

(a) (b)

Figure 2.16: Modulation depth of the excess horizontal micromotion on the 370 nm cooling transition as a function of voltage on the two compensation electrodes (a) C1 (used to compensate the motion visible to the 'horizontal' beam) and (b) C2 (used to compensate the motion visible to the 'up' and 'down' beams). The cooling laser was red-detuned to the half maximum of the lineshape, which was below saturation to prevent power broadening of the transition lineshape.

2.8 Summary

This chapter has outlined the necessary design considerations for an ion trap that will operate as part of an optical frequency standard. Various sources of the rf trapping voltage were discussed, with the final choice being a small toroidal resonator. The concept of Doppler cooling was introduced, along with the techniques for minimising the residual motion of the ion in the trap.

The following chapter contains a discussion of the lasers required to maintain effective cooling of the $^{171}Yb^{+}$ ion once it has been confined within the trap.

Chapter 3

Laser cooling of ytterbium ions

All of the candidates under investigation for a future re-definition of the SI second have advantages and disadvantages. This chapter will discuss the relative merits and drawbacks of singly-ionised ytterbium-171 as a candidate, along with the technical details of its cooling. The benefits of an alkali-like electronic structure are given in section 3.1. In section 3.2, a brief introduction is given to the j,L coupling that is required to understand some of the low-lying electronic configurations in $^{171}Yb^{+}$. Section 3.3 focuses on the advantages of $^{171}Yb^{+}$ as an optical frequency standard. The various lasers required for cooling and repumping the ion are discussed in section 3.4. Section 3.5 gives the details of a basic re-locking scheme that has been implemented on the cooling laser. Section 3.6 describes how the ion is 'initialised' into the correct hyperfine component of the ground state, ready to be interrogated by one of the probe lasers. The overall experimental setup is detailed in section 3.7.

3.1 Alkali-like ions

The main advantage of alkali-like ions (species with a single valence electron) over group III ions (such as aluminium and indium) is their strong S→P transitions, which generally occur in the visible region of the electromagnetic spectrum where lasers are readily available. This makes them straightforward to laser cool, and provides a means of detection of the ion's internal state (as explained in section 6.1.2). The corresponding transition in Al^+ lies at 167 nm, which is far beyond the reach of current continuous wave (cw) laser technology.

The disadvantage of an alkali-like electronic structure is that states with total angular momentum quantum number $J \geq 1$ have an electric quadrupole moment, which leads to a systematic shift of the clock transition frequency when coupled to any electric field gradients in the trap. This shift is not present in singly-ionised group III atoms as the clock transitions are singlet $\rightarrow$ triplet transitions ($^1S_0 \rightarrow {}^3P_0$) and hence do not involve states with high angular momentum.

3.2 A note on atomic structure

Ytterbium is a heavy lanthanide metal with atomic number $Z = 70$. When singly ionised, its ground state electronic structure consists of a full 4f shell with a single 'outer' 6s electron making it alkali-like. However, many of its low-lying excited states consist of a hole in the 4f shell ($4f^{13}$) and two valence electrons. Under these conditions the total angular momentum of the core j_c can couple to the orbital angular momentum of the valence electrons L_v, giving the quantum number $K = j_c + L_v$. The overall angular momentum of the system J is given by $J = K \pm S_v$. The resulting term symbol has the form:

$$^{2S_v+1}[K]_J \tag{3.1}$$

3.3 Ytterbium-171

Ytterbium is unique amongst the ions currently being investigated as optical frequency standards in that its lowest-lying excited state is $^2F_{7/2}$ (see figure 3.1). Once populated, this state can only decay to the ground state via electric octupole (E3) radiation at

467 nm with a natural lifetime of around 6 years [62]. This leads to the $^2\mathrm{S}_{1/2}$ $(F = 0) \rightarrow$ $^2\mathrm{F}_{7/2}$ $(F' = 3)$ transition possessing an extremely narrow natural linewidth on the order of 1 nHz; this results in a resolvable transition linewidth that depends primarily on the linewidth of the probe laser used to interrogate the transition (in the absence of other line broadening mechanisms). As reference cavities used for probe laser stabilisation continue to improve, so will the achievable short-term stability of the standard. The fact that the properties of the state are determined by the hole in the 4f orbital means that there is screening by the full 6s orbital, resulting in lower sensitivity to external electromagnetic fields. Two recent, independent measurements of the absolute frequency of this transition are in very good agreement [18, 17].

In addition to the E3 transition, $^{171}\mathrm{Yb}^+$ also possesses two electric quadrupole (E2) transitions from the ground state. The $^2\mathrm{S}_{1/2}$ $(F = 0) \rightarrow$ $^2\mathrm{D}_{3/2}$ $(F' = 2)$ transition at 436 nm has a natural linewidth of 3.02(14) Hz [63], and the 411 nm $^2\mathrm{S}_{1/2}$ $(F = 0) \rightarrow$ $^2\mathrm{D}_{5/2}$ $(F = 2)$ transition has a natural linewidth of 23(1) Hz [63]. It is the former transition that is of interest as an optical frequency standard due to its narrower natural linewidth. Two independent measurements of the absolute frequency of this transition are in good agreement [28, 64], and it is already accepted as a secondary representation of the SI second [65].

One of the biggest advantages of the odd isotopes of Yb^+ are their half-integer nuclear spin, I. When coupled to states with half-integer electronic spin S, this leads to hyperfine states with integer F and hence $m_F = 0$ sublevels. To first order, $m_F = 0 \rightarrow m_{F'} = 0$ transitions are insensitive to magnetic fields. It is these components of the 436 and 467 nm transitions that are investigated in this work. The isotope investigated in this work, $^{171}\mathrm{Yb}^+$, has $I = 1/2$. This results in the presence of an $F = 0$ sublevel in the ground state. This greatly simplifies initialisation of the ion before probing, as optical pumping into this state can be achieved with near-100% efficiency with π-polarised light.

Ytterbium also has the advantage of a low reaction rate with any background gas. Chemistry with residual gas in the vacuum system is one of the primary loss mechanisms

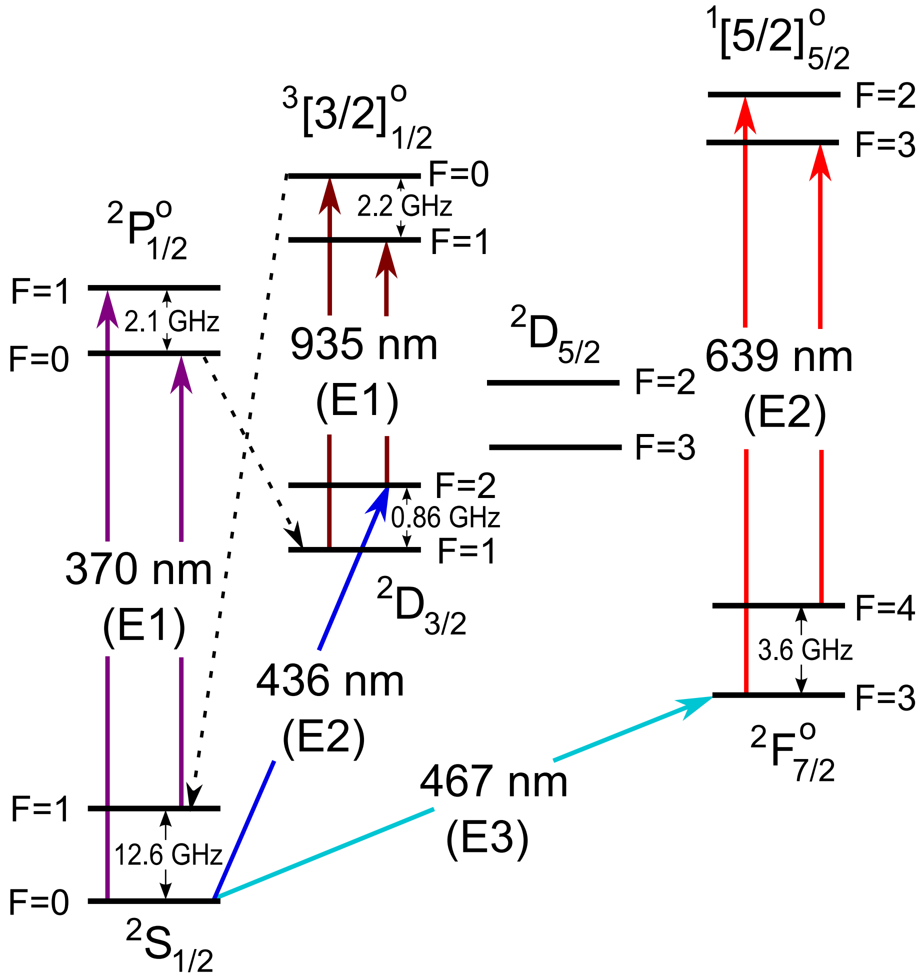

Figure 3.1: Partial term scheme of singly-ionised ytterbium-171 (not to scale). The vacuum wavelengths of all the transitions utilised in this experiment are shown, along with the hyperfine splittings of some important states. A dashed line indicates a spontaneous decay.

Term	**Configuration**	**Energy / cm^{-1}**
$^2S_{1/2}$	$4f^{14}$ 6s	0
$^2F_{7/2}$	$4f^{13}$ ($^2F^{\circ}_{7/2}$) $6s^2$	21419
$^2D_{3/2}$	$4f^{14}$ 5d	22961
$^2D_{5/2}$	$4f^{14}$ 5d	24333
$^2P_{1/2}$	$4f^{14}$ 6p	27062
$^3[3/2]_{1/2}$	$4f^{13}$ ($^2F^{\circ}_{7/2}$) 5d6s (^{3}D)	33654
$^1[5/2]_{5/2}$	$4f^{13}$ ($^2F^{\circ}_{7/2}$) 5d6s (^{1}D)	37078

Table 3.1: Summary of the low-lying energy levels in Yb^+ utilised in this experiment, with their electron configurations and energy above the ground state in cm^{-1} [66].

for ions, as the cooling lasers are not resonant with any transitions in the resulting charged molecule. As mentioned in section 2.6, it has been observed that the dissociation wavelength for YbH^+ lies near 370 nm [59, 60], a wavelength that is present during normal operation and hence any molecules are broken up shortly after formation.

3.4 Laser cooling of $^{171}Yb^+$

The $^2S_{1/2} \rightarrow {}^2P_{1/2}$ transition at 370 nm is a strong, dipole-allowed transition with a natural linewidth of 19.6 MHz [67]. The ion is cooled on the $F = 1 \rightarrow F' = 0$ component of this transition as decay from $^2P_{1/2}$ $(F = 0)$ to $^2S_{1/2}$ $(F = 0)$ is forbidden by the dipole transition selection rules, creating a 'closed' cooling cycle. However, occasionally the $F = 1 \rightarrow F' = 1$ component can be off-resonantly driven. From this state, there is an allowed decay to $^2S_{1/2}$ $(F = 0)$. This state was previously cleared out by using microwaves to mix the two hyperfine components of the ground state. Although simple, this setup has the disadvantage that the ion spends an equal time in each of the two hyperfine levels, thus reducing the maximum 370 nm scatter rate by a factor of two. This setup has since been superseded by modulating the cooling laser at the sum of the hyperfine splittings of the $^2S_{1/2}$ and $^2P_{1/2}$ states. The frequency modulation at 14.7 GHz (= 12.6 GHz + 2.1 GHz) is added as the second order sideband of a resonant electro-optic modulator operating at 7.35 GHz. This sideband can then drive the $F = 0 \rightarrow F' = 1$ component of the 370 nm transition. The ion can then either decay to $^2S_{1/2}$ $(F = 1)$, returning to the cooling cycle, or decay back to $^2S_{1/2}$ $(F = 0)$ from which it is again cleared out. The optimum modulation depth was determined to be 0.35 rad by maximising the detected fluorescence. This corresponds to a fractional power of 2×10^{-4} in the second order sideband.

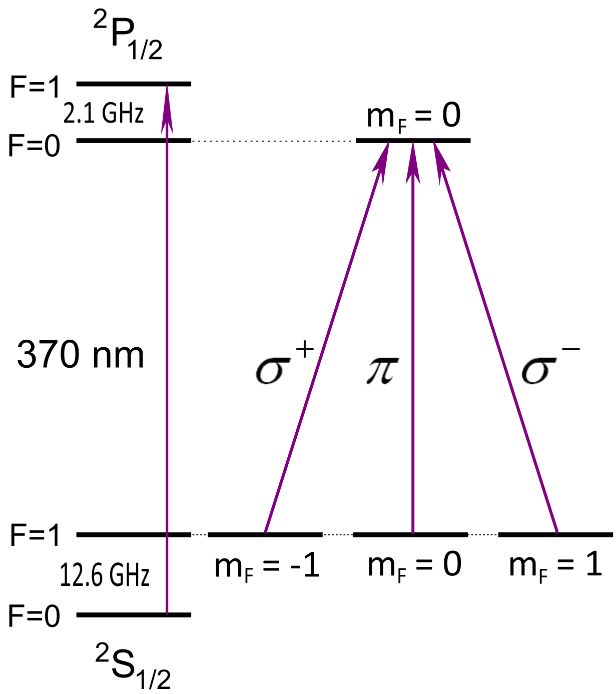

Figure 3.2: Depiction of the various laser frequencies and polarisations required to prevent population trapping in $^2\mathrm{S}_{1/2}$ $(F = 1, m_F = \pm 1)$ and optical pumping to $^2\mathrm{S}_{1/2}$ $(F = 0)$. The three transitions from $^2\mathrm{S}_{1/2}$ $(F = 1)$ are unresolved in the low residual magnetic field of ≈ 10 μT.

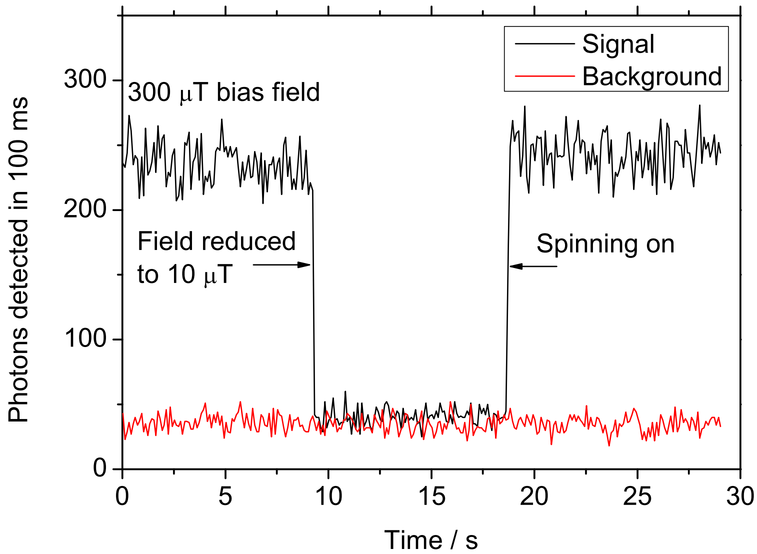

Figure 3.3: Demonstration of the cessation of fluorescence when a bias field of 300 μT is turned off, and resumption when polarisation spinning is turned on. The laser power and detuning was chosen to make the fluorescence in the two cases near-equal for clarity. In general, the fluorescence with polarisation spinning is around a factor of four lower than for a 300 μT field.

3.4.1 Destabilising the dark states

Previously, a single linearly-polarised cooling beam (consisting of the carrier and 14.7 GHz sidebands) was used to cool the ion and a magnetic field of around 300 μT destabilised the dark $m_F = \pm 1$ components of the $F = 1$ ground state by inducing a precession of the ion's dipole moment [68]. The polarisation of the laser relative to the applied magnetic field was set by optimising the detected fluorescence*. If present during the probe phase of the clock cycle, this large field would induce a large and potentially unstable second-order Zeeman shift on the clock transition. It was therefore switched off to leave a small residual field to define the ion's quantisation axis (which is determined using a well-calibrated orthogonal set of field coils driven by stable current supplies). The decay of the field to a suitable level was limited to a few tens of milliseconds by eddy currents in the vacuum chamber, and this constituted a large fraction of the dead time within the clock cycle. The need for switching has since been eliminated by the implementation of a technique involving the simultaneous application of two cooling beams with an angular separation of 45° in their propagation vectors. The polarisation of one beam is rapidly switched between left- and right-circular by a resonant electro-optic modulator operating near 1.3 MHz with a modulation depth of $3\pi/2$, and the second beam has a fixed linear polarisation. This arrangement prevents coherent population trapping in the ground state by introducing a time-dependent relative phase between the three polarisations [68], and is summarised in figure 3.2. The EOM is driven in a similar way to the ion traps (see subsection 2.1.5), using a resonant transformer based on a toroidal iron powder core. A lack of available windows on the vacuum chamber makes it impossible to have the fixed and spun beams meet at a right angle, which is the condition for optimum cooling as the time-evolution of the three polarisation components can be made to be completely orthogonal [68]. In [68], they repeat their calculations for a 63° relative angle between the polarisation of the fixed beam and the propagation vector of the spun beam. In this

*This angle must be oblique, as population trapping occurs if the laser is polarised either in the plane of or perpendicular to the applied magnetic field.

case, they find that the fluorescence is reduced by around a factor of three from optimum conditions. In this experiment the relative angle is limited to 45° and the fluorescence is around a factor of four lower than with a large magnetic field applied. A demonstration of the fluorescence increase in a low bias field is shown in figure 3.3.

3.4.2 Implementation of ECDL at 370 nm

Producing a coherent light source at 370 nm to cool Yb^+ is a technical challenge, owing to the absence of laser diodes at both 739 nm and 370 nm. NPL has traditionally produced this wavelength using an Ar^+-pumped Ti:sapphire laser at 739 nm which is resonantly frequency-doubled in lithium triborate (LiB_3O_5) as shown in figure 3.4. The wavelength of the Ti:sapphire laser is coarsely tuned by a Lyot filter in the laser cavity. An optical diode in the cavity prevents bi-directional operation, and a thin etalon allows the selection of a single longitudinal mode. To prevent mode-hops of the laser, the etalon tilt angle is stabilised to the lasing mode. The etalon is dithered at 6.4 kHz using a piezoelectric transducer (PZT) ring, and the reflection from its surface is focused onto a photodiode. This signal is fed to a lock-in amplifier along with a sample of the drive voltage in order to generate an error signal, which is integrated and fed back to a galvanometer to which the etalon is attached. This locks the mode of the etalon to the relevant laser cavity mode.

The laser frequency is stabilised by using a side-of-fringe lock as shown in figure 3.6a. The error signal is generated by an non-confocal ultra-low-expansivity (ULE) external reference cavity. The lock can be performed using a single photodiode to monitor the transmission of the cavity in the TEM_{00} mode, and then subtracting a fixed reference voltage from this in order to produce an error signal that crosses zero when the laser is tuned to the half-maximum of the transmission fringe (where the sensitivity of the cavity to frequency fluctuations is maximised). However, as the cavity transmission converts frequency noise to noise on the transmitted intensity and vice-versa, any amplitude noise on the laser is interpreted by the photodiode as frequency noise and hence additional

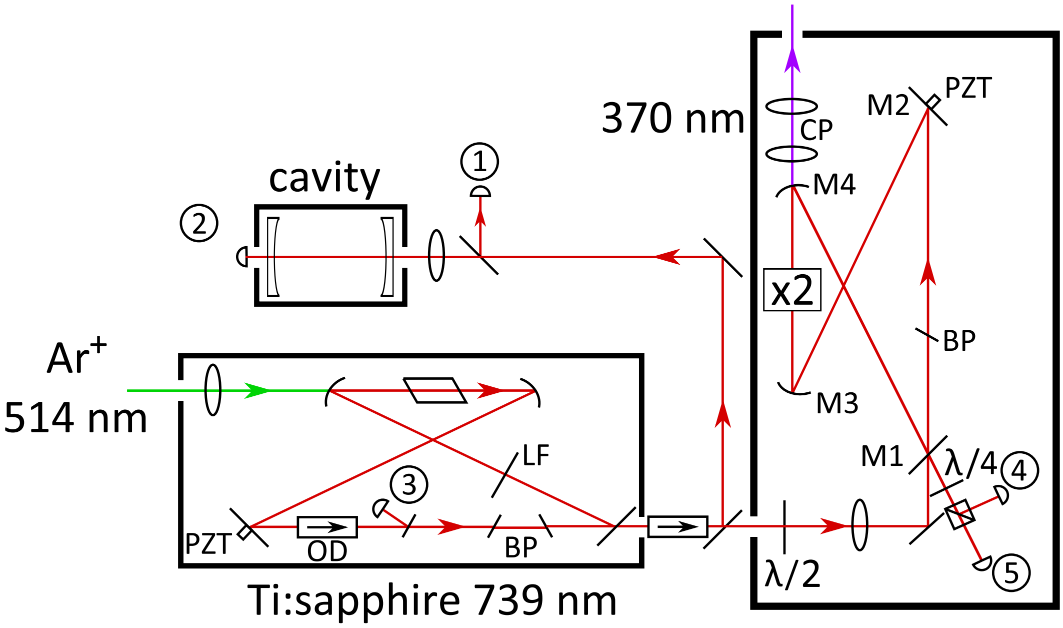

Figure 3.4: Ti:sapphire laser system used previously for laser cooling of $^{171}Yb^{+}$. Photodiodes 1 and 2 are used for the side-of-fringe lock, photodiode 3 is used to monitor the reflection from a thin intracavity etalon, and photodiodes 4 and 5 generate the error signal for the resonant doubling cavity. The laser itself and the doubling cavity are enclosed within foam-lined wooden boxes for acoustic isolation and additional temperature stability. A pair of oppositely-oriented cylindrical lenses corrects for astigmatism after the doubling cavity. LF = Lyot filter, OD = optical diode, BP = Brewster plates, CP = cylindrical pair.

frequency noise can be erroneously added to the laser in order to reduce the noise on the transmitted signal. To reduce this effect, a signal from a second photodiode before the cavity is used instead of the fixed reference voltage. Any amplitude noise is common to both signals, and hence is cancelled by the differencing process. Fast feedback from the servo loop is applied to a PZT to which one of the Ti:sapphire laser cavity mirrors is attached. The correction signal is integrated and used to feed back to two oppositely-oriented Brewster plates, which change the length of the laser cavity without displacing the beam. Once stabilised to the reference cavity, the laser frequency can be tuned by changing the separation of the reference cavity mirrors using electrostrictive actuators*.

The 739 nm light is converted to 369.5 nm by frequency doubling in a lithium triborate (LBO) crystal. To increase the conversion efficiency, the crystal is placed at the focus of a bow-tie build-up cavity. The cavity length is stabilised to the input laser using the Hänsch-Couillard technique [69], with fast and slow feedback applied to a PZT and Brewster plate respectively. Typically around 1 mW of ultraviolet (UV) light is generated from 500 mW of the fundamental. The low conversion efficiency is limited by the low finesse of the buildup cavity.

Overall, this laser system is very complex, requiring a large amount of electronics and four separate servo loops. In addition, the Ar^+ laser is very expensive to run and requires specialist laboratory infrastructure including a three-phase electrical supply and its own separate closed-circuit cooling water supply. The Ti:sapphire laser itself was prone to mode-hops, which often caused the loss of ions from the trap. It was therefore decided to implement a diode-based system to replace this setup, given the recent advances in this technology.

Nichia Corporation manufacture laser diodes at 375 nm for biomedical applications, but unfortunately the incredibly weak temperature tuning of these diodes (0.05 nm/°C) means that groups wishing to tune the lasers to 369.5 nm have had to resort to cooling the diodes to temperatures as low as -40°C [70, 71]. To extend the tuning range further, they

*Electrostrictive actuators are used instead of PZTs as they exhibit lower hysteresis.

have also anti-reflection coated the diode output facet. Both of these are risky processes for such sensitive diodes.

Fortunately, we have been able to purchase a selection of diodes with nominal emission wavelengths near 371 nm. Following the repeated failure of AR coating runs, the diodes are now operated without coating and with the protective can still intact. Feedback from a diffraction grating was achieved by mounting a diode in a commercial* extended cavity. By tuning the grating angle, the wavelength was found to tune smoothly over a range of almost 4 nm (see figure 3.5a). The threshold current was reduced by around 2 mA at the desired wavelength (see figure 3.5b).

The laser was stabilised to a tunable, low-expansivity etalon with a finesse of 150 using the side-of-fringe technique. The two photodiode signals were sent to transimpedance amplifiers before being subtracted. The resulting in-loop noise on the error signal is shown in figure 3.6b. The drift rate of the cavity is around 200 kHz/day and the laser remains in lock for many days at a time. There were initial concerns that the broad amplified spontaneous emission (ASE) plateau of the diode could drive the ion to the $^2\mathrm{P}_{3/2}$ state, thus potentially shelving the ion in the $^2\mathrm{D}_{5/2}$ state during the cooling cycle as has been observed in Ca^+ ions [72]. However, one major difference between the two species is that the transitions to the two P-states are separated by around 40 nm in Yb^+, as opposed to only 4 nm in Ca^+. There was some residual concern than the ASE could interact with the $F = 0 \rightarrow F' = 1$ component of the 370 nm transition, which would affect any attempted optical pumping to $^2\mathrm{S}_{1/2}$ $(F = 0)$ (see section 3.6). However, removal of the 14.7 GHz sideband still caused a loss of fluorescence, indicating that any interaction of the ASE with the other hyperfine states remains very weak.

The 370 nm transition lineshape resolved by this new cooling laser is shown in figure 3.7. The observed broadening from the natural half-linewidth of 10 MHz is likely due to a combination of the residual linewidth of the laser, imperfect micromotion compensation, Zeeman splitting in the residual bias field, and power broadening.

*Toptica DL 100 'Pro' Design.

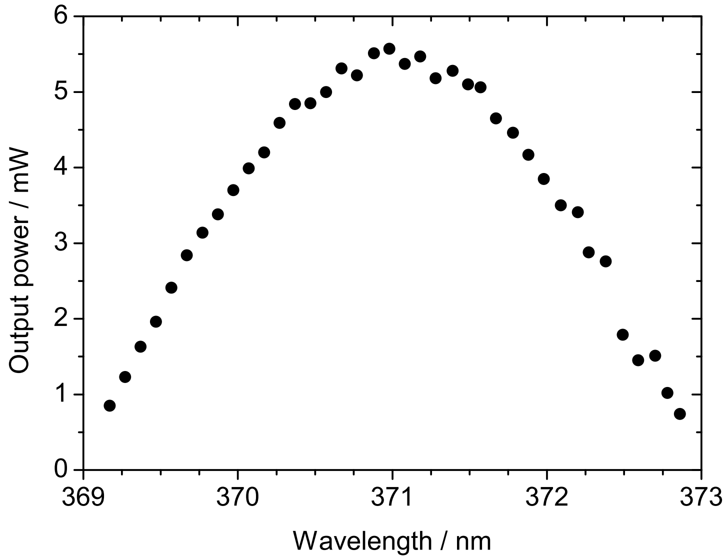

(a)

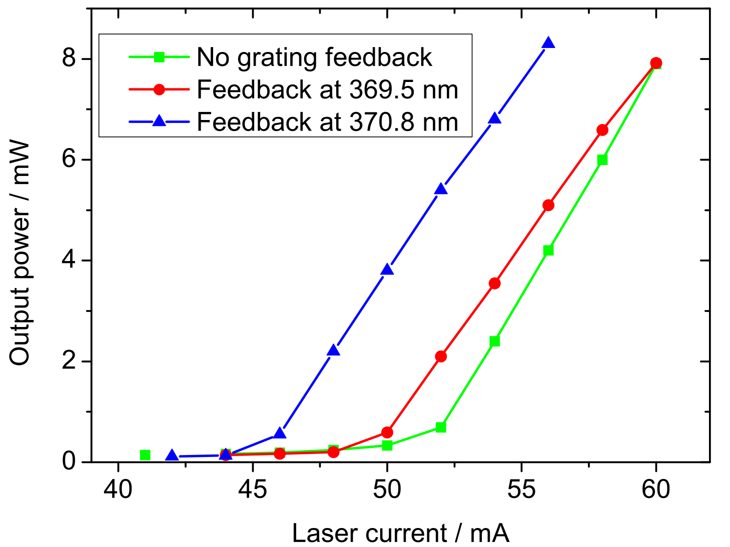

(b)

Figure 3.5: (a) Output power of the UV ECDL as a function of wavelength when the diffraction grating angle is tuned. (b) Power vs current curve for the cases of no feedback from the diffraction grating, with the laser operating at the desired wavelength of 369.5 nm, and with the laser operating near its centre wavelength of 371 nm. All data was taken with the temperature of the diode stabilised at 20°C. Above threshold, the gradient is nearly 1 mW/mA at 369.5 nm.

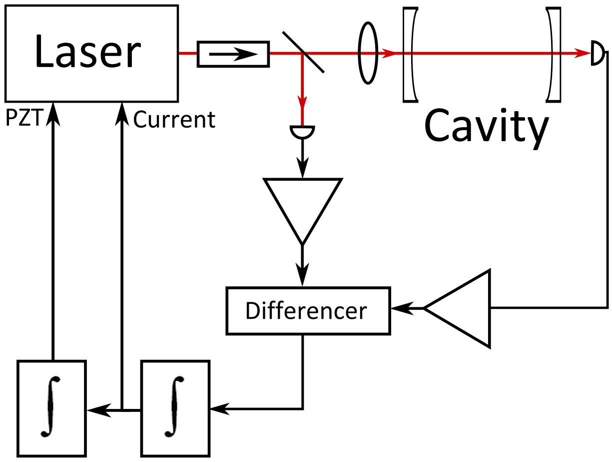

(a)

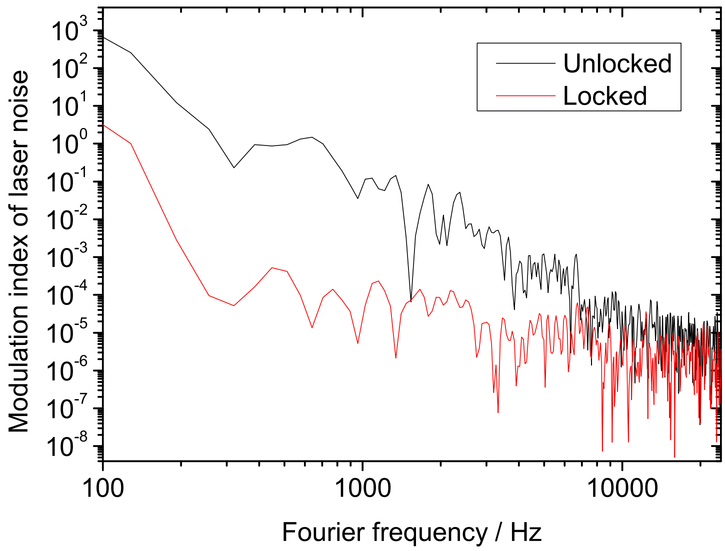

(b)

Figure 3.6: (a) Depiction of the side-of-fringe locking technique. (b) In-loop noise on the side-of-fringe error signal for the stabilisation of the 370 nm diode laser to an optical cavity. The lock bandwidth is primarily limited by the low transimpedance amplifier bandwidth of around 10 kHz.

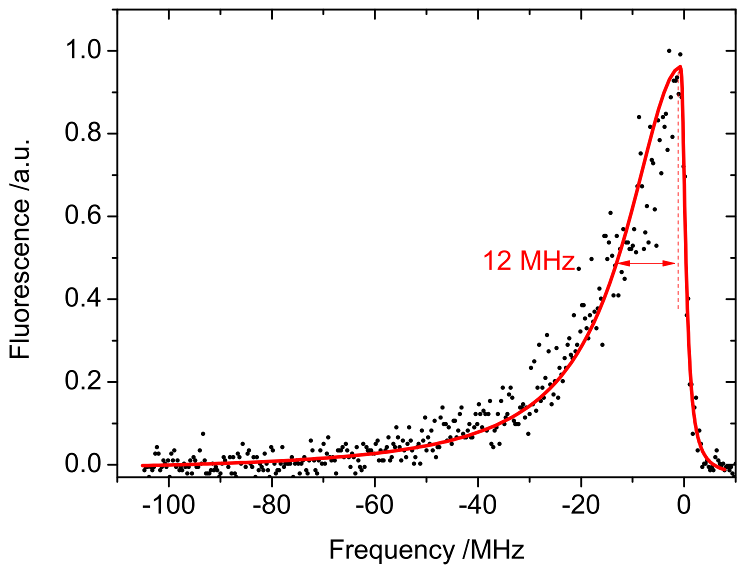

Figure 3.7: 370 nm transition as resolved by the new diode laser. The observed broadening from the natural half-linewidth of 10 MHz is explained in the text.

3.4.3 935 nm repumper lasers

Unfortunately, the 370 nm transition is not a completely closed cycle. The $^2P_{1/2}$ state can decay to $^2D_{3/2}$ with a branching ratio of around 5×10^{-3} [73, 63, 74]. The $^2D_{3/2}$ state can only decay via electric quadrupole radiation (a fact that we exploit in its role in the 436 nm clock transition, as will be mentioned later) and so therefore has a long natural lifetime of 52.7 ms before decay back to $^2S_{1/2}$. The resulting low level of fluorescence would be very hard to detect (especially with an overall detection efficiency of $\approx 0.1\%$) and the low number of scattered photons will have an impact on the efficiency of the laser cooling. To prevent 'shelving' of the ion in this state, the ion is also illuminated with light at 935 nm from a commercial* extended cavity diode laser (ECDL) stabilised to another tuneable optical cavity using the Pound-Drever-Hall (PDH) method (see subsection 4.2.2). This drives the ion to the $^3[3/2]_{1/2}$ state, which has a strong dipole-allowed decay to the ground state with a branching ratio of 98.2% [75] and natural lifetime of 37.7 ns [76]. As this repumper laser must be present to observe fluorescence, care must be taken that it does

*New Focus Vortex

not interact with the other hyperfine component of the $^2\mathrm{D}_{3/2}$ state. This state is part of the 436 nm clock transition, and to detect shelving of the ion in this state after a probe pulse requires an attempted detection of fluorescence*. The intensity and spectral purity of the 935 nm ($F = 1 \rightarrow F' = 0$) laser are therefore critical, as any broadband emission (e.g. from amplified spontaneous emission) can act to shorten the lifetime of the $F = 2$ clock state during the detection phase. This could lead to an increase in the error rate in the readout of the ion's state. For this reason, it is the transmission of the optical cavity that is delivered to the ion. The 3 MHz linewidth of the cavity acts to filter out most of the unwanted frequency noise on the transmitted beam. Some measurements of the clock state lifetime are shown in figure 3.8. It is clear that filtering greatly reduces the rate at which the laser off-resonantly clears out the clock state. The observed lifetime of of 51.8(17) ms agrees well with a natural lifetime of 52.7 ms [63].

With such a low intensity in this repumper laser, the ion is optically pumped to the $^2\mathrm{D}_{3/2}$ ($F = 2$) state during cooling (although it is much more rarely populated than ($F = 1$) as it can only be populated by a spontaneous decay from $^2\mathrm{P}_{1/2}$ ($F = 1$) [77]). The ion is observed to be shelved in this state around five times per second when the cooling transition is saturated. Therefore a second commercial† ECDL operating at 935 nm is tuned to the ($F = 2 \rightarrow F' = 1$) component of the 935 nm transition during laser cooling. This laser is absent during detection of the ion's state for the same reasons detailed above, but can also be used for rapid repumping of the ion from the $^2\mathrm{D}_{3/2}$ ($F = 2$) state after a successful probe on the 436 nm transition. Due to its much higher intensity and low drift rate, this laser does not need to be actively stabilised. With both 935 nm lasers present, the ion spends around 97% of its time cycling around the S $\rightarrow$ P transition [54].

From this point on, the convention 935 nm ($F = 1$) will be used to specify the laser that clears out the $F = 1$ component of the $^2\mathrm{D}_{3/2}$ state.

It is important to note that coherent population trapping also occurs on the 935 nm transitions in low magnetic fields. This limits the minimum operating field to around

*See subsection 6.1.2 for further information on the detection of quantum jumps.
†New Focus Vortex

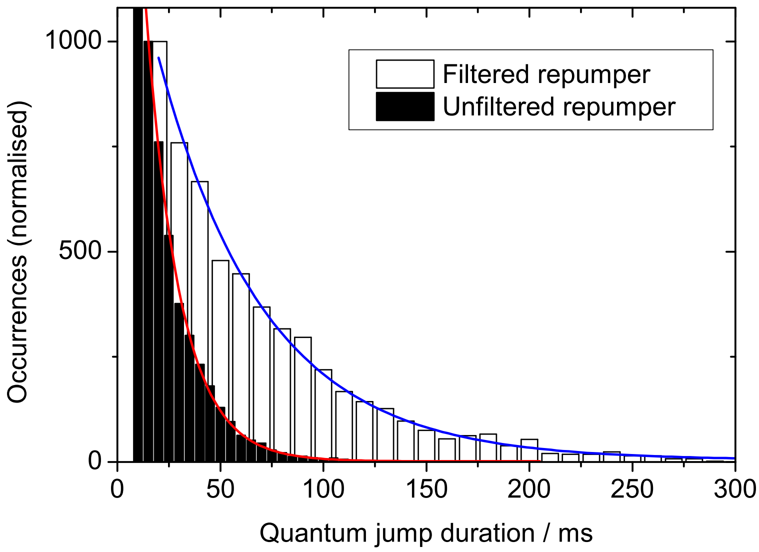

Figure 3.8: Measurements of the lifetime of the $^2D_{3/2}$ ($F = 2$) state with the 935 nm ($F = 2 \rightarrow F' = 1$) repumper laser not present. It is clear that an unfiltered 935 nm ($F = 1 \rightarrow F' = 0$) laser greatly reduces the lifetime. Exponential fits to the data show that the observed 1/e lifetime is extended from 16.5(1) ms to 51.8(17) ms when the repumper laser is filtered through an optical cavity.

10 μT, below which fluorescence rates start to fall.

3.4.4 639 nm repumper lasers

Once the 467 nm transition has been driven, the ion is trapped in the $^2F_{7/2}$ ($F = 3$) state, which has a natural lifetime of 6 years. Hence, if we want to drive the transition more than once every 6 years, we must either re-load the trap or repump the ion to the ground state. Repumping is currently performed using a commercial* ECDL to drive an electric quadrupole transition at 639 nm, exciting the ion to the $^1[5/2]_{5/2}$ ($F = 2$) state from which there are dipole-allowed decays to $^2D_{3/2}$ and $^2D_{5/2}$. The $^2D_{5/2}$ state can decay to the ground state via an electric quadrupole transition with a natural lifetime of around 7 ms [63]. However, there exists a dipole-allowed decay to $^2F_{7/2}$ with a branching ratio of 83% [78], and hence the ion can make several trips around this cycle before successfully decaying to the ground state. Due to the weakness of the 639 nm transition,

*New Focus Vortex

a narrow laser is required with high power, and must be accurately tuned to the transition frequency. However, if the laser is too narrow then it will not be able to clear out all of the Zeeman sublevels of the lower state which may be populated by spontaneous decay. Hence the laser is stabilised using the PDH method to a relatively low finesse, tunable optical cavity. In addition, the $^2F_{7/2}$ ($F = 4$) state can be populated by spontaneous decay, and so a second repumper wavelength is required to drive the ion to the $^1[5/2]_{5/2}$ ($F = 3$) state. This is provided by a second commercial* ECDL, which is offset-locked to the first (see figure 3.9). The lasers are overlapped on a fast photodiode, and the beat signal is mixed down to 30 MHz with a microwave synthesiser operating at 3.89 GHz. This signal is fed to a phase-frequency comparator operating at 30 MHz to generate the error signal for the lock. This extra stage is required as the unlocked laser is too noisy for a simple phase lock.

In addition, the $^2F_{7/2}$ state can be populated by a collisional de-excitation [78] of the ion from $^2D_{3/2}$ to $^2F_{7/2}$, or from $^2P_{1/2}$ to $^2D_{5/2}$, from which it can decay to $^2F_{7/2}$. This is observed to occur a few times per day.

Recently, PTB have implemented an alternative repumping scheme which drives the ion to $^1[3/2]_{3/2}$ via an electric quadrupole transition at 760 nm [17]. The upper state of this transition has a strong dipole-allowed decay to the ground state and should allow shorter repumping times after the 467 nm transition has been driven. With the relatively long detection times currently required in this experiment, the small reduction in dead time is outweighed by the prohibitive cost.

3.5 Automatic re-locking

To minimise the downtime associated with lasers falling out of lock, it is desirable to design a system that can automatically re-lock the lasers after a disturbance on the optical table. This has been initally tested on the 370 nm cooling laser as the side-of-fringe lock has

*New Focus Vortex

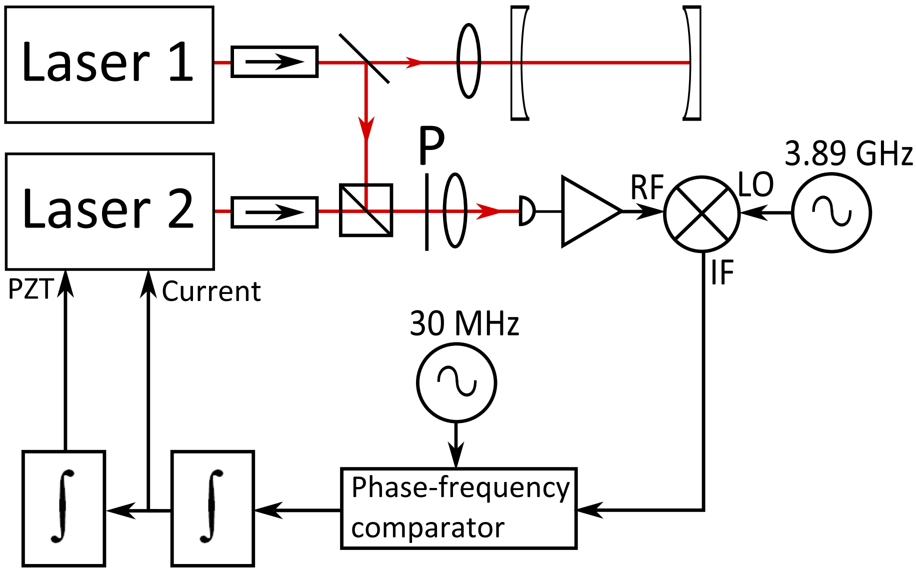

Figure 3.9: Depiction of the technique used to stabilise the frequency of a second 639 nm laser to the first. P = polariser.

proved easier to restore than a Pound-Drever-Hall lock.

Instead of using an analogue servo for the 'slow' feedback to the laser piezo, the output of the fast integrator (used to lock the laser to the optical cavity via its current modulation input) is continually monitored by an analogue input of a National Instruments NI-6008 DAQ card. A LabVIEW virtual instrument (VI) is then used to integrate this signal, and the result is fed back to the laser piezo using an analogue output of the same card. This forces the output of the current lock to zero. The error signal for the fast lock is continually monitored by a second analogue input on the card. If it falls outside certain pre-set bounds, the laser is deemed to be out of lock. A digital output on the card is then used to trigger a relay which shorts the input to the fast integrator, thus resetting the lock. The last known 'good' value of the voltage applied to the piezo is restored, and the lock re-engaged. The re-lock procedure is summarised in figure 3.10, and the front panel of the LabVIEW VI is shown in figure 3.11. This rather crude method of restoring the lock has proved to be extremely reliable, provided that the laser does not also mode-hop.

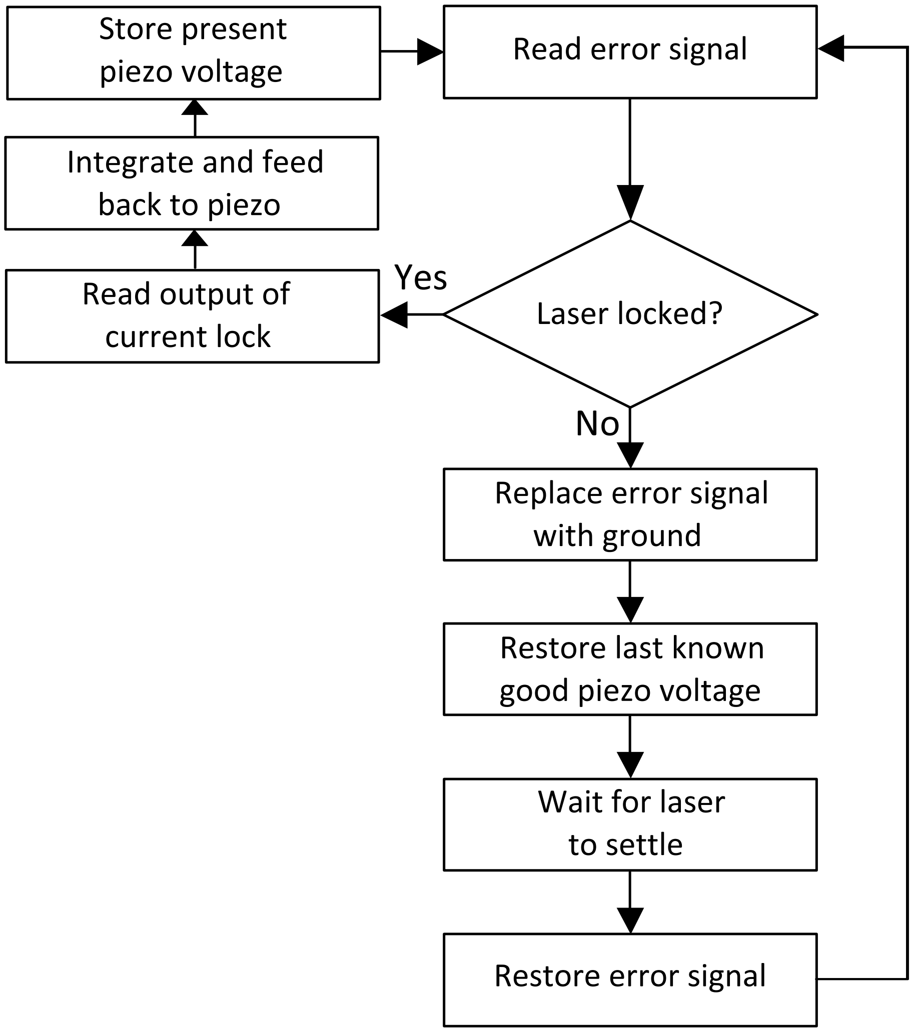

Figure 3.10: Flowchart showing the operation of the monitoring program that controls the laser piezo and re-locks the laser after a disturbance.

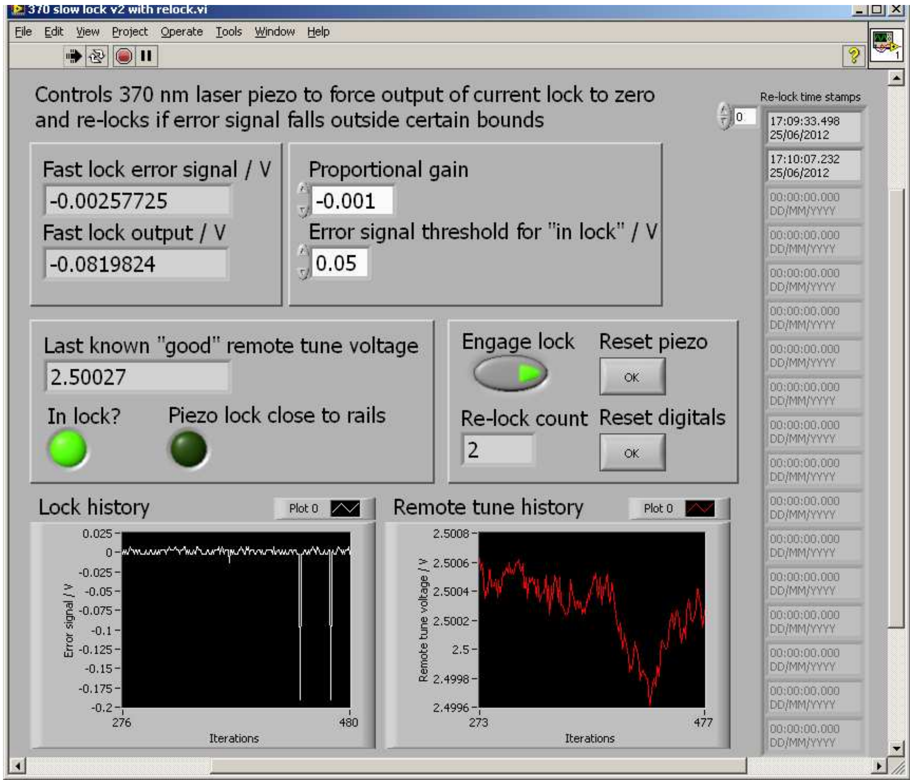

Figure 3.11: Screen capture of the LabVIEW VI used to monitor and relock the 370 nm cooling laser. The bottom left chart shows the in-loop error signal as a function of time, and shows two disturbances that have knocked the laser out of lock. The laser is re-locked and the disturbances timestamped.

3.6 State preparation

Before probing the clock transitions, the ion must be prepared in the $^2S_{1/2}$ ($F = 0$) state by driving the $F = 1 \rightarrow F' = 1$ component of the 370 nm transition. Light at this frequency is generated by modulating the 370 nm ($F = 1 \rightarrow F' = 0$) light at the hyperfine splitting frequency of the $^2P_{1/2}$ level using a second resonant EOM* operating at 2.1 GHz. The modulation depth of 2.5 rad is strong enough that the carrier is almost completely depleted, and so the ion is rapidly pumped to the desired state via $^2P_{1/2}$ ($F = 1$). State preparation has been observed for pulses as short as 2 ms, which is the shortest pulse duration allowed by the response of the mechanical shutter in the beam path. Previously [79], state initialisation was achieved by removal of the 12.6 GHz microwaves and waiting for an off-resonant excitation of the $F = 1 \rightarrow F' = 1$ transition which could take up to 20 ms.

3.7 Optical setup

The completed optical sestup is shown in figure 3.12. The two 639 nm lasers are combined on a polarising beamsplitter, and the two 935 nm lasers are combined similarly. The two wavelengths are then combined on a dichroic beamsplitter and coupled into the same optical fibre. The output of the fibre, a half waveplate and an achromatic focusing lens are all mounted on a 3D translation stage. The lasers are focused to a waist of 75 μm at the ion and are delivered in the horizontal plane.

The 399 nm photoionisation and 467 nm clock laser beams are combined on a dichroic beamsplitter. The 436 nm clock laser and a fraction of the 14.7 GHz-modulated 370 nm light are made to be collinear but can be switched between by using a mirror on a flip mount. All of the UV and blue beams are then combined using a polarising Glan-Thompson prism and coupled into the same polarisation-maintaining fibre. There is another fibre / achromatic lens arrangement for this fibre, and the lasers are focused to

*New Focus 4431.

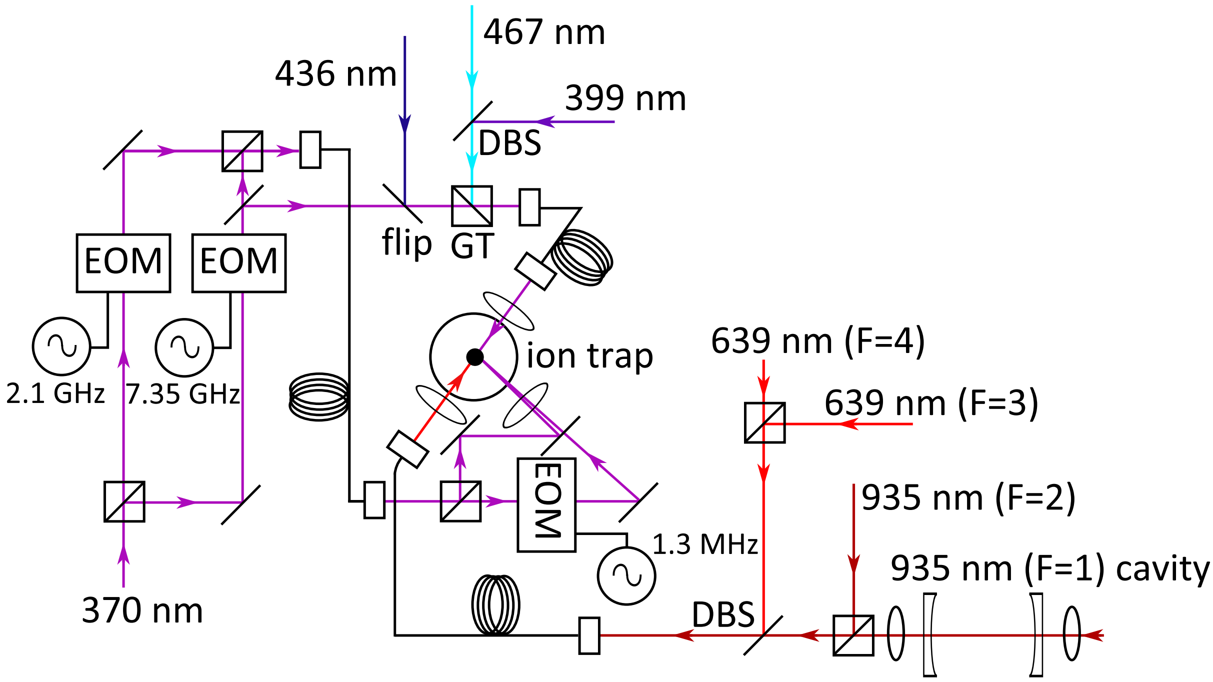

Figure 3.12: Depiction of the experimental setup, including all laser wavelengths applied to the ion. The two 370 nm beams entering the trap from the lower right as depicted are oppositely tipped out of the horizontal plane by 22.5°. DBS: dichroic beamsplitter. GT: Glan-Thompson polariser.

a waist of 5 μm and delivered in the horizontal plane.

There are two further beam paths for the 370 nm light, containing both modulation frequencies. The light is split into two paths, one passes through an EOM for polarisation rotation and the other is linearly polarised. The rotated light is delivered from below, tipped out of the horizontal plane by 22.5°. The linearly polarised light is 22.5° out of the plane in the opposite direction.

3.8 Summary

This chapter has detailed the various laser systems required for cooling, repumping and state preparation of single $^{171}Yb^{+}$ ions. In the following chapter, details are given about the two remaining wavelengths required in order to operate the system as an optical frequency standard: the two clock lasers at 871 and 934 nm. The challenges involved in their production and stabilisation are dicussed, along with their conversion to the probe wavelengths of 436 and 467 nm respectively.

Chapter 4

Ultrastable lasers

Free-running diode lasers have emission linewidths on the order of 1 MHz. Feedback comes from internal reflections within the diode structure, or from the output facet of the can in which the diode is enclosed. The resulting linewidth is described by the Schawlow-Townes equation [80]:

$$\Delta\nu = \frac{\pi h\nu}{P_{\text{out}}}(\Delta\nu_C)^2 \tag{4.1}$$

where $\Delta\nu$ is the laser linewidth, h is Planck's constant, ν is the laser frequency, P_{out} is the power coupled out of the laser cavity, and $\Delta\nu_C$ is the linewidth of the cavity defining the lasing mode. For semiconductor lasers, this expression is multiplied by a linewidth broadening factor of $(1+\alpha^2)$ to account for the coupling between intensity and phase noise ($4 < \alpha < 8$ [81, 82]). The laser cavity linewidth is given by [81]:

$$\Delta\nu_C = \frac{c}{2\pi L_{\text{eff}}}(\kappa L_D - \ln\sqrt{R_1 R_2}) \tag{4.2}$$

where c is the speed of light, L_{eff} is the effective length of the laser cavity, L_D is the length of the lasing medium, κ is the transmission of the lasing medium, and R_1 and R_2 are the reflectivities of the two surfaces defining the cavity. Since the linewidth of the laser cavity is inversely proportional to its length, the laser linewidth can be passively reduced by retro-reflecting a fraction of the laser power back into the diode from a more distant location. Such a retro-reflection can be realised using a diffraction grating, oriented such that the first order diffraction is coupled back into the diode. This arrangement is referred to as an external cavity diode laser (ECDL) and often suffers from competition between the modes of the diode and the modes of the external cavity. The diode can be anti-reflection (AR) coated to improve its tuneability by suppressing reflections from the diode facet, leaving the lasing mode defined only by the length of the external cavity. For a typical cavity length of 25 mm compared to a laser diode length of 0.25 mm, the contribution of the cavity length to the laser linewidth is reduced by a factor of 10^4. Technical noise such as ripple on the diode current supply or mechanical instability on the cavity length can then form the limit to the linewidth, and so attention must be paid to minimise these contributions. With careful design, free-running laser linewidths of below 100 kHz have been realised [83].

Linewidths on this order are acceptable for driving broad transitions such as those used for laser cooling, but the frequency of the laser is subject to long-term drift from temperature fluctuations and changes in the diode lasing medium. To eliminate this drift, the frequency of the laser can be referenced to a more stable anchor and an active feedback loop employed to maintain this frequency. This is referred to as laser locking. The reference oscillator often takes the form of an atomic absorption, an optical cavity, or another laser.

The most extreme requirements for narrow and stable lasers come from the optical clock transitions. These transitions often have natural linewidths substantially narrower than 1 Hz, and so the spectral intensity of the laser in the frequency range of the transition will be extremely small. This makes the transitions extremely difficult, if not impossible, to drive. Two options exist, the first is to increase the intensity of the probe laser. This is not always feasible, and does not improve the resolution with which the transition is probed. In the case of very high optical intensities, it will also introduce an unwanted ac Stark shift on the transition frequency*. The second option is to reduce the linewidth of the probe laser, which will increase the spectral intensity and improve the resolution of the probe. In order to resolve transition linewidths at the Hz-level, it is necessary to produce a probe laser with a linewidth below that level on the timescale of the interrogation sequence. This requires a reference cavity with fractional instability at the 10^{-15} level, which corresponds to a length stability on the order of 0.1 fm on a 10 cm cavity. Reference oscillators that achieve these extremely high stabilities are discussed in section 4.1. Locking the laser to any form of reference requires the derivation of an 'error signal', which can be used to feed back to the laser to correct any fluctuations in its frequency and the techniques involved are outlined in section 4.2. Details of the constructed laser systems and their performance are given in section 4.3.

4.1 Clock laser reference cavities

The realisation of ultrastable cavities is based around reducing the sensitivity of the system to external perturbations, in particular fluctuations in the number density of gas atoms in the cavity bore and thermal and vibrational noise.

4.1.1 Vibrational sensitivity

Vibrations on the floor or optical table can be transferred to the cavity mirrors and spacer, causing a change in the optical path length. A typical sensitivity of the length of an optical cavity to acceleration is $3 \times 10^{-9}/g$ [84] where g is the acceleration due to gravity at the surface of the Earth in ms^{-1}. It is possible to damp out these vibrations actively or passively, but even greater suppression can be achieved by careful design. The cavities used to stabilise the clock lasers consist of a cylindrical spacer with sections cut out at right-angles [84], as shown in figure 4.1. When mounted correctly, the net response of the cavity to vibrations is close to zero along its optical axis (the acceleration response

*This effect is especially important for the 467 nm clock transition, as will be discussed in chapter 7.

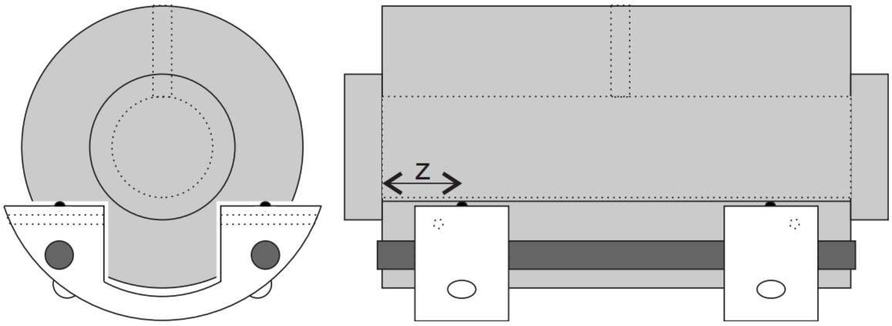

Figure 4.1: Cavity design used to stabilise the clock lasers [85]. The cut-out design is visible. Correct positioning z of the Viton balls used to support the cavity leads to a null in its vibrational response.

of a spacer of this design has been measured to be on the order of $10^{-12}/g$ in the vertical direction and $10^{-10}/g$ in the horizontal plane [84]). The spacer is mounted on four Viton balls, which approximate point contacts and also provide a degree of thermal isolation from the inner heat shield (see subsection 4.1.3). In addition, the optical table is mounted on pneumatic legs for passive isolation from the floor, and an active vibration isolation platform* (AVI) is used to isolate the cavity from the table†.

It is worth noting that alternative cavity designs exist that reduce the sensitivity of the cavity to vertical accelerations. Mounting the cavity vertically means that the response of the two mirrors to any vibration should be equal, leading to no net change in the cavity length. However, this mounting system becomes impractical for longer cavities and hence most cavities of more than 10 cm length are mounted horizontally.

4.1.2 Pressure sensitivity

Changes in the local air pressure P change the optical path length inside the cavity via the refractive index n of the air in the cavity bore. The increase in path length is given by:

$$\frac{\Delta\nu}{\nu} = \frac{\mathrm{d}n}{\mathrm{d}P}\Delta P \tag{4.3}$$

*Table Stable AVI350M.

†The AVI provides an attenuation of 35 dB for vibration frequencies between 10 and 200 Hz but does not provide any isolation at frequencies below $\sim$2 Hz.

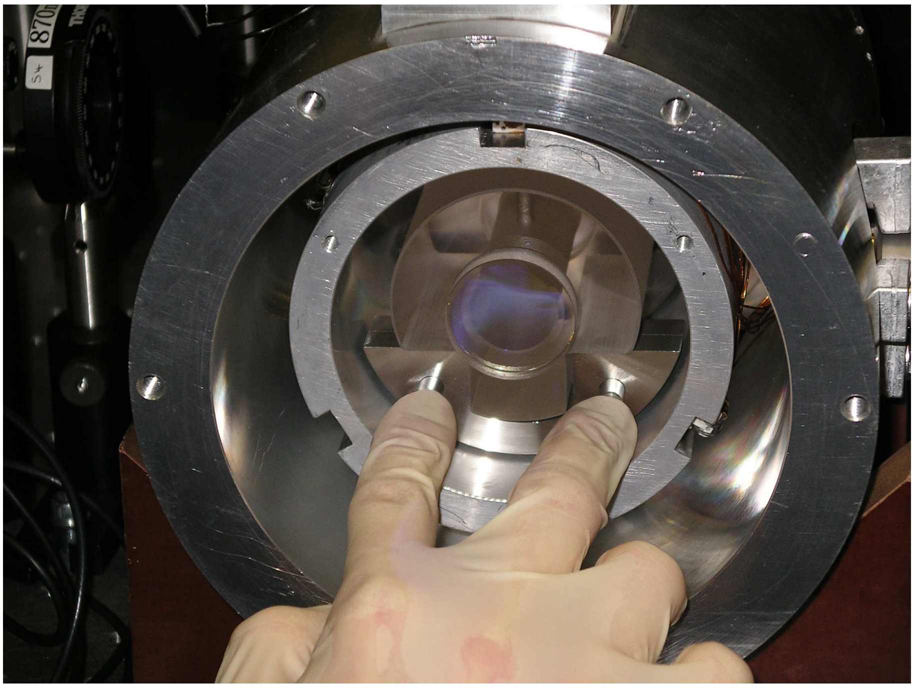

Figure 4.2: Cut-out cavity being mounted into its vacuum chamber. The inner heat shield is visible.

with $\mathrm{d}n/\mathrm{d}P$ roughly equal to 3×10^{-11} Torr^{-1} for dry air at room temperature [86]. The effect of fluctuations, whether thermal in nature or simply a disturbance, is minimised by mounting the cavities in vacuum. The vacuum chambers are aluminium cylinders, with indium seals on the large end flanges and copper gaskets on the smaller knife-edge flanges. After a week-long bake at 60°C, a 2 l/s ion pump is able to maintain a pressure of around 10^{-7} Torr. The maximum temperature for the bake was limited by the damage threshold of the Peltier coolers inside the vacuum system. At this pressure, a fluctuation of 10% would produce a frequency excursion at the level of 3×10^{-17}.

It is worth noting that for a completely sealed vacuum system, pressure fluctuations (e.g. due to temperature variations) do not lead to fluctuations in the number density of atoms in the cavity bore and hence the refractive index is unaffected. In most cases, changes in pressure are attributable to outgassing of elements inside the vacuum system or leakage through seals, and hence the atom number density is not conserved.

Mounting the cavity in vacuum has the added advantage that it can be cooled below the dew point of the laboratory without fear of condensation. The reasons for cooling are outlined in the following subsection.

4.1.3 Thermal sensitivity

The main contribution to the sensitivity of the cavity mode to any thermal perturbation is proportional to the coefficient of thermal expansivity (CTE), α, of the material used for the mirror spacer:

$$\frac{\Delta\nu}{\nu} = \alpha\Delta T \tag{4.4}$$

where $\Delta\nu/\nu$ is the fractional frequency change due to a temperature perturbation of ΔT. Silica glass typically has a CTE of around 6×10^{-7}/K, which would require a temperature stability of around 2 nK across the whole spacer in order to realise a fractional frequency instability of 10^{-15}. There are, however, some materials for which the thermal expansivity is temperature dependent ($\alpha(T)$), and in some cases this can even cross zero. One of these materials is TiO_2-doped silica glass, also known as ultra-low expansivity (ULE) glass. ULE has a CTE of around $1.5\times10^{-9}(T-T_C)/\mathrm{K}^2$, with the dopant level chosen to produce the zero crossing T_C near room temperature.

It is necessary to introduce a combination of both active and passive temperature control to maintain the cavity as close as possible to its CTE zero-crossing. The cavity is shielded from the outside environment by an aluminium heat shield inside a large aluminium vacuum chamber. This two-layer system, coupled with the evacuation of the vacuum chamber and the isolation of the cavity from the heat shield using Viton balls, leads to a passive low-pass filter for external temperature fluctuations with a time constant of several hours. The temperature of the inner shield is monitored using a Wheatstone bridge (see figure 4.3), which generates an error signal based on the impedance imbalance between its two arms. An ac bridge is used as it is insensitive to the low frequency $1/f$ noise of the circuit, and does not suffer from the creation of a thermal emf in the thermistor which allows more accurate measurements of the bridge imbalance. The majority of the bridge circuit is enclosed within the vacuum chamber to ensure the best possible thermal stability for the fixed resistors. An ac voltage is applied across the terminals, and the output voltage is fed into a lock-in amplifier with the drive voltage as the reference. Any temperature fluctuations are detected by a thermistor in good thermal contact with the inner heat shield. This results in a change in the output voltage, which is converted to an error signal by the lock-in amplifier. This signal is fed to an integrator, which produces a correction signal. This signal is applied to the base of a transistor with a 15 V, 3 A supply on its collector. The emitter is connected to a set of Peltier thermoelectric coolers, which actively pump heat out of the inner heat shield into the outer vacuum chamber. Heat is then removed from the system by a further set of thermoelectric coolers between the vacuum chamber and a large heat sink, with the error signal for this control provided by a dc Wheatstone bridge.

It has been suggested that radiative transfer can cause heating of the spacer and mirrors [87], and hence the inner heat shield is designed to block most of the line-of-sight between the vacuum chamber and the cavity. However, it is impossible to screen the cavity totally, especially along its optical axis (as a clear line-of-sight must be available for the laser).

With a good choice of gain settings, temperature stabilities of better than 50 μK are achieved on the inner heat shield for averaging times of ten seconds to several hours (see figure 4.4).

In order to locate the CTE zero-crossing, it is necessary to compare the frequency of a laser locked to the cavity to a fixed frequency reference. In this case, the pre-stabilised laser was then locked another (reference) cavity by feeding back to an acousto-optic modulator that bridges the frequency gap between the nearest TEM_{00} modes of the two cavities. The temperature of the test cavity is swept by changing the temperature set-point to be far below the suspected zero-crossing, whilst logging both the resistance of a monitor thermistor on the inner heat shield and the frequency offset between the modes of the two cavities. The end point for the sweep is chosen to be far from the suspected zero-crossing to prevent the temperature control servo damping the temperature sweep in

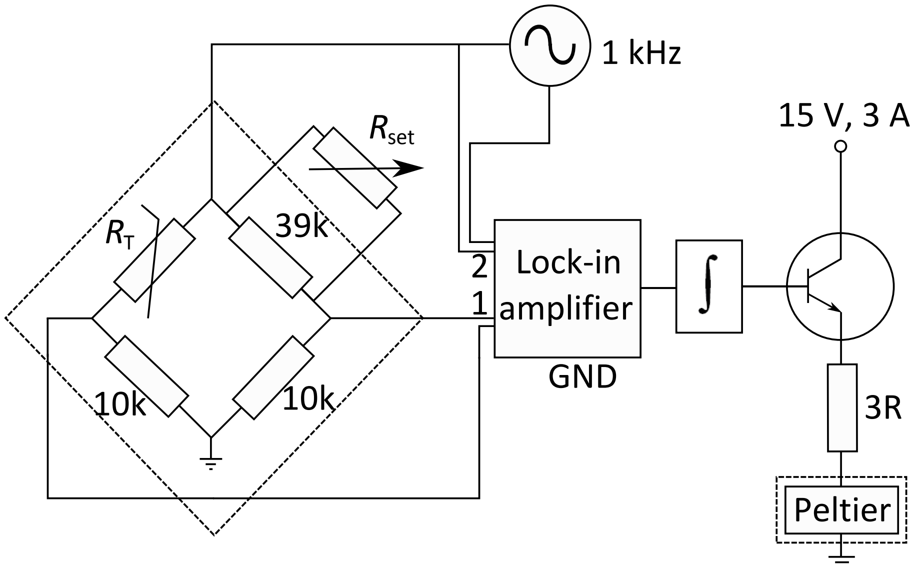

Figure 4.3: ac Wheatstone bridge used to stabilise the heat shield internal to the vacuum chamber. The cavity temperature is monitored using a 5.6 kΩ thermistor, R_T. R_{set} can be changed to determine the temperature set point, at which $R_T = (1/39\text{k} + 1/R_{set})^{-1}$. The lock-in time constant is set to 1 s. Dashed lines indicate the portion of the circuit internal to the vacuum chamber. A 3 Ω power resistor is placed on the transistor emitter to limit the maximum output current, thus preventing any outgassing by the Peltier when the control is first switched on.

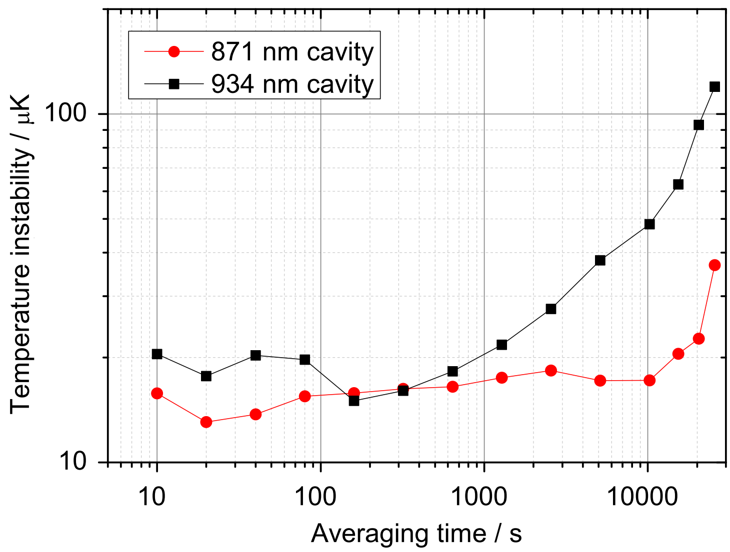

Figure 4.4: Temperature instability (in μK) of the inner heat shields around the two clock laser reference cavities, once under active two-layer control.

the region of interest. Once the cavity temperature has settled, it is then swept upwards to the other side of the zero-crossing.

The zero-crossing cannot be directly determined from the two offset frequency versus temperature curves as the damping due to the passive isolation causes a temperature lag between the heat shield and the cavity. This is manifested as an offset in temperature between the two curves. A simple heat exchange model must be applied to the data:

$$\frac{d}{dt}T_{\text{cav}}(t) = c(T_{\text{hs}}(t) - T_{\text{cav}}(t)) \tag{4.5}$$

where T_{cav} and T_{hs} are the temperatures of the cavity and the heat shield respectively, and c is the inverse of the time constant for the heat transfer. By multiplying both sides by an integrating factor of e^{ct}, this equation can be integrated between the limits $\infty < t < \tau$ to give an expression for $T_{\text{cav}}(t)$:

$$T_{\text{cav}}(\tau) = c\int_{-\infty}^{\tau} e^{c(t-\tau)}T_{\text{hs}}(t)dt. \tag{4.6}$$

Assuming that the temperature of the heat shield has settled before the temperature sweep begins (i.e. $\frac{d}{dt}T_{\text{hs}}(-\infty < t < 0) = 0$), this simplifies to:

$$T_{\text{cav}}(\tau) = T_{\text{hs}}(0)e^{-c\tau}\int_{0}^{\tau} e^{c(t-\tau)}T_{\text{hs}}(t)dt \tag{4.7}$$

The value of c can be iterated until the turning points on the heating and cooling curves are made to be coincident [88], as shown in figure 4.5a. The best reconciliation came with a time constant of 7.3 hours between the inner heat shield and the cavity. The resulting curve can then be differentiated to give the CTE as a function of temperature (as shown in figure 4.5b), and the cavity temperature set appropriately. Due to the rather simplistic nature of this model (including the absence of thermal gradients that might exist within the cavity), the two curves do not give the exact same CTE response to temperature. It is not possible to get closer to the zero-crossing by slightly adjusting the cavity temperature and monitoring changes in the offset frequency as the frequency response is now dominated by the much faster response of the mirror coatings to the temperature change [87], obscuring the response of the spacer.

Once under control at a temperature close to T_C, the dominant source of long-term noise on the cavity frequency is the slow isothermal 'creep' associated with the aging of the ULE material. The drift rate for both cavities is near to 60 mHz/s, corresponding to a fractional drift of 2×10^{-16}/s.

4.1.4 Thermal noise

Even after all of the listed design and environmental control considerations have been implemented, it is still not possible to reduce the length instability to an arbitrary level. The fundamental limit to the stability of a cavity is the thermal noise (Brownian motion) of the mirror substrate. The amplitude of motion can be reduced by cooling the cavity to cryogenic temperatures, but this would leave the cavity far from its CTE zero-crossing. The contribution of individual atoms can be averaged over by using large laser spot sizes on the mirrors (or similarly by coupling into higher order modes* than TEM_{00} as they sample more of the mirror surface†), but for a typical ULE substrate on a 10 cm spacer this corresponds to flicker-noise fluctuations of around 1×10^{-15} at all timescales [89]. Fused silica has thermal noise levels around a factor of three lower due to lower mechanical loss, but its CTE has no zero-crossing. Hence any temperature fluctuations can cause deformations on the cavity length, even at the zero crossing of the ULE. Nevertheless, an effective zero-crossing can still exist for the simple design of mirrors contacted directly to the spacer, though this is often at temperatures that are inconvenient for control. Through careful design modifications, the zero crossing can be moved closer to room temperature. Such designs include placing a ULE annulus on the back of the silica

*A non-confocal cavity geometry is generally used for high quality reference cavities to deliberately lift the degeneracy of the transverse modes in a controlled way. If the confocal geometry was not realised exactly then the cavity fringe may be broadened by the slight non-degeneracy.

†The TEM_{00} mode is generally used as it is easy to achieve a high coupling efficiency with an input Gaussian beam.

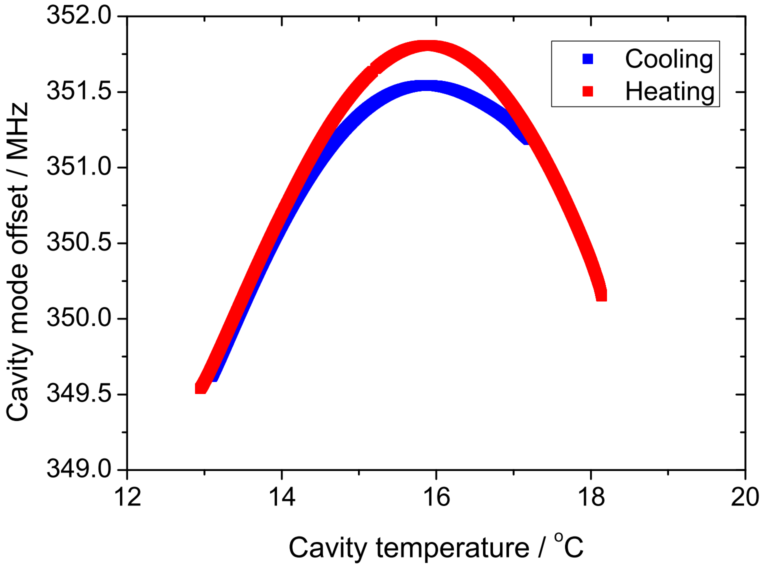

(a)

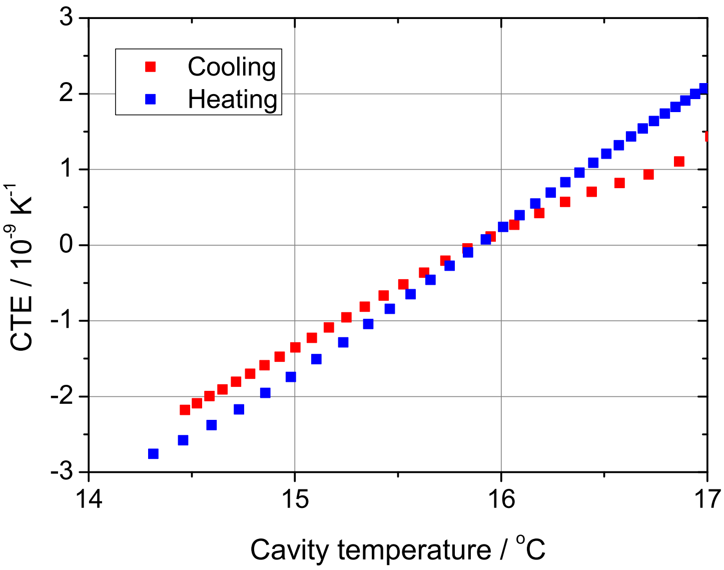

(b)

Figure 4.5: Determination of the zero-crossing temperature for the 871 nm clock laser reference cavity. (a) Offset between the 871 nm cavity mode and a fixed reference cavity, which is then differentiated and combined with the laser frequency to give (b) CTE as a function of temperature. The gradient is close to $1.5 \times 10^{-9}/\mathrm{K}^2$ as expected.

mirror substrate [90] or by making the mirrors re-entrant into the bore of the spacer [91]. In this manner, frequency instabilities approaching 1×10^{-16} have been reached at averaging times from 1 to 30 s [92, 93]. A cavity with such low short-term instability is arguably unnecessary for the 871 nm probe laser as it should be possible to realise the natural linewidth of the transition using the present optical cavity. However, it would be desirable to implement a cavity of such a design for the 934 nm laser to allow the realisation of even narrower linewidths (and smaller ac Stark shifts) on the 467 nm transition.

4.1.5 Alternative approaches

Some groups are working on alternative cavity designs and materials that will circumvent the issue of thermal noise. A cryogenic cavity constructed from a single crystal of silicon has been demonstrated [94], though it is only transparent at infra-red wavelengths near 1.5 μm. In order to generate a reference for the lasers required for optical clocks, it would be necessary to use an optical frequency comb to transfer the stability to a slave laser at the desired wavelength.

An alternative design of reference is a narrow spectral hole burned into a cryogenically cooled doped crystal. A short-term instability of 6×10^{-16} [95] has been demonstrated, and the holes appear to have competitive or lower sensitivities to environmental perturbations than even the best optical cavities.

One might also consider generating a clock laser using the same atomic species as used as the clock reference. The atoms could be confined in an optical lattice within a high finesse optical cavity [96, 97], and made to interact in a collective manner with a light field in the cavity (superradiant laser). The resulting radiation would have an incredibly narrow linewidth, but would be of very low power and therefore would need to be used to injection-lock a slave laser which could then be used to interrogate the reference atoms.

4.2 Achieving narrow linewidths

Once a stable design has been found for the spacer, attention must turn to the mirrors that form the optical cavity in order to produce a reference oscillator with a high Q-factor. The laser must then be actively stabilised to this reference at a level where the dominant source of noise is the noise on the cavity mode itself.

4.2.1 Generating a steep discriminant

In order to stabilise the laser, it is necessary to produce a discriminant that acts to convert laser frequency noise into an error signal. Assuming that the laser lock is limited to some fraction of the error signal by detection shot noise, a steeper discriminant will allow a more sensitive detection of the laser frequency fluctuations and hence a more stable laser after locking*. One way to generate a steep frequency discriminant is to use an absorption feature with a very narrow linewidth $\delta\nu$, for example a mode of an optical cavity. The linewidth of an optical cavity mode is determined by the finesse $\mathcal{F}$ of the mirrors and by the free spectral range (FSR) of the cavity. The finesse is given by:

$$\mathcal{F} = \frac{\mathrm{FSR}}{\delta\nu} \approx \frac{\pi R^{1/2}}{1-R} \tag{4.8}$$

where R is the power reflectivity of the mirrors, which are made to be equal (this is not necessary, but allows observation of transmission through the cavity as an indicator of the lock quality). By using highly reflective mirrors ($R \geq 99.997\%$) and keeping losses due to absorption and scatter much smaller than the transmission loss, it is possible to generate a finesse of more than 10^5. The cavity linewidth is determined by performing a 'ring down' measurement, which involves allowing the laser power to build up inside the cavity and then measuring the decay of the light leaked out of the cavity once the input light is switched off. The linewidth of the cavity mode is related to the $1/e$ decay time τ

*But in many locking schemes, this comes at the expense of a reduced capture range and lock bandwidth.

of the light by:

$$\delta\nu = \frac{1}{2\pi\tau} \tag{4.9}$$

The decays for the two cavities and the fitted exponential models are shown in figure 4.6. Combined with their 10 cm length (giving their FSR $= c/2L \approx 1.5$ GHz), this leads to a finesse of 10^5 for the 871 nm cavity and 2×10^5 for the 934 nm cavity.

4.2.2 Pound-Drever-Hall locking

Various schemes exist for stabilising lasers to optical cavities. One technique is to generate an error signal from the transmission of the cavity, such as with the side-of-fringe lock discussed in subsection 3.4.2. However, since the cavity fringe acts to convert frequency noise to noise on the transmitted intensity and vice-versa, any residual AM on the laser has the potential to disturb the zero of the error signal, leading to frequency fluctuations on the locked laser that are not visible in-loop. In addition, due to the high finesses ($> 10^5$) of the cavities involved, the strong frequency response of the optical cavity limits the detection bandwidth for the transmitted light. The low-frequency noise on the detectors and elsewhere in the control loop can also degrade the lock performance.

A technique that provides an error signal that is not sensitive to laser AM and is more weakly sensitive to the frequency response of the cavity is the Pound-Drever-Hall (PDH) lock [98, 99]. By stabilising the laser to the centre of the cavity resonance, the sensitivity to laser AM is minimised near the lock point. The detected signal used for the lock is typically in the MHz range, and therefore away from the low frequency noise on the detector. The focus of this subsection is the experimental implementation of such a lock.

The key elements of the PDH lock are shown in figure 4.7. Frequency sidebands are applied to the laser by an electro-optic modulator (EOM), driven resonantly in a similar configuration to the EOMs used to spin the polarisation of the cooling beam (see section 3.4.1). The modulation depth is such that the sidebands are around 20% of the intensity

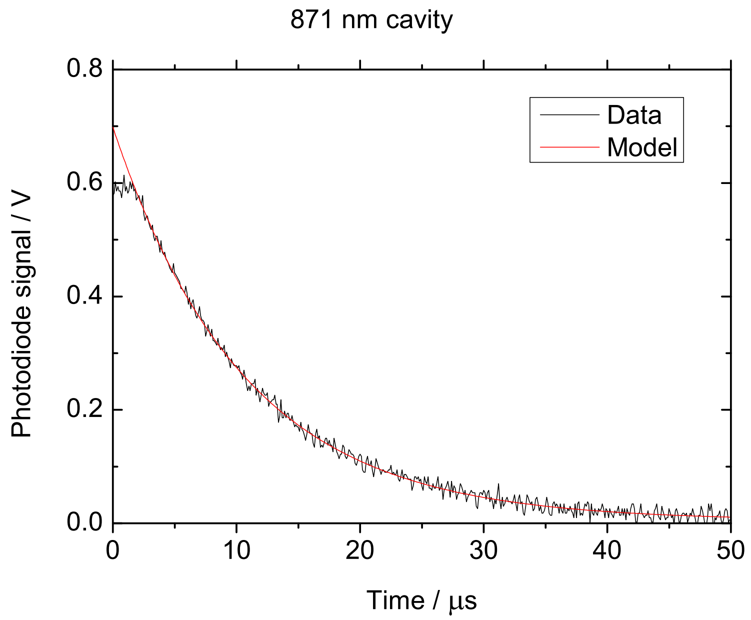

(a)

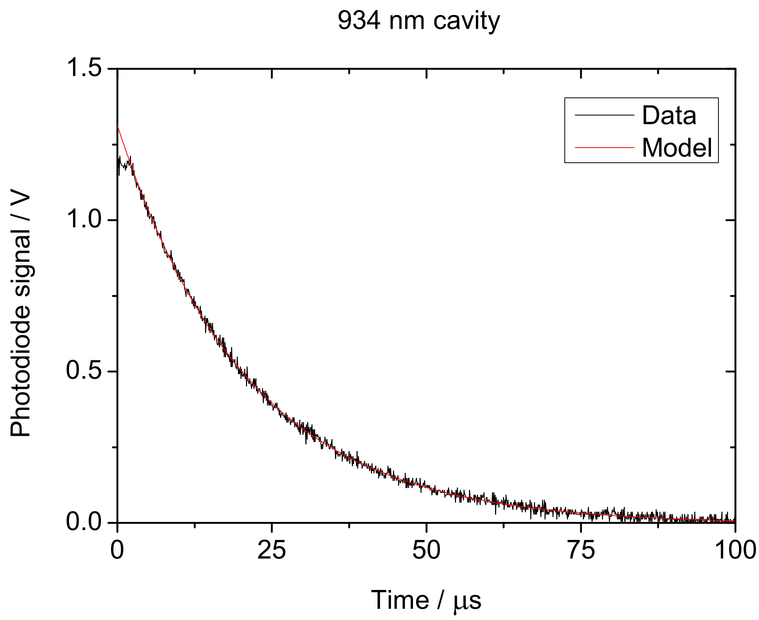

(b)

Figure 4.6: Decay of the light leaked from the two optical cavities at their operating wavelengths. The input light was shuttered off using an AOM shortly after $t = 0$. Exponential fits to the data lead to decay time constants of 11 μs and 21 μs for the 871 nm and 934 nm cavities respectively, and hence cavity linewidths of 15 kHz and 7.6 kHz.

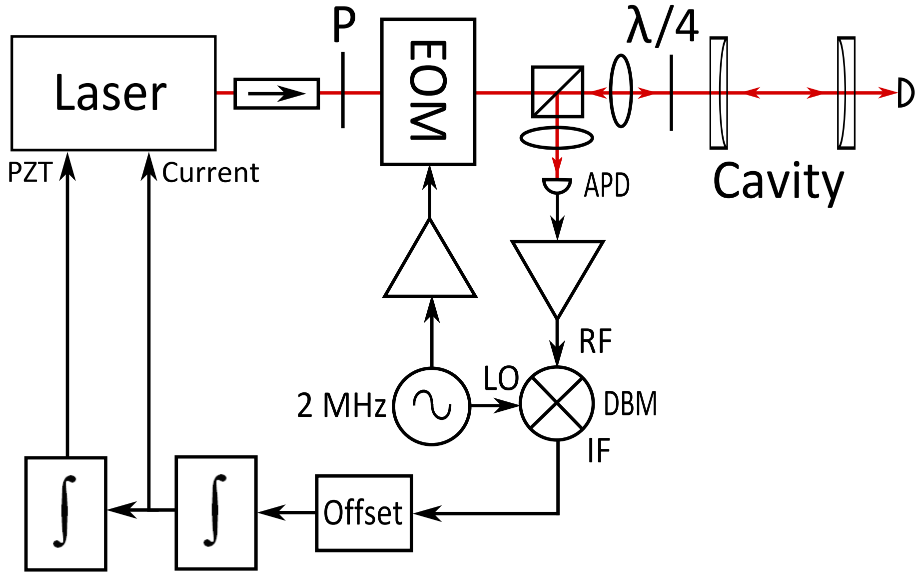

Figure 4.7: Simplified depiction of the probe laser setup, focusing on the main elements for the PDH lock. Optical paths are shown in red, and electrical signals are in black. The three doubled-balanced mixer (DBM) ports are explained in the text. P = polariser, APD = avalanche photodiode.

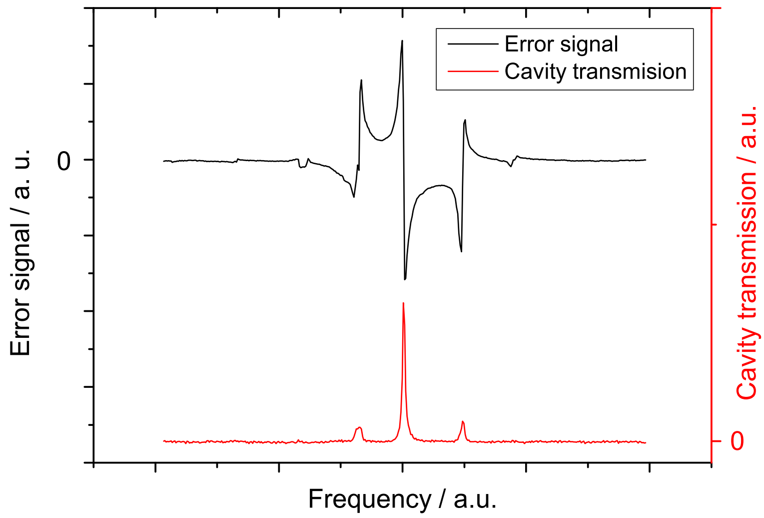

Figure 4.8: Example PDH error signal generated by sweeping a modulated 639 nm laser over an optical cavity mode. The slight asymmetry in the sidebands is attributable to the frequency modulation being applied via the laser diode injection current, which results in a considerable amount of amplitude modulation. For the clock lasers, the modulation is applied by an EOM to prevent this effect.

of the carrier. The modulation frequency is chosen to be many times the linewidth of a cavity mode to ensure that the sidebands and carrier cannot simultaneously be on resonance, and care is taken to ensure that the sidebands are not close to any higher-order transverse cavity modes. The nature of the modulation means that the first-order sidebands are of opposite phase. When the carrier is on resonance with the optical cavity, the sidebands are reflected by the front mirror of the cavity. This light is collinear with any leakage from the cavity, which consists of a superposition of waves that have made an integer number of round trips. This counter-propagating beam is picked off using a beamsplitter and monitored on an avalanche photodiode capable of resolving a signal at the modulation frequency, which arises from the high and low frequency sidebands beating with the leakage from the cavity. The two beats are exactly 180° out of phase when the laser is on resonance with the cavity and will perfectly destructively interfere, resulting in no signal on the photodiode. If the laser is not exactly on resonance, the beats will not be exactly 180° out of phase. The phase of the resulting signal carries information about which side of resonance the laser is on. Due to the long storage time of the cavity, the leakage beam contains information about the average frequency and phase of the laser over that time and hence provides a stable reference for fast fluctuations in the input laser frequency (and hence the lock bandwidth is not limited by the optical cavity).

This signal is amplified* and sent to the radio frequency (RF) port of a double-balanced mixer† (DBM), and a sample of the EOM drive signal is sent to the local oscillator (LO) port. An appropriate delay line is added before the LO port to ensure that the two signals entering the mixer are 90° out of phase, resulting in the maximum possible discriminator slope. The output of the mixer intermediate frequency (IF) port consists of a signal at twice the modulation frequency and a dc signal, which is normally isolated by either using a low-pass filter or using a notch filter to suppress the signal at twice the modulation frequency. However, in this experiment the desired lock bandwidth

*Minicircuits MAN-1LN.

†Minicircuits ZAD-6+.

is close to the modulation frequency and hence a low-pass filter would introduce additional phase delays which would limit the bandwidth, as would a notch filter of insufficient Q. The signal at twice the modulation frequency should not pose a problem as it is outside the desired lock bandwidth (see subsection 4.2.5), but care must be taken that it does not adversely affect the operational amplifier (op-amp) used for the lock.

The dc signal from the mixer is dispersion-like in shape as shown in figure 4.8, but the zero-crossing can be offset from the centre of the cavity resonance. This offset can be electrical or optical in nature. The electrical offset arises from the various components used to process the error signal, and can be observed by monitoring the error signal level when no light is present on the photodiode. A simple voltage adder connected to a stable, tunable reference voltage can then be used to null the offset. Offsets that are present only when light falls on the photodiode are generally attributable to amplitude modulation. If the laser is not polarised along the EOM crystal axis, the polarisation of the transmitted beam is also modulated at the EOM drive frequency. This polarisation modulation manifests as amplitude modulation once the beam has been through any polarising optics, and cannot be distinguished from the required frequency modulation signal by the photodiode. This signal can be reduced by using a Brewster-cut crystal, which prevents randomisation of the polarisation by scatter from the crystal surfaces, and then carefully aligning the polarisation of the light entering the EOM to be along the crystal axis. This is achieved by passing the light through a calcite polarising beam splitter with a 10^5:1 extinction ratio. When aligning the light reflected from the cavity onto an avalanche photodiode* (APD), the signal is first maximised by looking at the Fourier transform of the signal from the photodiode when the laser frequency is scanned across a cavity resonance. The laser is then tuned away from resonance, and the residual signal at the EOM drive frequency is minimised by fine-tuning of the polariser angle. Some residual amplitude modulation (RAM) remains, and this has a spatial dependence across the beam. The signal attributable to RAM causes a laser power-dependent offset

*Excelitas C30902EH.

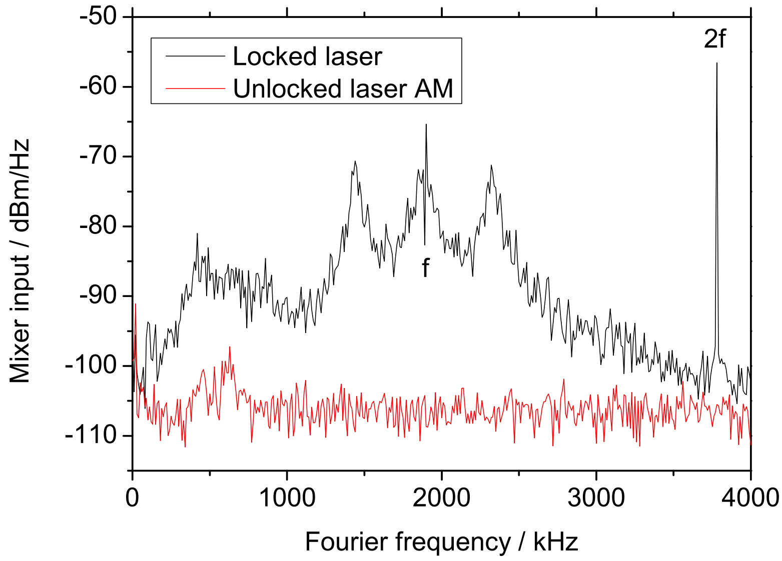

Figure 4.9: Sample in-loop signal on the avalanche photodiode used to monitor the retroreflection from the cavity. There is a pronounced feature at twice the modulation frequency of 1.9 MHz, which may be problematic if there is a substantial signal at the same frequency on the mixer LO port.

on the PDH error signal. This offset cannot be eliminated, but can be prevented from varying by actively stabilising the power entering the cavity by monitoring a small pickoff of the light before the cavity and feeding back to the drive power of an AOM between the laser and the cavity.

Another source of an offset on the mixer output comes from the second-order sidebands produced by the EOM. As they are in-phase, their beats with the leaked carrier will constructively interfere and hence the APD will also detect a signal at twice the modulation frequency (as shown in figure 4.9). This signal is problematic if there is any signal at the same frequency on the LO port of the mixer, as this will generate a DC signal on the mixer output. In this experiment the LO signal is produced by a filtered synthesiser*, resulting in a second harmonic smaller than -60 dBc and so any DC signal produced in this way should be small.

*Marconi 2022E with Minicircuits BLP-2.5.

4.2.3 Coupled power

Due to the high buildup factor of the cavity ($B = \mathcal{F}/\pi$), the circulating power in the cavity can be very high even for a relatively modest input power. This leads to local heating on the mirror surfaces, changing the cavity length and hence its resonant frequency. This sensitivity has been measured to be on the order of 10 Hz/μW of coupled light [87] for cavities similar to those used in this experiment, and so the input power is kept to 5 μW or below. The active stabilisation of the power entering the cavity also acts to stabilise this shift.

4.2.4 Parasitic etalons

Any optics in the beam path have the potential to set up a low-finesse cavity. The laser and its sidebands sit inside a mode of this 'parasitic' etalon. The upper and lower sidebands sample different parts of the etalon mode, and so develop different phases and amplitudes. This means that the beats of the upper and lower sidebands with the leaked carrier cannot perfectly destructively interfere at the centre of the cavity mode, leading to a laser frequency-dependent offset on the PDH error signal as shown in figure 4.10. Any changes in the refractive index of the air in the beam path (e.g. changes in the lab temperature) will change the central frequency of the mode of the parasitic etalon, leading to changes in the lock offset. Care is taken to ensure there are no normal faces in the beam path by using Brewster-cut optics where possible and ensuring that all other optics are not normal to the beam. However, scattered light can also lead to the formation of an etalon and so all optics are AR-coated where possible. We have found that one of the most likely places for an etalon to form is between the APD and the front mirror of the high-finesse cavity, and so the APD is tilted at an angle of around 30° with respect to the incident light and the beam is defocused from its surface to prevent the coupling of scatter back into the beam path. Another likely normal surface is that of the retro-reflector used to double-pass the light through an AOM before it reaches the optical cavity. Right-angle prisms with an AR-coated hypotenuse are now used as retro-relectors instead of mirrors in all of the double-pass setups.

The laser modulation frequency has also been reduced to around 2 MHz to reduce the separation of the sidebands within the parasitic etalon (but without restricting the achievable lock bandwidth), thus reducing the difference between the phase and amplitude of the two sidebands. Even with all these precautions, the effect of parasitic etalons could still be seen on the laser stability and so the entire probe laser setup is now enclosed within a wooden box lined with acoustic foam. There is a large heat load inside the box from the optical cavity temperature control, which has the potential to disturb beam alignments as the temperature rises. The internal temperature is maintained close to room temperature by passing water at 15°C through a heat exchanger in good thermal contact with a large heat sink recessed into the box lid. To minimise the coupling of any resulting vibrations to the cavity, the box has no physical contact to the AVI and the flow rate is limited to around 2 l/min.

4.2.5 PID Controllers

Once the error signal has been generated it can be used to feed back to the laser diode to correct for any frequency noise. The gain profile of the servo loop has to compensate for the phase lags inherent in the system, in particular the high-frequency response of the laser diode. As the servo is designed to provide negative feedback, a phase lag of more than 180° corresponds to positive feedback. The frequency and gain profiles of the servo are chosen to shift this critical point to as high a frequency as possible, resulting in a high lock bandwidth. Reducing the total system gain below unity at this frequency will suppress any signals that would cause positive feedback. The 871 nm clock laser system is used as an example in this subsection.

Throughout this subsection, 'gain' refers to the *amplitude* gain (as opposed to the *power* gain, which is the square of the amplitude gain).

Transfer functions (Bode plots) of an idealised PID controller are shown in figure

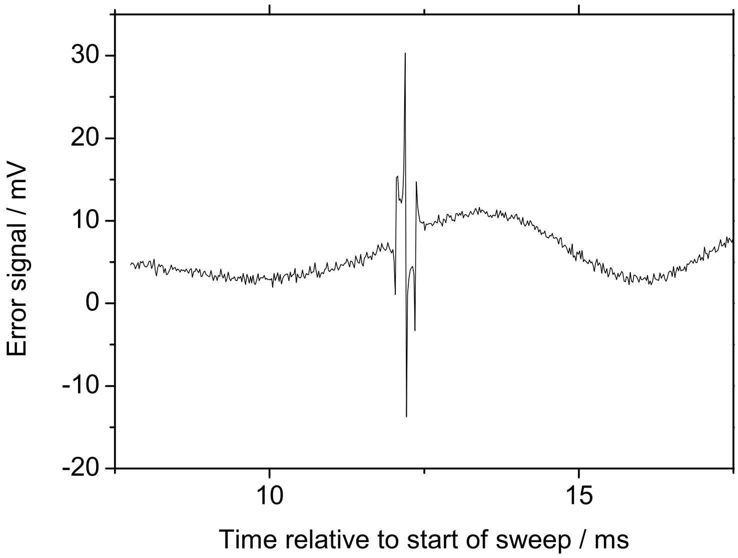

Figure 4.10: Pronounced modulation on the PDH error signal for the 871 nm laser as the laser frequency is swept over a cavity mode. The FSR of the parasitic etalon can be determined by calibration of the sweep against the FSR of the cavity. In this case the FSR was approximately 150 MHz, indicating an etalon length of around 1 m. This was identified to be between the front mirror of the optical cavity and the retro-reflecting mirror for a double-pass AOM setup, and was eliminated by replacing the mirror with an AR-coated right-angle prism.

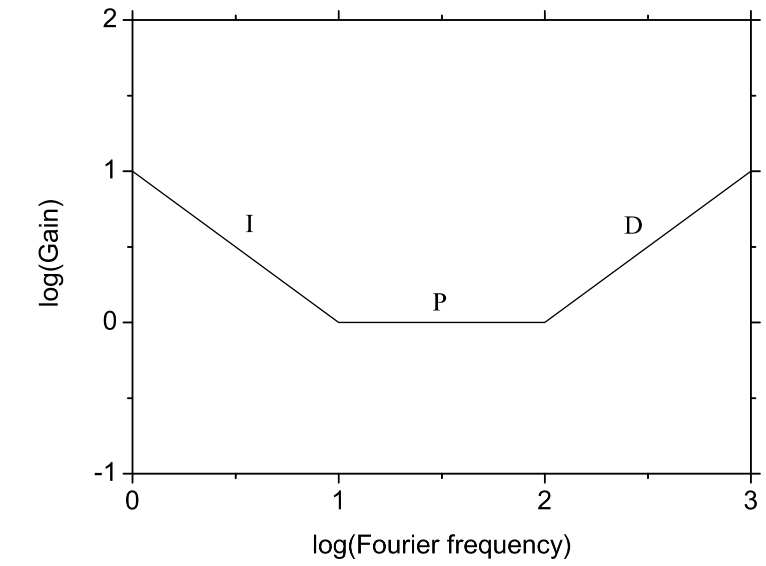

(a)

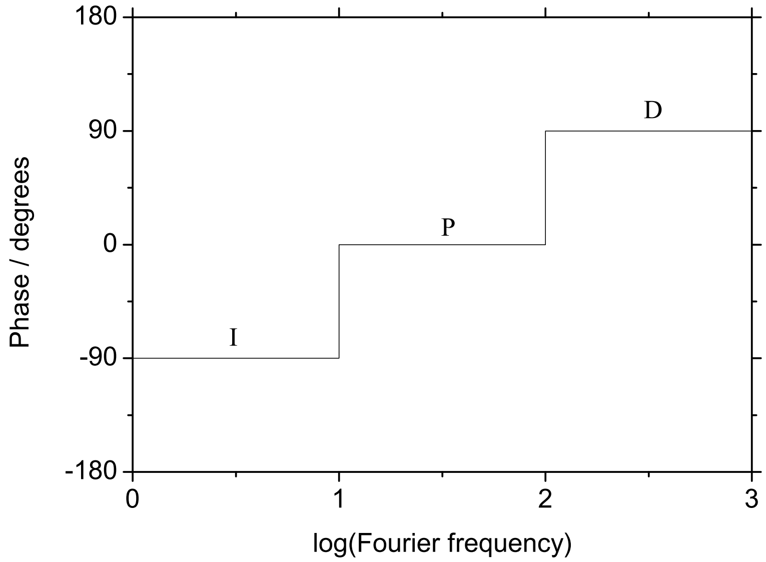

(b)

Figure 4.11: Idealised transfer function for a PID controller. (a) Gain vs frequency, (b) phase lag vs frequency. In a real system, the sharp discontinuities are smoothed.

4.11. In order to reduce the system gain below unity at the point where the system phase lag exceeds 180°, it is desirable to have a servo whose gain decreases with increasing frequency. This is characteristic of an integrator, in which gain rolls off as $1/f$. However, as there are other gain rolloffs associated with the system, the total system gain will fall as $1/f^2$, which renders the system unstable if unity gain is passed with this slope. Hence the servo gain must make a transition to f^0 behaviour before the unity gain point, which is characteristic of proportional gain.

Even though care has been taken in its design, the free-running linewidth of the 871 nm laser is still around 100 kHz and therefore a servo bandwidth of at least this figure is required to suppress the residual noise. In general, the components that limit the achieveable lock bandwidth are the optical cavity, the laser diode and the servo itself. As the optical cavity has a linewidth of around 15 kHz, it behaves as a low-pass filter with the gain rolling off at -3dB/octave for frequencies above 15 kHz. The laser diode also has a roll-off in gain and phase for frequencies higher than 100 kHz as shown in figure 4.12b, due to a change in the modulation mechanism [100]. These phase roll-offs can be compensated for (to some extent) by introducing a phase lead at high frequencies, which is characteristic of derivative gain. However, as the phase lag continues to increase with frequency, the bandwidth cannot be extended indefinitely.

The servo designed to optimise the lock bandwidth is shown in figure 4.12c. The phase lead at high frequencies is clearly visible.

The circuit design chosen to realise these frequency and phase parameters is shown in figure 4.13. The second integrator stage was added to provide additional gain at low frequencies, which allows further narrowing of the laser linewidth by an increased suppression of noise at frequencies below 1 kHz where the noise is greatest.

The phase and frequency response of the final system is shown in figure 4.14. The observed response agrees well with our desired behaviour at low frequencies, but departs from ideal behaviour for frequencies > 100 kHz. The rapid phase roll-off at high frequencies is characteristic of a time delay somewhere in the feedback loop. The total phase lag

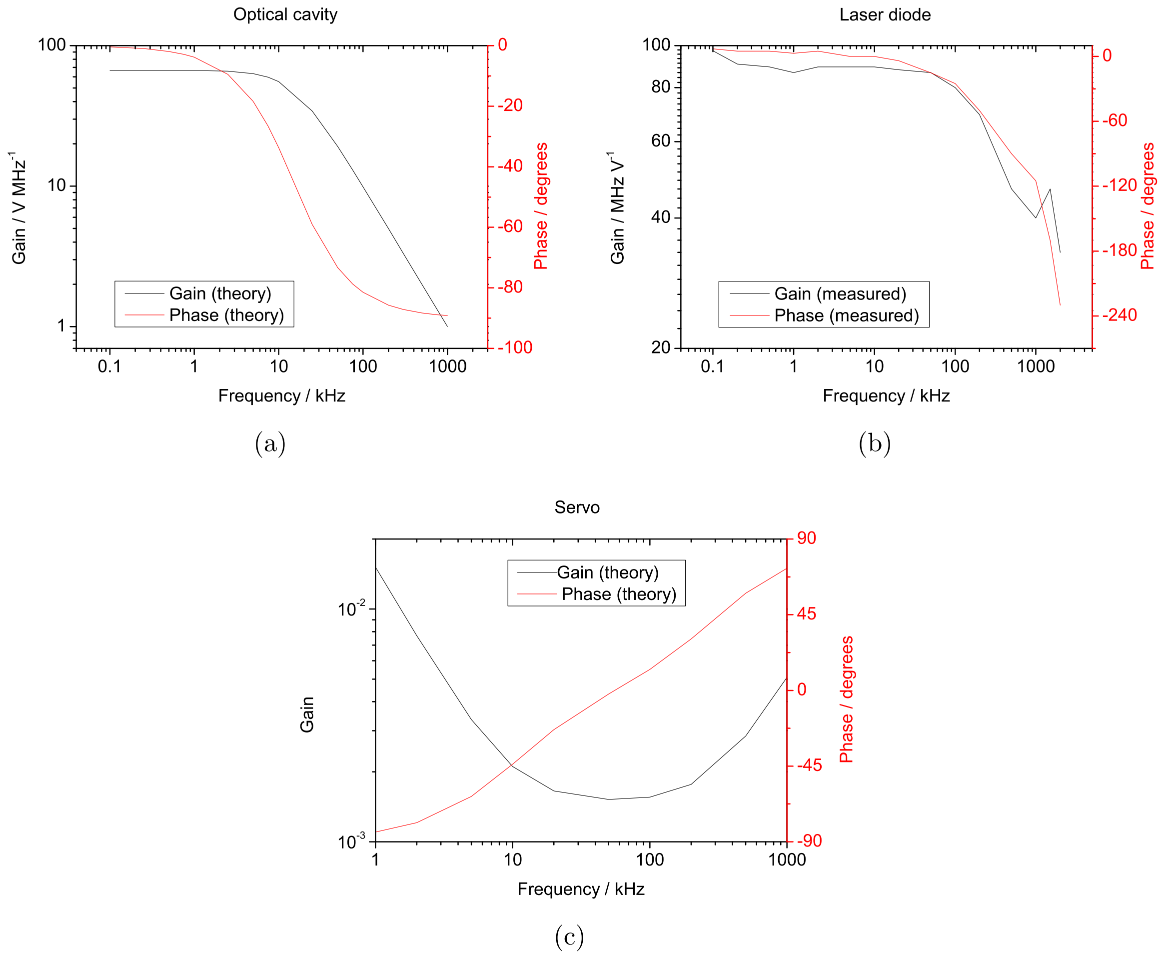

Figure 4.12: Gain and phase responses of the main components of the PDH feedback system for the 871 nm laser. (a) Optical cavity (error signal) response, assuming a mode FWHM of 15 kHz and 1 V p-p error signal. (b) Laser diode, showing the gain and phase rolloffs beyond 100 kHz. (c) Servo designed to optimise the lock bandwidth.

exceeds 180° at around 400 kHz, and our servo gain must be correspondingly reduced to prevent oscillation at this frequency.

4.2.6 Evaluation of lock quality

By analysing the Fourier transform of the in-loop error signal with the laser in lock, it is possible to determine the fidelity of the lock to the reference cavity. As a first approximation, the modulation index of the in-loop noise can provide this information.

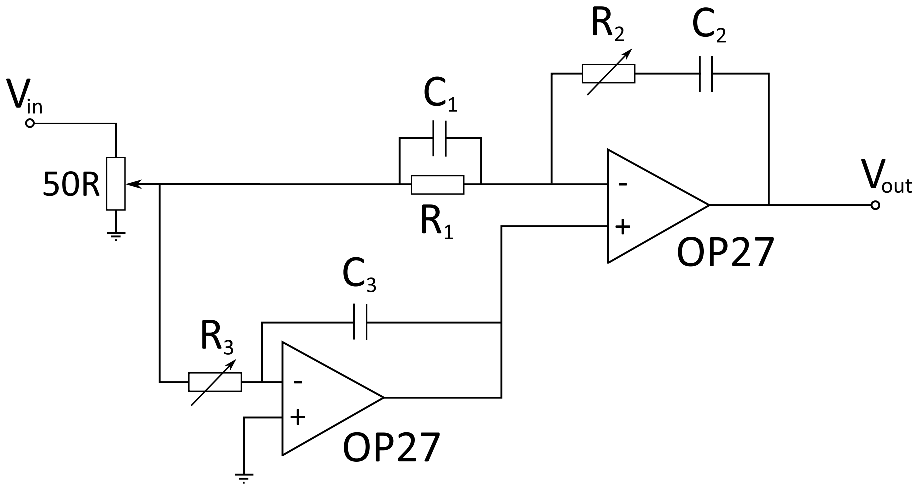

Figure 4.13: Schematic of the feedback portion of the loop filter used to stabilise the 871 nm laser to its high finesse cavity. The potential divider at the input allows for coarse tuning of the overall gain. Careful selection of the capacitor and resistance values allows optimisation of the feedback. Wire-wound potentiometers were avoided due to their stray inductance and capacitance. The OP27 op-amps have a gain-bandwidth product of 8 MHz, which does not limit the achievable servo bandwidth.

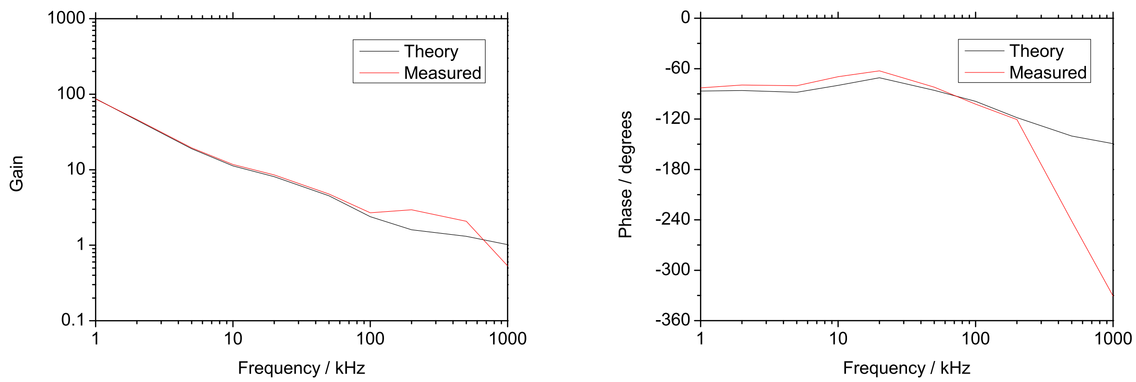

(a)

(b)

Figure 4.14: Gain and phase responses of the completed PDH feedback system. The achieved response compares very well to the desired response at low frequencies, but departs from ideal behaviour for frequencies > 100 kHz.

The frequency excursions $\delta\nu(f)$ are given by:

$$\delta\nu(f) = \left(\frac{\text{FWHM}}{\text{E}}\right)\delta V(f) \tag{4.10}$$

where $\delta V(f)$ is the voltage noise at a given Fourier frequency, FWHM is the full-width at half-maximum of the optical cavity mode and E is the peak-to-peak amplitude of the error signal. For a modulation index $\beta = \delta\nu(f)/f = 1$, significant power is transferred into the first order sideband and hence $\beta \ll 1$ is desirable. The linewidth of the laser relative to the cavity mode is roughly given by the point at which $\beta = 1$ [101].

4.3 Lasers

4.3.1 436 nm probe laser

Light at 871 nm is generated by an external cavity diode laser (ECDL) in the Littrow configuration. The laser cavity is of a monolithic design, which provides good mechanical stability and also good passive temperature control. The diode is removed from its protective can and AR-coated ($R = 5\times10^{-5}$) to prevent competition with the external cavity modes. Feedback of 9% is provided by a blazed diffraction grating with 1800 lines mm^{-1}, leading to a threshold current of 35 mA. Around 50 mW of light is coupled out of the cavity at the operating current of 115 mA. To prevent linewidth broadening due to noise on the current supply, a low-noise controller was built with residual noise of < 60 nA rms. The diode temperature is actively stabilised to 15°C by a Peltier cooler, with the error signal generated by a dc Wheatstone bridge. Out-of-loop temperature fluctuations are at the mK level over several hours.

The laser is split into three main beam paths, as shown in figure 4.15. A small fraction of the light (5 μW) is sent to the high finesse cavity for locking. En-route it passes through an acousto-optic modulator (marked AOM A in figure 4.15) in order to offset the light from the nearest cavity mode and to stabilise its power. It then passes through an electro-optic modulator, which adds frequency sidebands at 2.2 MHz to the light. The laser is locked to the cavity by feeding back to the diode current using a PID servo with two integrator stages (as shown in figure 4.13). The in-loop noise is shown in figures 4.16a and 4.16b. In figure 4.16a, the 'servo humps' either side of the modulation frequency indicate the lock bandwidth is near 400 kHz (i.e. at this frequency the phase delay has exceeded 180°, and the servo is adding noise to the laser). The low-frequency laser noise is resolved in figure 4.16b, and is compared to the shot noise limit. The modulation index of the in-loop noise on the error signal at frequencies above 0.1 Hz has been reduced to much less than unity (as shown in figure 4.17), indicating a highly faithful lock to the reference cavity. The correction signal is integrated again and used to feedback to the output of the piezo driver controlling the angle of the diffraction grating, thus correcting for large frequency excursions on long timescales. It was observed that the ThorLabs piezo driver added a large amount of noise to the laser at frequencies below 10 kHz, with some pronounced spikes around 3 kHz which may correspond to system mechanical resonances. As the piezo is only required for slow feedback, a low-pass filter with a cutoff frequency of 150 Hz was added to the output of the driver, and a large reduction in the in-loop noise was observed as shown in figure 4.18.

The majority of the laser power is focused through a 7 mm-long crystal of potassium niobate ($KNbO_3$) in order to frequency double the light to 436 nm in a single pass. Non-critical phase matching is achieved by heating the crystal to 53°C in a small oven in good thermal contact with a power resistor, powered by a ThorLabs standalone temperature controller. Around 4 μW of 436 nm light is produced for a fundamental power of 30 mW when the size and position of the beam focus is optimised within the crystal. This light is then coupled into an optical fibre and delivered to part of the optical table where it is combined with the other three blue and UV wavelengths.

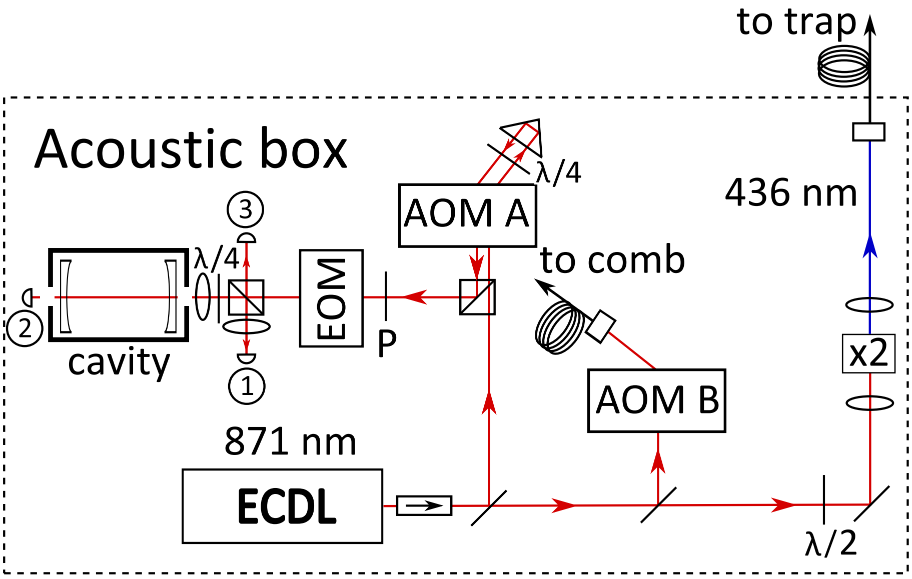

Figure 4.15: Schematic of the 871 / 436 nm laser system showing the principal components. AOM A offsets the laser from the nearest optical cavity mode. Frequency sidebands are added to the laser using an EOM, and detector 1 monitors the PDH error signal. Detector 2 monitors the cavity transmission as a measure of the lock status. Detector 3 is used for the stabilisation of the power incident on the optical cavity. AOM B is used to cancel the noise introduced by the optical fibre between the laser and the frequency comb (see section 7.9), and the majority of the 871 nm light is passed through a doubling crystal in order to produce the probe wavelength of 436 nm.

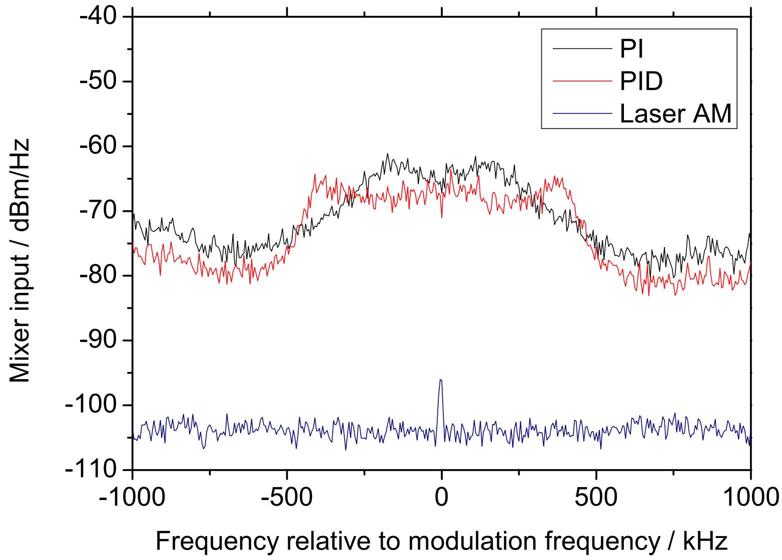

(a)

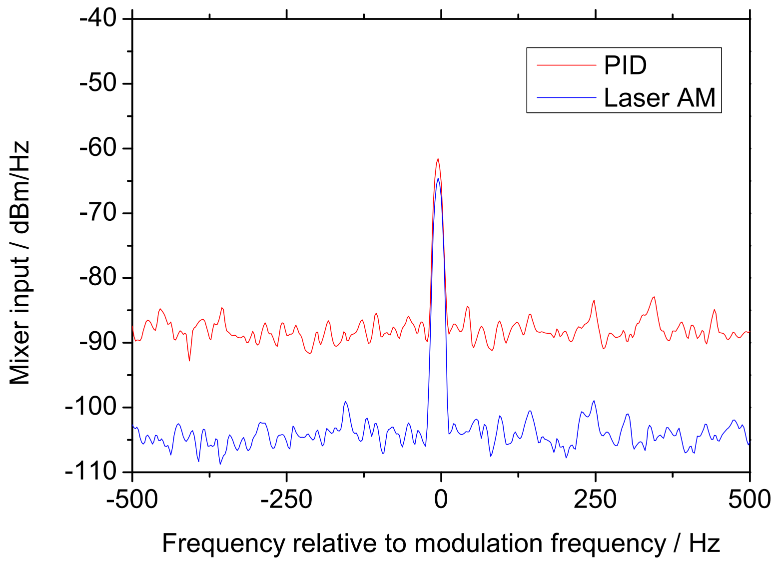

(b)

Figure 4.16: Signal on the avalanche photodiode with the 871 nm laser locked to its cavity, focusing on frequencies close to the modulation frequency. (a) 2 MHz span, showing the lock bandwidth increasing from 170 kHz to 410 kHz with the introduction of derivative gain to the servo. The central dip in noise at the modulation frequency is unresolved. (b) 1 kHz span, resolving the central dip and showing that the in-loop noise has been reduced to close to the shot noise limit. The relatively large RAM peak was attributable to imperfect alignment of the beam through the EOM, and has since been reduced to a level similar to that in figure 4.19b.

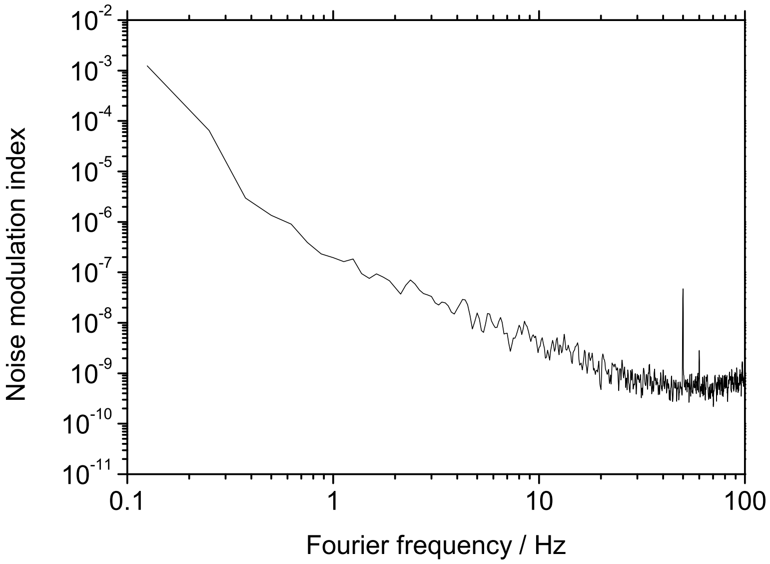

Figure 4.17: Modulation index of the noise on the in-loop error signal for the 871 nm laser at frequencies below 100 Hz. The peak at 50 Hz is the UK mains frequency.

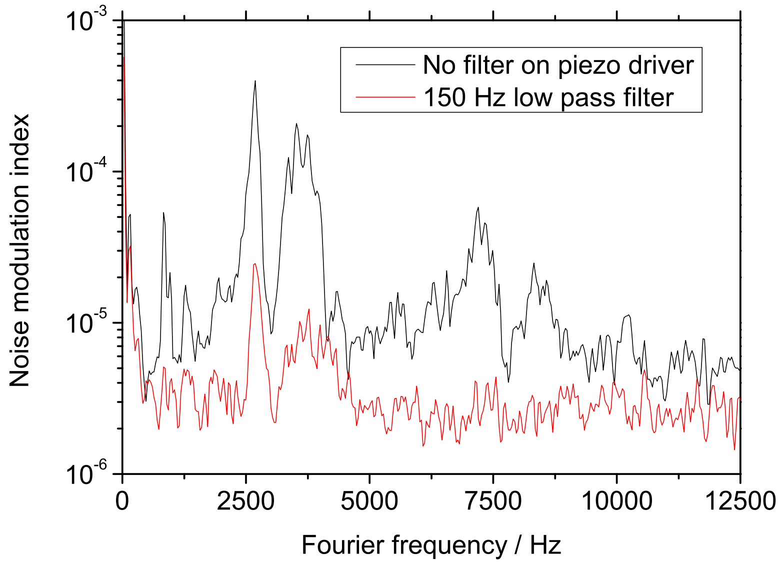

Figure 4.18: Modulation index of the noise on the in-loop error signal for the locked 871 nm laser. It can be seen that adding a low-pass filter to the output of the piezo controller causes a large reduction in the noise at frequencies below 10 kHz.

4.3.2 467 nm probe laser

Achieving a narrow linewidth is even more important for the 467 nm laser as the ac Stark shift of the transition scales as the square of the resolved transition linewidth.

Light at 934 nm was previously generated using an Ar^+-pumped Ti:sapphire laser. The laser cavity was of a monolithic design to improve its mechanical stability, and single-mode operation was achieved by placing a thin etalon in the cavity. Coarse tuning was provided by a Lyot filter. An error signal was derived from a reflection from the face of the thin etalon to keep the laser fixed on this single cavity mode. The laser was pre-stabilised to a small temperature-controlled optical cavity by feeding back to a piezoelectric actuator attached to the rear of one of the cavity mirrors, and then further stabilised to the high finesse cavity using an AOM. The benefits of using a Ti:sapphire laser were its inherently narrow linewidth before stabilisation, and high output power. The disadvantages included frequent mode hops of the laser which limited the uptime of the experiment, and the extremely high power consumption of the Ar^+ laser. This system was replaced with a commercial ECDL and tapered amplifier system*, with all its supplied electronics including current and temperature controllers.

The laser is stabilised to its high finesse cavity in a near-identical way to the 871 nm laser, but using a slightly lower modulation frequency of 1.9 MHz. The laser is stabilised to the cavity by a Toptica-supplied FALC 100 module, which acts as a PID servo with three integrator stages, and feeding back to the high-bandwidth input on the laser head. The in-loop noise is shown in figures 4.19a and 4.19b, showing the lock bandwidth of 420 kHz and low-frequency laser noise close to the shot noise limit. The modulation index of the laser noise at frequencies below 100 Hz is much smaller than unity, as shown in figure 4.20.

The stability of the locked laser was determined by delivering a small fraction of the phase-modulated light to the cavity used to stabilise the 871 nm laser. As the modes of the two cavities do not coincide, the light was double-passed through an AOM operating

*Toptica TA 100 'Pro' design.

near 315 MHz. The already-stable 934 nm laser was then locked to the second cavity using a PDH lock, with the feedback being delivered to the synthesiser used to drive the AOM. This drive frequency was monitored using a Stanford SR 620 counter, which also provided the stability at short (< 1 s) averaging times using its own internal algorithms. The drive frequency was then recorded with a gate time of 1 s (see figure 4.21a) and subsequently processed to give the longer-term stability (figure 4.21b).

Light at 467 nm is produced by resonantly doubling the 934 nm light in potassium niobate (see figure 4.22). The build-up cavity has a finesse of around 100, and phase matching is achieved by angle-tuning of the crystal. The temperature of the crystal is actively stabilised slightly above room temperature. The length of the cavity is stabilised to the input laser by feeding back to the voltage applied to a piezoelectric actuator attached to one of the plane mirrors in the cavity (M2 in figure 4.22), with the error signal derived using the Hänsch-Couillard technique [69]. The output voltage from the ThorLabs controller is passed through a "twin-T" notch filter to prevent feedback at the mechanical resonance near 44.6 kHz, as shown in figure 4.23a (in fact, it was found that the presence of this resonance causes large amplitude modulation, and is enough to prevent coherent excitation of the 467 nm transition). Once the build-up cavity is locked, around 80 mW of 467 nm light is obtained when 400 mW of the fundamental is incident on the input mirror. The residual amplitude modulation is shown in figure 4.23b.

4.4 Summary

Ultrastable lasers at 871 and 934 nm have been developed in order to drive the two highly forbidden clock transitions in $^{171}Yb^+$. The precise environmental control for the two optical cavities has been detailed, along with the measurements of their finesse and the active stabilisation of the two probe lasers to these cavities. Preliminary measurements indicate that the 934 nm laser stability is limited by the thermal noise of the mirror substrates.

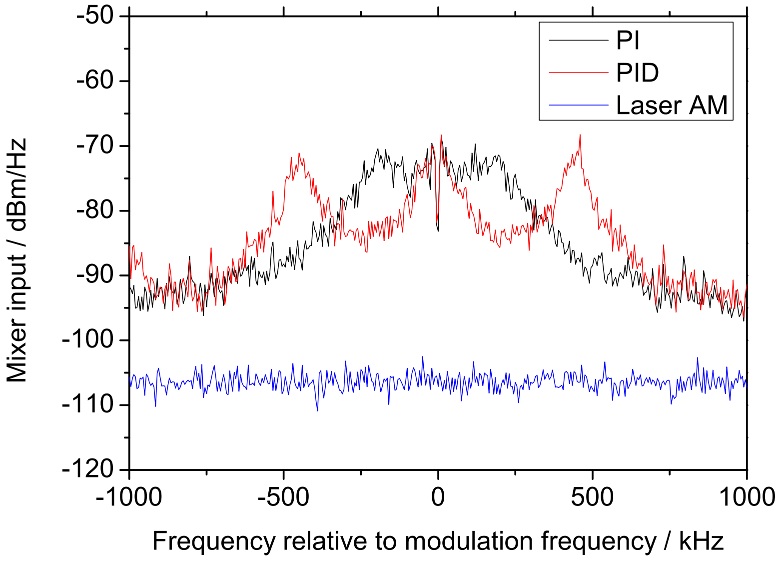

(a)

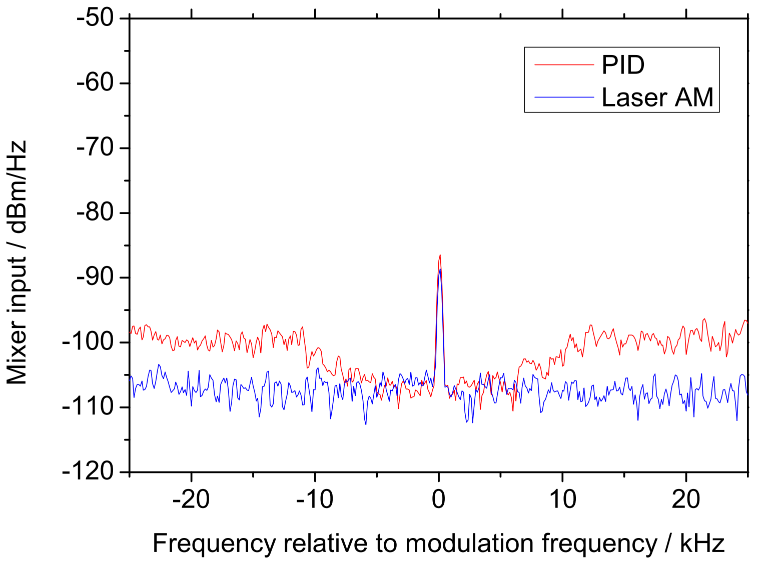

(b)

Figure 4.19: Signal on the avalanche photodiode with the 934 nm laser locked to its cavity, focusing on frequencies close to the modulation frequency. (a) 2 MHz span, showing the lock bandwidth increasing from 200 kHz to 420 kHz with the introduction of derivative gain to the servo. The central dip in noise at the modulation frequency is unresolved. (b) 50 kHz span, resolving the central dip and showing that the in-loop noise has been reduced to close to the shot noise limit for frequencies < 5 kHz.

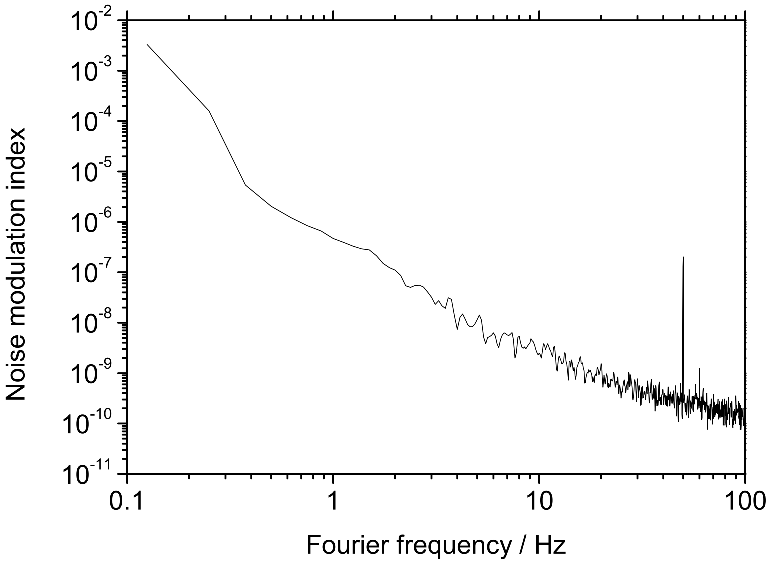

Figure 4.20: Modulation index of the noise on the in-loop error signal for the 934 nm laser at frequencies below 100 Hz. The peak at 50 Hz is the UK mains frequency.

The following chapter will discuss the interaction of these lasers with the ion.

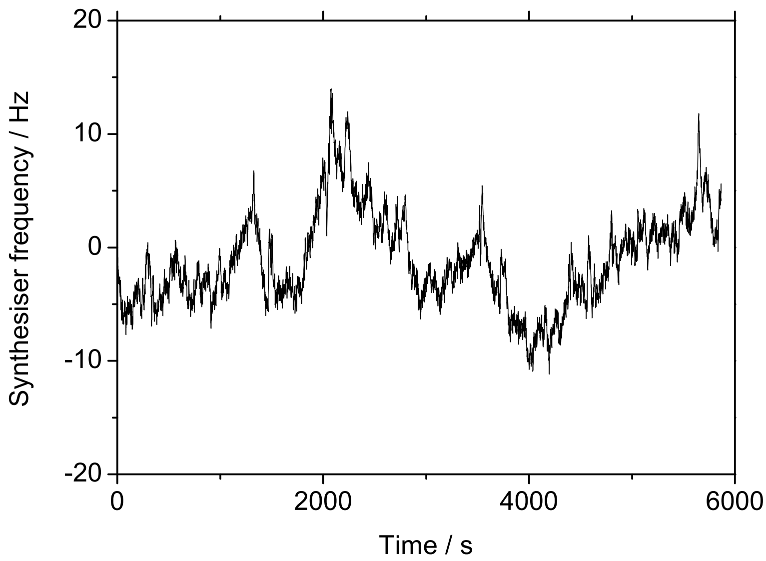

(a)

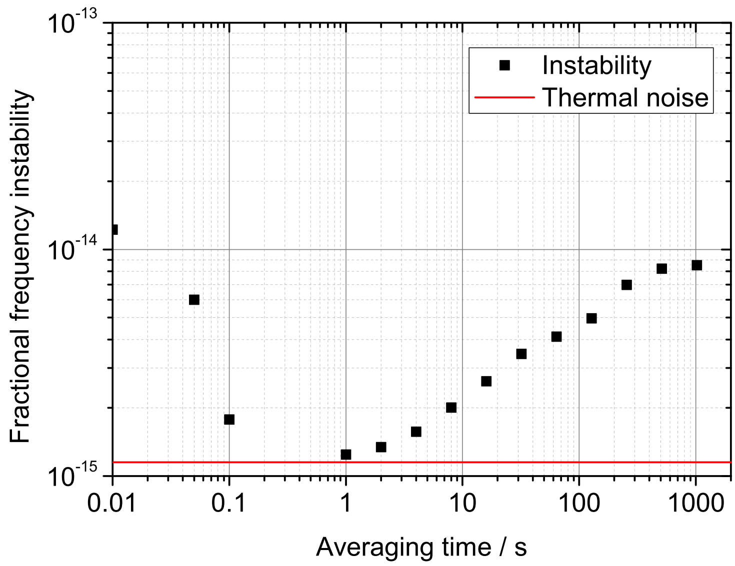

(b)

Figure 4.21: (a) Offset between the two optical cavities as a function of time, with a linear relative drift of 5 mHz/s removed. (b) Allan deviation of the cavity offset frequency. An estimate of the thermal noise based on [84] is also shown.

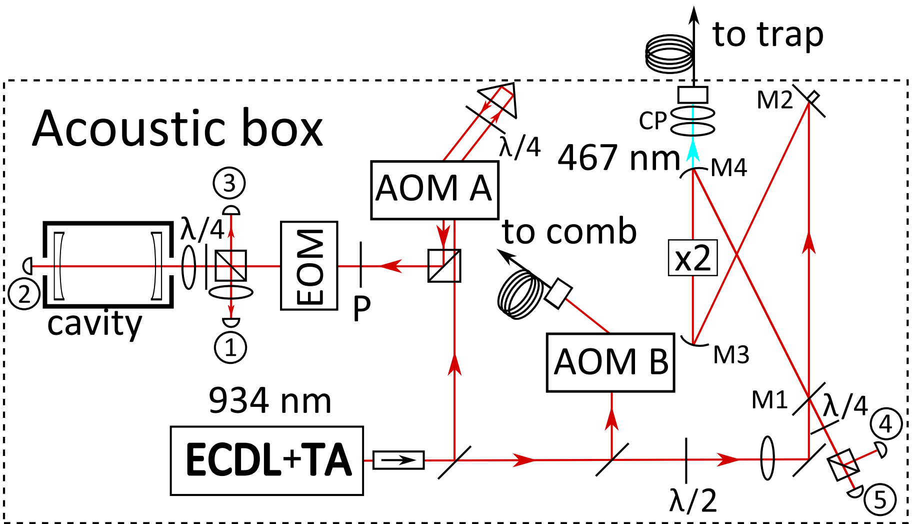

Figure 4.22: Optical setup used to generate the stable light at 467 nm. The setup is largely similar to that used to generate light at 436 nm, except the frequency doubling occurs inside a resonant cavity to increase the conversion efficiency. Light enters the cavity through mirror M1, and the doubled light leaves through mirror M4. The cavity length is stabilised to the input laser by a piezoelectric actuator attached to mirror M2, with the error signal derived using photodetectors 4 and 5. The 467 nm light is corrected for astigmatism immediately outside the doubling cavity using a pair of cylindrical lenses (CP).

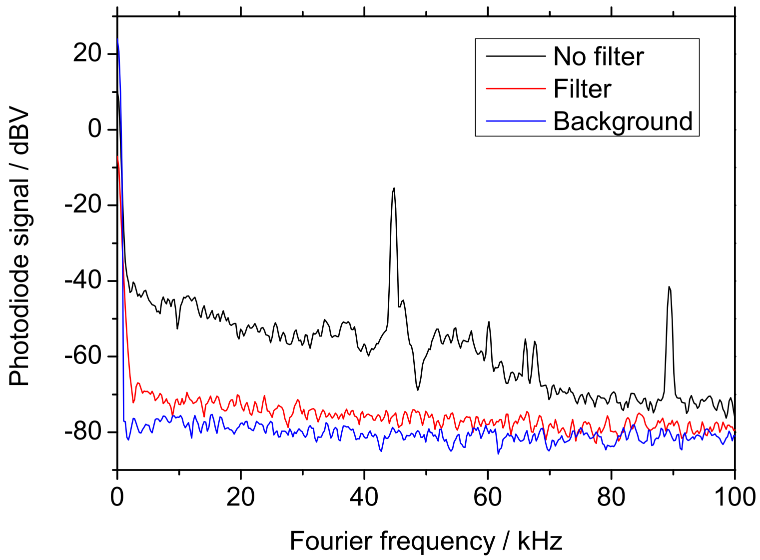

(a)

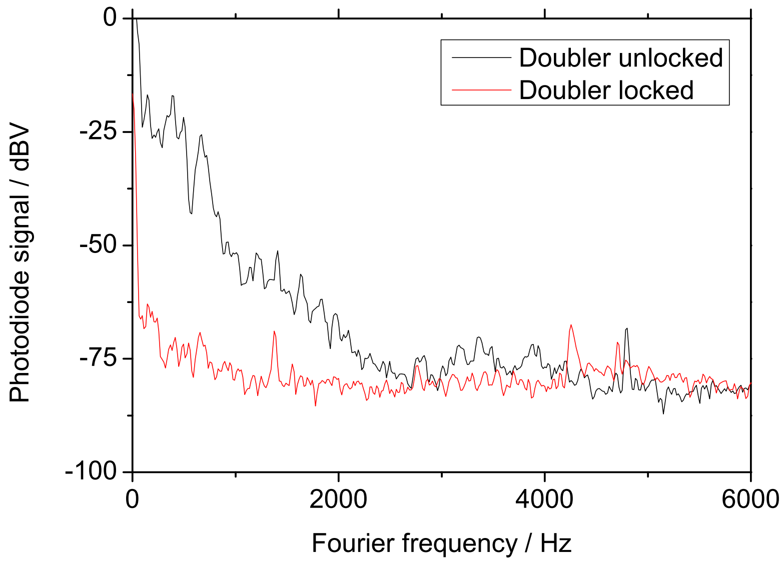

(b)

Figure 4.23: (a) Amplitude noise on the output of the unlocked doubling cavity as measured on a photodiode. It can be seen that the introduction of a notch filter on the output of the ThorLabs piezo controller prevents feedback at the system resonance at 44.6 kHz. The reduction in broadband noise from the piezo driver may be attributable to the low Q of the filter. (b) Amplitude noise on the cavity output with its length unstabilised and stabilised.

Chapter 5

Spectroscopy of highly forbidden transitions

This chapter covers the coherent interaction of monochromatic light with optical transitions in atoms. Recent improvements to the two clock lasers has increased their coherence time to a level that allows observation of coherent excitation profiles for probe pulse lengths exceeding 100 ms. Section 5.1 gives a brief overview of coherent atom-light interaction, detailing Rabi excitation of the forbidden transitions. Section 5.2 contains details of the two main decoherence mechanisms that limit excitation efficiency. Section 5.3 discusses the effect of motional quantum number on the excitation rate, and section 5.4 covers the experimental considerations that act to maximise the coupling between the atom and probe laser.

5.1 Atom-light interaction

The wavefunction of a two-level system in a superposition of ground and excited states $|0\rangle$ and $|1\rangle$ is given by:

$$\Phi = c_0 |0\rangle + c_1 |1\rangle \tag{5.1}$$

where c_0 and c_1 are the complex amplitudes of the two states. The population in each state is given by $|c_i|^2$. This section will discuss the transfer of population between the two states by near-resonant radiation. The case of a dipole transition is used for simplicity, though in reality the clock transitions are dipole-forbidden.

5.1.1 Rabi excitation

The simplest means of exciting a forbidden transition in a two-level system involves the application of a single pulse of resonant radiation with duration t. The coupling of the radiation to the transition depends on the transition dipole moment $\mathbf{d}$ and the electric field vector $\mathbf{E}$ of the exciting radiation. Solving the time-dependent Schrödinger equation for this situation [102] shows that for a system initially prepared in the ground state ($|c_0|^2 = 1$):

$$|c_1(t)|^2 = \sin^2 \Omega_0 t/2 \tag{5.2}$$

with $\Omega_0 = \mathbf{d} \cdot \mathbf{E}/\hbar$, which is referred to as the Rabi frequency. The population therefore oscillates sinusoidally between the ground and excited states. This equation introduces the π-pulse - this is the pulse area ($\Omega_0 t_\pi$) that deterministically excites the electron from the ground to the excited state with 100% probability. For a clock transition, this is the pulse that maximises the signal-to-noise.

This expression can be generalised to include the effect of near-resonant radiation:

$$\Omega(\omega) = \sqrt{\Omega_0^2 + (\omega - \omega_0)^2} \tag{5.3}$$

where ω is the frequency of the probe laser and ω_0 is the centre of the atomic resonance.

If the excitation π-pulse is rectangular in time (i.e. probe power rises sharply to a constant value before being sharply extinguished) then the excitation profile in frequency space [103] is given by the Fourier transform:

$$|c_1(\omega)|^2 = \left(\frac{\Omega_0 t_\pi}{2}\right)^2 \operatorname{sinc}^2\left(\frac{\Omega(\omega) t_\pi}{2}\right) \tag{5.4}$$

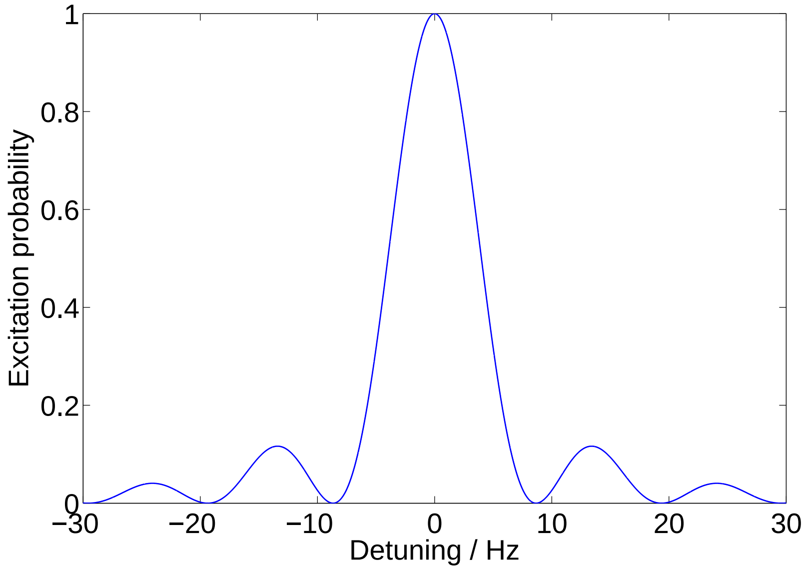

Figure 5.1: Simulated Rabi lineshape generated by a pulse length of 100 ms.

This assumes that the linewidth of the interrogating laser is much less than the resulting 'Fourier-limited' sinc^2 lineshape. A sample lineshape is shown in figure 5.1, and an experimentally-obtained lineshape is shown in figure 5.2.

5.1.2 Ramsey excitation

An alternative form of excitation which is in common usage in microwave frequency standards is the application of two rectangular pulses separated in time. We have not yet implemented this form of interrogation sequence, though it is a desirable upgrade as discussed in appendix A.

5.2 Decoherence

Unfortunately, it is not possible to apply longer and longer pulses and maintain 100% excitation. There are always sources of decoherence that act to damp out the Rabi oscillations at long timescales. In the presence of a decoherence rate Γ, the evolution of a two-level system can be described by the optical Bloch equations. In general, these equations can only be solved numerically. However, in a two-level atom the solution is

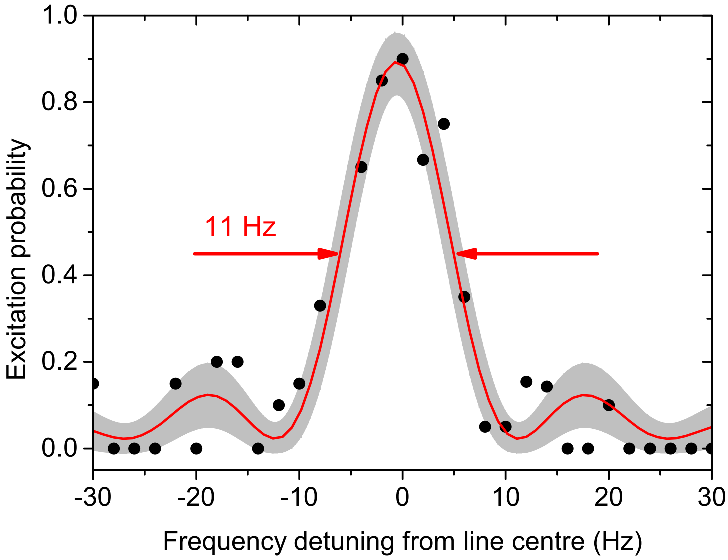

Figure 5.2: Typical scan over the octupole transition with a probe time of 100 ms, with twenty interrogations at each frequency detuning. A lineshape has been fitted in accordance with the rectangular nature of the probe pulse, showing a linewidth close to the Fourier limit of 9 Hz. The shaded area indicates the quantum projection noise at each point on the fit. There is a small probability of incorrectly registering a quantum jump, which leads to an offset on the y-axis and a non-zero error at the excitation minima. Decoherence has also been included, leading to a supression of the maximum excitation and a non-zero error at that point.

well described by [104]:

$$|c_1(t)|^2 = \frac{1}{2}\left(1 + e^{-\Gamma t}\sin^2 \Omega t\right) \tag{5.5}$$

hence the Rabi 'flopping' is suppressed exponentially to a steady-state excitation rate of 50%.

5.2.1 Atomic lifetime

The $^2F_{7/2}$ state has a lifetime τ_F of six years, and therefore the spontaneous decay out of the state during the probe time of 100 ms is likely to be negligible. On the other hand, the $^2D_{3/2}$ state has a lifetime τ_D of only 52.7 ms. If we consider spontaneous decay to the initial state to be the only source of decoherence, after a resonant π-pulse of duration t the resulting population in the excited state will be:

$$|c_1|^2 = P_D = \frac{1}{2}\left(1 + e^{-t/\tau_D}\right) \tag{5.6}$$

For a pulse of 30 ms duration on the 436 nm transition, the maximum excitation on resonance would be 78%. However, decay to the $F = 1$ component of the ground state is also possible, meaning that the ion can be completely lost from the clock transition. Hence there is a trade-off between the Fourier-limited transition linewidth and the maximum excitation efficiency. The linewidth yielding the lowest possible instability for this transition has been evaluated to be 10 Hz by PTB [105].

5.2.2 Laser noise

The decoherence mechanism that is more relevant to the 467 nm transition is the finite coherence time of the probe laser. For a Lorentzian laser noise spectrum, the coherence time [106] is given by:

$$\tau = \frac{1}{\pi \delta \nu} \tag{5.7}$$

where $\delta\nu$ is the linewidth of the laser. For a typical probe laser linewidth of 1 Hz this gives a coherence time of 318 ms, and hence after a probe of 100 ms duration the maximum excitation rate is 86.5%.

5.3 Heating rate and dephasing

The ion occupies a distribution of vibrational levels after Doppler cooling, as discussed in section 2.2. An increase in vibrational quantum number corresponds to a weaker coupling of the probe laser to the transition. In one dimension, the Rabi frequency for the transition that changes the vibrational quantum number from n' to n [107] is given by:

$$\Omega_{n',n} = \Omega e^{(-\eta^2/2)} \left(\frac{n_<!}{n_>!}\right)^{1/2} \eta^{|n'-n|} L_n^{|n'-n|}(\eta^2) \tag{5.8}$$

where Ω is the total coupling strength, η is the Lamb-Dicke parameter of the ion, $n_>$ and $n_<$ are the greater and lesser of the two vibrational quantum numbers respectively, and L_n^α is the generalised Laguerre polynomial of rank n. In the case of a pulse which does not change the vibrational quantum number ($n_> = n_<$) [104], this expression simplifies to:

$$\Omega_n = \Omega e^{(-\eta^2/2)} L_n(\eta^2) \tag{5.9}$$

In a cylindrically symmetric trap, the ion has two motional degrees of freedom. The Rabi frequency for the motional state n_r, n_z is given by:

$$\Omega_{n_r,n_z} = \Omega e^{(-\eta_r^2/2)} L_{n_r}(\eta_r^2) e^{(-\eta_z^2/2)} L_{n_z}(\eta_z^2) \tag{5.10}$$

When considering the effect of the probe laser, it is necessary to average over all of the possible vibrational states with their probability of occupancy P_n [104]. In practice, the result of this averaging is a suppression of the Rabi oscillations:

$$|c_1|^2 = \frac{1}{2}\left(1 - \sum_{n_r,n_z} P_{n_r} P_{n_z} \cos(\Omega_{n_r,n_z} t)\right) \tag{5.11}$$

This suppression is shown in figure 5.3a. In addition, the ion is heating through the vibrational levels during the probe pulse. This leads to a smearing of the thermal distribution during the interrogation that acts to suppress the observed average excitation, as shown in figure 5.3b.

5.4 Choosing the quantisation axis direction

The coupling of the driving laser to the transition depends on the angle θ between the quantisation axis $\mathbf{B}$ and the propagation vector $\mathbf{k}$ of the interrogating laser, and the angle ϕ between the Bk plane and the polarisation of the laser $\mathbf{E}$ (as shown in figure 5.4). When measuring the transition frequencies, it is necessary to drive the transition under three perpendicular orientations of $\mathbf{B}$ in order to average out any quadrupole shift (this is elaborated upon in chapter 7).

5.4.1 Quadrupole transitions

The optimum values of θ and ϕ are different for different values of Δm_F. The geometric factor governing the line strength of the $\Delta m_F = 0$ component is [108, 109, 110]:

$$Q_0 = 6 \sin^2 \theta \cos^2 \theta \cos^2 \phi \tag{5.12}$$

It is acceptable to have very different line strengths for the three orientations of $\mathbf{B}$ due to the comparatively high optical power available from the probe laser. For directions with a weaker line strength, the laser power can usually be increased accordingly without introducing an appreciable ac Stark shift on the transition frequency.

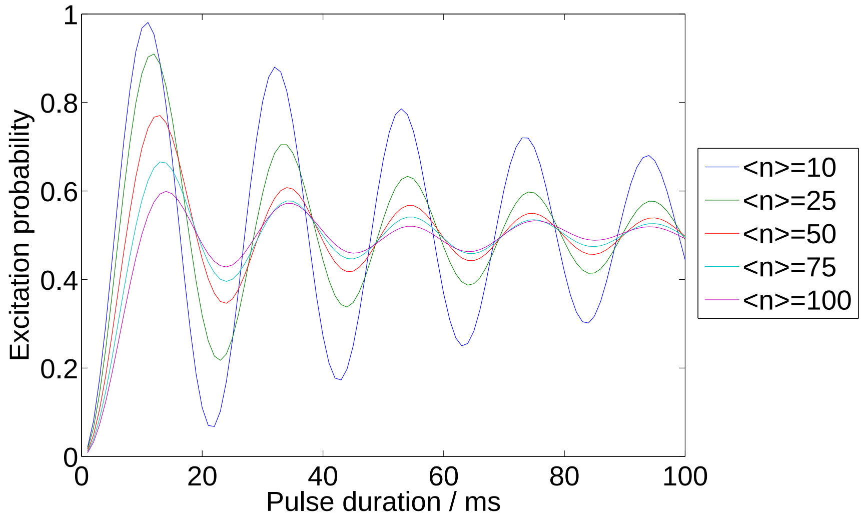

(a)

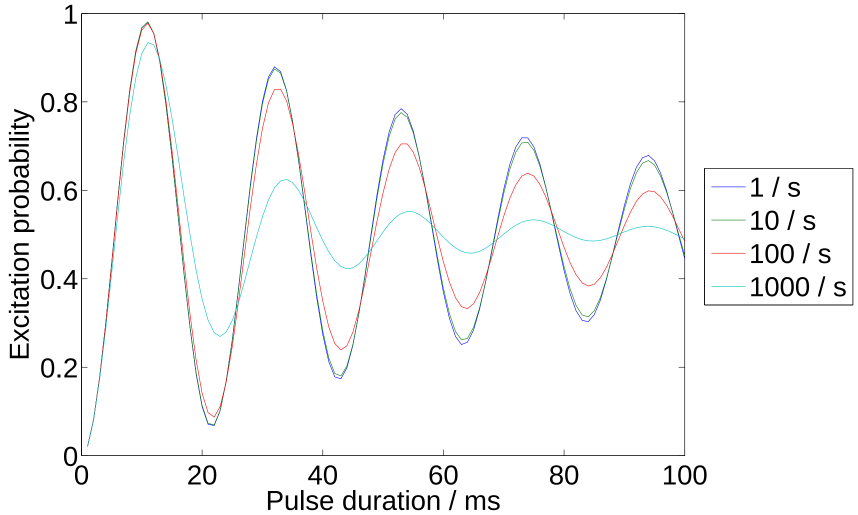

(b)

Figure 5.3: Comparison of a) dephasing as a function of Doppler cooling efficiency and b) dephasing as a function of radial heating rate in quanta/s, with a constant starting value of $\langle n \rangle = 10$. The Rabi frequency for $\langle n \rangle = 0$ is $2\pi \times 50$ Hz, the radial secular frequency is 1 MHz and the radial Lamb-Dicke parameter $\eta_r = 0.092$. The ion is probed in the horizontal plane and so $\eta_z = 0$.

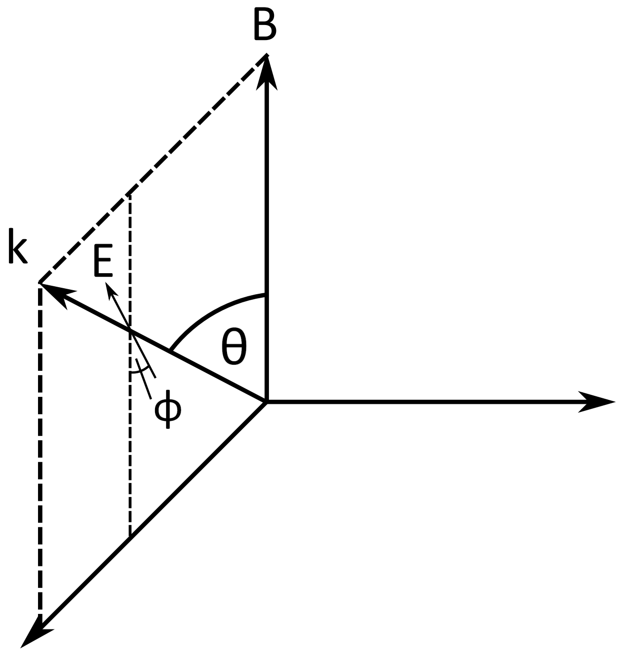

Figure 5.4: Depiction of the co-ordinate system used in the geometric factors governing the transition line strengths. θ is the angle between **B** and **k**, and ϕ is the angle at which **E** is tipped out of the Bk plane.

5.4.2 Octupole transitions

The equivalent expression for the line strength of the $\Delta m_F = 0$ component of an electric octupole transition is:

$$R_0 = 12 \sin^2\theta \cos^2\phi (1 - 5\cos^2\theta) \tag{5.13}$$

For the 467 nm transition, it is no longer acceptable to utilise a magnetic field direction for which the coupling of the probe laser to the transition is weak. The intensities required to drive octupole transitions are orders of magnitude higher than for quadrupole transitions, and are high enough to cause an easily resolvable ac Stark shift on the transition frequency (see chapter 7). Increasing the laser power would cause an increase in the ac Stark shift, and hence a set of three fields must be found with high line strength (and ideally, the strength should be near-equal in the three directions). The problem is complicated further by the lack of correspondence between the spherical polar co-ordinate system of the trap and the definitions of θ and ϕ in the expression for the line strength.

The two parameters that are known about the 467 nm probe beam are that it is angled at 25° to the trap y-axis, and the polarisation is tipped out of the xy plane by 15°. From this, it is possible to derive a propagation vector $\mathbf{k}$ in the co-ordinate system of the trap, and also a vector for the polarisation $\mathbf{E}$ in the same basis.

Now, an arbitrary B-field $\mathbf{B}$ is chosen in the same basis. θ can now be extracted from the dot product of $\mathbf{B}$ with $\mathbf{k}$.

The Bk plane is defined by the cross product of $\mathbf{B}$ with $\mathbf{k}$. It is now possible to extract a value of ϕ from the dot product of this vector with $\mathbf{E}$.

Now that the problem is set up, it can be solved numerically to find a set of three fields with the following constraints:

- Fields should be orthogonal to within 1°
- Line strengths should be equal to within 20%
- Line strengths should be as high as possible

From this, it is possible to find a set of three fields with line strengths greater than 25%, that are mutually orthogonal to better than 1°.

5.4.3 Optimising both transitions at once

The octupole and quadrupole probe lasers are orthogonally polarised at the ion, and so the above calculation must be repeated for the quadrupole probe laser with the laser polarisation tipped out of the xy plane by 105°. Remarkably, it is possible to find a set of three fields that produce near-equal, high line strengths for the octupole transition whilst simultaneously having reasonable (greater than 1%) line strengths for the quadrupole transition, that differ by no more than a factor of ten. This is invaluable for measuring a single-ion frequency ratio as the quadrupole shift can then be averaged out on both transitions simultaneously.

5.4.4 Realising the desired magnetic field direction

The three sets of Helmholtz field coils were calibrated by applying various currents through each of them in turn, and then observing the Zeeman splittings of the $\Delta m_F = \pm 1, \pm 2$ components of the 436 nm transition. The coils could then be characterised through:

$$B_x = \frac{dB_x}{dI_x}(I_x - I_{0x}) \tag{5.14}$$

$$B_y = \frac{dB_y}{dI_y}(I_y - I_{0y}) \tag{5.15}$$

$$B_z = \frac{dB_z}{dI_z}(I_z - I_{0z}) \tag{5.16}$$

$$|B| = \sqrt{B_x^2 + B_y^2 + B_z^2} \tag{5.17}$$

This allowed a precise determination of the currents I_{0i} required to cancel any ambient fields (such as that of the Earth, or that of the ion pump magnet attached to the trap vacuum system), and the rate of change of the field with respect to current through the coil pairs dB_i/dI_i. Small additional currents I_i could then be applied to generate a bias field in the required direction. This makes the implicit assumption that the fields from the three pairs of coils are completely orthogonal, but this has been found to be a good assumption.

5.5 Summary

The coherent interaction between a monochromatic probe laser and forbidden transition has been discussed. The transition line strengths and excitation rates can now be optimised and the decoherence rates understood. The following chapter discusses how the excitation profiles can be used to stabilise the frequency of the probe laser even further.

Chapter 6

Locking to the clock transition

The optical cavity provides the system with a short-term instability of around 1×10^{-15}. The stability is degraded at longer timescales due to linear drift and random walks in the cavity frequency. This is where an atomic reference transition becomes necessary.

Previous absolute measurements of the clock transition frequencies relied on scanning the probe laser over the absorption line and fitting the observed lineshape to determine its centre. This resulted in very low data rates that limited the statistical uncertainty on the measurements.

The data rate can be hugely increased by actively steering the laser back onto the transition. With a well-tuned feedback loop, and if the systematic shifts on the transition are well controlled, the long-term stability of the transition frequency is transferred onto the laser.

Section 6.1 contains the details of the experimental sequence and how the ion's internal state is read out. Section 6.2 discusses the generation of the error signal for the lock, and section 6.3 introduces the servo algorithm used to correct the laser frequency. The behaviour of the servo was simulated, with the details in section 6.4. The servo was then implemented, and the preliminary data are discussed in section 6.5. Some preliminary measurements of the servo performance by the comparison of two independent traps are given in section 6.6.

6.1 Experimental control

The experiment is mostly automated using specially written LabVIEW software. This interfaces with the experimental hardware including frequency synthesisers, laser shutters and a fluorescence counter.

6.1.1 Experimental sequence

Cooling, state preparation, probing, and detection of the ion's state all require different combinations of lasers to be applied to the ion. The unwanted beams can be extinguished using AOMs, though a good rf switch is required and there may still be some small residual diffraction or scatter. Even extremely low light levels from the other resonant lasers can cause appreciable ac Stark shifts on the clock transitions, hence all other lasers must be completely extinguished during the probe phase. This is performed by small mechanical shutters, constructed by spot-welding small razor blades onto electronic relay switches. The state of the shutter can be changed using a 5 V TTL signal. There is a typical delay of 2 ms between a signal being sent to the shutter and the blade beginning to move. The lasers are focused to waists of around 50 μm at the shutter, so the transit time is short (< 100 μs) and the extinction is high. Care was taken to make sure there is no 'ringing' of the shutter once it has reached the end of its travel. As the lasers are delivered to the trap via optical fibre, there is also a suppression of any small levels of scatter from the blade itself.

The relays are controlled using a National Instruments digital output card. An array of digital values, one corresponding to each laser shutter, is pre-programmed for each stage of the experimental cycle. As the experiment runs through its various phases, the appropriate lasers are applied to the ion. A summary of the various stages is shown in table 6.1.

Laser / Device	Cooling	State prep	Delay	Probe	Detection	Repump
370 nm (14 GHz)	1	0	0	0	1	1
370 nm (2 GHz)	0	1	0	0	0	0
935 nm (F=1)	1	1	0	0	1	1
935 nm (F=2)	1	1	0	0	0	1
639 nm	1	1	0	0	0	1
Probe laser	0	0	0	1	0	0
Counter	0	0	0	0	1	1

Table 6.1: Digital signals sent to the laser shutters and fluorescence counter during the experimental sequence.

6.1.2 Detection of the ion's internal state

It is relatively easy to determine whether an atom has undergone a strong cycling optical transition as it will scatter millions of fluorescence photons every second, some of which can be detected by a camera or photomultiplier tube. It is more difficult to determine whether the atom has undergone a weak 'clock' transition as the photon scatter rate can be more than six orders of magnitude smaller for an electric quadrupole transition. Instead, the state of the atom is detected indirectly by making use of the long lifetime of the upper state of the clock transition. The 'probe' radiation is briefly applied to the atom to interrogate the clock transition. Afterwards, a strong 'detection' laser (which couples to the strong cycling transition) is applied. If the clock transition has been driven, then the detection laser will not interact with the atom and no photons are scattered (as shown in figures 6.1 and 6.2). This is referred to as electron shelving or the quantum jump technique [111].

After probing, the ion's state is determined by applying the lasers necessary to observe fluorescence and starting the counter to detect the number of photons scattered by the ion in a given detection time. The state of the ion is decided by comparing the number of detected photons against a threshold level T set to be an equal number k of standard deviations away from the expected number of signal (n_s) and background (n_b) photons:

$$T = n_b + k\sqrt{n_b} = n_s - k\sqrt{n_s} \tag{6.1}$$

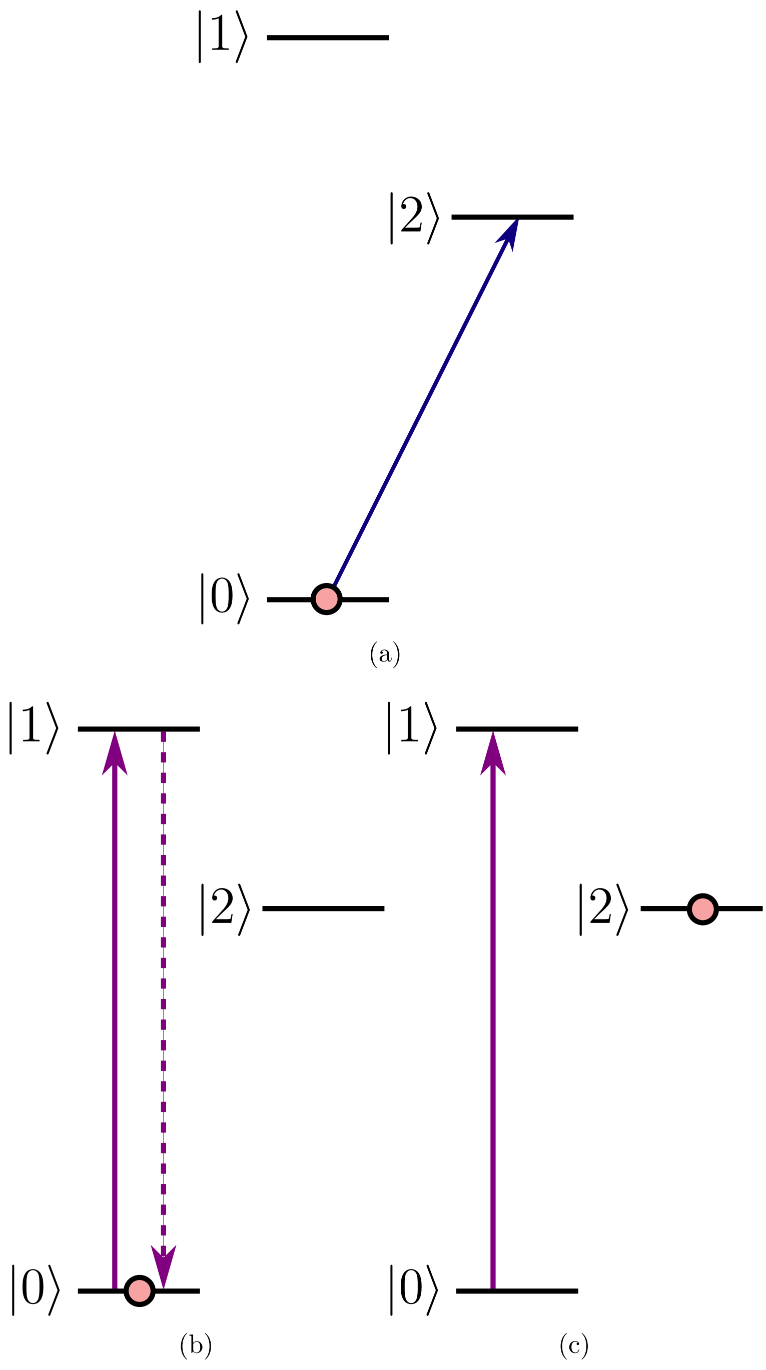

Figure 6.1: Simplified depiction of quantum jump detection in a 3-level system. a) ‘Clock’ probe, attempting to excite the ion from the ground state $|0\rangle$ to the metastable state $|2\rangle$. b) Probe unsuccessful, ion scatters many photons from the detection laser tuned to the strong $|0\rangle \rightarrow |1\rangle$ transition. c) Probe successful, ion is shelved in $|2\rangle$ and hence scatters no photons from the detection laser.

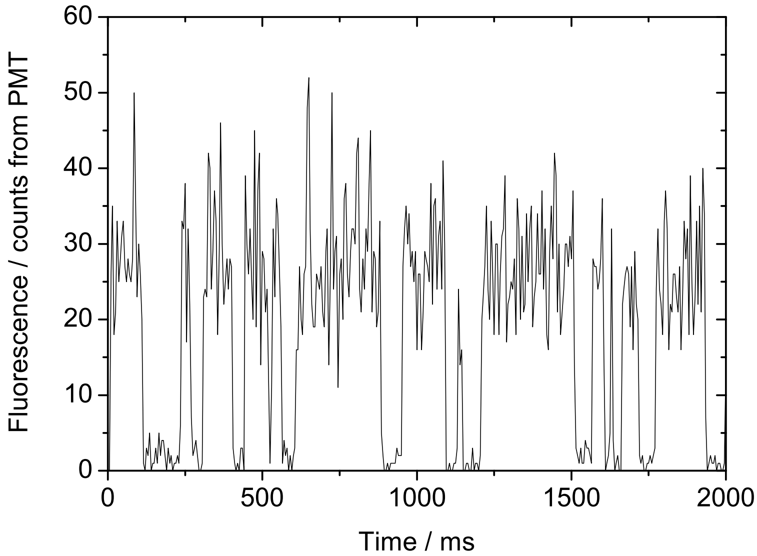

Figure 6.2: Quantum jumps observed with a single ion of $^{171}Yb^{+}$. During the dark times, the ion has been 'shelved' in the $^{2}D_{3/2}$ state and does not interact with the strong detection laser.

$$k = \frac{n_s - n_b}{\sqrt{n_s} + \sqrt{n_b}} \tag{6.2}$$

If fluorescence was observed during the detection phase, the sequence is reset. However, if a quantum jump was observed then the appropriate repumper lasers are applied and the ion's state detected again. Only once fluorescence is observed does the sequence reset. This ensures that the ion has been re-initialised before the next probe, as there may be a finite delay between the repump lasers being applied and the ion returning to the ground state (for example, if the 639 nm lasers are not tuned correctly it may take some time for the $^{2}F_{7/2}$ state to be cleared out).

The detection pulse length is typically chosen to give $n_s \approx 20$, in order to give a readout fidelity $> 95\%$. During the measurement of the E2 transition frequency, a bias field of 300 μT was applied during state detection. The necessary detection pulse duration was 5 ms to achieve the desired readout fidelity. However, the fluorescence rate during detection was substantially reduced after the introduction of polarisation spinning for the measurement of the E3 transition frequency. Initally, a 50 ms detection window was

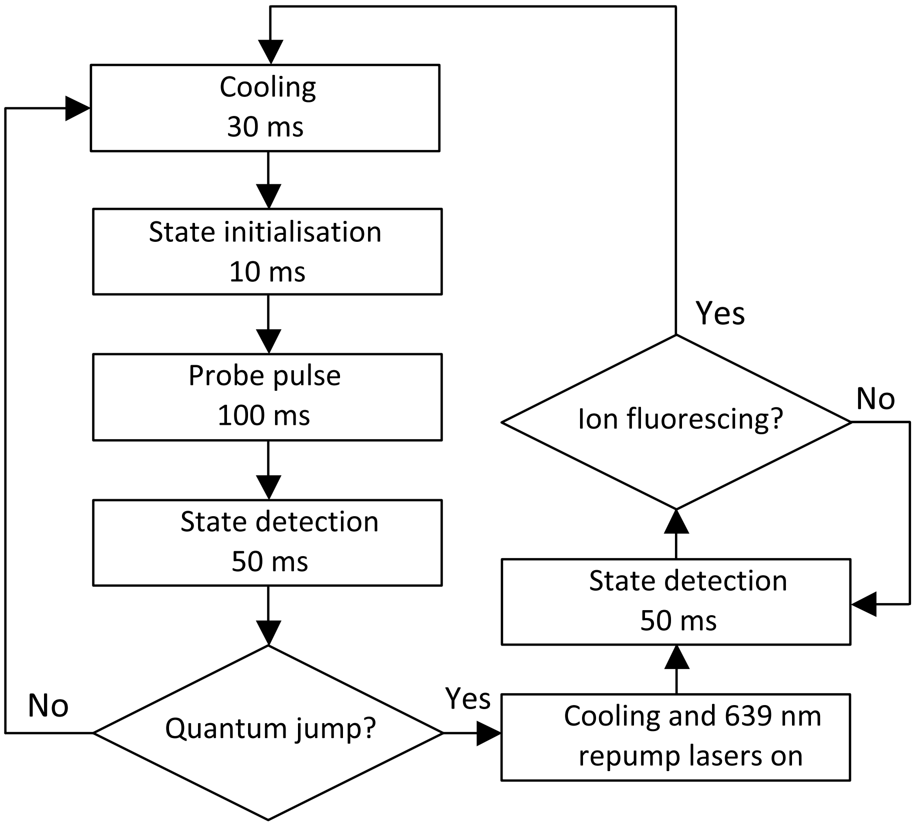

Figure 6.3: Depiction of the experimental sequence during the 2011 absolute measurement of the 467 nm transition frequency. If no fluorescence is detected after the probe phase, the bin is marked as a quantum jump and the repumper lasers are applied. Repeated measurements of the state of the ion are made until a fluorescence signal is detected, at which point the sequence is reset. Typically only one repump phase is required. The sequence runs several times at each probe laser frequency in order to determine an excitation probability.

required in order to give a good readout fidelity. This did not present a problem during this absolute frequency measurement, as the long lifetime of the F-state meant that excited state population was not lost during this time. However, such a long detection pulse is unacceptable for future measurements of the E2 transition frequency for which the excited state lifetime is comparable to this duration. The necessary detection window has since been reduced to around 10 ms by reducing the background scatter rate.

6.1.3 Tuning the probe lasers

The probe lasers can be tuned in frequency by sending GPIB commands to change the frequency of the synthesiser used to drive an AOM between the laser and the ion. Typically when initially scanning for the clock transition, twenty interrogations are made at each frequency detuning and then the synthesiser is updated.

6.2 Generating the error signal

As is the case when stabilising the laser to an optical cavity, it is not possible to determine the direction in which the laser has drifted from an excitation rate at line centre. Sitting on the side of the fringe produces a directional discriminant, with a signal-to-noise given by:

$$\mathrm{S/N} = \frac{np}{\sqrt{np(1-p)}} \tag{6.3}$$

where $\sqrt{np(1-p)}$ is the standard deviation of a binomial distribution. This discriminant is maximised for $p = p_{\mathrm{max}}/2$. Sitting on one side of the fringe requires excellent knowledge of the transition half-width, as this must be added to the observed frequency in post processing. Fluctuations in the width of the transition due to fluctuations in the laser intensity (power broadening) or changes in the probe time or pulse shape (changing the Fourier-limited full-width, $\delta\nu$) can lead to systematic uncertainties. In addition, the capture range of the servo is limited as the error signal will change sign for detunings on the other side of line centre. This makes the servo vulnerable to sudden changes in laser frequency (due to a sudden glitch in the cavity frequency or lock point), or when it received 'bad' information from the ion due to the low signal-to-noise.

It is more sensible to directly derive the centre frequency of the transition. This is done by probing on *both* sides of the transition and using the imbalance in the quantum jump rates to determine how far the average frequency of the two probes is from the line centre [105, 112] (as shown in figure 6.4). This probe method has the added advantage of a much larger capture range and no loss in directionality - even if both probes are on the same side of the transition, the imbalance still correctly predicts the direction in which the laser frequency should be changed. The error signal, e, is given by:

$$e = \frac{n_+ - n_-}{n_+ + n_-} \qquad \text{for } n_+ + n_- \neq 0 \tag{6.4}$$

where n_+ and n_- are the numbers of quantum jumps on the high and low frequency sides respectively. The denominator acts as a normalisation term, making the error signal appropriate for any number of interrogations or excitation rates. This makes the error signal equivalent to:

$$e = \frac{p_+ - p_-}{p_+ + p_-} \qquad \text{for } p_+ + p_- \neq 0 \tag{6.5}$$

where p_+ and p_- are the transition probabilities at the two detunings. For $n_+ = n_- = 0$, e is indeterminate and so a value of $e = 0$ is used in this situation.

If it were possible to probe an unlimited number of times at each of the two frequency detunings, then the error signal would look like figure 6.5. Under more reasonable experimental conditions of ten pulses at each detuning, and assuming quantum projection noise is absent, the error signal takes the form of figure 6.6.

6.3 Feeding back to the laser

As the probe laser frequency is drifting at a rate of around 0.1 Hz/s, simply using a servo algorithm with proportional gain to steer the laser back onto the line will lead to an offset between the laser and line centre given by:

$$\nu - \nu_0 = r t_{servo} \tag{6.6}$$

where r is the laser drift rate and t_{servo} is the time between successive corrections to the laser frequency via an acousto-optic modulator. This offset can be eliminated by the

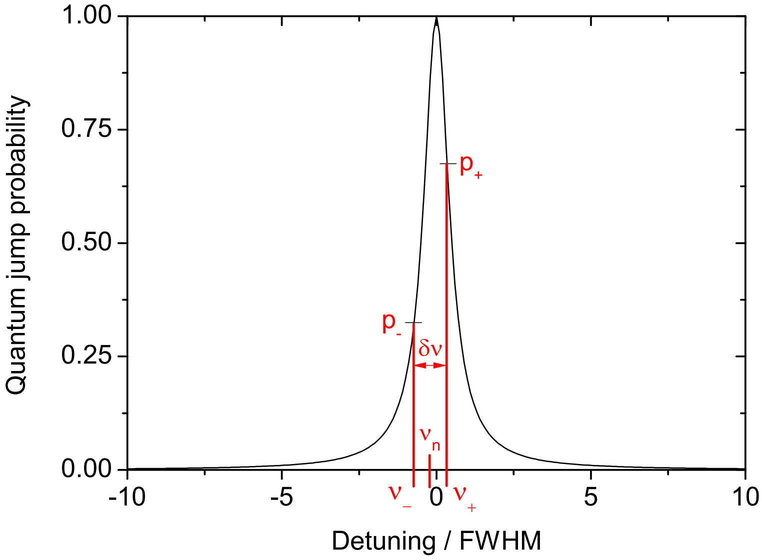

Figure 6.4: Illustration of how the quantum jump imbalance is derived. A Lorentzian function has been used for simplicity.

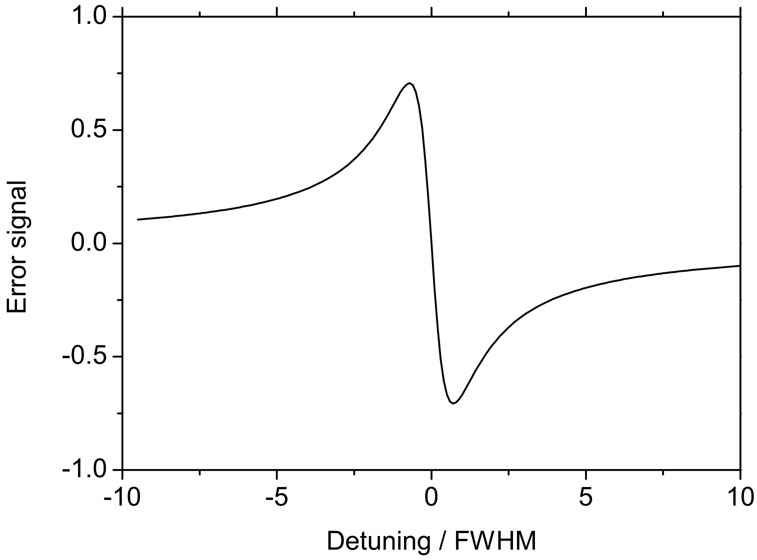

Figure 6.5: Error signal generated if it were possible to probe a large number of times at each frequency detuning.

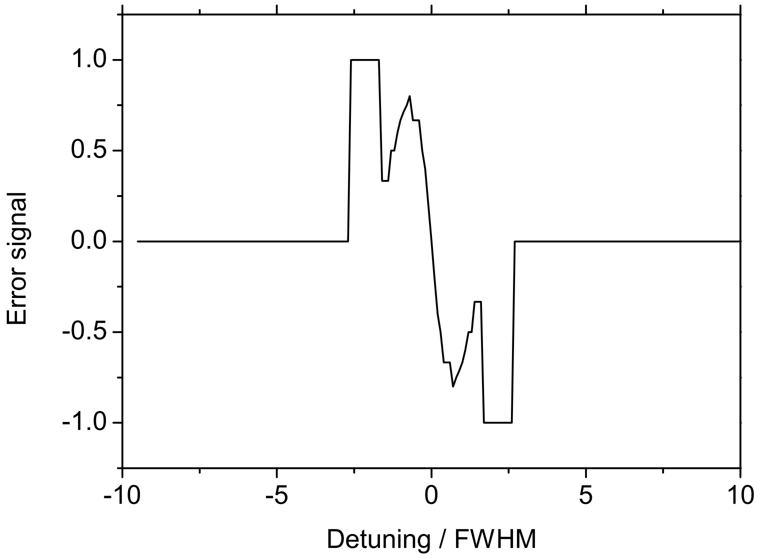

Figure 6.6: Error signal generated when the ion is interrogated ten times at each frequency detuning. The associated quantum projection noise is not shown. The somewhat odd behaviour at detunings of ± 2 FWHM is a consequence of the normalisation process: observing a single quantum jump leads to an error signal of ± 1 from equation 6.4.

introduction of an integral term to the feedback loop. As the detection cycle contains an additional delay after each quantum jump to ensure a return of the ion to the cooling cycle, each servo iteration will differ in length based on the total number of jumps in that iteration. Therefore it is sensible to include a time-dependence to the integral correction to allow for the fluctuations in cycle length. The final servo looks like:

$$\nu_{n+1} = \nu_n + \delta\nu \left(g_1 e_n + g_2 \Delta t_n \sum_{i=0}^{n} e_i \right) \tag{6.7}$$

where $\delta\nu$ is the FWHM of the resonance, g_1 and g_2 are gain parameters* and Δt_n is the time that has elapsed since the previous correction.

*Note that g_1 is dimensionless, but g_2 has dimensions of s^{-1}.

6.4 Simulations of lock behaviour

A LabVIEW routine was written to simulate the behaviour of the lock servo with a selection of values for the gain parameters g_1 and g_2. Initially laser noise was included in the model using LabVIEW's built-in $1/f^2$ and $1/f$ noise generators, but this greatly slowed down the simulation and so was removed. The inclusion of laser noise was likely to be unnecessary as during normal operation the laser should have a much narrower linewidth than the Fourier transform limit for the pulse durations used. The parameters that could be entered into the simulation were:

- transition linewidth (Fourier transform-limited)
- dead time
- number of pulses on each side of transition
- peak excitation rate
- optical cavity drift rate
- gain parameters

The simulation was designed to loop several times with each set of parameters and plot:

- servo offset from line centre
- Allan deviation of the offset

The stability of the lock was optimised by varying g_1, and the servo offset was minimised by varying g_2. Data for the optimised set of servo parameters $g_1 = 0.2$, $g_2 = 0.003$ s^{-1} are shown in figure 6.7. Naïvely, one might expect the optimum value of g_1 to be slightly above unity (i.e. an error signal of $\sim$0.75 corresponds to a frequency detuning of $\delta\nu/2$, and hence a correction of $\delta\nu/2$ is required to correct the fluctuation in a single step) but it was found that increasing g_1 makes the servo more unstable as it will react more violently to 'bad' quantum jump data from the low signal-to-noise extracted from the ion.

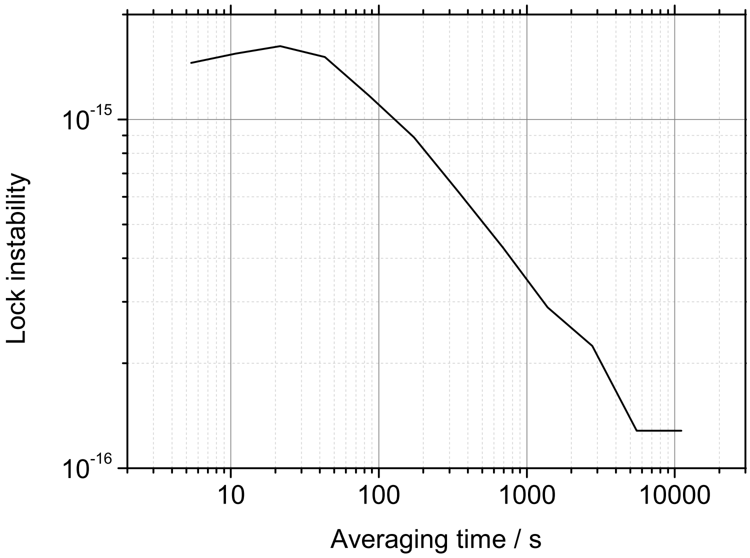

Figure 6.7: Simulation of the statistical noise on the 436 nm transition with typical experimental parameters (35% excitation probability, 30 ms pulse duration and 30 interrogations on each side of the transition leading to a total cycle time of 5.4 s). The optimised gain parameters were $g_1 = 0.2$, $g_2 = 0.003$. From a total of 100 simulations with lock periods of 48600 s, the mean fractional offset from line centre was $(3.8 \pm 8.0) \times 10^{-17}$.

6.5 Locking to the transition

The lock was first implemented on the 436 nm transition. The gain parameters derived from the simulation proved to be a good starting point, and the laser remained in lock for several hours. Some sample data is shown in figure 6.8.

6.6 Trap-trap comparisons

Once a stable lock had been demonstrated in a single trap, a second identical ion trap was brought online. Using two AOMs in independent beam paths, the probe laser was stabilised to the 436 nm transition in each trap. To avoid obscuring any laser noise by operating the locks in a synchronous way, the two lock routines were operated entirely independently using separate control programs. The gain parameters for the two locks and the number of interrogations on each side of the absorption line were iterated to make the locks more stable. It was found that increasing g_1 or decreasing the number of

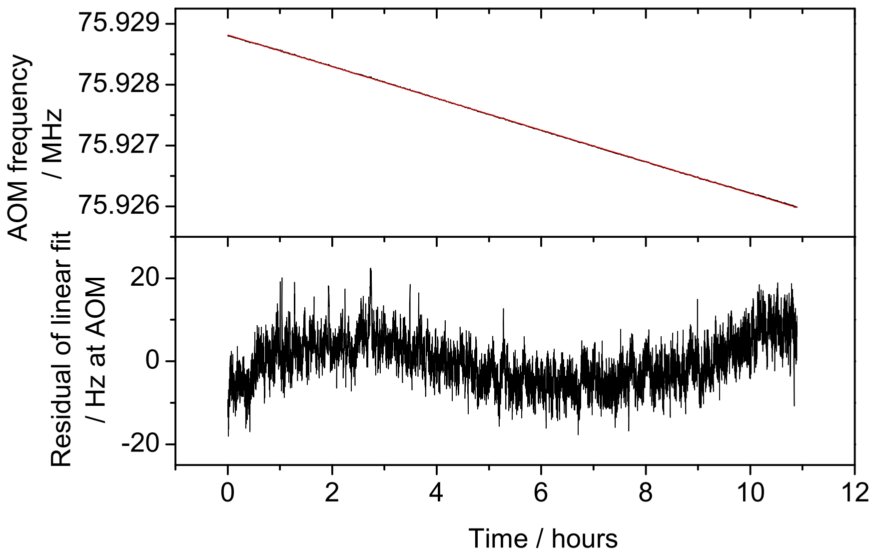

Figure 6.8: Upper: AOM drive frequency required to remain locked to the 436 nm transition. There is a factor of four between the AOM drive frequency and the resulting shift of the probe laser frequency. The drift of the ULE cavity is extremely linear over the entire dataset. Lower: Residuals to a linear fit to the AOM frequency as a function of time. Assuming a residual thermal expansivity of 1×10^{-9}/K, the resulting sinusoid could be explained by a thermal oscillation of the optical cavity with a peak deviation of 100 μK.

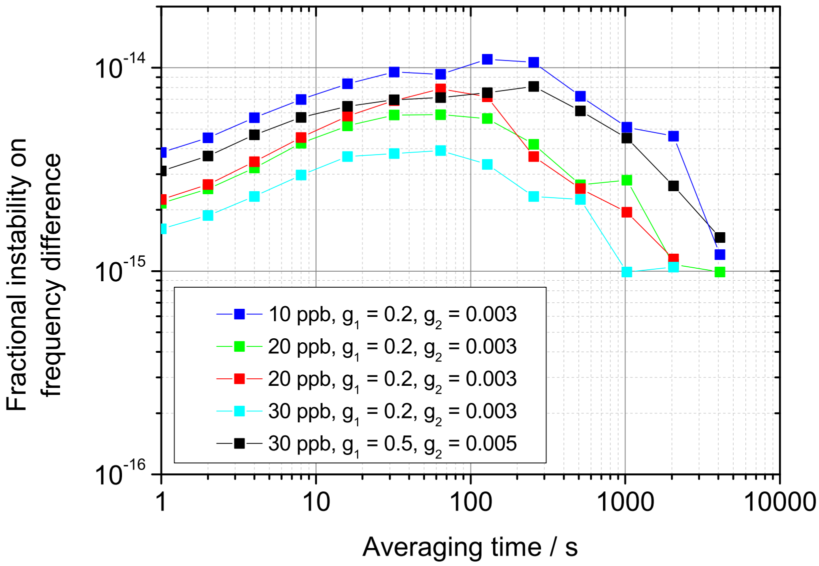

Figure 6.9: Allan deviation of the difference between the 436 nm transition frequency as measured using two separate ion traps. It can be seen that increasing the gain or decreasing the number of interrogations degrades the short-term stability. ppb = pulses per bin.

pulses on each side of the absorption line caused the servo 'hump' to rise (see figure 6.9), which is consistent with the servo becoming more unstable as discussed in section 6.4. Figure 6.10 shows the stability of the difference between the two AOM drive frequencies as a function of averaging time for the optimised servo parameters.

In the future, this technique will allow much faster evaluation of the systematic shifts on the two transition frequencies as the statistical noise on the data is around a factor of 100 below the caesium fountain at a given averaging time. It will also allow tests of the reproducibility of the $^{171}Yb^+$ standard by analysis of the level of agreement between the 'unperturbed' frequencies in the two traps.

6.7 Summary

The 436 and 467 nm lasers have been successfully locked to their respective transitions for extended periods. This has greatly reduced the statistical uncertainty on the measure-

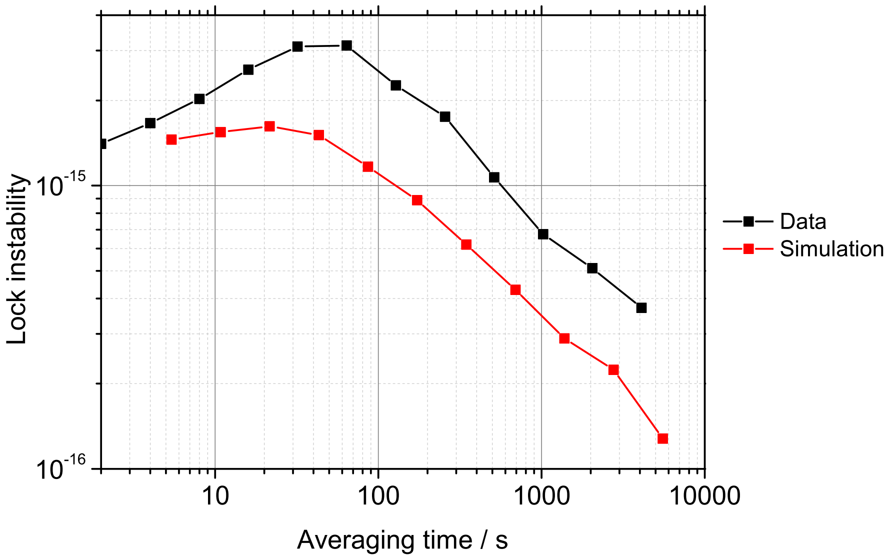

Figure 6.10: Stability of the E2 transition frequency in a single trap with optimised servo parameters, obtained by dividing the stability of the frequency difference between two separate ion traps by a factor of $\sqrt{2}$. Interrogating the transition 30 times on each side with a 30 ms pulse duration led to a cycle time of around 6 s. The servo 'hump' is clearly visible. The instability appears to be around a factor of two higher than that predicted by the statistical noise simulation. This may be due to additional short-term laser instability that is not included in the model.

ment of the transition frequencies, and has also allowed a reduction in the instability of the probe laser indicating successful operation as a frequency standard. Now that the statistical contribution to any measurement has been reduced, the systematic contributions must be considered.

Chapter 7

Systematic frequency shifts

Unfortunately it is not possible to isolate the ion from the environment completely. The presence of electric and magnetic fields acts to perturb the frequency of the clock transition through the Stark and Zeeman effects respectively. Relativistic effects associated with the ion's motion and presence in a gravitational potential must also be taken into account. Furthermore, there are shifts associated with the experimental apparatus including AOMs and optical fibres. In order to produce a frequency standard of the highest accuracy, it is necessary to evaluate all of these shifts thoroughly in order to determine the unperturbed transition frequency.

7.1 ac Stark shift (Light shift)

Presently, the largest systematic shift on the 467 nm transition is the ac Stark shift from the probe laser itself. On resonance, a probe laser does not shift its own transition. However, it can off-resonantly shift all of the other optical transitions within the ion, leading to a net shift $\Delta\nu_{\mathrm{Light}}$ [113] of the two states involved in the transition:

$$\Delta\nu_{\mathrm{Light}} = -\frac{1}{2h}\Delta\alpha_{\mathrm{sc}}^{\mathrm{ac}}\left\langle E^2\right\rangle - \frac{1}{4h}\alpha_{\mathrm{ten}}^{\mathrm{ac}}\frac{3m_F^2 - F(F+1)}{F(2F-1)}\left(3\cos^2\theta - 1\right)\left\langle E^2\right\rangle \tag{7.1}$$

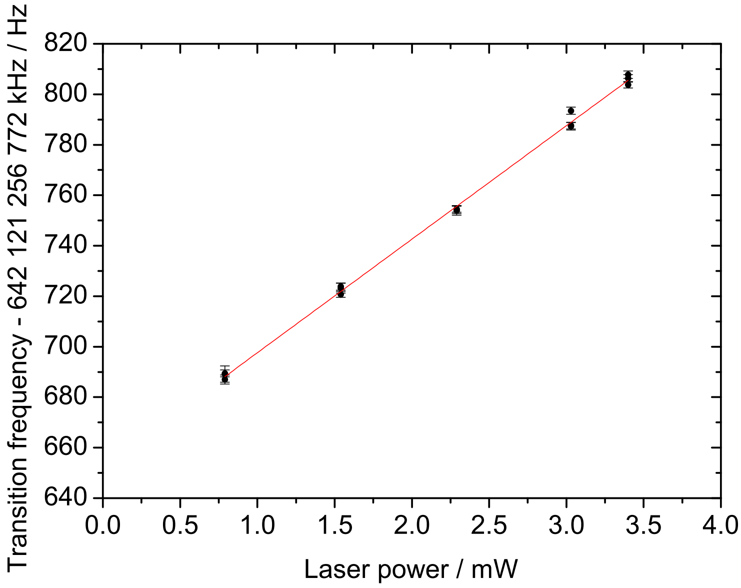

Figure 7.1: 467 nm transition frequency (not absolute) as a function of probe laser power measured against a hydrogen maser. The data were taken with probe pulses of 100 ms duration with the laser focused to a $1/e^2$ beam waist of 23 μm. The shift is approximately 45 Hz / mW, with a π-pulse requiring a laser power of 1 mW for this probe duration. The fractional uncertainty on the intercept at zero power is limited by the maser to 2×10^{-15}.

where $\Delta\alpha_{\rm sc}^{\rm ac}$ is the difference between the scalar dynamic polarisabilities of the two states involved in the transition, $\alpha_{\rm ten}^{\rm dc}$ is the tensor dynamic polarisability of the upper clock state*, $\langle E^2 \rangle$ is the mean-square electric field experienced by the ion and θ is the angle between the probe laser polarisation and the ion's quantisation axis. This effect is illustrated in figure 7.1.

The simplest method of correcting for the ac Stark shift is by driving the transition with two different probe intensities (I_1 and I_2) with a well-known relationship to one another (for example, a factor of two). The frequency of the transition at each intensity (f_1 and f_2) can then be used to determine the unperturbed frequency:

$$f_0 = f_2 - I_2 \frac{f_2 - f_1}{I_2 - I_1} = f_2 - \frac{f_2 - f_1}{1 - I_1/I_2} \tag{7.2}$$

This method requires excellent knowledge of I_1/I_2. It is usually not possible to measure

*The ground state has no tensor polarisability [113].

the laser spot size well enough to extract the intensity with the required level of uncertainty, but if the spot size is assumed to be constant then $I_1/I_2 \equiv P_1/P_2$. The laser power therefore must be stepped between two well known and well controlled values. The laser power in each case was measured using a calibrated photodetector.

7.1.1 Intensity noise due to frequency doubling

As the 467 nm light is produced by frequency doubling in a resonant cavity, care must be taken to avoid introducing amplitude noise onto the light. One of the main sources of noise in this process is the accidental driving of a system mechanical resonance, in particular the piezoelectric actuator and its mirror possess a resonance at 44.6 kHz. This resonance also limits the amount of gain that can be applied in the feedback loop, leading to imperfect suppression of low-frequency noise. Adding a notch filter to the output of the piezo's voltage source ('tweeter trap') prevents feedback at this system resonance and allows suppression of the low-frequency noise (as discussed in subsection 4.3.2).

7.1.2 467 nm power servo

Air currents and temperature drifts in the laboratory can cause changes in the beam alignment into the optical fibre that delivers the light to the trap, leading to additional low-frequency fluctuations in the probe laser power. For this reason, a small sample of the probe light is picked off after the fibre and monitored using a photodiode. This signal is compared against a stable reference voltage (see figure 7.2) in order to produce an error signal. This is then integrated and fed back to the drive power of the AOM that is also used to stabilise the laser to the clock transition frequency. When referenced to a home-built stable voltage source, the resulting probe power has a fractional instability below 1×10^{-3} as shown in figure 7.3. The long-term stability of the laser power was improved by replacing the voltage divider with a National Instruments analog input/output card (NI-6008).

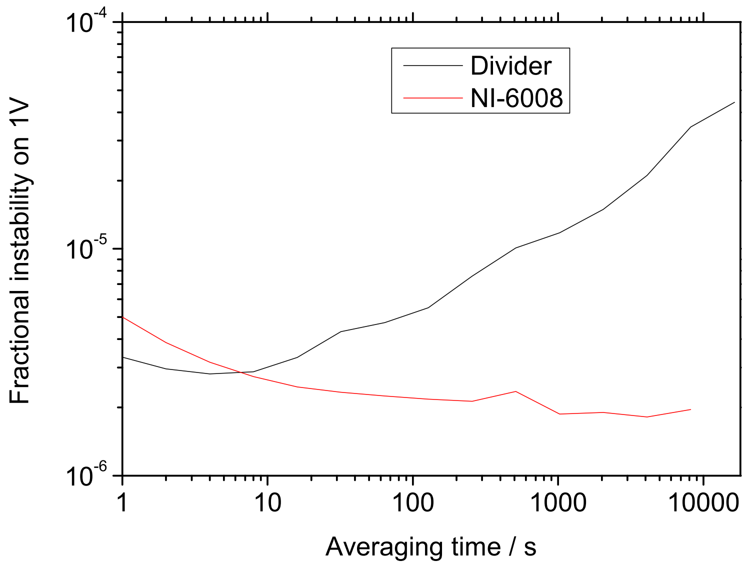

Figure 7.2: Allan deviation of the output of two voltage references used to stabilise the 467 nm laser power: a home-built voltage divider and a National Instruments NI-6008 analog/digital input/output card.

This servo must have as small a capture and settle time as possible, as any offset from the set point will become an inaccuracy in the power ratio, and any settle time will lead to a distortion of the transition lineshape. The servo capture time was reduced to around 150 μs (as shown in figure 7.4) by careful optimisation of the servo gain profile, with the capture time now limited by the bandwidth of the AM input of the synthesiser used to drive the AOM.

7.1.3 Tensorial shift

In addition to the scalar shift which depends purely on laser intensity, there is a tensorial component that depends upon the relative angle between the laser polarisation and the quantisation axis of the ion. If the ambient magnetic field were to drift, then the absolute magnitude of the shift would be observed to change. For this reason it is desirable to carry out the extrapolation to zero intensity on as short a timescale as possible. Previously the laser power was changed every twenty minutes, resulting in an extrapolated frequency every forty minutes. Recent changes to the interrogation sequence now allow alternate

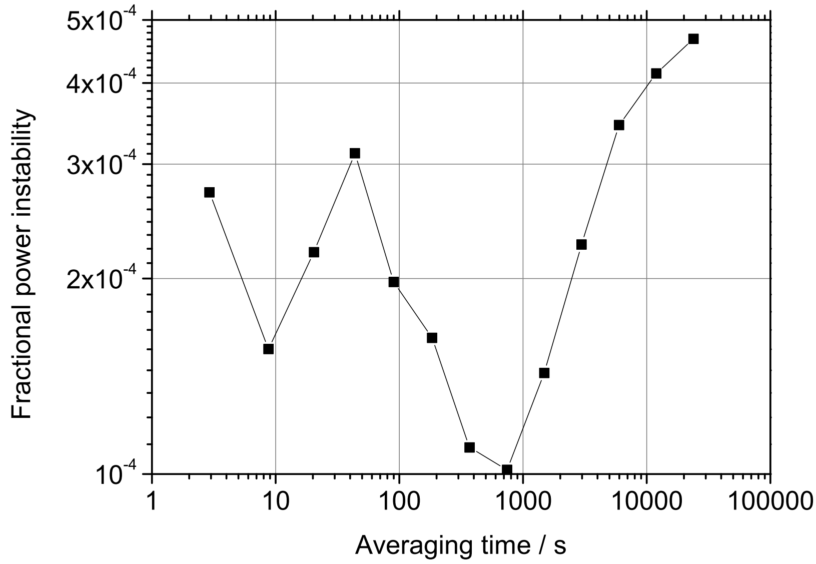

Figure 7.3: Allan deviation of the 467 nm laser power at the output of the optical fibre delivering the light to the trap, as measured by a photodiode, with the power stabilisation loop closed.

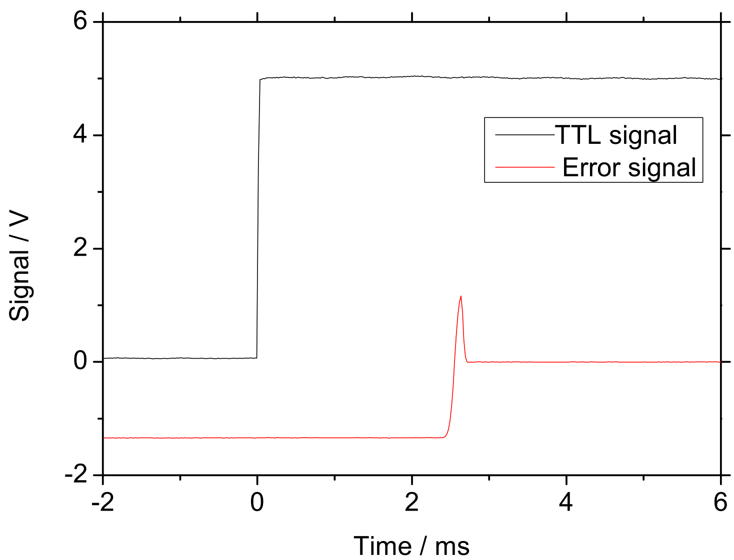

Figure 7.4: Capture time of the servo used to stabilise the 467 nm laser power at the trap. The 2.5 ms delay after the TTL edge is the response time of the mechanical shutter, with the servo settling in the following 150 μs.

interrogations at high and low laser power settings, resulting in an extrapolated frequency to be produced every cycle (~ 10 s).

7.2 Second-order Zeeman shift

The first-order Zeeman shift of a particular state with quantum numbers F and m_F is given by:

$$\Delta\nu_B = \frac{\mu_B B g_F m_F}{h} \tag{7.3}$$

where μ_B is the Bohr magneton, B is the magnetic flux density, h is Planck's constant and g_F is the g-factor of the state. The $m_F = 0 \rightarrow m_{F'} = 0$ components of the two clock transitions will therefore have no linear dependence on the magnetic flux density (and more importantly, on its fluctuations), which is an attractive feature for an optical frequency standard.

Treating the problem with second-order perturbation theory [114] reveals that there is a second-order shift that depends on the square of the magnetic flux density:

$$\Delta\nu_{B^2} = k\left(\left\langle B_{\mathrm{dc}}^2\right\rangle + \frac{\langle B_{\mathrm{ac}}^2\rangle}{2}\right) \tag{7.4}$$

It is possible to derive an expression for the shift coefficient k by assuming LS-coupling is valid for the clock states, though this has been shown to be in disagreement with the experimentally derived values in the case of the 467 nm transition. The values of k have been calculated to be 52 mHz/μT^2 for the 436 nm transition [115] and measured to be $-1.72(3)$ mHz/μT^2 for the 467 nm transition [116].

The magnitude of the magnetic field flux density was measured by observing the Zeeman splittings of the $\Delta m = \pm 1$ and ± 2 components of the 436 nm transition, as shown in figure 7.5.

There is likely to be some ac field noise (in the μT range) in the laboratory arising from the mains circuits, or ripple on the current supplies used to drive the magnetic

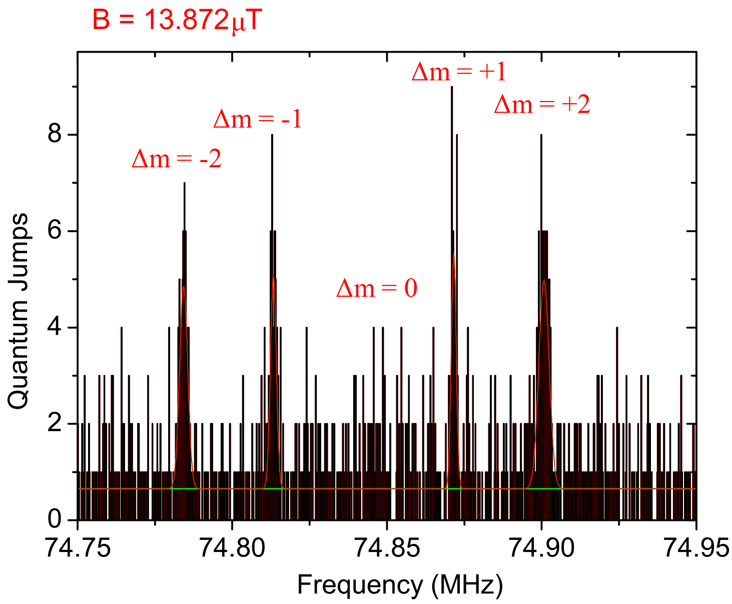

Figure 7.5: Measurement of the magnetic field flux density by observing the Zeeman splitting on the 436 nm transition. The comparatively broader linewidths are likely due to a combination of power broadening, ripple on the current supplies for the magnetic field coils and some ac magnetic field noise in the lab. The $\Delta m = 0$ component is too narrow to be seen on this scale.

field coils. This acts to broaden the $\Delta m = \pm 1, 2$ components of the clock transitions significantly (as seen in figure 7.5), which can be used to quantify the magnitude of the fluctuations and thereby correct for the resulting shift on the $\Delta m = 0$ components.

It is possible to shield the ion from drifts in the ambient field by enclosing the trap within a shield of high permeability metal. This also acts to stabilise the direction of the ion's quantisation axis. At the current time there is no shield present around the traps, but this is a necessary upgrade in the immediate future.

7.3 Quadrupole shift

As $^{171}\mathrm{Yb}^{+}$ has an alkali-like electron structure, the clock transitions involve states with $J \geq 1$. These states have an electric quadrupole moment [117], which will interact with any electric field gradients in the trap:

$$\Delta\nu_{\text{quad}} \propto Q_{DC}\Theta(L,J)\left(\frac{F(F+1)}{3} - m_F^2\right)(3\cos^2\beta - 1) \tag{7.5}$$

where Q_{DC} is the stray field gradient, $\Theta(L,J)$ is the quadrupole moment of the state and β is the angle between the stray field gradient and the ion's quantisation axis, defined by an applied magnetic field. In systems with a linear Zeeman shift it is common to lock to different pairs of Zeeman components to determine the shift as a function of m_F (m_J in ions with no hyperfine structure) and remove it in post-processing. The quadrupole moments for the $^2\text{D}_{3/2}$ and $^2\text{F}_{7/2}$ states have been measured to be $2.08(11)ea_0^2$ [118] and $-0.041(5)ea_0^2$ [17] respectively, where e is the elementary charge and a_0 is the Bohr radius.

In the case of $^{171}\text{Yb}^+$ it is desirable to remain only on the $m_F = 0 \to m_{F'} = 0$ component, and so the quadrupole shift must be identified in another manner. The geometry of the shift means that it averages to zero when the transition frequency is measured in three orthogonal magnetic field orientations [117]. This is the approach taken in this work.

Alternatively, the shift coefficient vanishes for the $m_F = \pm 2$ sublevels of the $^2\text{F}_{7/2}$ ($F = 3$) state. Therefore, by averaging over the $\Delta m_F = \pm 2$ components of the 467 nm transition, it is possible to measure a quadrupole shift-free transition frequency. The downside of this method is a re-introduction of a linear Zeeman shift, which then places more strict requirements on magnetic shielding and current supplies for the bias field.

7.4 Gravitational redshift

According to General Relativity, clocks experiencing a stronger gravitational potential will run slower. The Schwarzschild metric gives the gravitational redshift, z, to be:

$$z_{\text{grav}} = \frac{\Delta\nu_{\text{grav}}}{\nu_0} = \frac{1}{\sqrt{1 - \frac{2GM}{rc^2}}} - 1 \tag{7.6}$$

where G is Newton's gravitational constant, M is the mass of the body causing the redshift, r is the distance from the centre of mass of the body, and c is the speed of light. $2GM/c^2$ is referred to as the Schwarzschild radius, r_S. Under the limit $r \gg r_S$ (as is usually the case, unless the clock is orbiting particularly close to the event horizon of a black hole), the first term in the binomial expansion can be used:

$$z_{\mathrm{grav}} \approx \frac{GM}{rc^2} \tag{7.7}$$

When comparing clocks, it is the difference between the gravitational redshifts in the two systems that we observe. For two clocks at radii r_1 and r_2 from the centre of mass respectively:

$$\frac{\Delta\nu_{grav,1}}{\nu_0} - \frac{\Delta\nu_{grav,2}}{\nu_0} = \frac{GM}{r_1c^2} - \frac{GM}{r_2c^2} = \frac{GM}{c^2}\left(\frac{r_2 - r_1}{r_1r_2}\right) \tag{7.8}$$

Under the condition $r_1 \approx r_2 \approx r_E$ (as would be the case for clocks near the surface of the Earth, at radius r_E), and then making the substitutions $GM_E/r_E^2 = g$ (where M_E is the mass of the Earth and g is the acceleration due to gravity at its surface) and $r_2 - r_1 = \Delta r$, this simplifies even further to:

$$\frac{\Delta\nu_{grav,1}}{\nu} - \frac{\Delta\nu_{grav,2}}{\nu} \approx \frac{g\Delta r}{c^2} \tag{7.9}$$

This corresponds to a fractional shift of around 1×10^{-18} per centimetre of height difference.

However, this equation is only valid when comparing clocks that are reasonably local geographically, as local distortions to the gravitational potential become important when comparing clocks from further afield. Therefore, any long-distance comparison between frequency standards requires a precise determination of the relative gravitational potentials experienced by the two clocks.

7.5 Second-order Doppler shift

As the ion is in a reference frame moving at velocity $v_{\parallel} = v\cos\theta$ along the line-of-sight of an interrogating laser that is stationary in the laboratory reference frame (θ is measured in the laboratory frame [119]), it observes the laser to be Doppler shifted:

$$\frac{f_{\text{ion}}}{f_{\text{lab}}} = \gamma(1 - v_{\parallel}/c) \tag{7.10}$$

where f_{ion} is the laser frequency observed by the ion, f_{lab} is the laser frequency in the laboratory rest frame, and:

$$\gamma = \frac{1}{\sqrt{1 - v^2/c^2}} \tag{7.11}$$

is the contribution from special relativity (time dilation). Subsitituting the expression for γ into equation 7.10 gives:

$$\frac{f_{\text{ion}}}{f_{\text{lab}}} = \left(1 - v^2/c^2\right)^{-1/2}(1 - v_{\parallel}/c) \tag{7.12}$$

Performing a Taylor expansion of the first bracketed term leads to:

$$\frac{f_{\text{ion}}}{f_{\text{lab}}} = \left(1 + \frac{v^2}{2c^2} + O\left(\frac{v^4}{c^4}\right)\right)(1 - v_{\parallel}/c) \tag{7.13}$$

Since the ion is oscillating on timescales much shorter than our interrogation period of $\sim$ 100 ms, it is appropriate to use the time-averaged velocity $\left\langle v_{\parallel}\right\rangle$ and mean-square velocity $\langle v^2\rangle$, and therefore:

$$\frac{f_{\text{ion}}}{f_{\text{lab}}} \approx \left(1 - \frac{\left\langle v_{\parallel}\right\rangle}{c} + \frac{\langle v^2\rangle}{2c^2}\right) \tag{7.14}$$

to second order. As $\left\langle v_{\parallel} \right\rangle = 0$ for simple harmonic motion, the second term vanishes*. The expression for the fractional shift of the resonance observed in the laboratory frame is thus:

$$\frac{\Delta f}{f_{\text{lab}}} = \frac{f_{\text{lab}} - f_{\text{ion}}}{f_{\text{lab}}} = 1 - \frac{f_{\text{ion}}}{f_{\text{lab}}} \approx -\frac{\langle v^2 \rangle}{2c^2} \tag{7.15}$$

If we assume the laser is perfectly locked to the unperturbed atomic resonance in the frame of the ion, we can make the subsitution $f_{\text{ion}} = f_0$:

$$f_0 \approx \left(1 + \frac{\langle v^2 \rangle}{2c^2} \right) f_{\text{lab}} \tag{7.16}$$

7.5.1 Secular motion

The mean-square velocity associated with the slow secular motion in the axial direction is given by:

$$\left\langle v^2 \right\rangle = z_a^2 \omega_z^2 \tag{7.17}$$

For an ion cooled to the Doppler limit, $z_a \approx 20$ nm and $\omega_z \approx 2\pi \times 2$ MHz, leading to a fractional shift on the order of 10^{-19}.

7.5.2 Excess micromotion

The major contribution to the ion's velocity comes from the excess micromotion. This only occurs if the ion is displaced from the centre of the trap, and can be expressed in terms of the mean displacement:

$$\left\langle v^2 \right\rangle = \frac{q_z \Omega}{2\sqrt{2}} \left[\left(\frac{\bar{x}}{2} \right)^2 + \left(\frac{\bar{y}}{2} \right)^2 + \bar{z}^2 \right] \tag{7.18}$$

*A fractional first-order Doppler shift at the level of 10^{-17} has been observed in the Al^+ clocks at NIST by measuring a frequency difference when the clock transition is driven alternately by counter-propagating probe laser beams. This shift is not due to the ion's harmonic motion, but rather due to a slow drift (~ 2 nm/s) which may be attributable to slow charging of dielectric surfaces near the ion [120].

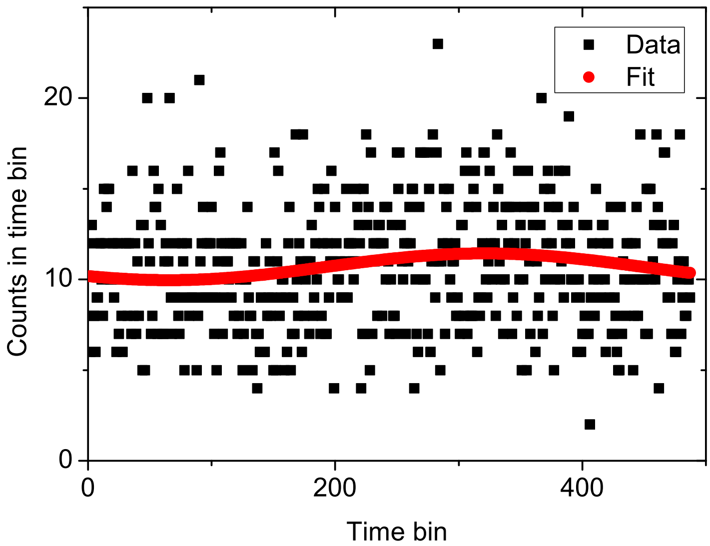

Figure 7.6: Typical rf-photon correlation signal for the horizontal cooling beam, with the depth of modulation corresponding to a peak Doppler shift of 1 MHz. This gives a peak residual velocity of 38 cm/s in this direction.

These mean displacements can be measured by analysis of the rf-photon correlation signal, as discussed in section 2.7. A typical rf-photon correlation signal is shown in figure 7.6. Analysis of this signal for three non-planar beams can determine the magnitude and direction of the ion's residual velocity.

7.6 Blackbody Stark shift

7.6.1 Static blackbody shift

Perhaps the most challenging systematic uncertainty for many optical frequency standards is the Stark shift imposed by thermal radiation. At room temperature, over 99% of the photons in the field are at wavelengths longer than 5 μm and hence theoretical models often approximate the black body field as a dc field and hence the static polarisability of the states is used. The shift $\Delta\nu_{\mathrm{BB}}$ is then given by:

$$\Delta\nu_{\mathrm{BB}} = -\frac{1}{2h}\Delta\alpha_{\mathrm{sc}}^{\mathrm{dc}}\left\langle E^2\right\rangle - \frac{1}{4h}\alpha_{\mathrm{ten}}^{\mathrm{dc}}\frac{3m_F^2 - F(F+1)}{F(2F-1)}\left(3\cos^2\theta - 1\right)\left\langle E^2\right\rangle \qquad (7.19)$$

where $\Delta\alpha_{\mathrm{sc}}^{\mathrm{dc}}$ is the difference between the scalar static polarisabilities of the two states involved in the transition, $\alpha_{\mathrm{ten}}^{\mathrm{dc}}$ is the tensor static polarisability of the upper clock state, $\langle E^2 \rangle$ is the mean-square electric field experienced by the ion and θ is the angle between the electric field and the ion's quantisation axis. Direct measurements of the scalar and tensor polarisabilities have been made for the 436 nm [118] and 467 nm [17] transitions. These values are combined with the electric field associated with the blackbody radiation using Planck's law for electromagnetic energy density:

$$\frac{\epsilon_0}{2}\left\langle E^2 \right\rangle + \frac{1}{2\mu_0}\left\langle B^2 \right\rangle = \frac{\hbar}{\pi^2 c^3}\int_0^\infty \frac{\omega^3}{\exp(\hbar\omega/k_B T) - 1} d\omega = \frac{\pi^2 k_B^4 T^4}{15\hbar^3 c^3} \tag{7.20}$$

where B is the Boltzmann constant and T is the temperature of the environment. The energies in the electric and magnetic fields are equal, so the left hand side simplifies to $\epsilon_0 \langle E^2 \rangle$.

In general, it is simply the scalar polarisability that is used to estimate the black body shift. However, this rather bold assumption does not take into account any anisotropy of the black body field, which would couple to the tensor polarisability of the clock state. The thermal environment inside the trap vacuum chamber is unlikely to be isotropic, as there may be some resistive heating of the electrodes and some rf heating of the dielectric spacers used in its construction. However, this shift has a similar functional dependence to the quadrupole shift and hence averages to zero when the transition frequency is averaged over three orthogonal magnetic field directions.

As it stands, it is not possible to make a measurement of the temperatures inside the vacuum chamber as the windows are opaque at thermal wavelengths. In a future rebuild of the vacuum system it would be sensible to use a window that allows thermal imaging of the trap, and to confirm these measurements by building an identical 'dummy' trap in which thermal sensors can be placed in the areas most liable to heating.

The problem is further compounded by the fact that the emissivity of the surfaces exposed to the ion are very poorly known. Even if their temperatures are precisely measured, their emission will not follow a blackbody spectrum. A high quality polish on the electrode surfaces could reduce the emissivity significantly, thus reducing their contribution to the radiation seen by the ion.

Alternatively, the as the shift scales as T^4 then cooling the trap to cryogenic temperatures would reduce the shift to a negligible level. This technique is used in the Hg^+ standard at NIST [22], though the use of a cryostat is primarily required to reduce the background pressure of Hg atoms. A cryostat would add an undesirable extra layer of complexity to the experiment and so efforts are ongoing to characterise the trap at room temperature.

7.6.2 Dynamic blackbody shift

The approximation that the blackbody shift is scalar in nature is only good up to a point. There will be some small fraction of the thermal radiation that can couple to the infrared transitions from the two clock states, leading to a temperature-dependent dynamic (ac) shift. This correction makes no contribution to our systematic error budget at the present time, as the uncertainty on the scalar correction is dominant. As this uncertainty is reduced, it will be become important to either directly measure the dynamic shift as a function of wavelength, or model the shift (which requires good knowledge of the oscillator strengths of all the relevant transitions).

7.7 Stark shift due to trapping field

If the ion is not located at the exact centre of the trap, there will be a Stark shift arising from the residual rf field. As the trap frequency is much lower than the optical frequency of the transition, it can be approximated as a dc field and the appropriate dc polarisability used as for the blackbody shift. The ion's net displacement from trap centre in the three principal trap axes can be derived from its residual velocity:

$$\bar{z} = \frac{v_{max}^{(z)}}{\sqrt{2}\omega_z} \tag{7.21}$$

The electric field as a function of displacement is known from the applied rf voltage and the geometry of the electrodes. For the residual micromotion shown in figure 7.6, the corresponding fractional frequency shift is -1.4×10^{-17} on the 467 nm transition.

7.8 AOM phase chirp

As the probe laser is shuttered on and off using an AOM, there will be some heating of the crystal shortly after application of the rf drive. This leads to a change in the optical phase leaving the crystal, which translates to a frequency offset when the light is applied to the ion. This effect has been measured using a self-heterodyne setup. One beam path was repeatedly shuttered on and off with the same duty cycle as the experiment, and the other was left permanently on. This was compared against the case where both beams were permanently on. The results are shown in figure 7.7, which shows no visible offset at the level of 4 mHz ($= 6 \times 10^{-18}$).

7.9 Fibre-induced phase noise

In order to measure the frequency of the probe laser, a small sample of the light must be delivered to the frequency comb via an optical fibre. Although the emission of the laser diode has a linewidth of around 1 Hz, by the time the light reaches the comb it can have been broadened significantly. Even more worrying is the prospect of a systematic frequency offset introduced by the fibre. This noise must therefore be actively cancelled.

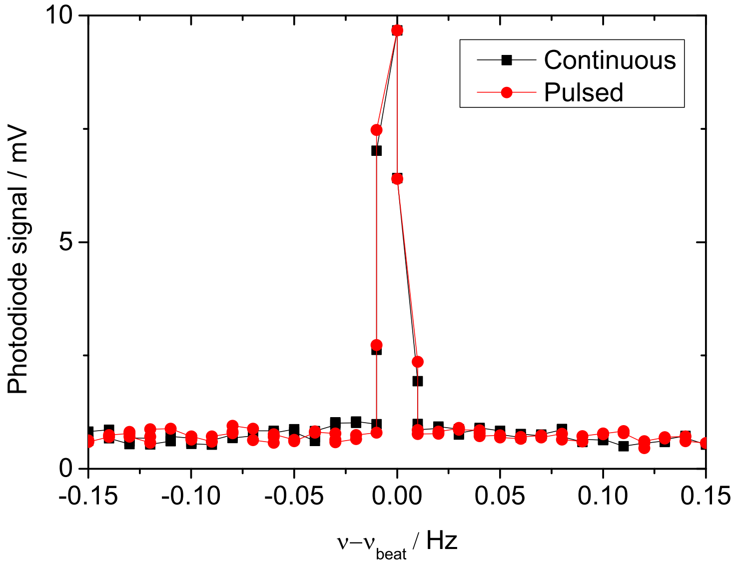

Figure 7.7: Frequency spectrum of the self-heterodyne beat with the AOM continuously on and with the AOM shuttered at the experimental duty cycle. There is no visible offset down to the level of 4 mHz. The shuttered signal is a factor of two smaller due to the 50% experimental duty cycle, and has been scaled up for easier comparison.

7.9.1 Noise sources

Acoustic vibrations couple to the light in the fibre by modifying the optical path length via the refractive index of the fibre. These act to broaden the linewidth of the laser at the comb end of the fibre, greatly degrading its stability. As the fibre is routed via a duct within the building, it is vulnerable to vibrations from the air conditioning system.

Fluctuations in the fibre temperature can lead to a systematic offset between the sent and received light. Using an estimate of 1.8×10^{-5} / K [121] for the thermal expansivity of the jacketed fibre, it is possible to place a limit on the shift that can be imposed by a thermal ramp acting to stretch the fibre. The ion clock and comb laboratories are actively stabilised to better than 0.1 K. Therefore, over a typical 6 hour dataset, the maximum temperature drift is 0.2 K. Assuming the temperature ramps linearly, for a 50 m fibre this corresponds to a length change of 8 nm/s. This would lead to a fractional first-order Doppler shift of 4×10^{-17}. In reality, temperature drifts in the air duct through which the fibre is routed could produce far larger shifts.

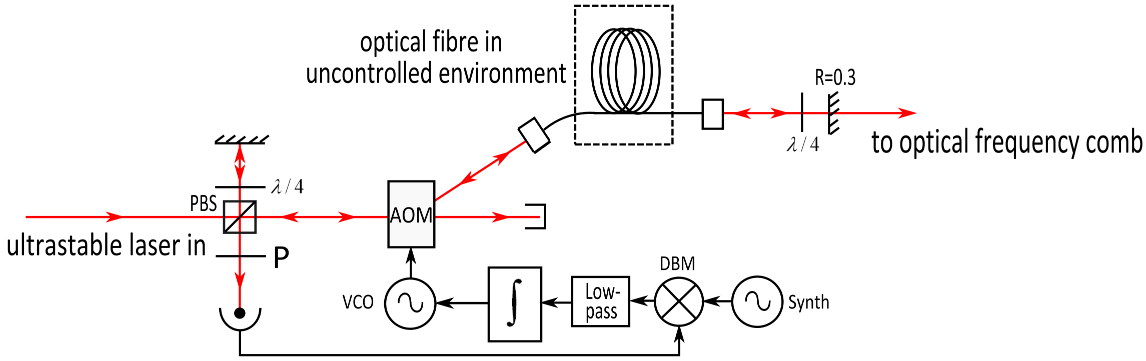

Figure 7.8: Apparatus used to cancel the phase noise introduced by an optical fibre. The first order diffraction from the AOM is coupled into the optical fibre. At the output end, a sample of the light is retroreflected back down the fibre, where it again is diffracted into the first order of the AOM before being overlapped with a reference arm to form a Michelson interferometer.

7.9.2 Cancelling the noise

Cancellation of the noise requires the derivation of an error signal, and the introduction of an actuator that can compensate for the noise. We have decided to take the approach of Ma *et al.* [122]. In this case the feedback is provided by a voltage-controlled oscillator (VCO) that drives an acousto-optic modulator at the input of the fibre, as shown in figure 7.8.

The first order diffraction of the AOM is coupled into the (polarisation-maintaining) optical fibre. Care is taken to align the polarisation of the light along one of the axes of the fibre, as the use of polarisating optics mean that polarisation noise is converted to amplitude noise on the light. Once it exits the fibre, a sample (30%) of the light is retroreflected using a normal-incidence beamsplitter. Double-passing the light through a quarter wave plate rotates the polarisation by 90°, and it is transmitted down the other axis of the fibre. It then passes through the AOM a second time, with the first order diffraction being collinear with the input beam. The polarising beamsplitter then picks off this counterpropagating beam. The light is overlapped with a 'reference' path to form a Michelson interferometer, and their 45° polarised components are selected using a polariser. The resulting self-heterodyne beat contains information about the noise introduced by the optical fibre. By mixing this beat with a reference synthesiser, an error signal is generated that can be used to lock the frequency of the beat to be the same as that of the synthesiser (with the VCO operating at half of this frequency).

The level of cancellation cannot be completely ascertatined from the noise on the in-loop beat, as this only carries information about how the light that has made a round-trip compares to the short arm of the interferometer (which may itself be noisy). It also relies on the noise being introduced equally on the forward and return trips through the fibre, as the output light comes from the half-way point in the loop. To monitor the performance of the servo, it is necessary to make an out-of-loop measurement of the frequency noise on the light. By attaching a second fibre to the end of the first, the resulting 100 m long fibre has its output in the same lab as the input. The output light can be overlapped with some of the input light, with the frequency spectrum of the beat providing a measure of the noise. The results are shown in figure 7.9. With the VCO locked, the resulting frequency offset on the transmitted laser was $(0.3 \pm 1.0) \times 10^{-19}$ at an averaging time of 10^4 s, and no cycle slips were observed throughout the entire dataset.

7.9.3 Noise on the laboratory fibres and free-space

At a lower level, it is of interest to know whether the optical fibres and free-space that link the probe laser to the ion trap introduce any noise onto the probe laser. By building an identical setup to that used to cancel the noise on the fibres leading to the optical frequency comb and looking at the self-heterodyne beat, it could be seen that the fibres (total length 4 m) had a negligible effect on the linewidth of the 467 nm light. However, it was observed that 2 m of free space introduced intermittent noise that occasionally broadened the beat by as much as 20 Hz. The noise was observed to decrease drastically when a simple cardboard box was placed over the exposed beam path, indicating that the noise was primarily due to air currents across the optical table.

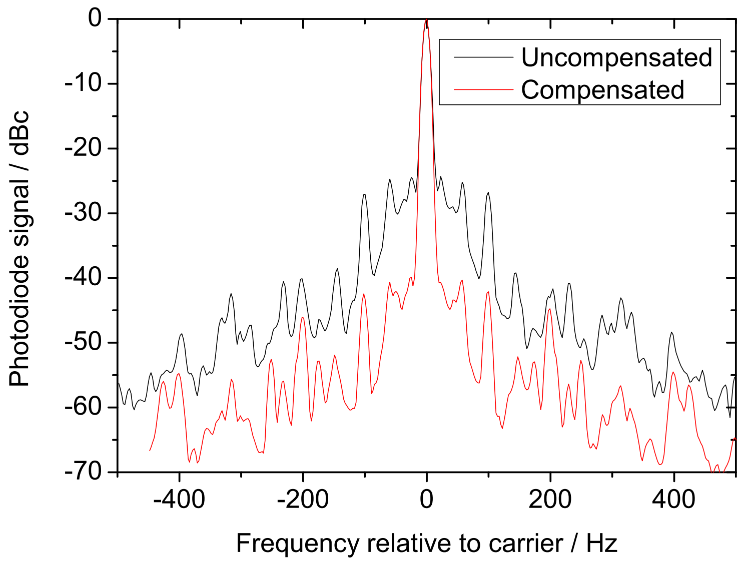

(a)

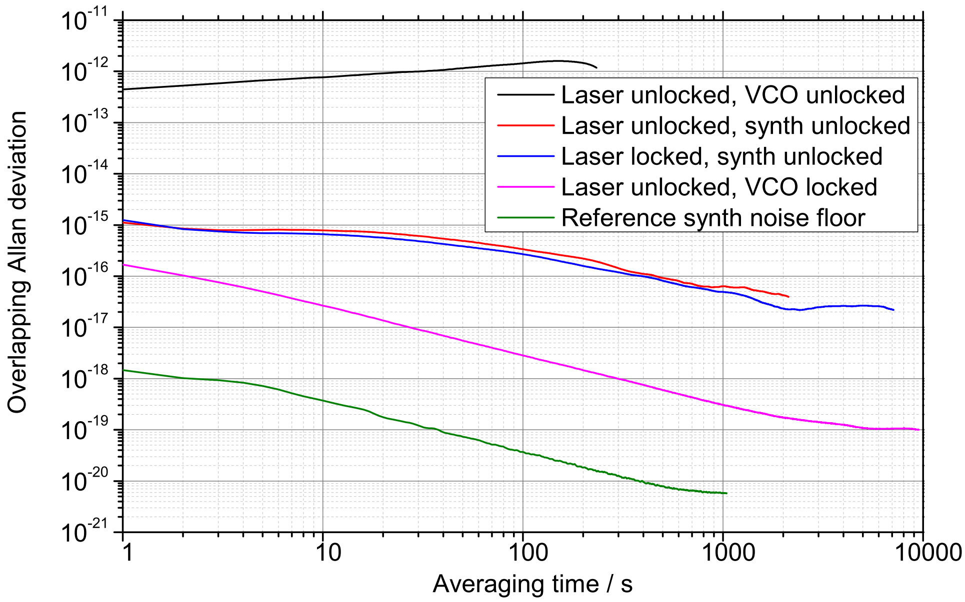

(b)

Figure 7.9: Out-of-loop performance of the noise cancellation servo for the 934 nm laser over 100 m of optical fibre. (a) Beat spectrum between the transmitted light and a reference path, demonstrating the suppression of the noise once the servo is locked. The central feature is resolution-limited (10 Hz FWHM) on the spectrum analyser used. (b) Stability of the beat as a function of averaging time measured with a bandwidth of 50 kHz. The locked instability is still around two orders of magnitude above the instability of the synthesiser used as the reference for the lock, most likely due to a combination of instability in the reference arm and the shot noise on the detector used to measure the beat.

7.10 Summary

All of the major systematic shifts associated with the two clocks transitions have now been discussed, along with their measurement and corrections. It is therefore possible to derive unperturbed transition frequencies by making absolute frequency measurements against the caesium primary standard. In the following chapter, table 8.1 summarises the systematic shifts relevant to the 2009 measurement of the frequency of the E2 transition at 436 nm, and table 8.2 contains a more comprehensive summary of the shifts relevant to the 2011 measurement of the frequency of the E3 transition at 467 nm.

Chapter 8

Absolute frequency measurements

Using the two new probe lasers, absolute frequency measurements were made of the two clock transitions using NPL's two optical frequency combs. Details of the workings of a frequency comb are given in section 8.1. A preliminary measurement of the 436 nm transition frequency is discussed in section 8.2, and a more advanced measurement of the 467 nm transition frequency is discussed in section 8.3.

8.1 Optical frequency combs

When considering a frequency standard based on an atomic transition, it is the periods of the radiation associated with the transition that are 'counted'. This raises an immediate problem as optical frequencies are much too high to be counted directly by current technology. Optical frequency combs are invaluable tools that allow these optical frequencies to be referenced to a microwave source [123, 124]. If the SI second is re-defined in terms of an optical transition, optical frequency combs will also allow the dividing down of the signal to the frequencies required by end-users. Such has been the impact of optical frequency combs, one half of the 2005 Nobel Prize in Physics was jointly awarded to Theodor Hänsch and John Hall 'for their contributions to the development of laser-based precision spectroscopy, including the optical frequency comb technique'. At NPL, there are two operational optical frequency combs: one based on a mode-locked Ti:sapphire laser [125, 126] and one based on an erbium-doped fibre laser [127].

8.1.1 Mode-locked lasers

An essential component of the modern frequency comb is a mode-locked laser, which provides the link between a microwave reference signal and optical frequencies. Mode-locked lasers with pulse lengths of a few picoseconds have been available since the 1970s in the form of dye lasers, but their gain bandwidth was not wide enough to make them a viable option as an optical frequency comb. The development of high-bandwidth Ti:sapphire lasers in the early 1990s [128] was one of the pivotal points in the development of this field, allowing the measurement of many optical frequencies using a single set of apparatus. Previously, the only method to count optical frequencies involved complex frequency chains [129] which were capable of targeting only one optical frequency at a time.

In a mode-locked laser, at a certain time all the modes will interfere constructively; this leads to brief, highly energetic output pulses that can be as short as a few femtoseconds [130], separated in time by the round-trip time of the cavity $\tau = 2L/c$. The modes within the pulse are separated in frequency by $f_{\mathrm{r}} = 1/\tau$, which is commonly referred to as the repetition rate.

It is possible to mode-lock a laser by passively and selectively attenuating cw light. One of these methods relies on the use of a nonlinear optical element within the cavity that has an absorption coefficient dependent on the intensity of incident light. This effect arises from the material having only a small number of absorbers (a 'saturable absorber'), which is the case in semiconductors. As the intensity of the light at the position of the saturable absorber varies, it attenuates the beam more strongly during periods of low intensity than high intensity, leading to preferential transmission of in-phase modes. This process repeats until out-of-phase modes are completely suppressed and the laser is mode-locked.

Another passive method utilises the optical Kerr effect. Some materials have a refractive index dependent on the intensity of the incident light. For a Gaussian beam, this means light near the centre of the beam sees a higher refractive index than light at the edge of the beam, leading to a lens-like effect known as self-focusing. Apertures can be placed inside the laser cavity near the focus of the pulsed mode, leading to the attenuation of the larger cw mode.

It is the process of mode-locking that allows a microwave reference signal to be brought up to optical frequencies. The separation of the comb modes in frequency space is compared to an rf source from the microwave standard, and then stabilised using a piezoelectric transducer attached to an intra-cavity mirror to feed back to the length of the cavity and hence stabilise the round-trip time.

Group velocity dispersion (GVD) in the laser cavity (especially in the lasing medium) acts to reduce the achievable comb bandwidth. In order to counteract this effect, an element with anomalous dispersion (e.g. a prism pair plus retroreflector, or a set of chirped mirrors) is introduced into the cavity. Careful adjustment of the compensation can lead to an order-of-magnitude reduction in the pulse length [130]. In a comb based on a mode-locked fibre laser, the dispersion is compensated for by stretching or squeezing the fibre.

8.1.2 Creating a supercontinuum

Ti:sapphire lasers are a common choice for generating frequency combs as they have a high peak power and a large gain bandwidth. They have the added advantage that the Ti:sapphire crystal itself acts as a Kerr lens, thus reducing the number of intracavity elements. It is often desirable to broaden the spectrum even further in order to allow the comb to deal with a wide range of unknown frequencies, as well as allowing the determination of the offset frequency via the self-referencing technique, as described in the following subsection. Supercontinuum generation can be achieved by passing the comb light through micro-structured crystal fibre [131]. The small diameter of the core results in strong four-wave mixing between the comb modes, producing a wide spectrum of coherent modes at the output of the fibre that can span an octave or more. The dispersion of the fibre is also made to be close to zero at the centre frequency of the input comb, which ensures that the pulse length (and hence comb bandwidth) is maintained through the fibre.

8.1.3 Stabilising the offset frequency

As mentioned earlier, one side effect of commonly used laser crystals is that longer wavelengths have a higher phase velocity through the crystal. This difference results in an increasing phase slip between the carrier peak and the pulse envelope from pulse to pulse (as shown in figure 8.1). This corresponds to a frequency offset f_0 of the comb modes from integer multiples of the repetition rate:

$$f_0 = \frac{f_{\mathrm{r}}\Delta\phi}{2\pi} \tag{8.1}$$

where $\Delta\phi$ is the phase offset between pulse n and pulse $n+1$ and f_{r} is the previously defined repetition rate of the mode-locked laser. The frequencies of the comb modes are given by:

$$f_n = nf_{\mathrm{r}} + f_0 \tag{8.2}$$

If the comb spectrum has been broadened to span an octave, it is possible to determine the offset frequency using a technique known as self-referencing [124]. The low frequency end of the comb is frequency doubled and overlapped spatially and temporally with the high frequency end on a photodiode. In this case, the lowest beat frequency will be f_0, which can easily be measured on a fast photodiode as shown in figure 8.1. The strength of the beat depends on the number of beating modes, which is limited by the bandwidth of the doubling crystal.

Under normal conditions, the cavity dispersion is free to drift due to temperature variations causing changes in the refractive index of the crystal and laser power (which alters the phase shifts associated with the Kerr effect). The offset frequency is compared

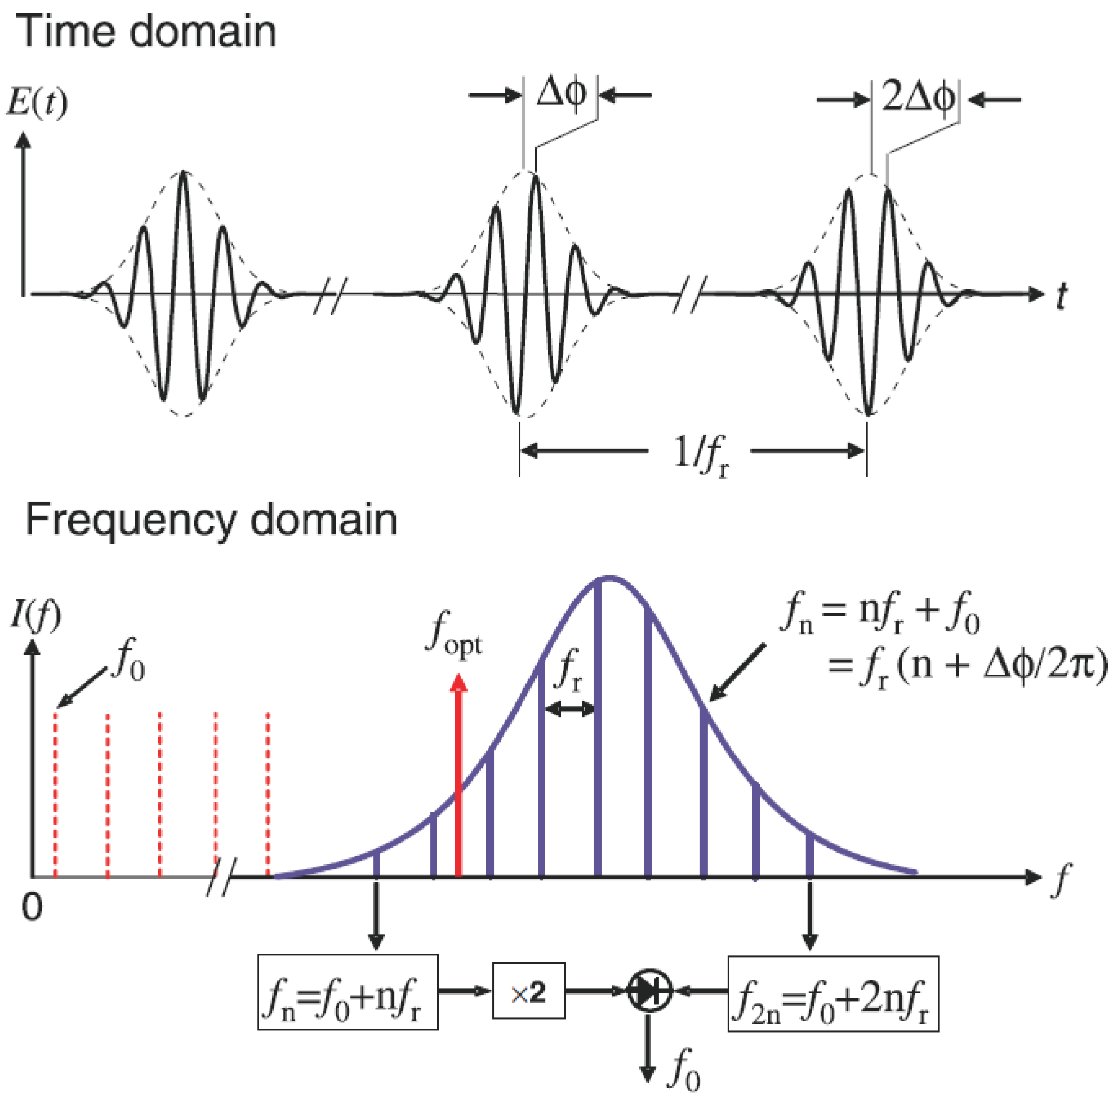

Figure reproduced with author's permission from [129]

Figure 8.1: Upper: Illustration of the pulse-to-pulse carrier-envelope phase slip resulting from dispersion within the laser cavity. Lower: Measurement of the offset frequency using the self-referencing technique (reproduced from [129]).

against a stable rf source in order to produce an error signal. If a prism pair is used to compensate for the cavity dispersion, any deviations can be corrected for by a retroreflecting mirror after the second prism, which allows changes to be made to the relative path lengths of the different modes. If chirped mirrors are used for dispersion compensation, feedback is provided by an AOM which controls the power of the pump laser [130].

8.1.4 Determination of unknown wavelengths

When developing optical frequency standards it is necessary to make high-precision measurements of the clock laser frequency. This can be either to determine the behaviour of the system, for example to determine the long- and short-term drift of reference oscillators, or to make an absolute measurement of the transition frequency. As mentioned previously, the repetition rate of the comb is locked to a stable microwave frequency, usually provided by the 10 MHz signal produced by a hydrogen maser which is itself referenced to the SI second as realised by a caesium fountain. The light of unknown frequency is overlapped with the comb light on a fast photodiode and a beat frequency extracted (the choice of comb mode is largely down to preference, it need not be the nearest to the unknown wavelength). The beat note f_{beat} and the frequency of the unknown laser f_{L} are related by:

$$f_{\mathrm{L}} = nf_{\mathrm{r}} \pm f_0 \pm f_{\mathrm{beat}} \tag{8.3}$$

where n represents the selected comb mode number. If the laser frequency is not initially known to better than one mode spacing, then this can be found by making changes to the repetition rate. The sign of the beat frequency and offset frequency can then be found by making small changes to the offset frequency and the repetition rate. In this way, measurements can be made of the absolute frequency of transitions that are only limited by the stability of the standard to which the repetition rate is referenced.

8.2 436 nm transition - January 2009

After construction of the 871 nm / 436 nm laser system, an absolute measurement of the $^2S_{1/2}$ $(F = 0) \rightarrow$ $^2D_{3/2}$ $(F' = 2)$ transition frequency was made [64]. As a lock to the clock transition had not been achieved at that time, the line centre was determined by repeatedly scanning the probe laser over the transition using a syntheisiser driving an AOM. The probe beam had a power of 1 nW focused to a $1/e^2$ radius of 50 μm. Each quantum jump was time-stamped and the synthesiser frequency at which it occurred was recorded as shown in figure 8.2a. The frequency of the optical cavity mode was continuously measured by the fibre comb. As the comb was referenced to a hydrogen maser, there was a large amount of statistical noise on the measured optical frequency (around 5×10^{-13} for a 1 second gate time). Once the data was averaged, it was clear that there was some residual nonlinear drift in the laser frequency, and so it was necessary to make a polynomial fit to the data in order to extract the cavity mode frequency more precisely at a given time and prevent any bias that would be introduced by the assumption of a linear drift. Some typical data are shown in figure 8.2b. Combining the synthesiser frequency with the absolute cavity frequency as measured by the comb allowed the absolute frequency at which each quantum jump occurred to be determined. Typically ten scans over the transition (a total of thirty minutes of data) were combined to improve the signal-to-noise ratio for the lineshape fit. Correcting each quantum jump to its absolute frequency also prevents any distortion of the lineshape that would occur if the raw scans were simply corrected using the apparent movement of line centre.

Line centre for this compound dataset was determined by fitting a Lorentzian lineshape to the observed profile. The linewidth appeared to be around 61 Hz for a 30 ms probe pulse (see figure 8.3), which is roughly twice the Fourier-limited width for a probe of this duration. This was attributed to a broadening of the probe laser linewidth to this level on the 30 minute timescale of the experiment, most likely due to random walks in the laser frequency.

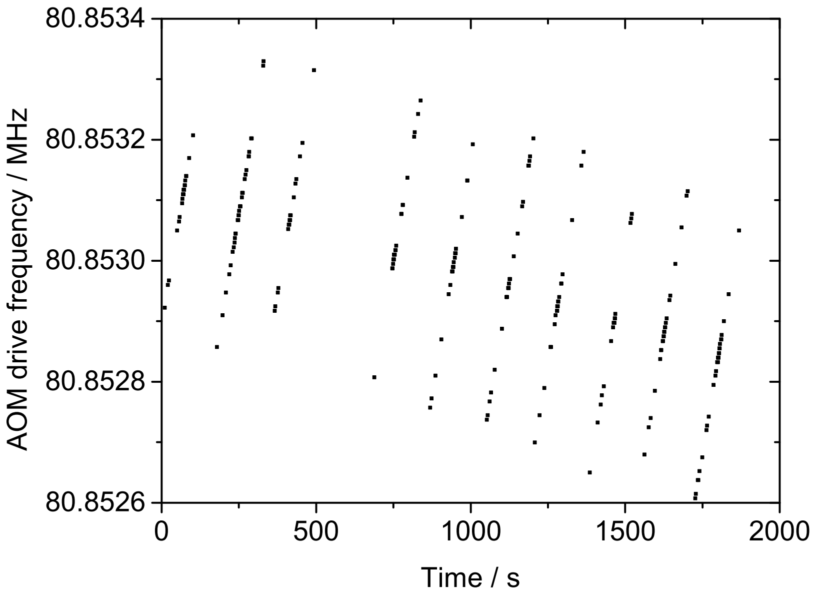

(a)

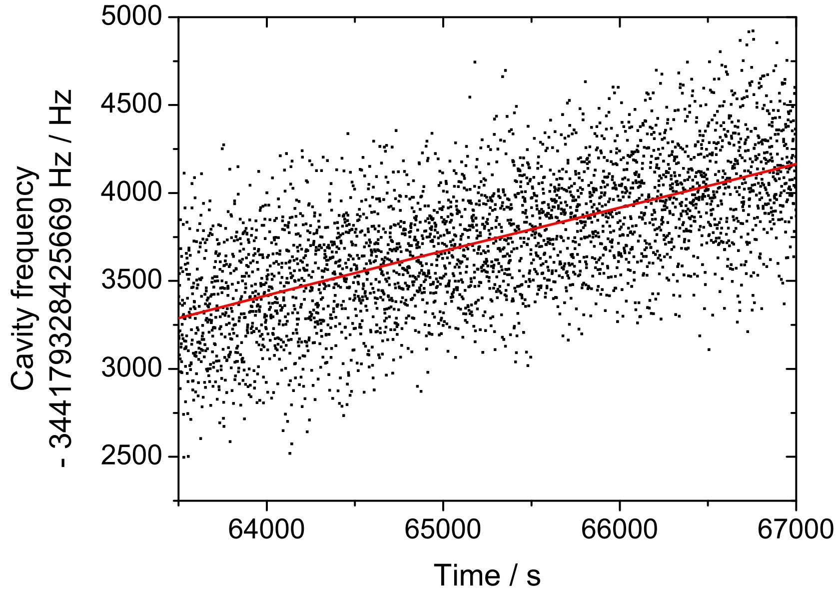

(b)

Figure 8.2: (a) AOM drive frequency at which quantum jumps were observed as a function of time for ten scans across the transition. Line centre is observed to move between scans due to isothermal drift of the ULE cavity. (b) Data from the frequency comb, showing the drift of the ULE cavity mode frequency. A third-order polynomial fit was made to the data before combining it with the synthesiser frequencies.

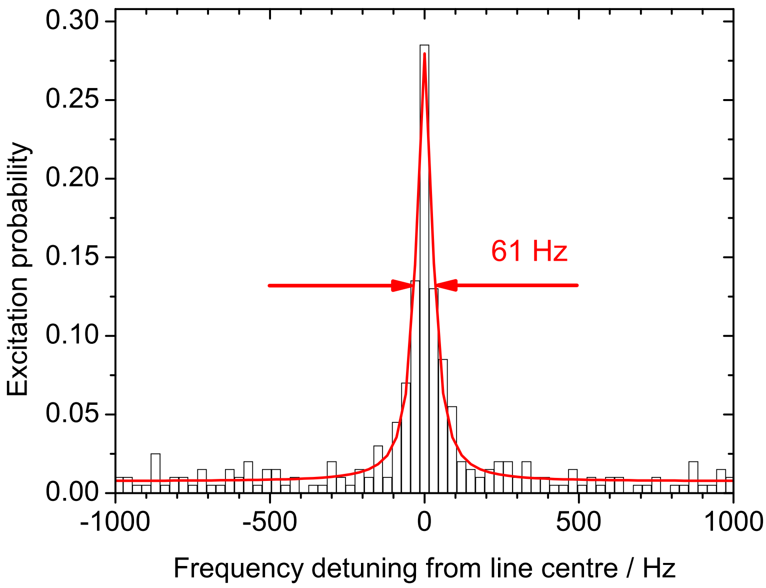

Figure 8.3: Typical compound scan over the 436 nm transition, composed of ten individual sweeps. The linewidth of 61 Hz is likely to be attributable to the linewidth of the probe laser on the thirty minute timescale of the repeated sweeps.

The magnetic field was rotated between three orthogonal orientations in order to average out any quadrupole shift. The shift was unresolved at the level of 10^{-14}, and so the uncertainty was taken to be the maximum possible value assuming a stray offset of 1 V between the endcaps (this value was chosen as it is approximately the voltage required on the outer endcaps to compensate the excess axial micromotion). The expected effect of this stray field gradient was calculated using a measured value for the quadrupole moment of the state [118].

In order to link the frequency measurement to the SI second, it was necessary to compare the frequency of the maser against Coordinated Universal Time (UTC). The International Bureau of Weights and Measures (BIPM) disseminate UTC and International Atomic Time (TAI) monthly in an update known as Circular T. The uncertainty in this correction was the dominant systematic uncertainty for the measurement.

Data were taken under a variety of magnetic field strengths, with their magnitude determined by observation of the Zeeman splittings of the $\Delta m_F = \pm 1, 2$ components of

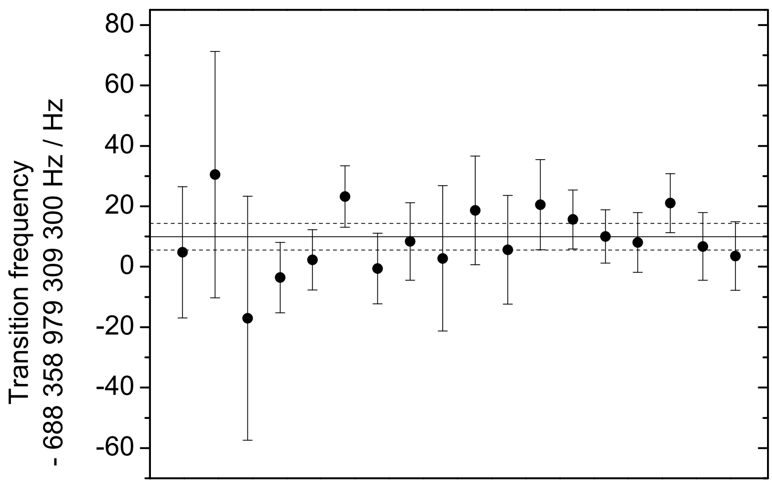

Figure 8.4: Summary of the data taken in January 2009. Each point represents a dataset, of varying duration and systematic conditions. The weighted mean and associated 1σ error bars are also shown.

the transition. The overall corrections and uncertainties are shown in table 8.1 and the data are summarised in figure 8.4.

After correction for these systematic shifts, the transition frequency was determined to be 688 358 979 309 309.9 (8.9) Hz, corresponding to a relative standard uncertainty of 1.3×10^{-14}. This is in good agreement with the previously reported value of 688 358 979 309 306.62 (73) Hz from PTB [28].

8.3 467 nm transition - June 2011

As one of the dominant uncertainties in the previous absolute measurements for both the 436 and 467 nm transitions was the statistical uncertainty from the relatively low amount of data obtained, for this measurement it was decided to lock to the atomic absorption. The replacement of the Ti:sapphire probe laser by the new diode-based system allowed the laser to operate for extended periods without mode-hops, a key attribute if it is to be stabilised to the clock transition. This allowed the acquisition of a total of 90 hours of

Source	Fractional shift / 10^{-15}	σ / 10^{-15}
Second-order Zeeman Fit to comb data Fit to peak	-	6.48
Black-body Stark	-0.52	0.10
Quadrupole shift	0	5.23
Maser correction to UTC	-12.90	9.99
TOTAL	-	**13.0**

Table 8.1: Summary of the systematic shifts of the 2009 measurement of the absolute frequency of the 436 nm transition, with their associated uncertainties. The shift and uncertainties for the second-order Zeeman shift, the uncertainty from the fit to the absorption line centre and the uncertainty on the polynomial fit to the comb data varied between datasets and hence have been quoted as a statistically-weighted average.

frequency data, with the longest dataset being 34 hours in duration (with some occasional dead time due to the laser losing lock to the transition, or periods during which the phase-locks in the comb laboratory were lost). The new probe laser has also enabled observation of a narrower clock transition linewidth and hence a reduced ac Stark shift and associated uncertainty. Improved amplitude stability of the probe laser has allowed observation of coherent excitation of the transition, increasing signal-to-noise via increased excitation efficiencies.

A further improvement over our measurement of the 436 nm transition frequency (and earlier measurements of the 467 nm transition) was the introduction of polarisation rotation on the cooling light, removing the need to switch off the 300 μT field that was required for laser cooling (see chapter 3) and therefore reducing the dead time in the interrogation cycle.

To eliminate the uncertainty in the correction of the maser frequency to UTC, the maser was frequently referenced to NPL's primary frequency standard CsF2.

As mentioned earlier, the largest systematic shift of the 467 nm transition at this time is the ac-Stark shift arising from the off-resonant coupling of the probe laser to other optical transitions from the two states involved in the clock transition, with the magnitude of the shift being directly proportional to the optical intensity at the ion. This shift was compensated for during the absolute measurement of the transition frequency.

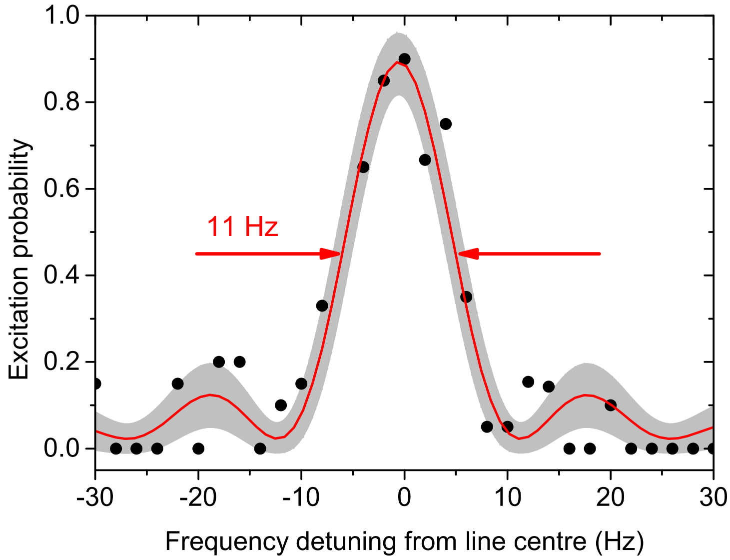

Figure 8.5: Typical scan over the octupole transition with a probe time of 100 ms. A lineshape has been fitted in accordance with the rectangular nature of the probe pulse, showing a linewidth close to the Fourier limit of 9 Hz. The shaded area indicates the predicted quantum projection noise.

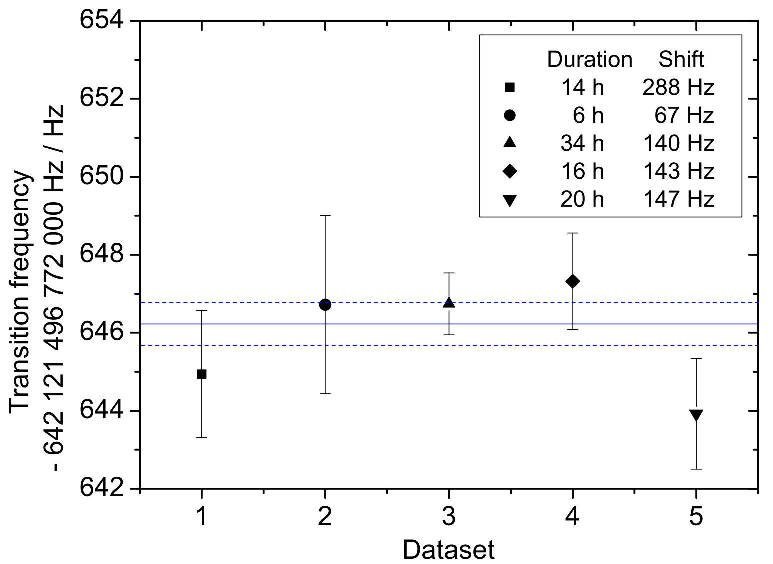

Figure 8.6: Summary of the data taken in July 2011. The legend shows the duration of each dataset and the differential scalar ac Stark shift between the high and low power probe phases. Error bars are one standard deviation and include statistical contributions only. The weighted mean and associated uncertainty is also shown as the solid and dashed lines respectively.

Source	**Fractional shift** / 10^{-15}	σ / 10^{-15}
ac Stark shift (467 nm)	0	0.49
Second-order Zeeman	-1.07	0.21
Black-body Stark	-0.23	0.11
Quadrupole shift	0	0.05
ac Stark shift (370 and 639 nm)	0	0.04
Second-order Doppler	0	0.02
Gravitational redshift	-0.08	0.02
Servo error	0	0.02
Trapping rf field Stark shift	-0.02	0.01
AOM phase chirp	0	0.01
SUBTOTAL	**-1.40**	**0.55**
Cs fountain systematics	0	0.23
10 MHz distribution and rf-synthesiser	0	0.10
50 m fibre frequency offset	0	0.04
Frequency comb	0	0.01
SUBTOTAL	**0**	**0.25**
Statistical	0	0.86
TOTAL	**-1.4**	**1.0**

Table 8.2: Summary of the systematic shifts of the 2011 measurement of the absolute frequency of the 467 nm transition, with their associated uncertainties. The individual uncertainties are quoted to two significant figures for clarity of presentation, but to avoid rounding errors the total has been evaluated from more precise values.

The laser was stepped between two power levels (approximately 4 and 8 mW) and the two perturbed transition frequencies were repeatedly measured by the optical frequency comb for periods of 20 min each. This period was chosen as it corresponded to the flicker floor of the maser to which the comb was referenced, meaning any additional data would not result in a decrease in the statistical uncertainty. Five different probe beam radii between 16 and 34 μm were used to allow the extrapolation to zero intensity to be carried out for different magnitudes of the ac Stark shift, while maintaining a constant intensity ratio by choosing to operate with the same two power levels throughout. The uncertainty on the extrapolation to zero intensity is dominated by how accurately the ratio between the two probe powers can be determined. The uncertainty on this figure is largely due to the instability of the probe power itself, but also includes contributions from the resolution and linearity of the detector used for the power measurement and the (at that time) 2 ms capture time of the power servo, which needed to be activated each time the probe pulse began.

The quadrupole shift was once again averaged over by rotating the magnetic field between three orthogonal orientations, each with a magnitude of 20 μT. To reduce the influence of short-term variations in the direction of the stray field gradient [28], the magnetic field was rotated after each extrapolation to zero intensity, corresponding to a 40 min interval. The three field directions were carefully chosen to give near-equal line strengths for the transition ($\approx 30\%$). The reduction in line strength from optimum conditions meant that higher probe intensities were required in order to give equivalent excitation efficiencies. The uncertainty in the cancellation of the shift due to inaccuracies in the orthogonality of the fields was estimated by analysis of the tensor component of the 467 nm ac Stark shift. This has a similar functional dependence to the quadrupole shift, but with the angle of interest being between the polarisation of the probe laser and the quantisation axis. By comparing the relative magnitudes of the easily-resolved tensor light shift in the three selected quantisation axes against the angular dependence predicted by theory, it is possible to place a limit on how well the desired field directions were realised and hence their orthogonality. This led to a 10% fractional uncertainty on the quadrupole shift in each of the three directions, and was taken to be 10% of the statistical uncertainty on each data point since the shift was unresolved.

The magnetic field flux density at the ion in each field orientation was determined by observation of the Zeeman splitting of the $\Delta m_F = \pm 1, 2$ components of the 436 nm transition. In this way, the residual field was determined to be 20 (2) μT throughout the entire dataset for all three orientations.

The second-order Doppler shift arising from excess micromotion was measured by analysis of the depth of modulation of the rf-photon correlation signal. This also allowed an estimate to be made of the Stark shift from any residual rf field experienced by the ion, which can arise from either a displacement of the ion from the centre of the trap or a phase offset between the endcaps. The corresponding shifts attributable to intrinsic micromotion and secular motion are two orders of magnitude lower and have therefore been neglected.

At the time of the measurement, the best estimate for the coefficient for the black-body Stark shift came from using a calculated value for the polarizability of the $^2F_{7/2}$ state [132] and an experimental measure of the dc-Stark shift of the 436 nm transition [118]. It was the uncertainty in this coefficient that dominated the uncertainty on the shift rather than uncertainties in the assumed 294 K temperature.

The frequency measurement has a gravitational redshift correction arising from the ytterbium ion being (75 ± 15) cm lower than the time-averaged height of the caesium atoms between the two interrogations (Ramsey time) of the caesium fountain.

Further systematic effects were considered that do not give rise to shifts at the uncertainty levels quoted in table 8.2. These included ac-Stark shifts from leakage of cooling and repumper laser beams through their shutters, phase chirp from the AOM used to shutter the 467 nm light and any offset in the steering of the clock laser frequency onto the line centre of the atomic transition.

This measurement was performed before the implementation of phase noise cancellation on the 50 m long optical fibre between the laser and the frequency comb, and so an estimate of any systematic offset had to be made as mentioned in section 7.9. If the temperature of the fibre were drifting by 0.2 K during the measurement period, this would lead to a first-order Doppler shift of 4×10^{-17} on the light received at the comb.

Apart from the ion-related systematic frequency shifts, it was also necessary to consider uncertainty arising from the frequency measurement process. For most of the data-taking period, the transition frequency was measured simultaneously using both of NPL's operational frequency combs, which were referenced to a common (maser-referenced) rf-synthesiser. The values obtained using the two combs agreed to within one part in 10^{17}, demonstrating that the combs themselves introduce negligible uncertainty. Any uncertainty introduced by the rf frequency distribution and synthesiser must also be accounted for. This was assessed by dividing the output of the rf synthesiser down to 10 MHz and comparing this signal with the original 10 MHz signal distributed from the hydrogen maser, and was found to result in a mean systematic frequency shift of less than one part in 10^{16} for the period of the measurements.

The systematic uncertainty arising from the caesium fountain primary standard itself was recently evaluated to be 2.3×10^{-16} [8], which also makes a contribution to the uncertainty of the frequency measurement.

After correcting for all of the systematic shifts, the transition frequency was determined to be 642 121 496 772 646.22 (67) Hz, corresponding to a relative standard uncertainty of 1.0×10^{-15} [18]. This is in good agreement with the previously published value [79] but has an uncertainty more than an order of magnitude lower. It is also in excellent agreement with another recently published value from PTB with a similar level of uncertainty [17]. To our knowledge, this represents the best international agreement between trapped ion optical frequency standards to date.

8.4 Summary

Absolute frequency measurements have been presented for the 436 and 467 nm optical clock transitions. Both obtained values agree well with other published results, demonstrating that $^{171}Yb^{+}$ has the potential to become a highly reproducible optical frequency standard.

Chapter 9

Conclusions

An external cavity diode laser has been constructed in order to drive the $^2S_{1/2} \rightarrow {}^2D_{3/2}$ electric quadrupole transition in a single ion of $^{171}Yb^+$. After locking the laser to an ultrastable optical cavity to improve its spectral purity, the highly forbidden clock transition has been successfully driven. The first measurement of the absolute frequency of this transition at NPL has been presented, with a total relative standard uncertainty of 1.3×10^{-14}. This agrees well with the previously reported value from PTB [28].

Since this measurement, the experimental duty cycle has been improved by the introduction of active state preparation, and fluorescence rates from the ion have been increased by adding a frequency sideband to the cooling laser rather than using microwaves to mix the hyperfine components of the ground state. Dead time in the interrogation cycle has been further reduced by the introduction of polarisation rotation on the cooling light, removing the need to switch off the 300 μT magnetic field that was required for laser cooling.

With these improvements, an improved measurement has been made of the absolute frequency of the 467 nm electric octupole transition, with more than an order of magnitude lower uncertainty than the previously published value [79]. The improved statistical uncertainty is largely attributable to the ability to stabilise the laser to the atomic absorption, along with greatly improved reliability due to the replacement of the

Ti:sapphire probe laser with a new diode-based system allowing stabilisation to the clock transition for extended periods. The new probe laser has also enabled observation of a narrower clock transition linewidth and hence a reduced ac Stark shift and associated uncertainty. Improved amplitude stability of the probe laser has allowed observation of coherent excitation of the transition, increasing signal-to-noise via increased excitation efficiencies. The total relative standard uncertainty on the frequency measurement was 1×10^{-15}, which places this transition amongst the most accurately measured optical frequencies. The recent measurement from PTB [17] at a similar level of uncertainty agrees with our obtained value to within their respective 1σ error bars, which represents the best international agreement between trapped ion optical frequency standards to date.

Other improvements to the experimental setup include the replacement of the the cooling laser (Ar^+-pumped Ti:sapphire and doubling cavity) with a diode laser operating at 370 nm, which can remain locked to its reference cavity for many days at a time. Noise introduced by the optical fibres between the experiment and the frequency comb is now actively cancelled, eliminating any possible systematic shift associated with the transfer. Finally, the noisy self-resonant circuit used to drive the ion traps has been replaced by a pure drive voltage from a resonant transformer driven by an amplified synthesiser.

9.1 Suggestions for future work

9.1.1 436 nm (E2) transition

The largest contribution to the uncertainty in the measurement of this transition frequency was the uncertainty in the correction of the frequency of the hydrogen maser to UTC. This could easily be improved in future by frequently calibrating the maser against the caesium primary standard during the measurement period, and stabilising the laser to the transition as was done for the 467 nm transition.

For this transition, the most troublesome systematic shift to stabilise is likely to be the quadrupole shift. Fluctuations in the ambient magnetic field in the laboratory will affect the orthogonality of the three operating fields and hence affecting the cancellation of the shift. In future, the trap should be shielded from fluctuations in the ambient field by enclosing it within a high-permeability magnetic shield (e.g. mu-metal). Protecting the ion from these variations should also act to reduce drift in the magnetic flux density and therefore the uncertainty in the second-order Zeeman shift.

9.1.2 467 nm (E3) transition

The largest contribution to the uncertainty in the measurement of this transition frequency was the level of statistical noise on the data (albeit more than an order-of-magnitude lower than for the measurement of the 436 nm transition frequency). However, the overall uncertainty would not be greatly improved by acquiring more data because the uncertainty in the ac Stark shift accounts for a large fraction of the systematic error budget. This shift and its uncertainty can be reduced by narrowing the observed linewidth even further (since the ac Stark shift scales inversely with the square of the resolvable linewidth [79]), and this may be possible by actively cancelling the noise on the optical fibres between the laser and the ion trap and replacing the high finesse cavity to which the laser is locked with one with a lower short-term instability. Alternatively, the introduction of an interrogation scheme that produces a near-perturbation-free transition frequency could reduce the uncertainty on the shift by two to three orders of magnitude [133]. Some details are given in appendix A.

The uncertainty in the second-order Zeeman shift correction also makes a large contribution to the error budget. The magnetic field was observed to drift from week-to-week by around 1–2 μT. In future, the ion should be shielded from this drift in the ambient field by the introduction of a mu-metal shield around the trap as discussed above. The uncertainty on the field would then be dominated by the uncertainty in the fits to the 436 nm Zeeman components, reducing it by one to two orders of magnitude. In addition, introduction of an EOM for polarisation rotation on the 935 nm transition would allow operation in even lower residual magnetic fields.

One of the limiting systematic uncertainties of the $^{171}Yb^{+}$ standard may be the black-body radiation shift of the clock transition, which is not easily measured. It can be reduced to negligible levels by cooling the trap to cryogenic temperatures [22]. In the case of the measurement presented in this thesis, it was the uncertainty in the shift coefficient for the 467 nm transition that dominated the uncertainty. This may be improved by utilising a recent measurement of the polarisability of the $^{2}F_{7/2}$ state [17], at which point the error would be dominated by knowledge of the thermal radiation profile. Imaging of the trap in the thermal region of the IR spectrum or the installation of sensors within the vacuum system would allow the identification of any hotspots, which could then be appropriately shielded or modified to minimise their temperature rise. Alternatively, since the 436 and 467 nm transitions experience the same thermal environment, a system has been proposed that combines the frequencies of the two transitions to produce a 'synthetic' frequency that can suppress this fractional shift to the 10^{-18} level [134].

With these further improvements, the fractional systematic uncertainty for this transition could be reduced below 1×10^{-17}.

9.1.3 Optical frequency ratios

As mentioned in section 8.1, the stability of the measured frequency is limited by the rf reference (namely a hydrogen maser) to which the repetition rate of the comb is stabilised. In order to improve the short-term stability of the beat (for example to compare two optical frequency standards with high precision) it is necessary to find a way to either further stabilise the repetition rate, or find a way to eliminate the noise introduced by the maser. For the optical frequency ratio measurements at NPL, we have decided to take the latter approach.

Two optical frequencies incident on the same frequency comb will see common-mode noise from the reference used to lock the repetition rate. The ratio of two laser frequencies incident on the same comb will have the form:

$$\frac{f_{\mathrm{L2}}}{f_{\mathrm{L1}}} = \frac{n_2 f_{\mathrm{r}} + f_0 + f_{\mathrm{beat}_2}}{n_1 f_{\mathrm{r}} + f_0 + f_{\mathrm{beat}_1}} \tag{9.1}$$

Making the substitution $\Delta_i = (f_0 + f_{\mathrm{beat_i}})/n_i f_{\mathrm{r}}$ leads to:

$$\frac{f_{\mathrm{L2}}}{f_{\mathrm{L1}}} = \frac{n_2}{n_1}\frac{1+\Delta_2}{1+\Delta_1} \tag{9.2}$$

and so, to first order, the repetition rate has been cancelled. This greatly suppresses the noise imposed by the microwave reference. For highly precise ratio measurements at the level of one part in 10^{17} or better, Δ_i must be known with an absolute uncertainty of 1×10^{-17}. Δ_i is typically on the order of 10^{-6}, so the fractional uncertainty in Δ_i need only be 10^{-11} which is well within the accuracy achievable with current technology.

The approach taken to measure the ratio is similar to that used at PTB [135] and is depicted in figure 9.1. The beat between laser 1 and the comb is added to the carrier-envelope offset frequency, and the resulting signal is divided by some integer N. The beat between laser 2 and the comb is also added to the carrier-envelope offset frequency, and the resulting signal is used to clock a direct digital synthesis (DDS) signal generator, which approximates a multiplication of $(n_1/\mathrm{N}n_2)$ using a 48-bit tuning word. Subtracting these two derived frequencies from one another yields a signal at the difference between the two laser frequencies, which is notably free of any terms that depend on the photodiode signals.

9.1.4 Time-variation of the fine structure constant?

With the two probe lasers simultaneously stabilised to their respective transitions, it should be possible to measure a single-ion optical frequency ratio. As discussed in section 1.5, measuring a ratio in this way should eliminate some common-mode systematic frequency shifts. Repeated measurements of this ratio should allow a stringent limit to

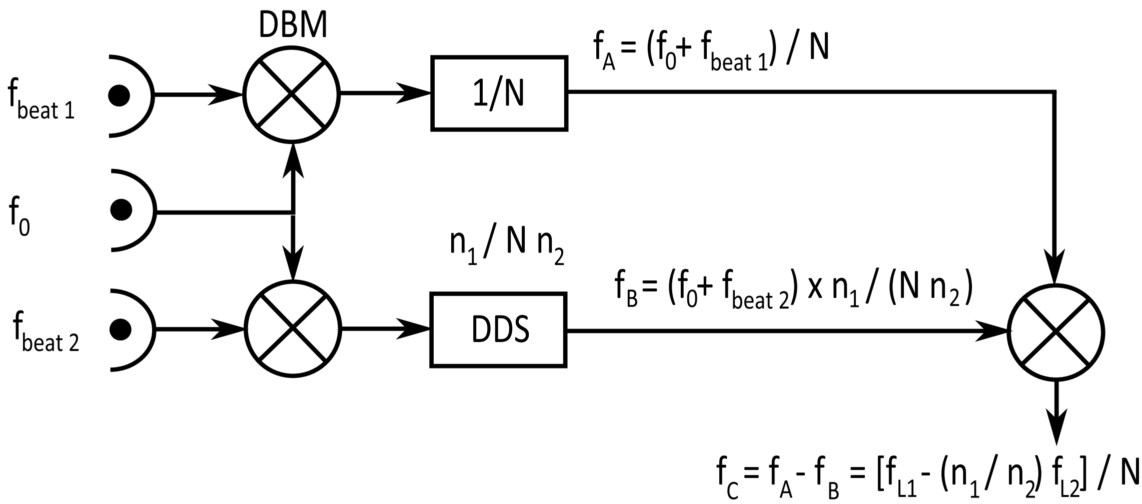

Figure 9.1: Extraction of the beat signal between lasers L1 and L2 (modified from [135]).

be placed on any time-variation of the fine-structure constant α.

Appendix A

Hyper-Ramsey Spectroscopy

As discussed in subsection 5.1.2, the excitation of the clock transition is currently performed using a single π-pulse (Rabi excitation). There is another technique in common practice amongst microwave frequency standards, referred to as Ramsey's method of separated oscillatory fields. The advantages of this technique are discussed in section A.1, and the difficulties in its implementation on the 467 nm transition are elaborated upon in section A.2.

A.1 Ramsey excitation

The excitation profile is modified if instead of a single π-pulse of duration t, two $\pi/2$ pulses of duration $\tau = t/2$ are separated by a dark time T ('Ramsey time'). In this case, the excitation profile takes the form:

$$|c_2|^2 = (\Omega_0\tau)^2 \operatorname{sinc}^2\left(\frac{\Omega\tau}{2}\right)\left[\cos\left(\frac{\Omega\tau}{2}\right)\cos\left(\frac{\delta T}{2}\right) - \frac{\delta}{\Omega}\sin\left(\frac{\Omega\tau}{2}\right)\sin\left(\frac{\delta T}{2}\right)\right] \tag{A.1}$$

where $\delta = \omega_0 - \omega$ is the laser detuning from line centre. Some sample fringes are shown in figure A.1b, with the corresponding Rabi lineshape shown in figure A.1a. In this case, the width of the central fringe is not determined by the duration of the pulse but rather by the time between the pulses. The resulting fringe is approximately twice as narrow as the lineshape generated by a single pulse with duration T. To make the most of this improvement, it is desirable to make $\tau \ll T$ in order to minimise dead time.

A.2 Hyper-Ramsey sequence

Unfortunately, the 467 nm (E3) transition does not lend itself to conventional Ramsey interrogation. As the two probe pulses are required to be as short as possible, an increase in intensity is required to in order to drive the transition. This leads to a much higher Stark shift from the probe laser. During the Ramsey time the frequency perturbation due to the laser is absent, and hence the ion and probe laser (which was resonant during the probe) will develop a phase offset, leading to a frequency bias on the Ramsey fringes as shown in figure A.2.

Recently, a new form of Ramsey interrogation has been proposed [133]. The phase offset during the Ramsey time (resulting in the frequency offset shown in figure A.2) is avoided by stepping the probe laser back to the Stark shift-free transition frequency, and an additional π-pulse of opposite phase to the two $\pi/2$-pulses acts to cancel any inaccuracies in the pulse areas, and also to cancel any phase errors associated with the different vibrational levels of the trap having different Rabi frequencies. Some simulated fringes are shown in figure A.3.

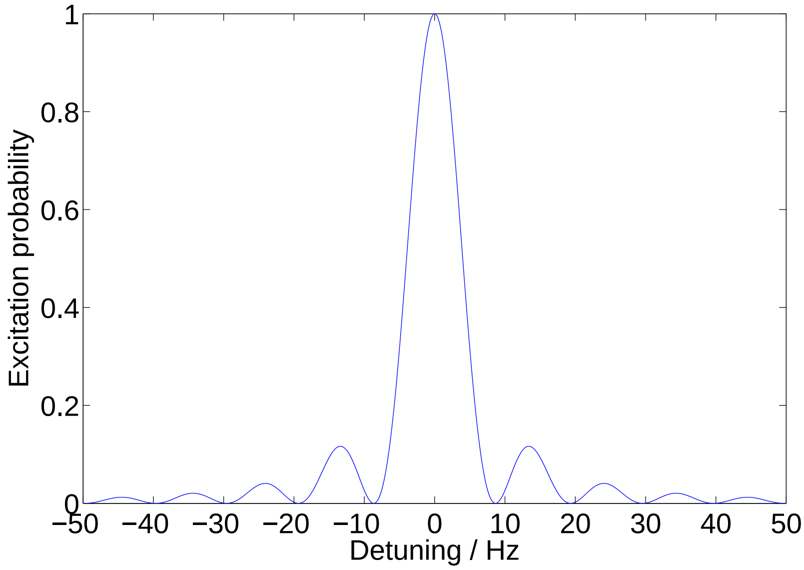

(a)

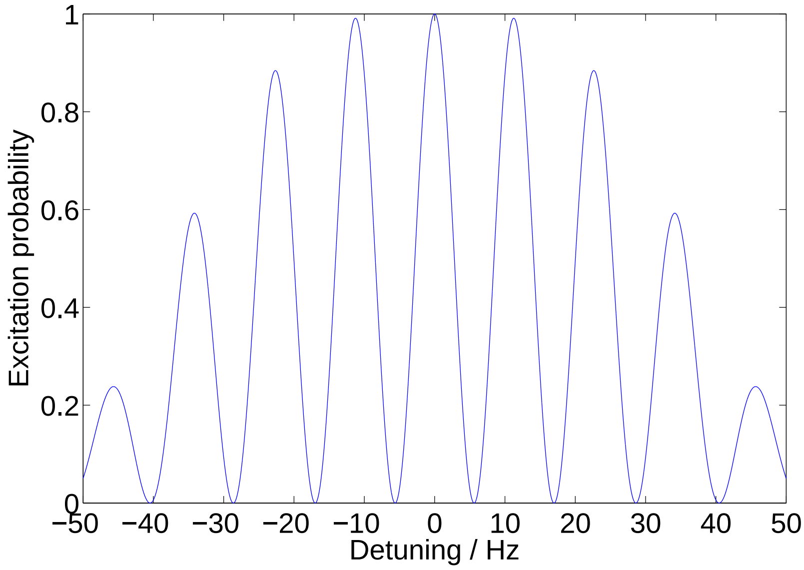

(b)

Figure A.1: Comparison of two interrogation schemes with the same total cycle time of 100 ms. a) Rabi profile generated by a single 100 ms π-pulse, and b) Ramsey profile generated by two 15 ms $\pi/2$-pulses separated by a Ramsey time of 70 ms.

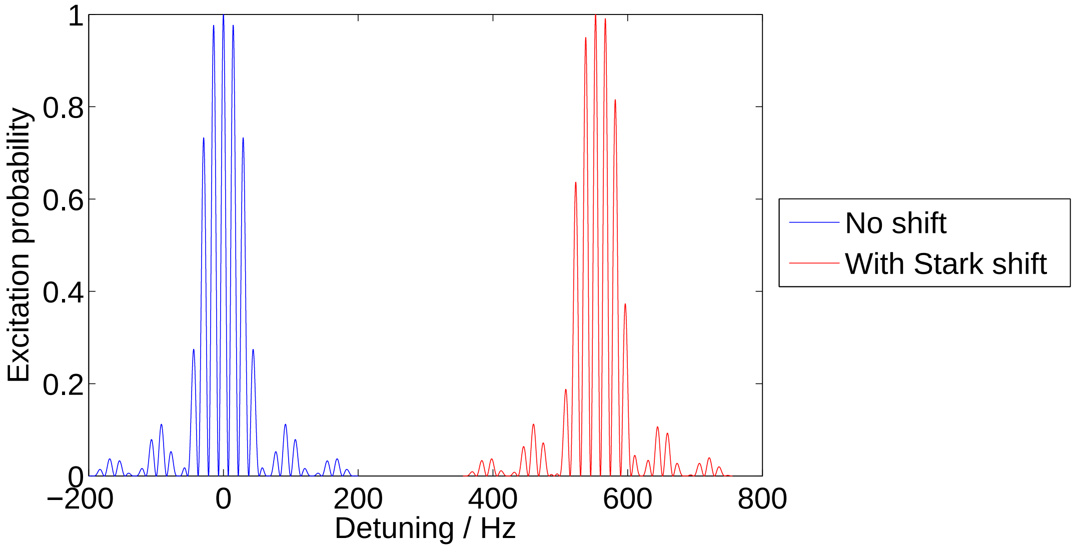

Figure A.2: Comparison of Ramsey fringes generated with pulse durations of 15 ms and a Ramsey time of 50 ms, with and without a probe laser-induced Stark shift of 555 Hz present. There is an obvious distortion to the perturbed fringes, and there is an additional offset of -3.2 Hz on the central fringe.

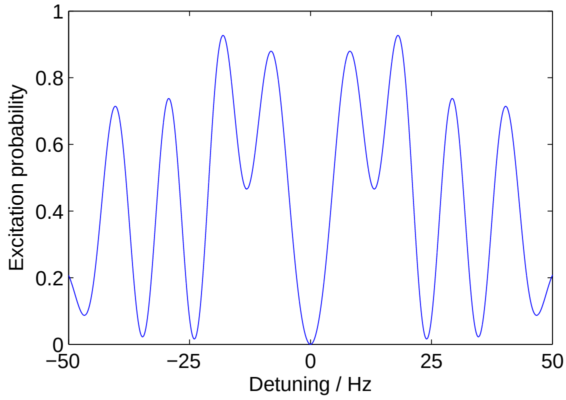

(a)

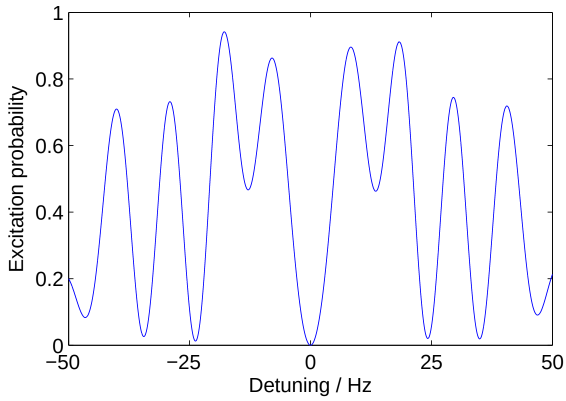

(b)

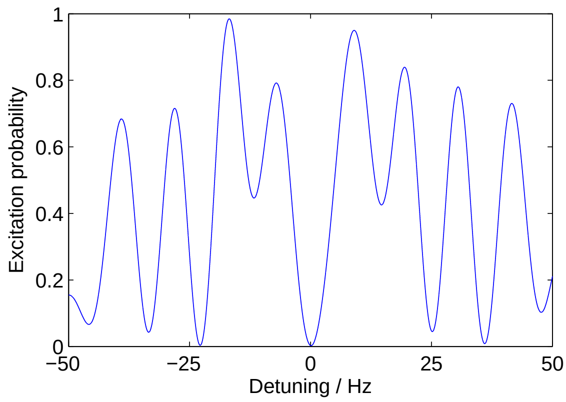

(c)

Figure A.3: Hyper-Ramsey fringes generated with pulse durations of 15 ms and a Ramsey time of 50 ms. In (a) the Stark shift of 555 Hz has been compensated perfectly. In (b) the correction has been misjudged by 0.1% with the central fringe having a resulting offset of only 1 mHz from the unperturbed frequency, an improvement of nearly three orders of magnitude over the uncertainty associated with the conventional extrapolation technique. In (c) the correction has been misjudged by 0.5% with the central fringe having an offset of 118 mHz, an improvement of more than a factor of twenty over the extrapolation technique.

Appendix B

Publications, presentations and posters

Peer reviewed publications that have resulted from this work:

- S. A. Webster, R. M. Godun, S. A. King *et al.*, "Frequency measurement of the $^2S_{1/2}$ – $^2D_{3/2}$ electric quadrupole transition in a single $^{171}Yb^+$ ion", *IEEE Trans. Ultra. Ferro. Freq. Contr.* **57** p 592, 2010.

- S. A. King, R. M. Godun, S. A. Webster, H. S. Margolis, L. A. M. Johnson, K. Szymaniec, P. E. G. Baird and P. Gill, "Absolute frequency measurement of the $^2S_{1/2}$ – $^2F_{7/2}$ electric octupole transition in a single $^{171}Yb^+$ ion with 10^{-15} fractional uncertainty", *New J. Phys.* **14** 013045, 2012.

Conference proceedings:

- P. Gill, S. A. Webster, G. Huang, K. Hosaka, A. Stannard, S. N. Lea, R. M. Godun, S. A. King, B. R. Walton and H. S. Margolis, "A trapped $^{171}Yb^+$ ion optical frequency standard based on the $^2S_{1/2}$ - $^2F_{7/2}$ transition", *Frequency Standards and Metrology (Proceedings of the 7th Symposium)*, pp 250-258, 2008.

Oral presentations:

Sep 2011	Quantum, Atomic, Molecular and Plasma Physics, University of Oxford
May 2011	European Frequency and Time Forum, San Francisco

Poster presentations:

Sep 2009	Quantum, Atomic, Molecular and Plasma Physics, University of Leeds
Aug 2010	Photon10, University of Southampton
Jul 2011	Atomic Clocks and Fundamental Constants, Physikzentrum Bad Honnef
Apr 2012	European Frequency and Time Forum, Gothenburg

Bibliography

[1] *Comptes Rendus de la 11e CGPM (1960)*, p. 87, 1961.

[2] *Resolutions adopted by the General Conference on Weights and Measures (24th meeting)*, p. 23, 2011.

[3] *Report of the 17th meeting to the International Committee for Weights and Measures*, p. 40, 2006.

[4] *Report of the 18th meeting to the International Committee for Weights and Measures*, p. 58, 2007.

[5] L. Essen and J. Parry *Nature*, vol. 176, pp. 280–2, 1955.

[6] *Comptes Rendus de la 13e CGPM (1967/8)*, p. 103, 1969.

[7] K. Szymaniec, S. E. Park, G. Marra, and W. Chalupczak *Metrologia*, vol. 47, pp. 363–76, 2010.

[8] R. Li, K. Gibble, and K. Szymaniec *Metrologia*, vol. 48, pp. 283–9, 2011.

[9] *Proceedings of the IEEE*, vol. 54, no. 2, pp. 221–30, 1966.

[10] W. M. Itano *et al. Phys. Rev. A*, vol. 47, pp. 3554–70, 1993.

[11] H. S. Margolis *Contemp. Phys.*, vol. 51, pp. 37–58, 2010.

[12] Updated from an image provided by H. S. Margolis

[13] C. Champenois *et al. Phys. Lett. A*, vol. 331, pp. 298–311, 2004.

[14] K. Matsubara *et al. Appl. Phys. Express*, vol. 1, p. 067011, 2008.

[15] M. Chwalla *et al. Phys. Rev. Lett.*, vol. 102, p. 023002, 2009.

[16] Y. Huang *et al. Phys. Rev. A*, vol. 85, p. 030503, 2012.

[17] N. Huntemann *et al. Phys. Rev. Lett.*, vol. 108, p. 090801, 2012.

[18] S. A. King *et al. New J. Phys.*, vol. 14, p. 013045, 2012.

[19] H. S. Margolis *et al. Science*, vol. 306, pp. 1355–8, 2004.

[20] G. Barwood, P. Gill, G. Huang, and H. Klein *IEEE Trans. Instr. Meas.*, vol. 47, pp. 226–9, 2007.

[21] P. Dubé *et al. Phys. Rev. Lett.*, vol. 95, p. 033001, 2005.

[22] W. H. Oskay *et al. Phys. Rev. Lett.*, vol. 97, p. 020801, 2006.

[23] J. von Zanthier *et al. Opt. Lett.*, vol. 25, pp. 1729–31, 2000.

[24] Y. H. Wang *et al. Opt. Commun.*, vol. 273, pp. 526–31, 2007.

[25] M. O. E. Peik, K. Zimmermann and C. Tamm *arXiv:0812.3548v2*, 2009.

[26] W. G. Rellergert *et al. arXiv:1011.0769v1*, 2010.

[27] C. W. Chou *et al. Phys. Rev. Lett.*, vol. 104, p. 070802, 2010.

[28] C. Tamm, S. Weyers, B. Lipphardt, and E. Peik *Phys. Rev. A*, vol. 80, p. 043403, 2009.

[29] T. Rosenband *et al. Science*, vol. 319, pp. 1808–12, 2008.

[30] H. Katori *Nature Photonics*, vol. 5, pp. 203–210, 2011.

[31] A. Derevianko and H. Katori *Rev. Mod. Phys.*, vol. 83, pp. 331–47, 2011.

[32] M. D. Swallows *et al. Science*, vol. 331, pp. 1043–6, 2011.

[33] G. Wilpers *et al. Metrologia*, vol. 44, p. 146, 2007.

[34] T. Kohno *et al. Applied Physics Express*, vol. 2, p. 072501, 2009.

[35] N. D. Lemke *et al. Phys. Rev. Lett.*, vol. 103, p. 063001, 2009.

[36] G. K. Campbell *et al. Metrologia*, vol. 45, p. 539, 2008.

[37] X. Baillard *et al. The European Physical Journal D - Atomic, Molecular, Optical and Plasma Physics*, vol. 48, pp. 11–7, 2008.

[38] F.-L. Hong *et al. Opt. Lett.*, pp. 692–4, 2009.

[39] S. Falke *et al. Metrologia*, vol. 48, p. 399, 2011.

[40] I. R. Hill in preparation

[41] S. Ruhmann and others. *Lasers and Electro-Optics Europe (CLEO EUROPE/EQEC), 2011 Conference on and 12th European Quantum Electronics Conference*, 2011.

[42] M. Petersen *et al. Phys. Rev. Lett.*, vol. 101, p. 183004, 2008.

[43] L. Yi *et al. Phys. Rev. Lett.*, vol. 106, p. 073005, 2011.

[44] T. Akatsuka, M. Takamoto, and H. Katori *Nature Physics*, vol. 4, pp. 954–9, 2008.

[45] N. Poli *et al. Phys. Rev. A*, vol. 77, p. 050501, 2008.

[46] J.-P. Uzan *Rev. Mod. Phys.*, vol. 75, pp. 403–55, 2003.

[47] S. N. Lea *Rep. Prog. Phys.*, vol. 70, pp. 1473–1523, 2007.

[48] V. A. Dzuba and V. V. Flambaum *Phys. Rev. A*, vol. 77, p. 012515, 2008.

[49] V. A. Dzuba, V. V. Flambaum, and M. V. Marchenko *Phys. Rev. A*, vol. 68, p. 022506, 2003.

[50] C. Schrama, E. Peik, W. Smith, and H. Walther *Opt. Commun.*, vol. 101, pp. 32–6, 1993.

[51] Image provided by H. A. Klein

[52] D. Leibfried, R. Blatt, C. Monroe, and D. Wineland *Rev. Mod. Phys.*, vol. 75, pp. 281–324, 2003.

[53] H. G. Dehmelt *Adv. At. Mol. Phys.*, vol. 3, pp. 53–72, 1967.

[54] M. Roberts, *Spectroscopy of a single Ytterbium Ion.* PhD thesis, Imperial College, University of London, UK, 1997.

[55] B. Bleaney and B. Bleaney, *Electricity and Magnetism.* Oxford Science Publications, 1989.

[56] C. Champenois, *Trapping and cooling of ions.* Les Houches predoc-school, 2008.

[57] D. J. Wineland, W. M. Itano, J. C. Bergquist, and R. G. Hulet *Phys. Rev. A*, vol. 36, pp. 2220–32, 1987.

[58] A. Safavi-Naini, P. Rabl, P. F. Weck, and H. R. Sadeghpour *Phys. Rev. A*, vol. 84, p. 023412, 2011.

[59] A. Bauch, D. Schnier, and C. Tamm *J. Mod. Opt.*, vol. 39, p. 389, 1992.

[60] K. Sugiyama and J. Yoda *Phys. Rev. A*, vol. 55, pp. R10–R13, 1997.

[61] T. Rosenband *et al. Phys. Rev. Lett.*, vol. 98, p. 220801, 2007.

[62] M. Roberts *et al. Phys. Rev. A*, vol. 62, p. 020501, 2000.

[63] N. Yu and L. Maleki *Phys. Rev. A*, vol. 61, p. 022507, 2000.

[64] S. Webster *et al. IEEE Trans. Ultrason. Ferroelec. Freq. Cont.*, vol. 57, pp. 592–9, 2010.

[65] P. Gill and F. Riehle *Proc. 20th European Frequency and Time Forum (Braunschweig)*, pp. 282–8, 2006.

[66] Source: NIST Atomic Spectra Database (ASD)

[67] S. Olmschenk *et al. Phys. Rev. A*, vol. 80, p. 022502, 2009.

[68] D. J. Berkeland and M. G. Boshier *Phys. Rev. A*, vol. 65, p. 033413, 2002.

[69] T. W. Hansch and B. Couillaud *Opt. Commun.*, vol. 35, pp. 441–444, 1980.

[70] D. Kielpinski *et al. Opt. Lett.*, vol. 31, pp. 757–9, 2006.

[71] A.-T. Nguyen *et al. Rev. Sci. Inst.*, vol. 81, p. 053110, 2010.

[72] K. Toyoda *et al. Opt. Lett.*, vol. 26, pp. 1897–1899, 2001.

[73] J. Migdalek *J. Phys. B: At. Mol. Opt. Phys.*, vol. 13, pp. L169–Ll74, 1980.

[74] S. Olmschenk *et al. Phys. Rev. A*, vol. 76, p. 052314, 2007.

[75] E. Biémont, J.-F. Dutrieux, I. Martin, and P. Quinet *J. Phys. B: At. Mol. Opt. Phys.*, vol. 31, pp. 3321–3333, 1998.

[76] R. W. Berends, E. H. Pinnington, B. Guo, and Q. Ji *J. Phys. B: At. Mol. Opt. Phys.*, vol. 26, pp. L701–L704, 1993.

[77] D. Engelke and C. Tamm *Europhys. Lett.*, vol. 33, pp. 347–352, 1996.

[78] M. M. Schauer *et al. Phys. Rev. A*, vol. 79, p. 062705, 2009.

[79] K. Hosaka *et al. Phys. Rev. A*, vol. 79, p. 033403, 2009.

[80] A. L. Schawlow and C. H. Townes *Phys. Rev.*, vol. 112, pp. 1940–9, 1958.

[81] M. W. Fleming and A. Mooradian *IEEE J. Quantum Electron.*, vol. 17, pp. 44–59, 1981.

[82] C. Henry *IEEE J. Quantum Electron.*, vol. 18, pp. 259–264, 1982.

[83] S. D. Saliba and R. E. Scholten *Appl. Optics*, vol. 48, pp. 6961–6, 2009.

[84] S. A. Webster, M. Oxborrow, and P. Gill *Phys. Rev. A*, vol. 75, p. 011801(R), 2007.

[85] Image provided by S. A. Webster

[86] G. W. C. Kaye and T. H. Laby, *Tables of Physical and Chemical Constants.*

[87] J. Alnis, A. Matveev, N. Kolachevsky, T. Udem, and T. W. Hänsch *Phys. Rev. A*, vol. 77, p. 053809, 2008.

[88] S. Pugla, *Ultrastable High Finesse Cavities for Laser Frequency Stabilization.* PhD thesis, Imperial College, University of London, UK, 2008.

[89] S. A. Webster *et al. Phys. Rev. A*, vol. 77, p. 033847, 2008.

[90] T. Legero, T. Kessler, and U. Sterr *J. Opt. Soc. Am. B*, vol. 27, pp. 914–9, 2010.

[91] S. A. Webster and P. Gill *Proceedings of the 2010 IEEE International Frequency Control Symposium*, pp. 470–3, 2010.

[92] B. C. Young, F. C. Cruz, W. M. Itano, and J. C. Bergquist *Phys. Rev. Lett.*, vol. 82, pp. 3799–802, 1999.

[93] Y. Y. Jiang *et al. Nature Photonics*, vol. 5, pp. 158–61, 2011.

[94] T. Kessler *et al. ArXiv e-prints*, Dec. 2011.

[95] M. Thorpe, L. Rippe, T. Fortier, M. Kirchner, and T. Rosenband *Nature Photonics*, vol. 5, pp. 688–93, 2011.

[96] D. Meiser, J. Ye, D. R. Carlson, and M. J. Holland *Physical Review Letters*, vol. 102, p. 163601, Apr. 2009.

[97] J. Bohnet, Z. Chen, J. Weiner, D. Meiser, M. Holland, and J. Thompson *Nature*, vol. 484, pp. 78–81, 2012.

[98] R. W. P. Drever *et al. Appl. Phys. B*, vol. 31, pp. 97–105, 1983.

[99] E. D. Black *Am. J. Phys.*, vol. 69, pp. 79–87, 2000.

[100] R. Fox, C. Oates, and L. Hollberg *Experimental Methods in the Physical Sciences: Cavity-Enhanced Spectroscopies*, vol. 40, pp. 1–46, 2003.

[101] C. Salomon, D. Hills, and J. Hall *J. Opt. Soc. Am.*, vol. B5, pp. 1576–87, 1988.

[102] C. J. Foot, *Atomic Physics.* Oxford University Press, 2005.

[103] N. F. Ramsey, *Molecular Beams.* Oxford University Press, 1956.

[104] V. Letchumanan, P. Gill, E. Riis, and A. G. Sinclair *Phys. Rev. A*, vol. 70, p. 033419, 2004.

[105] E. Peik *et al. J. Phys. B: At. Mol. Opt. Phys.*, vol. 39, pp. 145–158, 2006.

[106] R. Paschotta, *Encyclopedia of Laser Physics and Technology.* Wiley-VCH, 2008.

[107] D. J. Wineland *et al. J. Res. Natl. Inst. Stand. Technol.*, vol. 103, p. 259, 1998.

[108] I. I. Sobelman, *Atomic Spectra and Radiative Transitions.* Springer-Verlag, 1979.

[109] R. D. Cowan, *The Theory of Atomic Structure and Spectra.* University of California Press, 1981.

[110] G. P. Barwood private communication

[111] H. Dehmelt *Bull. Am. Phys. Soc.*, vol. 20, p. 60, 1975.

[112] E. Riis and A. G. Sinclair *J. Phys. B: At. Mol. Opt. Phys.*, vol. 37, p. 4719, 2004.

[113] J. R. P. Angel and P. G. H. Sandars *Proc. Roy. Soc. A.*, vol. 305, pp. 125–38, 1968.

[114] G. Woodgate, *Elementary Atomic Structure.* Oxford Science Publications, 1992.

[115] C. Tamm, D. Engelke, and V. Bühner *Phys. Rev. A*, vol. 61, p. 053405, 2000.

[116] K. Hosaka *et al. IEEE Trans. Inst. Meas.*, vol. 54, pp. 759–62, 2005.

[117] W. Itano *J. Res. Natl. Inst. Stand. Technol.*, vol. 105, pp. 829–37, 2000.

[118] T. Schneider, E. Peik, and C. Tamm *Phys. Rev. Lett.*, vol. 94, p. 230801, 2005.

[119] H. Muirhead, *The Special Theory of Relativity.* Macmillan Publishers Limited, 1973.

[120] C. W. Chou *et al. Science*, vol. 329, p. 5999, 2010.

[121] N. Lagakos, J. A. Bucaro, and J. Jarzynski *Appl. Optics*, vol. 20, pp. 2305–8, 1981.

[122] L.-S. Ma, P. Jungner, J. Ye, and J. L. Hall *Opt. Lett.*, vol. 19, pp. 1777–9, 1994.

[123] T. Udem, J. Reichert, R. Holzwarth, and T. W. Hänsch *Phys. Rev. Lett.*, vol. 82, pp. 3568–71, 1999.

[124] D. J. Jones *et al. Science*, vol. 288, pp. 635–639, 2000.

[125] H. S. Margolis *et al. Phys. Rev. A*, vol. 67, p. 032501, 2003.

[126] V. Tsatourian, H. S. Margolis, G. Marra, D. T. Reid, and P. Gill *Opt. Lett.*, vol. 35, pp. 1209–11, 2010.

[127] B. R. Walton, H. S. Margolis, V. Tsatourian, and P. Gill *IET Optoelectron.*, vol. 2, pp. 182–7, 2008.

[128] D. E. Spence, P. N. Kean, and W. Sibbett *Opt. Lett.*, vol. 16, pp. 42–44, 1991.

[129] L. Hollberg *et al. Metrologia*, vol. 42, pp. 105–24, 2005.

[130] S. T. Cundiff and J. Ye *Rev. Mod. Phys.*, vol. 75, pp. 325–42, 2003.

[131] J. Knight *Nature*, vol. 424, pp. 847–51, 2003.

[132] S. N. Lea, S. A. Webster, and G. P. Barwood *Proceedings of the 20th European Frequency and Time Forum (Braunschweig)*, pp. 302–7, 2006.

[133] V. I. Yudin *et al. Phys. Rev. A*, vol. 82, p. 011804(R), 2010.

[134] V. I. Yudin *et al. Phys. Rev. Lett.*, vol. 107, p. 030801, 2011.

[135] H. Schnatz *Meas. Sci. Technol.*, vol. 14, pp. 1200–15, 2003.