Optical frequency measurement and ground state cooling of single trapped Yb^+ ions

Peter Blythe

March 2004
with updates and corrections
August 2005

Submitted in partial fulfilment of the requirements
for the degree of Doctor of Philosophy
of the University of London

Abstract

The thesis describes experiments on single laser-cooled ions of Yb^+ in a radiofrequency ion trap.

The ion is laser-cooled on the 369 nm $^2\mathrm{S}_{1/2}\rightarrow{}^2\mathrm{P}_{1/2}$ electric dipole transition, and high resolution spectroscopy of the 467 nm $^2\mathrm{S}_{1/2}\rightarrow{}^2\mathrm{F}_{7/2}$ electric octupole 'clock' transition has been performed.

The 467 nm transition can be used as an optical frequency standard. To this end, several absolute optical frequency measurements of the $F = 0, m_F = 0 \rightarrow F = 3, m_F = 0$ component in $^{171}\mathrm{Yb}^+$ have been made with a femtosecond laser frequency comb generator. The comb was referenced to a hydrogen maser which forms part of the clock ensemble used to generate the UTC(NPL) timescale. During the work described in this thesis, the linewidth of the 467 nm probe laser has been narrowed from 4 kHz to 200 Hz, increasing the measurement resolution.

The frequency measurements have been supported by a full investigation of the systematic frequency shifts of the octupole transition, including the AC Stark, second-order Zeeman, quadrupole, DC Stark, second-order Doppler and blackbody shifts.

Cooling an ion of $^{172}\mathrm{Yb}^+$ to the ground motional state of the ion trap by 'EIT cooling' on the 369 nm transition has been investigated, both by numerical simulation and experiment. A new technique for measuring the temperature of a trapped ion without the use of a narrow transition has been developed. A new experimental control system has been implemented to assist in the EIT cooling experiments.

Acknowledgements

There are a number of people I must thank for their help, support, and encouragement, scientific or otherwise, over the past few years.

First of all, I must thank my thesis supervisors, Richard Thompson at Imperial College and Stephen Webster at NPL, particularly for all their help in the preparation of this thesis. Both have also been the source of many illuminating discussions, and have been a real help to me all through my studies. My first months at NPL were supervised by Paul Taylor, whom I must also thank. I have worked alongside Paul Taylor, Stephen Webster, Sang-Kyung Choi, and Kazumoto Hosaka on the experiment, and all have been a pleasure to work with.

Many other people at NPL have been helpful in a variety of ways. Patrick Gill, Helen Margolis, Stephen Lea, Hugh Klein, Geoff Barwood, Dick Rowley, Chris Edwards, Alastair Sinclair and Mark Oxborrow have all made significant contributions to the work described in this thesis, either through numerous suggestions and discussions or by direct intervention, as did previous student Alex Wilson, and Danny Segal at Imperial College. Thanks should also go to Tim Freegarde for his support and encouragement, and John Smith, for the same.

Many thanks must go to my fellow student Valliappan Letchumanan, with whom I have shared an office, a lab, a house, and the despair of thesis writing. He also deserves credit for innumerable useful scientific discussions.

On a personal note, my parents have always been incredibly supportive, and had far more belief in me than I have had myself, and I can only thank them for that. Thanks also to my sister, Marie, for her encouragement.

Finally, thanks to Esther, for everything.

Contents

Chapter 1

Introduction

A single atom, unperturbed by external fields, at rest, in a vacuum is a model system for studying atomic physics. A transition to a long-lived excited state in such an atom would be, in many ways, an ideal standard of frequency. The ideal is unattainable, but a single atomic ion trapped in a radiofrequency ion trap represents a very good approximation to it. The work described in this thesis is part of the process of improving that approximation, and of turning such a system – the 467 nm electric octupole transition in a single trapped ^{171}Yb^{+} ion – into a new standard of time and frequency.

Frequency, and therefore time, can currently be measured with greater accuracy than any other physical quantity. The highest levels of accuracy are now available across the electromagnetic spectrum from very low frequencies through the radio and microwave regions up to the optical and ultra-violet. The measurement of frequency underpins our daily lives in innumerable ways, from basic time-keeping through to radio and microwave communications (for television broadcasts, or mobile telephone systems), global positioning systems (such as GPS or the proposed european Galileo project), high-bandwidth optical communications, and the precision measurement of length and distance by interferometry, through the fixed definition of the speed of light. The benefits of frequency metrology are not confined merely to the technologies of modern life. The most immediate beneficiary of advances at the leading edge of frequency measurement is science. The ability to measure physical quantities with very high accuracy is the foundation on which fundamental physical theories rest. High-accuracy frequency metrology has been the means by which quantum electrodynamics (QED) has attained its current status as the most rigourously tested of our current basic theories, and it is in the fifteeth or sixteenth significant digits of such measurements that 'new' physics, beyond what we currently know, is expected to be found. One example of such new physics is the possible time-variation of the fundamental constants. Astronomical measurements have implied a variation of the fine-structure constant over the age of the universe [1], though these measurements remain controversial. Laboratory measurements have been able to put upper limits on present-day variation of the fine structure constant [2], but have found no evidence of actual variation as yet. The ytterbium octupole transition may be a good system to search for such variation [3, 4].

The whole of frequency metrology is underpinned by the existence of a relatively small number of ultra-high accuracy primary standards, which exist in national standards laboratories around the world. Currently, these are based on a microwave transition in neutral caesium atoms. Searches for new physics are approaching the limits of the accuracy that the caesium primary standard can provide [2]. Frequency standards based on atomic transitions at optical frequencies have the potential to improve on the accuracy of microwave standards by several orders of magnitude. The following chapters describe the progress so far, and the prospects for the future, of the development of the ytterbium electric octupole transition as a next-generation frequency standard.

This thesis covers work carried out at the National Physical Laboratory, Teddington, between October 2000 and November 2003. All of the experiments described here are based around single trapped ions of ytterbium. There are two distinct themes to the work, though they share common technology. The first of these is the development of an 'optical clock' based on the frequency of the 467 nm electric octupole transition in ^{171}Yb^{+}. The second is the study of a method of cooling a trapped particle to its motional ground state by exploiting quantum interference effects ('EIT cooling').

The experimental arrangement used for these experiments, and the techniques for control and data acquisition, are detailed in Chapter 3.

1.1 The caesium clock

The caesium beam clock was the first succesful atomic clock, building on earlier work on molecular beams (e.g. [5]). In 1955 it became the first atomic clock to contribute to a national time and frequency standard at NPL [6, 7], while a commercial caesium clock was developed and sold by a team at MIT. In 1967, the second was redefined 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,' and caesium has remained the primary standard of time and frequency ever since. While the early clocks were less complex than modern devices, they introduced all the essential features of atomic standards which remain with us to this day. Figure 1.1 shows a schematic representation of a caesium beam clock. The essential components of this, and all other atomic standards, are the local oscillator, the atomic reference, a feedback loop between the two, and a counter for output (the 'clock face'). In a beam clock, atoms in one of the two hyperfine levels ($F = 3$ and $F = 4$ – the $m_F = 0$ magnetic states in each are used for the reference) are selected from a beam by means of a magnetic gradient. The upper and lower states of the transition have different magnetic moments, and are physically separated by the field. The ground-state atoms then pass into the microwave cavity. Here, they interact with radiation from the local oscillator. This coherently excites a fraction of the atoms to the upper state, dependent on the frequency and intensity of the radiation. After the interaction region, the atoms are again state-selected and those in the upper state are detected. By locking the local oscillator to the centre of the fringe pattern thereby obtained, the transition frequency is transferred to the oscillator. Each cycle of the oscillator output is therefore one 'clock tick', and the counter then displays the passage of local time, directly referenced to the definition of the second.

The interaction region actually consists of two, separated interaction re-

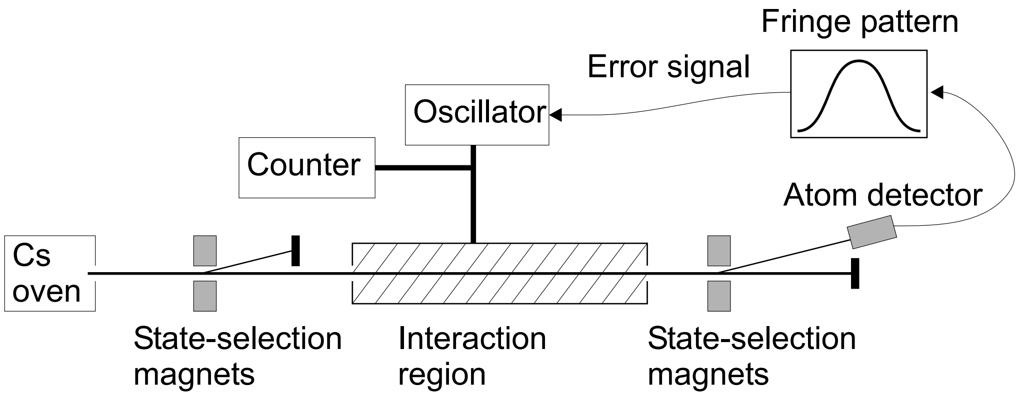

Figure 1.1: Schematic diagram of a caesium beam clock.

gions which are generally short compared to the distance (or time) between them. This “method of separated oscillatory fields” is also known as Ramsey spectroscopy [8], and has many useful properties as a spectroscopic technique. In particular, the width of the central fringe is inversely proportional to the time between the two interactions, without the need to maintain homogeneity and phase coherency of the microwave field over the entire size of the apparatus.

Modern caesium clocks use optical pumping to select the initial atomic state, and state-selective laser-induced flourescence to detect the proportion in either state after the microwave interrogation. The atoms are laser cooled, and launched vertically under gravity for periods of up to one second to maximise the interaction time and attain the highest possible frequency resolution, and to minimise systematic shifts of the transition frequency due to the atoms’ motion. The atoms pass through the microwave interaction region as they rise and the same region again as they fall. Such atomic fountain clocks have achieved accuracies of 1 part in 10^{15} [9], and the challenge for optical standards is to go beyond this level.

1.2 Trapped-ion frequency standards

The development of the ytterbium ion octupole transition as an optical frequency standard exists in the context of a much wider effort around the world to develop standards based on forbidden optical transitions in a number of trapped ion species. There are significant efforts to develop standards based on electric quadrupole transitions in Sr^+ [10, 11], Ca^+ [12], Hg^+ [13] and Yb^+ [14]. There is a growing interest in using narrow transitions in singly-ionised group III elements as standards [15], and an indium-based standard has been under development for quite some time [16]. The unique attribute of the 467 nm electric octupole transition in Yb^+ among these is its long natural lifetime, estimated at about 10 years. This affords the octupole transition the advantage that the spectroscopic resolution will always be limited by laser technology, and other technical issues, rather than the fundamental limit of the natural linewidth. This is analysed more fully in Chapter 2.

What the octupole transition shares with the other trapped-ion frequency standard transitions, is the prospect of extremely low systematic frequency shifts. The systematic shifts of the octupole transition are analysed in detail in Chapter 4.

The challenge currently faced by all trapped-ion frequency standard experiments is in realising the full potential of the standards. Traditionally one of the major obstacles to this was the difficulty in relating optical frequencies (hundreds of terahertz) to the frequency of the 9.2 GHz hyperfine transition in ^{133}Cs, which is the current primary standard of time and frequency. A major driving force behind the progress of optical frequency standards in the past few years has been the development of femtosecond laser frequency comb generators, or ‘femtosecond combs’, described in Chapter 2. They enable a broad range of optical frequencies to be coherently related to microwave frequencies in a single, relatively simple step. The has brought about an increase in the ease with which optical to microwave comparisons can be carried out, and reduced the uncertainties of such measurements by a large factor. Chapter 5 describes a series of measurements of the absolute frequency of the octupole transition in $^{171}Yb^+$, using the NPL femtosecond comb (described in Chapter 3).

1.3 Ground-state cooling

The motion of an ion in a radiofrequency trap is quantised, and can be treated as simple-harmonic motion. The ion has a semi-infinite ladder of motional states it might occupy, equally spaced in energy. The states can be labelled by the integer quantum number n, from $n = 0$ to ∞.

Ground-state cooling is the reduction of the ion's motion such that it is mostly in its ground motional state, $n = 0$. The general theory of ground-state cooling, and various techniques for achieving it are described in Chapter 2.

Cooling a trapped ion to its motional ground state can benefit many different types of experiment. For trapped-ion frequency standards, the focus of this thesis, it provides a method of reducing the size and uncertainty of systematic shifts such as the second-order Doppler shift, and quadratic DC Stark shift (Chapter 4). Cooling to the ground state also increases the contrast of coherent oscillations on narrow optical transitions. This improves the contrast, and hence obtainable frequency stability, when using Ramsey spectroscopy to interrogate a transition.

Because it prepares an ion in a known, pure, quantum state, ground-state cooling is a starting point for many other experiments. Much work on ground-state cooling has been motivated by its utility in trapped-ion quantum information processing (QIP) experiments. By initialising the motional state in $n = 0$, it is possible to use (for instance) the $n = 0$ and $n = 1$ states as a qubit ('quantum bit') for storing quantum information. Many interesting QIP-related experiments have been performed using single [17, 18] and multiple [19, 20] ground-state cooled ions.

'EIT cooling' is one method of cooling a particle confined in a harmonic trap to the ground motional state of the trapping potential. It was described [21], and then implemented [22], in late 2000 and exploits quantum intereference effects [23, 24] (such as electromagnetically induced transparency or EIT) in the cooling process. It has various advantages over other cooling

Isotope	**Abundance**	**Nuclear Spin**
^{168}Yb	0.1%	0
^{170}Yb	3.0%	0
^{171}Yb	14.3%	1/2
^{172}Yb	21.8%	0
^{173}Yb	16.1%	5/2
^{174}Yb	31.8%	0
^{176}Yb	12.8%	0

Table 1.1: Naturally occurring isotopes of ytterbium.

schemes, such as potentially being simpler, faster and more robust. The basic theory of EIT cooling is discussed in more depth in Chapter 2, along with possible way of exploiting the same underlying mechanism for measuring the temperature of a trapped ion.

In order to understand the EIT cooling scheme better, numerical modelling of the cooling process has been undertaken, and the details and results of this are given in Chapter 6. A scheme for implementing EIT cooling for ^{172}Yb^{+} ions has been developed, and the first results from the implementation are described in Chapter 7. These include single-ion EIT fluorescence spectra, and the observation of vibrational sidebands in the spectra.

1.4 Ytterbium

Ytterbium is a Lanthanide element, with an atomic number of 70. It is a soft, malleable metal at room temperature, with a melting point of 1097 K. It was first identified as an element in its own right in 1907, and is named after the Swedish village of Ytterby (along with yttrium, erbium and terbium) [25]. There are 7 naturally-occurring stable isotopes of ytterbium, shown in Table 1.1.

Both ^{171}Yb^{+} and ^{172}Yb^{+} are studied in the following experiments.

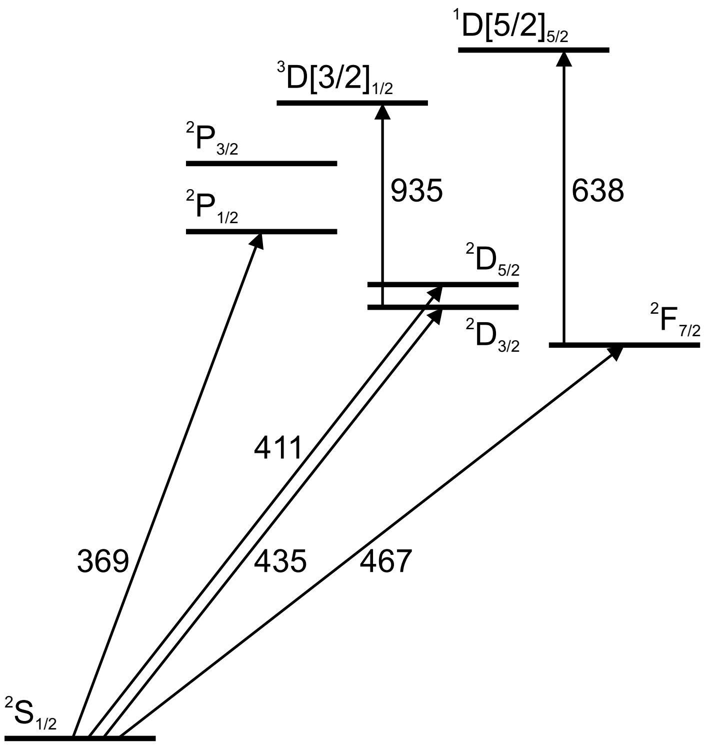

Figure 1.2: Gross energy level structure of the Yb$^+$ ion, approximately to scale. Arrows denote important transitions, with wavelengths shown in nanometres.

1.4.1 Gross energy level structure

Two isotopes of singly-ionised ytterbium are used in the work in this thesis. Both share a similar gross energy level structure, shown in Figure 1.2, and can therefore utilise the same laser systems. The unique feature of the ytterbium ion, compared to other commonly-used species, is that its lowest-lying excited state is the $^2F_{7/2}$ state. The only natural decay route open to this level is an electric octupole (E3) transition to the ground state, with a wavelength of 467 nm. Due to the angular momentum difference, this transition is strongly suppressed, and the $^2F_{7/2}$ state has an estimated lifetime on the order of ten years [26]. The natural linewidth of the transition is therefore on the order of nanohertz, and the decay has a $Q \sim 10^{23}$ - the highest known Q for an optical transition.

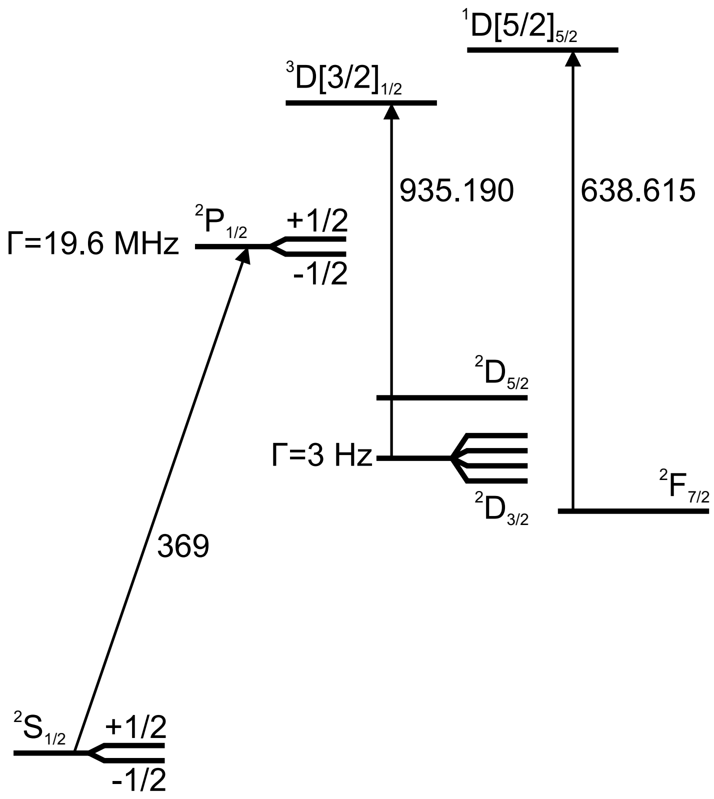

Figure 1.3: Energy levels of ^{172}Yb^{+}. The Zeeman structure of the ^{2}S$_{1/2}$, ^{2}P$_{1/2}$ and ^{2}D$_{3/2}$ levels is indicated.

1.4.2 ^{172}Yb^{+}

The even isotope of the ytterbium ion, ^{172}Yb^{+}, is the simpler of the two used in this work as it has no hyperfine structure, and simple Zeeman structure, shown in Figure 1.3. Due to its simple structure, ^{172}Yb^{+} can be laser-cooled (Section 2.4) with a relatively simple laser arrangement and without additional magnetic fields above the earth's field. A laser on the 369 nm transition is required for Doppler cooling, and a laser at 935 nm is necessary to maintain the cooling cycle when the ion falls from the ^{2}P$_{1/2}$ state to the ^{2}D$_{3/2}$ state, which occurs with a branching ratio of 1:150 [27, 28]. With the presence of the 935 nm laser, the cooling cycle is closed. However, if the ion undergoes a collision with a background gas particle during the cooling, it can be transferred to the ^{2}F$_{7/2}$ state. Once there the ion is effictively lost, due to the state's long lifetime. For this reason, a laser at 638 nm can be used to promote the ion to a higher state from which it will return to the cooling cycle via either of the D-states.

1.4.3 ^{171}Yb$^+$

The structure of the energy levels of ^{171}Yb$^+$ is shown in Figure 1.4. The attraction of ^{171}Yb$^+$ as a frequency standard is that it has half-integer nuclear spin, and so the ^{2}S$_{1/2}$ – ^{2}F$_{7/2}$ transition has a $m_F = 0 \rightarrow m_F = 0$ component, which is free from linear Zeeman shift. This isotope is preferred over ^{173}Yb$^+$ due to its simpler Zeeman structure.

Due to the hyperfine structure, ^{171}Yb$^+$ is more complex to laser cool than ^{172}Yb$^+$. The ^{2}S$_{1/2}(F = 1)$ – ^{2}P$_{1/2}(F = 0)$ transition is chosen for Doppler cooling, as the ^{2}P$_{1/2}(F = 0)$ to ^{2}S$_{1/2}(F = 0)$ decay is strictly forbidden. This, along with the 935 nm laser on the ^{2}D$_{3/2}(F = 1)$ to ^{3}D[3/2]$_{1/2}(F = 0)$ transition, should be sufficient to maintain a closed cooling cycle, as the ^{2}P$_{1/2}(F = 0)$ level can only decay to the $F = 1$ D-state, and the ^{3}D[3/2]$_{1/2}(F = 0)$ to ^{2}S$_{1/2}(F = 0)$ is forbidden. However, there is a non-zero probability of the 369 nm laser exciting a spontaneous Raman transition via the ^{2}P$_{1/2}(F = 1)$ state, to either the ^{2}D$_{3/2}(F = 2)$ level or the ^{2}S$_{1/2}(F = 0)$ level. The 935 nm laser is not resonant with transitions from the ^{2}D$_{3/2}(F = 2)$ level, and so the ion ceases fluorescing for the 52 ms [28] lifetime of this state. This complicates the state-detection scheme for this isotope (Section 3.5). The ion can also leave the cooling cycle by excitation to the ^{2}S$_{1/2}(F = 0)$ level. It is therefore necessary to return the ion to the cooling cycle, and this can be done either with 12.6 GHz microwaves, or with a second laser tuned to the other hyperfine transition. In this experiment, microwaves are used. The presence of hyperfine structure also means that two wavelengths (around 638 nm) are required to return the ion to the ground state from the ^{2}F$_{7/2}(F = 3)$ state which is used as a frequency standard. This is because the intermediate ^{1}D[5/2]$_{5/2}$ state can decay to the ^{2}F$_{7/2}(F = 4)$ level.

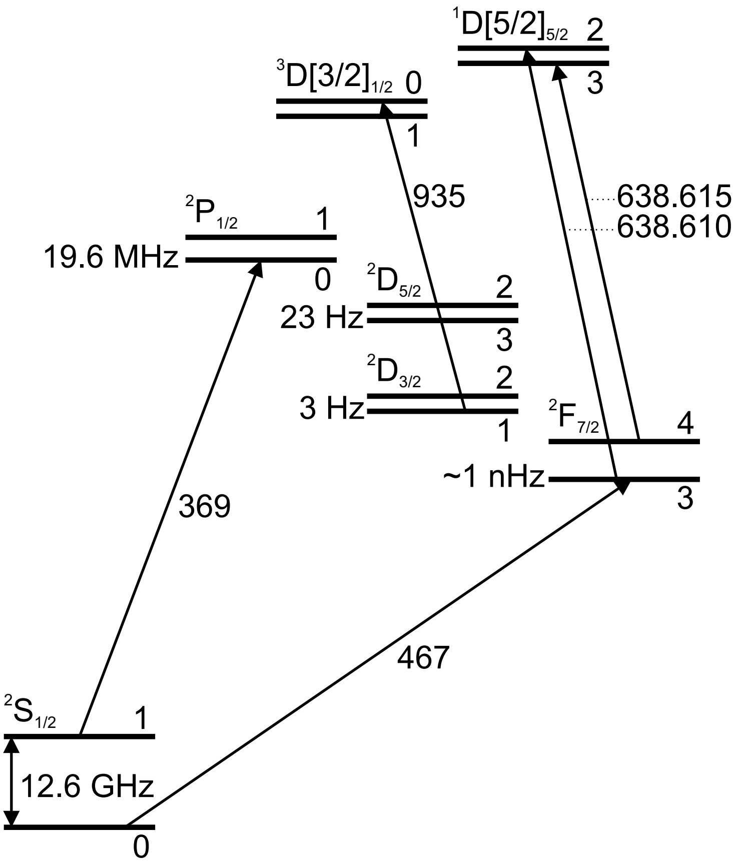

Figure 1.4: Hyperfine structure of ^{171}Yb$^+$. Levels are designated by their total angular momentum F. Transition wavelengths are in nanometres. Natural spontaneous decay linewidths are indicated for the ^{2}P$_{1/2}$, ^{2}D$_{3/2}$, and ^{2}F$_{7/2}$ states.

Chapter 2

Theoretical background

This chapter covers the basic ideas and theory behind the main themes of this thesis — ion trapping, optical frequency standards and ground–state cooling.

2.1 Ion trap theory

2.1.1 Radiofrequency ion trapping

To manipulate and study a single ion for a significant period of time, it is obviously necessary to confine it. However, it is not possible to construct an electrostatic potential which is confining in all three dimensions due to the requirement on the electric potential Φ that

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

The energy of a negatively charged particle, with charge $-e$ in such a potential is

$$V = -e\Phi. \tag{2.2}$$

To be a trapping potential

$$\frac{\partial^2 V}{\partial x_i^2} > 0, \tag{2.3}$$

for $x_i = \{x, y, z\}$, so a potential may only be confining in two dimensions simultaneously due to Equation 2.1. To overcome this problem, radiofrequency (Paul) traps have been developed, in which the potential oscillates rapidly between trapping in the z and x, y directions.

One of the simplest solutions to Equation 2.1 is a quadrupolar potential with cylindrical symmetry. In cylindrical coordinates:

$$\Phi = \frac{C}{2r_0^2}\left(r^2 - 2z^2\right), \tag{2.4}$$

where C represents the strength of the potential, and r_0 represents the characteristic size of the electrode structure producing the potential. The simplest electrode structure is the original Paul trap [29] design, in which the surfaces of a ring electrode and two hyperbolic 'endcaps' follow the equipotential surfaces of a pure quadrupole potential. The structure of the trap used in this thesis is described in Section 3.1. Dependent on the sign of C, the potential confines in either the radial (r) direction or the axial (z) direction. If C is a radio-frequency potential, such as

$$C = \mathcal{E}\left(U_{\mathrm{DC}} + U_{\mathrm{RF}}\cos(\omega_{\mathrm{RF}}t)\right), \tag{2.5}$$

then the potential oscillates between confining in the r and z directions. The parameter $\mathcal{E}$ is characteristic to a particular trap geometry and is known as the efficiency of the trap. The efficiency relates the potentials applied to the trap electrodes U_{DC} and U_{RF} to the field experienced by an ion near the centre of the trap. The efficiency of the original Paul trap is 1. The equations of motion of a charged particle, of mass m and charge Q, in such a field can be described in terms of the a and q parameters [30, 27]:

$$a_z = -2a_r = -\frac{8Q\mathcal{E}U_{\mathrm{DC}}}{mr_0^2\omega_{\mathrm{RF}}^2} \tag{2.6}$$

$$q_z = -2q_r = +\frac{4Q\mathcal{E}U_{\mathrm{RF}}}{mr_0^2\omega_{\mathrm{RF}}^2} \tag{2.7}$$

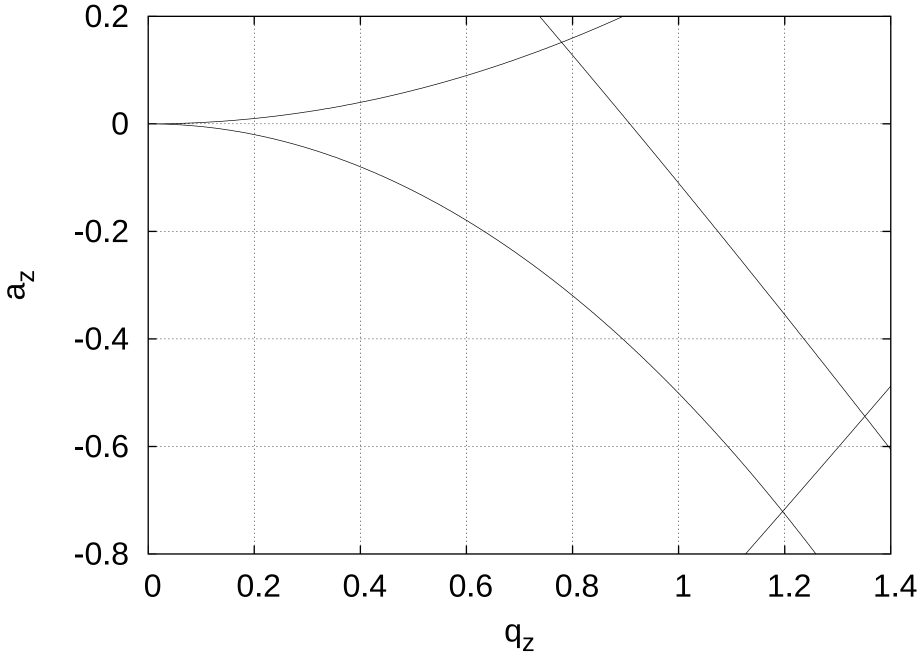

Figure 2.1: The region of stability for a cylindrically symmetric quadrupole ion trap. The motion of a trapped ion is stable within the region bounded by the four lines.

The equations of motion of the particle are of the form of Mathieu equations

$$\frac{\mathrm{d}^2 x_i}{\mathrm{d}t^2} + \frac{\omega_{\mathrm{RF}}^2}{4}(a_i - 2q_i \cos\omega_{\mathrm{RF}} t)x_i = 0 \tag{2.8}$$

where $x_i = \{z, r\}$. These equations have stable solutions for a range of a and q parameters, shown in Figure 2.1. A far more thorough treatment is available in [31].

In the region of stability, the motion of the ion can be separated into several components. For the z co-ordinate, the motion of an ion with mean position $\bar{z}$ is, assuming that z_a and $\bar{z}$ are ‘small’,

$$z = \bar{z} + z_a \cos\omega_z t - \frac{z_a q_z}{4}\left[\cos(\omega_{\mathrm{RF}} + \omega_z)t + \cos(\omega_{\mathrm{RF}} - \omega_z)t\right] - \frac{\bar{z} q_z}{2}\cos\omega_{\mathrm{RF}} t \tag{2.9}$$

Taking the terms in order, the first term $\bar{z}$ simply represents the mean position of the ion in the quadrupole potential, which is centred on $z = 0$. The second term is the *secular motion* (or thermal motion) of the ion. The frequency ω_z is known as the *secular frequency.* It is much smaller than the drive frequency ω_{RF}, and it corresponds to the motion (with amplitude z_a) of the ion in an effective harmonic potential known as the pseudopotential. The pseudopotential arises from the time-averaged force experienced by the ion in the trap, and is described in more detail in Section 2.1.2 below. The third term represents the *intrinsic micromotion* of the ion, a driven motion (by the RF trapping field) whose amplitude is related to the amplitude of the secular motion. In fact, this motion plays the role of the potential for the secular motion. The energy of the ion oscillates between the kinetic energy of the secular motion and the energy (kinetic and potential) of the intrinsic micromotion. The final term represents the *excess micromotion* of the ion, and is a purely driven motion. It occurs when a DC potential displaces the ion from the centre of the RF quadrupole field. This motion is examined further in Section 2.1.4.

2.1.2 The harmonic oscillator pseudopotential

The pseudopotential can be approximated to be a harmonic oscillator potential in the region of the ion's motion. That is $V(z) = \frac{1}{2}m\omega_{\mathrm{z}}^2 z^2$ in the axial direction, and similarly in the radial plane. It is the pseudopotential which confines the ion in a radiofrequency ion trap. The atom's motion is then quantised and only able to occupy states of the harmonic oscillator, which form an evenly spaced ladder of energy levels:

$$E_n = \hbar\omega_{\mathrm{sec}}\left(n + \frac{1}{2}\right). \tag{2.10}$$

Here ω_{sec} represents either ω_r or ω_z. The secular frequencies are determined by the relation

$$\omega_{\mathrm{sec}} = \frac{\beta\omega_{\mathrm{RF}}}{2}. \tag{2.11}$$

The parameter β is a function of the a and q parameters (Equations 2.6 and 2.7). The solution is a non-terminating expansion in powers of q. To fourth order in q, for high-precision calculations (such as those in Section 4.3.1), the expression is [32]

$$\beta = \left(a - \frac{a-1}{2(a-1)^2 - q^2} q^2 - \frac{5a+7}{32(a-1)^3(a-4)} q^4 \right)^{\frac{1}{2}} \qquad (2.12)$$

For most purposes, the lowest order approximation is sufficient:

$$\beta = \sqrt{a + \frac{q^2}{2}}. \qquad (2.13)$$

Using Equations 2.6 and 2.7 with zero applied DC potential, the ratio of the trap frequencies ω_z/ω_r is expected to be 2 using the low-order approximation.

The motion of the atom in this harmonic potential leads to sidebands in its spectrum, spaced at the secular frequency. In the classical picture, these sidebands are due to the modulation of the transition frequency by the Doppler shift due to the motion of the ion in the pseudopotential. At low vibrational quantum numbers this picture breaks down as discussed in Section 2.4.6. An example spectrum is shown in Figure 2.2. There is in fact a whole spectrum of sidebands at multiples of the secular frequency from the carrier, but these are small when the ion is in the Lamb-Dicke regime (Section 2.1.3).

2.1.3 The Lamb-Dicke Parameter

The Lamb-Dicke parameter occurs frequently in the theory of trapped ions, and is related to classical modulation theory. It is a special case of the classical modulation index β_{mod}, for the ground vibrational state of a trapped particle.

Classically, the height of any optical sidebands relative to their carrier transition depends on the amplitude of the ion's motion relative to the wavelength of the light probing the transition. This is expressed in a modulation index, β_{mod}, defined by

$$\beta_{mod} = kx_0. \qquad (2.14)$$

where x_0 is the amplitude of the ion's motion, and k is the wavevector of the

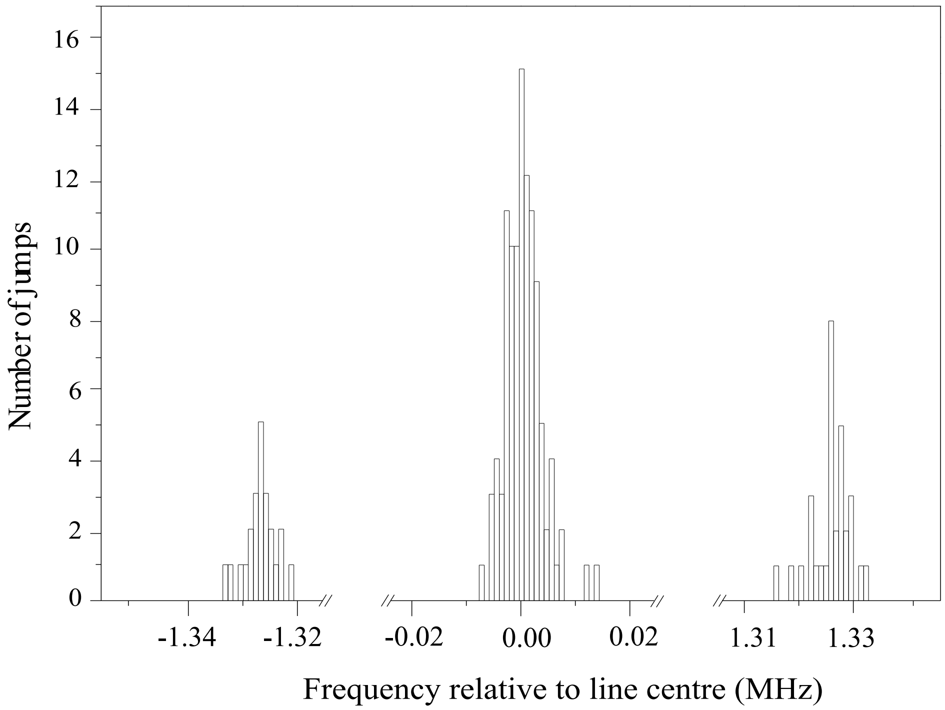

Figure 2.2: Spectrum of the radial sidebands of the 467 nm transition in a single trapped Yb^+ ion. The spectrum was recorded using the 'quantum jumps' technique (Section 2.3). A DC voltage had been applied to the endcaps to equalise the axial and radial frequencies. Figure taken from [33].

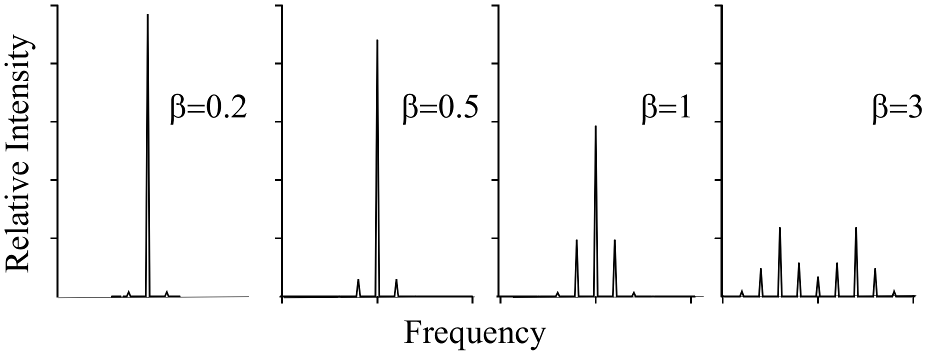

Figure 2.3: Variation of classical sideband height with modulation index β_{mod}. Figure taken from [34].

interrogating light. The height of the sidebands relative to the carrier varies with β_{mod} as shown in Figure 2.3.

If β_{mod} is less than 1, the ion is said to be in the *Lamb-Dicke regime.* In this regime, only first-order sidebands are of a significant size. Also, the carrier is a Doppler-free feature and so the first-order Doppler effect is eliminated.

In the Lamb-Dicke regime, the quantisation of the ion's secular motion becomes important. The heights of the sidebands then correspond to the transition strengths between different vibrational energy levels in the harmonic pseudopotential. The sideband structure shown in Figure 2.3 is modified slightly in this regime, as described in Section 2.4.6. Most obviously, when in the ground state, there is no red sideband, as there is no lower vibrational state to couple to.

The Lamb-Dicke parameter η as defined in this thesis, and as commonly used in the literature, is proportional to the classical modulation index of the ground vibrational state[1]. It is a function of particle mass, trap frequency, and the wavelength of the interrogating light:

$$\eta = k\sqrt{\frac{\hbar}{2m\omega_{\text{sec}}}}, \tag{2.15}$$

---

[1] Based on the most common definitions of both these quantities, $\eta^2 = \beta_{mod}^2(n=0)/2$.

where ω_{sec} is the secular frequency of the trap, m the mass of the ion and k the wavevector of the incoming light. It is not a function of the temperature of the particle, or its vibrational state. If the light is at an angle θ to the direction of motion, k is replaced by $|k| \cos\theta$.

2.1.4 Micromotion

Micromotion is the motion of the ion driven by the oscillating potential, and was introduced in Section 2.1.1. It arises from the ion's instantaneous position not coinciding with the saddle point of the oscillating trapping fields ($r = z = 0$), either because of the secular motion of the ion (intrinsic micromotion) or because of stray DC potentials displacing the equilibrium point of the ion's motion (excess micromotion). Micromotion cannot be reduced by laser cooling, though the intrinsic micromotion is a function of the secular motion, which can be laser-cooled. Excess micromotion can be compensated by application of DC potentials to compensation electrodes around the ion trap, to move the ion back to the centre of the RF field by nulling the stray fields. A phase mismatch of the RF trapping field between the electrodes will also cause micromotion which is not compensatable with external DC potentials, but can be compensated by applying an extra variable-phase AC potential to one of the endcaps.

One manifestation of the micromotion is sidebands in the spectrum of the ion, spaced at the driving frequency of the trap away from the carrier. These occur as the transition is modulated in frequency by the Doppler shift due to the ion's motion. The size of the micromotion can be more accurately measured (and hence minimised) using an rf-photon correlation technique. There are several ways to detect micromotion, detailed in Section 3.3.1.

2.2 Optical frequency standards

For any frequency standard, there are three quantities to be considered – stability, reproducibility and accuracy. To summarise briefly, stability is the variation in frequency of a single source, accuracy is a measure of how closely the source realises some ideal (such as an unperturbed atom in a perfect vacuum), and reproducibility is the typical frequency difference between two sources of the same design.

For many applications, a highly stable frequency source is sufficient, but a frequency *standard* should also have excellent accuracy and reproducibility.

2.2.1 Frequency stability

Allan deviation

The frequency stability (or instability) of any source may be characterised by its Allan deviation $\sigma_y(\tau)$, which provides a measure of the frequency variation between successive independent measurements with an averaging time τ [35]. For a mean frequency ν_0, and independent measurements $i = 0, 1 \ldots N-1$, the fractional frequency instability is

$$\sigma_y^2(\tau) = \frac{1}{2(N-1)} \sum_{i=1}^{N-1} \frac{(\bar{\nu}_i - \bar{\nu}_{i-1})^2}{\nu_0^2} \tag{2.16}$$

The dependence on τ arises because for an individual measurement time T_R, $\tau = NT_R$. For any frequency standard it is obviously desirable to decrease the measurement-to-measurement variation as much as possible. This is the path which the 9.2 GHz caesium standard (the primary time standard) has been following since its inception. However, both systematic frequency shifts, particularly collisional shifts, and the difficulty of obtaining long interaction times, have begun to limit the accuracy and stability (respectively) of the caesium standard [9]. One obvious way to improve the available fractional stability is to use an atomic transition with a higher ν_0. This is the rationale behind the development of optical frequency standards. It should be possible to achieve the same fractional stability for shorter averaging times. Another motivating factor is that with a wise choice of reference transition it should be possible to obtain much smaller fractional systematic frequency shifts.

Typical Allan deviation curves

An idealised Allan deviation curve has several regions, each characterised by a particular slope. Typically, for atomic frequency standards, Ramsey spectroscopy [36, 37] is used to measure the transition frequency. A measurement with averaging time τ can therefore be a single interrogation with Ramsey interrogation time $T_R = \tau$, or the average of several measurements of interrogation time $T_R < \tau$. For short timescales, where the standard is interrogation-time limited, $\sigma_y(\tau)$ decreases as τ^{-1}. For longer averaging times, where the Ramsey interrogation time T_R is no longer increasing, the increase in stability comes from averaging over many measurements, and the Allan deviation varies as $\tau^{-1/2}$. At a certain level, most standards become limited by 'flicker' noise, with a $1/f$ noise spectrum, possibly due to fluctuating systematic shifts of the frequency. This produces a flat line on an Allan deviation plot. At longer timescales still, the Allan deviation will increase, due to longer-term drifts in the systematic frequency shifts.

Atomic standards

For the case of an atomic transition, interrogated by time-domain Ramsey spectroscopy [36, 37], the quantum-limited frequency stability is [13]

$$\sigma_y(\tau) = \frac{1}{2\pi\nu_0\sqrt{N_a T_R \tau}}, \tag{2.17}$$

where N_a is the number of atoms being interrogated (with assumed 100% detection efficiency) and T_R is the time period between Ramsey pulses. The work described in this thesis utilises single trapped ytterbium ions, for which $N_a = 1$ and $\nu_0 = 642$ THz (Section 1.4.1). If the atoms are entangled in a spin-squeezed state, the Allan deviation reduces linearly with N_a [15].

There are several possible limitations to the Ramsey interrogation time T_R. It should not exceed the timescale of any decoherence processes, for optimum stability, because the fringe contrast is reduced. This means that both the atomic state lifetime and laser linewidth are limitations for optical standards. The natural lifetime is a more fundamental limitation, while the laser linewidth is a technical limitation. In an atomic fountain, where the atoms are launched vertically, the time between microwave interactions is limited by gravity and the size of the equipment, to typically 1 s. An ion trap can confine a single particle for very long periods of time (at least several days) without cooling, and essentially indefinitely with regular laser cooling.

Long-timescale Fourier-limited optical standards

A Fourier-limited standard is one in which there is no loss of coherence due to natural decay or laser-induced decoherence. All standards are Fourier-limited over short timescales due to the short interrogation times available. For standards with very long excited-state lifetimes, this regime can persist to long measurement times. For all practical purposes, the lifetimes of the excited states in the caesium microwave and ytterbium octupole optical standards are infinite. In this case, the Ramsey interrogation time T_R is confined only to be less than the total measurement time τ ($T_R \leq \tau$). Figure 2.4 shows approximate projected stabilities for a caesium fountain microwave standard, an optical transition with a limited interrogation time and an optical transition with a long lifetime . The frequency instabilities of all the standards initially decrease as τ^{-1}. The optical standard with a limited interaction time breaks from the τ^{-1} line at 1 s, and continues with a $\tau^{-1/2}$ slope, characteristic of averaging over multiple interrogations. The long lifetime standard, however, continues to decrease as τ^{-1} until it reaches the floor due to systematic shifts. The graph clearly demonstrates the attraction of a long-lifetime standard as the primary definition of time, as full advantage can always be taken of any reductions in systematic shifts, provided that the laser linewidth is not a limiting factor. Standards with limited interaction times can only improve as $\tau^{-1/2}$ or $N^{-1/2}$ as systematic shifts are minimised.

The graph also shows the benefits of using an optical, or even higher frequency, reference over a microwave standard. It requires billions of atoms (with unit detection efficiency) for a microwave standard to approach the short-term stability of a single-particle optical standard. A long-lifetime optical standard (with equally narrow probe laser) has an intrinsic uncertainty

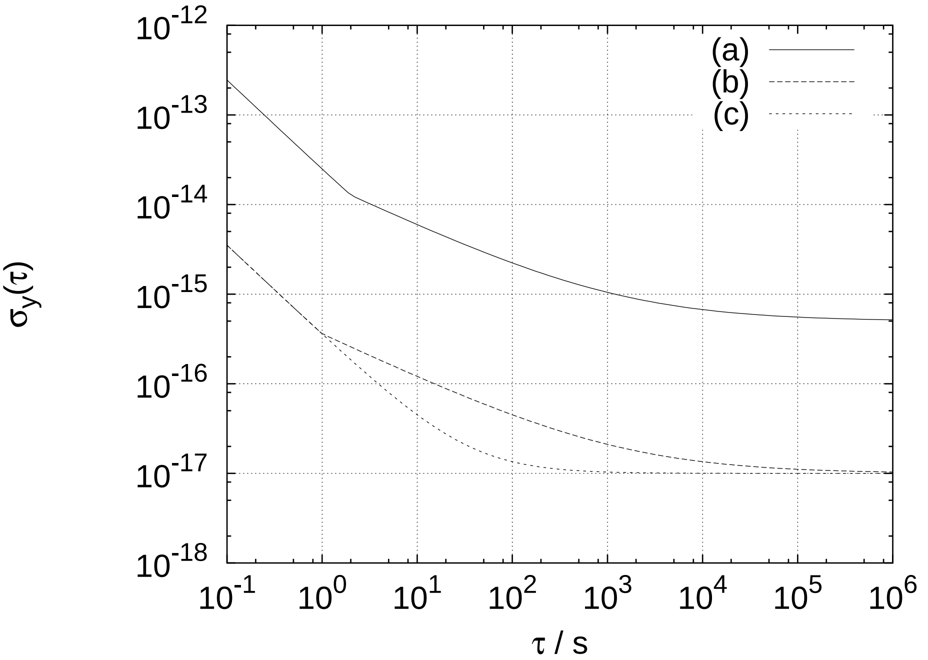

Figure 2.4: Approximate stabilities of frequency standards. (a) A 9.2 GHz caesium standard with 10^6 atoms, and unit state-detection efficiency. (b) A single-atom optical standard at 467 nm with an interrogation time limited to 0.5 s. (c) An infinite-lifetime single-atom optical standard at 467 nm. The microwave standard is limited to a 1 s interrogation time due to the difficulty in obtaining longer interaction times in fountain-type standards. The optical standards have baseline systematic uncertainties of 1×10^{-17} in this graph, and the caesium standard 5×10^{-16}. Laser decoherence is ignored as a limiting factor for the optical standards.

equal to the baseline systematic shifts (taken as 1 part in 10^{17} in this example) at around 1 minute, the interaction-time limited optical standard reaches the same level in 1000 seconds ($\sim$ quarter of an hour), while the many-atom microwave standard requires more than 10^6 s (almost two weeks) to reach that level.

It should be noted that there is a significant movement towards building optical clocks using large numbers of neutral atoms confined in an optical lattice, and these could again offer improvements in short-term stability over single-particle clocks, if high detection efficiency could be achieved.

For a primary standard of time (or, equivalently, frequency), the short-term stability offered by Fourier-limited optical standards is less important than having very small baseline systematic shifts, and therefore excellent accuracy and long-term stability. It is always possible to use some other secondary standard or ‘local flywheel’ to obtain short-term stability where it is needed. The biggest attraction of trapped-ion optical frequency standards in general is the prospect of very low fractional systematic shifts (Chapter 4). It is still philosophically appealing, however, to be able to use a single reference standard for all (or most) timescales without resort to a secondary standard.

2.2.2 Obtaining ‘ticks’ - the femtosecond comb

One of the fundamental problems with using an optical frequency standard as a primary definition of frequency (or, equivalently, time) until recently was the difficulty in counting the ‘ticks’ of the clock. There are no electronic methods of directly counting an optical frequency, unlike in the microwave domain.

This problem necessitated the construction of frequency chains (e.g [38]), in which microwave frequencies were coherently related to optical frequencies via a series of harmonic-generation and frequency-mixing processes. The devices were extremely complicated, and were built to measure specific optical frequencies.

The situation improved dramatically with the development of the ‘femtosecond comb’. The heart of such a device is a mode-locked femtosecond laser. The output is a series of regularly spaced, phase-coherent pulses, of typically a few tens of femtoseconds duration. This is illustrated in Figure 2.5. The output can be characterised by three parameters – the repetition rate of the laser, f_R (with associated time t_R), the pulse length t_{pulse}, and a third parameter known as the *carrier envelope offset frequency* f_{ceo}.

The Fourier transform of a series of regularly-spaced delta functions is another series of regularly spaced delta functions. By analogy, it is easy to see that the Fourier transform of the output of a mode-locked femtosecond laser

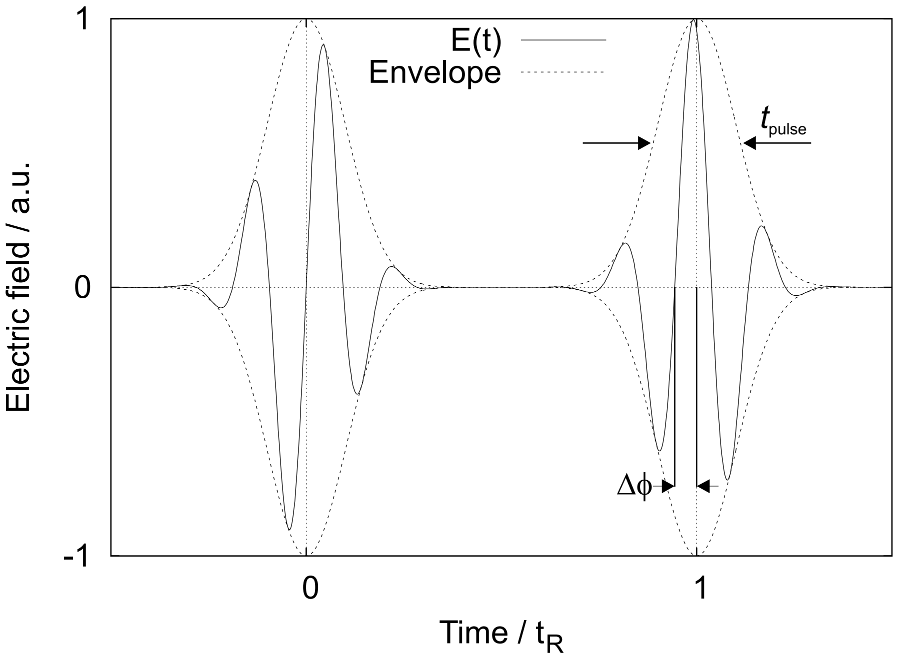

Figure 2.5: Output of a femtosecond comb. $\Delta\phi$ is the phase slippage between successive pulses.

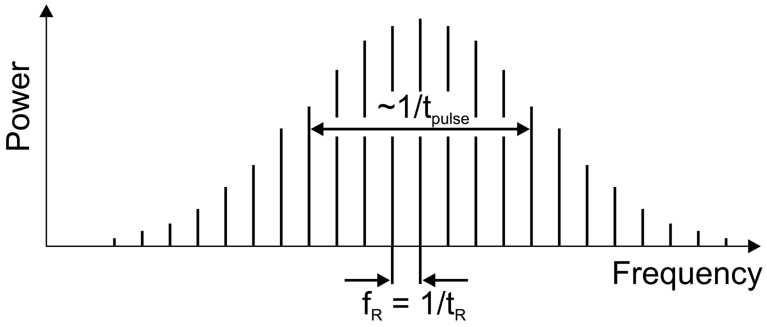

Figure 2.6: The spectral output of a mode-locked femtosecond laser is a series of modes spaced by the repetition frequency f_R, with an overall frequency width set by the duration of an individual laser pulse t_{pulse}, centred on the optical frequency of the laser output.

(in which each pulse can be thought of as approximating a delta function) will be a series of regular peaks in frequency space. The idealised 'delta function' picture is modified by the non-zero pulse width of the laser, and the result is shown in Figure 2.6. Centred on the optical frequency of the laser output, a 'comb' of frequencies spaced by the repetition frequency f_R is obtained, within an envelope of width on the order of $1/t_{\text{pulse}}$. This is typically a width of around 30 nm.

The light is focused into a section of 'holey fibre' - an optical fibre in which the core is surrounded by a microstructured arrangement of air holes. The fibre is highly non-linear, and it facilitates a variety of mixing processes between the modes of the laser output. In this way, the 30 nm span of the laser output can be broadened by the mixing to cover a complete octave in optical frequency - from wavelengths longer than 1 μm to shorter than 500 nm. This new spectrum is still composed of discrete, coherent, modes with the same frequency spacing as before.

It is then possible to measure the frequency of a laser anywhere in this bandwidth by beating the light from the laser to be measured against the femtosecond comb and measuring the beat frequency f_Δ. The process is

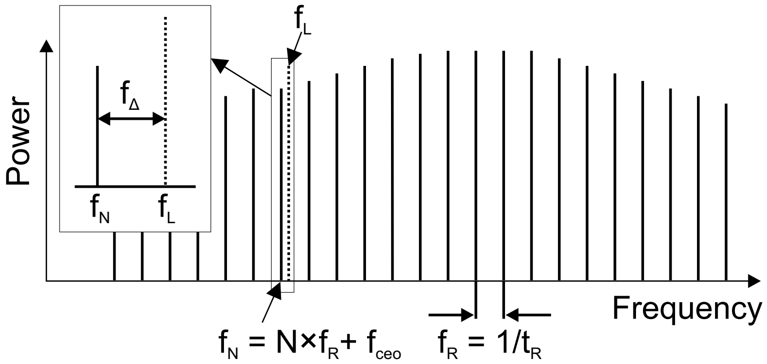

Figure 2.7: Measurement of a laser frequency f_L using a femtosecond comb. The beat frequency between the comb mode and the laser is f_Δ.

illustrated in Figure 2.7. The frequency of the N^{th} comb mode is given by

$$f_N = N \times f_R + f_{\text{ceo}} \tag{2.18}$$

It can be seen that f_{ceo} is the 'zero-offset' of the frequency comb. Its source is the phase slippage $\Delta\phi$ between successive laser pulses (Figure 2.5). The frequency of the laser is given by

$$f_L = N \times f_R + f_{\text{ceo}} \pm f_\Delta \tag{2.19}$$

If the frequency of the laser is known to better than f_R from some other method of measurement, then f_L can be uniquely determined in terms of f_R. If the frequency is not roughly known beforehand, then it can still be determined by varying f_R and f_{ceo}. If f_R and f_{ceo} are either stabilised to, or counted in reference to, a stable frequency derived from, say, a caesium fountain, then it is possible to determine optical frequencies directly in terms of the primary definition of the SI second in a single step. Methods for measuring f_R and f_{ceo} are described in Section 3.11.

2.3 Quantum jumps

To observe a narrow (and therefore weak) transition in a single ion is very difficult with direct laser fluorescence or absorption techniques. Once excited to the metastable state, the atom will typically decay after a few hundred milliseconds, and then will emit only a single photon. Hence, the fluorescence rate of the transition is extremely low. Also only a fraction of the solid angle into which photons are emitted can practicably be covered by the detection optics, and it is extremely difficult to detect the emitted light over a background of scattered light at the same frequency from the probe laser.

The solution to this problem is the technique of electron shelving, or *quantum jumps* [39, 40]. In this scheme, the atom is probed with a narrow linewidth, tunable laser. To detect whether the atom has made a transition to the metastable state, a laser corresponding to an allowed electric dipole transition also illuminates the atom. Resonance fluorescence from this transition, which can easily be detected, is monitored. If the atom does ‘jump’ to the metastable state, the electron is no longer available for cycling through the dipole transition, and so the resonance fluorescence ceases. Once the atom is returned to the ground state, the fluorescence reappears. A sequence of quantum jumps is shown in Figure 2.8. By monitoring the frequency of these jumps as the narrow probe laser is tuned, the spectrum of the weak transition may be recorded. Examples of these spectra are shown in Figures 2.2, 3.7, and 3.8

2.4 Ground-state cooling

2.4.1 A definition of ground-state cooling

Any definition of ground-state cooling is somewhat arbitrary, but it is usually taken to mean that there is a substantial probability of a particle being in the $n = 0$ vibrational state. For the following discussions, $\bar{n} < 1$ is taken as defining the regime of ground-state cooling.

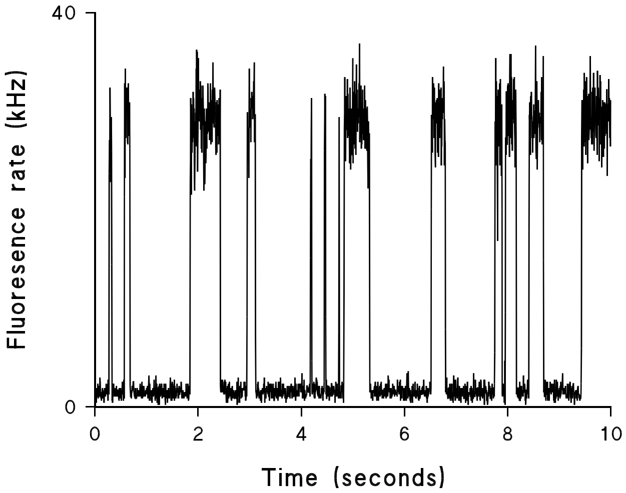

Figure 2.8: Quantum jumps of a single ion, taken from [34]. These are jumps to the $^2F_{7/2}$ state in $^{172}Yb^+$ with a repumper laser at 638 nm constantly illuminating the ion to depopulate the state.

2.4.2 Doppler cooling

To understand the attraction of ground-state cooling, it is necessary to appreciate the performance of alternative cooling methods. The simplest (and most used) method of laser cooling is Doppler cooling. This forms the baseline from which ground-state cooling can progress. A laser is tuned to the red (low-frequency) side of an allowed dipole transition, with a linewidth $\Gamma \gg \omega_{\text{sec}}$ (this is known as the weak binding regime). When the atom is moving towards the laser, the photons are blue-shifted (in the atom's frame) by the Doppler effect and the ion absorbs photons more frequently. On absorption, the photon's momentum is transferred to the atom, slowing it down. When the atom decays, the corresponding photon is emitted in a random direction, on average transferring no momentum to the atom. When the atom is moving away from the laser, any absorbed photons increase the atom's momentum, but the rate of absorption decreases due to the increased detuning of the photons from the transition caused by the Doppler shift. On average, the atom is cooled. For a transition of (FWHM) linewidth Γ, optimum cooling occurs for a detuning $\Gamma/2$. In this situation, the lowest energy per degree of freedom that can be reached is [41, 42]

$$E_{\text{Doppler}} = \frac{\hbar\Gamma}{4} \tag{2.20}$$

For a trapped ion in a harmonic potential, with average vibrational quantum number $\bar{n}$, the total (kinetic and potential) energy is

$$E_{\text{Trap}} = \left(\bar{n} + \frac{1}{2}\right)\hbar\omega_{\text{sec}} \tag{2.21}$$

The kinetic energy of the Doppler limit can be equated with the kinetic energy of a trapped particle to give the limit in terms of n. Then, the average vibrational excitation number for each degree of freedom is given by

[42]

$$\bar{n} = \frac{2E_{\text{Doppler}}}{\hbar\omega_{\text{sec}}} \tag{2.22}$$

$$= \frac{\Gamma}{2\omega_{\text{sec}}} \tag{2.23}$$

For singly-ionised ytterbium, $\Gamma \sim 2\pi \times 20$ MHz, and typical secular frequencies in rf traps are $\omega_{\text{sec}} \sim 2\pi \times 1$ MHz. Doppler cooling can therefore reach a minimum mean vibrational quantum number $\bar{n} \sim 10$, which does not satisfy any sensible criterion for ground-state cooling. It is possible to achieve near ground-state Doppler cooling in a sufficiently strong trap, however [43].

2.4.3 Direct sideband cooling

Conceptually the simplest form of ground-state cooling, sideband cooling is the basis for all other forms of ground-state cooling described here. As explained in Section 2.1.2, sidebands appear in the spectrum of a narrow transition when a particle is confined. The sidebands on the optical spectrum correspond to transitions between the vibrational levels of the trap. Specifically, if a two-level atom has a ground state $|g\rangle$ and excited state $|e\rangle$, with vibrational states labelled by n, then exciting on a sideband corresponds to the transition $|g, n\rangle \rightarrow |e, n'\rangle$. Hence, if a laser is tuned to the first lower sideband, it will excite the $|g, n\rangle \rightarrow |e, n-1\rangle$ transition. When the atom spontaneously decays, it has the highest probability of decaying on the carrier (i.e. with no change in n), and so on average one quantum ($\hbar\omega_{\text{sec}}$) of vibrational energy is taken from the ion's motion in each cycle. The cooling process is illustrated in Figure 2.9.

To sideband cool, it is necessary to be able to resolve the sidebands. This imposes the *strong binding condition* on the atomic transition to be used for the cooling, which is

$$\Gamma \ll \omega_{\text{sec}}, \tag{2.24}$$

where Γ is the linewidth of the transition and ω_{sec} the secular frequency of the trap. For this reason, sideband cooling cannot use the electric dipole

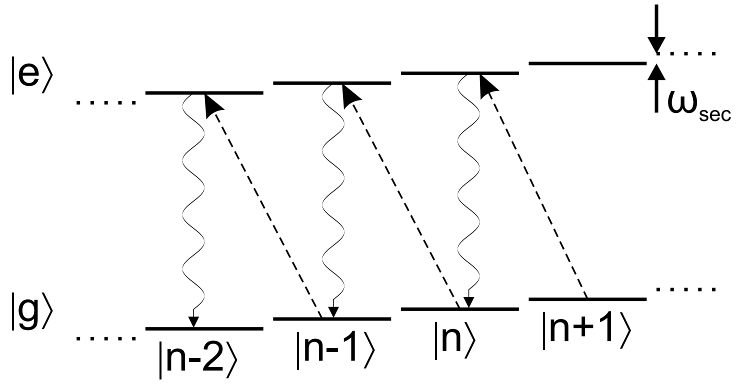

Figure 2.9: Spontaneous decay is an essential part of sideband cooling, as it is all cooling processes.

transitions typically used for Doppler cooling. Usually, an electric quadrupole transition is used, with a lifetime of the order of 1 s. Most traps have secular frequencies of a few MHz, and so the strong binding condition is satisfied and the sidebands are resolved.

Because the transitions used for sideband cooling have narrow linewidths in order to satisfy the strong binding condition, the rate of spontaneous decay out of the excited state is small. Hence, using a simple two-level system as described does not lead to efficient cooling in most cases.

Most ions used for trapping and cooling experiments are alkali-like, and hence have a relatively simple term structure. The sideband cooling transition is often chosen to be an $S \rightarrow D$ quadrupole transition, the Doppler cooling being performed on the dipole-allowed $S \rightarrow P$. In general, the ion has hyperfine structure and may also be in a magnetic field, so the transitions are between specific Zeeman sublevels or hyperfine levels. By coupling the D state to another sublevel of the P state with a 'quenching' laser, the effective decay rate of the D state can be increased because the P state has a rapid decay to the ground state. This offers the opportunity to 'design' an effective two-level system with arbitrary linewidth for optimal cooling from a three-level system [44]. There is a trade-off to be made in increasing the effective decay rate Γ between the metastable level and the ground state, however. As the effective linewidth is increased, the rates of off-resonant driving of the carrier and blue sideband increase. The lowest value of $\bar{n}$ (the average vibrational quantum number) which can be reached is given by [45]

$$\bar{n} = \left(\frac{\Gamma}{2\omega_{\text{sec}}}\right)^2 . \tag{2.25}$$

The intensity of the quenching laser is chosen as a compromise between the cooling rate and the final value of $\bar{n}$. It is also possible to quickly ramp the intensity of the quenching laser to optimise both speed and the final $\bar{n}$ [46].

Sideband cooling has been implemented experimentally several times, with good results [47, 48, 49, 50].

More detailed information on the theory of sideband cooling can be found in Chapter 6, which details the theory behind the numerical modelling of various cooling processes.

2.4.4 Raman sideband cooling

Raman sideband cooling is conceptually similar to direct sideband cooling, in terms of exciting transitions between different vibrational levels of two internal atomic states. The difference is that the sideband is resolved on a Raman transition. Like direct sideband cooling, it depends critically on spontaneous decay (or, more accurately, a spontaneous Raman transition).

Figure 2.10 shows the basic Raman sideband cooling system. It requires three laser beams, though typically both the Raman beams will be derived from the same laser via an acousto-optic modulator (AOM). The cooling process depends on resolving the sidebands, and so the strong binding condition applies. The linewidth of the Raman transition, though, is very small, and in practice this is not a limiting factor on the cooling. The limit is due to off-resonant heating (driving other sidebands), and is given by

$$\bar{n} \approx \left(\frac{\Omega}{\omega_{\text{sec}}}\right)^2 , \tag{2.26}$$

where Ω is the Rabi coupling strength of the Raman beam pair [51].

In a Raman process, at least two of the beams are usually derived from a

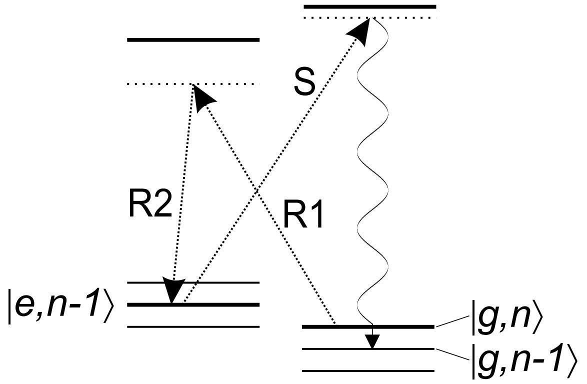

Figure 2.10: Raman cooling scheme. Beams R1 and R2 form a Raman pair, tuned such that they excite $|g, n\rangle$ to $|e, n-1\rangle$ transitions. The excited state $|e\rangle$ is repumped by the beam S, via a spontaneous Raman transition. This resets the cycle, allowing cooling to continue.

single source by a coherent process (such as in an AOM), and therefore have very high mutual coherence.

Raman cooling has been experimentally implemented with good performance [52, 53].

2.4.5 EIT Cooling

Recently, another method of ground state cooling has been proposed [21, 54], which exploits electromagnetically induced transparency (EIT) on a strong transition ($\Gamma \gg \omega_{\mathrm{sec}}$). EIT is a quantum interference effect, whereby the absorption of an atom is cancelled at one frequency by strongly driving it at another. It belongs to a class of multiple path interference effects in atoms which includes coherent population trapping (CPT) and stimulated Raman adiabatic passage (STIRAP) [23].

The EIT scheme is shown in Figure 2.11. A strong coupling laser links the excited state $|e\rangle$ and metastable state $|r\rangle$. A much weaker probe beam (typically, $\Omega_g \approx 0.1\Omega_r$) links $|g\rangle$ and $|e\rangle$. Both lasers are *blue* detuned from resonance with $|e\rangle$. This generates an absorption profile similar to that shown in Figure 2.12 [55], with a dip in absorption at $\Delta_g = \Delta_r$ (both lasers equally detuned from resonance). This is the induced transparency. In EIT cooling,

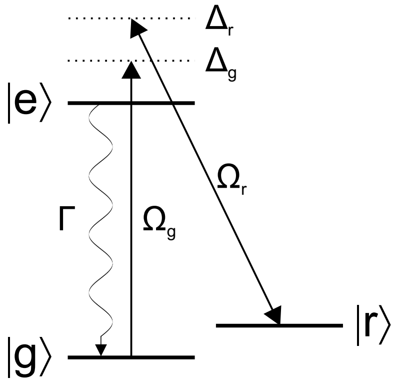

Figure 2.11: EIT cooling scheme. The coupling laser and probe laser have Rabi coupling strengths Ω_r and Ω_g, respectively. The upper state has a linewidth Γ. EIT is induced when the detunings of the lasers Δ_g and Δ_r above resonance with $|e\rangle$ are equal.

both lasers are tuned to this point. The detuning is of the order of a few natural linewidths of $|e\rangle$.

In addition to the transparency, there is a strong, narrow ($\Gamma < \omega_{\text{sec}}$) absorption feature close to the frequency of the coupling laser. This absorption is narrow enough that its red sideband can be resolved (it is harder to resolve the blue sideband due to the asymmetry of the profile). The separation of the transparency and the narrow resonance is given by the AC Stark shift (light shift) Δ_{AC} of $|e\rangle$ created by the coupling laser.

$$\Delta_{AC} = \frac{\sqrt{\Delta_r^2 + \Omega_r^2} - |\Delta_r|}{2}. \tag{2.27}$$

By tuning the intensity (and hence Rabi frequency) of the coupling laser, the position of this narrow resonance can be altered. By setting it such that $\Delta_{AC} = \omega_{\text{sec}}$, the cooling laser is brought into resonance with the red sideband of the narrow transition, while absorption on the carrier is (close to) zero. This eliminates one of the major sources of heating in ion-trap laser cooling experiments, whereby an atom is excited off-resonantly on the carrier and then decays on the blue sideband, increasing the ion's motional state.

Once excited on the red sideband, the atom decays with natural linewidth

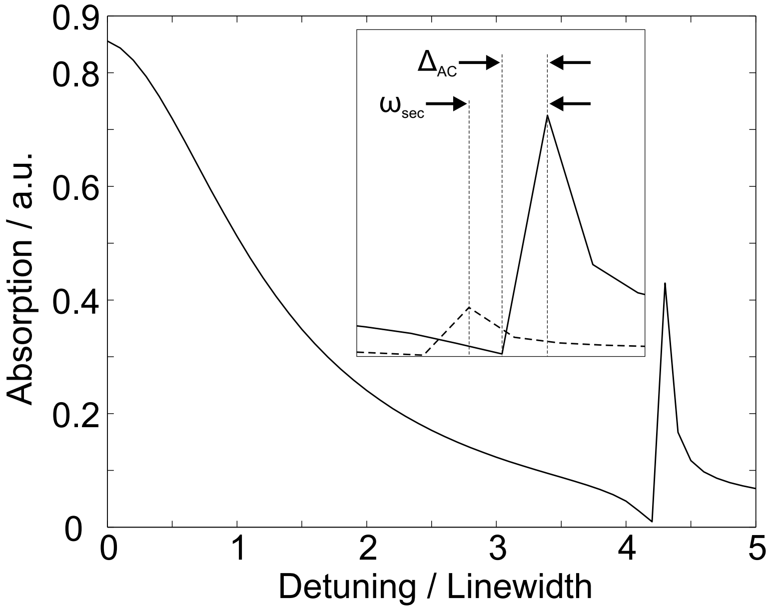

Figure 2.12: Main figure: EIT absorption profile for a free atom. This was generated by the model described in Chapter 6. Inset: The same profile, showing the first-order red sideband (dashed line) which appears when the atom is confined in a harmonic trap of frequency ω_{sec}. The absorption peak is split from the transparency by the AC Stark shift Δ_{AC}.

Γ. Hence the decay is fast and the cooling rate can be high. For a common detuning from resonance Δ, the cooling limit is given by [54]:

$$\bar{n} = \frac{\Gamma}{4\Delta^2} \tag{2.28}$$

The theory of EIT cooling is discussed in further detail in Chapter 6, which also covers numerical modelling of the EIT cooling process. EIT cooling has been experimentally demonstrated just once [22].

2.4.6 Temperature measurement

The thermal motion of a trapped ion produces sidebands in the spectrum of a narrow optical transition, as discussed in Section 2.1.3. These sidebands encode information about the extent of the thermal motion of the ion, or its temperature. Specifically, the heights of the sidebands in the spectrum of a trapped ion with ground-state Lamb-Dicke parameter η and average vibrational quantum number $\bar{n}$ are, to first order in η^2:

$$\text{First order red sideband: } \eta^2\bar{n} \tag{2.29}$$

$$\text{Carrier: } 1 - \eta^2(2\bar{n}+1) \tag{2.30}$$

$$\text{First order blue sideband: } \eta^2(\bar{n}+1) \tag{2.31}$$

This is illustrated in Figure 2.13. The height of an individual sideband in relation to the carrier allows determination of $\bar{n}$ provided that η is already known. The asymmetry of the red and blue sidebands provides a method of measuring the temperature without knowledge of η. The ratio of the height of the red sideband to the height of the blue sideband is $r = \bar{n}/(\bar{n}+1)$, and therefore

$$\bar{n} = \frac{r}{1-r} \tag{2.32}$$

Quantum jump spectroscopy of the red and blue sidebands of a narrow transition is therefore a good method of determining the temperature of a trapped ion, independent of knowledge of any trap parameters. It is particularly sensitive for low $\bar{n}$, where the red sideband vanishes. For higher excitation

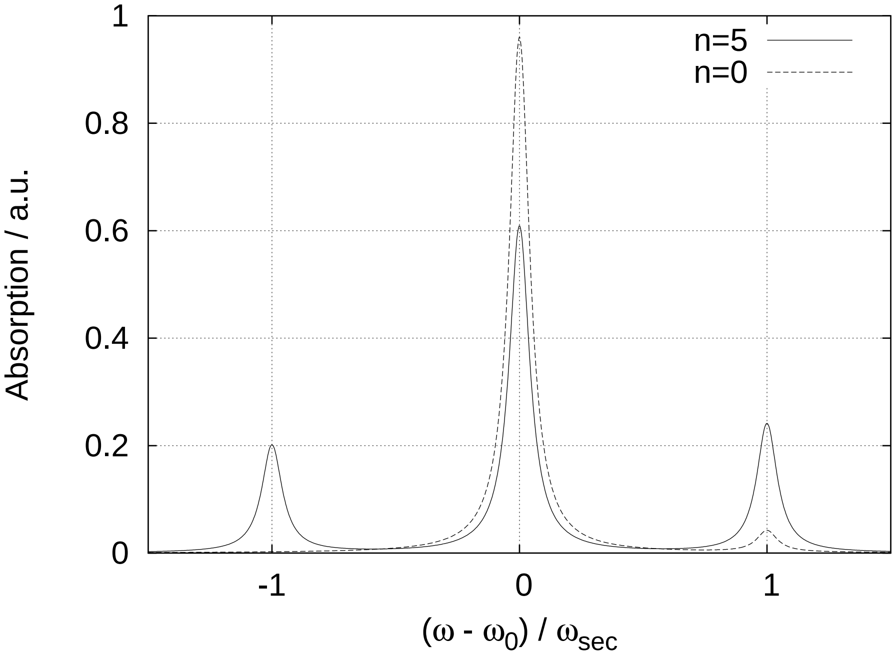

Figure 2.13: Asymmetry of red and blue vibrational sidebands for $\eta = 0.2$, with $n = 5$ and $n = 0$.

numbers, the ratio of the sidebands to the carrier can be a more accurate measure of $\bar{n}$, as the fractional asymmetry of the sidebands shrinks and is masked by experimental noise or counting statistics.

The coherent interaction strength of the sidebands is a function of the vibrational state n (note we are no longer discussing the average $\bar{n}$). The Rabi frequency of an ion, with vibrational quantum number n and ground-state Lamb-Dicke parameter η, driven on its motional sidebands is, for a light intensity which gives Rabi frequency Ω_0 in a free atom:

$$\text{First order red sideband: } \Omega_0 \eta \sqrt{n} \tag{2.33}$$

$$\text{Carrier: } \Omega_0 \left[1 - \eta^2 \left(n + \frac{1}{2}\right)\right] \tag{2.34}$$

$$\text{First order blue sideband: } \Omega_0 \eta \sqrt{n+1} \tag{2.35}$$

Typically, an ion will be in a thermal distribution over all n. Each part of the probability distribution will contribute its own frequency of oscillation when the transition is coherently driven. This leads to dephasing, and the visibility of the oscillation will decrease as a function of time [49]. This can be used to measure the temperature of the ion. Given a sufficiently high-quality data set, the entire population distribution of the ion over vibrational states could in principle be established. It is usually easier, however, to assume an appropriate distribution function (typically a thermal distribution) and then fit that to the observed data.

The coherent oscillation method of temperature determination can be more accurate than the simple spectroscopic method in practice, largely because the data rate is significantly higher. Coherent oscillations can be driven very quickly, and this is necessary to avoid laser-induced decoherence. Direct spectroscopy, meanwhile, should generally be performed below saturation.

2.4.7 EIT temperature measurement

For the Yb^+ system in particular, it is interesting to consider alternative techniques of measuring the temperature of a trapped ion to those discussed in Section 2.4.6. The octupole transition is weak, and cannot be driven at a high rate. Coherent excitation of the transition is not an option, currently. The detection efficiency of quantum jump spectroscopy is very high, however, and essentially background-free. This makes spectroscopy of the vibrational sidebands of the octupole transition an accurate and reliable method of temperature measurement, especially at low temperatures. Despite these benefits, it is not very quick, adds to the complexity of experimental control sequences, and requires another laser system to be operational during experiments. Therefore, a method of inferring the temperature from EIT spectra could be useful.

It should be possible to infer the temperature of a trapped ion by observing the fluorescence from the ion when illuminated with an EIT beam pair. The simplest experiment is one in which after a few milliseconds of EIT cooling, by which time a steady state should have been reached, the fluorescence is simply integrated over a suitable time period, without altering the frequencies of the lasers from the EIT cooling configuration. The fluorescence is then a measure of the magnitude of the red sideband of the transition. By comparing this with the fluorescence from the peak of the EIT spectrum, and with knowledge of the Lamb-Dicke parameter, the temperature of the ion (in one dimension) can be found.

As the act of probing the spectrum alters the height of the sidebands, this is not applicable as a general measurement technique. It can be extended to more general applicability by time-gating the EIT fluorescence collection. This allows the generation (by repeating the experimental sequence many times) of a curve of the fluorescence rate as a function of time. This can be extrapolated to zero time, and hence the original height inferred. The technique can therefore be used to measure the temperature of an ion after any arbitrary experimental procedure.

The limiting factor in this method is the fluorescence background. There is always a certain level of background scatter, due to the trap electrodes, vacuum windows and so on. There may also be off-resonant scatter from energy levels not involved in the three-level EIT system. Because of this, photon-counting statistics set an upper limit on the temperature which can be measured for a given integration time. A similar method was used in experiments with trapped indium ions [48], where direct single-photon scattering from a moderately narrow transition (a few hundred kilohertz) was used both for sideband cooling and to measure the final temperature.

Chapter 3

Experimental Details

This chapter covers the basic layout and equipment used for the majority of the experiments described in this thesis. After covering the overall scheme, the details of the trap, basic experimental procedures, lasers, the control system and the femtosecond comb are given. Any deviations from this layout for a particular experiment are mentioned in the relevant chapter. In particular, the cooling laser system was altered for the EIT experiments, and the details of this are given in Chapter 7.

The arrangement of the experimental apparatus is illustrated in Figure 3.1. The 467 nm and 369 nm beams are combined in a polarising beamsplitter, before being focused into the trap by an achromatic lens system. The ion fluorescence is collected by a large-aperture lens, and focused through a pinhole onto a photomultiplier tube (PMT). The 635 nm and 935 nm beams are both fibre-coupled to the trap, before being combined in a dichroic beam splitter and focused onto the ion. The wavelengths of all the beams are measured on a wavemeter. The 12.6 GHz microwaves required to maintain the Doppler cooling cycle in ^{171}Yb^{+} are fed into the trap via an external rectangular waveguide.

The probe and cooling lasers were measured using a NIST air-track wavemeter, with a central carriage carrying two retroreflectors. This is used to obtain fringe patterns from both the light being measured and a reference helium-neon laser, the ratio of the fringe count rates being measured

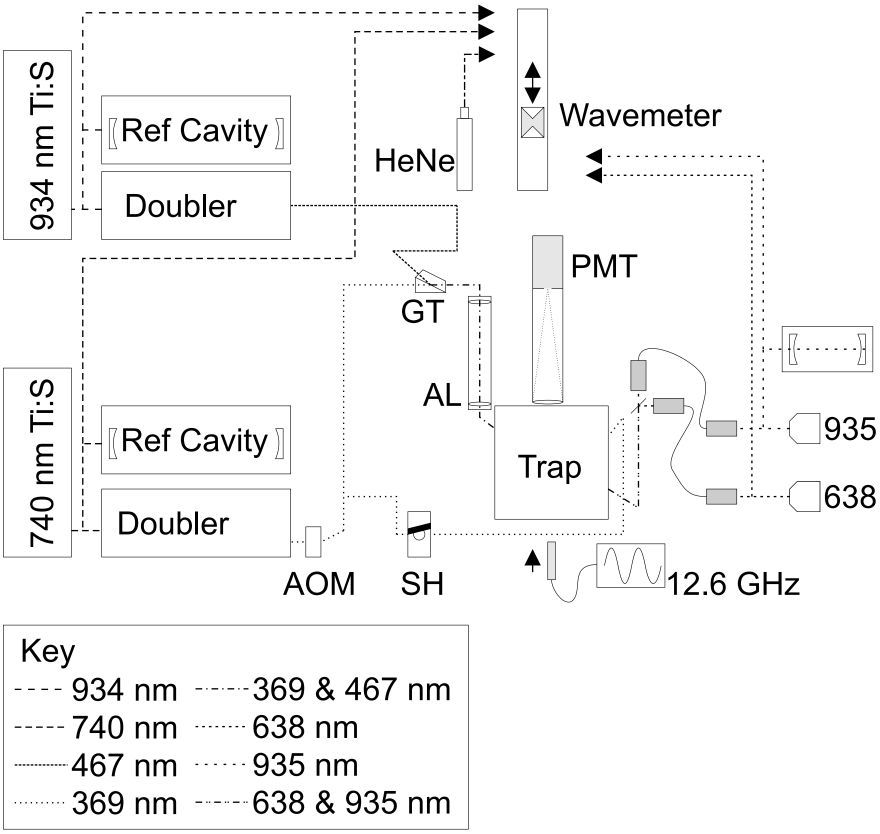

Figure 3.1: Schematic layout of the ytterbium ion trap experiment. The probe laser is shown in more detail in Figure 3.6. AOM = Acousto optic modulator, SH = Mechanical shutter, PMT = photomultiplier tube, GT = Glan-Thompson polarising beam splitter, AL = Achromatic lens, HeNe = Helium-Neon reference laser.

with a 1 s gate time on a HP5131A frequency counter. This is accurate to approximately 3 parts in 10^8. Initially, the 638 nm and 935 nm lasers were measured on a 6-digit Burleigh wavemeter, but were later transferred to the NIST wavemeter.

3.1 Endcap ion trap

There are several designs of radiofrequency trap in common usage. The original design (Paul trap) uses electrodes with hyperbolic surfaces - two ‘endcaps’ and a ring electrode in between them - to produce a very pure quadrupole potential. A D.C. voltage is applied to the endcaps, and the oscillating voltage to the ring. This design has the problem of poor optical access to the ion, though it can hold very large numbers of ions due to the purity of the quadrupole field.

The design used in the ytterbium ion traps at NPL is a variant known as the ‘endcap trap’. The geometry is illustrated in Figure 3.2. This design was originally described in [56], and has the benefit of good optical access to the ion. The characteristic dimension of such a trap is z_0, illustrated in Figure 3.2. For the ion trap used in this work, $z_0 = 0.56$ mm. The trap is constructed from tantalum, chosen partly because of its low contact potential with ytterbium. In addition to the electrodes shown in the simple schematic in Figure 3.2, there are two orthogonal compensation electrodes in the radial plane of the trap for micromotion compensation, shown in the photograph in Figure 3.3. It is also possible to adjust the potentials of the upper and lower outer endcaps independently both for axial micromotion compensation and to alter the axial and radial secular frequencies (Section 2.1). The trap is driven by a transistor oscillator, with $V_{RF} = 300$ V, and $\omega_{RF} = 12.8$ MHz. The efficiency $\mathcal{E}$ of the trap has been measured to be 0.63 ± 0.3, and with the outer endcaps grounded, the secular frequencies for a single ytterbium ion are $\omega_z = 2\omega_r = 1.84$ MHz. The trap has ovens for both ^{171}Yb and ^{172}Yb. These each use a piece of isotopically enriched ($\sim 80\%$ purity) ytterbium foil as the atom source.

The trap is loaded by electron impact ionisation. A current of 7.8 A is passed through the oven for 60 s, and during the final 15 s of this, a current of 4 A is passed through a coiled tungsten filament near the trap, which is not biased. This emits a thermionic electron current, and these electrons ionise the neutral ytterbium atoms passing nearby. Generally, this procedure loads several ions into the trap. By switching off the laser cooling light for a few minutes, it is usually possible to obtain a single ion.

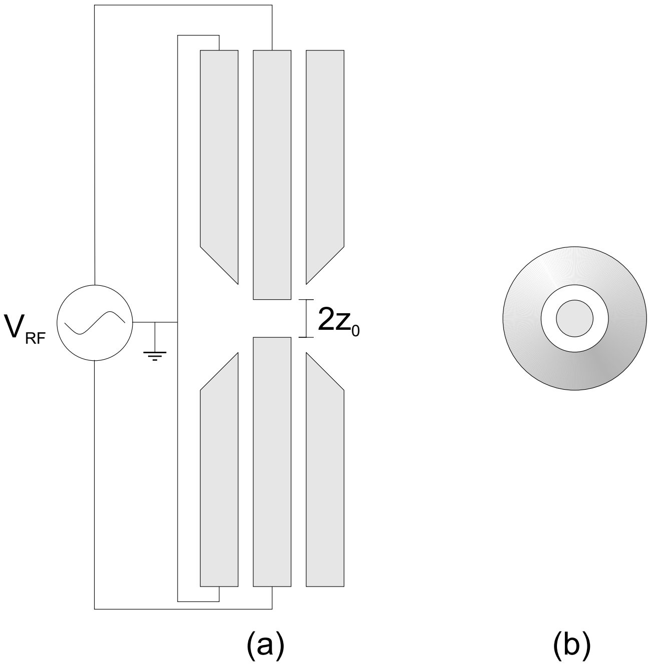

Figure 3.2: Schematic diagram of an endcap trap. (a) Viewing from the radial direction. (b) Viewing along the axis. A radiofrequency voltage V_{RF} is applied to the central pin, while the outer endcaps are, in the simplest case, grounded.

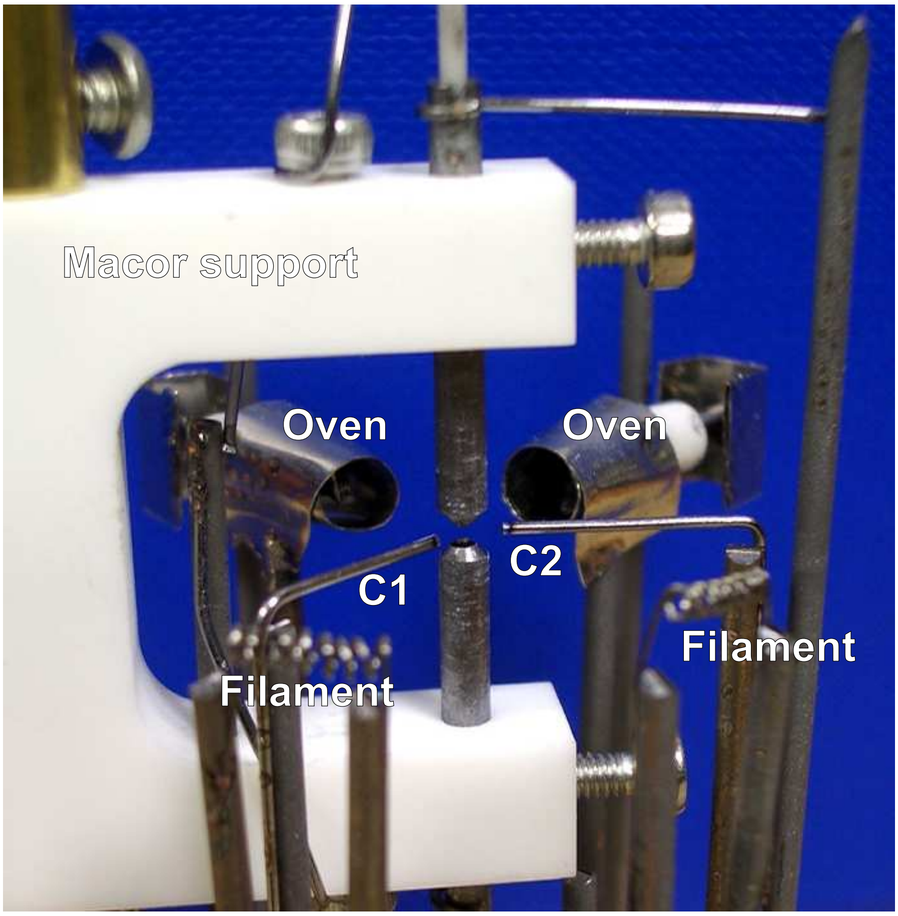

Figure 3.3: Photograph of a model of the ytterbium endcap ion trap. The filaments (out-of-focus) are used as electron sources for ionising the atoms released from the ovens. C1 and C2 are micromotion compensation electrodes for the radial plane. The inner and outer endcaps of the upper and lower trapping electrodes can be seen in the centre of the picture.

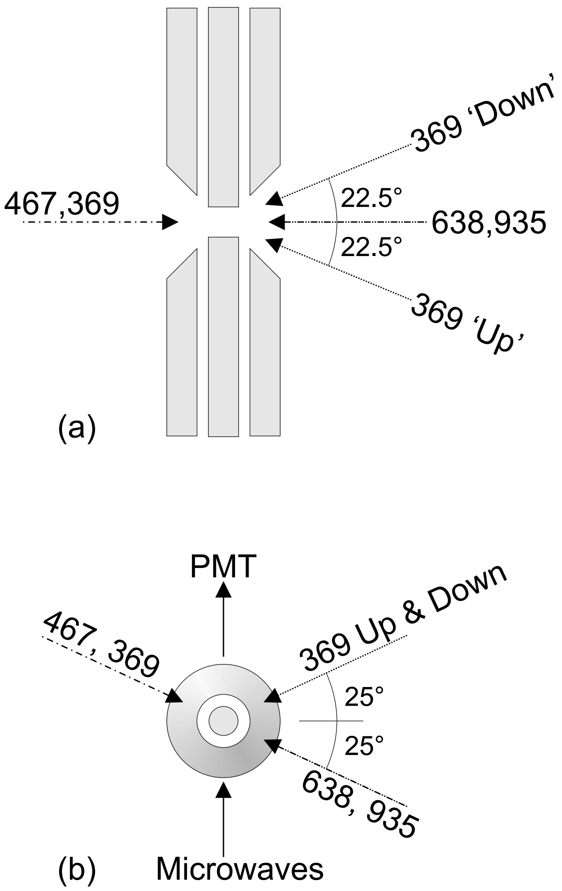

Figure 3.4: Beam paths. (a) Side view, cross-section, looking towards the PMT. (b) Top view, looking down the z axis of the trap. Beams are labelled by their wavelength in nanometres.

3.2 Beam paths

The beam paths used in the experiment are shown in Figure 3.4. The three Doppler cooling beams at 369 nm allow for three-dimensional micromotion compensation (Section 3.4). The 'horizontal' 369 nm beam is used as a tracer for aligning the 467 nm beam on the ion, and as the coupling beam for EIT experiments. The 'up' beam is used for Doppler cooling, and the 'down' beam as the probe beam in EIT experiments. The microwaves are sent into the trap from an external rectangular waveguide pointing towards the PMT.

The transition rate on the octupole transition is very low, and it does not provide a large enough signal to align the 467 nm beam on the ion (certainly not until it is properly aligned and the transition has been located in frequency relative to the reference cavity). For this reason, the 369 nm beam is used as a ‘tracer’ - it is aligned on the ion by monitoring the ion’s fluorescence signal, and then the 467 nm beam is aligned with it. Firstly, the 369 nm and 467 nm beams in the horizontal beam path are combined in a calcite Glan-Thompson polarising beamsplitter. Using a temporary mirror after the beamsplitter, they are then made collinear over a ~ 5 m beam path, by multiple reflections along the optical table. This done, the beams pass into an achromatic doublet lens (for 369 nm and 467 nm), designed to focus the beams to a 5 μm spot at the trap centre. A temporary mirror is inserted after this lens, and the foci of the two beams are checked for coincidence using a micrometer-mounted knife edge. The alignment is verified by checking the transmission of both the beams through a 5 μm pinhole. By following this procedure, aligning the 369 nm beam on the ion using a final mirror before the ion trap and adjusting the focus using a micrometer mounted beneath the achromatic lens assembly, it is possible to have confidence in the alignment of the 467 nm beam on the ion.

3.3 Fluorescence detection

For all the complexity and number of inputs to the experiment, there is just a single output - the ion’s fluorescence signal at 369 nm. There are several different methods of counting this fluorescence, each with their own uses, which are described below.

For all of these methods, what is actually being counted are discrete pulses from the photomultiplier tube (PMT) attached to the trap. This signal has three sources:

Fluorescence Photons emitted by the ion in the direction of the collection optics. A filter is in place such that only photons at around 369 nm are counted. The longer wavelengths involved in the experiment, at which the ion can also scatter significant numbers of photons (935 nm) are blocked. The peak fluorescence rates observed in the experiment are $\sim 35,000$ counts per second for $^{172}\mathrm{Yb}^+$, and $\sim 6,000$ counts per second for $^{171}\mathrm{Yb}^+$.

Environmental light Photons in the passband of the filter which do not come from the ion. The biggest contribution to this is scattered laser light at 369 nm. Light reflects from the trap electrodes, which are close to the ion and therefore close to the 'field of view' of the collection system. This is mostly due to the spatial extent of the beam (the tightly-focused beam has practically no background of this sort) as the wings clip the trap electrodes, and partly due to scatter from the vacuum windows onto the trap. Most of the 369 nm windows are at Brewster's angle to minimise this, or are anti-reflection coated. This contribution to the signal is proportional to the laser power, but is typically 100 counts per second for the 'up' Doppler cooling beam, and unmeasurable for the tightly focused beam. Any room light which enters the trap and scatters in the direction of the PMT can also increase the signal, but this is a very small contribution.

PMT noise The PMT also has a 'dark signal' even when no light falls on it. In our experiment this is very low, below 1 count per second.

The output of the PMT can therefore be broken down into two components - signal (fluorescence) and background (everything else). In normal operation, the signal is much larger than the background, but there are obvious problems when trying to detect small signals.

3.3.1 Scattering rate measurement

A basic rate measurement is the simplest form of photon detection. The method involves recording the number of events in a set time period, and extracting an average scattering rate from this. All timing information about the individual photons is lost. This is the method used for obtaining spectra of the cooling transition, and for quantum-jump detection (Section 2.3).

3.3.2 Correlated measurement

In a correlated measurement, accurate timing information about the detection of each photon is obtained. This can be used to obtain correlation data either within the data set itself, or against some other signal.

RF-photon correlation

In this technique, the arrival of a photon is measured relative to the phase of a radiofrequency signal, often the ion trap drive frequency (12.8 MHz). There are several ways of implementing this, either with counter-timer electronics, or a time-to-amplitude convertor (TAC) and multichannel analyser hardware. In these experiments, it was almost always the latter, due to the higher time resolution available. The TAC is started at a particular phase of the trap drive cycle, and stopped on the detection of a photon. In this way, any correlation of the fluorescence rate with the trap drive becomes clear when the fluorescence rate is plotted against arrival time relative to the rf signal. Micromotion (Section 2.1.4) gives rise to exactly such a correlation, as the fluorescence rate is modified by the Doppler shift due to the ion's motion towards or away from the cooling laser. The rf-photon correlation technique is a very sensitive method of detecting micromotion [30].

Photon-photon correlation

In this technique, it is the intervals between successive photons which are recorded. A histogram of the intervals can be constructed, which can reveal correlations at any frequency, as opposed to the single frequency sensitivity of the rf-photon correlation technique. In particular, the Fourier transform of the histogram will reveal any correlations within the limits set by the time resolution, total time span, and the noise level, which is set by the number of signal and background photons recorded. The technique can be useful for detecting the secular vibrational frequencies of an ion. For this, it is necessary for the ion to be relatively hot, so the signal is sufficiently strong to be detected. By inducing large micromotion, it is possible to tune the laser to a frequency where there is 'local heating', but above a certain temperature the ion is cooled again. Monitoring the fluorescence of the ion in this regime will reveal peaks in the Fourier transform of the fluorescence histogram at the ion's secular frequencies (subject to the beam having an appropriate direction), and a large peak at the trap drive frequency.

Triggered photon detection

Triggered photon detection involves measuring the arrival times of photons relative to some external trigger signal, for example a digital I/O signal from the experimental control system. The system then measures the interval between the trigger and the first photon, and the intervals between successive photons. It is then possible to reconstruct the entire time sequence with reference to the trigger. This enables measurement of the fluorescence evolution over short timescales synchronised with the switching of beams, or other events, in the experiment.

3.4 Micromotion minimisation

In order to minimise micromotion, there are two orthogonal 'compensation electrodes' in the radial plane of the trap, and it is also possible to apply potentials to the upper and lower outer endcaps of the trap. Each electrode (or pair in the radial direction) corresponds to a particular cooling beam in the minimisation procedure. The rf-photon correlation technique (Section 3.3.2) is used to detect the amplitude of micromotion along one beam direction, and then the relevant compensation electrode is adjusted until there is no detectable micromotion in that direction, or until it is below some reasonable threshold. The next direction is then adjusted, and the procedure is iterated either in two dimensions, using the 'Up' and 'Horizontal' Doppler cooling beams, or in 3D using the 'Down' beam as a probe to adjust the outer endcap voltages. In this way it is possible to minimise micromotion in all three directions simultaneously. Typical compensation voltages, some hours after loading the trap, are 5–20 V for the radial compensation electrodes and 0–0.5 V on one of the endcaps. For most measurements, the endcaps were left at zero volts, and minimisation done in two dimensions only (the 'horizontal' 467 nm probe beam direction and the 'up' Doppler cooling direction).

3.5 Setting the fluorescence threshold

Fluorescence from the Doppler cooling transition at 369 nm is monitored to observe quantum jumps. The simple scattering rate method (Section 3.3.1) is used to measure the number of photons collected in a specific time interval. The control system must then decide whether or not a quantum jump has occurred. This is achieved by comparing the measured rate to a pre-computed threshold value. To do this, both the fluorescence rate, and the background count rate must be known, so they are measured beforehand. If the expected number of counts in a given period when the ion is fluorescing is N_F, then the standard deviation on this count is $\sqrt{N_F}$, assuming the fluorescence obeys Poissonian statistics. Similarly, the background has an expected count N_B and standard deviation $\sqrt{N_B}$. The threshold T should be set an equal number of standard deviations S (the safety factor) from each of these expectation values for the minimum probability of mis-identifying a quantum jump. That is

$$T = N_B + S\sqrt{N_B} = N_F - S\sqrt{N_F} \tag{3.1}$$

The safety factor is given by

$$S = \frac{N_F - N_B}{\sqrt{N_F} + \sqrt{N_B}} \tag{3.2}$$

For a low probability of mis-identifying a quantum jump, S should be no less than 3. The detection time should be set so that enough photons are collected to make this true.

The above is based on the assumption of Poissonian statistics, and a constant underlying count rate. While this holds true for $^{172}\mathrm{Yb}^+$, it is not true for $^{171}\mathrm{Yb}^+$. During the cooling cycle, it is possible, via an off-resonant excitation, for the ion to decay to the $^2\mathrm{D}_{3/2}(F = 2)$ level (Section 1.4.3).

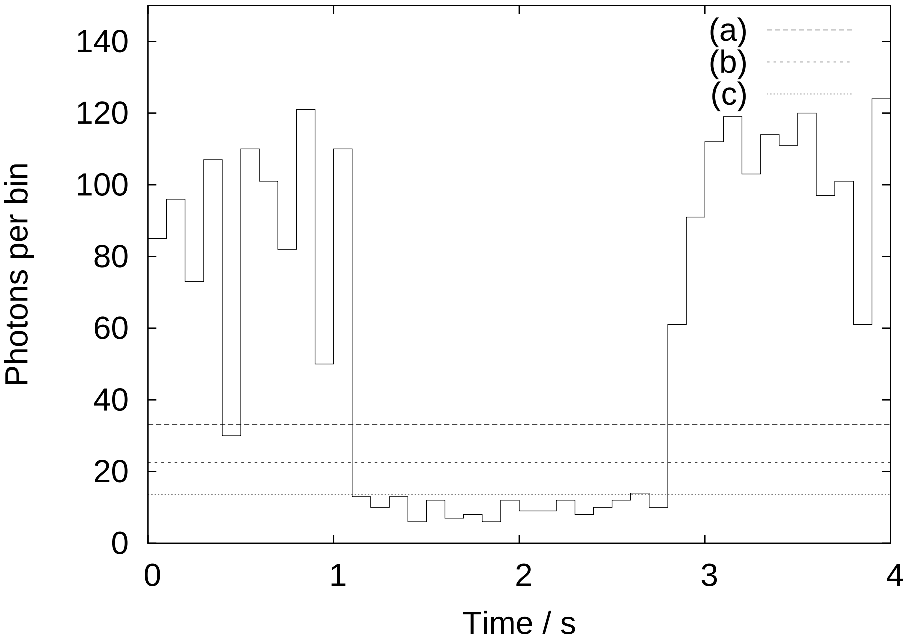

Figure 3.5: Simulated fluorescence data set for ^{171}Yb$^+$. A quantum jump to the ^{2}F$_{7/2}$ state occurs between about 1 s and 3 s. The parameters were a 1000 s^{-1} count rate from the 369 nm transition, on a 100 s^{-1} background scatter count rate, a ^{2}D$_{3/2}(F = 2)$ state lifetime of 50 ms, and a 5 s^{-1} decay rate into the ^{2}D$_{3/2}(F = 2)$-state. The simulated data was integrated over 100 ms bins. Line (a) shows the standard threshold calculation, (b) is 4σ above the background level, while (c) is set close to the background level. Both (a) and (c) would incorrectly register two quantum jumps for this data set.

Once there, it is not resonant with a repumper laser, and fluorescence ceases for the lifetime of the state (52 ms). The ion has undergone a quantum jump, but not to the state of interest. These jumps can be distinguished from those to the F state by their lifetime.

Several issues affect the choice of fluorescence threshold. Figure 3.5 shows a simulated fluorescence data set, with several choices of threshold. The errors which can occur are as follows. When the threshold is set:

Below the background level No jumps are detected.

- **Close to the background level** Shown as line (c) in Figure 3.5. Two different errors occur in this situation. Firstly, a jump may be missed, as the statistical fluctuations of the background count prevent detection. In this case, it may be detected on the next interrogation and assigned to the wrong probe frequency. Secondly, while waiting for fluorescence to return, the level may briefly rise above the threshold. In this case the jump is counted twice, or more, with the extra jumps possibly being attributed to the wrong probe laser frequency. The longer it takes for fluorescence to return, the higher the probability of this occurring.

- **Close to the minimum 'D-jump' level** Shown as line (a) in Figure 3.5. Assuming that the detection time is significantly longer than the lifetime of the $^2\mathrm{D}_{3/2}(F = 2)$ level, then any detection bin in which such a jump occurs will have a noticeably lower level than the average fluorescence rate. It is then possible to identify D-jumps as F-jumps if the threshold is incorrectly set.

When using the Visual Basic software to control the experiment (Section 3.10) the threshold was set manually. The newer software used two thresholds to minimise the error rate. The first threshold was used in the first detection cycle only, and can be thought of as the jump detection threshold. This was set to a certain number of standard deviations (four) above the background level (line (b) in Figure 3.5), ensuring that jumps are hardly ever missed without creating a D-jump background. The second is the return threshold. This is set much higher (using the standard algorithm described above — line (a) in Figure 3.5), and detects the return of the ion's fluorescence. This minimises the probability of double-counting any quantum jumps.

3.6 467 nm probe laser

The optical layout of the 467 nm probe laser system is shown in Figure 3.6. The Ti:Sapphire laser is a monolithic block resonator (MBR) design [57, 58], constructed from aluminium. It is pumped by 12 to 13 W of light from an argon-ion laser, and typically produces 800 mW of 934 nm light. This is

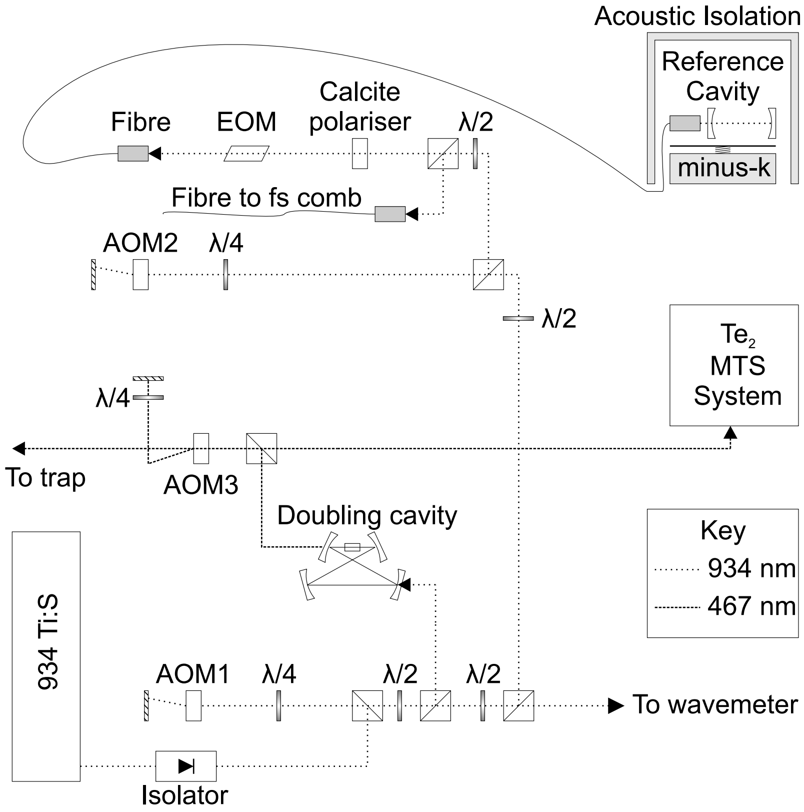

Figure 3.6: The 467 nm probe laser system. Beam-steering mirrors are not drawn.

frequency-doubled in a resonant cavity with a crystal of potassium niobate, with an output of 40 mW at 467 nm.

The blue light can optionally be sent, via a double-passed AOM, to a molecular tellurium modulation transfer spectroscopy (MTS) system [34]. This can be used as a local frequency reference, to aid in locating the octupole transition. The reference cavity is unsuitable for this purpose due to long-term isothermal drifts, while the wavemeter does not have sufficient accuracy. The MTS system can locate the octupole transition to within approximately 1 MHz at 467 nm. With a narrower probe laser, it becomes impractical to search this entire range, and so for later measurements, the femtosecond comb was used to make a preliminary determination of the absolute frequency of the reference cavity, and the tellurium system was not used.

3.6.1 Frequency stabilisation

The commercial Ti:sapphire laser is stabilised by its own electronics to an internal tunable low-finesse reference cavity [58], and has a frequency width of around 200 kHz at 934 nm.

The laser linewidth is narrowed further by a Pound-Drever-Hall [59, 60] lock to a high-finesse ($\mathcal{F} = 250000$) reference cavity, an avalanche photodiode being used for detection of the light reflected from the cavity. The error signal is fed back to AOM1, via a 'bypass topology' loop filter [59] and the frequency-modulation input of the voltage-controlled oscillator which drives AOM1. The output of the loop filter is compared to a reference voltage, and integrated, and this signal is fed back to the Ti:sapphire laser's internal reference cavity piezo, so that AOM1 does not go out of range. The cavity is made from an ultra-low-expansivity (ULE) glass spacer, with high-reflectance mirrors (on a ULE substrate) optically contacted to the spacer. The cavity is mounted on three diamond tips (record player styli) inside a temperature controlled cylinder. The cylinder is itself mounted inside an evacuated vacuum chamber, which is also temperature controlled. The dual layer control keeps the temperature stable to approximately 0.1 mK over a daily cycle, resulting in low drift rates of the cavity resonance (< 0.1 Hz/s once the system has settled, which takes a few days).

The light is modulated at 4.7 MHz by a Brewster-cut electro-optic modulator (EOM), driven by a frequency synthesiser, an amplifier, and a resonant circuit attached directly to the EOM electrodes. If the polarisation of the incoming light is not precisely aligned with the optic axis of the EOM crystal, there is slight residual amplitude modulation (RAM), which can cause an offset in the Pound-Drever-Hall error signal. Changes in intensity can then move the locking point, reducing the stability of the laser. To prevent this, a calcite polariser is used prior to the EOM, which is aligned by inspection of the error signal.

For later experiments, the intensity of the 467 nm light was also stabilised. A small fraction of the light was reflected onto a photodiode, and the signal level compared to a stable reference. This error signal then fed back to the amplitude modulation input of AOM1. This removes significant flicker noise at frequencies up to a few kilohertz.

Initially, the light was then input directly, through free-space, into the high-finesse cavity, which sat directly on the optical table. This set-up produced laser linewidths, observed by quantum-jump spectroscopy of the octupole transition in ytterbium, of around 4 kHz, as shown in Figure 3.7(a). This was the status of the laser for the initial femtosecond comb frequency measurements (Chapter 5).

The cavity was then mounted on a “minus-k” passive vibration isolation system, reducing the linewidth to 2 kHz, as shown in Figure 3.7(b). The light was fibre-coupled to the isolated platform, to prevent misalignment due to motion of the platform relative to the table.

Currently the reference cavity, and vibration isolation platform, are enclosed inside an acoustically-isolating box. This is constructed from 15 mm thick plywood, the interior covered with lead-lined acoustically absorbing foam. The underside of the box is not enclosed, and is bolted down to the optical table, with thin rubber seals around the bottom edges. Optical fibre and electronic connections are run underneath these seals, which are deep enough to accommodate them without significant gaps. This enclosure reduced the linewidth of the laser further, to approximately 150 Hz, shown

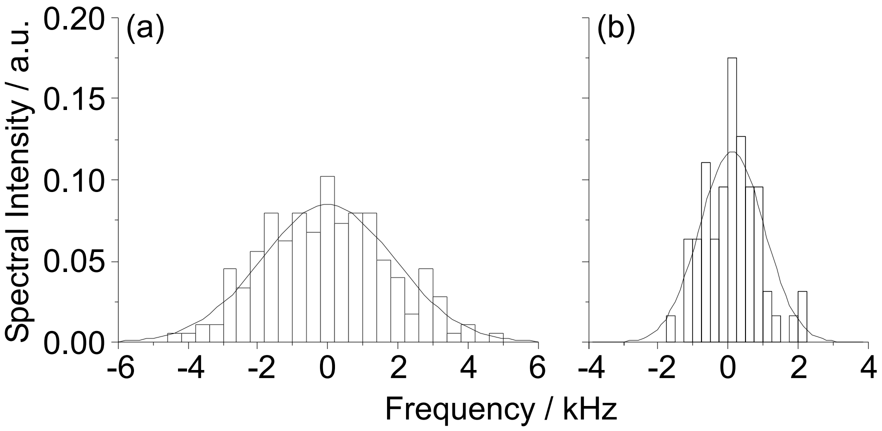

Figure 3.7: Spectra of the octupole transition. (a) Before and (b) after mounting the high-finesse reference cavity on a passive vibration-isolation platform. The smooth curves are Gaussian fits to the data, as a guide to the eye. For both of these situations, the Gaussian model appears to be good.

in Figure 3.8. The laser was in this state for the later femtosecond comb measurements (Chapter 5).

3.7 369 nm laser

The 369 nm laser is a frequency-doubled Ti:sapphire laser, using parts from a Schwarz Electro Optics laser, and is very similar to the 467 nm laser. The Ti:sapphire laser is pumped by 6.5 W of light from an argon-ion laser, and produces typically 400 mW of light at around 740 nm. It is not a monolithic block design, the components being mounted on an optical breadboard. The locking scheme is similar to that for the 467 nm laser, except the electronics are home-made. The laser is side-of-fringe locked to a tunable NPL low-drift etalon. There is no AOM lock in the system, and the linewidth is estimated to be 800 kHz. The laser can be tuned, in increments of approximately 1.5 MHz by a manual ‘click box’ which controls the voltage on the reference cavity piezo. The frequency of the 740 nm light is doubled in a crystal of LBO, and the output is a few milliwatts at 369 nm, which is sufficient to strongly saturate the cooling transition.

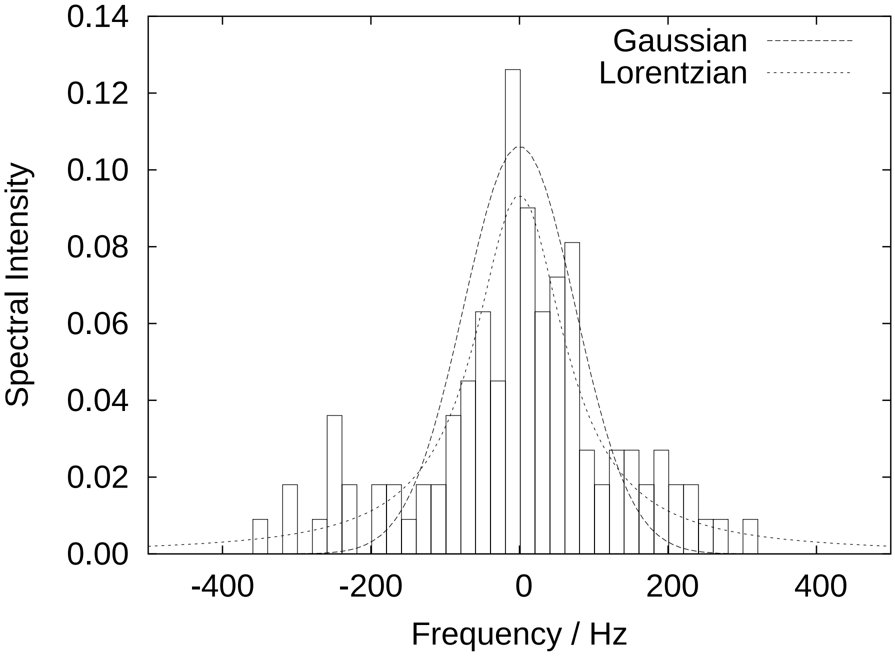

Figure 3.8: Spectrum of the octupole transition with the ULE reference cavity mounted on a vibration-isolation platform, inside an acoustically isolating enclosure. A Gaussian fit to the data gives a 180 Hz FWHM, while a Lorentzian fit gives a 150 Hz FWHM. The Lorentzian appears a better fit by eye, but there is no underlying physics reason (given the extent of our knowledge about the system) to choose either model, so the figures are a rough guide only.

3.8 Repumper lasers

The 638 nm and 935 nm lasers are both commercial extended-cavity diode lasers of the same design [61]. The 935 nm laser is stabilised to a tunable ULE reference cavity by a side-of-fringe lock feeding back to the extended-cavity grating piezo. The 638 nm laser is not externally stabilised, but is tuned to the correct wavelength using a wavemeter. An extra voltage can be switched in to the cavity piezo on this laser for rapid switching between the two wavelengths required for $^{171}\mathrm{Yb}^{+}$(638.610 nm and 638.615 nm).

3.9 12.6 GHz microwave source

The microwaves at 12.6 GHz are required only for cooling of $^{171}\mathrm{Yb}^{+}$, as a repumper between the ground hyperfine states, as discussed in Section 1.4.3. Two different sources have been used for the 12.6 GHz microwaves during the following experiments. Initially, a HP3732B microwave synthesiser was used, but now a 12.5 GHz Atlantic Microwave oscillator coupled with a mixer and RF synthesiser is used. The signal is amplified by a 1 W microwave amplifier, and sent into the trap via a piece of rectangular waveguide, pointing towards a large window in the trap vacuum chamber.

3.10 Control systems

The experiment is computer controlled, and two systems have been used during the period covered by this thesis. Initially, a control routine written using Visual Basic (VB) [62] was used to collect frequency measurement data. Later, a new system was developed using National Instruments' hardware and LabVIEW software. This system gave more flexibility in the experiments which could be performed. For spectroscopy of the octupole transition, several improvements were necessary to achieve the required level of accuracy in the measurement process. The major changes are detailed below.

Time-stamping The VB program recorded the time at the start of a probe scan. The times of the quantum jumps then had to be inferred - usually by the experimenter recording the end time of the scan in the lab book, and interpolating. The new system allows for time-stamping of each individual quantum jump. This allows better matching of the quantum-jump data with the femtosecond comb measurements of the cavity frequency, down to 1 s. Previously, the accuracy was only around one minute, and this led to errors up to 50 Hz in the determination of the frequency at which a quantum jump occurred.

Experimental range The VB program was geared to quantum-jump spectroscopy. The new system uses more flexible hardware, in particular a digital input-output (DIO) board (National Instruments DIO-32-HS) which can output arbitrary pulse patterns across 32 TTL channels, at up to 20 MHz. This allows the system to perform a much wider range of experiments, including the EIT experiments in Chapter 7. This is due both to the number of outputs, allowing the addition of new equipment to the experiment, and the pattern-output flexibility.

Photon-counting The VB program implemented one type of photon counting - recording the number in a particular time interval. The new counter-timer card (National Instruments PCI-6602) allows more flexibility. In particular, it is possible to record the intervals between successive photons (effectively recording the arrival time of each photon), and to precisely synchronise the counting with the DIO pattern output.

The upgrade allowed the addition of new shutters to the experiment (helpful in eliminating AC Stark shifts from repumper lasers), and the large expansion in digital outputs required to satisfy the fast switching requirements of the EIT experiments.

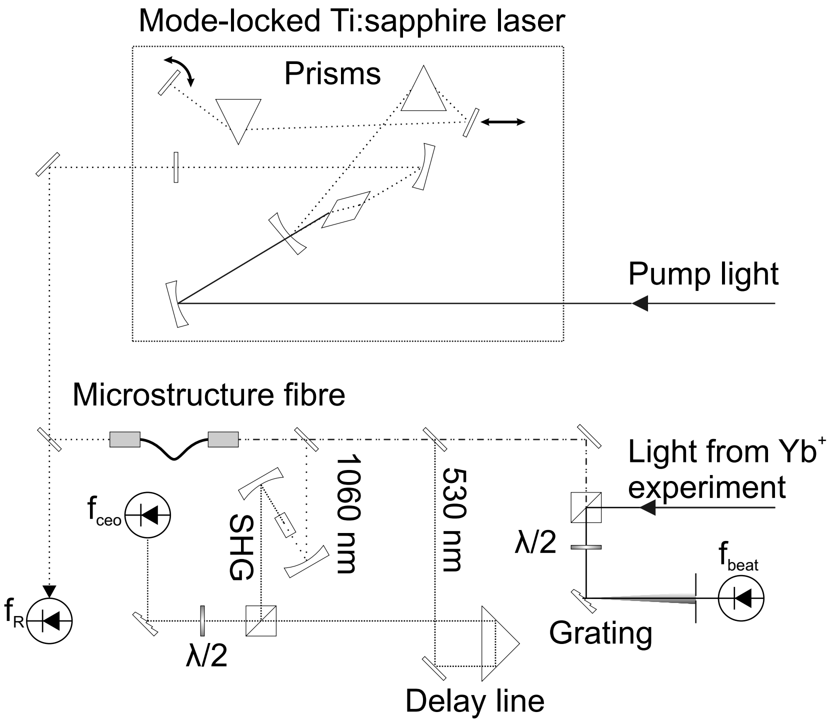

Figure 3.9: The NPL femtosecond comb. The repetition rate f_R of the laser is controlled by translating the cavity mirror before the dispersion-compensation prisms, and the carrier envelope offset frequency f_{ceo} is controlled by the tilt of the mirror after the prisms. The delay line is necessary to ensure that the femtosecond pulses overlap temporally on the detector.

3.11 Femtosecond comb

The NPL femtosecond laser (Figure 3.9) is a mode-locked Ti:sapphire laser constructed from a kit supplied by KMLabs [63], of a linear cavity design with a repetition rate (f_{R}) of 87 MHz, and a pulse length of a few tens of femtoseconds. The laser is pumped by a frequency-doubled Y:VO$_4$ laser (a 10 W Coherent Verdi laser). The microstructure optical fibre used to broaden the comb to cover the optical spectrum has come from two sources in these experiments - Lucent [64] and Bath University. The choice of fibre had no significant impact on the measurement process.

The carrier envelope offset frequency (f_{ceo}) is measured by the self-referencing technique [65, 66]. In this, light is picked off from the beam after the microstructure fibre at around 1060 nm using a dichroic beamsplitter, and is frequency doubled. This is then combined, using a polarising beam splitter and waveplate, with light from the high-frequency end of the comb spectrum, around 530 nm, and a beat is detected on a fast photodiode. The frequencies of these components are

$$f_a = 2(nf_R + f_{\mathrm{ceo}}) \qquad \text{(Doubled light)} \tag{3.3}$$

$$f_b = 2nf_R + f_{\mathrm{ceo}} \qquad \text{(Direct comb light)} \tag{3.4}$$

Therefore

$$(f_a - f_b) = f_{\mathrm{ceo}} \tag{3.5}$$

In this way, f_{ceo} can be directly measured on a photodiode. The beat frequency of the comb with the probe laser can be measured in a similar way. The repetition rate can also be monitored simply by detecting the amplitude modulation of the laser output directly on a photodiode.

The repetition rate, probe laser beat and carrier envelope offset frequency are then all measured with frequency counters referenced to a a signal derived from a stable microwave source, such as a hydrogen maser. For the measurements described in this thesis, the 10 MHz output of one of NPL's hydrogen masers is used. The maser forms part of the clock ensemble used to generate the UTC(NPL) timescale, providing traceability to the SI second at the 2.5×10^{-14} level. This is the limiting factor in comb performance for these measurements.

Because the repetition rate appears in the equation for determining the frequency of the probe laser (Equation 2.19) with a very high multiplier, the achievable accuracy in measuring the probe laser is limited by the accuracy of the repetition rate measurement. To achieve a good signal-to-noise ratio, a tracking oscillator is locked to the beat, and this is counted. To achieve the highest possible accuracy from the frequency counters, the repetition rate is first mixed down to a lower frequency with a signal derived from the 10 MHz output of a hydrogen maser. This frequency is then multiplied up to take full advantage of the available digits of resolution on the HP frequency counters. The beat with the probe laser and the carrier envelope offset frequency both have tracking oscillators locked to them, and the frequencies are counted directly using maser-referenced frequency counters. The repetition rate is also stabilised (with a weak feedback mechanism) using the maser as a reference.

Chapter 4

Systematic frequency shifts

One of the main attractions of single trapped ions as optical frequency standards is that they have extremely low susceptibility to environmental perturbations. These perturbations include low-frequency electric and magnetic fields, and the light fields used in cooling and interrogating the ion. These external fields can all lead to systematic shifts of the transition frequency. These must be characterised and understood in order that the unperturbed transition frequency may be realised.

4.1 AC Stark shift

Currently, the most significant shift for the octupole transition in $^{171}\text{Yb}^{+}$ is the light shift, or AC Stark shift. It is caused by the electric field of the light, and is proportional to $\langle E^2 \rangle$, which is also proportional to the light intensity. Because the transition is very weak, a relatively high light intensity is required to drive the transition at a reasonable rate (on the order of a few times per second). Because the transition is inherently very narrow (nanohertz), the required intensity for a given transition rate is also a function of the linewidth of the probe laser light. The quantity which determines the transition rate is the spectral intensity. The AC Stark shift, meanwhile, is not due to the interaction of the light with the octupole transition, but the off-resonant interaction of the probe light with transitions to other levels from the $^2\text{S}_{1/2}$

and $^2\mathrm{F}_{7/2}$ states. It is therefore dependent on the total light intensity, and largely independent of the probe laser linewidth.

The light shift $\Delta\omega_{LS}$ has the form

$$\Delta\omega_{LS} = \alpha_{LS}(\omega_L) I_{\omega_L}, \tag{4.1}$$

where I_{ω_L} is the intensity of light with angular frequency ω_L incident on the ion, and $\alpha_{LS}(\omega_L)$ is the frequency (wavelength) dependent coefficient of the shift.

4.1.1 Theory

The coefficient can be calculated from first principles by summing the dipole matrix elements of all transitions involving the $^2\mathrm{S}_{1/2}$ and $^2\mathrm{F}_{7/2}$ states. The following calculation was performed by S. Webster. For an initial state $|\gamma IJFM_F\rangle$, summing over transitions of angular frequency ω to $|\gamma' IJ'F'M_F\rangle$, assuming that the electric field of the light is defining the quantisation axis (and hence there is no change in M_F) [67, 68]:

$$\alpha_{LS}(\omega_L) = -\frac{1}{c\epsilon_0\hbar^2}\sum_{\gamma' J' F'} \frac{\omega |\langle \gamma IJFM_F | d_z | \gamma' IJ'F'M_F\rangle|^2}{\omega^2 - \omega_L^2} \tag{4.2}$$

The operator is the z-component of the dipole operator $d_z = -ez$. The matrix elements can be written in terms of weighted oscillator strengths gf and 3j and 6j angular momentum factors.

$$\begin{aligned}\alpha_{LS}(\omega_L) = &-\frac{1}{c\epsilon_0\hbar}\frac{3e^2}{2m_e}\sum_{\gamma' J' F'} \frac{\omega g f}{\omega^2 - \omega_L^2}(2F+1)(2F'+1) \\ &\times \begin{pmatrix} F & 1 & F' \\ -M_F & 0 & M_F \end{pmatrix}^2 \begin{Bmatrix} J & I & F \\ J' & 1 & F' \end{Bmatrix}^2\end{aligned} \tag{4.3}$$

The gf values for transitions in singly-ionised ytterbium have been calculated [69]. Using these values, and summing over the contributions from all transitions with $gf > 0.01$, for the 467 nm probe light gives the shifts for the

$^2\mathrm{S}_{1/2}(F=0)$ and $^2\mathrm{F}_{7/2}(F=3, M_F=0)$ states as

$$\alpha_{LS}(467\ \mathrm{nm}) = 2\pi \times -5.15 \times 10^{-4}\ \mathrm{HzW^{-1}m^2} \quad \text{For the } {}^2\mathrm{S}_{1/2} \text{ state} \tag{4.4}$$

$$\alpha_{LS}(467\ \mathrm{nm}) = 2\pi \times -4.67 \times 10^{-4}\ \mathrm{HzW^{-1}m^2} \quad \text{For the } {}^2\mathrm{F}_{7/2} \text{ state} \tag{4.5}$$

It is the difference between these two level shifts which gives rise to a shift of the transition frequency. Fortuitously, the difference is relatively small, and the calculated coefficient of the AC Stark shift for 467 nm light on the $^2\mathrm{S}_{1/2}(F=0)$ to $^2\mathrm{F}_{7/2}(F=3, M_F=0)$ transition in $^{171}\mathrm{Yb}^+$ is

$$\alpha_{LS} = 2\pi \times 47\ \mu\mathrm{Hz\,W^{-1}\,m^2} \tag{4.6}$$

4.1.2 Measurement

The shift has been observed and measured in our experiments. The most accurate measurement of the coefficient of the shift was performed in 2002. The frequency of the transition was measured (as described in Chapter 5) as a function of the probe laser power, which is then converted to intensity via a measurement of the probe beam diameter at its focus. The beam diameter was measured by observing the power transmitted past a micrometer-mounted knife edge as a function of micrometer position. This was done in both the vertical and horizontal directions. The measured transmission was fitted with a theoretical transmission function to obtain the beam diameters. It is this step which limits the accuracy of the determination of the coefficient of the shift. The beam is approximately 5 μm diameter, and the micrometers have at best 0.5 μm accuracy. The measurements were repeated three times for each direction, but the final uncertainty on the beam size was around 5% in each direction, leading to an 8% uncertainty in the estimated intensity, before considering other sources of error. The uncertainty of the alignment of the UV tracer beam on the ion, and of the overlap of the 467 nm and 369 nm beams, is asymmetric, as the intensity can only decrease as the alignment drifts from the original optimal setting. The uncertainty here is $^{+0}_{-17}\%$. The systematic uncertainty in intensity calibration is $^{+8}_{-15}\%$.

The slope of the transition frequency as a function of power for this measurement was 0.880±0.076 kHz/mW. Combining this with the measured spot size and related uncertainties gives a measurement of the AC Stark shift coefficient as

$$\alpha_{LS} = 2\pi \times 48 \pm 4 \text{ (stat.) } {}^{+7}_{-4} \text{ (syst.) } \mu\mathrm{Hz\,W^{-1}\,m^{2}} \tag{4.7}$$

This is in good agreement with the theoretical calculation (Equation 4.6). It is, perhaps, better than could reasonably be expected given that the theoretical calculation is the difference between two large, similar numbers.

4.1.3 Correcting for the shift

Two methods of correcting for the AC Stark shift have been used in these experiments. The first, used only for the preliminary measurement of September 2001, is to measure the beam size and power used in the experiment, and to calculate a correction using a previously measured coefficient. This leads to large uncertainties, of around 20% in the calculated intensity, largely due to the beam size measurement. This error also appears twice — first in the measurement of the coefficient and then in calculating the correction at a later date. This makes this method unsatisfactory for high-precision work.

The second method is more direct, and misses out the beam size measurement step. Here, the transition frequency is measured directly as a function of probe laser power, and it is assumed that the beam size does not change between measurements. The frequency can then be extrapolated back to zero probe laser power, and this represents the unperturbed transition frequency. The uncertainties involved in this method are otherwise common with the first method. These are the beam pointing stability, the initial alignment accuracy of the 467 nm beam with the 369 nm tracer beam, and the accuracy of the alignment of the tracer beam. These errors become fractionally smaller as the beam diameter increases, and so defocusing the probe laser beam is a good method of controlling the uncertainties, provided that there is enough power in the beam to drive the transition at the required rate.

4.1.4 Scaling of the AC Stark shift

While the AC Stark shift is currently the largest shift, it scales with the probe laser linewidth in an advantageous manner. This scaling is not because the AC Stark shift is dependent on the linewidth (to first order), but because the two are interlinked by the experimental method.

The AC Stark shift is proportional to the intensity I. The transition rate is proportional to the spectral intensity S, itself proportional to I/Γ for a laser of linewidth Γ. The excitation probability for a single interrogation of the ion, of length t, is proportional to $St \propto It/\Gamma$, assuming the interrogation is well below saturation. By setting this quantity to be conserved, whatever the experimental configuration, generalisations about the scaling of the AC Stark shift with linewidth can be deduced. Currently, the excitation probability is around 25% per interrogation at line centre, for a pulse length of 300 ms, a beam waist of 10 μm, and 1 mW of laser power in a 150 Hz bandwidth. Typical pulse lengths for the current state of the experiment are around 100-300 ms. There is little to be gained from reducing this further, as the experimental cycle becomes dominated by the state-detection time. In the situation where t is fixed, the AC Stark shift (proportional to I) is therefore directly proportional to the linewidth. As the linewidth shrinks, the probe time t must be extended, otherwise the linewidth of the laser will be Fourier-limited. In this regime, the pulse length is inversely proportional to linewidth ($t \propto 1/\Gamma$). Hence, for very narrow linewidths $I \propto \Gamma^2$. This is illustrated in Figure 4.1, and possible combinations of pulse length, beam power and beam size to obtain this are shown in Figure 4.2.

Because of the extreme weakness of the octupole transition, and the currently large size of the AC Stark shift it is often thought of as a major weakness of the octupole transition amongst candidate frequency standards. Currently, it must be accepted, it is a major limitation on our ability to make absolute frequency measurements. However, the probe laser linewidth must be narrowed significantly to take full advantage of the octupole transition. The quadratic scaling of the AC Stark shift with linewidth means that the shift is vanishing for small linewidths. For example, if the linewidth can be

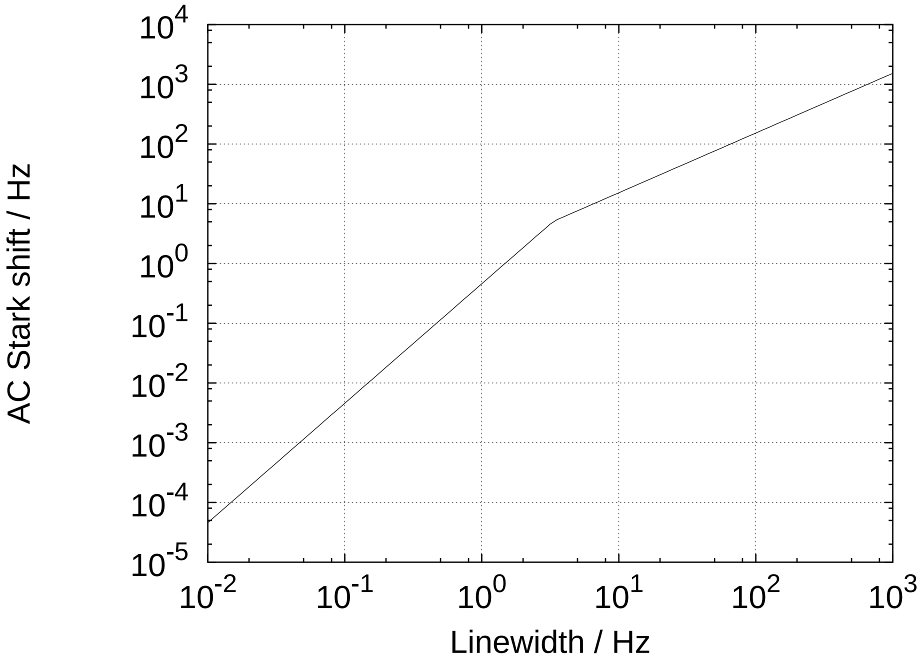

Figure 4.1: Approximate relationship between the AC Stark shift and probe laser linewidth for typical operating conditions. The data are scaled from the current conditions of $\Gamma = 200$ Hz, $I = 10^7$ Wm^{-2} and $t = 0.3$ s. A summary of the operating conditions for this curve is shown in Figure 4.2.

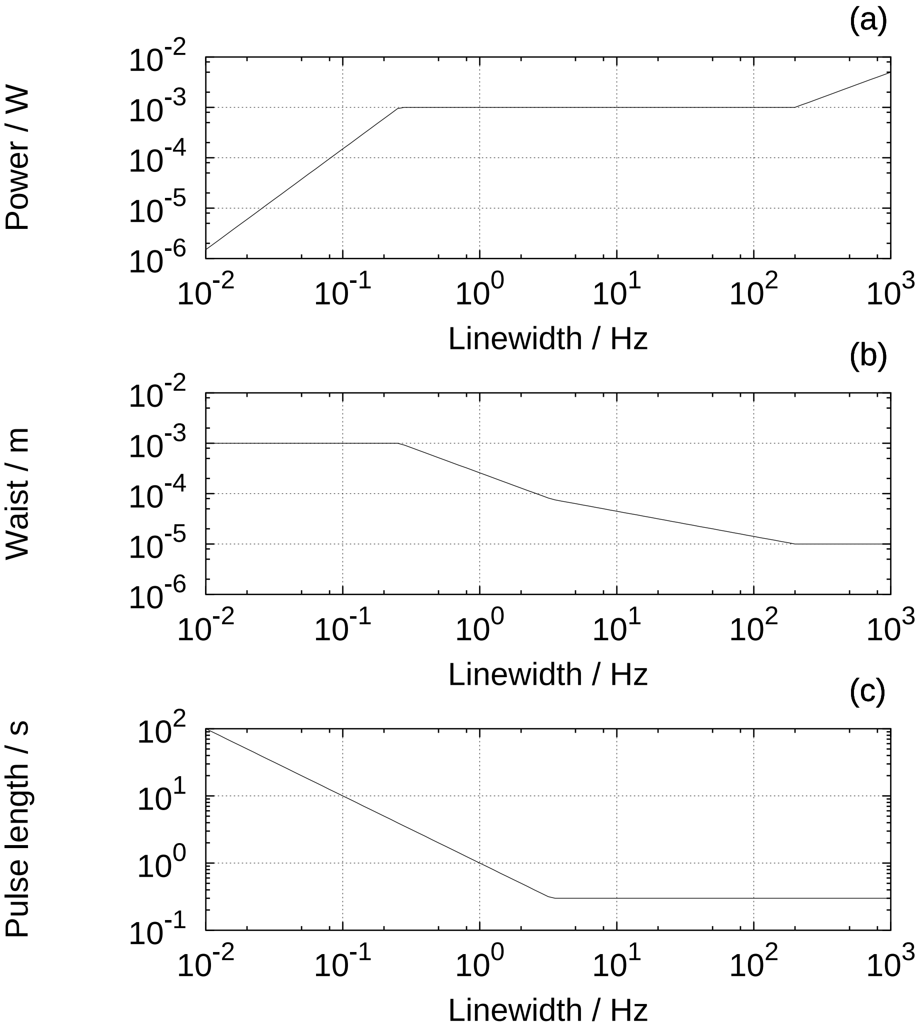

Figure 4.2: Possible operating conditions, as a function of probe laser linewidth, for the AC Stark shift shown in Figure 4.1.(a) Probe beam power. (b) Probe beam diameter at the ion. (c) Probe pulse length. All the conditions are extrapolated from the current operating conditions at $\Gamma = 200$ Hz.

reduced to 0.1 Hz, then the AC Star shift could be reduced to around 5 mHz.

4.2 Second-order Zeeman effect

The second-order Zeeman effect arises from the interaction of a magnetic field with the $^2\mathrm{S}_{1/2}$ and $^2\mathrm{F}_{7/2}$ states. The transition frequency is not affected by the linear Zeeman shift as the $m_F = 0 \rightarrow m_F = 0$ component is used as the frequency reference. However, this component is still affected by the second-order Zeeman shift.

4.2.1 Calculation

The shift $\Delta\omega_{Z2}$ induced by a magnetic field strength B has the form

$$\Delta\omega_{Z2} = \alpha_{Z2} B^2 \tag{4.8}$$

The coefficient α_{Z2} can be calculated as follows. This calculation was originally performed by G. P. Barwood, and later by S. Webster. The coefficient can be found using second-order perturbation theory, to be

$$\alpha_{Z2}(F, M_F) = -\left(\frac{\mu_B g_J}{\hbar}\right)^2 \sum_{F'} \frac{|\langle F' M_F | J_z | F M_F \rangle|^2}{\omega_{F'} - \omega_F} \tag{4.9}$$

This is written in angular frequency units, with $\omega_F = E_F/\hbar$, and assumes that the quantisation axis is set by the applied field, so the matrix elements are zero for states of different M_F. The matrix elements can be written in terms of angular momentum factors

$$\begin{aligned} |\langle F' M_F | J_z | F M_F \rangle|^2 =& I(I+1)(2I+1)(2F+1)(2F'+1) \\ &\times \begin{pmatrix} F & 1 & F' \\ -M_F & 0 & M_F \end{pmatrix}^2 \begin{Bmatrix} I & F & J \\ F' & J & 1 \end{Bmatrix}^2 \end{aligned} \tag{4.10}$$

For the $^2\mathrm{S}_{1/2}(F=0)$ state in $^{171}\mathrm{Yb}^+$, the only non-zero matrix element is that with the $^2\mathrm{S}_{1/2}(F=1, M_F=0)$ state. The coefficient for this state is

$$\alpha_{Z2}(0,0) = -\left(\frac{g_J\mu_B}{\hbar}\right)^2 \frac{|\langle 10|J_z|00\rangle|^2}{\Delta\omega_{HFS}} \tag{4.11}$$

The square of the matrix element is equal to 1/2, $g_J = 2$ and the hyperfine splitting $\Delta\omega_{HFS} = 2\pi \times 12.643$ GHz. This makes the shift coefficient for the $^2\mathrm{S}_{1/2}(F=0)$ state $\alpha_{Z2}(0,0) = -15.49$ mHz $(\mu\mathrm{T})^{-2}$. The upper state of the reference transition is the $^2\mathrm{F}_{7/2}(F=3, M_F=0)$ state. The only non-zero matrix element for the perturbation of this state is that with the $^2\mathrm{F}_{7/2}(F=4, M_F=0)$ state. For this, the squared matrix element is 1/2, $g_j = 8/7$ and $\Delta\omega_{HFS} = 2\pi \times 3.620$ GHz, giving a coefficient $\alpha_{Z2}(3,0) = -17.62$ mHz $(\mu\mathrm{T})^{-2}$. The shift of the transition frequency is the difference between these two coefficients and hence the predicted second-order Zeeman shift for the octupole transition in $^{171}\mathrm{Yb}^+$ is

$$\alpha_{Z2} = 2\pi \times -2.12 \text{ mHz } (\mu\mathrm{T})^{-2} \tag{4.12}$$

The shift has been measured (Section 5.4.3) as -1.72 ± 0.03 mHz $(\mu\mathrm{T})^{-2}$. The theoretical calculation does not precisely match the measured coefficient. Despite the discrepancy of about 20% in the shift coefficient, the difference is equivalent to a less than 3% error in the calculations of the individual level shifts. The agreement observed between theory and experiment is of the level which can be expected from calculations of this nature. The measured coefficient has been used for calculating all frequency corrections in this thesis.

4.2.2 AC Second-order Zeeman shifts

Because the shift is second-order, an oscillating perturbation will not average to zero. For low frequencies $\omega \ll \omega_{HFS}$, the coefficient of the shift is the same as for a DC field. For a field composed of static and oscillating components,

$$\langle B^2 \rangle = \langle B^2_{DC} \rangle + \langle B^2_{AC} \rangle \tag{4.13}$$

The shifts can therefore be calculated separately and independently for the different frequency components of a field. The shift caused by a low-frequency oscillating field will be equal to that produced by a static field of the same strength as the rms (root-mean-square) oscillating field strength. The most obvious source of such a low-frequency magnetic field is the mains electricity supply.

For higher frequency fields, the shift is calculated in a similar manner to the AC Stark shift. The coefficient for a field of angular frequency ω_L interacting with a state $|\gamma IJFM_F\rangle$ is

$$\alpha_{ACZ2}(\gamma, I, J, F, M_F) = -\left(\frac{\mu_B g_J}{\hbar}\right)^2 \sum_{J'F'} \frac{(\omega' - \omega)|\langle \gamma I J' F' M_F | J_z | \gamma I J F M_F \rangle|^2}{(\omega' - \omega)^2 - \omega_L^2} \tag{4.14}$$

It is expected that this shift is much smaller than the DC shift due to the scarcity of magnetic dipole transitions at optical frequencies, and the interaction with microwave magnetic dipole transitions is smaller than for DC by a factor $[(\omega' - \omega)/\omega_L]^2$. While the coefficient of the optical second-order Zeeman shift has not been fully calculated, this is not an issue for the experiment. The shift scales with B^2, and is therefore a linear function of the intensity of an optical field, as with the AC Stark shift. It is therefore automatically compensated for by the experimental method.

4.2.3 Reducing the shift

It is, at present, necessary to apply an external DC magnetic field to the trapped $^{171}\mathrm{Yb}^+$ ion. This is because there is a fluorescence null on the $F = 1 \rightarrow 0$ component of the 369 nm transition in zero field [70], and this is the component used for Doppler cooling. This field is typically 200 μT, giving a shift on the order of 80 Hz. Due to the accuracy of the calibration, the uncertainty on this is small ($\sim$ 3 Hz). This can still be improved by working in a lower, or zero, magnetic field. To do this it is necessary to spin the polarization of the Doppler cooling beam to prevent loss of fluorescence [70]. The technique involves projecting two beams into the trap at an angle to each other, and modulating the phase of the polarisation components of the light such that all three polarisations have a different time-dependence. This has been tested in ^{171}Yb^{+} and the technique does maintain fluorescence in magnetic fields down to a few tens of nanotesla. Below this, a narrower fluorescence null on the 935 nm repumper transition begins to affect the fluorescence rate. This can also be overcome by polarisation spinning, enabling work in zero (or very close to zero) DC magnetic field. The field itself can be nulled very accurately by observation of the Zeeman components of the octupole transition. The lower limit to the shift is expected to be set by the 50 Hz field from the mains electricity supply. This has been measured at around 500 nT in the current laboratory, and it is expected that a mu-metal shield around the ion trap could reduce this to around 50 nT. This shifts the transition by just a few microhertz.

4.3 Quadrupole shift

The electric quadrupole shift is described in detail in Appendix A, and the results are reproduced here. The shift occurs only for states with $J \geq 2$, and does not affect the ^{2}S$_{1/2}$ state. The physical cause is the interaction of the electric quadrupole moment of an electronic state with a quadrupolar electric field, such as might be produced by the electrodes of an ion trap. The shift for the 467 nm transition is caused by any DC electric quadrupole field which may be present in the trap interacting with the quadrupole moment of the ^{2}F$_{7/2}(F=3, M_F=0)$ state in ^{171}Yb^{+}. The shift is

$$\langle H_Q \rangle = +\frac{5}{7} A \left(3\cos^2\beta - 1\right) \Theta(\gamma, J) \tag{4.15}$$

The quantity A represents the quadrupole field strength for a potential $\phi\,(x', y', z') = A\,(x'^2 + y'^2 - 2z'^2)$. The quadrupole moment of the state is calculated (in the Appendix) to be $\Theta(4f, 7/2) = -0.22ea_0^2$ (a_0 is the Bohr radius). For the ion trap used in these experiments, the angle between the

B-field and the trap axis is $\beta = 90°$. The shift is then

$$\langle H_Q \rangle = +7.02 \times 10^{-41} A \text{ J} \tag{4.16}$$

$$\frac{\langle H_Q \rangle}{h} = +1.05 \times 10^{-7} A \text{ Hz} \tag{4.17}$$

By assuming that the field is caused by a common offset potential V_{Offset} on the upper and lower inner endcaps (see Section 3.1 for details of the electrode structure), relative to the outer endcap electrodes, typical values for A can be estimated using the equation

$$A = -\frac{\mathcal{E} V_{\mathit{Offset}}}{2z_0^2} \tag{4.18}$$

The 'efficiency' of the trap is $\mathcal{E} = 0.63$ and the inner electrode separation is $2z_0 = 1.12$ mm (Section 3.1). When no potential is deliberately applied, V_{Offset} is expected to be in the range ±1 V, due to offset potentials on the electrodes. Possible sources of these potentials include contact potentials between atomic ytterbium ejected from the trap ovens and the tantalum of the electrode structure [71]. The 1 V estimate is backed up by the fact that micromotion compensation voltages required in the z direction are typically $\pm 0 \rightarrow 0.2$ V. Using 1 V as a conservative estimate gives typical values of A in the range $\pm 1.0 \times 10^6$ Vm^{-2}. The uncorrected quadrupole shift is therefore expected to be approximately ±0.1 Hz. By assuming that the field is produced by the trap electrodes, it is possible to express the quadrupole shift in "hertz per volt on the inner endcap". This is an artificial, but practically useful, way to express the shift and is unique to each individual trap. When this is done, the shift is 0.11 HzV^{-1}.

4.3.1 Nulling the shift

It may be possible to control and reduce both the shift and its uncertainty by various means. The first of these is that if frequency measurements are made for any three orthogonal directions of the magnetic field, the quadrupole shift sums to zero over the three measurements [72].

A second method is to use the information provided by the secular frequencies of the trapped ion to either infer the size of the quadrupole field, or to null it. Equations 2.11 and 2.12 describe the axial and radial secular frequencies of the trap as a function of the a and q parameters, which are a function of the DC and AC quadrupole fields, respectively. For example, for our trap parameters $V_{ac} = 291$ V, $\mathcal{E} = 0.63$ and $\omega_{RF} = 2\pi \times 12.8$ MHz, the predicted ratio of the trap frequencies for zero DC quadrupole potential is $r = \omega_z/\omega_r = 2.054$. By adjusting the potentials applied to the trap electrodes such that this ratio is realised, with good micromotion control, then it is known that there is no DC quadrupole field in the trap, and the shift has been nulled[1] [73]. The accuracy with which this ratio may be set is limited by the knowledge of the trap efficiency $\mathcal{E}$ which, from experimental characterization of the trap is 5%. This corresponds to an uncertainty in r of 0.005 for $r = 2.054$. The gradient $\mathrm{d}r/\mathrm{d}V_{DC}$ for our operating conditions is calculated to be 0.10 V^{-1}. Combining these gives an uncertainty in locating $V_{DC} = 0$ of 0.02 V, corresponding to a quadrupole shift of 7 mHz.

The trap frequencies can also be used to infer the quadrupole potential and therefore can be used to measure the quadrupole moment. This will give a more reliable measure of the quadrupole moment than the atomic structure calculations. This method has recently been used at NPL to measure the quadrupole moment of the upper state of the 674 nm electric quadrupole transition in $^{88}\mathrm{Sr}^+$ [74].

It is interesting to note that for isotopes with nuclear spin $I = 7/2$, the electric quadrupole shift of the $^2\mathrm{F}_{7/2}$ state is zero. There are no stable isotopes of ytterbium with this nuclear spin, though. The unstable isotopes are $^{155}\mathrm{Yb}$ (1.8 s lifetime), $^{169}\mathrm{Yb}$ (32 days), and $^{175}\mathrm{Yb}$ (4 days). They are not naturally occurring, and not easy to obtain.

---

[1] Assuming the potentials on the outer endcaps are being adjusted, it is the common potential of these which gives rise to the quadrupole shift, while the difference potential causes micromotion.

4.3.2 Related shifts

In a similar vein to the quadrupole shift, the $^2\mathrm{F}_{7/2}$ state may also have an electric octupole shift, caused by any DC electric octupolar field in the trap interacting with the electric octupole moment of the state. It is expected that this is much smaller than the quadrupole shift both because the shift itself is weaker and because any octupole field in the trap will be comparatively weaker than the quadrupole field.

There is also the possibility of a second-order quadrupole shift affecting the $^2\mathrm{F}_{7/2}$ state. Even in zero DC quadrupole field, there is still a strong ($A \sim 3 \times 10^8$) AC quadrupole field, confining the ion. While the lowest-order quadrupole shift calculated here averages to zero for this AC field, any second-order shift would not. It is expected that this will again be smaller than the shift due to the DC quadrupole field, or than that due to the DC and low frequency Stark effect, as the coupling with other atomic states occurs via the electric quadrupole interaction and the interaction is a much weaker one than the electric dipole interaction.

Another shift which could occur, but has not been investigated in detail, is a magnetic quadrupole (M2) shift. This would only affect the $^2\mathrm{F}_{7/2}$ state, and is expected to be very small, for several reasons. The applied magnetic field is small, and will have a correspondingly small magnetic quadrupole component. Also, the magnetic quadrupole interaction is a weak one.

4.4 DC and low-frequency Stark shift

The DC quadratic Stark shift is simply the static-field equivalent of the AC Stark shift, discussed in Section 4.1. The shift has the form

$$\Delta\omega_{DCS} = \alpha_{DCS} E^2 \tag{4.19}$$

In a low frequency field, where the frequency is much lower than any electric dipole transitions, the shift has the same coefficient, the relevant physical quantity being the mean-square electric field. The calculation of the shift is similar to that for the AC Stark shift, with the difference that the frequency-dependence of the shift is not considered. The shift of a level $|\gamma IJFM_F\rangle$ is

$$\alpha_{DCS} = -\frac{1}{\hbar^2}\sum_{\gamma' J' F'} \frac{|\langle \gamma IJFM_F|d_z|\gamma' IJ'F'M_F\rangle|^2}{\omega} \tag{4.20}$$

As for the AC shift, the matrix elements may be calculated in terms of tabulated weighted oscillator strengths gf. When this is done, the calculated shift of the 467 nm transition is $\alpha_{DCS} = -1.44 \times 10^{-7}$ HzV^{-2}m^2, in frequency units. The calculation assumes that the electric field is defining the quantisation axis. It is actually the applied magnetic field which defines the quantisation axis of our experiment, but this is expected only to make a small (10% or so) difference to the calculated shift.

One source of the DC Stark shift is the displacement of the ion from the centre of the trap, where it experiences a non-zero electric field at the trap drive frequency. This electric field causes micromotion, and so good micromotion control is required to minimise this shift. For recent experiments, micromotion minimisation has been performed in two dimensions (one of them being the probe laser direction). In one experiment, the micromotion was minimised, and the experiment left to run for one hour, and then the micromotion was measured using rf-photon correlation. This represents the typical maximum level of micromotion during frequency measurements, as it was done at a time of relatively high micromotion drift. The measured mean offset of the ion from the centre of the trapping potential was $\bar{x} = +70$ nm, $\bar{y} = +150$ nm, $\bar{z} = -310$ nm. The electric field experienced by the ion (mass m, charge e in a trap drive frequency ω_{RF}) is

$$\left\langle E^2 \right\rangle = \left(\frac{m\omega_{RF}}{e}\right)^2 \left((\omega_x\bar{x})^2 + (\omega_y\bar{y})^2 + (\omega_z\bar{z})^2\right) \tag{4.21}$$

This gives us an estimated field strength of 5.5×10^5 V^{2}m^{-2}, and an estimated shift of -0.08 Hz. The uncertainty on the calculation is unknown, as there has been no opportunity to calibrate this shift experimentally, but it is expected to be of the correct order of magnitude, as the same calculation for the AC Stark shift gives a coefficient which agrees well with the measured value, as discussed in Section 4.1.

4.4.1 Reducing the shift

If the micromotion minimisation is performed in three dimensions, such that the modulation of the fluorescence signal at the drive frequency is less than 1% peak-to-peak (equivalent to an offset from the trap centre of a few nanometres in each direction), then the micromotion-induced DC Stark shift is much smaller than the DC Stark shift due to the secular (thermal) motion of the ion in the trap. For a Doppler-cooled ion, the amplitude of the thermal motion is of the order 20 nm along the trap z-axis and 40 nm along the x and y axes, leading to a DC Stark shift of -0.8 mHz. If the ion is cooled to the ground vibrational state of the trap, then the shift is -0.3 mHz. It is unlikely that this limit will be achieved for long probe times, due to ion heating, but we do not know how large an effect this is for the current ytterbium ion traps.

4.5 Second-order Doppler shift

The second order Doppler shift, like the DC Stark shift, is directly related to the motion of the ion. If the ion is observed at an angle θ to its direction of motion, and is moving with a velocity v, the observed transition frequency ν_{Obs} is related to the emission frequency ν_{Em} by the relationship [75]

$$\frac{\nu_{Obs}}{\nu_{Em}} = \frac{\sqrt{1 - v^2/c^2}}{1 - (v/c)\cos\theta} \tag{4.22}$$

This can be expanded in powers of (v/c)

$$\begin{aligned} \frac{\nu_{Obs}}{\nu_{Em}} &= \left(1 - \frac{v^2}{2c^2} + \dots\right)\left(1 + \frac{v}{c}\cos\theta + \frac{v^2}{c^2}\cos^2\theta + \dots\right) && (4.23) \\ &= 1 + \frac{v}{c}\cos\theta + \left(\cos^2\theta - \frac{1}{2}\right)\frac{v^2}{c^2} + \dots && (4.24) \\ &= 1 + \frac{v}{c}\cos\theta + \frac{v^2}{c^2}\cos 2\theta + \dots && (4.25) \end{aligned}$$

The fractional shift in the transition frequency due to the second-order Doppler shift is $\frac{1}{2}(v^2/c^2)\cos 2\theta$. At present, where micromotion minimisation is performed for 2 dimensions only, the dominant motion of the ion is perpendicular to the direction of observation, and so the transverse Doppler shift (time dilation) is observed. This gives a shift of -0.05 Hz for the measured micromotion amplitude in the direction where the micromotion is not controlled. If the micromotion is well minimised, then the thermal motion of the ion is the dominant motion. The motion in this case is random, with no well-defined angle to the direction of observation, and the shift has an expectation value of 0. The uncertainty on the shift is then around 0.3 mHz for a Doppler cooled ion (and 35 μHz for an ion in the vibrational ground state of the trap).

4.6 Blackbody shifts

The stainless steel vacuum chamber, its viewports, and the trap electrodes all contribute to the blackbody radiation field experienced by the ion. The electric and magnetic components of this field respectively lead to Stark and second-order Zeeman shifts. The larger of these is the Stark shift. The peak of the blackbody spectrum at room temperature lies well below the frequencies of electric dipole resonance transitions from the S and F states. At 297 K, half of the intensity is at wavelengths longer than 14 μm, and 99% is longer than 5 μm. In Equation 4.2, the denominator $(\omega^2 - \omega_L^2) \approx \omega^2$ for blackbody radiation. Therefore, the blackbody radiation field can be approximated by a DC field of the same mean-square field strength for the calculation of the AC Stark shift.

The electric field strength is given by $\langle E^2 \rangle = 6.29 \times 10^5 (T_{BB}/293)^4\ \mathrm{V^2 m^{-2}}$, where T_{BB} is the temperature of the radiation field. This gives a blackbody Stark shift of -0.09 ± 0.03 Hz at 297 K. The coefficient of the shift is -1.23×10^{-11} HzK^{-4}.

The energy in the magnetic field is equal to the energy in the electric field, and the magnetic field strength is given by $\langle B^2 \rangle = 6.97 \times 10^{-12} (T_{BB}/293)^4\ \mathrm{T}^2$. This gives a blackbody second-order Zeeman shift of -13 mHz, using the DC shift coefficient. However, this is not appropriate, as around 99% of the radiation is at frequencies higher than 1.8 THz (175 μm), and this is a much higher frequency than the 12.6 GHz and 3.6 GHz transitions taken into

Shift	**Coefficient**	**Units**
AC Stark	$(48 \pm 8) \times 10^{-6}$	Hz W^{-1} m^2
DC Stark	-1.44×10^{-7}	Hz V^{-2} m^2
Second-order Zeeman	$(-1.716 \pm 0.025) \times 10^{9}$	Hz T^{-2}
Quadrupole	1.05×10^{-7}	Hz V^{-1} m^2
Blackbody AC Stark	-1.23×10^{-11}	Hz K^{-4}

Table 4.1: The well-characterised atomic systematic shifts of the octupole transition in ^{171}Yb$^+$. The uncertainties shown are from experimental calibrations.

account in calculating the DC shift coefficient. Hence their contribution to the blackbody second-order Zeeman shift is very small. There are magnetic dipole transitions in ^{171}Yb$^+$ with wavelengths in the 1–3μm range, and it is likely that these will make a larger contribution to the coefficient, but this has not been calculated.

4.7 Gravitational redshift

The gravitational redshift is a general relativistic effect, whereby clocks run comparatively slower in stronger gravitational fields. It is necessary to take into account the height of a clock above some defined surface of constant gravitational field strength when coordinating clocks in different locations to produce a global timescale. The relative shift is approximately 1 part in 10^{16} per metre of height above the mean geoid (a ‘mean earth’) at sea level.

4.8 Summary of coefficients

A summary of the well-known and characterised coefficients of the atomic systematic shifts for the octupole transition is shown in Table 4.1. Shifts such as the second-order Doppler shift and gravitational redshift are well-characterised shifts which are independent of the atomic physics of ytterbium ions. Coefficients not included, because they are not well known, include the blackbody and optical second-order Zeeman shifts, the octupole shift and the second-order quadrupole shift.

4.9 Projected limits

Using the measured and calculated coefficients of the systematic shifts, it is possible to estimate the typical sizes of all the shifts for current operating conditions. It is also possible to project the levels to which the shifts may be reduced by improvements in the experiment. The shifts and their uncertainties obtained by this process are tabulated in Table 4.2, and plotted in Figures 4.3 and 4.4. The assumptions for the projected shifts are:

AC Stark shift Moving from 1 mW of light, in a 200 Hz bandwidth, with a spot size of 10 μm and a pulse length of 0.3 s, to 20 μW of light in a 0.1 Hz bandwidth, with a 1 mm spot size and a 10 s interrogation time.

DC Stark shift The current estimate is based on 'worst case' measurements of the motion of the ion, obtained by rf-photon correlation. These give mean displacements of the ion from the centre of the rf field as +150 nm, +70 nm and -310 nm along the x, y, z trap axes, respectively. The projection assumes that micromotion can be reduced to below the level at which it causes a 1% modulation of the rf-photon correlation signal for all three directions. This is easily achieved for one or two dimensions currently. In this situation, the secular motion of the ion is much larger than the micromotion. The projected limit is for an ion cooled to the ground vibrational state of the trap.

DC second-order Zeeman shift The current applied DC field is 188 μT, and it is anticipated that this can be reduced to near zero by using polarisation spinning on the 369 nm and 935 nm transitions. A small field will be required to set a quantisation axis for the ion, and separate the Zeeman components of the 467 nm transition. The projected uncertainty assumes that the field can be nulled only with the currently available accuracy of about 1 μT. In fact, the field could be set much more accurately by observation of the Zeeman components of the octupole transition.

50 Hz second-order Zeeman shift Based on measurements made on the NPL strontium-ion frequency standard, the 50 Hz (mains electricity) field is expected to be 500 nT at the present time, and this can be reduced by a factor of 10 with the addition of mu-metal shields around the trap.

Blackbody stark shift The shift is expected to remain the same size, as the operating temperature of the experiment will not change in the near future. The uncertainty is a rough estimate of our confidence in the theoretical calculation of the DC Stark shift coefficient.

Quadrupole shift The effective quadrupole potential on the trap electrodes is assumed to be rectangularly distributed between +1 V and -1 V at the present time. It is expected that the uncertainty can be reduced to 0.02 V by monitoring the trap frequencies.

Second-order Doppler shift The current shift is estimated from measured micromotion amplitudes, as described for the DC Stark shift. The projected shift assumes an ion cooled to the ground vibrational state of the trap.

The biggest single shift is currently the AC Stark shift, but as the linewidth of the probe laser narrows, and interrogation times are extended, this shrinks to below the blackbody Stark shift, and has a much smaller uncertainty.

Polarisation spinning of the cooling light at 369 nm has already been demonstrated, and this should allow a large reduction of the second-order Zeeman shift.

After the second-order Zeeman shift, the largest systematic shift (in terms of its uncertainty) is currently the quadrupole shift. This can be reduced to the projected levels by monitoring the trap frequencies and adjusting the dc potentials applied to the trap.

By minimizing micromotion in three dimensions the DC Stark and second-order Doppler shifts can be reduced to very low levels, assuming that the ion

Systematic	**Current**		**Projected**	
Shift	**Shift / Hz**	**± / Hz**	**Shift / Hz**	**± / Hz**
AC Stark	$+3.1 \times 10^{2}$	1.2×10^{2}	$+4.6 \times 10^{-3}$	4.6×10^{-4}
DC Stark	-7.9×10^{-2}	2.4×10^{-2}	-2.9×10^{-4}	8.8×10^{-5}
DC Zeeman	-6.1×10^{1}	2.6	0	1.72×10^{-3}
50 Hz Zeeman	-4.3×10^{-4}	4.3×10^{-5}	-4.3×10^{-4}	4.3×10^{-5}
BB Stark	-9.1×10^{-2}	2.7×10^{-2}	-9.1×10^{-2}	2.7×10^{-2}
Quadrupole	0	1.2×10^{-1}	0	6.7×10^{-3}
Doppler	0	5.0×10^{-2}	0	3.4×10^{-5}

Table 4.2: Current and projected systematic shifts and associated uncertainties for the octupole transition. The assumptions behind these are outlined in the text.

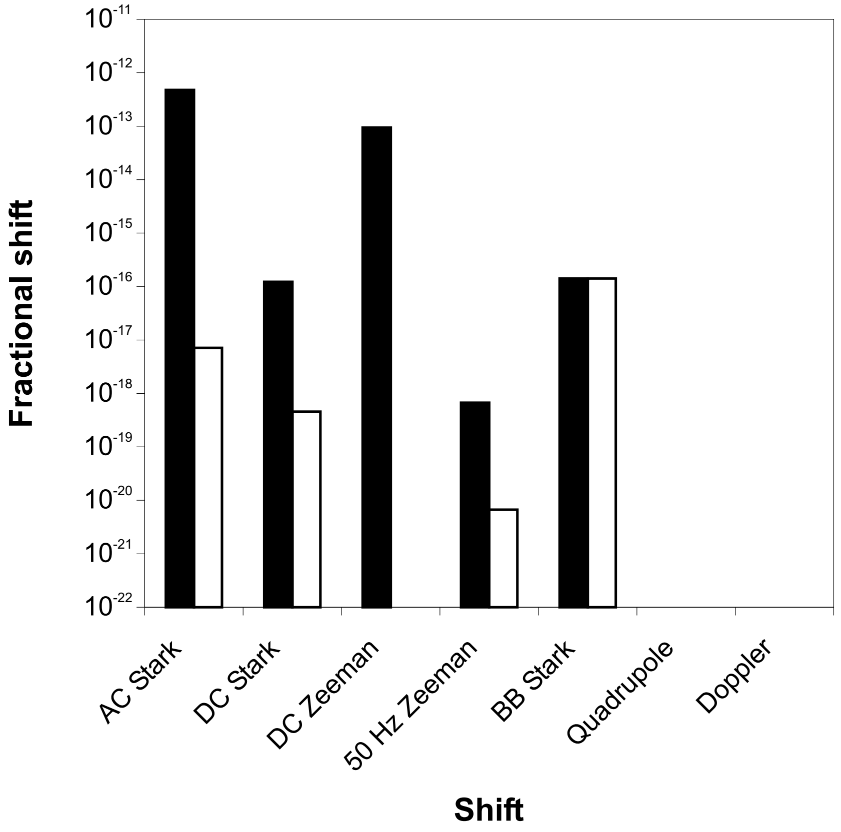

Figure 4.3: Current (black) and projected (white) systematic shifts of the octupole transition, as a fraction of the transition frequency.

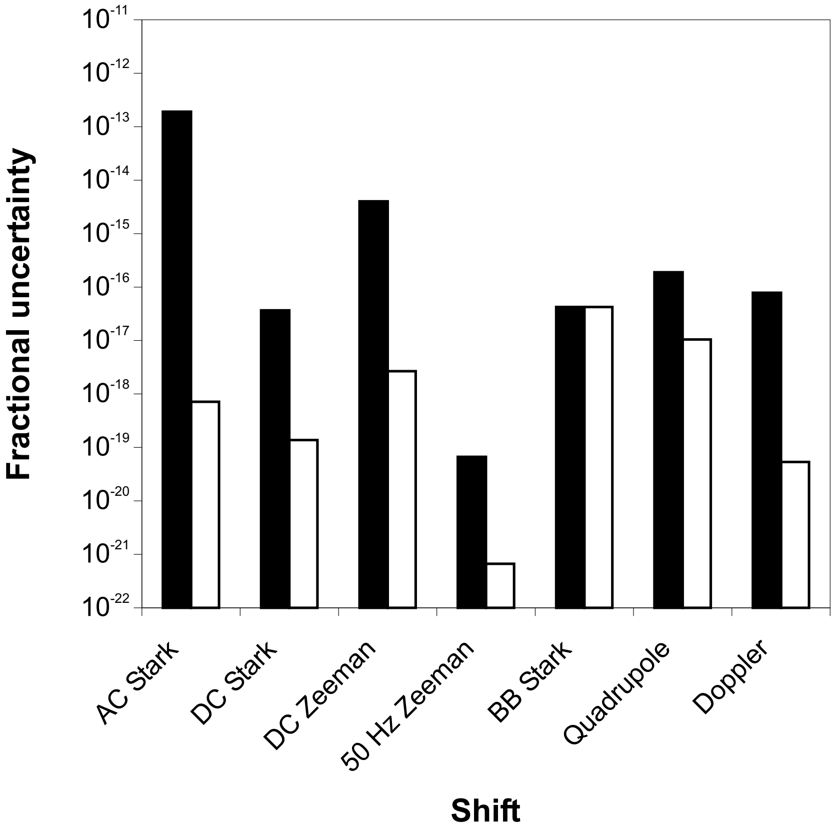

Figure 4.4: The uncertainties of current (black) and projected (white) systematic shifts, as a fraction of the octupole transition frequency.

is cooled close to the Doppler limit. The uncertainty of the blackbody Stark shift is mostly due to the uncertainty in the DC Stark shift coefficient. Without improving this, it appears that operation of the standard at the 5×10^{-17} level of frequency instability should be possible. There is potential for this to be improved by direct measurement of the DC Stark shift coefficient (for example, by measuring the transition frequency as a function of micromotion amplitude in a known direction), but this is not taken into account. Another, slightly more questionable, method of reducing the blackbody Stark shift is to define an arbitrary temperature (near room temperature) as being the 'operating temperature' of the environment of the trapped-ion frequency standard. The shift, and its uncertainty, would then be very small, but this is a philosophically unappealing scheme.

If the uncertainty on the blackbody Stark shift could be reduced by calibration or measurement of the coefficient, the limiting factor in the projected systematic shifts would then be the quadrupole shift. This is projected to be at the 1×10^{-17} level, though careful measurement and calibration could reduce this uncertainty further. With a 0.1 Hz laser linewidth, the line profile must be subdivided by a factor of 20 to achieve a 1×10^{-17} frequency measurement uncertainty. To reduce it further below the projected levels, the scheme of averaging frequencies measured for three perpendicular directions of magnetic field can be applied [72]. The resolution of this technique is limited by the laser linewidth and variations of other systematic shifts, and will require a laser much narrower than 0.1 Hz. For the near-to medium-term, it therefore appears that the achievable performance will be around this level. This is based on the (large) assumption that shifts not dealt with here, such as the blackbody second-order Zeeman shift, octupole shift, etc, are not significant on these scales.

For absolute frequency measurements, this analysis ignores issues involved in relating the transition frequency to the SI second via caesium. Caesium standards have not reached, and may not reach, this level of accuracy. Once the projected performance levels are achieved, as they are expected to be over the next few years by a number of trapped-ion frequency standards, the case for re-definition of the second in terms of an optical standard will grow.

The low systematic shifts and clear potential for further development offered by the octupole transition place it as one of the strong candidates to be the next primary standard.

Chapter 5

Frequency measurements

This chapter covers all the absolute frequency measurements of the 467 nm transition in ^{171}Yb$^+$ made using the NPL femtosecond comb during the period August 2001 to July 2003. As the experiment evolved during this period, there are significant differences in the experimental setup between these measurements, and these are noted in the relevant sections. As will become clear, there have been some inconsistencies in the results of these measurements, and this chapter also attempts to investigate the possible causes of this.

It should also be noted that in this chapter the most accurate estimates available of the coefficients of systematic shifts are used. These may be different to our best estimates at the time the measurements were made, and so results quoted here are not all identical to those published at the time.

5.1 Measurement method

The measurement of the absolute frequency of the transition is a multi-step process. For all the measurements detailed here, the basic procedure was the same:

1. The spectrum of the octupole transition at 467 nm is recorded by quantum jump spectroscopy, using the high finesse ULE cavity as a local reference.

2. Light locked to the ULE cavity at 934 nm is sent via a 50 m optical fibre to the femtosecond comb laboratory.

3. The femtosecond comb measures the frequency of the 934 nm light with reference to a stable 10 MHz reference signal, derived from a hydrogen maser.

4. The H-maser frequency is known, to a certain degree of accuracy, with reference to a caesium primary standard and hence the SI definition of the second.

There are several steps in the process, each with associated uncertainties. For the measurements here, however, it is the ion trap system which is the limiting factor. The femtosecond comb, operated by H. S. Margolis for these experiments, is sufficiently good to be treated as a 'black box' for the purposes of the following measurements. More details of the comb itself are given in Section 3.11.

5.1.1 Trap operation

The ion trap measurement procedure requires careful attention. There are numerous experimental parameters which drift with time, and must be monitored to ensure that systematic shifts are kept as low as possible. The AC Stark shift is the largest shift in these experiments, and is particularly susceptible to drift. This is because the probe beam is focused tightly on the ion, and any drift of the optics along the beam path can produce large fractional changes in the 467 nm intensity on the ion.

Therefore, the alignment of the probe beam on the ion, along with at least 2D micromotion compensation (along the probe beam direction and along the Doppler cooling beam) was carried out every few scans. For the majority of the data, it was every four scans, which is about every 40 minutes. This frequency ensured that large changes to the alignment or micromotion were not necessary, and the variation of these parameters could be estimated.

5.1.2 Data taking

The ion trap frequency measurement data was taken in the form of quantum jump spectra of the octupole transition. These were simple, linear scans. At each frequency, a number of probe-detect cycles were completed, before the frequency of the probe laser, relative to the ULE reference cavity, was stepped by a fixed amount (of much less than the laser linewidth), and the cycle repeated. The total scan range was typically 5 times the FWHM laser linewidth. If a large enough range is not picked, then the results are artificially narrowed and offset because the wings of the spectrum will be clipped. There is essentially no background level of quantum jumps (1 to 2 a day, collisionally induced during the cooling cycle) assuming the fluorescence threshold has been set correctly. Because there is no background, the choice of scan range does not affect the result provided it is large enough. If there were a background, and the range was not centred on the transition, the mean frequency would be pulled towards the centre of the scan range.

Scans were taken in alternate directions (scanning upwards in frequency, then scanning downwards), to give some resilience against drift of the ULE cavity frequency. If it is in the same direction as the cavity drift, the transition will appear wider than it really is, and more jumps will be observed. When scanning against the drift, the opposite happens. Scanning in both directions provides some insurance against this systematically skewing the results.

5.2 September 2001

The first measurement of the $^2\mathrm{S}_{1/2}(F=0)$ to $^2\mathrm{F}_{7/2}(F=1, m_F=0)$ transition in $^{171}\mathrm{Yb}^+$ using the femtosecond comb was made on 5 September 2001 [76, 77]. The crucial difference between this and subsequent measurements lies in the way the AC Stark shift was accounted for. For these data, a measurement of the size of the probe beam focus was made using a micrometer-mounted knife edge. This was accurate to approximately 10%. The power in the 467 nm probe beam was kept constant during the measurements, and

Date	5 September 2001
Probe laser linewidth	~ 4.5 kHz
Focusing	Tight focus (5 μm)
Magnetic field	$617 \pm 4\mu$T
Software	Visual Basic
Mechanical shutters	467 nm beam only
AOM switched	369 nm, 638 nm beams
Not switched	935 nm beam

Table 5.1: Experimental conditions for the September 2001 frequency measurement.

so a single correction could be calculated for all the data using a previous measurement of the AC Stark shift coefficient [67]. This measurement of the coefficient was itself limited by the accuracy in the measurement of the beam size at the time.

The experimental parameters are shown in Table 5.1. The uncertainty of the AC Stark shift correction completely dominates the other sources of uncertainty in this measurement, including the fact that the frequency comb was not as stable then as it was in future measurements.

The absolute frequencies for this measurement were found by taking the mean frequency of each individual scan relative to the reference cavity, and combining that value with a fitted measurement of the femtosecond-comb measurement of the cavity frequency at the mid-point of the scan. Corrections for cavity drift were not done on a per-jump or per-frequency-bin basis. The results are shown in Figure 5.1.

The uncorrected mean frequency of the transition is $642\ 121\ 496\ 776.1 \pm 0.3$ kHz. This must then be corrected for systematic shifts, in particular the AC Stark shift. The beam power was measured, and kept at 4.5 ± 0.5 mW during the scans, measured before the focussing lens. The lens and final beam-steering mirror attenuate the beam, so the actual power on the ion was 2.8 mW. The beam waist was measured in the horizontal and vertical directions to be $w_x = 5.18 \pm 0.41$ μm and $w_y = 3.55 \pm 0.34$ μm respectively.

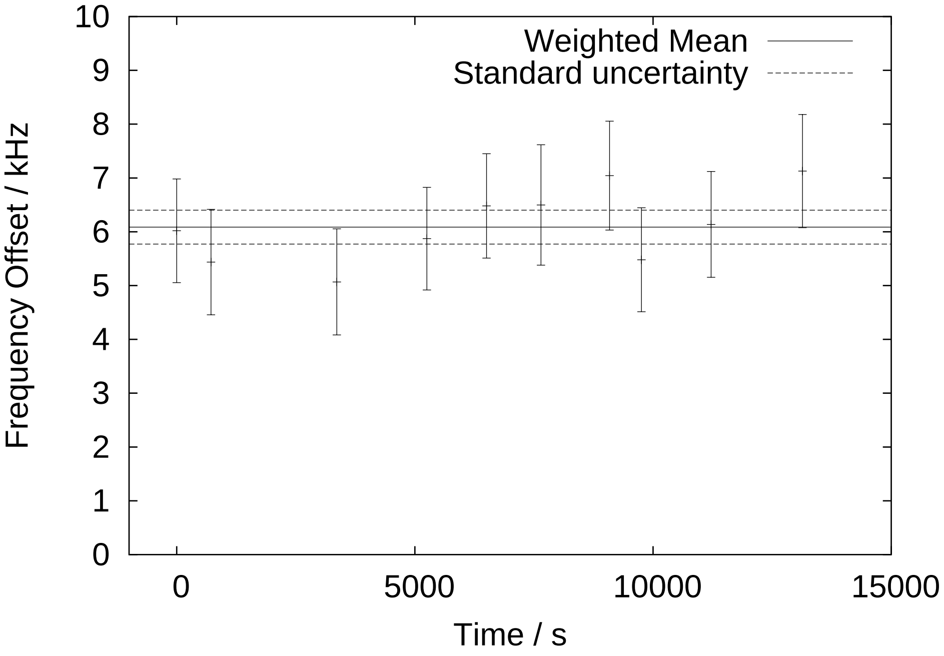

Figure 5.1: Absolute frequency measurements made on 5 September 2001. The frequency axis shows the offset from 642 121 496 770 kHz. Each point represents the mean and corresponding uncertainty of a single scan. The weighted mean and standard uncertainty of all the data are also indicated. The data are plotted against the time of the mid-point of each scan, as a check on systematic variations through the day.

The peak intensity I in the beam is given by

$$I = \frac{2P}{\pi w_x w_y} \tag{5.1}$$

So for our parameters $I = 9.7 \pm 2.5$ W m^{-2}, after taking into account uncertainties of the alignment of the beam on the ion. The calculated AC Stark shift is $+4.64 \pm 1.27$ kHz. The second-order Zeeman shift for the applied magnetic field is -660 ± 10 Hz.

Applying these corrections (all other shifts are on a much smaller scale than the uncertainties associated with these shifts) gives a final result for the September 2001 measurement as:

$$f_{\mathrm{Sep01}} = 642\ 121\ 496\ 772.1 \pm 1.3 \text{ kHz} \tag{5.2}$$

5.3 March 2002

The March 2002 measurement was a far more thorough measurement of the transition frequency, made over several days with better control of the systematic shifts. The original results were published in [78]. The results presented here have been re-analysed, and the significant differences in the analysis will be discussed later.

The major difference between this measurement (and all subsequent measurements) and the September 2001 experiment is in the calibration of the AC Stark shift. The uncertainty in the coefficient of the shift (largely due to the measurement of the beam size at the time of calibration), combined with the uncertainty in measuring the beam size during the experiment leads to a very large uncertainty in the AC Stark shift correction.

This measurement bypasses the uncertainties in the beam size measurement by measuring the transition frequency for a range of probe beam powers and extrapolating to zero power. If the variation of the beam size is smaller than the uncertainty in its measurement (which it is) then this method provides a more accurate indication of the AC Stark shift-corrected frequency.

The parameters of the measurement are shown in Table 5.2. They are

Date	13–15 March 2002
Probe laser linewidth	~ 4.5 kHz
Focusing	Tight focus (5 μm)
Magnetic field	617 ± 4 μT
Software	Visual Basic
Mechanical Shutters	467 nm beam only
AOM Switching	369 nm, 638 nm
Not switched	935 nm

Table 5.2: Experimental conditions for the March 2002 frequency measurements.

essentially the same as for the September 2001 measurement. The synchronisation with the femtosecond comb data was improved. The time was recorded at the beginning and end of each scan. The comb frequency was interpolated for each frequency bin of the scan by linear interpolation. This is not perfect, as the scan is not a constant speed in frequency, but takes longer the more quantum jumps occur. The linear approximation is good to better than 1 minute, however, and this makes very little difference on the scale of the uncertainties in this measurement.

The femtosecond comb measurements of the ULE cavity frequency are shown in Figure 5.2. There is a day-to-day variation due to the isothermal long-term drift of the cavity, plus a daily variation due to the temperature stability of the laboratory and the control system. The data from 14 March is shown, expanded, in Figure 5.3. This plot illustrates the interpolation used to estimate the cavity frequency for each quantum jump observed.

Some typical scans of the octupole transition are shown in Figure 5.4. Dependent on probe laser power and pulse length, there are around 10 quantum jumps in a scan. With such a low number of jumps, there is no reason to expect a ‘nice’ lineshape in every scan. To estimate the linewidth of the laser better, the scans must be accumulated. This has been done in Figure 5.5. The best estimate of the FWHM linewidth from this data is 4.9 kHz, assuming a Gaussian profile ($1\sigma = 2$ kHz).

Combining the interpolated femtosecond comb data with the scan data gives the absolute frequency measurements shown as a function of probe laser

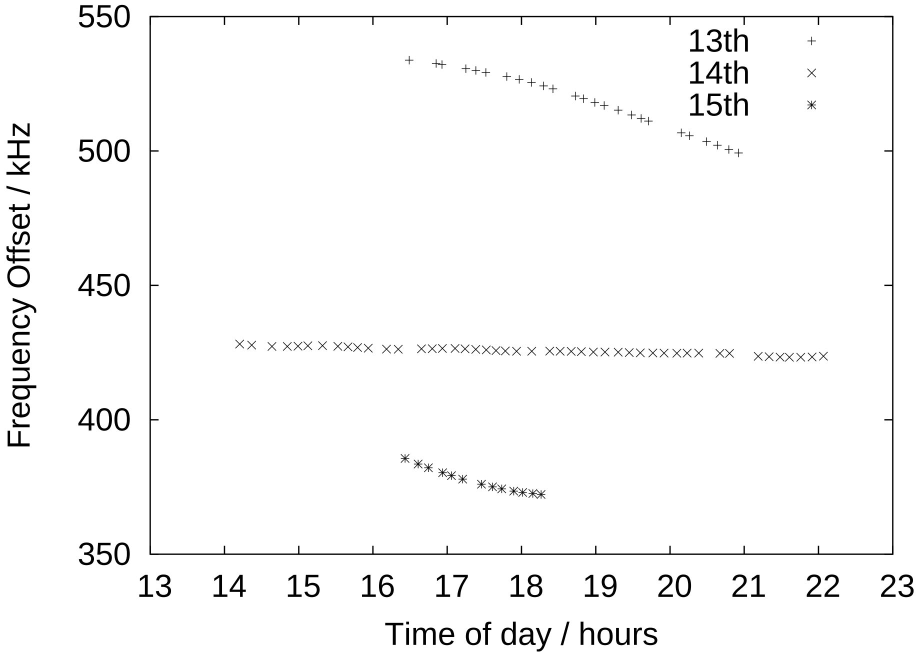

Figure 5.2: Femtosecond comb measurement data for 13,14, and 15 March 2002. There is a large drift from day to day. The linear drift rates for the three days are roughly -2 Hz, -0.15 Hz, and -2 Hz respectively. A magnified plot of the 14 March cavity frequency data is shown in Figure 5.3.

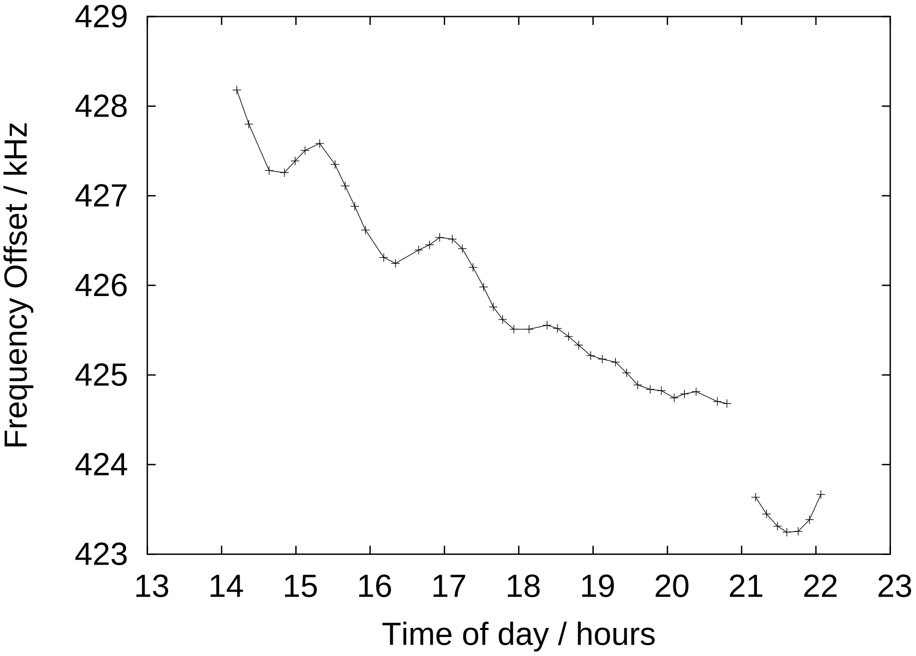

Figure 5.3: Femtosecond comb measurements of the ULE cavity frequency on 14 March 2002. Oscillations can be seen on the cavity frequency, but the overall drift rate is close to the isothermal drift of the cavity. The break in the data at 21 hours is due to the laser losing lock and the lock offset being adjusted. The line shows the interpolated frequency used to estimate the comb frequency for each frequency bin in a scan of the ion.

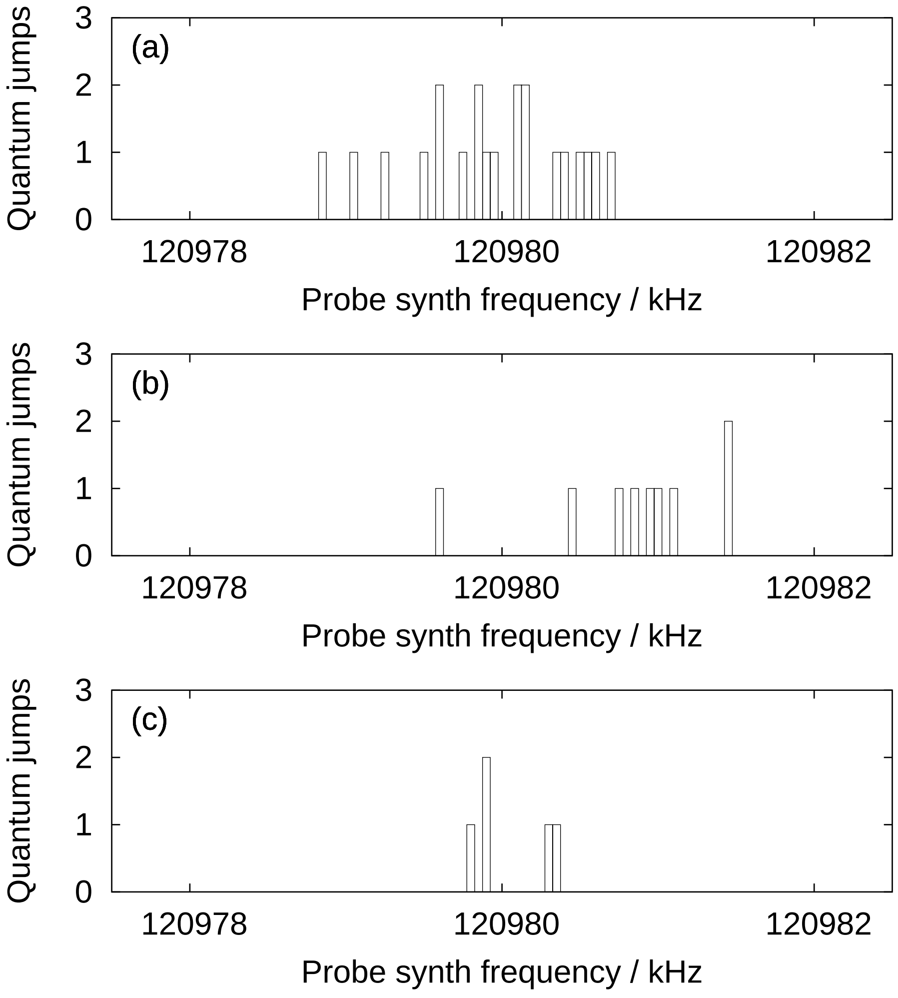

Figure 5.4: Sample data sets from 14 March 2002. (a) Scan ‘e’ — a 4 mW data point. (b) Scan ‘y’ — a 2 mW data point. (c) Scan ‘k’ - a 1 mW data point.

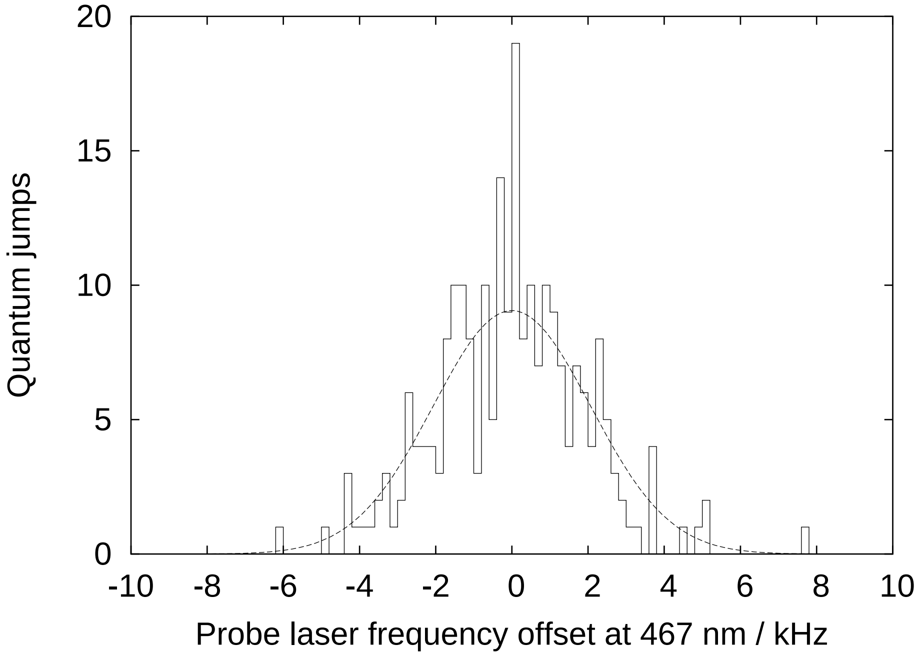

Figure 5.5: Accumulated spectrum of the 467 nm transition, using all data (235 quantum jumps) taken on 14 March 2002. The frequency axis represents the offset of each quantum jump from the centre of its scan. The scans are accumulated by aligning their centres, rather than by correcting for the AC Stark shift. This spectrum therefore reflects the probe laser linewidth accurately, but neglects systematic shifts. Assuming a Gaussian profile (shown), the best estimate of the FWHM laser linewidth is 4.9 kHz.

power in Figure 5.6. The frequency uncertainties are simply the statistical uncertainties on the means of the scans (after combination with the comb data). The x-error bars are our best estimates of the uncertainties arising from focussing, alignment, and laser power variations. Most of these uncertainties are specific to the alignment process. Because of this, the data shown in Figure 5.6 are grouped by alignment - all scans of the same power, between re-alignments of the probe beam, are combined, and estimated alignment uncertainties are attributed to the grouped points. A weighted least-squares fit to the data, as a function of power, provides a best estimate of the unperturbed (by the AC Stark shift) transition frequency. This is 642 121 496 771.02 $\pm$ 0.28 kHz. This number then needs correcting for other systematic shifts. The only significant systematic shift is the second order Zeeman shift. An uncertainty (although not a correction to the measured frequency) also comes from the accuracy with which the frequency of the hydrogen maser used as the reference frequency for these measurements is related to the caesium primary standard. This uncertainty is 16 Hz on the measured optical frequency. All other shifts and uncertainties (the quadrupole shift, AC Stark shifts from 369 nm and 935 nm light, DC Stark shifts and blackbody shifts) are estimated to be below 1 Hz [73]. The second-order Zeeman shift is calculated to be -660 ± 10 Hz. Combining these with the measurement data, the absolute frequency of the ytterbium octupole transition, as measured in March 2002 is

$$f_{\mathrm{Mar02}} = 642\ 121\ 496\ 771.68 \pm 0.28\ \mathrm{kHz} \tag{5.3}$$

This number disagrees with the published number [78] by about 400 Hz, and has a slightly larger uncertainty. The reason for the discrepancy in the measured frequency is an error in the published paper. An incorrect calibration of the magnetic field coils was used. This can be traced back to an earlier paper [67] which presented two numbers (one in the text, one in a figure caption) for the calibration of our magnetic field coils – 153 μT A^{-1} and 308.5 μT A^{-1}. The difference between the two is that one is based on an incorrect calculation of the g-factor for the upper state of the ^{2}S$_{1/2}$(F=0

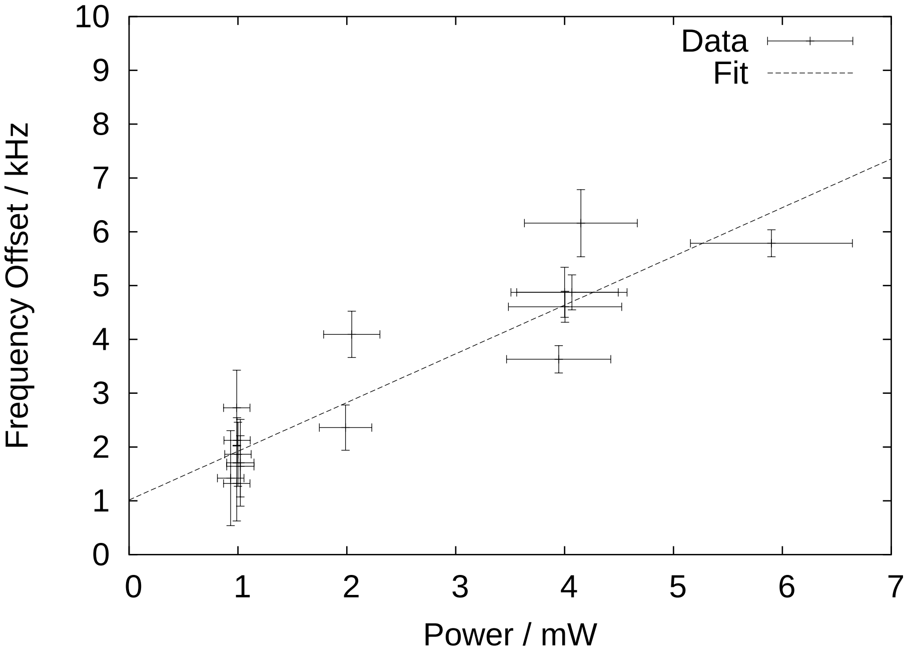

Figure 5.6: Measurement data for the 13–15 March 2002 scans. The frequency axis shows the offset of the measured absolute frequency from 642 121 496 770 kHz. The error bars show the statistical frequency uncertainty and estimated power uncertainty. The straight line is a weighted least-squares fit taking both power and frequency uncertainties into account. Data are grouped by probe beam alignment and power.

→ F=1) transition used in that paper to calibrate the field. The correct number is the larger one – 308.5 μT A^{-1} – while the smaller number was used in error in the published measurement [78]. This explains the frequency discrepancy.

The uncertainty is also larger due to the manner in which probe laser alignment errors are accounted for. In the published measurement, the estimated uncertainties were attributed to each individual scan, and the data set was then fitted. For the number presented here, scans between re-alignments of the probe beam were grouped, and then the power uncertainties were attributed. It was felt that this gave a more realistic estimate of the real uncertainties in the measurement. The weighted fitting procedure effectively 'averages down' the power uncertainties, and so, with less points to average over, the newer measurement gives a slightly larger uncertainty on the final number - 280 Hz as opposed to 230 Hz.

It is also worthwhile checking that this uncertainty is reasonable. The χ^2/ν value of the fit (ν is the number of degrees of freedom) gives the ratio of the residuals of the fitting procedure (the distance of points from the line) to their attributed uncertainties. For purely statistical data, the expected value is $\chi^2/\nu = 1$. For these data, $\chi^2/\nu = 1.44$. For 14 degrees of freedom the probability, $P(\chi^2/\nu > 1.44) = 0.12$. That is, the probability of observing a reduced χ^2 value this bad or worse, due to random statistical variation with standard deviations equal to those we have attributed to the data points, is approximately 1 in 8. This is not strong enough evidence to say that we have under-estimated any of the errors in the measurement, nor is it strong evidence of any non-statistical variation of the transition frequency. To err on the side of safety, the uncertainties may be scaled by the square root of the reduced chi-squared, giving a final uncertainty estimate of 340 Hz.

5.4 May 2003

A number of changes were made to the experiment between the March 2002 measurement period and the May 2003 measurements. The new LabVIEW experimental control system was used both to control the experiment and

Date	May2003
Probe laser linewidth	$\sim$ 150 Hz
Focusing	Defocused ($\sim$ 10 μm)
Magnetic field	188 ± 4 μT
Software	LabVIEW
Mechanical Shutters	467 nm, 369 nm, and 935 nm beams
AOM Switching	369 nm, 638 nm
Not switched	None

Table 5.3: Experimental conditions for the May 2003 frequency measurements.

acquire data. Mechanical shutters were added to the 369 nm and 935 nm beams, to avoid any possibility of leakage light from AOMs giving rise to AC Stark shifts. A lower magnetic field was used, reducing the second order Zeeman shift, and its uncertainty.

There was considerable development work on the 467 nm probe laser system between the measurements. The largest changes were that the ULE reference cavity was mounted on a vibration isolation platform, surrounded by an acoustically isolating enclosure. This reduced the laser linewidth, as observed on scans of the octupole transition, to approximately 150 Hz (See Section 3.6 and Figure 3.8). This reduction in linewidth, with corresponding increase in spectral intensity, allowed a reduction in the 467 nm intensity while keeping the same transition rate. This reduced the AC Stark shifts by a factor of about 10, in line with the reduction in linewidth. To reduce the intensity, the beam was defocused 1–2 Rayleigh ranges, increasing the spot size on the ion rather than reducing the power in the beam. This gave a fractional decrease in alignment uncertainties, and hence intensity uncertainties.

Several types of measurement were made in the May 2003 experimental period. The first of these was frequency measurement data with a defocused probe laser beam, and this represents an attempt to reach the highest possible frequency accuracy. Consistency checks were also carried out by operating the experiment in a variety of different conditions, and a calibration of the second-order Zeeman shift was performed.

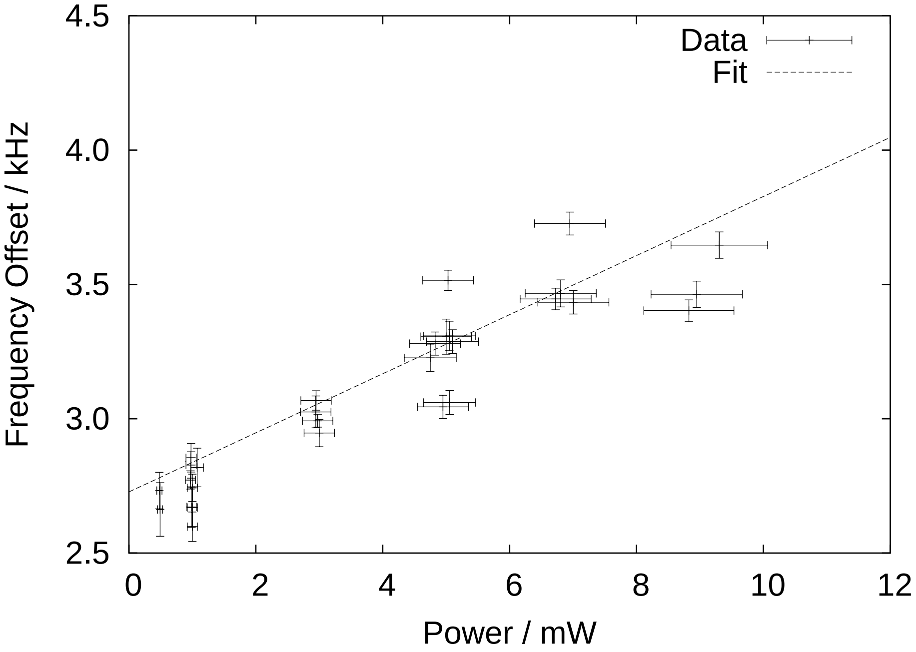

Figure 5.7: Measurement data for 7,8 and 15 May 2003, with a defocused probe laser beam. The frequency axis shows the offset of the measured absolute frequency from 642 121 496 770 kHz

5.4.1 Defocused probe beam data

The data with the defocused probe beam were taken on 7,8, and 15 May 2003, and are shown in Figure 5.7. The measurement procedure was the same as for the March 2002 measurement. The slope of the graph is much smaller than for the March 2002 data, due to the defocusing.

A weighted least-squares fit, taking into account alignment, power, and statistical frequency errors gives an intercept of (2.728 ± 0.041) kHz. However, the reduced chi-square value $\chi^2/\nu = 3.4$ for 27 degrees of freedom, which is vanishingly improbable (10^{-8}). As is clear from Figure 5.7, the distribution of data points is inconsistent with the uncertainties ascribed to them. Figure 5.8 shows the same data, corrected for the AC Stark shift by the slope of the least-squares fit.

The AC Stark shift corrected data does not appear to be statistically

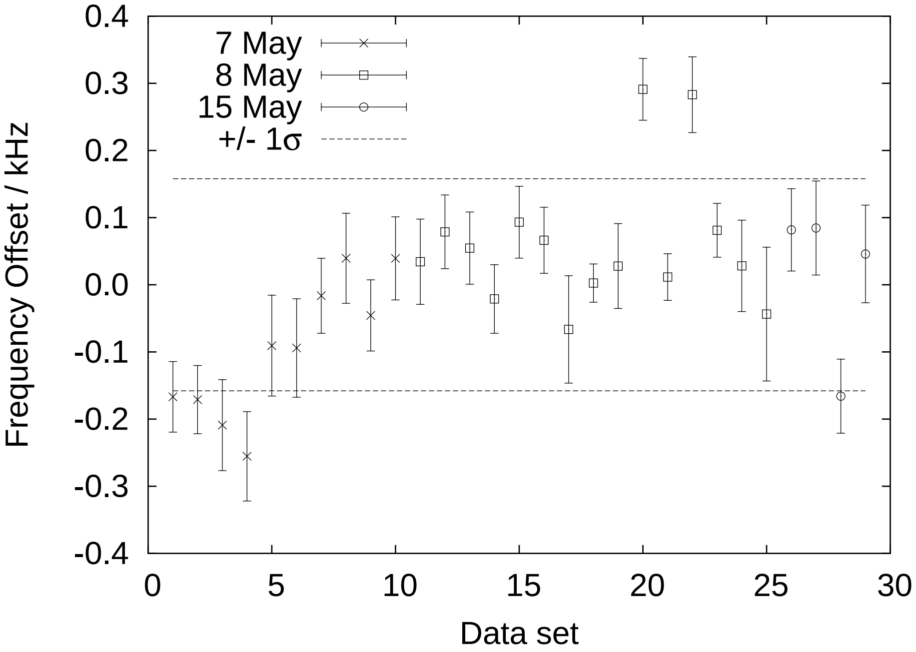

Figure 5.8: AC Stark-shift corrected measurement data. The data points do not appear to be statistically distributed. As plotted, the mean frequency is zero. The dotted lines indicate $\pm 1\sigma$ estimated using a rectangular distribution (see text).

distributed. In particular, the first few data points on 7 May and a (non-contiguous) pair of data points on 8 May appear to be substantially distant from the main body of measurements. The original data and operating conditions for these points were checked very carefully, but no reason could be found to discard them, or any other points in this data set. It would seem that some systematic shift is fluctuating and not being accounted for in our error estimates. However, no plausible source for this has been found. It takes a factor of five increase in the estimate of the intensity uncertainties to obtain a reduced chi-squared value of one. All of the data we have about the fluctuations of alignment and laser power is inconsistent with such a large uncertainty. As the statistical uncertainty on the weighted mean of 41 Hz on this measurement is clearly unreasonable, the error can be estimated by using a rectangular distribution. This is a worst-case model of the fluctuations which are causing the scatter we observe. Taking the highest and lowest data points, the data has a range of $R = 550$ Hz. If this is taken as the full range of a rectangular distribution, the standard deviation is

$$\sigma = \frac{R}{2\sqrt{3}} = 160 \text{ Hz} \tag{5.4}$$

In this situation, where statistical errors no longer dominate, it is unreasonable to use the statistically-weighted fit as an estimate of the intercept of Figure 5.7. An unweighted fit gives an intercept of 2.67 kHz.

A magnetic field of 188 ± 4 μT was used in the experiment, and so the second-order Zeeman shift correction is -60.79 ± 0.04 Hz. The maser reference frequency contributes an uncertainty of 16 Hz. Therefore, the best estimate of the transition frequency, as measured in May 2003 is

$$f_{\text{May03}} = 642\ 121\ 496\ 772.73 \pm 0.16 \text{ kHz} \tag{5.5}$$

This is not in good agreement with the measurement made in March 2002. There is a 2.9σ difference in the measured frequencies[1]. This led us to perform

---

[1] It was this measurement which lead to the discovery of the incorrect magnetic field calibration in March 2002 – there is a 6σ discrepancy with the published frequency [78].

experiments searching for a possible cause of the discrepancy.

5.4.2 The search for systematic effects

Many details of the experiment had changed between March 2002 and May 2003, and so these were investigated as possible causes of the discrepancy observed between the measured frequencies at those times. A large number of small tests were performed, detailed in order below, in addition to a more thorough investigation of the largest systematic shift – the AC Stark shift – which is described at the end of this section.

Possible causes of systematic shifts

Control software The control software was changed from the Visual Basic system to a new LabVIEW system between the measurements. Data was acquired using the VB system as a check on the accuracy of the new software. No change in the measured frequency was observed.

Mechanical shutters Between the experiments, mechanical shutters had been added to the 369 nm and 935 nm beams, to avoid any AC Stark shifts which may occur from light ‘leaking’ through AOMs. These were removed, but no change in the measured frequency was observed within the statistical uncertainties, as expected.

Synthesiser referencing In the March 2002 measurement, the probe beam offset synthesiser, which sets the offset of the probe laser from the ULE reference cavity, took its reference from the 10 MHz output of an off-air frequency standard. This is a device which obtains a reference frequency from the 60 kHz signal broadcast from Droitwich. The records of the stability of this signal during the 2002 measurement period were checked, and no signal outages which would affect the measurement were found. In 2003, the maser-derived signal was used as a reference for the synthesiser. This was swapped for an off-air standard, and no significant difference in measured frequency was observed. Tests were also done in which the off-air standard lost it radio signal and became unlocked. In these cases, the synthesisers reverted back to their internal references, which were accurate to a few hertz at 100 MHz.

Microwave source The source of the 12.6 GHz microwaves had been changed from a microwave synthesiser to a stable oscillator plus radiofrequency synthesiser and mixer. This was changed back, but no change in the measured frequency was observed.

638 nm F-state repumper The 638 nm repumper beam, which pumps out of the F-state, was left on during a scan, to see if this shifted the frequency significantly. It did not. The repumper transition itself is a weak transition, and the laser is far off-resonant from other transitions in the ion.

UV power meter A mains-powered UV power meter sits on the top of the trap support structure. In case magnetic fields from this were affecting the experiment, measurements were made both with and without the meter switched on. No difference was measured.

Micromotion Micromotion is associated with several systematic effects. To see if this could have caused the shift between measurements, it was deliberately made bad by adjusting the trap compensation electrodes. As expected from our calculations, no difference was measurable within the experimental resolution.

The ion If a dark ion is present in the trap along with the $^{171}\mathrm{Yb}^+$ ion, it should be possible to detect that either in the bright ion's micromotion characteristics or in its fluorescence spectrum (Doppler cooling profile), which should show evidence of crystallisation. To ensure that we had not missed a dark ion which was somehow biasing the frequency measurement, the ion we had kept trapped for several days was removed from the trap (the trap was switched off), and a new ion loaded. No change in the measured frequency was observed.

Probe beam focusing and linewidth

The largest changes to the experiment in terms of systematic shifts between March 2002 and May 2003 were the reduction in linewidth and the defocusing of the probe laser beam. In order to see if these had caused a systematic shift of the frequency, measurements were made in two configurations, otherwise identical to the configuration used for the narrow-linewidth, defocused data presented in the previous section. Firstly, the probe beam was re-focused on the ion, to a 5 μm spot. The second set of data also removed the acoustically isolating box from around the ULE reference cavity, increasing the laser linewidth to approximately 2 kHz. The results of these measurements are shown, along with the defocused data, in Figure 5.9. The fitted intercepts for the three data sets are (2.65 ± 0.20) kHz, (2.57 ± 0.17) kHz and (2.67 ± 0.16) kHz for the focused, narrow linewidth data, the focused, broad linewidth data and the defocused data, respectively. There is no measurable difference in frequency between the different sets of experimental conditions within the limits of the measurement accuracy. The discrepancy with the March 2002 measurement remains.

5.4.3 Second-order Zeeman shift calibration

A new calibration of the second-order Zeeman shift was performed during the May 2003 measurement period, both to check that previous measurements [67] were correct and to obtain a more accurate value of the coefficient of the shift. The data are shown in Figure 5.10. The least-squares fit to the measurement data gives a coefficient of the second-order Zeeman shift $C_{\mathrm{Zeeman}} = -1.72 \pm 0.03$ mHz$(\mu\mathrm{T})^{-2}$. This is in good agreement with the previously measured value of -1.9 ± 0.3 mHz$(\mu\mathrm{T})^{-2}$. The new value has been used for all calculations of the second-order Zeeman shift in this thesis.

5.5 July 2003

Another measurement was made in July 2003 in basically identical conditions to during May of that year. The measurement data is shown in Figure 5.11.

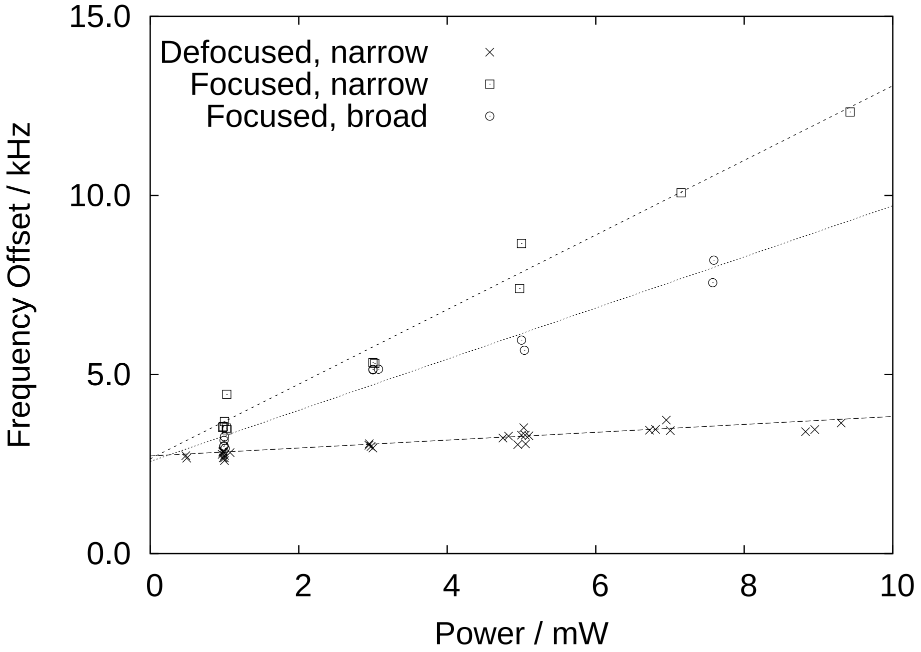

Figure 5.9: Comparison of three measurement data sets. Firstly, with a defocused probe beam and narrow linewidth probe laser, then with a focused probe beam and narrow linewidth laser, and finally with a focused probe beam and broad linewidth probe laser. Error bars are omitted for clarity, but all the points are reasonably consistent with their estimated uncertainties. The dotted lines indicate the weighted least-squares fits for the three data sets.

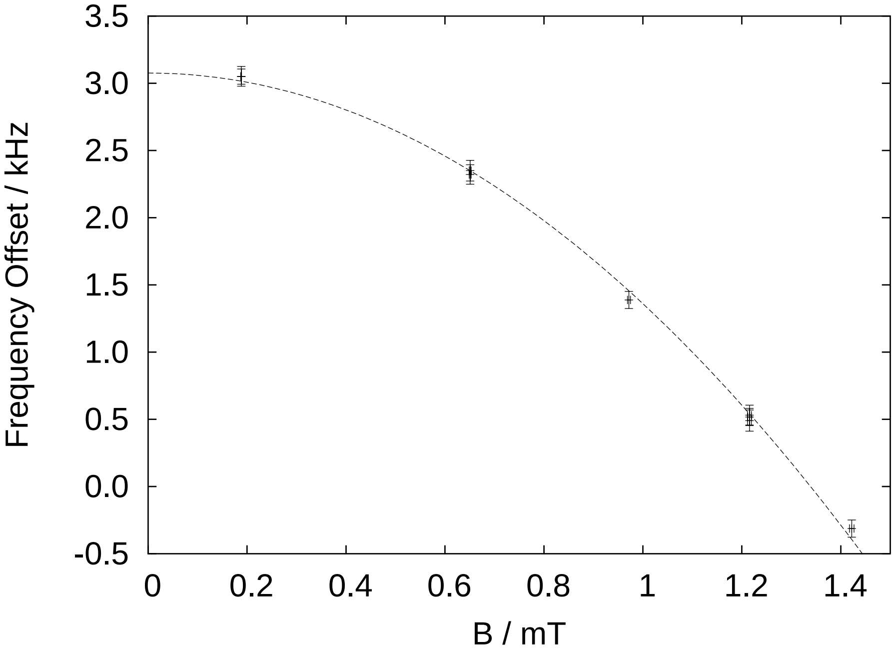

Figure 5.10: Second-order Zeeman shift of the octupole transition, and least-squares fit.

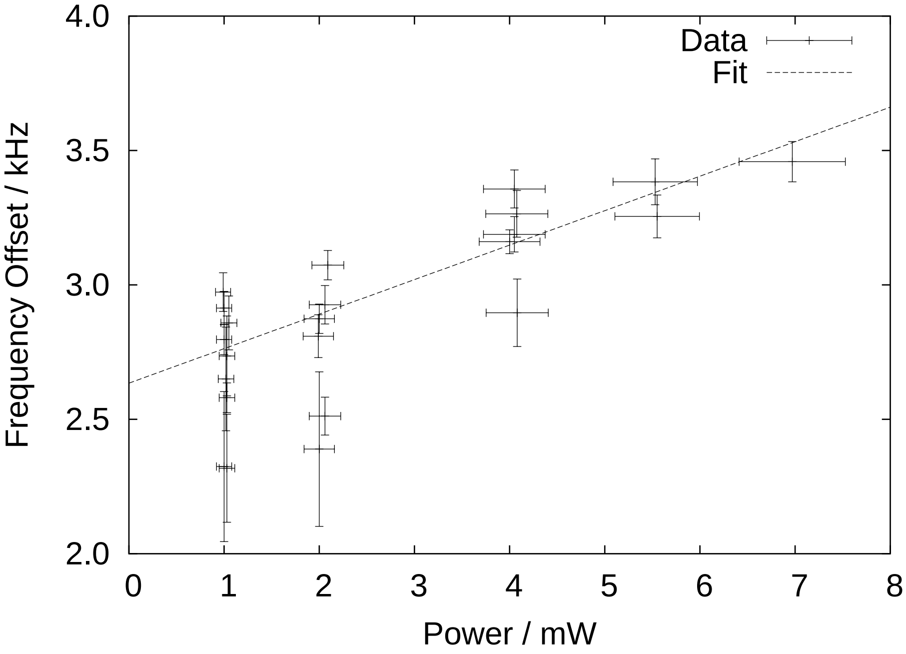

Figure 5.11: Frequency measurement data and weighted least-squares fit for July 2003.

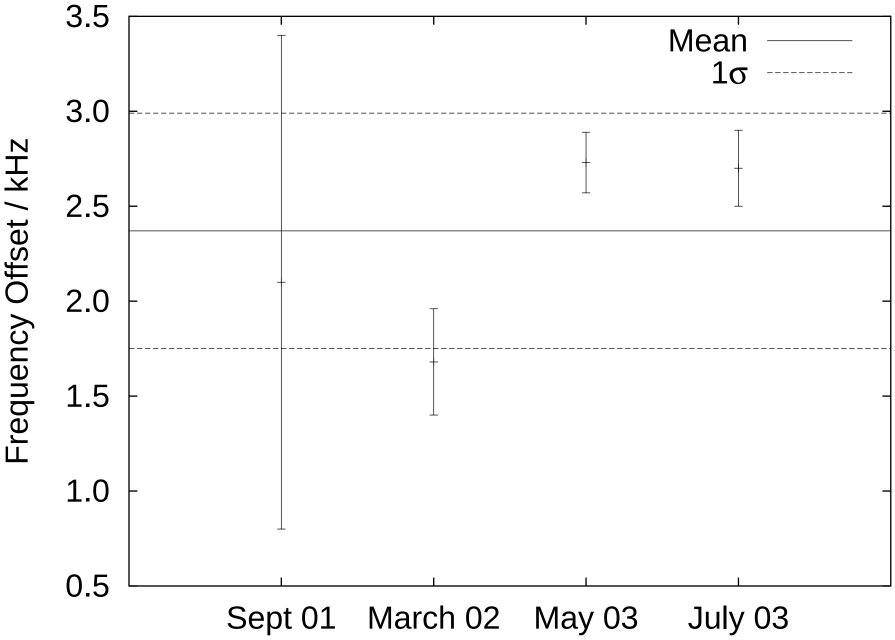

Figure 5.12: Summary of all femtosecond comb measurements of the frequency of the octupole transition. The unweighted mean, and standard deviation estimated from the spread of the points are shown.

A weighted least-squares fit gives an intercept of 2.634 ± 0.031 kHz frequency offset. The distribution of the points is again non-Gaussian, and so the uncertainty calculated from the error estimates is not reliable. Taking a rectangular distribution over the AC Stark-shift corrected data, a more reasonable uncertainty estimate is 200 Hz. The magnetic field was the same as in May 2003, and so the correction remains 61 Hz. Combining these numbers gives the best estimate for the frequency in July 2003 as

$$f_{\mathrm{Jul03}} = 642\ 121\ 496\ 772.70 \pm 0.20\ \mathrm{kHz} \tag{5.6}$$

5.6 Summary of measurements

A summary of the measurements and their uncertainties is shown in Figure 5.12. As the situation stands, the March 2002 measurement is uncomfortably distant from subsequent measurements. Though there was a thorough search for a possible cause of this discrepancy, none could be found. It remains possible that a combination of a large number of small effects conspired to produce the difference. In the absence of knowing the source of the discrepancy, the best estimate of the transition frequency can be found from an unweighted mean of the three most recent measurements. Doing this, and estimating the standard deviation of the underlying distribution by multiplying the range of the measurements by an appropriate weighting factor[2] [79] gives a best estimate of the $^2S_{1/2}$(F=0) to $^2F_{7/2}$(F=0) transition in $^{171}Yb^+$ as

$$f_{\mathrm{Jul03}} = 642\ 121\ 496\ 772.37 \pm 0.62\ \mathrm{kHz} \tag{5.7}$$

5.7 Future development

The relative uncertainty of 1×10^{-12} on the measured frequency can be compared with the uncertainties of measurements made of other trapped ion clock transition frequencies. The $^{171}Yb^+$ 435 nm quadrupole transition has been measured to an accuracy of 9×10^{-15} [14], the 282 nm quadrupole transition in $^{199}Hg^+$ to 1×10^{-14} [80], the 674 nm quadrupole transition in $^{88}Sr^+$ to 2×10^{-13} [10], and the 236 nm clock transition in $^{115}In^+$ to 1×10^{-13} [16]. The ytterbium octupole transition frequency measurement is limited in respect to these primarily by the linewidth of the probe laser.

The future direction of the experiment is therefore clear – the biggest challenge is to continue to reduce the linewidth of the probe laser system. The increased precision this offers will help to make more accurate measurements, and, if there is a real systematic problem, to diagnose it. Several projects are ongoing to reduce the linewidth of the system. The first is to fit an intra-cavity electro-optic modulator to the 934 nm Ti:sapphire laser. This should remove high-frequency noise which is beyond the bandwidth of the current AOM servo system. Improved vibration and acoustic isolation for the reference cavity can also make a very large difference, as has been observed recently in experiments at NPL with Nd:YAG lasers in acoustically quiet

---

[2] This factor is 0.591 for three data points.

environments [81].

A second ytterbium ion trap has been built, and once operational this will allow inter-trap comparisons. These are a much faster and more accurate way of determining relative systematic shifts, and should allow for very accurate calibrations of various shifts. It will also enable the reproducibility of the ion-trap environment to be tested at a high level.

To fully reap the benefits of a narrower linewidth and having two traps, it will be necessary to lock the probe laser to the octupole transition, rather than scanning across it. This will increase the data rate by an order of magnitude or more. With the new control system, performing this locking should be relatively simple.

In the longer term, it will be necessary to use the NPL caesium fountain as a reference for the femtosecond comb, rather than a hydrogen maser. The uncertainty which the maser contributes is similar to the statistical uncertainty of our best measurements. The fountain is a direct realisation of the SI second, and a much better reference for absolute frequency measurements. The ultimate aim, though is to surpass the caesium fountain in terms of absolute frequency stability and reproducibility, and possibly to become the new primary standard. These absolute frequency measurements represent progress towards that goal.

Chapter 6

EIT cooling simulation

6.1 Introduction

Ground-state cooling can bring several benefits to trapped-ion experiments (particularly frequency standards and quantum information processing systems), in addition to being an interesting and fundamental topic in its own right. The basic idea was covered in Section 2.4.5. To recap quickly, a strong coupling laser is blue-detuned from one Zeeman component of the 369 nm transition in $^{172}\mathrm{Yb}^+$. A weaker coupling laser is tuned to the same frequency above resonance on a transition to the same upper state from the other magnetic sublevel of the ground state. The strong laser modifies the absorption profile of the transition, creating a narrow resonance for the probe laser. This resonance is light-shifted from the detuning of the coupling laser by the shift induced by the coupling laser. By setting this shift equal to the vibrational frequency of the trap, the weak probe laser is resonant with the lower vibrational sideband of the narrow resonance, and can selectively excite it. This excitation, followed by a natural decay on the 369 nm transition, leads to cooling of the ion below the Doppler limit.

The modelling described in this chapter was carried out during 2001, and was based on the theory set out in reference [21]. This theory was simple, and had some shortcomings, such as being valid only in the regime where the probe is much weaker than the coupling laser. The simulations described in this chapter were motivated partly by a wish to understand the behaviour of the cooling process in regimes not covered by the theory in [21], but these are now covered by the more general theory [54] published after this modelling was complete. However, the density matrix model described here is much more general than the tightly defined conditions of EIT cooling. This allowed exploration of the behaviour of the cooling process in more physically realistic conditions, such as mis-matched laser detunings and imperfect light-shift adjustment, which are encountered in a real experiment.

6.2 Rate-equation model

Initially, a rate equation model, based on the theory in [21] was developed using a spreadsheet. This model calculates the rate coefficients for the vibrational state changing transitions using

$$A_{\pm} = \frac{\Omega_g^2}{\gamma} \frac{\gamma^2 \omega_{\mathrm{sec}}^2}{\gamma^2 \omega_{\mathrm{sec}}^2 + 4\left[\Omega_r^2/4 - \omega_{\mathrm{sec}}(\omega_{\mathrm{sec}} \mp \Delta)\right]^2}, \tag{6.1}$$

where A_+ is the rate coefficient for $|n\rangle \rightarrow |n{+}1\rangle$ transitions, A_- the coefficient for $|n\rangle \rightarrow |n-1\rangle$ transitions and γ is the linewidth of the upper state. These coefficients are derived from the standard theory of EIT cooling in [21]. The model then takes an initial distribution over vibrational states, and simply calculates the rates of transitions in and out of each level, producing a distribution over vibrational states at each point in time.

This model assumes the ideal three-level system for the cooling, and makes no attempt to include the effects of using a four (or more) level system in the experiment.

The model includes various macros which can optimise the detuning of the pair of lasers from resonance, or plot the dependence of the final $\bar{n}$ on a background heating rate included in the model. Figures 6.1 and 6.2 show the dependence of the cooling rate and final $\bar{n}$ on the detuning of the lasers from resonance.

It is also worth considering the effects of a background heating rate on

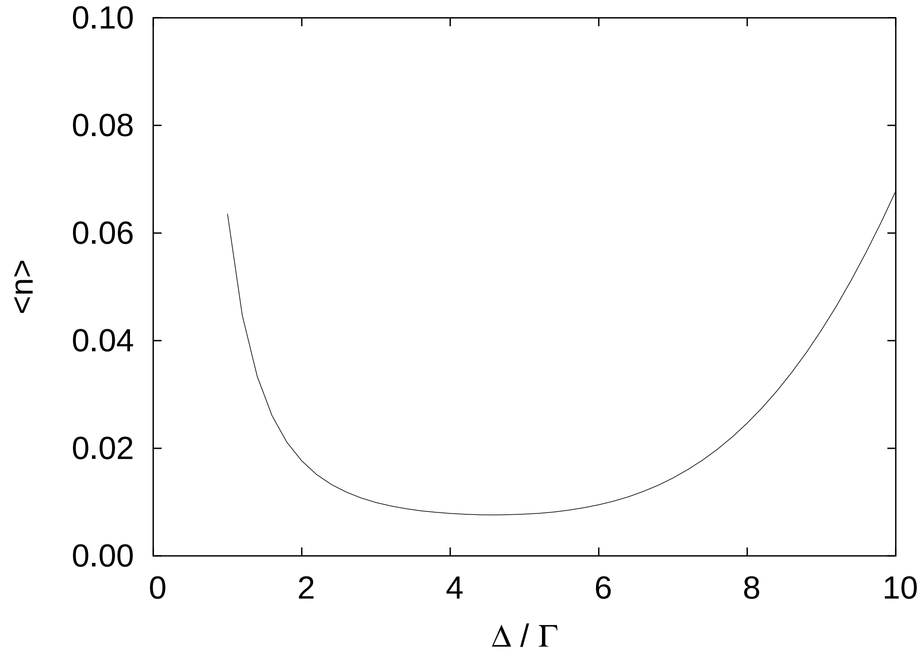

Figure 6.1: Dependence of the achieved $\bar{n}$ for EIT cooling on the common detuning of lasers from resonance. The parameters in the model were for the ytterbium system, with a Lamb-Dicke parameter of $\eta = 0.18$.

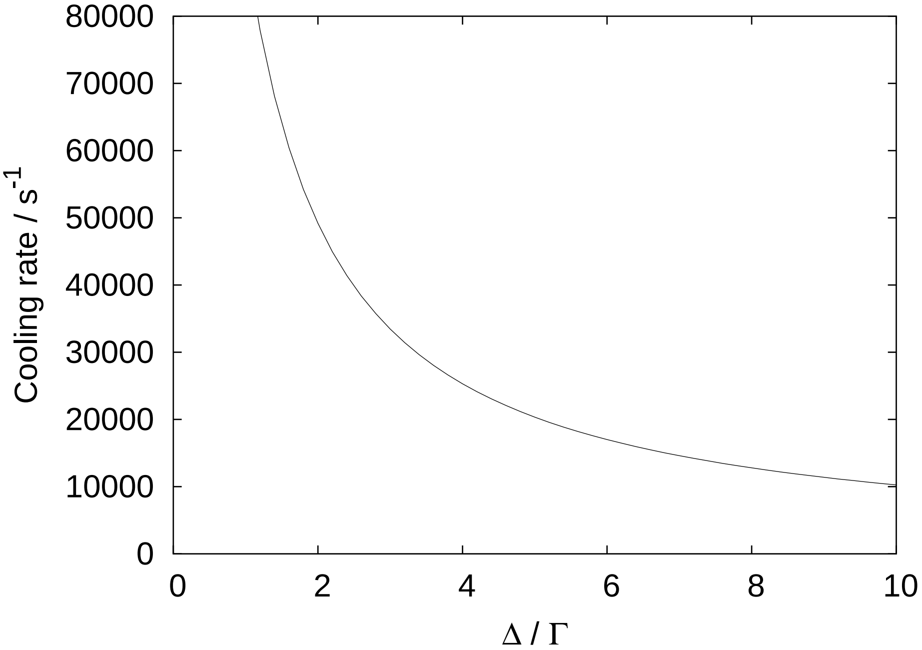

Figure 6.2: Dependence of the cooling rate for EIT cooling on the common detuning of lasers from resonance. The parameters in the model were for the ytterbium system, with a Lamb-Dicke parameter of $\eta = 0.18$. The coupling laser power was optimised for each detuning, and $\Omega_g = 0.1\Omega_r$.

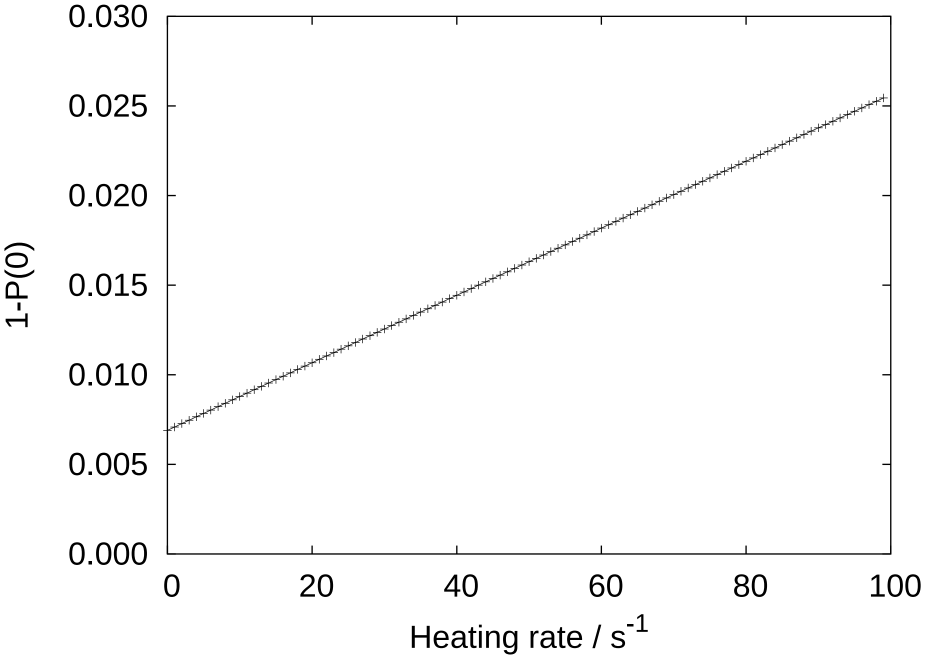

Figure 6.3: Dependence of final ground state occupancy P(0) on heating rate. The parameters in the model were for the ytterbium system, with a Lamb-Dicke parameter of $\eta = 0.18$.

the final value of $\bar{n}$ achieved. Figure 6.3 shows this. The dependence is linear, with a coefficient (in this case) of 1.9×10^{-4} s (phonons per [phonon per second]). The source of the heating is not described in any way - it is implemented simply as the rate at which ions would heat in the trap if no lasers were present. It is interesting to note that in these circumstances a ground state occupation of 99% or better can be maintained with a background heating rate of up to 20 s^{-1}. This number is larger than quoted for most ion traps of the ring or endcap type (e.g. [82, 83]), though much smaller than the heating rates measured in the NIST miniature traps. If the NPL Yb^{+} trap has similar characteristics, this implies that good ground state cooling should be possible.

6.3 Density matrix model

Following the rate equation model, a more complex, more general model of the quantum dynamics of a trapped ion with several electronic energy levels interacting with several laser fields was made.

6.3.1 Density matrices

The model represents the state of the system in a density matrix. This is a matrix which records the populations of the states of the system, and the coherences between them. The density matrix is a representation of the components of the density operator

$$\hat{\rho} = \sum_{a,b} \rho_{ab} |a\rangle\langle b|. \tag{6.2}$$

For example, if a system can take a finite set of states $|1\rangle$, $|2\rangle \ldots |n\rangle$, then for a state represented by

$$\psi = c_1|1\rangle + c_2|2\rangle + \cdots + c_n|n\rangle, \tag{6.3}$$

the density matrix is

$$\rho = \begin{pmatrix} c_1c_1^* & c_1c_2^* & \cdots & c_1c_n^* \\ c_2c_1^* & c_2c_2^* & & c_2c_n^* \\ \vdots & & \ddots & \vdots \\ c_nc_1^* & c_nc_2^* & \cdots & c_nc_n^* \end{pmatrix}. \tag{6.4}$$

This is a coherent superposition of states. For a statistical mixture of states the density matrix is simply

$$\rho = \begin{pmatrix} c_1c_1^* & 0 & \cdots & 0 \\ 0 & c_2c_2^* & & 0 \\ \vdots & & \ddots & \vdots \\ 0 & 0 & \cdots & c_nc_n^* \end{pmatrix}. \tag{6.5}$$

The diagonal elements of the density matrix represent the populations of the states $|1\rangle \ldots |n\rangle$, and so the matrix has the property that

$$\mathrm{Tr}(\rho) = 1. \tag{6.6}$$

This expresses the rule that the sum of the probabilities of being in the individual states must be unity. The matrix is also Hermitian. These properties serve as useful checks on the basic functioning of the model. If the model does not preserve these properties, it cannot be correct.

6.3.2 The time evolution of the density matrix

If the system represented by the density matrix has a Hamiltonian $\hat{H}$, the time evolution of the system is given by [55, 84]

$$\frac{\mathrm{d}\rho}{\mathrm{d}t} = -\frac{i}{\hbar}[\hat{H}, \rho] + L_{damp}(\rho), \tag{6.7}$$

where L_{damp} is a damping operator which takes account of dissipative processes in the system, such as spontaneous decay.

Thus, to build a model it is necessary to numerically integrate the above equation. The difficult part of this is not in the integration, but in finding the correct Hamiltonian for the system under consideration, and expressing it in a way which is integrable in a reasonable time.

It is also necessary to find a finite set of basis states to consider. As an atom or a harmonic oscillator has an infinite set of energy levels, it is necessary to truncate this set in a way which will not affect the dynamics of the system, under the circumstances to be considered, in any significant way.

6.3.3 The Hamiltonian

The Hamiltonian for the trapped atom may be considered as consisting of three separable parts [55, 85]:

- The internal electronic energy levels of the atom, $\hat{H}_{atom}$,

- The energy levels of the atom in the trapping potential, $\hat{H}_{trap}$,
- The energy of interaction of the atom with the incident light field, $\hat{H}_{int}$.

The internal atomic Hamiltonian

The form of H_{atom} may be written down simply. If an atom has internal energy levels $|a\rangle$, $|b\rangle$, $|c\rangle$, with energies E_a, E_b, E_c, then

$$\begin{aligned} \hat{H}_{atom}|a\rangle &= E_a|a\rangle, \\ \hat{H}_{atom}|b\rangle &= E_b|b\rangle, \\ \hat{H}_{atom}|c\rangle &= E_c|c\rangle. \end{aligned} \tag{6.8}$$

The following discussion is in terms of this three-level system, but it can be trivially extended to more states in the model. Following [21], the states of the three-level atom are labelled $|g\rangle$ (ground), $|e\rangle$ (excited) and $|r\rangle$ (metastable). The electronic Hamiltonian can then be written as

$$\hat{H}_{atom} = \begin{pmatrix} E_g & 0 & 0 \\ 0 & E_e & 0 \\ 0 & 0 & E_r \end{pmatrix}, \tag{6.9}$$

in a basis

$$|g\rangle = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \; |e\rangle = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \; |r\rangle = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}. \tag{6.10}$$

This Hamiltonian operates only on the internal states of the atom, and not on the vibrational states.

The trap Hamiltonian

The trapping potential is approximated as the harmonic oscillator pseudopotential, and its true time-varying nature is neglected. This should have little effect on the simulation as the time variation of the potential is much faster than other processes in the system. Neglecting the time variation means that micromotion is not accounted for in the model. As micromotion is minimised in all experiments before any cooling is attempted, this should have only a small effect on the cooling dynamics.

Only one dimension of the trapping potential is considered here. This simplifies the Hamiltonian, and reduces the computational requirements, but should retain most of the essential physics of the cooling processes.

If the trap has a secular frequency ω_{sec}, then the energy levels are given by

$$E_n = \left(n + \frac{1}{2}\right) \hbar\omega_{\text{sec}}. \tag{6.11}$$

This can also be written in terms of the harmonic oscillator raising and lowering operators $\hat{a}^\dagger$ and $\hat{a}$. In matrix form,

$$\hat{a} = \begin{pmatrix} 0 & \sqrt{1} & 0 & \dots \\ 0 & 0 & \sqrt{2} & \\ 0 & 0 & 0 & \ddots \\ \vdots & & & \end{pmatrix}, \tag{6.12}$$

in a purely vibrational basis. Using these operators,

$$\hat{H}_{trap} = \left(\hat{a}^\dagger\hat{a} + \frac{1}{2}\right) \hbar\omega_{\text{sec}}. \tag{6.13}$$

This Hamiltonian produces an infinite spectrum of levels, but in the model this is truncated to the lowest few levels. The higher levels should not be involved in the cooling to any significant extent if they are not initially occupied, and so can be neglected. The validity of this assumption can be checked by increasing the number of levels under consideration and comparing results. This Hamiltonian operates only on the vibrational states of the atom, and not on the electronic states.

The interaction Hamiltonian

The most important part of the Hamiltonian for the purposes of the model is the interaction Hamiltonian. This introduces coupling between the incident light field, the motional states and the internal electronic states. It is this coupling which leads to sidebands on optical transitions, and allows cooling processes to occur.

The energy of interaction of a light beam with an atomic dipole can be written

$$\hat{H} = e\mathbf{D}.\mathbf{E}, \tag{6.14}$$

where $\mathbf{D}$ is the dipole moment of the atom and $\mathbf{E}$ the electric field of the incident radiation.

The electric field has the form

$$\mathbf{E} = \mathbf{E_0}\cos(\omega t - kx). \tag{6.15}$$

It is important to include the spatial dependence of the electric field, as it is this which gives rise to the coupling between the motional and electronic states.

Between any two states, the interaction energy can be defined by a Rabi frequency Ω_{12}, where

$$\Omega_{12} = \frac{eE_0X_{12}}{\hbar}, \tag{6.16}$$

given the dipole matrix element between the two states

$$X_{12} = \langle\psi_1|\mathbf{D}.\hat{\mathbf{z}}|\psi_2\rangle, \tag{6.17}$$

and assuming the electric field acts in the z direction only [86].

The interaction Hamiltonian then takes the form

$$\hat{H} = \hbar\Omega_{12}\cos(\omega t - kx). \tag{6.18}$$

This, however, is sufficient only for a two-level system. For a multiple-level system, the Hamiltonian must also specify which states it is linking. This is achieved through use of the internal state raising and lowering operators, σ_+ and σ_-, for the two levels under consideration. Using these operators, the

Hamiltonian becomes [87]

$$\hat{H}_{int} = \frac{1}{2}\hbar\Omega_{12}(\sigma_+ + \sigma_-)\left(e^{i(kx-\omega t)} + e^{-i(kx-\omega t)}\right). \tag{6.19}$$

Here the cosine term has been expanded as a sum of exponentials. It is worth noting that the Rabi frequency used here is as defined in Loudon [86] and agrees with the value used by the Innsbruck group generally [87], but is twice the value of the Rabi frequency used by NIST in their papers. That is $\Omega = 2\Omega_{NIST}$.

The position co-ordinate $\hat{x}$ can be written as a sum of the harmonic oscillator raising and lowering operators, and hence the kx term in equation 6.19 can be written as

$$e^{ikx} = e^{i\eta\left(\hat{a}+\hat{a}^\dagger\right)}. \tag{6.20}$$

Here η is the Lamb-Dicke parameter

$$\eta = k\sqrt{\frac{\hbar}{2m\omega_{\text{sec}}}} \tag{6.21}$$

Thus we have all the ingredients for a numerical simulation of the atom-trap-laser system, treating the laser field classically. However, the interaction Hamiltonian contains terms which oscillate at optical frequencies, and is unamenable to numerical integration in a reasonable time. For this, it is necessary to transform into an *interaction picture*. How to do this is explained in a following section, but first it is necessary to explain how the basis states for the calculation are constructed.

The basis states of the model

The system under consideration has two independent degrees of freedom, and these must be united in some way for the numerical calculations. Firstly, there are the internal atomic states $|g\rangle$,$|e\rangle$ and $|r\rangle$. The density matrix for these states is a simple 3×3 matrix.

Secondly, the vibrational energy levels, truncated at n_{max} have their 'own' density matrix of size $(n_{max}+1) \times (n_{max}+1)$, independent of the internal atomic states.

In the model, these are united in a larger structure. The combined density matrix is a square of dimension $3 \times (n_{max} + 1)$ elements. The raising and lowering operators for the atomic state and for the vibrational state are modified such that they remain orthogonal when operating on the combined density matrix. The advantage of this is that the unified structure means that standard matrix operations can be used in the simulation, avoiding the need for custom numerical routines for low-level operations.

Transforming to the interaction picture

To transform to the interaction picture, a unitary operator $\hat{U}$ is applied to the *total* Hamiltonian $(\hat{H}_{atom} + \hat{H}_{trap} + \hat{H}_{int})$ and the density matrix to give

$$\hat{H}' = \hat{U}\hat{H}\hat{U}^\dagger - i\hbar\hat{U}\frac{\mathrm{d}\hat{U}^\dagger}{\mathrm{d}t} \tag{6.22}$$

$$\rho' = \hat{U}\rho\hat{U}^\dagger \tag{6.23}$$

This is a general result for any system. For the interaction picture specifically, $U = \exp(-i(\hat{H}_0/\hbar)t)$, where $\hat{H}_0$ in this case is $\hat{H}_{atom} + \hat{H}_{trap}$. All calculations are then carried out in the basis of ρ'.

After transforming to the interaction picture and making a rotating wave approximation (neglecting terms which oscillate at $\omega + \omega_0$),

$$\hat{H}'_{int} = \frac{1}{2}\hbar\Omega_{12}\left(\sigma_+ e^{i\eta(\hat{a}e^{i\omega_{\mathrm{sec}}t} + \hat{a}^\dagger e^{-i\omega_{\mathrm{sec}}t})} e^{-i\Delta t} + \sigma_- e^{-i\eta(\hat{a}e^{i\omega_{\mathrm{sec}}t} + \hat{a}^\dagger e^{-i\omega_{\mathrm{sec}}t})} e^{i\Delta t}\right). \tag{6.24}$$

Here, Δ is the detuning of the laser field from resonance. This is now the only Hamiltonian which requires integration - the H_0 part of the Hamiltonian is transformed to a zero matrix by the second term in equation 6.22. This Hamiltonian evolves no faster than the fastest of either the trap frequency ω_{sec} or the detuning. Both these frequencies are much lower than optical frequencies, and so the simulation is more manageable in this basis. It is worth noting that the harmonic oscillator raising and lowering operators acquire a time dependence in the transformation.

6.3.4 The damping operator

Having found the appropriate form for the interaction Hamiltonian, it is necessary to know how to construct the damping term in equation 6.7. The general form of this operator is the Liouville operator [55]

$$L_{damp}(\rho) = -\frac{1}{2}\sum_{m}\left[\hat{C}_m^{\dagger}\hat{C}_m\rho + \rho\hat{C}_m^{\dagger}\hat{C}_m - 2\hat{C}_m\rho\hat{C}_m^{\dagger}\right], \qquad (6.25)$$

where each m labels an independent damping process. The $\hat{C}$ operators are defined for spontaneous emission as

$$\hat{C}_{eg} = \sqrt{\Gamma_{eg}}|g\rangle\langle e|. \qquad (6.26)$$

Γ_{eg} is the natural linewidth of the decay route from level $|e\rangle$ to level $|g\rangle$. The decoherence due to finite laser linewidths may also be included in the model using

$$\hat{C}_{laser} = \sqrt{2\Gamma_{laser}}|x\rangle\langle x|, \qquad (6.27)$$

where x is state $|g\rangle$ or state $|r\rangle$ in the three-level lambda system, and Γ_{laser} is the linewidth of the laser coupling this level to another.

The damping operator is unchanged in transforming to the interaction picture.

6.3.5 Implementation

With the form of the Hamiltonian and the damping operators known, it is necessary to implement the equations of motion of the system. This was done using the MATLAB package (v4 and v6), which provides a reasonably fast and robust set of numerical methods optimised for dealing with matrices. Various integration algorithms were used, with no observed change in results, but of varying speed dependent on the situation. Generally a second/third order Runge-Kutta method was used.

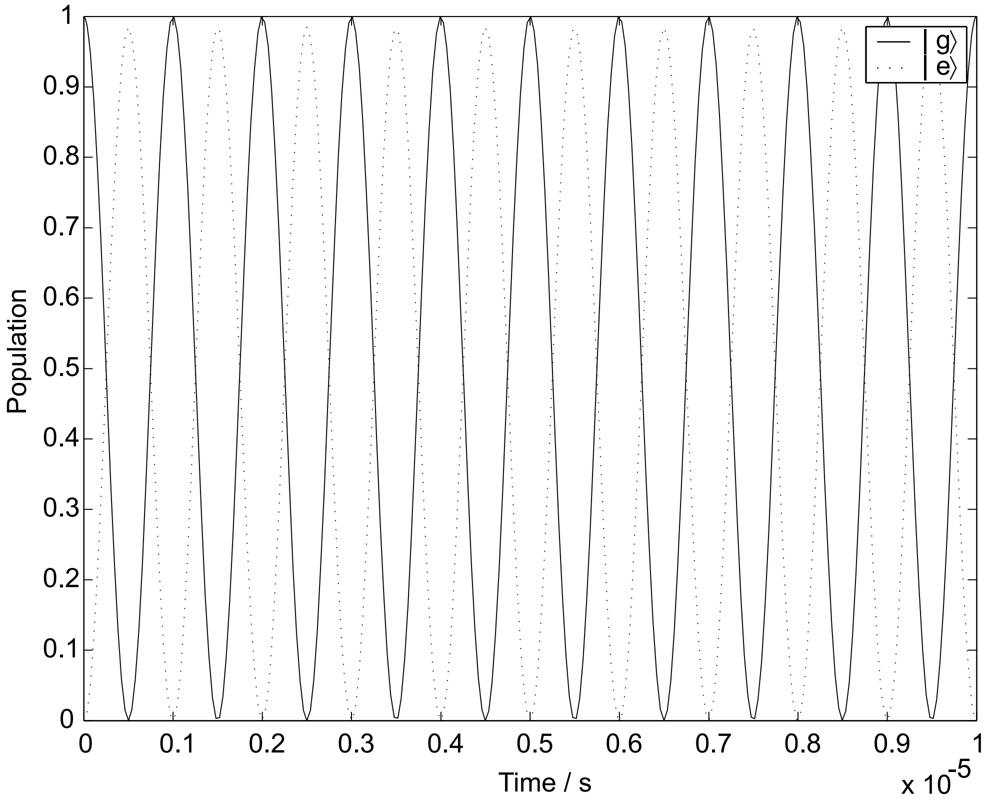

Figure 6.4: Rabi flopping. $\Omega = 10$ MHz.

6.3.6 Results

Verification

A quick test of one aspect of the basic operation of the model is provided by a Rabi flopping experiment. Figure 6.4 shows the results of tuning a single laser to resonance, in a system with no decay routes and no vibrational structure, with a Rabi frequency of $2\pi\times1$ MHz. Over 10 μs, 10 complete Rabi cycles are observed. There appears to be no mismatch between the model and theory.

A test of the interaction Hamiltonian is provided by the sideband Rabi flopping frequency. If a laser with carrier Rabi frequency Ω is tuned to the first red sideband of an atom in state $|n\rangle$, the expected Rabi flopping

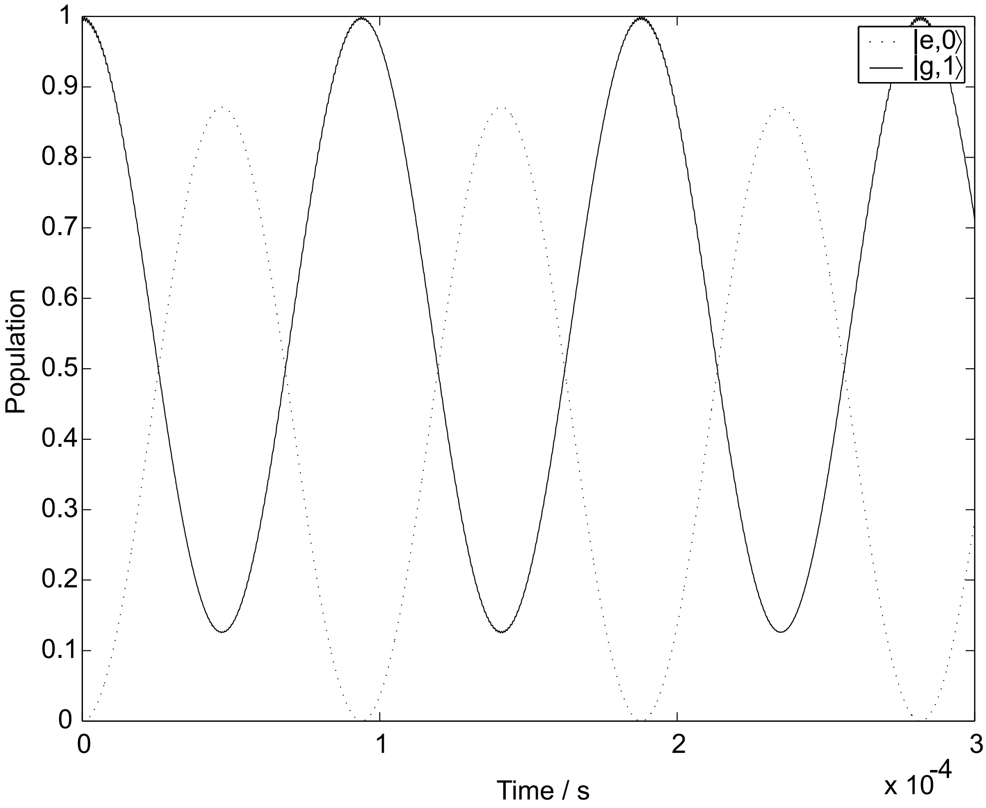

Figure 6.5: Rabi oscillations on the first red sideband.

frequency between $|n\rangle$ and $|n-1\rangle$ is [87]

$$\Omega_{n,n-1} = \eta\sqrt{n}\Omega. \tag{6.28}$$

Figure 6.5 shows Rabi oscillations on the red sideband of an atom initially in the $n = 1$ state. The parameters are $\Omega = (2\pi) \times 100$ kHz, $\Delta = -\omega_{\mathrm{sec}} = 1$ MHz, $\eta = 0.1$. Hence the sideband Rabi frequency is expected to be 10 kHz, and the atom should complete three cycles between the vibrational levels in the time span covered by the graph.

The actual flopping frequency is slightly faster than is naïvely expected, and the population transfer between vibrational levels is not complete. This behaviour is expected if the laser is not exactly on resonance. Although the detuning of the laser has been set to the ‘ideal’ position of the sideband, no allowance has been made for the AC Stark shift induced by the off-resonant interaction of the laser with the carrier, and so the laser will be slightly off-resonant with the sideband. In this respect, the model behaves as expected.

Sideband cooling

Sideband cooling calculations are theoretically more simple than EIT cooling calculations. There is only one laser and two atomic levels involved in the cooling, at a basic level. In reality, sideband cooling is performed with a laser coupling the metastable state to the ground state via an intermediate level. Following [44], this system can be approximated to a two level system with an effective linewidth for the metastable state.

It is possible to generate a spectrum of the sidebands using the model. This allows optimisation for AC Stark shifts, if they are large (which is usually not the case), and acts as a check on the validity of the model - are the sidebands in the right place? Figure 6.6 shows a vibrational evolution spectrum.

The spectrum does show that the sidebands are in their expected positions, and also confirms that the blue sideband is more intense than the red. The relationship between sideband heights in Figure 6.6 is not a simple one, as the points in the spectrum represent end-points of cooling curves, and not initial rates of heating or cooling.

EIT absorption profiles

By ignoring the effect of the trapping potential, it is possible to generate a spectrum of the EIT absorption feature in a free atom.

The lineshape shows the cancellation of absorption when the two lasers are detuned to the same point above resonance, and displays the narrow resonance peak close by. This is for the case of zero laser linewidth. To implement in a practical experiment, it is useful to know the laser linewidth which will allow the EIT condition to hold well. It is the relative laser linewidth, rather than absolute linewidth which causes decoherence. If both coupling and probe beams are derived from the same laser via some coherent method

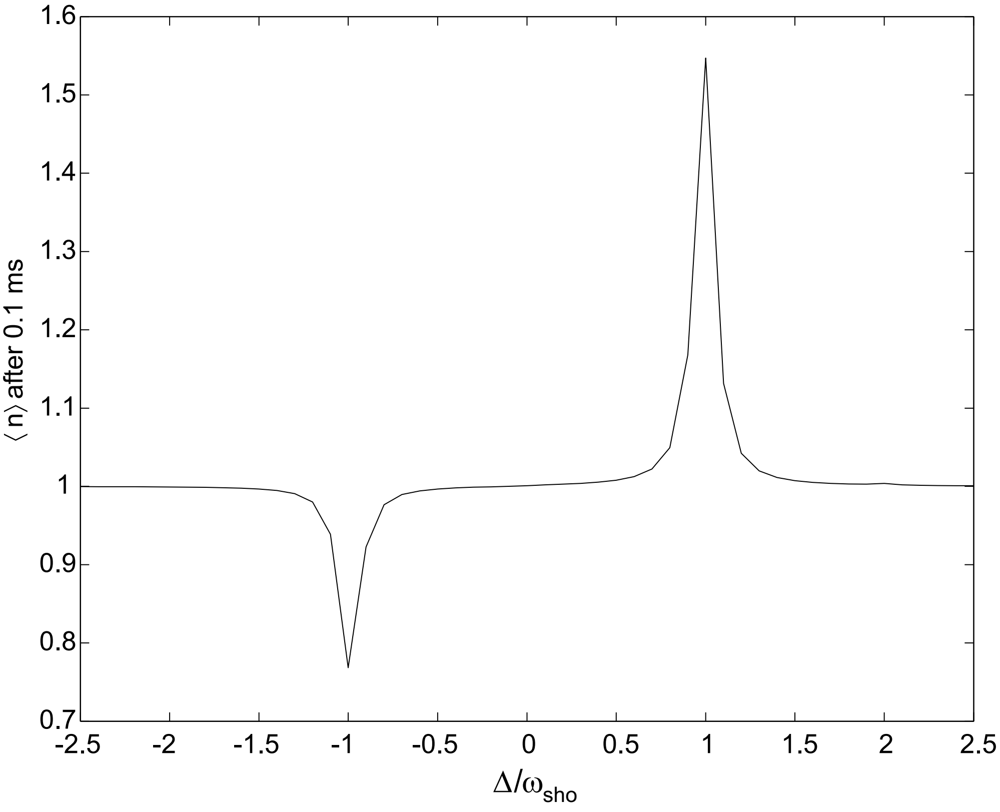

Figure 6.6: Vibrational evolution spectrum. The first-order red and blue sidebands are visible. The system was allowed to evolve for 0.1 ms at each detuning, and the value of $\bar{n}$ recorded. The trap frequency was 1.3 MHz, and the effective linewidth of the cooling transition 150 kHz.

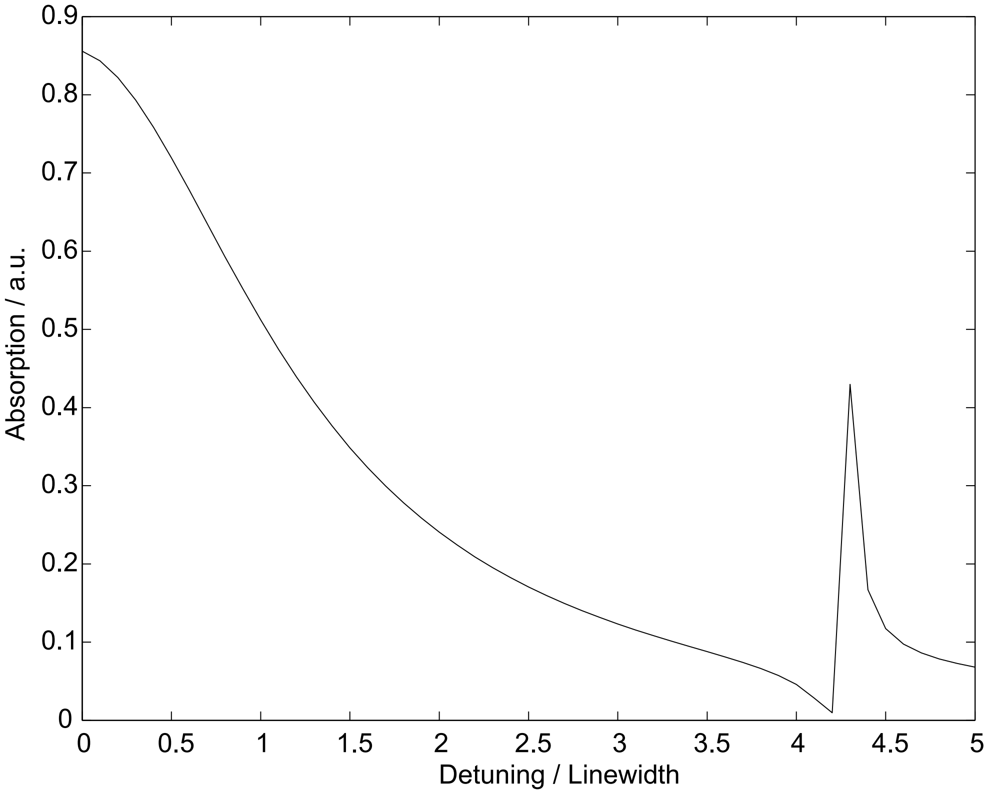

Figure 6.7: EIT absorption profile. Parameters are for $\Gamma = (2\pi) \times 20$ MHz. The detuning of the coupling laser is 4.2 linewidths.

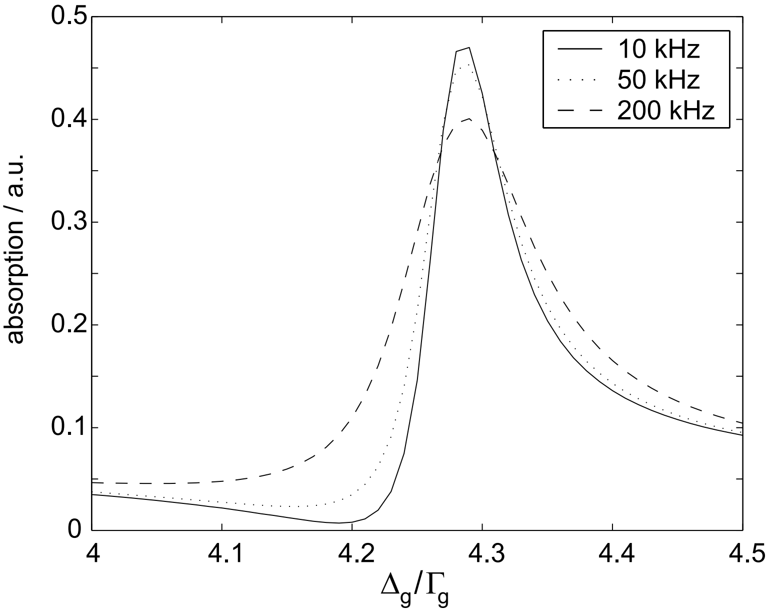

Figure 6.8: EIT absorption profiles for various laser linewidths. Both the lasers had the same linewidth, and the lasers were assumed to be independent. The detuning of the coupling laser is 4.2 atomic linewidths.

(e.g. an AOM), then this is unlikely to be an experimental issue. Figure 6.8 shows the effect of the relative laser linewidth on the absorption profile. As the linewidth increases, the cancellation of absorption disappears and the narrow resonance broadens. This situation has been studied analytically in [88].

The figure shows that 10 kHz laser linewidth has little (though some) impact on the overall shape of the absorption curve, but linewidths larger than this would swamp the first order sideband. The height of the sideband is roughly η^2 times the height of the carrier. For reasonable Lamb-Dicke parameters, the sideband would not be larger than the carrier at the point of equal laser detunings for a linewidth larger than around 10 kHz. Hence, cooling limits are not expected to be as low for larger linewidths.

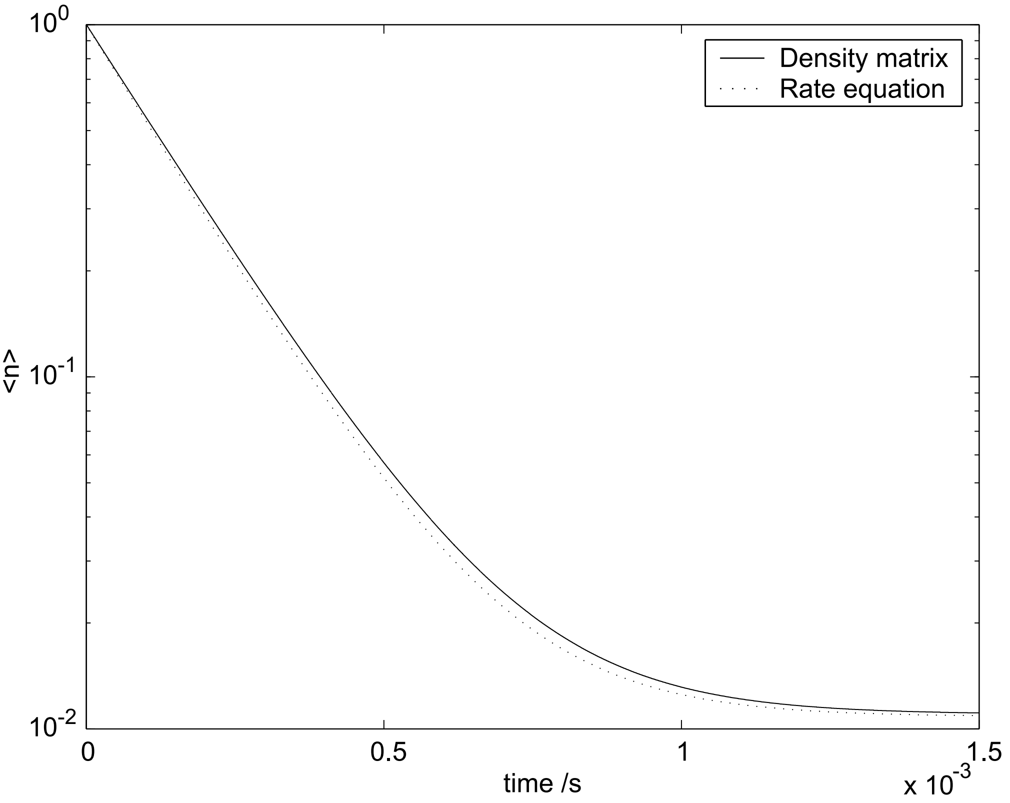

Figure 6.9: Comparison of density matrix and rate-equation models for $\Gamma = (2\pi)20\ \mathrm{MHz}, \Omega_r = \Gamma, \Omega_g = \Gamma/10, \eta = 0.145, \Delta_g = \Delta_r = 2.5\Gamma$.

EIT cooling - comparison with rate equations

The form of the Hamiltonian and method of implementing the density matrix model have been described in earlier sections. Firstly, the model should be able to reproduce the theoretical results in the original EIT theory paper [21]. Figure 6.9 shows a comparison between the rate equation model described in Section 6.2 and the density matrix model for the system described in [21]. Both systems start in a pure $n = 1$ state. The parameters are approximately those of the S-P-D system in calcium.

The agreement between the two models is very close. Much of the difference can be attributed to the first few microseconds of the cooling. The rate equation model is insensitive to the internal states of the atom, and assumes steady-state excitation values for the internal states. The density matrix model, however, is initialised in a pure $n = 1$ state of the $^2\mathrm{S}_{1/2}$ level. In the first few microseconds, the cooling is slower while coherences between levels build up. Both models predict the same final ground state occupation of 99%.

Coupling Rabi frequency

For optimal cooling, the EIT cooling scheme is dependent on the AC Stark shift induced by the coupling laser matching the vibrational frequency of the trapped ion. If it is to be implemented in experiment, it is necessary to know how stable and accurate the intensity of the coupling laser must be. Figure 6.10 shows the dependence of the cooling rate on the Rabi frequency of the coupling laser.

The fractional width of the peak is roughly 20% (FWHM). Hence the requirements on intensity stability are not severe. If both noise and drift are below the 5% level over the period of a typical experimental period — at least tens of minutes — then there should be few intensity-related problems.

Detuning from resonance

Assuming that both lasers are equally detuned from resonance ($\Delta_g = \Delta_r$), the cooling speed will depend on the size of this detuning. Figure 6.11 shows this dependence.

The cooling is strong over a relatively broad range of detunings, and optimising the detuning should not be a hard problem in practice. Heating only occurs for a small range of detuning near resonance, and so there is little risk of accidentally heating the ion. The heating occurs because as $\Delta \to 0$, the condition that the AC Stark shift should equal the trap frequency can no longer be satisfied, and the probe laser is no longer tuned to the red sideband, and in fact starts to off-resonantly excite the blue sideband.

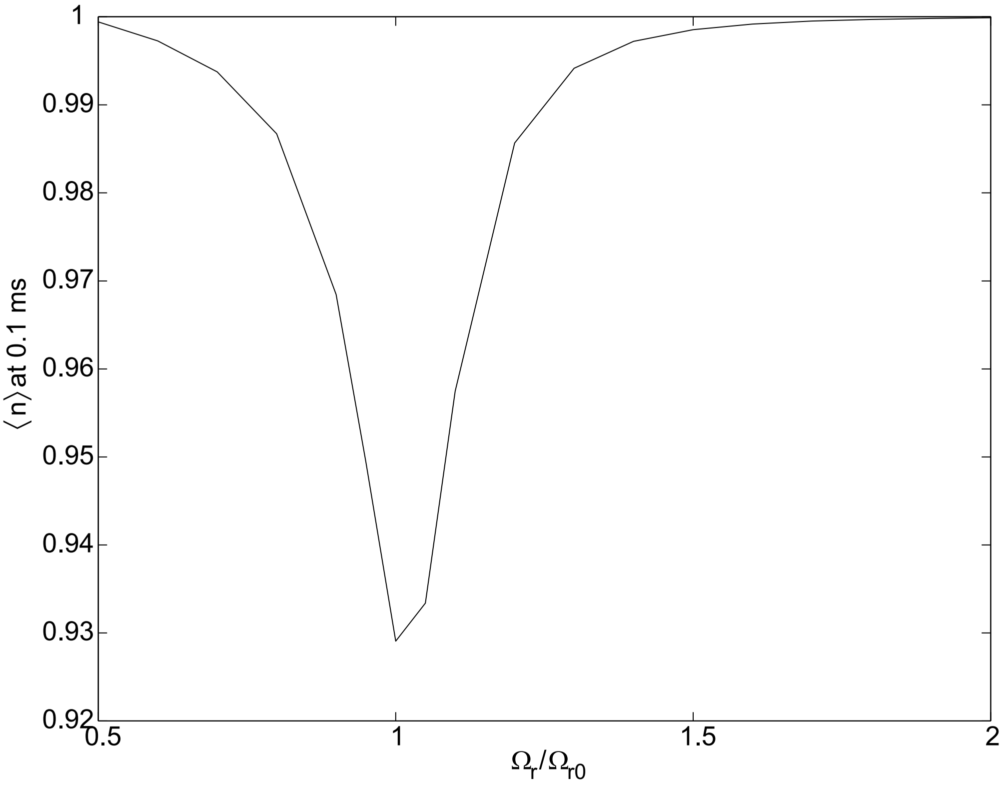

Figure 6.10: Effects of coupling laser power on EIT cooling. The system is optimised for $\Omega_r = \Omega_{r0}$.

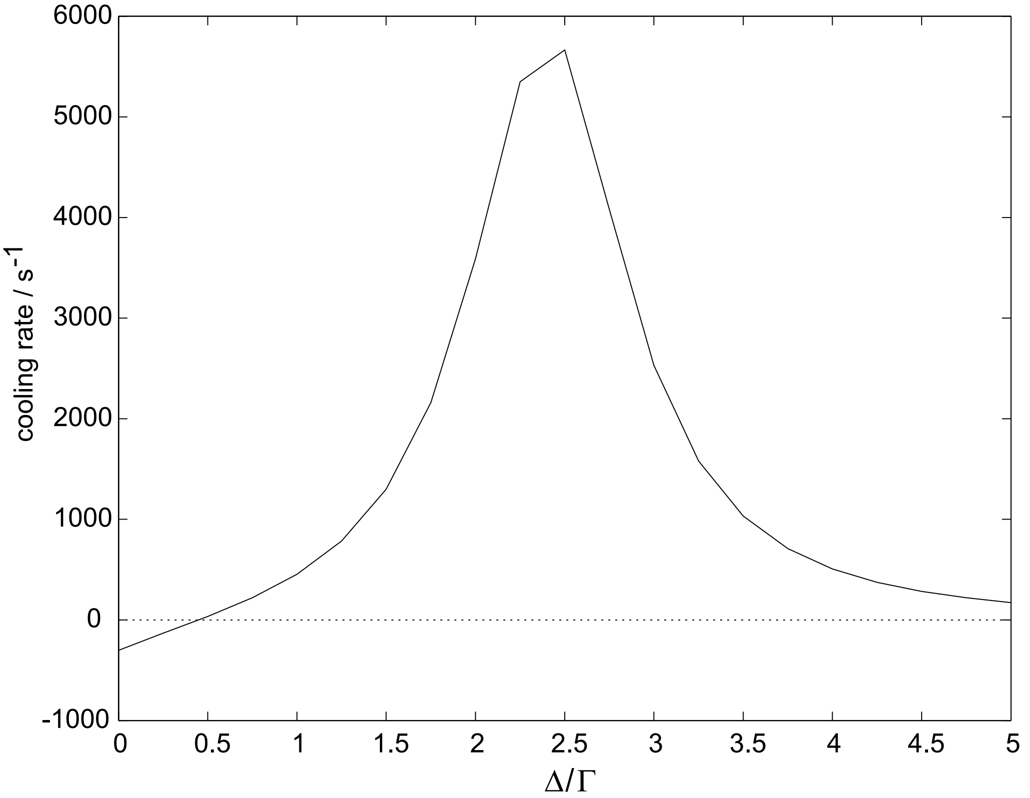

Figure 6.11: Cooling rate as a function of laser detuning from resonance, for $\Delta_g = \Delta_r$. The laser powers are optimised for $\Delta_g = \Delta_r = 2.5\Gamma$.

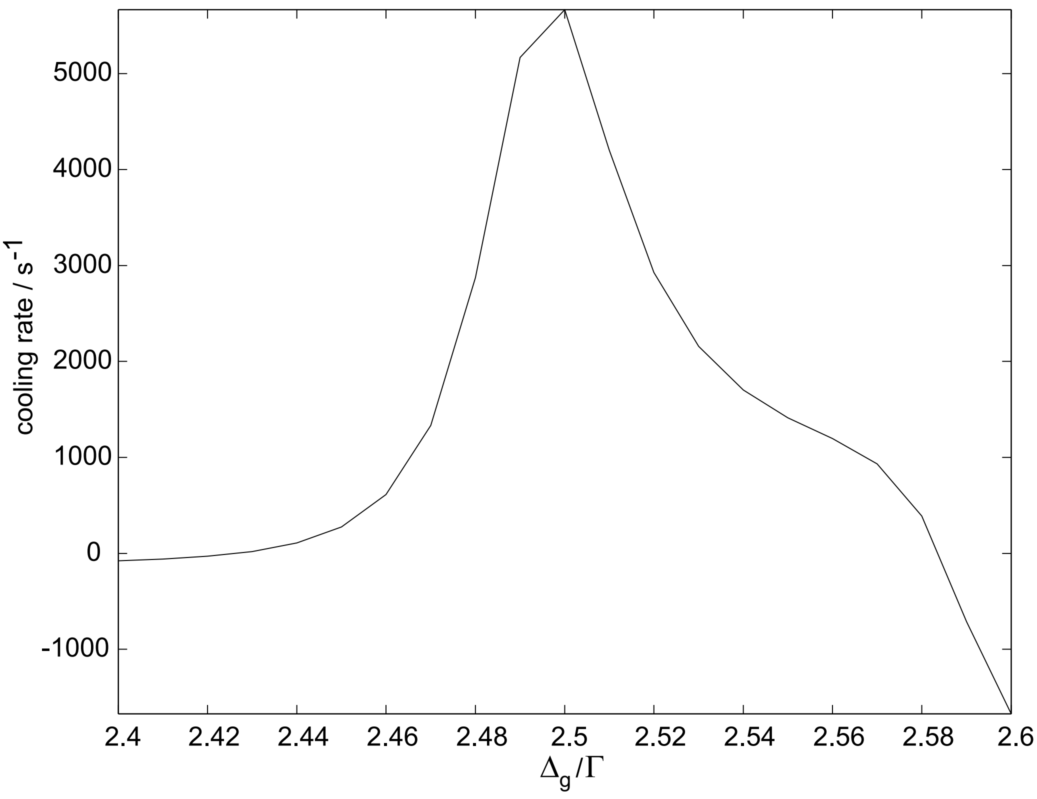

Figure 6.12: Variation of cooling rate with cooling beam frequency. The coupling laser is at $\Delta_r = 2.5\Gamma$.

Relative detuning

The accuracy of the relative detuning of the two lasers is also an issue. Figure 6.12 shows the effect on the cooling rate of scanning the probe laser relative to the coupling laser.

The control required for the relative frequency of the two beams is no greater than is already standard on most experiments. Good cooling occurs over a relatively wide region (a few hundred kHz). If the frequency is too high, the blue sideband is driven and the ion heats rapidly, but this is not a major risk. In a situation where the coupling and probe beams are derived from the same laser via an AOM, controlling the relative frequency to the desired accuracy should not pose any problem.

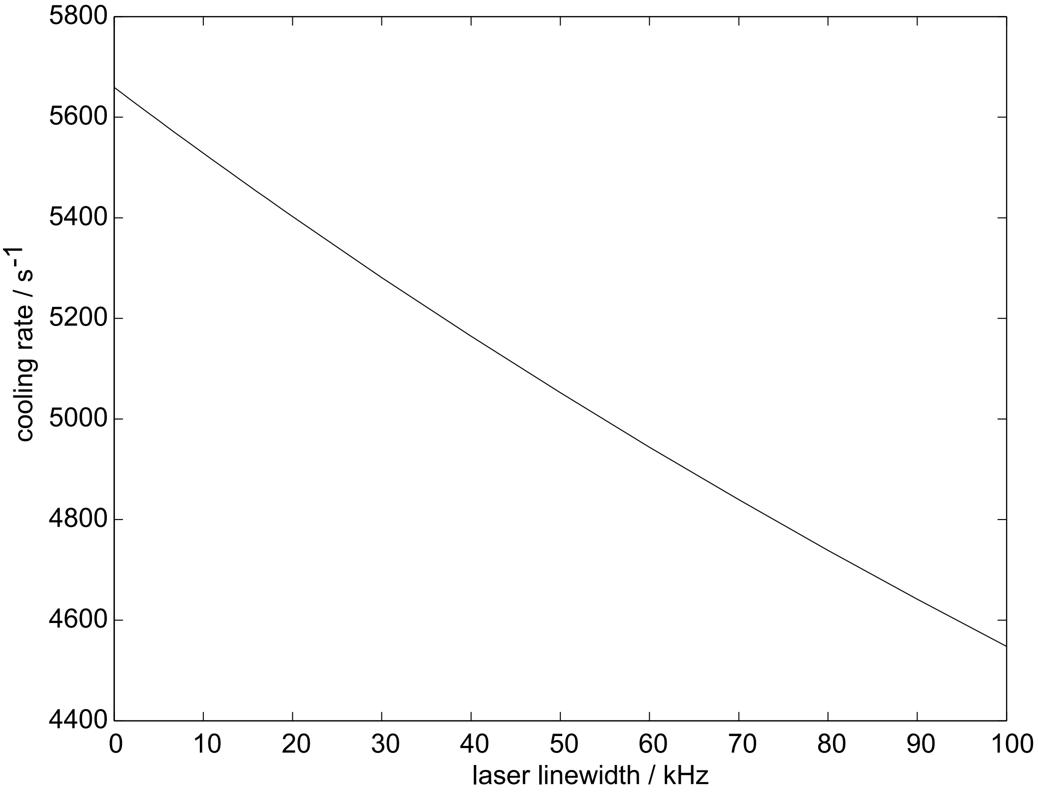

Figure 6.13: Effects of laser linewidth on EIT cooling.

Laser linewidth

For Raman cooling or sideband cooling, it is generally necessary to have a small laser linewidth (relative linewidth in the case of a Raman beam pair). In order to find out what the constraints on EIT cooling may be, a plot of cooling rate against relative linewidth was obtained from the model. This is shown in Figure 6.13. The linewidth of the probe laser was set to the value indicated on the graph, while the linewidth of the coupling laser was zero. The same results were obtained if the linewidth was transferred to the coupling laser instead of the probe.

It is interesting to note that the dependence of the cooling rate on linewidth is not very strong. The model predicts good cooling even with a linewidth of 500 kHz, which is half the secular frequency. The reason for this may well be a problem with the model. While the absorption spectrum is modified by the trap potential in the model, the emission spectrum is not, and decay only occurs on the carrier. Because of this, the heating process of absorption on the carrier and emission on the red sideband (*emission* on the red sideband heats the ion, while *absorption* on the red sideband cools) does not occur in the model. This would probably be the major source of heating as the laser linewidth was increased. Cooling limits may be affected more strongly by linewidth than initial cooling rates, but this has not been tested.

In the experiments of Chapter 7, both the coupling and probe beams are derived from the same laser via an AOM, and their relative linewidth is expected to be very narrow. It is unlikely that those experiments are affected by the relative linewidth at all.

6.3.7 Problems with the model

There are several minor problems with the model. One problem is that it does not include the effects of the trapping potential on the emission spectrum (L_{damp} is not modified by the trap potential in this model), and so most off-resonant heating effects are absent from these simulations.

Two effects which the model ignores, but are not planned to be incorporated in the near future, are the effects of making the rotating wave approximation in the derivation of the interaction Hamiltonian, and the effects of laser fields on far off-resonant levels. The latter can be included in the model by defining a set of operators linking off-resonant states via the laser fields. This is not done automatically, but is trivial and could be automated. Neither of these effects is expected to be significant in the current work. If they are deemed necessary at some time in the future, they can be included.

Another issue is the speed of the simulation. Sideband cooling is simulated more quickly than EIT cooling. Simulating EIT cooling for a period of 10^{-5} s takes approximately two-and-a-half minutes for typical parameters (a 12×12 matrix) on an 800 MHz computer.To generate a high quality spectrum of some sort requires many such runs, over much longer time periods (up to several milliseconds). The computing time required quickly becomes impractical as the problems scale. Also, as more levels become involved in the simulation, the time required to perform all the matrix operations grows quickly (as n^2 or n^3, typically). While the simulation could be optimised further, other methods of simulation may be more efficient than the straight numerical integration used here.

6.3.8 Future developments

There are obvious improvements which can be made to the model, such as including the effects of the trap on the spontaneous emission and optimising the simulation with regard to scalability.

Once this is done, there are many interesting numerical experiments which can be done with the model. Raman cooling, for example, can be studied and compared against the alternatives for the cases of specific ions. Also more realistic systems, such as the four-level $^2\mathrm{S}_{1/2}$—$^2\mathrm{P}_{1/2}$ system in $^{172}\mathrm{Yb}^+$, could be simulated. A simulation of this complete with the effects of laser polarisation could shed light on the experimental results of Chapter 7.

It is also possible to run a complete simulation of the EIT temperature measure scheme, described in Section 2.4.7. This might help us to better understand the results of the experiments of Chapter 7.

Another interesting possibility is simulating state manipulation pulse sequences. It should be possible to implement simple quantum gates in the system, and determine what the experimental requirements would be for implementation in various real systems.

Chapter 7

EIT cooling experiment

This chapter details the progress to date of experiments investigating EIT effects on the 369 nm transition in trapped $^{172}\mathrm{Yb}^{+}$. The project has several aims. EIT cooling has been demonstrated only once, in $^{40}\mathrm{Ca}^{+}$ [22], and was not investigated in great detail. It would therefore be interesting to implement the scheme again, in another experiment. It would then be possible to explore the parameter space of the cooling some more, with a view to obtaining (or understanding how to obtain) the theoretical performance levels of the scheme, which are much better than the demonstrated performance.

The experimental objectives are, in approximate sequence

- To observe EIT spectra in a single ion,
- To characterise the AC Stark shift using these spectra,
- To attempt temperature measurement using the EIT spectra,
- To compare the EIT temperature measurement with measurements using spectra of the octupole transition,
- To demonstrate EIT cooling, and
- To investigate how various parameters, including the Zeeman splitting and polarisation purity, affect the EIT cooling process.

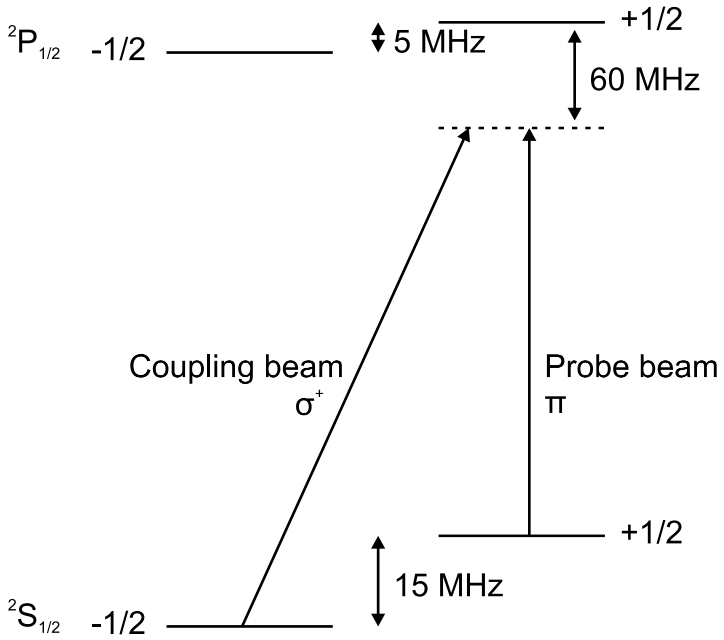

Figure 7.1: Energy levels involved in EIT experiments on the 369 nm transition in ^{172}Yb$^+$. The situation shown here is for σ^+ coupling light, red-detuned, as in the calibration scans. The level splittings shown are approximate, for the typical field strengths used.

The progress made towards these objectives is detailed in the following sections. The 467 nm octupole transition laser was unavailable during the experiments described here, and so comparison of the two temperature measurement methods has not yet been possible. Only the EIT temperature measurement technique was available to attempt a demonstration of EIT cooling.

7.1 Experimental details

The experiments are carried out on the $^2S_{1/2} \rightarrow {}^2P_{1/2}$ transition at 369 nm in ^{172}Yb$^+$. The structure of this transition, and the lasers involved are shown in Figure 7.1. A magnetic field was applied which produced a Zeeman splitting of the $^2S_{1/2}$ state of approximately 15 MHz.

7.1.1 Polarisation control

It is important for the EIT spectra that the coupling beam is purely circularly polarised. For this to be possible, the magnetic field axis must be aligned with the direction of propagation of the light. This was achieved through an iterative procedure. Firstly, the polarisation of the coupling beam was set to be approximately σ^+, by rotating a quarter-wave plate following a linear polariser. The direction of the field was then adjusted to minimise fluorescence, and the polarisation re-adjusted on the same criterion. This was repeated until a minimum fluorescence level had been found. Once the polarisation is pure, no fluorescence should be observed, because the atom will be optically pumped to the $^2\mathrm{S}_{1/2}(m_J = +1/2)$ state. The procedure did not work well until the final mirror before the trap had been replaced. Originally, this was an aluminium mirror, and the reflected polarisation was elliptical. It was replaced with a dielectric mirror, which reflects the polarisation without introducing ellipticity, allowing better alignment of the magnetic field. It also increased the size of the bright resonance by a large factor, indicating a strong dependence of the size of the EIT effect on polarisation purity. The coupling beam still passes through a quartz vacuum window after this mirror, and this could be a limiting factor on the achieved polarisation purity.

The geometry of the vacuum canister around the trap is such that there is not a right angle between any pair of optical windows. This makes it impossible to illuminate the ion with both π- and σ-polarised light simultaneously, in the atom's frame. The best purity which can be achieved for the weak probe beam is approximately 60% π-polarisation (by intensity).

7.1.2 Laser system

For these experiments, the ability to rapidly and precisely switch beams on or off, or change their frequency is essential. This required some changes to the 369 nm laser system from the setup described in Chapter 3. The modified laser setup is shown in Figure 7.2. The beams shown were switched on and off using a network of RF switches, shown in Figure 7.3. The frequency and amplitude of the source RF can be controlled via GPIB (IEEE 488.2). The

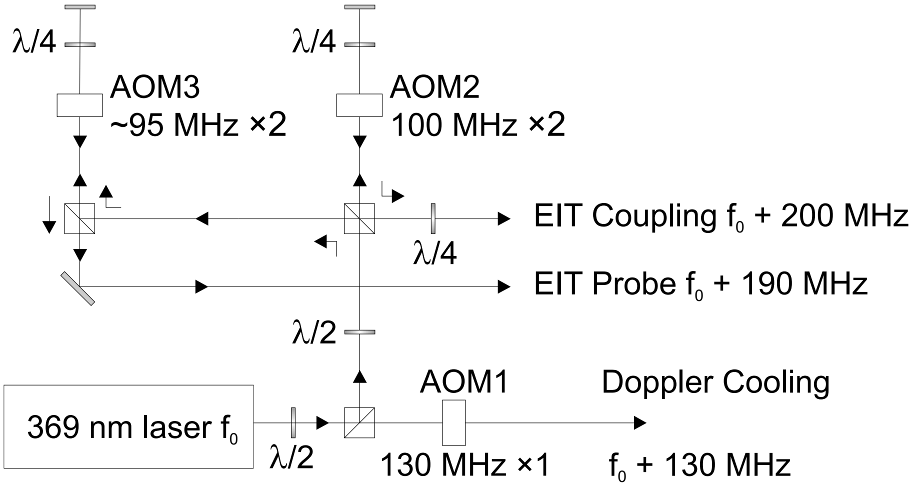

Figure 7.2: Schematic arrangement of the 369 nm beam paths for the EIT experiments.

RF switches are each controlled by four TTL lines; each line switches one of the outputs. One of the outputs of each switch is 50 Ω terminated, to enable signals to be switched off with minimum RF leakage. The other three outputs are connected to three power combiners. The output of each of these combiners is then amplified, and fed to an AOM. This system allows beams to be switched on and off, or have their frequencies changed, very quickly.

For example, it is possible to have synthesiser 1 connected to AOM2, providing the EIT coupling beam frequency, and synthesiser 2 connected to AOM3, at the correct frequency for EIT cooling, leaving synthesiser 3 and AOM1 disconnected. The RF output of synthesiser 3 can be pre-set to any frequency and amplitude, while being terminated by a 50 Ω load. It is then possible to swap the output of synthesiser 3 to AOM2, and terminate the output of synthesiser 2. This probes the EIT spectrum at another frequency immediately following a period of EIT cooling. This can be accurately synchronised with the photon-counting hardware. The limiting factor in the speed of all these operations is the time of propagation of the acoustic wave through the AOM. The reason the switches are required is that although all of these operations can be achieved by GPIB commands, these are neither fast enough nor consistent enough in their timing. GPIB commands for

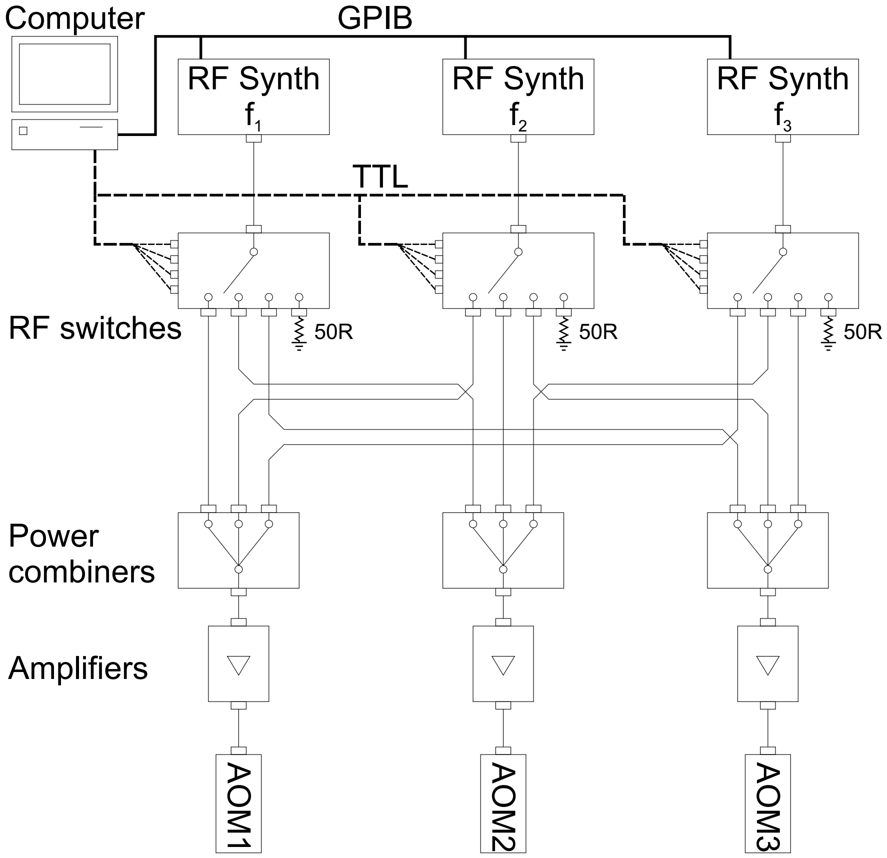

Figure 7.3: RF switching arrangement for the EIT experiments. The frequency and amplitude of the source RF can be controlled by GPIB (IEEE 488.2). The switches are each controlled by four TTL logic signals. Only one (or no) synthesiser is connected to an AOM at any one time. The switch/combiner network allows arbitrary routing of synthesisers to AOMs.

changing the frequency of a synthesiser, or switching the RF output on or off were timed to take anywhere between 20 ms and 100 ms to complete. This is longer than a full experimental cycle in certain circumstances.

7.2 EIT spectra

This section covers the spectra of the EIT feature which have been observed, and the results which may be derived from them. The results are discussed in approximate chronological order.

7.2.1 Initial results

The first stage of the experiment was to obtain spectra of the EIT feature on the red side of the transition. All the beams were red-detuned from line centre by approximately 60 MHz, and the fluorescence recorded for several seconds at each frequency as the probe beam frequency was scanned. The coupling beam polarisation was σ^+.

Some of these first spectra are shown in Figure 7.4. The size and position of the EIT peak vary with the power of the coupling and probe beams. This graph allows a calibration of the AC Stark shift of the feature and of the Zeeman splitting. The spectra all show the EIT peak having a sharp edge.

Figure 7.5 shows the positions of both the peak and the edge of the EIT feature as a function of coupling beam power. Fitting to the data gives the Zeeman splitting as the difference between the coupling laser frequency and zero-power position of the EIT peak, which is (2×)93.0 MHz. The Zeeman splitting of the $^2S_{1/2}$ state is therefore 14 MHz (the probe AOM is double-passed). The slopes of the two lines are -0.0246 for the edge and -0.0255 for the peak. This allows prediction of the optimum parameters for EIT cooling, which occur when the AC Stark shift of the peak is equal to the trap frequency (900 kHz for the radial motions in our trap). This calibration, performed on the red side of the transition, can then be transferred to the blue side of the transition for the actual cooling experiments. It is necessary to calibrate red-detuned so that normal photon-scattering events outside the

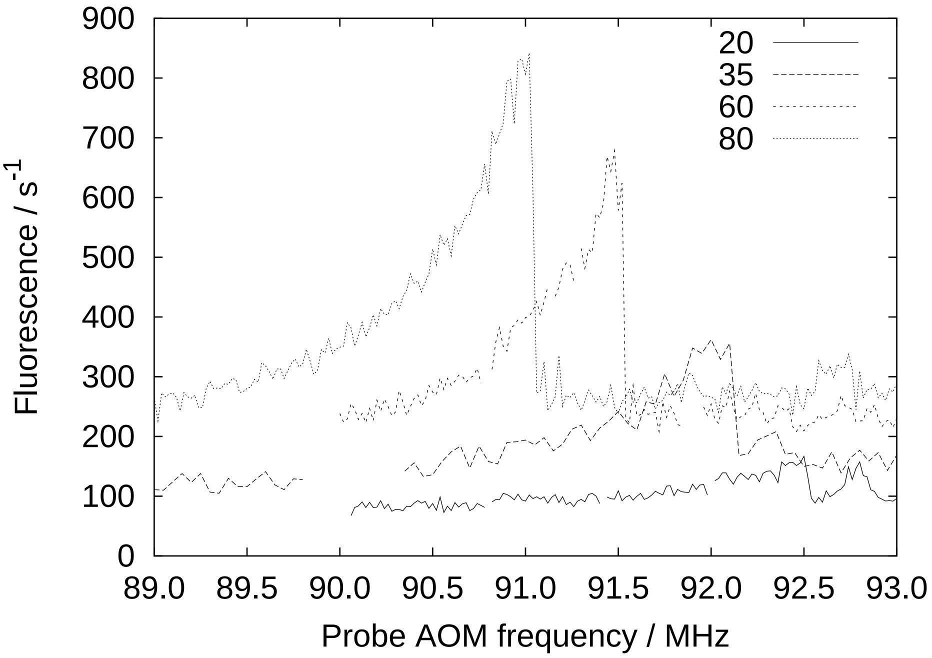

Figure 7.4: EIT spectra taken on 10 October. The spectra are labelled by the power of the coupling beam, in arbitrary units. The intensities of the probe and coupling beams were kept in a constant ratio of roughly 1:10. The empty periods are when the doubling cavity lost lock. Scans were taken in the direction of increasing probe AOM frequency.

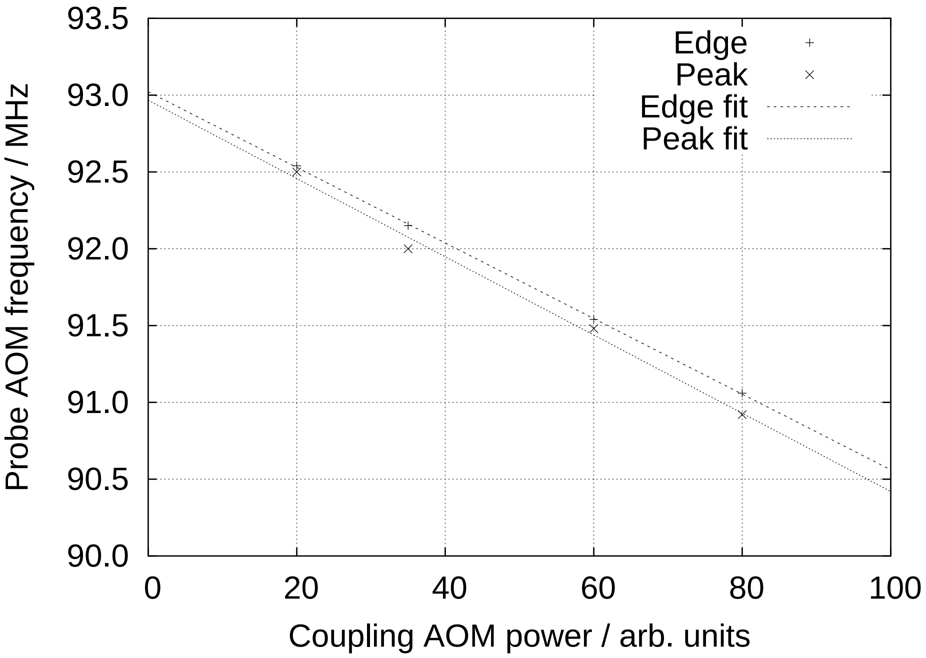

Figure 7.5: Calibration of the AC Stark shift and Zeeman splitting using the spectra in Figure 7.4.

region of EIT cooling do not lead to heating (and loss of fluorescence) of the ion. It was later possible to acquire these spectra while blue-detuned using a pulsed method, described in Section 7.2.5.

7.2.2 Observation of secular sidebands

Secular sidebands were also observed on the EIT transition. Figure 7.6 shows a spectrum with clear sidebands. The reason that sidebands are seen on some spectra and not others is that they are most obvious when the AC Stark shift and trap secular frequency are coincident. The sideband fluorescence is then visible above the minimum of the carrier absorption.

It is not possible to distinguish between axial sidebands and higher-order radial sidebands, as their frequencies are in a 2:1 ratio. Any splitting of the radial frequencies is far too small to resolve on this transition. In fact, no splitting of the radial frequencies has ever been conclusively observed in this

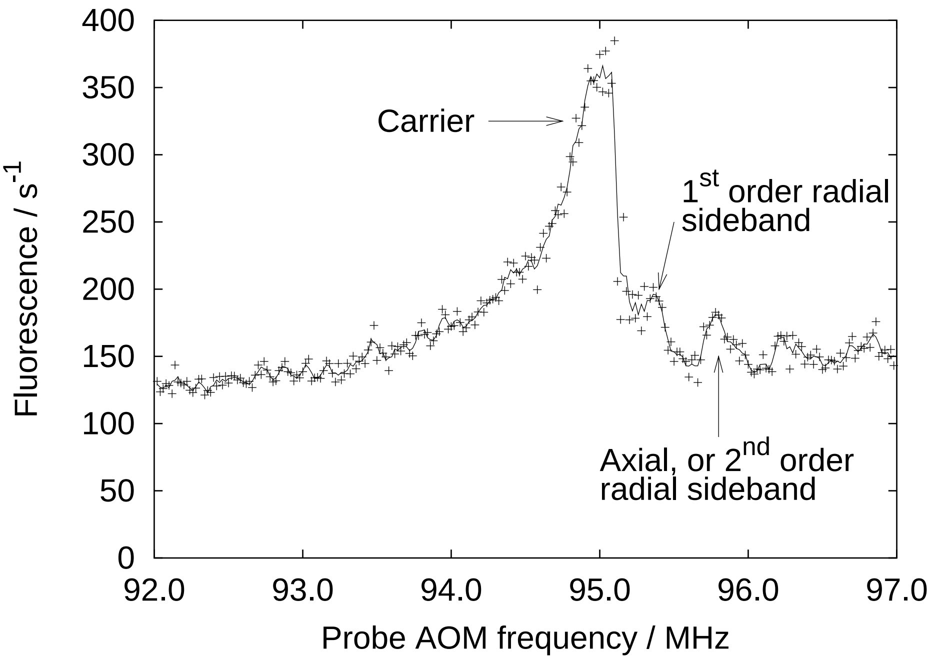

Figure 7.6: Secular sidebands of the EIT feature. This data was obtained on 9 October, red-detuned 80 MHz from line centre, with a coupling laser power of about 50, on the same power scale as Figure 7.4.

trap by any method.

7.2.3 Background fluorescence

All of the spectra in Figure 7.4 have a background count rate which is fairly constant over the range of the scans. It is larger than the sum of the count rates of the two beams individually, by a factor of more than four. It would appear that the EIT 'transparency' only dips down to the level of this background, and not to an overall zero count rate. This is shown more clearly in Figure 7.7. This spectrum was obtained at a high (unmeasured) coupling beam power, and with a smaller magnetic field than the earlier spectra, red-detuned 60 MHz. The Zeeman splitting is estimated at 7 MHz for this field. The peak is observed at (2×)93.5 MHz, 13 MHz from the coupling beam frequency. If 7 MHz of this splitting is due to the Zeeman shift, then the

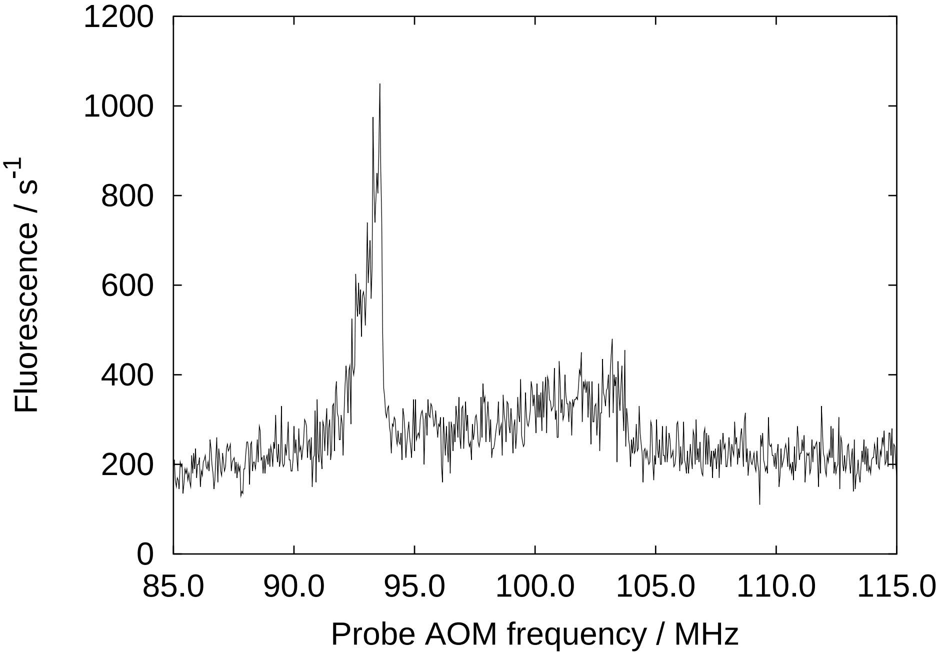

Figure 7.7: A wider scan of the EIT spectrum, taken on 9 October.

induced AC Stark shift is 6 MHz. The ‘background’ of the above scans appears to have an edge at a frequency of 104 MHz on the probe AOM, which is 8 MHz (in ‘real’ frequency, not AOM frequency) above the coupling laser and 20 MHz above the large peak.

The most likely explanation of the spectrum in Figure 7.7 appears to be that it has three main components. The most obvious component is the EIT profile created by the σ^+ coupling light and the π-polarised probe light. This is the component we are interested in, the other two can be classed as ‘noise’.

One of these components is the fairly flat, constant offset on the spectrum. It seems likely that this is normal single-photon scattering from the wings of the Lorentzian frequency distribution of the 369 nm transition. However, when the beams are individually scanned across the transition, the probe beam fluorescence is only 10–20% of the count rates shown in Figure 7.4, and the coupling beam fluorescence is negligible[1], because the polarisation is

---

[1] The σ^+ fluorescence is $\sim 30\ \mathrm{s}^{-1}$ observed close to line centre, compared to a maximum

very pure. When both probe and coupling beams illuminate the ion, though, the optical pumping efficiency will be destroyed, and so the fluorescence is expected to rise, though maybe not by as large a factor as is observed.

The remaining component is the smaller peak at about 103.7 MHz probe AOM frequency. This might also be a weak coherent resonance involving the unwanted polarisation components present in the 369 nm beams. This is illustrated in Figure 7.8. Any σ^- impurity in the σ^+ polarisation of the coupling beam can become resonant when the π-polarised component of the probe light is tuned to the appropriate frequency. The Zeeman splitting for the applied magnetic field is estimated at roughly 7 MHz. In the absence of the AC Stark shift, resonances would be expected at about ± 7 MHz relative to the coupling laser (96.5 MHz and 103.5 MHz probe AOM frequency). The resonance of the σ^+ component of the coupling beam is strongly AC Stark-shifted. The AC Stark shift of the σ^- resonance will be much lower, because the beam is much weaker. Given the uncertainty of the Zeeman splitting and the likely AC Stark shift, this seems a plausible explanation for the peak at 103.7 MHz.

7.2.4 Low-coherence spectra

Immediately following the initial spectra, we were unable to reproduce the results shown in Section 7.2.1. An ion was lost from the trap and it took five attempts to load another. Following this, all the spectra we acquired resembled those in Figure 7.9.

Many sidebands of the feature are visible in Figure 7.9, on both sides of the carrier. The feature is also very small, compared to spectra obtained earlier in the experiment, though the background level is similar. This supports the conclusion that the ion is hot (hence the visibility of the sidebands), and possibly that the coherence inherent to the EIT process is somehow being destroyed. The evidence for this second claim is seen more clearly in Figure 7.10, which shows spectra taken in similar conditions (both have the same coupling laser power) on two different days. While the more recent

---

observed count rate for π polarisation of $\sim 40,000$ s^{-1}

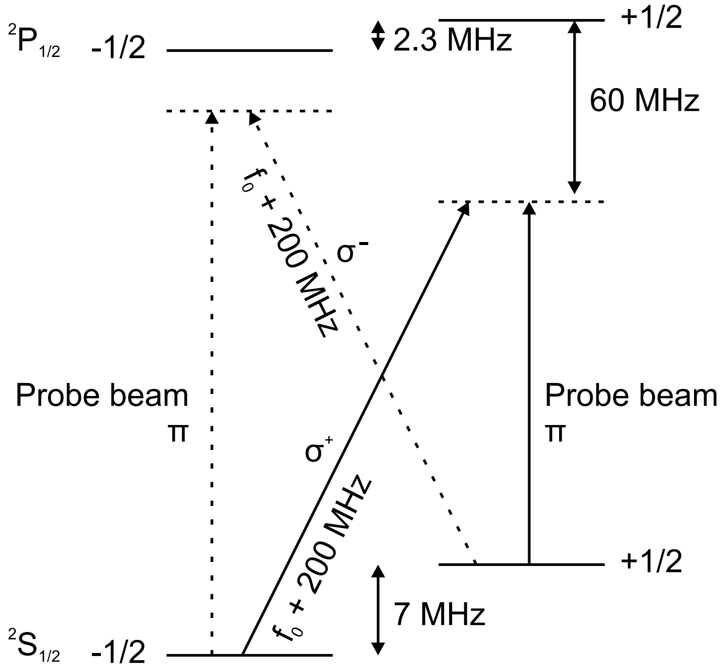

Figure 7.8: The presence of unwanted polarisations can produce secondary coherent resonances.

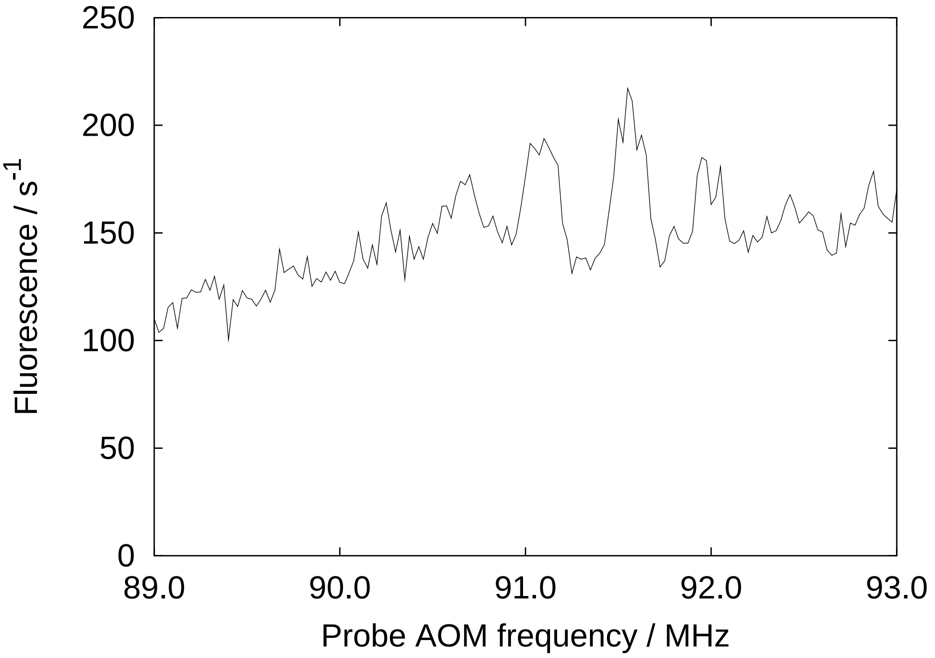

Figure 7.9: A medium-power (52) spectrum of the EIT feature, from Friday 17 October.

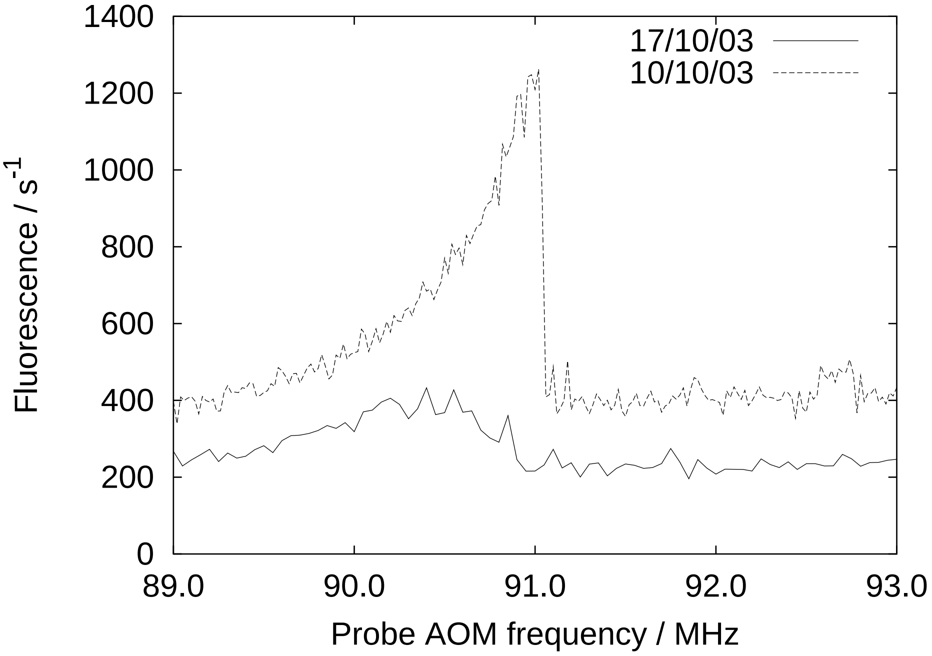

Figure 7.10: Spectra obtained on 10 and 17 October, both with a coupling laser power of 80. The probe laser power may have been different in the two cases.

data have a smaller background level (perhaps implying a lower probe laser power), they also has a different form. The main peak is in the same position, roughly, and therefore the AC Stark shift must be the same for the two curves. The peak does not, however, have a sharp edge, and resembles what might be expected if the coupling and probe beams were not coherent. It is expected that they are coherent, as they are both derived from the same source beam via synthesiser-driven AOMs. It is possible that some other process is destroying the coherence between the ground states of the system, and this would produce the same result. Might this be attributed to a rapid heating rate in the trap? This would also be consistent with the fact that it was difficult to keep ions for long periods without laser cooling during these few days of the experiment, whereas an ion had stayed in the trap without any laser cooling over the previous weekend.

In a later experimental run, there was no difficulty in reproducing the original results, so whatever the problem may have been, it appears to be temporary.

7.2.5 Blue-detuned spectra

It is desirable to acquire fluorescence spectra, and perform AC Stark shift calibrations, while blue-detuned if at all possible. This avoids the problems involved in transferring calibration data from the red side of the spectrum to the blue, such as uncertainties in the detuning from the 369 nm resonance.

It proved not to be possible to acquire blue-detuned 'slow accumulation' spectra such as those detailed in Section 7.2.1. These involved accumulating fluorescence for up to several seconds continuously at each probe frequency. When blue-detuned, the fluorescence disappeared when this was attempted, and the ion became very hot.

The shorter timescale 'time evolution' spectra, such as those in Section 7.3, have a very low data rate, due to the large overhead of cycling the experimental state for every 10 ms of data acquisition. So while it is possible to acquire blue-detuned calibration data using this technique, it is quite slow.

A compromise between these two methods was found to work well for obtaining such data. By collecting fluorescence for 100 ms, and then Doppler cooling for 100 ms, it was possible to acquire fluorescence data without the overheads of the time-resolved method, while keeping in check the rapid heating that can occur, such that the ion is always cool for the start of the next interrogation. Figure 7.11 shows the sum of 20 such spectra. One large peak is clearly visible above the background, and there are no obvious sidebands. The spectrum resembles the 'end-point' spectrum obtained from time-evolution experiments in Figure 7.13. This method is probably the best method for calibrating the AC Stark shift prior to attempting EIT cooling.

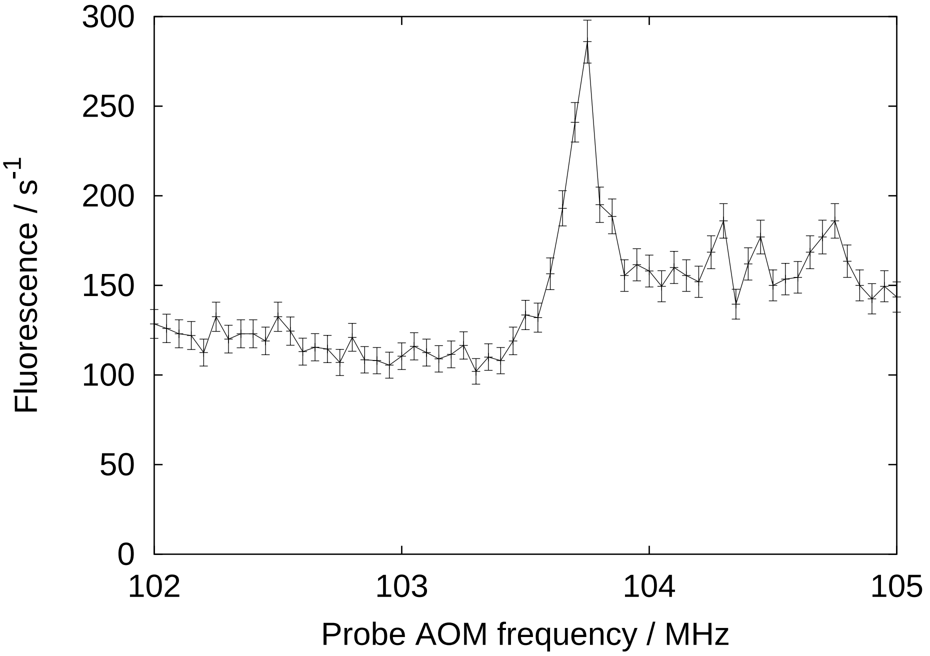

Figure 7.11: Blue-detuned fluorescence spectrum. This spectrum is the sum of 20 individual spectra. To acquire each of these spectra, fluorescence was accumulated for 100 ms at each probe AOM frequency, with 100 ms of Doppler cooling between each interrogation.

7.3 Time evolution of the fluorescence

7.3.1 EIT Fluorescence evolution

The second set of experiments, conducted on Monday 13 October, involved measuring the evolution of the fluorescence over short (millisecond) timescales. The experimental sequence was as follows:

1. Doppler cool the ion for 50 ms.
2. Turn off the Doppler cooling beam, and turn on the EIT beams, with the coupling AOM set to 100 MHz and the probe AOM set to the frequency of interest. This happens at time $t = 0$.
3. Measure the arrival times of each photon detected, up to a maximum time of 10 ms.
4. Repeat this cycle 5000 times.
5. Set a new probe AOM frequency, repeat.

The photon arrival times were sorted into 0.1 ms bins, and a fluorescence curve, such as that shown in Figure 7.12 could then be plotted. Figure 7.12 clearly shows a rapid initial decline in the fluorescence level, followed by the fluorescence reaching a plateau. The probe AOM frequency at which these data were obtained has been identified (tentatively) as the frequency of the red sideband of the EIT feature, from the spectra shown in Figure 7.13, and so this *may* be interpreted as reflecting the time evolution of the strength of the red sideband, and therefore evidence for EIT cooling.

By stepping across several probe AOM frequencies, a spectrum may be plotted. Each curve collected was fitted with the function $F(t) = Ae^{-Rt} + B$, where A represents the initial level above the background, R the decay rate, B the background level and F the observed fluorescence. Figure 7.13 shows $(A + B)$, representing the fluorescence at $t = 0$, and B, representing the final fluorescence, as a function of probe AOM frequency. The ‘initial’ spectrum clearly shows vibrational sideband structure, while the ‘final’ spectrum shows just a single peak.

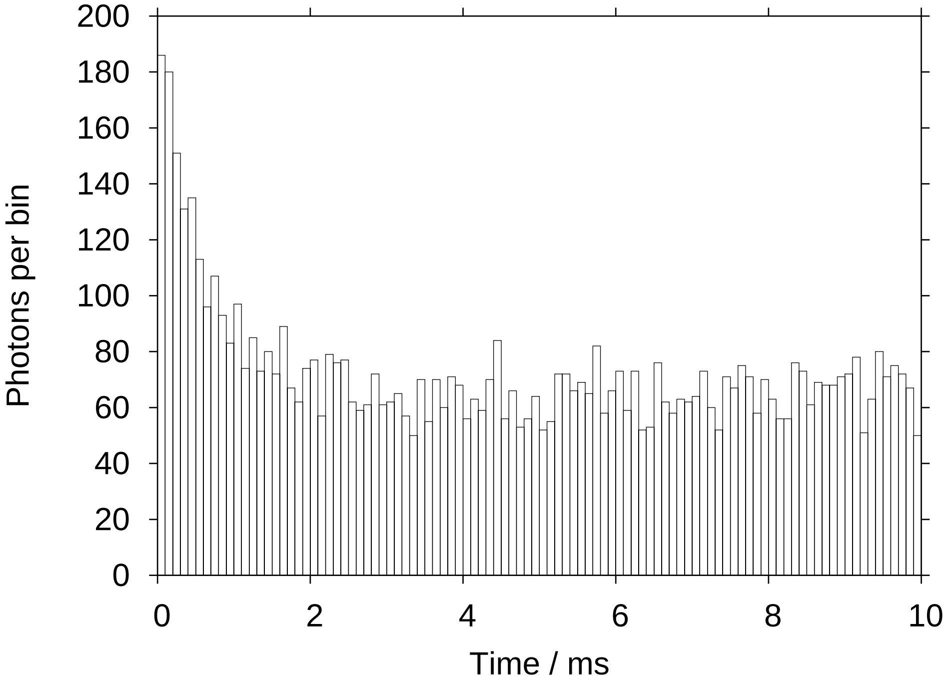

Figure 7.12: Evolution of the fluorescence level, over 10 ms. Taken on 13 October, with a probe AOM frequency of 93.6 MHz.

The initial spectrum is a single-temperature spectrum, representing the Doppler-cooled state of the ion. It is difficult to make a precision measurement from it, but it is clear that the red and blue sidebands are very similar in strength, and almost as high as the carrier. The modulation index is greater than 1, and so the average vibrational excitation number cannot be accurately determined. It appears that the ion is not cooled to the Doppler limit, but to a temperature a few times higher, at least during these experiments.

The final spectrum shows asymmetry at the positions of the red and blue sidebands, but this cannot be used as a measurement of temperature. At each individual point, the spectrum is the end-point of a process which modifies the vibrational state, in a frequency-dependent way. The spectrum does not represent a single temperature.

One interpretation of Figure 7.13 is that the central peak of the spectrum is the carrier, that the fluorescence decreases on the red sideband because EIT cooling has been observed, and decreases on the blue sideband because the ion is heating. The obvious problem with this is that the experiment cannot prove whether a fluorescence decrease signifies heating or cooling when the first-order blue sideband is so large (probably near its maximum value) in the initial state.

The ‘carrier’ feature grows slightly over the observation period, which might be expected if the ion were becoming cooler at that particular frequency, having started out relatively hot. It also seems to shift down in frequency by around 100 kHz. It is not clear why this occurs, or if indeed it is a real effect or the result of drift of experimental parameters. Assuming it is real, a putative explanation can be offered in terms of the heating and cooling dynamics. While on the red side of the carrier, there is a residual cooling effect from the ‘wings’ of the red sideband, and so over the period of observation the ion cools down, leading to a slight increase in the height of the carrier. Slightly to the blue side of the carrier, there is expected to be a residual heating effect. Given that the modulation index of the ion’s vibrational spectrum is greater than one at the start of the EIT interrogation, any further heating will only decrease the height of the carrier. Hence the carrier appears to shift due to the effect the cooling dynamics have on the

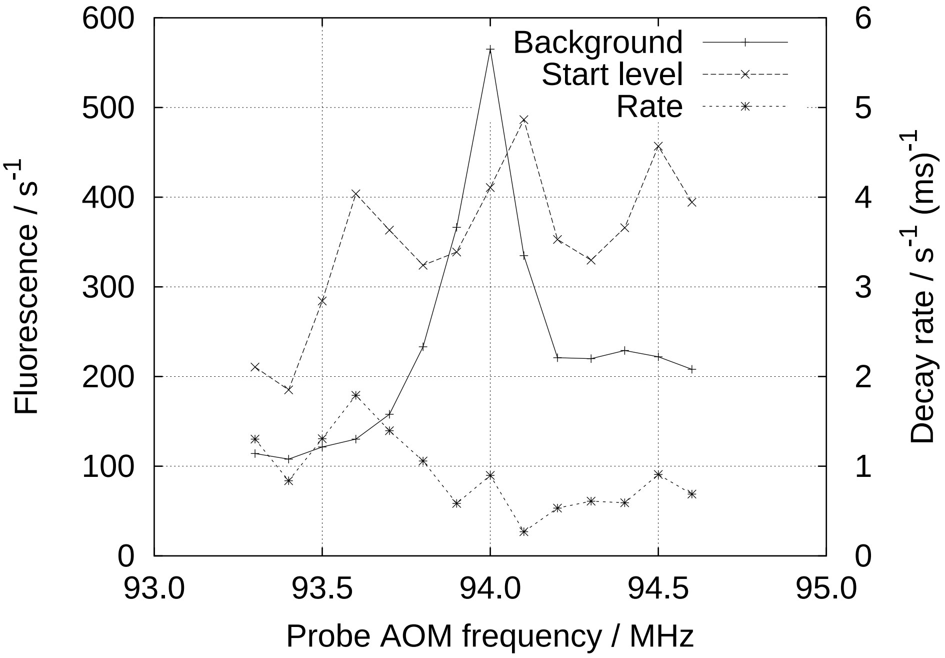

Figure 7.13: Fitted start and end points of the fluorescence level, over 10 ms. The characteristic rate of the fluorescence evolution is plotted on the right-hand axis. The peak at around 94.0 MHz has been identified as the carrier.

observed fluorescence.

There is a characteristic fluorescence evolution rate R associated with each point in Figure 7.13, and this is also plotted in Figure 7.13. The decay rate is highest on what is identified here as the red sideband (though that identification is not definitive), with a small peak on the carrier, and on the blue sideband. It should be noted that the rates at 93.9 MHz and 94.0 MHz correspond to an increase in fluorescence, while all others correspond to a decrease.

This evidence is entirely consistent with the observation of EIT cooling. However, it falls short of proof as there is a high level of interpretation involved. A convincing proof would involve comparing this experiment, with initial Doppler cooling only, to one with an additional EIT cooling stage in the preparation phase. Only what is plotted here as the 'initial' spectrum would be required for these two cases. Then it would be possible to see whether the blue sideband was still present after EIT cooling (which it should be), while the red sideband disappears. This would provide clear evidence that there was in fact a cooling process occurring. This may be difficult to observe using EIT spectra with the current level of background fluorescence. This could mask the small first-order blue sideband which is expected to remain after cooling to near the ground vibrational state. Unless this background can be reduced, a convincing observation of EIT cooling may require spectroscopy of the vibrational sidebands of the octupole transition.

7.3.2 Time evolution of single beam fluorescence

A small series of experiments on the time evolution of the ion's fluorescence was also conducted using just a single laser beam, rather than the EIT pair, to drive the transition. Figure 7.14 shows the evolution of the fluorescence with the EIT probe beam tuned 10 MHz ($\Gamma/2$) above line centre. The intensity of the probe beam was similar to that during the EIT experiments. The peak in the first bin may be due to a switching issue with the RF. There was no 'blank' period inserted into the digital IO pattern between the Doppler cooling and the blue detuned beam , and the first bin may contain remnants

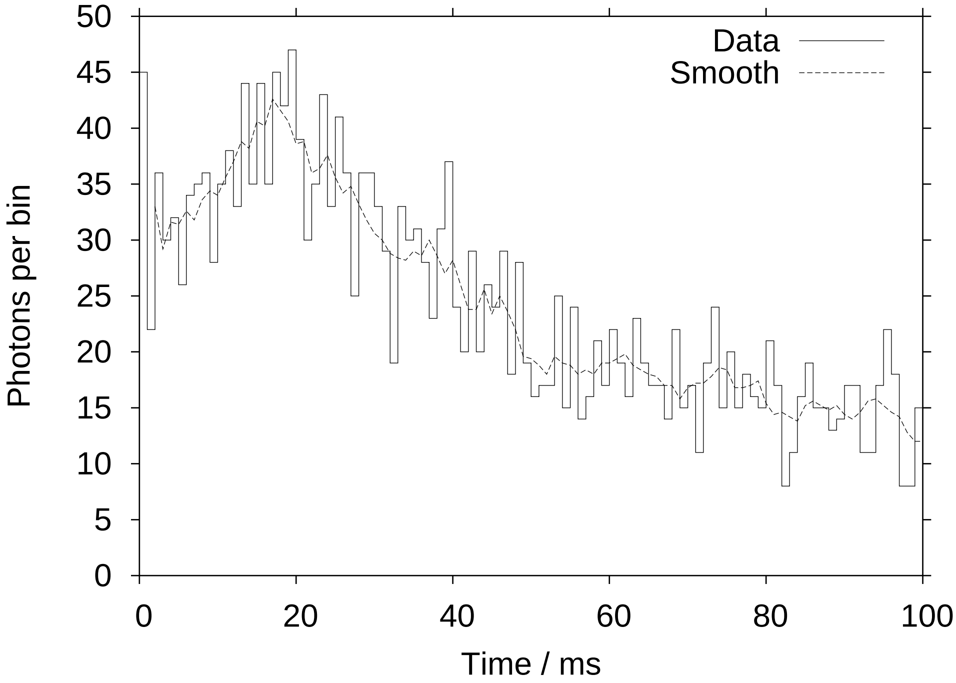

Figure 7.14: Evolution of the fluorescence with the single probe beam tuned 60 MHz above line centre of the cooling transition. The 'smooth' curve is a 5-sample moving average, as a guide to the eye. Data from 13 October.

of the Doppler cooling fluorescence which preceded the probe period. This has since been rectified, and does not affect the other experiments described in this Chapter.

As might be expected, the fluorescence initially rises, and then falls, as the ion is heated. At this detuning, the probe is close to the 10^{th} blue sideband of the cooling transition, though the sidebands are not resolved. The timescale for the fluorescence evolution is ~ 20 ms, much longer than the timescale of the fluorescence changes observed in Section 7.3 (of order 1 ms). Experiments observing the fluorescence over shorter timescales showed no evidence of any significant changes. That the timescales for fluorescence evolution in the single- and dual-beam experiments are so different lends weight to the idea that different mechanisms are operating in the two cases.

7.4 Summary

The initial results from the EIT experiments show promise, but several issues remain to be resolved. The source of the background fluorescence is still unclear. It may be that more stringent control of the polarisation of the coupling beam would reduce the background, or that detuning further from resonance would help. The time evolution experiments have produced some interesting results. The spectra show clear evidence of manipulation of the ion's vibrational state by sideband excitation, as the large sidebands present after Doppler cooling disappear through probing the spectrum at various frequencies. It it difficult to pin down exactly what changes are occurring though. Further investigations with an EIT cooling stage in the preparation phase of the experiment will be necessary to understand the dynamics in more depth. At some point a comparison of the data obtained from EIT spectra with the spectrum of the octupole transition will be necessary to be able to have full confidence in the EIT temperature measurement method. If the background fluorescence rate can be reduced, probably by more sophisticated and precise control of the coupling beam polarisation, the prospects for achieving EIT cooling in $^{172}\mathrm{Yb}^{+}$ appear good.

Chapter 8

Conclusion

The work described here has been directed towards developing the 467 nm electric octupole transition in $^{171}\mathrm{Yb}^+$ as an optical frequency standard. Within that framework, there have been several specific areas of focus.

The work with the most immediate impact has been the series of absolute frequency measurements of the transition, and the development of lasers, optical systems, electronics and control systems supporting this. While the overall uncertainty of the frequency of the octupole transition has been reduced, there is some evidence that we do not fully understand the systematic shifts of the transition. The most recent data, with a 150 Hz probe laser linewidth, have uncertainties much larger than those that we would ascribe to them based on our understanding of the experiment and its associated systematic shifts. Further measurements with improved control over experimental parameters may improve our understanding of the situation.

Despite these problems, the absolute frequency of the octupole transition has been measured to below 1 part in 10^{12}, and there is a clear path to improve this figure. There are several obvious limitations on the experiment currently. The most obvious is the linewidth of the probe laser. Construction of a new 934 nm Ti:sapphire laser with an intra-cavity EOM to further reduce the probe laser linewidth is underway. Improvements in the vibrational and acoustic isolation of the ULE reference cavity also have the potential to reduce the linewidth further. Another limitation is the data rate from spectroscopy of the transition. Scanning across the transition to obtain spectra is a less efficient method of obtaining frequency measurement data than probing the spectrum at two points (at $\pm\Gamma/2$) and using feedback from the control system to lock the probe laser to the transition. This will increase the data rate by an order of magnitude, and is currently being developed in the new LabVIEW control system.

The absolute frequency measurements have been underpinned by an evaluation of the systematic shifts of the transition. In addition to providing useful information about the current performance of the standard, this has highlighted areas which must be worked on in the medium term in order to realise the potential of the ytterbium octupole transition. After reducing the laser linewidth, which will remain the first priority for some time, a better calibration of the blackbody stark shift, possibly via a measurement of the coefficient of the DC quadratic Stark shift, will reduce the uncertainties of the systematic shifts of the transition below the 1 part in 10^{17} level. After that, better control of the quadrupole shift will become a priority. What the evaluation of the systematic shifts has demonstrated most clearly is the potential of the ytterbium octupole transition to be a true next-generation optical frequency standard.

In the longer term, a move to use ground-state cooling in the experiment will be required to reduce the uncertainties of the second-order Doppler and DC Stark shifts. The investigations undertaken here into ground-state cooling in general, and EIT cooling in particular, will help to inform the route that the experiment takes.

The numerical modelling of the EIT cooling process has shown it to be robust against many likely sources of experimental imperfection, such as imperfect calibrations, or instabilities of beam frequencies and intensities.

This has been confirmed, to a certain extent, by the experimental investigations of EIT in $^{171}\mathrm{Yb}^{+}$. It has been possible to calibrate AC Stark shifts and beam frequencies with the accuracy which is necessary for EIT cooling.

The observation of vibrational sidebands in the EIT spectra is a significant milestone on the path to EIT cooling. The spectra have also illuminated issues with the current experimental arrangement. The background fluorescence level from the ion is higher than expected. This is at least partially due to imperfect control of the coupling beam polarisation, and this is an area for future improvement.

The time-resolved spectra of the EIT spectrum have revealed some interesting behaviour, and show strong evidence of the modification of the vibrational state of the ion by the driving of the vibrational sidebands of the EIT resonance peak. By acquiring such spectra immediately following a period of EIT cooling, it should be possible to demonstrate EIT cooling. This could then be confirmed by quantum-jump spectroscopy of the vibrational sidebands of the octupole transition. Once the cooling process has been demonstrated, a thorough investigation of the experimental characteristics of the cooling process can be undertaken. This could include the dependence of the achieved temperature on polarisation purity, beam intensity stability, and similar factors. These would all allow optimisation of the cooling process for real experimental situations, and hopefully the full potential of the EIT cooling scheme can be realised. This work would allow a full comparison of the EIT cooling scheme against direct and Raman sideband cooling schemes, and inform decisions about the most appropriate cooling scheme to use in future experiments at NPL and elsewhere.

Appendix A

The quadrupole shift

The electric quadrupole shift is one of the more significant systematic shifts for trapped-ion frequency standards. This Appendix covers the details of calculating the shift in depth. In doing this, it follows the treatment of the shift by W. Itano for the $^{199}\mathrm{Hg}^{+}$ ion [72]. The shift itself is caused by the interaction of the quadrupolar spatial component of a DC electric field with the quadrupole moment of an atomic state. The $^{2}\mathrm{S}_{1/2}$ state has no quadrupole moment, so the shift of the 467 nm transition is entirely due to the shift of the $^{2}\mathrm{F}_{7/2}$ state.

A.1 The quadrupole shift

A.1.1 Reference frames

Two co-ordinate systems are used in the following calculation. The first, 'laboratory' frame has co-ordinates (x, y, z), and has its z-axis aligned with the direction of the magnetic field B. The 'trap' frame, denoted by primed co-ordinates (x', y', z'), has its z-axis aligned with the trap z-axis, along the endcaps. The assumption is made that any DC quadrupole field is being produced by DC offset potentials on the trap electrodes.

A.1.2 The Hamiltonian

The Hamiltonian for the quadrupole shift is

$$H_Q = \nabla E^{(2)} \cdot \Theta^{(2)} = \sum_{q=-2}^{2} (-1)^q \nabla E_q^{(2)} \Theta_{-q}^{(2)} \tag{A.1}$$

where the $E_q^{(2)}$ are the quadrupole components of the electric field, and $\Theta_{-q}^{(2)}$ are the components of the electric quadrupole moment of the atom.

A.1.3 The electric field

The quadrupolar field has the form, in the trap frame

$$\phi\left(x', y', z'\right) = A\left[\left(x'^2 + y'^2 - 2z'^2\right) + \epsilon\left(x'^2 - y'^2\right)\right]. \tag{A.2}$$

Henceforth it will be assumed that the trap is radially symmetric and that $\epsilon = 0$. This simplifies the derivation from that in [72] without losing the essential physics. In particular, both the electric field gradient tensors and quadrupole moment operator have several components. Setting $\epsilon = 0$ allows us to ignore all but the $q = 0$ components. The relevant component of the electric field gradient tensor is

$$\nabla E_0^{(2)} = -\frac{1}{2}\frac{\mathrm{d}E_z}{\mathrm{d}z} \tag{A.3}$$

The electric field gradient for this field, in the trap frame, is

$$\nabla E_0^{(2)\prime} = -2A \tag{A.4}$$

The Hamiltonian in the trap frame is simply

$$H_Q = -2A\Theta_0^{(2)\prime} \tag{A.5}$$

Note that the quadrupole moment is expressed in the trap frame here, and not in the ‘lab’ frame defined by the magnetic field.

A.1.4 Matrix elements

The shift of a state $|\gamma JFM_F\rangle$ is given by the expectation value of the Hamiltonian interacting with that state:

$$\Delta E(\gamma JFM_F) = \langle \gamma JFM_F | H_Q | \gamma JFM_F \rangle \tag{A.6}$$

This is the expression in the laboratory frame. The states $|\gamma JFM_F\rangle$ in the lab frame are related to the states $|\gamma JFM'\rangle$ in the trap frame by rotation matrix elements $D^{(F)}_{M'M_F}(\beta)$, where (β) is the angle of rotation relating the two frames. Only one angle is required for calculating the $q = 0$ components due to the symmetry of the situation.

$$|\gamma JFM'\rangle = \sum_{M_F} D^{(F)}_{M'M_F}(\beta) |\gamma JFM_F\rangle \tag{A.7}$$

The matrix element which must be calculated to evaluate the shift is

$$\langle \gamma JFM_F | \Theta^{(2)\prime}_0 | \gamma JFM_F \rangle \tag{A.8}$$

It can be shown [72] that this can be written

$$\begin{aligned} &\langle \gamma JFM_F | \Theta^{(2)\prime}_0 | \gamma JFM_F \rangle \\ &\quad = (-1)^{F-M_F} (\gamma JF || \Theta^{(2)} || \gamma JF) \begin{pmatrix} F & 2 & F \\ -M_F & 0 & M_F \end{pmatrix} D^{(2)*}_0(\beta) \end{aligned} \tag{A.9}$$

The rotation matrix element is $D^{(2)*}_0(\beta) = \frac{1}{2}\left(3\cos^2\beta - 1\right)$. The shift can therefore be written

$$\Delta E(\gamma JFM_F) = -A\left(3\cos^2\beta - 1\right) K_1 (\gamma JF || \Theta^{(2)} || \gamma JF) \tag{A.10}$$

With the constant K_1 defined by

$$K_1 = (-1)^{F-M_F} \begin{pmatrix} F & 2 & F \\ -M_F & 0 & M_F \end{pmatrix} \tag{A.11}$$

For the $^2\mathrm{F}_{7/2}(F=3, M_F=0)$ state, $K_1 = -\frac{2}{3}\sqrt{\frac{3}{35}}$.

The remaining undefined quantity is the reduced matrix element $(\gamma JF||\Theta^{(2)}||\gamma JF)$. This can be related to the quadrupole moment of the state $\Theta(\gamma, J)$. In the IJ–coupling approximation,

$$\begin{aligned}&(\gamma(IJ)F||\Theta||\gamma(IJ)F) = \\ &\qquad(-1)^{I+J+F}(2F+1)\begin{Bmatrix} J & 2 & J \\ F & I & F \end{Bmatrix}\begin{pmatrix} J & 2 & J \\ -J & 0 & J \end{pmatrix}^{-1}\Theta(\gamma, J)\end{aligned} \tag{A.12}$$

Hence the shift (Equation (A.10)) can be written

$$\langle H_Q\rangle = -A(3\cos^2\beta - 1)K_1K_2\Theta(\gamma, J) \tag{A.13}$$

With the constant K_2 being defined by

$$K_2 = (-1)^{I+J+F}(2F+1)\begin{Bmatrix} J & 2 & J \\ F & I & F \end{Bmatrix}\begin{pmatrix} J & 2 & J \\ -J & 0 & J \end{pmatrix}^{-1} \tag{A.14}$$

For the $^2\mathrm{F}_{7/2}(F=3, M_F=0)$ state, $K_2 = +\frac{5}{2}\sqrt{\frac{15}{7}}$.

A.1.5 The quadrupole shift for the ^{2}F$_{7/2}$ state

Using the constants calculated above, $K_1 = -\frac{2}{3}\sqrt{\frac{3}{35}}$ and $K_2 = +\frac{5}{2}\sqrt{\frac{15}{7}}$, the quadrupole shift for the $^2\mathrm{F}_{7/2}(F=3, M_F=0)$ state in $^{171}\mathrm{Yb}^+$ is

$$\langle H_Q\rangle = +\frac{5}{7}A\left(3\cos^2\beta - 1\right)\Theta(\gamma, J) \tag{A.15}$$

A.2 The quadrupole moment

The quadrupole moment $\Theta(\gamma, J)$ is defined as the matrix element of the quadrupole operator for the state of maximum M_J. That is

$$\Theta(\gamma, J) = \langle \gamma J J | \Theta_0^{(2)} | \gamma J J \rangle \tag{A.16}$$

The quadrupole operator is

$$\Theta_0^{(2)} = -\frac{e}{2} \sum_j (3z_j^2 - r_j^2) \tag{A.17}$$

The summation is over all the atomic electrons, and z_j and r_j are the electronic position operators. A closed atomic shell has no quadrupole moment, and so only unfilled shells need be considered. Writing in terms of spherical harmonics, this is

$$\Theta(\gamma, J) = -e\sqrt{\frac{4\pi}{5}} \langle \gamma J J | Y_{20} | \gamma J J \rangle. \tag{A.18}$$

But

$$\langle 4f\ ^2F_{7/2}, m_j = 7/2 | Y_{20} | 4f\ ^2F_{7/2}, m_j = 7/2 \rangle = \langle 4f, m_l = 2 | Y_{20} | 4f, m_l = 2 \rangle \tag{A.19}$$

Therefore, the quadrupole moment can be written

$$\Theta(\gamma, J) = -e\sqrt{\frac{4\pi}{5}} \langle \gamma | r^2 | \gamma \rangle \int_0^{2\pi} \int_0^{\pi} Y_{LL}^* Y_{20} Y_{LL} \sin\theta d\theta d\phi. \tag{A.20}$$

The integrals over spherical harmonics can be written in terms of 3j symbols and normalisation factors. In the general case,

$$\int_0^{2\pi}\int_0^{\pi} Y^*_{l_1 m_1} Y_{l_2 m_2} Y_{l_3 m_3} \sin\theta d\theta d\phi =$$
$$(-1)^{m_1}\sqrt{\frac{(2l_1+1)(2l_2+1)(2l_3+1)}{4\pi}}\begin{pmatrix} l_1 & l_2 & l_3 \\ 0 & 0 & 0 \end{pmatrix}\begin{pmatrix} l_1 & l_2 & l_3 \\ -m_1 & m_2 & m_3 \end{pmatrix} \tag{A.21}$$

Therefore

$$\Theta(\gamma, J) = -e\langle\gamma|r^2|\gamma\rangle(-1)^L(2L+1)\begin{pmatrix} L & 2 & L \\ 0 & 0 & 0 \end{pmatrix}\begin{pmatrix} L & 2 & L \\ -L & 0 & L \end{pmatrix} \tag{A.22}$$

For the ${}^2\mathrm{F}_{7/2}$ state in ${}^{171}\mathrm{Yb}^+$ the quadrupole moment is due to the *absence* of an electron , or hole in, an otherwise full f-shell. This changes the sign of the quadrupole moment. From that calculated above, which was treated as a single electron in the 4f state, rather than a single hole. The orbital angular momentum is $L = 3$ for the ${}^2\mathrm{F}_{7/2}$ state. The quadrupole moment is therefore

$$\Theta(4f, 7/2) = -e\frac{1}{3}\langle 4f|r^2|4f\rangle \tag{A.23}$$

A.2.1 Matrix elements of $\langle r^2\rangle$

The matrix elements of $\langle r^2\rangle$ can be estimated from atomic structure calculations. The Cowan atomic structure code [89] calculates atomic orbitals, and related properties, for many-electron atoms. From these calculations, the expectation values in Table A.1 have been found.

The value for ${}^{199}\mathrm{Hg}^+$ agrees with that calculated in [72], using the same code. The value for ${}^{88}\mathrm{Sr}^+$ does not agree with the published value in [11]. However, if the calculation is performed for the 4d orbital of neutral strontium, there seems to be agreement [75]. This perhaps points to an error in the initialisation of the Cowan code, though this has not been confirmed with the authors of [11].

Species	**Orbital**	$\langle r^2 \rangle / a_0^2$
^{171}Yb$^+$	$4f$	0.657
^{171}Yb$^+$	$5d$	10.8
^{199}Hg$^+$	$5d$	2.32
^{88}Sr$^+$	$4d$	10.6
^{40}Ca$^+$	$3d$	6.64

Table A.1: Expectation values of the squared electronic co-ordinate for various atomic orbitals, in units of the Bohr radius a_0.

The size of the $4f$-orbital in ^{171}Yb$^+$ is particularly small. In common with the $5d$ orbital in mercury, it is not the 'outermost' occupied orbital. In both cases the $6s^2$ orbital is full, and the source of the quadrupole moment is a hole in an inner shell. The d-orbitals in ytterbium, strontium and calcium ions are occupied by a single electron, 'outside' completely full inner shells, which gives some indication as to why those orbitals are relatively larger.

Using the calculated $\langle r^2 \rangle$, the quadrupole moment of the ^{2}F$_{7/2}(F = 3, M_F = 0)$ state in ^{171}Yb$^+$is

$$\begin{aligned} \Theta(4f, 7/2) &= -0.22 e a_0^2 \\ &= -6.13 \times 10^{-22} e \text{ m}^2 \\ &= -9.82 \times 10^{-41} \text{Cm}^2 \end{aligned} \tag{A.24}$$

Publications

References [73, 78, 90, 76, 77] have been published by the author during the course of this thesis. In addition to these papers, the following conference proceedings have been published, but not necessarily peer-reviewed: [91, 92, 93, 4, 94].

Bibliography

[1] J. K. Webb, M. T. Murphy, V. V. Flambaum, V. A. Dzuba, J. D. Barrow, C. W. Churchill, J. X. Prochaska, and A. M. Wolfe. Further evidence for cosmological evolution of the fine structure constant. *Physical Review Letters*, 87:091301, 2001.

[2] S. Bize, S. A. Diddams, U. Tanaka, C. E. Tanner, W. H. Oskay, R. E. Drullinger, T. E. Parker, T. P. Heavner, S. R. Jefferts, L. Hollberg, W. M. Itano, and J. C. Bergquist. Testing the stability of fundamental constants with the ^{199}Hg^{+} single-ion optical clock. *Physical Review Letters*, 90:150802, 2003.

[3] V. A. Dzuba, V. V. Flambaum, and M. V. Marchenko. Relativistic effects in sr, dy, yb ii, and yb iii and search for variation of the fine-structure constant. *Physical Review A*, 68:022506, 2003.

[4] S.N. Lea, P.J. Blythe, S.A. Webster, K. Hosaka, H.S. Margolis, J.L. Flowers, and P. Gill. Proposed laboratory test of time-variation of alpha using octupole and quadrupole transitions in ^{171}Yb^{+}, 2003. Heraeus Seminar on Astrophysics, Clocks and Fundamental Constants, Bad Honnef.

[5] I. I. Rabi, S. Millman, P. Kusch, and J. R. Zacharias. The molecular beam resonance method for measuring nuclear magnetic moments. *Phys.Rev.*, 55:526–535, March 1939.

[6] Norman F. Ramsey. Fifty years of atomic frequency standards. In P. Gill, editor, *Proceedings of the 6th Symposium on Frequency Standards and Metrology*, pages 8–17, Singapore, 2002. World Scientific.

[7] W. Markowitz, R. Glenn Hall, L. Essen, and J. V. L. Parry. Frequency of cesium in terms of ephemeris time. *Physical Review Letters*, 1(3):105–107, August 1958.

[8] N. F. Ramsey. Molecular beam resonances in oscillatory fields of nonuniform amplitudes and phases. *Phys.Rev.*, 109(3):822–825, February 1958.

[9] S. R. Jefferts, T. P. Heavner, J. Shirley, and T. E. Parker. Systematic frequency shifts and quantum projection noise in NIST-F1. In P. Gill, editor, *Proceedings of the 6th Symposium on Frequency Standards and Metrology*, pages 361–368, Singapore, 2002. World Scientific.

[10] H. S. Margolis, G. Huang, G. P. Barwood, S. N. Lea, H. A. Klein, W. R. C. Rowley, P. Gill, and R. S. Windeler. Absolute frequency measurement of the 674-nm $^{88}\mathrm{Sr}^+$ clock transition using a femtosecond optical frequency comb. *Physical Review A*, 67:032501, March 2003.

[11] J. E. Bernard, A. A. Madej, L. Marmet, B. G. Whitford, K. J. Siemsen, and S. Cundy. Cs-based frequency measurement of a single, trapped ion transition in the visible region of the spectrum. *Physical Review Letters*, 82(16):3228–3231, April 1999.

[12] M. Knoop, C. Champenois, P. Courteille, M. Herbane, M. Houssin, M. Vedel, and F. Vedel. Single Ca^+ ions in a Paul-Straubel trap. In P. Gill, editor, *Proceedings of the 6th Symposium on Frequency Standards and Metrology*, pages 492–494, Singapore, 2002. World Scientific.

[13] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Laser-cooled mercury ion frequency standard. *Physical Review Letters*, 80:2089–2092, March 1998.

[14] J. Stenger, Chr. Tamm, N. Haverkamp, S. Weyers, and H.R. Telle. Absolute frequency measurement of the 435.5-nm $^{171}\mathrm{Yb}^+$ clock transition

with a kerr-lens mode-locked femtosecond laser. *Opt. Lett.*, 26:1589–1591, October 2001.

[15] D. J. Wineland, J. C. Bergquist, J. J. Bollinger, R. E. Drullinger, and W. M. Itano. Quantum computers and atomic clocks. In P. Gill, editor, *Proceedings of the 6th Symposium on Frequency Standards and Metrology*, pages 361–368, Singapore, 2002. World Scientific.

[16] Th. Becker, J. von Zanthier, A. Yu. Nevsky, Ch. Schwedes, M. N. Skvortsov, H. Walther, and E. Peik. High-resolution spectroscopy of a single In^+ ion: Progress towards an optical frequency standard. *Physical Review A*, 63:051802(R), May 2001.

[17] S. Gulde, M Riebe, G. P. T. Lancaster, C. Becher, J. Eschner, H. H affner, F. Schmidt-Kaler, I. L. Chuang, and R. Blatt. Implementation of the Deutsch-Jozsa algorithm on an ion-trap quantum computer. *Nature*, 421:48–50, 2003.

[18] D. M. Meekhof, C. Monroe, B. E. King, W. M. Itano, and D. J. Wineland. Generation of nonclassical motional states of a trapped atom. *Physical Review Letters*, 76(11):1796–1799, March 1996.

[19] D. Leibfried, B. DeMarco, V. Meyer, D. Lucas, M. Barrett, J. Britton, W. M. Itano, B. Jelenkovic, C. Langer, T. Rosenband, and D. J. Wineland. Experimental demonstration of a robust, high-fidelity geometric two ion-qubit phase gate. *Nature*, 422:412–414, 2003.

[20] F. Schmidt-Kaler, H. H affner, M Riebe, S. Gulde, G. P. T. Lancaster, T. Deuschle, C. Becher, C. F. Roos, J. Eschner, and R. Blatt. Realization of the Cirac-Zoller controlled-NOT quantum gate. *Nature*, 422:408–411, 2003.

[21] Giovanna Morigi, Jürgen Eschner, and Christoph H. Keitel. Ground state laser cooling using electromagnetically induced transparency. *Physical Review Letters*, 85(21):4458–4461, November 2000.

[22] C. F. Roos, D. Leibfried, A. Mundt, F. Schmidt-Kaler, J. Eschner, and R. Blatt. Experimental demonstration of ground state laser cooling with electromagnetically induced transparency. *Physical Review Letters*, 85(26):5547–5550, December 2000.

[23] J. P. Marangos. *Electromagnetically Induced Transparency*, volume IV of *Handbook of Optics*, chapter 23. Optical Society of America / McGraw-Hill Education, November 2000. ISBN 0071364560.

[24] Y. Stalgies, I. Siemers, B. Appasamy, and P. E. Toschek. Light shift and Fano resonances in a single cold ion. *Journal of the Optical Society of America B*, 15(10):2505–2514, 1998.

[25] I. Asimov. *The Stars in their Courses*. Panther Books Ltd., St Albans, U.K., 1975.

[26] M. Roberts, P. Taylor, W. R. C. Rowley, G. P. Barwood, and P. Gill. Observation of the $^2S_{1/2} -^2 F_{7/2}$ electric octupole transition in a single ^{171}Yb$^+$ ion. *Physical Review A*, 62, July 2000.

[27] Matthew Roberts. *Spectroscopy of a single ytterbium ion*. PhD thesis, Imperial College, University of London, and National Physical Laboratory, Teddington, January 1997.

[28] D. Engelke and Chr. Tamm. Dark times in the resonance fluorescence of trapped ^{171}Yb$^+$ ions caused by spontaneous quantum jumps to the ^{2}D$_{3/2}$ (F=2) state. *Europhysics Letters*, 33(5):347–352, February 1996.

[29] Wolfgang Paul. Electromagnetic traps for charged and neutral particles. *Reviews of Modern Physics*, 62(3):531–540, 1990. Nobel Prize Lecture 1990.

[30] D. J. Berkeland, J. D. Miller, J. C. Bergquist, W. M. Itano, and D. J. Wineland. Minimization of ion micromotion in a Paul trap. *Journal of Applied Physics*, 83(10):5025–5033, May 1998.

[31] P. K. Ghosh. *Ion traps*. Clarendon Press, Oxford, 1995.

[32] D. J. Bate, K. Dholakia, R. C. Thompson, and D. C. Wilson. Ion oscillation frequencies in a combined trap. *Journal of Modern Optics*, 39(2):305–316, February 1992.

[33] P. Taylor, S. A. Webster, M. Roberts, W. R. C. Rowley, G. P. Barwood, and P. Gill. Progress towards a single ytterbium ion optical clock. In *Proceedings of LASE 2001, High Power Lasers and Applications, Photonics West*. SPIE, 2001.

[34] Paul Taylor. *Observation of an Ultra-High Q Resonance in a Single Ion of* $^{172}Yb^{+}$. PhD thesis, Department of Phyics, University of Oxford and National Physical Laboratory, Teddington, September 1996.

[35] D. W. Allan. Time and frequency (time-domain) characterization, estimation, and prediction of precision clocks and oscillators. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, UFFC-34(6):647–654, November 1987.

[36] N. F. Ramsey. A molecular beam resonance method with separated oscillating fields. *Phys. Rev.*, 78:695–699, 1950.

[37] N. F. Ramsey. Experiments with separated oscillatory fields and hydrogen masers. *Rev. Mod. Phys.*, 62:541–552, 1990.

[38] S. N. Lea, G. M. Macfarlane, G. Huang, and P. Gill. Progress toward an optical frequency chain at NPL. *IEEE Trans. Instrum. Meas.*, 48:578–582, 1999.

[39] H. Dehmelt. Proposed $10^{14}\delta\nu < \nu$ laser fluorescence spectroscopy on Tl^{+} mono-ion oscillator II (spontaneous quantum jumps). *Bull. Am. Phys. Soc.*, 20:60, 1975.

[40] Warren Nagourney, Jon Sandberg, and Hans Dehmelt. Shelved optical electron amplifier: Observation of quantum jumps. *Physical Review Letters*, 56(26):2797–2799, June 1986.

[41] D. J. Wineland and Wayne M. Itano. Laser cooling of atoms. *Physical Review A*, 20(4):1521–1540, October 1979.

[42] Wayne M. Itano and D. J. Wineland. Laser cooling of ions stored in harmonic and Penning traps. *Physical Review A*, 25(1):35–54, January 1982.

[43] S. R. Jefferts, C. Monroe, E. W. Bell, and D. J. Wineland. Coaxial-resonator-driven rf (Paul) trap for strong confinement. *Physical Review A*, 51(4):3112–3116, April 1995.

[44] I. Marzoli, J. I. Cirac, R. Blatt, and P. Zoller. Laser cooling of trapped three-level ions: Designing two-level systems for sideband cooling. *Physical Review A*, 49(4):2771–2779, April 1994.

[45] D. J. Wineland, C. Monroe, W. M. Itano, D. Leibfried, B. E. King, and D. M. Meekhof. Experimental issues in coherent quantum-state manipulation of trapped atomic ions. *J. Res. Natl. Inst. Stand. Technol.*, 103(3):259–328, May-June 1998.

[46] V. Letchumanan. Sideband cooling of strontium. First year report, Imperial College / National Physical Laboratory, April 2001.

[47] F. Diedrich, J. C. Bergquist, Wayne M. Itano, and D. J. Wineland. Laser cooling to the zero-point energy of motion. *Physical Review Letters*, 62(4):403–406, January 1989.

[48] E. Peik, J. Abel, Th. Becker, J. von Zanthier, and H. Walther. Sideband cooling of ions in radio-frequency traps. *Physical Review A*, 60(1):439–449, January 1999.

[49] C. F. Roos, Th. Zeiger, A. Nägerl, D. Leibfried, F. Schmidt-Kaler, and R. Blatt. Quantum state engineering on an optical transition and decoherence in a Paul trap. *Physical Review Letters*, 83(23):4713–4716, December 1999.

[50] V. Letchumanan. *Coherent control and ground state cooling of a single trapped strontium ion*. PhD thesis, Imperial College London / National Physical Laboratory, 2004.

[51] B. E. King. *Quantum State Engineering and Information Processing with Trapped Ions.* PhD thesis, University of Colorado, 1999.

[52] C. Monroe, D. M. Meekhof, B. E. King, S. R. Jefferts, W. M. Itano, D. J. Wineland, and P. Gould. Resolved-sideband Raman cooling of a bound atom to the 3D zero-point energy. *Physical Review Letters*, 75(22):4011–4014, November 1995.

[53] D. Reiß, K. Abich, W. Neuhauser, Ch. Wunderlich, and P. E. Toschek. Raman cooling and heating of two trapped Ba^+ ions. *Physical Review A*, 65:053401, April 2002.

[54] Giovanna Morigi. Cooling atomic motion with quantum interference. *Physical Review A*, 67:033402, 2003.

[55] C. Raab. *Interference Experiments with the Fluorescence Light of Ba^+ ions.* PhD thesis, Institut für Experimentalphysik, Innsbruck, January 2001.

[56] C. A. Schrama, E. Peik, W. W. Smith, and H. Walther. Novel miniature ion traps. *Opt. Comm.*, 101:32–36, March 1993.

[57] Microlase, now Coherent. The laser is the equivalent of the current Coherent MBR-110 model.

[58] C. S. Adams. *Narrow linewidth operation of a Ti:sapphire laser: Applications in high resolution spectroscopy.* PhD thesis, Department of Physics and Applied Physics, University of Strathclyde, December 1991.

[59] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward. *Applied Physics B*, 31:97, 1983.

[60] Eric D. Black. An introduction to Pound-Drever-Hall laser frequency stabilisation. *American Journal of Physics*, 69(1):79–87, January 2001.

[61] New Focus Vortex laser.

[62] The Visual Basic ion trap control software was written and maintained by G. P. Barwood.

[63] Kapteyn-Murnane Laboratories, 4699 Nautilus Court South, #204, Boulder, CO 80301, USA.

[64] Jinendra K. Ranka, Robert S. Windeler, and Andrew J. Stentz. Visible continuum generation in air–silica microstructure optical fibers with anomalous dispersion at 800 nm. *Opt. Lett.*, 25(1):25–27, January 2000.

[65] D. J. Jones, S. A. Diddams, J. K. Ranka, A. Stentz, R. S. Windeler, J. L. Hall, and S. T. Cundiff. Carrier-envelope phase control of femtosecond mode-locked lasers and direct optical frequency synthesis. *Science*, 288:635–639, April 2000.

[66] R. Holzwarth, Th. Udem, T. W. Hänsch, J. C. Knight, W. J. Wadsworth, and P. St J. Russell. Optical frequency synthesizer for precision spectroscopy. *Physical Review Letters*, 85(11):2264–2267, September 2000.

[67] S. A. Webster, P. Taylor, M. Roberts, G. P. Barwood, and P. Gill. Kilohertz-resolution spectroscopy of the $^2\mathrm{S}_{1/2}$ – $^2\mathrm{F}_{7/2}$ electric octupole transition in a single $^{171}\mathrm{Yb}^+$ ion. *Physical Review A*, 65:052501, April 2002.

[68] Bruce W. Shore. *The Theory of Coherent Atomic Excitation.* Wiley, New York, 1990.

[69] E. Biemont and P. Quinet. Theoretical study of the $4f^{14}$ $6s$ $^2S_{1/2}$ – $4f^{13}$ $6s^2$ $^2F^0_{7/2}$ E3 transition in Yb II. *Physical Review Letters*, 81(16):3345–3346, October 1998.

[70] D. J. Berkeland and M. G. Boshier. Destabilization of dark states and optical spectroscopy in Zeeman-degenerate atomic systems. *Physical Review A*, 65:033413, 2002.

[71] R. G. DeVoe and C. Kurtsiefer. Experimental study of anomalous heating and trap instabilities in a microscopic $^{137}\mathrm{Ba}$ ion trap. *Physical Review A*, 65:063407, 2002.

[72] Wayne M. Itano. External-field shifts of the ^{199}Hg^{+} optical frequency standard. *J. Res. Natl. Inst. Stand. Technol.*, 105(6):829–837, November – December 2000.

[73] P. J. Blythe, S. A. Webster, K. Hosaka, and P. Gill. Systematic frequency shifts of the 467 nm electric octupole transition in ^{171}Yb^{+}. *J. Phys. B: At. Mol. Opt. Phys.*, 36:981–989, February 2003.

[74] G. P. Barwood, H. S. Margolis, G. Huang, P. Gill, and H. A. Klein. Measurement of the electric quadrupole moment of the 4d $^{2}\mathrm{D}_{5/2}$ level in ^{88}Sr^{+}. 2004. Submitted.

[75] G. P. Barwood. Private communication.

[76] S. A. Webster, P. Taylor, M. Roberts, G. P. Barwood, P. Blythe, and P. Gill. A frequency standard using the $^{2}\mathrm{S}_{1/2}$– $^{2}\mathrm{F}_{7/2}$ octupole transition in ^{171}Yb^{+}. In P. Gill, editor, *Proceedings of the 6th Symposium on Frequency Standards and Metrology*, pages 115–122, Singapore, 2002. World Scientific.

[77] S. N. Lea, H. S. Margolis, G. Huang, W. R. C. Rowley, D. Henderson, G. P. Barwood, H. A. Klein, S. A. Webster, P. Blythe, P. Gill, and R. S. Windeler. Femtosecond optical frequency comb measurement of lasers stabilised to transitions in ^{88}Sr^{+}, ^{171}Yb^{+}, and I$_{2}$ at NPL. In P. Gill, editor, *Proceedings of the 6th Symposium on Frequency Standards and Metrology*, pages 144–151, Singapore, 2002. World Scientific.

[78] P. J. Blythe, S. A. Webster, H. S. Margolis, S. N. Lea, G. Huang, S.-K. Choi, W. R. C. Rowley, P. Gill, and R. S. Windeler. Subkilohertz absolute-frequency measurement of the 467-nm electric octupole transition in ^{171}Yb^{+}. *Phys. Rev. A*, 67:020501(R), February 2003. 4 pages.

[79] British standard BS2846: Part 1 (23), 1975.

[80] R. J. Rafac, B. C. Young, J. A Beall, W. M. Itano, D. J. Wineland, and J. C. Bergquist. Sub-dekahertz ultraviolet spectroscopy of ^{199}Hg^{+}. *Physical Review Letters*, 85(12):2462, 2000.

[81] S. Webster, M. Oxborrow, and P. Gill. Sub-hertz linewidth Nd:YAG laser. 2004. Submitted.

[82] Q. A. Turchette, D. Kielpinski, B. E. King, D. Leibfried, C. J. Myatt, M. A. Rowe, C. A. Sackett, C. S. Wood, W. M. Itano, C. Monroe, and D. J. Wineland. Heating of trapped ions from the quantum ground state. *Physical Review A*, 61:063418, May 2000.

[83] H. Rohde, S. T. Gulde, C. F. Roos, P. A. Barton, D. Leibfried, J. Eschner, F. Schmidt-Kaler, and R. Blatt. Sympathetic ground-state cooling and coherent manipulation with two-ion crystals. *J. Opt. B: Quantum Semiclass. Opt.*, 3:S34–S41, 2001.

[84] P. Meystre and M. Sargent III. *Elements of Quantum Optics.* Springer-Verlag, second edition, 1991.

[85] W. M. Itano, C. Monroe, D. M. Meekhof, D. Leibfried, B. E. King, and D. J. Wineland. Quantum harmonic oscillator state synthesis and analysis. *Proc. SPIE*, 2995(43), 1997.

[86] Rodney Loudon. *The Quantum Theory of Light.* Oxford Science Publications, second edition, 1995.

[87] C. F. Roos. *Controlling the quantum state of trapped ions.* PhD thesis, Institut für Experimentalphysik, Innsbruck, February 2000.

[88] M. McDonnell, D. N. Stacey, and A. M. Steane. Laser linewidth effects in quantum state discrimination by electromagnetically induced transparency. *Physical Review A*, 70:053802, 2004.

[89] R. D. Cowan. *The Theory of Atomic Structure and Spectra.* Univ. California Press, Berkeley, CA, USA, 1983.

[90] P. Gill, G. P. Barwood, H. A. Klein, G. Huang, S. A. Webster, P. J. Blythe, K. Hosaka, S. N. Lea, and H. S. Margolis. Trapped ion optical frequency standards. *Meas. Sci. Technol.*, 14:1174–1186, July 2003.

[91] S.N. Lea, H.S. Margolis, G. Huang, G.P. Barwood, H.A. Klein, P.J. Blythe, S.A. Webster, W.R.C. Rowley, and P. Gill. Femtosecond comb measurements of optical frequency standards at NPL, 2002. Proceedings of The 15th Annual Meeting of the IEEE Lasers and Electro-Optics Society, Glasgow.

[92] P. Gill, G.P. Barwood, G. Huang, H.A. Klein, P.J. Blythe, K. Hosaka, R.C. Thompson, S.A. Webster, S.N. Lea, and H.S. Margolis. Trapped ion optical frequency standards, 2003. 35th Conference of the European Group for Atomic Spectroscopy, Brussels.

[93] Stephen Lea, Helen Margolis, Guilong Huang, Peter Blythe, Stephen Webster, Geoff Barwood, Hugh Klein, Dick Rowley, Patrick Gill, and Robert S. Windeler. Femtosecond comb measurements of optical frequency standards at NPL, 2002. Conference on Precision Electromagnetic Measurements, Ottowa.

[94] H. S. Margolis, G. P. Barwood, P. J. Blythe, P. Gill, K. Hosaka, G. Huang, H. A. Klein, S. N. Lea, and S. A. Webster. Trapped ion optical frequency standards and metrology, 2003. IEEE LEOS Photonic Time/Frequency Measurement and Control Meeting, Vancouver.